CSE 150A 250A
AI: Probabilistic Models

Fall 2025

Instructor: Trevor Bonjour

Slides adapted from previous versions of the course (Prof. Lawrence, Prof. Alvarado, Prof Berg-Kirkpatrick)

Agenda

- Course Logistics
- Syllabus Review
- Course Overview
- Probability Review

Course Website

https://tbonjour-courses.github.io/cse150a250a-fa25/

We are here to help!

It’s a BIG class! We need to work together!

Course enrollments are controlled by the department.

- They don’t let me as an instructor directly approve EASY requests.
- Please submit one – department will try to approve as many as they can.
- The course is limited by TA capacity and room size. Can’t add more seats.

I have another class right after this, so I won’t be able to stick around after the lectures :/

Prerequisites

- Programming
  - Most HW assignments will involve coding in **Python**.
  - Also, basic data analysis and visualization.
  - We can help with algorithmic and conceptual issues.
  - We cannot help with installing, compiling, plotting, etc.

Non-CS backgrounds are welcome.

Prerequisites

- Elementary probability:
  - Random variables — discrete and continuous
  - Expected values (via sums and integrals)
- Multivariable calculus:
  - Chain rule
  - Gradients and partial derivatives
  - Computing maxima and minima
  - Constrained optimization with Lagrange multipliers

Prerequisites

- Linear algebra
  - Vectors and matrices
  - Matrix multiplication, inverses, determinants
  - Systems of linear equations
- Mathematical maturity
  - Patience and persistence go a long way
  - Willingness to fill in gaps

Course Overview

- What we do cover:
  - Inference and learning in Bayesian networks
  - Markov decision processes for reinforcement learning (RL)
- What we don't cover (not exhaustive):
  - Neural architectures (though we will talk about deep learning)
  - Purely logical reasoning
  - Heuristic search (A*)
  - Theorem proving
  - Genetic algorithms Philosophy of AI

Course Overview

- What we do cover:
  - Inference and learning in Bayesian networks
  - Markov decision processes for reinforcement learning (RL)

Why these topics?

Skyscraper

Beams/Columns

Foundation

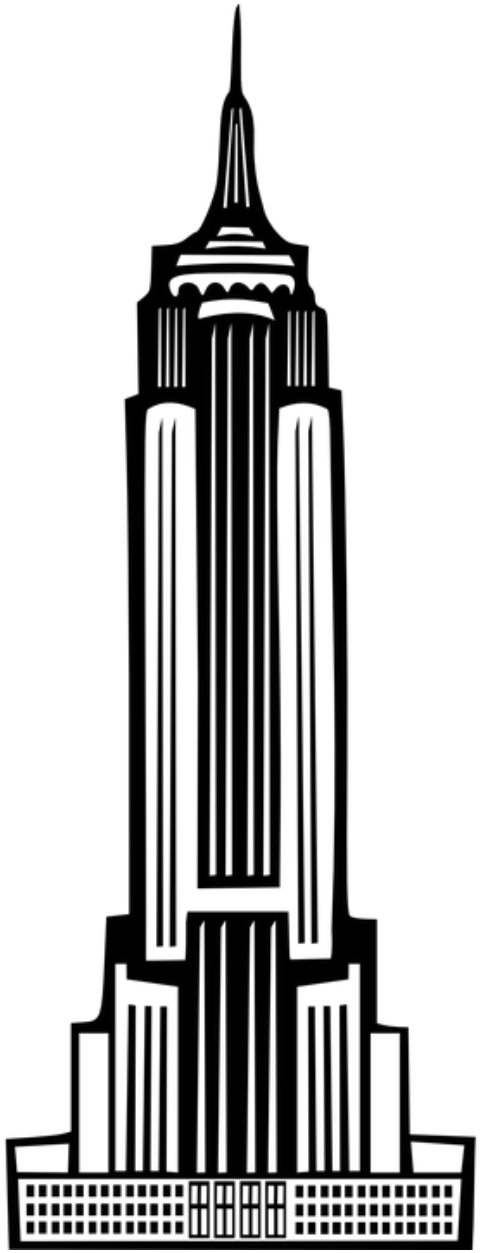

Modern AI
(ChatGPT, LLMs)

Neural Networks
(Deep Learning)

(THIS COURSE)

Probabilistic Models
(Bayesian Networks)

Turing Award 2011

Turing Award Citation:

"Judea Pearl is credited with the invention of Bayesian networks, a mathematical formalism for defining complex probability models, as well as the principal algorithms used for inference in these models. This work not only revolutionized the field of AI but also became an important tool for many other branches of engineering and the natural sciences."

Image Source: Lex Freidman Podcast

Probability and Neural Nets

“a classy synthwave apartment on mars, digital art”

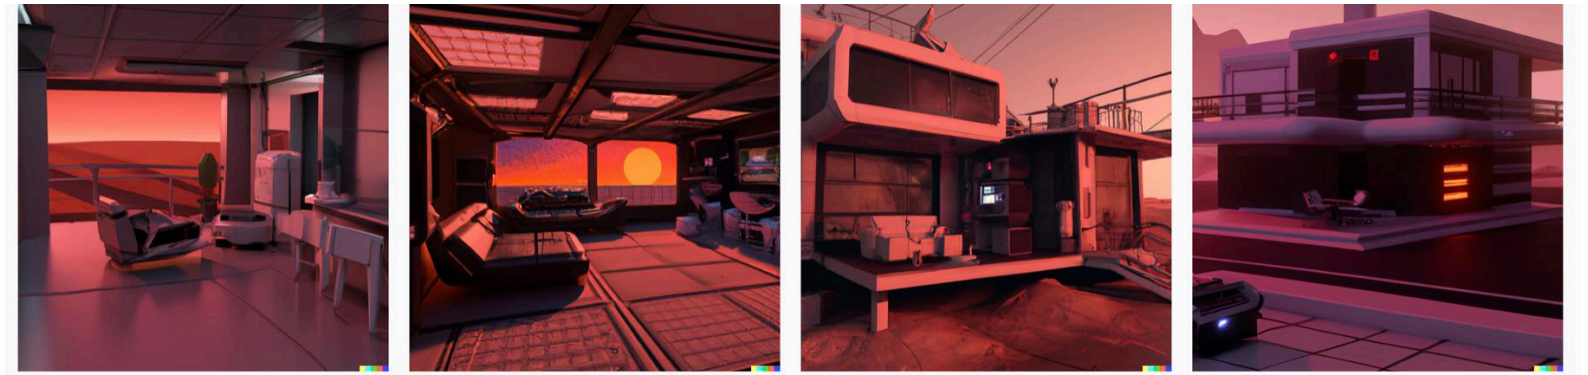

“an intricate line drawing of new your subway station full of trumpet players”

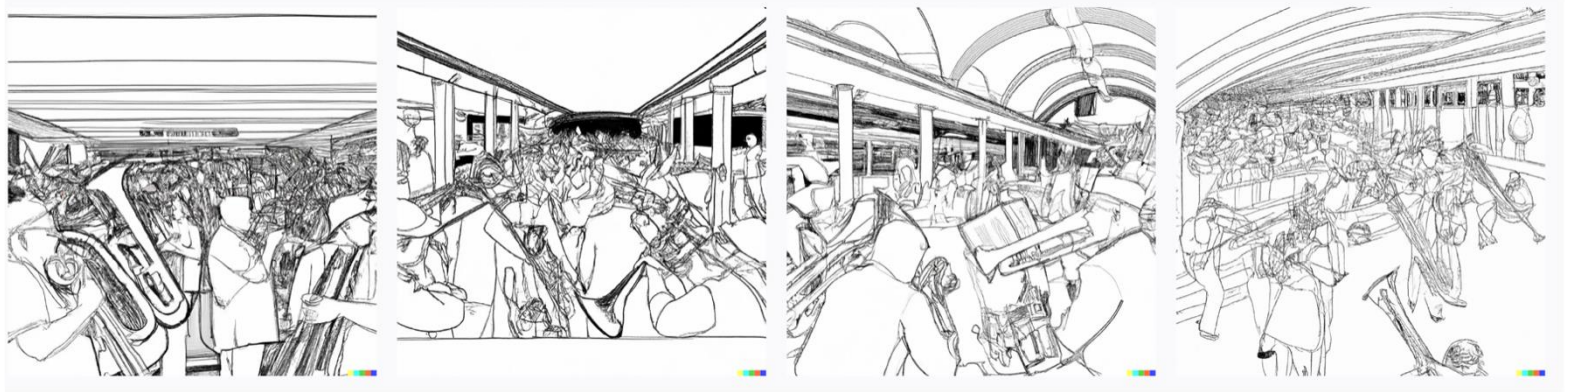

Image Source: Prof Berg-Kirkpatrick

Probability and Neural Nets

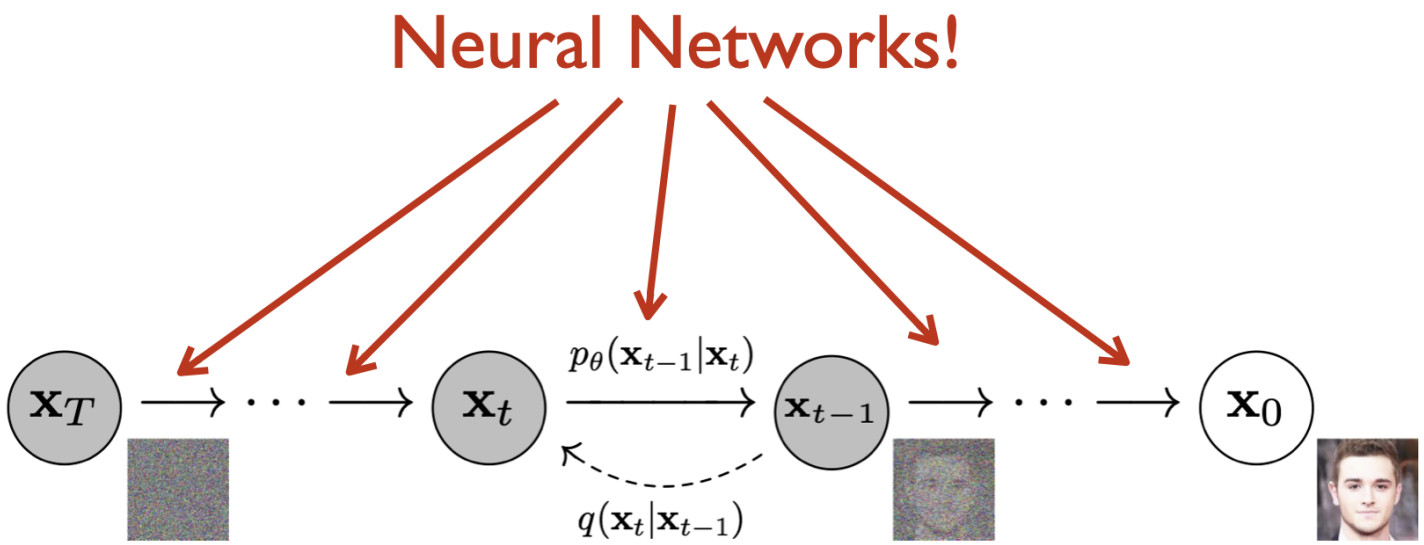

$$\mathbb{E}_q\Big[\underbrace{D_{\mathrm{KL}}(q(\mathbf{x}_T|\mathbf{x}_0)\parallel p(\mathbf{x}_T))}_{L_T}+\sum_{t>1}\underbrace{D_{\mathrm{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t,\mathbf{x}_0)\parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))}_{L_{t-1}}\underbrace{-\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)}_{L_0}\Big]$$

source: Ho et al. 2020

Breakthrough in RL

HOW THE ARTIFICIAL-INTELLIGENCE PROGRAM ALPHAZERO MASTERED ITS GAMES

By James Somers

December 28, 2018

In 2016, a Google program soundly defeated Lee Sedol, the world's best Go player, in a match viewed by more than a hundred million people. Photograph by Ahn Young-joon / AP

Chess commentators have praised AlphaZero, declaring that the engine "plays like a human on fire." Photograph Courtesy DeepMind Technologies

That’s all folks!