CSE 150 4/2

Course segments:

Probabilistic Reasoning

- Ex: medical diagnosis
  - Knowledge representation: diseases <u>cause</u> symptoms
  - Modeling uncertainty: some diseases, some symptoms <u>more likely</u> than others
  - Reasoning: <u>infer</u> diseases <u>from</u> symptoms

- Probability: quantitative, self-consistent framework that captures commonsense patterns of reasoning

Graphical Model

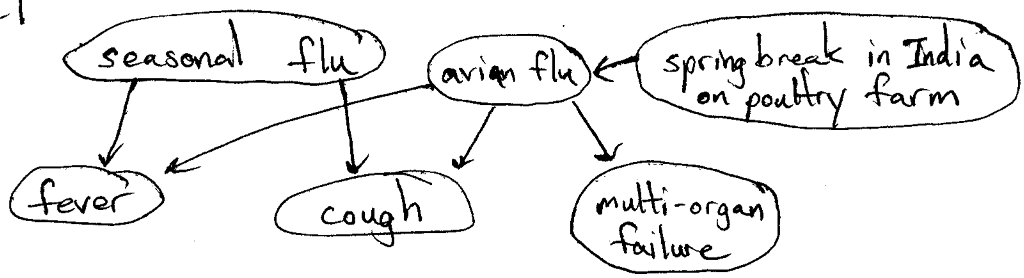

How do graphs represent correlation, causation, statistical independence?
Marriage of probability and graph theory.

Classification

- Ex: spam filtering
  - inputs: email messages
  - output: {spam, non-spam}
- How to represent input?

  Convert text to vector of word counts (or indicators)

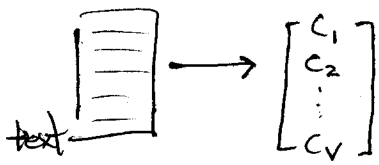

V = vocabulary size

c_i = # times that ith word appears in text

- Graphical Model

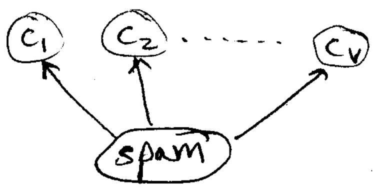

Certain words more likely in spam.
How to quantify? estimate? classify?

Sequential Modeling

- How to model systems whose "state" changes over time (or has a similarly extended representation)?

- Ex: text (written language)
  "states" = words
  which sentence is more likely?
  1. Mary had a little lamb.
  2. Colorless green ideas sleep furiously.

$\Longrightarrow$ Markov models for statistical language processing.
W_l = word at the l^{th} position in sentence (with V possible values)

- Graphical Model

(A)

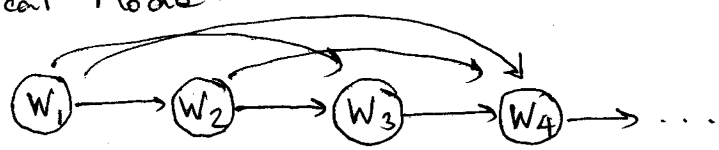

"The green..." banana?
"Some green..." bananas?

(B)

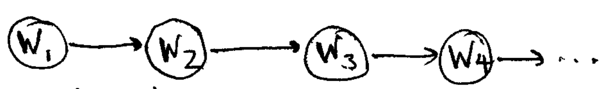

This model rules out imprabable sentence,
even though it is simpler...

Model (A) is richer but harder to estimate.
Model (B) is wrong but easier to estimate.

| Tradeoff: power vs tractability, learnability |

- Ex: speech (spoken language)
  states = words (or syllables or smaller units of speech)
  observations = sounds, waveforms

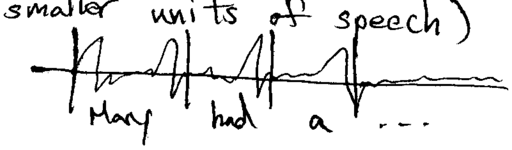

How to infer words from waveforms?
$\Longrightarrow$ hidden Markov models for speech recognition

Planning and Decision-making

- Ex: robot navigation

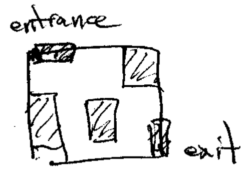

  - 2D grid world
  - "states" = cells on 2D grid
  - actions = north, south, east, west
  - noisy dynamics in world
  - rewards = feedback from environment
    - delayed vs. immediate
    - evaluative vs. instructive
- More generally: How can autonomous agents learn from experience?

  $\Longrightarrow$ Markov decision process
  Reinforcement learning

- Other "embodied" agents: helicopter, elevators
- other "embedded" agents: game-playing (e.g. backgammon), spoken dialog system

Core Ideas of Modern AI

① Probabilistic modeling of uncertainty

② Learning as optimization

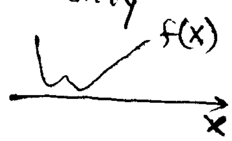

compute $x^* = \arg\min_x f(x)$

- modern AI ...
  objective function f measure agent's performance
  variables x describe/parameterize agent's behavior

③ Knowledge as predictions (dynamic)
not facts (static)

- old, classical AI
  fact 1: a canary is a bird
  fact 2: a book is on the table
- modern, agent-centric AI
  prediction: if action a, then observe consequence with probability p

Themes of class

- power vs tractability: how to develop compact representations of complex worlds

- principles vs. heuristics
  {optimization, calculation} vs. rules-of-thumb

- synergies of AI
  inference and learning,
  perception and action,
  theory and practice