Course segments

1) Probabilistic reasoning
   Ex: medical diagnosis
   Knowledge representation : diseases cause symptoms
   Modeling uncertainty : some diseases, symptoms more likely than others

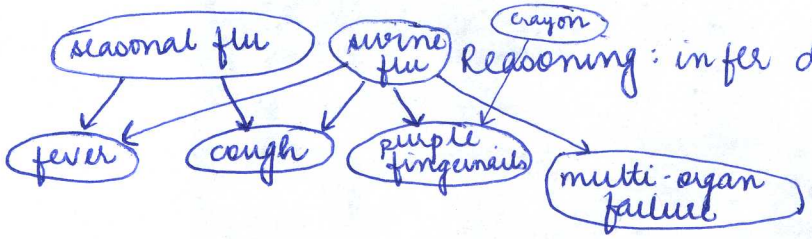

Reasoning : infer diseases from symptoms

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 in e-mail.
input : email message
output : {spam, non-spam}
How to represent the input?
Convert text to vector of word counts :

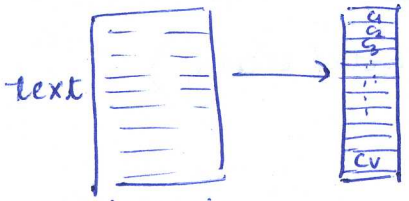

c_i = count # times i^{th} word in email

V = vocabulary size

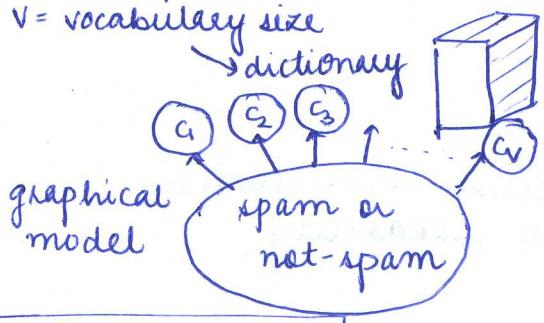

Certain words more likely in spam.
How to quantify? Estimate?

Sequential modeling
How to model systems where "state" changes over time (or has a similarly extended representation)?
Ex : text (written language)
"state" = word
Which sentence is more likely?

① Mary had a little lamb.
② Colorless green ideas sleep furiously.

⇒ Markov models of statistical language processing

W_ℓ = word at ℓ^{th} position in sentence (with V possible values)

Graphical models

A) 

Model A is rich but harder to estimate.

B) 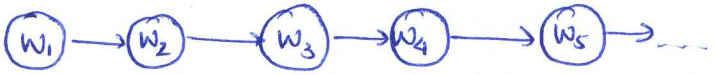

Model B is obviously wrong, over-simplified but easier to estimate.

Tradeoff: { power, expressiveness } vs { tractability, learning }

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

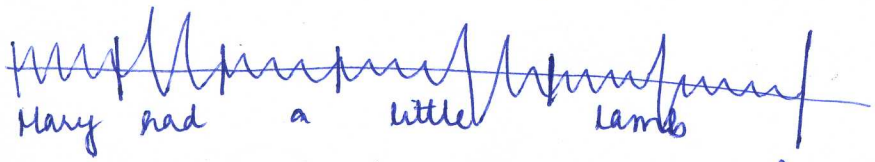

How to infer words from waveforms?
⇒ hidden Markov models ~~from~~ for speech recognition

Planning and decision-making

Ex: robot navigation

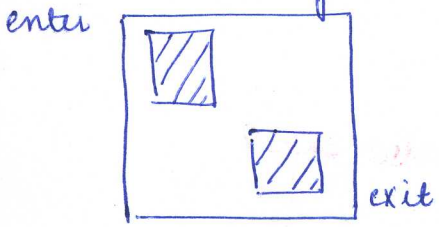

2D grid world
- "states" = cells of grid
- actions = north, south, east, west
- noisy dynamics in world

Learning = feedback from environment { delayed vs immediate; evaluative vs instructive }

More generally: how can autonomous agents learn from experience?
⇒ Markov decision processes
Reinforcement learning

Other "embodied" agents: helicopter, elevators
Other "embedded" agents: game playing ~~agents~~ AI, telephone operators

Core ideas of modern AI

1) Probabilistic models of uncertainty
2) Learning as optimization

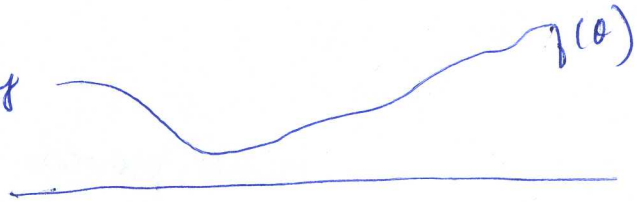

Variable θ parametrizes agent's model / predictions / behaviors.
Function $f(\theta)$ measures agent's performance.

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

- classical AI

  fact #1 : canaries are birds
  fact #2 : a bird is on the table

- modern, agent-centric AI

  prediction : if action a, then observe consequence with same probability.