CS 6347

Causal Inference

(based on slides of David Sontag and Uri Shalit)

What We've Done

- Compactly representable models
  - Bayesian networks, MRFs, CRFs, artificial neural nets
- Inference
  - Variable elimination
  - Loopy belief propagation & the Bethe free energy
  - Mean-field methods
  - Approximate MAP inference: MAP LP and duality
  - Sampling methods: importance, Gibbs sampling
- Learning
  - Maximum likelihood & psuedolikelihood
  - Expectation maximization
  - Structure learning

Limits of the Theory

- Our thought process:
  - Collect data
  - Build model
  - Do inference
- What are the limitations of this approach?
  - That is, what kinds of questions can't we answer with this approach?

Causal Inference

- Our models are not really capable of answering questions about causation
  - Recall that, in Bayesian networks, it may be tempting to infer causation from the directions of the arrows, but this isn't justified
    - The arrows only indicate which conditional probabilities are being modeled
- A philosophical question: how do we determine whether or not X causes Y (called causal inference)?
  - Can we tell just by looking at data?

Causality

- An example
  - Suppose a patient goes to their doctor and is diagnosed with high blood pressure
  - Further, suppose the doctor can give her one of two possible drugs to lower her blood pressure
  - How should the doctor select the appropriate drug?

Causality

- An example
  - Suppose a patient goes to their doctor and is diagnosed with high blood pressure
  - Further, suppose the doctor can give her one of two possible drugs to lower her blood pressure
  - How should the doctor select the appropriate drug?
    - One approach: find similar patients and look at their outcomes under the two different drugs

Causality

- An example
  - Suppose a patient goes to their doctor and is diagnosed with high blood pressure
  - Further, suppose the doctor can give her one of two possible drugs to lower her blood pressure
  - How should the doctor select the appropriate drug?
    - Another approach: build a model for classification of blood pressure given patient features

Causality

- Build a model for classification of blood pressure given patient features
  - Patient features include age, weight, etc. plus which drug they are taking
  - Goal is to predict blood pressure
    - E.g., could use naive Bayes
- Issues:
  - This really isn't the correct model as it is trained to predict blood pressure, not to predict the influence of the possible drugs

Causality

- Build a model for classification of blood pressure given patient features
  - Patient features include age, weight, etc. plus which drug they are taking
  - Goal is to predict blood pressure
    - E.g., could use naive Bayes
- Issues:
  - What happens if the model basically ignores the features corresponding to the drugs?

Causality

- Build a model for classification of blood pressure given patient features
    - Patient features include age, weight, etc. plus which drug they are taking
    - Goal is to predict blood pressure
        - E.g., could use naive Bayes
- Issues:
    - What happens if the training data doesn't contain patients like the current patient?

Causality

- Build a model for classification of blood pressure given patient features
  - Patient features include age, weight, etc. plus which drug they are taking
  - Goal is to predict blood pressure
    - E.g., could use naive Bayes
- Issues:
  - What happens if there are reasons why the training data doesn’t contain patients like the current patient (e.g., you can’t give a certain drug to a certain type of patient)?

Causality

- One thought: causal relationships may not, in general, be able to be learned only from our training data
    - Can use observations to rule out possibilities and formulate hypotheses about which variables may be causally linked
    - Need an intervention or experiment to actually test those hypotheses
- Challenges:
    - We can't go back in time, change treatments, and observe the effects

Causality

- One thought: causal relationships may not, in general, be able to be learned only from our training data
  - Can use observations to rule out possibilities and formulate hypotheses about which variables may be causally linked
  - Need an intervention or experiment to actually test those hypotheses
- Challenges:
  - In practice, their may not be any “true zeros”, that is, almost all variables may have some small effect on each other

Causality

- One thought: causal relationships may not, in general, be able to be learned only from our training data
  - Can use observations to rule out possibilities and formulate hypotheses about which variables may be causally linked
  - Need an intervention or experiment to actually test those hypotheses
- Challenges:
  - Confounding variables may prevent us from accurately assessing a causal relationship

Confounding Variables

- In the blood pressure example socioeconomic status is a confounding variable
  - Maybe one medication is given disproportionally based on wealth
- Other examples:
  - Does smoking cause cancer?
  - Do stricter gun laws make communities safer?
  - Will a particular ad campaign increase sales?
  - Does a company discriminate in its hiring practices?

Causal Inference in Practice

- Randomized trials remain the gold standard method for determining causality
  - Drawbacks:
    - Can’t try all possible outcomes in practice
      - Does asbestos cause cancer?
      - Does a particular drug cause heart disease?

Causal Inference in Practice

- Randomized trials remain the gold standard method for determining causality
  - Drawbacks:
    - Can only assess the effects of confounding variables that are part of the controlled experiment
    - Difficult to populate a trial with a uniform sample of the desired population

Causal Inference in Practice

- Randomized trials remain the gold standard method for determining causality
  - Drawbacks:
    - Study could fail to generalize from one locale to the next
    - Conclusions apply at a population level, not the individual level

Causal Inference

- Goal: build a mathematical model of causal inference
  - Many, many challenges and disagreements (both practical and philosophical) about the right way to model causality
    - At one extreme: causality can only be inferred under very strict modelling assumptions
    - The middle: causality can be inferred from appropriately designed randomized trials
    - At the other extreme: many causal relationships should be able inferred from observational data (we do it all the time!)

Causal Inference

- Goal: build a mathematical model of causal inference
  - In practice, assumptions needed to make causal inference doable
    - All confounders must be part of the model
    - Outcomes should be independent of the treatments given the features
    - Difficult to assess whether or not assumptions hold in practice
- Ongoing area of research