Bayes net wrapup

- Exact inference algorithms
  - Use to compute P(X1, …, Xn)
    or P(X1, …, Xn | Y1, …, Ym)
- Approximate inference algorithms
  - Direct sampling
  - Rejection sampling
  - Likelihood weighting

Direct Sampling

- To estimate P(X1, …, Xn), sample from the bayes net and count how many samples match the query.
- Divide by the total number of samples.

Rejection Sampling

- We want to estimate P(X1, …, Xn | Y1, …, Ym).
- P(X1, …, Xn | Y1, …, Ym)
  = P(X1, …, Xn, Y1, …, Ym) / P(Y1, …, Ym)
- To estimate this conditional probability, sample from the bayes net and count how many samples match the numerator, divide by the number of samples that match the denominator.
- Called rejection sampling because if the Y variables are very restrictive, not many samples will match (we reject them).

Likelihood weighting

- We want to estimate P(X1, …, Xn | Y1, …, Ym).
- We avoid generating samples that don't fit the Y variables.
  - Rejection sampling wastes time by throwing these out.
- Instead, we directly generate samples that fit the Y variables along with a corresponding **weight**.
- After generating a bunch of samples, count them proportionally to the sum of their weights.