Reasoning with Uncertainty

Probability

Probability

- Bayesian probabilities summarize the effects of uncertainty on the state of knowledge
  - Probabilities represent the values of statistics
    - *P(o) = (# of times of outcome o) / (# of outcomes)*
  - All types of uncertainty are incorporated into a single number

    P(H | E)
  - Probabilities follow a set of strict axioms

Probability

- Random variables define the entities of probability theory
  - Propositional random variables:
    - E.g.: IsRed, Earthquake
  - Multivalued random variables:
    - E.g.: Color, Weather
  - Potentially Real-Valued
    - E.g.: Height, Weight

Axioms of Probability

- Probability follows a fixed set of rules
  - Propositional random variables:
  - $P(A) \in [0..1]$
  - $P(T) = 1$, $P(F) = 0$
  - $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$
  - $P(A \wedge B) = P(A)\, P(B|A)$
  - $\sum_{x \in Values(X)} P(X=x) = 1$

Probability Syntax

- Unconditional or prior probabilities represent the state of knowledge before new observations or evidence
  - *P(H)*
- A probability distribution gives values for all possible assignments to a random variable
- A joint probability distribution gives values for all possible assignments to all random variables

Conditional Probability

- Conditional probabilities represent the probability after certain observations or facts have been considered
  - $P(H|E)$ is the posterior probability of H after evidence E is taken into account
  - Bayes rule allows to derive posterior probabilities from prior probabilities
    - $P(H \mid E) = P(E \mid H) P(H)/P(E)$

Conditional Probability

- Probability calculations can be conditioned by conditioning all terms
  - Often it is easier to find conditional probabilities
- Conditions can be removed by marginalization
  - $P(H) = \sum_E P(H|E)\, P(E)$

Joint Distributions

- A joint distribution defines the probability values for all possible assignments to all random variables
  - Exponential in the number of random variables
  - Conditional probabilities can be computed from a joint probability distribution
    - $P(A|B) = P(A \wedge B)/P(B)$

Inference

- Inference in probabilistic representation involves the computation of (conditional) probabilities from the available information
  - Most frequently the computation of a posterior probability *P(H/E)* form a prior probability *P(H)* and new evidence *E*

Probabilistic Inference

- Benefits:
  - Precise quantitative measure of uncertainty
  - Inference can be automated
- Problems:
  - *Worst case time complexity:*
    - $O(d^n)$:d = arity, n = #of random variables
  - *Worst case space complexity:*
    - $O(d^n)$