

Informatics 2D: Reasoning and Agents

Alex Lascarides

School of
informatics



Lecture 29b: Decision Making under Uncertainty
Decision Networks


Where are we?

Last time. . .

- Maximise Expected Utility=Rational behaviour
- Axioms of preferences, utility functions
- Today: **Decision Networks**
Combine utility with Bayesian Net representing (uncertain) beliefs


Decision networks

- What we now need is a way of integrating utilities into our view of probabilistic reasoning
- **Decision networks (influence diagrams)** combine BNs with additional node types for actions and utilities
- Illustrate with airport siting problem:




Decision networks

- What we now need is a way of integrating utilities into our view of probabilistic reasoning
- **Decision networks (influence diagrams)** combine BNs with additional node types for actions and utilities
- Illustrate with airport siting problem:



Decision networks

- What we now need is a way of integrating utilities into our view of probabilistic reasoning
- **Decision networks (influence diagrams)** combine BNs with additional node types for actions and utilities
- Illustrate with airport siting problem:




Representing decision problems with DNs

- **Chance nodes** (ovals) represent random variables with CPTs, parents can be decision nodes
- **Decision nodes** represent decision-making points at which actions are available
- **Utility nodes** represent utility function connected to all nodes that affect utility directly
- Often nodes describing outcome states are omitted and expected utility associated with actions is expressed (rather than states) – **action-utility tables**

Representing decision problems with DNs

- Simplified version with action-utility tables
- Less flexible but simpler (like pre-compiled version of general case)



Evaluating decision networks

- Evaluation of a DN works by setting decision node to every possible value
- “Algorithm”:
 - 1 Set evidence variables for current state
 - 2 For each value of decision node:
 - 1 Set decision node to that value
 - 2 Calculate posterior probabilities for parents of utility node
 - 3 Calculate resulting (expected) utility for action
 - 3 Return action with highest (expected) utility
- Using any algorithm for BN inference, this yields a simple framework for building agents that make single-shot decisions

Summary

- Foundations for rational decision making under uncertainty
- Decision networks nicely blend with our BN framework
- Only looked at one-shot decisions so far
- Next time: **Markov Decision Processes**