Decision Networks

Decision networks: the concept

- A framework for rational decision making.
- Better known as influence diagrams.
- Example:

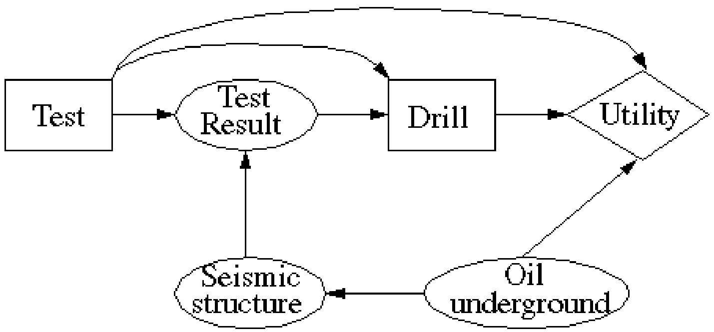

- DAG with three types of nodes:
  - Decision nodes (rectangles): decisions to make.
  - Random nodes (ovals): random quantities relevant to the decisions.
  - Utility nodes (diamonds): the utility function.

Semantics of arcs:

- Arcs into random nodes: probabilistic dependence.
- Arcs into decision nodes: information availability. Also called *informational arcs*.

  The arcs into *drill* means that the drill decision is to be made based on *test* and *test-result*.

Decision Networks

Numerical information:

- Each random node is associated with a conditional probability table.

P(oil)

dry	*wet*	*soaking*
.500	.300	.200

P(structure|oil)

	dry	*wet*	*soaking*
no-structure	.600	.300	.100
open	.300	.400	.400
closed	.100	.300	.500

P(test-result|test, structure)

	not-test	*test*		
		no-structure	*open*	*closed*
no-result	1.0	0.0	0.0	0.0
no-structure	0.0	0.9	0.1	0.0
open	0.0	0.1	0.8	0.1
closed	0.0	0.0	0.1	0.9

Decision Networks

Numerical information (cont'd):

- The utility node is associated with a utility (value) function.

$$f(\text{test, drill, oil})$$

not-test				*test*			
not-drill	*drill*			*not-drill*	*drill*		
	dry	*wet*	*soaking*		*dry*	*wet*	*soaking*
$0K	-$70K	$50K	$200K	-$10K	-$80K	$40K	$190K

Utility

- Money is a basis for the utility function.
- Amounts of money usually cannot be used as the utility function directly.
- Counter example: Suppose you have won some lucky draw. If the host offers you two choices:
  1. A grand prize of $1,000,000 for sure.
  2. A game where you flip a coin. You get $3,000,000 if the coin comes up head and you get nothing if the coin comes up tail.

  Which one will you take?
- Most people would choose 1. However, the expected value of the game is

$$0.5 \times 3,000,000 = 1,500,000 > 1,000,000.$$

  Why?
  - Let c be one's current asset, $c > 0$.
  - $U(c + 1,000,000) > 0.5U(c) + 0.5U(c + 3,000,000)$
  - One's utility function is concave.
  - In other words, people are usually **risk-averse** with positive wealth.

Utility

- Now assume you are $10,000,00 in debt and the choices are
  1. A grand prize of $1,000,000 for sure.
  2. A game where you flip a coin. You get $10,000,000 if the coin comes up head and you lose $10,000,000 the coin comes up tail.

  Which one will you take?
- Some people would choose 2. However, the expected value of the game is

  $$0.5 \times 10,000,000 - 0.5 \times 10,000,000 = 0 < 1,000,000.$$

  Why?
  - $U(-9,000,000) < 0.5U(0) + 0.5U(-20,000,000)$
  - One's utility function is convex.
  - In other words, people are usually **risk-seeking** with negative wealth.
- If an agent's untility function is linear, s/he is said to be **risk-neutral**. People are risk-neutral with **small** changes in wealth.

Utility

Insurance and lotteries (gambling)

- Most people are risk-averse. This is why the insurance industry can flourish.
- Why people buy lotteries or gamble?
  - Fun.
  - Risk-seeking when poor. (Many poor people do not buy insurance.)
  - How government make money by issue lotteries?

Decision Networks

Solution to a decision network: an optimal policy.

- *Decision rule (table)*: A mapping from information variables (parents) to decision variables.
- Example: $\text{rule}_{\text{drill}}$: $\{$*test, test-result*$\} \to$ *drill*

	not-test	*test*		
	no-result	*no-structure*	open	closed
$\text{rule1}_{\text{drill}}$	*drill*	*not-drill*	*drill*	*drill*
$\text{rule2}_{\text{drill}}$	*drill*	*not-drill*	*not-drill*	*drill*
:	:	:	:	:

- Example: Since the variable *test* has no parents, a decision rule for *test* is simply one of its values.

$\text{rule1}_{\text{test}}$	*not-test*
$\text{rule2}_{\text{test}}$	*test*

- *Policy*: a tuple of decision rules consisting of one decision rule for each decision node, e.g. Δ=($\text{rule1}_{\text{test}}$, $\text{rule1}_{\text{drill}}$).

Decision Networks

- Given a decision rule, a decision variable can be viewed as a random variable.
  - Given $\text{rule1}_{\text{drill}}$, *drill* can be viewed as a random variable with conditional probability:

$$P(\text{drill}|\text{test, test-result})$$

	not-test	*test*		
	no-resulty	*no-structure*	*open*	*closed*
not-drill	0.0	0.0	0.0	0.0
drill	1.0	0.0	1.0	1.0

  - Given $\text{rule1}_{\text{test}}$, *test* can be viewed as a random variable with probability:

$$P(\text{test})$$

not-test	1.0
test	0.0

- Given policy Δ=($\text{rule1}_{\text{test}}$, $\text{rule1}_{\text{drill}}$), we can compute

$$P_\Delta(\text{test, drill, oil})$$

  by belief network inference.

- *Expected utility of policy* Δ=($\text{rule1}_{\text{test}}$, $\text{rule1}_{\text{drill}}$):

$$\sum_{\text{test, drill, oil}} P_\Delta(\text{test, drill, oil}) f(\text{test, drill, oil})$$

- **Optimal policy**: A policy with highest expected utility.

Evaluating Decision Networks

A Naive algorithm.

- For each possible policy, compute its expected utility.
- Return the one with maximum expected utility.
- Consider full decision tables and pick the best one. Question being asked: Is this rule good?

	not-test	*test*		
	no-result	*no-structure*	open	closed
$\text{rule1}_{\text{drill}}$	*drill*	*not-drill*	*drill*	*drill*