Influence Diagrams
Based on the HUGIN Training Course by Anders L. Madsen

Prof. Dr. Galia Weidl
Basics of Decision-Making
TH Aschaffenburg
SoSe 2023

Outline

- More decision problems
- Influence diagrams
- Semantics of influence diagrams
- Information constraints
- Solving influence diagrams

Mildew

Two months before harvest H of a wheat field we observe the state of the crop and we observe whether it has been attacked by mildew. If there is an attack we should decide on a treatment A with fungicides having a cost. There is a utility associated with harvest.

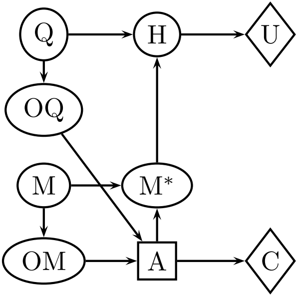

- $\mathrm{EU}(\mathrm{A} \mid \epsilon) = \mathrm{C}(\mathrm{A}) + \sum_{\mathrm{H}} \mathrm{U}(\mathrm{H})\mathrm{P}(\mathrm{H} \mid \mathrm{A}, \mathrm{e})$
- A is an intervening decision

Mildew

Two months after the decision on a treatment with fungicides, the farmer has to make a decision on the time of harvest T

T : now, wait 1 week, wait 2 weeks

A number of different information scenarios are possible w.r.t. T:

- No further information is available
- A perfect observation on H is known in addition to OQ, OM and A
- An imperfect observation OH on H is known in addition to OQ, OM and A
- Only an imperfect observation OH on H is known.
- ...

Insemination

see eBook, section 6.2.2, Example 6.5, Fig, 6.1, Fig 6.2

Six weeks after insemination (success rate 0.87) of a cow there are three tests for the result: blood test ($\text{TP} = 0.7, \text{TN} = 0.9$), urine test ($\text{TP} = 0.8, \text{TN} = 0.9$), and scanning ($\text{TP} = 0.9, \text{TN} = 0.99$). The results of the blood test and the urine test are mediated through the hormonal state ($\text{TP} = 0.9, \text{TN} = 0.99$) which is affected by a possible pregnancy.

Assume that you have the options to repeat the insemination or to wait for another six-weeks period. The cost of repeating the insemination is 65 no matter the pregnancy state of the cow. If the cow is pregnant, and you wait, it will cost you nothing, but if the cow is not pregnant, and you wait, it will cost you further 30 (that makes a total of 95 for waiting).

Influence Diagram

The influence diagram $\mathcal{N} = (V, G, \mathcal{P}, \mathcal{U})$ is a probabilistic graphical model for representing and solving sequential decision problems:

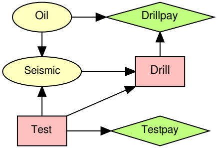

- A single decision maker with a total order on decision nodes
- Non-forgetting assumption — perfect recall
  - Recalls all past observations and decisions
  - Referred to as perfect recall influence diagrams (PRIDs)
- Finite horizon

Influence diagram $\mathcal{N}$ represents an expected utility function:

$$EU(U_C, U_D) = \prod_{i=1}^{n} P(X_i \mid pa(X_i)) * \sum_{U \in \mathcal{U}} U$$

We need to find a strategy Δ with one policy δ_i for each D_i

Influence Diagram Semantics

An influence diagram is a model of a decision scenario with a fixed sequence of decisions and a single decision maker

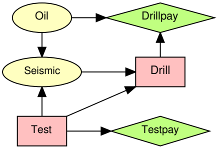

$X \rightarrow D, D_i \rightarrow D$ - Information links into a decision node D: The parents of D are known prior to making decision D

$X \rightarrow Y, D \rightarrow Y$: Links into chance node Y specifies probabilistic dependence

$X \rightarrow U, D \rightarrow U$: Links into utility node U specifies functional dependence

- The utility function is assumed to decompose additively
- Non-forgetting: The decision maker has perfect recall
- Total order on decisions

Strategy

A strategy Δ consists of one policy δ_{D} for each decision D

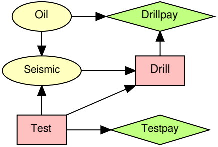

In Oil Wildcatter a strategy $\Delta = \{\hat{\delta}_{\mathrm{Test}}, \hat{\delta}_{\mathrm{Drill}}\}$ has two policies $\hat{\delta}_{\mathrm{Test}} = \mathrm{yes}$ and

$$\hat{\delta}_{\mathrm{Drill}}(\mathrm{Seismic}, \mathrm{Test}) = \begin{cases} \mathrm{yes} & \mathrm{Seismic} = \mathrm{closed}, \mathrm{Test} = \mathrm{no} \\ \ldots & \\ \mathrm{yes} & \mathrm{Seismic} = \mathrm{closed}, \mathrm{Test} = \mathrm{yes} \\ \mathrm{yes} & \mathrm{Seismic} = \mathrm{open}, \mathrm{Test} = \mathrm{yes} \\ \mathrm{no} & \mathrm{Seismic} = \mathrm{diffuse}, \mathrm{Test} = \mathrm{yes} \end{cases}$$

A policy δ_{D} is a function from past D to states of D.

To solve a decision problem is to find a (optimal) strategy Δ

Strategy

A policy δ_{D} can be encoded as a probability distribution $\mathrm{P}(\mathrm{D} \mid \mathrm{pa}(\mathrm{D}))$

$$\delta_{\mathrm{Drill}}(\mathrm{Seismic}, \mathrm{Test}) = \begin{cases} \mathrm{yes} & \mathrm{Seismic} = \mathrm{closed}, \mathrm{Test} = \mathrm{no} \\ \dots & \\ \mathrm{yes} & \mathrm{Seismic} = \mathrm{closed}, \mathrm{Test} = \mathrm{yes} \\ \mathrm{yes} & \mathrm{Seismic} = \mathrm{open}, \mathrm{Test} = \mathrm{yes} \\ \mathrm{no} & \mathrm{Seismic} = \mathrm{diffuse}, \mathrm{Test} = \mathrm{yes} \end{cases}$$

$\mathrm{P}(\mathrm{D} = \mathrm{d} \mid \mathrm{pa}(\mathrm{D}) = \mathrm{k}) = 1$
if option d selected in parent state k

$\mathrm{P}(\mathrm{Drill} \mid \mathrm{Seismic}, \mathrm{Test})$

		Drill	
Seismic	Test	no	yes
closed	no	0	1
...			
closed	yes	0	1
open	yes	0	1
diffuse	yes	1	0

We can compute probability of future events under strategy Δ

Probability of Future Decision

Influence diagram $\mathcal{N}$ represents an expected utility function:

$$\mathrm{EU}(\mathrm{U}_{\mathrm{C}} \mid \Delta) = \prod_{i=1}^{n} \mathrm{P}(\mathrm{X}_i \mid \mathrm{pa}(\mathrm{X}_i)) * \sum_{\mathrm{U} \in \mathcal{U}} \mathrm{U}$$

Δ is a strategy with one policy δ_i for each decision D_i

$$\mathrm{P}(\mathrm{V} \mid \Delta) = \prod_{\mathrm{X} \in \mathrm{V}_{\mathrm{C}}} \mathrm{P}(\mathrm{X} \mid \mathrm{pa}(\mathrm{X})) * \prod_{i=1}^{m} \delta_i(\mathrm{D}_i \mid \mathrm{pa}(\mathrm{D}_i))$$

If the policies are encoded as CPTs, $\mathcal{N}$ supports the calculation of probability of future decisions

Monty Hall As A Decision Problem

Game show setting. There are 3 doors, behind one of which is a prize. Monty Hall, the host, asks you to pick a door, any door. You pick door A (say). Monty opens door B (say) and shows voila there is nothing behind door B. Monty gives you the choice of either sticking with your original choice of door A, or switching to door C.

- Should you switch, i.e., what door will you choose ?

Extended Poker Game

During a game of poker I take a series of decision. Before deciding on whether to call or fold I also take decisions on the number of cards to change. That is, when I decide on my first change of cards I only know my initial hand, and when I decide for the second change of cards I know my two hands, my opponent's first change of cards and my own first change of cards. Accordingly, I decide to fold or to call.

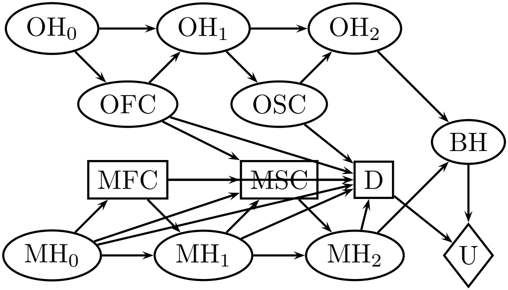

A DAG for determining the optimal decisions in poker

Information Constraints

In the perfect recall influence diagram we have a complete ordering on the decision variables $\mathrm{D}_1, \ldots, \mathrm{D}_\mathrm{n}$

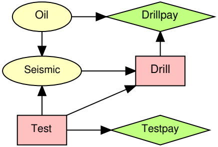

- Partition chance variables relative to decision variables:

$$\mathcal{I}_0 \prec \mathrm{D}_1 \prec \cdots \prec \mathcal{I}_{\mathrm{i}-1} \prec \mathrm{D}_\mathrm{i} \prec \mathcal{I}_\mathrm{i} \prec \mathrm{D}_{\mathrm{i}+1} \prec \cdots \prec \mathrm{D}_\mathrm{n} \prec \mathcal{I}_\mathrm{n}$$

- $\mathcal{I}_0$: the variables observed initially
- $\mathcal{I}_\mathrm{i}$: the variables observed after D_i and before $\mathrm{D}_{\mathrm{i}+1}$
- $\mathcal{I}_\mathrm{n}$: the variables never observed or observed after D_n

This partial ordering is important for the process of solving an influence diagram

Oil Wildcatter $\emptyset \prec \text{Test} \prec \{\text{Seismic}\} \prec \text{Drill} \prec \{\text{Oil}\}$

Information Constraints

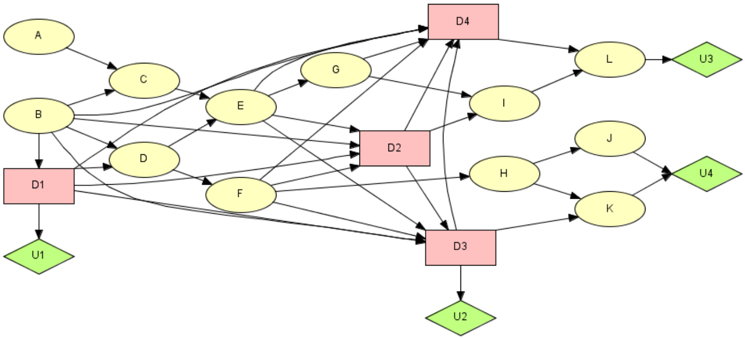

$\mathcal{I}_0 = \{\mathrm{B}\}, \;\; \mathcal{I}_1 = \{\mathrm{E}, \mathrm{F}\}, \;\; \mathcal{I}_2 = \emptyset, \;\; \mathcal{I}_3 = \{\mathrm{G}\},$
$\mathcal{I}_4 = \{\mathrm{A}, \mathrm{C}, \mathrm{D}, \mathrm{H}, \mathrm{I}, \mathrm{J}, \mathrm{K}, \mathrm{L}\}$

$\{\mathrm{B}\} \prec \mathrm{D}_1 \prec \{\mathrm{E}, \mathrm{F}\} \prec \mathrm{D}_2 \prec \emptyset \prec \mathrm{D}_3 \prec \{\mathrm{G}\} \prec \mathrm{D}_4 \prec \{\mathrm{A}, \mathrm{C}, \mathrm{D}, \mathrm{H}, \mathrm{I}, \mathrm{J}, \mathrm{K}, \mathrm{L}\}$

Policy for D_4 is, in principle, a function of $\mathrm{B}, \mathrm{D}_1, \mathrm{E}, \mathrm{F}, \mathrm{D}_2, \mathrm{D}_3, \mathrm{G}$ (the past)

The Chain Rule

The chain-rule says that an influence diagram is a representation of a joint expected utility function

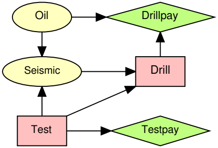

Let $U_C = \{X_1, \ldots, X_n\}$ be a universe of chance variables and let U_D be a universe of decision variables. Then:

$$P(U_C \mid U_D) = \prod_{i=1}^{n} P(X_i \mid pa(X_i))$$

where $pa(X)$ are the parents of X. Also:

$$U(U_C, U_D) = \sum_{U \in \mathcal{U}} U$$

Hence

$$EU(U_C, U_D) = P(U_C \mid U_D) * U(U_C, U_D) = \prod_{i=1}^{n} P(X_i \mid pa(X_i)) * \sum_{U \in \mathcal{U}} U$$

Oil Wildcatter $EU(U_D, U_C) = P(O)P(S|O,T)(U(T) + U(D,O))$

Solving an Influence Diagram

To solve an influence diagram $\mathcal{N}$ is to determine an optimal strategy $\hat{\Delta}$ and to compute $\mathrm{EU}(\hat{\Delta})$

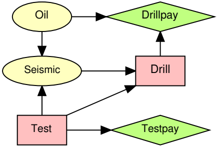

- The optimal strategy $\hat{\Delta} = (\delta_1, \ldots, \delta_n)$ consists of one optimal policy δ_i for each decision D_i
- Each policy δ_i is, in principle, a mapping from $\mathrm{pa}(D_i)$ to D_i

An influence diagram is solved by variable elimination

- Eliminate decision variables by maximization
- Eliminate chance variables by summation
- Solve for one decision at a time in reverse order (roll-back and collapse)

Oil Wildcatter: eliminate O, S by $\sum$, eliminate T, D by max

Solving Influence Diagrams

The optimal strategy $\hat{\Delta} = (\delta_1, \ldots, \delta_n)$ consists of one optimal decision policy δ_i for each decision D_i

An influence diagram can be solved by variable elimination considering one decision at a time in reverse time order

$$\mathrm{EU}(\hat{\Delta}) = \sum_{\mathcal{I}_0} \max_{D_1} \sum_{\mathcal{I}_1} \cdots \max_{D_n} \sum_{\mathcal{I}_n} \mathrm{EU}(U_D, U_C)$$

Various algorithms for solving influence diagrams exist — a strong elimination order is identified based on

$$\mathcal{I}_0 \prec D_1 \prec \cdots \prec \mathcal{I}_{i-1} \prec D_i \prec \mathcal{I}_i \prec \cdots \prec D_n \prec \mathcal{I}_n$$

A strong elimination order can be identified based on operations on the graph of the influence diagram

Solving Oil Wildcatter

Chain rule

$\mathrm{EU}(\mathrm{U_D}, \mathrm{U_C}) = \mathrm{P(O)P(S|O,T)(U(T){+}U(D,O))}$

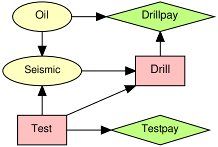

Partial order of information

$$\emptyset \prec \mathrm{Test} \prec \{\mathrm{Seismic}\} \prec \mathrm{Drill} \prec \{\mathrm{Oil}\}$$

The strong elimination order is

$$\sigma = (\mathrm{Oil}, \mathrm{Drill}, \mathrm{Seismic}, \mathrm{Test})$$

Expected utility and optimal decisions:

$$\begin{aligned} \mathrm{EU}(\hat{\Delta}) &= \max_{\mathrm{Test}} \sum_{\{\mathrm{Seismic}\}} \max_{\mathrm{Drill}} \sum_{\{\mathrm{Oil}\}} \mathrm{EU}(\mathrm{U_D}, \mathrm{U_C}) \\ &= \max_{\mathrm{T}} \sum_{\{\mathrm{S}\}} \max_{\mathrm{D}} \sum_{\{\mathrm{O}\}} \mathrm{P(O)P(S|O,T)(U(T) + U(D,O))} \end{aligned}$$

Strong Junction Tree

Expected utility and optimal decisions:

$$
\begin{aligned}
\mathrm{EU}(\hat{\Delta}) &= \max_{\mathrm{Test}} \sum_{\{\mathrm{Seismic}\}} \max_{\mathrm{Drill}} \sum_{\{\mathrm{Oil}\}} \mathrm{EU}(\mathrm{U_D}, \mathrm{U_C}) \\
&= \max_{\mathrm{T}} \sum_{\{\mathrm{S}\}} \max_{\mathrm{D}} \sum_{\{\mathrm{O}\}} \mathrm{P(O)P(S|O,T)(U(T)+U(D,O))} \\
&= \max_{\mathrm{T}} (\mathrm{U(T)} + \sum_{\{\mathrm{S}\}} \max_{\mathrm{D}} \sum_{\{\mathrm{O}\}} \mathrm{P(O)P(S|O,T)U(D,O)})
\end{aligned}
$$

Graphical representation of calculations:

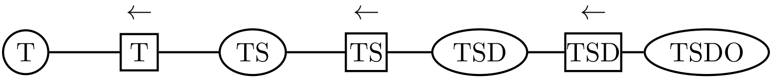

(Strong) junction tree representation:

The strong junction tree is constructed based on the strong triangulation

Oil Wildcatter Solution

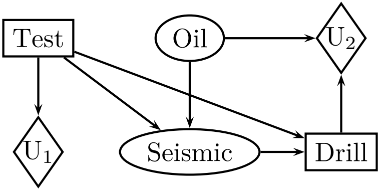

The optimal strategy Δ consists of a policy δ_{T} for T and a policy δ_{D} for D given T and S:

$$\delta_{\mathrm{T}} = \text{test} \qquad \delta_{\mathrm{D}}(\mathrm{S},\mathrm{T}) = \begin{cases} \text{drill} & \mathrm{S} = \text{closed}, \mathrm{T} = \text{test} \\ \text{drill} & \mathrm{S} = \text{open}, \mathrm{T} = \text{test} \\ \neg\text{drill} & \mathrm{S} = \text{diffuse}, \mathrm{T} = \text{test} \\ \text{drill} & \mathrm{S} = \text{closed}, \mathrm{T} = \neg\text{test} \\ \text{drill} & \mathrm{S} = \text{open}, \mathrm{T} = \neg\text{test} \\ \text{drill} & \mathrm{S} = \text{diffuse}, \mathrm{T} = \neg\text{test} \end{cases}$$

Decision Tree

Decision tree representation of Oil Wildcatter

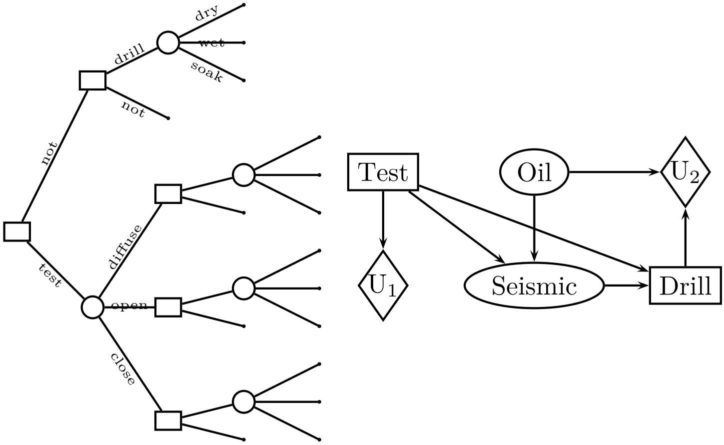

The influence diagram was introduced as a high level specification language for decision trees

Solving an Influence Diagram

The influence diagram N represents:

$$\mathrm{EU}(\mathrm{U_C}, \mathrm{U_D}) = \prod_{i=1}^{n} \mathrm{P}(\mathrm{X_i} \mid \mathrm{pa}(\mathrm{X_i})) \sum_{j=1}^{o} \mathrm{U_j}$$

To solve N, we use the $\sum$-max-$\sum$-Rule:

First average over unknown variables, then maximize over the decisions, and finally average over the variables known by the decision maker, but not known by you, the analyst

We solve N by considering one decision at a time in reverse time order

$$\mathrm{EU}(\Delta) = \sum_{\mathcal{I}_0} \max_{\mathrm{D_1}} \sum_{\mathcal{I}_1} \cdots \max_{\mathrm{D_n}} \sum_{\mathcal{I}_n} \mathrm{EU}(\mathrm{U_D}, \mathrm{U_C}),$$

This is a roll-back and collapse algorithm

Summary

- More decision problems
- Influence diagrams
- Semantics of influence diagrams
- Information constraints
- Solving influence diagrams