Decision Theory

(CIV6540 - Probabilistic Machine Learning for Civil Engineers)

Professor: James-A. Goulet

Département des génies civil, géologique et des mines
Polytechnique Montréal

Chapter 14 – Goulet (2020)
Probabilistic Machine Learning for Civil Engineers
MIT Press

Chapter 16 – Russell, S. and Norvig, P. (1995)
Artificial Intelligence, A modern approach
Prentice-Hall

Making rational decisions - Fire alarm

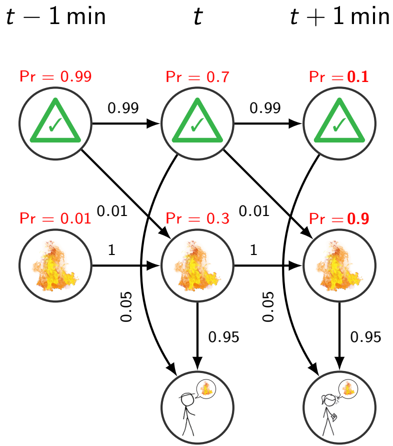

Making rational decisions - Soil contamination

We have $1\,\mathrm{m}^3$ of soil from an industrial site

What should we do?

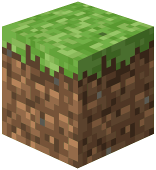

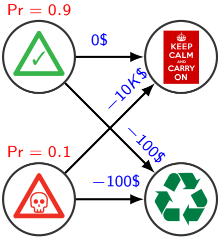

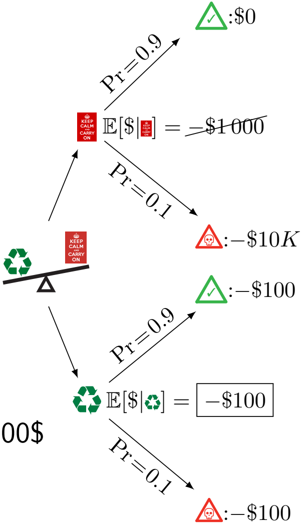

$$\mathbb{E}[\$|\text{KEEP CALM AND CARRY ON}] = (0\$ \times 0.9) + (-10K\$ \times 0.1) = -1K\$$$

$$\mathbb{E}[\$|\text{♻}] = (-100\$ \times 0.9) + (-100\$ \times 0.1) = -100\$$$

Optimal action: ♻

Nomenclature

$\mathcal{A} = \{a_1, a_2, \cdots, a_{\mathtt{A}}\}$	A set of possible actions
$x \in \mathbb{Z}$ or $\in \mathbb{R}$	An outcome in a set of possible states
$\Pr(x)$	Probability of a state x
$\mathbb{U}(a, x)$	Utility given a state x and an action a
$\mathbb{L}(a, x) \equiv -\mathbb{U}(a, x)$	Loss given a state x and an action a

Soil contamination example

$a_i \in \{0, 1\}$

$x \in \{0, 1\}$

$\Pr(x) = \{0.9, 0.1\}$

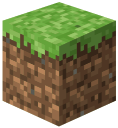

$$\mathbb{U}(a,x) = \mathbb{U} \equiv \begin{bmatrix} 0\$ & -10K\$ \\ -100\$ & -100\$ \end{bmatrix}$$

$$\mathbb{L}(a,x) = \mathbb{L} \equiv \begin{bmatrix} 0\$ & 10K\$ \\ 100\$ & 100\$ \end{bmatrix}$$

Rational decisions → Expected utility maximization

The perceived benefit of an outcome x_i given an action a_i is measured by the **expected utility** or **expected loss**

$$\overline{\mathbb{U}}(a) \equiv \mathbb{E}[\mathbb{U}(a, X)] = \sum_{i=1}^{\mathtt{X}} \mathbb{U}(a, x_i) \cdot \Pr(x_i)$$

$$\overline{\mathbb{L}}(a) \equiv \mathbb{E}[\mathbb{L}(a, X)] = \sum_{i=1}^{\mathtt{X}} \mathbb{L}(a, x_i) \cdot \Pr(x_i)$$

The **optimal** action a^* is the one that **maximizes the expected utility** or **minimizes the expected loss**

$$a^* = \arg\max_a \mathbb{E}[\mathbb{U}(a, X)] = \arg\min_a \mathbb{E}[\mathbb{L}(a, X)]$$

Module #9 Outline

Intro

Utility theory

Utility & Loss Functions

Value of Information

Topics organization

Background	1	Revision probability & linear algebra
	2	Probability distributions
Machine Learning Basics	0	Introduction
	3	Bayesian Estimation $p(A\|B) = \frac{p(B\|A)p(A)}{p(B)}$
	4	MCMC sampling & Newton
Supervised learning	5	Regression
	6	Classification
	7	LSTM networks for time series
Unsupervised learning	7	State-space model for time-series
Decision Making & RL	8	Decision Theory
	9	AI & Sequential decision problems

Section Outline

Utility theory

2.1 Lotteries

2.2 Axioms of utility theory

Nomenclature for ordering preferences

A lottery: $L_i = [\{p_1, x_1\}; \{p_2, x_2\}; \cdots; \{p_{\mathtt{X}}, x_{\mathtt{X}}\}]$

$$
\begin{aligned}
L_{♻} &= [\{1.0, (♻, \triangle)\}; \{0.0, (♻, \triangle)\}] \\
L &= [\{0.9, (\,, \triangle)\}; \{0.1, (\,, \triangle)\}]
\end{aligned}
$$

A decision maker

$L_i \succ L_j$ prefers L_i over L_j

$L_i \sim L_j$ is indifferent between L_i and L_j

$L_i \succeq L_j$ prefers L_i over L_j or is indifferent

Axioms of utility theory

What is defining a **rational behaviour**?

Orderability: Exactly one of $(L_i \succ L_j)$, $(L_j \succ L_i)$, $(L_i \sim L_j)$ holds

Transitivity: if $(L_i \succ L_j)$ and $(L_j \succ L_k)$, then $(L_i \succ L_k)$

Continuity: if $(L_i \succ L_j \succ L_k)$, then $\exists p : [\{p, L_i\}; \{1 - p, L_k\}] \sim L_j$

Substitutability:
if $(L_i \sim L_j)$, then $[\{p, L_i\}; \{1 - p, L_k\}] \sim [\{p, L_j\}; \{1 - p, L_k\}]$

Monotonicity:
if $L_i \succ L_j$, then $(p \geq q \Leftrightarrow [\{p, L_i\}; \{1 - p, L_j\}] \succ [\{q, L_i\}; \{1 - q, L_j\}])$

Decomposability: ...no fun in gambling

Section Outline

Utility & Loss Functions

3.1 Utility

3.2 Non-linear utility functions

3.3 Utility and Loss functions $\mathbb{U}(v)$ & $\mathbb{L}(v)$

3.4 $\mathbb{E}[\mathbb{U}(v(a_i, X))]$ and risk aversion

Axioms $\rightarrow$ utility

Existence of a utility function:

$$\begin{array}{ccccccc} \mathbb{U}(L_i) & > & \mathbb{U}(L_j) & \Leftrightarrow & L_i & \succ & L_j \\ \mathbb{U}(L_i) & = & \mathbb{U}(L_j) & \Leftrightarrow & L_i & \sim & L_j \end{array}$$

Expected utility of a lottery:

$$\mathbb{E}[\mathbb{U}([\{p_1, x_1\}, \{p_2, x_2\}, \cdots, \{p_{\mathtt{X}}, x_{\mathtt{X}}\}])] = \sum_{i=1}^{\mathtt{X}} p_i \mathbb{U}(x_i)$$

Invariance to linear transformation:

$$\mathbb{U}^{\mathtt{tr}}(x) = w\mathbb{U}(x) + b, \quad w > 0$$

$$\mathbb{L}(a, x) \equiv -\mathbb{U}(a, x) \begin{cases} a^* & = \arg\max_a \mathbb{E}[\mathbb{U}(a, X)] \\ & = \arg\min_a \mathbb{E}[\mathbb{L}(a, X)] \end{cases}$$

Do you want to take the lottery? [YouTube]

$$L = [\{\tfrac{1}{2}, +\ \}, \{\tfrac{1}{2}, -\ \}]$$
$$L = [\{1, +0\$\}]$$

Which lottery do you choose? Why?

$$\mathbb{E}[\$(L)] = \tfrac{1}{2} \times +200\$ + \tfrac{1}{2} \times -100\$ = +50\$$$
$$\mathbb{E}[\$(L)] = 0\$$$

Are you being irrational?
For individuals $\mathbb{U}(\$)$ and $\mathbb{L}(\$)$ are non-linear...

Risk aversion and utility functions $\mathbb{U}(v)$

$\mathbb{U}(v)$: An utility function weight monetary value (v) as a function of risk aversion/propension
(**± 1\$ not the same effet if you have 1\$ or 1M\$**)

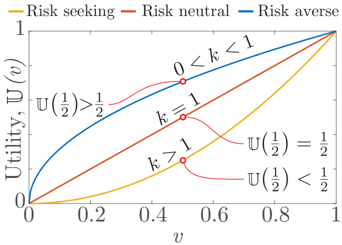

People/organizations are risk averse

$$\mathbb{U}(v) = v^k \begin{cases} k > 1 & \text{Risk seeking} \\ k = 1 & \text{Neutral} \\ 0 < k < 1 & \text{Risk averse} \end{cases}$$

Risk aversion and loss functions $\mathbb{L}(v)$

$\mathbb{L}(v)$: A loss function weight monetary value (v) as a function of risk aversion/propension

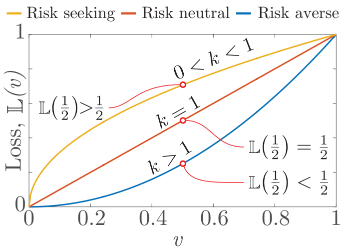

People/organizations are risk averse

$$\mathbb{L}(v) = v^k \begin{cases} k > 1 & \text{Risk averse} \\ k = 1 & \text{Neutral} \\ 0 < k < 1 & \text{Risk seeking} \end{cases}$$

Attitude toward risks

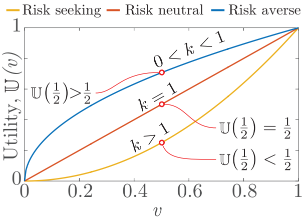

$$\mathbb{U}(v) = v^k \begin{cases} k > 1 & \text{Risk seeking} \\ k = 1 & \text{Neutral} \\ 0 < k < 1 & \text{Risk averse} \end{cases}$$

A **neutral** attitude toward risks **maximizes/minimizes the expected value/cost over a multiple decisions**

- Insurance compagnies: **neutral** attitude toward risks
- Insured people: **risk averse**; they pay a premium not to be in a risk neutral position (i.e. expected costs are higher over multiple decisions)

Expected Utility

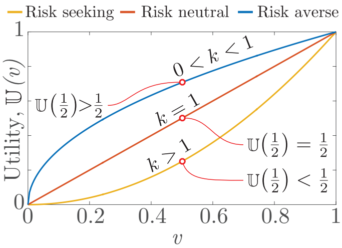

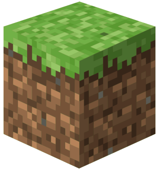

Value → Utility:

Value	$x=$ △	$x=$ ⚠	→	**Utility**	$x=$ △	$x=$ ⚠	$\mathbb{E}[\mathbb{U}(v(a,X))]$
$a=$ ♻	$v($♻, △$)$	$v($♻, ⚠$)$		$a=$ ♻	$\mathbb{U}(v($♻, △$))$	$\mathbb{U}(v($♻, ⚠$))$	$\mathbb{E}[\mathbb{U}(v($♻$, X))]$
$a=$ ■	$v($■, △$)$	$v($■, ⚠$)$		$a=$ ■	$\mathbb{U}(v($■, △$))$	$\mathbb{U}(v($■, ⚠$))$	$\mathbb{E}[\mathbb{U}(v($■$, X))]$

$\mathbb{E}[\mathbb{U}(v(a_i, X))]$ and risk aversion (ex. discrete) [.m]

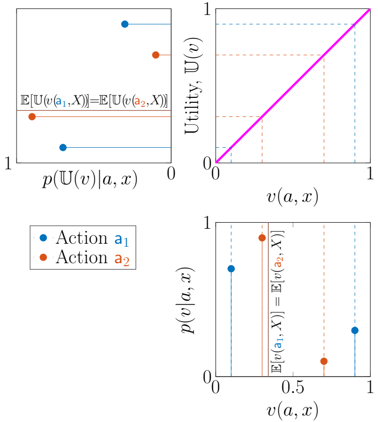

Risk perception ¬neutral: $\mathbb{E}[\mathbb{U}(v(a_1, X))] \neq \mathbb{E}[\mathbb{U}(v(a_2, X))]$

[CIV_ML/Decision/PCgA_discrete.m]

$\mathbb{E}[\mathbb{U}(v(a_i, X))]$ and risk aversion (ex. discrete) [.m]

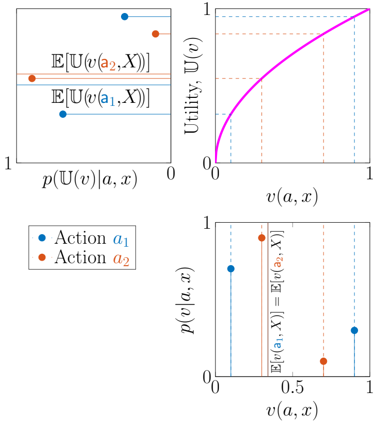

Risk perception ¬neutral: $\mathbb{E}[\mathbb{U}(v(a_1, X))] \neq \mathbb{E}[\mathbb{U}(v(a_2, X))]$

[CIV_ML/Decision/PCgA_discrete.m]

$\mathbb{E}[\mathbb{U}(v(a_i, X))]$ and risk aversion (ex. continuous) [.m]

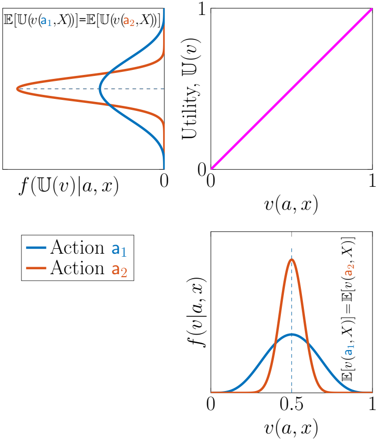

Risk perception ¬neutral: $\mathbb{E}[\mathbb{U}(v(a_1, X))] \neq \mathbb{E}[\mathbb{U}(v(a_2, X))]$

[CIV_ML/Decision/PCgA_continuous.m]

$\mathbb{E}[\mathbb{U}(v(a_i, X))]$ and risk aversion (ex. continuous) [.m]

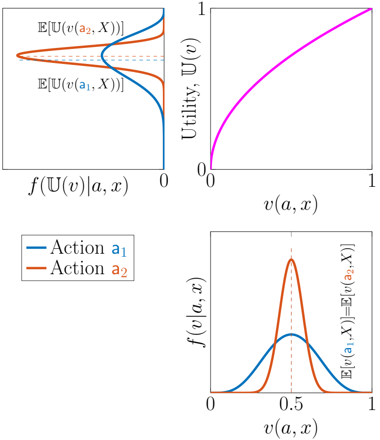

Risk perception ¬neutral: $\mathbb{E}[\mathbb{U}(v(a_1, X))] \neq \mathbb{E}[\mathbb{U}(v(a_2, X))]$

[CIV_ML/Decision/PCgA_continuous.m]

Section Outline

Value of Information

4.1 Value of perfect information
4.2 Value of imperfect information
4.3 Exemple

Expected utility of collecting information

In cases where the value of a state x **is imperfectly known**, one possible action is to **collect information about X**

$$\overline{\mathbb{U}}(a^*) \equiv \mathbb{E}[\mathbb{U}(a^*, X)] = \max_a \sum_{i=1}^{\mathtt{X}} \mathbb{U}(a, x_i) \cdot \Pr(x_i)$$

$$\mathbb{U}(a^*, x = y) = \max_a \mathbb{U}(a, x = y)$$

Because y has not been observed yet, **we must consider all possibilities $Y = X_i$ according to their probability**

$$\overline{\mathbb{U}}(\tilde{a}^*) \equiv \mathbb{E}[\mathbb{U}(\tilde{a}^*, X)] = \sum_{i=1}^{\mathtt{X}} \max_a [\mathbb{U}(a, x_i)] \cdot \Pr(x_i)$$

Value of perfect information

$$VPI(y) = \mathbb{E}[\mathbb{U}(\tilde{a}^*, X)] - \mathbb{E}[\mathbb{U}(a^*, X)] \geq 0$$

Soil contamination example

$\mathbb{U}(a,x)$	$x=$ ✓	$x=$ ⚠
$a=$ ♻	$-100\$$	$-100\$$
$a=$ KEEP CALM AND CARRY ON	$0\$$	$-10K\$$

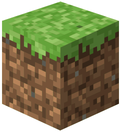

Current expected utility conditional on actions

$$\mathbb{E}[\mathbb{U}(\text{KEEP CALM AND CARRY ON},X)] = (0\$\times 0.9) + (-10K\$\times 0.1) = -1K\$$$

$$\mathbb{E}[\mathbb{U}(\text{♻},X)] = (-100\$\times 0.9) + (-100\$\times 0.1) = -100\$\ \boxed{-100\$} = \mathbb{E}[\mathbb{U}(a$$

Expected utility conditional on perfect information

$$\begin{aligned}\mathbb{E}[\mathbb{U}(\tilde{a}^*,X)] &= \sum_{i=1}^{\mathtt{X}} \min_a(\mathbb{U}(a,x_i))\cdot \Pr(x_i)\\ &= \underbrace{0\$\times 0.9}_{y=x=\text{✓}} + \underbrace{-100\$\times 0.1}_{y=x=\text{⚠}} = \boxed{-10\$}\end{aligned}$$

Value of perfect information

$$VPI(y) = \mathbb{E}[\mathbb{U}(\tilde{a}^*,X)] - \mathbb{E}[\mathbb{U}(a^*,X)] = \boxed{90\$}$$

VPI – Decision Tree Representation

$\mathbb{U}(a,x)$	$x=$ △	$x=$ △
$a=$ ♻	$-100\$$	$-100\$$
$a=$ KEEP CALM AND CARRY ON	$0\$$	$-10K\$$

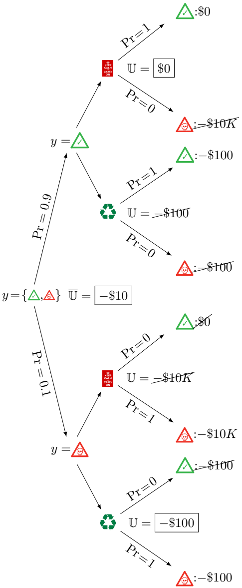

Value of information

The **value of information** represents how much you are willing to pay for an information.

What if the information is not perfect?

Value of imperfect information

$$\overline{\mathbb{U}}(\tilde{a}^*) \equiv \mathbb{E}[\mathbb{U}(\tilde{a}^*, X)] = \sum_{y \in \mathcal{X}} \max_{a} \Big(\sum_{x \in \mathcal{X}} \mathbb{U}(a, x) \cdot \underbrace{p(y|x) \cdot p(x)}_{p(y,x)} \Big)$$

VOI – Soil contamination example

$$\overline{\mathbb{U}}(\tilde{a}^*) = \sum_{y\in\mathcal{X}} \max_{a}\left(\sum_{x\in\mathcal{X}} \mathbb{U}(a,x)\cdot \underbrace{p(y|x)\cdot p(x)}_{p(x,y)}\right)$$

$$= \underbrace{\underbrace{\$0\times 0.9}_{x=\text{△}} + \underbrace{\text{-}\$10K\times 0.005}_{x=\text{⚠}}}_{y=\text{△}} + \underbrace{\underbrace{\text{-}\$100\times 0}_{x=\text{△}} + \underbrace{\text{-}\$100\times 0.095}_{x=\text{⚠}}}_{y=\text{⚠}}$$

$$= \boxed{-\$59.5}$$

$\mathbb{U}(a,x)$	$x=$ △	$x=$ ⚠
$a=$ ♻	$-\$100$	$-\$100$
$a=$ KEEP CALM CARRY ON	$0\$$	$-\$10K$

$$\begin{aligned}\Pr(y=\text{△}|\text{△}) &= 1\\ \Pr(y=\text{⚠}|\text{⚠}) &= 0.95\end{aligned}$$

$p(x,y)$	$y=$ △	$y=$ ⚠
$x=$ △	$0.9\cdot 1 = 0.9$	$0.9\cdot 0 = 0$
$x=$ ⚠	$0.1\cdot 0.05 = 0.005$	$0.1\cdot 0.95 = 0.095$
$p(y)$	0.905	0.095

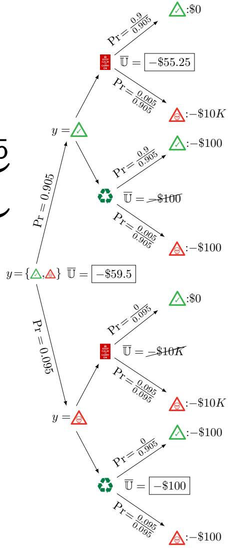

VOI – Soil contamination example (cont.)

Expected utility for the optimal action

$$\overline{\mathbb{U}}(a^*) \equiv \mathbb{E}[\mathbb{U}(a^*, X)] = -100\$$$

Expected utility conditional on imperfect information

$$\overline{\mathbb{U}}(\tilde{a}^*) = -\$59.5$$

Value of imperfect information

$$VOI(y) = \overline{\mathbb{U}}(\tilde{a}^*) - \overline{\mathbb{U}}(a^*) = \boxed{40.5\$}$$

Summary

Rational Decision:

Choose the **action** a_i^* which minimize the expected **loss** $\mathbb{L}(a,x)$ or maximizes the expected **utility** $\mathbb{U}(a,x)$

$$a^* = \arg\min_a \mathbb{E}[\mathbb{L}(a,X)] = \arg\max_a \mathbb{E}[\mathbb{U}(a,X)]$$

$\mathbb{L}(v(a,x))$ & $\mathbb{U}(v(a,x))$: Subjective weight on value as a function of the attitude toward risks
($\pm$ **1\$ not the same effect if you have 1\$ or 1M\$**)

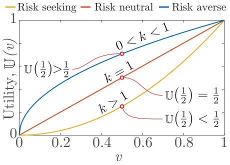

$$\mathbb{U}(v) = v^k \begin{cases} k > 1 & \text{Risk seeking} \\ k = 1 & \text{Neutral} \\ 0 < k < 1 & \text{Risk averse} \end{cases}$$

Value of information:

Value you should be willing to pay for information

$$VOI(y) = \mathbb{E}[\mathbb{U}(\tilde{a}^*,X)] - \mathbb{E}[\mathbb{U}(a^*,X)] \geq 0$$

Value of perfect information:

$$\overline{\mathbb{U}}(\tilde{a}^*) \equiv \mathbb{E}[\mathbb{U}(\tilde{a}^*,X)] = \sum_{x\in\mathcal{X}} \max_a [\mathbb{U}(a,x)] \cdot \Pr(x)$$

Value of imperfect information:

$$\overline{\mathbb{U}}(\tilde{a}^*) \equiv \mathbb{E}[\mathbb{U}(\tilde{a}^*,X)] = \sum_{y\in\mathcal{X}} \max_a \Big(\sum_{x\in\mathcal{X}} \mathbb{U}(a,x) \cdot \underbrace{p(y|x)\cdot p(x)}_{p(y,x)}\Big)$$