

About this class

Basics of utility theory (Chapter 16 of Russell & Norvig)

Cost-sensitive classification

Utility Functions

A utility function is a mapping from states of the world to real numbers

Under uncertainty, we typically work with expected utilities

Actions have nondeterministic outcomes, but essentially:

$$S \times A \rightarrow S$$

(the actual function would map to a set of $\langle S, \mathbb{R} \rangle$ tuples).

$$EU(A) = \sum_i P(\text{Result}_i|A)U(\text{Result}_i)$$

Principle of maximum expected utility says that the rational action for an agent to take is the

one that maximizes the agent's expected utility.

Does that solve one branch of AI?

No, we still have to define the problem appropriately and solve the utility maximization problem

Why MEU?

Why not minimize worst possible loss?

Why do we need a utility function?

What defines a reasonable utility function?

Constraints on rational preferences.

Notation: $A \succ B$ means A is strictly preferred to B , $A \sim B$ means the agent is indifferent between A and B , $A \succeq B$ means A is weakly preferred to B

In the general case, A and B are *lotteries* which are probability distributions over sets of outcomes $[p_1, S_1; p_2, S_2, \dots, p_n, S_n]$

The axioms of utility theory:

1. **Orderability:** $A \succ B$ or $B \succ A$ or $A \sim B$
2. **Transitivity:** If $A \succ B$ and $B \succ C$ then $A \succ C$
3. **Continuity:** If $A \succ B \succ C$ then $\exists p$ such that $B \sim [p, A; 1 - p, C]$
4. **Substitutability:** If $A \sim B$ then $[p, A; 1 - p, C] \sim [p, B; 1 - p, C]$
5. **Monotonicity:** If $A \succ B$ then $p \geq q \iff [p, A; 1 - p, B] \succsim [q, A; 1 - q, B]$
6. **Decomposability:** Compound lotteries can be reduced to simpler ones using the laws of probability.

$$[p, A; 1 - p, [q, B; 1 - q, C]] \sim [p, A; (1 - p)q, B; (1 - p)(1 - q), C]$$

Von-Neumann and Morgenstern showed that given the axioms, two things follow:

1. There exists a real-valued function U over states such that $U(A) > U(B)$ iff $A \succ B$ and $U(A) = U(B)$ iff $A \sim B$
2. The utility of a lottery is the sum of the probability of each outcome times the utility of that outcome

Human-Beings and Utility

Common model: log utility function (note that this gives us risk aversion naturally)

Linear utility functions imply risk neutrality

Rationality? Choose a bet: A is 80% chance of \$4000 and B is 100% chance of \$3000.

Now choose another: C is 20% chance of \$4000 and D is 25% chance of \$3000

Consistent?

Cost-Sensitive Classification

Making optimal decisions based on a cost or utility matrix.

	Pred. Negative	Pred. Positive
Act. Negative	U_{00}	U_{01}
Act. Positive	U_{10}	U_{11}

Suppose we compute, for a given x , the probability that $Y = 1$, call this p

When should we predict that $Y = 1$?

$$\text{If } pU_{11} + (1 - p)U_{01} > pU_{10} + (1 - p)U_{00}$$

Example, suppose my utility matrix for Spam prediction is as follows:

	Pred. NotSpam	Pred. Spam
Act. NotSpam	5	-5
Act. Spam	0	2

$$pU_{11} + (1 - p)U_{01} = 2p + (1 - p)(-5) = 7p - 5$$

$$pU_{10} + (1 - p)U_{00} = (1 - p)5$$

$$7p - 5 > 5 - 5p$$

$$\Rightarrow p > 10/12$$