Computing Science (CMPUT) 455

Search, Knowledge, and Simulations

James Wright

Department of Computing Science
University of Alberta
james.wright@ualberta.ca

Fall 2021

455 Today - Lecture 3

Topics:

- Introduction to decision-making
- Optimal decision-making
- Some models of human decision-making

Coursework

- TA Office hours began this week
  - See website for times and meeting links
- Start coding Assignment 1
- Quiz 2, review Reinforcement Renaissance
- Read Heingartner, Maybe We Should Leave That Up to the Computer.
- Activities for Lecture 3

Lecture 3:
Problem Solving and Decision-making

Lecture Topics (1)

- What is decision-making?
- Models of the world, reward, and utility
- How to evaluate alternatives in decision-making?
- Exact evaluation, expected values

Lecture Topics (2)

- How do humans make decisions?
- Heuristics, Bounded Rationality, and Satisficing
- What is the “right” decision for a program to make?
- Kahneman and Tversky experiments, criticism of utility theory

Decision Making in Humans and Machines

Image source:

`blogs-images.forbes.com/mikemyatt/files/`

`2012/11/decision-making-processes1.jpg`

Decision making is studied in many fields

- Business
- Psychology
- Advertising
- Computing Science
- AI
- ...

Decision Making in Humans and Machines (2)

- Decision making in politics can have far-reaching consequences (war, peace, prosperity, ...)
- Decision making is big business -
  what to buy, sell, produce,...
- Decision making is studied by many people in many different ways
  - “Common sense”
  - Academic and industry research
  - Popular “how to” books
- We make decisions every day. How and why?

Decision Making in Humans and Machines (3)

Some big questions:

- Can we make better decisions?
- Can we understand and influence other people's decisions?
- Can we teach decision-making -
  to children, students, employees?
- Can we model decision-making in a computer program?

Decision Theory (Theory of Choice)

Two main strands of research:

- *Normative* decision theory
  - Analyze decision problem
  - Tell user what is best action
  - Example:
    "You should play e4, it is the best move"

Decision Theory (Theory of Choice)

Two main strands of research:

- *Normative* decision theory
  - Analyze decision problem
  - Tell user what is best action
  - Example:
    "You should play e4, it is the best move"
- *Descriptive* decision theory
  - Analyze how real agents (people, programs?) make decisions
  - Example:
    "In our user study, 55% played move d4, 32% played e4, and 13% played some other move.
    The reasons are: ..."

Game Theory and Expected Value

- Classical game theory (e.g. von Neumann and Morgenstern 1947)
- Selfish players, try to maximize their money (*)
- Simplest case: two player zero sum games
- Zero sum - my win is your loss
- Actions can involve random outcomes, but with known probabilities
- Goal: maximize *expected value*

Expected Value

- Concept from probability theory
- Random event, with n different outcomes
- Each outcome evaluated by a number value v_i (reward, money, ...)
- Probability p_i of each outcome known
  - $\sum_{i=1}^{n} p_i = 1$
- Expected value (EV):
  - $\sum_{i=1}^{n} p_i v_i$

Expected Value Example

- Throw a six-sided fair die
- Value = the number rolled
- $n = 6$
- $p_1 = p_2 = p_3 = p_4 = p_5 = p_6 = 1/6$
- $v_1 = 1, v_2 = 2, v_3 = 3, v_4 = 4, v_5 = 5, v_6 = 6$
- Expected value
- $\sum_{i=1}^{n} p_i v_i = 1/6(1 + 2 + 3 + 4 + 5 + 6) = 21/6 = 3.5$

Expected Value Example

- Throw a six-sided fair die
- Value = the number rolled
- $n = 6$
- $p_1 = p_2 = p_3 = p_4 = p_5 = p_6 = 1/6$
- $v_1 = 1, v_2 = 2, v_3 = 3, v_4 = 4, v_5 = 5, v_6 = 6$
- Expected value
- $\sum_{i=1}^{n} p_i v_i = 1/6(1 + 2 + 3 + 4 + 5 + 6) = 21/6 = 3.5$
- Question: what is the EV when rolling two dice?

Example for Expected Value - Fold or Bid?

- Assume you play a simple card game, and all you care about is maximizing money
- Assume you have two possible actions, fold or bid
- Fold, you lose $1 for sure
- Bid, you either win $5 or lose $3

Example for Expected Value - Fold or Bid?

- Assume you play a simple card game, and all you care about is maximizing money
- Assume you have two possible actions, fold or bid
- Fold, you lose $1 for sure
- Bid, you either win $5 or lose $3
- Question: What should you do?

Example for Expected Value - Fold or Bid?

- Assume you play a simple card game, and all you care about is maximizing money
- Assume you have two possible actions, fold or bid
- Fold, you lose $1 for sure
- Bid, you either win $5 or lose $3
- Question: What should you do?
- Question: What *additional information* do you need to answer that question?

Example for Expected Value - Fold or Bid?

- Assume you play a simple card game, and all you care about is maximizing money
- Assume you have two possible actions, fold or bid
- Fold, you lose $1 for sure
- Bid, you either win $5 or lose $3
- Question: What should you do?
- Question: What *additional information* do you need to answer that question?
- The *probability of winning* if you bid

Fold or Bid?

- Set p to be your probability of winning if you bid
- Fold: your value is -1.
- Bid: your value is
  - +5 with probability p
  - -3 with probability $1 - p$
- Which action is better *in expectation*?

Fold or Bid? - Analysis

- Bid: +5 with probability p
- Bid: -3 with probability $1 - p$
- Expected value after bid: $5p + (-3)(1 - p) = 8p - 3$
- When is this better, worse, or equal to -1 (Fold)?
- Depends on p, compare $8p - 3$ with -1
- When are they equal? Solve equation $8p - 3 = -1$
- Solution $p = 1/4$

Fold or Bid? - Solution

- When $p = 1/4$, you are *indifferent*
  - *Expected value (EV)* of both actions, fold and bid, is the same
  - Confirm EV for bid:

$$5 \times \frac{1}{4} + (-1) \times (1 - \frac{1}{4}) = \frac{5}{4} - \frac{9}{4} = -1$$

Fold or Bid? - Solution

- When $p = 1/4$, you are *indifferent*
  - *Expected value (EV)* of both actions, fold and bid, is the same
  - Confirm EV for bid:

$$5 \times \frac{1}{4} + (-1) \times (1 - \frac{1}{4}) = \frac{5}{4} - \frac{9}{4} = -1$$

- When p grows, $8p - 3$ also grows
- For $p > 1/4$, bidding is better
- For $p < 1/4$, folding is better

Fold or Bid? - Example

Example: when $p = 1/3$, you want to bid

- EV(fold) = -1
- EV(bid) = $1/3 \times 5 + 2/3 \times -3 = 5/3 - 6/3 = -1/3$, better than folding

Fold or Bid? - Scaling Up

- The analysis before was probably reasonable for most people, describes the "most rational" choice
- What happens if we scale it up?
  - Instead of -1, +5, -3 dollars,
    play with -10000, 50000, -30000

Fold or Bid? - Scaling Up

- The analysis before was probably reasonable for most people, describes the “most rational” choice
- What happens if we scale it up?
  - Instead of -1, +5, -3 dollars,
    play with -10000, 50000, -30000
- Question: Is optimizing expected value the “most rational” strategy *now*?

Fold or Bid? - Scaling Up

- Some people would hate to lose \$10000 without even trying
- Some people can lose \$10000 without horrible consequences, but not \$30000
- Some people would value winning \$50000 very highly
- Our **utility** of money does not always scale linearly with the amount of money
- It depends on how it affects our life

St. Petersburg Paradox (Nicolas Bernoulli, 1713)

A paradox about expected value vs. actual behavior of people

- Play a game against the bank:
- The bank puts $2 in the pot originally
- Each round you flip a coin
- If head, the bank doubles the pot
- If tail, the game ends and you win the whole pot

St. Petersburg Paradox (Nicolas Bernoulli, 1713)

A paradox about expected value vs. actual behavior of people

- Play a game against the bank:
- The bank puts $2 in the pot originally
- Each round you flip a coin
- If head, the bank doubles the pot
- If tail, the game ends and you win the whole pot
- Q1: What is your expected value for this game?

St. Petersburg Paradox (Nicolas Bernoulli, 1713)

A paradox about expected value vs. actual behavior of people

- Play a game against the bank:
- The bank puts $2 in the pot originally
- Each round you flip a coin
- If head, the bank doubles the pot
- If tail, the game ends and you win the whole pot
- Q1: What is your expected value for this game?
- Q2: How much would you pay to be allowed to play this game?

Let's Play St. Petersburg Paradox...

- Start with $2
- Head - double
- Tail - game over
- See Python code `petersburg.py`, `petersburg2.py`
- Short demo now.

St. Petersburg Paradox Analysis

- probability 1/2, win $2 - tail
- probability 1/4, win $4 - head, tail
- probability 1/8, win $8 - head, head, tail
- probability 1/16, win $16 - head, head, head, tail
- ...
- Expected value of your win
  $1/2 \times 2$
  $+1/4 \times 4$
  $+...$
  $= 1 + 1 + ... = \infty$
- How much would *you* pay to play?

St. Petersburg Paradox vs. Reality

- The expected value is infinity but:
- It includes mostly extremely unlikely events
- Example:
  - Chance of 1/1,024 to win $1,024
  - Chance of 1/1,048,576 to win $1,048,576
  - Chance of 1/1,099,511,627,776 to win $1,099,511,627,776
- How to evaluate those in practice?

Utility

- Utility is a concept from economics
- Measures satisfaction of a consumer with an outcome (e.g., receiving a specific good)
- Utility of a good? It determines the price that a consumer is willing to pay
- But - what is the utility of money?
- Q: Is twice the money twice as desirable?

Utility

- Utility is a concept from economics
- Measures satisfaction of a consumer with an outcome (e.g., receiving a specific good)
- Utility of a good? It determines the price that a consumer is willing to pay
- But - what is the utility of money?
- Q: Is twice the money twice as desirable?
- In general, no.

Utility Function and Risk

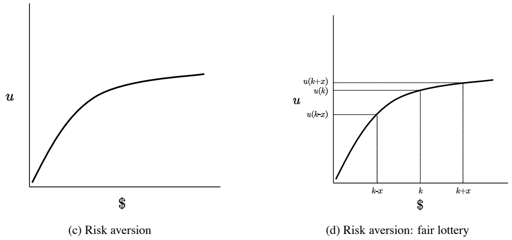

Image source: (Shoham & Leyton-Brown, 2008)

- *Utility function for money* is a mapping:
  - From: Monetary outcome
  - To: utility scale reflecting personal preferences
- Linked with types of behavior:
  - Risk-averse (conservative)
    - Utility function grows slower than linear
  - Risk-neutral
  - Risk-seeking (e.g. playing lottery)
    - Utility function grows faster than linear

Marginal Utility

Image source:

`http://s3.crackedcdn.com/articleimages/dan/rags/gates3.jpg`

- Marginal utility: increase in consumer satisfaction from having one unit more of a good
- Example: what is the value of having $100 more?
  - Very high if you are broke
  - Very low if you are Bill Gates
  - Marginal utility of money generally decreases with wealth

Car Example

Image source: `shedsunlimited.net`

Another example:

- What is the marginal utility of owning one more car?
- High if you have no car
- Much lower if you already own 3

Car Example

Image source: `shedsunlimited.net`

Another example:

- What is the marginal utility of owning one more car?
- High if you have no car
- Much lower if you already own 3

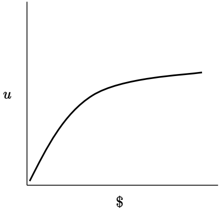

(c) Risk aversion

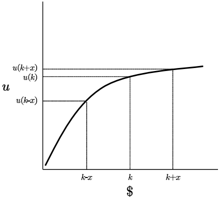

(d) Risk aversion: fair lottery

Image source: (Shoham & Leyton-Brown, 2008)

Maximum Expected Utility (MEU)

- Principle of maximum expected utility (Ramsey; von Neumann / Morgenstern):
- Choose action which maximizes your expected utility:

Maximum Expected Utility (MEU)

- Principle of maximum expected utility
  (Ramsey; von Neumann / Morgenstern):
- Choose action which maximizes your expected utility:
- With known probabilities, we can compute
  expected utilities just as we computed expected value
- Just replace the values with the utilities in the formula

$$\begin{aligned} a^* &= \arg\max_{a \in \mathcal{A}} \mathbb{E}[u(v_i) \mid a] \\ &= \arg\max_{a \in \mathcal{A}} \sum_{v_i} u(v_i) \Pr[v_i \mid a] \end{aligned}$$

Maximum Expected Utility (MEU)

- Principle of maximum expected utility (Ramsey; von Neumann / Morgenstern):
- Choose action which maximizes your expected utility:
- With known probabilities, we can compute expected utilities just as we computed expected value
- Just replace the values with the utilities in the formula

$$
\begin{aligned}
a^* &= \arg\max_{a \in \mathcal{A}} \mathbb{E}[u(v_i) \mid a] \\
&= \arg\max_{a \in \mathcal{A}} \sum_{v_i} u(v_i) \Pr[v_i \mid a]
\end{aligned}
$$

- Expected Utility Hypothesis: under certain conditions, people behave in that way...
- Q: Do you think people *actually* act like this?

Maximum Expected Utility (MEU)

- Principle of maximum expected utility (Ramsey; von Neumann / Morgenstern):
- Choose action which maximizes your expected utility:
- With known probabilities, we can compute expected utilities just as we computed expected value
- Just replace the values with the utilities in the formula

$$
\begin{aligned}
a^* &= \arg\max_{a \in \mathcal{A}} \mathbb{E}[u(v_i) \mid a] \\
&= \arg\max_{a \in \mathcal{A}} \sum_{v_i} u(v_i) \Pr[v_i \mid a]
\end{aligned}
$$

- Expected Utility Hypothesis: under certain conditions, people behave in that way...
- Q: Do you think people *actually* act like this?
- Q: Do you think expected utility maximization is necessarily the *best* (most rational) way to act? Why or why not?

Estimating Probabilities, Risk and Insurance

- Insurance companies charge an insurance fee...
- ... in return for promising to reimburse you for a low-probability large loss
- How to come up with a “fair” insurance premium?
- Need to know all the risks - bad things that could happen - and their probabilities
- Typically, only specific types of risk are covered by a policy
  - Home insurances usually exclude war, water damage, some types of natural and human-made disasters, . . .
- Impossible to estimate singular events, “black swans”

Heuristics

Image source:

`http://archimedespalimpsest.`

`org/images/kaltoon/4.php`

- From Greek word “find” or “discover” (wikipedia)
- Any practical problem solving method
- “Mental shortcut”, rule of thumb, educated guess, common sense rule, a rough model, using a similar case for guidance, ...

Heuristic vs Exact

Image source:

`http://archimedespalimpsest.org/images/kaltoon/4.php`

- Opposites of heuristic approach: exact solution, exhaustive analysis, precise theory
- We rely on heuristics all the time
- Most of life is too complex to "solve exactly"
- Heuristic decision-making and exact methods are both used in computer programs

Heuristics in Computing Science

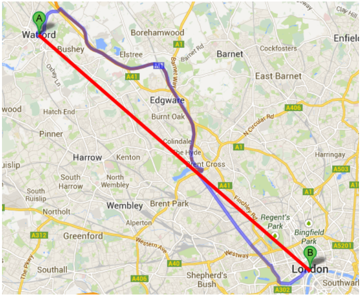

Image source:

`https://stackoverflow.com/`

`questions/17594924/`

- Typical heuristic in CS:
  solution to simplified problem
- Example problem:
  find shortest path from A to B
- Real solution:
  follow roads, avoid obstacles
- Heuristic: straight-line distance

Polya - How to Solve It

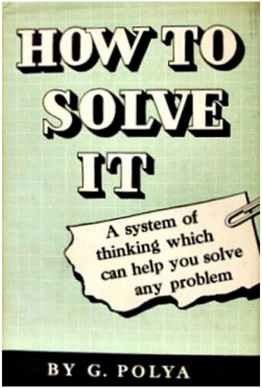

Image source:

`en.wikipedia.org/wiki/File:`

`HowToSolveIt.jpg`

- A system for human problem-solving
  - Classic book by mathematician George Polya
  - Published in 1945
  - Still popular and influential
- Four principles
- Large set of heuristics

How to Solve It - Four Principles

1. Understand the problem
2. Devise a plan
   - Find connection between data and unknown
3. Carry out the plan
4. Looking back
   - Examine the solution, review/extend

- Polya's book focuses on problem-solving mathematics
- We "translate" some ideas for CS

Polya - How to Solve It - Inside Cover

HOW TO SOLVE IT

First.

You have to *understand* the problem.

Second.

Find the connection between the data and the unknown. You may be obliged to consider auxiliary problems if an immediate connection cannot be found. You should obtain eventually a *plan* of the solution.

Third.

Carry out your plan.

Fourth.

Examine the solution obtained.

UNDERSTANDING THE PROBLEM

- *What is the unknown? What are the data? What is the condition?*
- Is it possible to satisfy the condition? Is the condition sufficient to determine the unknown? Or is it insufficient? Or redundant? Or contradictory?
- Draw a figure. Introduce suitable notation.
- Separate the various parts of the condition. Can you write them down?

DEVISING A PLAN

- Have you seen it before? Or have you seen the same problem in a slightly different form?
- *Do you know a related problem?* Do you know a theorem that could be useful?
- *Look at the unknown!* And try to think of a familiar problem having the same or a similar unknown.
- *Here is a problem related to yours and solved before. Could you use it?* Could you use its result? Could you use its method? Should you introduce some auxiliary element in order to make its use possible?
- Could you restate the problem? Could you restate it still differently?
- Go back to definitions.
- If you cannot solve the proposed problem try to solve first some related problem. Could you imagine a more accessible related problem? A more general problem? A more special problem? An analogous problem? Could you solve a part of the problem? Keep only a part of the condition, drop the other part; how far is the unknown then determined, how can it vary? Could you derive something useful from the data? Could you think of other data appropriate to determine the unknown? Could you change the unknown or the data, or both if necessary, so that the new unknown and the new data are nearer to each other?
- Did you use all the data? Did you use the whole condition? Have you taken into account all essential notions involved in the problem?

CARRYING OUT THE PLAN

- Carrying out your plan of the solution *check each step*. Can you see clearly that the step is correct? Can you prove that it is correct?

LOOKING BACK

- Can you *check the result?* Can you check the argument?
- Can you derive the result differently? Can you see it at a glance?
- Can you use the result, or the method, for some other problem?

Principle 1: Understand the Problem

Format: *Polya's text in italic* - comments below

- *What are the data?*
  - What is given? What is the input?
- *What is the unknown?*
  - What is the output?
- *What is the condition?*
  - What are the requirements/constraints for the solution?

Principle 1: Understand the Problem (continued)

- *Draw a figure. Introduce suitable notation.*
  - Draw or write down the important concepts in the problem and their relations.

- *Separate the various parts of the condition*
  - Find smaller parts, functions that make up the required solution

Principle 2: Devise a Plan

Polya gives a list of general approaches to try. Examples:

- *Find connection between data and unknown*
  - How do you compute the output as a function of the input?

- *Have you seen it before? Do you know a related problem?*
  - Can you re-use the previous solution?

- *Could you restate the problem?*
  - Is there a different way to write it,
    which is more similar to things you know?

Principle 2: Devise a Plan (continued)

- *If you cannot solve the proposed problem try to solve first some related problem*
  - Solve a special case
  - Solve a concrete example
  - Drop the complicated parts for now

- *Did you use all the data?*
  - Are you using everything you know?
  - The whole specification?
  - All properties of the input?

Principle 3: Carry out the Plan

- *Carrying out your plan of the solution check each step*
  - Write functions to implement your program, test each one separately.
  - Use unit tests to help verify that each function works as expected, at least on the test cases.

- *Can you see clearly that the step is correct?*
  Can you prove that it is correct?
  - Use assertions in your code to make sure input and output are as you expect. For really tricky code, you can even try a formal proof with pre- and postconditions and loop invariants (Cmput 204 stuff).

Principle 4: Looking Back

- *Can you check the result?*
  - Examine the solution
  - Review the problem in all details, check with your solution

- *Can you use the result, or the method, for some other problem?*
  - Refactor code, simplify functions, clean up
  - Extend or generalize functions for other problems
  - Organize into modules

How to Solve It - Heuristics

- Dictionary of heuristics is largest part of book
- Over 60 entries
- Some are specific to mathematical problem solving
- Most are generally useful
- Next two slides show examples

Polya - Example of Heuristic

Auxiliary problem

- Find an easier problem that will help solve the original
- Example: useful helper function
- Solve problem in several small steps, each implemented in a simpler function

Polya - Example of Heuristic

Decomposing and recombining

- Break a big problem into parts
- Find out which parts are important
- Solve parts
- Put together solutions of parts
- Examples:
    - Separate UI from engine
    - Floodfill - separate scan of full board from what to do in each area
    - Separate tree search algorithm from details of what to do in each node

Herb Simon and Bounded Rationality

Image source:

`http://www.cs.cmu.edu/simon/`

- Herb Simon (1916 - 2001)
  - One of founders of AI (and other disciplines)
  - Nobel-prize winner
  - Professor at Carnegie-Mellon
- Original background: decision-making in business, economics
- Criticized “perfect rationality” assumption of previous theorists
- Developed influential concept of “Bounded Rationality”

Activity: Watch Herb Simon Videos

See activities course page

- Video 1: The Limits or Bounds of Rationality
- Video 2: What is Intuition?
- Optional - read more about Herb Simon: `https://en.wikipedia.org/wiki/Herbert_A._Simon`

Decision-making and Optimization

- Mathematical optimization:
  find the **best possible** solution to a problem

Decision-making and Optimization

- Mathematical optimization:
  find the **best possible** solution to a problem
- Simple example:
  what is the minimum of function $x^2 - 5x + 3$?

Decision-making and Optimization

- Mathematical optimization:
  find the **best possible** solution to a problem
- Simple example:
  what is the minimum of function $x^2 - 5x + 3$?
- Harder example:
  How many regular size soccer balls can we pack into a standard shipping container?

Decision-making and Optimization

- Mathematical optimization:
  find the **best possible** solution to a problem
- Simple example:
  what is the minimum of function $x^2 - 5x + 3$?
- Harder example:
  How many regular size soccer balls can we pack into a standard shipping container?
- Much harder example:
  what should my company produce to maximize its profit?

Decision-making and Optimization

- Mathematical optimization:
  find the **best possible** solution to a problem
- Simple example:
  what is the minimum of function $x^2 - 5x + 3$?
- Harder example:
  How many regular size soccer balls can we pack into a standard shipping container?
- Much harder example:
  what should my company produce to maximize its profit?
- Even harder:
  what should the company produce to make the most people the happiest?

Decision-making vs solving an Optimization Problem

- Is Decision-making an Optimization Problem?
- Answer: sometimes...

Decision-making vs solving an Optimization Problem

- Is Decision-making an Optimization Problem?
- Answer: sometimes...
- Yes:
  - Decision means optimizing some quantity such as money, grade average, “reward”

Decision-making vs solving an Optimization Problem

- Is Decision-making an Optimization Problem?
- Answer: sometimes...
- Yes:
  - Decision means optimizing some quantity such as money, grade average, “reward”
- No:
  - Decision involves many factors that are hard to compare

Decision-making vs solving an Optimization Problem

- Is Decision-making an Optimization Problem?
- Answer: sometimes...
- Yes:
  - Decision means optimizing some quantity such as money, grade average, “reward”
- No:
  - Decision involves many factors that are hard to compare
- Example:
  study more, or get more sleep?

Can We Make Perfect Decisions?

- Herb Simon:
  - Most often, no.
- Why?
- Humans (and computers) have:
  - Limited memory
  - Limited time to make a decision
  - Incomplete, or wrong, information about actions and results
  - Limited powers of logic, deduction, lookahead
  - Limited imagination to come up with new approaches
  - Limited everything ...

Bounded Rationality - Discussion

- Perfect decisions often not possible in practice
- How can we act well,
  while acknowledging our limitations?
- How can we use what we know,
  and even what we don't know?
- How do we deal with multiple, conflicting goals?
- When should we use heuristics,
  and when a more systematic search?
- What is the "best" thing to do, given our limitations?
  Is that even well-defined?

Exact vs Good Enough, Satisficing

- Humans use heuristics as shortcuts
- Concept of satisficing (Herb Simon)
- Trying to optimize is often too hard
- More reasonable:
    - Define criteria for “good enough”
- A satisficing solution is one that fulfills these criteria
- Example in games:
  play a “good” or “strong” move,
  even if we cannot prove it is the best

Herb Simon on Satisficing

"Decision makers can satisfice
either by finding optimum solutions
for a simplified world,
or
by finding satisfactory solutions
for a more realistic world.
Neither approach, in general, dominates the other, and
both have continued to co-exist..."

Optimum Solutions for a Simplified World

Image source: `www.`

`simpleitsolutions.com`

- Example: what is the cost of buying a new computer?
- Simple answer: look at the price tag
- More complete answer: add tax, cost of new software, cost of time for upgrades, electricity, insurance, carrying bag, ...
- In practice, we ignore or roughly summarize many of these details and make a decision for a simplified problem
- Some costs are not known anyway, e.g. future costs of electricity, repairs, ...

Satisfactory solutions for a more realistic world

Image source:

`thehealthysubstitute.`

`wordpress.com`

- Example: what to eat for lunch?
- A myriad of choices
- Many small or large variations are possible (seasoning, extras, ...)
- In real life, we only consider a small number of choices
- We (usually) satisfice, not optimize

Comment - Model vs Direct Observation

- Remember Herb Simon's quote:

  Decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world.

- Two approaches to reasoning for decision-making:
- Reasoning based on a **model** of the world
- Reasoning from **direct observations** of the world
- Big topic in Reinforcement Learning
- In games, we have a *perfect* model. But that is the exception, not true for real world

Kahneman and Tversky

Image source: `www.vanityfair.com/news/2016/11/`

- Very influential psychologists
- Kahneman won Nobel prize in economics
- Humans have systematic *cognitive biases*
- Most are averse to loss and ambiguity, “losses loom larger than gains”
- Activity: Watch Daniel Kahneman Videos

Kahneman and Tversky - Anchoring

Image source: `https://www.jeremysaid.com/blog/anchoring-effect-power-conversion-optimization/`

- People tend to “anchor” on first impressions
- Later decisions made relative to this, not in absolute terms
- People focus more on *changes* in their utility than on *absolute* utilities

Anchoring - Another Car Example

Anchoring - Another Car Example

Scenario A

- Monday, I offer to sell you my car for \$30000.
- Tuesday, I offer it to you for \$20000.

Anchoring - Another Car Example

Scenario A

- Monday, I offer to sell you my car for $30000.
- Tuesday, I offer it to you for $20000.

Scenario B (same car)

- Monday, I offer to sell you my car for $10000.
- Tuesday, I offer it to you for $20000.

Anchoring - Another Car Example

Scenario A

- Monday, I offer to sell you my car for $30000.
- Tuesday, I offer it to you for $20000.

Scenario B (same car)

- Monday, I offer to sell you my car for $10000.
- Tuesday, I offer it to you for $20000.

Question: In which scenario are you more likely to accept the offer on Tuesday?

Summary

- Quick tour of theories and experiments in human decision-making
- How do we make decisions?
- Limits to making "perfect" decisions
- Bounded rationality and satisficing
- Expected value, expected utility
- Cognitive biases
- Next time:
    - Formal models of decision-making in sequential games
    - Representing games