# DMR Worked Examples

Yordan Hristov

# Preliminaries

# Preliminaries

- Easy to digest set of problems for the sake of understanding the underlying principle
- All problems are about decision making (some under uncertainty).
- We always have a way to measure decision outcomes - utility function.
- The approaches shown are not mutually exclusive.

# Dynamic Programming Recap

# Problem 1 - Medics

# Problem 1 - Medics

The *World Health Council* is devoted to improving healthcare in the underdeveloped countries in the world. If now has five medical teams available to allocate among three such countries to improve their medical care, health education and training programs. Therefore the Council needs to determine how many teams (if any) to allocate to each of these countries to maximise the total effectiveness of the five teams.

The measure of effectiveness being used is *additional man-years of life*. (For a particular country, this measure equals the country's *increased life expectancy* in years times its population.) The table on the next slide fives the estimated additional man-years of life (in multiples of 1000) for each country for each possible allocation of medical teams.

# Problem 1 - Medics

| # medical teams | Thousand of additional man-years of life | | |
|---|---|---|---|
| | Country 1 | Country 2 | Country 3 |
| 0 | 0 | 0 | 0 |
| 1 | 45 | 20 | 50 |
| 2 | 70 | 45 | 70 |
| 3 | 90 | 75 | 80 |
| 4 | 105 | 110 | 100 |
| 5 | 120 | 150 | 130 |

- **What is the best split of teams?**

- Three stages for 3 countries: $n = 3$
- State of the system: $s$ = # of teams still available
- Decision variable: $x_n$ = # of teams assigned to country n (state n)
- $p_i(x_i)$ measures the effectiveness from assigning $x_i$ medical teams to country i - in the table.

# Problem 2 - Scientists

# Problem 2 - Scientists

A government space project is conducting research on a certain engineering problem that must be solved before man cal fly to Mars. Three research teams are currently trying three different approaches for solving this problem. The estimate has been made that, under present circumstances, the probability that the respective teams - call them 1, 2 and 3 - will not succeed is 0.40, 0.60, and 0.80 respectively. Thus, the current probability that all three teams will fails is (0.4)x(0.6)x(0.8) = 0.192. Since the objective is the minimize this probability, the decision has been made to assign two more top scientist among the three teams to lower it as much as possible.

The table in the next slide gives the estimated probabilities that the respective teams will fails when 0, 1 or 2 additional scientists are added to that team. The problem is to determine how to allocate the two additional scientists to minimize the probability that all three teams will fail.

# Problem 2 - Scientists

| # of new scientists | Probability of failure | | |
|---|---|---|---|
| | Team 1 | Team 2 | Team 3 |
| 0 | 0.4 | 0.6 | 0.8 |
| 1 | 0.2 | 0.4 | 0.5 |
| 2 | 0.15 | 0.2 | 0.3 |

- **What is the best split of scientists between teams?**

- P(all teams fail) = 0.192
- Three stages for 3 teams: n = 3
- State of the system: s = # of scientists still available
- Decision variable: $x_n$ = # of scientists assigned to team n (state n)
- $p_i(x_i)$ measures the probability of failure for team i if $x_i$ scientists are assigned that team - in the table.

# Problem 3 - Carnival Planning

# Problem 3 - Carnival Planning

A carnival is scheduled to appear in a city on a given date. The profits obtained are heavily dependent on the weather. In particular, if the weather is rainy, the carnival looses \$15,000; if cloudy, the carnival looses \$5,000 and if sunny, the carnival makes a profit of \$10,000.

The carnival has to set up equipment for its show but can cancel the show prior ti setting up its equipment. This action results in a loss of \$1000. Furthermore, the carnival can postpone its decision until the day before the scheduled performance and request a local weather report at the additional cost of \$1,000. The Weather Bureau has compiled data based on its predictions (bottom left table on the next slide). Furthermore, the Bureau has compiled a prior distribution of the weather - top right table on the next slide.

Note that there are two actions - set-up ($a_1$) and not set-up ($a_2$). There are three states of the weather (w) and the forecast (f): Rain (r), Clouds (c), and Sun (s). The respective utilities for each action in the context of a particular weather condition are presented in the top left table on the next slide.
The task is to find the sequence of actions which yield the highest expected utility (in this case equivalent to profit).

# Problem 3 - Carnival Planning

| Utility Function for the Carnival Organisers | | | |
|---|---|---|---|
| U(a, w) | w = r | w = c | w = s |
| $a_1$ - set up | -$15K | -$5K | $10K |
| $a_2$ - not set up | -$1K | -$1K | -$1K |

| | P(w) |
|---|---|
| w = r | 0.1 |
| w = c | 0.3 |
| w = s | 0.6 |

| Weather Bureau Data | | | |
|---|---|---|---|
| P(f \| w) | w = r | w = c | w = s |
| f = r | 0.7 | 0.2 | 0.1 |
| f = c | 0.2 | 0.6 | 0.2 |
| f = s | 0.1 | 0.2 | 0.7 |

- Request weather report = -$1K
- What is the best plan of action?
- Calculate expected utility of
  - No report - $U_{nwr}$
  - Report - $U_{wr}$

  and decide how to act

# Decision Trees Recap

# Decision Trees

- A classical way of representing decision problems with several decisions is with decision trees. A decision tree is a model that encodes the structure of the decision problem by representing all possible sequences of decisions and observations explicitly in the model.
- The non-leaf nodes in a decision tree are decision nodes (rectangular boxes) or chance nodes (circles or ellipses), and the leaves are utility nodes (diamond shaped). The links in the tree have labels. A link from a decision node is labeled with the action chosen, and a link from a chance node is labeled by a state.
- A decision tree is read from the root downward. When you pass a decision node, the label tells you what the decision is, and when you pass a chance node, the label tells you the state of the node. If a decision node follows a chance node, then the chance node is observed before the decision is made. Hence the sequence in which we visit the nodes corresponds to the sequence of observations and decisions.


Decision node (non-leaf) | Chance node (non-leaf) | Utility node (leaf) | Link labelled with action a

# Decision Trees

- Each path from the root to a leaf specifies a complete sequence of observations and decisions, and we call such a sequence a decision scenario. Furthermore, we require decision trees to be complete: from a chance node there must be a link for each possible state, and from a decision node there must be a link for each possible decision option. This also means that a decision tree specifies all the possible scenarios in the decision problem.
- A solution to a decision tree is a strategy that specifies how we should act at the various decision nodes. The optimal strategies in the following examples are illustrated by the boldfaced links. Strategies are compared based on their expected utilities, and finding an optimal strategy amounts to finding a strategy with highest expected utility.
- One way for calculating an optimal strategy is a procedure known as “average-out and fold-back”.


Decision node (non-leaf) | Chance node (non-leaf) | Utility node (leaf) | Link labelled with action a

# Decision Trees

**Average-out and fold-back algorithm:**

Let X be a node in a decision tree T . To calculate an optimal strategy and the maximum expected utility for the subtree rooted at X, do:

1. If X is a utility node, then return U(X).
2. If X is a chance node, then return

   $EU(X) = \sum_{x \in sp(X)} P(X = x \mid past(X))\, EU(N(X = x))$.

3. If X is a decision node, then return

   $EU(X) = \max_{x \in sp(X)} EU(N(X = x))$,

   and mark the arc labeled $x' = \arg\max_{x \in sp(X)} EU(N(X = x))$.

A decision tree for the carnival problem is provided in the solutions.

# Problem 4 - Car Start Problem

# Problem 4 - Car Start Problem

In the morning, my car will not start. There are three possible faults:

- the spark plugs may be dirty, with probability 0.3;
- the ignition system may be malfunctioning, with probability 0.2;
- or there is some other cause, with probability 0.5.

I can perform two repair actions myself:

- SP, which at the cost of 4 minutes always fixes spark plugs;
- and IS, which takes 2 minutes and fixes the ignition system with probability 0.5.
- I can perform a test T , to check the charge on the spark plugs when starting. It takes half a minute, and it says ok if and only if the ignition system is okay.
- Finally, I can call road service RS, which at the cost of 15 minutes fixes everything.

What is the best plan for action - the one in which the expected delay is minimal?

# Acknowledgements

Examples and images were taken from the following resources:

1. Operations Research, Frederick S. Hiller, Second Edition
2. Bayesian Decision and Decision Graphs, Finn V. Jensen
3. Probabilistic Robotics, Sebastian Thrun

# Thanks. If you have questions:

yordan.hristov@ed.ac.uk