EXPLOITING STRUCTURE IN DECISION MAKING UNDER THE LENS OF RECENT ADVANCES IN STARAI

MARCEL GEHRKE[1], NAZLI NUR KARABULUT, FLORIAN MARWITZ, AND TANYA BRAUN[2]

[1]Institute of Humanities-Centered Artificial Intelligence, University of Hamburg
[2]Computer Science Department, University of Münster

AGENDA

1. Introduction to Relational Models and Online Decision Making [Marcel]
   - Relational models under uncertainty
   - Lifted inference in decision-theoretic models (online decision making)
   - Markov Decision Process (Offline decision making)
2. Lifting Offline Decision Making [Flo]
3. Lifting Multi-Agent Decision Making [Nazlı Nur]
4. Summary [Marcel]

GENERAL AGENT SETTING

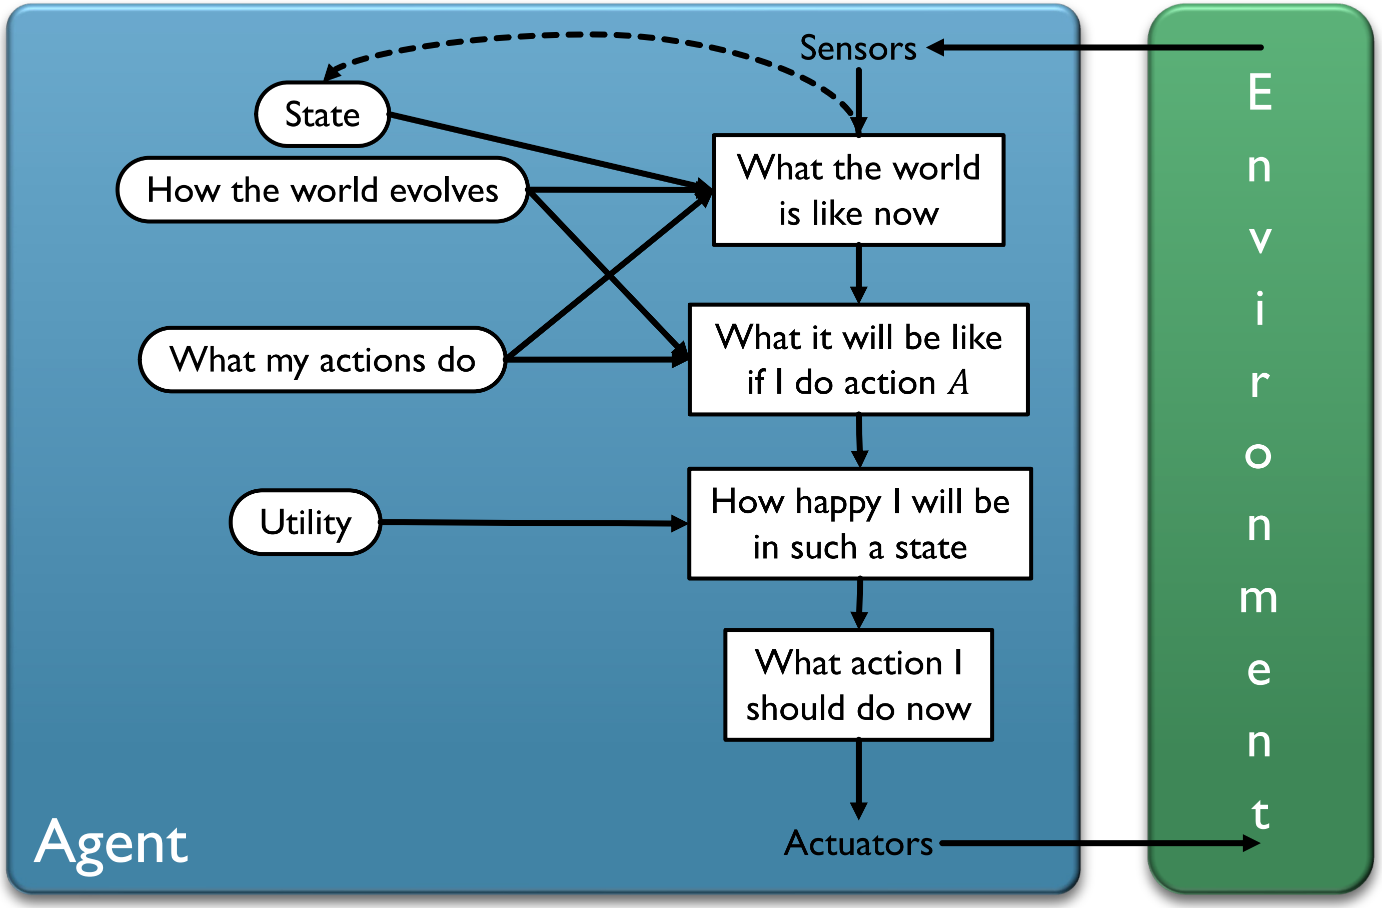

AIMA, Russell/Norvig

RELATIONAL MODELS UNDER UNCERTAINTY

INTRODUCTION

WHY RELATIONAL MODELS?

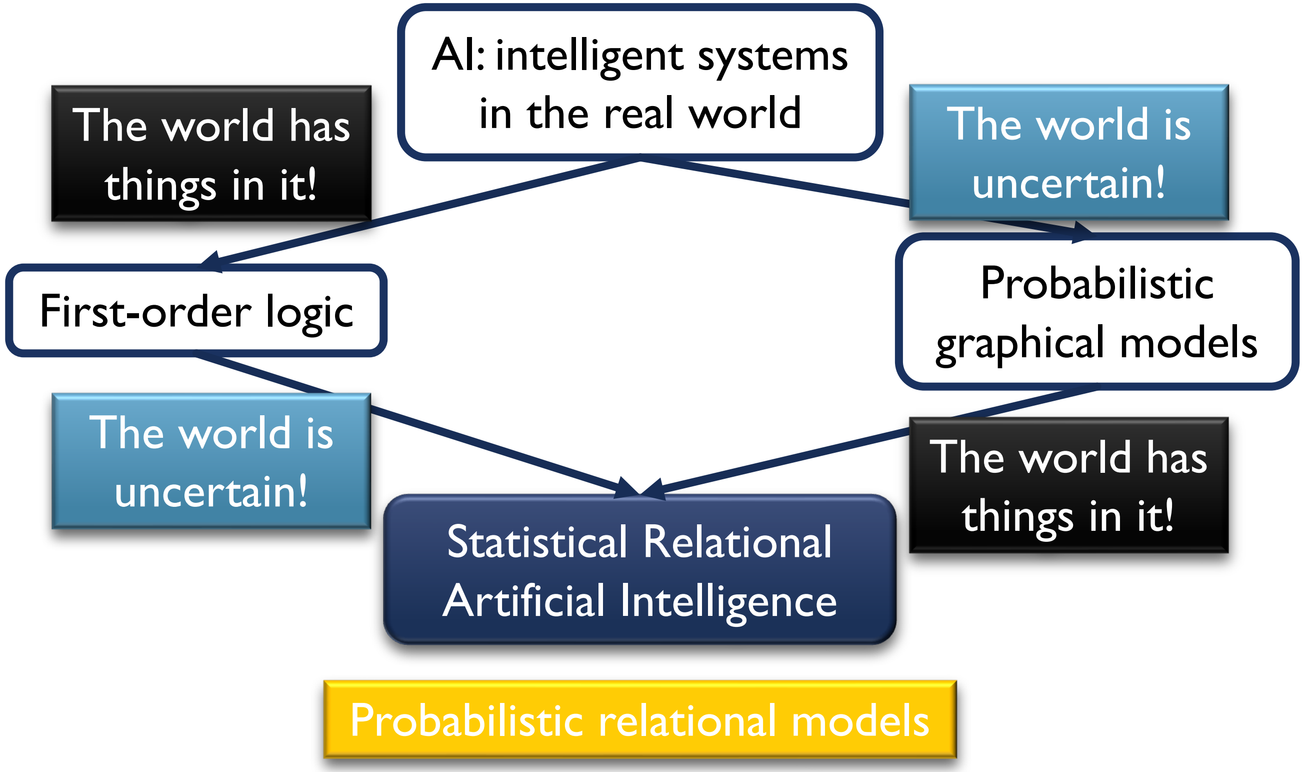

Figure based on a figure by Stuart Russell

LOGICAL VARIABLES IN RANDOM VARIABLES

- Atoms: Parameterised random variables = PRVs
  - With logical variables
    - E.g., X, M
    - Possible values (domain):

$$dom(X) = \{alice, eve, bob\}$$
$$dom(M) = \{injection, tablet\}$$

  - With range
    - E.g., Boolean, but any discrete, finite set possible
    - $ran\big(Travel(X)\big) = \{true, false\}$
  - Represent sets of *indistinguishable* random variables

$$Nat(D) = natural\ disaster\ D$$
$$Acc(A) = accident\ A$$

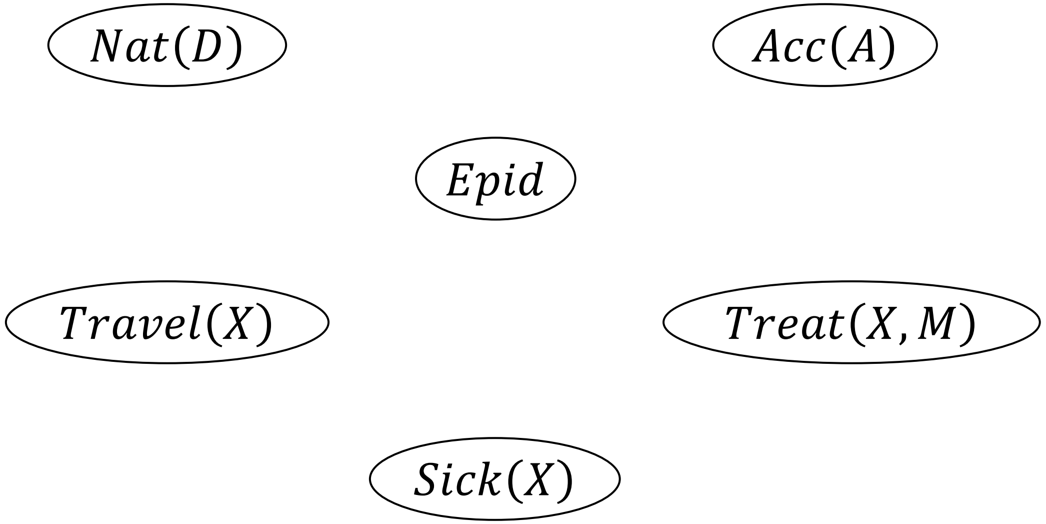

Poole 03

PARFACTORS

- Factors with PRVs = parfactors

$Travel(X)$	$Epid$	$Sick(X)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

Potentials

- In parfactors, just like in factors, no probability distribution as factors required

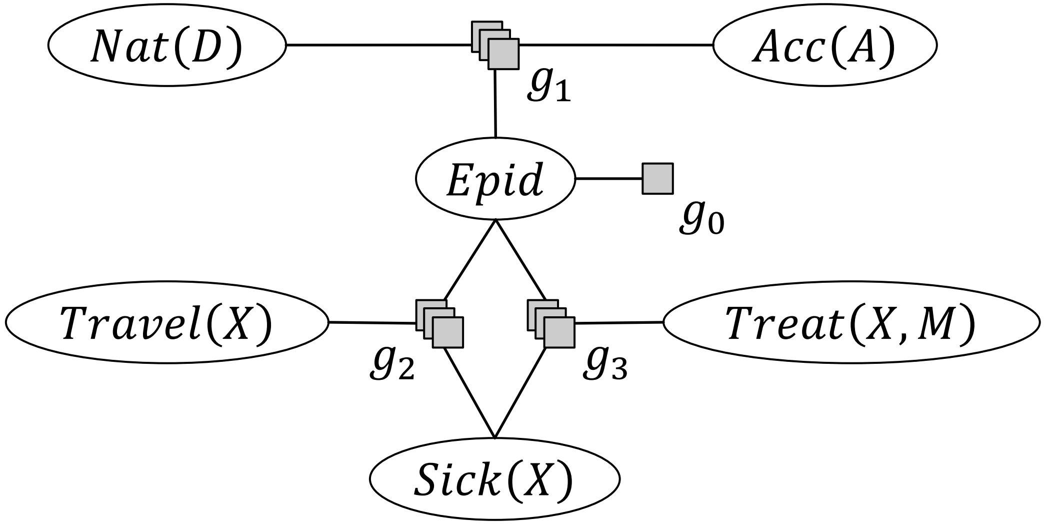

Poole 03

FACTORS

- Grounding

$Travel(X)$	$Epid$	$Sick(X)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

$Travel(eve)$	$Epid$	$Sick(eve)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$fals$		
$true$	tru		
$true$	tru		

$Travel(bob)$	$Epid$	$Sick(bob)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
		$true$	6
		$false$	4
		$true$	6
		$false$	2
		$true$	9

$Travel(alice)$	$Epid$	$Sick(alice)$	g_2
$false$	$false$	$false$	5
$false$	$false$	$true$	0
$false$	$true$	$false$	4
$false$	$true$	$true$	6
$true$	$false$	$false$	4
$true$	$false$	$true$	6
$true$	$true$	$false$	2
$true$	$true$	$true$	9

$reat(X,M)$

Poole 03

GROUNDED MODEL

- Given domains
  - $dom(X) = \{alice, eve, bob\}$
  - $dom(M) = \{m_1, m_2\}$
  - $dom(D) = \{flood, fire\}$
  - $dom(W) = \{virus, war\}$
- Symmetry in
  - Graph structure (isomorphism)
  - Factors (identity)

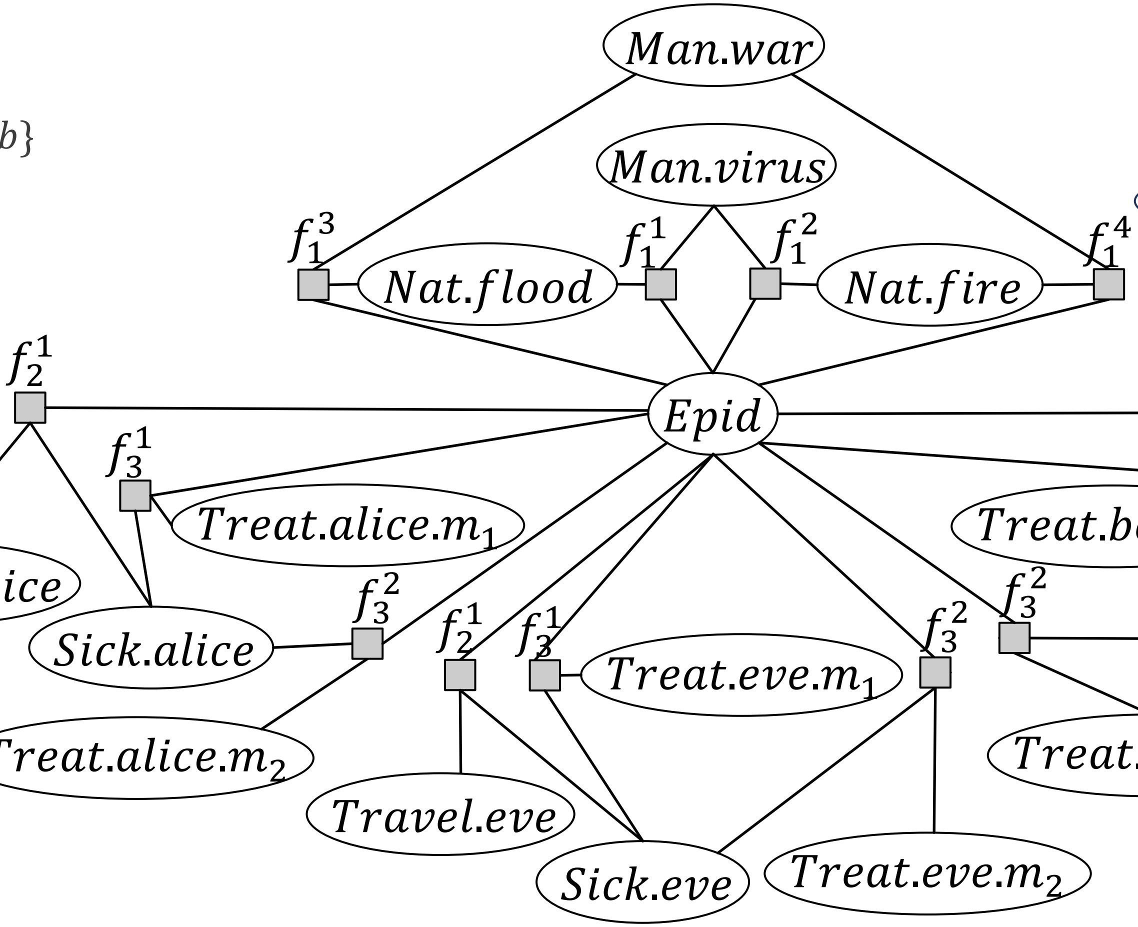

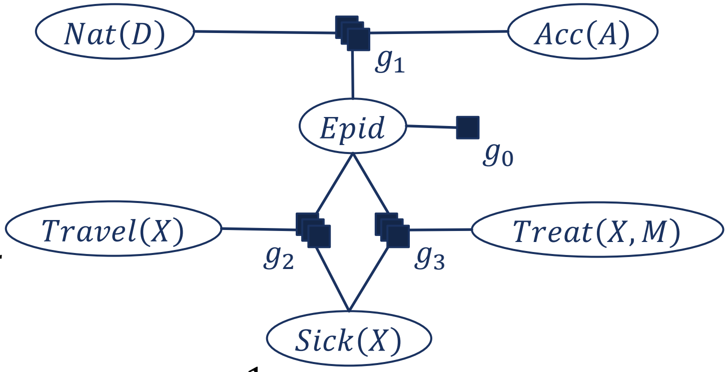

Poole 03

INFERENCE PROBLEMS WITH AND WITHOUT EVIDENCE

- Query answering problem given a model:
  - Probability of events
    - E.g., $P(Att(eve, ki) = true), P(Epid = true)$
  - Conditional (marginal) probability distributions
    - E.g., $P(Att(ev, ki) | FarAway(ki)), P(Epid | sick(alice), sick(eve))$
  - Assignment queries:
    - Most probable states of random variables
    - Most-probable explanation (MPE), Maximum a posteriori (MAP)
- Lifted inference:
  Work with representatives for
  exchangeable random variables
  - Avoid grounding for as long as possible

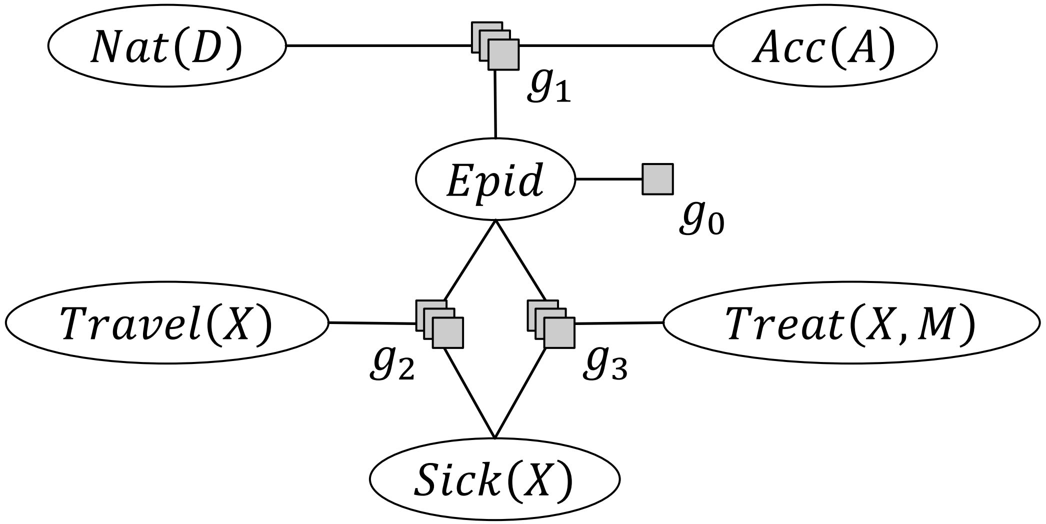

Poole 03

QA IN PARFACTOR MODELS: LIFTED VARIABLE ELIMINATION (LVE)

- Eliminate all variables not appearing in query
  - [Poole 03, de Salvo Braz et al. 05, 06, Milch et al. 08, Taghipour et al. 13, 13a, B & Möller 18]
- Lifted summing out
  - Sum out *representative* instance as in propositional variable elimination
  - Exponentiate result for exchangeable instances
- Correctness: Equivalent ground operation
  - Each instance is summed out
  - Result: factor f that is identical for all instance
  - Multiplying indistinguishable results
    → exponentiation of one representative f

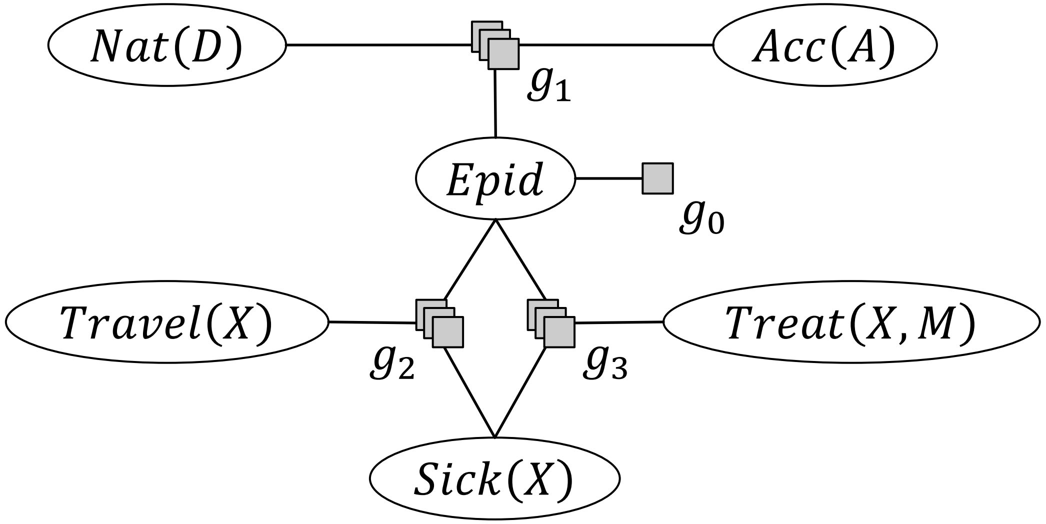

QA: LVE IN DETAIL

- Eliminate $Treat(X, M)$
  - Appears in only one g: g_3
  - Contains all logical variables of g_3: X, M
  - For each X constant: the same number of M constants
  - ✓ Preconditions of lifted summing out fulfilled, lifted summing out possible

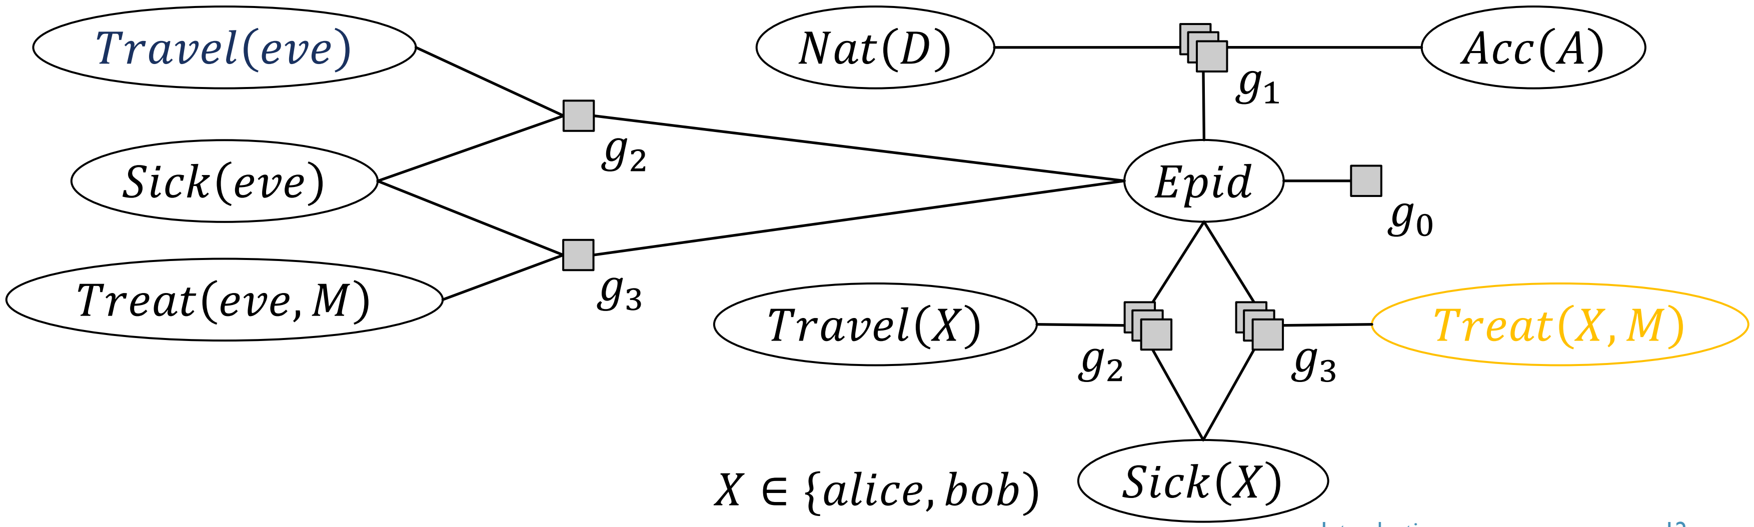

Poole 03

LVE IN DETAIL: LIFTED SUMMING OUT

- Eliminate $Treat(X, M)$ by lifted summing out
  1. Sum out representative

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X, M) = t)\right)^{\#M|X}$$

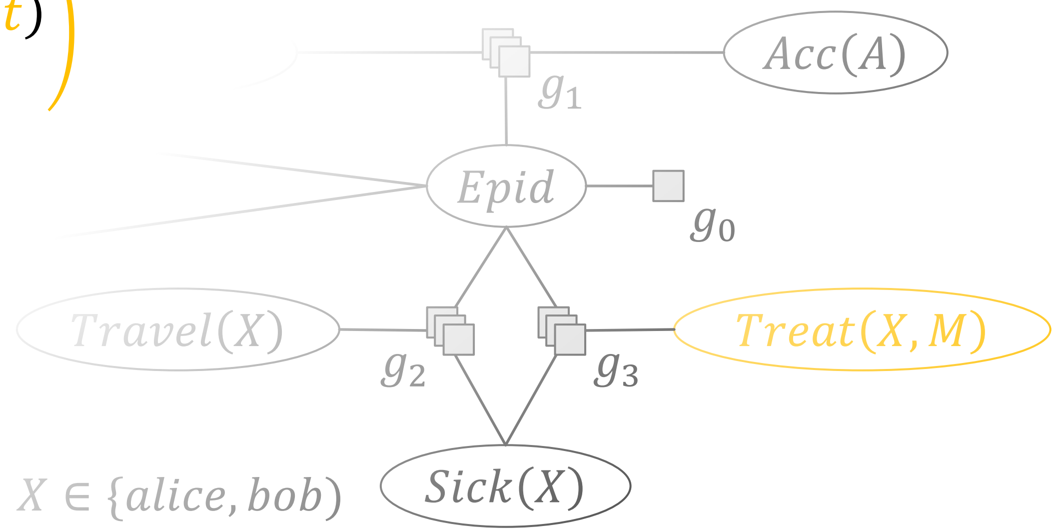

Poole 03

LVE IN DETAIL: LIFTED SUMMING OUT

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X,M) = t)\right)^{\#M|X}$$

$Epid$	$Sick(X)$	$Treat(X,M)$	g_3
$false$	$false$	$false$	9
$false$	$false$	$true$	1
$false$	$true$	$false$	6
$false$	$true$	$true$	3
$true$	$false$	$false$	7
$true$	$false$	$true$	5
$true$	$true$	$false$	4
$true$	$true$	$true$	8

$Epid$	$Sick(X)$	Σ
$false$	$false$	10
$false$	$true$	9
$true$	$false$	12
$true$	$true$	12

$Epid$	$Sick(X)$	^
$false$	$false$	10^2
$false$	$true$	9^2
$true$	$false$	12^2
$true$	$true$	12^2

SYMMETRIES WITHIN

- Assume four epidemics with identical characteristics
  - $Epid_1, Epid_2, Epid_3, Epid_4$
  - Reasonable to model the epidemics such that it does not matter which $Epid$ variables specifically are $true$ or $false$, i.e., they are interchangeable

    - All $false$ maps to 8
    - 1 $true$, 3 $false$ maps to 6
    - 2 $true$, 2 $false$ maps to 4
    - 3 $true$, 1 $false$ maps to 2
    - All $true$ maps to 0

    → Five lines enough to describe

$\#_{true}$, $\#_{false}$

[0,4]	8
[1,3]	6
[2,2]	4
[3,1]	2
[4,0]	0

$Epid_1$	$Epid_2$	$Epid_3$	$Epid_4$	ϕ
$false$	$false$	$false$	$false$	8
$false$	$false$	$false$	$true$	6
$false$	$false$	$true$	$false$	6
$false$	$false$	$true$	$true$	4
$false$	$true$	$false$	$false$	6
$false$	$true$	$false$	$true$	4
$false$	$true$	$true$	$false$	4
$false$	$true$	$true$	$true$	2
$true$	$false$	$false$	$false$	6
$true$	$false$	$false$	$true$	4
$true$	$false$	$true$	$false$	4
$true$	$false$	$true$	$true$	2
$true$	$true$	$false$	$false$	4
$true$	$true$	$false$	$true$	2
$true$	$true$	$true$	$false$	2
$true$	$true$	$true$	$true$	0

Milch 08

COUNTING RANDOM VARIABLE

- New PRV type:
  (Parameterised) counting random variable ((P)CRV)
  $\#_X[A_{|C}]$
  - $A_{|C}$ a PRV under constraint C
  - $X \in lv(A)$
  - Range values: Histogram $h = \{(v_i, n_i)\}_{i=1}^m$
    - $m = |ran(A)|$ (number of buckets)
    - $n = \sum_{i=1}^m n_i = |gr(A_{|\pi_X(C)})|$ (number of instances to distribute into buckets)
    - $v_i \in ran(A)$ (buckets)
    - $n_i \in \mathbb{N}$ (number of instances in bucket v_i)
    - Shorthand: $[n_1, \ldots, n_m]$
- Range of a (P)CRV = space of histograms fulfilling the conditions on the histograms
  - (All possible ways of distributing n interchangeable instances into m buckets)
- Single histogram encodes several interchangeable assignments at once
  - Given by multinomial coefficient $Mul(h)$

$$Mul(h) = \frac{n!}{\prod_{i=1}^m n_i!}$$

    - If $m = 2$, binomial coefficient:

$$\binom{n}{n_1} = \frac{n!}{(n-n_1)!\,n_1!} = \frac{n!}{n_2!\,n_1!}$$

Milch 08

CRV: EXAMPLE

- (P)CRV $\#_X[A_{|C}]$
  - Range values: Histogram $h = \{(v_i, n_i)\}_{i=1}^m$
    - $m = |ran(A)|$ (number of buckets)
    - $n = \sum_{i=1}^m n_i = |gr(A_{|C})|$ (number of instances to distribute into buckets)
    - $v_i \in ran(A)$ (buckets)
    - $n_i \in \mathbb{N}$ (number of instances in bucket v_i)
    - Shorthand: $[n_1, \dots, n_m]$
  - Single histogram encodes several interchangeable assignments at once:

$$Mul(h) = \frac{n!}{\prod_{i=1}^m n_i!}$$

- E.g., CRV: $\#_E[Epid(E)]$
  - $ran(Epid(E)) = \{true, false\} \rightarrow m = 2$
  - $dom(E) = \{e_1, e_2, e_3, e_4\} \rightarrow n = 4$
  - Range values and multiplicities

$$\{(true, 0), (false, 4)\} = [0,4] \quad Mul([0,4]) = \frac{4!}{0!\cdot 4!} = 1$$

$$\{(true, 1), (false, 3)\} = [1,3] \quad Mul([1,3]) = \frac{4!}{1!\cdot 3!} = 4$$

$$\{(true, 2), (false, 2)\} = [2,2] \quad Mul([2,2]) = \frac{4!}{2!\cdot 2!} = 6$$

$$\{(true, 3), (false, 1)\} = [3,1] \quad Mul([3,1]) = \frac{4!}{3!\cdot 1!} = 4$$

$$\{(true, 4), (false, 0)\} = [4,0] \quad Mul([4,0]) = \frac{4!}{4!\cdot 0!} = 1$$

Milch 08

CRV: EXAMPLE

- E.g., (continued)
  - CRV: $\#_E[Epid(E)]$
    - Range values
      [0,4], [1,3], [2,2], [3,1], [4,0]
      1 4 6 4 1
      how many assignments encoded
  - $g' = \phi(\#_E[Epid(E)])$

$\#_E[Epid(E)]$	ϕ'
[0,4]	8
[1,3]	6
[2,2]	4
[3,1]	2
[4,0]	0

$Epid_1$	$Epid_2$	$Epid_3$	$Epid_4$	ϕ
false	false	false	false	8
false	false	false	true	6
false	false	true	false	6
false	false	true	true	4
false	true	false	false	6
false	true	false	true	4
false	true	true	false	4
false	true	true	true	2
true	false	false	false	6
true	false	false	true	4
true	false	true	false	4
true	false	true	true	2
true	true	false	false	4
true	true	false	true	2
true	true	true	false	2
true	true	true	true	0

Milch 08

CRVS CONTINUED

- (P)CRV $\#_X[A_{|C}]$ with
  - $m = |ran(A)|$ (number of buckets)
  - $n = \sum_{i=1}^{m} n_i = |gr(A_{|\pi_X(C)})|$ (number of instances to distribute into buckets)
- Instead of m^n mappings in the ground factor, the counted factor has

$$\binom{n+m-1}{n-1}$$

mappings
  - Upper bound of range size of a CRV:

$$\binom{n+m-1}{n-1} \leq n^m$$

Exponential in number of random variables n

Polynomial in number of random variables n

Milch 08

LIFTED INFERENCE IN DECISION-THEORETIC MODELS (ONLINE DECISION MAKING)

INTRODUCTION

DECISION PRVS

$R(X)$	I	ϕ_4
$free$	$false$	1
$free$	$true$	0
ban	$false$	0
ban	$true$	1

$R(X)$	$Tl(X)$	ϕ_1
$free$	$false$	1
$free$	$true$	1
ban	$false$	1
ban	$true$	0

- Decision PRV D
  - Range $\text{ran}(D) = \{a_i\}_{i=1}^K$ set of possible actions
    - Actions a_i mutually exclusive (consistent with range definition)
    - Always have to get a value assigned
      - Cannot not make a decision!
  - Depicted by a rectangle in a graphical representation
  - E.g., travel restrictions for people X: $Restrict(X)$
    - Range values: $ban, free$
- Set of decision PRVs $\boldsymbol{D}$ in a model, i.e., $\boldsymbol{R} = \boldsymbol{D} \cup \boldsymbol{V}$
  - $\boldsymbol{D}$ can occur as arguments to any parfactor
  - Example:
    - $\phi_1(Restrict(X), Travel(X)), \phi_4(Restrict(X), Interference)$

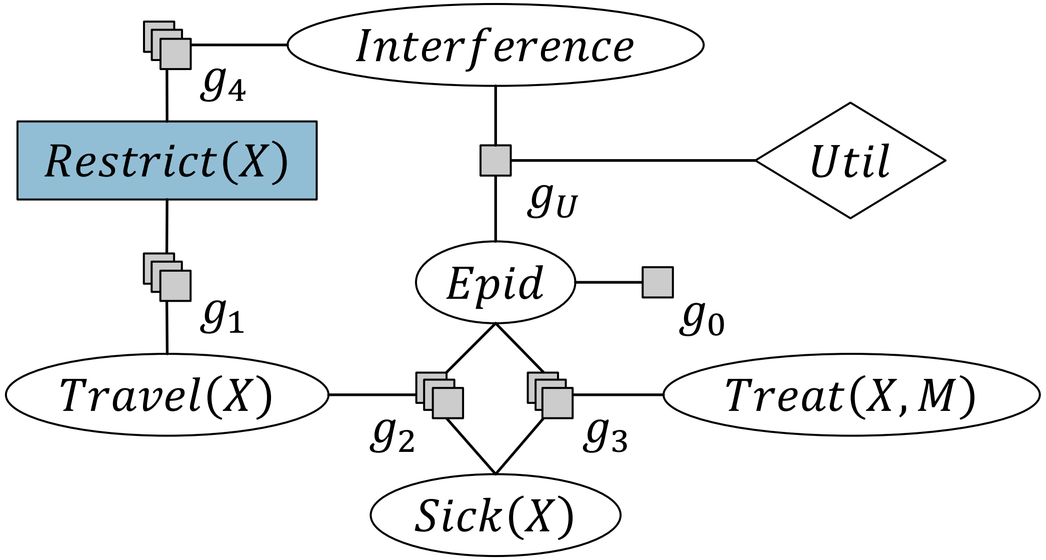

G. et al. 19

EXPECTED UTILITY QUERIES

- Given a decision model $G = \{g_i\}_{i=1}^{n} \cup \{g_U\}$
  - One can ask queries for (conditional) marginal distributions or events as before given an action assignment $\boldsymbol{d}$ based on the semantics, $P_G[\boldsymbol{d}]$
  - New query type: query for an expected utility (EU)
    - What is the expected utility of making decisions $\boldsymbol{d}$ in G?

$$eu(\boldsymbol{e}, \boldsymbol{d}) = \sum_{\boldsymbol{r} \in \mathrm{ran}(\mathrm{gr}(\mathrm{rv}(g_U) \setminus \boldsymbol{E} \setminus \boldsymbol{D}))} P(\boldsymbol{r} | \boldsymbol{e}, \boldsymbol{d}) \cdot \phi_U(\boldsymbol{r}, \boldsymbol{e}, \boldsymbol{d})$$

    - $P(\boldsymbol{r} | \boldsymbol{e}, \boldsymbol{d})$ means that the PRVs not occurring in this expression need to be eliminated accordingly
      - I.e., $\boldsymbol{V} = \mathrm{rv}(G) \setminus \boldsymbol{D} \setminus \boldsymbol{E} \setminus \mathrm{rv}(g_U)$

G. et al. 19

MEU PROBLEM

- Given a decision model G and evidence $\boldsymbol{e}$, maximum Expected Utility (MEU) problem:
  - Find the action assignment that yields the highest expected utility in G
  - Formally, $$\mathrm{meu}(G|\boldsymbol{e}) = \left(\boldsymbol{d}^*, eu(\boldsymbol{E}, \boldsymbol{d}^*)\right)$$
    $$\boldsymbol{d}^* = \arg\max_{\boldsymbol{d} \in \mathrm{ran}(\boldsymbol{D})} eu(\boldsymbol{e}, \boldsymbol{d})$$

Additive semantics with inner sum and outer max: Sum up utilities, then pick maximum → Max-sum algorithms

  - For an exact solution, $\mathrm{meu}(G|\boldsymbol{e})$ requires an algorithm to go through all $\boldsymbol{d} \in \mathrm{ran}(\boldsymbol{D})$
    - Size of $\mathrm{ran}(\boldsymbol{D})$ exponential in $|\boldsymbol{D}|$

Alternative specification
$$\mathrm{meu}(G|\boldsymbol{e}) = \left(\arg\max_{\boldsymbol{d} \in \mathrm{ran}(\boldsymbol{D})} eu(\boldsymbol{e}, \boldsymbol{d}), \max_{\boldsymbol{d} \in \mathrm{ran}(\boldsymbol{D})} eu(\boldsymbol{e}, \boldsymbol{d})\right)$$

AIMA, Russell/Norvig

DECISION MAKING OVER TIME

- Basis: a sequential model $(G^0, G^{\rightarrow})$
  - Describe behaviour over time using interslice parfactors
  - Within a slice, describe intra-slice (episodic) behaviour
- → Extend intra-slice parts with decision + utility PRVs
  - Intra-slice behaviour described using a decision model
  - Inter-slice behaviour allows for predicting effect of decision on next step

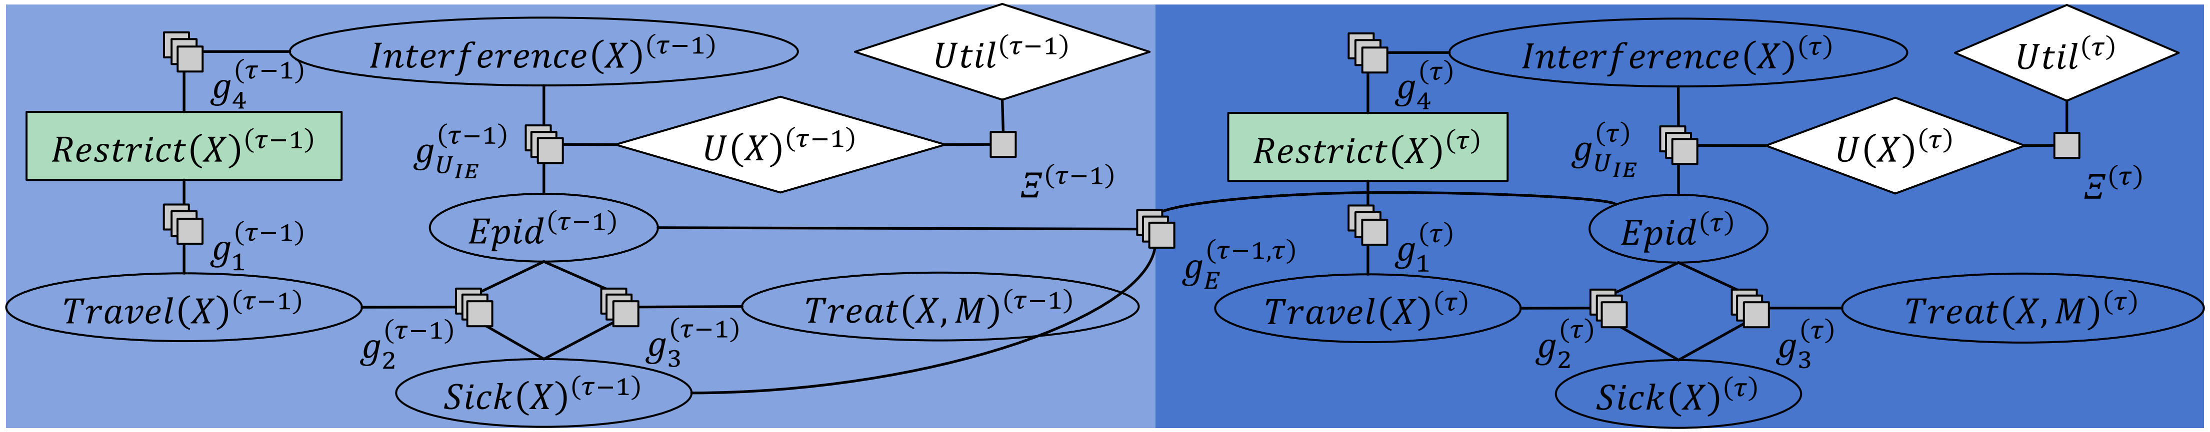

G., Braun & Möller 19

MARKOV DECISION PROCESS (OFFLINE DECISION MAKING)

INTRODUCTION

MARKOV DECISION PROCESS / PROBLEM (MDP)

- *Sequential* decision problem
  for a fully observable, stochastic environment
  with a Markovian transition model
  and additive rewards (next slide)
- MDP is a four-tuple (S, A, T, R) with
  - S a random variable whose domain is a set of states (with an initial state s^0)
  - For each $s \in \mathrm{dom}(S)$
    - a set $A(s)$ of actions
    - a transition model $T(s', s, a) = P(s'|s, a)$
    - a reward function $R(s)$ (also with a possible)
- Robot navigation example to the right

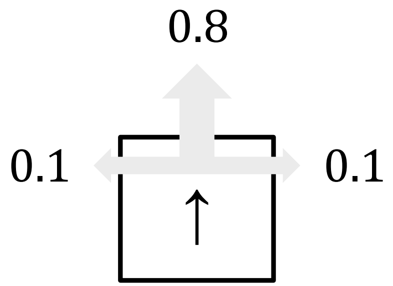

U, D, L, R

each move costs 0.04

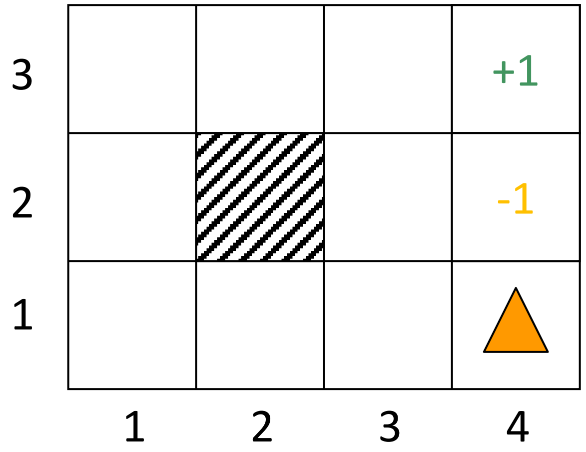

AIMA, Russell/Norvig

ADDITIVE UTILITY

- History $h = \left(s^{(0)}, s^{(1)}, \ldots, s^{(T)}\right)$
- In each state s, agent receives **reward** $R(s)$
- Utility of h is **additive** iff

$$U\left(s^{(0)}, s^{(1)}, \ldots, s^{(T)}\right) = R\left(s^{(0)}\right) + U\left(s^{(1)}, \ldots, s^{(T)}\right)$$
$$= \sum_{t=0}^{T} R\left(s^{(t)}\right)$$

  - **Discount** factor $\gamma \in]0,1]$:

$$U\left(s^{(0)}, s^{(1)}, \ldots, s^{(T)}\right) = \sum_{t=0}^{T} \gamma^t R\left(s^{(t)}\right)$$

    - Close to 0: future rewards insignificant
    - Corresponds to interest rate ${}^{1-\gamma}/_{\gamma}$

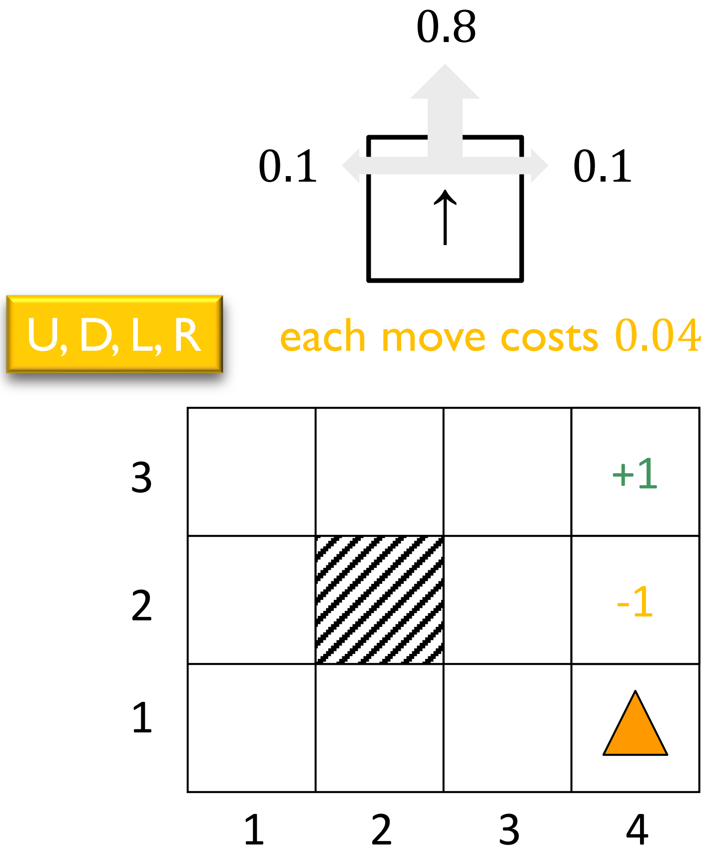

AIMA, Russell/Norvig

PRINCIPLE OF MEU

- Bellman equation:

$$U(s) = R(s) + \gamma \max_{a \in A(s)} \sum_{s' \in \mathrm{dom}(S)} P(s'|a,s)U(s')$$

- Optimal policy:

$$\pi^*(s) = \operatorname*{argmax}_{a \in A(s)} \sum_{s' \in \mathrm{dom}(S)} P(s'|a,s)U(s')$$

  - Bellman equation for $[1,1]$ with $\gamma = 1$ as discount factor
    - $U(1,1) = -0.04 + \gamma \max_{U,L,D,R}$

$$\begin{aligned} \{\, & 0.8U(1,2) + 0.1U(2,1) + 0.1U(1,1), && \text{(U)} \\ & 0.8U(1,1) + 0.1U(1,1) + 0.1U(1,2), && \text{(L)} \\ & 0.8U(1,1) + 0.1U(2,1) + 0.1U(1,1), && \text{(D)} \\ & 0.8U(2,1) + 0.1U(1,2) + 0.1U(1,1) \,\} && \text{(R)} \end{aligned}$$

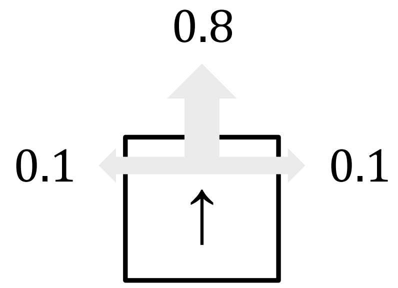

U, D, L, R

each move costs 0.04

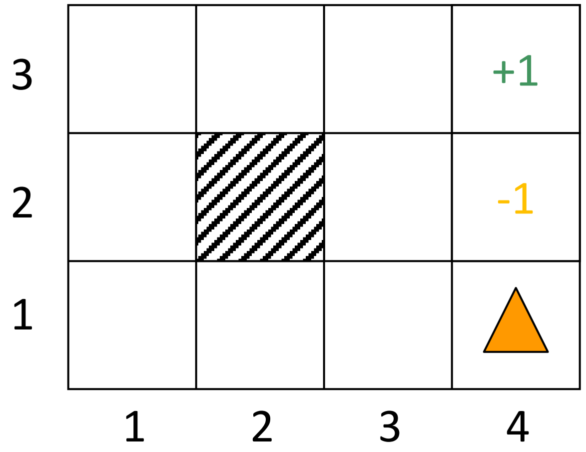

AIMA, Russell/Norvig

VALUE ITERATION

- Initialise the utility of each non-terminal state s to $U^{(0)}(s) = 0$
- For $t = 0, 1, 2, \ldots,$ do

$$U^{(t+1)}(s) \leftarrow R(s) + \gamma \max_{a \in A(s)} \sum_{s' \in \mathrm{dom}(S)} P(s'|a,s) U^{(t)}(s')$$

  - So called Bellman update

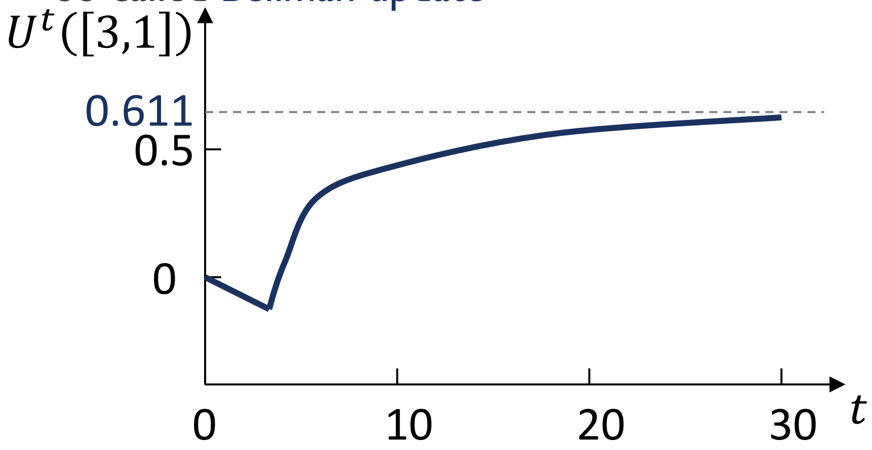

Note the importance of terminal states and connectivity of the state-transition graph

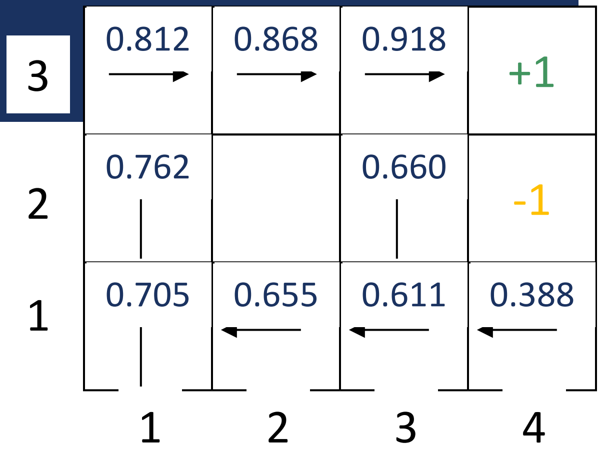

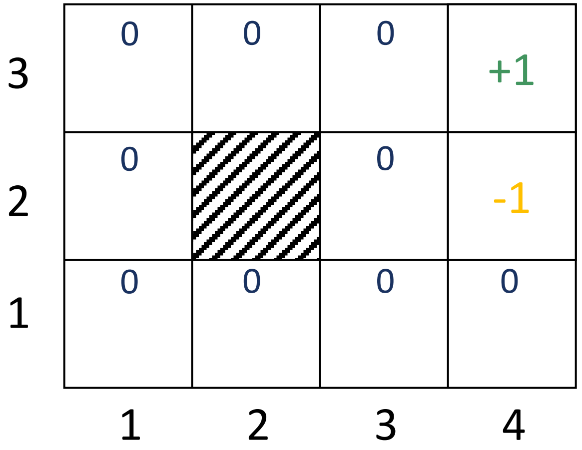

AIMA, Russell/Norvig

VALUE ITERATION: ALGORITHM

- Returns a policy π that is optimal
- Inputs
  - MDP mpd
    - Set of states S
    - For each $s \in S$
      - Set $A(s)$ of applicable actions
      - Transition model $P(s'|s,a)$
      - Reward function $R(s)$
  - Maximum error allowed ϵ

```
function value-iteration(mdp,ϵ)
    U′ ← 0, π ← ⟨⟩
    repeat
        U ← U′
        δ ← 0
        for each state s ∈ S do
            U′[s] ← R(s) + γ max_{a∈A(s)} Σ_{s′} P(s′|a.s) U[s′]
            if |U′[s] - U[s]| > δ then
                δ ← |U′[s] - U[s]|
    until δ < ϵ(1-γ)/γ
    for each state s ∈ S do
        π(s) ← argmax_{a∈A(s)} Σ_{s′} P(s′|a.s) U[s′]
    return π
```

- Local variables
  - U, U' vectors of utilities for states in S
  - δ maximum change in utility of any state in an iteration

AIMA, Russell/Norvig

POMDP

- POMDP = Partially Observable MDP
  - Sensing operation returns multiple states, with a probability distribution
    - Sensor model Ω that encodes $P(o|s)$ (or $P(o|s,a)$)
      - Probability of observing o given state s (and action a)
    - Example:
      - Sensing number of adjacent walls (1 or 2)
      - Return correct value with probability 0.9
  - Formally, POMDP is a six-tuple (S, A, T, R, O, Ω)
    - MDP (S, A, T, R) extended with a set of observations O and a sensor model Ω
  - Choosing action that maximizes expected utility of state distribution assuming "state utilities" computed as before not good enough → Does not make sense (not rational)
- POMDP agent: Constructing a new MDP in which the current probability distribution over states plays the role of the state variable

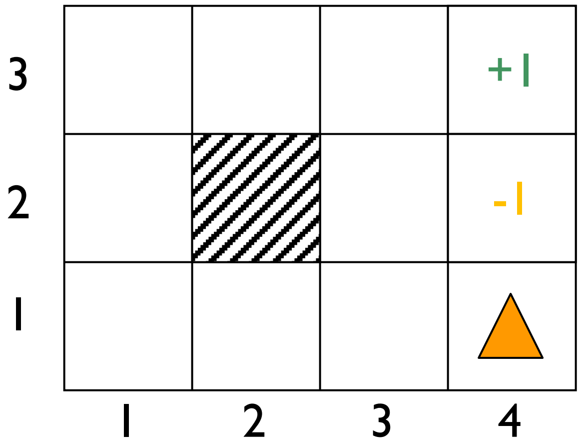

AIMA, Russell/Norvig

BELIEF STATE & UPDATE

- Belief state $b(s)$ is the probability assigned to the actual state s by belief state b
- Initial belief state
  - Probability of 0 for terminal states
  - Uniform distribution for rest
  - Robot navigation example:
    - $b = \left(\frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{9}, 0, 0\right)$

	1	2	3	4
3	$0.\overline{1}$	$0.\overline{1}$	$0.\overline{1}$	0.0
2	$0.\overline{1}$		$0.\overline{1}$	0.0
1	$0.\overline{1}$	$0.\overline{1}$	$0.\overline{1}$	$0.\overline{1}$

AIMA, Russell/Norvig

BELIEF STATE & UPDATE

- Update $b' = SE(b, a, o)$

$$b'(s') = P(s'|o,a,b) = \frac{P(o|s',a)\sum_{s\in \mathrm{dom}(S)} P(s'|s,a)b(s)}{\sum_{s''\in \mathrm{dom}(S)} P(o|s'',a)\sum_{s\in \mathrm{dom}(S)} P(s''|s,a)b(s)}$$

- Consider as two stage-update: (1) Update for the action (2) Update for the observation

b

	1	2	3	4
3	$0.\bar{1}$	$0.\bar{1}$	$0.\bar{1}$	0.0
2	$0.\bar{1}$		$0.\bar{1}$	0.0
1	$0.\bar{1}$	$0.\bar{1}$	$0.\bar{1}$	$0.\bar{1}$

Move L once

$b^{(1)}$

	1	2	3	4
3	0.2	$0.\bar{1}$	$0.0\bar{2}$	0.0
2	$0.\bar{1}$		$0.\bar{1}$	$0.0\bar{1}$
1	0.2	$0.\bar{1}$	$0.\bar{1}$	$0.0\bar{1}$

Perceive 1 wall

$b^{(2)} = b'$

	1	2	3	4
3	0.06569	0.03650	0.06569	0.0
2	0.03650		0.32847	0.03285
1	0.06569	0.03650	0.32847	0.00365

AIMA, Russell/Norvig

AGENDA

1. Introduction to Relational Models and Online Decision Making [Marcel]
   - Relational models under uncertainty
   - Lifted inference in decision-theoretic models (online decision making)
   - Markov Decision Process (Offline decision making)
2. Lifting Offline Decision Making [Flo]
3. Lifting Multi-Agent Decision Making [Nazlı Nur]
4. Summary [Marcel]

AGENDA

1. Introduction to Relational Models and Online Decision Making [Marcel]
2. Lifting Offline Decision Making [Flo]
   - Factored Markov Decision Processes
   - First-order Markov Decision Processes
   - Lifted Factored Markov Decision Processes
3. Lifting Multi-Agent Decision Making [Nazlı Nur]
4. Summary [Marcel]

BIBLIOGRAPHY

ORDERED ALPHABETICALLY

BIBLIOGRAPHY

- AMAI, Russel/Norvig
  Russell & Norvig: Artificial Intelligence: A Modern Approach. 2020
- Braun & Möller 18
  Tanya Braun and Ralf Möller: Parameterised Queries and Lifted Query Answering. In *IJCAI-18 Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018.
- De Salvo Braz et al. 05
  Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth: Lifted First-order Probabilistic Inference. In IJCAI-05 Proceedings of the 19th International Joint Conference on Artificial Intelligence, 2005.
- De Salvo Braz et al. 06
  Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth: MPE and Partial Inversion in Lifted Probabilistic Variable Elimination. In AAAI-06 Proceedings of the 21st Conference on Artificial Intelligence, 2006.

BIBLIOGRAPHY

- G. et al. 19

  Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In Proceedings of Artificial Intelligence in Health, 2019.

- G., Braun & Möller 19

  Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Maximum Expected Utility. In Canadian AI-19 Proceedings of the 32nd Canadian Conference on Artificial Intelligence, 2019.

- Milch et al. 08

  Brian Milch, Luke S. Zettelmoyer, Kristian Kersting, Michael Haimes, and Leslie Pack Kaelbling: Lifted Probabilistic Inference with Counting Formulas. In AAAI-08 Proceedings of the 23rd AAAI Conference on Artificial Intelligence, 2008.

BIBLIOGRAPHY

- Poole 03

  David Poole: First-order Probabilistic Inference. In IJCAI-03 Proceedings of the 18th International Joint Conference on Artificial Intelligence, 2003.

- Taghipour et al. 13

  Nima Taghipour, Daan Fierens, Guy Van den Broeck, Jesse Davis, and Hendrik Blockeel. Completeness Results for Lifted Variable Elimination. In AISTATS-13 Proceedings of the 16th International Conference on Artificial Intelligence and Statistics, 2013.

- Taghipour et al. 13a

  Nima Taghipour, Daan Fierens, Jesse Davis, and Hendrik Blockeel. Lifted Variable Elimination: Decoupling the Operators from the Constraint Language. Journal of Artificial Intelligence Research, 47(1):393–439, 2013.