Model-Based Reinforcement Learning

- **Incrementally Learn a Model of $P(s'|s,a)$ and $R(s'|s,a)$**
- **Apply Prioritized Sweeping after Each Action**

Data structures

- $n[s,a,s']$ = number of times we visited state s, performed action a, and moved to state s'.
- $N[s,a]$ = number of times we visited s and performed a.
- $r[s,a,s']$ = total reward received when we visited s, performed a, and moved to state s'
- Estimate $P(s'|s,a) = n[s,a,s']/N[s,a]$
- Estimate $R(s'|s,a) = r[s,a,s']/n[s,a,s']$

Online Algorithm

initialize Q(s,a) arbitrarily

n[s,a,s'] = 0

N[s,a] = 0

r[s,a,s'] = 0

repeat

choose action a according to exploration policy

perform a, observe r and s'

$n[s, a, s'] + +$, $N[s, a] + +$, $r[s, a, s'] + = r$

PQ.push(s, a, ∞)

PSweep(n, ϵ)

end

Model-Based vs. Model-Free Algorithms

- **Model Based Methods** are better when each experience is expensive and when randomized exploration methods have a low probability of reaching the goal.
- **Model Free Methods** are better when experience is cheap (e.g., simulated), and when randomized exploration can reach the goal.

Learning Complex Models

Often, there are important regularities in the models, that can allow them to be simplified.

- **number of cars at location 1** depends only on the action taken, the demand at location 1, and the number of returns at location 1.
- **number of cars at location 2** depends only on the action taken, the demand at location 2, and the number of returns at location 2.
- **demand** and **returns** are independent
- **reward** depends only on the sum of a term for the action, a term for location 1 and a term for location 2.

Learning will be much faster if we can learn a compact ("sparse") model of the environment, instead of a separate model for each (s, a, s').

This is why we will study Belief Networks and Influence Diagrams.

Goal: Find Sparse Models of Actions and Rewards

- **Define a set of random variables**
  - $n1$ number of cars at location 1
  - $n2$ number of cars at location 2
  - $d1$ demand for cars at location 1
  - $d2$ demand for cars at location 2
  - $s1$ cars returned at location 1
  - $s2$ cars returned at location 2
- **Identify the "causal" interactions among these variables**

  Express this as a graph structure.

Influence Diagram for Jack's Car Rental

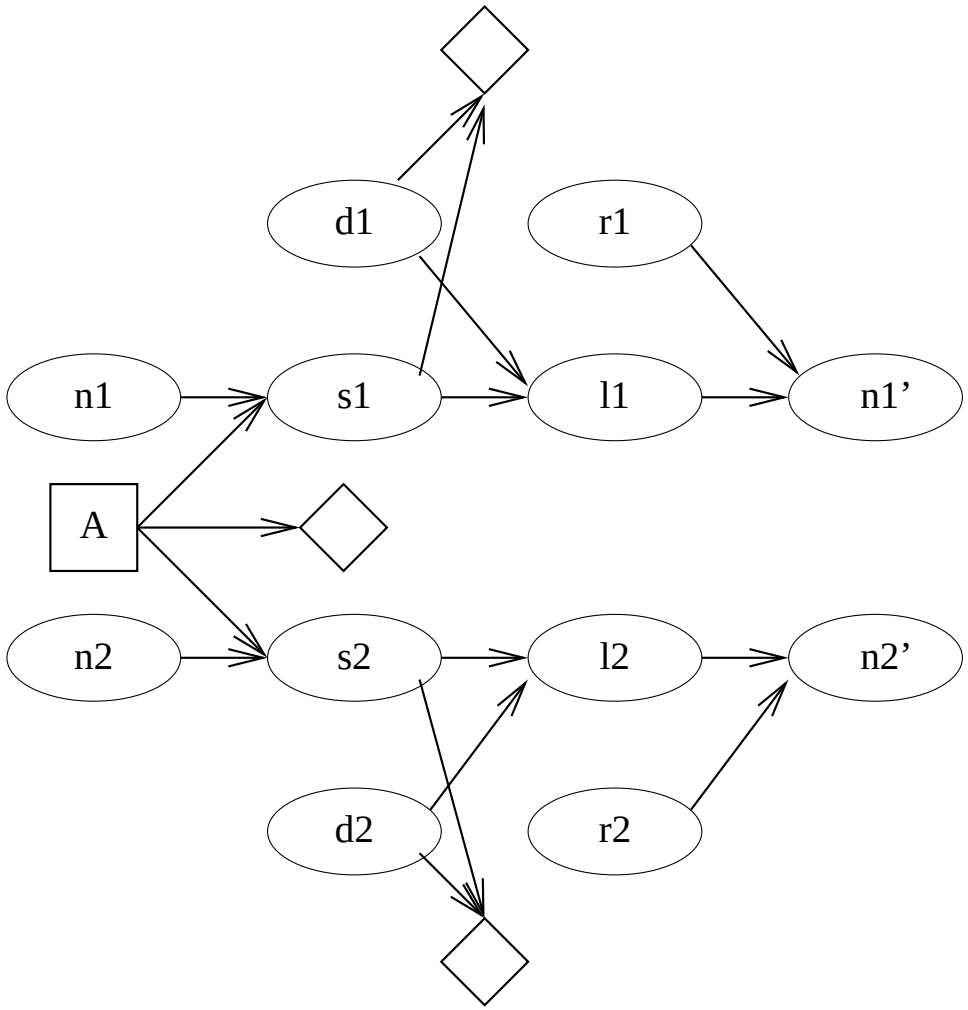

Compare number of parameters for $P(s'|s,a)$: 4 versus 73,205.

Number of parameters for $R(s'|s,a)$: 247 versus 73,205

Review of Probability: The Joint Distribution

When there are multiple random variables, the fundamental mathematical object is the joint probability distribution:

	Starts		¬Starts	
	Gas	¬Gas	Gas	¬Gas
GasGauge	.5040	.0006	.2160	.0594
¬GasGauge	.0560	.0014	.0240	.1386

We will use three notations:

- **Single Cell:** $P(\mathsf{Gas}, -\mathsf{GasGauge}, \mathsf{Starts}) = P(\mathsf{Gas} \wedge \neg\mathsf{GasGauge} \wedge \mathsf{Starts}) = .0560$.
- **Whole Table:** $\mathbf{P}(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts})$ stands for the entire joint distribution.
- **Summation Shorthand:** $\sum_{\mathsf{Gas}} P(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) = P(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) + P(-\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts})$

  For the distribution, we write this as $\sum_{\mathsf{Gas}} \mathbf{P}(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) = \mathbf{P}(\mathsf{GasGauge}, \mathsf{Starts})$

Marginal Probabilities

Suppose we want to know the probability that GasGauge is true. We obtain this by summing over the entire top row of the table.

$$P(\mathsf{GasGauge}) = .5040 + .0006 + .2160 + .0594 = 0.78$$

Another example: the probability that Gas is false:

$$P(-\mathsf{Gas}) = .0006 + .0014 + .0594 + .1386 = 0.2$$

This is called *marginalizing*. We write it as

$$P(-\mathsf{Gas}) = \sum_{GasGauge, Starts} P(-\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts})$$

$\mathbf{P}(\mathsf{Gas})$ is called the *marginal distribution* of Gas.

Conditional Probability: Belief Updating

Suppose we want to know the probability that the car will start given that we have observed GasGauge and found it to be true.

$$P(\mathsf{Starts}|\mathsf{GasGauge}) = \frac{P(\mathsf{Starts}, \mathsf{GasGauge})}{P(\mathsf{GasGauge})} = \frac{.5040 + .0006}{.5040 + .0006 + .2160 + .0594} = .6469$$

	Starts		–Starts	
	Gas	–Gas	Gas	–Gas
GasGauge	.5040	.0006	.2160	.0594
–GasGauge	.0560	.0014	.0240	.1386

Since GasGauge is true, one of the four situations from the top row of the table must hold. We *renormalize* the probabilities and then marginalize to compute $\mathbf{P}(\mathsf{Starts})$.

Every time the agent learns the value of some random variable, it can update the joint distribution to include only the remaining relevant variables.

Conditional Probability Distribution

The conditional probability distribution $\mathbf{P}(\mathsf{GasGauge}|\mathsf{Gas})$:

	GasGauge	–GasGauge
Gas	.9	.1
–Gas	.3	.7

Each row must sum to 1.

Chain Rule of Probability

$$P(A, B, C) = P(A|B, C) \cdot P(B|C) \cdot P(C)$$

for any number of variables in any order.

Proof:

$$\begin{aligned} P(A, B, C) &= \frac{P(A, B, C)}{P(B, C)} \cdot P(B, C) \\ &= P(A|B, C) \cdot P(B, C) \\ &= P(A|B, C) \cdot \frac{P(B, C)}{P(C)} \cdot P(C) \\ &= P(A|B, C) \cdot P(B|C) \cdot P(C) \end{aligned}$$

This holds for distributions as well:

$$\mathbf{P}(A, B, C) = \mathbf{P}(A|B, C) \cdot \mathbf{P}(B|C) \cdot \mathbf{P}(C)$$

Belief Networks (1): Independence

Two random variables X and Y are *independent* if $P(X, Y) = P(X) \cdot P(Y)$ for all values of X and Y. This can be written as $\mathbf{P}(X, Y) = \mathbf{P}(X) \cdot \mathbf{P}(Y)$. Intuitively, there is no causal connection betwen X and Y.

Example:

- **X**: coin with probability $X = true$ of .4, and $X = false$ of .6.
- **Y**: coin with probability $Y = true$ of .8, and $Y = false$ of .2.

Joint distribution:

	$X = true$	$X = false$
$Y = true$	.32	.48
$Y = false$	.08	.12

Belief Networks (2): Conditional Independence

Two random variables X and Y are *conditionally independent given* Z if $P(X, Y|Z) = P(X|Z) \cdot P(Y|Z)$ for all values of X, Y, and Z.

Example:

$$\mathbf{P}(\mathsf{Starts}, \mathsf{GasGauge}|\mathsf{Gas}) = \mathbf{P}(\mathsf{Starts}|\mathsf{Gas}) \cdot \mathbf{P}(\mathsf{GasGauge}|\mathsf{Gas})$$

	Starts		–Starts	
	GasGauge	–GasGauge	GasGauge	–GasGauge
Gas	.630	.070	.270	.030
–Gas	.003	.007	.297	.693

	GasGauge	–GasGauge
Gas	.9	.1
–Gas	.3	.7

	Starts	–Starts
Gas	.7	.3
–Gas	.01	.99

Simplifying the Joint Distribution by Conditional Independence

By the definition of conditional probability, we have

$$\mathbf{P}(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) = \mathbf{P}(\mathsf{GasGauge}, \mathsf{Starts}|\mathsf{Gas})\mathbf{P}(\mathsf{Gas})$$

But by conditional independence, the first factor can be rewritten as

$$\mathbf{P}(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) = \mathbf{P}(\mathsf{GasGauge}|\mathsf{Gas})\mathbf{P}(\mathsf{Starts}|\mathsf{Gas})\mathbf{P}(\mathsf{Gas})$$

Bayes Network Diagram

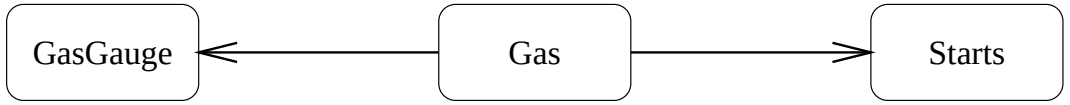

- **One node for each random variable**
- **Nodes with only outgoing arrows store marginal distributions**
  The Gas node stores $\mathbf{P}(\mathsf{Gas})$.
- **Nodes with incoming arrows store conditional probability tables**
  The Starts node stores $\mathbf{P}(\mathsf{Starts}|\mathsf{Gas})$.
- **Only direct dependencies are shown.** If two nodes are not directly connected, then they are conditionally independent on any set of nodes that disconnects them.
  Starts and GasGauge are conditionally independent given Gas.
- **Joint distribution is product of the distributions of the nodes.**

$$\mathbf{P}(\mathsf{Gas}, \mathsf{GasGauge}, \mathsf{Starts}) = \mathbf{P}(\mathsf{GasGauge}|\mathsf{Gas})\mathbf{P}(\mathsf{Starts}|\mathsf{Gas})\mathbf{P}(\mathsf{Gas})$$

Conformal Products

Given two probability tables, their **conformal product** is constructed by allocating space for the resulting table and then filling in each cell with the product of the corresponding cells.

	Starts		–Starts	
	GasGauge	–GasGauge	GasGauge	–GasGauge
Gas	$a11 \cdot b11$	$a10 \cdot b11$	$a11 \cdot b10$	$a10 \cdot b10$
–Gas	$a01 \cdot b01$	$a00 \cdot b01$	$a01 \cdot b00$	$a00 \cdot b00$

	GasGauge	–GasGauge
Gas	$a11$	$a10$
–Gas	$a01$	$a00$

	Starts	–Starts
Gas	$b11$	$b10$
–Gas	$b01$	$b00$

Commutativity and Associativity of Conformal Products

We can perform conformal products in any order. The only essential thing is that in each cell $x[i, j, k]$ in the final joint distribution, we have the product of the corresponding cells in each of the conditional probability tables.

$$P(\text{–Gas}, \text{GasGauge}, \text{–Starts}) = P(\text{–Starts}|\text{–Gas}) \cdot P(\text{GasGauge}|\text{–Gas}) \cdot P(\text{Gas})$$