

Sampling

Suppose we want to evaluate $P(Q|E)$ where Q are the query variables and E are the evidence variables.

Prior Sampling: Draw samples from the Bayes net by sampling the parents and then sampling the children given the parents. $P(Q|E) \approx \frac{\text{count}(Q \text{ and } E)}{\text{count}(E)}$.

Rejection Sampling: Like prior sampling, but ignore all samples that are inconsistent with the evidence.

Likelihood Weighting: Fix the evidence variables, and weight each sample by the probability of the evidence variables given their parents.

Gibbs Sampling:

1. Fix evidence.
2. Initialize other variables randomly
3. Repeat:
 - (a) Choose non-evidence variable X .
 - (b) Resample X from $P(X|\text{markovblanket}(X))$

Decision Networks

- **Chance nodes** - Chance nodes in a decision network behave identically to Bayes' nets. Each outcome in a chance node has an associated probability, which can be determined by running inference on the underlying Bayes' net it belongs to. We'll represent these with ovals.
- **Action nodes** - Action nodes are nodes that we have complete control over; they're nodes representing a choice between any of a number of actions which we have the power to choose from. We'll represent action nodes with rectangles.
- **Utility nodes** - Utility nodes are children of some combination of action and chance nodes. They output a utility based on the values taken on by their parents, and are represented as diamonds in our decision networks.

The **expected utility** of taking an action $A = a$ given evidence $E = e$ and n chance nodes is computed with the following formula:

$$EU(A = a|E = e) = \sum_{x_1, \dots, x_n} P(X_1 = x_1, \dots, X_n = x_n|E = e)U(A = a, X_1 = x_1, \dots, X_n = x_n)$$

where each x_i represents a value that the i^{th} chance node can take on. The **maximum expected utility** is the expected utility of the action that has the highest expected utility:

$$MEU(E = e) = \max_a EU(A = a|E = e).$$

Value of Perfect Information

Value of perfect information (VPI) quantifies the amount an agent's maximum expected utility is expected to increase if it were to observe some new evidence. Usually observing new evidence comes at a cost. If we observed some new evidence $E' = e'$ before acting, the maximum expected utility of our action at that point would become

$$MEU(E = e, E' = e') = \max_a \sum_x P(X = x | E = e, E' = e') U(X = x, A = a).$$

However, note that *we don't know what new evidence we'll get*. Because we don't know what what new evidence e' we'll get, we must represent it as a random variable E' . We will compute the expected value of the maximum expected utility:

$$MEU(E = e, E') = \sum_{e'} P(E' = e' | E = e) MEU(E = e, E' = e').$$

Observing a new evidence variable yields a different MEU with probabilities corresponding to the probabilities of observing each value for the evidence variable, and so by computing $MEU(E = e, E')$ as above, we compute what we expect our new MEU will be if we choose to observe new evidence. The VPI is the expected maximum expected utility if we were to observe the new evidence, minus the maximum expected utility if we were not to observe the new evidence:

$$VPI(E' | E = e) = MEU(E = e, E') - MEU(E = e).$$

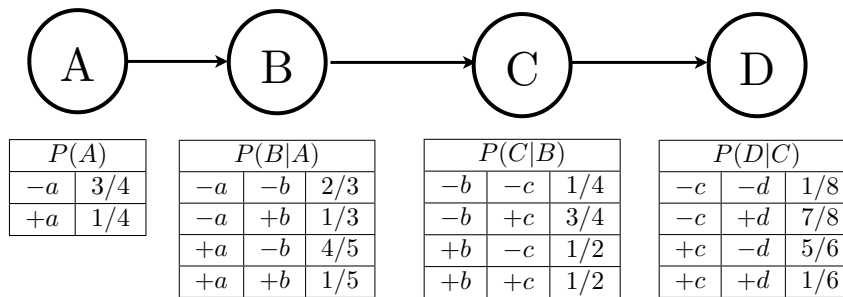
Properties of VPI

The value of perfect information has several very important properties, namely:

- **Nonnegativity.** $\forall E', e VPI(E' | E = e) \geq 0$
Observing new information always allows you to make a *more informed* decision, and so your maximum expected utility can only increase (or stay the same if the information is irrelevant for the decision you must make).
- **Nonadditivity.** $VPI(E_j, E_k | E = e) \neq VPI(E_j | E = e) + VPI(E_k | E = e)$ in general.
This is probably the trickiest of the three properties to understand intuitively. It's true because generally observing some new evidence E_j might change how much we care about E_k ; therefore we can't simply add the VPI of observing E_j to the VPI of observing E_k to get the VPI of observing both of them. Rather, the VPI of observing two new evidence variables is equivalent to observing one, incorporating it into our current evidence, then observing the other. This is encapsulated by the order-independence property of VPI, described more below.
- **Order-independence.** $VPI(E_j, E_k | E = e) = VPI(E_j | E = e) + VPI(E_k | E = e, E_j) = VPI(E_k | E = e) + VPI(E_j | E = e, E_k)$
Observing multiple new evidences yields the same gain in maximum expected utility regardless of the order of observation. This should be a fairly straightforward assumption - because we don't actually take any action until after observing any new evidence variables, it doesn't actually matter whether we observe the new evidence variables together or in some arbitrary sequential order.

Q1. Bayes' Nets Sampling

Assume the following Bayes' net, and the corresponding distributions over the variables in the Bayes' net:



(a) You are given the following samples:

$+a$	$+b$	$-c$	$-d$	$+a$	$-b$	$-c$	$+d$
$+a$	$-b$	$+c$	$-d$	$+a$	$+b$	$+c$	$-d$
$-a$	$+b$	$+c$	$-d$	$-a$	$+b$	$-c$	$+d$
$-a$	$-b$	$+c$	$-d$	$-a$	$-b$	$+c$	$-d$

(i) Assume that these samples came from performing Prior Sampling, and calculate the sample estimate of $P(+c)$.

(ii) Now we will estimate $P(+c \mid +a, -d)$. Above, clearly cross out the samples that would **not** be used when doing Rejection Sampling for this task, and write down the sample estimate of $P(+c \mid +a, -d)$ below.

(b) Using Likelihood Weighting Sampling to estimate $P(-a \mid +b, -d)$, the following samples were obtained. Fill in the weight of each sample in the corresponding row.

Sample	Weight
$-a$ $+b$ $+c$ $-d$	_____
$+a$ $+b$ $+c$ $-d$	_____
$+a$ $+b$ $-c$ $-d$	_____
$-a$ $+b$ $-c$ $-d$	_____

(c) From the weighted samples in the previous question, estimate $P(-a \mid +b, -d)$.

(d) Which query is better suited for likelihood weighting, $P(D \mid A)$ or $P(A \mid D)$? Justify your answer in one sentence.

(e) Recall that during Gibbs Sampling, samples are generated through an iterative process.

Assume that the only evidence that is available is $A = +a$. Clearly fill in the circle(s) of the sequence(s) below that could have been generated by Gibbs Sampling.

Sequence 1				
1 :	+a	-b	-c	+d
2 :	+a	-b	-c	+d
3 :	+a	-b	+c	+d


Sequence 2				
1 :	+a	-b	-c	+d
2 :	+a	-b	-c	-d
3 :	-a	-b	-c	+d

Sequence 3				
1 :	+a	-b	-c	+d
2 :	+a	-b	-c	-d
3 :	+a	+b	-c	-d

Sequence 4				
1 :	+a	-b	-c	+d
2 :	+a	-b	-c	-d
3 :	+a	+b	-c	+d

2 Decision Networks and VPI

A used car buyer can decide to carry out various tests with various costs (e.g., kick the tires, take the car to a qualified mechanic) and then, depending on the outcome of the tests, decide which car to buy. We will assume that the buyer is deciding whether to buy car c and that there is time to carry out at most one test which costs \$50 and which can help to figure out the quality of the car. A car can be in good shape (of good quality $Q = +q$) or in bad shape (of bad quality $Q = \neg q$), and the test might help to indicate what shape the car is in. There are only two outcomes for the test T : pass ($T = \text{pass}$) or fail ($T = \text{fail}$). Car c costs \$1,500, and its market value is \$2,000 if it is in good shape; if not, \$700 in repairs will be needed to make it in good shape. The buyer's estimate is that c has 70% chance of being in good shape. The Decision Network is shown below.



(a) Calculate the expected net gain from buying car c , given no test.

(b) Tests can be described by the probability that the car will pass or fail the test given that the car is in good or bad shape. We have the following information:

$$P(T = \text{pass} | Q = +q) = 0.9$$

$$P(T = \text{pass} | Q = \neg q) = 0.2$$

Calculate the probability that the car will pass (or fail) its test, and then the probability that it is in good (or bad) shape given each possible test outcome.

(c) Calculate the optimal decisions given either a pass or a fail, and their expected utilities.

(d) Calculate the value of (perfect) information of the test. Should the buyer pay for a test?