## Stochastic Simulation

- Idea: probabilities $\leftrightarrow$ samples
- Get probabilities from samples:

| $X$ | count |
| :---: | :---: |
| $x_{1}$ | $n_{1}$ |
| $\vdots$ | $\vdots$ |
| $x_{k}$ | $n_{k}$ |
| total | $m$ |$\leftrightarrow$| $X$ | probability |
| :---: | :---: |
| $x_{1}$ | $n_{1} / m$ |
| $\vdots$ | $\vdots$ |
| $x_{k}$ | $n_{k} / m$ |

- If we could sample from a variable's (posterior) probability, we could estimate its (posterior) probability.


## Generating samples from a distribution

For a variable X with a discrete domain or a (one-dimensional) real domain:

- Totally order the values of the domain of $X$.
- Generate the cumulative probability distribution: $f(x)=P(X \leq x)$.
- Select a value $y$ uniformly in the range $[0,1]$.
- Select the $x$ such that $f(x)=y$.


## Cumulative Distribution



1


## Forward sampling in a belief network

- Sample the variables one at a time; sample parents of $X$ before you sample $X$.
- Given values for the parents of $X$, sample from the probability of $X$ given its parents.


## Rejection Sampling

- To estimate a posterior probability given evidence $Y_{1}=v_{1} \wedge \ldots \wedge Y_{j}=v_{j}:$
- Reject any sample that assigns $Y_{i}$ to a value other than $v_{i}$.
- The non-rejected samples are distributed according to the posterior probability:

$$
P(\alpha) \approx \frac{\sum_{\text {sample } \models \alpha} 1}{\sum_{\text {sample }} 1}
$$

where we consider only samples consistent with observations.

## Rejection Sampling Example: $P(t a \mid s m, r e)$

|  | Ta | Fi | Al | Sm | Le | Re |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $s_{1}$ | true | false | true | false | - | - | $\boldsymbol{X}$ |


$s_{2}$
$S_{3}$
$s_{4}$ true true true true true true
$s_{1000}$ false false false false
$P(s m)=0.02$
$P(r e \mid s m)=0.32$
How many samples are rejected?
How many samples are used?

## Importance Sampling

- Samples have weights: a real number associated with each sample that takes the evidence into account.
- Probability of a proposition is weighted average of samples:

$$
P(\alpha \mid \text { observations }) \approx \frac{\sum_{\text {sample } \models \alpha} \text { weight }(\text { sample })}{\sum_{\text {sample }} \text { weight }(\text { sample })}
$$

- If we can compute $P$ (evidence|sample) we can weight the (partial) sample by this value.
- Mix exact inference with sampling: don't sample all of the variables, but weight each sample appropriately.
- Sample according to a proposal distribution, as long as the samples are weighted appropriately.


## Importance Sampling Example: $P(t a \mid s m, r e)$

|  | Ta | Fi | Al | Le | Weight |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $s_{1}$ | true | false | true | false | $0.01 \times 0.01$ |


$s_{2} \quad$ false true false false $0.9 \times 0.01$
$s_{3}$ false true true true $0.9 \times 0.75$
$s_{4}$ true true true true $0.9 \times 0.75$
$s_{1000}$ false false true true $0.01 \times 0.75$
$P(s m \mid f i)=0.9$
$P(s m \mid \neg f i)=0.01$
$P(r e \mid l e)=0.75$
$P(r e \mid \neg l e)=0.01$

## Particle Filtering

- Suppose the evidence is $e_{1} \wedge e_{2}$ $P\left(e_{1} \wedge e_{2} \mid\right.$ sample $)=P\left(e_{1} \mid\right.$ sample $) P\left(e_{2} \mid e_{1} \wedge\right.$ sample $)$
- After computing $P\left(e_{1} \mid\right.$ sample $)$, we may know the sample will have an extremely small probability.
- Idea: we use lots of samples: "particles". A particle is a sample on some of the variables.
- Based on $P\left(e_{1} \mid\right.$ sample $)$, we resample the set of particles. We select from the particles according to their weight.
- Some particles may be duplicated, some may be removed.

