## Review

- Probability
- Random Variables
- Joint and Marginal Distributions
- Conditional Distribution
- Product Rule, Chain Rule, Bayes' Rule
- Inference
- Independence


## Random Variables

- A random variable is some aspect of the world about which we (may) have uncertainty
- $\mathrm{R}=\mathrm{Is}$ it raining?
- $\mathrm{D}=$ How long will it take to drive to work?
- L = Where am I?
- We denote random variables with capital letters
- Like variables in a CSP, random variables have domains
- $R$ in $\{$ true, false (sometimes write as $\{+r, \neg r\}$ )
- D in [0, $\infty$ )
- $L$ in possible locations, maybe $\{(0,0),(0,1), \ldots\}$


## Probability Distributions

- Unobserved random variables have distributions

| $P(T)$ |  | $P(W)$ |  |
| :---: | :---: | :---: | :---: |
| T | P | W | P |
| warm | 0.5 | sun | 0.6 |
| cold | 0.5 | rain | 0.1 |
|  |  | fog | 0.3 |
|  |  | meteor | 0.0 |

- A distribution is a TABLE of probabilities of values
- A probability (lower case value) is a single number

$$
P(W=\text { rain })=0.1 \quad P(\text { rain })=0.1
$$

- Must have: $\forall x P(x) \geq 0$

$$
\sum_{x} P(x)=1
$$

## Joint Distributions

- A joint distribution over a set of random variables: $X_{1}, X_{2}, \ldots X_{n}$ specifies a real number for each assignment (or outcome):

$$
\begin{aligned}
& P\left(X_{1}=x_{1}, X_{2}=x_{2}, \ldots X_{n}=x_{n}\right) \\
& P\left(x_{1}, x_{2}, \ldots x_{n}\right)
\end{aligned}
$$

- Must obey:

$$
\begin{aligned}
P\left(x_{1}, x_{2}, \ldots x_{n}\right) & \geq 0 \\
\sum_{\left(x_{1}, x_{2}, \ldots x_{n}\right)} P\left(x_{1}, x_{2}, \ldots x_{n}\right) & =1
\end{aligned}
$$

$P(T, W)$

| T | W | P |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

- For all but the smallest distributions, impractical to write out


## Probabilistic Models

- A probabilistic model is a joint distribution over a set of random variables
- Probabilistic models:
- (Random) variables with domains Assignments are called outcomes
- Joint distributions: say whether assignments (outcomes) are likely
- Normalized: sum to 1.0
- Ideally: only certain variables directly interact
- Constraint satisfaction probs:
- Variables with domains
- Constraints: state whether assignments are possible
- Ideally: only certain variables directly interact

Distribution over T,W

| T | W | P |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

Constraint over T,W

| T | W | P |
| :---: | :---: | ---: |
| hot | sun | T |
| hot | rain | F |
| cold | sun | F |
| cold | rain | T |

## Events

- An event is a set E of outcomes

$$
P(E)=\sum_{\left(x_{1} \ldots x_{n}\right) \in E} P\left(x_{1} \ldots x_{n}\right)
$$

- From a joint distribution, we can calculate the probability of any event

| $T$ | $W$ | $P$ |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

- Probability that it's hot AND sunny?
- Probability that it's hot?
- Probability that it's hot OR sunny?
- Typically, the events we care about are partial assignments, like $\mathrm{P}(\mathrm{T}=$ hot $)$


## Marginal Distributions

- Marginal distributions are sub-tables which eliminate variables
- Marginalization (summing out): Combine collapsed rows by adding

| $P(T, W)$ |  |  |  | $P(T)$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | T | P |
| T | W | P | $P(t)=\sum_{s} P(t, s)$ | hot | 0.5 |
| hot | sun | 0.4 |  | cold | 0.5 |
| hot | rain | 0.1 |  | $P(W)$ |  |
| cold | sun | 0.2 | $\overrightarrow{P(s)=\sum_{t} P(t, s)}$ | W | P |
| cold | rain | 0.3 |  | sun | 0.6 |
|  |  |  |  | rain | 0.4 |

$$
P\left(X_{1}=x_{1}\right)=\sum_{x_{2}} P\left(X_{1}=x_{1}, X_{2}=x_{2}\right)
$$

## Conditional Probabilities

- A simple relation between joint and conditional probabilities
- In fact, this is taken as the definition of a conditional probability

$$
P(a \mid b)=\frac{P(a, b)}{P(b)}
$$

$$
P(T, W)
$$



| T | W | P |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

## Conditional Distributions

- Conditional distributions are probability distributions over some variables given fixed values of others

Conditional Distributions

$\left\{\right.$| $P(W \mid T=$ hot $)$ |  |
| :---: | :---: |
| W P <br> sun 0.8 <br> rain 0.2 <br> $P(W \mid T=$ cold $)$  <br> $\begin{array}{\|c\|c\|}\hline \mathrm{W} & \mathrm{P} \\ \hline \text { sun } & 0.4 \\ \hline \text { rain } & 0.6 \\ \hline\end{array}$ $. \begin{array}{l}\end{array}$ |  |


| $P(T, W)$ |  |  |
| :---: | :---: | :---: |
| T | W | P |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

## Normalization Trick

- A trick to get a whole conditional distribution at once:
- Select the joint probabilities matching the evidence
- Normalize the selection (make it sum to one)
$P(T, W)$

| T | W | P |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |


|  | $P(T, r)$ |  |  |  | $P(T \mid r)$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | T | R | P |  | T | P |
| Select | hot | rain | 0.1 | Normalize | hot | 0.25 |
|  | cold | rain | 0.3 |  | cold | 0.75 |

## Probabilistic Inference

- Probabilistic inference: compute a desired probability from other known probabilities (e.g. conditional from joint)
- We generally compute conditional probabilities
- P (on time $\mid$ no reported accidents) $=0.90$
- These represent the agent's beliefs given the evidence
- Probabilities change with new evidence:
- P (on time | no accidents, 5 a.m.) $=0.95$
- $\quad \mathrm{P}$ (on time | no accidents, 5 a.m., raining) $=0.80$
- Observing new evidence causes beliefs to be updated


## Inference by Enumeration

- General case:
- Evidence variables: $E_{1} \ldots E_{k}=e_{1} \ldots e_{k}$
- Query* variable: $Q$
- Hidden variables: $H_{1} \ldots H_{r}$
$X_{1}, X_{2}, \ldots X_{n}$
All variables
- We want: $P\left(Q \mid e_{1} \ldots e_{k}\right)$
- First, select the entries consistent with the evidence
- Second, sum out H to get joint of Query and evidence:

$$
P\left(Q, e_{1} \ldots e_{k}\right)=\sum_{h_{1} \ldots h_{r}} \underbrace{P\left(Q, h_{1} \ldots h_{r}, e_{1} \ldots e_{k}\right)}_{X_{1}, X_{2}, \ldots X_{n}}
$$

- Finally, normalize the remaining entries to conditionalize
- Obvious problems:
- Worst-case time complexity O(dn)
- Space complexity $O\left(d^{n}\right)$ to store the joint distribution


## The Product Rule

- Sometimes have conditional distributions but want the joint

$$
P(x \mid y)=\frac{P(x, y)}{P(y)}
$$



$$
P(x, y)=P(x \mid y) P(y)
$$

- Example:
$P(D \mid W)$
$P(D \mid W)$

| D | W | P |
| :---: | :---: | :---: |
| wet | sun | 0.1 |
| dry | sun | 0.9 |
| wet | rain | 0.7 |
| dry | rain | 0.3 |


$P(W)$

| $R$ | $P$ |
| :---: | :---: |
| sun | 0.8 |
| rain | 0.2 |

$P(D, W)$

| D | W | P |
| :---: | :---: | :---: |
| wet | sun | 0.08 |
| dry | sun | 0.72 |
| wet | rain | 0.14 |
| dry | rain | 0.036 |

## The Chain Rule

- More generally, can always write any joint distribution as an incremental product of conditional distributions

$$
\begin{aligned}
& P\left(x_{1}, x_{2}, x_{3}\right)=P\left(x_{1}\right) P\left(x_{2} \mid x_{1}\right) P\left(x_{3} \mid x_{1}, x_{2}\right) \\
& P\left(x_{1}, x_{2}, \ldots x_{n}\right)=\prod_{i} P\left(x_{i} \mid x_{1} \ldots x_{i-1}\right)
\end{aligned}
$$

## Bayes' Rule

- Two ways to factor a joint distribution over two variables:

$$
P(x, y)=P(x \mid y) P(y)=P(y \mid x) P(x)
$$

- Dividing, we get:

$$
P(x \mid y)=\frac{P(y \mid x)}{P(y)} P(x)
$$

- Why is this at all helpful?
- Lets us build one conditional from its reverse
- Often one conditional is tricky but the other one is simple
- Foundation of many systems
- In the running for most important AI equation!


## Independence

- Two variables are independent in a joint distribution if:

$$
\begin{gathered}
P(X, Y)=P(X) P(Y) \\
\forall x, y P(x, y)=P(x) P(y)
\end{gathered}
$$

- Says the joint distribution factors into a product of two simple ones
- Usually variables aren't independent!
- Can use independence as a modeling assumption
- Independence can be a simplifying assumption
- Empirical joint distributions: at best "close" to independent

