

# Artificial Intelligence

CS 165A

Apr 13, 2023

Instructor: Prof. Yu-Xiang Wang

T  
o  
d  
a  
y

- Continuous optimization
- Wrapping up ML
- Starting PGM

# Announcement

- Project 1 out on the course website. Please start early!
- Notes from the discussion section are available on the course website.
- We have weekly quiz on Ed discussion. All you need to do is to attempt it. We will decide on the bonus participation credit based on your participation there.


# Recap: Last lecture

- Data splitting
  - Holdout and Cross validation for tuning hyperparameters
- The problem of distribution shift
- Learning a classifier:
  - From 0–1 loss to surrogate losses
  - Logistic loss as an example
- Continuous Optimization with Gradient Descent

# Recap: How do we optimize a continuously differentiable function in general?

- The problem:  $\min_{\theta} f(\theta)$
- Let's just optimize it anyway!
  - With gradient descent.
- Assumption: The objective function is differentiable almost everywhere.

$$\theta_{t+1} = \theta_t - \eta_t \nabla f(\theta_t)$$



# Plan for today

- Stochastic Gradient Descent
- Making sense of the SGD updates
- Multiclass classification

# Gradient of logistic loss for learning a linear classifier

$$\nabla f(w) = \frac{1}{n} \sum_{i=1}^n \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} (-y_i x_i)$$

# Gradient of logistic loss for learning a linear classifier

$$\nabla f(w) = \frac{1}{n} \sum_{i=1}^n \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} (-y_i x_i)$$

- Question: What is the time complexity of computing this gradient?

# Gradient of logistic loss for learning a linear classifier

$$\nabla f(w) = \frac{1}{n} \sum_{i=1}^n \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} (-y_i x_i)$$

- Question: What is the time complexity of computing this gradient?



# Gradient of logistic loss for learning a linear classifier

$$\nabla f(w) = \frac{1}{n} \sum_{i=1}^n \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} (-y_i x_i)$$

- Question: What is the time complexity of computing this gradient?

$$O(n \cdot d)$$

$$O(n \cdot \text{nnz}(X))$$

**Can we do better?**

# Stochastic Gradient Descent (Robbins-Monro 1951)

- Gradient descent

$$\theta_{t+1} = \theta_t - \eta_t \nabla f(\theta_t)$$


- Stochastic gradient descent

$$\theta_{t+1} = \theta_t - \eta_t \hat{\nabla} f(\theta_t)$$

- Using a stochastic approximation of the gradient:

$$\mathbb{E}[\hat{\nabla} f(\theta_t) | \theta_t] = \nabla f(\theta_t) \quad \text{“unbiased”}$$

$$\mathbf{Var}[\hat{\nabla} f(\theta_t) | \theta_t] \leq \sigma^2 \quad \underbrace{\mathbb{E}[\|\hat{\nabla} f(\theta_t) - \nabla f(\theta_t)\|^2]}_{\leq \sigma^2}$$



Herbert Robbins  
1915 - 2001

Question: What's the time complexity of each iteration in SGD?

# One natural stochastic gradient to consider in machine learning

$$\min_{\theta} f(\theta)$$

- Recall that

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \ell(\theta, (x_i, y_i))$$

# One natural stochastic gradient to consider in machine learning

- Recall that

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \ell(\theta, (x_i, y_i))$$

- Pick a **single** data point  $i$  uniformly at random

- Use  $\nabla_{\theta} \ell(\theta, (x_i, y_i))$

# One natural stochastic gradient to consider in machine learning

- Recall that

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \ell(\theta, (x_i, y_i))$$


- Pick a **single** data point  $i$  uniformly at random

- Use  $\nabla_{\theta} \ell(\theta, (x_i, y_i))$


- Show that this is an unbiased estimator!

$$\left[ \mathbb{E}_{i \sim \text{Uniform}\{1, 2, 3, \dots, n\}} \nabla_{\theta} \ell(\theta, (x_i, y_i)) \stackrel{?}{=} \nabla_{\theta} \left( \frac{1}{n} \sum_{i=1}^n \ell_i(\theta) \right) \right]$$
$$\sum_i p_{i1} \cdot \nabla_{\theta} \ell(\theta, (x_i, y_i)) = \sum_i \frac{1}{n} \nabla_{\theta} \ell(\theta, (x_i, y_i))$$

# Illustration of GD vs SGD




**Batch Gradient Descent**



**Stochastic Gradient Descent**


**Observation:** With the time gradient descent taking one step. SGD would have already moved many steps.

# Intuition of the SGD algorithm on the “Spam Filter” example



- $\text{Score}(\mathbf{x}) = w_0 + w_1 * 1(\text{hyperlinks}) + w_2 * 1(\text{contact list}) + w_3 * \text{misspelling} + w_4 * \text{length}$

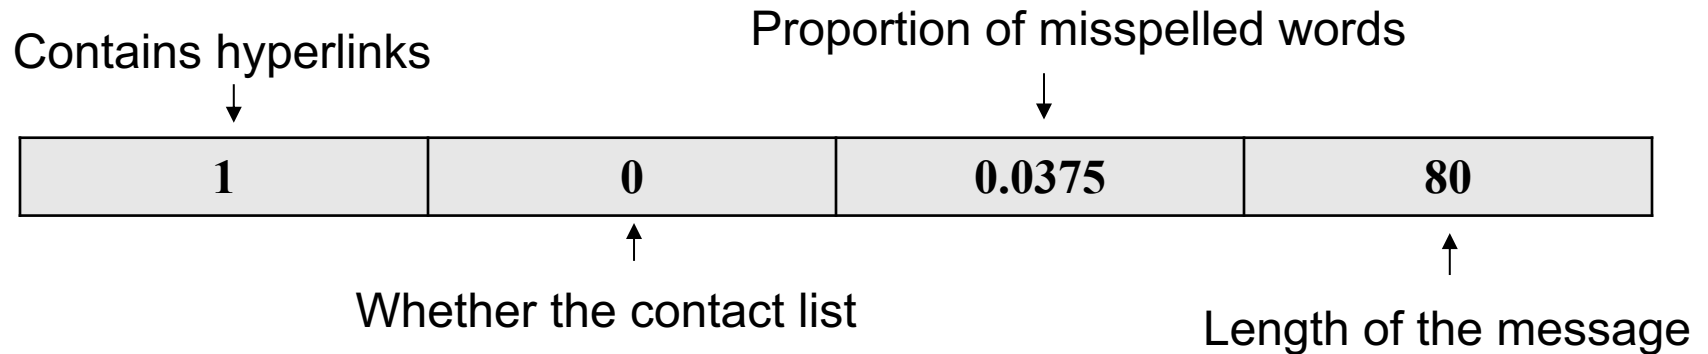
# Intuition of the SGD algorithm on the “Spam Filter” example



- $\text{Score}(\mathbf{x}) = w_0 + w_1 * 1(\text{hyperlinks}) + w_2 * 1(\text{contact list}) + w_3 * \text{misspelling} + w_4 * \text{length}$
- **Meaning of these weights?**




# Intuition of the SGD algorithm on the “Spam Filter” example



- $\text{Score}(\mathbf{x}) = w_0 + w_1 * 1(\text{hyperlinks}) + w_2 * 1(\text{contact list}) + w_3 * \text{misspelling} + w_4 * \text{length}$
- **Meaning of these weights?**
  - The more positive, the more we think the feature is associated with Spam email.

# Intuition of the SGD algorithm on the “Spam Filter” example



- $\text{Score}(x) = w_0 + w_1 * 1(\text{hyperlinks}) + w_2 * 1(\text{contact list}) + w_3 * \text{misspelling} + w_4 * \text{length}$
- **Meaning of these weights?**
  - The more positive, the more we think the feature is associated with Spam email.
  - The more negative, the less that we think the feature is associated with Spam email

# Intuition of the SGD algorithm on the “Spam Filter” example

$$\nabla \ell(w, (x_i, y_i)) = \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} (-y_i x_i)$$

# Intuition of the SGD algorithm on the “Spam Filter” example

$$\nabla \ell(w, (x_i, y_i)) = \frac{\exp(-y_i \cdot x_i^T w)}{\underbrace{1 + \exp(-y_i \cdot x_i^T w)}} (-y_i x_i)$$

Scalar  $> 0$ :

$\approx 0$  if the prediction  
is correct

$\approx 1$  otherwise

# Intuition of the SGD algorithm on the “Spam Filter” example

$$\nabla \ell(w, (x_i, y_i)) = \underbrace{\frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)}}_{\text{Scalar}} \underbrace{(-y_i x_i)}_{\text{Vector}}$$

Scalar > 0:  
≈ 0 if the prediction  
is correct  
≈ 1 otherwise

Vector of dimension d:  
provides the direction  
of the gradient

# Intuition of the SGD algorithm on the “Spam Filter” example

$$\nabla \ell(w, (x_i, y_i)) = \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} \underbrace{\left( -y_i x_i \right)}$$

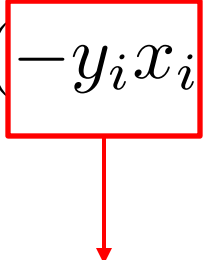
Scalar > 0:  
≈ 0 if the prediction  
is correct  
≈ 1 otherwise

Vector of dimension d:  
provides the direction  
of the gradient

If we receive an example [1, 0, 0.0375, 80] like the one before.  
And a label  $y = 1$  saying that this is a spam.

How will the SGD update change the weight vector?

# Intuition of the SGD algorithm on the “Spam Filter” example

$$\nabla \ell(w, (x_i, y_i)) = \frac{\exp(-y_i \cdot x_i^T w)}{1 + \exp(-y_i \cdot x_i^T w)} \underbrace{\left( -y_i x_i \right)}$$


Scalar  $> 0$ :  
 $\approx 0$  if the prediction  
is correct  
 $\approx 1$  otherwise

Vector of dimension  $d$ :  
provides the direction  
of the gradient

If we receive an example  $[1, 0, 0.0375, 80]$  like the one before.  
And a label  $y = 1$  saying that this is a spam.

**How will the SGD update change the weight vector?**

Then by moving  $w$  towards the negative gradient direction, we are changing the weight vector by increasing the weights. i.e., increasing the amount they contribute to the score function (if currently the classifier is making a mistake on this example)

# How to choose the step sizes / learning rates?

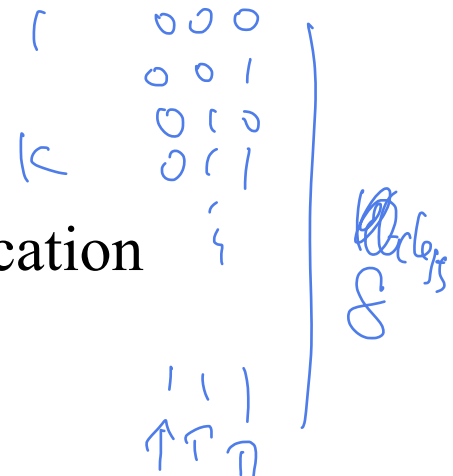
- In theory:
  - Gradient decent:  $1/L$  where  $L$  is the **Gradient Lipschitz constant** of the function we minimize.
  - SGD:  $\sum_t \eta_t = \infty, \sum_t \eta_t^2 < \infty$ 
    - e.g.  $\eta_t \in [1/t, 1/\sqrt{t})$
- In practice:
  - Use cross-validation on a subsample of the data.
  - Fixed learning rate for SGD is usually fine.
  - If it diverges, decrease the learning rate.
  - If for extremely small learning rate, it still diverges, check if your gradient implementation is correct.



# The power of SGD

- Extremely general:
  - Specify an end-to-end differentiable score function, e.g., a complex neural network.
  - Beyond the context of machine learning
- Extremely simple: a few lines of code.
- Extremely scalable
  - Just a few pass of the data, no need to store the data
- People are continuing to discover that many methods are special cases of SGD.

# Multiclass classification problem



- You will be implementing a multi-class classification algorithm using linear logistic regression.


$$\mathcal{Y} = \{1, 2, \dots, k\}$$

- Recall that in binary classification you have one score function, the multi-class has one score function for each class.



$$\text{Sign}(x^T(w_1 - w_2))$$

$$h(x) = \underset{j \in \{1, \dots, k\}}{\text{argmax}} \text{Score}_{w_j}(x) = \underset{j \in \{1, \dots, k\}}{\text{argmax}} (x^T w_j)$$

- Now you have k linear score functions. Represent them as a matrix of weights  $W$  (**what's the dimension?**)
- You use the softmax-cross-entropy loss instead of the logistic loss.



# Decision boundaries of multi-class linear logistic regression



map image x to digit y

# From linear logistic regression to neural networks




Fig. 3.4.1 Softmax regression is a single-layer neural network.




Fig. 4.1.1 An MLP with a hidden layer of 5 hidden units.

- Same input, same loss function
- The only difference is in how the score function is computed

# From linear logistic regression to neural networks (side by side comparison)

$$\text{Score}_w(x) = W^T x = \sum_{i=1}^d x_i w_i$$

Parameters       $w$

$$\text{Score}_w(x) = \underbrace{w_2^T \sigma(W_1 x)}_{w_2 \in \mathbb{R}^h}$$

$$\theta = \underline{\underline{(W_1, W_2)}}$$


# Summary: steps to build a classifier agent

- 0. Collect a large labeled dataset.
- 1. Data splitting into training-validation-testing
- 2. Feature extraction
- 3. Specifying a hypothesis class
- 4. Learning with SGD using logistic loss
- 5. Validating on the validation set.
- 6. Testing on the test set.
- 7. Deploy the classifier agent.



Return to Step 2 or 3  
If not happy with the  
validation error.  
Do model selection  
(e.g. hyperparameter  
tuning)

# Objected Oriented Design of a “ClassifierAgent” class (as you will in the project)

 ClassifierAgent
□ <b>params</b> : numpy array
■ <b>score</b> (input feature X): predicted scores ■ <b>feature_extractor</b> (raw input text): feature vector X
● <b>train</b> (labeled data, algorithm: GD or SGD): update <b>params</b> ● <b>predict</b> (raw input text): predicted labels ● <b>loss</b> (labeled data): average loss ● <b>error</b> (labeled data): classification error

- Considerations:
  - Lightweight. It does not store data. Can update parameters with daily new data by calling “train”.
  - Extendable. Inherit with a different “feature extractor” / “score” function.
  - Easy to modify training algorithms / loss functions / performance metrics

# So far, we've learned everything about a classifier agent

- **Modeling**
  - Design meaningful feature extractors
  - Specify a family of classifiers on how the agent is going to classify examples using the features (indexed by free parameters)
- **Inference:** Trivially follows from the specification
- **Learning:** minimize errors (surrogate losses) over “free parameters”

This is called “discriminative” approach in modelling.



# “Generative” modeling: Modeling the world with a (joint) probability distribution

- The “Discriminative approach” doesn’t care about the underlying processes that give rise to the observed data.
- **Modeling:** “Generative” modelling:
  - Label is a random variable  $Y$ , features are a vector random variables  $X = [X_1, \dots, X_d]^T$
  - Model the world with a (family of) joint probability distributions.
- **Inference:** Then we can make principled inference by calculating  $P(Y|X)$ 
  - You don’t just get a prediction, but also confidence.
- **Learning:** Find the distribution that fits the data well.

# Side-by-side comparisons of two modelling principles: “Discriminative” vs “Generative”

	“Discriminative”	“Generative”
Modelling	<ul style="list-style-type: none"> <li>• Extract features</li> <li>• Hypothesis class</li> <li>• <math>h_{\theta}: X \rightarrow Y</math></li> </ul>	<ul style="list-style-type: none"> <li>• How data are generated</li> <li>• Family of probability distributions <math>P_{\theta}(X, Y)</math></li> </ul>
Inference	Just apply the classifier: $h(X)$	Calculate $P(Y X)$ <del><math>P(Y X)</math></del> Predict with: $\max_y P(Y=y   X)$
Learning	Minimize error over “free parameters”: $\theta$	Maximize likelihood over “free parameters”: $\theta$

See Ng and Jordan (NIPS-2001):

<https://ai.stanford.edu/~ang/papers/nips01-discriminativegenerative.pdf>

# Example: Spam filter

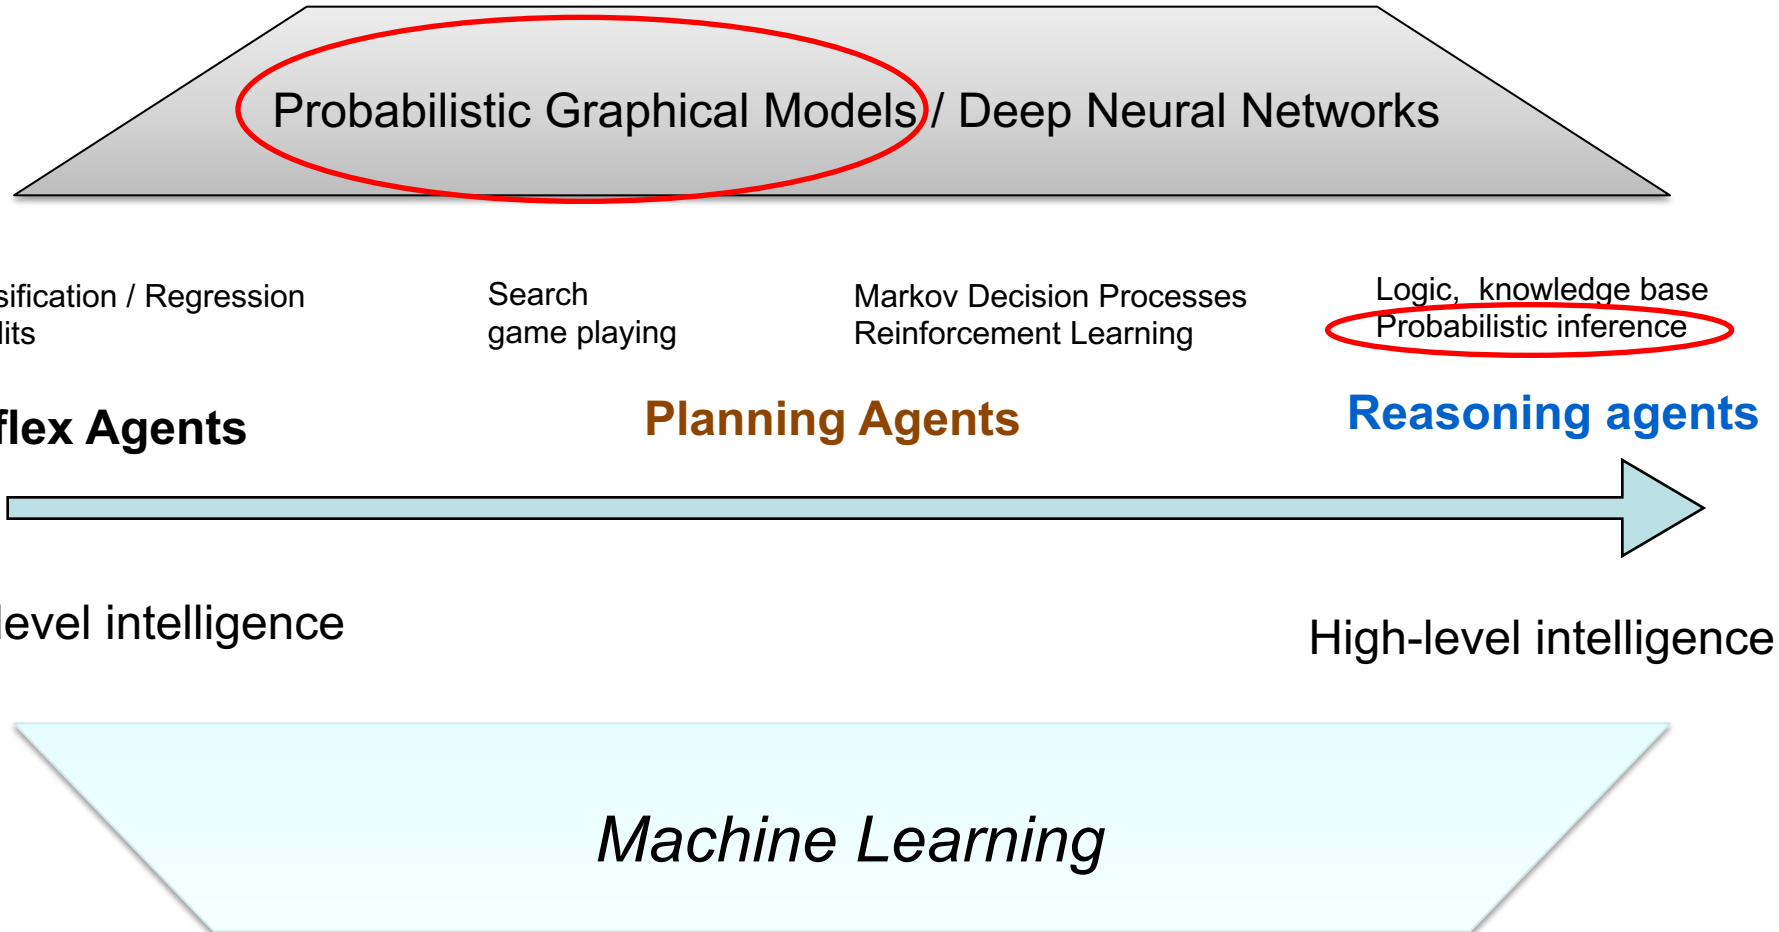
- “Discriminative modeling”:
  - Think about features that has discriminative power
  - Think about hypothesis class (e.g., shape of the decision boundaries, functional form of the score function)
- “Generative modelling”:
  - Think about how “spammers” write (generate) their emails.
  - And also how “non-spammers” write their emails.
  - Model the probability of certain words appear / the probability of certain word combinations appear.

# Example: Spam filter

- “Discriminative modeling”:
  - Think about features that has discriminative power
  - Think about hypothesis class (e.g., shape of the decision boundaries, functional form of the score function)
- “Generative modelling”:
  - Think about how “spammers” write (generate) their emails.
  - And also how “non-spammers” write their emails.
  - Model the probability of certain words appear / the probability of certain word combinations appear.

Once we learned probabilistic graphical model, we will be able to do this effectively.

# Structure of the course



(Again this idea is adapted from Percy Liang's teachings)

# Remaining time today and the next two lectures

- Probability notations
- Joint distributions, marginal, conditional
- Representing these quantities as arrays / matrices
- Conditional independences
- BayesNet, and examples
- d-separation, reasoning and inference, probabilistic modelling

# Probability notation and notes

- Probabilities of *propositions*
  - $P(A)$ ,  $P(\text{the sun is shining})$
- Probabilities of *random variables (r.v.)*
  - $P(X = x_1)$ ,  $P(Y = y_1)$ ,  $P(x_1 < X < x_2)$

# Probability notation and notes


- Probabilities of *propositions*
  - $P(A)$ ,  $P(\text{the sun is shining})$
- Probabilities of *random variables (r.v.)*
  - $P(X = x_1)$ ,  $P(Y = y_1)$ ,  $P(x_1 < X < x_2)$
- $P(A)$  usually means  $P(A = \text{True})$  (**A is a proposition, not a variable**)
  - This is a probability **value**
  - Technically,  $P(A)$  is a probability *function*
- $P(X = x_1)$ 
  - This is a probability **value** ( $P(X)$  is a probability *function*)
- $P(X)$ 
  - This is a probability mass function (**p.m.f.**) for discrete r.v. or a probability density function (**p.d.f.**) for continuous r.v.



# Probability notation and notes

- Probabilities of *propositions*
  - $P(A)$ ,  $P(\text{the sun is shining})$
- Probabilities of *random variables (r.v.)*
  - $P(X = x_1)$ ,  $P(Y = y_1)$ ,  $P(x_1 < X < x_2)$
- $P(A)$  usually means  $P(A = \text{True})$  (**A is a proposition, not a variable**)
  - This is a probability **value**
  - Technically,  $P(A)$  is a probability *function*
- $P(X = x_1)$ 
  - This is a probability **value** ( $P(X)$  is a probability *function*)
- $P(X)$ 
  - This is a probability mass function (**p.m.f.**) for discrete r.v. or a probability density function (**p.d.f.**) for continuous r.v.
- Technically, if  $X$  is an r.v., we should not write  **$P(X) = 0.5$** 
  - But rather  **$P(X = x_1) = 0.5$**

# Discrete probability distribution



# Continuous probability distribution


0.4

$$\int_{-\infty}^{\infty} p(X) = 1$$

$p(X)$

0.2

0



$X$

# Joint Probabilities

- A **complete probability model** is a single joint probability distribution over all propositions/variables in the domain
  - $P(X_1, X_2, \dots, X_i, \dots)$
- A particular instance of the world has the probability
  - $P(X_1=x_1 \wedge X_2=x_2 \wedge \dots \wedge X_i=x_i \wedge \dots) = p$

# Joint Probabilities

- A **complete probability model** is a single joint probability distribution over all propositions/variables in the domain
  - $P(X_1, X_2, \dots, X_i, \dots)$
- A particular instance of the world has the probability
  - $P(X_1=x_1 \wedge X_2=x_2 \wedge \dots \wedge X_i=x_i \wedge \dots) = p$
- Rather than stating knowledge as
  - $\text{Raining} \Rightarrow \text{WetGrass}$

# Joint Probabilities

- A **complete probability model** is a single joint probability distribution over all propositions/variables in the domain
  - $P(X_1, X_2, \dots, X_i, \dots)$
- A particular instance of the world has the probability
  - $P(X_1=x_1 \wedge X_2=x_2 \wedge \dots \wedge X_i=x_i \wedge \dots) = p$
- Rather than stating knowledge as
  - $\text{Raining} \Rightarrow \text{WetGrass}$
- We can state it as
  - $P(\text{Raining}, \text{WetGrass}) = 0.15$
  - $P(\text{Raining}, \neg\text{WetGrass}) = 0.01$
  - $P(\neg\text{Raining}, \text{WetGrass}) = 0.04$
  - $P(\neg\text{Raining}, \neg\text{WetGrass}) = 0.8$

# Joint Probabilities

- A **complete probability model** is a single joint probability distribution over all propositions/variables in the domain
  - $P(X_1, X_2, \dots, X_i, \dots)$
- A particular instance of the world has the probability
  - $P(X_1=x_1 \wedge X_2=x_2 \wedge \dots \wedge X_i=x_i \wedge \dots) = p$
- Rather than stating knowledge as
  - $\text{Raining} \Rightarrow \text{WetGrass}$
- We can state it as
  - $P(\text{Raining}, \text{WetGrass}) = 0.15$
  - $P(\text{Raining}, \neg\text{WetGrass}) = 0.01$
  - $P(\neg\text{Raining}, \text{WetGrass}) = 0.04$
  - $P(\neg\text{Raining}, \neg\text{WetGrass}) = 0.8$

	$\neg\text{WetGrass}$	$\text{WetGrass}$
$\neg\text{Raining}$	0.8	0.04
$\text{Raining}$	0.01	0.15

# Joint Probabilities

- A **complete probability model** is a single joint probability distribution over all propositions/variables in the domain
  - $P(X_1, X_2, \dots, X_i, \dots)$
- A particular instance of the world has the probability
  - $P(X_1=x_1 \wedge X_2=x_2 \wedge \dots \wedge X_i=x_i \wedge \dots) = p$
- Rather than stating knowledge as
  - $\text{Raining} \Rightarrow \text{WetGrass}$
- We can state it as
  - $P(\text{Raining}, \text{WetGrass}) = 0.15$
  - $P(\text{Raining}, \neg\text{WetGrass}) = 0.01$
  - $P(\neg\text{Raining}, \text{WetGrass}) = 0.04$
  - $P(\neg\text{Raining}, \neg\text{WetGrass}) = 0.8$

	$\neg\text{WetGrass}$	$\text{WetGrass}$
$\neg\text{Raining}$	0.8	0.04
$\text{Raining}$	0.01	0.15



# Marginal and Conditional Probability

- Marginal Probability
  - Marginal probability (distribution) of X:  $P(X) = \sum_Y P(X, Y)$
  - **Bayesian interpretation:** Probabilities associated with one proposition or variable, **prior** to any evidence
  - E.g.,  $P(\text{WetGrass})$ ,  $P(\neg\text{Raining})$
- Conditional Probability
  - $P(A | B)$  – “The probability of A given that we know B”
  - **Bayesian interpretation:** After (**posterior** to) procuring evidence
  - E.g.,  $P(\text{WetGrass} | \text{Raining})$

$$P(X | Y) = \frac{P(X, Y)}{P(Y)} \quad \text{or} \quad P(X | Y) P(Y) = P(X, Y)$$

Assumes  $P(Y)$  nonzero

# The chain rule: factorizing a joint distribution into marginal and conditionals

$$P(X, Y) = P(X | Y) P(Y)$$

By the Chain Rule

$$P(X, Y, Z) = P(X | Y, Z) P(Y, Z)$$

$$= P(X | Y, Z) P(Y | Z) P(Z)$$

*or, equivalently*

$$= P(X) P(Y | X) P(Z | X, Y)$$

# The chain rule: factorizing a joint distribution into marginal and conditionals

$$P(X, Y) = P(X | Y) P(Y)$$

## By the Chain Rule

$$P(X, Y, Z) = P(X | Y, Z) P(Y, Z)$$

$$= P(X | Y, Z) P(Y | Z) P(Z)$$

*or, equivalently*

$$= P(X) P(Y | X) P(Z | X, Y)$$

- Notes:
- Precedence: ‘|’ is lowest
  - E.g.,  $P(X | Y, Z)$  means which?  
 $P((X | Y), Z)$   
 $P(X | (Y, Z))$

# The chain rule: factorizing a joint distribution into marginal and conditionals

$$P(X, Y) = P(X | Y) P(Y)$$

## By the Chain Rule

$$P(X, Y, Z) = P(X | Y, Z) P(Y, Z)$$

$$= P(X | Y, Z) P(Y | Z) P(Z)$$

*or, equivalently*

$$= P(X) P(Y | X) P(Z | X, Y)$$

- Notes:
- Precedence: ‘|’ is lowest
  - E.g.,  $P(X | Y, Z)$  means which?  
 $P((X | Y), Z)$   
 $P(X | (Y, Z)) \leftarrow$

# Chain Rule implies Bayes' Rule

- Since  $P(X, Y) = P(X | Y) P(Y)$   
and  $P(X, Y) = P(Y | X) P(X)$
- Then  $P(X | Y) P(Y) = P(Y | X) P(X)$

$$P(X | Y) = \frac{P(Y | X) P(X)}{P(Y)}$$

Thomas Bayes: 1701 - 1761



# Chain Rule implies Bayes' Rule

Thomas Bayes: 1701 - 1761



- Since  $P(X, Y) = P(X | Y) P(Y)$   
and  $P(X, Y) = P(Y | X) P(X)$
- Then  $P(X | Y) P(Y) = P(Y | X) P(X)$

$$P(X | Y) = \frac{P(Y | X) P(X)}{P(Y)}$$

Bayes' Rule

# Chain Rule implies Bayes' Rule

Thomas Bayes: 1701 - 1761



- Since  $P(X, Y) = P(X | Y) P(Y)$   
and  $P(X, Y) = P(Y | X) P(X)$
- Then  $P(X | Y) P(Y) = P(Y | X) P(X)$

$$P(X | Y) = \frac{P(Y | X) P(X)}{P(Y)}$$

Bayes' Rule

**Funny fact:** Thomas Bayes is arguably a frequentist.

Stephen Fienberg. "When did Bayesian inference become 'Bayesian'?" *Bayesian analysis* 1.1 (2006): 1-40.  
<https://projecteuclid.org/euclid.ba/1340371071>

# Representing Probability Distributions using linear algebraic data structures (in python)

	<u>Continuous vars</u>	<u>Discrete vars</u>
$P(X)$	Function (of one variable)	m vector
$P(X=x)$	Scalar*	Scalar
$P(X,Y)$	Function of two variables	m×n matrix
$P(X Y)$	Function of two variables	m×n matrix
$P(X Y=y)$	Function of one variable	m vector
$P(X=x Y)$	Function of one variable	n vector
$P(X=x Y=y)$	Scalar*	Scalar

\* - actually zero. Should be  $P(x_1 < X < x_2)$



# Example: Joint probability distribution

From  $P(X, Y)$ , we can always calculate:

$P(X)$	$P(X=x_1)$
$P(Y)$	$P(Y=y_2)$
$P(X Y)$	$P(X Y=y_1)$
$P(Y X)$	$P(Y X=x_1)$
	$P(X=x_1 Y)$
	etc.

# Example: Joint probability distribution

From  $P(X, Y)$ , we can always calculate:

$P(X)$	$P(X=x_1)$
$P(Y)$	$P(Y=y_2)$
$P(X Y)$	$P(X Y=y_1)$
$P(Y X)$	$P(Y X=x_1)$
	$P(X=x_1 Y)$
	etc.

		<b>X</b>		
		$x_1$	$x_2$	$x_3$
<b>Y</b>	$y_1$	0.2	0.1	0.1
	$y_2$	0.1	0.2	0.3

**P(X,Y)**

$x_1$

$x_2$

$x_3$

$y_1$

0.2

0.1

0.1

$y_2$

0.1

0.2

**0.3**

**P(X,Y)**

$x_1$

$x_2$

$x_3$

$y_1$

0.2

0.1

0.1

$y_2$

0.1

0.2

**0.3**

**P(X)**

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(Y)**

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**



**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

**P(Y|X)**

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

$P(X=x_1, Y=y_2) = ?$

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

$P(X=x_1, Y=y_2) = ?$

$P(X=x_1) = ?$

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

$P(X=x_1, Y=y_2) = ?$

$P(X=x_1) = ?$

$P(Y=y_2) = ?$

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

$P(X=x_1, Y=y_2) = ?$

$P(X=x_1) = ?$

$P(Y=y_2) = ?$

$P(X|Y=y_1) = ?$

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75

**P(X,Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.2	0.1	0.1
$y_2$	0.1	0.2	0.3

**P(X)**

$x_1$	$x_2$	$x_3$
0.3	0.3	0.4

**P(Y)**

$y_1$	0.4
$y_2$	0.6

**P(X|Y)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.5	0.25	0.25
$y_2$	0.167	0.333	0.5

$P(X=x_1, Y=y_2) = ?$

$P(X=x_1) = ?$

$P(Y=y_2) = ?$

$P(X|Y=y_1) = ?$

$P(X=x_1|Y) = ?$

**P(Y|X)**

	$x_1$	$x_2$	$x_3$
$y_1$	0.667	0.333	0.25
$y_2$	0.333	0.667	0.75



# Quick checkpoint

- Probability notations
  - $P(A)$  is a number when  $A$  is an event / predicate.
  - $P(X)$  is a vector/function when  $X$  is a random variable.
- Joint probability distribution
  - Enumerating all combinations of events.
  - All values the random variables can take.
  - Assign a non-negative value to each.
- Marginals, conditionals
  - How they are related: Chain rule, Bayes rule

# How can a joint-distribution help us?

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

$$= \operatorname{argmax}_y P(Y = y | X)$$

- We can make Bayesian inference (probabilistic reasoning)
  - Sherlock Holmes:  $P(\text{Murderer} | \text{Observed Evidence})$
  - Doctor:  $P(\text{Disease} | \text{Symptoms})$ ,  $P(\text{Effect} | \text{Treatment})$
  - Parenting:
    - $P(\text{Dirty Diaper, Hungry, Lonely} | 5 \text{ a.m., Baby crying})$
    - $P(\text{Baby crying at 5 a.m.} | \text{feeding at 2 a.m.})$
    - $P(\text{Baby crying at 5 a.m.} | \text{feeding at 1 a.m.})$

# How can a joint-distribution help us?

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

- With the joint distribution, we can do anything we want

$$= \operatorname{argmax}_y P(Y = y | X)$$

- We can make Bayesian inference (probabilistic reasoning)
  - Sherlock Holmes:  $P(\text{Murderer} | \text{Observed Evidence})$
  - Doctor:  $P(\text{Disease} | \text{Symptoms})$ ,  $P(\text{Effect} | \text{Treatment})$
  - Parenting:
    - $P(\text{Dirty Diaper, Hungry, Lonely} | 5 \text{ a.m., Baby crying})$
    - $P(\text{Baby crying at 5 a.m.} | \text{feeding at 2 a.m.})$
    - $P(\text{Baby crying at 5 a.m.} | \text{feeding at 1 a.m.})$

# How can a joint-distribution help us?

*argmax<sub>x</sub> P(Y=y | X=x)*

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

- With the joint distribution, we can do anything we want

- Design classifiers

$$y = \operatorname{argmax}_y P(Y = y | X)$$

- We can make Bayesian inference (probabilistic reasoning)

- Sherlock Holmes: P(Murderer | Observed Evidence)
- Doctor: P( Disease | Symptoms ), P(Effect | Treatment)
- Parenting:

- P( Dirty Diaper, Hungry, Lonely | 5 a.m., Baby crying)

- P(Baby crying at 5 a.m. | feeding at 2 a.m.)

- P(Baby crying at 5 a.m. | feeding at 1 a.m.)

# How can a joint-distribution help us?

*argmax<sub>x</sub> P(x, y) P(y)*

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

- With the joint distribution, we can do anything we want
  - Design classifiers
  - We can make Bayesian inference (probabilistic reasoning)
    - Sherlock Holmes:  $P(\text{Murderer} \mid \text{Observed Evidence})$
    - Sherlock Holmes:  $P(\text{Murderer} \mid \text{Observed Evidence})$
    - Doctor:  $P(\text{Disease} \mid \text{Symptoms})$ ,  $P(\text{Effect} \mid \text{Treatment})$
    - Parenting:
      - $P(\text{Dirty Diaper, Hungry, Lonely} \mid 5 \text{ a.m., Baby crying})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 2 a.m.})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 1 a.m.})$

# How can a joint-distribution help us?

*argmmaa x y xx x y yy x y PP Y=y YY=yy Y=y XX)*

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

- With the joint distribution, we can do anything we want
  - Design classifiers
  - We can make Bayesian inference (probabilistic reasoning)
    - Sherlock Holmes:  $P(\text{Murderer} \mid \text{Observed Evidence})$
    - Doctor:  $P(\text{Disease} \mid \text{Symptoms})$ ,  $P(\text{Effect} \mid \text{Treatment})$
    - Doctor:  $P(\text{Disease} \mid \text{Symptoms})$ ,  $P(\text{Effect} \mid \text{Treatment})$
    - Parenting:
      - $P(\text{Dirty Diaper, Hungry, Lonely} \mid 5 \text{ a.m., Baby crying})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 2 a.m.})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 1 a.m.})$

# How can a joint-distribution help us?

*argmmaa x y xx x y yy x y PP Y=y YY=yy Y=y XX)*

A principled way to model the world

- Handles missing data/variables
- Easy to incorporate prior knowledge

- With the joint distribution, we can do anything we want
  - Design classifiers
  - We can make Bayesian inference (probabilistic reasoning)
    - Sherlock Holmes:  $P(\text{Murderer} \mid \text{Observed Evidence})$
    - Doctor:  $P(\text{Disease} \mid \text{Symptoms})$ ,  $P(\text{Effect} \mid \text{Treatment})$
    - Parenting:
      - $P(\text{Dirty Diaper, Hungry, Lonely} \mid 5 \text{ a.m., Baby crying})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 2 a.m.})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 1 a.m.})$
      - $P(\text{Baby crying at 5 a.m.} \mid \text{feeding at 1 a.m.})$

# (3 min discussion) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000



# (3 min discussion) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- Questions:
  - What is the dimension(s) of the joint probability table?
  - How many **independent** parameters are needed to represent this distribution?

# (3 min discussion) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- Questions:
  - What is the dimension(s) of the joint probability table?
  - How many **independent** parameters are needed to represent this distribution?
- Answers:  $3000^N \times 15$  and  $3000^N \times 15 - 1$

# (3 min discussion) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- Questions:
  - What is the dimension(s) of the joint probability table?
  - How many **independent** parameters are needed to represent this distribution?
- Answers:  $3000^N \times 15$  and  $3000^N \times 15 - 1$ 

**Note: Exponential in the number of words per document...**

# Statistical Independences

## (Marginal / absolute) Independence

- X and Y are independent iff
  - $P(X, Y) = P(X) P(Y)$  [by definition]

# Statistical Independences

## (Marginal / absolute) Independence

- X and Y are independent iff
  - $P(X, Y) = P(X) P(Y)$  [by definition]
  - $P(X | Y) = P(X)$

# Statistical Independences

## (Marginal / absolute) Independence

- X and Y are independent iff
  - $P(X, Y) = P(X) P(Y)$  [by definition]
  - $P(X | Y) = P(X)$     Since  $P(X | Y) = P(X, Y)/P(Y) = P(X) P(Y)/P(Y)$

# Statistical Independences

## (Marginal / absolute) Independence

- X and Y are independent iff
  - $P(X, Y) = P(X) P(Y)$  [by definition]
  - $P(X | Y) = P(X)$  Since  $P(X | Y) = P(X, Y)/P(Y) = P(X) P(Y)/P(Y)$

## Conditional Independence

- If X and Y are (conditionally) independent given Z, then
  - $P(X | Y, Z) = P(X | Z)$

# Statistical Independences

## (Marginal / absolute) Independence

- X and Y are independent iff
  - $P(X, Y) = P(X) P(Y)$  [by definition]
  - $P(X | Y) = P(X)$     Since  $P(X | Y) = P(X, Y)/P(Y) = P(X) P(Y)/P(Y)$

## Conditional Independence

- If X and Y are (conditionally) independent given Z, then
  - $P(X | Y, Z) = P(X | Z)$
  - Example:
    - $P(\text{WetGrass} | \text{Season}, \text{Rain}) = P(\text{WetGrass} | \text{Rain})$



# Example of Conditional Independence

- In practice, conditional independence is more common than marginal independence.

# Example of Conditional Independence

- In practice, conditional independence is more common than marginal independence.
  - $P(\text{Final exam grade} \mid \text{Weather}) \neq P(\text{Final exam grade})$ 
    - i.e., they are not independent

# Example of Conditional Independence

- In practice, conditional independence is more common than marginal independence.
  - $P(\text{Final exam grade} \mid \text{Weather}) \neq P(\text{Final exam grade})$ 
    - i.e., they are not independent
  - $P(\text{Final exam grade} \mid \text{Weather, Effort}) = P(\text{Final exam grade} \mid \text{Effort})$ 
    - But they are conditionally independent given Effort

# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000

# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- In addition, assume that: *Word 1, ..., Word N* are **mutually independent** given *Author*
  - $P(\text{Word } 1, \dots, \text{Word } N \mid \text{Author}) = P(\text{Word } 1 \mid \text{Author}) \times \dots \times P(\text{Word } N \mid \text{Author})$

# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- In addition, assume that: *Word 1, ..., Word N* are **mutually independent** given *Author*
  - $P(\text{Word } 1, \dots, \text{Word } N \mid \text{Author}) = P(\text{Word } 1 \mid \text{Author}) \times \dots \times P(\text{Word } N \mid \text{Author})$

- Answer:  $14 + 2999 \times 15 \times N$ . ( **Linear in N** )

# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- In addition, assume that: *Word 1, ..., Word N* are **mutually independent** given *Author*
  - $P(\text{Word } 1, \dots, \text{Word } N \mid \text{Author}) = P(\text{Word } 1 \mid \text{Author}) \times \dots \times P(\text{Word } N \mid \text{Author})$
- Questions:
- Answer:  $14 + 2999 \times 15 \times N$ . ( **Linear in N** )

# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- In addition, assume that: *Word 1, ..., Word N* are **mutually independent** given *Author*
  - $P(\text{Word } 1, \dots, \text{Word } N \mid \text{Author}) = P(\text{Word } 1 \mid \text{Author}) \times \dots \times P(\text{Word } N \mid \text{Author})$
- Questions:
  - What is the dimension(s) of the joint probability table?
- Answer:  $14 + 2999 \times 15 \times N$ . ( **Linear in N** )



# (Example continued) Modeling the world with probability distribution

- Example: Author attribution
- Variables: *Word 1, Word 2, Word 3, ..., Word N, Author*
  - 15 authors in total: {Dickens, Shakespeare, Jane Austen, Tolkien, George RR. Martin, ... , Xueqin Cao, Douglas Adams}
  - A vocabulary of size 3000
- In addition, assume that: *Word 1, ..., Word N* are **mutually independent** given *Author*
  - $P(\text{Word } 1, \dots, \text{Word } N \mid \text{Author}) = P(\text{Word } 1 \mid \text{Author}) \times \dots \times P(\text{Word } N \mid \text{Author})$
- Questions:
  - What is the dimension(s) of the joint probability table?
  - How many **independent** parameters are needed to represent this distribution?
- Answer:  $14 + 2999 \times 15 \times N$ . ( **Linear in N** )

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world
  - Let the variables be binary. How many numbers (“free parameters”) does it take to define the joint distribution?

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world
  - Let the variables be binary. How many numbers (“free parameters”) does it take to define the joint distribution?
    - Defined by  $2^N - 1$  parameters

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world
  - Let the variables be binary. How many numbers (“free parameters”) does it take to define the joint distribution?
    - Defined by  $2^N - 1$  parameters
- If the variables are independent, then
$$P(X_1, X_2, \dots, X_N) = P(X_1) P(X_2) \dots P(X_N)$$

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world
  - Let the variables be binary. How many numbers (“free parameters”) does it take to define the joint distribution?
    - Defined by  $2^N - 1$  parameters
- If the variables are independent, then
$$P(X_1, X_2, \dots, X_N) = P(X_1) P(X_2) \dots P(X_N)$$
  - How many numbers does it take to define the joint distribution?

# Quiz time: Representing a joint Probability

- Joint probability:  $P(X_1, X_2, \dots, X_N)$ 
  - Defines the probability for any possible state of the world
  - Let the variables be binary. How many numbers (“free parameters”) does it take to define the joint distribution?
    - Defined by  $2^N - 1$  parameters
- If the variables are independent, then
$$P(X_1, X_2, \dots, X_N) = P(X_1) P(X_2) \dots P(X_N)$$
  - How many numbers does it take to define the joint distribution?
    - Just  $N$  parameters!

# Tradeoffs in our model choices

## Fully Independent

$$P(X_1, X_2, \dots, X_N) \\ = P(X_1) P(X_2) \dots P(X_N)$$

$$O(N)$$

## Fully general

$$P(X_1, X_2, \dots, X_N)$$

$$O(e^N)$$

Space / computation efficiency




Expressiveness

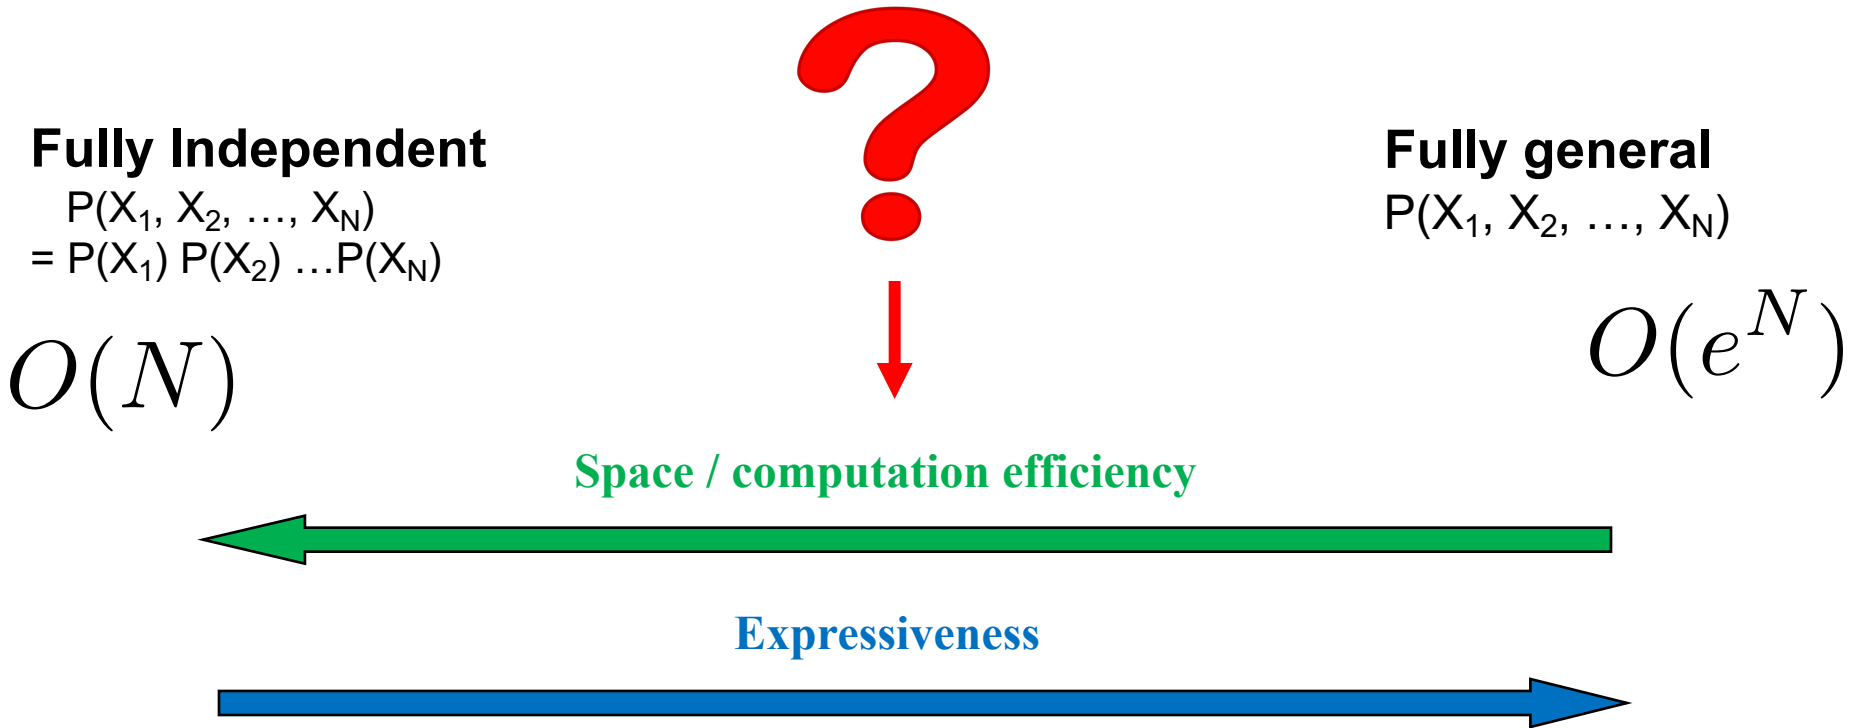




# Tradeoffs in our model choices

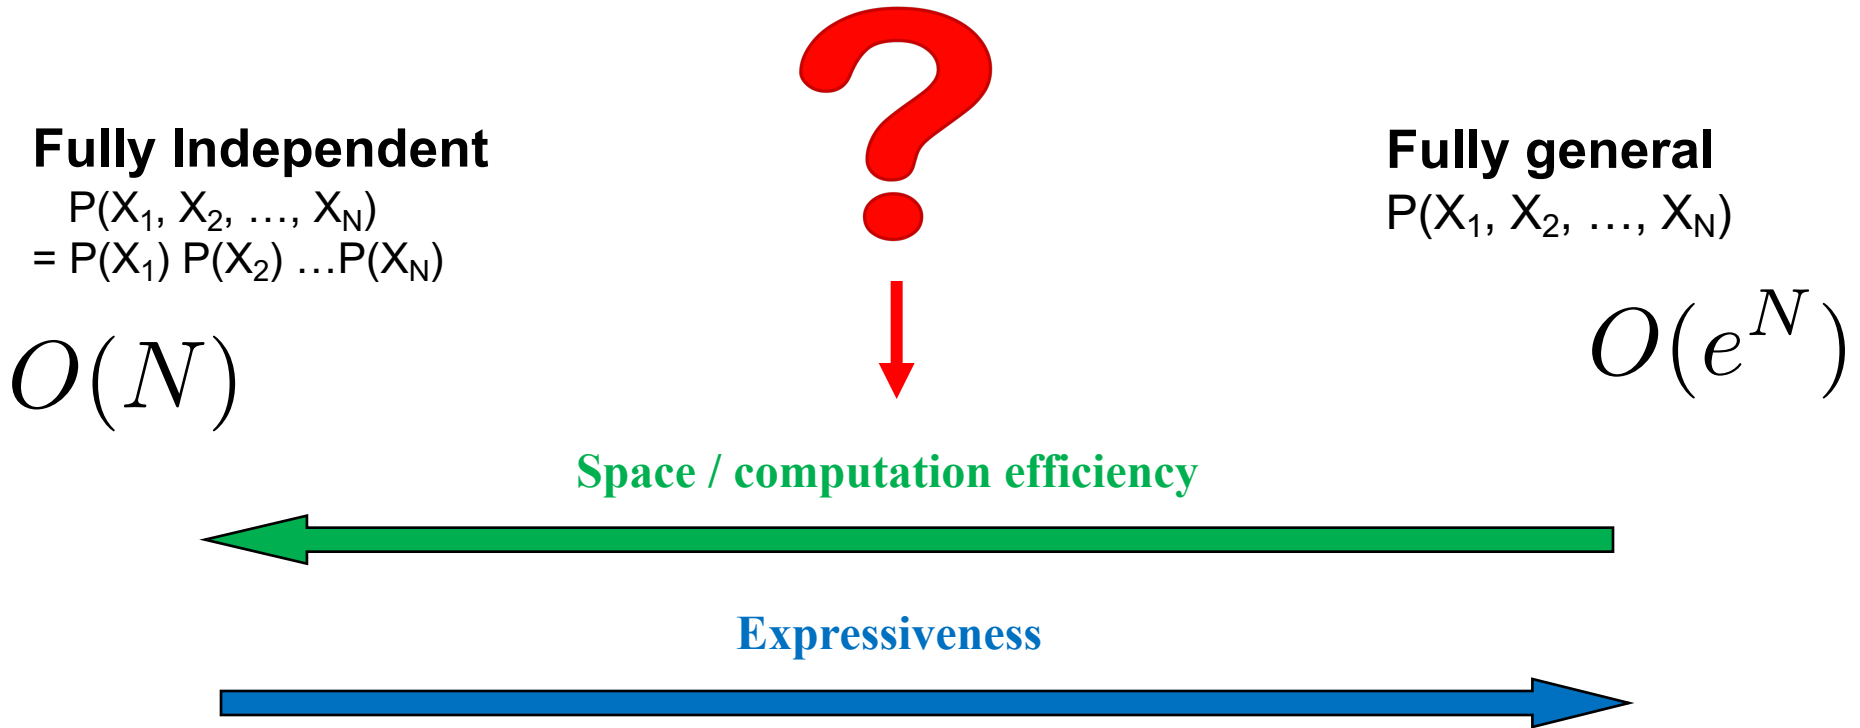


# Tradeoffs in our model choices



**Idea:**


# Tradeoffs in our model choices



## Idea:

1. Independent groups of variables?

# Tradeoffs in our model choices



## Idea:

1. Independent groups of variables?
- 2. Conditional independences?**

# Key points of today's lecture

- Probability notations
  - Distinguish between events and random variables, apply rules of probabilities
- Representing a joint-distribution
  - number of parameters exponential in the number of variables
  - Calculating marginals and conditionals from the joint-distribution.
- Conditional independences and factorization of joint-distributions
  - Saves parameters, often exponential improvements

## Next lectures

- Bayes networks, directed graphical models.
- Read off conditional independences from the graph!
- More examples