15-251: Great Ideas in Theoretical Computer Science

Lecture 18: Introduction to Randomness and Probability Theory Review

October 30th, 2018

Randomness and the Universe

Randomness and the Universe

Does the universe have "true" randomness?

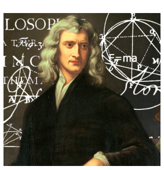

Newtonian physics:

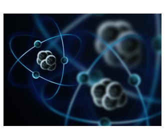

Quantum physics:

Randomness is an essential tool in
modeling and analyzing nature.

It also plays a key role in **computer science**.

Randomness and Computer Science

Statistics via Sampling

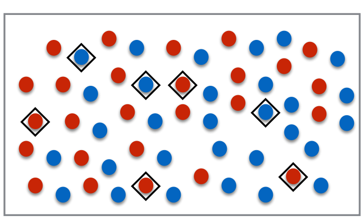

Population: 300m **Random sample size**: 2000

Theorem:

Randomized Algorithms

Dimer Problem:

Given a region, in how many different ways can you tile it with 2x1 rectangles (dominoes)?

e.g.

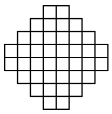

→ 1024 tilings

Captures thermodynamic properties of matter.

Distributed Computing

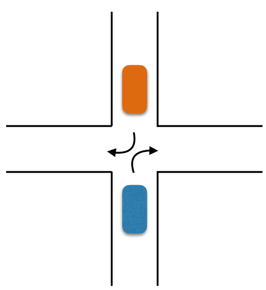

Nash Equilibria in Games

The Chicken Game

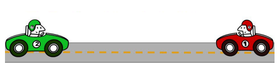

	Swerve	Straight
Swerve	1 1	0 2
Straight	2 0	-3 -3

Theorem (Nash):

Cryptography

Adversary
Eavesdropper

"loru23n8uladjkfb!#@"

"Meet me at 5."

encryption

"loru23n8uladjkfb!#@"

"loru23n8uladjkfb!#@"

decryption

"Meet me at 5."

Shannon:

Error-Correcting Codes

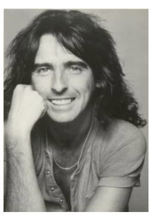

"bit.ly/vrxUBN"

noisy channel

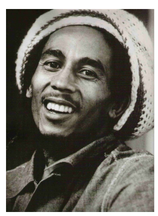

Alice

Bob

Each symbol can be corrupted with a certain probability.

How can Alice still get the message across?

Communication Complexity

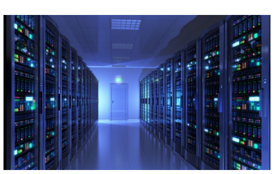

Want to check if the contents of two databases are exactly the same.

How many bits need to be communicated?

Quantum Computing

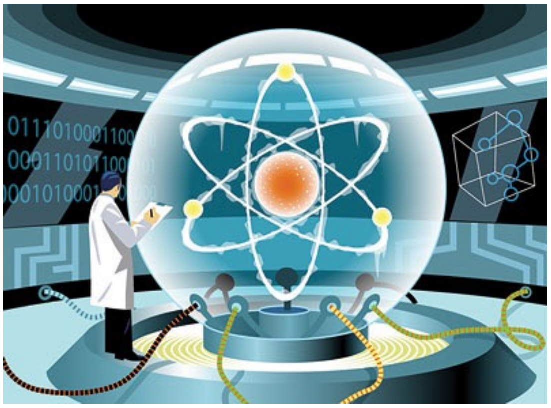

Probability Theory:
The CS Approach

The Big Picture

The Non-CS Approach

Real World → Mathematical Model

(random) experiment/process

probability space

The Big Picture

Real World → Mathematical Model

Flip a coin.

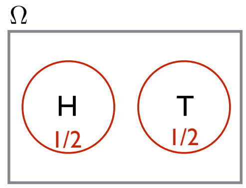

Ω = "sample space"
= set of all possible outcomes

$\Pr : \Omega \to [0, 1]$ prob. distribution

$\sum_{\ell \in \Omega} \Pr[\ell] = 1$ (why?)

The Big Picture

Real World → Mathematical Model

Flip a coin.

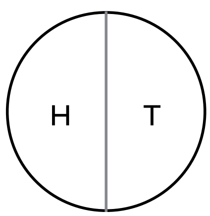

unit pie, **area = 1**

$$\Pr[\text{outcome}] = \text{area of outcome}$$

$$= \frac{\text{area of outcome}}{\text{area of pie}}$$

The Big Picture

Real World → Mathematical Model

Flip two coins.

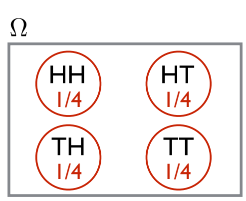

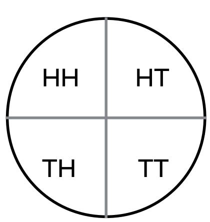

The Big Picture

Real World → Mathematical Model

Flip a coin.
If it is Heads, throw a 3-sided die.
If it is Tails, throw a 4-sided die.

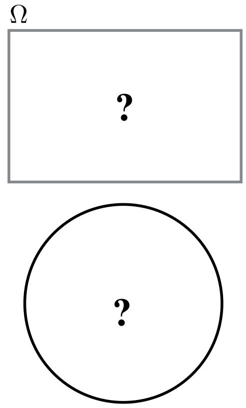

The Big Picture

The CS Approach

The Big Picture

Flip a coin.
If it is Heads, throw a 3-sided die.
If it is Tails, throw a 4-sided die.

```
flip <— Bernoulli(1/2)
if flip = 1: # i.e. Heads
    die <— RandInt(3)
else:
    die <— RandInt(4)
```

Probability Tree

```
flip <— Bernoulli(1/2)
if flip = H:
    die <— RandInt(3)
else:
    die <— RandInt(4)
```

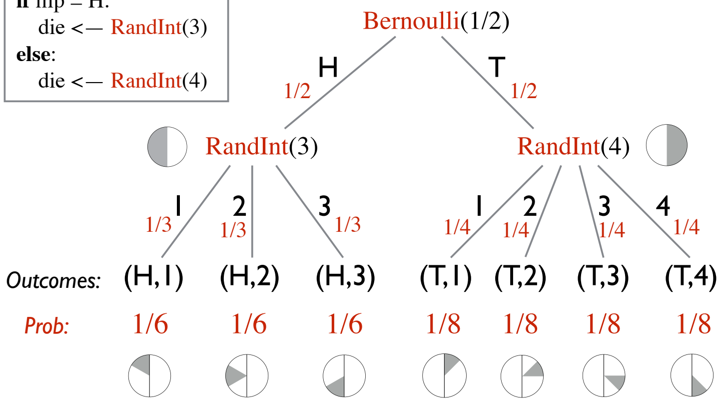

Events

Real World → Code → Probability Tree

Flip a coin.
If it is Heads, throw a 3-sided die.
If it is Tails, throw a 4-sided die.
→

```
flip <— Bernoulli(1/2)
if flip = H:
    die <— RandInt(3)
else:
    die <— RandInt(4)
```
→

What is the probability $\boxed{\text{die roll is} \geq 3}$?
"event" →

Events

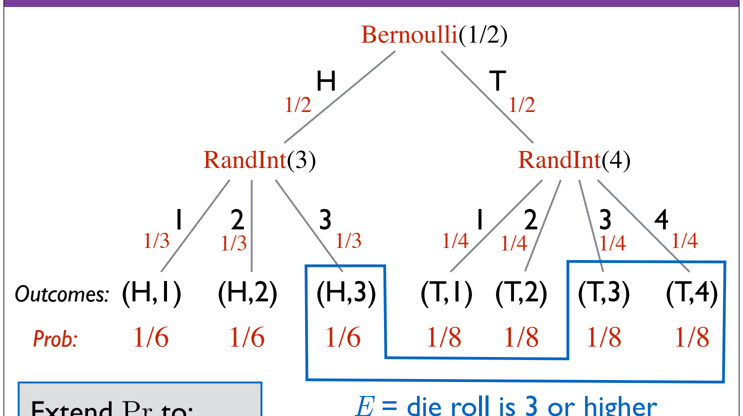

Extend $\Pr$ to:
$\Pr : \mathcal{P}(\Omega) \to [0,1]$

Conditional Probability

Real World → Code → Probability Tree

Flip a coin.
If it is Heads, throw a 3-sided die.
If it is Tails, throw a 4-sided die.

```
flip <— Bernoulli(1/2)
if flip = H:
    die <— RandInt(3)
else:
    die <— RandInt(4)
```

What is the probability of flipping Heads $\boxed{\text{given the die roll is } \geq 3}$?

conditioning on partial information

Conditional Probability

Revising probabilities based on '*partial information*'.

'*partial information*' = event E

Conditioning on E = Assuming/promising E has happened

Conditional Probability

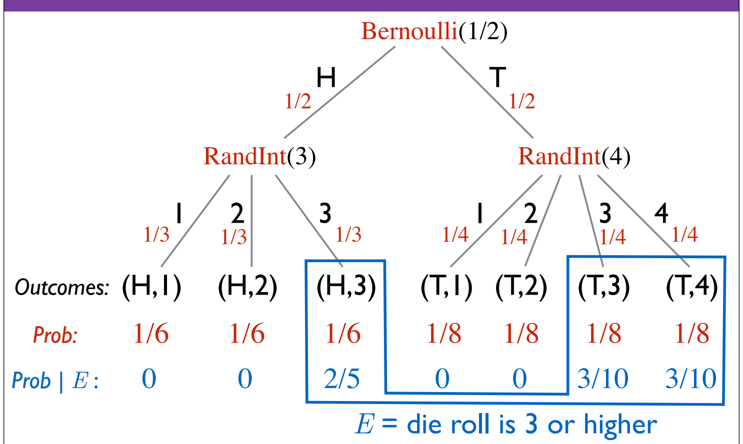

Conditioning

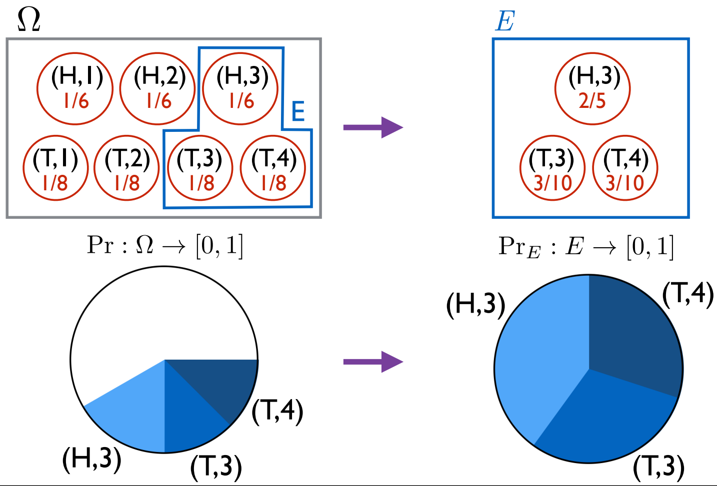

Conditioning

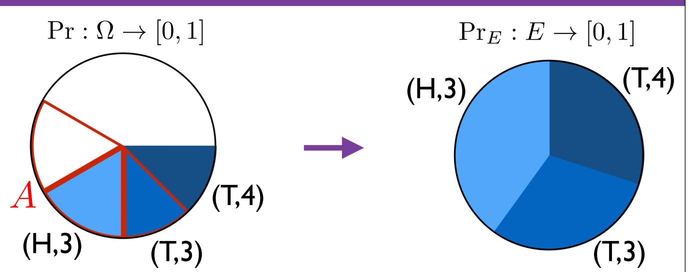

$\Pr[A \mid E] \quad =$

(cannot condition on an event with prob. 0)

Conditional Probability –> Chain Rule

$\Pr[A \cap B] =$

"For A and B to occur:
- first A must occur
- then B must occur given that A occured"

Generalizes to more than two events.
e.g.

$\Pr[A \cap B \cap C] =$

Conditional Probability –> LTP

LTP = Law of Total Probability

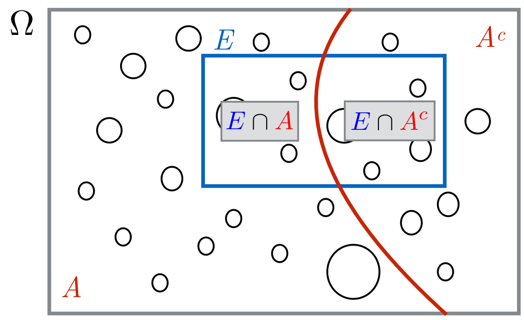

$\Pr[E] =$

SUMMARY SO FAR

Real World → Code → Probability Tree
=
Mathematical Model
- set of outcomes Ω
- a prob. associated with each outcome.

Events

Conditional probability:
$\Pr[A \mid B] = \Pr[A \cap B] / \Pr[B]$

Chain rule:
$\Pr[A \cap B] = \Pr[A] \cdot \Pr[B \mid A]$

Law of total probability:
$\Pr[B] = \Pr[A] \cdot \Pr[B \mid A] + \Pr[A^c] \cdot \Pr[B \mid A^c]$

Independent events:
$\Pr[A \cap B] = \Pr[A] \cdot \Pr[B]$

Union bound:
$\Pr[A \cup B] \leq \Pr[A] + \Pr[B]$

Random Variables

What is a Random Variable?

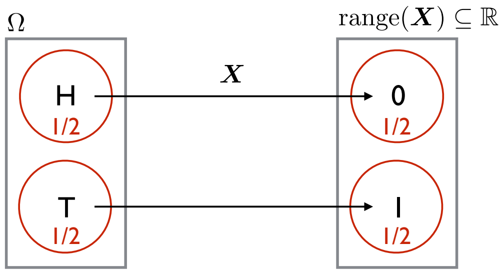

<u>typical description:</u> $\boldsymbol{X}$ = number of Tails

What is a Random Variable?

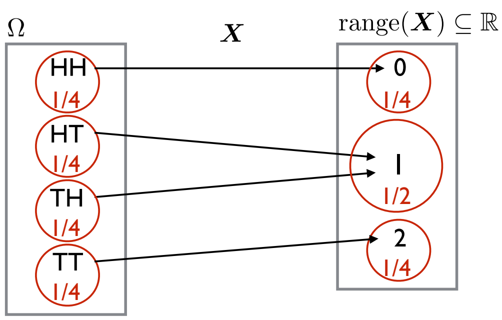

<u>typical description:</u> $\boldsymbol{X}$ = number of Tails

Why?

Often we are interested in numerical outcomes
(e.g. # Tails we see if we toss *n* coins)

but initially outcomes are best expressed non-numerically.
(e.g. an outcome is a sequence of *n* coin tosses)

We like talking about mean values (averages), variance, etc..

What is a Random Variable?

2nd Definition:

Example:

S <— RandInt(6) + RandInt(6)
if S = 12: I <— 1
else: I <— 0

Random variables:

What is a Random Variable?

S <— RandInt(6) + RandInt(6)
if S = 12: I <— 1
else: I <— 0

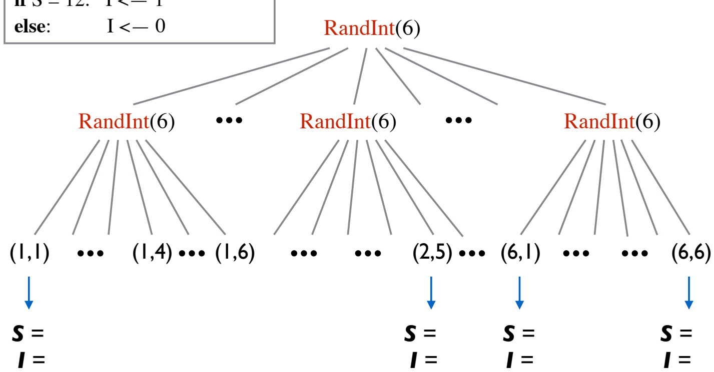

Expectation of a Random Variable

Expected Value = Mean = (Weighted) Average

$$\text{Weighted Average} = \sum_{\text{elements } e} \text{value}(e) \cdot \text{weight}(e)$$

Example:

Weight	Value
30% Final	85
20% Midterm	75
50% Homework	82

Weighted Average = $0.3 \cdot 85 + 0.2 \cdot 75 + 0.5 \cdot 82 = 81.5$

Expectation of a Random Variable

Expected value of a random variable X:

$$\mathbf{E}[X] \stackrel{\text{def}}{=}$$

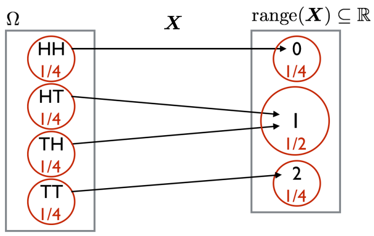

Expectation of a Random Variable

Example

Let X = Bernoulli(6), Y = Bernoulli(6), Z = Bernoulli(6)

Let $S = X + Y + Z$

$\mathbf{E}[S]$

$= 3 \cdot \Pr[S = 3] + 4 \cdot \Pr[S = 4] + \cdots + 18 \cdot \Pr[S = 18]$

lot's of arithmetic :-(

$= 10.5$

Most Useful Equality in Probability Theory:

Linearity of Expectation

Linearity of Expectation

Example

Let X = Bernoulli(6), Y = Bernoulli(6), Z = Bernoulli(6)

Let $S = X + Y + Z$

$$\mathbf{E}[S] =$$

<u>Next Time:</u>

Introduction to Randomized Algorithms