Uncertainty (Ch. 13)

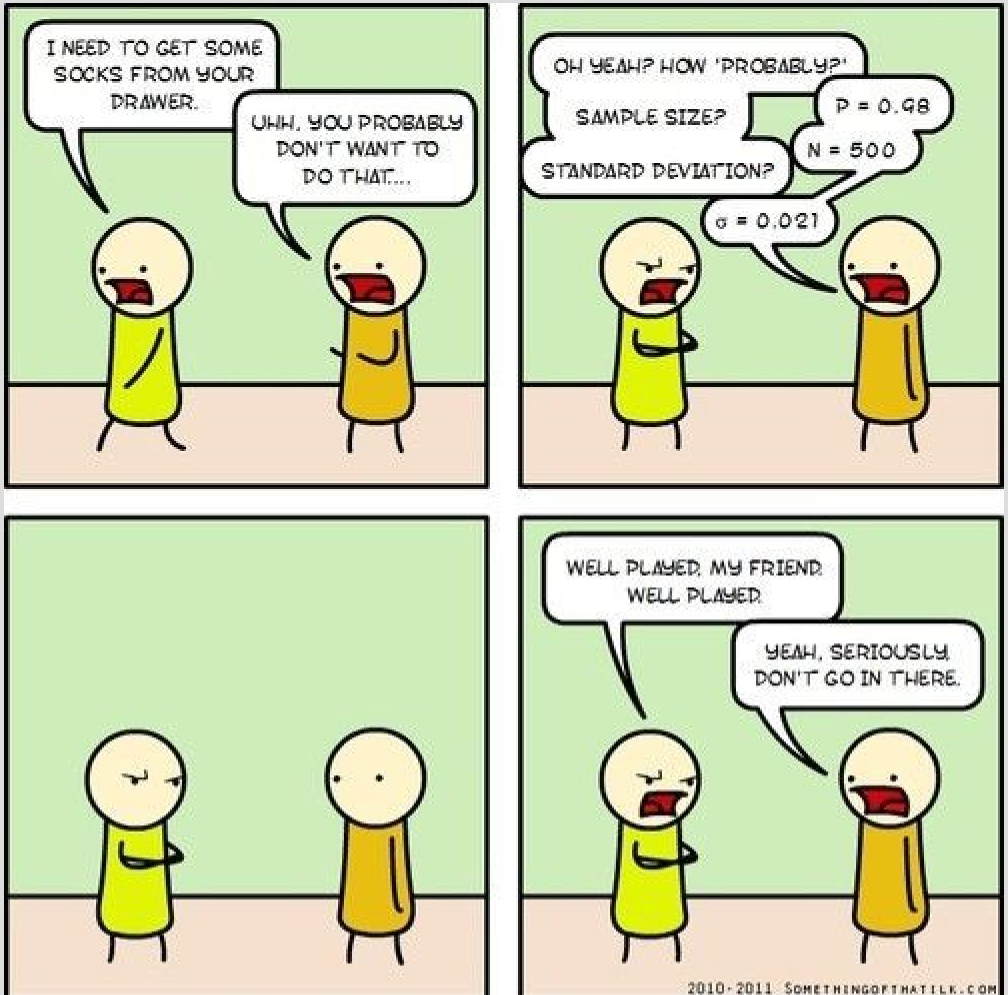

Representation

Robots quite often do not know everything about problem (uncertainty):

- partial observability
- non-deterministic actions

For example, if you were making a poker AI:

1. You cannot see the other player’s cards, so you have to reason without that info
2. When you draw/exchange cards, you do not know what new card you will get

Representation

One simple way is to use belief states and track each possible outcome

This is quite often too burdensome as:

1. Large number of possible states
2. Would need to plan/decide for each state
3. Possible that no single plan is guaranteed to exist (very true in "games of chance")

1 & 2 especially annoying for low probability

Representation

We also need to reason on affects of actions or the info that we do have

Logic would be one possibility, but often does not work well with uncertainty

Consider: your friend sat down next to you and has wet hair... you guess they got out of the shower recently

$$WetHair \Rightarrow TookShower$$

Representation

This however is a bit simplistic

$$WetHair \Rightarrow TookShower$$

There could be other reasons for wet hair...

$$WetHair \Rightarrow TookShower \vee Raining \vee LikesToRollInPuddles \ldots$$

This should include all possible outcomes, yet be able to combine knowledge:

If everyone has wet hair, probably rain

Representation

Explicitly writing out all possibilities:

1. Makes it more difficult to reason/deduce (for tractability, ignore “unlikely” reasons)
2. Some rules or the exact requirements of rules might not be known

For all these reasons, using logical inference with uncertainty can be cumbersome

Instead, probabilistic reasoning works better

Probability

Quite often when dealing with probability,
it is useful to evaluate how good outcomes are

For example, studying for tests:
You do not know what will be asked,
so you have to guess what topics to review

At some point, you feel “confident enough”
about the material and stop

Probability

Often it is not even possible to have a 100% chance of success (e.g. cannot win every hand of poker or ace every test)

Instead, if we have a utility or value for states, we will try to achieve the maximum expected utility

	Percent	Utility/Value
Gamble	99% 1%	0 100
Go home	100%	10

Probability

The maximum expected utility can be thought of as the "best on average" (expectation of a random variable)

For the rest of today, we will go over some probability basics (will use a lot in this class)

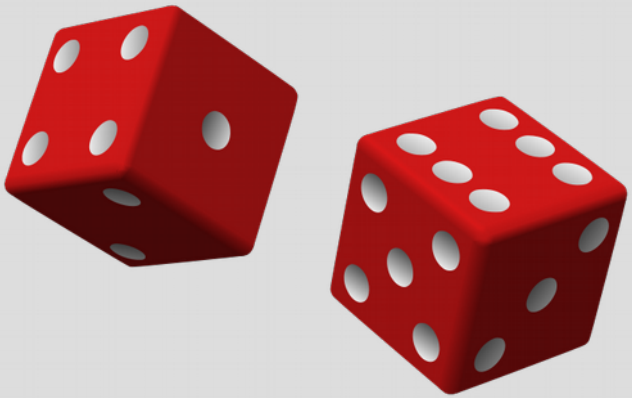

Probability: the basics

A probability of an event (or proposition) is:

$$P(x) = \frac{\text{number of times } x \text{ happens}}{\text{number of possibilities}}$$

For example, the probability that a 6-sided die rolls up odd is:

Possible rolls:	1	2	3	4	5	6
Is odd?	Y	N	Y	N	Y	N

P(die = odd) = 3 / 6 = 0.5

Probability: notation

Some notation blah-blah (from the book):

ω - one possible state/outcome

Ω - all possible outcomes

ϕ - an “event” or subset of possible outcomes

(I will quite often just call this “A”)

Probability: notation

Some notation blah-blah (from the book):

ω - one possible state/outcome

Ω - all possible outcomes

ϕ - an “event” or subset of possible outcomes
(I will quite often just call this “A”)

So in the dice example:

ω - The die is 2 (one possibility)

Ω - <1, 2, 3, 4, 5, 6> (all possibilities)

ϕ - <1, 3, 5> (the die is odd)

Probability: the basics

So in the dice example:

ω - The die is 2 (one possibility)

Ω - <1, 2, 3, 4, 5, 6> (all possibilities)

ϕ - <1, 3, 5> (the die is odd) $\mathrm{P}(\phi) = \sum_{\omega \in \phi} P(\omega)$

(or: $P(Die = odd)$)

Probabilities also need to:

- Be between zero and one:

  $$0 \leq P(\omega) \leq 1$$

- Add up to 100%:

  $$1 = \sum_{\omega \in \Omega} P(\omega)$$

Probability: the basics

Beyond these properties of probability, we only really need three more facts:

1. Conditional probability

$$P(A|B) = \frac{P(A,B)}{P(B)}$$
$$P(A,B) = P(A|B)P(B)$$

(this is definition)

2. Probability of opposite happening

$$P(A) + P(\neg A) = 1$$

3. Definition of "or"

$$P(A \text{ or } B) = P(A) + P(B) - P(A,B)$$

Probability: terminology :(

Terminology side note:

P(A) is "unconditional" or "prior"

P(A|B) is "conditional" or "posterior"

P(A,B) is "joint" probability $P(A,B) = P(A \cap B) = P(AB)$

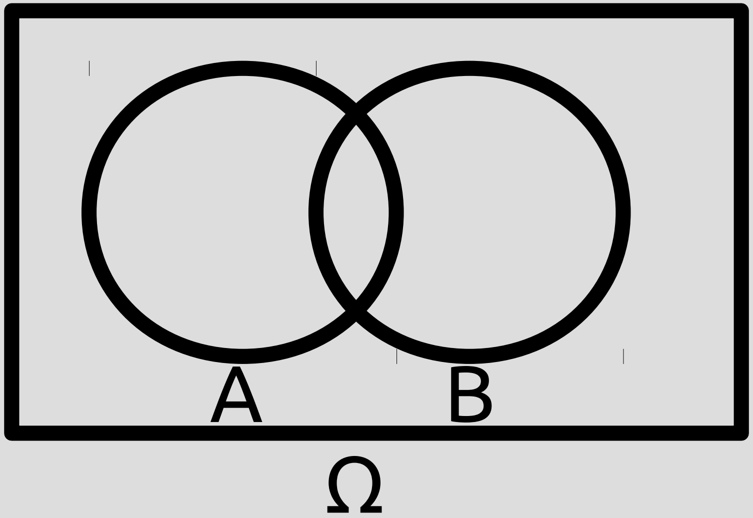

$$P(A \text{ or } B) = P(A) + P(B) - P(A,B)$$

Proof by picture

Probability: the basics

Proof:

$$\begin{aligned}
P(A) &= \sum_{\omega \in A} P(\omega) \\
&= \sum_{\omega \in A} P(\omega) + \left(\sum_{\omega \in \neg A} P(\omega) - \sum_{\omega \in \neg A} P(\omega) \right) \\
&= \left(\sum_{\omega \in A} P(\omega) + \sum_{\omega \in \neg A} P(\omega) \right) - \sum_{\omega \in \neg A} P(\omega) \\
&= 1 - \sum_{\omega \in \neg A} P(\omega) \\
&= 1 - P(\neg A)
\end{aligned}$$

Probability: example

I showed earlier (brute force) that if {A = die roll}, then $P(A = odd) = 0.5$

Why don't you try to compute the following:
(B, C, D, etc. are other die rolls)

1. Sum of two dice is odd: $P(A + B = odd)$
2. Sum of three dice is odd: $P(A + B + C = odd)$
3. 20 dice: $P(A + B + C + \cdots + T = odd)$
   (can you prove this rather than guess?)

Probability: example

To get some intuition, let's brute force the 2-dice example: $P(A+B=odd)$

$$= P(A+B=3)+P(A+B=5)+P(A+B=7)+P(A+B=9)+P(A+B=11)$$
$$= \frac{2}{36}+\frac{4}{36}+\frac{6}{36}+\frac{4}{36}+\frac{2}{36}$$
$$= \frac{18}{36} = 0.5$$

At this point you might guess what the other answers are

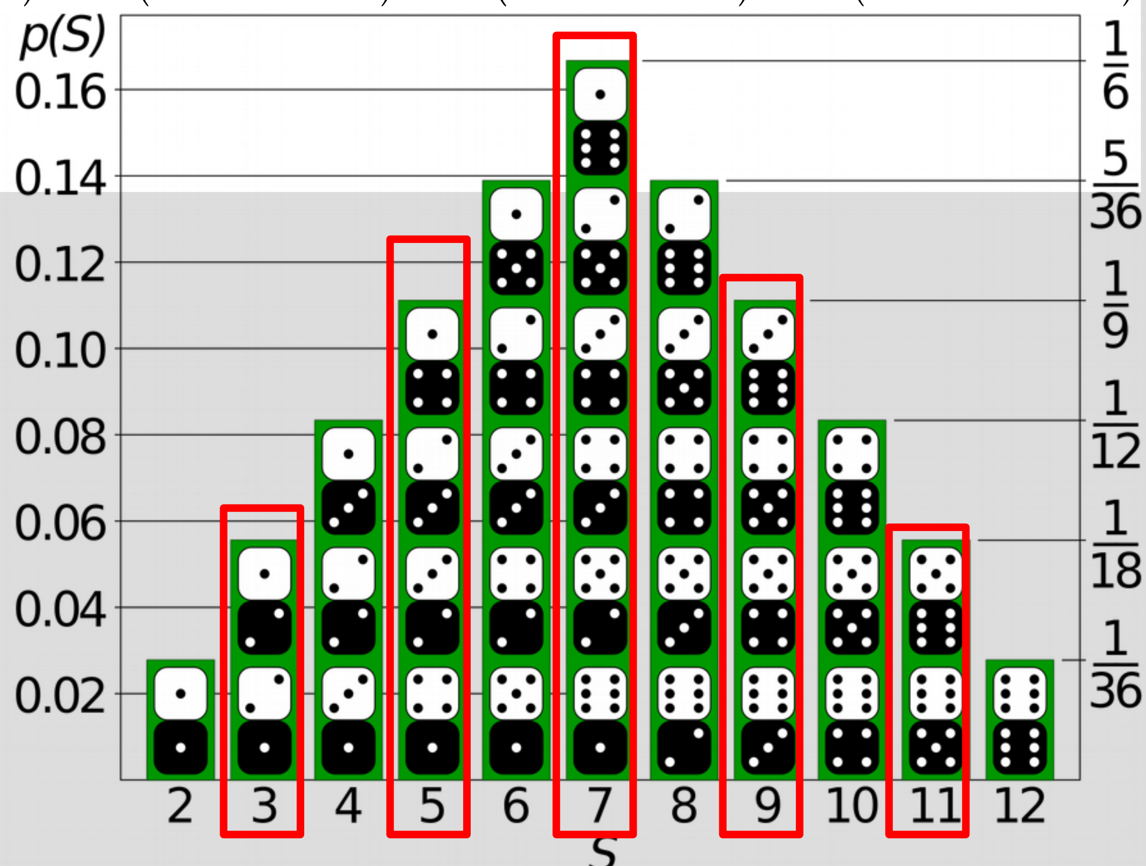

Probability: example

You might be able to brute force the 3-dice example but the 20-dice... probably not

$$P(A+B=odd) = P(A=odd, B=even) + P(A=even, B=odd)$$

We can break this down into to cases:

1. Original die is odd, then next must be even
2. Original die is even, then next must be odd

The "then" part of both are 50% chance, since regardless of which case we are in there is a 50% chance means overall probability=0.5

Probability: example

You can then use induction from this argument to generalize it:

Inductive step (by cases):
1. Sum of n dice is odd, “n+1” die is even
2. Sum of n dice is even, “n+1” die is odd
“n+1” die is just a single die, so 50% chance

Base case: we showed by brute force 50% for single die

Probability: example

You might try to prove this with independence (talk about next time), which you could

But you might notice that this proof actually says something stronger, as we never actually use the probability of the cases happening

So regardless of your original probabilities for odd/even, if you add a 6-sided die you will end up 50/50 split for odd/even

Random Variable: basics

<u>Random variables</u> are a set of value-probability pairs

You could think of our 6-sided die as a random variable with the following value-probabilities:

Prob.	1/6	1/6	1/6	1/6	1/6	1/6
Value	1	2	3	4	5	6

As I mentioned earlier, we often want to associate values/utilities with probabilities

Random Variable: basics

The expected value of a random variable is just the sum of the value*probability

So if a variable X is our die:

Prob.	1/6	1/6	1/6	1/6	1/6	1/6
Value	1	2	3	4	5	6

... then the expectation of X is:

$$E[X] = \sum_i X_i \cdot P(X_i) = 1\frac{1}{6} + 2\frac{1}{6} + 3\frac{1}{6} + 4\frac{1}{6} + 5\frac{1}{6} + 6\frac{1}{6} = 3.5$$

Random Variable: basics

This makes some sense, as the “average” value of a die is between 3 and 4 (1,2,3...4,5,6)

It is more interesting to look at more complex cases, like sum of 2 or 3 dice:

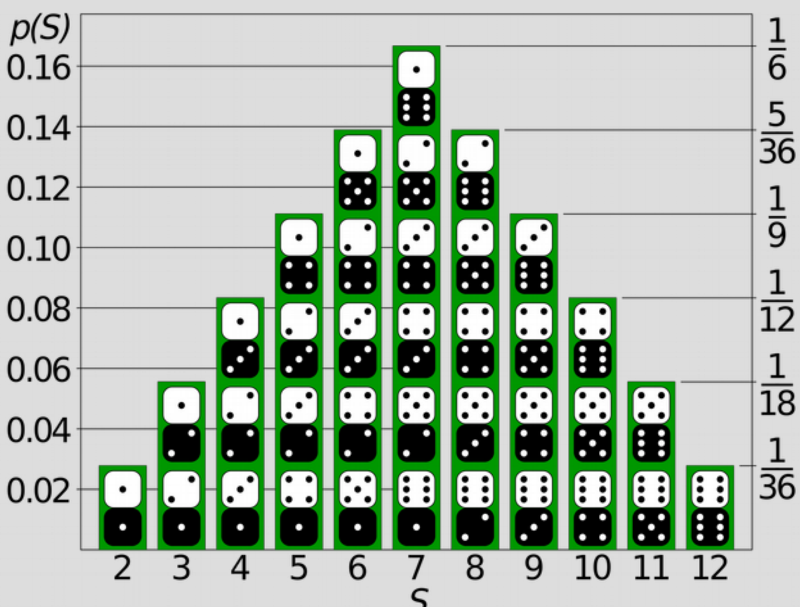

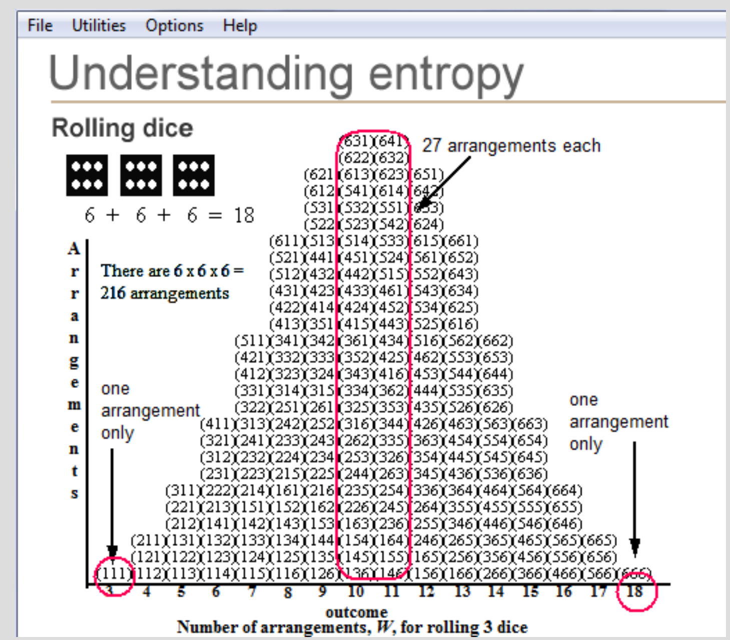

Random Variable: basics

$E[X+Y] = ??$

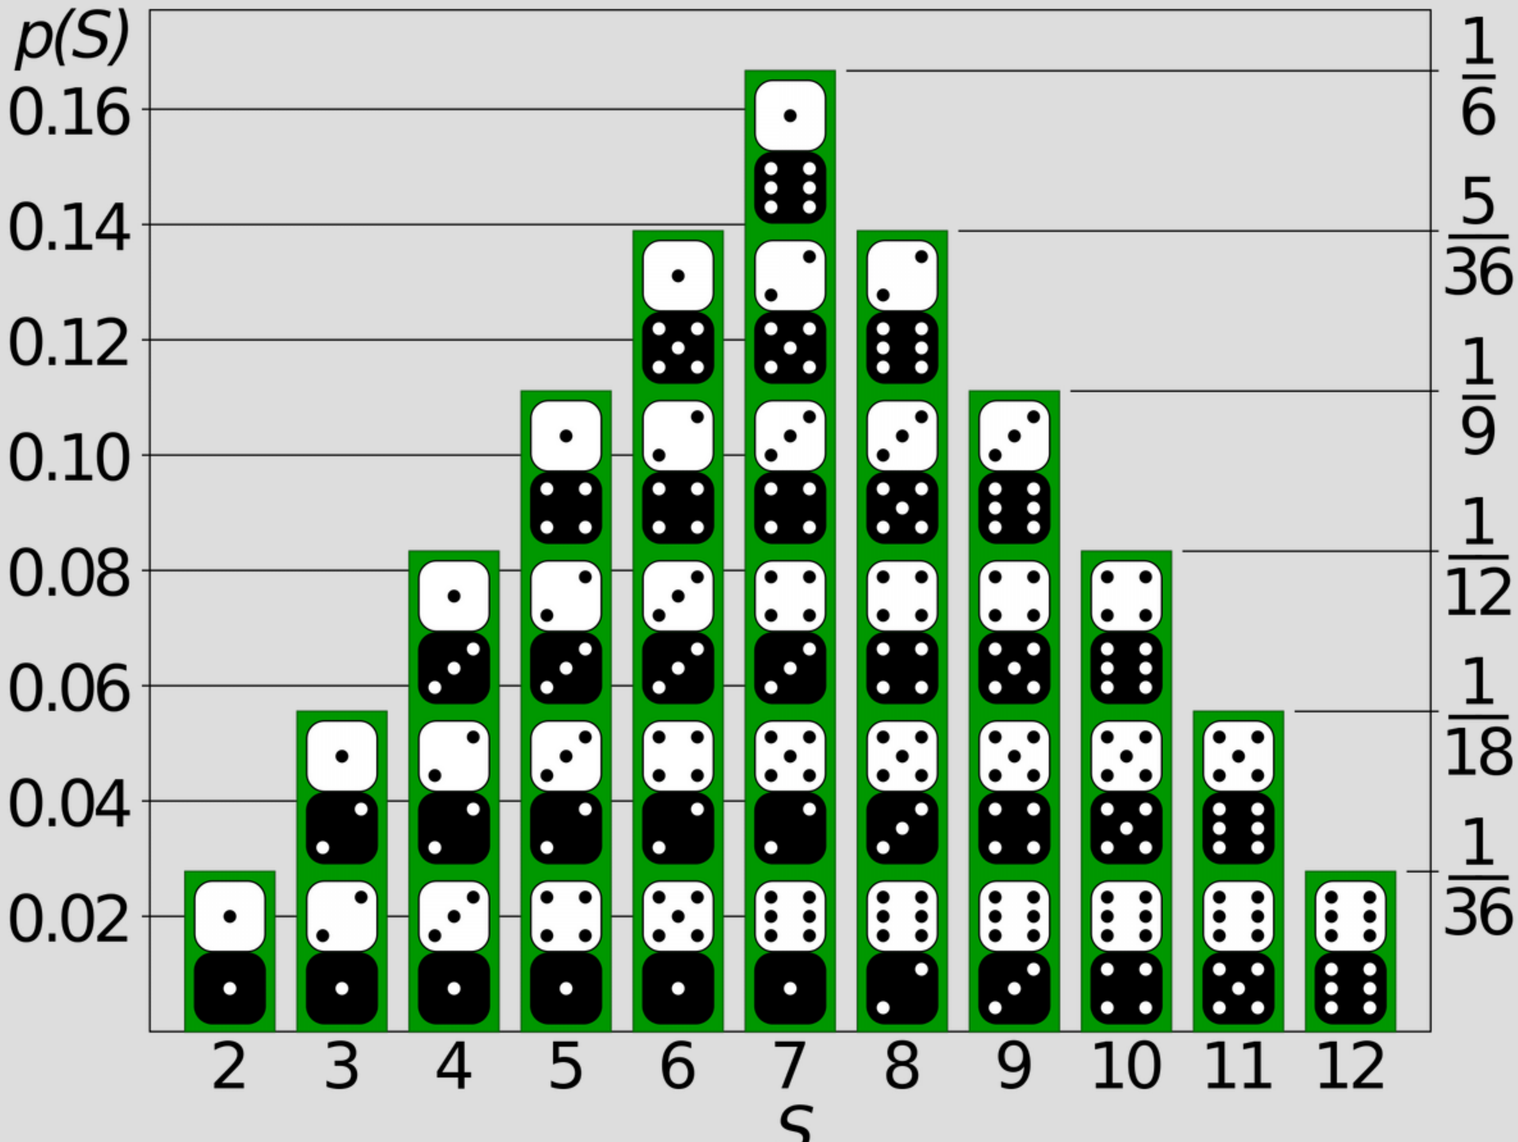

Random Variable: basics

$E[X+Y]=7$

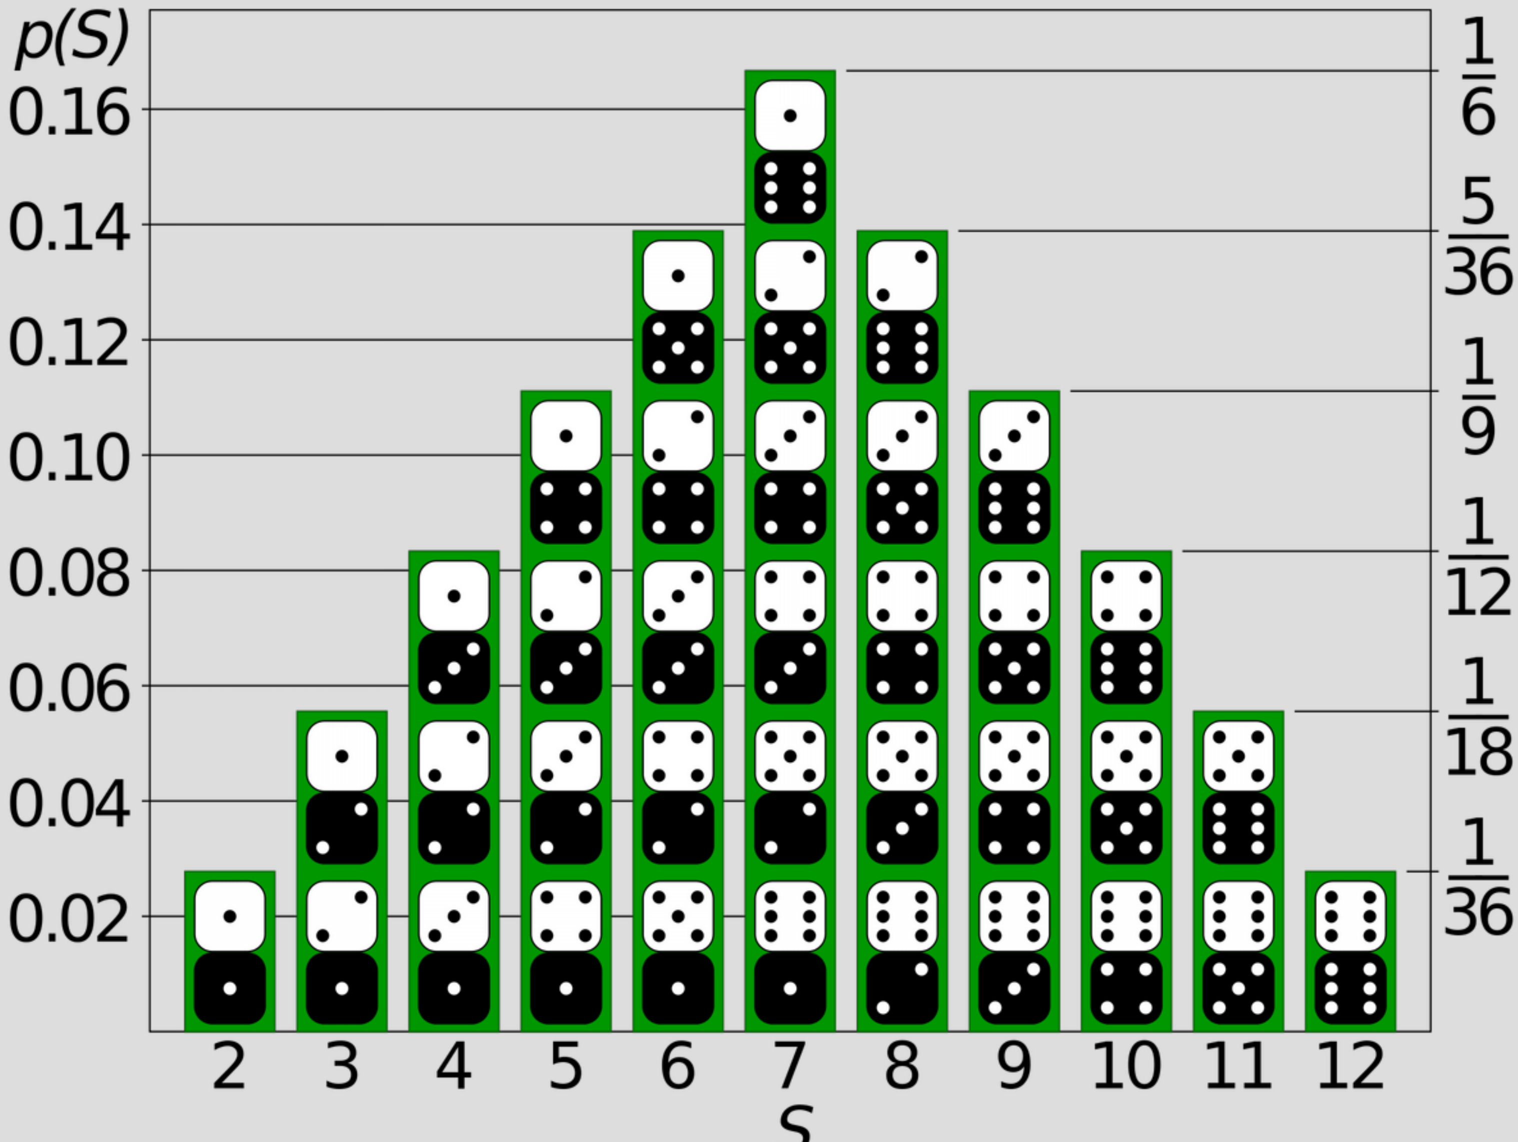

Random Variable: basics

$E[X + Y + Z] = ??$

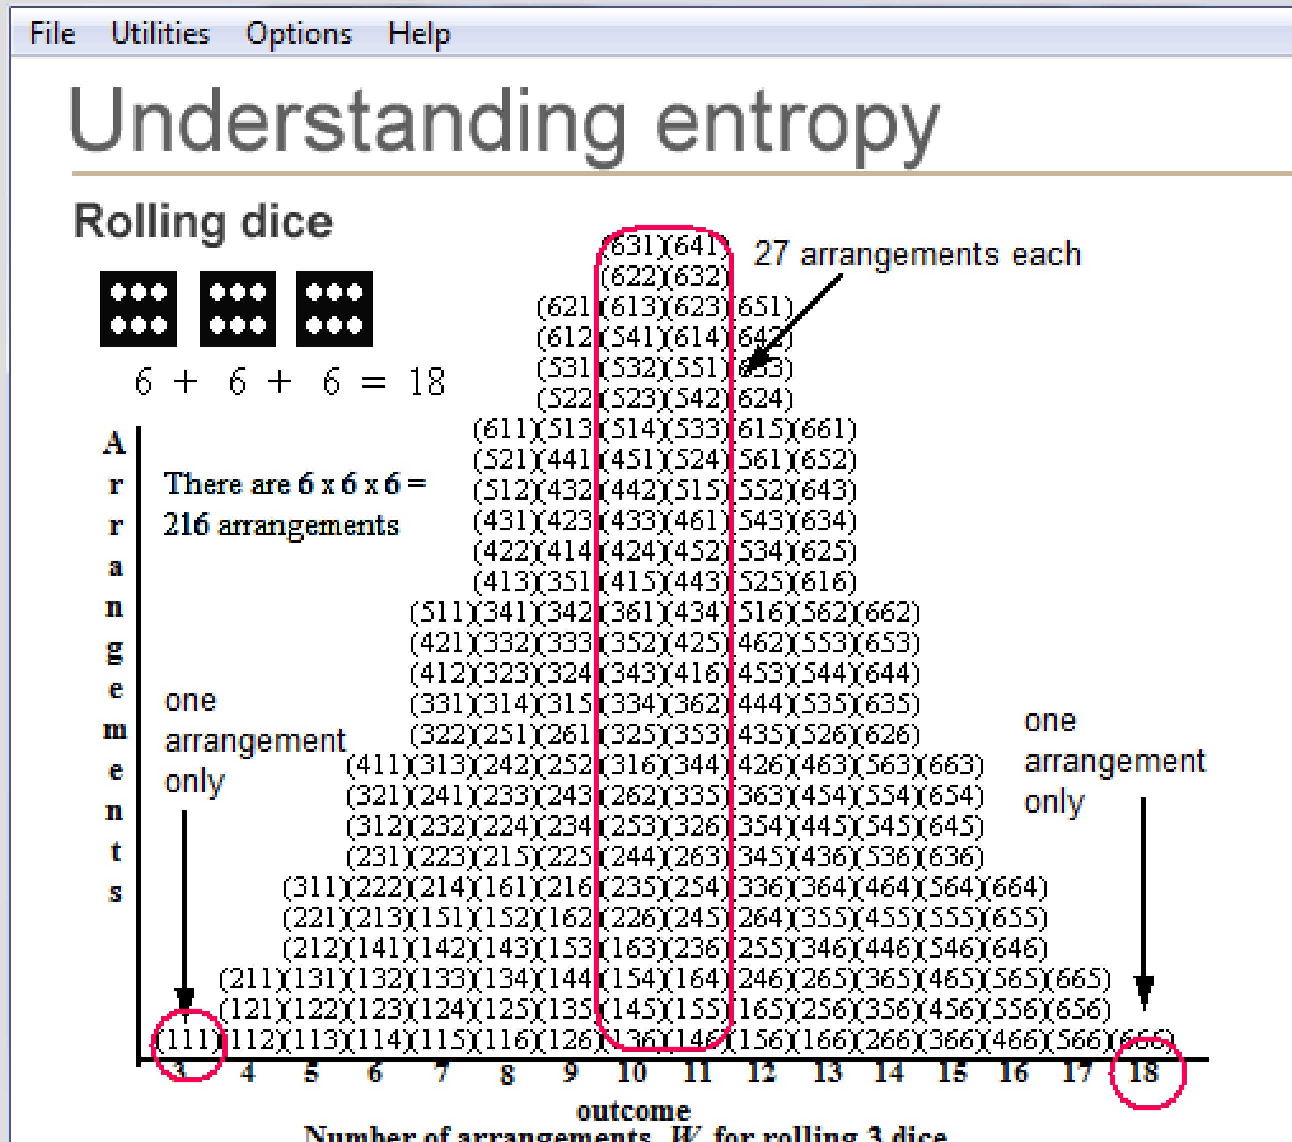

Random Variable: basics

$$E[X + Y + Z] = 10.5$$

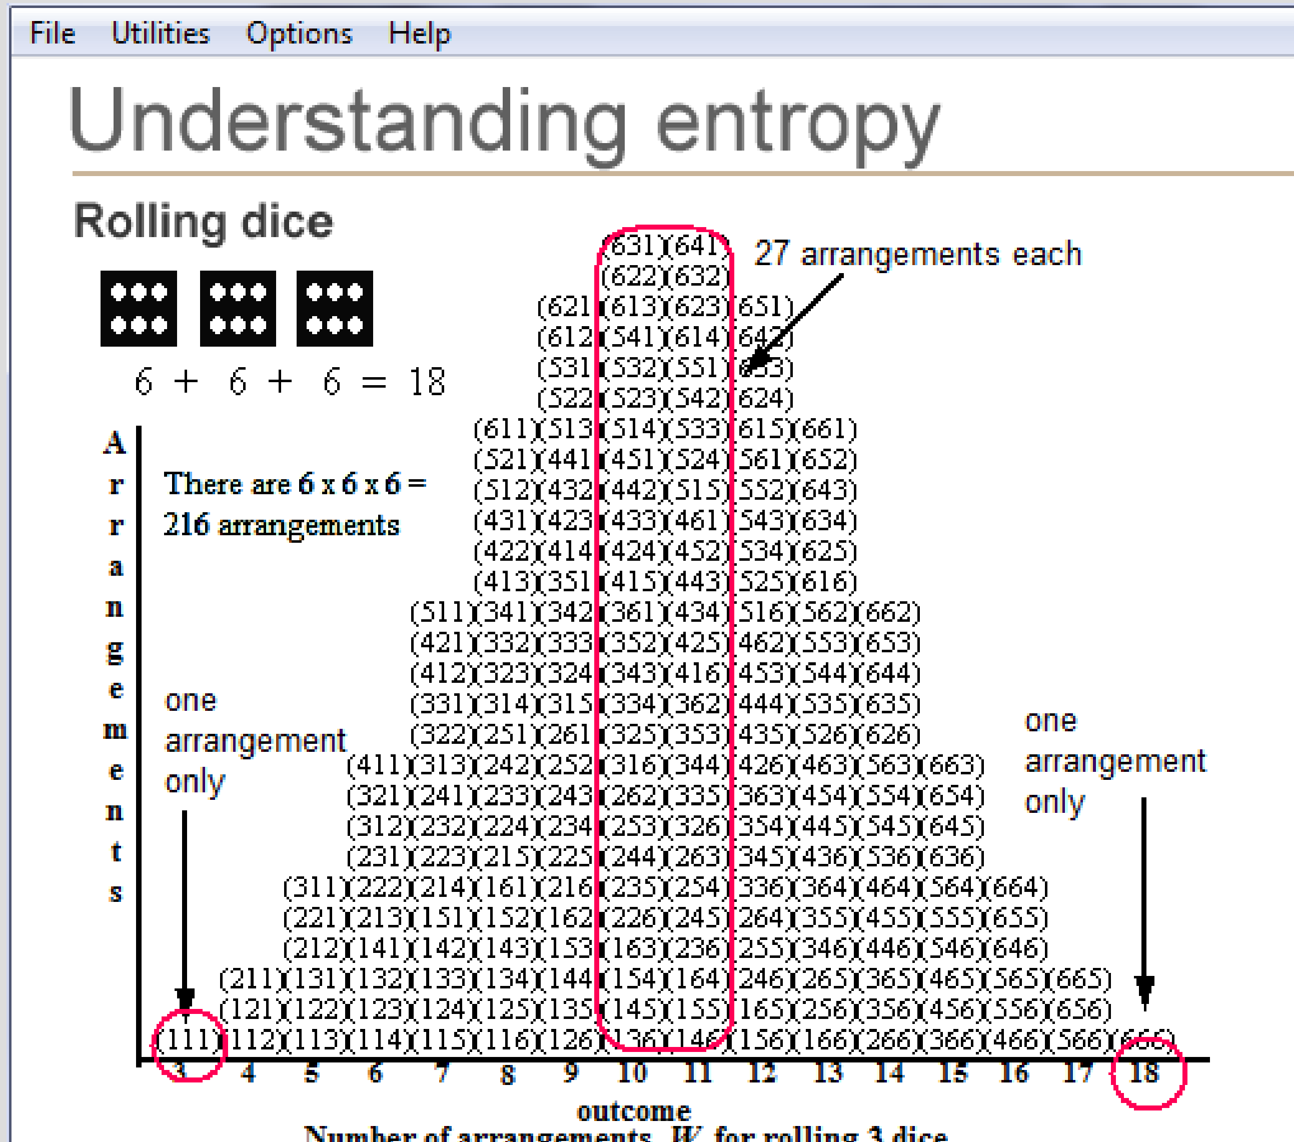

Random Variable: basics

Just like probabilities, random variables have their own set of properties

(also for scalar “a”)

$$E[aX] = a \cdot E[X]$$

One of which is:

$$E[X + Y] = E[X] + E[Y]$$

Since for a single die, E[X] = 3.5...

E[X+Y] = E[X] + E[Y] = 3.5 + 3.5 = 7

So 3 dice is 3*3.5 = 10.5

4 dice is 4*3.5 = 14

Continuous spaces

Dice are an easy example as they are discrete, but sometimes probabilities/random variables are not nice (continuous)

Consider:

$$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

(above is the probability density function)

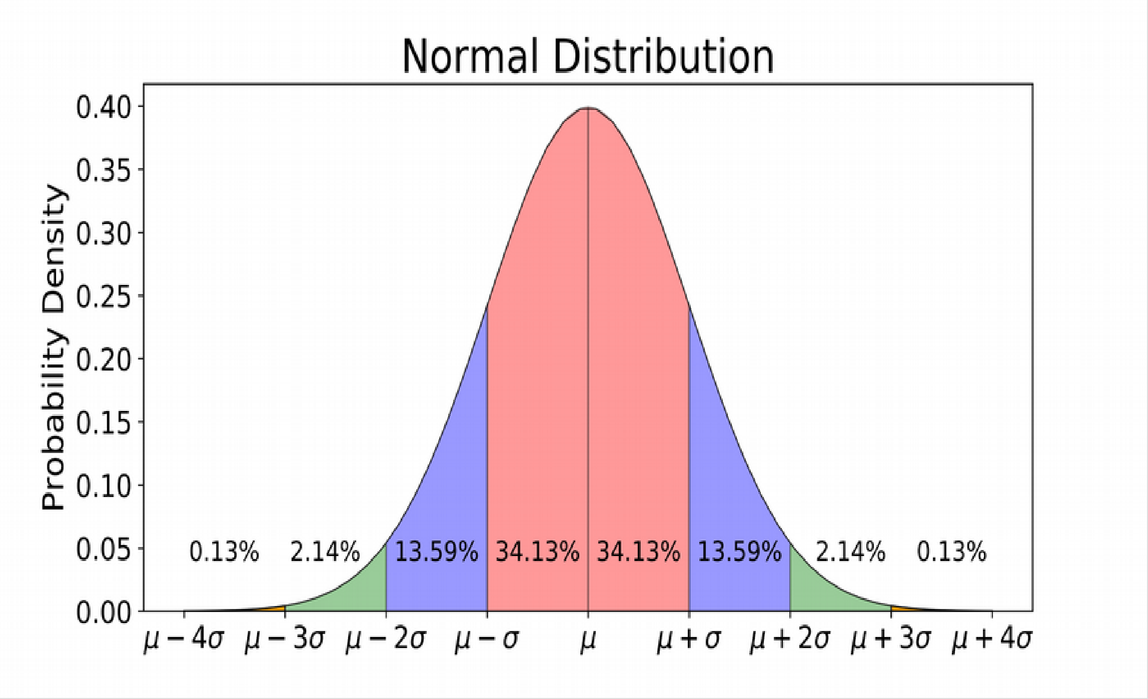

Continuous spaces

For continuous spaces, the probability that a specific value is taken is always zero:

$P(x = 3) = 0$

Instead, we have to work over a range:

$P(x \leq 0) = 0.5$

... which unfortunately requires integration:

$$P(x < 0) = P(x \leq 0) = \int_{-\infty}^{0} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx = 0.5$$

Continuous spaces

We will use the following distributions:

Uniform | Normal | Poisson

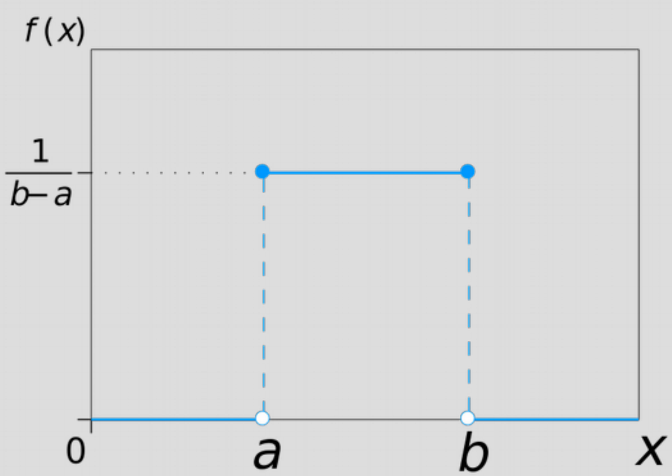

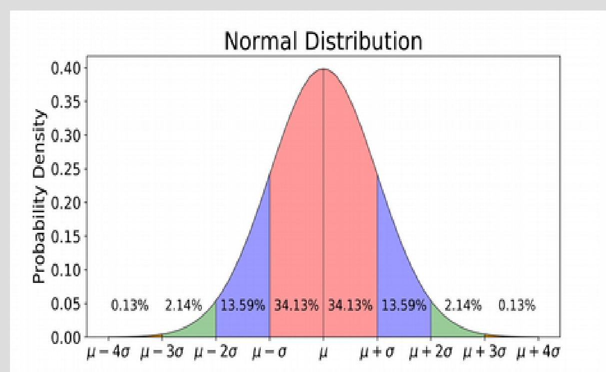

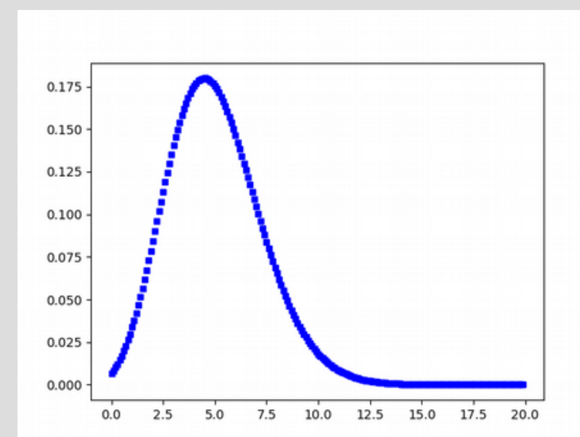

Probability distribution functions:

$$\frac{1}{b-a} \text{ if in [a,b]}$$
$$0 \text{ otherwise}$$

$$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

$$\frac{\lambda^k e^{-\lambda}}{k!}$$