Summary of Lecture 1

- An experiment can have a number of possible outcomes, but which outcome will occur is not known in advance
- Events are combinations of outcomes which behave as sets, and can be combined or modified using the operations not (^{c}), or ($\cup$) and and ($\cap$).
- Different outcomes may be more (or less) likely than others to be the result of the experiment.
- We quantify how likely each outcome is by defining a probability for each outcome (and by implication, each event)
- Probability must obey the Axioms of Probability (and their consequences), ie probability is defined to be a real number in $[0, 1]$ with larger values indicating more likely events.
- There are simple formulæ for the probabilities of the compound events A^c and $A \cup B$

Summary of Lecture 2

- When learning that event B has occurred changes the probability that event A will occur, then we call this *updated* probability the conditional probability of A given B — $\mathrm{P}[A|B]$
- If learning that event B has occurred does *not* change the probability that event A will occur, then we say A and B are independent
- In some problems, it is natural to divide the sample space into a collection of disjoint events - a partition - which can be used to simplify more complex probabilities via the Partition Theorem
- We can express the conditional probability $\mathrm{P}[A|B]$ in terms of $\mathrm{P}[B|A]$ via Bayes Theorem
- When the results of our experiment give numerical outcomes, we call the uncertain outcome a random variable

Summary of Lecture 3

- A discrete random variable X can only take a finite number of possible values, and we describe its probabilistic behaviour by a table of probabilities – its probability distribution
- A continuous random variable Y takes possible values from the real line, and we describe its probabilistic behaviour by a probability density function $f(x)$ over the possible values
- We must integrate the pdf to get probabilities, or we use the cumulative distribution function $F(x) = \int_{-\infty}^{x} f(x)\,dx$
- When we have more than one random variable, we describe the behaviour of the pair via a joint distribution (picture!)
- We can summarise distributions, one such summary is the expectation (or mean, or average)
- Expectation has a number of handy properties

Summary of Lecture 4

- Expectation summarises the location of a distribution, it is the “centre of probability”
- Variance summarises the spread of a distribution, and is a measure of how far on average we would expect values to be from the expectation
- We calculate expectations differently depending on whether the random variable is discrete or continuous
- There are two ways to express variance as an expectation
- Expectations and variances of linear combinations of random variables can be simplified

Summary of Lecture 5

- A common example of a discrete random variable occurs when we are interested in the counts of things
- Under certain circumstances, such variables may follow the Binomial or the Poisson distributions
- We have a Binomial situation when the problem can be thought of as equivalent to counting the number of successes in n independent trials, each with constant probability p of success
- The Poisson applies when we count the number of random events which occur in a time period s where the random events occur with a rate of λ per unit time and their times of occurrence are independent of each other.
- We can use the Poisson distribution to approximate Binomial probabilities when n is large and p is small

Summary of Lecture 6

- The Normal distribution is a useful continuous distribution which is parametrised by its mean μ and its variance σ^2
- The standard normal random variable $Z \sim N(0,1)$ is one example, and has a cumulative distribution function $\Phi(z)$
- The cumulative probabilities for the Z (ie the values of $\Phi(z)$) are given in standard normal tables for certain values of z
- Any $X \sim N(\mu, \sigma^2)$ can be transformed into an equivalent $Z \sim N(0,1)$ allowing us to use standard Normal probabilities to answer questions about X
- We can exploit properties of the Normal distribution (such as symmetry) to determine probabilities which are not tabulated
- A Binomial rv $X \sim Bin(n,p)$ with n large and $p \simeq 0.5$ can be well approximated by $X' \sim N(np, np(1-p))$.