Basic Mathematical Techniques for Computer Scientists

Introduction to Discrete Probability

January 21, 2013

Playing cards

Shuffled decks

- Recall the deck of 52 playing cards
- In how many ways can you order the cards in this deck?
- In a *shuffled* deck, *all* these orderings are *equally likely*
- Suppose you deal a shuffled deck of cards to four players
  - 13 cards each
  - Cards 1 to 13 to the first player, 14 to 26 to the second, ...
- What is the *probability* that each player gets a King?

How do we find this probability?

One way

- Each ordering of the cards is equally likely
  - By assumption
- The probability we want is the *proportion* of “good” orderings
- Define “success” as: *All four players get a King*
- A “good” or “favourable” ordering of the cards
  - An ordering which results in success
- The probability of success:

$$\frac{\text{Number of favourable orderings}}{\text{Total number of orderings}}$$

- This makes sense because each ordering is *equally likely*

How do we find this probability?

Another way

- Only the *positions* of the Kings matters
  - One King has to be present in each 13-block
  - The order among these Kings doesn't matter
- We can express the probability in terms of positions of the Kings
  - A favourable position has one King in each 13-block
- The probability of success:

$$\frac{\text{Number of favourable positions}}{\text{Total number of positions}}$$

- Again, this makes sense because each ordering is *equally likely*

Experiment, outcome, event

- Experiment: Some process which involves chances
  - E.g: Shuffling the deck and dealing 13 cards each to 4 players
- Outcome: The result of *one* “trial” of the experiment
  - E.g: The hands which the four players get after one deal
- Event: Some *set* of outcomes which we are interested in
  - E.g: “Each player gets a King”
  - Why does this match the definition of an event?

A slightly different experiment

- Deal a shuffled deck equally among four players
- Suppose you are a player, and you get exactly one King
- Given this extra information,
  - What is the probability that each player has a King?
- The experiment is different:
  - Shuffle and deal 13 cards each
  - *Such that* a specified player (you) gets exactly one King
- The set of outcomes is different
  - Those deals in which you get no King are not outcomes
  - A proper subset of the previous set of outcomes
- The event is the same
  - "Each player gets a King"

A slightly different experiment

- Deal a shuffled deck equally among four players
- Suppose you are a player, and you get exactly one King
- What is the probability that each player has a King?
  - Will this be less, equal, or greater than before?

Another variant

- Deal a shuffled deck equally among four players
  - In a *round-the-table* fashion
  - Player 1 gets card 1, player 2 card 2, ..., player 4 card 4,
  - Player **1** gets card 5, ...
- What is the probability that each player gets a King:
  1. In the absence of any extra information?
  2. Given that a specific player got exactly one King?

The six-sided die

A device commonly found in probability textbooks

- Plural: dice

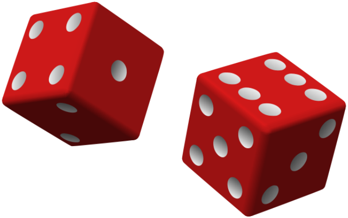

- These are “rolled” or “thrown”
- A “fair” die is one which is equally likely to turn up each of $1, 2, \ldots, 6$

The origin of modern probability theory

Letter from a bemused gambler

- Letter from gambler de Méré to mathematician Pascal (1654):
  - I used to bet 50-50 that I would get at least one 6 in four rolls of a fair die
  - The probability of this is four times the probability of getting a 6 in a roll, namely, $4 \times \frac{1}{6} = \frac{2}{3}$
  - Clearly this was in my favour, and indeed I was making money with this
  - I started betting 50-50 that I would get at least one $(6, 6)$ in twentyfour rolls of a pair of dice
  - This has the same advantage ($\frac{24}{6^2} = \frac{2}{3}$), but now I am losing money!
  - Why?!!

A domestic example

Distribution of kids

- In a certain country, a newborn child is equally likely to be a girl or a boy
  - This is *not* necesssarily true in the real world
- In a family with six children, what is the probability that exactly three are girls?
- The experiment:
  - Picking a family with six children and tallying their genders
- The set of outcomes:
  - All 6-length strings over the alphabet $\{G, B\}$
- The event we are interested in:
  - The subset of strings which have an equal number of Gs and Bs

Dice, again

- Consider rolling a fair die six times
- What is the probability that all the six numbers are different?

Summarizing these examples

- Suppose an experiment has n **equally likely** outcomes
  - Let S denote the set of all outcomes
- Let $E \subseteq S$ be an event
  - Let $|E| = m$
- The probability of event E occurring
  - Is the probability that *one* of the outcomes in E happens
  - Is equal to $\frac{m}{n}$
- *Intuitively*, the probability of an event E
  - Is the proportion of times that one *expects* the event to happen
    - This happens when *one* of the outcomes in the event happens
  - In a very large number of repetitions of the experiment
- We can put our counting skills to good use!

More examples

- We toss four fair coins
  - What is the probability that **exactly two** turn up heads?
- We roll two fair dice
  - What is the probability that the sum of the two numbers is 7?
- We shuffle a deck of cards. What is the probability that
  1. The first card in the deck is a King?
  2. All the spades appear consecutively somewhere in the deck?
  3. All cards of every suite appear together?

Some terminology

- When we say that we make a random choice
  - We mean we make a choice which is *uniformly* random
  - Each choice is equally likely

Some terminology

- A *population* is just a collection of some things
  - E.g: The collection of cards which make up a deck
- The process of *sampling* from a population
  - Choosing an object at random from the population
  - Checking what its properties are
    - E.g: Which colour/suite is a card?

Some terminology

- Sampling *with* replacement
  - Before picking each new sample, ...
  - ...we put back the previous object back into the population
- Sampling *without* replacement
  - We just pick an object when needed
  - And *do not* put it back after sampling

Sampling with and without replacement

Examples

- Seven pieces of paper in a bag
  - Written on them: **C, C, E, S, S, S, U**
- Sample seven pieces of paper from the bag, without replacement
  - Probability that the letters spell **SUCCESS**?
- Sample seven pieces of paper from the bag, *with* replacement
  - Probability that the letters spell **SUCCESS**?

Sampling with and without replacement

Examples

- A basket in the grocery store has 20 bell peppers
  - 10 red, and 10 orange
- You pick five peppers at random from the basket
- What is the probability that exactly two are red?
  1. If you pick without replacement?
  2. If you pick with replacement?

Sample spaces, events, and probability functions

For formalizing these notions of probability

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- What are the sample spaces of the following experiments?
  - Tossing one coin and looking at its result
  - Tossing two coins and looking at the results
  - Drawing two cards from a shuffled deck and checking what they are

Sample spaces, events, and probability functions

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- An **event** is a subset of the sample space
  - Denoted E
  - $E \subseteq \Omega$
- Experiment: tossing three coins
  - What is the sample space?
  - What is the event “at least two coins turn up heads”?

Sample spaces, events, and probability functions

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- An *event* is a subset of the sample space
  - Denoted E
  - $E \subseteq \Omega$
- Experiment: tossing three coins and counting the number of heads which occur
  - What is the sample space?
  - What is the event “at least two coins turn up heads”?

Sample spaces, events, and probability functions

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- An *event* is a subset of the sample space
  - Denoted E
  - $E \subseteq \Omega$
- Experiment: Drawing three cards from a shuffled deck and checking what they are
  - What is the sample space?
  - What is the event “the first card is an Ace”?

Sample spaces, events, and probability functions

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- An *event* is a subset of the sample space
  - Denoted E
  - $E \subseteq \Omega$
- An event E is said to *occur* as the result of an experiment
  - If the outcome of the experiment is an element of E

Sample spaces, events, and probability functions

- The *sample space* of an experiment is the set of *all* its possible outcomes
  - Denoted Ω
- An *event* is a subset of the sample space
  - Denoted E
  - $E \subseteq \Omega$
- A **probability function** assigns a real number to each event
  - Denoted Pr()
  - $\Pr : 2^{\Omega} \rightarrow \mathbb{R}$
- This number is usually the "physical" probability that we associate with each event
- The function must obey certain simple, intuitive laws

Axioms of Discrete Probability

For finite sample spaces

A function $\Pr : 2^{\Omega} \to \mathbb{R}$ defined on a sample space Ω is said to be a probability function if and only if it satisfies the following axioms:

1. **(Non-negativity)** For each $S \subseteq \Omega$, $\Pr(S) \geq 0$
2. **(Normalization)** $\Pr(\Omega) = 1$
3. **(Additivity)** If $S_1, S_2, \ldots, S_n$ are pairwise disjoint events, then

$$\Pr\left(\bigcup_{1}^{n} S_i\right) = \sum_{i=1}^{n} \Pr(S_i)$$

- Probabilities of events range from 0 to 1
  - Events with probability zero are said to be “impossible” events
  - Events with probability one are said to be “certain” events

Axioms of Discrete Probability

For finite sample spaces

A function $\Pr : 2^\Omega \to \mathbb{R}$ defined on a sample space Ω is said to be a probability function if and only if it satisfies the following axioms:

1. **(Non-negativity)** For each $S \subseteq \Omega$, $\Pr(S) \geq 0$
2. **(Normalization)** $\Pr(\Omega) = 1$
3. **(Additivity)** If $S_1, S_2, \ldots, S_n$ are pairwise disjoint events, then

$$\Pr(\bigcup_1^n S_i) = \sum_{i=1}^n \Pr(S_i)$$

- Such a pair $(\Omega, \Pr)$ is called a *probability space*

Probability spaces

Examples

- For rolling a fair die:
  - $\Omega = \{1, 2, 3, 4, 5, 6\}$
  - For $S \subseteq \Omega, \Pr(S) = \frac{|S|}{6}$

Probability spaces

Examples

- For randomly picking five bell peppers from a bag of ten red and ten orange bell peppers
  - $\Omega = \{R, O\}^5$
  - Alternatively, $\Omega =$
    - $\{0, 1, \ldots, 5\}$
    - $\{(r, o) \mid r, o \in \mathbb{N}, r + o = 5\}$
  - The choice of Ω depends ...
  - ... on the aspect of the experiment which we want to study
  - For $S \subseteq \Omega$, $\Pr(S) =$
    - $\frac{|S|}{2^5}$ for the first Ω
    - More involved in the other two cases

Thank You!