Dynamic Programming Lecture #4

Outline:

- Probability Review, cont
  - Bernoulli trials
  - Random variables
  - Expectation
  - Conditional expectation
  - Independence
  - Laws of large numbers

Bernoulli Trials

- Setup:
  - N experiements
  - Each experiment has 2 outcomes (H/T).
  - Perform multiple independent identical experiments.
  - "Identical" means identical probabilities, not outcomes.
- Assume $\mathrm{P}[H] = p$ for any toss (independent & identical experiments).
- $\Omega =$ binary sequences of length N, e.g.,

$$\omega = \underbrace{HHTHTT\ldots HTTH}_{N \text{ terms}}$$

- Q1: Probability of a specific $\omega \in \Omega$?
- Example: $N = 3$, $\omega = HHT$ is the event:
  - First toss: H, AND
  - Second toss: H, AND
  - Third toss: T
- Independence assumption: $\mathrm{P}[A \mathbin{\&} B \mathbin{\&} C] = \mathrm{P}[A]\,\mathrm{P}[B]\,\mathrm{P}[C]$
- Identical experiment assumption $\Rightarrow \mathrm{P}[HHT] = p \times p \times (1-p)$
- Ans1: For a specific ω : $p^{(\#\text{ heads})}(1-p)^{(\#\text{ tails})}$

Bernoulli Trials, cont

- Q2: Probability that N tosses results in k heads?
- How many ways can k heads occur? $\binom{N}{k} = \frac{N!}{(N-k)!k!}$
- Ans2: Each of these is a disjoint event, therefore

$$\mathrm{P}\left[k \text{ heads}\right] = \underbrace{p^k(1-p)^{(N-k)} + \ldots + p^k(1-p)^{(N-k)}}_{\text{sum of individual probabilities}}$$

$$\Rightarrow$$

$$\mathrm{P}\left[k \text{ heads}\right] = \binom{N}{k} p^k(1-p)^{(N-k)}$$

- Example: $p = 0.7$, $N = 20$:

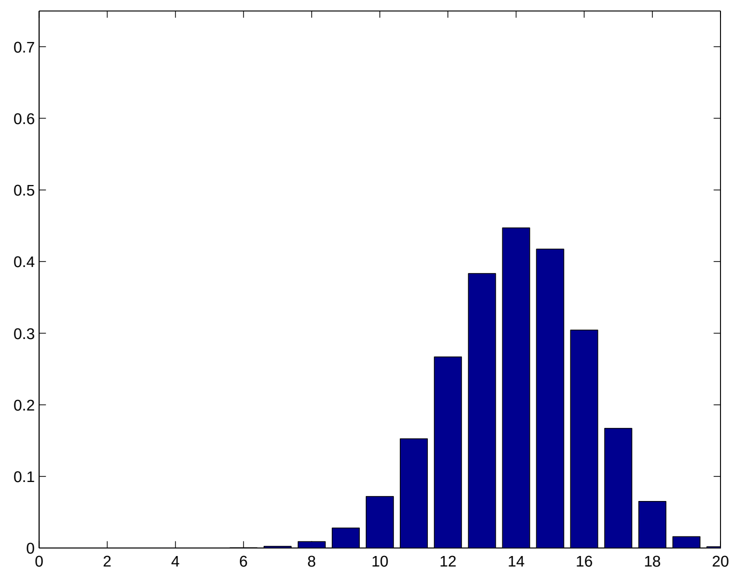

- Notice peak at $14 = p \times N$
- Important note: Nobody tossed a coin...derivation based on assumptions and probability axioms.

Discrete Random Variables

- Given:
  - Countable sample space Ω
  - Probability assignment:
  $$\mathrm{P}[A] = \sum_{\omega \in A} \mathrm{P}[\omega]$$
  i.e., probability of an event, A, is sum of probabilities of individual elements of A.
- DEFINE: A "random variable" (RV), X, is a FUNCTION mapping Ω to real numbers.
- Notation: $x = X(\omega)$ (must distinguish between the function and the value.)
- Examples:
  - $X =$ number of heads in a Bernoulli trial
  $$X(HHTH) = 3, \quad X(TTHT) = 1$$
  - $X =$ toss resulting in first heads
  $$X(HHTH) = 1, \quad X(TTHT) = 3$$
  - $X =$ sum of rolls of dice
  $$X(i,j) = i + j$$
  - $X =$ maximum roll of dice
  $$X(i,j) = \max(i,j)$$
- Same Ω yields different RV's.

Discrete Random Variables, cont

- Random variables provide a convenient and more expressive way to represent events.
- Example: Two dice, $X(i,j) = i + j$.
  - $X = 2$: The event $(1,1)$
  - $X \geq 11$: The event $(5,6) \cup (6,5) \cup (6,6)$
  - $X \neq 12$: The event $(i,j) \neq (6,6)$

$$\mathrm{P}\left[X \neq 12\right] = 1 - p_{66}$$

- Example: Bernoulli trials, N tosses, $X(\cdot)$ is number of heads.
  - $X = k$: The event k heads.
  - $X < k$: Fewer than k heads:

$$\mathrm{P}\left[X < k\right] = \sum_{j=0}^{k-1} \binom{N}{j} p^j (1-p)^{(N-j)}$$

- As before, translate condition on X to a subset of Ω.
- As before... "nothing happened".

Probability Mass Function (pmf)

- DEFINE: The probability mass function (pmf) of a discrete random variable as

$$p_X(x) = \mathrm{P}[X = x]$$

- How to compute pmf?
  - For each x, find all ω such that $X(\omega) = x$
  - Sum their probabilities

  (Can be clumsy.)

- Example: Two dice, $X(i, j) = i + j$

$$p_X(1) = 0$$
$$p_X(2) = p_{11}$$
$$p_X(3) = p_{12} + p_{21}$$
$$p_X(4) = p_{13} + p_{22} + p_{31}$$
$$\text{etc}$$

  where $p_{ij} = \mathrm{P}[(i, j)]$.

Expectation

- Define: The "expected value" of X

$$\mathrm{E}[X] = \sum_x x p_X(x)$$

- Expected value is a weighted sum:
  - All values $X(\omega)$ can take...
  - Weighted by the probability it takes that value
- Fact: An alternative formula (and usually more convenient) is

$$\mathrm{E}[X] = \sum_{\omega \in \Omega} X(\omega) \mathrm{P}[\omega]$$

- Example: Bernoulli trials, $N = 2$, $X =$ number of heads:

$$\mathrm{E}[X] = (\,2 \times p_{HH}\,) + (\,1 \times (p_{HT} + p_{TH})\,) + (\,0 \times p_{TT}\,)$$

$$\text{vs}$$

$$\mathrm{E}[X] = (\,2 \times p_{HH}\,) + (\,1 \times p_{HT}\,) + (\,1 \times p_{TH}\,) + (\,0 \times p_{TT}\,)$$

- Example: $\Omega = \{-2, -1, 0, 1, 2\}$, uniform probability:
  - $X(\omega) = \omega$:

$$\mathrm{E}[X] = -2/5 + -1/5 + 0/5 + 1/5 + 2/5 = 0$$

  - Let $Y = X^2$. This defines a a new RV:

$$\mathrm{E}[Y] = (-2)^2/5 + (-1)^2/5 + 0^2/5 + 1^2/5 + 2^2/5 = 5$$

Examples

- Example: Pair of dice:
  - RV$_1$: $X_1 = i + j$ = sum of rolls.

$$\begin{aligned} \mathrm{E}[X_1] &= \sum_i \sum_j x_1(i,j) p_{ij} \\ &= (1+1)p_{11} + (1+2)p_{12} + (1+3)p_{13} + \ldots + \\ &\quad (2+1)p_{21} + (2+2)p_{22} + \ldots \end{aligned}$$

  - RV$_2$: $X_2 = \max(i,j)$ = maximum of rolls.

$$\begin{aligned} \mathrm{E}[X_2] &= \sum_i \sum_j x_2(i,j) p_{ij} \\ &= \max(1,1)p_{11} + \max(1,2)p_{12} + \max(1,3)p_{13} + \ldots + \\ &\quad \max(2,1)p_{21} + \max(2,2)p_{22} + \ldots \end{aligned}$$

- Example: Bernoulli trials, $N = \infty$ (what is Ω?), X = number of tosses until first head:

$$\mathrm{E}[X] = 1p + 2(1-p)p + 3(1-p)^2 p + 4(1-p)^3 p + \ldots = \sum_{k=1}^{\infty} k(1-p)^{(k-1)} p$$

- Example: Heavy tails
  - X takes on values $1, 2, 3, \ldots$
  - $p_X(k) = \frac{c}{k^2}$

$$\mathrm{E}[X] = \sum k \frac{c}{k^2} = \sum \frac{c}{k} = \infty??!!!$$

  - Not all RV's have expected values!
  - Heavy tailed pmf's used to model rare (and catastrophic) events.

Conditional Expectation

- Define conditional pmf:

$$p_{X|A}(x) = \mathrm{P}\left[X = x|A\right]$$

- Define conditional expectation:

$$\mathrm{E}\left[X|A\right] = \sum_x x p_{X|A}(x)$$

- Example: Fair dice.

$i \backslash j$	1	2	3	4	5	6
1	·	·	·	·	·	·
2	·	·	·	·	·	·
3	·	·	·	·	·	·
4	·	·	·	·	·	·
5	·	·	·	·	·	·
6	·	·	·	·	·	·

$$X = i + j$$

$$\begin{aligned}
&\mathrm{E}\left[X\right] = \sum_x x\mathrm{P}\left[X = x\right] \\
&\quad = (2)(1/36) + (3)(2/36) + (4)(3/36) + ... + (11)(2/36) + (12)(1/36) = 252/36 = 7 \\
&\mathrm{E}\left[X|X \le 3\right] \\
&\quad = (2)\mathrm{P}\left[X = 2|X \le 3\right] + (3)\mathrm{P}\left[X = 3|X \le 3\right] + 0 + ... + 0 \\
&\quad = (2)\left(\tfrac{1/36}{3/36}\right) + (3)\left(\tfrac{2/36}{3/36}\right) = (2)(1/3) + (3)(2/3) = 3\left(\tfrac{1}{3}\right) \\
&\mathrm{E}\left[X|\mathrm{even}\right] \\
&\quad = (2)\mathrm{P}\left[X = 2|\mathrm{even}\right] + (3)\mathrm{P}\left[X = 3|\mathrm{even}\right] + ... \\
&\qquad + (11)\mathrm{P}\left[X = 11|\mathrm{even}\right] + 12\mathrm{P}\left[X = 12|\mathrm{even}\right] = 7
\end{aligned}$$

Conditional Expectation & Independence

- Can condition on another random variable, e.g.,

$$\mathrm{E}\left[X_1|X_2=x_2\right]=\sum_i x_1(\omega_i)\mathrm{P}\left[\omega_i|X_2=x_2\right]$$

- Dice Example: $\mathrm{E}\left[X_1|X_2=2\right]=$?

$$\begin{aligned}&(1+1)P(i=1,j=1|\max(i,j)=2)\\&\quad+(1+2)P(i=1,j=2|\max(i,j)=2)+(2+1)P(i=2,j=1|\max(i,j)=2)+(2+2)P(i=2,j=2|\max(i,j)=2)\end{aligned}$$

$$=$$

$$\begin{aligned}&(1+1)\frac{P(i=1,j=1,\max(i,j)=2)}{P(\max(i,j)=2)}\\&\quad+(1+2)\frac{P(i=1,j=2,\max(i,j)=2)}{P(\max(i,j)=2)}+(2+1)\frac{P(i=2,j=1,\max(i,j)=2)}{P(\max(i,j)=2)}+(2+2)\frac{P(i=2,j=2,\max(i,j)=2)}{P(\max(i,j)=2)}\end{aligned}$$

$$=$$

$$(1+1)\frac{0}{p_{12}+p_{21}+p_{22}}+(1+2)\frac{p_{12}}{p_{12}+p_{21}+p_{22}}+(2+1)\frac{p_{21}}{p_{12}+p_{21}+p_{22}}+(2+2)\frac{p_{22}}{p_{12}+p_{21}+p_{22}}$$

- Independence: RV's X_1 and X_2 are independent if

$$\mathrm{P}\left[X_1=x_1\&X_2=x_2\right]=p_{X_1}(x_1)p_{X_2}(x_2)$$

- Similar definition for multiple RV's

What's the connection?

- Informal probability setup:

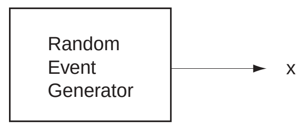

- Suppose random event is coin toss: H=0 vs T=1
- Define "sample mean"

$$M_N = \frac{\#\text{ tails}}{N}$$

  NOTE: M_N is also random...don't expect same result for every set of samples.

- Informal probability DEFINITION:

$$\lim_{N\to\infty} M_N \stackrel{\text{def}}{=} P(\mathsf{T})$$

- *Is this consistent with formal probability framework?*
  - Define independent RV's:

$$x_i = \begin{cases} 1 & \text{w/ prob } p \text{ (tails);} \\ 0 & \text{w/ prob } 1-p \text{ (heads)} \end{cases}$$

  - Define "sample mean"

$$M_N = \frac{x_1 + \ldots + x_N}{N}$$

  - M_N is a RV over

$$\Omega_N = \Omega \times \Omega \times \ldots \times \Omega$$

  - Example $N = 4$: Ω_N is set of H/T sequences of length 4.

$$P(H, H, T, H) = (1-p)^3 p$$

    as a CONSEQUENCE of independence.

What's the connection? cont.

$$M_N = \frac{x_1 + \ldots + x_N}{N}$$

- FACT: $\mathrm{E}[M_N] = \mathrm{E}[x] = p$.
- WEAK LAW OF LARGE NUMBERS:

$$P(|M_N - p| > \varepsilon) \to 0 \text{ as } N \to \infty$$

- Expected value corresponds to "taking averages"...but is NOT defined as average.
- Still allows M_N to deviate from p infinitely often with diminishing probability.
- STRONG LAW OF LARGE NUMBERS:

$$P(\lim_{N\to\infty} M_N = p) = 1$$

- This is CONSEQUENCE of formulation...not definition of probability.
- Probability space in strong law is space of sequences.
- Conforms to informal frequency intuition!