

Lecture 2: January 9, 2025

Lecturer: Madhur Tulsiani

1 Joint Entropy

We have two random variables X and Y . The joint distribution of the two random variables (X, Y) takes values (x, y) with probability $p(x, y)$. Merely by using the definition, we can write down the entropy of $Z = (X, Y)$ trivially. However what we are more interested in is seeing how the entropy of (X, Y) , the joint entropy, relates to the individual entropies, which we work out below:

$$\begin{aligned}
 H(X, Y) &= \sum_{x,y} p(x, y) \log \frac{1}{p(x, y)} \\
 &= \sum_{x,y} p(x)p(y|x) \log \frac{1}{p(x)} + \sum_{x,y} p(x)p(y|x) \log \frac{1}{p(y|x)} \\
 &= \sum_x p(x) \log \frac{1}{p(x)} \sum_y p(y|x) + \sum_{x,y} p(x)p(y|x) \log \frac{1}{p(y|x)} \\
 &= H(X) + \sum_x p(x) H(Y|X = x) \\
 &= H(X) + \mathbb{E}_x [H(Y|X = x)]
 \end{aligned}$$

Denoting $\mathbb{E}_x [H(Y|X = x)]$ as $H(Y|X)$, this can simply be written as

$$H(X, Y) = H(X) + H(Y|X)$$

If we were to redo the calculations, we could similarly obtain:

$$H(X, Y) = H(Y) + H(X|Y)$$

This is called the *Chain Rule* for Entropy. Note that in the calculations above, we treat $(Y|X = x)$ as a random variable, with distribution given by $\mathbb{P}[Y = y | X = x] = p(y|x)$. Also note that $H(Y|X)$ is simply a shorthand for the *expected* entropy of $(Y|X = x)$, with the expectation taken over the values for X .

Example 1.1. Consider the random variable (X, Y) with $X \vee Y = 1$ and $X \in \{0, 1\}$ and $Y \in \{0, 1\}$ such that:

$$(X, Y) = \begin{cases} 01 & \text{with probability } 1/3 \\ 10 & \text{with probability } 1/3 \\ 11 & \text{with probability } 1/3 \end{cases}$$

Now, let us calculate the following:

1. $H(X) = H(Y) = \frac{1}{3} \log 3 + \frac{2}{3} \log \frac{3}{2}$
2. $H(Y|X = 0) = 0$
3. $H(Y|X = 1) = \frac{1}{2} \log \frac{1}{2} + \frac{1}{2} \log \frac{1}{2} = 1$
4. $H(Y|X) = \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3}$
5. $H(X, Y) = \frac{1}{3} \log 3 + \frac{1}{3} \log 3 + \frac{1}{3} \log 3 = \log 3$

From the above we see that:

$$H(Y) \geq H(Y|X)$$

this is actually *always* true and we prove this fact below.

Proposition 1.2. $H(Y) \geq H(Y|X)$

Proof: We want to show that $H(Y|X) - H(Y) \leq 0$. Consider the quantity on the left hand side.

$$\begin{aligned} H(Y|X) - H(Y) &= \sum_x p(x) \sum_y p(y|x) \log \frac{1}{p(y|x)} - \sum_y p(y) \log \frac{1}{p(y)} \\ &= \sum_x p(x) \sum_y p(y|x) \log \frac{1}{p(y|x)} - \sum_y p(y) \log \frac{1}{p(y)} \sum_x p(x|y) \\ &= \sum_{x,y} p(x,y) \left(\log \frac{1}{p(y|x)} - \log \frac{1}{p(y)} \right) \\ &= \sum_{x,y} p(x,y) \left(\log \frac{p(x)p(y)}{p(x,y)} \right) \end{aligned}$$

Now consider a random variable Z that takes value $\frac{p(x)p(y)}{p(x,y)}$ with probability $p(x,y)$. Then we can use Jensen's inequality to get:

$$\sum_{x,y} p(x,y) \left(\log \frac{p(x)p(y)}{p(x,y)} \right) \leq \log \left(\sum_{x,y} \frac{p(x)p(y)}{p(x,y)} p(x,y) \right) = \log(1) = 0.$$

■

Note however the fact that conditioning on X reduces the entropy of Y is only true *on average over all fixings of X* . In particular, in the above example we have $H(Y|X = 1) = 1 > H(Y)$. But $H(Y|X)$, which is an average over all fixings of X , is indeed smaller than $H(Y)$. Also, check that above inequality is tight only when X and Y are independent.

Exercise 1.3. Show that $H(Y) = H(Y|X)$ if and only if X and Y are independent.

Using induction, we can use the chain rule to show that the following also holds for a tuple of random variables (X_1, \dots, X_m) .

$$H(X_1, X_2, \dots, X_m) = H(X_1) + H(X_2|X_1) + H(X_3|X_1, X_2) \dots H(X_m|X_1, \dots, X_{m-1}).$$

Combining this with the fact that conditioning (on average) reduces the entropy, we get the following inequality which is referred to the sub-additivity property of entropy.

$$H(X_1, X_2, \dots, X_m) \leq H(X_1) + H(X_2) + H(X_3) + \dots + H(X_m).$$

2 Source Coding Theorem

We begin by recalling the Shannon Code. We considered a random variable X that took on values a_1, a_2, \dots, a_n with probabilities p_1, p_2, \dots, p_n . We wanted to encode the values of X such that the expected number of bits needed is small. If $\ell_1, \ell_2, \dots, \ell_n$ are the number of bits needed to encode a_1, a_2, \dots, a_n , then we saw that a prefix free code exists iff:

$$\sum_{i=1}^n 2^{-\ell_i} \leq 1$$

Furthermore, we saw that the expected length of the encoding is lower bounded by $H(X)$ and upper bounded by $H(X) + 1$ (a code as specified as above, the Shannon code may be constructed by setting $\ell_i = \lceil \log(1/p_i) \rceil$).

We will now try to improve this upper bound and we will do so by considering multiple copies of X . The idea is that by amortizing the loss over many symbols, we can start to approach an expected length equal to the lower bound i.e. the entropy.

The design may be done as follows: Consider m copies of the random variable X , $\{X_1, \dots, X_m \in U\}$ and a code $C : \mathcal{X}^m \rightarrow \{0, 1\}^*$. Let $|\mathcal{X}|^m = N$. Now, we know that:

$$H(X_1, \dots, X_m) \leq \sum_{i=1}^N p_i \left[\log \frac{1}{p_i} \right] \leq H(X_1, \dots, X_m) + 1$$

Let us also assume that the m copies of X are drawn i.i.d. Using this assumption we try to work out the quantity $H(X_1, \dots, X_m)$. Which may be expanded using the chain rule and independence:

$$\begin{aligned} H(X_1, \dots, X_m) &= H(X_1) + H(X_2|X_1) + \dots + H(X_m|X_1, \dots, X_{m-1}) \\ &= H(X_1) + H(X_2) + \dots + H(X_m) \\ &= m \cdot H(X) \end{aligned}$$

Therefore, we get

$$\mathbb{E} [|C(X_1, \dots, X_m)|] \leq m \cdot H(X) + 1.$$

Thus, we used $H(X) + \frac{1}{m}$ bits on average per copy of X . This leads us to the source coding theorem.

Theorem 2.1 (Fundamental Source Coding Theorem (Shannon)). *For all $\epsilon > 0$ there exists a n_0 such that for all $n \geq n_0$ and given n copies of X , X_1, \dots, X_n sampled i.i.d., it is possible to communicate (X_1, \dots, x_n) using at most $H(X) + \epsilon$ bits per copy on average.*

3 Bounding binomial sums using binary entropy

We use the subadditivity property to obtain an upper bound on the number of subsets of $[n] = \{1, \dots, m\}$ of size at most k i.e., we need to bound size of the following set

$$S = \{(x_1, \dots, x_n) \in \{0, 1\}^n \mid x_1 + \dots + x_n \leq k\}.$$

Of course we can write the following expression for the size of S

$$|S| = \binom{n}{0} + \binom{n}{1} + \dots + \binom{n}{k},$$

but how much is the value of the above sum? We will estimate it in terms of the binary entropy function defined as

$$H_2(p) := p \cdot \log\left(\frac{1}{p}\right) + (1-p) \cdot \log\left(\frac{1}{1-p}\right)$$

Note that $h(p)$ is the entropy of a random variable X , which takes value 1 with probability p and 0 with value $1-p$ (or vice-versa). This immediately tells us that the maximum possible value of $H_2(p)$ is 1, which is achieved at $p = 1/2$. The function $H_2(p)$ can also easily be shown to be concave. It is also increasing for $p \in (0, 1/2)$ and decreasing for $p \in (1/2, 1)$

Claim 3.1. *The function $H_2(p)$ is concave.*

Proof: We would like to prove that for any $p_1, p_2, \alpha \in [0, 1]$, we have

$$H_2(\alpha \cdot p_1 + (1-\alpha) \cdot p_2) \geq \alpha \cdot H_2(p_1) + (1-\alpha) \cdot H_2(p_2).$$

Consider a pair of binary random variables (X, Y) distributed as follows. X is 0/1 with probability $1-\alpha$ and α respectively. Conditioned on $X = 1$, we take Y to be 0/1 with probability $1-p_1$ and p_1 , and conditioned on $X = 0$, we take Y to be 0/1 with probability $1-p_2$ and p_2 . Now, we have $H(Y|X)$ as

$$H(Y|X) = \mathbb{P}[X = 1] \cdot H(Y|X = 1) + \mathbb{P}[X = 0] \cdot H(Y|X = 0) = \alpha \cdot H_2(p_1) + (1-\alpha) \cdot H_2(p_2).$$

Also, note that Y is a binary random variable with $\mathbb{P}[Y = 1] = \alpha \cdot p_1 + (1 - \alpha) \cdot p_2$. So we can write

$$H(Y) = H_2(\alpha \cdot p_1 + (1 - \alpha) \cdot p_2).$$

Using $H(Y) \geq H(Y|X)$ now proves the claim. ■

Exercise 3.2. *Prove that the function $H_2(p)$ is increasing when $p \in (0, 1/2)$. You can also prove this using $H(Y|X) \leq H(Y)$ if you like.*

We now return to the estimation problem. Let (X_1, \dots, X_n) be a uniformly distributed over S . Thus, we have that $H(X_1, \dots, X_n) = \log |S|$. We can also use sub-additivity to say that

$$\log |S| = H(X_1, \dots, X_n) \leq H(X_1) + \dots + H(X_n) = n \cdot H(X_1),$$

where the last equality used the symmetry of the variables X_1, \dots, X_n . Now since X_1 was an indicator variable, let us say that it takes value 1 with probability p and value 0 with probability $1 - p$. Then $H(X_1) = H_2(p)$. Also, we have that $X_1 + \dots + X_n \leq k$, which gives by symmetry that $p = \mathbb{E}[x_1] \leq k/n$. Finally, we note that since the function $H_2(p)$ is increasing for $p \leq 1/2$, we get $H(X_1) = H_2(p) \leq H_2(k/n)$. This gives the bound

$$\log |S| \leq n \cdot H_2(k/n) \quad \Rightarrow \quad |S| \leq 2^{n \cdot H_2(k/n)}.$$

You can check that the upper bound obtained here is not too bad since the sum is approximately equal to $\frac{2^{n \cdot H_2(k/n)}}{\sqrt{2\pi \cdot k \cdot (1 - k/n)}}$.


4 Cauchy-Schwarz and some generalizations

We now give an information-theoretic proof of the Cauchy-Schwarz inequality. While this inequality can of course be proved in many other ways, the method described here was used by Friedgut [Fri04] to prove several interesting generalizations. I highly recommend taking a look at his paper!

Recall that the (finite version of) Cauchy-Schwarz inequality states that for real numbers a_1, \dots, a_n and b_1, \dots, b_n , we have that

$$\left(\sum_{i=1}^n a_i \cdot b_i \right) \leq \left(\sum_{i=1}^n a_i^2 \right) \cdot \left(\sum_{i=1}^n b_i^2 \right).$$

Note that we can assume numbers are non-negative, since the LHS can only decrease otherwise, while the RHS remains unchanged. Also, by continuity of the expressions on both sides (in the numbers a_i, b_i) it suffices to prove the inequality for rational numbers. Finally, since we can scale both sides by the same number, it suffices to only consider natural numbers a_1, \dots, a_n and b_1, \dots, b_n .



Consider disjoint subsets $A_1, \dots, A_n \subseteq \mathbb{N}$, where $|A_i| = a_i$, and similarly, disjoint subsets B_1, \dots, B_n with $|B_i| = b_i$. Let r_i denote the rectangle $A_i \times B_i$ with $a_i \cdot b_i$ points. We think of the rectangles being in the x, y plane with x coordinates in the sets A_i and y coordinates in the sets B_i . We pick two points random (X_1, Y_1) and (X_2, Y_2) as follows:

- Pick a rectangle R with probability proportional to its area i.e.,

$$\mathbb{P}[R = r_i] = \frac{a_i \cdot b_i}{\sum_j a_j \cdot b_j}$$

- Pick two points (X_1, Y_1) and (X_2, Y_2) independently from R .

Note that for the random point (X_1, Y_1) , the distribution is uniform over the set of all points (and similarly for (X_2, Y_2)). Also, since the sets are disjoint, specifying any of the variables X_1, Y_1, X_2, Y_2 reveals R , which means that

$$H(X_1, Y_1, R) = H(X_2, Y_2, R) = H(X_1, Y_1) = H(X_2, Y_2) = \log \left(\sum_i a_i \cdot b_i \right)$$

Finally, we use the fact that *given* the choice of the rectangle R , all four random variables

X_1, Y_1, X_2, Y_2 are independent. We get

$$\begin{aligned}
H(X_1, Y_1, R) + H(X_2, Y_2, R) &= 2H(R) + H(X_1, Y_1|R) + H(X_2, Y_2|R) \\
&= 2H(R) + H(X_1|R) + H(Y_1|R) + H(X_2|R) + H(Y_2|R) \\
&= 2H(R) + H(X_1, X_2|R) + H(Y_1, Y_2|R) \\
&= H(X_1, X_2, R) + H(Y_1, Y_2, R) \\
&= H(X_1, X_2) + H(Y_1, Y_2) \\
&\leq \log \left(\sum_i a_i^2 \right) + \log \left(\sum_i b_i^2 \right).
\end{aligned}$$

Note that the last inequality used the fact that the number of choices for (X_1, X_2) is $\sum_i a_i^2$ and that for (Y_1, Y_2) is $\sum_i b_i^2$. Combining the above, we get

$$2 \log \left(\sum_i a_i \cdot b_i \right) \leq \log \left(\sum_i a_i^2 \right) + \log \left(\sum_i b_i^2 \right),$$

which yields the Cauchy-Schwarz inequality. This method can also be used to prove interesting generalizations such as

$$\left(\sum_{i,j,k} a_{ij} \cdot b_{jk} \cdot c_{ki} \right)^2 = \left(\sum_{i,j} a_{ij}^2 \right) \cdot \left(\sum_{j,k} b_{jk}^2 \right) \cdot \left(\sum_{k,i} c_{ki}^2 \right).$$

Take a look at [\[Fri04\]](#) for details.

5 Shearer's Lemma and Combinatorial Applications

The sub-additivity property of entropy lets us bound the entropy of the tuple (X_1, \dots, X_n) in terms of the individual entropies $H(X_1), \dots, H(X_n)$. Shearer's lemma can be viewed as a generalization of this statement which lets us obtain better bounds in case we can estimate the entropy of subsets of random variables containing more than one random variable.

Lemma 5.1 (Shearer's Lemma). *Let $\{X_1, \dots, X_n\}$ be a set of random variables. For any $S \subset [n]$, let us denote $X_S = \{X_i : i \in S\}$. Let $\mathcal{F} \subseteq 2^{[n]}$ be a collection of subsets of $[n]$ with the property that for all $i \in [n]$, we have that $|\{S \in \mathcal{F} \mid S \ni i\}| \geq t$. Then*

$$t \cdot H(X_1, \dots, X_n) \leq \sum_{S \in \mathcal{F}} H(X_S).$$

We will actually prove a more general version of the lemma which can be stated in terms of a distribution over subsets of $[m]$ such that for each $i \in [n]$, we have a lower bound on the probability that a random subset from the distribution includes i . The lemma below can easily be seen to imply the version above, by using the uniform distribution on the collection \mathcal{F} .

Lemma 5.2 (Shearer's Lemma: distribution version). *Let $\{X_1, \dots, X_n\}$ be a set of random variables. For any $S \subset [n]$, let us denote $X_S = \{X_i : i \in S\}$. Let D be an arbitrary distribution on $2^{[n]}$ (set of all subsets of $[n]$) and let μ be such that $\forall i \in [n] \mathbb{P}_{S \sim D}[i \in S] \geq \mu$. Then*

$$\mu \cdot H(X_1, \dots, X_n) \leq \mathbb{E}_{S \sim D} [H(X_S)] .$$

Exercise 5.3. *Check that Lemma 5.2 implies Lemma 5.1. Also check that both these lemmas imply sub-additivity.*

A generalization of Shearer's lemma was also used in the paper by Friedgut [Fri04], which is also a great one to read for many interesting applications to analysis and inequalities. We sketch the proof of one such inequality below, for bounding volumes of convex bodies. Please take a look at the excellent surveys by Radhakrishnan [Rad03] and Galvin [Gal14] for a number of interesting combinatorial applications of entropy and Shearer's lemma. If you are interested, there is also a delightful paper of Kontoyiannis [?], which shows how to obtain a somewhat weaker version of the prime number theorem (estimating the number of primes between 1 and n) using entropy arguments.

5.1 Bounding volumes using projections

Consider a set of points S in (say) three dimensions, such that the projections in the xy , yz and zx plain contain n_1 , n_2 and n_3 points respectively. How many points can there be in the set S ? Note that since many points in S can have the same projection on a plane, the numbers n_1 , n_2 and n_3 can each be much smaller than $|S|$. However, since two different points cannot have the same projection in *all three* planes, we know that each triple of projections must determine a unique point. This gives

$$|S| \leq n_1 \cdot n_2 \cdot n_3 .$$

It turns out that we can significantly improve this bound using Shearer's lemma. Let (X, Y, Z) be a triple of random variables denoting the coordinates of a uniformly sampled point from S . Thus, we have that $H(X, Y, Z) = \log |P|$. Moreover, using Shearer's lemma, we also get that

$$2 \cdot H(X, Y, Z) \leq H(X, Y) + H(Y, Z) + H(Z, X) ,$$

since the family of pairs on the right includes each random variable twice. Also, since (X, Y) denotes the projection of a random point from S in the xy plane, and total number of projections is n_1 , we get that $H(X, Y) \leq \log n_1$. Similarly, $H(Y, Z) \leq \log n_2$ and $H(Z, X) \leq \log n_3$. Combining these estimates gives

$$2 \cdot \log |S| \leq \log n_1 + \log n_2 + \log n_3 \quad \Rightarrow \quad |S| \leq \sqrt{n_1 \cdot n_2 \cdot n_3}.$$

Note that there is nothing special about three dimensions. One can also prove the following d -dimensional analogue using the same argument.

Proposition 5.4. *Let $S \subseteq \mathbb{R}^d$ be a finite set of points in d dimensions, and let S_1, \dots, S_d denote the set of projections orthogonal to each of the d coordinate axes. Then we have*

$$|S| \leq \left(\prod_{i=1}^d |S_i| \right)^{1/(d-1)}.$$

This can also be used to bound the volume of a body B in d dimensions in terms of the $(d-1)$ -dimensional volumes of its projections. One can consider the body to be a union of axis parallel cubes, with a point at the center of each cube. Then, a limiting argument combined with the above estimate gives the following result known as the Loomis-Whitney inequality.

Proposition 5.5 (Loomis-Whitney inequality). *Let $B \subseteq \mathbb{R}^d$ be a measurable body and let B_1, \dots, B_d denote its projections orthogonal to each of the coordinate axes. Then, we have*

$$\text{Vol}_d(B) \leq \left(\prod_{i=1}^d \text{Vol}_{d-1}(B_i) \right)^{1/(d-1)}.$$

5.2 Proof of Shearer's lemma

We now prove [Lemma 5.2](#)

Proof: The proof of the lemma follows simply from the chain rule for entropy and the fact that conditioning reduces entropy (on average).

$$\begin{aligned}
\mathbb{E}_{S \sim D} [H(X_S)] &= \mathbb{E}_{S \sim D} \left[\sum_{i \in S} H(X_i | X_{S \cap [i-1]}) \right] && \text{by Chain rule} \\
&\geq \mathbb{E}_{S \sim D} \left[\sum_{i \in S} H(X_i | X_{[i-1]}) \right] && H(X_i | X_A) \geq H(X_i | X_B) \text{ for } A \subset B \\
&= \mathbb{E}_{S \sim D} \left[\sum_{i \in [n]} \mathbb{1}_{\{i \in S\}} \cdot H(X_i | X_{[i-1]}) \right] \\
&= \sum_{i \in [n]} \mathbb{P}_{S \sim D} [i \in S] \cdot H(X_i | X_{[i-1]}) \\
&\geq \mu \cdot \sum_{i \in [n]} H(X_i | X_{[i-1]}) \\
&= \mu \cdot H(X_1, \dots, X_m)
\end{aligned}$$

■

References

- [Fri04] Ehud Friedgut, *Hypergraphs, entropy, and inequalities*, The American Mathematical Monthly **111** (2004), no. 9, 749–760. [5](#), [7](#), [8](#)
- [Gal14] David Galvin, *Three tutorial lectures on entropy and counting*, arXiv preprint arXiv:1406.7872 (2014). [8](#)
- [Kon08] Ioannis Kontoyiannis, *Counting the primes using entropy*, IEEE Information Theory Society Newsletter (2008), 6–9. [8](#)
- [Rad03] Jaikumar Radhakrishnan, *Entropy and counting*, Computational mathematics, modelling and algorithms **146** (2003). [8](#)