

Mathematical Foundations of Information Theory

Discrete Stochastics and Information Theory

Wolfgang Woess, Joshua Erde

Revised by Stephan Wagner

Institute of Discrete Mathematics, TU Graz

Contents

1	Introduction to Probability Theory	4
1.1	Probability Spaces	4
1.2	Random Variables	9
1.3	Markov's and Chebyshev's inequality	13
1.4	Convergence of Random Variables and the Law of Large Numbers	16
2	Discrete Entropy	20
2.1	Hartley's formula and Shannon's formula	20
2.2	Entropy	22
2.3	Kullback-Leibler Divergence and Mutual Information	34
3	Entropy Rate and Asymptotic Equipartition	48
3.1	Entropy Rate	48
3.2	Time-homogeneous Markov chains	50
3.3	The Asymptotic Equipartition Property	67
4	Data compression and Codes	72

4.1	Block codes	72
4.2	Variable length codes	74
4.3	Huffman Codes	80
5	Information Channels	85
5.1	Shannon's channel coding theorem	89
5.2	Source-channel separation theorem	101
6	Differential Entropy	103
6.1	Differential Entropy	103
6.2	Discretization	104
6.3	Joint and conditional differential entropy	105
6.4	The Gaussian Channel	110

Preface

The course is based on the lecture notes of Wolfgang Woess, themselves based partially on the book “Elements of Information Theory” by Cover and Thomas.

1 Introduction to Probability Theory

Probability is a mathematical phenomenon that we see in everyday life that we perhaps intuitively understand. As a motivating example, consider what is called the *law of large numbers*—if we toss a fair coin 1000 times every day, then each day we will get heads *about* 500 times. Of course, we won't get exactly 500 heads, but the *deviations* we observe, over the repeated trials, should be *small*. Similarly, if we roll a fair die many times, the relative frequency of the outcome “6” will be approximately $1/6$. From a certain philosophical viewpoint, this is what we mean when we say “The probability of rolling a 6 is $1/6$ ”.

More generally, the law of large numbers says that if we have some random experiment whose outcome is a real number, and we repeat the experiment many times, then the average of the outcomes should *converge* to some specific, deterministic number, which is the *expected* outcome of the experiment.

In some ways this is intuitive, in other ways almost tautological, but what we want then from a theory of probability is a set of axioms which behaves like how we experience probability in the real world, and so in particular statements like the law of large numbers should follow as a mathematical theorem from these axioms.

1.1 Probability Spaces

Definition 1.1. A *probability space* is a triple $(\Omega, \mathcal{A}, \mathbb{P})$, where

1. Ω is a non-empty set, the *sample space*,
2. \mathcal{A} is a σ -algebra, that is, a collection of subsets of Ω such that
 - (i) $\Omega \in \mathcal{A}$,
 - (ii) $A \in \mathcal{A} \Rightarrow A^c = \Omega \setminus A \in \mathcal{A}$,
 - (iii) if $A_n \in \mathcal{A}$ for $n = 1, 2, \dots$, then $\bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$.
3. \mathbb{P} is a *probability measure* on \mathcal{A} , that is, a function $\mathbb{P} : \mathcal{A} \rightarrow [0, 1]$ such that
 - (i) $\mathbb{P}(\emptyset) = 0$ and $\mathbb{P}(\Omega) = 1$,
 - (ii) if $A_n \in \mathcal{A}$ are pairwise disjoint for $n = 1, 2, \dots$, that is, $A_n \cap A_m = \emptyset$ for all $m \neq n$, then

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mathbb{P}(A_n).$$

We should think of the sample space Ω as consisting of all possible outcomes of some random experiment.

Example 1.2. (a) Our experiment is a coin toss. Our two outcomes are “Heads” and “Tails” and so our sample space is $\Omega = \{\text{Heads}, \text{Tails}\}$. We could equally ‘encode’ the outcomes as Heads = 1 and Tails = 0, in which case our sample space is $\Omega = \{0, 1\}$.

- (b) Our experiment is again a coin toss, but we don't just measure the side that the coin lies on, but also its position on the ground, which is some point (x, y) in the plane with the coin-tosser standing at the origin, as well as the number m of times that the coin rotates whilst in the air. Then, a possible sample space would be

$$\Omega = \{(\ell, x, y, m) : \ell \in \{0, 1\}, (x, y) \in \mathbb{R}^2, m \in \mathbb{N}_0\}.$$

- (c) Our experiment is a sequence of n coin tosses, and we measure the sequence of outcomes. We can take our sample space to be

$$\Omega = \{0, 1\}^n$$

sequences of 0s and 1s of length n , which we call *bit strings*, where the k th element of the sequence is the outcome of the k th coin toss.

- (d) Our (theoretical) experiment is an infinite sequence of coin tosses. Our sample space is then

$$\Omega = \{0, 1\}^{\mathbb{N}},$$

i.e., all infinite sequences of 0s and 1s (an uncountable set!).

The function \mathbb{P} then tells us, for an *event*, a particular subset of the possible outcomes, how likely it is that this event occurs, that is, how likely it is that the outcome lies in this subset.

It turns out, for complicated mathematical reasons, even if the sample space Ω is something familiar like the real numbers \mathcal{R} or the unit interval $[0, 1]$, there is not always a consistent notion of measure that will assign a probability to *every* subset of Ω while also satisfying certain natural requirements (e.g., in the case $\Omega = [0, 1]$ that the measure of an interval should be equal to its length). Very weird sets exist! For this reason we have to restrict ourselves to some 'well-behaved' collection of sets, the σ -algebra \mathcal{A} . However, this is no great restriction, as we can choose \mathcal{A} such that any event that you can actually describe will lie inside \mathcal{A} .

When Ω is *countable*, one can usually take $\mathcal{A} = \mathcal{P}(\Omega)$, the *power set* of Ω , consisting of all subsets of Ω . However, when Ω is uncountable, such as $\Omega = \mathbb{R}$, then it might be necessary to consider a smaller σ -algebra.

For the most part we will work with *discrete probability spaces*, those where Ω is countable, and so avoid these difficulties. In this case, if Ω is countable and $\mathcal{A} = \mathcal{P}(\Omega)$, then the probability measure is determined by the measure of the *elementary events* $\omega \in \Omega$ since

$$\mathbb{P}(A) = \sum_{\omega \in A} \mathbb{P}(\{\omega\}) \quad \text{for all } A \in \mathcal{P}(\Omega).$$

Definition 1.3. Given a logical expression Φ concerning the elements of Ω , we will write $[\Phi]$ for the event

$$A = \{\omega \in \Omega : \Phi \text{ is true for } \omega\},$$

and if $A \in \mathcal{A}$ (which will usually be the case) we define

$$\mathbb{P}[\Phi] := \mathbb{P}(A).$$

Example 1.4. (a) Our experiment is to roll two fair dice. Our sample space is

$$\Omega = \{(i, j) : 1 \leq i, j \leq 6\},$$

and we can take our σ -algebra to be $\mathcal{A} = \mathcal{P}(\Omega)$.

An event we could consider is the event that the total value of the two dice is 11, that is, we consider the event $A = [\text{the total value of the dice is 11}]$, or in other words

$$A = \{(i, j) \in \Omega : i + j = 11\}.$$

(b) Our experiment is as in Example 1.2 (b). We consider the event

$$\begin{aligned} A &= [\text{the coin lands at distance at most } r \text{ from the coin tosser}] \\ &= \{(\ell, x, y, m) : x^2 + y^2 \leq r^2\}. \end{aligned}$$

Similarly the event

$$\begin{aligned} A &= [\text{the coin spins at least 3 times and lands on Heads}] \\ &= \{(\ell, x, y, m) : \ell = 1, m \geq 3\}. \end{aligned}$$

One can easily deduce the following properties from Definition 1.1.

Proposition 1.5. (i) $\emptyset \in \mathcal{A}$,

(ii) If $A_n \in \mathcal{A}$ for $n = 1, 2, \dots$, then $\bigcap_{n=1}^{\infty} A_n \in \mathcal{A}$,

(iii) If $A, B \in \mathcal{A}$ then $A \cup B, A \cap B \in \mathcal{A}$ and

$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B),$$

(iv) If $A \in \mathcal{A}$, then $\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$,

(v) If $A, B \in \mathcal{A}$ and $A \subseteq B$, then $\mathbb{P}(A) \leq \mathbb{P}(B)$.

Proof. (i) Indeed, $\Omega \in \mathcal{A}$ and so $\emptyset = \Omega^c \in \mathcal{A}$.

(ii) Since $A_n \in \mathcal{A}$ for all n , it follows that $A_n^c \in \mathcal{A}$ for all n , and hence $\bigcup_{n=1}^{\infty} A_n^c \in \mathcal{A}$. However, De Morgan's laws say $(\bigcap_{n=1}^{\infty} A_n)^c = \bigcup_{n=1}^{\infty} A_n^c$, and hence $\bigcap_{n=1}^{\infty} A_n = (\bigcup_{n=1}^{\infty} A_n^c)^c \in \mathcal{A}$.

(iii) Note that, by taking the sequence $A = A_1$ and $B = A_2 = A_3 = \dots$ we can conclude that $A \cup B = \bigcup_{n=1}^{\infty} A_i \in \mathcal{A}$, and by the previous claim also $A \cap B = \bigcap_{n=1}^{\infty} A_i \in \mathcal{A}$. Furthermore, by the same reasoning it follows that $B \setminus A = B \cap A^c \in \mathcal{A}$. Now,

$$\begin{aligned} \mathbb{P}(A \cup B) &= \mathbb{P}(A) + \mathbb{P}(B \setminus A) \\ &= \mathbb{P}(A) + \mathbb{P}(B \setminus A) + \mathbb{P}(A \cap B) - \mathbb{P}(A \cap B) \\ &= \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B), \end{aligned}$$

where we used that $A \cup B = A \cup (B \setminus A)$ and $B = (A \cap B) \cup (B \setminus A)$ are both disjoint unions.

(iv) We apply the previous claim with $B = A^c$, so that $\mathbb{P}(A \cup B) = \mathbb{P}(\Omega) = 1$ and $\mathbb{P}(A \cap B) = \mathbb{P}(\emptyset) = 0$. Hence

$$1 = \mathbb{P}(A) + \mathbb{P}(A^c) - 0,$$

which rearranges to the desired identity.

(v) Since $B = A \cup (B \setminus A)$ is a disjoint union, it follows that

$$\mathbb{P}(B) = \mathbb{P}(A) + \mathbb{P}(B \setminus A) \geq \mathbb{P}(A).$$

□

Remark 1.6. From the identity in (iii), i.e.,

$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B),$$

it follows that $\mathbb{P}(A \cup B) \leq \mathbb{P}(A) + \mathbb{P}(B)$. In general (as can be shown by induction), we have

$$\mathbb{P}\left(\bigcup_{k=1}^n A_k\right) \leq \sum_{k=1}^n \mathbb{P}(A_k),$$

which is known as the union bound.

The next lemma is fundamental.

Lemma 1.7 (Continuity of the probability measure). *If $(A_n : n \in \mathbb{N})$ is an increasing sequence, that is $A_n \subseteq A_{n+1}$ for all $n \in \mathbb{N}$, with $A_n \in \mathcal{A}$ for all n , then*

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \lim_{N \rightarrow \infty} \mathbb{P}(A_N).$$

Similarly, if $(A_n : n \in \mathbb{N})$ is a decreasing sequence, that is $A_n \supseteq A_{n+1}$ for all $n \in \mathbb{N}$, with $A_n \in \mathcal{A}$ for all n , then

$$\mathbb{P}\left(\bigcap_{n=1}^{\infty} A_n\right) = \lim_{N \rightarrow \infty} \mathbb{P}(A_N).$$

Proof. Let us start by proving the first statement, the second statement follows by taking complements.

Let us define a new sequence of sets $B_1 = A_1$, $B_2 = A_2 \setminus A_1$, and in general

$$B_n = A_n \setminus A_{n-1} (= A_n \cap A_{n-1}^c).$$

Note that $B_n \in \mathcal{A}$ since $A_n, A_{n-1}^c \in \mathcal{A}$. It is easy to see that the sets B_n are pairwise disjoint, and for any N

$$A_N = \bigcup_{n=1}^N B_n \quad \text{and} \quad \bigcup_{n=1}^{\infty} A_n = \bigcup_{n=1}^{\infty} B_n,$$

Hence

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \mathbb{P}\left(\bigcup_{n=1}^{\infty} B_n\right) = \sum_{n=1}^{\infty} \mathbb{P}(B_n) = \lim_{N \rightarrow \infty} \sum_{n=1}^N \mathbb{P}(B_n) = \lim_{N \rightarrow \infty} \mathbb{P}(A_N).$$

We note that if A_n is decreasing, then A_n^c is increasing. Hence, we can apply the first part to say that

$$\mathbb{P}\left(\bigcap_{n=1}^{\infty} A_n\right) = \lim_{N \rightarrow \infty} \mathbb{P}(A_N^c) = \lim_{N \rightarrow \infty} 1 - \mathbb{P}(A_N) = 1 - \lim_{N \rightarrow \infty} \mathbb{P}(A_N).$$

However, by De Morgan's laws, $(\bigcup_{n=1}^{\infty} A_n^c)^c = \bigcap_{n=1}^{\infty} A_n$ and hence

$$\mathbb{P}\left(\bigcap_{n=1}^{\infty} A_n\right) = 1 - \mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n^c\right) = \lim_{N \rightarrow \infty} \mathbb{P}(A_N).$$

□

Definition 1.8 (Conditional probability, Independence). Given two events $A, B \in \mathcal{A}$ we define

$$\mathbb{P}(A \mid B) = \begin{cases} \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}, & \text{if } \mathbb{P}(B) > 0, \\ 0, & \text{if } \mathbb{P}(B) = 0. \end{cases}$$

Given two logical expressions Φ_1 and Φ_2 , where $A = \{\omega \in \Omega : \Phi_1 \text{ is true for } \omega\}$ and $B = \{\omega \in \Omega : \Phi_2 \text{ is true for } \omega\}$, we will write

$$\mathbb{P}[\Phi_1 \mid \Phi_2] = \mathbb{P}(A \mid B).$$

We say A and B are *independent* if $\mathbb{P}(A \cap B) = \mathbb{P}(A) \cdot \mathbb{P}(B)$. A sequence (finite or infinite) of events $(A_n : n \in I)$ is called *(mutually) independent* if for all choices of indices $J \subseteq I$

$$\mathbb{P}\left(\bigcap_{j \in J} A_j\right) = \prod_{j \in J} \mathbb{P}(A_j).$$

We note that this condition is stronger than asking for *pairwise* independence of the events A_i, A_j for $i, j \in I$.

Example 1.9. Suppose we flip two fair coins, so that $\Omega = \{0, 1\}^2$, $\mathcal{A} = \mathcal{P}(\Omega)$. We can check that each outcome is equally likely, and so for all events $A \subseteq \Omega$

$$\mathbb{P}(A) = \frac{|A|}{|\Omega|} = \frac{|A|}{4}.$$

Let us consider the events

$$\begin{aligned} A_1 &= [\text{the first coin lands Heads}], \\ A_2 &= [\text{the second coin lands Heads}], \\ A_3 &= [\text{both coins land on the same side}], \end{aligned}$$

so that $A_1 = \{(1, 0), (1, 1)\}$, $A_2 = \{(0, 1), (1, 1)\}$ and $A_3 = \{(0, 0), (1, 1)\}$, and $A_i \cap A_j = \{(1, 1)\}$ for all i, j .

In particular, $\mathbb{P}(A_i) = \frac{2}{4} = \frac{1}{2}$ for all $i \leq 3$ and $\mathbb{P}(A_i \cap A_j) = \frac{1}{4} = \mathbb{P}(A_i) \cdot \mathbb{P}(A_j)$ for all $i, j \leq 3$, and hence all pairs of events are independent.

However, $A_1 \cap A_2 \cap A_3 = \{(1, 1)\}$, and so $\mathbb{P}(A_1 \cap A_2 \cap A_3) = \frac{1}{4}$, but

$$\prod_{i=1}^3 \mathbb{P}(A_i) = \left(\frac{1}{2}\right)^3 = \frac{1}{8} \neq \frac{1}{4}.$$

Hence, whilst these three events are pairwise independent, we shouldn't intuitively think of the sequence as being independent. Indeed, if we know that both A_1 and A_2 happen, then it is already determined that A_3 must happen!

1.2 Random Variables

Definition 1.10 (Discrete Random Variable, Distribution). Given a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, a *discrete random variable* is a function $X: \Omega \rightarrow \mathcal{X}$, where $|\mathcal{X}|$ is countable, such that for every $B \subseteq \mathcal{X}$ the set

$$X^{-1}(B) = \{\omega \in \Omega: X(\omega) \in B\} = [X \in B]$$

is a member of the σ -algebra \mathcal{A} .

That is, a discrete random variable is some function of our sample space which takes values in some discrete set \mathcal{X} , such that for any possible subset of \mathcal{X} , the probability that X lies in this subset is well-defined.

The *distribution* of X is the function $P_X: \mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$ given by

$$P_X(B) = \mathbb{P}[X \in B].$$

One can check that $(\mathcal{X}, \mathcal{P}(\mathcal{X}), P_X)$ is a probability space. If $\mathcal{X} \subseteq \mathbb{R}$ we say that X is a discrete *real* random variable.

We can think of random variables as “functions of chance”. When our probability space models the outcome of some random experiment, a random variable extracts some aspect of the experiment which can be measured.

A lot of very natural random variables are not discrete, for example when the observable is not a discrete, but a continuous quantity, and there is a corresponding theory of *continuous* random variables. However, this won’t be relevant until the very last part of the course, and dealing with them formally is a bit more involved. In particular, unless otherwise stated, every random variable we consider in the course will be discrete, and we will only state the relevant results for discrete random variables. However, in almost all cases, analogous statements can be shown to hold for continuous random variables.

So, a random variable X assigns to each outcome ω in the sample space an element $X(\omega) = x \in \mathcal{X}$. It is important then to keep track of the difference between the random variable X and one of the possible values x that X can take.

Example 1.11. Suppose we toss a sequence of n fair coins, so that we have a sample space $\Omega = \{0, 1\}^n$ (and since Ω is finite we can take $\mathcal{A} = \mathcal{P}(\Omega)$). Since the coin is fair, each possible outcome, each sequence $\omega \in \Omega$, is equally likely to occur, and so $\mathbb{P}(\{\omega\}) = 2^{-n}$ for all ω , and $\mathbb{P}(A) = |A|2^{-n}$ for all $A \in \mathcal{A}$.

Now we can look at some random variables, functions from Ω to \mathbb{R} , which are observable quantities from this experiment. For example I could define $X_k(\omega)$ to be the k th element in the sequence ω , the outcome of the k th coin toss.

Then X_k is a discrete random variable, it takes values in $\mathcal{X}_k = \{0, 1\}$, and we can calculate the distribution

$$P_{X_k}(\{0\}) = \mathbb{P}[X_k = 0] = \mathbb{P}[\text{The } k\text{th coin toss is tails}] = \frac{1}{2},$$

and similarly $P_{X_k}(\{1\}) = \frac{1}{2}$, $P_{X_k}(\emptyset) = 0$, $P_{X_k}(\{0, 1\}) = 1$.

Or we could define $S_n(\omega)$ to be the sum of the elements of ω , or in other words the number of heads thrown. Again, S_n is a discrete random variable, taking values in $\mathcal{S}_n = \{0, \dots, n\}$. In this case for each $k \in \{0, \dots, n\}$ we can calculate $P_{S_n}(\{k\}) = \mathbb{P}[S_n = k]$. Indeed, this is just a combinatorial exercise:

$$\begin{aligned}\mathbb{P}[S_n = k] &= |[S_n = k]| 2^{-n} \\ &= |\{\omega : \omega \text{ has precisely } k \text{ zeroes}\}| 2^{-n} \\ &= \binom{n}{k} 2^{-n}.\end{aligned}$$

It is then easy to see that for every $A \subseteq \{0, \dots, n\}$

$$P_{S_n}(A) = \sum_{k \in A} P_{S_n}(\{k\}).$$

Here, note that $S_n(\omega) = \sum_{k=1}^n X_k(\omega)$, and so we might write $S_n = \sum_{k=1}^n X_k$.

More generally, given real random variables X_1, \dots, X_k and some real function f in k variables, we can consider the random variable Y whose value is given by

$$Y(\omega) = f(X_1(\omega), X_2(\omega), \dots, X_k(\omega)),$$

as long as the domain of f contains the product of the ranges of the X_i . In this way the random variables on a probability space form an *algebra*—for many of the common algebraic operations like $f(X_1, X_2) = X_1 + X_2$ or $X_1 \cdot X_2$, the standard properties, e.g. associativity or commutativity of these operations remain valid for random variables.

Definition 1.12 (Discrete density function). Given a discrete random variable X taking values in \mathcal{X} the *discrete density function* $p_X : \mathcal{X} \rightarrow [0, 1]$ is defined by

$$p_X(x) = \mathbb{P}[X = x] = P_X(\{x\}).$$

Thus,

$$1 = \mathbb{P}[X \in \mathcal{X}] = \sum_{x \in \mathcal{X}} p_X(x),$$

and for any $B \subseteq \mathcal{X}$

$$P_X(B) = \sum_{x \in B} p_X(x),$$

and so the discrete density function and the distribution of X determine one another. For this reason, we will often refer to the discrete density function *as* the distribution of the random variable.

We will often think of the discrete density function as a vector $\mathbf{p} \in \mathbb{R}^{\mathcal{X}}$ with $\|\mathbf{p}\|_1 = 1$. Conversely, for any such vector \mathbf{p} there is a random variable X whose density function satisfies $p_X = \mathbf{p}$.

During the course we will normally just introduce random variables by specifying their distributions or discrete density functions, rather than making reference to any specific probability space.

Example 1.13. (a) Given $q \in [0, 1]$ a *Bernoulli random variable* $\text{Ber}(q)$ takes values in $\{0, 1\}$ and has distribution given by

$$p_{\text{Ber}(q)}(1) = q \quad \text{and} \quad p_{\text{Ber}(q)}(0) = 1 - q,$$

so that $p_{\text{Ber}(q)} = (1 - q, q)$. We can think of this random variable as the outcome of a biased coin flip.

(b) Given an event A , the *indicator random variable* of the event A is

$$\mathbb{1}_A(\omega) = \begin{cases} 0 & \text{if } \omega \notin A, \\ 1 & \text{if } \omega \in A. \end{cases}$$

In particular, if $\mathbb{P}(A) = q$, then $p_{\mathbb{1}_A} = (1 - q, q)$, and so $\mathbb{1}_A$ has the same distribution as $\text{Ber}(q)$.

(c) Given $n \in \mathbb{N}$ and $q \in [0, 1]$, a *binomial random variable* $\text{Bin}(n, q)$ takes values in $\{0, \dots, n\}$ and has distribution given by

$$p_{\text{Bin}(n, q)}(k) = \binom{n}{k} q^k (1 - q)^{n-k}.$$

We can think of this random variable as counting the number of heads in a sequence of n consecutive flips of a random coin.

However, in this way, if I have two random variables X and Y , given just in terms of their distributions, this doesn't necessarily tell us 'the whole story'. Indeed, if X and Y are both observables from the same random experiment, then their values may be related in some way - if X is the height of a random person on the street and Y is the weight, then for most outcomes, most people, the values of X and Y will be *positively correlated*, if X is large then Y is more likely to be large and vice versa.

Definition 1.14 (Joint distribution). If we have two discrete random variables X and Y defined on the same probability space, the *joint discrete density function* (which again we will usually refer to as the joint distribution) is defined as

$$p_{X,Y}(x, y) = \mathbb{P}[X = x, Y = y],$$

and from the joint density function we can reconstruct the *marginal* density functions of X and Y , which are given by

$$p_X(x) = \sum_{y \in \mathcal{Y}} p_{X,Y}(x, y) \quad \text{and} \quad p_Y(y) = \sum_{x \in \mathcal{X}} p_{X,Y}(x, y).$$

Note that there can be many different joint density functions $p_{X,Y}$ with the same marginal density functions p_X and p_Y .

More generally, if X_1, \dots, X_n are all discrete random variables, defined on the same probability space, taking values in sets $\mathcal{X}_1, \dots, \mathcal{X}_n$, then the 'vector' of random variables (X_1, \dots, X_n) is also a discrete random variable, which takes values in some subset of the product set $\mathcal{X}_1 \times \dots \times \mathcal{X}_n$. In this case, the *joint discrete density function* is defined as

$$p_{X_1, \dots, X_n}(x_1, \dots, x_n) = \mathbb{P}[X_1 = x_1, X_2 = x_2, \dots, X_n = x_n],$$

and the marginal distributions are defined in the obvious way.

Definition 1.15 (Conditional distribution). Given jointly distributed discrete random variables X and Y , and some value $x \in \mathcal{X}$ with $p_X(x) > 0$, the *conditional density function* (conditional distribution) of Y , given that $X = x$, is

$$p_{Y|X}(y|x) = \mathbb{P}[Y = y|X = x] = \frac{p_{X,Y}(x, y)}{p_X(x)} = \frac{p_{X,Y}(x, y)}{\sum_{y' \in \mathcal{Y}} p_{X,Y}(x, y')}.$$

Note that p_X and $p_{Y|X}$ together determine the joint distribution $p_{X,Y}$ and hence also the marginal density function p_Y .

When the random variables that we are dealing with are clear from the context, we will often drop the subscripts in the notation above and simply write expressions like $p(x)$, $p(x, y)$ or $p(y|x)$.

Example 1.16. Suppose we toss three coins and we let X be the number of heads in the first two coin tosses and Y be the number of heads in the last two coin tosses.

Then we can calculate the joint distribution of X and Y :

(x, y)	0	1	2
0	1/8	1/8	0
1	1/8	1/4	1/8
2	0	1/8	1/8

The marginal distribution p_X is given by the sum of the rows, which is $p_X = (1/4, 1/2, 1/4)$, and the marginal distribution p_Y is given by the sum of the columns, which is $p_Y = (1/4, 1/2, 1/4)$. Note that, as expected, both are distributed as $\text{Bin}(3, 1/2)$.

The conditional distribution $p_{X|Y}$ is then :

$(x y)$	0	1	2
0	1/2	1/4	0
1	1/2	1/2	1/2
2	0	1/4	1/2

Definition 1.17 (Independence). A sequence of discrete random variables X_1, \dots, X_n are *independent*, if for any sequence of subsets $B_1 \subseteq \mathcal{X}_1, \dots, B_n \subseteq \mathcal{X}_n$ the events

$$[X_1 \in B_1], \dots, [X_n \in B_n]$$

are independent. In particular

$$\mathbb{P}[X_1 \in B_1, \dots, X_n \in B_n] = \prod_{i=1}^n \mathbb{P}[X_i \in B_i].$$

One can check that it is equivalent to show that the joint density of the sequence is equal to the product of the marginal distributions, that is

$$p_{X_1, \dots, X_n}(x_1, \dots, x_n) = \prod_{i=1}^n p_{X_i}(x_i) \quad \text{whenever } x_i \in \mathcal{X}_i \text{ for all } i. \quad (1.1)$$

An infinite sequence $(X_n)_{n \in \mathbb{N}}$ of discrete random variables is *independent* if the sequence X_1, \dots, X_n is independent for all n , or equivalently if (1.1) holds for all $n \in \mathbb{N}$.

1.3 Markov's and Chebyshev's inequality

Definition 1.18 (Expectation). The *expectation* or *expected value* or *mean* of a discrete real random variable X is

$$\mathbb{E}(X) := \sum_{x \in \mathcal{X}} x \cdot \mathbb{P}[X = x] = \sum_{x \in \mathcal{X}} x \cdot p_X(x),$$

if the sum converges. Otherwise we informally say that the expectation is infinite.

If $(\Omega, \mathcal{P}(\Omega), \mathbb{P})$ is the underlying probability space, it can sometimes be simpler to use the formula

$$\mathbb{E}(X) = \sum_{\omega \in \Omega} X(\omega) \mathbb{P}(\{\omega\}).$$

When $\mathcal{X} \not\subseteq \mathbb{R}$, so that X is not a real random variable, it does not make sense to talk about the expectation of X . However, for any function $g: \mathcal{X} \rightarrow \mathbb{R}$, we have that $g(X)$ is a real random variable (defined as $g(X)(\omega) = g(X(\omega))$ for all ω) whose expectation we can compute as

$$\mathbb{E}(g(X)) = \sum_{g \in \mathcal{X}} g(x) \cdot p_X(x).$$

Lemma 1.19 (Linearity of expectation). *Let X and Y be jointly distributed discrete real random variables with finite expectations and let $c \in \mathbb{R}$. Then*

- (i) $\mathbb{E}(c) = c$,
- (ii) $\mathbb{E}(X + Y) = \mathbb{E}(X) + \mathbb{E}(Y)$,
- (iii) $\mathbb{E}(c \cdot X) = c \cdot \mathbb{E}(X)$,
- (iv) $x \geq 0$ for all $x \in \mathcal{X} \Rightarrow \mathbb{E}(X) \geq 0$.

Proof. For the first, we think of c as a constant random variable C with $p_C(c) = 1$, so that

$$\mathbb{E}(c) = \mathbb{E}(C) = c \cdot p_C(c) = c.$$

More generally, given $a, b \in \mathbb{R}$ and random variables X and Y we see that

$$\begin{aligned} \mathbb{E}(aX + bY) &= \sum_{\omega \in \Omega} (aX + bY)(\omega) \mathbb{P}(\{\omega\}) \\ &= \sum_{\omega \in \Omega} aX(\omega) \mathbb{P}(\{\omega\}) + bY(\omega) \mathbb{P}(\{\omega\}) \\ &= a\mathbb{E}(X) + b\mathbb{E}(Y). \end{aligned}$$

Finally, if $x \geq 0$ for all $x \in \mathcal{X}$, then since $p_X(x) \geq 0$ it follows that $\mathbb{E}(X)$ is the sum of non-negative terms, and so is also non-negative. \square

However, it is not true in general that $\mathbb{E}(X \cdot Y) = \mathbb{E}(X) \cdot \mathbb{E}(Y)$! This does however hold in the special case where X and Y are independent.

Lemma 1.20. *Let X and Y be independent discrete real random variables with finite expectations. Then $\mathbb{E}(X \cdot Y) = \mathbb{E}(X) \cdot \mathbb{E}(Y)$.*

Proof. Suppose $X \cdot Y$ takes values in \mathcal{Z} . We can write

$$\begin{aligned}
\mathbb{E}(X \cdot Y) &= \sum_{z \in \mathcal{Z}} z \cdot \mathbb{P}[X \cdot Y = z] \\
&= \sum_{z \in \mathcal{Z}} z \sum_{\substack{x \in \mathcal{X}, y \in \mathcal{Y} \\ x \cdot y = z}} \mathbb{P}[X = x, Y = y] \\
&= \sum_{z \in \mathcal{Z}} \sum_{\substack{x \in \mathcal{X}, y \in \mathcal{Y} \\ x \cdot y = z}} x \cdot y \cdot \mathbb{P}[X = x, Y = y] \\
&= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} x \cdot y \cdot \mathbb{P}[X = x, Y = y] \\
&= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} x \cdot y \cdot \mathbb{P}[X = x] \cdot \mathbb{P}[Y = y] \\
&= \sum_{x \in \mathcal{X}} x \cdot \mathbb{P}[X = x] \sum_{y \in \mathcal{Y}} y \cdot \mathbb{P}[Y = y] \\
&= \mathbb{E}(X) \cdot \mathbb{E}(Y).
\end{aligned}$$

□

Another important quantity comes from considering how far a random variable deviates from its expectation.

Definition 1.21 (Variance). Let X be a discrete real random variable with $\mu = \mathbb{E}(X) < \infty$. The *variance* of X is defined as

$$\mathbb{V}(X) = \mathbb{E}((X - \mu)^2).$$

It can easily be shown, using the linearity of expectation, that

$$\mathbb{V}(X) = \mathbb{E}(X^2) - (\mathbb{E}(X))^2.$$

In general, the variance is also not linear, but again for independent random variables, we do have a nice formula.

Lemma 1.22. *Let X and Y be independent discrete real random variables with finite expectations and variances and let $c \in \mathbb{R}$. Then $\mathbb{V}(cX) = c^2 \mathbb{V}(X)$ and $\mathbb{V}(X \pm Y) = \mathbb{V}(X) + \mathbb{V}(Y)$.*

Proof. It is clear that $\mathbb{E}(cX) = c\mathbb{E}(X)$ and $\mathbb{E}((cX)^2) = c^2\mathbb{E}(X^2)$. Hence

$$\mathbb{V}(cX) = \mathbb{E}((cX)^2) - (\mathbb{E}(cX))^2 = c^2 (\mathbb{E}(X^2) - (\mathbb{E}(X))^2) = c^2 \mathbb{V}(X).$$

Since X and Y are independent, it follows from Lemma 1.20 that $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$, and we can expand $(X \pm Y)^2 = X^2 \pm 2XY + Y^2$. Hence

$$\begin{aligned}\mathbb{V}(X \pm Y) &= \mathbb{E}(X^2 \pm 2XY + Y^2) - (\mathbb{E}(X \pm Y))^2 \\ &= \mathbb{E}(X^2) + \mathbb{E}(Y^2) \pm 2\mathbb{E}(XY) - (\mathbb{E}(X) \pm \mathbb{E}(Y))^2 \\ &= \mathbb{E}(X^2) - (\mathbb{E}(X))^2 + \mathbb{E}(Y^2) - (\mathbb{E}(Y))^2 \pm 2(\mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y)) \\ &= \mathbb{V}(X) + \mathbb{V}(Y).\end{aligned}$$

□

It turns out that we can get some control over the distribution of a random variable X just by controlling its expectation or variance. In practice, since many of the random variables that arise are ‘simple’ combinations of ‘simple’ random variables, it is often possible to calculate, estimate or bound the expectation or variance of these random variables, and in this way obtain information about their distributions.

In particular, given an event A , it is easy to calculate the expectation and variance of the indicator random variable $\mathbb{1}_A$. Indeed

$$\mathbb{E}(\mathbb{1}_A) = 1 \cdot \mathbb{P}[\mathbb{1}_A = 1] + 0 \cdot \mathbb{P}[\mathbb{1}_A = 0] = \mathbb{P}(A),$$

and since $\mathbb{1}_A^2 = \mathbb{1}_A$, we see that

$$\mathbb{V}(\mathbb{1}_A) = \mathbb{E}(\mathbb{1}_A^2) - (\mathbb{E}(\mathbb{1}_A))^2 = \mathbb{E}(\mathbb{1}_A) - (\mathbb{E}(\mathbb{1}_A))^2 = \mathbb{P}(A) - \mathbb{P}(A)^2.$$

Lemma 1.23 (Markov’s inequality). *Let X be a non-negative discrete real random variable such that $0 < \mathbb{E}(X) < \infty$ and let $a > 0$. Then*

$$\mathbb{P}[X \geq a] \leq \frac{\mathbb{E}(X)}{a}.$$

Proof. Let $A = [X \geq a]$ and let us consider the indicator random variable $\mathbb{1}_A$.

Note that $X \geq X \cdot \mathbb{1}_A \geq a \cdot \mathbb{1}_A$, pointwise on Ω , and so by Lemma 1.19

$$\mathbb{E}(X) \geq \mathbb{E}(a \cdot \mathbb{1}_A) = a\mathbb{P}[X \geq a],$$

which re-arranges to the desired inequality. □

A simple, but powerful consequence of Markov’s inequality is Chebyshev’s inequality.

Corollary 1.24 (Chebyshev’s inequality). *Let X be a real random variable such that $\mathbb{E}(X), \mathbb{V}(X) < \infty$ and let $a > 0$. Then*

$$\mathbb{P}[|X - \mathbb{E}(X)| \geq a] \leq \frac{\mathbb{V}(X)}{a^2}.$$

Proof. We apply Markov’s inequality (Lemma 1.23) to the random variable $Y = (X - \mathbb{E}(X))^2$. By definition $\mathbb{E}(Y) = \mathbb{V}(X)$ and

$$\mathbb{P}[|X - \mathbb{E}(X)| \geq a] = \mathbb{P}[(X - \mathbb{E}(X))^2 \geq a^2] = \mathbb{P}[Y \geq a^2] \leq \frac{\mathbb{E}(Y)}{a^2} = \frac{\mathbb{V}(X)}{a^2}.$$

□

The proof follows by applying Markov’s inequality to the random variable $(X - \mathbb{E}(X))^2$.

1.4 Convergence of Random Variables and the Law of Large Numbers

Definition 1.25 (Convergence of random variables). Let $(X_n)_{n \in \mathbb{N}}$ be a sequence of real random variables and X a random variable, all defined on the same probability space.

(i) $X_n \longrightarrow X$ *in probability* if for every $a > 0$

$$\lim_{n \rightarrow \infty} \mathbb{P}[|X_n - X| > a] = 0.$$

(ii) $X_n \longrightarrow X$ *almost surely* if

$$\mathbb{P}[X_n \rightarrow X] = 1,$$

that is, if

$$\mathbb{P}(\{\omega \in \Omega : \lim_{n \rightarrow \infty} X_n(\omega) = X(\omega)\}) = 1.$$

We will see that convergence almost surely implies convergence in probability, and so the second is a stronger notion of convergence. We note that the converse is not true! There exist sequences of random variables which converge in probability but not almost surely.

Example 1.26. Let us consider a sequence of independent random variables $(X_n)_{n \in \mathbb{N}}$ each taking values in $\{0, 1\}$ such that

$$\mathbb{P}(X_n = i) = \begin{cases} \frac{1}{n} & \text{if } i = 1, \\ 1 - \frac{1}{n} & \text{if } i = 0. \end{cases}$$

For all $a > 0$, it is clear that

$$\mathbb{P}(|X_n - 0| \geq a) \leq \mathbb{P}(X_n = 1) = \frac{1}{n}$$

and hence, X_n tends to 0 in probability.

On the other hand, $X_n \rightarrow 0$ if and only if there is some N such that $X_n = 0$ for all $n \geq N$. However for any fixed N , since the X_n are independent,

$$\begin{aligned} \mathbb{P}[\forall n \geq N : X_n = 0] &= \lim_{M \rightarrow \infty} \mathbb{P}[\forall N \leq n \leq M : X_n = 0] \\ &= \lim_{M \rightarrow \infty} \prod_{n=N}^M \left(1 - \frac{1}{n}\right) \\ &= \lim_{M \rightarrow \infty} \frac{N-1}{N} \cdot \frac{N}{N+1} \cdots \frac{M-2}{M-1} \cdot \frac{M-1}{M} \\ &= \lim_{M \rightarrow \infty} \frac{N-1}{M} = 0. \end{aligned}$$

Hence,

$$\mathbb{P}[X_n \rightarrow 0] \leq \sum_{N=1}^{\infty} \mathbb{P}[\forall n \geq N : X_n = 0] = 0.$$

Whilst convergence in probability is not enough to guarantee convergence almost surely, it does guarantee the existence of an almost surely convergent subsequence.

Theorem 1.27. *Let $(X_n)_{n \in \mathbb{N}}$ and X be as above. If $X_n \rightarrow X$ in probability, then there is a subsequence $(n_k)_{k \in \mathbb{N}}$ such that $X_{n_k} \rightarrow X$ almost surely.*

Theorem 1.28 (Weak law of large numbers). *Let $(X_n)_{n \in \mathbb{N}}$ be a sequence of independent random variables with finite mean $\mathbb{E}(X_n) = \mu < \infty$ and finite variance $\mathbb{V}(X_n) = \sigma^2 < \infty$ (the same for all n). Then*

$$\bar{X}_n = \frac{1}{n} (X_1 + X_2 + \dots + X_n) \rightarrow \mu \text{ in probability.}$$

Proof. Since the X_i are independent, we can calculate the mean and variance of \bar{X}_n easily.

$$\mathbb{E}(\bar{X}_n) = \mathbb{E}\left(\frac{1}{n} (X_1 + X_2 + \dots + X_n)\right) = \frac{1}{n} (\mathbb{E}(X_1) + \mathbb{E}(X_2) + \dots + \mathbb{E}(X_n)) = \frac{1}{n} \cdot n\mu = \mu$$

and similarly

$$\mathbb{V}(\bar{X}_n) = \frac{\sigma^2}{n}.$$

In particular, Chebyshev's inequality implies that for any $a > 0$

$$\lim_{n \rightarrow \infty} \mathbb{P}[|\bar{X}_n - \mu| \geq a] \leq \lim_{n \rightarrow \infty} \frac{\sigma^2}{a^2 n} = 0.$$

□

As we saw, convergence in probability is weaker than convergence almost surely, and at times we will need the stronger statement that the *sample mean* converges to the mean almost surely.

Theorem 1.29 (Strong law of large numbers). *Let $(X_n)_{n \in \mathbb{N}}$ be a sequence of independent, identically distributed (i.i.d.) random variables (that is, each X_n has the same distribution) with finite mean $\mathbb{E}(X_n) = \mu < \infty$. Then*

$$\bar{X}_n = \frac{1}{n} (X_1 + X_2 + \dots + X_n) \rightarrow \mu \text{ almost surely.}$$

The proof of this theorem is a bit beyond the focus of the course, however we will at time want to relate the two notions of convergence, and will prove a few short lemmas on this topic.

One useful thing for a notion of convergence is that limits should be unique, although here we only have uniqueness up to a set of measure 0.

Lemma 1.30. *Let $X, (X_n)_{n \in \mathbb{N}}$ be jointly distributed real random variables. If $X_n \rightarrow X$ and $X_n \rightarrow X'$ in probability, then $\mathbb{P}[X = X'] = 1$.*

Proof. We note that, by the triangle inequality, for all n

$$|X - X'| = |X - X_n - (X' - X_n)| \leq |X - X_n| + |X' - X_n|,$$

where this inequality holds pointwise on the probability space.

In particular, for all n and for all $a > 0$, we have the inclusion of events

$$\left[|X - X'| > a\right] \subseteq \left[|X - X_n| > \frac{a}{2}\right] \cup \left[|X' - X_n| > \frac{a}{2}\right],$$

since for $|X - X'|$ to be large, at least one of $|X - X_n|$ or $|X' - X_n|$ must be large.

Hence,

$$\mathbb{P}[|X - X'| > a] \leq \lim_{n \rightarrow \infty} \mathbb{P}\left[|X - X_n| > \frac{a}{2}\right] + \mathbb{P}\left[|X' - X_n| > \frac{a}{2}\right] = 0$$

In particular, we can take the events $A_r = \left[|X - X'| > \frac{1}{r}\right]$ and see that A_r is a decreasing sequence of events, and hence

$$\mathbb{P}[X \neq X'] = \mathbb{P}[|X - X'| > 0] = \mathbb{P}\left(\bigcup_r A_r\right) = \lim_{r \rightarrow \infty} \mathbb{P}(A_r) = 0.$$

□

Finally, let us prove that convergence almost surely implies convergence in probability.

Theorem 1.31. *Let $X, (X_n)_{n \in \mathbb{N}}$ be jointly distributed real random variables. If $X_n \rightarrow X$ almost surely, then $X_n \rightarrow X$ in probability.*

Moreover, if we write $U_k = \sup\{|X_n - X| : n \geq k\}$, then $X_n \rightarrow X$ almost surely if and only if $U_k \rightarrow 0$ in probability.

Proof. Let us start by noting that for all $n \in \mathbb{N}$, $|X_n - X| \leq U_n$, and so for all $a > 0$

$$\lim_{n \rightarrow \infty} \mathbb{P}[|X_n - X| > a] \leq \lim_{n \rightarrow \infty} \mathbb{P}[U_n > a] = \lim_{n \rightarrow \infty} \mathbb{P}[|U_n - 0| > a].$$

Hence, if $U_n \rightarrow 0$ in probability, then also $X_n \rightarrow X$ in probability.

Now, by definition,

$$\mathbb{P}[X_n \rightarrow X] = \mathbb{P}\left[\forall r \in \mathbb{N}, \exists k \in \mathbb{N}, \forall n \geq k : |X_n - X| \leq \frac{1}{r}\right] = \mathbb{P}\left[\forall r \in \mathbb{N}, \exists k \in \mathbb{N} : U_k \leq \frac{1}{r}\right].$$

Let us write $A_r = [\exists k \in \mathbb{N} : U_k \leq \frac{1}{r}]$. We notice that A_r is a decreasing sequence of events, and so by the continuity of the probability measure

$$\mathbb{P}[X_n \rightarrow X] = \mathbb{P}\left(\bigcap_{r=1}^{\infty} A_r\right) = \lim_{r \rightarrow \infty} \mathbb{P}(A_r).$$

Hence, since the A_r are decreasing, $\mathbb{P}[X_n \rightarrow X] = 1$ if and only if $\mathbb{P}(A_r) = 1$ for all r .

If we fix r and let $B_k = [U_k \leq \frac{1}{r}]$ then similarly it is easy to check that B_k is an increasing sequence of events, and so

$$\mathbb{P}(A_r) = \mathbb{P}\left(\bigcup_{k=1}^{\infty} B_k\right) = \lim_{k \rightarrow \infty} \mathbb{P}(B_k) = \lim_{k \rightarrow \infty} \mathbb{P}\left[U_k \leq \frac{1}{r}\right] = \lim_{k \rightarrow \infty} \mathbb{P}\left[|U_k - 0| \leq \frac{1}{r}\right].$$

Hence, $\mathbb{P}[X_n \rightarrow X] = 1$ if and only if for all r , $\mathbb{P}\left[|U_k - 0| \leq \frac{1}{r}\right] \rightarrow 1$, or in other words

$$\mathbb{P}\left[|U_k - 0| > \frac{1}{r}\right] \rightarrow 0,$$

which is equivalent to the statement that U_k tends to 0 in probability. □

2 Discrete Entropy

2.1 Hartley's formula and Shannon's formula

Information theory deals with the mathematical problems that arise in the *storage*, *transformation* and *transmission* of information.

We would like to have some sort of theory that measures the *informational content* of some data, which in some way should not depend on the particular form the data takes.

Example 2.1. Suppose I have written down secretly a number from 0 to 31 and you wish to identify the number asking only yes/no questions.

It is, intuitively, clear that the ‘best’ question to start with is “Is your number at most 15?”, since either answer will reduce the number of possibilities by $\frac{1}{2}$. In a similar fashion each question can reduce the number of possibilities by $\frac{1}{2}$, and so after 5 questions you can also identify the number.

Considering a question as a unit of information, we might say that this hidden number then contains 5 of these units, which we will call *bits*, of information.

If we think of encoding the numbers from 0 to 31 in binary, each number corresponds to a sequence in $\{0, 1\}^5$, and the questions that we ask correspond to asking about the value of the k th digit in the sequence.

Hartley made this idea formal in 1928 when he defined the notion of the *uncertainty* of a uniform random sample.

Definition 2.2 (Hartley's formula). Suppose some element is chosen from a collection U_N of N different elements, with each being equally likely. The *uncertainty* of this random element (which one can think of the *informational cost* to identify the element) is given by

$$H(U_N) = \log_2 N.$$

This can be justified in terms of the following (heuristic) *axiomatic requirements*

- (A) $H(U_2) = 1$,
- (B) $H(U_{N+1}) \geq H(U_N)$,
- (C) $H(U_{N \cdot M}) = H(U_N) + H(U_M)$.

The first two are relatively intuitive—the amount of information needed to identify one of two elements is a single question or bit (alternatively, this is just some arbitrary choice to normalise this measure with respect to the units we have chosen). Furthermore, clearly there is more information needed to identify an element from a larger set.

For the third we can think of grouping our elements into N disjoint groups consisting of M elements

$$U_{N \cdot M} = U_M^{(1)} \cup \dots \cup U_M^{(N)}.$$

In order to identify one element from $U_{N \cdot M}$ we could identify first the group $U_M^{(i)}$ that the element lies in, and so identify a uniformly chosen unknown group from a collection of N many groups, and then identify the unknown element of $U_M^{(i)}$, which is equally likely to be any of these elements. Hence, the cost to identify this element is at most $H(U_N) + H(U_M)$.

However, conversely, suppose we choose our random element by first choosing a random group, and then choosing a random element of our group. If we can identify the random element, we can identify both of these random choices, and so the cost to identify this element must be at least $H(U_N) + H(U_M)$.

In fact, Rényi showed that these three properties uniquely determine Hartley's formula.

Lemma 2.3. *The function $H(U_N) = \log_2 N$ is the unique function satisfying properties (A)–(C).*

Proof. Given $k \in \mathbb{N}$ let us define $s(k) \in \mathbb{N}$ such that $2^{s(k)} \leq N^k < 2^{s(k)+1}$ (i.e., $s(k) = \lfloor k \log_2 N \rfloor$). Hence,

$$\frac{s(k)}{k} \leq \log_2 N < \frac{s(k) + 1}{k}.$$

Taking limits as $k \rightarrow \infty$, we see that $\frac{s(k)}{k} \rightarrow \log_2 N$.

However, for any $k \in \mathbb{N}$

$$\begin{aligned} s(k) &\stackrel{(A)}{=} s(k)H(U_2) \stackrel{(C)}{=} H(U_{2^{s(k)}}) \stackrel{(B)}{\leq} H(U_{N^k}) \stackrel{(C)}{=} kH(U_N) \\ &\stackrel{(B)}{\leq} H(U_{2^{s(k)+1}}) \stackrel{(C)}{=} (s(k) + 1)H(U_2) \stackrel{(A)}{=} s(k) + 1. \end{aligned}$$

and so

$$\frac{s(k)}{k} \leq H(U_N) \leq \frac{s(k) + 1}{k}$$

and taking limits as $k \rightarrow \infty$ implies that $H(U_N) = \lim_{k \rightarrow \infty} \frac{s(k)}{k} = \log_2 N$. □

Suppose now that our elements are not equally likely to be chosen, but that the k th element is instead chosen with some probability p_k . Can we justify, using the previous heuristic, what the informational cost of identifying the chosen element is?

Well, in some sense the ‘cost’ to identify the hidden element does not change, we still might need to identify any one of N elements. However, if one of the p_k s were much larger than all the others, so in almost every case the k th element is chosen, it would be much more sensible to start by asking “is the hidden element the k th element”? In the worst case we would have to ask more questions, but *on average* we would identify the element with many fewer questions!

So, it makes sense instead to consider the *expected uncertainty*, or the expected informational cost to identify the unknown element. The following is a *heuristic* argument for how we should define this quantity.

Let us assume that the probabilities p_k are all rational, otherwise we can take some rational approximations and argue “in the limit”. Instead of an element from U_N where the k th element is chosen with probability p_k , I could choose an element from a larger set

$$U_M = U_{M_1} \cup \dots \cup U_{M_N}$$

where U_{M_i} contains $p_i M$ elements, for some large M such that all these numbers are integers. There is then a clear equivalence between identifying the group U_{M_i} in which this element lies, and of identifying the element in the original problem.

By a similar argument as before, the expected informational cost of identifying the random element of U_M , which should be $\log_2 M$ by Hartley’s formula, should be given by the expected cost to identify the correct group U_{M_i} , which is the quantity we are interested in, which we denote by H_1 , plus the expected cost to identify the correct element of this group, which we denote by H_2 . Now the element lies in M_i with probability p_i , and if the element lies in U_{M_i} then the informational cost to identify it is $\log_2 M_i$ by Hartley’s formula. Hence, the expected cost is

$$H_2 = \sum_i p_i \log_2 M_i = \sum_i p_i \log_2 p_i M = \sum_i p_i \log_2 p_i + \sum_i p_i \log_2 M = \log_2 M + \sum_i p_i \log_2 p_i.$$

Since $H_1 + H_2 = \log_2 M$, it follows that

$$H_1 = - \sum_i p_i \log_2 p_i$$

which is known as *Shannon’s formula*. In the following section we will make this informal discussion mathematically rigorous.

2.2 Entropy

The idea of entropy originated in statistical mechanics. Roughly, given a thermodynamic system, such as a gas or a liquid, if we know some global properties of the system, e.g. temperature, volume, energy, there are many different *microstates*, that is configurations of the individual particles within the system, which are consistent with these measurements.

As an example imagine flipping 1000 coins. We have a global measurement, the number of heads, but for each particular value for this, there are many different configurations of the specific states each of the 1000 coins landed in which achieve this number of heads.

Under a broad assumption that each of these microstates is equally likely, Boltzmann defined entropy of the system to be $k_B \log(\# \text{ of microstates})$ where k_B is some constant. Gibbs generalized this to microstates with unequal probabilities and gave the formula

$$S = -k_B \sum p_i \log(p_i),$$

where S is the *entropy*, p_i is the probability of the i th microstate, and the sum ranges over all the microstates. This reduces to Boltzmann’s formula when the p_i are equal.

The second law of thermodynamics states that the entropy of an isolated system never decreases, and so such systems naturally ‘tend’ towards the state with maximum entropy, known

as thermodynamic equilibrium. This was an attempt to formalise the idea that there is a natural ‘direction’ to natural processes, for example to explain why heat is transferred from hotter objects to cooler objects, rather than the other way round (which would not by itself contradict the conservation of energy in a process).

In the early 20th century Hartley and Shannon found that similar equations arise naturally in the study of information theory, and at the suggestion of Von Neumann, Shannon also named it *entropy*.

“You should call it entropy, for two reasons. In the first place your uncertainty function has been used in statistical mechanics under that name, so it already has a name. In the second place, and more important, nobody knows what entropy really is, so in a debate you will always have the advantage.” – John von Neumann

Definition 2.4 (Entropy). Let X be a discrete random variable taking values in a *finite* set \mathcal{X} , and let

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

be the distribution of X . The *entropy* of X , which we will also call the entropy of the distribution p , is defined as

$$H(X) = H(p) = - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x). \quad (2.1)$$

If we enumerate $\mathcal{X} = \{x_1, \dots, x_n\}$ and set $p_k = p(x_k)$ then we might also use the following notation, that $p = (p_1, \dots, p_n)$ and

$$H(p) = H(p_1, \dots, p_n) = - \sum_{i=1}^n p_i \log_2 p_i.$$

Remark 2.5. For ease of notation, it will often be convenient to define

$$0 \log 0 := 0,$$

whenever it appears in such a sum.

Example 2.6. Suppose that X is uniformly distributed on a set $\mathcal{X} = \{x_1, \dots, x_n\}$ of size n , so that $p_k = p(x_k) = \frac{1}{n}$ for each k .

In this case

$$H(X) = H(p) = H\left(\frac{1}{n}, \dots, \frac{1}{n}\right) = - \sum_{i=1}^n p_i \log_2 p_i = - \sum_{i=1}^n \frac{1}{n} \log_2 \frac{1}{n} = - \log_2 \frac{1}{n} = \log_2 n.$$

Another way to think of entropy is as a measure of the expected amount of information we gain from learning the value of X . Indeed, suppose we have some function $g(x)$ which measures the information we gain from learning that X takes the value x . We clearly gain more information by knowing that a low probability event happens, so this function $g(x)$ should be a decreasing function of $p(x)$. In fact, as we will see later, there are other natural assumptions to make about $g(x)$ which, similar to Hartley’s formula, imply that the only ‘reasonable’ choice for this function g is to take $g(x) = -\log_2 p(x)$.

In this way, we can view (2.1) as an expectation—we have the weighted sum over some probability distribution of a quantity, where this quantity is the function $g : \mathcal{X} \rightarrow \mathbb{R}$ given by $g(x) = -\log_2 p(x)$ (note that this is a deterministic function, even though it encodes the probability distribution of the random variable X), then (2.1) can be rewritten as

$$H(X) = \sum_{x \in \mathcal{X}} p(x)g(x) = \mathbb{E}(g(X)) = \mathbb{E}(-\log_2 p(X)),$$

and it represents the expected amount of information we gain from learning the value of X .

Let us collect a few basic facts about the entropy function, some of which are obvious and some of which we will prove formally later.

Remark 2.7. (1) $H(X) \geq 0$, with equality if and only if X is constant.

(2) $H(X)$ doesn't depend on the values of the random variable X , just the distribution of probabilities between these values. In other words, if we relabel the outcomes, that is, if we take some bijection $f : \mathcal{X} \rightarrow \mathcal{X}'$ and let $X' = f(X)$, then $H(X') = H(X)$.

In other words if (p_1, \dots, p_n) and (p'_1, \dots, p'_n) are the same up to some permutation, then

$$H(p_1, \dots, p_n) = H(p'_1, \dots, p'_n).$$

(3) The function $p_1 \mapsto H(p_1, 1 - p_1)$ is continuous for $p_1 \in [0, 1]$. Furthermore this function is symmetric, takes values 0 at $p_1 = 0, 1$ and is maximised at $p_1 = \frac{1}{2}$ where it takes the value 1.

(4) More generally, for fixed n ,

$$\max\{H(p) : p = (p_1, \dots, p_n)\} = H\left(\frac{1}{n}, \dots, \frac{1}{n}\right) = \log_2 n,$$

which the last equality is equivalent to Hartley's formula.

In other words, the uniform distribution has the maximum expected uncertainty, or the maximum expected information.

A discrete random variable X determines another jointly distributed random variable Y if there is an function $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that $Y = f(X)$.

Lemma 2.8. Let X and Y be jointly distributed discrete random variables taking finitely many values such that X determines Y . Then $H(Y) \leq H(X)$.

Proof. Suppose that X takes values in \mathcal{X} and Y takes values in \mathcal{Y} . By definition there is some function $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that $Y = f(X)$. Then, for every $y \in \mathcal{Y}$ we have that $p(y) = \sum_{x: f(x)=y} p(x)$ and $p(x) \leq p(f(x))$. Hence

$$\begin{aligned} H(Y) &= - \sum_{y \in \mathcal{Y}} p(y) \log_2 p(y) \\ &= - \sum_{y \in \mathcal{Y}} \sum_{x: f(x)=y} p(x) \log_2 p(y) \\ &= - \sum_{x \in \mathcal{X}} p(x) \log_2 p(f(x)) \\ &\leq - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x) = H(X). \end{aligned}$$

□

In particular, if X determines Y and Y determines X , then $H(X) = H(Y)$. In fact, if this holds, there must be some bijection $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that $f(X) = Y$, and so this holds by part (2) of Remark 2.7.

Given jointly distributed discrete random variables X and Y , taking values in finite sets \mathcal{X} and \mathcal{Y} , as mentioned the random vector $Z = (X, Y)$ is again a discrete random variable and we can define the *joint entropy* of X and Y as the entropy of Z . That is, since for any $z = (x, y) \in \mathcal{X} \times \mathcal{Y}$

$$\mathbb{P}[Z = z] = \mathbb{P}[X = x, Y = y] = p_{X,Y}(x, y),$$

we can calculate

$$H(X, Y) := H(Z) = - \sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p(x, y) \log_2 p(x, y).$$

Similarly, given $x \in \mathcal{X}$ we can consider the conditional distribution of Y , given that $X = x$, which we recall is

$$p_{Y|X}(y|x) = \frac{p_{X,Y}(x, y)}{p_X(x)}$$

assuming that $p_X(x) > 0$. We can thus write the entropy of this conditional distribution as

$$H(Y | X = x) := - \sum_{y \in \mathcal{Y}} p(y|x) \log_2 p(y|x) = - \sum_{y \in \mathcal{Y}} \frac{p(x, y)}{p(x)} \log_2 \frac{p(x, y)}{p(x)}.$$

Example 2.9. Suppose X and Y are both distributed on $\{1, 2, 3\}$ and have joint distribution given by

	$X = 1$	$X = 2$	$X = 3$
$Y = 1$	1/8	1/4	0
$Y = 2$	1/8	1/8	1/4
$Y = 3$	0	1/8	0

so that $p_X = (1/4, 1/2, 1/4)$ and $p_Y = (3/8, 1/2, 1/8)$. In this case we can calculate

$$H(X) = \frac{1}{4} \log_2 4 + \frac{1}{2} \log_2 2 + \frac{1}{4} \log_2 4 = \frac{3}{2},$$

$$H(Y) = \frac{3}{8} \log_2 \frac{8}{3} + \frac{1}{2} \log_2 2 + \frac{1}{8} \log_2 8 = 2 - \frac{3}{8} \log_2 3,$$

$$H(X, Y) = 4 \cdot \frac{1}{8} \log_2 8 + 2 \cdot \frac{1}{4} \log_2 4 = \frac{5}{2},$$

$$H(Y|X = 1) = \frac{1}{2} \log_2 2 + \frac{1}{2} \log_2 2 + 0 \log_2 0 = 1.$$

Definition 2.10 (Conditional entropy). The conditional entropy of a discrete random variable Y given a discrete random variable X , both taking finitely many values, is the average value of $H(Y|X = x)$ with respect to the possible values of X

$$\begin{aligned} H(Y | X) &= \sum_{x \in \mathcal{X}} p_X(x) H(Y | X = x) \\ &= - \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log_2 p(y|x) \\ &= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 \frac{p(x, y)}{p(x)}. \end{aligned}$$

So in the previous example we can calculate

$$H(Y | X) = \frac{1}{8} \log_2 2 + \frac{1}{8} \log_2 2 + \frac{1}{4} \log_2 2 + \frac{1}{8} \log_2 4 + \frac{1}{8} \log_2 4 + \frac{1}{4} \log_2 1 = 1.$$

We can think of $H(Y|X)$ as the expected amount of information contained in the random variable Y if we already know the value of X . In particular, if $H(X, Y)$ represents the expected total information in both X and Y , then since discovering the value of X and Y is the same as first discovering the value of X , and then discovering the value of Y , heuristically it should be the case that $H(X, Y) = H(X) + H(Y | X)$, and indeed in the example above

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

It should heuristically be true that conditioning can only decrease the entropy, and indeed this is the case. Later we will show a far more general statement.

Lemma 2.11. *For any two jointly distributed discrete random variables X and Y taking finitely many values $H(Y) \geq H(Y | X)$.*

Proof. We use the elementary inequality $t \leq e^{t-1}$ which holds for all $t \in \mathbb{R}$ to deduce that $\log_2 t \leq (t - 1) \log_2 e$ for all $t > 0$. Hence

$$\begin{aligned} H(Y | X) &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 \frac{p(x)}{p(x, y)} \\ &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \left(\log_2 \frac{1}{p(y)} + \log_2 \frac{p(x)p(y)}{p(x, y)} \right) \\ &\leq \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \left(\log_2 \frac{1}{p(y)} + \left(\frac{p(x)p(y)}{p(x, y)} - 1 \right) \log_2 e \right) \\ &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \left(\log_2 \frac{1}{p(y)} \right) \\ &= - \sum_{y \in \mathcal{Y}} p(y) \log_2 p(y) \\ &= H(Y), \end{aligned}$$

since $\sum_{x,y} p(x, y) = \sum_{x,y} p(x)p(y) = 1$. □

Theorem 2.12 (Chain rule). *For any two jointly distributed discrete random variables X and Y taking finitely many values*

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

Proof.

$$\begin{aligned}
H(X) + H(Y | X) &= - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x) - \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log_2 p(y|x) \\
&= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 p(x) - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 p(y|x) \\
&= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) (\log_2 p(x) + \log_2 p(y|x)) \\
&= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 (p(x) \cdot p(y|x)) \\
&= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 p(x, y) \\
&= H(X, Y).
\end{aligned}$$

The second equality holds then by symmetry. \square

More generally, using Theorem 2.12 it can be shown by induction that the following holds.

Theorem 2.13 (Chain rule). *For any sequence of discrete random variables X_1, X_2, \dots, X_n taking finitely many values*

$$H(X_1, \dots, X_n) = \sum_{k=1}^n H(X_k | X_{k-1}, \dots, X_1).$$

In fact, the chain rule holds in a slightly more general form, for conditional entropies, which can be proved in much the same way.

Lemma 2.14 (Conditional Chain Rule). *For any three jointly distributed random variables X, Y and Z taking finitely many values*

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

Arguably, the notion of conditional entropy is the more fundamental one, since one can view the entropy $H(X)$ as just the entropy $H(X | \emptyset)$ of X conditioned on some trivial random variable (or on an empty set of random variables). All the normal rules for entropy follow then from the rules from conditional entropy in this manner, but the converse is not true.

Example 2.15. Suppose Z takes values in $\{1, \dots, n\}$ and has probability distribution (p_1, \dots, p_n) . Let us define two new random variables: $X = Z + \mathbb{1}_{[Z=1]}$ and $Y = \mathbb{1}_{[Z=1]}$. In particular, $p_X = (p_1 + p_2, p_3, \dots, p_n)$ and $p_Y = (1 - p_1, p_1)$.

Now, since $Z = X - Y$, the pair (X, Y) determines Z , and clearly X and Y are determined by Z , and so

$$H(X, Y) = H(Z) = H(p_1, \dots, p_n), \quad H(X) = H(p_1 + p_2, p_3, \dots, p_n), \quad H(Y) = H(1 - p_1, p_1).$$

Now, if $X = x \geq 3$, then $Y = 0$ and so $H(Y | X = x) = 0$. If $X = 2$, which happens with probability $p_X(2) = p_1 + p_2$, then Z is either 1 or 2, with probabilities p_1 and p_2 , and so Y is

either 1 or 0 with the same probabilities, that is,

$$p_{Y|X}(1 | 2) = \frac{p_1}{p_1 + p_2} \quad \text{and} \quad p_{Y|X}(0 | 2) = \frac{p_2}{p_1 + p_2}.$$

Hence we can calculate

$$H(Y|X) = \sum_{i=2}^n p_X(i) H(Y|X = i) = (p_1 + p_2) H\left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2}\right),$$

and the chain rule $H(X, Y) = H(X) + H(Y|X)$ in this case implies

$$H(p_1, \dots, p_n) = H(p_1 + p_2, p_3, \dots, p_n) + (p_1 + p_2) H\left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2}\right), \quad (2.2)$$

which holds for any probability distribution (p_1, \dots, p_n) (if we interpret the second term as 0 when $p_1 + p_2 = 0$).

For Hartley's formula, there was a heuristic collection of axioms that determine what we should expect from a measure of uncertainty that in fact determined Hartley's formula as the unique way to capture these axioms mathematically. It turns out that there is a similar axiomatic basis for Shannon's formula, in terms of some natural axioms that any measure of expected uncertainty should satisfy.

Theorem 2.16. *Let $\mathcal{B} = \{(p_1, \dots, p_n) : n \in \mathbb{N}, p_k \geq 0 \text{ for all } k \leq n, p_1 + \dots + p_n = 1\}$ be the set of all finite probability distributions. Suppose that we have some function $H : \mathcal{B} \rightarrow \mathbb{R}$ which satisfies the following axioms:*

(I) *H is transposition invariant : if $1 \leq i < j \leq n$ then*

$$H(p_1, \dots, p_i, \dots, p_j, \dots, p_n) = H(p_1, \dots, p_j, \dots, p_i, \dots, p_n).$$

(It is easy to show that this implies H is permutation invariant).

(II) *Normalisation : $H(1/2, 1/2) = 1$.*

(III) *Continuity : The function $p_1 \rightarrow H(p_1, 1 - p_1)$ is continuous.*

(IV) *Equation (2.2) holds for all $p \in \mathcal{B}$ with $n \geq 2$.*

Then

$$H(p_1, \dots, p_n) = - \sum_{k=1}^n p_k \log_2 p_k.$$

Proof. (For mathematics students only.) □

In order to do this we will need the following variant of Lemma 2.3.

Proposition 2.17. *The function $H(U_N) = \log_2 N$ is the unique function satisfying properties (A),(C) of Lemma 2.3 and the following variant of property (B):*

$$(B^*) \lim_{N \rightarrow \infty} H(U_{N+1}) - H(U_N) = 0.$$

Let us start by showing that Theorem 2.16 follows from Proposition 2.17.

Proof of Theorem 2.16. We start by considering $H(1, 0)$. By Property (IV) we have that

$$H(1, 0) = H(1) + 1 \cdot H(1, 0) \Rightarrow H(1) = 0. \quad (2.3)$$

Similarly, Properties (I) and (IV) imply that

$$H(p_1, \dots, p_n, 0) = H(0, p_1, \dots, p_n) = H(p_1, \dots, p_n) + p_1 H(0, 1) = H(p_1, \dots, p_n) + p_1 H(1, 0). \quad (2.4)$$

But equally, using Property (I) to rearrange, we see by symmetry that

$$H(p_1, \dots, p_n, 0) = H(p_1, \dots, p_n) + p_i H(1, 0)$$

for any i . In particular, by choosing some non-constant p , we see that

$$H(1, 0) = 0. \quad (2.5)$$

In particular, (2.4) and (2.5) imply that for any $(p_1, \dots, p_n) \in \mathcal{B}$

$$H(p_1, \dots, p_n, 0) = H(p_1, \dots, p_n). \quad (2.6)$$

By Property (I), a similar statement holds if any coordinate is 0, and so we may assume without loss of generality that $p_i > 0$ for all i .

Claim 1. For all $N \geq 0$ and $m \geq 2$

$$H(p_1, \dots, p_m, p_{m+1}, \dots, p_{N+m}) = H(q, p_{m+1}, \dots, p_{N+m}) + qH\left(\frac{p_1}{q}, \dots, \frac{p_m}{q}\right)$$

where $q = \sum_{k=1}^m p_k > 0$.

Proof. We prove this by induction on m . For $m = 2$ this is just the statement of (IV).

Suppose the claim holds for some m . Given (p_1, \dots, p_{N+m+1}) , and writing $p' = p_1 + p_2$ by Property (IV) we have

$$H(p_1, \dots, p_{N+m+1}) = H(p', p_3, \dots, p_{N+m+1}) + p' H\left(\frac{p_1}{p'}, \frac{p_2}{p'}\right). \quad (2.7)$$

Applying the induction hypothesis to the first term we see that

$$H(p', p_3, \dots, p_{N+m+1}) = H(q, p_{m+1}, \dots, p_{N+m+1}) + qH\left(\frac{p'}{q}, \frac{p_3}{q}, \dots, \frac{p_m}{q}\right). \quad (2.8)$$

However, by Property (IV)

$$H\left(\frac{p_1}{q}, \frac{p_2}{q}, \dots, \frac{p_m}{q}\right) = H\left(\frac{p'}{q}, \frac{p_3}{q}, \dots, \frac{p_m}{q}\right) + \frac{p'}{q} H\left(\frac{p_1}{p'}, \frac{p_2}{p'}\right). \quad (2.9)$$

By combining (2.7)–(2.9) we see that

$$\begin{aligned}
H(p_1, \dots, p_{N+m+1}) &= H(p', p_3, \dots, p_{N+m+1}) + p' H\left(\frac{p_1}{p'}, \frac{p_2}{p'}\right) \\
&= H(q, p_{m+1}, \dots, p_{N+m+1}) + q H\left(\frac{p'}{q}, \frac{p_3}{q}, \dots, \frac{p_m}{q}\right) + p' H\left(\frac{p_1}{p'}, \frac{p_2}{p'}\right) \\
&= H(q, p_{m+1}, \dots, p_{N+m+1}) + q H\left(\frac{p_1}{q}, \frac{p_2}{q}, \dots, \frac{p_m}{q}\right)
\end{aligned}$$

as claimed. \square

Claim 2. Given probability vectors $p^{(1)}, \dots, p^{(n)}$ of lengths N_1, \dots, N_n respectively, and a probability vector q of length n , let

$$u = (q_1 \cdot p^{(1)}, q_2 \cdot p^{(2)}, \dots, q_n \cdot p^{(n)})$$

be the probability vector of length $\sum_{k=1}^n N_k$ obtained by concatenating the vectors $q_k \cdot p^{(k)}$ for $k = 1, \dots, n$. Then

$$H(u) = H(q) + \sum_{k=1}^n q_k H(p^{(k)}).$$

Proof. If $n = 1$, then $q = 1$, and there is nothing to prove. We induct on n and suppose that $n \geq 2$.

We apply the previous claim recursively, using Property (I) to rearrange coordinates, to see

$$\begin{aligned}
H(u) &= H(q_1 \cdot p^{(1)}, q_2 \cdot p^{(2)}, \dots, q_n \cdot p^{(n)}) \\
&= H(q_1, q_2 \cdot p^{(2)}, q_3 \cdot p^{(3)}, \dots, q_n \cdot p^{(n)}) + q_1 H(p^{(1)}) \\
&= H(q_2 \cdot p^{(2)}, q_1, q_3 \cdot p^{(3)}, \dots, q_n \cdot p^{(n)}) + q_1 H(p^{(1)}) \\
&= H(q_2, q_1, q_3 \cdot p^{(3)}, \dots, q_n \cdot p^{(n)}) + q_1 H(p^{(1)}) + q_2 H(p^{(2)}) \\
&= \dots \\
&= H(q_n, q_{n-1}, \dots, q_1) + \sum_{k=1}^n q_k H(p^{(k)}) \\
&= H(q) + \sum_{k=1}^n q_k H(p^{(k)}).
\end{aligned}$$

\square

Claim 3. The function $f(U_N) = H(1/N, \dots, 1/N)$ satisfies Properties (A), (B^*) and (C).

Proof. Property (A) says that $H(1/2, 1/2) = 1$, which is Property (II).

Property (C) says that $H(1/(N \cdot M), \dots, 1/(N \cdot M)) = H(1/N, \dots, 1/N) + H(1/M, \dots, 1/M)$. This follows from Claim 2 applied with $p^{(1)} = \dots = p^{(M)} = (1/N, \dots, 1/N)$ and $q = (1/M, \dots, 1/M)$.

So it remains to show Property (B*). By Claim 1, applied to the first n coordinates

$$\begin{aligned} f(U_{N+1}) &= H\left(\frac{1}{N+1}, \dots, \frac{1}{N+1}\right) \\ &= H\left(1 - \frac{1}{N+1}, \frac{1}{N+1}\right) + \left(1 - \frac{1}{N+1}\right) H\left(\frac{1}{N}, \dots, \frac{1}{N}\right) \\ &= H\left(1 - \frac{1}{N+1}, \frac{1}{N+1}\right) + \left(1 - \frac{1}{N+1}\right) f(U_N). \end{aligned}$$

Rearranging the above we see that

$$d_N := f(U_{N+1}) - f(U_N) = H\left(1 - \frac{1}{N+1}, \frac{1}{N+1}\right) - \frac{f(U_N)}{N+1} := \delta_N - \frac{f(U_N)}{N+1}. \quad (2.10)$$

Now, since H is continuous and $H(1, 0) = 0$, we have that

$$\lim_{N \rightarrow \infty} \delta_N = \lim_{N \rightarrow \infty} H\left(1 - \frac{1}{N+1}, \frac{1}{N+1}\right) = H(1, 0) = 0.$$

On the other hand, since $f(U_1) = H(1) = 0$, by (2.10) we can write δ_N as a telescoping sum

$$\begin{aligned} \delta_N &= d_N + \frac{f(U_N)}{N+1} \\ &= d_N + \frac{1}{N+1} \sum_{k=1}^{N-1} d_k, \end{aligned} \quad (2.11)$$

from which it follows that

$$\sum_{k=1}^N (k+1) \delta_k = (N+1) \sum_{k=1}^N d_k. \quad (2.12)$$

We use the following fact, whose proof is a simple exercise in analysis: if $a_n \rightarrow a$ and $b_n > 0$ with $\sum_{n=1}^{\infty} b_n = \infty$, then

$$\lim_{N \rightarrow \infty} \frac{\sum_{k=1}^N a_k b_k}{\sum_{k=1}^N b_k} = a.$$

Hence, since $\sum_{k=1}^N (k+1) = \frac{N(N+3)}{2}$, we see that

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N d_k &= \lim_{N \rightarrow \infty} \frac{1}{N(N+1)} \left((N+1) \sum_{k=1}^N d_k \right) \\ &= \lim_{N \rightarrow \infty} \frac{N+3}{N+1} \cdot \frac{(N+1) \sum_{k=1}^N d_k}{2 \sum_{k=1}^N (k+1)} \\ &= \frac{1}{2} \lim_{N \rightarrow \infty} \frac{\sum_{k=1}^N (k+1) \delta_k}{\sum_{k=1}^N (k+1)} \\ &= \frac{1}{2} \lim_{N \rightarrow \infty} \delta_N = 0, \end{aligned}$$

where we applied the above fact with $a_k = \delta_k$ and $b_k = k+1$.

Therefore, by (2.11), it follows that

$$\lim_{N \rightarrow \infty} d_N = 0,$$

which is Property (B*). □

Hence, we may assume from Claim 3 and Proposition 2.17 that $f(U_N) = H(1/N, \dots, 1/N) = \log_2 N$. Hence, for any positive integers $k < N$, by Claim 2 applied with $p^{(1)} = (1/k, \dots, 1/k)$, $p^{(2)} = (1/(N-k), \dots, 1/(N-k))$ and $q = (k/N, (N-k)/N)$

$$\begin{aligned} \log_2 N &= H\left(\frac{1}{N}, \dots, \frac{1}{N}\right) \\ &= H\left(\frac{k}{N}, \frac{N-k}{N}\right) + \frac{k}{N} H\left(\frac{1}{k}, \dots, \frac{1}{k}\right) + \frac{N-k}{N} H\left(\frac{1}{N-k}, \dots, \frac{1}{N-k}\right) \\ &= H\left(\frac{k}{N}, \frac{N-k}{N}\right) + \frac{k}{N} \log_2 k + \frac{N-k}{N} \log_2(N-k). \end{aligned}$$

Rearranging the above we see

$$\begin{aligned} H\left(\frac{k}{N}, \frac{N-k}{N}\right) &= \log_2 N - \frac{k}{N} \log_2 k - \frac{N-k}{N} \log_2(N-k) \\ &= -\frac{k}{N} \log_2 \frac{k}{N} - \frac{N-k}{N} \log_2 \frac{N-k}{N}. \end{aligned}$$

In particular, it follows that for all rational probability vectors $p = (p_1, p_2) \in \mathbb{Q}^2$

$$H(p_1, p_2) = -p_1 \log_2 p_1 - p_2 \log_2 p_2,$$

and hence by Property (III), this holds for all probability vectors of length two.

Finally we can use Property (IV) to recursively show that the claim holds for all vectors of length n , indeed

$$\begin{aligned} H(p_1, \dots, p_n) &= (p_1 + p_2) H\left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2}\right) + H(p_1 + p_2, \dots, p_{n-2}, p_{n-1}, p_n) \\ &= (p_1 + p_2) \left(-\frac{p_1}{p_1 + p_2} \log_2 \frac{p_1}{p_1 + p_2} - \frac{p_2}{p_1 + p_2} \log_2 \frac{p_2}{p_1 + p_2} \right) \\ &\quad - (p_1 + p_2) \log_2(p_1 + p_2) - \sum_{k=3}^n p_k \log_2 p_k \\ &= -\sum_{k=1}^n p_k \log_2 p_k. \end{aligned}$$

□

Let us then briefly prove Proposition 2.17.

Proof of Proposition 2.17. Let $q \geq 2$ be an integer. Consider the quantity

$$g(N) := H(U_N) - \frac{H(U_q) \log_2 N}{\log_2 q},$$

which we hope to prove is 0. It is easy to check that g also satisfies (C), since both H and \log_2 satisfy (C). In particular, $g(1) = g(1 \cdot 1) = g(1) + g(1) = 2g(1)$, and so $g(1) = 0$. Furthermore,

$$g(q) = 0 \quad \text{and} \quad g(q \cdot N) = g(q) + g(N) = g(N) \text{ for any } N.$$

The idea is to express $g(N)$ as a telescoping sum

$$g(N) = (g(N) - g(N-1)) + (g(N-1) - g(N-2)) + \dots + (g(2) - g(1)) + g(1) =: \sum_{j=1}^{N-1} \varepsilon_j.$$

However we can reduce the number of terms in this sum greatly using the following observation: if $N' = \lfloor \frac{N}{q} \rfloor$, so that $N = qN' + r$, where $r \leq q-1$, then

$$\begin{aligned} g(N) - g(N') &= g(N) - g(qN') \\ &= (g(N) - g(N-1)) + (g(N-1) - g(N-2)) + \dots + (g(qN'+1) - g(qN')) \\ &= \sum_{j=qN'}^{N-1} \varepsilon_j, \end{aligned}$$

where there are then at most $r < q$ terms in the sum. Hence if we define recursively $N^{(0)} = N$ and $N^{(k+1)} = (N^{(k)})' = \lfloor \frac{N^{(k)}}{q} \rfloor$, then $N^{(k)} < \frac{N}{q^k}$ and hence if $k_N = \lfloor \log_q N \rfloor$ then $N^{(k_N)} = 1$ and so

$$\begin{aligned} g(N) &= g(N) - g(N^{(k_N)}) + g(N^{(k_N)}) \\ &= g(1) + \sum_{k=0}^{k_N-1} g(N^{(k)}) - g(N^{(k+1)}) \\ &= \sum_{k=0}^{k_N-1} \sum_{j=qN^{(k)}}^{N^{(k+1)}-1} \varepsilon_j. \end{aligned} \tag{2.13}$$

However, if we consider the terms

$$\begin{aligned} \varepsilon_j &= g(j+1) - g(j) = H(U_{j+1}) - \frac{H(U_q) \log_2(j+1)}{\log_2 q} - H(U_j) + \frac{H(U_q) \log_2 j}{\log_2 q} \\ &= H(U_{j+1}) - H(U_j) - \frac{H(U_q)}{\log_2 q} \log_2 \frac{j+1}{j}, \end{aligned}$$

then by assumption $H(U_{j+1}) - H(U_j) \rightarrow 0$ as $j \rightarrow \infty$, and by continuity of the logarithm $\log_2 \frac{j+1}{j} \rightarrow 0$ as $j \rightarrow \infty$ and so $\varepsilon_j \rightarrow 0$ as $j \rightarrow \infty$.

Hence, in the expression (2.13) there are less than q terms in each inner sum, and in the outer sum there are less than $k_N = \lfloor \log_q N \rfloor$ terms, and the individual terms tend to 0, and

hence we can conclude that

$$\begin{aligned}
0 &= \lim_{N \rightarrow \infty} \frac{g(N)}{q \log_q N} \\
&= \lim_{N \rightarrow \infty} \frac{g(N)}{\log_q N} \\
&= \lim_{N \rightarrow \infty} \frac{g(N)}{\log_2 N} \\
&= \lim_{N \rightarrow \infty} \frac{H(U_N)}{\log_2 N} - \frac{H(U_q)}{\log_2 q}.
\end{aligned}$$

However, our choice of q was arbitrary. Hence, for any $q > 1$

$$\lim_{N \rightarrow \infty} \frac{H(U_N)}{\log_2 N} = \frac{H(U_q)}{\log_2 q},$$

and so $\frac{H(U_q)}{\log_2 q} = c$ is a constant. Since $\frac{H(U_2)}{\log_2 2} = 1$, we can conclude that

$$H(U_q) = \log_2 q$$

as claimed. □

2.3 Kullback-Leibler Divergence and Mutual Information

Definition 2.18. Let p and q be probability distributions on the same finite set \mathcal{X} . The *relative entropy* or *Kullback-Leibler Divergence* of p with respect to q is

$$D(p \parallel q) = \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{p(x)}{q(x)} = \mathbb{E}_X \left(\log_2 \frac{p(X)}{q(X)} \right),$$

where X is some random variable with distribution p .

Remark 2.19. Again here we need some convention to deal with the cases where the quantity $p(x) \log_2 \frac{p(x)}{q(x)}$ is not defined. If $p(x) = 0$ then we define

$$0 \log_2 \frac{0}{q(x)} = 0 \text{ for any } q(x) \geq 0,$$

and if $p(x) \neq 0, q(x) = 0$ we define

$$p(x) \log_2 \frac{p(x)}{q(x)} = \infty \text{ for any } p(x) > 0.$$

In particular, if there is any $x \in \mathcal{X}$ such that $p(x) > 0$ and $q(x) = 0$, then $D(p \parallel q) = \infty$.

This quantity is also sometimes called the *Kullback-Leibler distance*, however one should be careful that this function does not behave as we would expect a distance function to behave - in particular, it is not always *symmetric* and it does not satisfy the *triangle inequality*. In fact, it is not even obvious that this quantity is *non-negative*, although we will later show that this is the case.

Example 2.20. Let $\mathcal{X} = \{0, 1\}$, $p = (p_1, p_2)$ and $q = (q_1, q_2)$, with $p_1 + p_2 = q_1 + q_2 = 1$. Then

$$D(p \parallel q) = p_1 \log_2 \frac{p_1}{q_1} + p_2 \log_2 \frac{p_2}{q_2}.$$

For example, for $p = (1/2, 1/2)$ and $q = (1/4, 3/4)$ we can compute

$$D(p \parallel q) = \frac{1}{2} \log_2 2 + \frac{1}{2} \log_2 \frac{2}{3} = 1 - \frac{1}{2} \log_2 3$$

and

$$D(q \parallel p) = \frac{1}{4} \log_2 \frac{1}{2} + \frac{3}{4} \log_2 \frac{3}{2} = \frac{3}{4} \log_2 3 - 1.$$

Definition 2.21. Let X and Y be two jointly distributed discrete random variables taking values in finite sets \mathcal{X} and \mathcal{Y} . The *mutual information* of X and Y is defined as

$$I(X ; Y) = D(p_{X,Y} \parallel p_X \otimes p_Y)$$

where $p_X \otimes p_Y(x, y) = p_X(x)p_Y(y) = \mathbb{P}[X = x]\mathbb{P}[Y = y]$.

In other words, if we are thinking of the Kullback-Leibler divergence as a distance between distributions, the mutual information of X and Y measures how far their joint distribution is from the joint distribution of independent copies of X and Y , and so we can think of the mutual information as a measure of dependence between random variables.

Plugging Definition 2.18 into Definition 2.21 we get the following explicit formula for the mutual information

$$I(X ; Y) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_{X,Y}(x, y) \log_2 \frac{p_{X,Y}(x, y)}{p_X(x)p_Y(y)}.$$

In particular, if X and Y are independent, then the term inside the log is always 1, and so $I(X ; Y) = 0$. The larger the mutual information, the further in some sense X and Y are from being independent.

Lemma 2.22. Let X and Y be two jointly distributed discrete random variables taking finitely many values. Then

$$I(X ; Y) = H(X) + H(Y) - H(X, Y) = H(Y) - H(Y | X) = H(X) - H(X | Y) = I(Y ; X).$$

In particular, $I(X ; X) = H(X)$.

Proof. We prove the first identity, the rest follows by the chain rule $H(X, Y) = H(X) + H(Y | X) = H(Y) + H(X | Y)$.

$$\begin{aligned} I(X ; Y) &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 \frac{p(x, y)}{p(x)p(y)} \\ &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log_2 p(x, y) - p(x, y) \log_2 p(x) - p(x, y) \log_2 p(y) \\ &= -H(X, Y) - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x) + \sum_{y \in \mathcal{Y}} p(y) \log_2 p(y) \\ &= H(X) + H(Y) - H(X, Y). \end{aligned}$$

□

Example 2.23 (Secure encryption). Suppose we have a set of *messages* \mathcal{M} that we wish to encrypt and a set of *keys* \mathcal{K} that we can use to encrypt these messages. That is, every pair $m \in \mathcal{M}$ and $k \in \mathcal{K}$ of a message and a key can be used to generate some encrypted text $c \in \mathcal{C}$, or *ciphertext*.

Normally we have some (pseudo)-random method of generating keys $k \in \mathcal{K}$, which determines some random variable K on \mathcal{K} , and there is some underlying distribution M on the messages \mathcal{M} . An *encryption scheme* for M is a pair of random variables K and C , representing the key and the encrypted text such that K and C together determine M . This last condition is just saying that we can decrypt the message given the key and the ciphertext.

A *classical encryption scheme* would consist of some deterministic function $e : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{C}$ such that for each $k \in \mathcal{K}$ the function $e(\cdot, k) \rightarrow \mathcal{C}$ is injective, and then taking $C = e(M, K)$.

What does it mean for an encryption scheme to be *secure*? We want that someone who doesn't know the key cannot infer any information about the message from the ciphertext. To put this in terms of entropy, we want that there is no mutual information between C and M .

Definition 2.24 (Perfectly secure encryption scheme). An encryption scheme K, C for M is *perfectly secure* if $I(M ; C) = 0$.

There is an obvious example of a perfectly secure encryption scheme which is known as a *one-time pad*. We assume (essentially without loss of generality) that $\mathcal{M} = \{0, 1\}^n$ and that we have a uniformly distributed set of keys on the same set $\mathcal{K} = \{0, 1\}^n$. We take a classical encryption scheme where $e(m, k) = m + k$ and addition is taken in \mathbb{Z}_2^n .

Theorem 2.25. *The one-time pad is perfectly secure.*

Proof. It will be sufficient to show that M and $C = M + K$ are independent.

We first note that C is uniformly distributed on $\{0, 1\}^n$, since for any $x \in \{0, 1\}^n$

$$\begin{aligned} \mathbb{P}[C = x] &= \sum_{(y,z): y+z=x} \mathbb{P}[M = y, K = z] \\ &= \sum_y \mathbb{P}[M = y] \mathbb{P}[K = x - y] \\ &= \sum_y \mathbb{P}[M = y] 2^{-n} = 2^{-n}, \end{aligned}$$

where the second line follows since M and K are independent.

However, then

$$\begin{aligned} \mathbb{P}[M = x, C = y] &= \mathbb{P}[M = x] \mathbb{P}[C = y | M = x] \\ &= \mathbb{P}[M = x] \mathbb{P}[K = y - x] \\ &= \mathbb{P}[M = x] 2^{-n} \\ &= \mathbb{P}[M = x] \mathbb{P}[C = y]. \end{aligned}$$

It follows that $I(M ; C) = 0$. □

However this clearly isn't a very efficient method of encryption, since it requires the two parties to share a key which is as large as the message itself. However Shannon showed that this is essentially necessary for a secure encryption scheme, in the sense that, in a perfectly secure encryption scheme the set of keys must contain as least as much information as the messages.

Theorem 2.26. *If K, C is a perfectly secure encryption scheme for M then $H(K) \geq H(M)$.*

Proof. Since K, C is a perfectly secure encryption scheme for M , by assumption $I(M; C) = 0$. Furthermore, since K and C together determine M , by Lemma 2.8 $H(K, C, M) = H(K, C)$.

Hence, since $I(M; C) = H(M) + H(C) - H(M, C)$,

$$\begin{aligned} H(M) &= I(M; C) - H(C) + H(M, C) \\ &= H(M, C) - H(C) \\ &\leq H(M, C, K) - H(C) \\ &= H(K, C) - H(C) \\ &= H(K|C) \leq H(K). \end{aligned}$$

□

As a more concrete example, if both M and K are uniformly distributed then Theorem 2.26 says that

$$\log_2 |\mathcal{K}| = H(K) \geq H(M) = \log_2 |\mathcal{M}|.$$

That is, $|\mathcal{K}| \geq |\mathcal{M}|$ and so we need at least as many different keys as we have messages.

As with entropy, we can extend the concept of mutual information to conditional spaces.

Definition 2.27 (Conditional mutual information). Let X, Y, Z be three jointly distributed discrete random variables taking finitely many values. The *conditional mutual information* of X and Y given Z is defined as

$$I(X; Y | Z) = \sum_{z \in \mathcal{Z}} p_Z(z) I(X; Y | Z = z)$$

Let us briefly clarify the meaning of the above definition. Suppose $p_{X,Y,Z}(x, y, z)$ is the joint distribution of the three random variables. We can define the joint distribution of X and Y conditioned on Z as

$$p_{X,Y|Z}(x, y|z) = \frac{p_{X,Y,Z}(x, y, z)}{p_Z(z)}.$$

Then

$$I(X; Y | Z = z) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_{X,Y|Z}(x, y|z) \log_2 \frac{p_{X,Y|Z}(x, y, z)}{p_{X|Z}(x|z)p_{Y|Z}(y|z)},$$

so that

$$I(X; Y | Z) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}, z \in \mathcal{Z}} p_{X,Y,Z}(x, y, z) \log_2 \frac{p_{X,Y|Z}(x, y, z)}{p_{X|Z}(x|z)p_{Y|Z}(y|z)}.$$

Lemma 2.28. *Let X, Y, Z be three jointly distributed discrete random variables taking finitely many values. Then*

$$\begin{aligned} I(X; Y | Z) &= H(X | Z) + H(Y | Z) - H(X, Y | Z) \\ &= H(X | Z) - H(X | Y, Z) = H(Y | Z) - H(Y | X, Z). \end{aligned}$$

Proof. Exercise—as in Lemma 2.22. □

As a consequence it is easy to deduce the following variant of the chain rule for mutual information.

Theorem 2.29 (Chain rule for mutual information). *Let X_1, \dots, X_n and Y be jointly distributed discrete random variables taking finitely many values. Then*

$$I(X_1, \dots, X_n; Y) = \sum_{k=1}^n I(X_k; Y | X_{k-1}, \dots, X_1).$$

Proof. Exercise—use Theorem 2.13 and induct on n . □

So far we have only proved various equalities about entropy, just by rearranging the formulas. At various points it will be useful to be able to *estimate*, that is, bound from above or below, entropies and related quantities, and a particularly useful tool for this comes from *Jensen's inequality*. To state this inequality we will require a little background from analysis.

Definition 2.30 (Convex and concave). Let $I \subseteq \mathbb{R}$ be an open interval. A function $f: I \rightarrow \mathbb{R}$ is *convex* if for every $x, y \in I$ and $\lambda \in (0, 1)$

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y). \quad (2.14)$$

We can think of this as taking a weighted average $z = \lambda x + (1 - \lambda)y$ of the points x and y , which lies somewhere between x and y . Now f is convex if the value of the function at this point is smaller than the same weighted average of $f(x)$ and $f(y)$.

Geometrically, this asserts that the line between $f(x)$ and $f(y)$ lies above the graph of the function f between x and y .

We say f is *strictly convex* if (2.14) is strict for any $x \neq y$. Similarly we say f is *concave* or *strictly concave* if the inequality in (2.14) is reversed.

Theorem 2.31 (Jensen's inequality). *Let $f: I \rightarrow \mathbb{R}$ be a convex function on an open interval I and let X be a real random variable taking values in I . If $\mathbb{E}(X)$ and $\mathbb{E}(f(X))$ exist, then*

$$\mathbb{E}(f(X)) \geq f(\mathbb{E}(X)).$$

Furthermore, if f is strictly convex then the inequality is strict unless X is almost surely constant.

Proof. This holds for continuous random variables as well, but we will just prove it in the case where X is a discrete random variable taking finitely many values, since this is all we will use in the main part of the course.

Indeed, suppose X takes values in \mathcal{X} . If $\mathcal{X} = \{x_1, x_2\}$ and X has distribution (p_1, p_2) then $p_2 = 1 - p_1$ and so, since f is convex,

$$\mathbb{E}(f(X)) = p_1 f(x_1) + p_2 f(x_2) \geq f(p_1 x_1 + p_2 x_2) = f(\mathbb{E}(X)).$$

We can see that the general case is equivalent to the statement that for any probability distribution p on \mathcal{X} , i.e., a positive $p = (p_1, \dots, p_n)$ with $p_1 + \dots + p_n = 1$, we have

$$\sum_{i=1}^n p_i f(x_i) \geq f\left(\sum_{i=1}^n p_i x_i\right), \quad (2.15)$$

where the case $n = 2$ is the definition of convexity. Suppose (2.15) holds for some n . Given $p_1 + \dots + p_{n+1} = 1$ we can write

$$\begin{aligned} f\left(\sum_{i=1}^{n+1} p_i x_i\right) &= f\left((1 - p_{n+1}) \sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} x_i + p_{n+1} x_{n+1}\right) \\ &\leq (1 - p_{n+1}) f\left(\sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} x_i\right) + p_{n+1} f(x_{n+1}), \end{aligned} \quad (2.16)$$

where the inequality holds since f is convex. Then since

$$\sum_{i=1}^{n+1} \frac{p_i}{1 - p_{n+1}} = \frac{1 - p_{n+1}}{1 - p_{n+1}} = 1,$$

we can apply (2.15) to see that

$$f\left(\sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} x_i\right) \leq \sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} f(x_i). \quad (2.17)$$

It follows from (2.16) and (2.17) that

$$f\left(\sum_{i=1}^{n+1} p_i x_i\right) \leq (1 - p_{n+1}) \sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} f(x_i) + p_{n+1} f(x_{n+1}) = \sum_{i=1}^{n+1} p_i f(x_i),$$

and so (2.15) holds for all n by induction.

Finally, if f is strictly convex, we may assume by induction that equality holds in (2.17) only in the case that $x_1 = \dots = x_n$, and then equality holds in (2.16) only when $\sum_{i=1}^n \frac{p_i}{1 - p_{n+1}} x_i = x_{n+1}$, but since $\sum_{i=1}^n p_i = 1 - p_{n+1}$ it follows that $x_1 = \dots = x_n = x_{n+1}$. \square

Probably the most important application of Jensen's inequality in information theory is the following:

Theorem 2.32 (Information Inequality). *Let p and q be distributions on a finite set \mathcal{X} . Then $D(p \parallel q) \geq 0$ with equality if and only if $p = q$.*

Proof. We recall the definition of $D(p \parallel q)$,

$$\begin{aligned} D(p \parallel q) &= \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{p(x)}{q(x)} \\ &= \mathbb{E}_Y \left(\log_2 \frac{p(Y)}{q(Y)} \right), \end{aligned}$$

where Y is some random variable with distribution p (where we are taking the expected value only over x such that $p(x) > 0$).

Now, $\log_2 x$ is a concave function, not a convex function, but we could consider instead $-\log_2 x$ which is then strictly convex on $(0, \infty)$ (this can be proved by considering the second derivative on $(0, \infty)$).

Hence, by Jensen's inequality, applied with f mapping $x \mapsto \log_2 x$ and the random variable $X = \frac{q(Y)}{p(Y)}$ we see that

$$\begin{aligned} D(p \parallel q) &= \mathbb{E}_Y \left(-\log_2 \frac{q(Y)}{p(Y)} \right) \\ &\geq -\log_2 \mathbb{E}_Y \left(\frac{q(Y)}{p(Y)} \right) \\ &= -\log_2 \left(\sum_{x: p(x) > 0} p(x) \frac{q(x)}{p(x)} \right) \\ &= -\log_2 \left(\sum_{x: p(x) > 0} q(x) \right) \\ &\geq 0. \end{aligned}$$

Suppose then that we obtain equality through this argument. In particular, $\sum_{x: p(x) > 0} q(x) = 1$ and so p and q have the same support. Also, we have equality in the application of Jensen's inequality, and so there is some constant c such that, whenever $p(x) > 0$, $\frac{p(x)}{q(x)} = c$. However, since $\sum_{x: p(x) > 0} q(x) = 1 = \sum_{x: p(x) > 0} p(x)$, it follows that $c = 1$ and so $p = q$. \square

Let us note some immediate, and incredibly useful, corollaries of this theorem.

Corollary 2.33. *Let X, Y, Z and X_1, \dots, X_n be jointly distributed discrete random variables taking values in a finite set.*

1. $I(X ; Y) \geq 0$, with equality if and only if X and Y are independent,
2. $H(X | Y) \leq H(X)$, with equality if and only if X and Y are independent,
3. $I(X ; Y | Z) \geq 0$, with equality if and only if X and Y are independent conditional upon Z ,
4. $H(X_1, \dots, X_n) \leq H(X_1) + \dots + H(X_n)$ with equality if and only if the X_k are mutually independent.

Proof. The first is clear since $I(X ; Y) = D(p_{X,Y} \parallel p_X \otimes p_Y) \geq 0$, and equality holds when the two distributions are equal—i.e., when $p(x, y) = p(x)p(y)$ for all x, y .

Now the second follows from the equality

$$I(X ; Y) = H(X) - H(X | Y).$$

For the third we note that

$$I(X ; Y | Z) = \sum_{z \in \mathcal{Z}} p(z) I(X ; Y | Z = z).$$

If the left hand side is 0, and each term on the right hand side is positive, they must also be 0, and so $I(X ; Y | Z = z) = 0$ whenever $p(z) > 0$. In particular, by the first part

$$p(x, y | z) = p(x | z) p(y | z)$$

whenever $p(z) > 0$ (which is the definition of conditional independence).

The last one can be proved by induction. The case $n = 2$ follows from the second, since $H(X, Y) = H(Y) + H(X | Y) \leq H(Y) + H(X)$ with equality if and only if X and Y are independent. For the general case we note that

$$H(X_1, \dots, X_n) \leq H(X_n) + H(X_1, \dots, X_{n-1} | X_n) \leq H(X_1) + H(X_2) + \dots + H(X_n)$$

If we have equality in the first line, then by the $n = 2$ case we have that $p_{X_1, \dots, X_n} = p_{X_1, \dots, X_{n-1}} p_{X_n}$ and if we have equality in the second line then by the induction hypothesis $p_{X_1, \dots, X_{n-1}} = p_{X_1} p_{X_2} \dots p_{X_{n-1}}$ and hence if equality holds altogether we have

$$p_{X_1, \dots, X_n} = p_{X_1} p_{X_2} \dots p_{X_n},$$

i.e., the joint distribution is the product of the marginals, and the random variables are mutually independent. \square

Another important consequence of Theorem 2.32 is the following.

Lemma 2.34 (Log sum inequality). *Let $a_1, b_1, \dots, a_n, b_n \geq 0$ and let $a = \sum_{k=1}^n a_k$ and $b = \sum_{k=1}^n b_k$. Then*

$$\sum_{k=1}^n a_k \log_2 \frac{a_k}{b_k} \geq a \log_2 \frac{a}{b}.$$

Proof. We may assume that $a_i, b_i > 0$ for all i (exercise). Let us define a function

$$f(t) = \begin{cases} t \log_2 t & \text{if } t > 0, \\ 0 & \text{if } t = 0. \end{cases}$$

It can be checked (exercise) that f is strictly convex on $[0, \infty)$.

We define a probability distribution p by letting $p_i = \frac{b_i}{b}$ and set $t_i = \frac{a_i}{b_i}$. Since f is convex, by Jensen's inequality

$$\begin{aligned} \sum_{k=1}^n \frac{a_i}{b} \log_2 \frac{a_i}{b_i} &= \sum_{k=1}^n p_i f(t_i) \\ &\leq f\left(\sum_{k=1}^n p_i t_i\right) \\ &= f\left(\sum_{k=1}^n \frac{a_i}{b}\right) \\ &= f\left(\frac{a}{b}\right) \\ &= \frac{a}{b} \log_2 \frac{a}{b}. \end{aligned}$$

Dividing both sides by b we obtain the inequality.

One can check here that we get equality if and only if the t_i s are constant, i.e., if a_i and b_i are always in the same ratio. \square

As a corollary we get a weird looking statement asserting a sort of multivariable concavity of the Kullback-Leibler divergence.

Corollary 2.35. *Let $p^{(1)}, p^{(2)}, q^{(1)}$ and $q^{(2)}$ be probability distributions on the same finite set \mathcal{X} and let $\lambda \in (0, 1)$. We have*

$$D(\lambda \cdot p^{(1)} + (1 - \lambda) \cdot p^{(2)} \parallel \lambda \cdot q^{(1)} + (1 - \lambda) \cdot q^{(2)}) \leq \lambda D(p^{(1)} \parallel q^{(1)}) + (1 - \lambda) D(p^{(2)} \parallel q^{(2)}).$$

Proof. Let us define probability distributions

$$p = \lambda \cdot p^{(1)} + (1 - \lambda) \cdot p^{(2)} \quad \text{and} \quad q = \lambda \cdot q^{(1)} + (1 - \lambda) \cdot q^{(2)}.$$

For all $x \in \mathcal{X}$ consider the following quantities

$$a_1 = \lambda p^{(1)}(x), a_2 = (1 - \lambda) p^{(2)}(x), b_1 = \lambda q^{(1)}(x), b_2 = (1 - \lambda) q^{(2)}(x),$$

which are all positive. We apply the Log sum inequality to deduce that

$$\lambda p^{(1)}(x) \log_2 \frac{p^{(1)}(x)}{q^{(1)}(x)} + (1 - \lambda) p^{(2)}(x) \log_2 \frac{p^{(2)}(x)}{q^{(2)}(x)} \geq p(x) \log_2 \frac{p(x)}{q(x)}.$$

Summing this inequality over all $x \in \mathcal{X}$ leads to

$$\lambda D(p^{(1)} \parallel q^{(1)}) + (1 - \lambda) D(p^{(2)} \parallel q^{(2)}) \geq D(p \parallel q),$$

as claimed. \square

As a simple corollary we find that the entropy function is also concave.

Corollary 2.36. *Let p, q be probability distributions on a finite set \mathcal{X} and let $\lambda \in (0, 1)$. Then*

$$H(\lambda p + (1 - \lambda) q) \geq \lambda H(p) + (1 - \lambda) H(q).$$

Proof. Let u be the uniform distribution on \mathcal{X} . We note that

$$D(p \parallel u) = \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{p(x)}{u(x)} = \sum_{x \in \mathcal{X}} p(x) \log_2 |\mathcal{X}| \cdot p(x) = \log_2 |\mathcal{X}| - H(p),$$

and similarly for q . Applying Corollary 2.35 with $p^{(1)} = p, p^{(2)} = q$ and $q^{(1)} = q^{(2)} = u$ we conclude that

$$D(\lambda p + (1 - \lambda) q \parallel u) \leq \lambda D(p \parallel u) + (1 - \lambda) D(q \parallel u)$$

and hence

$$H(\lambda p + (1 - \lambda) q) \geq \lambda H(p) + (1 - \lambda) H(q).$$

\square

Remark 2.37. From the inequality

$$D(p \parallel u) = \log_2 |\mathcal{X}| - H(p) \geq 0,$$

we can conclude from Theorem 2.32 that $H(p) \leq \log_2 |\mathcal{X}|$ with equality if and only if p is the uniform distribution.

Suppose we are given the conditional distribution of a random variable Y with respect to a random variable X , but not the distribution of X or Y . That is, we have the function

$$p_{Y|X}(y, x) = p(y|x),$$

where for each x we can think of the function $p(\cdot|x)$ as a distribution on \mathcal{Y} .

Then, any probability distribution p_X on \mathcal{X} gives rise to a joint distribution $p_{X,Y}$ via

$$p_{X,Y}(x, y) = p_X(x)p_{Y|X}(y, x),$$

from which we can also derive the distribution $p(y) = \mathbb{E}_{p_X} p(y|X) = \sum_{x \in \mathcal{X}} p(x)p(y|x)$.

In other words, if we fix ahead of time the conditional distribution, then the joint distribution (X, Y) , and so in particular the distribution of Y , is determined by the distribution of X . We can think of this as the evolution of a random process in time, where X and Y represent the outcome at two specific times. In order to calculate the probability that we saw the outcome (x, y) we need to know the probability that at the first time we had the outcome x , and then the outcome that the process developed from x to become y at the second time.

We can think of this in terms of a *transition matrix* P , whose columns are indexed by \mathcal{X} and rows by \mathcal{Y} and whose entry $P_{xy} = p(y|x)$ is the probability that we observe y after observing x , the probability that our process evolves from x to become y . The rows of this matrix $p(\cdot|x)$ correspond to the conditional distribution of Y given that $X = x$, and so given a distribution p_X on \mathcal{X} we can compute the marginal distribution on \mathcal{Y} as $p_Y = p_X P$ and the joint distribution (as an $\mathcal{X} \times \mathcal{Y}$ matrix) can be seen to be the product $\text{diag}(p_X) \cdot P$ of a diagonal matrix with entries $p_X(x)$ together with P , so that the x th row is $p_X(x)p(\cdot|x) = p_{X,Y}(x, \cdot)$.

Theorem 2.38. Suppose $p_{Y|X}$ is the conditional distribution of some discrete random variable Y taking values in a finite set \mathcal{Y} with respect to some unknown discrete random variable X taking values in a finite set \mathcal{X} . Then the function $f(p_X) = I(X; Y) = D(p_{X,Y} \parallel p_X \otimes p_Y)$ is concave, that is, for any two distributions p and q on \mathcal{X} and $\lambda \in (0, 1)$

$$f(\lambda \cdot p + (1 - \lambda) \cdot q) \geq \lambda f(p) + (1 - \lambda) f(q).$$

Conversely, if p_X is known, then the function $F(p_{Y|X}) = I(X; Y)$ is convex.

Proof. Let us start by proving the first statement. We first note that

$$f(p_X) = I(X; Y) = H(Y) - H(Y | X) = H(p_Y) - \sum_{x \in \mathcal{X}} p_X(x) H(p(\cdot|x)).$$

Suppose $p_X^{(1)}$ and $p_X^{(2)}$ are probability distributions on \mathcal{X} and $\lambda \in (0, 1)$. Let

$$p_X = \lambda p_X^{(1)} + (1 - \lambda) p_X^{(2)}.$$

Each $p_X^{(i)}$ leads to a marginal distribution on \mathcal{Y} , denoted $p_Y^{(i)}$. By the linearity of matrix multiplication, we see that

$$p_Y = \lambda p_Y^{(1)} + (1 - \lambda) p_Y^{(2)}.$$

Hence we can calculate

$$\begin{aligned} f(p_X) &= H(p_Y) - \sum_{x \in \mathcal{X}} p_X(x) H(p(\cdot|x)) \\ &= H\left(\lambda p_Y^{(1)} + (1 - \lambda) p_Y^{(2)}\right) - \sum_{x \in \mathcal{X}} \left(\lambda p_X^{(1)}(x) + (1 - \lambda) p_X^{(2)}(x)\right) H(p(\cdot|x)) \\ &\geq \lambda H\left(p_Y^{(1)}\right) + (1 - \lambda) H\left(p_Y^{(2)}\right) - \lambda \sum_{x \in \mathcal{X}} p_X^{(1)}(x) H(p(\cdot|x)) - (1 - \lambda) \sum_{x \in \mathcal{X}} p_X^{(2)}(x) H(p(\cdot|x)) \\ &= \lambda f\left(p_X^{(1)}\right) + \lambda f\left(p_X^{(2)}\right), \end{aligned}$$

where we used that the entropy function is concave.

The proof of the second part follows the same idea. Given two conditional distributions $p_{Y|X}^{(1)}$ and $p_{Y|X}^{(2)}$, if we take a convex combination $p_{Y|X} = \lambda p_{Y|X}^{(1)} + (1 - \lambda) p_{Y|X}^{(2)}$, then again by linearity of matrix multiplication the corresponding joint distributions satisfy the relation

$$p_{X,Y} = \lambda p_{X,Y}^{(1)} + (1 - \lambda) p_{X,Y}^{(2)},$$

and a similar relation holds for the product of p_X with the marginal distributions

$$p_X \otimes p_Y = \lambda p_X \otimes p_{Y|X}^{(1)} + (1 - \lambda) p_X \otimes p_{Y|X}^{(2)}.$$

Hence, we can write

$$F(p_{Y|x}) = D(p_{X,Y} \parallel p_X \otimes p_Y) = D(\lambda p_{X,Y}^{(1)} + (1 - \lambda) p_{X,Y}^{(2)} \parallel \lambda p_X \otimes p_{Y|X}^{(1)} + (1 - \lambda) p_X \otimes p_{Y|X}^{(2)}),$$

and it follows from Corollary 2.35 that the function F is convex. \square

Definition 2.39. Let X, Y and Z be jointly distributed discrete random variables taking values in a finite set. The triple (X, Y, Z) is called a *Markov(ian) triple*, which we write as $X \rightarrow Y \rightarrow Z$, if for all x, y with $\mathbb{P}[X = x \mid Y = y] > 0$

$$\mathbb{P}[Z = z \mid X = x, Y = y] = \mathbb{P}[Z = z \mid Y = y]$$

If we think of $X \rightarrow Y \rightarrow Z$ as the evolution of some random process over time, then this says that if we know the “present”, the value of Y , then the future evolution does not depend on the past.

For a Markov triple we can write the joint distribution as the product

$$p_{X,Y,Z}(x, y, z) = p_X(x) \cdot p_{Y|X}(y \mid x) \cdot p_{Z|X,Y}(z \mid x, y) = p_X(x) \cdot p_{Y|X}(y \mid x) \cdot p_{Z|Y}(z \mid y). \quad (2.18)$$

There is another nice equivalent description in terms of the conditional distributions.

Lemma 2.40. *(X, Y, Z) is a Markovian triple if and only if X and Z are conditionally independent, given Y . That is, if whenever $p_Y(y) > 0$,*

$$p_{X,Z|Y}(x, z \mid y) = p_{X|Y}(x \mid y) p_{Z|Y}(z \mid y).$$

Proof. Using the definition of conditional probability

$$p_{X,Y,Z}(x, y, z) = p_{X,Z|Y}(x, z | y)p_Y(y) \text{ and } p_{X|Y}(x | y) = \frac{p_{X,Y}(x, y)}{p_Y(y)} = \frac{p_X(x)p_{Y|X}(y | x)}{p_Y(y)}.$$

Hence, it follows from (2.18) that

$$\begin{aligned} p_{X,Z|Y}(x, z | y) &= \frac{p_{X,Y,Z}(x, y, z)}{p_Y(y)} \\ &= \frac{p_X(x) \cdot p_{Y|X}(y | x) \cdot p_{Z|Y}(z | y)}{p_Y(y)} \\ &= p_{X|Y}(x | y)p_{Z|Y}(z | y). \end{aligned}$$

□

Note that the condition in Lemma 2.40 is symmetric in X and Z . In other words, we see that $X \rightarrow Y \rightarrow Z$ is a Markovian triple if and only if $Z \rightarrow Y \rightarrow X$ is as well.

If we think about our random process as some method of processing some data, our input data X is encoded or transmitted say, and we have as output data Y . Then, if we forget about our original data X , there should be no way to increase the amount of information about X that the output Y contains by further processing (via some deterministic, or random, process).

As an example, we can think about the transmission of some message X from Alice to Bob, who receives the transmitted message Y (perhaps some random noise has been added in the channel). Without access to the message X , there should be no way that Bob can deduce more information about X than what is contained in Y .

Theorem 2.41 (Data processing inequality). *If $X \rightarrow Y \rightarrow Z$, then $I(X ; Z) \leq I(X ; Y)$.*

Proof. This follows quickly from the chain rule for mutual information (Theorem 2.29). Namely, we know that

$$I(X ; Y, Z) = I(X ; Z) + I(X ; Y | Z) = I(X ; Y) + I(X ; Z | Y).$$

However, we also know that X and Z are conditionally independent given Y , and hence

$$I(X ; Z | Y) = 0,$$

and also mutual information is non-negative (Corollary 2.33) and so $I(X ; Y | Z) \geq 0$ and hence

$$I(X ; Z) \leq I(X ; Z) + I(X ; Y | Z) = I(X ; Y) + I(X ; Z | Y) = I(X ; Y).$$

□

Note that, since $Z \rightarrow Y \rightarrow X$ is also a Markovian triple, we can also deduce from Theorem 2.41 that $I(Z ; Y) \geq I(Z ; X) = I(X ; Z)$, and so

$$I(X ; Z) \leq \max\{I(X ; Y), I(Y ; Z)\}.$$

Suppose, in the previous example, Alice transmits a message X to Bob, who receives Y , and Bob wishes to reconstruct the message X from Y , via some process. Perhaps Alice is sending pictures of some letters, and Bob receives slightly perturbed pictures, and so has to guess which letter fits the picture best. This might be some deterministic process, or maybe when Bob is unsure he makes some random choice, weighted by how likely he thinks each letter is. In this way Bob makes a (potentially random) guess $Z = \hat{X}$ based on Y , and we have a Markov triple $X \rightarrow Y \rightarrow Z$.

How accurate can Bob be? There are many ways we could measure this, but one way would be to look at the probability that Bob's guess is correct, the probability that $\hat{X} = X$.

Now, from the Data processing inequality, we know that if lots of information is lost in transmission, and so $I(X ; Y)$ is small, it shouldn't be the case that \hat{X} is well-correlated with X , since $I(X ; \hat{X})$, which is a measure of dependence, is also small. So, we should expect to be able to bound the probability of success, as some function of the mutual information $I(X ; Y)$, and the following inequality makes this precise (in fact, we bound instead the probability of failure, as a function of the conditional entropy $H(X | Y)$).

Theorem 2.42 (Fano's inequality). *Let $X \rightarrow Y \rightarrow \hat{X}$ be a Markov triple, where we think of \hat{X} as an estimation of X on the basis of Y . Define $p_{\text{err}} = \mathbb{P}[\hat{X} \neq X]$. Then*

$$H(p_{\text{err}}, 1 - p_{\text{err}}) + p_{\text{err}} \log_2 |\mathcal{X}| \geq H(X | \hat{X}) \geq H(X | Y).$$

In particular

$$p_{\text{err}} \geq \frac{H(X | \hat{X}) - 1}{\log_2 |\mathcal{X}|} \geq \frac{H(X | Y) - 1}{\log_2 |\mathcal{X}|}.$$

Proof. We first note that, since $I(X ; Y) = H(X) - H(X | Y)$, the data processing inequality implies that

$$H(X) - H(X | Y) \geq H(X) - H(X | \hat{X}) \quad \text{and so} \quad H(X | \hat{X}) \geq H(X | Y).$$

Let us define $E = \mathbb{1}_{[\hat{X} \neq X]}$, so that $\mathbb{P}[E = 1] = p_{\text{err}}$ and $\mathbb{P}[E = 0] = 1 - p_{\text{err}}$, and so $H(E) = H(p_{\text{err}}, 1 - p_{\text{err}})$.

We now apply the conditional chain rule (Lemma 2.14) in two ways to get the equality

$$H(E, X | \hat{X}) = H(X | \hat{X}) + H(E | X, \hat{X}) = H(E | \hat{X}) + H(X | E, \hat{X}).$$

Now, since E is determined by X and \hat{X} , Lemma 2.8 implies that $H(E | X, \hat{X}) = 0$ and so, since conditioning reduces entropy

$$H(X | \hat{X}) = H(X | \hat{X}) + H(E | X, \hat{X}) = H(E | \hat{X}) + H(X | E, \hat{X}) \leq H(E) + H(X | E, \hat{X}). \quad (2.19)$$

Then, by the definition of conditional entropy

$$H(X | E, \hat{X}) = \mathbb{P}[E = 1]H(X | \hat{X}, E = 1) + \mathbb{P}[E = 0]H(X | \hat{X}, E = 0).$$

However, under the event $E = 0$, \hat{X} and X coincide, and so X determines \hat{X} and $H(X | \hat{X}, E = 0) = 0$. Finally, since entropy is maximised by the uniform distribution (Remark 2.37)

$$H(X | E, \hat{X}) = p_{\text{err}}H(X | \hat{X}, E = 1) \leq p_{\text{err}} \log_2 |\mathcal{X}|. \quad (2.20)$$

Combining (2.19) and (2.20) we obtain

$$H(X | \hat{X}) \leq H(E) + H(X | E, \hat{X}) \leq H(p_{\text{err}}, 1 - p_{\text{err}}) + p_{\text{err}} \log_2 |\mathcal{X}|.$$

Finally, since $H(p_{\text{err}}, 1 - p_{\text{err}}) \leq 1$, it follows that

$$H(X | Y) \leq H(X | \hat{X}) \leq 1 + p_{\text{err}} \log_2 |\mathcal{X}|,$$

which rearranges to the second inequality. □

It is reasonable to ask why we mention also the weaker bounds in terms of $H(X | Y)$ rather than $H(X | \hat{X})$. This is useful if we are interested in an *a priori estimate*, one that only depends on the message X and the transmitted message Y , and not the method of reconstruction \hat{X} . This bound then holds for *every* possible \hat{X} - no matter how Bob attempts to guess the message X , he must always have a failure probability of at least this quantity.

Remark 2.43. *All the material in this section extends straightforwardly to arbitrary discrete random variables, respectively probability distributions, taking values in countable sets.*

That is, given a random variable X taking values in $\mathcal{X} = \{x_k : x \in \mathbb{N}\}$ with distribution $p = (p_1, p_2, \dots)$, we can define the entropy

$$H(X) = H(p) = - \sum_{k=1}^{\infty} p_k \log_2 p_k,$$

which may also take the value $+\infty$.

3 Entropy Rate and Asymptotic Equipartition

3.1 Entropy Rate

Definition 3.1 (Stochastic process, state space). A *stochastic process in discrete time* is a sequence $(X_n)_{n \in \mathbb{N}}$ of jointly distributed random variables. The *state space* of the process is the set \mathcal{X} of possible values which the X_n can take.

Example 3.2. Pick a random page of a book and let X_n be the n th letter on the page.

Let $X_1 = 100$ be constant, and let X_n be the bankroll after n spins of a roulette wheel of a gambler who bets his entire stake on red each time.

Let $X_1 = (0, 0) \in \mathbb{Z}^2$ and let X_n be the position after n steps of a ‘random walk’, where in each time step we choose uniformly at random one of the four neighbours in the grid of our current position and move there.

Given a stochastic process $(X_n)_{n \in \mathbb{N}}$ we can think of $H(X_1, \dots, X_n)$ as the total amount of information in the system at time n . This is clearly an increasing function of n . What we are interested in is the rate at which this function is increasing.

Definition 3.3 (Entropy rate). If $(X_n)_{n \in \mathbb{N}}$ is a stochastic process whose state space is finite, then the *entropy rate* or *asymptotic entropy* of the stochastic process is defined as

$$h := \lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, \dots, X_n),$$

if the limit exists. The unit of h is *bits per time unit*.

We will see that the entropy rate represents a theoretical limit on how efficiently we can encode the data stream $(X_n)_{n \in \mathbb{N}}$. Conversely, we will find that for a broad class of processes, we can achieve this theoretical limit asymptotically, using the idea of asymptotic equipartition.

We can think of h as the average amount of new information introduced in each step of the stochastic process. Indeed, by the chain rule (Theorem 2.13)

$$\frac{1}{n} H(X_1, \dots, X_n) = \frac{1}{n} \sum_{k=1}^n H(X_k \mid X_1, \dots, X_{k-1}), \quad (3.1)$$

where $H(X_k \mid X_1, \dots, X_{k-1})$ is the amount of information introduced at the k th step.

What we will find is that the entropy rate represents a theoretical limit on how efficiently we can encode the data stream $(X_n)_{n \in \mathbb{N}}$. Conversely, the idea of asymptotic equipartition is that for a broad class of processes, we can achieve this theoretical limit asymptotically.

Lemma 3.4. *If $(X_n)_{n \in \mathbb{N}}$ is a stochastic process whose state space is finite and the limit $h' = \lim_{k \rightarrow \infty} \frac{1}{k} H(X_k \mid X_1, \dots, X_{k-1})$ exists, then h exists and $h = h'$.*

Proof. This follows from the following basic analytic fact: if $a_n \rightarrow a$ as $n \rightarrow \infty$, then $\frac{1}{n} \sum_{k=1}^n a_k \rightarrow a$, whose proof we leave as an exercise.

Then, letting $a_k = H(X_k | X_1, \dots, X_{k-1})$, our assumption is that $\lim_{k \rightarrow \infty} a_k = h'$, and hence by (3.1)

$$h = \lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, \dots, X_n) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n a_k = h'.$$

□

For i.i.d. sequences, it is trivial to compute h using Lemma 3.4.

Lemma 3.5. *If $(X_n)_{n \in \mathbb{N}}$ is a sequence of i.i.d. discrete random variables taking values in a finite set, then the entropy rate h exists and is equal to $H(X_1)$.*

Proof. Indeed, since the X_n are independent and identically distributed,

$$H(X_k | X_1, \dots, X_{k-1}) = H(X_k) = H(X_1)$$

for all k and hence $\lim_{k \rightarrow \infty} H(X_k | X_1, \dots, X_{k-1})$ exists and is equal to $H(X_1)$. Hence, by Lemma 3.4, $h = H(X_1)$. □

Definition 3.6. A stochastic process $(X_n)_{n \in \mathbb{N}}$ is *stationary* if for every $\ell, k \in \mathbb{N}$ the two random vectors

$$(X_1, \dots, X_\ell) \quad \text{and} \quad (X_{k+1}, \dots, X_{k+\ell})$$

have the same distribution. In other words, for every choice of elements $x_1, \dots, x_\ell \in \mathcal{X}$ in the state space,

$$\mathbb{P}[X_1 = x_1, \dots, X_\ell = x_\ell] = \mathbb{P}[X_{k+1} = x_1, \dots, X_{k+\ell} = x_\ell].$$

Example 3.7. The simplest example of stationary processes are sequences of independent and identically distributed random variables. For example, if we repeatedly roll a dice and let X_n be the value of the n th roll.

As another example, suppose we have two biased coins with different probabilities p and q of heads, and I choose randomly, say with probability $\frac{1}{2}$, one of the coins to flip, and then let X_n be the (bit) value of the n th coin toss. Then the X_n are not independent: if p is very close to 1 and q is very close to 0, then if $X_1 = 1$, it's very likely that I picked the first coin and so very likely that $X_2 = 1$ as well. However, it is easy to show that this process is stationary.

The random walk on \mathbb{Z}^2 from the previous example is not stationary— X_0 is deterministic, whereas X_1 is uniformly distributed on $\{(\pm 1, 0), (0, \pm 1)\}$.

Lemma 3.8. *If $(X_n)_{n \in \mathbb{N}}$ is a stationary process with a finite state space, then the entropy rate h exists.*

Proof. We note that (X_k, \dots, X_2) has the same joint distribution as (X_{k-1}, \dots, X_1) . In particular,

$$H(X_k | X_1, \dots, X_{k-1}) \leq H(X_k | X_2, \dots, X_{k-1}) = H(X_{k-1} | X_1, \dots, X_{k-2}).$$

Hence the sequence $a_k = H(X_k | X_1, \dots, X_{k-1})$ is a decreasing sequence of positive reals, and therefore the limit $\lim_{k \rightarrow \infty} a_k = h'$ exists, and so by Lemma 3.4, the entropy rate h exists and $h = h'$. □

Whilst Lemma 3.8 asserts the existence of the entropy rate for stationary processes, it is non-constructive - it does not provide us a formula to calculate h . It is reasonable to ask if there is a broader class of stochastic processes (than i.i.d.) for which we can compute h explicitly.

3.2 Time-homogeneous Markov chains

Definition 3.9. A stochastic process $(X_n)_{n \geq 0}$ with finite state space \mathcal{X} is a *Markov chain* (MC) if for all $n \in \mathbb{N}$ and for all $x_0, \dots, x_n \in \mathcal{X}$,

$$\mathbb{P}[X_n = x_n \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}] = \mathbb{P}[X_n = x_n \mid X_{n-1} = x_{n-1}] := p_n(x_n | x_{n-1}),$$

whenever $\mathbb{P}[X_1 = x_0, \dots, X_{n-1} = x_{n-1}] > 0$.

A Markov chain is *time-homogeneous* if $p_n(y|x) := p(y|x)$ does not depend on n . In this case the matrix $P = \left(p(y|x) \right)_{x,y \in \mathcal{X}}$ is the *transition matrix* of the time-homogeneous Markov chain and the distribution of X_0

$$\nu(x) = \mathbb{P}[X_0 = x]$$

is the *initial distribution* or *starting distribution*.

We note that the transition matrix P is always a *stochastic* matrix—the entries are all non-negative and each row sums to 1, that is for all $x \in \mathcal{X}$

$$\sum_{y \in \mathcal{X}} p(y|x) = 1.$$

We can think of a Markov chain as a *memoryless* stochastic process—given the state of the process at some time n , the future distribution does not depend on the past. In particular, it is easy to check that each consecutive triple is Markovian and so

$$X_0 \rightarrow X_1 \rightarrow \dots \rightarrow X_n.$$

Example 3.10. Suppose we are playing some board game with a number of possible states \mathcal{X} . Each turn we roll a dice and play according to some fixed strategy, so that the probability that we move from a state x to a state y in any particular turn is fixed. The state X_n of some player is then a time-homogeneous Markov chain.

A random walk is also an example of a time-homogeneous Markov chain—if we are currently at a vertex x the probability that we move to a vertex y only depends on the current state, and not the history of the walk.

Lemma 3.11. *If $(X_n)_{n \geq 0}$ is a Markov chain, then for any $k \in \mathbb{N}$ $((X_n, X_{n+1}, \dots, X_{n+k}))_{n \geq 0}$ is a Markov chain.*

Proof. Exercise. □

Definition 3.12 (The (di)graph of a Markov chain). Given a time-homogeneous Markov chain $(X_n)_{n \geq 0}$ with state space \mathcal{X} , we can draw an associated (weighted) (di)graph whose vertex set is \mathcal{X} and for any two states x and y we draw an arc from x to y with weight $p(y|x)$ if $p(y|x) > 0$.

We can think of the Markov chain as a simple random walk on this graph, where the probability of moving from state x to y is given by the weight of the arc from x to y and the distribution of the starting vertex is given by X_0 .

Example 3.13. We can think of the following simplified model of the evolution of the weather. Our stochastic process has three states $\mathcal{X} = \{\text{sun, rain, snow}\} = \{N, R, S\}$

Our transition matrix is given as follows:

$$P = \begin{array}{c|ccc} & N & R & S \\ \hline N & 0 & 1/2 & 1/2 \\ R & 1/4 & 1/2 & 1/4 \\ S & 1/4 & 1/4 & 1/2 \end{array}$$

So, we never have two sunny days in a row—if a day is sunny then the next day is equally likely to be rainy or snowy. On rainy or snowy days the next day has probability $1/2$ to have the same weather, and probability $1/2$ to change to one of the other options uniformly.

In this case the digraph of this Markov chain has vertex set $V = \{N, R, S\}$ and arcs

$$\begin{aligned} e_1 &= (N, R), e_2 = (N, S), e_3 = (R, N), e_4 = (R, R), \\ e_5 &= (R, S), e_6 = (S, N), e_7 = (S, R), e_8 = (S, S). \end{aligned}$$

with weights

$$\begin{aligned} w(e_1) &= 1/2, w(e_2) = 1/2, w(e_3) = 1/4, w(e_4) = 1/2, \\ w(e_5) &= 1/4, w(e_6) = 1/4, w(e_7) = 1/4, w(e_8) = 1/2. \end{aligned}$$

By the Markovian property it is relatively easy to write down the joint distribution of (X_0, \dots, X_n) as

$$\mathbb{P}[X_0 = x_0, X_1 = x_1, \dots, X_n = x_n] = \nu(x_0)p(x_1|x_0)p(x_2|x_1) \cdots p(x_n|x_{n-1}), \quad (3.2)$$

and from this it is also clear what the conditional distribution of X_n , given X_0 , is.

Lemma 3.14. *Let $(X_n)_{n \geq 0}$ be a Markov chain with initial distribution ν and transition matrix P . Then*

$$p^{(n)}(y|x) := \mathbb{P}[X_n = y \mid X_0 = x] = (P^n)_{xy},$$

that is, the matrix given by $\left(p^{(n)}(y|x)\right)_{x,y \in \mathcal{X}}$ is the n th power of P .

If we consider ν as a row vector, then $p_{X_n} = \nu P^n$, that is

$$p_{X_n}(y) = \sum_{x \in \mathcal{X}} \nu(x) p^{(n)}(y|x).$$

Proof. The first part follows from (3.2) and the definition of matrix multiplication, noting that

$$\begin{aligned} \mathbb{P}[X_n = y \mid X_0 = x] &= \sum_{x_1, \dots, x_{n-1} \in \mathcal{X}} \mathbb{P}[X_1 = x_1, \dots, X_n = y \mid X_0 = x] \\ &= \sum_{x_1, \dots, x_{n-1} \in \mathcal{X}} p(x_1|x) p(x_2|x_1) \cdots p(y|x_{n-1}) \\ &= \sum_{x_1, \dots, x_{n-1} \in \mathcal{X}} (P)_{xx_1} (P)_{x_1x_2} \cdots (P)_{x_{n-1}y} \\ &= (P^n)_{xy}. \end{aligned}$$

It is also easy to prove via induction (exercise).

The second part is just the law of conditional probability

$$\mathbb{P}[X_n = y] = \sum_{x \in \mathbb{X}} \mathbb{P}[X_n = y \mid X_0 = x] \cdot \mathbb{P}[X_0 = x] = \sum_{x \in \mathcal{X}} \nu(x) p^{(n)}(y|x).$$

□

We will show that for a natural class of time-homogeneous Markov chains the entropy rate exists, and can be easily calculated, and furthermore is independent of the choice of the initial distribution ν .

The existence of the entropy rate would be clear if the Markov chain were stationary, by Lemma 3.8. However, whilst it is easy to verify that every stationary Markov chain is time-homogeneous (exercise), the converse is not true in general, but will hold for sensible choices of initial distribution.

Lemma 3.15. *Let $(X_n)_{n \geq 0}$ be a time-homogeneous Markov chain with initial distribution ν and transition matrix P . Then the Markov chain is stationary if and only if $\nu P = \nu$, that is, only if ν is an eigenvector of P with eigenvalue 1.*

Proof. Clearly time-homogeneity implies that $p_{X_1} = \nu P = \nu = p_{X_0}$, and so this condition is obviously necessary.

Conversely, if $\nu P = \nu$, given $k, \ell \in \mathbb{N}$, by Lemma 3.14 the distribution of X_k is given by $\nu P^k = \nu$. It is relatively easy to check that the Markov property then implies that $(X_k, X_{k+1}, \dots, X_\ell)$ has the same distribution as $(X_0, X_1, \dots, X_{\ell-1})$. □

In this case we call ν a *stationary distribution* for P , or for the Markov chain.

Lemma 3.16. *Let $(X_n)_{n \geq 0}$ be a time-homogeneous Markov chain with a stationary initial distribution ν . Then the entropy rate exists and is given by*

$$h = \sum_{x \in \mathcal{X}} \nu(x) H(p(\cdot|x)),$$

where

$$H(p(\cdot|x)) = - \sum_{y \in \mathcal{X}} p(y|x) \log_2 p(y|x)$$

is the entropy of the probability vector which is the row of the transition matrix P indexed by x .

Proof. By the Markov property and the fact that the chain is stationary

$$\begin{aligned}
a_k &:= H(X_k \mid X_0, \dots, X_{k-1}) \\
&= H(X_k \mid X_{k-1}) \\
&= H(X_1 \mid X_0) \\
&= \sum_{x \in \mathcal{X}} \mathbb{P}[X_0 = x] H(X_1 \mid X_0 = x) \\
&= \sum_{x \in \mathcal{X}} \nu(x) H(p(\cdot \mid x)) \\
&= h.
\end{aligned}$$

In particular a_k is a constant sequence with value h , and so by Lemma 3.4 the entropy rate is equal to h . \square

We note that there is a trivial right eigenvector of P with eigenvalue 1 given by the all-ones vector $\mathbf{1} = (1, \dots, 1)$. Indeed, since P is stochastic, for all $x \in \mathcal{X}$

$$(P\mathbf{1})_x = \sum_{y \in \mathcal{X}} p(y \mid x) = 1.$$

More generally, a function $f : \mathcal{X} \rightarrow \mathbb{R}$ is called *harmonic* (with respect to P) if, when viewed as a column vector, it satisfies $Pf = f$. Since

$$(Pf)_x = \sum_{y \in \mathcal{X}} p(y \mid x) f(y) = \sum_{y: p(y \mid x) > 0} p(y \mid x) f(y),$$

we can think of this as saying that the weighted average of the function f over the neighbourhood of x in the digraph of the Markov chain is equal to $f(x)$.

Hence, since the left and right eigenvalues of a matrix agree, there must be *some* vector ν which is a left eigenvector of P with eigenvalue 1. It remains to show that ν is a probability vector. Arranging that ν sums to one is trivial, any linear scaling of an eigenvector lies in the same eigenspace, however it is not obvious that ν is non-negative.

Lemma 3.17. *Let $(X_n)_{n \geq 0}$ be a time-homogeneous Markov chain with a finite state space \mathcal{X} and transition matrix P . Then there is at least one stationary probability distribution ν for P .*

Proof. Given an arbitrary distribution μ on \mathcal{X} , consider the vector

$$\mu_n = \frac{1}{n} (\mu + \mu P + \mu P^2 + \dots + \mu P^{n-1}).$$

We can think of μ as the average distribution of a state in the Markov chain in the first n steps with initial distribution μ .

Since P is stochastic, each μP^i is a probability vector, and hence μ_n is also a probability distribution on \mathcal{X} .

Thinking about μ_n as a sequence of vectors in \mathbb{R}^n , we see that this sequence is bounded, since they are all probability vectors, and hence it has some convergent subsequence, by the

Heine-Borel theorem. Let us say $\mu_{n_k} \rightarrow \mu$ as $k \rightarrow \infty$, that is

$$\lim_{k \rightarrow \infty} \mu_{n_k}(x) = \mu(x) \text{ for all } x \in \mathcal{X}.$$

Furthermore, since $\sum_{x \in \mathcal{X}} \mu_{n_k}(x) = 1$ for all k , and this is a finite sum, it follows that

$$\sum_{x \in \mathcal{X}} \mu(x) = \sum_{x \in \mathcal{X}} \lim_{k \rightarrow \infty} \mu_{n_k}(x) = \lim_{k \rightarrow \infty} \sum_{x \in \mathcal{X}} \mu_{n_k}(x) = 1$$

and similarly it is easy to see that $\mu(x) \in [0, 1]$ for all $x \in \mathcal{X}$, and hence μ is a probability distribution on \mathcal{X} .

Let us consider

$$\begin{aligned} \mu_n P - \mu_n &= \frac{1}{n} (\mu P + \mu P^2 + \mu P^2 + \cdots + \mu P^n) - \frac{1}{n} (\mu + \mu P + \mu P^2 + \cdots + \mu P^{n-1}) \\ &= \frac{1}{n} (\mu P^n - \mu). \end{aligned}$$

However, since μP^n and μ are both probability vectors and thus bounded, it follows that

$$\lim_{n \rightarrow \infty} \mu_n P - \mu_n = \mathbf{0},$$

and again since P is a finite matrix, and multiplication by a finite matrix is a continuous function, we can conclude that

$$\mathbf{0} = \lim_{n_k \rightarrow \infty} \mu_{n_k} P - \mu_{n_k} = \mu P - \mu,$$

and so μ is the desired stationary distribution. \square

Note that the same argument would apply to any accumulation point μ of the sequence μ_n . Can we say when this stationary distribution is unique?

Definition 3.18. Let $(X_n)_{n \geq 0}$ be a time-homogeneous Markov chain with a finite state space \mathcal{X} and transition matrix P . The Markov chain, and transition matrix, are called *irreducible* if for every pair $x, y \in \mathcal{X}$ there is some $n \in \mathbb{N}$ such that $p^{(n)}(y|x) > 0$. That is, for any pair of states, there is some n such that we can transition from one state to the other in n steps with positive probability.

If we think about the associated digraph of the Markov chain, then irreducibility is equivalent to the property that this digraph is strongly connected—for any pair of vertices x and y , there is a directed path from x to y .

Proposition 3.19. Let $(X_n)_{n \geq 0}$ be an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} and transition matrix P . Then there is a unique stationary distribution ν and furthermore $\nu(x) > 0$ for all $x \in \mathcal{X}$.

Proof. We will use the following elementary claim:

Claim 4. If $f : \mathcal{X} \rightarrow \mathbb{R}$ is a harmonic function with respect to an irreducible P , that is, when viewed as a column vector $Pf = f$, then $f = c\mathbf{1}$ is constant.

Proof. Let $m = \min_{x \in \mathcal{X}} f(x)$. If f is not constant and P is irreducible, there exist some $x_0, y_0 \in \mathcal{X}$ such that $m = f(x_0) < f(y_0)$ and $p(y_0|x_0) > 0$. Equivalently, since the digraph of the Markov chain is strongly connected, there is some x_0 with $f(x_0) = m$ which is adjacent to some y_0 with $f(y_0) > m$.

Then, we can write

$$\begin{aligned} m &= f(x_0) \\ &= (Pf)_{x_0} \\ &= \sum_{y \in \mathcal{X}} p(y|x_0)f(y) \\ &\geq (1 - p(y_0|x_0)) \cdot m + p(y_0|x_0) \cdot f(y_0) \\ &> m, \end{aligned}$$

contradicting our assumptions. \square

Now, by Lemma 3.17 there exists at least one stationary distribution ν , which satisfies for all $x \in \mathcal{X}$

$$\nu(x) = (\nu P)_x = \sum_{y \in \mathcal{X}} \nu(y)p(x|y). \quad (3.3)$$

Now, there must exist some $x_0 \in \mathcal{X}$ such that $\nu(x_0) > 0$ and then for any $y \in \mathcal{X}$, since P is irreducible, there exists $n \in \mathbb{N}$ such that $p^{(n)}(y|x_0) > 0$ and hence,

$$\nu(y) = (\nu P^n)_y = \sum_{x \in \mathcal{X}} \nu(x)p^{(n)}(y|x) \geq \nu(x_0)p^{(n)}(y|x_0) > 0.$$

Hence, we can divide by $\nu(x)$ in (3.3) to deduce that

$$1 = \sum_{y \in \mathcal{X}} \frac{\nu(y)}{\nu(x)} p(x|y) := \sum_{y \in \mathcal{X}} q(y|x),$$

where $q(y|x)$ is some (new) conditional probability distribution on \mathcal{X} . We can build a matrix $(Q)_{x,y} = q(y|x)$, which is then related closely to P^T .

This is, in fact, the transition matrix of another stationary Markov chain, often called the *reverse Markov chain*. Indeed, if $X_0 \rightarrow \cdots \rightarrow X_n$, then one can see that $X_n \rightarrow X_{n-1} \rightarrow \cdots \rightarrow X_0$, and it is easy to calculate that if we take ν as the ‘initial’ distribution on X_n , then for all $x, y \in \mathcal{X}$, $\mathbb{P}[X_{n-1} = y | X_n = x] = q(y|x)$. For example,

$$\mathbb{P}[X_0 = y | X_1 = x] = \mathbb{P}[X_1 = x | X_0 = y] \frac{\mathbb{P}[X_0 = y]}{\mathbb{P}[X_1 = x]} = p(x|y) \frac{\nu(y)}{\nu(x)}.$$

In particular, Q is also irreducible. Indeed, the quickest way to see this is via the associated digraph. It is easy to see that $q(y|x) > 0$ if and only if $p(x|y) > 0$, and so the digraph for the reverse Markov chain is given by reversing all the arcs, and hence since the original chain is irreducible, the digraph associated with the reverse chain is also strongly connected.

Suppose μ is a stationary distribution for P , that is, $\mu P = \mu$. Let us consider the function

$$f(x) = \frac{\mu(x)}{\nu(x)},$$

which we wish to show is **1**. We can consider

$$\begin{aligned}
(Qf)_x &= \sum_{y \in \mathcal{X}} q(y|x) f(y) \\
&= \sum_{y \in \mathcal{X}} \frac{\nu(y)}{\nu(x)} p(x|y) \frac{\mu(y)}{\nu(y)} \\
&= \frac{1}{\nu(x)} \sum_{y \in \mathcal{X}} \mu(y) p(x|y) \\
&= \frac{1}{\nu(x)} (\mu P)_x \\
&= \frac{\mu(x)}{\nu(x)} = f(x).
\end{aligned}$$

That is, f is harmonic with respect to Q , and so by the claim $f = c \cdot \mathbf{1}$ for some constant c , or in other words $\mu(x) = c\nu(x)$ for all $x \in \mathcal{X}$. However, since $\sum_{x \in \mathcal{X}} \mu(x) = \sum_{x \in \mathcal{X}} \nu(x) = 1$, it follows that $c = 1$. \square

Corollary 3.20. *Let $(X_n)_{n \geq 0}$ be an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} and transition matrix P . Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} (\mu + \mu P + \mu P^2 + \dots + \mu P^{n-1})$$

exists and is equal to the unique stationary distribution ν .

Proof. The proof of Lemma 3.17 shows that every accumulation point of the sequence is a stationary distribution for P , and so by Lemma 3.19 there is a unique accumulation point. However it is an exercise in analysis to show that a bounded sequence with a unique accumulation point is convergent. \square

Corollary 3.21. *Let $(X_n)_{n \geq 0}$ be an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} and transition matrix P . Then for any initial distribution μ , the entropy rate of the Markov chain exists and is equal to the entropy rate of the Markov chain under its unique stationary distribution (see Lemma 3.16).*

Proof. It is easy to calculate that for $k \geq 1$

$$\begin{aligned}
H(X_k | X_0, \dots, X_{k-1}) &= H(X_k | X_{k-1}) \\
&= \sum_{x \in \mathcal{X}} p_{X_{k-1}}(x) H(X_k | X_{k-1} = x) \\
&= \sum_{x \in \mathcal{X}} \left(\mu P^{k-1} \right)_x H(p(\cdot|x)).
\end{aligned}$$

Now, by the chain rule (Theorem 2.13)

$$\begin{aligned}
\frac{1}{n}H(X_0, \dots, X_{n-1}) &= \frac{1}{n} \sum_{k=0}^{n-1} H(X_k \mid X_0, \dots, X_{k-1}) \\
&= \frac{1}{n}H(X_0) + \frac{1}{n} \sum_{k=1}^{n-1} \sum_{x \in \mathcal{X}} \left(\mu P^{k-1} \right)_x H(p(\cdot \mid x)) \\
&= \frac{1}{n}H(\mu) + \sum_{x \in \mathcal{X}} H(p(\cdot \mid x)) \frac{1}{n} \sum_{k=0}^{n-1} \left(\mu P^{k-1} \right)_x.
\end{aligned}$$

Now, as $n \rightarrow \infty$ we have that $\frac{1}{n}H(\mu) \rightarrow 0$ and by Corollary 3.20 $\frac{1}{n} \sum_{k=0}^{n-1} (\mu P^{k-1})_x \rightarrow \nu(x)$. Hence

$$h = \lim_{n \rightarrow \infty} \frac{1}{n}H(X_0, \dots, X_{n-1}) = \sum_{x \in \mathcal{X}} \nu(x) H(p(\cdot \mid x)).$$

□

Definition 3.22 (Return time). If $(X_n)_{n \geq 0}$ is a Markov chain with a finite state space \mathcal{X} , then for any $x \in \mathcal{X}$ we define

$$\tau^x = \inf\{n \geq 1 : X_n = x\},$$

which is the first time the chain is in state x after the start (where we define $\inf \emptyset = \infty$). Note that τ_x is a random variable! If $X_0 = x$ then we call τ^x the *return time* to x . A state x is called *recurrent* if the Markov chain returns to x almost surely, that is, if

$$\mathbb{P}[\tau^x < \infty \mid X_0 = x] = 1,$$

and it is *positive recurrent* if in addition the return time has finite expectation, that is,

$$\mathbb{E}(\tau^x \mid X_0 = x) < \infty.$$

Theorem 3.23 (Ergodic Theorem for Markov chains). *Let $(X_n)_{n \geq 0}$ be an irreducible, time-homogeneous Markov chain with a finite state space \mathcal{X} . Then every state $x \in \mathcal{X}$ is positive recurrent, and the (unique) stationary distribution ν is given by*

$$\nu(x) = \frac{1}{\mathbb{E}(\tau^x \mid X_0 = x)}.$$

Furthermore, for any initial distribution μ and any function $f: \mathcal{X} \rightarrow \mathbb{R}$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) = \sum_{x \in \mathcal{X}} f(x) \nu(x) \quad \text{almost surely.}$$

The expression on the right hand side is a deterministic quantity, the *space average* of the value $f(x)$ —the average of f over the state space \mathcal{X} under the stationary distribution ν . On the left hand side we have a *random* quantity, the *time average* of f over the trajectory of the random process X_n , and the theorem asserts that the two are almost surely equal.

In applications, we are often interested in the value of the right hand side, however if the state space is large then it can be hard, or inefficient to compute the space average directly. On the other hand, the time average can be approximated by simulating the Markov chain for a large

number of steps and calculating the time average. In this way we get a tool for approximating the sum on the right hand side, which is known as the *Markov chain Monte Carlo* method.

There are a few ways to approach the theorem, we will give a proof using generating functions.

As well as the return time, we also define a similar quantity $s^x = \inf\{n \geq 0: X_n = x\}$. We note that both these quantities are *stopping times*. A stopping time is a random variable such that the event $[\tau^x = n]$ is determined by the initial sequence X_0, \dots, X_n .

Note that if $X_0 = y \neq x$, then $s^x = \tau^x$ and conversely if $X_0 = x$ then $s^x = 0$ and τ^x is the return time to x . Recalling that $p^{(n)}(y|x) = \mathbb{P}[X_n = y \mid X_0 = x]$, we also define the functions

$$f^{(n)}(y|x) = \mathbb{P}[s^y = n \mid X_0 = x] \quad \text{and} \quad u^{(n)}(y|x) = \mathbb{P}[\tau^y = n \mid X_0 = x].$$

Note that if $y \neq X_0$, then $s^y = \tau^y$ and so $f^{(n)}(y|x) = u^{(n)}(y|x)$, and that $u^{(n)}(x|x)$ is the probability that if we start at x then we first return to x at time n . In particular, $u^{(0)}(x|x) = 0$.

Definition 3.24. Given $x, y \in \mathcal{X}$, we define the following formal power series in the variable z :

$$\begin{aligned} P[x, y](z) &= \sum_{n=0}^{\infty} p^{(n)}(y|x) z^n, \\ F[x, y](z) &= \sum_{n=0}^{\infty} f^{(n)}(y|x) z^n, \\ U[x, x](z) &= \sum_{n=0}^{\infty} u^{(n)}(x|x) z^n. \end{aligned}$$

Since $p^{(n)}$, $f^{(n)}$ and $u^{(n)}$ are probabilities, they are bounded by 1, and so the radius of convergence for all these series is at least 1. Moreover, since the events determining $f^{(n)}$ and $u^{(n)}$ are disjoint, their sum is at most 1, and so $F[x, y]$ and $U[x, x]$ converge even on $[-1, 1]$.

Let us note some interpretations of particular values:

$$\begin{aligned} F[x, y](1) &= \sum_{n=0}^{\infty} f^{(n)}(y|x) = \sum_{n=0}^{\infty} \mathbb{P}[s^y = n \mid X_0 = x] = \mathbb{P}[s^y < \infty \mid X_0 = x], \\ U[x, x](1) &= \mathbb{P}[\tau^x < \infty \mid X_0 = x]. \end{aligned}$$

Whilst $P[x, y]$ might not converge at $z = 1$, we see that the sum has the following interpretation:

$$\begin{aligned} P[x, y](1) &= \sum_{n=0}^{\infty} p^{(n)}(y|x) \\ &= \sum_{n=0}^{\infty} \mathbb{P}[X_n = y \mid X_0 = x] \\ &= \text{Expected number of times the chain visits } y \text{ if it starts at } x. \end{aligned}$$

It is also relatively simple to relate the three functions to each other.

Lemma 3.25. (a) For all $z \in [0, 1)$, $P[x, x](z) = \frac{1}{1 - U[x, x](z)}$,

(b) For all $z \in [0, 1]$, $P[x, y](z) = F[x, y](z)P[y, y](z)$,

(c) For all $z \in [0, 1]$, $U[x, x](z) = \sum_{y \in \mathcal{X}} p(y|x)F[y, x](z) \cdot z$,

(d) If $x \neq y$, then for all $z \in [0, 1]$, $F[x, y](z) = \sum_{w \in \mathcal{X}} p(w|x)F[w, y](z) \cdot z$,

Remark 3.26. Note that if $x = y$, then $s^y = 0$ and so $F[y, y](z) = 1$.

Proof. (a) We note that, for all $n \geq 1$,

$$\begin{aligned}
p^{(n)}(x|x) &= \mathbb{P}[X_n = x \mid X_0 = x] \\
&= \mathbb{P}[X_n = x, \tau^x \leq n \mid X_0 = x] \\
&= \sum_{k=1}^n \mathbb{P}[X_n = x, \tau^x = k \mid X_0 = x] \\
&= \sum_{k=1}^n \mathbb{P}[X_n = x \mid \tau^x = k, X_0 = x] \mathbb{P}[\tau^x = k \mid X_0 = x] \\
&= \sum_{k=1}^n p^{(n-k)}(x|x) u^{(k)}(x|x) \\
&= \sum_{k=0}^n p^{(n-k)}(x|x) u^{(k)}(x|x).
\end{aligned}$$

On the other hand, if $n = 0$, then $p^{(0)}(x|x) = 1$, whereas $p^{(0)}(x|x)u^{(0)}(x|x) = 0$.

Hence

$$\begin{aligned}
P[x, x](z) &= \sum_{n=0}^{\infty} p^{(n)}(x|x) z^n \\
&= 1 + \sum_{n=0}^{\infty} \sum_{k=0}^n p^{(n-k)}(x|x) z^{n-k} u^{(k)}(x|x) z^k \\
&= 1 + P[x, x](z)U[x, x](z),
\end{aligned}$$

since the second sum is the product formula for power series, which holds inside the radius of convergence. This re-arranges to the desired equality.

(b) Similarly, for all $n \geq 0$,

$$\begin{aligned}
p^{(n)}(y|x) &= \mathbb{P}[X_n = y \mid X_0 = x] \\
&= \mathbb{P}[X_n = y, s^y \leq n \mid X_0 = x] \\
&= \sum_{k=0}^n \mathbb{P}[X_n = y, s^y = k \mid X_0 = x] \\
&= \sum_{k=0}^n \mathbb{P}[X_n = y \mid s^y = k, X_0 = x] \mathbb{P}[s^y = k \mid X_0 = x] \\
&= \sum_{k=0}^n p^{(n-k)}(y|y) f^{(k)}(y|x).
\end{aligned}$$

Hence,

$$\begin{aligned}
P[x, y](z) &= \sum_{n=0}^{\infty} p^{(n)}(y|x) z^n \\
&= \sum_{n=0}^{\infty} \sum_{k=0}^n p^{(n-k)}(y|y) f^{(k)}(y|x) z^n \\
&= P[y, y](z) F[x, y](z).
\end{aligned}$$

(c) Similarly, for $n \geq 1$,

$$\begin{aligned}
u^{(n)}(x|x) &= \mathbb{P}[\tau^x = n \mid X_0 = x] \\
&= \sum_{y \in X} p(y|x) \mathbb{P}[s^x = n-1 \mid X_0 = y] \\
&= \sum_{y \in X} p(y|x) f^{(n-1)}(x|y).
\end{aligned}$$

Furthermore $u^{(0)}(x|x) = 0$.

Hence,

$$\begin{aligned}
U[x, x](z) &= \sum_{n=0}^{\infty} u^{(n)}(x|x) z^n \\
&= 0 + \sum_{n=1}^{\infty} \sum_{y \in X} p(y|x) f^{(n-1)}(x|y) z^n \\
&= \sum_{y \in X} p(y|x) \left(\sum_{n=1}^{\infty} f^{(n-1)}(x|y) z^{n-1} \right) \cdot z \\
&= \sum_{y \in X} p(y|x) F[y, x](z) \cdot z.
\end{aligned}$$

(d) Similarly, if $x \neq y$ and $n \geq 1$,

$$\begin{aligned}
f^{(n)}(y|x) &= \mathbb{P}[s^y = n \mid X_0 = x] \\
&= \sum_{w \in X} p(w|x) \mathbb{P}[s^y = n-1 \mid X_0 = w] \\
&= \sum_{w \in X} p(w|x) f^{(n-1)}(y|w).
\end{aligned}$$

Furthermore, for $n = 0$, $f^{(0)}(y|x) = 0$.

Hence,

$$\begin{aligned}
F[x, y](z) &= \sum_{n=0}^{\infty} f^{(n)}(y|x) z^n \\
&= 0 + \sum_{n=1}^{\infty} \sum_{w \in X} p(w|x) f^{(n-1)}(y|w) z^n \\
&= \sum_{w \in X} p(w|x) \left(\sum_{n=1}^{\infty} f^{(n-1)}(y|w) z^{n-1} \right) \cdot z \\
&= \sum_{w \in X} p(w|x) F[w, y](z) \cdot z
\end{aligned}$$

□

Recall that a state $x \in \mathcal{X}$ is recurrent if $\mathbb{P}[\tau^x < \infty \mid X_0 = x] = U[x, x](1) = 1$. Let us write $P[x, y](1^-) := \lim_{z \rightarrow 1^-} P[x, y](z)$. Then the following follows immediately from Lemma 3.25 (a).

Lemma 3.27. *$x \in \mathcal{X}$ is recurrent if and only if*

$$P[x, x](1^-) = +\infty.$$

Proof. By Lemma 3.25 (a)

$$P[x, x](1^-) = \lim_{z \rightarrow 1^-} P[x, x](z) = \lim_{z \rightarrow 1^-} \frac{1}{1 - U[x, x](z)}.$$

However, U converges on $[-1, 1]$, so $\lim_{z \rightarrow 1^-} U[x, x](z) = U[x, x](1)$. In particular, since $U[x, x](z) < 1$ for all $z < 1$, it follows that $U[x, x](1) = 1$ if and only if $P[x, x](1^-) = +\infty$. □

Lemma 3.28. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain,*

$$P[x, y](1^-) = \infty \text{ for some } x, y \in \mathcal{X} \iff P[x, y](1^-) = \infty \text{ for all } x, y \in \mathcal{X}.$$

Proof. Let $x, y, x', y' \in \mathcal{X}$. Since the chain is irreducible, there exist $k, \ell \in \mathbb{N}$ such that $p^{(k)}(x'|x) > 0$ and $p^{(\ell)}(y|y') > 0$. Hence

$$p^{(k)}(x'|x)p^{(n)}(y'|x')p^{(\ell)}(y|y') \leq p^{(k+n+\ell)}(y|x).$$

It follows that, for fixed k and ℓ and $z > 0$

$$\begin{aligned} P[x, y](z) &\geq \sum_{n=0}^{\infty} p^{(k+n+\ell)}(y|x) z^{k+n+\ell} \\ &\geq \sum_{n=0}^{\infty} p^{(k)}(x'|x) z^k p^{(n)}(y'|x') z^n p^{(\ell)}(y|y') z^\ell \\ &= p^{(k)}(x'|x) z^k P[x', y'](z) p^{(\ell)}(y|y') z^\ell, \end{aligned}$$

and so if $P[x', y'](1^-) = +\infty$, then

$$P[x, y](1^-) \geq p^{(k)}(x'|x) p^{(\ell)}(y|y') P[x', y'](1^-) = +\infty.$$

□

The following is then a consequence of Lemmas 3.25 and 3.27.

Theorem 3.29. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} , then the following are equivalent:*

- (i) *There exists an $x \in \mathcal{X}$ which is recurrent,*
- (ii) *For all $x \in \mathcal{X}$, x is recurrent,*

(iii) There exist $x, y \in \mathcal{X}$ such that $P[x, y](1^-) = +\infty$,

(iv) For all $x, y \in \mathcal{X}$, $P[x, y](1^-) = +\infty$,

(v) For all $x, y \in \mathcal{X}$, $F[x, y](1) = 1$.

Proof. By Lemma 3.27, (i) implies $P[x, x](1^-) = +\infty$ and hence also implies (iii), which by Lemma 3.28 is equivalent to (iv). However, conversely by Lemma 3.27, (iv) implies (ii), which clearly implies (i). Hence, the first four are equivalent, and it remains to show that the last property is also equivalent.

By Lemma 3.25 (c), if (v) holds, then

$$U[x, x](1) = \sum_{y \in \mathcal{X}} p(y|x)F[y, x](1) = \sum_{y \in \mathcal{X}} p(y|x) = 1,$$

and so (v) implies (i). Conversely, assume that (ii) holds and fix $y \in \mathcal{X}$. For any $x \neq y$, by Lemma 3.25 (d),

$$F[x, y](1) = \sum_{w \in \mathcal{X}} p(w|x)F[w, y](1),$$

which essentially says that the function $f(x) = F[x, y](1)$ is harmonic (except perhaps at the point y). However, since the chain is irreducible, this is still sufficient to conclude that $f(x)$ is constant (exercise). However, by (ii) and Lemma 3.25 (c),

$$1 = U[y, y](1) = \sum_{w \in \mathcal{X}} p(w|y)F[w, y](1) = \sum_{w \in \mathcal{X}} p(w|y)f(w),$$

and hence $f(x)$ must be identically 1. Since our choice of y was arbitrary, this implies (v). \square

We say that a time-homogeneous Markov chain $(X_n)_{n \geq 0}$ is *recurrent* if every state x is recurrent.

Theorem 3.30. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} , then it is recurrent.*

Proof. In some sense this theorem is obvious. We have a Markov chain that takes values in a finite state space. Clearly some $x \in \mathcal{X}$ must be visited infinitely often, and hence is recurrent, and by Theorem 3.29 this implies that every $y \in \mathcal{X}$ is recurrent, and so the chain is recurrent.

Let $0 \leq z < 1$ and fix $x \in \mathcal{X}$. We consider the sum

$$\sum_{y \in \mathcal{X}} P[x, y](z) = \sum_{y \in \mathcal{X}} \sum_{n=0}^{\infty} p^{(n)}(y|x)z^n = \sum_{n=0}^{\infty} z^n \sum_{y \in \mathcal{X}} p^{(n)}(y|x) = \sum_{n=0}^{\infty} z^n = \frac{1}{1-z}. \quad (3.4)$$

Hence $\lim_{z \rightarrow 1^-} \sum_{y \in \mathcal{X}} P[x, y](z) = +\infty$ and since \mathcal{X} is finite there must be some $y \in \mathcal{X}$ such that $P[x, y](1^-) = +\infty$ and hence by Lemma 3.27, y is recurrent, and since the chain is irreducible, by Theorem 3.29 the chain is recurrent. \square

Theorem 3.31. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} , then the unique stationary distribution is given by*

$$\nu(x) = \frac{1}{\mathbb{E}(\tau^x | X_0 = x)}.$$

Proof. From the equality (3.4) together with Lemma 3.25, we can deduce the equality

$$1 = (1 - z) \sum_{y \in \mathcal{X}} P[x, y](z) = (1 - z) \sum_{y \in \mathcal{X}} F[x, y](z) \frac{1}{1 - U[y, y](z)} = \sum_{y \in \mathcal{X}} F[x, y](z) \frac{1 - z}{1 - U[y, y](z)}.$$

Since this equality holds for all $z < 1$, it also holds in the limit as $z \rightarrow 1^-$. However, if the chain is recurrent then

$$F[x, y](1) = \mathbb{P}[s^y < \infty | X_0 = x] = 1,$$

and so the limit

$$\lim_{z \rightarrow 1^-} \sum_{y \in \mathcal{X}} \frac{1 - z}{1 - U[y, y](z)} = 1$$

must exist and be equal to 1. Since U is differentiable inside the radius of convergence, we can conclude by L'Hospital's rule that

$$1 = \lim_{z \rightarrow 1^-} \sum_{y \in \mathcal{X}} \frac{1}{U'[y, y](z)} = \sum_{y \in \mathcal{X}} \frac{1}{U'[y, y](1^-)},$$

and we can give a combinatorial interpretation of this derivative. Indeed

$$U'[y, y](z) = \sum_{n=1}^{\infty} n u^{(n)}(x|x) z^{n-1},$$

and so in the limit as $z \rightarrow 1^-$ this will tend to

$$\sum_{n=1}^{\infty} n \mathbb{P}[\tau^y = n | X_0 = y] = \mathbb{E}(\tau^y | X_0 = y).$$

Hence, by the above the values satisfy

$$1 = \sum_{y \in \mathcal{X}} \frac{1}{U'[y, y](1^-)} = \sum_{y \in \mathcal{X}} \frac{1}{\mathbb{E}(\tau^y | X_0 = y)},$$

or in other words, they define a probability distribution on \mathcal{X} , $\nu(x) = \frac{1}{\mathbb{E}(\tau^x | X_0 = x)}$.

Let us write $P(z)$ for the matrix given by $(P(z))_{xy} = P[x, y](z) = \sum_{n=0}^{\infty} z^n P^n$. Note that

$$P(z) = \sum_{n=0}^{\infty} z^n P^n = I + \sum_{n=1}^{\infty} z^n P^n = I + \left(\sum_{n=0}^{\infty} z^n P^n \right) zP = I + P(z) \cdot zP. \quad (3.5)$$

By a similar argument, $P(z) = I + zP \cdot P(z)$ and so, since the state space is finite, we can conclude that

$$P(z) = (I - zP)^{-1},$$

but we will not need this.

Looking at the diagonal entries of $P(z)$, we see using (3.5) that

$$P[y, y](z) = 1 + \sum_{x \in \mathcal{X}} P[y, x](z) p(y|x) z,$$

and so by Lemma 3.25

$$\frac{1}{1 - U[y, y](z)} = 1 + \sum_{x \in \mathcal{X}} F[y, x](z) \frac{1}{1 - U[x, x](z)} p(y|x) z.$$

Multiplying through by $1 - z$

$$\frac{1 - z}{1 - U[y, y](z)} = 1 - z + \sum_{x \in \mathcal{X}} F[y, x](z) \frac{1 - z}{1 - U[x, x](z)} p(y|x) z.$$

Taking limits as $z \rightarrow 1^-$ we see again by L'Hospital's rule

$$\nu(y) = \sum_{x \in \mathcal{X}} \nu(x) p(y|x),$$

or in other words, $\nu = \nu P$ and ν is a stationary distribution. However, by Proposition 3.19 an irreducible time-homogeneous Markov chain has a unique stationary distribution. \square

Corollary 3.32. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} , then the chain is positive recurrent.*

Proof. By Theorem 3.30 the chain is recurrent, and by Theorem 3.31 the unique stationary distribution is given by

$$\nu(x) = \frac{1}{\mathbb{E}(\tau^x | X_0 = x)}.$$

However, by Proposition 3.19, for an irreducible time-homogeneous Markov chain the stationary distribution satisfies $\nu(x) > 0$ for all x , and hence

$$\mathbb{E}(\tau^x | X_0 = x) < \infty$$

for all x , and so the chain is positive recurrent. \square

So, it remains to prove that the time average of any function $f: \mathcal{X} \rightarrow \mathbb{R}$ over the Markov chain converges almost surely to the space average $\sum_{x \in \mathcal{X}} \nu(x) f(x)$.

Let us consider first the special case where $f(x) = \mathbb{1}_x$, i.e., $f(x) = 1$ and $f(y) = 0$ otherwise. In this case

$$\frac{1}{n} \sum_{k=0}^{n-1} f(X_k) = \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_x(X_k) = \text{Proportion of time spent in state } x,$$

which should then be equal to $\nu(x)$. In fact, heuristically, the general case should follow from this—if on average the chain spends proportionally $\nu(x)$ of the time in each $x \in \mathcal{X}$, then the time average of f will tend to $\sum_{x \in \mathcal{X}} \nu(x) f(x)$.

Lemma 3.33. *If $(X_n)_{n \geq 0}$ is an irreducible time-homogeneous Markov chain with a finite state space \mathcal{X} and $x \in \mathcal{X}$, then*

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_x(X_k) \rightarrow \nu(x)$$

almost surely, where $\nu(x) = \frac{1}{\mathbb{E}(\tau^x | X_0 = x)}$ is the unique stationary distribution.

Proof. Let us fix $x \in \mathcal{X}$ and suppose that $X_0 = x_0$ (which may or may not be equal to x). We can define a sequence of times τ_0, τ_1, \dots which are all the times at which the chain is in state x , where $\tau_0 = s^x$ is the first time the chain enters the state x , and $\tau_k = \inf\{n > \tau_{k-1} : X_n = x\}$. Note that, since almost surely τ^x is finite, almost surely this infimum will be an achieved minimum for each k and so almost surely every τ_k will be finite. Note that again each τ_k is a stopping time.

We are interested then in the length of the *excursions* $\Delta_k = \tau_k - \tau_{k-1}$, between visits to x . The key point here is that, by the Markov property, the quantities Δ_k are identically and independently distributed (at least for $k \geq 1$). Indeed, at time τ_{k-1} we are in state x and τ_k is the first return time, and so by the memoryless property of the Markov chain Δ_k has the same distribution as $(\tau^x | X_0 = x)$, which we note by Corollary 3.32 has finite expectation.

Hence, we can view

$$\tau_k - \tau_0 = \sum_{j=1}^k \Delta_j$$

as a sum of i.i.d. random variables with finite expectation, and so by the strong law of large numbers (Theorem 1.29)

$$\frac{1}{k} \sum_{j=1}^k \Delta_j \rightarrow \mathbb{E}(\Delta_1) = \mathbb{E}(\tau^x | X_0 = x) = \frac{1}{\nu(x)} \quad \text{almost surely.}$$

Finally, since τ_0 is almost surely finite, it follows that

$$\frac{1}{k} \tau_k = \frac{1}{k} (\tau_k - \tau_0) \rightarrow \frac{1}{\nu(x)} \quad \text{almost surely.}$$

Moreover, it follows that

$$\frac{\tau_{k+1}}{\tau_k} = \frac{k+1}{k} \frac{\frac{1}{k+1} \tau_{k+1}}{\frac{1}{k} \tau_k} \rightarrow 1 \quad \text{almost surely.}$$

Given $n \in \mathbb{N}$, there is some $k(n)$ such that $\tau_{k(n)} < n \leq \tau_{k(n)+1}$ (note that $k(n)$ is a random variable here), and so

$$1 < \frac{n}{\tau_{k(n)}} \leq \frac{\tau_{k(n)+1}}{\tau_{k(n)}}.$$

Hence, taking limits as $n \rightarrow \infty$,

$$\frac{n}{\tau_{k(n)}} \rightarrow 1 \quad \text{almost surely.}$$

However, we have already shown that $\frac{\tau_k}{k} \rightarrow \frac{1}{\nu(x)}$ almost surely, and so combining the two we see that

$$\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_x(X_i) = \frac{k(n)}{n} \rightarrow \nu(x) \quad \text{almost surely,}$$

as claimed. □

Finally we can prove the ergodic theorem.

Proof of Theorem 3.23. Given some function $f: \mathcal{X} \rightarrow \mathbb{R}$, we can decompose f as a finite sum of weighted indicator functions

$$f = \sum_{x \in \mathcal{X}} f(x) \mathbb{1}_x.$$

By Lemma 3.33,

$$\begin{aligned} \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) &= \frac{1}{n} \sum_{k=0}^{n-1} \sum_{x \in \mathcal{X}} f(x) \mathbb{1}_x(X_k) \\ &= \sum_{x \in \mathcal{X}} f(x) \left(\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_x(X_k) \right) \\ &\rightarrow \sum_{x \in \mathcal{X}} f(x) \nu(x) \quad \text{almost surely.} \end{aligned}$$

□

Some short remarks to end here. If the state space is not finite, but countable, then an irreducible Markov chain can be transient (not recurrent), or even recurrent but not positive recurrent.

For example, consider a biased random walk on \mathbb{Z} which moves right with probability p and left with probability $1 - p$. This chain is clearly irreducible—there is a positive probability to move from any state x to any other state y in a finite number of steps, however for $p \neq \frac{1}{2}$, it can be seen that the Markov chain is not recurrent.

Indeed, if we let $(Y_n)_{n \geq 0}$ be an i.i.d. sequence of random variables with $\mathbb{P}[Y = 1] = p$ and $\mathbb{P}[Y = -1] = 1 - p$, then it can be seen that the random walk $(X_n)_{n \geq 0}$ described above is given by $X_n = \sum_{i=1}^n Y_i$. By the law of large numbers, if $p \neq \frac{1}{2}$, then $\mathbb{E}(Y_i) = p - q \neq 0$, and almost surely $\frac{1}{n} X_n \rightarrow p - q$. Thus, almost surely X_n tends to $+\infty$ or $-\infty$.

In the case $p = q = \frac{1}{2}$, one can calculate explicitly the return probabilities and one finds that the chain is recurrent but the expected return time is infinite.

When the chain is positive recurrent, a version of the ergodic theorem still holds, but only for a restricted class of functions f .

3.3 The Asymptotic Equipartition Property

Let $(X_n)_{n \geq 1}$ be a stochastic process with a finite state space \mathcal{X} . For each $n \in \mathbb{N}$ we can consider the joint distribution $p_n = p_{X_1, \dots, X_n}$ on \mathcal{X}^n , that is

$$p_n(x_1, \dots, x_n) = \mathbb{P}[X_1 = x_1, X_2 = x_2, \dots, X_n = x_n].$$

Note that the sequence of distributions $(p_n)_{n \geq 1}$ (which are deterministic functions on \mathcal{X}^n) determines all the probabilistic characteristics of the stochastic process.

Suppose that the entropy rate of the stochastic process

$$h = \lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, \dots, X_n) = \lim_{n \rightarrow \infty} \frac{1}{n} H(p_n)$$

exists. Since for any random variable X we can write

$$H(X) = \mathbb{E}(-\log_2 p_X(X)),$$

as the expected value of the deterministic function $-\log_2 \circ p_X$ applied to the random variable X , we can apply this to the random vector (X_1, \dots, X_n) to conclude that

$$\frac{1}{n} H(X_1, \dots, X_n) = \mathbb{E} \left(-\frac{1}{n} \log_2 p_n(X_1, \dots, X_n) \right).$$

Then the entropy rate h is the limit of the expected value of the random variables $Y_n = -\frac{1}{n} \log_2 p_n(X_1, \dots, X_n)$. A much stronger property than the limit of the expectation existing would be that the random variables themselves converge (in probability or almost surely) to some limiting random variable with finite expectation h .

Definition 3.34 (Asymptotic equipartition property). Let $(X_n)_{n \geq 1}$ be a stochastic process with a finite state space \mathcal{X} whose entropy rate h exists. We say X has the *asymptotic equipartition property* (AEP), if

$$-\frac{1}{n} \log_2 p_n(X_1, \dots, X_n) \longrightarrow h \text{ almost surely, as } n \rightarrow \infty.$$

The asymptotic equipartition property makes a very strong *prediction* about the observed outcome $(x_n)_{n \geq 1}$ of the stochastic process: namely, with very high probability the quantity $-\frac{1}{n} \log_2 p_n(x_1, \dots, x_n)$ will be close to the entropy rate h , where the error in probability and in approximation to h is tending to 0 with n .

Example 3.35. Let $\mathcal{X} = \{0, 1\}$, and let $(X_n)_{n \geq 1}$ be i.i.d. with distribution $\text{Ber}(\theta)$, that is,

$$\mathbb{P}[X_n = 1] = \theta \quad \text{and} \quad \mathbb{P}[X_n = 0] = 1 - \theta,$$

where $0 < \theta < 1$. For any bit string $(x_1, \dots, x_n) \in \{0, 1\}^n$, let

$$s_n = x_1 + \dots + x_n$$

be the total number of ones, so that $n - s_n$ is the total number of zeroes. Let us write S_n for the random variable $X_1 + \dots + X_n$.

Since the sequence is i.i.d., we can compute

$$p_n(x_1, \dots, x_n) = \theta^{s_n} (1 - \theta)^{n-s_n} \quad \text{and so} \quad p_n(X_1, \dots, X_n) = \theta^{S_n} (1 - \theta)^{n-S_n}.$$

Hence,

$$-\frac{1}{n} \log_2 p_n(X_1, \dots, X_n) = -\frac{1}{n} \log_2 (\theta^{S_n} (1 - \theta)^{n-S_n}) = -\frac{S_n}{n} \log_2 \theta - \left(1 - \frac{S_n}{n}\right) \log_2 (1 - \theta).$$

On the other hand, since the X_n are i.i.d., by Lemma 3.5 the entropy rate h exists and is equal to

$$h = H(X_1) = H(\theta, 1 - \theta) = -\theta \log_2 \theta - (1 - \theta) \log_2 (1 - \theta).$$

So, in this case, the statement that $(X_n)_{n \geq 1}$ satisfies the AEP would be that almost surely

$$-\frac{S_n}{n} \log_2 \theta - \left(1 - \frac{S_n}{n}\right) \log_2 (1 - \theta) \longrightarrow -\theta \log_2 \theta - (1 - \theta) \log_2 (1 - \theta).$$

Or, in other words, the AEP is equivalent to the statement that $\frac{S_n}{n} \rightarrow \theta$ almost surely, which is the strong law of large numbers.

In fact the argument above works in general for i.i.d. sequences.

Lemma 3.36. *Let $(X_n)_{n \geq 1}$ be an i.i.d. stochastic process with a finite state space \mathcal{X} . Then $(X_n)_{n \geq 1}$ satisfies the AEP, where the entropy rate is $h = H(X_1)$.*

Proof. Let $p = p_{X_1} = p_{X_n}$ for all n , so that $p_n(x_1, \dots, x_n) = p(x_1) \dots p(x_n)$. Then we can calculate

$$-\frac{1}{n} \log_2 p_n(X_1, \dots, X_n) = -\frac{1}{n} \log_2 p(X_1) \dots p(X_n) = \frac{1}{n} \sum_{k=1}^n -\log_2 p(X_k),$$

where we may assume that X_1 takes all values in \mathcal{X} with strictly positive probability, and hence that (X_1, \dots, X_n) takes all values in \mathcal{X}^n with strictly positive probability, so that this sum is well-defined.

However, since the X_k are i.i.d., so are the random variables $-\log_2 p(X_k)$ and so by the strong law of large numbers (Theorem 1.29)

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n -\log_2 p(X_k) = \mathbb{E}(-\log_2 p(X_1)) = H(X_1) \quad \text{almost surely.}$$

□

There are perhaps two natural questions to ask at this point:

- (I) Which classes of stochastic process have the AEP (does this include a nice large natural class)?

(II) What is the practical application of knowing that we have the AEP?

Theorem 3.37. *Let $(X_n)_{n \geq 0}$ be an irreducible, time-homogeneous Markov chain with a finite state space \mathcal{X} . Then for any initial distribution, $(X_n)_{n \geq 0}$ satisfies the AEP, where the entropy rate h is given by the formula in Corollary 3.21 / Lemma 3.16.*

Proof. Recall that given an initial distribution μ , we can write the joint distribution of (X_0, \dots, X_{n-1}) as

$$p_n(x_0, \dots, x_{n-1}) = \mu(x_0)p(x_1|x_0)p(x_2|x_1) \dots p(x_{n-1}|x_{n-2}).$$

Hence we can write

$$\begin{aligned} -\frac{1}{n} \log_2 p(X_0, \dots, X_{n-1}) &= -\frac{1}{n} \log_2 \mu(X_0)p(X_1|X_0)p(X_2|X_1) \dots p(X_{n-1}|X_{n-2}) \\ &= -\frac{1}{n} \log_2 \mu(X_0) - \frac{1}{n} \sum_{k=1}^{n-1} \log_2 p(X_k|X_{k-1}). \end{aligned} \quad (3.6)$$

Since μ is a fixed distribution on a finite set \mathcal{X} , the first term $-\frac{1}{n} \log_2 \mu(X_0)$ tends to 0 almost surely. For the terms in the sum, we can think of $-\log_2 p(X_k|X_{k-1})$ as a function $f(X_k, X_{k-1})$, given by $f(x, y) = -\log_2 p(y|x)$.

Now, since $(X_n)_{n \geq 0}$ is a Markov chain, by Lemma 3.11 the sequence $((X_n, X_{n+1}))_{n \geq 0}$ is also a Markov chain. The state space of this Markov chain is given by $\tilde{\mathcal{X}} = \{(x, y) \in \mathcal{X}^2 : p(y|x) > 0\}$, whose elements correspond to the edges in the digraph D of the original Markov chain, with distribution

$$\tilde{p}_n((x_2, y_2)|(x_1, y_1)) = \begin{cases} p(y_2|x_2) & \text{if } x_2 = y_1 \\ 0 & \text{if } x_2 \neq y_1, \end{cases}$$

where we note that, since this doesn't depend on n the Markov chain is also time-homogeneous. Furthermore it is easy to see that this chain is also irreducible—the digraph \tilde{D} of this Markov chain is the line digraph of D , which can be seen to be strongly connected since D is strongly connected (exercise).

Hence, we may apply the Ergodic Theorem (Theorem 3.23), specifically to the function $f : \tilde{\mathcal{X}} \rightarrow \mathbb{R}$ defined above, to conclude that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^{n-1} -\log_2 p(X_k|X_{k-1}) = \lim_{n \rightarrow \infty} \frac{n-1}{n} \left(\frac{1}{n-1} \sum_{k=1}^{n-1} f(X_k, X_{k-1}) \right) = \sum_{(x,y) \in \tilde{\mathcal{X}}} f(x, y) \tilde{\nu}(x, y), \quad (3.7)$$

where $\tilde{\nu}$ is the stationary distribution on $\tilde{\mathcal{X}}$. Unfortunately, we don't know yet what the stationary distribution is!

However, since the stationary distribution is unique, if we can find (or in fact, guess) a distribution $\tilde{\nu}$ such that $\tilde{\nu} \tilde{P} = \tilde{\nu}$, where \tilde{P} is the transition matrix of the edge Markov chain, then we can conclude that $\tilde{\nu}$ is the stationary distribution.

What would be a reasonable guess? Well, we know that the stationary distribution ν for P represents the ‘time-average’ state of the chain, the long term probability that the state of the chain corresponds to a particular vertex of the digraph. However, if we know the long

term probability of being in a particular vertex, and the probability that we traverse each edge incident to this vertex, then we can calculate the long term probability of traversing each edge!

Explicitly, if ν is the stationary distribution for P , then we take $\tilde{\nu}(x, y) = \nu(x)p(y|x)$, and it is easy to verify that $\tilde{\nu}$ is the stationary distribution for \tilde{P} . Indeed,

$$\begin{aligned} (\tilde{\nu}\tilde{P})_{(x,y)} &= \sum_{(x',y') \in \tilde{\mathcal{X}}} \tilde{\nu}(x', y') \tilde{p}((x, y)|(x', y')) \\ &= \sum_{(x',y') : y'=x} \nu(x') p(y'|x') p(y|x) \\ &= \sum_{x' \in \tilde{X}} \nu(x') p(x|x') p(y|x) \\ &= \nu(x) p(y|x) = \tilde{\nu}(x, y). \end{aligned}$$

Hence, by (3.6) and (3.7), it follows that

$$\begin{aligned} \lim_{n \rightarrow \infty} -\frac{1}{n} \log_2 p(X_0, \dots, X_n) &= 0 + \sum_{(x,y) \in \tilde{\mathcal{X}}} f(x, y) \tilde{\nu}(x, y) \\ &= \sum_{(x,y) \in \tilde{\mathcal{X}}} f(x, y) \tilde{\nu}(x, y) \\ &= - \sum_{(x,y) \in \tilde{\mathcal{X}}} \nu(x) p(y|x) \log_2 p(y|x) \\ &= \sum_{x \in \mathcal{X}} \nu(x) \left(\sum_{y \in \mathcal{X}} -p(y|x) \log_2 p(y|x) \right) \\ &= \sum_{x \in \mathcal{X}} \nu(x) H(p(\cdot|x)), \end{aligned}$$

which is the formula for the entropy rate of the Markov chain. □

What can we conclude from the fact that the AEP holds? Since convergence almost surely implies convergence in probability, if the AEP holds then for any $\varepsilon > 0$

$$\mathbb{P} \left[\left| -\frac{1}{n} \log_2 p_n(X_1, \dots, X_n) - h \right| < \varepsilon \right] \rightarrow 1. \quad (3.8)$$

In other words, there is some *deterministic set* inside the set of trajectories \mathcal{X}^n , which we can specify ahead of time in terms of the deterministic function p_n and the quantity h , such that with very high probability the trajectory of the process lies inside this set.

Definition 3.38 (Typical set). Let $(X_n)_{n \geq 1}$ be a stochastic process that satisfies the AEP with entropy rate h . For every n and (small) $\varepsilon > 0$ the *typical set* is given by

$$A_\varepsilon^{(n)} = \left\{ \mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n : \left| -\frac{1}{n} \log_2 p_n(x_1, \dots, x_n) - h \right| < \varepsilon \right\}.$$

The following properties of typical sets follow immediately from the definitions.

Proposition 3.39. *Let $(X_n)_{n \geq 1}$ be a stochastic process that satisfies the AEP with entropy rate h . For all (small) $\varepsilon > 0$ the typical set has the following properties:*

(a) *There exists $N(\varepsilon)$ such that for all $n \geq N(\varepsilon)$*

$$\mathbb{P}[(X_1, \dots, X_n) \in A_\varepsilon^{(n)}] > 1 - \varepsilon.$$

(b) *For all $\mathbf{x} = (x_1, \dots, x_n) \in A_\varepsilon^{(n)}$,*

$$2^{-n(h+\varepsilon)} < p_n(\mathbf{x}) < 2^{-n(h-\varepsilon)}.$$

(c) *The size of the typical set satisfies*

$$(1 - \varepsilon)2^{n(h-\varepsilon)} < |A_\varepsilon^{(n)}| < 2^{n(h+\varepsilon)},$$

where the second inequality holds for all n , and the first for all $n \geq N(\varepsilon)$.

Proof. The first follows immediately from (3.8). The second is the definition of $A_\varepsilon^{(n)}$, since

$$p_n(\mathbf{x}) = 2^{-n(h \pm \varepsilon)} \iff -\frac{1}{n} \log_2 p_n(\mathbf{x}) = h \pm \varepsilon.$$

The third follows then since the first property implies that

$$1 - \varepsilon < \mathbb{P}[(X_1, \dots, X_n) \in A_\varepsilon^{(n)}] = \sum_{\mathbf{x} \in A_\varepsilon^{(n)}} p_n(\mathbf{x}) \leq 1,$$

and the second implies that

$$|A_\varepsilon^{(n)}| 2^{-n(h+\varepsilon)} < \sum_{\mathbf{x} \in A_\varepsilon^{(n)}} p_n(\mathbf{x}) < |A_\varepsilon^{(n)}| 2^{-n(h-\varepsilon)}$$

□

So, from (a) we see that p_n is almost concentrated on the set $A_\varepsilon^{(n)}$, and from (b) and (c) we see furthermore that it is almost *equidistributed* on this set (and this is where the name asymptotic equipartition property comes from)!

In general, the fact that p_n is concentrated on this ‘smaller’ set $A_\varepsilon^{(n)}$ will be most useful when $|A_\varepsilon^{(n)}| \ll |\mathcal{X}^n|$, which in light of (c) will be the case when

$$2^{n(h+\varepsilon)} \ll |\mathcal{X}^n| \iff h + \varepsilon < \log_2 |\mathcal{X}|,$$

in which case $A_\varepsilon^{(n)}$ will be exponentially smaller than \mathcal{X}^n . When the X_n are i.i.d. and uniformly distributed, then $h = \log_2 |\mathcal{X}|$, and the typical set consists of almost all of the possible trajectories.


Whilst there are possibly many more possible trajectories in \mathcal{X}^n than typical sequences in $A_\varepsilon^{(n)}$, it is vanishingly unlikely that the observed trajectory lies outside of $A_\varepsilon^{(n)}$, and so these non-typical trajectories play no significant role in the analysis of the process.

4 Data compression and Codes

4.1 Block codes

Suppose we have a set of elements \mathcal{X} , for example the alphabet of some language, and we wish to encode an element $x \in \mathcal{X}$, using a binary string or in general the elements of some finite set Σ . It is clear that we can represent each element of $x \in \mathcal{X}$ by a unique binary string of length $n = \lceil \log |\mathcal{X}| \rceil$ and so we can encode an arbitrary element using at most n bits of information, but equally we cannot do so with $n - 1$ bits.

However if there is some distribution, given by a random variable X on \mathcal{X} which we call a *source*, in which some elements are more likely to appear than others, then it might be that we can exploit this to find an encoding whose length is shorter on average, or one which is deterministically shorter, but has some (small) probability of error.



Definition 4.1 (Encoding scheme). An encoding scheme is a triple (X, C, g) where X is some discrete random variable taking values in a finite set \mathcal{X} , C is a *code*, that is a mapping

$$C: \mathcal{X} \rightarrow \bigcup_{n=1}^{\infty} \{0, 1\}^n := \{0, 1\}^+,$$

and $g: \{0, 1\}^+ \rightarrow \hat{\mathcal{X}} := \mathcal{X} \cup \{\perp\}$ is a *decoding function*, where \perp represents an error.

That is, the source X is encoded as $C(X)$, and is decoded by the decoding function to $\hat{X} = g(C(X))$. If C is injective, then we can take the decoding function g to be any function such that $C \circ g$ is the identity on \mathcal{X} and then $X = \hat{X}$ and this is a *lossless* encoding, otherwise we will also be interested in the *error probability* $p_{\text{err}} = \mathbb{P}[\hat{X} \neq X]$.

In general, we might wish to encode not just one element $x \in \mathcal{X}$, but a string of elements (x_1, \dots, x_n) , which are then generated according to some stochastic process $(X_n)_{n \geq 1}$. In this case we can ask whether it is more efficient to encode each element individually, or instead to encode longer strings.

To begin with, we will consider the case where we encode the whole string, and so in general our source will be a random vector (X_1, \dots, X_n) , and we are interested in finding a coding $C^{(n)}: \mathcal{X}^n \rightarrow \{0, 1\}^+$, either lossless or with error probability tending to zero, which is particularly *efficient*, either in terms of the average number of symbols we use in a lossless encoding, or perhaps in terms of the maximum number of symbols we use in a lossy encoding.

More precisely, for a lossless encoding we want to minimise

$$L^{(n)} = L^{(n)}(C^{(n)}) = \mathbb{E} \left(\frac{1}{n} \ell(C^{(n)}(X_1, \dots, X_n)) \right),$$

where $\ell(\cdot)$ measures the length of a binary string. The AEP gives us a powerful method of *data compression*, which allows us to encode messages with a close to optimal rate.

Theorem 4.2. *Let $(X_n)_{n \geq 1}$ be a stochastic process which satisfies the AEP with rate h . Then there is a lossless encoding $C^{(n)}: \mathcal{X}^n \rightarrow \{0, 1\}^+$ such that*

$$\lim_{n \rightarrow \infty} L^{(n)} \leq h.$$

Proof. Let us fix $\varepsilon' > 0$, let n be sufficiently large and let ε be such that

$$\varepsilon(1 + \log_2 |\mathcal{X}|) + \frac{1 + \varepsilon}{n} \leq \varepsilon'.$$

We will define an encoding $C^{(n)}: \mathcal{X}^n \rightarrow \{0, 1\}^+$ with $L^{(n)} \leq h + \varepsilon'$.

Since $(X_n)_{n \geq 1}$ satisfies the AEP, we can consider the typical sets $A_\varepsilon^{(n)}$ as defined in Definition 3.38. By Proposition 3.39 (c), $|A_\varepsilon^{(n)}| < 2^{n(h+\varepsilon)}$, and so there is an injective mapping $C^{(n)}: A_\varepsilon^{(n)} \rightarrow \{0, 1\}^{k_n}$ where $k_n = \lceil n(h + \varepsilon) \rceil + 1$, and we insist that the image $C(\mathbf{x})$ begin with a 0.

For the remaining elements $\mathbf{x} \notin A_\varepsilon^{(n)}$ we define the code $C^{(n)}$ as some arbitrary injection to $\{0, 1\}^{k'_n}$ where $k'_n = n \log_2 |\mathcal{X}| + 1$, and we insist that the image $C^{(n)}(\mathbf{x})$ begin with a 1.

(These restrictions on the first bit of the images are not strictly necessary for the theorem, but they guarantee that the code is even *prefix-free*, which we will later see is an important property.)

Clearly this code is injective, what is its expected length? Well, by Proposition 3.39 (a)

$$\begin{aligned} n \cdot L_C^n &= \sum_{\mathbf{x} \in A_\varepsilon^{(n)}} p(\mathbf{x}) \ell(C(\mathbf{x})) + \sum_{\mathbf{x} \notin A_\varepsilon^{(n)}} p(\mathbf{x}) \ell(C(\mathbf{x})) \\ &\leq k_n \sum_{\mathbf{x} \in A_\varepsilon^{(n)}} p(\mathbf{x}) + k'_n \sum_{\mathbf{x} \notin A_\varepsilon^{(n)}} p(\mathbf{x}) \\ &\leq (\lceil n(h + \varepsilon) \rceil + 1) \cdot 1 + (n \log_2 |\mathcal{X}| + 1) \cdot \varepsilon \end{aligned}$$

and hence

$$L^{(n)} \leq h + \varepsilon(1 + \log_2 |\mathcal{X}|) + \frac{1 + \varepsilon}{n} \leq h + \varepsilon'.$$

□

Of course, if we just insist the encoding should be lossless, then one can construct an encoding function minimising $L^{(n)}$ by greedily assigning the most likely strings $\mathbf{x} \in \mathcal{X}^n$ to the shortest strings in $\{0, 1\}^+$. However, the code we constructed in Theorem 4.2 has some extra properties which we will see later are useful.

On the other hand, if we consider lossy encoding, we can hope to minimise even the maximum number of symbols we use. The simplest case would be to try to minimise the number of symbols used in some *block code*, which encodes sequences of length n as bit strings of some fixed length m . In this case, given a block code $C^{(n)}: \mathcal{X}^n \rightarrow \Sigma^m$ we say it has *rate* $r^{(n)} = \frac{m}{n}$.

Theorem 4.3 (Shannon's source coding theorem for block codes). *Let $(X_n)_{n \geq 1}$ be a stochastic process which satisfies the AEP with rate h . Suppose $r^{(n)}$ is a sequence of rates with $\lim_{n \rightarrow \infty} r^{(n)} = r$.*

- If $r < h$ then for any sequence of block codes $C^{(n)} : \mathcal{X}^n \rightarrow \{0, 1\}^m$ with rate $r^{(n)} \rightarrow r$, $\lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} > 0$;
- If $r > h$ then there exists some sequence of block codes $C^{(n)} : \mathcal{X}^n \rightarrow \{0, 1\}^m$ with rate $r^{(n)} \rightarrow r$ such that $\lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} = 0$.

Proof. Let us just sketch the ideas. The encoding in the second case is the same as in Theorem 4.2, we simply map the non-typical sequences to the error symbol \perp . It is a simple exercise to check that this satisfies the condition on the theorem.

For the first case, letting $X = (X_1, \dots, X_n)$, $Y = C^{(n)}(X)$ and $\hat{X} = g(Y)$, one could view the triple $X \rightarrow Y \rightarrow \hat{X}$ as a Markovian triple and use Fano's inequality.

We know that $p_{\text{err}}^{(n)} \geq \frac{H(X|Y)-1}{\log_2 |\mathcal{X}|}$, and since X determines Y ,

$$H(X|Y) = H(X, Y) - H(Y) = H(X) - H(Y).$$

Now, since Y is distributed on $\{0, 1\}^m$, $H(Y) \leq m$, and by assumption $\lim_{n \rightarrow \infty} \frac{1}{n} H(X) = h$. If $r = h - 4\varepsilon$, then for large enough n

$$H(X|Y) = H(X) - H(Y) \geq n(h - \varepsilon) - m = n(h - \varepsilon - r^{(n)}) \geq 2\varepsilon n.$$

Putting this together we see that

$$p_{\text{err}} \geq \frac{H(X|Y) - 1}{\log_2 |\mathcal{X}^n|} \geq \frac{2\varepsilon n - 1}{n \log_2 |\mathcal{X}|} \geq \frac{\varepsilon}{|\mathcal{X}|} > 0,$$

where the lower bound here does not depend on n , but just r, h and \mathcal{X} .

This shows that the error rate cannot tend to 0, but with a bit more care one can actually show it must be tending to 1. Roughly, one can show that the best decoding function g maps a bit string $w \in \{0, 1\}^m$ to the most likely preimage of w , which we denote by $\mathbf{x}(w)$. In particular, the error probability is then at least $1 - \sum_w p(\mathbf{x}(w))$. However, if $r < h$, then since the size of the typical set $A_\varepsilon^{(n)}$ is exponentially larger than 2^m , and p is approximately uniform on $A_\varepsilon^{(n)}$, the contribution to $\sum_w p(\mathbf{x}(w))$ from typical \mathbf{x} is negligible, and clearly the sum from non-typical \mathbf{x} is also negligible. \square

4.2 Variable length codes

However, Theorem 4.2 does not lead to efficient construction of codes, and the codes are not particularly useful in application. More useful are codes which assign codewords to each element of \mathcal{X} and encode a string (X_1, \dots, X_n) by concatenating the codewords. It is much easier to encode messages using such a code, and under certain conditions, also much easier to decode. However, we will see it is still possible to construct codes of this sort which are essentially as efficient as that of Theorem 4.2.

In what follows we will deal with more general alphabets than binary. Given a set Σ let us write Σ^* for the set of finite strings (words) of elements of Σ , i.e.,

$$\Sigma^* = \{w = a_1 a_2 \dots a_n : n \geq 0, a_i \in \Sigma\},$$

where $\ell(w) := n$ is the *length* of the word w . When $n = 0$, we have the empty word $\varepsilon \in \Sigma^*$, and we will write $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$ for the set of non-empty words. Given a word $w = a_1 \dots a_n$ for any $k \leq n$ we say $v = a_1 \dots a_k$ is a *prefix* of w .

Definition 4.4 (Source code). A *source code* is a mapping

$$C: \mathcal{X} \rightarrow \Sigma^+,$$

where $C(x)$ is the *codeword* of $x \in \mathcal{X}$. We define an extension of C to \mathcal{X}^+ , which we still denote by $C: \mathcal{X}^+ \rightarrow \Sigma^+$, by concatenation via

$$C(x_1 \dots x_k) = C(x_1) \dots C(x_k).$$

Given a source code and a discrete random variable X taking values in \mathcal{X} , the *expected code length* is defined as

$$L_C = \mathbb{E}(\ell(C(X))) = \sum_{x \in \mathcal{X}} \ell(C(x)) p_X(x).$$

Example 4.5. Let $\mathcal{X} = \{a, b, c, d\}$ and $\Sigma = \{0, 1\}$, and suppose X has distribution

$$p(a) = \frac{1}{2}, \quad p(b) = \frac{1}{4}, \quad p(c) = p(d) = \frac{1}{8}.$$

One possible source code would be

$$C(a) = 00, \quad C(b) = 10, \quad C(c) = 01, \quad C(d) = 11.$$

In this case all codewords have length 2, and so it is clear that $L_C = 2$.

However, a better code would be as follows:

$$C^*(a) = 0, \quad C^*(b) = 10, \quad C^*(c) = 110, \quad C^*(d) = 111,$$

where we can compute that

$$L_C = \frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{8} \cdot 3 + \frac{1}{8} \cdot 3 = \frac{7}{4}.$$

Of course, minimising L_C is a trivial optimisation problem—we simply assign the shortest codewords to the most likely elements of \mathcal{X} . However, the extension of C to Σ^+ might cause ambiguities when there are distinct sequences of elements (x_1, x_2, \dots, x_k) and $(x'_1, x'_2, \dots, x'_k) \in \mathcal{X}^+$ such that

$$C(x_1) \dots C(x_k) = C(x'_1) \dots C(x'_k).$$

Definition 4.6 (Unique decodability). (a) A source code $C: \mathcal{X} \rightarrow \Sigma^+$ is called *non-singular* if the mapping is injective, that is, if different elements have different codewords.

(b) C is called *uniquely decodable* if the extension $C: \mathcal{X}^+ \rightarrow \Sigma^+$ is non-singular.

(c) C is called *prefix-free*, or *instantaneous*, if no codeword is a prefix of another codeword, that is, whenever $x \neq y \in \mathcal{X}$ then $C(x)$ is not a prefix of $C(y)$.

It is clear that a uniquely decodable code must be non-singular. It is also easy to see that a prefix-free code is uniquely decodable—if $C(x_1) \dots C(x_k) = C(x_1 \dots x_k) = C(x'_1 \dots x'_m) = C(x'_1) \dots C(x'_m)$ then since $C(x_1)$ is not a prefix of $C(x'_1)$ and vice versa, it follows that $C(x'_1) = C(x_1)$ and hence $x_1 = x'_1$, since C is non-singular. The result then follows by induction.

Prefix-free codes are called instantaneous as they can be decoded in ‘real time’. If Alice encodes a word $x_1 \dots x_n$ by $C(x_1 \dots x_n)$ and transmits the codeword to Bob, then for any k Bob can decode the prefix $x_1 \dots x_k$ as soon as he has received the prefix $C(x_1) \dots C(x_k)$ of the codeword.

Example 4.7. Let $\mathcal{X} = \{a, b, c, d\}$ and $\Sigma = \{0, 1\}$ as before.

(a) The code

$$C(a) = 0, \quad C(b) = 010, \quad C(c) = 01, \quad C(d) = 10,$$

is non-singular, but it is not uniquely decodable, since $C(ad) = 010 = C(b)$.

(b) If C is a prefix code and \overleftarrow{C} is the code obtained by reversing each codeword, then \overleftarrow{C} is still uniquely decodable (since reversing a string is an involution on Σ^+). However, \overleftarrow{C} is not in general prefix-free.

For example, the code C^* from Example 4.5 is prefix-free, and so

$$C(a) = 0, \quad C(b) = 01, \quad C(c) = 011, \quad C(d) = 111,$$

is uniquely decodable, however it is not prefix-free, as $C(a) = 0$ is a prefix of $C(b) = 01$.

(c) Let

$$C(a) = 10, \quad C(b) = 00, \quad C(c) = 11, \quad C(d) = 110.$$

It can be shown (exercise) that C is uniquely decodable, but it is not *instantaneous*, in that there are arbitrarily long words $w \in \mathcal{X}^+$ such that Bob cannot decode any prefix of w before having received the entire codeword $C(w)$. For example $cb \dots b$ and $db \dots b$ both encode to $110 \dots 0$, with the only difference being the parity of the string of 0s, which Bob cannot discover until the last element is transmitted.

So, rather than asking to minimise the expected length of *any* code, it is reasonable to restrict our attention to codes that are uniquely decodable, so that sequences of codewords can be unambiguously decoded. Stronger still would be to insist that our code is prefix-free, in which case we should expect our code to have to be longer. Rather surprisingly, it turns out that the expected length of the shortest uniquely decodable code in fact coincides with the expected length of the shortest prefix-free code, and that both are controlled by the entropy of X .

A useful tool will be Kraft-McMillan inequality, which bounds from below the length of the codewords in a prefix-free code.

Lemma 4.8. [*Kraft-McMillan inequality*]

(1) Let $C: \mathcal{X} \rightarrow \Sigma^+$ be a prefix-free code with $|\Sigma| = D \geq 2$. Then

$$\sum_{x \in \mathcal{X}} D^{-\ell(C(x))} \leq 1.$$

(2) Let $\{\ell_x : x \in \mathcal{X}\}$ be a (multi)-set of numbers such that

$$\sum_{x \in \mathcal{X}} D^{-\ell_x} \leq 1$$

Then there exists a prefix-free code C such that $\ell(C(x)) = \ell_x$ for each $x \in \mathcal{X}$.

Remark 4.9. In fact, the bound in (1) can be shown to hold for uniquely decodable codes.

Proof. Let us start by showing (1). Let N be the maximum length of a codeword, that is, $N = \max_{x \in \mathcal{X}} \{\ell(C(x))\}$, and let $L_N \subseteq \Sigma^*$ be the words of length exactly N . It is clear that $|L_N| = |\Sigma|^N = D^N$.

For every $x \in \mathcal{X}$ let us define the set

$$W_x = \{w \in L_N : C(x) \text{ is a prefix of } w\} \subseteq L_N.$$

We note that for any $x, y \in \mathcal{X}$ the sets W_x and W_y are disjoint, since if $C(x)$ and $C(y)$ are both prefixes of a word w , then either $C(x)$ is a prefix of $C(y)$ or vice versa.

However, it is clear that $|W_x| = D^{N-\ell(C(x))}$ and hence

$$D^N = |L_N| \leq \sum_{x \in \mathcal{X}} |W_x| = \sum_{x \in \mathcal{X}} D^{N-\ell(C(x))}.$$

Dividing both sides by D^N leads to the desired inequality.

An alternative proof comes from choosing an element $X \in \Sigma^N$ uniformly at random and looking at $\mathbb{P}[C(x) \text{ a prefix of } X]$. This probability can be seen to be $D^{-\ell(C(x))}$, and the events are mutually exclusive, since no $C(x)$ is a prefix of $C(y)$. Hence $1 \geq \sum_x \mathbb{P}[C(x) \text{ a prefix of } X] = \sum_x D^{-\ell(C(x))}$.

We now prove (2). Let $n_j = |\{x : \ell_x = j\}|$, so that $\sum_{j \geq 0} n_j D^{-j} \leq 1$. We will choose our code ‘greedily’, defining the codewords $C(x)$ of length ℓ_x for all x such that $\ell_x = k$ for each $k \in \mathbb{N} \cup \{0\}$ in turn.

For $k = 1$ we note that $n_1 D^{-1} \leq 1$, and so it is always possible to choose n_1 codewords of length 1. Suppose we have already chosen n_ℓ codewords of length ℓ for each $\ell < k$. For each codeword $C(x)$ of length ℓ_x , there are exactly $D^{k-\ell_x}$ words w of length k such that $C(x)$ is a prefix of w . Hence, there are at least

$$|L_k| - \sum_{\substack{x \in \mathcal{X} \\ \ell_x < k}} D^{k-\ell_x} = D^k - \sum_{j=0}^{k-1} n_j D^{k-j} = D^k \left(1 - \sum_{j=0}^{k-1} n_j D^{-j} \right) \geq D^k \cdot n_k D^{-k} = n_k$$

words w of length k such that no $C(x)$ of length ℓ_x with $\ell_x < k$ is a prefix of w . Furthermore, since L_k itself is prefix-free, in particular so is this subset. Hence, it is always possible to choose inductively n_k codewords of length k .

□

One useful thing to notice is that Kraft's inequality allows us to rephrase the problem of finding codes of minimal expected length as an integer optimisation problem. Given the distribution $p: \mathcal{X} \rightarrow [0, 1]$ we wish to find $\ell = (\ell_x)_{x \in \mathcal{X}} \in \mathbb{N}_0^{\mathcal{X}}$ which

$$\text{minimises } \sum_{x \in \mathcal{X}} \ell_x p(x) \text{ subject to the constraint } \sum_{x \in \mathcal{X}} D^{-\ell_x} \leq 1.$$

This problem can be solved algorithmically in a number of ways.

For example, we can solve the corresponding continuous optimisation problem, noting that we may strengthen the constraint to $\sum_{x \in \mathcal{X}} D^{-\ell_x} = 1$ in this case, using Lagrange multipliers. This gives an approximate solution ℓ_x , and if you 'round up', the values $\lceil \ell_x \rceil$ will satisfy the constraint from Kraft's inequality and you can use the implicit algorithm therein to build a code which is close to optimal.

An alternative method uses entropy, and can also give a nearly matching upper bound. Since we are working over a general alphabet Σ of size D , which might not always be equal to 2, it makes sense to consider a slightly different notion of entropy, namely

$$H_D(X) = - \sum_{x \in \mathcal{X}} p(x) \log_D p(x) = - \frac{1}{\log_2 D} H(X).$$

Theorem 4.10. *[Source coding theorem for symbol codes] Let $C: \mathcal{X} \rightarrow \Sigma^+$ be a prefix-free source code to an alphabet Σ of size D , and let X be a discrete random variable taking values in \mathcal{X} . Then*

$$L_C \geq H_D(X),$$

with equality if and only if $p(x) = D^{-\ell(C(x))}$ for all $x \in \mathcal{X}$.

Conversely there exists a prefix-free code C such that

$$L_C \leq H_D(X) + 1.$$

Proof. By definition we have that

$$\begin{aligned} H_D(X) - L_C &= - \sum_{x \in \mathcal{X}} p(x) \log_D p(x) - \mathbb{E}(\ell(C(X))) \\ &= - \sum_{x \in \mathcal{X}} p(x) (\log_D p(x) + \ell(C(x))) \\ &= \sum_{x \in \mathcal{X}} p(x) \log_D \left(\frac{1}{D^{\ell(C(x))} p(x)} \right). \end{aligned}$$

Using Jensen's inequality (Theorem 2.31), it follows that

$$H_D(X) - L_C \leq \log_D \left(\sum_{x \in \mathcal{X}} \frac{p(x)}{D^{\ell(C(x))} p(x)} \right) = \log_D \left(\sum_{x \in \mathcal{X}} D^{-\ell(C(x))} \right),$$

and hence by Kraft's inequality (4.8) we can conclude that

$$H(X) - L_C \leq \log_D 1 = 0,$$

with equality if and only if we had equality in Kraft's inequality and Jensen's inequality, so that $\sum_{x \in \mathcal{X}} D^{-\ell(C(x))} = 1$ and $D^{\ell(C(x))}$ is proportional to $p(x)$, in which case they must be equal.

For the upper bound let us define, for every $x \in \mathcal{X}$,

$$\ell_x = \left\lceil \log_D \left(\frac{1}{p(x)} \right) \right\rceil.$$

It follows that

$$\sum_{x \in \mathcal{X}} D^{-\ell_x} \leq \sum_{x \in \mathcal{X}} p(x) = 1.$$

Hence by Kraft's inequality (Lemma 4.8) there exists a prefix-free code C such that $\ell(C(x)) = \ell_x$ for all $x \in \mathcal{X}$. However this code now satisfies

$$\begin{aligned} L_C &= \sum_{x \in \mathcal{X}} p(x) \ell(C(x)) \\ &= \sum_{x \in \mathcal{X}} p(x) \ell_x \\ &\leq \sum_{x \in \mathcal{X}} p(x) \left(\log_D \left(\frac{1}{p(x)} \right) + 1 \right) \\ &= H_D(X) + \sum_{x \in \mathcal{X}} p(x) \\ &= H_D(X) + 1. \end{aligned}$$

□

Remark 4.11. *Since Kraft's inequality holds for uniquely decodable codes, the first inequality in Theorem 4.10 holds for any uniquely decodable codes. In particular, Theorem 4.2 is optimal if we insist the code is uniquely decodable (and in fact, it is easy to check that the code we constructed there is even prefix-free).*

If instead we defined ℓ_x according to an incorrect distribution $q(x)$, so that

$$\ell_x = \left\lceil \log_D \left(\frac{1}{q(x)} \right) \right\rceil,$$

then $-\log_D q(x) + 1 \geq \ell_x \geq -\log_D q(x)$, and so the expected code length will satisfy

$$\begin{aligned} L_C &= \sum_{x \in \mathcal{X}} p(x) \ell_x \\ &\geq - \sum_{x \in \mathcal{X}} p(x) \log_D q(x) \\ &= - \sum_{x \in \mathcal{X}} p(x) \log_D \frac{q(x)}{p(x)} - \sum_{x \in \mathcal{X}} p(x) \log_D p(x) \\ &= H_D(p) + D_D(p \parallel q), \end{aligned}$$

and a similar argument will show that $L_C \leq H_D(p) + D_D(p \parallel q) + 1$.

This gives us a nice concrete interpretation for the relative entropy: it is essentially the penalty we pay when encoding if we incorrectly estimate the underlying distribution p as q .

If we are transmitting a sequence $(X_n)_{n \geq 1}$ of i.i.d. elements distributed according to X_1 , then we can see that the average number of symbols used to encode each element of \mathcal{X} is given by L_C ,

and so Theorem 4.10 leads to a code with the same asymptotic rate as that given by Theorem 4.2 (in the case where $D = |\Sigma| = 2$), since the entropy rate of an i.i.d. sequence is given by $H(X_1)$.

However, in general, if we are transmitting a sequence of elements from \mathcal{X} , which now come from some stochastic process $(X_n)_{n \geq 1}$, even if the process is time-homogeneous, so that each X_i has the same distribution, it might not be the case that a code that minimises the expected length of each individual codeword also minimises the length of an encoded message (X_1, \dots, X_n) of greater length. Moreover, we may be able to achieve a smaller expected length of codeword *per symbol transmitted*, so an encoding scheme with a smaller rate, if we group our symbols together and encode the elements of \mathcal{X}^n rather than of \mathcal{X} .

In this case, given a joint distribution (X_1, \dots, X_n) on \mathcal{X} and a code $C^{(n)}: \mathcal{X}^n \rightarrow \Sigma^+$, we could ask about the expected codeword length per symbol, or the rate

$$L^{(n)} = \mathbb{E} \left(\frac{1}{n} \ell(C^{(n)}(X_1, \dots, X_n)) \right).$$

If we insist that the code $C^{(n)}$ is prefix-free, then Theorem 4.10 implies that the optimal expected length L^* satisfies

$$H_D(X_1, \dots, X_n) \leq n \cdot L^{(n)} \leq H_D(X_1, \dots, X_n) + 1,$$

and hence

$$\frac{1}{n} H_D(X_1, \dots, X_n) \leq L^{(n)} \leq \frac{1}{n} H_D(X_1, \dots, X_n) + \frac{1}{n}.$$

Therefore, if the limit

$$h_D := \lim_{n \rightarrow \infty} \frac{1}{n} H_D(X_1, \dots, X_n)$$

exists, then we have a natural bound for this quantity.

However, this limit is just precisely, up to a multiplicative factor of $\log_2 D$, the entropy rate of the process $(X_n)_{n \geq 1}$. Hence we see that the code from Theorem 4.2 does indeed have an optimal rate, and we get another process by which we can construct such codes.

Theorem 4.12. *Let $(X_n)_{n \geq 1}$ be a stochastic process whose entropy rate h exists. Then the minimal expected rate of a prefix-free code satisfies*

$$\lim_{n \rightarrow \infty} L^{(n)} = h_D := \frac{h}{\log_2 D}.$$

4.3 Huffman Codes

Given a source X , Theorem 4.10 tells us that for any prefix-free source code C on an alphabet Σ of size D , $L_C \leq H_D(X)$, and conversely, gives us via Kraft's inequality (Lemma 4.8) a way to construct a prefix-free C such that $L_C \leq H_D(X) + 1$, however in general these codes will not be optimal.

However, Huffman gave a simple algorithm which, given a distribution p on \mathcal{X} with $|\mathcal{X}| \geq 2$, produces an optimal prefix-free binary code.

Whilst this works for all alphabet sizes, let us focus on the case $\Sigma = \{0, 1\}$. We can think of Σ^* as the *infinite binary tree*, whose root corresponds to the empty string ε , and where each vertex labelled w has two *children* labelled $w0$ and $w1$, which we call *siblings*. Given a string $w \in \Sigma^+$, let us write w' for its sibling.

When is a string w a prefix of another string v ? Precisely when the unique path from the root ε to v passes through w . In particular, if we have a prefix-free code, we can build a finite subtree T whose leaves are precisely the codewords by taking the union of these paths from the codewords to the root. For any other vertex in T , at least one child also lies in T .

Huffman's algorithm works by constructing, for each distribution p on \mathcal{X} , an appropriate subtree $T(p)$ whose leaves are labelled by the elements of \mathcal{X} .

The algorithm is *recursive*: if $N := |\mathcal{X}| = 2$, then we take $T(p)$ to be the tree consisting of the root and its two children.

Otherwise, we start by sorting the elements $\mathcal{X} = \{x_1, \dots, x_N\}$ such that $p(x_1) \geq p(x_2) \geq \dots \geq p(x_N)$. Note that this ordering is not necessarily unique, and this may change the output of the algorithm.

We define a new set $\mathcal{X}' = \{x'_1, \dots, x'_{N-1}\}$, and a probability distribution p' on \mathcal{X}' , by

$$p'(x'_i) = \begin{cases} p(x_i) & \text{if } i \leq N-2, \\ p(x_{N-1}) + p(x_N) & \text{if } i = N-1. \end{cases}$$

We build $T(p'(x'_1), \dots, p'(x'_{N-1}))$, and we form $T(p(x_1), \dots, p(x_N))$ by taking the leaf labelled x'_{N-1} and adding its two children as leaves, labelled x_{N-1} and x_N . All other leaves labelled x'_i with $i \leq N-2$ we label with x_i .

We call a code constructed in this way a *Huffman code*.

Alternatively we can think of building the tree starting from the leaves up—we start with an independent set of vertices labelled $p(x_1)$ to $p(x_N)$, and recursively we choose the two vertices in the forest which have no parent and have the smallest labels p_1 and p_2 . Then we add a new vertex, which is joined as a parent to these two vertices, and has label $p_1 + p_2$. We continue until there is a unique vertex in the forest with no parent. By construction this graph is a binary tree, and by choosing an arbitrary $\{0, 1\}$ -labelling of the edges from a vertex to its children we can assign to each leaf a string in $\{0, 1\}^+$, giving us a prefix-free code.

Example 4.13. Suppose $\mathcal{X} = \{x_1, x_2, x_3, x_4, x_5\}$ and

$$p(x_1) = 0.4, \quad p(x_2) = 0.2, \quad p(x_3) = 0.15, \quad p(x_4) = 0.15, \quad p(x_5) = 0.1.$$

In the first step, the vertices with the smallest labels are x_4 and x_5 , and so we would merge them to a new vertex, which we might call $x_{4,5}$ with label $p(x_{4,5}) = p(x_4) + p(x_5) = 0.25$. Now the vertices without parents are x_1, x_2, x_3 and $x_{4,5}$, with labels

$$p(x_1) = 0.4, \quad p(x_{4,5}) = 0.25, \quad p(x_2) = 0.2, \quad p(x_3) = 0.15,$$

and so the vertices with the smallest labels are x_2 and x_3 . In the next step, we merge them to a new vertex $x_{2,3}$ with label $p(x_{2,3}) = p(x_2) + p(x_3) = 0.35$. Now the vertices without parents

are $x_1, x_{2,3}$ and $x_{4,5}$ with labels

$$p(x_1) = 0.4, \quad p(x_{2,3}) = 0.35, \quad p(x_{4,5}) = 0.25,$$

and so the vertices with the smallest labels are $x_{2,3}$ and $x_{4,5}$. In the next step, we merge them to a new vertex $x_{2,3,4,5}$ with label $p(x_{2,3,4,5}) = p(x_{2,3}) + p(x_{4,5}) = 0.6$. Now the vertices without parents are x_1 and $x_{2,3,4,5}$ with labels

$$p(x_{2,3,4,5}) = 0.6, \quad p_1 = 0.4,$$

and so the vertices with the smallest labels are $x_{2,3,4,5}$ and x_1 . In the last step, we merge them to a new vertex $x_{1,2,3,4,5}$ with label $p(x_{1,2,3,4,5}) = 1$.

If we label the edges so that the edge labelled 1 goes to the child with the smaller label, then we end up with the code $C: \mathcal{X} \rightarrow \Sigma^*$ given by

$$C(x_1) = 1, \quad C(x_2) = 000, \quad C(x_3) = 001, \quad C(x_4) = 010, \quad C(x_5) = 011.$$

In the example above we can calculate that the expected code length is then

$$L_C = 0.4 + 3 \cdot (0.2 + 0.15 + 0.15 + 0.1) = 2.2,$$

and the entropy of p is ≈ 2.15 . So, this code is definitely close to the theoretical limit of $H(p)$, and in fact, it can be shown that no other binary code will do better than the Huffman code. Given a source X , let us say a prefix-free binary code C is *optimal* if $L_C \leq L_{C'}$ for any other prefix-free binary code C' .

Theorem 4.14. *Huffman codes are optimal binary codes.*

Let us start by showing the following lemma.

Lemma 4.15. *Let X be a source on \mathcal{X} , and let $C: \mathcal{X} \rightarrow \{0,1\}^+$ be an optimal prefix-free binary code. Then*

(1) *If $p(x) > p(y)$, then $\ell(C(x)) \leq \ell(C(y))$;*

(2) *Let $\ell_{\max} = \max\{\ell(w) : w \in C(\mathcal{X})\}$ and*

$$W = \{w \in C(\mathcal{X}) : \ell(w) = \ell_{\max}\}.$$

Then for every $w \in W$, its sibling $w' \in W$.

Proof. The first is relatively clear—if we build a new code C' by swapping the codewords for x and y , then

$$\begin{aligned} 0 \leq L_{C'} - L_C &= \sum_{x \in \mathcal{X}} p(x) \ell(C'(x)) - p(x) \ell(C(x)) \\ &= p(x) (\ell(C(y)) - \ell(C(x))) + p(y) (\ell(C(x)) - \ell(C(y))) \\ &= (p(x) - p(y)) (\ell(C(y)) - \ell(C(x))). \end{aligned}$$

However, $p(x) - p(y) > 0$, and hence $\ell(C(y)) - \ell(C(x)) \geq 0$.

For the second suppose $w \in W$ has parent v , so that $w \in \{v0, v1\}$, and suppose for contradiction that $w' \notin C(\mathcal{X})$. Then, since no codeword has length $> \ell(v) + 1$, the only codeword which v is a prefix of is w , and hence we may form a new prefix-free code by taking the $x \in \mathcal{X}$ such that $C(x) = w$, setting $C'(x) = v$ and setting $C'(y) = C(y)$ for all $y \neq x$. However then

$$L'_C - L_C = p(x)(\ell(C'(x)) - \ell(C(x))) = -p(x) < 0.$$

□

Proposition 4.16. *Let X be a source on \mathcal{X} , and let $C: \mathcal{X} \rightarrow \{0,1\}^*$ be an optimal prefix-free binary code. If $\mathcal{X} = \{x_1, \dots, x_N\}$ is some ordering of \mathcal{X} such that $p(x_1) \geq p(x_2) \dots \geq p(x_N)$ then there is some permutation π on \mathcal{X} such that $C' = C \circ \pi$ is an optimal prefix-free binary code such that*

$$C'(x_{N-1}), C'(x_N) \in W \quad \text{and they are siblings,}$$

where W is as in Lemma 4.15 (2).

Proof. We first note that if $p(x) = p(y)$, or $\ell(C(x)) = \ell(C(y))$, then exchanging $C(x)$ and $C(y)$ does not change L_C .

Now, by Lemma 4.15 (2) there exist x_i, x_j such that $C(x_i)$ and $C(x_j)$ are siblings in W . Let us suppose without loss of generality that $i < j$. Then $p(x_j) \leq p(x_N)$ and $\ell(C(x_j)) \leq \ell(C(x_N))$ and so by Lemma 4.15 (1) that either $p(x_j) = p(x_N)$ or $\ell(C(x_j)) = \ell(C(x_N))$. A similar argument holds for x_i and x_{N-1} .

Hence, we can take the permutation which maps x_j to x_N and x_i to x_{N-1} (which may be the identity), which will produce a code C' with $L_C = L_{C'}$, so that C' is optimal. □

We call codes satisfying the conclusion of Proposition 4.16 *canonical* for the ordering $\mathcal{X} = \{x_1, \dots, x_N\}$.

At this point we are ready to prove that Huffman codes are optimal.

Proof of Theorem 4.14. We prove the theorem by induction on $N = |\mathcal{X}|$. For $N = 2$ the Huffman codes give $C(x_1) = 0$, $C(x_2) = 1$, which is clearly optimal.

Suppose we know that Huffman codes on sets of size $N - 1$ are optimal. Let the set $\mathcal{X} = \{x_1, \dots, x_N\}$ be ordered such that $p(x_1) \geq \dots \geq p(x_N)$. Let us define, as in the construction of Huffman codes, a new probability distribution p' on the set $\mathcal{X}' = \{x'_1, \dots, x'_{N-2}, x'_{N-1}\}$ given by $p'(x'_i) = p(x_i)$ for all $i \leq N - 2$ and $p'(x'_{N-1}) = p(x_{N-1}) + p(x_N)$. Let C'_{N-1} be a Huffman code for p' on \mathcal{X}' , which by assumption is optimal.

By construction, the Huffman code C_N for p on \mathcal{X} is then given by

$$C_N(x_j) = \begin{cases} C'_{N-1}(x'_j) & \text{if } j \leq N - 2, \\ C'_{N-1}(x'_{N-1})0 & \text{if } j = N - 1, \\ C'_{N-1}(x'_{N-1})1 & \text{if } j = N. \end{cases}$$

We can calculate

$$L_{C_N} = L_{C'_{N-1}} + p(x_{N-1}) + p(x_N).$$

Let \tilde{C}_N be a canonical code for the ordering $\mathcal{X} = \{x_1, \dots, x_N\}$. Let w be the parent of $\tilde{C}_N(x_N)$ and $\tilde{C}_N(x_{N-1})$. We can define a new code \tilde{C}'_{N-1} via the ‘reverse’ of the Huffman construction, that is,

$$\tilde{C}'_{N-1}(x'_j) = \begin{cases} \tilde{C}_N(x_j) & \text{if } j \leq N-2, \\ w & \text{if } j = N-1. \end{cases}$$

In this case we can calculate

$$\begin{aligned} L_{\tilde{C}'_{N-1}} &= L_{\tilde{C}_N} - p(x_N) - p(x_{N-1}) \\ &\leq L_{C_N} - p(x_N) - p(x_{N-1}) \\ &= L_{C'_{N-1}}, \end{aligned}$$

using the assumption that \tilde{C}_N was optimal. However, since C'_{N-1} is optimal by the induction hypothesis, we must have equality throughout, and so in particular, $L_{\tilde{C}_N} = L_{C_N}$ and C_N is also optimal.

However, there is a delicate issue with this proof, which is skipped over in many computer science textbooks—we don’t necessarily know that \tilde{C}_N exists! Indeed, Proposition 4.16 says that if an optimal code exists, then a canonical code exists, however, it is not a priori obvious that an optimal code exists at all. That is, the infimum

$$\inf\{L_C : C \text{ a prefix-free code}\}$$

might not be an attained minimum.

This can be solved in a number of ways, perhaps the simplest is the following: if we view a prefix code as a binary tree, then for any code C that has a vertex with a unique child, there is a code C' of strictly shorter expected length with fewer vertices (exercise). Note that any tree with $|\mathcal{X}|$ leaves in which every vertex has either two or zero children has depth at most $|\mathcal{X}|$ (exercise).

In particular, for any prefix-free code C , there is a tree (corresponding to a code C') of depth at most $|\mathcal{X}|$ with $L_{C'} \leq L_C$, and hence

$$\inf\{L_C : C \text{ a prefix-free code}\} = \min\{L_C : C \text{ a prefix-free code with } \ell_{\max} \leq |\mathcal{X}|\},$$

where the latter is an attained minimum, since it is over a finite set. □


5 Information Channels

A *channel* is a way to model the transmission of some message. We have some set \mathcal{X} of messages which are to be transmitted, potentially in some encrypted form, and a set \mathcal{Y} of possible outputs, messages received by the other party, where the output might not depend deterministically on the message due to some inherent *noise* in the channel, which might randomly change the output, or even some randomness in the encryption process.

Definition 5.1 (Discrete channel). A *discrete (memoryless) channel*

$$\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$$

consists of two finite sets \mathcal{X} and \mathcal{Y} and a stochastic transition matrix $P = (p(y|x))_{x \in \mathcal{X}, y \in \mathcal{Y}}$.



That is, the rows of P are indexed by \mathcal{X} , the columns by \mathcal{Y} and each row is a conditional probability distribution $p(\cdot|x)$ on \mathcal{Y} , that is $p(y|x) \geq 0$ for all $y \in \mathcal{Y}$ and

$$\sum_{y \in \mathcal{Y}} p(y|x) = 1.$$

We think of the distribution $p(\cdot|x)$ as being the distribution of the output of the channel, the received message $y \in \mathcal{Y}$ when x is the input.

Example 5.2. (a) The *binary symmetric channel* has $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, where each message results in the wrong output with probability ε . Thus $P = \begin{pmatrix} 1 - \varepsilon & \varepsilon \\ \varepsilon & 1 - \varepsilon \end{pmatrix}$.

(b) The *binary erasure channel* has $\mathcal{X} = \{0, 1\}$ and $\mathcal{Y} = \{0, 1, \perp\}$ and each message has $(1 - \varepsilon)$ chance of being transmitted correctly and ε chance of being ‘lost’, resulting in the output \perp . Thus $P = \begin{pmatrix} 1 - \varepsilon & 0 & \varepsilon \\ 0 & 1 - \varepsilon & \varepsilon \end{pmatrix}$.

Definition 5.3 (Channel extension). Given a channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ the *nth channel extension* is the channel

$$\mathcal{C}^n = (\mathcal{X}^n, P_n, \mathcal{Y}^n)$$

where

$$(P_n)_{\mathbf{x}, \mathbf{y}} = p_n(\mathbf{y}|\mathbf{x}) = p_n(y_1, \dots, y_n | x_1, \dots, x_n) = \prod_{k=1}^n p(y_k | x_k).$$

In other words, the *nth* channel extension is the channel we get by sending n consecutive, independent messages over the channel \mathcal{C} .

Typically, we are interested in the behaviour of the channel when the input arrives as some \mathcal{X} -valued random variable X , with some distribution p_X . In this case the output Y is also a

random variable, which inherits the randomness from X , as well as some of the randomness inherent in the channel described by P .

For a fixed channel, and so a fixed P , we might hope to choose the input distribution p_X in some optimal way. Let $\mathcal{M}(\mathcal{X})$ be the collection of all probability distributions on \mathcal{X} . In other words, if $\mathcal{X} = \{x_1, \dots, x_n\}$, then we can think of \mathcal{M} as consisting of all vectors $(p_1, \dots, p_n) \in \mathbb{R}^n$ with non-negative entries and sum 1, with $p(x_i) = p_i$.

One measure of the ‘quality’ of the channel is how well the message is preserved, and one way to measure this would be to measure how much information about the message X is contained in the random variable Y , motivating the following definition.

Definition 5.4 (Channel capacity). Given a channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ the *channel capacity* is defined as

$$\text{cap}(\mathcal{C}) = \max\{I(X ; Y) : p_X \in \mathcal{M}(\mathcal{X})\}.$$

Remark 5.5. Note that p_X and P together determine $p_{X,Y}$ and hence p_Y and $I(X ; Y)$. Indeed,

$$p_{X,Y}(x, y) = p_X(x)p_{Y|X}(y|x) = p_X(x)(P)_{x,y} \quad \text{and} \quad p_Y(y) = \sum_{x \in \mathcal{X}} p_{X,Y}(x, y).$$

Hence, we can consider the function from $\mathcal{M}(\mathcal{X})$ to \mathbb{R} given by $p_X \mapsto I(X ; Y)$. This is then a continuous function on a compact set $\mathcal{M}(\mathcal{X}) \subseteq \mathbb{R}^{\mathcal{X}}$ and so it indeed achieves some maximum, which is then the channel capacity. Note that this maximum is not necessarily achieved by a unique distribution p_X !

Note that, by Corollary 2.33 and Lemma 2.22,

$$0 \leq I(X ; Y) = H(X) - H(X | Y) \leq H(X) \leq \log_2 |\mathcal{X}|,$$

and so

$$0 \leq \text{cap}(\mathcal{C}) \leq \log_2 |\mathcal{X}|.$$

Example 5.6. (a) *Noiseless binary channel* : Suppose we take the binary symmetric channel with $\varepsilon = 0$, that is $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and $P = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$.

It is easy to verify that in this case $X = Y$ and so for any p_X , $I(X ; Y) = I(X ; X) = H(p_X)$, and in particular

$$\text{cap}(\mathcal{C}) = \max\{H(p_X) : p_X \in \mathcal{M}(\mathcal{X})\} = \log_2 |\mathcal{X}| = 1,$$

which is achieved only for the uniform distribution $(1/2, 1/2)$. Hence $\text{cap}(\mathcal{C}) = 1$.

(b) *Channel with non-overlapping outputs* : More generally, for any transition matrix P where X is determined by Y , we have by Lemma 2.8

$$I(X ; Y) = H(X) - H(X|Y) = H(X) \leq \log_2 |\mathcal{X}|,$$

where equality is again achieved uniquely by the uniform distribution on X . Hence $\text{cap}(\mathcal{C}) = \log_2 |\mathcal{X}|$.

- (c) *Noisy typewriter* : Suppose $\mathcal{X} = \{x_1, \dots, x_{2N}\}$ is a set of letters on a rather improbable circular typewriter, where when I go to type a letter x_i , I miss and hit the letter x_{i+1} with probability $\frac{1}{2}$ (with addition mod $2N$). Hence $\mathcal{Y} = \mathcal{X}$, and the transition matrix is

$$P = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & \dots & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & \dots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ \frac{1}{2} & 0 & 0 & \dots & \frac{1}{2} \end{pmatrix}.$$

Given any distribution p_X we can calculate

$$H(Y|X) = \sum_{j=1}^{2N} p_X(x_j) H(Y | X = x_j) = \sum_{j=1}^{2N} p_X(x_j) H(1/2, 1/2) = 1,$$

and so

$$I(X ; Y) = H(Y) - H(Y | X) = H(Y) - 1 \leq \log_2 2N - 1 = \log_2 N,$$

and equality is achieved whenever p_Y is uniform. Hence $\text{cap}(\mathcal{C})$ will be $\log_2 N$ if there is some p_X such that p_Y is uniform, and it is easy to verify that when p_X is uniform, so is p_Y . Hence $\text{cap}(\mathcal{C}) = \log_2 N$.

However, in this case there are multiple optimal distributions. For example if X is uniformly distributed on the odd elements, or uniformly distributed on the even elements, then Y is again uniformly distributed on \mathcal{Y} . More generally, any convex combination of these two distributions achieves the optimal capacity.

A small remark here is that, when X is distributed on the even or odd elements, then this example reduces to a channel with non-overlapping outputs.

- (d) *Binary symmetric channel* : Recall that $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and $P = \begin{pmatrix} 1 - \varepsilon & \varepsilon \\ \varepsilon & 1 - \varepsilon \end{pmatrix}$.

In this case we can again compute

$$\begin{aligned} I(X ; Y) &= H(Y) - H(Y | X) \\ &= H(Y) - H(\varepsilon, 1 - \varepsilon) \\ &\leq 1 - H(\varepsilon, 1 - \varepsilon), \end{aligned}$$

and it is easy to verify that if X is uniformly distributed, then so is Y , and so equality is achieved. Hence $\text{cap}(\mathcal{C}) = 1 - H(\varepsilon, 1 - \varepsilon)$.

- (e) *Binary erasure channel* : Recall that $\mathcal{X} = \{0, 1\}$ and $\mathcal{Y} = \{0, 1, \perp\}$ and $P = \begin{pmatrix} 1 - \varepsilon & 0 & \varepsilon \\ 0 & 1 - \varepsilon & \varepsilon \end{pmatrix}$.

It is easy to see that, for any distribution p_X ,

$$\begin{aligned} \mathbb{P}[Y = \perp] &= \mathbb{P}[X = 0] \cdot \mathbb{P}[Y = \perp | X = 0] + \mathbb{P}[X = 1] \cdot \mathbb{P}[Y = \perp | X = 1] \\ &= \varepsilon p_X(1) + \varepsilon p_X(0) \\ &= \varepsilon \end{aligned}$$

and for any $i \in \{0, 1\}$

$$\mathbb{P}[X = i | Y = \perp] = \frac{\mathbb{P}[X = i, Y = \perp]}{\mathbb{P}[Y = \perp]} = p_X(i).$$

Furthermore, $(X|Y = i)$ is constant for $i \in \{0, 1\}$. Hence

$$\begin{aligned} I(X ; Y) &= H(X) - H(X|Y) \\ &= H(X) - \mathbb{P}[Y = 1]H(X|Y = 1) - \mathbb{P}[Y = 0]H(X|Y = 0) - \mathbb{P}[Y = \perp]H(X|Y = \perp) \\ &= H(X) - 0 - 0 - \varepsilon H(X) \\ &= (1 - \varepsilon)H(X) \leq 1 - \varepsilon. \end{aligned}$$

Again, it is easy to verify that equality is achieved only when $H(X) = 1$, and so when p_X is uniform. Hence $\text{cap}(\mathcal{C}) = 1 - \varepsilon$.

- (f) *Non-symmetric binary channel* : Suppose $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, but now the probability of error is different for 0 and 1, so that $P = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}$.

The distribution of X can be described by a single parameter $p_X = (1 - t, t)$, and so for fixed α and β the mutual information $I(X ; Y)$ is a (continuous) function of t , where $t \in [0, 1]$, which can be maximised using, for example, Lagrange multipliers (exercise).

Definition 5.7 (Weakly symmetric channel). A channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ is called *weakly symmetric* if

- (i) All rows $p(\cdot|x)$ for $x \in \mathcal{X}$ are permutations of each other;
- (ii) All columns $p(y|\cdot)$ for $y \in \mathcal{Y}$ have the same (constant) sum.

Proposition 5.8. If $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ is a weakly symmetric channel, then

$$\text{cap}(\mathcal{C}) = \log_2 |\mathcal{Y}| - H(p(\cdot|x_0)),$$

where $x_0 \in \mathcal{X}$ is arbitrary.

Proof. By definition, for any distribution p_X

$$\begin{aligned} I(X ; Y) &= H(Y) - H(Y|X) \\ &= H(Y) - \sum_{x \in \mathcal{X}} p_X(x) H(p(\cdot|x)) \end{aligned}$$

However, by Property (i) $p(\cdot|x) = p(\cdot|x_0)$ for all $x \in \mathcal{X}$. Hence

$$\begin{aligned} I(X ; Y) &= H(Y) - \sum_{x \in \mathcal{X}} p(x) H(p(\cdot|x)) \\ &= H(Y) - \sum_{x \in \mathcal{X}} p(x) H(p(\cdot|x_0)) \\ &\leq \log_2 |\mathcal{Y}| - H(p(\cdot|x_0)), \end{aligned}$$

with equality when Y is uniformly distributed.

On the other hand, by Property (ii), there is some constant c such that for all $y \in \mathcal{Y}$, $\sum_{x \in \mathcal{X}} p(y|x) = c$. In particular, if p_X is uniform, then for all $y \in \mathcal{Y}$

$$p_Y(y) = \sum_{x \in \mathcal{X}} p_X(x) p(y|x) = \frac{c}{|\mathcal{X}|}.$$

In particular, since $\sum_{y \in \mathcal{Y}} p_Y(y) = 1$, it follows that $c = \frac{|\mathcal{X}|}{|\mathcal{Y}|}$. Hence, if p_X is uniform, then for all $y \in \mathcal{Y}$

$$p_Y(y) = \frac{1}{|\mathcal{Y}|},$$

and p_Y is also uniform. Hence equality is achieved and $\text{cap}(\mathcal{C}) = \log_2 |\mathcal{Y}| - H(p(\cdot|x_0))$ as claimed. \square

Another nice property of the channel capacity is that the capacity of the n th channel extension is determined by the channel itself.

Lemma 5.9. *Let $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ be a channel and let \mathcal{C}^n be the n th channel extension. Then*

$$\text{cap}(\mathcal{C}^n) = n \cdot \text{cap}(\mathcal{C}).$$

Proof. Given a sequence (X_1, \dots, X_n) of independent random variables, the output of the channel \mathcal{C}^n with input (X_1, \dots, X_n) is given by a sequence (Y_1, \dots, Y_n) of independent random variables.

In particular, for any input (X_1, \dots, X_n) , by the chain rule (Theorem 2.13)

$$\begin{aligned} I(X_1, \dots, X_n ; Y_1, \dots, Y_n) &= H(Y_1, \dots, Y_n) - H(Y_1, \dots, Y_n | X_1, \dots, X_n) \\ &= \sum_{i=1}^n H(Y_i) - \sum_{i=1}^n H(Y_i | X_1, \dots, X_n, Y_1, \dots, Y_{i-1}) \\ &= \sum_{i=1}^n H(Y_i) - \sum_{i=1}^n H(Y_i | X_i) \\ &= \sum_{i=1}^n I(X_i ; Y_i) \\ &\leq n \cdot \text{cap}(\mathcal{C}^n) \end{aligned}$$

On the other hand, there is some random variable X distributed on \mathcal{X} such that $\text{cap}(\mathcal{C}) = I(X ; Y)$. If we take our input (X_1, \dots, X_n) to be i.i.d. each with distribution p_X , then the output (Y_1, \dots, Y_n) is also i.i.d. with distribution Y , and hence, by the same argument as above

$$I(X_1, \dots, X_n ; Y_1, \dots, Y_n) = n \cdot I(X ; Y) = n \cdot \text{cap}(\mathcal{C}).$$

Hence

$$\text{cap}(\mathcal{C}^n) = n \cdot \text{cap}(\mathcal{C}).$$

\square

5.1 Shannon's channel coding theorem

Let us try to give a practical meaning to the channel capacity.

Suppose Alice wishes to transmit a message from some input set \mathcal{W} to Bob through a noisy channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$. For example, $\mathcal{W} \subseteq \mathcal{X}^+$ might be some set of phrases over the alphabet

\mathcal{X} and Alice will independently transmit the symbols one by one through the channel, where the output $\mathcal{Y} = \mathcal{X}$ that Bob receives is a symbol that may or may not agree with the transmitted symbol.

Now, Bob would like to recreate the message, so he should have some function $g: \mathcal{Y}^+ \rightarrow \mathcal{W}$ which represents his *guess* of what the input was, given the observed output.

Since the channel is noisy, there is some chance Bob's guess will be incorrect. We could reduce this chance by sending a longer sequence of symbols, perhaps by sending the entire message twice, or via some more complicated *encoding* scheme. However, this comes then at the cost of a longer transmission. Ideally we would like to keep this error of failure small, using as few transmissions as possible.

Definition 5.10 ((M, n) -codes). An (M, n) -code for the channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ consists of the following:


- A set \mathcal{W} of *messages*, with $|\mathcal{W}| = M$,
- A mapping $\mathbf{x}^{(n)}: \mathcal{W} \rightarrow \mathcal{X}^n$, which we call the *codebook*,
- A function $g: \mathcal{Y} \rightarrow \widehat{\mathcal{W}}$, where $\widehat{\mathcal{W}} = \mathcal{W}$ or $\widehat{\mathcal{W}} = \mathcal{W} \cup \{\perp\}$, with the *erasure* (or dummy) symbol \perp .

Remark 5.11. Note the word *code* here does not have the same meaning as in Section 4.

The *rate* of an (M, n) -code is given by

$$R = \frac{\log_2 M}{n}.$$

The rate is then measured in bits per transmission.



For example, if we have a binary channel, with $\mathcal{X} = \{0, 1\}$, then in order to encode the elements of \mathcal{W} by distinct *codewords*, binary sequences of length n , we would need $n = \lceil \log_2 M \rceil$. In this case the rate would be equal to 1.

However, since these binary sequences may be corrupted by the channel, the likelihood that Bob's guess is correct will be closely related to the probability that any bit is incorrectly transmitted, which may be very large. Hence, in order to make more accurate guesses, it might be necessary to take a larger n , and so longer transmissions, and encode the messages in some *robust* manner, at the expense of decreasing the rate.

Definition 5.12 (Probability of error). Given a channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ and an (M, n) -code $(\mathcal{W}, \mathbf{x}^{(n)}, g)$ we define the *conditional probability of error*, given $w \in \mathcal{W}$, as the probability λ_w

that if we encode the message w , transmit the encoded message across the channel and then decode, that we do not recover w . That is,

$$\lambda_w^{(n)} = \sum_{\mathbf{y} \in \mathcal{Y}^n : g(\mathbf{y}) \neq w} p(\mathbf{y} | \mathbf{x}^{(n)}(w)).$$

The *maximal probability of error* is defined as

$$\lambda_{\max}^{(n)} = \max \left\{ \lambda_w^{(n)} : w \in \mathcal{W} \right\}.$$

The *average probability of error* is defined as

$$p_{\text{err}}^{(n)} = \frac{1}{M} \sum_{w \in \mathcal{W}} \lambda_w^{(n)}.$$

In the above, it can be useful to think of a random variable W distributed uniformly on \mathcal{W} (which is independent of the channel). The codebook transforms W into a random vector $\mathbf{x}^{(n)}(W) = (X_1, \dots, X_n)$, which is then transmitted through the channel, with output the random vector (Y_1, \dots, Y_n) . This output is then decoded as $\widehat{W} = g(Y_1, \dots, Y_n)$, which is a random variable distributed on $\widehat{\mathcal{W}}$. In this case we can express

$$\lambda_w^{(n)} = \mathbb{P} \left[\widehat{W} \neq w \mid W = w \right] \quad \text{and} \quad p_{\text{err}} = \mathbb{P} \left[\widehat{W} \neq W \right].$$

Note that in the above

$$W \rightarrow (X_1, \dots, X_n) \rightarrow (Y_1, \dots, Y_n) \rightarrow \widehat{W}$$

is a Markovian quadruple.

The specific question we will be interested in is how *efficiently*, in terms of the rate of the code, we can achieve (arbitrarily) small maximum error probability. That is, given $\varepsilon > 0$, for what rate R can we achieve $\lambda_{\max}^{(n)} \leq \varepsilon$?

It is not a priori obvious that we would not need increasing rate, as a function of ε , to achieve smaller and smaller error probabilities. Indeed, if we fix M and n , then there are only finitely many choices of codebook $\mathbf{x}^{(n)}$ and decoding function g , and (for non-trivial channels) each leads to a *strictly positive* maximum probability of error $\lambda_{\max}^{(n)}$. In particular, the minimum over all (M, n) -codes of $\lambda_{\max}^{(n)}$ will also be strictly positive, and so if we fix M and n we cannot reduce the error probability arbitrarily.

However, it will turn out that for certain rates, if we allow the length of our codewords to grow, we can achieve *any* maximum error probability, however small.

Definition 5.13 (Achievable rates). A real number $R > 0$ is an *achievable rate* for the channel \mathcal{C} if there is a sequence of (M_n, n) -codes such that the rate $R_n = \frac{\log_2 M_n}{n}$ and the maximal probability of error $\lambda_{\max}^{(n)}$ satisfy

$$\lim_{n \rightarrow \infty} R_n = R \quad \text{and} \quad \lim_{n \rightarrow \infty} \lambda_{\max}^{(n)} = 0.$$

In other words, R is achievable if for any $\varepsilon > 0$ there is some (M_n, n) -code with rate $R_n > R - \varepsilon$ and $\lambda_{\max}^{(n)} < \varepsilon$.

It turns out that the capacity of a channel controls the achievable rates.

Definition 5.14. The *achievable capacity* of a channel \mathcal{C} is defined as

$$R^* = \sup\{R: R \text{ is an achievable rate for } \mathcal{C}\}.$$

Remark 5.15. We note that this supremum is in fact an attained maximum. Indeed, there is some sequence $R^{(k)}$ of achievable rates such that $\lim_{k \rightarrow \infty} R^{(k)} = R^*$, and since each $R^{(k)}$ is achievable, there is a sequence of $(M_n^{(k)}, n)$ -codes whose rate tends to $R^{(k)}$ and whose maximum probability of error tends to 0.

In this case we can build a ‘diagonal’ sequence of (M, n) -codes, whose rates will tend to R^* , and if we choose the sequence carefully, we can guarantee that the maximum probability of error also tends to 0.

Explicitly, suppose that for each k there is some sequence of $(M_n^{(k)}, n)$ -codes with rate $R_n^{(k)}$ such that $\lim_{n \rightarrow \infty} \lambda_{\max}^{(n)} \rightarrow 0$ and $\lim_{n \rightarrow \infty} R_n^{(k)} = R^{(k)}$.

Given $\varepsilon > 0$, we can choose K sufficiently large such that for all $k \geq K$, $|R^{(k)} - R^*| \leq \frac{\varepsilon}{2}$ and given some fixed $k_0 \geq K$ we can choose N such that for all $n \geq N$, $|R_n^{(k_0)} - R^{(k_0)}| \leq \frac{\varepsilon}{2}$ and $\lambda_{\max}^{(n)} \leq \varepsilon$.

Hence, for some fixed $n_0 \geq N$, we have that the $(M_{n_0}^{(k_0)}, n_0)$ code with rate $R_{n_0}^{(k_0)}$ is such that

$$|R_{n_0}^{(k_0)} - R^*| \leq |R_{n_0}^{(k_0)} - R^{(k_0)}| + |R^{(k_0)} - R^*| \leq \varepsilon$$

and $\lambda_{\max}^{(n)} \leq \varepsilon$.

Theorem 5.16 (Shannon’s channel coding theorem). *For any channel \mathcal{C} its achievable capacity R^* is equal to the channel capacity $\text{cap}(\mathcal{C})$.*

We will find that it is relatively easy to show that $R^* \leq \text{cap}(\mathcal{C})$, that is, every achievable rate R satisfies $R \leq \text{cap}(\mathcal{C})$. It will be rather more difficult to show that if $R < \text{cap}(\mathcal{C})$ then R is achievable, since to do so we will have to construct sequences of (M_n, n) -codes with rate tending to R .

Shannon’s original proof, while mathematically ingenious, is merely an existence proof—it does not provide an explicit algorithm to construct such (M_n, n) -codes. It is one of the earliest examples of a proof via the *probabilistic method*.

Let us start by showing the ‘easy’ half of Shannon’s channel coding theorem, that the achievable capacity of a channel is at most the channel capacity.

Proof that $R^ \leq \text{cap}(\mathcal{C})$.* Since R^* is defined as the supremum of the achievable rates R , saying that $R^* \leq \text{cap}(\mathcal{C})$ is equivalent to saying that $\text{cap}(\mathcal{C})$ is an upper bound for the achievable rates, or in other words, that for every achievable rate R , $R \leq \text{cap}(\mathcal{C})$.

So, suppose that R is an achievable rate, so that there is a sequence of (M_n, n) -codes $(\mathcal{W}_n, \mathbf{x}^{(n)}, g_n)$ with rates R_n and maximal/average probabilities of error $\lambda_{\max}^{(n)}/p_{\text{err}}^{(n)}$ such that

$$\lim_{n \rightarrow \infty} R_n = R \quad \text{and} \quad \lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} \leq \lim_{n \rightarrow \infty} \lambda_{\max}^{(n)} = 0.$$

Let us consider a random variable $W^{(n)}$ which is uniformly distributed on \mathcal{W}_n . We can then also define the Markovian quadruple

$$W^{(n)} \rightarrow \mathbf{X}^{(n)} \rightarrow \mathbf{Y}^{(n)} \rightarrow \widehat{W}^{(n)},$$

where

- $\mathbf{X}^{(n)} = \mathbf{x}^{(n)}(W^{(n)})$,
- $\mathbf{Y}^{(n)}$ is the output of the channel \mathcal{C}^n with input $\mathbf{X}^{(n)}$,
- $\widehat{W}^{(n)} = g^{(n)}(\mathbf{Y}^{(n)})$.

The heuristic idea is that if $\mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)}] = p_{\text{err}}^{(n)}$ is small, then we can deduce a lot of information about $W^{(n)}$ from $\widehat{W}^{(n)}$. However, this information has been passed through the channel, and so we can bound it from above in some way by the channel capacity.

To make this precise, let us try to bound the mutual information between W^n and \widehat{W}^n in two ways.

Since subsequences of Markovian sequences, and their inverses, are also Markovian, we can apply the data processing inequality (Theorem 2.41) to the Markovian triples $\mathbf{Y}^{(n)} \rightarrow \mathbf{X}^{(n)} \rightarrow W^{(n)}$ and $W^{(n)} \rightarrow \mathbf{Y}^{(n)} \rightarrow \widehat{W}^{(n)}$ to deduce that

$$I(\mathbf{Y}^{(n)} ; \mathbf{X}^{(n)}) \geq I(\mathbf{Y}^{(n)} ; W^{(n)}) = I(W^{(n)} ; \mathbf{Y}^{(n)}) \geq I(W^{(n)} ; \widehat{W}^{(n)}). \quad (5.1)$$

On the other hand, we can also express

$$I(W^{(n)} ; \widehat{W}^{(n)}) = H(W^{(n)}) - H(W^{(n)} | \widehat{W}^{(n)}) = \log_2 M_n - H(W^{(n)} | \widehat{W}^{(n)}). \quad (5.2)$$

Furthermore, since the triple $W^{(n)} \rightarrow \mathbf{Y}^{(n)} \rightarrow \widehat{W}^{(n)}$ is also Markovian, we can apply Fano's inequality (Theorem 2.42) to deduce that

$$p_{\text{err}}^{(n)} \geq \frac{H(W^{(n)} | \widehat{W}^{(n)}) - 1}{\log_2 M_n},$$

or in other words

$$H(W^{(n)} | \widehat{W}^{(n)}) \leq 1 + p_{\text{err}}^{(n)} \cdot \log_2 M_n. \quad (5.3)$$

Combining equations (5.1)–(5.3) we see that

$$\begin{aligned} I(\mathbf{Y}^{(n)} ; \mathbf{X}^{(n)}) &\geq I(W^{(n)} ; \widehat{W}^{(n)}) \\ &\geq \log_2 M_n - 1 - p_{\text{err}}^{(n)} \cdot \log_2 M_n \\ &= \log_2 M_n (1 - p_{\text{err}}^{(n)}) - 1. \end{aligned} \quad (5.4)$$

On the other hand, by Lemma 5.9

$$I(\mathbf{Y}^{(n)} ; \mathbf{X}^{(n)}) \leq \text{cap}(\mathcal{C}^n) = n \cdot \text{cap}(\mathcal{C}),$$

and hence

$$n \cdot \text{cap}(\mathcal{C}) \geq \log_2 M_n \left(1 - p_{\text{err}}^{(n)}\right) - 1,$$

or in other words

$$\text{cap}(\mathcal{C}) \geq \frac{\log_2 M_n}{n} \left(1 - p_{\text{err}}^{(n)}\right) - \frac{1}{n} = R_n \left(1 - p_{\text{err}}^{(n)}\right) - \frac{1}{n}.$$

Since this holds for all n , we can take limits to see that

$$\text{cap}(\mathcal{C}) \geq \lim_{n \rightarrow \infty} R_n \left(1 - p_{\text{err}}^{(n)}\right) - \frac{1}{n} = R.$$

□

To prove the other direction, $R^* \geq \text{cap}(\mathcal{C})$, we have to produce a good sequence of (M_n, n) -codes, so we have to build clever codebooks and decoding functions that work well with the channel \mathcal{C} .

We start with the following lemma, which tells us that it is sufficient to bound the *average* probability of error, which is easier to work with, due to the nice equality $p_{\text{err}} = \mathbb{P}[\widehat{W} \neq W]$.

Lemma 5.17. *A real number $R > 0$ is an achievable rate for a channel \mathcal{C} if and only if there is a sequence of (M_n, n) -codes such that the rate $R_n = \frac{\log_2 M_n}{n}$ and the average probability of error $p_{\text{err}}^{(n)}$ satisfy*

$$\lim_{n \rightarrow \infty} R_n = R \quad \text{and} \quad \lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} = 0.$$

Proof. Since $0 \leq p_{\text{err}} \leq \lambda_{\text{max}}$, if R is an achievable rate, then there is a sequence of (M_n, n) -codes such that $\lim_{n \rightarrow \infty} \lambda_{\text{max}}^{(n)} = 0$ and hence for this same sequence of codes $\lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} = 0$.

Conversely, suppose that there exists a sequence of (M_n, n) -codes with rates $R_n \rightarrow R$ such that $p_{\text{err}}^{(n)} \rightarrow 0$. We will construct a sequence of (M'_n, n) -codes with rates R'_n such that $R'_n \rightarrow R$ and $\lambda_{\text{max}}'^{(n)} \rightarrow 0$.

Since

$$p_{\text{err}}^{(n)} = \frac{1}{M_n} \sum_{w \in \mathcal{W}_n} \lambda_w,$$

it follows that $\lambda_w \leq 2p_{\text{err}}^{(n)}$ for at least half of the $w \in \mathcal{W}_n$.

Let us define \mathcal{W}'_n to be the set of $w \in \mathcal{W}_n$ such that $\lambda_w \leq 2p_{\text{err}}^{(n)}$, so that $M'_n = |\mathcal{W}'_n| \geq \frac{M_n}{2}$, and take our (M'_n, n) code to be the restriction of the (M_n, n) code to \mathcal{W}'_n . That is, the codebook $\mathbf{x}'^{(n)}$ is the restriction of $\mathbf{x}^{(n)}$ to \mathcal{W}'_n , and the decoding function g' is such that $g'(\mathbf{y}) = g(\mathbf{y}) = w$ if $w \in \mathcal{W}'_n$ and $g'(\mathbf{y}) = \perp$ otherwise.

Note that, in the code $(\mathcal{W}'_n, \mathbf{x}'^{(n)}, g')$, the encoding and transmission is the same as in the code $(\mathcal{W}_n, \mathbf{x}^{(n)}, g)$ and, whilst the decoding differs, it only differs on \mathbf{y} such that $g(\mathbf{y}) \notin \mathcal{W}'_n$, and so everything which is correctly decoded in $(\mathcal{W}_n, \mathbf{x}^{(n)}, g)$ is also correctly decoded in $(\mathcal{W}'_n, \mathbf{x}'^{(n)}, g')$.

It follows that for every $w \in \mathcal{W}'_n$, the conditional probability of error $\lambda_w'^{(n)}$ in the new code is the same as the conditional probability of error $\lambda_w^{(n)}$ in the old code, which by assumption is at most $2p_{\text{err}}^{(n)}$.

Hence, for this code $\lambda_{\max}'^{(n)} \leq 2p_{\text{err}}^{(n)}$, and the rate is given by

$$R'_n = \frac{\log_2 M'_n}{n} = \frac{\log_2 \frac{M_n}{2}}{n} = R_n - \frac{1}{n}.$$

In particular, for this sequence of (M'_n, n) -codes, the rates satisfy

$$\lim_{n \rightarrow \infty} R'_n = \lim_{n \rightarrow \infty} R_n - \frac{1}{n} = R,$$

and the maximum error probabilities satisfy

$$\lambda_{\max}'^{(n)} \leq 2p_{\text{err}}^{(n)} \rightarrow 0,$$

so that R is an achievable rate. □

Let us assume without loss of generality that $\mathcal{W} = \{1, \dots, M_n\}$. We can think of the codebook $\mathbf{x}^{(n)}: \mathcal{W} \rightarrow \mathcal{X}^n$ as a large $(M_n \times n)$ matrix:

$$\mathbf{x}^{(n)} = \begin{pmatrix} x_1(1) & x_2(1) & \dots & x_n(1) \\ x_1(2) & x_2(2) & \dots & x_n(2) \\ \vdots & \vdots & \ddots & \vdots \\ x_1(M) & x_2(M) & \dots & x_n(M) \end{pmatrix}$$

where $\mathbf{x}^{(n)}(i) = (x_1(i), x_2(i), \dots, x_n(i))$ is the codeword corresponding to the i th message in \mathcal{W} . One could also think of this as a physical book, where on each page there is some vector $(x_1(i), x_2(i), \dots, x_n(i))$, which is the codeword corresponding to some message.

Alice will transmit one of these codewords across the channel \mathcal{C}^n , and Bob will receive a transmission (y_1, y_2, \dots, y_n) . He then has to ‘choose’ which of the possible codewords $(x_1(i), x_2(i), \dots, x_n(i))$ he thinks was transmitted, which corresponds to the decoding function g .

So, our aim is to choose a sensible codebook so that for a ‘typical’ transmission (y_1, y_2, \dots, y_n) there is a uniquely identifiable codeword $(x_1(i), x_2(i), \dots, x_n(i))$ which is the likely input resulting in the output (y_1, y_2, \dots, y_n) .

Shannon’s ingenious idea was to choose a *random* codebook. We let \mathcal{B}_n be the set of all possible codebooks, where it is easy to see that

$$|\mathcal{B}_n| = |\mathcal{X}|^{n \cdot M_n},$$

since we get to choose the $n \cdot M_n$ elements of the matrix, each of which is an element of \mathcal{X} .

Our codebook will then be a random variable $B^{(n)}$, which is distributed on \mathcal{B}_n . In other words, $B^{(n)}$ is a random $(M_n \times n)$ matrix:

$$B^{(n)} = \begin{pmatrix} X_1(1) & X_2(1) & \dots & X_n(1) \\ X_1(2) & X_2(2) & \dots & X_n(2) \\ \vdots & \vdots & \ddots & \vdots \\ X_1(M) & X_2(M) & \dots & X_n(M) \end{pmatrix}$$

where each $X_i(w)$ is a random variable taking values in \mathcal{X} . In this way, given a fixed message $w \in \mathcal{W}$ the input and output to the channel \mathcal{C}^n are random variables

$$\mathbf{X}^{(n)}(w) = (X_1(w), \dots, X_n(w)) \quad \text{and} \quad \mathbf{Y} = (Y_1, \dots, Y_n)^{(n)}.$$

In order to decode the transmission, we now need some way to identify which pairs $(\mathbf{x}, \mathbf{y}) \in \mathcal{X}^n \times \mathcal{Y}^n$ are likely to occur, or in other words: what are the ‘typical’ values taken by the random vector $(\mathbf{X}(w), \mathbf{Y})$?

Suppose we have a pair of jointly distributed random variables X, Y taking values in \mathcal{X} and \mathcal{Y} respectively with joint distribution $p_{X,Y}$ and marginal distributions p_X and p_Y . Given $\mathbf{x} \in \mathcal{X}^n$ and $\mathbf{y} \in \mathcal{Y}^n$ we write

$$p_{X,Y}^{(n)}(\mathbf{x}, \mathbf{y}) = \prod_{k=1}^n p_{X,Y}(x_k, y_k)$$

for the joint distribution of n i.i.d. copies $(X_1, Y_1), \dots, (X_n, Y_n)$ of (X, Y) . Analogously we have the (marginal) distributions of X_1, \dots, X_n and Y_1, \dots, Y_n

$$p_X^{(n)}(\mathbf{x}) = \prod_{k=1}^n p_X(x_k) \quad \text{and} \quad p_Y^{(n)}(\mathbf{y}) = \prod_{k=1}^n p_Y(y_k).$$

Since i.i.d. sequences of random variables have the AEP, we know from Lemma 3.36 that

$$\begin{aligned} -\frac{1}{n} \log_2 p_{X,Y}^{(n)}(X_1, \dots, X_n, Y_1, \dots, Y_n) &\rightarrow H(X, Y), \\ -\frac{1}{n} \log_2 p_X^{(n)}(X_1, \dots, X_n) &\rightarrow H(X), \text{ and} \\ -\frac{1}{n} \log_2 p_Y^{(n)}(Y_1, \dots, Y_n) &\rightarrow H(Y), \text{ almost surely.} \end{aligned}$$

Let us thus define the set of ‘typical’ sequences in $\mathcal{X}^n \times \mathcal{Y}^n$ which match these predictions up to some small deviation.

Definition 5.18 (Jointly typical sequences). Given X, Y and $(X_1, Y_1), \dots, (X_n, Y_n)$ as above and $\varepsilon > 0$, the set of *jointly typical sequences* is given by

$$\tilde{A}_\varepsilon^{(n)} = \left\{ (\mathbf{x}, \mathbf{y}) \in \mathcal{X}^n \times \mathcal{Y}^n : \begin{cases} \left| -\frac{1}{n} \log_2 p_{X,Y}^{(n)}(X_1, \dots, X_n, Y_1, \dots, Y_n) - H(X, Y) \right| < \varepsilon \\ \left| -\frac{1}{n} \log_2 p_X^{(n)}(X_1, \dots, X_n) - H(X) \right| < \varepsilon \\ \left| -\frac{1}{n} \log_2 p_Y^{(n)}(Y_1, \dots, Y_n) - H(Y) \right| < \varepsilon \end{cases} \right\}$$

We will use the jointly typical sequences to decode. Indeed, suppose Bob has access to the codebook $\mathbf{x}^{(n)}$ and he also knows which sequences are jointly typical.

Bob receives some transmission \mathbf{y} , and he knows that it is very likely that the input \mathbf{x} is such that (\mathbf{x}, \mathbf{y}) is jointly typical. So, Bob can look through the codebook and see whether the pair $(\mathbf{x}^{(n)}(k), \mathbf{y})$ is jointly typical.

If there is a unique such pair, then it is reasonable to guess that the input was $\mathbf{x}^{(n)}(k)$, since given any other input it’s very unlikely that the output was \mathbf{y} .

Of course, even if the codebook is chosen very carefully, it may still be the case that there is no codeword such that $(\mathbf{x}^{(n)}(k), \mathbf{y})$ is jointly typical, or there are multiple. In these cases, Bob cannot make an unambiguous guess, and so he can just decode to the erasure symbol \perp .

What we will find is that the properties of jointly typical sequences mean that, when Bob makes a guess it is very accurate, and that if we choose our codebook at random, then as long as the rate is not too large, it is in fact unlikely that Bob cannot make an unambiguous guess.

Proposition 5.19. *Let $X, Y, (X_1, Y_1), \dots, (X_n, Y_n)$, $\varepsilon > 0$ and $\tilde{A}_\varepsilon^{(n)}$ be as above. Then there exists $N(\varepsilon)$ such that:*

(a) For all $n \geq N(\varepsilon)$

$$\mathbb{P} \left[(X_1, \dots, X_n, Y_1, \dots, Y_n) \in \tilde{A}_\varepsilon^{(n)} \right] > 1 - \varepsilon,$$

(b)

$$(1 - \varepsilon)2^{n(H(X, Y) - \varepsilon)} \leq |\tilde{A}_\varepsilon^{(n)}| \leq 2^{n(H(X, Y) + \varepsilon)},$$

where the first inequality holds for all $n \geq N(\varepsilon)$ and the second for all n .

(c) If $(X'_1, Y'_1), \dots, (X'_n, Y'_n)$ are i.i.d. random variables where X'_i and Y'_i are independently distributed as X and Y , then

$$(1 - \varepsilon)2^{-n(I(X; Y) + 3\varepsilon)} \leq \mathbb{P} \left[(X'_1, \dots, X'_n, Y'_1, \dots, Y'_n) \in \tilde{A}_\varepsilon^{(n)} \right] \leq 2^{-n(I(X; Y) - 3\varepsilon)},$$

where the first inequality holds for all $n \geq N(\varepsilon)$ and the second for all n .

Proof. The proofs of (a) and (b) are exactly as in Proposition 3.39, and are left as an exercise.

For (c) we note that

$$\begin{aligned} \mathbb{P} \left[(X'_1, \dots, X'_n, Y'_1, \dots, Y'_n) \in \tilde{A}_\varepsilon^{(n)} \right] &= \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{A}_\varepsilon^{(n)}} \mathbb{P} \left[(X'_1, \dots, X'_n, Y'_1, \dots, Y'_n) = (\mathbf{x}, \mathbf{y}) \right] \\ &= \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{A}_\varepsilon^{(n)}} p_X^{(n)}(\mathbf{x}) \cdot p_Y^{(n)}(\mathbf{y}). \end{aligned}$$

However, for all $(\mathbf{x}, \mathbf{y}) \in \tilde{A}_\varepsilon^{(n)}$, by Proposition 3.39

$$\begin{aligned} 2^{-n(H(X) + \varepsilon)} &\leq p_X^{(n)}(\mathbf{x}) \leq 2^{-n(H(X) - \varepsilon)}, \\ 2^{-n(H(Y) + \varepsilon)} &\leq p_Y^{(n)}(\mathbf{y}) \leq 2^{-n(H(Y) - \varepsilon)}. \end{aligned}$$

Hence, using (b), it follows that

$$\begin{aligned} \mathbb{P} \left[(X'_1, \dots, X'_n, Y'_1, \dots, Y'_n) \in \tilde{A}_\varepsilon^{(n)} \right] &= \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{A}_\varepsilon^{(n)}} p_X^{(n)}(\mathbf{x}) \cdot p_Y^{(n)}(\mathbf{y}) \\ &\leq |\tilde{A}_\varepsilon^{(n)}| 2^{-n(H(X) - \varepsilon)} 2^{-n(H(Y) - \varepsilon)} \\ &\leq 2^{n(H(X, Y) - H(X) - H(Y) + 3\varepsilon)} \\ &= 2^{n(I(X; Y) + 3\varepsilon)}, \end{aligned}$$

and the lower bound follows in a similar manner. \square

It is now apparent why this encoding scheme should be effective. If we choose a message $w \in \mathcal{W}$, and transmit the (randomly chosen) codeword $\mathbf{X}^{(n)}(w)$ across the channel, then it is very likely, by (a), that $(\mathbf{X}^{(n)}(w), \mathbf{Y}^{(n)})$ forms a typical pair, and so Bob will have at least one candidate codeword.

However, whilst the output $\mathbf{Y}^{(n)}$ of the channel may depend on the input $\mathbf{X}^{(n)}(w)$, it is independent of the other (randomly chosen) codewords. In particular, for any $w' \neq w$ the pair $\mathbf{X}^{(n)}(w')$ and $\mathbf{Y}^{(n)}$ are independent, and so by (c) it is very unlikely that the pair $(\mathbf{X}^{(n)}(w'), \mathbf{Y}^{(n)})$ forms a typical pair, and so the candidate codeword will likely be unique!

Proof that $R^ \geq \text{cap}(\mathcal{C})$.* Let $R < \text{cap}(\mathcal{C})$. We have to prove that R is an achievable rate. By definition, there is some random variable X distributed on \mathcal{X} such that if Y is the output of the channel \mathcal{C} with input X , then

$$\text{cap}(\mathcal{C}) = I(X; Y).$$

Given $\delta > 0$, we will show that there is some $N(\delta)$ such that for any $n \geq N(\delta)$ there is an (M_n, n) -code with $M_n = 2^{\lfloor nR \rfloor}$ and with $p_{\text{err}}^{(n)} < \delta$. In particular, there is a sequence of (M_n, n) -codes such that $\lim_{n \rightarrow \infty} R_n = \lim_{n \rightarrow \infty} \frac{\log_2 M_n}{n} = R$ and $\lim_{n \rightarrow \infty} p_{\text{err}}^{(n)} = 0$. Then Lemma 5.17 will imply that R is an achievable rate.

Let us assume without loss of generality that $\mathcal{W}_n = \{1, \dots, M_n\}$. As in the discussion prior, we let \mathcal{B}_n be the set of all possible codebooks $\mathbf{x}^{(n)}: \mathcal{W} \rightarrow \mathcal{X}^n$, which we think of as $(M_n \times n)$ -matrices $(x_i(w))_{w \in \mathcal{W}, 1 \leq i \leq n}$. We will choose a random codebook $B \in \mathcal{B}_n$ by choosing independently each element $x_i(w)$ as a random variable $X_i(w)$ with distribution p_X .

Now, since $R < I(X; Y)$, if $\varepsilon > 0$ is sufficiently small then $I(X; Y) - R > 4\varepsilon$. We now define the *jointly typical decoding function* $g: \mathcal{Y}^n \rightarrow \hat{\mathcal{W}}_n = \mathcal{W} \cup \{\perp\}$, which will depend on the codebook B (and so is also a random variable!), as follows:

$$g(\mathbf{y}) = \begin{cases} w & \text{if there exists a unique } w \in \mathcal{W} \text{ with } (B(w), \mathbf{y}) \in \tilde{A}_\varepsilon^{(n)}, \\ \perp & \text{otherwise.} \end{cases}$$

So, we have defined a random (M_n, n) -code (\mathcal{W}, B, g) , whose average error probability $p_{\text{err}}^{(n)}(B)$ depends on the random choice of codebook B .

If we can show that the *expected* average error probability

$$\mathbb{E}_B \left(p_{\text{err}}^{(n)} \right) := \sum_{\mathbf{x}^{(n)} \in \mathcal{B}_n} \mathbb{P} \left[B = \mathbf{x}^{(n)} \right] p_{\text{err}}^{(n)} \left(\mathbf{x}^{(n)} \right)$$

is small, say $\mathbb{E}_B \left(p_{\text{err}}^{(n)} \right) \leq \alpha$, then there must be *some* codebook $\mathbf{x}^{(n)} \in \mathcal{B}_n$ for which the average error probability $p_{\text{err}}^{(n)}(\mathbf{x}^{(n)})$ is at most α . But this is then some deterministic codebook we can use, whose average error probability is small.

So, recall that we can express the average error probability $p_{\text{err}}^{(n)}$ as the probability that when we encode, transmit and decode a uniformly random chosen message, we end up with the correct message. That is, if we take a random variable $W^{(n)}$ which is uniformly distributed on \mathcal{W}_n , and is independent from the random codebook B , and define the random variables

- $\mathbf{X}^{(n)} = B(W^{(n)}) = (X_1(W^{(n)}), X_2(W^{(n)}), \dots, X_n(W^{(n)}))$,
- $\mathbf{Y}^{(n)}$ is the output of the channel \mathcal{C}^n with input $\mathbf{X}^{(n)}$,
- $\widehat{W}^{(n)} = g(\mathbf{Y}^{(n)})$, where g is the jointly typical decoding function,

then by the law of total probability

$$\begin{aligned} \mathbb{E}_B(p_{\text{err}}^{(n)}) &= \sum_{\mathbf{x}^{(n)} \in \mathcal{B}_n} \mathbb{P}[B = \mathbf{x}^{(n)}] p_{\text{err}}^{(n)}(\mathbf{x}^{(n)}) \\ &= \sum_{\mathbf{x}^{(n)} \in \mathcal{B}_n} \mathbb{P}[B = \mathbf{x}^{(n)}] \mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)} \mid B = \mathbf{x}^{(n)}] \\ &= \mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)}], \end{aligned} \quad (5.5)$$

where this probability is taken over the choice of codebook B and random message $W^{(n)}$.

Let us look at this probability in a different way. Again by the law of total probability,

$$\begin{aligned} \mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)}] &= \sum_{w \in \mathcal{W}_n} \mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)} \mid W^{(n)} = w] \cdot \mathbb{P}[W^{(n)} = w] \\ &= \frac{1}{M_n} \sum_{w \in \mathcal{W}_n} \mathbb{P}[W^{(n)} \neq \widehat{W}^{(n)} \mid W^{(n)} = w]. \end{aligned} \quad (5.6)$$

So, given a fixed message w , we want to estimate the probability that w is incorrectly transmitted.

Now, w is encoded to $B(w) = (X_1(w), \dots, X_n(w))$, and then $B(w)$ is transmitted to (Y_1, \dots, Y_n) across the n th extension channel \mathcal{C}^n , and we are interested in the probability that $g(Y_1, \dots, Y_n) \neq w$.

This can fail to happen for two reasons:

- The pair $(X_1(w), \dots, X_n(w), Y_1, \dots, Y_n)$ is not an element of $\tilde{A}_\varepsilon^{(n)}$;
- There is another $w' \in \mathcal{W}_n$ such that $(X_1(w'), \dots, X_n(w'), Y_1, \dots, Y_n)$ is an element of $\tilde{A}_\varepsilon^{(n)}$.

Since $X_1(w), \dots, X_n(w)$ are i.i.d. with distribution p_X , and each symbol is transmitted across the channel independently, it follows that the sequence $(X_1(w), Y_1), (X_2(w), Y_2), \dots, (X_n(w), Y_n)$ is i.i.d. with joint distribution $p_{X,Y}$. In particular, by Proposition 5.19 (a), if $n \geq N(\varepsilon)$,

$$\mathbb{P}[(X_1(w), \dots, X_n(w), Y_1, \dots, Y_n) \notin \tilde{A}_\varepsilon^{(n)}] < \varepsilon.$$

On the other hand, for any $w' \neq w$ the sequence $(X_1(w'), Y_1), \dots, (X_n(w'), Y_n)$ is i.i.d., where for each i , $X_i(w')$ and Y_i are independently distributed as X and Y . Hence, by Proposition 5.19 (c), if $n \geq N(\varepsilon)$,

$$\mathbb{P} \left[(X_1(w'), \dots, X_n(w'), Y_1, \dots, Y_n) \in \tilde{A}_\varepsilon^{(n)} \right] < 2^{-n(I(X; Y) - 3\varepsilon)} = 2^{-n(R + \varepsilon)},$$

so by the union bound

$$\begin{aligned} \mathbb{P} \left[\text{There exists } w' \neq w \text{ with } (X_1(w'), \dots, X_n(w'), Y_1, \dots, Y_n) \in \tilde{A}_\varepsilon^{(n)} \right] &\leq |\mathcal{W}_n| 2^{-n(R + \varepsilon)} \\ &\leq 2^{-n(R + \varepsilon) + \lfloor nR \rfloor} \\ &\leq 2^{-\varepsilon n}. \end{aligned}$$

Hence, for any fixed $w \in \mathcal{W}_n$,

$$\mathbb{P} \left[W^{(n)} \neq \widehat{W}^{(n)} \mid W^{(n)} = w \right] \leq \varepsilon + 2^{-\varepsilon n}.$$

Therefore by (5.5) and (5.6)

$$\mathbb{E}_B \left(p_{\text{err}}^{(n)} \right) = \mathbb{P} \left[W^{(n)} \neq \widehat{W}^{(n)} \right] = \frac{1}{M_n} \sum_{w \in \mathcal{W}_n} \mathbb{P} \left[W^{(n)} \neq \widehat{W}^{(n)} \mid W^{(n)} = w \right] \leq \varepsilon + 2^{-\varepsilon n}.$$

So, to recap, we have shown that for any sufficiently small $\varepsilon > 0$ and any $n \geq N(\varepsilon)$ there is an (M_n, n) -code, where $M_n = 2^{\lfloor nR \rfloor}$, with average error probability $p_{\text{err}}^{(n)} \leq \varepsilon + 2^{-\varepsilon n}$. In particular, if $\varepsilon < \frac{\delta}{2}$ and $n \geq \max \left\{ N(\varepsilon), \frac{\log_2 \frac{2}{\delta}}{\varepsilon} \right\} := N(\delta)$, then

$$p_{\text{err}}^{(n)} \leq \varepsilon + 2^{-\varepsilon n} < \delta.$$

□

This is perhaps in some ways unsatisfactory, since whilst Theorem 5.16 asserts the existence of a good sequence of (M_n, n) -codes, whose rate is tending to R and whose maximum error probability can be made arbitrarily small, the proof is *non-constructive*.

Firstly, we used a probabilistic argument to deduce the existence of a sequence of codes whose average error probability is small, and secondly one can check that in Lemma 5.17 we also used a probabilistic argument to show that we can use a code whose average error probability is small to construct a code whose maximum error probability is small.

Indeed, in Theorem 5.16 we essentially used the following ‘trivial’ statement:

Claim. If $a \in \mathbb{R}$, X is a real discrete random variable taking values in \mathcal{X} , $f: \mathcal{X} \rightarrow \mathbb{R}$ and $\mathbb{E}(f(X)) \leq a$, then there is some $x \in \mathcal{X}$ such that $f(x) \leq a$.

Note that this follows from Lemma 1.19 (iv), by considering the random variable $a - f(X)$. This tells us that such an x exists, but gives us no *algorithmic* way to construct such an x .

Similarly in Lemma 5.17 we used the slightly more complicated claim:

Claim. If a, X, \mathcal{X}, f are as above, then

$$\sum_{\substack{x \in \mathcal{X} \\ f(x) \leq 2a}} p_X(x) = \mathbb{P}[f(X) \leq 2a] \geq \frac{1}{2}.$$

This is essentially immediate from Markov's inequality (Lemma 1.23), since


$$\mathbb{P}[f(X) \geq 2a] \leq \frac{\mathbb{E}(f(X))}{2a} \leq \frac{1}{2}.$$

However, again this does not lead to any particular algorithm to identify $\{x \in \mathcal{X} : f(x) \leq 2a\}$.

One could of course run a brute-force search for appropriate (M_n, n) -codes, and there are also ways to *derandomize* the argument to give a constructive algorithm to find these codes, but neither are computationally efficient. Explicit constructions of such codes were a major open problem for a long time, finally being settled in the 90s with the invention of *turbo codes*, however we will not discuss these codes in any detail.

5.2 Source-channel separation theorem


Suppose we don't want to transmit a single message, but the (partial) output of some stochastic process $(V_n)_{n \in \mathbb{N}}$ (which satisfies the AEP with rate h) across the m th extension of some channel $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$, where we wish to minimise the *rate* $\frac{m}{n}$. Let us write $V^{(n)} = (V_1, \dots, V_n)$.



As before, we transmit the message $V^{(n)}$ by first encoding it as some sequence X^m in \mathcal{X}^m of length m , transmitting it across the channel with output Y^m , and then decoding the message to some estimate $\widehat{V}^{(n)}$.

If the capacity of the channel $\text{cap}(\mathcal{C}^m) = m \cdot \text{cap}(\mathcal{C})$ is significantly smaller than the entropy rate nh of the process, then a similar argument as in Theorem 5.16 will tell us that we cannot transmit $V^{(n)}$ with vanishing error probability $p_{\text{err}}^{(n)} = \mathbb{P}[V^{(n)} \neq \widehat{V}^{(n)}]$, however we choose our channel encoding and decoding.

Conversely, if $m \cdot \text{cap}(\mathcal{C}) > nh$, then one can combine Theorem 4.3 and Theorem 5.16 to transmit the message with vanishing error probability by encoding the source and the channel separately, which we refer to as *source-channel separation*.



That is, we choose first an encoding of the source $C : \mathcal{V}^n \rightarrow \{0,1\}^{nh}$ which we can decode with a vanishing probability of error (for ease of presentation here we have written nh , although in practice we need to take $n(h + \varepsilon)$ for some sufficiently small ε), which exists by Theorem 4.3. Since $R = \frac{nh}{m} < \text{cap}(\mathcal{C})$ is an achievable rate, by Theorem 5.16 there is a channel encoding $\mathbf{x}^{(m)} : \{0,1\}^{nh} \rightarrow \mathcal{X}^m$ which can be transmitted across the channel and correctly decoded with a vanishing probability of error. The total probability of error is then (at most) the sum of the errors in the source and channel encoding, and so is also vanishing.

The following is a semi-formal statement of the above discussion.

Theorem 5.20 (Shannon source-channel separation theorem). *Let $(V_n)_{n \in \mathbb{N}}$ be a stochastic process which satisfies the AEP with rate h . Set $V^{(n)} = (V_1, \dots, V_n)$, and let $\mathcal{C} = (\mathcal{X}, P, \mathcal{Y})$ be a channel.*

- *If $m \cdot \text{cap}(\mathcal{C}) < nh$, then $V^{(n)}$ cannot be transmitted across the m th extension channel \mathcal{C}^m with a vanishing probability of error.*
- *If $m \cdot \text{cap}(\mathcal{C}) > nh$, then $V^{(n)}$ can be transmitted across the m th extension channel \mathcal{C}^m with a vanishing probability of error, and the source and channel coding can be done separately.*

6 Differential Entropy

6.1 Differential Entropy

Definition 6.1. Let X be a real, continuous random variable with density function $f(x) = f_X(x)$, that is, for any $B \subseteq \mathbb{R}$

$$P_X(B) := \mathbb{P}[X \in B] = \int_B f(x) dx.$$

The *differential entropy* of X , respectively of the density function f , is defined as

$$\mathfrak{h}(X) = \mathfrak{h}(f) = - \int_{\mathbb{R}} f(x) \log_2 f(x) dx = \mathbb{E}(-\log_2 f(X)),$$

whenever the integral exists (in the sense of Lebesgue integration).

As with discrete entropy, we use the convention that $0 \log_2 0 := 0$, and so we can think of the integral as being taken over the set $\{x: f(x) > 0\}$. An immediate observation is that, by the translation invariance of the Lebesgue measure, for any $a \in \mathbb{R}$,

$$\mathfrak{h}(X - a) = \mathfrak{h}(X).$$

Before, we had that entropy was preserved by *any* bijections on the random variable, but we will see this cannot be the case in the continuous setting. We will need to make some further restrictions on the maps.

Example 6.2. (a) *Continuous equidistribution on an interval $[a, b]$.*

In this case,

$$f_X = \frac{1}{b-a} \mathbb{1}_{[a,b]},$$

and so we can calculate

$$\mathfrak{h}(X) = - \int_a^b \frac{1}{b-a} \log_2 \frac{1}{b-a} dx = \log_2(b-a).$$

Already here we see some fundamental differences to the discrete entropy function. When $b-a=1$, the entropy is $\log_2 1 = 0$, and if $b-a < 1$, the entropy is negative!

(b) *Normal distribution $N(\mu, \sigma^2)$.*

In general, $f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$.

By the comment above about translation invariance, we may assume that $\mu = 0$, in which case

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad \text{and} \quad \log_2 f(x) = -\frac{1}{\ln 2} \left(\frac{x^2}{2\sigma^2} + \ln(\sqrt{2\pi}\sigma) \right),$$

and we can calculate

$$\begin{aligned}
\mathfrak{h}(X) &= \frac{1}{\ln 2} \int_{\mathbb{R}} f(x) \left(\frac{x^2}{2\sigma^2} + \ln(\sqrt{2\pi}\sigma) \right) dx \\
&= \frac{1}{\ln 2} \left(\int_{\mathbb{R}} f(x) \frac{x^2}{2\sigma^2} dx + \int_{\mathbb{R}} \ln(\sqrt{2\pi}\sigma) f(x) dx \right) \\
&= \frac{1}{\ln 2} \left(\frac{\sigma^2}{2\sigma^2} + \ln(\sqrt{2\pi}\sigma) \right) \\
&= \frac{1}{\ln 2} \left(\frac{1}{2} + \ln(\sqrt{2\pi}\sigma) \right) \\
&= \frac{1}{2} \log_2(2\pi e\sigma^2).
\end{aligned}$$

Lemma 6.3. *Let X be a real, continuous random variable, and let $a \in \mathbb{R}$ be non-zero. Then*

$$\mathfrak{h}(aX) = \mathfrak{h}(X) + \log_2 |a|.$$

Proof. Let us assume $a > 0$, the other case is analogous. Let $Y = aX$, so that Y has density function

$$f_Y(y) = \frac{1}{a} f_X\left(\frac{y}{a}\right).$$

Then we can calculate

$$\begin{aligned}
\mathfrak{h}(aX) &= - \int_{\mathbb{R}} f_Y(y) \log_2 f_Y(y) dy \\
&= - \int_{\mathbb{R}} \frac{1}{a} f_X\left(\frac{y}{a}\right) \log_2 \left(\frac{1}{a} f_X\left(\frac{y}{a}\right) \right) dy.
\end{aligned}$$

Making a change of variables $x = \frac{y}{a}$, we see

$$\begin{aligned}
\mathfrak{h}(aX) &= - \int_{\mathbb{R}} f_X(x) \log_2 \left(\frac{1}{a} f_X(x) \right) dx \\
&= - \int_{\mathbb{R}} f_X(x) \log_2 f_X(x) dx + \log_2 a \\
&= \mathfrak{h}(X) + \log_2 a.
\end{aligned}$$

□

6.2 Discretization

Let us ‘compare’ in a way the differential entropy to the discrete entropy by way of *discretization*.

Suppose we have a random variable X on \mathbb{R} which has a ‘well-behaved’ density function, say one that is continuous on some open (bounded or unbounded) interval and is zero outside of the closure of that interval. In this case we can subdivide this interval into finitely or countably many disjoint intervals I_k , each of length $\delta > 0$.

Now, by the Mean Value Theorem for integrals, there is some $x_k = x_k(\delta) \in I_k^\circ$ (the interior of the interval) such that

$$\int_{I_k} f(x) dx = \delta f(x_k).$$

We can then define a discrete approximation to X , which we denote by X_δ , defined as

$$X_\delta(\omega) = x_k, \text{ if } X(\omega) \in I_k.$$

X_δ is then a discrete random variable, where for each k , we have

$$\mathbb{P}[X_\delta = x_k] = \mathbb{P}[X \in I_k] = \delta f(x_k).$$

We can then calculate the discrete entropy of the random variable X_δ .

$$\begin{aligned} H(X_\delta) &= \sum_k -\delta f(x_k) \log_2(\delta f(x_k)) \\ &= \sum_k -\delta f(x_k) \log_2 \delta - \delta f(x_k) \log_2 f(x_k) \\ &= -\log_2 \delta - \sum_k \delta f(x_k) \log_2 f(x_k). \end{aligned}$$

However, the latter sum is a Riemann sum for the integral $\int_{\mathbb{R}} f(x) \log_2 f(x)$, and so as $\delta \rightarrow 0$

$$H(X_\delta) \approx -\log_2 \delta + \mathfrak{h}(X),$$

where we have glossed over some technical details of convergence if the interval is unbounded.

In particular, taking $\delta = \frac{1}{n}$ we see that

$$H(X_{\frac{1}{n}}) \approx \log_2 n + \mathfrak{h}(X).$$

So, it is not the case that the differential entropy is merely the limit of the discrete entropy of a sequence of sufficiently fine discrete approximations to our random variable, the entropy of these approximations will grow unboundedly. However, we can recover an approximation to the differential entropy

$$\mathfrak{h}(X) \approx H(X_\delta) + \log_2 \delta,$$

by taking into account the approximation parameter δ .

6.3 Joint and conditional differential entropy

More generally, if we have a collection of jointly distributed random variables X_1, \dots, X_n , we can think of them as a random vector $X = (X_1, \dots, X_n)^T$ distributed on \mathbb{R}^n .

We then have a corresponding density function¹ $f_X : \mathbb{R}^n \rightarrow \mathbb{R}$ with respect to the Lebesgue measure.

¹Note that in general, even if X_1, \dots, X_n have density functions, it may be that X does not.

Definition 6.4. Given a random vector $X = (X_1, \dots, X_n)^T$ with joint density function f_X , the *joint differential entropy* is defined as

$$\mathfrak{h}(X) = \mathfrak{h}(f_X) := - \int_{\mathbb{R}^n} f(\mathbf{x}) \log_2 f(\mathbf{x}) d\mathbf{x} \quad \text{where } \mathbf{x} = (x_1, \dots, x_n)^T,$$

when the integral exists in the sense of Lebesgue integration. As before, it is easy to see that the differential entropy is translation invariant.

Given two random vectors X and Y , distributed on \mathbb{R}^m and \mathbb{R}^n respectively, with a joint density function $f_{X,Y}$ and a conditional density function $f_{X|Y}$, we have in general²

$$f_{X|Y}(\mathbf{x} | \mathbf{y}) = \frac{f_{X,Y}(\mathbf{x}, \mathbf{y})}{f_Y(\mathbf{y})},$$

and we can define

$$\mathfrak{h}(X | Y = \mathbf{y}) := \mathfrak{h}(f_{X|Y}(\cdot | \mathbf{y}))$$

and then the *conditional differential entropy* as

$$\mathfrak{h}(X | Y) := \int_{\mathbb{R}^n} \mathfrak{h}(X | Y = \mathbf{y}) f_Y(\mathbf{y}) d\mathbf{y},$$

where $f_Y(\mathbf{y}) = \int_{\mathbb{R}^m} f_{X,Y}(\mathbf{x}, \mathbf{y}) d\mathbf{x}$, as always if the integral exists.

It is relatively easy to verify that, if all the involved integrals are finite, then

$$\mathfrak{h}(X | Y) = - \int_{\mathbb{R}^m \times \mathbb{R}^n} f_{X,Y}(\mathbf{x}, \mathbf{y}) \log_2 f_{X|Y}(\mathbf{x} | \mathbf{y}) d\mathbf{x} d\mathbf{y} = \mathfrak{h}(X, Y) - \mathfrak{h}(Y).$$

The following lemma then follows inductively from the definitions as in the proof of Theorem 2.13.

Lemma 6.5 (Chain rule for differential entropy). *Given a continuous random vector $X = (X_1, \dots, X_n)^T$,*

$$\mathfrak{h}(X) = \sum_{i=1}^n \mathfrak{h}(X_i | X_1, X_2, \dots, X_{i-1}).$$

We also have the following multidimensional version of Lemma 6.3 for how the entropy of a random vector scales under invertible linear transformations.

Lemma 6.6. *Let $X = (X_1, \dots, X_n)$ be a real, continuous random vector, and let A be a non-singular $(n \times n)$ -matrix and $\mathbf{b} \in \mathbb{R}^n$. Then*

$$\mathfrak{h}(AX + \mathbf{b}) = \mathfrak{h}(X) + \log_2 |\det A|.$$

Proof. By translation invariance, we may assume without loss of generality that $\mathbf{b} = \mathbf{0}$.

Let us start by considering the density function of $Y = AX$. Given a Borel set B ,

$$\mathbb{P}[Y \in B] = \mathbb{P}[AX \in B] = \mathbb{P}[X \in A^{-1}B] = \int_{A^{-1}B} f_X(\mathbf{x}) d\mathbf{x},$$

²Glossing over some technicalities.

where $A^{-1}B = \{A^{-1}\mathbf{b} : \mathbf{b} \in B\}$.

We would like to make the change of variables $\mathbf{y} = A\mathbf{x}$, so that $d\mathbf{y} = |\det A|d\mathbf{x}$, and it is clear that $\mathbf{x} \in A^{-1}B$ if and only if $\mathbf{y} = A\mathbf{x} \in B$. Hence

$$\mathbb{P}[Y \in B] = \int_B \frac{1}{|\det A|} f_X(A^{-1}\mathbf{y}) d\mathbf{y},$$

so that $f_Y(\mathbf{y}) = \frac{1}{|\det A|} f_X(A^{-1}\mathbf{y})$.

Hence, making the same change of variables in reverse, we can calculate

$$\begin{aligned} \mathfrak{h}(Y) &= - \int_{\mathbb{R}} \frac{1}{|\det A|} f_X(A^{-1}\mathbf{y}) \log_2 \left(\frac{1}{|\det A|} f_X(A^{-1}\mathbf{y}) \right) d\mathbf{y} \\ &= - \int_{\mathbb{R}} f_X(\mathbf{x}) \log_2 \left(\frac{1}{|\det A|} f_X(\mathbf{x}) \right) d\mathbf{x} \\ &= - \int_{\mathbb{R}} f_X(\mathbf{x}) \log_2 (f_X(\mathbf{x})) d\mathbf{x} + \int_{\mathbb{R}} f_X(\mathbf{x}) \log_2 (|\det A|) d\mathbf{x} \\ &= \mathfrak{h}(X) + \log_2 |\det A|. \end{aligned}$$

□

Example 6.7. The general n -dimensional normal distribution is determined by two parameters: $\boldsymbol{\mu} \in \mathbb{R}^n$, and a positive definite $(n \times n)$ -matrix Σ . Its density function is

$$f(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \Sigma}} \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right).$$

From this it can be calculated that $N(\boldsymbol{\mu}, \Sigma) = (X_1, \dots, X_n)^T$ is such that

$$\mathbb{E}(X) = \boldsymbol{\mu} \text{ and } \Sigma = (\text{Cov}(X_i, X_j))_{i \in [n], j \in [n]}.$$

This distribution can also be obtained as follows: We start with a vector $Y = (Y_1, \dots, Y_n)^T$ of i.i.d. standard normal random variables, i.e., $Y_i \sim N(0, 1)$, a non-singular $(n \times n)$ -matrix A and some vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)^T$. If we let

$$X = (X_1, \dots, X_n)^T = AY + \boldsymbol{\mu},$$

then we can compute

$$\mathbb{E}(X) = (\mathbb{E}(X_1), \dots, \mathbb{E}(X_n))^T = \boldsymbol{\mu},$$

and the covariance matrix is given by

$$\Sigma = (\text{Cov}(X_i, X_j))_{i \in [n], j \in [n]} = AA^T.$$

Finally, a computation will show that the density of X is given by

$$f_X(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \Sigma}} \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right),$$

and so $X \sim N(\boldsymbol{\mu}, \Sigma)$.

Hence, it follows from Lemma 6.6 that

$$\mathfrak{h}(N(\boldsymbol{\mu}, \Sigma)) = \mathfrak{h}(Y_1, \dots, Y_n) + \log_2 |\det A|.$$

Now, since $\Sigma = AA^T$, it follows that $|\det A| = \sqrt{\det \Sigma}$, and we will show shortly that for independent random variables, $\mathfrak{h}(Y_1, \dots, Y_n) = \sum_{i=1}^n \mathfrak{h}(Y_i)$ and hence

$$\mathfrak{h}(N(\boldsymbol{\mu}, \Sigma)) = \frac{n}{2} \log_2(2\pi e) + \frac{1}{2} \log_2 \det \Sigma = \frac{1}{2} \log_2((2\pi e)^n \det \Sigma).$$

Definition 6.8. Given two density functions f and g on \mathbb{R}^n we can define the *Kullback-Leibler Divergence* as

$$D(f \parallel g) := \int_{\mathbb{R}^n} f(\mathbf{x}) \log_2 \left(\frac{f(\mathbf{x})}{g(\mathbf{x})} \right) d\mathbf{x}$$

under usual conventions about the value of $a \log_2 b$ when a or b is equal to 0.

It can be checked that if P_f and P_g are the associated probability measures, then $D(f \parallel g)$ will be finite only if $g > 0$ holds P_f -almost everywhere.

Theorem 6.9. [Information Inequality] Let f and g be density functions on \mathbb{R}^n . Then

$$D(f \parallel g) \geq 0$$

with equality if and only if $f = g$ almost everywhere.

Proof. The proof follows the exact same lines as the proof of Theorem 2.32.

Let $\text{Supp}(f) = \{x : f(x) > 0\}$ denote the support of a function f . Since $\log_2 t$ is a strictly concave function of t on $[0, \infty)$, we can apply Jensen's inequality in its general form to see that

$$\begin{aligned} -D(f \parallel g) &:= \int_{\mathbb{R}^n} f(\mathbf{x}) \log_2 \left(\frac{g(\mathbf{x})}{f(\mathbf{x})} \right) d\mathbf{x} \\ &= \int_{\text{Supp}(f)} f(\mathbf{x}) \log_2 \left(\frac{g(\mathbf{x})}{f(\mathbf{x})} \right) d\mathbf{x} \\ &\leq \log_2 \left(\int_{\text{Supp}(f)} f(\mathbf{x}) \left(\frac{g(\mathbf{x})}{f(\mathbf{x})} \right) d\mathbf{x} \right) \\ &\leq \log_2 1 = 0, \end{aligned}$$

with equality only if $\int_{\text{Supp}(f)} g(\mathbf{x}) d\mathbf{x} = 1$, and $\frac{g}{f}$ is constant almost everywhere on $\text{Supp}(f)$. Finally, since f and g are probability densities, it follows that this constant must be 1, and hence f and g are equal almost everywhere. \square

Definition 6.10. Given two random vectors X and Y , with a joint density function $f_{X,Y}$ we define

$$I(X ; Y) = D(f_{X,Y} \parallel f_X \otimes f_Y).$$

One can argue in precisely the same way as in the discrete case that the following holds (whenever all the relevant integrals are finite).

- Lemma 6.11.**
1. $I(X ; Y) = \mathfrak{h}(X) - \mathfrak{h}(X | Y) = \mathfrak{h}(Y) - \mathfrak{h}(Y | X) = \mathfrak{h}(X) + \mathfrak{h}(Y) - \mathfrak{h}(X, Y);$
 2. $I(X ; Y) \geq 0$ with equality if and only if X and Y are independent;
 3. $\mathfrak{h}(X, Y) \leq \mathfrak{h}(X) + \mathfrak{h}(Y)$ with equality if and only if X and Y are independent.

Remark 6.12. Whilst the differential entropy does not satisfy some of the useful properties of the discrete entropy function, the mutual information is much better behaved.

Indeed, considering just the one-dimensional case, given jointly distributed random variables X and Y we could consider the discrete approximations X_δ and Y_δ . Then, one can similarly show that

$$I(X_\delta ; Y_\delta) = H(X_\delta) - H(X_\delta | Y_\delta) \approx \mathfrak{h}(X) - \log_2 \delta - (\mathfrak{h}(X | Y) - \log_2 \delta) = I(X ; Y).$$

In the general case, one can show that the mutual information of X and Y can be obtained as the limit of the mutual information of a sequence of finer and finer discrete approximations to the pair (X, Y) .

Theorem 6.13. Let $X = (X_1, \dots, X_n)^T$ be a continuous real random vector with $\mathbb{E}(X) = \mathbf{0}$ and with a positive definite covariance matrix $\Sigma = (\text{Cov}(X_i, X_j))_{i \in [n], j \in [n]}$. Then

$$\mathfrak{h}(X) \leq \mathfrak{h}(N(\mathbf{0}, \Sigma)),$$

with equality if and only if $X \sim N(\mathbf{0}, \Sigma)$.

Proof. Let us write ψ_Σ for the density function of $N(\mathbf{0}, \Sigma)$ and f for the density function of X . Then, by Theorem 6.9 we know that

$$0 \leq D(f \parallel \psi_\Sigma) = -\mathfrak{h}(X) - \int_{\mathbb{R}^n} f(\mathbf{x}) \log_2 \psi_\Sigma(\mathbf{x}) d\mathbf{x}.$$

Now,

$$\log_2 \psi_\Sigma(\mathbf{x}) = -\frac{1}{2} \log_2 ((2\pi)^n |\det \Sigma|) - \frac{1}{2} \mathbf{x}^T \Sigma^{-1} \mathbf{x},$$

and so

$$\begin{aligned} \mathfrak{h}(X) &\leq \frac{1}{2} \int_{\mathbb{R}^n} f(\mathbf{x}) (\log_2 ((2\pi)^n |\det \Sigma|) + \mathbf{x}^T \Sigma^{-1} \mathbf{x}) d\mathbf{x} \\ &= \frac{1}{2} \int_{\mathbb{R}^n} f(\mathbf{x}) (\log_2 ((2\pi)^n |\det \Sigma|)) d\mathbf{x} + \frac{1}{2} \int_{\mathbb{R}^n} f(\mathbf{x}) (\mathbf{x}^T \Sigma^{-1} \mathbf{x}) d\mathbf{x}. \end{aligned}$$

Now, the first integral is a constant integrated over a probability distribution, so it makes no difference *which* probability distribution we take. However, if we consider the second integral, we see that it is of the form

$$\int_{\mathbb{R}^n} \sum_{i,j} \alpha_{i,j} f(\mathbf{x}) x_i \cdot x_j d\mathbf{x} = \sum_{i,j} \alpha_{i,j} \text{Cov}(X_i, X_j).$$

which only depends on the covariance matrix Σ of f . In particular, this second integral is unchanged if we replace the distribution f with another distribution with the same covariance matrix, for example ψ_Σ .


So, explicitly

$$\begin{aligned}
\mathfrak{h}(X) &\leq \frac{1}{2} \int_{\mathbb{R}^n} f(\mathbf{x}) (\log_2 ((2\pi)^n |\det \Sigma|) + \mathbf{x}^T \Sigma^{-1} \mathbf{x}) d\mathbf{x} \\
&= \frac{1}{2} \int_{\mathbb{R}^n} \psi_{\Sigma}(\mathbf{x}) (\log_2 ((2\pi)^n |\det \Sigma|) + \mathbf{x}^T \Sigma^{-1} \mathbf{x}) d\mathbf{x} \\
&= \frac{1}{2} \int_{\mathbb{R}^n} \psi_{\Sigma}(\mathbf{x}) \log_2 \psi_{\Sigma}(\mathbf{x}) d\mathbf{x} \\
&= \mathfrak{h}(N(\mathbf{0}, \Sigma)).
\end{aligned}$$

□

6.4 The Gaussian Channel

This is a simple channel where the input is some continuous, real random variable X . During transmission some Gaussian random noise Z with distribution $N(0, \sigma^2)$ is added to the signal, and this noise is *additive*, so that the output Y is given by $X + Z$.



As with discrete channels, we can consider the n th *channel extension*, where the input consists of a sequence X_1, \dots, X_n of independent random variables (which we will consider to be i.i.d. copies of some fixed X) and there is an independent sequence of i.i.d. copies Z_1, \dots, Z_n of the Gaussian noise Z , and the output is then a sequence Y_1, \dots, Y_n of i.i.d. random variables where $Y_i = X_i + Z_i$ for each $i \in [n]$.

A typical assumption on the input distribution X is that $\mathbb{E}(X^2) \leq \tau^2$ for some pre-determined constant $\tau > 0$, which we can think of as an analogue of restricting the size $|\mathcal{X}|$ of the input alphabet in a discrete channel.

Definition 6.14. The *capacity* of the Gaussian channel with parameters σ^2 and τ^2 is defined as

$$\text{cap}(\sigma^2, \tau^2) = \max \left\{ I(X; Y) : \begin{array}{l} X \text{ a continuous, real random variable with } \mathbb{E}(X^2) \leq \tau^2 \text{ and} \\ Y = X + Z, \text{ where } Z \sim N(0, \sigma^2) \text{ is independent of } X \end{array} \right\}.$$

In some ways we can think of this channel as being *weakly symmetric*—if we think of the conditional distribution function $f_{Y|X}$, which is just the distribution of Z , we see that it is symmetric in the sense that all the functions $f_{Y|X}(\cdot|x)$ and $f_{Y|X}(y|\cdot)$ are interchangeable, just translations of the distribution of Z .

Theorem 6.15. For any $\sigma^2, \tau^2 > 0$,

$$\text{cap}(\sigma^2, \tau^2) = \frac{1}{2} \log_2 \left(1 + \frac{\tau^2}{\sigma^2} \right),$$

and the maximum is achieved when $X \sim N(0, \tau^2)$.

Proof. By Lemma 6.11 we can write

$$\begin{aligned} I(X ; Y) &= \mathfrak{h}(Y) - \mathfrak{h}(Y | X) \\ &= \mathfrak{h}(Y) - \mathfrak{h}(X + Z | X) \\ &= \mathfrak{h}(Y) - \mathfrak{h}(Z) \\ &= \mathfrak{h}(Y) - \mathfrak{h}(N(0, \sigma^2)), \end{aligned}$$

where we used the fact that differential entropy is translation invariant and the independence of X and Z , i.e.,

$$\begin{aligned} \mathfrak{h}(X + Z | X) &= \int_{\mathbb{R}^n} f_X(\mathbf{x}) \mathfrak{h}(X + Z | X = \mathbf{x}) d\mathbf{x} = \int_{\mathbb{R}^n} f_X(\mathbf{x}) \mathfrak{h}(\mathbf{x} + Z) d\mathbf{x} \\ &= \int_{\mathbb{R}^n} f_X(\mathbf{x}) \mathfrak{h}(Z) d\mathbf{x} = \mathfrak{h}(Z). \end{aligned}$$

Since the second term is constant, it remains to maximise $\mathfrak{h}(Y)$ under the constraint that $\mathbb{E}(X^2) \leq \tau^2$. Now, again by translation invariance, we may assume furthermore without loss of generality that $\mathbb{E}(X) = 0$: otherwise, subtract $\mu = \mathbb{E}(X)$ to get a random variable $X' = X - \mu$ with

$$\mathbb{E}(X'^2) = \mathbb{E}((X - \mu)^2) = \mathbb{V}(X) = \mathbb{E}(X^2) - \mathbb{E}(X)^2 \leq \mathbb{E}(X^2) \leq \tau^2.$$

Now, if we don't control the moments of X , but the moments of Y , we have already solved this optimisation problem in Theorem 6.13. However, we can note that, since $\mathbb{E}(X) = \mathbb{E}(Z) = 0$ we have $\mathbb{E}(Y) = \mathbb{E}(X + Z) = 0$, and since X and Z are independent

$$\mathbb{E}(Y^2) = \mathbb{E}((X + Z)^2) = \mathbb{E}(X^2) + \mathbb{E}(Z^2) \leq \tau^2 + \sigma^2.$$

Hence, by Theorem 6.13

$$\mathfrak{h}(Y) \leq \mathfrak{h}(N(0, \tau^2 + \sigma^2)) = \frac{1}{2} \log_2 (2\pi e(\tau^2 + \sigma^2))$$

with equality if and only if $Y \sim N(0, \tau^2 + \sigma^2)$.

However, if we take $X \sim N(0, \tau^2)$, it is a simple exercise to show that

$$X + Z \sim N(0, \tau^2 + \sigma^2).$$

Hence, we have shown that

$$I(X ; Y) \leq \mathfrak{h}(N(0, \tau^2 + \sigma^2)) - \mathfrak{h}(N(0, \sigma^2))$$

and that equality is achieved when $X \sim N(0, \tau^2)$, and consequently

$$\begin{aligned} \text{cap}(\sigma^2, \tau^2) &= \mathfrak{h}(N(0, \tau^2 + \sigma^2)) - \mathfrak{h}(N(0, \sigma^2)) \\ &= \frac{1}{2} \log_2(2\pi e(\tau^2 + \sigma^2)) - \frac{1}{2} \log_2(2\pi e\sigma^2) \\ &= \frac{1}{2} \log_2 \left(1 + \frac{\tau^2}{\sigma^2} \right). \end{aligned}$$

□