## 1 Class Administrivia

15-859 "ITCS" is Information Theory + Application in TCS. The instructors are Mahdi Cheraghchi and Venkat Guruswami. The requirements for the course are:

- 3-5 problem sets (every 2-3 weeks, "on the hardish side")
- Some amount of scribe notes, depending on how many people are actually enrolled

The impetus for this class is Shannon's seminal 1948 paper and the increasing role of informationtheoretic mindset and tools in the theory of computing.

Essentially, information theory gives a clean calculus in various settings, especially when it comes to offerings of lower bounds.

The text is Cover and Thomas's Elements of Information Theory, but it's not required. Hopefully the scribe notes will be very good so those will offer a good reference.

The prerequisite for the class is to be comfortable with discrete probability and random variables.

## 2 Entropy

Let $X$ be a discrete random variable, say,

$$
X \sim \begin{cases}a & \frac{2}{3} \\ b & \frac{1}{6} \\ c & \frac{1}{6}\end{cases}
$$

and define $H(X)$ to be the entropy of the random variable $X$. We'd like the entropy to measure the amount of information conveyed by the value of the random variable $X$, where the amount of information is, roughly speaking, the amount of surprise we get from this random variable.

Suppose we had such a function $S:[0,1] \rightarrow \mathbb{R}^{+}$that takes a probability and maps it to a non-negative real. Let's think about some natural properties of $S$.

Suppose we told you that $X$ is $a, b$, or $c$. Are you surprised? No-because this occurs with probability 1 and we already know this. So therefore we'd like $S(1)=0$ : flipping a two-headed coin doesn't give us any surprise.

Now suppose we told you that $X=a$. This is a little bit surprising, but $X=b$ is more surprising because this is a rarer event. So we'd like $S$ to satisfy $S(p)>S(q)$ if $p<q$.

We'd also like $S(x)$ to be a continuous function of $x$, because we'd like the surprise to not exhibit jumps (if we had instead

$$
X \sim \begin{cases}a & \frac{2}{3}+10^{-6} \\ b & \frac{1}{6}-10^{-6} \\ c & \frac{1}{6}\end{cases}
$$

our impression of the "surprise" of $a$ should not be much different than the original $X$ ).
Now suppose $X_{1}$ and $X_{2}$ are two independent instantiations of $X$. It's natural for our surprise from $X_{1}=a$ and $X_{2}=b$ to be additive (as each event "adds" to our surprise):

$$
S\left(\frac{2}{3} \frac{1}{3}\right)=S\left(\frac{2}{3}\right)+S\left(\frac{1}{3}\right)
$$

which means we must have $S(p q)=S(p)+S(q)$.
Exercise. The only $S$ that satisfies the above requirements is $S(p)=c \log _{2}\left(\frac{1}{p}\right), c>0$. A convenient normalization constant is that we are unit surprised by a fair coin flip $(S(1 / 2)=1)$, which sets $c=1$.

Definition 1 (Entropy). The entropy of a discrete random variable $X$, denoted $H(X)$, is the average surprise for an outcome of $X$ :

$$
\underset{x}{\mathbb{E}}\left(\log _{2} \frac{1}{P(x=x)}\right)=\sum_{x \in \operatorname{supp}(X)} p(x) \log _{2} \frac{1}{p(x)},
$$

where we define $0 \log _{2} \frac{1}{0}=0$.
It is simple to derive the following facts and examples:
Fact. $H(X) \geq 0 ; H(X)=0$ if and only if $X$ is deterministic.
Example. Let $X \sim\left\{\begin{array}{ll}a & \frac{2}{3} \\ b & \frac{1}{6} \\ c & \frac{1}{6}\end{array}\right.$ as before. Then $H(X)=\frac{2}{3} \log _{2} \frac{3}{2}+\frac{1}{6} \log _{2} 6+\frac{1}{6} \log _{2} 6$.
Example. Let $X$ be uniform on $\left\{a_{1}, \ldots, a_{n}\right\}$. Then $H(X)=\log _{2} n$.
Fact. If $|\operatorname{supp}(x)|=n$, then $H(X) \leq \log _{2} n$.
Communication. Let's see if we can derive the binary entropy function in a different setting. Suppose $X \leftarrow \operatorname{Ber}(1 / 2)$, then $H(X)=1$. And if $X \leftarrow \operatorname{Ber}(1)$, then $H(X)=0$. Now suppose $X \leftarrow \operatorname{Ber}(p)$, and we have $X_{1}, X_{2}, \ldots, X_{n}$ iid samples from this distribution, represented as a bitstring $\{0,1\}^{n}$. A way to communicate this in an efficient way is that we first send $k$, the number of 1 bits in the string, and then use $\left\lceil\log _{2}\binom{n}{k}\right\rceil$ bits to reveal which subset those 1 s go in. The expected number of bits communicated with this scheme is

$$
\mathbb{E} \text { (number of bits) }=\underbrace{\left\lceil\log _{2} n\right\rceil}_{\text {sending } k}+\underbrace{\sum_{k=0}^{n}\binom{n}{k} p^{k}(1-p)^{n-k}\left\lceil\log _{2}\binom{n}{k}\right\rceil}_{\text {number of bits required for each } k} .
$$

What we are interested is the average fraction of bits we need as the number of bits as $n$ grows large: $\lim _{n \rightarrow \infty} \frac{1}{n} \mathbb{E}$ (number of bits). Dividing our expression by $n$, we can omit the first term ( $\log n / n \rightarrow 0$ as $n \rightarrow \infty$ ) and we can use Stirling's approximation (or the intuitive concept that if $k$ is far from $p n$ the term in the sum is close to zero) to obtain that

$$
\lim _{n \rightarrow \infty} \frac{1}{n} \mathbb{E}(\text { number of bits })=h(p)=p \log _{2} \frac{1}{p}+(1-p) \log _{2} \frac{1}{1-p},
$$

which tells us that we can take $H(\operatorname{Ber}(p))$ bits, on average, to communicate a $\operatorname{Ber}(p)$ random variable.
Non-binary random variable. We can extend this definition of entropy in the communications context to non-binary random variables. Suppose $X \leftarrow\left\{a_{1}, \ldots, a_{n}\right\}$. Suppose we have related variables

$$
Z_{1}= \begin{cases}1 & X=a_{1} \\ 0 & \text { else }\end{cases}
$$

and

$$
Z_{2}=\left\{\begin{array}{ll}
a_{2} & \frac{p_{2}}{1-p_{1}} \\
\vdots & \\
a_{n} & \frac{p_{n}}{1-p_{1}}
\end{array} .\right.
$$

We can break

$$
H(X)=H\left(Z_{1}\right)+\operatorname{Pr}\left(Z_{1}=0\right) H\left(Z_{2}\right)
$$

as we need to communicate $Z_{1}$ and then, if $Z_{1}$ is 0 , we need to communicate $Z_{2}$. Since we know how to compute $H\left(Z_{1}\right)$ from the binary case, we can recursively use this to obtain $H(X)=\sum_{i=1}^{n} p_{i} \log 1 / p_{i}$, which is the same expression as we got before.

Lemma 1. $|\operatorname{supp}(x)|=n$ implies that $H(X) \leq \log _{2} n$ (presented as a fact earlier).
Proof. We use Jensen's inequality: for convex functions $f$, we have $\mathbb{E}[f(x)] \leq f(\mathbb{E}[x])$ (this implies $\mathbb{E}\left[X^{2}\right] \geq$ $\mathbb{E}[X]^{2}$ and $\mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}}[X]$.) Since we have

$$
\begin{aligned}
H(X) & =\underset{x \leftarrow X}{\mathbb{E}}\left[\log \frac{1}{p(x)}\right] \\
& \leq \log \left(\underset{x \leftarrow X}{\mathbb{E}} \frac{1}{p(x)}\right)
\end{aligned}
$$

where we have applied Jensen's inequality to the expression with $f(x)=\log x$ while using the random variable $Z \sim 1 / p(x)$ w.p. $p(x)$ :

$$
H(X)=\underset{Z}{\mathbb{E}}[\log Z] .
$$

## 3 Lossless Data Compression aka Source Coding

Suppose $X \in\{a, b, c, \ldots, z\}$. We can communicate this with $\left\lceil\log _{2}(26)\right\rceil=5$ bits or use ASCII and use 7 bits. But suppose the distribution of $X$ is not uniform. Then we might be able to leverage that for a more efficient bit representation.

We introduce the concept of a prefix-free code or instantaneous code. The simple rule for such a code is that no codeword is a prefix of another codeword. Suppose we need to encode $A, B$, and $C$. We can't say $A$ maps to 0 and $B$ maps to 1 because then $C$ can neither start with 0 or 1 . This also makes decoding easy, because you can decode as soon as your buffer matches a codeword (it can't be a prefix of another code). A potential prefix-free code for $A, B$, and $C$ is $A \mapsto 0, B \mapsto 10$ and $C \mapsto 11$.

Our question now becomes: what are the lengths of codewords for which prefix-free codes exist? We can visualize a prefix-free code as a binary tree where the letters sit at the leaves (see figure 1)


Figure 1: The tree for the example $A=0, B=10$, and $C=11$
Now suppose we have a probability distribution on $X$ :

$$
X \leftarrow \begin{cases}a_{1} & p_{1} \\ a_{2} & p_{2} \\ \vdots & \vdots \\ a_{n} & p_{n}\end{cases}
$$

and we can use a greedy construction (combine the least likely two letters and put them at the bottom of the tree, and then recurse on the $n-1$-size support, as seen in figure 2 ) to minimize $\sum_{i} p_{i} l_{i}$ where $l_{i}$ is the length of the codeword that presents $i$. This construction is called Huffman coding.


Figure 2: We combine the two least likely letters and recurse on the rest of the triangle
Fact. Huffman codes minimize the expected depth $\sum_{i} p_{i} l_{i}$. However, they have time to construct that is polynomial in the size of the support of the underlying random variable, which may be very large in the case of product probability distributions.
Fact (Kraft's Inequality). A prefix-free code for $n$ symbols exists with lengths $\ell_{1}, \ell_{2}, \ldots, \ell_{n}$ if and only if

$$
\sum_{i=1}^{n}|\Sigma|^{-\ell_{i}} \leq 1
$$

where $|\Sigma|$ is the size of the alphabet for the code ( 2 for a binary alphabet, 3 if were coding over $\{0,1,2\}$, etc).
Proof (only if). This will proceed probabilistically. Take a random walk down the tree until you hit a leaf or fall off the tree. The

$$
\begin{aligned}
& 1 \geq \operatorname{Pr}(\text { we hit a leaf }) \\
& \stackrel{(a)}{=} \sum_{i=1}^{n} \operatorname{Pr}(\text { hit leaf } i) \\
& \stackrel{(b)}{=} \sum_{i=1}^{n} \frac{1}{2^{\ell_{i}}}
\end{aligned}
$$

where (a) follows because we can't hit multiple leaves and therefore the probability is over a disjoint union, and (b) follows because there is a unique path of length $\ell_{i}$ that leads to any particular leaf. This gives the result that we desire.

Proof (if). Without loss of generality, we have $\ell_{1} \leq \ell_{2} \leq \ell_{3} \ldots \leq \ell_{n}=L$, where $L$ is the total depth of the tree. Since there is a codeword at $\ell_{1}$, that eliminates $2^{L-\ell_{1}}$ possible leaves from the full binary tree. Since the second codeword can't have a prefix of the first codeword, we can argue the same for $\ell_{2}$. This process is illustrated in 3. But if we have

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

then we can't possibly "eat up" the entirety of the tree until the last step.


Figure 3: The codewords of length $\ell_{1}$ and $\ell_{2}$ are colored in their respective colors. The red triangles indicate the parts of the tree that have been cut. Due to the prefix-free nature of the code, they are non-overlapping.

