# A Gentle Tutorial on <br> Information Theory and Learning 

Roni Rosenfeld

Carnegie Mellon University
Carnegie Mellon

## Outline

- First part based very loosely on [Abramson 63].
- Information theory usually formulated in terms of information channels and coding - will not discuss those here.

1. Information
2. Entropy
3. Mutual Information
4. Cross Entropy and Learning

## Information

- information $\neq$ knowledge

Concerned with abstract possibilities, not their meaning

- information: reduction in uncertainty

Imagine:
\#1 you're about to observe the outcome of a coin flip
\#2 you're about to observe the outcome of a die roll There is more uncertainty in \#2

Next:

1. You observed outcome of $\# 1 \rightarrow$ uncertainty reduced to zero.
2. You observed outcome of $\# 2 \rightarrow$ uncertainty reduced to zero.
$\Longrightarrow$ more information was provided by the outcome in \#2

## Definition of Information

(After [Abramson 63])
Let $E$ be some event which occurs with probability $P(E)$. If we are told that $E$ has occurred, then we say that we have received

$$
I(E)=\log _{2} \frac{1}{P(E)}
$$

bits of information.

- Base of log is unimportant - will only change the units We'll stick with bits, and always assume base 2
- Can also think of information as amount of "surprise" in $E$ (e.g. $P(E)=1, P(E)=0$ )
- Example: result of a fair coin flip $\left(\log _{2} 2=1 \mathrm{bit}\right)$
- Example: result of a fair die roll $\left(\log _{2} 6 \approx 2.585\right.$ bits)


## Information is Additive

- $I(k$ fair coin tosses $)=\log \frac{1}{1 / 2^{k}}=k$ bits
- So:
- random word from a 100,000 word vocabulary: $I($ word $)=\log 100,000=16.61$ bits
- A 1000 word document from same source: $\mathrm{I}($ document $)=16,610$ bits
- A 480x640 pixel, 16-greyscale video picture: $I($ picture $)=307,200 \cdot \log 16=1,228,800$ bits
$\bullet \Longrightarrow A(V G A)$ picture is worth (a lot more than) a 1000 words!
- (In reality, both are gross overestimates.)


## Entropy

A Zero-memory information source $S$ is a source that emits symbols from an alphabet $\left\{s_{1}, s_{2}, \ldots, s_{k}\right\}$ with probabilities $\left\{p_{1}, p_{2}, \ldots, p_{k}\right\}$, respectively, where the symbols emitted are statistically independent.

What is the average amount of information in observing the output of the source $S$ ?

Call this Entropy:

$$
H(S)=\sum_{i} p_{i} \cdot I\left(s_{i}\right)=\sum_{i} p_{i} \cdot \log \frac{1}{p_{i}}=E_{P}\left[\log \frac{1}{p(s)}\right]
$$

## Alternative Explanations of Entropy

$$
H(S)=\sum_{i} p_{i} \cdot \log \frac{1}{p_{i}}
$$

1. avg amt of info provided per symbol
2. avg amount of surprise when observing a symbol
3. uncertainty an observer has before seeing the symbol
4. avg \# of bits needed to communicate each symbol (Shannon: there are codes that will communicate these symbols with efficiency arbitrarily close to $H(S)$ bits/symbol; there are no codes that will do it with efficiency $<H(S)$ bits/symbol)

## Entropy as a Function of a Probability Distribution

Since the source $S$ is fully characterized by $P=\left\{p_{1}, \ldots p_{k}\right\}$ (we don't care what the symbols $s_{i}$ actually are, or what they stand for), entropy can also be thought of as a property of a probability distribution function $P$ : the avg uncertainty in the distribution. So we may also write:

$$
H(S)=H(P)=H\left(p_{1}, p_{2}, \ldots, p_{k}\right)=\sum_{i} p_{i} \log \frac{1}{p_{i}}
$$

(Can be generalized to continuous distributions.)

## Properties of Entropy

$$
H(P)=\sum_{i} p_{i} \cdot \log \frac{1}{p_{i}}
$$

1. Non-negative: $H(P) \geq 0$
2. Invariant wrt permutation of its inputs:

$$
H\left(p_{1}, p_{2}, \ldots, p_{k}\right)=H\left(p_{\tau(1)}, p_{\tau(2)}, \ldots, p_{\tau(k)}\right)
$$

3. For any other probability distribution $\left\{q_{1}, q_{2}, \ldots, q_{k}\right\}$ :

$$
H(P)=\sum_{i} p_{i} \cdot \log \frac{1}{p_{i}}<\sum_{i} p_{i} \cdot \log \frac{1}{q_{i}}
$$

4. $H(P) \leq \log k$, with equality iff $p_{i}=1 / k \forall i$
5. The further $P$ is from uniform, the lower the entropy.

## Special Case: $k=2$

Flipping a coin with $P($ "head" $)=p, P($ "tail" $)=1-p$

$$
H(p)=p \cdot \log \frac{1}{p}+(1-p) \cdot \log \frac{1}{1-p}
$$

Notice:

- zero uncertainty/information/surprise at edges
- maximum info at 0.5 (1 bit)
- drops off quickly


## Special Case: $k=2$ (cont.)

Relates to: " 20 questions" game strategy (halving the space).

So a sequence of (independent) 0's-and-1's can provide up to 1 bit of information per digit, provided the 0's and 1's are equally likely at any point. If they are not equally likely, the sequence provides less information and can be compressed.

## The Entropy of English

27 characters (A-Z, space).

100,000 words (avg 5.5 characters each)

- Assuming independence between successive characters:
- uniform character distribution: $\log 27=4.75$ bits/character
- true character distribution: 4.03 bits/character
- Assuming independence between successive words:
- unifrom word distribution: $\log 100,000 / 6.5 \approx 2.55$ bits/character
- true word distribution: 9.45/6.5 $\approx 1.45$ bits/character
- True Entropy of English is much lower!


## Two Sources

Temperature $T$ : a random variable taking on values $t$

$$
\begin{aligned}
& P(T=\text { hot })=0.3 \\
& P(T=\text { mild })=0.5 \\
& P(T=\text { cold })=0.2 \\
& \Longrightarrow H(T)=H(0.3,0.5,0.2)=1.48548
\end{aligned}
$$

huMidity $M$ : a random variable, taking on values $m$

$$
\begin{aligned}
& \mathrm{P}(\mathrm{M}=\text { low })=0.6 \\
& \mathrm{P}(\mathrm{M}=\text { high })=0.4 \\
& \Longrightarrow H(M)=H(0.6,0.4)=0.970951
\end{aligned}
$$

$T, M$ not independent: $P(T=t, M=m) \neq P(T=t) \cdot P(M=m)$

Joint Probability, Joint Entropy

|  | cold | mild | hot |  |
| ---: | :--- | :--- | :--- | :--- |
| low | 0.1 | 0.4 | 0.1 | 0.6 |
| high | 0.2 | 0.1 | 0.1 | 0.4 |
|  | 0.3 | 0.5 | 0.2 | 1.0 |

- $H(T)=H(0.3,0.5,0.2)=1.48548$
- $H(M)=H(0.6,0.4)=0.970951$
- $H(T)+H(M)=2.456431$
- Joint Entropy: consider the space of $(t, m)$ events $H(T, M)=$ $\sum_{t, m} P(T=t, M=m) \cdot \log \frac{1}{P(T=t, M=m)}$ $H(0.1,0.4,0.1,0.2,0.1,0.1)=2.32193$

Notice that $H(T, M)<H(T)+H(M)!!!$

## Conditional Probability, Conditional Entropy

$$
P(T=t \mid M=m)
$$

|  | cold | mild | hot |  |
| ---: | :--- | :--- | :--- | :--- |
| low | $1 / 6$ | $4 / 6$ | $1 / 6$ | 1.0 |
| high | $2 / 4$ | $1 / 4$ | $1 / 4$ | 1.0 |

## Conditional Entropy:

- $H(T \mid M=$ low $)=H(1 / 6,4 / 6,1 / 6)=1.25163$
- $H(T \mid M=$ high $)=H(2 / 4,1 / 4,1 / 4)=1.5$
- Average Conditional Entropy (aka equivocation):

$$
H(T / M)=\sum_{m} P(M=m) \cdot H(T \mid M=m)=
$$

$$
0.6 \cdot H(T \mid M=\text { low })+0.4 \cdot H(T \mid M=h i g h)=1.350978
$$

How much is $M$ telling us on average about $T$ ?

$$
H(T)-H(T \mid M)=1.48548-1.350978 \approx 0.1345 \text { bits }
$$

## Conditional Probability, Conditional Entropy

$$
P(M=m \mid T=t)
$$

|  | cold | mild | hot |
| ---: | :--- | :--- | :--- |
| low | $1 / 3$ | $4 / 5$ | $1 / 2$ |
| high | $2 / 3$ | $1 / 5$ | $1 / 2$ |
|  | 1.0 | 1.0 | 1.0 |

Conditional Entropy:

- $H(M \mid T=$ cold $)=H(1 / 3,2 / 3)=0.918296$
- $H(M \mid T=$ mild $)=H(4 / 5,1 / 5)=0.721928$
- $H(M \mid T=h o t)=H(1 / 2,1 / 2)=1.0$
- Average Conditional Entropy (aka Equivocation):

$$
\begin{aligned}
& H(M / T)=\sum_{t} P(T=t) \cdot H(M \mid T=t)= \\
& 0.3 \cdot H(M \mid T=\text { cold })+0.5 \cdot H(M \mid T=\text { mild })+0.2 \cdot H(M \mid T= \\
& h o t)=0.8364528
\end{aligned}
$$

How much is $T$ telling us on average about $M$ ?

$$
H(M)-H(M \mid T)=0.970951-0.8364528 \approx 0.1345 \text { bits }
$$

## Average Mutual Information

$$
\begin{aligned}
I(X ; Y) & =H(X)-H(X / Y) \\
& =\sum_{x} P(x) \cdot \log \frac{1}{P(x)}-\sum_{x, y} P(x, y) \cdot \log \frac{1}{P(x \mid y)} \\
& =\sum_{x, y} P(x, y) \cdot \log \frac{P(x \mid y)}{P(x)} \\
& =\sum_{x, y} P(x, y) \cdot \log \frac{P(x, y)}{P(x) P(y)}
\end{aligned}
$$

Properties of Average Mutual Information:

- Symmetric (but $H(X) \neq H(Y)$ and $H(X / Y) \neq H(Y / X)$ )
- Non-negative (but $H(X)-H(X / y)$ may be negative!)
- Zero iff $X, Y$ independent
- Additive (see next slide)


## Mutual Information Visualized

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

## Three Sources

From Blachman:
(" /" means " given". ";" means "between". "," means "and".)

- $H(X, Y / Z)=H(\{X, Y\} / Z)$
- $H(X / Y, Z)=H(X /\{Y, Z\})$
- $I(X ; Y / Z)=H(X / Z)-H(X / Y, Z)$

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

$\Longrightarrow$ Can be negative!

- $I(X ; Y, Z)=I(X ; Y)+I(X ; Z / Y)$ (additivity)
- But: $I(X ; Y)=0, I(X ; Z)=0$ doesn't mean $I(X ; Y, Z)=0!!!$


## A Markov Source

Order- $k$ Markov Source: A source that "remembers" the last $k$ symbols emitted.

Ie, the probability of emitting any symbol depends on the last $k$ emitted symbols: $P\left(s_{T=t} \mid s_{T=t-1}, s_{T=t-2}, \ldots, s_{T=t-k}\right)$

So the last $k$ emitted symbols define a state, and there are $q^{k}$ states.

First-order markov source: defined by $q X q$ matrix: $P\left(s_{i} \mid s_{j}\right)$

Example: $S_{T=t}$ is position after $t$ random steps

## Approximating with a Markov Source

A non-Markovian source can still be approximated by one.

Examples: English characters: $C=\left\{c_{1}, c_{2}, \ldots\right\}$

1. Uniform: $H(C)=\log 27=4.75$ bits/char
2. Assuming 0 memory: $H(C)=H(0.186,0.064,0.0127, \ldots)=$ 4.03 bits/char
3. Assuming 1st order: $H(C)=H\left(c_{i} / c_{i-1}\right)=3.32$ bits/char
4. Assuming 2nd order: $H(C)=H\left(c_{i} / c_{i-1}, c_{i-2}\right)=3.1$ bits/char
5. Assuming large order: Shannon got down to $\approx 1$ bit/char

## Modeling an Arbitrary Source

Source $\mathcal{D}(Y)$ with unknown distribution $P_{\mathrm{D}}(Y)$
$\left(\right.$ recall $\left.H\left(P_{\mathrm{D}}\right)=E_{P_{\mathrm{D}}}\left[\log \frac{1}{P_{\mathrm{D}}(Y)}\right]\right)$

Goal: Model (approximate) with learned distribution $P_{M}(Y)$

What's a good model $P_{M}(Y)$ ?

1. RMS error over D's parameters $\Rightarrow$ but $D$ is unknown!
2. Predictive Probability: Maximize the expected log-likelihood the model assigns to future data from $\mathcal{D}$

## Cross Entropy

$$
\begin{aligned}
M^{*} & =\underset{M}{\arg \max } E_{\mathrm{D}}\left[\log P_{M}(Y)\right] \\
& =\underset{M}{\arg \min } E_{\mathrm{D}}\left[\log \frac{1}{P_{M}(Y)}\right] \\
& =C H\left(P_{\mathrm{D}} ; P_{M}\right) \Longleftarrow \text { Cross Entropy }
\end{aligned}
$$

The following are equivalent:

1. Maximize Predictive Probability of $P_{M}$
2. Minimize Cross Entropy $C H\left(P_{\mathrm{D}} ; P_{M}\right)$
3. Minimize the difference between $P_{\mathrm{D}}$ and $P_{M}$ (in what sense?)

## A Distance Measure Between Distributions

Kullback-Liebler distance:

$$
\begin{aligned}
K L\left(P_{\mathrm{D}} ; P_{M}\right) & =C H\left(P_{\mathrm{D}} ; P_{M}\right)-H\left(P_{\mathrm{D}}\right) \\
& =E_{P_{\mathrm{D}}}\left[\log \frac{P_{\mathrm{D}}(Y)}{P_{M}(Y)}\right]
\end{aligned}
$$

Properties of KL distance:

1. Non-negative. $K L\left(P_{\mathrm{D}} ; P_{M}\right)=0 \Longleftrightarrow P_{\mathrm{D}}=P_{M}$
2. Generally non-symmetric

The following are equivalent:

1. Maximize Predictive Probability of $P_{M}$ for distribution D
2. Minimize Cross Entropy $C H\left(P_{\mathrm{D}} ; P_{M}\right)$
3. Minimize the distance $K L\left(P_{\mathrm{D}} ; P_{M}\right)$
