Information Theory

http://elec3004.com

ELEC 3004: **Systems**: Signals & Controls

Timothy Sherry

Lecture 23

(with material from D. MacKay, B. Lathi, K. Harris, D. Corrigan, and more!)

elec3004@itee.uq.edu.au

http://robotics.itee.uq.edu.au/~elec3004/

May 30, 2017

Lecture Schedule:

Week	Date	Lecture Title
1	28-Feb	Introduction
	2-Mar	Systems Overview
2	7-Mar	Systems as Maps & Signals as Vectors
	9-Mar	Systems: Linear Differential Systems
3	14-Mar	Sampling Theory & Data Acquisition
	16-Mar	Aliasing & Antialiasing
4	21-Mar	Discrete Time Analysis & Z-Transform
	23-Mar	Second Order LTID (& Convolution Review)
5	28-Mar	Frequency Response
	30-Mar	Filter Analysis
6	4-Apr	Digital Filters (IIR) & Filter Analysis
	6-Apr	Digital Filter (FIR)
7	11-Apr	Digital Windows
	13-Apr	FFT
	18-Apr	**Holiday**
	20-Apr	
	25-Apr	
8	27-Apr	Active Filters & Estimation
9	2-May	Introduction to Feedback Control
	4-May	Servoregulation/PID
10	9-May	PID & State-Space
	11-May	State-Space Control
11	16-May	Digital Control Design
	18-May	Stability
12	23-May	State Space Control System Design
	25-May	Shaping the Dynamic Response
13	30-May	System Identification & Information Theory
	1-Jun	Summary and Course Review

Follow Along Reading:

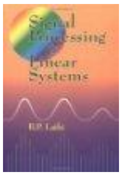

B. P. Lathi
Signal processing and linear systems
1998
TK5102.9.L38 1998

David J. C. MacKay
Information Theory, Inference and Learning Algorithms
2003

http://www.inference.phy.cam.ac.uk/itila/

Today

➔Information Theory 🡐

- Lathi Ch. 2 (?)
  - § 2.7-6 Time Constant and Rate of Information Transmission
- Information Theory!

- Final Exam 2015: http://robotics.itee.uq.edu.au/~elec3004/tutes.html#Final

Next Time

Information Theory!

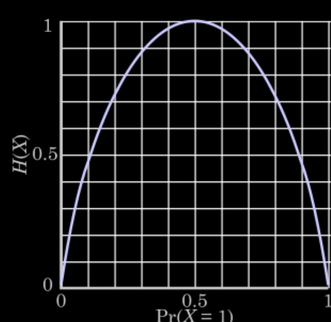

Shannon Information Theory

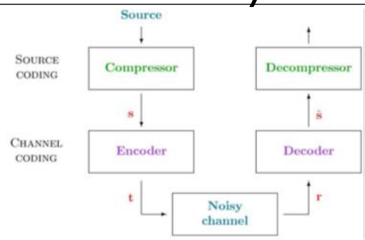

"The fundamental problem of communication is that of reproducing at one point, either exactly or approximately, a message selected at another point."

On the transmission of information over a noisy channel:

- An **information source** that produces a message
- A **transmitter** that operates on the message to create a signal which can be sent through a channel
- A **channel**, which is the medium over which the signal, carrying the information that composes the message, is sent
- A **receiver**, which transforms the signal back into the message intended for delivery
- A **destination**, which can be a person or a machine, for whom or which the message is intended

Information theory is…

- It all starts with **Probability Theory**!

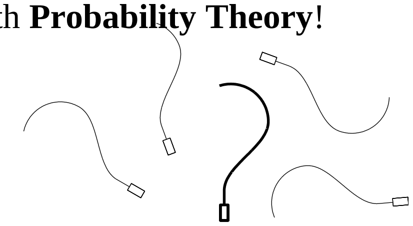

1. (Wikipedia) `Information theory is a branch of applied mathematics, electrical engineering, and computer science involving the quantification of information.`

2. Information theory is **probability theory** where you take logs to base 2.

Entropy

Entropy! ☺

Entropy

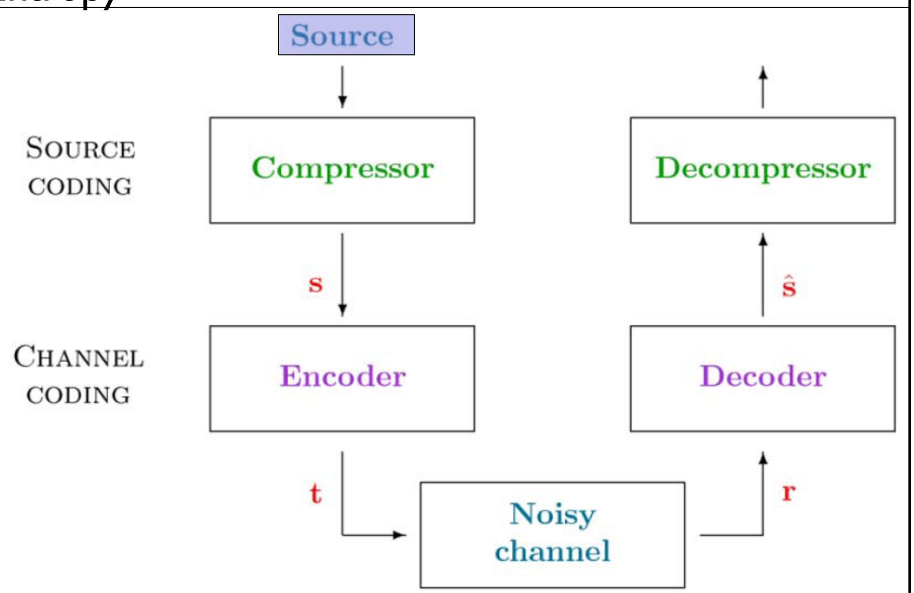

Entropy – a measure of randomness

- **Entropy**

The entropy of a random variable X with a probability mass function $p(x)$ is defined by

$$H(X) = -\sum_{x} p(x)\log_2 p(x). \tag{1.1}$$

Example 1.1.2 Suppose that we have a horse race with eight horses taking part. Assume that the probabilities of winning for the eight horses are $\left(\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{64}, \frac{1}{64}, \frac{1}{64}, \frac{1}{64}\right)$. We can calculate the entropy of the horse race as

$$H(X) = -\frac{1}{2}\log\frac{1}{2} - \frac{1}{4}\log\frac{1}{4} - \frac{1}{8}\log\frac{1}{8} - \frac{1}{16}\log\frac{1}{16} - 4\frac{1}{64}\log\frac{1}{64}$$
$$= 2 \text{ bits.} \tag{1.3}$$

Entropy

- Consider the Random Variable X, the outcome of each consecutive horse race.
- Assigning a binary number to each horse in the race would require $8 = 2^k \rightarrow k = 3\ bits/symbol$
- Shannon states that some encoding scheme on the sequence $X_1, X_2 \dots X_n$ can reduce the average number of bits -> $\mathrm{H}(X) = 2\ bits/symbol$
- We will show you how soon

Connection to Physics

- A macrostate is a description of a system by large-scale quantities such as pressure, temperature, volume.
- A macrostate could correspond to many different microstates i, with probability p_i.
- Entropy of a macrostate is
- $S = -k_B \sum_i p_i \ln p_i$
- Hydrolysis of 1 ATP molecule at body temperature: ~ 20 bits

Entropy: The bent coin

Example 2.1.1 Let

$$X = \begin{cases} 1 & \text{with probability } p, \\ 0 & \text{with probability } 1-p. \end{cases} \quad (2.4)$$

Then

$$H(X) = -p \log p - (1-p)\log(1-p) \stackrel{\text{def}}{=} H(p). \quad (2.5)$$

- Consider X to be the result of a coin toss.
- This coin has been modified to land a Head (1) with probability p and Tails (0) with probability 1-p.
- What is the entropy of X as we vary p?

Plot a graph of H(p) against p.

Entropy

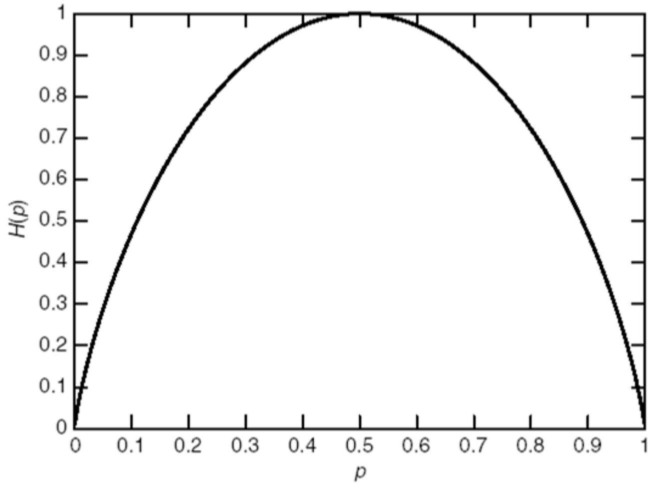

FIGURE 2.1. $H(p)$ vs. p.

Bent/Unfair Coin

- The most unpredictable coin is the fair coin.
- More generally for a variable of k states, a uniform distribution across all k states exhibits maximum entropy
- The observation of a coin with two tails is not random at all.

Example: Mystery Text

```
I.    Emma Woodh*use, hands*me, clever* and rich,*with a comiortab*e home an*
      happy di*position,*seemed to*unite som* of the b*st bless*ngs of
      e*istence;*and had *ived nea*ly twenty *ne year* in the*world w*th
      very*little *o distr*ss or vex*her. *he was*the yo*ngest *f the *wo
      dau*hters *f a most *ffect*onate* indu*gent *ather* and *ad, i* cons*quenc*
      of h*r si*ter'* mar*iage* bee* mis*ress*of h*s ho*se f*om a ver* ea*ly
      *eri*d. *er *oth*r h*d d*ed *oo *ong*ago*for*her to*ha*e *or*
      t*an*an*in*is*in*t *em*mb*an*e *f *er*ca*es*es* a*d*h*r*p*a*e*h*d*b*e*
      *u*p*i*d*b* *n*e*c*l*e*t*w*m*n*a* g**e**e**,**h**h**
      **l**n**i**l**s**r**o**a**o**e**i*
      a***c***n***S***e***y***s***d***s***a***r***e***n***
      W****o****s****i****l****a****g****n****t****a****e****v***

II.   Emma Woodhouse, handsome, clever, and rich, with a comfortable home and happy
      disposition, seemed to unite some of the best blessings of existence; and had lived
      nearly twenty one years in the world with very little to distress or Vex her. She was
      the youngest of the two daughters of a most affectionate, indolent father; and had ,
      in consequence of her sister's marriage , been mistress of his house from a very
      early period. Her mother had died too long ago for her to have more than an
      indistinct remembrance of her caresses; and her place had been supplied by an
      excellent woman as governess, who had fallen little short of a mother in affection.
      Sixteen years had Miss Taylor been in Mr Woodhouse's family , less as a eoverness
      than a friend , very
```

Entropy of English – 2.6 Bits/letter || Raw encoding 27+4 = 31sym = 5.1 Bits/letter

Source: MacKay VideoLectures 02, Slide 5

Source coding

Data Compression/Source coding

- **Removes** redundancy to reduce bit length of messages
- Save bits on common symbols/sequences
- Spend extra bits on the 'surprises', as they very rarely occur
- Lossless (Huffman, Algorithmic, LZ, DEFLATE)
- Lossy (JPEG, MP3, H.265 etc)
- Lossy techniques exploit perceptual dynamic range to expend bits where a human is sensitive, and save where they're not

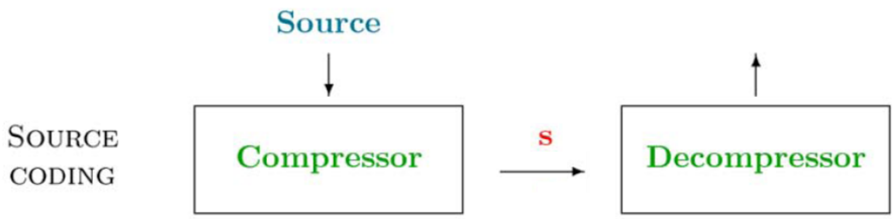

Huffman Coding

- Huffman is the simplest entropy coding scheme
  - It achieves average code lengths no more than 1 bit/symbol of the entropy

- A binary tree is built by combining the two symbols with lowest probability into a dummy node
- The code length for each symbol is the number of branches between the root and respective leaf

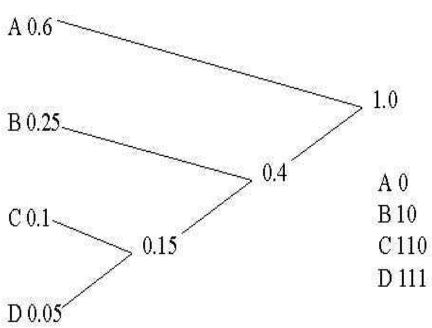

Ex: Bent/Unfair Coin

- Form symbols TT, HT, TH, HH
- P(TT) = 0.81 ; P(HT) = P(TH) = 0.09 ; P(HH) = 0.01
- $X_{TT} = 0 \mid p = 0.81$ $\quad X_{TH} = 10 \mid p = 0.09$
  $X_{HT} = 110 \mid p = 0.09$ $\quad X_{HH} = 111 \mid p = 0.01$

- Transmission cost $0.81 \cdot 1 + 0.09 \cdot 2 + 0.09 \cdot 3 + 0.01 \cdot 3 = 1.29 \frac{bits}{symbol}$
- $H(X) = \ 0.4690 \cdot 2 \ = \ 0.938 \frac{Bits}{symbol}$ – non-optimal code
- Huffman transmission cost $\mathrm{H(X)} \leq \mathrm{C} < H(X) + 1 \frac{bit}{symbol}$
- longer code words = better performance

Huffman performance -> Longer codewords

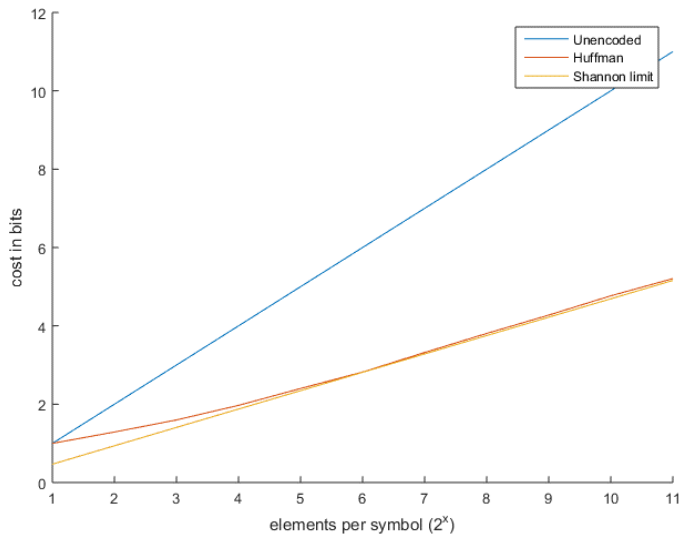

Lossless compression and entropy

- The lower the entropy of the source, the more we can gain via compression.
- The fair coin exhibits no exploitable means of **compressing** a sequence of observations from their direct description, without losing some information in the process
- All forms of compression exploit structure to reduce the average number of bits to describe a sequence.

Channel Coding

Channel and Source Coding

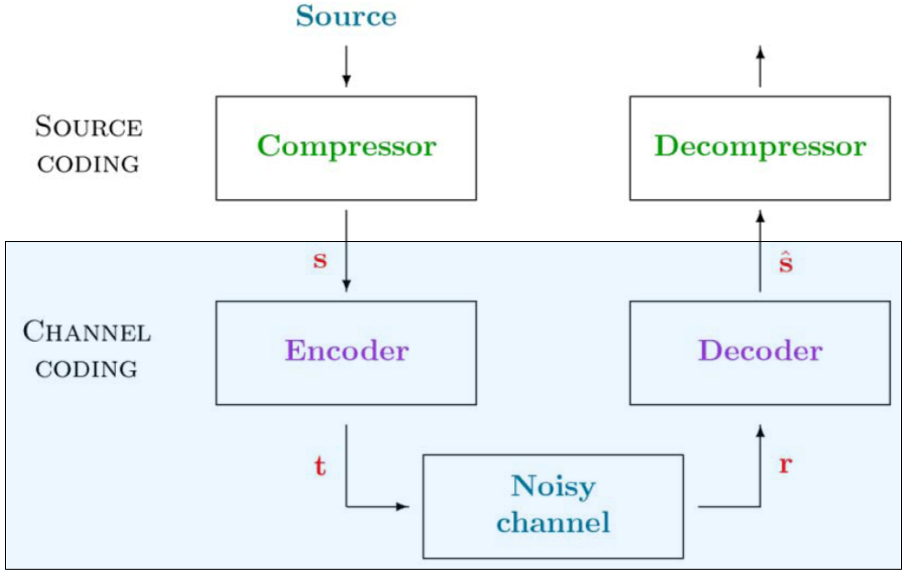

Channel coding

- All physical transmission is susceptible to noise.
- Noise results in errors when a transmission is converted back to the source space
- Channel coding **introduces** redundancy to reduce the probability an error is made in conversion back to the source.

Noisy channels: Binary Symmetric Channel

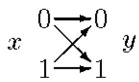

$$P(y=0\,|\,x=0) = 1-f; \quad P(y=0\,|\,x=1) = f;$$
$$P(y=1\,|\,x=0) = f; \quad P(y=1\,|\,x=1) = 1-f.$$

- A 1 bit digital channel where each bit experiences a probability, *f*, of being misinterpreted (flipped)
- Ex: A File of N=10,000 bits is transmitted with ***f*=0.1**

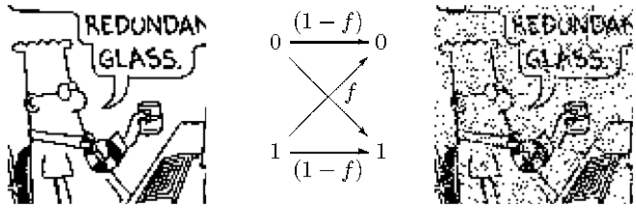

How many bits are flipped?

Assume a Binomial distribution

- Mean: $\mu = Np$
- Variance: $\sigma^2 = Npq$

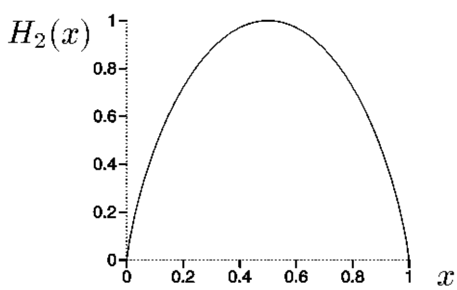

Then:

$$\mu = (0.1)10{,}000$$
$$\sigma^2 = (0.1)(0.9)10{,}000 = 900$$

Thus:

$$1000 \pm 30$$

How to beat this?

Filter like mad!

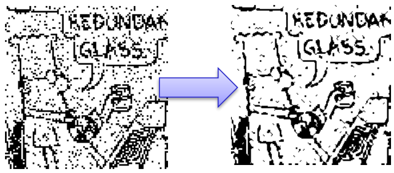

- But this will also affect "non-noisy" signal portions

Matlab: `medfilt2(noisy_im)`

Redundant Transmission!

- Repetition Code 'R3'
  $r = t \; \boxed{modulo\ 2} \; n$
- Does there exist some scheme where P(error) -> 0
- Can we quantify the required redundancy?

Repetition Code 'R3'

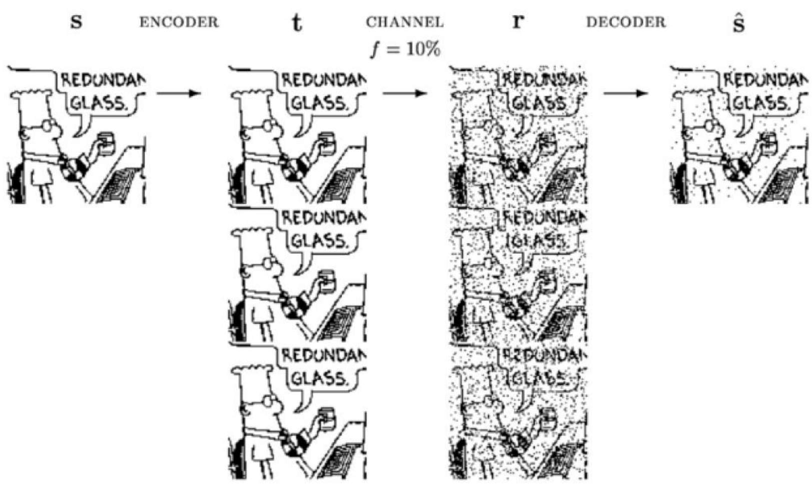

Source: MacKay VideoLectures 01, Slide 30

Hamming distance of R3

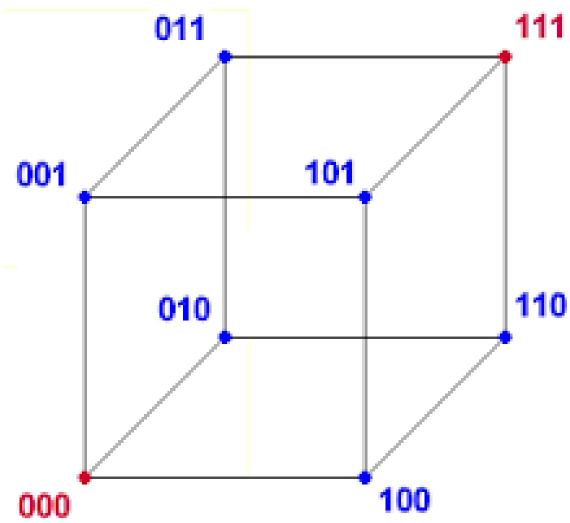

P_{error} of the R_3 code

- $P(r_{000}, s_1) = P(r_{111}, s_0) = 0.1^3 = 0.001$
- $P(r_{100}, s_1) = P(r_{010}, s_1) \dots = 0.1^2 = 0.01$

- $P_{err} = (0.001 + 0.01 \cdot 3) \cdot p(s_1) + (0.001 + 0.01 \cdot 3) \cdot p(s_0)$
- $P_{err} = 0.031$

- The number of errors $\mu = 310 \pm 17.3$
- Approximately 3x reduction in errors, but 3x as many bits required to send the message

Performance of Repetition Codes

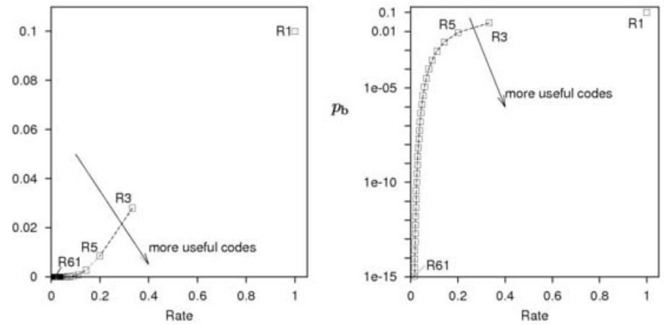

- Want to maximise rate, minimise Error
- Software (e.g. Checksums) can identify errors and correct to higher stringency
- Can we do better?

Source: MacKay VideoLectures 01, Slide 45

The Noisy-channel coding theorem

- The channel capacity for the general channel:

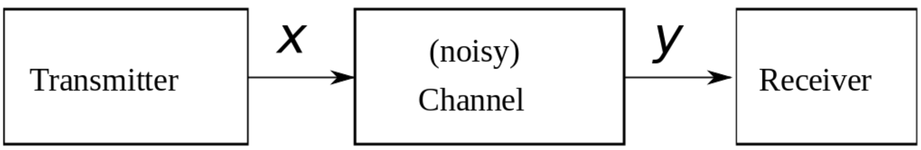

- Is the **mutual information** between X and Y:

$$Capacity = I(X;Y)$$

- We will show Capacity of the BSC is:

$$Capacity_{BSC} = I(X_{BSC};Y_{BSC}) = 1 - H(P_{BSC})$$

- Where H is the **Shannon entropy** of the bent coin with P_{BSC}

Conditional Entropy

- Suppose Alice wants to tell Bob the value of X
  - And they both know the value of a second variable Y.
- Now the optimal code depends on the conditional distribution $p(X|Y)$
- Code length for $X = i$ has length $-\log_2 p(X = i|Y)$
- Conditional entropy measures average code length when they know Y

$$H(X|Y) = -\sum_{X,Y} p(X,Y) \log_2 p(X|Y)$$

Mutual information

- How many bits do Alice and Bob save when they both know Y?

$$I(X;Y) = H(X) - \mathrm{H(X|Y)}$$

$$= \sum_{X,Y} p(X,Y)(-\log_2 p(X) + \log_2 p(X|Y))$$

$$= \sum_{X,Y} p(X,Y) \log_2 \left(\frac{p(X,Y)}{p(X)p(Y)}\right)$$

- Symmetrical in X and Y!
- Amount saved in transmitting X if you know Y equals amount saved transmitting Y if you know X.

Properties of Mutual Information

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

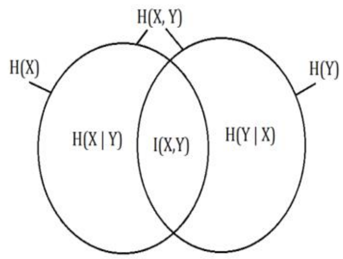

- If $X = Y$, $I(X;Y) = H(X) = H(Y)$

- If X and Y are independent, $I(X;Y) = 0$

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

- In straightforward terms, information is never negative, we are always at least as certain about X|Y vs X

Conditional Entropy & Mutual Information

- Note that the distributions $p(X), p(Y)$ and $p(X|Y)$ may be complex.
- computing $p(X|Y)$ may be intractable
- Mutual Information and Conditional entropy hold for all distributions, but quantifiying them may be intractable

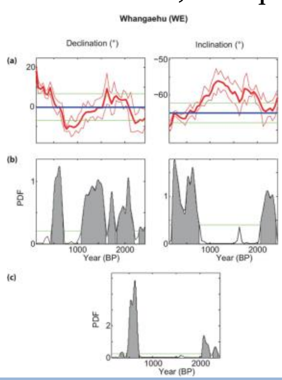

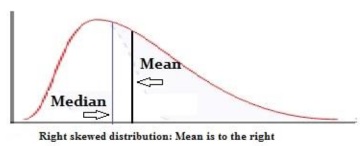

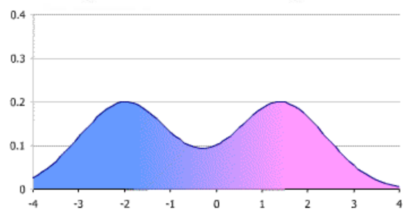

Mutual information of the $BSC_{p=0.1}$

$$P(X_0,Y_0) = P(X_1,Y_1) = 0.5 * 0.9 = 0.45$$
$$P(X_0,Y_1) = P(X_1,Y_0) = 0.5 * 0.1 = 0.05$$
$$P(X_0) = P(X_1) = P(Y_0) = P(Y_1) = 0.5$$

$$I(X;Y) = \sum_{X,Y} p(X,Y)\log_2\left(\frac{p(X,Y)}{p(X)p(Y)}\right)$$
$$I(Y;Y) = 2 \cdot 0.45\log_2\left(\frac{0.45}{0.5^2}\right) + 2 \cdot 0.05\log_2\left(\frac{0.05}{0.5^2}\right)$$
$$I(Y;Y) = 0.7632 - 0.2322 = 0.531$$

Ultimate error free performance

- H(0.1) = 0.4690
- Capacity = 0.531 Bits per Symbol
- Theoretically can obtain a $p_b = 0$ at that rate
- Does not tell us how to obtain that rate

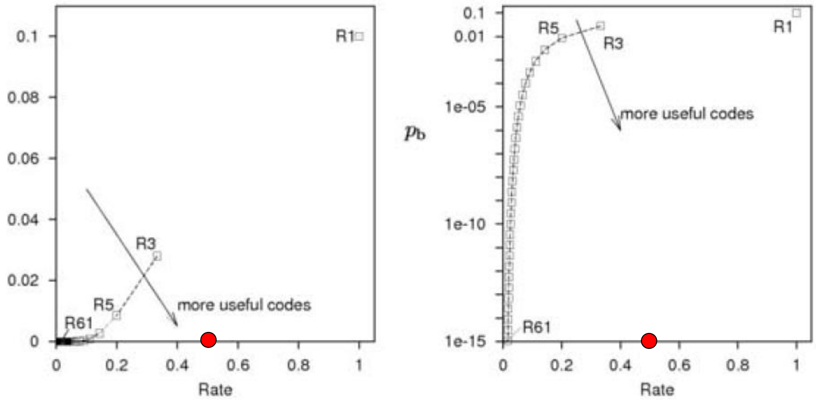

What's Practically Achievable?

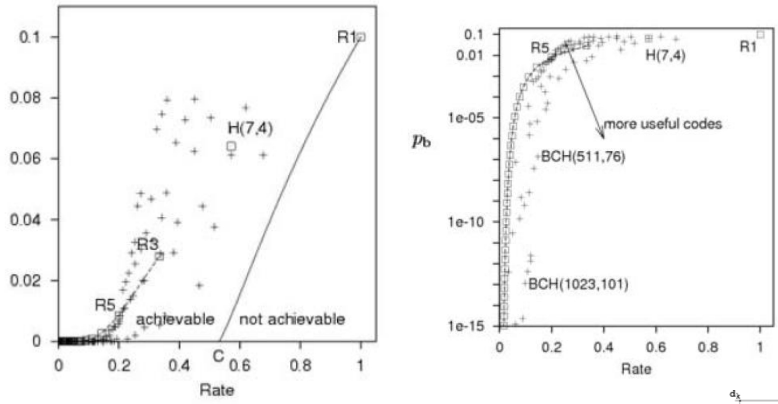

- Interative forward error correction (FEC) codes
- Shannon told us the limit, reaching it is not so easy

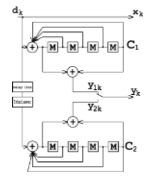

Source: MacKay VideoLectures 01, Slide 54

Break ☺

Communications

Physical layer – Modulation Schemes

- The final piece to transmitting over real communication channels
- Analog modulation (AM (SSB) ,FM , PM, QAM)
- Digital modulation(ASK,PSK,FSK,QAM)
- Both encode the signal on the amplitude (A) or phase (F,P,Q) of a carrier sine frequency
- Will only cover digital Quadrature Amplitude Modulation (QAM)

QAM

- A pulse of the carrier frequency encompasses a symbol.
- The symbol can take on a number of states, encoded as finite levels of amplitude and phase (quadrature)
- Here we need to window our sine carrier to form a pulse
- A rectangular window would produce spectral leakage, and heavy interference – use a raised cosine window
- A RCW results in no inter-symbol-interference

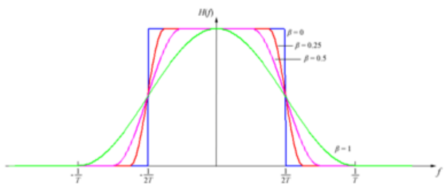

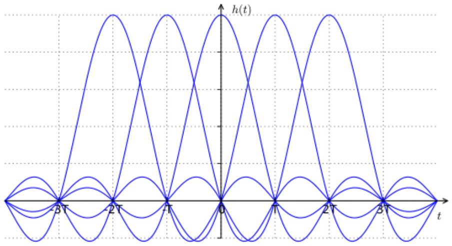

Constellations

- For QAM 16 – 16 discrete locations on the inphase / quadrature plane
- Red lines denote decision boundary

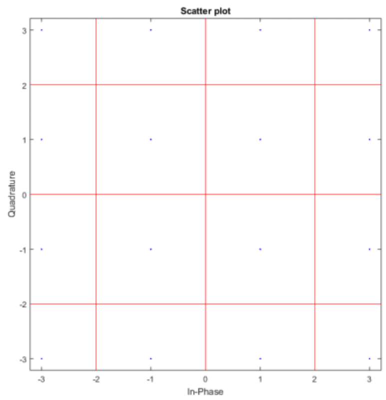

Doubling data rate – QAM4 to QAM16

- Bits per symbol of QAM4 = 2, QAM16 = 4
- Constellation diagrams
- Assuming receive power of $\frac{Eb}{No} = 10dB$, what is the probability that a message sent crosses the descision boundary?
  - $\frac{Eb}{No}$ is the energy per bit received, over the average noise energy at the receiver.

QAM4 v QAM16

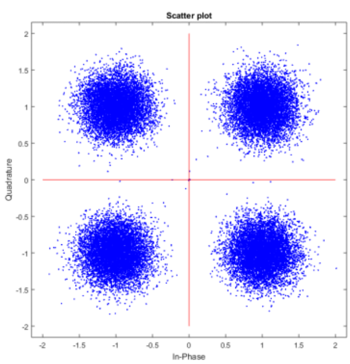

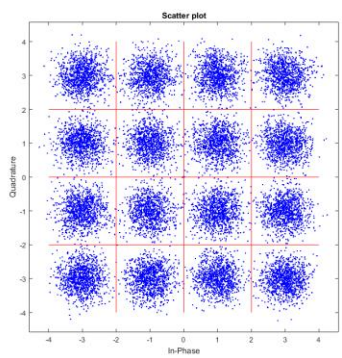

$Perr_4 = 8 \cdot 10^{-6}$ $\quad Perr_{16} = 8 \cdot 10^{-3}$

$\text{BER} = 4 \cdot 10^{-6}$ $\quad 2 \cdot 10^{-3}$

$\text{Rate} = 2 \frac{\text{Bits}}{\text{Symbol}}$ $\quad 4 \frac{\text{Bits}}{\text{Symbol}}$

QAM4 v QAM16

- Double the data rate
- 500 × more likely to flip a bit!
- Generally BER is a specification of the communication protocol
- For example Gigabit Ethernet adjusts the constellation to maintain a raw BER below 10^{-10} (one per 10 gigabits)
  - In other words, your Ethernet connection will produce ~ 7-8 one bit errors per hour of 4k Netflix streaming (~7GB/Hr)
- Source coding reduces this to 10^{-14} (one per 100 terabits)

QAM vs Theoretical maximum

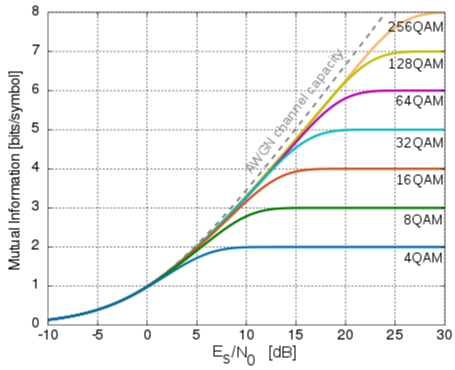

- Using a smaller constellation restricts the channel capacity
- Using too high a constellation requires complex channel coding to exploit

Shannon Capacity of the AWGN channel

- $C = B \log_2(1 + \frac{S}{N})$
- B is the bandwidth S the signal power in real terms, N noise in real terms, and C the capacity in bits.
- Upper bound on error free transmission
- By intelligently selecting QAM constellation and Channel coding schema in tandem capacity can be maximised.
- Requires some coordination between transmitter and sender

Shannon and Weaver: **Models of Communication**

Three "problems" in Communication:

- The technical problem: how accurately can the message be transmitted?
- The semantic problem: how precisely is the meaning "conveyed"?
- The effectiveness problem: how effectively does the received meaning affect behaviour?

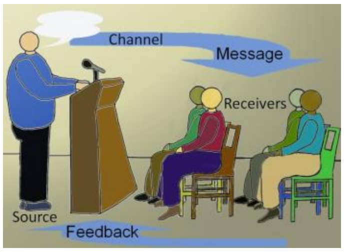

Source:
http://en.wikipedia.org/wiki/File:Transactional_comm_model.jpg

Ex: Morse code

- Code words are shortest for the most common letters
- This means that messages are, on average, sent more quickly.

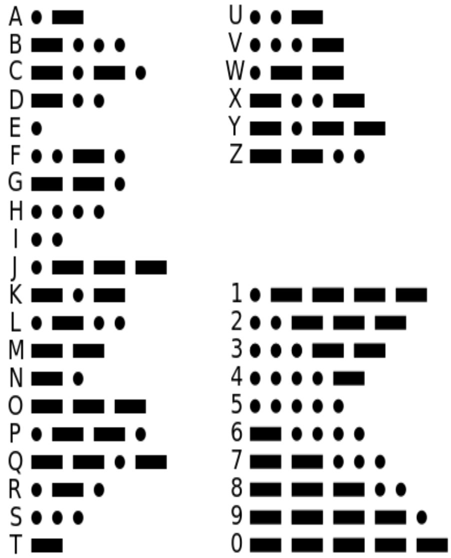

What is the "optimal code"?

- X is a random variable
- Alice wants to tell Bob the value of X (repeatedly)
- What is the best binary code to use?
- How many bits does it take (on average) to transmit the value of X?

Optimal code lengths

- In the optimal code, the word for $X = i$ has length
- $\log_2 \frac{1}{p(X=i)} = -\log_2 p(X = i)$
- For example:

Value of X	Probability	Code word
A	½	0
B	¼	10
C	¼	11

- ABAACACBAACB coded as 010001101110001110
- If code length is not an integer, transmit many letters together

Kullback-Leibler divergence

- Measures the difference between two probability distributions
  - (Mutual information was between two random variables)
- Suppose you use the wrong code for X. How many bits do you waste?

Length of codeword — Length of optimal codeword

- $D_{KL}(p||q) = \sum_x p(x) \left[\log_2 \frac{1}{q(x)} - \log_2 \frac{1}{p(x)}\right]$

$$= \sum_x p(x) \log_2 \frac{p(x)}{q(x)}$$

- $D_{KL}(p||q) \geq 0$, with equality when p and q are the same.
- $I(X;Y) = D_{KL}(p(x,y)||p(x)p(y))$

Continuous variables

- X uniformly distributed between 0 and 1.
- How many bits required to encode X to given accuracy?

Decimal places	Entropy
1	3.3219
2	6.6439
3	9.9658
4	13.2877
5	16.6096
Infinity	Infinity

- Can we make any use of information theory for continuous variables?

K-L divergence for continuous variables

- Even though entropy is infinite, K-L divergence is usually finite.
- Message lengths using optimal and non-optimal codes both tend to infinity as you have more accuracy. But their difference converges to a fixed number.

$$\sum_x p(x)\log_2\frac{p(x)}{q(x)} \to \int p(x)\log_2\frac{p(x)}{q(x)}dx$$

Calculating the Entropy of an Image

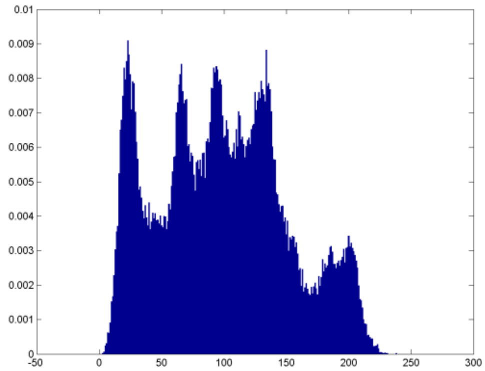

The entropy of lena is = 7.57 bits/pixel approx

Huffman Coding of Lenna

Symbol	Code Length
0	42
1	42
2	41
3	17
4	14
…	…

Average Code Word Length = $\sum_{k=0}^{255} p_k l_k = 7.59 \ bits/pixel$

So the code length is not much greater than the entropy

But this is not very good

- Why?
  - Entropy is not the minimum average codeword length for a source with memory
  - If the other pixel values are known we can predict the unknown pixel with much greater certainty and hence the effective (ie. conditional) entropy is much less.

- Entropy Rate
  - The minimum average codeword length for any source.
  - It is defined as

$$H(\chi) = \lim_{n\to\infty} \frac{1}{N} H(X_1, X_2, ..., X_n)$$

Coding Sources with Memory

- It is very difficult to achieve codeword lengths close to the entropy rate
  - In fact it is difficult to calculate the entropy rate itself – $P(X_1|X_2 \dots X_n)$ is described in R^{n} Space – for lenna n = 65536

- We looked at LZW as a practical coding algorithm
  - Average codeword length tends to the entropy rate if the file is large enough
  - Efficiency is improved if we use Huffman to encode the output of LZW
  - LZ algorithms used in lossless compression formats (eg. .tiff, .png, .gif, .zip, .gz, .rar…)

Efficiency of Lossless Compression

- Lenna (256x256) file sizes
  - Uncompressed tiff - 64.2 kB
  - LZW tiff – 69.0 kB
  - Deflate (LZ77 + Huff) – 58 kB

- Green Screen (1920 x 1080) file sizes
  - Uncompressed – 5.93 MB
  - LZW – 4.85 MB
  - Deflate – 3.7 MB

Differential Coding

- Key idea – code the differences in intensity.

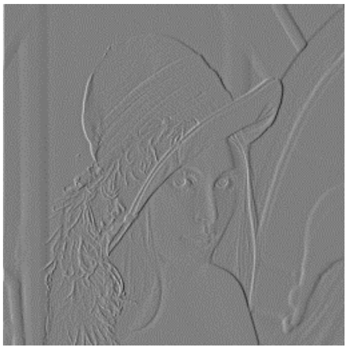

G(x,y) = I(x,y) – I(x-1,y)

Differential Coding

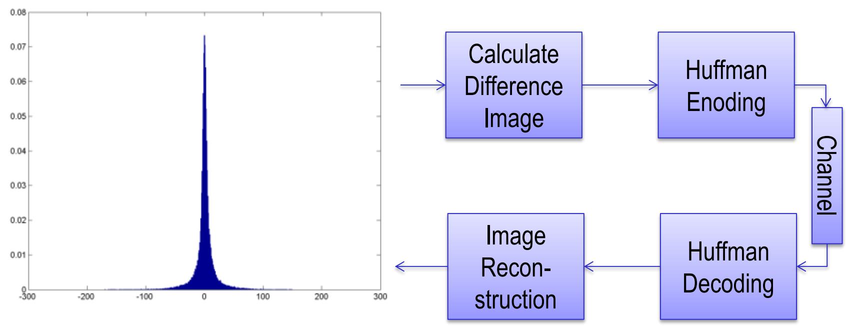

- The entropy is now 5.60 bits/pixel which is much less than 7.57 bits/pixel we had before (despite having twice as many symbols)

Entropy In General

- Entropy of a source is maximised when all signals are equiprobable and is less when a few symbols are much more probable than the others.

Entropy = 7.57 bits/pixel

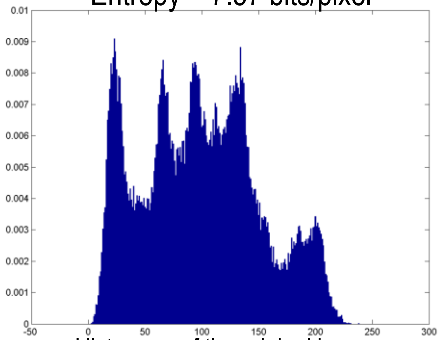

Histogram of the original image

Entropy = 5.6 bits/pixel

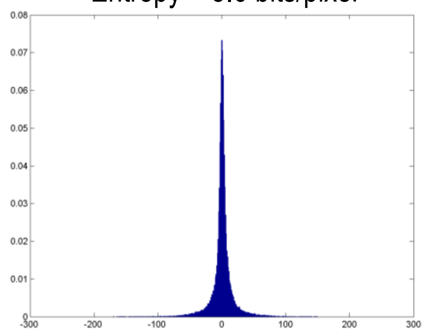

Histogram of the difference image

Lossy Compression

- But this is still not enough compression
  - Trick is to throw away data that has the least perceptual significance

Effective bit rate = 8 bits/pixel

Effective bit rate = 1 bit/pixel (approx)

Next Time…

- **Exam Review!!**

- Review:
  - Chapter 6 of FPW
  - Chapter 13 of Lathi

- Deeper Pondering??

Final Exam **Reviews**

1. June 10, 2016
   - 8-139
   - 4p-6p

2. June 8, 2017
   - 50-T203
   - 9a-12 [+ EBESS BBQ]

- Some Review Notes (from Course Textbooks)

➔ **http://robotics.itee.uq.edu.au/~elec3004/tutes.html**

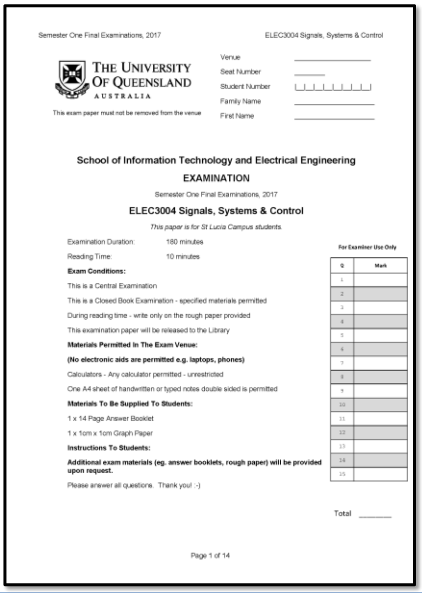
Semester One Final Examinations, 2017 ELEC3004 Signals, Systems & Control

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

This exam paper must not be removed from the venue

Venue ______
Seat Number ______
Student Number ______
Family Name ______
First Name ______

School of Information Technology and Electrical Engineering

EXAMINATION

Semester One Final Examinations, 2017

ELEC3004 Signals, Systems & Control

This paper is for St Lucia Campus students.

Examination Duration: 180 minutes

Reading Time: 10 minutes

Exam Conditions:

This is a Central Examination

This is a Closed Book Examination - specified materials permitted

During reading time - write only on the rough paper provided

This examination paper will be released to the Library

Materials Permitted In The Exam Venue:

(No electronic aids are permitted e.g. laptops, phones)

Calculators - Any calculator permitted - unrestricted

One A4 sheet of handwritten or typed notes double sided is permitted

Materials To Be Supplied To Students:

1 x 14 Page Answer Booklet

1 x 1cm x 1cm Graph Paper

Instructions To Students:

Additional exam materials (eg. answer booklets, rough paper) will be provided upon request.

Please answer all questions. Thank you! :-)

For Examiner Use Only

Q	Mark
1	
2	
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	
15	

Total ______

Page 1 of 14