

Information-Theoretic Modeling

Lecture 5: Source Coding: Practice

Teemu Roos

Department of Computer Science, University of Helsinki

Fall 2014



- 1 Codes
 - Decodable Codes
 - Prefix Codes
 - Kraft-McMillan Theorem

- 2 Optimal Codes
 - Entropy Lower Bound
 - Shannon-Fano Coding




Extension Code

A (binary) **symbol code** $C : \mathcal{X} \rightarrow \{0,1\}^*$ is a mapping from the alphabet \mathcal{X} to the set of finite binary sequences.

The **extension** of code C is the mapping $C^* : \mathcal{X}^* \rightarrow \{0,1\}^*$ obtained by concatenating the codewords $C(x_i)$ for each input symbol x_i :

$$C^*(x_1, x_2, \dots, x_n) = C(x_1)C(x_2) \dots C(x_n) .$$




Extension Code

A (binary) **symbol code** $C : \mathcal{X} \rightarrow \{0,1\}^*$ is a mapping from the alphabet \mathcal{X} to the set of finite binary sequences.

The **extension** of code C is the mapping $C^* : \mathcal{X}^* \rightarrow \{0,1\}^*$ obtained by concatenating the codewords $C(x_i)$ for each input symbol x_i :

$$C^*(x_1, x_2, \dots, x_n) = C(x_1)C(x_2) \dots C(x_n) .$$




Extension Code

A (binary) **symbol code** $C : \mathcal{X} \rightarrow \{0,1\}^*$ is a mapping from the alphabet \mathcal{X} to the set of finite binary sequences.

The **extension** of code C is the mapping $C^* : \mathcal{X}^* \rightarrow \{0,1\}^*$ obtained by concatenating the codewords $C(x_i)$ for each input symbol x_i :

$$C^*(x_1, x_2, \dots, x_n) = C(x_1)C(x_2) \dots C(x_n) .$$




Extension Code

A (binary) **symbol code** $C : \mathcal{X} \rightarrow \{0,1\}^*$ is a mapping from the alphabet \mathcal{X} to the set of finite binary sequences.

The **extension** of code C is the mapping $C^* : \mathcal{X}^* \rightarrow \{0,1\}^*$ obtained by concatenating the codewords $C(x_i)$ for each input symbol x_i :

$$C^*(x_1, x_2, \dots, x_n) = C(x_1)C(x_2) \dots C(x_n) .$$




Extension Code

A (binary) **symbol code** $C : \mathcal{X} \rightarrow \{0,1\}^*$ is a mapping from the alphabet \mathcal{X} to the set of finite binary sequences.

The **extension** of code C is the mapping $C^* : \mathcal{X}^* \rightarrow \{0,1\}^*$ obtained by concatenating the codewords $C(x_i)$ for each input symbol x_i :

$$C^*(x_1, x_2, \dots, x_n) = C(x_1)C(x_2) \dots C(x_n) .$$



(Other types of codes)

For reference, some example of codes that are *not* symbol codes:

Run Length Encoding (RLE): for example, encode aaaaccabbb as

$$(a, 4), (c, 2), (a, 1), (b, 3)$$

Adaptive codes, where the code of a symbol may change based on what symbols have appeared previously




Note that coding blocks of b bits for some constant b is still a symbol code, with alphabet size $|\mathcal{X}| = 2^b$.

Decodable Codes

Decodable Code

Code C is (uniquely) **decodable** iff its extension C^* is a one-to-one mapping, i.e., iff

$$(x_1, \dots, x_n) \neq (y_1, \dots, y_m) \Rightarrow C^*(x_1, \dots, x_n) \neq C^*(y_1, \dots, y_m) .$$

-  A code with codewords $\{0, 1, 10, 11\}$ is *not* uniquely decodable: What does 10 mean?
-  A code with codewords $\{00, 01, 10, 11\}$ is uniquely decodable: Each pair of bits can be decoded individually.
-  A code with codewords $\{0, 01, 011, 0111\}$ is also uniquely decodable: What does 0011 mean?

Prefix Codes

An important subset of decodable codes is the set of **prefix(-free) codes**.

Prefix Code

A code $C : \mathcal{X} \rightarrow \{0, 1\}^*$ is called a **prefix code** iff no codeword is a prefix of another.

It is easily seen that all prefix codes are uniquely decodable: each symbol can be decoded as soon as its codeword is read. Therefore, prefix codes are also called *instantaneous* codes.

- ✗ A code with codewords $\{0, 01, 011, 0111\}$ is uniquely decodable *but not prefix-free*: e.g., 0 is a prefix of 01.
- ✓ A code with codewords $\{0, 10, 110, 111\}$ is prefix-free.

Kraft Inequality

The codeword lengths of a prefix codes satisfy the following important property.

Kraft Inequality

The codeword lengths ℓ_1, \dots, ℓ_m of any (binary) prefix code satisfy

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

Conversely, given a set of codeword lengths that satisfy this inequality, there is a prefix code with these codeword lengths.

Kraft Inequality: Proof

Proof: Assume first we have a prefix code with codeword lengths ℓ_1, \dots, ℓ_m .

Create a binary tree with the code-words as leafs.

Start from the root and turn left (0) or right (1) with equal probability. Stop when you hit a code-word (or not at all).

Probability of hitting code-word $C(x_i)$ is $2^{-\ell_i}$. (Mind the prefix property!)

Sum of probabilities: $\sum_{i=1}^m \Pr[C(x_i)] = \sum_{i=1}^m 2^{-\ell_i} \leq 1$.

This concludes the first half of the proof.

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		010	0100		
			0101		
	011	0110			
		0111			
	1	10	100	1000	
				1001	
		101	1010		
			1011		
		110	1100		
			1101		
111	1110				
	1111				

✓ Codewords $\{0, 10, 110, 111\}$

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		01	010	0100	
				0101	
	011	0110			
		0111			
	1	10	100	1000	
				1001	
		101	1010		
			1011		
11		110	1100		
			1101		
111	1110				
	1111				

✗ Kraft inequality violated. \Rightarrow Not decodable.

Kraft Inequality

Total budget

0	00	000	0000	
			0001	
		001	0010	
	01		0011	
		010	0100	
		011	0101	
1	10		0110	
			0111	
		100	1000	
	11		1001	
		101	1010	
		110	1011	
		1100		
		1101		
		1110		
		1111		

✓ Fixed-length code

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		010	0100		
	01		0101		
		011	0110		
			0111		
	1	10	100	1000	
				1001	
		101	1010		
		1011			
11		110	1100		
			1101		
		111	1110		
		1111			
Kraft?	Decodable?	Prefix-free?			

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		010	0100		
			0101		
	1	01	011	0110	
				0111	
		10	100	1000	
				1001	
		101	1010		
	1011				
11	110	1100			
		1101			
	111	1110			
		1111			

Kraft? ✓ Decodable? ✓ Prefix-free? ✗

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		010	0100		
			0101		
	1	01	011	0110	
				0111	
		10	100	1000	
				1001	
		101	1010		
			1011		
11	110	1100			
		1101			
	111	1110			
		1111			

Kraft?

Decodable?

Prefix-free?

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
			001	0010	
			0011		
		01	010	0100	
				0101	
	011		0110		
			0111		
	1	10	100	1000	
				1001	
			101	1010	
			1011		
11		110	1100		
			1101		
	111	1110			
		1111			

Kraft? Decodable? Prefix-free?

Kraft Inequality: Proof

Assume now we are given a set of integers such that

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

Let $L = \max_i \ell_i$ be the maximum code-word length. We construct a prefix code with codeword lengths ℓ_i .

We can assume that the lengths are sorted so that $\ell_1 \leq \ell_2 \leq \dots \leq \ell_m$, and that $\ell_1 > 0$.


Kraft Inequality: Proof

Now create a binary tree of depth L . Associate codewords to nodes of the tree as previously (0 turn left, 1 turn right).

Repeat the following for $i = 1, \dots, m$.

- 1 Choose the leftmost node at depth ℓ_1 (depth of the root is 0). Assign the corresponding string as codeword i .
- 2 Remove the chosen node and its descendants from the tree.


Kraft Inequality: example



$l_1 = 1$ a
 $l_2 = 3$ b
 $l_3 = 2$ c
 $l_4 = 3$ d

$C(a) = \dots$
 $C(b) = \dots$
 $C(c) = \dots$
 $C(d) = \dots$


Kraft Inequality: example



$l_1 = 1$ a
 $l_2 = 2$ c
 $l_3 = 3$ b
 $l_4 = 3$ d

$C(a) = \dots$
 $C(c) = \dots$
 $C(b) = \dots$
 $C(d) = \dots$


Kraft Inequality: example



$l_1 = 1$ a
 $l_2 = 2$ c
 $l_3 = 3$ b
 $l_4 = 3$ d

$C(a) = 0$
 $C(c) = \dots$
 $C(b) = \dots$
 $C(d) = \dots$


Kraft Inequality: example



$l_1 = 1$ a
 $l_2 = 2$ c
 $l_3 = 3$ b
 $l_4 = 3$ d

$C(a) = 0$
 $C(c) = \dots$
 $C(b) = \dots$
 $C(d) = \dots$


Kraft Inequality: example



$l_1 = 1$ a
 $l_2 = 2$ c
 $l_3 = 3$ b
 $l_4 = 3$ d

$C(a) = 0$
 $C(c) = 10$
 $C(b) = \dots$
 $C(d) = \dots$

Kraft Inequality: example



$$l_1 = 1 \quad a$$

$$l_2 = 2 \quad c$$

$$l_3 = 3 \quad b$$

$$l_4 = 3 \quad d$$


$$C(a) = 0$$

$$C(c) = 10$$


$$C(b) = \dots$$

$$C(d) = \dots$$

Kraft Inequality: example

 $l_1 = 1$ a $l_2 = 2$ c $l_3 = 3$ b $l_4 = 3$ d $C(a) = 0$ $C(c) = 10$ $C(b) = 110$ $C(d) = \dots$

Kraft Inequality: example



$l_1 = 1$ a

$l_2 = 2$ c

$l_3 = 3$ b

$l_4 = 3$ d


$C(a) = 0$

$C(c) = 10$

$C(b) = 110$

$C(d) = \dots$

Kraft Inequality: example



$$l_1 = 1 \quad a$$

$$l_2 = 2 \quad c$$

$$l_3 = 3 \quad b$$

$$l_4 = 3 \quad d$$

$$C(a) = 0$$


$$C(c) = 10$$

$$C(b) = 110$$

$$C(d) = 111$$

Kraft Inequality: Proof

The order of using up the tree is important.



$\ell_1 = 3$ b
 $\ell_2 = 2$ c
 $\ell_3 = 2$ d
 $\ell_4 = 1$ a

$C(b) = 000$
 $C(c) = 01$
 $C(d) = 100$
 $C(a) = ?$

No code of length 1 left for a!

Kraft Inequality: Proof

Notice that since $\ell_i \leq \ell_{i+1}$, we keep moving deeper into the tree. Therefore the descendants of the node chosen for ℓ_i cannot include any descendants of nodes chosen for $\ell_1, \dots, \ell_{i-1}$.

Therefore, in step i , exactly $2^{L-\ell_i}$ leaves are removed.

Total number of removed leaves is

$$\sum_{i=1}^m 2^{L-\ell_i} = 2^L \sum_{i=1}^m 2^{-\ell_i} \leq 2^L,$$

so that the tree will not run out of leaves (2^L of them) before we are done.

This concludes the second direction of the proof. \square

Kraft Inequality

Question: What if the inequality is satisfied strictly, i.e., the sum of the terms in the sum equals *less* than one:

$$\sum_{i=1}^m 2^{-\ell_i} < 1 .$$

Then it is possible to make the codewords shorter and still have a decodable (prefix) code.

Kraft Inequality

Total budget	0	00	000	0000	
				0001	
			001	0010	
		01	010	0100	
			011	0101	
				0110	
	1	10	100	0111	
				1000	
				1001	
		11	101	1010	
				1011	
			110	1100	
				1101	
			111	1110	
	1111				

Not all of budget used. \Rightarrow Some codewords can be made shorter.

Kraft Inequality

Total budget

0	00	000	0000	
			0001	
		001	0010	
	01	010	0100	
			0101	
		011	0110	
1	10	100	1000	
			1001	
		101	1010	
	11	110	1100	
			1101	
		111	1110	
		1111		

“Kraft tight” / complete code.

Kraft–McMillan Theorem

The Kraft inequality restricts the codeword lengths of prefix codes. Could we do much better if we would only require decodability?

In fact it can be shown that we do not lose anything at all!

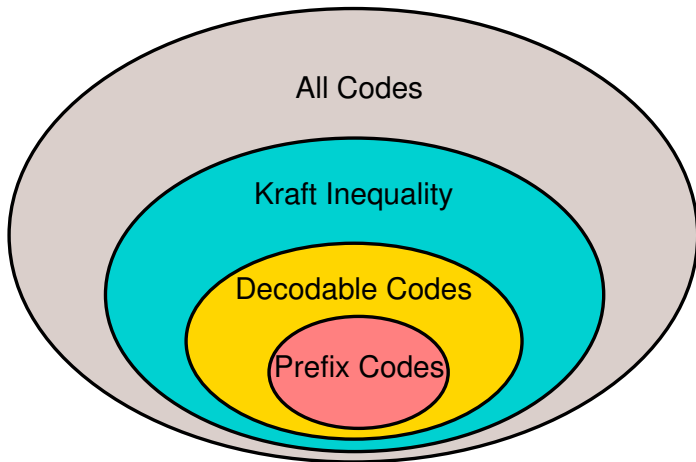
Kraft-McMillan Theorem

The codeword lengths ℓ_1, \dots, ℓ_m of any **uniquely decodable** (binary) code satisfy

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

Conversely, given a set of codeword lengths that satisfy this inequality, there is a uniquely decodable (prefix) code with these codeword lengths.

Kraft-McMillan Theorem & Codes



- 1 Codes
 - Decodable Codes
 - Prefix Codes
 - Kraft-McMillan Theorem

- 2 Optimal Codes
 - Entropy Lower Bound
 - Shannon-Fano Coding



Codelengths and Probabilities

Let l_1, \dots, l_m be the codeword lengths of a uniquely decodable code $C : \mathcal{X} \rightarrow \{0, 1\}^*$. By the Kraft-McMillan theorem we have

$$c = \sum_{i=1}^m 2^{-l_i} \leq 1 .$$

Define a probability mass function $p : \mathcal{X} \rightarrow [0, 1]$ as follows:

$$p_i = \frac{2^{-l_i}}{c} \quad \Leftrightarrow \quad l_i = \log_2 \frac{1}{cp_i} ,$$

where c is given above.

Function p is indeed a pmf:

- 1 Non-negative: $p(x) \geq 0$ for all $x \in \mathcal{X}$.
- 2 Sums to one: $\sum_{x \in \mathcal{X}} p(x) = \sum_{i=1}^m \frac{1}{c} 2^{-l_i} = \frac{c}{c} = 1 .$

Codelengths and Probabilities

Assuming that the code is “Kraft tight”, $c = 1$, then under the pmf p corresponding to the codeword lengths ℓ_1, \dots, ℓ_m , the expected codeword length is

$$\begin{aligned} E[\ell(X)] &= \sum_{i=1}^m 2^{-\ell_i} \ell_i \\ &= \sum_{i=1}^m p_i \log_2 \frac{1}{p_i} = H(X) . \end{aligned}$$

This is the best we can hope for:

The expected codelength of any uniquely decodable code is at least the entropy:

$$E[\ell(X)] \geq H(X) .$$

Entropy Lower Bound

$$E[\ell(X)] \geq H(X) .$$

Proof.

$$\begin{aligned} E[\ell(X)] - H(X) &= \sum_{x \in \mathcal{X}} p(x) \ell(x) - \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{1}{p(x)} \\ &= \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{1}{2^{-\ell_x}} - \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{1}{p(x)} \\ &= \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{p(x)}{2^{-\ell_x}} \\ &= \sum_{x \in \mathcal{X}} p(x) \left[\log_2 \frac{p(x)}{q(x)} + \log_2 \frac{1}{c} \right] \quad \boxed{q(x) = \frac{2^{-\ell(x)}}{c}} \\ &= D(p \parallel q) + \log_2 \frac{1}{c} \geq 0 . \end{aligned}$$

Entropy Lower Bound

So what have we learned? For decodable symbols codes:

- 1 $E[\ell(X)] - H(X) = D(p \parallel q) + \log_2 \frac{1}{c}$, where $q(x) = \frac{2^{-\ell(x)}}{c}$.
- 2 $E[\ell(X)] \geq H(X)$.
- 3 If $\ell(x) = \log_2 \frac{1}{p(x)}$, then $E[\ell(X)] = H(X)$. **Optimal!**

Note also that for a sequence X_1, \dots, X_n the expected codelength becomes

$$E[\ell(X_1, \dots, X_n)] = E\left[\sum_{i=1}^n \ell(X_i)\right] = \sum_{i=1}^n E[\ell(X_i)] = nH(X) .$$

! By Shannon's Noiseless Channel Coding Theorem, this is optimal among all codes, **not only symbol codes**. Fine print: only if X_i i.i.d.!

Codelengths and Probabilities

The only problem with the $\ell(x) = \log_2 \frac{1}{p(x)}$ codeword choice is the requirement that codeword lengths must be **integers** (try to think about a codeword with length 0.123, for instance).

The simplest solution is to round upwards:

Shannon's Code

Given a pmf, the **Shannon code** has the codeword lengths

$$\ell(x) = \left\lceil \log_2 \frac{1}{p(x)} \right\rceil \quad \text{for all } x \in \mathcal{X}.$$

Alice in Wonderland



Shannon's code: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$
■	a	0.0644	3.9	4
■	b	0.0108	6.5	7
■	c	0.0178	5.8	6
■	d	0.0359	4.7	5
■	e	0.0991	3.3	4
■	f	0.0147	6.0	7
■	g	0.0184	5.7	6
■	h	0.0535	4.2	5
■	i	0.0551	4.1	5
■	j	0.0011	9.8	10
■	k	0.0083	6.8	7
■	l	0.0343	4.8	5
		⋮		
■	y	0.0165	5.9	6
■	z	0.0005	10.7	11
■		0.2111	2.2	3

$$H(X) = 4.03$$

Shannon (1948):

- 1 Sort by probability.

Shannon's code: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$
█		0.2111	2.2	3
█	e	0.0991	3.3	4
█	t	0.0781	3.6	4
█	a	0.0644	3.9	4
█	o	0.0598	4.0	5
█	i	0.0551	4.1	5
█	h	0.0535	4.2	5
█	n	0.0516	4.2	5
█	s	0.0475	4.3	5
█	r	0.0401	4.6	5
█	d	0.0359	4.7	5
█	l	0.0343	4.8	5
		⋮		
	x	0.0011	9.8	10
	j	0.0011	9.8	10
	z	0.0005	10.7	11

$$H(X) = 4.03$$

Shannon (1948):

- 1 Sort by probability.
- 2 Choose codewords in order, avoiding prefixes. (same idea as in the proof of the Kraft inequality)

Shannon's code: Example

Total budget	0	00	000	0000	
				0001	
		001	0010		
			0011		
		010	0100		
	01	0101			
		0110			
		0111			
		1000			
		1001			
1	10	1010			
		1011			
	11	1100			
		1101			
		1110			
		1111			

Codeword lengths (3, 4, 4, 4, 5, 5, 5, 5, ..., 10, 10, 11)

Shannon's code: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$	$C(X)$
█		0.2111	2.2	3	000
█	e	0.0991	3.3	4	0010
█	t	0.0781	3.6	4	0011
█	a	0.0644	3.9	4	0100
█	o	0.0598	4.0	5	01010
█	i	0.0551	4.1	5	01011
█	h	0.0535	4.2	5	01100
█	n	0.0516	4.2	5	01101
█	s	0.0475	4.3	5	01110
█	r	0.0401	4.6	5	01111
█	d	0.0359	4.7	5	10000
█	l	0.0343	4.8	5	10001
		⋮			
	x	0.0011	9.8	10	1010111101
	j	0.0011	9.8	10	1010111110
	z	0.0005	10.7	11	101011111110

$$H(X) = 4.03$$

$$E[\ell(X)] = 4.60$$

$$E[\ell(X)] - H(X) = 0.57$$

Shannon's code

The expected codeword length of Shannon's code is

$$\begin{aligned} E[\ell(X)] &= E\left[\left\lceil \log_2 \frac{1}{p(X)} \right\rceil\right] \\ &\leq E\left[\log_2 \frac{1}{p(X)} + 1\right] = H(X) + 1 . \end{aligned}$$

In the Alice example we had

$$E[\ell(X)] - H(X) = 4.60 - 4.03 = 0.57 \leq 1 .$$

Fano: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$
■	a	0.0644	3.9
■	b	0.0108	6.5
■	c	0.0178	5.8
■	d	0.0359	4.7
■	e	0.0991	3.3
■	f	0.0147	6.0
■	g	0.0184	5.7
■	h	0.0535	4.2
■	i	0.0551	4.1
■	j	0.0011	9.8
■	k	0.0083	6.8
■	l	0.0343	4.8
	⋮		
■	y	0.0165	5.9
■	z	0.0005	10.7
■		0.2111	2.2

(Shannon-)Fano code:

- 1 Sort by probability.

Fano: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$
████████		0.2111	2.2
███	e	0.0991	3.3
██	t	0.0781	3.6
█	a	0.0644	3.9
█	o	0.0598	4.0
█	i	0.0551	4.1
█	h	0.0535	4.2
█	n	0.0516	4.2
█	s	0.0475	4.3
█	r	0.0401	4.6
█	d	0.0359	4.7
█	l	0.0343	4.8
		⋮	
	x	0.0011	9.8
	j	0.0011	9.8
	z	0.0005	10.7

(Shannon-)Fano code:

- 1 Sort by probability.
- 2 Divide in two equally probable parts (as equal as possible)
- 3 Add '0' to the codewords in the first part, '1' to the others.
- 4 Repeat recursively for both parts.

Fano: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$	$C(X)$
█		0.2111	2.2	2	0
█	e	0.0991	3.3	4	0
█	t	0.0781	3.6	4	0
█	a	0.0644	3.9	4	0
█	o	0.0598	4.0	4	0
█	i	0.0551	4.1	4	1
█	h	0.0535	4.2	4	1
█	n	0.0516	4.2	4	1
█	s	0.0475	4.3	5	1
█	r	0.0401	4.6	5	1
█	d	0.0359	4.7	5	1
█	l	0.0343	4.8	5	1
		⋮			
	x	0.0011	9.8	10	1
	j	0.0011	9.8	10	1
	z	0.0005	10.7	10	1

Fano: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$	$C(X)$
█		0.2111	2.2	2	00
█	e	0.0991	3.3	4	01
█	t	0.0781	3.6	4	01
█	a	0.0644	3.9	4	01
█	o	0.0598	4.0	4	01
█	i	0.0551	4.1	4	10
█	h	0.0535	4.2	4	10
█	n	0.0516	4.2	4	10
█	s	0.0475	4.3	5	10
█	r	0.0401	4.6	5	10
█	d	0.0359	4.7	5	11
█	l	0.0343	4.8	5	11
		⋮			
	x	0.0011	9.8	10	11
	j	0.0011	9.8	10	11
	z	0.0005	10.7	10	11

Fano: Example

	X	$p(X)$	$\log_2 \frac{1}{p(X)}$	$\ell(X)$	$C(X)$
█		0.2111	2.2	2	00
█	e	0.0991	3.3	4	0100
█	t	0.0781	3.6	4	0101
█	a	0.0644	3.9	4	0110
█	o	0.0598	4.0	4	0111
█	i	0.0551	4.1	4	1000
█	h	0.0535	4.2	4	1001
█	n	0.0516	4.2	4	1010
█	s	0.0475	4.3	5	10110
█	r	0.0401	4.6	5	10111
█	d	0.0359	4.7	5	11000
█	l	0.0343	4.8	5	11001
		⋮			
	x	0.0011	9.8	10	1111111101
	j	0.0011	9.8	10	1111111110
	z	0.0005	10.7	10	1111111111

$$H(X) = 4.03$$

$$E[\ell(X)] = 4.07$$

$$E[\ell(X)] - H(X) = 0.04$$

Shannon-Fano Code

The expected codeword length of the Shannon-Fano code is

$$\begin{aligned} E[\ell(X)] &\leq E\left[\left\lceil \log_2 \frac{1}{p(X)} \right\rceil\right] \\ &\leq E\left[\log_2 \frac{1}{p(X)} + 1\right] = H(X) + 1 . \end{aligned}$$

In the Alice example we had

$$E[\ell(X)] - H(X) = 4.06 - 4.03 = 0.04 \leq 1 .$$

Is this optimal? Not necessarily — Huffman!

Things to come

Next:

- Huffman Coding: the optimal symbol code
- Arithmetic Coding: even better!!!