

Natural Distributed Algorithms

- Lecture 3 -

Rumor Spreading in the Noisy PULL Model



Emanuele Natale
CNRS - UCA

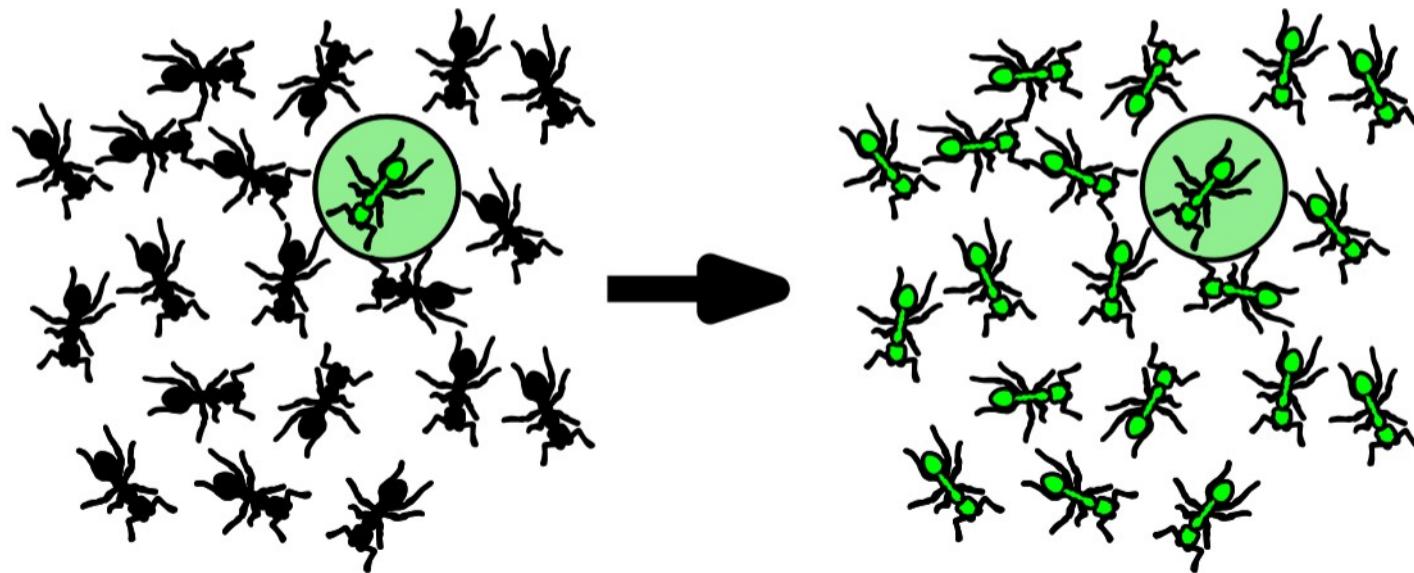
CdL in Informatica
Università degli Studi di Roma “Tor Vergata”



UNIVERSITÉ
CÔTE D'AZUR



Rumor Spreading Problem



- One **source** node in a **special state**
- **Goal configuration:** all agents in the **special state**

Stochastic Interactions: PULL Model



Desired features

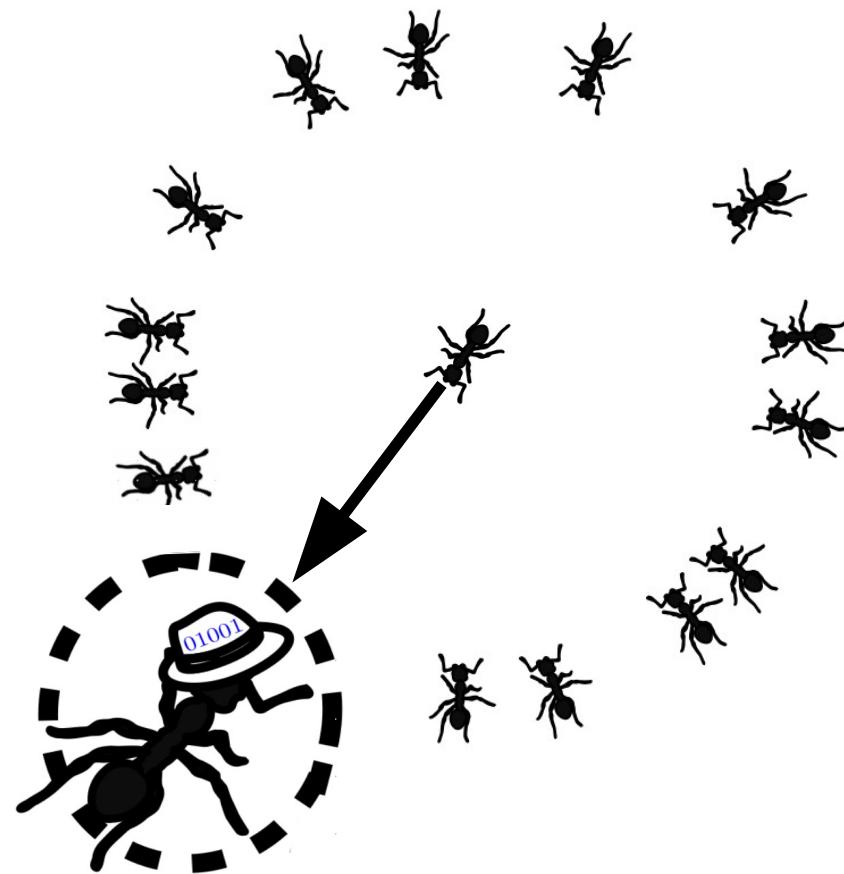
- Stochastic
- Parsimonious (Anonymous)
- **PASSIVE** (uni-directional)

(Uniform) PULL model

[Demers '88]

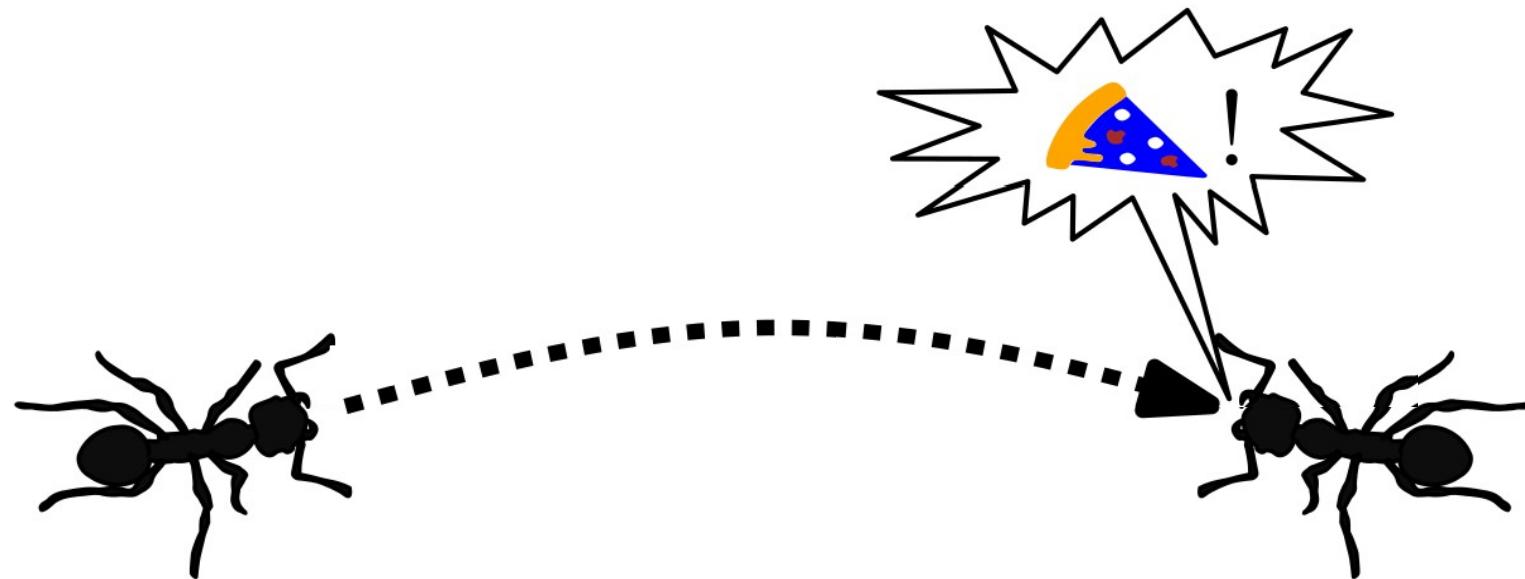
(single binary message)

- Discrete parallel time
- At each round each agent **shows one bit** message
- **The agent can observe the message shown by one agent chosen independently and uniformly at random**



Noisy Communication

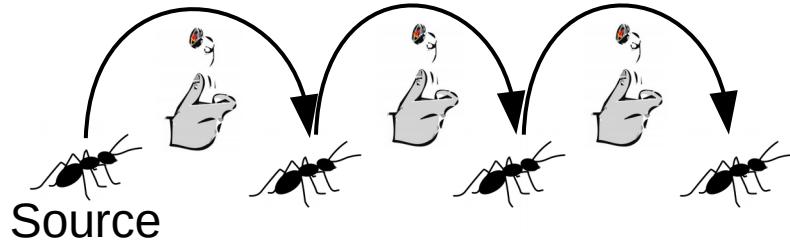
Before being received,
each bit is flipped with probability $\frac{1}{2} - \epsilon$



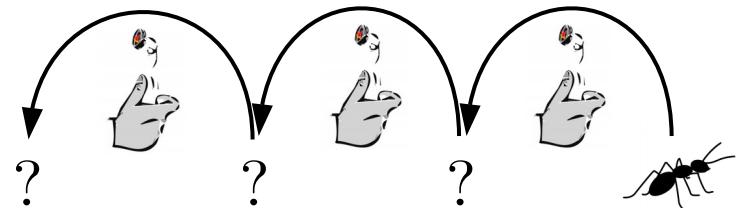
1 bit per message,
noise flips it with prob $\frac{1}{2} + \epsilon$

Rumor Spreading in the **Noisy PULL** Model: The Source is Lost in Too Much Noise

In the **PUSH** model we could estimate the distance that a message travelled from the source, thanks to the fact that in the B.B.S. algorithm agents can avoid communicating



In the **PULL** model we can only ask one bit to a random node:



What is the probability to receive the right opinion message from the source?

How many messages do we need then?

Broadcast-PULL: A Simpler Model

In the PULL Model each agent gets a message from another randomly-chosen agent



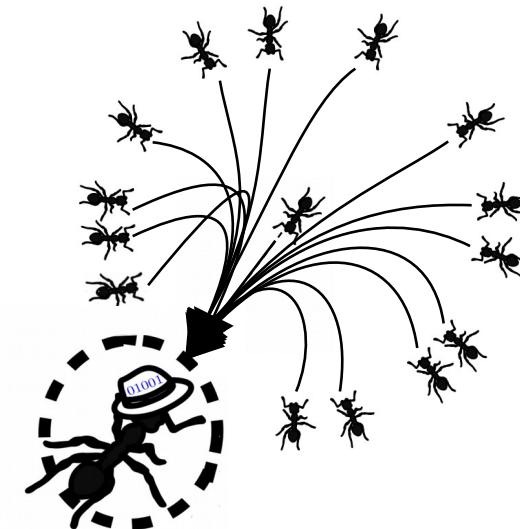
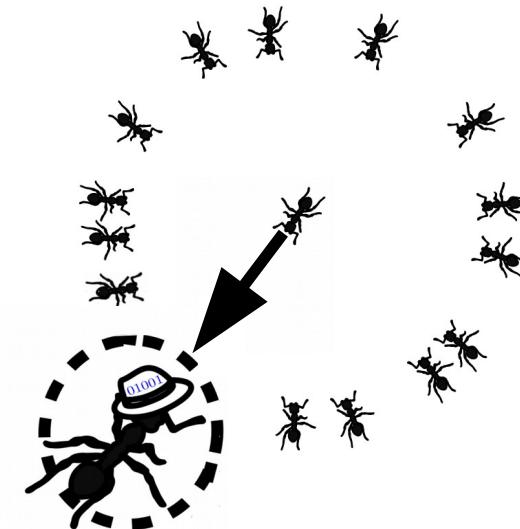
Different agents get different information...
Difficult to analyze!



Pòlya's principle:
"If you can't solve a problem, then there is an easier problem you can solve: find it."



Broadcast-PULL Model:
Each agent sees all messages...
and everything all other (non-source) agents see!



Distinguishing the Opinion of the Source

Assume w.l.o.g. that at each round the agents pull a message one-by-one (**sequential** version).

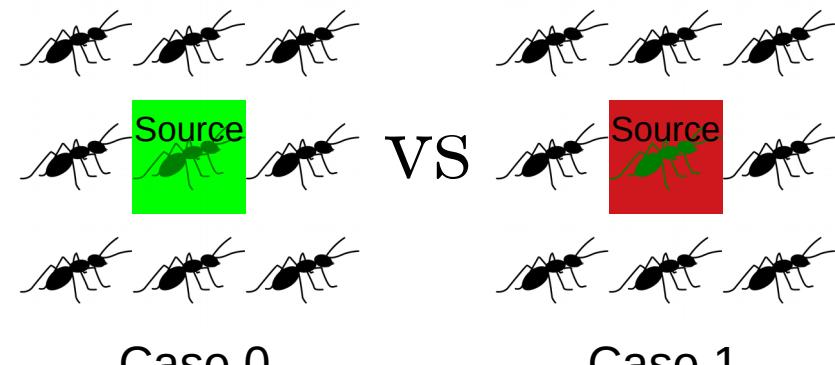
Recall: **every agent sees everything** that each non-source agents sees, even “random” choices (shared randomness)



The only difference between the two cases is the opinion of the source



Given **identical message histories**, the two systems behave identically



Lower Bound on Number of Messages via Relative Entropy: Preliminaries

The relative entropy between random variables X and Y is

$$D(X, Y) = \sum_x P(X = x) \log \frac{P(X = x)}{P(Y = x)}.$$

Important properties are the [Chain Rule](#)

$$\begin{aligned} D((X, X'), (Y, Y')) &= D(X', Y') \\ &\quad + \mathbb{E}_{x' \sim X'} [D((X|X' = x'), (Y|Y' = x'))] \end{aligned}$$

and the [Information Processing Inequality](#): for any function f

$$D(X, Y) \geq D(f(X), f(Y))$$

Lower Bound on the Number of Messages

Theorem

$X = (X_1, \dots, X_n)$ and $Y = (Y_1, \dots, Y_n)$ random variables,
 $f : \mathbb{R}^n \rightarrow \{0, 1\}$ a function that outputs 1 if it guesses that the
input is distributed like X , 0 if it is distributed like Y , and such
that $P(f(X) = 0), P(f(Y) = 1) \leq \delta$.

Denote $X_{<i} = (X_1, \dots, X_{i-1})$ and the same for $Y_{<i}$.

Then, it is necessary that

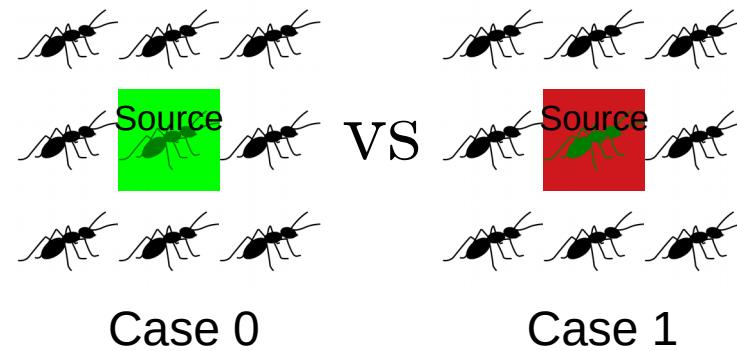
$$n \geq \frac{\log \frac{1}{2\delta}}{2 \max_i D((X_i|X_{<i}), (Y_i|Y_{<i}))}.$$

Proof 

Application of the Lower Bound 1/2

$$\begin{aligned} D((X_i|X_{<i}), (Y_i|Y_{<i})) \\ = \mathbb{E}_{Z_{<i} \sim X_{<i}} [D((X_i|X_{<i} = Z_{<i}), (Y_i|Y_{<i} = Z_{<i}))] \end{aligned}$$

Given the same history $Z_{<i}$, $p_{x,i}(z) := P(X_i = z|X_{<i} = z_{<i})$ and $p_{y,i}(z) := P(Y_i = z|Y_{<i} = z_{<i})$ are identical, except for the different sources.



Let $\epsilon_i(z) = p_{x,i}(z) - p_{y,i}(z)$.

$|\epsilon_i(z)| \leq \frac{1}{n}$ since $p_{x,i}(z)$ and $p_{y,i}(z)$ differ iff the message come from the source, which happens with probability $\frac{1}{n}$.

Note also that $\epsilon_i(1) = -\epsilon_i(0)$.

Application of the Lower Bound 2/2

$$\begin{aligned} D(p_{x,i}(z), p_{y,i}(z)) &= \sum_{z \in \{0,1\}} p_{x,i}(z) \log \frac{p_{x,i}(z)}{p_{y,i}(z)} \\ &= \sum_{z \in \{0,1\}} p_{x,i}(z) \log \frac{p_{x,i}(z)}{p_{x,i}(z) + \epsilon_i(z)} \\ &\stackrel{\text{(Taylor expansion of } \log \frac{1}{1-x})}{\leq} \sum_{z \in \{0,1\}} p_{x,i}(z) \left(-\frac{\epsilon_i(z)}{p_{x,i}(z)} + \left(\frac{\epsilon_i(z)}{p_{x,i}(z)} \right)^2 \right) \stackrel{\text{(constant noise: } p_{x,i}(z) \geq \frac{1}{3})}{\leq} O\left(\frac{1}{n^2}\right) \end{aligned}$$

that is, the required number of messages to attain the given probability of error is at least

$$\frac{\log \frac{1}{2\delta}}{2 \max_i D((X_i|X_{<i}), (Y_i|Y_{<i}))} \geq \frac{n^2}{2} \log \frac{1}{2\delta} \quad (\geq n \text{ rounds!})$$

Each round of PULL model corresponds to n observation of broadcast-PULL

Project Idea:

Propose algorithms to solve Rumor Spreading in the Noisy PULL Model on Erdős-Rényi graphs with parameter p and simulate them varying p and the noise parameter ϵ .

- Clementi, Andrea, Luciano Gualà, Emanuele Natale, Francesco Pasquale, Giacomo Scornavacca, and Luca Trevisan. 2018. “Consensus Needs Broadcast in Noiseless Models but Can Be Exponentially Easier in the Presence of Noise.” Report. CNRS. <https://hal.inria.fr/hal-01958994/document>
- Boczkowski, Lucas, Ofer Feinerman, Amos Korman, and Emanuele Natale. 2018. “Limits for Rumor Spreading in Stochastic Populations.” In 9th Innovations in Theoretical Computer Science Conference (ITCS 2018), edited by Anna R. Karlin, 94:49:1–49:21. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. <https://doi.org/10.4230/LIPIcs.ITCS.2018.49>

Consider only the binary-message case.

Simulations should be performed using open-source software with some effort to make them efficient (e.g. coded in Python using Numpy), and the source code should be made publicly available (e.g. on Gitlab) and GPL licensed.



Proof of the Lower Bound on the Number of Messages

$$\begin{aligned} D(f(X), f(Y)) &= \Pr(f(X) = 1) \log \frac{\Pr(f(X) = 1)}{\Pr(f(Y) = 1)} \\ &\quad + \Pr(f(X) = 0) \log \frac{\Pr(f(X) = 0)}{\Pr(f(Y) = 0)} \\ (\text{since } P(f(X) = 0), P(f(Y) = 1) &\leq \delta \text{ and } \lim_{x \rightarrow 0} x \log x = 0) \\ &\geq \frac{1}{2} \log \frac{\Pr(f(X) = 1)}{\Pr(f(Y) = 1)} \geq \frac{1}{2} \log \frac{1 - \delta}{\delta} \geq \frac{1}{2} \log \frac{1}{2\delta}. \end{aligned}$$

By the chain rule of the relative entropy, denoting $X_{<i} = (X_1, \dots, X_{i-1})$ and the same for $Y_{<i}$, we have

$$\begin{aligned} D(X, Y) &= \sum_i D((X_i | X_{<i}), (Y_i | Y_{<i})) \\ &\leq n \max_i D((X_i | X_{<i}), (Y_i | Y_{<i})). \end{aligned}$$

Combine the bounds with the information processing inequality.

Back to
Entropy
Bound

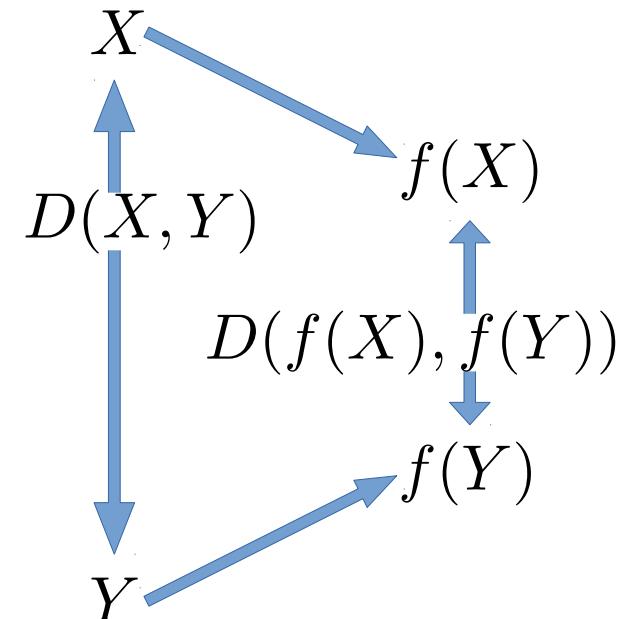
Information Processing Inequality for Relative Entropy

Let

- X and Y be two random variable on the same space Ω ,
- $f : \Omega \rightarrow \{0, 1\}$ be any binary function.

It holds

$$D(X, Y) \geq D(f(X), f(Y)).$$



Proof 

Back to
Relative
Entropy 

Sketch of Proof of the Information Processing Inequality

Notice that

$$\begin{aligned} P(X = z) &= P(X = z, f(X) = f(z)) \\ &= P(X = z | f(X) = f(z)) P(f(X) = f(z)) \end{aligned}$$

and the same for Y .

$$\begin{aligned} D(X, Y) &= \sum_{i \in \{0,1\}} \sum_{z: f(z)=i} P(f(X) = i) P(X = z | f(X) = i) \\ &\quad \cdot \log \frac{P(f(Y) = i) P(Y = z | f(Y) = i)}{P(f(X) = i) P(X = z | f(X) = i)} \\ &= \sum_{i \in \{0,1\}} P(f(X) = i) \log \frac{P(f(Y) = i)}{P(f(X) = i)} \\ &\quad + \sum_{i \in \{0,1\}} P(f(X) = i) D((X | f(X) = i), (Y | f(Y) = i)) \\ &\geq D(f(X), f(Y)) \end{aligned}$$

Back to
Relative
Entropy

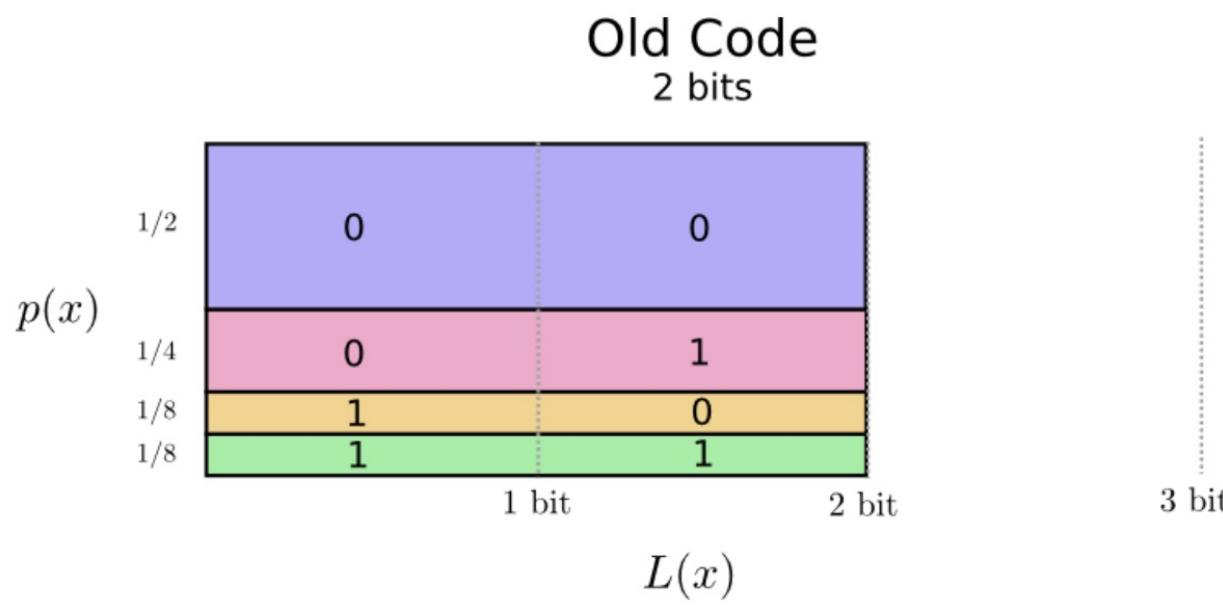
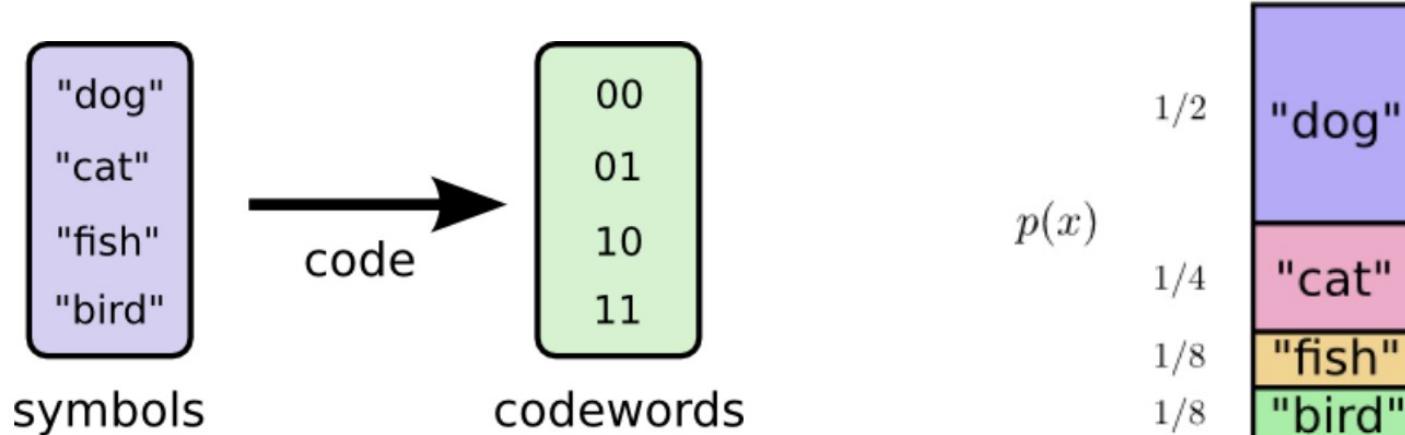
Chain Rule for Relative Entropy

$$\begin{aligned} D((X, X'), (Y, Y')) &= \sum_{x, x'} P(X = x, X' = x') \log \frac{P(X = x, X' = x')}{P(Y = y, Y' = y')} \\ &= \sum_{x, x'} P(X = x) P(X' = x' | X = x) \log \frac{P(X' = x' | X = x) P(X = x)}{P(Y' = x' | Y = x) P(Y = x)} \\ &= \sum_{x, x'} P(X = x) P(X' = x' | X = x) \log \frac{P(X = x)}{P(Y = x)} \\ &\quad + \sum_{x, x'} P(X = x) P(X' = x' | X = x) \log \frac{P(X' = x' | X = x)}{P(Y' = x' | Y = x)} \\ &= \sum_x P(X = x) \log \frac{P(X = x)}{P(Y = x)} \\ &\quad + \sum_x P(X = x) D((X' | X = x), (Y' | Y = x)) \end{aligned}$$

Back to
Relative
Entropy

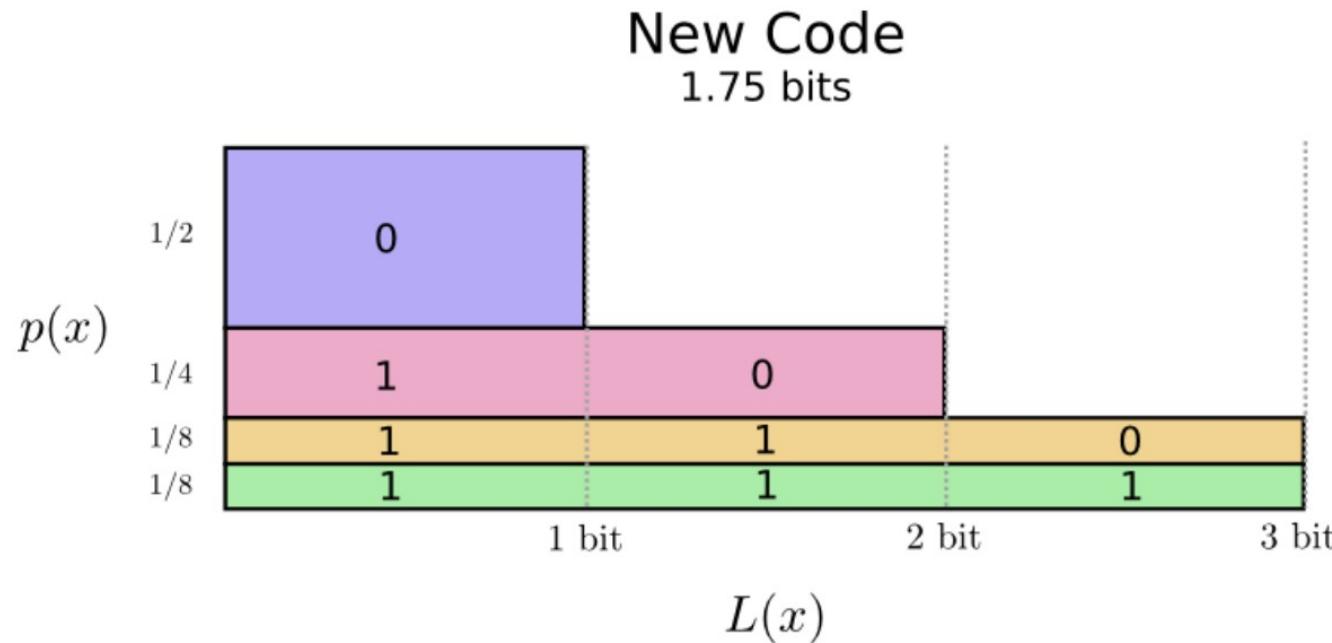
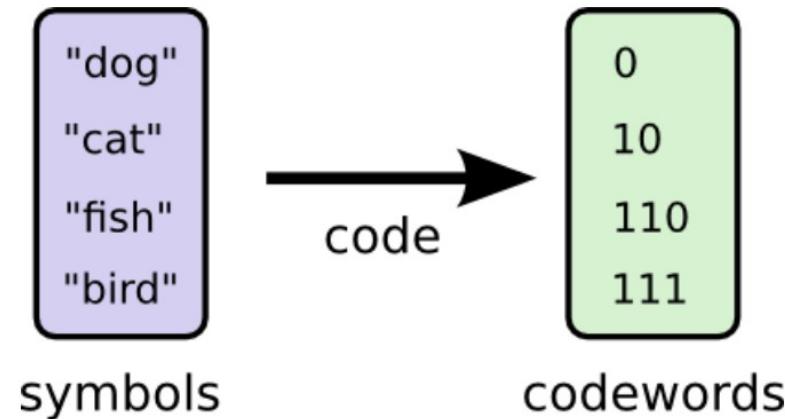
Visual Information Theory - Crash Course

Based on the [Colah's essay Visual Information Theory](#)
(all following images are from there)



Back to
Relative
Entropy

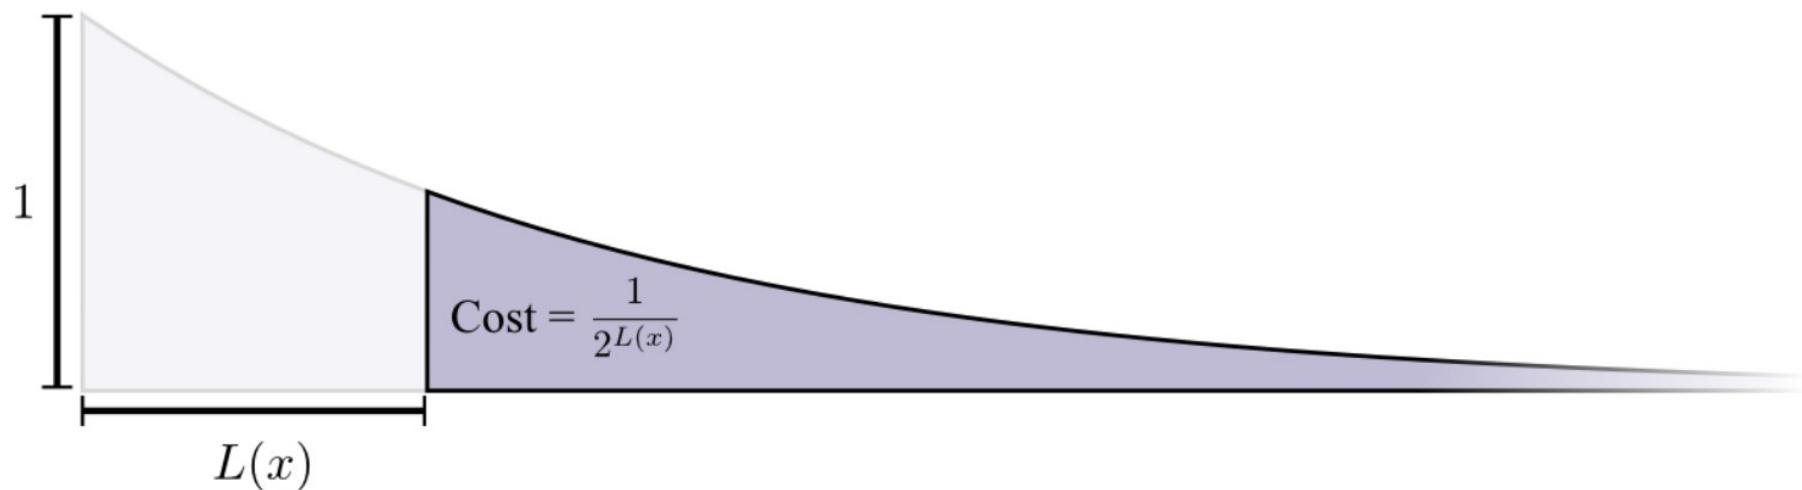
Minimizing Average Length



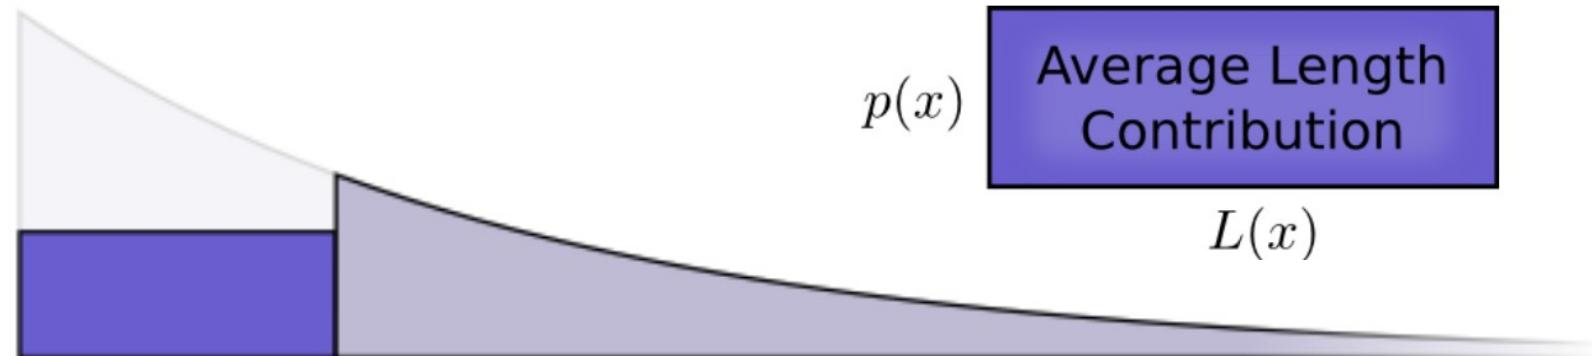
Prefix-free Codes

0	0	0	$\frac{1}{2^L} = \frac{1}{4}$
	1	1	
1	0	0	$\frac{1}{2^L} = \frac{1}{4}$
	1	0	

bit 1 bit 2 bit 3



Average Length vs Cost



Average Length
Contribution

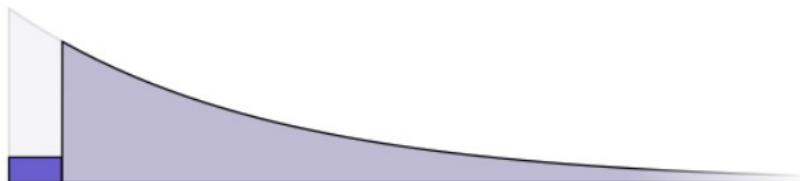
Codeword Cost

Average Length
Contribution

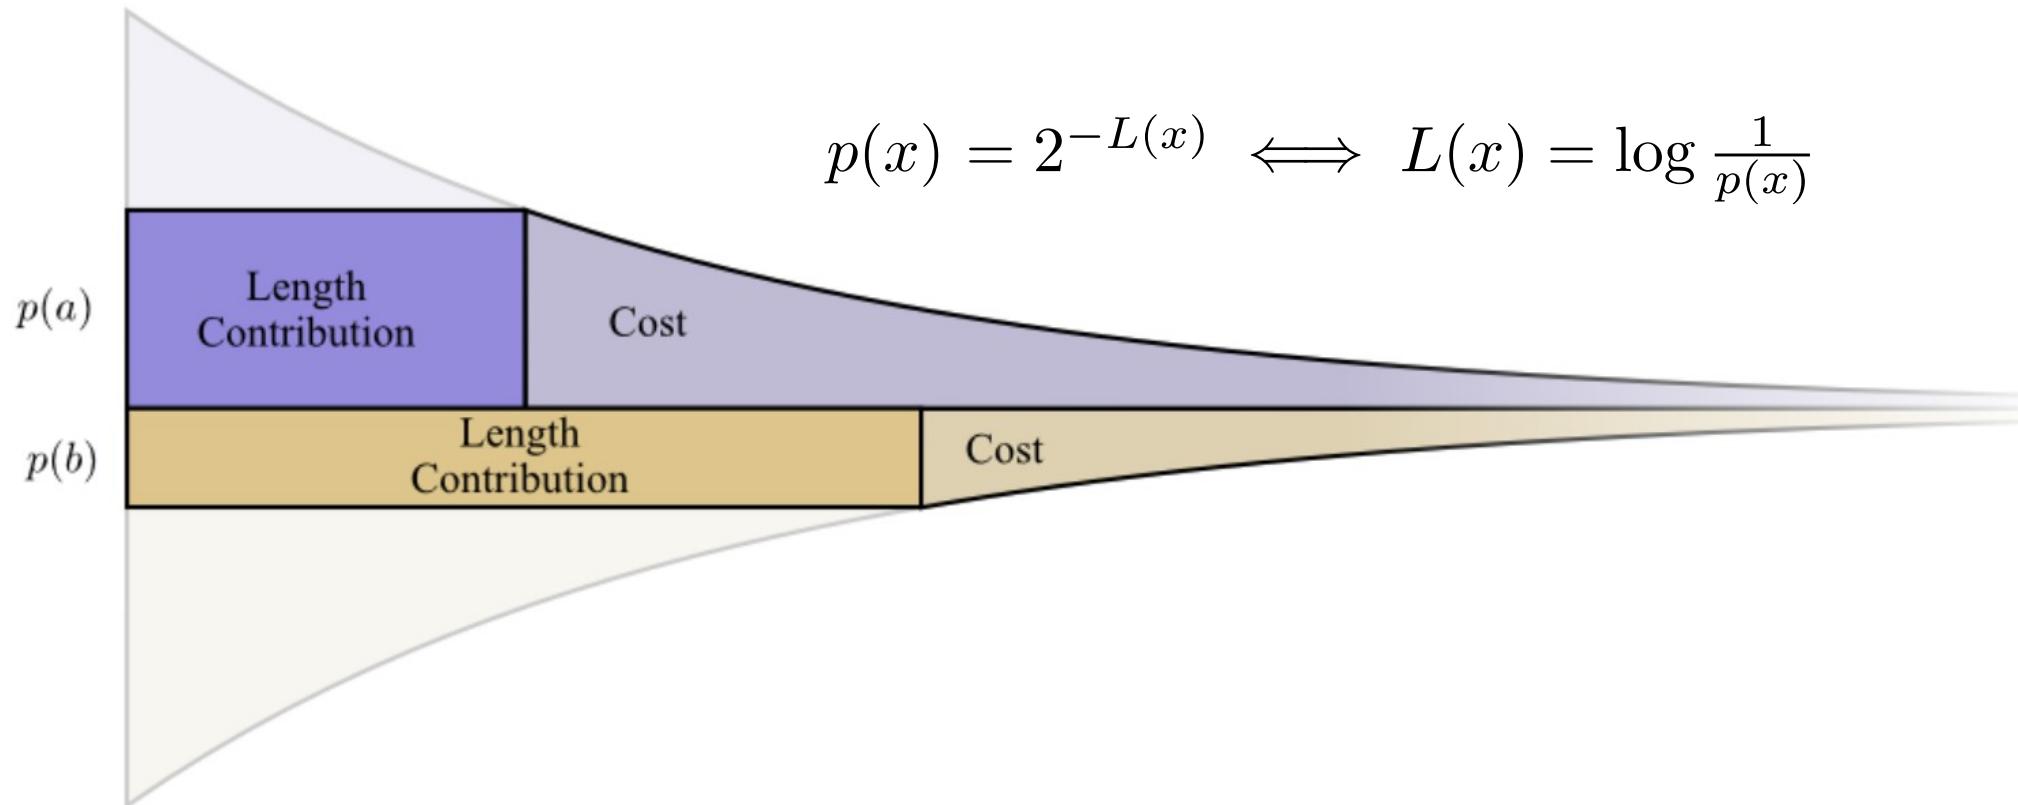
$L(x)$

Short Codeword,
High Cost

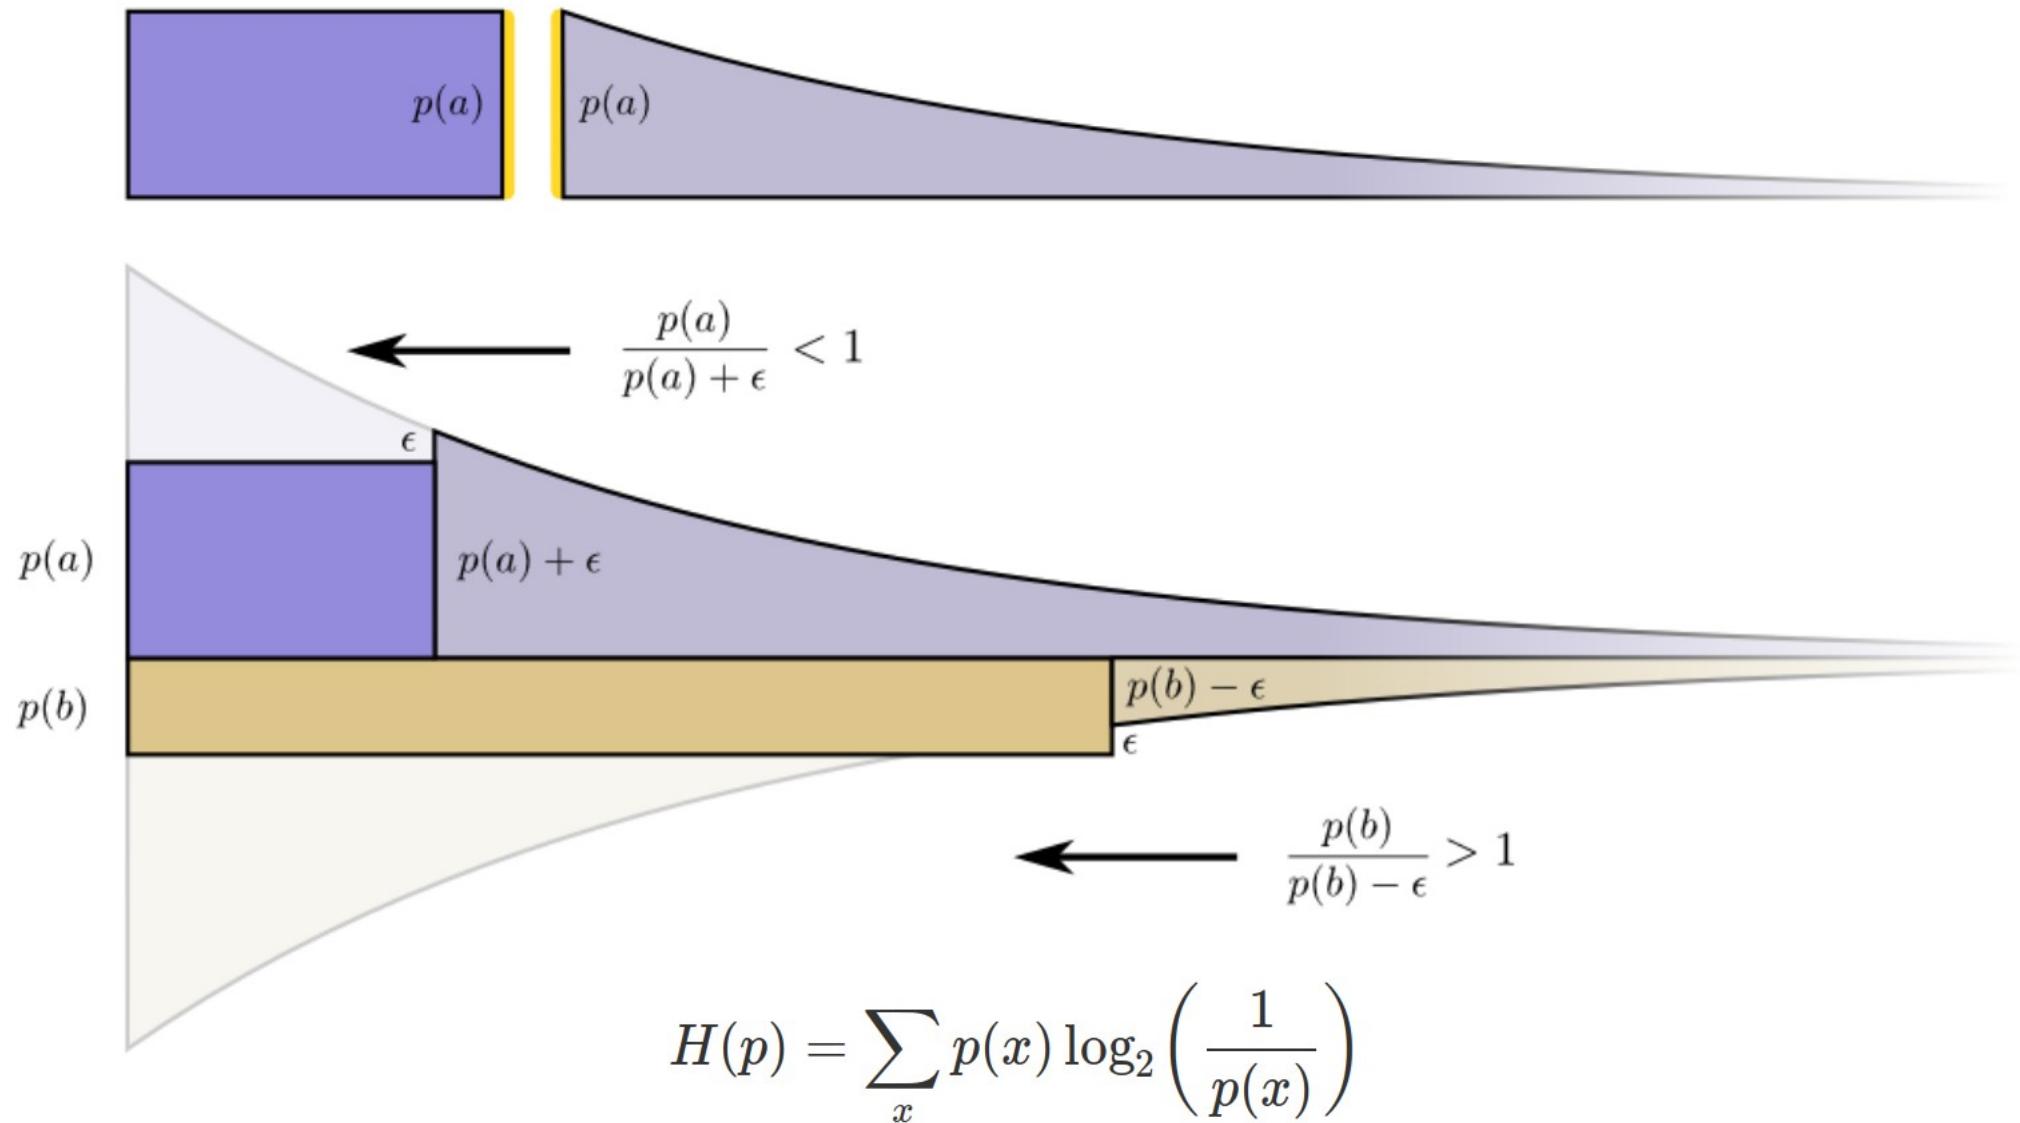
Long Codeword,
Small Cost



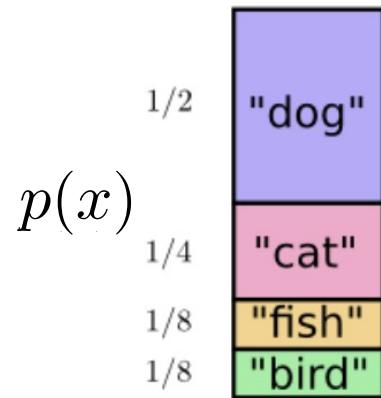
Information Content gives Optimal Length



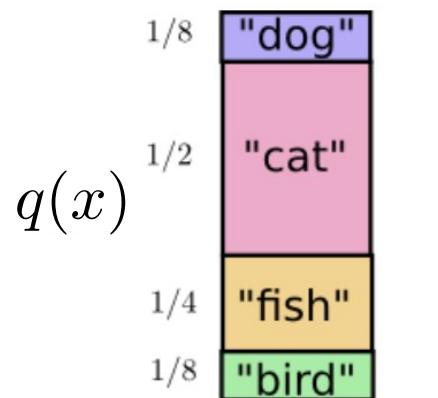
Proving Optimality



Cross and Relative Entropy



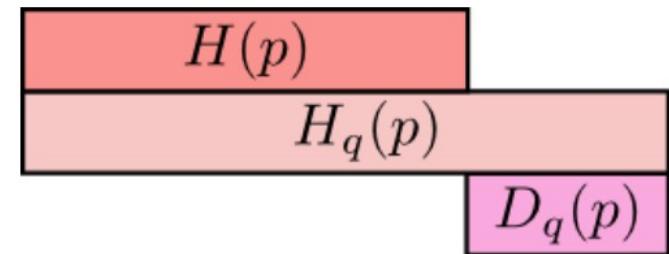
Dog Lover's
Word Frequency



Cat Lover's
Word Frequency

Cross Entropy: Average Length of message from $q(x)$ using code for $p(x)$

$$H_p(q) = \sum_x q(x) \log_2 \left(\frac{1}{p(x)} \right)$$



Relative Entropy: Average Extra Length of message from $q(x)$ using code for $p(x)$

$$\begin{aligned} D(q, p) &= \sum q(x) \log \frac{q(x)}{p(x)} \\ &= H_p(q) - H(q) \end{aligned}$$