

Objectives of Image Coding

- Representation of an image with acceptable quality, using as small a number of bits as possible

Applications:

- Reduction of channel bandwidth for image transmission
- Reduction of required storage

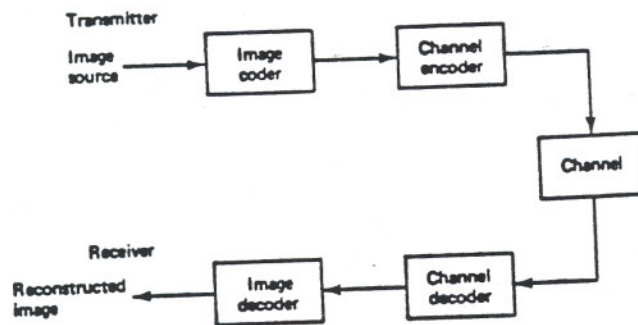


Figure 10.1 Typical environment for image coding.

Issues in Image Coding

1. What to code?
 - a. Image density
 - b. Image transform coefficients
 - c. Image model parameters

2. How to assign reconstruction levels
 - a. Uniform spacing between reconstruction levels
 - b. Non-uniform spacing between reconstruction levels

3. Bit assignment
 - a. Equal-length bit assignment to each reconstruction level
 - b. Unequal-length bit assignment to each reconstruction level




Figure 10.2 Three major components in image coding.

Methods of Reconstruction Level Assignments


Assumptions:

- Image intensity is to be coded
- Equal-length bit assignment


Scalar Case

1. Equal spacing of reconstruction levels (Uniform Quantization)

(Ex): Image intensity $f : 0 \sim 255$



Number of reconstruction levels: 4 (2 bits for equal bit assignment)







Figure 10.3 Example of uniform quantizer. The number of reconstruction levels is 4, f is assumed to be between 0 and 1, and \hat{f} is the result of quantizing f . The reconstruction levels and decision boundaries are denoted by r_i and d_i , respectively.

10.3

Scalar Case (cont.)

2. Spacing based on some error criterion

r_i : reconstruction levels (32, 96, 160, 224)

d_i : decision boundaries (0, 64, 128, 192, 256)

Optimally choose r_i and d_i .

To do this, assume $f_{\min} \leq f \leq f_{\max}$

$J \triangleq$ the number of reconstruction levels

$p(f)$: probability density function for f

Minimize

$$r_i, d_i : \varepsilon = E[(f - \hat{f})^2] = \int_{f=f_{\min}}^{f_{\max}} (f - \hat{f})^2 \cdot p(f) \cdot df$$

$$= \sum_{j=0}^{J-1} \int_{d_j}^{d_{j+1}} (f - r_j)^2 \cdot p(f) \cdot df$$

$$\implies \frac{\partial \varepsilon}{\partial r_j} = 0 \implies r_j = 2d_j - r_{j-1}$$

$$\frac{\partial \varepsilon}{\partial d_j} = 0 \quad r_j = \frac{\int_{d_j}^{d_{j+1}} f \cdot p(f) \cdot df}{\int_{d_j}^{d_{j+1}} p(f) \cdot df}$$

\implies Lloyd-Max Quantizer

These are not simple linear equations.

Scalar Case (cont.)

Solution to the Optimization Problem

TABLE 10.1 PLACEMENT OF RECONSTRUCTION AND DECISION LEVELS FOR LLOYD-MAX QUANTIZER. FOR UNIFORM PDF, $p_r(f_0)$ IS ASSUMED UNIFORM BETWEEN -1 AND 1. THE GAUSSIAN PDF IS ASSUMED TO HAVE MEAN OF 0 AND VARIANCE OF 1. FOR THE LAPLACIAN PDF,

$$p_r(f_0) = \frac{\sqrt{2}}{2\sigma} e^{-\frac{\sqrt{2}|f_0|}{\sigma}} \text{ with } \sigma = 1.$$

Bits	Uniform		Gaussian		Laplacian	
	r_i	d_i	r_i	d_i	r_i	d_i
1	-0.5000	-1.0000	-0.7979	$-\infty$	-0.7071	$-\infty$
	0.5000	0.0000	0.7979	0.0000	0.7071	0.0000
		1.0000		∞		∞
2	-0.7500	-1.0000	-1.5104	$-\infty$	-1.8340	$-\infty$
	-0.2500	-0.5000	-0.4528	-0.9816	-0.4198	-1.1269
	0.2500	0.0000	0.4528	0.0000	0.4198	0.0000
	0.7500	0.5000	1.5104	0.9816	1.8340	1.1269
	1.0000		∞		∞	
3	-0.8750	-1.0000	-2.1519	$-\infty$	-3.0867	$-\infty$
	-0.6250	-0.7500	-1.3439	-1.7479	-1.6725	-2.3796
	-0.3750	-0.5000	-0.7560	-1.0500	-0.8330	-1.2527
	-0.1250	-0.2500	-0.2451	-0.5005	-0.2334	-0.5332
	0.1250	0.0000	0.2451	0.0000	0.2334	0.0000
	0.3750	0.2500	0.7560	0.5005	0.8330	0.5332
	0.6250	0.5000	1.3439	1.0500	1.6725	1.2527
	0.8750	0.7500	2.1519	1.7479	3.0867	2.3769
	1.0000		∞		∞	
4	-0.9375	-1.0000	-2.7326	$-\infty$	-4.4311	$-\infty$
	-0.8125	-0.8750	-2.0690	-2.4008	-3.0169	-3.7240
	-0.6875	-0.7500	-1.6180	-1.8435	-2.1773	-2.5971
	-0.5625	-0.6250	-1.2562	-1.4371	-1.5778	-1.8776
	-0.4375	-0.5000	-0.9423	-1.0993	-1.1110	-1.3444
	-0.3125	-0.3750	-0.6568	-0.7995	-0.7287	-0.9198
	-0.1875	-0.2500	-0.3880	-0.5224	-0.4048	-0.5667
	-0.0625	-0.1250	-0.1284	-0.2582	-0.1240	-0.2664
	0.0625	0.0000	0.1284	0.0000	0.1240	0.0000
	0.1875	0.1250	0.3880	0.2582	0.4048	0.2644
	0.3125	0.2500	0.6568	0.5224	0.7287	0.5667
	0.4375	0.3750	0.9423	0.7995	1.1110	0.9198
	0.5625	0.5000	1.2562	1.0993	1.5778	1.3444
	0.6875	0.6250	1.6180	1.4371	2.1773	1.8776
	0.8125	0.7500	2.0690	1.8435	3.0169	2.5971
0.9375	0.8750	2.7326	2.4008	4.4311	3.7240	
	1.0000		∞		∞	




Figure 10.5 Example of a Lloyd-Max quantizer. The number of reconstruction levels is 4, and the probability density function for f is Gaussian with mean of 0 and variance of 1.

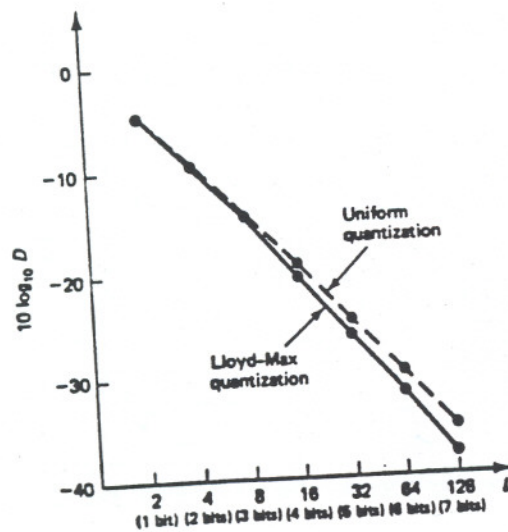


Figure 10.6 Comparison of average distortion $D = E[(\hat{f} - f)^2]$ as a function of L , the number of reconstruction levels, for a uniform quantizer (dotted line) and the Lloyd-Max quantizer (solid line). The vertical axis is $10 \log_{10} D$. The probability density function is assumed to be Gaussian with variance of 1.

Scalar Case (cont.)

For some densities, the optimal solution can be viewed as follows:




Figure 10.7 Nonuniform quantization by companding.

$p(g)$ is flat density equally likely between g_{\max} , g_{\min}

TABLE 8.1-2 Companding quantization transformation

	Probability Density	Forward Transformation	Inverse Transformation
Gaussian	$p(f) = (2\pi\sigma^2)^{-1/2} \exp\left\{-\frac{f^2}{2\sigma^2}\right\}$	$g = \frac{1}{\sqrt{2}} \operatorname{erf}\left\{\frac{f}{\sqrt{2}\sigma}\right\}$	$f = \sqrt{2}\sigma \operatorname{erf}^{-1}(2g)$
Rayleigh	$p(f) = \frac{f}{\sigma^2} \exp\left\{-\frac{f^2}{2\sigma^2}\right\}$	$g = \frac{1}{2} - \exp\left\{-\frac{f^2}{2\sigma^2}\right\}$	$f = [\sqrt{2}\sigma^2 \ln(1/(1/2 - g))]^{1/2}$
Laplacian	$p(f) = \frac{a}{2} \exp(-a f)$ $a = \frac{\sqrt{2}}{\sigma}$	$g = \frac{1}{2}[1 - \exp(-a f)] \quad f \geq 0$	$f = -\frac{1}{a} \ln[1 - 2g] \quad g \geq 0$
		$g = -\frac{1}{2}[1 - \exp(a f)] \quad f < 0$	$f = \frac{1}{a} \ln[1 + 2g] \quad g < 0$

where $\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z \exp(-y^2) dy$

Vector Case


f_1, f_2 : two image intensities

One approach: Separate level assignment for f_1, f_2

- previous discussions apply

Another approach: Joint level assignment for (f_1, f_2)

- typically more efficient than separate level assignment



Minimize reconstruction levels/decision boundaries:

$$\varepsilon = \int_{f_1} \int_{f_2} ((f_1 - \hat{f}_1)^2 + (f_2 - \hat{f}_2)^2) \cdot p(f_1, f_2) \cdot df_1 \cdot df_2$$

- efficiency depends on the amount of correlation between f_1 and f_2
- finding joint density is difficult

(Ex): Extreme case—vector level assignment for 256×256

Joint density $P(x_1, x_2, \dots, x_{256})$ is difficult to get

Law of diminishing returns comes into play

Codeword Design: Bit Allocation

- After quantization, need to assign *binary* codeword to each quantization symbol.
- Options:
 - uniform: equal number of bits to each symbol \rightarrow inefficient
 - non-uniform: short codewords to more probable symbols, and longer codewords for less probable ones.
- For non-uniform, code has to be uniquely decodable:
 - Example: $L = 4$, $r_1 \rightarrow 0$, $r_2 \rightarrow 1$, $r_3 \rightarrow 10$, $r_4 \rightarrow 11$
 - suppose we receive 100.
 - Can decode it two ways: either r_3r_1 or $r_2r_1r_1$.
 - Not uniquely decodable.
 - One codeword is a *prefix* of another one.
- Use Prefix codes instead to achieve unique decodability.


Overview

- Goal: Design variable length codewords c_i so that the average bit rate is minimized.
- Define entropy to be:

$$H \equiv \sum_{i=1}^L p_i \log_2 p_i$$

p_i is the probability that the i th symbol is a_i .

- Since $\sum_{i=1}^L p_i = 1$, we have $0 \leq H \leq \log_2 L$
- Entropy: average amount of information a message contains.
- Example: $L = 2$, $p_1 = 1$, $p_2 = 0 \rightarrow H = 0$ A symbol contains NO information.
- $p_1 = p_2 = 1/2 \rightarrow H = 1$ Maximum possible value.



Overview

- Information theory: H is the theoretically minimum possible average bite rate.
- In practice, hard to achieve
- Example: L symbols, $p_i = \frac{1}{L} \rightarrow H = \log_2 L$
uniform length coding can achieve this.
- Huffman coding tells you how to do non-uniform bit allocation to different codewords so that you get unique decodability and get pretty close to entropy.