Topics: Lossy image source coding

Spring 2008 Final: Problem 4 (rate-distortion))

Consider a discrete-time random process X_n with i.i.d. samples that are Gaussian with mean 0 and variance $\sigma^2 > 0$.

The rate distortion relation for this source is then given by

$$\begin{aligned} R(\Delta) &= \max\left\{\frac{1}{2}\log_2\left(\frac{\sigma^2}{\Delta^2}\right), 0\right\} \\ D(\Delta) &= \min\left\{\sigma^2, \Delta^2\right\} \end{aligned}$$

a) Plot the minimum possible rate (y-axis) versus distortion (x-axis) required to code this source when $\sigma^2 = 1$.

Solution:

$\sigma^2 = 1,\ D(\Delta) = \min\left\{1, \Delta^2\right\}$

$$\text{if } 0 < \Delta < 1, \text{ and } D(\Delta) = \Delta^2, \text{ and } \frac{\sigma^2}{\Delta^2} = \frac{1}{\Delta^2} > 1 \Rightarrow \frac{1}{2}\log_2\left(\frac{\sigma^2}{\Delta^2}\right) > \frac{1}{2}\log_2 1 = 0$$

$$\therefore R(\Delta) = \frac{1}{2}\log_2\frac{1}{\Delta^2} = \frac{1}{2}\log_2\frac{1}{D(\Delta)}$$

$$\text{if } \Delta \geq 1, \text{ and } D(\Delta) = \sigma^2 = 1, \text{ and } \frac{\sigma^2}{\Delta^2} \leq 1, \text{ and } \frac{1}{2}\log_2\left(\frac{\sigma^2}{\Delta^2}\right) < 0 \Rightarrow R(\Delta) = 0$$

Therefore, the plot is:

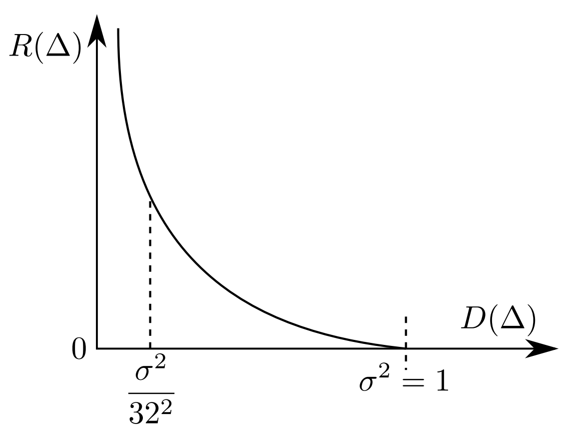

b) If we require that the distortion $D \leq \sigma^2$, then what is the minimum (lower bound) on the number of bits per sample that is required to transmit this signal?

Solution:

$D \leq \sigma^2$, the minimum number of bits per sample $R = 0$, since $\frac{1}{2}\log_2 \frac{\sigma^2}{D^2} \geq \frac{1}{2}\log_2 \frac{\sigma^2}{\sigma^2} = 0$

c) If we require that the distortion $D \leq \frac{\sigma^2}{(32)^2}$, then what is the minimum (lower bound) on the number of bits per sample that is required to transmit this signal?

Solution:

$$\text{When } D \leq \frac{\sigma^2}{32^2},\ R = \frac{1}{2}\log_2 \frac{\sigma^2}{D^2} \geq 5 \text{ bits}$$

d) How many bits per sample are required in order to achieve zero distortion?

Solution:

If $D = 0$ wanted, $R = \frac{1}{2}\log_2 \frac{\sigma^2}{D^2} \to \infty$

$\therefore$ An infinite number of bits is required for zero distortion, which is always impractical. Therefore, we generally consider lossy compression, i.e. distortion $\neq 0$.

e) Describe how you would design a lossy coder for this signal assuming that your objective is to achieve a bit rate of approximately 8 bits per sample.

Solution:

We want $R \approx 8$ bits/sample, then $D \approx \frac{\sigma^2}{256^2}$

1) We first clip the source samples in the range $[-T, T]$, where $T \gg \sigma$. For example, take 3σ; then about 99% of the samples will be in $[-3\sigma, 3\sigma]$.

2) Uniform quantization: since the distortion for a uniform quantization is $\frac{1}{12}\Delta^2$, where Δ is the quantization step.

$$D = \frac{1}{12}\Delta^2 = \frac{\sigma^2}{256^2} \ \Rightarrow \Delta = \sqrt{12}\frac{\sigma}{256}$$
$$\text{and } Q(X_n) = \text{round}\left(\frac{X_n}{\Delta}\right) = \text{round}\left(\frac{256 X_n}{\sqrt{12}\sigma}\right)$$

3) Entropy coding the quantized data, and get about 8 bits/sample.