Lecture 13 Rate Distortion Theory

Textbook 10.1-10.3

December 24 and 31, 2024

Outline

- In this section we motivate the elegant theory of rate distortion by showing how complicated it is to solve the quantization problem exactly for a single random variable.
- Since a continuous random source requires infinite precision to represent exactly, we cannot reproduce it exactly using a finite-rate code.
- The question is then to find the best possible representation for any given data rate.

- We first consider the problem of representing a single sample from the source.
- Let the random variable be represented be X and let the representation of X be denoted as $\hat{X}(X)$.
- If we are given R bits to represent X, the function $\hat{X}$ take on 2^R values.
- The problem is to find the optimum set of values for $\hat{X}$ (called the **reproduction points** or **code points**) and the regions that are associated with each value $\hat{X}$.

- For example, let $X \sim \mathcal{N}(0, \sigma^2)$, and assume a squared-error distortion measure.
- In this case we wish to find the function $\hat{X}(X)$ such that $\hat{X}$ takes on at most 2^R values and minimizes $\mathsf{E}(X - \hat{X}(X))^2$.
- If we are given one bit to represent X, it is clear that the bit should distinguish whether or not $X > 0$.
- To minimize squared error, each reproduced symbol should be the conditional mean of its region.

This is illustrated in the following figure.

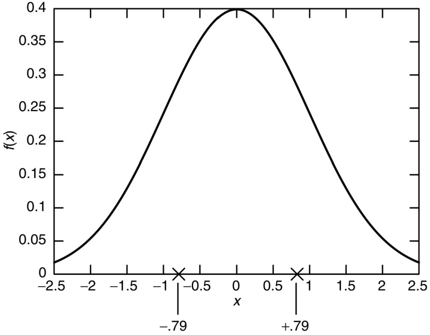

Thus,

$$\hat{X}(x) = \begin{cases} \sqrt{\frac{2}{\pi}}\sigma & \text{if } x \geq 0, \\ -\sqrt{\frac{2}{\pi}}\sigma & \text{if } x < 0. \end{cases}$$

- If we are given 2 bits to represent the sample, the situation is not as simple.
- Clearly, we want to divide the real line into four regions and use a point within each region to represent the sample.
- But it is no longer immediately obvious what the representation regions and the reconstruction points should be.
- We can, however, state two simple properties of optimal regions and reconstruction points for the quantization of a single random variable:

- Given a set $\{\hat{X}(w)\}$ of reconstruction points, the distortion is minimized by mapping a source random variable X to the representation $\hat{X}(w)$ that is closest to it. The set of regions of X defined by this mapping is called a **Voronoi** or **Dirichlet partition** defined by the reconstruction points.
- The reconstruction points should minimize the conditional expected distortion over their respective assignment regions.

- These two properties enable us to construct a simple algorithm to find a "good" quantizer.
- We start with a set of reconstruction points, find the optimal set of reconstruction regions (which are the nearest-neighbor regions with respect to the distortion measure), then find the optimal reconstruction points for these regions (the centroids of these regions if the distortion is squared error), and then repeat the iteration for this new set of reconstruction points.
- The expected distortion is decreased at each stage in the algorithm, so the algorithm will converge to a local minimum of the distortion.
- This algorithm is called the Lloyd algorithm (for real-valued random variables) or the generalized Lloyd algorithm (for vector-valued random variables) and is frequently used to design quantization systems.

- Instead of quantizing a single random variable, let us assume that we are given a set of n i.i.d. random variables drawn according to a Gaussian distribution.
- These random variables are to be represented using nR bits.
- Since the source is i.i.d., the symbols are independent, and it may appear that the representation of each element is an independent problem to be treated separately.
- But this is not true, as the results on rate distortion theory will show.
- We will represent the entire sequence by a single index taking 2^{nR} values.
- This treatment of entire sequences at once achieves a lower distortion for the same rate than independent quantization of the individual samples.

Outline

Definition

A **distortion function** or **ditortion measure** is a mapping

$$d : \mathcal{X} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}^{+}$$

from the set of source alphabet-reproduction alphabet pairs into the set of nonnegative real numbers. The distortion $d(x, \hat{x})$ is a measure of the cost of representing the symbol x by the symbol $\hat{x}$.

Definition

a distortion measure is said to be **bounded** if the maximum value of the distortion is finite:

$$d_{\max} := \max_{x \in \mathcal{X}, \hat{x} \in \hat{\mathcal{X}}} d(x, \hat{x}) < \infty.$$

Example (Hamming (probability of error) distortion)

The Hamming distortion is given by

$$d(x, \hat{x}) = \begin{cases} 0 & \textit{if } x = \hat{x} \\ 1 & \textit{if } x \neq \hat{x}. \end{cases}$$

which results in a probability of error distortion, since $Ed(X, \hat{X}) = P(X \neq \hat{X})$.

Example

The squared distortion,

$$d(x, \hat{x}) = (x - \hat{x})^2,$$

is the most popular distortion measure used for continuous alphabets. Its advantages are its simplicity and its relationship to least-squares prediction. But in applications such as image and speech coding, various authors have pointed out that the mean-squared error is not an appropriate measure of distortion for human observers. For example, there is a large squared-error distortion between a speech waveform and another version of the same waveform slightly shifted in time, even though both would sound the same to a human observer.

Definition

The distortion between sequences x^n and $\hat{x}^n$ is defined by

$$d(x^n, \hat{x}^n) = \frac{1}{n} \sum_{i=1}^{n} d(x_i, \hat{x}_i).$$

Definition

The **rate distortion region** for a source is the closure of the set of achievable rate distortion pairs (R, D).

Definition

The **rate distortion function** $R(D)$ is the infimum of rates R such that (R, D) is in the rate distortion region of the source for a given distortion D.

Definition

The **distortion rate function** $D(R)$ is the infimum of all distortions D such that (R, D) is in the rate distortion region of the source for a given rate R.

The **information rate distortion function** $R^{(I)}(D)$ for a source X with distortion measure $d(x, \hat{x})$ is defined as

$$R^{(I)}(D) = \min_{p(\hat{x}|x): \sum_{(x,\hat{x})} p(x)p(\hat{x}|x)d(x,\hat{x}) \le D} I(X : \hat{X}),$$

where the minimization is over all conditional distributions $p(\hat{x}|x)$ for which the joint distribution $p(x, \hat{x}) = p(x)p(\hat{x}|x)$ satisfies the expected distortion constraint.

Theorem

The rate distortion function for a Bernoulli(p) *source with Hamming distortion is given by*

$$R(D) = \begin{cases} H(p) - H(D), & 0 \le D \le \min\{p, 1-p\}, \\ 0, & D > \min\{p, 1-p\}. \end{cases}$$

Proof.

- Consider a binary source $X \sim \text{Bernoulli}(p)$ with a Hamming distortion measure.
- Without loss of generality, we may assume that $p < \frac{1}{2}$.
- We wish to calculate the rate distortion function,

$$R(D) = \min_{p(\hat{x}|x):\sum_{(x,\hat{x})} p(x)p(\hat{x}|x)d(x,\hat{x}) \leq D} I(X;\hat{X}).$$

Proof.

- Let $\oplus$ denote modulo 2 addition.
- Thus, $X \oplus \hat{X} = 1$ is equivalent to $X \neq \hat{X}$.
- We do not minimize $I(X; \hat{X})$ directly; instead, we find a lower bound and then show that this lower bound is achievable.
- For any joint distribution satisfying the distortion constraint, we have

$$\begin{aligned} I(X; \hat{X}) &= H(X) - H(X|\hat{X}) \\ &= H(p) - H(X \oplus \hat{X}|\hat{X}) \\ &\geq H(p) - H(X \oplus \hat{X}) \\ &\geq H(p) = H(D), \end{aligned}$$

since $\Pr(X \neq \hat{X}) \leq D$ and $H(D)$ increases with D for $D \leq \frac{1}{2}$.

Proof.

- We now show that the lower bound is actually the rate distortion function by finding a joint distribution that meets the distortion constraint and has $I(X;\hat{X}) = R(D)$.
- For $0 \le D \le p$, we can achieve the value of the rate distortion function above by choosing $(X, \hat{X})$ to have the joint distribution given by the binary symmetric channel shown in the next figure.

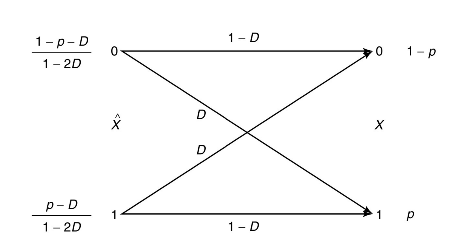

- We choose the distribution of $\hat{X}$ at the input of the channel so that the output distribution of X is the specified distribution. Let $r = \Pr(\hat{X} = 1)$.
- Then choose r so that

$$r(1-D) + (1-r)D = p,$$

or

$$r = \frac{p-D}{1-2D}.$$

- If $D \le p \le \frac{1}{2}$, then $\Pr(\hat{X} = 1) \ge 0$ and $\Pr(\hat{X} = 0) \ge 0$. We then have

$$I(X;\hat{X}) = H(X) - H(X|\hat{X}) = H(p) - H(D),$$

and the expected distortion is $\Pr(X \ne \hat{X}) = D$.

- If $D \geq p$, we can achieve $R(D) = 0$ by letting $\hat{X} = 0$ with probability 1.
- In this case, $I(X; \hat{X}) = 0$ and $D = p$.
- Similarly, if $D \geq 1 - p$, we can achieve $R(D) = 0$ by setting $\hat{X} = 1$ with probability 1.
- Hence, the rate distortion function for a binary source is

$$R(D) = \begin{cases} H(p) - H(D), & 0 \leq D \leq \min\{p - 1, -p\}, \\ 0, & D > \min\{p, 1 - p\}. \end{cases}$$

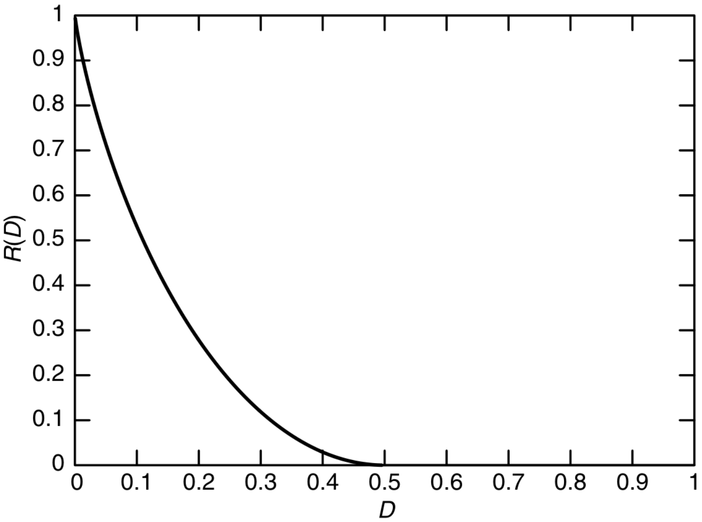
R(D)
D
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Theorem

The rate distortion for a $\mathcal{N}(0,\sigma^2)$ *source with squared-error distortion is*

$$R(D) = \begin{cases} \frac{1}{2}\log\frac{\sigma^2}{D}, & 0 \le D \le \sigma^2, \\ 0, & D > \sigma^2. \end{cases}$$

Proof.

- Let $X \sim \mathcal{N}(0, \sigma^2)$.
- By the rate distortion theorem extended to continuous alphabets, we have

$$R(D) = \min_{f(\hat{x}|x): E(\hat{X}-X)^2 \leq D} I(X; \hat{X}).$$

- As in the preceding example, we first find a lower bound for the rate distortion function and then prove that this is achievable.

Since $E(X-\hat{X})^2 \leq D$, we observe that

$$\begin{aligned}
I(X;\hat{X}) &= h(X) - h(X|\hat{X}) \\
&= \frac{1}{2}\log(2\pi e)\sigma^2 - h(X-\hat{X}|\hat{X}) \\
&\geq \frac{1}{2}\log(2\pi e)\sigma^2 - h(X-\hat{X}) \\
&\geq \frac{1}{2}\log(2\pi e)\sigma^2 - h(\mathcal{N}(0, E(X-\hat{X})^2)) \\
&= \frac{1}{2}\log(2\pi e)\sigma^2 - \frac{1}{2}\log(2\pi e)E(X-\hat{X})^2 \\
&\geq \frac{1}{2}\log(2\pi e)\sigma^2 - \frac{1}{2}\log(2\pi e)D \\
&= \frac{1}{2}\log\frac{\sigma^2}{D}.
\end{aligned}$$

- Hence, $R(D) \geq \frac{1}{2} \log \frac{\sigma^2}{D}$.
- To find the conditional density $f(\hat{x}|x)$ that achieves this lower bound, it is usually more convenient to look at the conditional density $f(x|\hat{x})$, which is sometimes called the **test channel** (thus emphsizing the duality of rate distortion with channel capacity).
- As in binary case, we construct $f(x|\hat{x})$ to achieve equality in the bound.
- We choose the joint distribution as shown in the next figure.

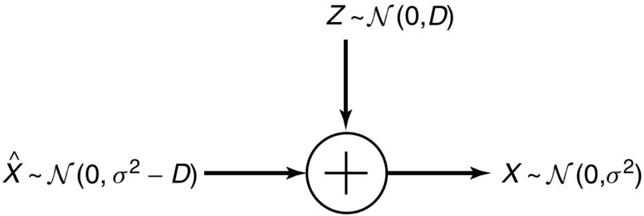
$Z \sim \mathcal{N}(0, D)$
$\hat{X} \sim \mathcal{N}(0, \sigma^2 - D)$
$X \sim \mathcal{N}(0, \sigma^2)$

- If $D \leq \sigma^2$, we choose

$$X = \hat{X} + Z, \;\; \hat{X} \sim \mathcal{N}(0, \sigma^2 - D), \;\; Z \sim \mathcal{N}(0, D),$$

  where $\hat{X}$ and Z are independent.

- For this joint distribution, we calculate

$$I(X; \hat{X}) = \frac{1}{2} \log \frac{\sigma^2}{D},$$

  and $\mathsf{E}(X - \hat{X})^2 = D$, thus achieving the bound.

- If $D > \sigma^2$, we choose $\hat{X} = 0$ with probability 1, achieving $R(D) = 0$.

Hence, the rate distortion function for the Gaussian source with squared-error distortion is

$$R(D) = \begin{cases} \frac{1}{2}\log\frac{\sigma^2}{D}, & 0 \le D \le \sigma^2, \\ 0, & D > \sigma^2, \end{cases}$$

as illustrated in the next figure. We can rewrite the above equation to express the distortion in terms of the rate,

$$D(R) = \sigma^2 2^{-2R}.$$

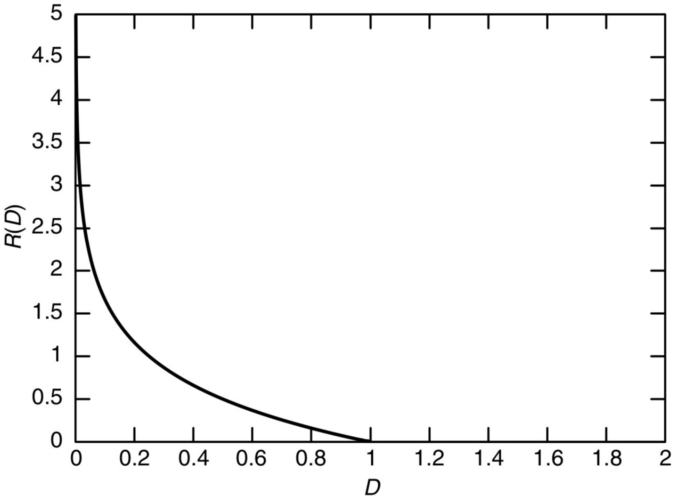
5
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
R(D)
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
D

Consider the case of representing m independent (but not identically distributed) normal random sources $X_1, \cdots, X_m$, where $X_i \sim \backslash(0, \sigma_i^2)$, with square-error distortion. Assume that we are given R bits with which to represent this random vector. The question naturally arises as to how we should allot these bits to the various components to minimize the total distortion. Extending the definition of the information rate distortion function to the vector case, we have

$$R(D) = \min_{f(\hat{x}|x^m):\mathsf{E}d(x^m,\hat{X}^n)\leq D} I(X^m; hatX^m),$$

where $d(x^m, \hat{x}^m) = \sum_{i=1}^m (x_i - \hat{x}_i)^2$.

Now using the argument in the preceding example, we have

$$\begin{aligned}
I(X^n;\hat{X}^m) &= h(X^m) - h(X^m|\hat{X}^m) \\
&= \sum_{i=1}^{m} h(X_i) - \sum_{i=1}^{m} h(X_i|x^{i-1},\hat{X}^m) \\
&\geq \sum_{i=1}^{m} h(X_i) - \sum_{i=1}^{m} h(X_i|\hat{X}_i) \\
&= \sum_{i=1}^{m} I(X_i;\hat{X}_i) \\
&\geq \sum_{i=1}^{m} R(D_i) \\
&= \sum_{i=1}^{m} \left(\frac{1}{2}\log\frac{\sigma_i^2}{D_i}\right)^+,
\end{aligned}$$

where $D_i = \mathsf{E}(X_i - \hat{X}_i)^2$.

We can achieve the equalities by choosing $f(x^m|\hat{x}^m) = \prod_{i=1}^m f(x_i|\hat{x}_i)$ and by choosing the distribution of each $\hat{X}_i \sim \mathcal{N}(0, \sigma_i^2 - D_i)$, as in the preceding example. Hence, the problem of finding the rate distortion function can be reduced to the following optimization (using nats for convenience):

$$R(D) = \min_{\sum D_i = D} \sum_{i=1}^{m} \max\{\frac{1}{2} \ln \frac{\sigma_i^2}{D_i}, 0\}.$$

Using lagrange multipliers, we construct the functional

$$J(D) = \sum_{i=1}^{m} \frac{1}{2} \ln \frac{\sigma_i^2}{D_i} + \lambda \sum_{i=1}^{m} D_i,$$

and differentiating with respect to D_i and setting equal to 0, we have

$$\frac{\partial J}{\partial D_i} = -\frac{1}{2}\frac{1}{D_i} + \lambda = 0$$

or

$$D_i = \lambda'.$$

Hence, the optimum allotment of the bits to the various descriptions results in an equal distortion for each random variable. This is possible if the constant λ' is less than σ_i^2 for all i. As the total allowable distortion D is increased, the constant λ' increases until it exceeds σ_i^2 allowable region of distortions. If we increase the total distortion, we must use the Kuhn-Tucker conditions to find minimum. In this case the Kuhn-Tucker condition yield

$$\frac{\partial J}{\partial D_i} = -\frac{1}{2}\frac{1}{D_i} + \lambda,$$

where λ is chosen so that

$$\frac{\partial J}{\partial D_i} \begin{cases} = 0 & \text{if } D_i < \sigma_i^2 \\ \leq 0 & \text{if } D_i \geq \sigma_i^2. \end{cases}$$

Theorem

Rate distortion for a parallel Gaussian source Let $X_i \sim \mathcal{N}(0, \sigma_i^2)$, $i = 1, 2, \cdots, m$, *be independent Gaussian random variables, and let the distortion measure be* $d(x^m, \hat{x}^m) = \sum_{i=1}^{m} (x_i - \hat{x}_i)^2$. *Then the rate distortion function is given by*

$$R(D) = \sum_{i=1}^{m} \frac{1}{2} \log \frac{\sigma_i^2}{D_i},$$

where

$$D_i = \begin{cases} \lambda & \text{if } \lambda < \sigma_i^2, \\ \sigma_i^2 & \text{if } \lambda \geq \sigma_i^2, \end{cases}$$

where λ *is chosen so that* $\sum_{i=1}^{m} D_i = D$.