Rate Distortion Theory & Quantization

- Rate Distortion Theory
- Rate Distortion Function
- $R(D^*)$ for Memoryless Gaussian Sources
- $R(D^*)$ for Gaussian Sources with Memory
- Scalar Quantization
- Lloyd-Max Quantizer
- High Resolution Approximations
- Entropy-Constrained Quantization
- Vector Quantization

Rate Distortion Theory

- Theoretical discipline treating data compression from the viewpoint of information theory.
- Results of rate distortion theory are obtained without consideration of a specific coding method.
- **Goal:** Rate distortion theory calculates minimum transmission bit-rate R for a given distortion D and source.

Transmission System

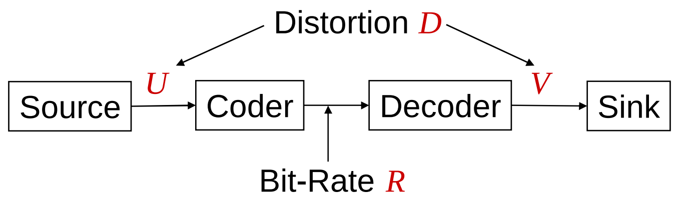

- Need to define U, V, Coder/Decoder, Distortion D, and Rate R.
- Need to establish functional relationship between $U, V, D,$ and R

Definitions

- **Source symbols** are given by the random sequence $\{U_k\}$
  - Each U_k assumes values in the discrete set $\mathcal{U} = \{u_0, u_1, \ldots, u_{M-1}\}$
    - For a binary source: $\mathcal{U} = \{0, 1\}$
    - For a picture: $\mathcal{U} = \{0, 1, \ldots, 255\}$
  - For simplicity, let us assume U_k to be independent and identically distributed (i.i.d.) with distribution $\{P(u),\ u \in \mathcal{U}\}$

- **Reconstruction symbols** are given by the random sequence $\{V_k\}$ with distribution $\{P(v),\ v \in \mathcal{V}\}$
  - Each V_k assumes values in the discrete set $\mathcal{V} = \{v_0, v_1, \ldots, v_{N-1}\}$
  - The sets $\mathcal{U}$ and $\mathcal{V}$ need not to be the same

Coder / Decoder

- Statistical description of **Coder/Decoder**, i.e. the mapping of the source symbols to the reconstruction symbols, via

$$Q = \left\{ Q(v \mid u),\, u \in \mathcal{U},\, v \in \mathcal{V} \right\}$$

- Q is the conditional probability distribution over the letters of the reconstruction alphabet $\mathcal{U}$ given a letter of the source alphabet $\mathcal{V}$
- Transmission system is described via

$$\text{Joint pdf}: \quad P(u,v)$$

$$P(u) = \sum_{v \in \mathcal{V}} P(u,v)$$

$$P(v) = \sum_{u \in \mathcal{U}} P(u,v)$$

$$P(u,v) = P(u) \cdot Q(v \mid u) \;\; \text{(Bayes' rule)}$$

Distortion

- To determine distortion, we define a non-negative cost function

$$d(u,v), \qquad d(\cdot,\cdot) : \mathcal{U} \times \mathcal{V} \to [0, \infty)$$

- Examples for d
  - Hamming distance: $$d(u,v) = \begin{cases} 0, & \text{for } u = v \\ 1, & \text{for } u \neq v \end{cases}$$
  - Squared error: $$d(u,v) = |u - v|^2$$
- Average **Distortion**

$$D(Q) = \sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} \underbrace{P(u) \cdot Q(v \mid u)}_{P(u,v)} \cdot d(u,v)$$

Mutual Information

- Shannon average mutual information

$$
\begin{aligned}
I &= H(U) - H(U|V) \\
&= -\sum_{u\in\mathcal{U}} P(u)\cdot \mathrm{ld} P(u) + \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}} P(u,v)\cdot \mathrm{ld} P(u|v) \\
&= -\sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}} P(u,v)\cdot \mathrm{ld} P(u) + \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}} P(u,v)\cdot \mathrm{ld}\frac{P(u,v)}{P(v)} \\
&= \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}} P(u,v)\cdot \mathrm{ld}\frac{P(u,v)}{P(u)\cdot P(v)}
\end{aligned}
$$

- Using Bayes' rule

$$
I(Q) = \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}} \underbrace{P(u)\cdot Q(v|u)}_{P(u,v)}\cdot \mathrm{ld}\frac{Q(v|u)}{P(v)},
$$

$$
\text{with}\quad P(v) = \sum_{u\in\mathcal{U}} P(u)\cdot Q(v|u)
$$

Rate

- Shannon average mutual information expressed via entropy

$$
I(U;V) = \underset{\substack{\uparrow \\ \text{Source entropy}}}{H(U)} - \underset{\substack{\uparrow \\ \text{Equivocation: conditional entropy}}}{H(U|V)}
$$

- Equivocation:
  - The conditional entropy (uncertainty) about the source U given the reconstruction V
  - A measure for the amount of missing [quantized] information in the received signal V

Rate Distortion Function

- Definition: $R(D^*) = \min_{Q:D(Q) \leq D^*} \{I(Q)\}$

- For a given maximum average distortion D, the rate distortion function $R(D^*)$ is the lower bound for the transmission bit-rate.
- The minimization is conducted for all possible mappings Q that satisfy the average distortion constraint
- $R(D^*)$ is measured in bits for ld

Discussion

- In information theory: maximize mutual information for efficient communication
- In rate distortion theory: minimize mutual information
- In rate distortion theory: source is given, not the channel
- Problem which is addressed:

  Determine the minimum rate at which information about the source must be conveyed to the user in order to achieve a prescribed fidelity.
- Another view: Given a prescribed distortion, what is the channel with the minimum capacity to convey the information.
- Alternative definition via interchanging the roles of rate and distortion

Distortion Rate Function

- Definition: $D(R^*) = \min_{Q:I(Q) \leq R^*} \{d(Q)\}$

- For a given maximum average rate R, the distortion rate function $D(R^*)$ is the lower bound for the average distortion.

- Here, we can set R^* to the capacity C of the transmission channel and determine the minimum distortion for this ideal communication system

Properties of the Rate Distortion Function, I

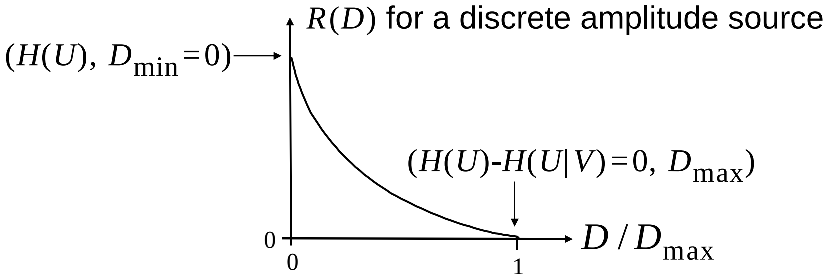

- $R(D)$ is well defined for $D \in (D_{min}, D_{max})$
- For discrete amplitude sources, $D_{min} = 0$
- $R(D) = 0$, if $D > D_{max}$

Properties of the Rate Distortion Function, II

- $R(D)$ is always positive

$$0 \leq I(U;V) \leq H(U)$$

- $R(D)$ is non-increasing in D
- $R(D)$ is strictly convex downward in the range (D_{min}, D_{max})
- The slope of $R(D)$ is continous in the range (D_{min}, D_{max})

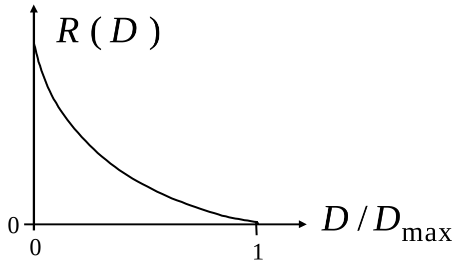

Shannon Lower Bound

- It can be shown that $\boxed{H(U-V|V) = H(U|V)}$
- Then we can write

$$\boxed{\begin{aligned} R(D^*) &= \min_{Q:D(Q)\leq D^*} \{H(U) - H(U|V)\} \\ &= H(U) - \max_{Q:D(Q)\leq D^*} \{H(U|V)\} \\ &= H(U) - \max_{Q:D(Q)\leq D^*} \{H(U-V|V)\} \end{aligned}}$$

- Ideally, the source coder would produce distortions u-v that are statistically independent from the reconstructed signal v (not always possible!).

- Shannon Lower Bound: $\boxed{R(D^*) \geq H(U) - \max_{Q:D(Q)\leq D^*} H(U-V)}$

$R(D^*)$ for a Memoryless Gaussian Source and MSE Distortion

- Gaussian source, variance σ^2
- Mean squared error (MSE) $D = E\{(u\text{-}v)^2\}$

$$R(D^*) = \frac{1}{2}\log\frac{\sigma^2}{D^*}; \qquad D(R^*) = \sigma^2 2^{-2R^*}, R \geq 0$$

$$SNR = 10 \cdot \log_{10}\frac{\sigma^2}{D} = 10 \cdot \log_{10} 2^{2R} \approx 6R \ [dB]$$

- Rule of thumb: 6 dB ~ 1 bit
- The $R(D^*)$ for non-Gaussian sources with the same variance σ^2 is always below this Gaussian $R(D^*)$ curve.

$R(D^*)$ Function for Gaussian Source with Memory I

- Jointly Gaussian source with power spectrum $S_{uu}(\omega)$
- MSE: $D = E\{(u\text{-}v)^2\}$
- Parametric formulation of the $R(D^*)$ function

$$D = \frac{1}{2\pi}\int_{\omega} \min[D^*, S_{uu}(\omega)]d\omega$$

$$R = \frac{1}{2\pi}\int_{\omega} \max[0, \frac{1}{2}\log\frac{S_{uu}(\omega)}{D^*}]d\omega$$

- $R(D^*)$ for non-Gaussian sources with the same power spectral density is always lower.

$R(D^*)$ Function for Gaussian Source with Memory II

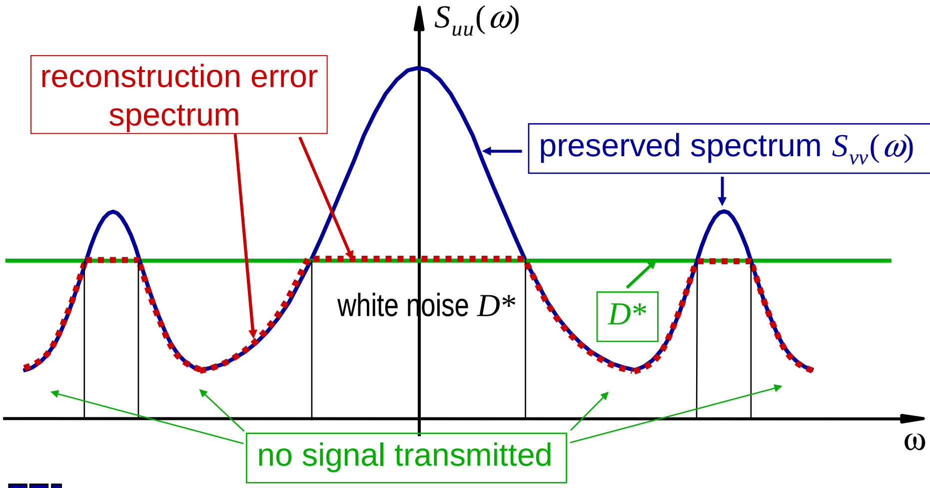

$R(D^*)$ Function for Gaussian Source with Memory III

- ACF and PSD for a first order AR(1) Gauss-Markov process: $U[n] = Z[n] + \rho\, U[n\text{-}1]$

$$R_{uu}(k) = \rho^{|k|}\sigma^2, \; S_{uu}(\omega) = \frac{\sigma^2(1-\rho^2)}{1-2\rho\cos\omega+\rho^2}$$

- Rate Distortion Function:

$$R(D^*) = \frac{1}{4\pi}\int_{-\pi}^{\pi}\log_2\frac{S_{uu}(\omega)}{D^*}d\omega, \; \frac{D^*}{\sigma^2} \le \frac{1-\rho}{1+\rho}$$

$$= \frac{1}{4\pi}\int_{-\pi}^{\pi}\log_2\frac{\sigma^2(1-\rho^2)}{D^*}d\omega - \frac{1}{4\pi}\int_{-\pi}^{\pi}\log_2\left(1-2\rho\cos\omega+\rho^2\right)d\omega$$

$$= \frac{1}{2}\log_2\frac{\sigma^2\left(1-\rho^2\right)}{D^*} = \frac{1}{2}\log_2\frac{\sigma_z^{\;2}}{D^*}.$$

$R(D^*)$ Function for Gaussian Source with Memory IV

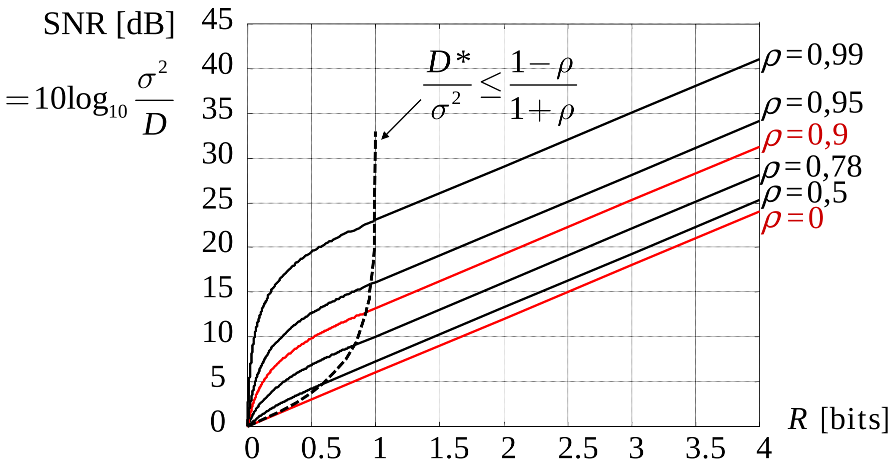

Quantization

- Structure

$$u \longrightarrow \boxed{\text{Quantizer}} \longrightarrow v$$

- Alternative: coder (α) / decoder (β) structure

$$u \longrightarrow \boxed{\alpha} \xrightarrow{i} \boxed{\beta} \xrightarrow{v}$$

- Insert entropy coding (γ) and transmission channel

$$u \longrightarrow \boxed{\alpha} \xrightarrow{i} \boxed{\gamma} \xrightarrow{b} \boxed{\textit{channel}} \xrightarrow{b} \boxed{\gamma^{-1}} \xrightarrow{i} \boxed{\beta} \xrightarrow{v}$$

Scalar Quantization

- Average distortion

$$D = E\{d(U,V)\}$$

$$= \sum_{k=0}^{N-1} \int_{u_k}^{u_{k+1}} d(u,v_k) \cdot f_U(u) \cdot \mathrm{d}u$$

- Assume MSE

$$d(u,v_k) = (u - v_k)^2$$

$$D = \sum_{k=0}^{N-1} \int_{u_k}^{u_{k+1}} (u - v_k)^2 \cdot f_U(u) \cdot \mathrm{d}u$$

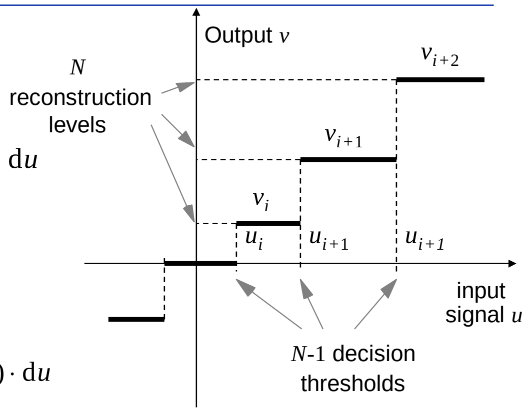

- Fixed code word length vs. variable code word length

$$R = \log N \quad vs. \quad R = -E\{\log P(v)\}$$

Lloyd-Max Quantizer

0: Given: a source distribution $f_U(u)$
a set of reconstruction levels $\{v_k\}$

1: Encode given $\{v_k\}$ (Nearest Neighbor Condition):

$$\alpha(u) = \text{argmin } \{d(u,v_k)\} \quad \rightarrow \quad u_k = (v_k + v_{k+1})/2 \text{ (MSE)}$$

2: Update set of reconstruction levels given $\alpha(u_k)$
(Centroid Condition):

$$v_k = \text{argmin } E\{\, d(u,v_k) \mid \alpha(u) = k\} \rightarrow v_k = \frac{\int_{u_k}^{u_{k+1}} u \cdot f_U(u)\,\mathrm{d}u}{\int_{u_k}^{u_{k+1}} f_U(u)\,\mathrm{d}u} \quad \text{(MSE)}$$

3: Repeat steps 1 and 2 until convergence

High Resolution Approximations

- Pdf of U is roughly constant over individual cells C_k

$$f_U(u) \approx f_k, \;\; u \in C_k$$

- The fundamental theorem of calculus

$$P_k = \Pr(u \in C_k) = \int_{u_k}^{u_{k+1}} f_U(u) \cdot \mathrm{d}u \approx (u_{k+1} - u_k) \cdot f_k = \Delta_k f_k$$

- Approximate average distortion (MSE)

$$D = \sum_{k=0}^{N-1} \int_{u_k}^{u_{k+1}} (u - v_k)^2 \cdot f_U(u) \cdot \mathrm{d}u = \sum_{k=0}^{N-1} f_k \int_{u_k}^{u_{k+1}} (u - v_k)^2 \,\mathrm{d}u$$

$$= \sum_{k=0}^{N-1} f_k \frac{\Delta_k^3}{12} = \frac{1}{12} \sum_{k=0}^{N-1} P_k \Delta_k^2$$

Uniform Quantization

- Reconstruction levels of quantizer $\{v_k\}, \; k \in K$ are uniformly spaced
- Quantizer step size, i.e. distance between reconstruction levels: Δ
- Average distortion

$$\sum_{k=0}^{N-1} P_k = 1, \;\; \Delta_k = \Delta$$

$$D = \frac{1}{12} \sum_{k=0}^{N-1} P_k \Delta_k^2 = \frac{\Delta^2}{12} \sum_{k=0}^{N-1} P_k = \frac{\Delta^2}{12}$$

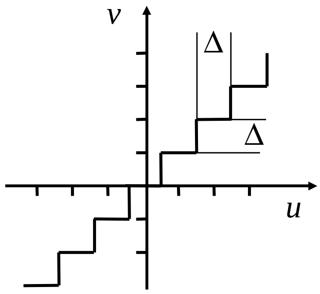

- Closed-form solutions for pdf-optimized uniform quantizers for Gaussian RV only exist for $N=2$ and $N=3$
- Optimization of Δ is conducted numerically

Panter and Dite Approximation

- Approximate solution for optimized spacing of reconstruction and decision levels
- Assumptions: high resolution and smooth pdf $\Delta(u)$

$$\Delta(u) = \frac{\text{const}}{\sqrt[3]{f_U(u)}}$$

- Optimal pdf of reconstruction levels is not the same as for the input levels
- Average Distortion $D \approx \frac{1}{12N^2}\left(\int_{\Re} f_U^{1/3}(u)\cdot \mathrm{d}u\right)^3$
- Operational distortion rate function for Gaussian RV

$$U \sim N(0,\sigma^2),\ D(R) \approx \frac{\pi\sqrt{3}}{2}\sigma^2 2^{-2R}$$

Entropy-Constrained Quantization

- So far: each reconstruction level is transmitted with fixed code word length
- Encode reconstruction levels with variable code word length
- Constrained design criteria:

$$\min D, \text{ s.t. } R < R_c \quad \text{or} \quad \min R, \text{ s.t. } D < D_c$$

- Pose as unconstrained optimization via Lagrangian formulation:

$$\min D + \lambda R$$

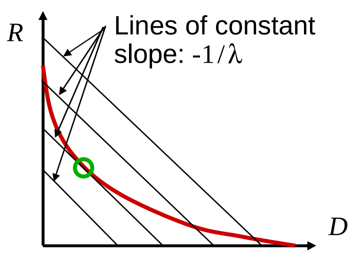

- For a given λ, an optimum is obtained corresponding to either R_c or D_c
- If λ small, then D small and R large if λ large, then D large and R small
- Optimality also for functions that are neither continuous nor differentiable

Chou, Lookabaugh, and Gray Algorithm*

0: Given: a source distribution $f_U(u)$
a set of reconstruction levels $\{v_k\}$
a set of variable length code (VLC) words $\{\gamma_k\}$
with associated length $|\gamma_k|$

1: Encode given $\{v_k\}$ and $\{\gamma_k\}$:
$$\alpha(u) = \text{argmin}\ \{d(u, v_k) + \lambda|\gamma_k|\ \}$$

2: Update VLC given $\alpha(u_k)$ and $\{v_k\}$
$$|\gamma_k| = -\log\ \text{P}(\alpha(u)=k)$$

3: Update set of reconstruction levels given $\alpha(u_k)$ and $\{v_k\}$
$$v_k = \text{argmin}\ E\ \{\ d(u, v_k)\ |\ \alpha(u)=k\}$$

4: Repeat steps 1 - 3 until convergence

*1989, has been proposed for Vector Quantization

Entropy-Constrained Scalar Quantization: High Resolution Approximations

- Assume: uniform quantization: $P_k = f_k \Delta$

$$R = -\sum_{k=0}^{N-1} P_k \log P_k = -\sum_{k=0}^{N-1} f_k \Delta \log(f_k \Delta)$$
$$= -\sum_{k=0}^{N-1} f_k \Delta \log(f_k) - \sum_{k=0}^{N-1} f_k \Delta \log(\Delta)$$
$$\sum \Delta \approx \int \mathrm{d}u \quad \downarrow$$
$$\approx \underbrace{\int_{\Re} f_U(u) \log(f_U(u)) \mathrm{d}u}_{\textit{Differential Entropy } h(U)} - \log \Delta \underbrace{\int_{\Re} f_U(u) \mathrm{d}u}_{1}$$
$$= h(U) - \log \Delta$$

- Operational distortion rate function for Gaussian RV
$$U \sim N(0, \sigma^2), \ D(R) \approx \frac{\pi\ e}{6} \sigma^2 2^{-2R}$$

- It can be shown that for high resolution:
Uniform Entropy-Constrained Scalar Quantization is optimum

Comparison for Gaussian Sources

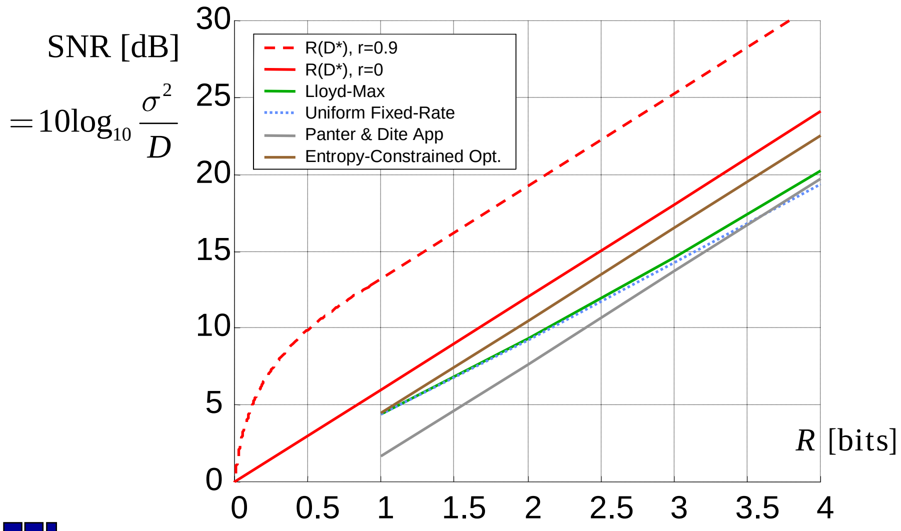

Vector Quantization

- So far: scalars have been quantized
- Encode vectors, ordered sets of scalars
- Gain over scalar quantization (Lookabaugh and Gray 1989)
  - Space filling advantage
    - Z lattice is not most efficient sphere packing in *K-D* $(K>1)$
    - Independent from source distribution or statistical dependencies
    - Maximum gain for $K \rightarrow \infty$: 1.53 dB
  - Shape advantage
    - Exploit shape of source pdf
    - Can also be exploited using entropy-constrained scalar quantization
  - Memory advantage
    - Exploit statistical dependencies of the source
    - Can also be exploited using DPCM, Transform coding, block entropy coding

Comparison for Gauss-Markov Source: $\rho = 0.9$

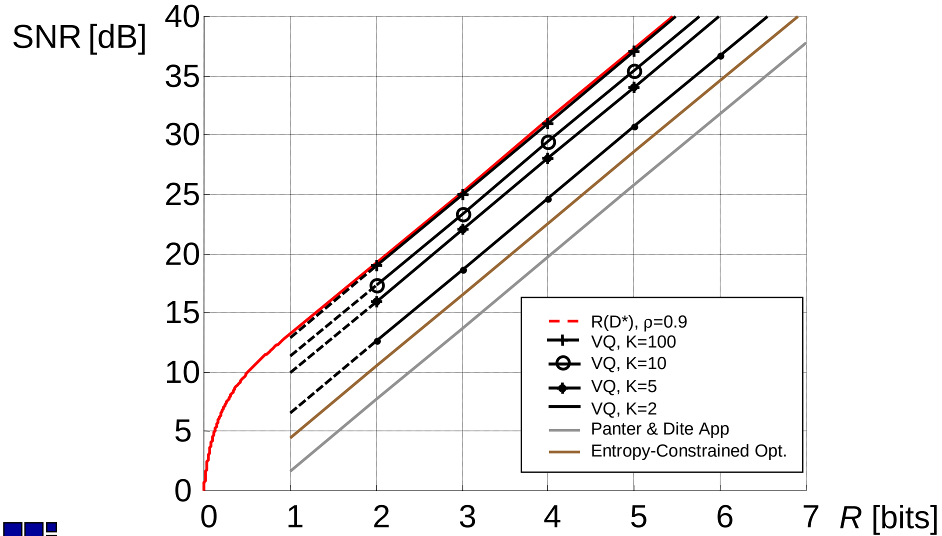

Vector Quantization II

- Vector quantizers can achieve $R(D^*)$ if $K \rightarrow \infty$
- Complexity requirements: storage and computation
- Delay
- Impose structural constraints that reduce complexity
- Tree-Structured, Transform, Multistage, etc.
- Lattice Codebook VQ

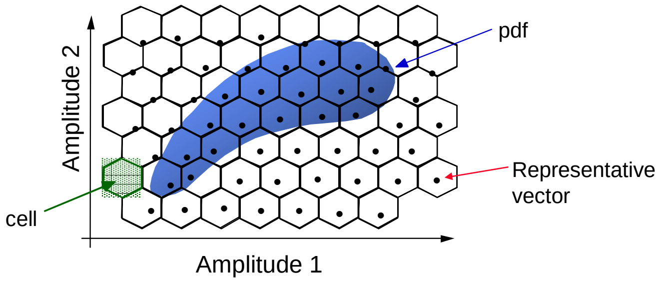

Summary

- Rate-distortion theory: minimum bit-rate for given distortion
- $R(D^*)$ for memoryless Gaussian source and MSE: 6 dB/bit
- $R(D^*)$ for Gaussian source with memory and MSE: encode spectral components independently, introduce white noise, suppress small spectral components
- Lloyd-Max quantizer: minimum MSE distortion for given number of representative levels
- Variable length coding: additional gains by entropy-constrained quantization
- Minimum mean squared error for given entropy: uniform quantizer (for fine quantization!)
- Vector quantizers can achieve $R(D^*)$ if $K \rightarrow \infty$ - Are we done ?
- No! Complexity of vector quantizers is the issue

Design a coding system with optimum rate distortion performance, such that the delay, complexity, and storage requirements are met.