Univ. Prof. Dr.-Ing. habil. Norbert Goertz
Institute of Telecommunications
TU Wien
norbert.goertz@tuwien.ac.at

Source Coding

Tutorial Problems

WS 2020/2021

Hints:

Parts of a problem that are marked by "$\Longrightarrow$" can be solved without the solutions of previous parts.

This material is supplementary but not relevant for the exam.

Problem 1

Consider uniform quantisation (with step-size Δ) of a source signal with unlimited amplitude range.

In the figure below both the characteristics of the granular distortion and the overload distortion are depicted versus the step-size Δ of the quantiser.

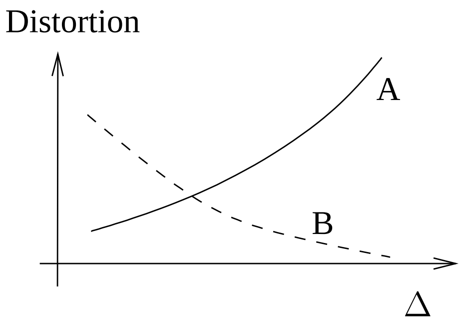

(a) Label the curves appropriately and give reasons for your answer.

⟹ (b) State, for a uniform quantiser with mid-riser characteristic, the reproducer values and the decision boundaries as functions of Δ and the rate R.

⟹ (c) Give the general formula for the power of the quantiser error and state how in principle the optimal value for Δ can be found (no calculations!).

Problem 2

In this problem we will consider high-rate R scalar uniform quantisation of a Laplacian source with the PDF

$$p(x) = \frac{1}{2}\, e^{-|x|} \ .$$

(a) Show how to compute the factor $K = \frac{P_x}{x_{\max}^2}$ (see lecture notes). Give the exact value for K dependent on the probability $p_1 \doteq \Pr\{|x| > x_{\max}\}$. Calculate the value of $x_{\max}$ for $p_1 = 0.001$.

(b) Compute the overload distortion for $p_1 = 0.001$. For this, assume approximately that the reproducer values that are used to quantise the overload regions are $\hat{x}_0 = -x_{\max}$ and $\hat{x}_{2^R-1} = x_{\max}$.

Is the true value of the overload-distortion over- or under-estimated by the approximation? (Why?)

(c) Is it justified to always neglect the overload distortion for $p_1 = 0.001$, or is that a matter of rate? To clarify this issue, assume that the granular distortion equals $P_Q = \Delta^2/12$, write the SNR in the form

$$\frac{SNR}{\text{dB}} = 6.02 \cdot R + 10 \log_{10}(3K) - 10 \log_{10}(1 + c) \ .$$

How large is $10 \log_{10}(1 + c)$ for $p_1 = 0.001$ and $R = 5$?

Problem 3

In this problem we will compute the SNR that results from logarithmic quantisation (A-law characteristic); the result is given in the lecture notes. We will assume a symmetric PDF $p(x)$ for the source signal samples.

(a) For the contribution to the total noise power that results from the i-th quantiser interval we obtain approximately

$$P_{Q_i} \approx \frac{\Delta x_i^2}{12} \int\limits_{\hat{x}_i - \Delta x_i/2}^{\hat{x}_i + \Delta x_i/2} p(x)\, dx = \frac{\Delta x_i^2}{12}\, p_i \ ,$$

where $\hat{x}_i$ is the i-th reproducer value and Δx_i is the size of the corresponding quantiser interval, which is *not* equal for all i.

Which approximation is used? Which quantity is described by p_i?

(b) The total quantiser noise power P_Q is given by the sum of all interval noise powers $P_{Q,i}$:

$$P_Q = \sum_{i=0}^{L-1} P_{Q_i} = 2 \sum_{i=L/2}^{L-1} P_{Q_i} \approx 2 \sum_{i=L/2}^{L-1} \frac{\Delta x_i^2}{12}\, p_i$$

Explain the second equality.

(c) The width of a quantiser interval follows from the compressor characteristic $y = g(x)$:

$$\Delta x_i \approx \frac{\Delta y}{g'(\hat{x}_i)}$$

Explain this by elementary differential calculus.

(d) In the derivation of the A-law compressor we required the interval-width (in the x-domain) to be proportional to the signal amplitude. From this and from (c) we obtain

$$\frac{1}{g'(x)} \stackrel{!}{=} c\, x \quad \text{with } c \text{ a constant}$$

Determine an expression for Δx_i from (c) and the expression above.

(e) For the total noise power we finally obtain

$$P_Q \approx \frac{\Delta y^2}{12}\, c^2\, P_x \quad ,$$

where P_x is the power of the input signal.

Derive this formula.

(f) The compressed samples y are uniformly quantised by R bits per sample. Assume that $x_{\max} = y_{\max} = 1$ (which means no loss of generality) and insert a relation for Δy into the solution of (e).

(g) Compute the SNR in dB according to

$$\frac{SNR}{\text{dB}} = 10\ \log_{10}\left(\frac{P_x}{P_Q}\right)$$

The result is

$$\frac{SNR}{\text{dB}} = R \cdot 20\ \log_{10}(2) + 10\ \log_{10}(3) - 20\ \log_{10}(c)$$

Give an interpretation of this result. Does it depend on the signal power?

(h) Compute the factor c for the A-law case in the amplitude range $\frac{1}{A} \leq |x| \leq 1$! (Hint: $1/g'(x) = c\,x$)

(i) As a result we obtain

$$\frac{SNR}{\text{dB}} = R \cdot 6.02 - 9.99 \quad \text{for} \quad \frac{1}{A} \leq |x| \leq 1$$

Compare this result with conventional uniform quantisation.

(j) For small signal amplitudes (we implicitly assumed large amplitudes above) the logarithmic quantiser behaves like a uniform quantiser (with a different bit rate).

Why and for which amplitudes of the input signal x is this the case?

(k) Compute the gradient of the A-law characteristic for small values of x.

(l) Determine the effective interval size of the logarithmic quantiser (A-law) for small values of x.

(m) Give the resulting SNR and compare it with the result of a conventional uniform quantiser with the same bit rate R.

Problem 4

Consider the probability density function (PDF)

$$p(x) = \begin{cases} |x|, & -1 \le x \le 1 \\ 0, & \text{else} \end{cases} .$$

Independent realisations of the random variable X with the PDF $p(x)$ are to be scalar quantised by $R = 2$ bits per sample.

(a) Sketch the PDF $p(x)$ and determine its mean μ_x, its variance σ_x^2, and the power P_x of a random process that has independent realisations, each with the PDF $p(x)$.

⟹ (b) Give the quantiser decision boundaries $g_i, i = 0, 1, ..., 4$, and the reproducer values $\hat{x}_i, i = 0, 1, ..., 3$, for uniform quantisation (mid-riser characteristic, $x_{\max} = 1$).

Sketch the quantiser characteristic; label the axes.

(c) Give the general formula to compute the power P_Q of the quantiser noise (dependent on the reproducer values and the decision boundaries)?

(d) Show that in our case

$$P_Q = \frac{3}{4}\hat{x}_0^2 + \frac{7}{6}\hat{x}_0 + \frac{1}{4}\hat{x}_1^2 + \frac{1}{6}\hat{x}_1 + \frac{1}{2}$$

holds. For this, simplify the general expression for P_Q in (c) by use of the symmetries of the PDF and insert numerical values only for the decision boundaries g_i.

(e) Compute the value of P_Q and give the quantiser SNR in dB.

In what follows the quantiser decision boundaries g_i will be adopted from (b), but the reproducer values $\hat{x}_i, i = 0, 1, ..., 3$, shall be selected such that the power P_Q of the quantiser error is minimised.

⟹ (f) Compute the new reproducer values $\hat{x}'_i$, $i = 0, 1, ..., 3$.

(g) Sketch the new quantiser characteristic and label the axes with variables and numerical values.

(h) Calculate the new power P'_Q of the quantiser error and the new SNR' in dB.

(i) Compare the SNR-values in (e) and (h) and explain potential differences.

(j) How could the quantiser performance be further improved?

Problem 5

We want to scalar quantise independent Gaussian samples X with the PDF

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-\frac{1}{2\sigma^2}x^2) .$$

Further, we assume that the rate R (bits per sample) is high.

(a) Determine the signal power P_x and give the quantiser SNR of a scalar uniform quantiser for an overload probability of $P_o = \Pr\{|x| > x_{\max}\} = 0.001$ (neglect the overload distortion in the SNR computation).

⟹ (b) How large is the SNR in dB of the best possible scalar *fixed-rate* quantiser?

⟹ (c) Which maximum SNR can be obtained by the best scalar *variable-rate* scalar quantiser?

⟹ (d) What is the maximum SNR in dB of *any* quantisation scheme without *any* restrictions.

In what follows we will assume that the samples X are quantised by a companding scheme, which uses the compressor-function

$$y = g(x) = \alpha \cdot \frac{2}{\pi} \cdot \arctan(\beta x) \quad \text{with} \qquad \alpha, \beta > 0 .$$

In the y-domain, the samples are uniformly quantised.

⟹ (e) Into which interval $(y_{\min}, y_{\max})$ is the support $(-\infty, +\infty)$ of the PDF $p(x)$ mapped by the compressor-characteristic $g(x)$?

(f) Compute the quantiser distortion $P_Q(R)$ as a function of α and β? Do both variables influence P_Q?
Hint:

$$\int_0^\infty x^n \cdot e^{-ax^2} dx = \frac{1 \cdot 3 \cdot ... \cdot (2k-1)\sqrt{\pi}}{2^{k+1} \cdot a^{k+1/2}} \quad \text{for} \quad n = 2k,\ k = 1, 2, ...$$

(g) Determine the variables α and β such that the quantiser distortion computed in (f) is minimised and calculate the optimal quantiser SNR in dB.

Now we assume a fixed *low* rate of $R = 1$ bit / sample, and we want do design the best possible fixed-rate scalar quantiser; we assume, however, symmetric reproducer values ($\hat{x}_0 = -\hat{x}_1$).

⟹ (h) Determine the decision boundaries g_0, g_1, g_2 (use the symmetry of the source PDF).

(i) Calculate the optimal reproducer values.

(j) Calculate the quantiser distortion P_Q and the corresponding SNR in dB.

(k) Discuss the difference in SNR of the results from (j) and (b).

Problem 6

In the figure (see overleaf), code-vectors of a two-dimensional vector quantiser are given.

(a) How many bits are required to transmit the code-vectors by a fixed bit rate per index? Are the bits efficiently used (in terms of source coding)?

$\Longrightarrow$ (b) Insert the Voronoi-regions into the figure; assume that the mean squared error is used as a distance measure.

Hint: consider first, how the partitioning of the space looks like, if only two code-vectors are used.

$\Longrightarrow$ (c) Is there any potential gain from using bits inefficiently from a source-coding point-of-view?

Hint: think about bit errors.

$\Longrightarrow$ (d) Assume the codevectors are uniformly distributed. How large is their entropy?

(e) Construct a Huffman code for the codevectors. How large is the average word-length? Compare the result with the entropy and with a fixed-rate index allocation.

(Problem 6)

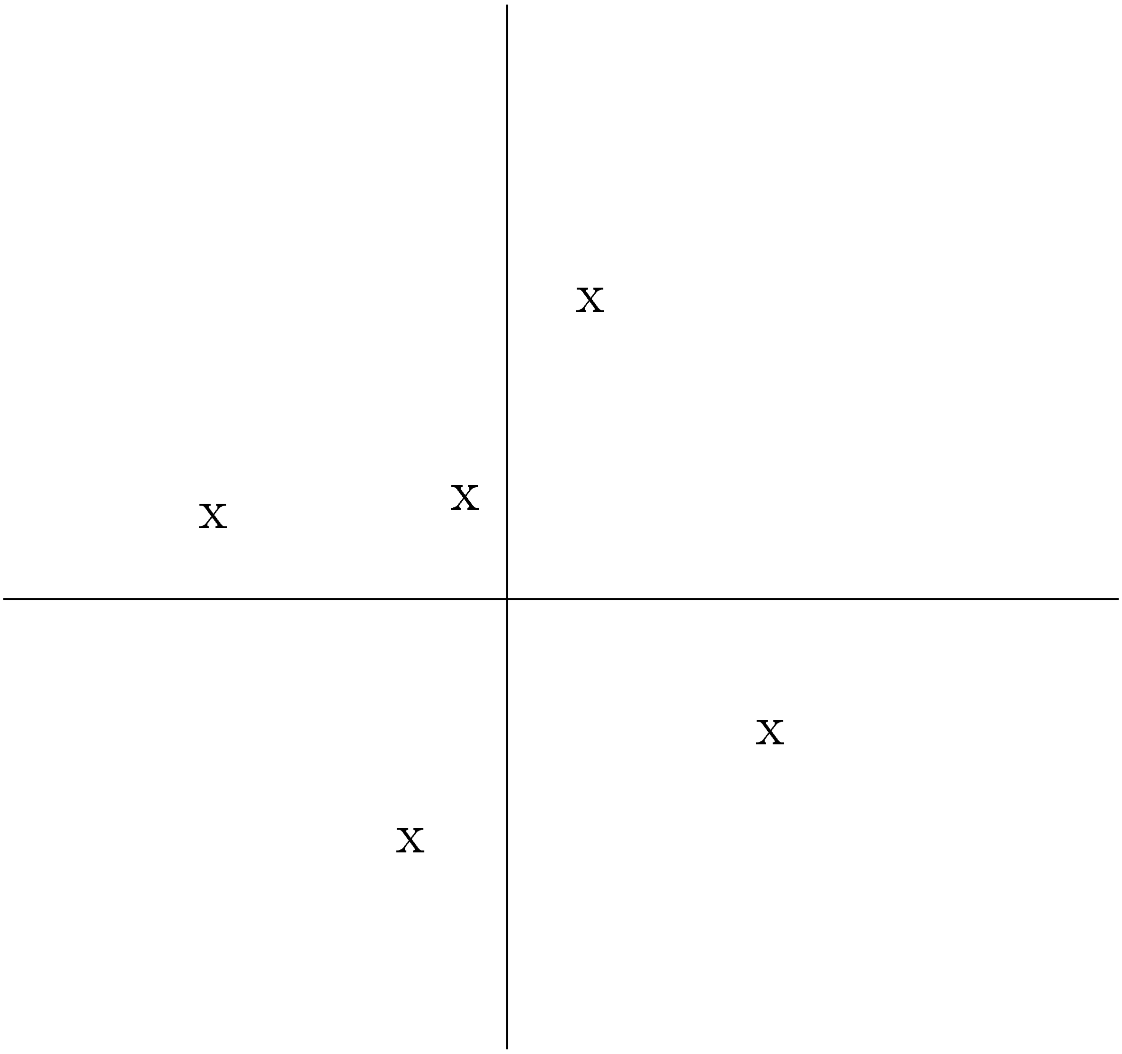

Problem 7

In this problem, we will illustrate how the LBG algorithm for vector quantiser codebook training works.

In the figure entitled "1. Iteration," initial code-vectors (marked by "x") are given; the mid-points of lines connecting them are marked by "o". Moreover, data points are also given by the dots.

We use the mean squared error (MSE) as a distance measure.

(a) Show analytically that the centroids are given by the average of the data-vectors that are allocated to a quantiser region (Voronoi region) by the first step of the iterative training procedure.

$\Longrightarrow$ (b) Add the Voronoi regions for the given code-vectors into the figure entitled "1. Iteration."

(c) Determine (graphical approximation, no calculations) the location of the "new" code-vectors in each Voronoi region.

(d) Copy the new code-vectors resulting from (c) into the figure entitled "2. Iteration" and determine the Voronoi regions for these code-vectors.

(e) Determine the location of the code-vectors after the second iteration.

Now we consider vector quantisation with the following modified distance measure:

$$d(\underline{x}_k, \underline{\hat{y}_i}) = f(\underline{x}_k) \cdot \sum_{j=0}^{N-1} (x_k(j) - \hat{y}_i(j))^2 \quad \text{with} \quad f(\underline{x}_k) > 0$$

The data vector to be quantised is given by $\underline{x}_k = \{x_k(j), j = 0, 1, ..., N-1\}$, the codevector with the index i is given by $\underline{\hat{y}}_i = \{\hat{y}_i(j), j = 0, 1, ..., N-1\}$.

$\Longrightarrow$ (f) Assume a codebook with the codevectors $\underline{\hat{y}}_i$ is given and some data vector $\underline{x}_k$ is to be quantised: will the use of the modified distance measure make a difference compared with the conventional MSE distance measure?

$\Longrightarrow$ (g) Compute, for a general function $f(\underline{x}_k)$, the "new" centroids of the partition regions in Step 4 of the LBG algorithm.

(h) We now assume

$$f(\underline{x}_k) = \frac{1}{||\underline{x}_k||^2} = \frac{1}{\sum\limits_{j=0}^{N-1} x_k^2(j)} .$$

Discuss the influence of the distance measure for the location of the code-vectors that result from LBG codebook training.

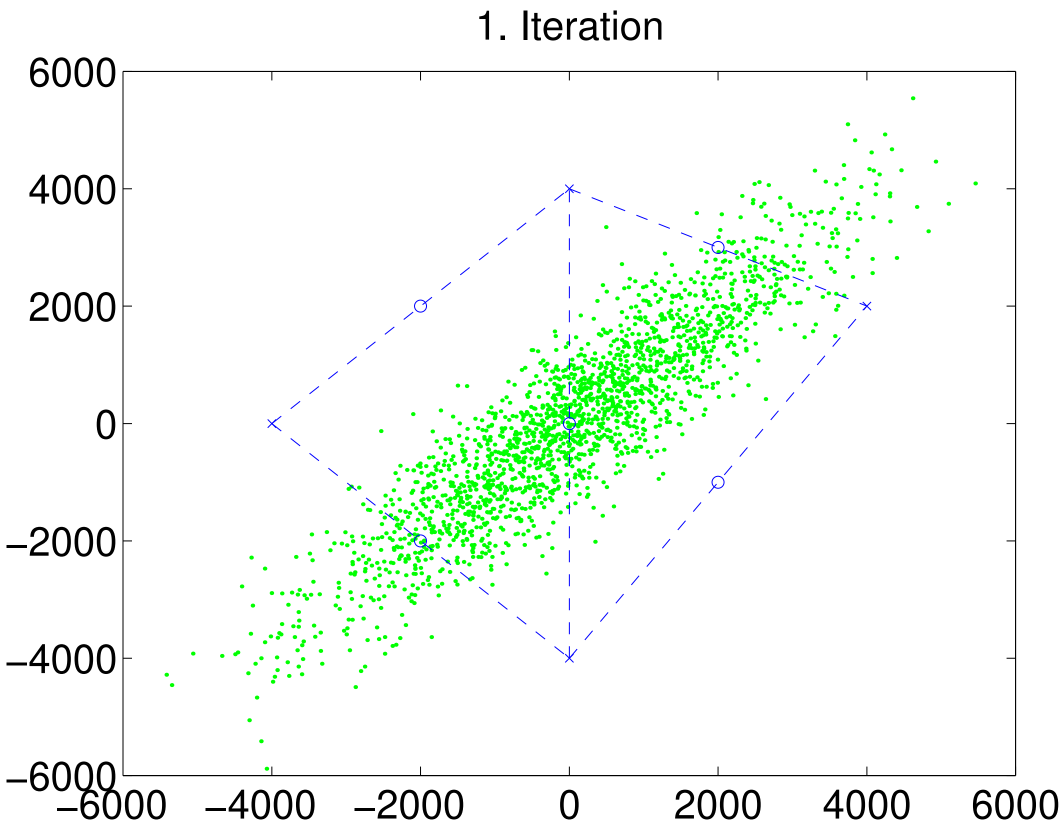
1. Iteration
6000
4000
2000
0
−2000
−4000
−6000
−6000
−4000
−2000
0
2000
4000
6000

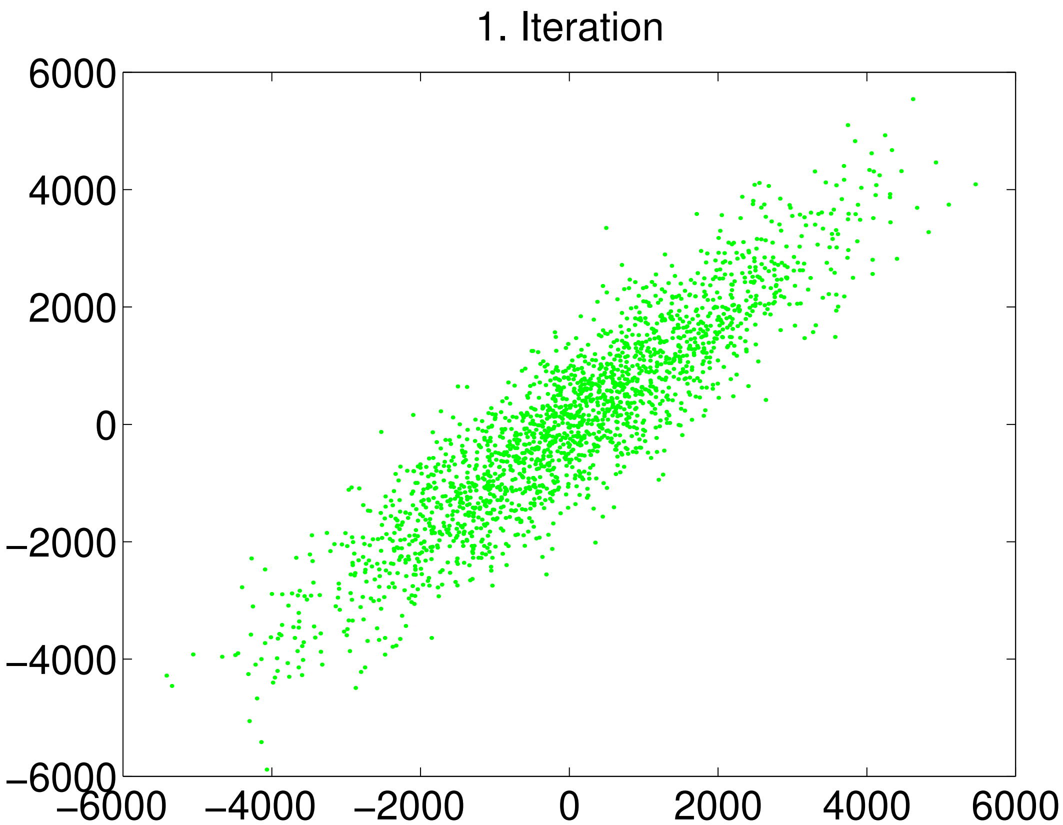
1. Iteration
6000
4000
2000
0
−2000
−4000
−6000
−6000
−4000
−2000
0
2000
4000
6000

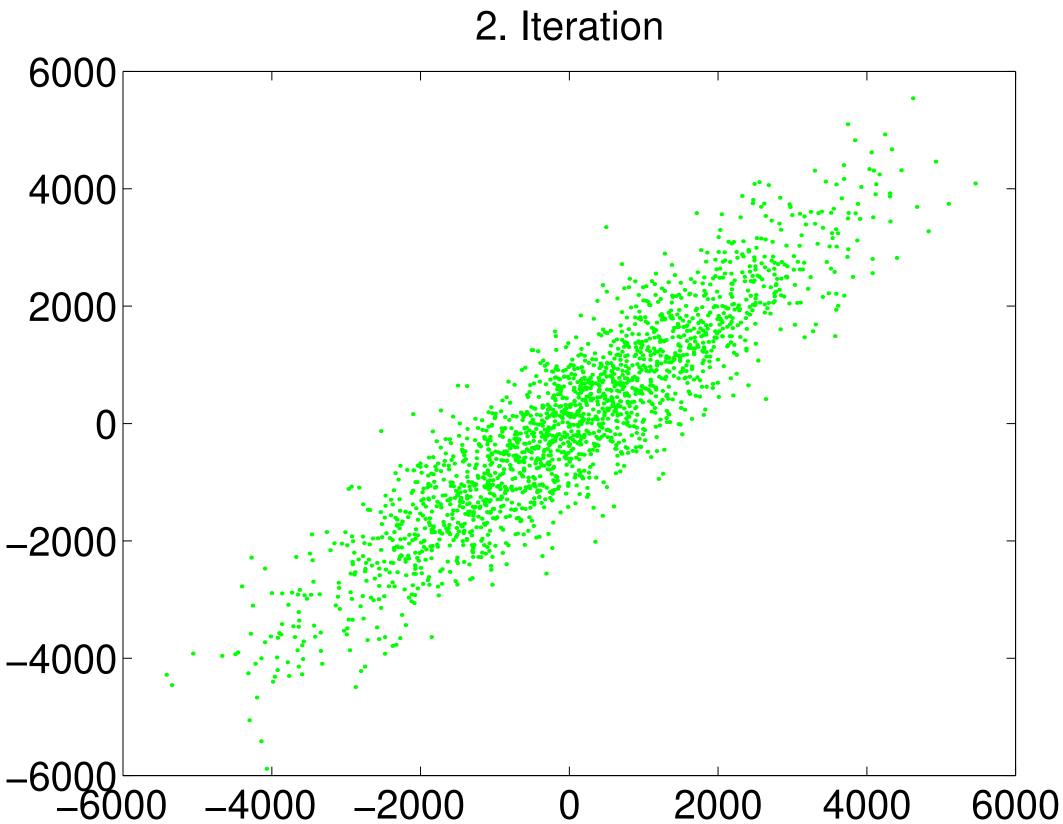
2. Iteration
6000
4000
2000
0
−2000
−4000
−6000
−6000
−4000
−2000
0
2000
4000
6000

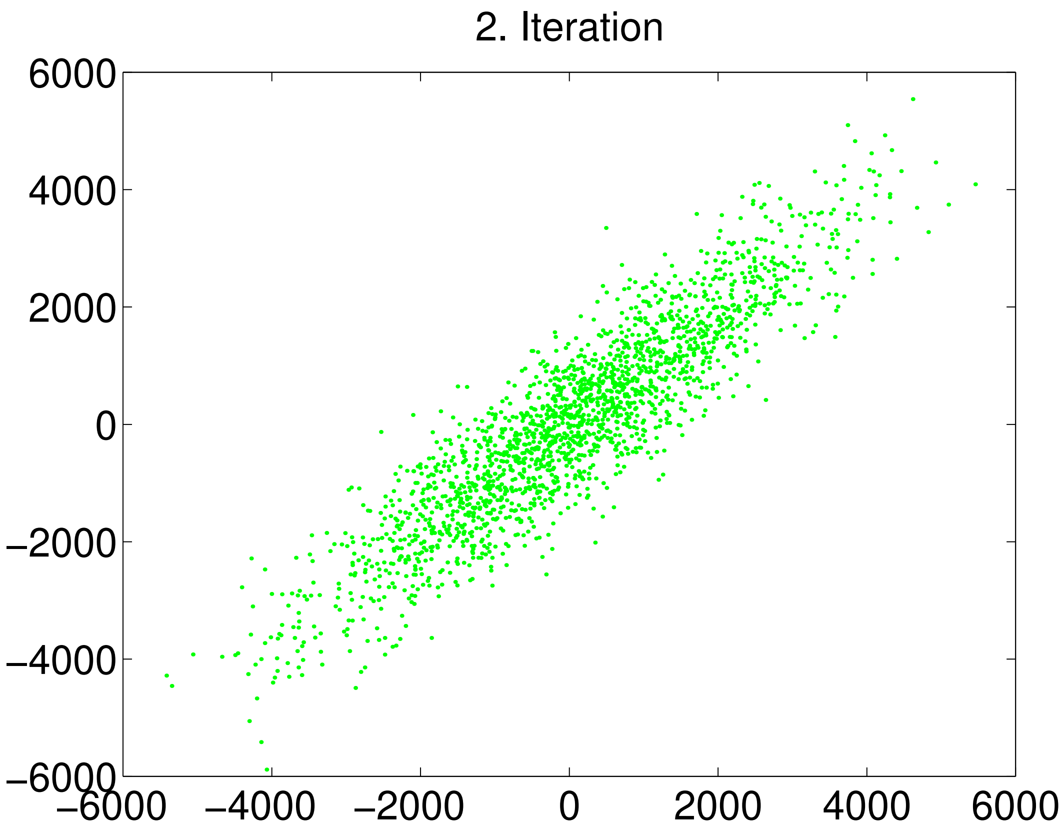
2. Iteration
6000
4000
2000
0
−2000
−4000
−6000
−6000
−4000
−2000
0
2000
4000
6000

Problem 8

In the lectures, we found the rate-distortion function for a memoryless Gaussian source and squared error distortion.

There is one more case in which an analytical result is known: the binary Bernoulli(p) source[1] with an expected portion of errors in the sequence less than or equal to D. The corresponding per-letter distortion measure (the Hamming distortion measure) is simply "1" if $X_i \neq \tilde{X}_i$ and "0" if $X_i = \tilde{X}_i$, where X_i is a letter from the source sequence and $\tilde{X}_i$ its reproduction.

In this problem we will prove the rate-distortion function

$$R(D) = \begin{cases} H_2(p) - H_2(D), & 0 \leq D \leq \min\{p, 1-p\} \\ 0, & D > \min\{p, 1-p\} \end{cases}$$

for the Bernoulli source and the Hamming (per-letter) distortion measure. The quantity $H_2(p) = -p \cdot \log_2(p) - (1-p) \cdot \log_2(1-p)$ is the so-called binary entropy function.

The steps of the proof are similar to the Gaussian case.

(a) State the general expression for the (information) rate-distortion function for sources with letter alphabets of limited size.

(b) Find a sequence of inequalities (as in the Gaussian case) to derive a lower bound for the mutual information $I(X; \tilde{X})$.
Hint: consider the conditional entropy of $X \oplus \tilde{X}$ given $\tilde{X}$, where "$\oplus$" denotes modulo-2 addition.

(c) In the second step of the proof, consider the following system:

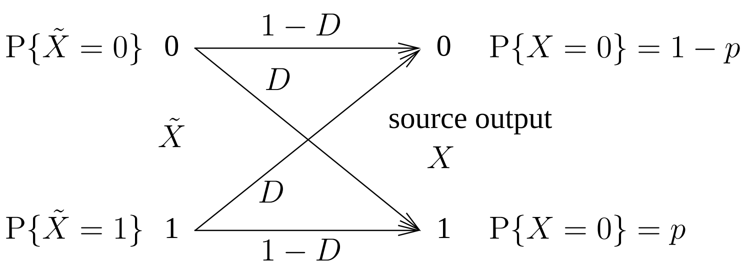

Compute the probabilities $\Pr\{\tilde{X} = 1\} = 1 - \Pr\{\tilde{X} = 0\}$.

Prove[2] that the system achieves the lower bound derived in (c).

[1]Such a source emits a "1" with probability p and a "0" with probability $1-p$.
[2]Only if really necessary, consider the book by Cover, Thomas.

Problem 9

We want to show that the inverse of the spectral flatness measure for a Gauss-Markov process, which is generated by linear filtering of a Gaussian source with the transfer function

$$H(z) = \left(1 + \sum_{i=1}^{N_z} b_i \cdot z^{-i}\right) \Bigg/ \left(1 + \sum_{i=1}^{N_p} a_i \cdot z^{-i}\right) ,$$

is given by

$$\frac{1}{\gamma_x^2} = \frac{\sigma_x^2}{\sigma_r^2} = \frac{1}{2\pi} \int_{-\pi}^{+\pi} |H(\Omega)|^2 \, d\Omega = \sum_{k=0}^{\infty} h_0^2(k)$$

where $h_0(k)$ is the discrete-time impulse response of the filter.

(a) What is the definition of the spectral flatness measure?

(b) Factorise the polynomials in the numerator and the denominator of $H(z)$ assuming general poles $z_{\infty,i}$ and zeros $z_{0,i}$, but for simplicity assume that all poles and zeros are real numbers.

(c) Write the square of the magnitude of a complex number by an alternative expression (that does not use "$|\cdot|$").

(d) Use the hint below to compute the spectral flatness measure. Consider that the filter must be stable and have minimum phase (what does that mean for the poles and zeros?).

(e) Explain the rightmost equality in the formula for the spectral flatness measure stated above.

Hint:

$$\int_{0}^{\pi} \log_e(a^2 - 2ab\cos(x) + b^2)dx = \begin{cases} 2\pi \log_e(a), & a \geq |b| > 0 \\ 2\pi \log_e(b), & b \geq |a| > 0 \end{cases}$$

Problem 10

We are given a memoryless discrete-time Gaussian source which emits the samples $r(k)$. The latter are filtered by a discrete-time causal filter with the transfer function

$$H(z) = \frac{z^2}{(z-a)(z-b)} \quad \text{with} \quad a = \frac{1}{4},\ b = \frac{3}{4}\ .$$

The resulting output signal is denoted by $x(k)$.

In what follows, we will assume high-rate coding.

(a) What is the maximum SNR we can obtain by coding the source with the samples $r(k)$ at some rate R with a coding scheme without any restrictions?

(b) Is the result (a) also true for low rates? Why?

(c) Compute the impulse response $h_0(k)$ of the filter.

(d) Determine the inverse $1/\gamma^2$ of the spectral flatness measure of $x(k)$.

(e) Which maximum SNR (in dB) can be achieved by any coding scheme (no restrictions) if the signal $x(k)$ is coded at some high rate R?

(f) Which PDF have the individual samples $x(k)$?

(g) Which SNR in dB would result from a scalar fixed-rate optimal quantisation of $x(k)$ without use of the signal correlation for coding?

(h) Which maximum SNR in dB would result from a scalar fixed-rate optimal quantisation if we use a certain coding system that exploits the signal correlation of $x(k)$? What are the names of such a system?

Hints:

$$\sum_{i=0}^{\infty} q^i = \frac{1}{1-q} \quad \text{for} \quad 0 < |q| < 1$$

$$Z\{\alpha^k \gamma_{-1}(k)\} = \frac{z}{z-\alpha} \quad \text{with} \quad \gamma_{-1}(k) = \begin{cases} 1, & k \geq 0 \\ 0, & k < 0 \end{cases}$$