Spring 2003: EE223 Homework 3
Solutions due Mar 6
Collaboration among students is encouraged

Problem 3.1 *Scalar quantization of a uniform*

Suppose that X is a uniform continuous random variable between $[-1, +1]$ and we want to do scalar quantization on it so as to minimize mean squared error. Give a two bit quantizer for it and prove that it is optimal in the sense that no other scalar quantizer can achieve better mean squared error using only two bits. (Hint: show that it uniquely satisfies the two conditions on optimal quantizers.) Can you generalize to N bits?

Problem 3.2 *High rate scalar quantization*

Suppose that we want to quantize X in the high rate regime, but instead of wanting to minimize the entropy rate of the quantized points, we are just interested in using M regions and in minimizing the resulting mean squared error. Our intuition suggests that we should use smaller regions where the probability is larger and larger regions where the probability is smaller. This problem is designed to help us get a quantitative handle on that.

Suppose that X has a simple non-uniform density:

$$f_X(x) = \begin{cases} 0 & \text{if } x < -L_1 \\ f_1 & \text{if } -L_1 \leq x < 0 \\ f_2 & \text{if } 0 \leq x < L_2 \\ 0 & \text{if } x \geq L_2 \end{cases}$$

where the f_1, f_2 are required to be such that the probability density integrates to 1.

Suppose also that we quantize in such a way as to use quantization regions of width Δ_1 on the side that has density f_1 and Δ_2 on the side that has density f_2.

a. *If we require M regions total, what is the constraint that the Δ_1 and Δ_2 must satisfy?*

b. *What is the mean squared error if we use this scheme with a particular choice of Δ_1, Δ_2?*

c. *Relax the integrality constraint on M and set up a Lagrange multiplier on the constraint from part a. Find the optimal choice of Δ_1, Δ_2 which minimizes mean squared error and has M regions where M need not be an integer.*

d. *For the above optimal solution, how does the mean squared error go down with increasing M in the limit of large M?*

e. *Compare the performance of the quantizer above with the performance of a uniform quantizer with M levels and also to the quantizer that just equalizes probability so that $f_1\Delta_1 = f_2\Delta_2$.*

f. *Calculate the entropy of the quantizer output of the quantizer in part c. and how does it increase with increasing M in the limit of large M?*

Problem 3.3 *Scalar quantization of noisy signals*

Suppose that instead of having access to X itself, we only have access to $Y = X + N$ where N is independent zero-mean Gaussian noise with variance σ^2. If we apply scalar quantization to Y, we get $\hat{Y}$ which we can interpret as $\hat{Y} = X + N_q$ where N_q is a combination of the effects of "quantization noise" and underlying noise. This often occurs when we are sampling noisy analog signals for further digital processing.

A natural choice for how to measure performance is to look at the effective signal energy to noise energy ratio (SNR) after quantization: $\frac{E[X^2]}{E[N_q^2]}$.

a. *Show that N_q is not generally independent of X.*

b. *Suppose that we are in the high rate regime. How does the effective SNR change with large M? (Notice that it does not tend to 0 and is lower bounded by the original pre-quantization SNR)*

c. *Assume that we are in the large noise regime where $E[X^2]$ is much smaller than σ^2. (The next problem will shed some light on when this can occur.) Consider a one-bit quantizer (just detect sign). Suppose that the underlying X only took on the values $\pm A$ each equally likely, where A is small. Given reconstruction points $-B$ and $+B$, what is the effective SNR like after one-bit quantization? (You can use any combination of simulation, analysis, etc. to get at an answer.)*

d. *How does the SNR vary with B and what is the optimal choice assuming you knew A?*

e. *In the large noise regime, what is the SNR price of using 1-bit scalar quantization as opposed to high-rate quantization?*

f. *Use a similar model and a combination of simulations and analysis to determine the rough SNR price of using 2-bit quantization in the large noise regime.*

g. *Comment.*

Problem 3.4 *Sampling of signals*

Our goal will be to sample $X(t)$. Sometimes we will have access to $X(t)$ directly, and at other times we have access to $Y(t) = X(t) + N(t)$ where $N(t)$ is white Gaussian noise with noise power σ^2 per Hz.

a. *Suppose that $X(t)$ is bandlimited to $\pm\frac{1}{2}$ Hz. Use the standard Shannon-Nyquist sampling theorem to tell us the minimum rate we need to sample the signal using bandlimited sampling (using sinc functions) in order to be able to reconstruct it perfectly.*

b. *If we sample $X(t)$ at a rate faster than that (using bandlimited sampling), is there any harm or benefit?*

c. *Please give a discrete time model for the samples of $Y(t)$ if those samples are taken using sinc functions spaced by t_s. How does the per sample SNR change with smaller t_s and thus faster sampling?*

d. *Is there any harm or benefit from sampling $Y(t)$ at a rate faster than the rate calculated in part a?*

e. *Use any combination of simulation and analysis to compare the quality of reconstructions of $X(t)$ using M-bit scalar quantization of the samples at the Shannon-Nyquist rate with using 1-bit quantization of the samples taken using M times the Shannon-Nyquist rate.*

f. *Suppose instead that $X(t)$ was not bandlimited. Instead, it was generated as follows: $X(t) = X_{\lfloor t \rfloor}$ where the X_i are i.i.d. Gaussian zero-mean random variables with variance 0.1. Show that $X(t)$ can be perfectly reconstructed if we sampled it at rate 1 by projecting onto the shifts of the unit pulse.*

g. *Using the $X(t)$ from part f., suppose instead that we sampled it using ideal low-pass filtering followed by impulsive sampling at rate 1. Use any arguments you would like to answer the question as to whether we can still perfectly reconstruct typical $X(t)$?*

h. *Suppose that we have $Y(t)$ from part f. and we use the sampling from part f. to get $\hat{Y}_i$. Give an expression for $\hat{Y}_i$ in terms of X_i and noise. If we were to use these samples $\hat{Y}_i$ to generate LLSE estimates of X_i, what is the resulting mean squared estimation error?*

i. *Now assume that we sample $Y(t)$ using sinc-sampling at rates $1, 2, 3, 4$. Use any arguments that you would like to estimate the resulting expected LLSE estimation error on X_i as a function of the sampling rate.*

j. *Comment.*