COM-202 - Signal Processing

Solutions for Homework 12

Exercise 1. Quantization

Consider a white random process $\mathbf{x}$ whose samples are uniformly distributed over the interval $[-1,1]$. The signal is sent through a uniform quantizer with step size Δ to obtain the signal $\hat{\mathbf{x}}$. The quantization error can be modeled as a white noise process whose samples are uniformly distributed over the interval $[-\Delta/2, \Delta/2]$; each quantized sample is thus

$$\hat{x}[n] = Q(x[n]) = x[n] + e[n]$$

where $e[n]$ is independent of $x[n]$.

After quantization, the signal is processed by a filter with impulse response

$$h[n] = \frac{a^n + (-a)^n}{2} u[n]$$

to obtain the signal $\mathbf{y} = \mathbf{h} * \hat{\mathbf{x}}$.

(a) compute the SNR of the quantized signal $\hat{\mathbf{x}}$

(b) compute the SNR of the filtered signal $\mathbf{y}$ (ie the power ratio between the filtered clean signal and the filtered quantization noise)

(c) could you improve the SNR using a different filter?

Solution:

(a) the signal is i.i.d. and uniformly distributed over $[-1,1]$ *so its variance is* $\sigma_x^2 = 1/3$. *The error is also i.i.d, and and uniformly distributed over* $[-\Delta/2, \Delta/2]$ *so its variance is* $\sigma_e^2 = \Delta^2/12$. *The SNR is thus*

$$SNR_x = \frac{\sigma_x^2}{\sigma_e^2} = \frac{4}{12}\frac{12}{\Delta^2} = \left(\frac{\Delta}{2}\right)^{-2}$$

(b) both signal and quantization error are white signals so their PSDs are

$$P_x(\omega) = \sigma_x^2$$
$$P_e(\omega) = \sigma_e^2$$

The PSD of the filter's output is

$$\begin{aligned} P_y(\omega) &= |H(\omega)|^2 P_{\hat{x}}(\omega) \\ &= |H(\omega)|^2 (P_x(\omega) + P_e(\omega)) \\ &= |H(\omega)|^2 P_x(\omega) + |H(\omega)|^2 P_e(\omega) \end{aligned}$$

where the last term is the PSD of the error at the filter's output. Since the variance of a random process is the integral of its PSD, the SNR at the ouptut of the filter is

$$\begin{aligned} SNR_y &= \frac{\int_{-\pi}^{\pi} |H(\omega)|^2 P_x(\omega) d\omega}{\int_{-\pi}^{\pi} |H(\omega)|^2 P_e(\omega) d\omega} \\ &= \frac{\sigma_x^2}{\sigma_e^2} \frac{\int_{-\pi}^{\pi} |H(\omega)|^2 d\omega}{\int_{-\pi}^{\pi} |H(\omega)|^2 d\omega} \\ &= SNR_x \end{aligned}$$

(c) *the SNR of the filter's output does not depend on the filter! Moral of the story: a linear filter cannot improve the SNR due to quantization if the quantizer matches the input.*

Exercise 2. 2D FIR filtering

Consider a two-dimensional discrete-space signal $x[n_1, n_2]$ that is nonzero only over a 8×8 square region. The following figure, called a *support representation*, shows as individual dots the locations of the nonzero values of the signal:

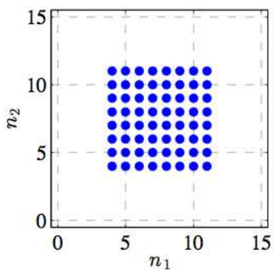

Consider now a set of 2D filters with impulse responses $h_i[n_1, n_2]$, for $i = 1, \ldots, 6$. All filters are 3×3 FIRs and the positions of their nonzero samples are shown in the following figure:

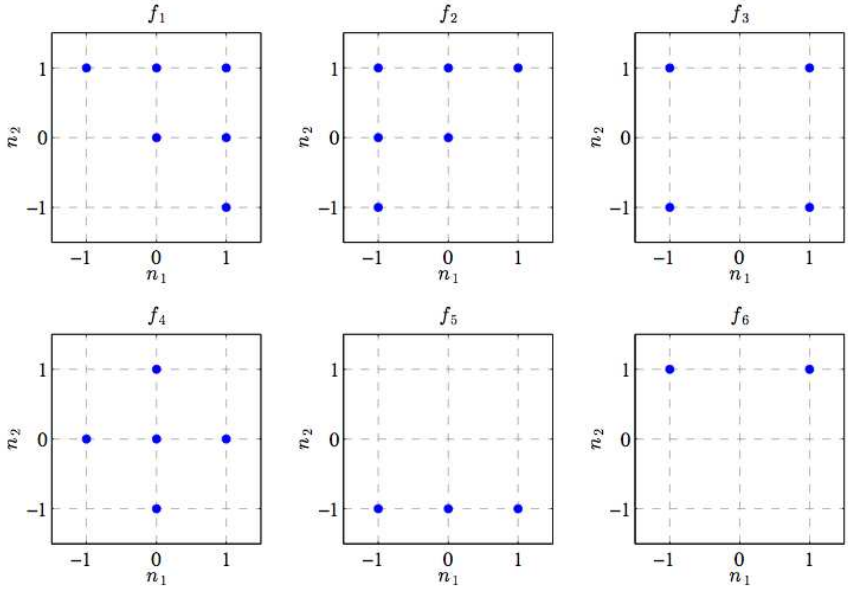

Call $y_i[n_1, n_2] = (h_i * x)[n_1, n_2]$ the signal obtained by filtering $x[n_1, n_2]$ with $f_i[n_1, n_2]$; since the input signal has finite support, $y_i[n_1, n_2]$ will also be a finite-support signal.

The following figure shows the support representations of $y_i[n_1, n_2]$ in random order; for each one of them, determine the correct value of i. If you prefer, you can also solve this exercise by writing some Python code.

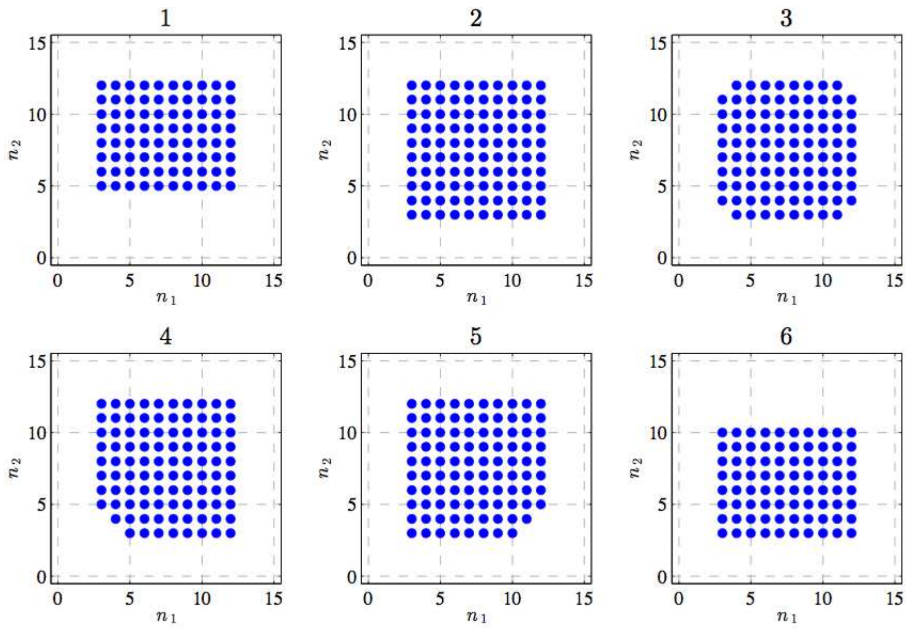

Solution:

the six (filter → plot) pairs are: $h_1 \to 4, h_2 \to 5, h_3 \to 2, h_4 \to 3, h_5 \to 6, h_6 \to 1$

You can determine these associations by mentally "sliding" the space-reversed impulse re-

sponses over the support of the input as in the computation of the convolution sum; but rememeber that in 2D you need to flip the impulse response along both dimensions.

Exercise 3. Bilinear interpolation

A square image of size $N \times N$ is built in the following way:

- the pixels in the upper left corner and in the lower right corner are set to a value of 255 (white)
- the pixels in the upper right corner and in the lower left corner are set to a value of 0 (black)
- every other pixel in the image is the result of a bilinear interpolation between the four corners

Which one of the following plots shows the correct image? Feel free to write a few lines of Python to help you find the answer.

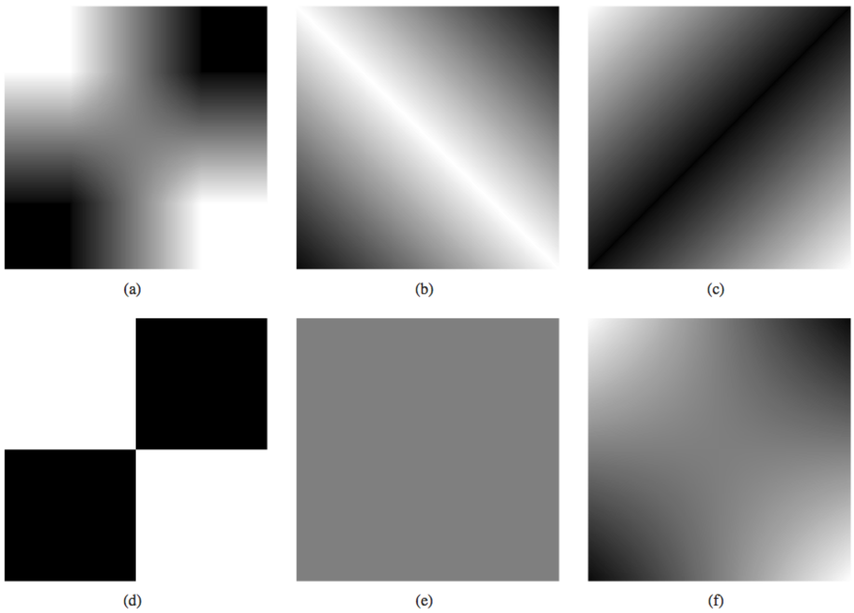

Solution:

Image (f). Here's a short code snippet to check the result

```
import matplotlib.pyplot as plt
import numpy as np

N = 128
```

```
x = np.zeros((N, N))
x[0,0] = x[N-1, N-1] = 255

# bilinear interpolation from four corners
for n1 in range(N):
  for n2 in range(N):
    b = x[0,0] * (N - n1) * (N - n2) +\
        x[0, N-1] * (N - n1) * n2  +\
        x[N-1, N-1] * n1 * n2 +\
        x[N-1, 0] * n1 * (N - n2)
        x[n1, n2] = b / (N * N)

plt.gray();
plt.matshow(x);
```

Exercise 4. Deadzone Quantization

We have seen that one of the fundamental ingredients of JPEG compression is the *deadzone* quantizer, i.e. a quantizer with a quantization interval centered around zero. To see the effects of deadzone quantization on SNR consider the following problem.

Assume $\mathbf{x}$ is an i.i.d. discrete-time random signal with $|x[n]| \leq 1$ for all n. Consider the following uniform 2-bit quantizers for the $[-1, 1]$ interval:

uniform quantizer

$$Q_u(x) = \begin{cases} 3/4 & \text{if } 1/2 \leq x \leq 1 \\ 1/4 & \text{if } 0 \leq x < 1/2 \\ -1/4 & \text{if } -1/2 \leq x < 0 \\ -3/4 & \text{if } -1 \leq x < -1/2 \end{cases}$$

deadzone quantizer

$$Q_d(x) = \begin{cases} 2/3 & \text{if } 1/3 \leq x \leq 1 \\ 0 & \text{if } |x| < 1/3 \\ -2/3 & \text{if } -1 \leq x \leq -1/3 \end{cases}$$

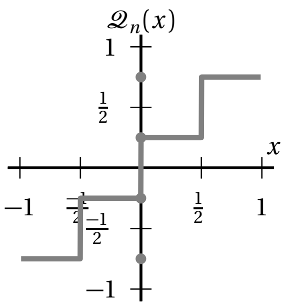

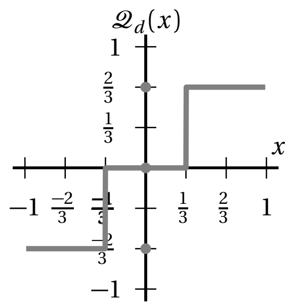

Both quantizers operate at two bits per sample but the deadzone quantizer "wastes" a fraction of a bit since it has only 3 quantization intervals instead of 4.

Assume that each sample $x[n]$ has the following probability distribution:

$$P[x[n]=\alpha] = \begin{cases} 0 & \text{if } |\alpha|>1 \\ p & \text{if } |\alpha|=0 \\ \dfrac{1-p}{2} & \text{otherwise} \end{cases}$$

In other words, each sample is either zero with probability p or drawn from a uniform distribution over the $[-1,1]$ interval; the probability density function for such a distribution can be expressed as:

$$f(x) = \frac{1-p}{2} + p\delta(x)$$

where the Dirac delta encodes the specific probability of a sample being equal to zero.

Determine the minimum value of p for which it is better to use the deadzone quantizer, i.e. the value of p for which the MSE of the deadzone quantizer becomes smaller than the MSE of the uniform quantizer.

Hint: remember that the formula for the MSE of a scalar quantizer over the $[-1,1]$ interval (under the hypotheses of iid samples with distribution $f(x)$) is

$$\sigma^2 = \int_{-1}^{1} [\mathcal{Q}(x)-x]^2 f(x)\,dx$$

For a uniform quantizer with M quantization levels (as the ones in this exercise) if the input distribution is uniform (i.e. $f(x)=1/2$) the above simplifies to:

$$\sigma^2 = \int_{-1}^{1} (Q(x)-x)^2 \frac{1}{2}\,dx = \frac{\Delta^2}{12} = \frac{(2/M)^2}{12} = \frac{1}{3M^2}$$

Solution:

The number of quantization levels in the two quantizers are $M = M_u = 4$ for the normal 2-bit quantizer and $M = M_d = 3$ for the deadzone quantizer. Let's compute the MSE for the normal quantizer using the composite pdf for the input

$$\begin{aligned}
\sigma_n^2 &= \int_{-1}^{1} (Q_u(x)-x)^2 \left(\frac{1-p}{2} + p\delta(x)\right) dx \\
&= (1-p)\int_{-1}^{1} (Q_u(x)-x)^2 \frac{1}{2}\,dx + p\int_{-1}^{1} (Q_u(x)-x)^2\,\delta(x)\,dx \\
&= (1-p)\frac{1}{3M_n^2} + p[Q_u(0)]^2 \\
&= (1-p)\frac{1}{48} + p\frac{1}{16}
\end{aligned}$$

where we have used the fact that the normal quantizer maps zero to $1/4$*; similarly, for the deadzone quantizer (which maps zero to zero):*

$$\sigma_d^2 = \int_{-1}^{1} (Q_d(x) - x)^2 \left(\frac{1-p}{2} + p\delta(x) \right) dx$$
$$= (1-p)\frac{1}{3M_d^2} + p[Q_d(0)]^2$$
$$= (1-p)\frac{1}{27}$$

from which we find

$$\sigma_d^2 < \sigma_n^2 \quad \textit{for} \quad p > \frac{21}{102} \approx 20\%$$