

ECE 161C: Introduction to Image Processing

Sampling and Acquisition of Images

Rahul Parhi
Assistant Professor, ECE, UCSD

Spring 2026

Outline

- Sampling Theory
 - Review 1D Sampling Theory
 - Sampling in Two Dimensions
- Acquisition Systems
 - Real Acquisition Systems
 - Aliasing Problems
- Image Quantization
 - Uniform Quantizer
 - Minimum-Error (Lloyd-Max) Quantizer
 - Grayscale vs. Spatial Resolution Tradeoff


Sampling Theory

- Review 1D Sampling Theory
- Sampling in Fourier Domain
- Sampling and Aliasing
- Shannon's Sampling Theorem
- Sampling in 2D

Review of 1D Sampling Theory

- Ideal sampling = multiplication with a **Dirac comb**

$$\mathbb{I}\mathbb{I}_T(x) = \sum_{k \in \mathbb{Z}} \delta(x - kT)$$




$$f_T(x) = f(x) \mathbb{I}\mathbb{I}_T(x)$$

$$= f(x) \sum_{k \in \mathbb{Z}} \delta(x - kT)$$

$$= \sum_{k \in \mathbb{Z}} f(x) \delta(x - kT)$$


$$= \sum_{k \in \mathbb{Z}} f(kT) \delta(x - kT)$$



Review of 1D Sampling Theory (cont'd)

- Ideal sampling = multiplication with a **Dirac comb**

$$f_T(x) = \sum_{k \in \mathbb{Z}} f(kT) \delta(x - kT)$$




- Sampling periodizes the Fourier domain

$$\begin{aligned} \hat{f}_T(\omega) &= \int_{-\infty}^{\infty} f_T(x) e^{-j\omega x} dx = \int_{-\infty}^{\infty} \sum_{k \in \mathbb{Z}} f(kT) \delta(x - kT) e^{-j\omega x} dx \\ &= \sum_{k \in \mathbb{Z}} f(kT) \int_{-\infty}^{\infty} \delta(x - kT) e^{-j\omega x} dx \\ &= \sum_{k \in \mathbb{Z}} f(kT) e^{-j\omega kT} \end{aligned}$$

This is a $(2\pi/T)$ -periodic function in ω

Dirac Comb Fourier Transform

$$\mathbb{I}\mathbb{I}_T(x) = \sum_{k \in \mathbb{Z}} \delta(x - kT) \quad \xleftrightarrow{\mathcal{F}} \quad ???$$



- T -periodic functions can be specified by their **Fourier series**


$$f_T(x) = \sum_{n \in \mathbb{Z}} c_n e^{jn\omega_0 x} \quad \text{with} \quad c_n = \frac{1}{T} \int_{-T/2}^{T/2} f_T(x) e^{-jn\omega_0 x} dx, \quad \omega_0 = \frac{2\pi}{T}$$

$$\begin{aligned} \mathbb{I}\mathbb{I}_T(x) &= \sum_{n \in \mathbb{Z}} \left(\frac{1}{T} \int_{-T/2}^{T/2} \mathbb{I}\mathbb{I}_T(x) e^{-jn\omega_0 x} dx \right) e^{jn\omega_0 x} \\ &= \sum_{n \in \mathbb{Z}} \left(\frac{1}{T} \int_{-T/2}^{T/2} \delta(x) e^{-jn\omega_0 x} dx \right) e^{jn\omega_0 x} \\ &= \sum_{n \in \mathbb{Z}} \frac{1}{T} e^{jn\omega_0 x} \quad \text{with} \quad \omega_0 = \frac{2\pi}{T} \end{aligned}$$

Dirac Comb Fourier Transform (cont'd)

$$\mathbb{I}\mathbb{I}_T(x) = \sum_{k \in \mathbb{Z}} \delta(x - kT) \quad \xleftrightarrow{\mathcal{F}} \quad ???$$

$$= \sum_{n \in \mathbb{Z}} \frac{1}{T} e^{jn\omega_0 x}$$



$$e^{jn\omega_0 x} \quad \xleftrightarrow{\mathcal{F}} \quad 2\pi\delta(\omega - n\omega_0) \quad (\text{by duality})$$

$$\mathbb{I}\mathbb{I}_T(x) = \sum_{n \in \mathbb{Z}} \frac{1}{T} e^{jn\omega_0 x} \quad \xleftrightarrow{\mathcal{F}} \quad \sum_{n \in \mathbb{Z}} \frac{1}{T} 2\pi\delta(\omega - n\omega_0)$$

$$= \frac{2\pi}{T} \sum_{n \in \mathbb{Z}} \delta\left(\omega - \frac{2\pi}{T}n\right)$$

The Fourier transform of a Dirac comb is a Dirac comb

Sampling in the Fourier Domain


- Sampling formula

$$f_T(x) = f(x) \cdot \sum_{k \in \mathbb{Z}} \delta(x - kT) \xleftrightarrow{\mathcal{F}} \frac{1}{2\pi} \left(\hat{f}(\omega) * \frac{2\pi}{T} \sum_{n \in \mathbb{Z}} \delta\left(\omega - \frac{2\pi n}{T}\right) \right)$$


$$\hat{f}_T(\omega) = \frac{1}{T} \sum_{n \in \mathbb{Z}} \hat{f}\left(\omega - \frac{2\pi n}{T}\right) = \sum_{k \in \mathbb{Z}} f(kT) e^{-j\omega kT}$$

- Sampling **periodizes** the spectrum

$\hat{f}(\omega)$: Fourier transform of $f(x)$
(continuous-domain function)



$\hat{f}_T(\omega)$: $(2\pi/T)$ -periodization of $\hat{f}(x)$
(continuous-domain function)



Sampling in the Fourier Domain

$$\hat{f}_T(\omega) = \frac{1}{T} \sum_{n \in \mathbb{Z}} \hat{f}\left(\omega - \frac{2\pi n}{T}\right) = \sum_{k \in \mathbb{Z}} f(kT) e^{-j\omega kT}$$


- Normalized sampling step ($T = 1$)

$$\hat{f}_{T=1}(\omega) = \sum_{n \in \mathbb{Z}} \hat{f}(\omega - 2\pi n) = \sum_{k \in \mathbb{Z}} f(k) e^{-j\omega k}$$

Poisson
summation
formula


- Define the discrete sequence $f[k] = f(k)$, $k \in \mathbb{Z}$

$\hat{f}(\omega)$: Fourier transform of $f(x)$
(continuous-domain function)



$\hat{f}_{T=1}(\omega)$: Discrete~~-time~~ ^{space} Fourier
transform of $f[k]$


$$F(e^{j\omega}) = \sum_{k \in \mathbb{Z}} f[k] e^{-j\omega k}$$





Sampling and Aliasing

- Sampled signal


$$f_T(x) = \sum_{k \in \mathbb{Z}} f(kT) \delta(x - kT) \quad \xleftrightarrow{\mathcal{F}} \quad \hat{f}_T(\omega) = \frac{1}{T} \sum_{n \in \mathbb{Z}} \hat{f}\left(\omega - \frac{2\pi n}{T}\right)$$



$$\omega_0 = \frac{2\pi}{T}$$




$$\omega'_0 = \frac{2\pi}{T'}$$



Undersampling causes **aliasing** which destroys the spectrum


Sampling and Reconstruction

- Continuous-domain input signal $f(x)$




- Sampling

$$\hat{f}_T(\omega) = \frac{1}{T} \sum_{n \in \mathbb{Z}} \hat{f}\left(\omega - \frac{2\pi n}{T}\right)$$



- Ideal low-pass filtering

$$\hat{f}_{\text{LP}}(\omega) = \hat{f}_T(\omega) T \operatorname{rect}\left(\frac{\omega T}{2\pi}\right)$$



$$= \begin{cases} \sum_{n \in \mathbb{Z}} \hat{f}\left(\omega - \frac{2\pi n}{T}\right), & |\omega| < \frac{\pi}{T} \\ 0, & \text{else} \end{cases}$$

When does $\hat{f}_{\text{LP}} = \hat{f}$?


When do we have perfect recovery?

Nyquist–Shannon Sampling Theorem


(Whittaker–Nyquist–Kotelnikov–Shannon Sampling Theorem)

- Ideal Reconstruction Process

Exercise:




$$\begin{aligned} \text{sinc}(x) &\xleftrightarrow{\mathcal{F}} \text{rect}\left(\frac{\omega}{2\pi}\right) \\ f\left(\frac{x}{T}\right) &\xleftrightarrow{\mathcal{F}} T \hat{f}(T\omega) \\ \text{sinc}\left(\frac{x}{T}\right) &\xleftrightarrow{\mathcal{F}} T \text{rect}\left(\frac{T\omega}{2\pi}\right) \end{aligned}$$



Nyquist–Shannon Sampling Theorem

(Whittaker–Nyquist–Kotelnikov–Shannon Sampling Theorem)

- Ideal Reconstruction Process



Sampling Theorem



A function $f(x)$ that is bandlimited to ω_{max} can be reconstructed exactly from its equidistant samples provided that the sampling step $T < \pi/\omega_{\text{max}}$. Specifically,

$$f(x) = \sum_{k \in \mathbb{Z}} f(kT) \text{sinc}\left(\frac{x - kT}{T}\right)$$

$$\text{where } \text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \xleftrightarrow{\mathcal{F}} \text{rect}\left(\frac{\omega}{2\pi}\right).$$

Sampling in 2D

How do we sample in 2D?



- 2D Dirac comb



$$\sum_{(k,l) \in \mathbb{Z}^2} \delta(x - kT_1, y - lT_2) \quad \xleftrightarrow{\mathcal{F}} \quad \frac{(2\pi)^2}{T_1 T_2} \sum_{(m,n) \in \mathbb{Z}^2} \delta\left(\omega_1 - \frac{2\pi m}{T_1}, \omega_2 - \frac{2\pi n}{T_2}\right)$$

- 2D sampling formula

$$f_{T_1, T_2}(x, y) = f(x, y) \sum_{(k,l) \in \mathbb{Z}^2} \delta(x - kT_1, y - lT_2) \\ \xleftrightarrow{\mathcal{F}} \quad \hat{f}_{T_1, T_2}(\omega_1, \omega_2) = \frac{1}{T_1 T_2} \sum_{(m,n) \in \mathbb{Z}^2} \hat{f}\left(\omega_1 - \frac{2\pi m}{T_1}, \omega_2 - \frac{2\pi n}{T_2}\right)$$

Sampling and Spectral Repetition

$$\hat{f}_{T_1, T_2}(\omega_1, \omega_2) = \frac{1}{T_1 T_2} \sum_{(m, n) \in \mathbb{Z}^2} \hat{f}\left(\omega_1 - \frac{2\pi m}{T_1}, \omega_2 - \frac{2\pi n}{T_2}\right)$$



Sampling in 2D

(For notational simplicity, let $T_1 = T_2$, i.e., the sampling grid is **regular**)

- Sampling function in two dimensions

$$\mathbb{I}\mathbb{I}_T(\boldsymbol{x}) = \sum_{\boldsymbol{k} \in \mathbb{Z}^2} \delta(\boldsymbol{x} - \boldsymbol{k}T) \quad \xleftrightarrow{\mathcal{F}} \quad \widehat{\mathbb{I}\mathbb{I}}_T(\boldsymbol{\omega}) = \frac{(2\pi)^2}{T^2} \sum_{\boldsymbol{n} \in \mathbb{Z}^2} \delta\left(\boldsymbol{\omega} - \frac{2\pi\boldsymbol{n}}{T}\right)$$

- Two-dimensional sampling formula

$$f_T(\boldsymbol{x}) = f(\boldsymbol{x}) \mathbb{I}\mathbb{I}_T(\boldsymbol{x}) \quad \xleftrightarrow{\mathcal{F}} \quad \hat{f}_T(\boldsymbol{x}) = \frac{1}{(2\pi)^2} (\hat{f} * \widehat{\mathbb{I}\mathbb{I}}_T)(\boldsymbol{\omega})$$

$$\Rightarrow \hat{f}_T(\boldsymbol{x}) = \frac{1}{T^2} \sum_{\boldsymbol{n} \in \mathbb{Z}^2} \hat{f}\left(\boldsymbol{\omega} - \frac{2\pi\boldsymbol{n}}{T}\right)$$

- Condition for perfect recovery (two-dimensional sampling theorem)


$$\omega_{\max} < \frac{\pi}{T}$$

$$f(\boldsymbol{x}) = \sum_{\boldsymbol{k} \in \mathbb{Z}^2} f(\boldsymbol{k}T) \operatorname{sinc}\left(\frac{\boldsymbol{x} - \boldsymbol{k}T}{T}\right)$$

$$\text{with } \operatorname{sinc}(\boldsymbol{x}) = \operatorname{sinc}(x_1)\operatorname{sinc}(x_2)$$

Sampling Lattices


- Cartesian lattice: $\mathbb{Z}^2$



$$t_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad t_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Are there other options for sampling lattices?

- General lattice



A lattice is completely determined by two vectors: **lattice vectors**

What are the lattice vectors of the Cartesian grid?

Sampling Lattices (cont'd)

- Lattice matrix: $\mathbf{T} = [\mathbf{t}_1 \quad \mathbf{t}_2] \in \mathbb{R}^{2 \times 2}$, with lattice vectors $\mathbf{t}_1, \mathbf{t}_2$


How do we use this matrix to define a lattice?

What is the lattice matrix of the Cartesian lattice?

Cartesian lattice $\implies \mathbf{T} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \mathbf{I}$ (identity matrix)

Cartesian lattice $\implies \mathbb{Z}^2$

- A general lattice is of the form $\mathbf{T}\mathbb{Z}^2 = \left\{ \mathbf{T} \begin{bmatrix} m \\ n \end{bmatrix} : \begin{bmatrix} m \\ n \end{bmatrix} \in \mathbb{Z}^2 \right\}$




Sampling Lattices (cont'd)

What would be the Dirac comb for a general lattice?

- Dirac comb for $\mathbb{Z}^2$

$$\sum_{\mathbf{k} \in \mathbb{Z}^2} \delta(\mathbf{x} - \mathbf{k}) \quad \xleftrightarrow{\mathcal{F}} \quad (2\pi)^2 \sum_{\mathbf{n} \in \mathbb{Z}^2} \delta(\boldsymbol{\omega} - 2\pi\mathbf{n})$$



Affine transformation: $f(\mathbf{T}\mathbf{x}) \xleftrightarrow{\mathcal{F}} \frac{1}{|\det(\mathbf{T})|} \hat{f}(\mathbf{T}^{-\top}\boldsymbol{\omega})$

- Dirac comb for $\mathbf{T}\mathbb{Z}^2$

$$\sum_{\mathbf{k} \in \mathbb{Z}^2} \delta(\mathbf{x} - \mathbf{T}\mathbf{k}) \quad \xleftrightarrow{\mathcal{F}} \quad \frac{(2\pi)^2}{|\det(\mathbf{T})|} \sum_{\mathbf{n} \in \mathbb{Z}^2} \delta(\boldsymbol{\omega} - 2\pi\mathbf{T}^{-\top}\mathbf{n})$$

$\mathbf{T}\mathbb{Z}^2$

$2\pi\mathbf{T}^{-\top}\mathbb{Z}^2$


“dual” or “reciprocal”
lattice

Sampling Lattices (cont'd)

- Adapted Poisson summation formula

$$\sum_{\mathbf{k} \in \mathbb{Z}^2} f(\mathbf{T}\mathbf{k}) e^{-j\boldsymbol{\omega}^\top \mathbf{T}\mathbf{x}} = \frac{1}{|\det(\mathbf{T})|} \sum_{\mathbf{n} \in \mathbb{Z}^2} \hat{f}(\boldsymbol{\omega} - 2\pi \mathbf{T}^{-\top} \mathbf{n})$$



What is a “good” choice of 2D lattice?




2D “sphere packing” corresponds to what lattice?

Hexagonal lattice

Nature's Favorite Lattice




Hexagonal Lattice



Exercise: Determine the lattice matrix.

$$\mathbf{T} = \begin{bmatrix} 1 & 1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}$$



Exercise: Determine the reciprocal lattice.

$$2\pi\mathbf{T}^{-\top} = 2\pi \begin{bmatrix} 1 & 0 \\ -\sqrt{3}/3 & 2\sqrt{3}/3 \end{bmatrix}$$

The reciprocal lattice of a hexagonal lattice is a hexagonal lattice

Acquisition Systems

- Real Acquisition Systems
- Acquisition Models
- Aliasing Problems

Real Acquisition Systems

In practice, sampling is **not ideal**

- Pixel measurement process

$$p[\mathbf{k}] = \int_{\mathbb{R}^2} \varphi_a(\mathbf{y} - \mathbf{k}T) f(\mathbf{y}) d\mathbf{y}$$



$p[\mathbf{k}]$: pixel value at location $\mathbf{k} = (k_1, k_2) \in \mathbb{Z}^2$

φ_a : sampling aperture (or integration window)


Typically: $\int_{\mathbb{R}^2} \varphi_a(\mathbf{x}) d\mathbf{x} = 1$ (normalized intensity)

- Example: CCD camera

$$\varphi_a(x, y) = \frac{1}{T^2} \text{rect}\left(\frac{x}{T}\right) \text{rect}\left(\frac{y}{T}\right)$$



Equivalent Pixel Measurement Model




Exercise: Show that this system implements the pixel measurement process.

$$p[\mathbf{k}] = \int_{\mathbb{R}^2} \varphi_a(\mathbf{y} - \mathbf{k}T) f(\mathbf{y}) \, d\mathbf{y}$$

How does this system recover ideal sampling?

$$\varphi_a(\mathbf{x}) = \delta(\mathbf{x})$$

Even More Real Acquisition Systems




Equivalent impulse response: $h(\mathbf{x})$

- Complete image-acquisition model:

$$p[\mathbf{k}] = (h * f)(\mathbf{k}T) + n[\mathbf{k}]$$



- $\varphi_0(\mathbf{x})$: point-spread function (image formation)
- $\varphi_a(\mathbf{x})$: sampling aperture (sensors)
- $h(\mathbf{x}) = (\varphi_a^\vee * \varphi_0)(\mathbf{x})$: equivalent LSI system ($\varphi_a^\vee(\mathbf{x}) = \varphi_a(-\mathbf{x})$)
- $n[\mathbf{k}]$: additive measurement noise



[source: dvinfo.net]

Aliasing Problems

- Aliasing = Spectral overlap induced by sampling
 - Alias-free condition: f must be bandlimited with $\omega_{\max} < \frac{\pi}{T}$



- Practical solutions?
 - Adapt the sampling step to the frequency content
 - Low-pass filtering prior to sampling (implicitly or explicitly)
- via φ_0 : imaging system;
 φ_a : sampling aperture

Example of Aliased Image



original “analog” image
(high-resolution digital image)

What will aliasing look like?

How do we simulate sampling?




sampled image
(2×2 downsampling)




prefiltered then sampled image

More “Bad” Aliasing



Especially a problem in printing applications

Rotogravure/Intaglio Printing



This is how banknotes, passports, and stamps are printed today


Aliasing Can Be “Good”




Outline

- Sampling Theory ✓
 - Review 1D Sampling Theory
 - Sampling in Two Dimensions
- Acquisition Systems ✓
 - Real Acquisition Systems
 - Aliasing Problems
- Image Quantization
 - Uniform Quantizer
 - Minimum-Error (Lloyd-Max) Quantizer
 - Grayscale vs. Spatial Resolution Tradeoff

Aliasing = Ambiguity in Analog Signal



Effect of Reducing Spatial Resolution



There is a trade-off between number of gray levels and resolution

Quantization



Image Quantization

- Quantizer Specification
- Histogram
- Uniform Quantization
- Minimum-Error (Lloyd-Max) Quantizer
- Grayscale vs. Spatial Resolution Tradeoff
- Dithering

Quantizer Specification

What even is a quantizer?

- Images have real-valued intensity values $f = f(\mathbf{x}) \in \mathbb{R}$



- Quantization thresholds: $t_k \quad k = 0, \dots, K$
- Quantized output: $r_k \quad k = 0, \dots, K - 1$

$$\tilde{f} = Q(f) = r_k \iff f \in [t_k, t_{k+1})$$

Exercise: Come up with a 256 gray-level quantizer for images with grayscale intensity values $f \in [0, 1]$.

Histograms: A Probabilistic Viewpoint


How do we know how to quantize?

Look at the **distribution** of gray levels

- Gray-level probability density function (p.d.f.)

$$p(f) = \lim_{\Delta \rightarrow 0} \left\{ \frac{1}{\Delta} \frac{\# \text{ of pixels with gray level } \in [f, f + \Delta)}{\# \text{ of total pixels}} \right\} \geq 0$$

- Normalization: $\int_{-\infty}^{\infty} p(f) \, df = 1$
- Mean: $\mu = \mathbf{E}[F] = \int_{-\infty}^{\infty} f p(f) \, df, \quad F \sim p(f)$
- Variance: $\sigma^2 = \mathbf{E}[(F - \mu)^2] = \int_{-\infty}^{\infty} (f - \mu)^2 p(f) \, df, \quad F \sim p(f)$





What is the probability that a gray level is in $[t_k, t_{k+1})$?

Histograms: A Probabilistic Viewpoint (cont'd)

Given a quantizer Q , what is the corresponding **quantized histogram**?

If $F \sim p(f)$, what is the probability distribution of $\tilde{F} = Q(F)$?



equivalent to a probability mass function (p.m.f.)

- Quantized histogram


$$p(f) \rightarrow p_Q(\tilde{f}) = \sum_{k=0}^{K-1} P_k \delta(\tilde{f} - r_k)$$

$$\text{where } P_k = \mathbf{P}(f \in \text{bin}_k) = \int_{t_k}^{t_{k+1}} p(f) \, df$$

How can we measure the performance of a quantizer?

Mean-squared error:
 $\mathbf{E}[(F - \tilde{F})^2], \quad F \sim p(f)$

Histograms: A Probabilistic Viewpoint (cont'd)



MSE Analysis of the Uniform Quantizer

- Setup

$$r_k = k \Delta + r_0$$

$$t_k = \frac{r_k + r_{k-1}}{2}$$

Typically:

0-255 (256 gray levels)

0-1 (binary)

Pixel budget:

8 bits

1 bit


Human visual system can only distinguish about 60 gray levels (allegedly)

Exercise: Estimate the **quantization error** $\mathbf{E}[(F - \tilde{F})^2]$, $F \sim p(f)$
(You may assume that K is large)

$$\mathbf{E}[(F - \tilde{F})^2] = \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} (f - r_k)^2 p(f) \, df \approx \sum_{k=0}^{K-1} \int_{-\Delta/2}^{\Delta/2} e^2 \frac{P_k}{\Delta} \, de = \frac{\Delta^2}{12}$$

large K hypothesis
(high gray-level resolution)

Example of Uniform Quantization




Can we do something **better**?

Nonuniform
quantization?

Minimize the
MSE directly?

Nonuniform Quantization and Segmentation

Search for the “optimal” threshold values to segment images



$K = 2$

$K = 4$

Minimum mean squared error solutions:

For a given K , find the MMSE thresholds = Lloyd-Max quantizer

MMSE/Lloyd-Max Quantization

Goal: For a fixed K , minimize

$$\varepsilon^2 = \mathbf{E}[(F - \tilde{F})^2]$$

When $p(f)$ is uniform, the uniform quantizer is optimal

$$= \int_{t_0}^{t_K} (f - \tilde{f})^2 p(f) \, df = \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} (f - r_k)^2 p(f) \, df$$

over t_k and r_k .

How do we do this?

Take the partial derivatives w.r.t. t_k and r_k , set them equal to 0, and solve

Hint: $\int_a^b g(x) \, dx = \int_{-\infty}^b g(x) \, dx - \int_{-\infty}^a g(x) \, dx = G(b) - G(a) \quad \Rightarrow \quad \frac{\partial}{\partial a} \int_a^b g(x) \, dx = -g(a)$

$$\bullet \quad \frac{\partial \varepsilon^2}{\partial t_k} = 0 \quad \Rightarrow \quad t_k = \frac{r_k + r_{k-1}}{2}$$

(midpoint solution is optimal)

$$\bullet \quad \frac{\partial \varepsilon^2}{\partial r_k} = 0 \quad \Rightarrow \quad r_k = \frac{\int_{t_k}^{t_{k+1}} f p(f) \, df}{\int_{t_k}^{t_{k+1}} p(f) \, df} = \mathbf{E}[F \mid F \in [t_k, t_{k+1})]$$


(conditional mean is optimal)

MMSE/Lloyd-Max Quantization Algorithm

How do we realize this with an algorithm?

Does something seem funny?

Solution: Iterate back and forth between the t_k and the r_k



- Efficient implementation


- r_k can be computed directly from the histogram
- Recursive update

- Generalizations


- Non-quadratic cost functions
- Multivariate extensions (e.g., RGB)
- This is just the K -means algorithm

Lloyd-Max Quantization

uniform quantization




Lloyd-Max quantization



- Other useful properties:

- Unbiased estimate: $\mathbf{E}[\tilde{F}] = \mathbf{E}[F]$
- Error $E = F - \tilde{F}$ is **orthogonal** to the quantized value: $\mathbf{E}[\tilde{F}E] = 0$

Grayscale vs. Spatial Resolution Tradeoff







These all use the same number of bits!

Need for Binary Images

- Some devices can only render binary output (e.g., printers, fax machines, etc.)
 - ink or no ink!
- Lucky coincidence: The human visual system locally integrates black-and-white information and sees the “average”
- Exploit tradeoff between spatial resolution and grayscale resolution
- Implemented by “Raster Image Processors” (RIPs) in printing systems

Can we do better than just thresholding?

Dithering



Dithering Example



grayscale image



fixed threshold



random dithering



ordered dithering



Summary

- Ideal sampling is modeled as a multiplication with a sequence of Dirac impulses (ideal sampling function). In the frequency domain, this corresponds to a convolution with a sequence of Diracs with a reciprocal spacing.
- Sampling periodizes the Fourier transform of the image.
- The periodization pattern can be predicted from the Fourier transform of the ideal sampling function (also a sequence of Dirac impulses).
- Perfect recovery is possible only if the image is sampled at or above the Nyquist frequency $\omega_{\max}/2$.
- Undersampling produces aliasing. It can be prevented by ideal lowpass prefiltering prior to sampling. True acquisition systems include a sampling aperture which acts as a lowpass prefilter, thereby reducing aliasing.
- During acquisition, the intensity values of the individual pixels are quantized (A-to-D conversion).
- Uniform quantization is the most common. Monochrome monitors typically display 256 gray levels.
- Alternatively, the quantization steps may be selected to minimize the mean squared error (Max-Lloyd quantizer). This also yields an effective segmentation algorithm.
- To some extent, grayscale resolution can be traded for spatial resolution. This tradeoff can be exploited via dithering.