ECE 172A: Introduction to Image Processing
Sampling and Acquisition of Images: Part II

Rahul Parhi
Assistant Professor, ECE, UCSD

Winter 2025

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

Effect of Reducing Spatial Resolution

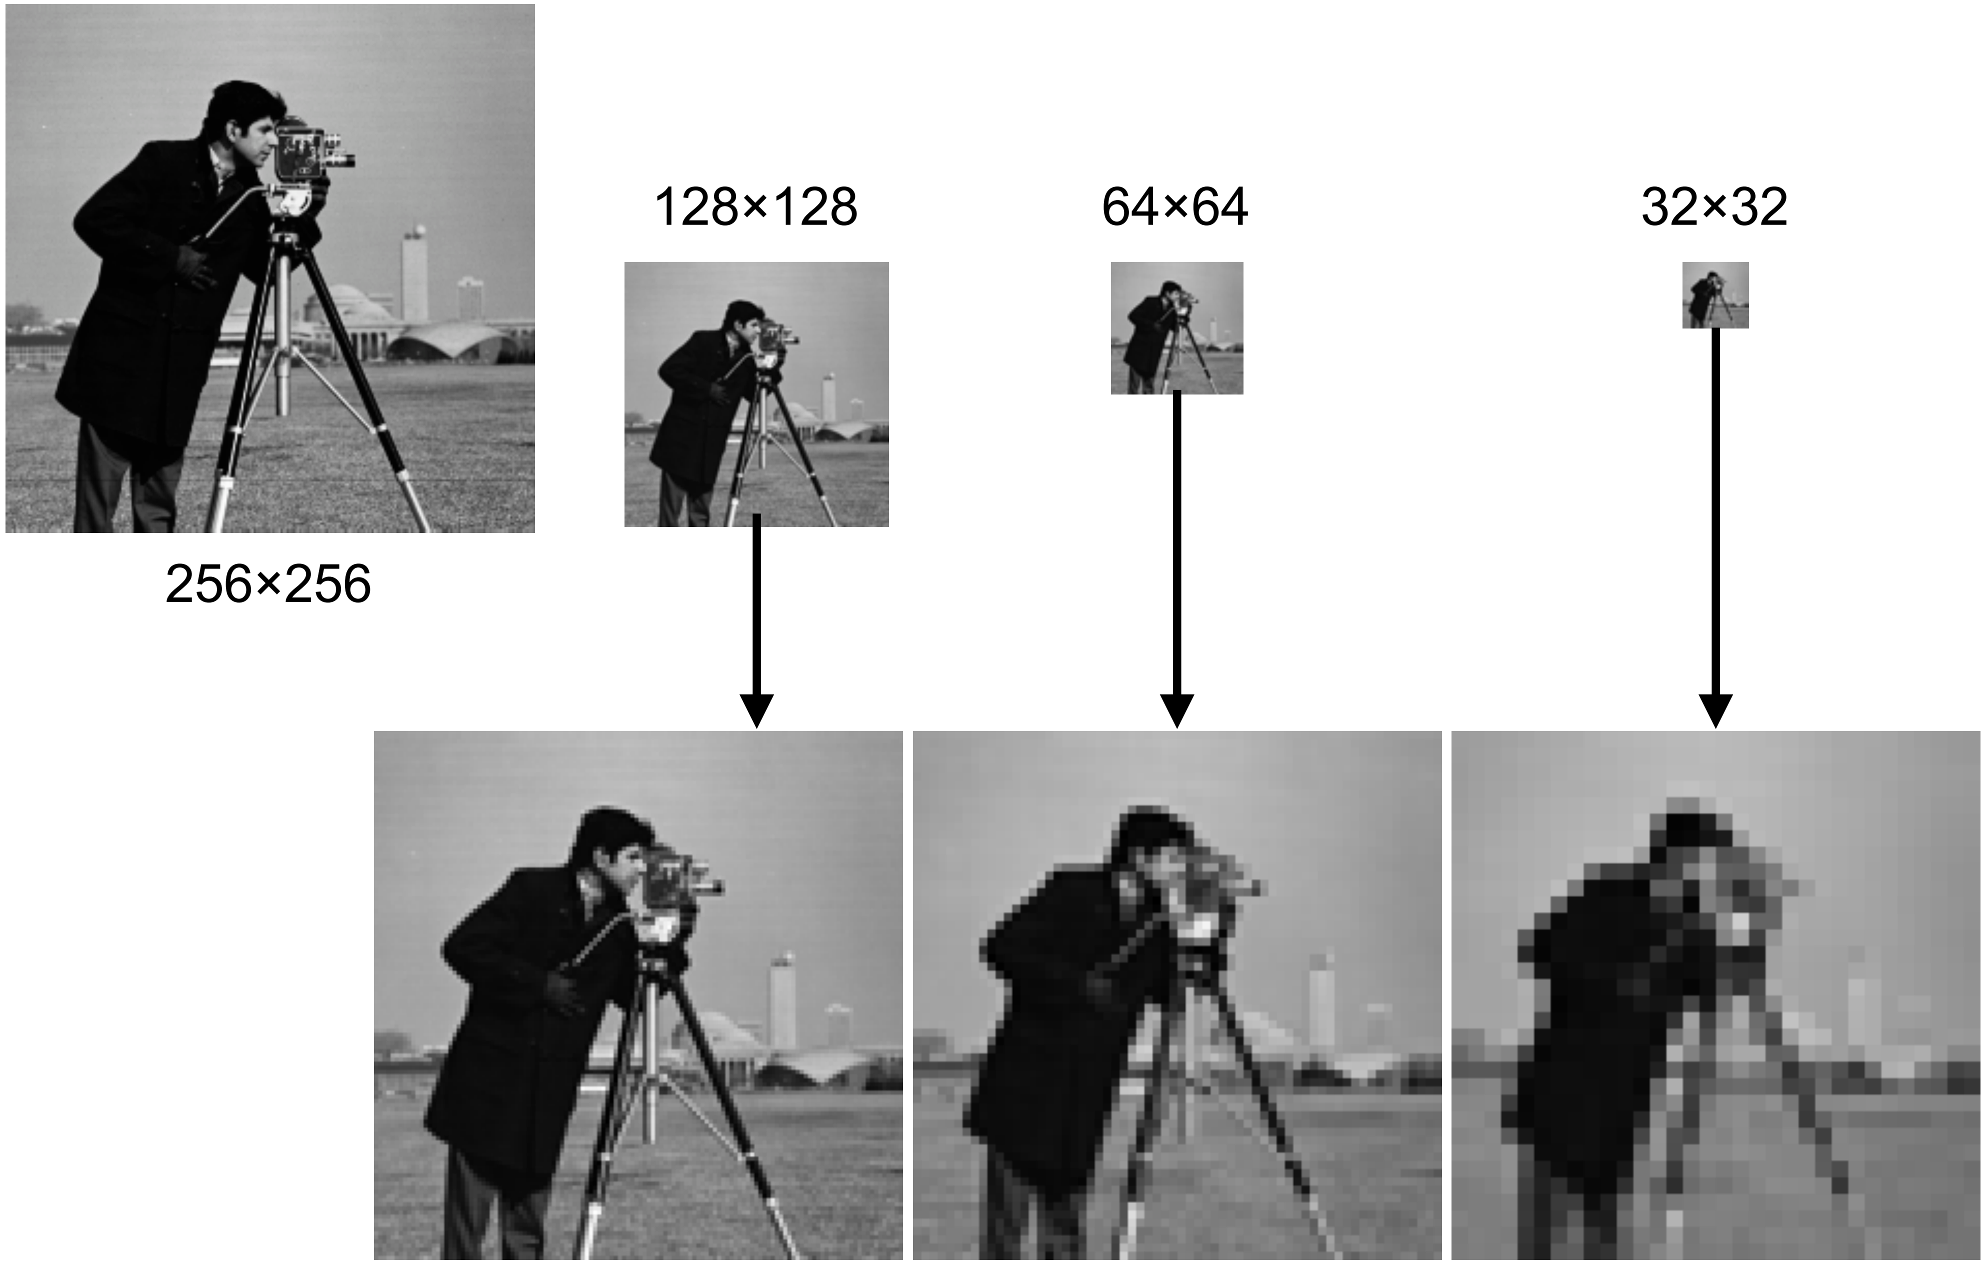

There is a trade-off between number of gray levels and resolution

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 $f = f(\boldsymbol{x}) \in \mathbb{R}$

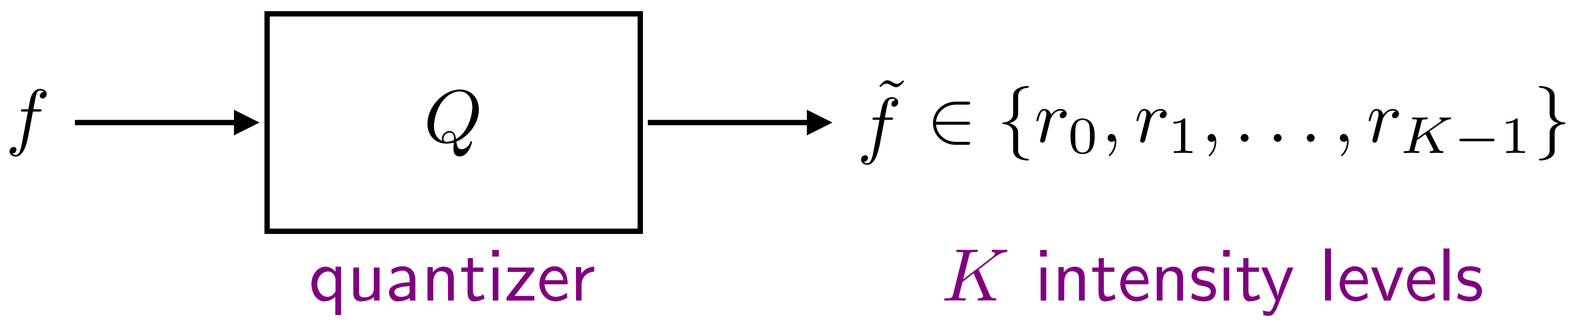

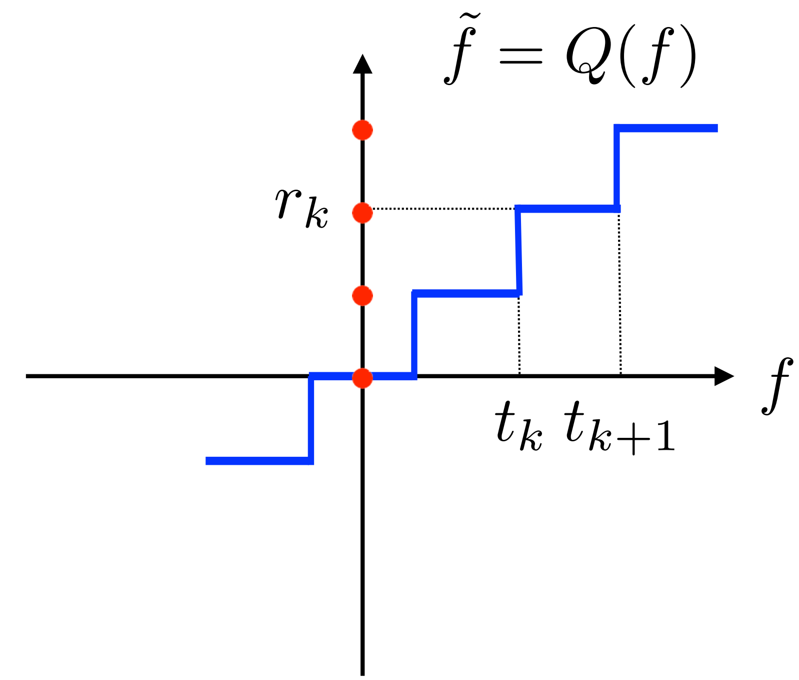

- Quantization thresholds: $t_k \qquad k = 0, \ldots, K$
- Quantized output: $r_k \qquad k = 0, \ldots, K-1$

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

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

Histograms: A Probabilistic Viewpoint

Look at the **distribution** of gray levels

How do we know how to quantize?

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

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

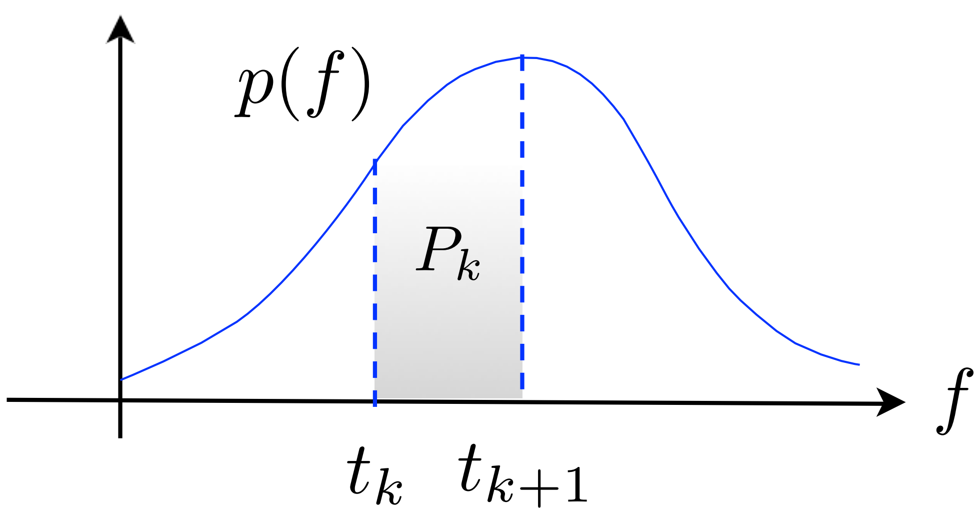

- Normalization: $\int_{-\infty}^{\infty} p(f)\,\mathrm{d}f = 1$
- Mean: $\mu = \mathbf{E}[F] = \int_{-\infty}^{\infty} f\,p(f)\,\mathrm{d}f, \quad F \sim p(f)$
- Variance: $\sigma^2 = \mathbf{E}[(F-\mu)^2] = \int_{-\infty}^{\infty} (f-\mu)^2\,p(f)\,\mathrm{d}f, \quad F \sim p(f)$

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 $\tilde{F} = Q(F)$?

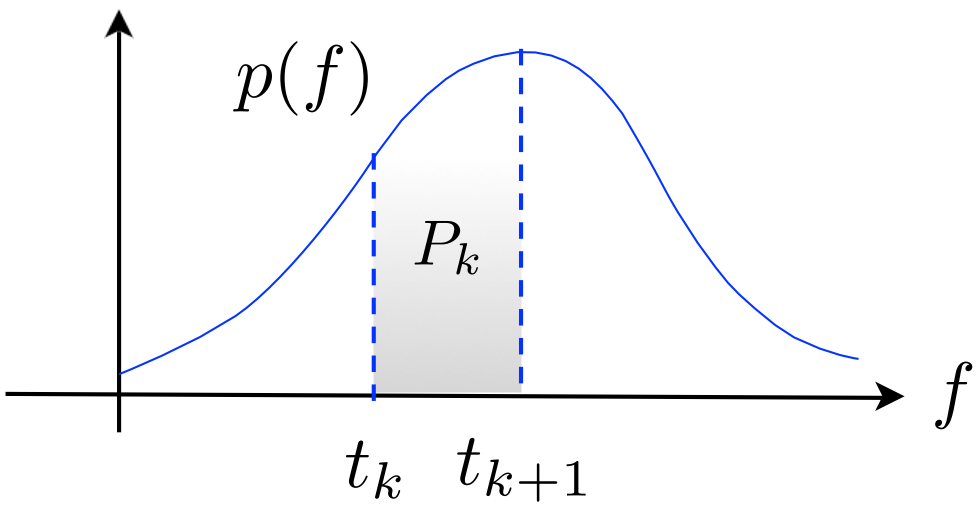

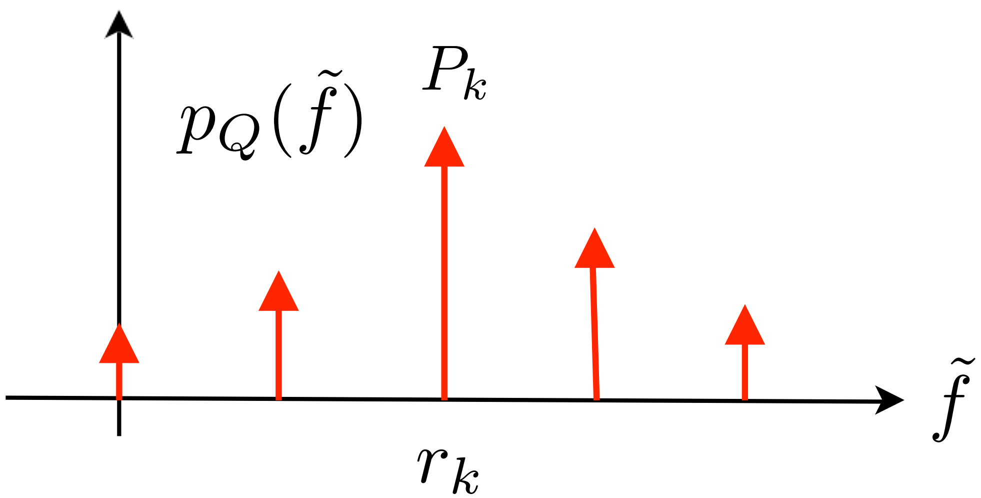

- Quantized histogram

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

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

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

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)

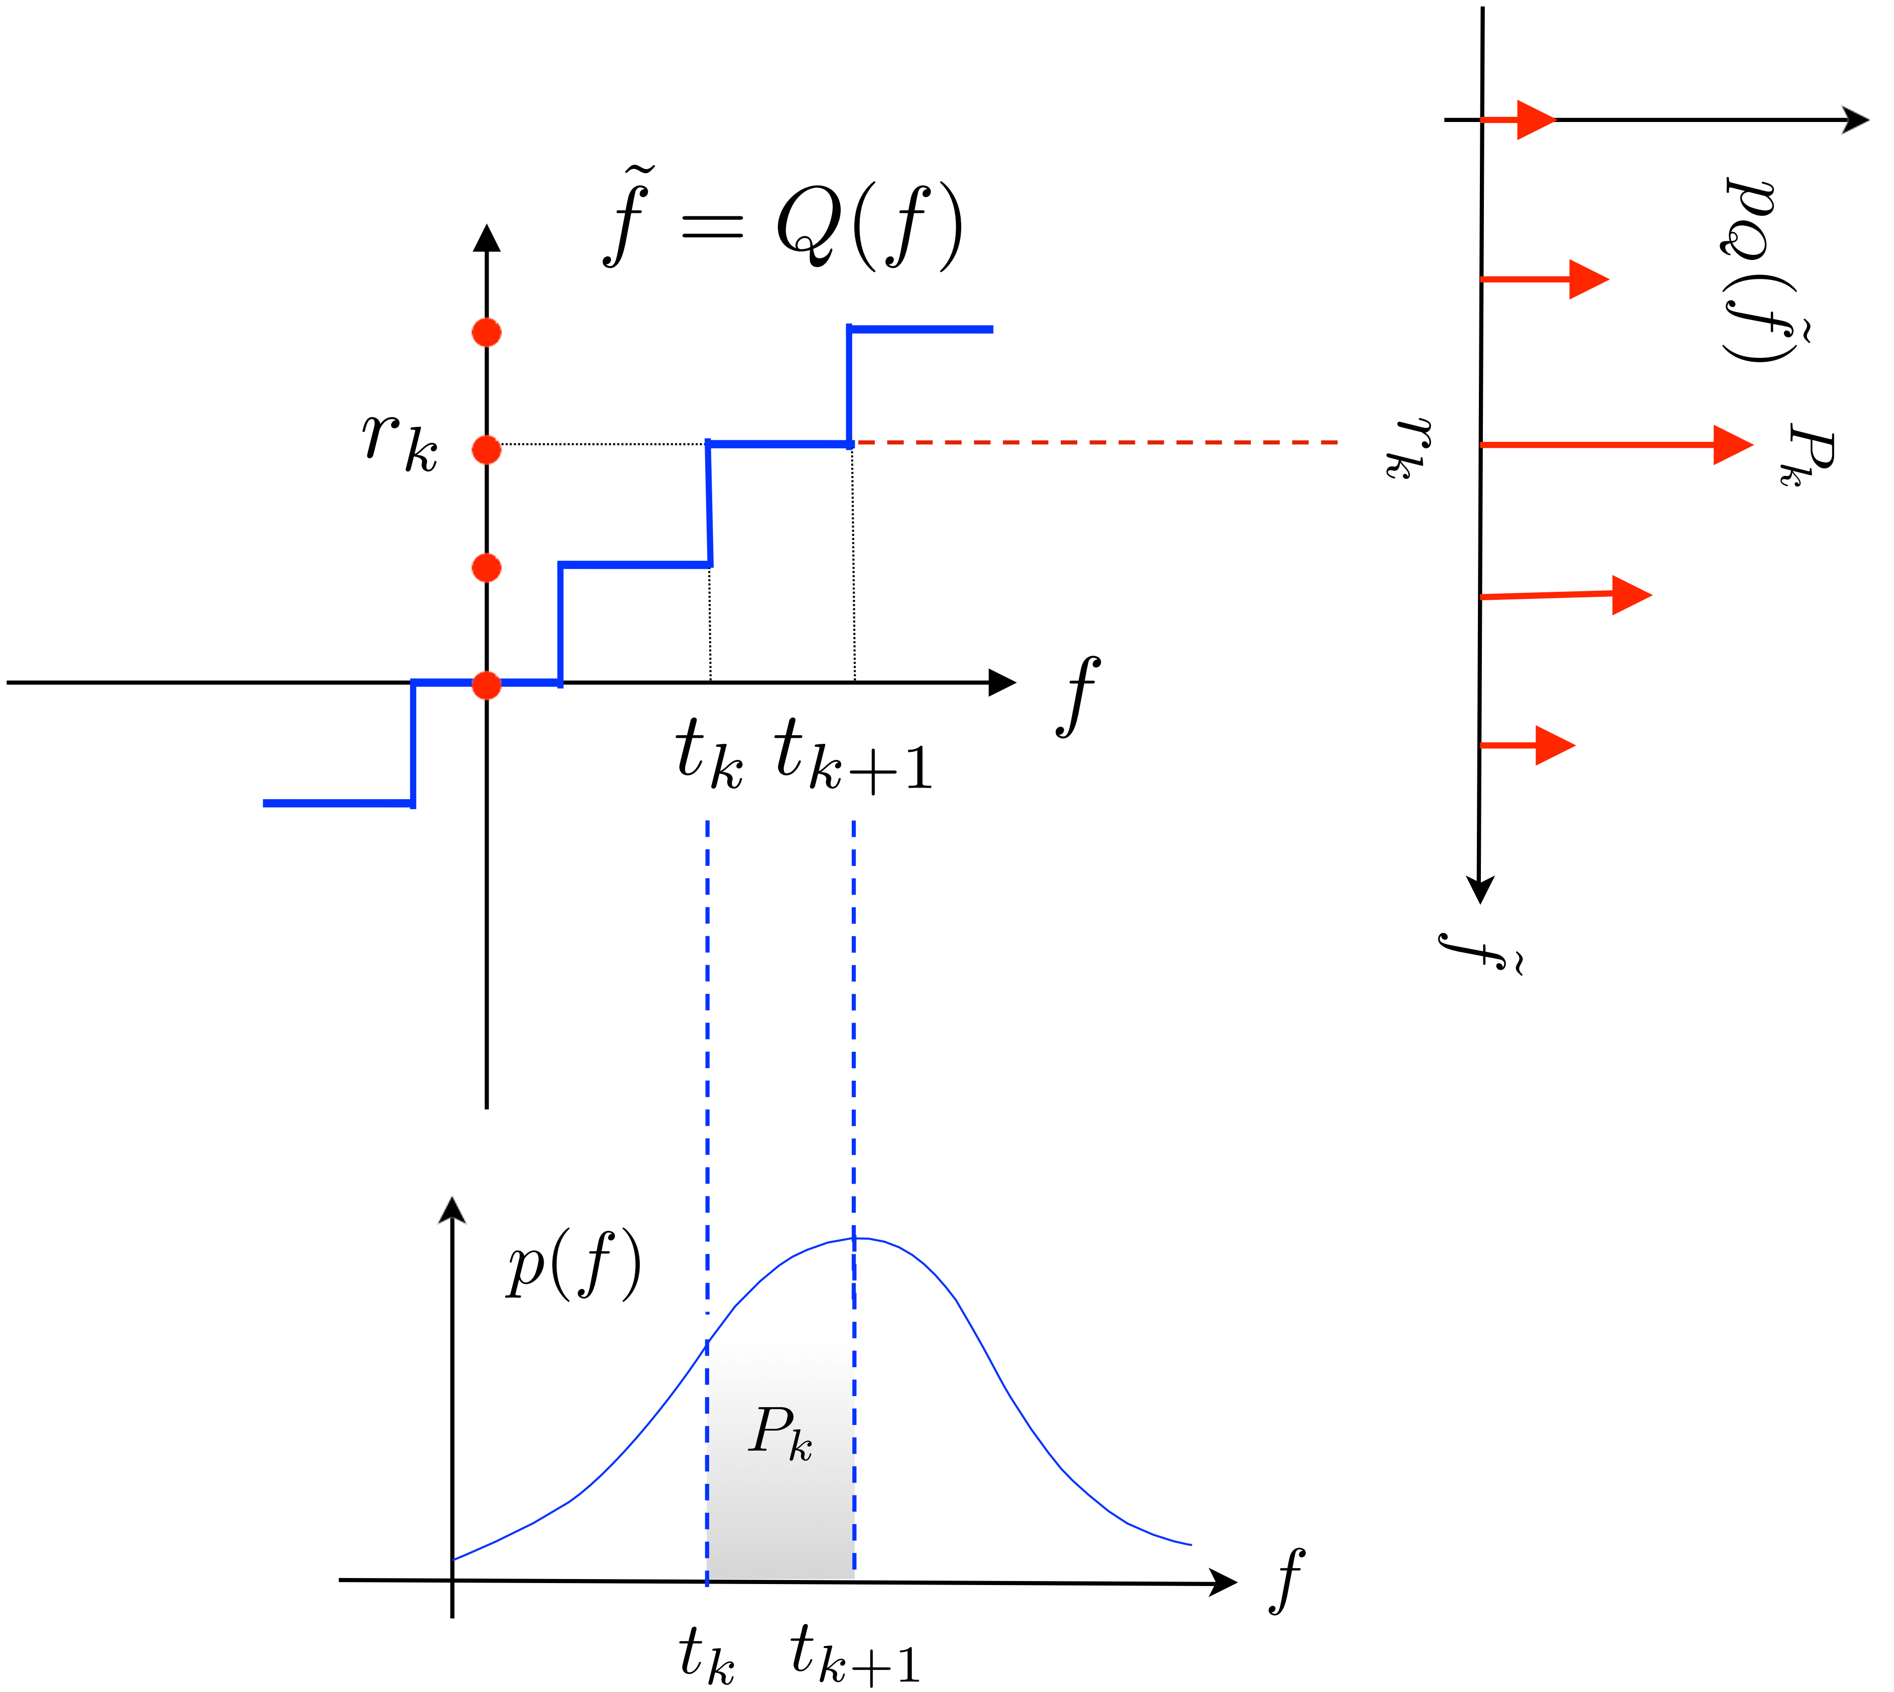

MSE Analysis of the Uniform Quantizer

- Setup

$$r_k = k\,\Delta + r_0 \qquad\qquad t_k = \frac{r_k + r_{k-1}}{2}$$

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)

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

large K hypothesis
(high gray-level resolution)

Example of Uniform Quantization

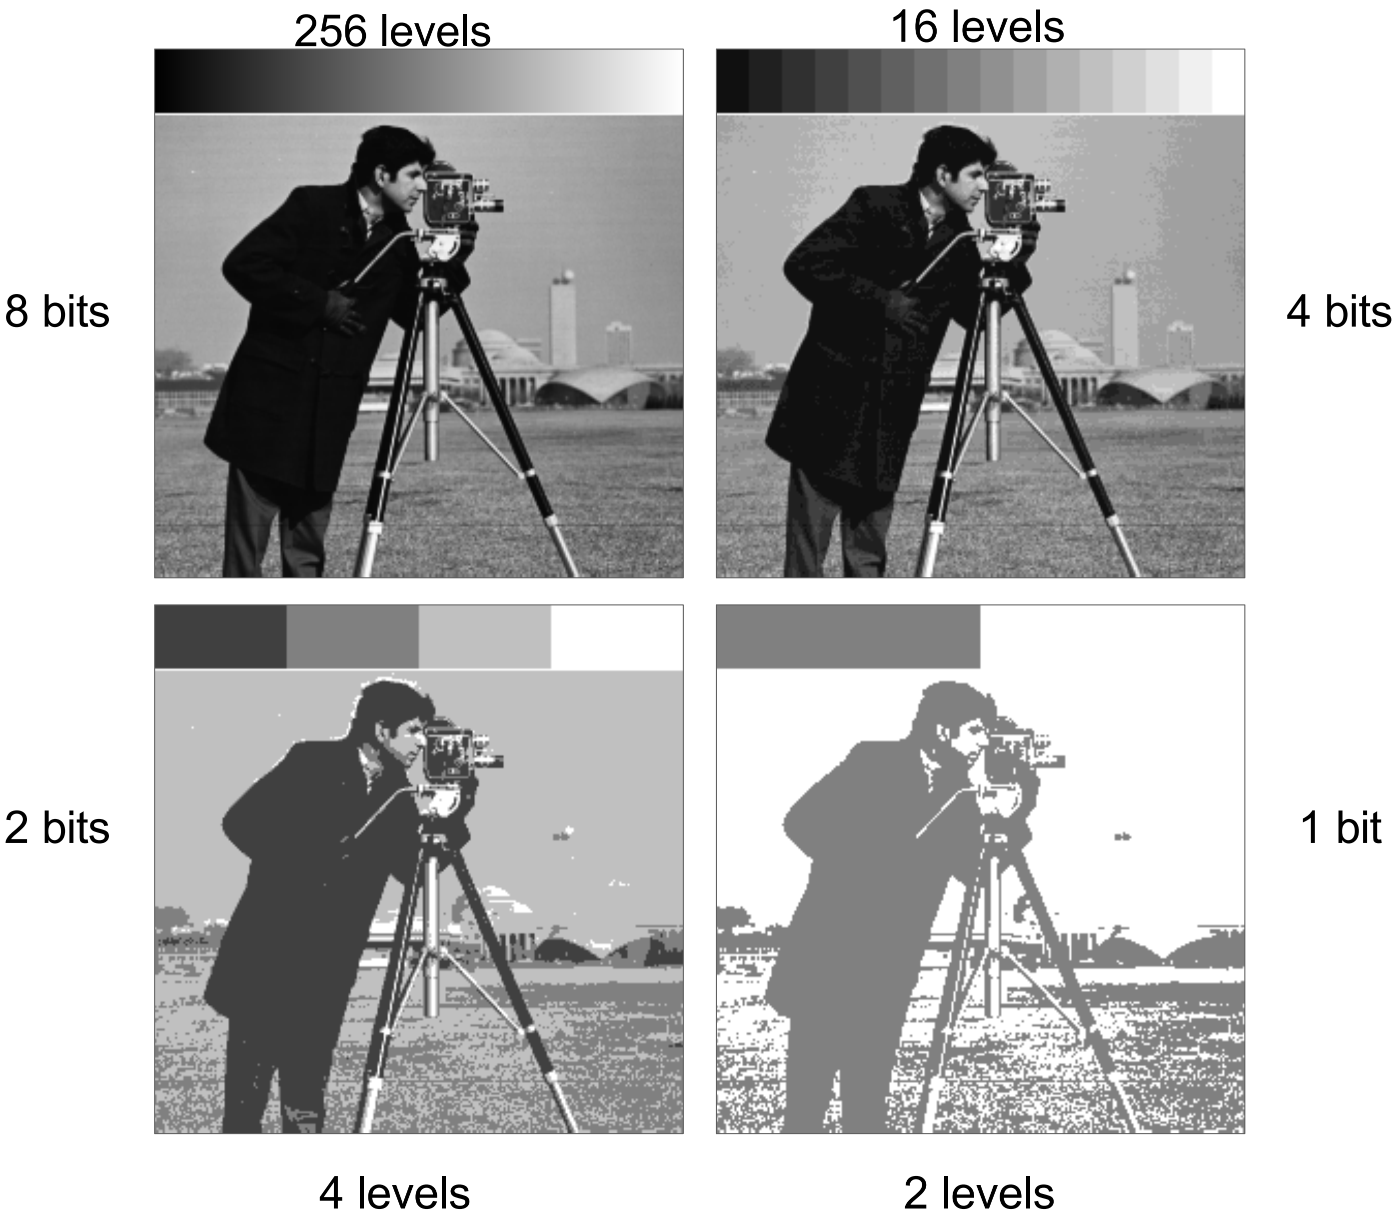

Can we do something **better**?

Nonuniform quantization?

Minimize the MSE directly?

Nonuniform Quantization and Segmentation

Search for the “optimal” threshold values to segment images

Minimum mean squared error solutions:

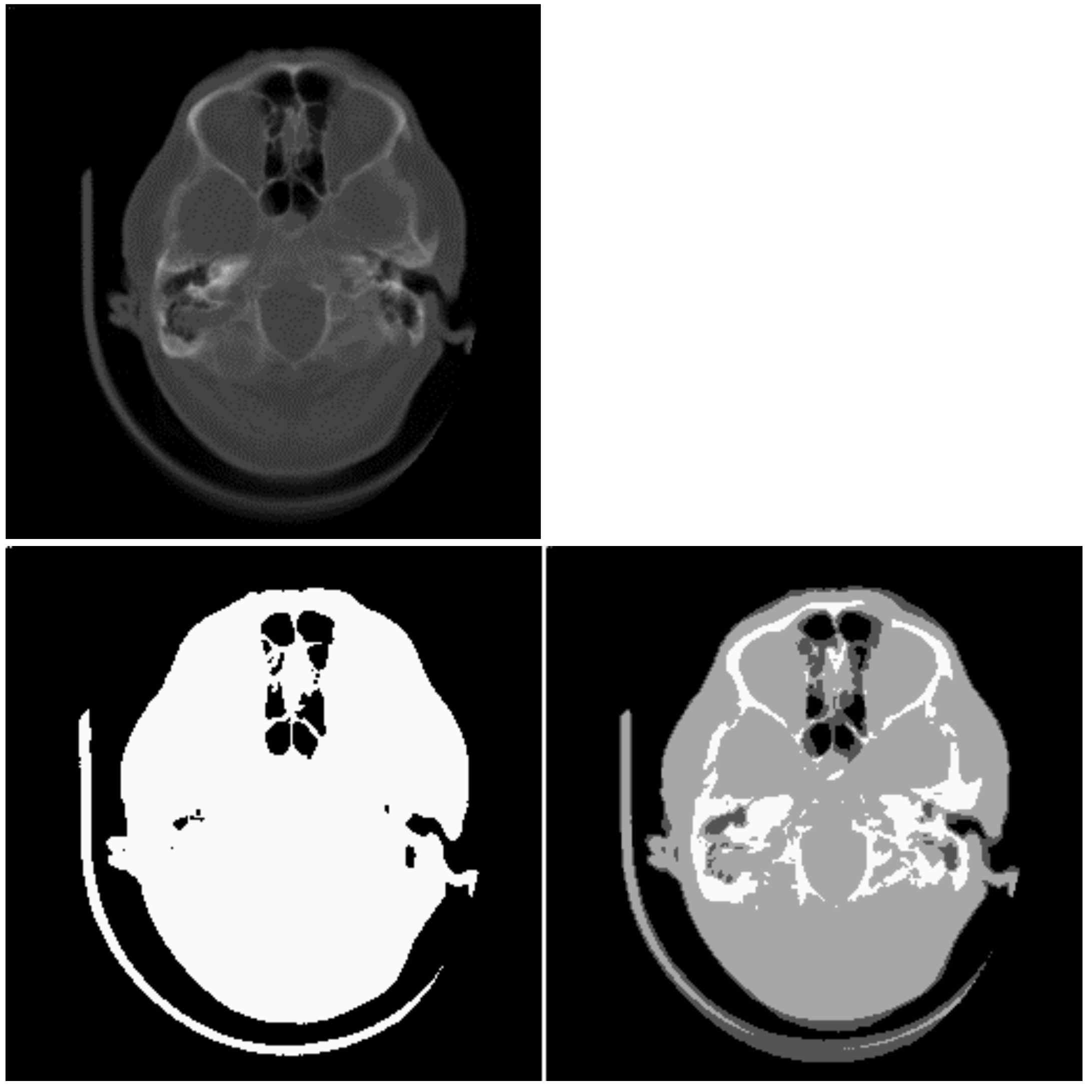

$K = 2$ $K = 4$

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

MMSE/Lloyd-Max Quantization

Goal: For a fixed K, minimize

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

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

over t_k and r_k.

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

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

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

  (midpoint solution is optimal)

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

  (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

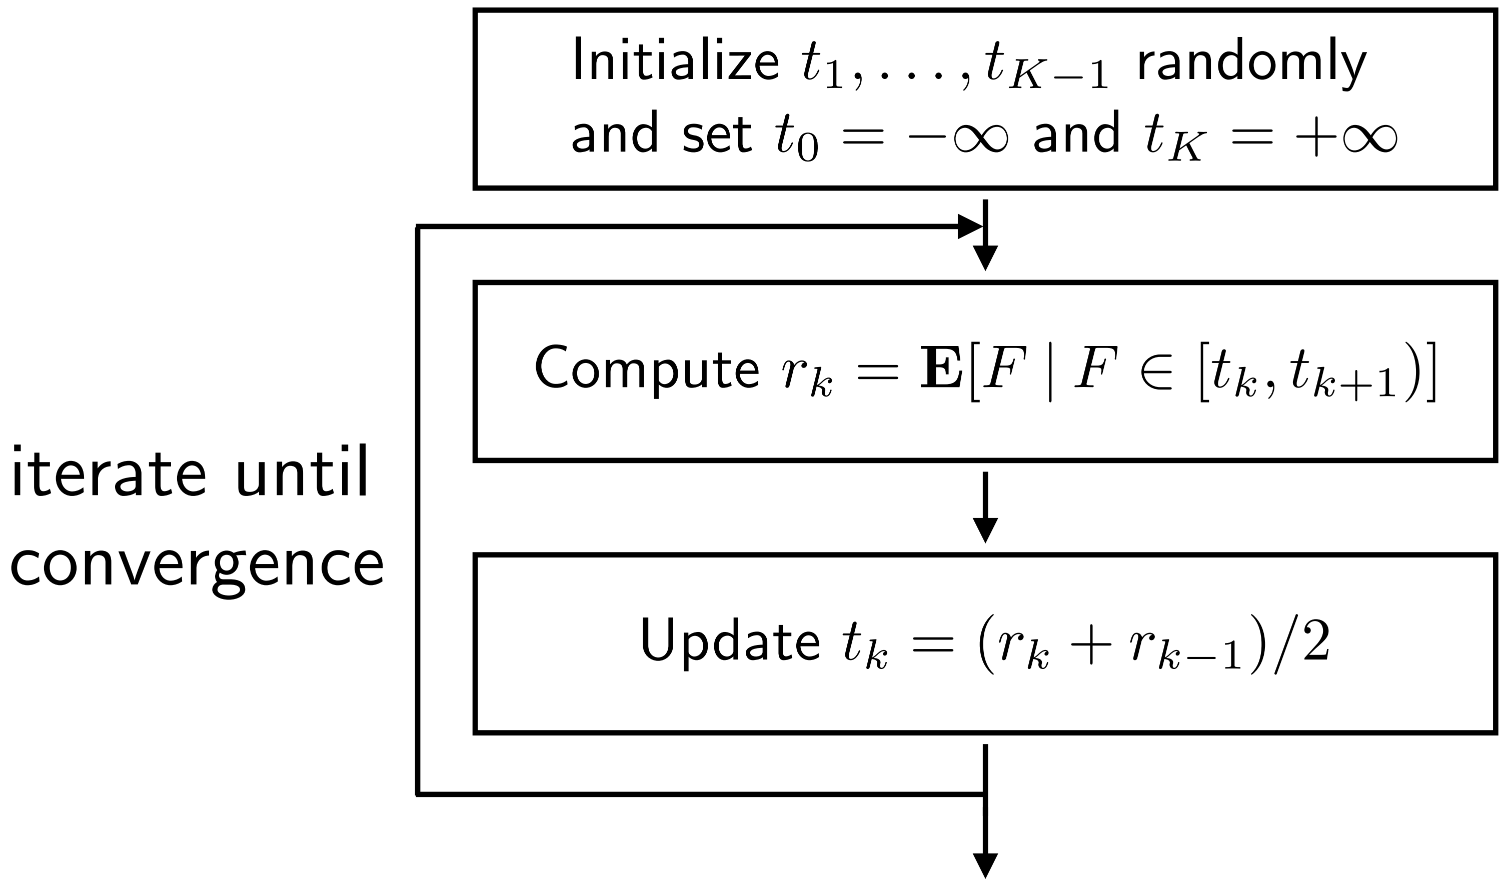

- Efficient implementation
  - r_k can be computed directly from the histogram
  - Recursive update
- Generalizations
  - Non-quadratic cost functions
  - Multivariate extensions (e.g., RGB)
  - This is just the K-means algorithm

Lloyd-Max Quantization

uniform quantization

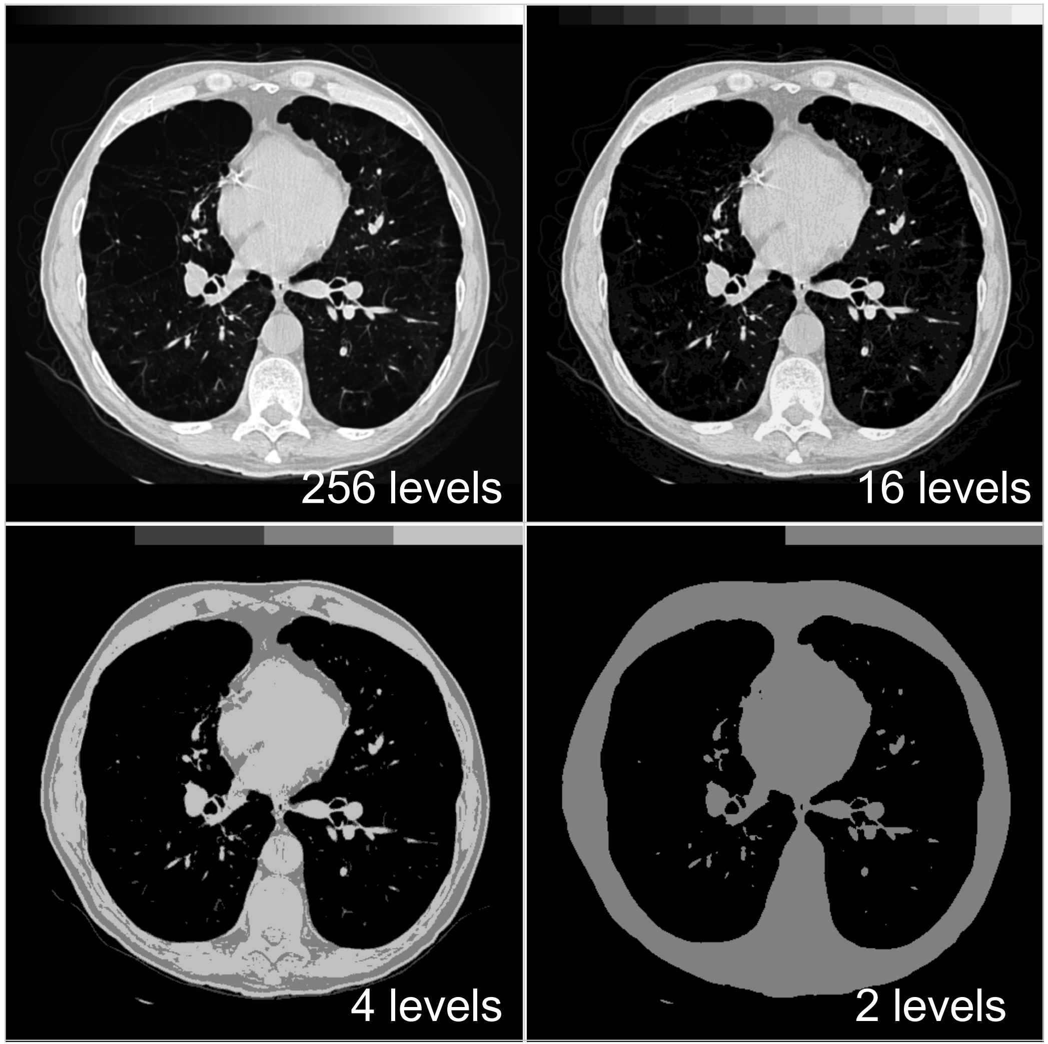

Lloyd-Max quantization

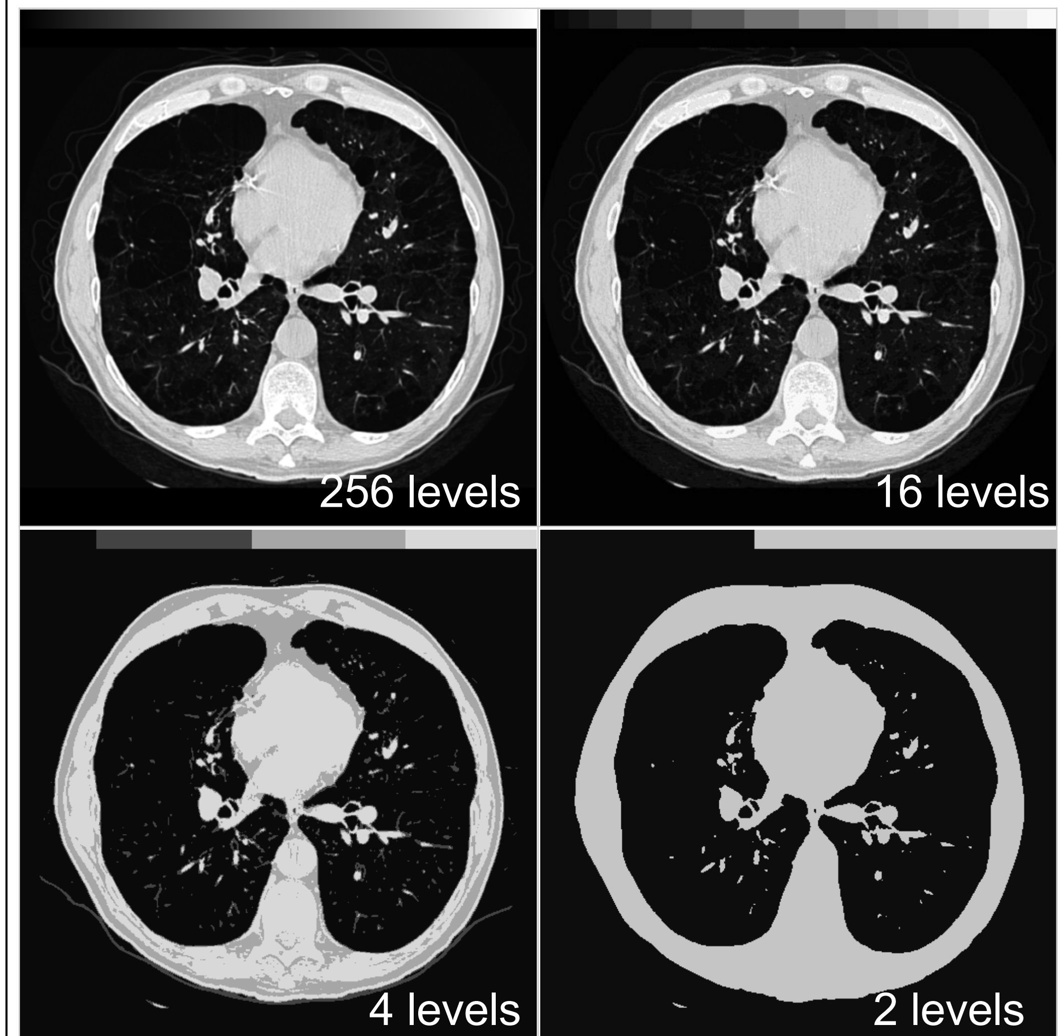

- Other useful properties:
  - Unbiased estimate: $\mathbf{E}[\tilde{F}] = \mathbf{E}[F]$
  - Error $E = F - \tilde{F}$ is **orthogonal** to the quantized value: $\mathbf{E}[\tilde{F}E] = 0$

Grayscale vs. Spatial Resolution Tradeoff

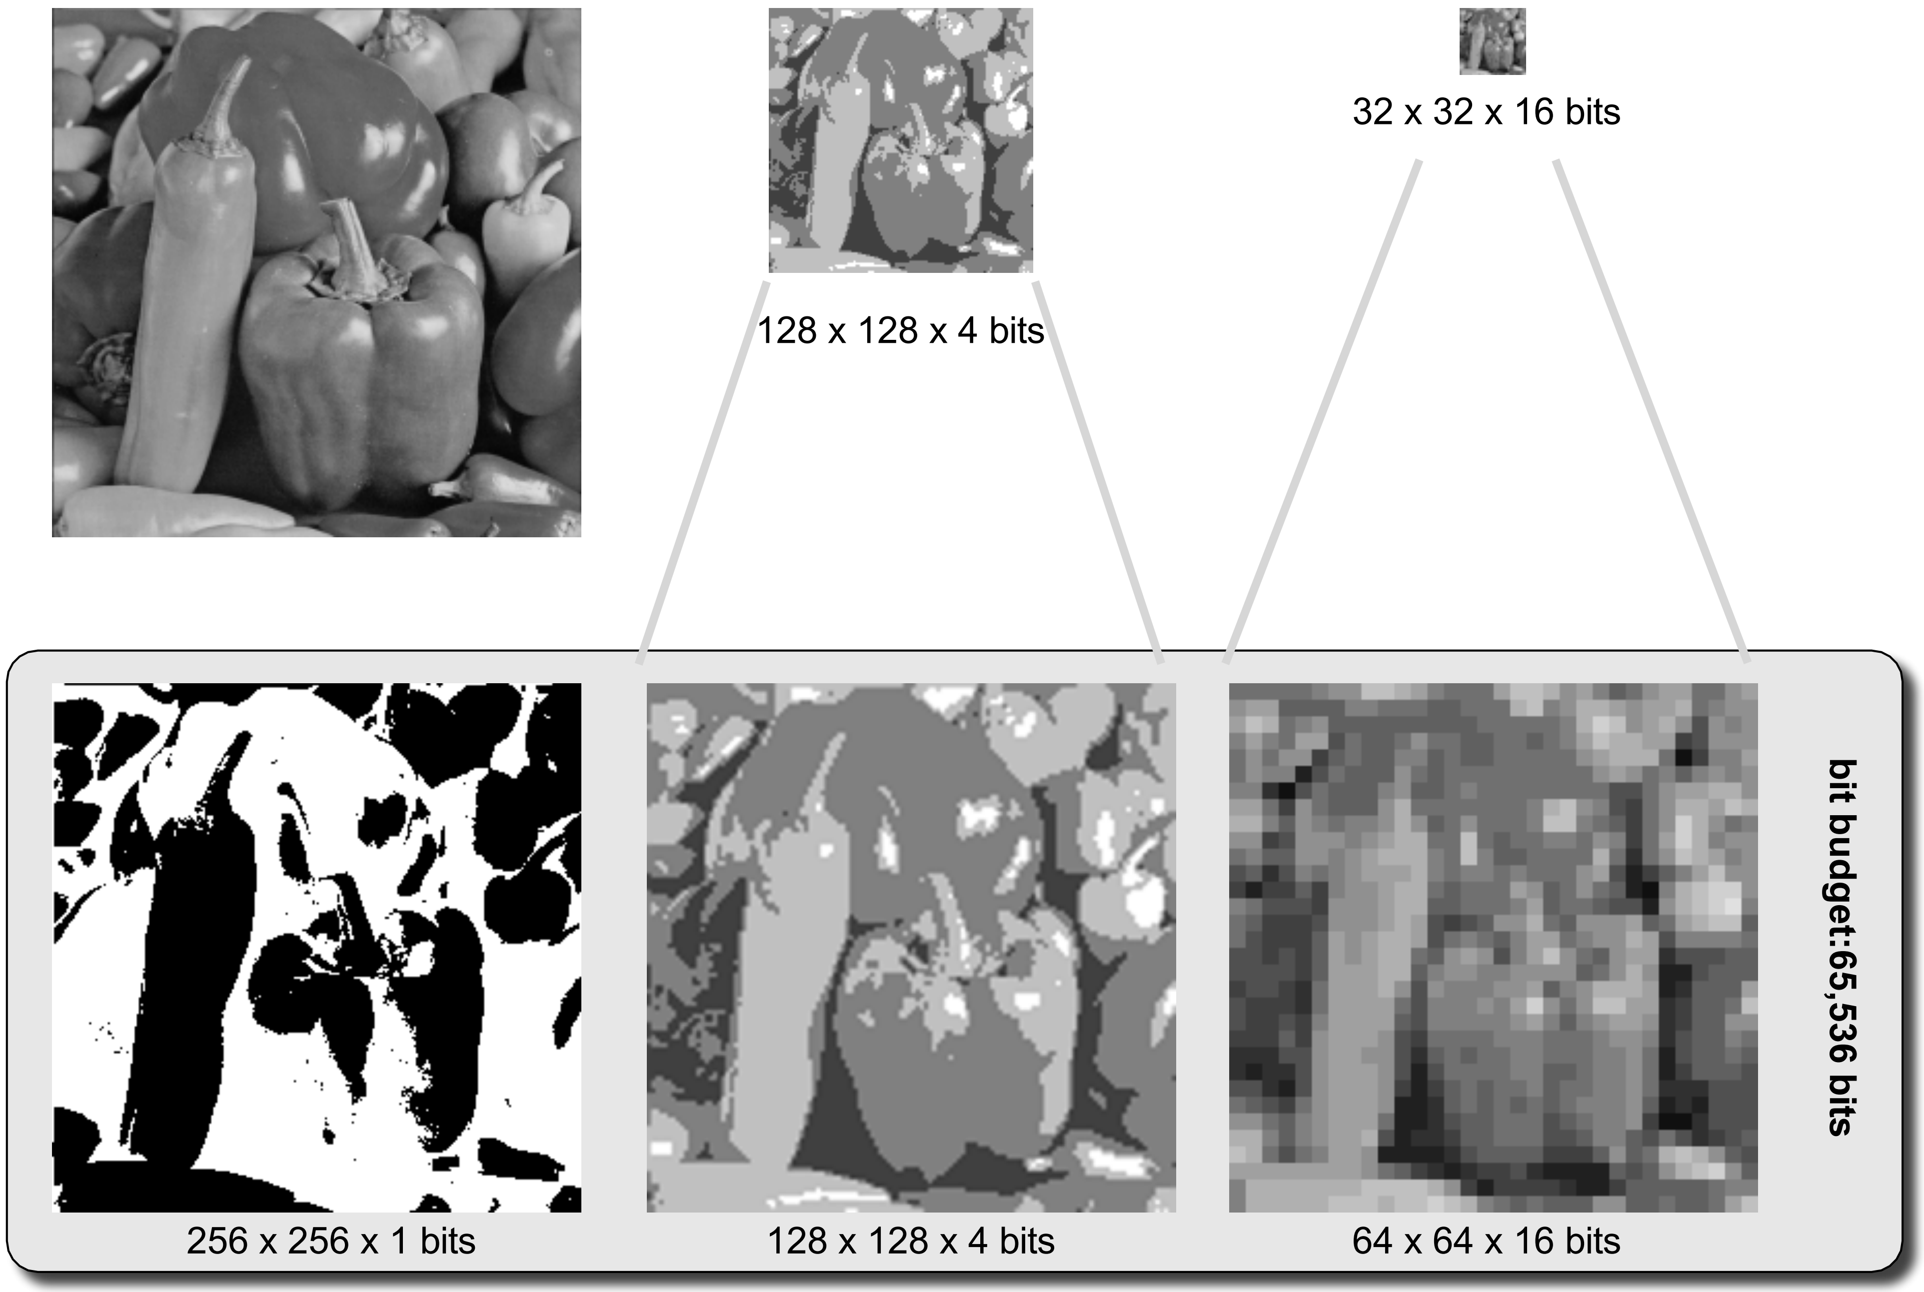

These all use the same number of bits!

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 Processers" (RIPs) in printing systems

Can we do better than just thresholding?

Dithering

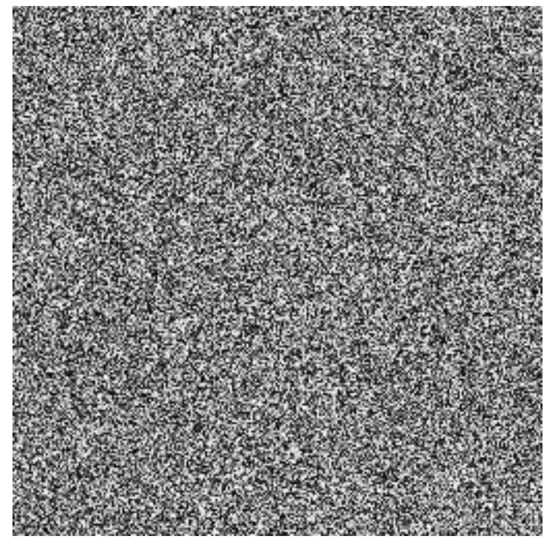

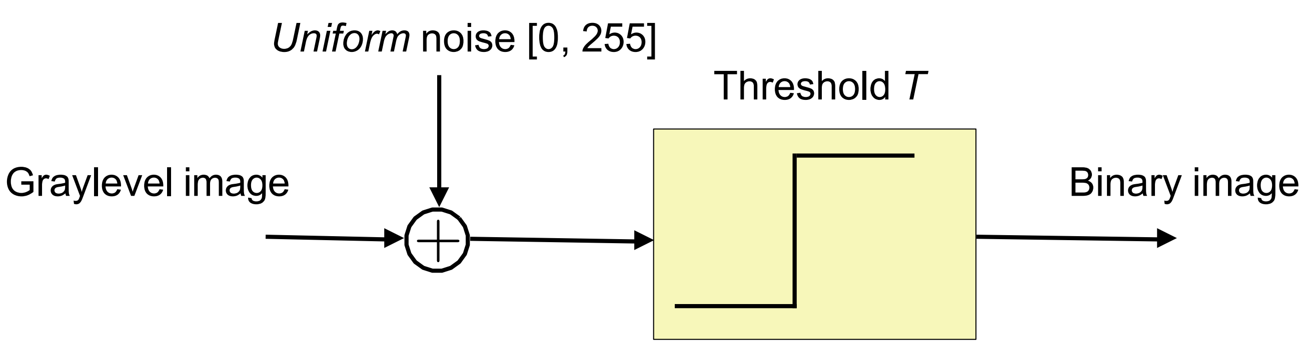

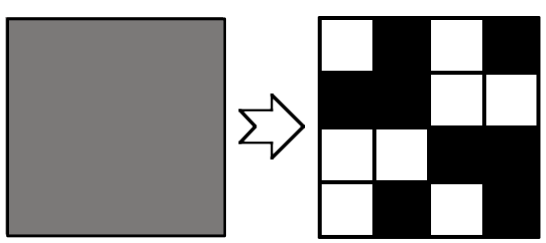

grayscale ramp

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.