Image Processing

Chapter 2

Image sampling and acquisition

Prof. Michael Unser, LIB `bigwww.epfl.ch`
Prof. Dimitri Van De Ville, MIP `miplab.epfl.ch`

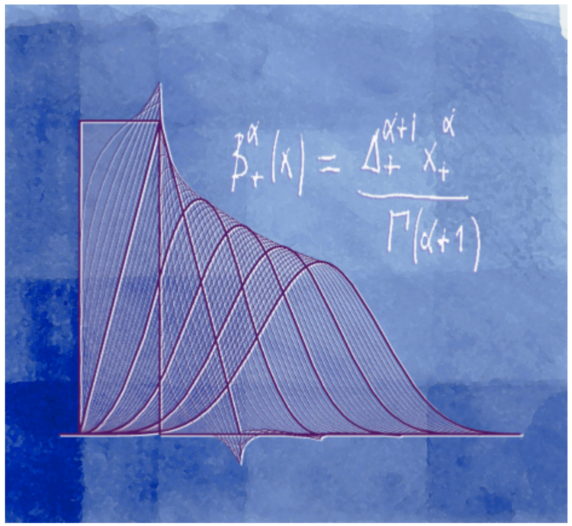

October 2024

CONTENT

- **2.1 Sampling theory**
  - Review of 1D sampling theory
  - Sampling in two and multiple dimensions
- **2.2 Acquisition systems**
  - Real acquisition systems
  - Aliasing problems
- **2.3 Quantization**
  - Uniform quantizer
  - Minimum error (Lloyd-Max) quantizer
  - Grayscale versus resolution

2.1 SAMPLING THEORY

- Review of 1D sampling theory
- Sampling in Fourier domain
- Sampling and aliasing
- Shannon’s sampling theorem
- Sampling in 2D
- Sampling in multiple dimensions

Review of 1D sampling theory

- Ideal sampling = multiplication with a Dirac comb

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

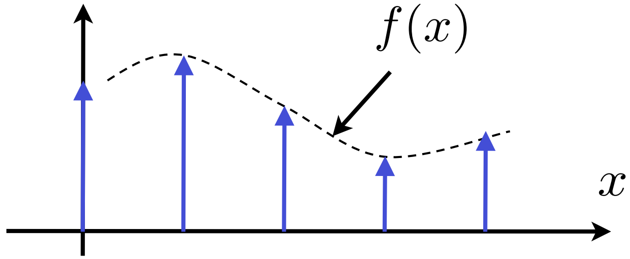

$$f(x)\delta(x - x_0) = f(x_0)\delta(x - x_0)$$

Review of 1D sampling theory

- Ideal sampling = multiplication with a Dirac comb

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

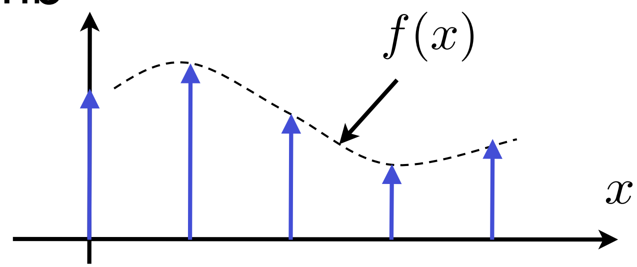

- Sampling implies periodization in the Fourier domain

Indeed, the Fourier transform of $f_T(x)$ is given by:

$$\begin{aligned} \hat{f}_T(\omega) &= \int_{\mathbb{R}} f_T(x) e^{-j\omega x}\,\mathrm{d}x = \sum_{k \in \mathbb{Z}} f(kT) \overbrace{\int_{\mathbb{R}} \delta(x - kT) e^{-j\omega x}\,\mathrm{d}x}^{e^{-j\omega x}|_{x=kT}} \\ &= \sum_{k \in \mathbb{Z}} f(kT) e^{-j\omega kT} \end{aligned}$$

which is a $\left(\frac{2\pi}{T}\right)$-periodic function of ω

Dirac sampling sequence

$$\sum_{k\in\mathbb{Z}} \delta(x-kT) \quad \stackrel{\mathcal{F}}{\longleftrightarrow} \quad ?$$

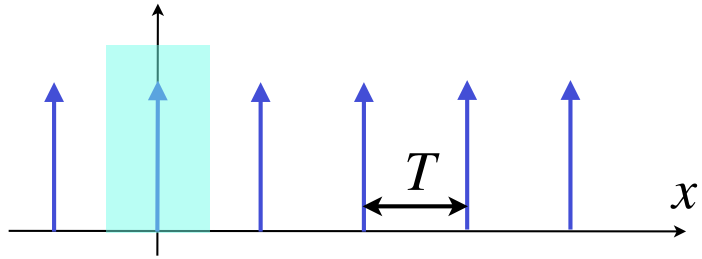

- Explicit determination of Fourier transform

1. Fourier series (in the sense of distributions)

   Justification: $\sum_{k\in\mathbb{Z}} \delta(x-kT)$ is T-periodic

$$\sum_{k\in\mathbb{Z}} \delta(x-kT) = \sum_{n\in\mathbb{Z}} e^{jn\omega_0 x} \underbrace{\frac{1}{T}\int_{-T/2}^{T/2} \delta(x)\, e^{-jn\omega_0 x} \mathrm{d}x}_{\frac{1}{T}}$$

$$\Rightarrow \qquad \sum_{k\in\mathbb{Z}} \delta(x-Tk) = \sum_{n\in\mathbb{Z}} \frac{1}{T} e^{jn\omega_0 x} \qquad \text{with} \qquad \omega_0 = \tfrac{2\pi}{T}$$

Reminder: Fourier series

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

Sampling in the Fourier domain

- Sampling formula

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

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

- Continuous-space to discrete-space Fourier transform

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

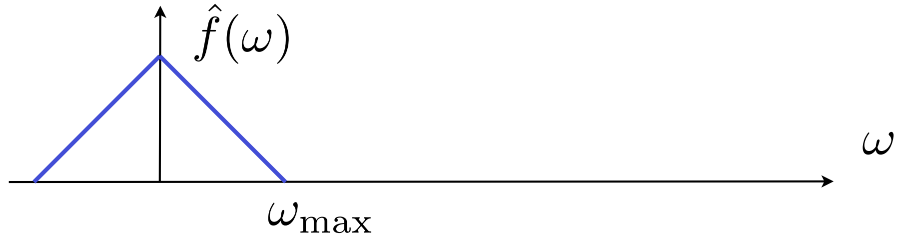

Normalized sampling step ($T = 1$)

$F(e^{j\omega}) = \hat{f}_{T=1}(\omega)$: Fourier transform of discrete sequence $\{f[k]\}_{k\in\mathbb{Z}}$

$$F(e^{j\omega}) = \boxed{\sum_{k\in\mathbb{Z}} f[k]e^{-j\omega k} = \sum_{n\in\mathbb{Z}} \hat{f}(\omega - 2\pi n)}$$

Poisson summation formula

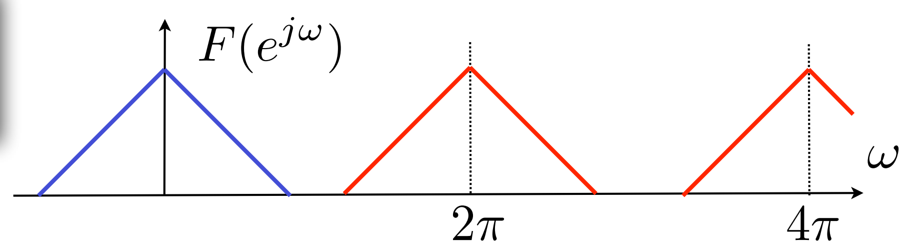

Sampling and aliasing

- Sampled signal

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

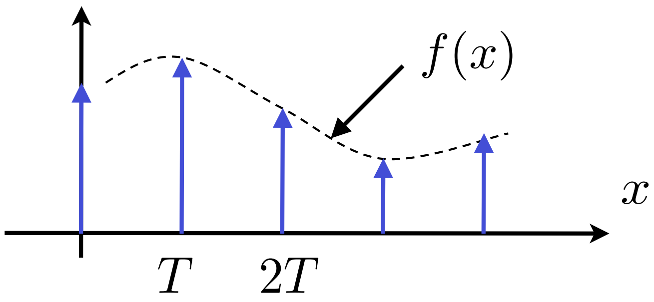

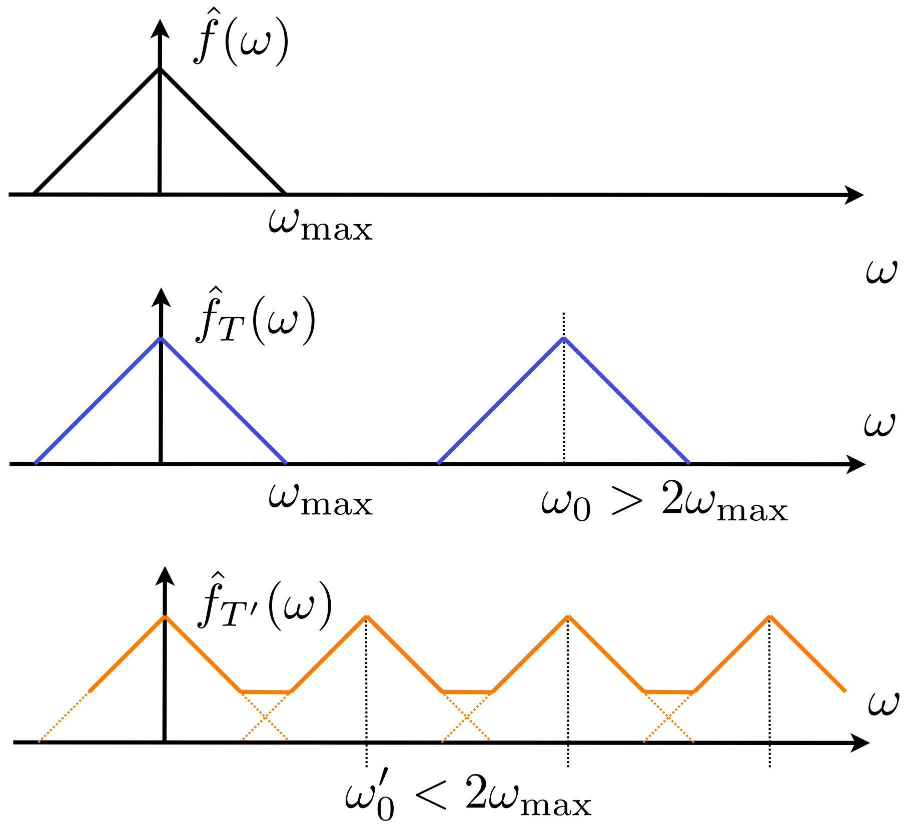

$$\omega_0 = \frac{2\pi}{T}$$

“under-sampling”

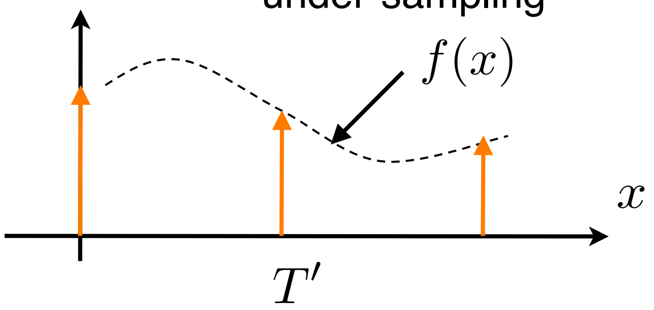

Sampling and reconstruction: Fourier-domain analysis

- Continuous-domain input signal

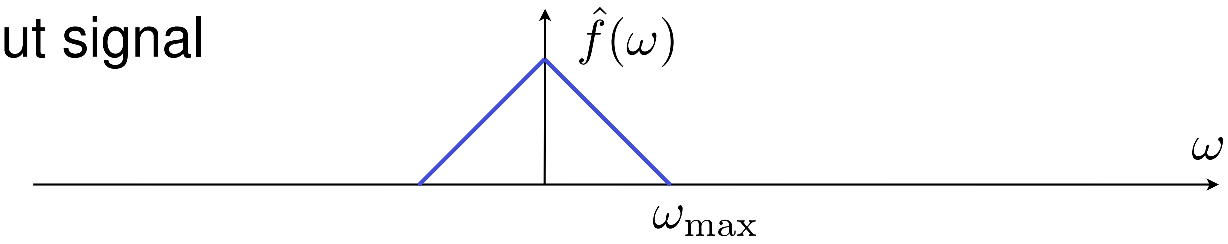

- Sampling

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

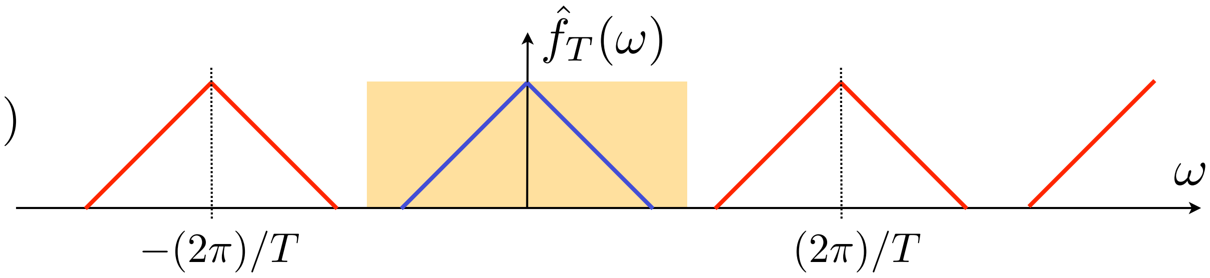

- Ideal lowpass filtering

$$\hat{f}_r(\omega) = \hat{f}_T(\omega) \cdot T \,\text{rect}\left(\frac{\omega T}{2\pi}\right)$$

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

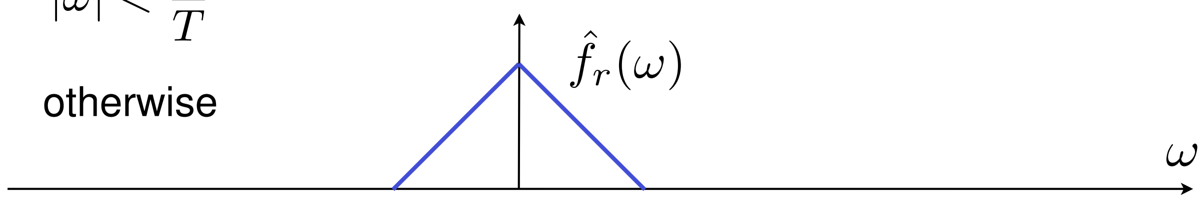

$\Rightarrow$ Condition for a perfect recovery: $\omega_{\max} < \frac{\pi}{T}$ (Nyquist criterion)

Shannon's sampling theorem

- Ideal reconstruction process

$$\delta(x)$$

$$\sum_{k\in\mathbb{Z}} f(kT)\,\delta(x-kT) \longrightarrow \boxed{\text{Ideal lowpass filter}} \longrightarrow$$

$$h_T(x) = \mathrm{sinc}\left(\frac{x}{T}\right) \quad \xleftrightarrow{\mathcal{F}} \quad T\cdot \mathrm{rect}\left(\frac{\omega}{2\pi/T}\right)$$

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

- $\mathrm{sinc}(x) \quad \xleftrightarrow{\mathcal{F}} \quad \mathrm{rect}\left(\frac{\omega}{2\pi}\right)$
- $f(x/T) \quad \xleftrightarrow{\mathcal{F}} \quad |T|\hat{f}(T\omega)$
- $\mathrm{sinc}(x/T) \quad \xleftrightarrow{\mathcal{F}} \quad T\mathrm{rect}\left(\frac{T\omega}{2\pi}\right)$

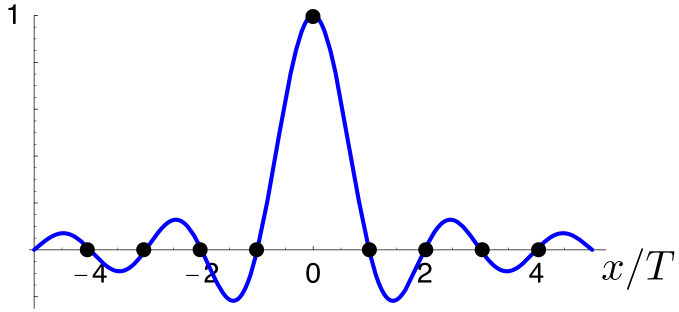

Theorem: A function $f(x)$ that is bandlimited to $\omega_{\max}$ can be reconstructed exactly from its equidistant samples provided that the sampling step $T < \pi/\omega_{\max}$. Specifically,

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

where $\quad \mathrm{sinc}(x) = \dfrac{\sin(\pi x)}{(\pi x)} \quad \xleftrightarrow{\mathcal{F}} \quad \mathrm{rect}\left(\dfrac{\omega}{2\pi}\right)$

Sampling in 2-D

- Sampling function in 2D (separable and regular)

$$\sum_{(k,l)\in\mathbb{Z}^2} \delta(x-kT_x, y-lT_y) \quad \stackrel{\mathcal{F}}{\longleftrightarrow} \quad \frac{(2\pi)^2}{T_xT_y} \sum_{(m,n)\in\mathbb{Z}^2} \delta(\omega_x - \frac{2\pi m}{T_x}, \omega_y - \frac{2\pi n}{T_y})$$

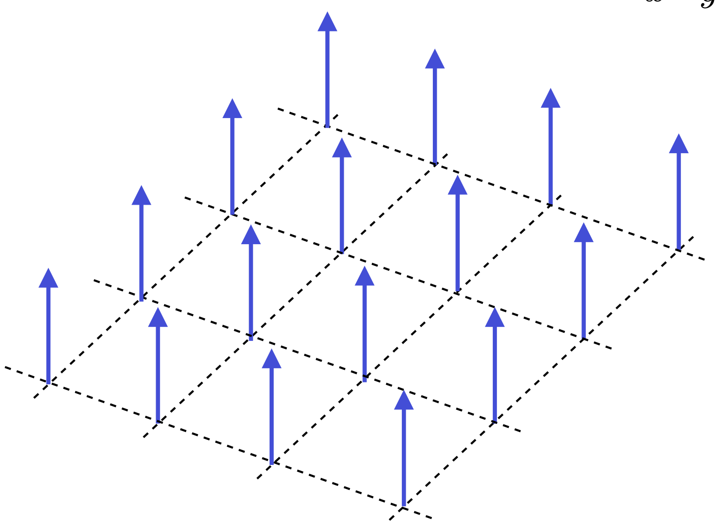

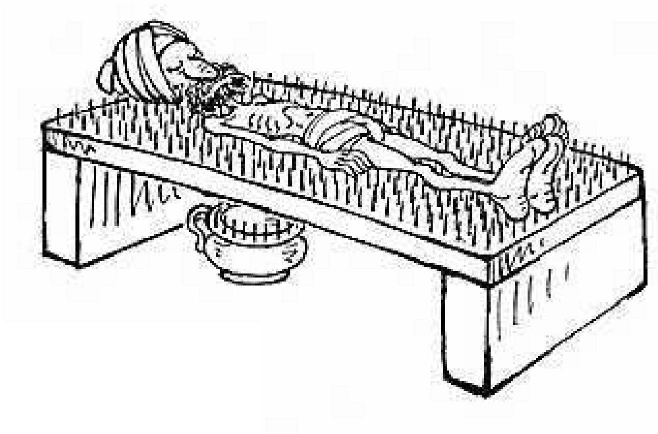

- 2D sampling formula

$$f_{T_x,T_y}(x,y) = f(x,y) \cdot \sum_{(k,l)\in\mathbb{Z}^2} \delta(x-kT_x, y-lT_y)$$

$$\Rightarrow \qquad \hat{f}_{T_x,T_y}(\omega_x,\omega_y) = \frac{1}{T_xT_y} \sum_{(m,n)\in\mathbb{Z}^2} \hat{f}(\omega_x - \frac{2\pi m}{T_x}, \omega_y - \frac{2\pi n}{T_y})$$

Sampling and spectral repetition

$$\hat{f}_{T_x,T_y}(\omega_x,\omega_y) = \frac{1}{T_x T_y}\sum_{m\in\mathbb{Z}}\sum_{n\in\mathbb{Z}}\hat{f}(\omega_x - \frac{2\pi m}{T_x}, \omega_y - \frac{2\pi n}{T_y})$$

$\hat{f}(\boldsymbol{\omega})$

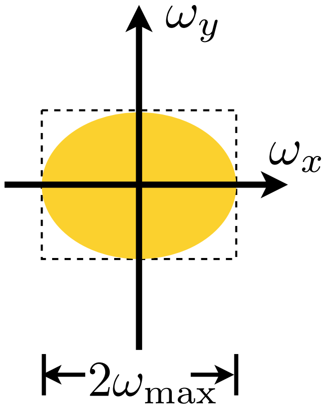

spectral support

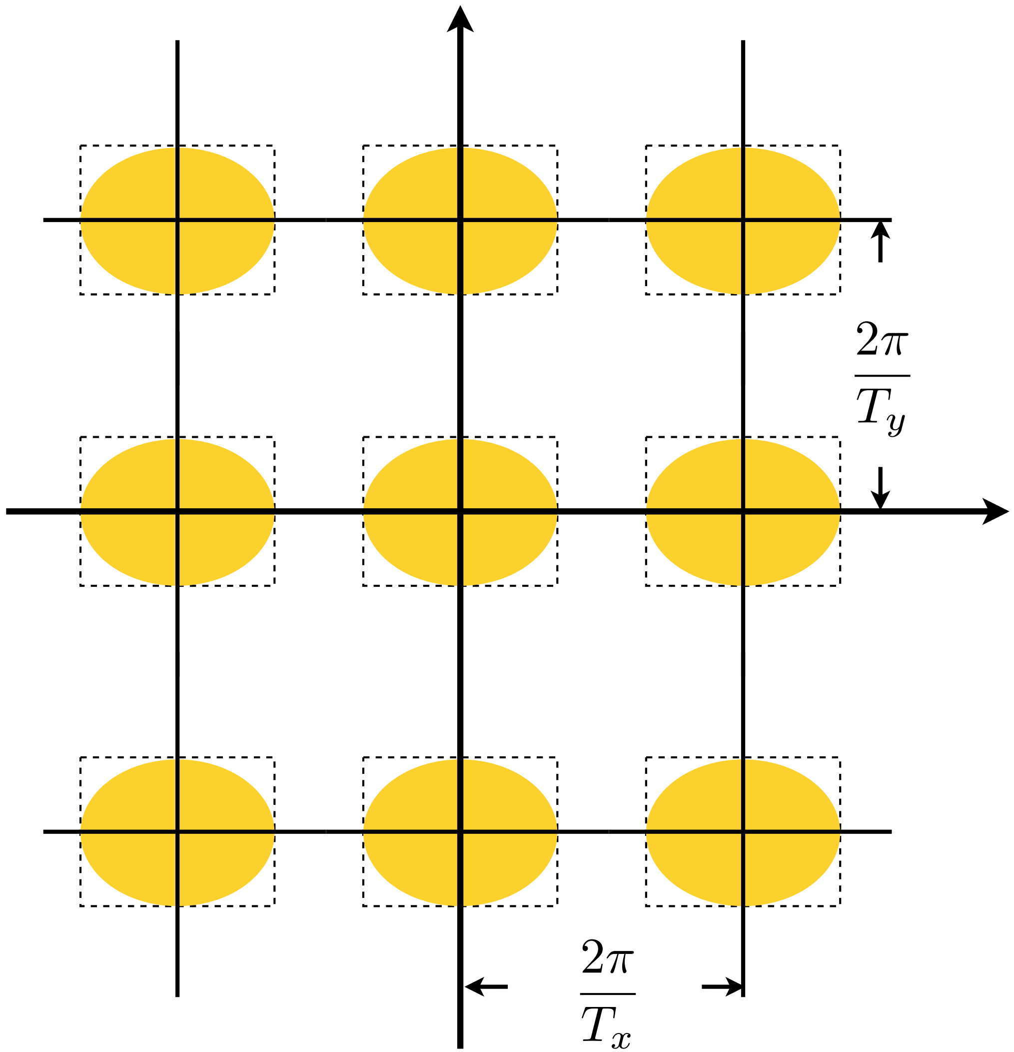

Sampling in multiple dimensions

For notational simplicity, we assume $T_1 = T_2 = \cdots = T$

- Sampling function in d dimensions (separable and regular)

$$s_T(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^d} \delta(\boldsymbol{x}-\boldsymbol{k}T) \quad \stackrel{\mathcal{F}}{\longleftrightarrow} \quad \hat{s}_T(\boldsymbol{\omega}) = \frac{(2\pi)^d}{T^d} \sum_{\boldsymbol{n}\in\mathbb{Z}^d} \delta(\boldsymbol{\omega} - \frac{2\pi\boldsymbol{n}}{T})$$

- Multidimensional sampling formula

$$f_T(\boldsymbol{x}) = f(\boldsymbol{x}) \cdot s_T(\boldsymbol{x}) \quad \stackrel{\mathcal{F}}{\longleftrightarrow} \quad \hat{f}_T(\boldsymbol{\omega}) = \frac{1}{(2\pi)^d} \left(\hat{f} * \hat{s}_T\right)(\boldsymbol{\omega})$$

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

- Condition for a perfect recovery: $\omega_{\max} < \dfrac{\pi}{T}$

$$f(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^d} f(\boldsymbol{k}T)\,\mathrm{sinc}\left(\frac{\boldsymbol{x}-\boldsymbol{k}T}{T}\right) \quad \text{with} \quad \mathrm{sinc}(\boldsymbol{x}) = \prod_{i=1}^{d} \mathrm{sinc}(x_i)$$

Sampling lattices

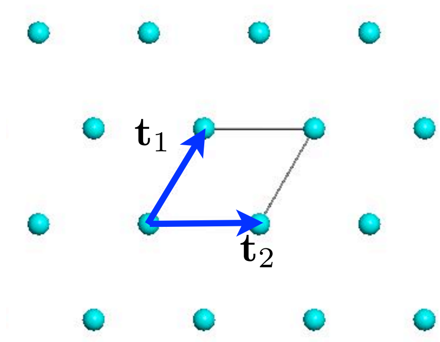

Cartesian lattice: $\mathbb{Z}^d$

Sampling lattice: $\mathbf{T}\mathbb{Z}^d$ with $\mathbf{T} = [\mathbf{t}_1 \cdots \mathbf{t}_d]$ (lattice vectors)

Affine transformation: $f(\mathbf{T}\boldsymbol{x}) \quad \overset{\mathcal{F}}{\longleftrightarrow} \quad \frac{1}{|\det(\mathbf{T})|}\hat{f}\left(\mathbf{T}^{-\mathsf{T}}\boldsymbol{\omega}\right)$

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

Dual (or reciprocal) lattice (Fourier domain): $2\pi\mathbf{T}^{-\mathsf{T}}\mathbb{Z}^d$

- Adapted Poisson summation formula

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

Sampling and spectral repeat (hexagonal case)

Reciprocal lattice: $2\pi\mathbf{T}^{-\mathsf{T}}\mathbb{Z}^d$

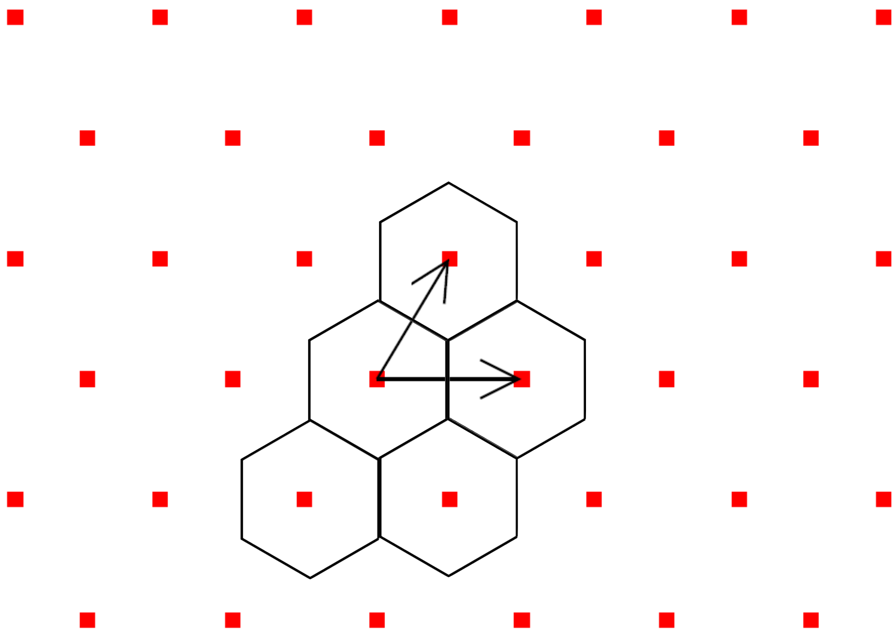

spatial domain

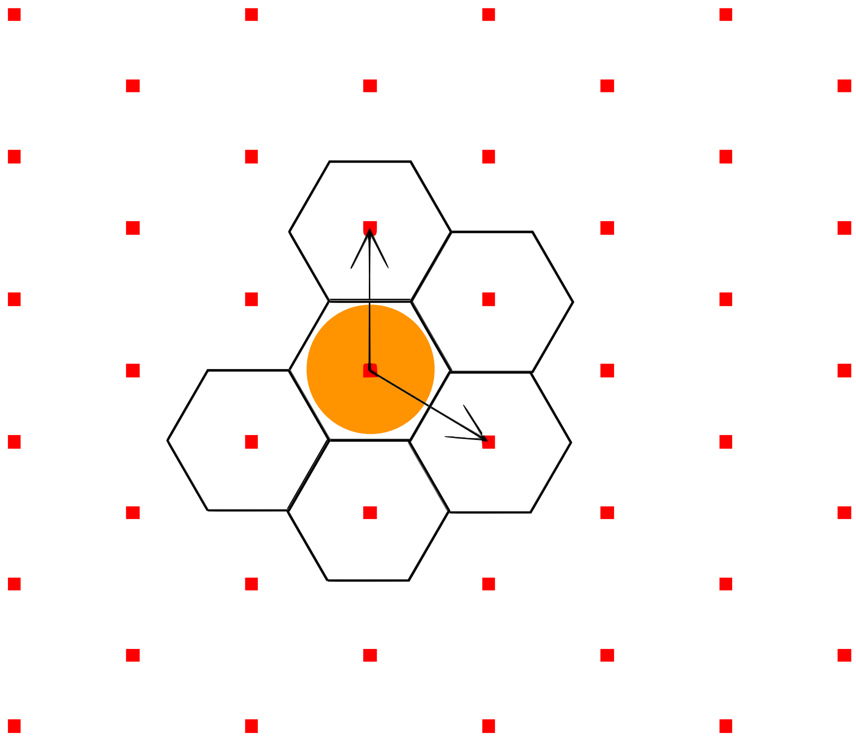

spectral domain

$$\mathbf{T} = [\mathbf{t_1}|\mathbf{t_2}] = \begin{bmatrix} 1 & 1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}$$

$$\hat{\mathbf{T}} = [\hat{\mathbf{t}}_1|\hat{\mathbf{t}}_2] = \mathbf{T}^{-T} = \begin{bmatrix} 1 & 0 \\ -\sqrt{3}/3 & 2\sqrt{3}/3 \end{bmatrix}$$

Note: regular hexagonal lattice is **optimal** for isotropic band-limited spectra (i.e., maximal frequency component lies within a disc)

Nature’s favorite

Maryna Viazovska

- Awarded Fields Medal for solving sphere-packing problem in 8 and 24 dimensions!

Annals of Mathematics **185** (2017), 991–1015
https://doi.org/10.4007/annals.2017.185.3.7

The sphere packing problem in dimension 8

By Maryna S. Viazovska

2.2 ACQUISITION SYSTEMS

- Real acquisition systems
- Acquisition models
- Aliasing problems

Real acquisition systems

In practice, the sampling is not ideal

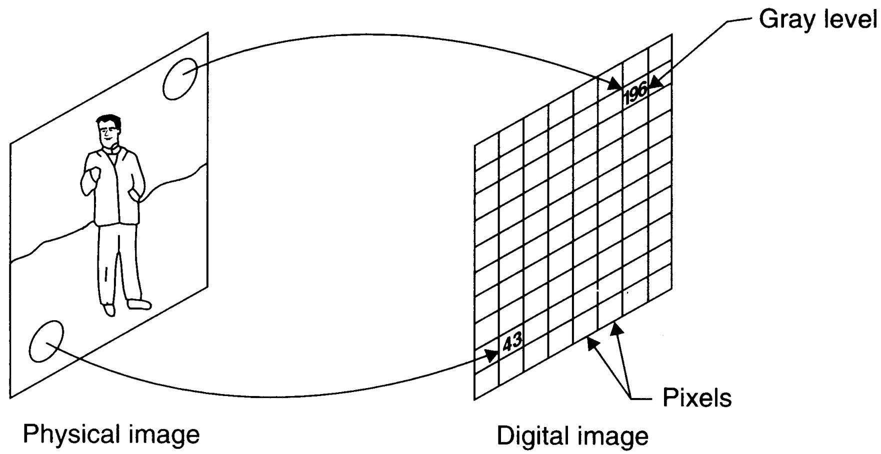

- Pixel measurement process

$$p[\boldsymbol{k}] = \int_{\mathbb{R}^d} \varphi_a(\boldsymbol{y} - \boldsymbol{k}T) f(\boldsymbol{y})\, \mathrm{d}y_1 \cdots \mathrm{d}y_d = \langle \varphi_a(\cdot - \boldsymbol{k}T), f \rangle_{L_2(\mathbb{R}^d)}$$

$p[\boldsymbol{k}]$: pixel value at location $\boldsymbol{k} = (k_1, \ldots, k_d)$

$\varphi_a(\boldsymbol{y})$: sampling aperture (or integration window)

Typically: $\int_{\mathbb{R}^d} \varphi_a(\boldsymbol{x})\, \mathrm{d}x_1 \cdots \mathrm{d}x_d = 1$ (normalized intensity)

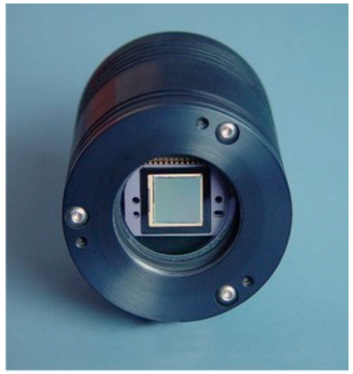

- Example: CCD camera

$$\varphi_a(x, y) = \frac{1}{T^2} \mathrm{rect}(x/T) \cdot \mathrm{rect}(y/T)$$

Equivalent pixel measurement model

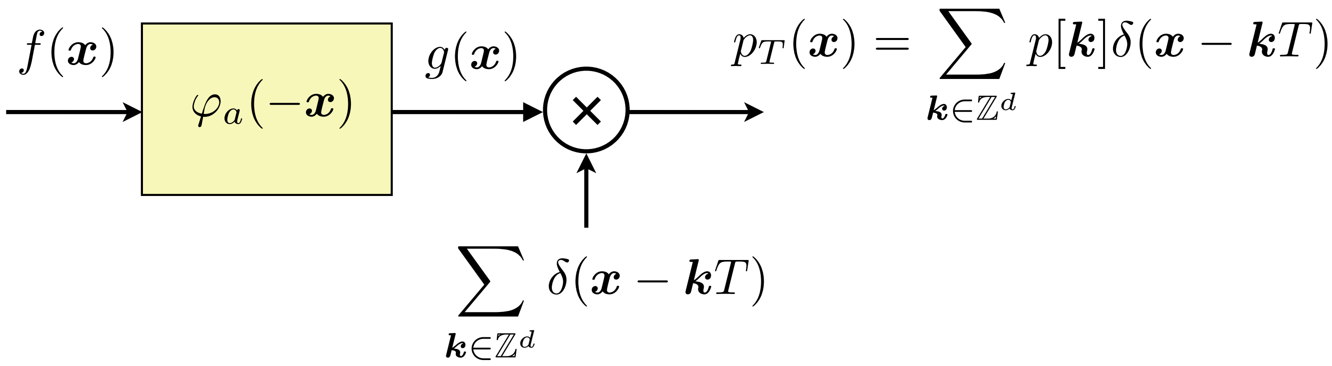

- Prefiltering $\Rightarrow$ auxiliary image function

$$g(\boldsymbol{x}) = \int_{\mathbb{R}^d} \varphi_a(\boldsymbol{y}-\boldsymbol{x}) f(\boldsymbol{y})\, \mathrm{d}y_1 \cdots \mathrm{d}y_d = (\varphi_a^\vee * f)(\boldsymbol{x})$$

where $\quad \varphi_a^\vee(\boldsymbol{x}) = \varphi_a(-\boldsymbol{x})$

- Ideal sampling: $\quad p[\boldsymbol{k}] = \langle \varphi_a(\cdot - \boldsymbol{k}T), f\rangle = g(\boldsymbol{x})\Big|_{\boldsymbol{x}=\boldsymbol{k}T}$

Realistic acquisition model

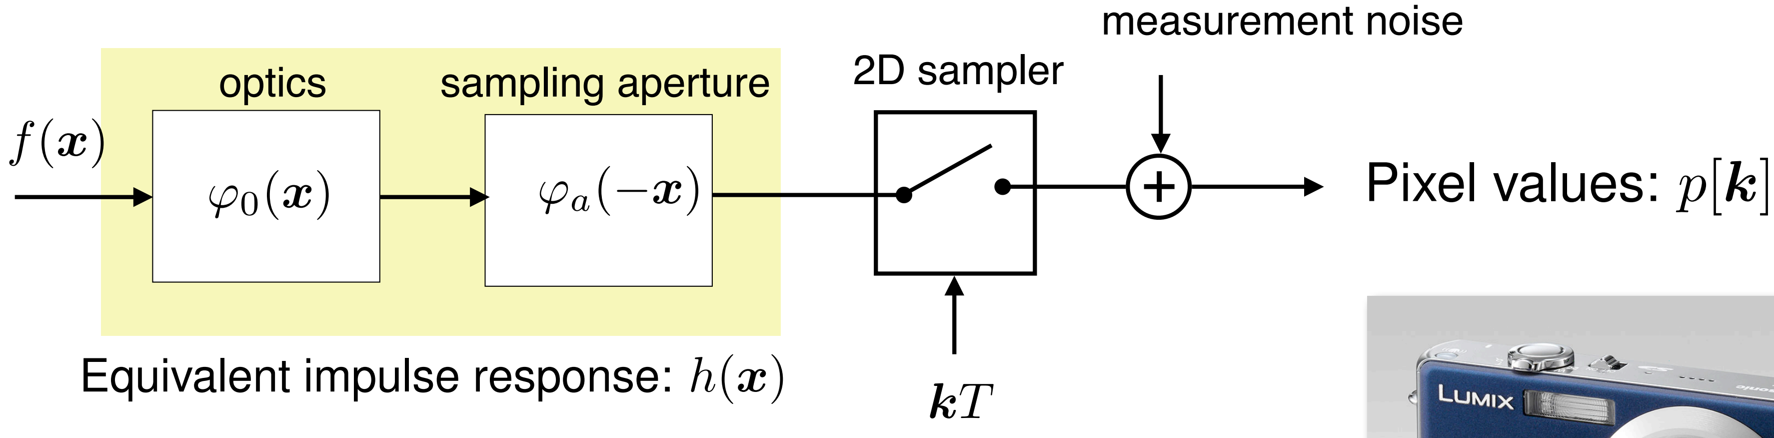

Complete image-acquisition model:

$$p[\boldsymbol{k}] = (h * f)(\boldsymbol{x})|_{\boldsymbol{x}=\boldsymbol{k}T} + n[\boldsymbol{k}]$$

$\varphi_0(\boldsymbol{x})$: point-spread function (image formation)

$\varphi_a(\boldsymbol{x})$: sampling aperture

$h(\boldsymbol{x}) = (\varphi_a^T * \varphi_0)(\boldsymbol{x})$: equivalent LSI system

$n(\boldsymbol{k})$: additive measurement noise

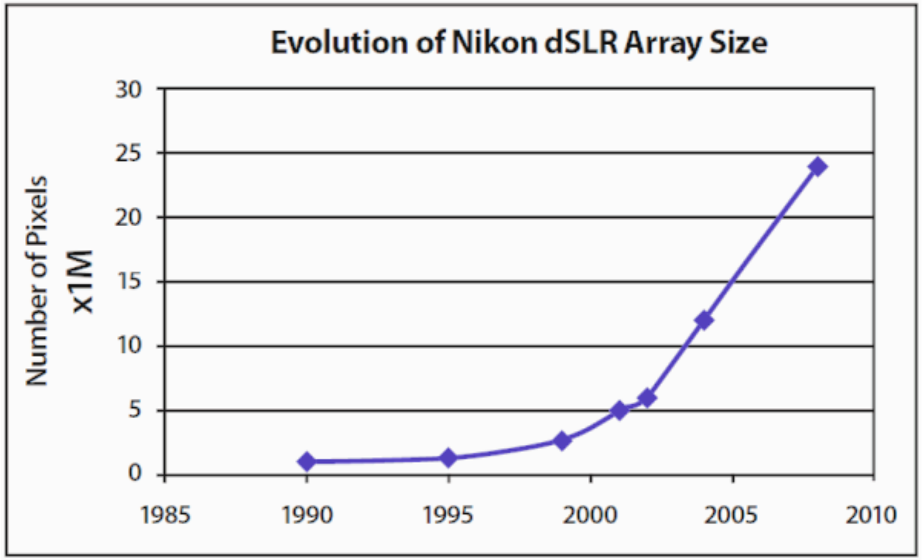

[source: dvinfo.net July 14, 2010]

Aliasing problems

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

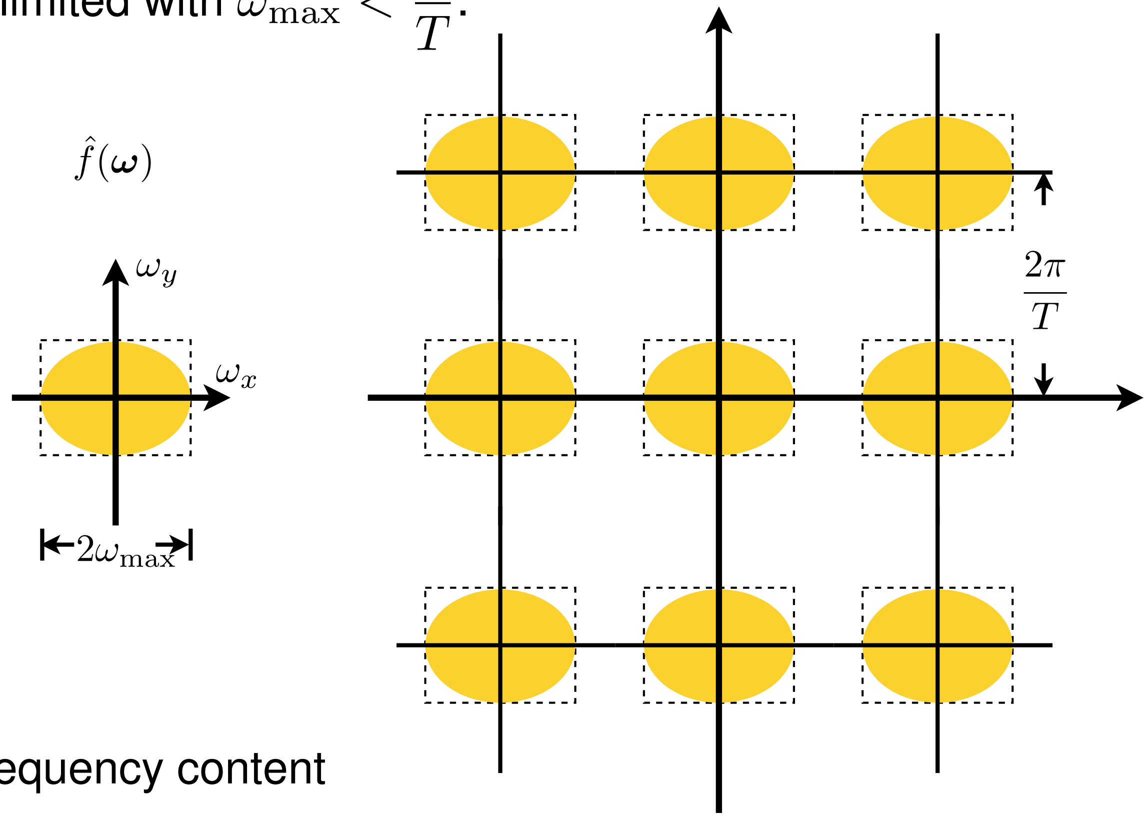

Practical solutions

- Adapt sampling step to spatial frequency content
- Implicit lowpass filtering prior to sampling

  via φ_0: imaging system;
  φ_a: sampling aperture

Example of aliased picture

Sampling

Prefiltering/sampling

“Bad” aliasing in printing applications

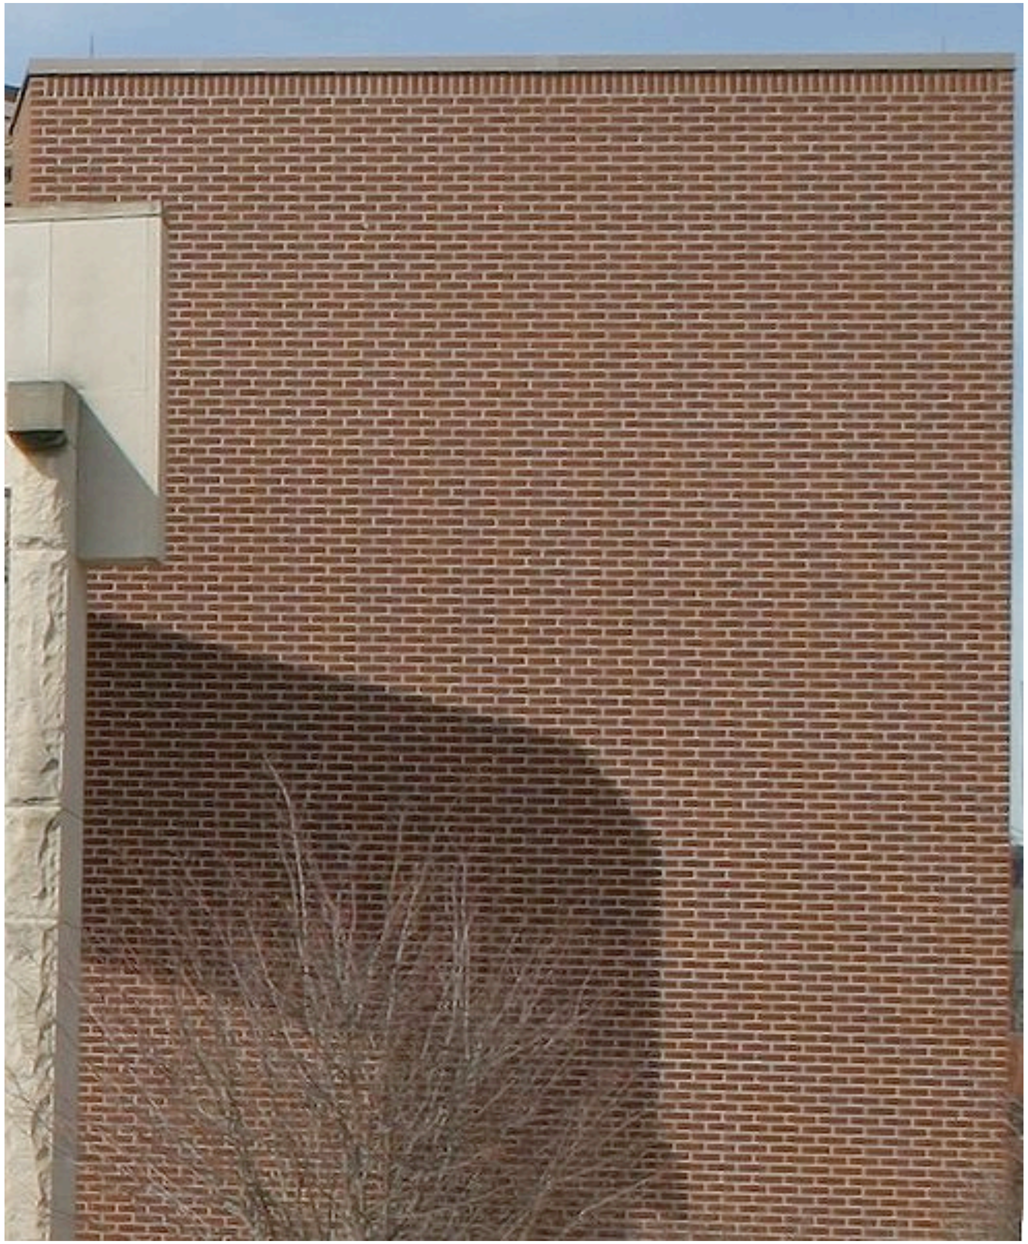

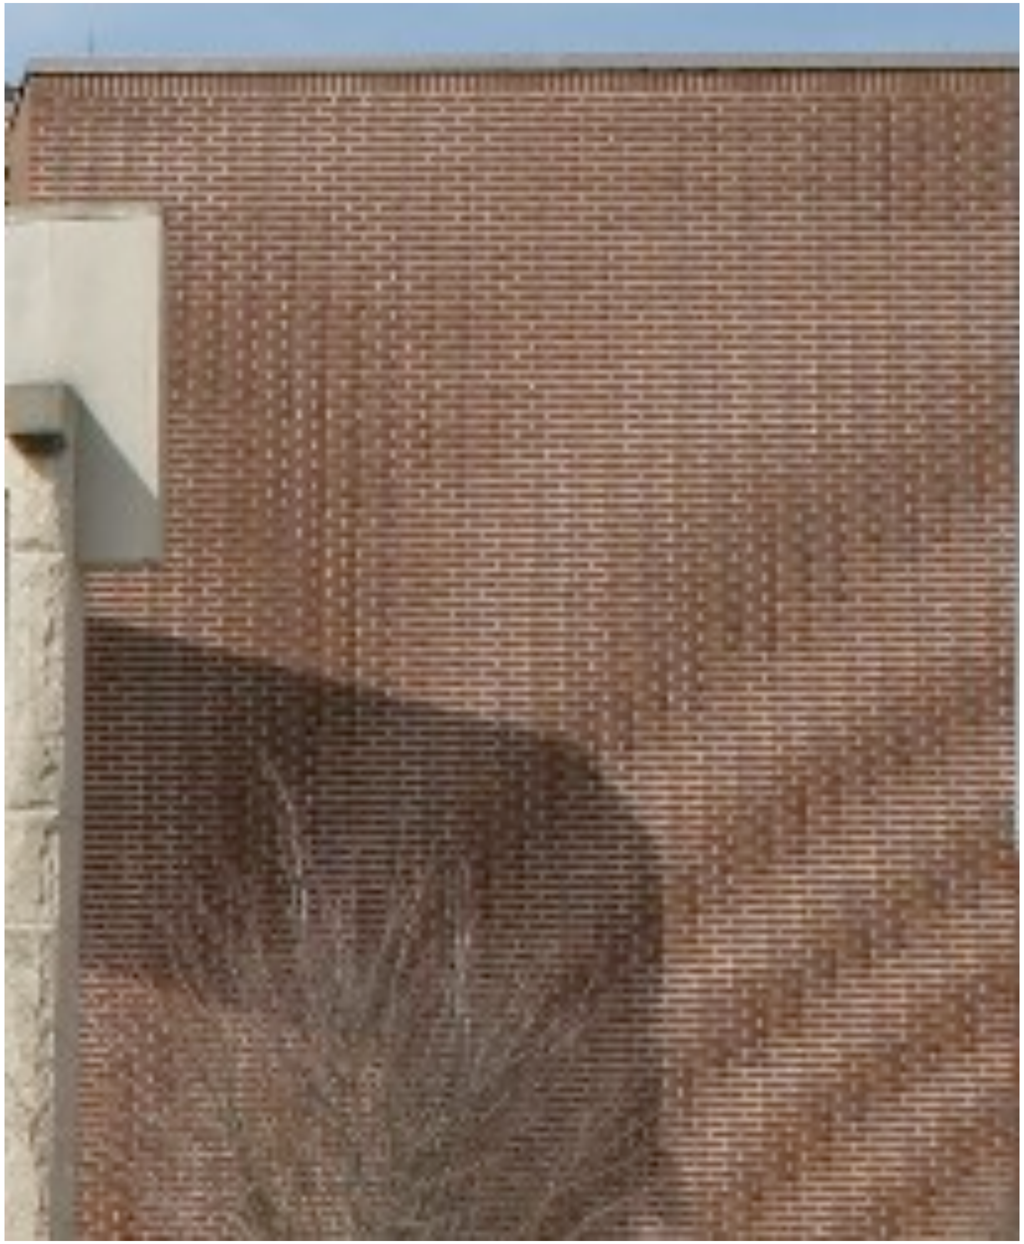

Intaglio (rotogravure printing)

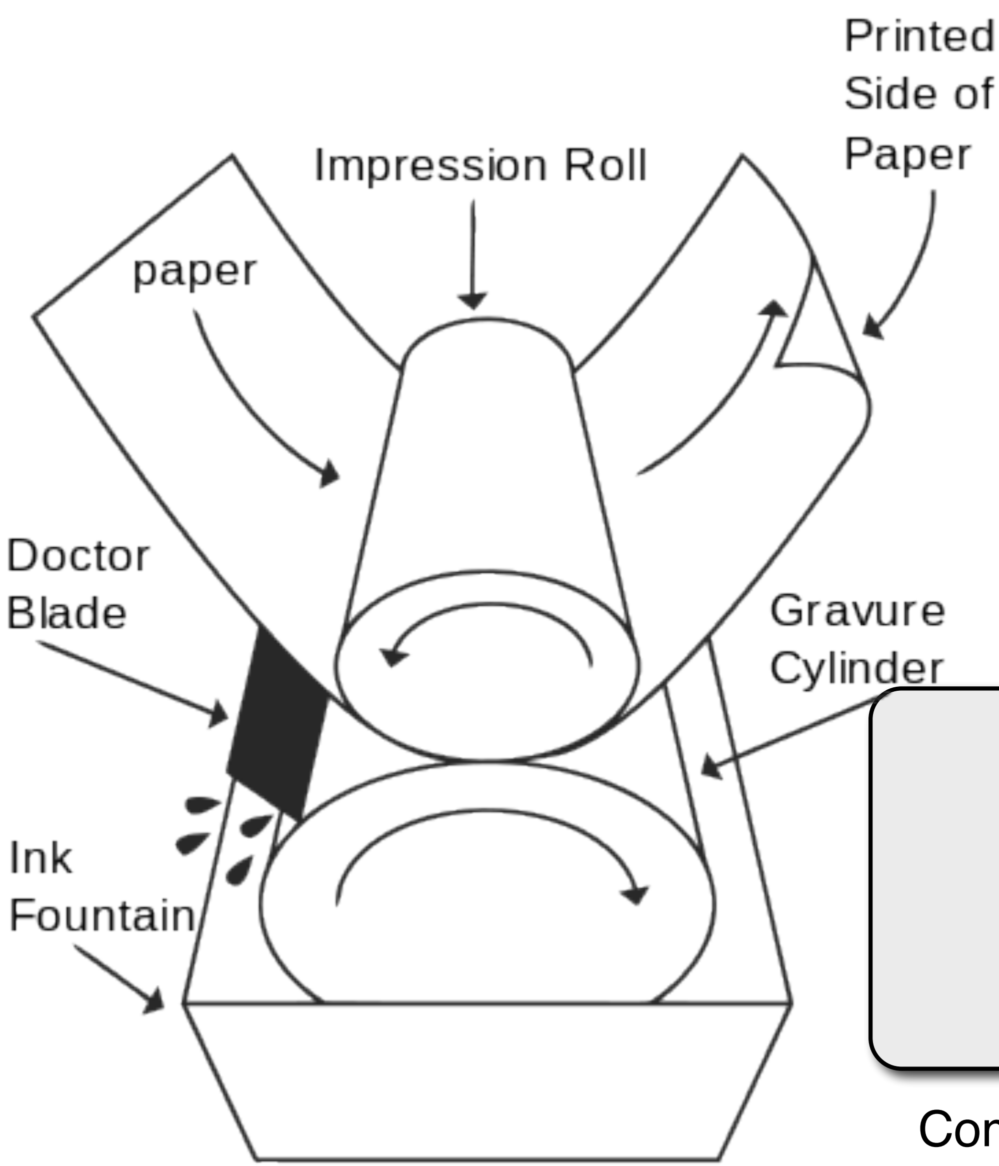

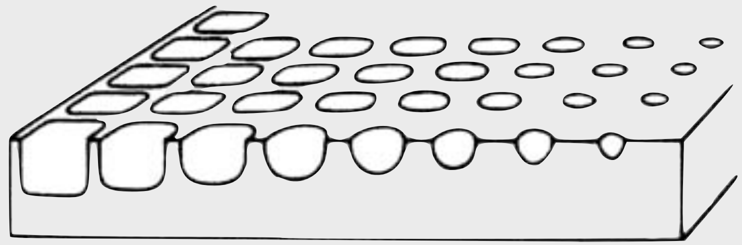

etching with diamond to make cells
with varying size and/or depth

Commonly used for mass-printed materials

"Bad" aliasing in printing applications

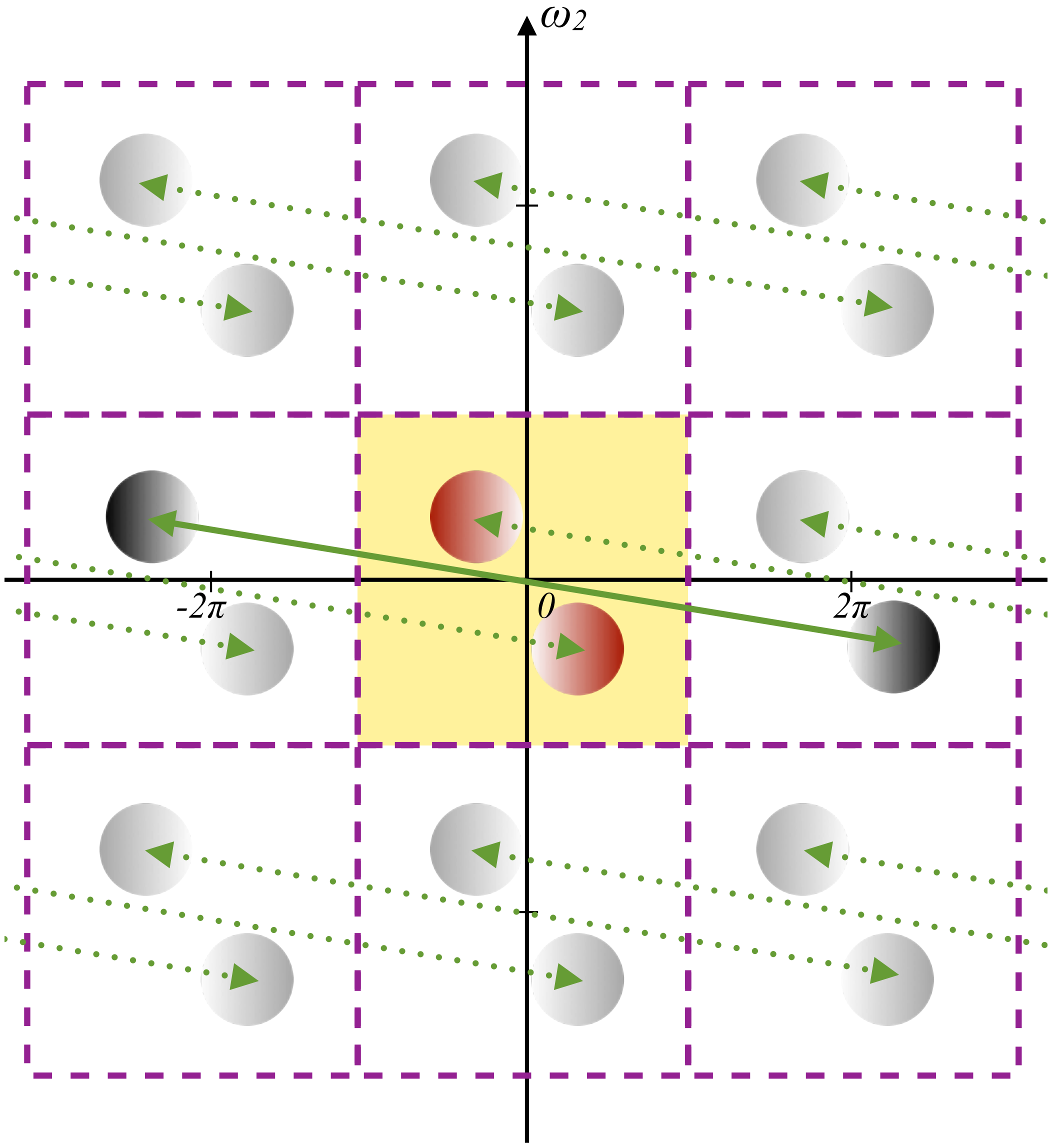

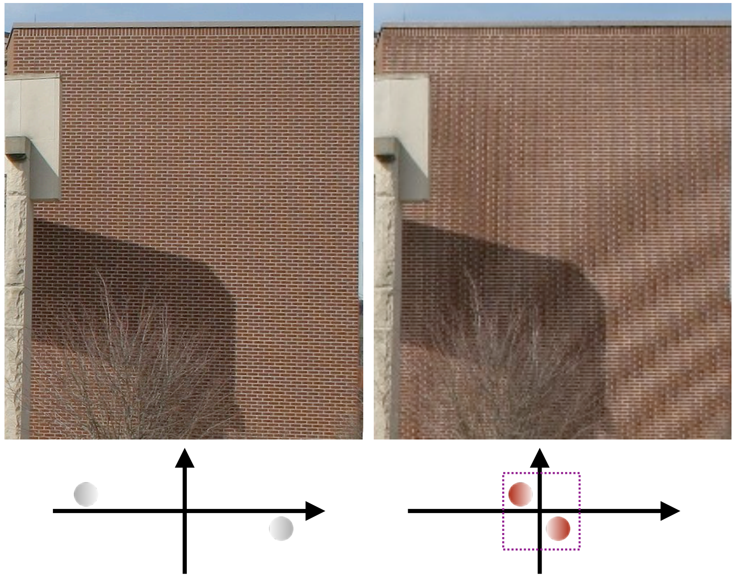

Aliasing can also be **good** such as for optical illusions, security printing, measurements of small displacements, ...

“Good” aliasing demo

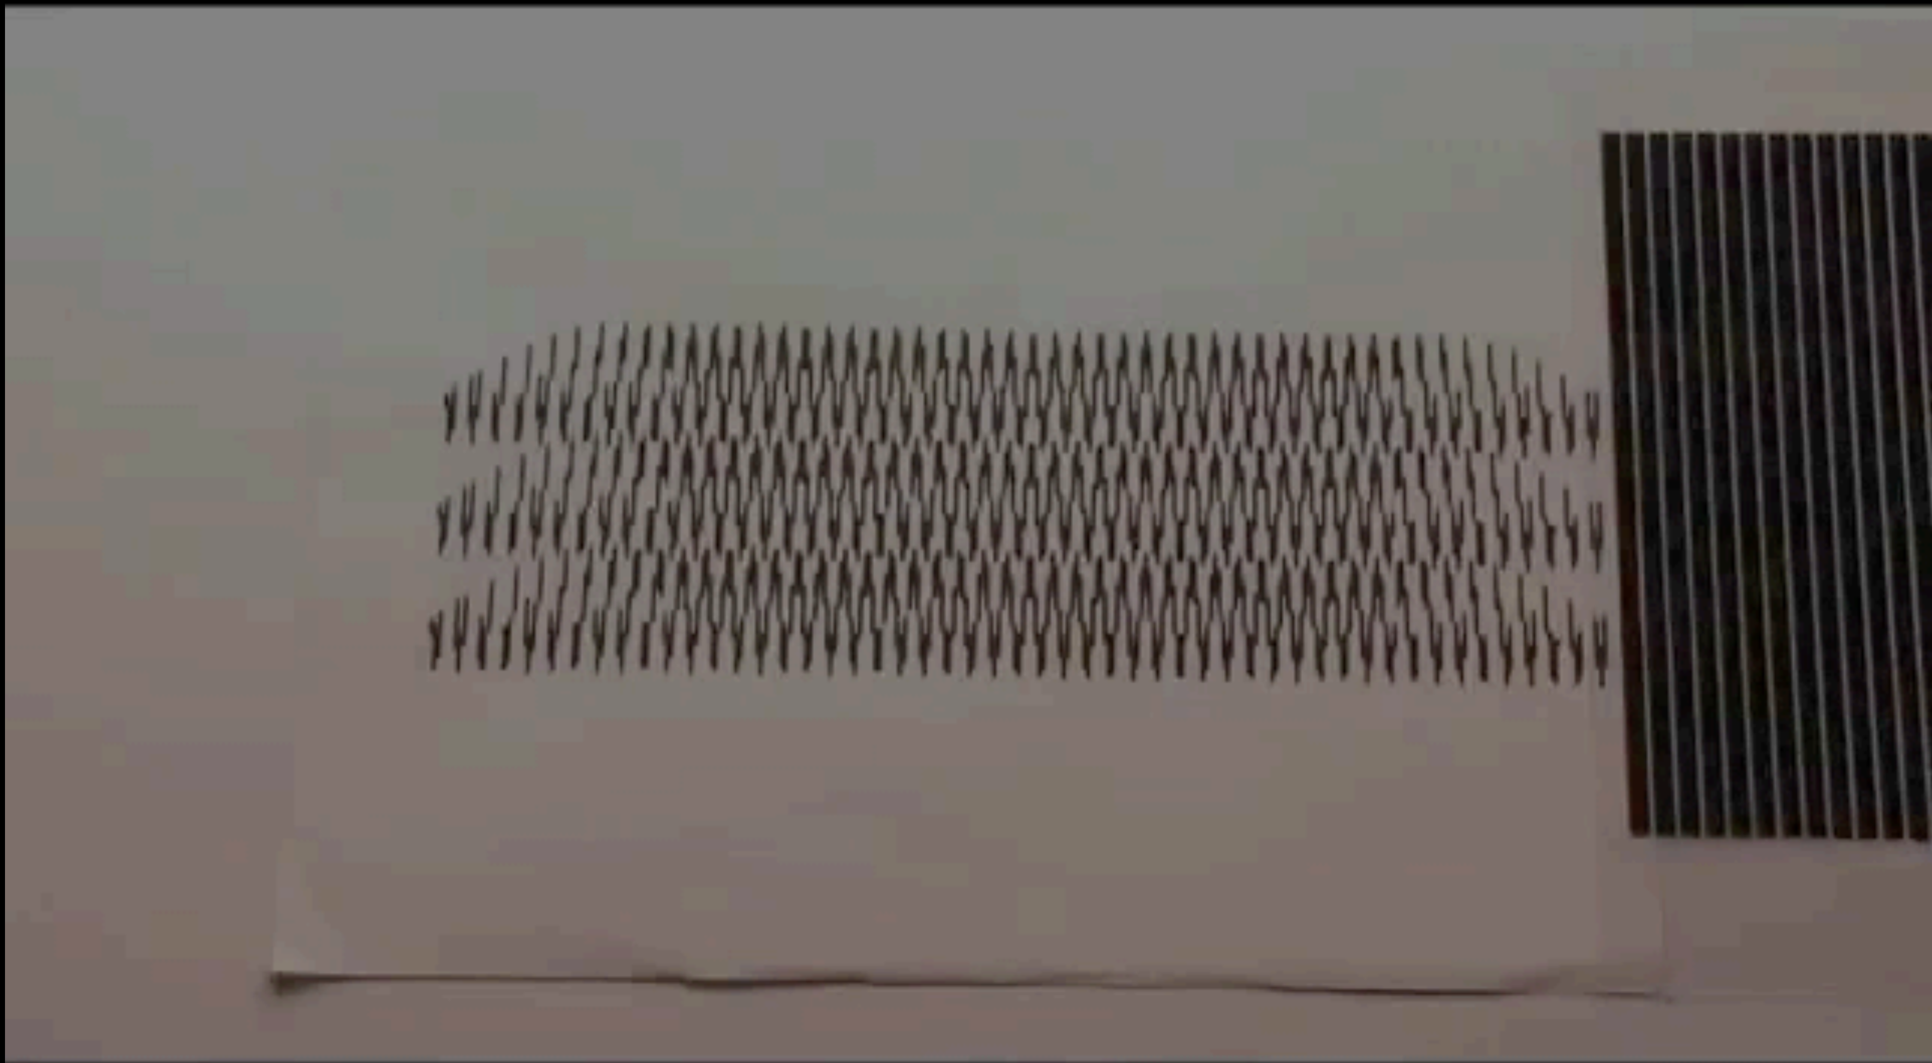

http://www.youtube.com/watch?v=lvvcRdwNhGM

Effect of reducing spatial resolution

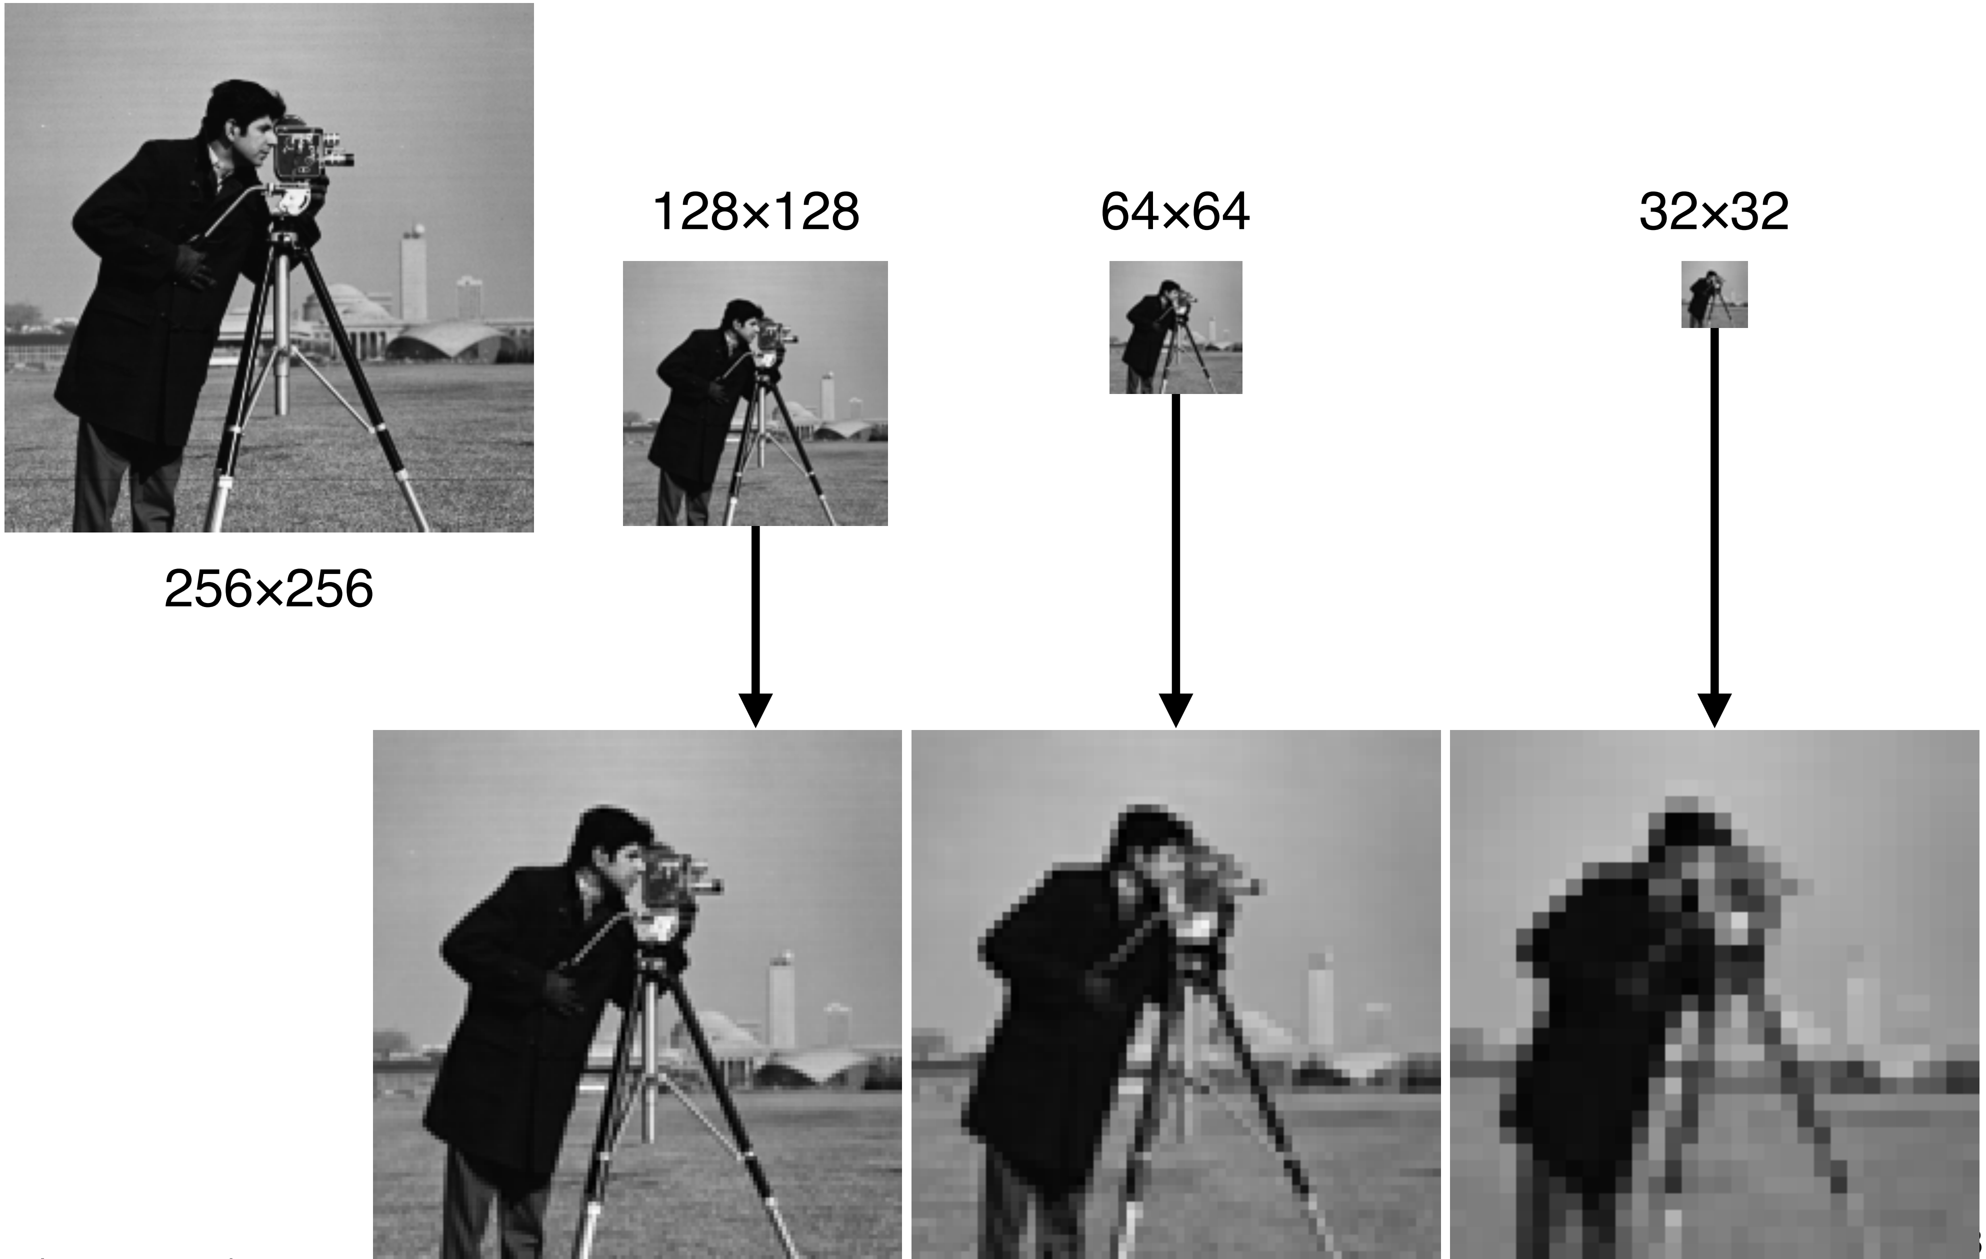

2.3 IMAGE QUANTIZATION

- Quantizer specification
- Histogram
- Uniform quantization
- Minimum error (Lloyd-Max) quantization
- Grayscale versus resolution
- Dithering

Quantizer specification

- Real-valued intensity value $f = f(\boldsymbol{x})$

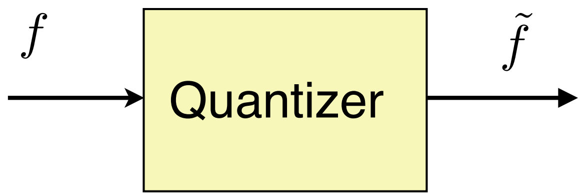

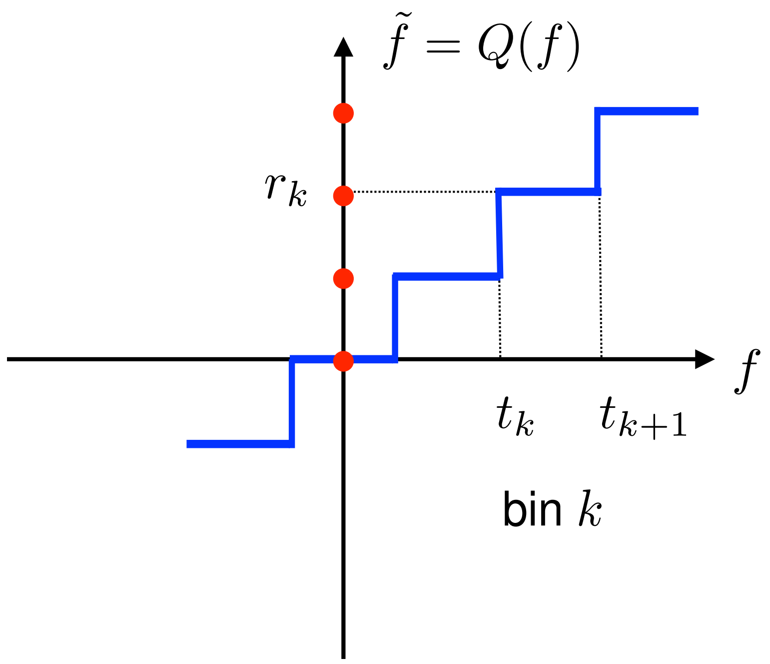

- Quantization thresholds: $t_k, \quad k = 0, \dots, K$

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

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

Histogram

- Gray-level probability-density function

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

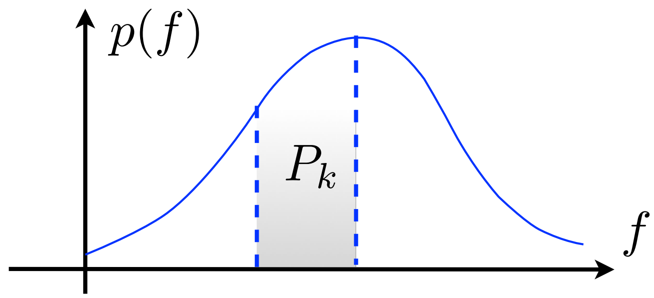

- Normalization: $\int_{-\infty}^{+\infty} p(f)\,\mathrm{d}f = 1$
- Mean: $\mu_f = E\{f\} = \int_{-\infty}^{+\infty} f \cdot p(f)\,\mathrm{d}f$
- Variance: $\sigma_f^2 = E\left\{(f-\mu_f)^2\right\} = \int_{-\infty}^{+\infty} (f-\mu_f)^2 p(f)\,\mathrm{d}f$

- Quantized histogram

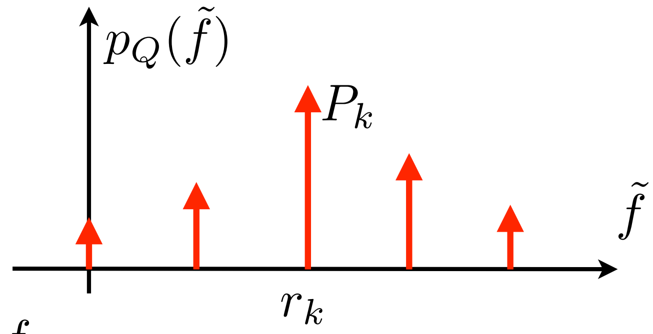

$$p(f) \quad \to \quad p_Q(\tilde{f}) = \sum_{k=0}^{K-1} P_k\,\delta(\tilde{f} - r_k)$$

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

- Entropy

Prior quantization: $H(f) = -\int_{-\infty}^{+\infty} p(f) \log\left(p(f)\right)\,\mathrm{d}f \quad \Rightarrow$ After quantization: $H_Q(\tilde{f}) = -\sum_k P_k \log_2 P_k$

Histogram

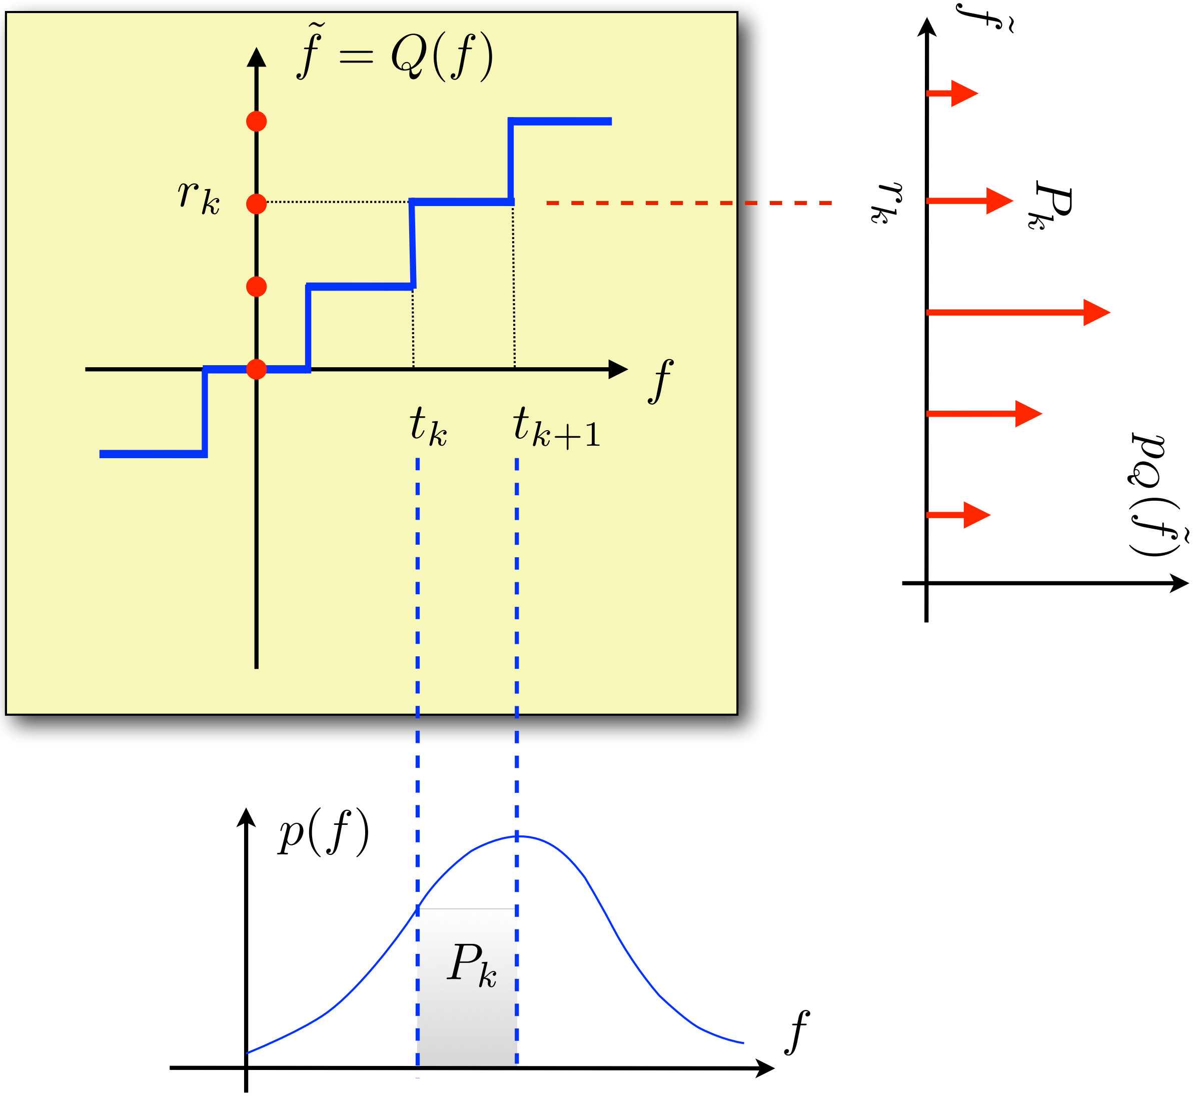

Uniform quantization

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

Typically:

0-255 (256 grayscale levels)

0-1 (binary)

Pixel budget:

8 bits

1 bit

Note: Human visual system can only distinguish about 60 shades of gray

- Quantization error

$$E\left\{(f-\tilde{f})^2\right\} = \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} (f - r_k)^2 p(f)\,\mathrm{d}f \cong \sum_{k=0}^{K-1} P_k \int_{-\Delta/2}^{+\Delta/2} e^2 \frac{1}{\Delta}\,\mathrm{de} = \frac{\Delta^2}{12}$$

High-resolution hypothesis (large K)

$$p(f) \cong \text{const} = \frac{P_k}{\underbrace{t_{k+1} - t_k}_{\Delta}} \quad \text{for} \quad t_k \leqslant f < t_{k+1}$$

Effect of reducing the number of gray levels

256 levels | 16 levels

8 bits | 4 bits

2 bits | 1 bit

4 levels | 2 levels

Non-uniform quantization and segmentation

Search for the "optimal" threshold to segment images

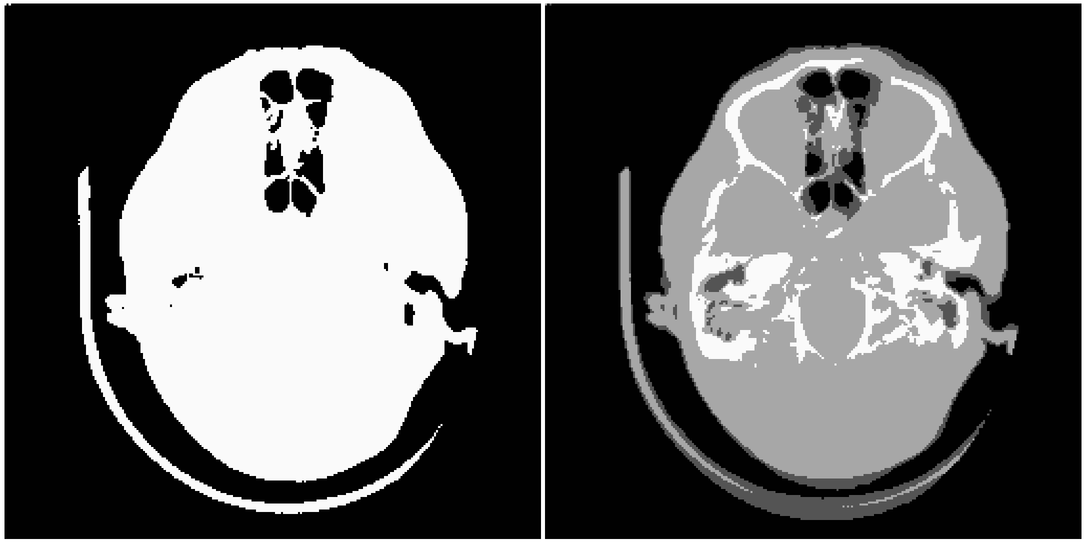

Minimum-error solution: K=2 K=4

Minimum-error (Lloyd-Max) quantization

Goal: minimize

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

- Differentiate with respect to t_k

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

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

(mid-point solution)

- Differentiate with respect to r_k

$$\frac{\partial \varepsilon^2}{\partial r_k} = -2\int_{t_k}^{t_{k+1}} (f - r_k)p(f)\,\mathrm{d}f = 0 \quad \Rightarrow \quad r_k = \frac{\displaystyle\int_{t_k}^{t_{k+1}} f\,p(f)\,\mathrm{d}f}{\displaystyle\int_{t_k}^{t_{k+1}} p(f)\,\mathrm{d}f}$$

(partial mean)

Iterative optimization algorithm (*K* means)

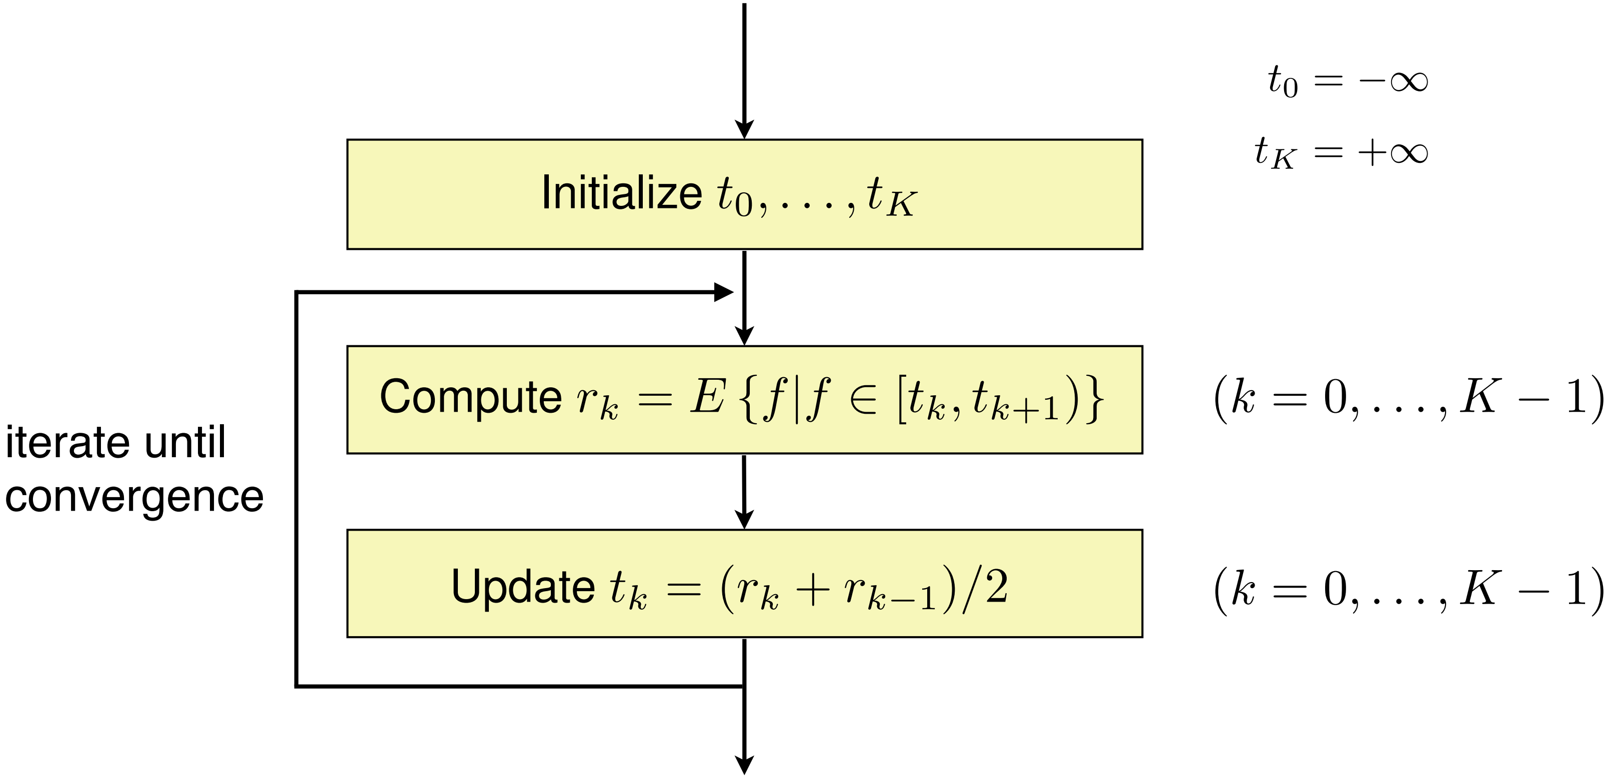

- **Efficient implementation**
  - Partial means can be computed directly from the histogram
  - Recursive update

- **Generalizations**
  - Non-quadratic cost function
  - Multivariate extension (e.g., RGB)
  - Iterative clustering (K-means algorithm)

Minimum-error quantization

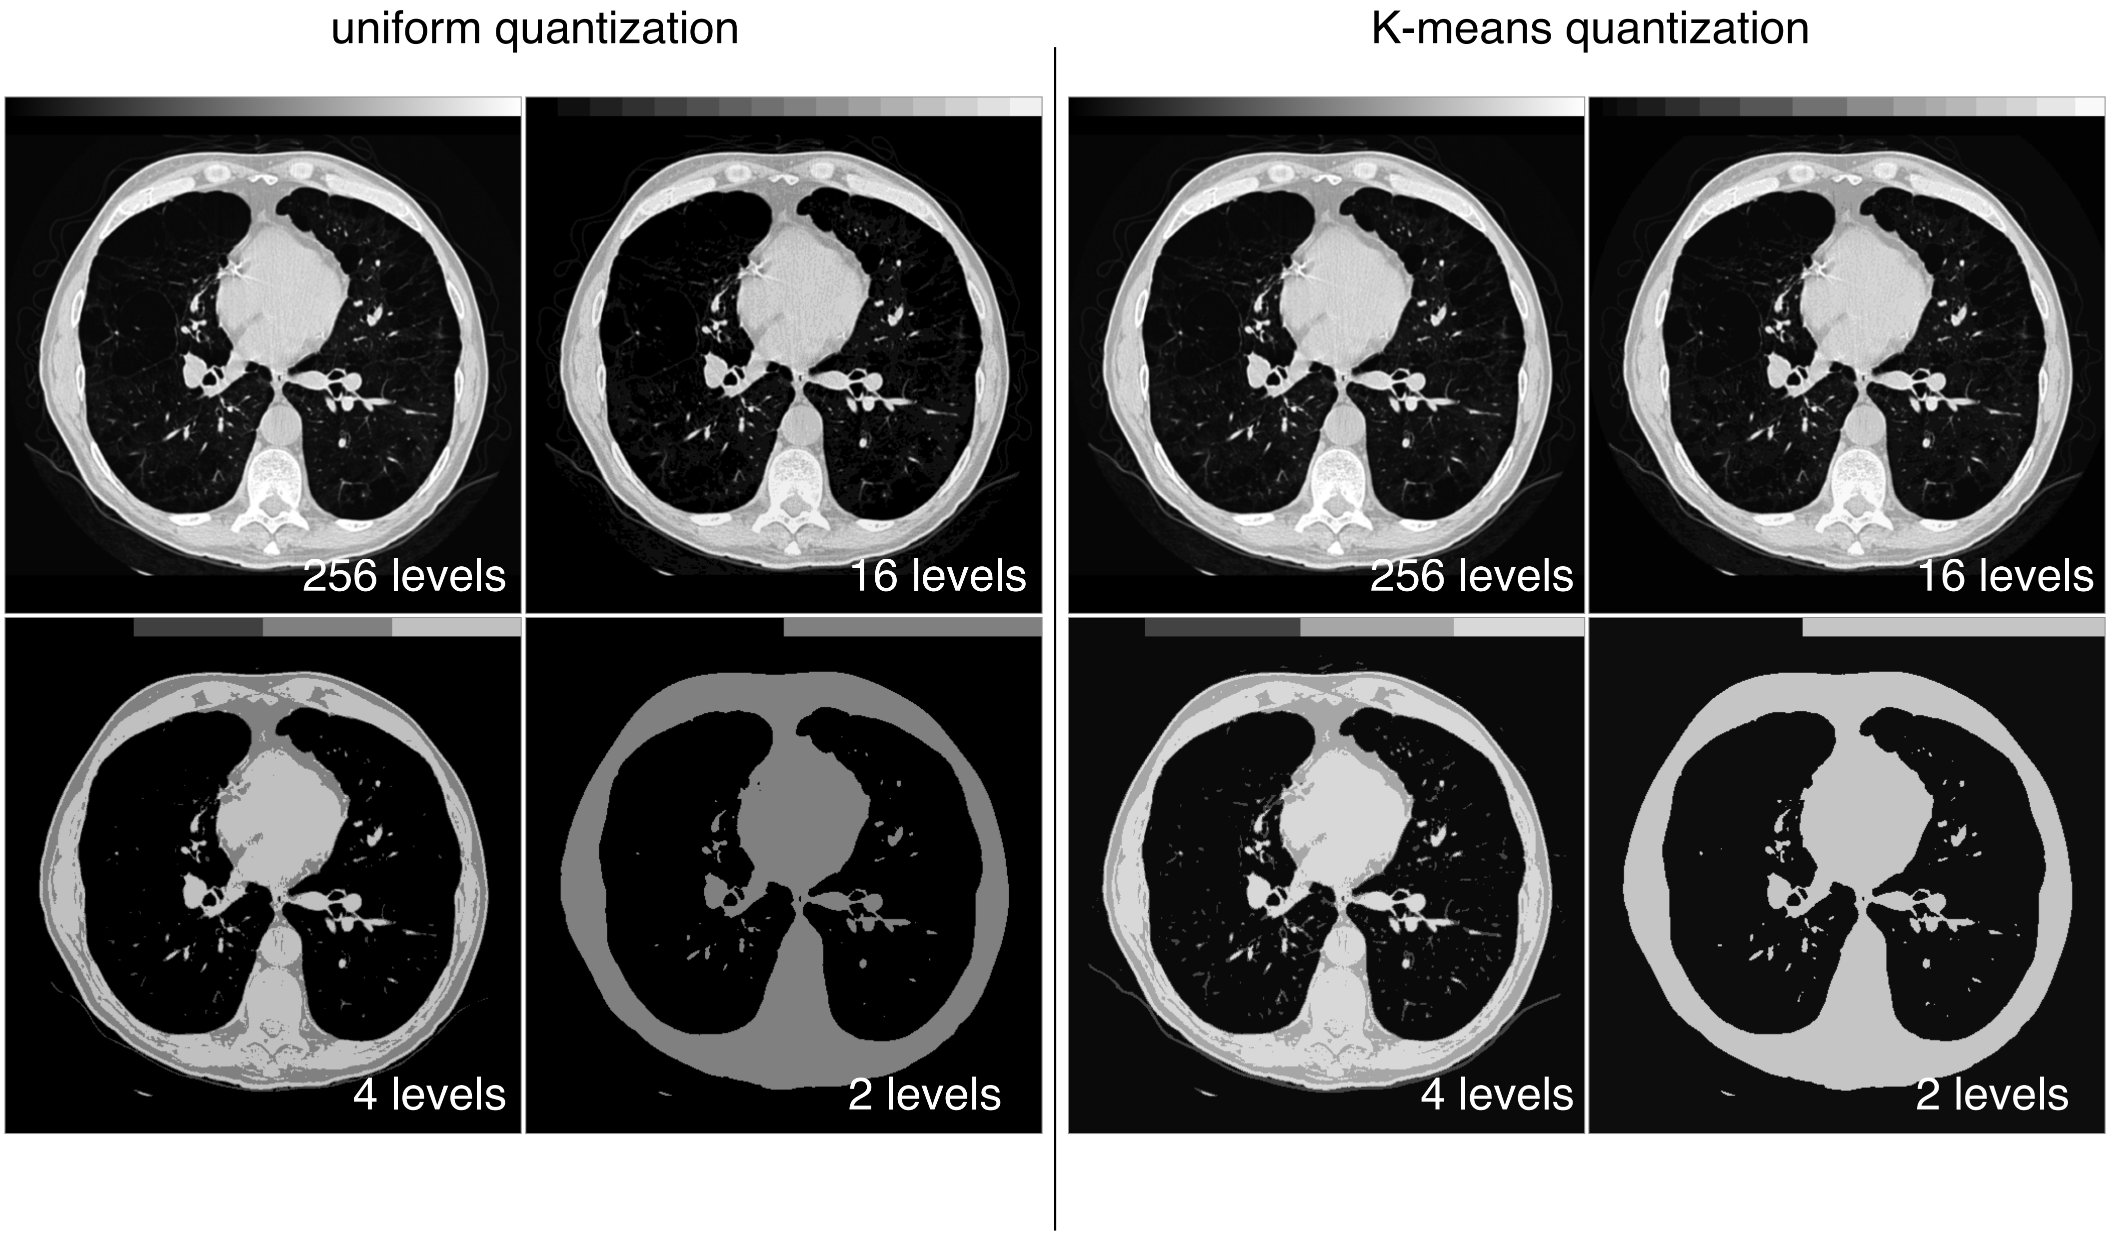

Minimum-error quantization

- Further properties of Lloyd-Max quantizer
  - Unbiased image estimate: $E\{\tilde{f}\} = E\{f\}$
  - Orthogonality of the approximation error: $E\{\tilde{f}(f - \tilde{f})\} = 0$

$$\Rightarrow \quad E\{(f-\tilde{f})^2\} = E\{f^2\} - E\{\tilde{f}^2\}$$

$$E\{f^2\} = E\{\tilde{f}^2\} + E\{(f-\tilde{f})^2\}$$

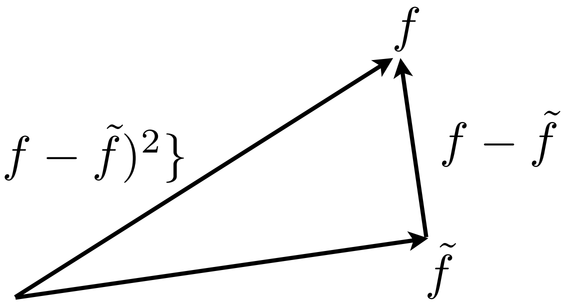

Proof:

$$\begin{aligned}
\sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \tilde{f}(f-\tilde{f})p(f)\,\mathrm{d}f &= \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} r_k(f-r_k)p(f)\,\mathrm{d}f \\
&= \sum_{k=0}^{K-1} r_k \underbrace{\int_{t_k}^{t_{k+1}} f p(f)\,\mathrm{d}f - r_k^2 \int_{t_k}^{t_{k+1}} p(f)\,\mathrm{d}f}_{\text{Lloyd-Max optimality}}
\end{aligned}$$

Grayscale versus resolution

128 x 128 x 4 bits

32 x 32 x 16 bits

256 x 256 x 1 bits

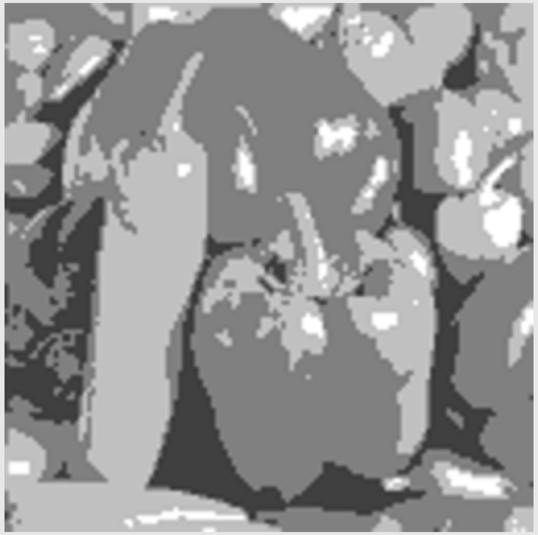

128 x 128 x 4 bits

64 x 64 x 16 bits

bit budget:65'536 bits

Need for binary images

- Some devices can only render binary output (e.g., printers, fax, etc)
  ink or no ink!
- Lucky coincidence: the human visual system locally integrates B&W information and sees the “average”
- Exploit trade-off between spatial resolution and grayscale resolution
- Implemented by “Raster Image Processor” (RIP) in printing systems

Can we do better than simple thresholding?

Dithering

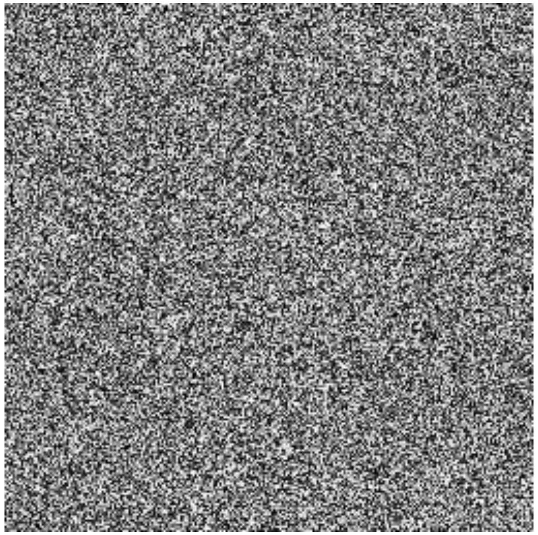

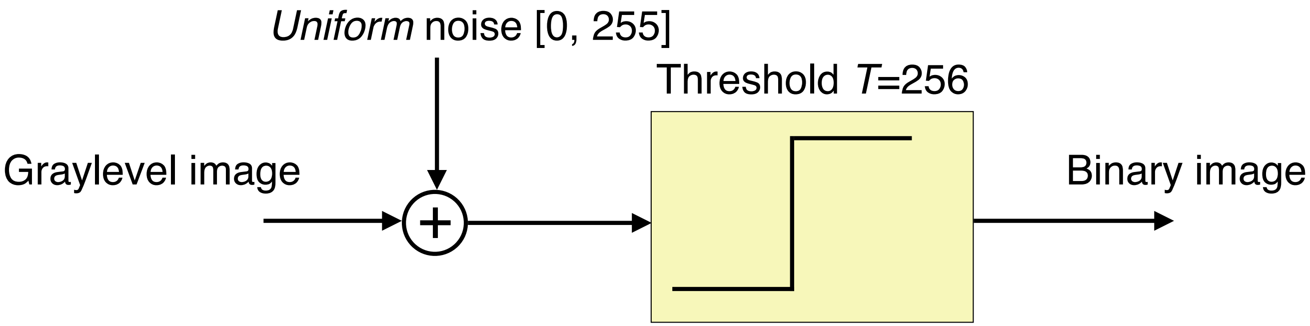

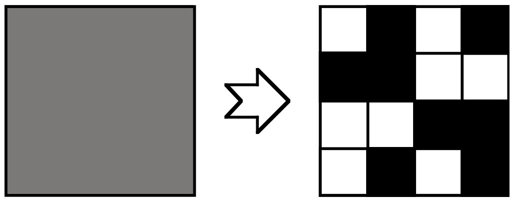

grayscale ramp

Ordered dithering

- Oversample (*e.g.*, factor 2: $256 \times 256 \quad \Rightarrow \quad 512 \times 512$) $\quad \Rightarrow \quad$ pixel replication

- Replace macropixel with dithering pattern according to grayscale value

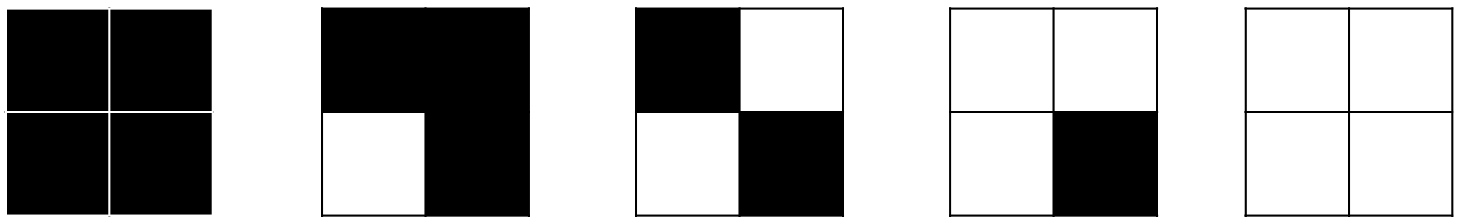

0 1 2 3 4

- Bayer index matrices are defined recursively

$$H_2 = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix} \quad H_{2N} = \begin{bmatrix} 4H_N + 2 & 4H_N + 1 \\ 4H_N & 4H_N + 3 \end{bmatrix}$$

corresponding threshold matrices for grayscale values in $[0, 255]$

$$T_N = 256 \frac{H_N + 0.5}{N^2} \quad \text{for example: } T_2 = \begin{bmatrix} 160 & 96 \\ 32 & 223 \end{bmatrix}$$

below threshold indicates where to put ink; e.g., value 190 for T_2 results into case 3

Dithering example

grayscale image

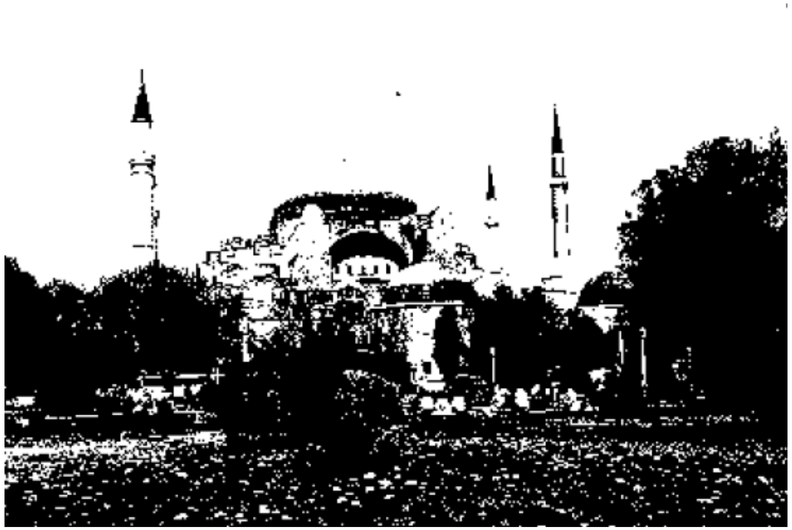

fixed threshold

random dithering

ordered dithering

[source images: Charles A. Bouman, Purdue University]

2.4 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 pulsation $\omega_{\max}/2$.
- 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 form. Monochrome monitors typically display 256 levels of gray.
- Alternatively, the quantization steps may be selected to minimize the approximation error (Max-Lloyd quantizer). This also yields a most effective segmentation algorithm (minimum error thresholding).
- To some extent, grayscale resolution can be traded for spatial resolution. This tradeoff is the basis for halftoning algorithms.

Bibliography

- C. E. Shannon, "Communication in the Presence of Noise," *Proceedings of the I.R.E.*, vol. 37, pp. 10-21, 1949.
- M. Unser, "Sampling—50 Years After Shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569-587, 2000.