THE UNIVERSITY OF BRITISH COLUMBIA

CPSC 425: Computer Vision

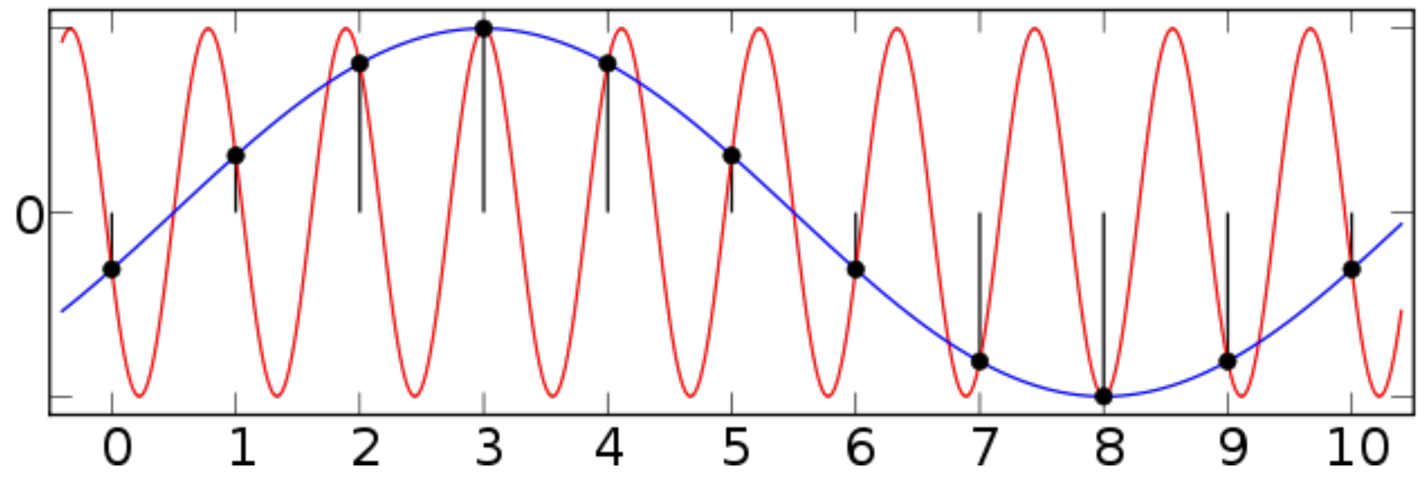

Image Credit: https://en.wikibooks.org/wiki/Analog_and_Digital_Conversion/Nyquist_Sampling_Rate

Lecture 7: Sampling

(unless otherwise stated slides are taken or adopted from **Bob Woodham, Jim Little** and **Fred Tung**)

Menu for Today (**September 23, 2020**)

Topics:

— **Sampling** theory

— **Nyquist** rate

— Color **Filter Arrays**

— **Bayer** patterns

Redings:

— **Today's** Lecture: Forsyth & Ponce (2nd ed.) 4.4

— **Next** Lecture: Forsyth & Ponce (2nd ed.) 4.5, 4.6

Reminders:

— **Assignment 1:** Image Filtering and Hybrid Images due **September 30**th

— **Quiz 1**, **Quiz 2**, **Quiz 3** dates are posted — **Quiz 1** is Friday

— We have a new TA - **Ruolan** taking over **Ariel** (TA times will remain the same)

Today's "**fun**" Example: Face on the Moon

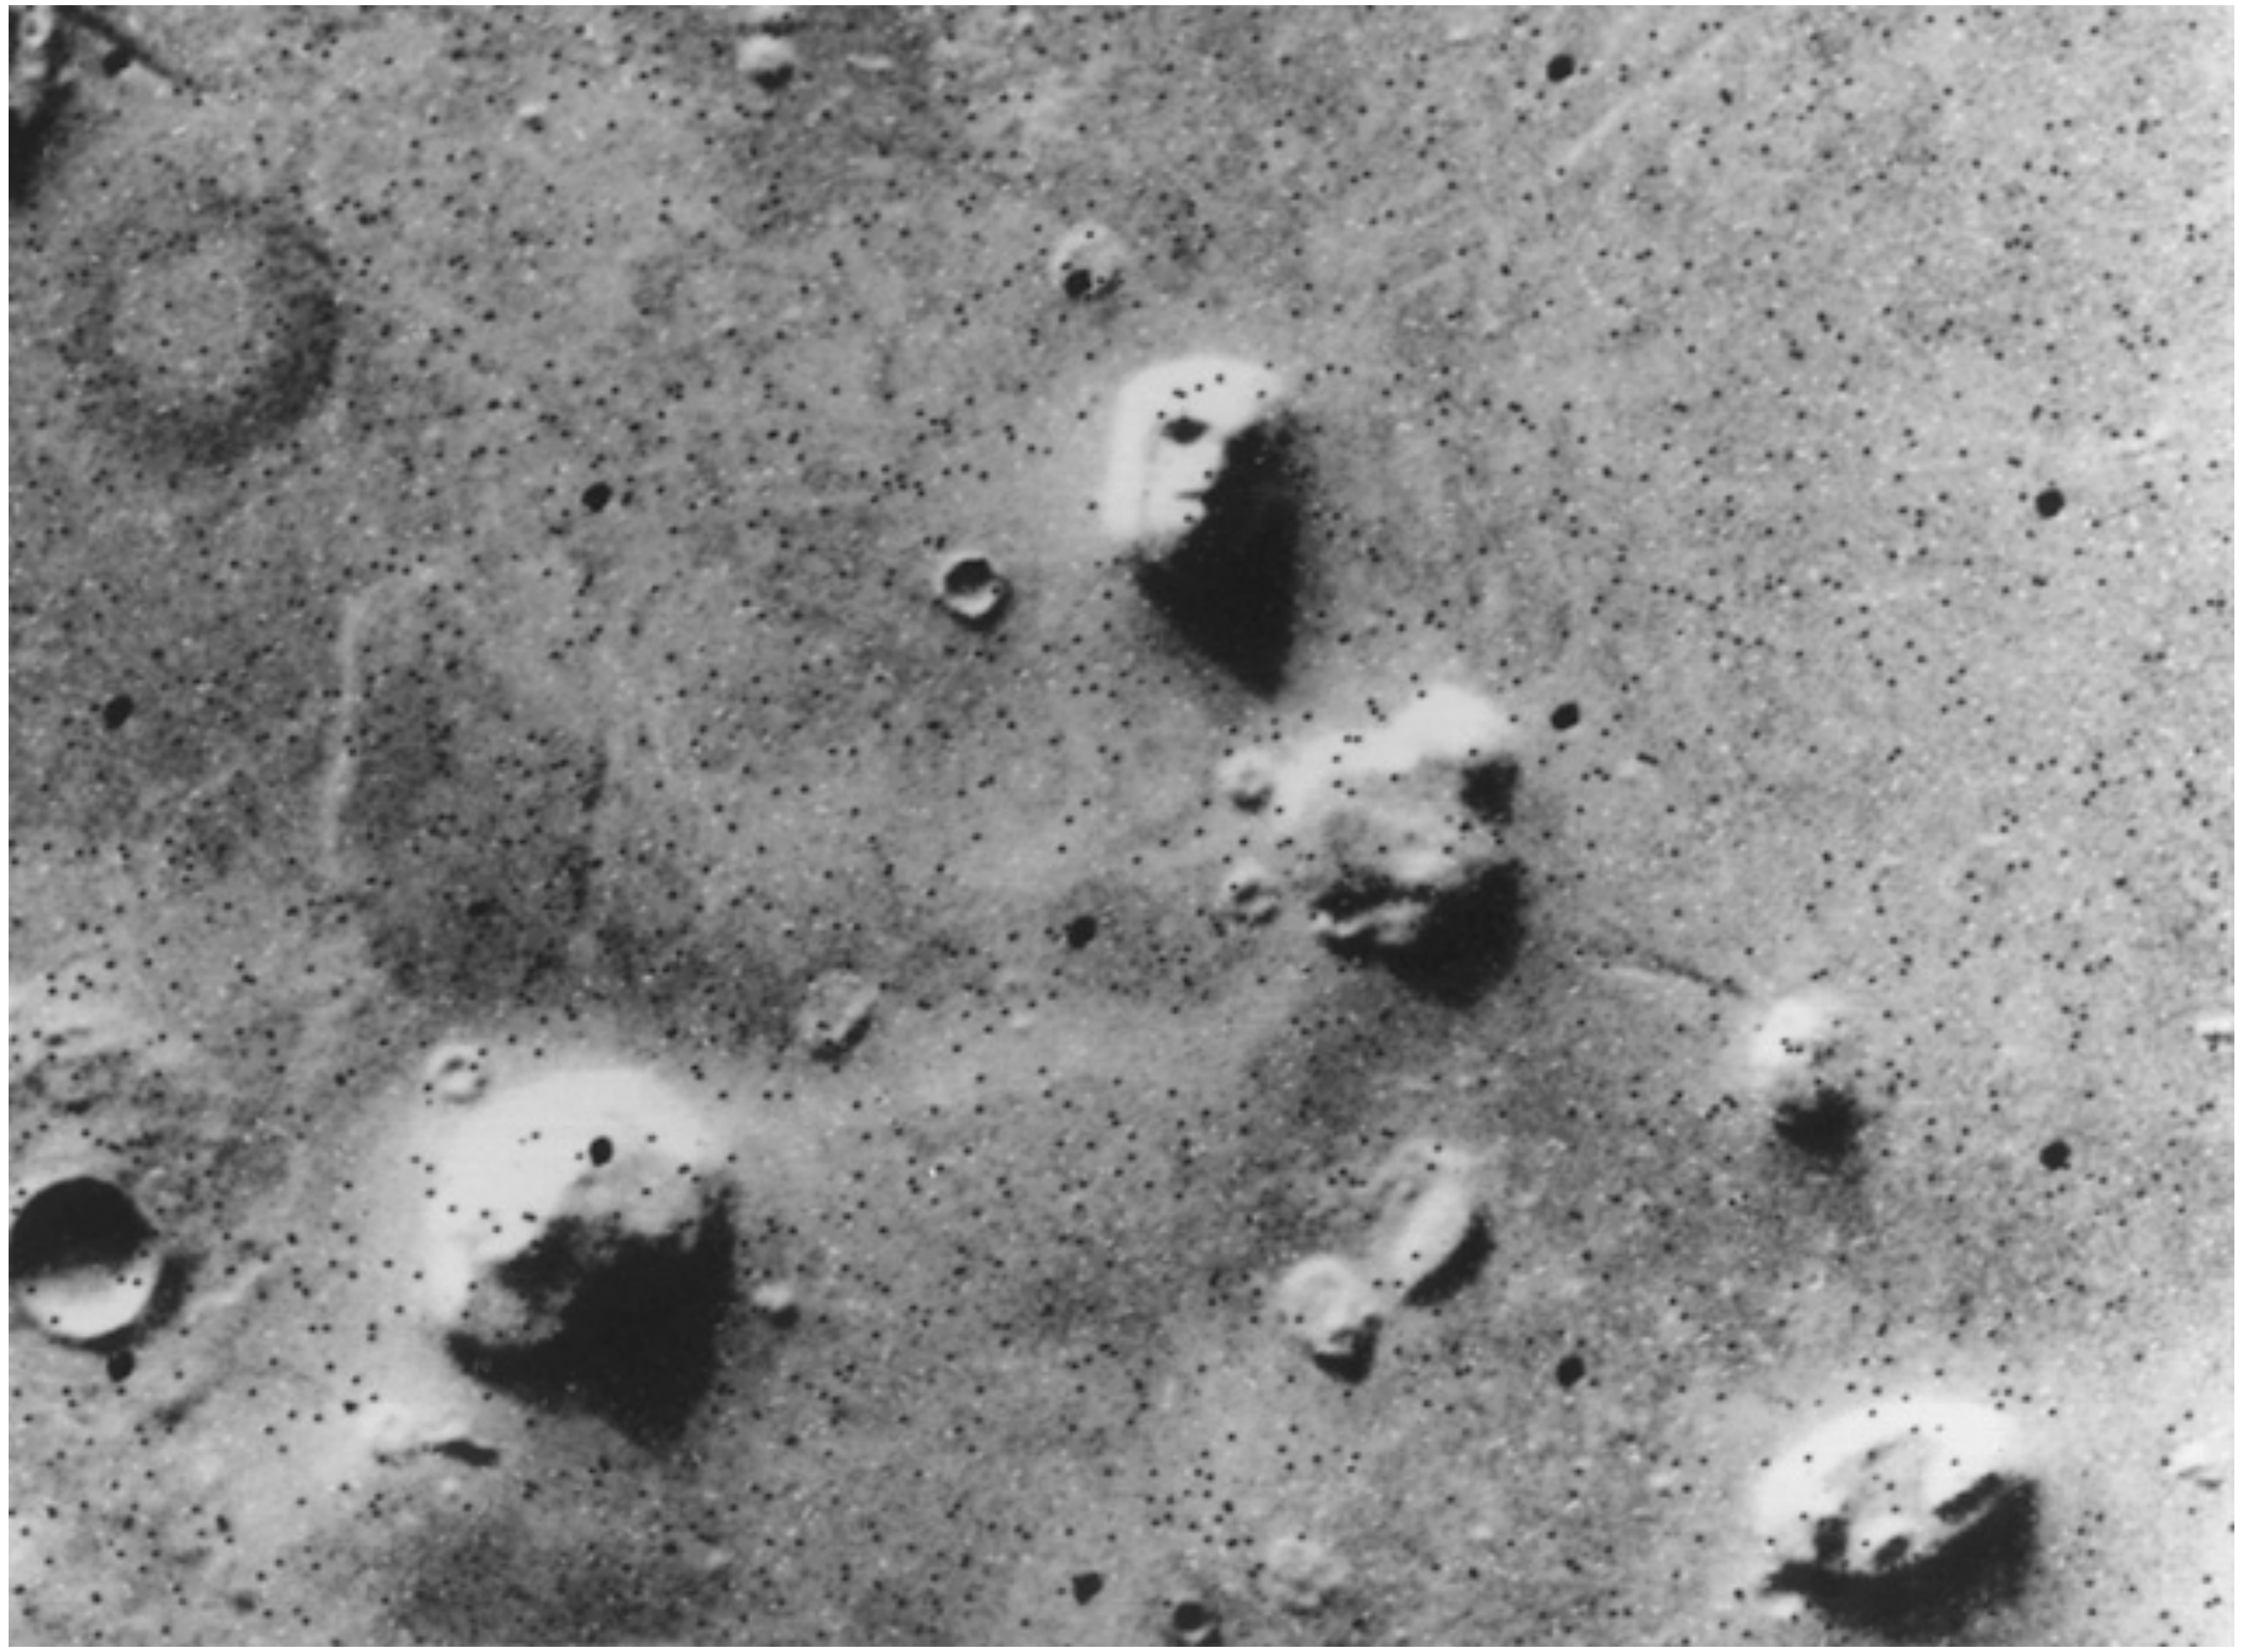

Image Credit: http://esamultimedia.esa.int/images/marsexpress/300-230906-3253-6-vk1-Cydonia_H.jpg

Today's "**fun**" Example: Face on the Moon

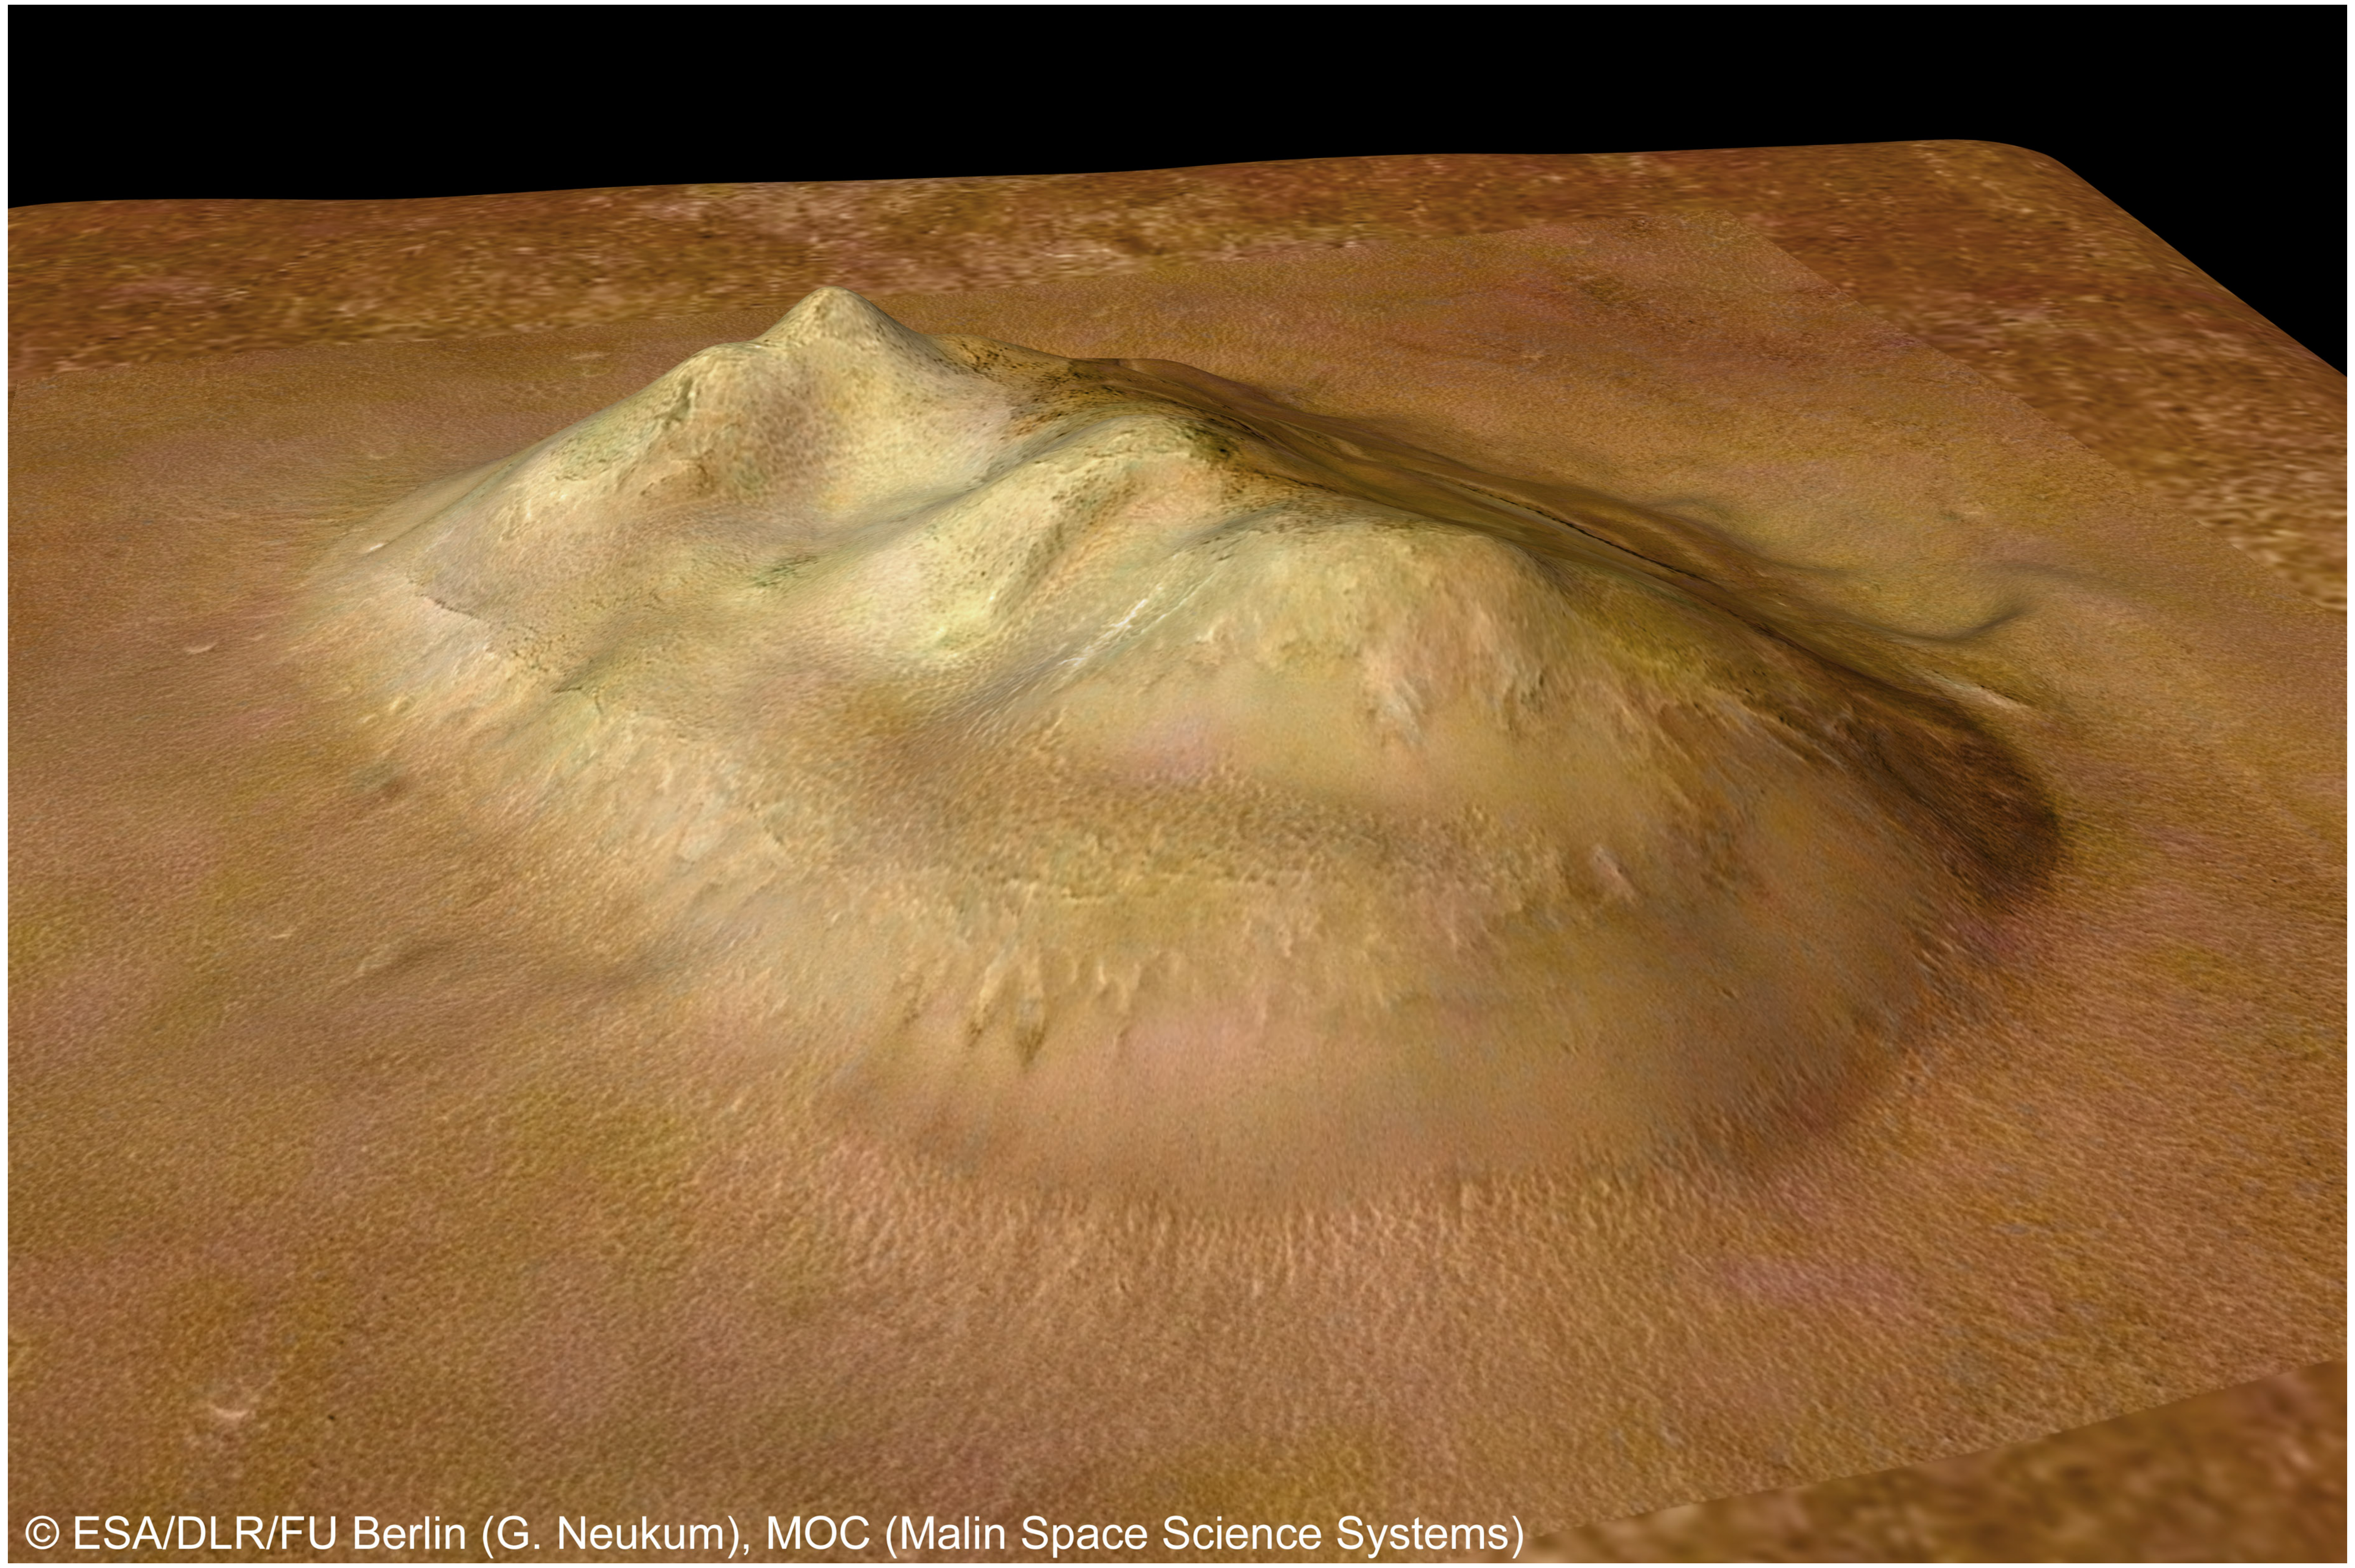

Image Credit: http://esamultimedia.esa.int/images/marsexpress/311-230906-3253-6-3d5-Cydonia_H.jpg

Today's "**fun**" Example: Tool for **Surrealists** Artists

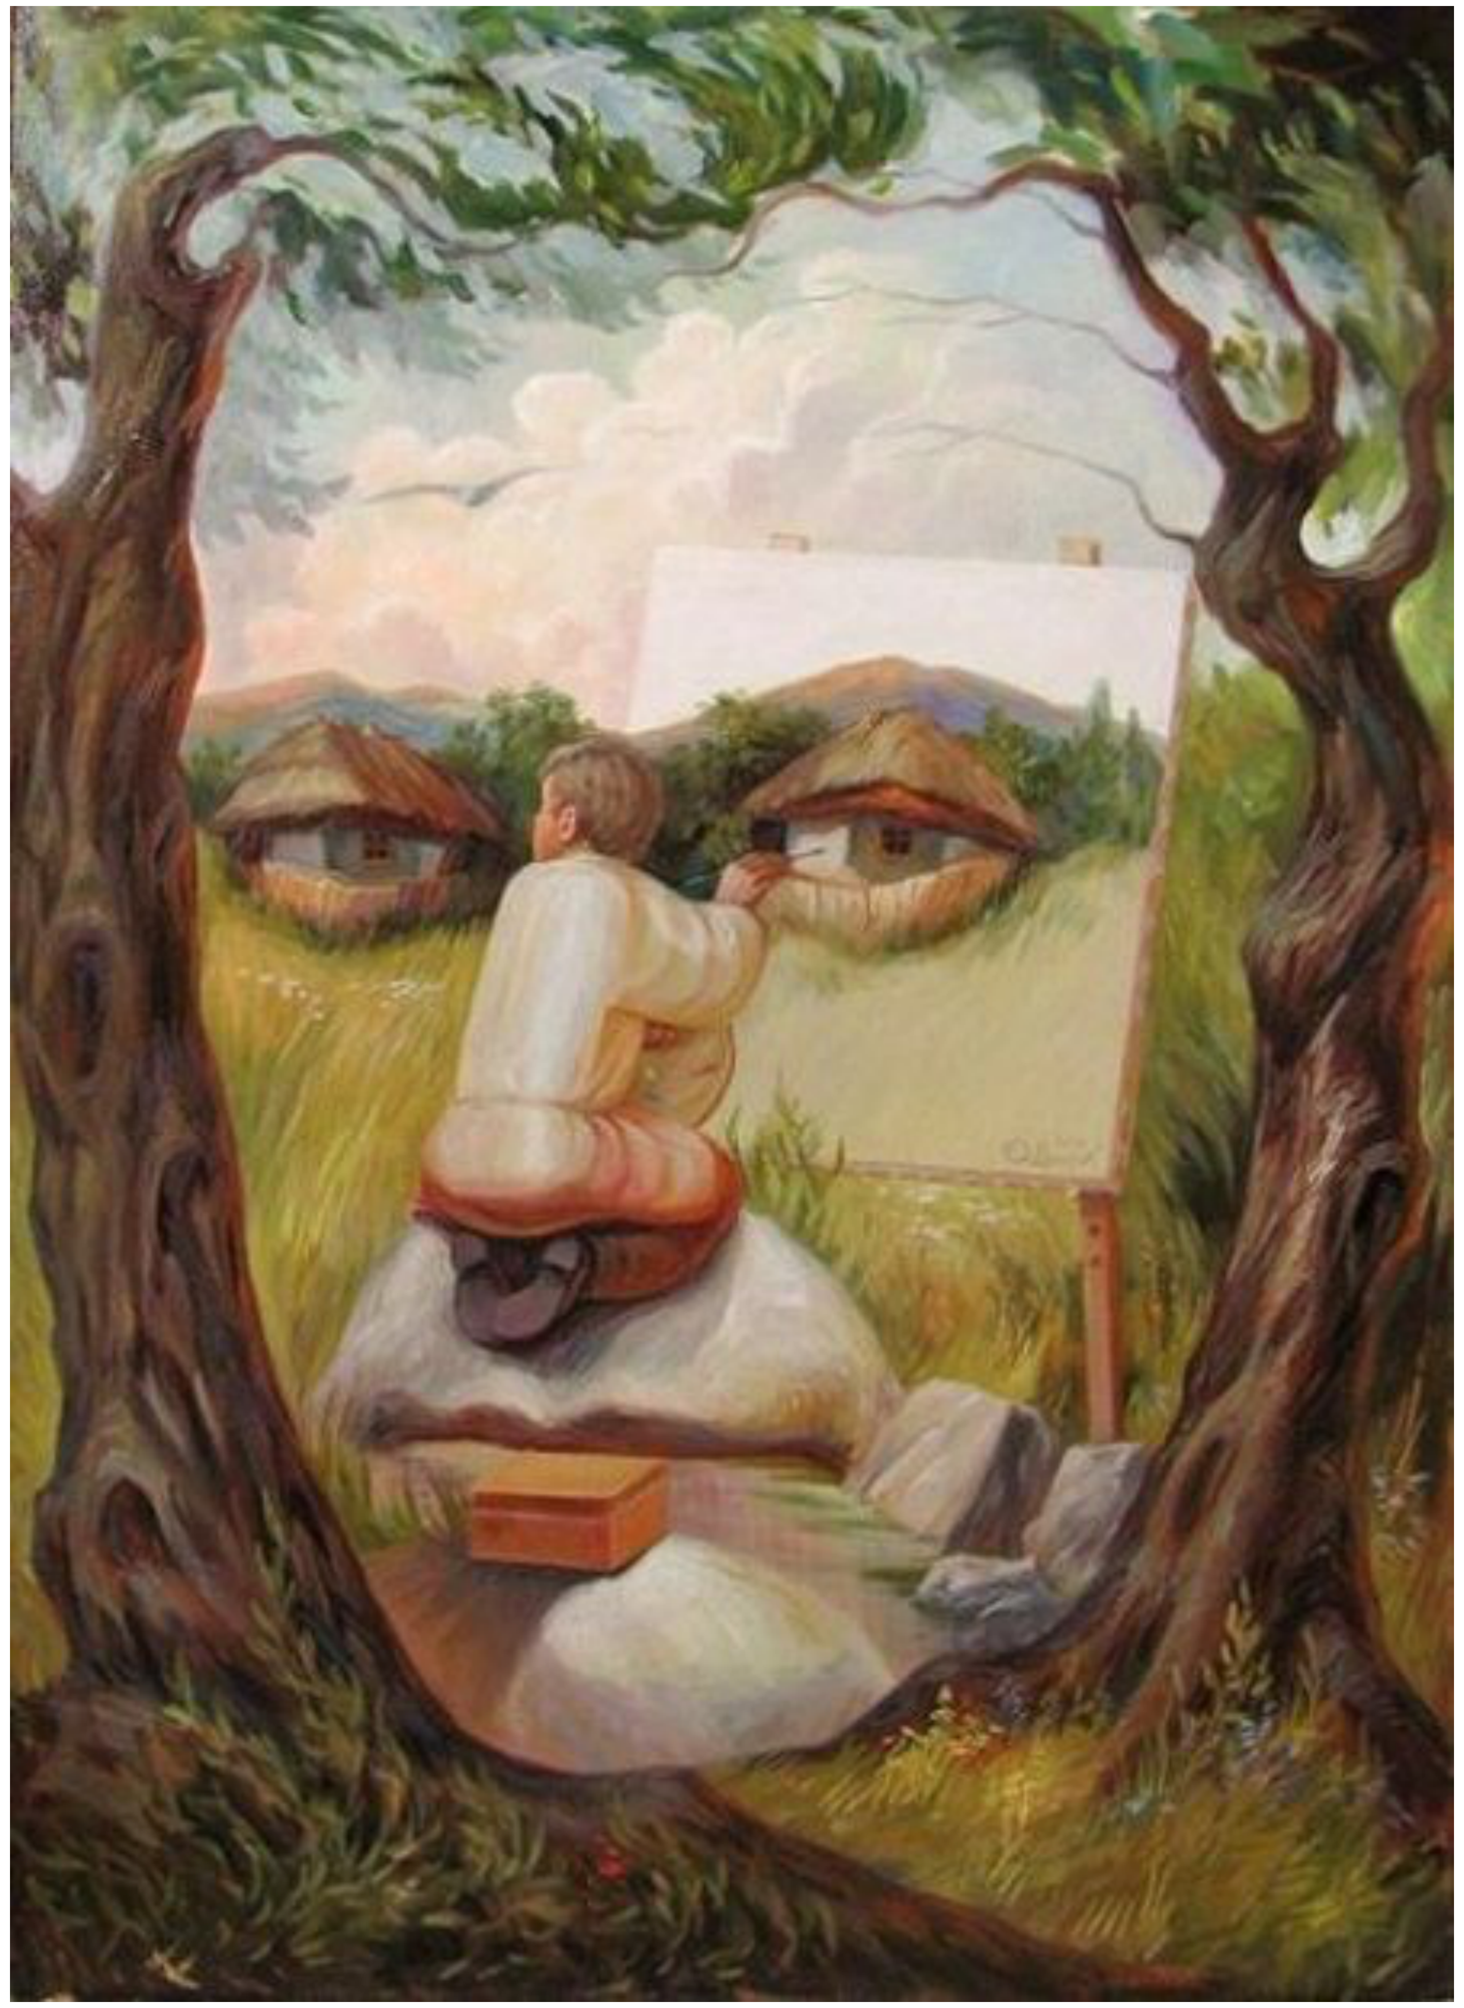

Oleg Shuplyak

Today's "**fun**" Example: Tool for **Surrealists** Artists

Artush Voskanyan

Lecture 5: Re-cap Non-linear Filters

We covered two three **non-linear filters**: Median, Bilateral, ReLU

Separability (of a 2D filter) allows for more efficient implementation (as two 1D filters)

Convolution is **associative** and **symmetric**

Convolution of a Gaussian with a Gaussian is another Gaussian

The **median filter** is a non-linear filter that selects the median in the neighbourhood

The **bilateral filter** is a non-linear filter that considers both spatial distance and range (intensity) distance, and has edge-preserving properties

Aside: Linear Filter with ReLU

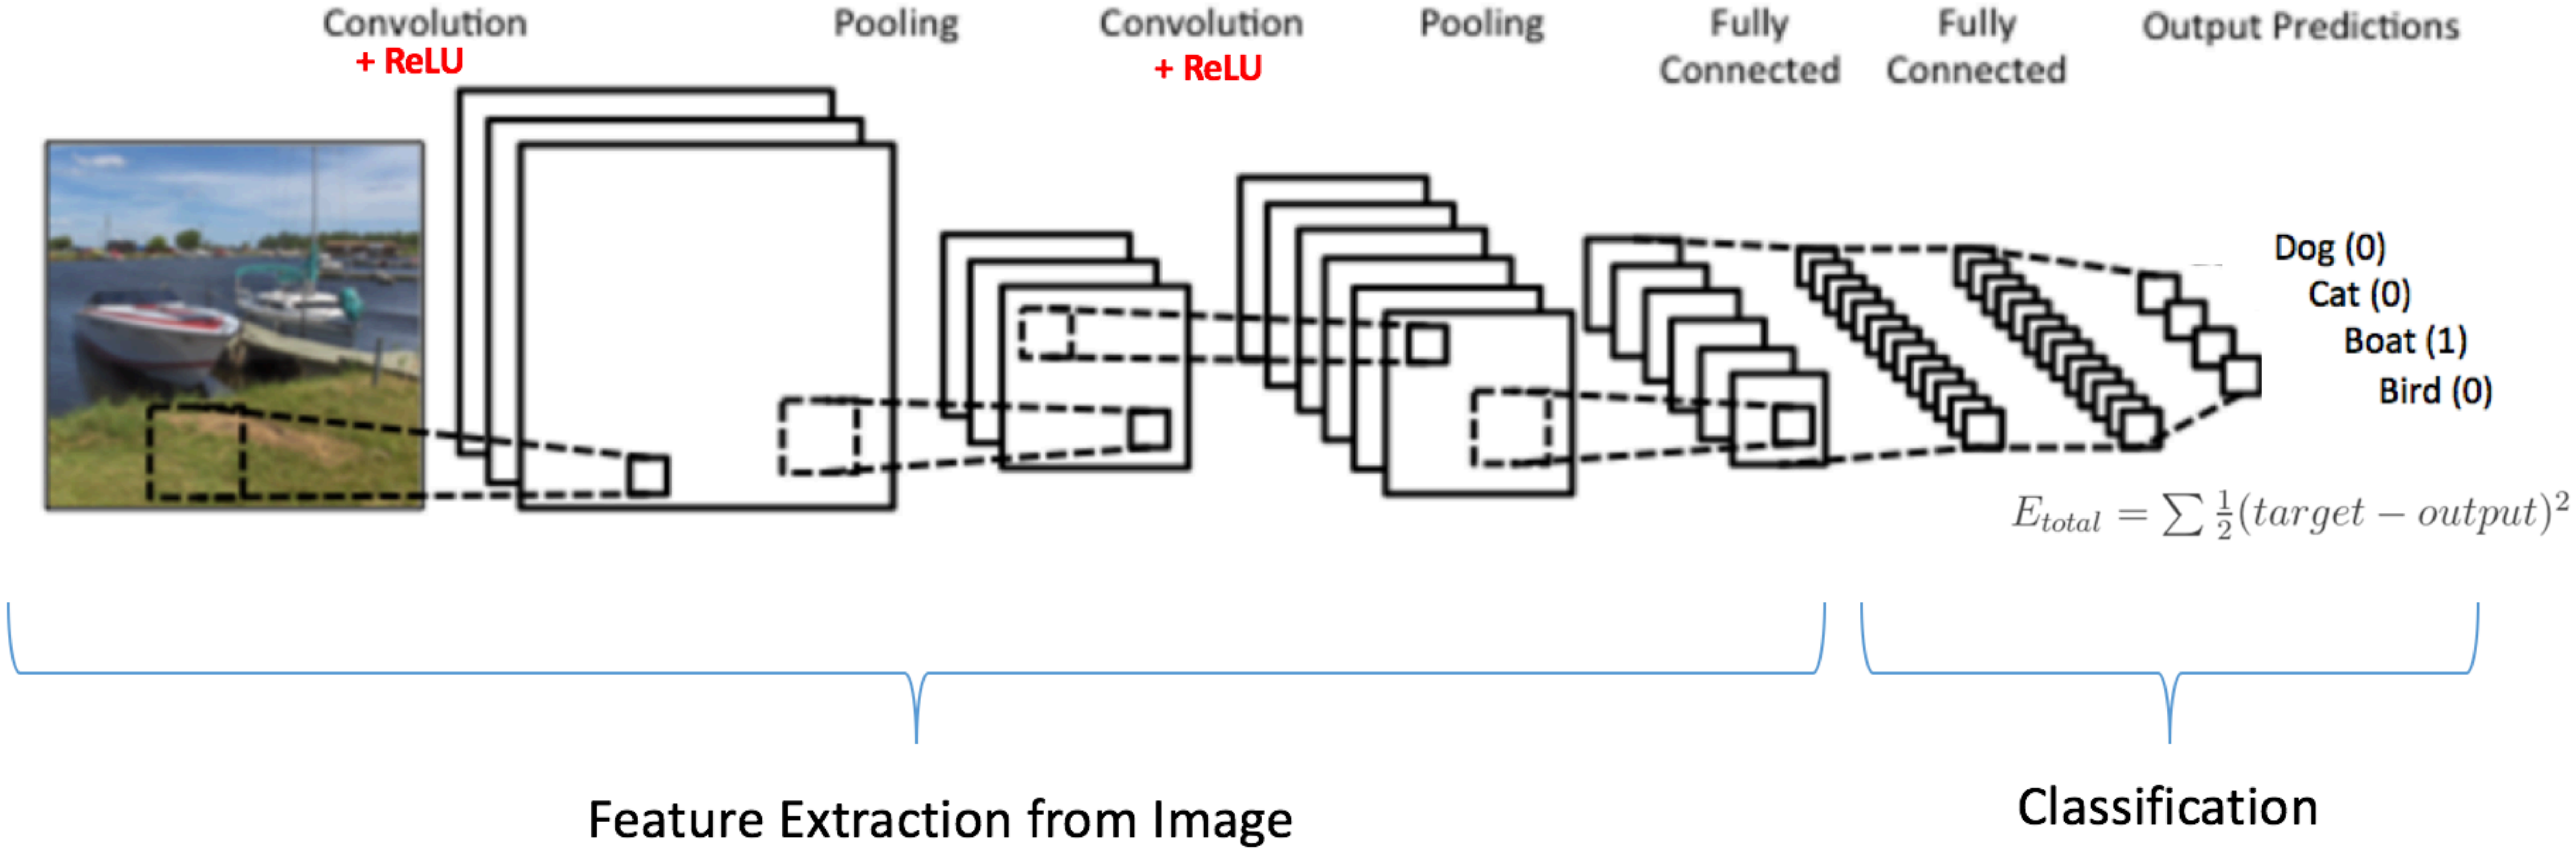

9	3	5	-8
-6	2	-3	1
1	3	4	1
3	-4	5	1

9	3	5	0
0	2	0	1
1	3	4	1
3	0	5	1

Result of: Linear Image Filtering → After Non-linear ReLU

Framework for Today's Topic

Problem: How do we go from the optics of image formation to digital images as arrays of numbers?

Key Idea(s): Sampling and the notion of band limited functions

Theory: Sampling Theory

Reminder

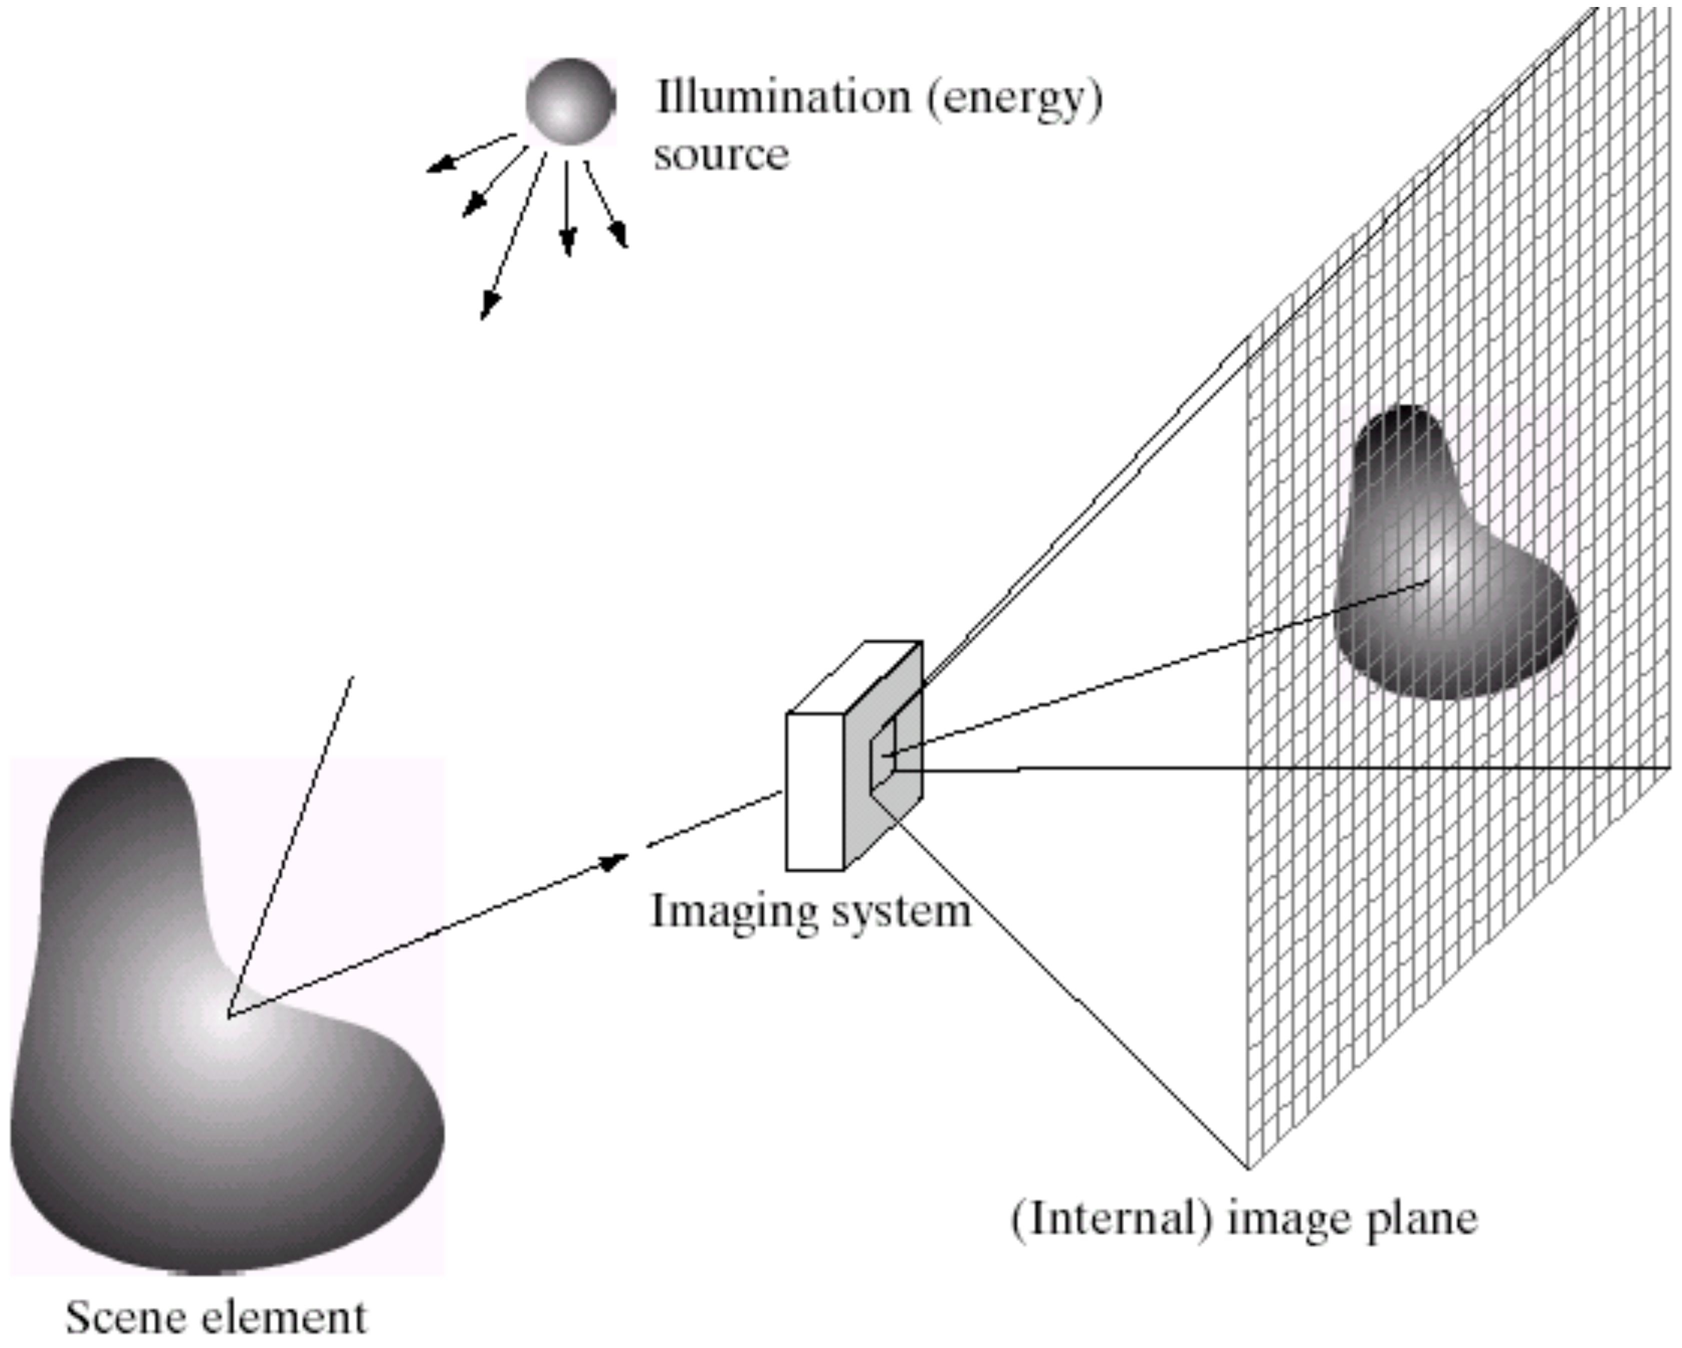

Images are a **discrete**, or **sampled**, representation of a continuous world

What is an **Image**?

Up to now provided a **physical characterization**

— image formation as a problem in physics/optics

— we also talked about simple image processing algorithms on image arrays

Now provide a **mathematical characterization**

— to understand how to represent images digitally

— to understand how to compute with images

Continuous Case

“**Image**” suggests a 2D surface whose appearance varies from point–to–point — the surface typically is a plane (but might be curved, e.g., as is with an eye)

Appearance can be **Grayscale** (Black and White) or **Colour**

In **Grayscale**, variation in appearance can be described by a single parameter corresponding to the amount of light reaching the image at a given point in a given time

Continuous Case

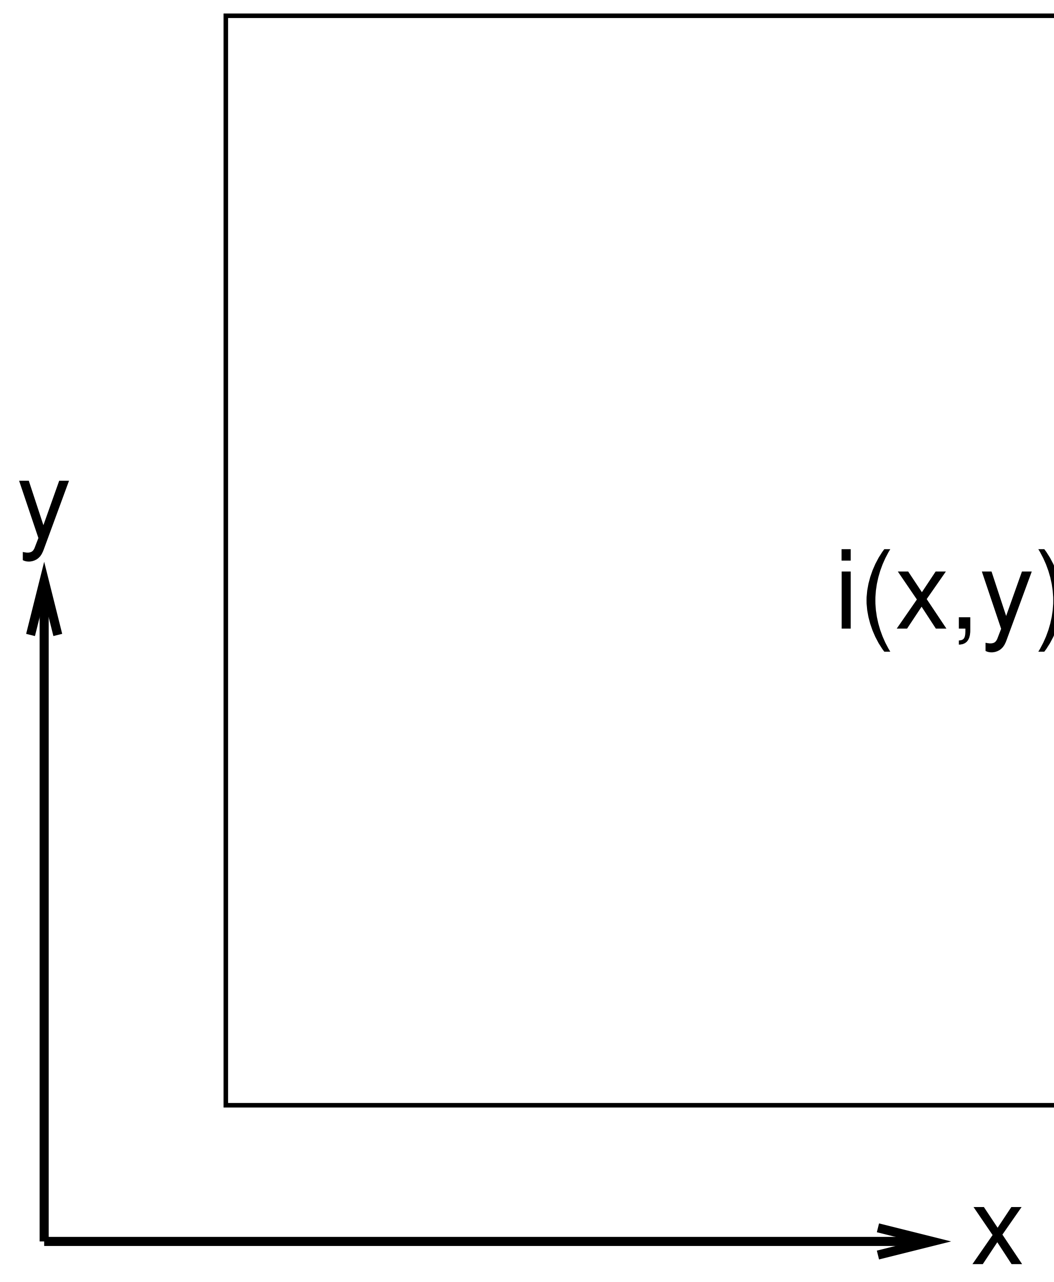

Denote the image as a function, $i(x, y)$, where x and y are spatial variables

Aside: The convention for this section is to use lower case letters for the continuous case and upper case letters for the discrete case

Continuous Case: Observations

– $i(x, y)$ is a **real-valued function** of **real spatial variables**, x and y

Recall: Pinhole Camera

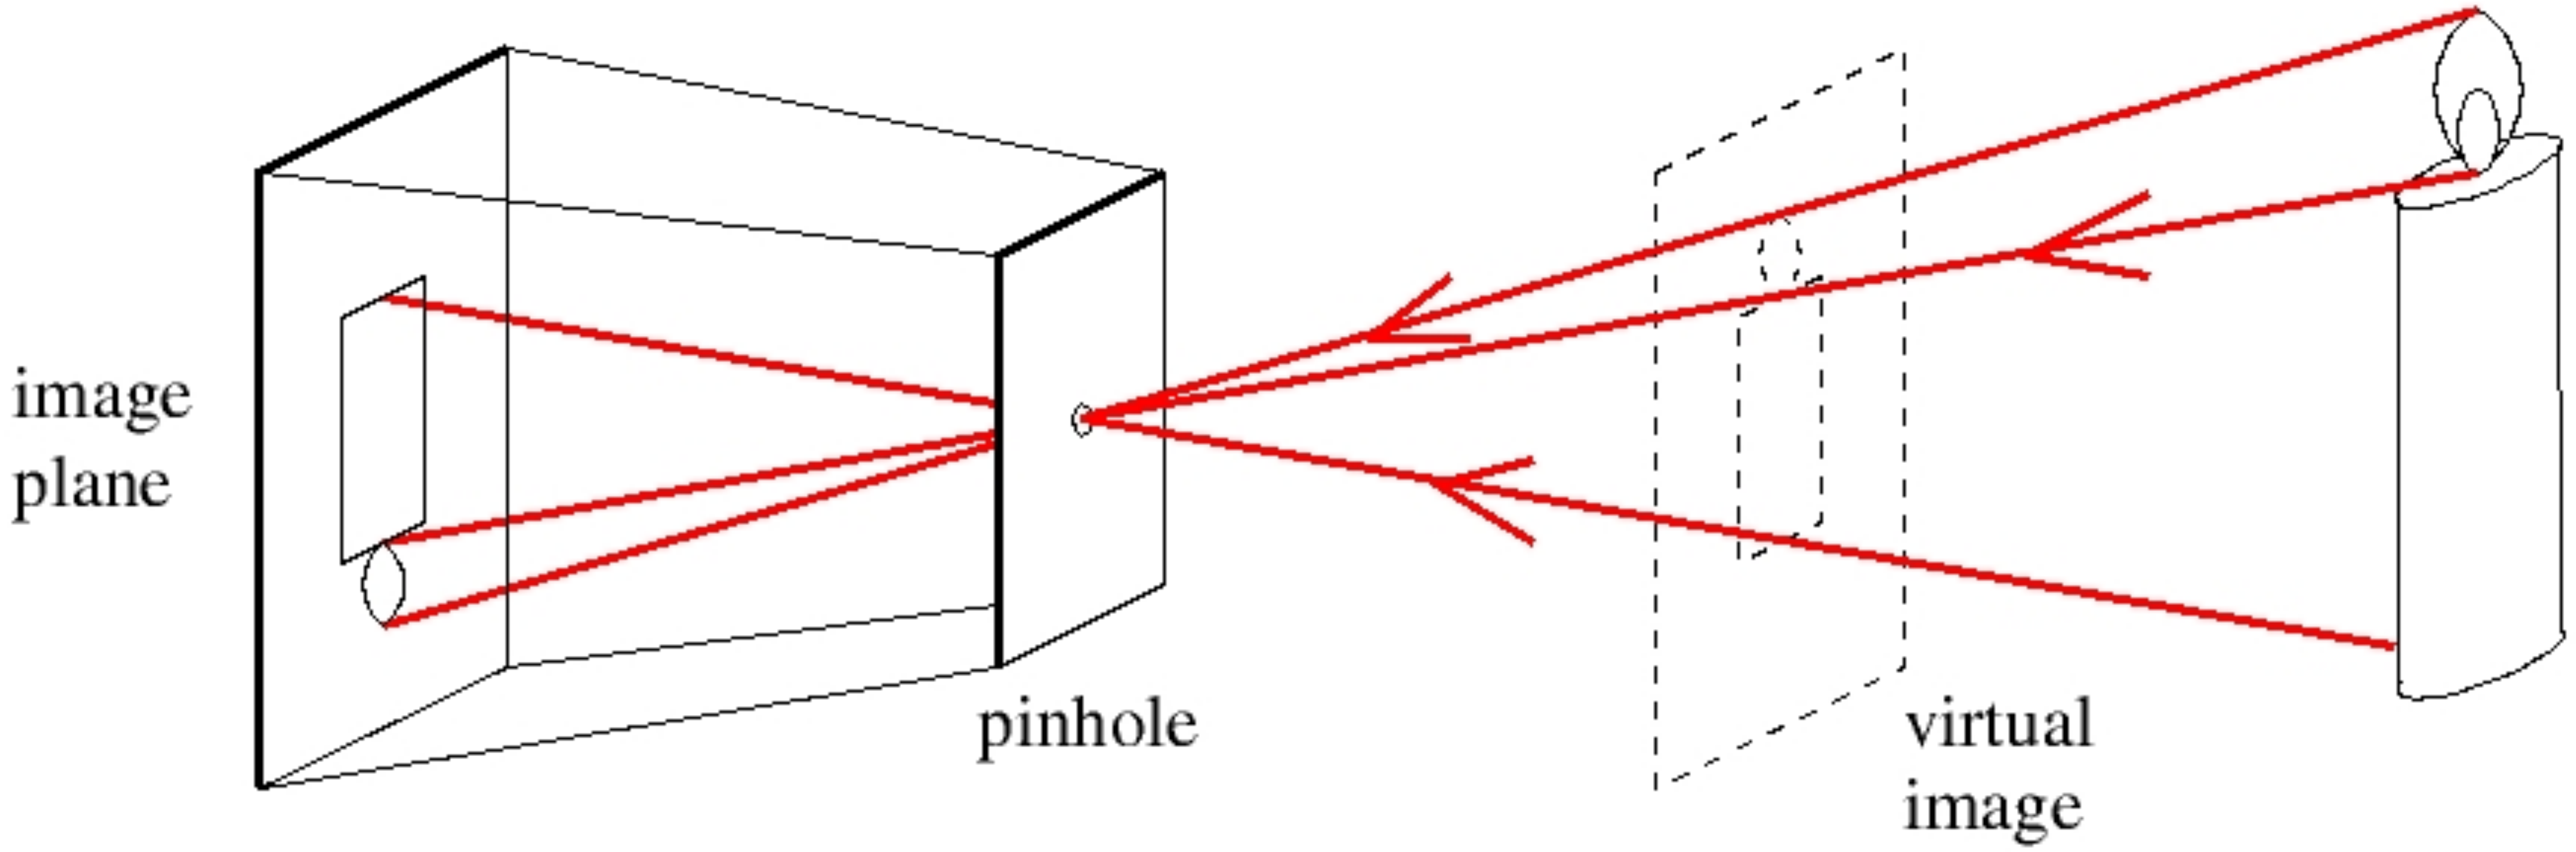

Forsyth & Ponce (2nd ed.) Figure 1.2

Continuous Case: Observations

- $i(x, y)$ is a **real-valued function** of **real spatial variables**, x and y

- $i(x, y)$ is **bounded above and below**. That is

$$0 \leq i(x, y) \leq M$$

for some maximum brightness M

Continuous Case: Observations

– $i(x, y)$ is a **real-valued function** of **real spatial variables**, x and y

– $i(x, y)$ is **bounded above and below**. That is

$$0 \leq i(x, y) \leq M$$

for some maximum brightness M

– $i(x, y)$ is **bounded in extent**. That is, $i(x, y)$ is non-zero (i.e., strictly positive) over, at most, a bounded region

Continuous Case

— Images also can be considered a function of time. Then, we write $i(x, y, t)$ where x and y are spatial variable and t is a **temporal variable**

— To make the dependence of brightness on wavelength explicit, we can instead write $i(x, y, t, \lambda)$ where x, y and t are as above and where λ is a **spectral variable**

— More commonly, we think of "color" already as discrete and write

$$\begin{gathered} i_R(x, y) \\ i_G(x, y) \\ i_B(x, y) \end{gathered}$$

for specific colour channels, R, G and B

Discrete Case

Idea: Superimpose (regular) grid on continuous image

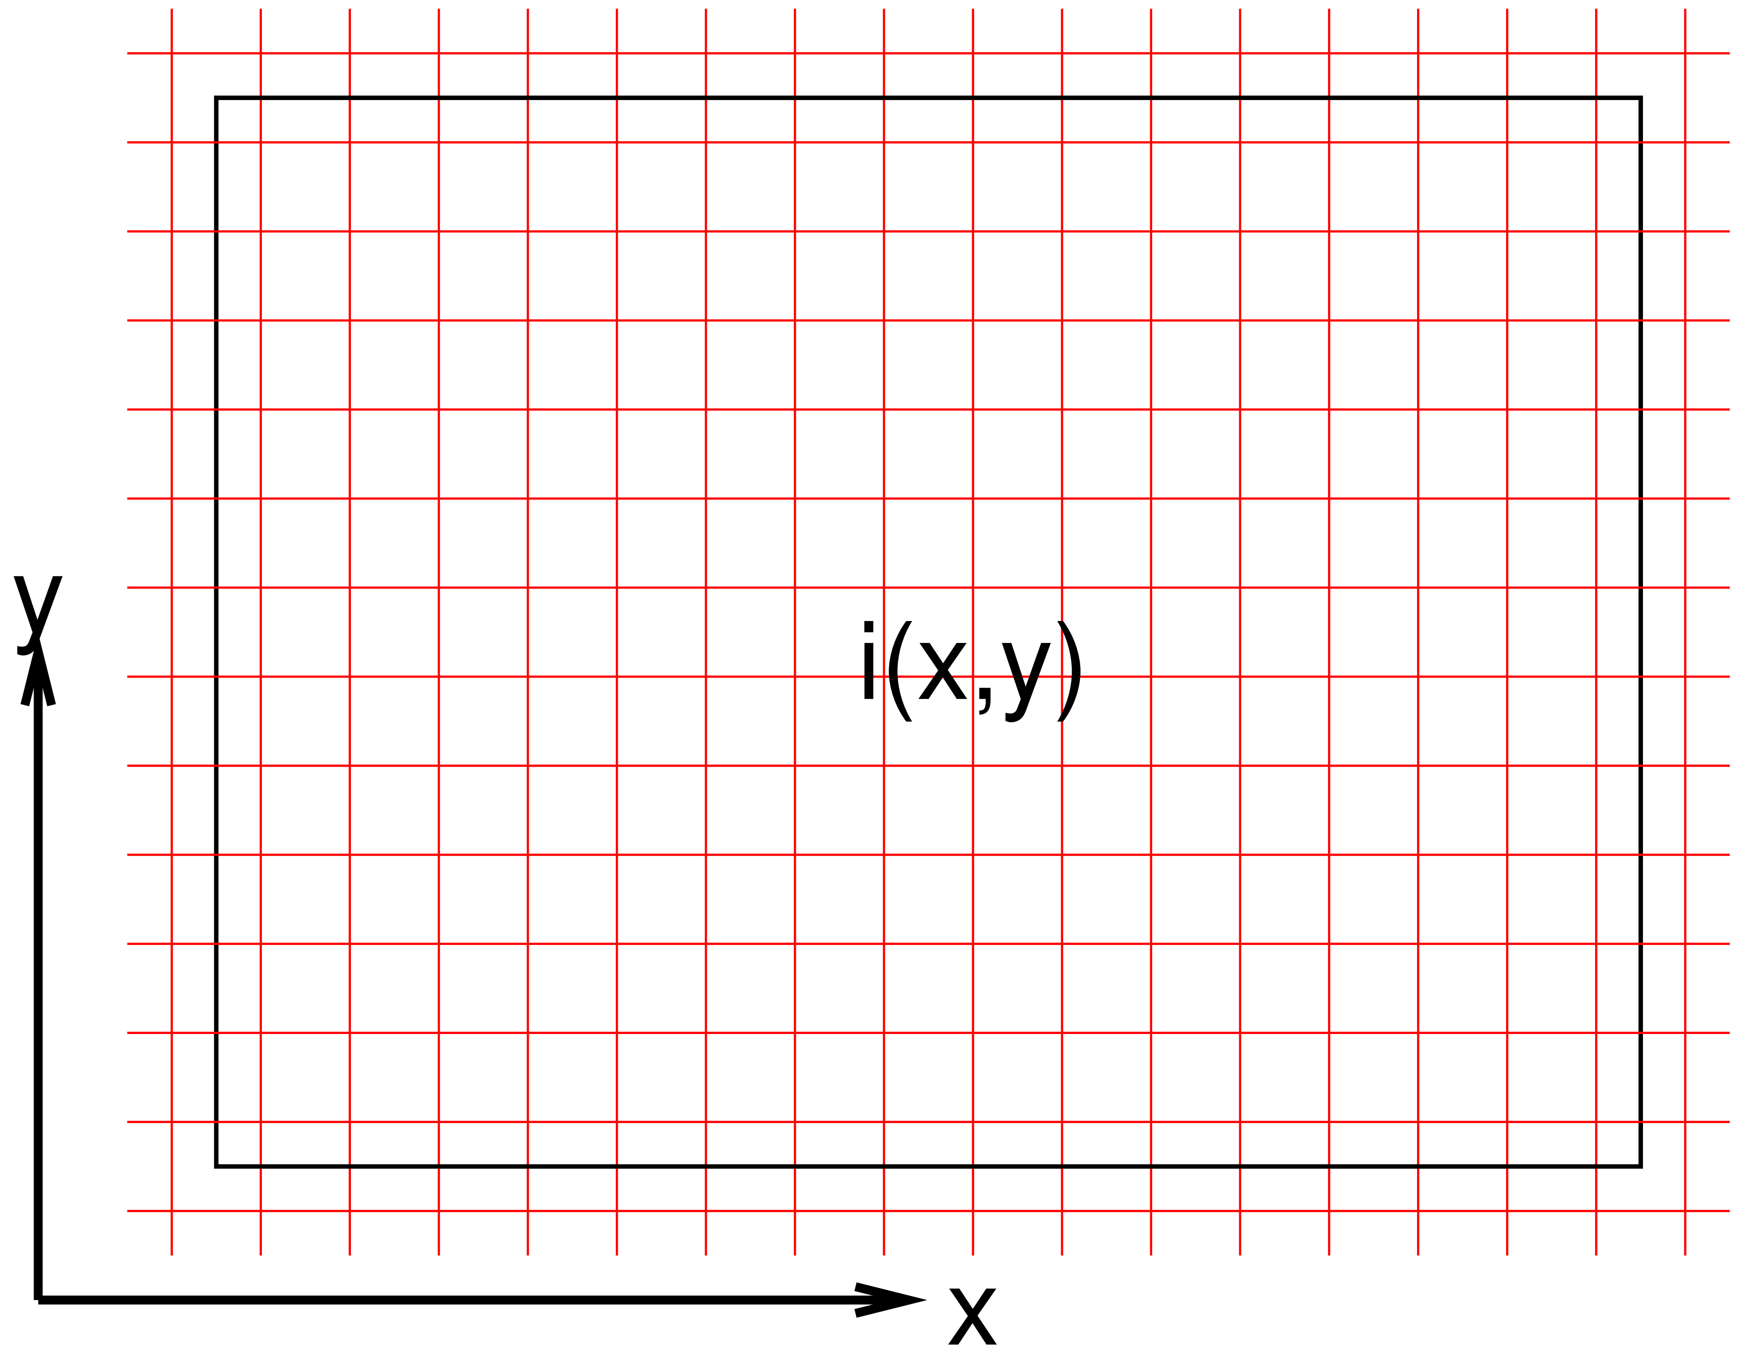

Sample the underlying continuous image according to the **tessellation** imposed by the grid

Discrete Case

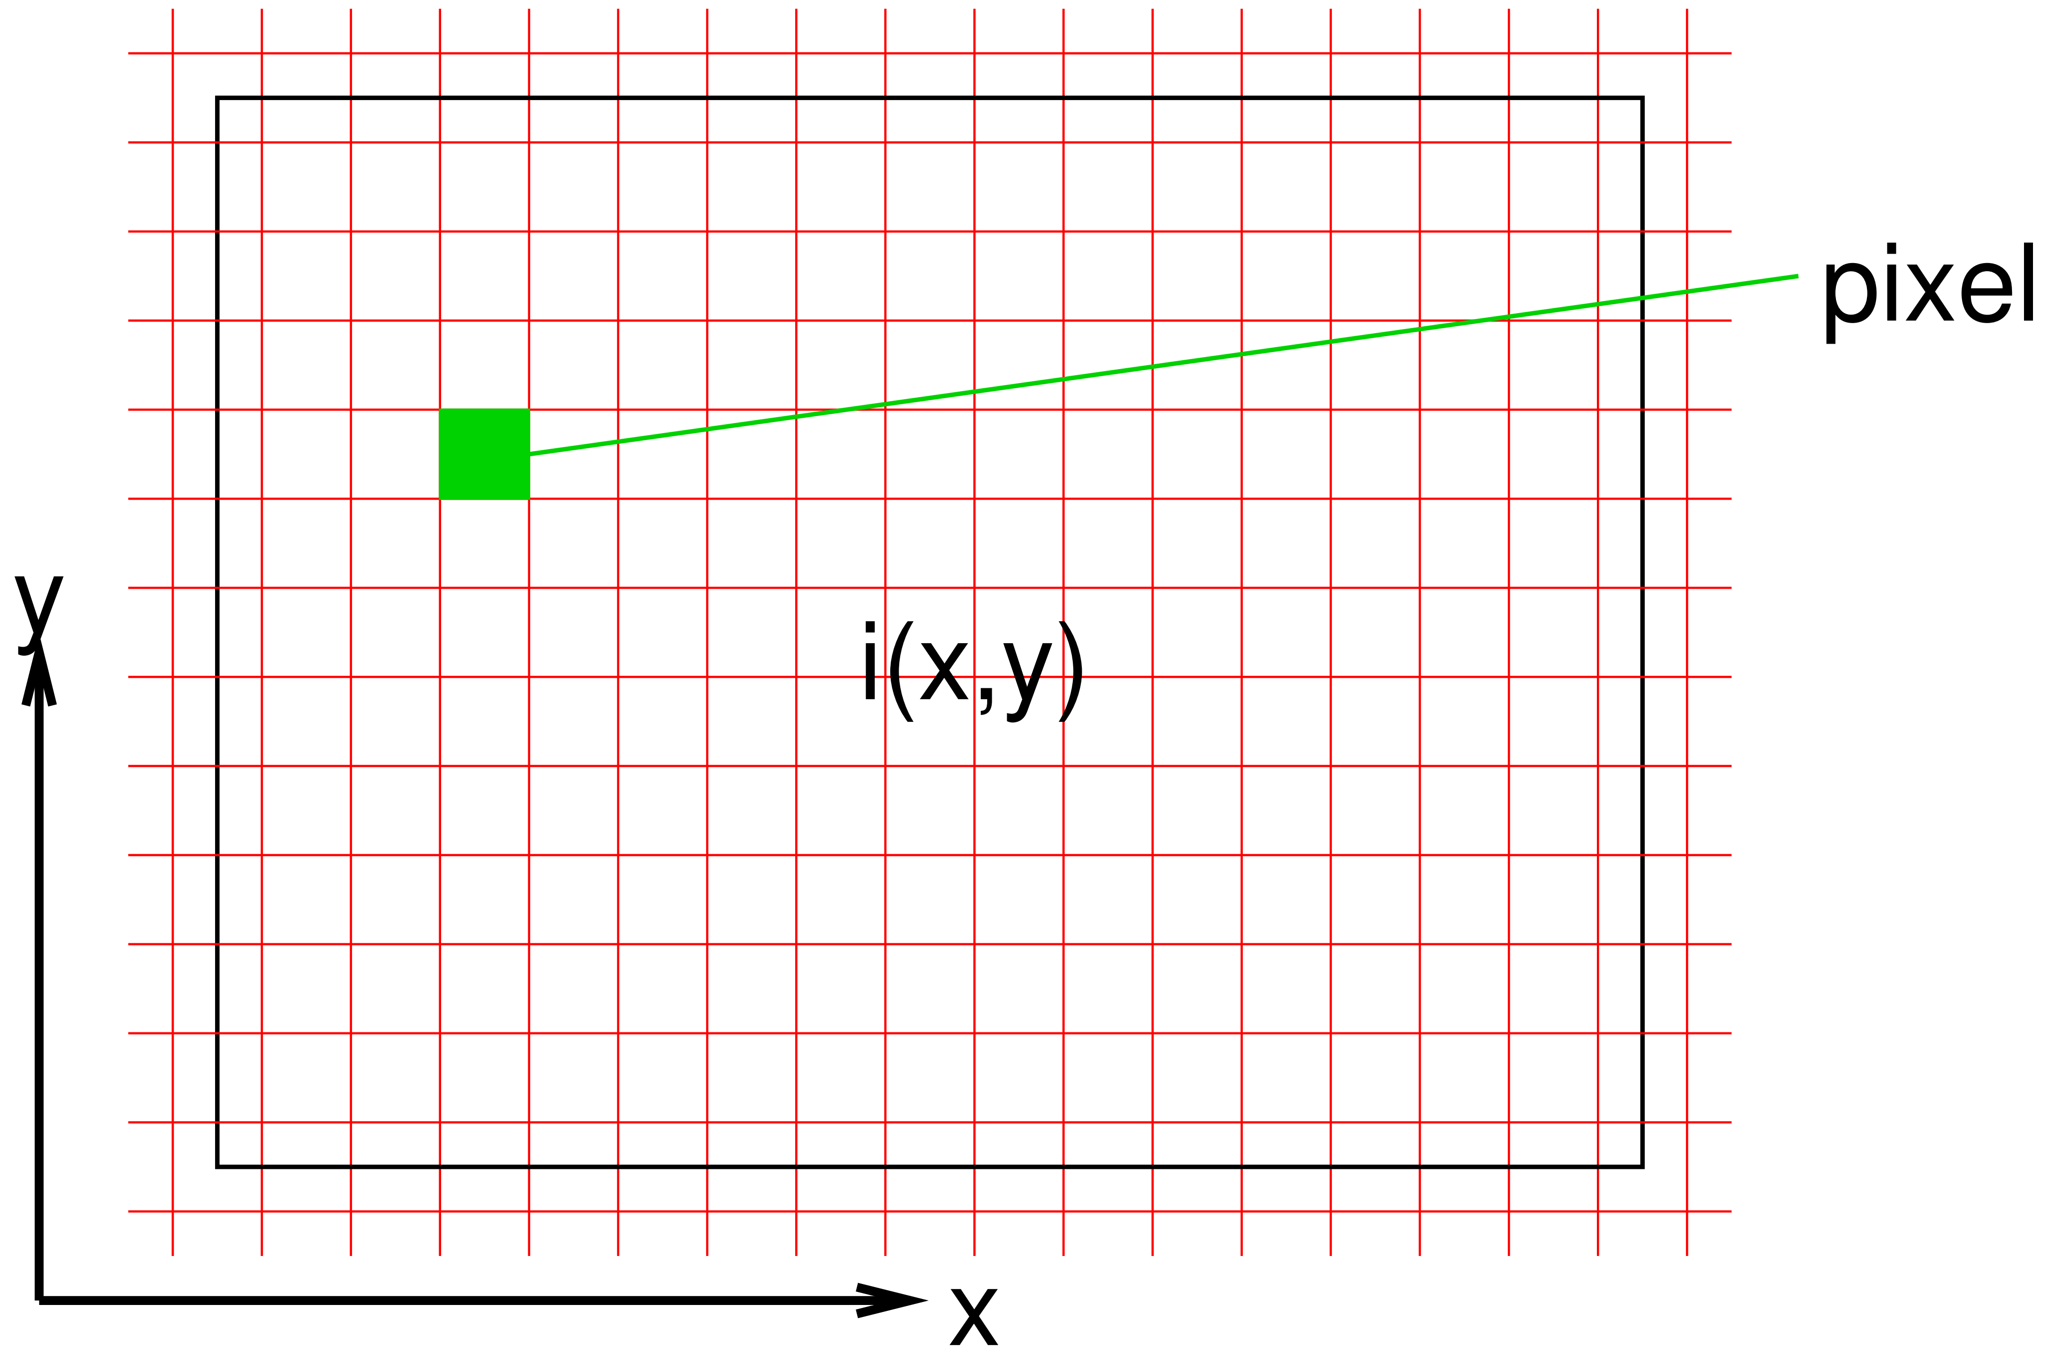

Discrete Case

Each grid cell is called a picture element (**pixel**)

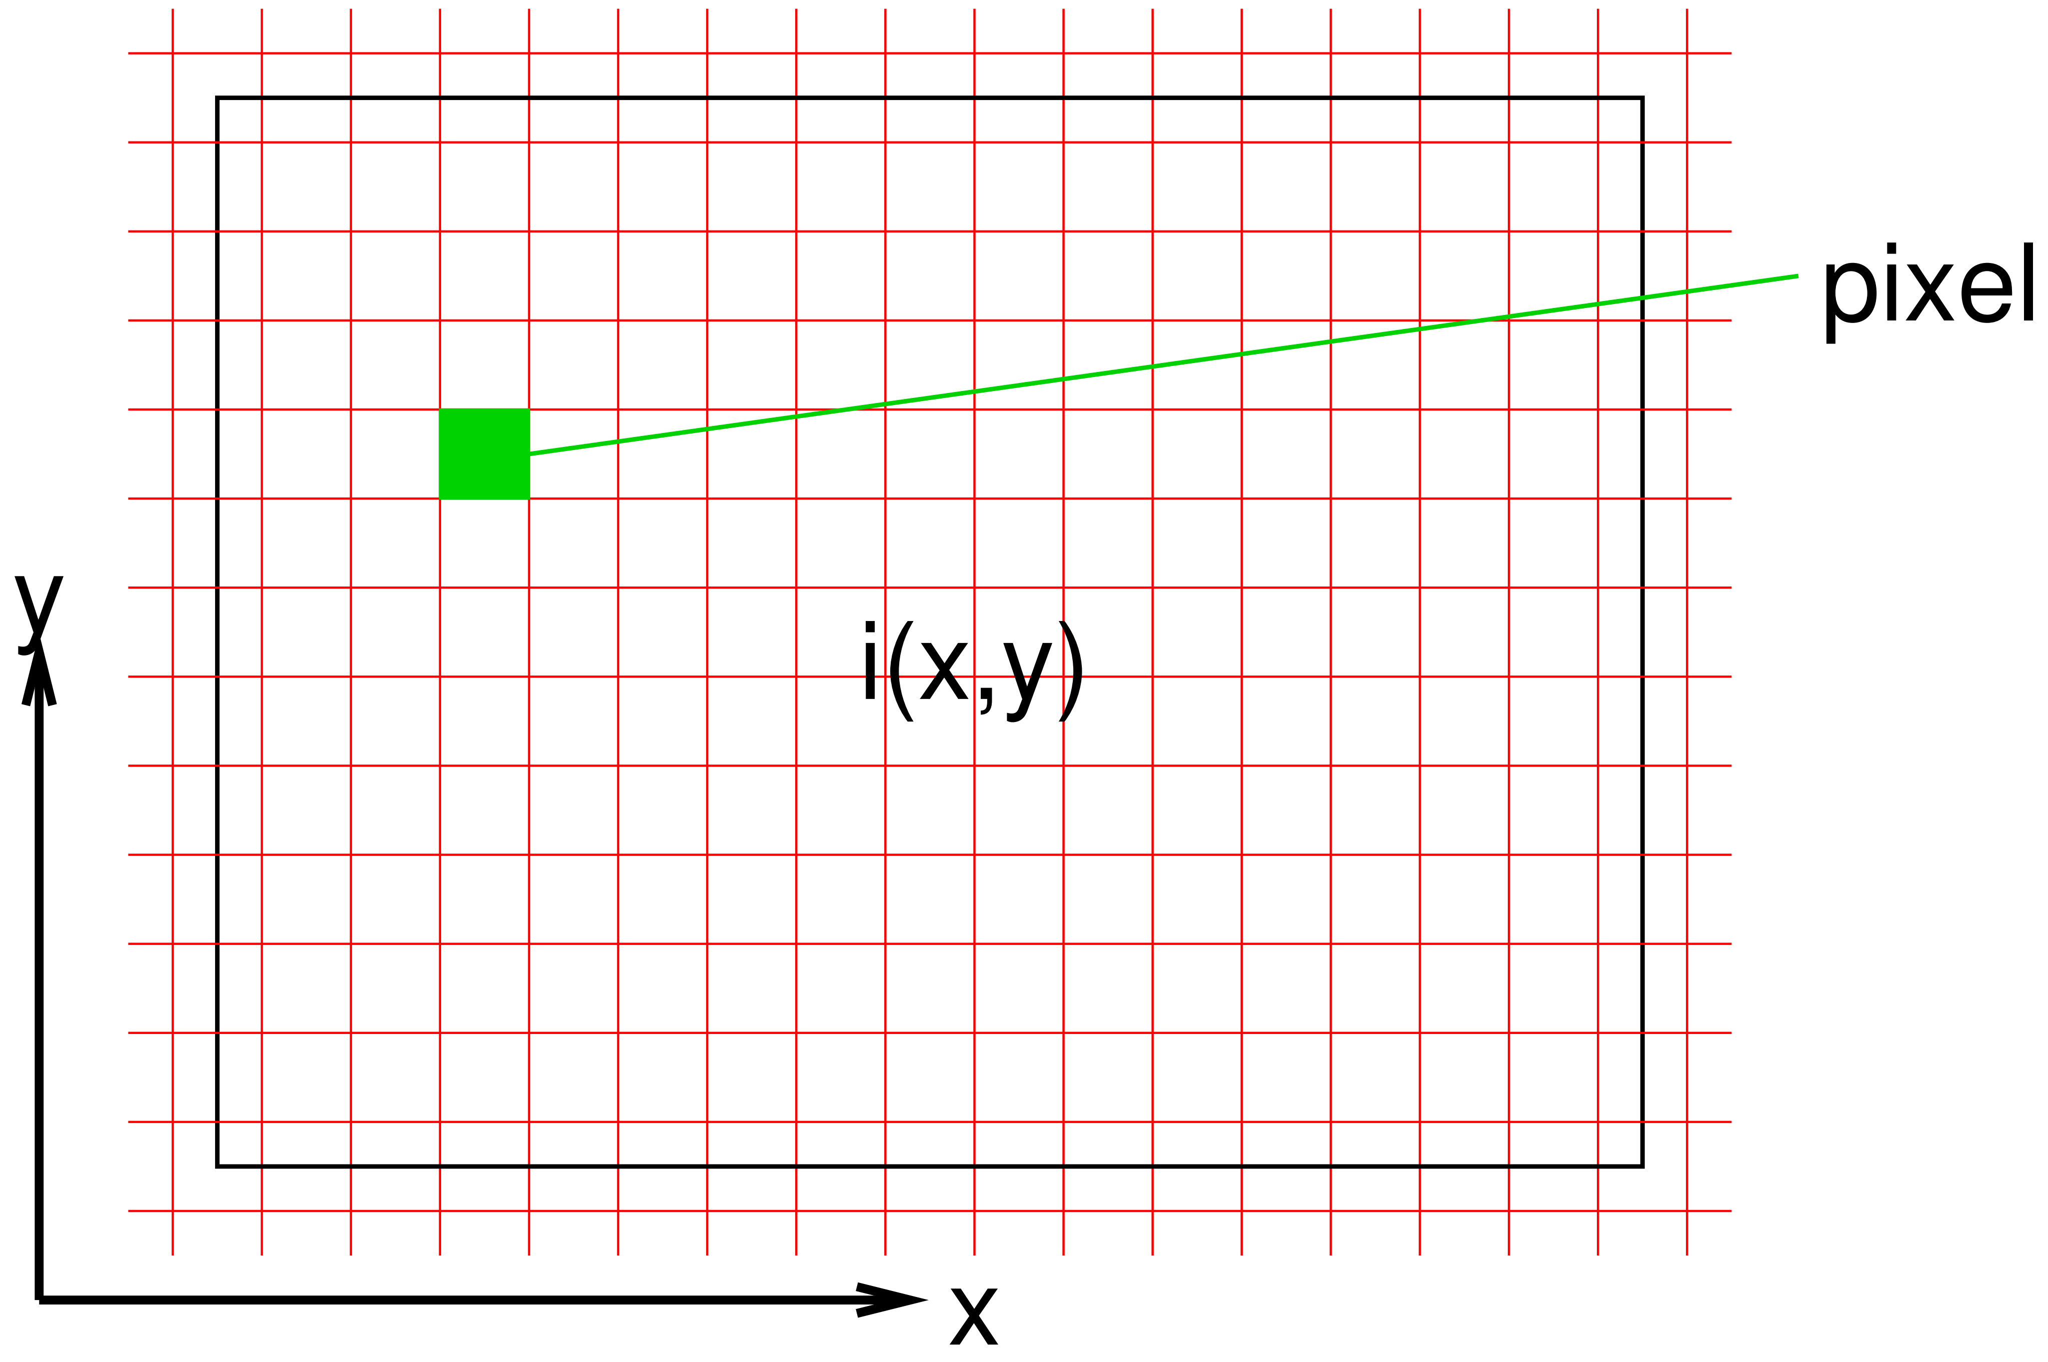

Denote the discrete image as $I(X,Y)$

We can store the pixels in a matrix or array

Discrete Case

Question: How to sample?

— Sample brightness at the point?

— “Average” brightness over entire pixel?

Answer:

— Point sampling is useful for theoretical development

— Area-based sampling occurs in practice

Discrete Case

Question: What about the brightness samples themselves?

Discrete Case

Question: What about the brightness samples themselves?

Answer: We make values of $I(X, Y)$ discrete as well

Recall: $0 \leq i(x, y) \leq M$

We divide the range $[0, M]$ into a finite number of equivalence classes. This is called **quantization**.

The values are called **grey-levels**.

Discrete Case

Quantization is a topic in its own right

For now, a simple linear scheme is sufficient

Suppose n bits-per-pixel are available. One can divide the range $[0, M]$ into evenly spaced intervals as follows:

$$i(x,y) \rightarrow \left\lfloor \frac{i(x,y)}{M}(2^n - 1) + 0.5 \right\rfloor$$

where $\lfloor\ \rfloor$ is floor (i.e., greatest integer less than or equal to)

Typically $n = 8$ resulting in grey-levels in the range $[0, 255]$

Sampling

It is clear that *some* information may be lost when we work on a discrete pixel grid.

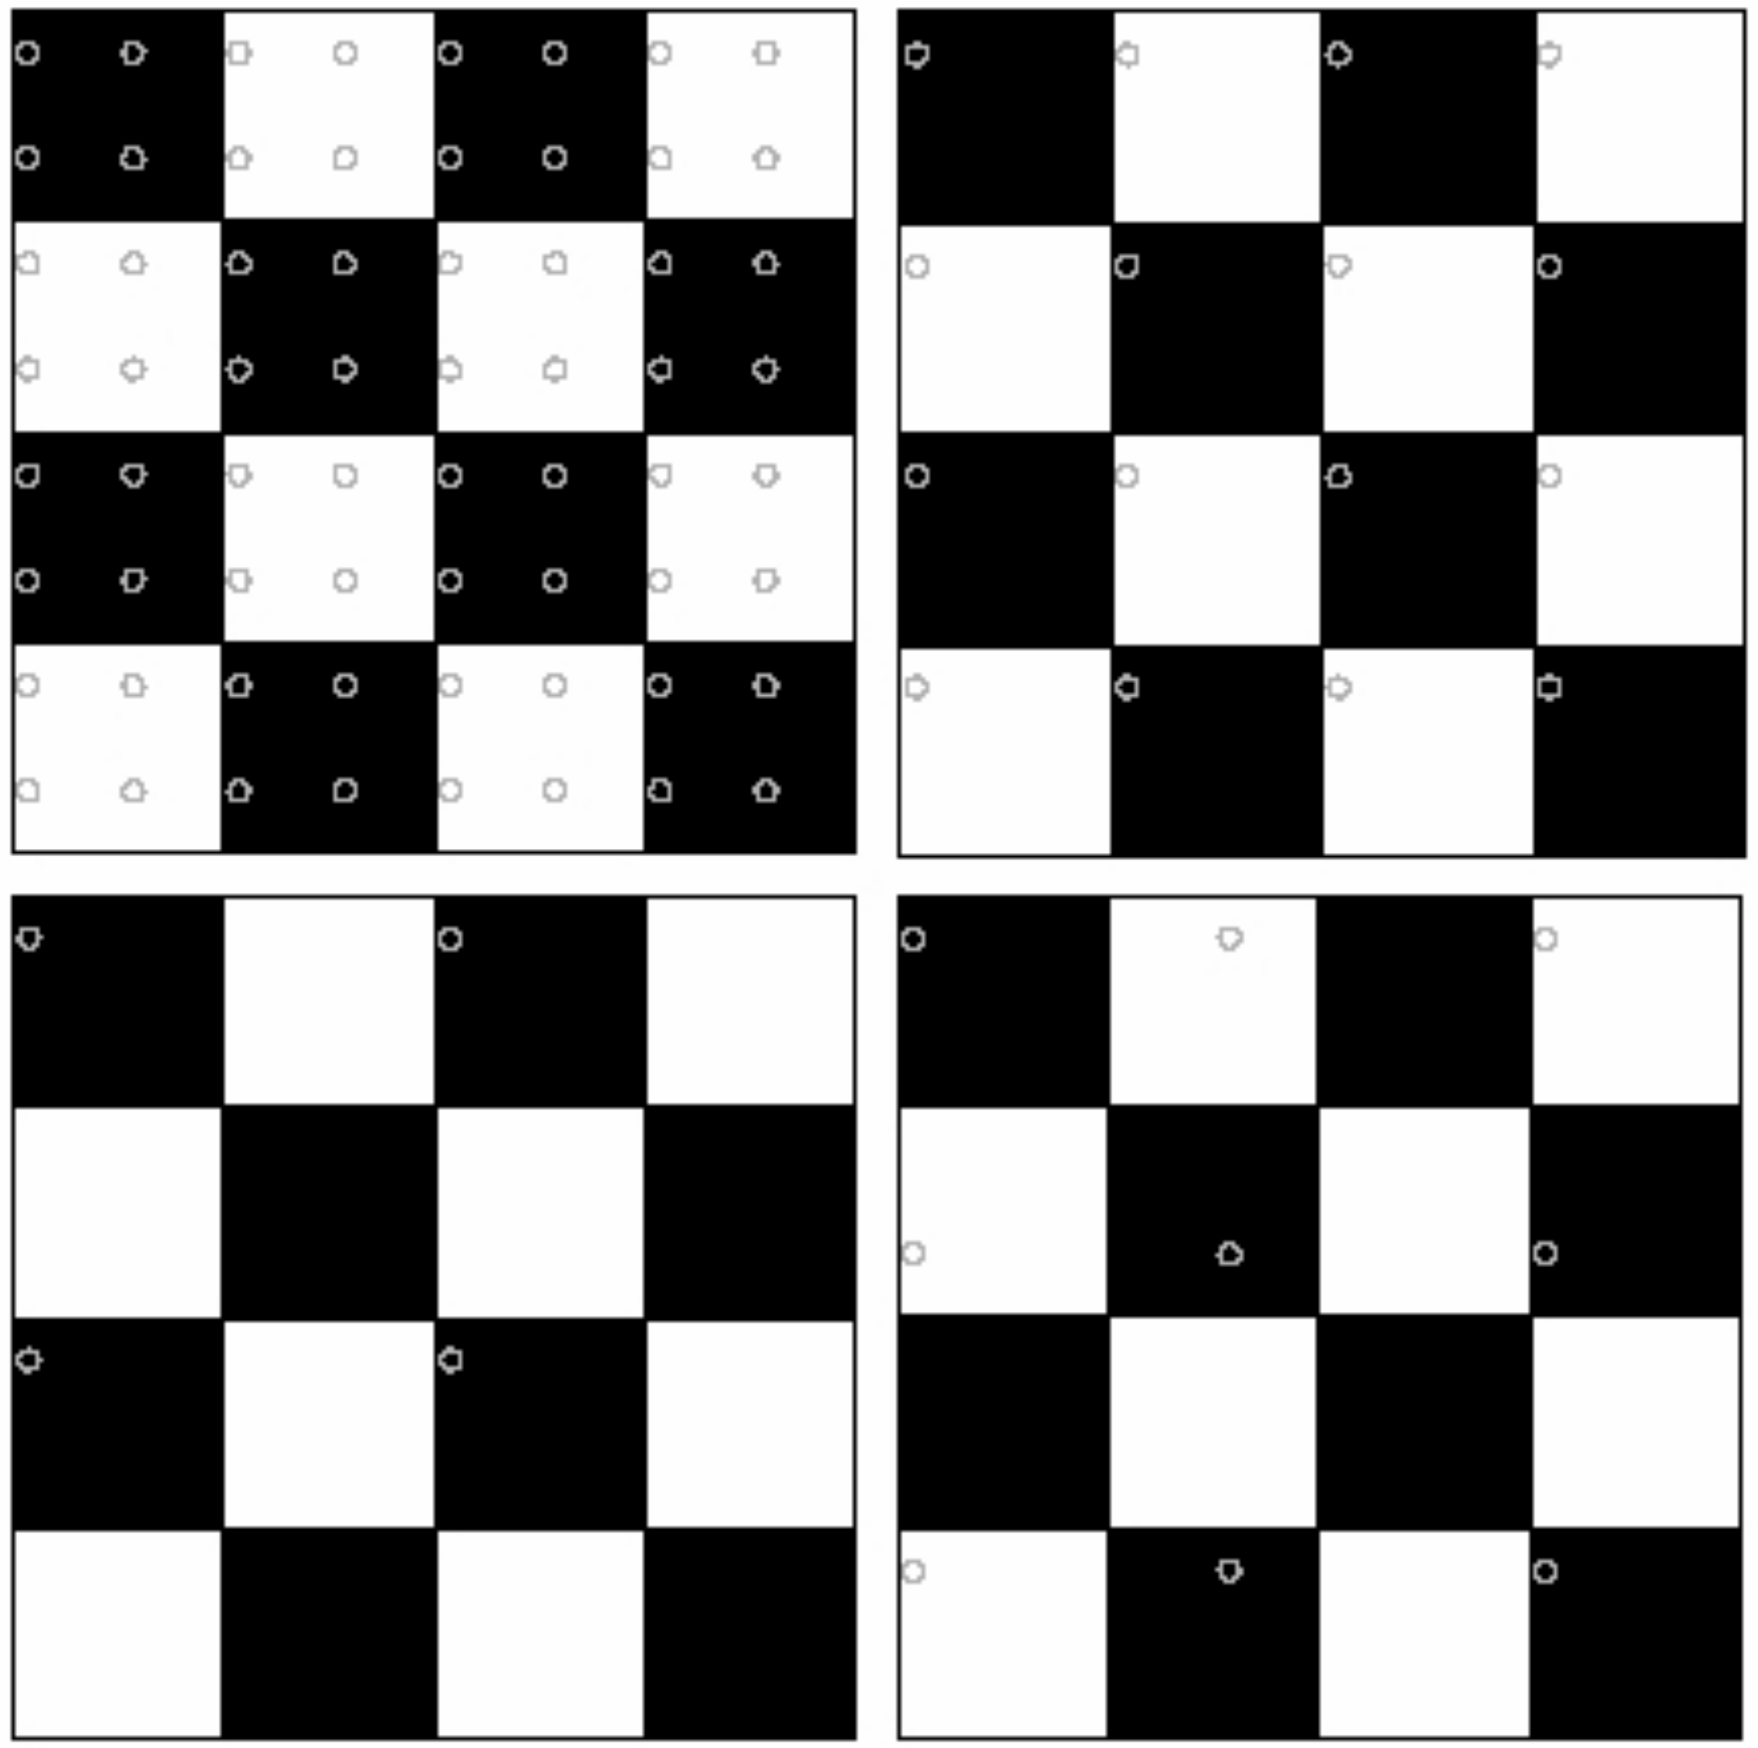

Forsyth & Ponce (2nd ed.) Figure 4.7

Sampling

It is clear that *some* information may be lost when we work on a discrete pixel grid.

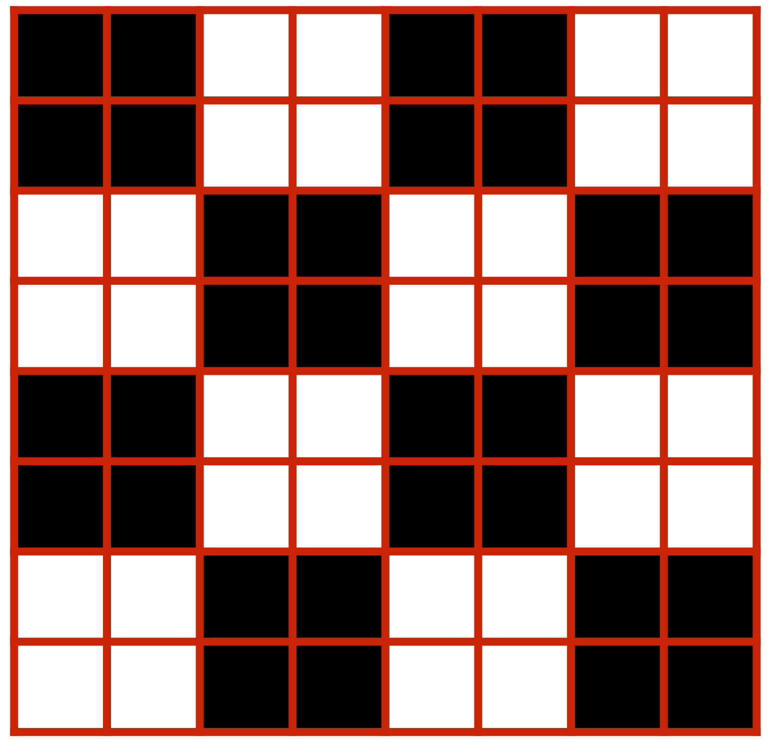

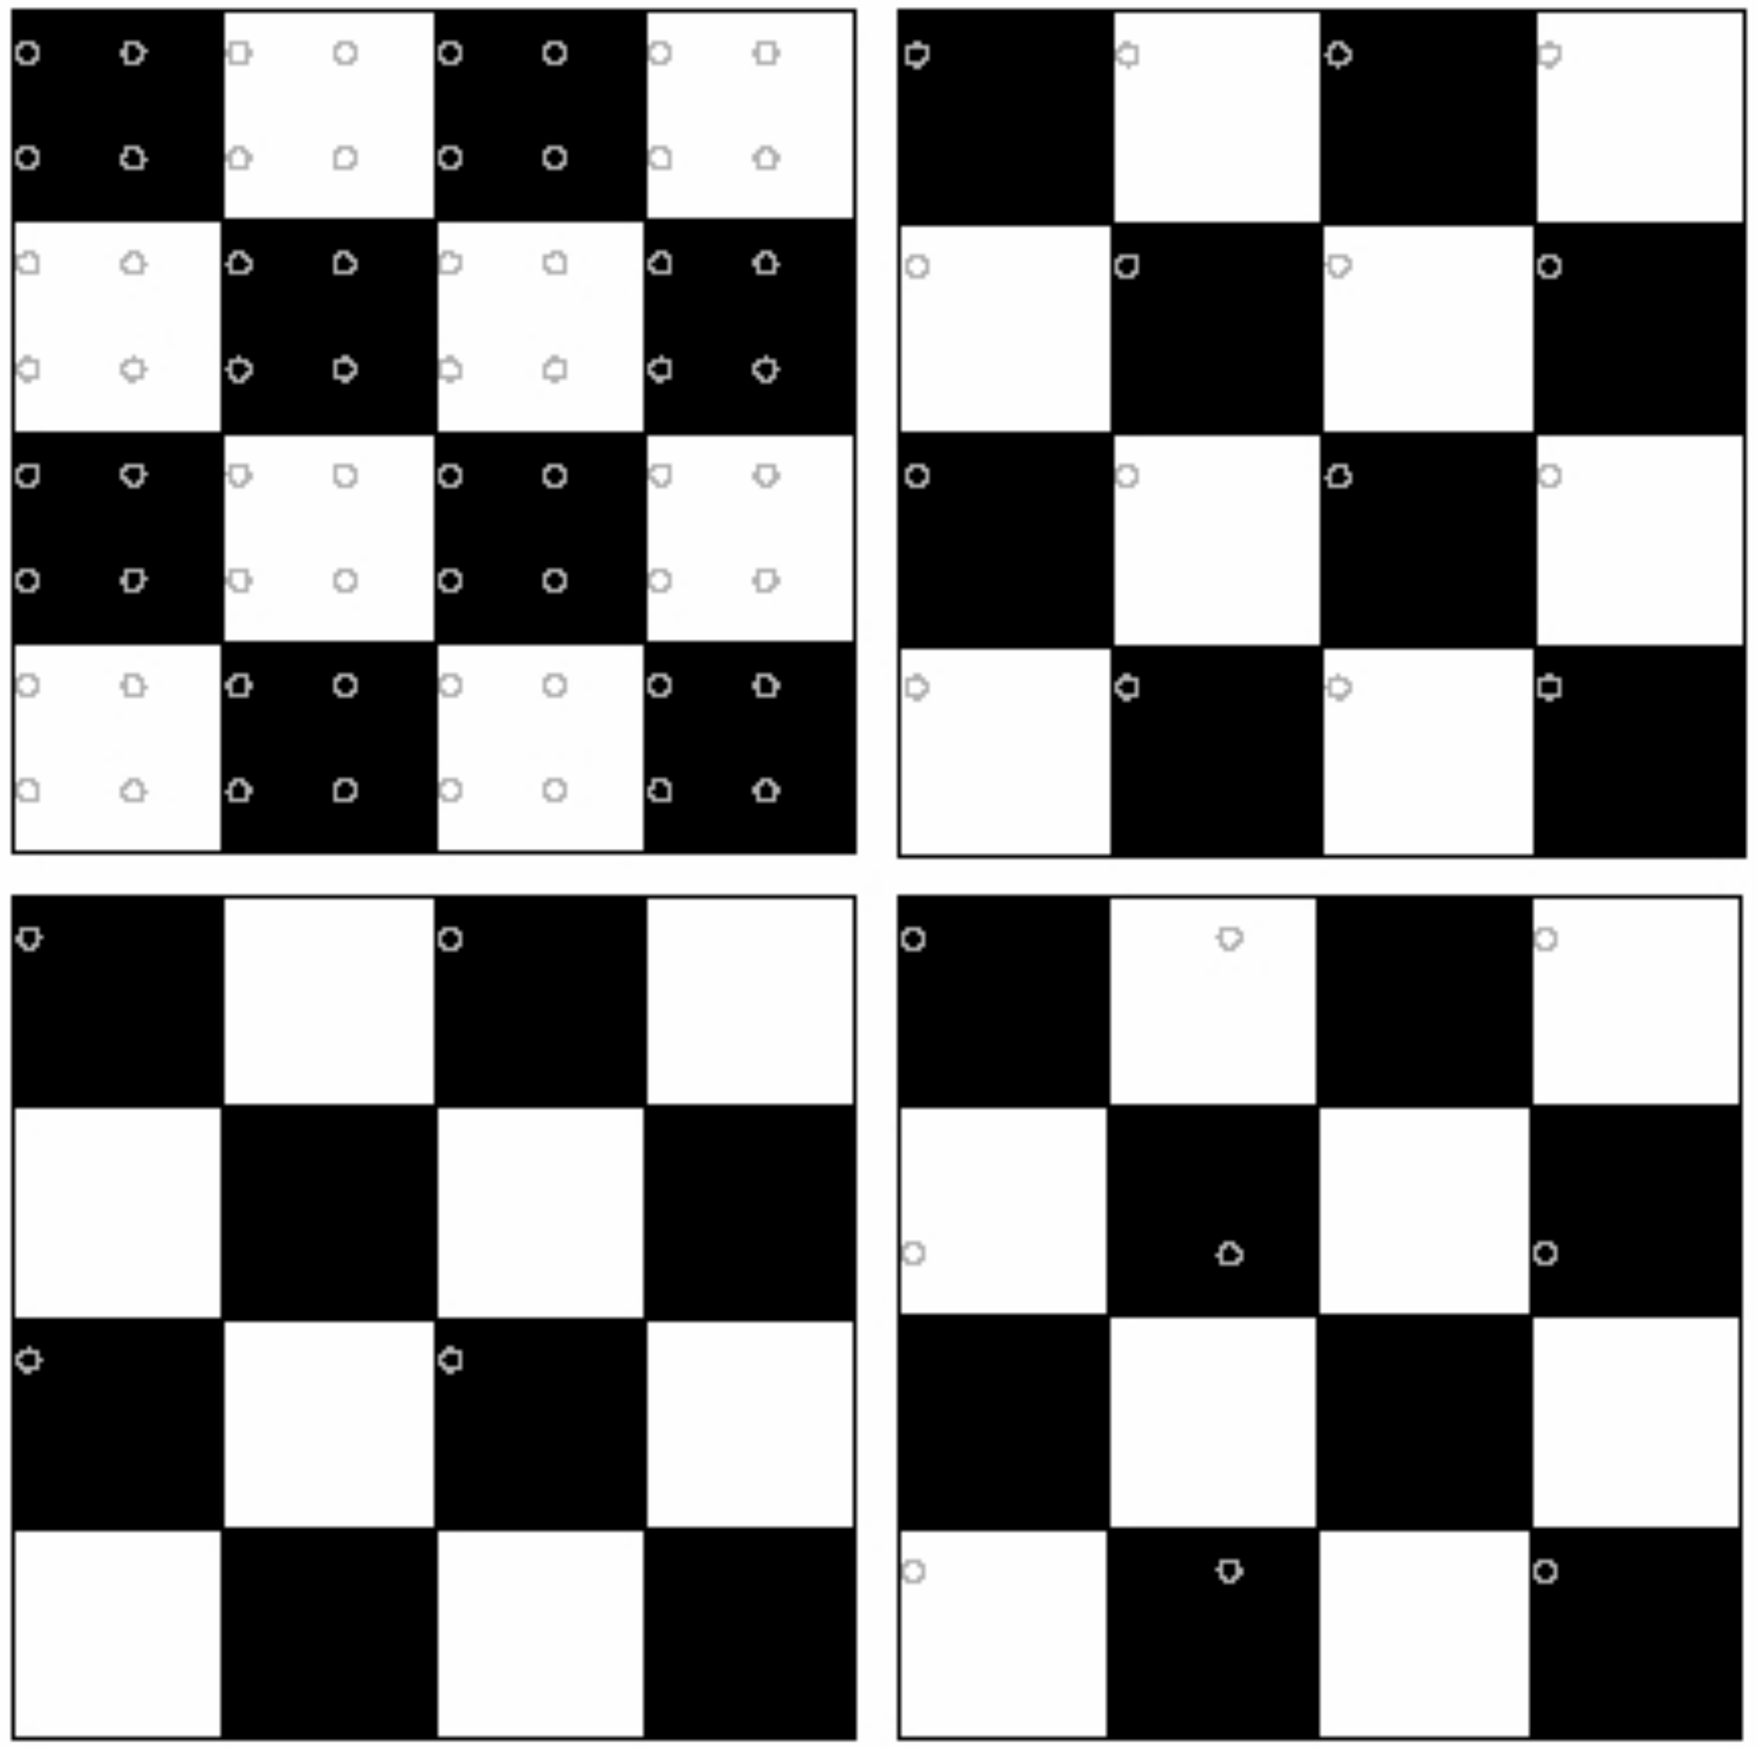

Forsyth & Ponce (2nd ed.) Figure 4.7

Sampling

It is clear that *some* information may be lost when we work on a discrete pixel grid.

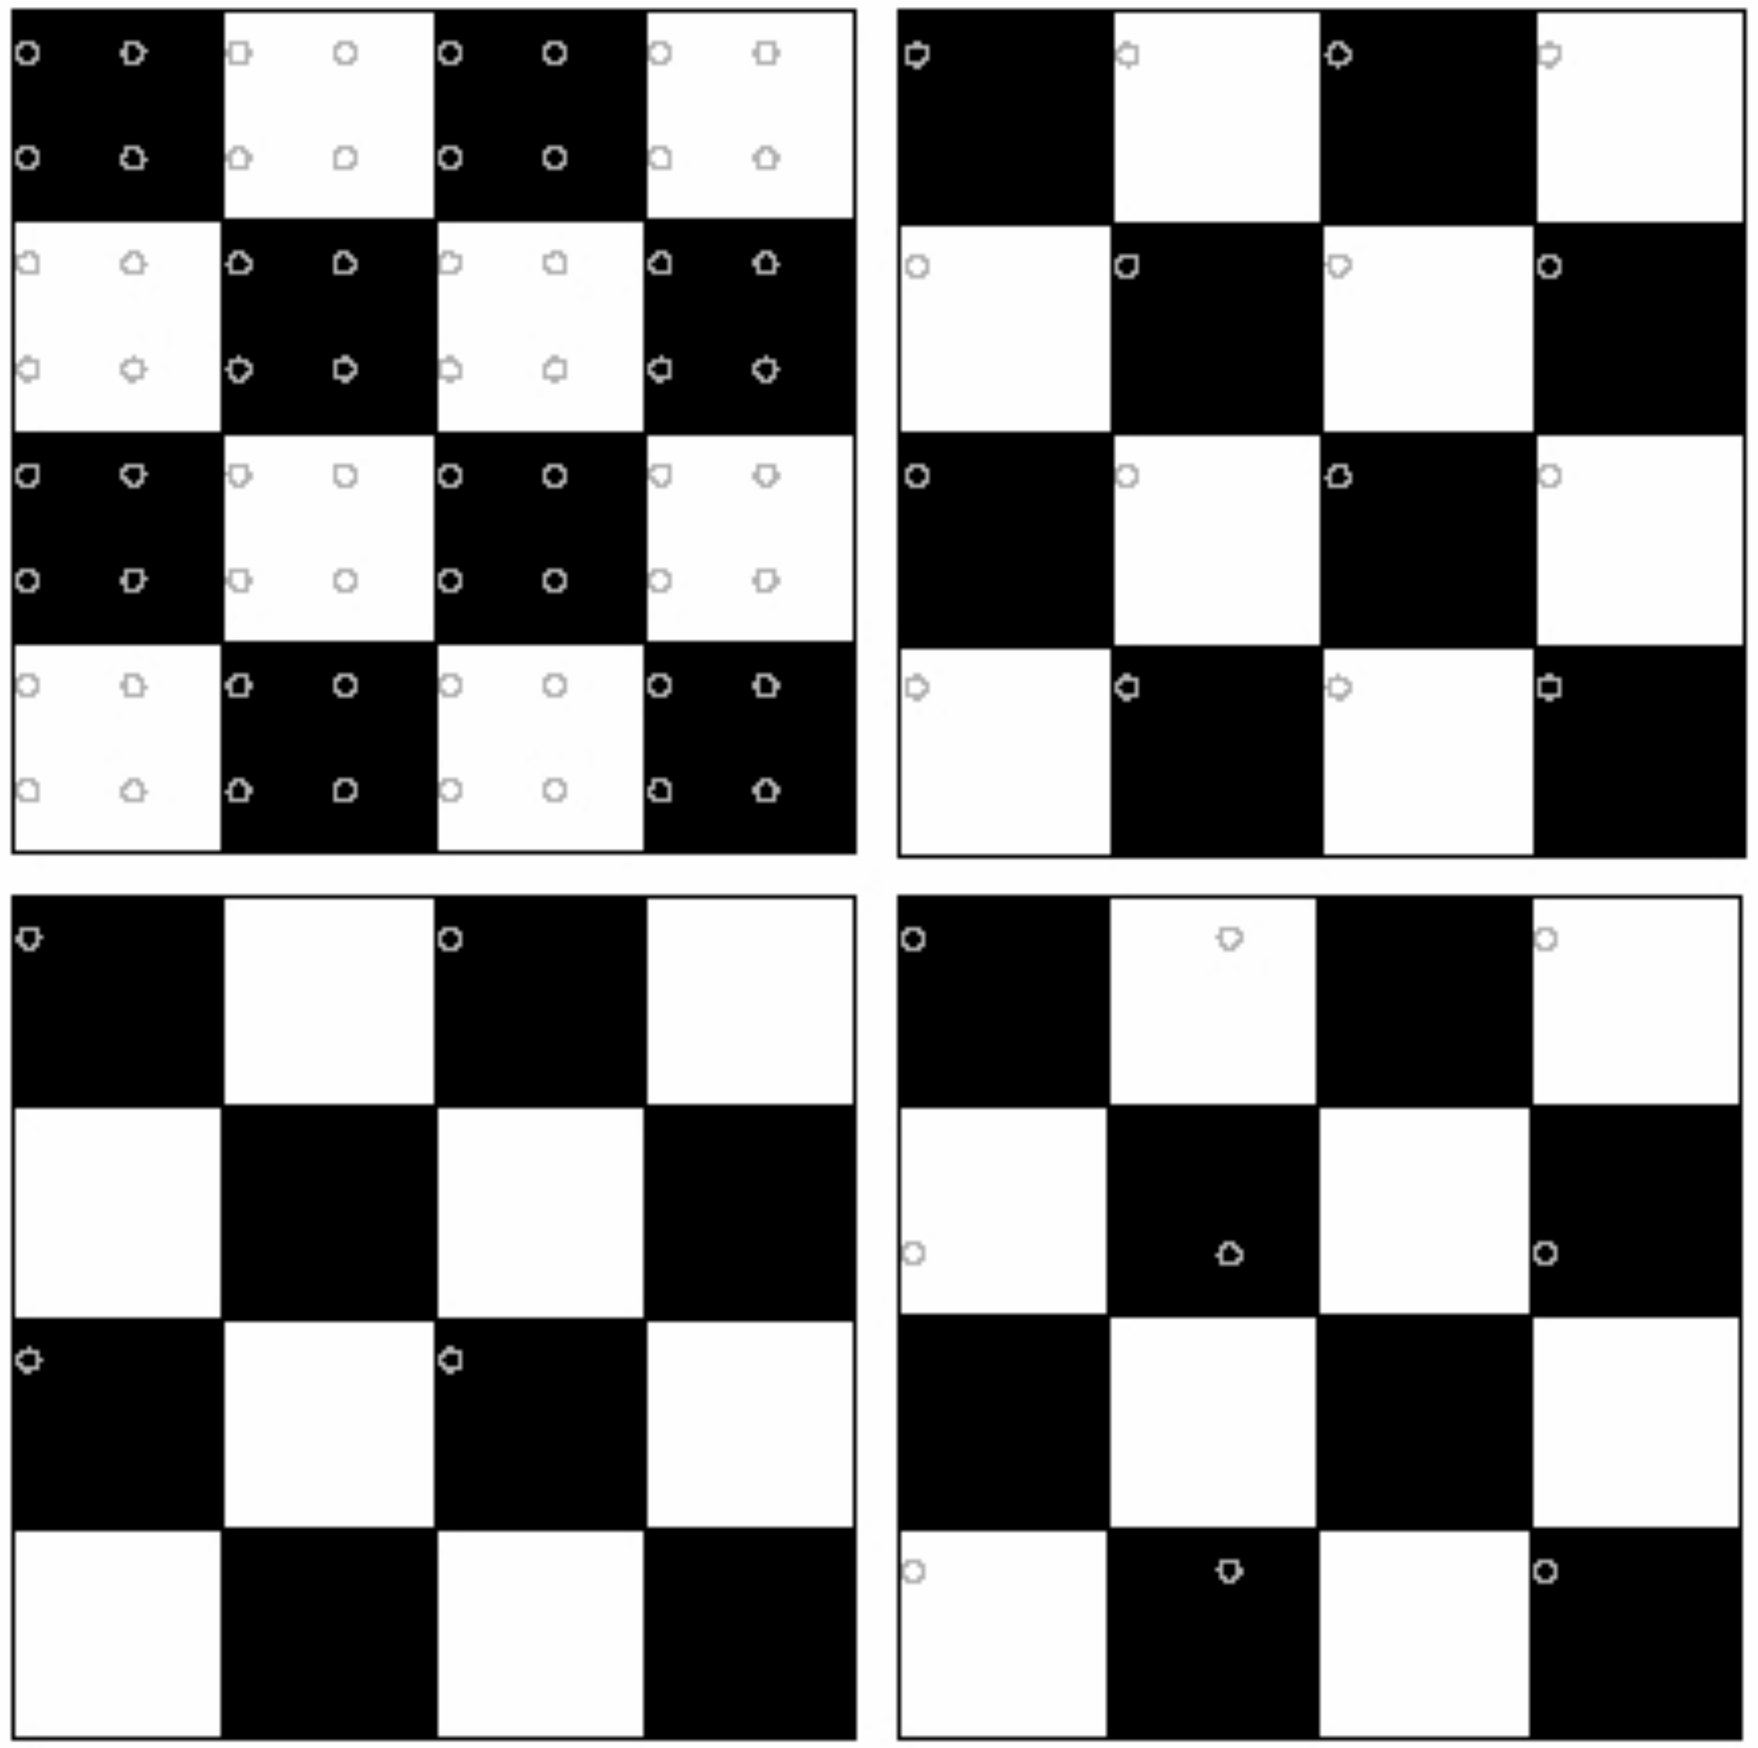

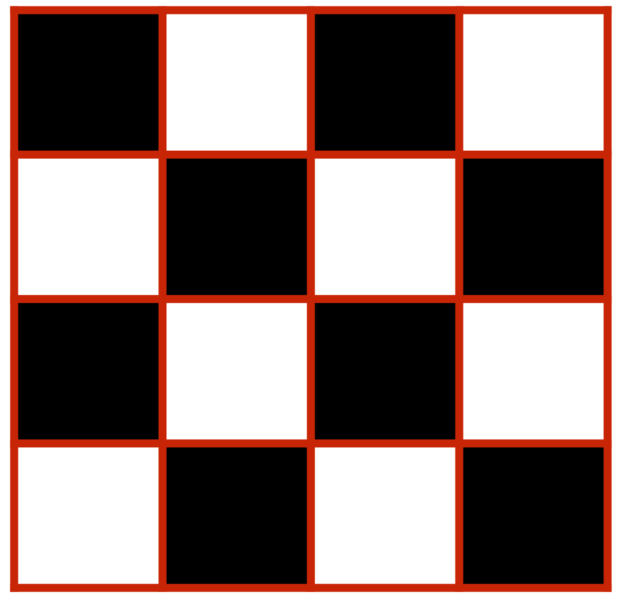

Forsyth & Ponce (2nd ed.) Figure 4.7

Sampling

It is clear that *some* information may be lost when we work on a discrete pixel grid.

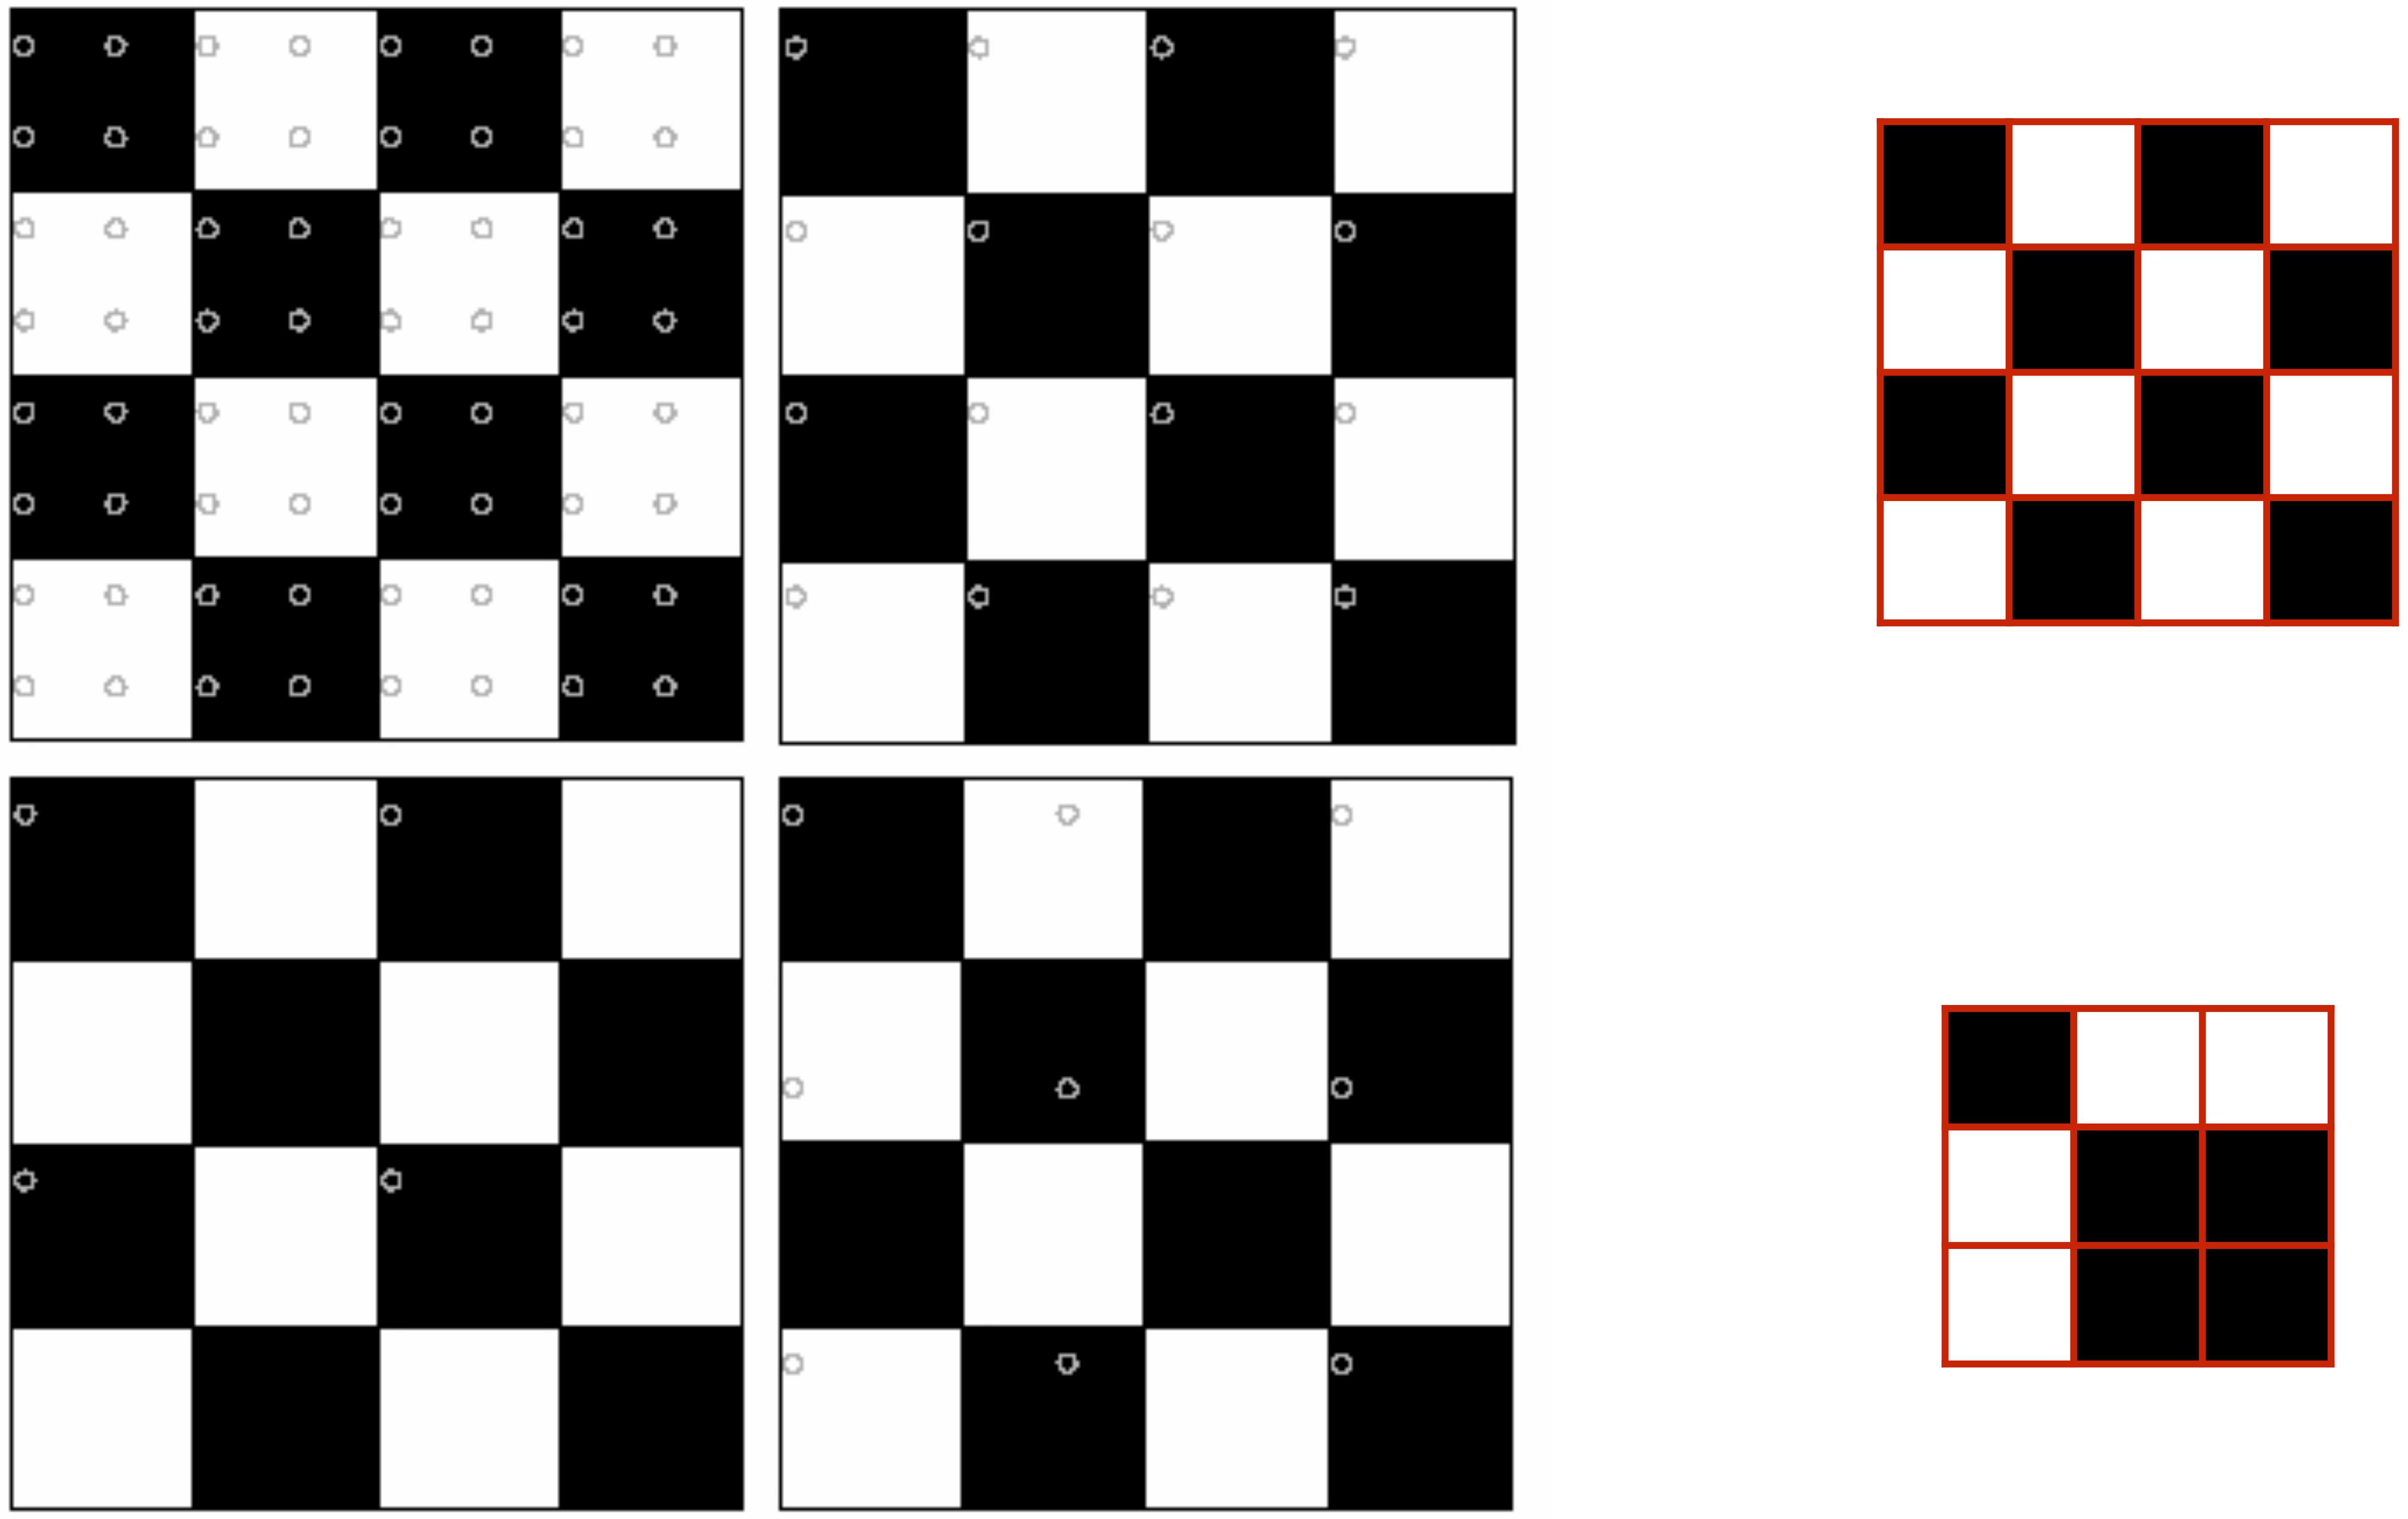

Forsyth & Ponce (2nd ed.) Figure 4.7

Sampling

It is clear that *some* information may be lost when we work on a discrete pixel grid.

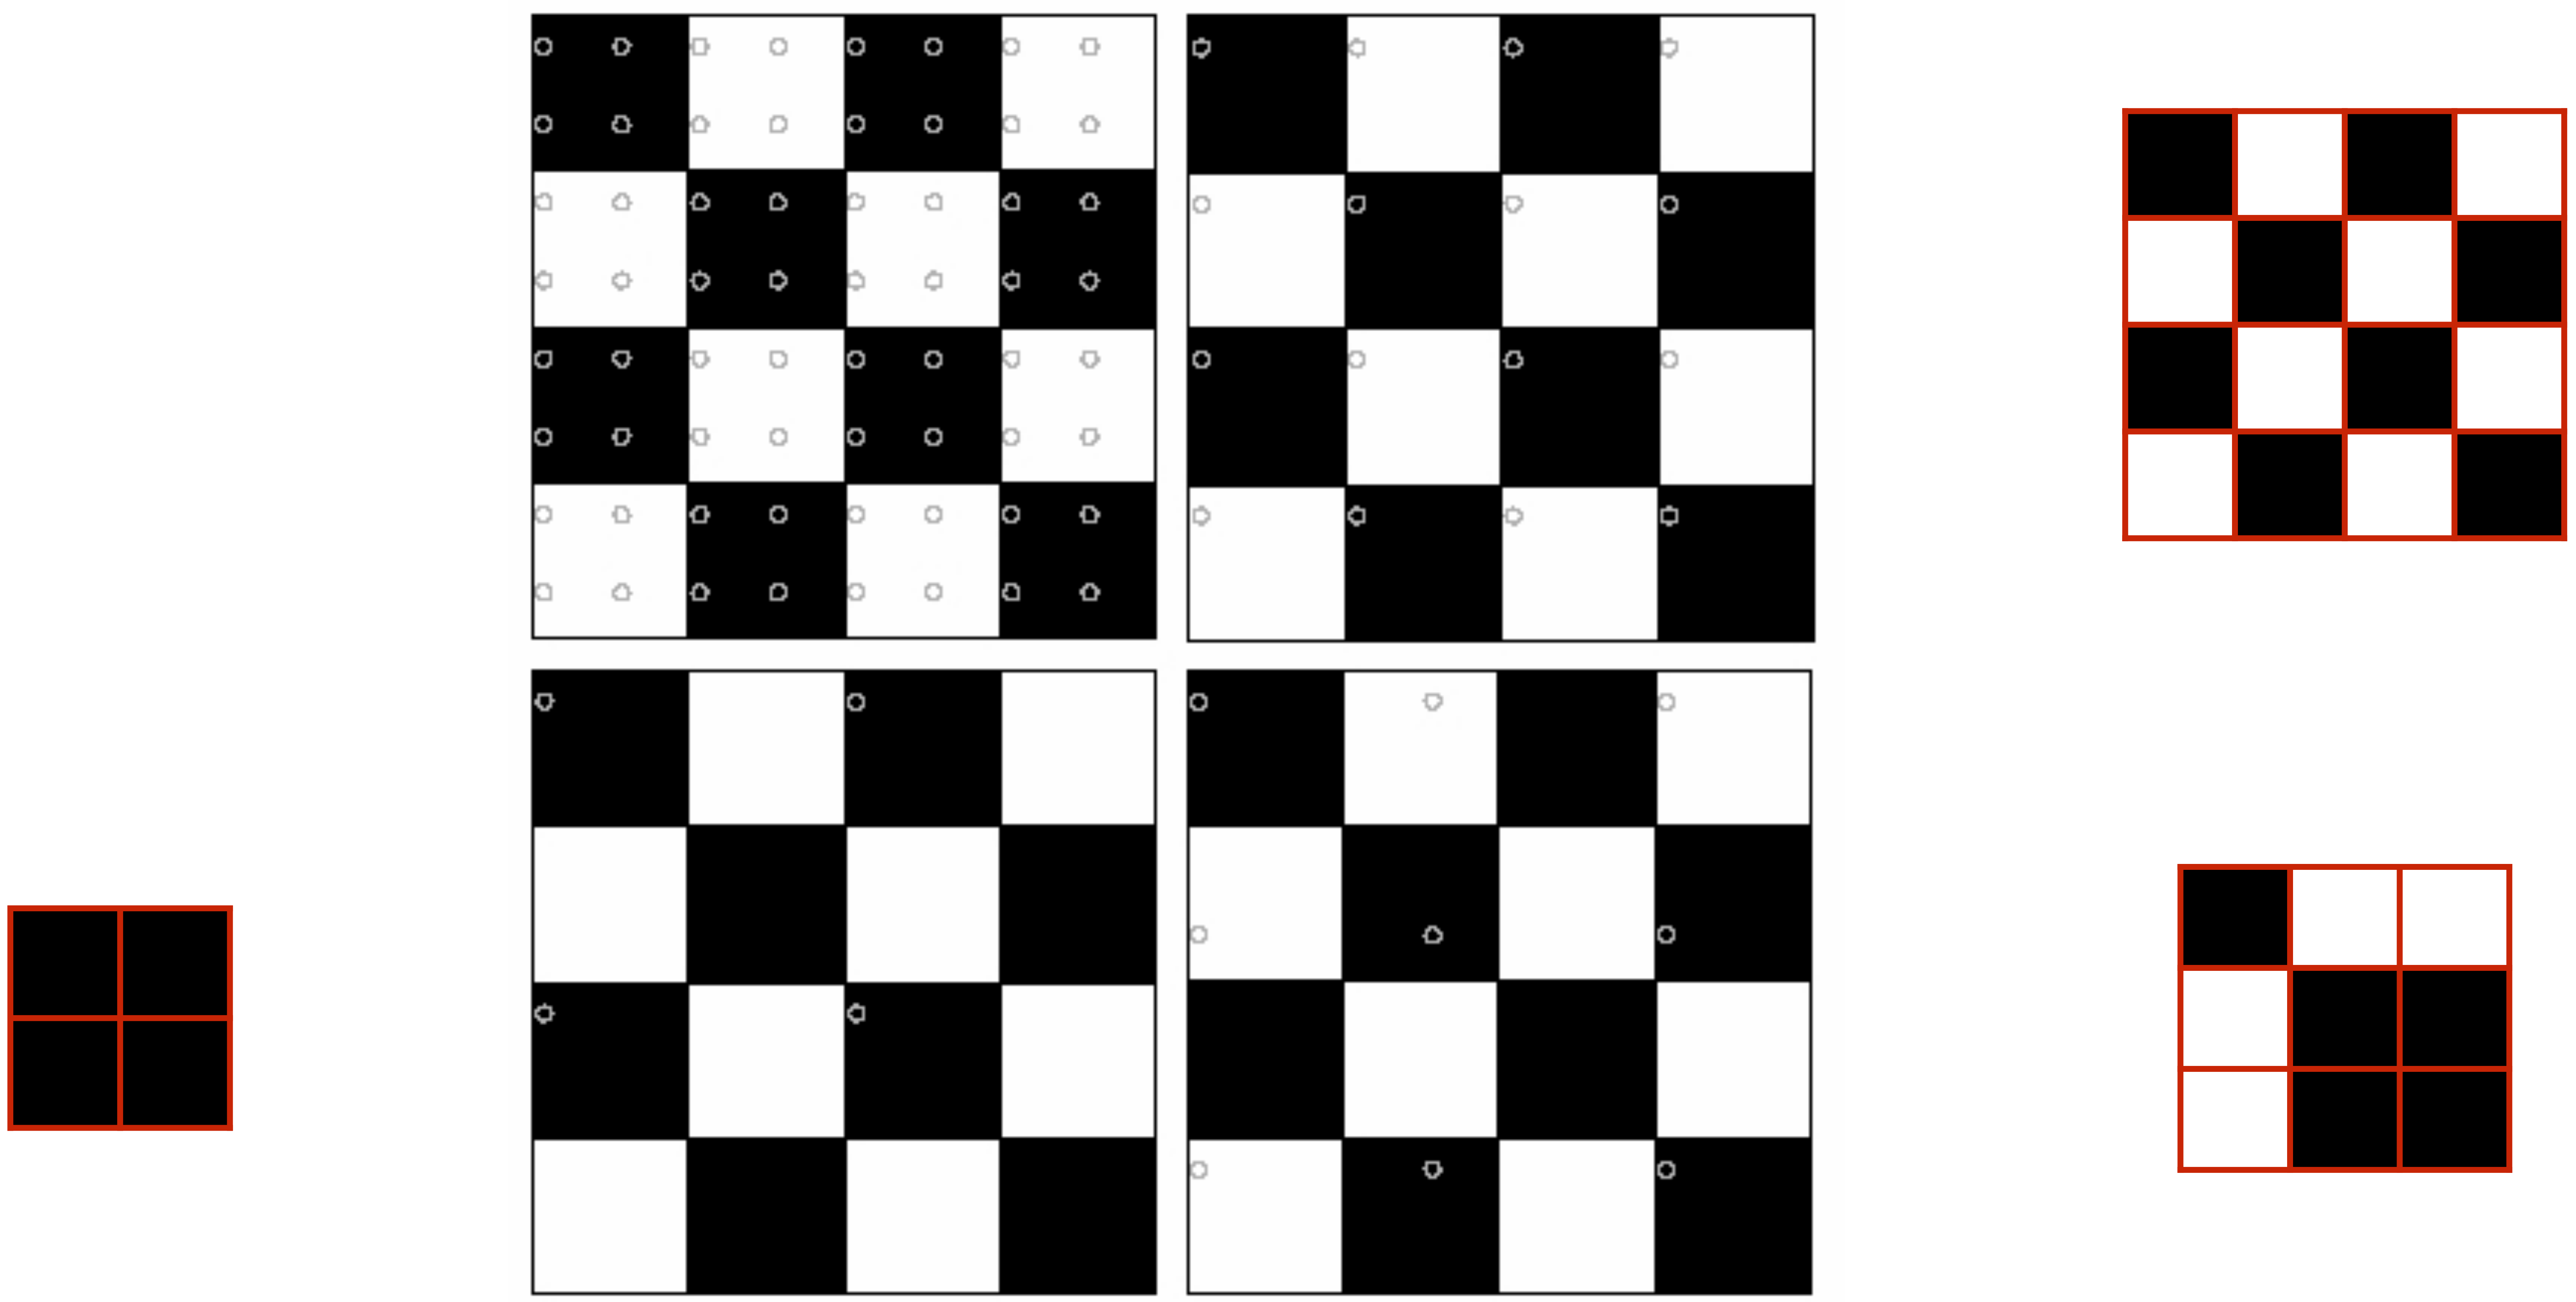

Forsyth & Ponce (2nd ed.) Figure 4.7

Sampling Theory (informal)

Question: When is $I(X, Y)$ an exact characterization of $i(x, y)$?

Sampling Theory (informal)

Question: When is $I(X, Y)$ an exact characterization of $i(x, y)$?

Question (modified): When can we reconstruct $i(x, y)$ exactly from $I(X, Y)$?

Sampling Theory (informal)

Question: When is $I(X,Y)$ an exact characterization of $i(x,y)$?

Question (modified): When can we reconstruct $i(x,y)$ exactly from $I(X,Y)$?

Intuition: Reconstruction involves some kind of **interpolation**

Sampling Theory (informal)

Question: When is $I(X, Y)$ an exact characterization of $i(x, y)$?

Question (modified): When can we reconstruct $i(x, y)$ exactly from $I(X, Y)$?

Intuition: Reconstruction involves some kind of **interpolation**

Heuristic: When in doubt, consider simple cases

Sampling Theory (informal)

Case 0: Suppose $i(x, y) = k$ (with k being one of our gray levels)

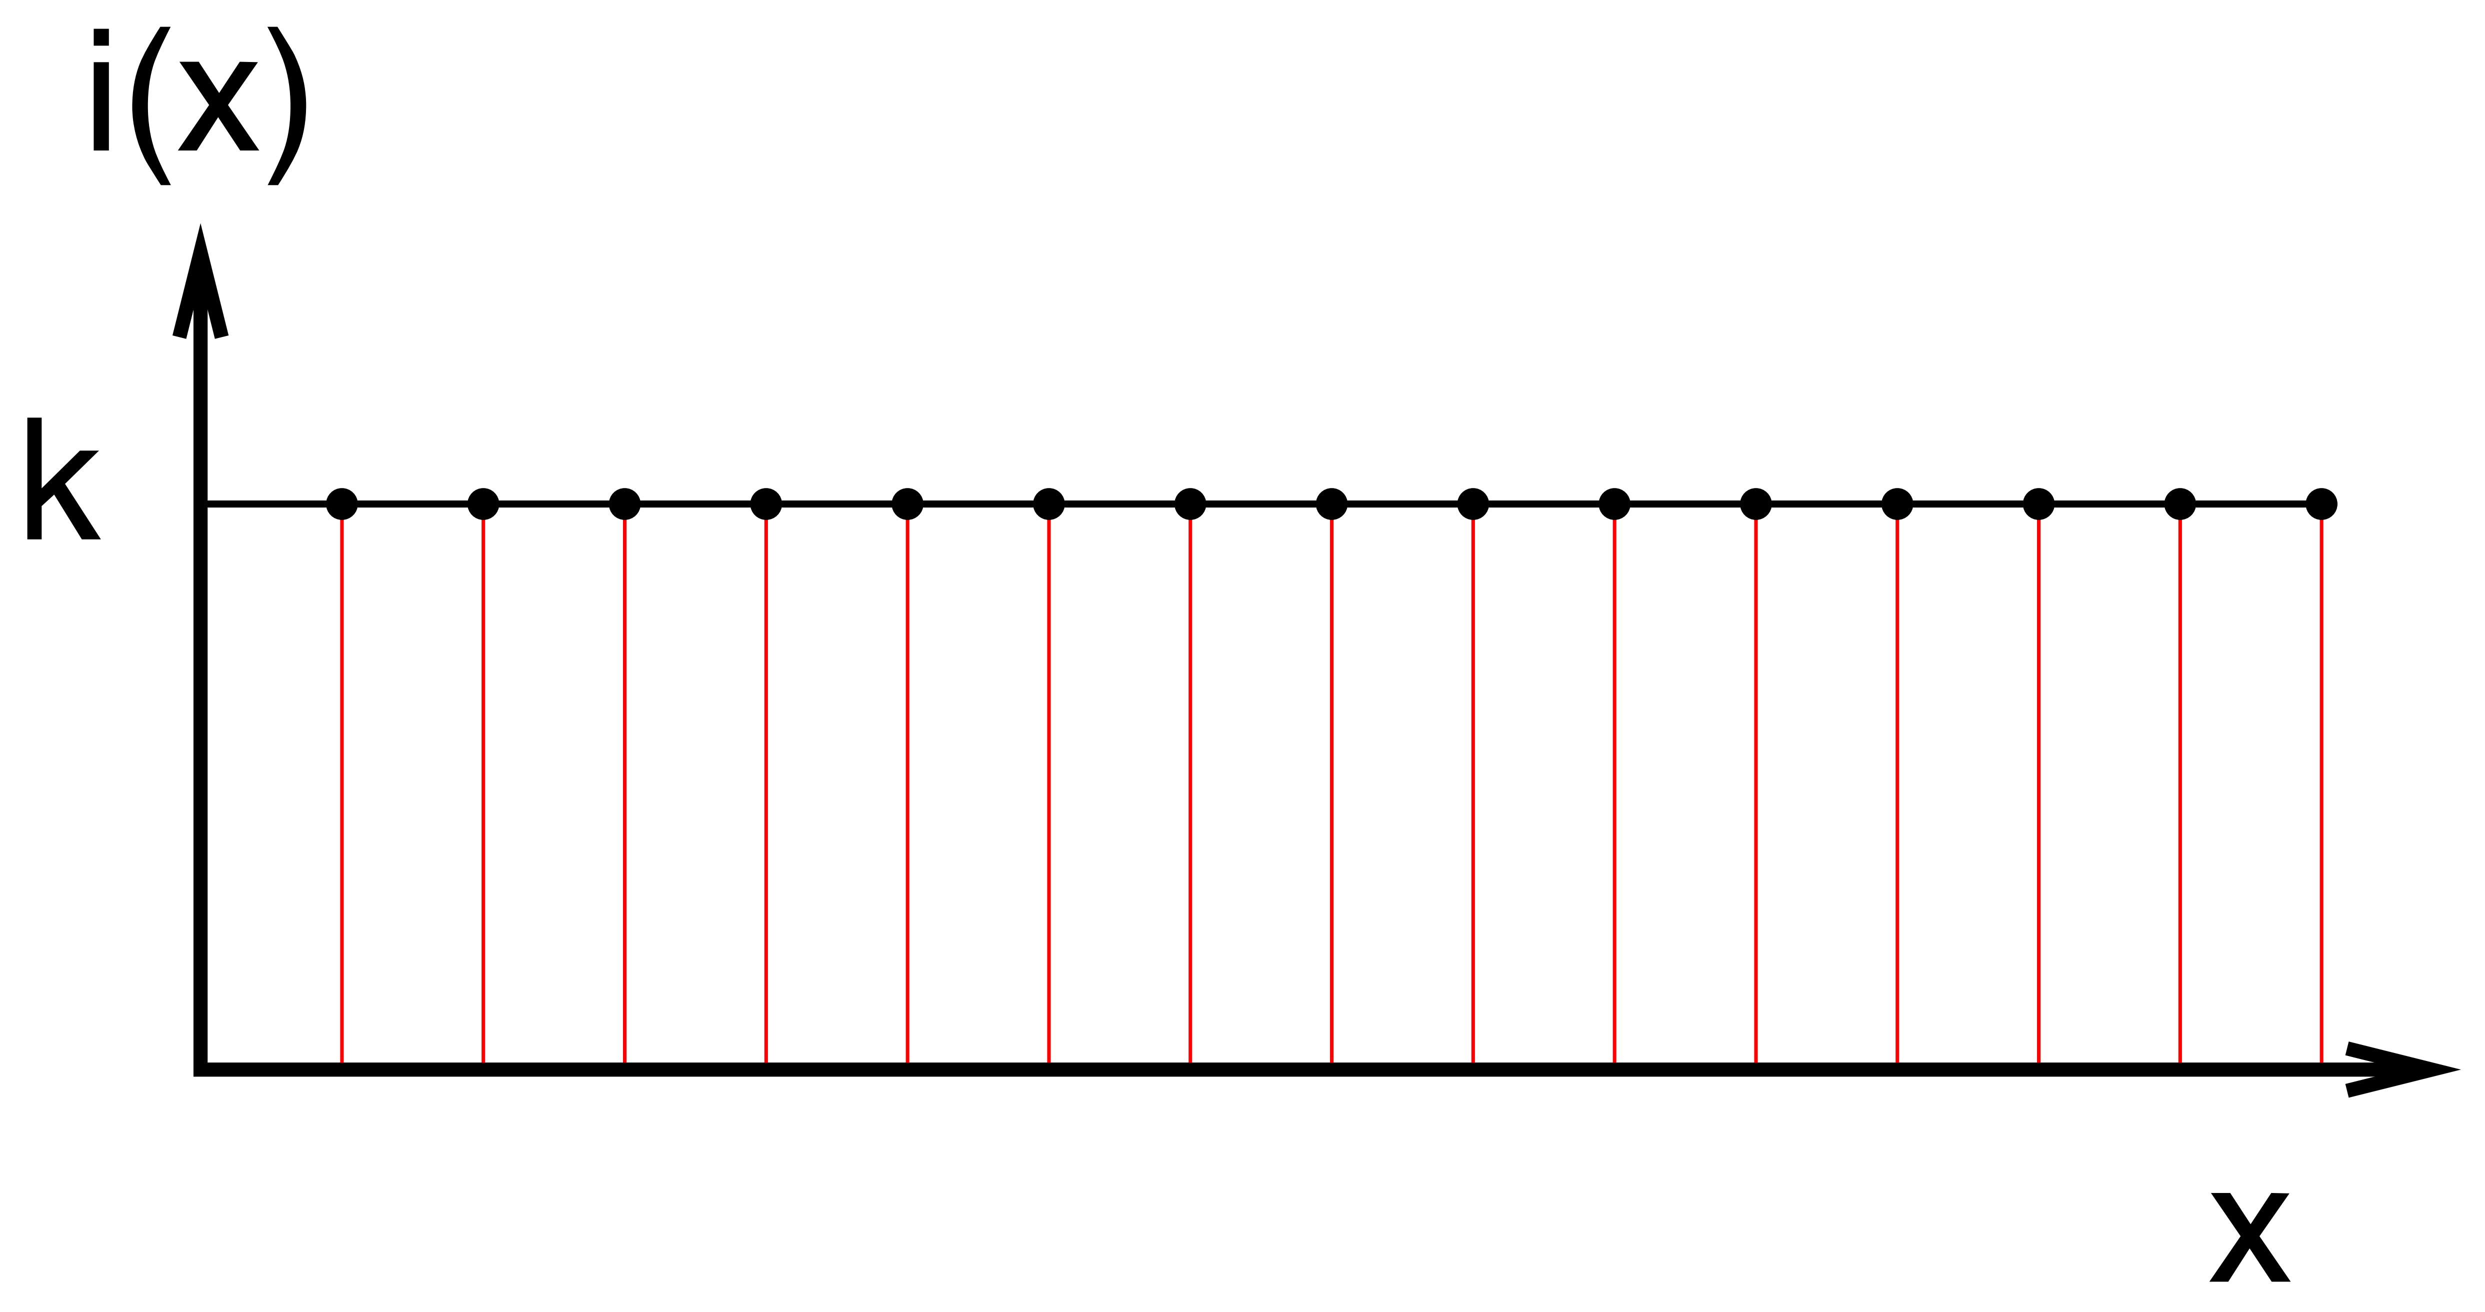

Note: we use equidistant sampling at integer values for convenience, in general, sampling doesn't need to be equidistant

Sampling Theory (informal)

Case 0: Suppose $i(x, y) = k$ (with k being one of our gray levels)

This is **easy**!

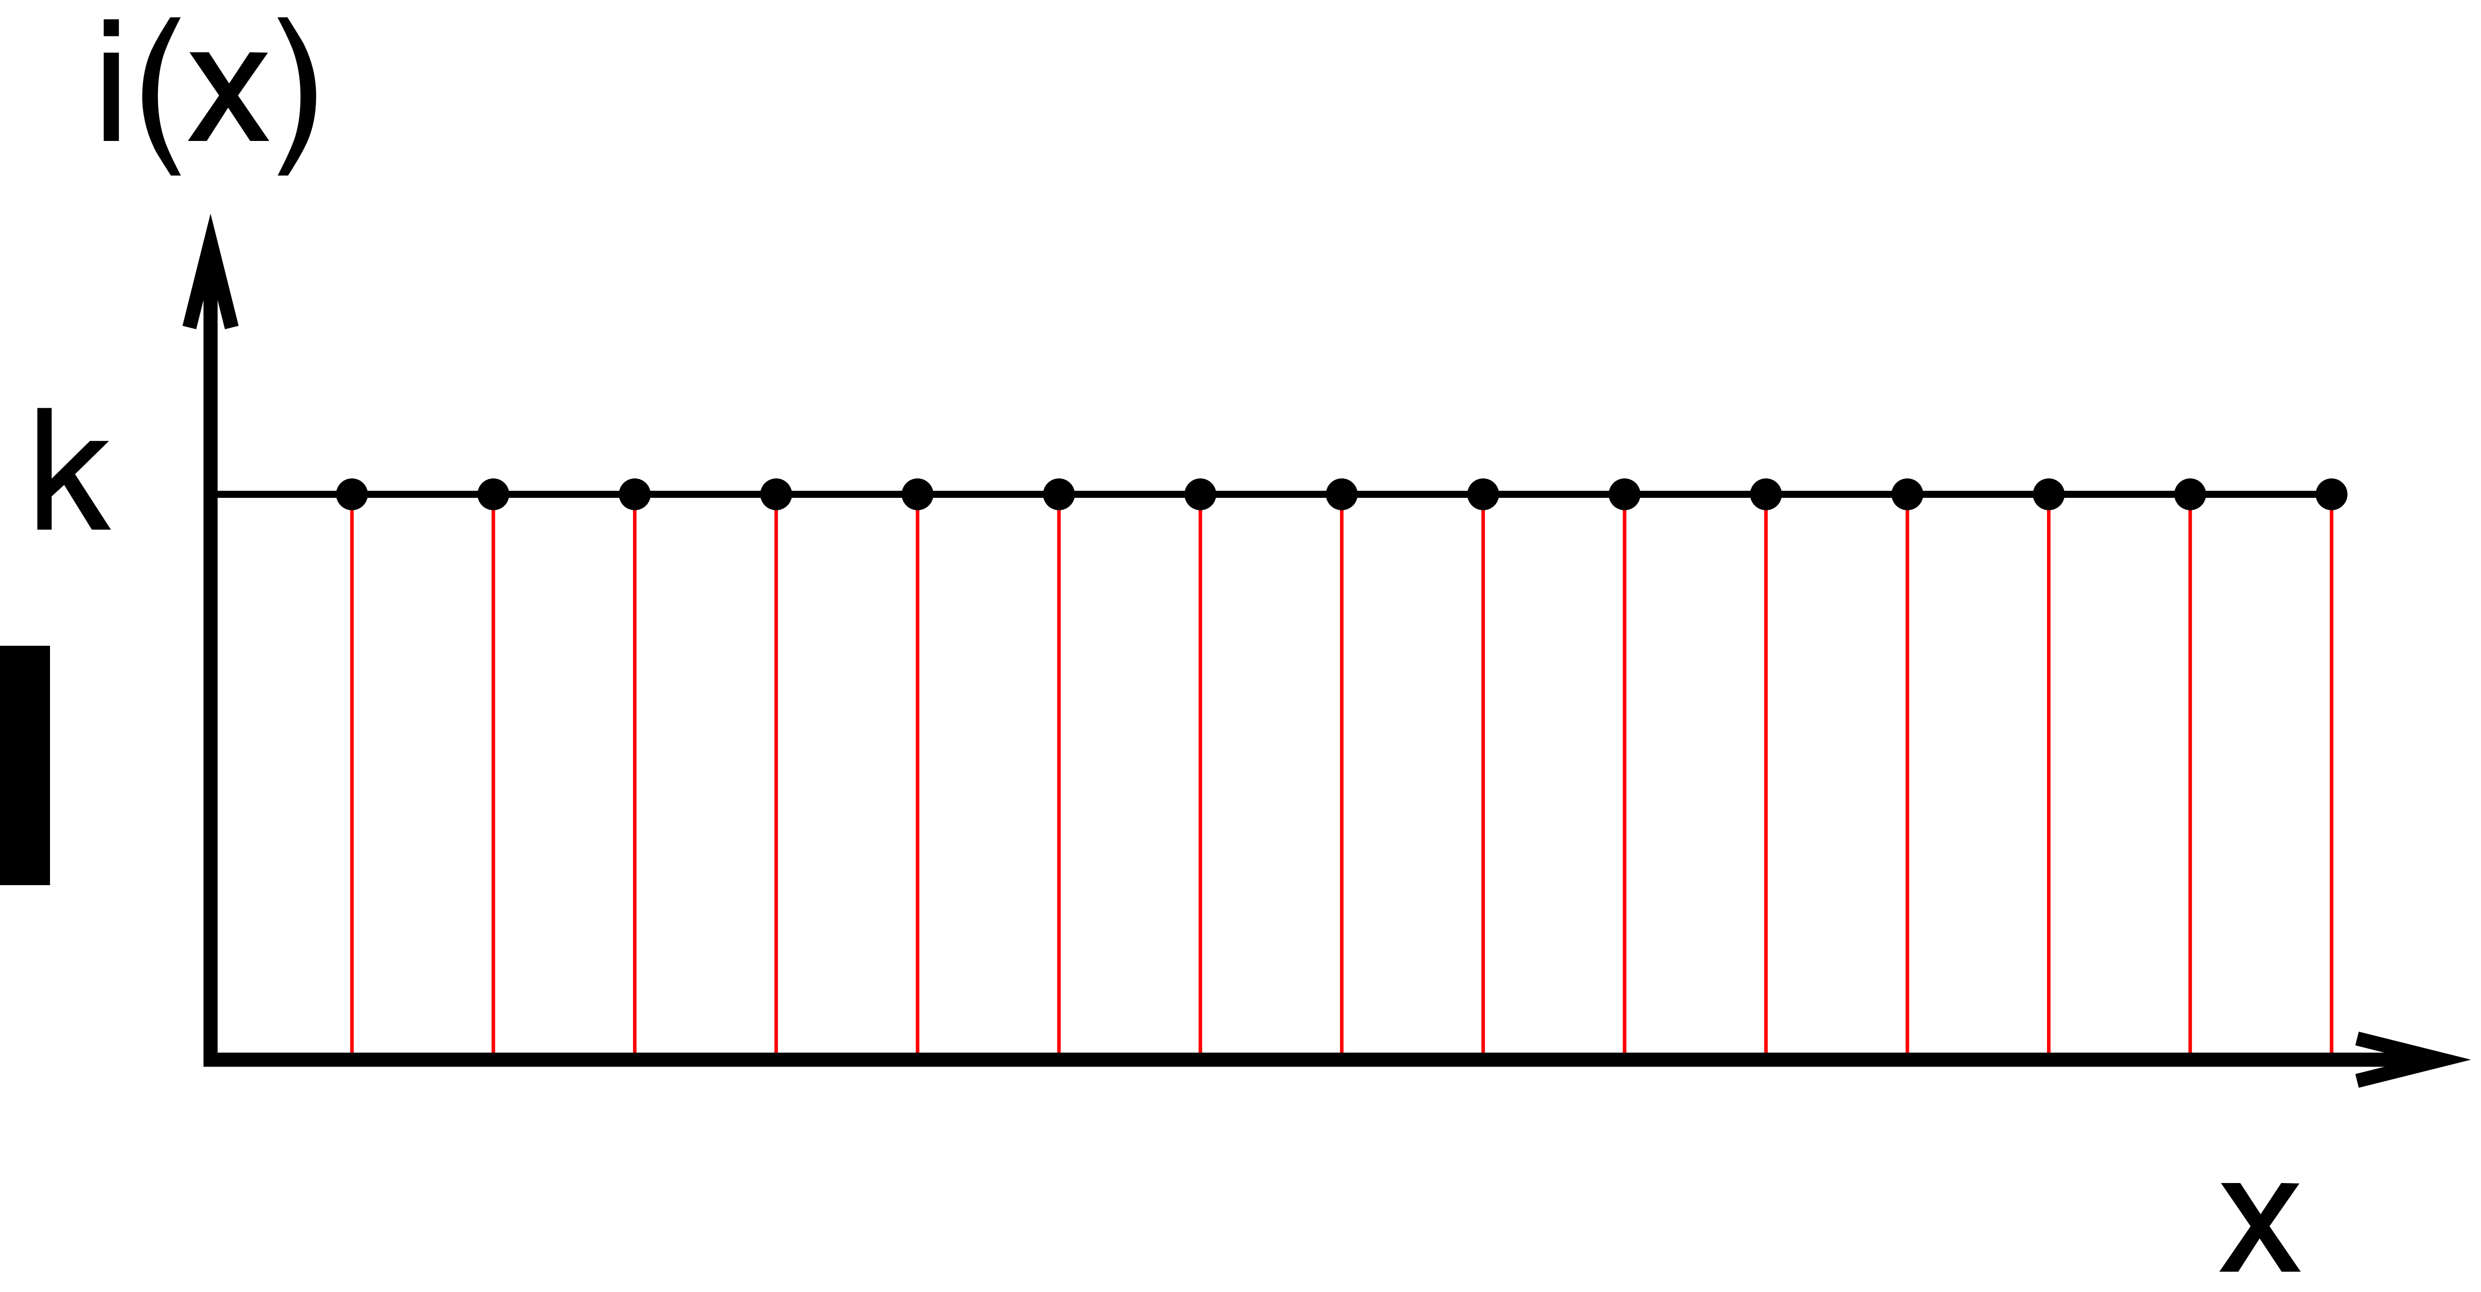

Sampling Theory (informal)

Case 0: Suppose $i(x, y) = k$ (with k being one of our gray levels)

This is **easy**!

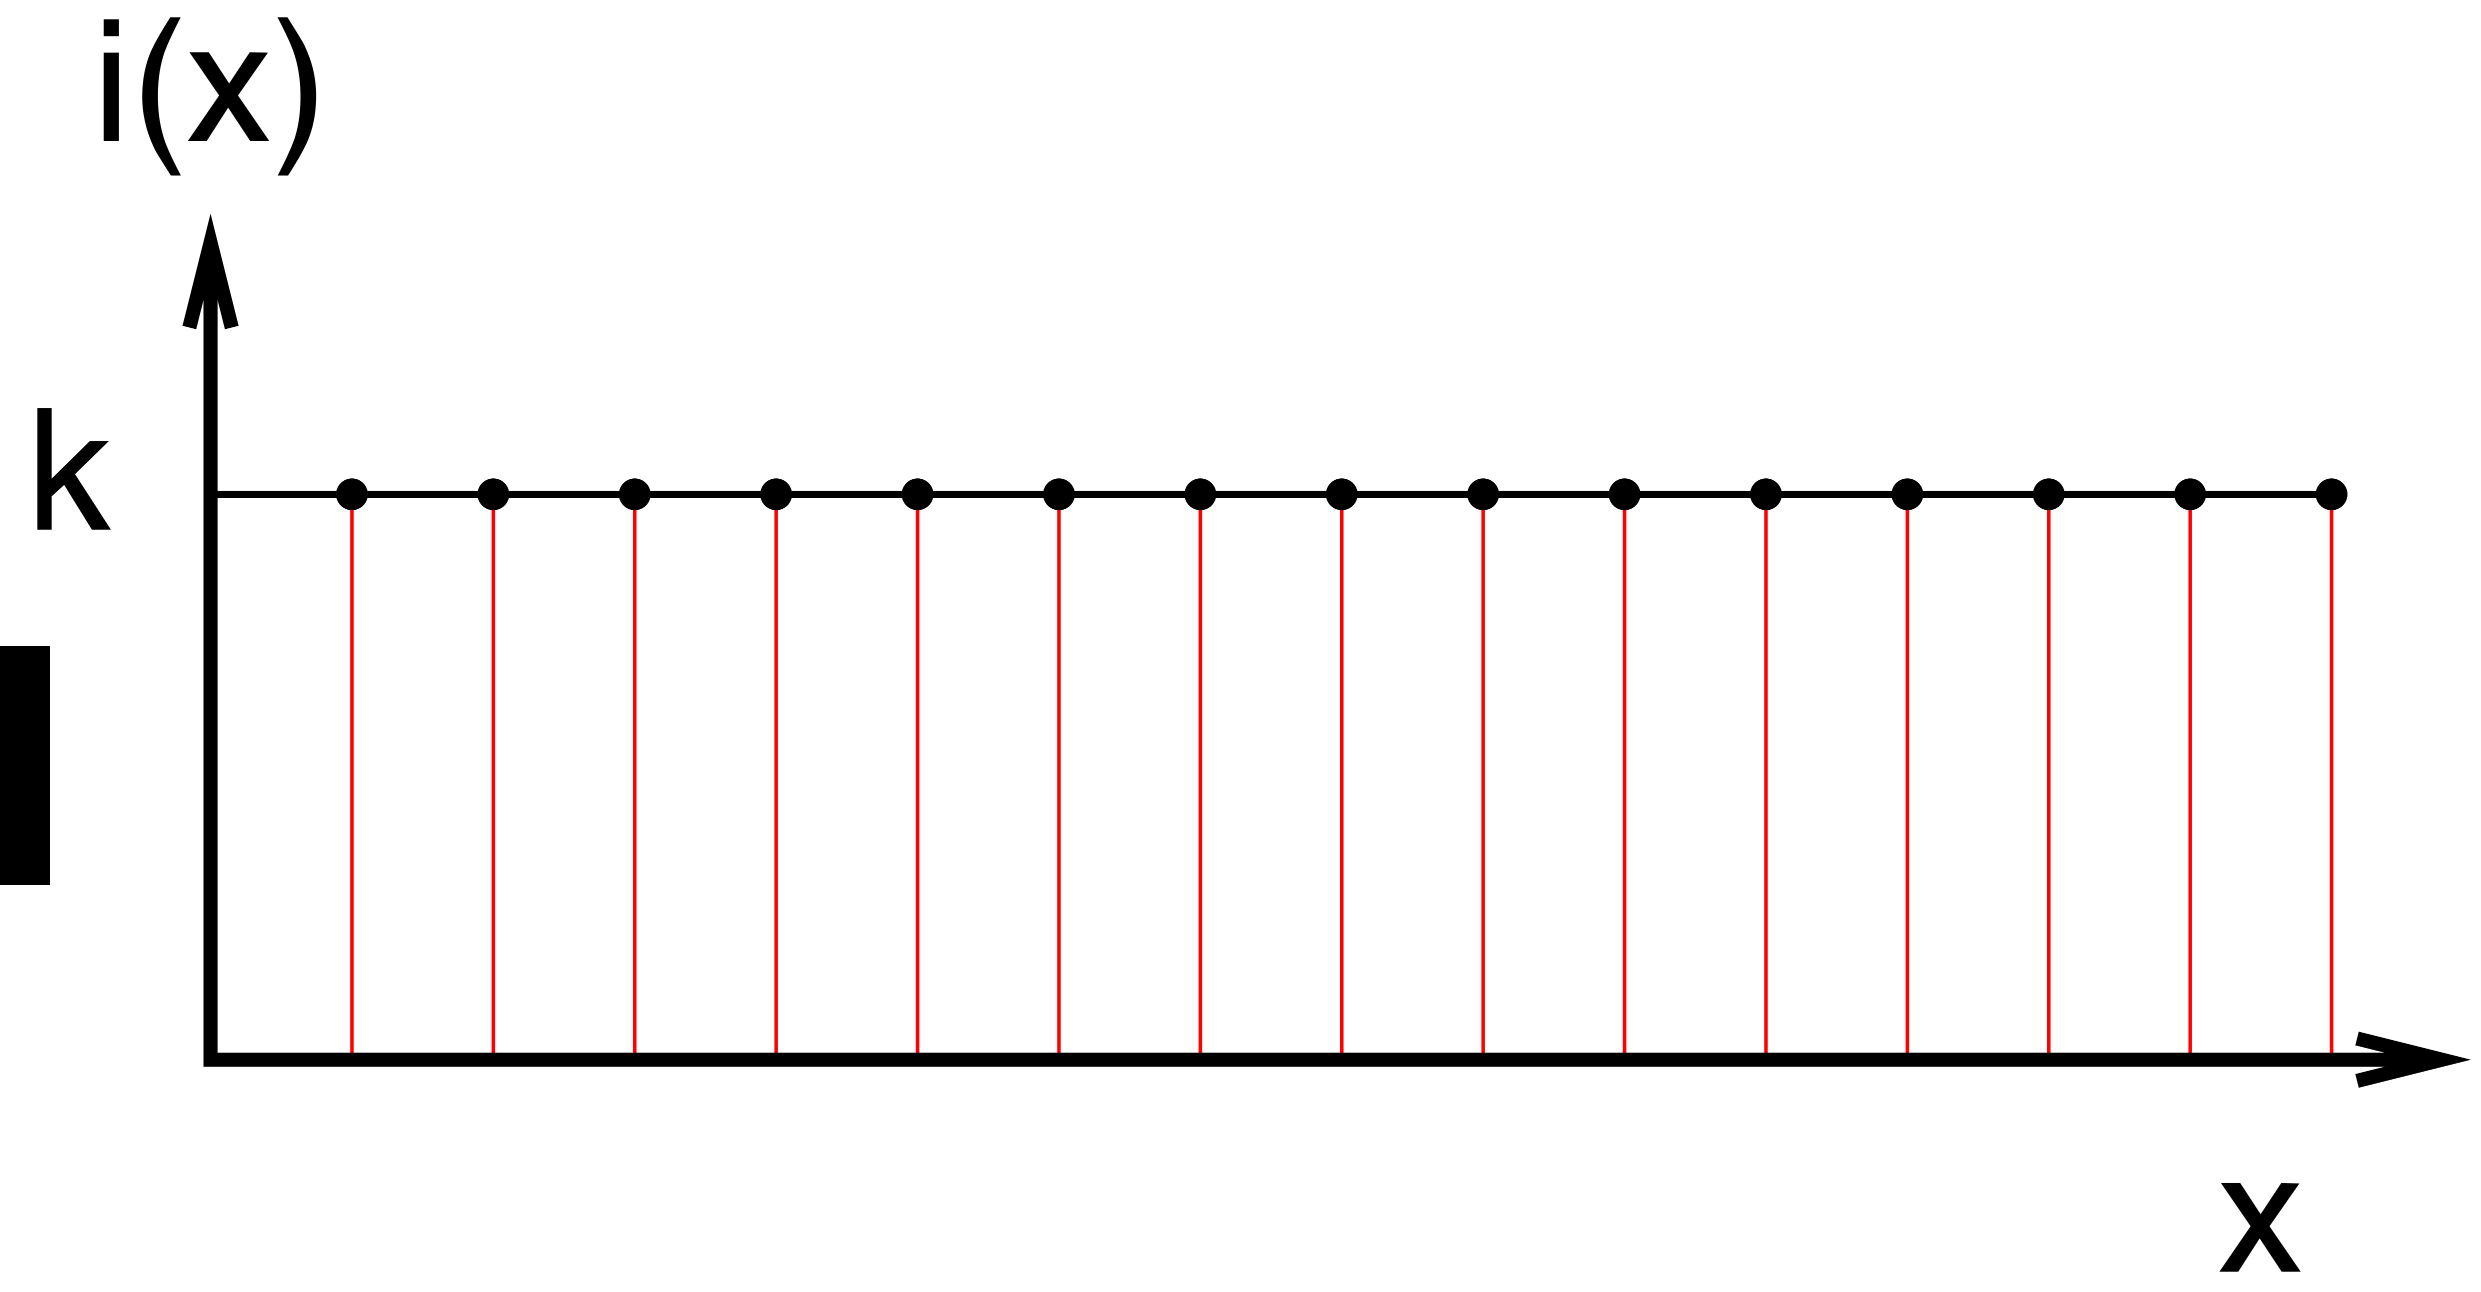

$I(X, Y) = k$. Any standard interpolation function would give $i(x, y) = k$ for non-integer x and y (irrespective oh how coarse the sampling is)

Sampling Theory (informal)

Case 1: Suppose $i(x, y)$ has a discontinuity not falling precisely at integer x, y

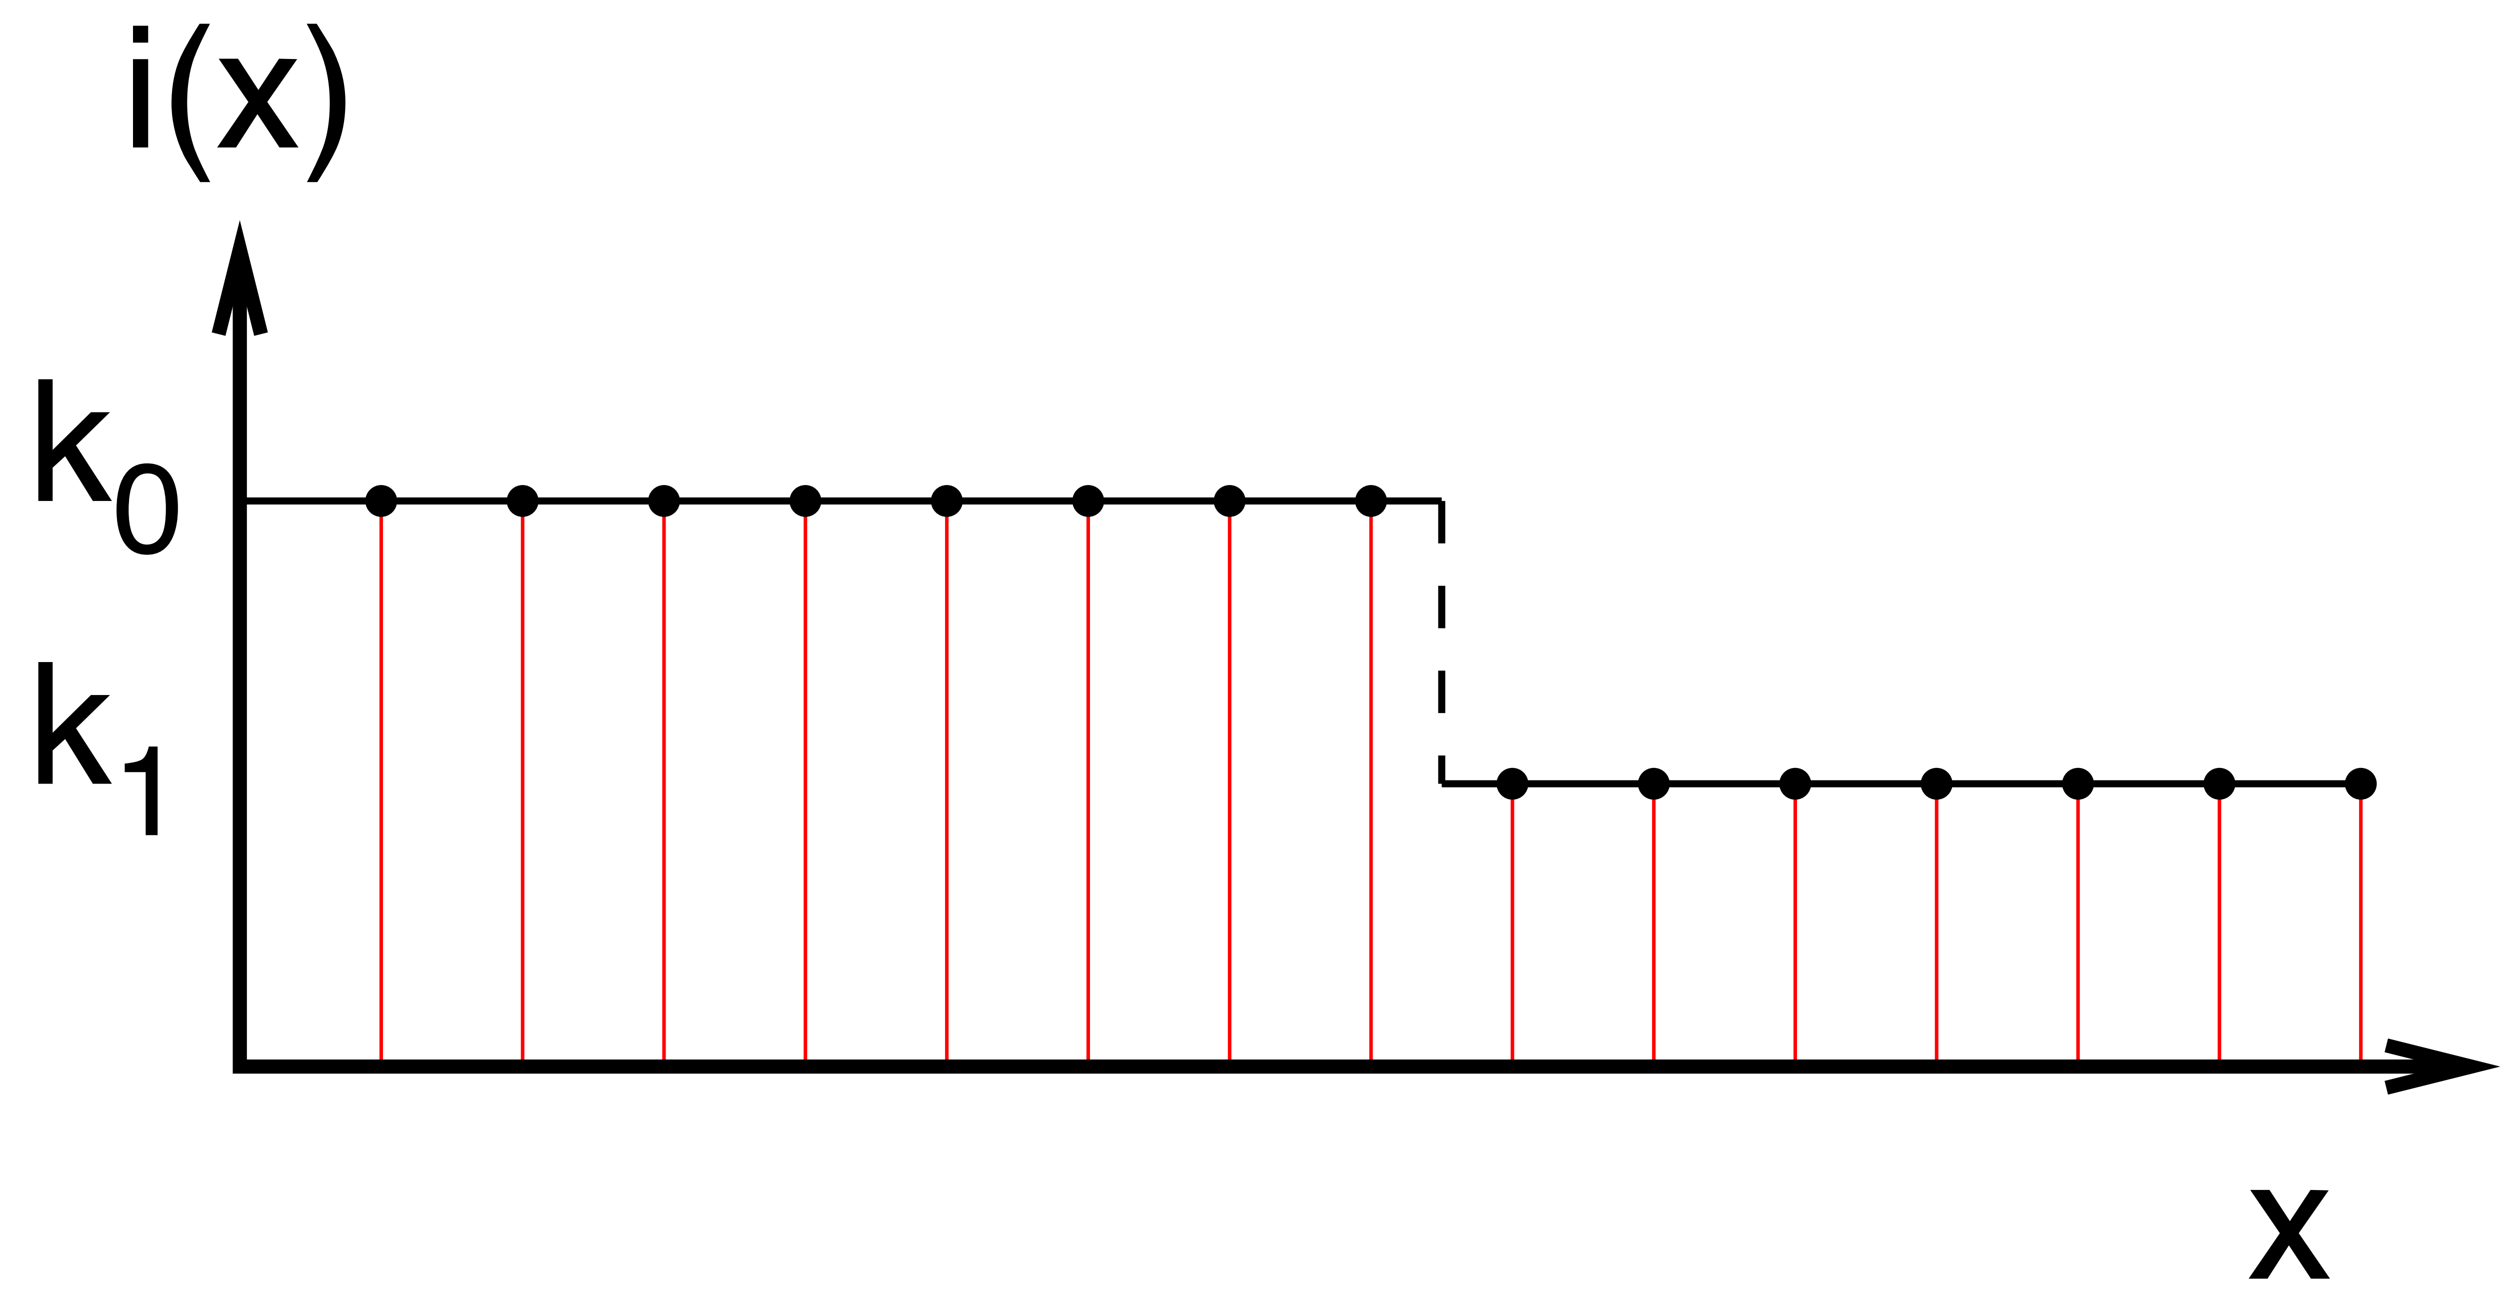

Sampling Theory (informal)

Case 1: Suppose $i(x, y)$ has a discontinuity not falling precisely at integer x, y

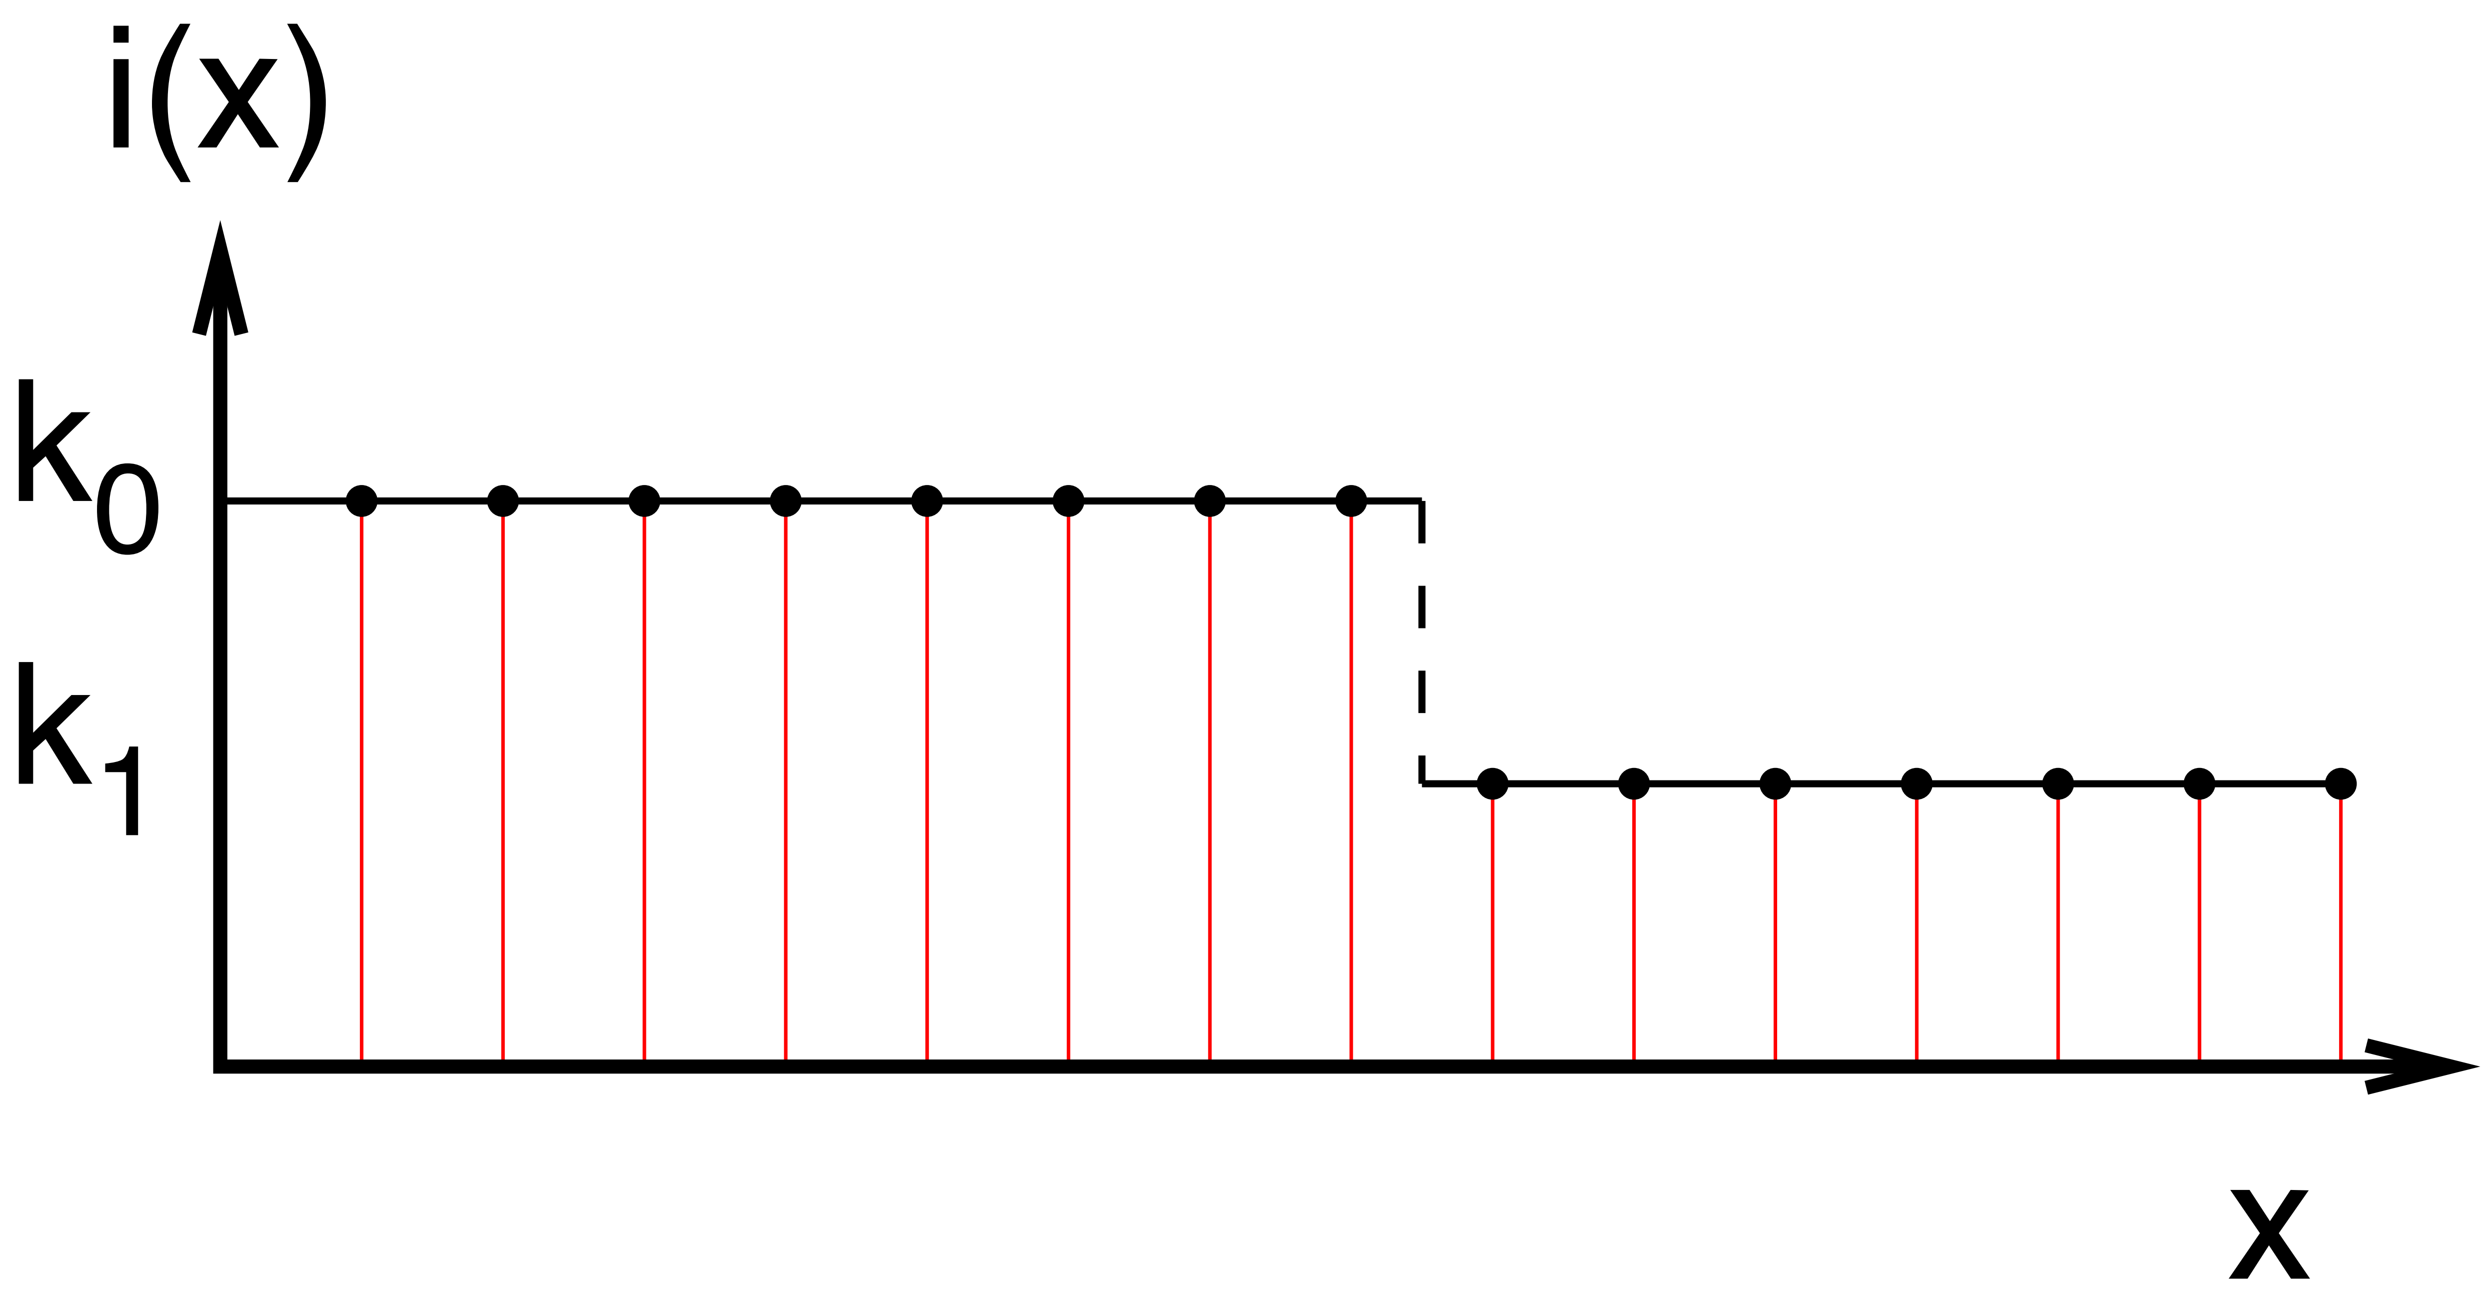

We cannot reconstruct $i(x, y)$ exactly because we can never know exactly where the discontinuity lies

Sampling Theory (informal)

Case 1: Suppose $i(x, y)$ has a discontinuity not falling precisely at integer x, y

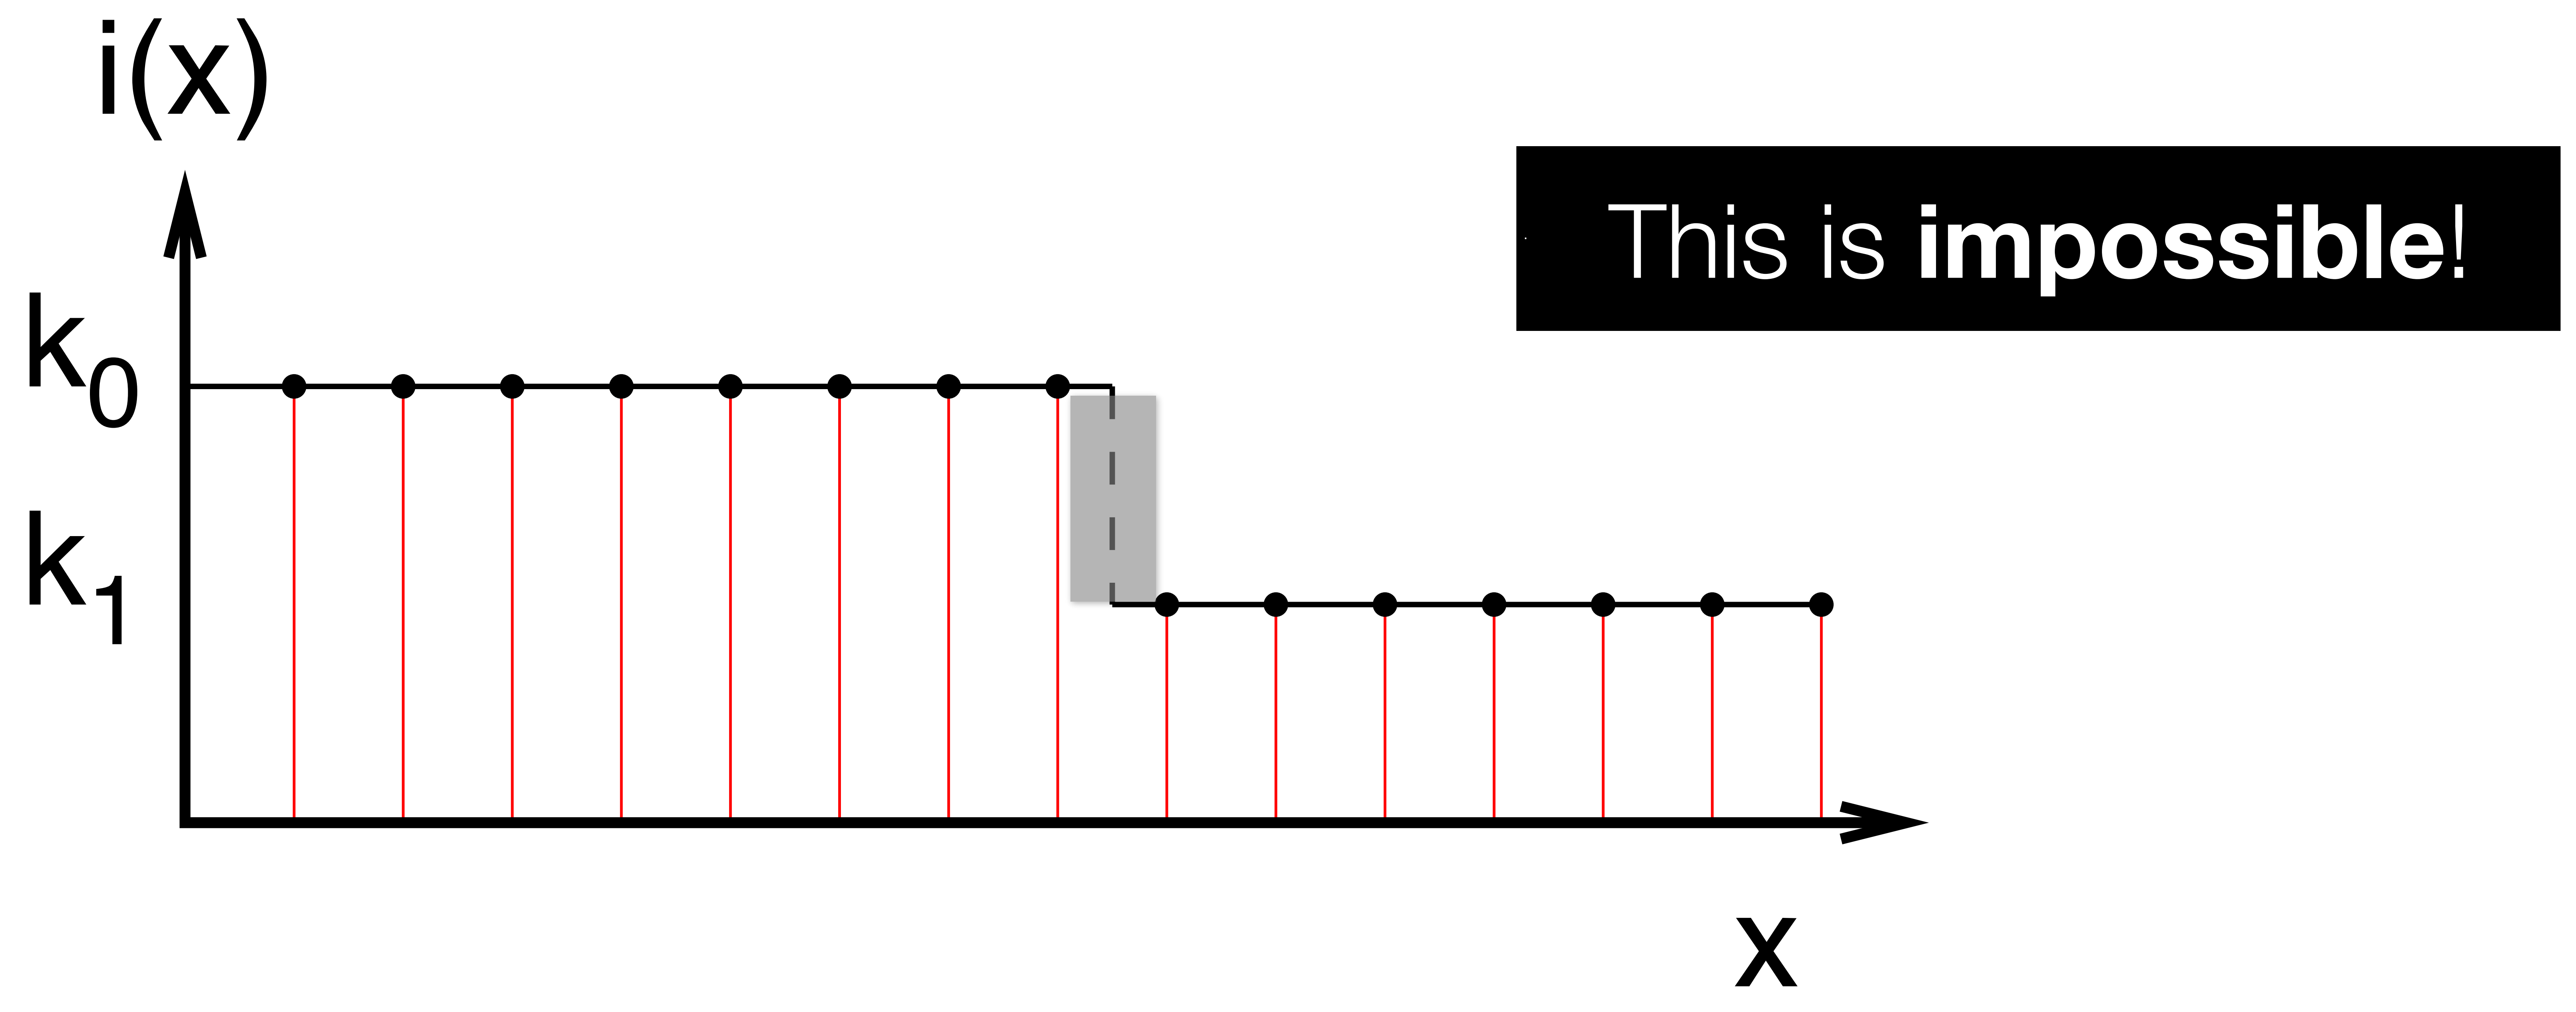

We cannot reconstruct $i(x, y)$ exactly because we can never know exactly where the discontinuity lies

Sampling Theory (informal)

Question: How do we close the gap between “**easy**” and “**impossible**?”

Next, we build intuition based on informal argument

Sampling Theory (informal)

Exact reconstruction requires constraint on the rate at which i(x,y) can change between samples

— “rate of change” means derivative

— the formal concept is **bandlimited signal**

— “bandlimit” and “constraint on derivative” are linked

Think of music

— bandlimited if it has some maximum **temporal frequency**

— the upper limit of human hearing is about 20 kHz

Think of imaging systems. Resolving power is measured in

— “line pairs per mm” (for a bar test pattern)

— “cycles per mm” (for a sine wave test pattern)

An image is bandlimited if it has some maximum **spatial frequency**

Example: A Simple Sine Wave

How do we discretize the signal?

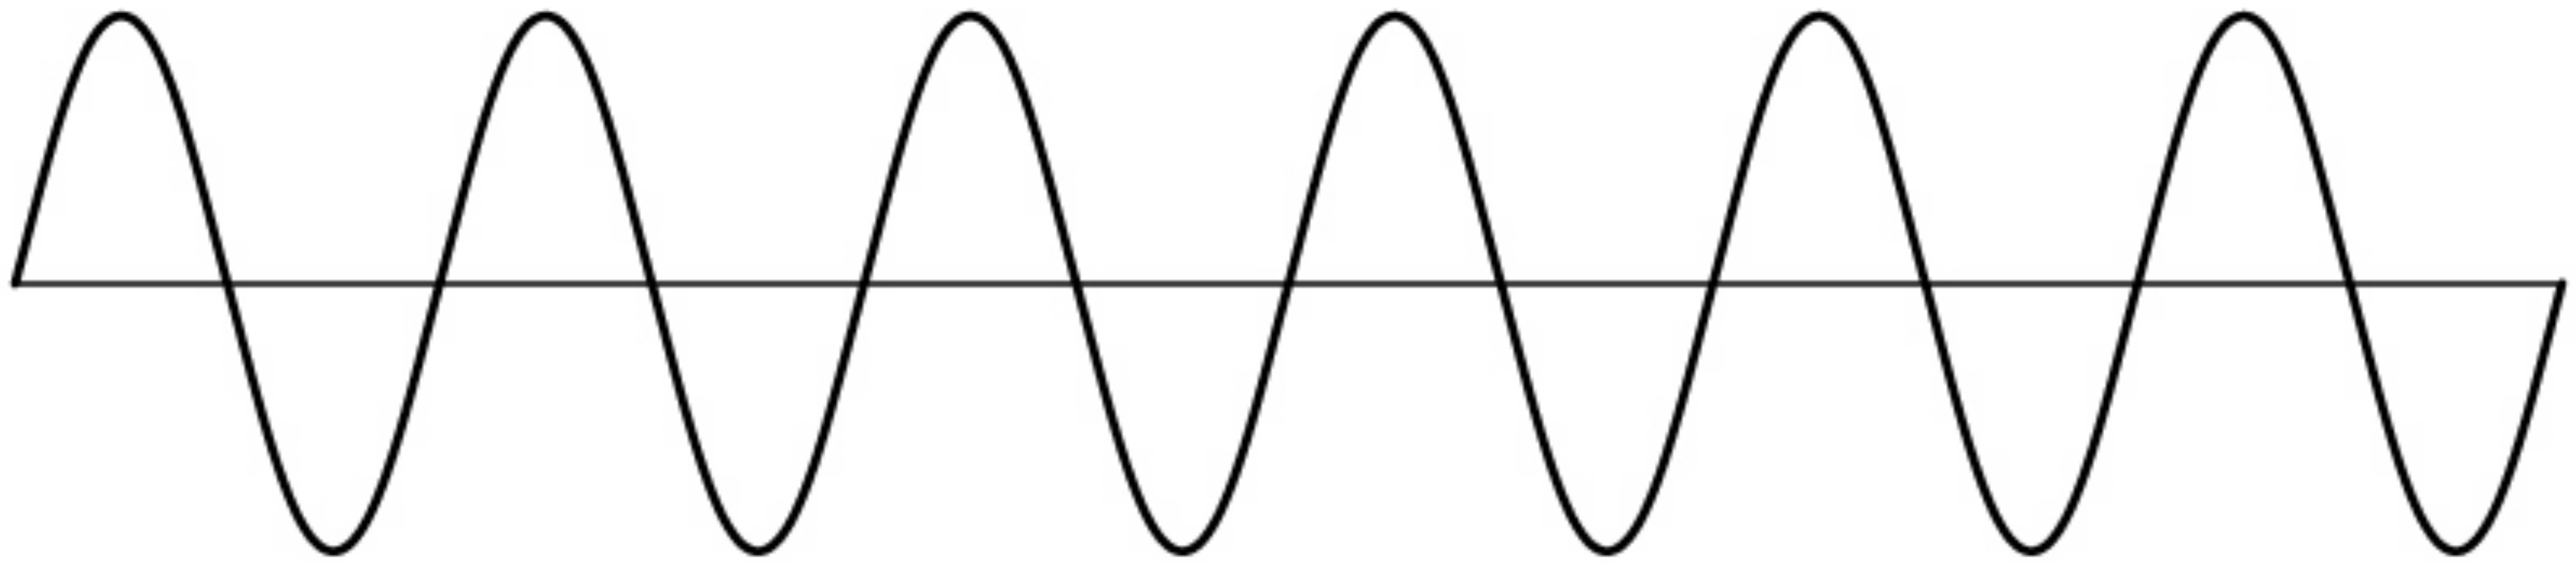

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

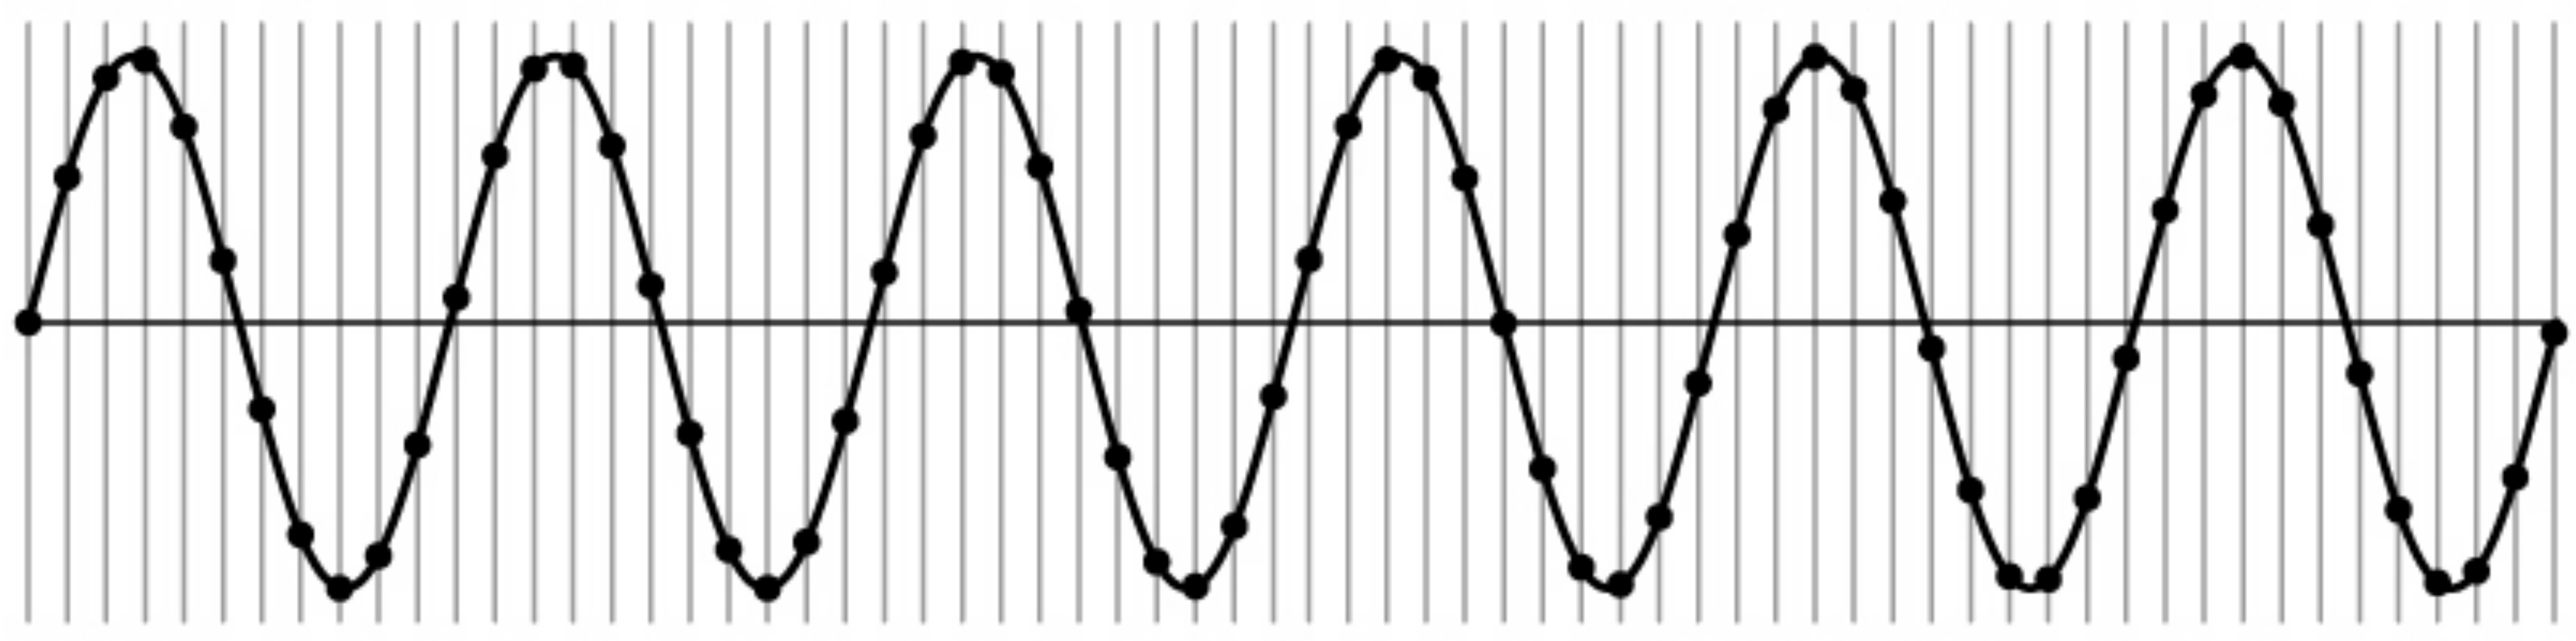

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

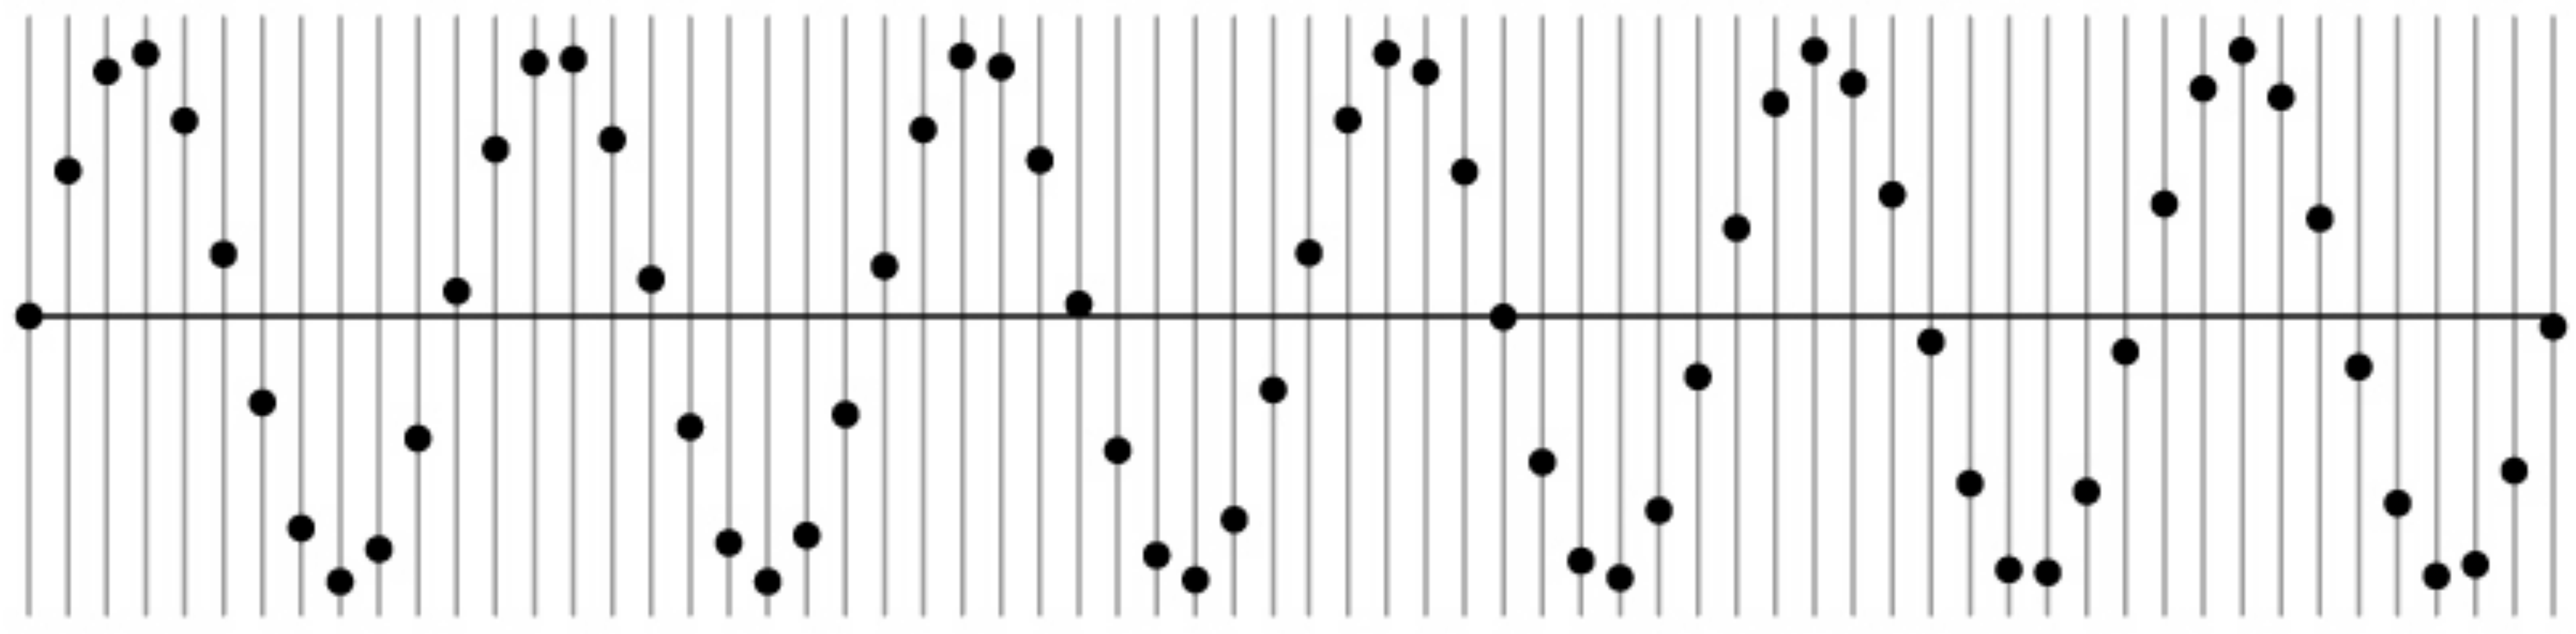

How many samples should I take?

Can I take as many samples as I want?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

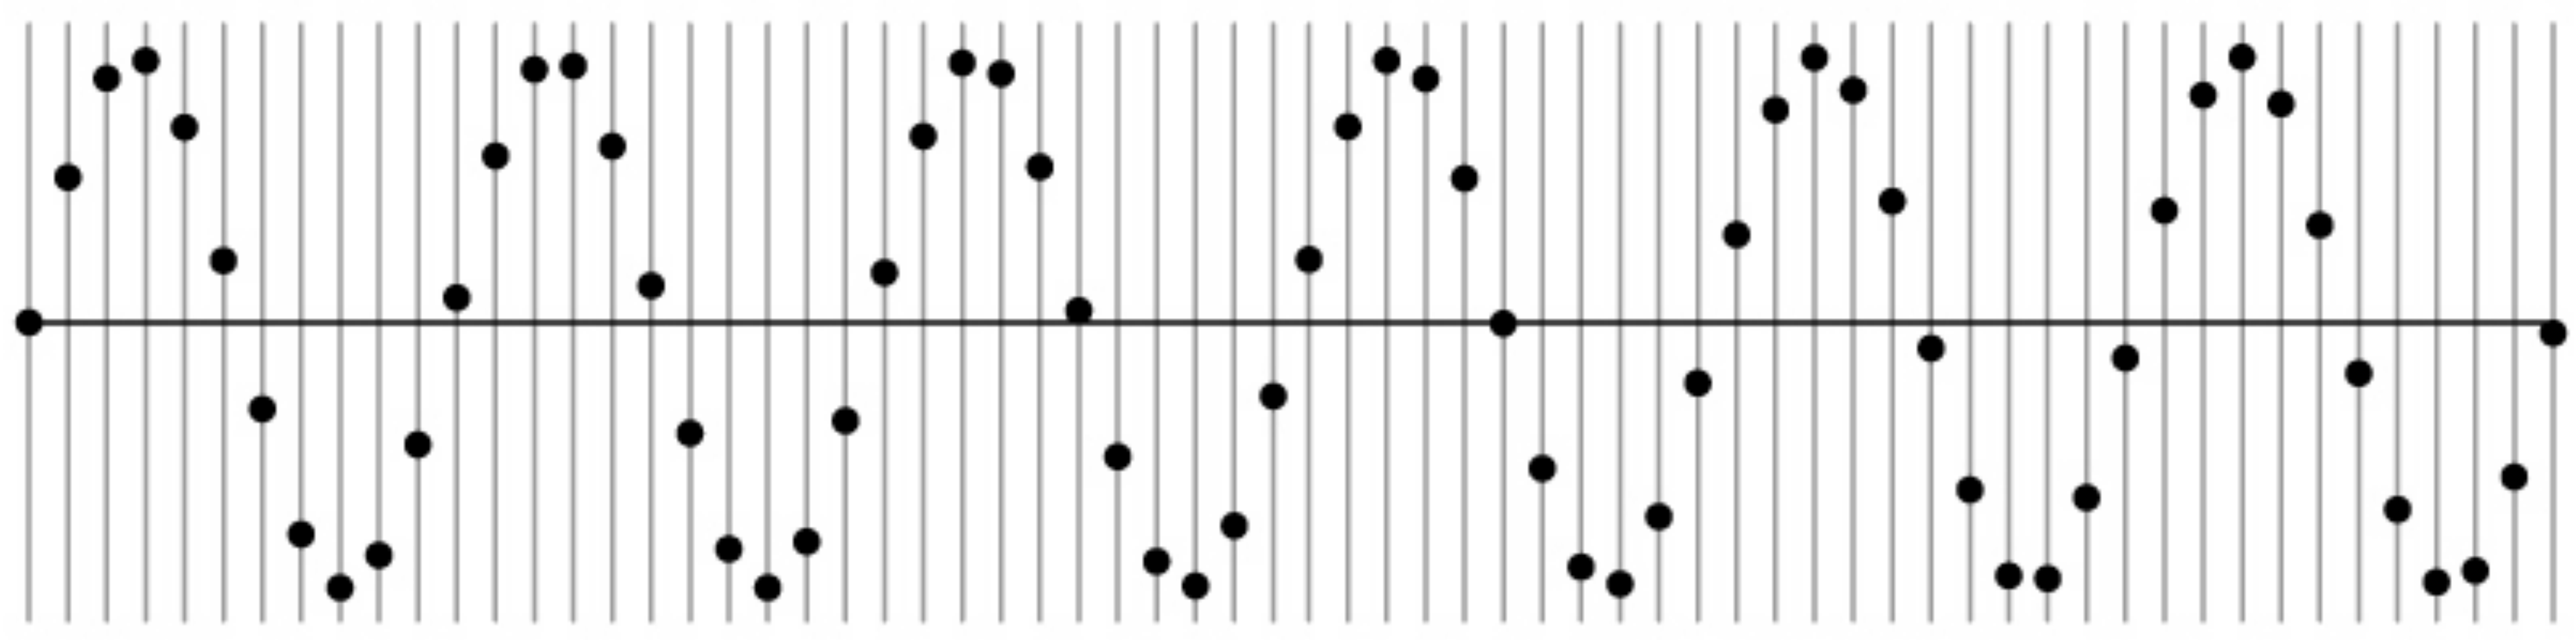

How many samples should I take?

Can I take as few samples as I want?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

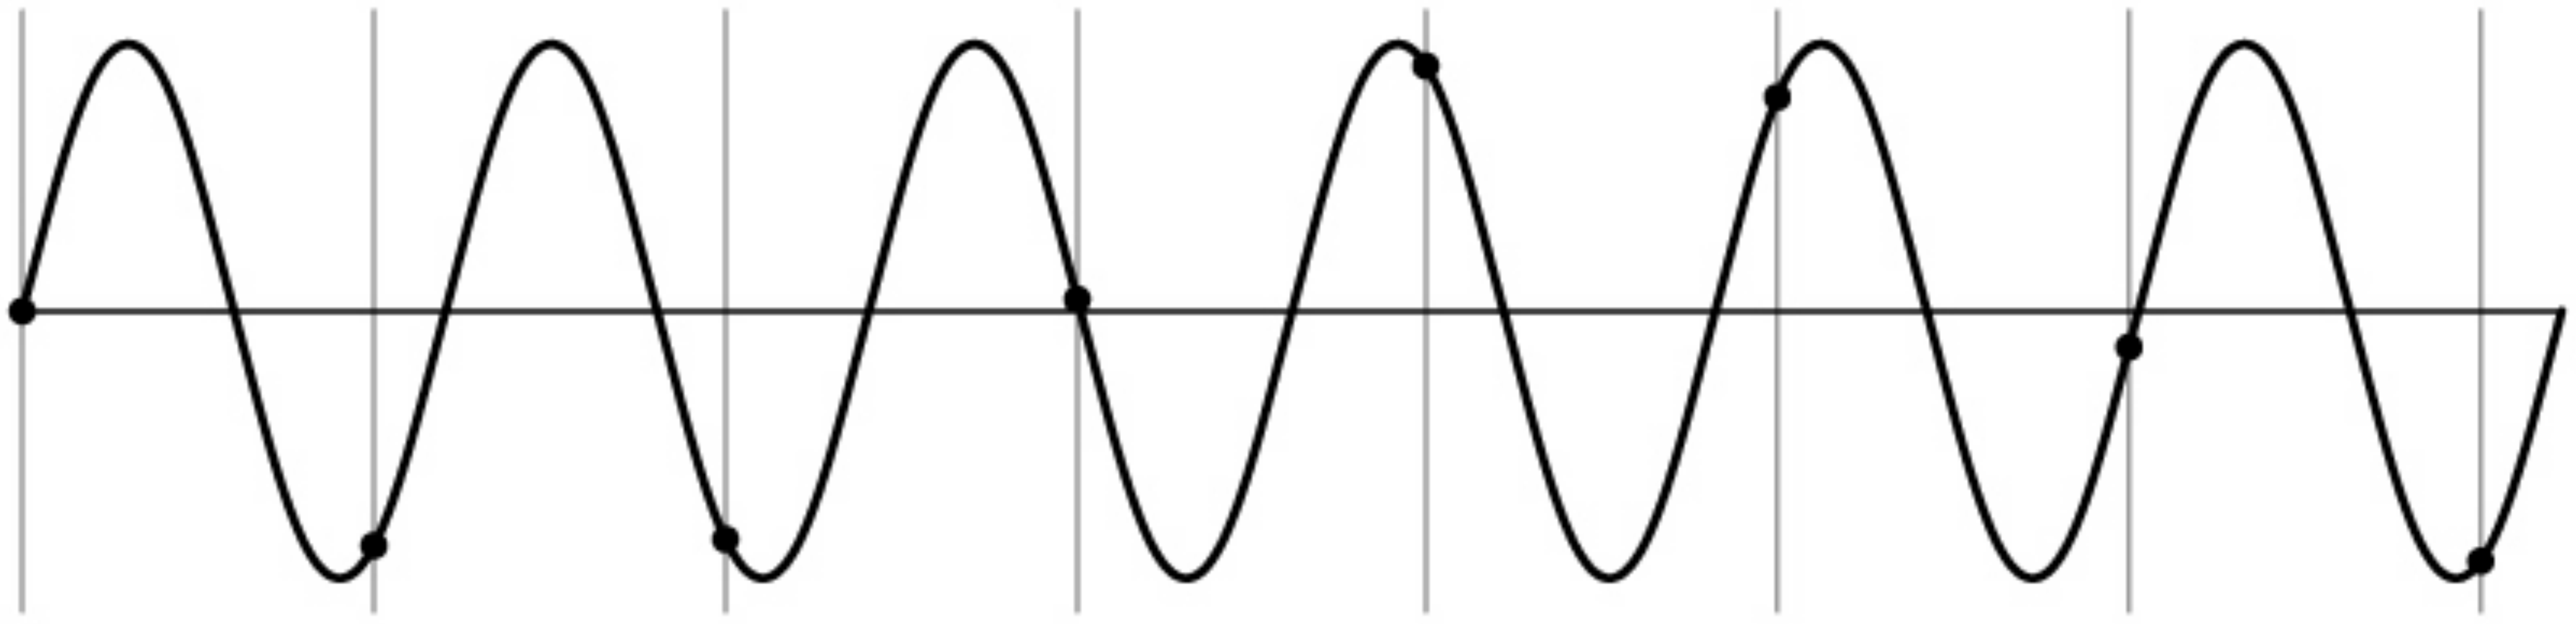

Signal can be confused with one at lower frequency

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

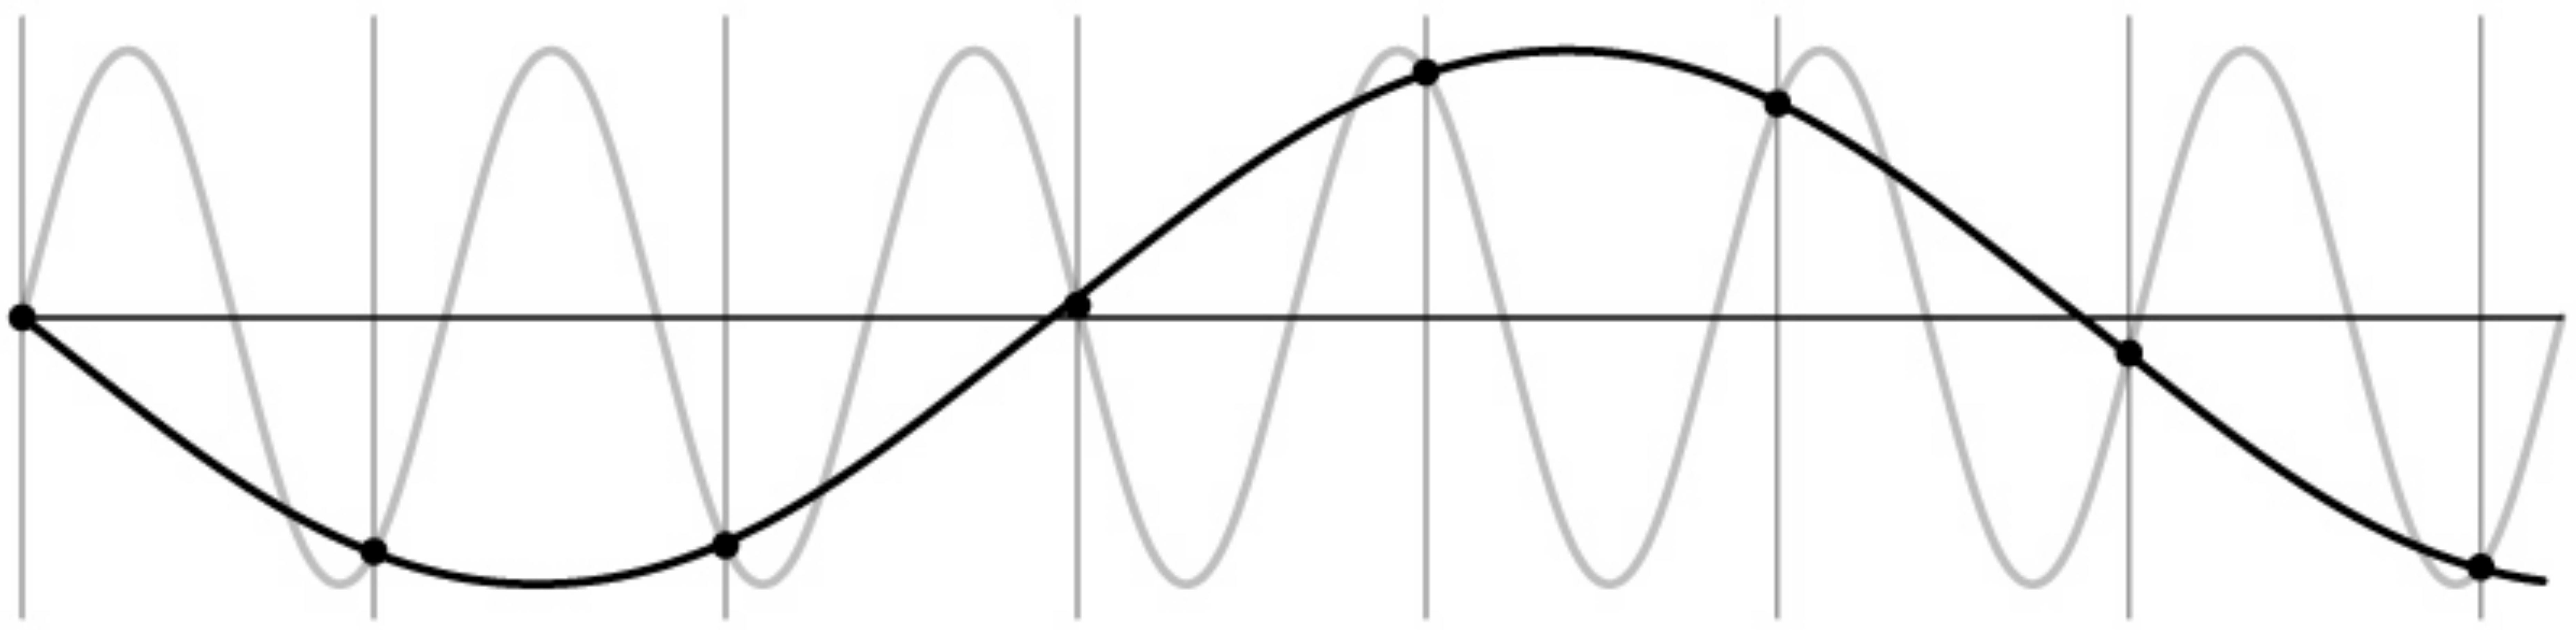

Signal can be confused with one at lower frequency

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Example: A Simple Sine Wave

How do we discretize the signal?

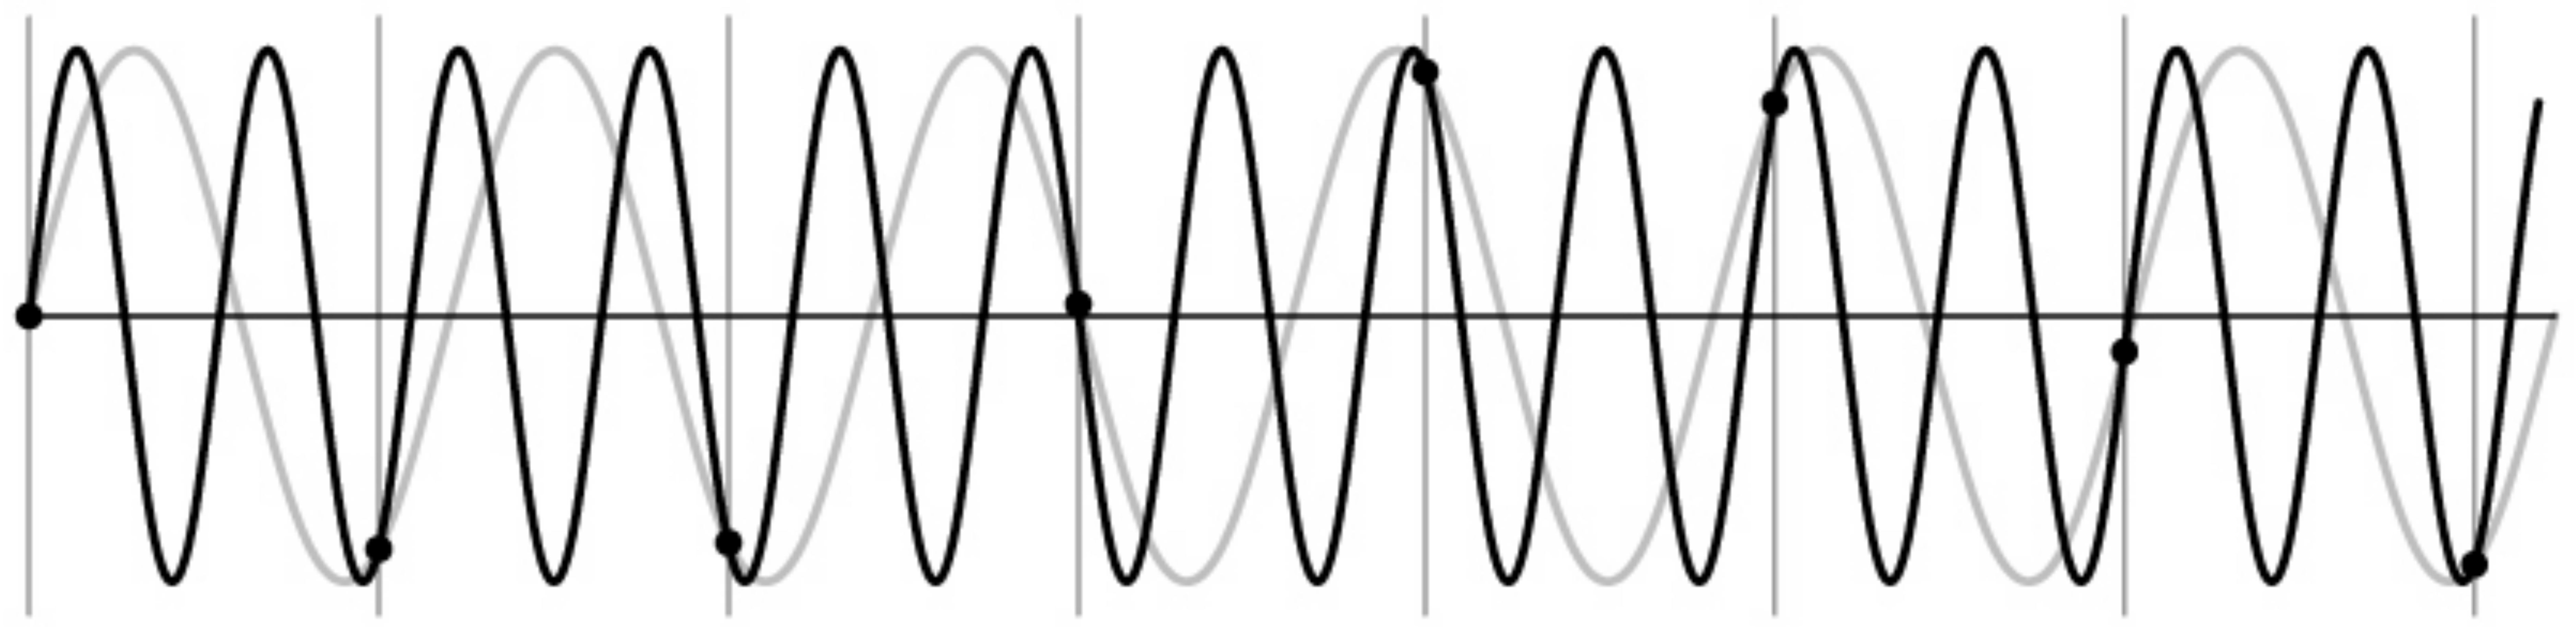

Signal can always be confused with one at higher frequency

Undersampling = **Aliasing**

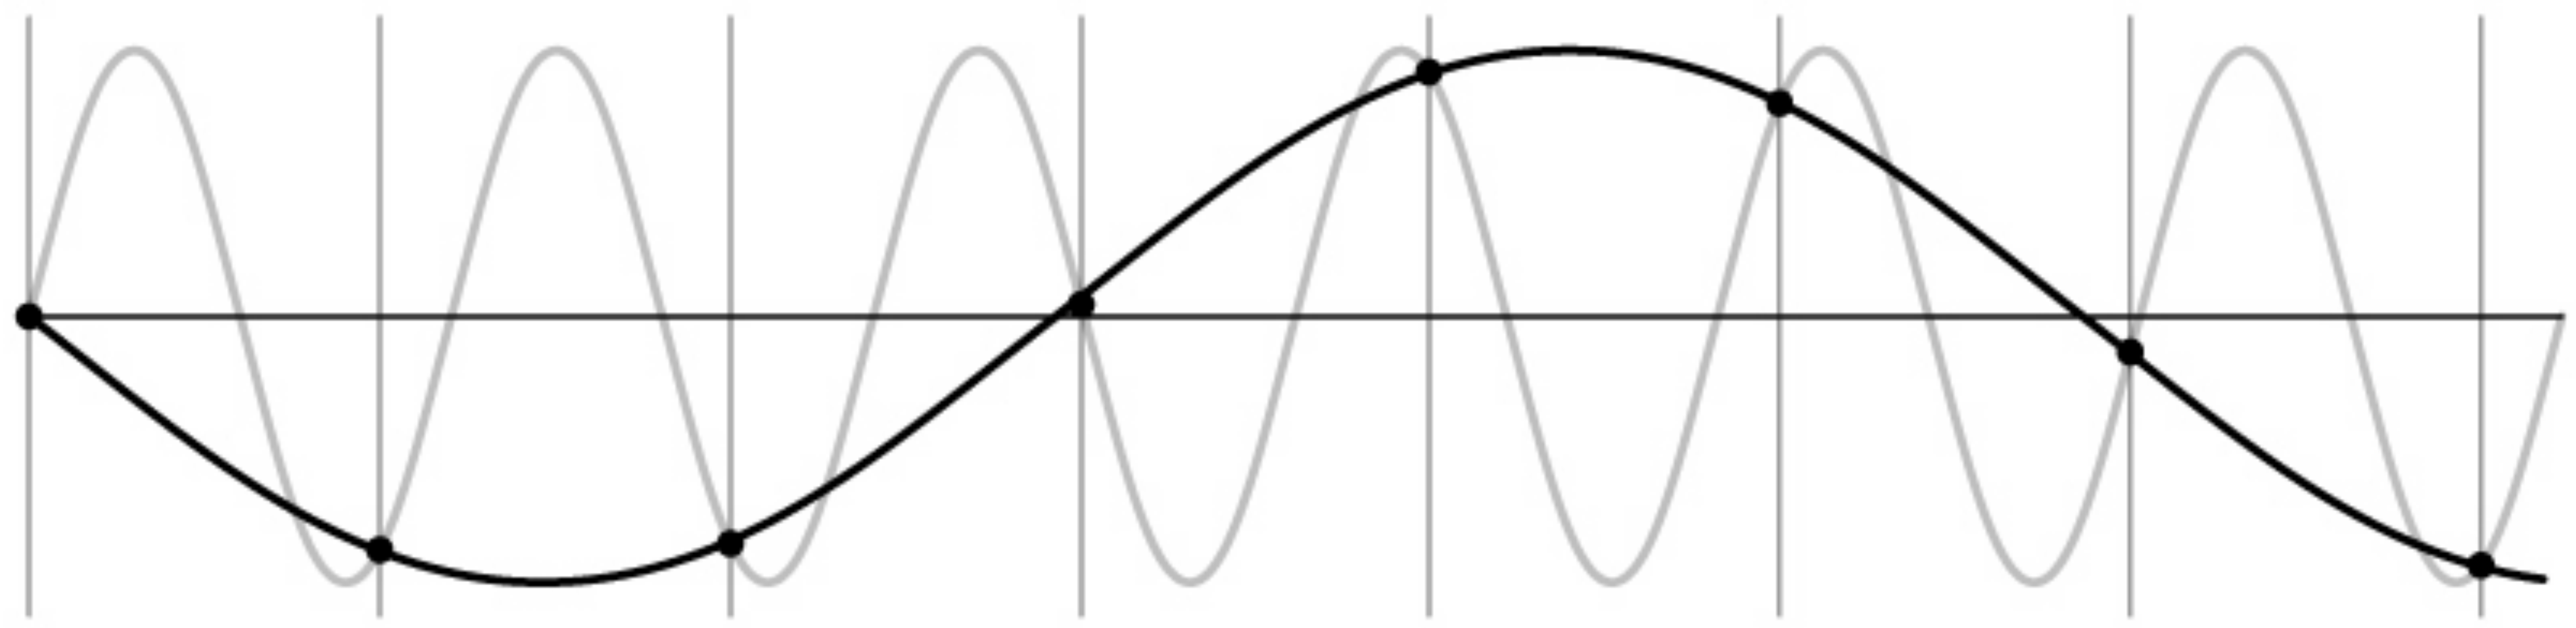

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Sampling Theory (informal)

The challenge to intuition is the fact that music (in the 1D case) and images (in the 2D case) can be represented as linear combinations of individual sine waves of differing frequencies and phases (remember discussion on FFTs)

A fundamental result (**Sampling Theorem**) is:

> For bandlimited signals, if you sample regularly at or above twice the maximum frequency (called the **Nyquist rate**), then you can reconstruct the original signal exactly

Sampling Theory (informal)

Question: For a bandlimited signal, what if you **oversample** (i.e., sample at greater than the Nyquist rate)

Sampling Theory (informal)

Question: For a bandlimited signal, what if you **oversample** (i.e., sample at greater than the Nyquist rate)

Answer: Nothing bad happens! Samples are redundant and there are wasted bits

Sampling Theory (informal)

Question: For a bandlimited signal, what if you **oversample** (i.e., sample at greater than the Nyquist rate)

Answer: Nothing bad happens! Samples are redundant and there are wasted bits

Question: For a bandlimited signal, what if you **undersample** (i.e., sample at less than the Nyquist rate)

Sampling Theory (informal)

Question: For a bandlimited signal, what if you **oversample** (i.e., sample at greater than the Nyquist rate)

Answer: Nothing bad happens! Samples are redundant and there are wasted bits

Question: For a bandlimited signal, what if you **undersample** (i.e., sample at less than the Nyquist rate)

Answer: Two bad things happen! Things are missing (i.e., things that should be there aren’t). There are artifacts (i.e., things that shouldn’t be there are)

Sampling Theory (informal)

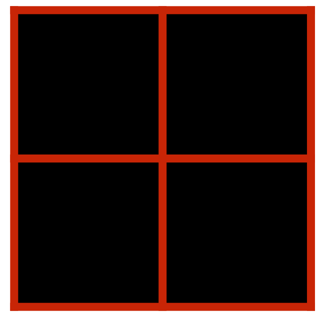

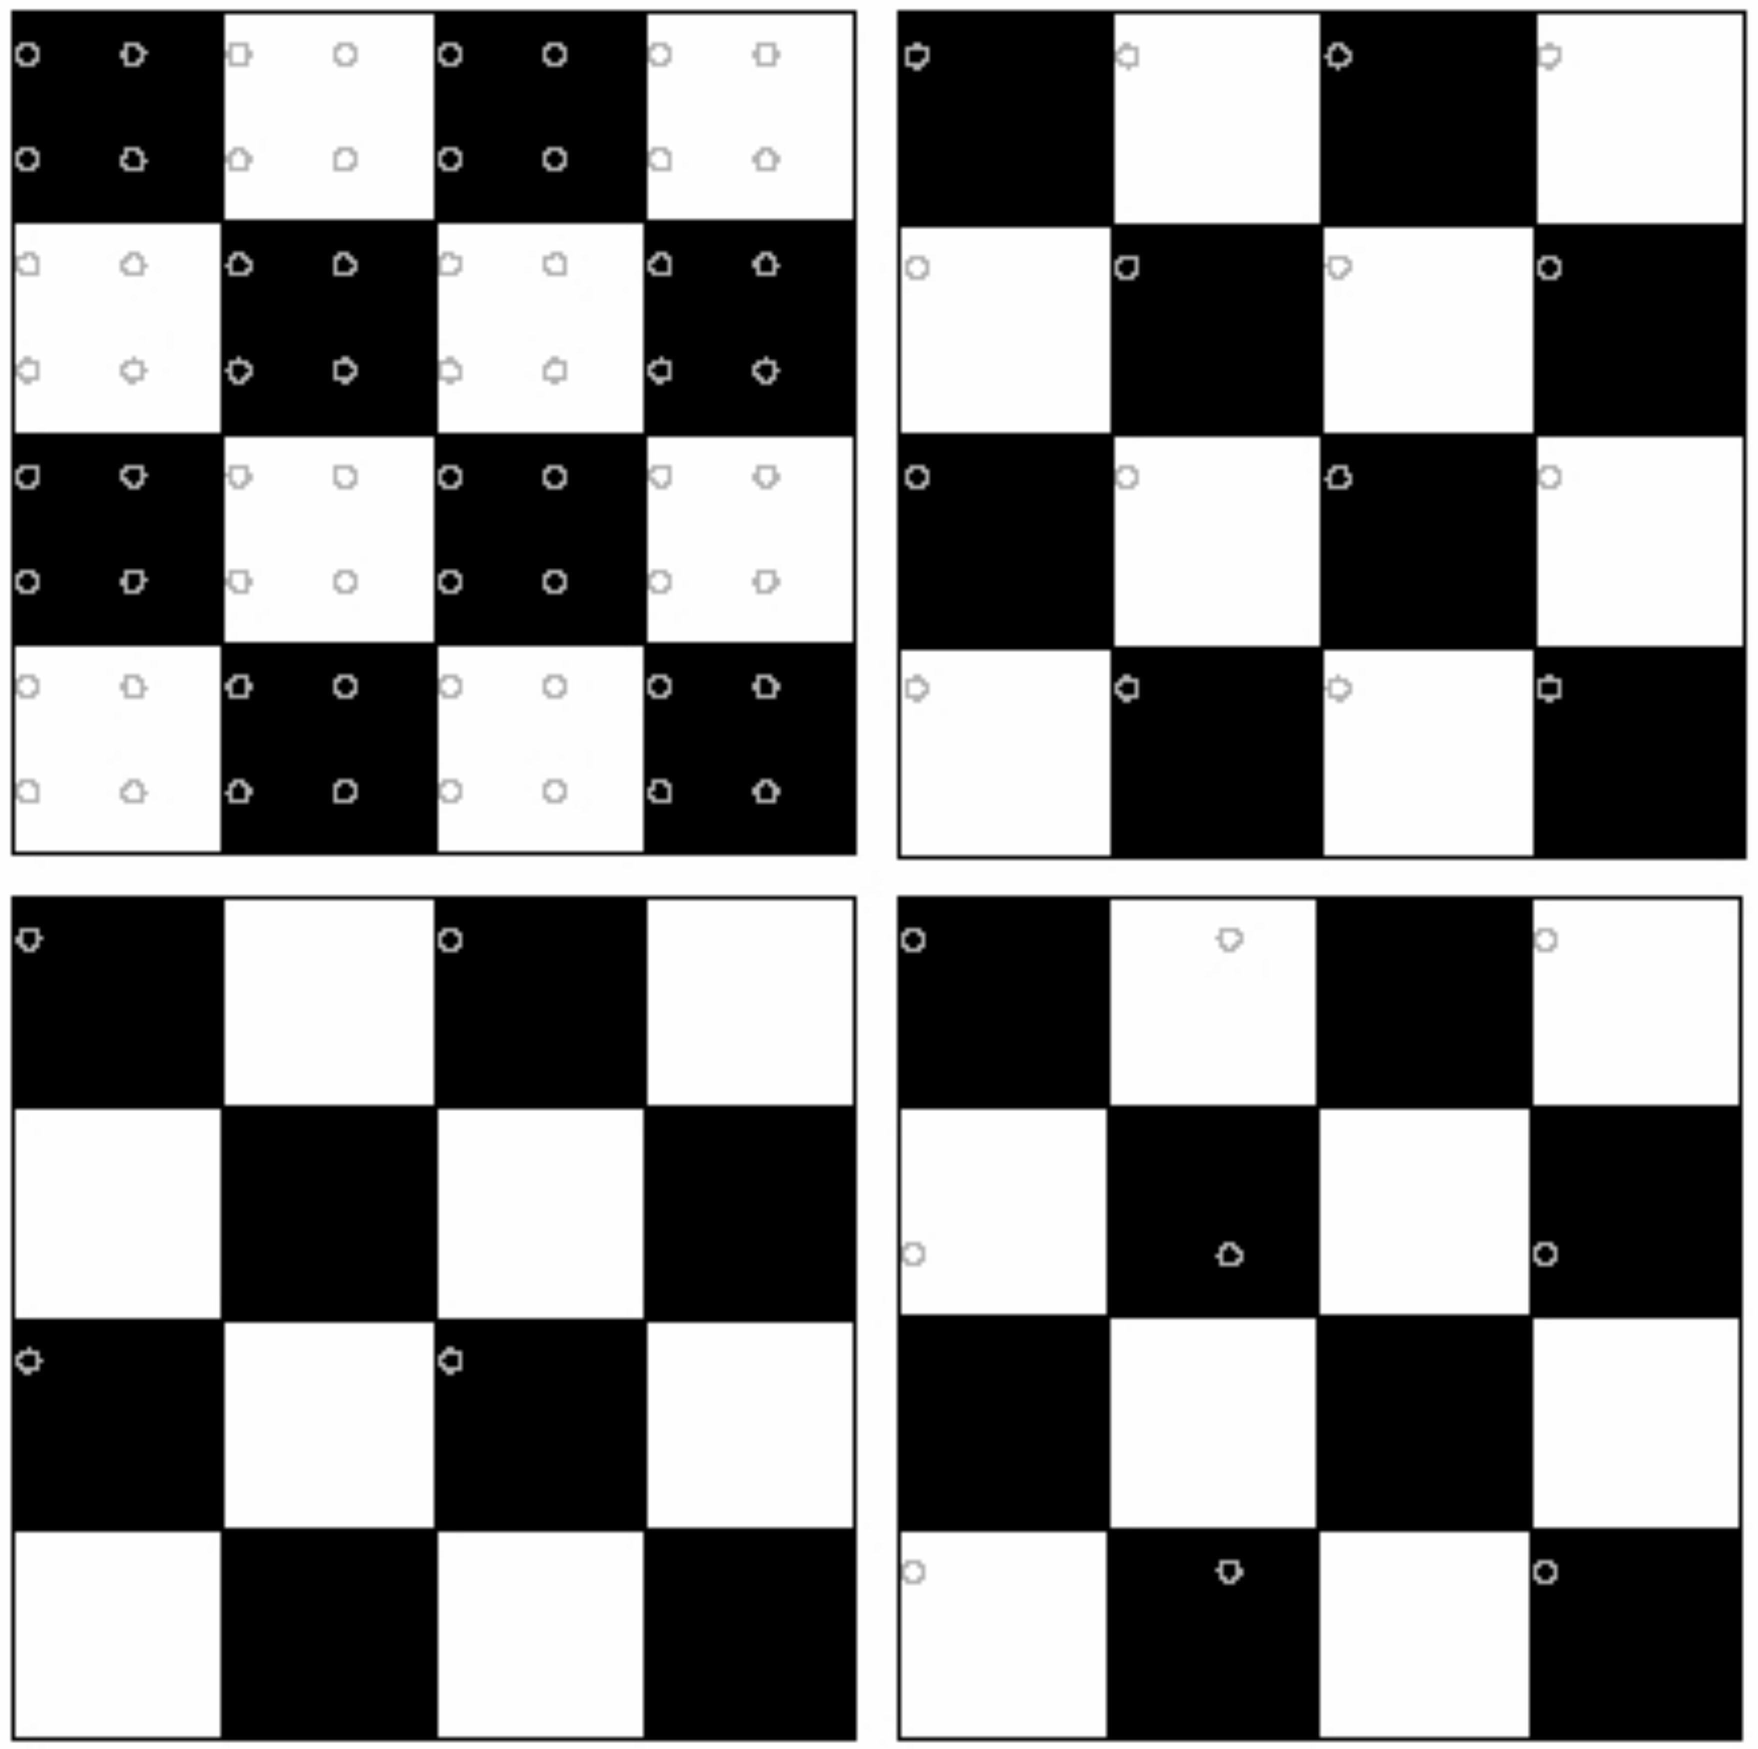

Forsyth & Ponce (2nd ed.) Figure 4.7

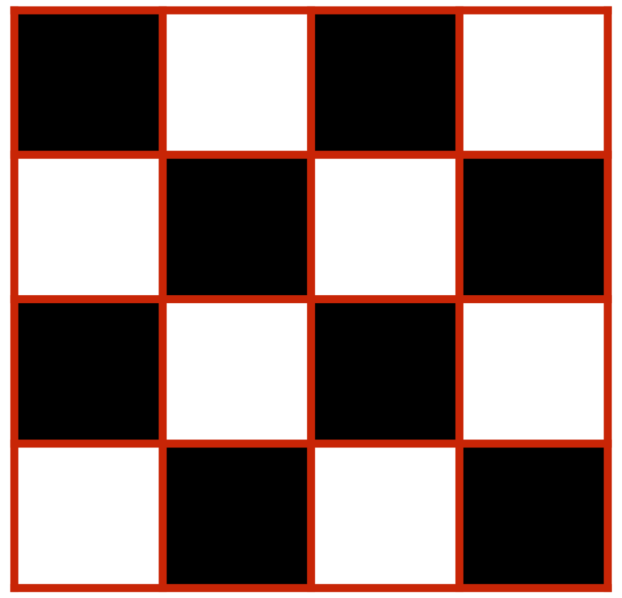

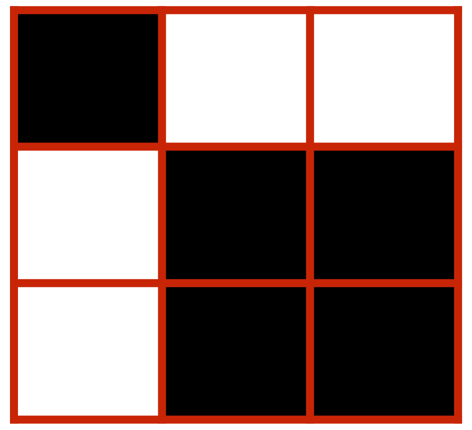

Sampling Theory (informal)

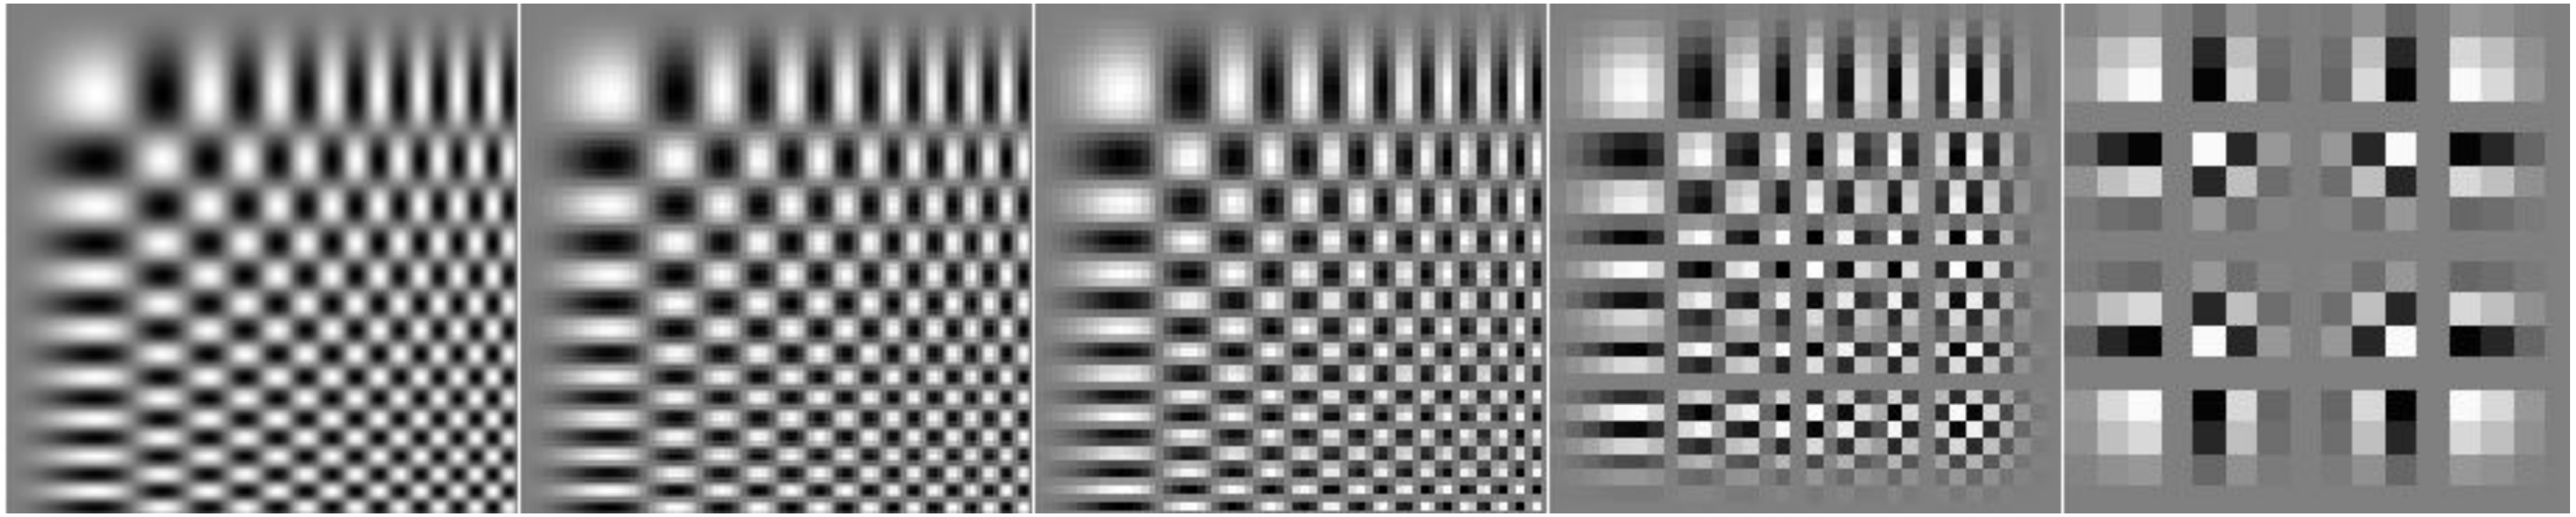

Forsyth & Ponce (2nd ed.) Figure 4.12

Reducing Aliasing Artifacts

1. **Oversampling** — sample more than you think you need and average (i.e., area sampling)

Aliasing

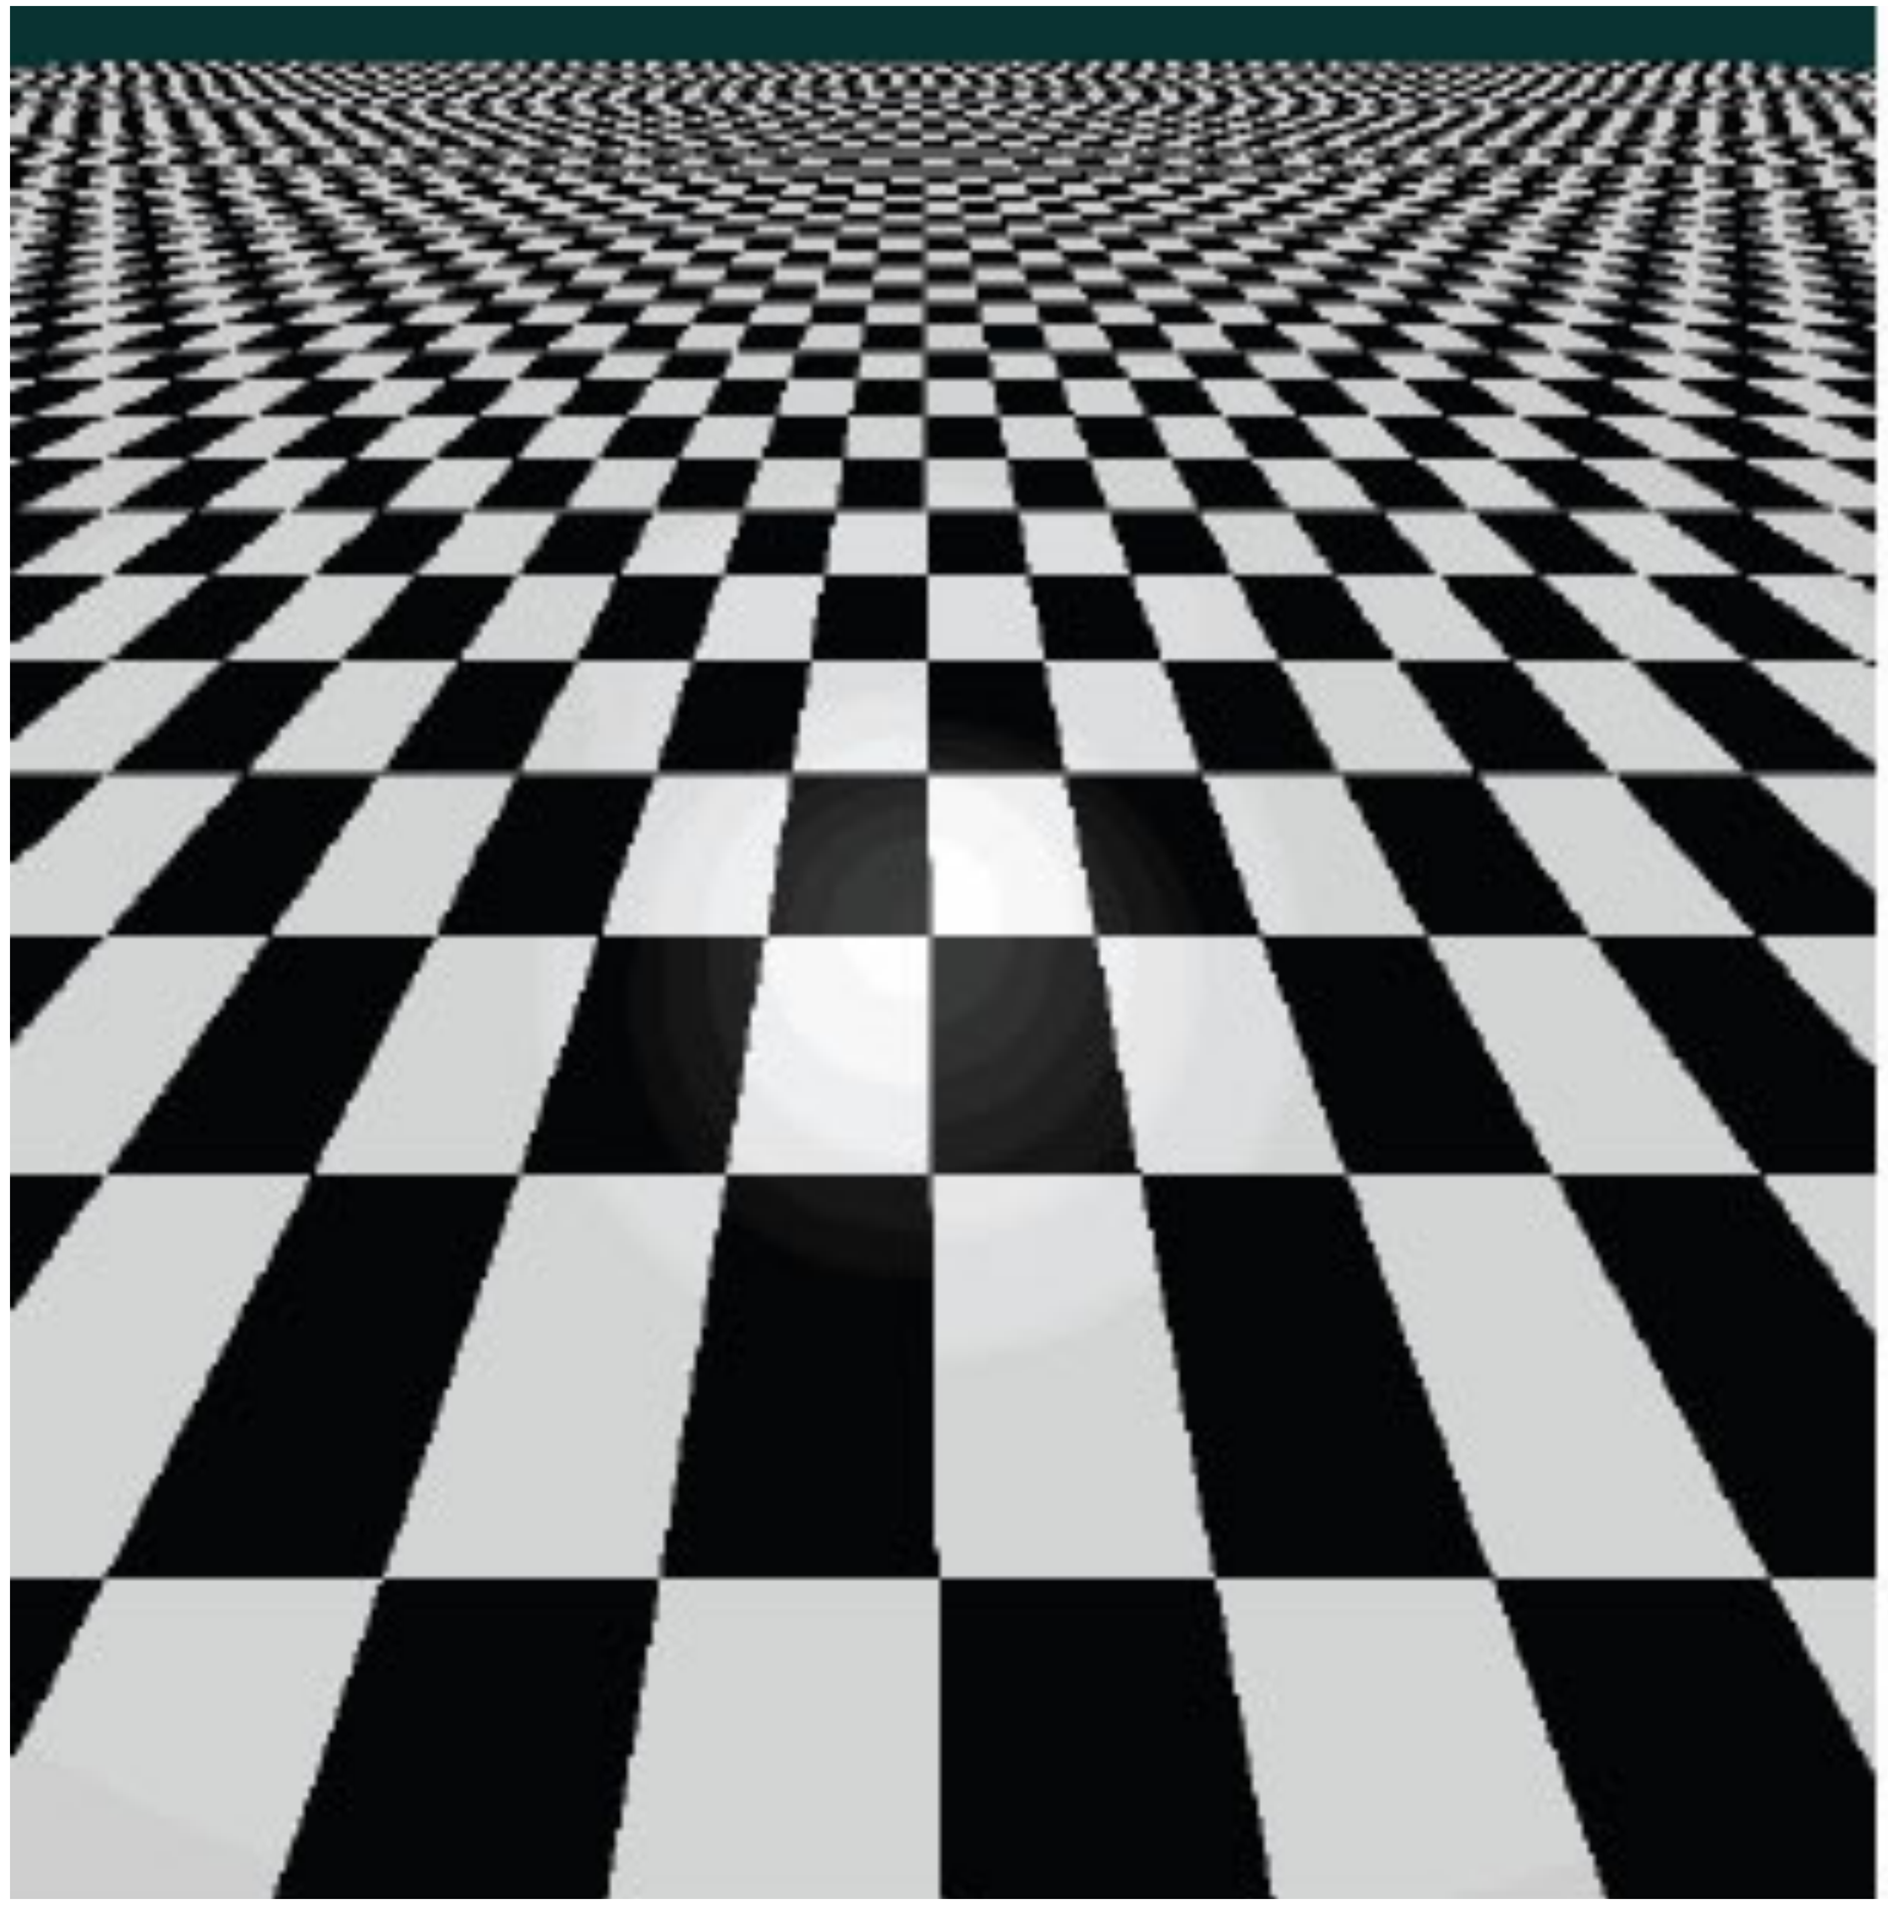

aliasing artifacts

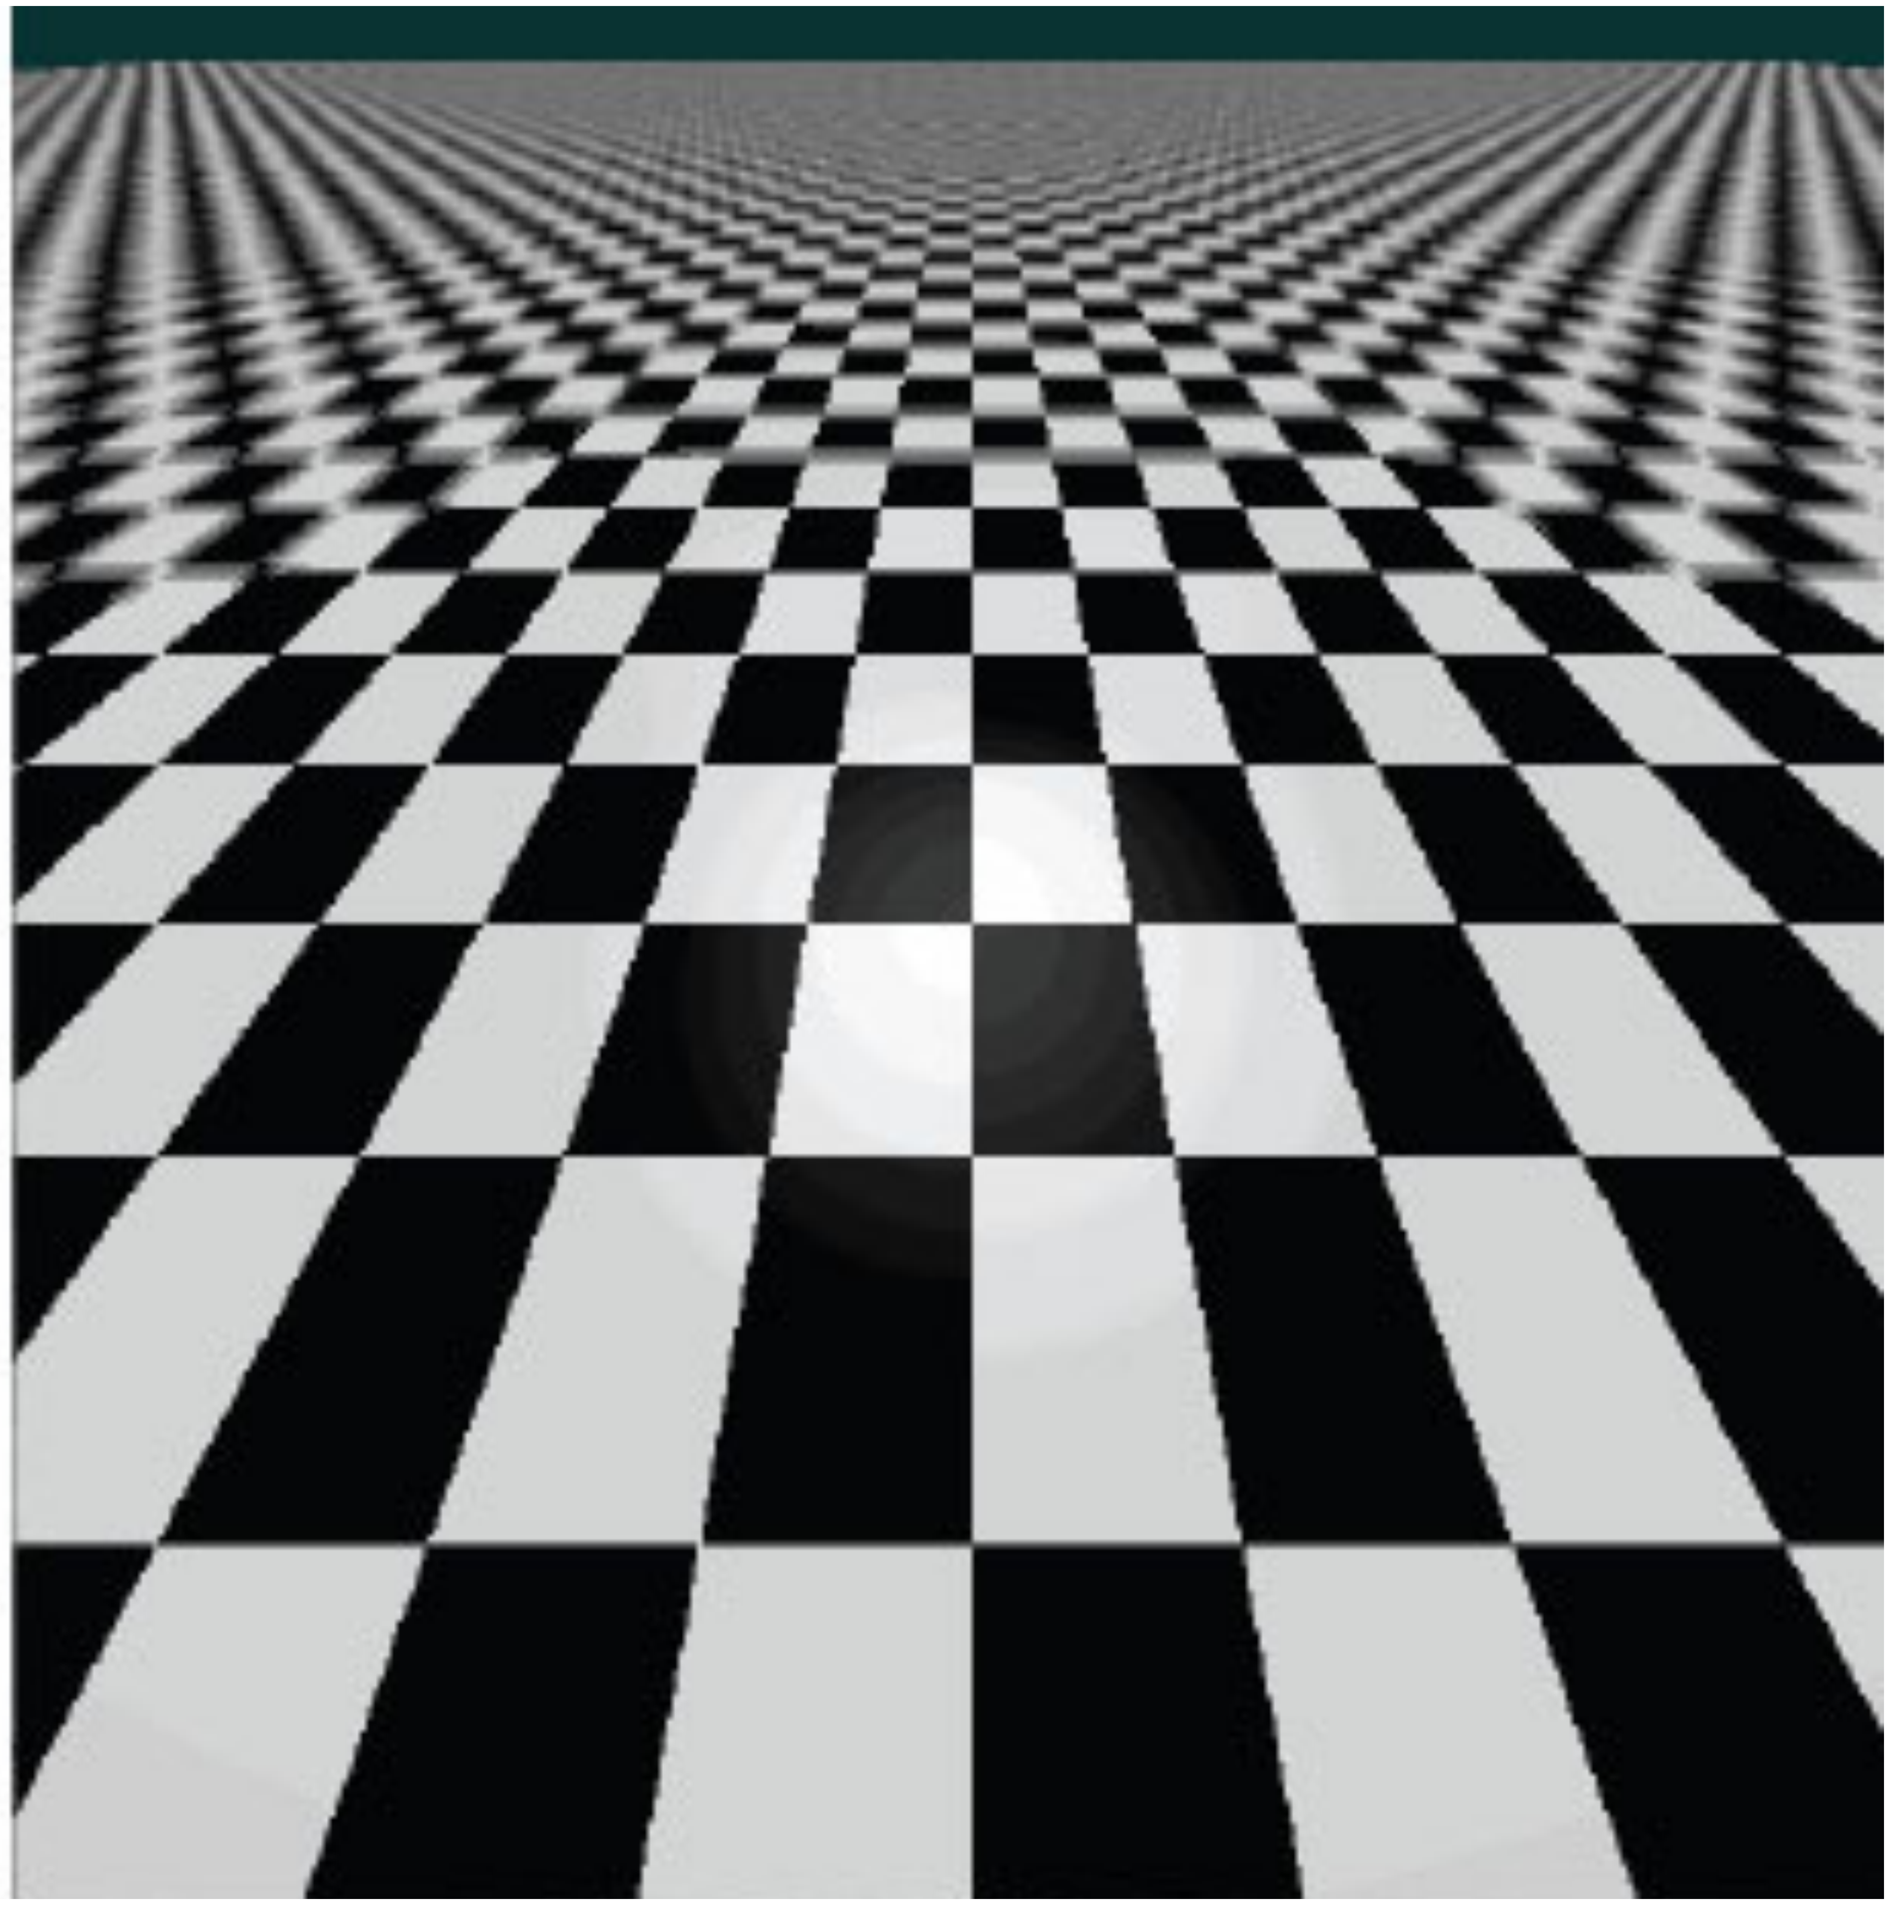

anti-aliasing by oversampling

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Reducing Aliasing Artifacts

1. **Oversampling** — sample more than you think you need and average (i.e., area sampling)

2. **Smoothing** before sampling. Why?

Aliasing in Photographs

This is also known as “moire”

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Imagine a spoked wheel moving to the right (rotating clockwise). Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):

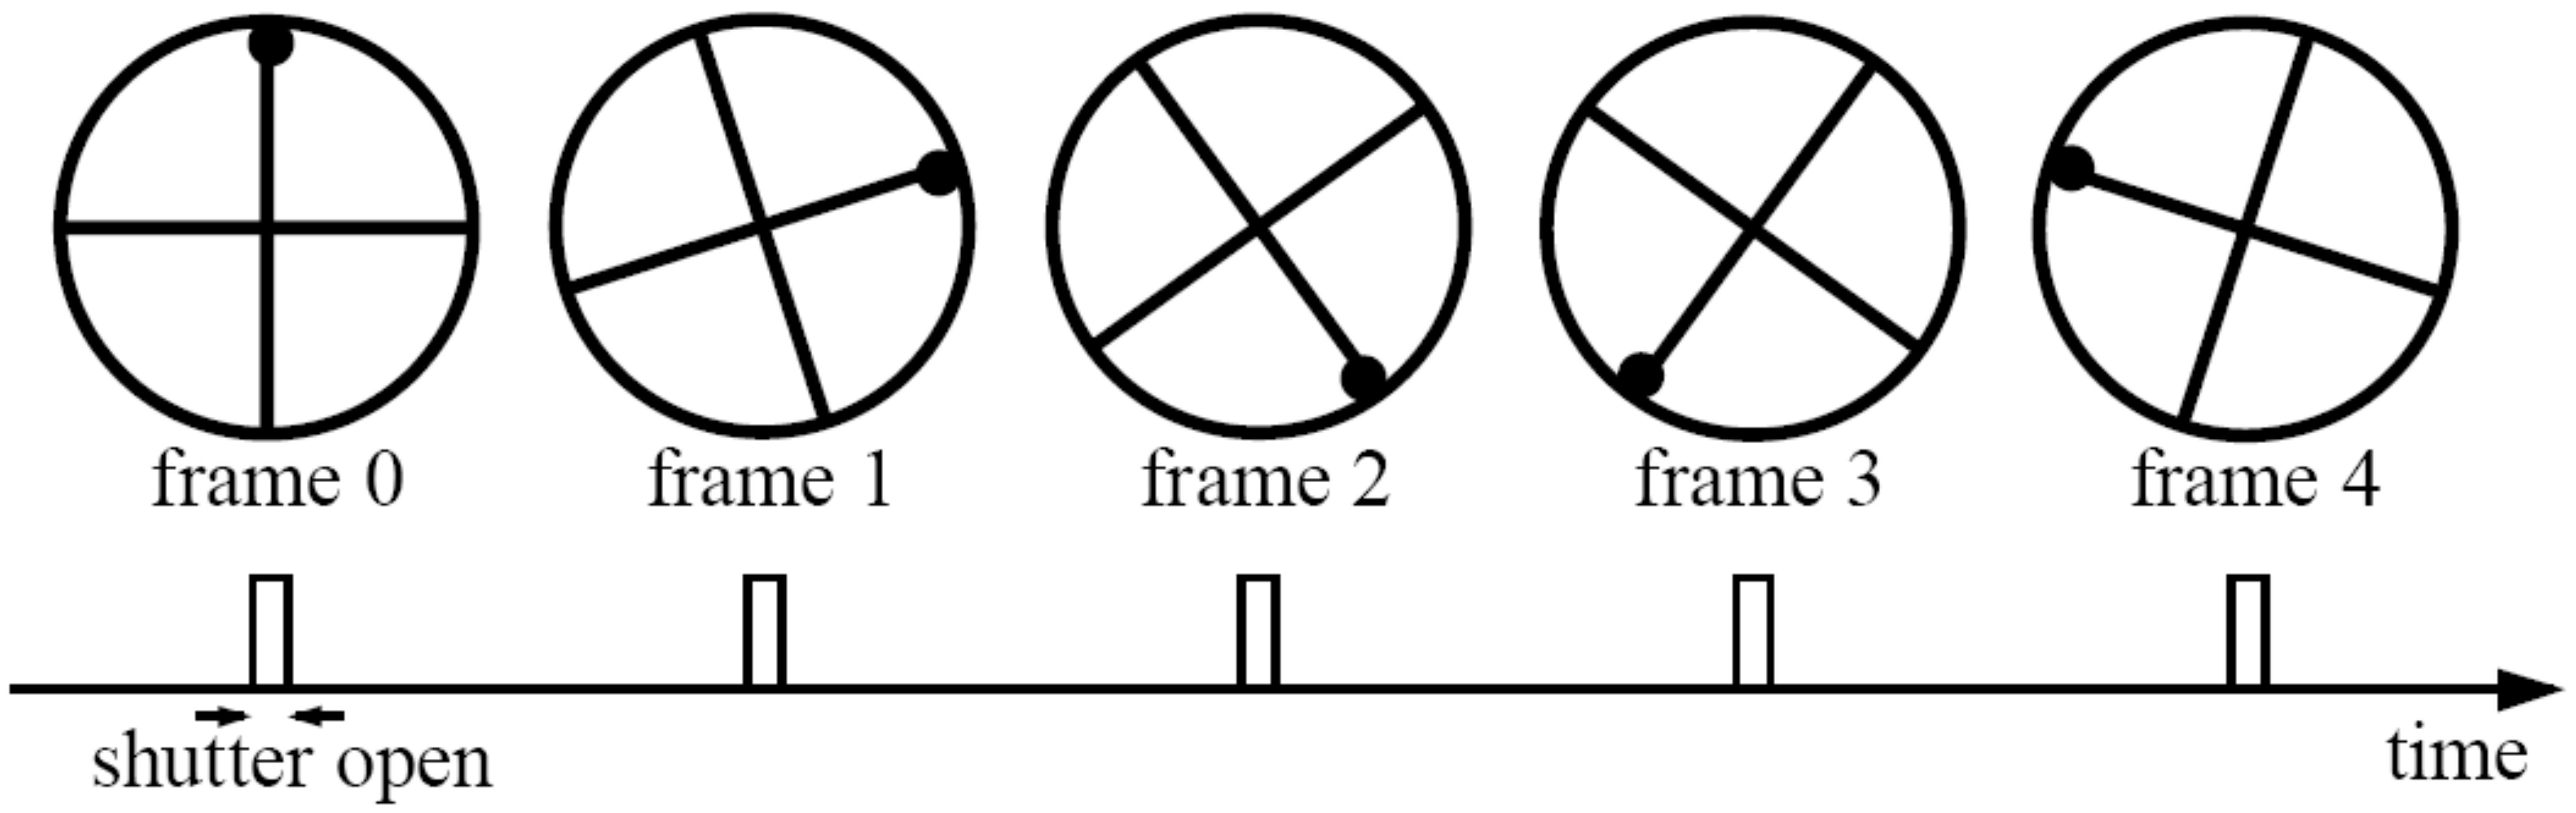

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Temporal Aliasing

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Sampling Theory (informal)

Sometimes **undersampling** is unavoidable, and there is a trade-off between "things missing" and "artifacts."

— **Medical imaging**: usually try to maximize information content, tolerate some artifacts

— **Computer graphics**: usually try to minimize artifacts, tolerate some information missing