Dr. George Karraz, Ph. D.

Computer Vision

Introduction and Overview

Dr. George Karraz, Ph.D.

This presentation is an overview of some of the ideas and techniques to be covered during the course.

Topics

1. Image formation
2. Point processing and equalization
3. Color correction
4. The Fourier transform
5. Convolution
6. Image sampling, warping, and stitching
7. Frequency Domain (FD) Filtering
8. Spatial filtering
9. Noise reduction
10. Mathematical morphology
11. High dynamic range imaging

1. Image Formation

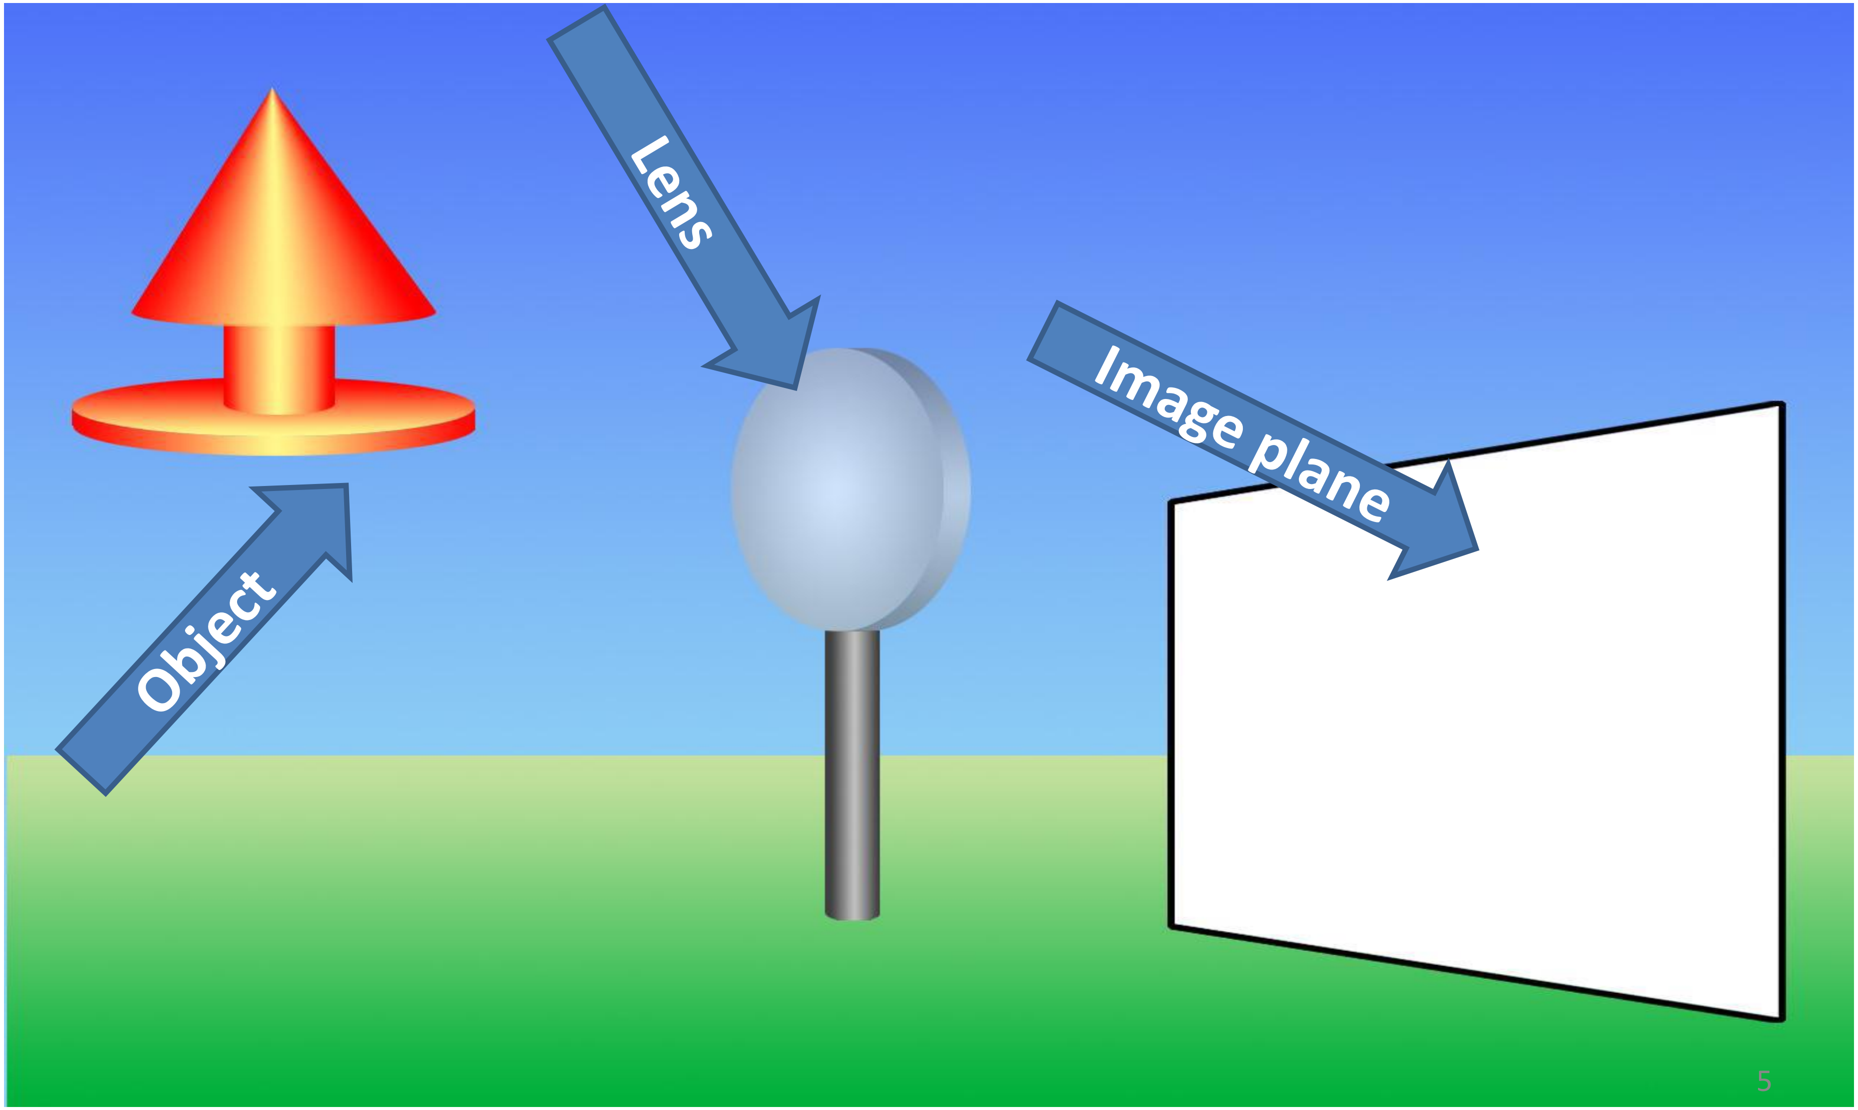

1. Image Formation

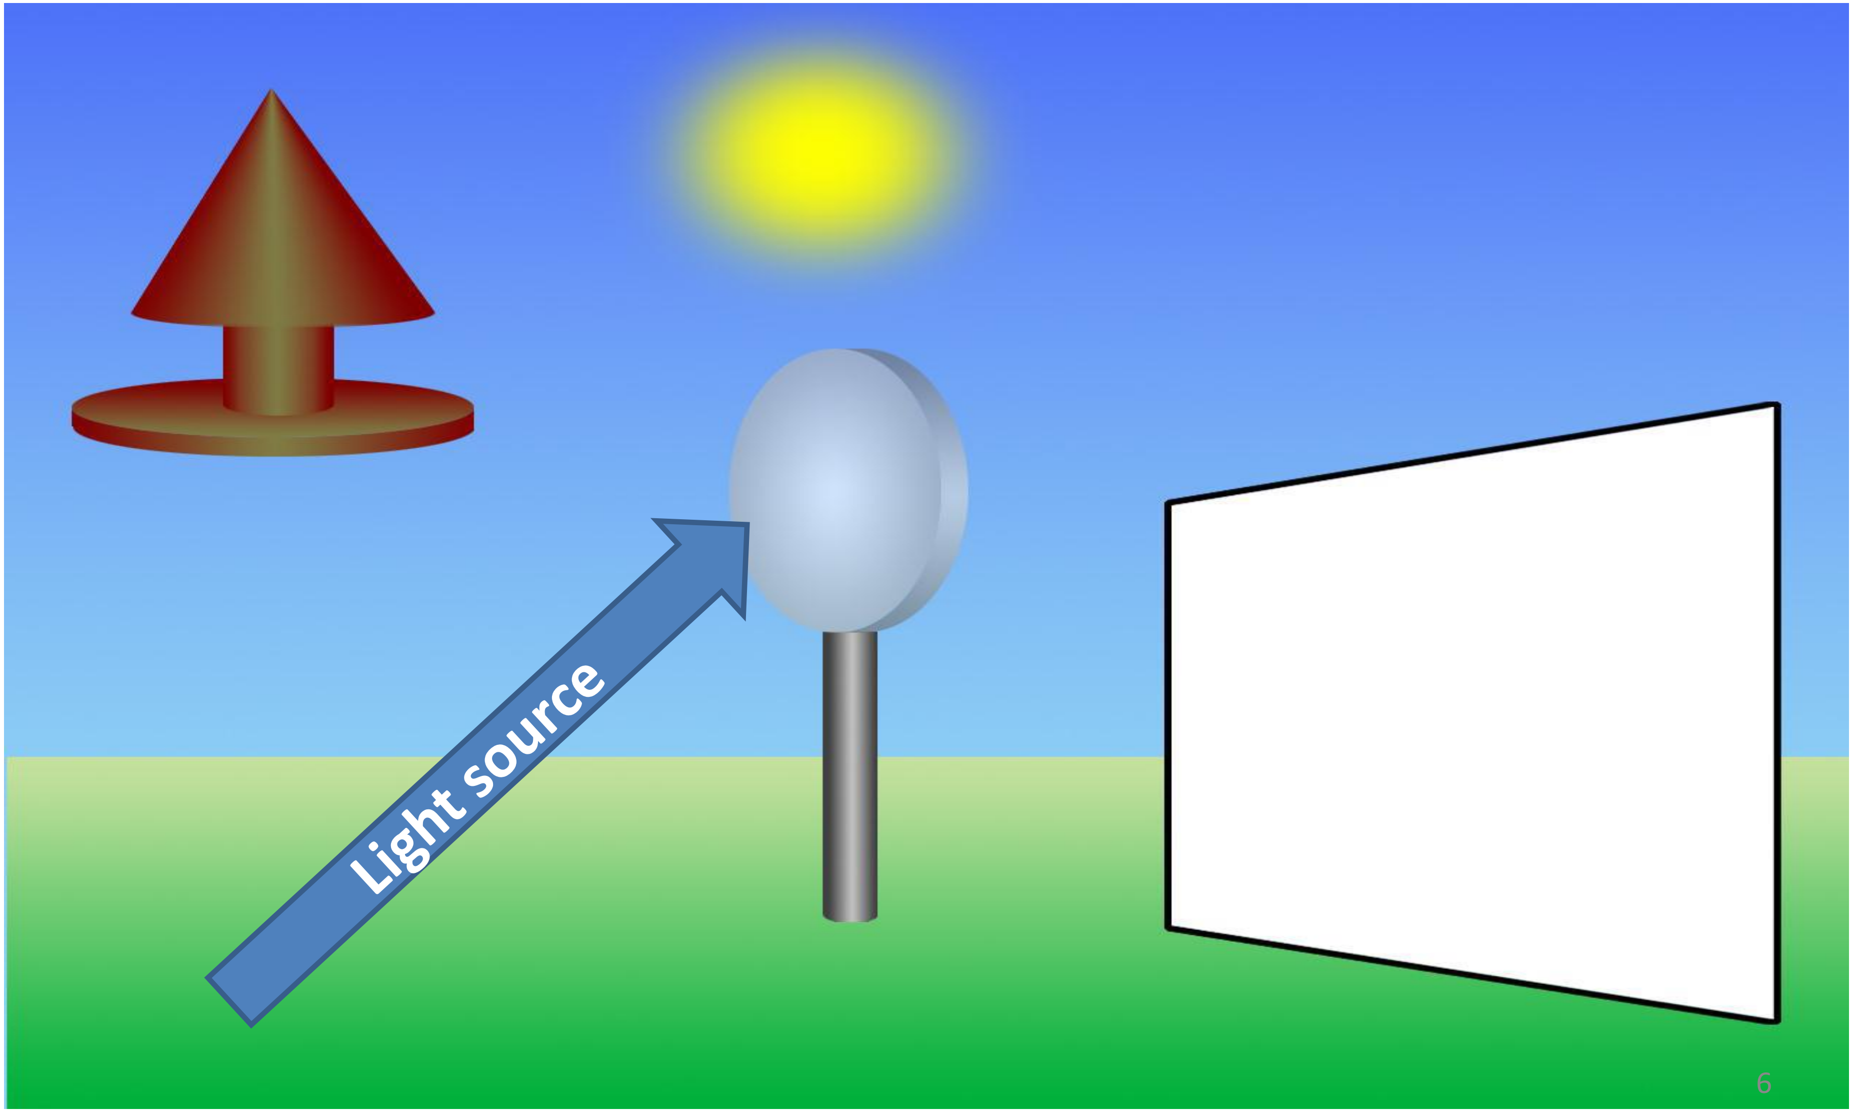

1. Image Formation

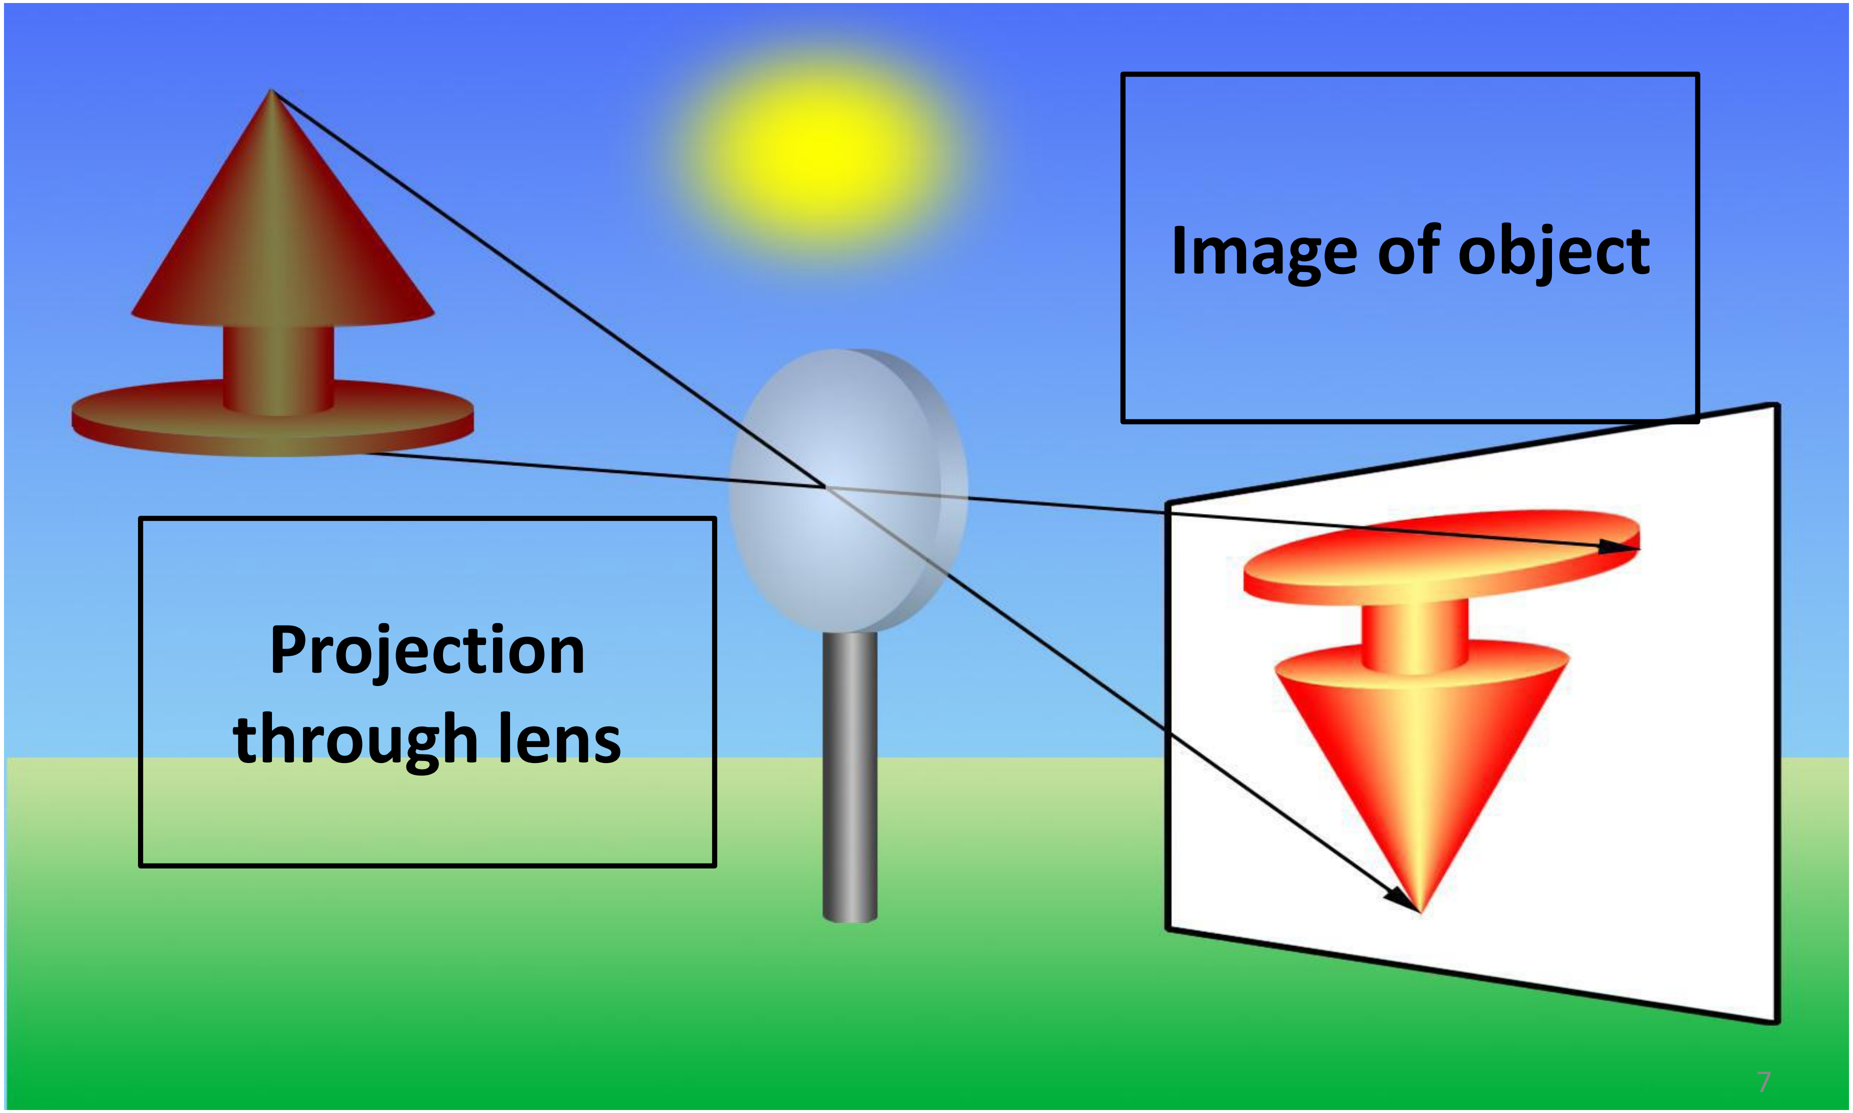

1. Image Formation

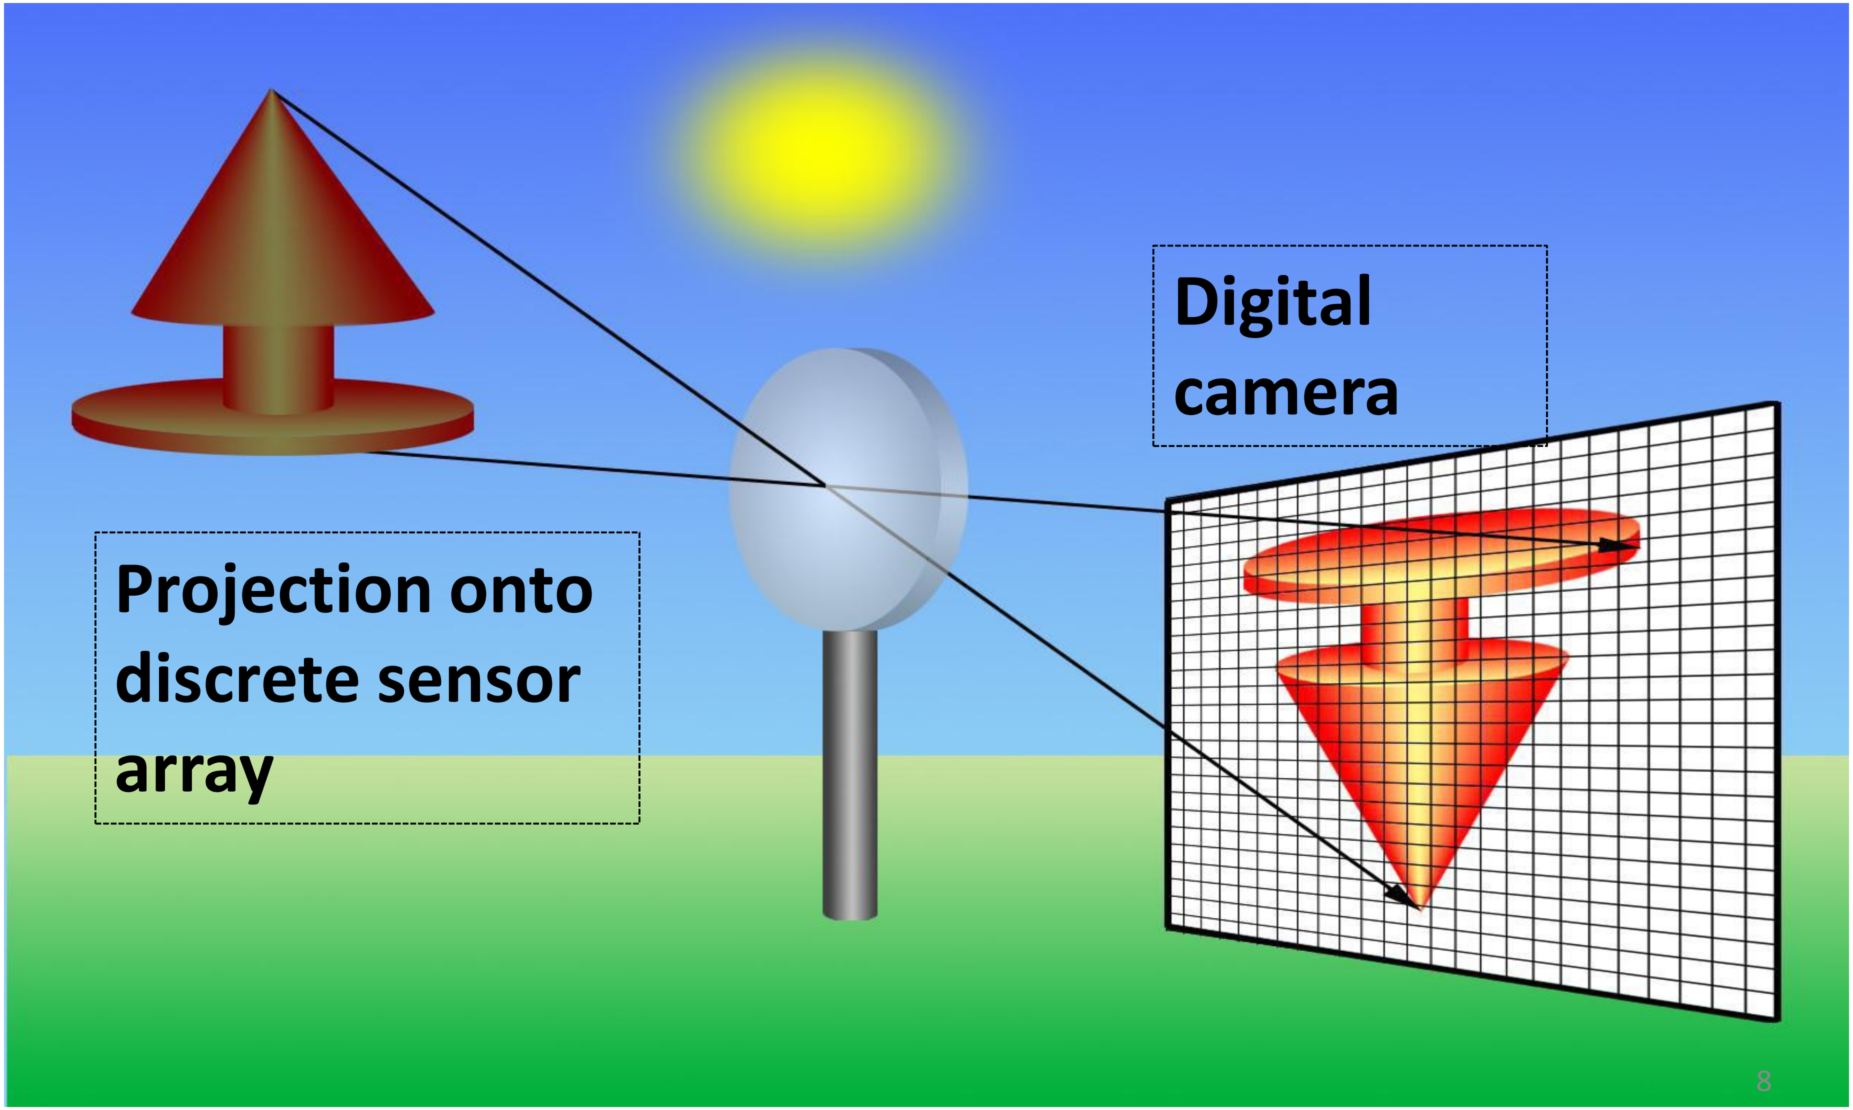

1. Image Formation

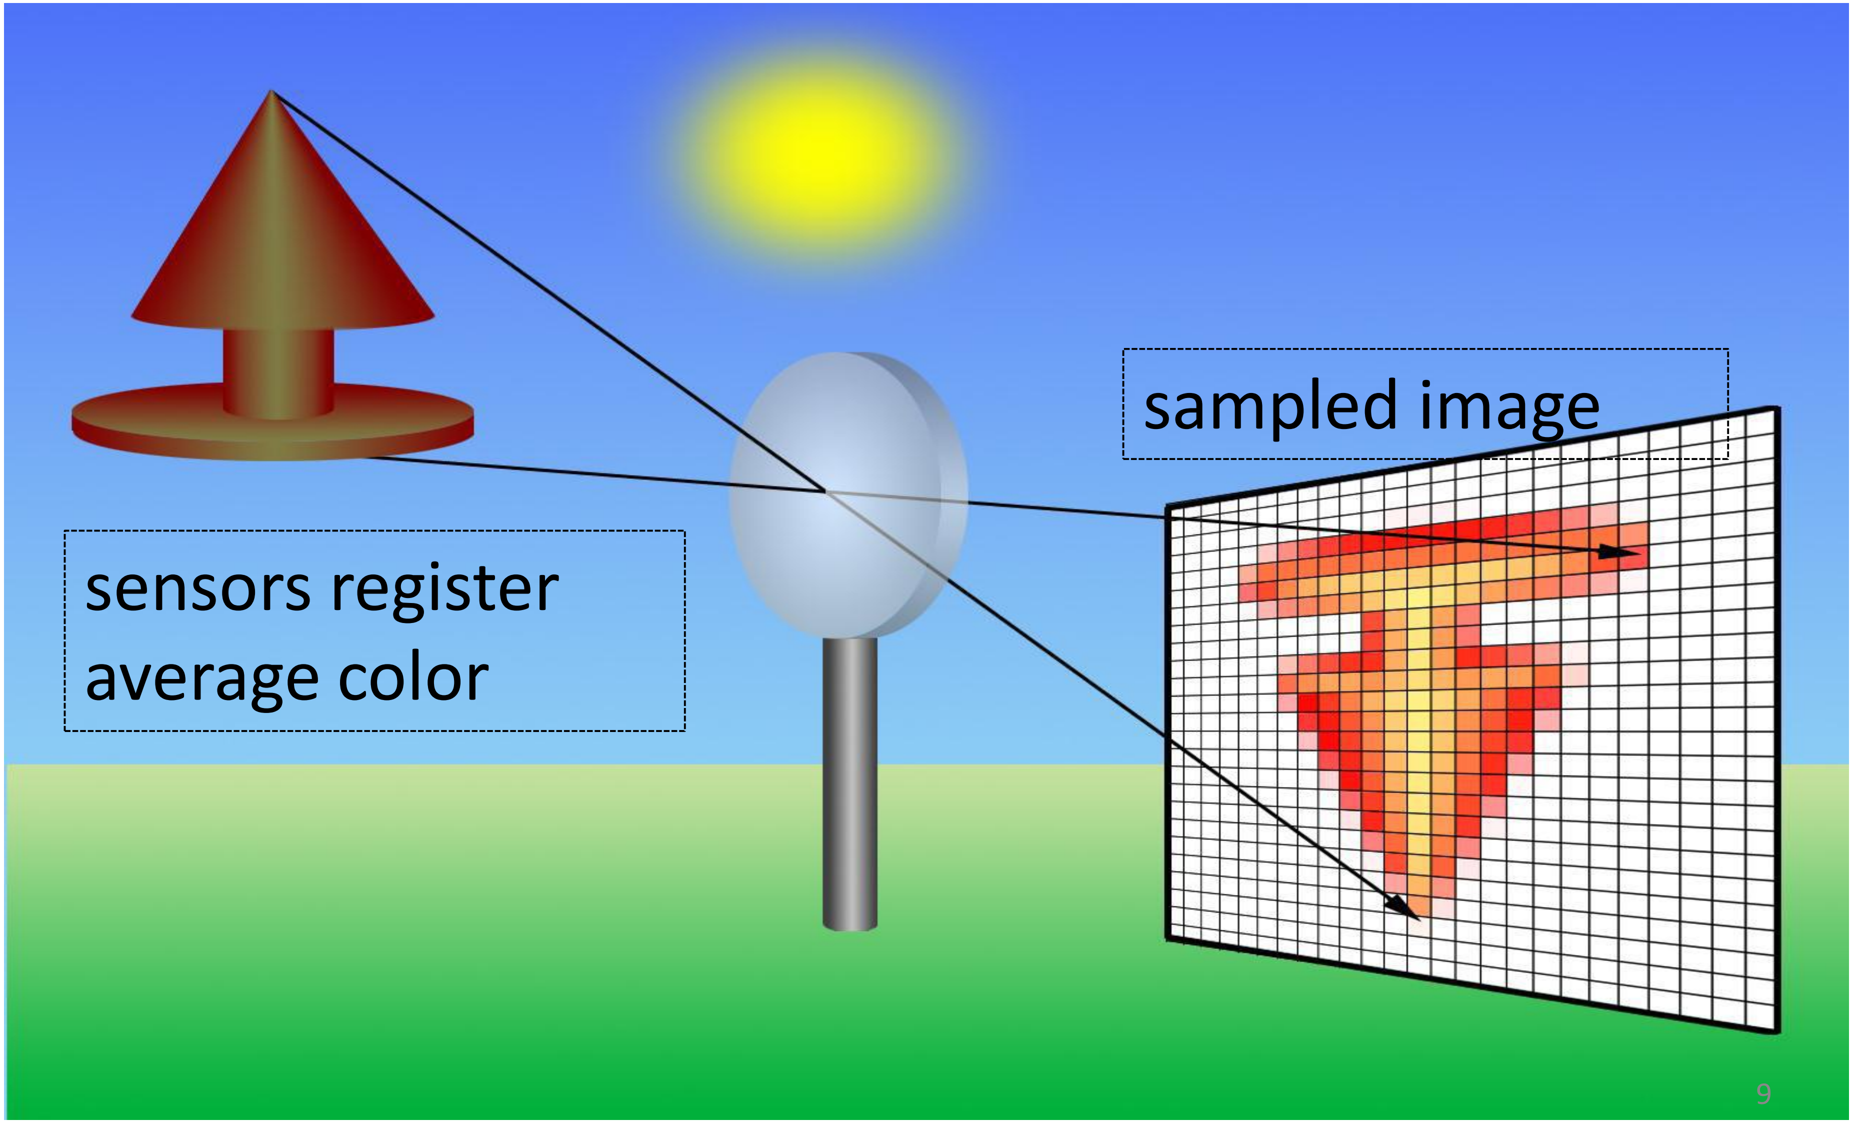

1. Image Formation

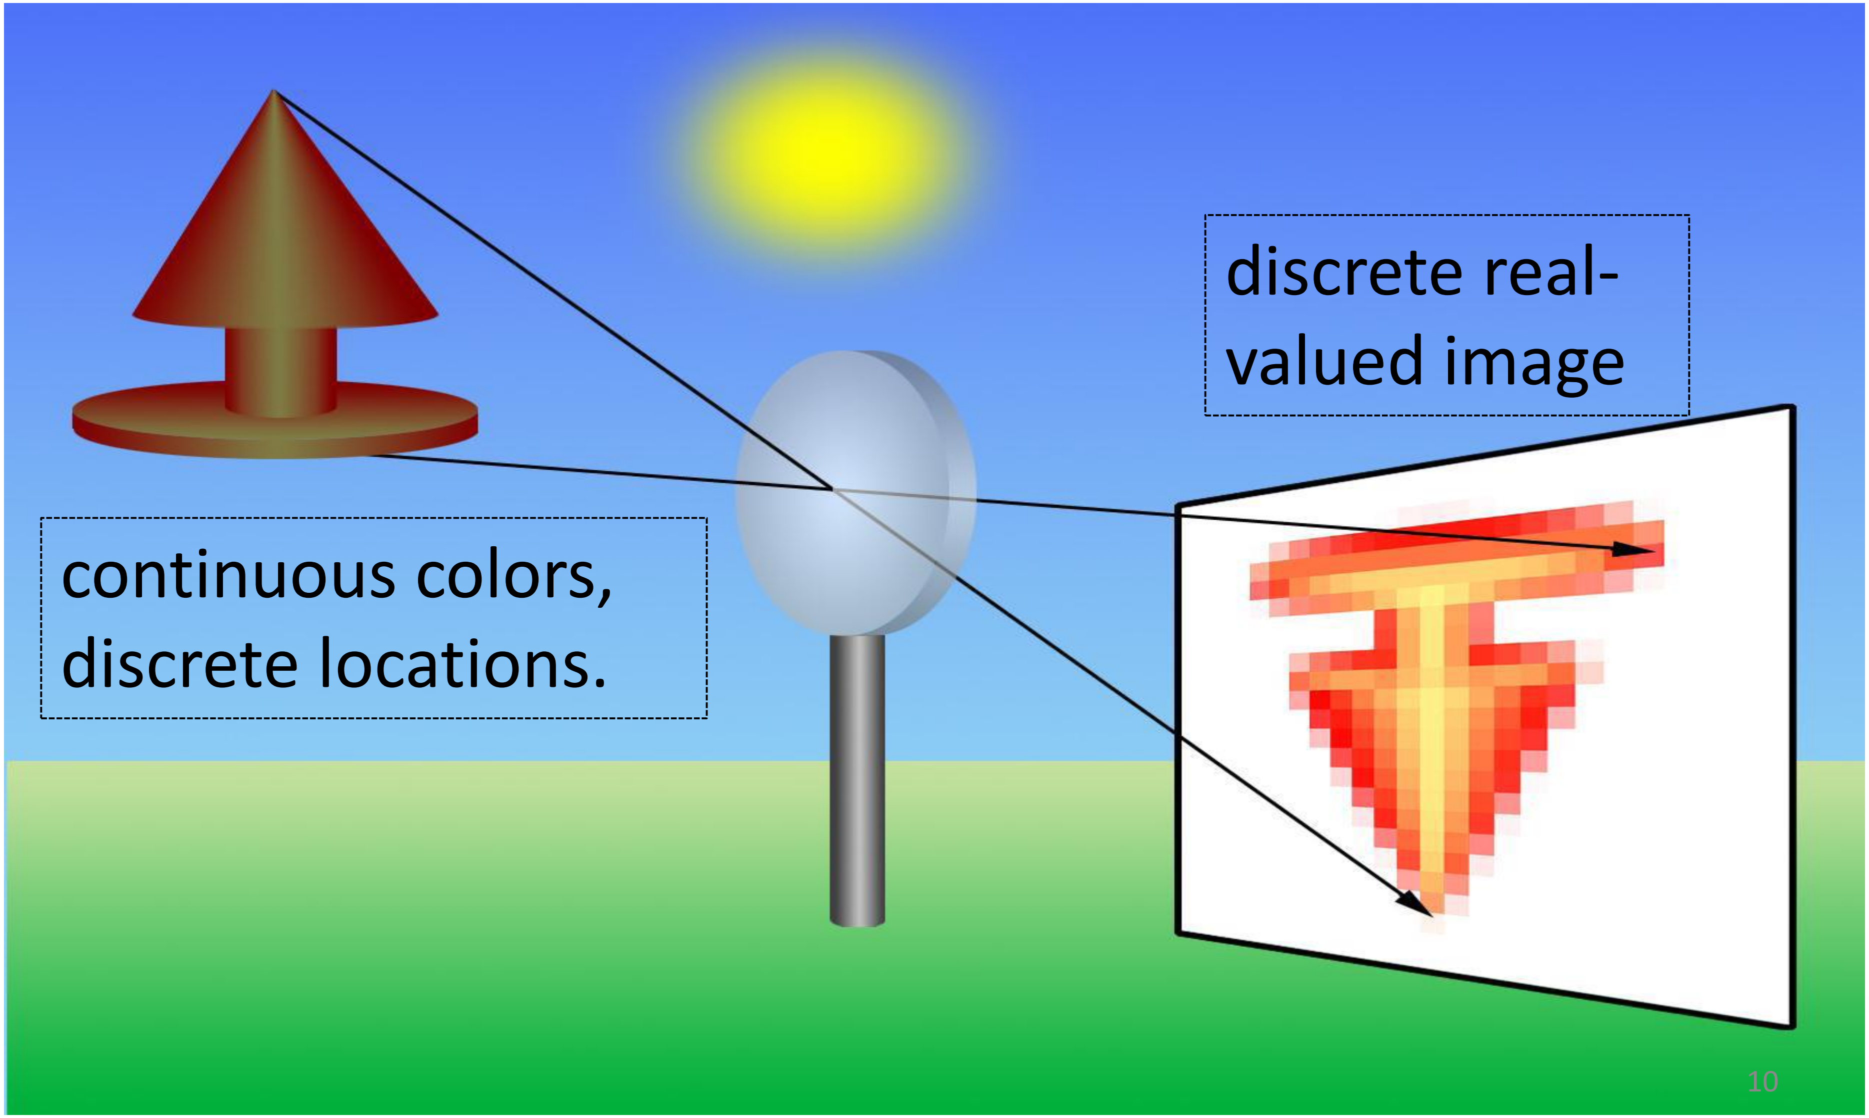

1. Image Formation (Quantization)

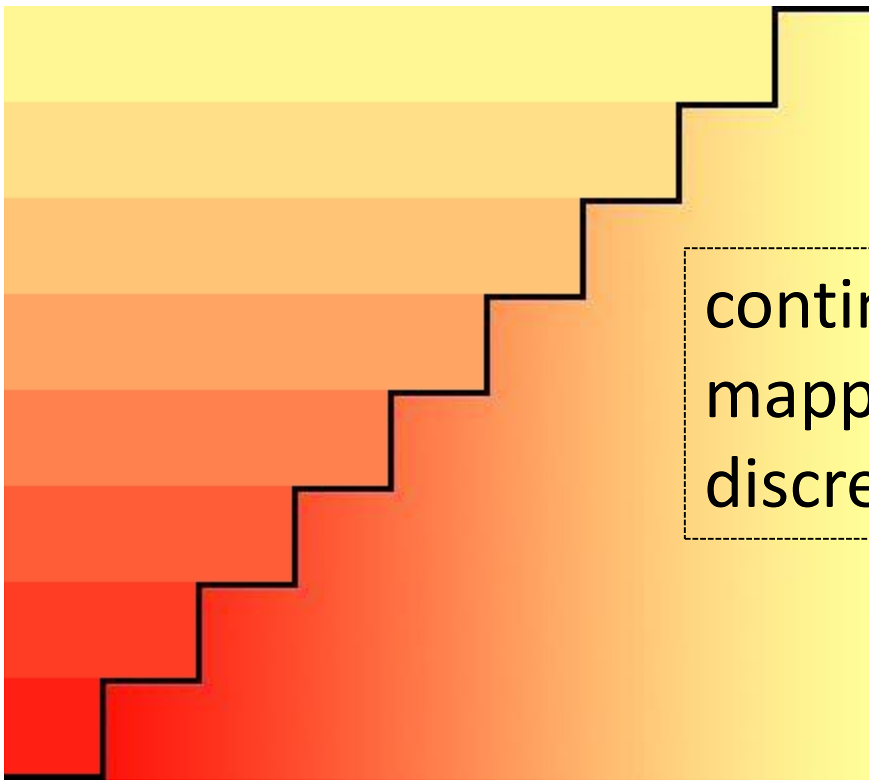

continuous colors mapped to a finite, discrete set of colors.

1. Image Formation (Sampling & Quantization)

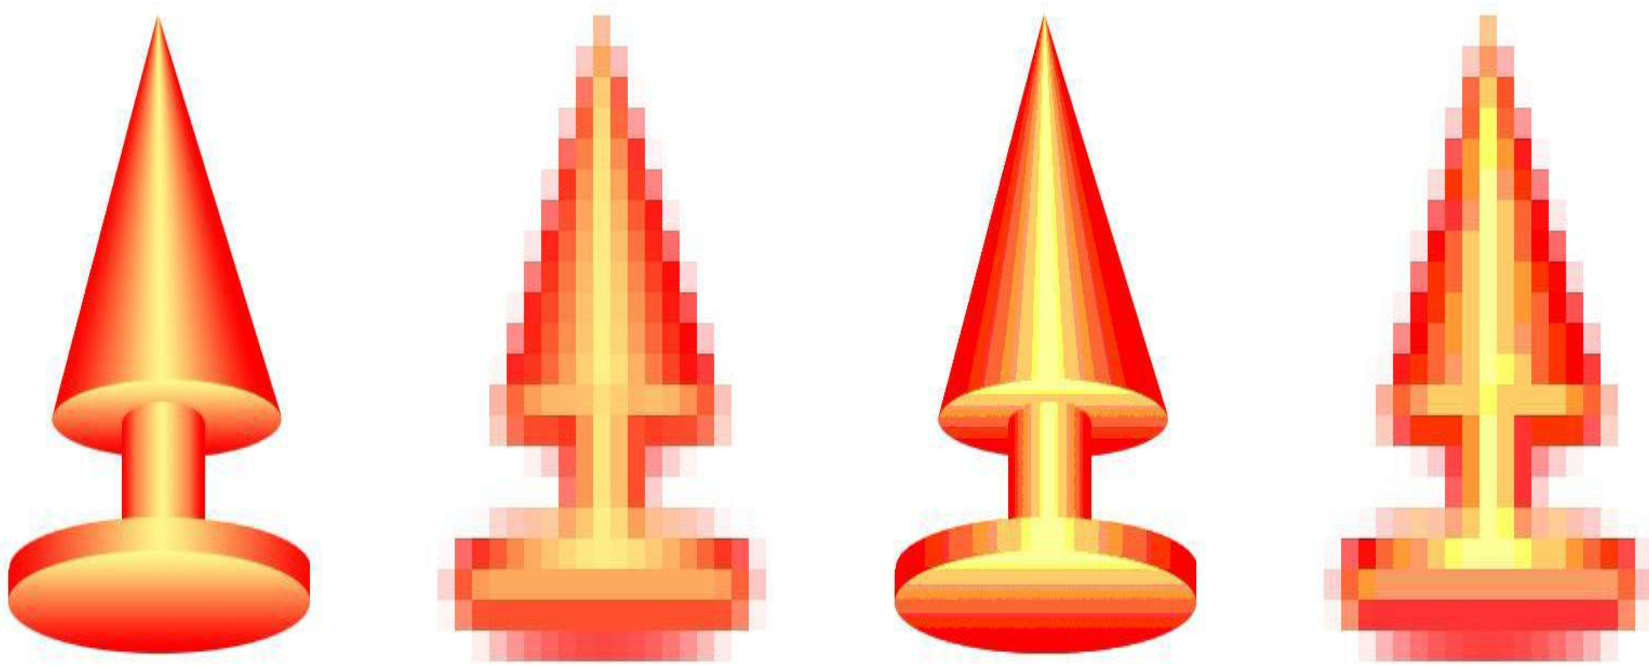

Real image   Sampled   quantized   sampled & quantized

1. Image Formation (Digital Image)

- A grid of squares, each of which contains a single color

- each square is called a pixel (for *picture element)*

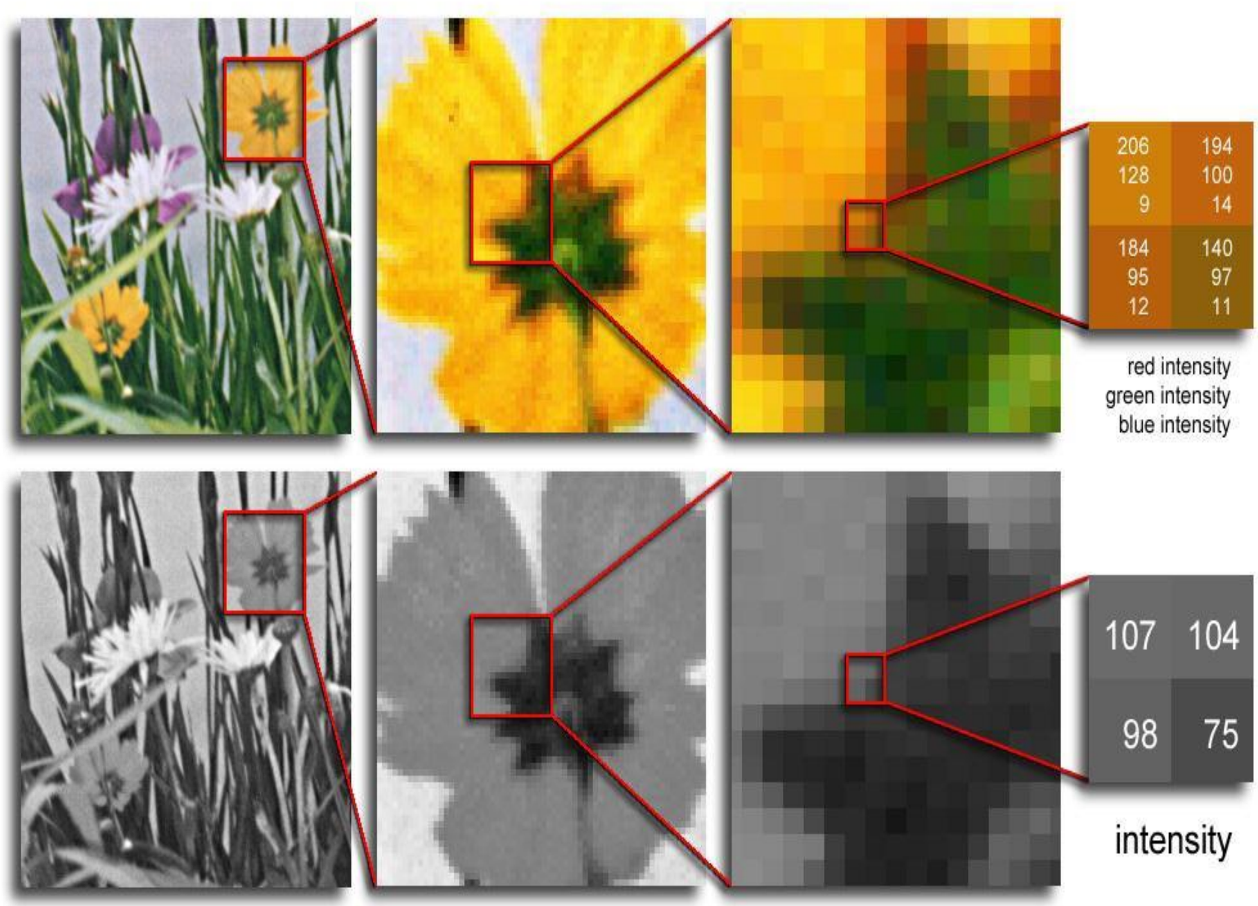

- Color images have 3 values per pixel; monochrome images have 1 value per pixel.

1. Image Formation (Color Images)

• Are constructed from three intensity maps.

• Each intensity map is projected through a color filter (*e.g., red, green, or blue, or cyan, magenta, or yellow) to create a monochrome image.*

• The intensity maps are overlaid to create a color image.

• Each pixel in a color image is a three element vector

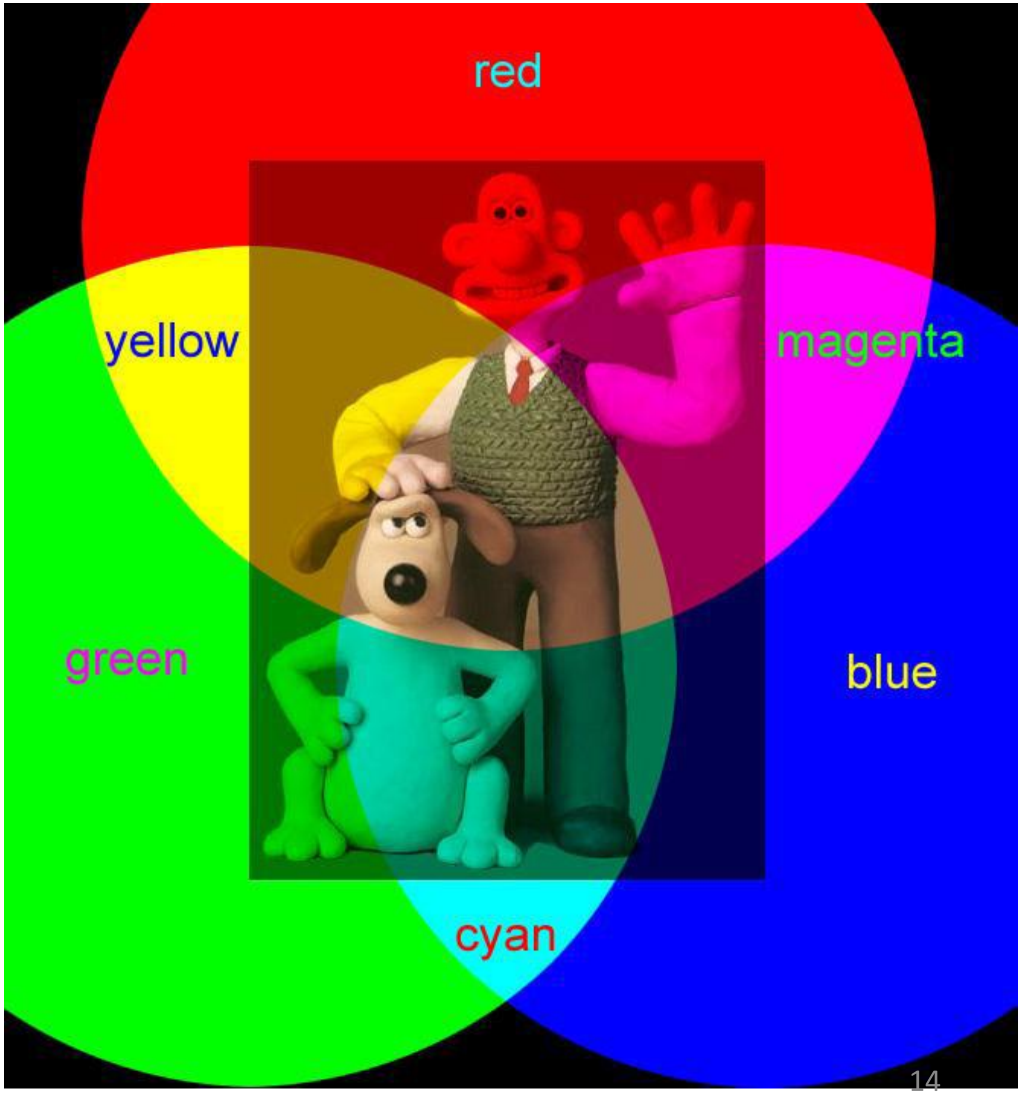

2. Point Processing

- gamma

- brightness

original

+ brightness

+ gamma

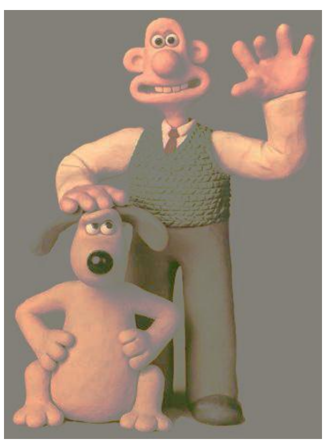

histogram mod

- contrast

original

+ contrast

histogram EQ

3. Color Processing

requires some knowledge of how we see colors

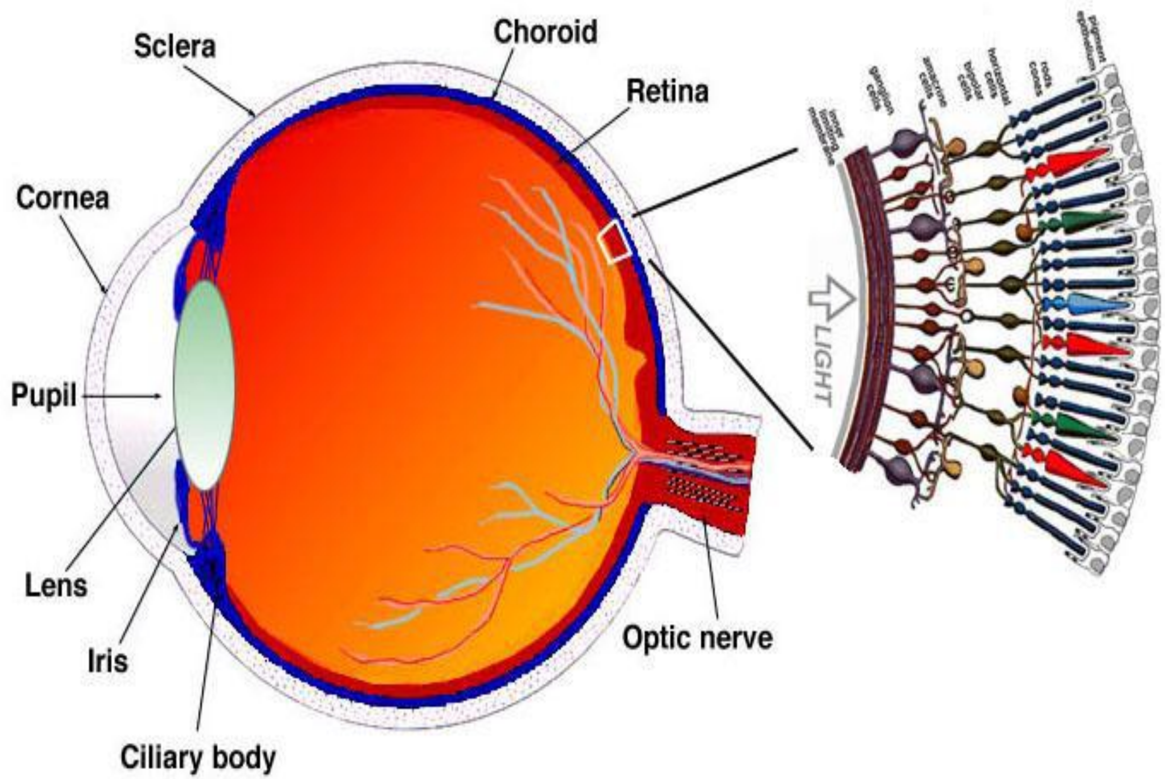

Fig. 1.1. A drawing of a section through the human eye with a schematic enlargement of the retina.

3. Color Processing

cone density near fovea

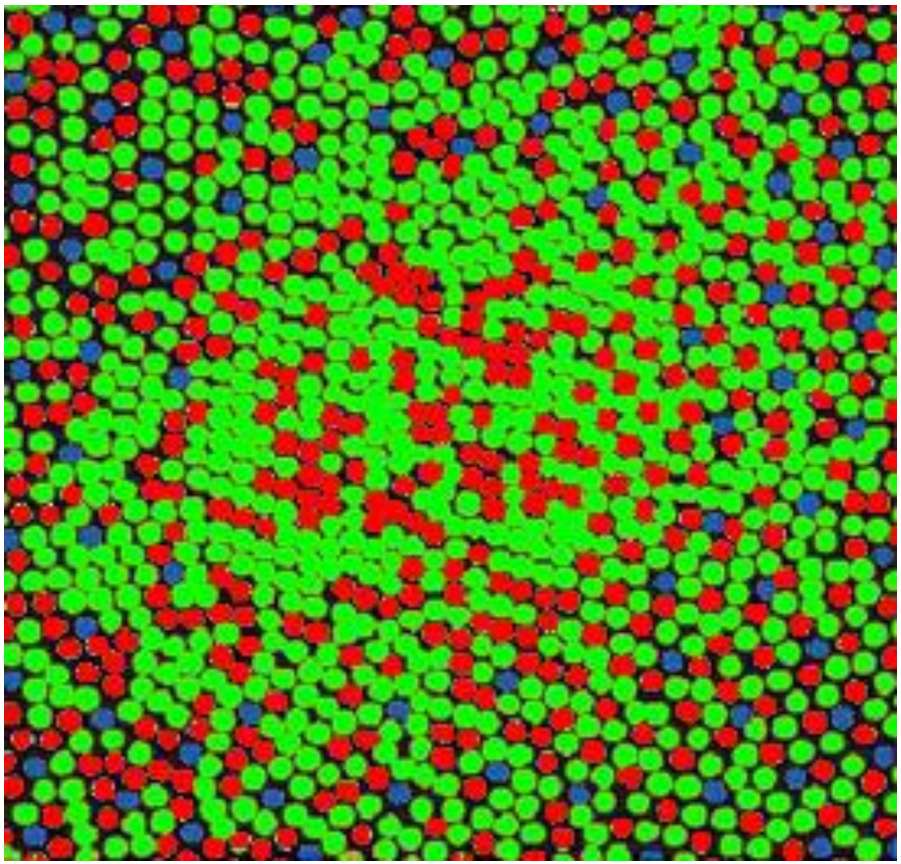

#(blue) << #(red) < #(green)

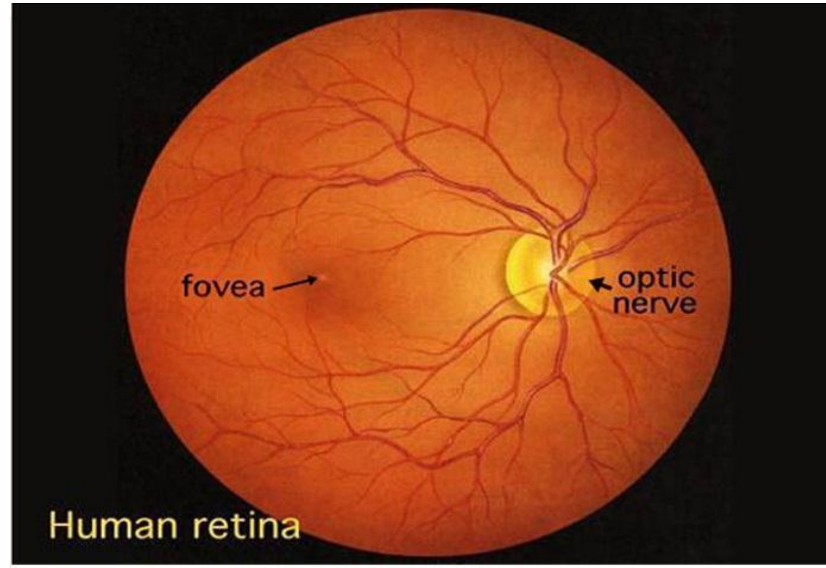

Fig. 1. Human retina as seen through an opthalmoscope.

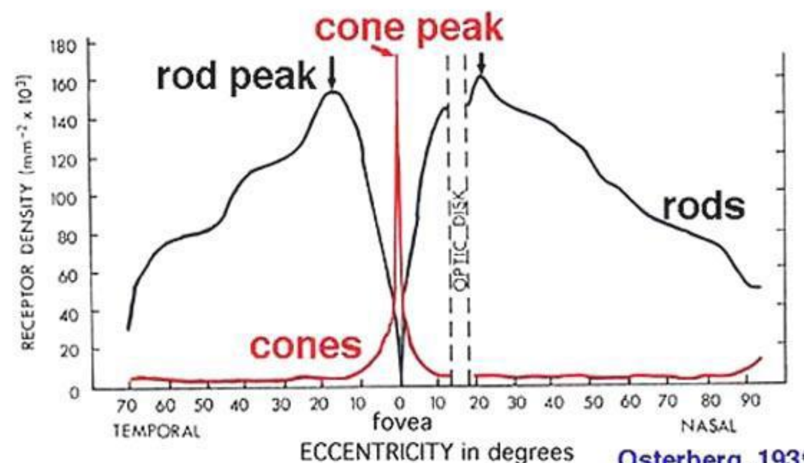

Fig. 20. Graph to show rod and cone densities along the horizontal meridian.

3. Color Processing

These are approximations of the responses to the visible spectrum of the “red”, “green”, and “blue” receptors of a typical human eye.

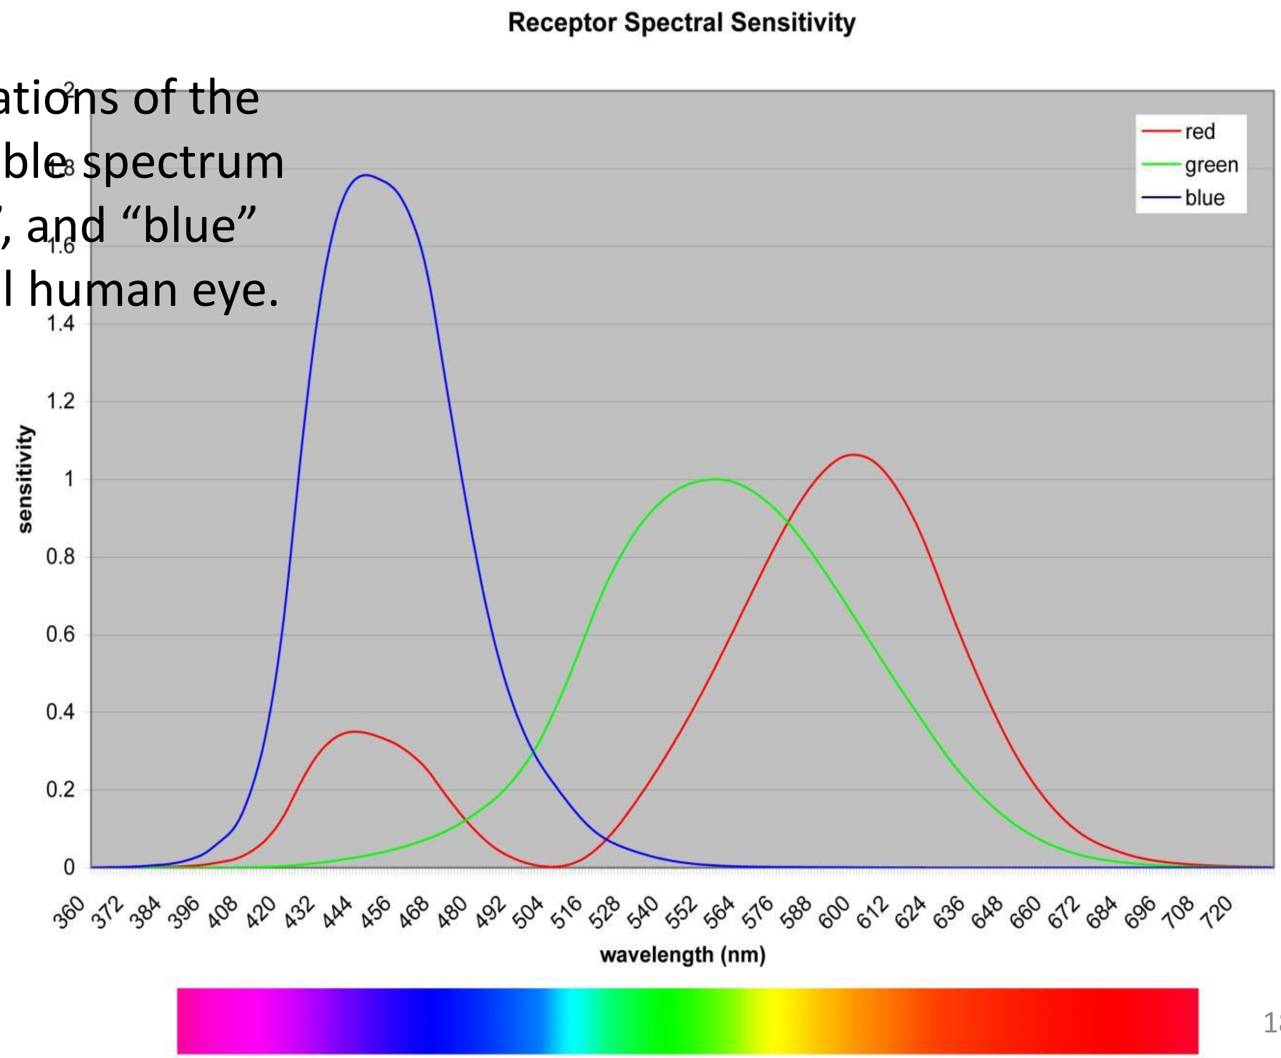

3. Color Processing

The eye has 3 types of photoreceptors: sensitive to red, green, or blue light.

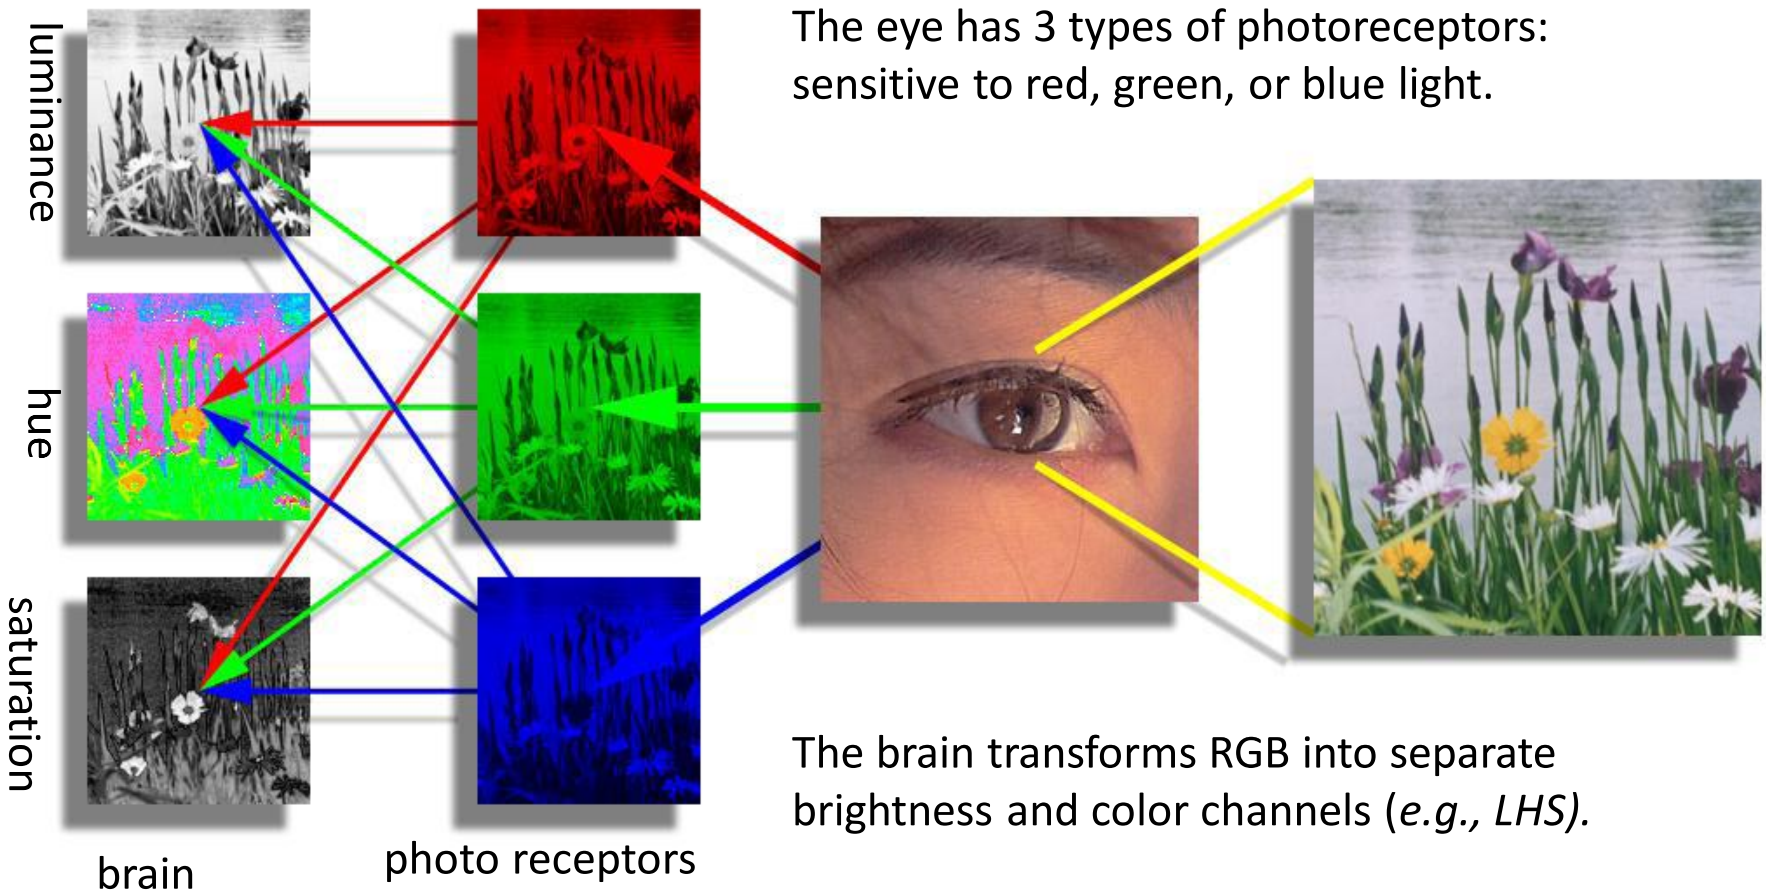

The brain transforms RGB into separate brightness and color channels (*e.g., LHS).*

3. Color Processing

luminance and chrominance (hue+saturation) are perceived with different resolutions, as are red, green and blue.

all bands

luminance

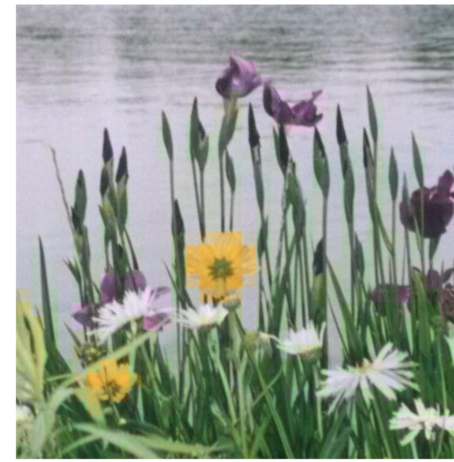

chrominance

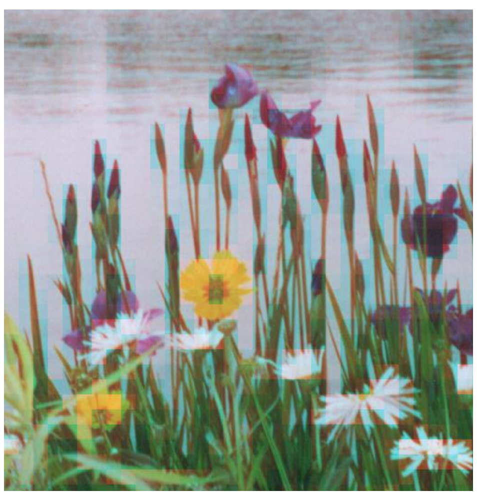

red

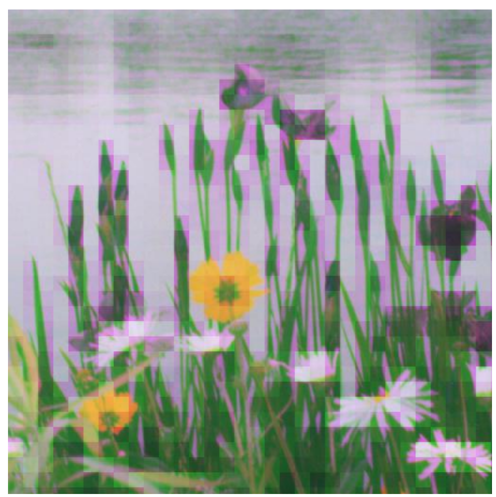

green

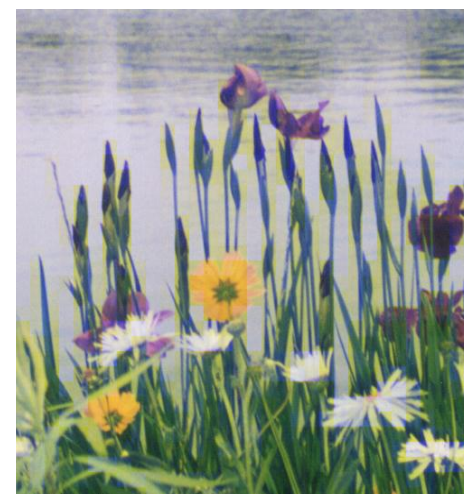

blue

3. Color Processing

Color Balance and Saturation

Uniform changes in color components result in change of tint.

If all G pixel values are multiplied by a > 1 then the image takes a green cast.

3. Color Processing

Color Transformations

Image aging: a transformation, Φ, that mapped:

$$\begin{bmatrix}17\\122\\114\end{bmatrix}=\Phi\left\{\begin{bmatrix}17\\121\\171\end{bmatrix}\right\} \quad \begin{bmatrix}222\\222\\185\end{bmatrix}=\Phi\left\{\begin{bmatrix}222\\222\\218\end{bmatrix}\right\} \quad \begin{bmatrix}240\\171\\103\end{bmatrix}=\Phi\left\{\begin{bmatrix}240\\171\\160\end{bmatrix}\right\} \quad \begin{bmatrix}236\\227\\106\end{bmatrix}=\Phi\left\{\begin{bmatrix}240\\230\\166\end{bmatrix}\right\}$$

4. The Fourier transform

Let $I(r,c)$ be a single-band (intensity) digital image with R rows and C columns. Then, $I(r,c)$ has Fourier representation

$$I(r,c) = \sum_{u=0}^{R-1} \sum_{v=0}^{C-1} \mathfrak{I}(u,v)\, e^{+i2\pi\left(\frac{ur}{R}+\frac{vc}{C}\right)},$$

these complex exponentials are 2D sinusoids.

where

$$\mathfrak{I}(u,v) = \frac{1}{RC} \sum_{r=0}^{R-1} \sum_{c=0}^{C-1} I(r,c)\, e^{-i2\pi\left(\frac{ur}{R}+\frac{vc}{C}\right)}$$

are the R x C Fourier coefficients.

4. The Fourier transform

2D Sinusoids:

$$I(r,c) = \frac{A}{2}\left\{\cos\left[\frac{2\pi}{\lambda}\left(\frac{c}{C}\cos\theta - \frac{r}{R}\sin\theta\right)+\varphi\right]+1\right\}$$

... are plane waves with grayscale amplitudes, periods in terms of lengths, ...

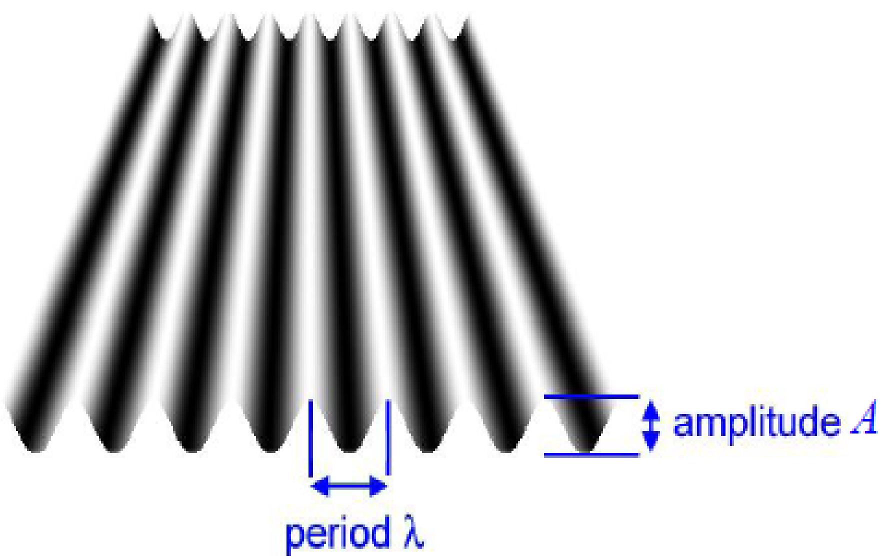

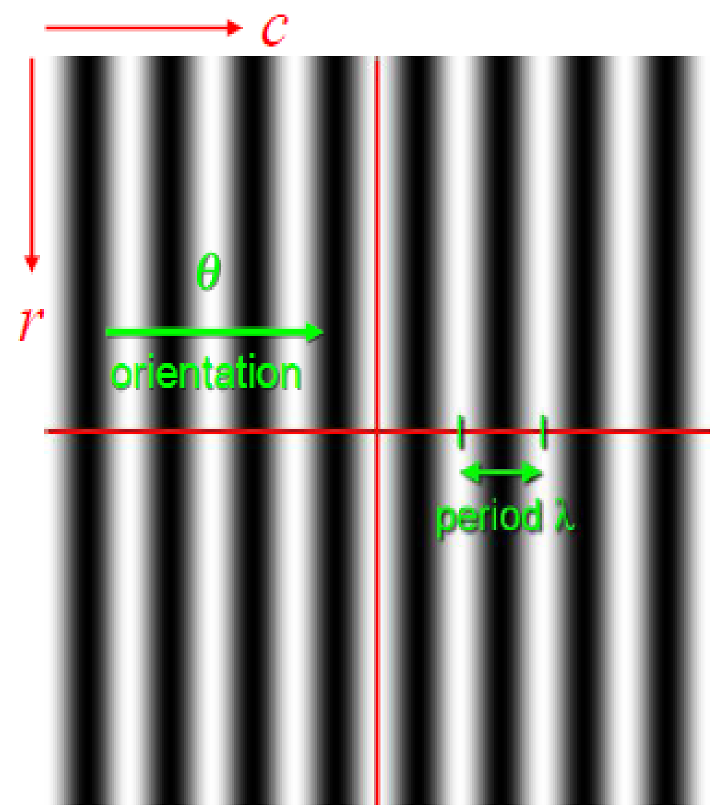

ϕ = phase shift

4. The Fourier transform

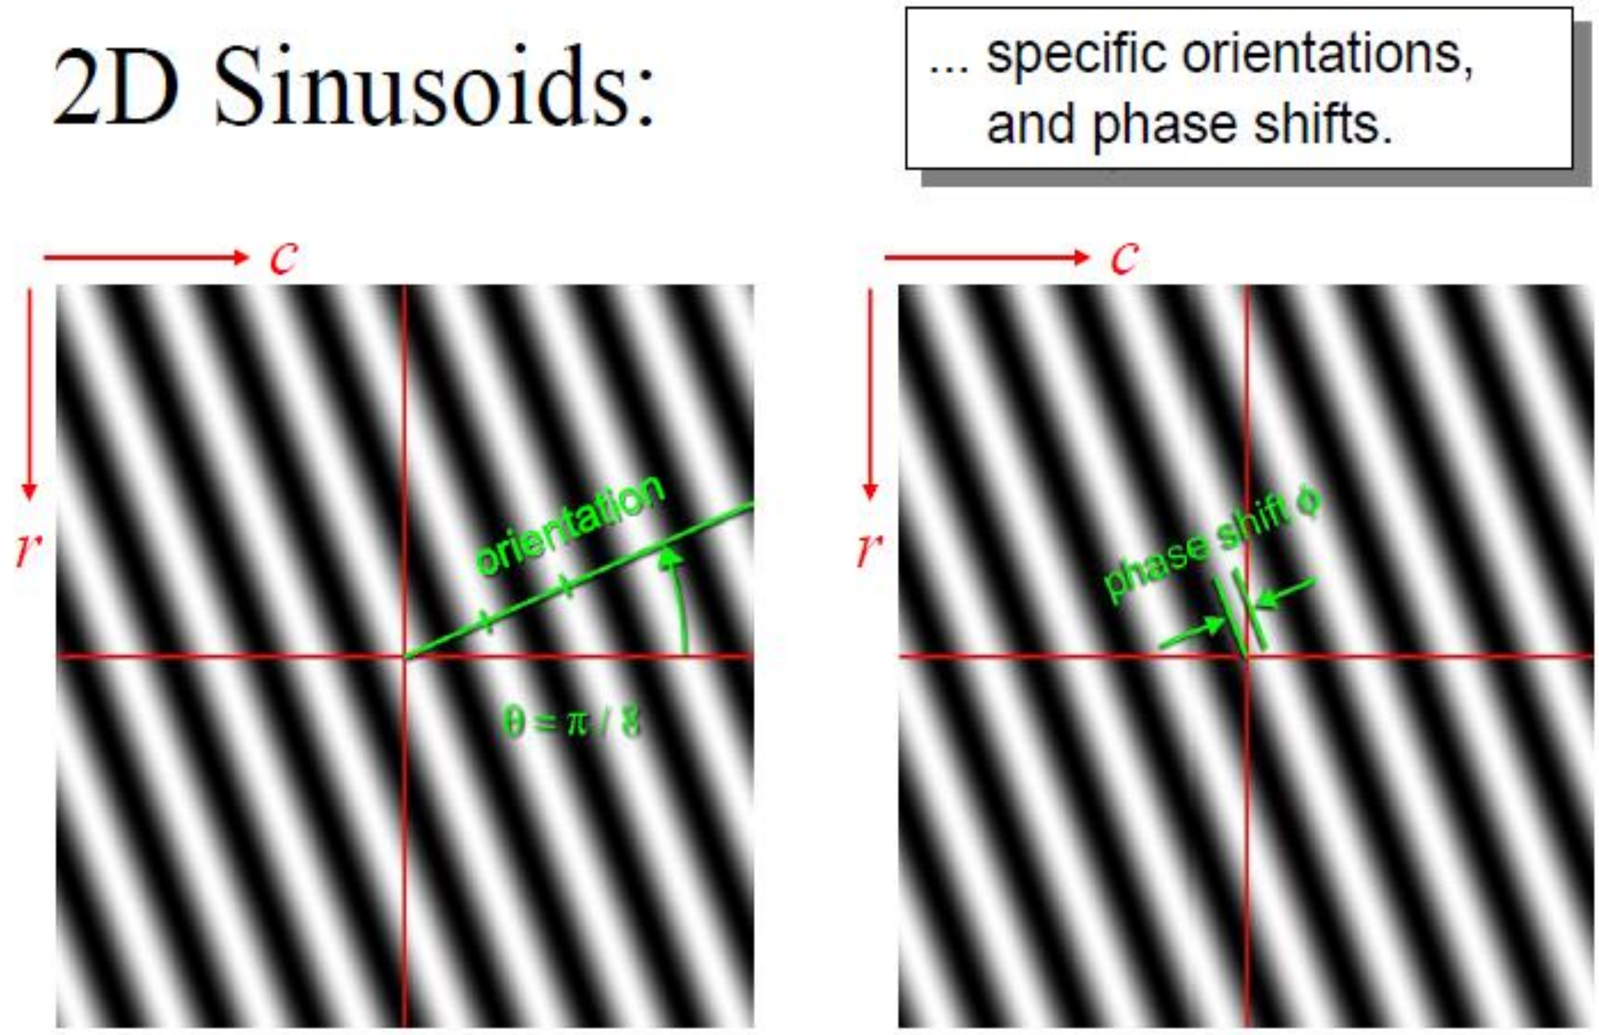

4. The Fourier transform

The Value of a Fourier Coefficient …

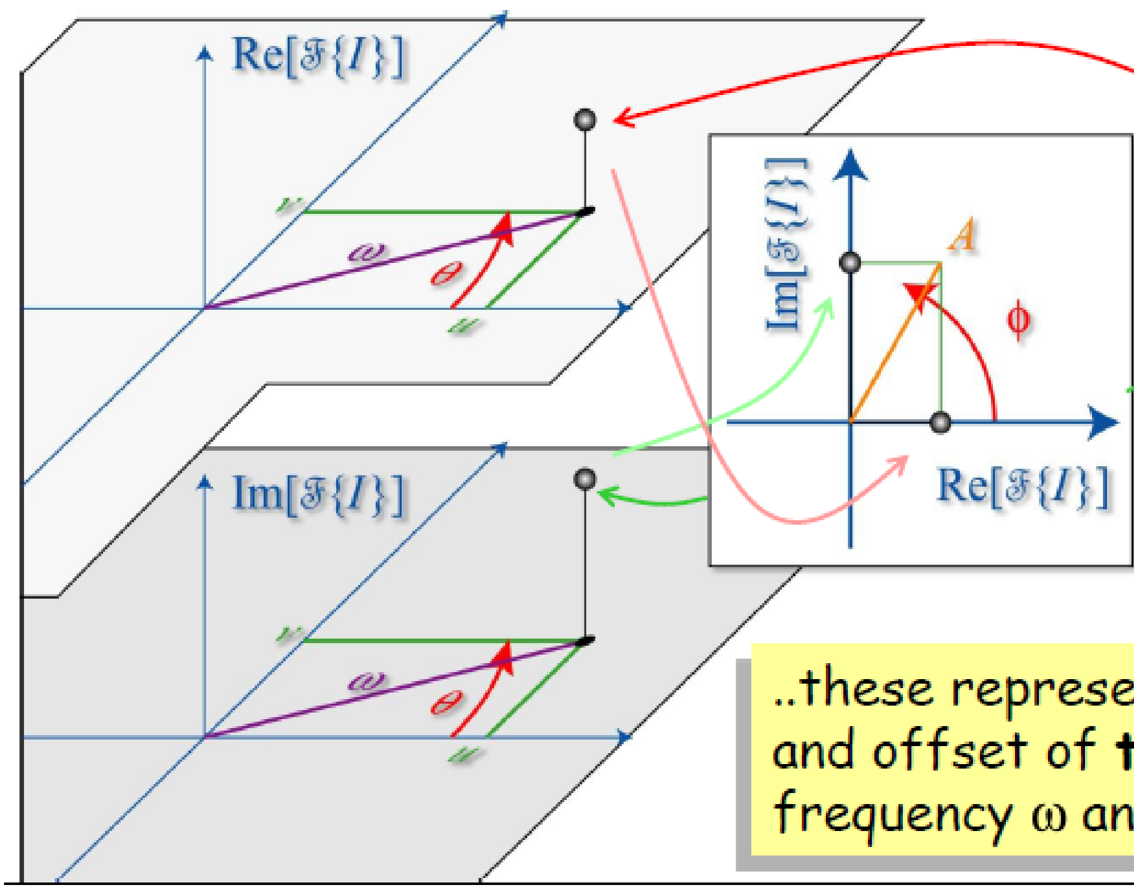

… is a complex number with a real part and an imaginary part.

If you represent that number as a magnitude, A, and a phase, ϕ, …

..these represent the amplitude and offset of **the** sinusoid with frequency ω and direction θ.

4. The Fourier transform

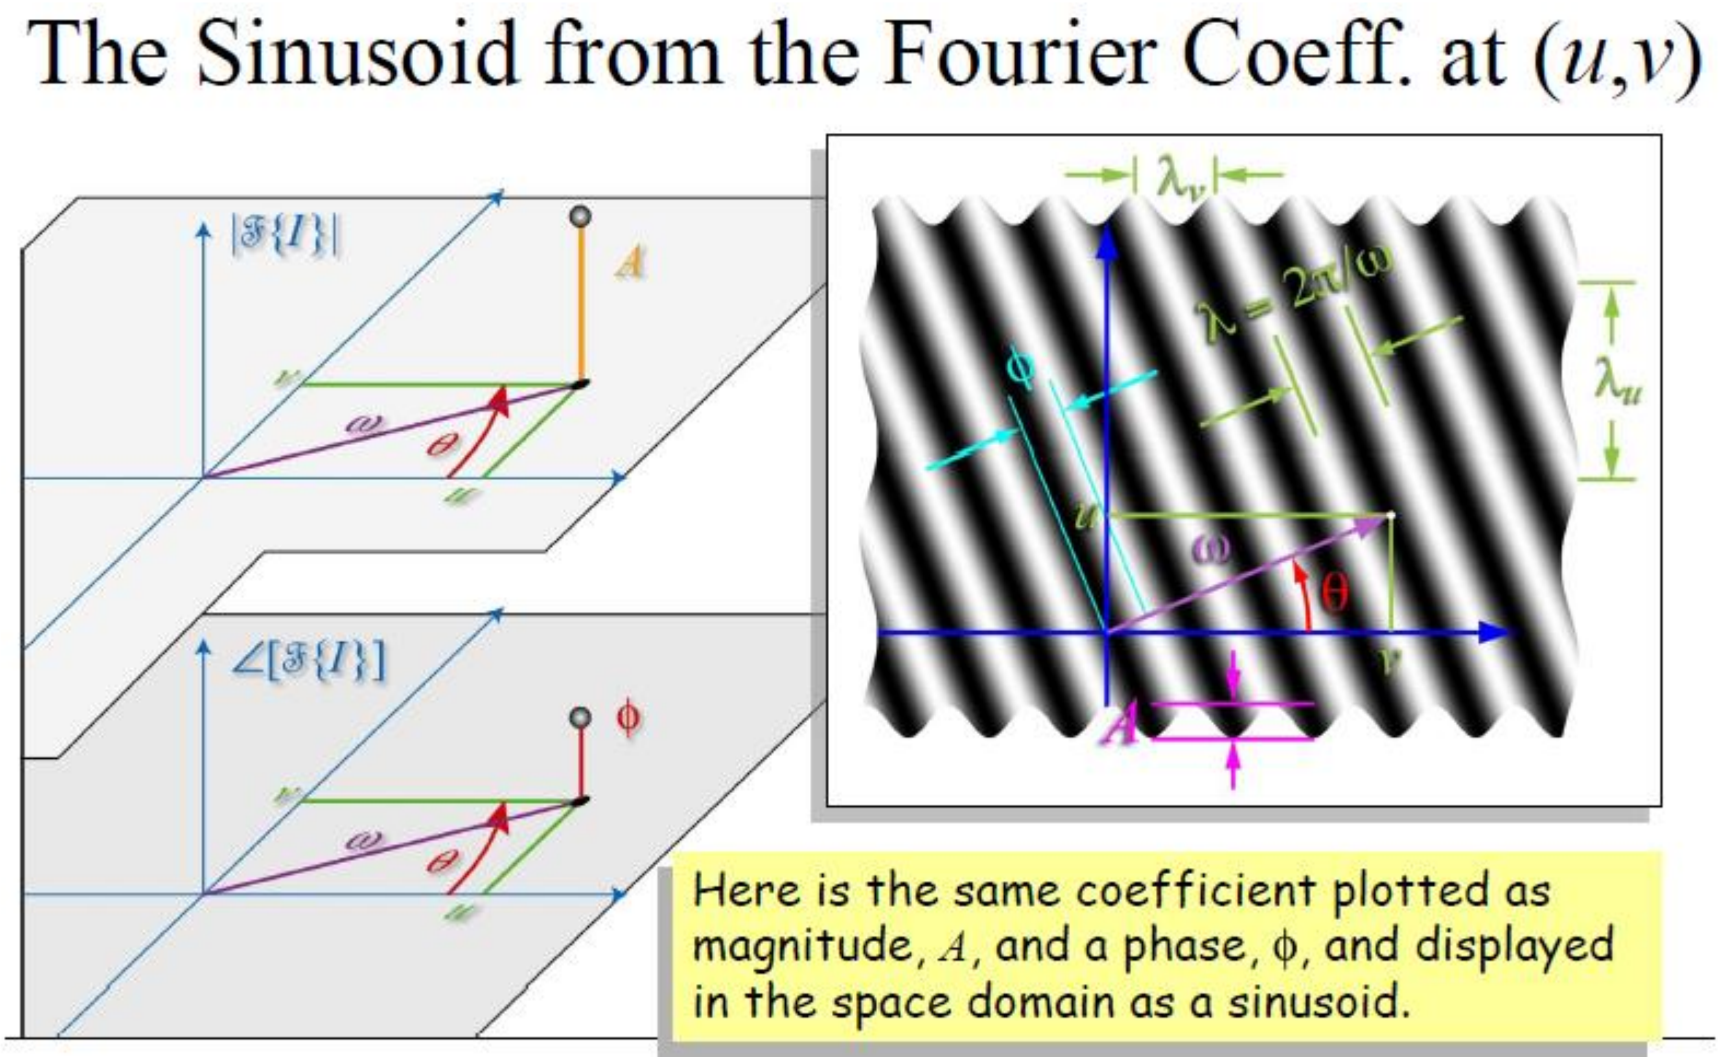

4. The Fourier transform

The Fourier Transform of an Image

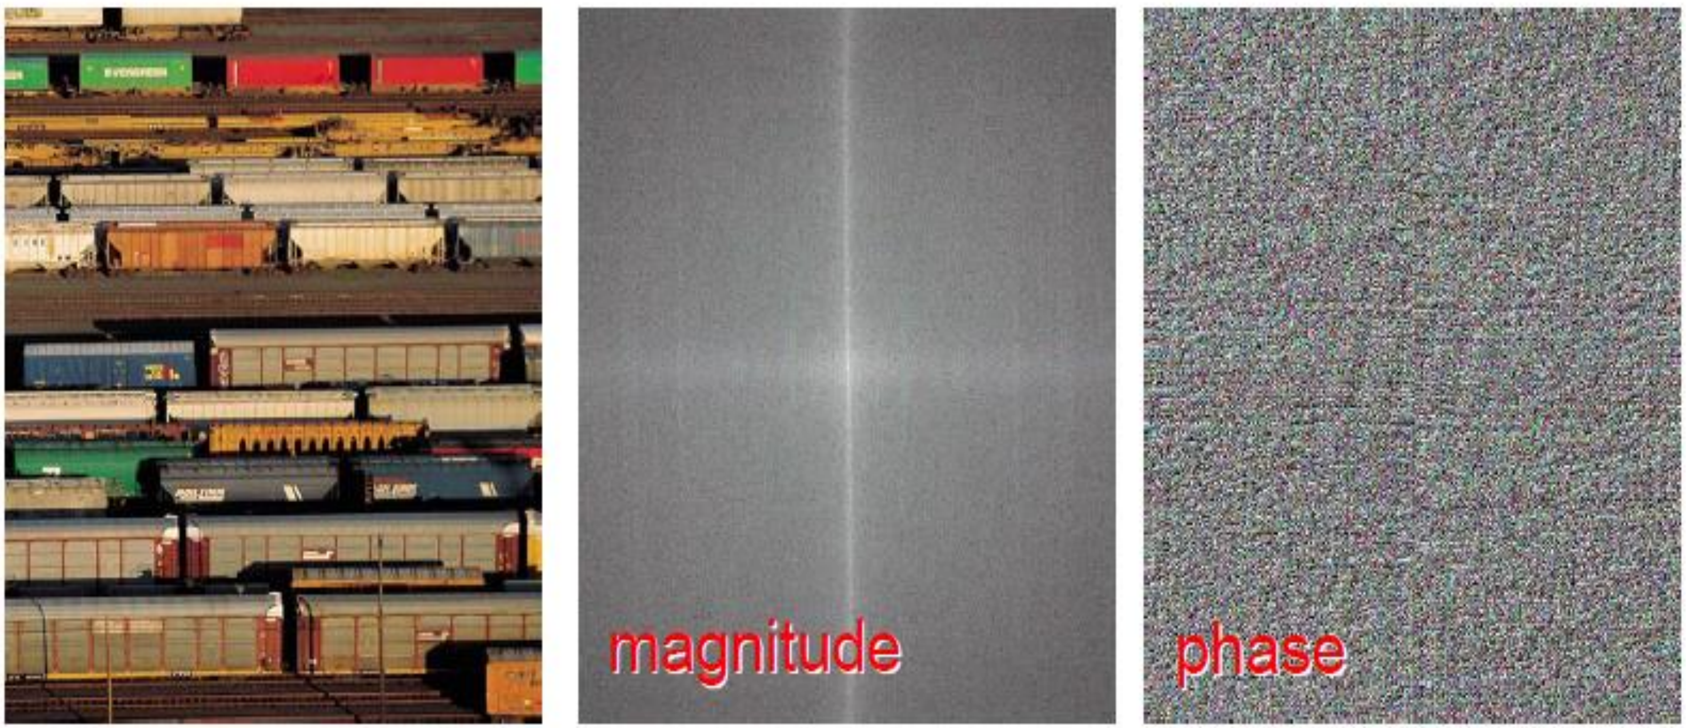

I $|\mathfrak{F}\{I\}|$ $\angle[\mathfrak{F}\{I\}]$

4. The Fourier transform

Continuous Fourier Transform

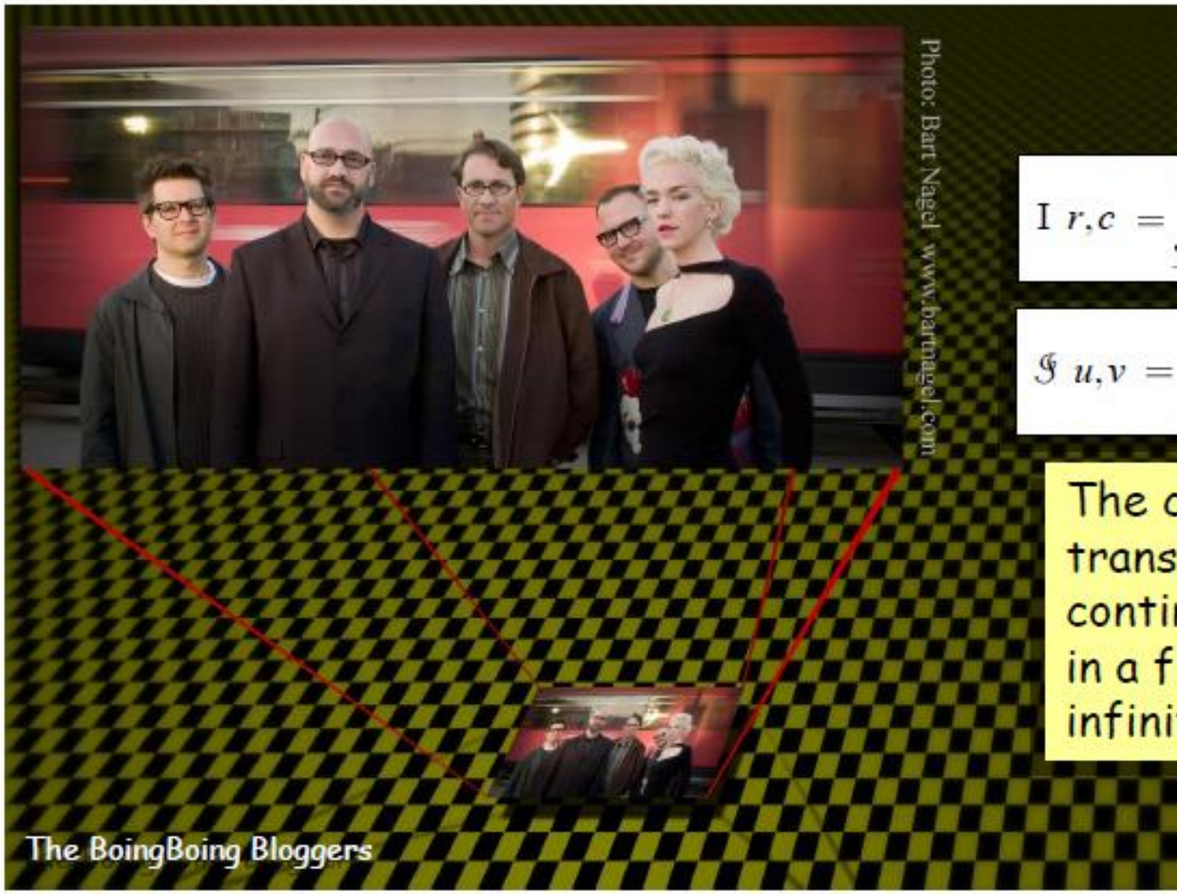

$$\mathrm{I}\ r,c\ = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \Im\ u,v\ \ e^{+i2\pi(uc+vr)}\, dudv$$

$$\Im\ u,v\ = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathrm{I}\ r,c\ \ e^{-i2\pi(uc+vr)}\, dcdr$$

The continuous Fourier transform assumes a continuous image exists in a finite region of an infinite plane.

4. The Fourier transform

Discrete Fourier Transform

The discrete Fourier transform assumes a digital image exists on a closed surface, a torus.

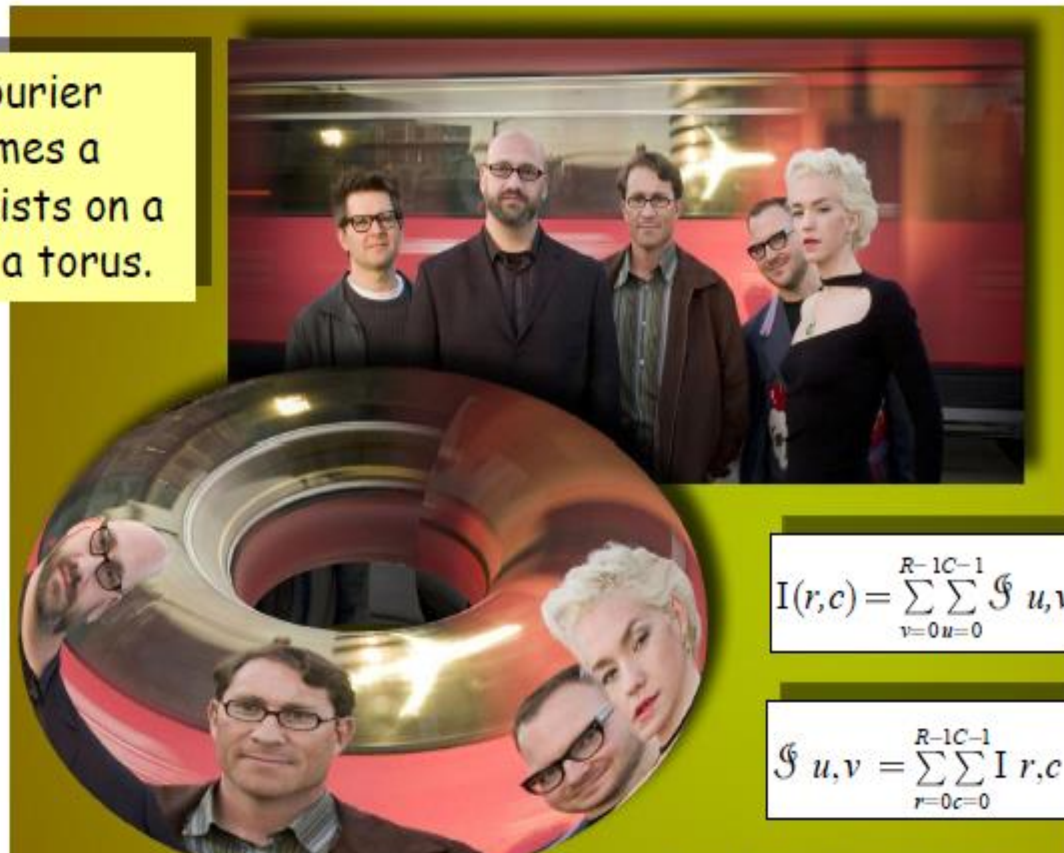

Photo: Bart Nagel www.bartnagel.com

$$I(r,c)=\sum_{v=0}^{R-1}\sum_{u=0}^{C-1}\Im(u,v)\,e^{+i2\pi\left(\frac{uc}{C}+\frac{vr}{R}\right)}$$

$$\Im(u,v)=\sum_{r=0}^{R-1}\sum_{c=0}^{C-1}I(r,c)\,e^{-i2\pi\left(\frac{cu}{C}+\frac{rv}{R}\right)}$$

5. Convolution

Sums of shifted and weighted copies of images or Fourier transforms.

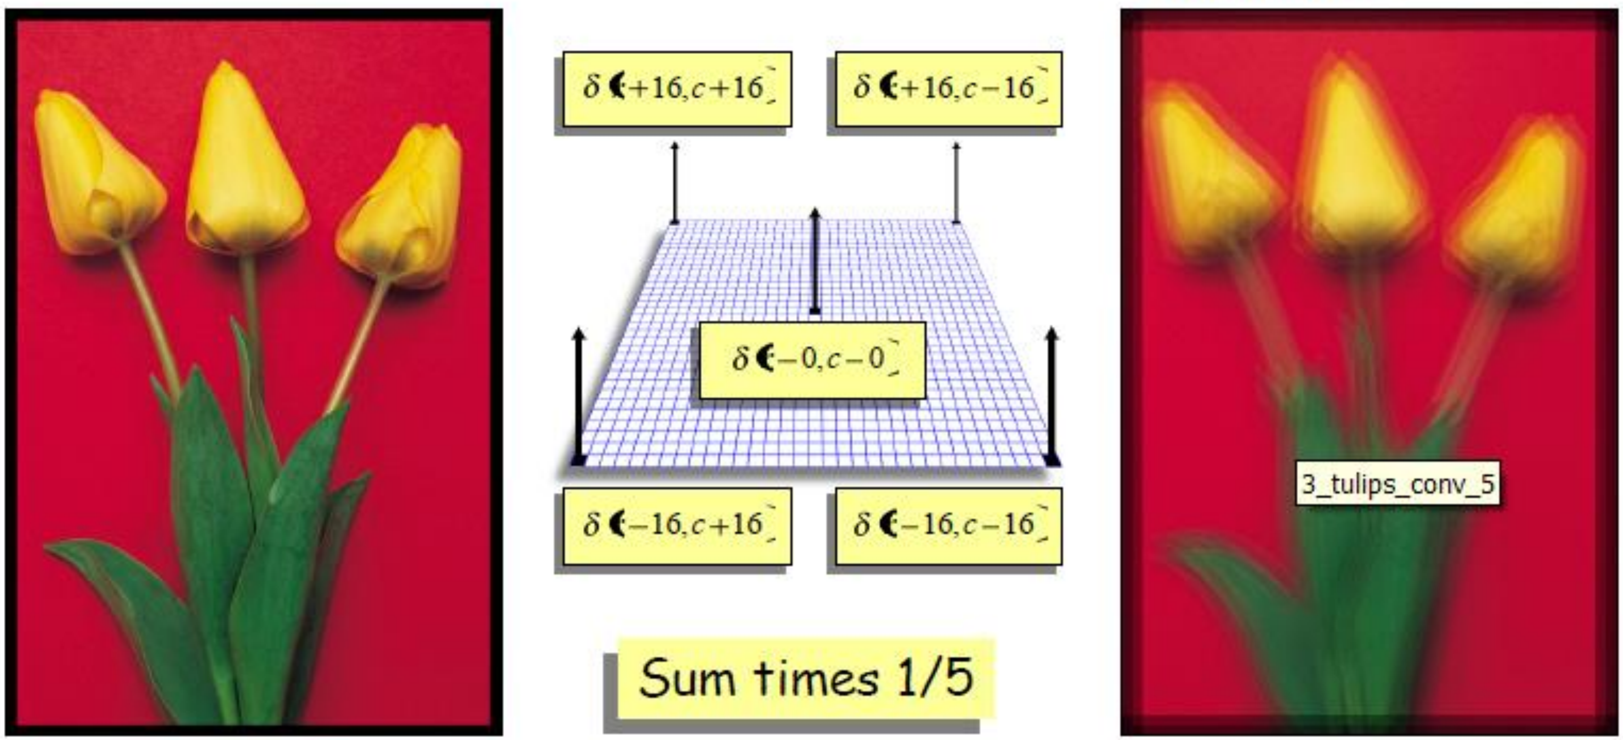

5. Convolution

Convolution Property of the Fourier Transform

Let functions $f(r,c)$ and $g(r,c)$ have Fourier Transforms $F(u,v)$ and $G(u,v)$.

Then,

$$\mathfrak{F}\{f * g\} = F \cdot G.$$

Moreover,

$$\mathfrak{F}\{f \cdot g\} = F * G.$$

$*$ represents convolution

$\cdot$ represents pointwise multiplication

Then, a spatial convolution can be computed by

$$f * g = \mathfrak{F}^{-1}[F \cdot G]$$

The Fourier Transform of a product equals the convolution of the Fourier Transforms. Similarly, the Fourier Transform of a convolution is the product of the Fourier Transforms

5. Convolution

Sampling, Aliasing, & Frequency Convolution

$$\text{samp}_{1/N}(u,v) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(u - \tfrac{j}{N})\,\delta(v - \tfrac{k}{N})$$

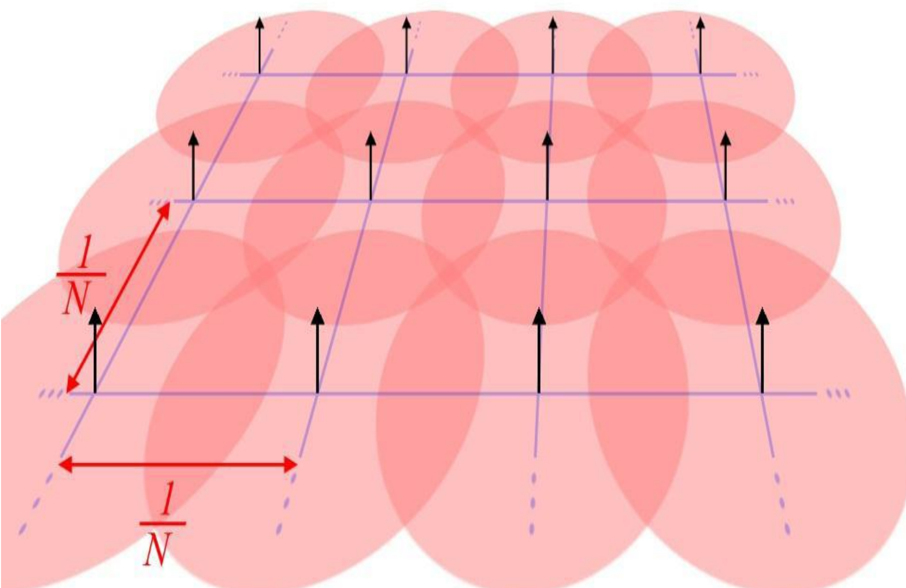

aliasing

$$\text{samp}_{1/N}(u,v) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(u - \tfrac{j}{N})\,\delta(v - \tfrac{k}{N})$$

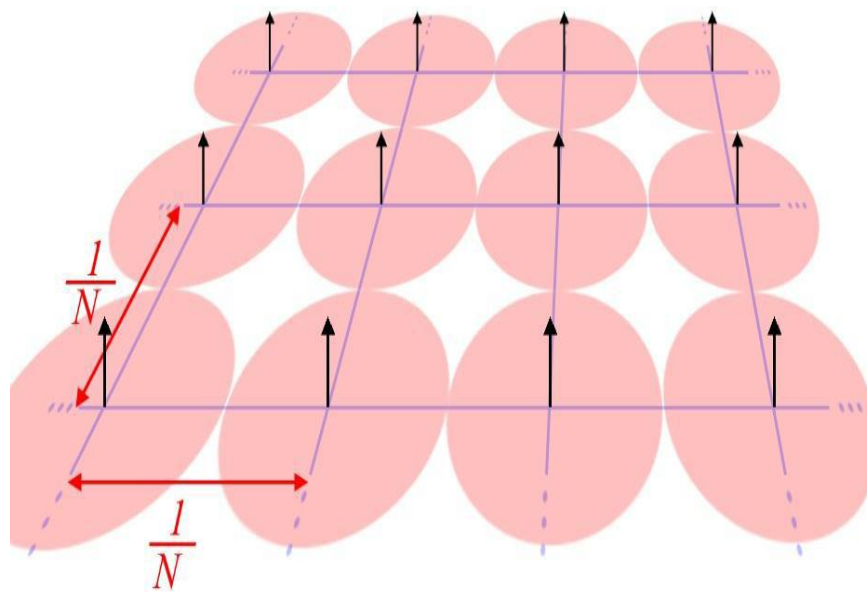

no aliasing (smooth lines)

5. Convolution

Sampling, Aliasing, & Frequency Convolution

(a) aliased
(b) power spectrum
(c) unaliased
(d) power spectrum

6. Image sampling, warping, and stitching

Resampling

(resizing)

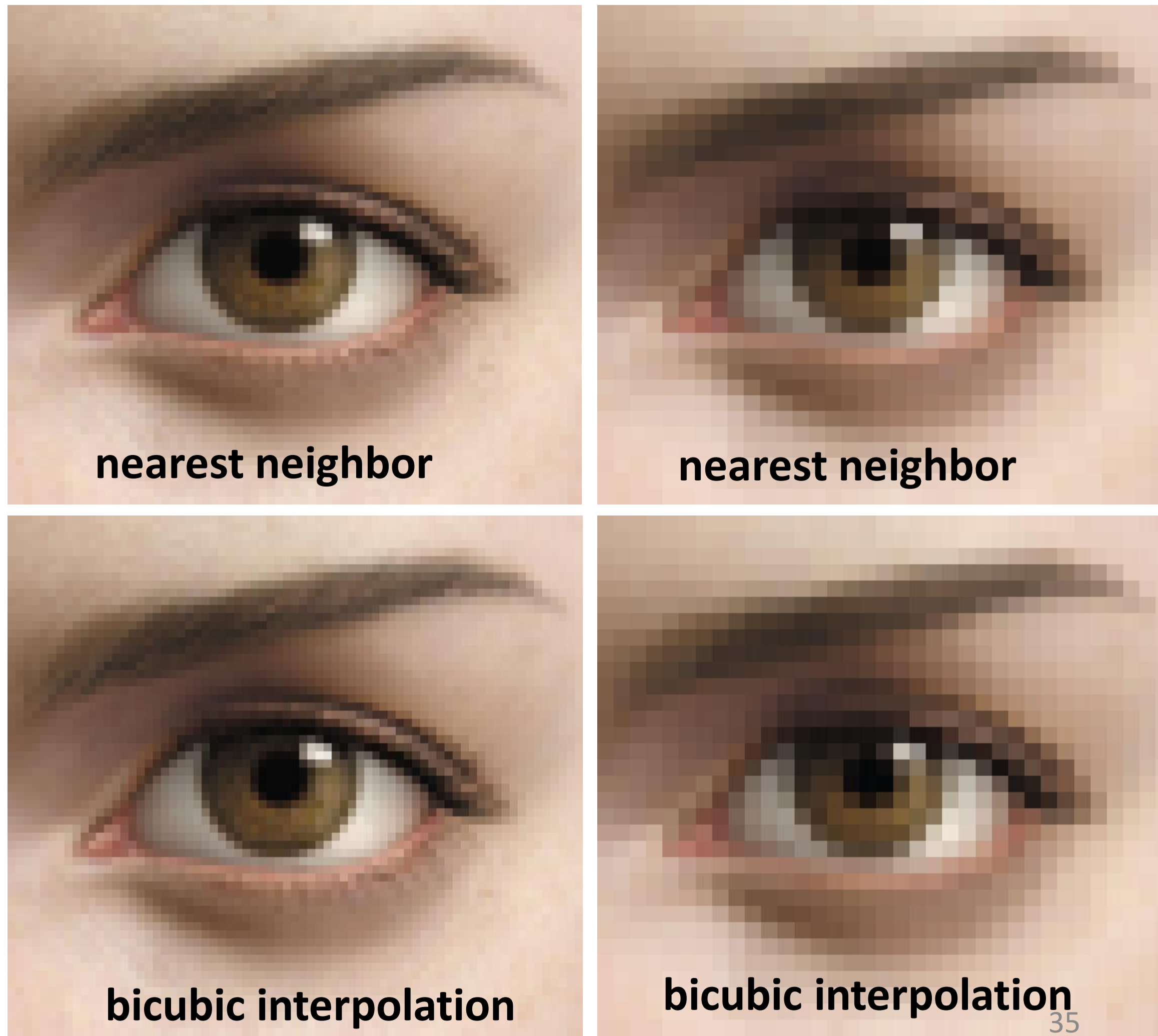

6. Image sampling, warping, and stitching

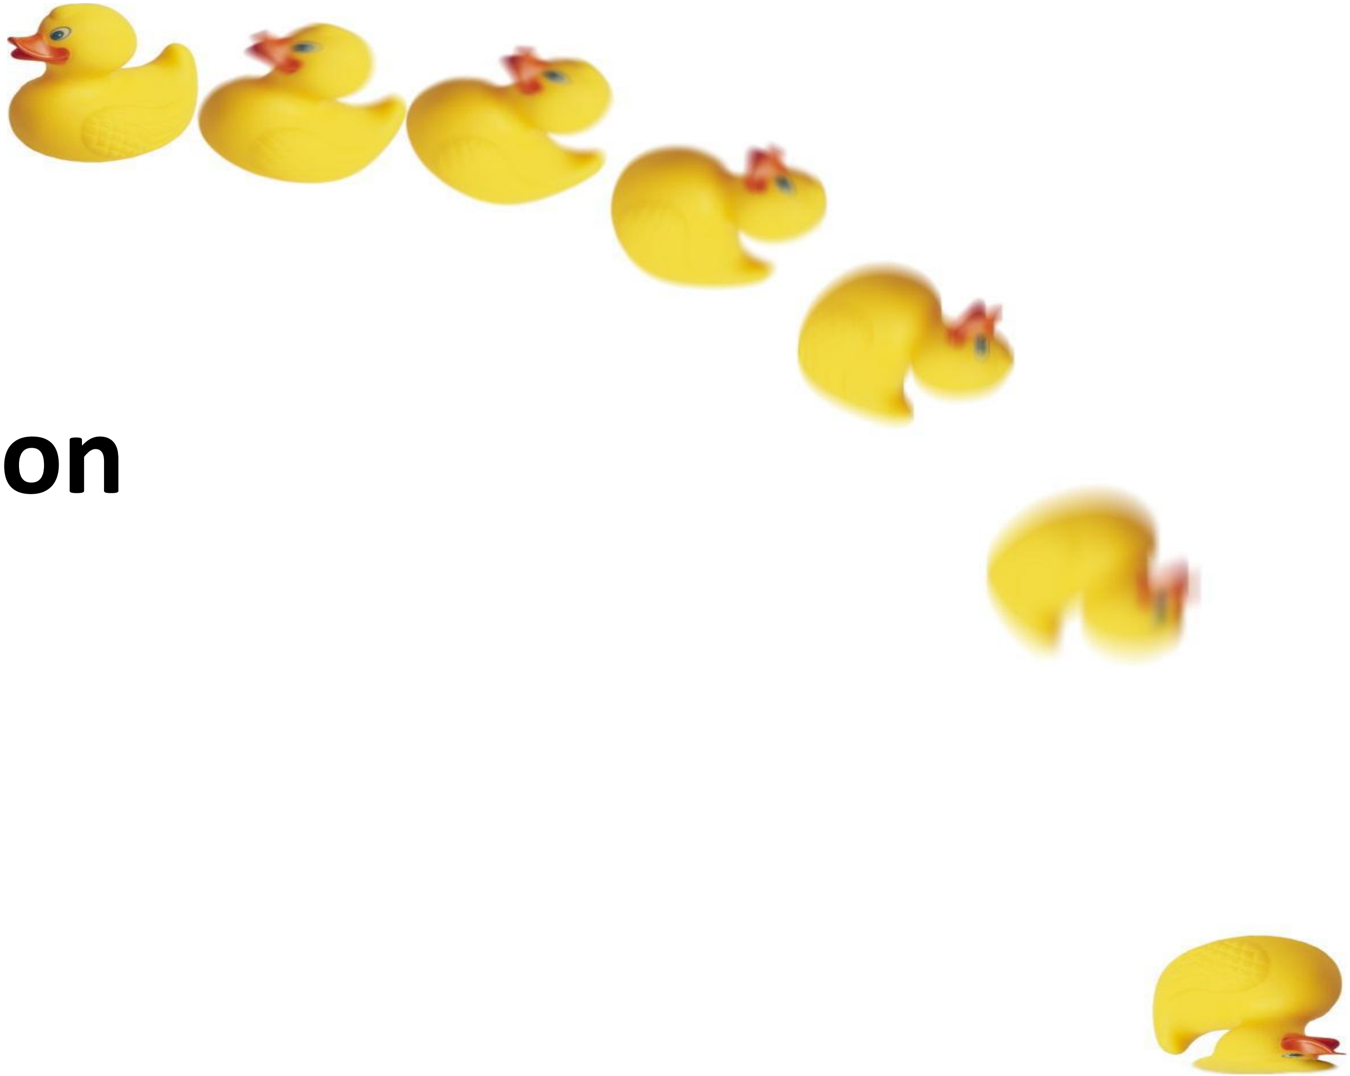

Rotation

6. Image sampling, warping, and stitching

7. Frequency Domain (FD) Filtering

Low-pass Filter

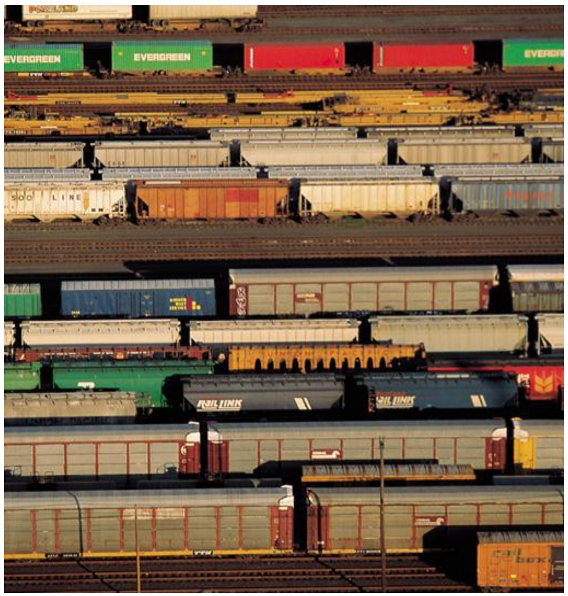

Original Image

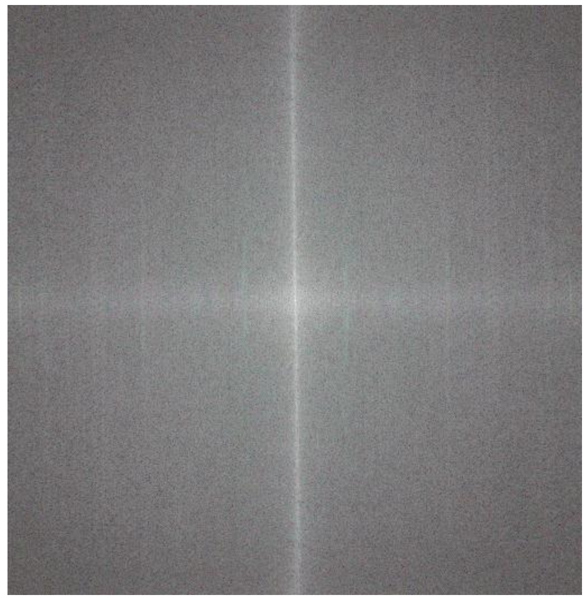

Power Spectrum

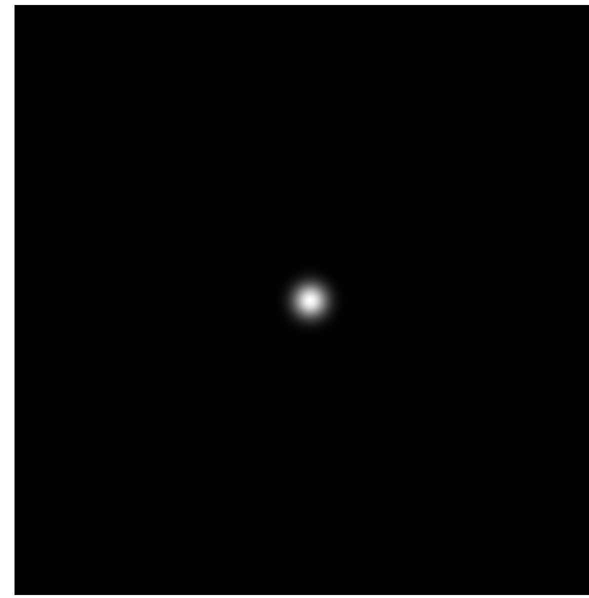

Gaussian LPF

7. Frequency Domain (FD) Filtering

Low-pass Filter

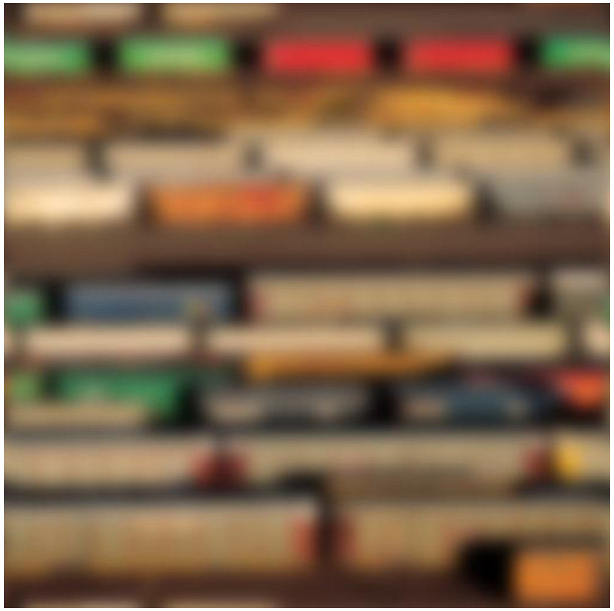

Filtered Image

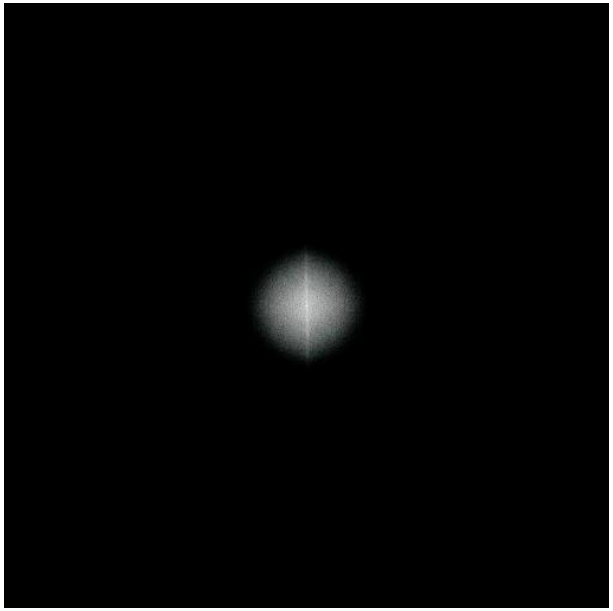

Filtered Power Spectrum

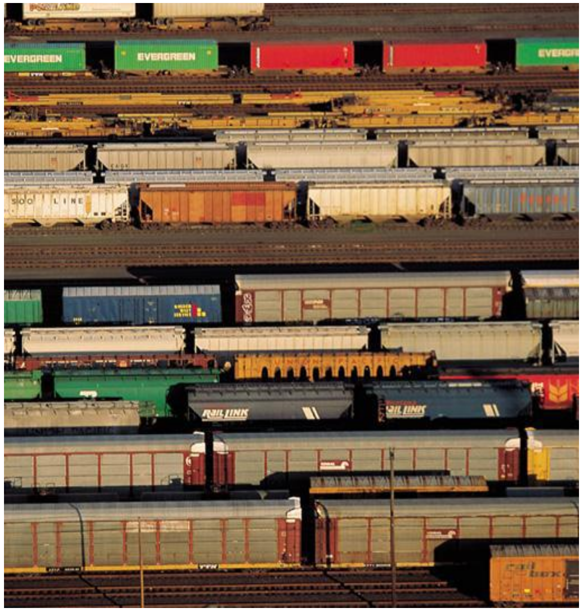

Original Image

7. Frequency Domain (FD) Filtering

High-pass Filter

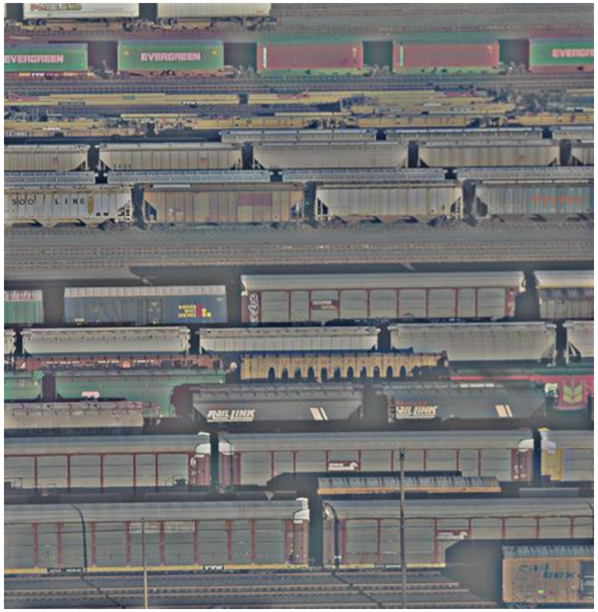

Filtered Image

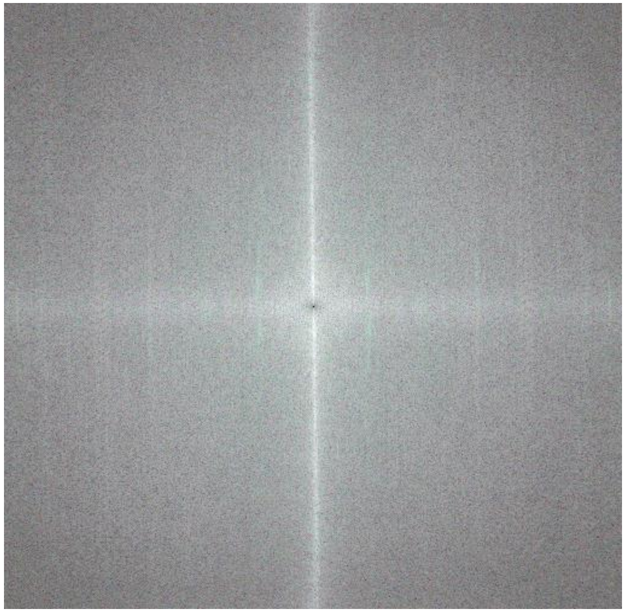

Filtered Power Spectrum

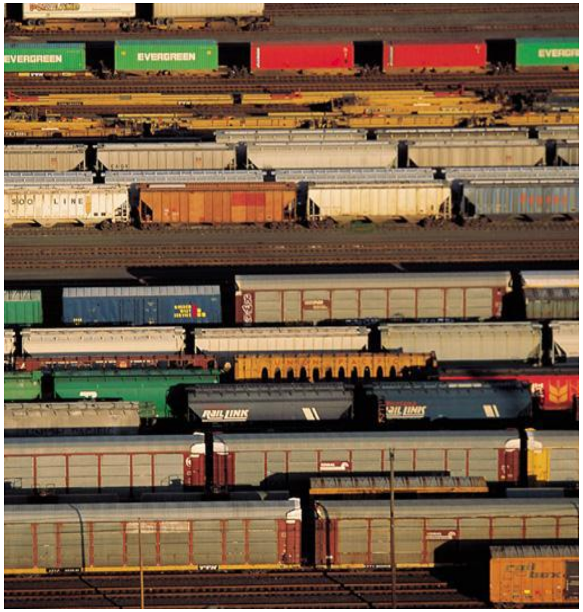

Original Image

8. Spatial Filtering

blurred

original

sharpened

band-pass filter

8. Spatial Filtering

regional

vertical

original

zoom

rotational

9. Noise Reduction

blurred image

color noise

color-only blur

9. Noise Reduction

blurred image

color noise

5x5 Wiener filter

9. Noise Reduction

periodic noise **original** **frequency tuned filter**

9. Noise Reduction

Shot Noise or Salt & Pepper Noise

+ shot noise

s&p noise

- shot noise

9. Noise Reduction

Nonlinear Filters: the Median

original

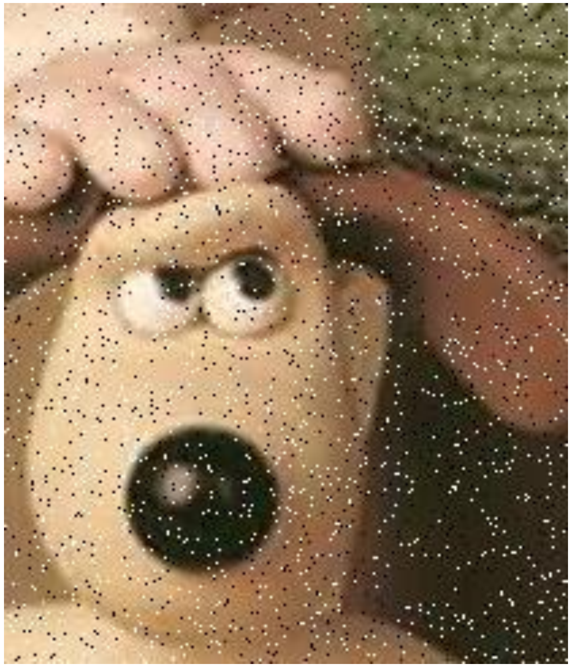

s&p noise

median filter

9. Noise Reduction

Nonlinear Filters: Min and Maxmin

+ shot noise **min filter** **maxmin filter**

Nonlinear Filters: Max and Minmax

- shot noise

max filter

min-max

10. Mathematical morphology

Nonlinear Processing: Binary Morphology

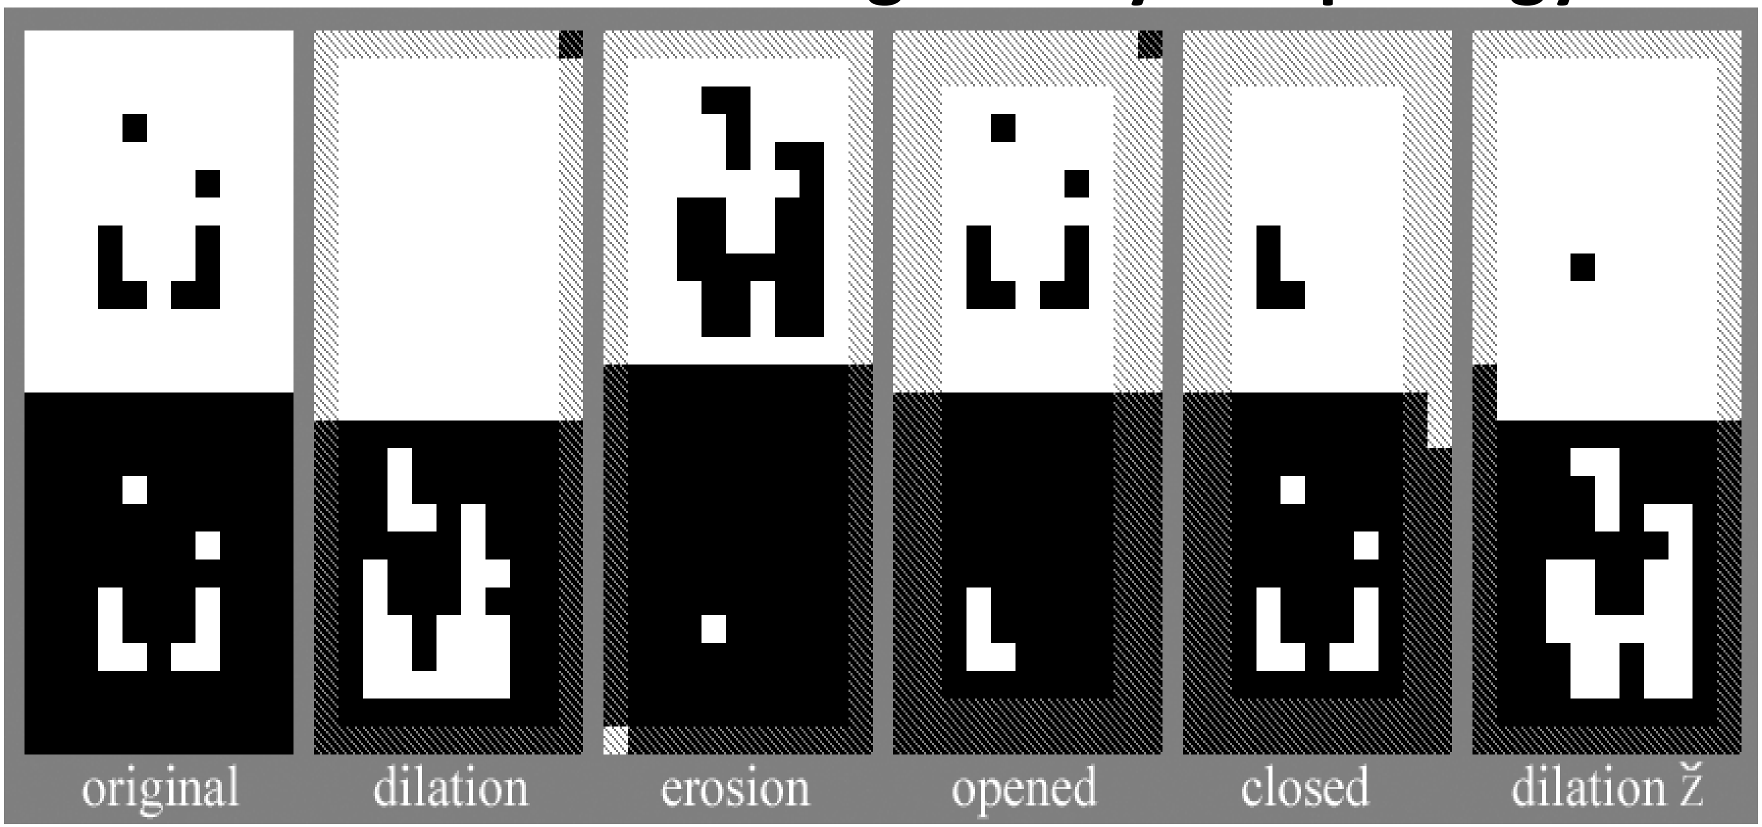

Foreground: white pixels
Background: black pixels

Cross-hatched pixels are indeterminate.

Nonlinear Processing: Binary Morphology

- Used after opening to grow back pieces of the original image that are connected to the opening.
- Permits the removal of small regions that are disjoint from larger objects without distorting the small features of the large objects.

original

opened

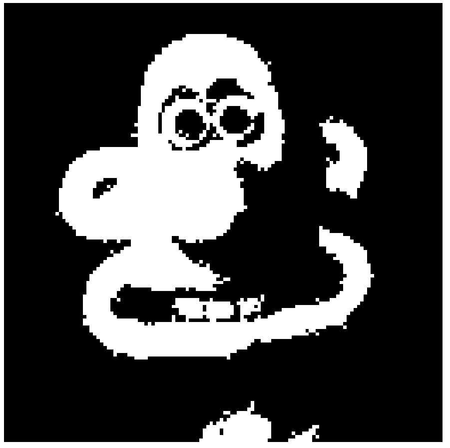

reconstructed

Nonlinear Processing: Grayscale Morphology

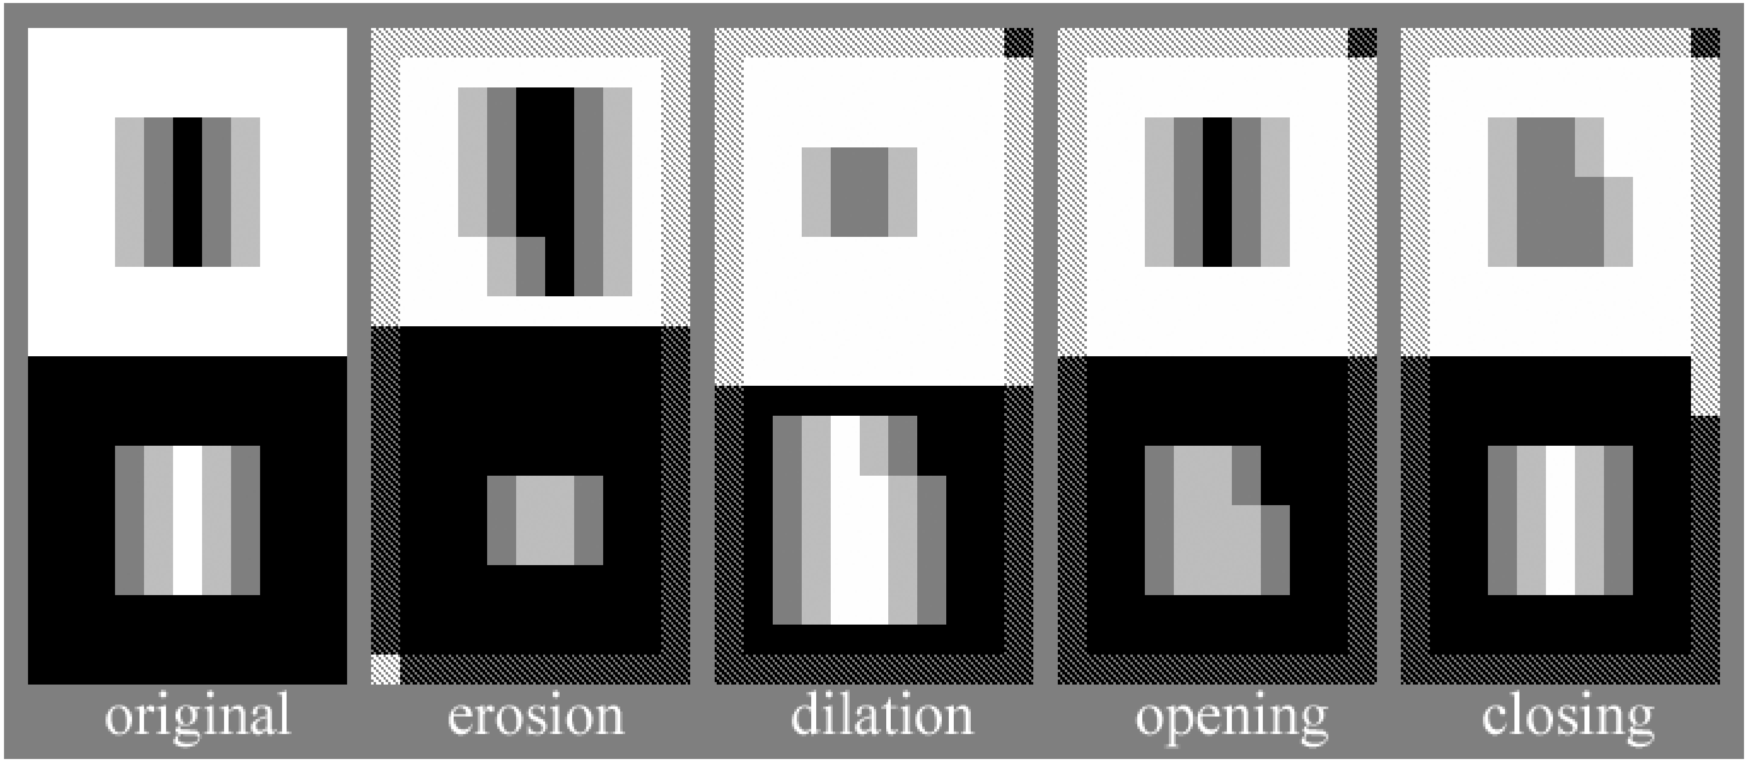

Foreground: white pixels
Background: black pixels

Cross-hatched pixels are indeterminate.

Grayscale Morphology: Opening

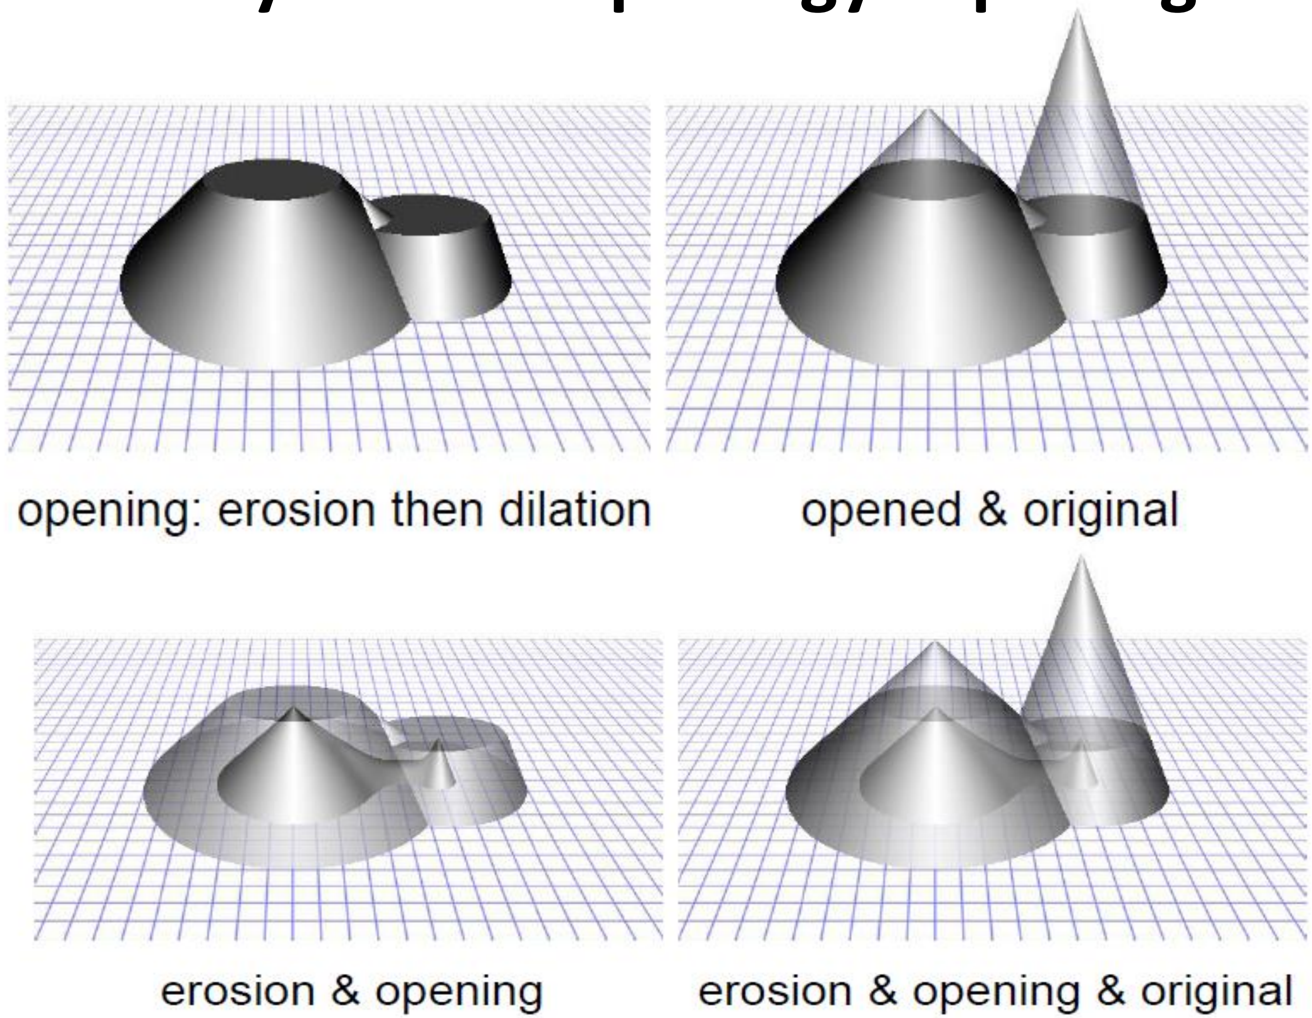

11. High Dynamic Range (HDR) Imaging

under exposed

11. High Dynamic Range (HDR) Imaging

default exposure

11. High Dynamic Range (HDR) Imaging

over exposed

11. High Dynamic Range (HDR) Imaging

combined

References

Textbooks

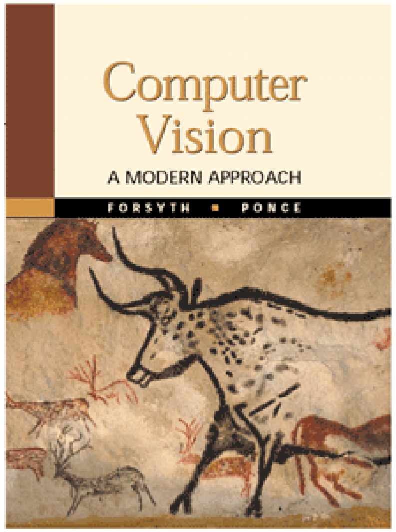

D. Forsyth, J. Ponce
Computer Vision – A Modern Approach
Prentice Hall, 2002

R. Hartley, A. Zisserman
Multiple View Geometry in Computer Vision
2nd Ed., Cambridge Univ. Press, 2004