⊕ ⊖ Computational ⊗ ⊘ Photography

Sensors

Jongmin Baek

CS 478 Lecture
Jan 18, 2012

Announcements

- Final Project: Description on the web.
- Piazza: Sign up at piazza.com
    - Ask questions here.
- Office hours: Tuesday slot removed.
    - Now W3:45-5, Th2:30-3:45

Tegra 3 Tablets

- Pick up: Friday afternoon at Rm. 360.

Looking Ahead

- Next week
  - Monday
    - Lecture on FCam (Reading assigned!)
    - “Hello Camera” Assignment out
  - Wednesday
    - Lecture on camera control algo.

Outline

- Background material, Part II
  - Perception
  - Sensor
  - Noise

Lot of slides stolen from Marc Levoy

Outline

- Background material, Part II
  - Perception
  - Sensor
  - Noise

Visible Light

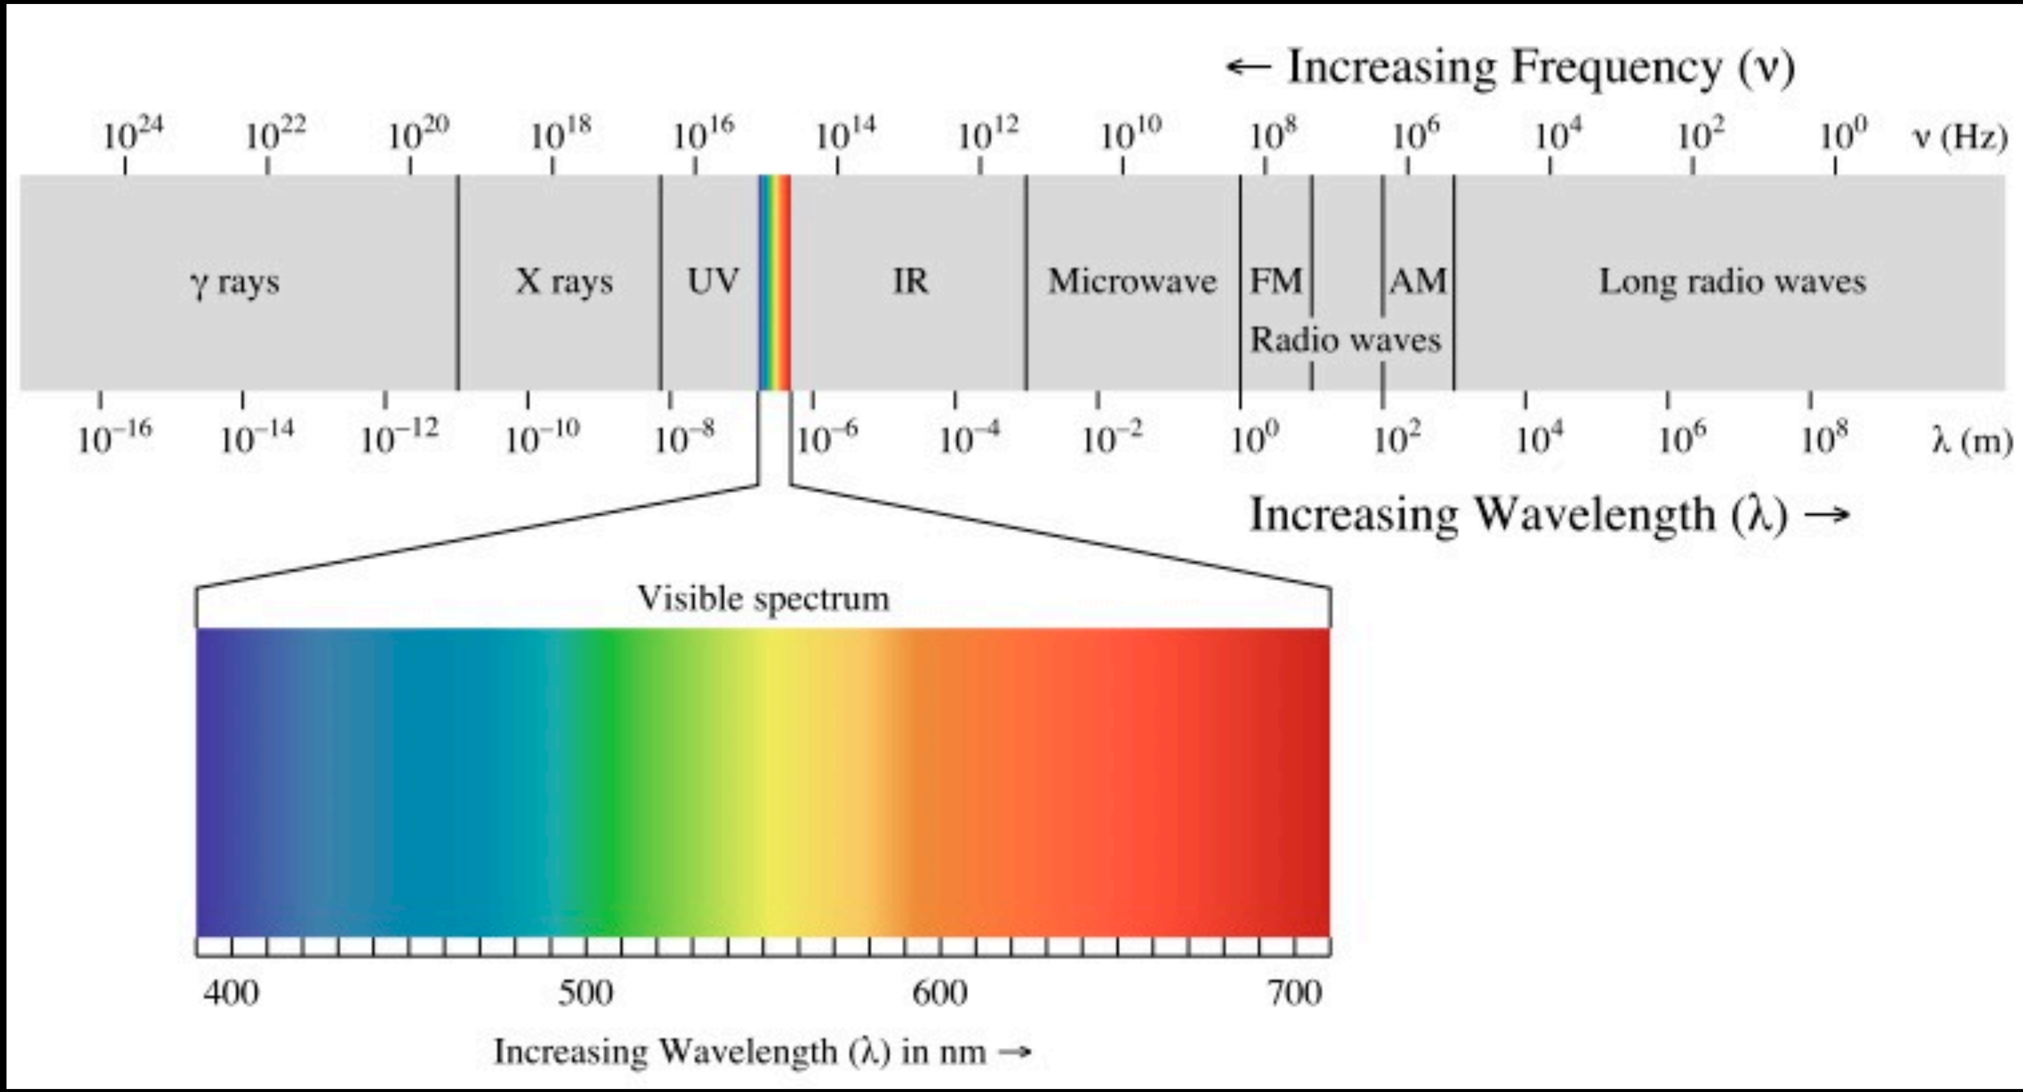

(wikipedia)

- wavelengths between 400nm and 700nm

Illumination

(LampTech)

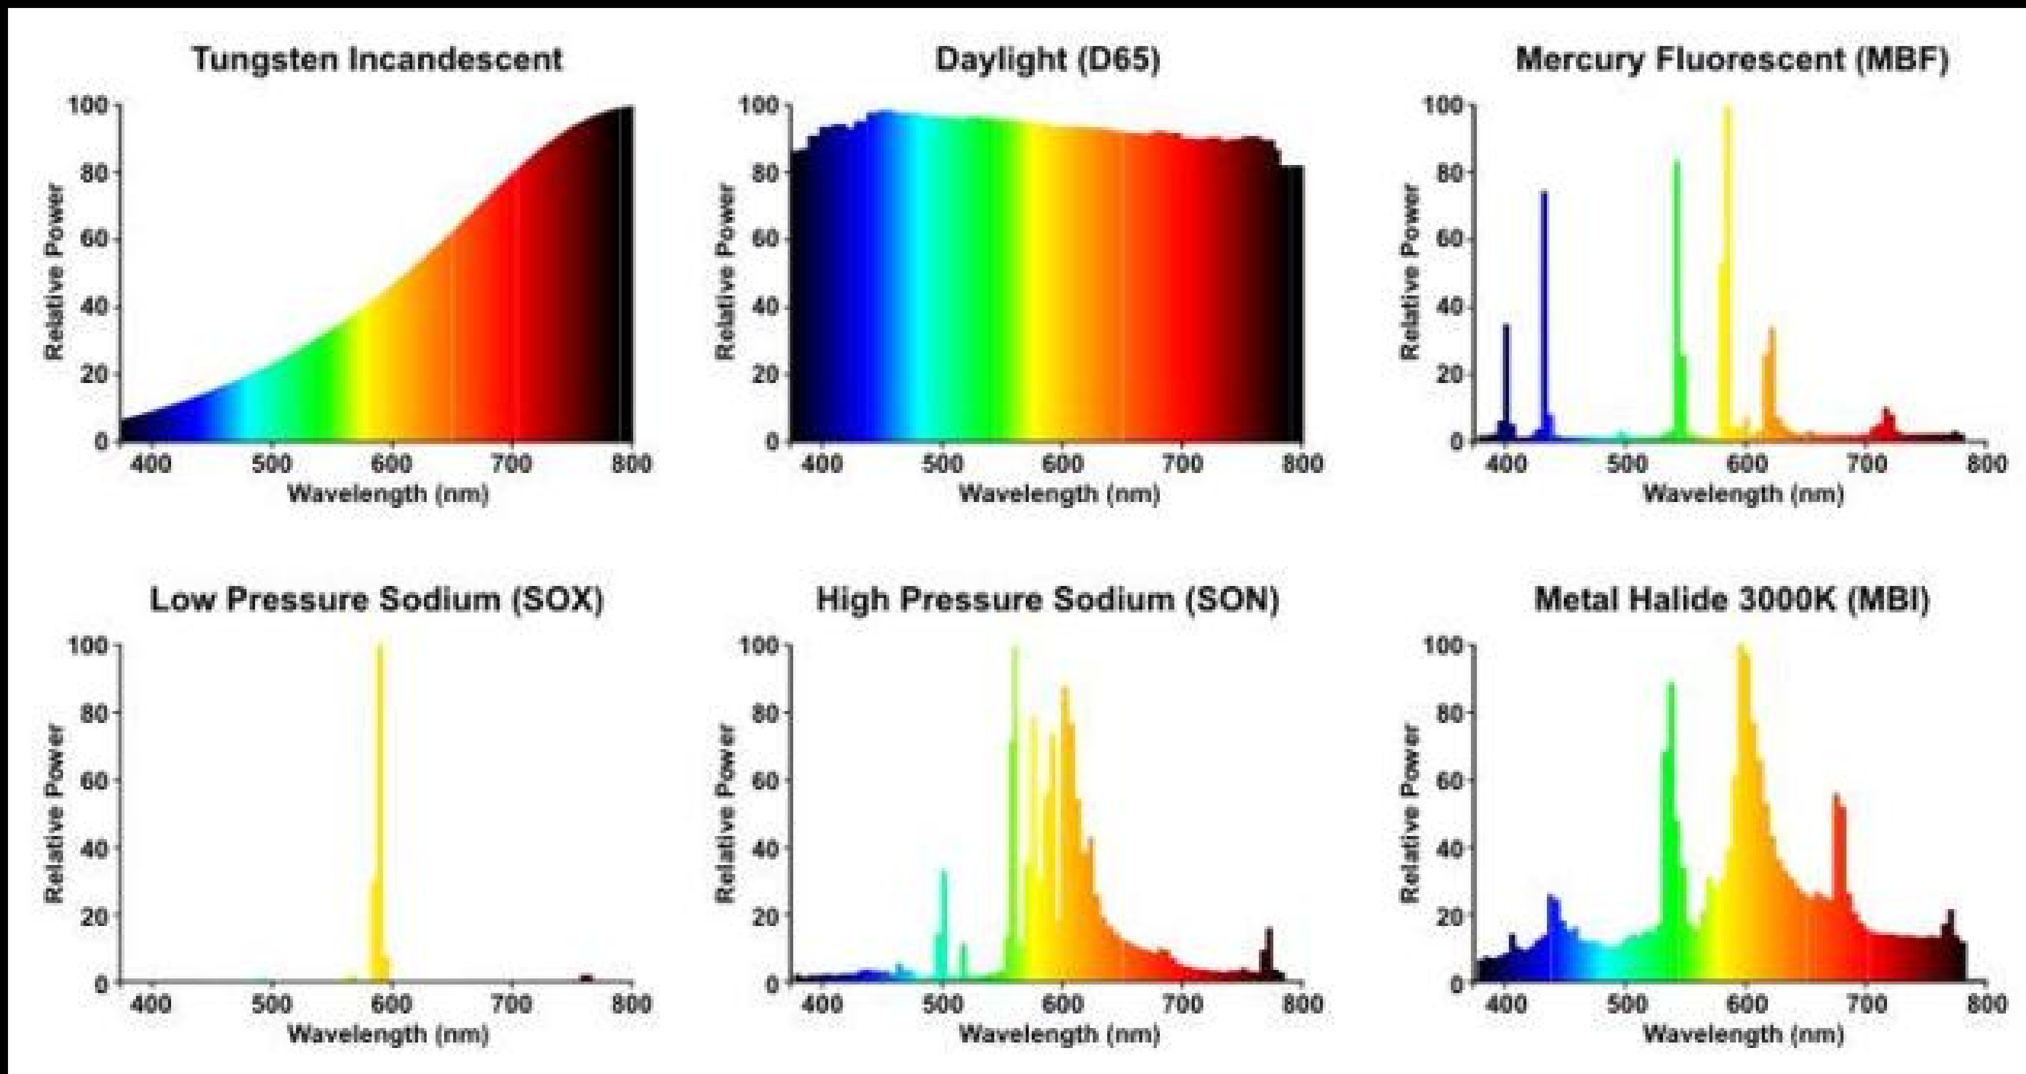

- wavelengths between 400nm and 700nm

Light Transport

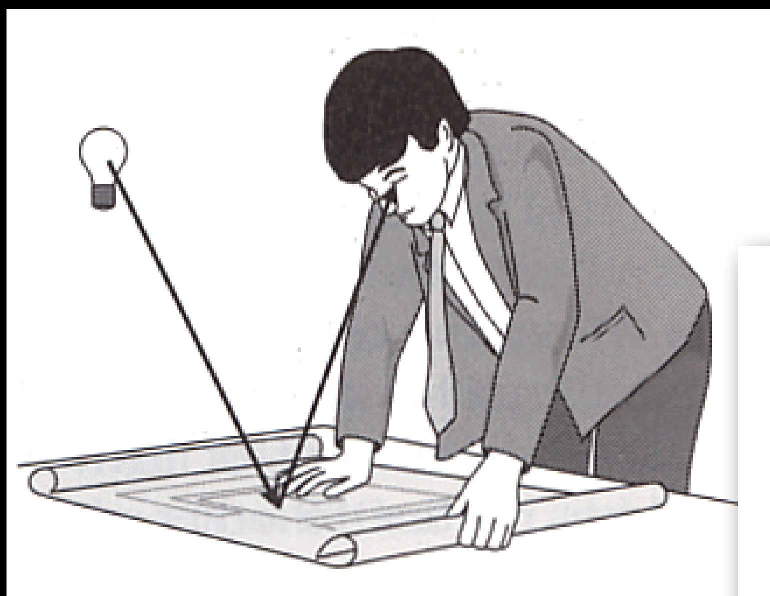

light is reflected by an object

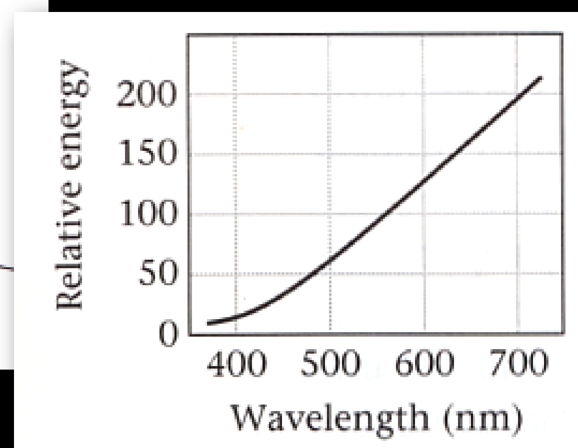

illumination
(Light)

×

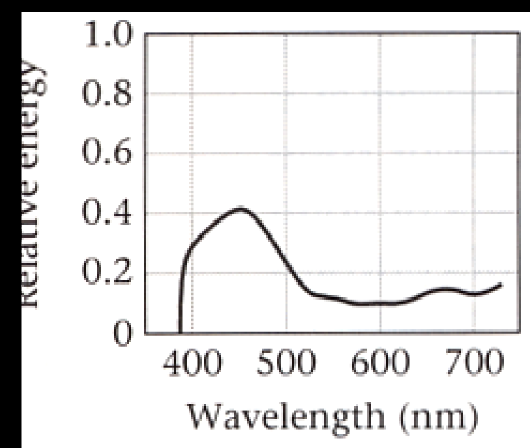

reflectance
(Object)

=

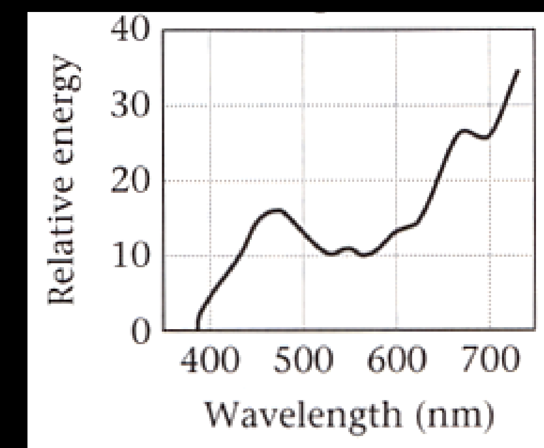

stimulus that enters your eye

- One cause for “metamers”

Trichromatic Vision

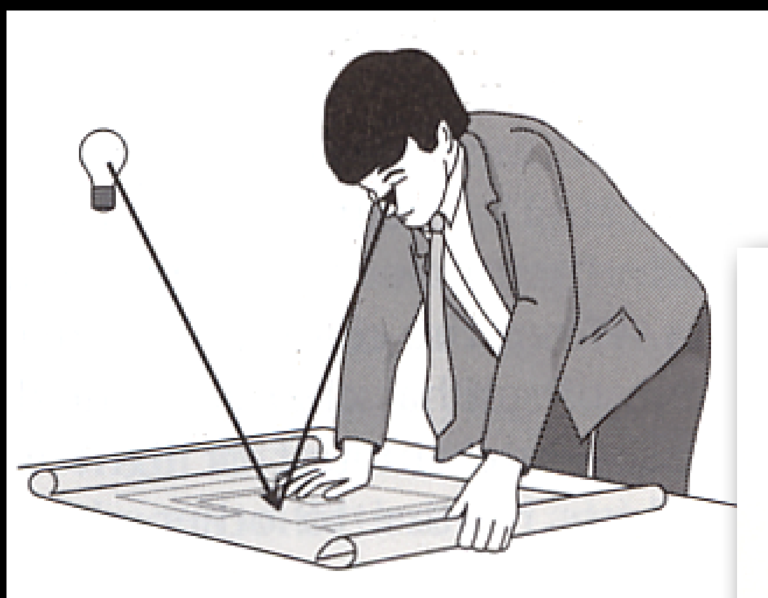

light is reflected by an object

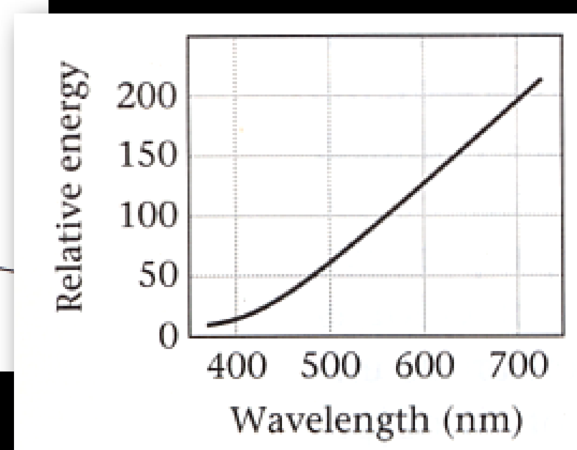

illumination (Light)

×

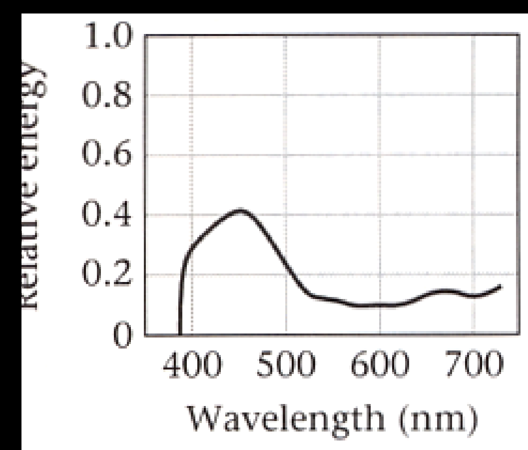

reflectance (Object)

=

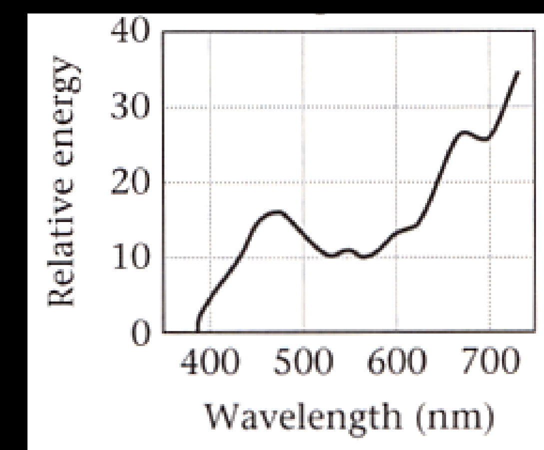

stimulus that enters your eye

Color Gamut

β

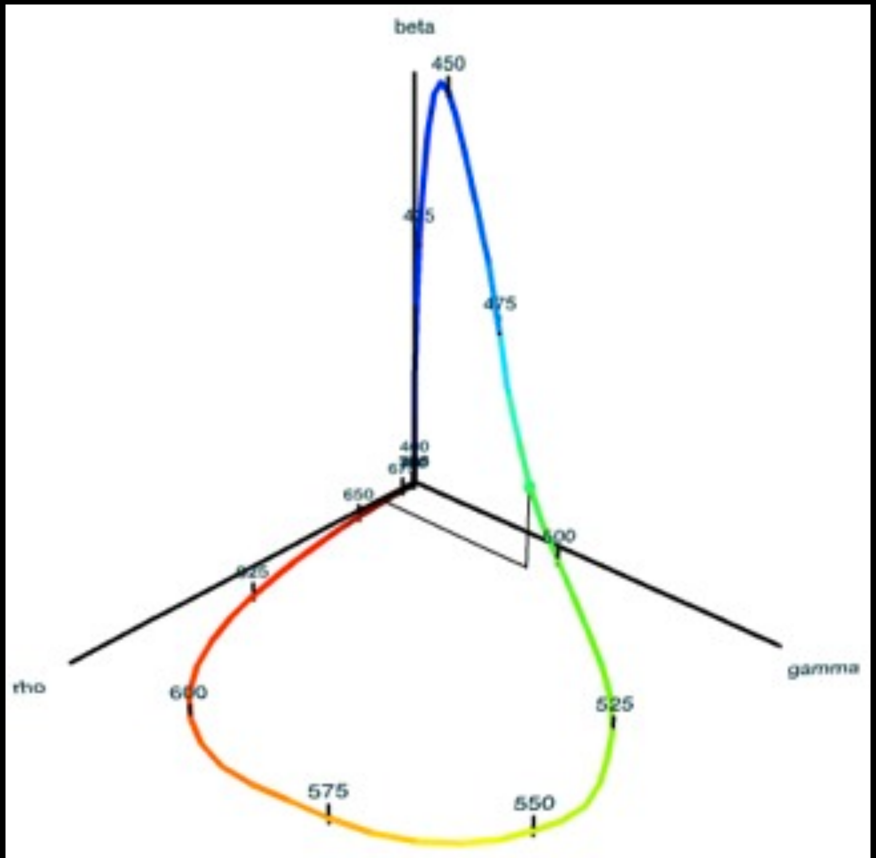

ρ γ

spectral locus

- Set of perceivable colors (distinct stimuli)
  - Convex hull of the responses for pure wavelengths.

Color Gamut: Consequences

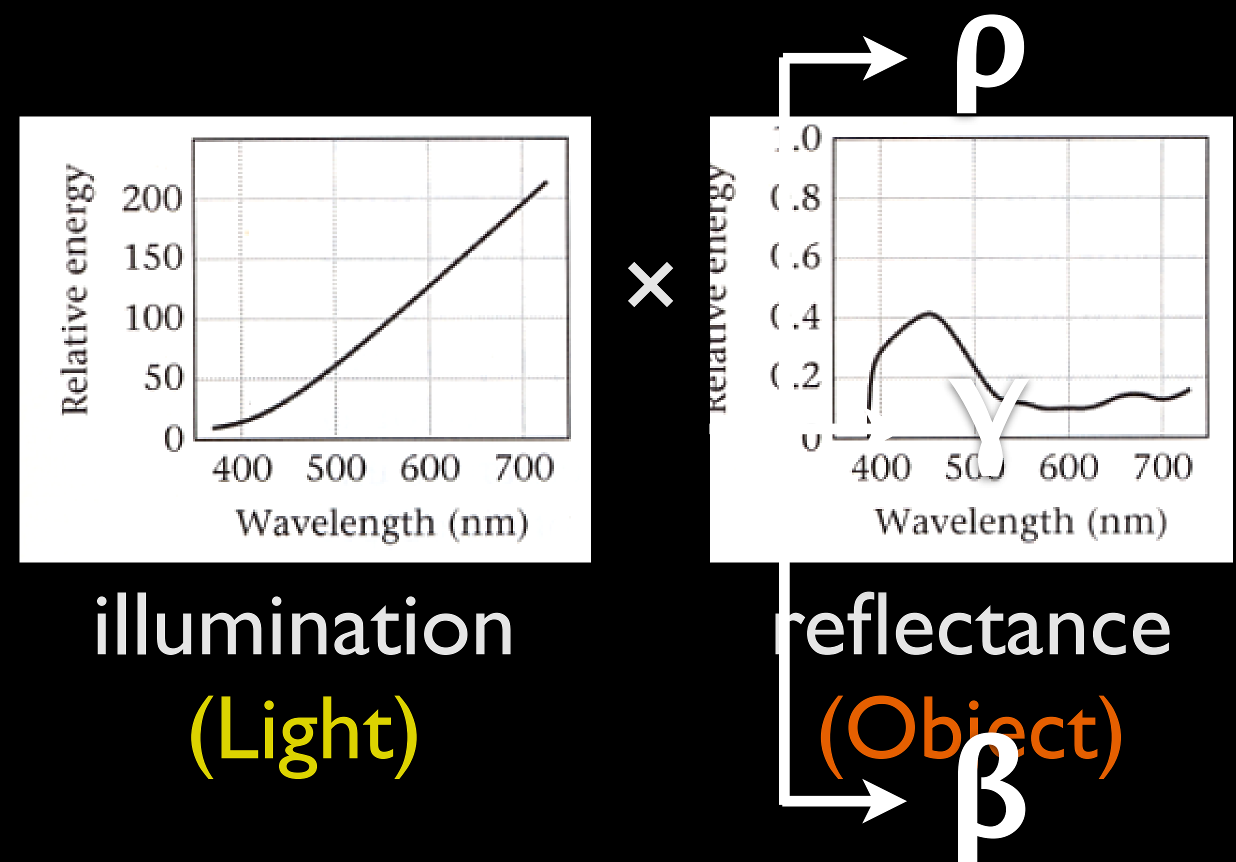

- Goal of photography: reproduce the sensation of seeing a scene.

Two Questions

- Given a point in the scene, how do we calculate the appropriate (ρ, γ, β)?
- Given (ρ, γ, β), how do we recreate the sensation in the viewer?

Color Primaries

- Given (ρ, γ, β), how do we recreate the sensation in the viewer?
  - Want to display a spectrum that would generate the desired (ρ, γ, β) in the viewer's eyes.
  - Pure wavelength is hard to isolate.
  - Instead, use primary colors.

Color Primaries

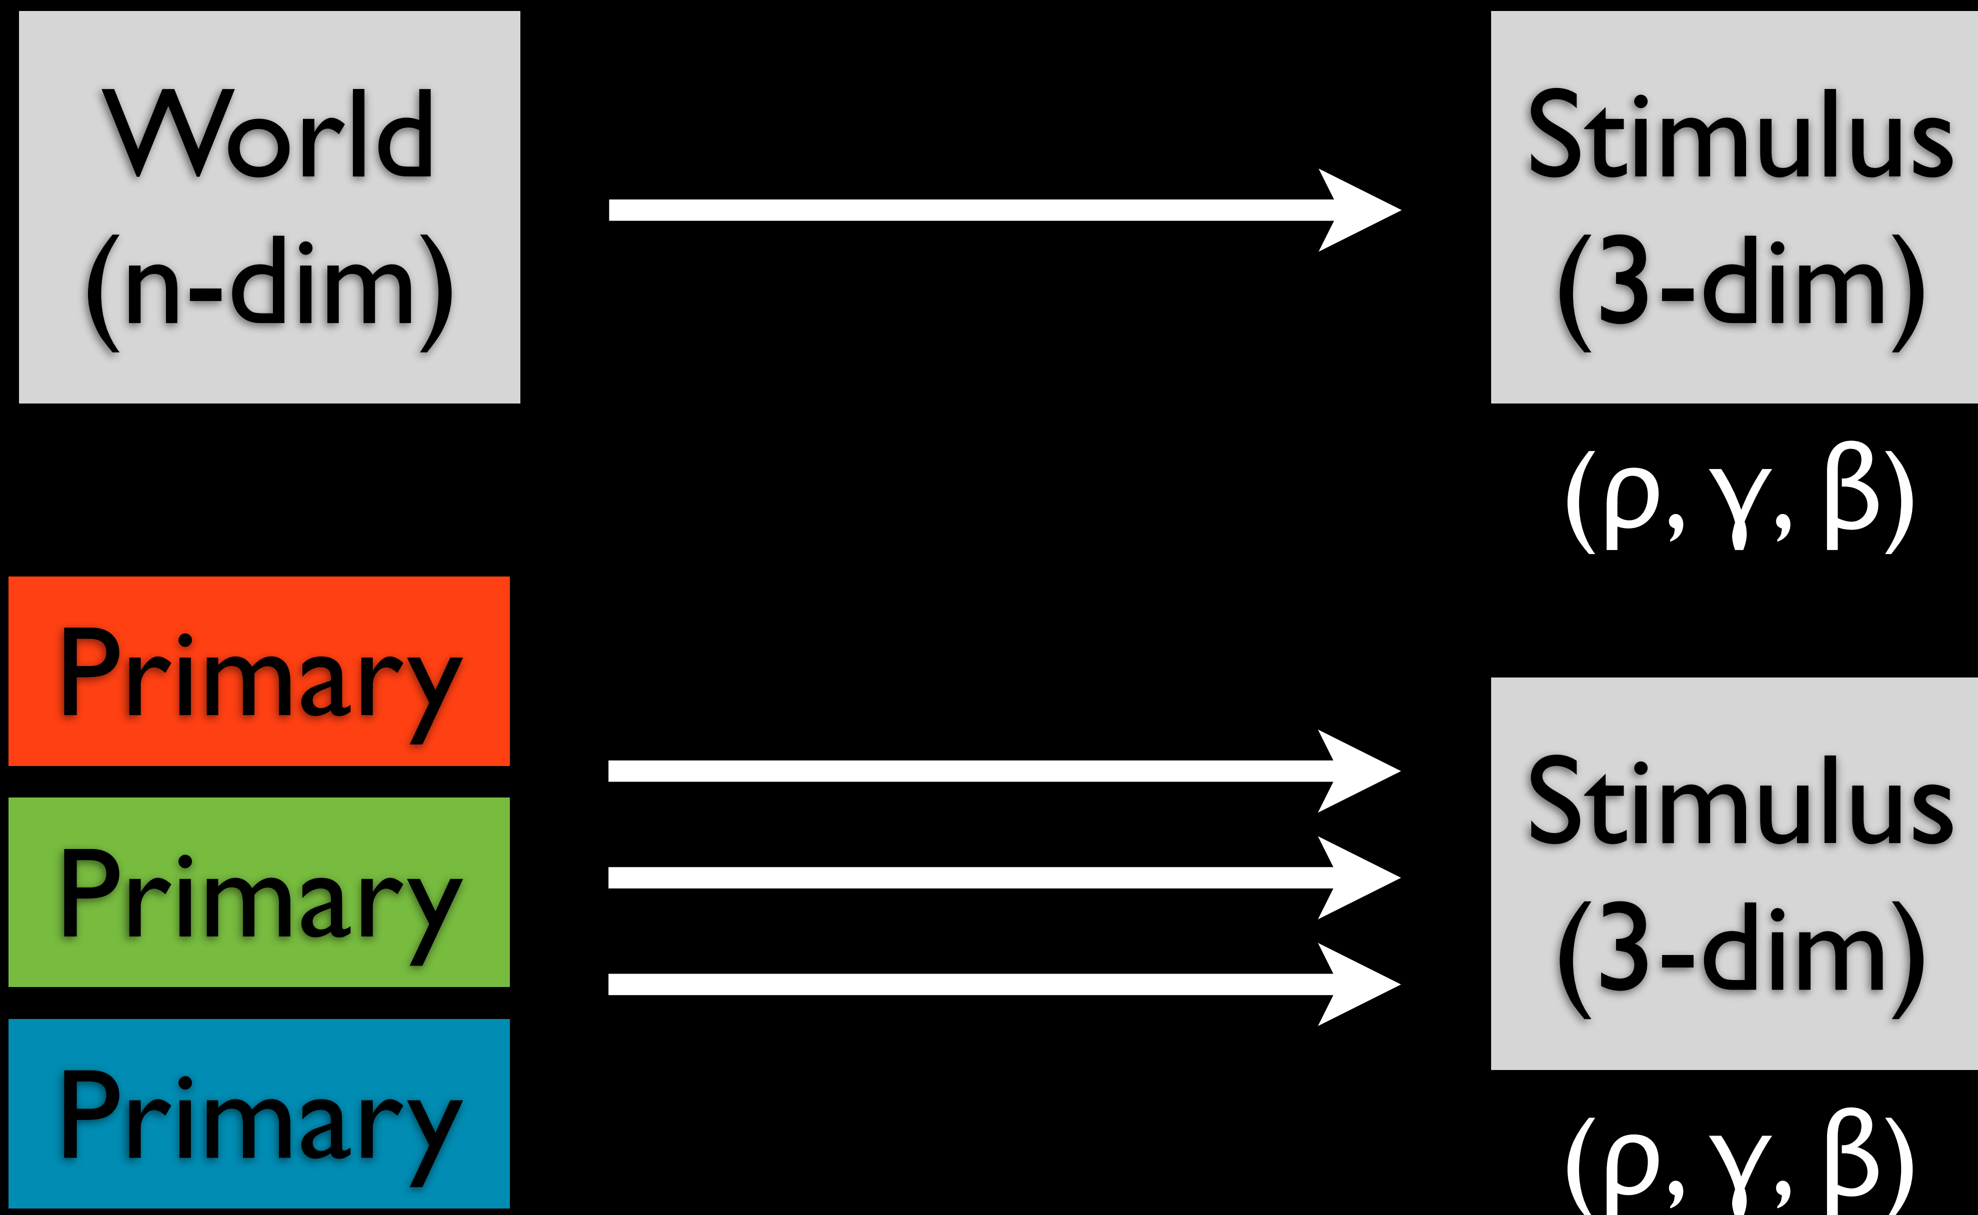

Color Primaries

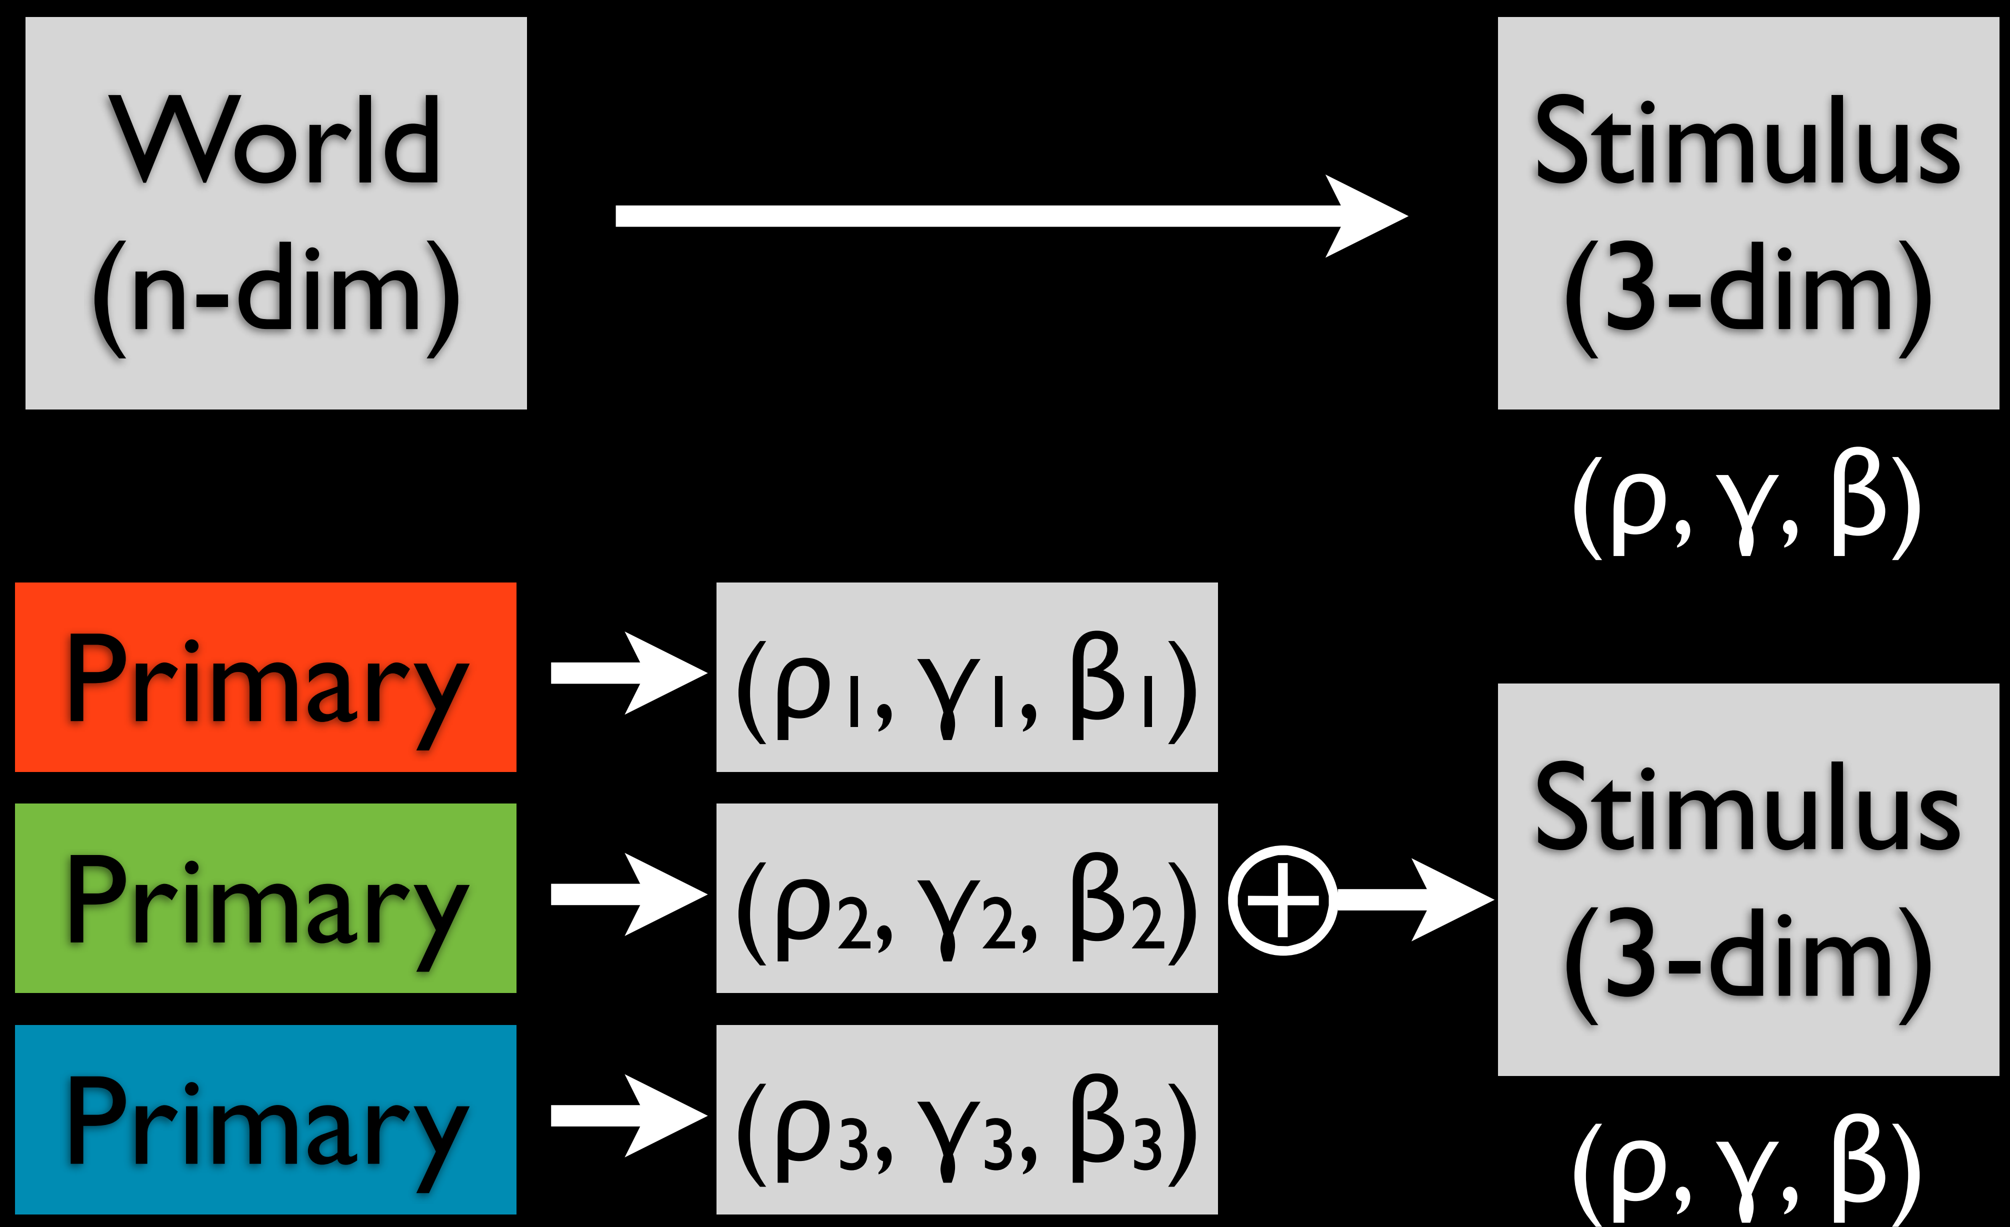

Wednesday, January 18, 12

Color Primaries

- Choose three primaries R, G, B.
    - Does not have to be pure wavelengths.

- Normalize to obtain a desired *reference white*
    - This yields an *RGB cube*

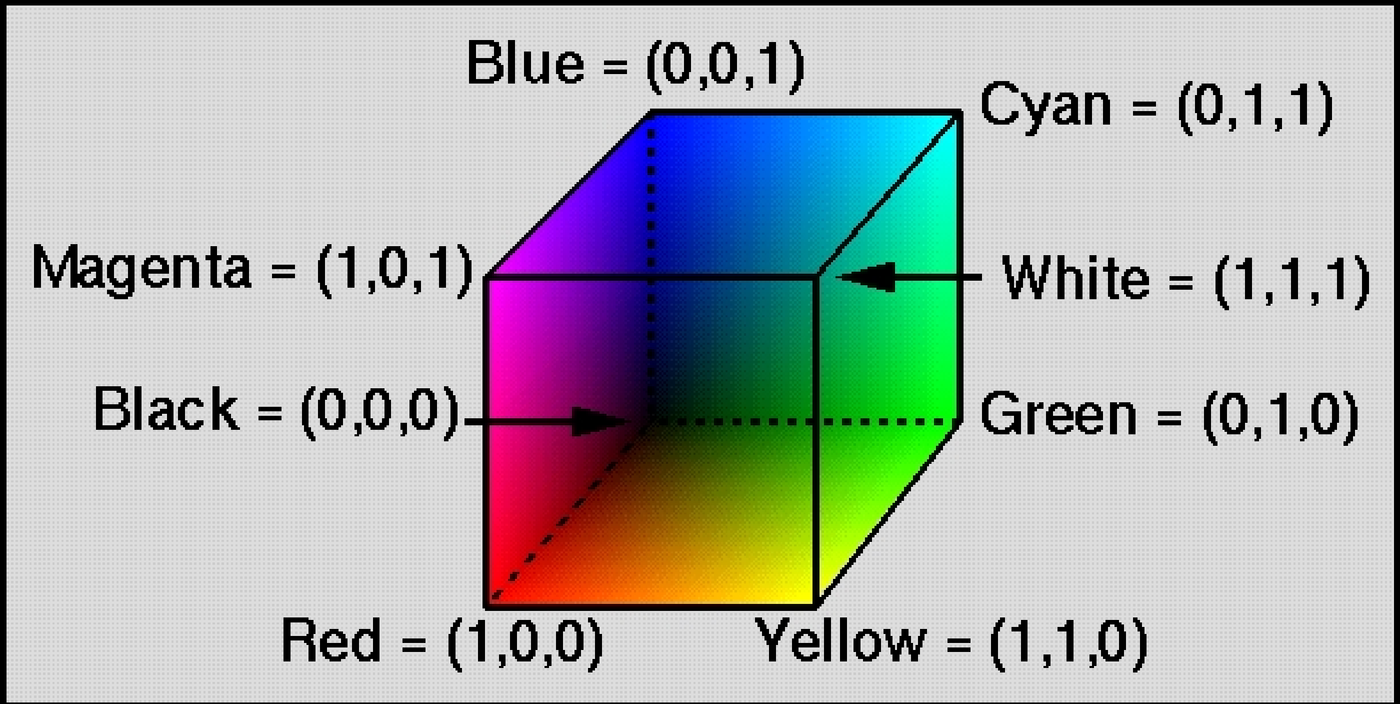

Color Primaries

- What exactly is R, G, B each?
  - Is there a specfic wavelength for each? No.
  - Is there a specific spectrum for each? Yes, but you can pick your own.

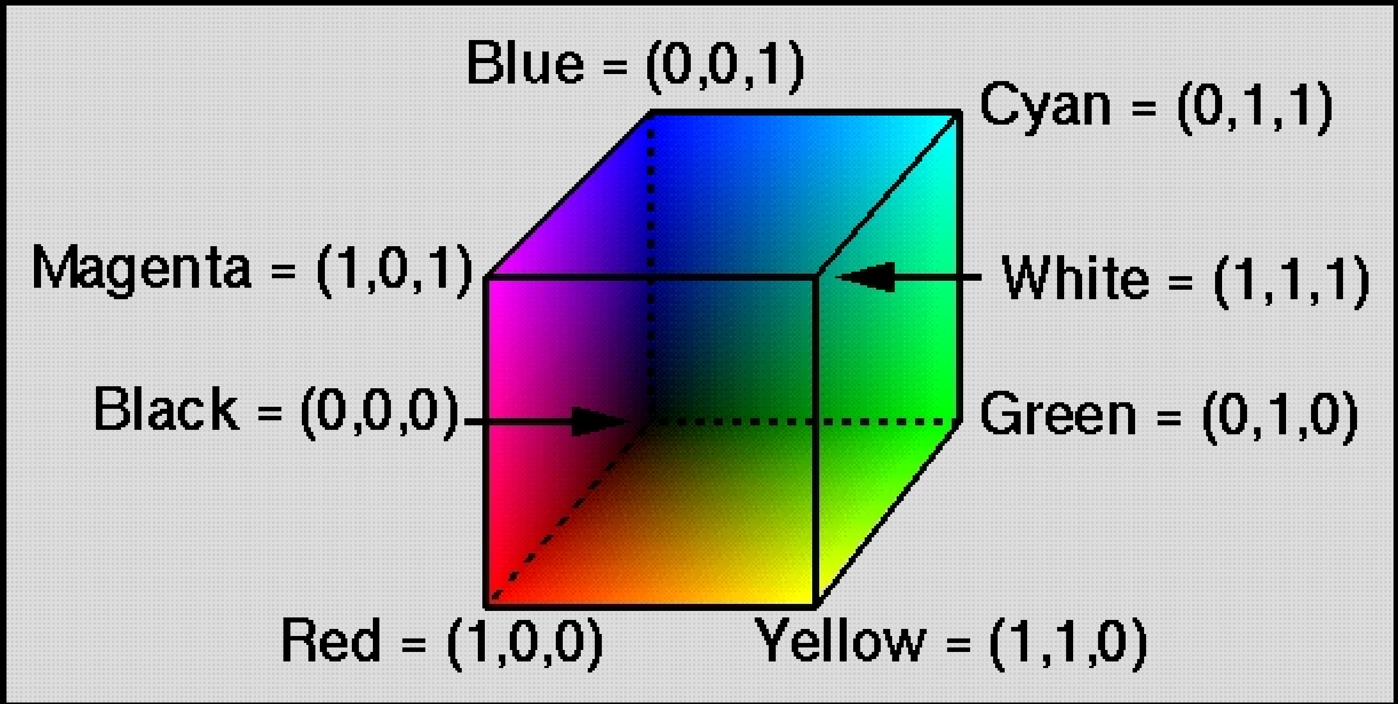

Choice of Primaries

- sRGB (HP, Microsoft, 1996)
- Adobe RGB (Adobe, 1998)
- Adobe Wide-Gamut RGB
- ...

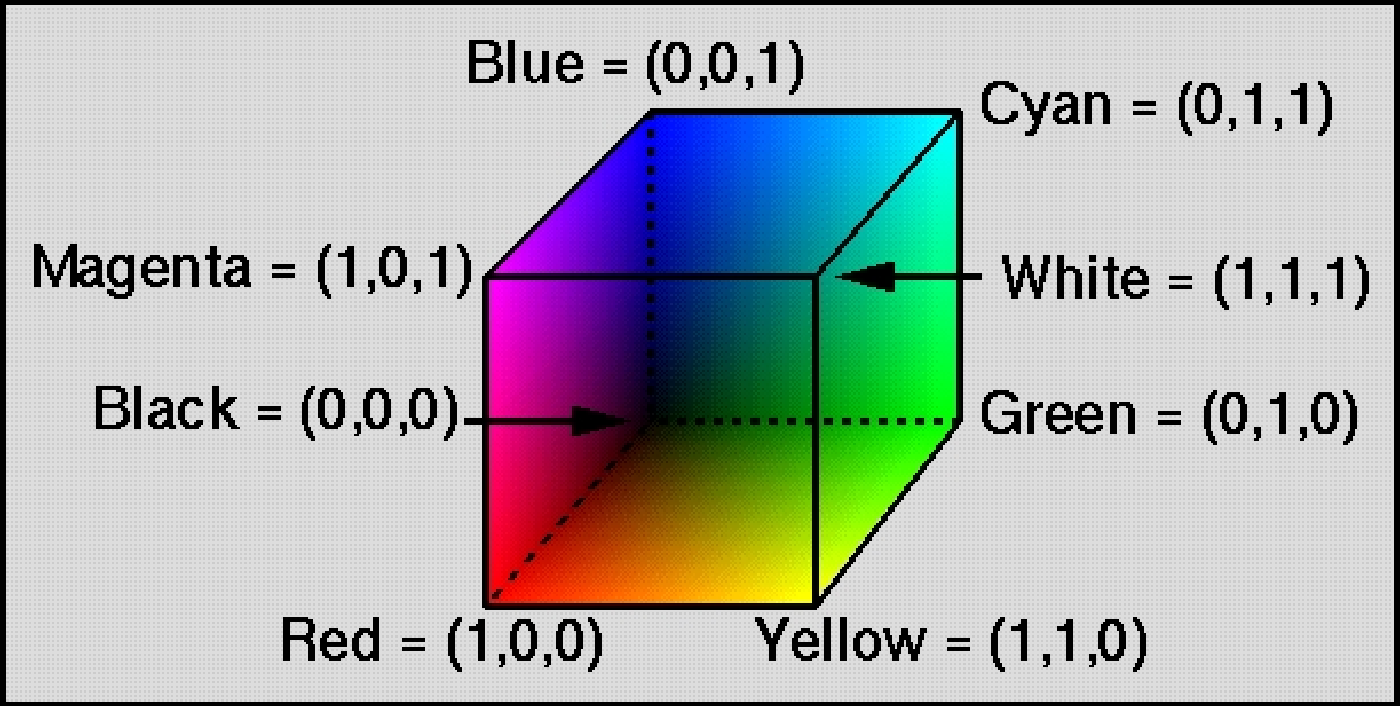

Chromaticity Diagram

- The color gamut diagram can be compressed into 2D by homogenizing the coordinates.
- Plot the primaries.
  - The convex hull is the extent of reproducible sensations.

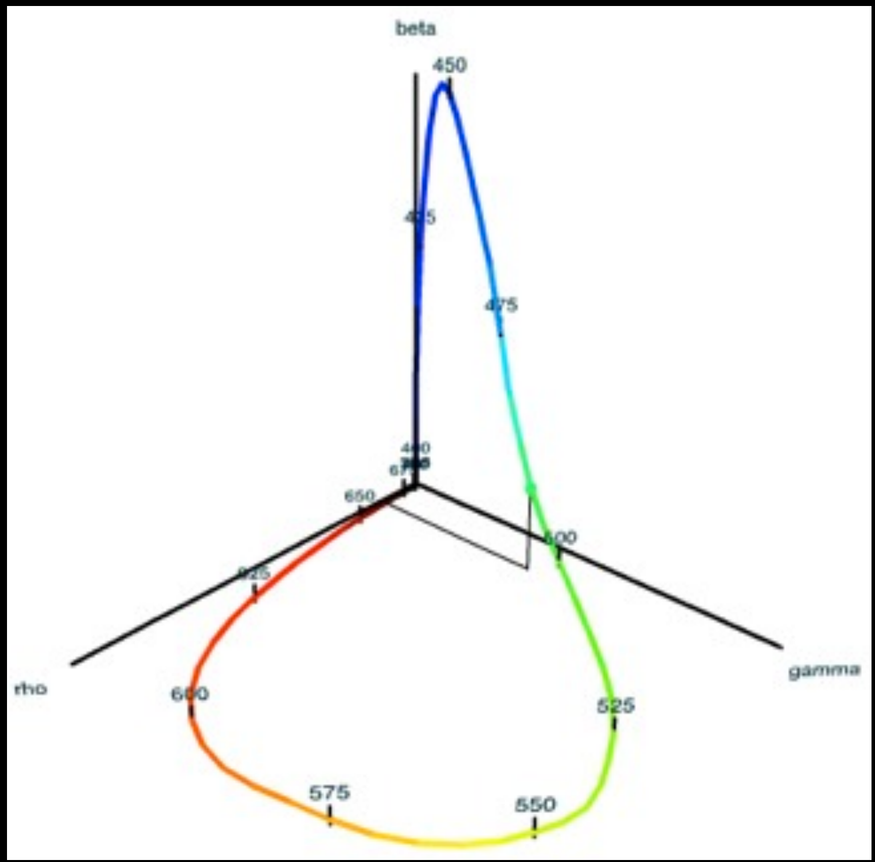

spectral locus

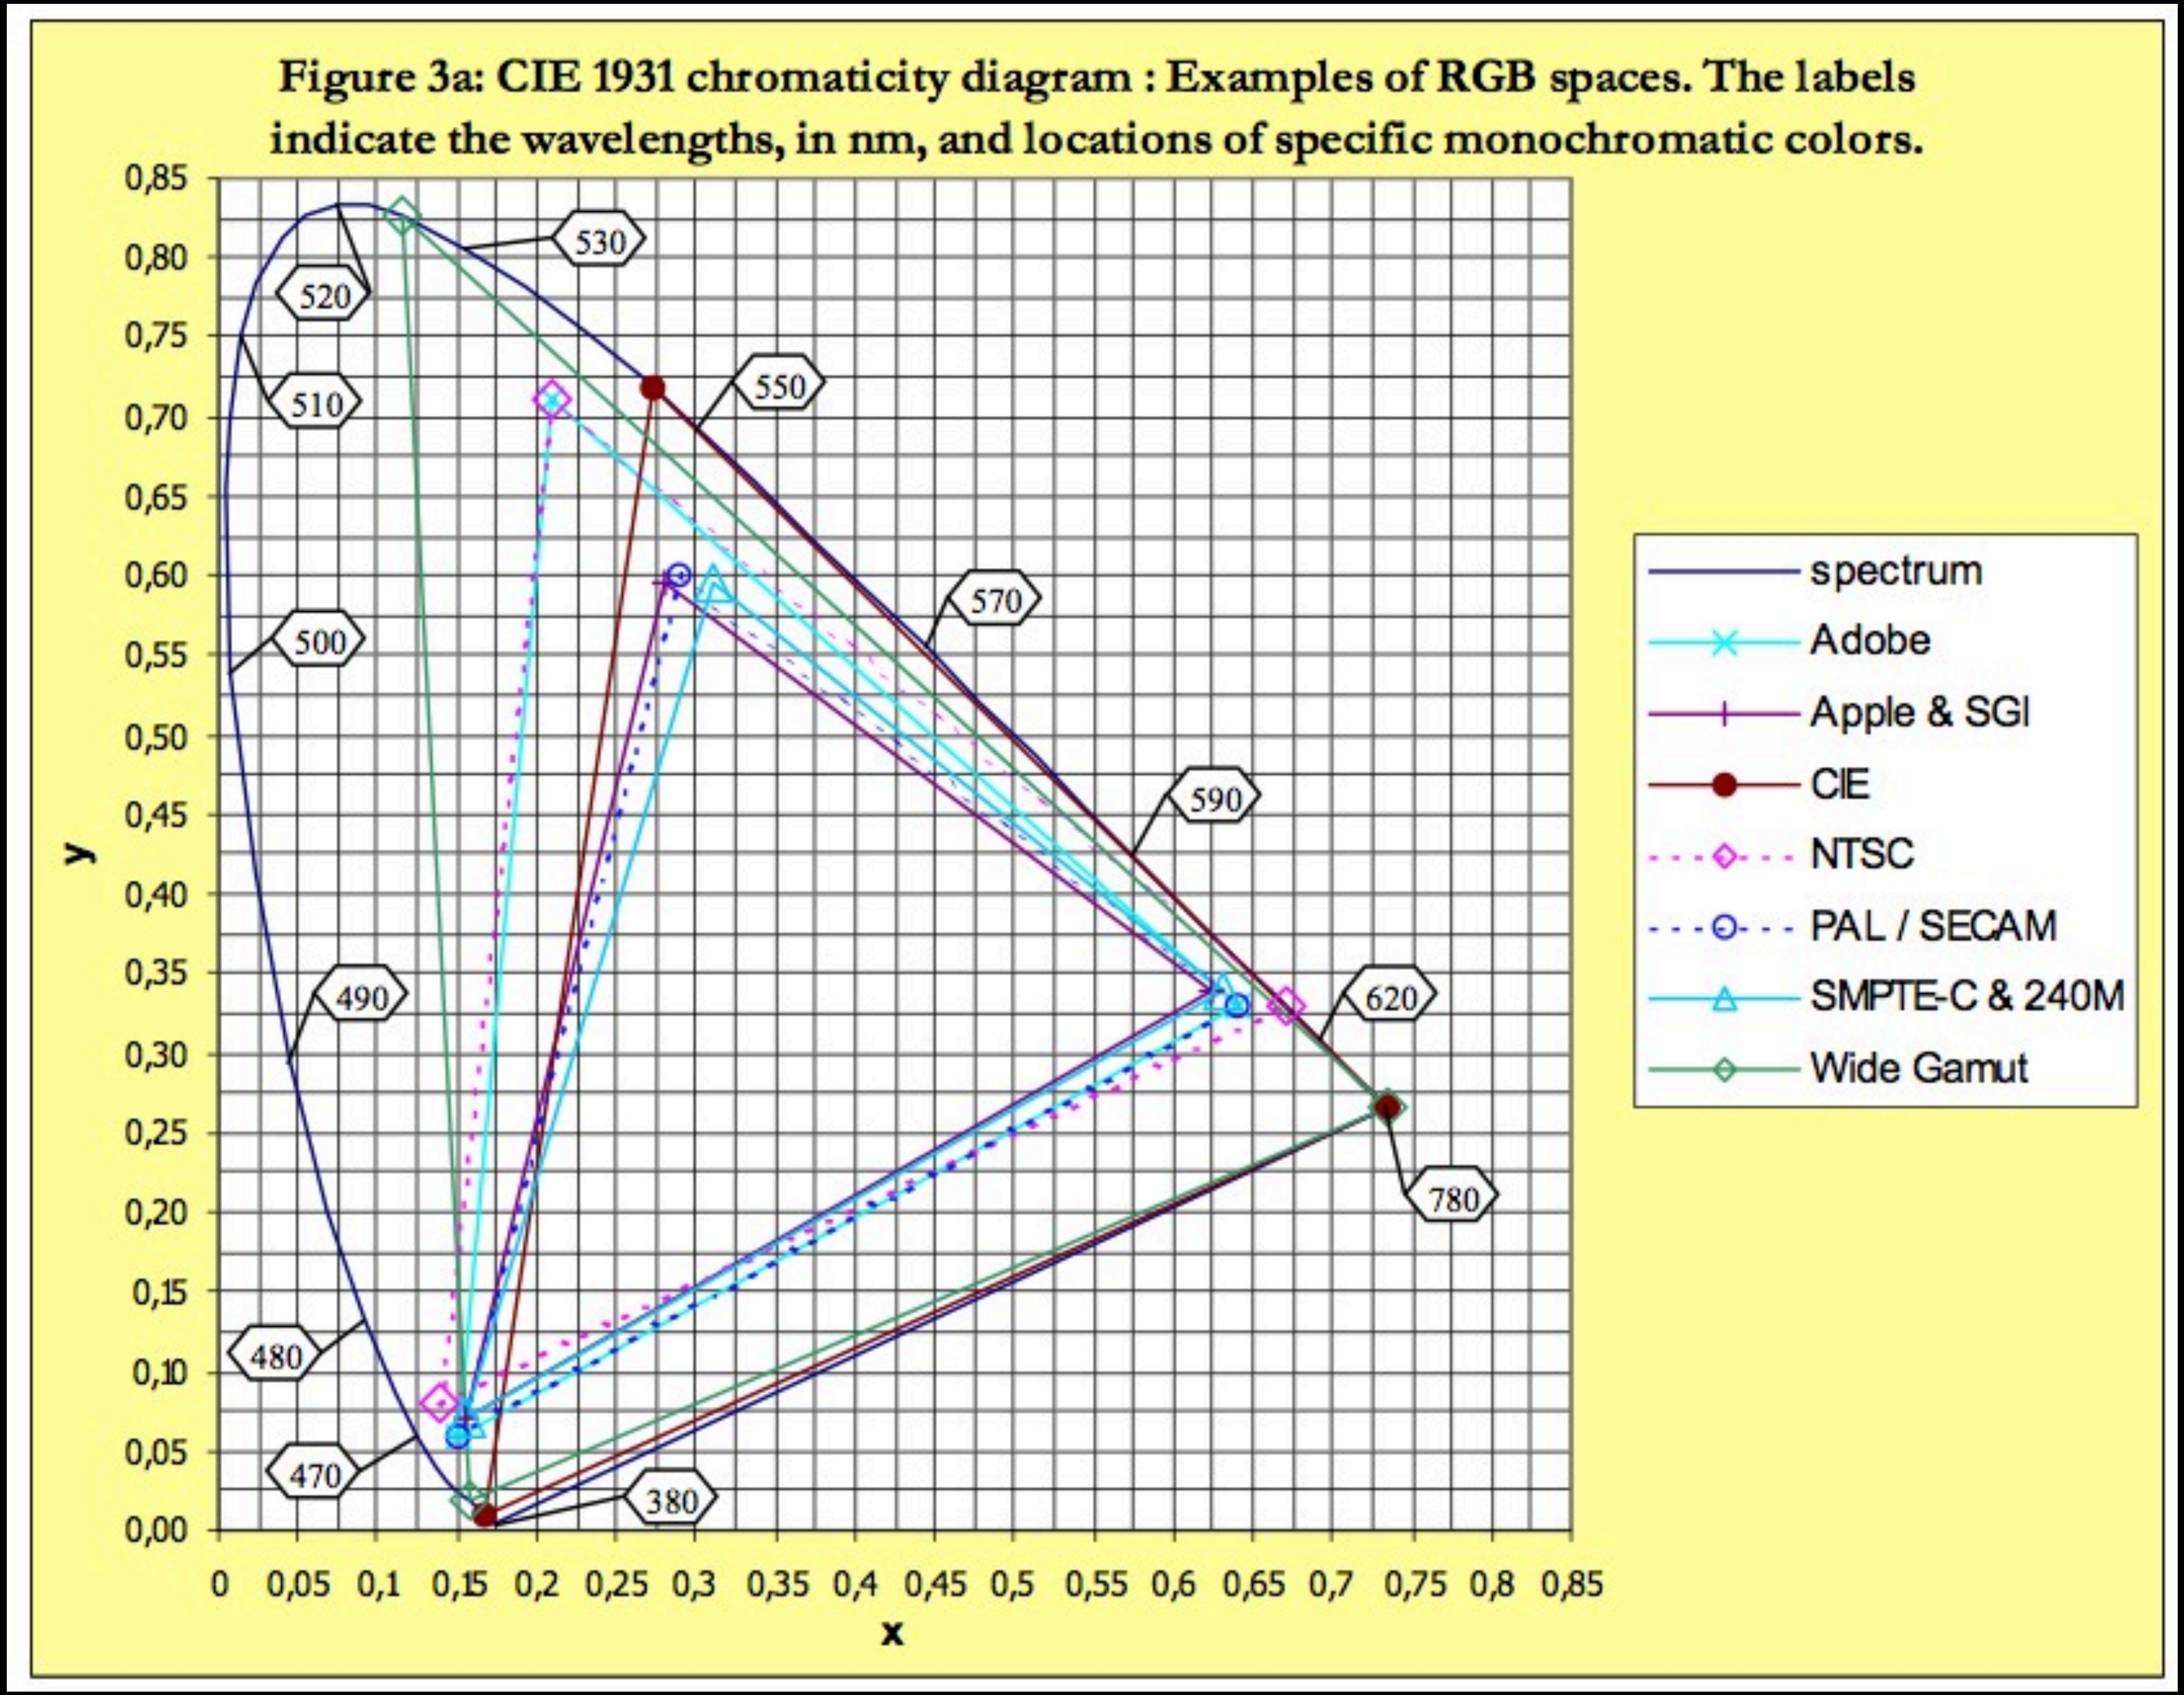
Figure 3a: CIE 1931 chromaticity diagram : Examples of RGB spaces. The labels indicate the wavelengths, in nm, and locations of specific monochromatic colors.
530
520
510
550
570
500
590
490
620
780
480
470
380
y
x
0,85
0,80
0,75
0,70
0,65
0,60
0,55
0,50
0,45
0,40
0,35
0,30
0,25
0,20
0,15
0,10
0,05
0,00
0 0,05 0,1 0,15 0,2 0,25 0,3 0,35 0,4 0,45 0,5 0,55 0,6 0,65 0,7 0,75 0,8 0,85
spectrum
Adobe
Apple & SGI
CIE
NTSC
PAL / SECAM
SMPTE-C & 240M
Wide Gamut

Color Models

- RGB (R, G, B)
- YUV (Y, U/U, V/V) or YCbCr
  - Y: Luminance; U,V: Chrominance
  - Used in video processing(incl. Tegra 3), JPEG
- HSV
  - H: Hue; S: Saturation; V: Value
- CMYK

Perceptual Color Model

- XYZ:
  - X,Y,Z primaries correspond to a spectral sensitivity of the three cones.
  - Unfortunately, the values are not perceptually spaced.
    - e.g. the difference between X=1 and X=2, and the difference between X=2 and X=3 are not equivalent.

Perceptual Color Model

- (CIE-)LAB:
  - Meant to be a perceptually correct metric.
  - $L^* = 116(Y/Y_w)^{1/3}-16$, for $Y/Y_w>0.008856$,
    $903.3\ (Y/Y_w)^{1/3}$, otherwise.
  - $a^* = 500((X/X_w)^{1/3}-(Y/Y_w)^{1/3})$
  - $b^* = 200((Y/Y_w)^{1/3}-(Z/Z_w)^{1/3})$

Perceptual Color Model

- (CIE-)LAB:
  - How do you convert from L*a*b* to RGB?
    - There's no fixed formula. It depends on the RGB primaries.
      - e.g. LAB-sRGB

Outline

- Background material, Part II
  - Perception
  - Sensor
  - Noise

Camera Pipeline

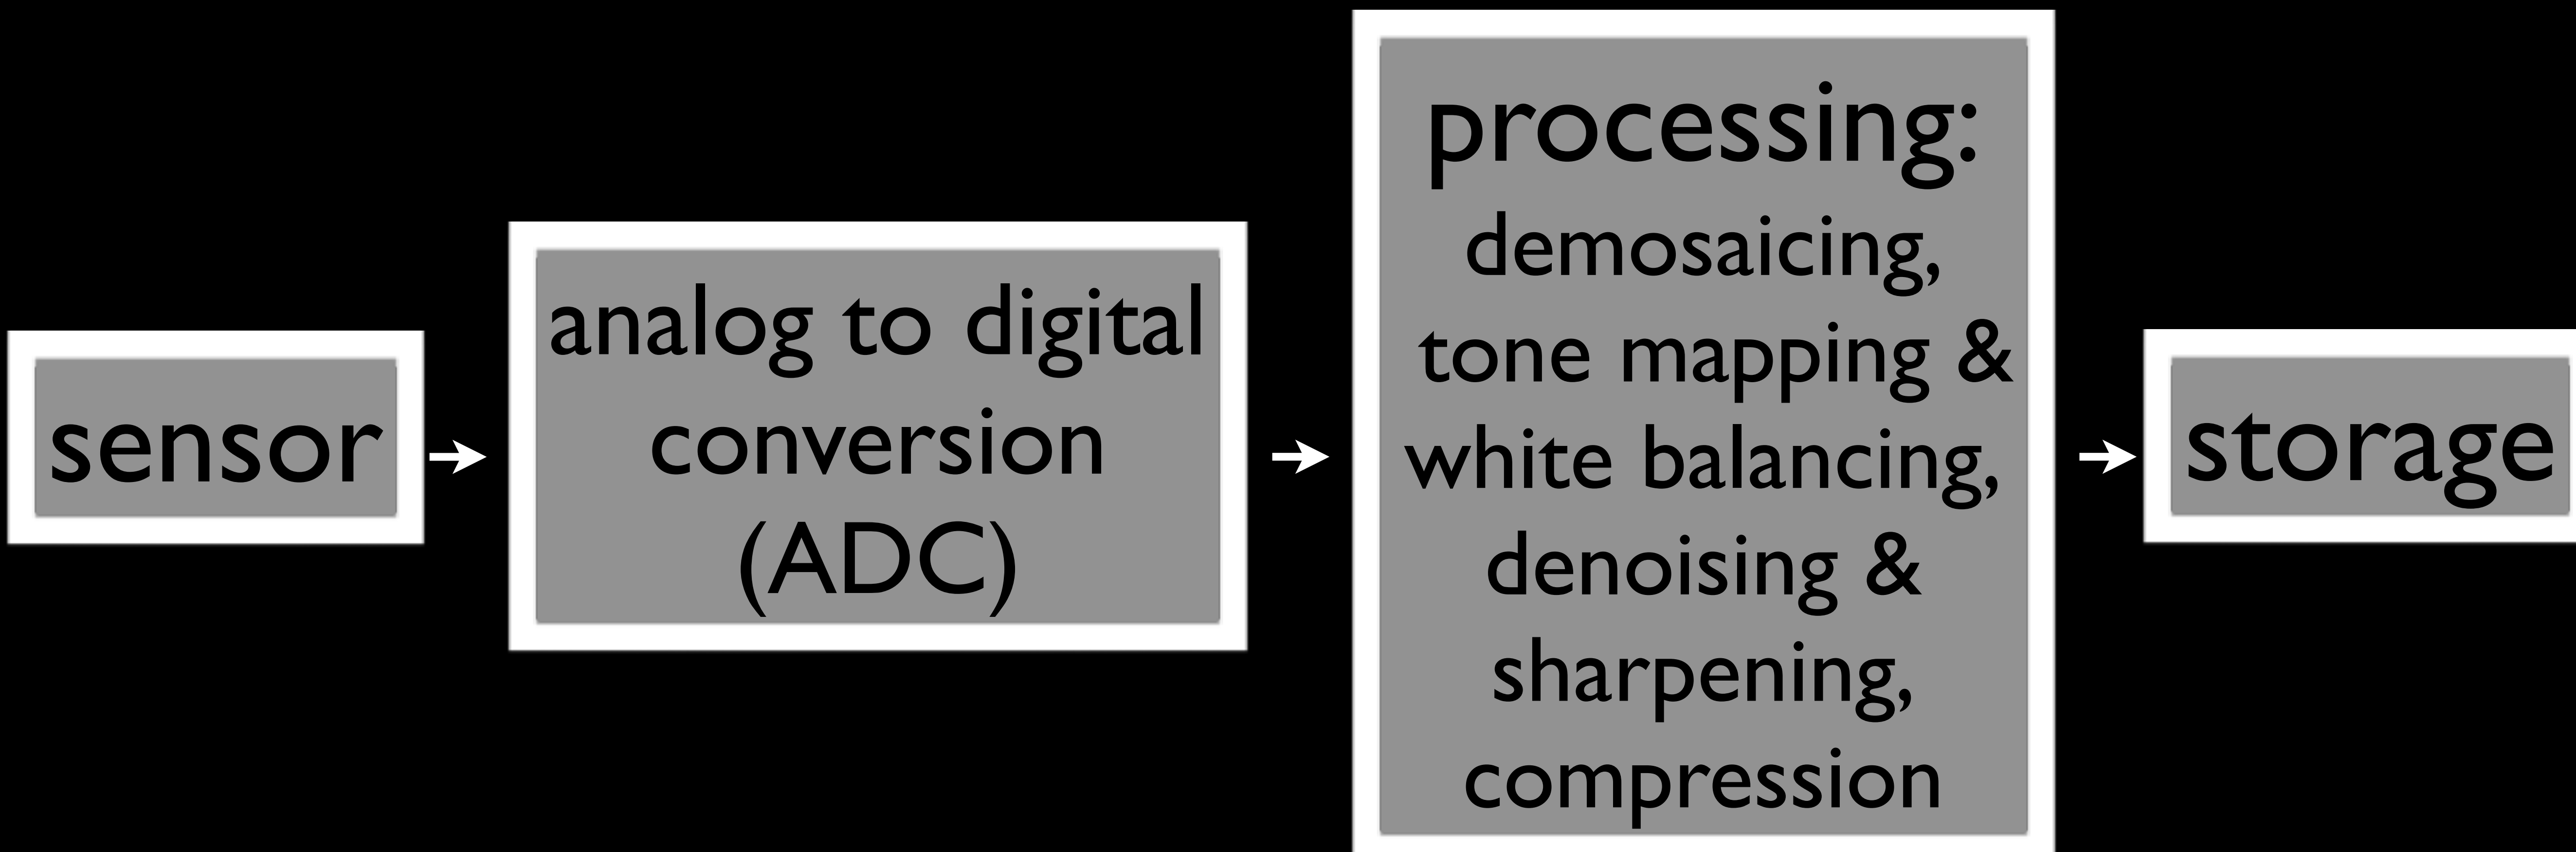

Example Pipeline

sensor → analog to digital conversion (ADC) → processing: demosaicing, tone mapping & white balancing, denoising & sharpening, compression → storage

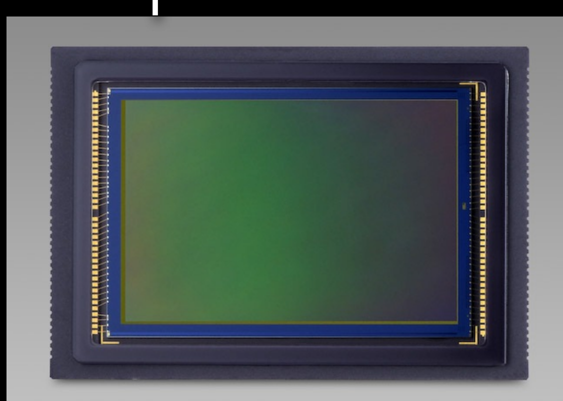

Canon 21 Mpix CMOS sensor

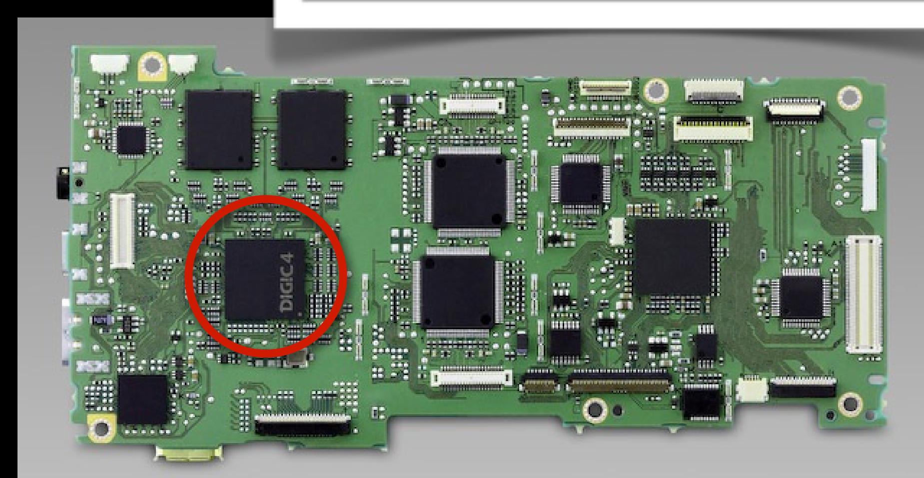

Canon DIGIC 4 processor

Compact Flash card

The Science

- Photoelectric Effect
  - Materials may generate electrons upon being hit by a photon.
- Quantum Efficiency
  - Not all photons will produce an electron.

back-illuminated CMOS (Sony)

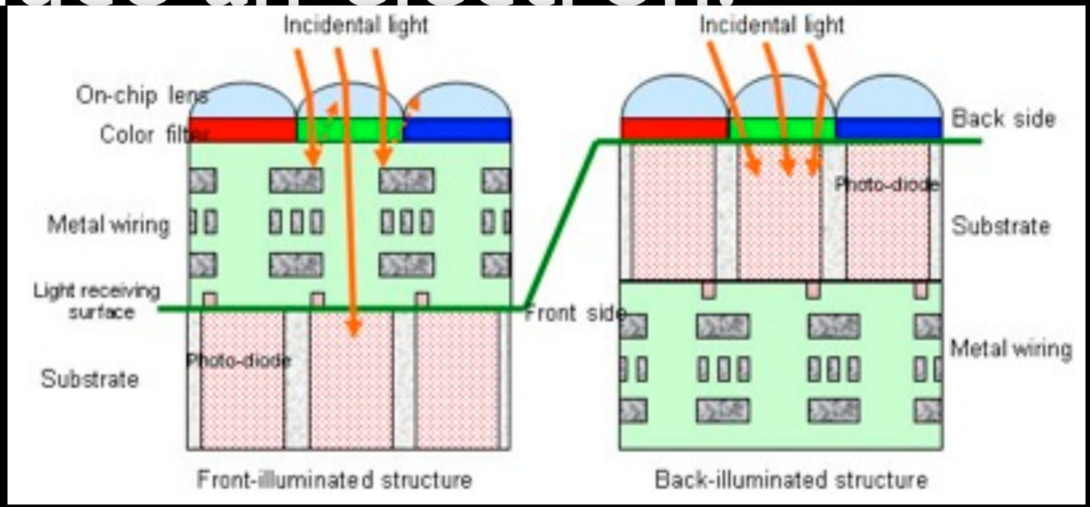

The Pixel

- Size matters
  - Casio EX-F1: 2.5µ x 2.5µ
  - Nokia N900: 3.1µ x 3.1µ
  - Canon 5D II: 6.4µ x 6.4µ
- Capacity matters

Blooming

(ccd-sensor.de)

CMOS vs. CCD

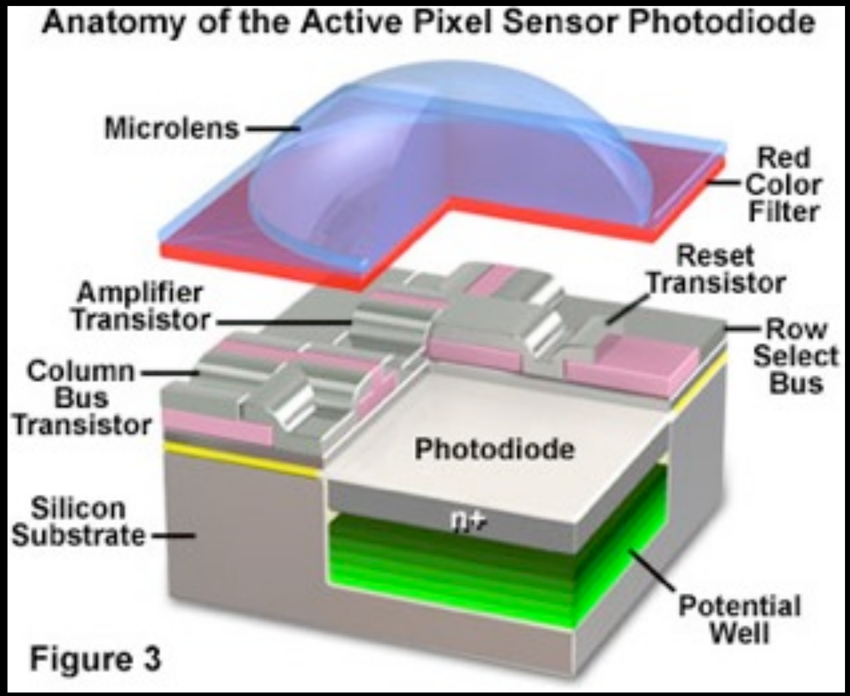

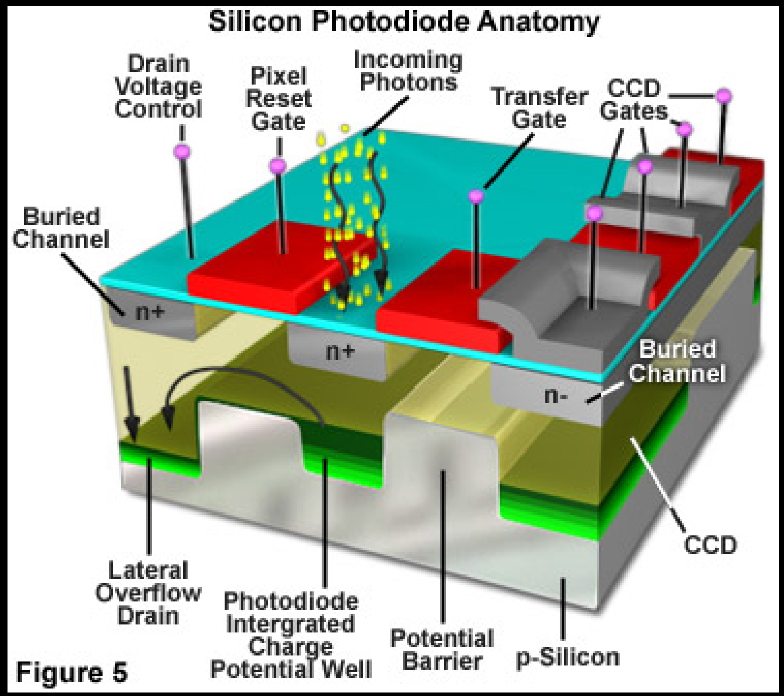

- Complimentary Metal-Oxide Semiconductor
  - per-pixel amplifier converts charges to voltage.
  - cheap, low-power but noisy
- Charge-Coupled Device
  - charge shifts along column to an amplifier
  - good but not as cheap.

Anatomy of a Pixel

(Sony)

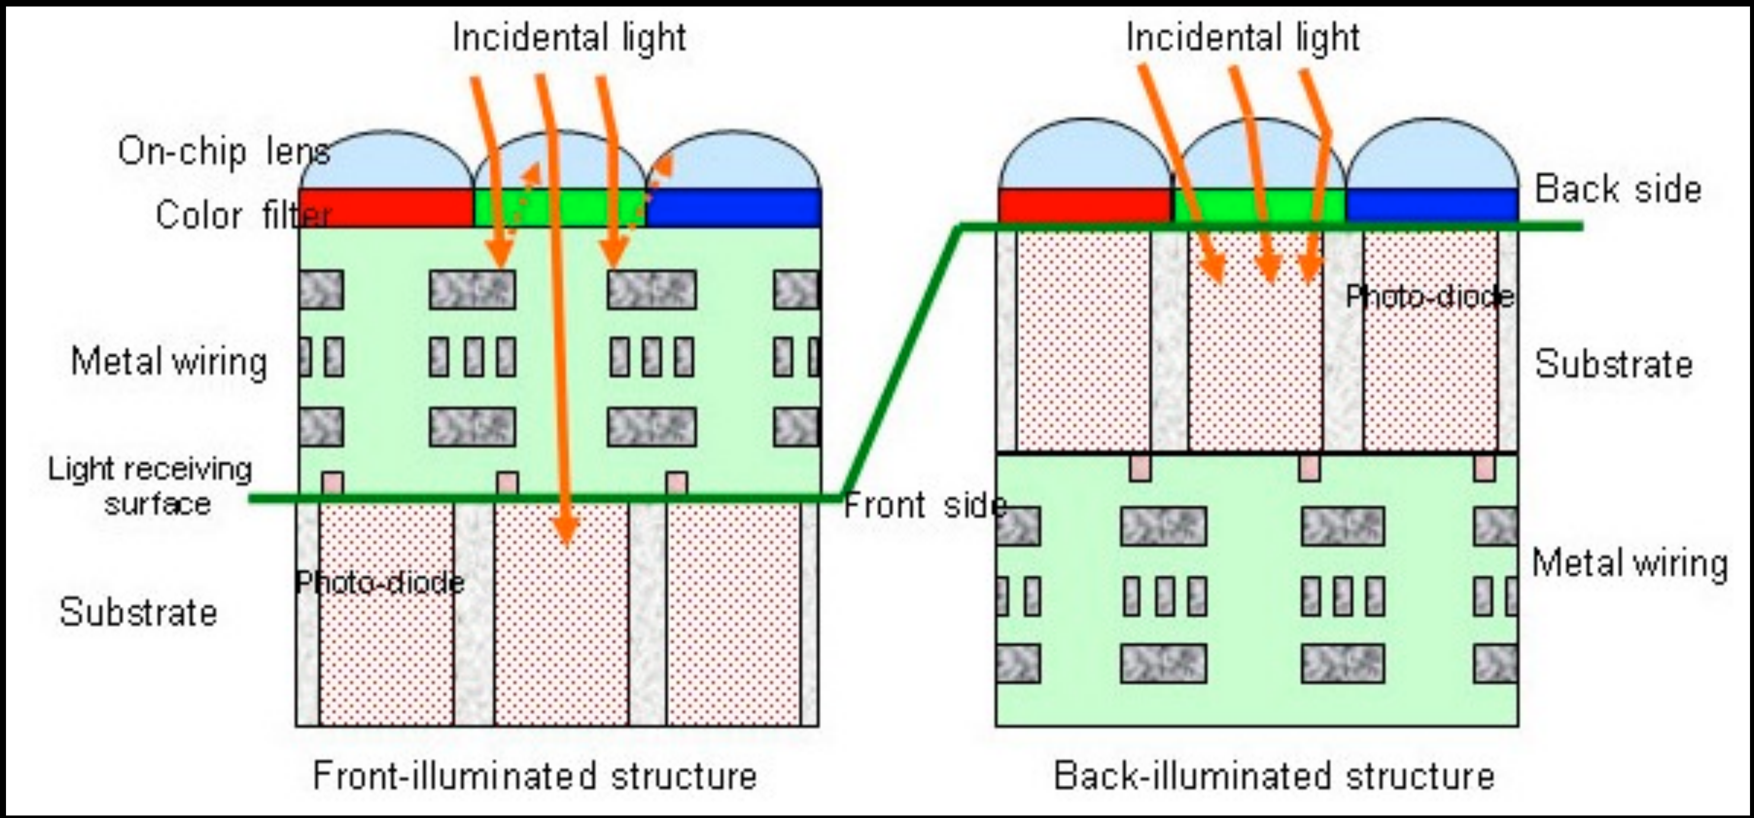

Anatomy of a Pixel

- Microlens
  - Improves fill factor
- Substrate
  - Science happens
- Circuitry
  - For reading / resetting

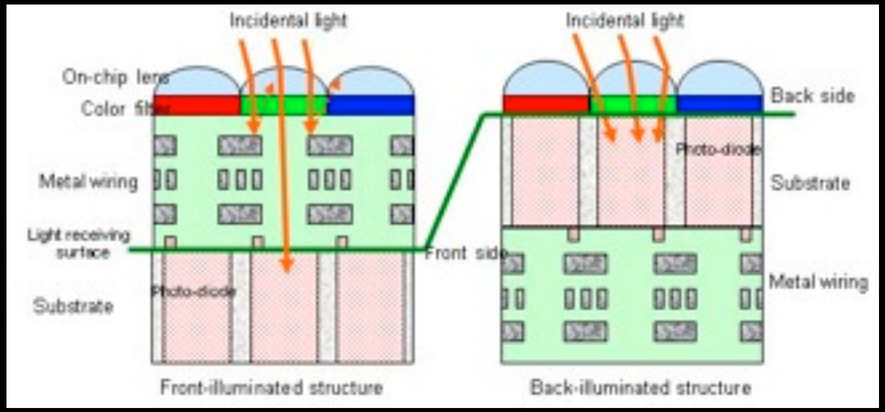

Antialiasing filters

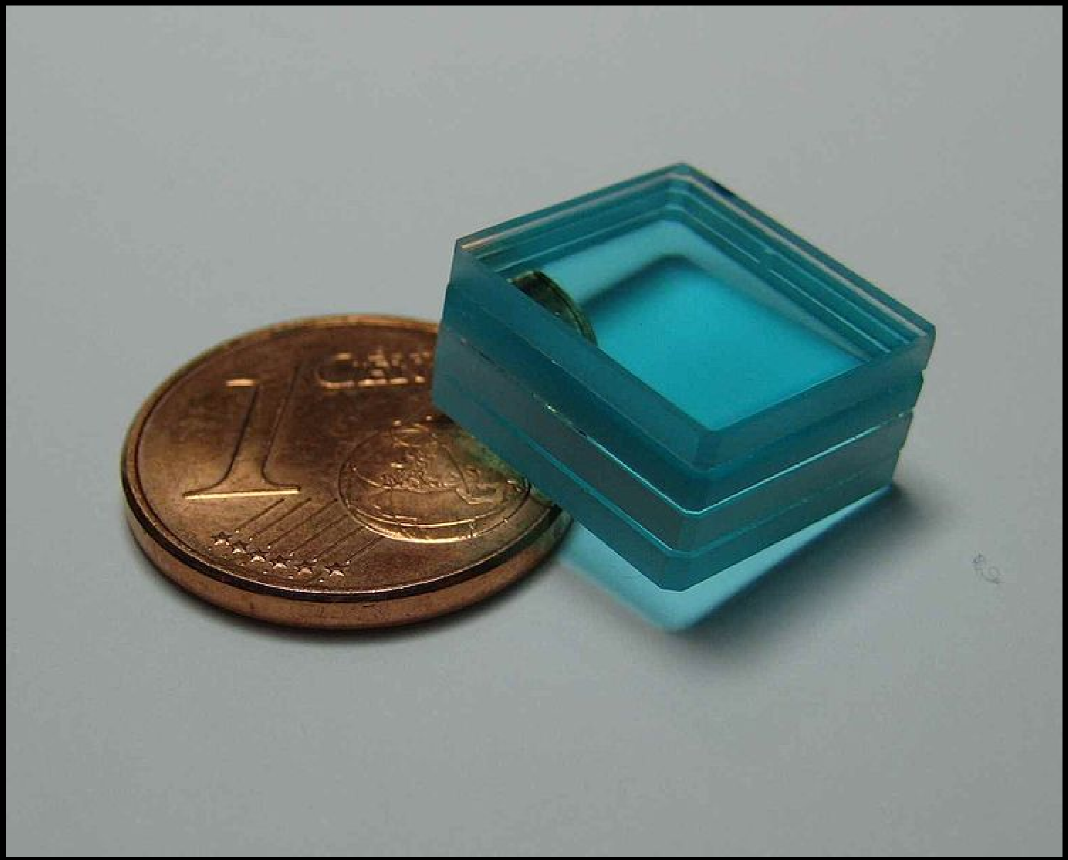

antialiasing filter

birefringence in a calcite crystal

- Typically two layers of birefringent material
  - splits 1 ray into 4 rays

Antialiasing filters

anti-aliasing filter removed normal

Antialiasing filters

anti-aliasing filter removed

normal

Color Filter Arrays

- Recall: we need information on (ρ, γ, β).
- Need discrimination among multiple wavelengths
  - Three types (of spectral sensitivity) of pixels would be sufficient.
- Color filter array: turns pixels into one of three types.

Bayer Pattern

- Checkered pattern of green and alternating red/blue
- Pretty much everywhere

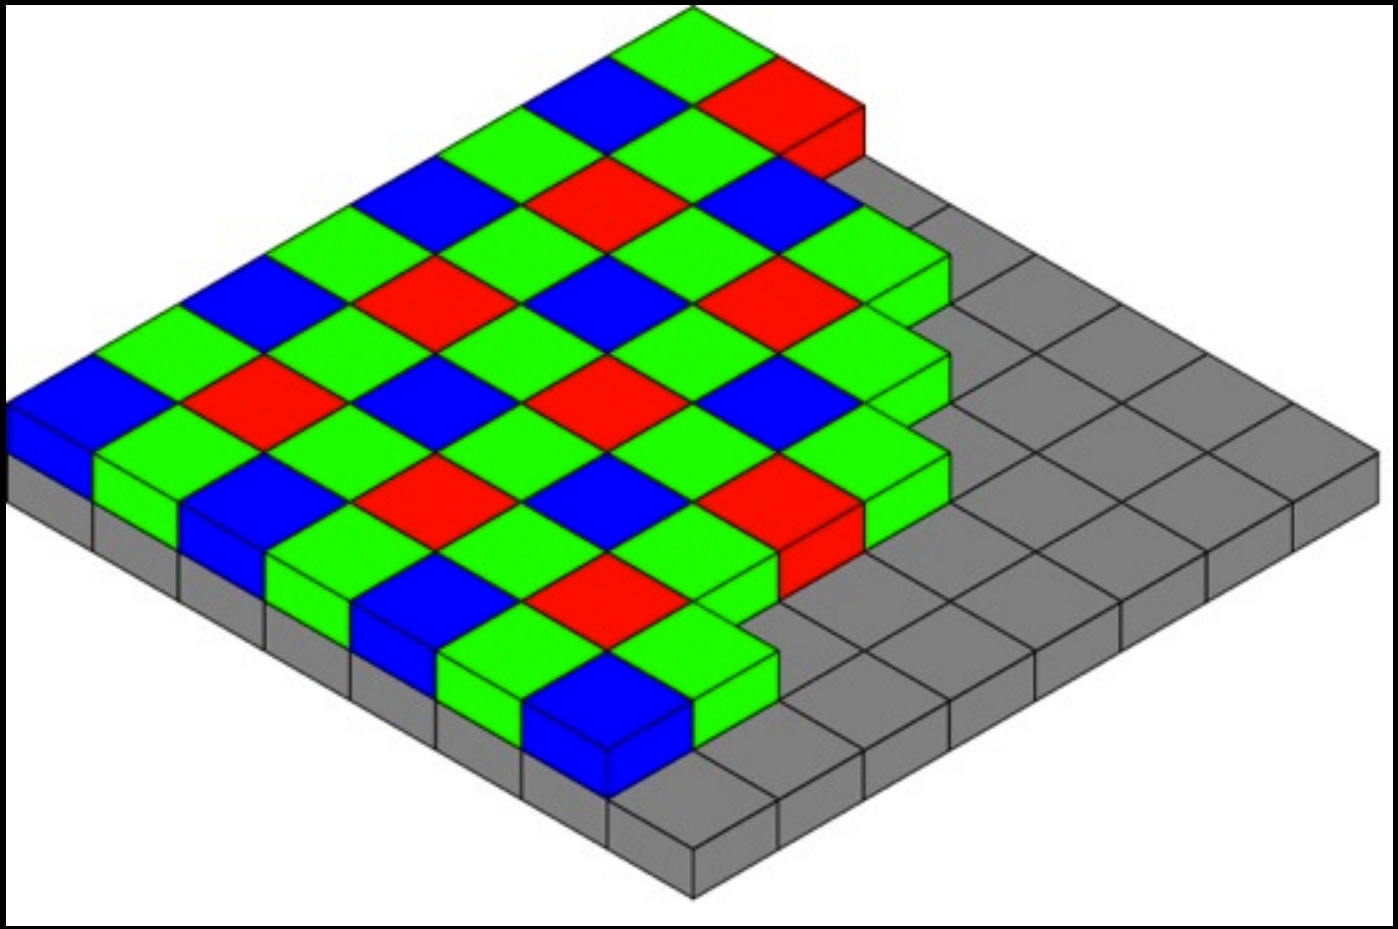

Bayer Pattern

- Checkered pattern of green and alternating red/blue
- Pretty much everywhere

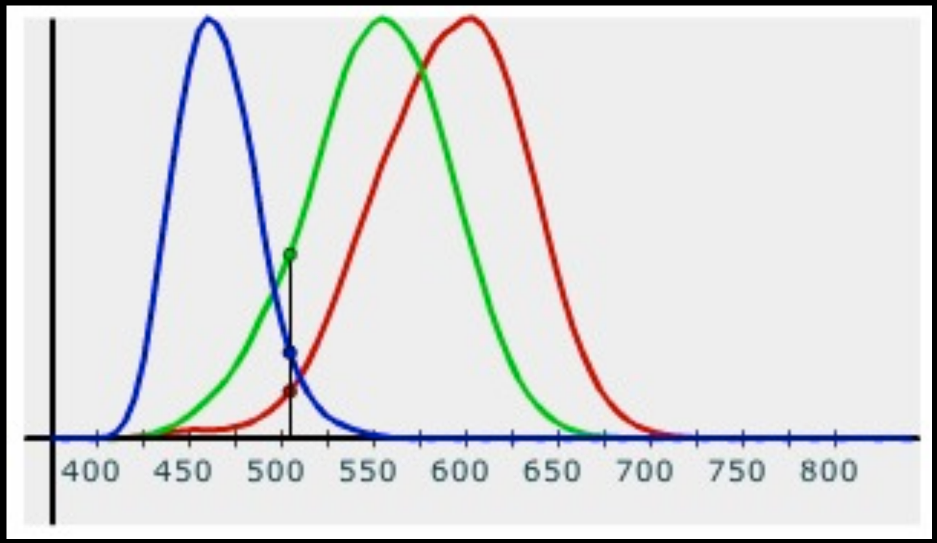

Cone cells

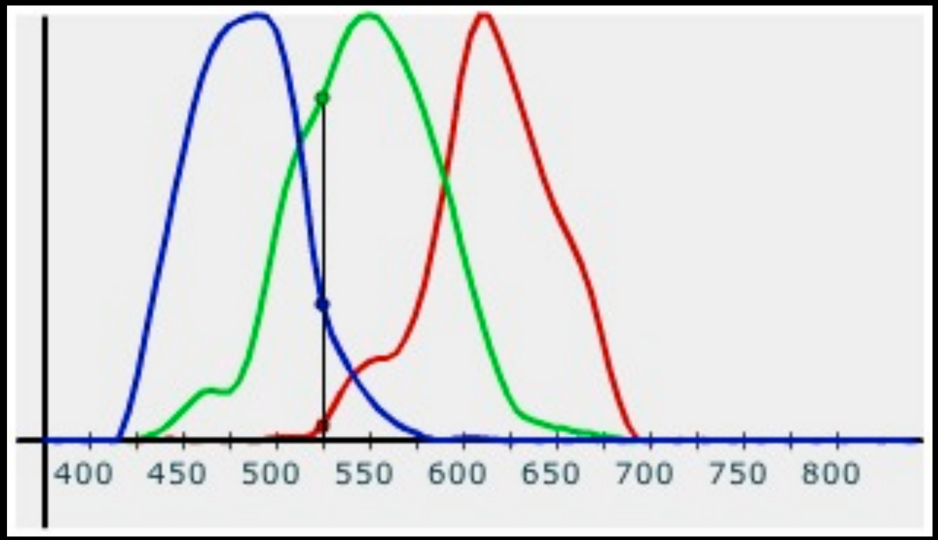

Color filters in Canon 30D

Other Patterns

RGBY

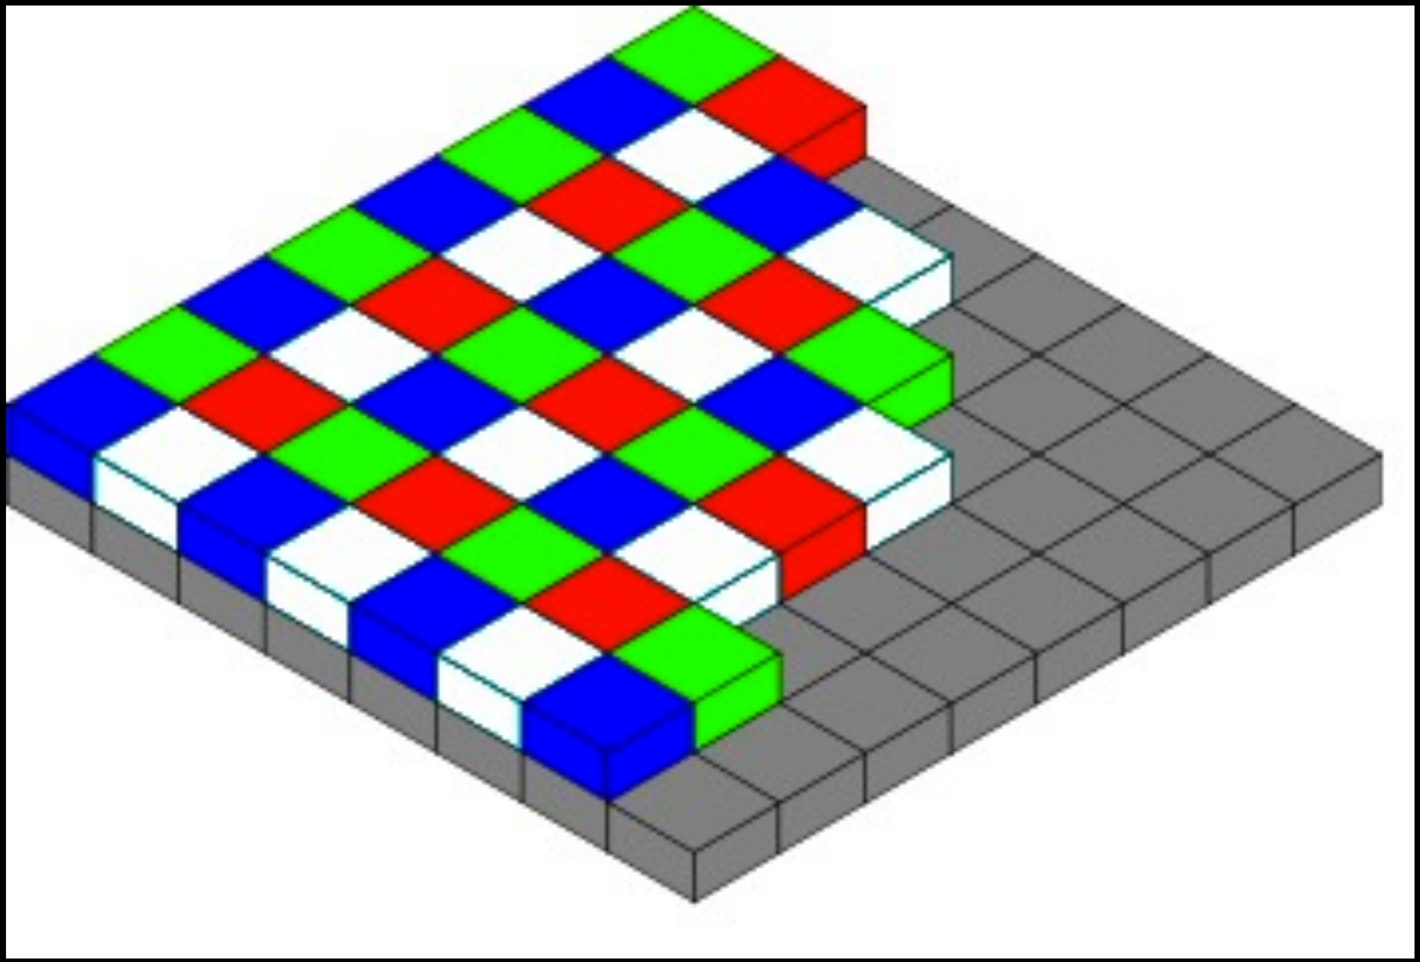

Other Patterns

"RGBE" (Sony)

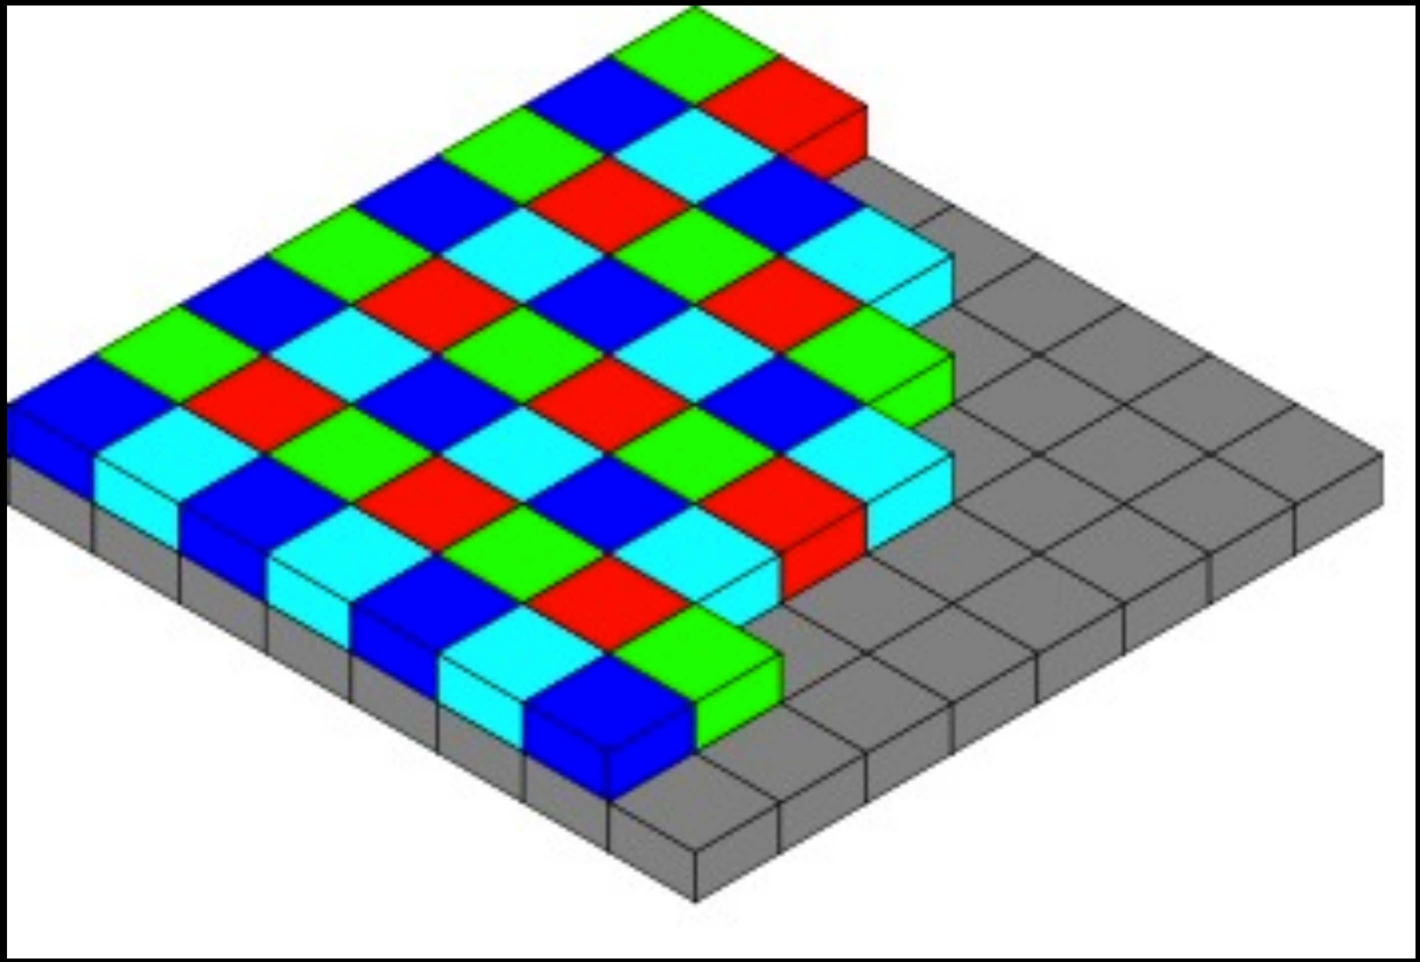

Other Patterns

Subtractive Colors

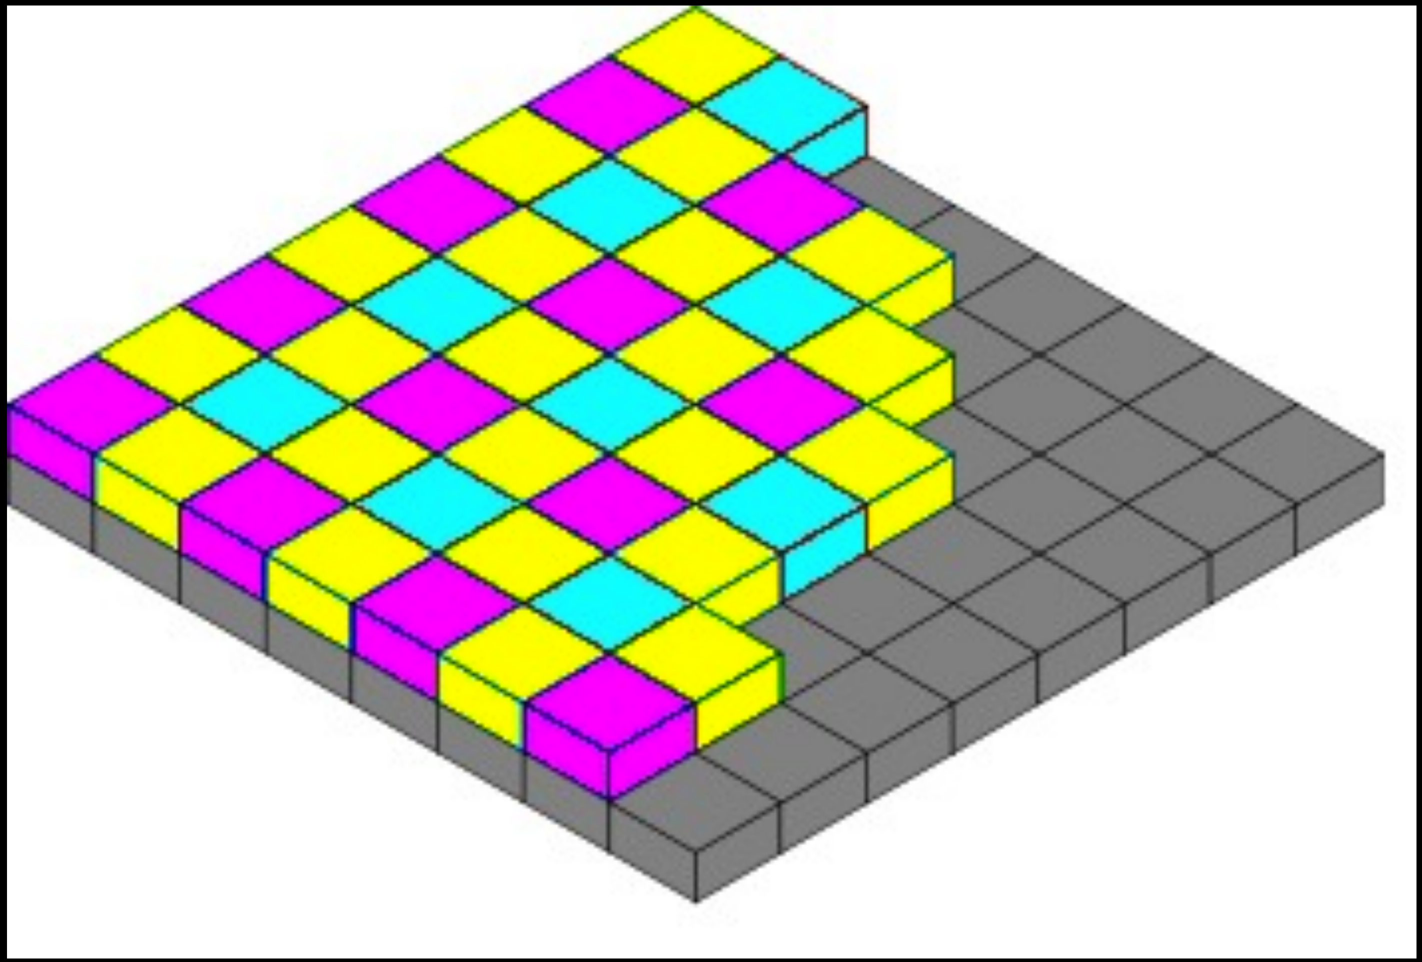

Other Patterns

"SuperCCD" (Fuji)

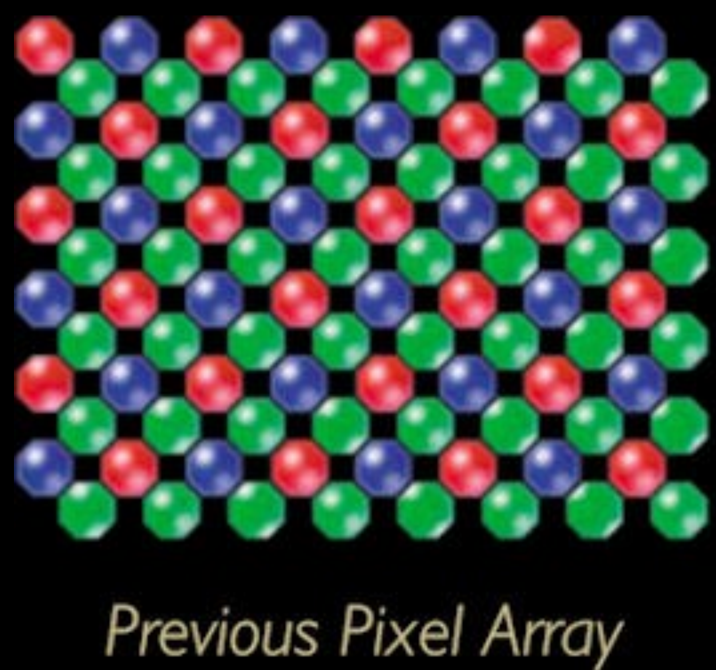

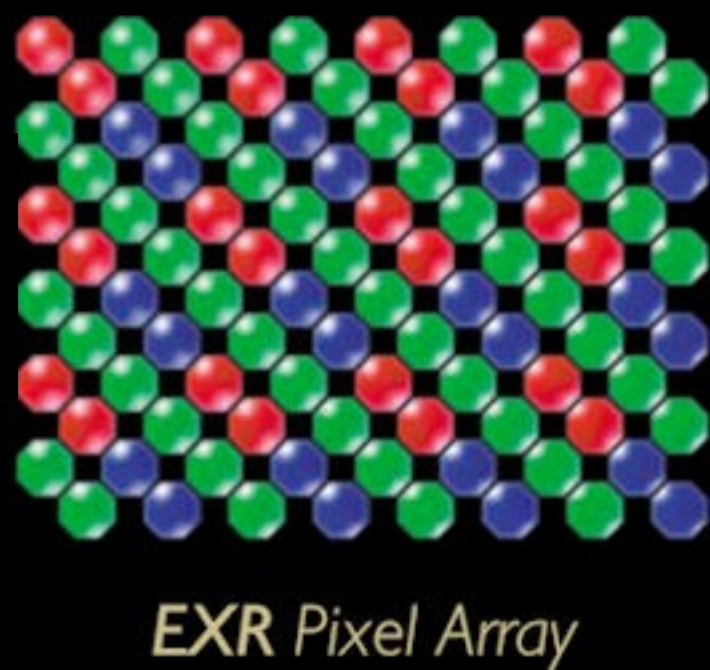

Foveon Sensor

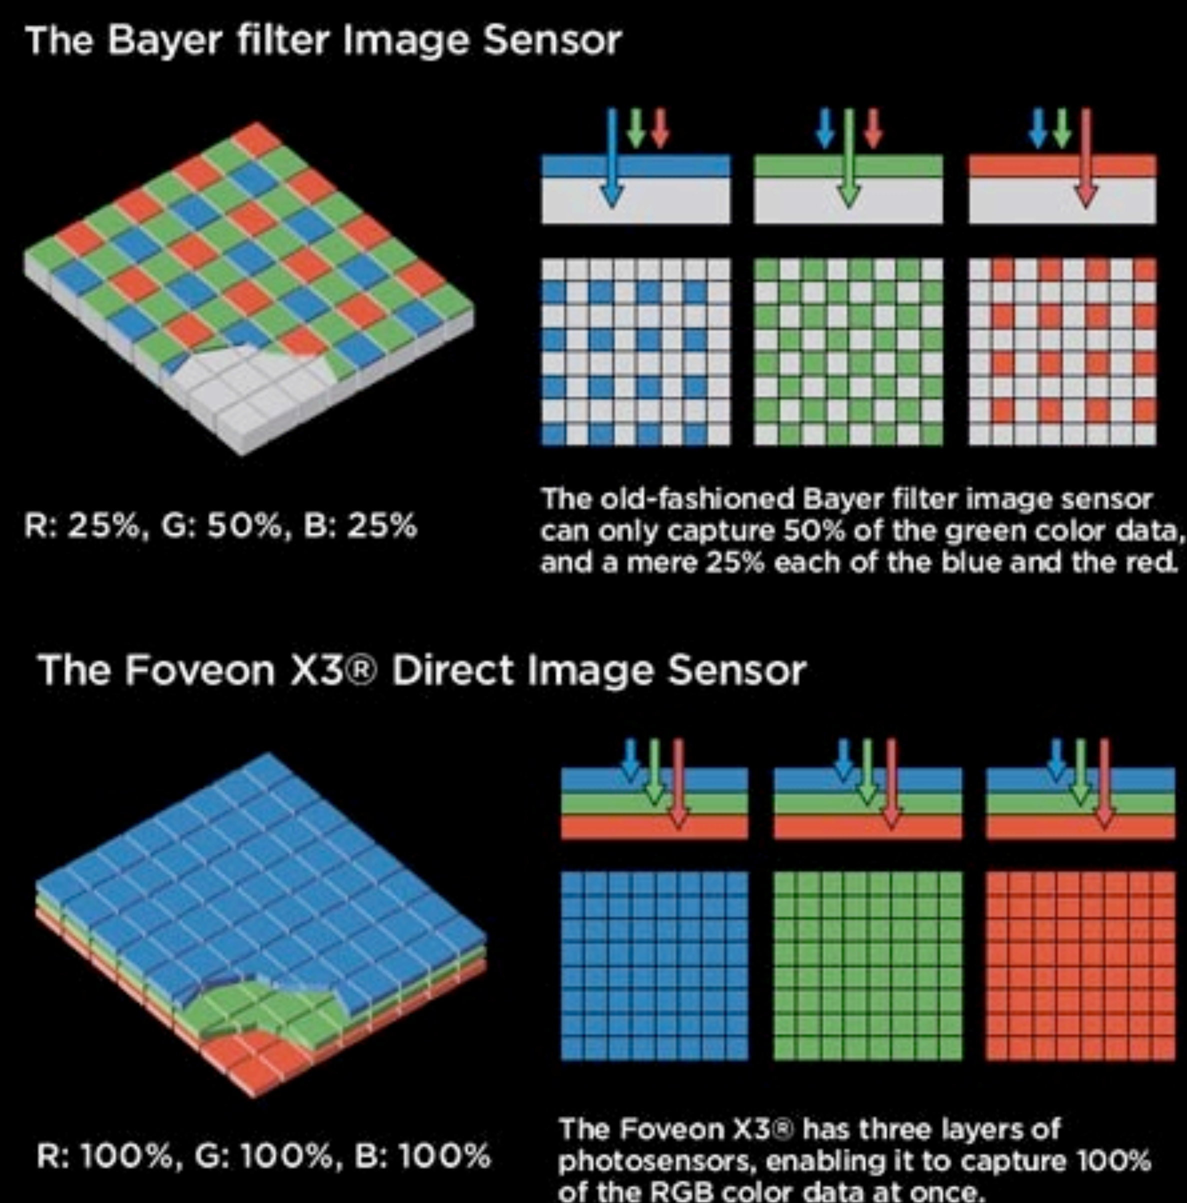

Reading Pixels (CCD)

- Each pixel has a “vertical” CCD
- A … m all p … riate “…
- V … nison.
  - … CD is … ges there move *right* in unison, and are read out at the end.

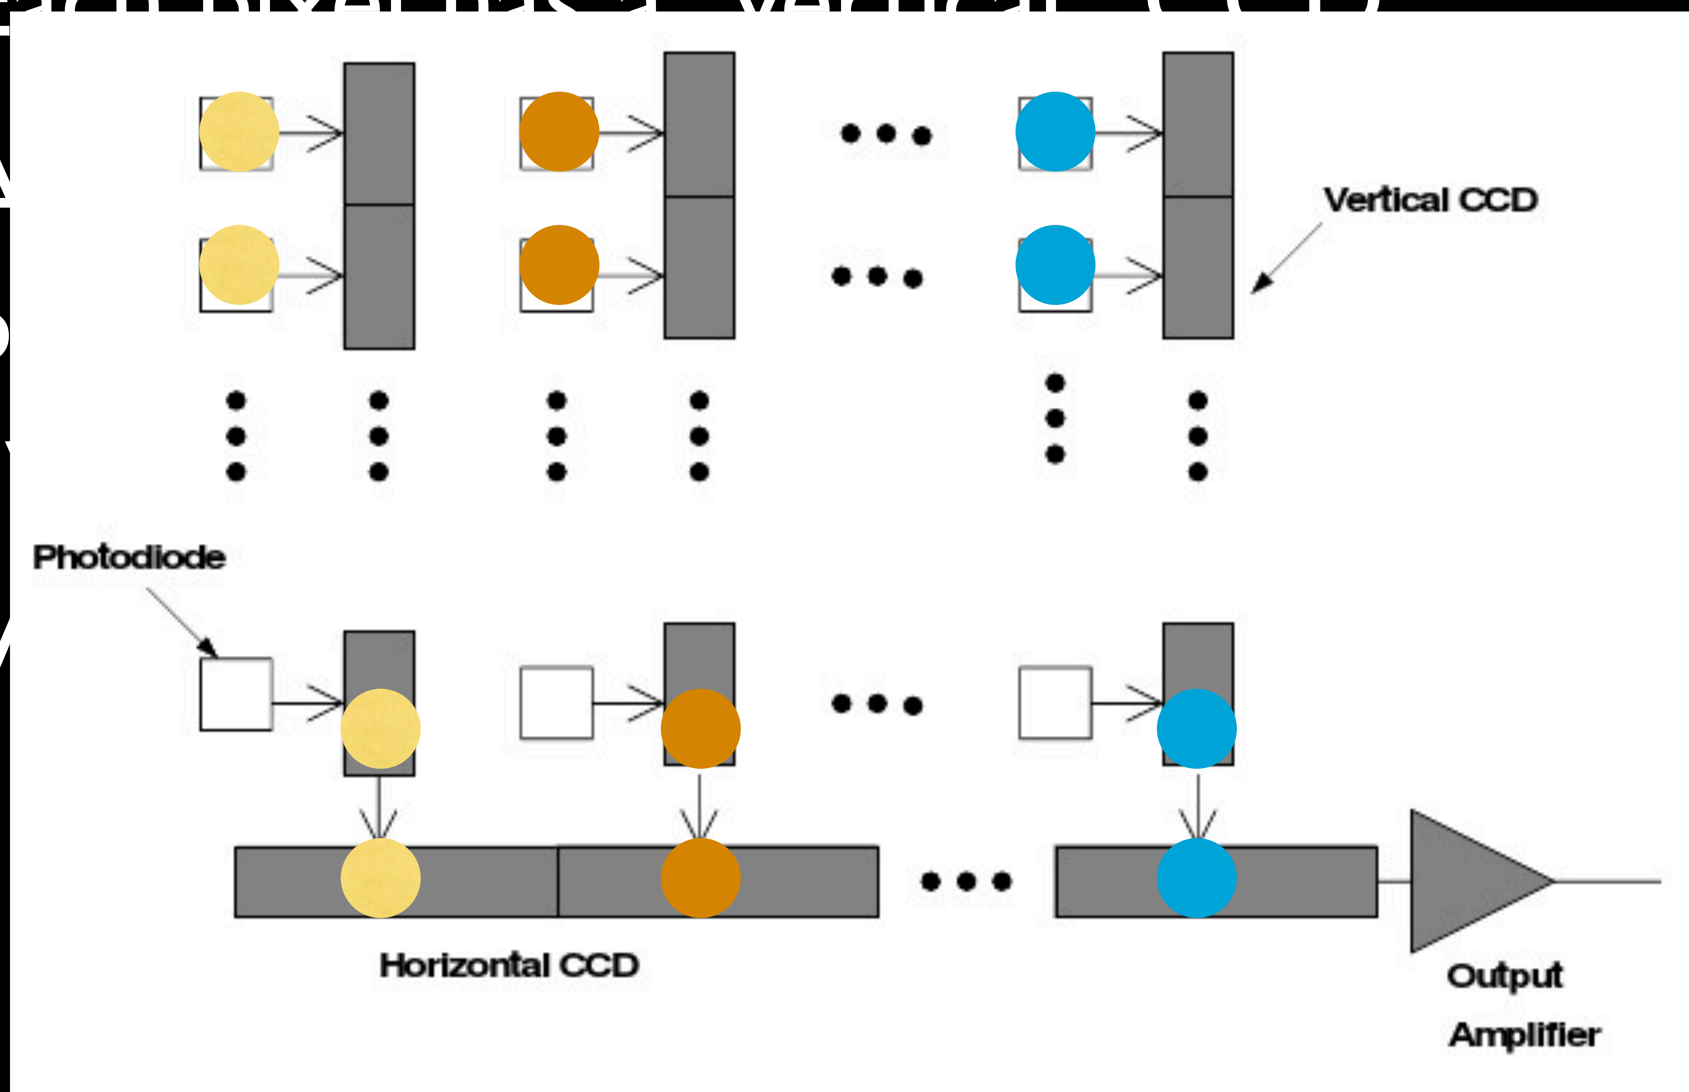

Reading Pixels (CMOS)

- One storage capacitor per column
  -
- F
  - umn

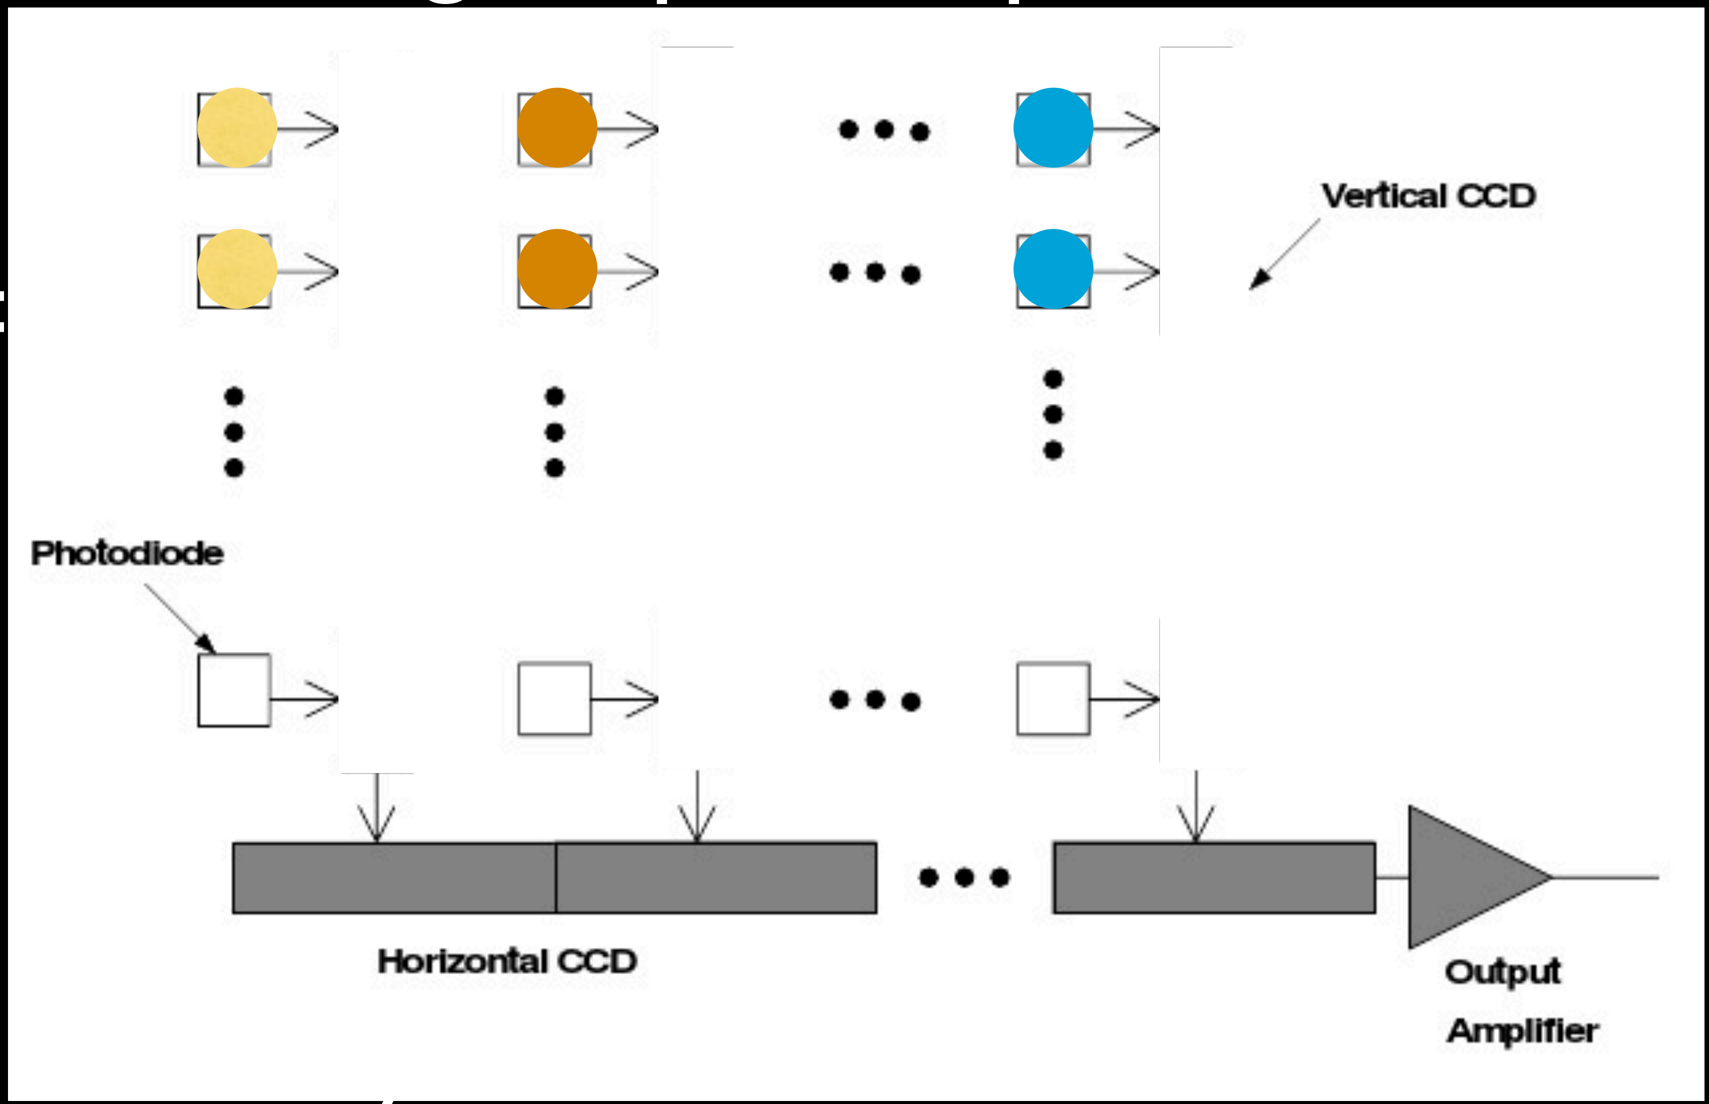

  - Read the storage capacitors.

Analog-to-Digital Conversion

- Convert analog voltage to discrete values.

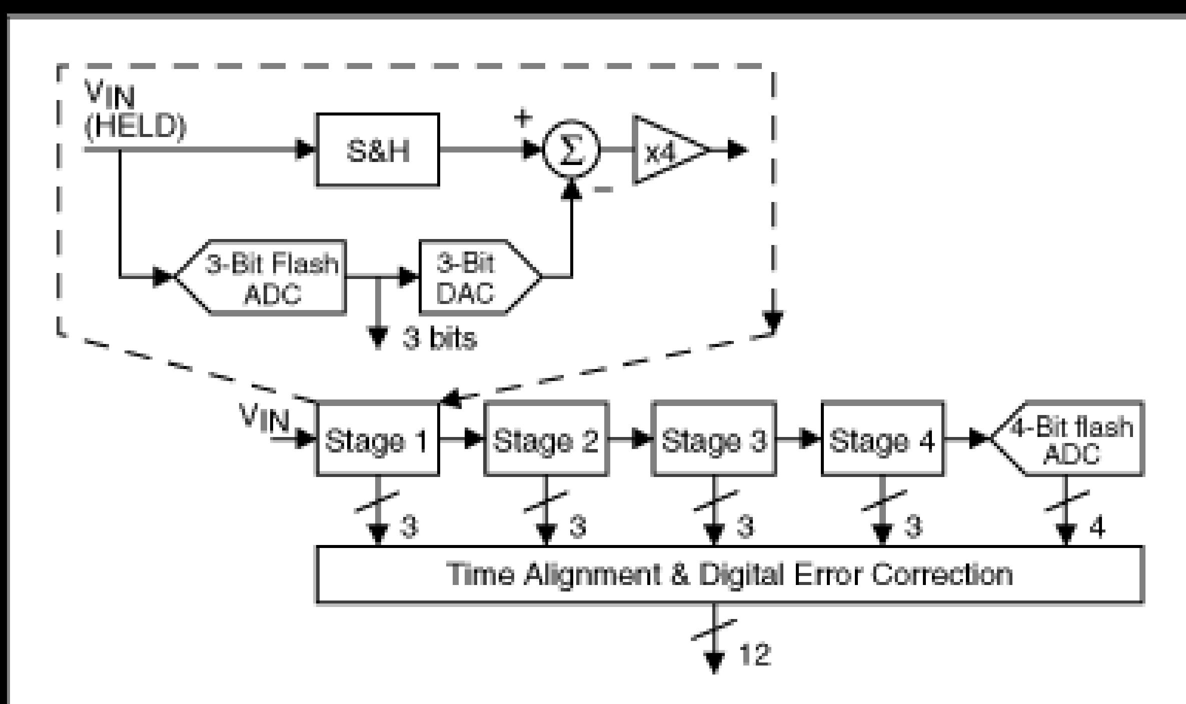

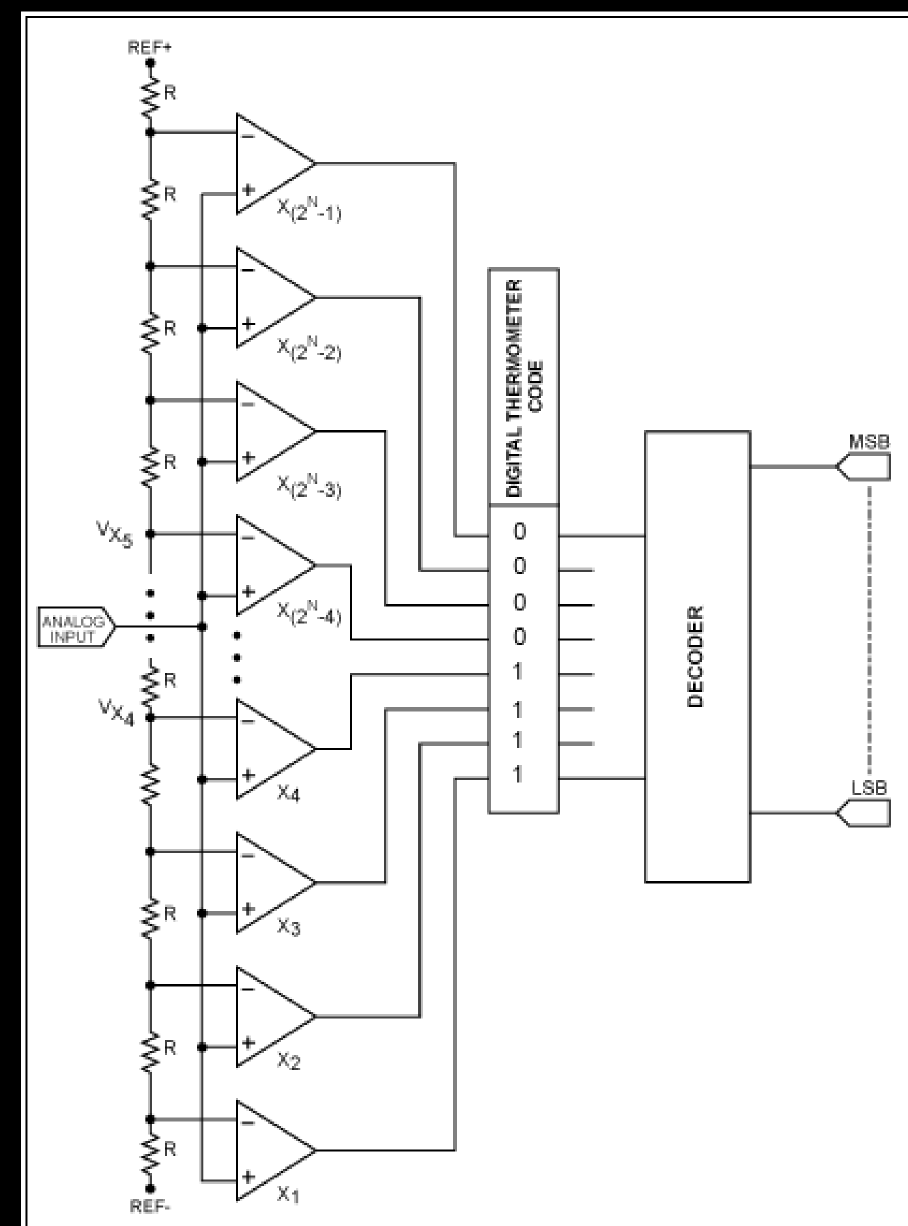

Dynamic Range

- Typical ADC work with 8-16 bits.
  - At n bits, *dynamic range* of $\mathbf{2^n:1}$
- Even the best ADC is only as good as the pixel well capacity.

Dynamic Range

- Human eye
  - Capable of **100:1**
  - With adaptation, **1,000,000:1**
- World
  - Typically**100,000:1**
  - Up to **100,000,000,000:1**

Analog Gain (“ISO”)

- Amplifies the analog voltage before ADC
  - Avoids amplifying quantization error + other noise post-ADC.

Other Types of Sensors

- Amplifies the analog voltage before ADC
  - Avoids amplifying quantization error + other noise post-ADC.

Outline

- Background material, Part II
  - Perception
  - Sensor
  - Noise

Photon Shot Noise

- Pixels measure the # of incident photons.
  - Upon a fixed area, during a fixed time.

- Varies from time to time.
- Varies from pixel to pixel.
  - Follows the Poisson distribution.

Photon Shot Noise

- Poisson distribution
  - $p(k; \lambda) = \lambda^k e^{-\lambda} / k!$
  - Mean = Variance = λ
- Typically approximated as a Gaussian

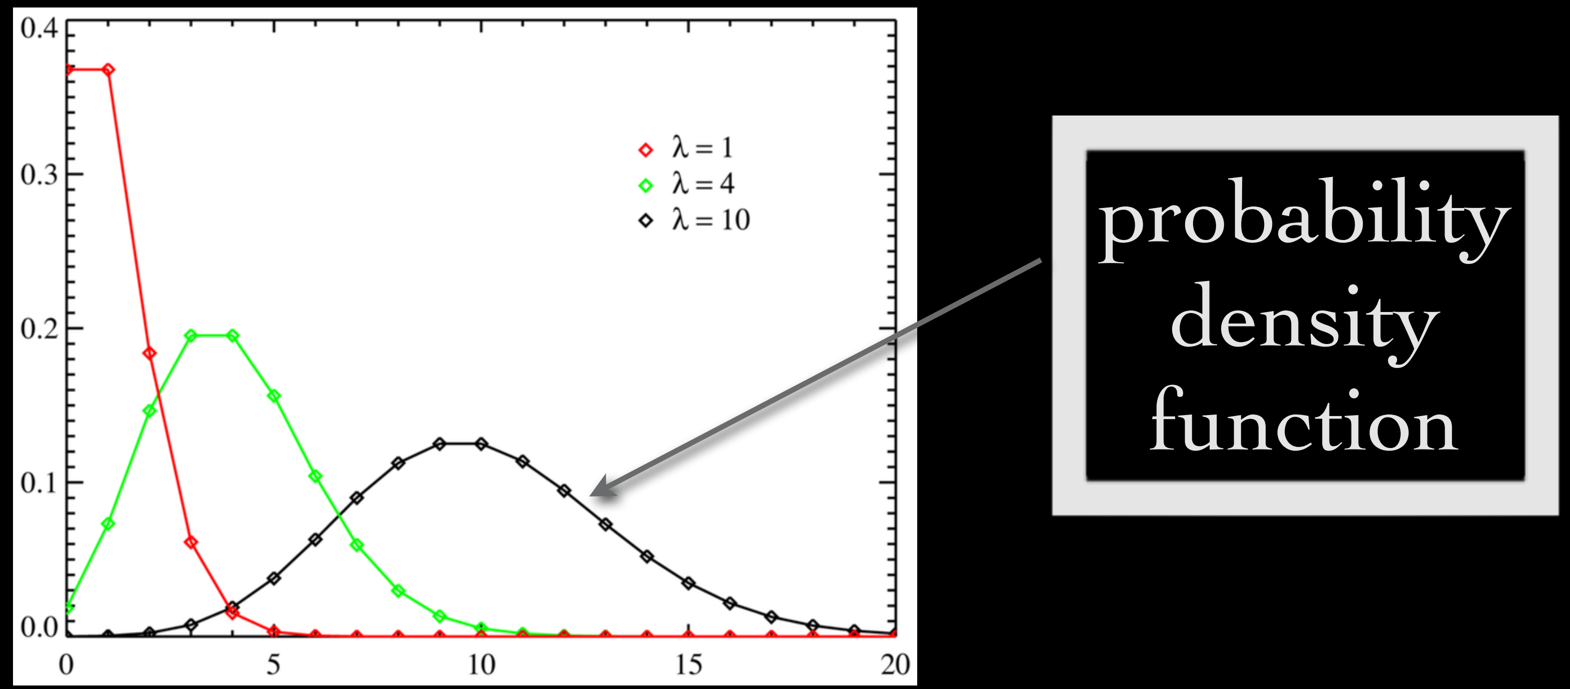

Signal v. Noise

- Assuming zero bias, we care about the ratio between
  - mean (signal)
  - standard dev. (noise)

Signal v. Noise

- e.g. in Poisson noise
  - Mean = λ
  - Standard deviation = $\lambda^{1/2}$
  - As the expected pixel value (mean) grows, the standard deviation grows slowly.
  - As signal grows, SNR rises.

Signal v. Noise

Test Chart

Captured by Canon 10D (ISO 1600)

Signal v. Noise

- SNR = 20 $\log_{10} (\mu/\sigma)$
  - Unit is dB
- If the ratio is 10-to-1, we achieve 20 dB.

Dark Current

- Electrons dislodged by random thermal activity.
  - Increases linearly with exposure time.
  - Increases exponentially with temperature..

Canon 20D, 612s exposure

Hot Pixels

- Electrons leaking into wells because of manufacturing defects
  - Increases linearly with exposure time.

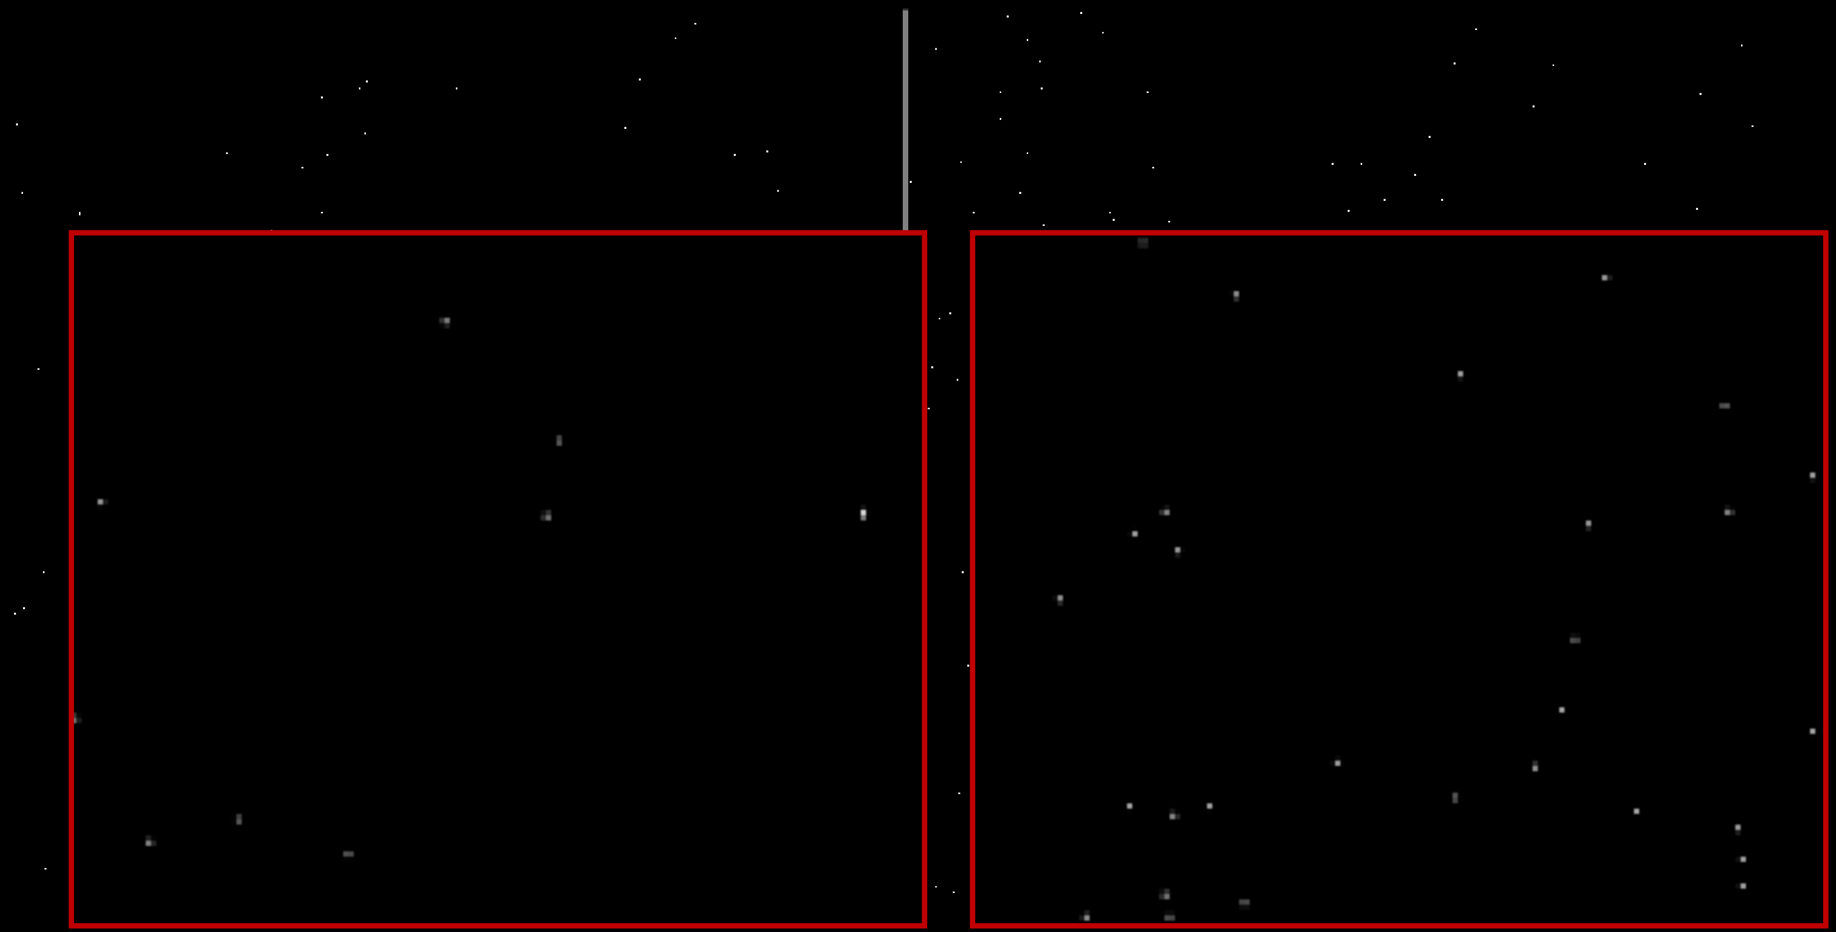

Canon 20D, 15s/30s exposure

Fixed Pattern Noise

- Manufacturing variations across pixels, columns, etc
  - Constant over time

Canon 20D, ISO 800, cropped

Read Noise

- Thermal noise in readout circuitry
  - Mainly in CMOS

Canon 1D Mk III, cropped

Pixel Response Non-Uniformity

- ~1% variance in the sensivity of pixels
  - Think about it as a per-pixel vignetting issue.

Quantization Error

- Any ADC process has quantization errors.
    - Depends on the bitdepth of the ADC.

Electronic Interference

- Interference from other circuitry
  - Exacerbated by poor insulation

Noise: Summary

- Photon shot noise
- Hot pixels
- Dark current
- Fixed pattern noise
- Read noise
- Pixel non-uniformity
- ...

Much of the literature treats these altogether as a Gaussian noise

Noise Game

		Signal	Shot noise variance	Read noise variance (incurred before analog gain)	Read noise variance (incurred after analog gain)
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1				
sharp, noisy	(C) exp 1, gain n				
	(D) Accum. n (A)'s				
multi-image denoising	(E) Average n (C)'s				

Slide stolen and adapted from Sung Hee Park

Noise Game

incurred before / after analog gain

		Signal	Shot noise variance	Read noise variance	Read noise variance
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1				
sharp, noisy	(C) exp 1, gain n				
	(D) Accum. n (A)'s				
multi-image denoising	(E) Average n (C)'s				

Slide stolen and adapted from Sung Hee Park

Noise Game

		Signal	Shot noise variance	Read noise variance	Read noise variance
				incurred before / after analog gain	
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1	nJ			
sharp, noisy	(C) exp 1, gain n	nJ			
	(D) Accum. n (A)'s	nJ			
multi-image denoising	(E) Average n (C)'s	nJ			

Slide stolen and adapted from Sung Hee Park

Noise Game

		Signal	Shot noise variance	Read noise variance (incurred before analog gain)	Read noise variance (incurred after analog gain)
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1	nJ	nJ		
sharp, noisy	(C) exp 1, gain n	nJ	n^2J		
	(D) Accum. n (A)'s	nJ	nJ		
multi-image denoising	(E) Average n (C)'s	nJ	nJ		

Slide stolen and adapted from Sung Hee Park

Noise Game

		Signal	Shot noise variance	Read noise variance	Read noise variance
				incurred before / after analog gain	
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1	nJ	nJ	σ_0^2	
sharp, noisy	(C) exp 1, gain n	nJ	n^2J	$n^2\sigma_0^2$	
	(D) Accum. n (A)'s	nJ	nJ	$n\sigma_0^2$	
multi-image denoising	(E) Average n (C)'s	nJ	nJ	$n\sigma_0^2$	

Slide stolen and adapted from Sung Hee Park

Noise Game

		Signal	Shot noise variance	Read noise variance	Read noise variance
				incurred before / after analog gain	
	(A) exp 1, gain 1	J	J	σ_0^2	σ_1^2
blurry, clean	(B) exp n, gain 1	nJ	nJ	σ_0^2	σ_1^2
sharp, noisy	(C) exp 1, gain n	nJ	n^2J	$n^2\sigma_0^2$	σ_1^2
	(D) Accum. n (A)'s	nJ	nJ	$n\sigma_0^2$	$n\sigma_1^2$
multi-image denoising	(E) Average n (C)'s	nJ	nJ	$n\sigma_0^2$	$1/n\ \sigma_1^2$

Slide stolen and adapted from Sung Hee Park

Questions?

- TO-DOs
  - Keep thinking about the final project.
  - Pick up tablets on Friday.
  - Do the assigned readings for Monday.