


Recap: Frequency Domain

Fourier Bases


Teases away fast vs. slow changes in the image.



This change of basis is the Fourier Transform


Review

1. Match the spatial domain image to the Fourier magnitude image




Sampling

Why does a lower resolution image still make sense to us? What do we lose?




Subsampling by a factor of 2



Throw away every other row and column to create a 1/2 size image


Aliasing problem

- 1D example (sinewave):



Aliasing problem


- 1D example (sinewave):



Aliasing problem

- Sub-sampling may be dangerous....
- Characteristic errors may appear:
 - “car wheels rolling the wrong way in movies”
 - “Checkerboards disintegrate in ray tracing”
 - “Striped shirts look funny on color television”

A form of spatial aliasing: Moire Pattern






Aliasing in video

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)

Aliasing in graphics



Sampling and aliasing


256x256

128x128

64x64


32x32

16x16



Nyquist-Shannon Sampling Theorem

- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{\max}$
- f_{\max} = max frequency of the input signal
- This will allow to reconstruct the original perfectly from the sampled version



Anti-aliasing

Solutions:

- Sample more often
- Get rid of all frequencies that are greater than half the new sampling frequency
 - Will lose information
 - But it's better than aliasing
 - Apply a smoothing filter

Algorithm for downsampling by factor of 2

1. Start with image(h, w)

2. Apply low-pass filter

```
im_blur = imfilter(image, fspecial('gaussian', 7, 1))
```

3. Sample every other pixel

```
im_small = im_blur(1:2:end, 1:2:end);
```

Anti-aliasing


256x256

128x128

64x64

32x32

16x16




256x256

128x128


64x64

32x32

16x16



Subsampling without pre-filtering



1/2




1/4 (2x zoom)



1/8 (4x zoom)

Subsampling with Gaussian pre-filtering



Gaussian $1/2$



G $1/4$



G $1/8$

Recap – Light, Cameras, Eyes









FIGURE 2.17 (a) Continuous image projected onto a sensor array. (b) Result of image sampling and quantization.




Electromagnetic Spectrum



Human Luminance Sensitivity Function

Visible Light

Why do we see light of these wavelengths?




...because that's where the Sun radiates EM energy

The Physics of Light

Any patch of light can be completely described physically by its spectrum: the number of photons (per time unit) at each wavelength 400 - 700 nm.


Photons
(per ms.)




The Physics of Light

Some examples of the spectra of light sources


A. Ruby Laser




B. Gallium Phosphide Crystal



C. Tungsten Lightbulb




D. Normal Daylight





The Physics of Light

Some examples of the reflectance spectra of surfaces



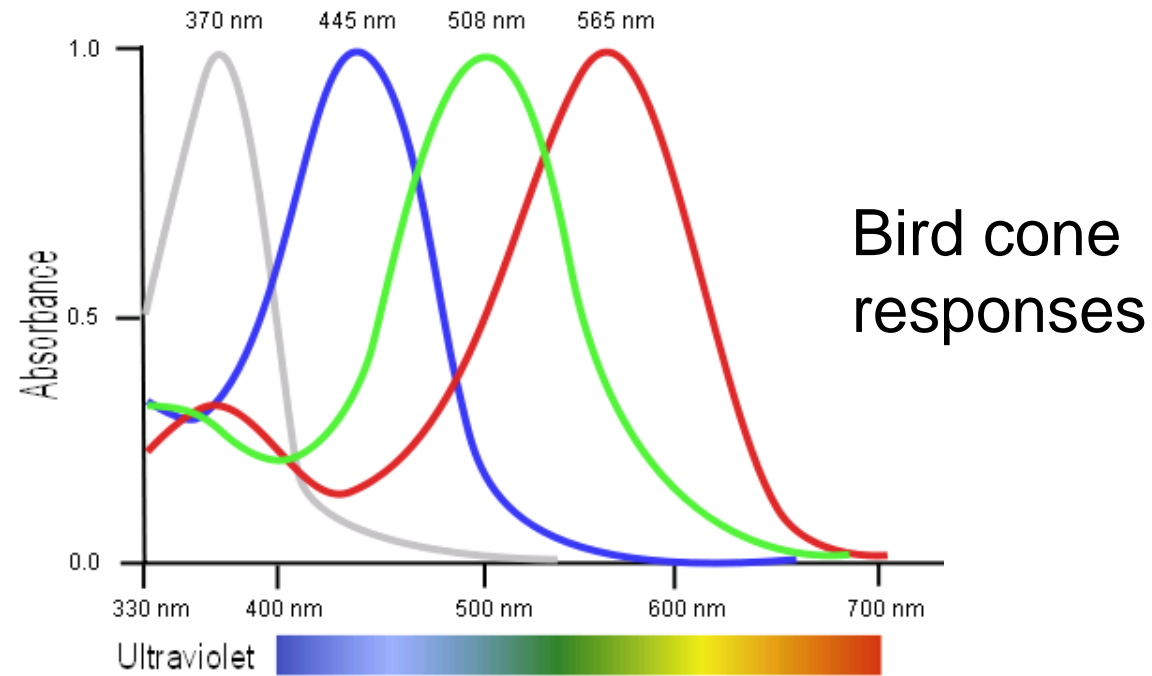
Physiology of Color Vision

Three kinds of cones:



- Why are M and L cones so close?
- Why are there 3?

Tetrachromacy



- Most birds, and many other animals, have cones for ultraviolet light.
- Some humans, mostly female, seem to have slight tetrachromatism.

Table 1.

Cone pigment complements in some eutherian mammals.


order	exemplars	SWS1 pigment ^a	LWS pigment ^a	reference
Rodentia	<i>Mus</i> (mouse)	UV	M	Jacobs et al. (1991)
	<i>Rattus</i> (rat)	UV	M	Jacobs et al. (1991)
	<i>Geomys</i> (gopher)	UV	M	Williams et al. (2005)
	<i>Cavia</i> (guinea pig)	S	M	Parry & Bowmaker (2002)
	<i>Spermophilus</i> (squirrel)	S	M	Jacobs et al. (1985)
	<i>Cricetomys</i> (African rat)	absent	ML ^b	Peichl & Moutairou (1998)
Lagomorpha	<i>Oryctolagus</i> (rabbit)	S	M	Nuboer et al. (1983)
primate	<i>Macaca</i> (macaque monkey)	S	M+L	Schnapf et al. (1988)
	<i>Saimiri</i> (squirrel monkey)	S	poly (3)	Mollon et al. (1984)
	<i>Aotus</i> (owl monkey)	absent	L	Jacobs et al. (1993b)
	<i>Alouatta</i> (howler monkey)	S	M+L	Jacobs et al. (1996a)
	<i>Galago</i> (bushbaby)	absent	L	Deegan II & Jacobs (1996)
	<i>Lemur</i> (ring-tailed)	S	L	Jacobs & Deegan II (1993)
	<i>Propithecus</i> (sifaka)	S	poly (2)	Tan & Li (1999)
Scandentia	<i>Tupaia</i> (tree shrew)	S	L	Jacobs & Neitz (1986)
Cetacea	<i>Eschrichtius</i> (whale)	absent	L	Levenson & Dizon (2003)
	<i>Tursiops</i> (dolphin)	absent	L	Fasick et al. (1998)
Artiodactyla	<i>Bos</i> (cow)	S	L	Jacobs et al. (1994)
	<i>Odocoileus</i> (deer)	S	M	Jacobs et al. (1994)
	<i>Sus</i> (pig)	S	L	Neitz & Jacobs (1989)
Perissodactyla	<i>Equus</i> (horse)	S	L	Carroll et al. (2001)
Carnivora	<i>Felis</i> (cat)	S	L	Loop et al. (1987)
	<i>Canis</i> (dog)	S	L	Jacobs et al. (1993a)
	<i>Mustela</i> (ferret)	S	L	Calderone & Jacobs (2003)
	<i>Ursus</i> (bear)	S	L	Levenson et al. (2006)
	<i>Crocuta</i> (hyena)	UV/S	ML	Calderone et al. (2003)
	<i>Phoca</i> (seal)	absent	L	Levenson et al. (2006)
	<i>Enhydra</i> (otter)	S	L	Levenson et al. (2006)

(b) Eutherians

Representatives from two cone opsin gene families appear in contemporary eutherian mammals and, with the exception of some primates, none of these animals derive more than a single photopigment type from each of their two gene families (*SWS1* and *LWS*; [figure 1](#)). Given that commonality, what cone pigments may have been present in the retinas of the early eutherians? Sequence comparisons of cone opsin genes have suggested that the ancestral mammalian pigments drawn from these two gene families peaked in the UV, at about 360 nm ([Hunt et al. 2001](#)), and in the long wavelengths at 560 nm ([Yokoyama et al. 2008](#)). Assuming the visual pigment chromophore of early mammals was 11-*cis*-retinal, the same as that of contemporary mammals, the cone pigments of early eutherian mammals would have had absorption properties similar to those sketched at the bottom of [figure 1](#).

If these deductions are correct, the retinas of the early eutherians were similar to those of the majority of contemporary mammals in containing two types of cone pigment. Such an arrangement could support dichromatic colour vision. Whether it did would additionally depend on there having been at least some degree of selective expression of the two pigment types in separate receptor classes, on these early mammals having a nervous system organized to allow a contrast of signals from the two types of cone, and on them at least occasionally encountering photic environments sufficiently intense to activate neural comparison circuits. The fact that this basic two cone pigment arrangement is largely conserved among contemporary mammals strongly attests to its adaptive utility in our eutherian ancestors and to its probable role in supporting dichromatic colour vision in these early mammals.

More Spectra



Color can be ambiguous



Color can be ambiguous



The dress



From Wikipedia, the free encyclopedia

For other uses, see [The Dress](#).

The dress is a photograph that became a [viral](#) internet sensation on 26 February 2015, when viewers disagreed over whether the dress pictured was coloured black and blue, or white and gold. The phenomenon revealed differences in human colour perception, which have been the subject of ongoing scientific investigations into [neuroscience](#) and [vision science](#), producing a number of papers published in peer-reviewed science journals.

The photo originated from a washed-out colour photograph of a [dress](#) posted on the [social networking service Facebook](#). Within the first week after the surfacing of the image, more than 10 million tweets mentioned the dress, using [hashtags](#) such as [#thedress](#), [#whiteandgold](#), and [#blackandblue](#). Although the colour of the dress was eventually confirmed as black and blue,^{[3][4]} the image prompted many discussions, with users discussing their different perceptions of the dress's colour. Members of the scientific community began to investigate the photo for fresh insights into human [colour vision](#).

The dress itself, which was identified as a product of the retailer Roman Originals, experienced a major surge in sales as a result of the incident. The retailer also produced a one-off version of the dress in white and gold as a charity campaign.

Contents [\[hide\]](#)


The dress





The original *The dress* picture

Designer

Roman Originals^[1]




Practical Color Sensing: Bayer Grid




- Estimate RGB at 'G' cells from neighboring values

Color Image




Color spaces

- How can we represent color?



Color spaces: RGB

Default color space



Some drawbacks

- Strongly correlated channels
- Non-perceptual



R
(G=0,B=0)





G
(R=0,B=0)



B
(R=0,G=0)

Color spaces: HSV

Intuitive color space



H
(S=1,V=1)




S
(H=1,V=1)



V
(H=1,S=0)

Color spaces: YCbCr

Fast to compute, good for compression, used by TV



Y
(Cb=0.5,Cr=0.5)




Cb
(Y=0.5,Cr=0.5)



Cr
(Y=0.5,Cb=0.5)

Color spaces: L*a*b*

“Perceptually uniform”* color space



L
(a=0,b=0)



a
(L=65,b=0)



b
(L=65,a=0)

If you had to choose, would you rather go without luminance or chrominance?

If you had to choose, would you rather go
without **luminance** or chrominance?

Most information in intensity



Only color shown – constant intensity

Most information in intensity



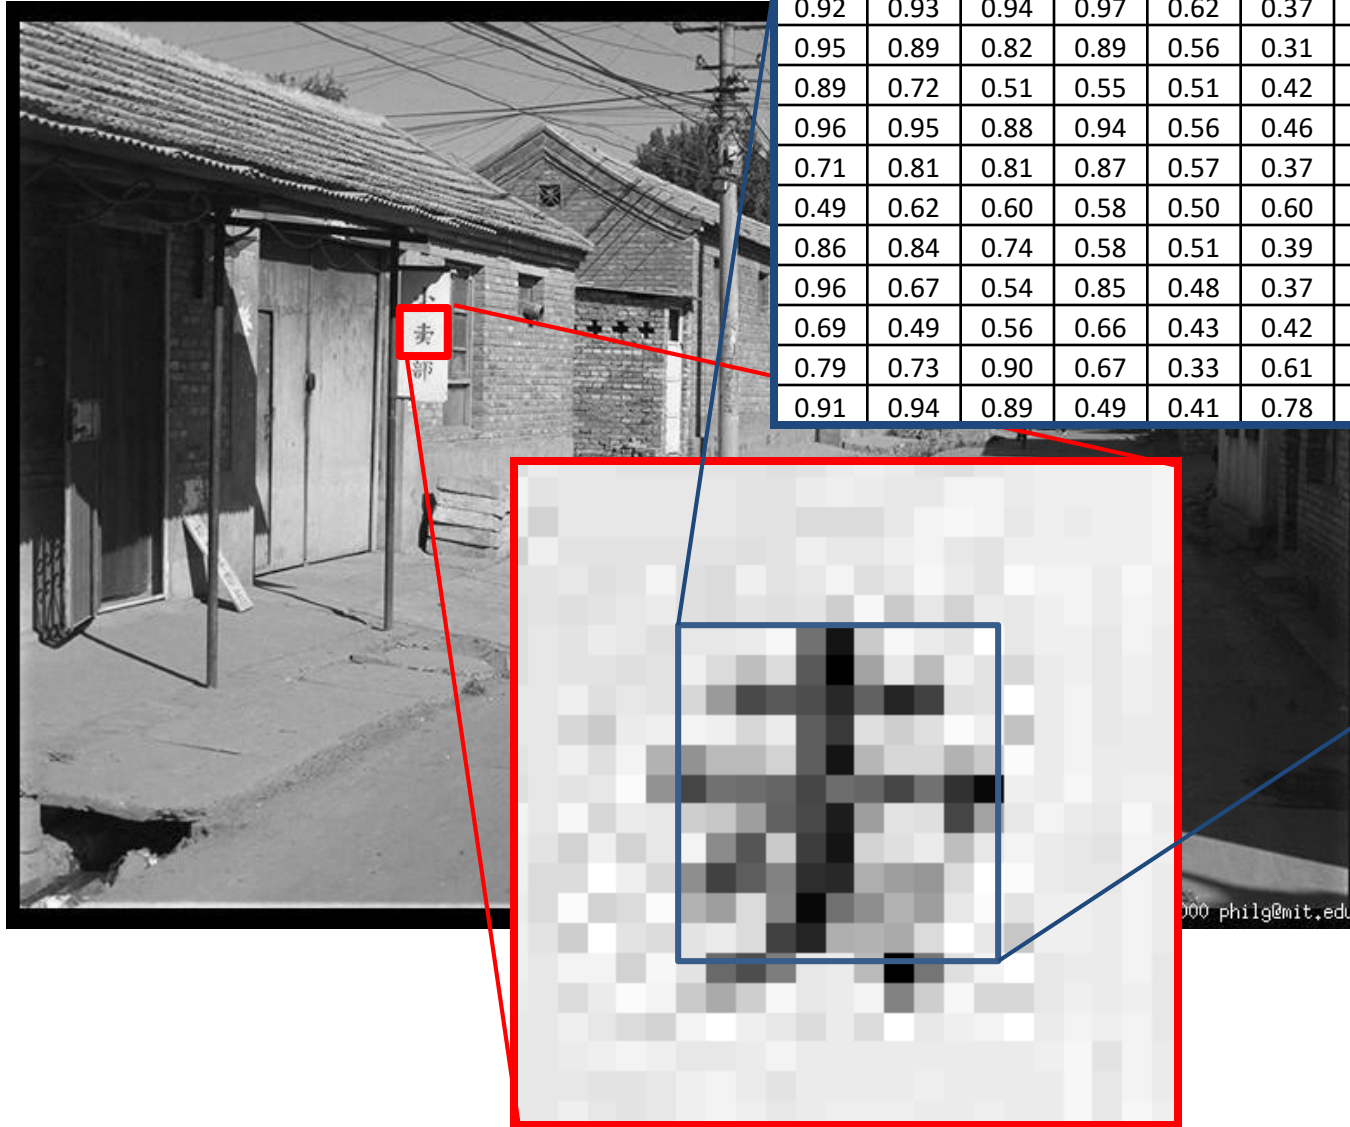
Only intensity shown – constant color

Most information in intensity



Original image

Back to grayscale intensity



0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Wrap up: Why do we care about cameras and eyes?

Inside Tesla as Elon Musk Pushed an Unflinching Vision for Self- Driving Cars

The automaker may have undermined safety in designing its Autopilot driver-assistance system to fit its chief executive's vision, former employees say.




Hardware choices have also raised safety questions. Within Tesla, some argued for pairing cameras with radar and other sensors that worked better in heavy rain and snow, bright sunshine and other difficult conditions. For several years, Autopilot incorporated radar, and for a time Tesla worked on developing its own radar technology. But three people who worked on the project said Mr. Musk had repeatedly told members of the Autopilot team that humans could drive with only two eyes and that this meant cars should be able to drive with cameras alone.

Schuyler Cullen, who oversaw a team that explored autonomous-driving possibilities at the South Korean tech giant Samsung, said in an interview that Mr. Musk's cameras-only approach was fundamentally flawed. "Cameras are not eyes! Pixels are not retinal ganglia! The F.S.D. computer is nothing like the visual cortex!" said Mr. Cullen, a computer vision specialist who now runs a start-up that is building a new kind of camera-based sensor.

Amnon Shashua, chief executive of Mobileye, a former Tesla supplier that has been testing technology that is similar to the electric-car maker's, said Mr. Musk's idea of using only cameras in a self-driving system could ultimately work, though other sensors may be needed in the short term. He added that Mr. Musk might exaggerate the capabilities of the company's technology, but that those statements shouldn't be taken too seriously.

Another driving system with human level eyesight








https://www.reddit.com/r/teslamotors/comments/nrs8kf/you_think_ice_cream_truck_stop_signs_are_a_problem/

- On the whole, cameras *are* a reasonable analogy for eyes. They do capture sufficient information for safe driving 99.9% of the time.



gets G's halfway over and then tries to go back to

- On the whole, cameras *are* a reasonable analogy for eyes. They do capture sufficient information for safe driving 99.9% of the time.
 - Imagine remote controlling a vehicle based on a camera feed.
- *But* the computer vision and machine learning methods that interpret the camera images *are not* yet a reasonable analogy for the human brain.

Next class: Interest points and corners