1950

Future Vision

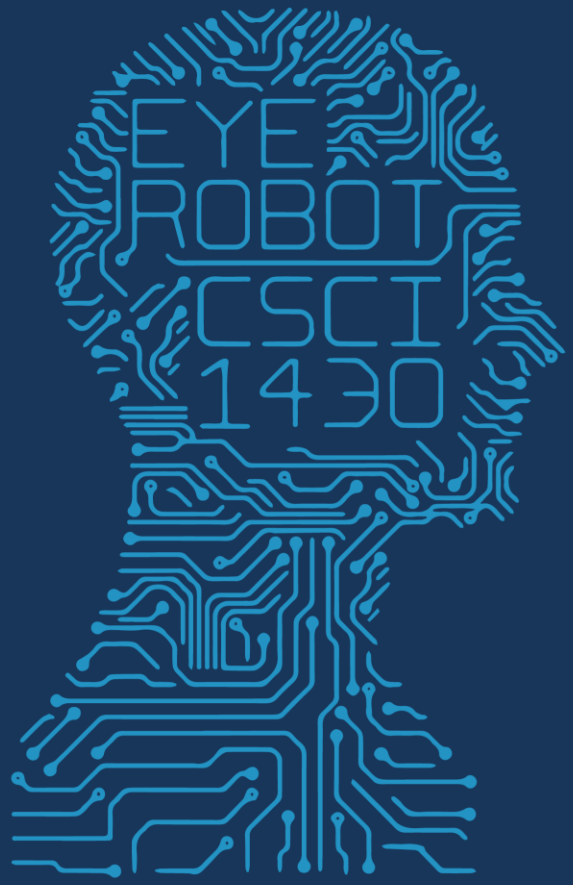

2020

Computer Vision

Project 0 Q3 – CV Superpowers

Instantly identify people, objects, places, plant and animal species, …
“I wouldn’t need to have a cheat sheet during exams”

Read people’s minds. Infer thoughts from emotions.

Take a picture of a product and it takes you to a site you can purchase the item, like an Amazon listing.

“Try to get on the Price is Right to become the most dominant Price is Right player ever, and win $10,000s worth of consumer products.”

Project 0 Q3 – CV Food Superpowers

"The ability to tell if a picture of an item is of a hot dog or not… I'd use it to distinguish things that are hot dogs from things that are not hot dogs."

Determine the recipe of a food by looking at it.
"Will use it to make food that my mom likes."

Reliably tell the difference between choco-chip cookies and oatmeal raisin cookies *from any distance*.

"I would never have my expectations let down from eating an oatmeal raisin cookie expecting it to be a chocolate cookie." :' (

Project 0 Q3 – Real Heroes

Overlay myself on objects like a Snapchat filter but in real life.

"Mess around and turn into different famous people to live life in their shoes…

Then master my powers …. use them for good
… scare politicians into doing right… like Batman."

"Detect a person's history by scanning their face, having access to any known records of them.

Would use it to travel the world and figure out who to be friends with.
Possibly help find criminals."

Recognize precise 3D movement of a body and transfer that data to a robot, e.g., for medial tasks.

Telesurgery

da Vinci robotic surgery system have been successfully tested remotely

Project 0 Q3 – Contentious Issues

- Hiring for jobs; interview screening
- Autonomous vehicle safety
- Face reconstruction / recognition.
  - Accuracy and bias
  - Real-time mass surveillance - in and out of towns, buildings, buying a cellphone, jaywalking
  - Face ‘swaps’, editing
- Policing; crime prediction
- Healthcare; medical imaging; privacy
- Industrial automation; jobs
- Delivery drones; privacy
- Military; autonomous weapons

Project 0 Q3 – Contentious Issues

Educators using CV to take attendance, improve campus security, remove from logistic burdens.

Immature --- might give inaccurate or biased results; questions of student privacy

Project 0 Q3 – Contentious Issues

Companion Labs in SF managed to use CV to train dogs and automatically *launch* treats.

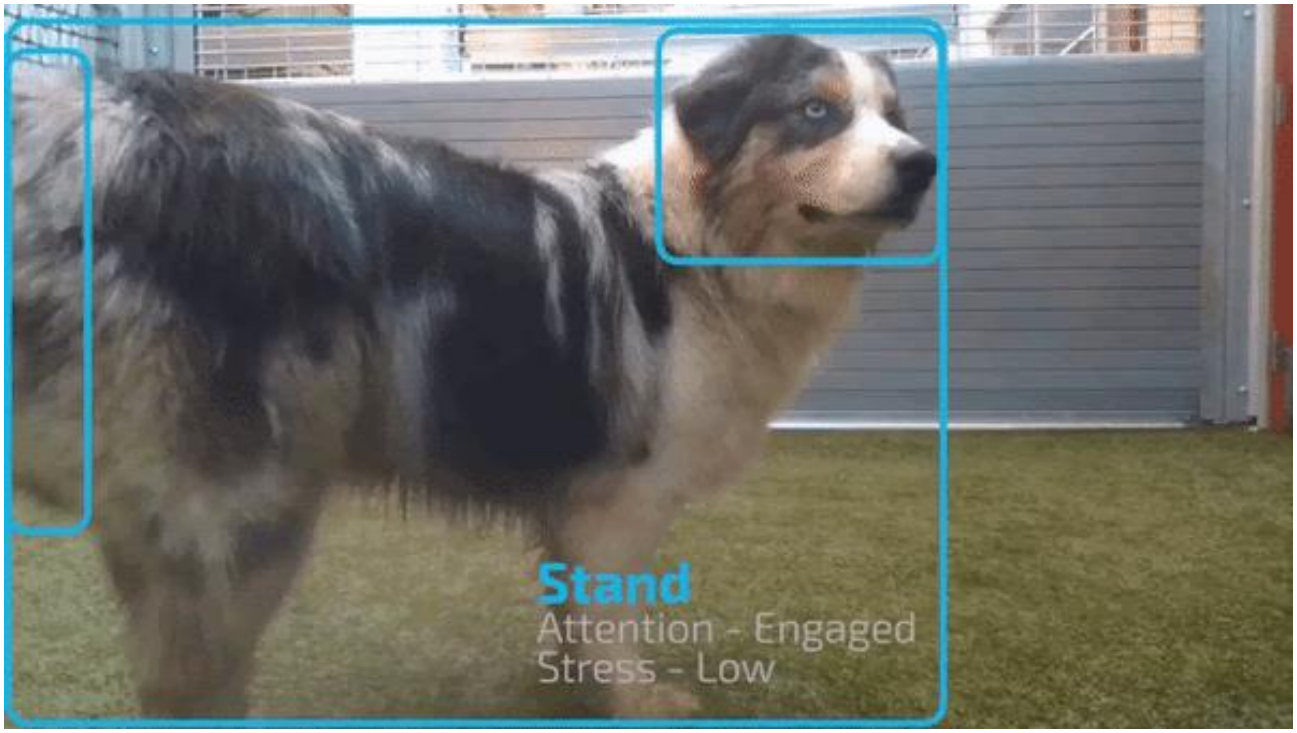

Next Classes

- Spatial frequency
- Fourier transform and frequency domain
  - Frequency view of filtering
  - Hybrid images
  - Sampling
- Reminder: Textbook
  - Today's lecture covers material in Klette 1, 2; Szeliski 3.4

Slide: Hoiem

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

Image: http://www.flickr.com/photos/igorms/136916757/

Slide: Hoiem

Sampling

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

Image: http://www.flickr.com/photos/igorms/136916757/

Subsampling by a factor of 2

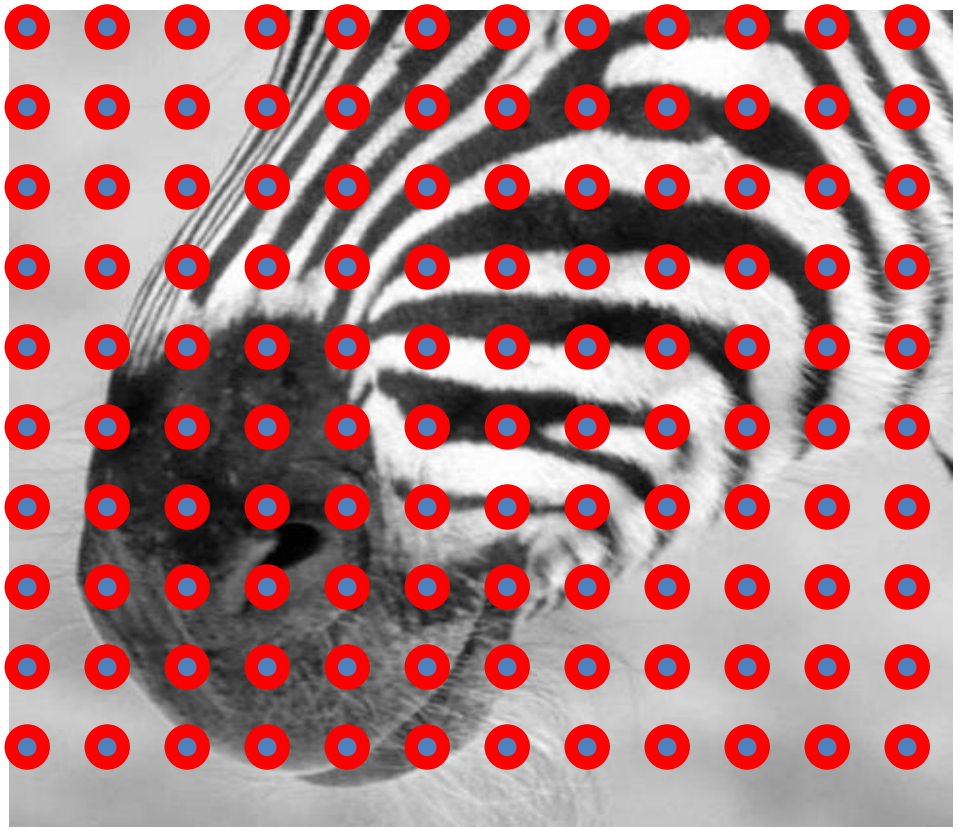

Throw away every other row and column to create a 1/2 size image

Figure from David Forsyth

A ‘bar’ in the big images is a hair on the zebra’s nose; in smaller images, a stripe; in the smallest, the animal’s nose

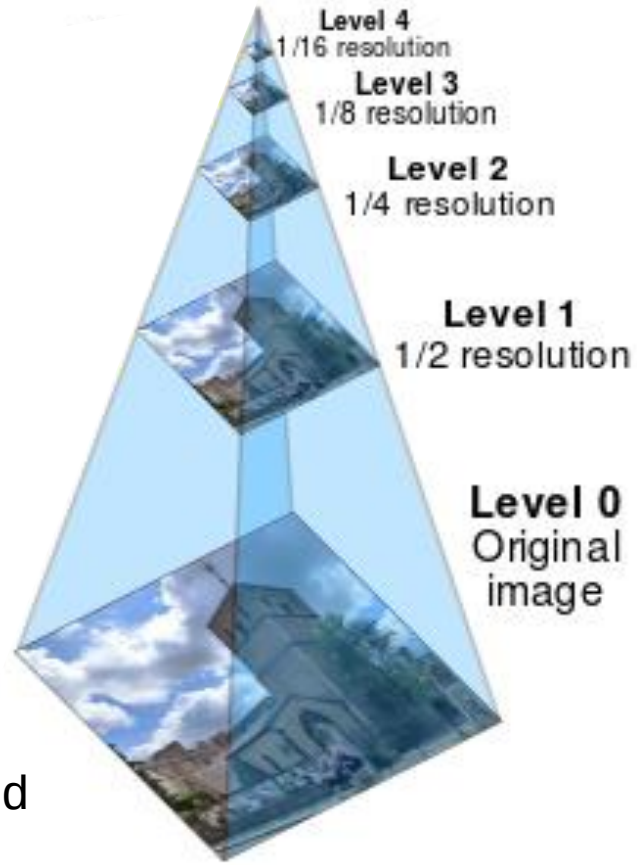

Image pyramid

Algorithm for downsampling by factor of 2

1. Start with image of w x h
2. Sample every other pixel
   - im_small = image[::2:, ::2]
3. Repeat until im_small is 1 pixel large.

Numpy syntax:
::2 -> start at 0, end at 'end', increase every 2, until the end.
e.g.,
0,2,4,6,…,w

(if w is not even, then this goes to w-1)

Image sub-sampling

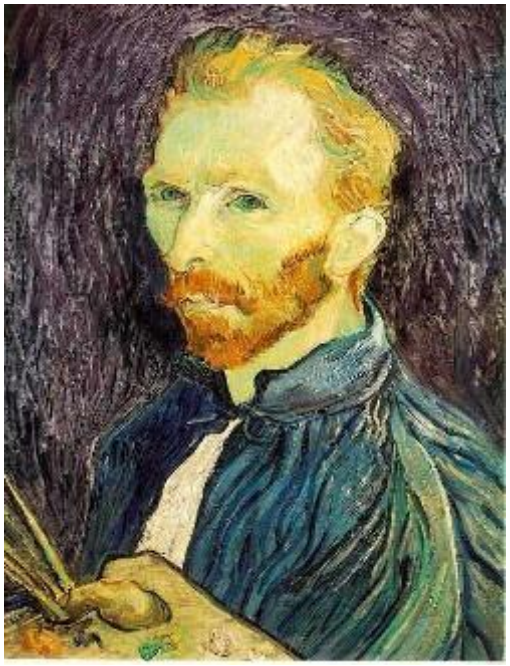

Throw away every other row and column to create a 1/2 size image.

1/4

1/8

Subsampling without filtering

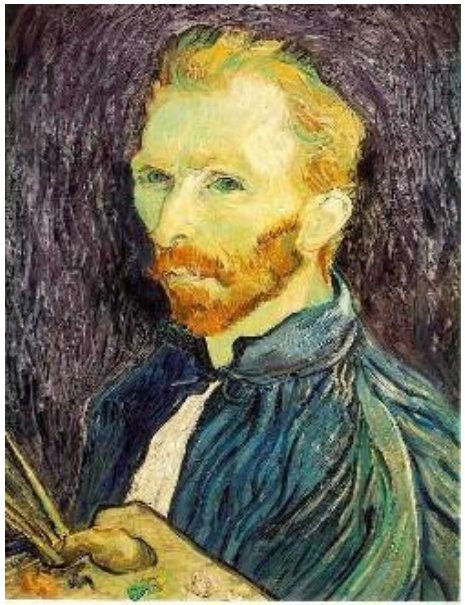

1/2

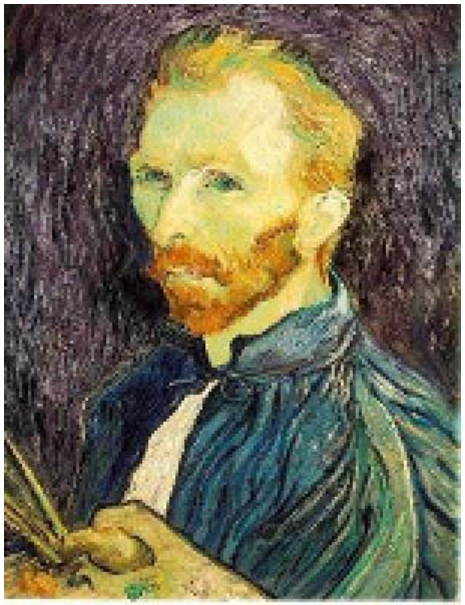

1/4 (2x subsample)

1/8 (4x subsample)

Steve Seitz

Sampling and aliasing

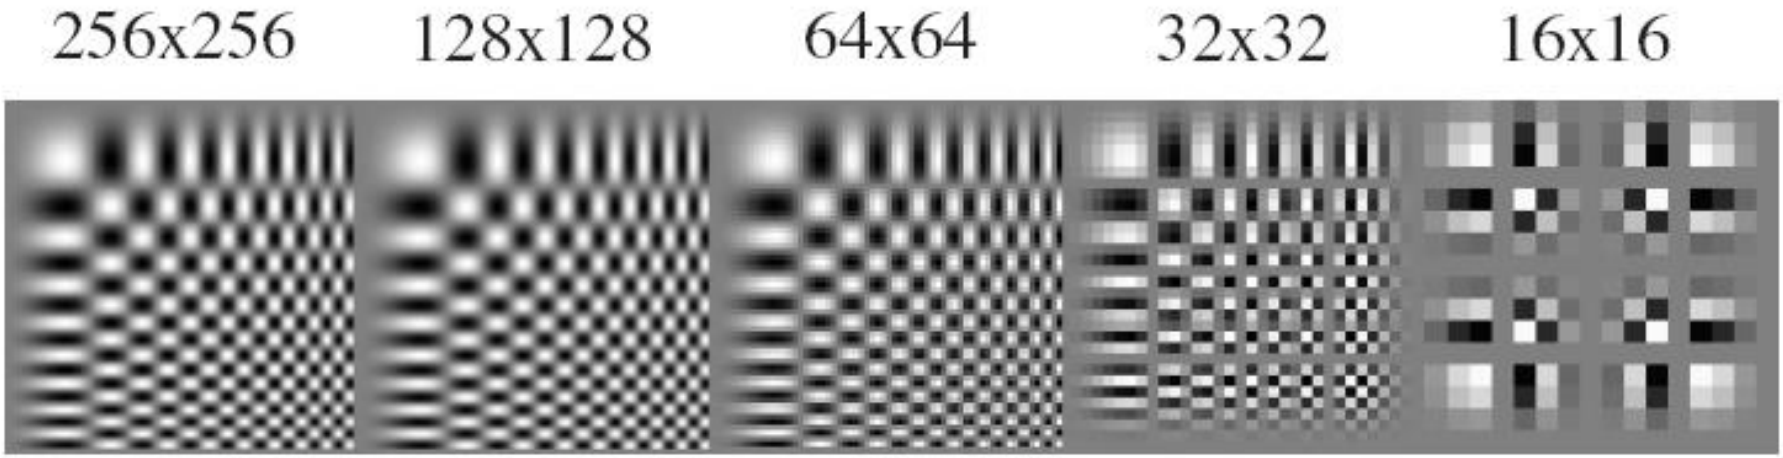

Aliasing problem

- 1D example (sinewave):

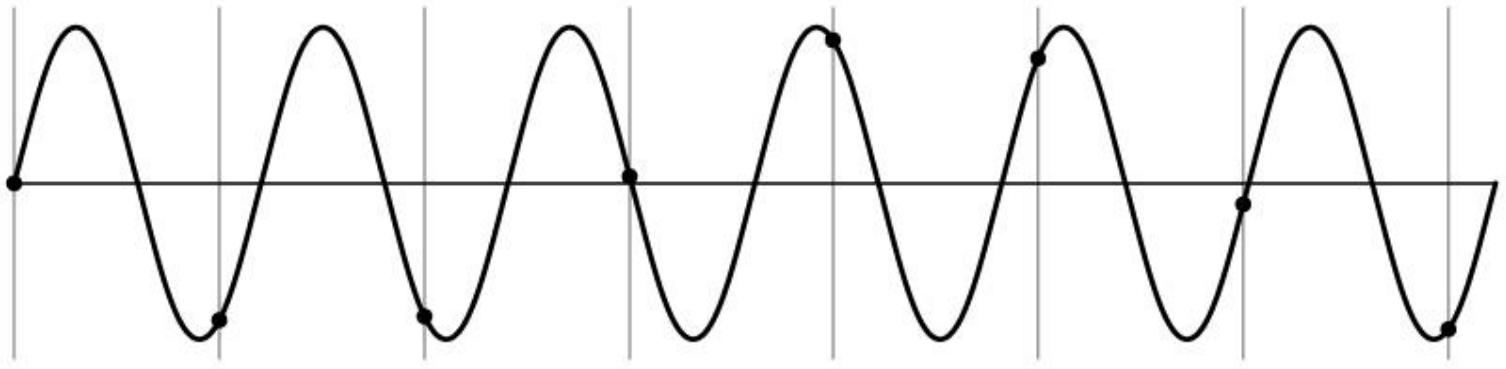

Source: S. Marschner

Aliasing problem

- 1D example (sinewave):

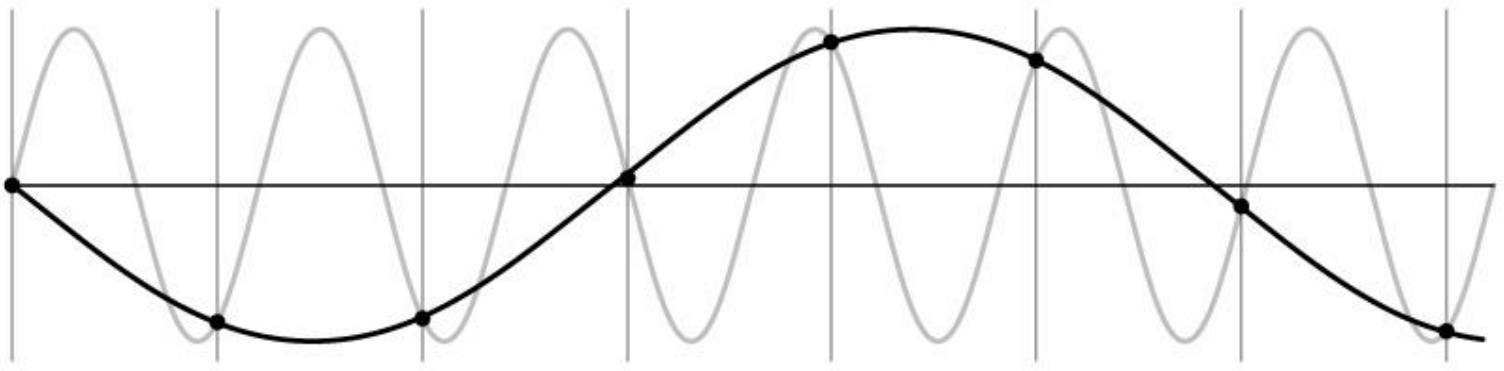

Source: S. Marschner

Aliasing problem

- Sub-sampling may be dangerous….
- Characteristic errors may appear:
  - “car wheels rolling the wrong way in movies”
  - “checkerboards disintegrate in graphics”
  - “striped shirts look funny on color television”
    - Moiré patterns

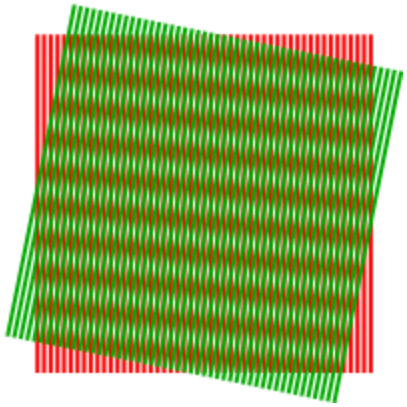

Source: D. Forsyth; Wikipedia

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):

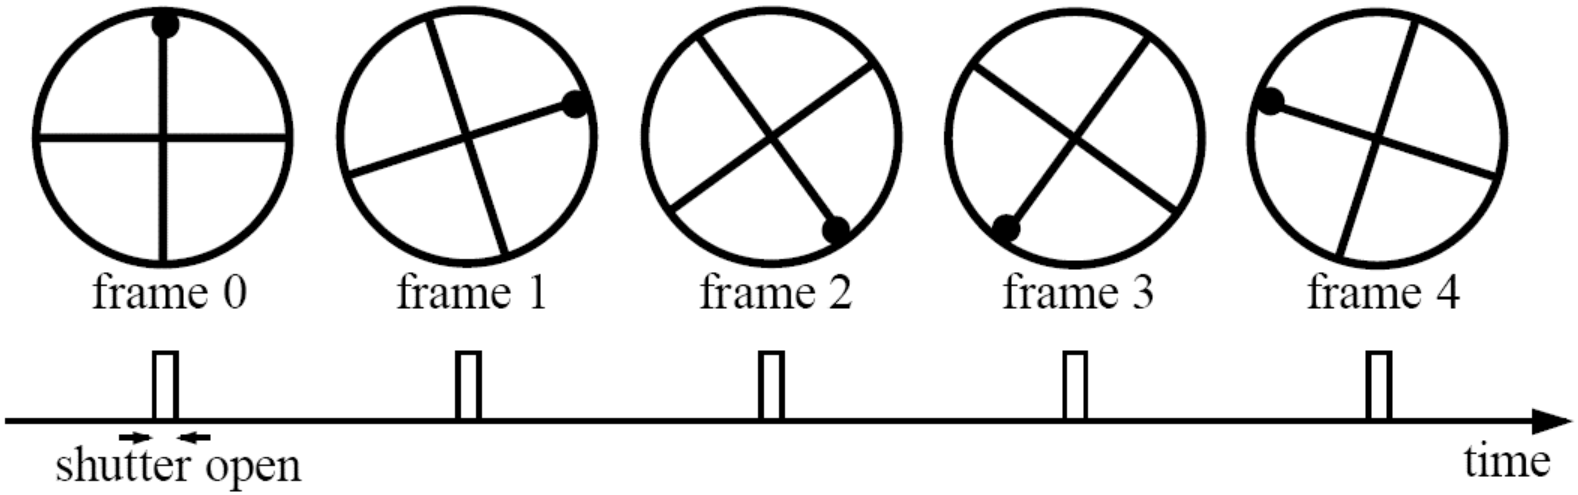

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Steve Seitz

Videos

[YouTube; JoinBuzzirk; phrancque]

Aliasing in graphics

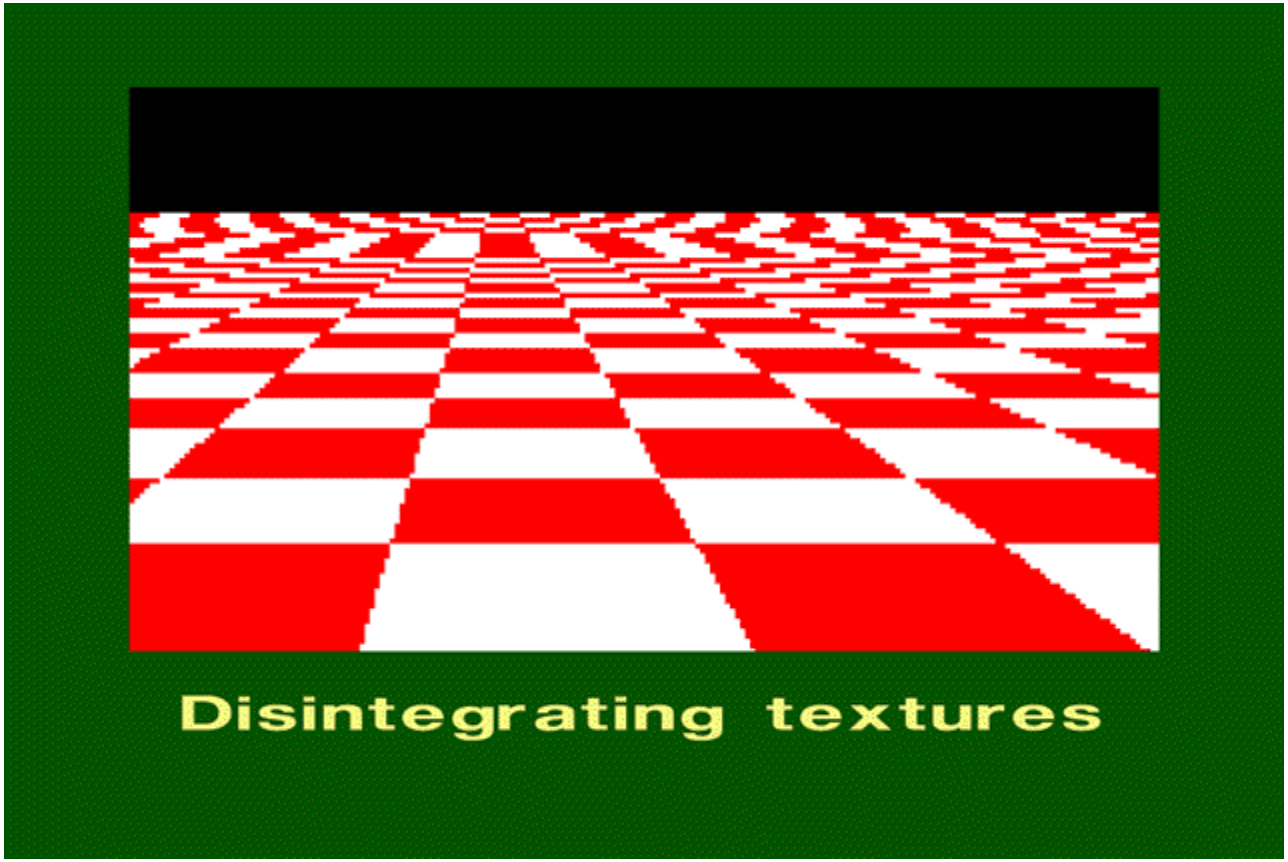

Source: A. Efros

Aliasing and Moiré patterns

Gong 96, 1932, Claude Tousignant, Musée des Beaux-Arts de Montréal

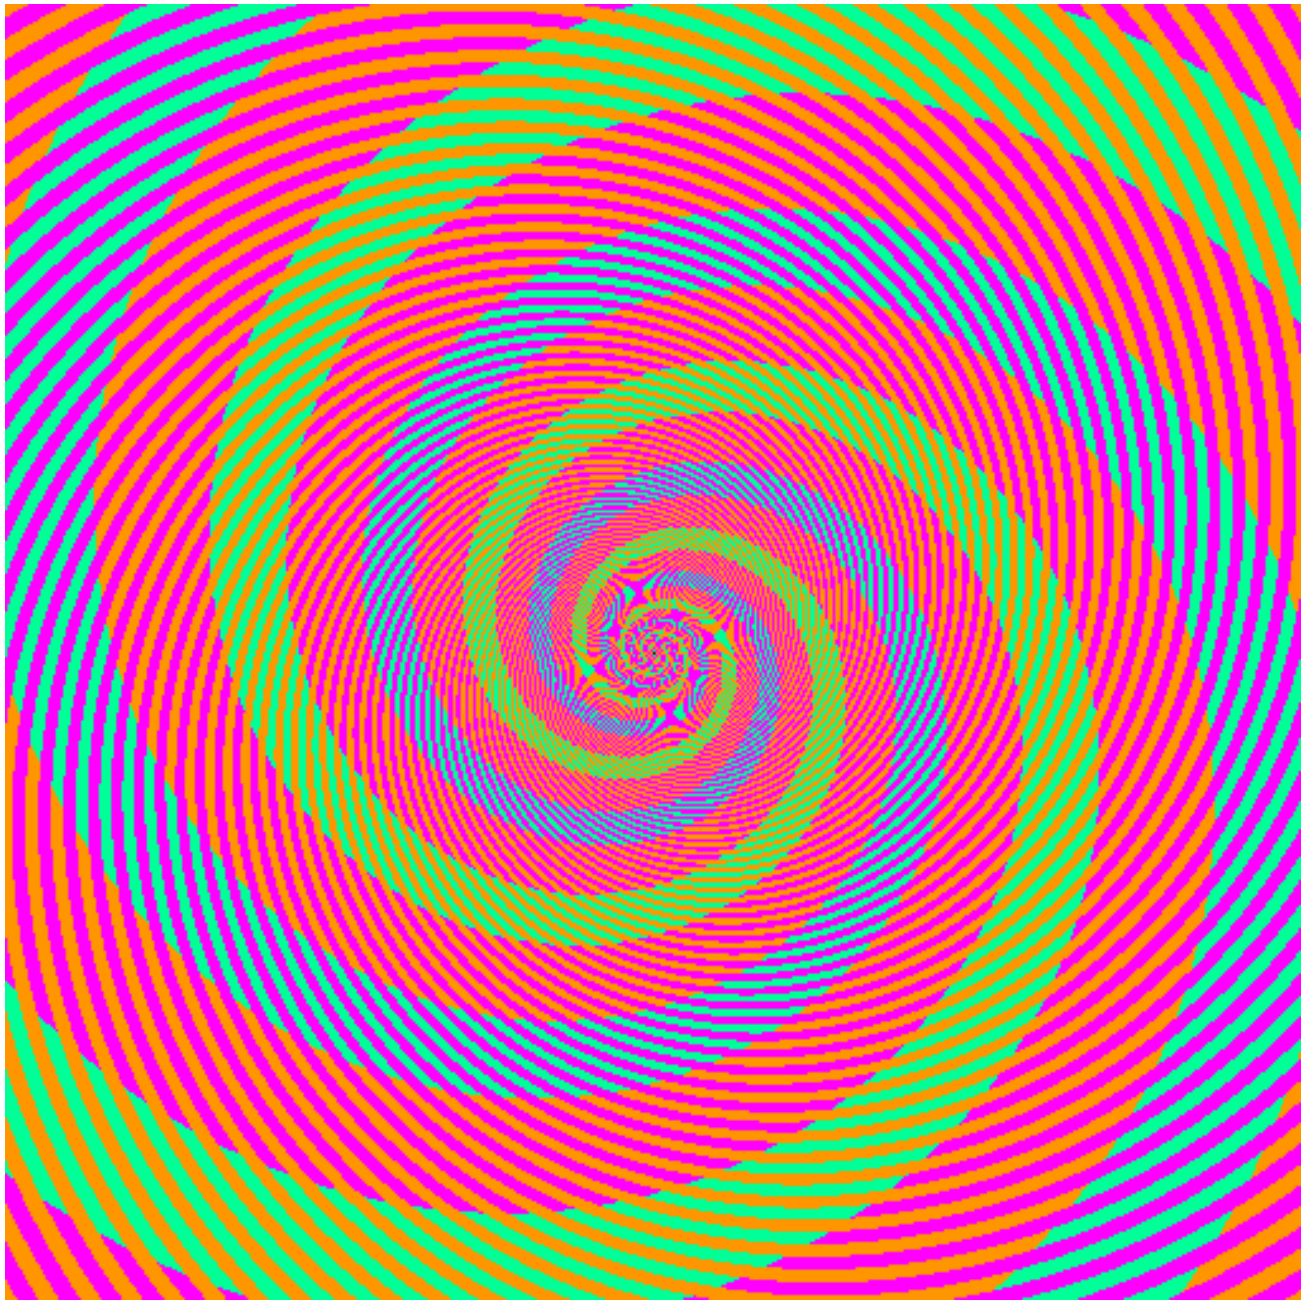

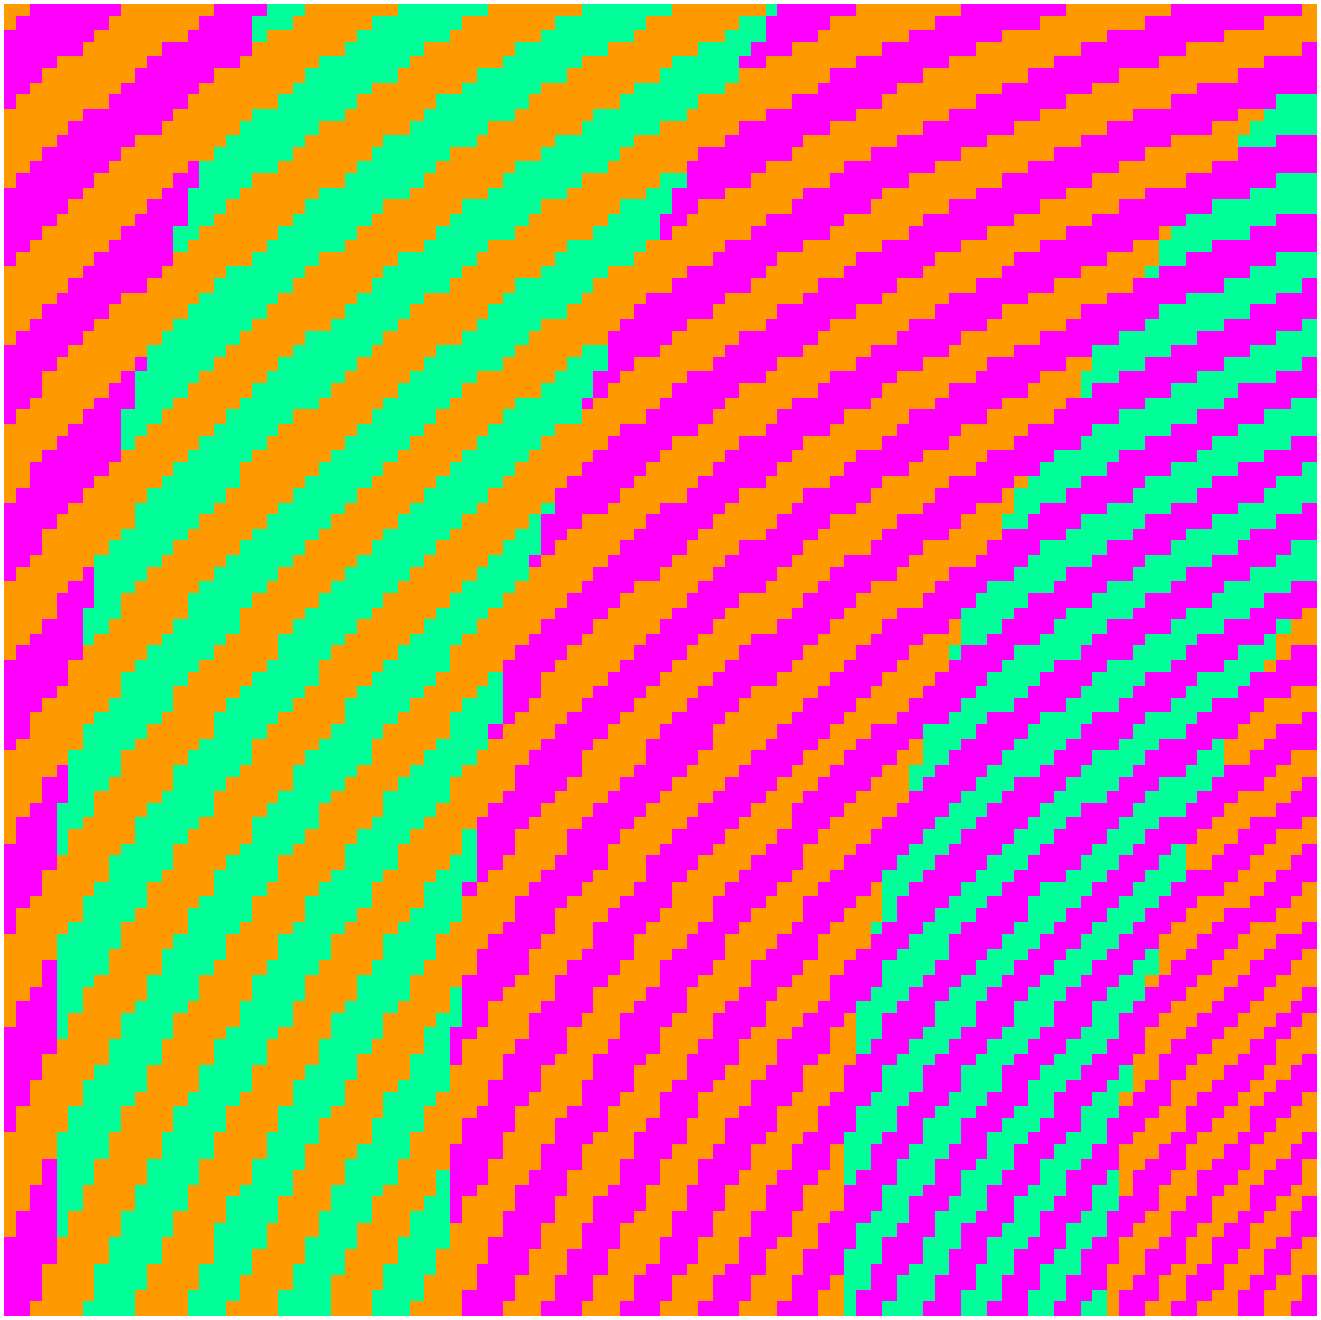

The blue and green colors are actually the same

http://blogs.discovermagazine.com/badastronomy/2009/06/24/the-blue-and-the-green/

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 allows us to reconstruct the original perfectly from the sampled version

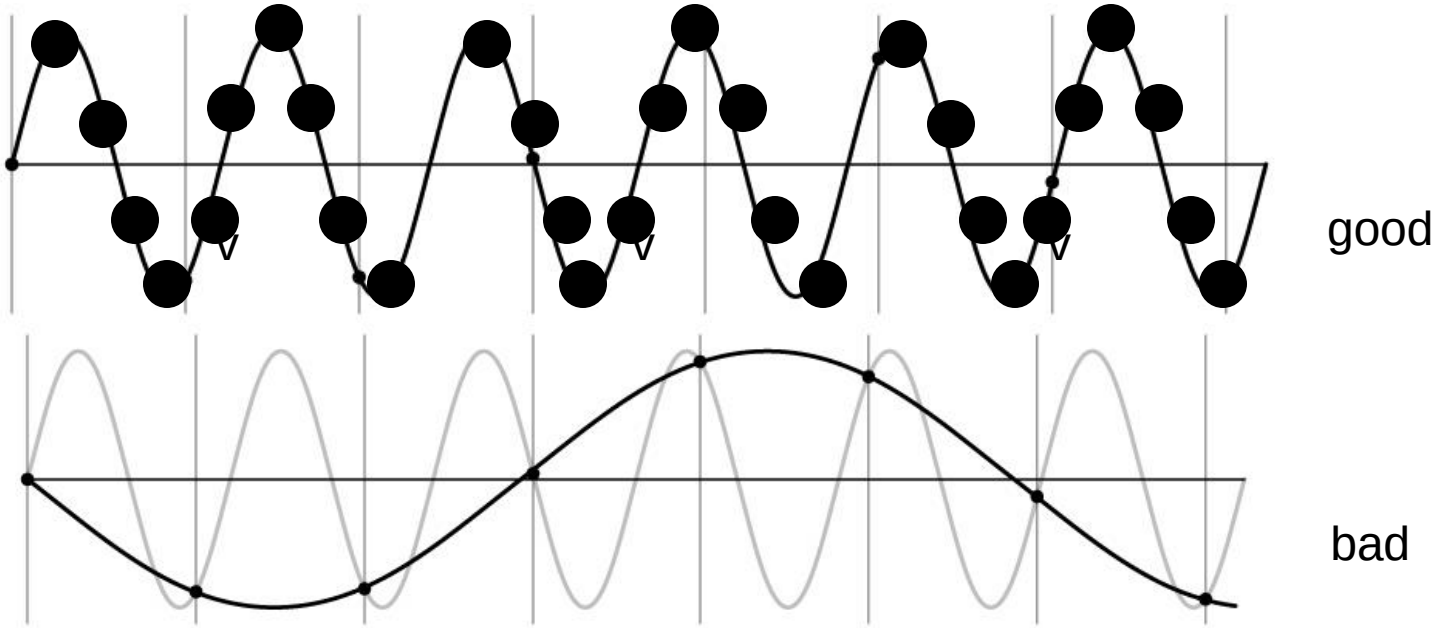

How to fix aliasing?

Solutions?

Better sensors

Solutions:

- Sample more often

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 (*low pass*) filter

Anti-aliasing

256x256 128x128 64x64 32x32 16x16

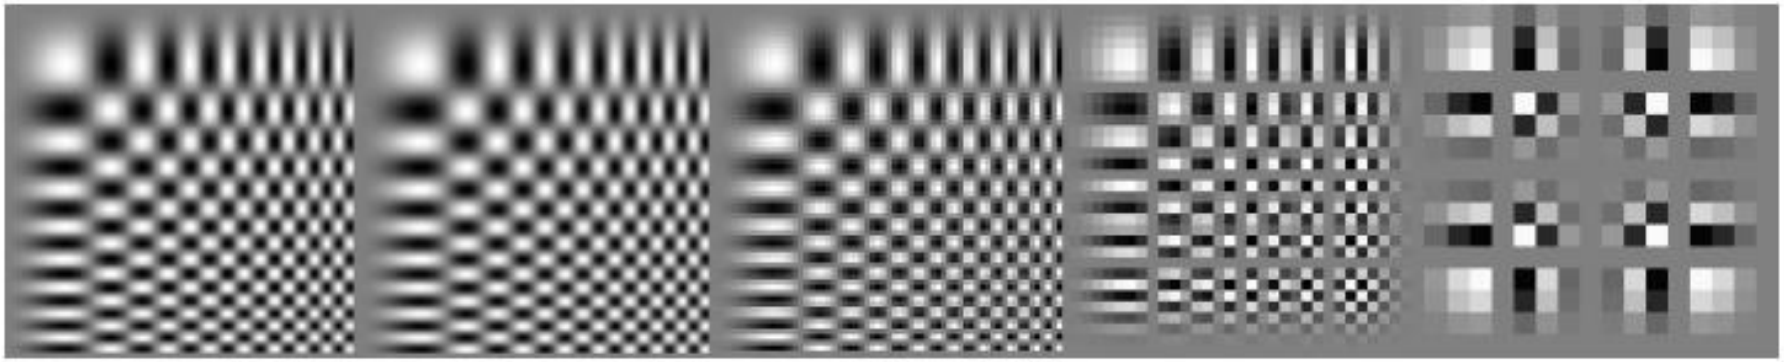

256x256 128x128 64x64 32x32 16x16

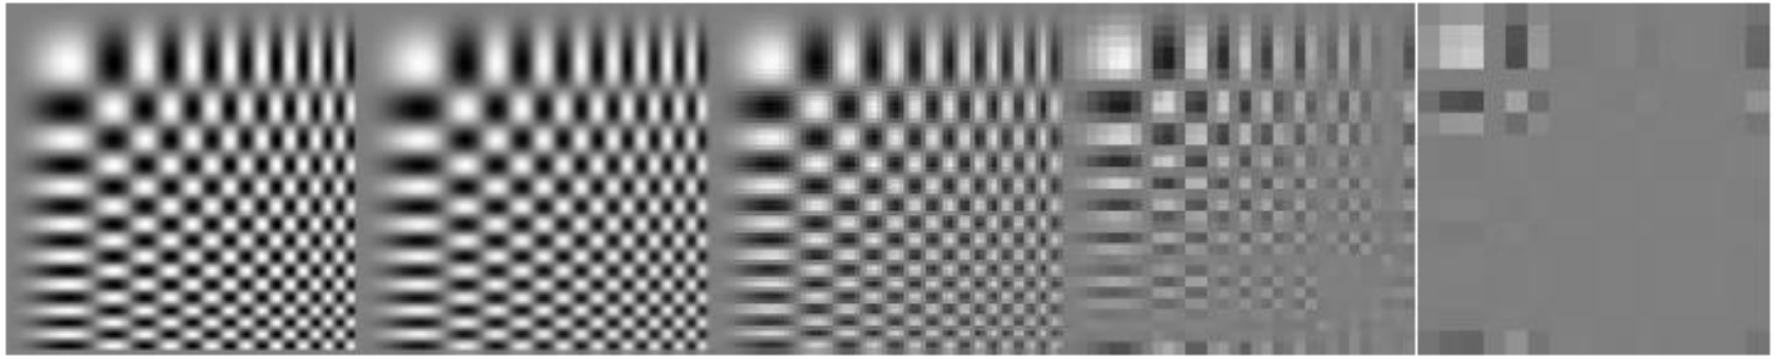

Forsyth and Ponce 2002

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);

Subsampling without filtering

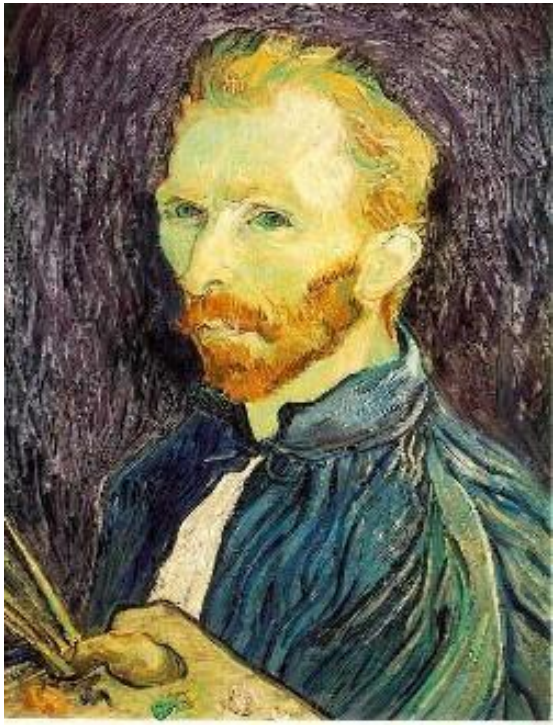

1/2

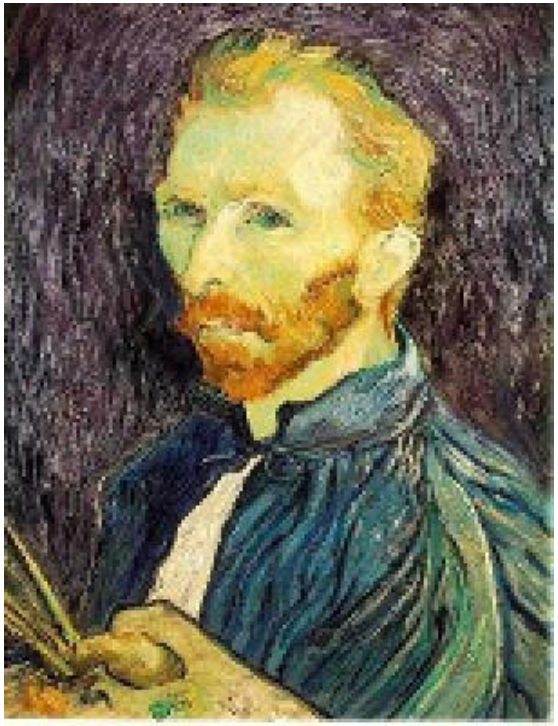

1/4 (2x subsample)

1/8 (4x subsample)

Steve Seitz

Subsampling with Gaussian pre-filtering

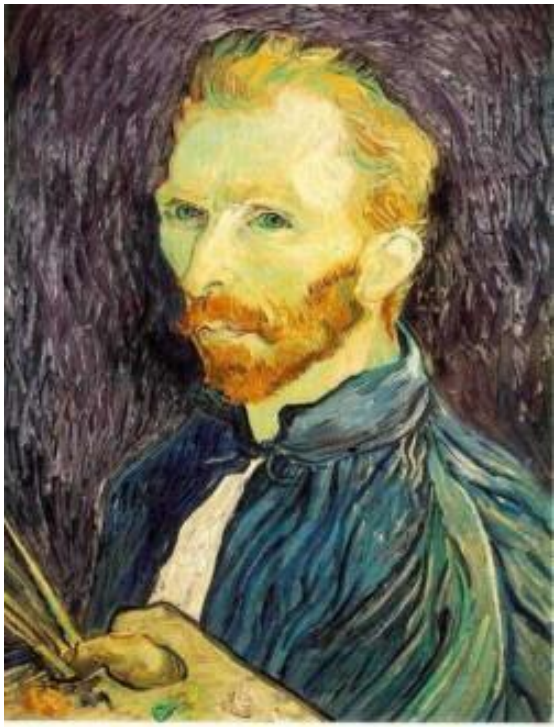

Gaussian 1/2

G 1/4

G 1/8

Gaussian Pyramid [Burt and Adelson, 1983]

Steve Seitz

Gaussian Pyramid [Burt and Adelson, 1983]

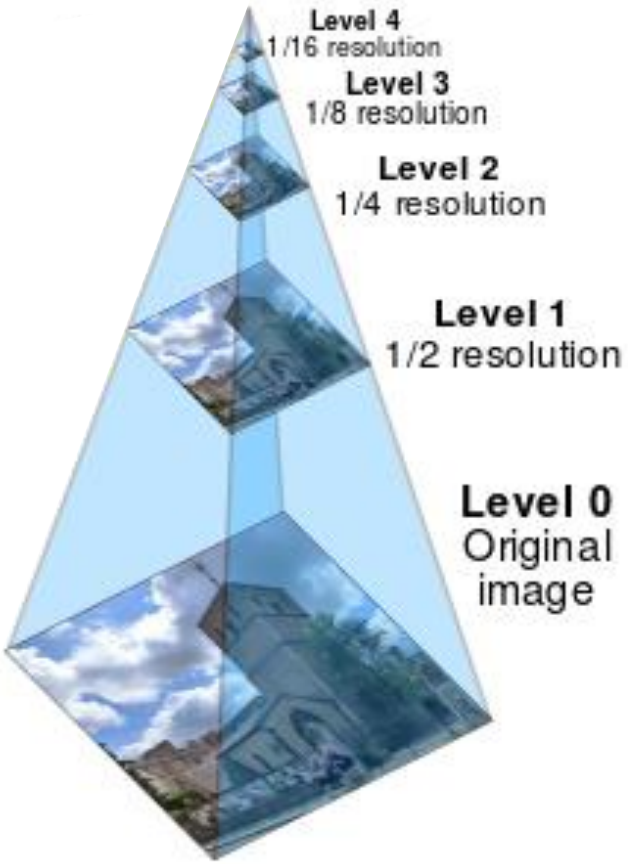

Hybrid Images

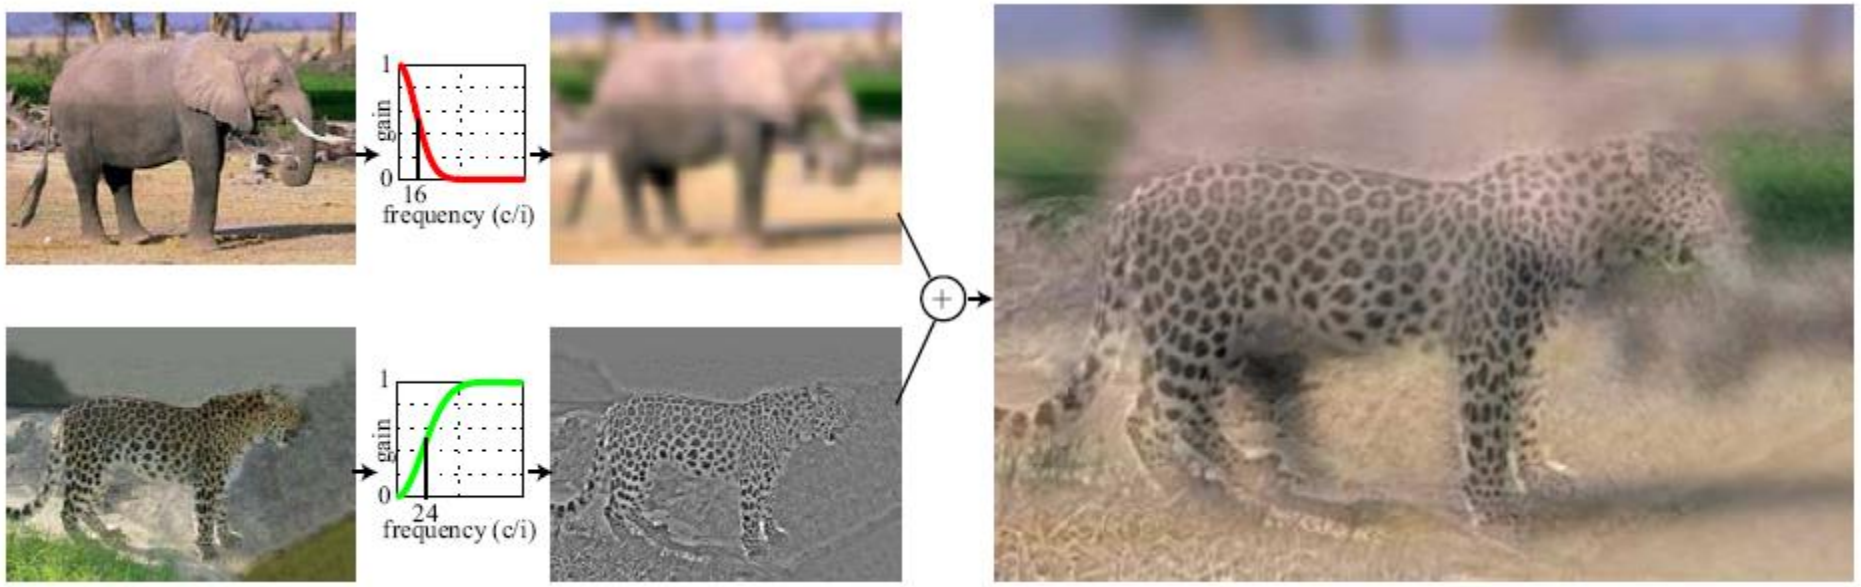

- A. Oliva, A. Torralba, P.G. Schyns, "Hybrid Images," SIGGRAPH 2006

Why do we get different, distance-dependent interpretations of hybrid images?

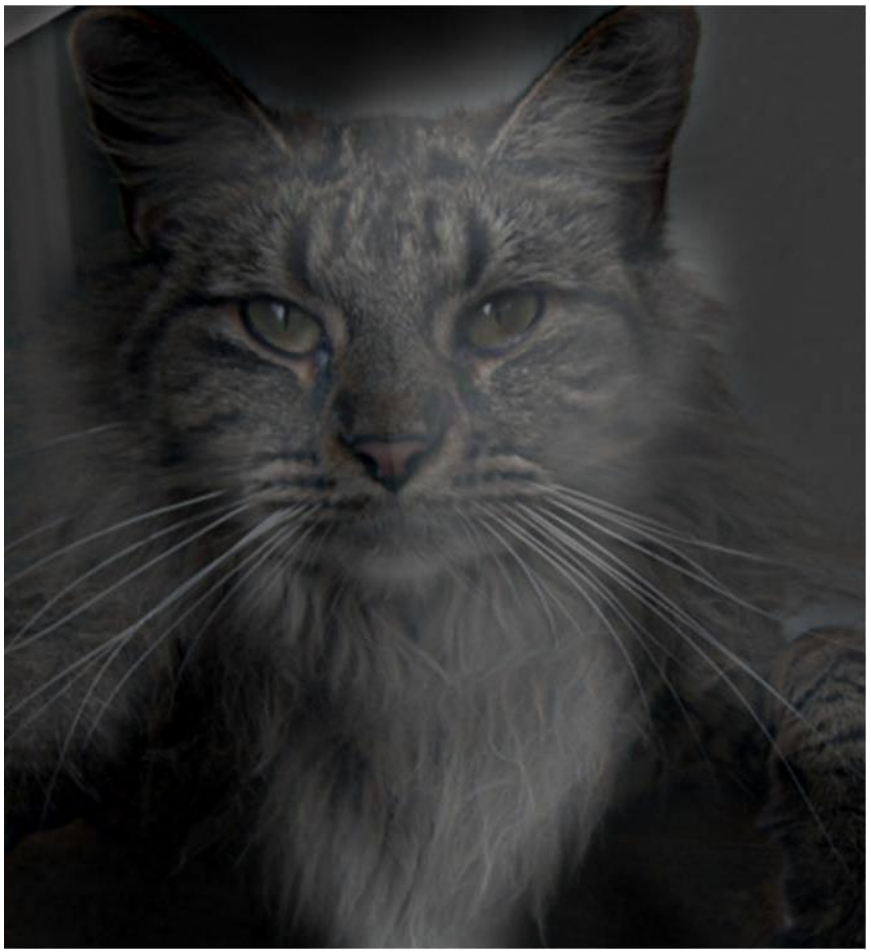

?

Slide: Hoiem

Clues from Human Perception

Early visual processing in human perception
filters for orientations and scales of frequency.

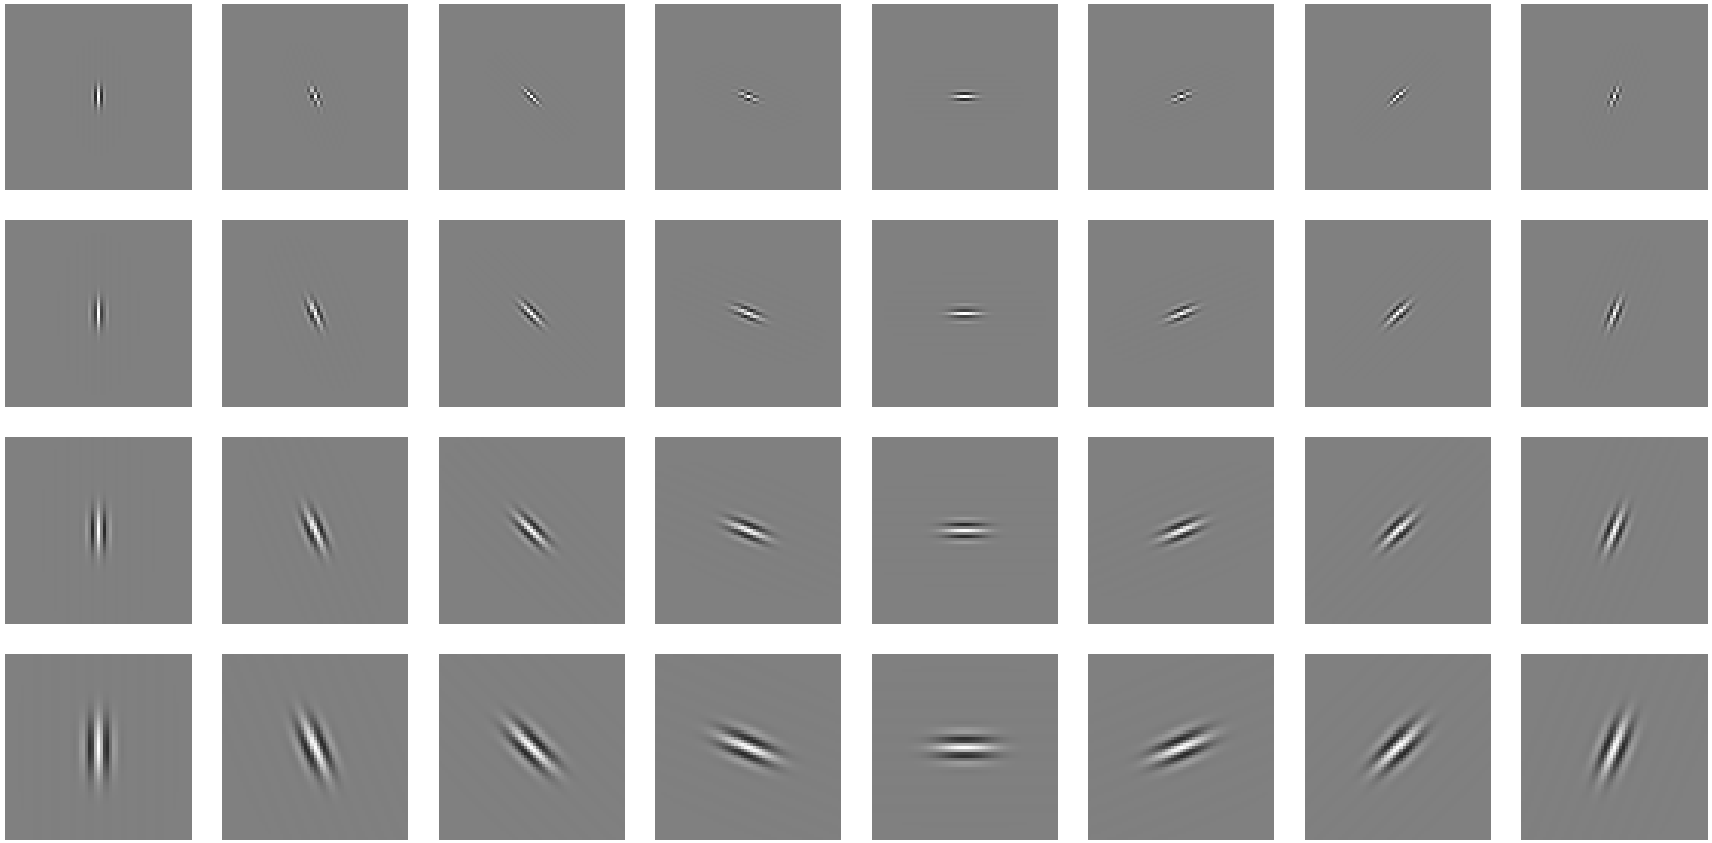

Related to *Gabor filters:* sinusoids convolved with Gaussians

Campbell-Robson contrast sensitivity curve

Perceptual cues in the mid-high frequencies dominate perception.

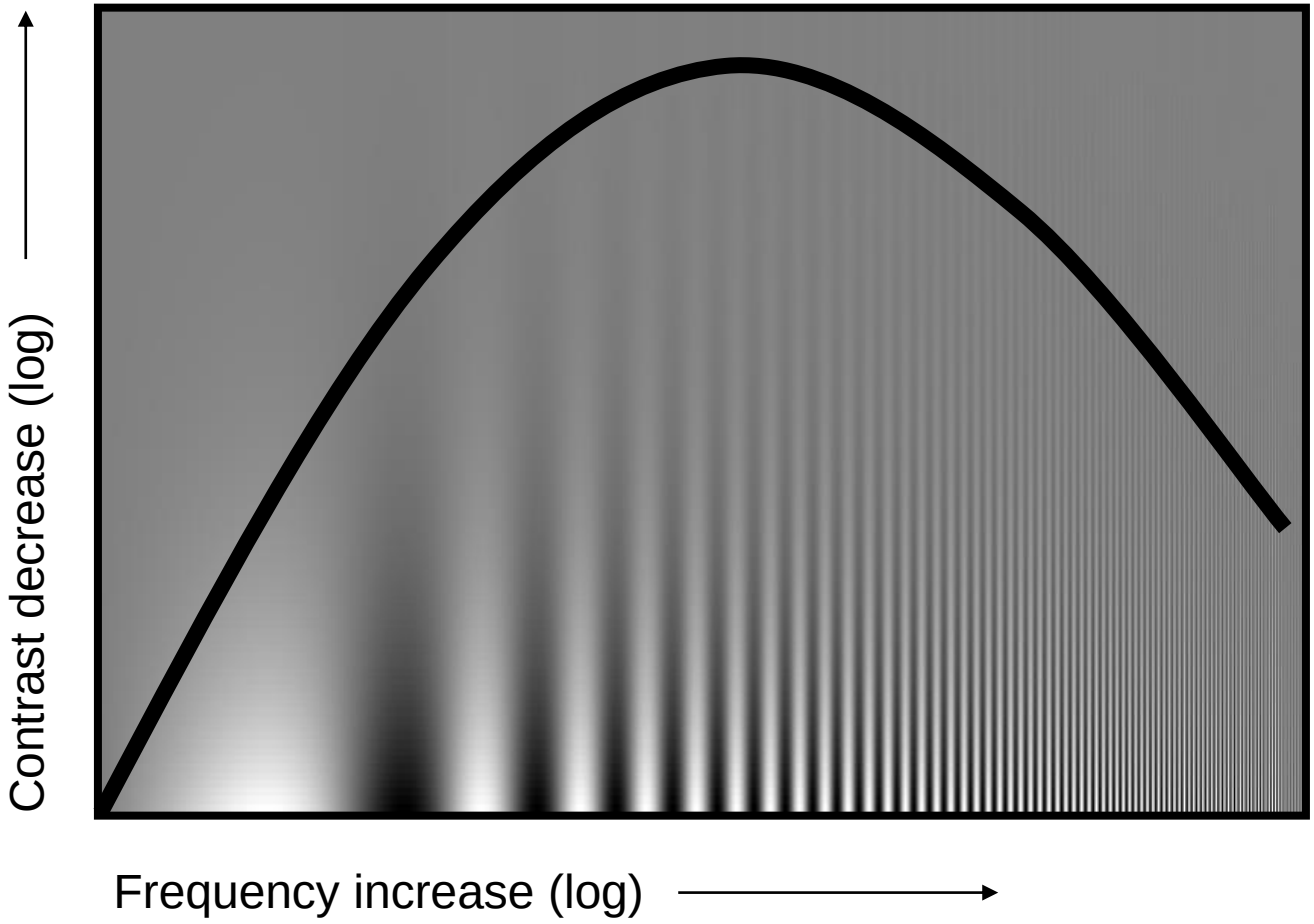

Application: Hybrid Images

When we see an image from far away, we are effectively subsampling it!

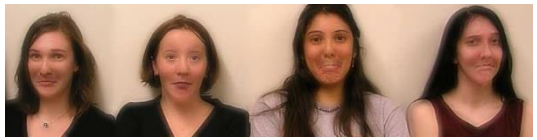

A. Oliva, A. Torralba, P.G. Schyns, SIGGRAPH 2006

Salvador Dali
"Gala Contemplating the Mediterranean Sea, which at 30 meters becomes the portrait of Abraham Lincoln", 1976

Salvador Dali invented Hybrid Images?

Salvador Dali
"Gala Contemplating the Mediterranean Sea, which at 30 meters becomes the portrait of Abraham Lincoln", 1976