Introduction to Computer Vision

Michael J. Black

Sept 2009

Lecture 7:

Linear filtering, smoothing and pyramids

Goals

- Linear filtering (cont.)
  - Foundations for asng1.
    - Problem 1, image pyramids
    - Problem 2, image derivatives
- Friday: correlation, features
  - Problems 3&4

Homework

- Assignment 1
  - Problems 1&2 due Monday Sept 28
  - Problems 3&4 due Oct 5.
  - Grad credit – do extra credit questions

Example: Smoothing by Averaging

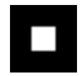

Linear systems

Basic properties. If T[.] is a linear operator, and a and b are scalars, then:

- homogeneity $T[a X] = a T[X]$
- Additivity $T[X_1+X_2] = T[X_1]+T[X_2]$
- superposition $T[aX_1+bX_2] = aT[X_1]+bT[X_2]$
- Linear system ⇔ superposition

- Examples:
  - matrix operations (additions, multiplication)
  - convolutions

Smoothing as Inference About the Signal

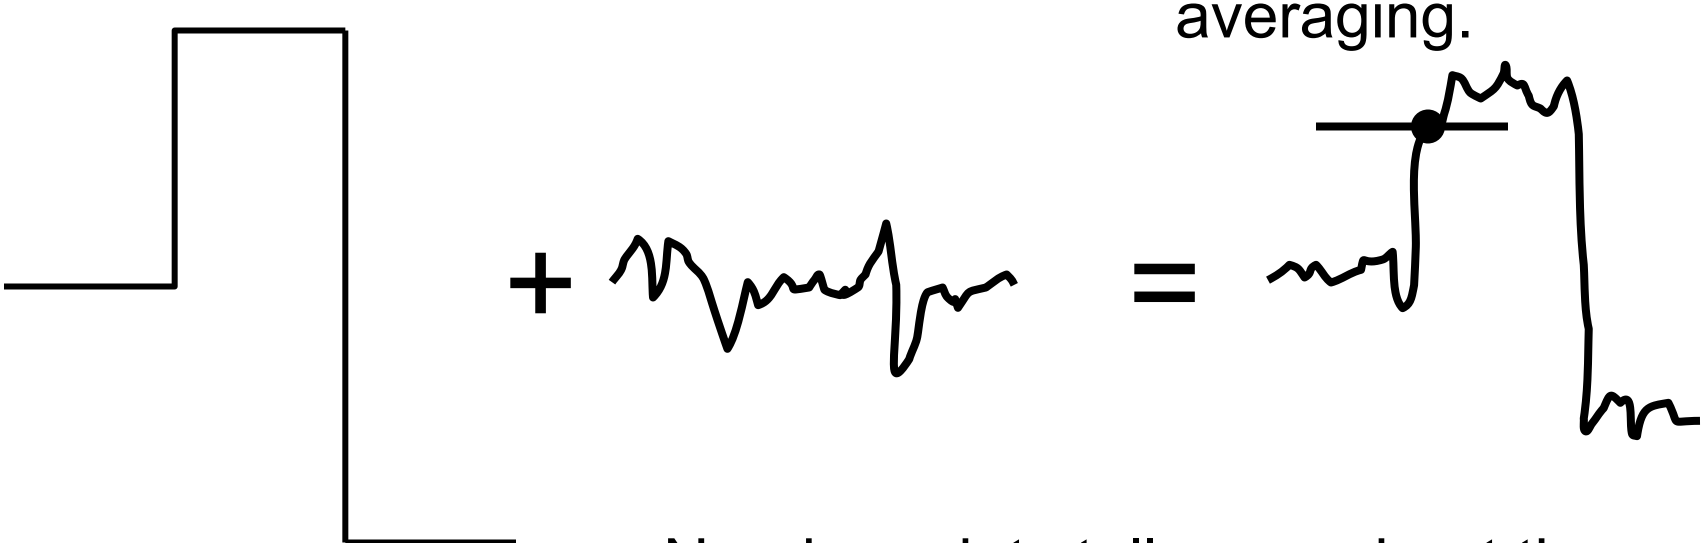

Nearby points tell more about the signal than distant ones.

Gaussian Averaging

- Rotationally symmetric.
- Weights nearby pixels more than distant ones.
  - This makes sense as probabilistic inference.

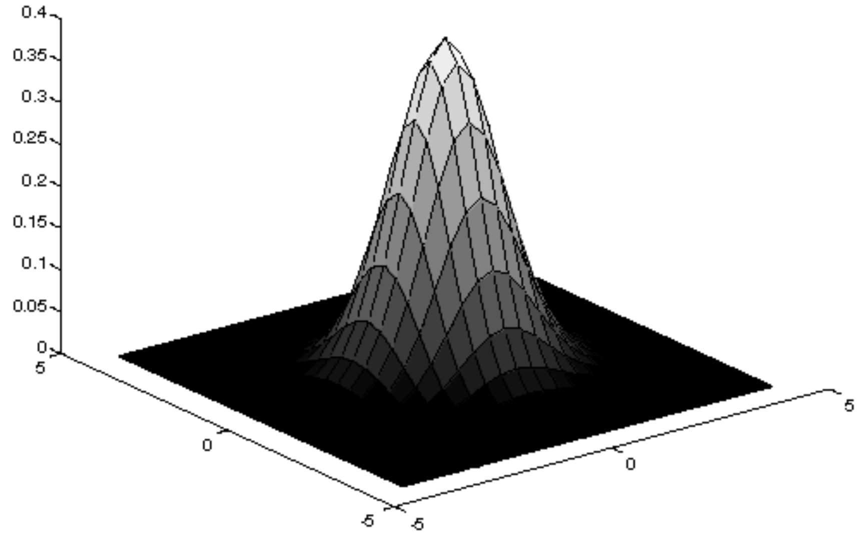

- A Gaussian gives a good model of a fuzzy blob

Ponce & Forsyth

An Isotropic Gaussian

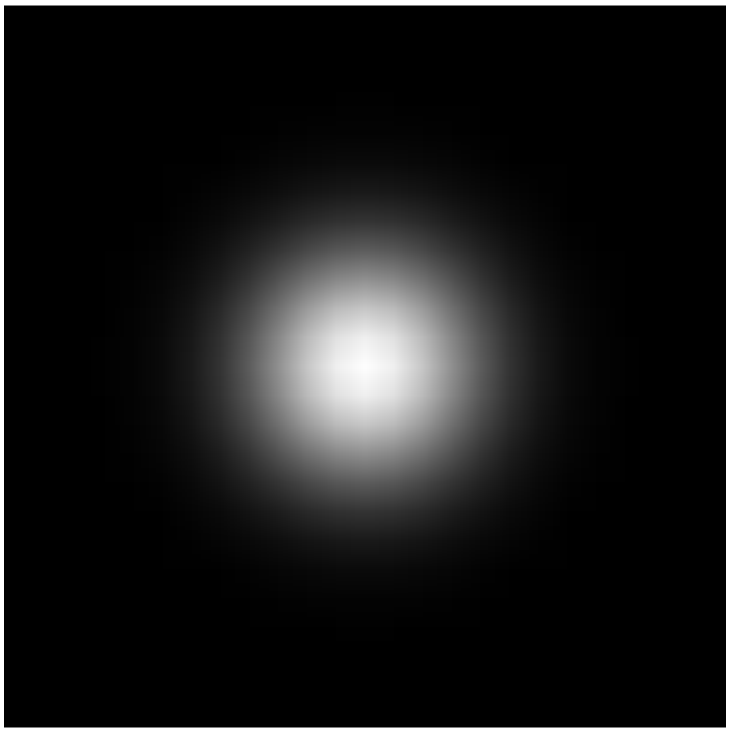

(0,0) is the center

- The picture shows a smoothing kernel proportional to

$$g(x,y) = \frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x^2+y^2)}{2\sigma^2}\right)$$

(which is a reasonable model of a circularly symmetric fuzzy blob)

Ponce & Forsyth

Smoothing with a Gaussian

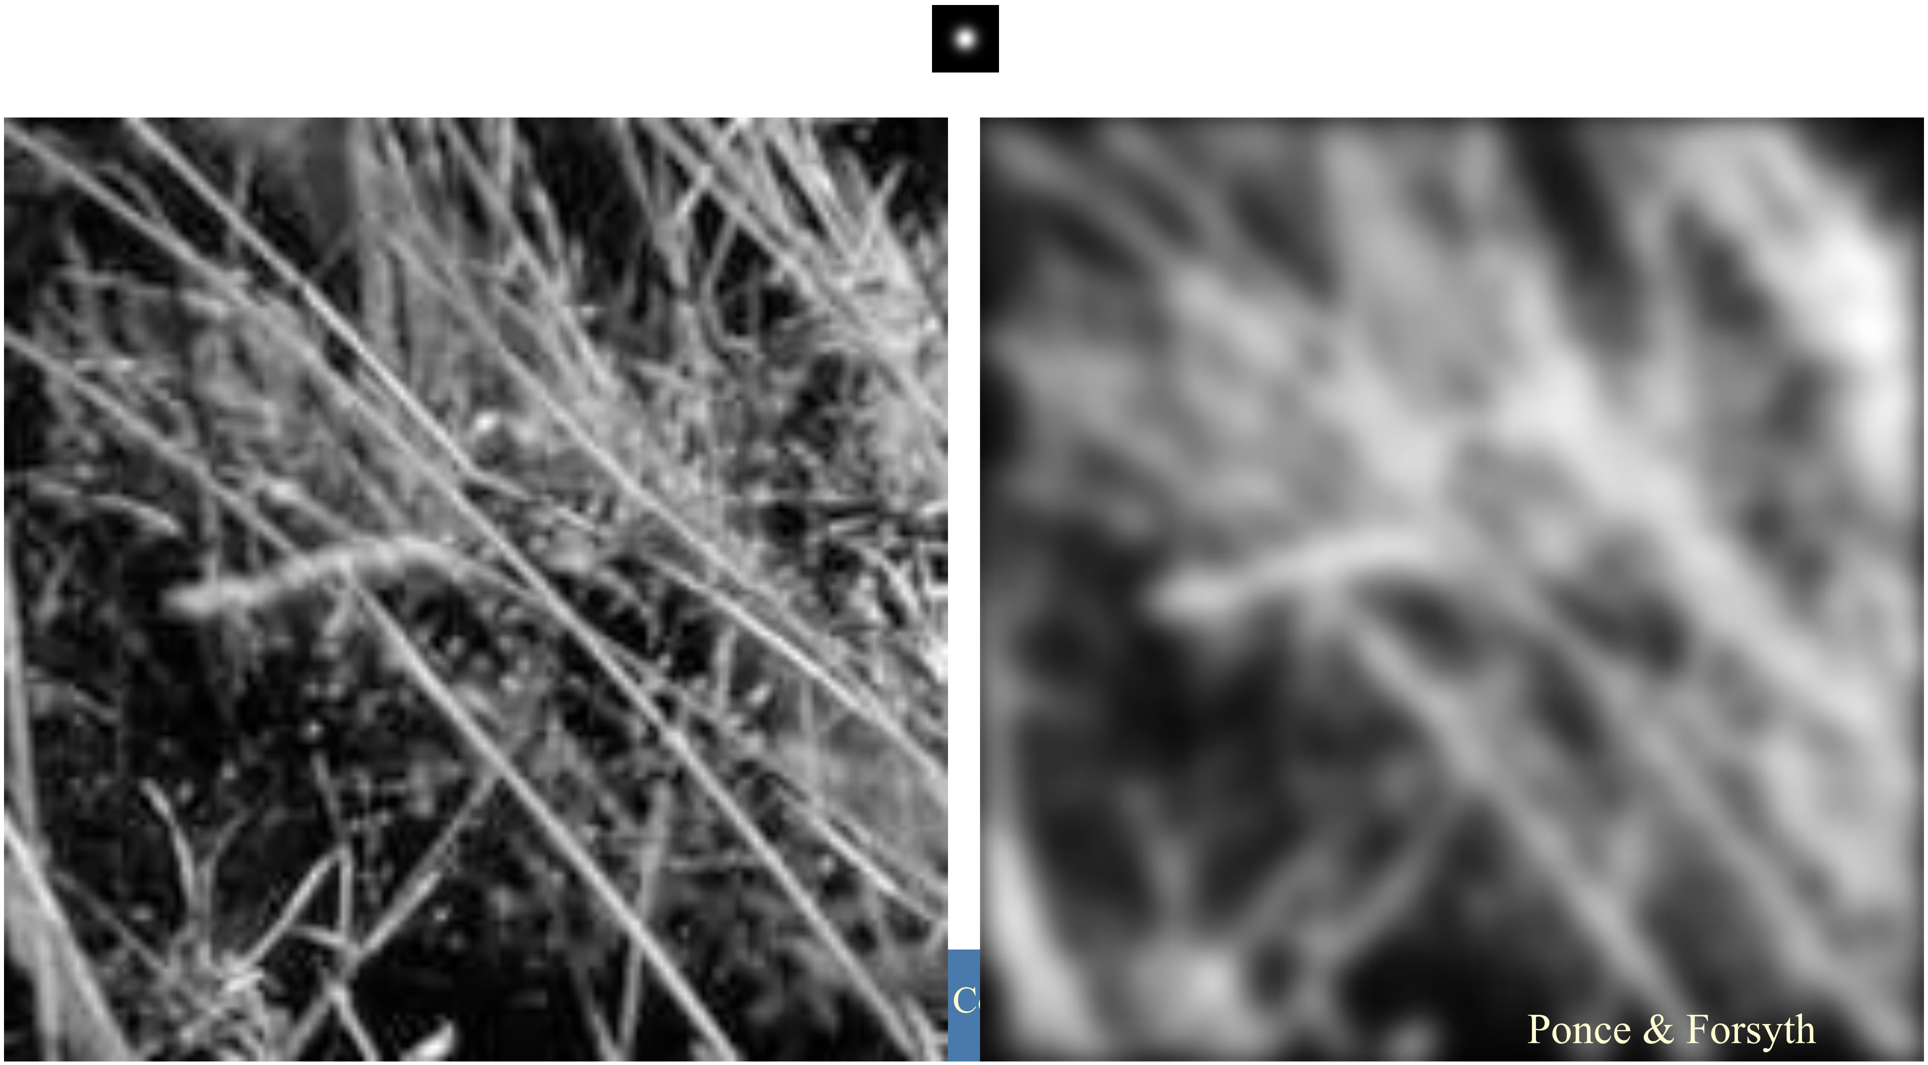

Homework

g=fspecial('gaussian',3,1)

g =

0.0751	0.1238	0.0751
0.1238	0.2042	0.1238
0.0751	0.1238	0.0751

```
B = imfilter(A,g,'symmetric','conv')
B = imfilter(A,g,'symmetric','corr')
% in the case of a symmetric filter, these are the same
```

Separable Gaussian

$$g(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-x^2/(2\sigma^2)) = G_x$$

$$g(y) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-y^2/(2\sigma^2)) = G_y$$

Product?

$$g(x,y) = \frac{1}{2\pi\sigma^2} \exp(-(x^2+y^2)/(2\sigma^2))$$

$$G_x \otimes (G_y \otimes I) = (G_x \otimes G_y) \otimes I = G_{xy} \otimes I$$

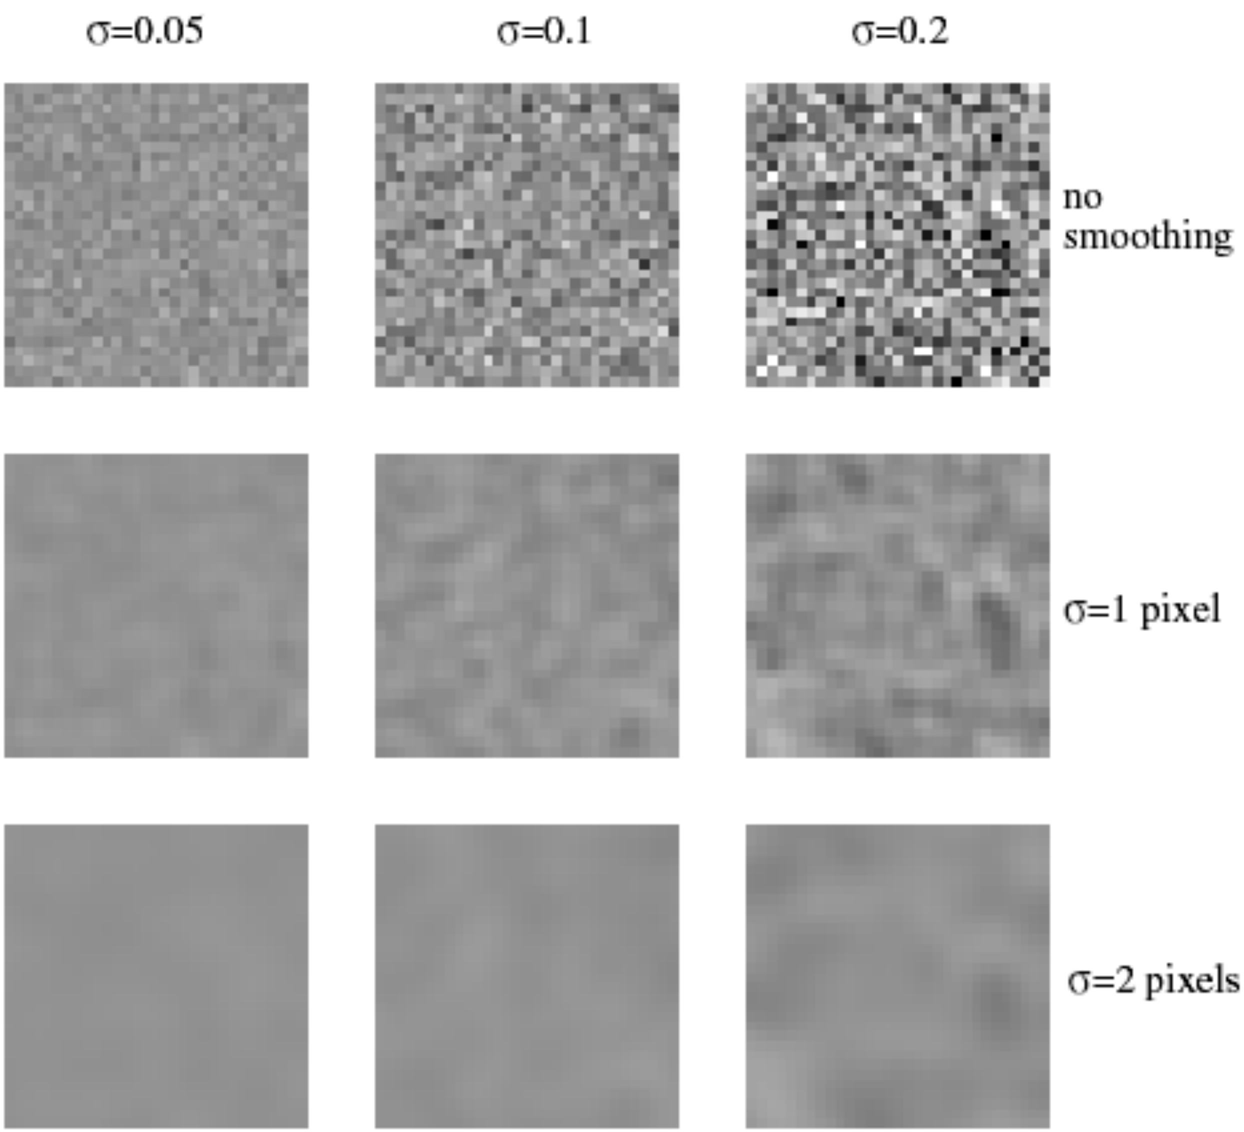

The effects of smoothing
Each row shows smoothing with Gaussians of different width; each column shows different realizations of an image of gaussian noise.

Ponce & Forsyth

Multi-Resolution Image Representation

- Gaussian pyramids
- Laplacian Pyramids

Source: Irani

Motivation for studying scale.

ELDER AND ZUCKER: LOCAL SCALE CONTROL FOR EDGE DETECTION AND BLUR ESTIMATION

IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, VOL. 20, NO. 7, JULY 1998

Motivation for studying scale.

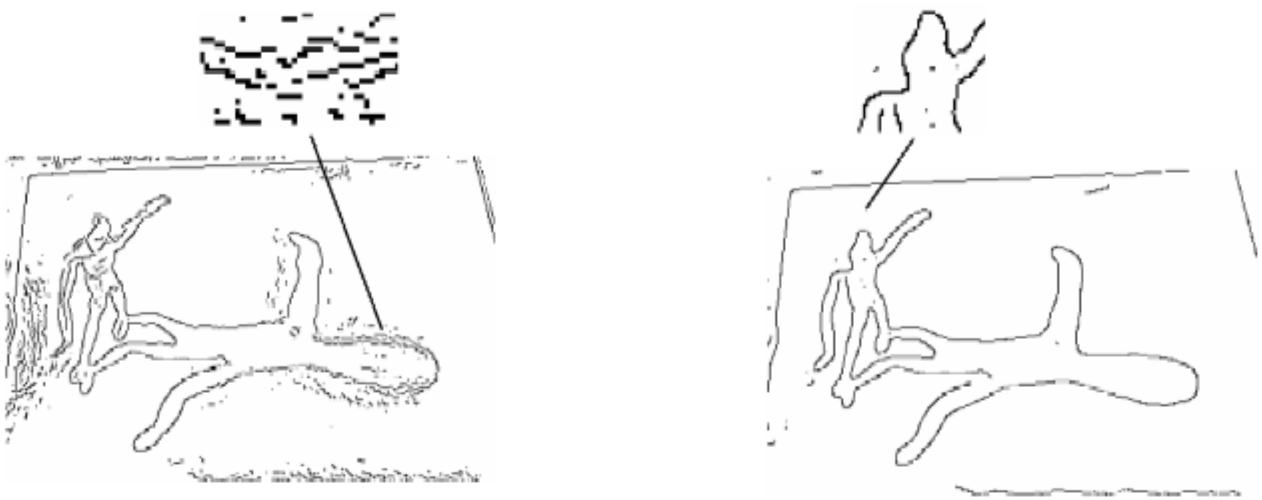

ELDER AND ZUCKER: LOCAL SCALE CONTROL FOR EDGE DETECTION AND BLUR ESTIMATION

IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, VOL. 20, NO. 7, JULY 1998

A little Fourier so we can talk about spatial frequency:

Fourier Transform in Picutres

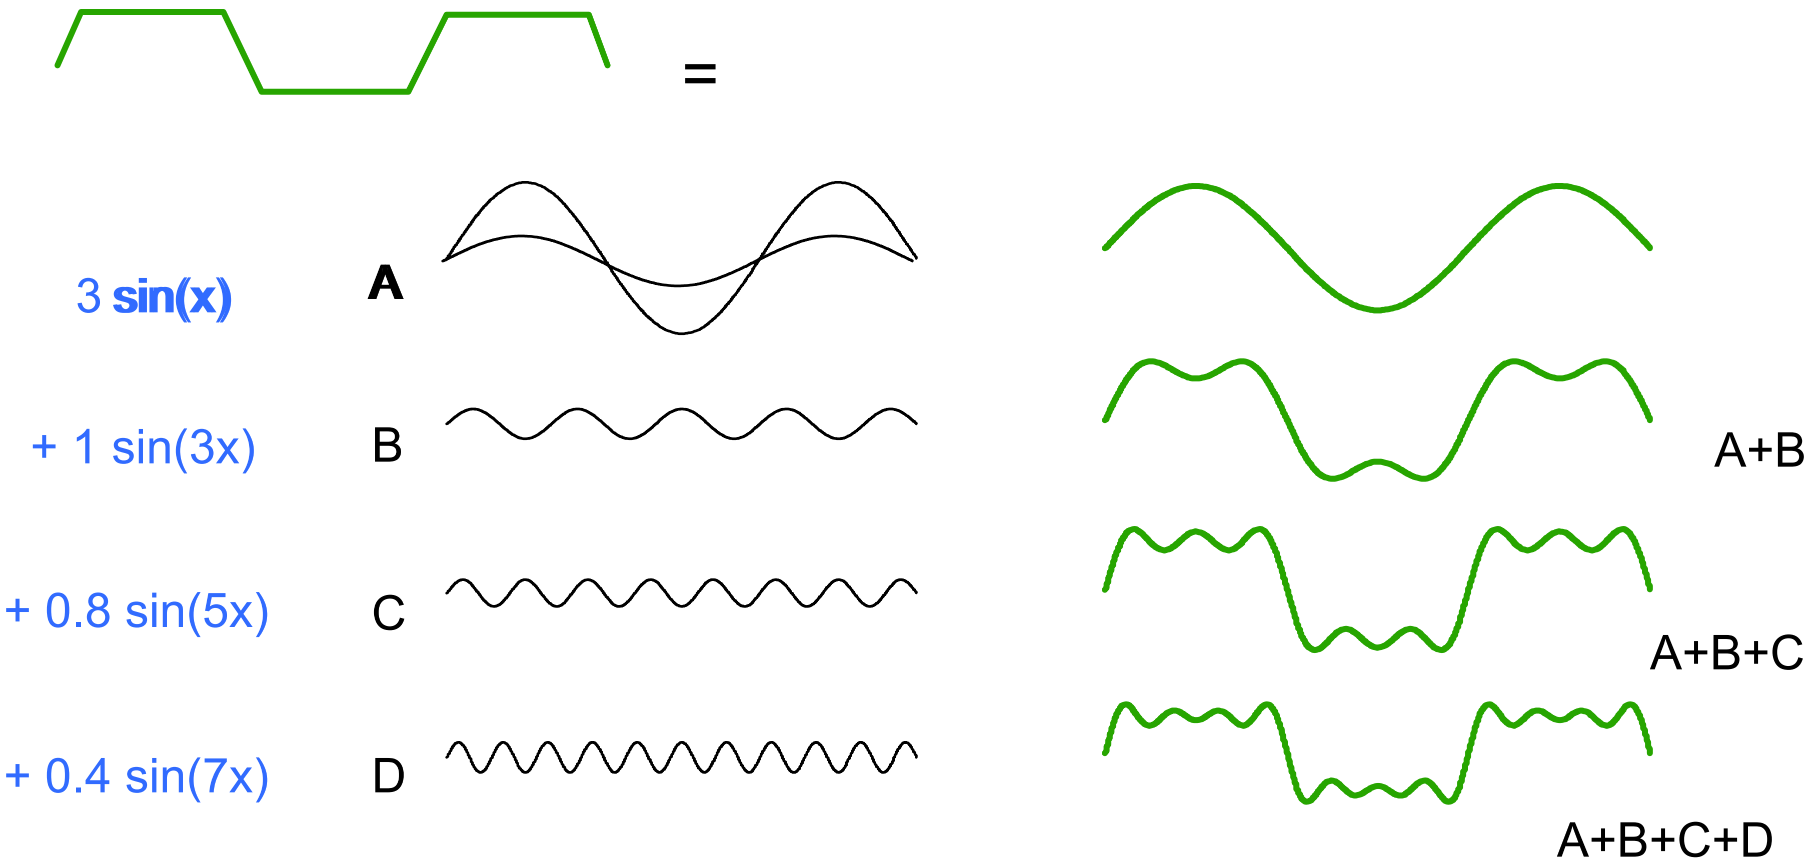

Hwk 1 Prob 1, hint

- Use structures and cell arrays
- Write a nice display function
- Don’t represent pyramid as an image and then filter this image – it will produce artifacts at the boundaries between the levels.

Gaussian Pyramid

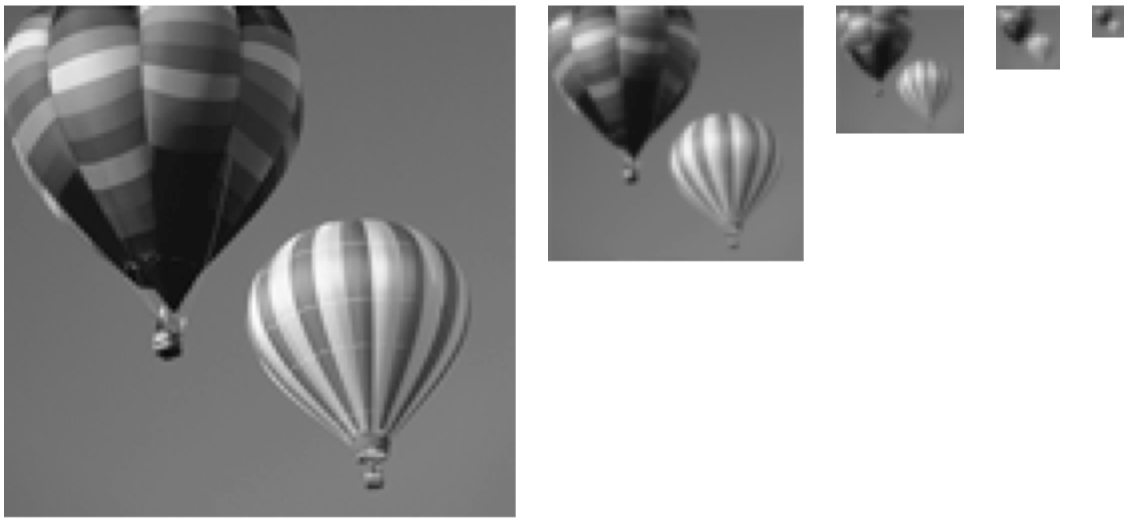

High resolution → **Low resolution**

Source: Irani

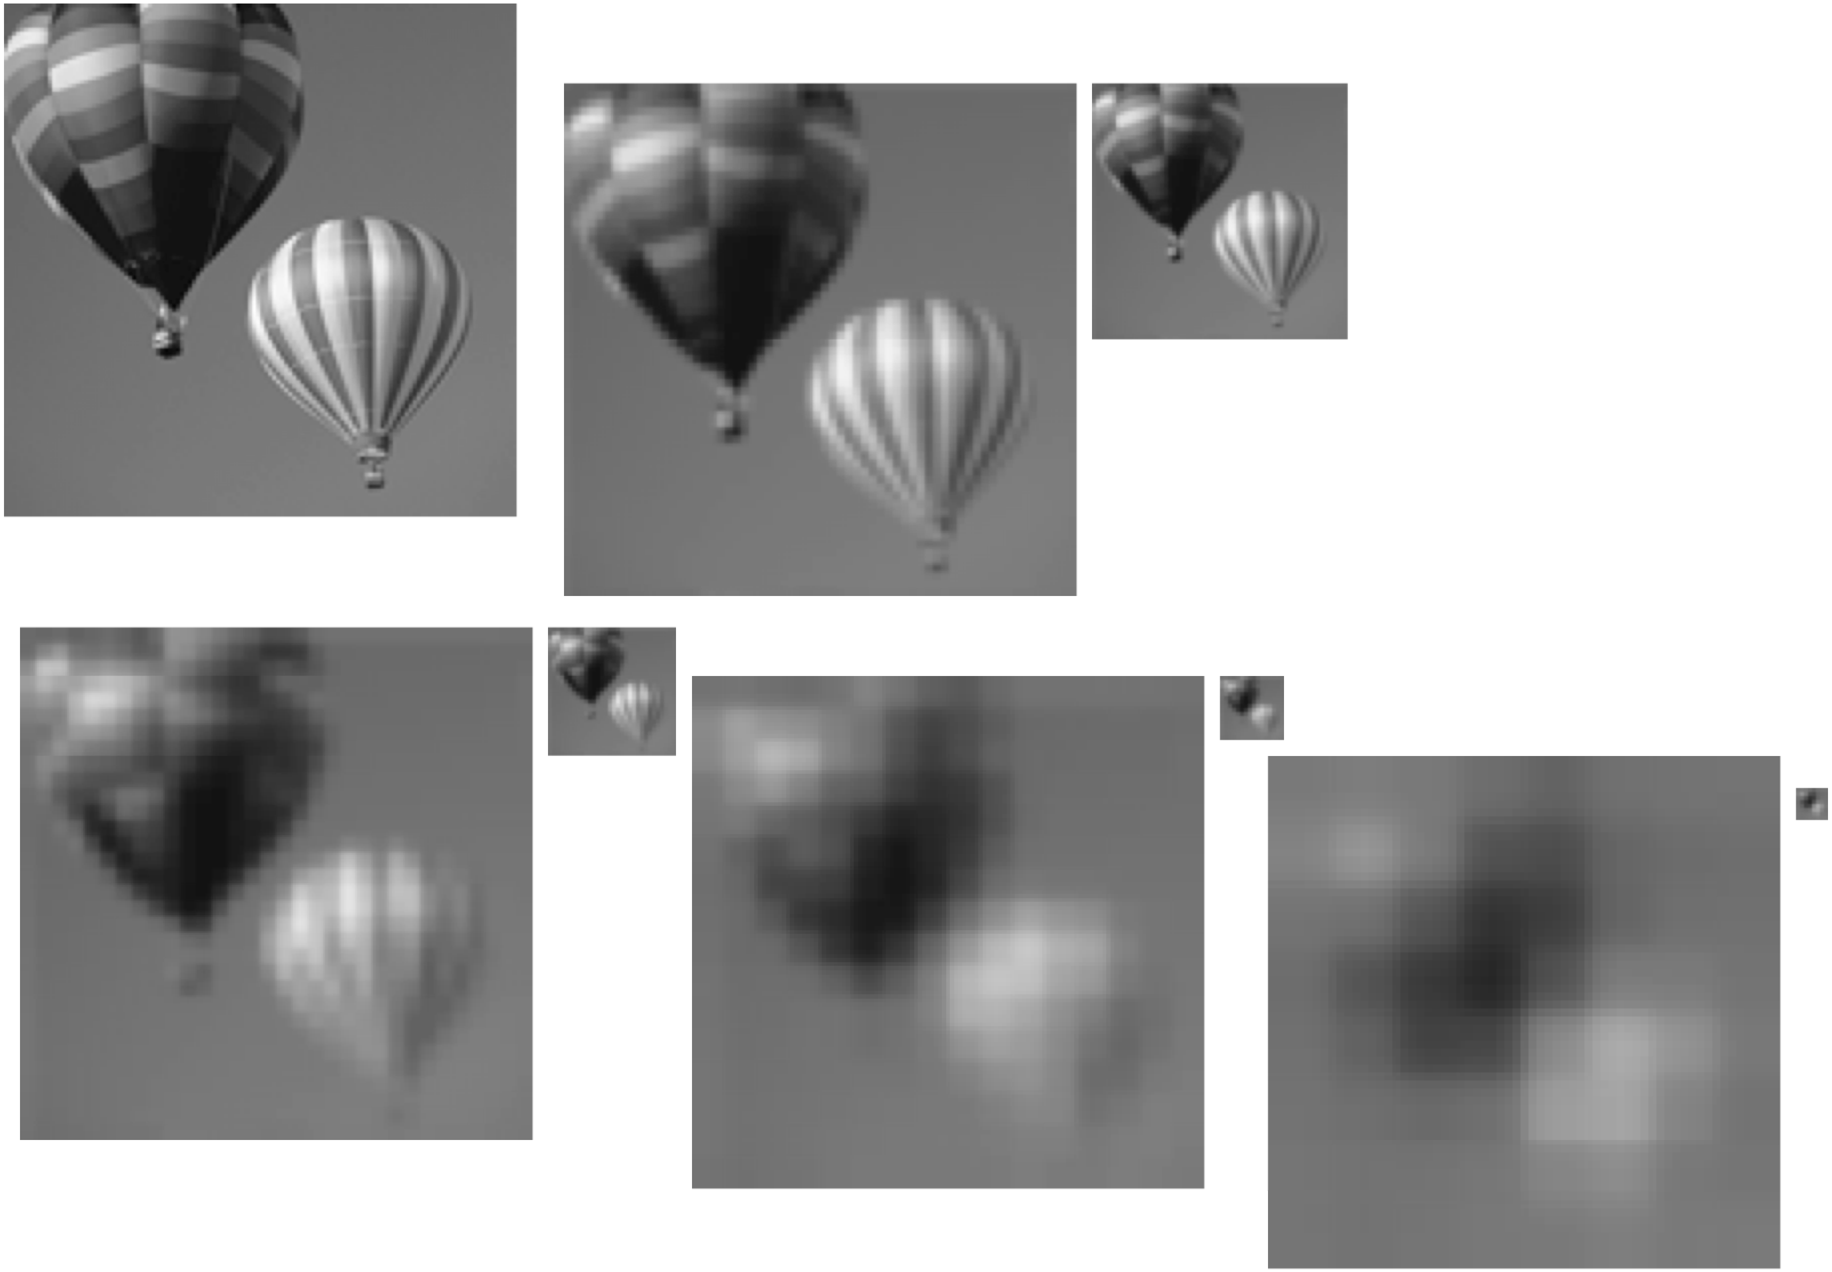

Source: Irani

An Isotropic Gaussian

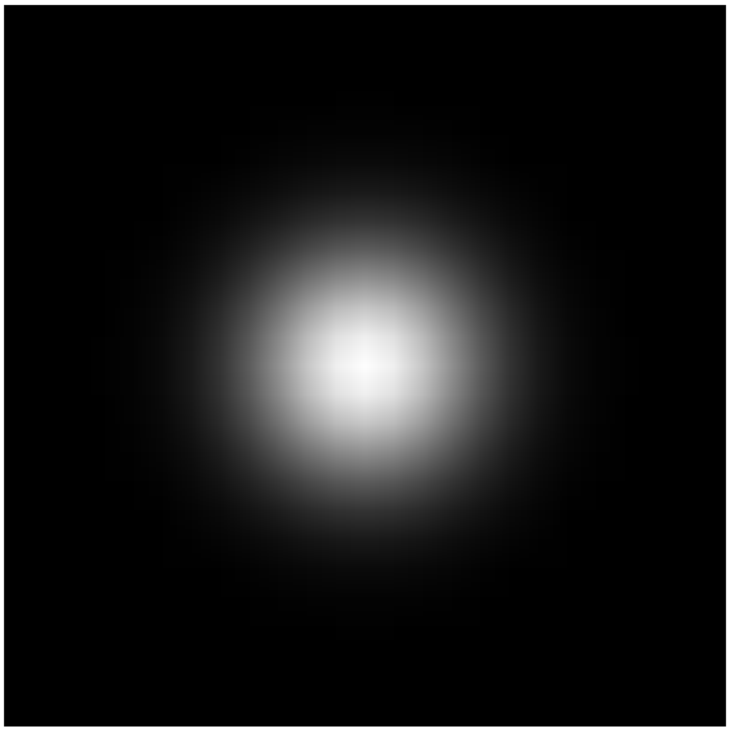

(0,0) in center

- The picture shows a smoothing kernel proportional to

$$g(x,y)=\frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x^2+y^2)}{2\sigma^2}\right)$$

To smooth image, convolve with this filter.

Ponce & Forsyth

The Gaussian Pyramid

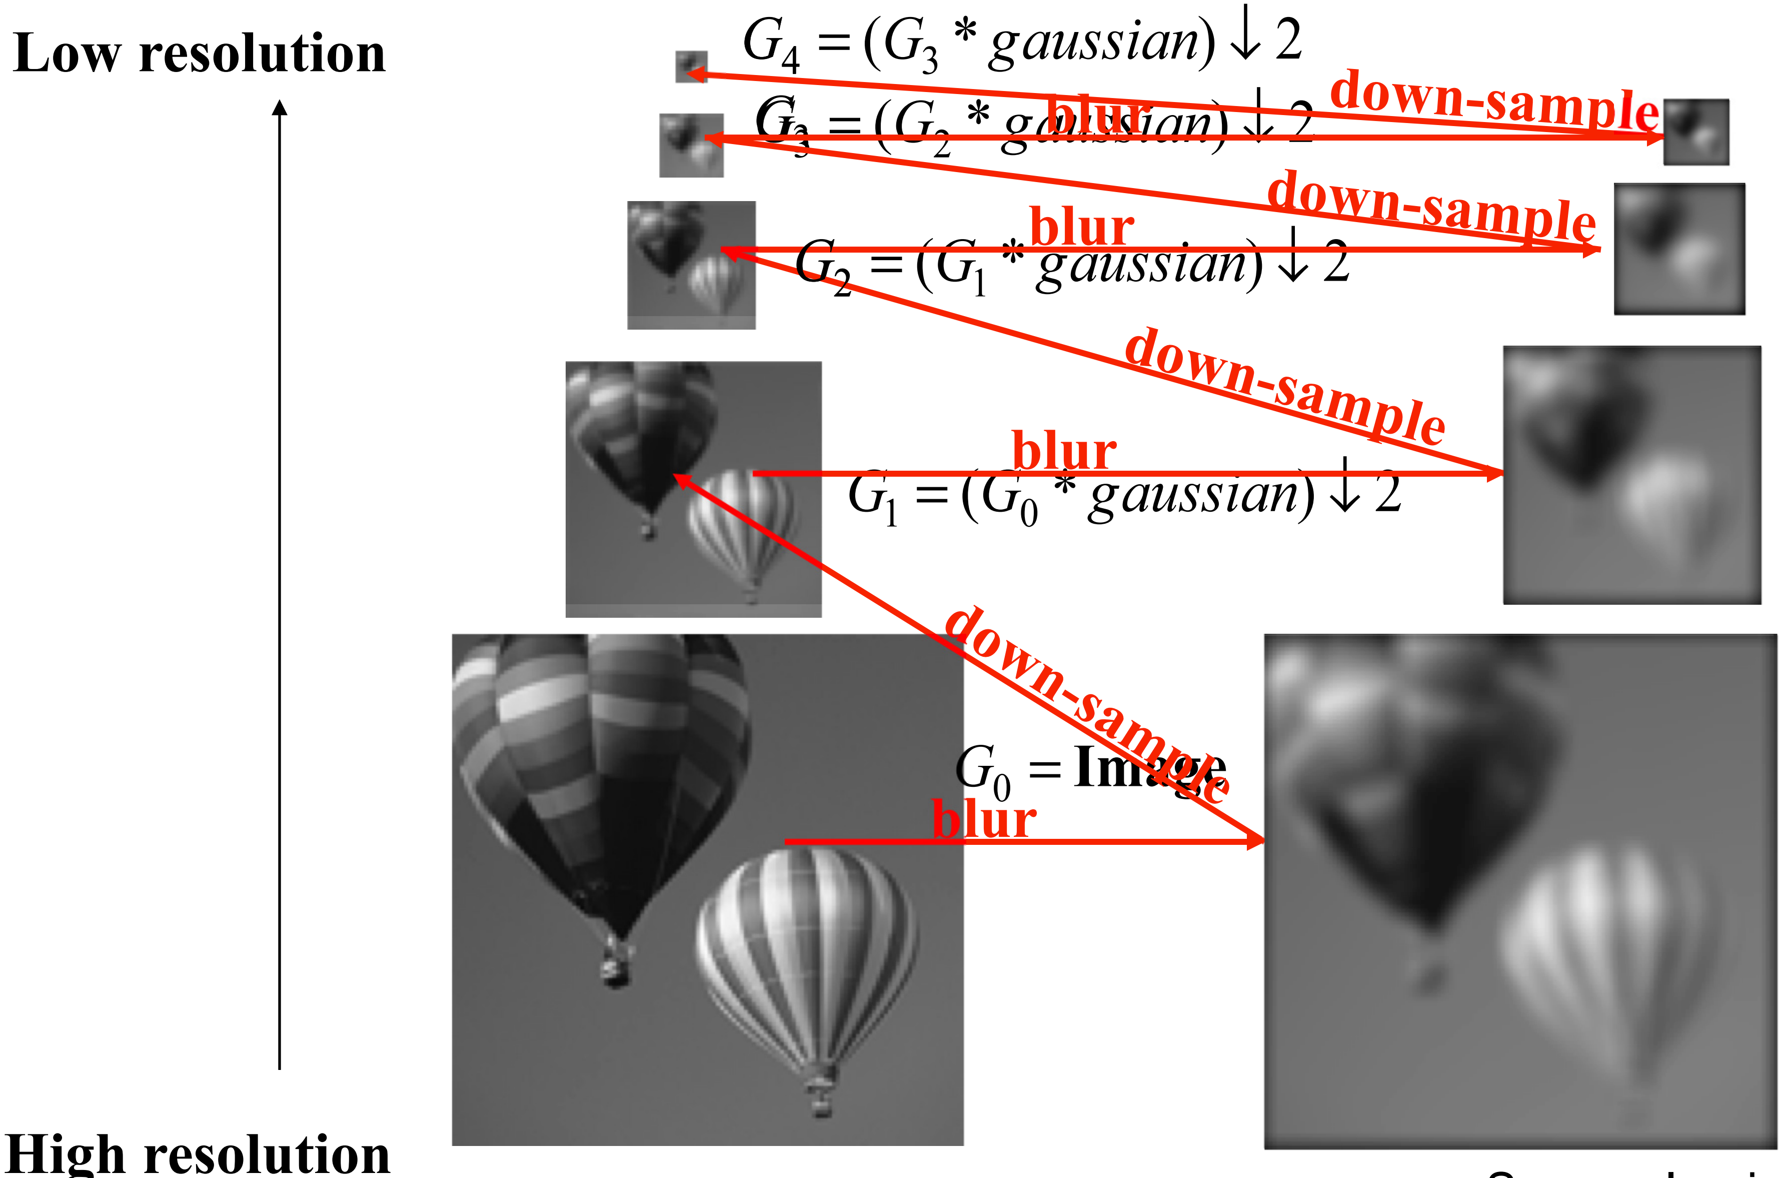

Source: Irani

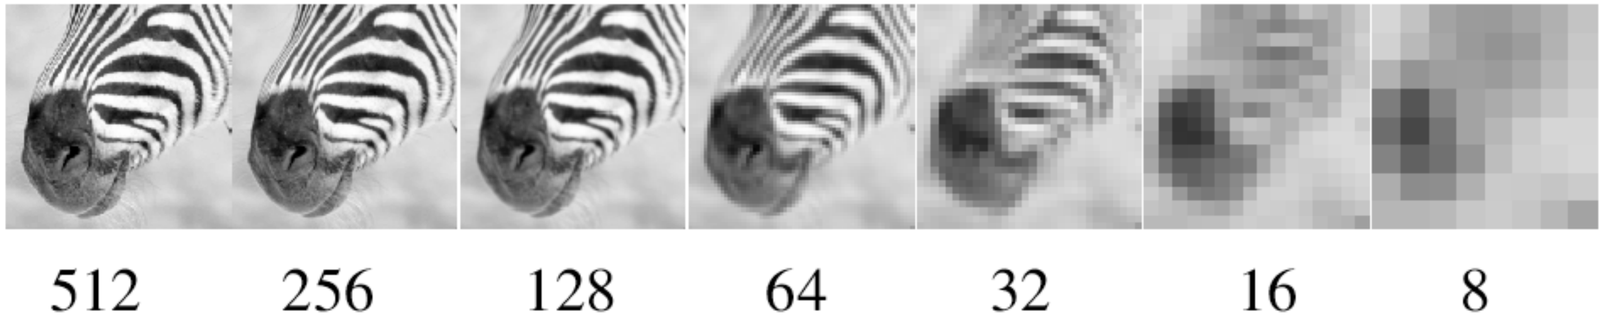

A bar in the big images is a hair on the zebra's nose; in smaller images, a stripe; in the smallest, the animal's nose

Source: Ponce&Forsyth

Foveal/peripheral vision

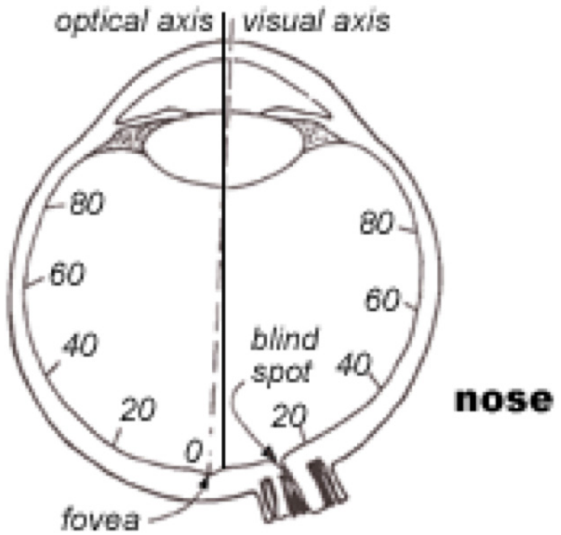

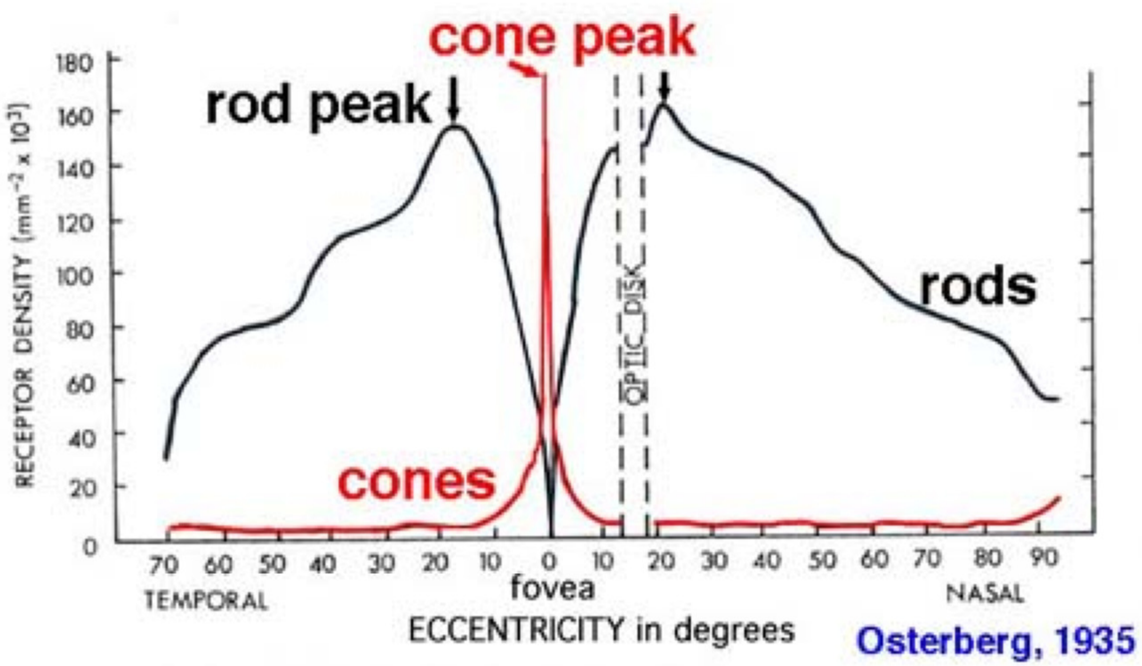

Mona Lisa Smile

Why is her smile so mysterious?
Why is this picture so fascinating?
What could spatial frequency have to do with it?

Mona Lisa Smile

Low frequency

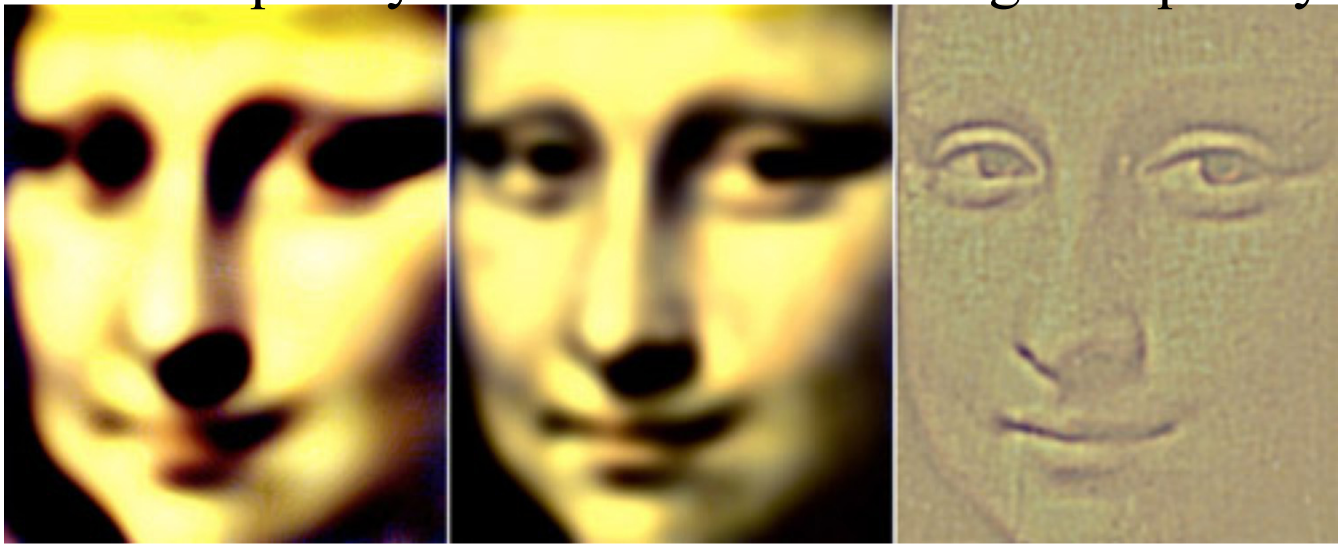

High frequency

Margaret Livingstone

Mona Lisa Smile

Fixate on eyes

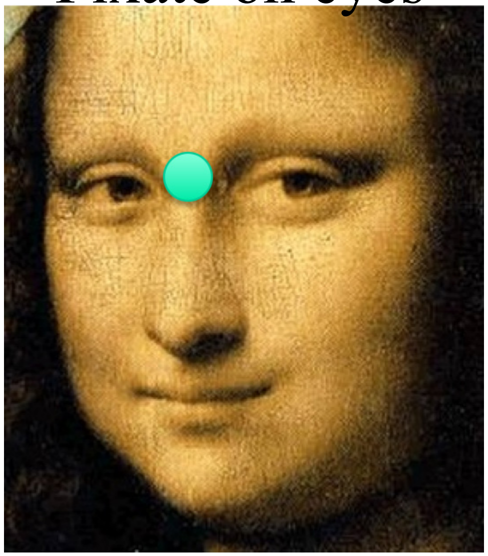

Low frequency

High frequency

Mouth seen in low resolution periphery

Margaret Livingstone

Mona Lisa Smile

Fixate on mouth

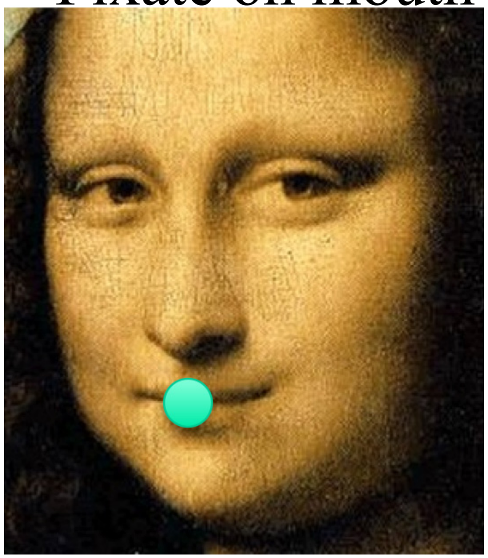

Low frequency High frequency

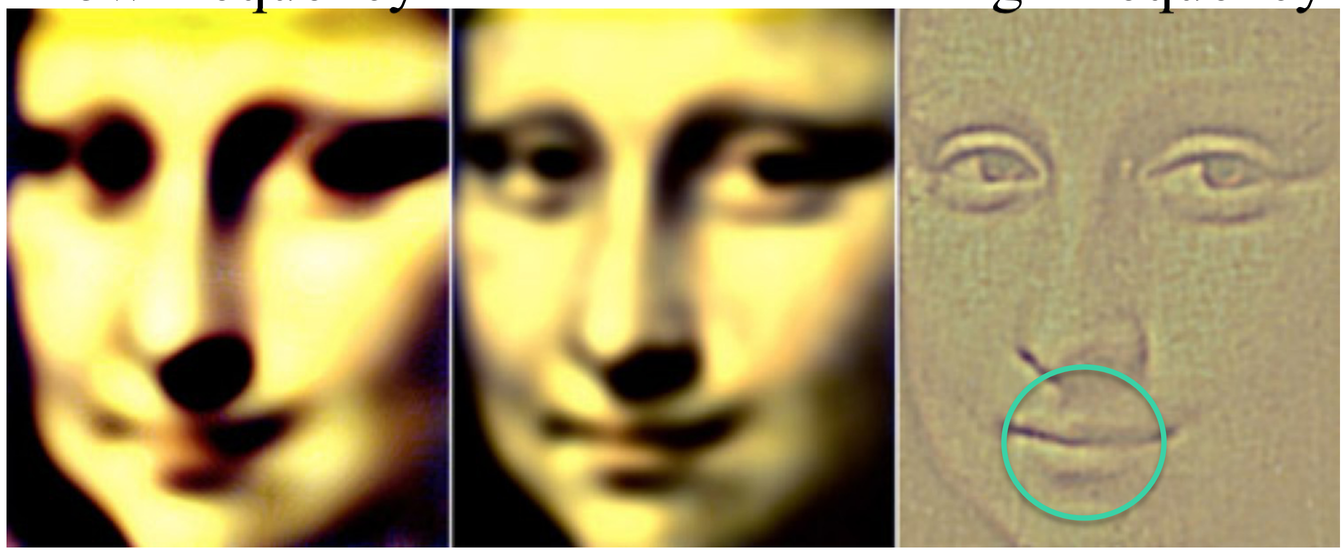

Mouth seen in high resolution fovea

Margaret Livingstone

Motivation: Search

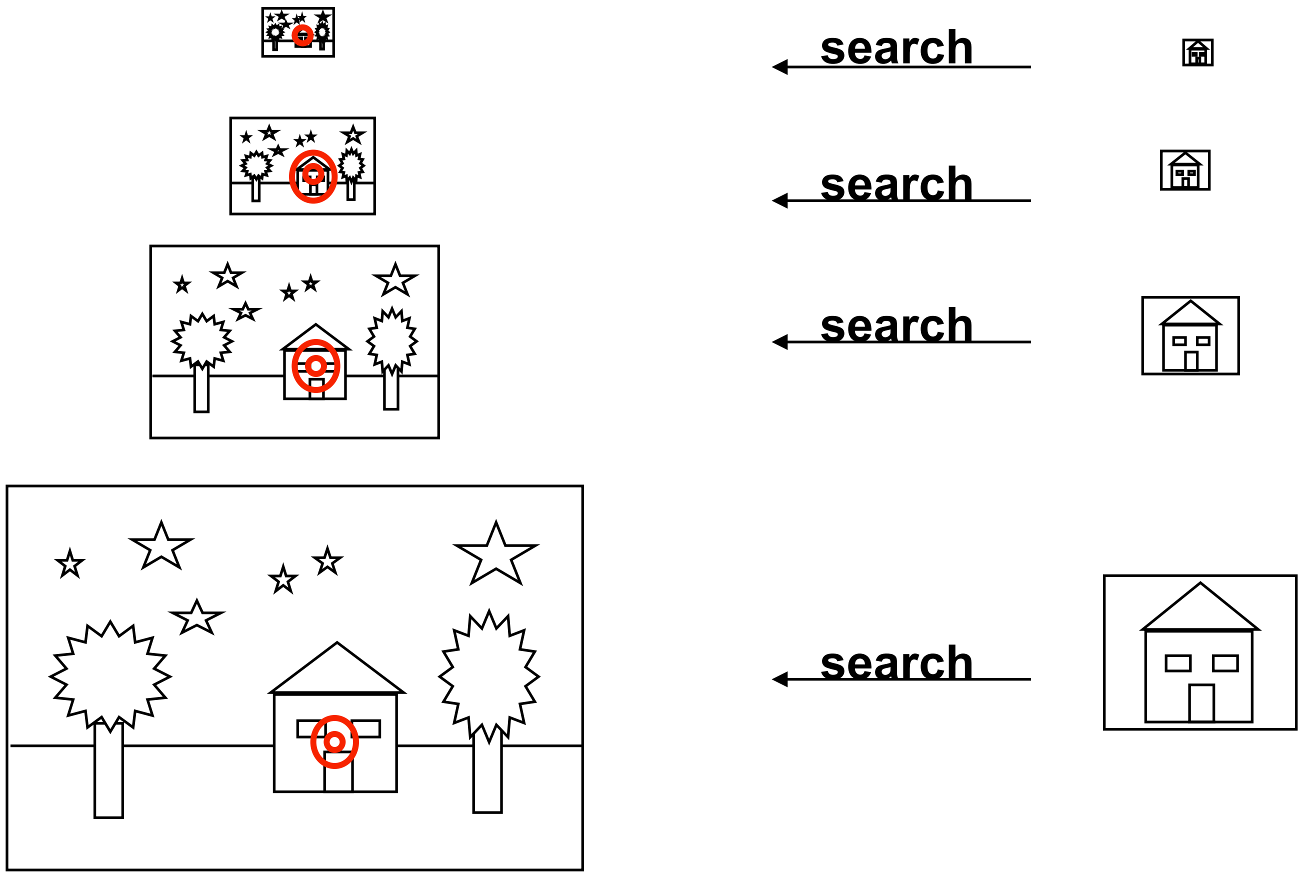

Irani & Basri

Sub-sampling

Why smooth before sub-sampling?

Subsampling

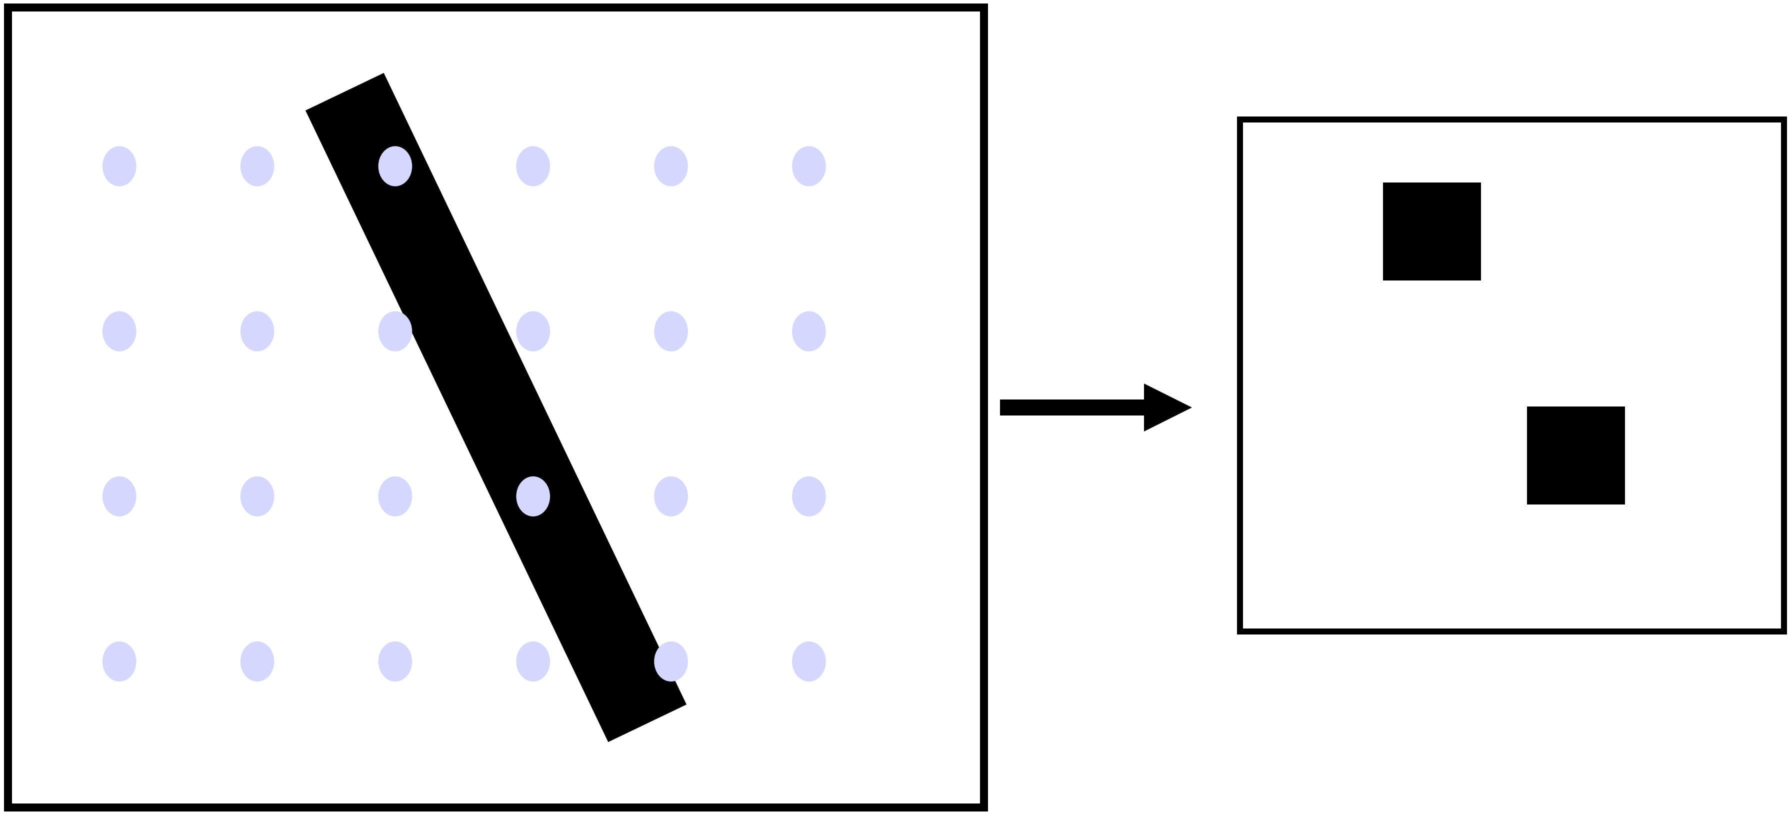

Aliasing

- Can’t shrink an image by taking every second pixel
- If we do, characteristic errors appear
  - Common phenomenon
    - Wagon wheels rolling the wrong way in movies
    - Checkerboards misrepresented in ray tracing

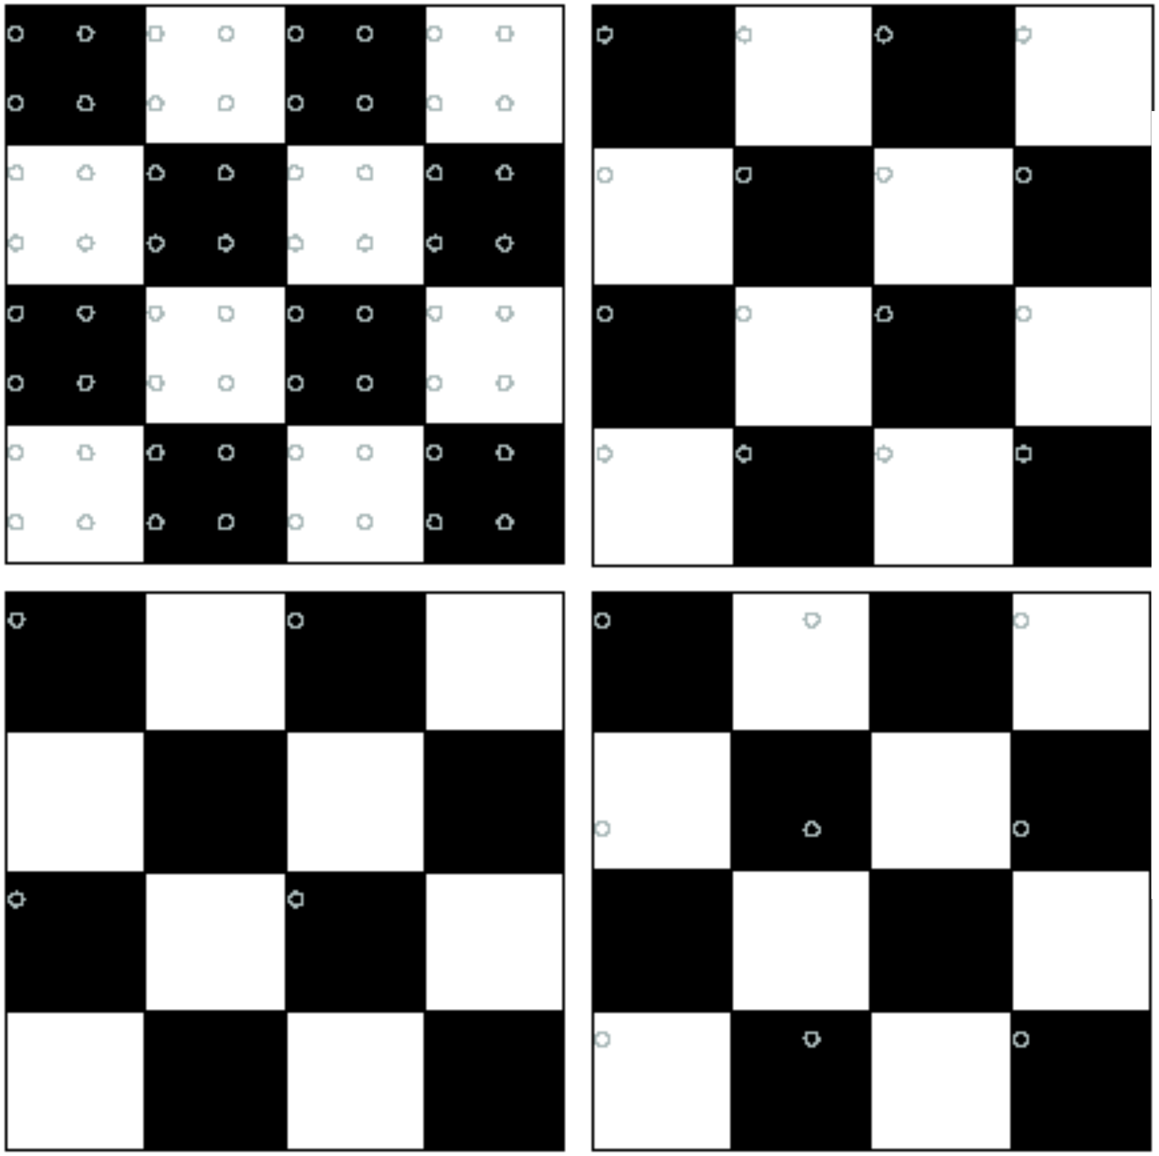

Resample the checkerboard by taking one sample at each circle. In the case of the top left board, new representation is reasonable. Top right also yields a reasonable representation. Bottom left is all black (dubious) and bottom right has checks that are too big.

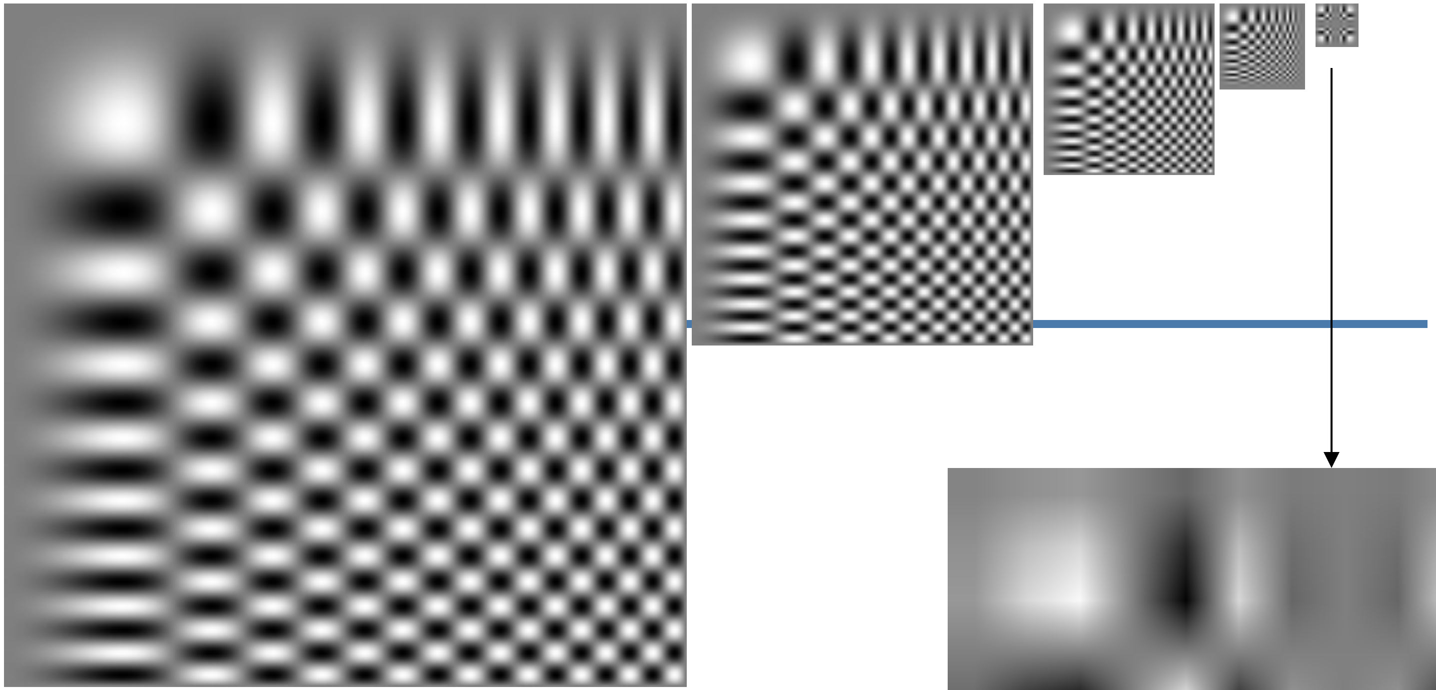

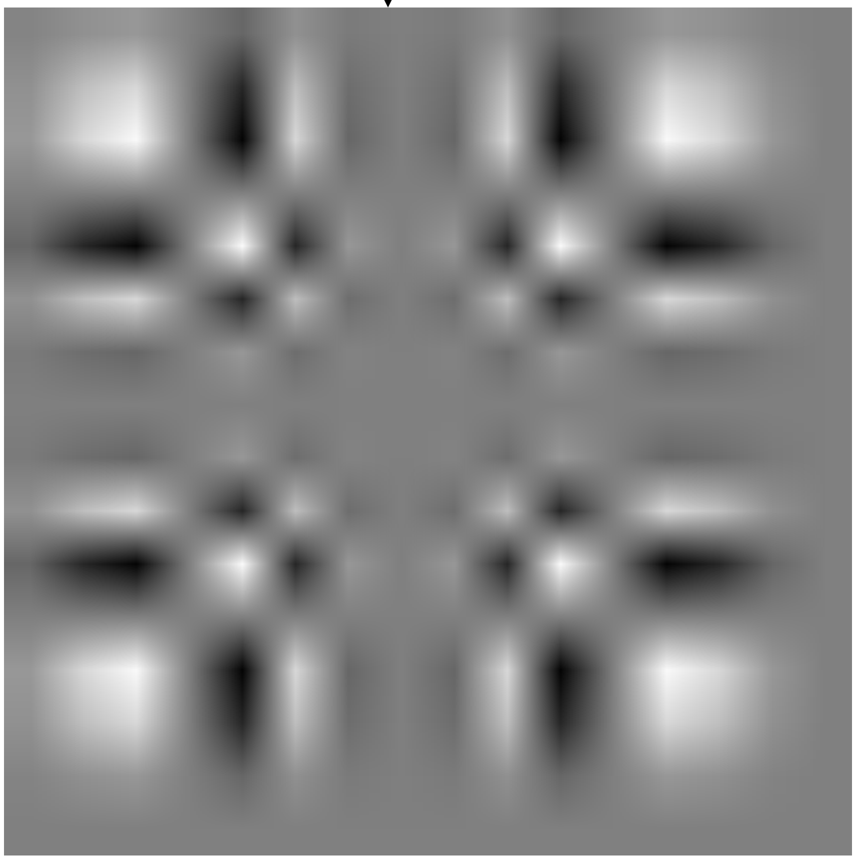

Constructing a pyramid by taking every second pixel leads to layers that badly misrepresent the top layer