THE UNIVERSITY OF BRITISH COLUMBIA

CPSC 425: Computer Vision

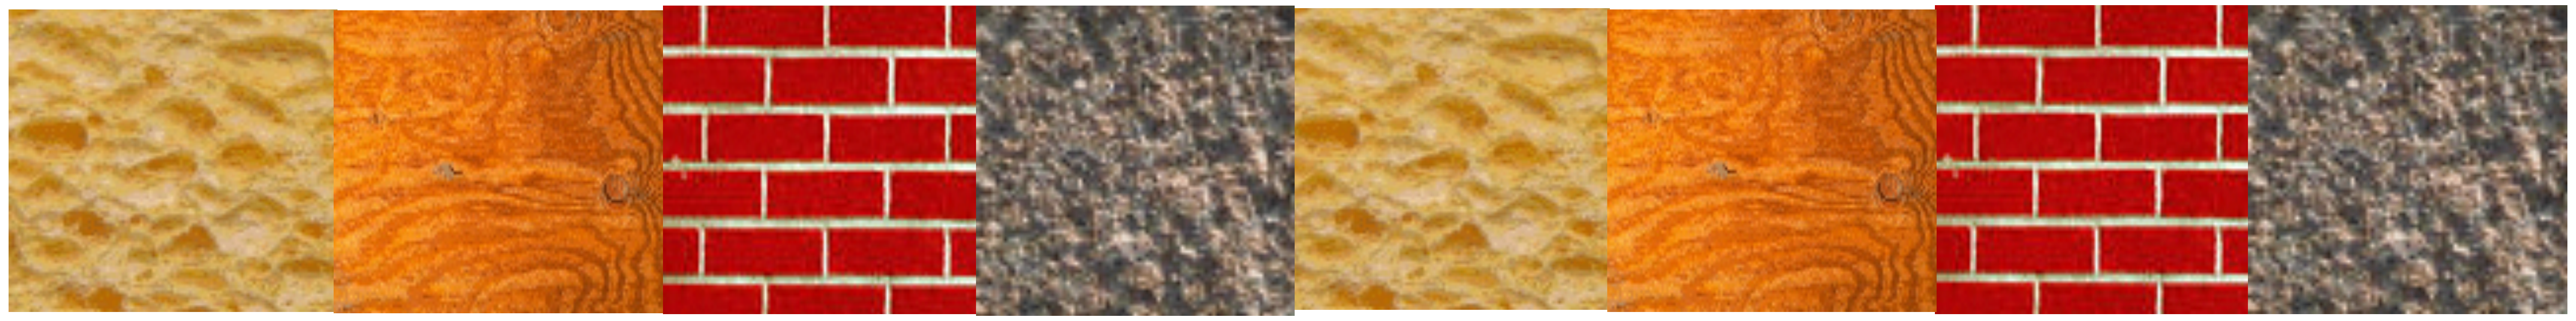

Lecture 15: Texture, Intro to Color

(unless otherwise stated slides are taken or adopted from **Bob Woodham, Jim Little** and **Fred Tung**)

Menu for Today (**October 10, 2018**)

Topics:

- Texture Analysis
- Laplacian and Oriented Pyramids
- iClicker Quiz
- Introduction to Color

Redings:

- **Today's** Lecture: Forsyth & Ponce (2nd ed.) 3.1-3.3
- **Next** Lecture: N/A

Reminders:

- **Assignment 2**: Face Detection in a Scaled Representation is **due today**
- **Assignment 3**: Texture Syntheis will be **out today**
- **Practice questions** are available now, Additional Office hours Fri, Mon

Today's "**fun**" Example: NCIS

Today's "**fun**" Example: LavaRAND

Today's "**fun**" Example: LavaRAND at Cloudflare

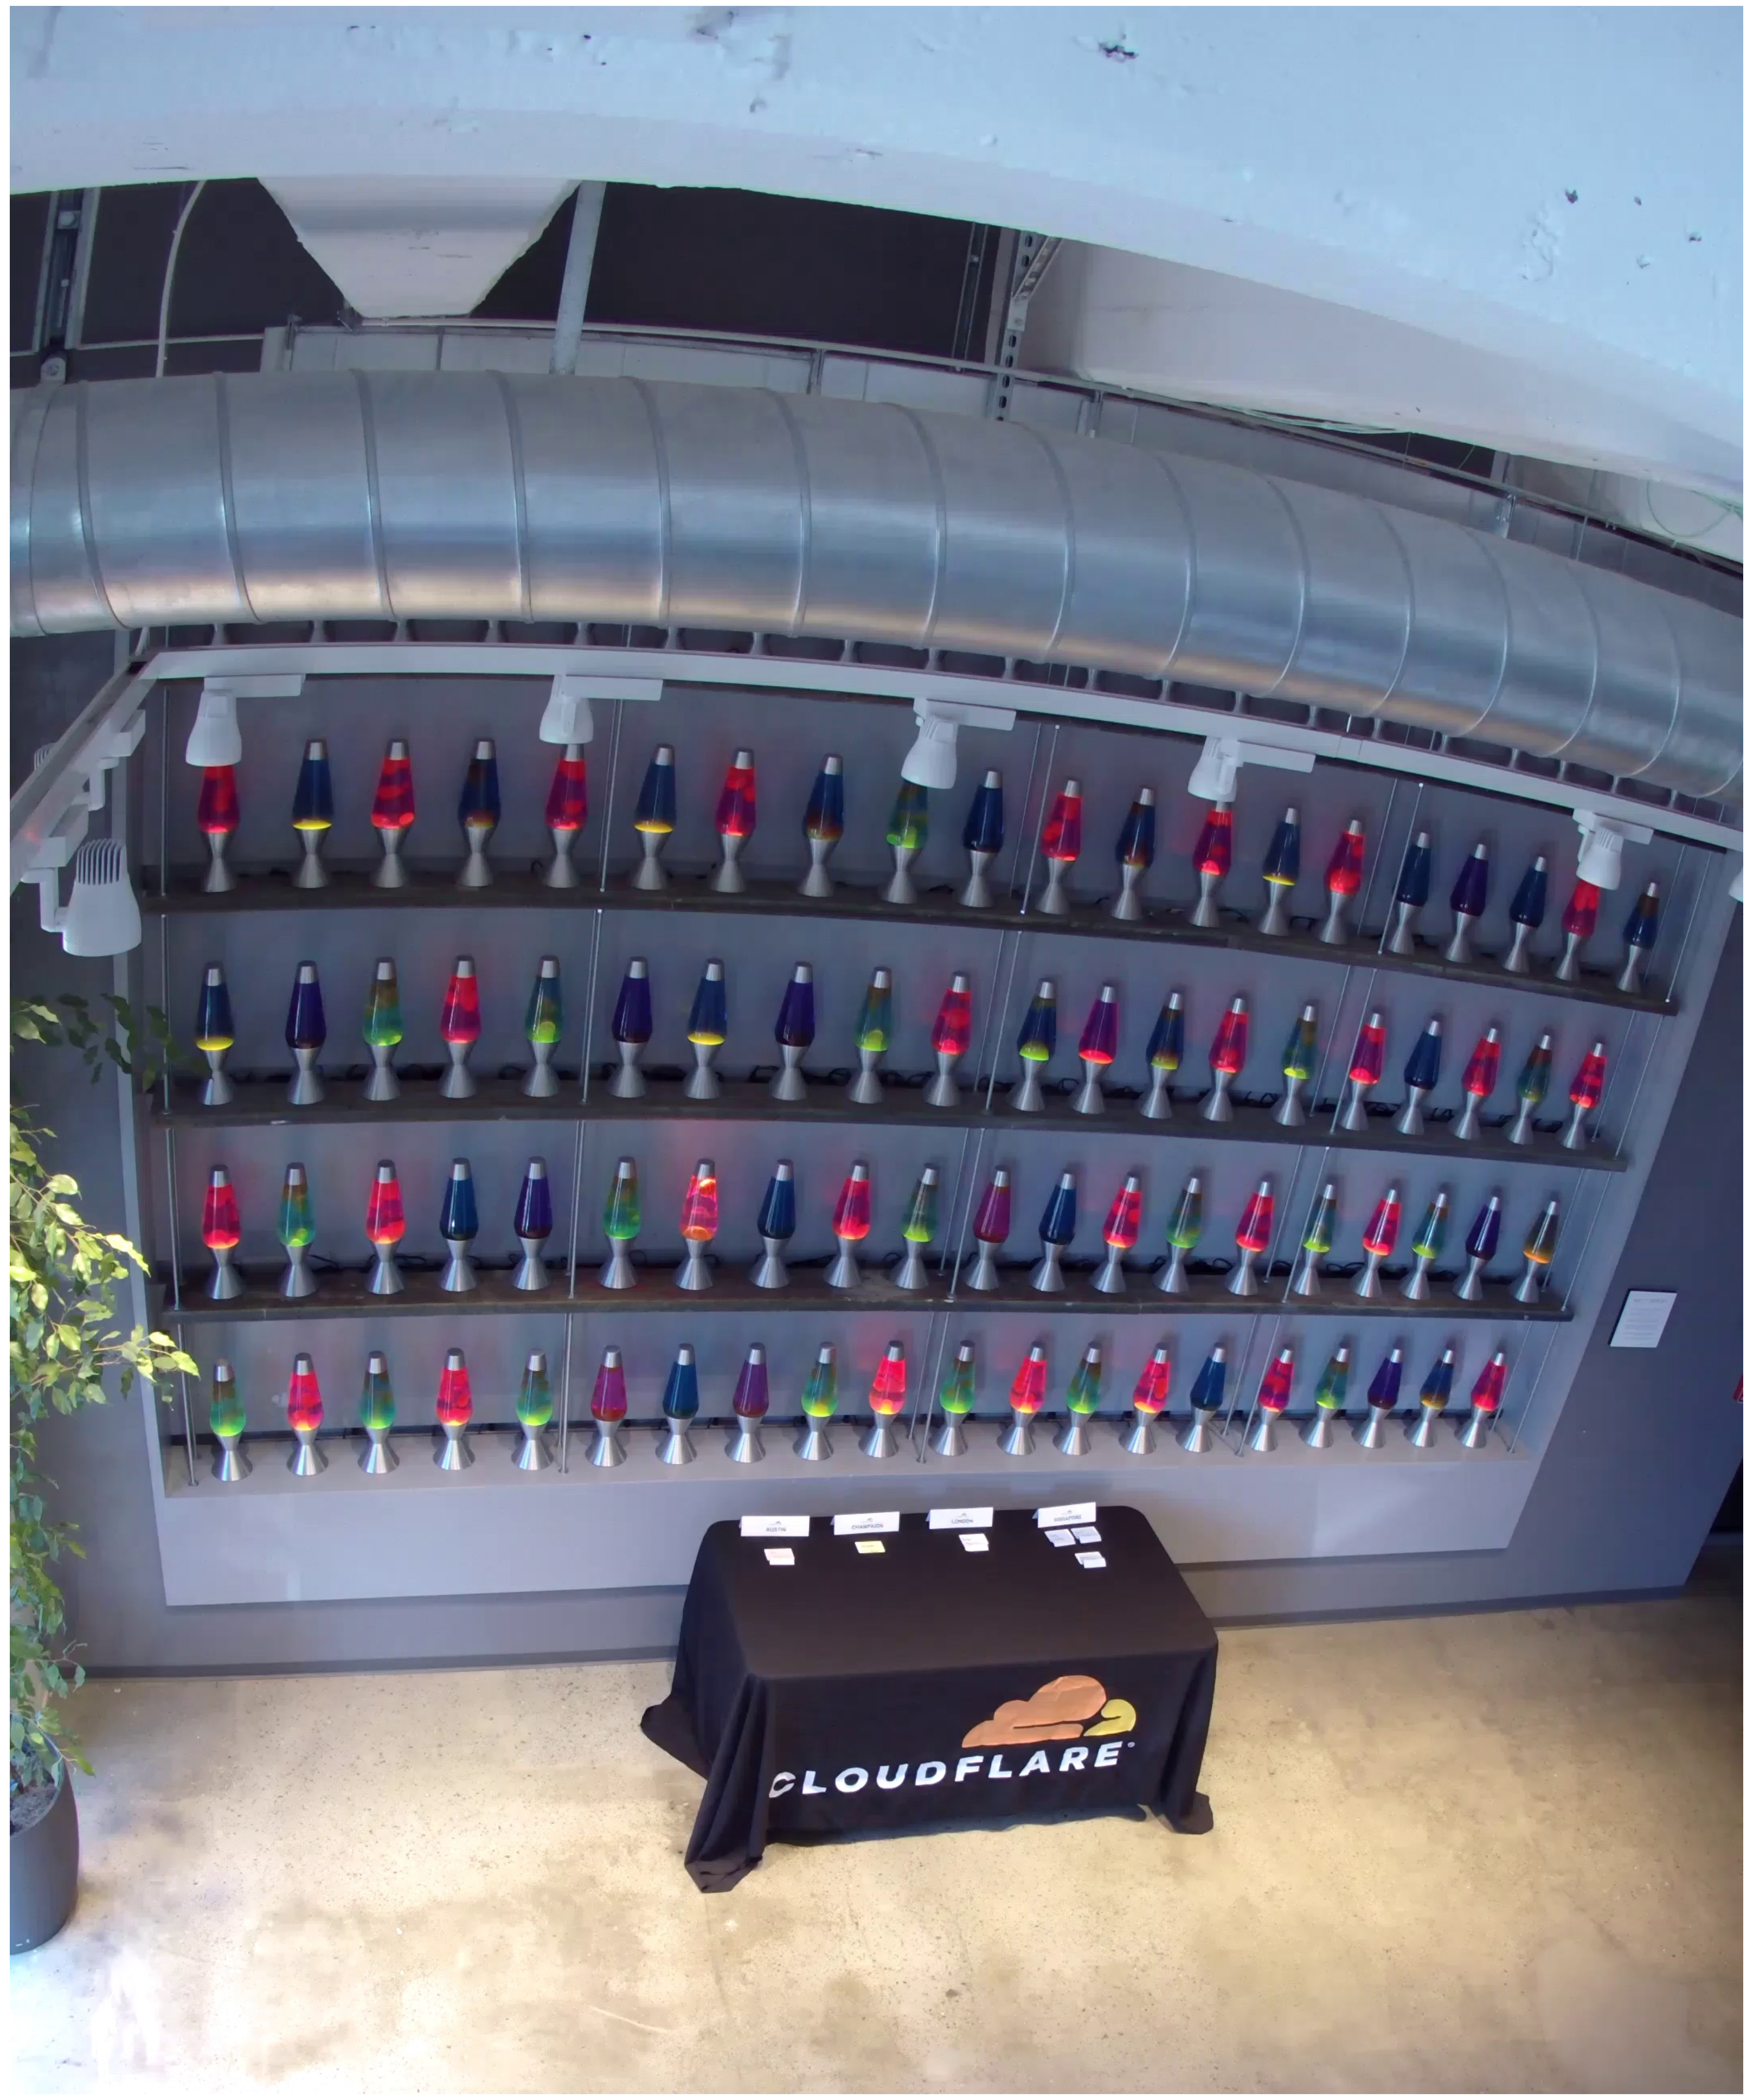

Lecture 14: Re-cap

Texture representation is hard
— difficult to define, to analyze
— texture synthesis appears more tractable

Objective of texture **synthesis** is to generate new examples of a texture
— Efros and Leung: Draw samples directly from the texture to generate one pixel at a time. A “data-driven" approach.

Approaches to texture embed assumptions related to human perception

Spots and Bars (Fine Scale)

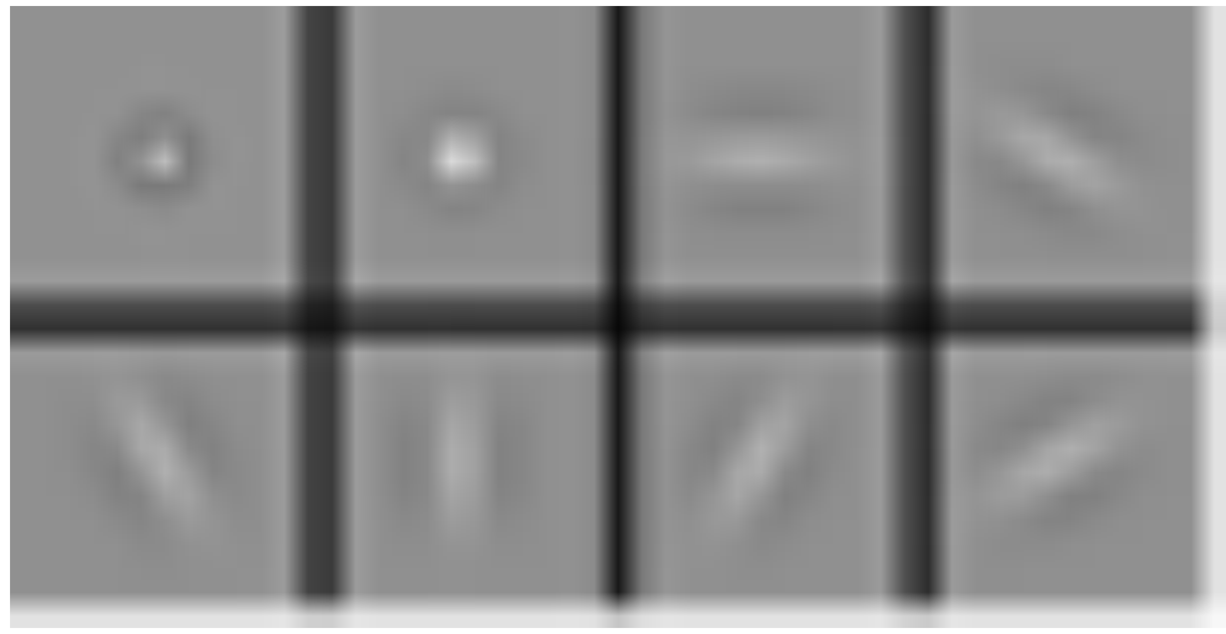

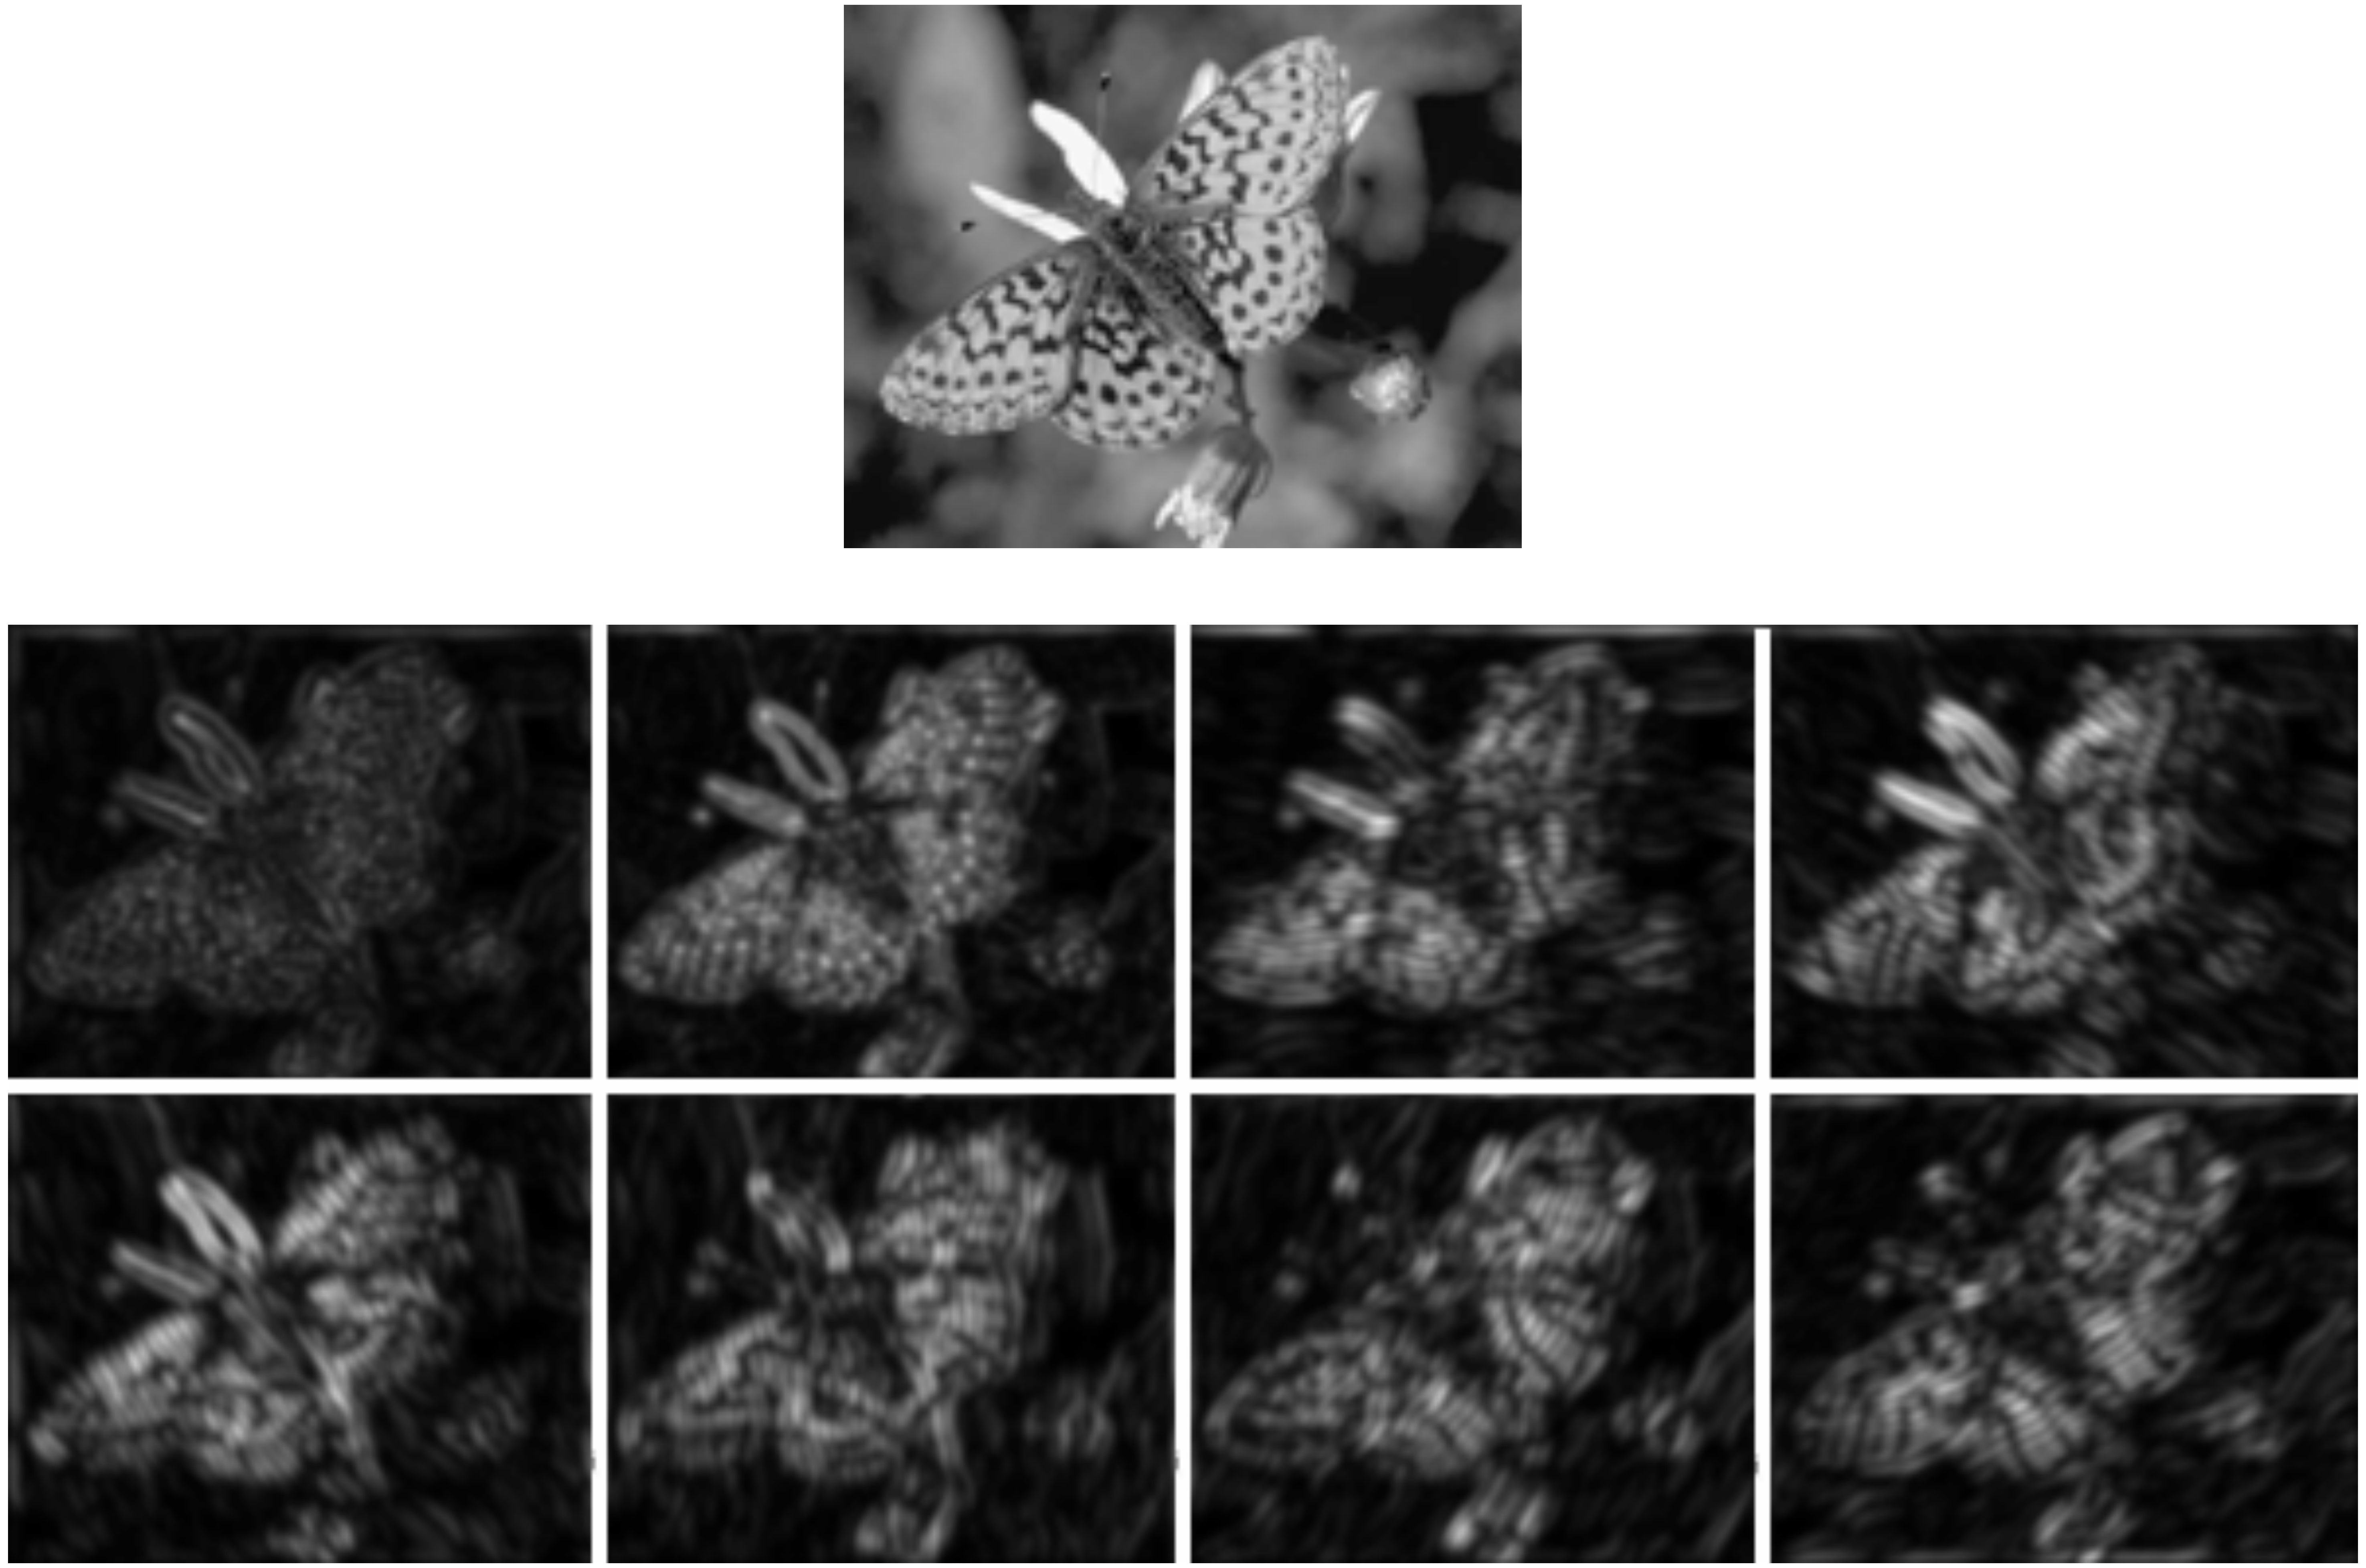

Forsyth & Ponce (1st ed.) Figures 9.3–9.4

Spots and Bars (Coarse Scale)

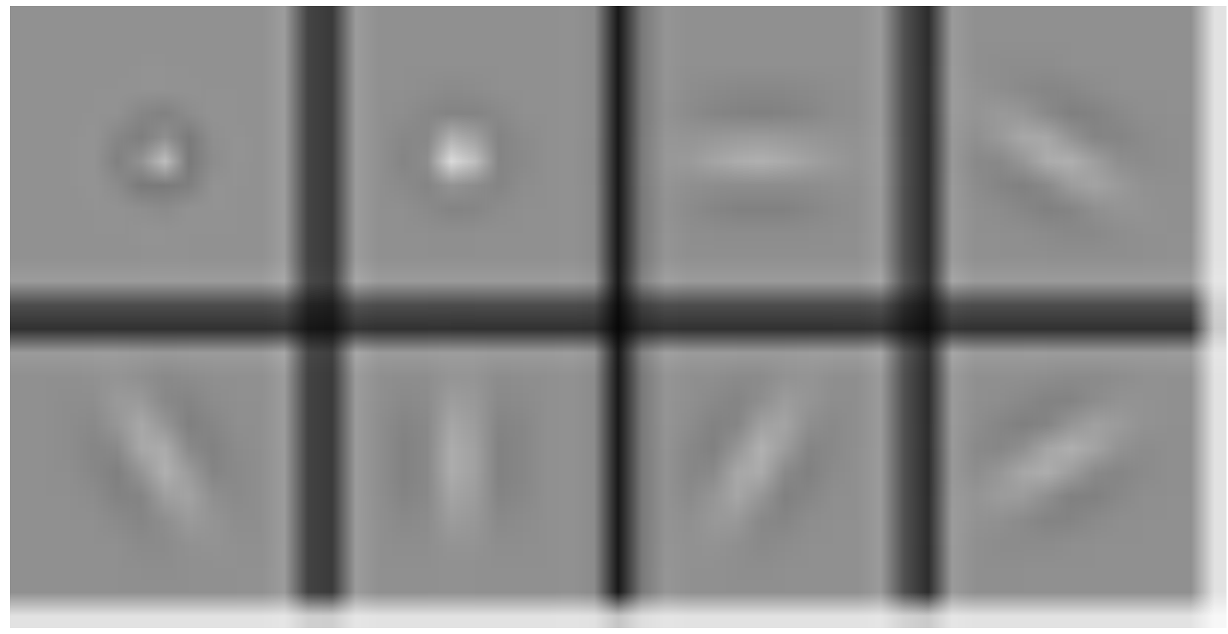

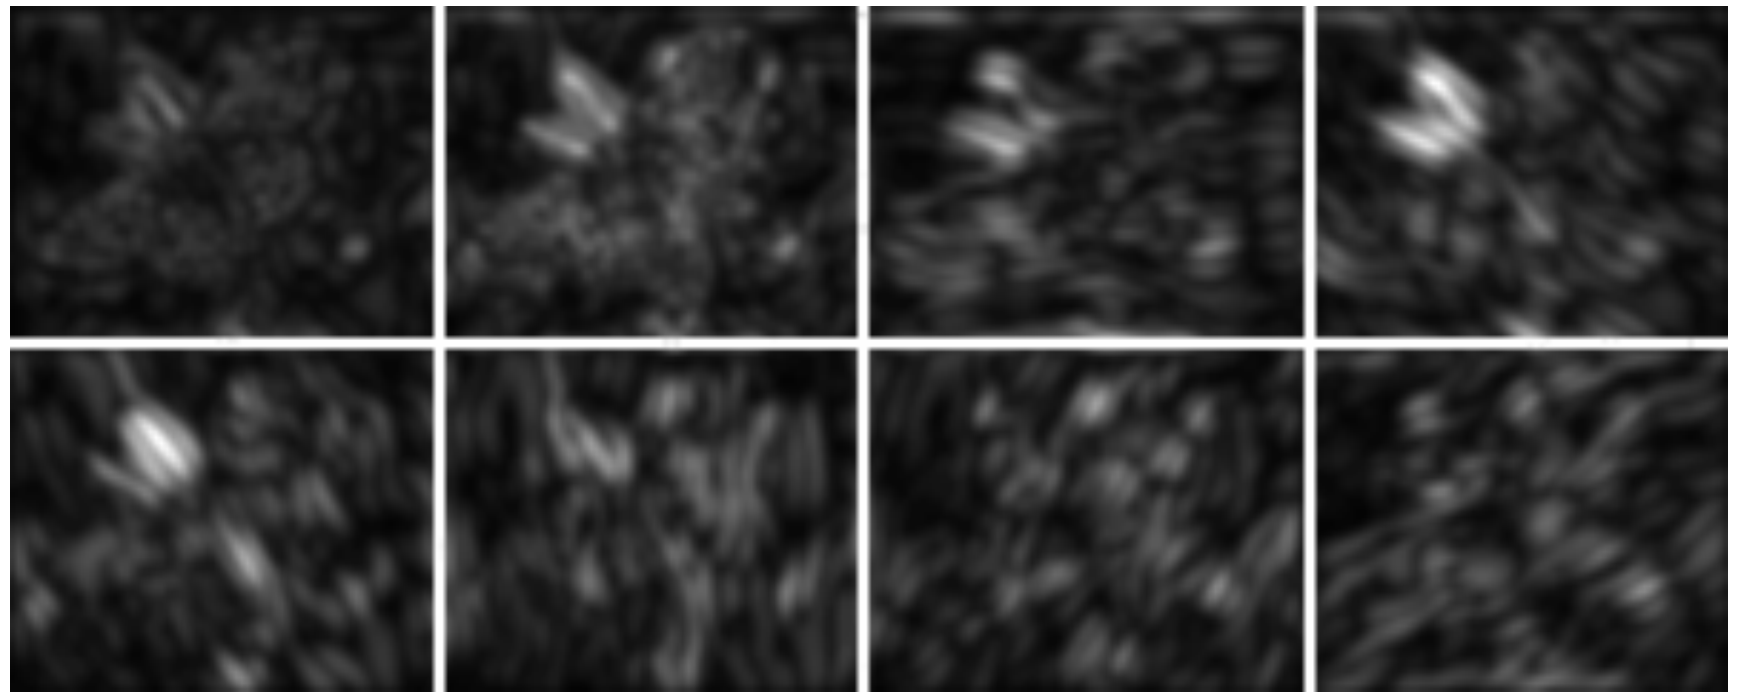

Forsyth & Ponce (1st ed.) Figures 9.3 and 9.5

Gaussian Pyramid

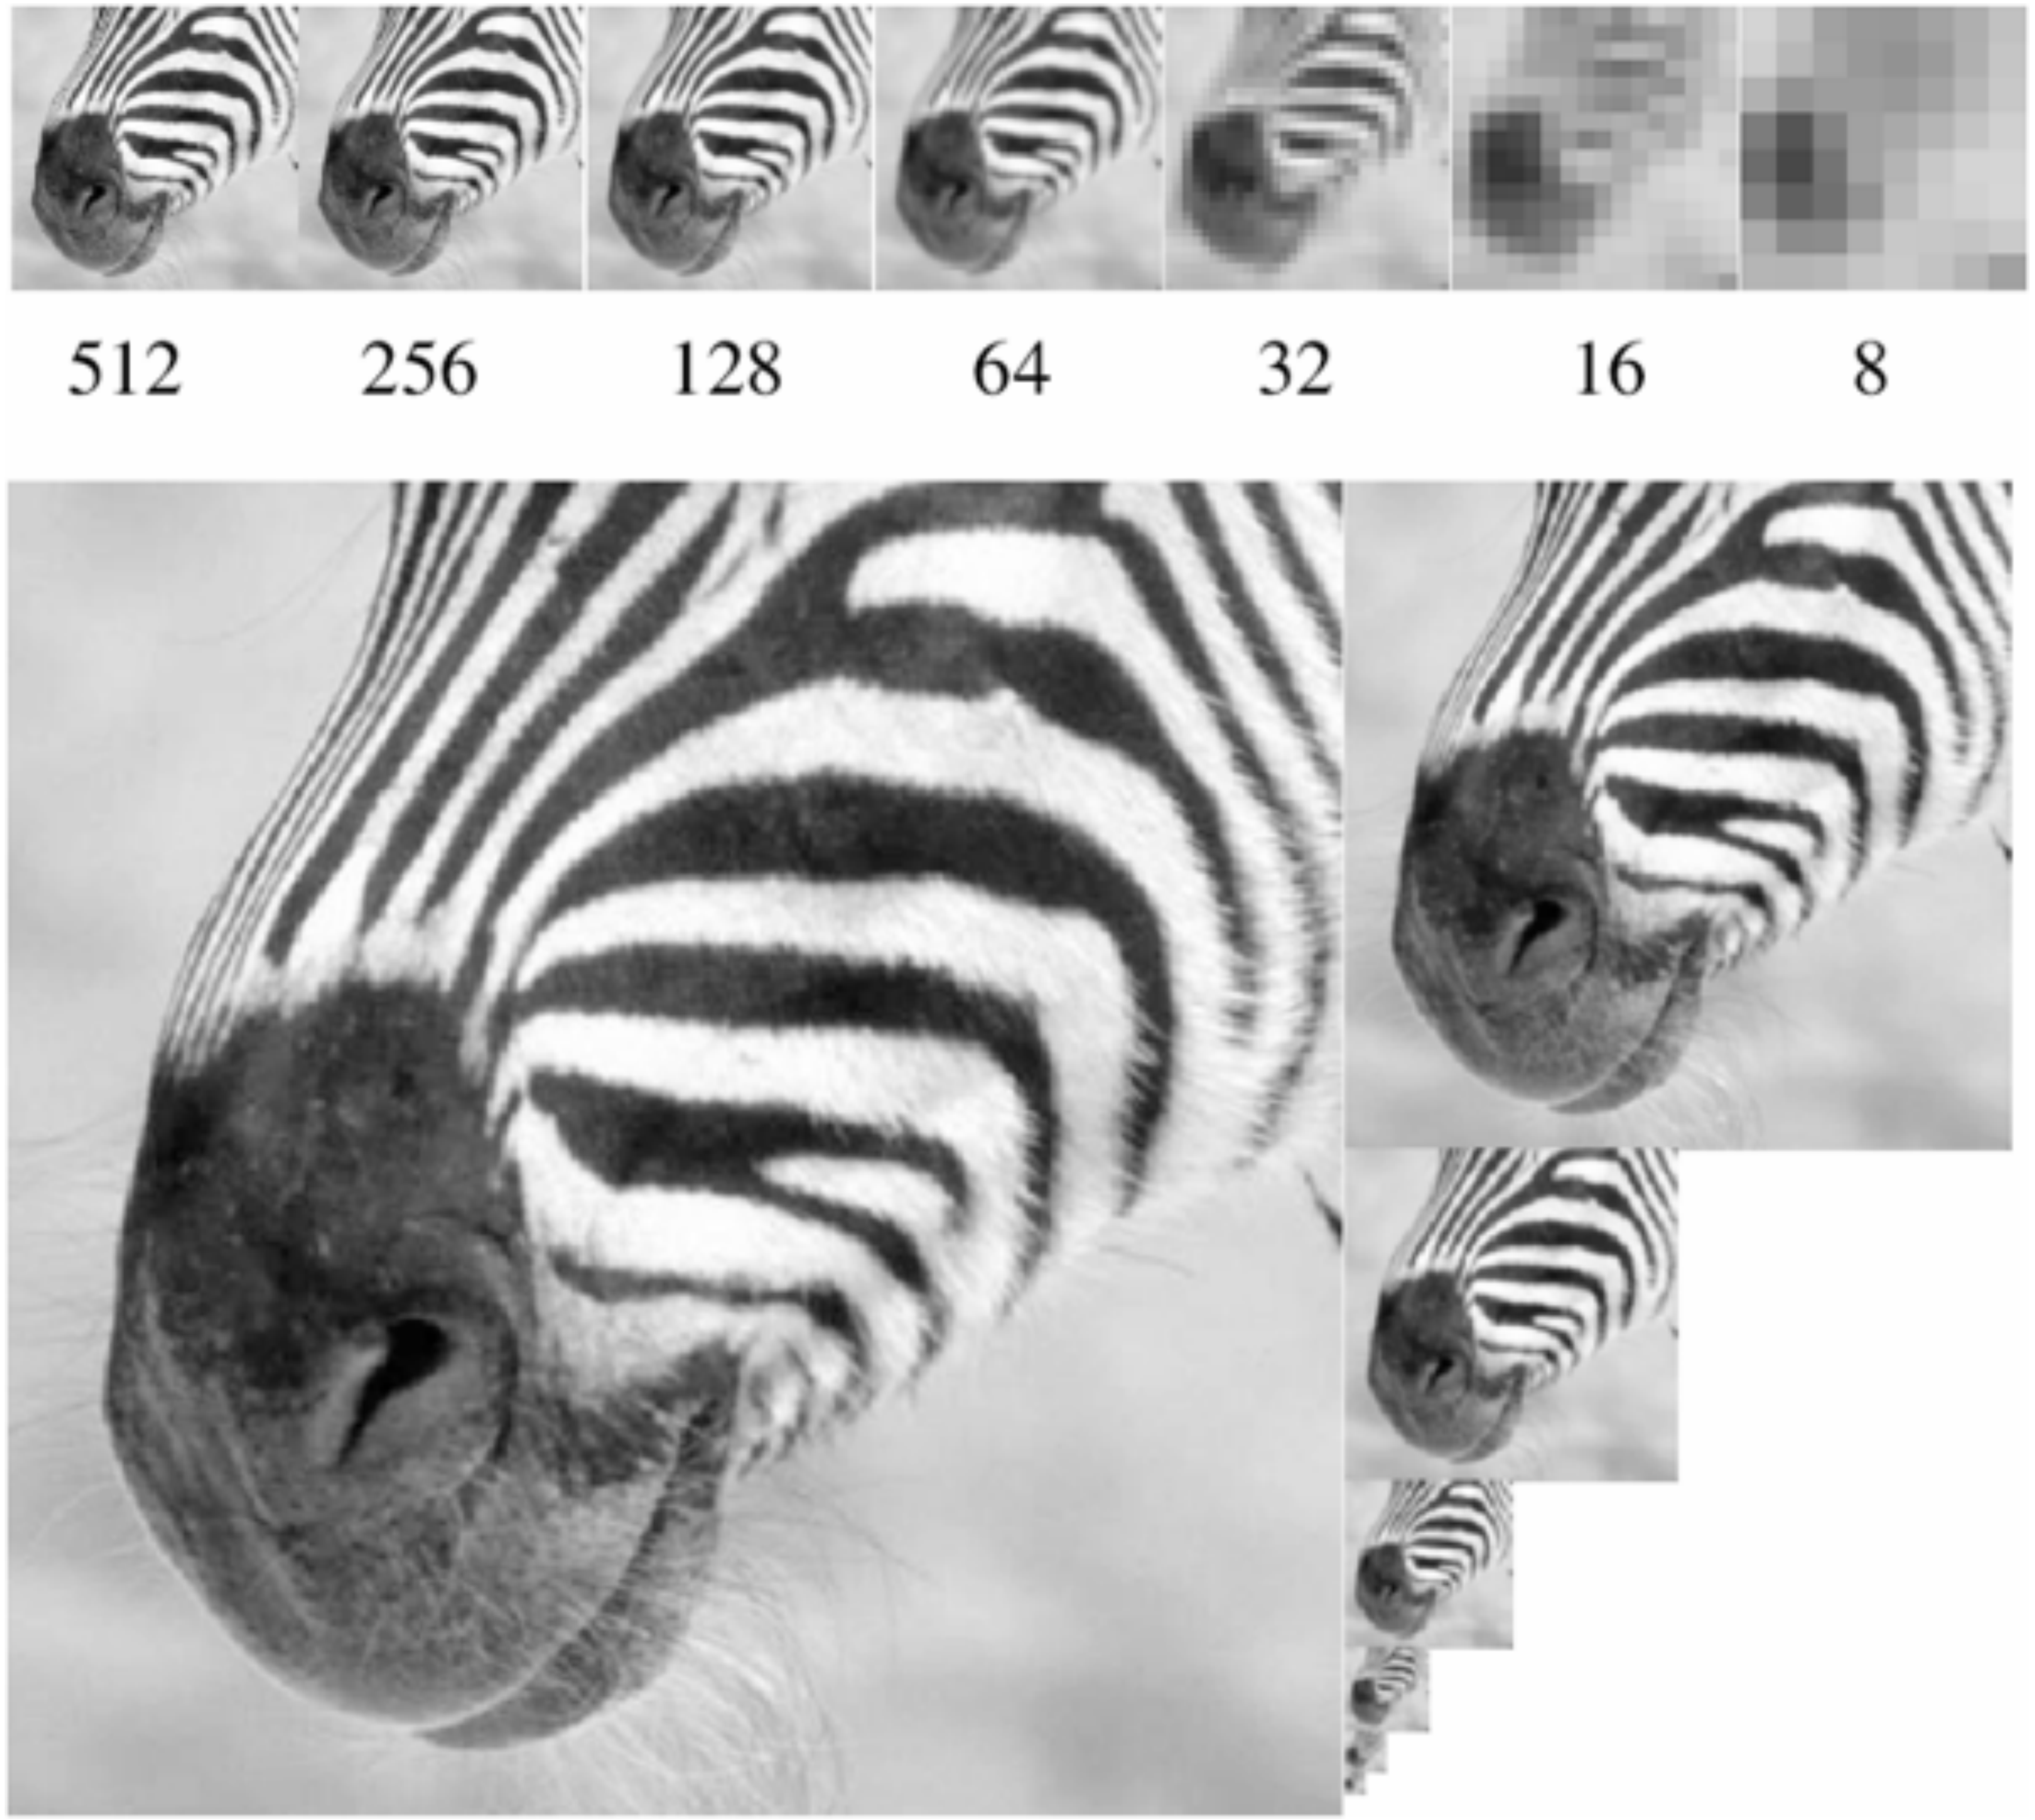

Forsyth & Ponce (2nd ed.) Figure 4.17

What happens to the details?

– They get smoothed out as we move to higher levels

What is preserved at the higher levels?

– Mostly large uniform regions in the original image

How would you reconstruct the original image from the image at the upper level?

– That's not possible

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Laplacian Pyramid

Building a **Laplacian** pyramid:

— Create a Gaussian pyramid

— Take the difference between one Gaussian pyramid level and the next (before subsampling)

Properties

— Also known as the difference-of-Gaussian (DOG) function, a close approximation to the Laplacian

— It is a band pass filter – each level represents a different band of spatial frequencies

Reconstructing the original image:

— Reconstruct the Gaussian pyramid starting at top

Laplacian Pyramid

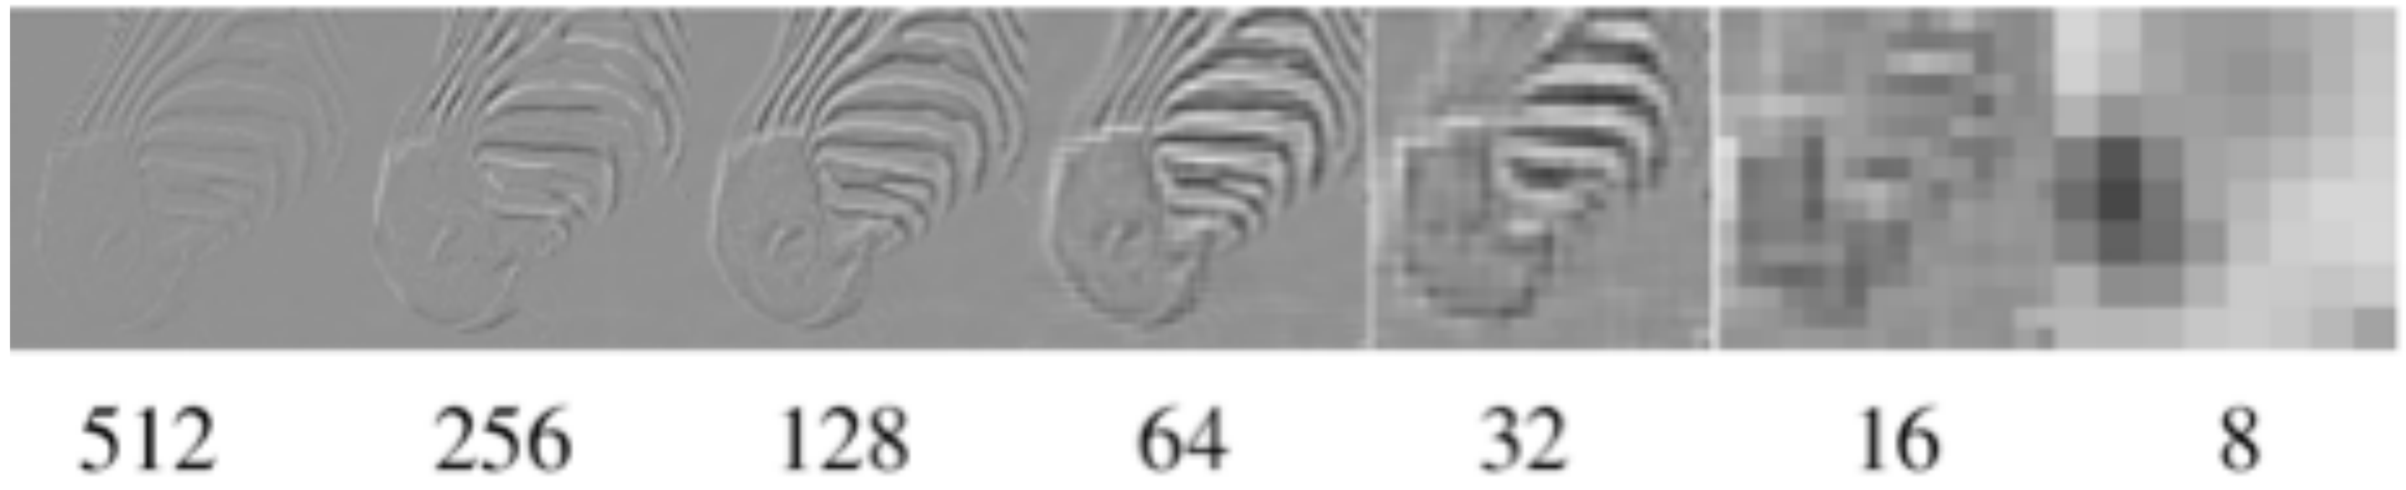

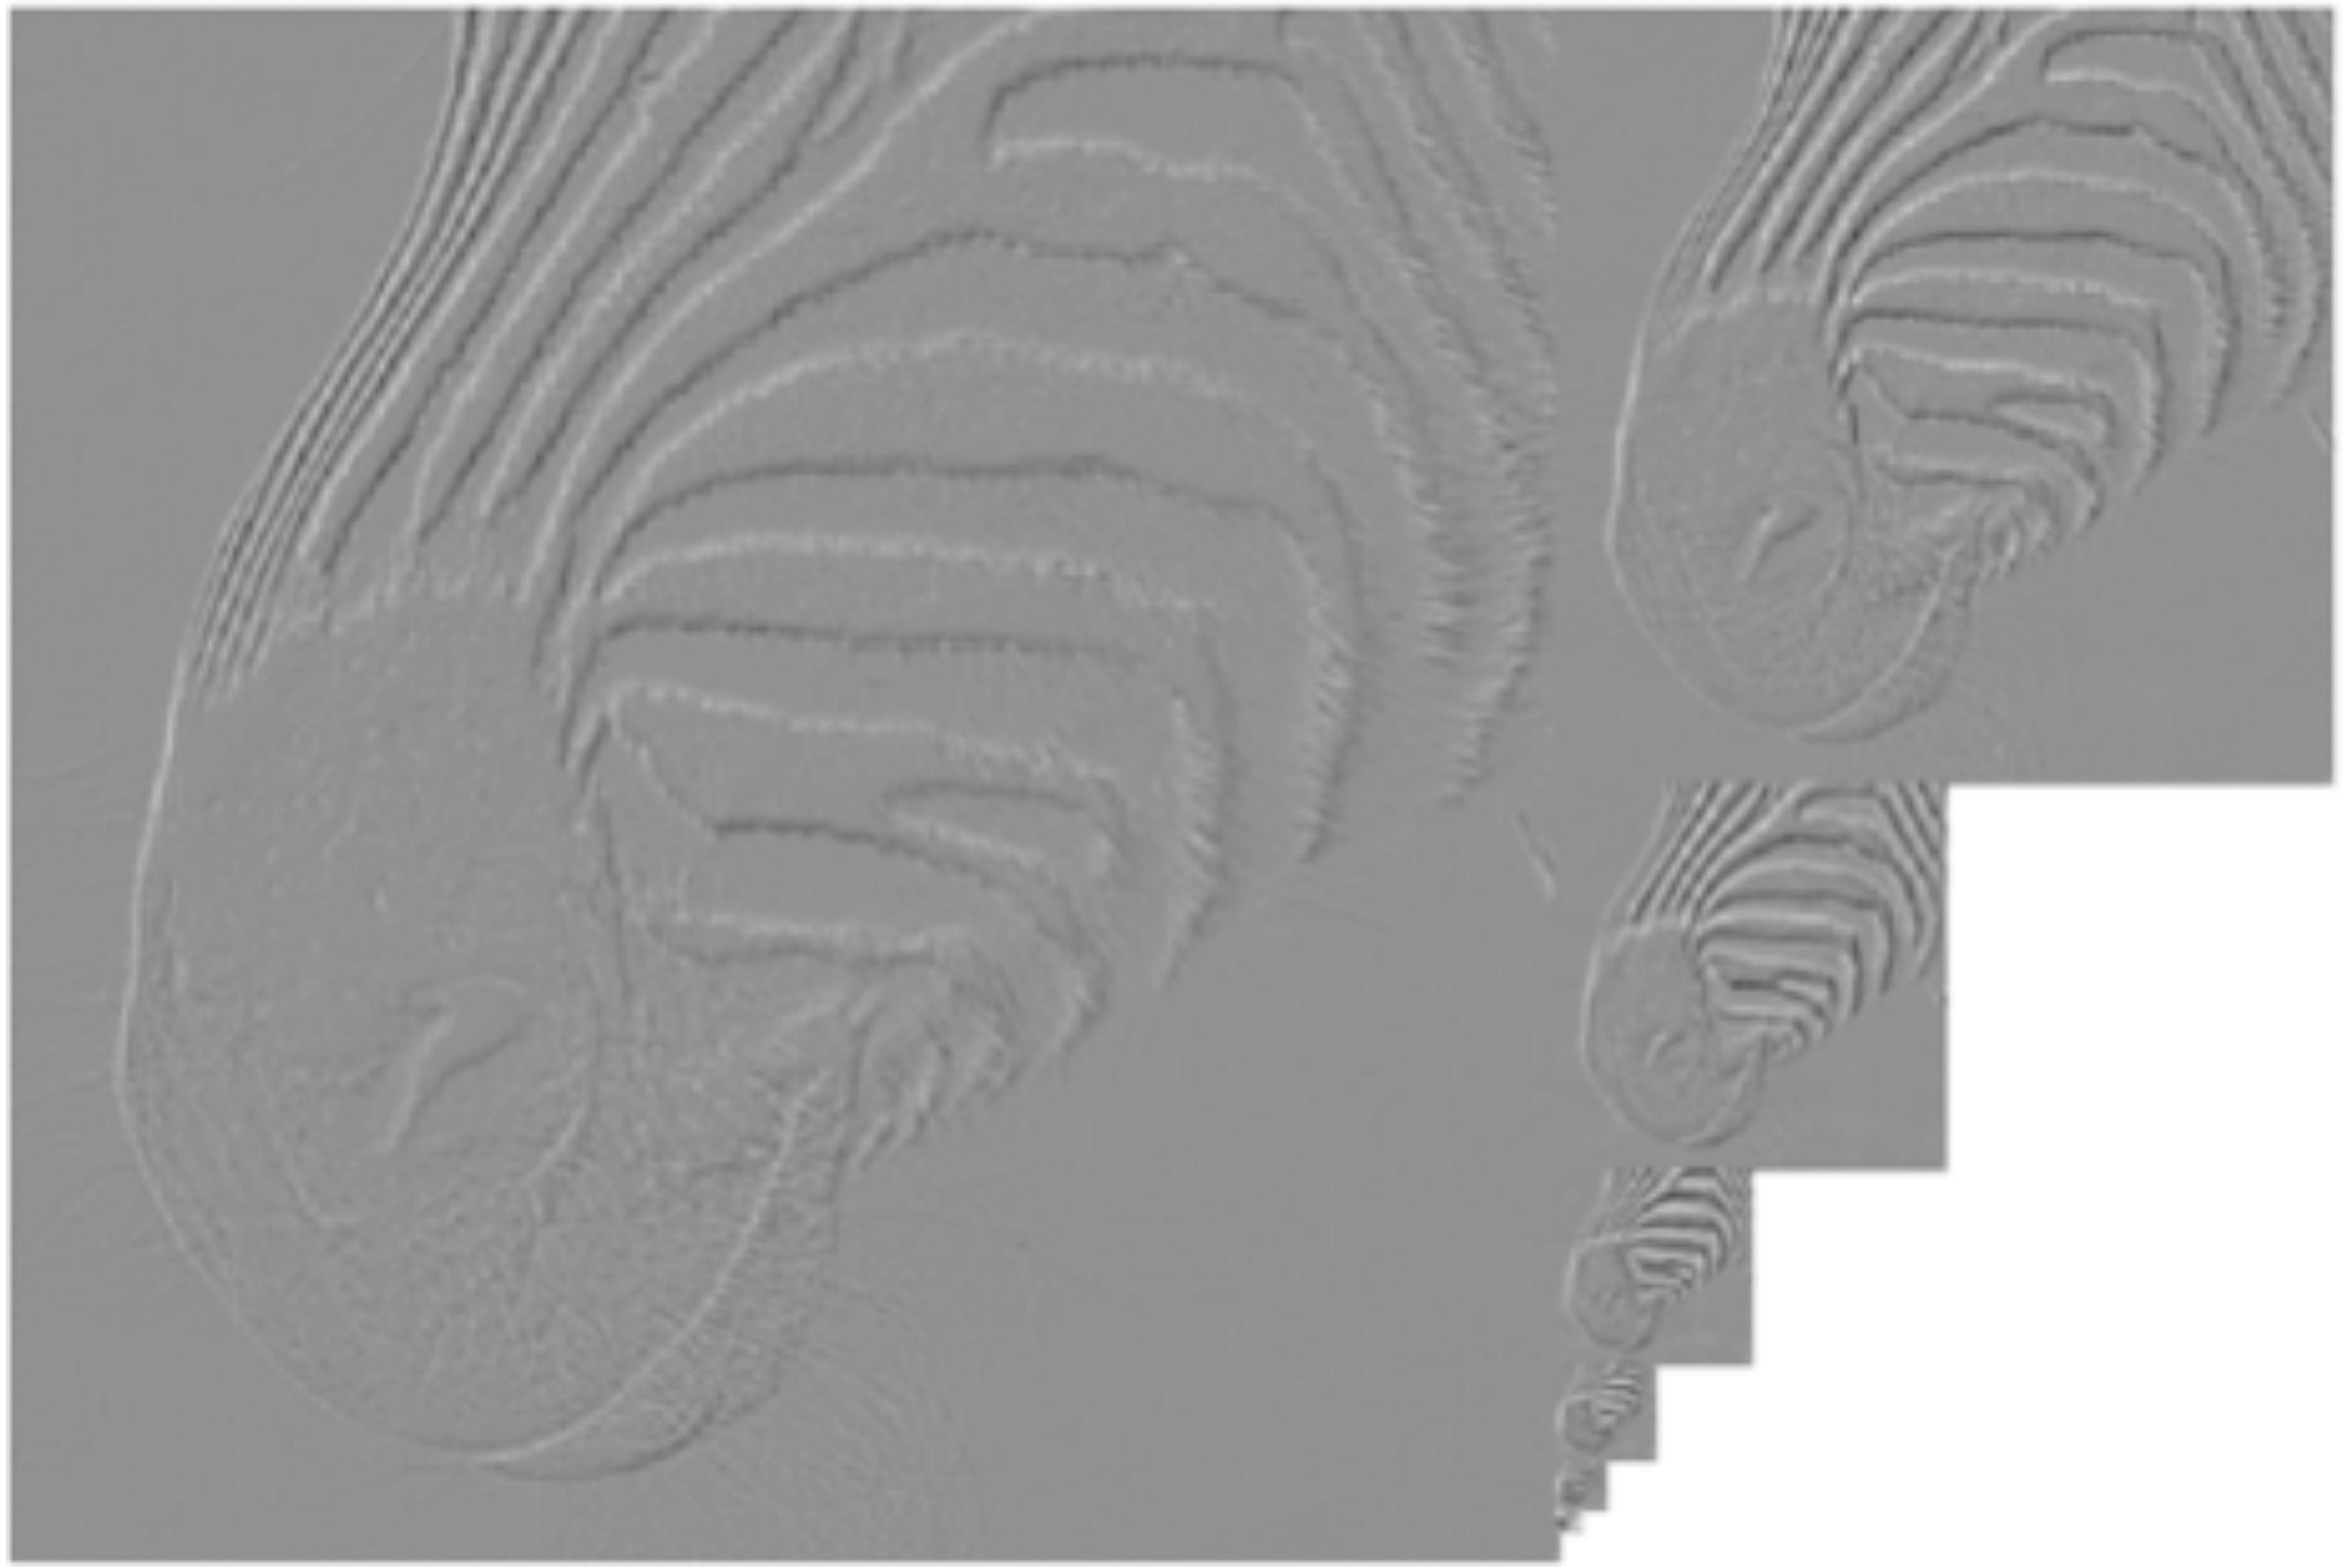

At each level, retain the residuals instead of the blurred images themselves.

Why is it called Laplacian Pyramid?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Why **Laplacian** Pyramid?

-

=

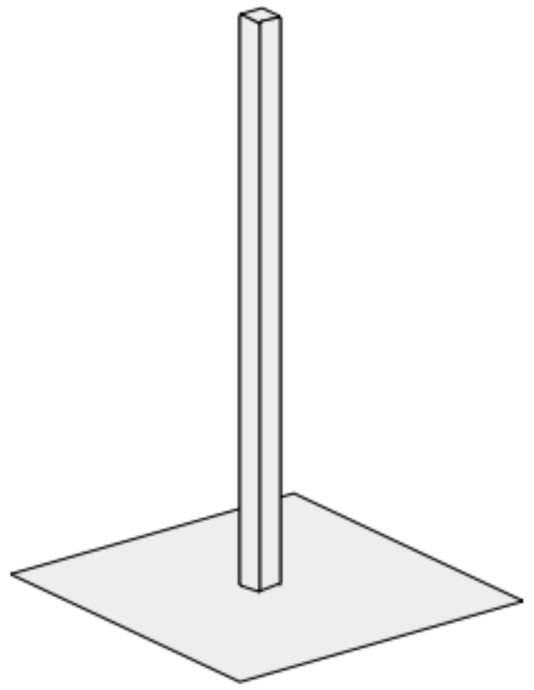

unit

-

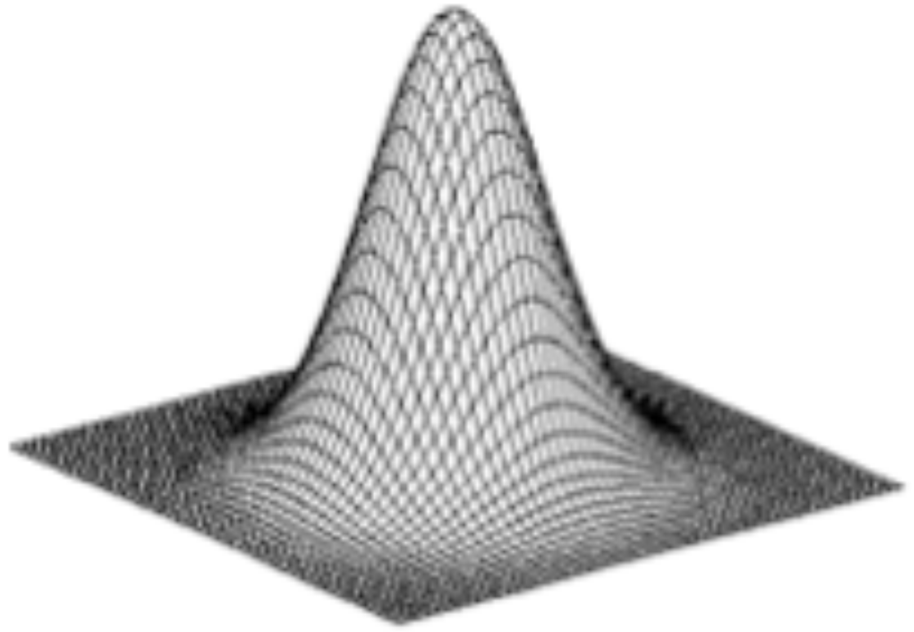

Gaussian

≈

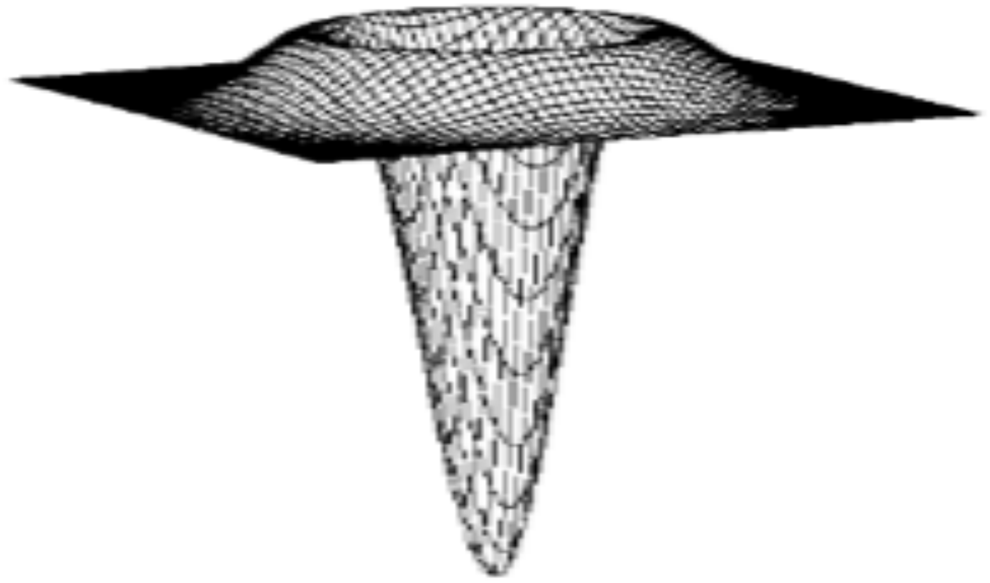

Laplacian

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Laplacian Pyramid

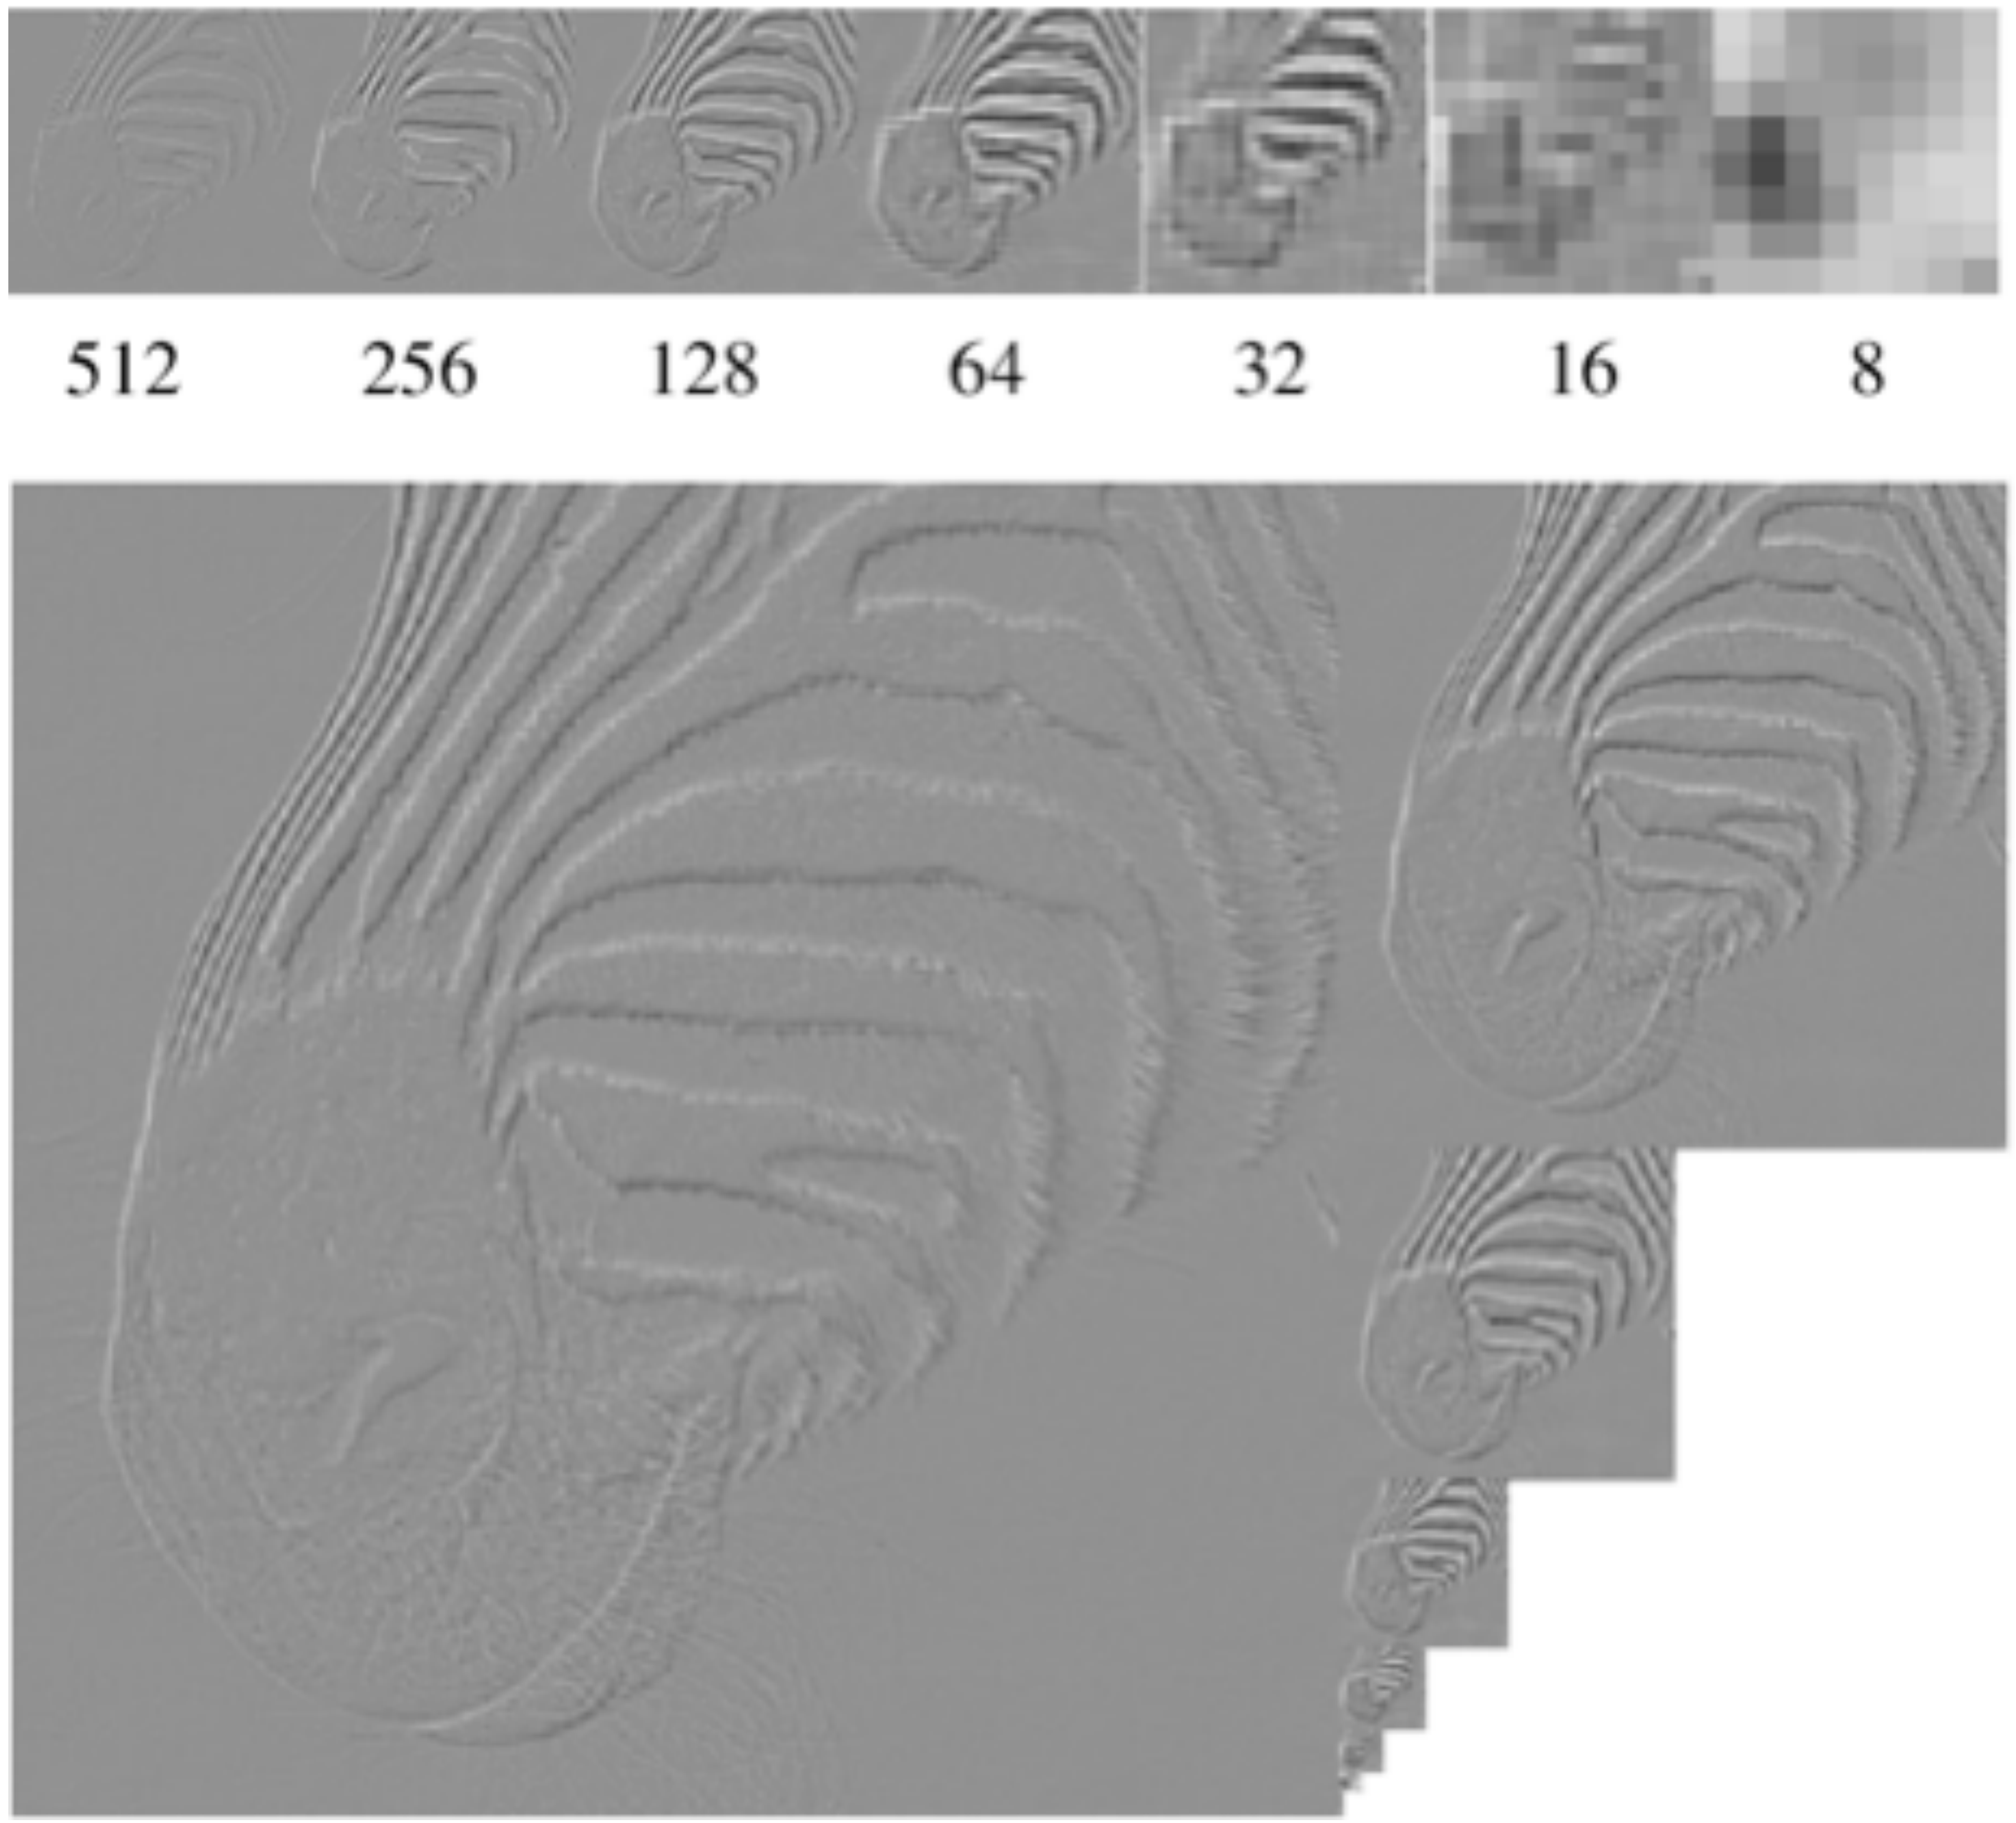

At each level, retain the residuals instead of the blurred images themselves.

Why is it called Laplacian Pyramid?

Can we reconstruct the original image using the pyramid?
— Yes we can!

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Laplacian Pyramid

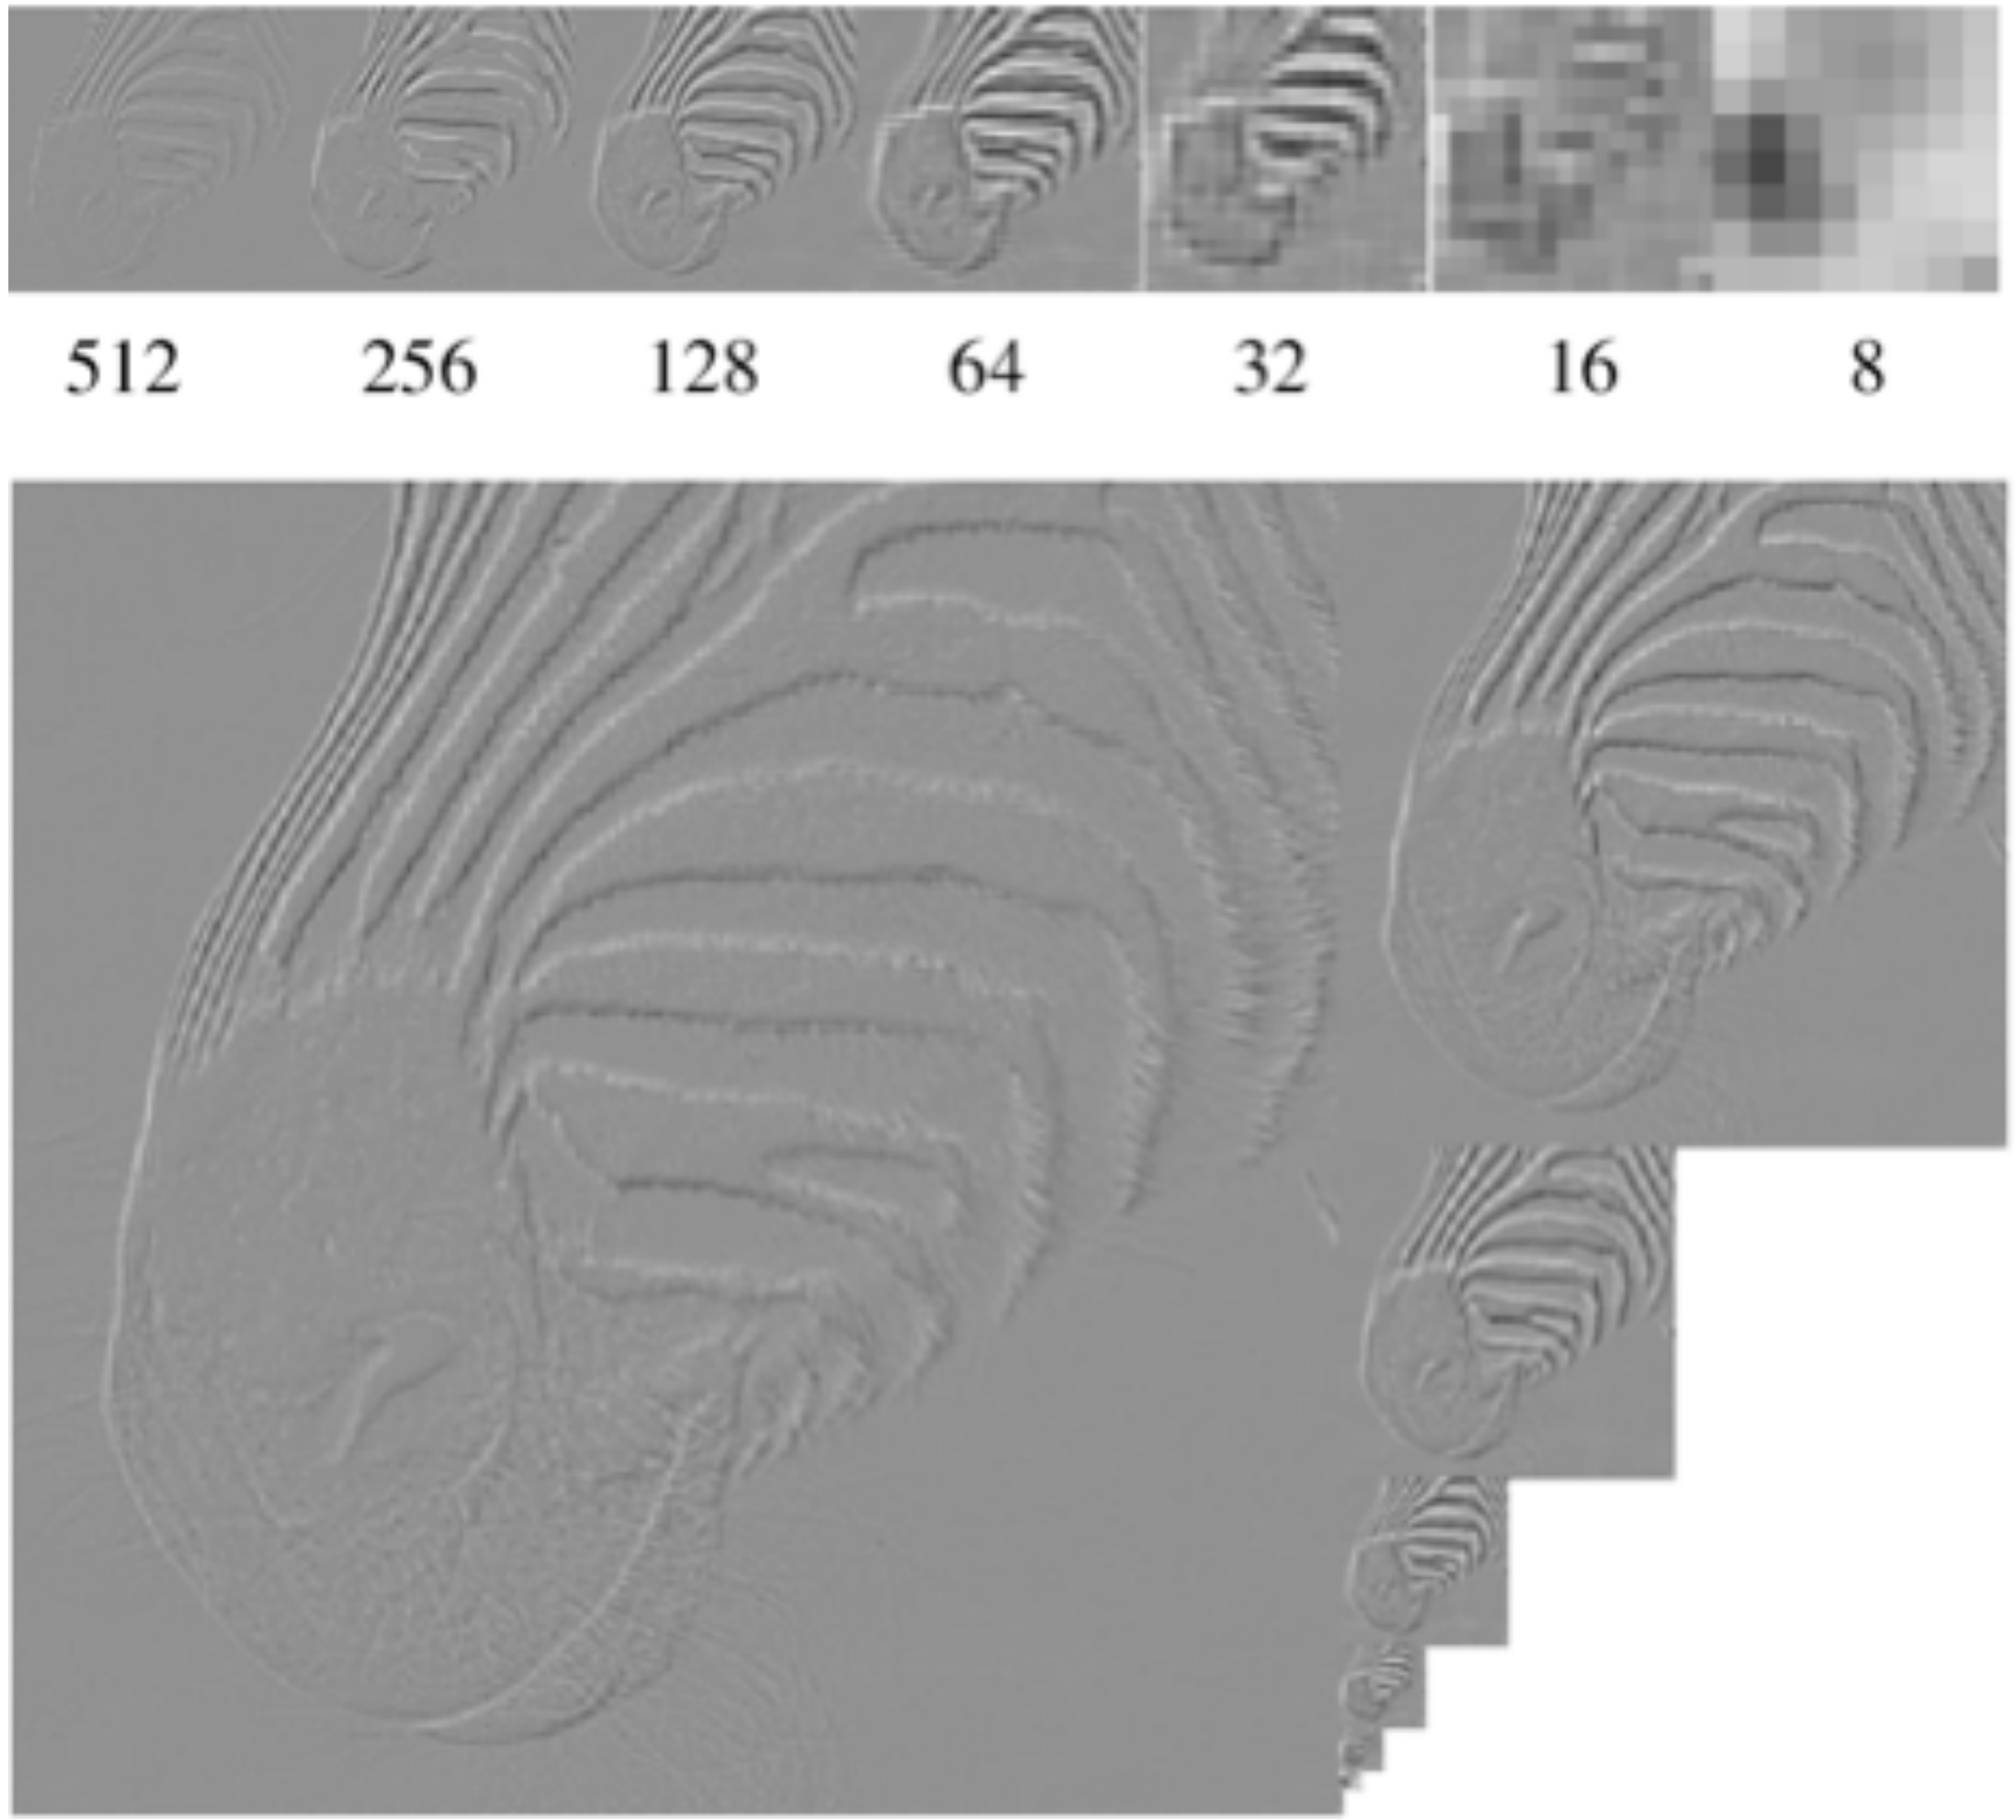

At each level, retain the residuals instead of the blurred images themselves.

Why is it called Laplacian Pyramid?

Can we reconstruct the original image using the pyramid?
— Yes we can!

What do we need to store to be able to reconstruct the original image?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Let's start by just looking at **one level**

=

+

level 0 | level 1 (upsampled) | residual

Does this mean we need to store both residuals and the blurred copies of the original?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Constructing a **Laplacian** Pyramid

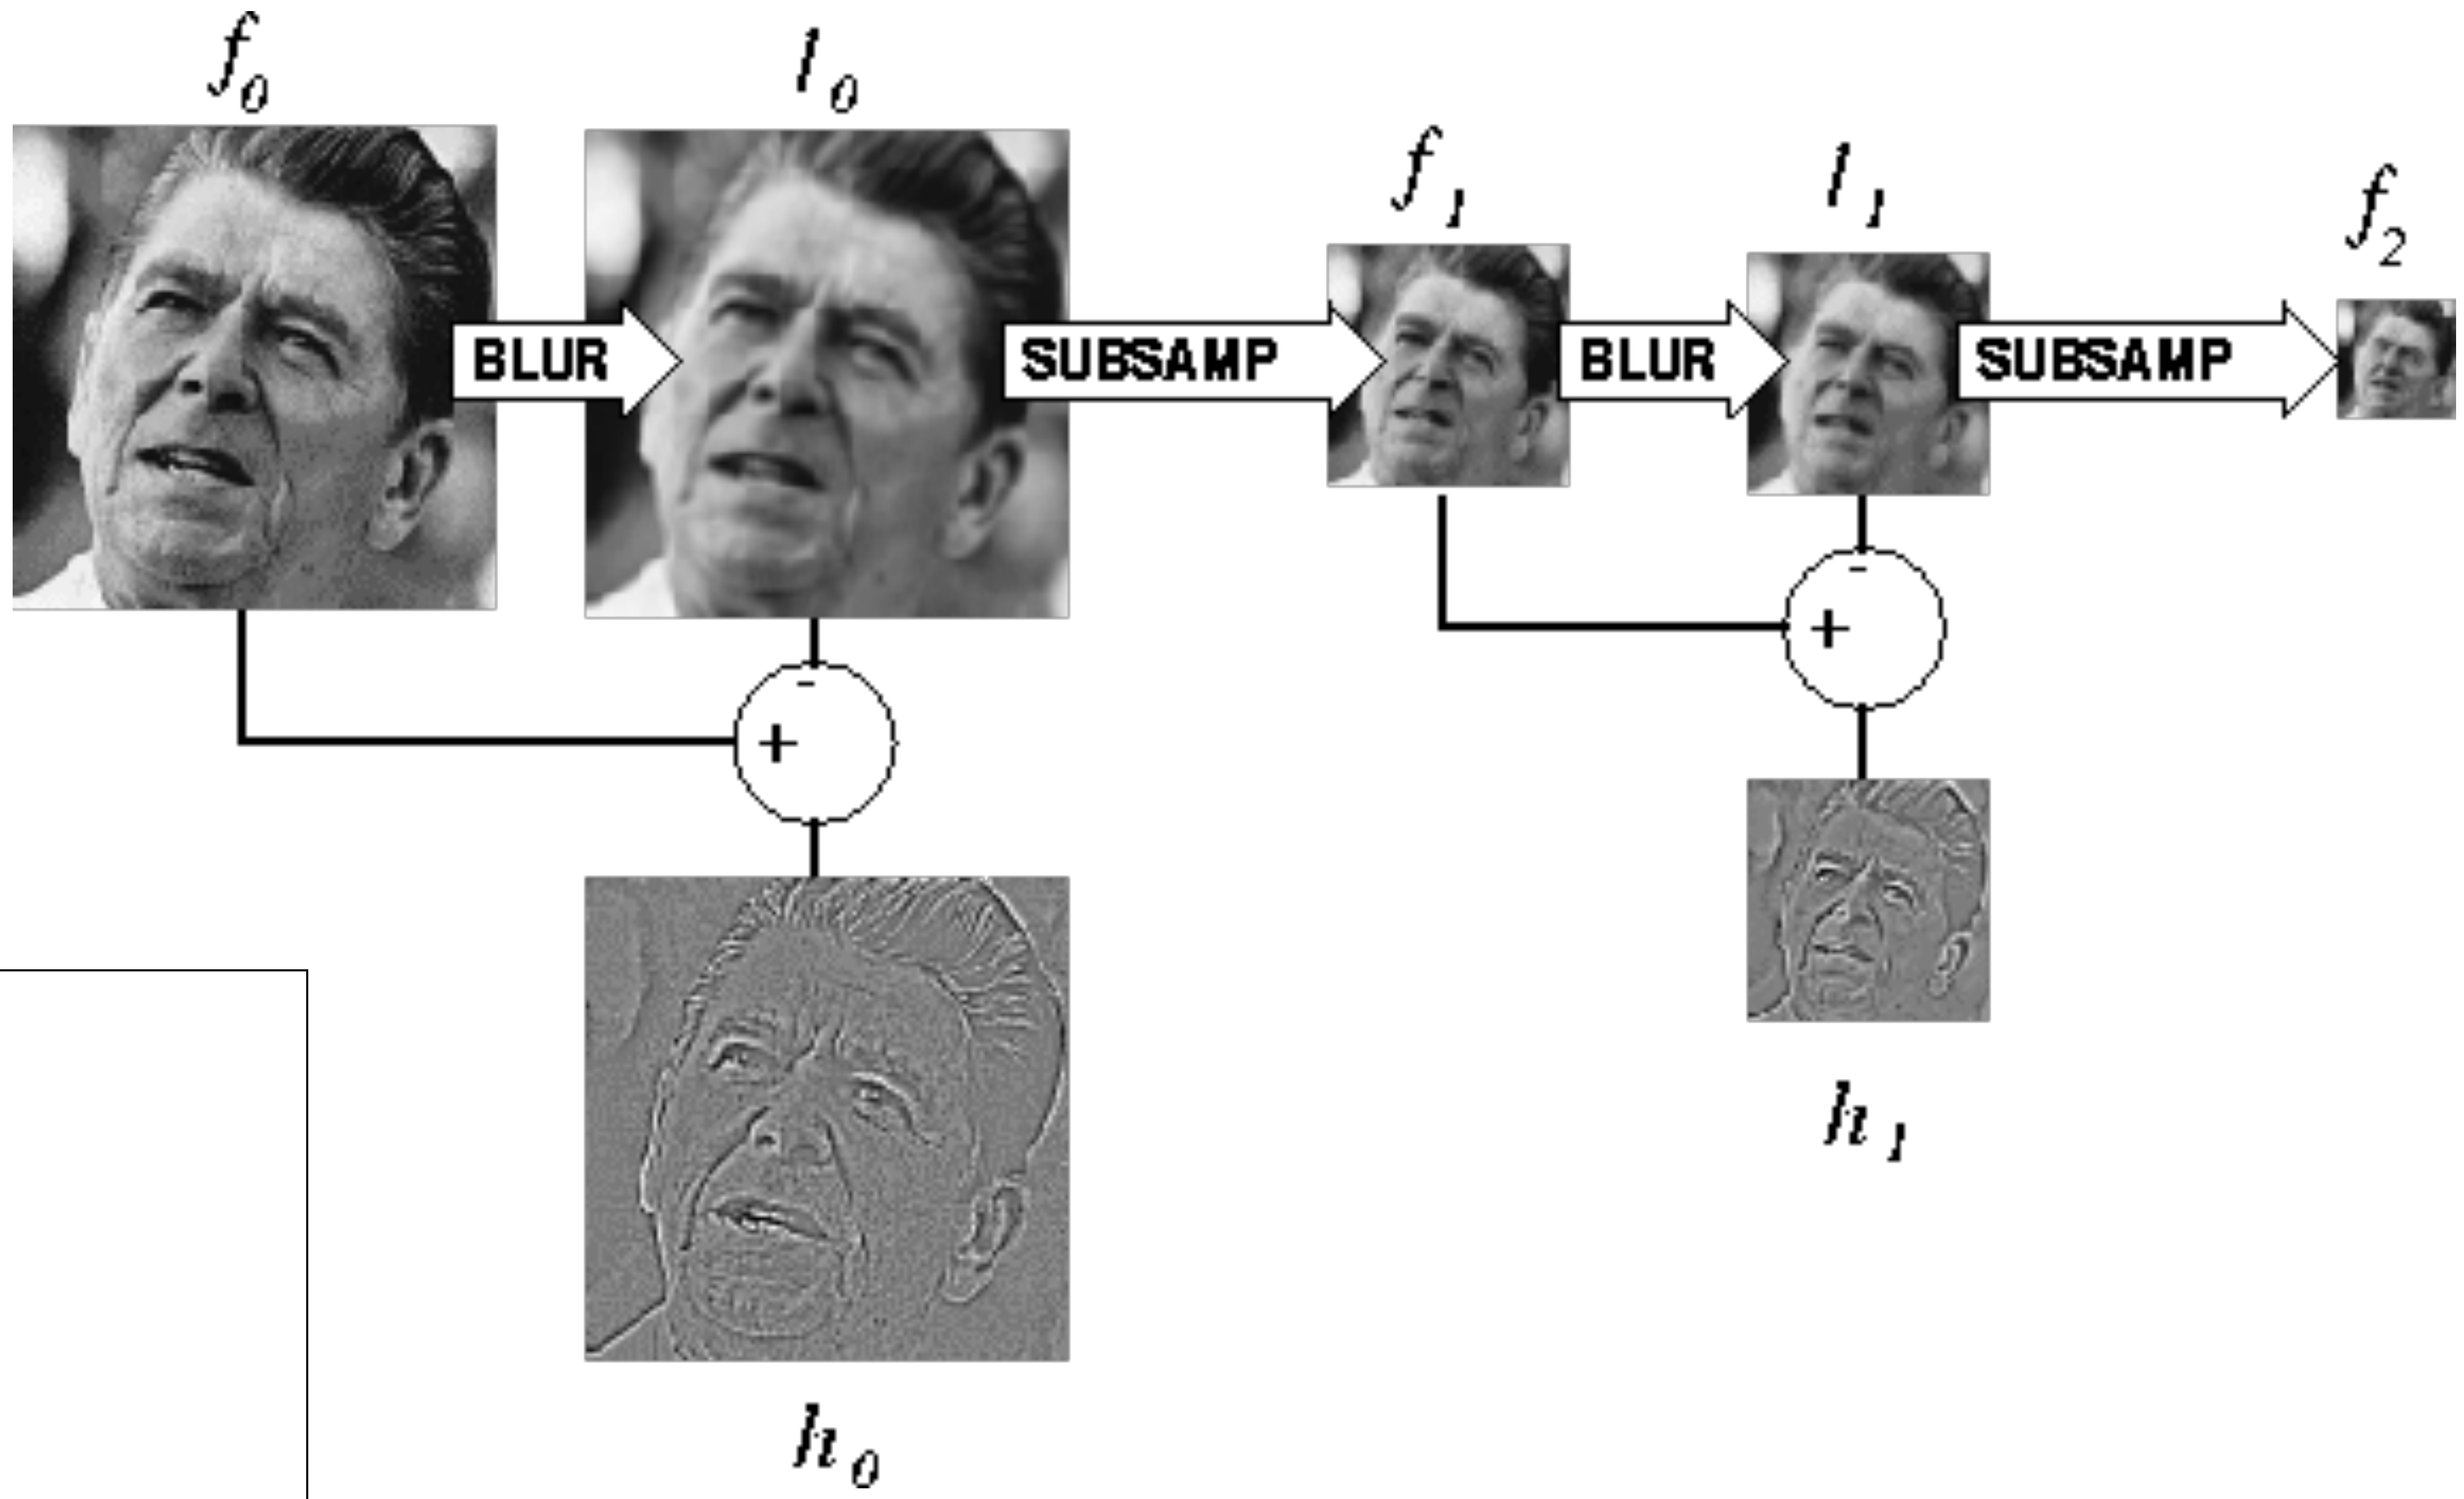

Algorithm

repeat:
    filter
    compute residual
    subsample
until min resolution reached

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Constructing a **Laplacian** Pyramid

What is this part?

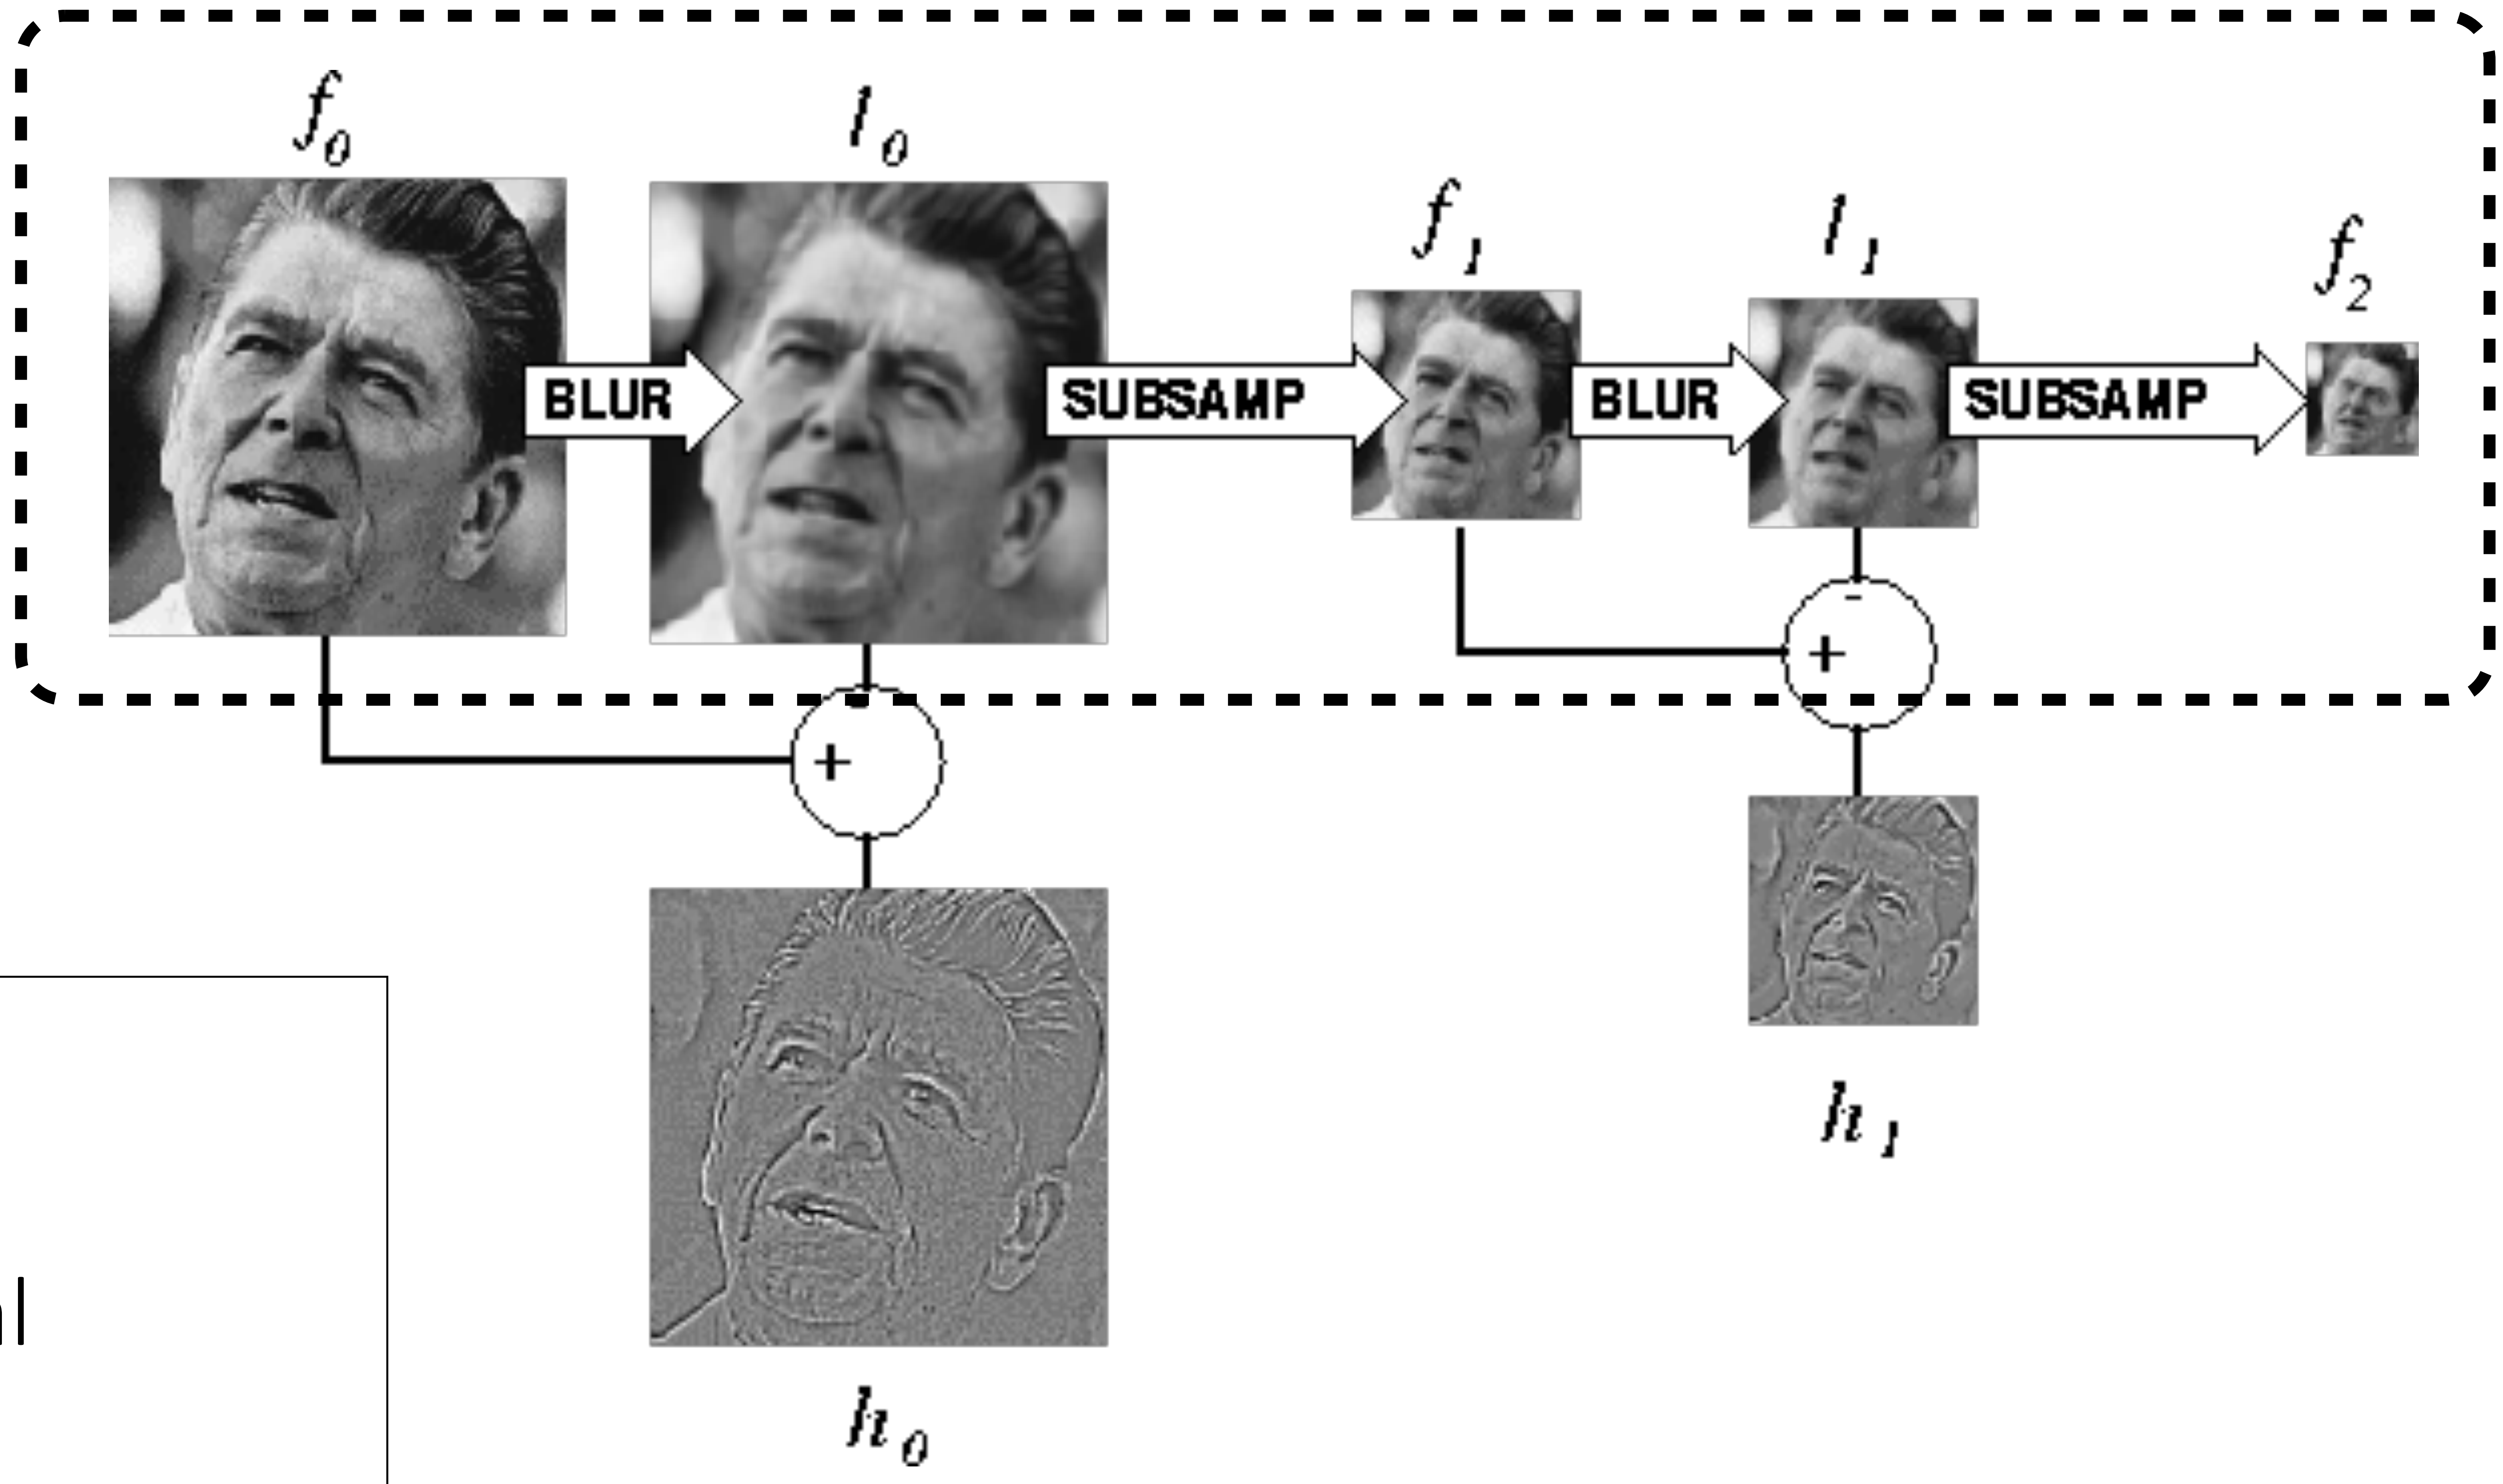

Algorithm

repeat:
  filter
  compute residual
  subsample
until min resolution reached

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Constructing a **Laplacian** Pyramid

It's a Gaussian Pyramid

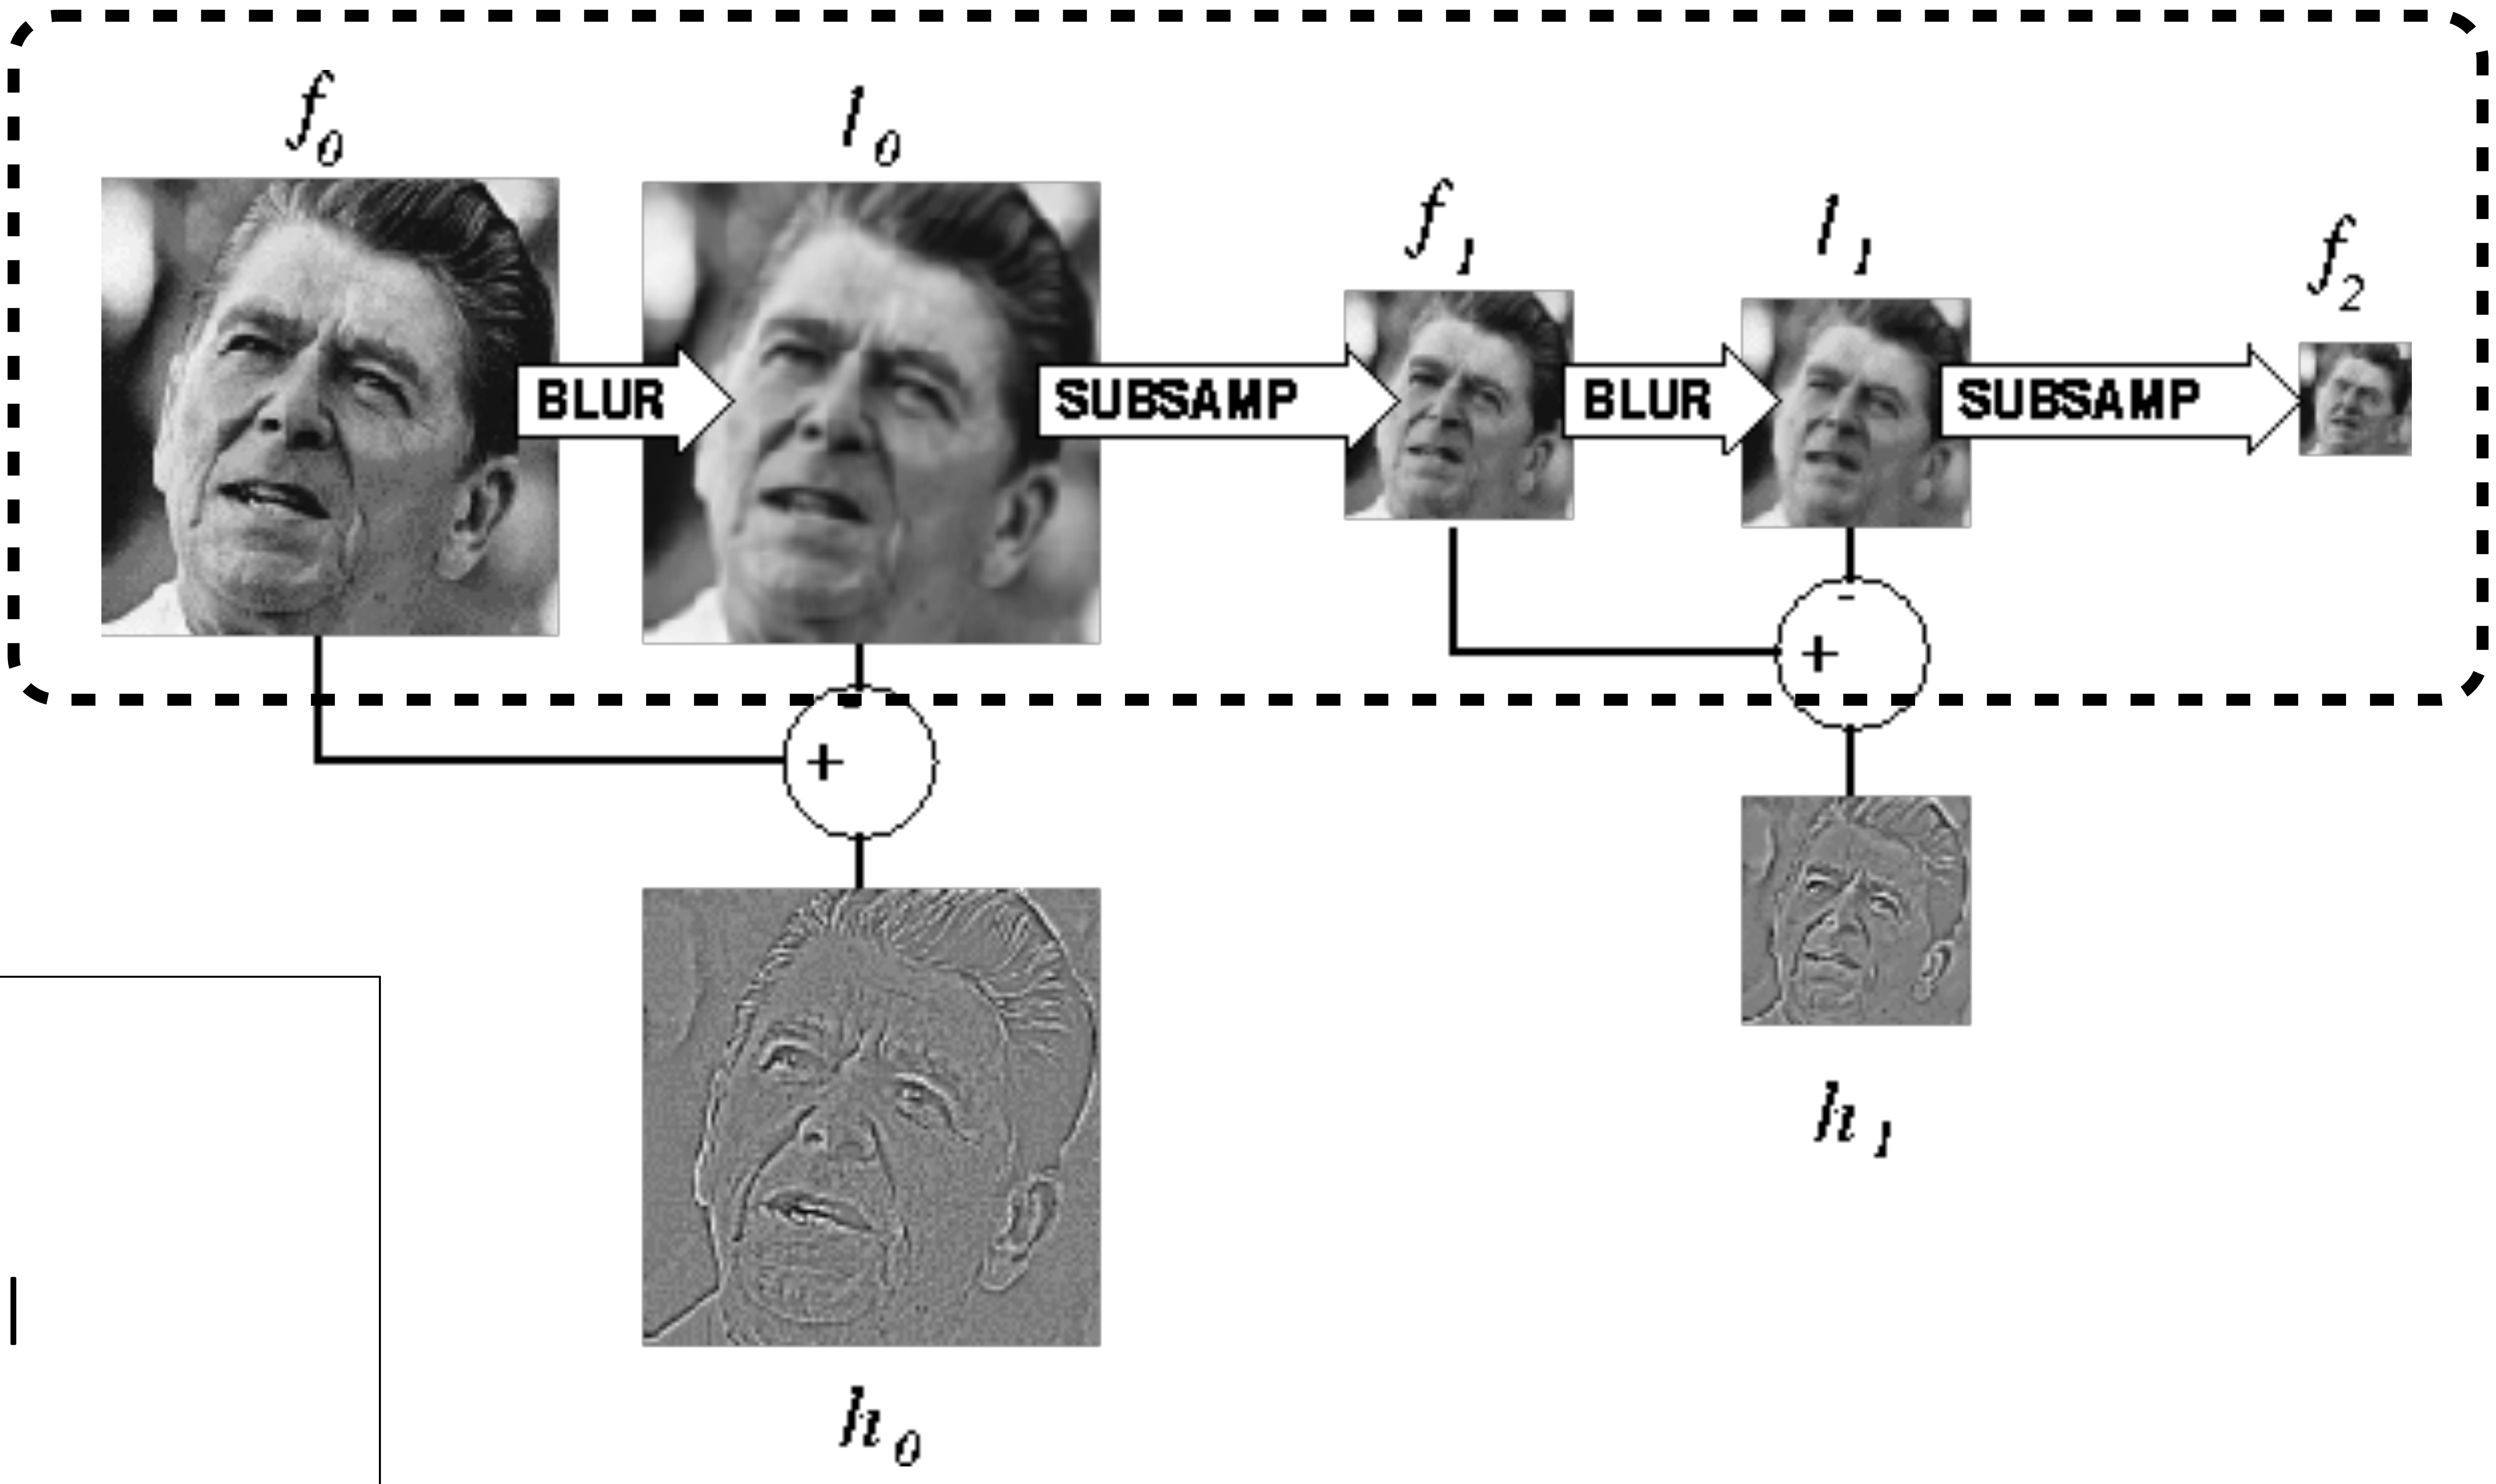

Algorithm

repeat:

filter

compute residual

subsample

until min resolution reached

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Reconstructing the Original Image

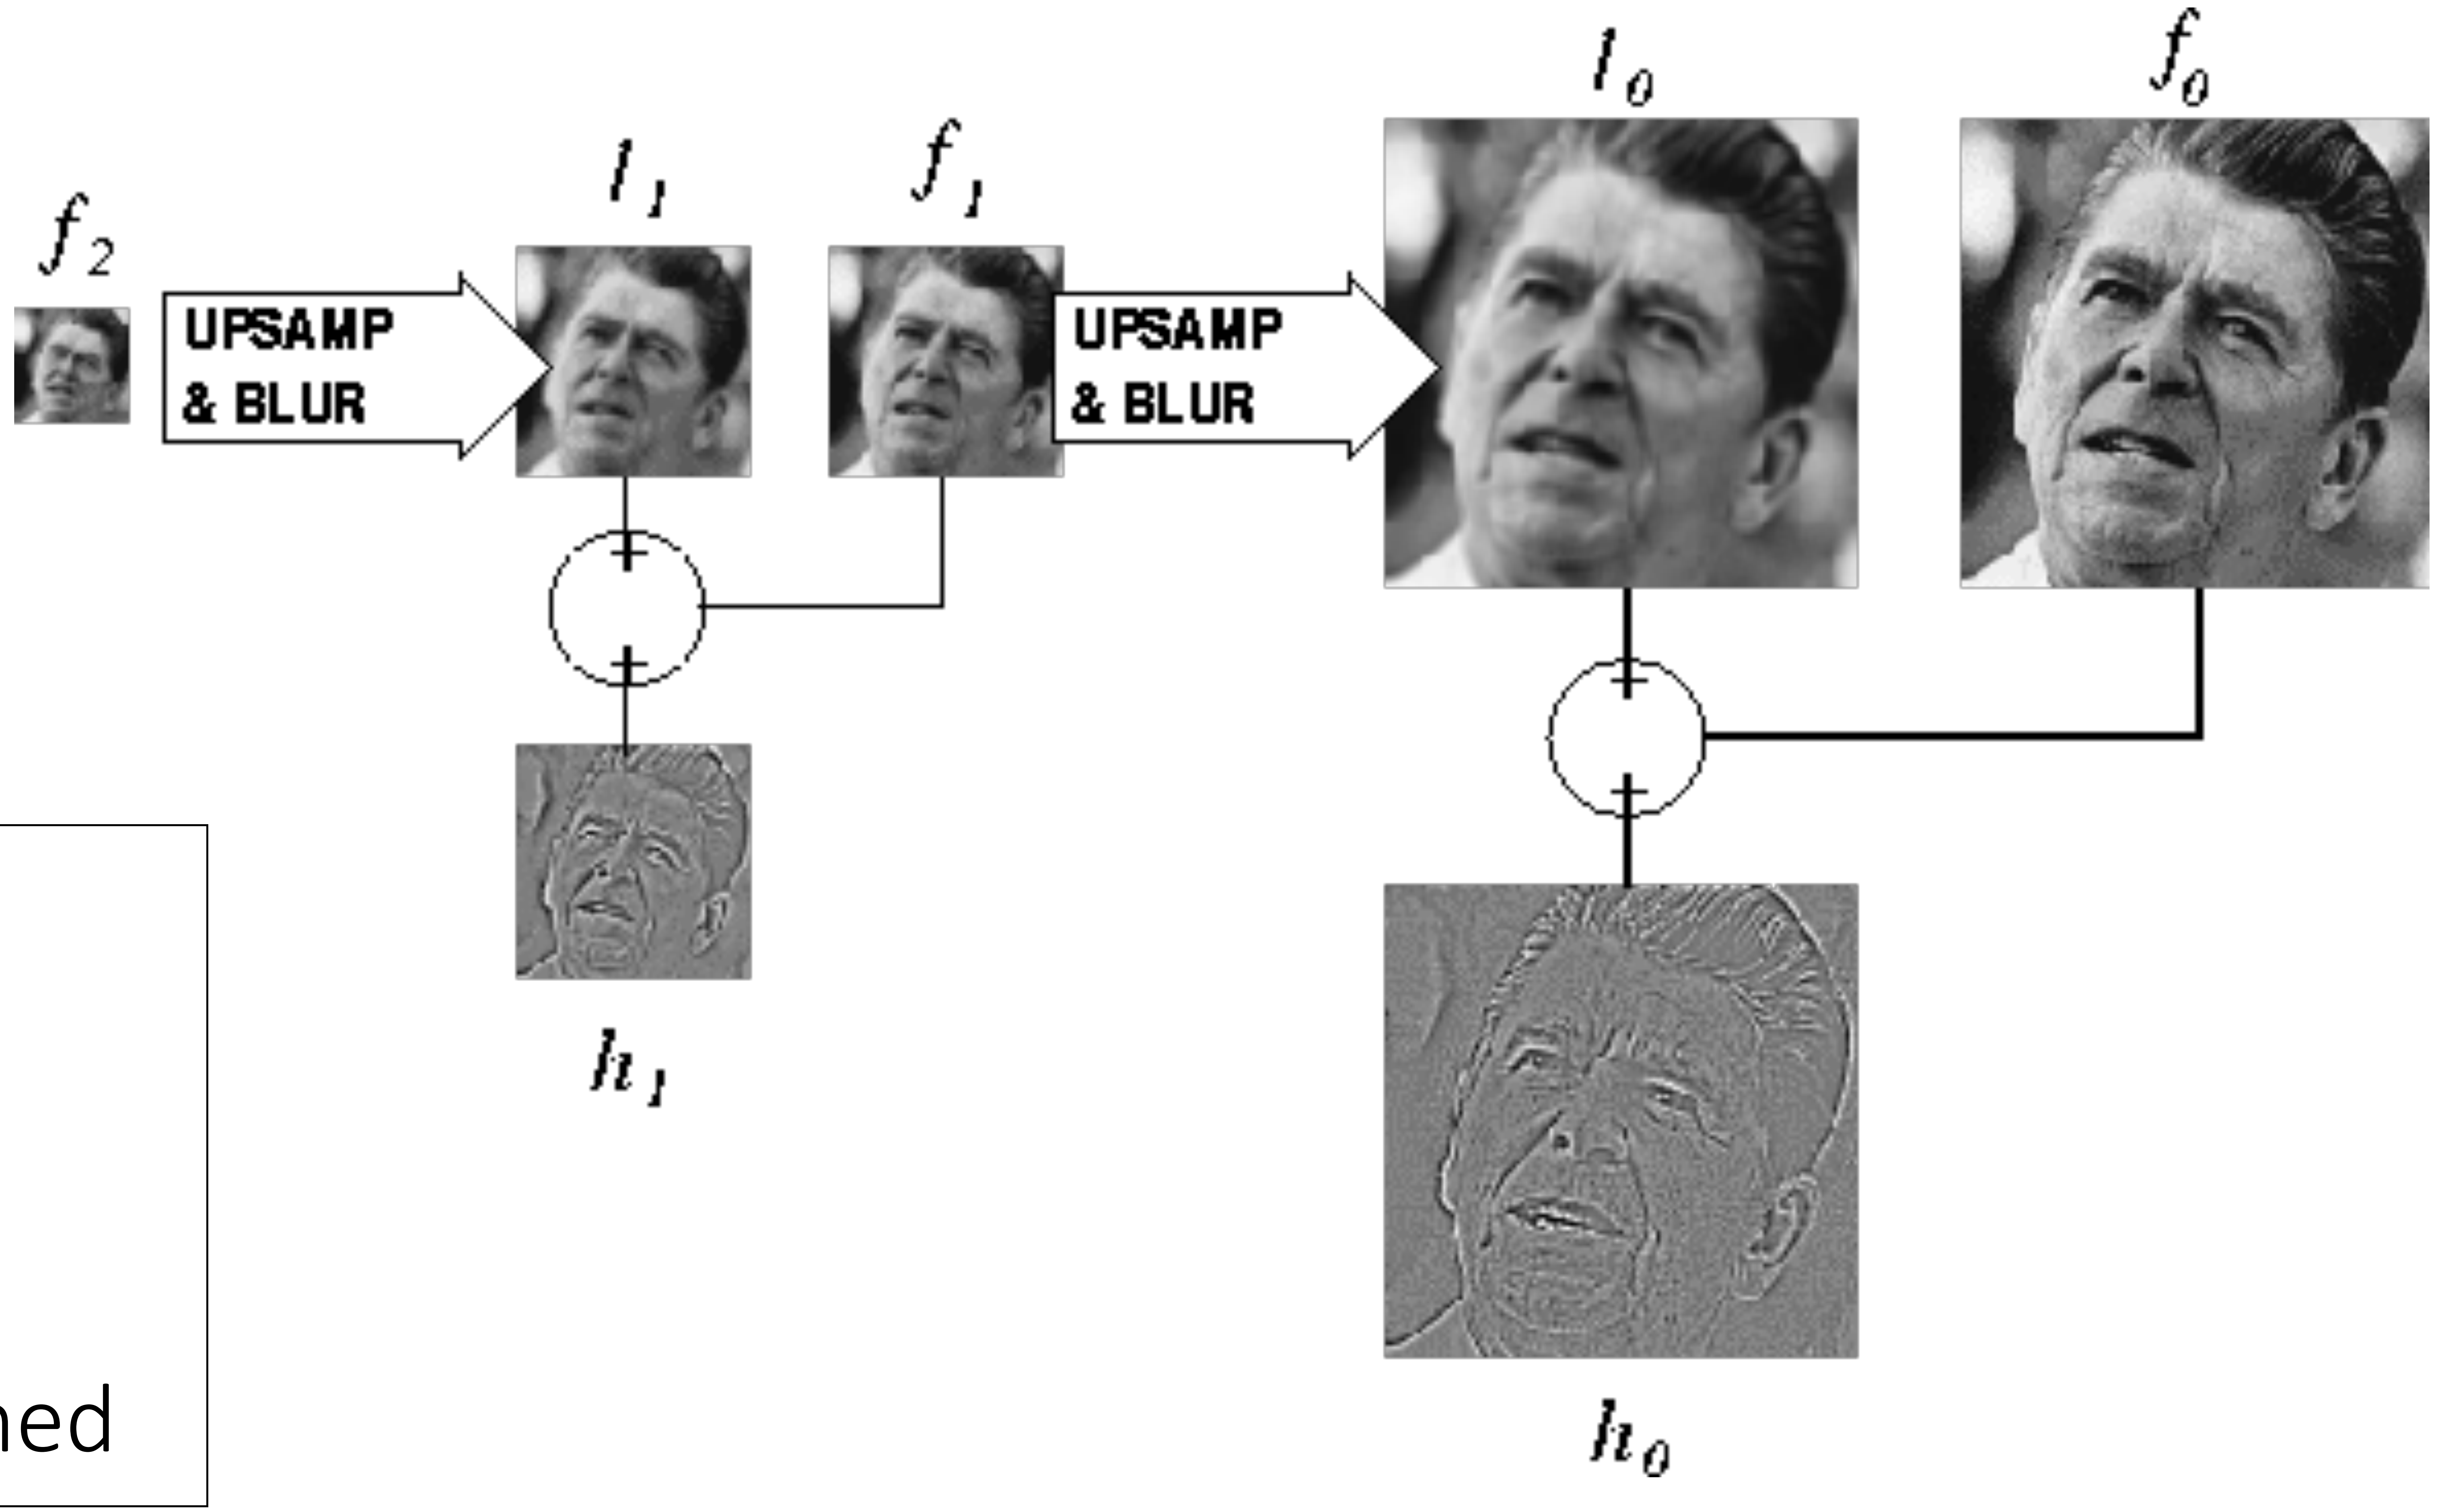

Algorithm

repeat:

upsample

sum with residual

until orig resolution reached

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Gaussian vs **Laplacian** Pyramid




Shown in opposite order for space

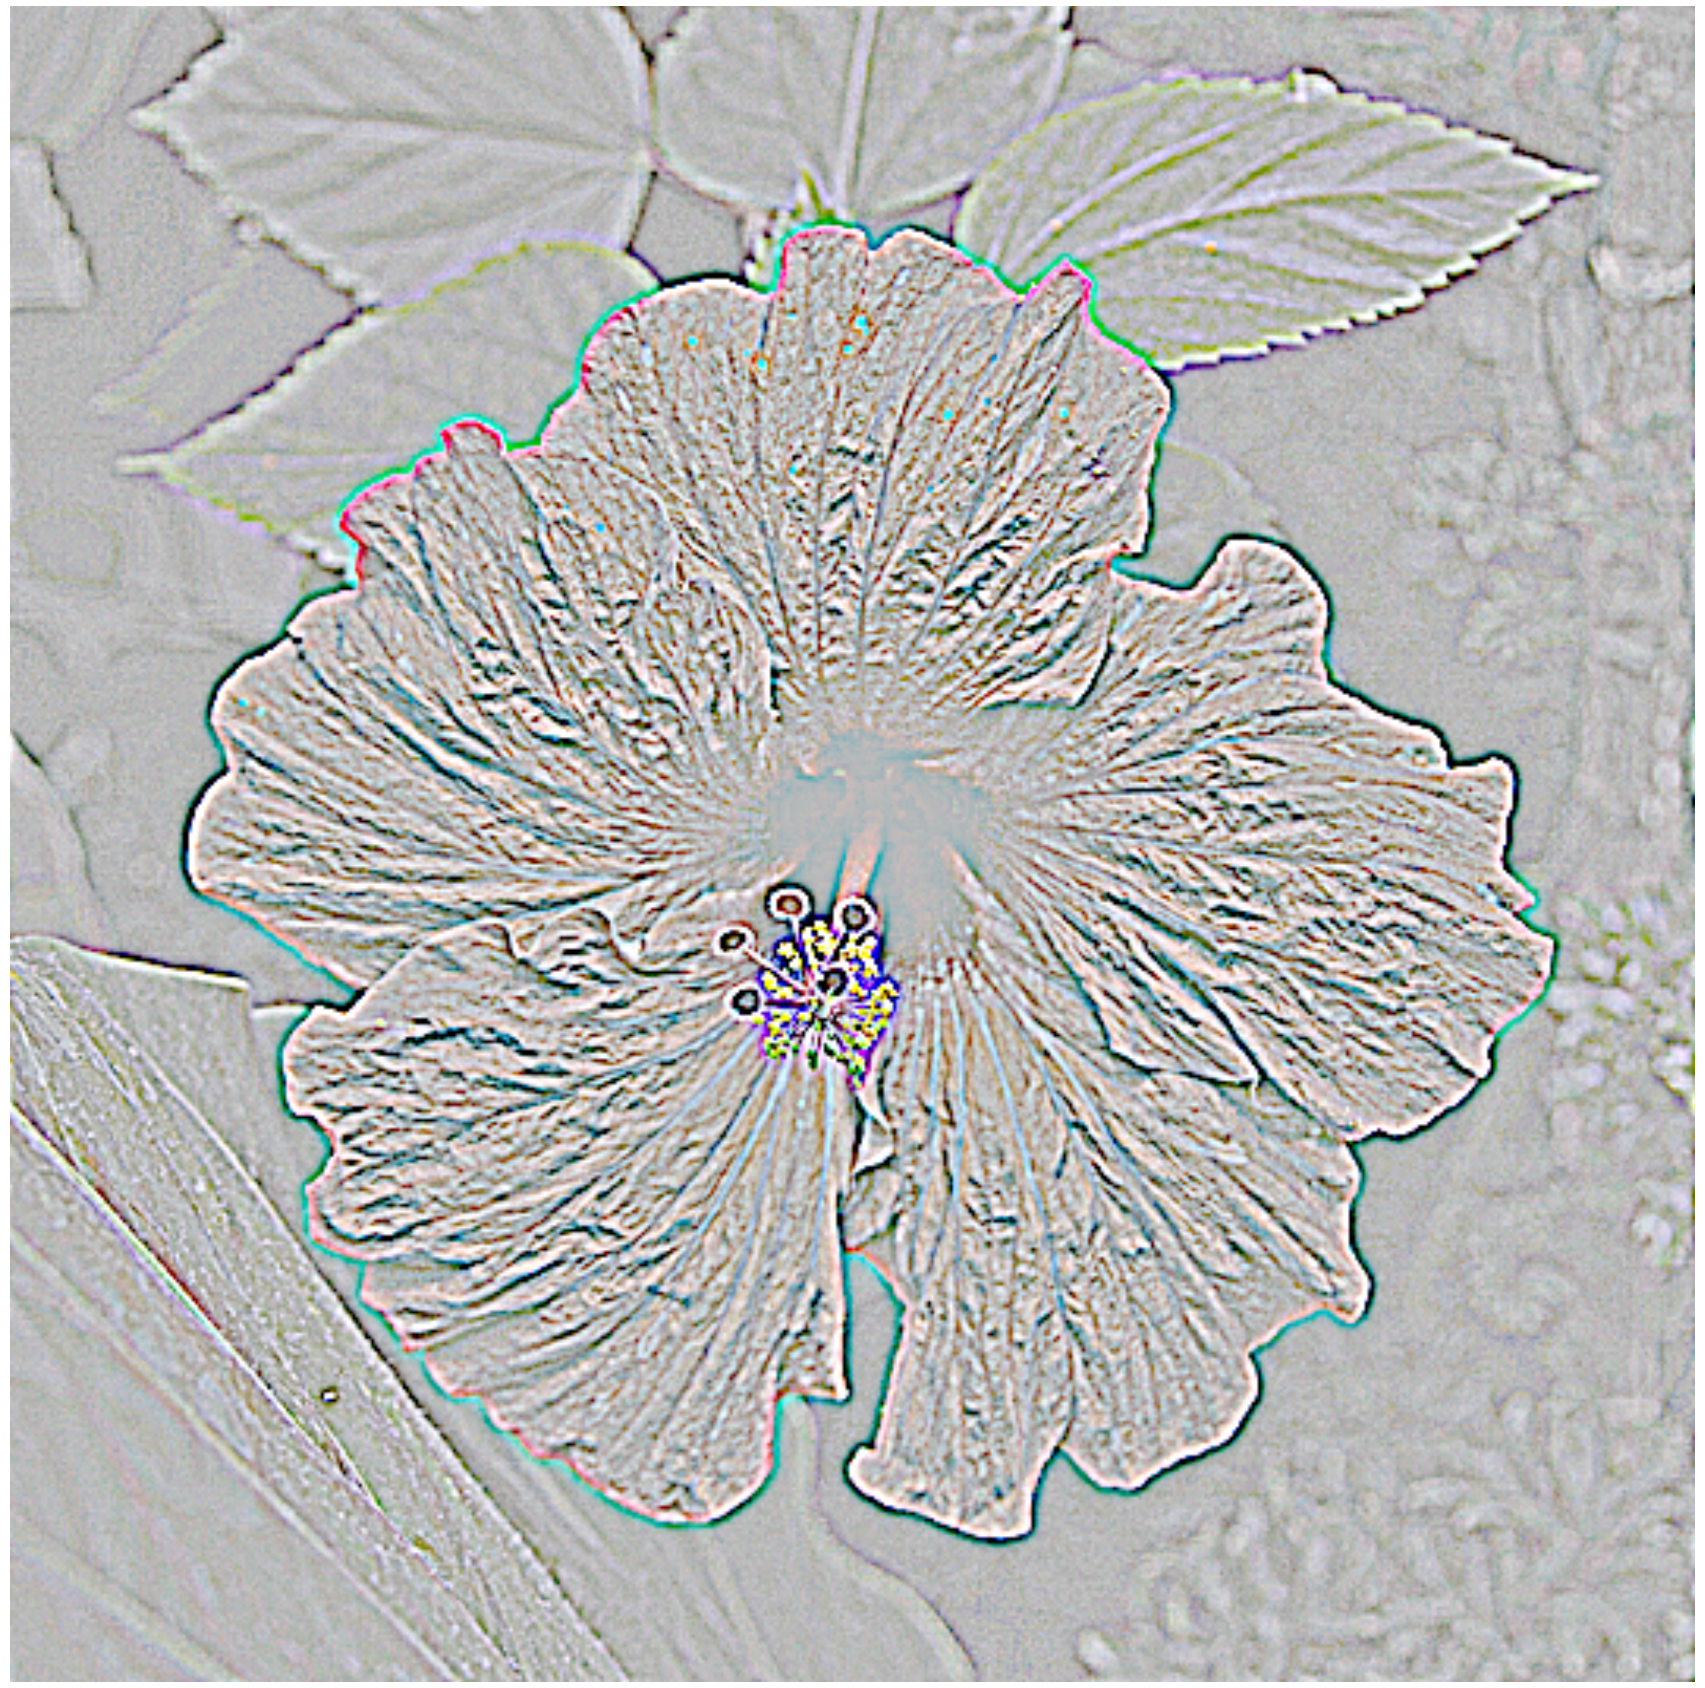

Which one takes more space to store?



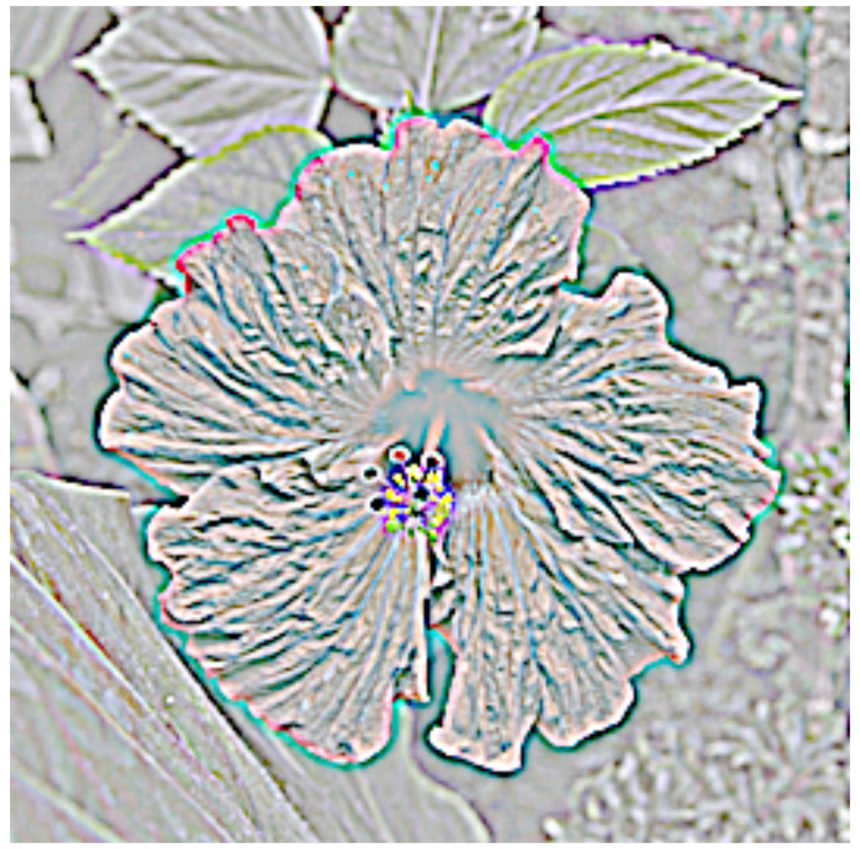

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Aside: Image Blending

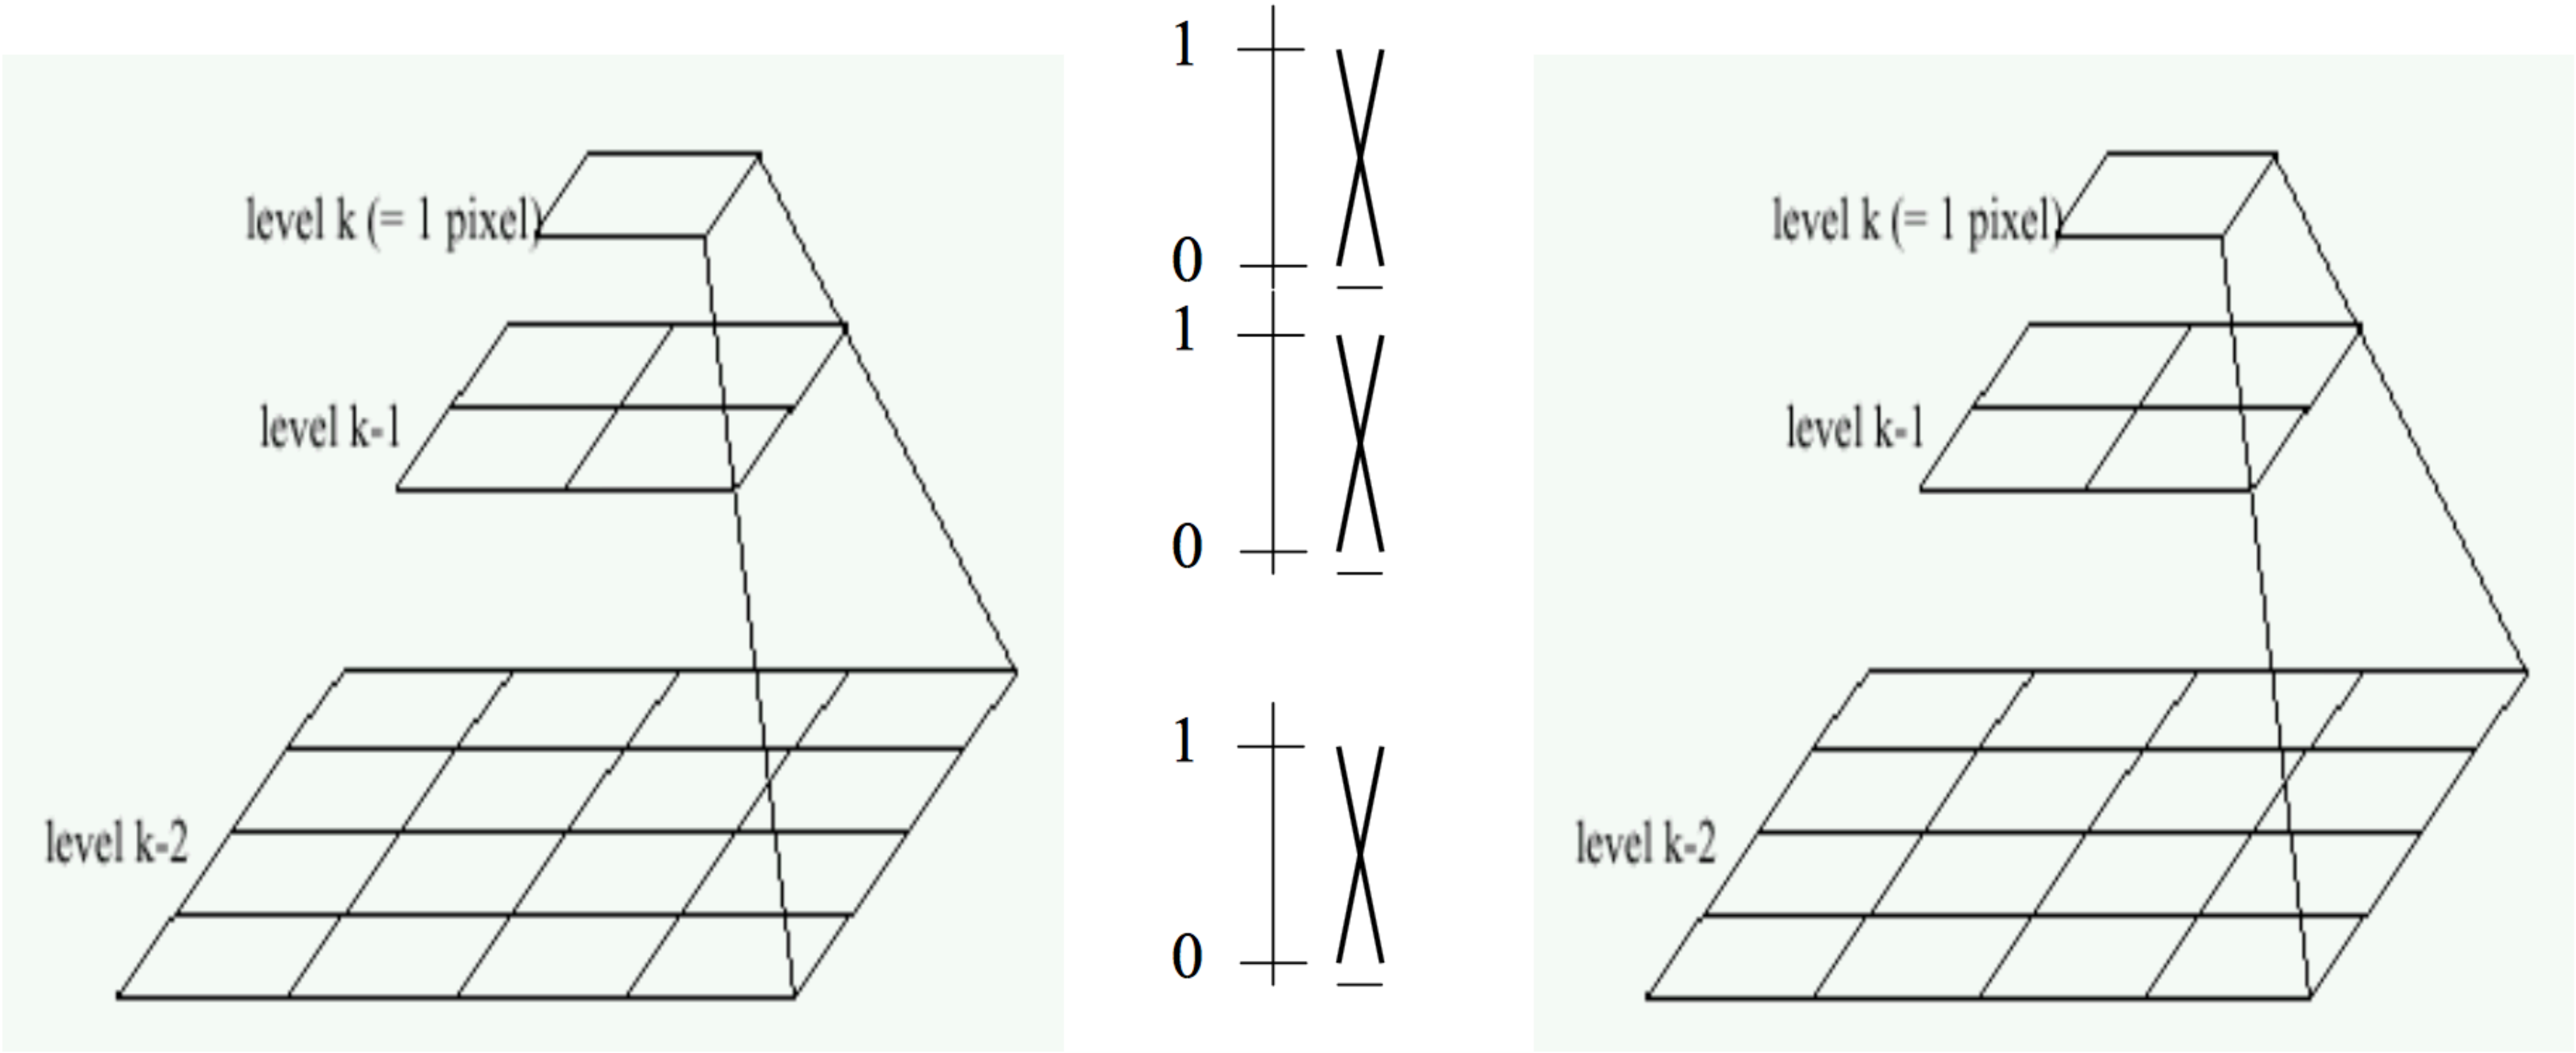

Left pyramid blend Right pyramid

Burt and Adelson, “A multiresolution spline with application to image mosaics,”ACM Transactions on Graphics, 1983, Vol.2, pp.217-236.

Aside: Image Blending

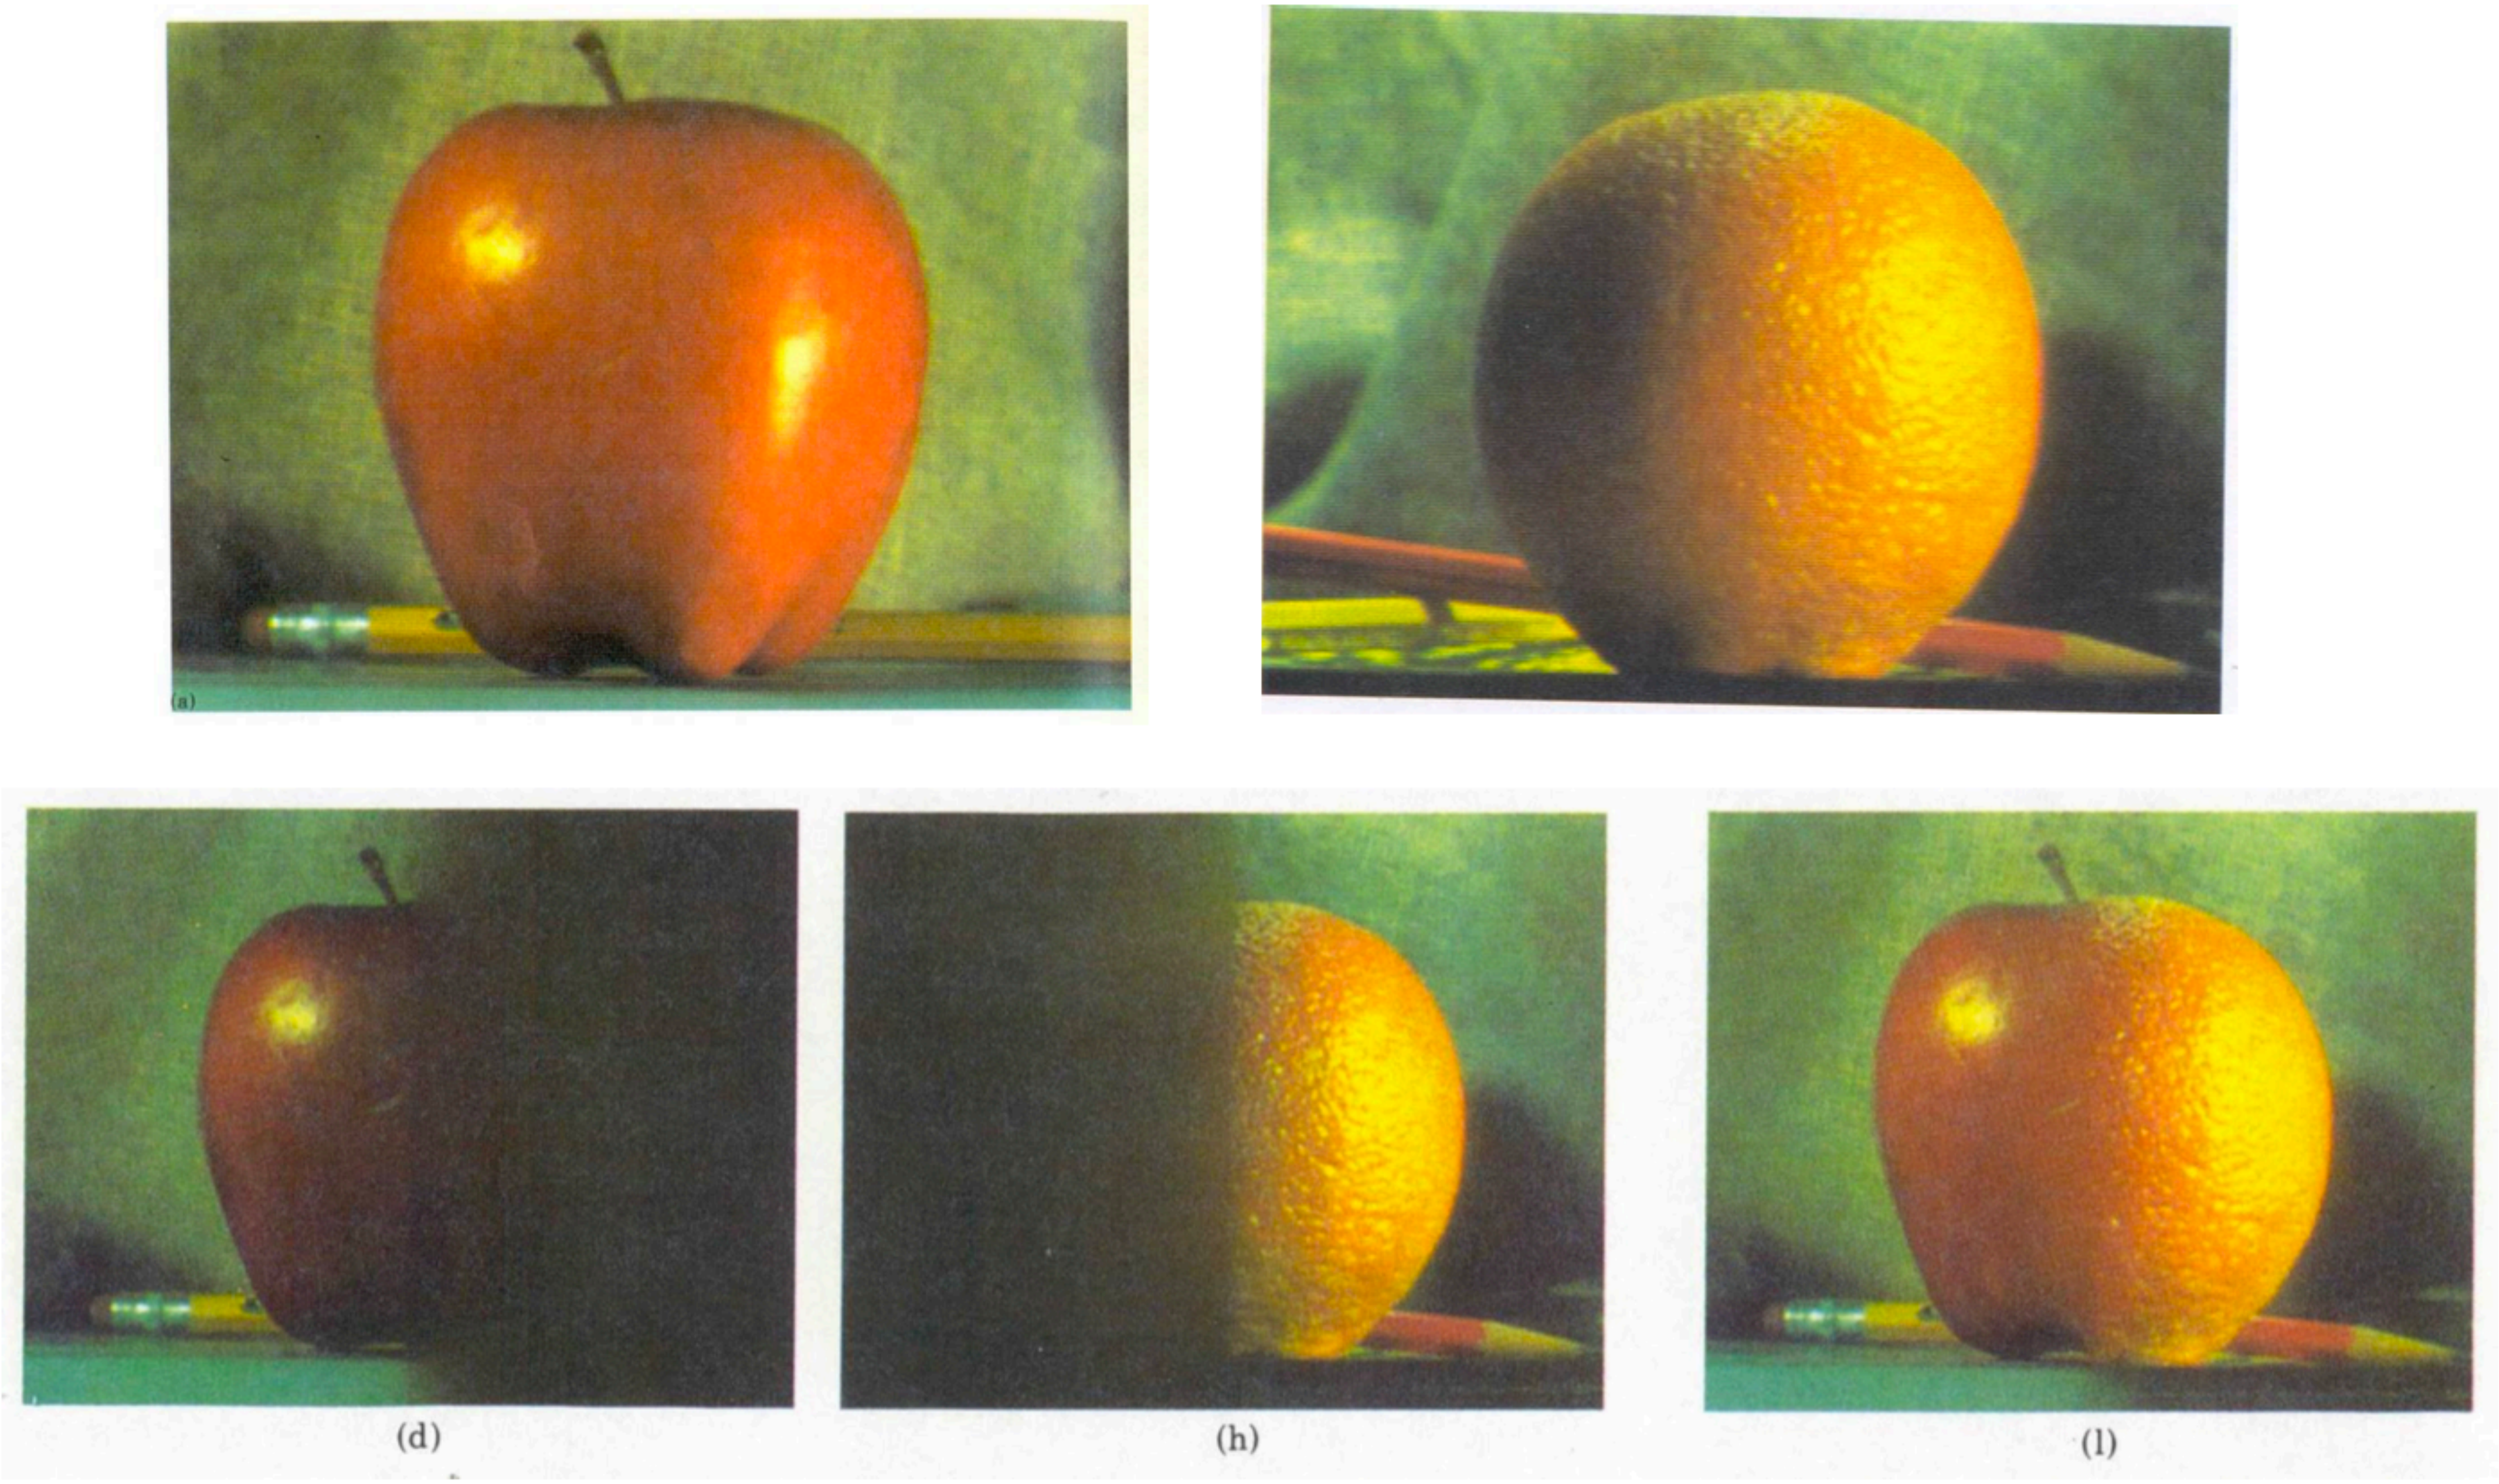

Burt and Adelson, “A multiresolution spline with application to image mosaics,”ACM Transactions on Graphics, 1983, Vol.2, pp.217-236.

Aside: Image Blending

Algorithm:

1. Build Laplacian pyramid LA and LB from images A and B

2. Build a Gaussian pyramid GR from mask image R (the mask defines which image pixels should be coming from A or B)

3. Forma a combined (blended) Laplacian pyramid LS, using nodes of GR as weights: $LS(i,j) = GR(i,j) * LA(i,j) + (1-GR(i,j)) * LB(i,j)$

4. Reconstruct the final blended image from LS

Aside: Image Blending

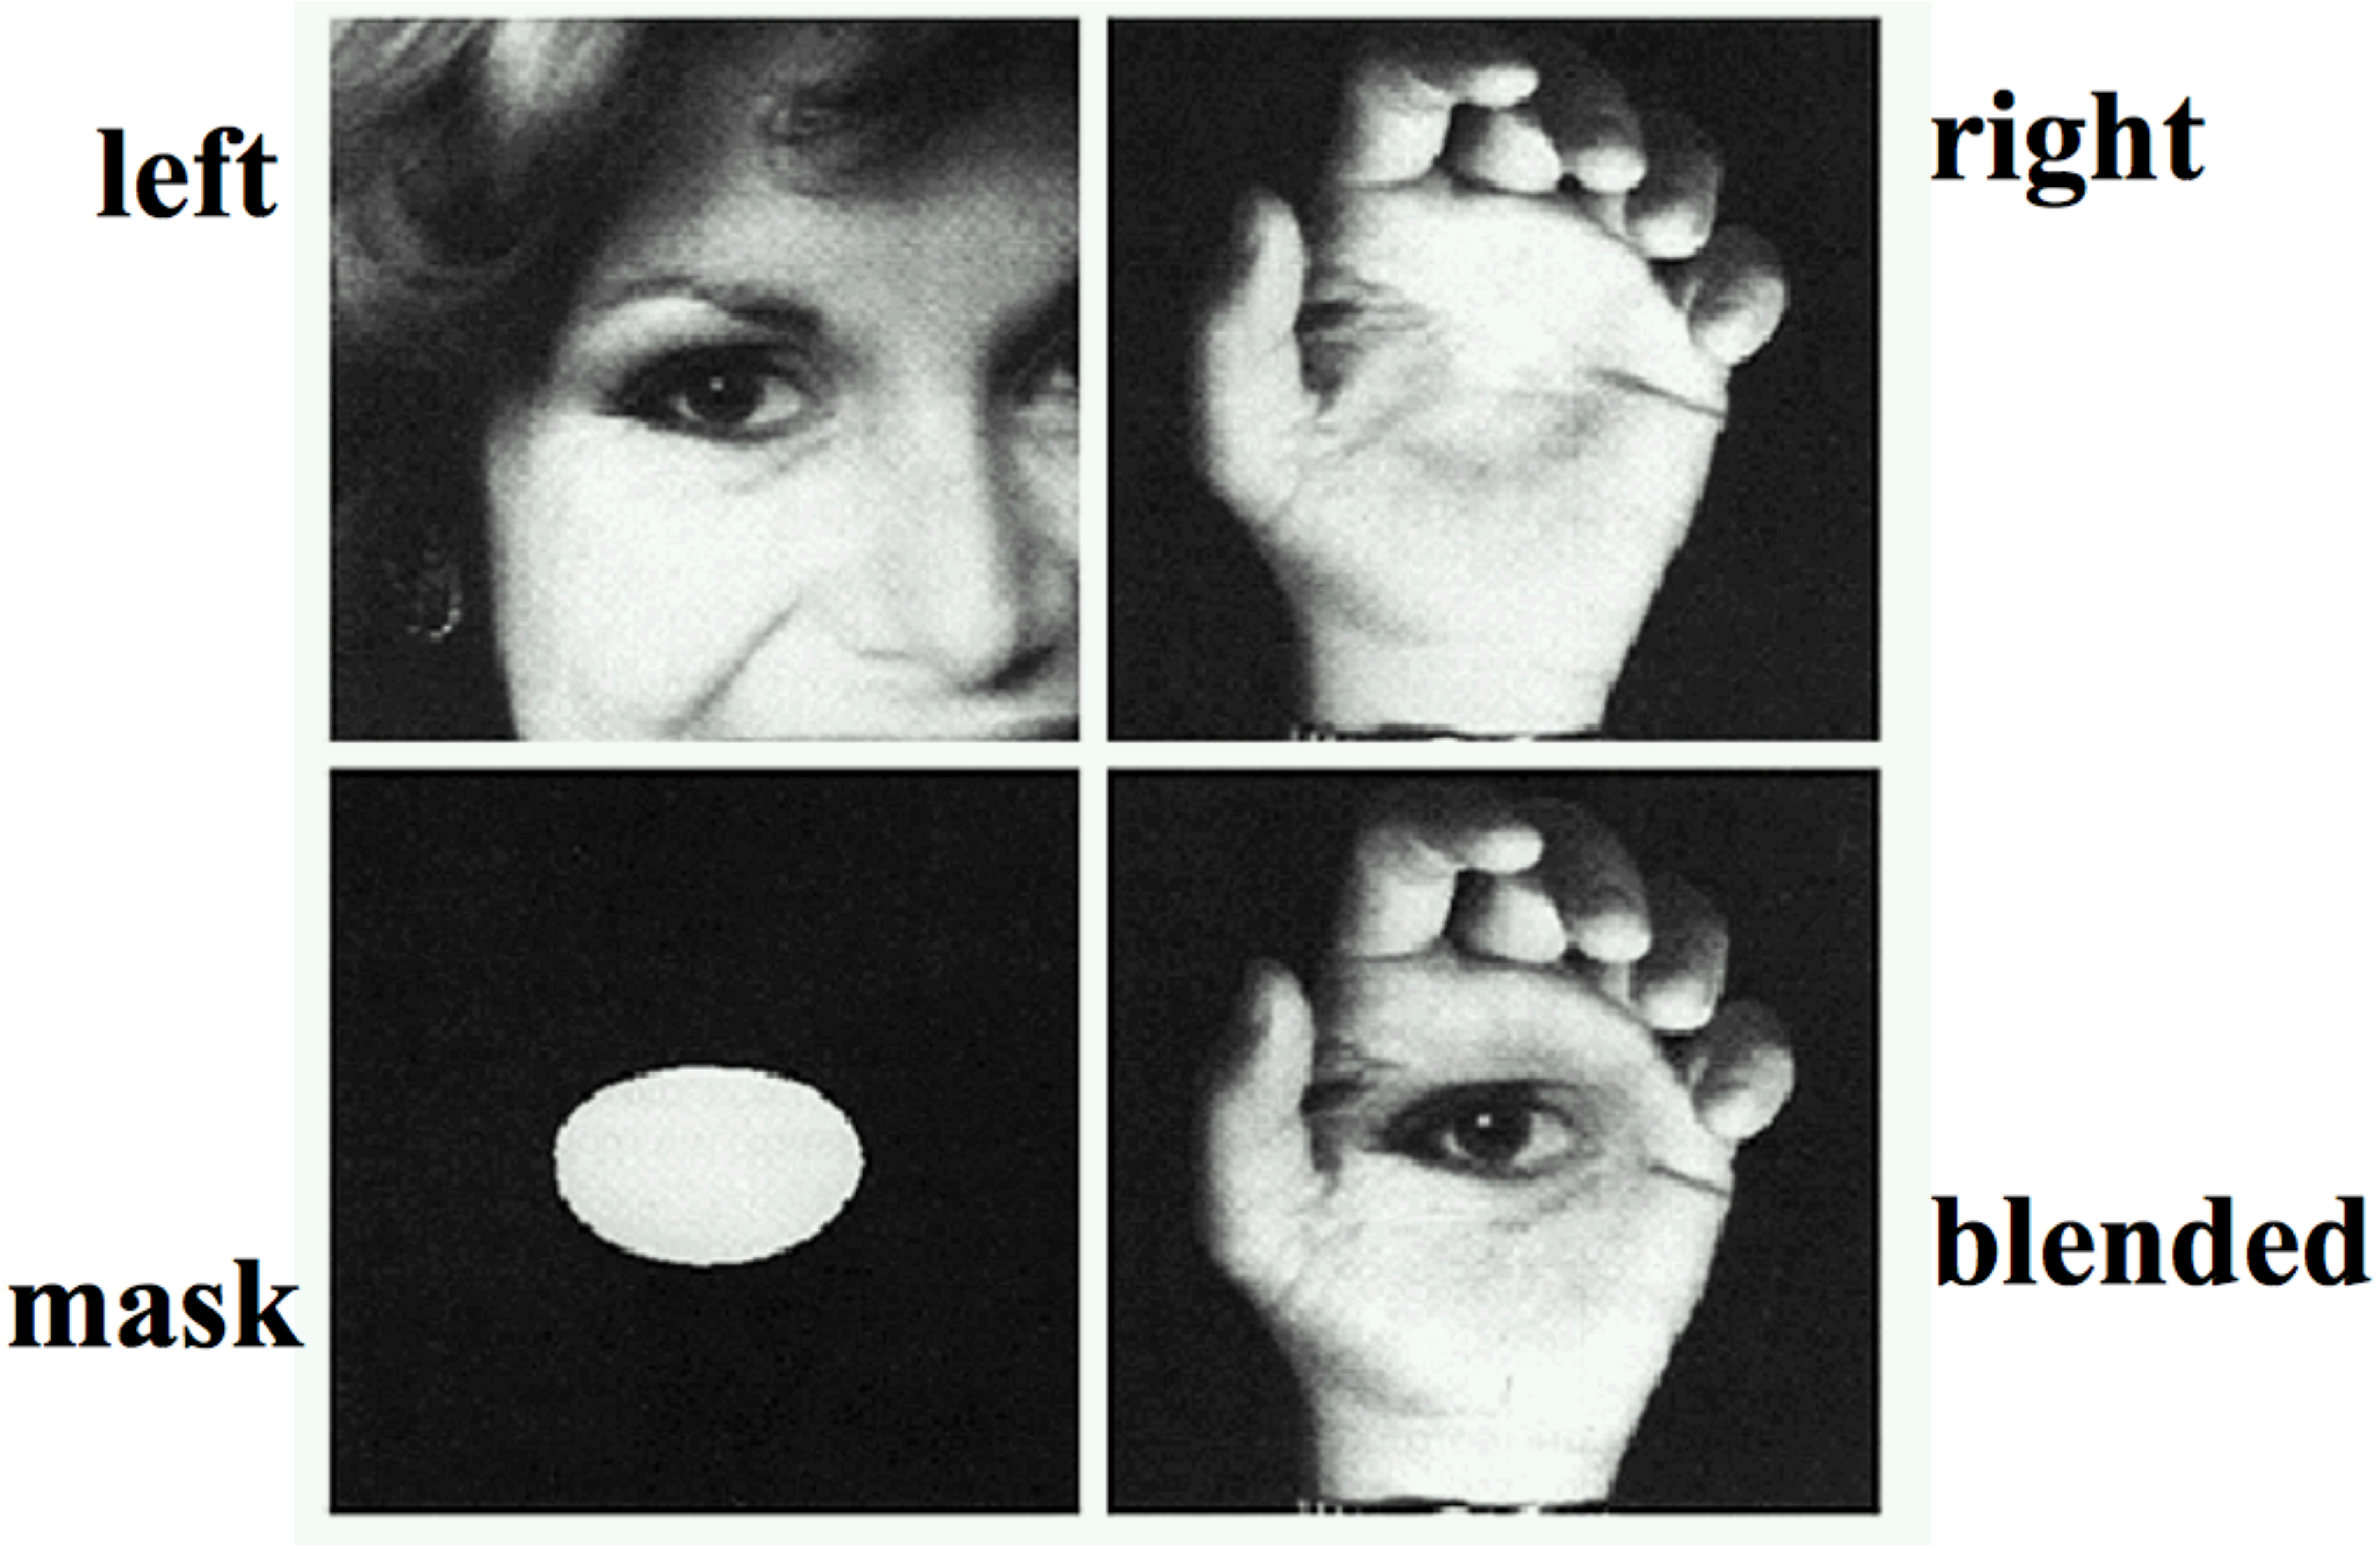

Aside: Image Blending

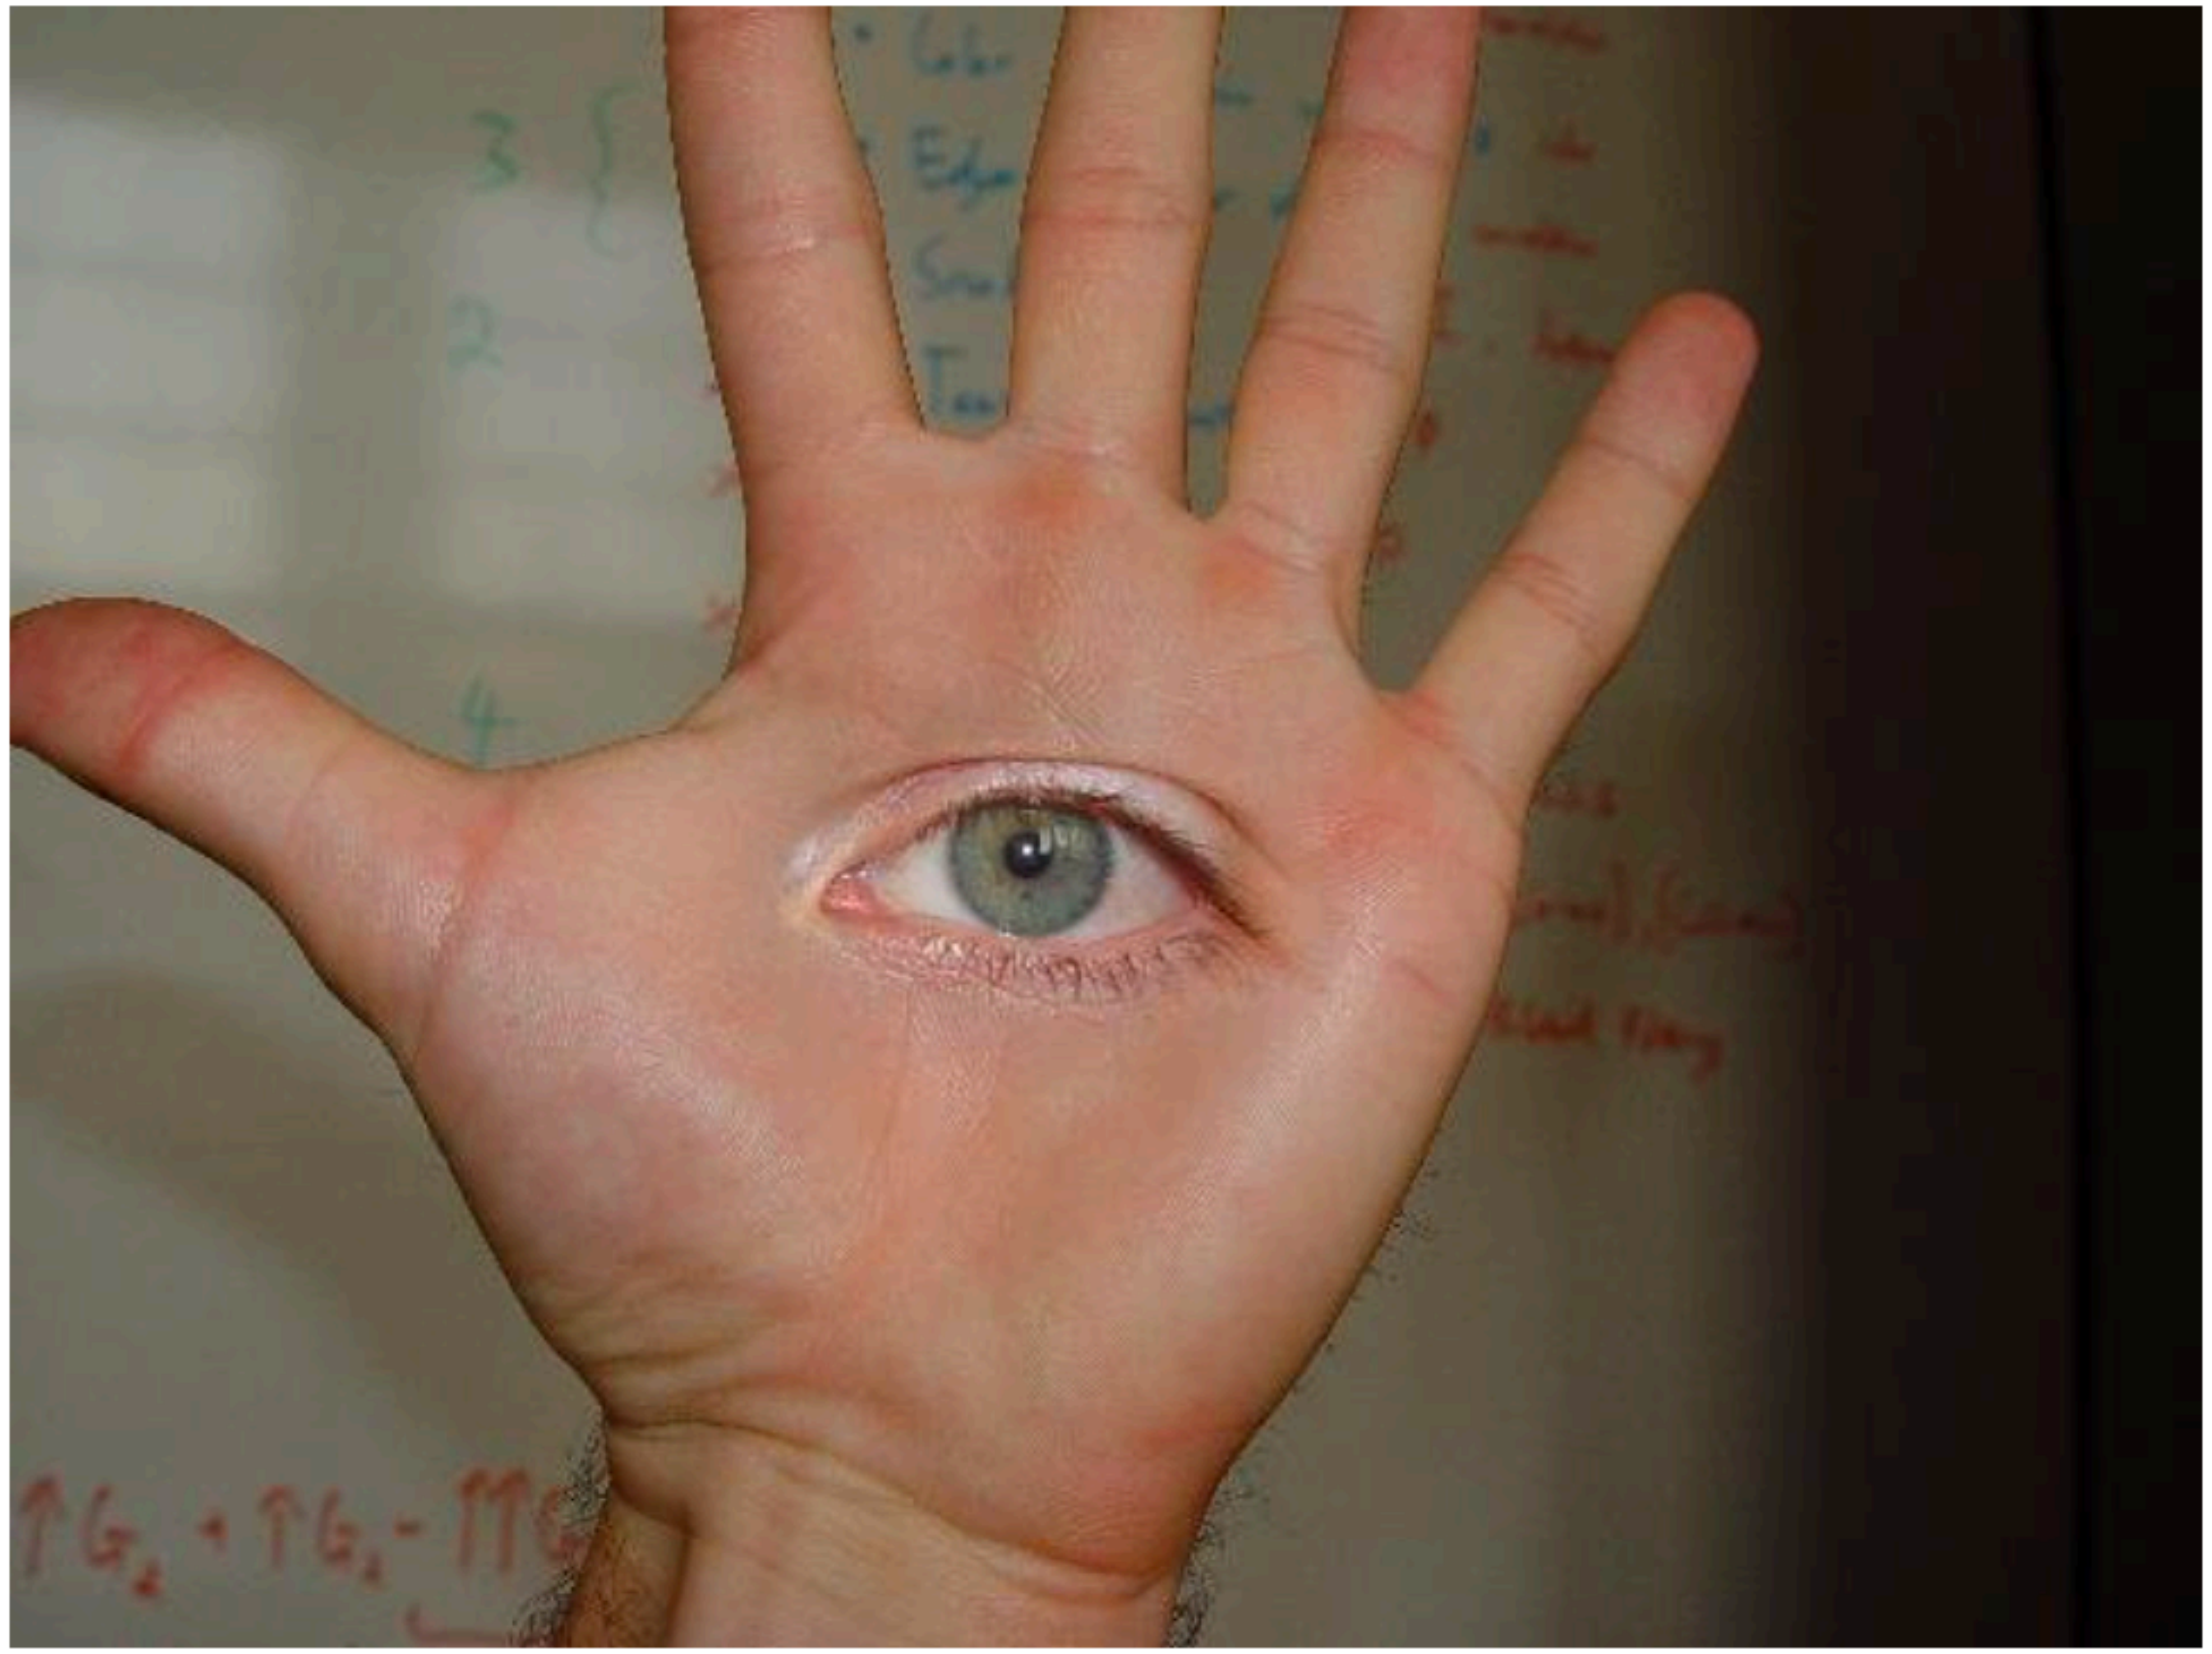

© david dmartin (Boston College)

Aside: Image Blending

© Chris Cameron

Oriented Pyramids

Laplacian pyramid is orientation independent

Idea: Apply an oriented filter at each layer

— represent image at a particular scale and orientation

— Aside: We do not study details in this course

Oriented Pyramids

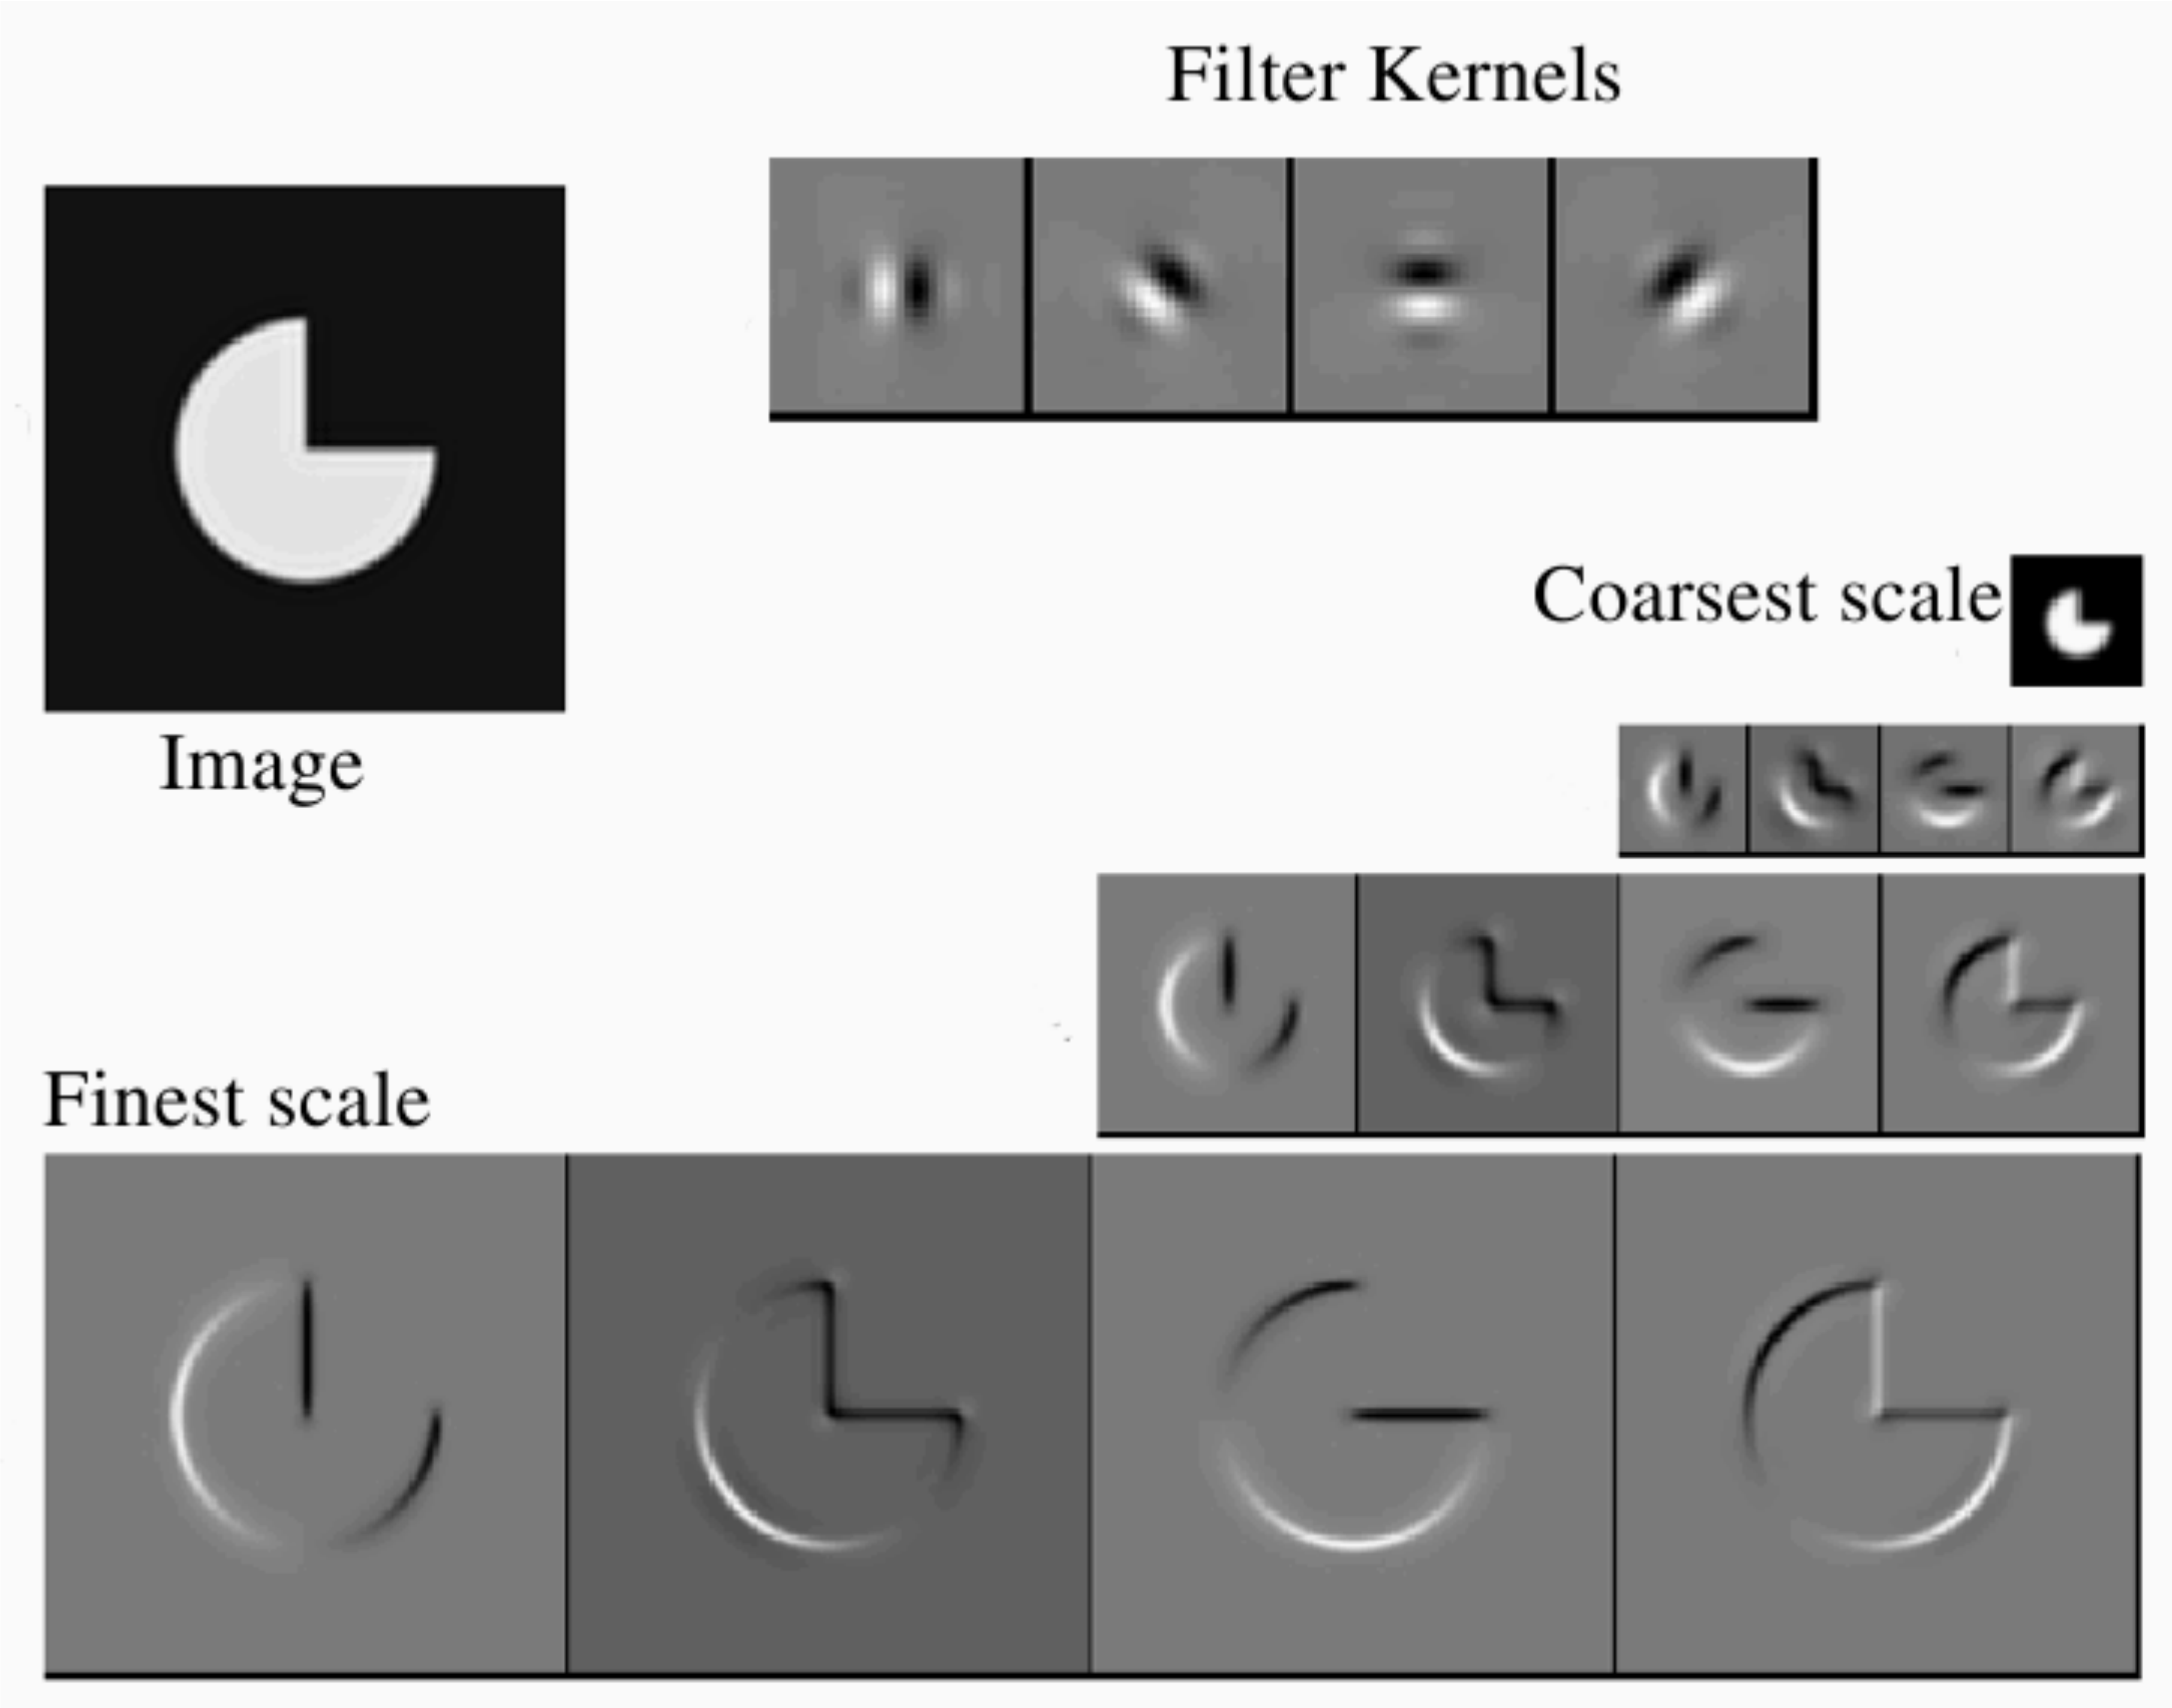

Forsyth & Ponce (1st ed.) Figure 9.13

Oriented Pyramids

Oriental Filters

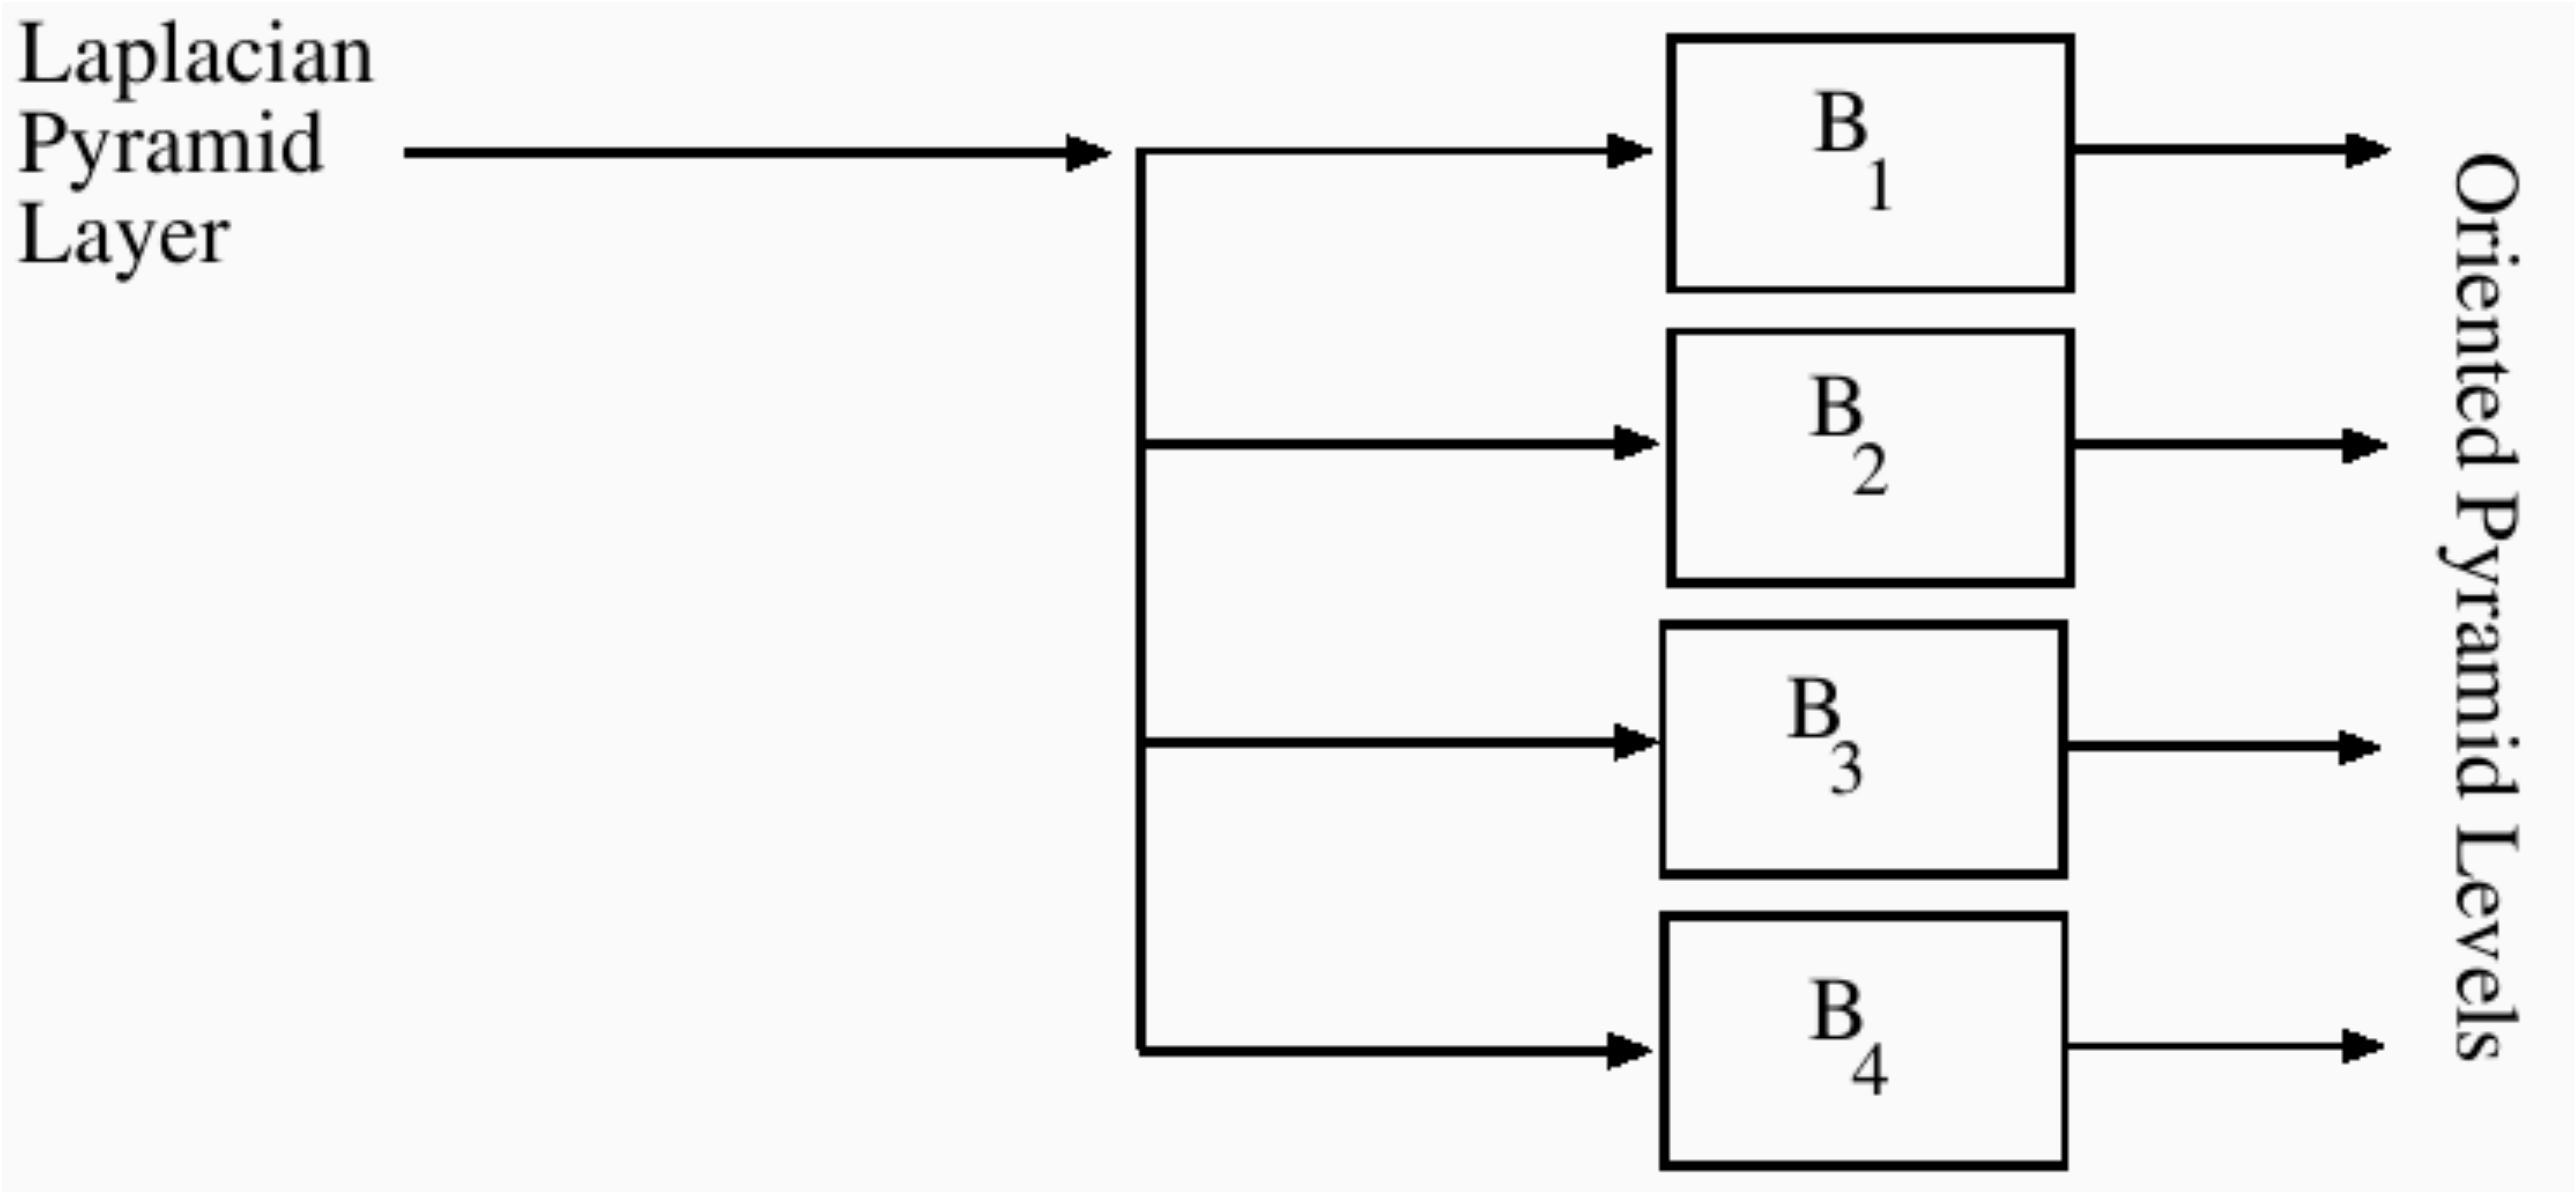

Forsyth & Ponce (1st ed.) Figure 9.14

Final Texture Representaation

Steps:

1. Form a Laplacian and oriented pyramid (or equivalent set of responses to filters at different scales and orientations)

2. Square the output (makes values positive)

3. Average responses over a neighborhood by blurring with a Gaussian

4. Take statistics of responses
— Mean of each filter output
— Possibly standard deviation of each filter

Please get your **iClickers** — Quiz