Texture

April 18th, 2019

Yong Jae Lee
UC Davis

Announcements

- PS1 out today
  - Due 5/1st, 11:59 pm
  - start early!

Review: last time

- Edge detection:
  - Filter for gradient
  - Threshold gradient magnitude, thin
- Chamfer matching to compare shapes (in terms of edge points)
- Binary image analysis
  - Thresholding
  - Morphological operators to “clean up”

Chamfer matching system

- Gavrila et al. http://gavrila.net/Research/Chamfer_System/chamfer_system.html

Slide credit: Kristen Grauman

Chamfer matching system

- Gavrila et al. http://gavrila.net/Research/Chamfer_System/chamfer_system.html

Slide credit: Kristen Grauman

ShadowDraw [Lee et al., SIGGRAPH 2011]

video

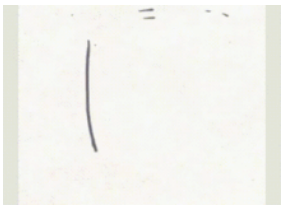

6

Today: Texture

What defines a texture?

7

Includes: more regular patterns

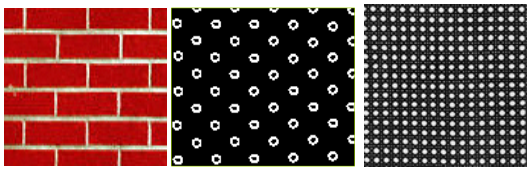

8

Slide credit: Kristen Grauman

Includes: more random patterns

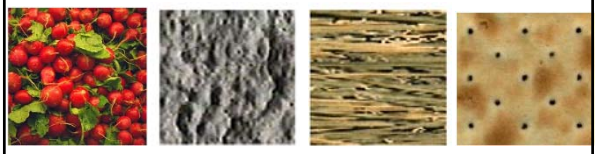

9

Slide credit: Kristen Grauman

Texture-related tasks

- **Shape from texture**
  - Estimate surface orientation or shape from image texture

10

Slide credit: Kristen Grauman

Shape from texture

- Use deformation of texture from point to point to estimate surface shape

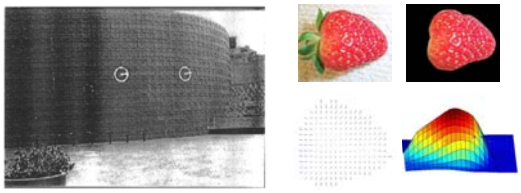

Pics from A. Loh: http://www.csse.uwa.edu.au/~angie/phdpics1.html

11

Slide credit: Kristen Grauman

Texture-related tasks

- **Shape from texture**
  - Estimate surface orientation or shape from image texture
- **Classification/segmentation** from texture cues
  - Analyze, represent texture
  - Group image regions with consistent texture
- **Synthesis**
  - Generate new texture patches/images given some examples

12

Slide credit: Kristen Grauman

Analysis vs. Synthesis

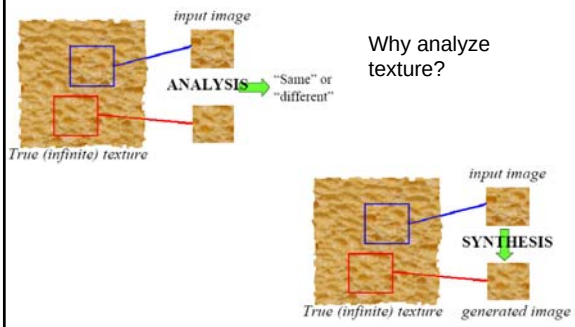

Why analyze texture?

13

Images:Bill Freeman, A. Efros

14

Image credit: D. Forsyth

Why analyze texture?

Importance to perception:

- Often indicative of a material's properties

15

Slide credit: Kristen Grauman

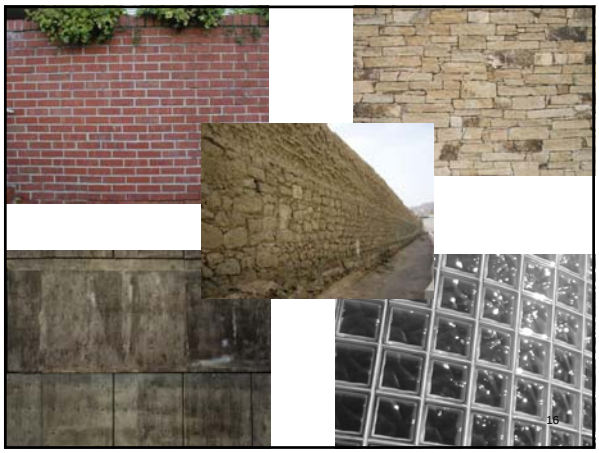

16

Why analyze texture?

Importance to perception:

- Often indicative of a material's properties
- Can be important appearance cue, especially if shape is similar across objects

17

Slide credit: Kristen Grauman

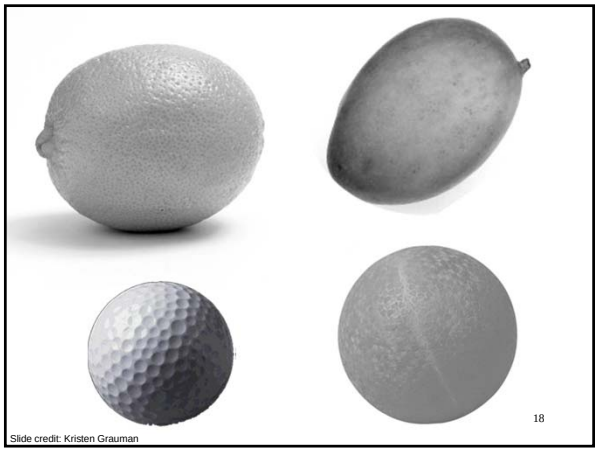

18

Slide credit: Kristen Grauman

Why analyze texture?

Importance to perception:

- Often indicative of a material's properties
- Can be important appearance cue, especially if shape is similar across objects
- Aim to distinguish between boundaries and texture

19

Slide credit: Kristen Grauman

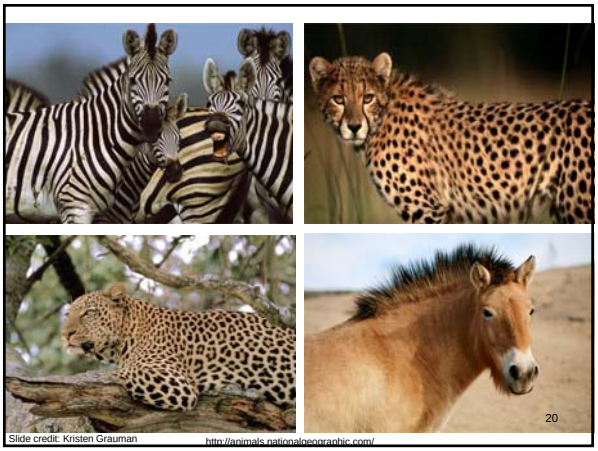

20

Slide credit: Kristen Grauman

http://animals.nationalgeographic.com/

Why analyze texture?

Importance to perception:

- Often indicative of a material's properties
- Can be important appearance cue, especially if shape is similar across objects
- Aim to distinguish between boundaries and texture

Technically:

- Representation-wise, we want a feature one step above "building blocks" of filters, edges.

21

Slide credit: Kristen Grauman

Images from Malik and Perona, 1990

22

Psychophysics of texture

- Some textures distinguishable with *preattentive* perception – without scrutiny, eye movements [Julesz 1975]

Same or different?

23

Slide credit: Kristen Grauman

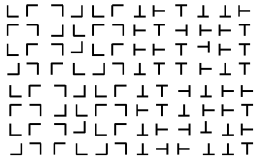

24

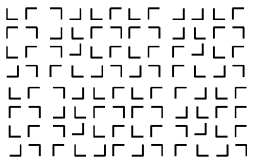

25

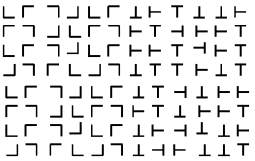

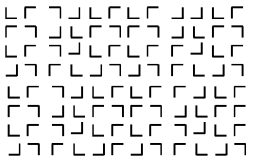

26

Capturing the local patterns with image measurements

[Bergen & Adelson, *Nature* 1988]

Scale of patterns influences discriminability

Size-tuned linear filters

27

Texture representation

- Textures are made up of repeated local patterns, so:
  - Find the patterns
    - Use filters that look like patterns (spots, bars, raw patches…)
    - Consider magnitude of response
  - Describe their statistics within each local window
    - Mean, standard deviation
    - Histogram of "prototypical" feature occurrences

28

Slide credit: Kristen Grauman

Texture representation: example

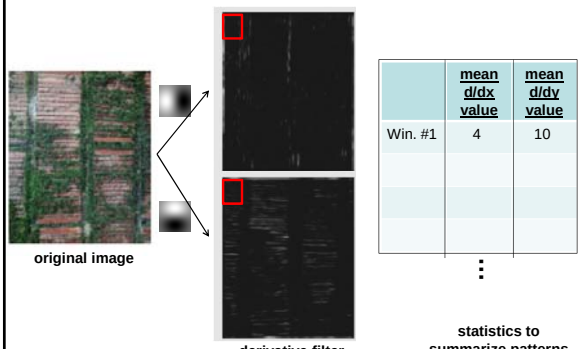

original image

derivative filter responses, squared

	mean d/dx value	mean d/dy value
Win. #1	4	10

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

29

Texture representation: example

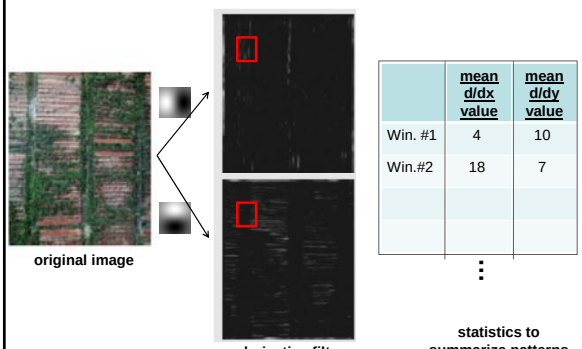

original image

derivative filter responses, squared

	mean d/dx value	mean d/dy value
Win. #1	4	10
Win.#2	18	7

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

30

Texture representation: example

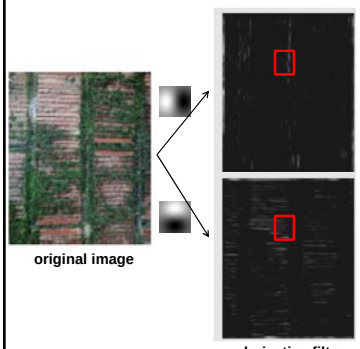

original image

derivative filter responses, squared

	mean d/dx value	mean d/dy value
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20

⋮

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

31

Texture representation: example

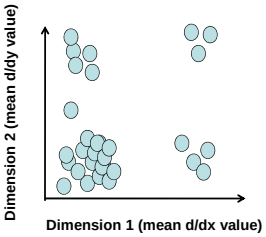

	mean d/dx value	mean d/dy value
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20

⋮

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

32

Texture representation: example

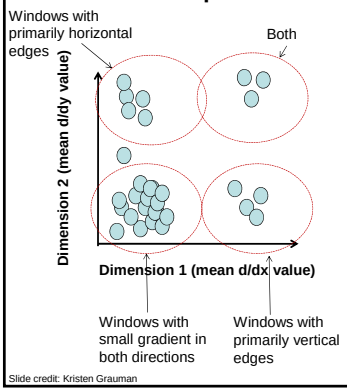

	mean d/dx value	mean d/dy value
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20

⋮

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

33

Texture representation: example

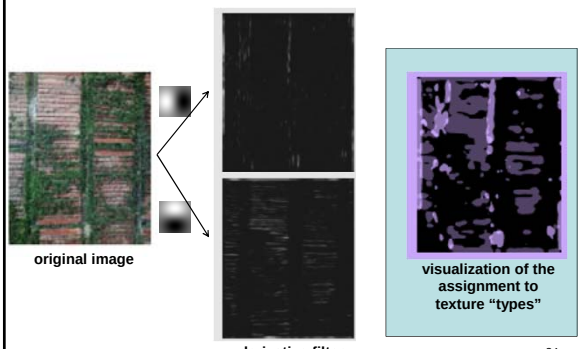

Slide credit: Kristen Grauman

Texture representation: example

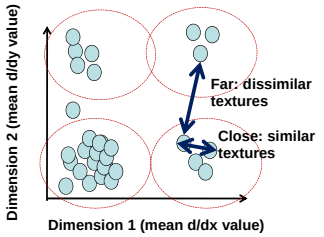

	mean d/dx value	mean d/dy value
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20
⋮		

statistics to summarize patterns in small windows

Slide credit: Kristen Grauman

Texture representation: example

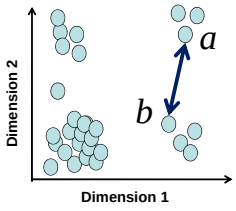

$$D(a,b) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2}$$

$$D(a,b) = \sqrt{\sum_{i=1}^{2} (a_i - b_i)^2}$$

Slide credit: Kristen Grauman

Texture representation: example

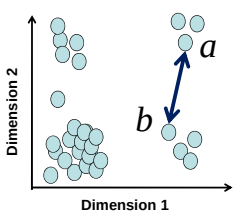

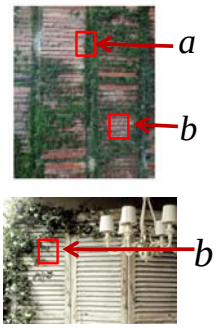

Distance reveals how dissimilar texture from window a is from texture in window b.

37

Slide credit: Kristen Grauman

Texture representation: window scale

- The window size (i.e., scale) for which we collect these statistics is important.

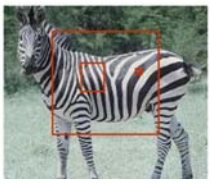

Possible to perform scale selection by looking for window scale where texture description not changing.

38

Filter banks

- Our previous example used two filters, and resulted in a 2-dimensional feature vector to describe texture in a window
  - x and y derivatives revealed something about local structure
- We can generalize to apply a collection of multiple (*d*) filters: a "filter bank"
- Then our feature vectors will be *d*-dimensional
  - still can think of nearness, farness in feature space

39

Slide credit: Kristen Grauman

Filter banks

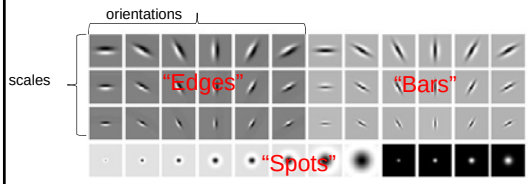

- What filters to put in the bank?
  - Typically we want a combination of scales and orientations, different types of patterns.

Matlab code available for these examples:
http://www.robots.ox.ac.uk/~vgg/research/texclass/filters.html

40

Slide credit: Kristen Grauman

Multivariate Gaussian

$$p(x;\mu,\Sigma) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu)\right)$$

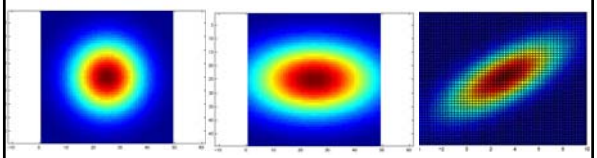

$$\Sigma = \begin{bmatrix} 9 & 0 \\ 0 & 9 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 16 & 0 \\ 0 & 9 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 10 & 5 \\ 5 & 5 \end{bmatrix}$$

41

Slide credit: Kristen Grauman

Filter bank

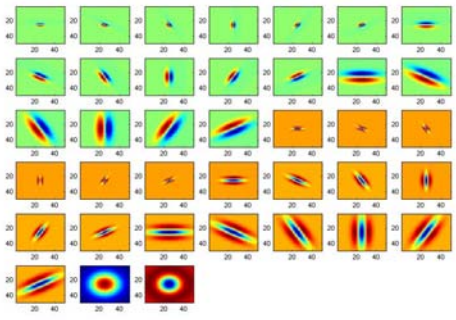

42

Slide credit: Kristen Grauman

Image from http://www.texasexplorer.com/austincap2.jpg

Slide credit: Kristen Grauman

43

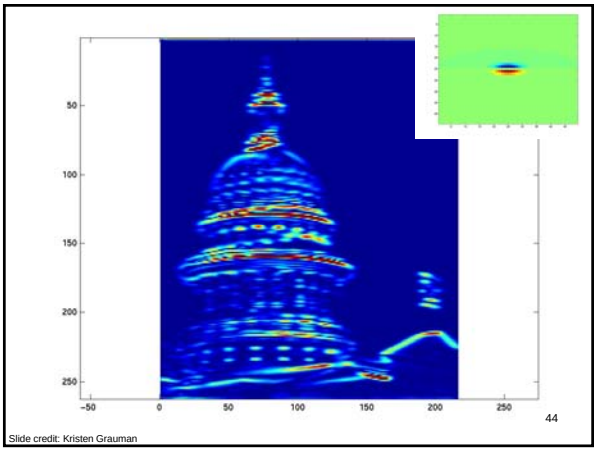

44

Slide credit: Kristen Grauman

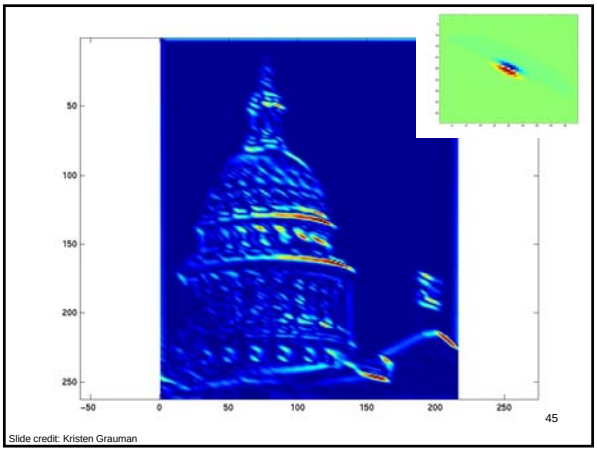

45

Slide credit: Kristen Grauman

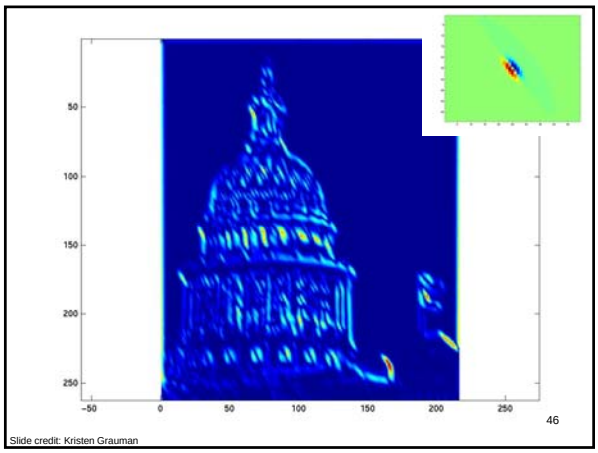

46

Slide credit: Kristen Grauman

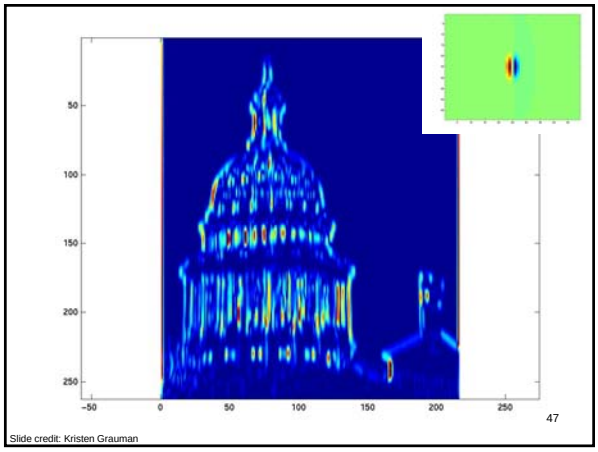

47

Slide credit: Kristen Grauman

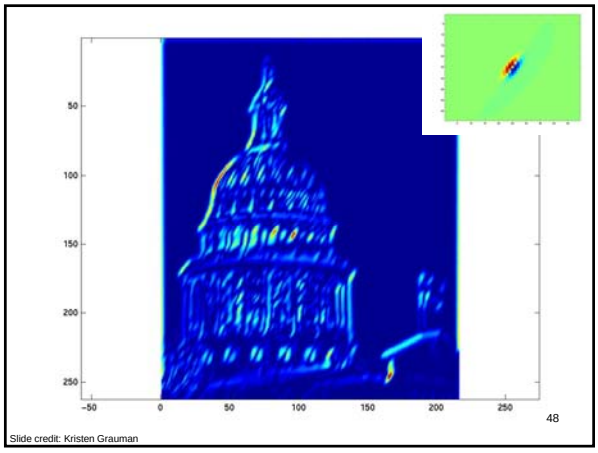

48

Slide credit: Kristen Grauman

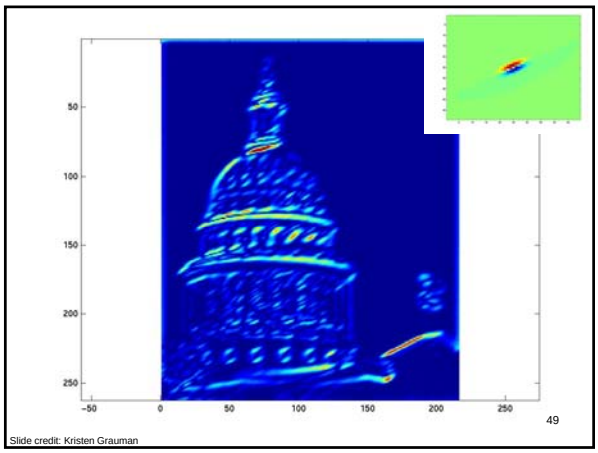

49

Slide credit: Kristen Grauman

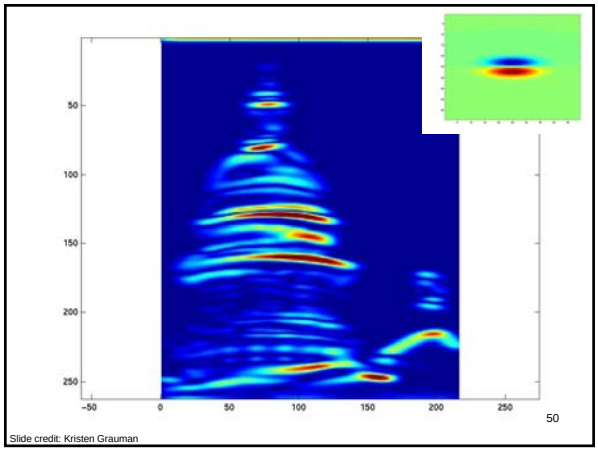

50

Slide credit: Kristen Grauman

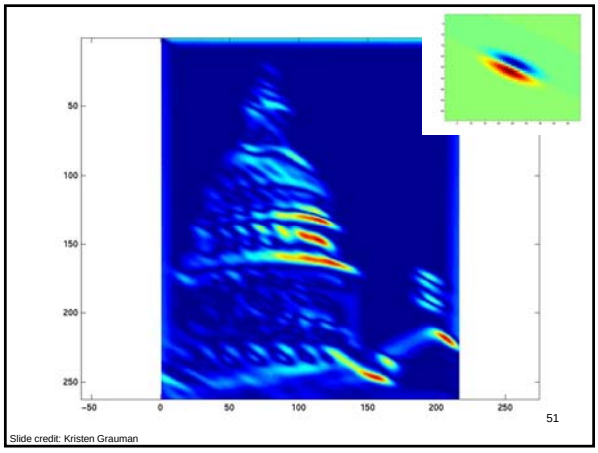

51

Slide credit: Kristen Grauman

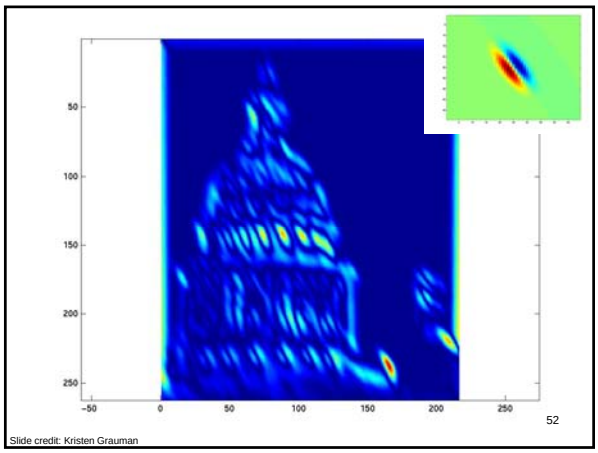

52

Slide credit: Kristen Grauman

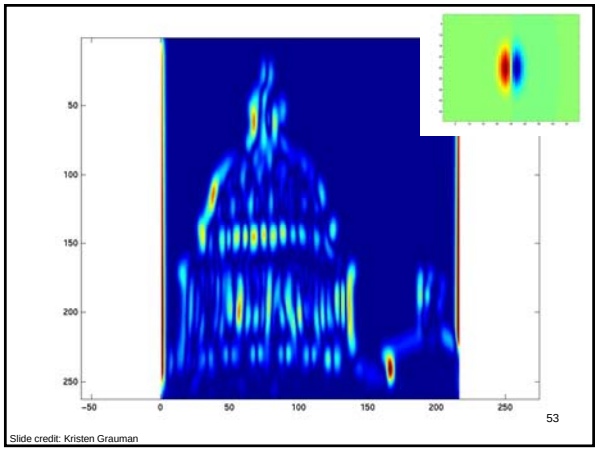

53

Slide credit: Kristen Grauman

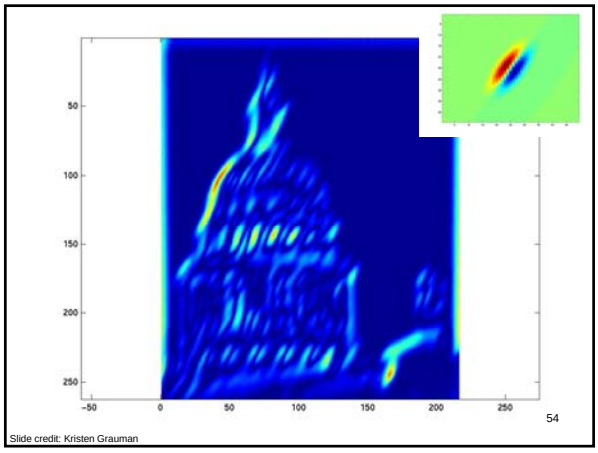

54

Slide credit: Kristen Grauman

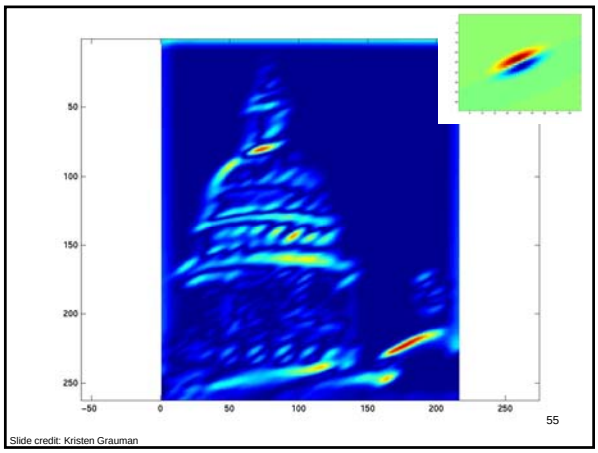

55

Slide credit: Kristen Grauman

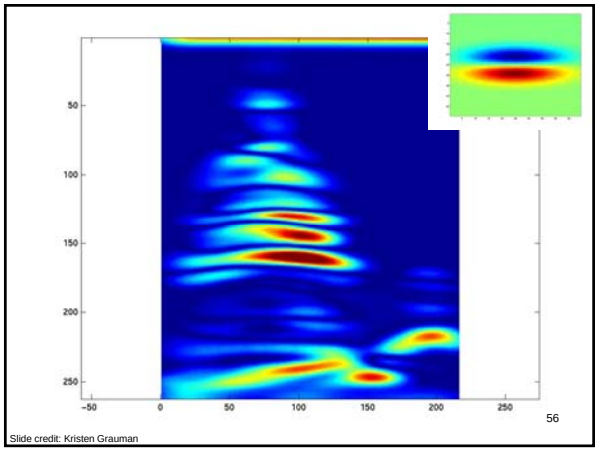

56

Slide credit: Kristen Grauman

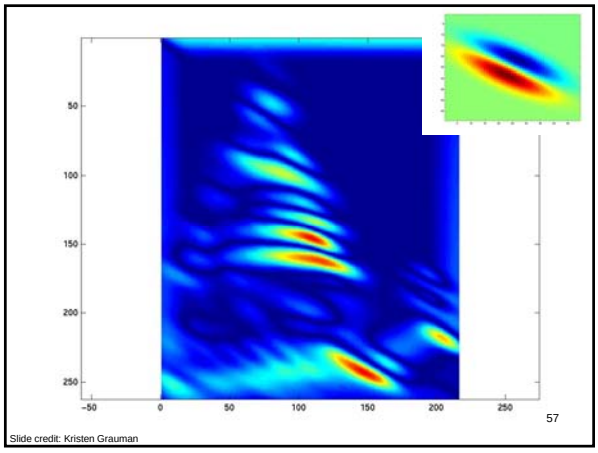

57

Slide credit: Kristen Grauman

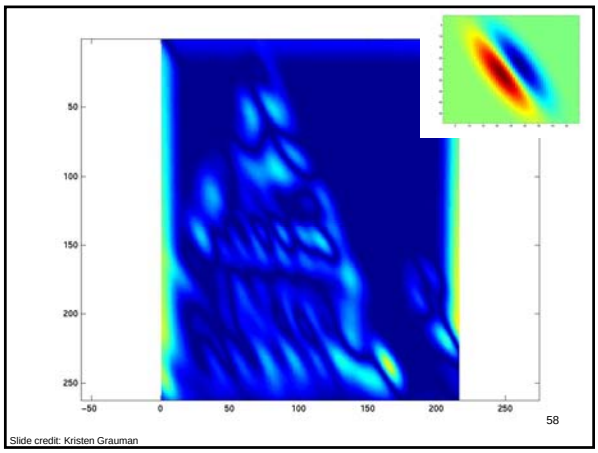

58

Slide credit: Kristen Grauman

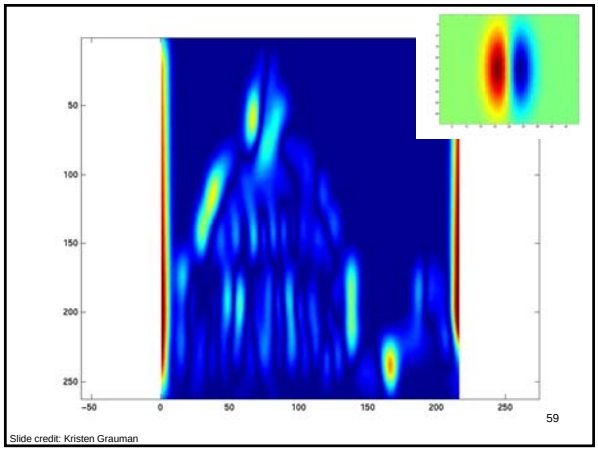

59

Slide credit: Kristen Grauman

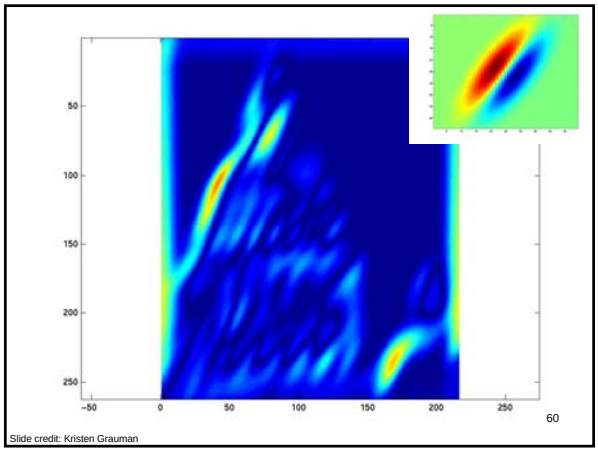

60

Slide credit: Kristen Grauman

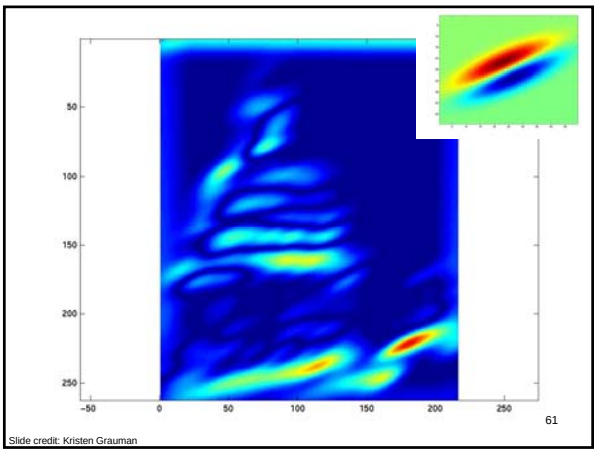

61

Slide credit: Kristen Grauman

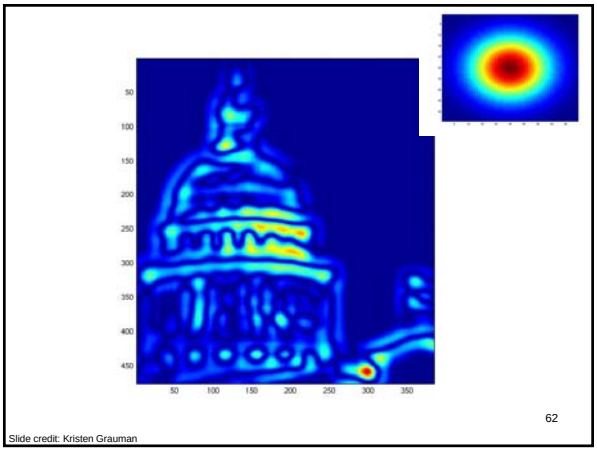

62

Slide credit: Kristen Grauman

You try: Can you match the texture to the response?

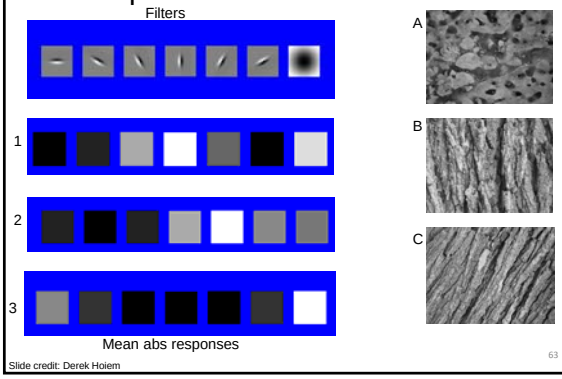

63

Slide credit: Derek Hoiem

Representing texture by mean abs response

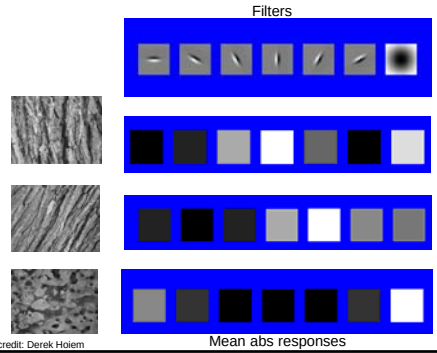

Slide credit: Derek Hoiem

64

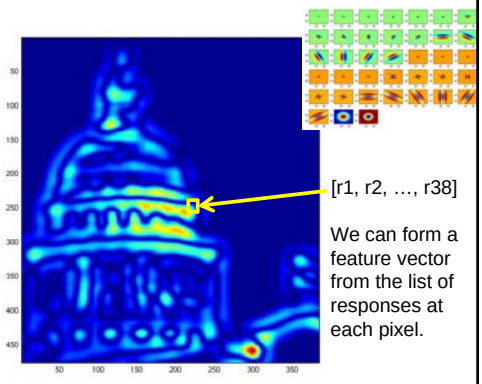

[r1, r2, …, r38]

We can form a feature vector from the list of responses at each pixel.

Slide credit: Kristen Grauman

65

d-dimensional features

$$D(a,b) = \sqrt{\sum_{i=1}^{d} (a_i - b_i)^2}$$

Euclidean distance (L_2)

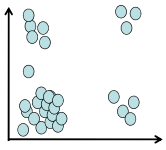

Slide credit: Kristen Grauman

66

Example uses of texture in vision: analysis

67

Classifying materials, “stuff”

Figure by Varma & Zisserman

68

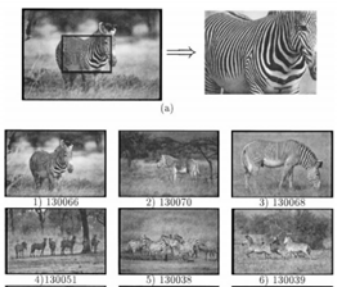

Texture features for image retrieval

Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99-121, November 2000,

69

Slide credit: Kristen Grauman

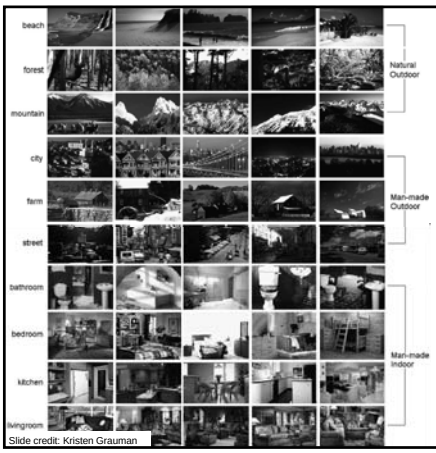

Characterizing scene categories by texture

L. W. Renninger and J. Malik. When is scene identification just texture recognition? Vision Research 44 (2004) 2301–2311

Slide credit: Kristen Grauman

70

Texture-related tasks

- **Shape from texture**
  - Estimate surface orientation or shape from image texture
- **Segmentation/classification** from texture cues
  - Analyze, represent texture
  - Group image regions with consistent texture
- **Synthesis**
  - Generate new texture patches/images given some examples

71

Slide credit: Kristen Grauman

Texture synthesis

- Goal: create new samples of a given texture
- Many applications: virtual environments, hole-filling, texturing surfaces

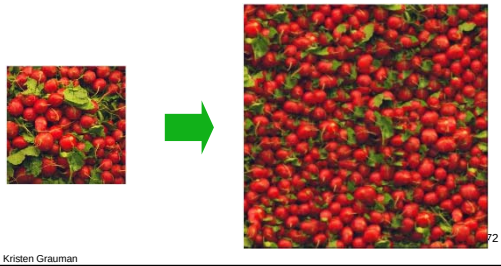

72

Slide credit: Kristen Grauman

The Challenge

- Need to model the whole spectrum: from repeated to stochastic texture

Alexei A. Efros and Thomas K. Leung, "Texture Synthesis by Non-parametric Sampling," Proc. International Conference on Computer Vision (ICCV), 1999.

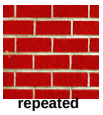

repeated

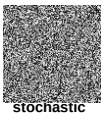

stochastic

Both?

73

Markov Chains

Markov Chain

- a *sequence* of random variables $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$
- $\mathbf{x}_t$ is the **state** of the model at time t

$$\boxed{\mathbf{x}_1} \rightarrow \boxed{\mathbf{x}_2} \rightarrow \boxed{\mathbf{x}_3} \rightarrow \boxed{\mathbf{x}_4} \rightarrow \boxed{\mathbf{x}_5}$$

74

Markov Chain Example: Text

"A dog is a man's best friend. It's a dog eat dog world out there."

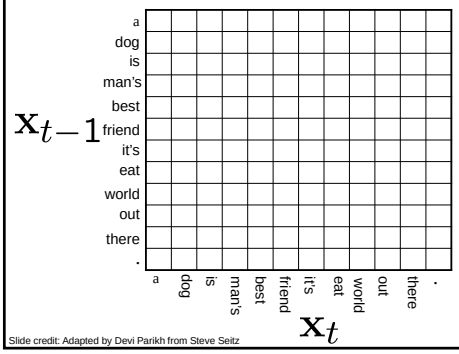

$p(\mathbf{x}_t|\mathbf{x}_{t-1})$

p(dog|a) = ?

75

Slide credit: Adapted by Devi Parikh from Steve Seitz

Markov Chain Example: Text

"A dog is a man's best friend. It's a dog eat dog world out there."

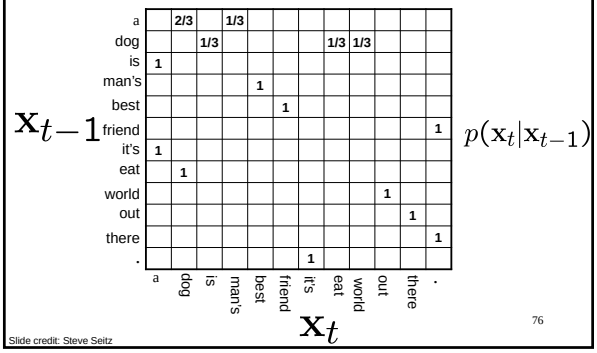

76

Slide credit: Steve Seitz

Text synthesis

Create plausible looking poetry, love letters, term papers, etc.

Most basic algorithm

1. Build probability histogram/table
   - find all blocks of N consecutive words/letters in training documents
   - compute probability of occurrence $p(\mathbf{x}_t|\mathbf{x}_{t-1},\ldots,\mathbf{x}_{t-(n-1)})$

WE NEED TO EAT CAKE

77

Slide credit: Steve Seitz

Text synthesis

- Results:
  - *"As I've commented before, really relating to someone involves standing next to impossible."*
  - *"One morning I shot an elephant in my arms and kissed him."*
  - *"I spent an interesting evening recently with a grain of salt"*

Dewdney, "A potpourri of programmed prose and prosody" *Scientific American,* 1989.

78

Slide from Alyosha Efros, ICCV 1999

Markov Random Field

A Markov random field (MRF)

- generalization of Markov chains to two or more dimensions.

First-order MRF:

- probability that pixel *X* takes a certain value given the values of neighbors *A*, *B*, *C*, and *D*:

$$P(\mathrm{X}|\mathrm{A},\mathrm{B},\mathrm{C},\mathrm{D})$$

79

Slide credit: Steve Seitz

Texture Synthesis [Efros & Leung, ICCV 99]

Can apply 2D version of text synthesis

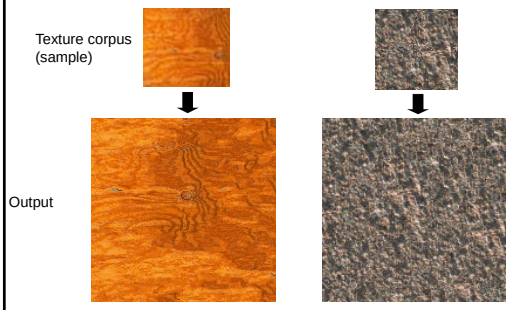

80

Slide from Alyosha Efros, ICCV 1999

Texture synthesis: intuition

Before, we inserted the next word based on existing nearby words…

Now we want to insert **pixel intensities** based on existing nearby pixel values.

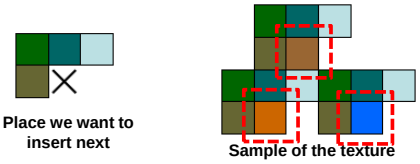

Distribution of a value of a pixel is conditioned on its neighbors alone.

81

Slide credit: Kristen Grauman

Synthesizing One Pixel

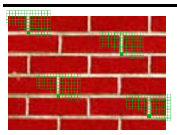

input image

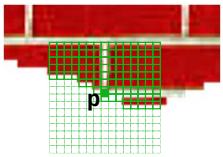

synthesized image

- What is $P(\mathrm{x}|\mathrm{neighborhood\ of\ pixels\ around\ x})$?
- Find all the windows in the image that match the neighborhood
- To synthesize **x**
  - pick one matching window at random
  - assign **x** to be the center pixel of that window
- An **exact** neighbourhood match might not be present, so find the **best** matches using **SSD error** and randomly choose between them, preferring better matches with higher probability

82

Slide from Alyosha Efros, ICCV 1999

Neighborhood Window

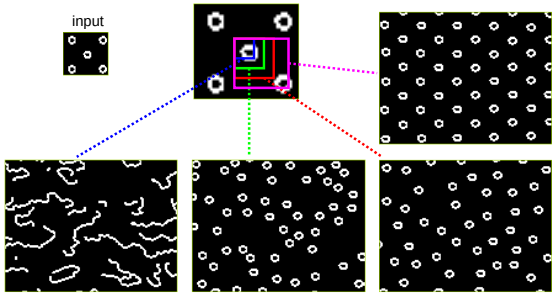

Slide from Alyosha Efros, ICCV 1999

83

Varying Window Size

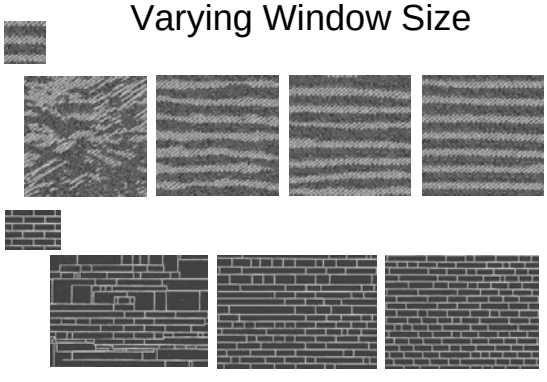

Increasing window size

84

Slide from Alyosha Efros, ICCV 1999

Growing Texture

- Starting from the initial image, "grow" the texture one pixel at a time

85

Slide from Alyosha Efros, ICCV 1999

Synthesis results

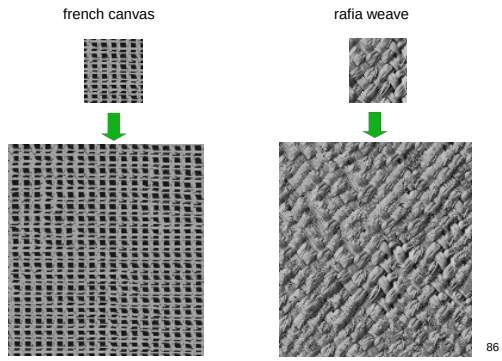

86

Slide from Alyosha Efros, ICCV 1999

Synthesis results

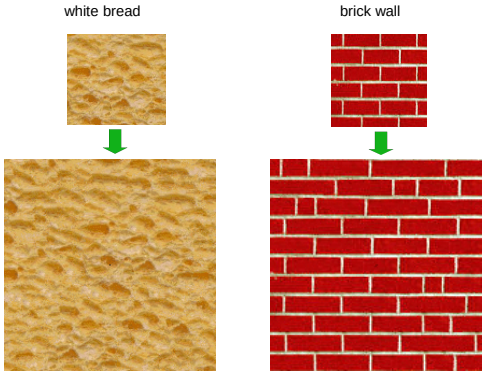

87

Slide from Alyosha Efros, ICCV 1999

Synthesis results

88

Slide from Alyosha Efros, ICCV 1999

Failure Cases

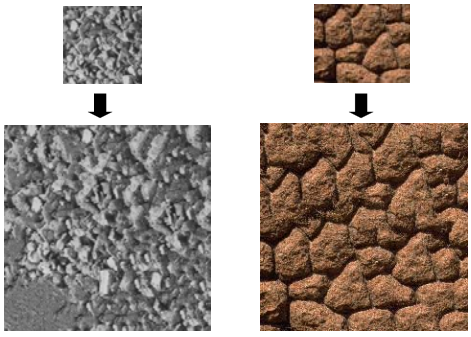

Growing garbage **Verbatim copying**

89

Slide from Alyosha Efros, ICCV 1999

Hole Filling

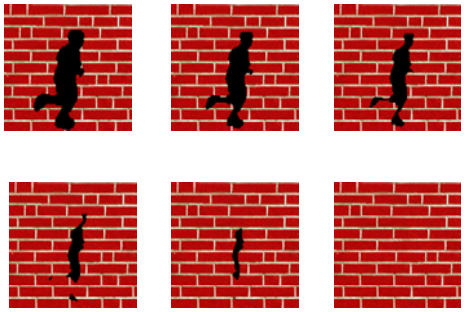

90

Slide from Alyosha Efros, ICCV 1999

Extrapolation

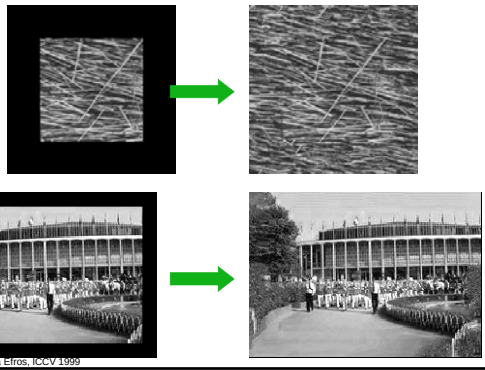

Slide from Alyosha Efros, ICCV 1999

91

- The Efros & Leung algorithm
  - Simple
  - Surprisingly good results
  - Synthesis is easier than analysis!
  - …but very slow

92

Image Quilting [Efros & Freeman 2001]

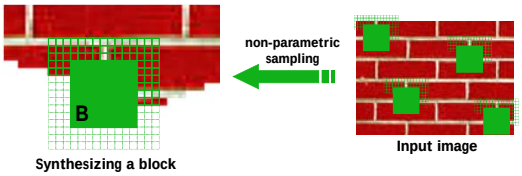

Synthesizing a block

- Observation: neighbor pixels are highly correlated

Idea: unit of synthesis = block

- Exactly the same but now we want P(**B**|N(**B**))
- Much faster: synthesize all pixels in a block at once

93

Slide credit: Alyosha Efros

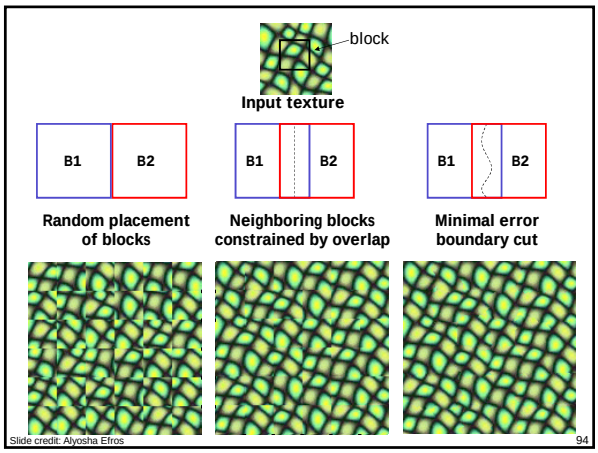
block
Input texture
B1
B2
B1
B2
B1
B2
Random placement of blocks
Neighboring blocks constrained by overlap
Minimal error boundary cut
Slide credit: Alyosha Efros
94

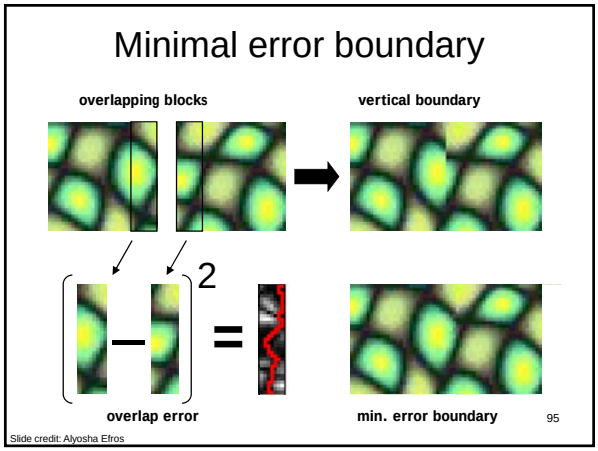
Minimal error boundary
overlapping blocks
vertical boundary
2
overlap error
min. error boundary
95
Slide credit: Alyosha Efros

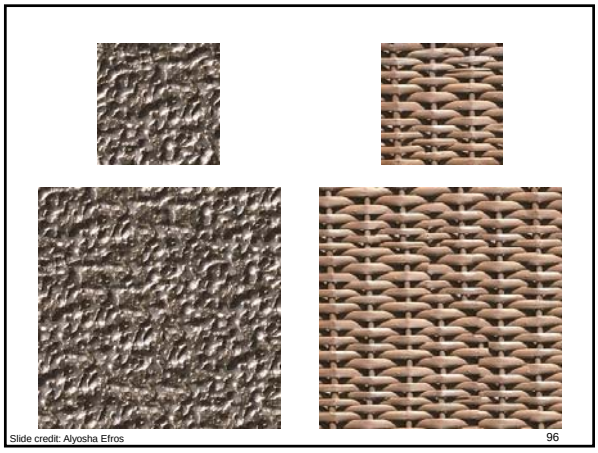
Slide credit: Alyosha Efros
96

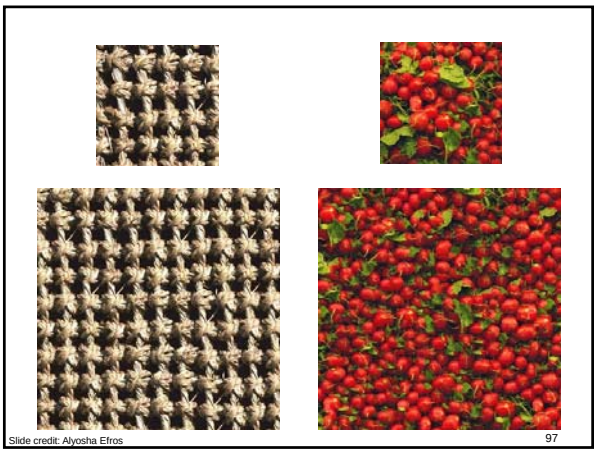

Slide credit: Alyosha Efros

97

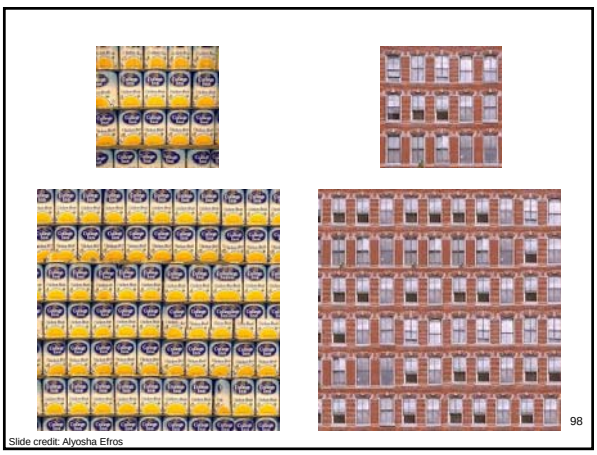

Slide credit: Alyosha Efros

98

Slide credit: Alyosha Efros

99

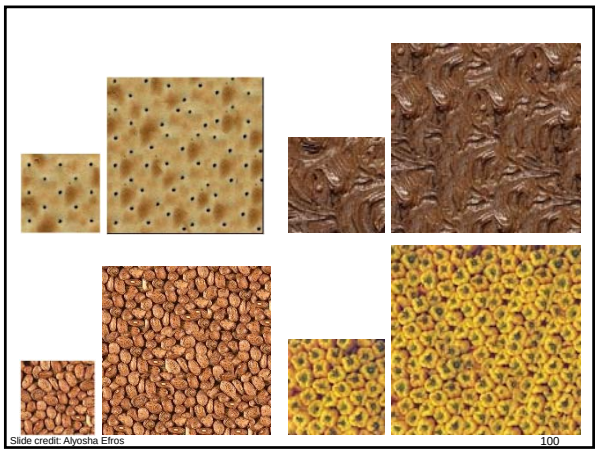

Slide credit: Alyosha Efros

100

Slide credit: Alyosha Efros

101

Texture Transfer

- Take the texture from one object and “paint” it onto another object
  - This requires separating texture and shape
  - That’s HARD, but we can cheat
  - Assume we can capture shape by boundary and rough shading

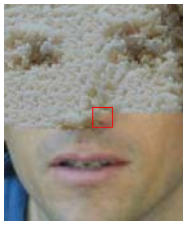

- Then, just add another constraint when sampling: similarity to underlying image at that spot

102

Slide credit: Alyosha Efros

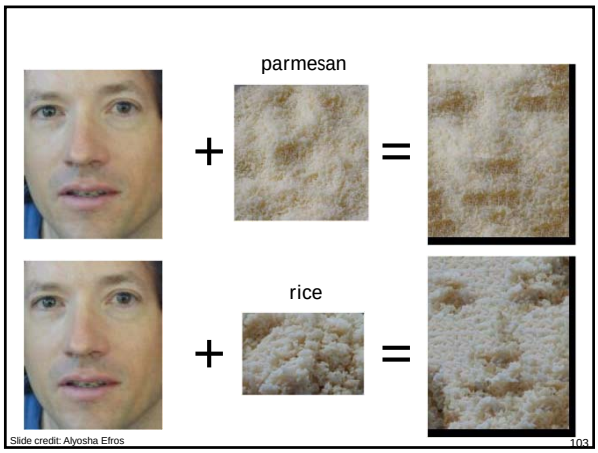

Slide credit: Alyosha Efros

103

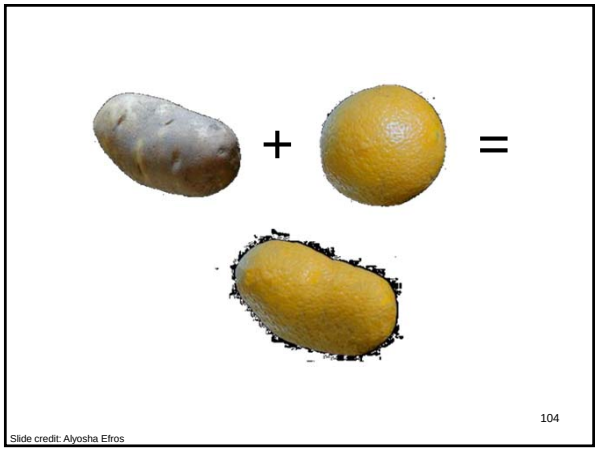

Slide credit: Alyosha Efros

104

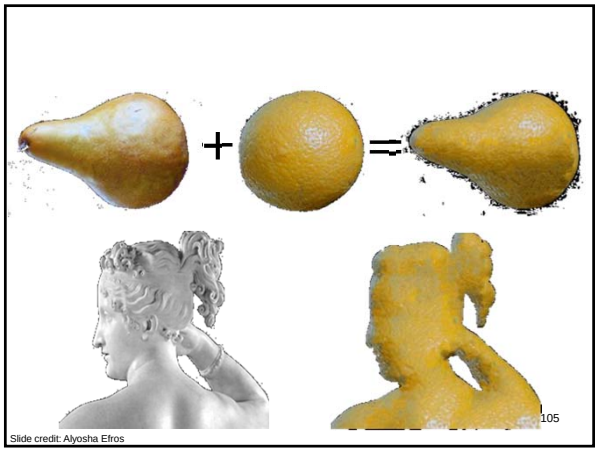

Slide credit: Alyosha Efros

105

Gatys et al., CVPR 2016

106

Gatys et al., CVPR 2016

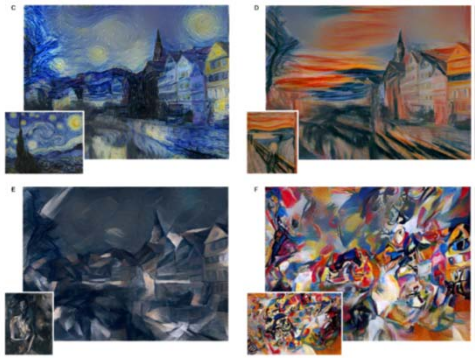

107

(Manual) texture synthesis in the media

108

Slide credit: Kristen Grauman

(Manual) texture synthesis in the media

109

Slide credit: Kristen Grauman

110

http://www.dailykos.com/story/2004/10/27/22442/878

Slide credit: Kristen Grauman

Summary

- Texture is a useful property that is often indicative of materials, appearance cues
- **Texture representations** attempt to summarize repeating patterns of local structure
- **Filter banks** useful to measure variety of structures in local neighborhood
  - Feature spaces can be multi-dimensional
- Neighborhood statistics can be exploited to "sample" or **synthesize** new texture regions
  - Example-based technique

111

Questions?

See you Tuesday!

112