



Texture
April 16th, 2020

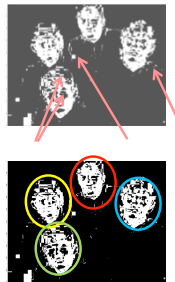
Krishna Kumar Singh
UC Davis

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”

Issues

- What to do with “noisy” binary outputs?
 - Holes
 - Extra small fragments
- How to demarcate multiple regions of interest?
 - Count objects
 - Compute further features per object



Slide credit: Kristen Grauman

Connected components

- Identify distinct regions of “connected pixels”

1	1	0	1	1	1	0	1
1	1	0	1	0	1	0	1
1	1	1	1	0	0	0	1
0	0	0	0	0	0	0	1
1	1	1	1	0	1	0	1
0	0	0	1	0	1	0	1
1	1	0	1	0	0	0	1
1	1	0	1	0	1	1	1

a) binary image

3	1	1	0	1	1	1	0	2
1	1	0	1	0	1	0	2	
1	1	1	1	0	0	0	2	
0	0	0	0	0	0	0	2	
3	3	3	3	0	4	0	2	
0	0	0	3	0	4	0	2	
5	5	0	3	0	0	0	2	
5	5	0	3	0	2	2	2	

b) connected components labeling

c) binary image and labeling, expanded for viewing

>> L = bwlabel(BW,conn) 4

Connectedness

- Defining which pixels are considered neighbors

	↑	
←	[i, j]	→
	↓	

4-connected

↖	↑	↗
←	[i, j]	→
↙	↓	↘

8-connected

Slide credit: Chaitanya Chandra 5

Connected components


connected components of 1's from thresholded image

connected components of cluster labels

Slide credit: Pinar Duygulu 6

Region properties

- Given connected components, can compute simple features per blob, such as:
 - Area (num pixels in the region)
 - Centroid (average x and y position of pixels in the region)
 - Bounding box (min and max coordinates)



Slide credit: Kristen Grauman 7

Binary image analysis: basic steps (recap)

- Convert the image into binary form
 - Thresholding
- Clean up the thresholded image
 - Morphological operators
- Extract separate blobs
 - Connected components
- Describe the blobs with region properties

Slide credit: Kristen Grauman 8

Matlab

```

• L = bwlabel (BW,8);
• STATS = regionprops (L,PROPERTIES) ;
  - 'Area'
  - 'Centroid'
  - 'BoundingBox'
  - 'Orientation', ...
• IM2 = imerode (IM,SE);
• IM2 = imdilate (IM,SE);
• IM2 = imclose (IM, SE);
• IM2 = imopen (IM, SE);
    
```

Slide adapted from Kristen Grauman 9

Binary images

- Pros
 - Can be fast to compute, easy to store
 - Simple processing techniques available
 - Lead to some useful compact shape descriptors
- Cons
 - Hard to get “clean” silhouettes
 - Noise common in realistic scenarios
 - Can be too coarse of a representation

Slide credit: Kristen Grauman

10


Summary

- Operations, tools
 - Derivative filters
 - Smoothing, morphology
 - Thresholding
 - Connected components
 - Matching filters
 - Histograms
- Features, representations
 - Edges, gradients
 - Blobs/regions
 - Local patterns
 - Textures (next)
 - Color distributions

Slide credit: Kristen Grauman

11


Chamfer matching system



- Gavrilu et al. http://gavrilu.net/Research/Chamfer_System/chamfer_system.html

Slide credit: Kristen Grauman

Chamfer matching system




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

Slide credit: Kristen Grauman


ShadowDraw [Lee et al., SIGGRAPH 2011]

[video](#)



14


Today: Texture



What defines a texture?

15


Includes: more regular patterns



Slide credit: Kristen Grauman

16

Includes: more random patterns



Slide credit: Kristen Grauman

17

Texture-related tasks

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

Slide credit: Kristen Grauman

18

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/~angle/phdpics1.html>

19

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

20

Slide credit: Kristen Grauman

Analysis vs. Synthesis

Why analyze texture?

Images: Bill Freeman, A. Efros

21



Image credit: D. Forsyth

22


Why analyze texture?

Importance to perception:

- Often indicative of a material's properties

Slide credit: Kristen Grauman

23



24


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

Slide credit: Kristen Grauman

25



Slide credit: Kristen Grauman

26


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

Slide credit: Kristen Grauman

27



Slide credit: Kristen Grauman <http://animals.nationalgeographic.com/>

Why analyze texture?

Importance to perception:

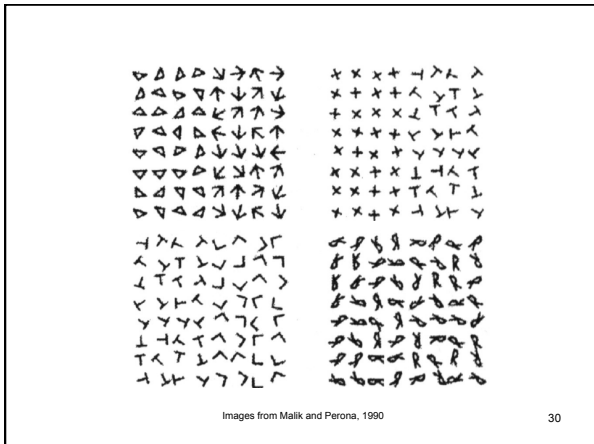
- Often indicative of a material's properties
- Can be important appearance cue, especially if shape is similar across objects
- Aim to distinguish between boundaries and texture

Technically:

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

29

Slide credit: Kristen Grauman



Images from Malik and Perona, 1990

30


Psychophysics of texture

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


Same or different?

Slide credit: Kristen Grauman


31




32



33



Capturing the local patterns with image measurements



[Bergen & Adelson, *Nature* 1988]

Scale of patterns influences discriminability

Size-tuned linear filters

35

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

Slide credit: Kristen Grauman

36

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

Side credit: Kristen Grauman 37

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

Side credit: Kristen Grauman 38

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

Side credit: Kristen Grauman 39

Texture representation: example

Dimension 2 (mean d/dy value)

Dimension 1 (mean d/dx value)

	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 40

Slide credit: Kristen Grauman

Texture representation: example

Windows with primarily horizontal edges

Both

Dimension 2 (mean d/dy value)

Dimension 1 (mean d/dx value)

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

Windows with small gradient in both directions

Windows with primarily vertical edges

statistics to summarize patterns in small windows 41

Slide credit: Kristen Grauman

Texture representation: example

original image

derivative filter responses, squared

visualization of the assignment to texture "types"

42

Slide credit: Kristen Grauman

Texture representation: example

Dimension 2 (mean d/dy value)

Dimension 1 (mean d/dx value)

Far: dissimilar textures

Close: similar textures

	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 43

Slide credit: Kristen Grauman

Texture representation: example

Dimension 2

Dimension 1

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

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

44

Slide credit: Kristen Grauman

Texture representation: example

Dimension 2

Dimension 1

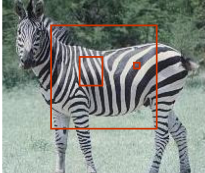
Distance reveals how dissimilar texture from window a is from texture in window b.

45

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.

46

Filter banks

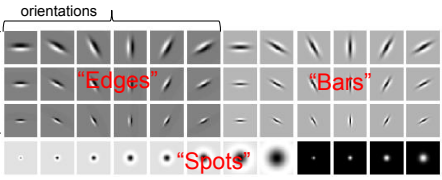
- Our previous example used two filters, and resulted in a 2-dimensional feature vector to describe texture in a window
 - x and y derivatives revealed something about local structure
- We can generalize to apply a collection of multiple (d) filters: a “filter bank”
- Then our feature vectors will be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

Side credit: Kristen Grauman

47

Filter banks

orientations



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

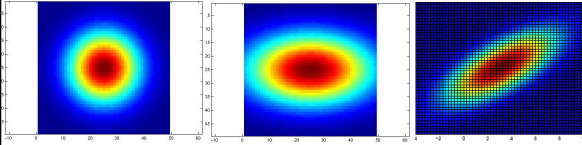
Matlab code available for these examples: <http://www.robots.ox.ac.uk/~vgg/research/texclass/filters.html>

Side credit: Kristen Grauman

48

Multivariate Gaussian

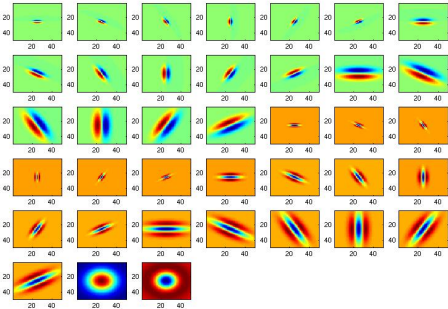
$$p(x; \mu, \Sigma) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1} (x - \mu)\right)$$



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

Side credit: Kristen Grauman 49

Filter bank





Side credit: Kristen Grauman 50





Image from <http://www.teasesplorer.com/austincap2.jpg>


Side credit: Kristen Grauman 51




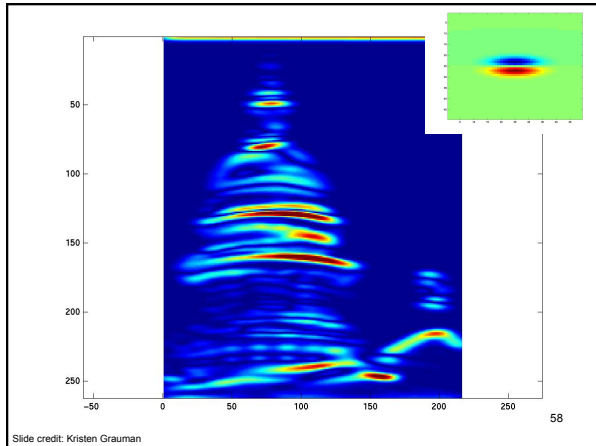



































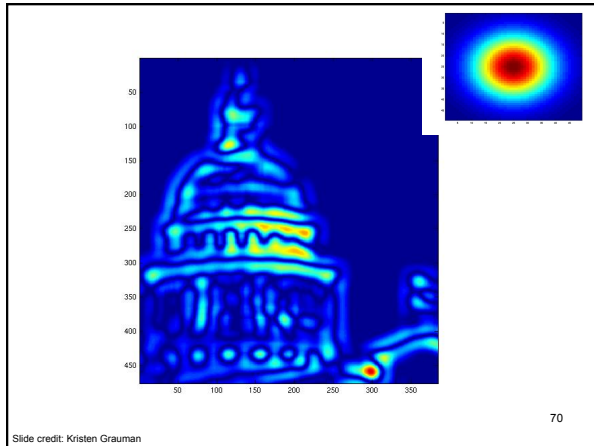


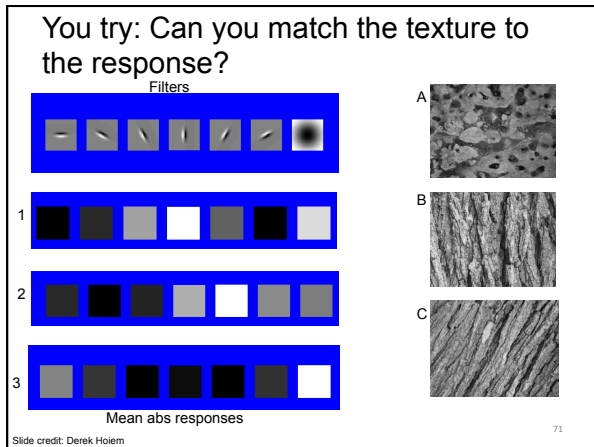


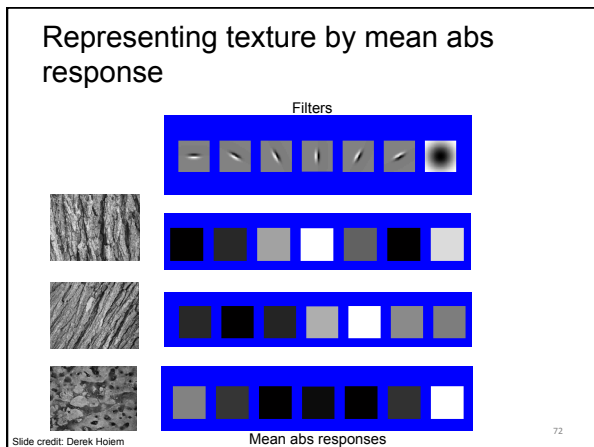












[r1, r2, ..., r38]

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

Slide credit: Kristen Grauman

73

d-dimensional features

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

Euclidean distance (L_2)

2d

Slide credit: Kristen Grauman

74

Example uses of texture in vision: analysis

Slide credit: Kristen Grauman

75

Classifying materials, "stuff"






Figure by Varma & Zisserman 76



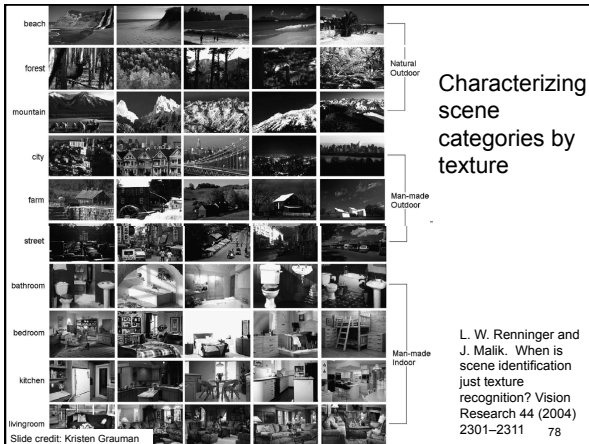
Texture features for image retrieval



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.

Slide credit: Kristen Grauman

77



Characterizing scene categories by texture

L. W. Renninger and J. Malik. When is scene identification just texture recognition? *Vision Research* 44 (2004) 2301-2311 78

Slide credit: Kristen Grauman

Texture-related tasks

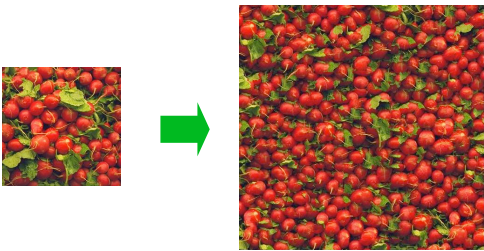
- **Shape from texture**
 - Estimate surface orientation or shape from image texture
- **Segmentation/classification** from texture cues
 - Analyze, represent texture
 - Group image regions with consistent texture
- **Synthesis**
 - Generate new texture patches/images given some examples

79

Slide credit: Kristen Grauman

Texture synthesis

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

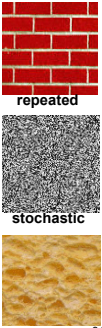


80

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.

81

Markov Chains

Markov Chain

- a sequence of random variables X_1, X_2, \dots, X_n
- X_t is the **state** of the model at time t

$$X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow X_4 \rightarrow X_5$$

82

Markov Chain Example: Text

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

a												
dog												
is												
man's												
best												
friend												
it's												
eat												
world												
out												
there												
.												
	a	dog	is	man's	best	friend	it's	eat	world	out	there	.

X_{t-1} $p(x_t|x_{t-1})$ $p(\text{dog}|a) = ?$

X_t

Slide credit: Adapted by Devi Parikh from Steve Seitz

83

Markov Chain Example: Text

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

a	2/3	1/3										
dog		1/3					1/3	1/3				
is	1											
man's				1								
best					1							
friend											1	
it's	1											
eat		1										
world									1			
out											1	
there												1
.						1						
	a	dog	is	man's	best	friend	it's	eat	world	out	there	.

X_{t-1} $p(x_t|x_{t-1})$

X_t

Slide credit: Steve Seitz

84

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 $p(X_t | X_{t-1}, \dots, X_{t-(n-1)})$

WE NEED TO EAT CAKE

Slide credit: Steve Seitz

85

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.

Slide from Alyosha Efros, ICCV 1999

86

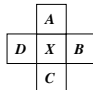
Markov Random Field

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:

$$P(X|A, B, C, D)$$


Slide credit: Steve Seitz

87

Texture Synthesis [\[Efros & Leung, ICCV 99\]](#)

Can apply 2D version of text synthesis

Texture corpus (sample)

Output

Slide from Alyosha Efros, ICCV 1999 88

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.

Place we want to insert next

Sample of the texture ("corpus")

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

Slide credit: Kristen Grauman 89


Synthesizing One Pixel

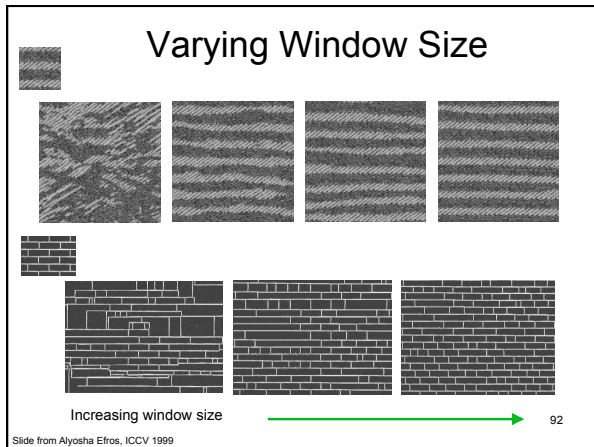
input image


synthesized image


- What is $P(x|\text{neighborhood of pixels around } x)$?
- 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


Slide from Alyosha Efros, ICCV 1999 90

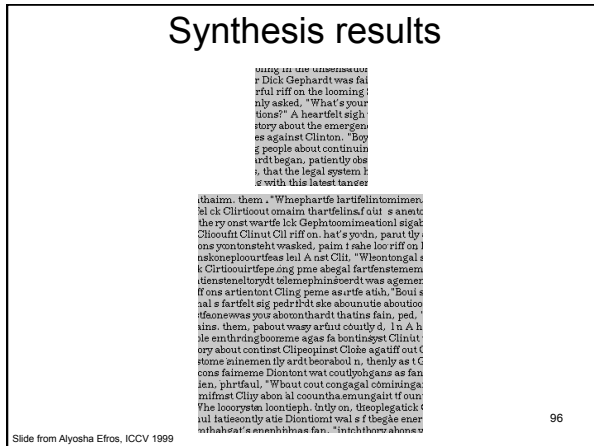





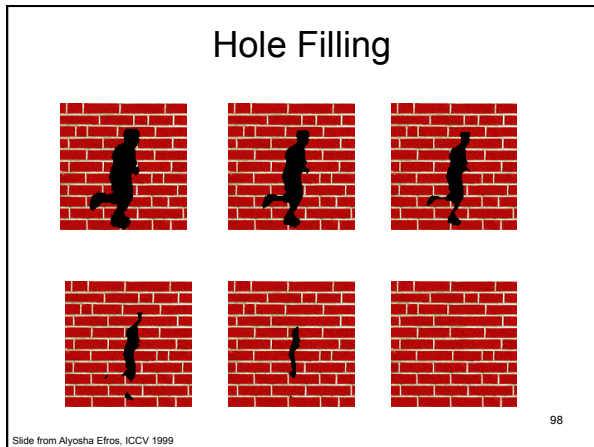















- The Efros & Leung algorithm
 - Simple
 - Surprisingly good results
 - Synthesis is easier than analysis!
 - ...but very slow

100

Image Quilting [Efros & Freeman 2001]

- **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

Slide credit: Alyosha Efros

101

Input texture

block

B1 B2

B1 B2

B1 B2

Random placement of blocks

Neighboring blocks constrained by overlap

Minimal error boundary cut

Slide credit: Alyosha Efros

102

Minimal error boundary

overlapping blocks vertical boundary

overlap error min. error boundary

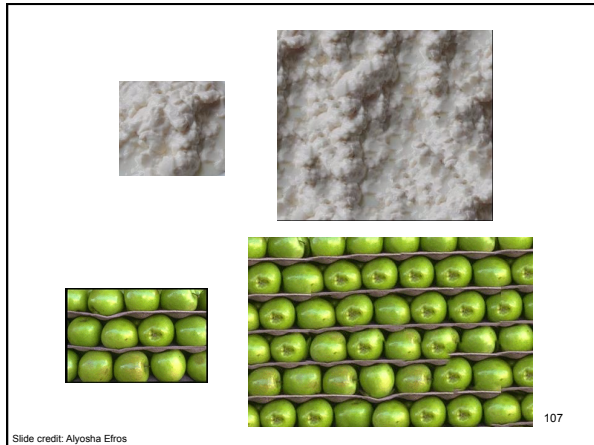
Side credit: Alyosha Efros 103

Side credit: Alyosha Efros 104


Side credit: Alyosha Efros 105



Slide credit: Alyosha Efros



Slide credit: Alyosha Efros



Slide credit: Alyosha Efros


Failures
(Chernobyl Harvest)



Slide credit: Alyosha Efros 109

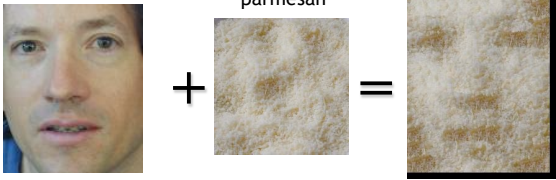
Texture Transfer

- 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

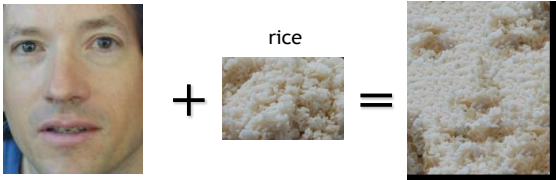


Slide credit: Alyosha Efros 110

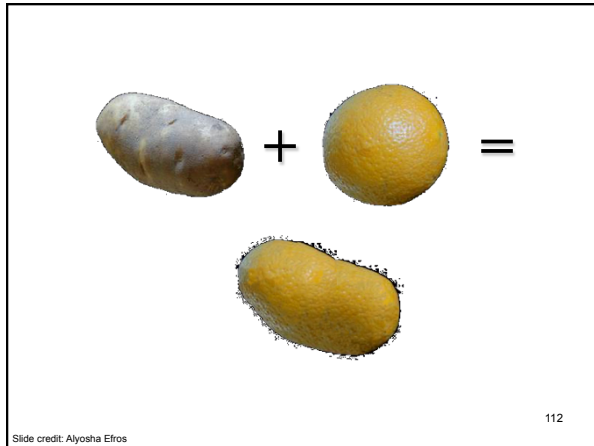
parmesan

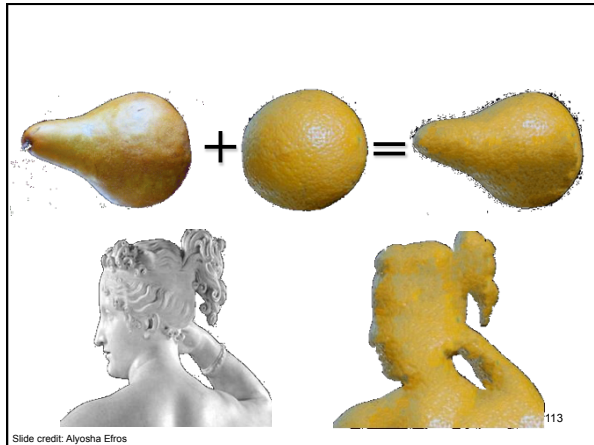


rice



Slide credit: Alyosha Efros 111







Gatys et al., CVPR 2016

115

(Manual) texture synthesis in the media

116

Slide credit: Kristen Grauman

(Manual) texture synthesis in the media

117

Slide credit: Kristen Grauman



Slide credit: Kristen Graunlich <http://www.dailykos.com/story/2004/10/27/22442/878>

118

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

119

Questions?

See you Tuesday!

120
