


Texture

COS 429

Princeton University




Texture

What is a texture?






Texture

What is a texture?




Texture

What is a texture?



Texture

- Texture is stochastic and stationary (same regardless of position)




stationary



non-stationary


Texture



Stochastic

Stationary

Texture




Stochastic

Stationary

Goal

Computational representation of texture

- Textures generated by same stationary stochastic process have same representation
- Perceptually similar textures have similar representations



5, 7, 34, 2, 199, 12

Hypothetical texture representation

Applications?

Applications

Segmentation

3D Reconstruction

Classification

Synthesis

etc.




Applications

Segmentation ←

3D Reconstruction

Classification

Synthesis




Applications

Segmentation

3D Reconstruction ←

Classification

Synthesis




Applications

Segmentation

3D Reconstruction

Classification ←

Synthesis




Applications

Segmentation


3D Reconstruction

Classification

Synthesis ←



Input



Output

Texture Representation?

What's a good texture representation?


- Textures generated by same stationary stochastic process have same representation
- Perceptually similar textures have similar representations



Texture Representation?

Julesz conjectured that the putative units of pre-attentive human texture perception are related to local features, such as edges, line ends, blobs, etc.

B. Julesz. Textons, the Elements of Texture Perception, and their Interactions. *Nature*, 290:91-97, March 1981.



Texture Representation?

Research suggests that the human visual system performs local spatial frequency analysis (Gabor filters)

J. J. Kulikowski, S. Marcelja, and P. Bishop.
Theory of spatial position and spatial frequency relations in the
receptive fields of simple cells in the visual cortex.
Biol. Cybern, 43:187-198, 1982.


Texture Representation

- Analyze textures based on the responses of linear filters
 - Use filters that look like patterns (spots, edges, bars, ...)
 - Compute magnitudes of filter responses
- Represent textures with statistics of filter responses within local windows
 - Histogram of feature responses for all pixels in window

Texture Representation Example



original image




derivative filter responses, squared

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

⋮

statistics to summarize patterns in small windows

Texture Representation Example




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

statistics to
summarize patterns
in small windows

Texture Representation Example

Windows with primarily horizontal edges

Both



Windows with small gradient in both directions

Windows with primarily vertical edges



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

statistics to summarize patterns in small windows


Texture Representation Example



original image




derivative filter
responses, squared



visualization of the
assignment to
texture "types"


Texture Representation Example



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


statistics to
summarize patterns
in small windows

Texture Representation Example




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

Texture Representation Example




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



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

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>

Multivariate Gaussian


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



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



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



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

Filter bank

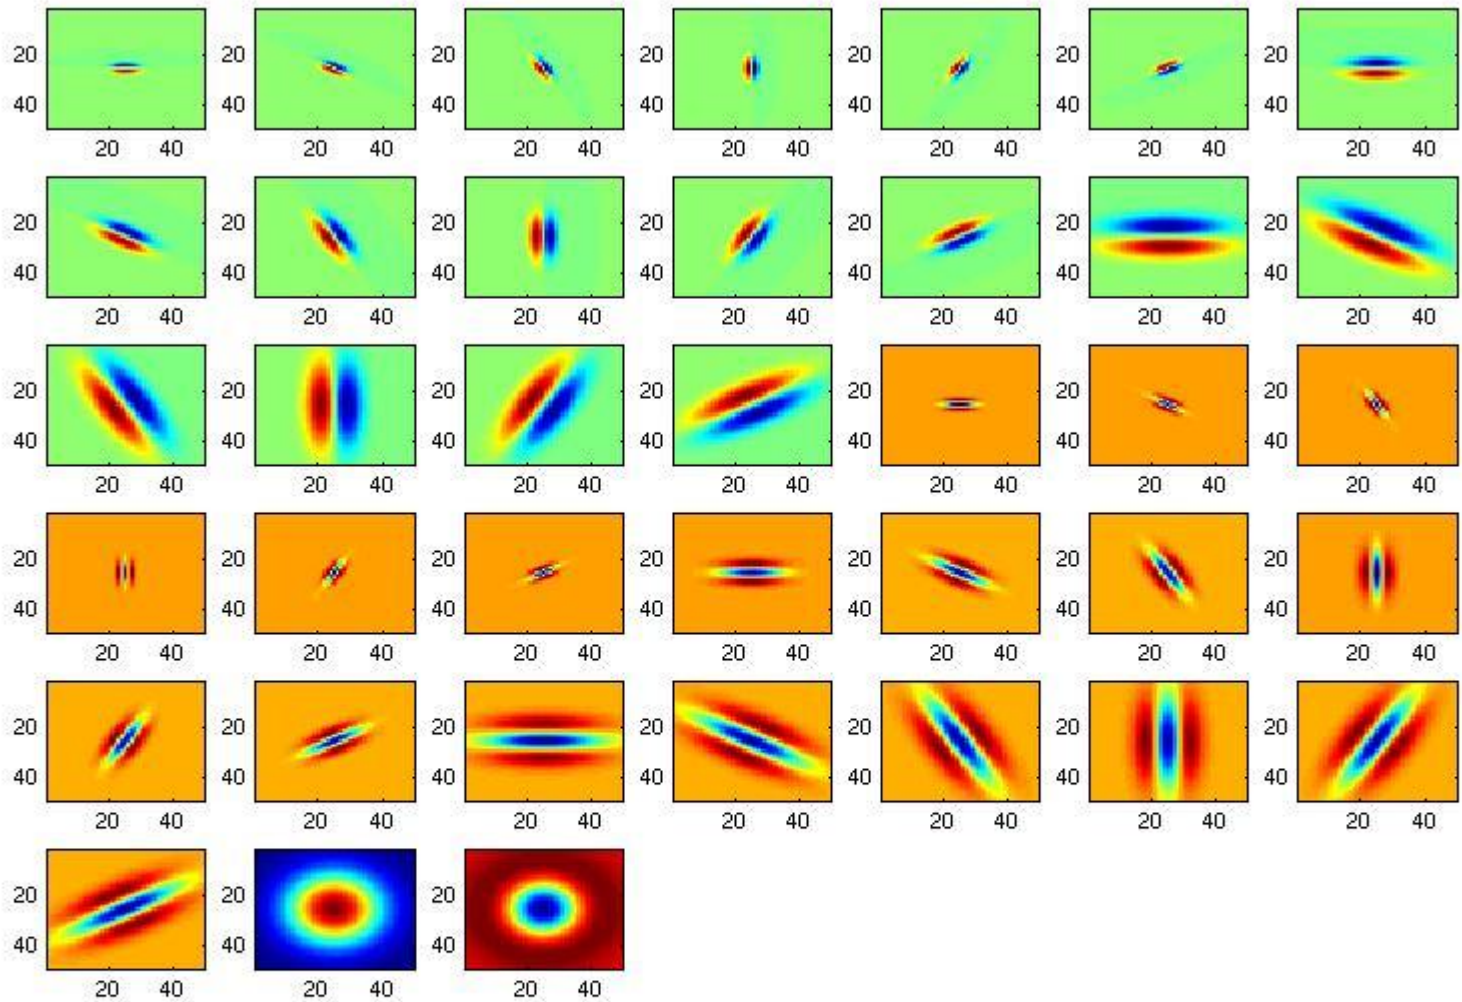





Image from <http://www.texasexplorer.com/austincap2.jpg>



Kristen Grauman























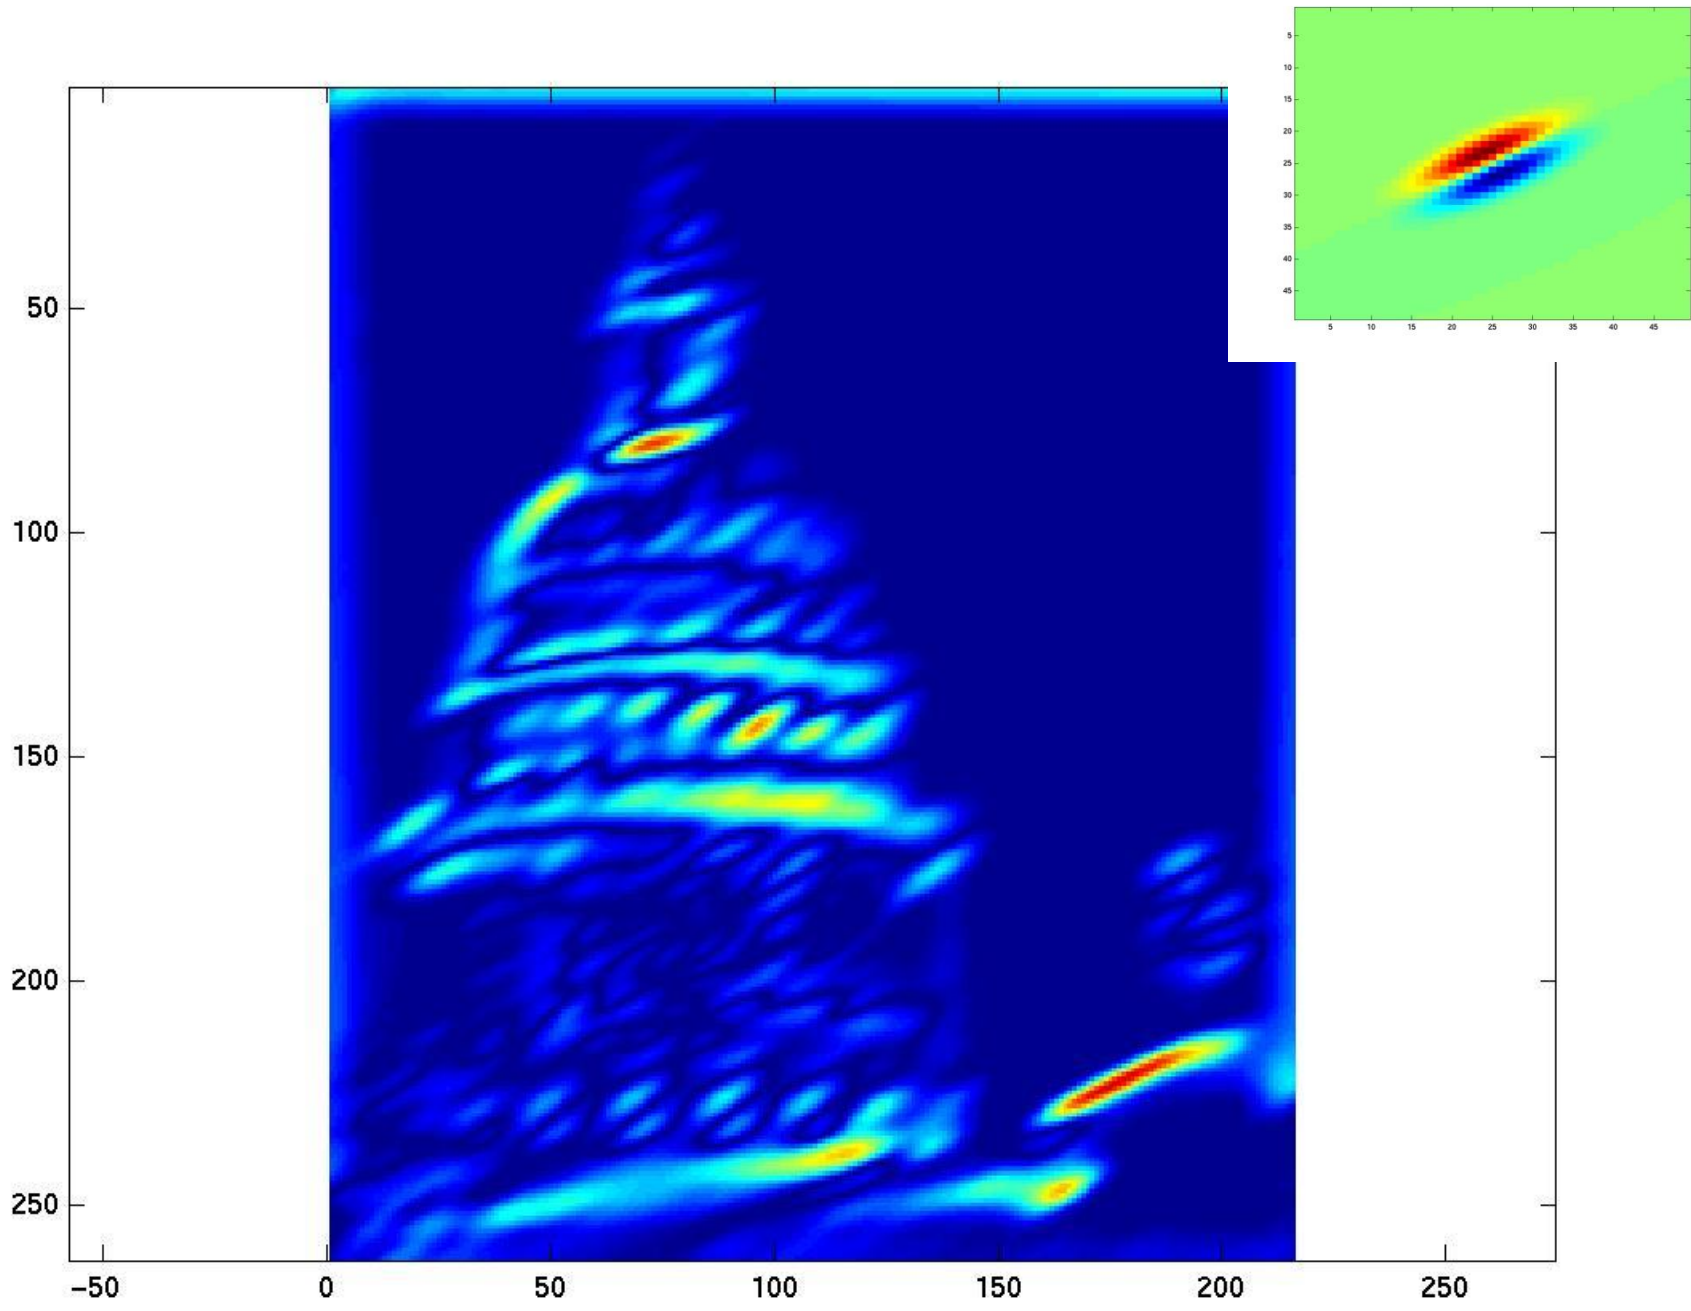




































You try: Can you match the texture to the response?


Filters




1



2




3



Mean abs responses

A



B








C

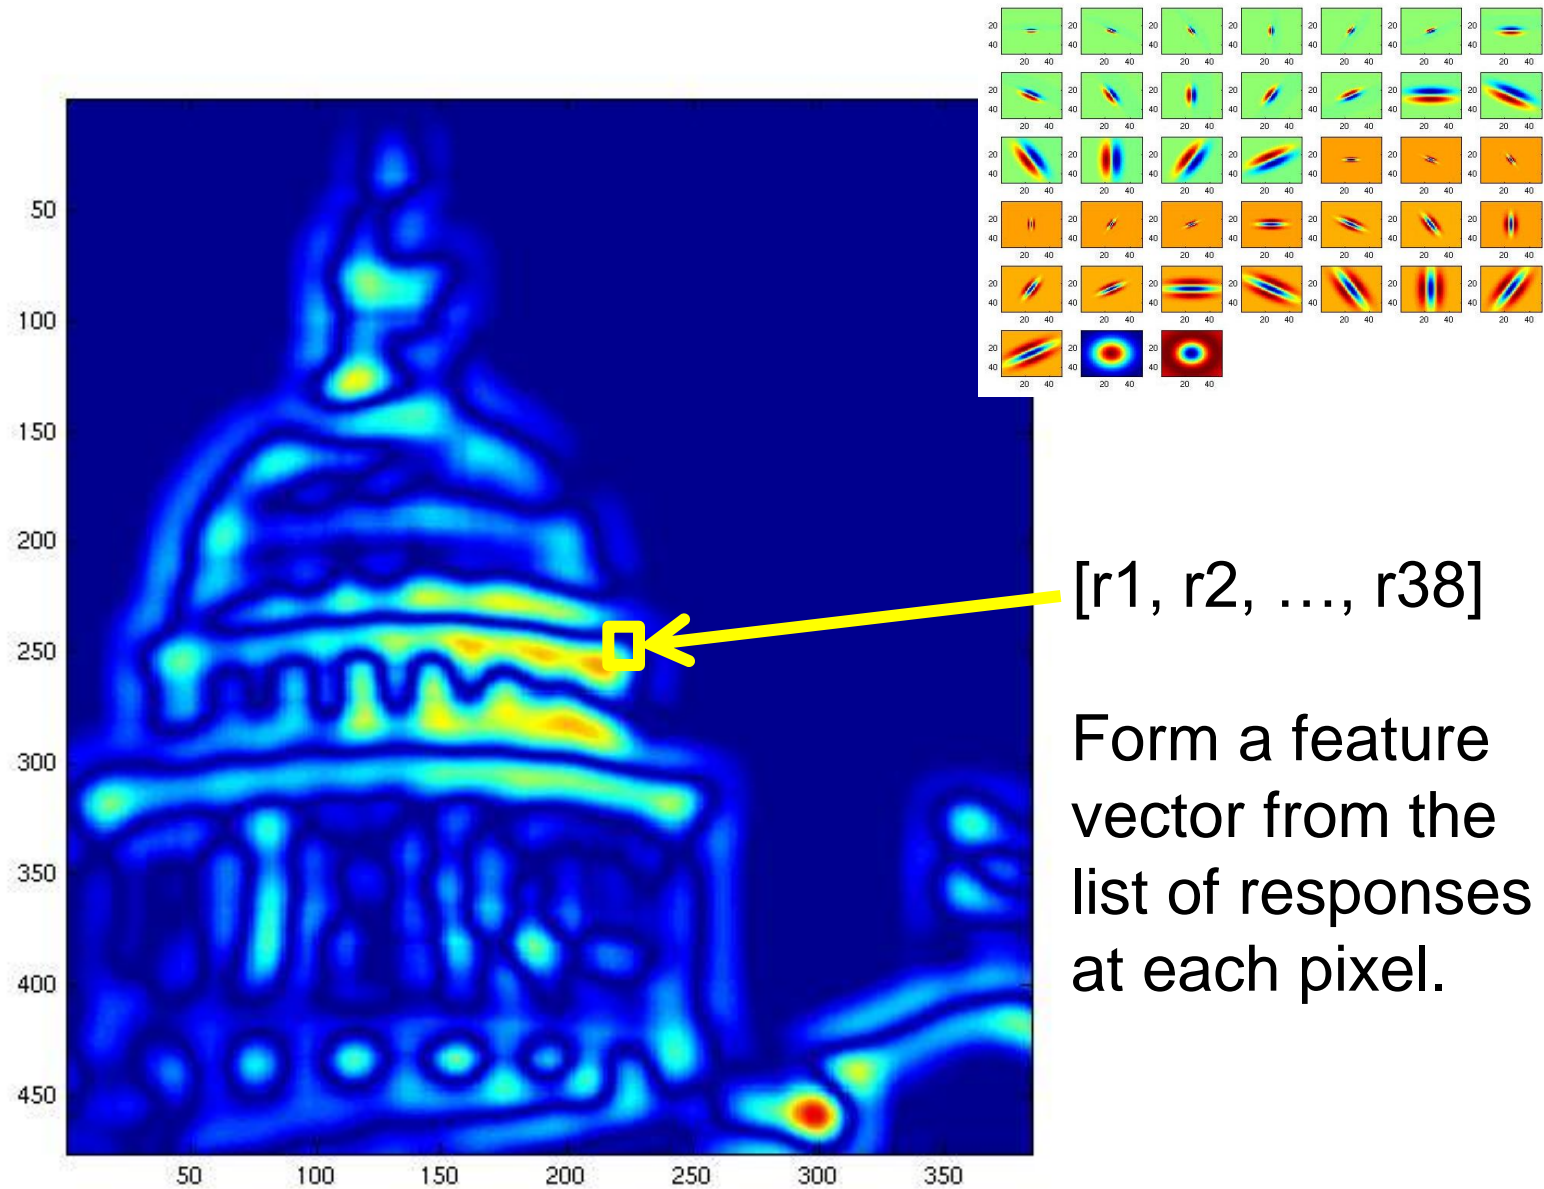


Representing texture by mean abs response

Filters



Mean abs responses




Texture Representation Summary

- Analyze textures based on the responses of linear filters
 - Use oriented filters at multiple scales
 - Compute magnitudes of filter responses
- Represent textures with statistics of filter responses within local windows
 - Histogram of feature responses for all pixels in window

Example uses of texture analysis in computer vision

Similarity

- Predict perceptual similarity of textures
 - based on Euclidean distance (L_2) in d-dimensional feature space



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

Segmentation



Segment
images with
color and
texture


BlobWorld

Segmentation

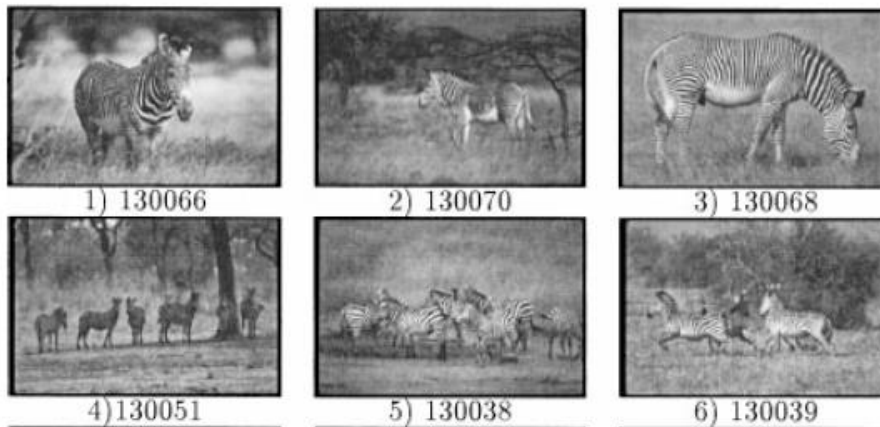


Segment
aerial imagery
by textures

Retrieval



Retrieve
similar images
based on texture



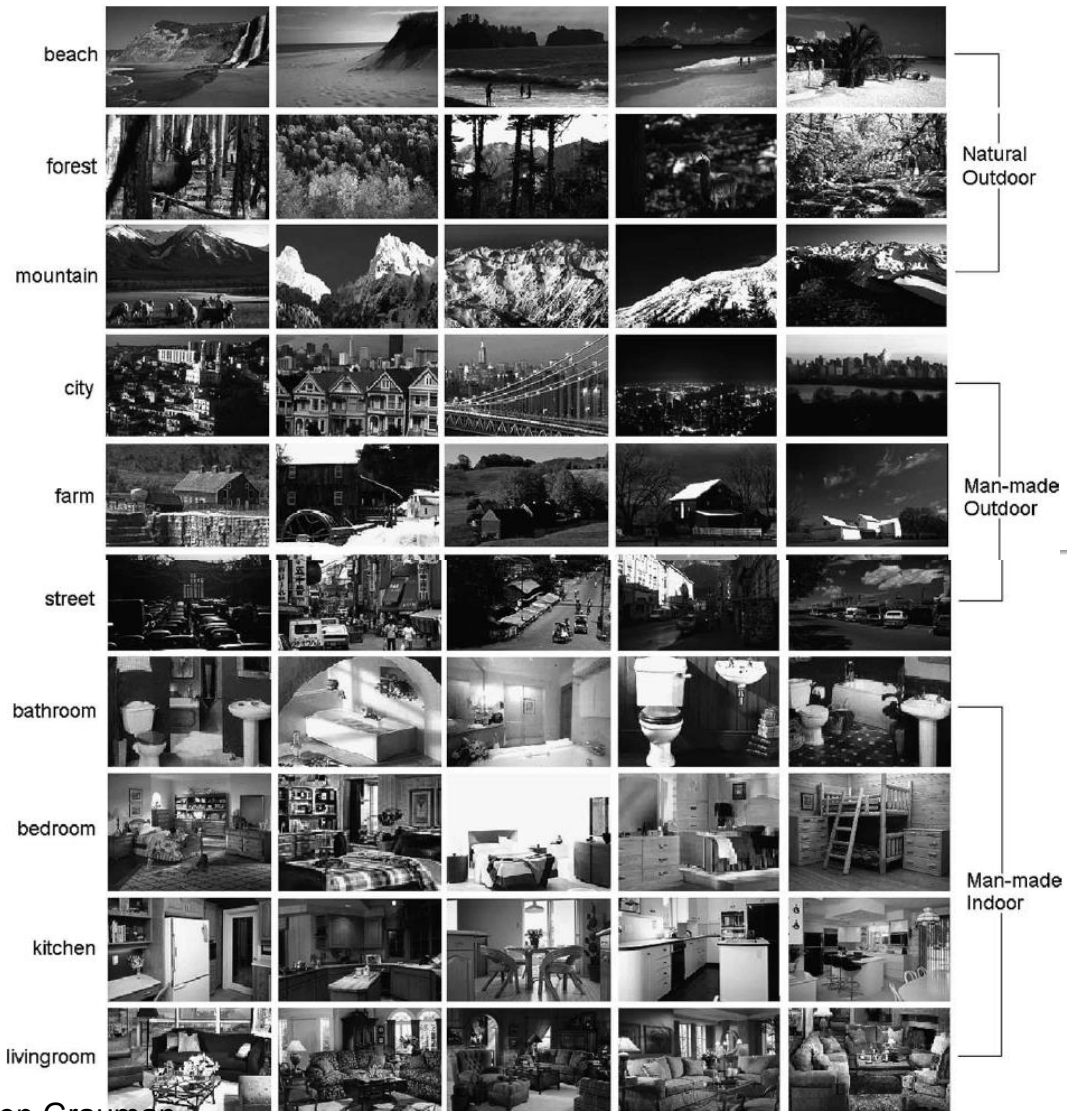
Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99-121, November 2000,

Classification



Figure by Varma & Zisserman

Classification



Characterizing
scene
categories by
texture


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

Texture synthesis

Create new image (e.g., of different size or shape) with texture of an input image



Input




Output

Texture synthesis

How can we do this?



Input




Output

Texture synthesis

Copying texture multiple times produces seams and repetitions




Input




Bad Output

Texture Synthesis

Can synthesize new texture by sampling from the probability distribution of local neighborhoods



input image




synthesized image

$$P(\mathbf{x} | \text{neighborhood of pixels around } \mathbf{x})$$

Texture Synthesis

Texture is stochastic and stationary

– $p(\text{pixel}) = p(\text{pixel} \mid \text{neighborhood})$



Motivation from Language

- Shannon (1948) proposed a way to synthesize new text using *N-grams*
 - Use a large text to compute probability distributions of each letter given $N-1$ previous letters
 - Starting from a seed repeatedly sample the conditional probabilities to generate new letters
 - Can do this with words too ...


Mark V. Shaney (Bell Labs)

- Results (using alt.singles corpus):
 - *“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.”*
- Notice how well local structure is preserved!

Now let's try this in 2D... [Efros](#)

Efros & Leung Algorithm



Compute output pixels in scanline order
(top-to-bottom, right-to-left)



output image

Efros & Leung Algorithm

For each pixel, find candidate pixels based on similarities of pixel features in neighborhoods




Efros & Leung Algorithm


Similarities of pixel neighborhoods can be computed with squared differences (SSD) of pixel colors and/or filter bank responses

$$\| \left(\begin{array}{cccc} \text{blue} & \text{light blue} & \text{light purple} & \text{purple} \\ \text{blue} & \text{light blue} & \text{light purple} & \text{purple} \\ \text{blue} & \text{light blue} & \text{black} & \text{black} \end{array} - \begin{array}{cccc} \text{light blue} & \text{light purple} & \text{light purple} & \text{purple} \\ \text{light blue} & \text{light purple} & \text{purple} & \text{purple} \\ \text{purple} & \text{light blue} & \text{black} & \text{black} \end{array} \right) \| ^2$$

Efros & Leung Algorithm

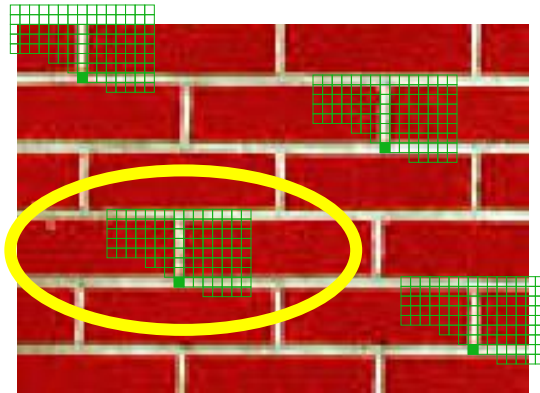


input image




- For each pixel x :
 - Find the best matching K windows from the input image
 - Pick one matching window at random
 - Assign x to be the center pixel of that window

Efros & Leung Algorithm




input image



- For each pixel x :
 - Find the best matching K windows from the input image
 - Pick one matching window at random
 - Assign x to be the center pixel of that window



Efros & Leung Algorithm




- For each pixel x :
 - Find the best matching K windows from the input image
 - Pick one matching window at random
 - Assign x to be the center pixel of that window

Synthesis results

french canvas



rafia weave





Synthesis results

white bread



brick wall





Synthesis results

...ing in the unsensadon
r Dick Gephardt was fai
rful riff on the looming
nly asked, "What's your
tions?" A heartfelt sigh
story about the emergen
es against Clinton. "Boy
g people about continuin
ardt began, patiently obs
s, that the legal system h
g with this latest tanger

ithaim. them . "Whnephartfe lartifelintomimen.
lel ck Clirtioout omaim thartfelins.f out s anetc
the ry onst wartfe lck Gephtoomimeationl sigab
Chiooufit Clinut Cll riff on. hat's yodn, parut tly:
ons yontonsteht waked, paim t sahe loo riff on l
nskoneploourtfeas leil A nst Clit, "Wleontongal s
k Cirtioouirtfepe.ong pme abegal fartfenstemem
itiensteneltorydt telemephminsverdt was agemer
ff ons artientont Cling peme as artfe atih, "Boui s
nal s fartfelt sig pedrltdt ske abounutie aboutioo
tfaonewas yow abowonthardt thatins fain, ped, '
ains. them, pabout wasy arfiut coutly d, l n A h
ole emthringbooreme agas fa bontinsyst Clinut
ory about continst Clipeopinst Cloke agatiff out C
stome zinemen tly ardt beorabou n, thenly as t C
cons faimeme Diontont wat coutlyohgans as fan
ien, phrtfaul, "Wbaut cout congagal comininga:
mifmst Cliiy abon al coountha.emungait tf our
The loocrysta loontieph. intly on, theoplegatick C
nul fatiezonly atie Diontiomt wal s f tbegeae ener
nthahgat's enenhhnmas fan. "intchthorv abons w

Failure Cases



Growing garbage





Verbatim copying



Example Applications

- Hole filling and extrapolation
 - Fill pixels in “onion skin” order
 - Within each “layer”, pixels with most neighbors are synthesized first
 - Normalize error by the number of known pixels
 - If no close match can be found, the pixel is not synthesized until the end

Hole Filling



Extrapolation



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 redundant 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