

CMPSCI 370HH: Intro. to Computer Vision

Texture synthesis

University of Massachusetts, Amherst

April 19, 2016

Instructor: Subhransu Maji

Slides credit: Kristen Grauman and others

Next class

- Presentation guidelines
 - 20 mins for each team (random order)
 - 15 mins presentation + 5 mins for questions
- Clearly describe
 - Problem statement
 - Preliminary results
 - What are you going to do the next week (write-up)
- 4-6 page writeup (May 6)
 - No deadline extension

2

Texture synthesis

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



3

The challenge

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



repeated



stochastic



Both?


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

4

Markov chains

Markov chain

- A **sequence** of random variables $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$
- \mathbf{X}_t is the **state** of the model at time t



- **Markov assumption:** each state is dependent only on the previous one
 - dependency given by a **conditional probability**:

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

- The above is actually a **first-order** Markov chain
- An **N 'th-order** Markov chain:

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

Source: S. Seitz 5

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		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					

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

Source: S. Seitz 6

Text synthesis

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

Most basic algorithm

1. Build probability histogram
 - find all blocks of N consecutive words/letters in training documents
 - compute probability of occurrence $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-(n-1)})$
2. Given words $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{k-1}$
 - compute \mathbf{x}_k by sampling from $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-(n-1)})$

WE NEED TO EAT CAKE

Source: S. Seitz 7

Text synthesis


- “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 8

Synthesizing Computer Vision text

- What do we get if we extract the probabilities from a chapter on Linear Filters, and then synthesize new statements?



Check out Yisong Yue's website implementing text generation: build your own text Markov Chain for a given text corpus. <http://www.yisongyue.com/shaney/index.php>

Kristen Grauman

9

Synthesized text

- This means we cannot obtain a separate copy of the best studied regions in the sum.
- All this activity will result in the primate visual system.
- The response is also Gaussian, and hence isn't bandlimited.
- Instead, we need to know only its response to any data vector, we need to apply a low pass filter that strongly reduces the content of the Fourier transform of a very large standard deviation.
- It is clear how this integral exist (it is sufficient for all pixels within a $2k + 1 \times 2k + 1 \times 2k + 1 \times 2k + 1$ — required for the images separately).

Kristen Grauman

10

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 :


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



Source: S. Seitz 11

Texture synthesis

Can apply 2D version of text synthesis





[Efros & Leung, ICCV 99](#)

12

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.

13


Synthesizing one pixel



- 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 14

Neighborhood window




Slide from Alyosha Efros, ICCV 1999 15

Varying window size



Slide from Alyosha Efros, ICCV 1999 16


Growing Texture



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


Slide from Alyosha Efros, ICCV 1999 17

Synthesis results



Slide from Alyosha Efros, ICCV 1999

Synthesis results



Slide from Alyosha Efros, ICCV 1999



Synthesis results

being in the unsensu-
r Dick Gephardt was fair
ful 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
rdt began, patiently obs
s, that the legal system h
g with this latest tanger

ithaim, them. "Whnephartfe lartifelintomimerv
el ck Clirtioout omaim thartfelins.f out s anet
the ry onst wartfe lck Gephtoomimeational sigat
Chiooufit Clinut Cil riff on. hat's yodn, parut ty:
ons ycontonsteht wasked, paim f sahe loo riff on l
nskoneploourtfeas leil A nst Clit. "Wleontongal s
k Clirtioouirtepe.ong pme abegal fartfenstemem
stenstenelorydt telemephinsboerdrt was agemer
ff ons artientont Cling peme asrtfe atih, "Boul s
aal s fartfelt sig pedrtirdt ske abounutie aboutioo
tfeonewas youz abownthardt thatins fain, ped, '
ains, them, pabout wasy arbut coutty d, l n A h
ple emthrdngbooreme agas fa bontinsyst Clinton
ory about continst Clieopinst Cloke agatiff out C
stome ainemen tly ardt beorabou n, thenly as t C
cons faimeme Diontont wat couttyohgans as fan
ien, phrtaul, "Wbaut cout congagal comininga
mifmst Clily abon al coountha.emungait tf oun
Vhe loocrystan loontieph. Inly on, theoplegatick t
ul latiesontly atie Diontiomi wal s f tbegae ener
enthahsat's anenhhmas fan. "intchthru ahonsu

Slide from Alyosha Efros, ICCV 1999

Failure Cases






Growing garbage



Verbatim copying

Slide from Alyosha Efros, ICCV 1999

Extrapolation




Slide from Alyosha Efros, ICCV 1999 22

Texture synthesis


- The Efros & Leung algorithm
 - Simple
 - Surprisingly good results
 - Synthesis is easier than analysis!
 - ... but can be very slow
 - $[n \ m]$ image synthesis from $[p \ q]$ image requires $n \times m \times p \times q$ patch comparisons

Image Quilting [Efros & Freeman 2001]





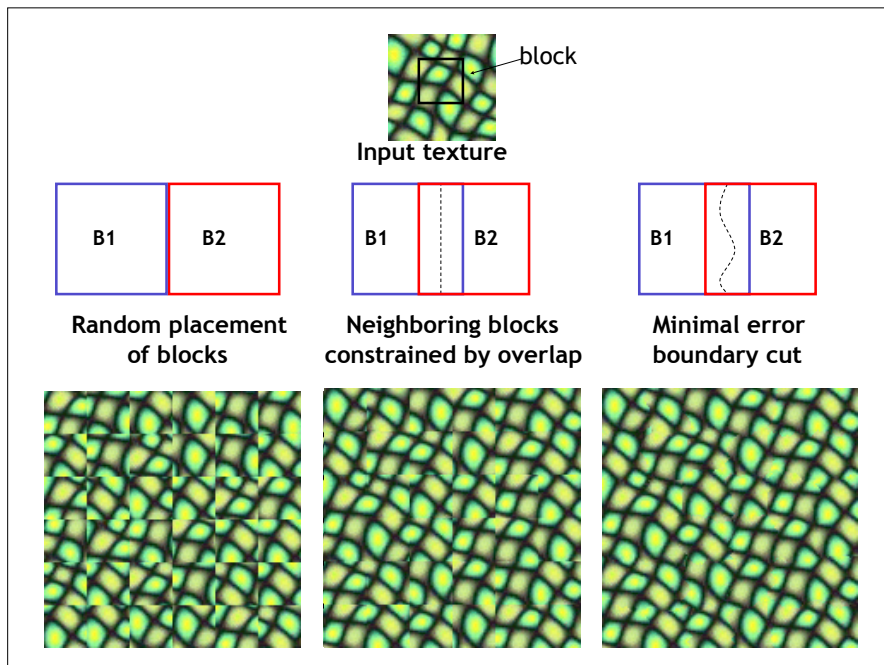
Synthesizing a block

non-parametric sampling






Input image

- 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






29



30



Failures
(Chernobyl Harvest)



31




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

32




(Manual) texture synthesis in the media



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

Style transfer using CNNs

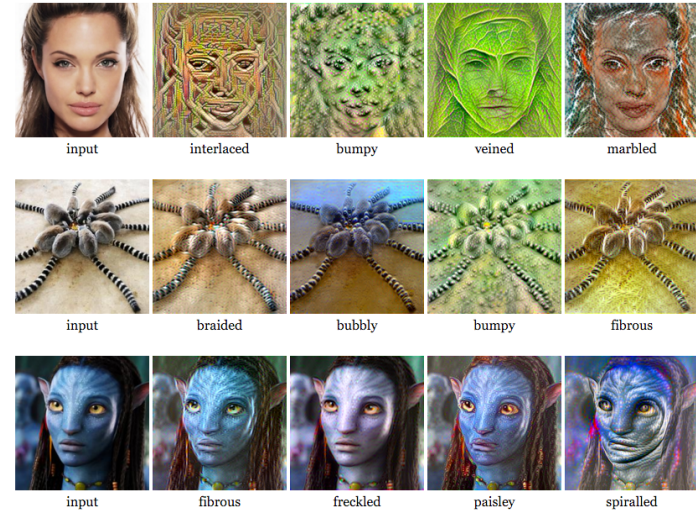


A Neural Algorithm of Artistic Style

Leon A. Gatys,^{1,2,3*} Alexander S. Ecker,^{1,2,4,5} Matthias Bethge^{1,2,4}

37

Style transfer with texture attributes



<http://vis-www.cs.umass.edu/texture/>

Tsung-Yu Lin, Subhransu Maji, CVPR 16

38