Announcements

- Guest lecture next Tuesday
  - **Dan Goldman: CV in special effects**
  - **held in Allen Center (room TBA)**
- Evals at the end of class today

Texture

Today's Reading

- Alexei A. Efros and Thomas K. Leung, "Texture Synthesis by Non-parametric Sampling," Proc. International Conference on Computer Vision (ICCV), 1999.
  - http://www.cs.berkeley.edu/~efros/research/NPS/efros-iccv99.pdf

Modeling Texture

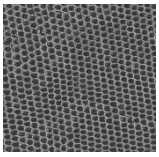

What is texture?

How can we model it?

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

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

- **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:

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

Markov Chain Example: Text

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

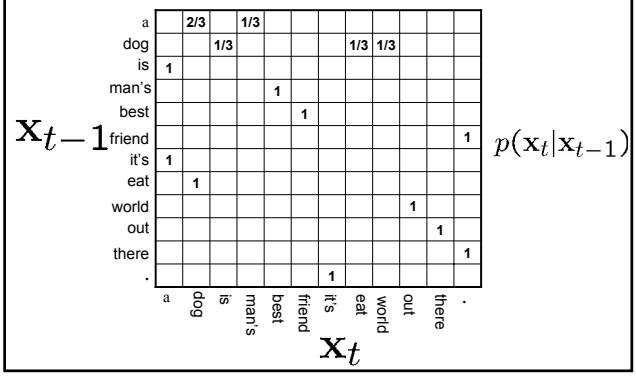

Text synthesis

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 occurance $p(\mathbf{x}_t|\mathbf{x}_{t-1},\ldots,\mathbf{x}_{t-(n-1)})$
2. Given words $\mathbf{x}_1,\mathbf{x}_2,\ldots,\mathbf{x}_{k-1}$
   - compute $\mathbf{x}_k$ by sampling from $p(\mathbf{x}_t|\mathbf{x}_{t-1},\ldots,\mathbf{x}_{t-(n-1)})$

Example on board...

[Scientific American, June 1989, Dewdney]

"I Spent an Interesting Evening Recently with a Grain of Salt"

- Mark V. Shaney

(computer-generated contributor to UseNet News group called net.singles)

Output of 2nd order word-level Markov Chain after training on 90,000 word philosophical essay:

"Perhaps only the allegory of simulation is unendurable--more cruel than Artaud's Theatre of Cruelty, which was the first to practice deterrence, abstraction, disconnection, deterritorialisation, etc.; and if it were our own past. We are witnessing the end of the negative form. But nothing separates one pole from the very swing of voting "rights" to electoral..."

Modeling Texture


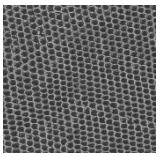
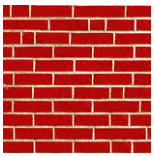

What is texture?

- An image obeying some statistical properties
- Similar structures repeated over and over again
- Often has some degree of randomness

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*:

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

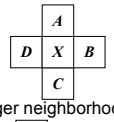

- Higher order MRF's have larger neighborhoods

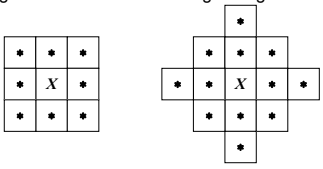

Texture Synthesis [Efros & Leung, ICCV 99]

Can apply 2D version of text synthesis

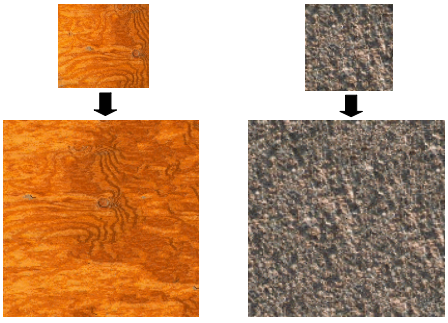

Synthesizing One Pixel

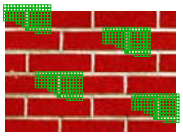

input image

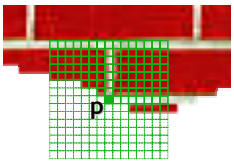

synthesized image

- What is $P(\mathrm{x}|\text{neighborhood of pixels around x})$?
- Find all the windows in the image that match the neighborhood
  - consider only pixels in the neighborhood that are already filled in
- To synthesize **x**
  - pick one matching window at random
  - assign **x** to be the center pixel of that window

Slides courtesy of Alyosha Efros

Really Synthesizing One Pixel

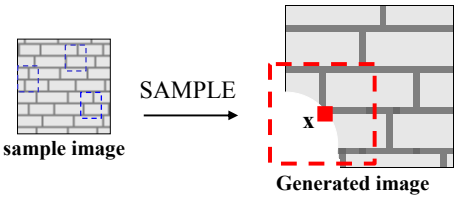

sample image

Generated image

- An exact neighbourhood match might not be present
- So we find the **best** matches using SSD error and randomly choose between them, preferring better matches with higher probability

Growing Texture

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

Window Size Controls Regularity

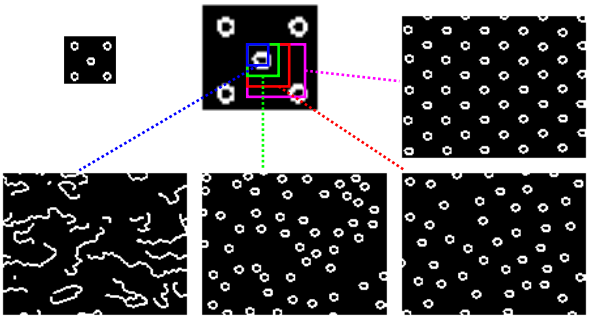

More Synthesis Results

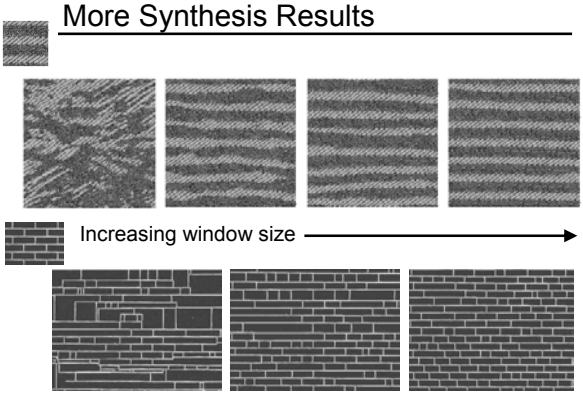

More Results

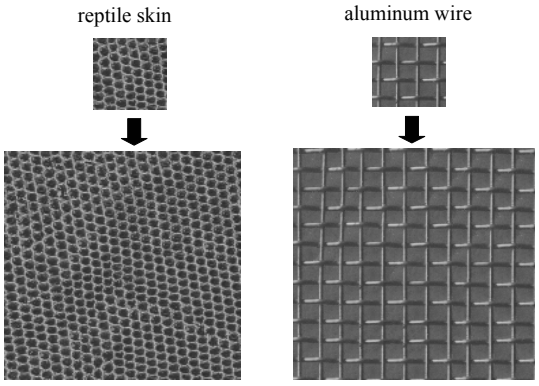

Failure Cases

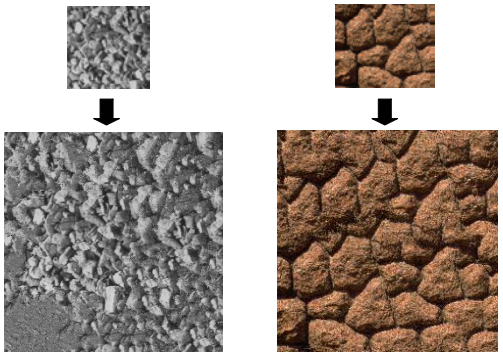

Growing garbage Verbatim copying

Image-Based Text Synthesis

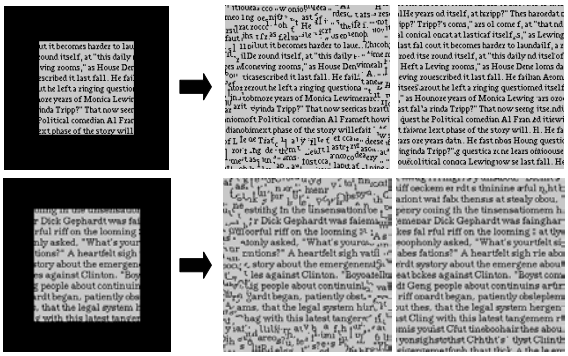

Speed

- Given: image of k^2 pixels
- Output: image of n^2 pixels
- how many window comparisons does this algorithm require?

Block-based texture synthesis

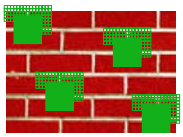

Input image

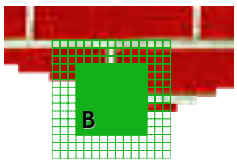

Synthesizing a block

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

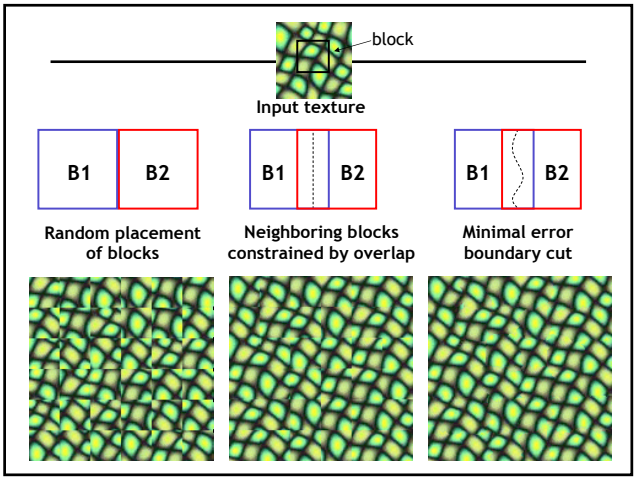

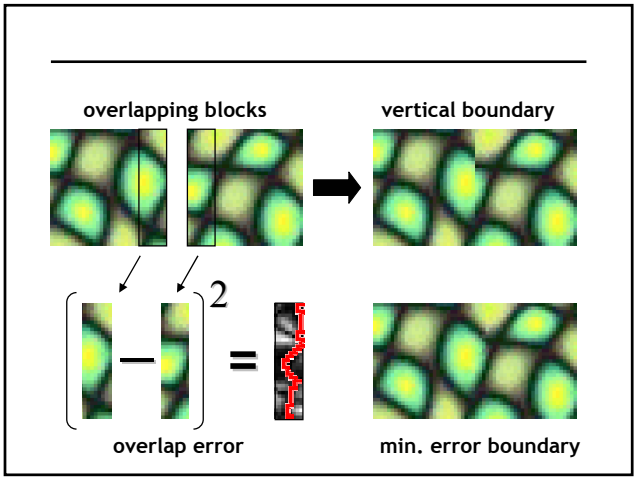

Texture Transfer

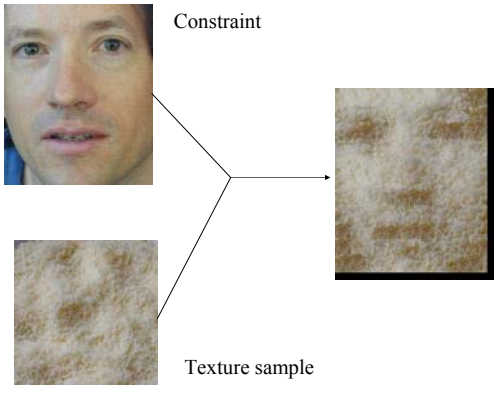

Texture Transfer

Take the texture from one image and "paint" it onto another object

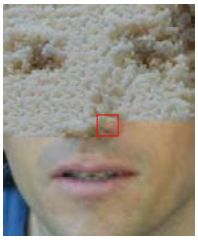

Same algorithm as before with additional term

- do texture synthesis on image1, create new image (size of image2)
- add term to match intensity of image2

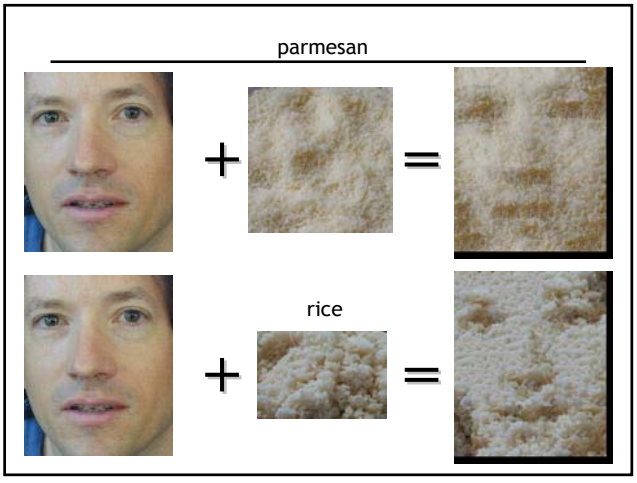

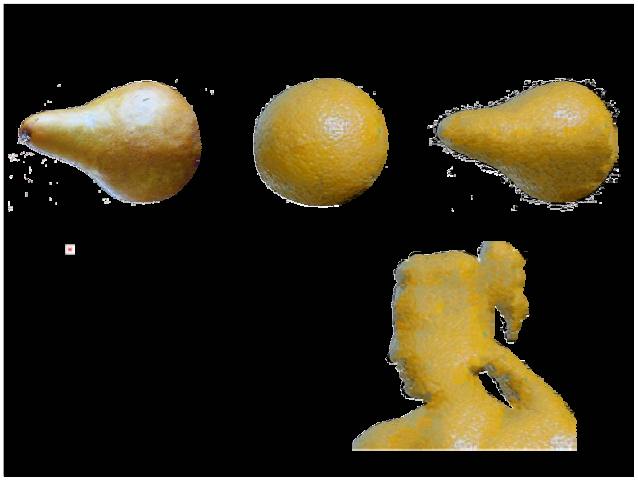

Combining two images

Graph cut setup

Graph cut texture synthesis: Video

Image Analogies (Hertzmann '01)


A


A'


B


B'

Artistic Filters


A


A'


B


B'

Texture-by-numbers


A


A’

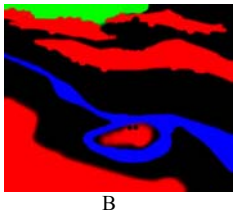
B

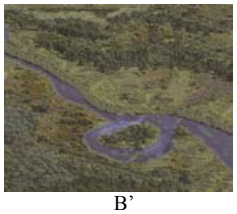
B’

Other applications of Image Analogies

- Texture synthesis
- Super-resolution
- Texture transfer
- Image colorization
- Simple filters (blur, emboss)
- More details: Hertzmann et al., SIGGRAPH 2001
  - http://mrl.nyu.edu/projects/image-analogies/

Applications of Texture Modeling

Super-resolution

- Freeman & Pasztor, 1999
- Baker & Kanade, 2000

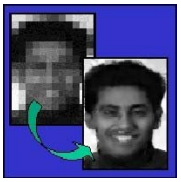

Image/video compression

Texture recognition, segmentation

- DeBonet

Restoration

- removing scratches, holes, filtering
- Zhu et al.

Art/entertainment