CMPSCI 670: Computer Vision

Texture continued …

University of Massachusetts, Amherst

October 8, 2014

Instructor: Subhransu Maji

Slides credit: Kristen Grauman and others

Administrivia

- Homework 1 grade posted:
  - you should have received an email.
  - questions? email me and I will resolve it with the graders.
- Today's office hours are cancelled
  - Instead having it tomorrow, i.e., Th 3:45 - 4:45 PM
  - DLS speaker Richard Sutton (go to his talk instead)

Distinguished Lecturer Series

Richard Sutton
University of Alberta
Department of Computing Science

Wednesday, October 8, 2014
4:00pm - 5:00pm
Computer Science Building, Room 151
Faculty Host: Andrew Barto

"Temporal-difference Learning and the Coming of Artificial Intelligence"

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

Recap: Filter bank

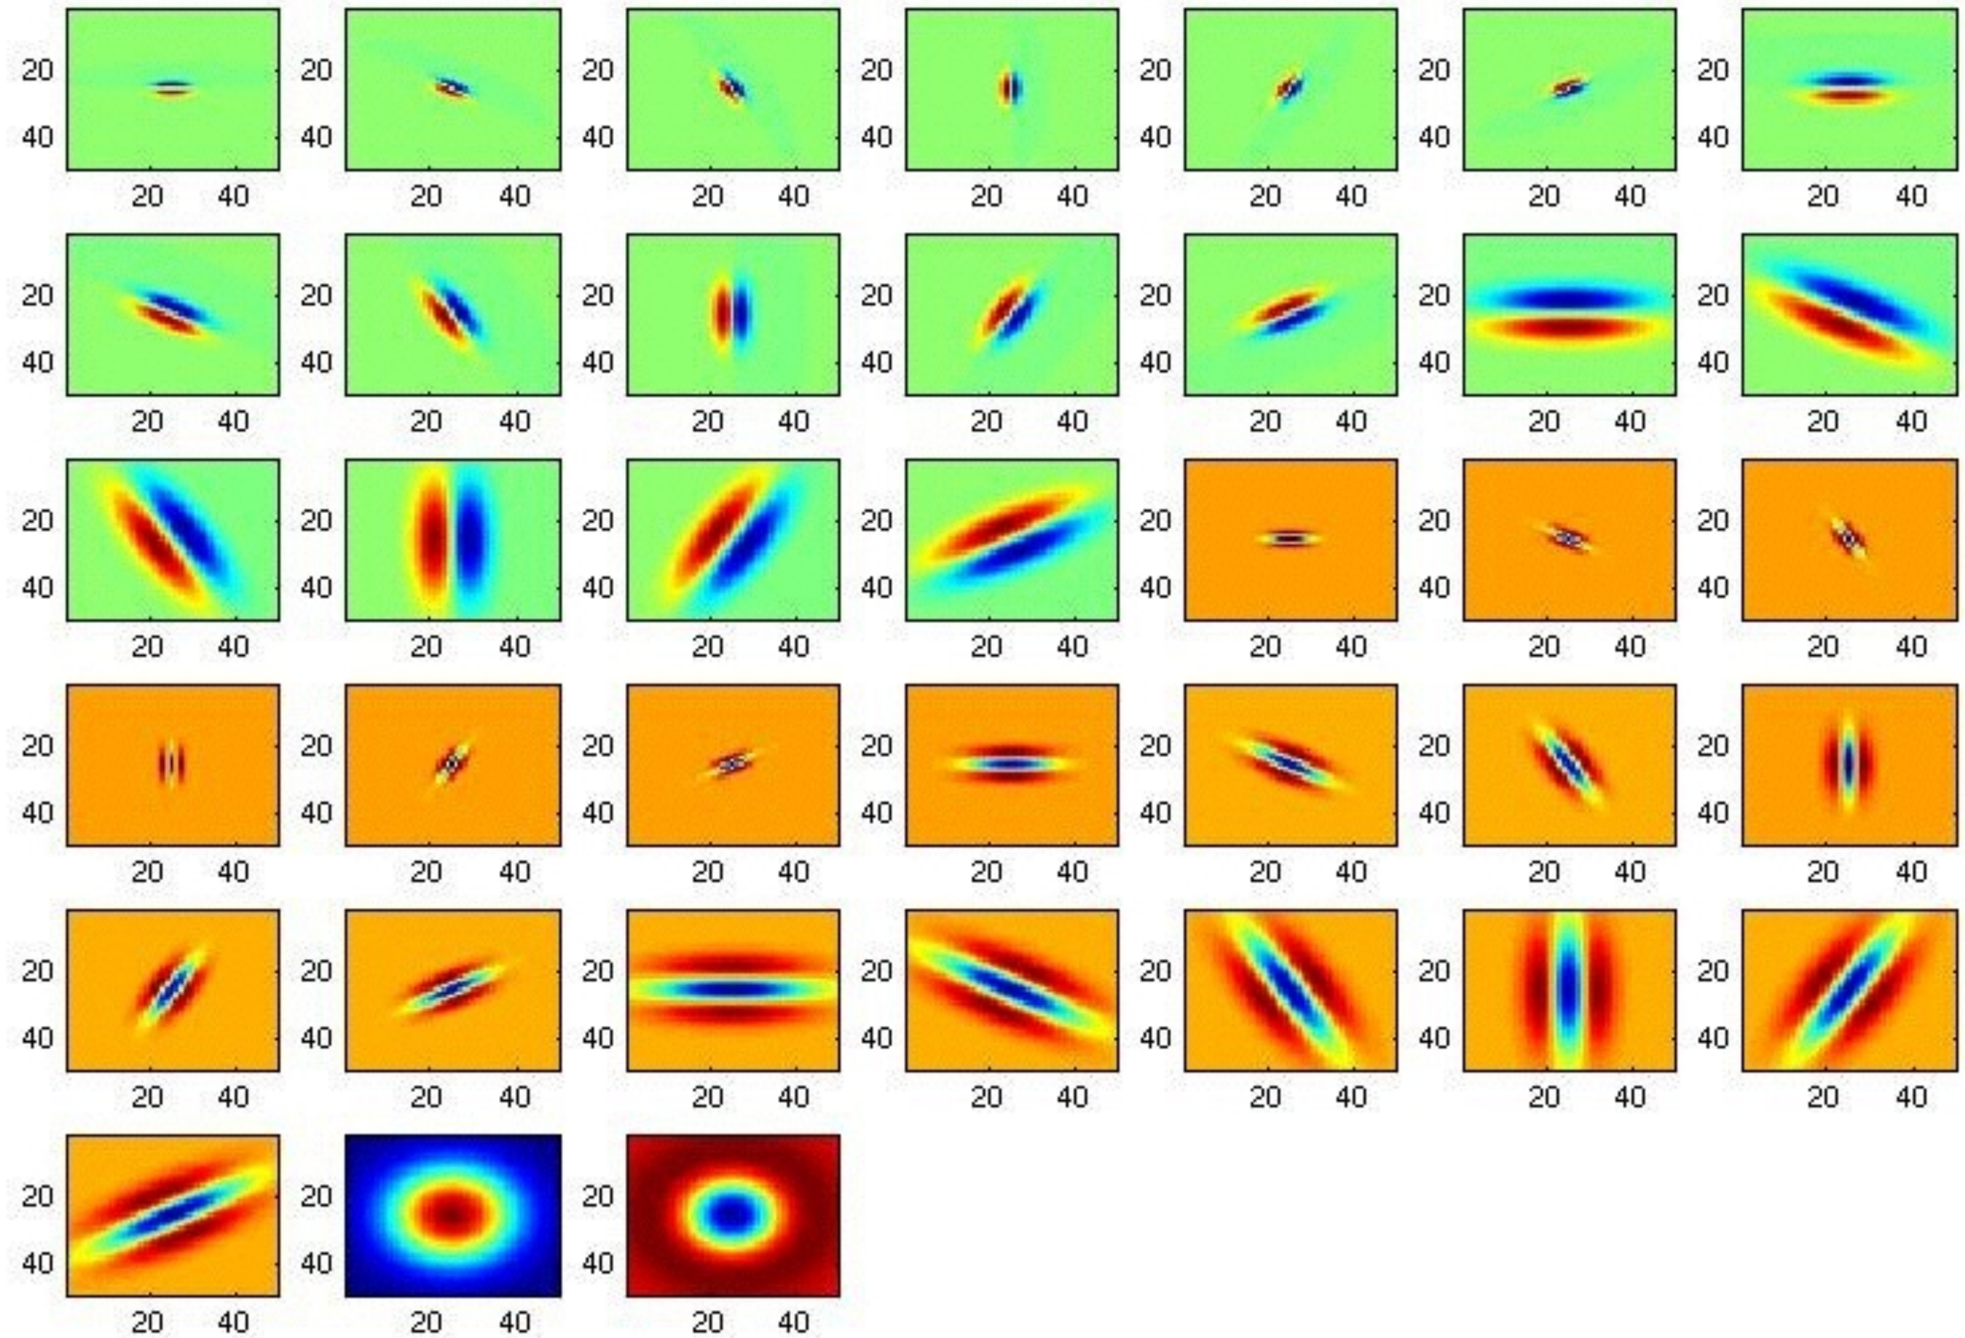

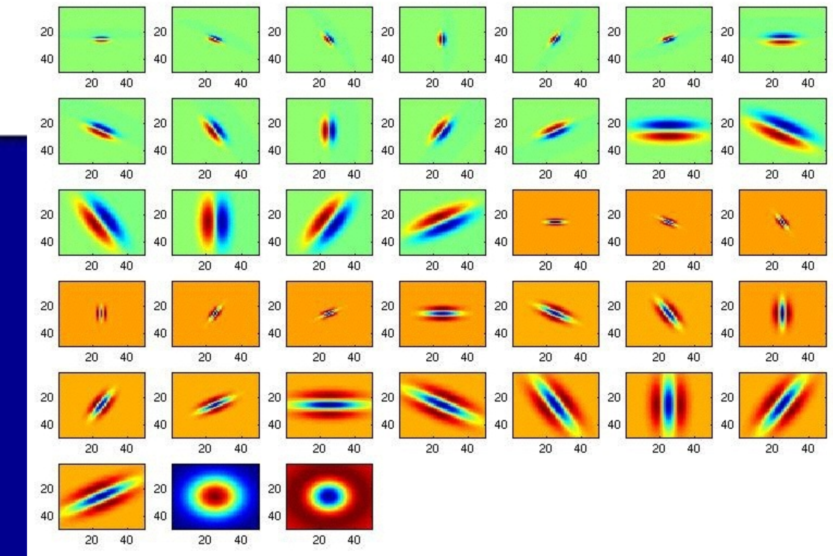

[r1, r2, …, r38]

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

K-means for vector quantization

Given a set of observations **($\mathbf{x}_1$, $\mathbf{x}_2$, …, $\mathbf{x}_n$)**, where each observation is a d-dimensional real vector, k-means clustering aims to partition the **n** observations into **k (≤ n)** sets **S =** *{S_1, S_2, …, S_k}* so as to minimize the within-cluster sum of squares **(WCSS)**. In other words, its objective is to find:

$$\underset{\mathbf{S}}{\arg\min} \sum_{i=1}^{k} \sum_{\mathbf{x} \in S_i} \left\| \mathbf{x} - \boldsymbol{\mu}_i \right\|^2$$

where $\boldsymbol{\mu}_i$ is the mean of points in S_i.

Easy to compute $\boldsymbol{\mu}$ given **S** and vice versa.

http://en.wikipedia.org/wiki/K-means_clustering

Lloyd's algorithm for k-means

- Initialize k centers by picking k-points randomly
- Repeat till convergence (or max iterations)
  - Assign each point to the nearest center (assignment step)
  - Estimate the mean of each group (update step)

MATLAB [idx, c] = **kmeans**(X, k)

- Simple, fast and works well in practice
- But can be unstable
  - Run multiple times and the best solution (one with the smallest WCSS)
  - Better initializations are possible (e.g. kmeans++)

Textons in images

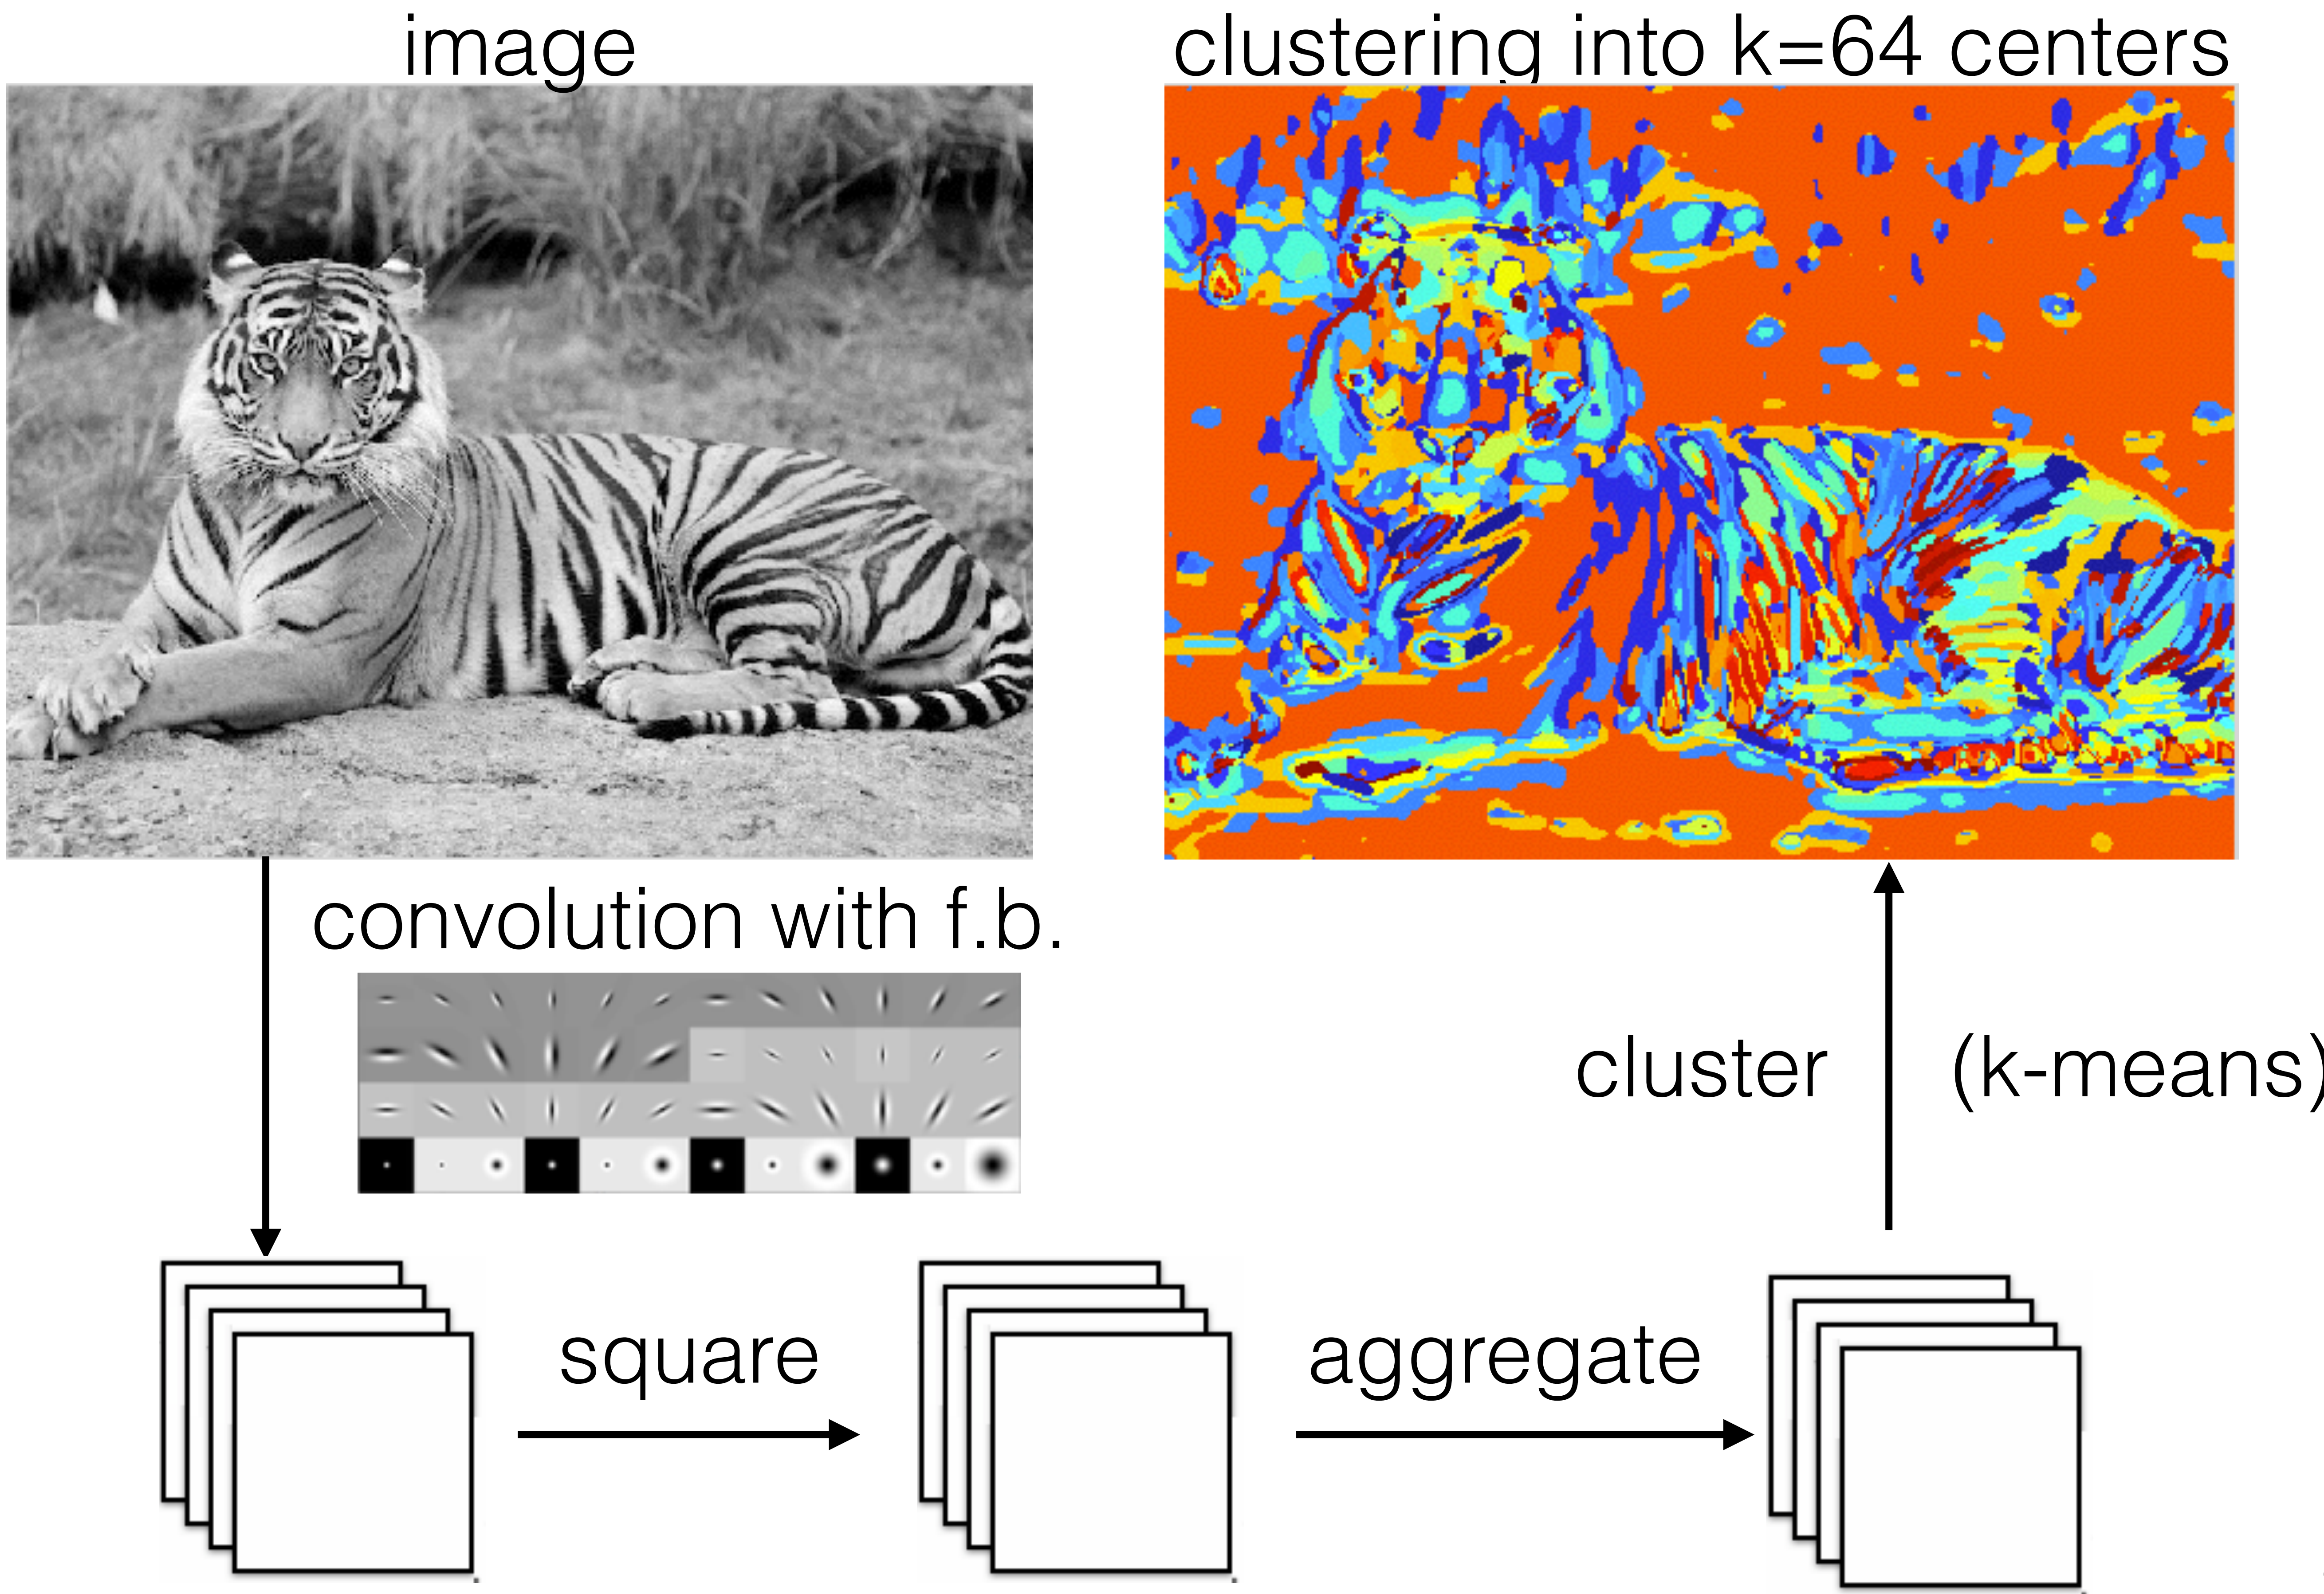

Classifying materials, “stuff”

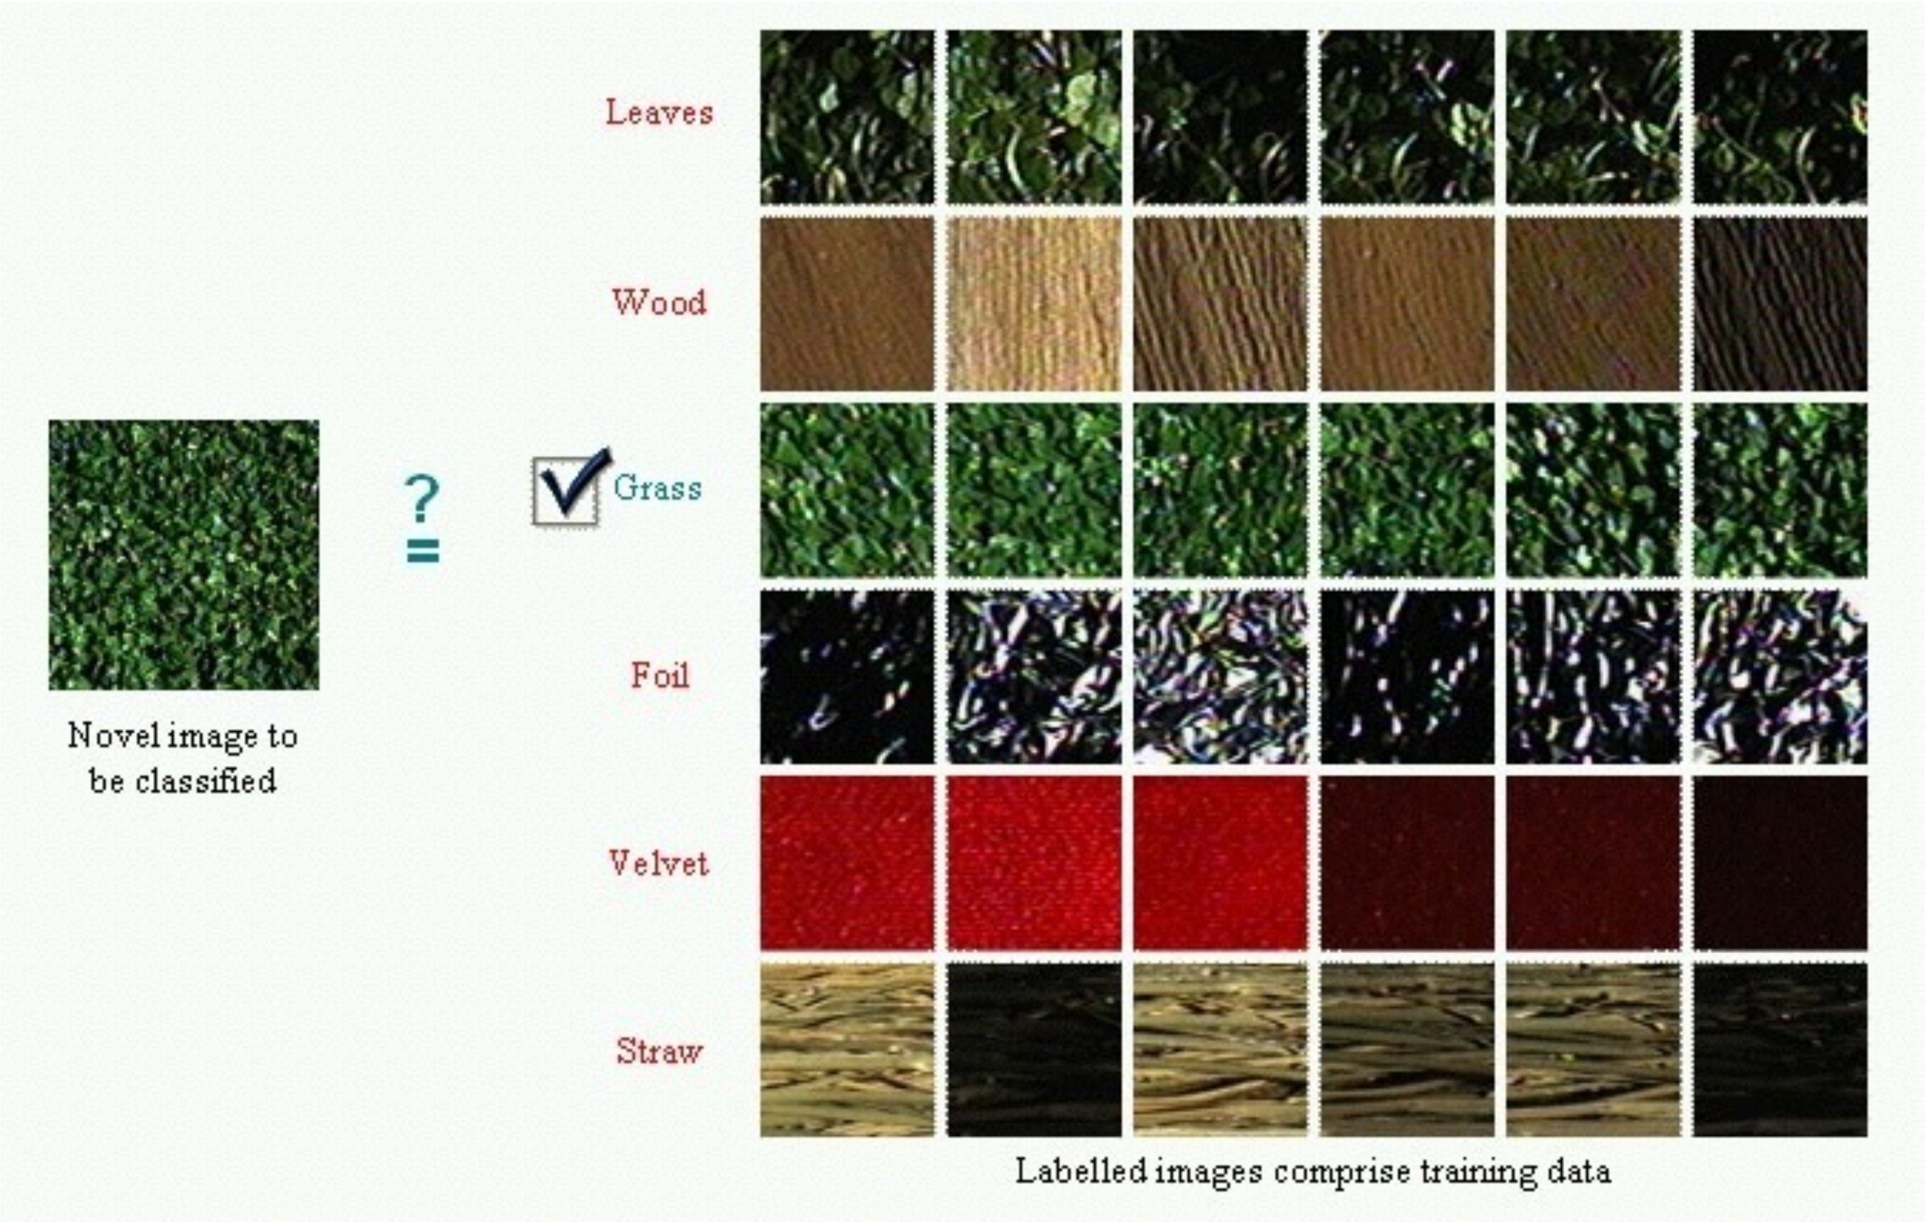

Figure by Varma & Zisserman

Global texton histogram is a good representation

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

Texture synthesis

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

The challenge

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

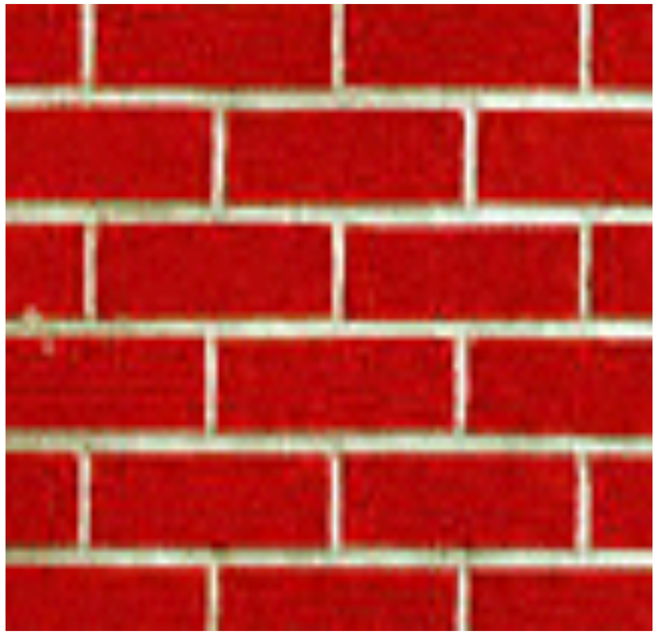

repeated

stochastic

Both?

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

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

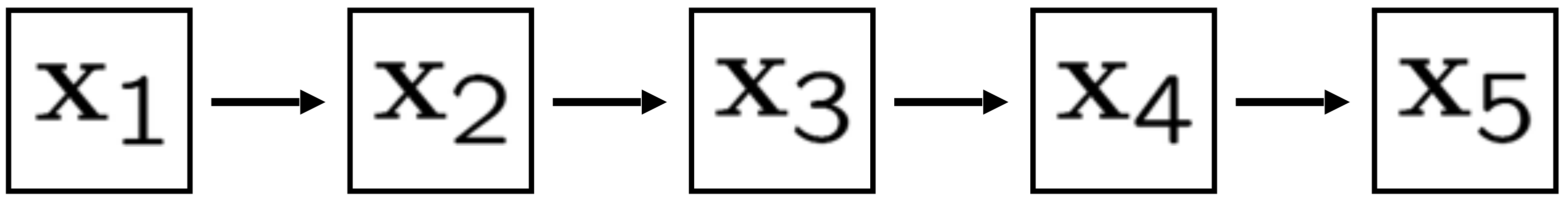

- **Markov assumption**: each state is dependent only on the previous one
  - dependency given by a **conditional probability**:

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

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

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

Source: S. Seitz

Markov Chain Example: Text

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

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

$\mathbf{x}_{t-1}$ \ $\mathbf{x}_t$	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					

Source: S. Seitz

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

WE NEED TO EAT CAKE

Source: S. Seitz

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.

Synthesizing Computer Vision text

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

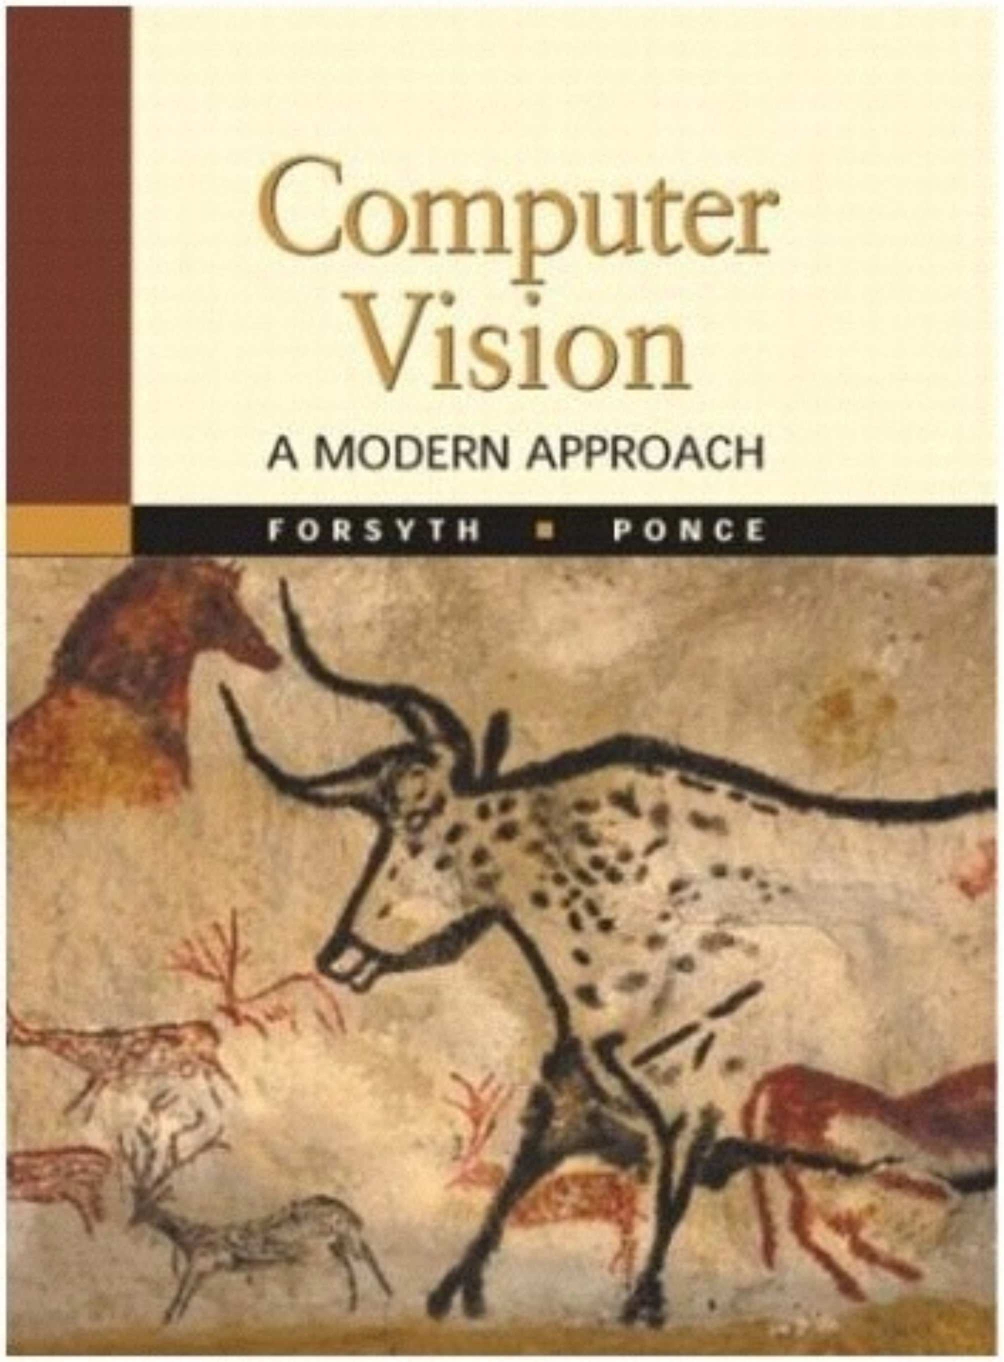

Check out Yisong Yue's website implementing text generation: build your own text Markov Chain for a given text corpus. http://www.yisongyue.com/shaney/index.php

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 × 2k +1 × 2k +1 × 2k + 1 — required for the images separately.

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*:

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

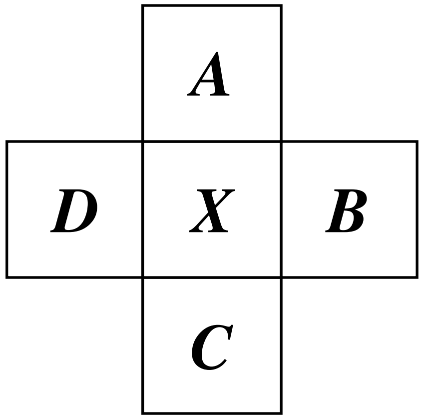

Source: S. Seitz

Texture synthesis

Can apply 2D version of text synthesis

Texture corpus
(sample)

Output

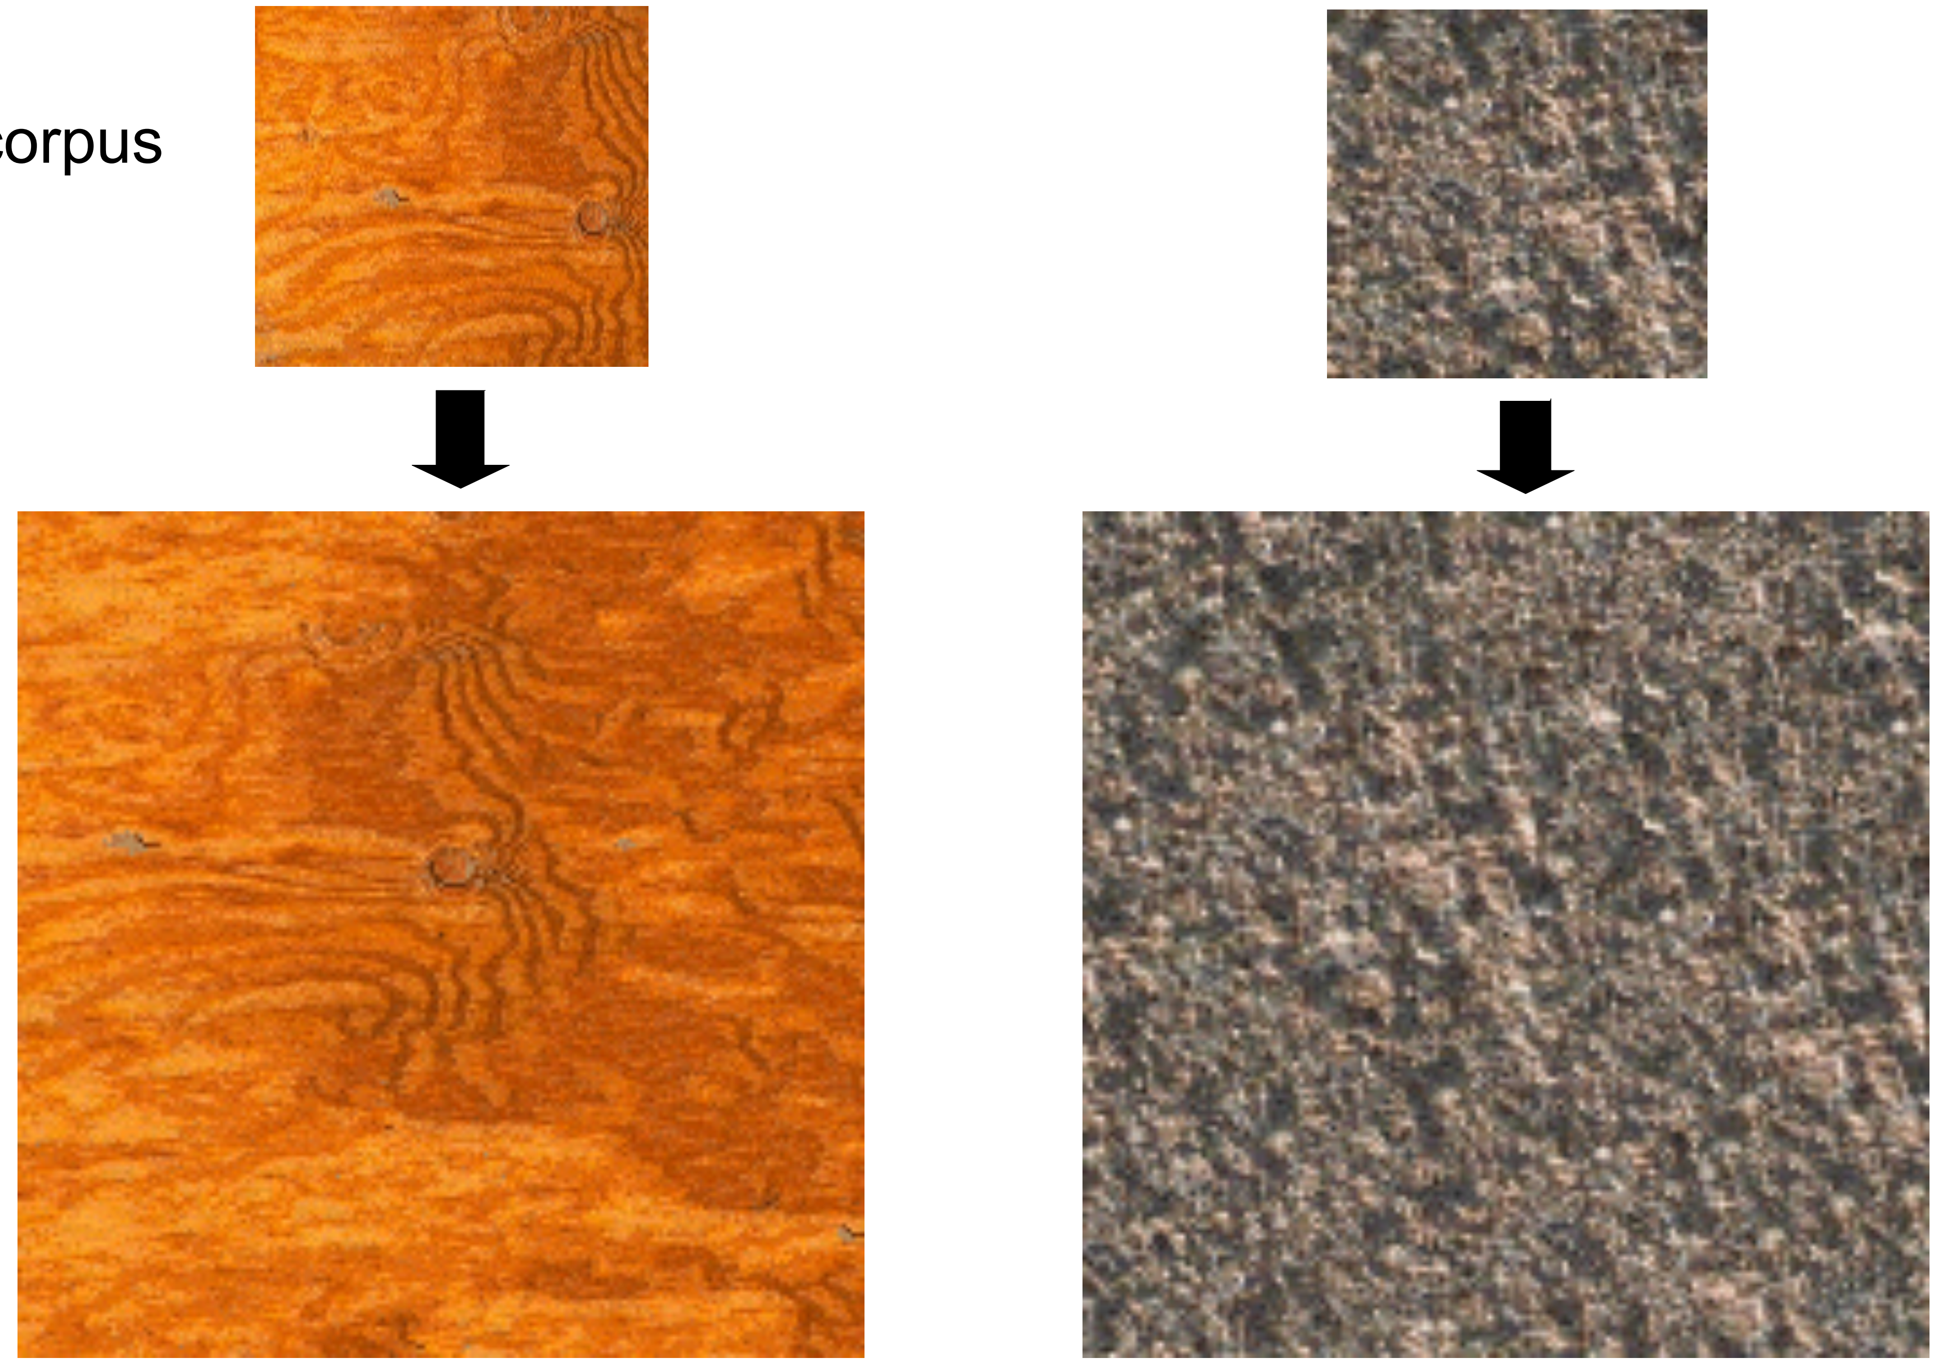

[**Efros & Leung, ICCV 99**]

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.

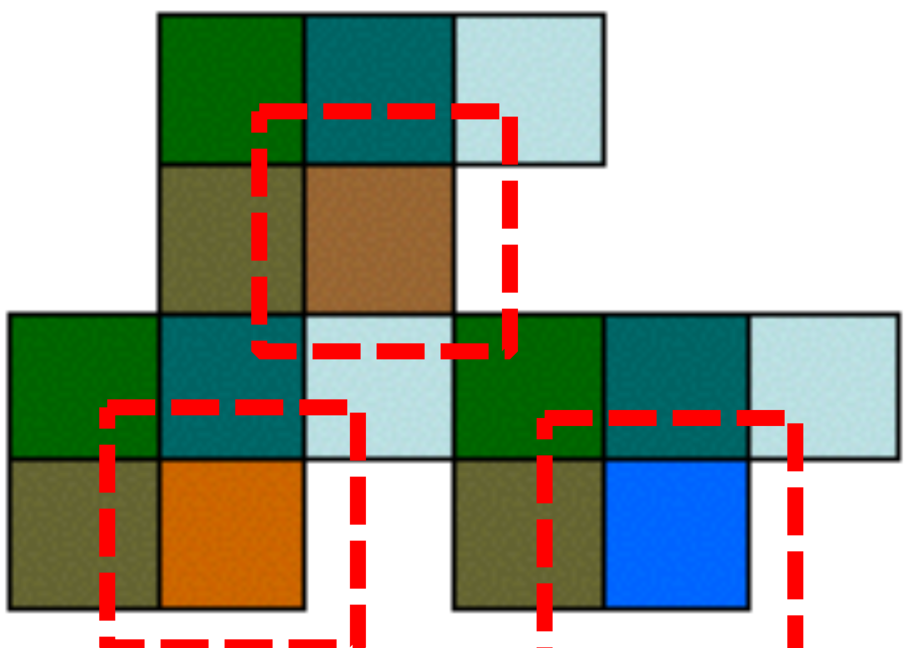

Sample of the texture (“corpus”)

Place we want to insert next

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

Synthesizing one pixel

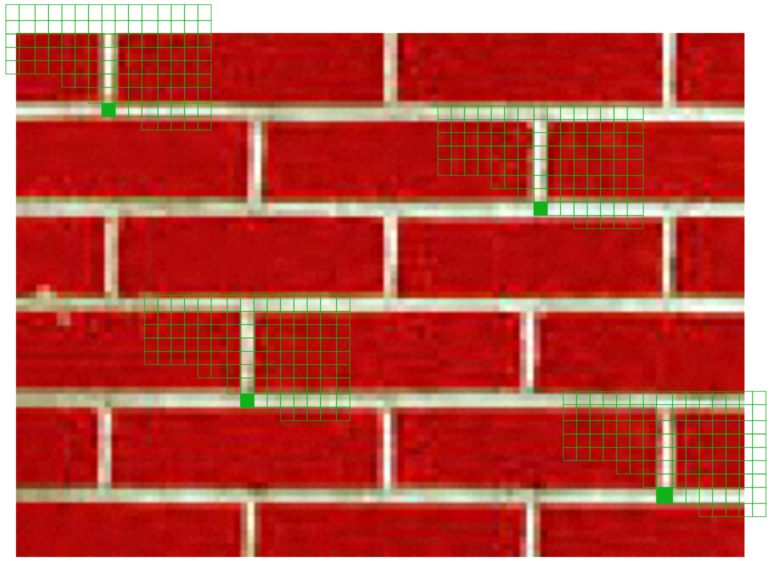

input image

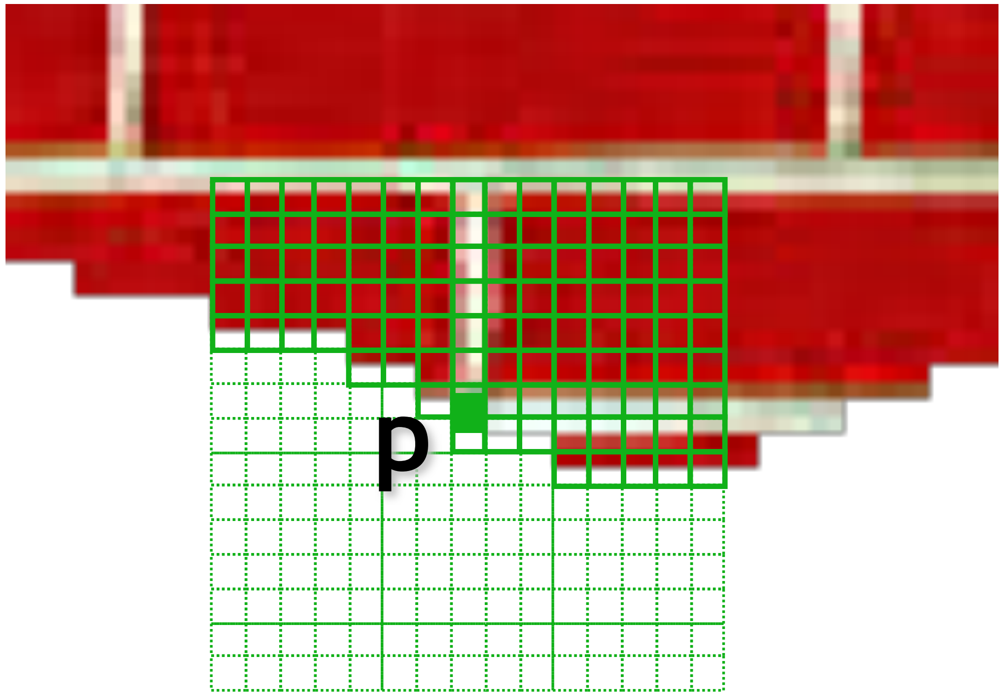

synthesized image

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

Neighborhood window

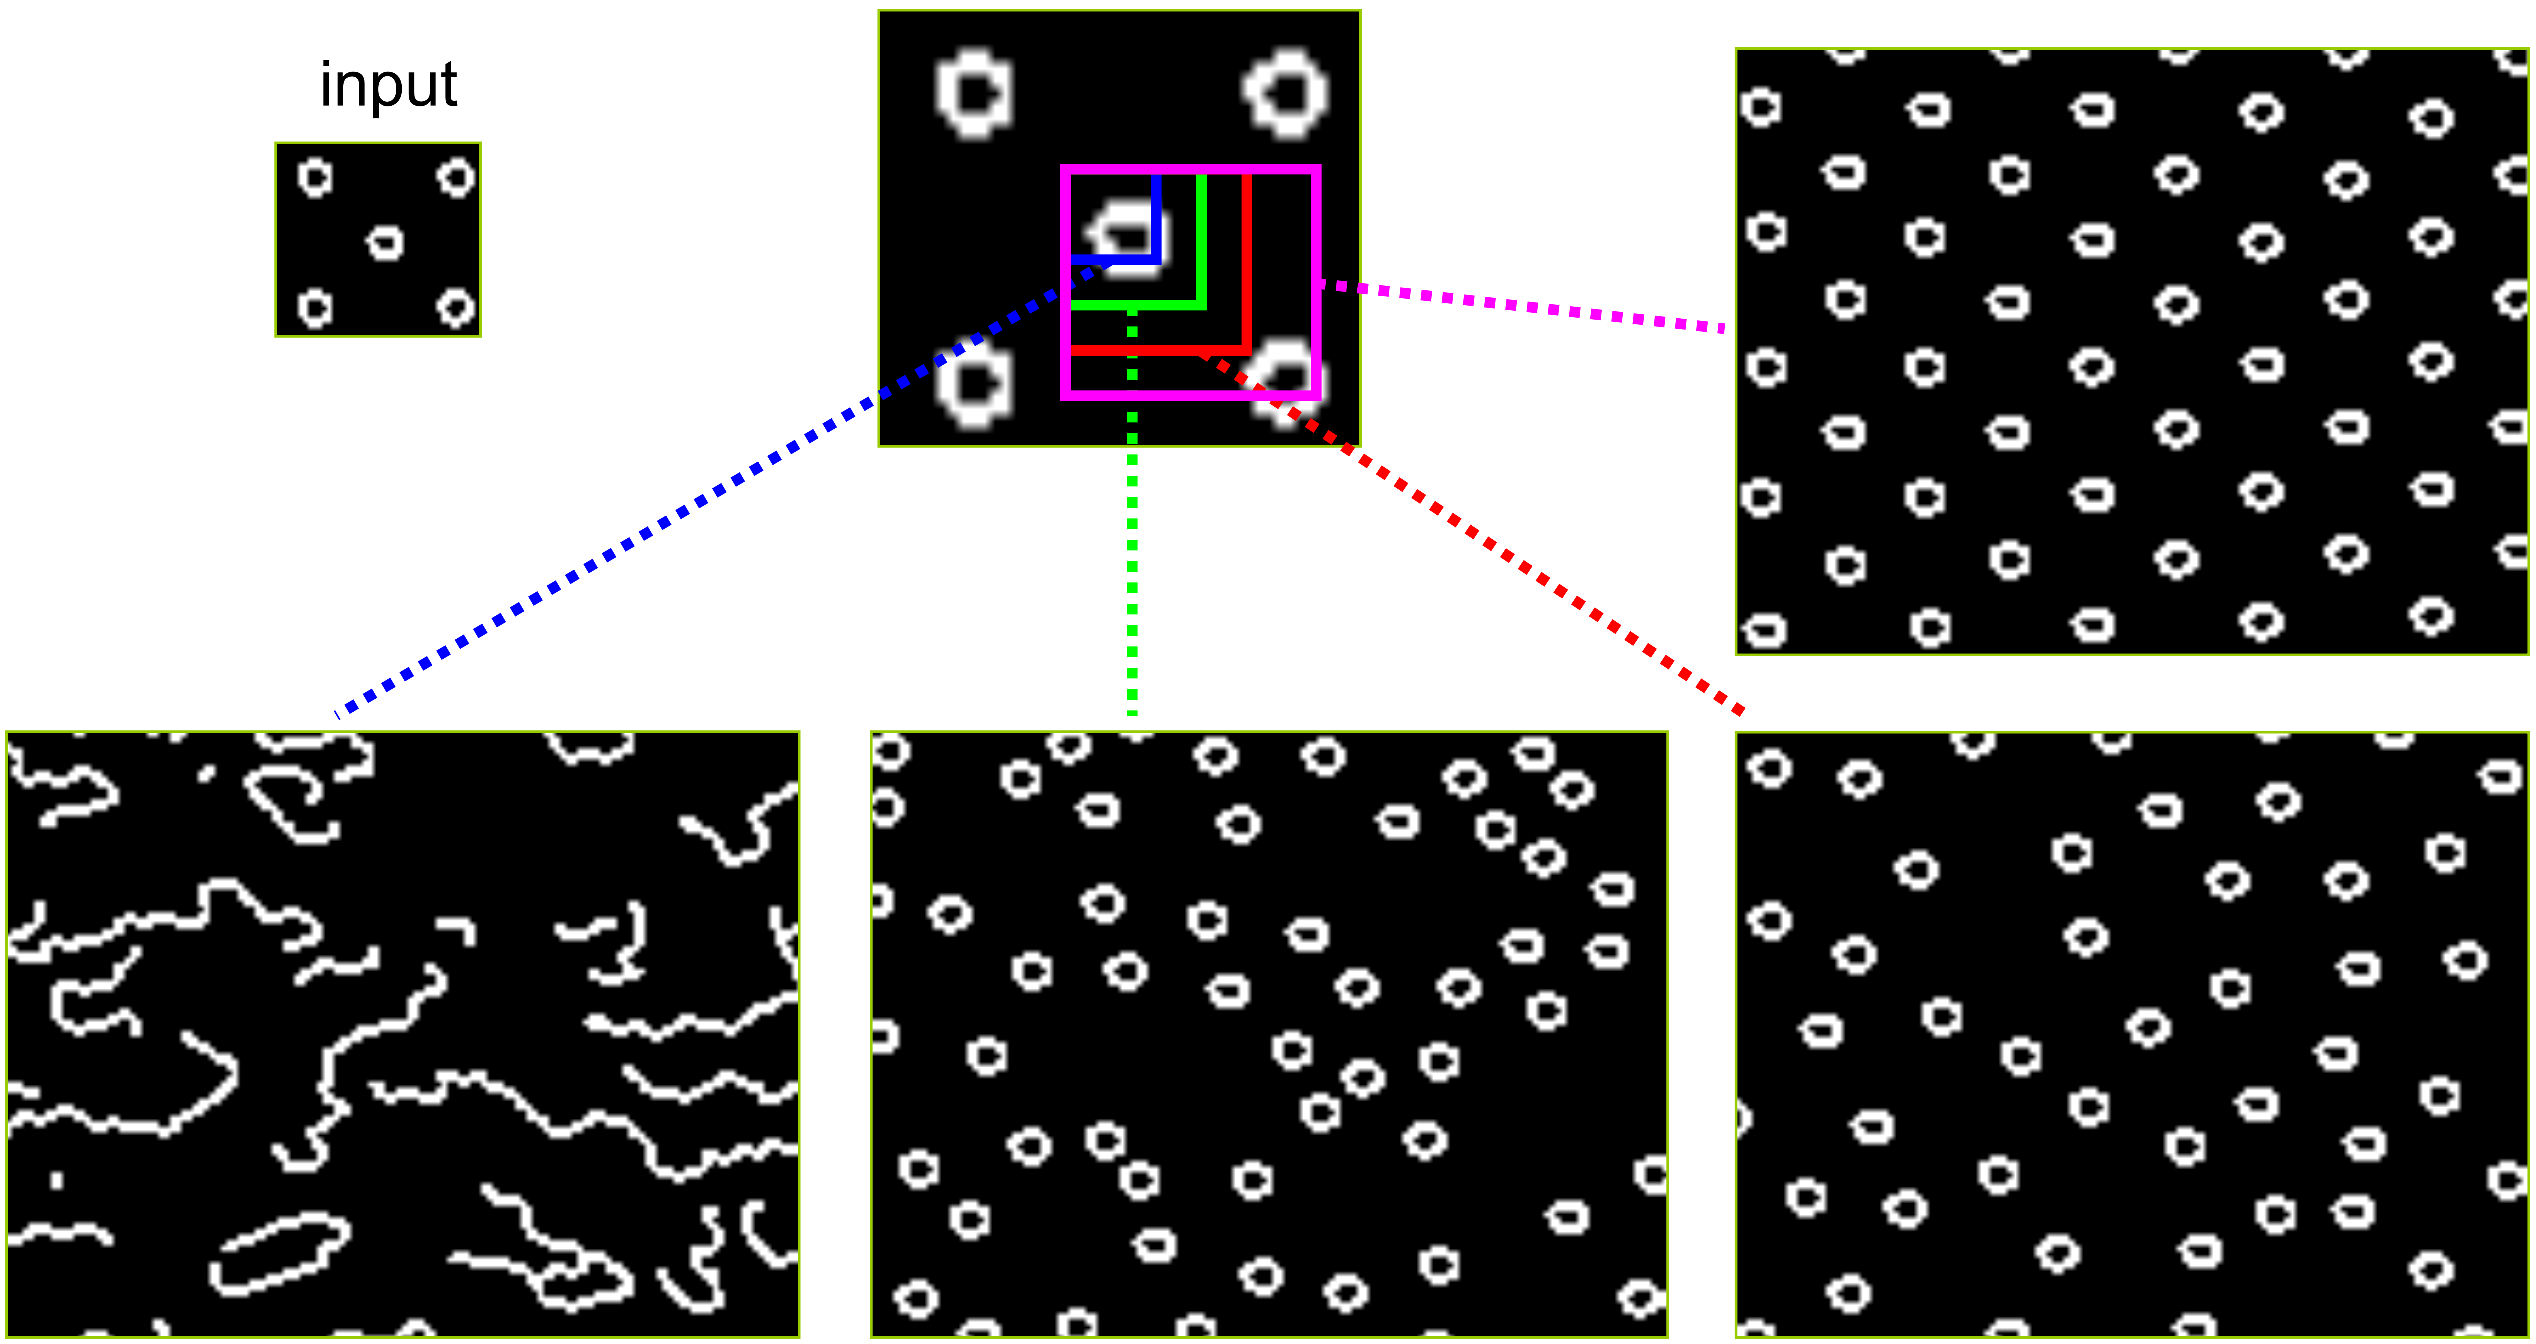

Slide from Alyosha Efros, ICCV 1999

Varying window size

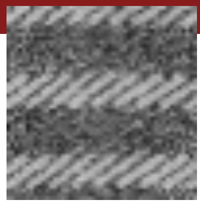

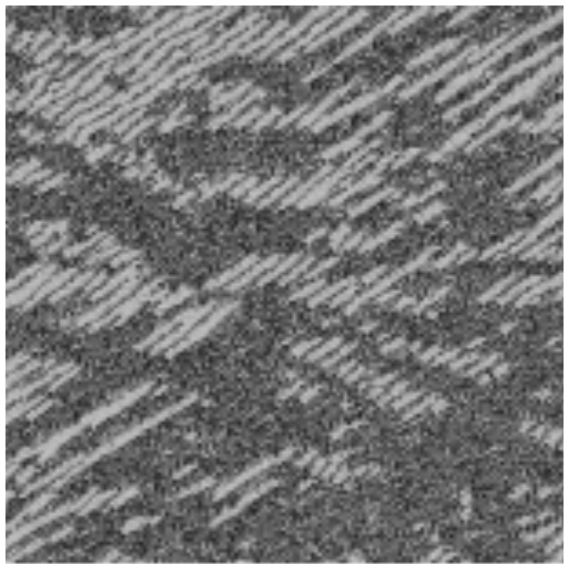

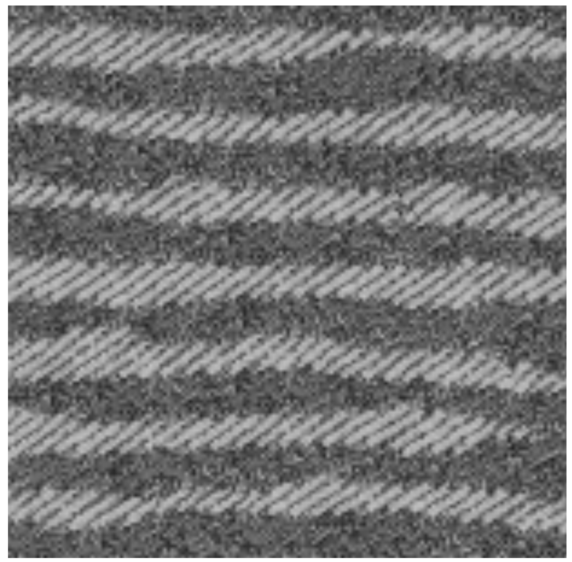

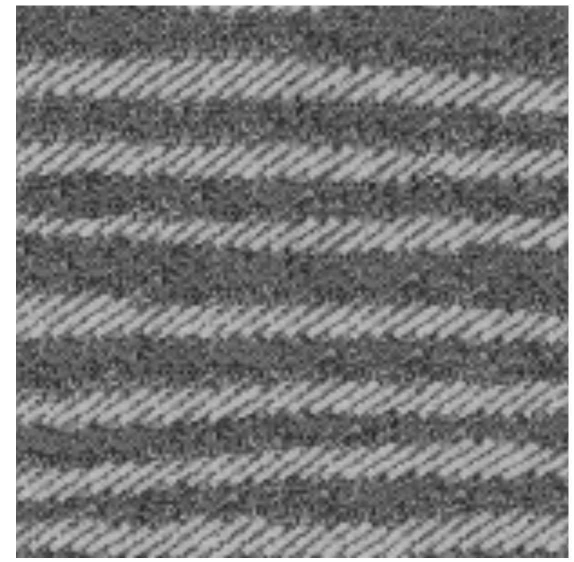

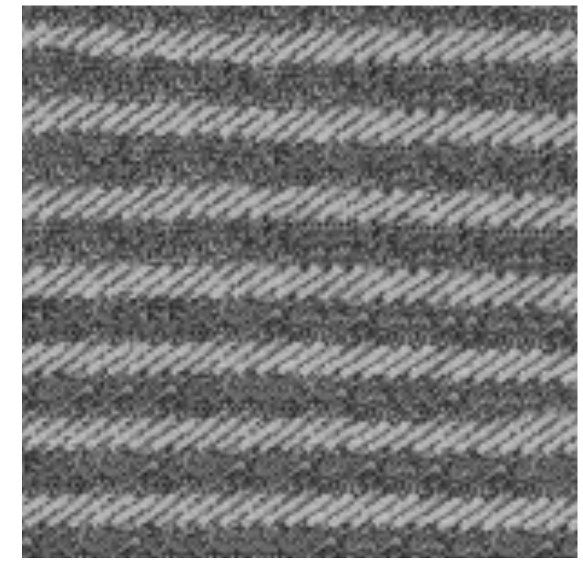

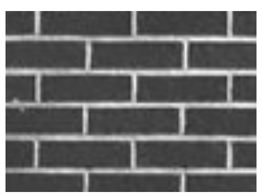

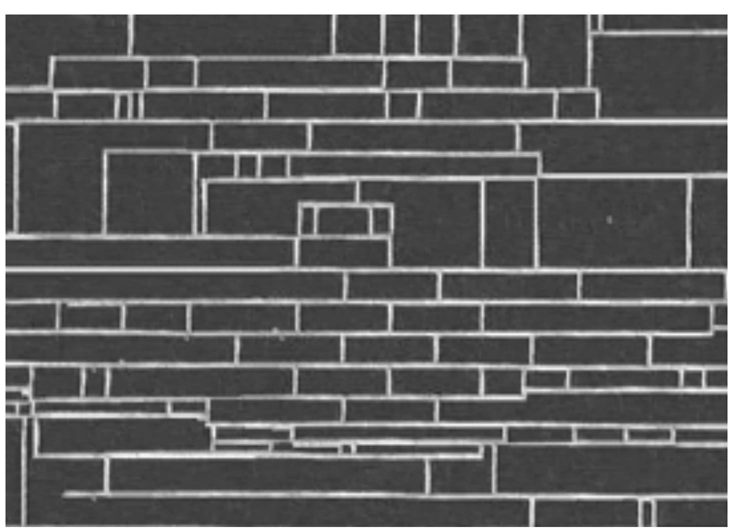

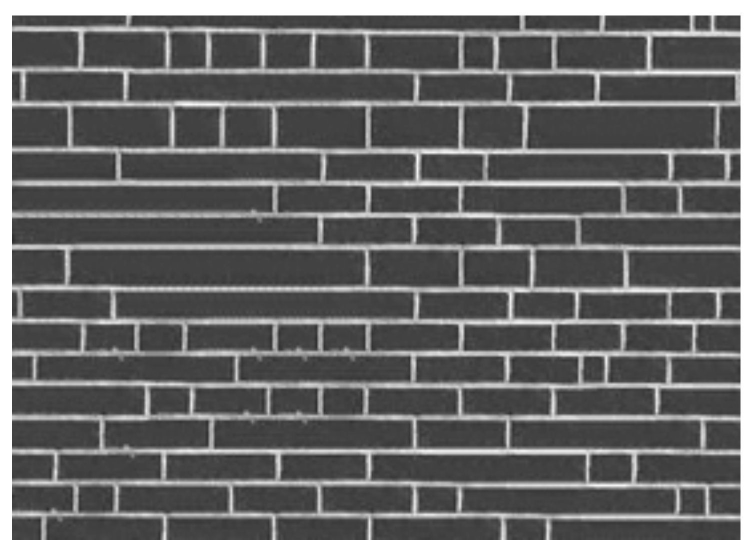

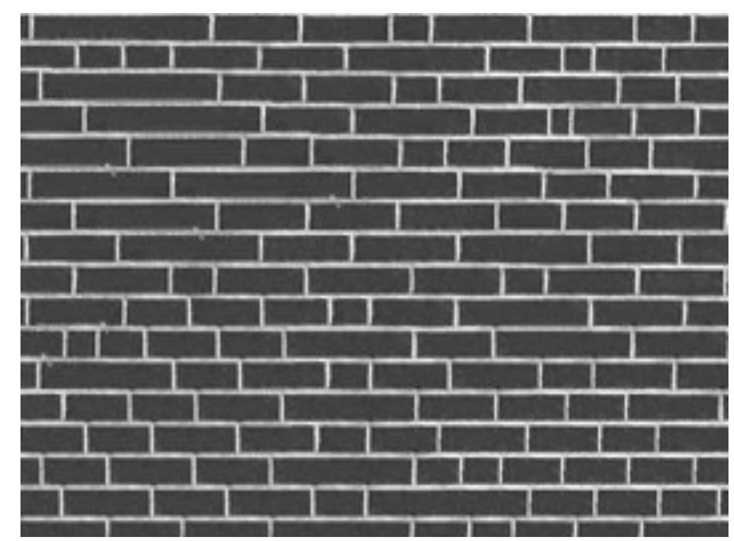

Increasing window size

Slide from Alyosha Efros, ICCV 1999

Growing Texture

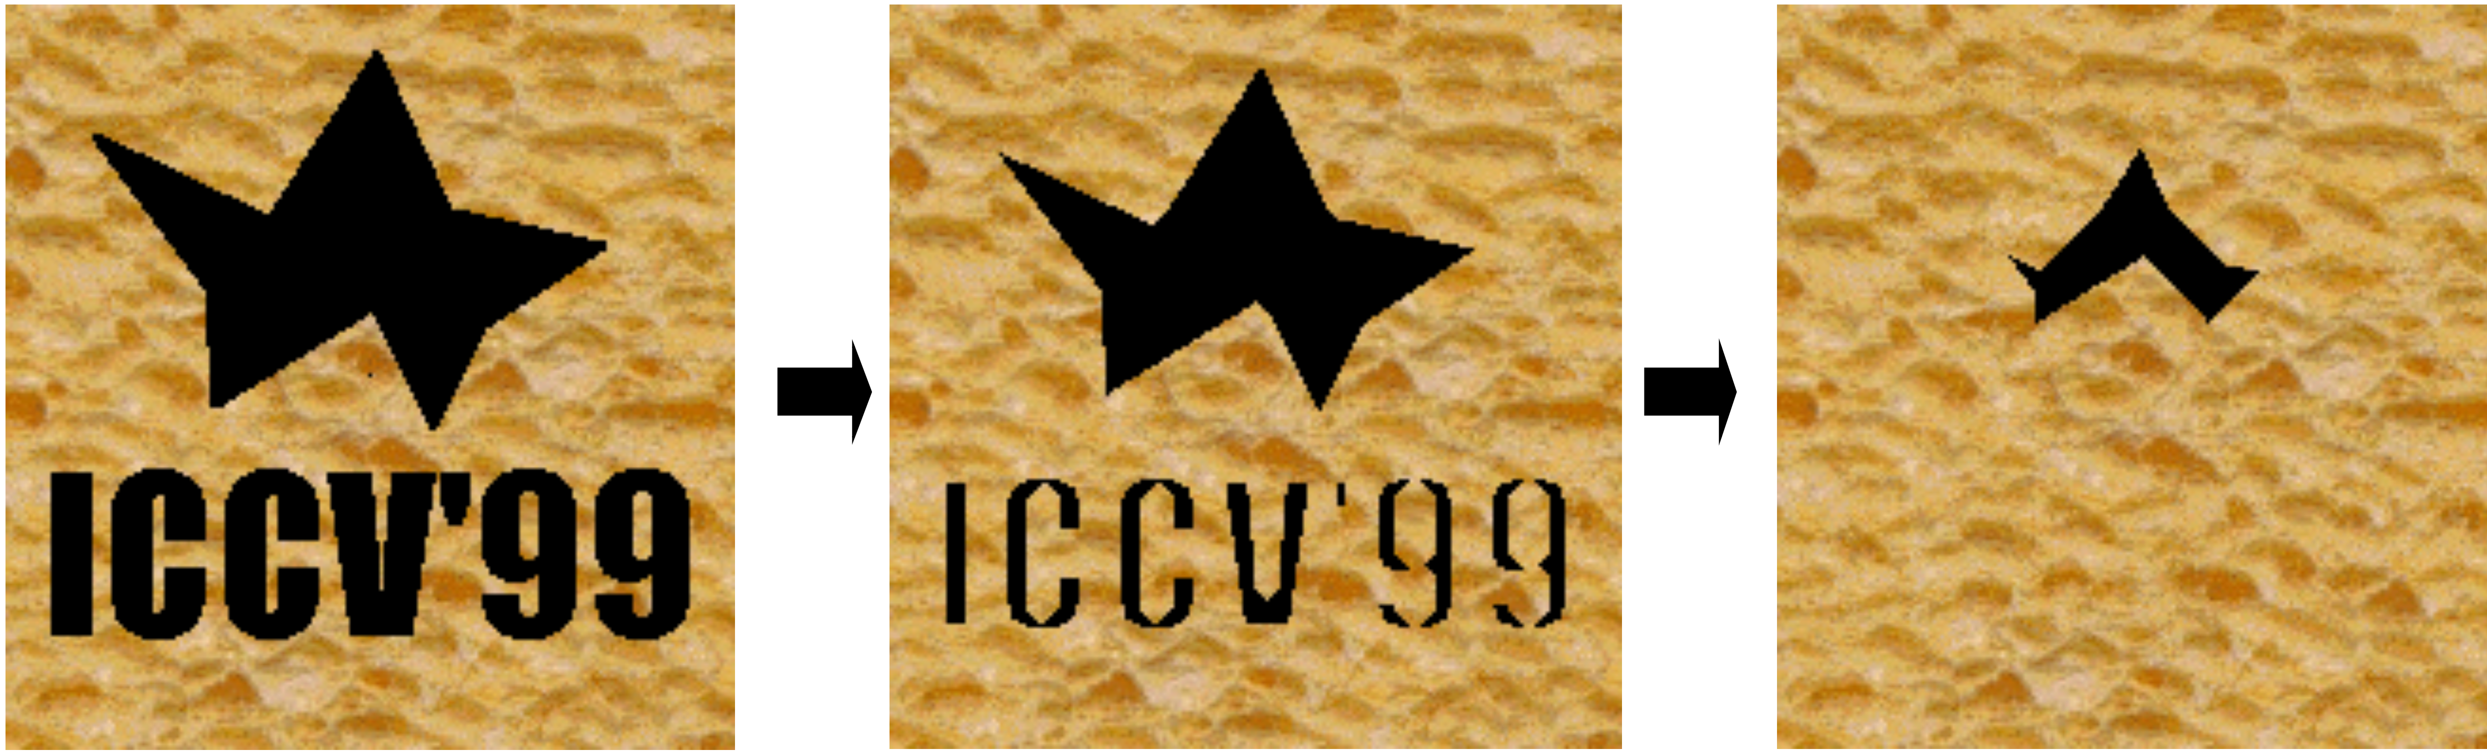

- Starting from the initial image, “grow” the texture one pixel at a time

Synthesis results

french canvas

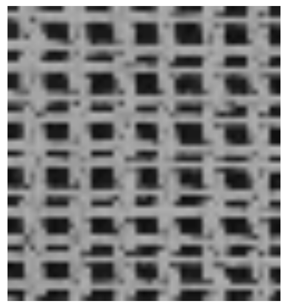

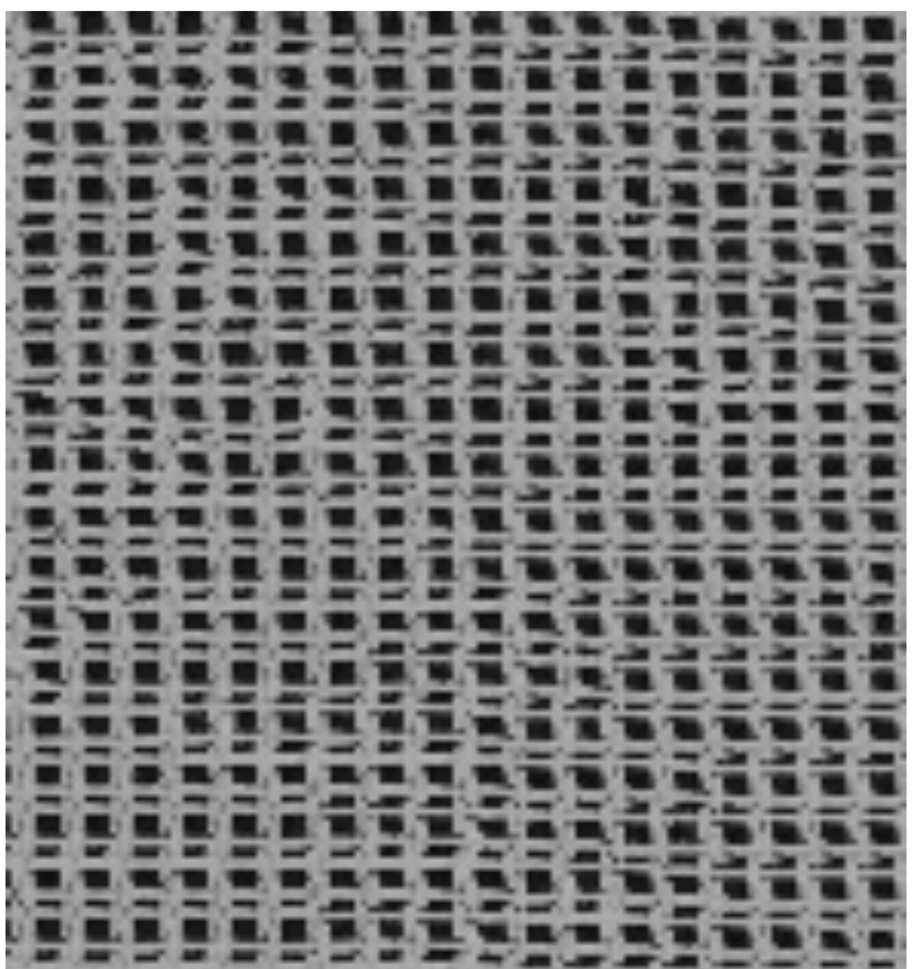

rafia weave

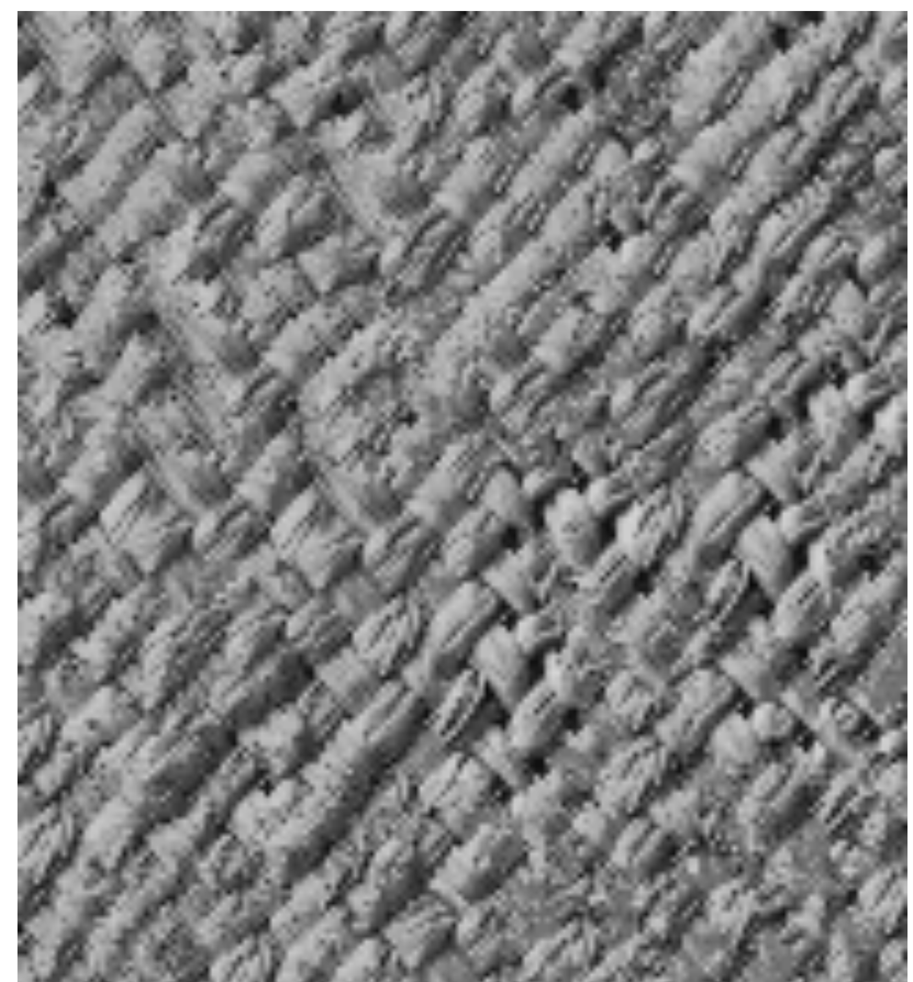

Slide from Alyosha Efros, ICCV 1999

Synthesis results

white bread

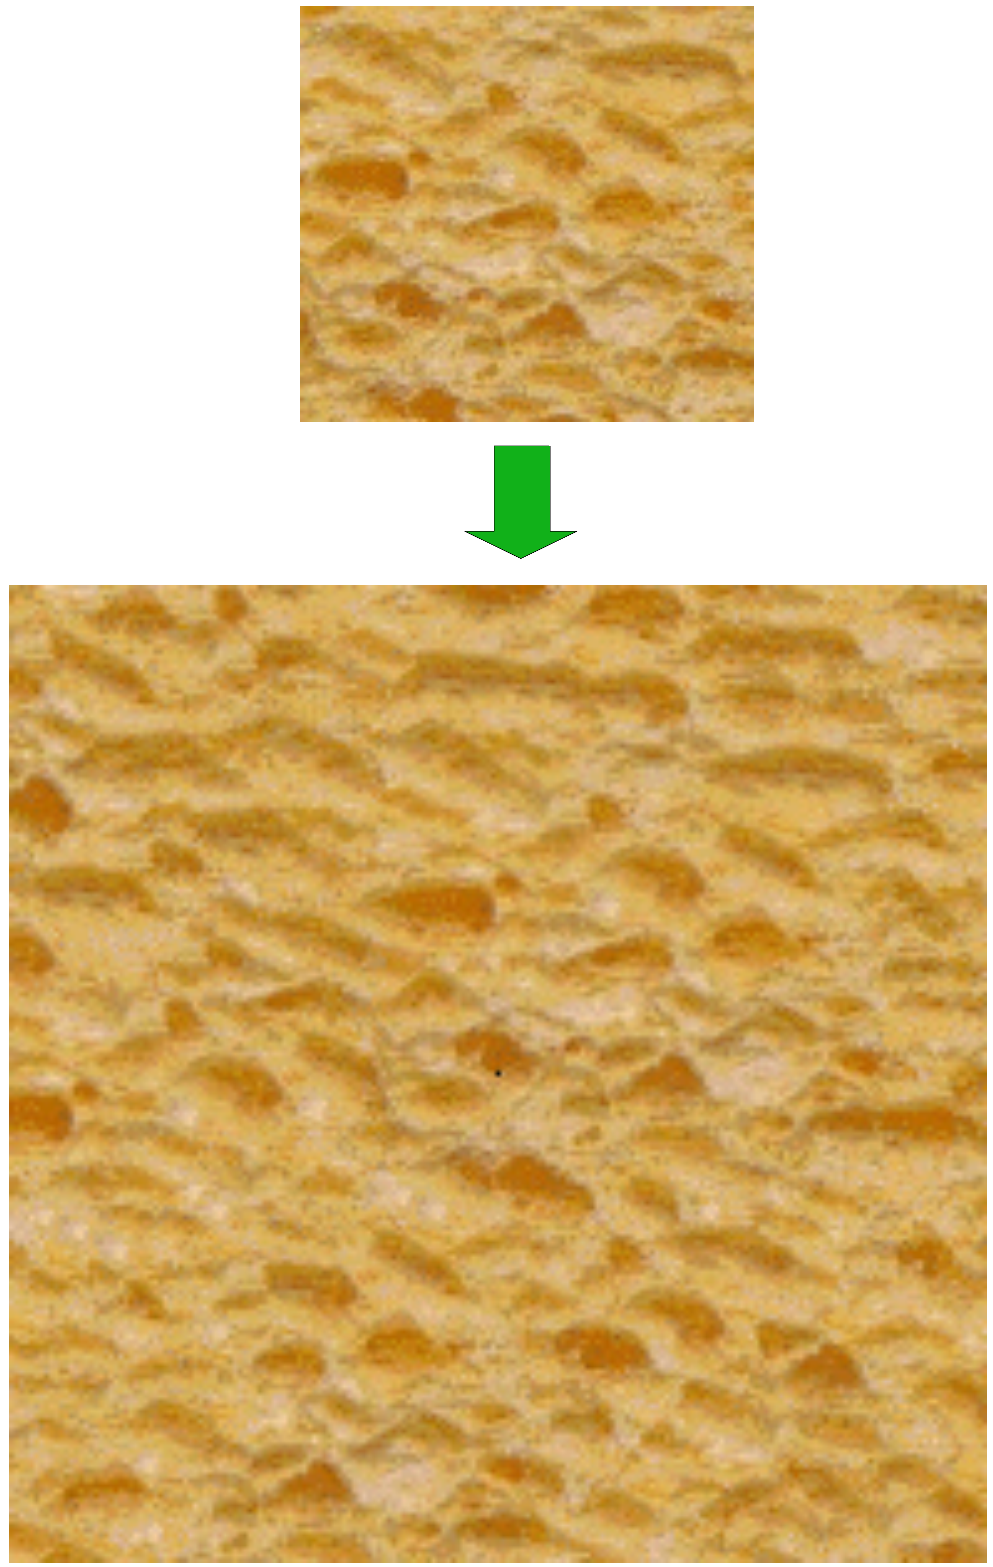

brick wall

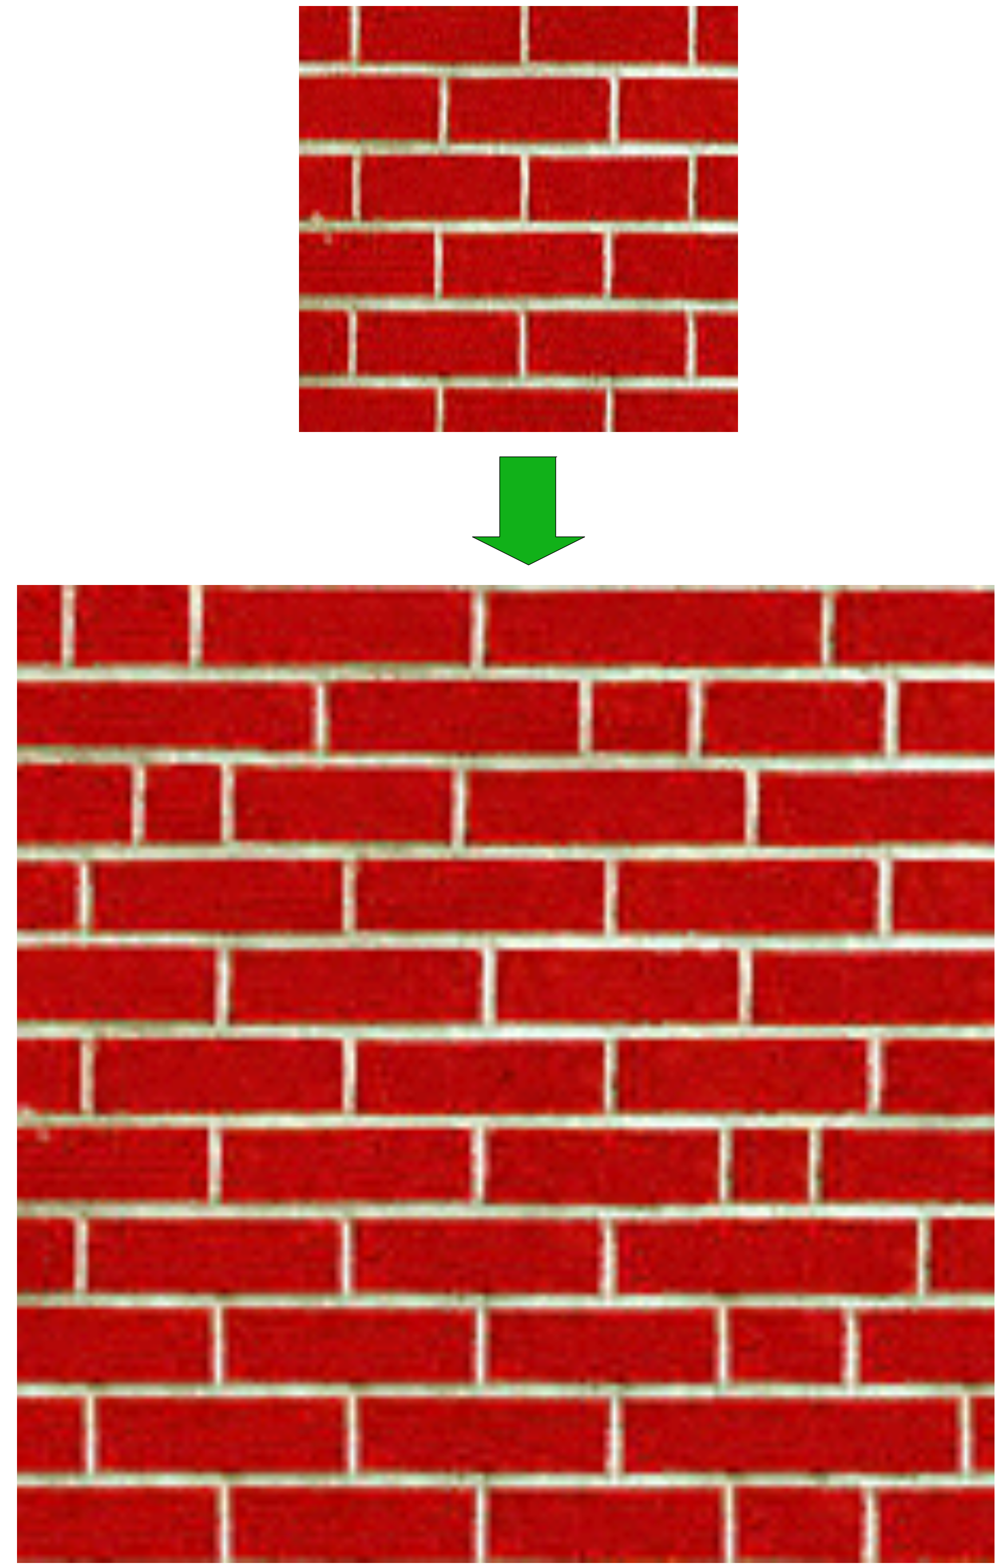

Slide from Alyosha Efros, ICCV 1999

Synthesis results

Slide from Alyosha Efros, ICCV 1999

Failure Cases

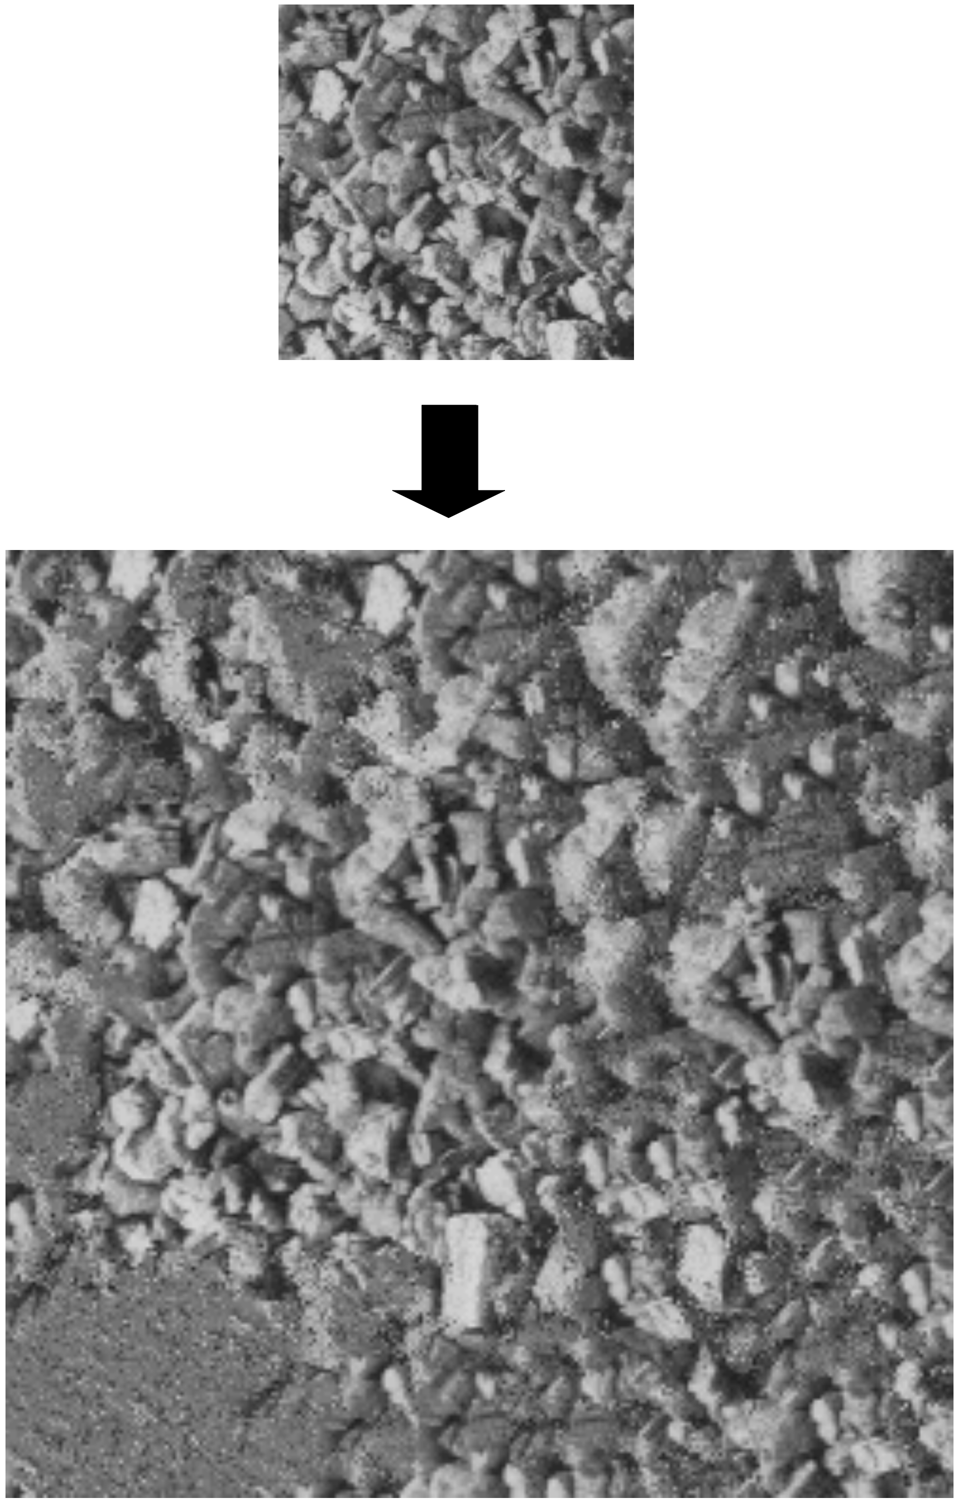

Growing garbage

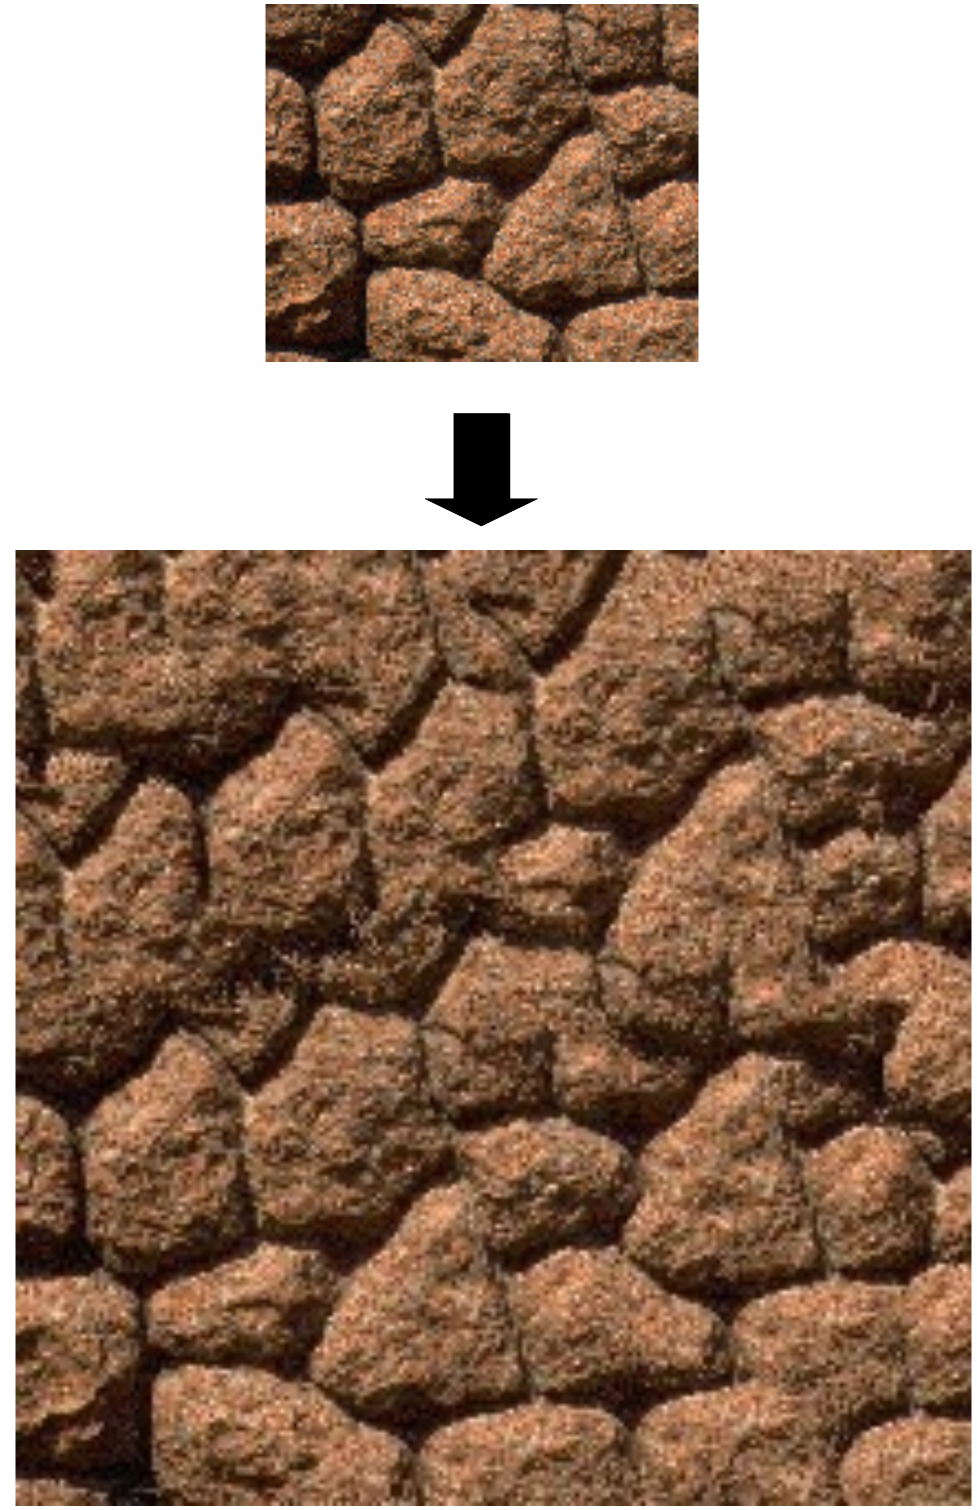

Verbatim copying

Slide from Alyosha Efros, ICCV 1999

Extrapolation

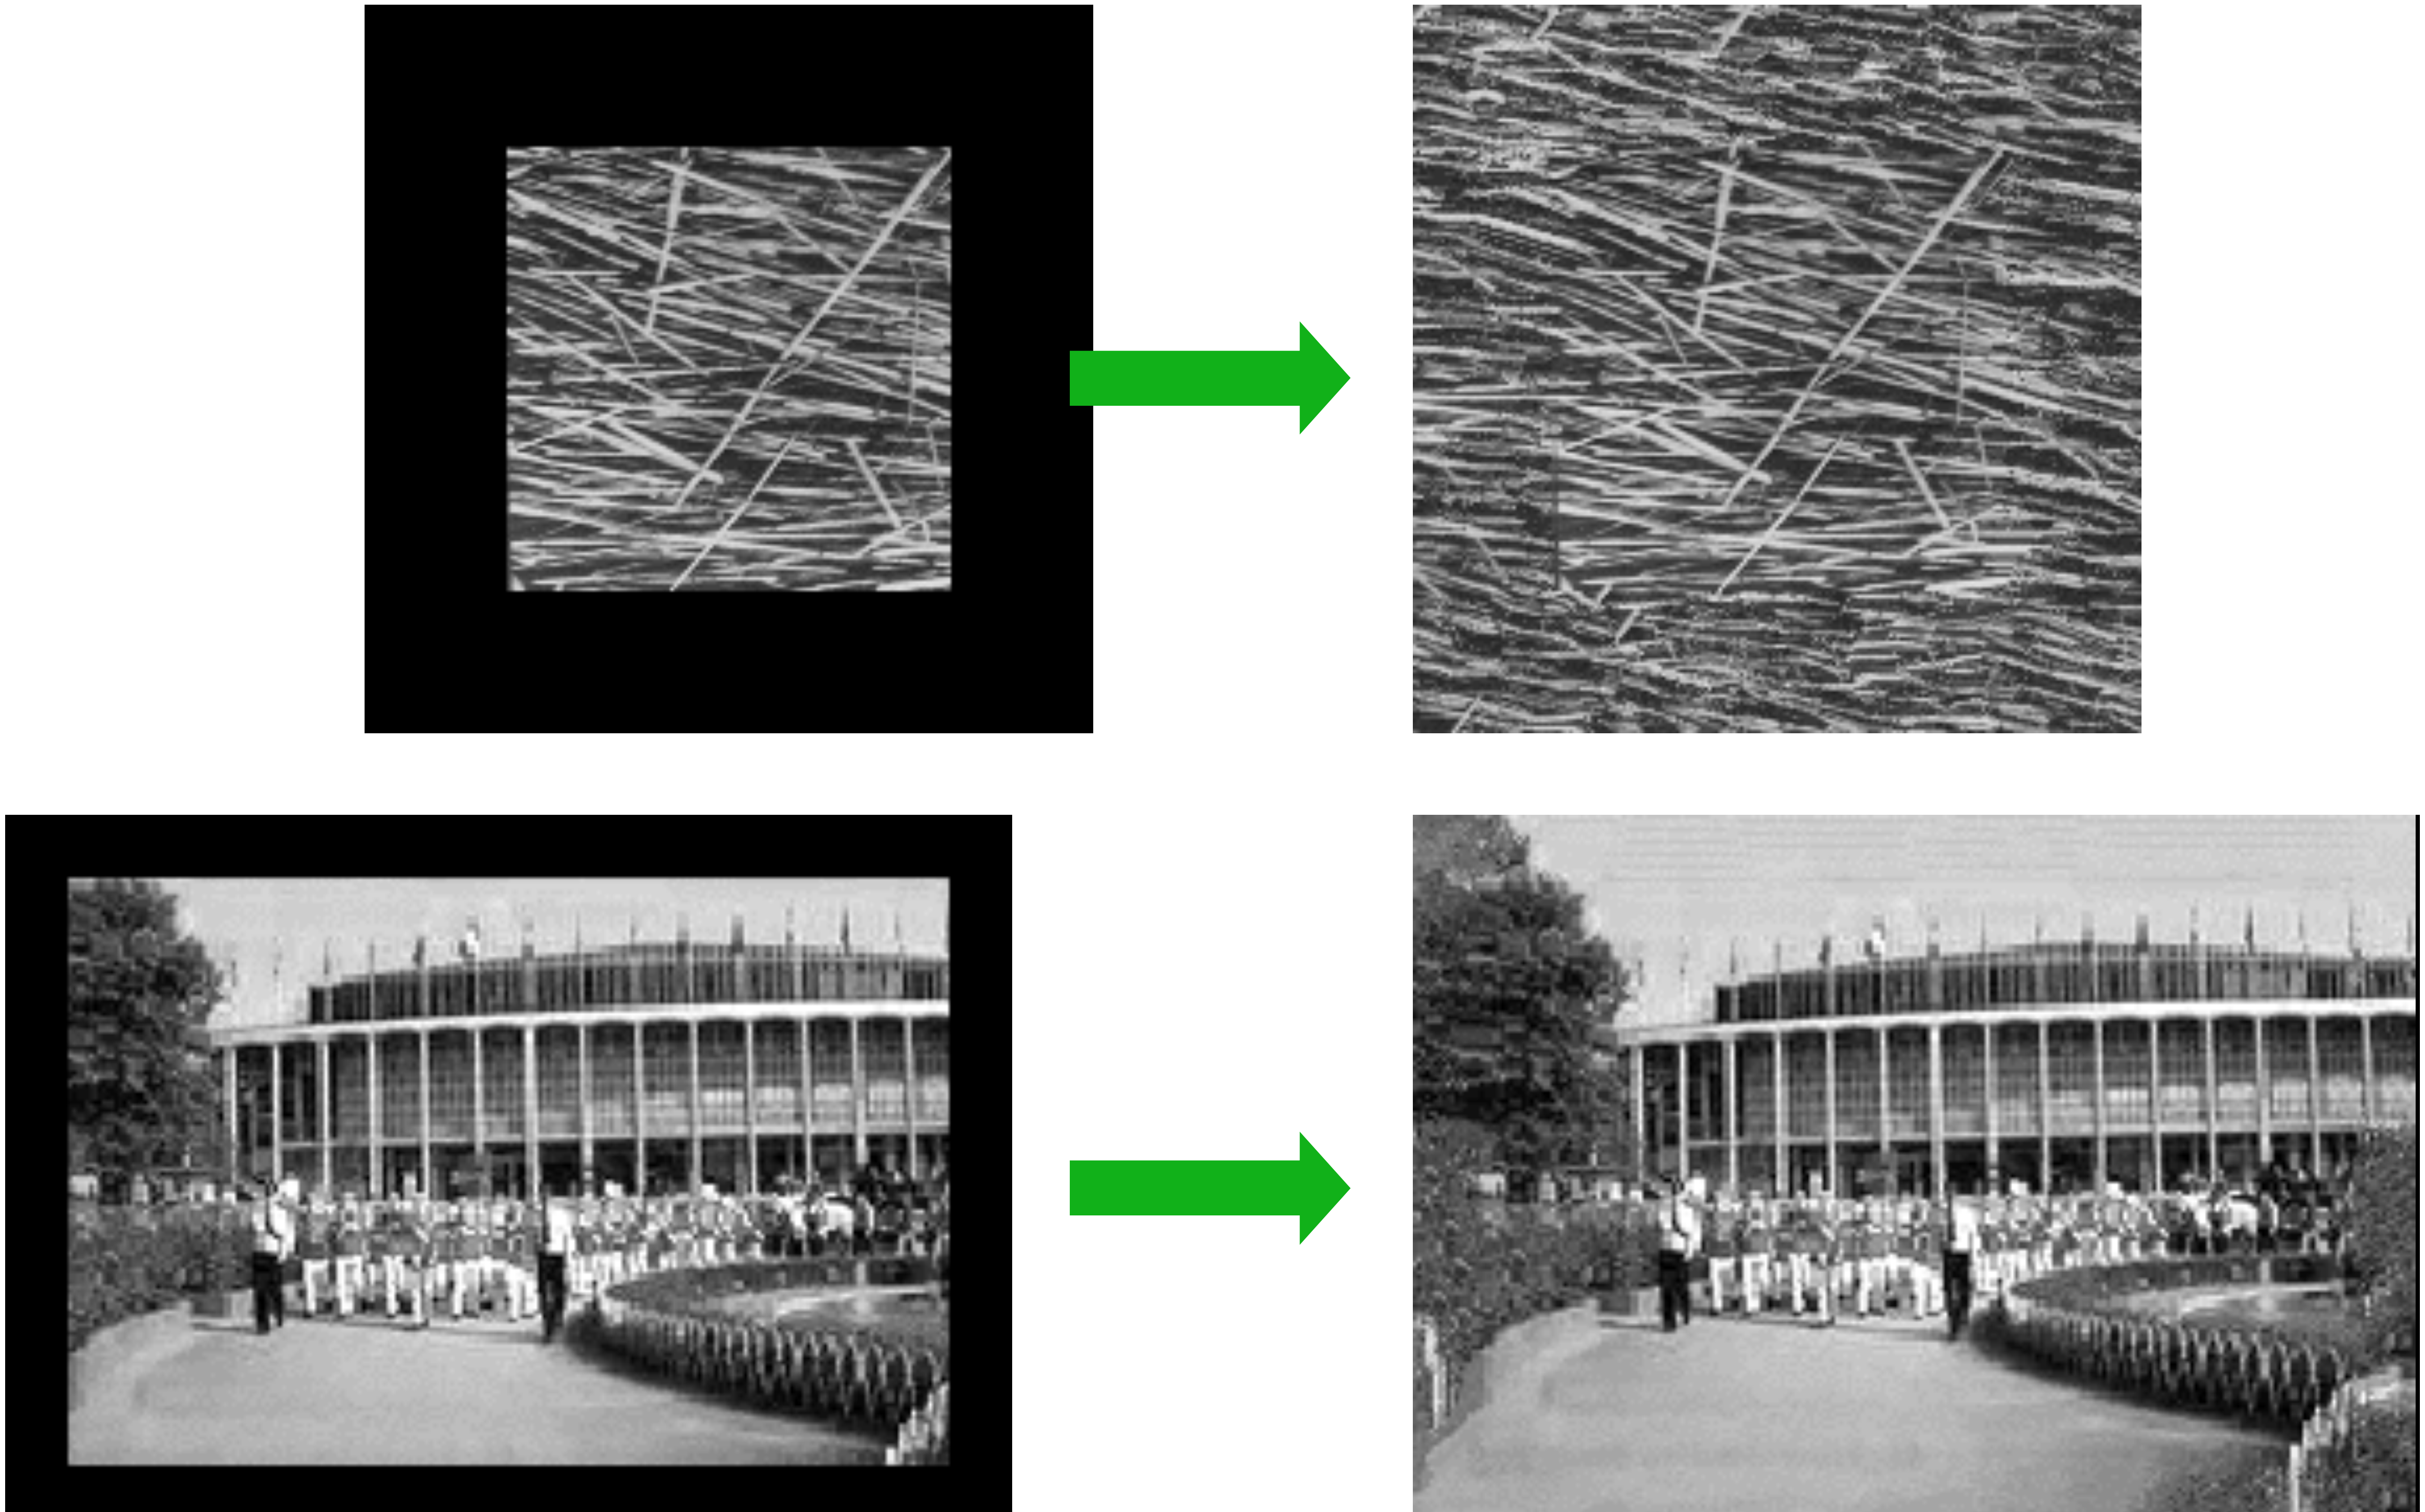

Slide from Alyosha Efros, ICCV 1999

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 nxmxpxq patch comparisons

Image Quilting [Efros & Freeman 2001]

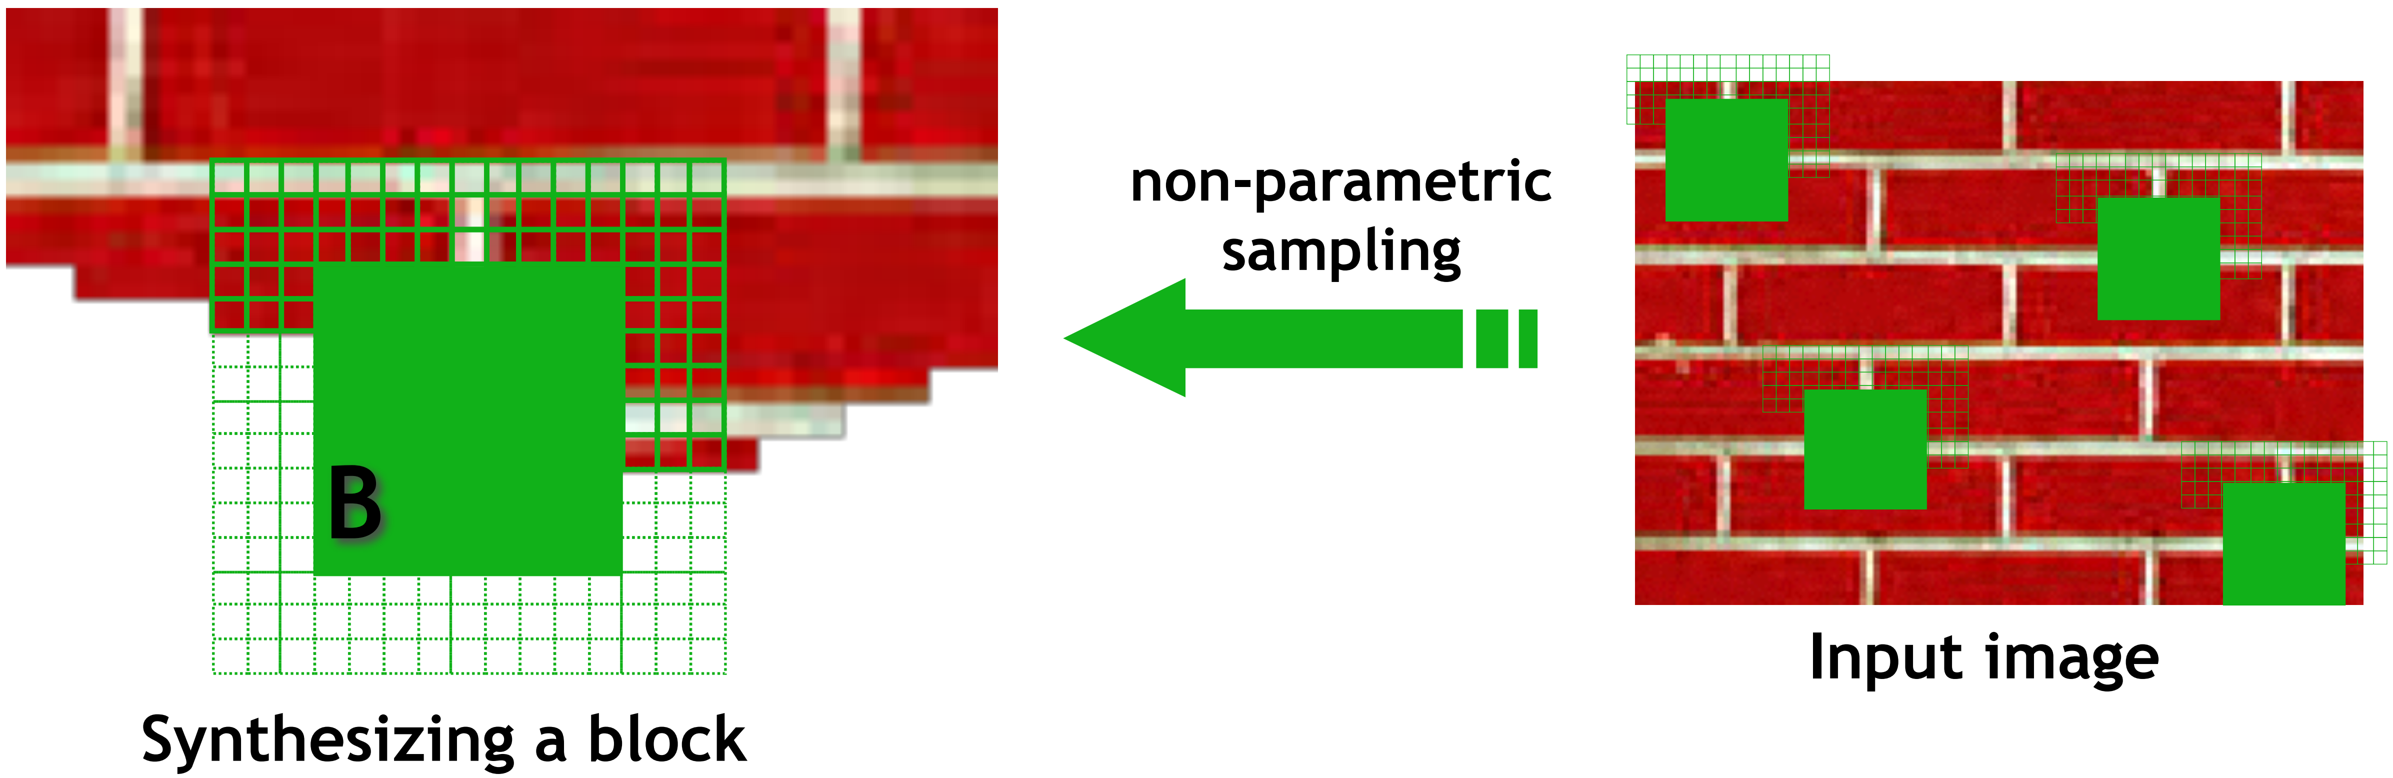

- Observation: neighbor pixels are highly correlated

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 from Alyosha Efros, ICCV 1999

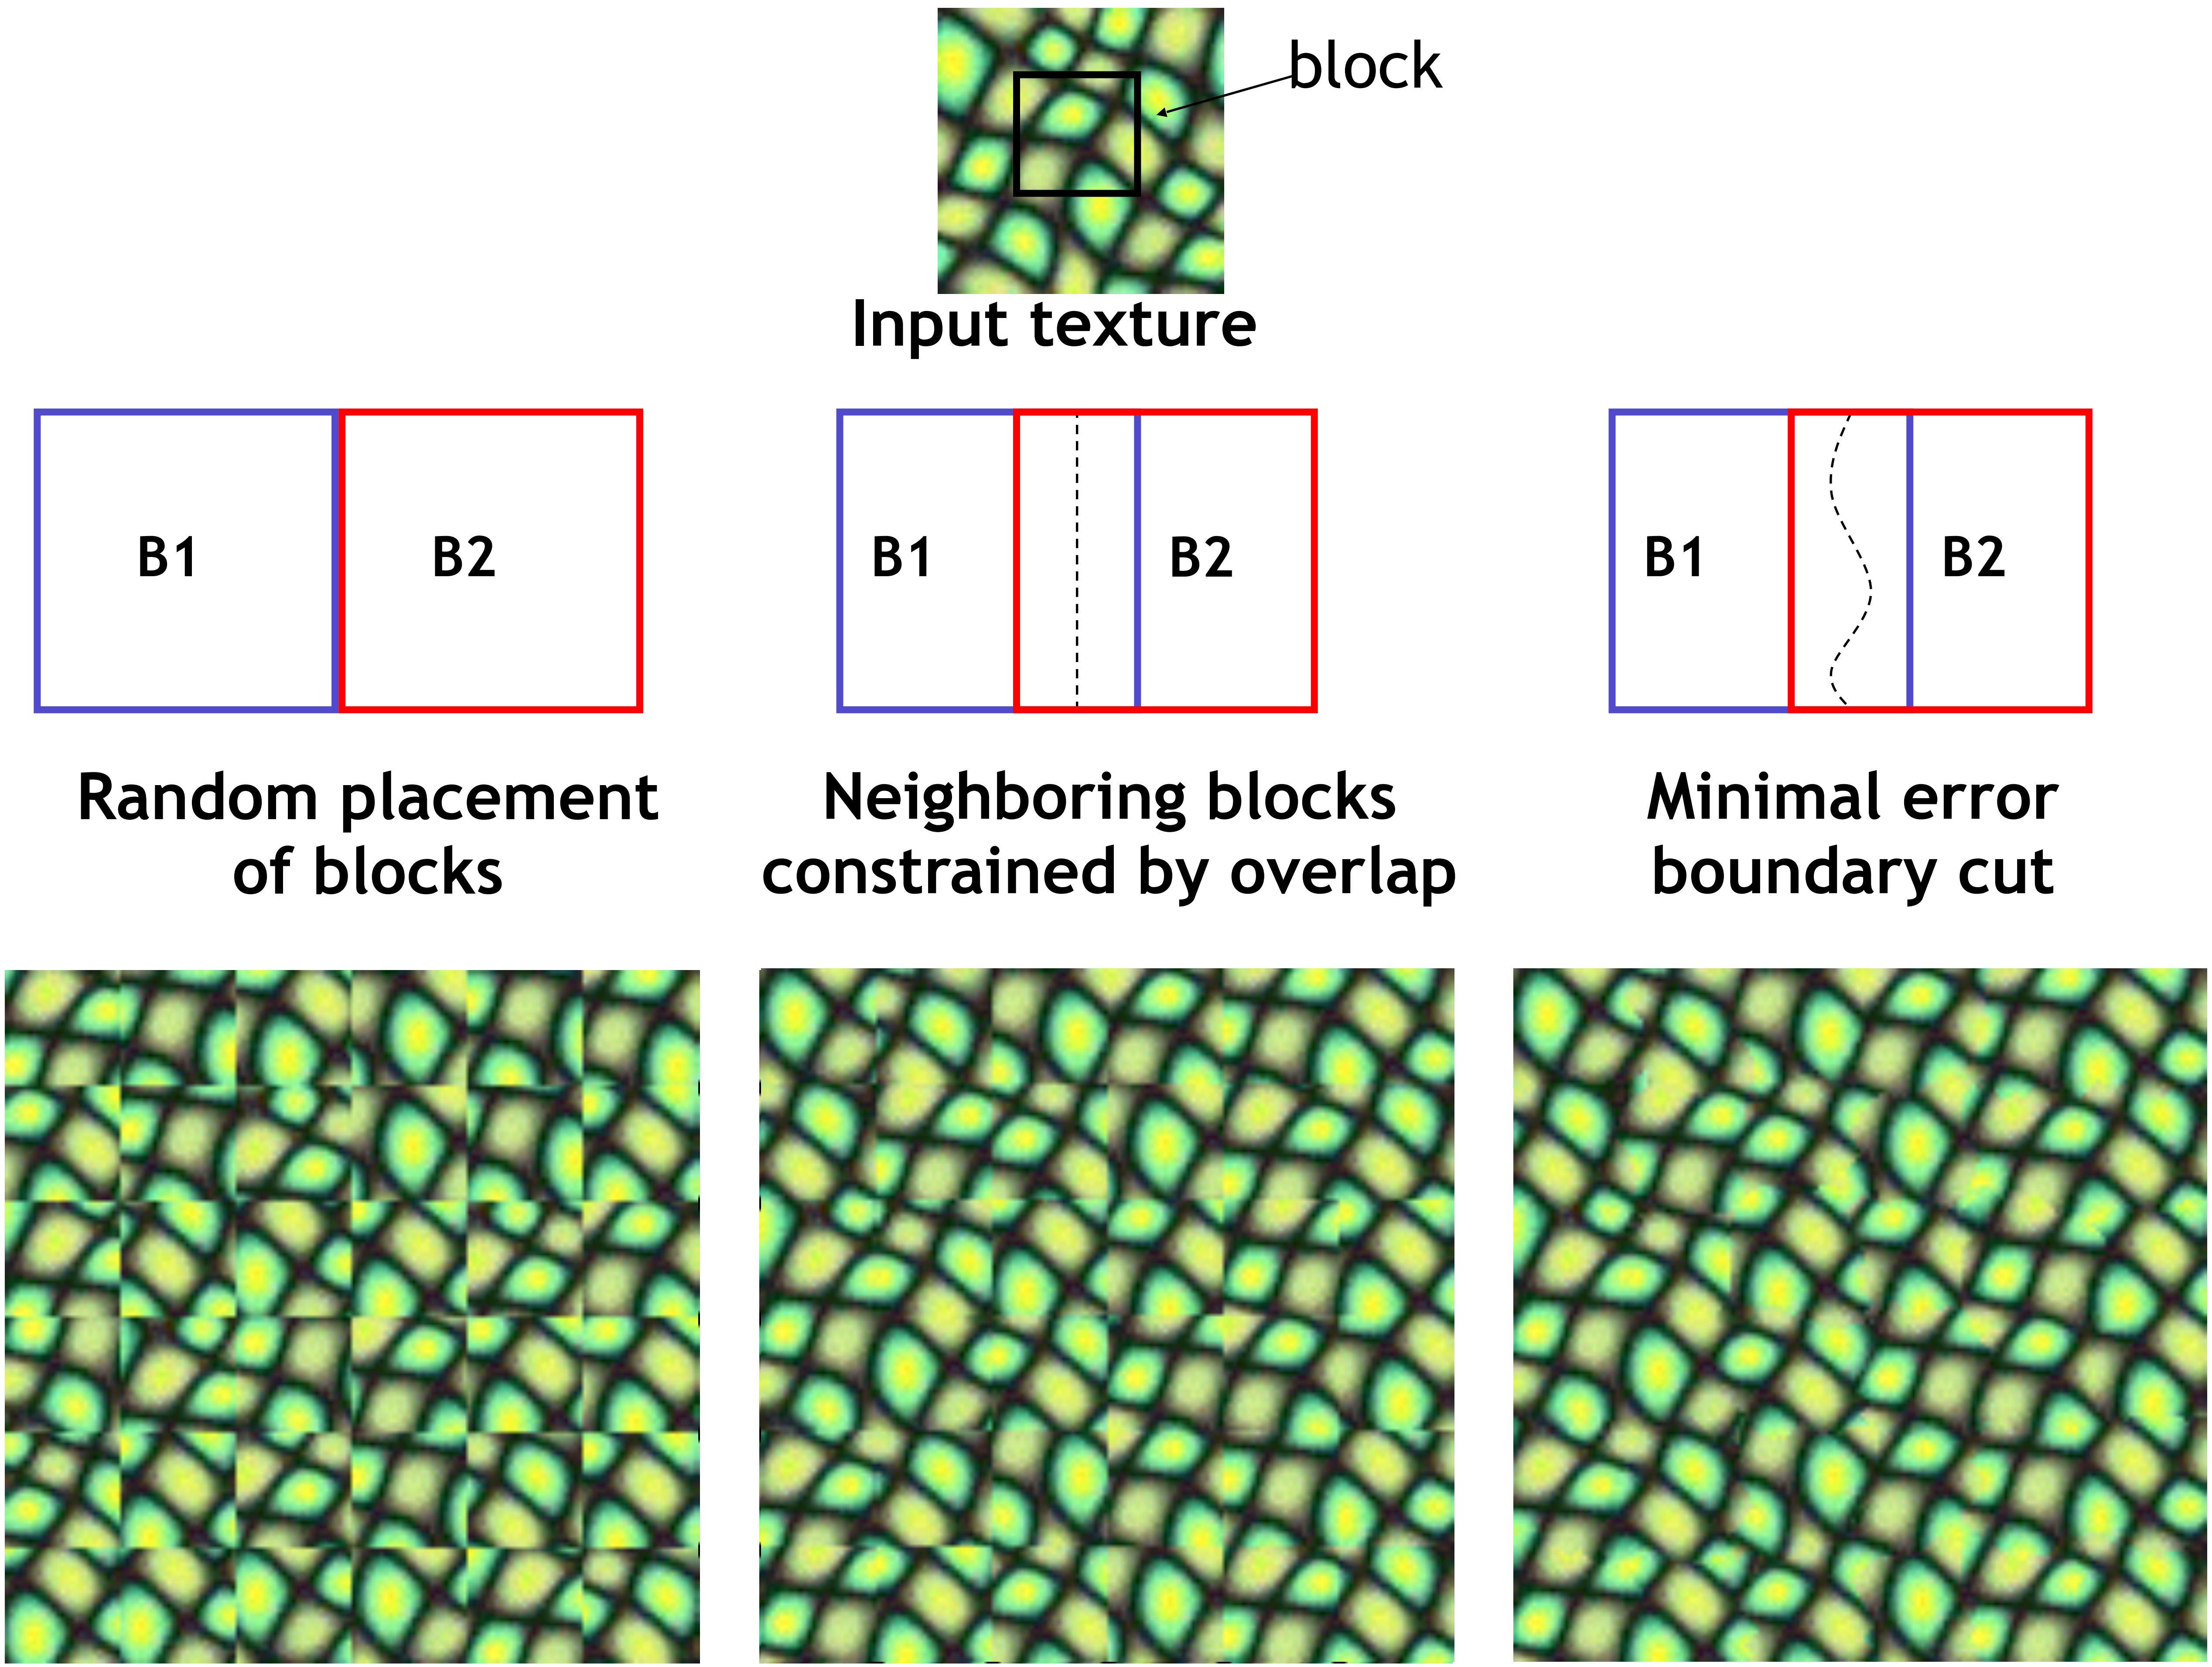
block
Input texture
B1
B2
B1
B2
B1
B2
Random placement of blocks
Neighboring blocks constrained by overlap
Minimal error boundary cut

Minimal error boundary

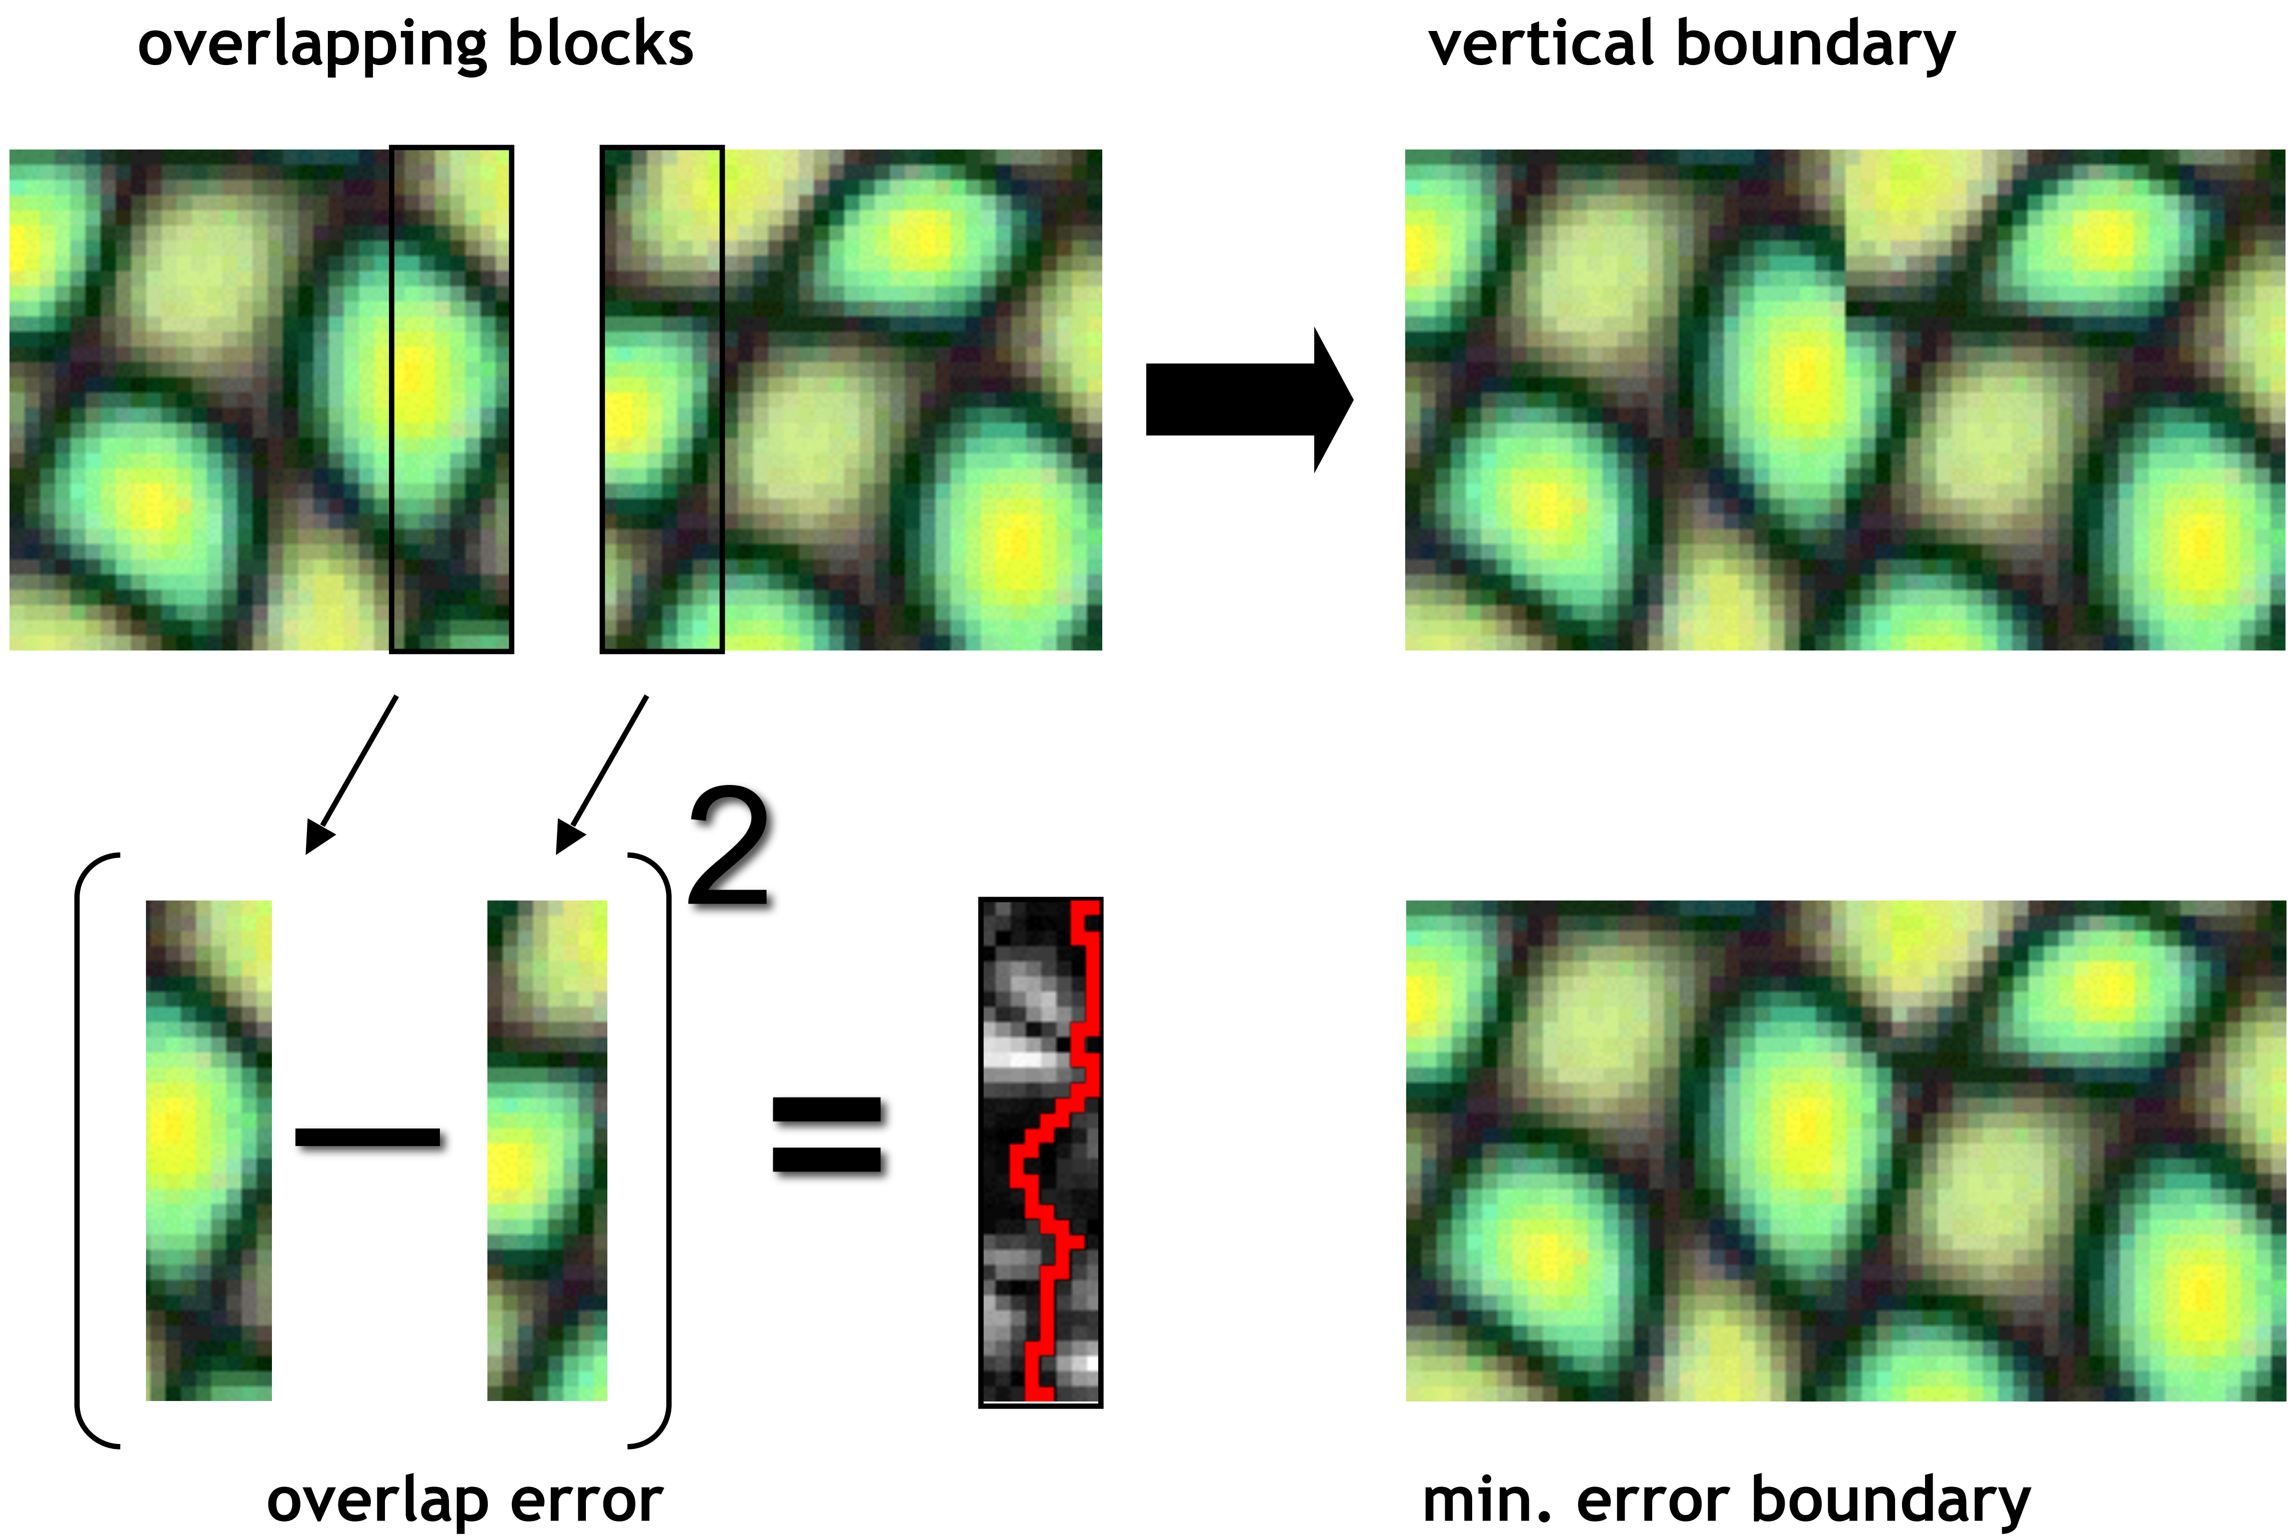

Slide from Alyosha Efros





College Inn
Chicken Broth

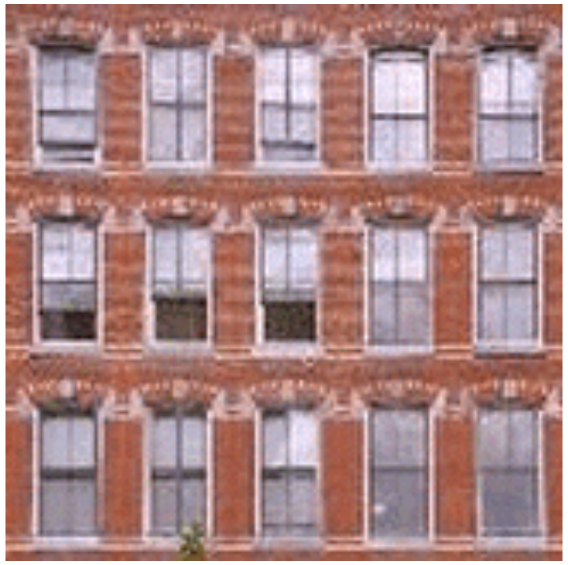


College Inn
Chicken Broth




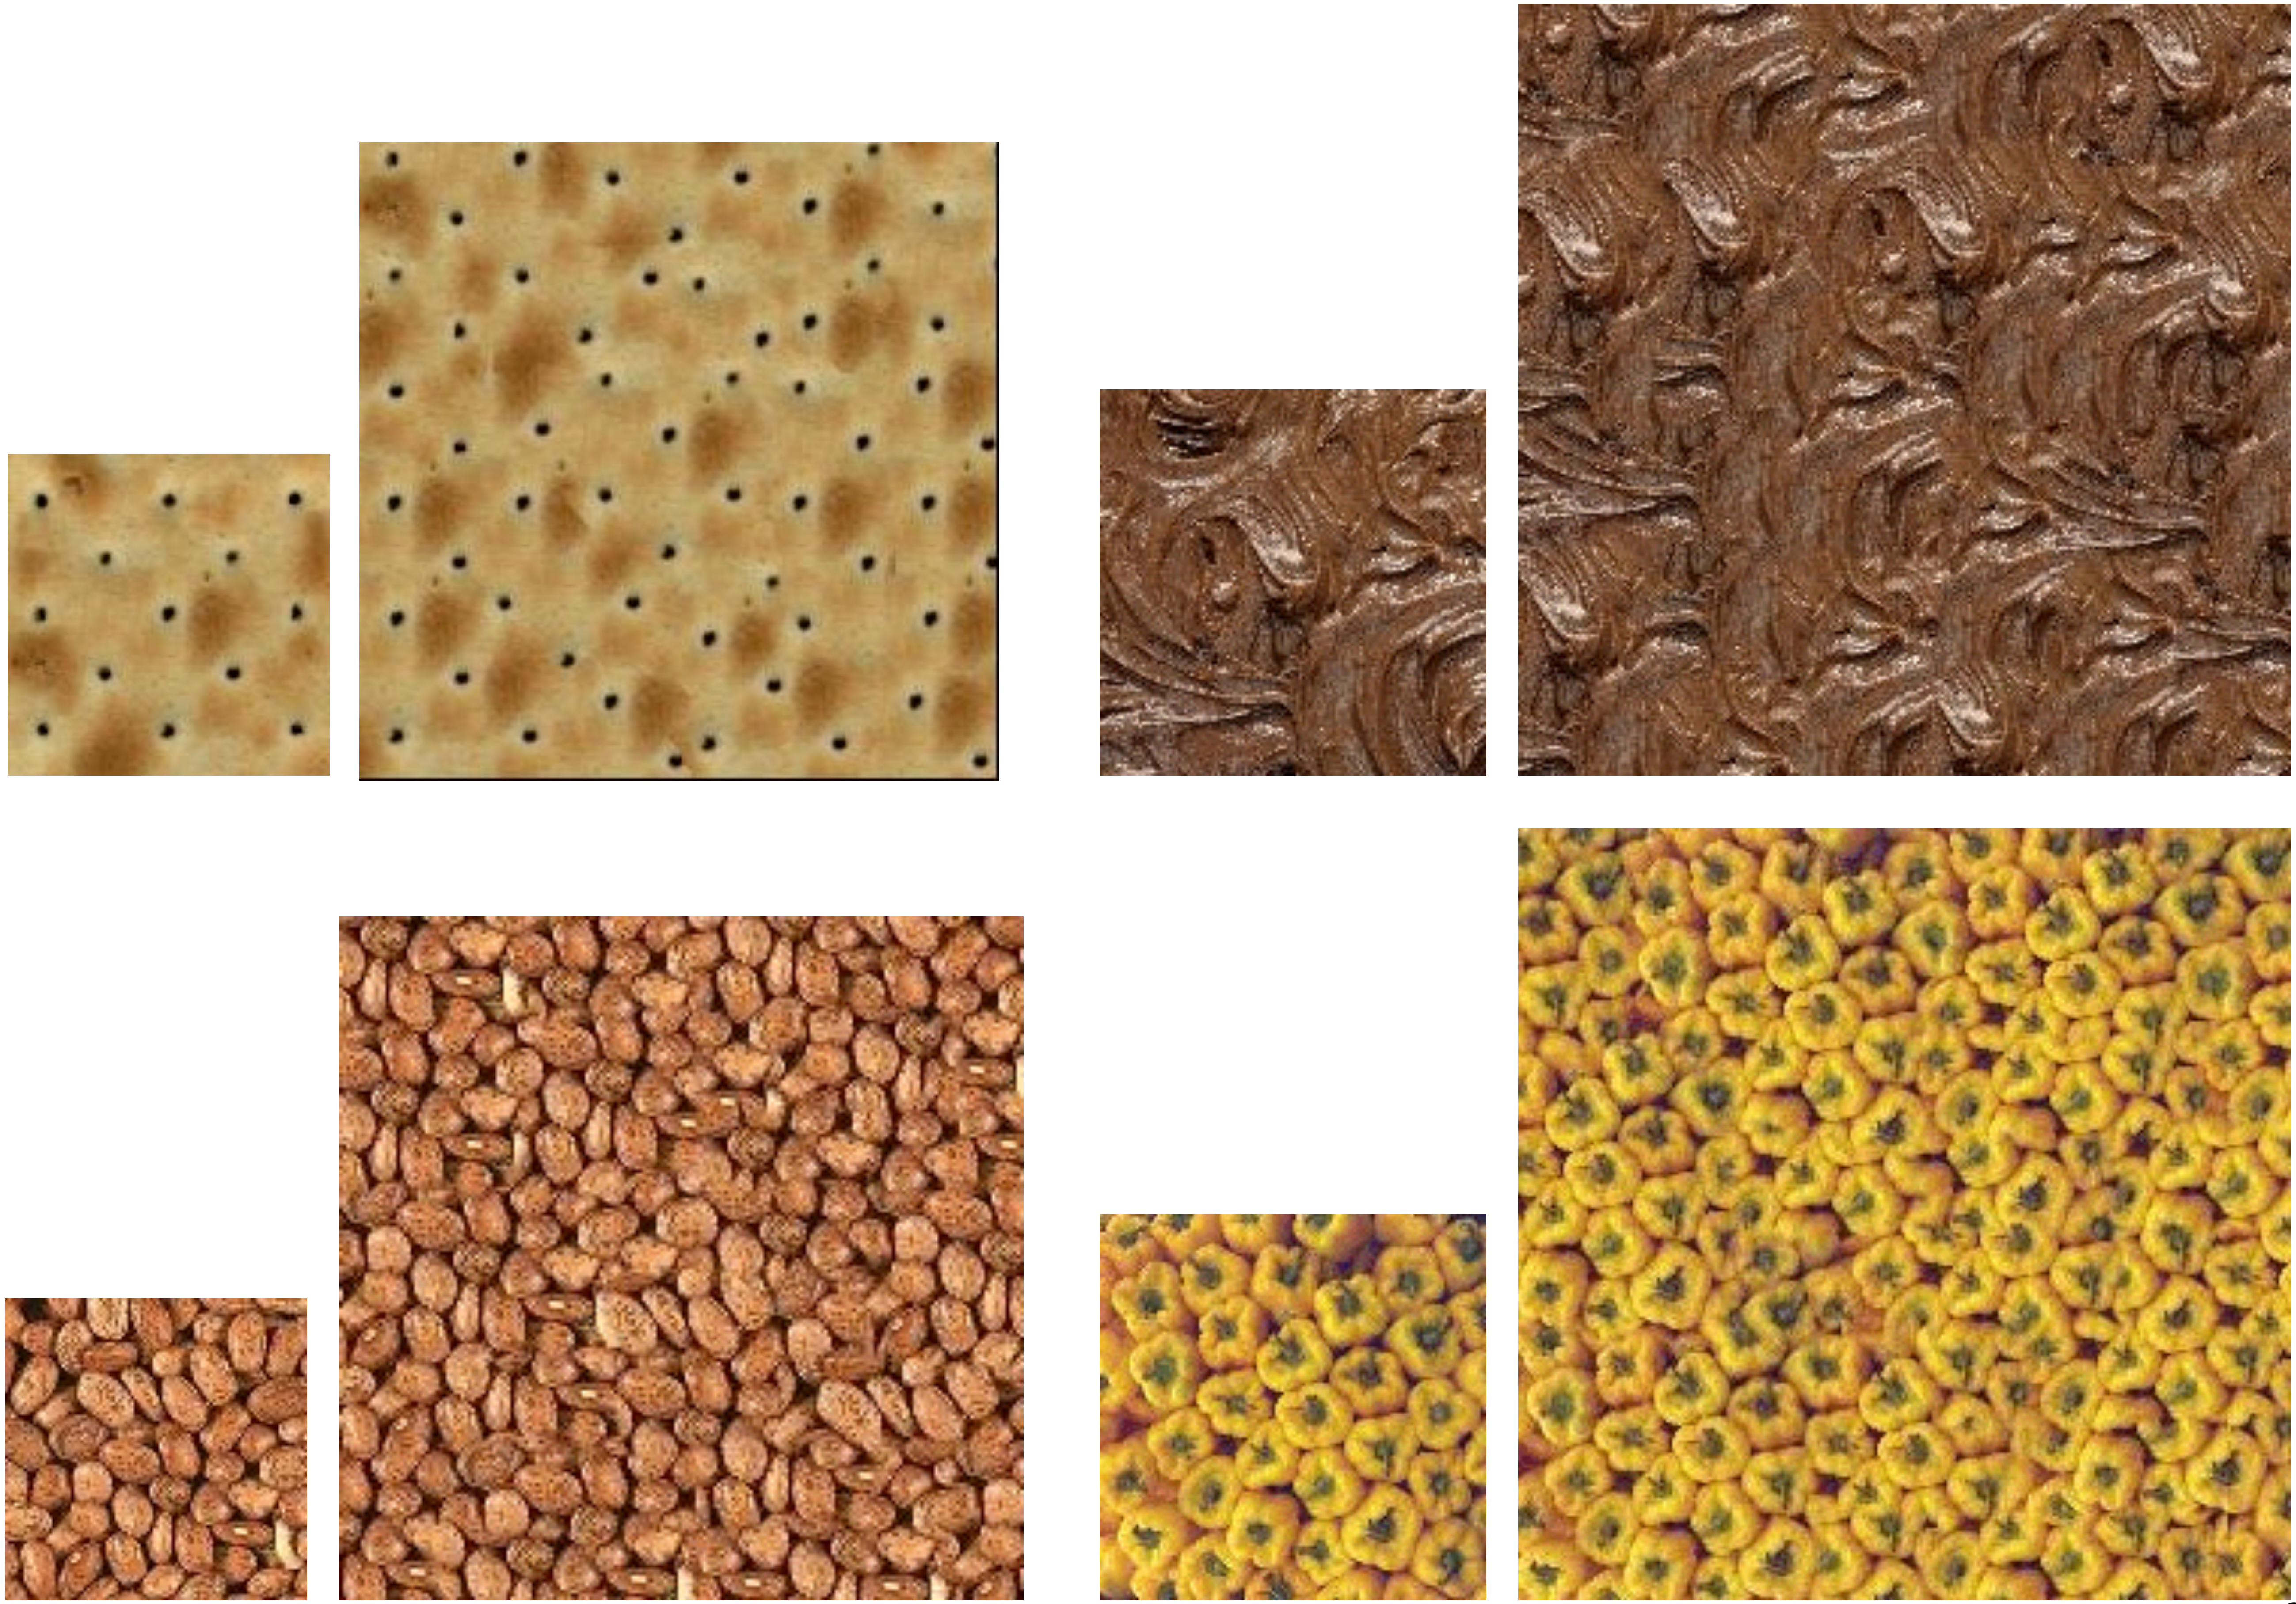

Failures
(Chernobyl Harvest)

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

+

parmesan

=

+

rice

=

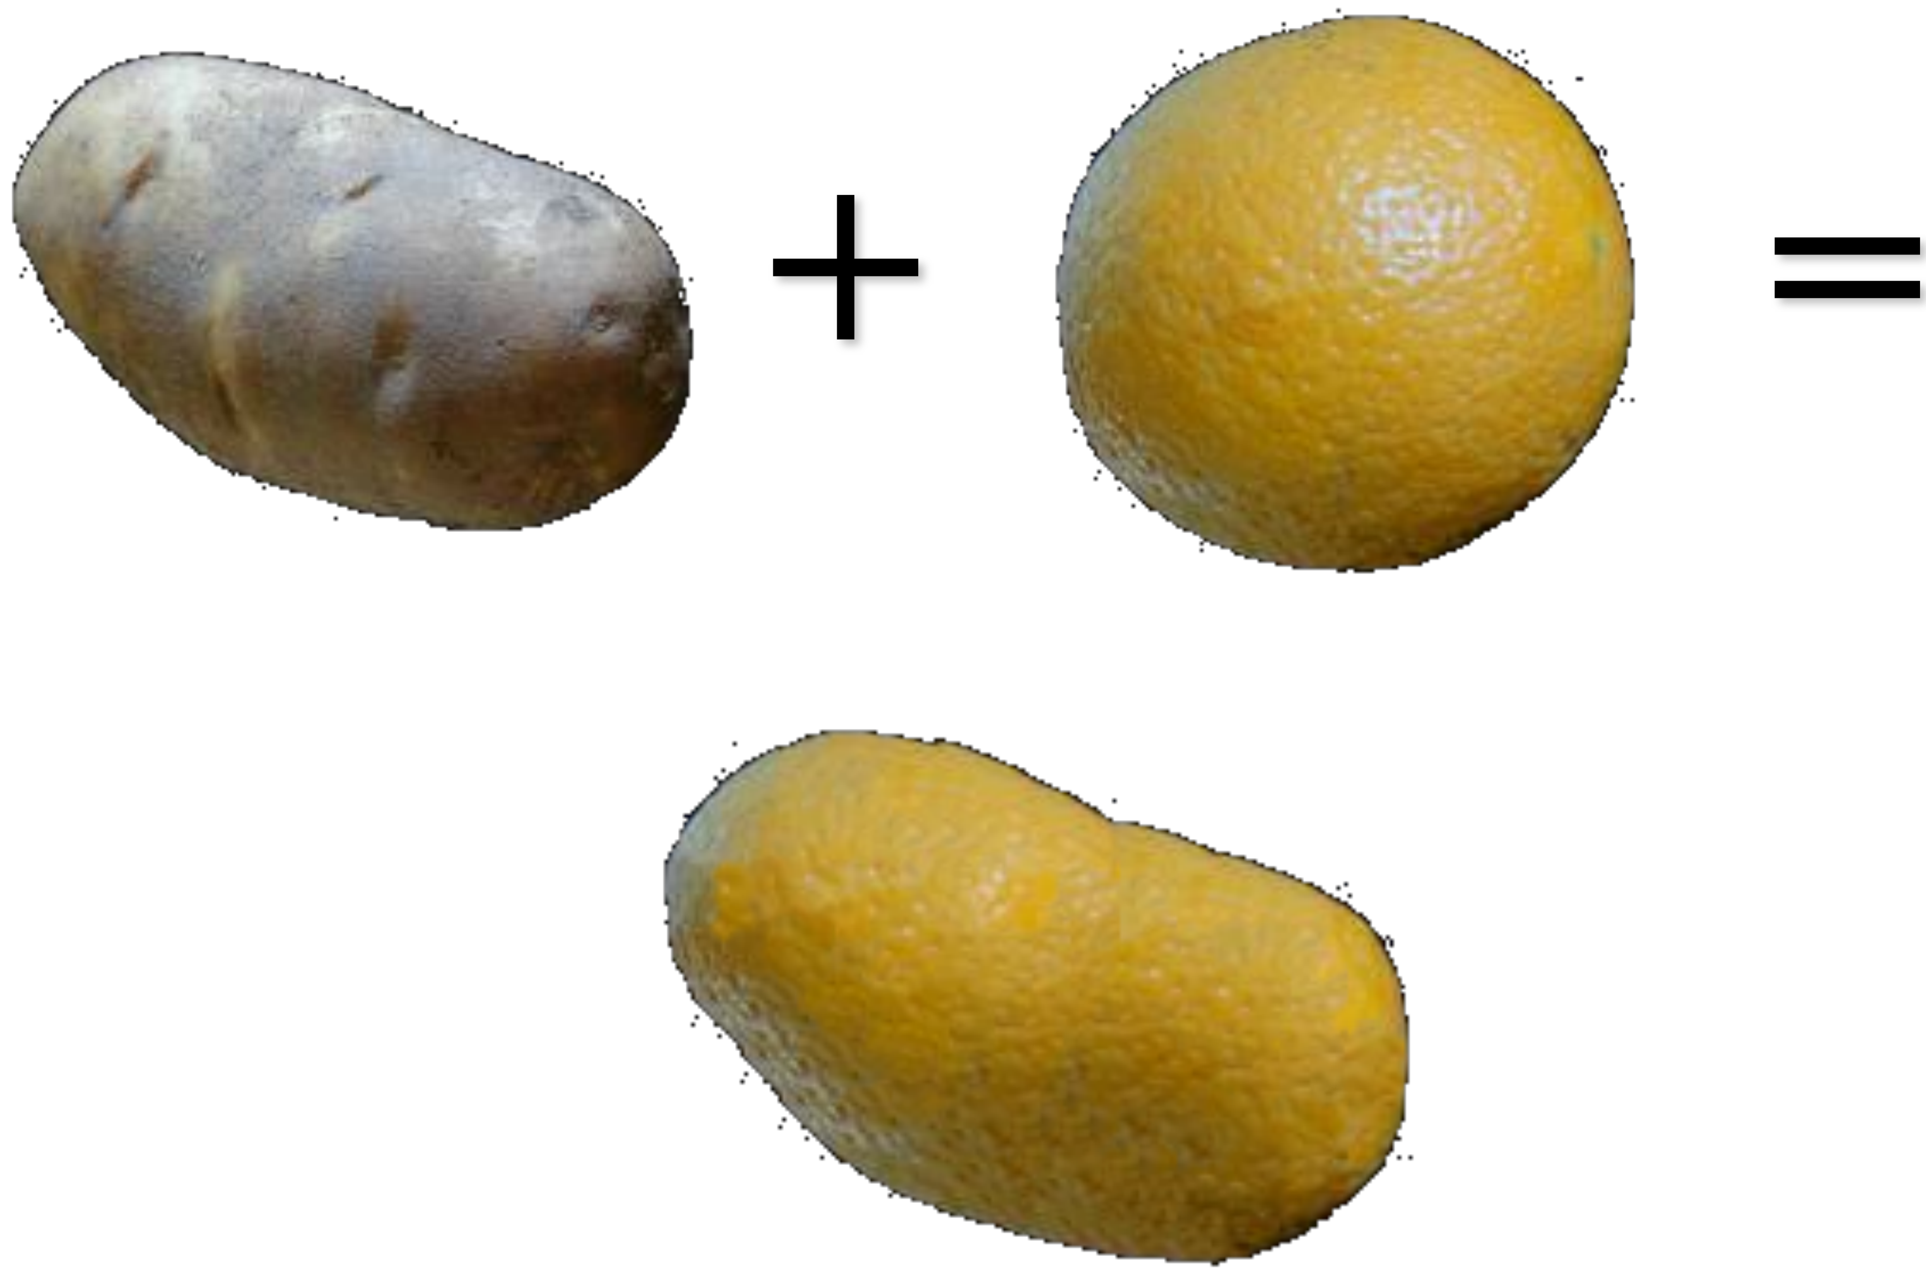

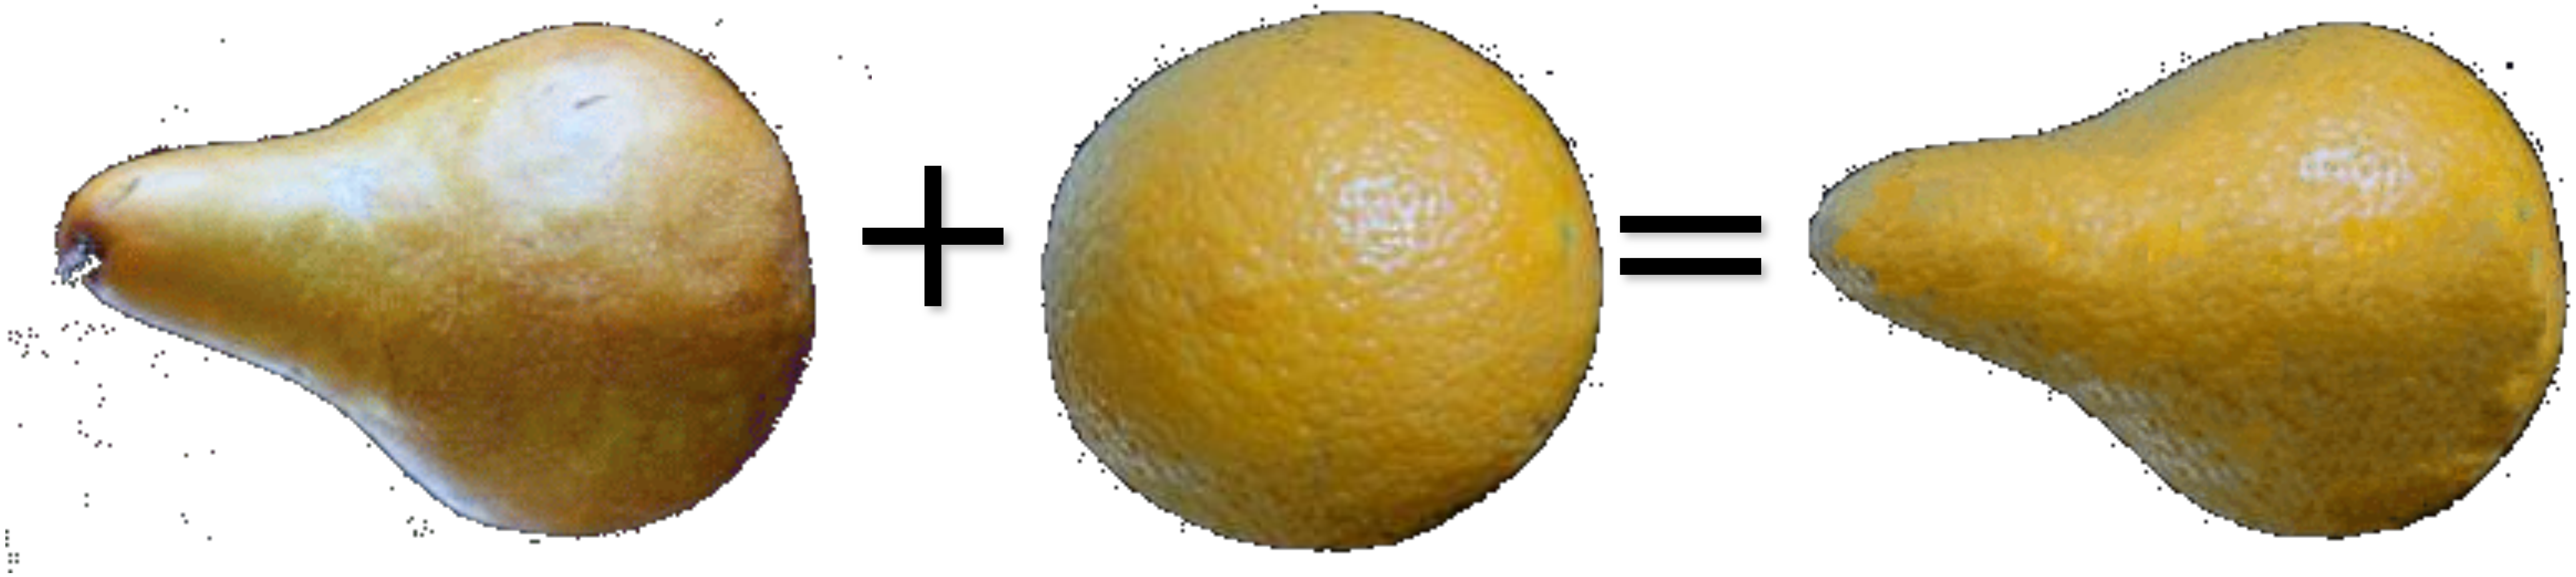

(Manual) texture synthesis in the media

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

Conclusion

- 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
  - Vector quantize to build histograms
- Neighborhood statistics can be exploited to “sample” or **synthesize** new texture regions
  - Example-based technique

Further thoughts and readings ...

- Texture and human psychophysics
  - *Bela Julesz*, Textons, the elements of texture perception and their interactions, Nature 1981 **pdf**
  - *N. Bhusan et al.*, The Texture Lexicon: Understanding the Categorization of Visual Texture Terms and Their Relationship to Texture Images **pdf**
- Texture representation
  - Are filter banks necessary? (Varma and Zisserman, CVPR 2003)
  - Local binary patterns (Ojala, Pietikainen, Maenpaa, PAMI 2002)
- State of the art in texture classification
  - **http://people.cs.umass.edu/~smaji/papers/textures-cvpr14.pdf**
  - Learning to detect describable attributes, e.g. lined, dotted, blotchy, striped, checkered, etc.