Texture Synthesis

Daniel Cohen-Or

Texture Weathering

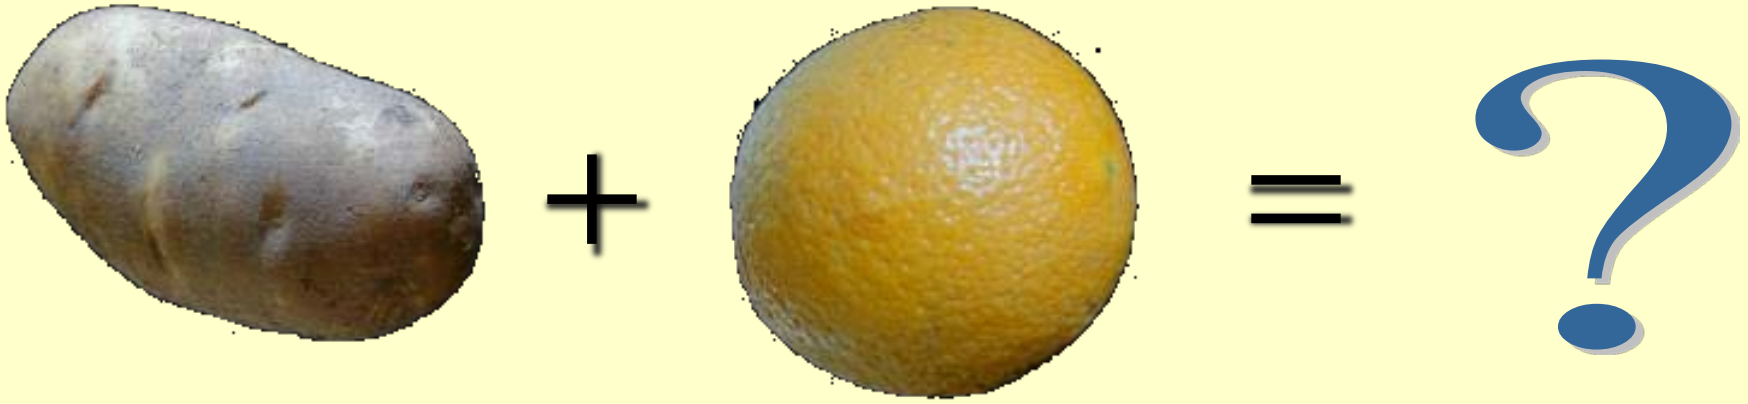

The slides are based of Efros and Freeman “Efros&Freeman Image Quilting“

+

=

?

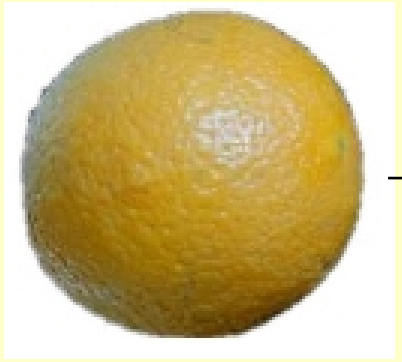

The Goal of Texture Synthesis

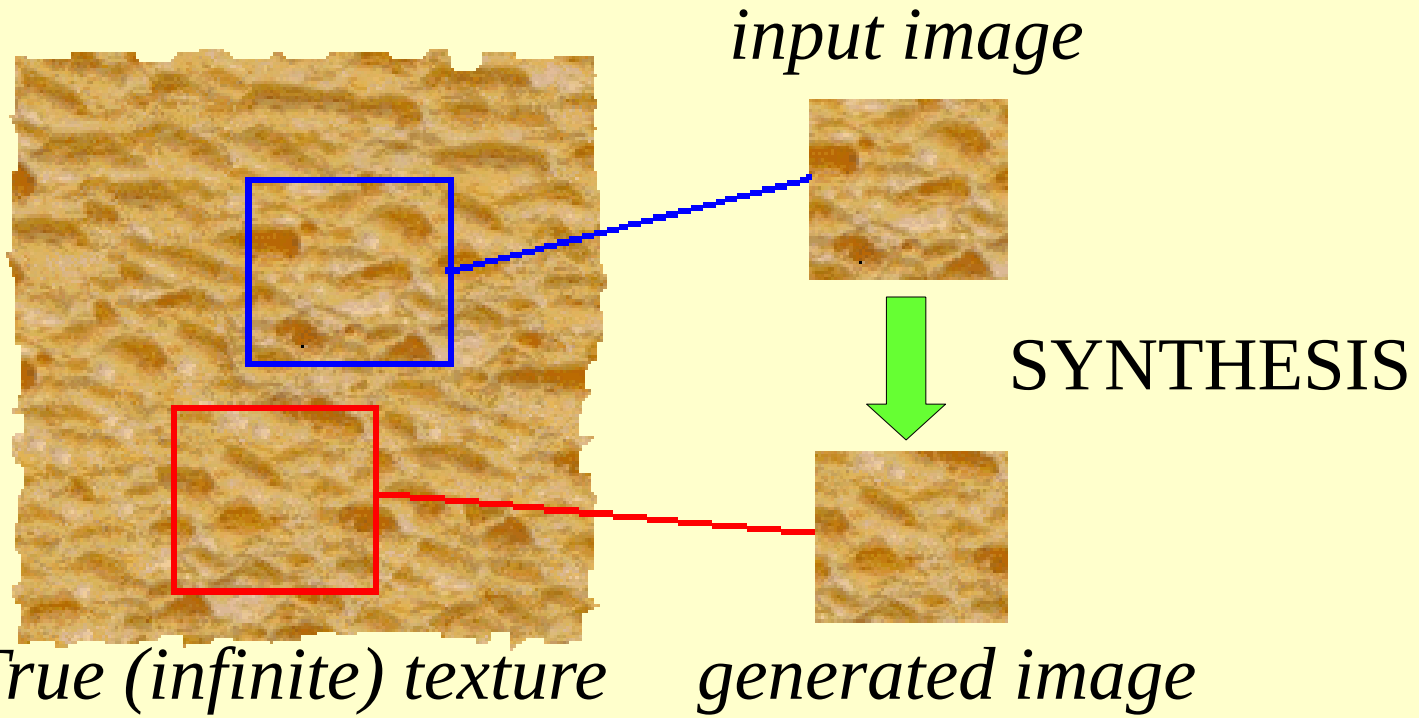

- Given a finite sample (large enough) of some texture, the goal is to synthesize other samples from that same texture.

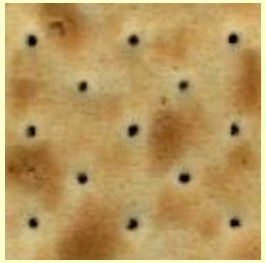

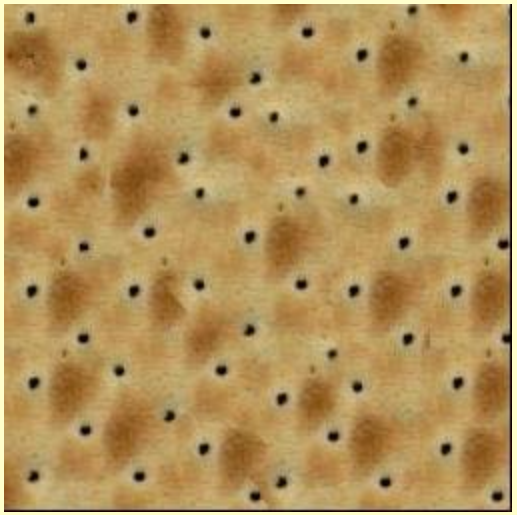

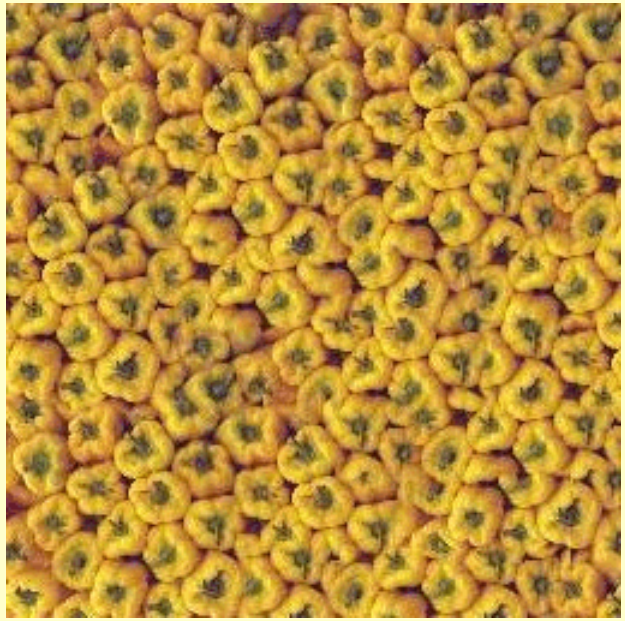

The Challenge

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

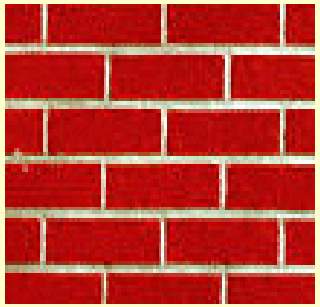

repeated

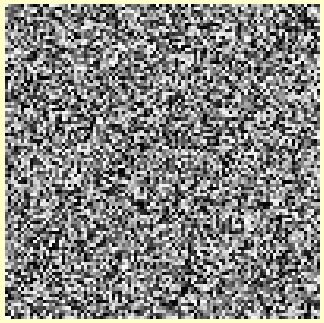

stochastic

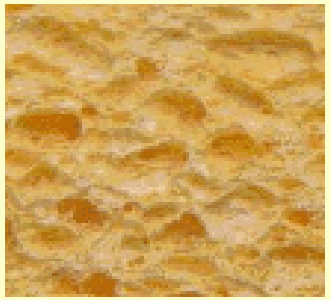

Both

Texture Types

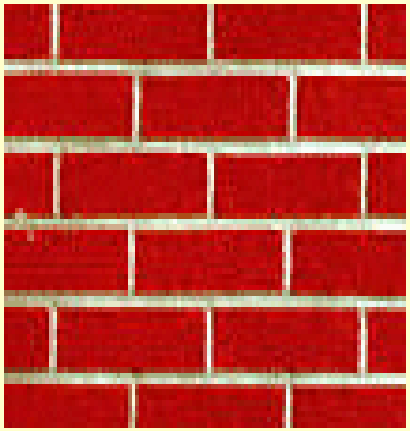

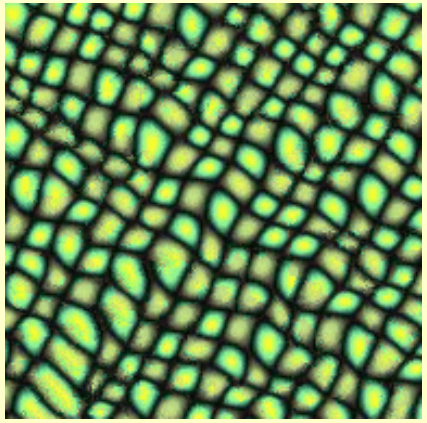

Texture model

Stationary - under a proper window size, the observable portion always appears similar.

Local - each pixel is predictable from a small set of neighboring pixels and independent of the rest of the image.

Non-Stationary

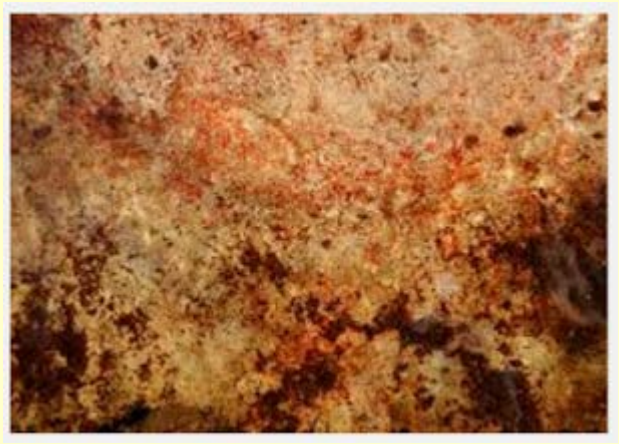

Non-Stationary

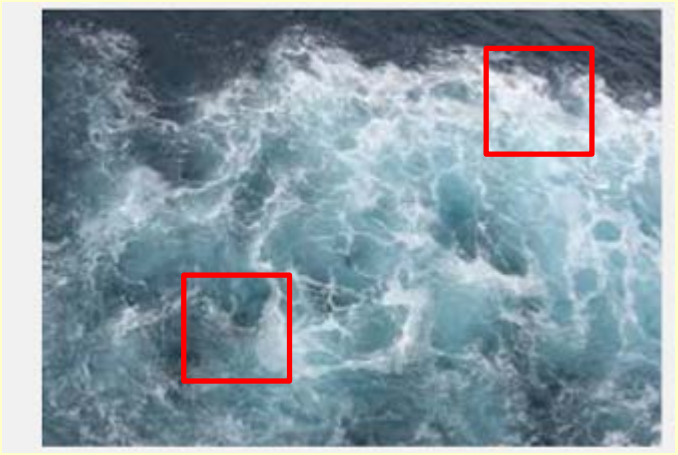

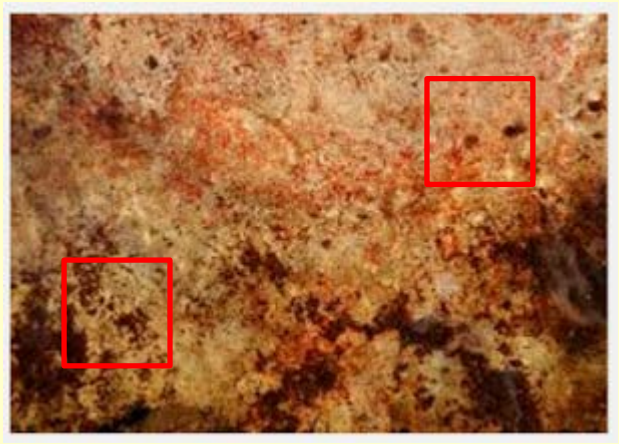

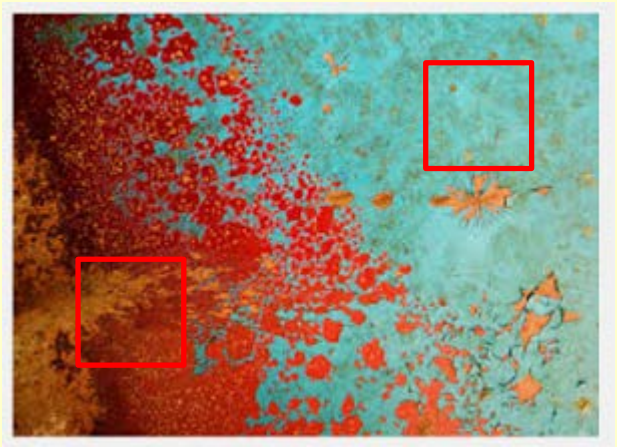

Texture Synthesis for Graphics

- Inspired by Texture Analysis and Psychophysics
  - [Heeger & Bergen,'95]
  - [DeBonet,'97]
  - [Portilla & Simoncelli,'98]
- …but didn't work well for structured textures
  - [Efros & Leung,'99]
    - (originally proposed by [Garber,'81])

“By Example” Texture Synthesis

- Input patch boundary.

- Input texture example.

- Fill boundary with texture.

Texture Synthesis by Non Parametric Sampling

- Generate English-looking text using N-grams, [Shannon,'48]

- Assuming Markov Chain on letters:
  - *P(letter | Proceeding n-letters)*

E N C Y C L O P E D I

Synthesizing English-looking text

- [Shannon,'48] proposed a way to generate English-looking text using N-grams:
  - Assume a generalized Markov model
  - Use a large text to compute prob. distributions of each letter given N-1 previous letters
  - Starting from a seed, repeatedly sample this Markov chain to generate new letters
  - Also works for whole words

WE NEED TO EAT CAKE

Unit of Synthesis

- Letter-by-letter: Used to name planets in early 80s game “Elite”.
- Word-by-word: M.V. Shaney (Bell Labs) 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“.*

Mark V. Shaney (Bell Labs)

- Letter-by-letter:
- Word-by-word:
- *"I spent an interesting evening recently with a grain of salt".*

Notice how well local structure is preserved!

Now, instead of letters let's try pixels…

Efros & Leung 99*

Assuming Markov property, compute $P(p \mid N(p))$.

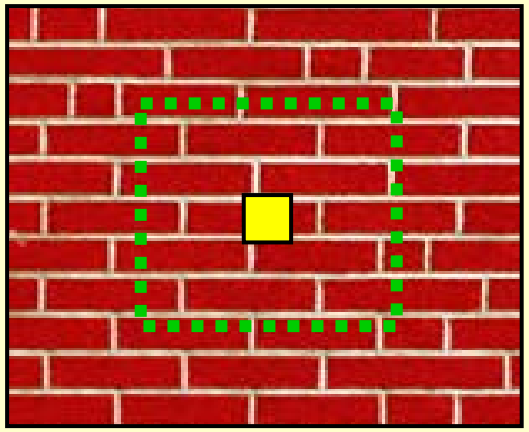

Non-parametric sampling

* A.A.Efros, T.K.Leung; "Texture synthesis by non-parametric sampling"; ICCV99.

(originally proposed by [Garber,'81])

Non-parametric Sampling

- *P(p | N(p))*

– Explicit probability tables infeasible.
– Instead, *search input image* for similar neighbourhoods - that's our histogram for *p*.

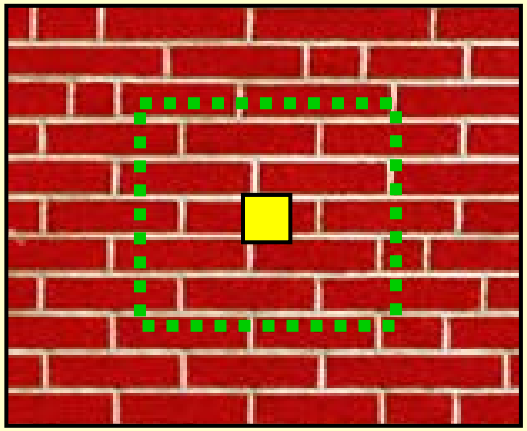

Efros & Leung 99 - Algorithm

Sample

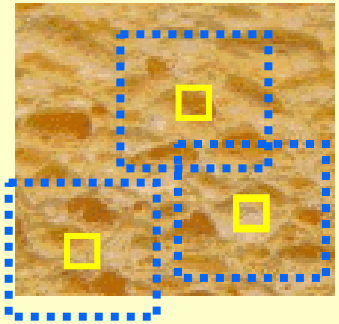

Output

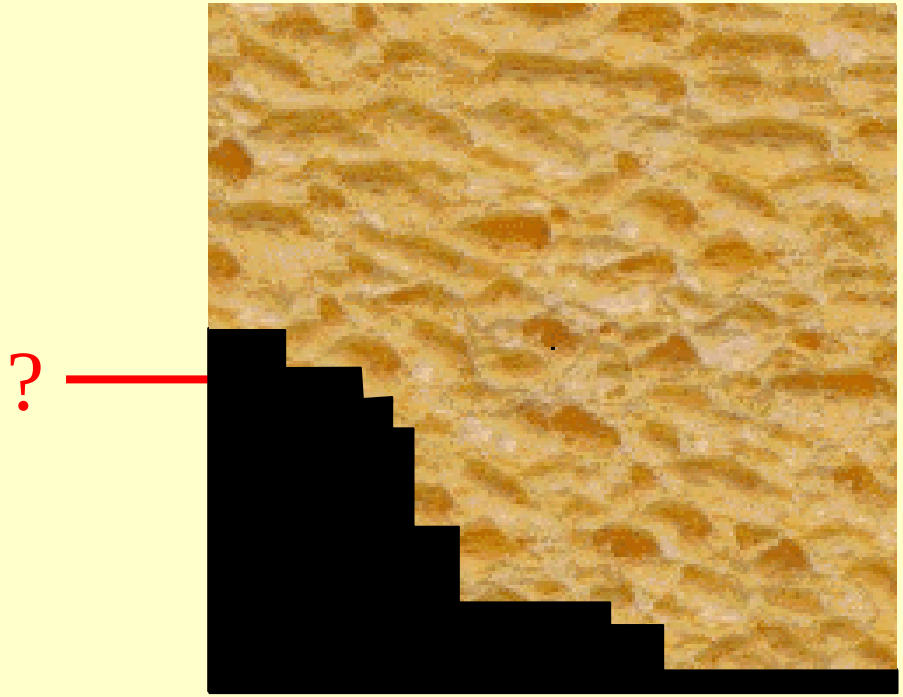

- Causal neighborhood – Neighboring pixels with known values.

Efros & Leung ’99

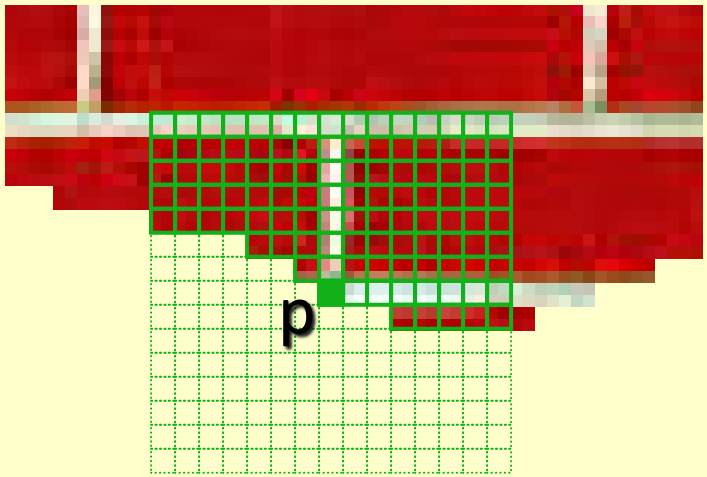

Synthesizing a pixel

non-parametric sampling

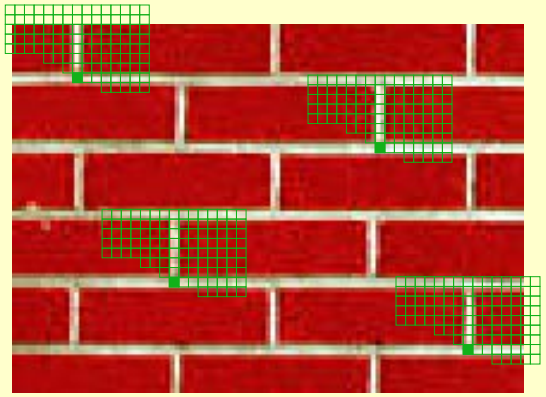

Input image

To synthesize p, just pick one match at random

Efros & Leung ’99

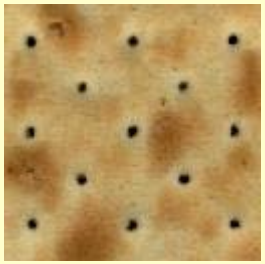

- The algorithm
  - Very simple
  - Surprisingly good results
  - Synthesis is easier than analysis!
  - …but very slow

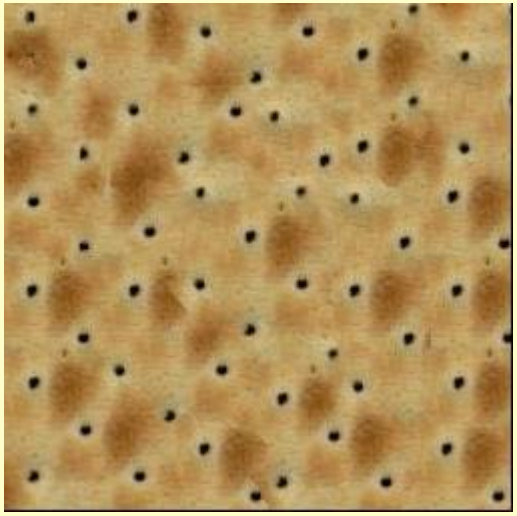

- Optimizations and Improvements
  - [Wei & Levoy,’00] (based on [Popat & Picard,’93])
  - [Harrison,’01]
  - [Ashikhmin,’01]
  - PatchMatch [Barnes et al. 2009]

Chaos Mosaic [Xu, Guo & Shum, ‘00]

input

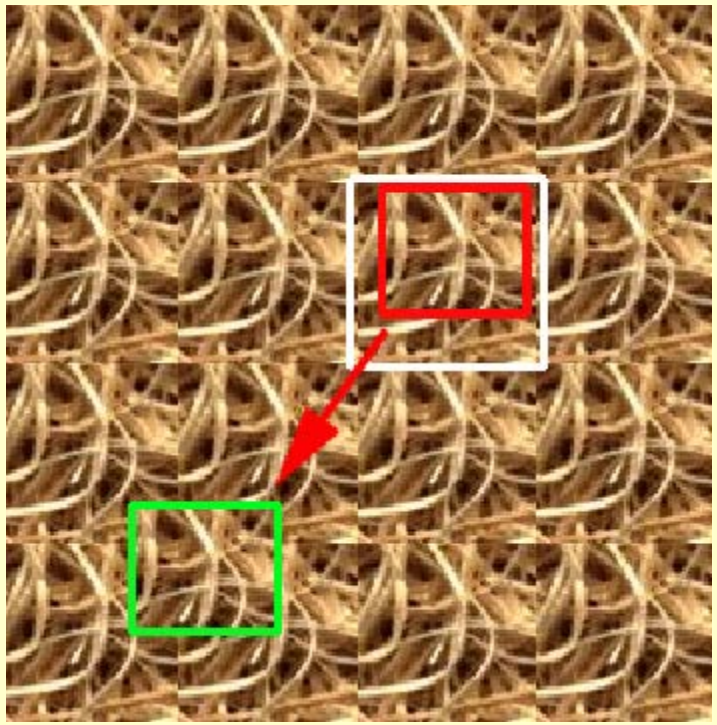

idea

result

- Process: 1) tile input image; 2) pick random blocks and place them in random locations 3) Smooth edges

Used in Lapped Textures [Praun et.al,’00]

Chaos Mosaic [Xu, Guo & Shum, ‘00]

input

result

Of course, doesn’t work for structured textures

Multi-Resolution Pyramids*

Example texture pyramid

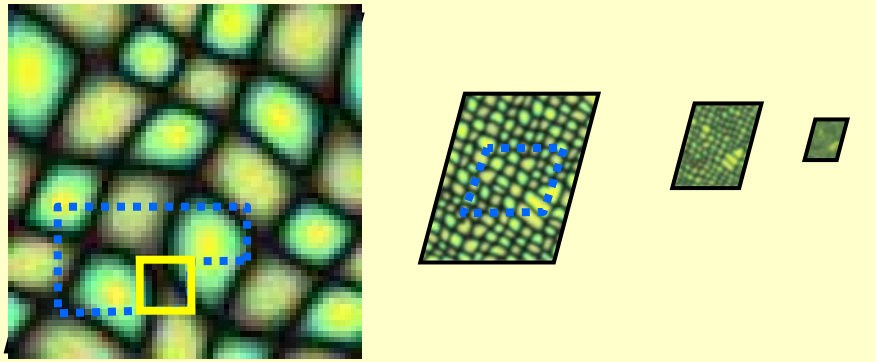

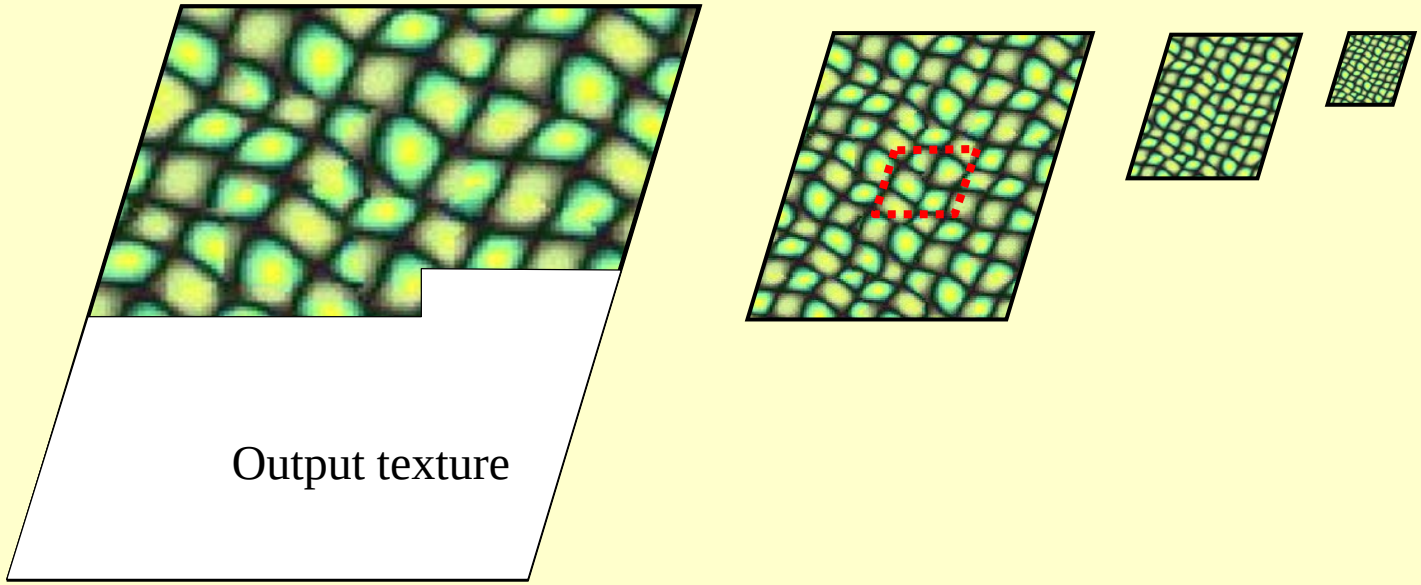

* L.-Y.Wei, M.Levoy; "*Fast Texture Synthesis using Tree-structured Vector Quantization*"; SIGGRAPH00.

Extension to 3D Textures

- Motion both in space and time
  - fire, smoke, ocean waves.
- How to synthesize?
  - extend 2D algorithm to 3D.

The Problems of Causal Scanning

- Scanning order:
  - Efros&Leung(1): Pixels with most neighbors.
  - Wei&Levoi(2): Raster scan.
- These are "causal" scans.

(1) A.A.Efros, T.K.Leung; "*Texture synthesis by non-parametric sampling*"; ICCV99.

(originally proposed by [Garber,'81])

(2) L.-Y.Wei, M.Levoy; "Fast Texture Synthesis using Tree-structured Vector Quantization"; SIGGRAPH00.

The Problems of Causal Scanning

- Can grow garbage.
- No natural means of refining synthesis.
- Cannot be parallelized.
- Problems are made worst for synthesis of 3D space-time volumes (a.k.a. video)...

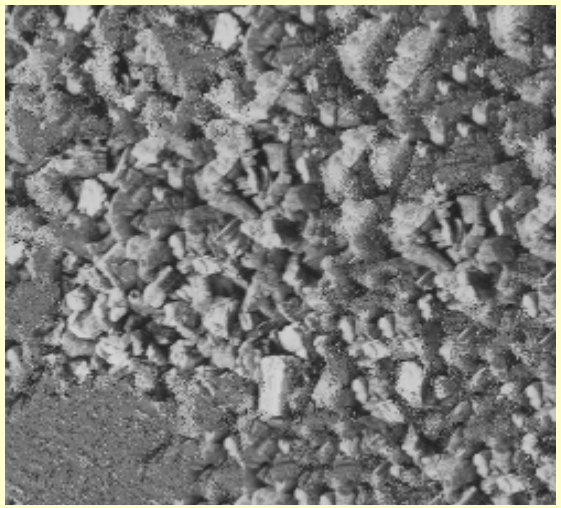

A.A.Efros, T.K.Leung; "*Texture synthesis by non-parametric sampling*"; ICCV99.

Image Quilting

- Observation: neighbor pixels are highly correlated

- Idea:
  - let's combine random block placement of Chaos Mosaic with spatial constraints of Efros & Leung.

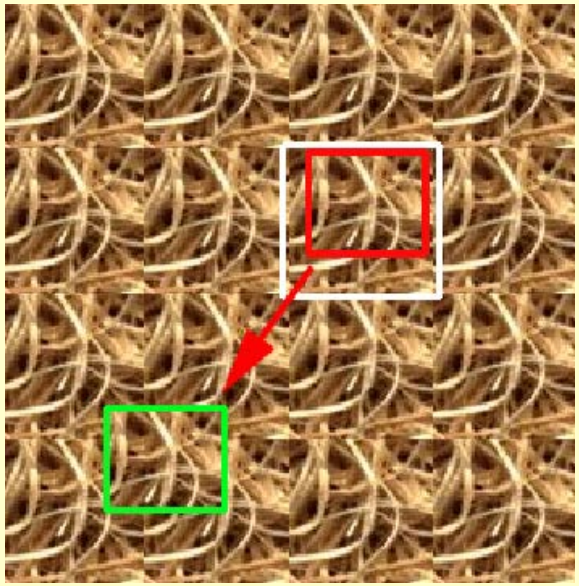

Efros & Leung ’99 extended

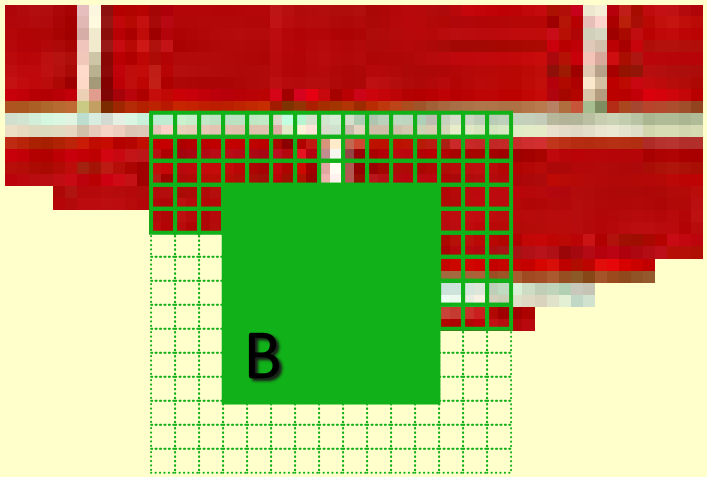

Synthesizing a block

non-parametric sampling

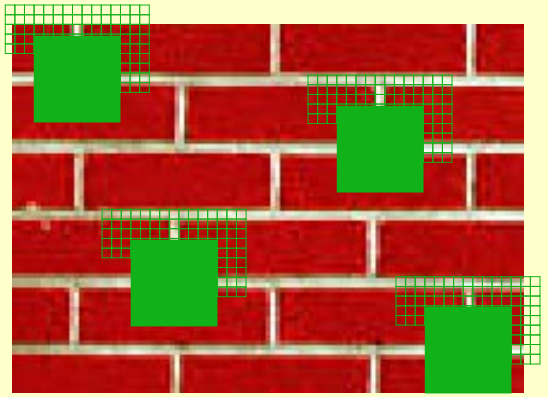

Input image

<u>Idea:</u> 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
- Not the same as multi-scale!

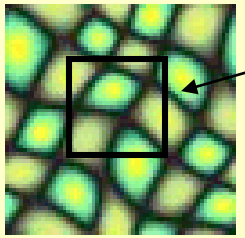

Input texture

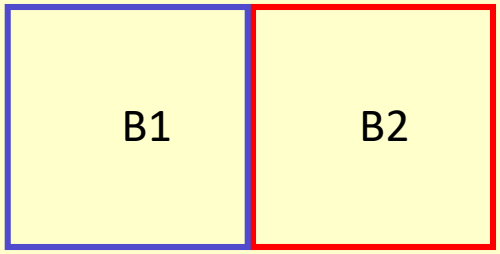

Random placement of blocks

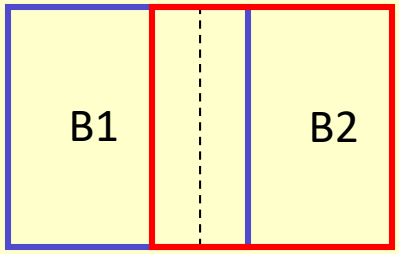

Neighboring blocks constrained by overlap

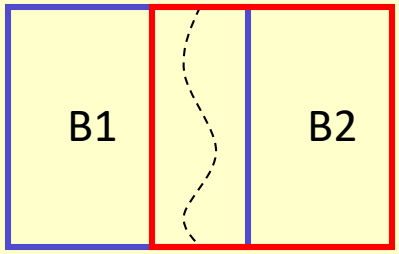

Minimal error boundary cut

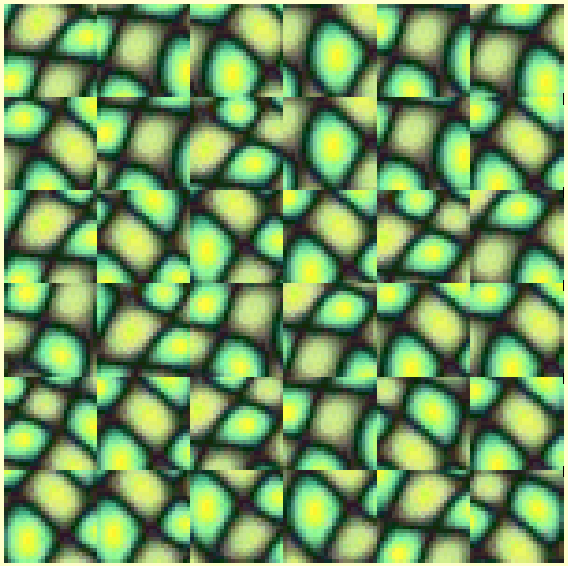

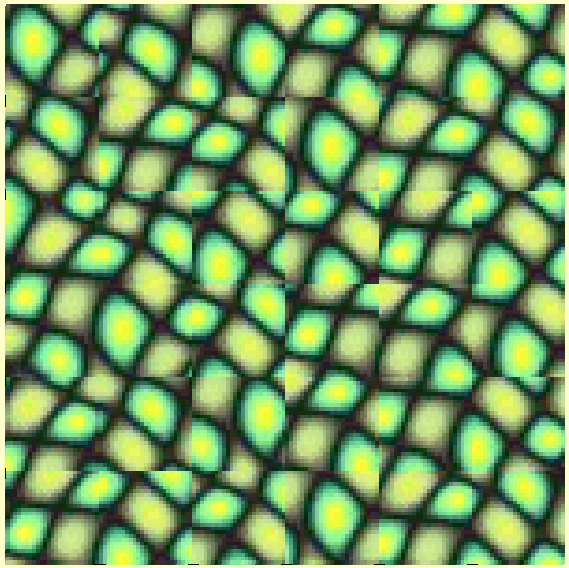

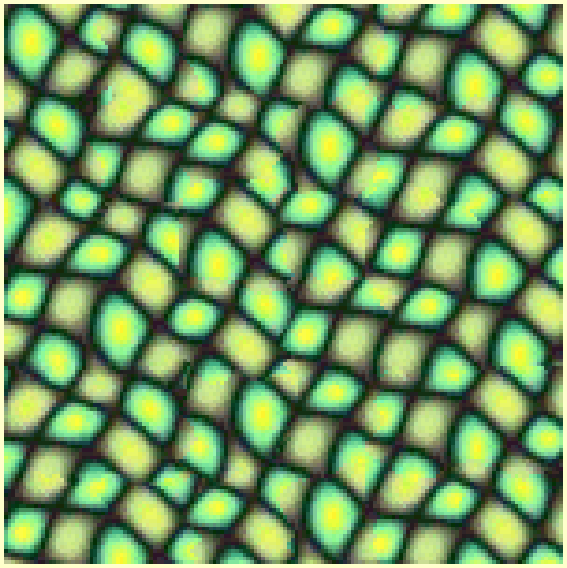

Minimal error boundary

overlapping blocks

vertical boundary

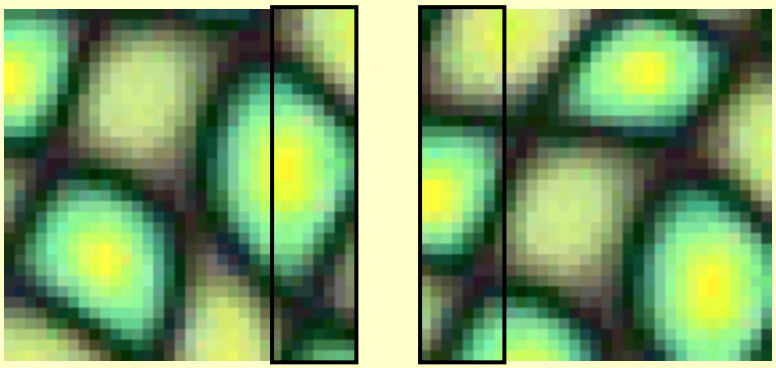

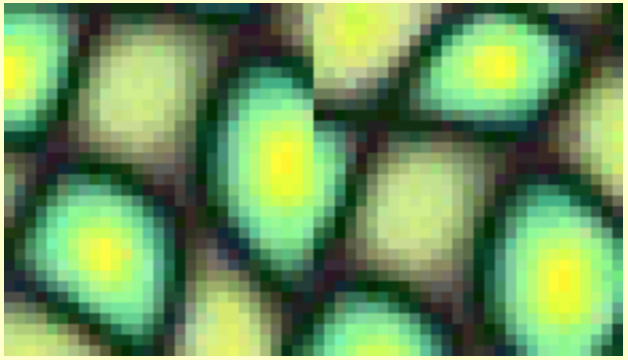

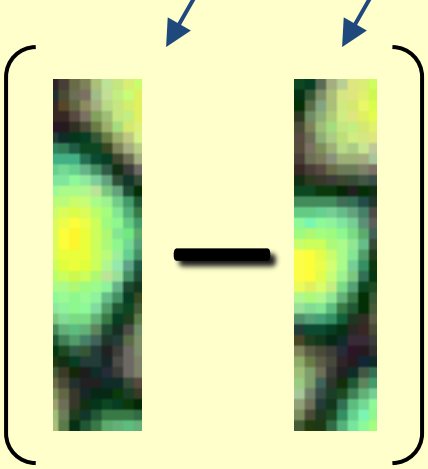

2

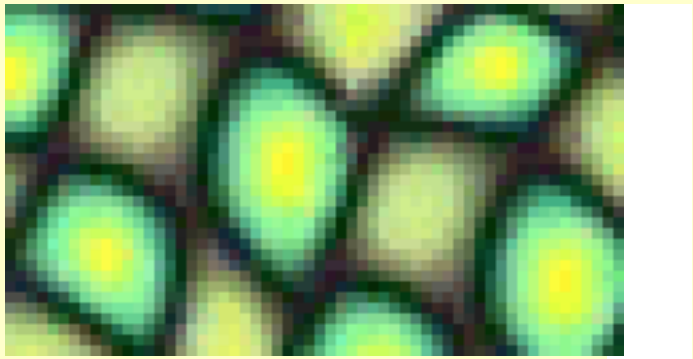

overlap error

min. error boundary

Our Philosophy

- The “Corrupt Professor’s Algorithm”:
  - Plagiarize as much of the source image as you can
  - Then try to cover up the evidence
- Rationale:
  - Texture blocks are by definition correct samples of texture so problem only connecting them together

Image Quilting Algorithm

- Pick size of block and size of overlap
- Synthesize blocks in raster order

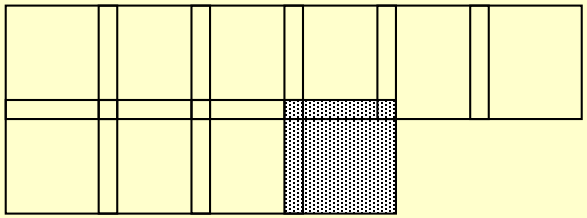

- Search input texture for block that satisfies overlap constraints (above and left)
  - Easy to optimize using NN search [Liang et.al., '01]
- Paste new block into resulting texture
  - use dynamic programming to compute minimal error boundary cut

See https://www.youtube.com/watch?v=t6DzioKuVEs

Video

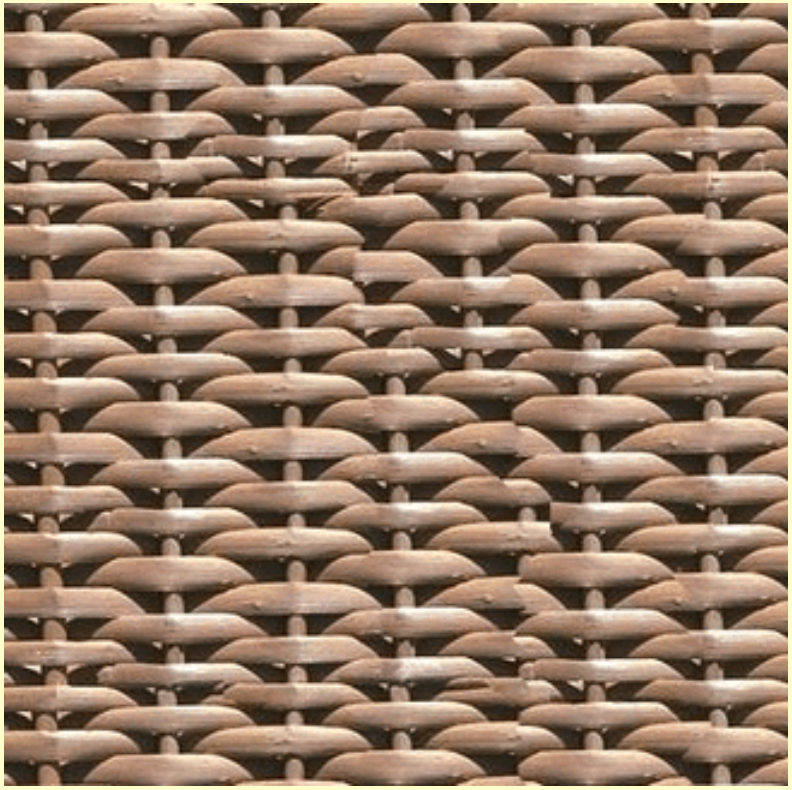

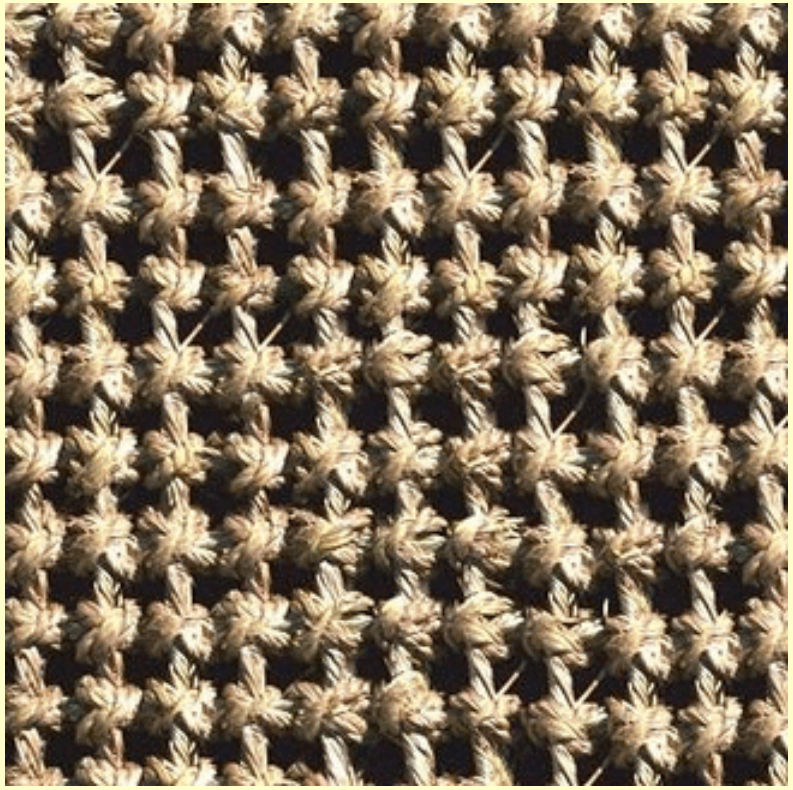

College Inn
Chicken Broth

College Inn
Chicken Broth

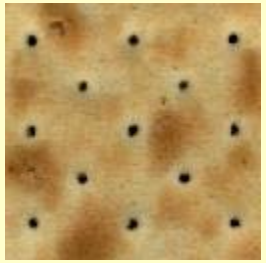

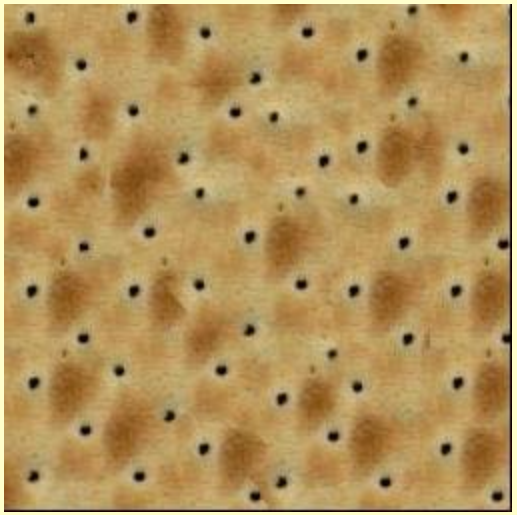

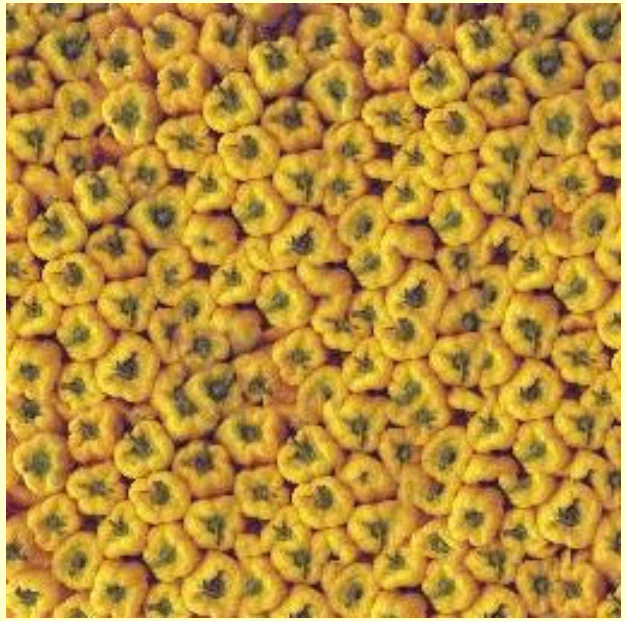

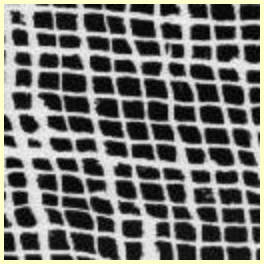

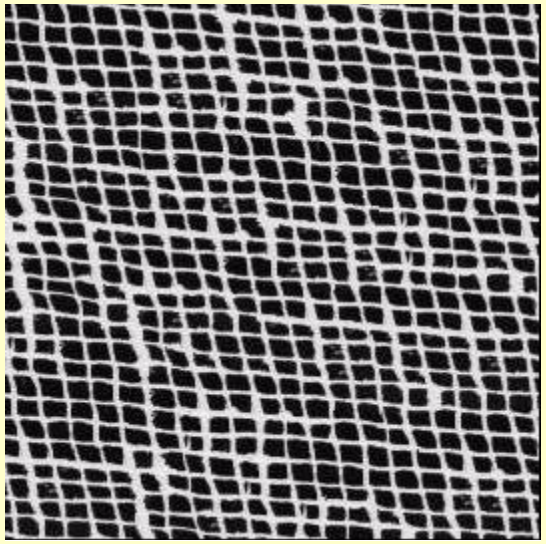

Failures

(Chernobyl Harvest)

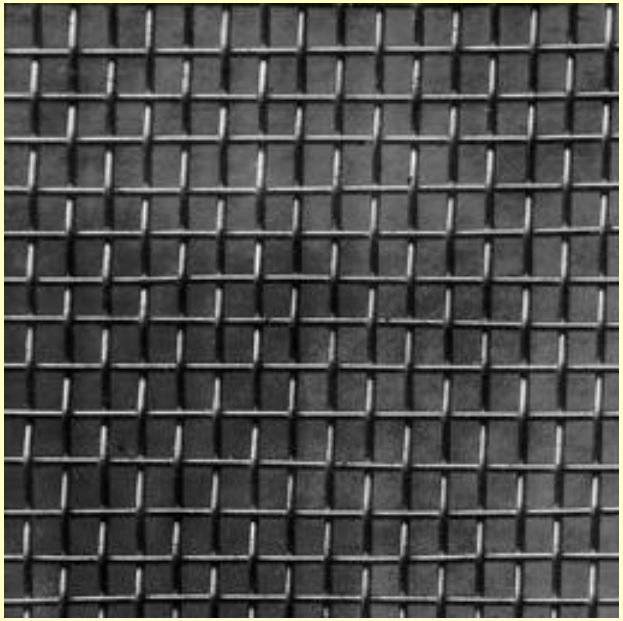

input image

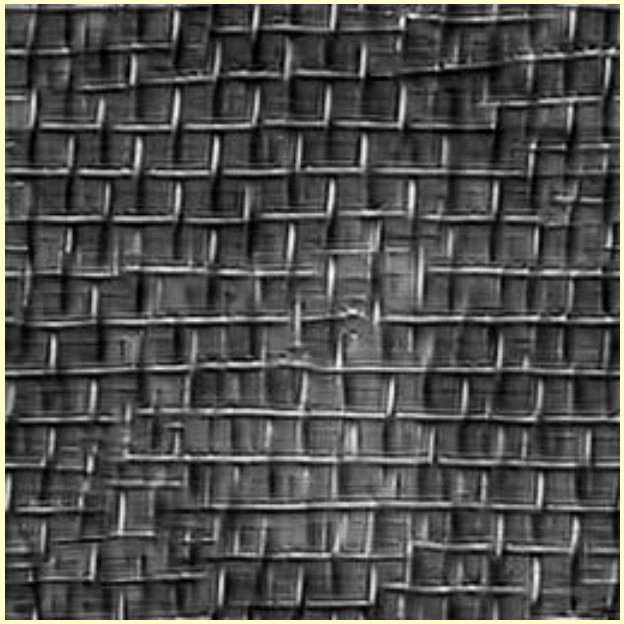

Portilla & Simoncelli

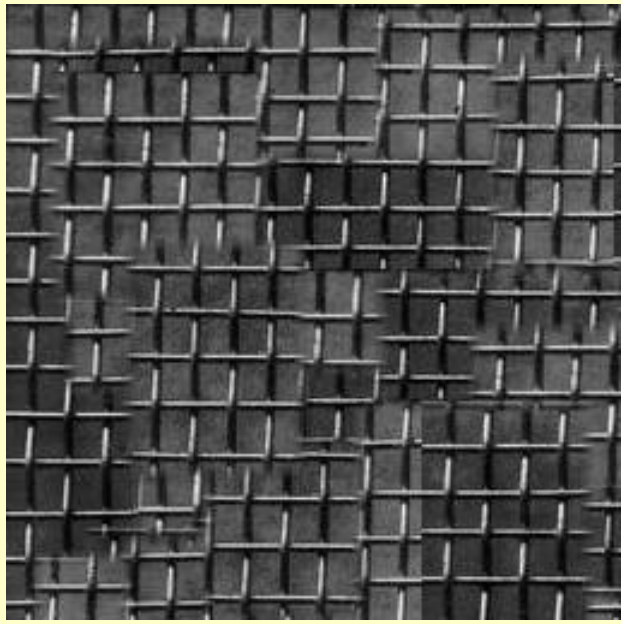

Xu, Guo & Shum

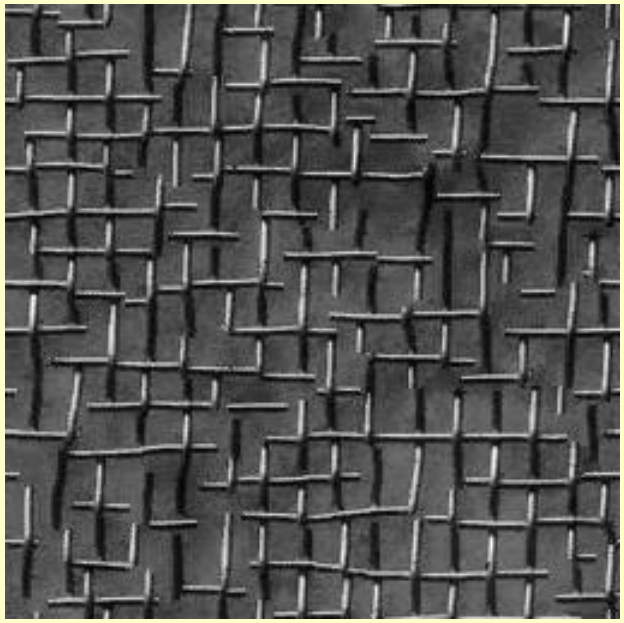

Wei & Levoy

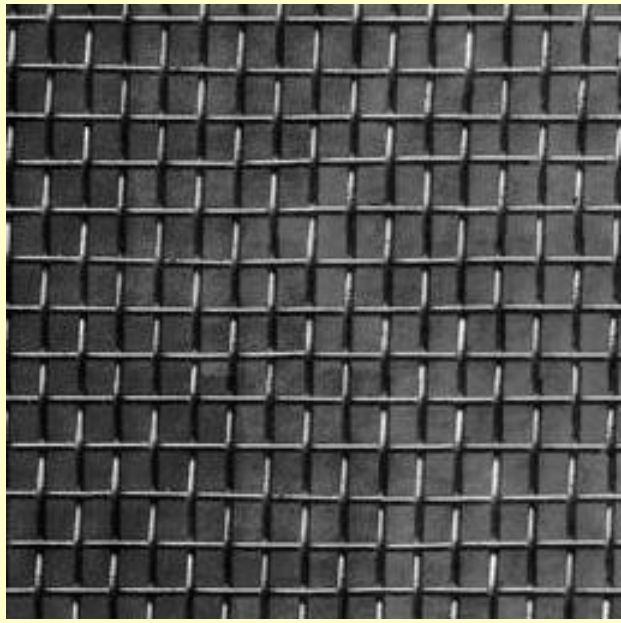

Image Quilting

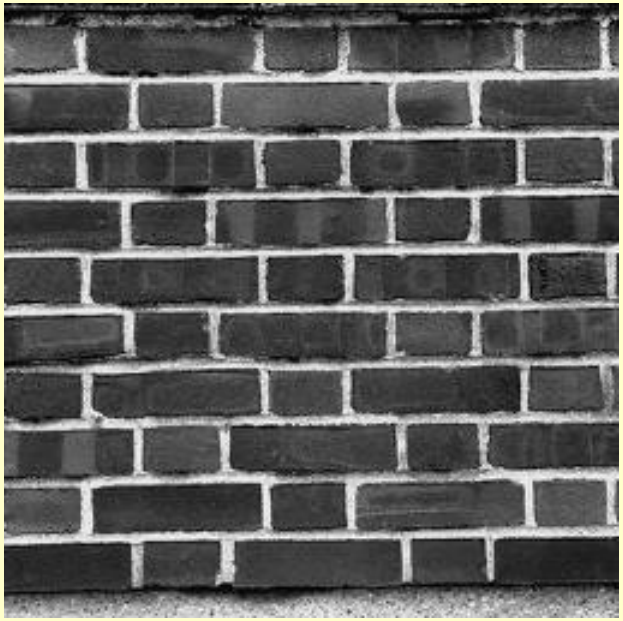

input image

Portilla & Simoncelli

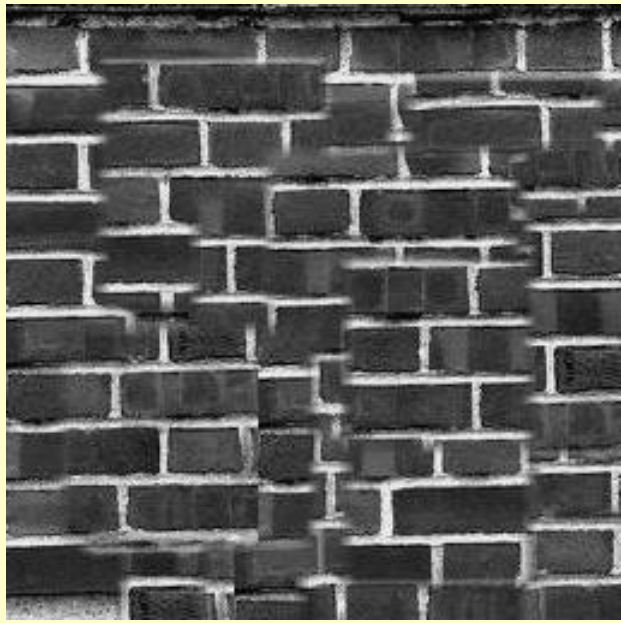

Xu, Guo & Shum

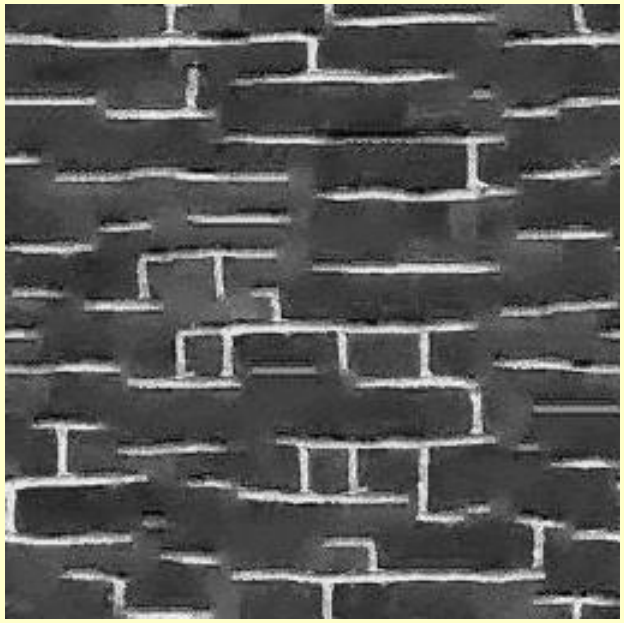

Wei & Levoy

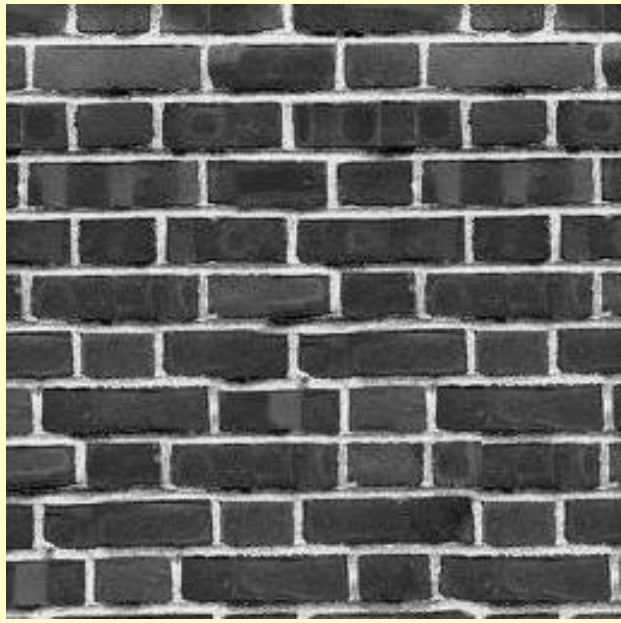

Image Quilting

Homage to Shannon!

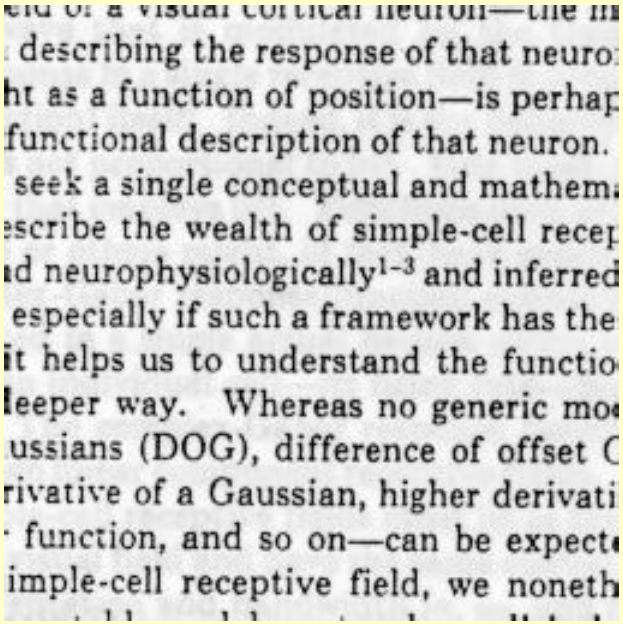

input image

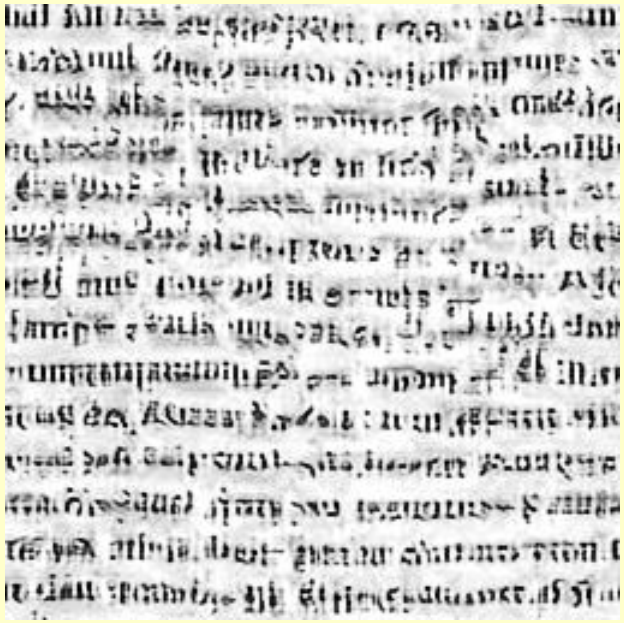

Portilla & Simoncelli

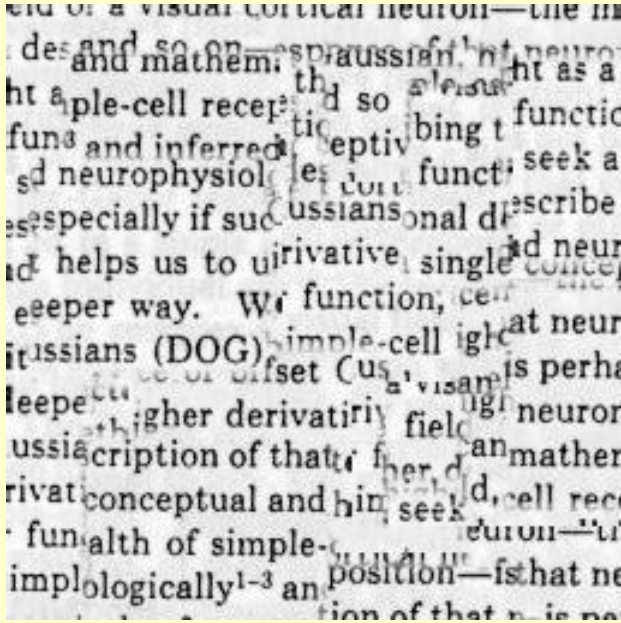

Xu, Guo & Shum

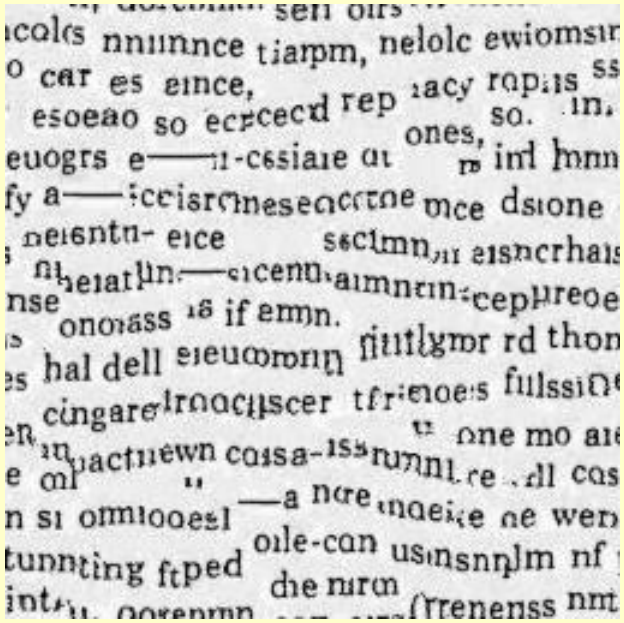

Wei & Levoy

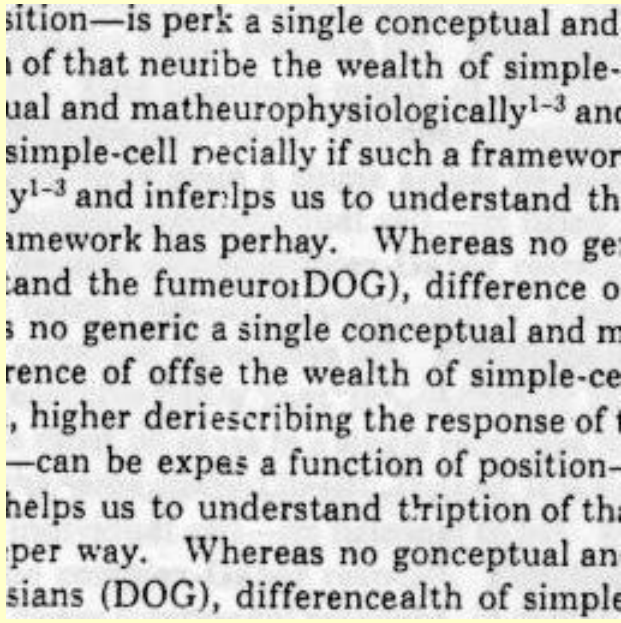

Image Quilting

Synthesis in Action

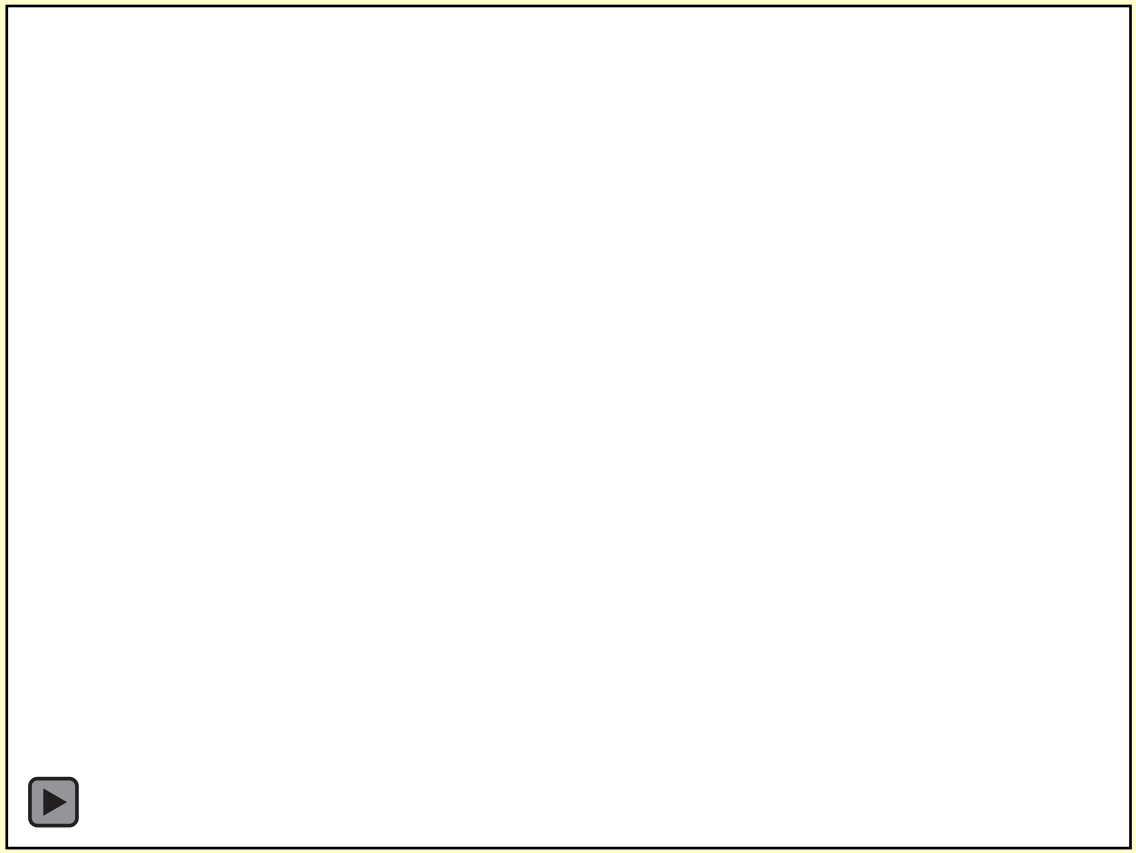

Bush campaign digitally altered TV ad

President Bush's campaign acknowledged Thursday that it had digitally altered a photo that appeared in a national cable television commercial. In the photo, a handful of soldiers were multiplied many times.

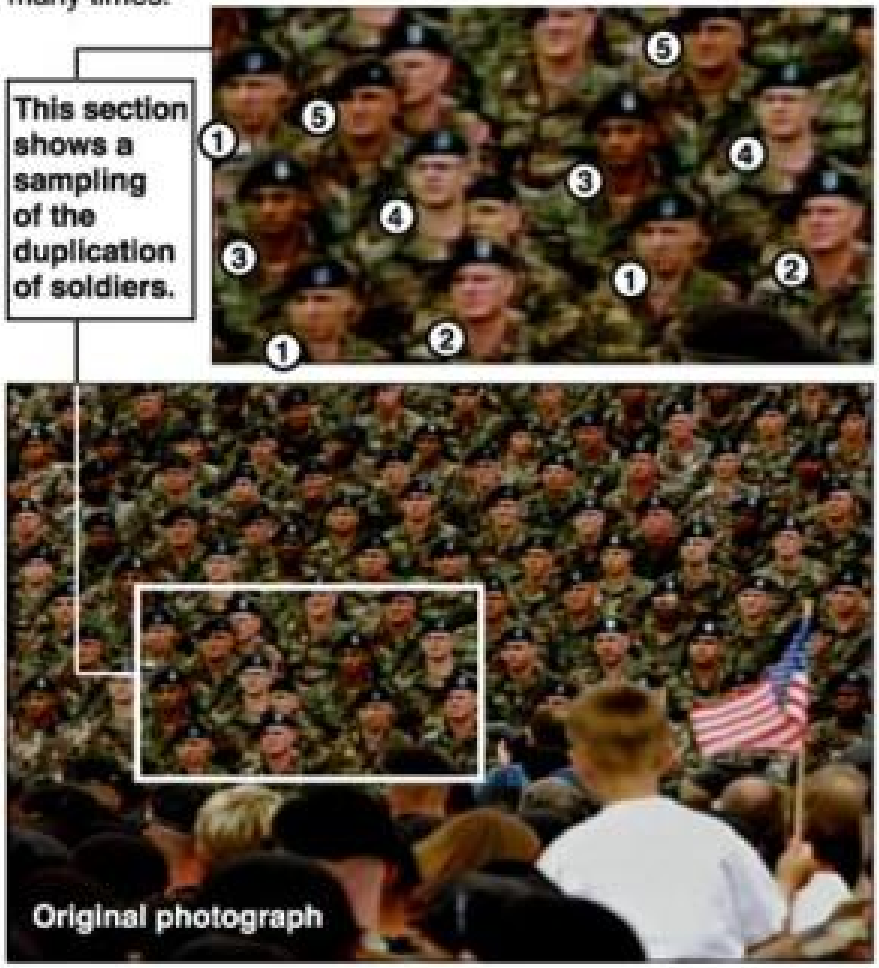

AP

Synthesis by Optimization

The pixels are all synthesized in parallel, not in a particular order

Iterate until convergence

$$E = \sum_{v} \sum_{i \in \{x,y,z\}} \left\| \mathbf{s}_{v,i} - \mathbf{e}_{v,i} \right\|^{r}$$

Y. Wexler E. Shechtman M. Irani;
“Space-Time Video Completion"; CVPR’04.

V.Kwatra I.Essa A.Bobick N.Kwatra;
“Texture Optimization for Example-based Synthesis"; SIGGRAPH’05.

Synthesis by Optimization

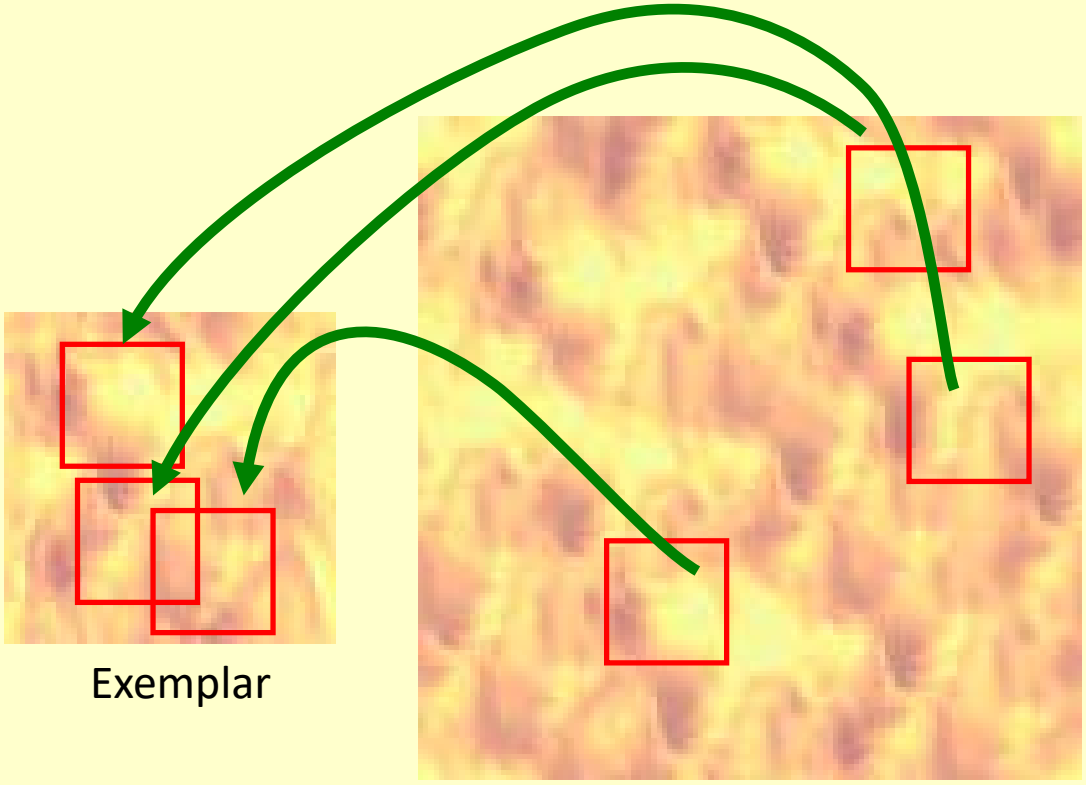

Synthesis by Optimization

Patches overlap!!

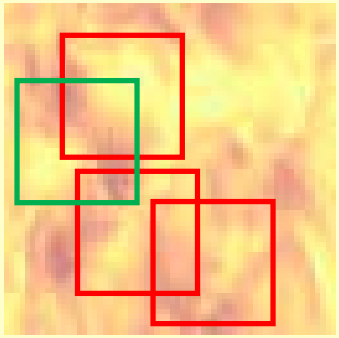

Exemplar

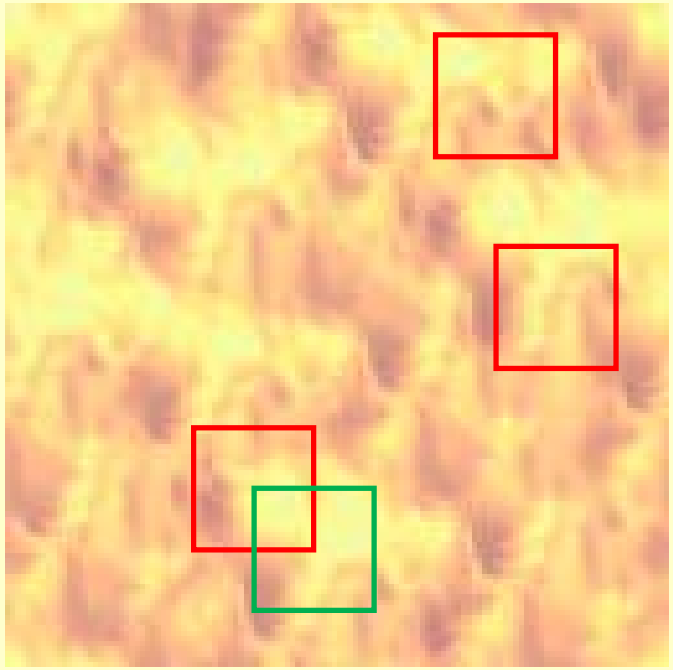

Synthesized Texture

Synthesis by Optimization

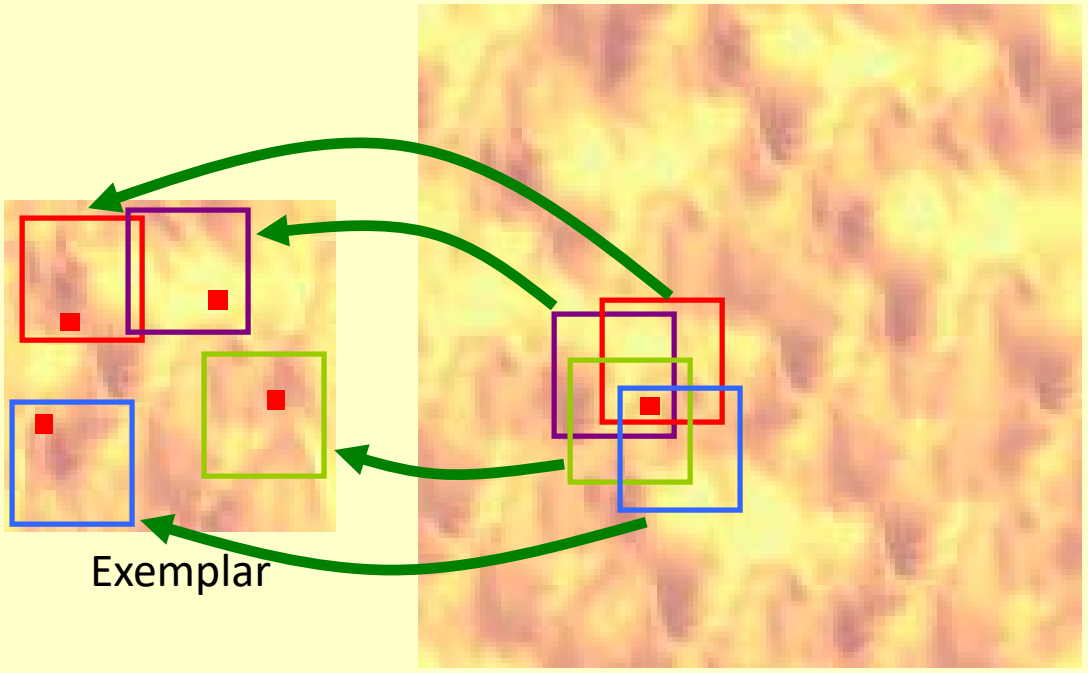

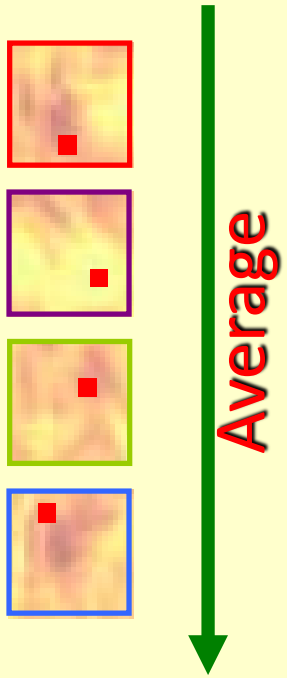

Histogram Matching, Kopf et. al SIG2006

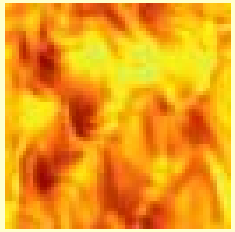

Exemplar

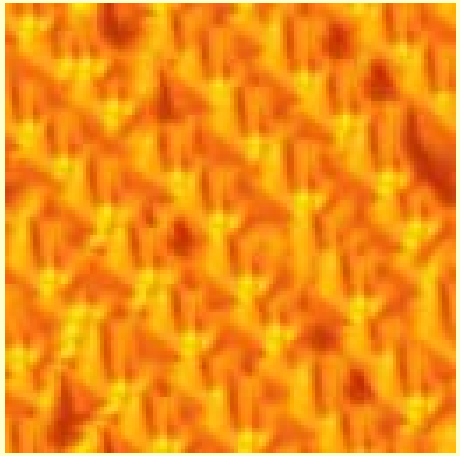

Synthesis

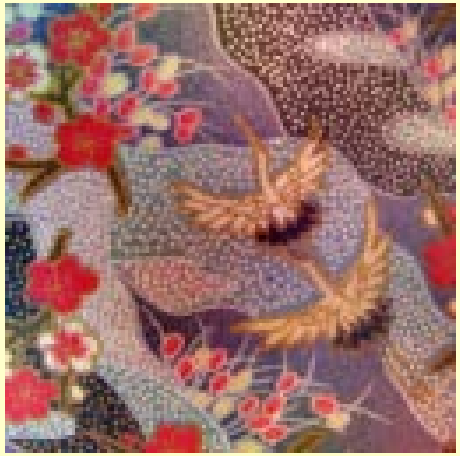

Exemplar

Synthesis

Histogram Matching

Exemplar

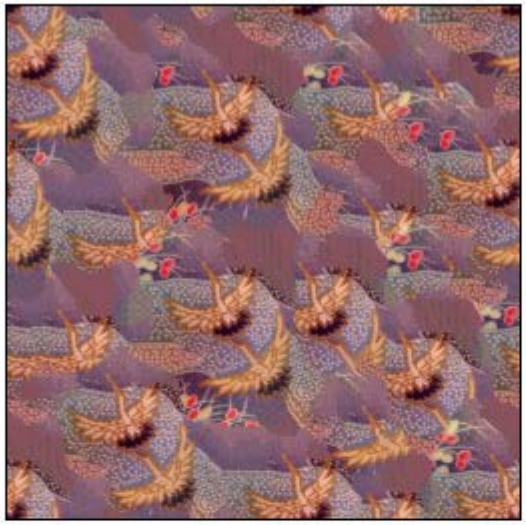

Without Histogram Matching

With Histogram Matching

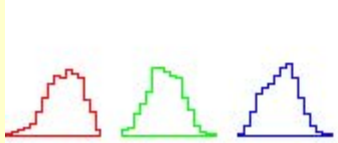

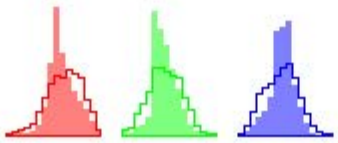

Histogram Matching

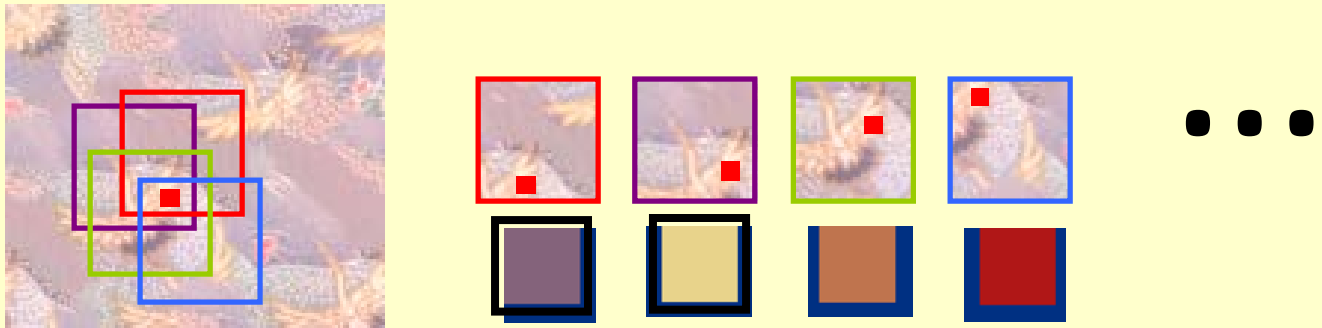

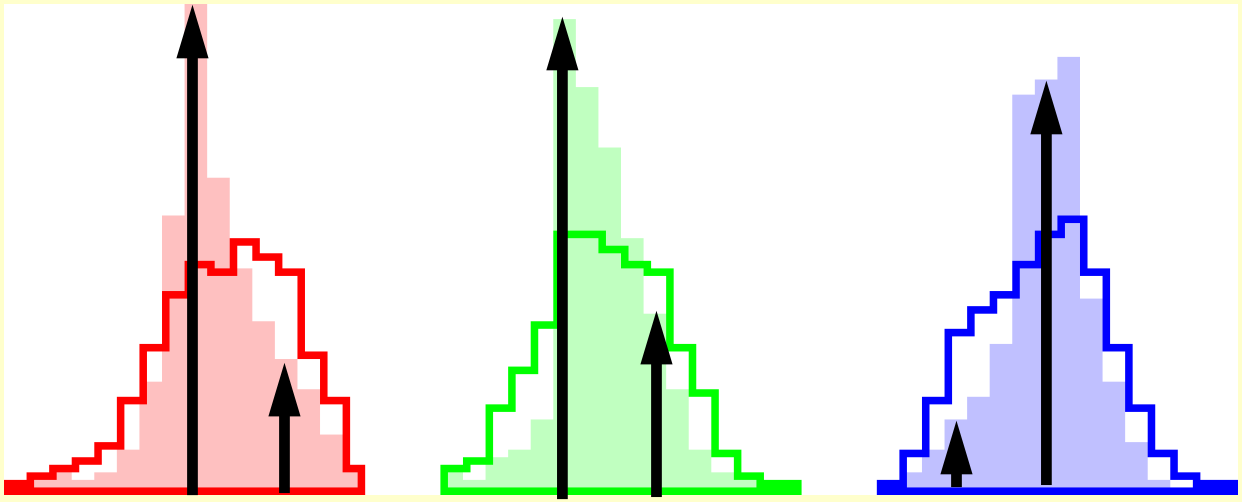

Histogram Matching

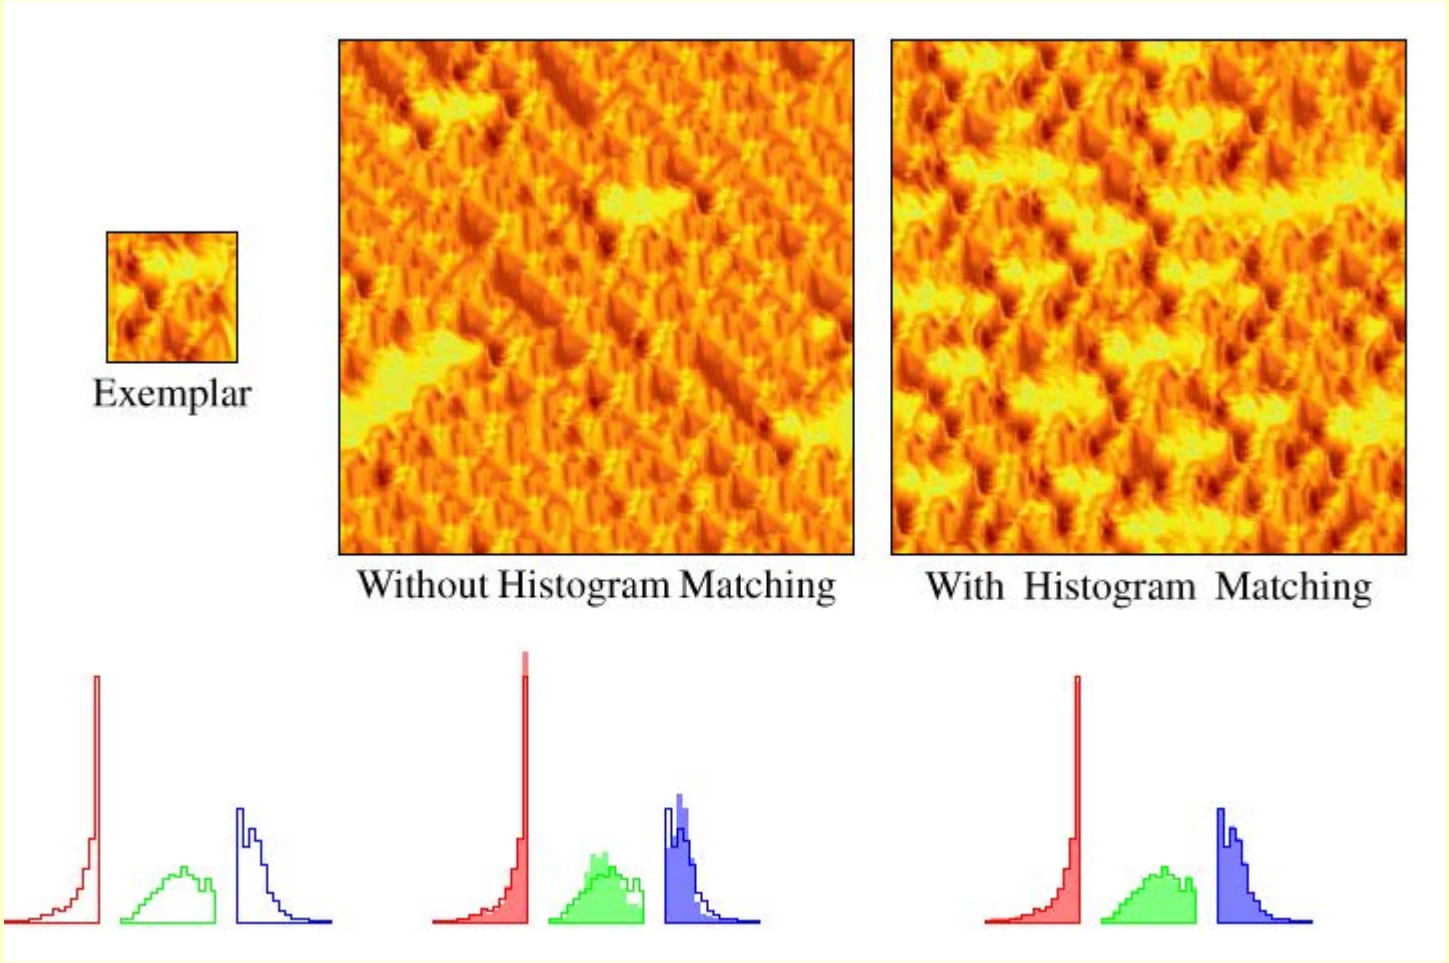

Histogram Matching

Exemplar

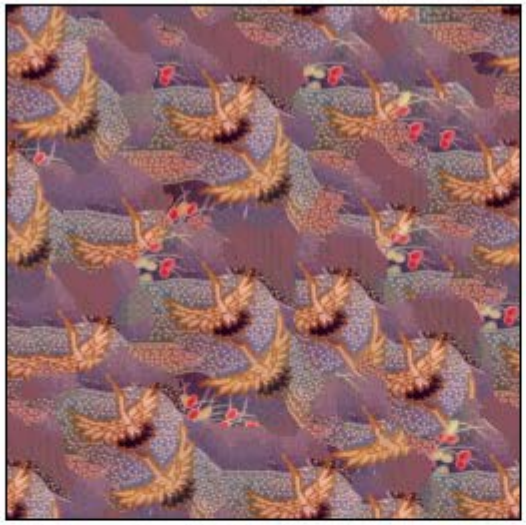

Without Histogram Matching

With Histogram Matching

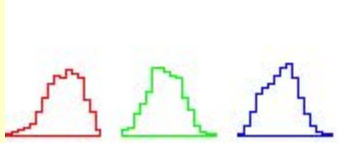

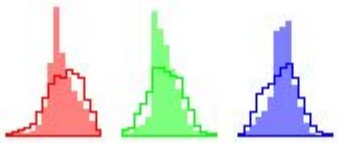

Application: Texture Transfer

- Try to explain one object with bits and pieces of another object:

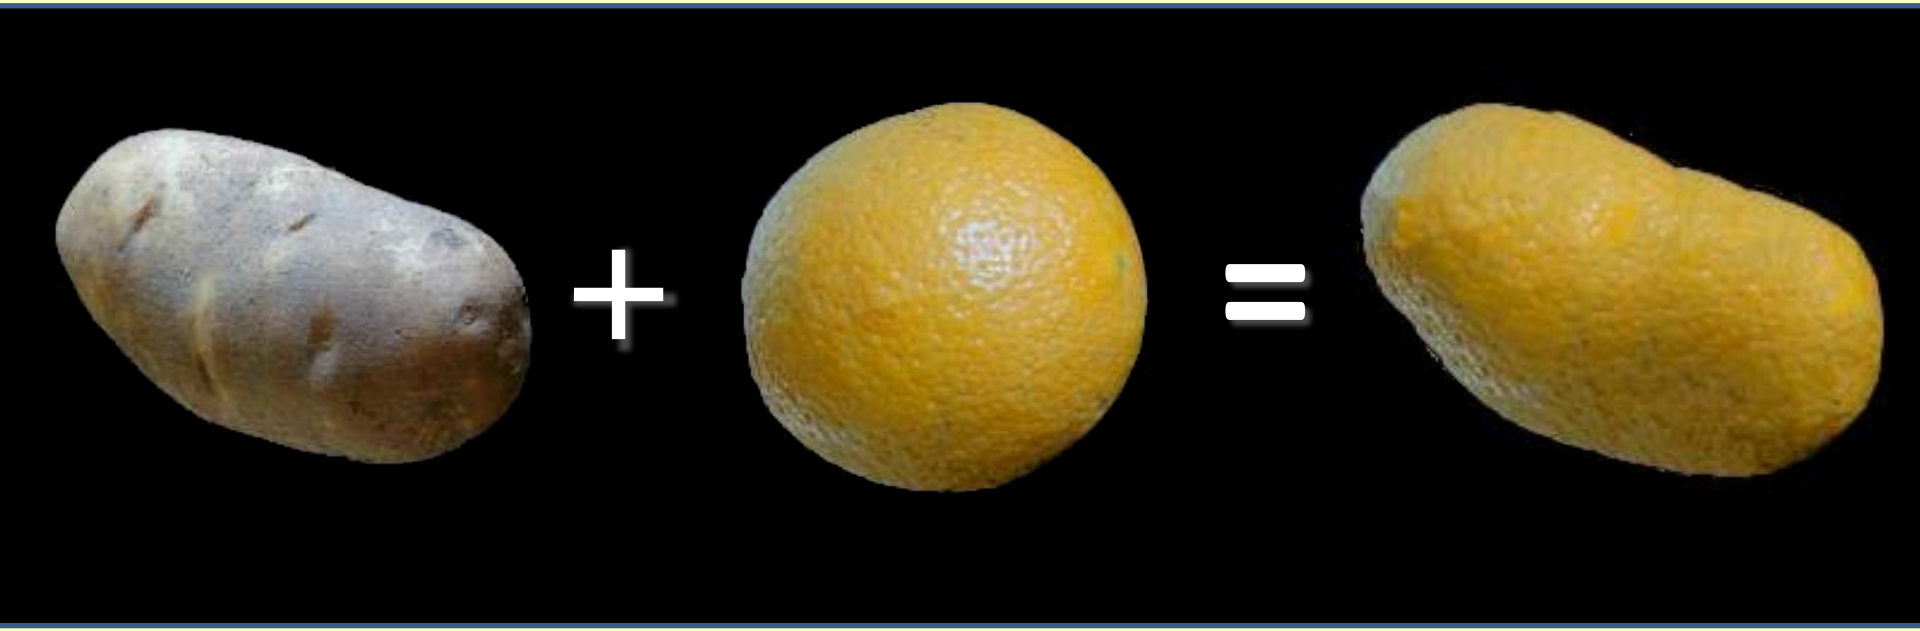

Texture Transfer

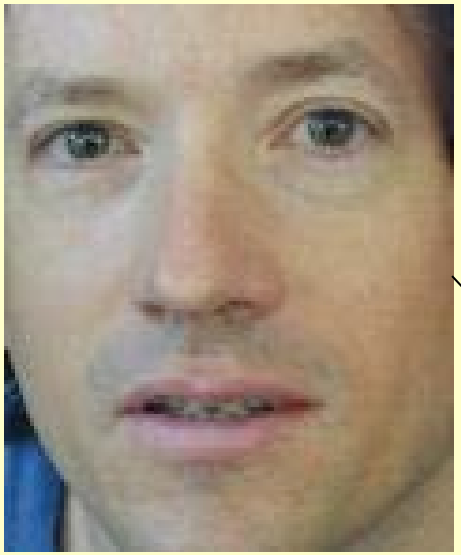

Constraint

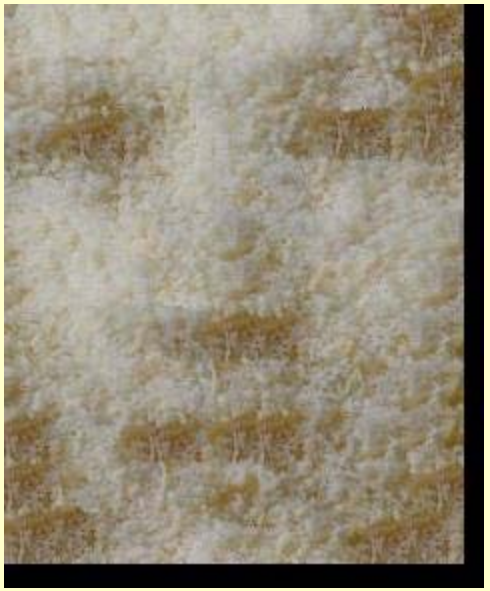

Texture sample

Texture Transfer

- Take the texture from one image and “paint” it onto another object

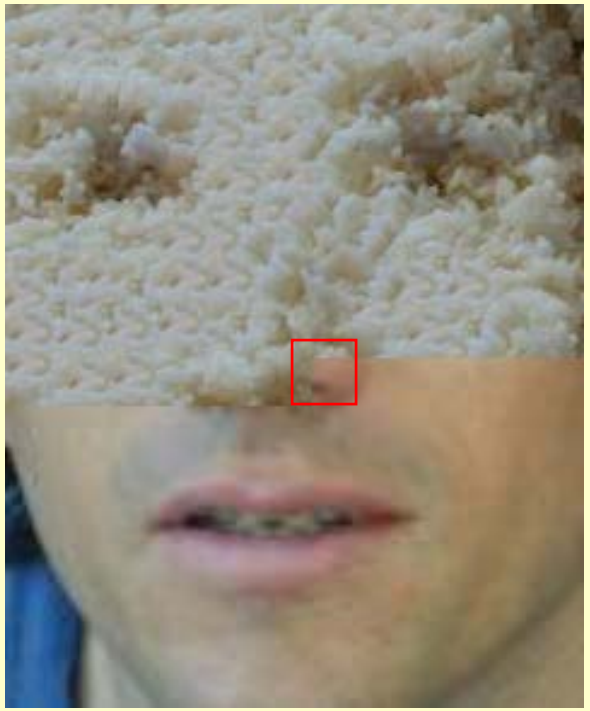

Same as texture synthesis, except an additional constraint:

1. Consistency of texture
2. Similarity to the image being “explained”

+

=

Source texture

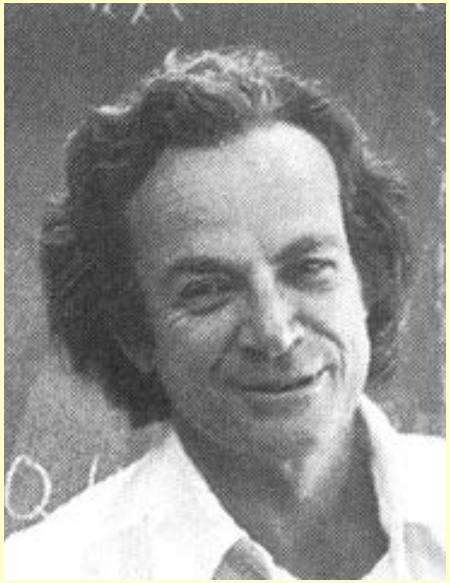

Target image

Source correspondence image

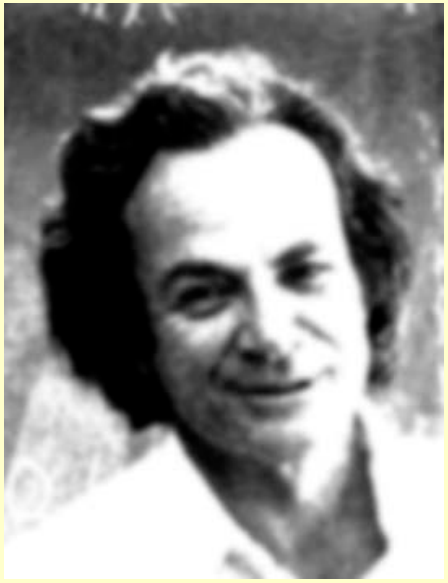

Target correspondence image

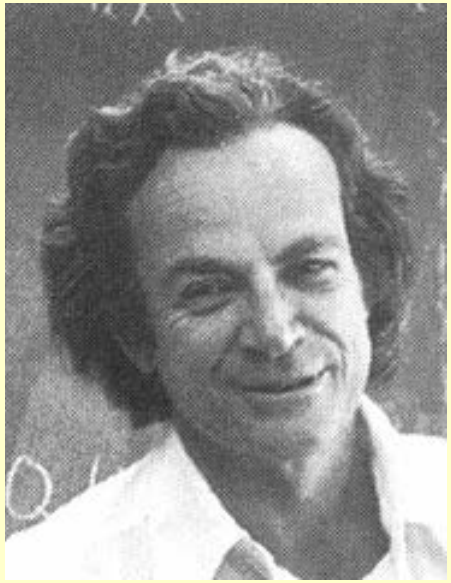
+

=

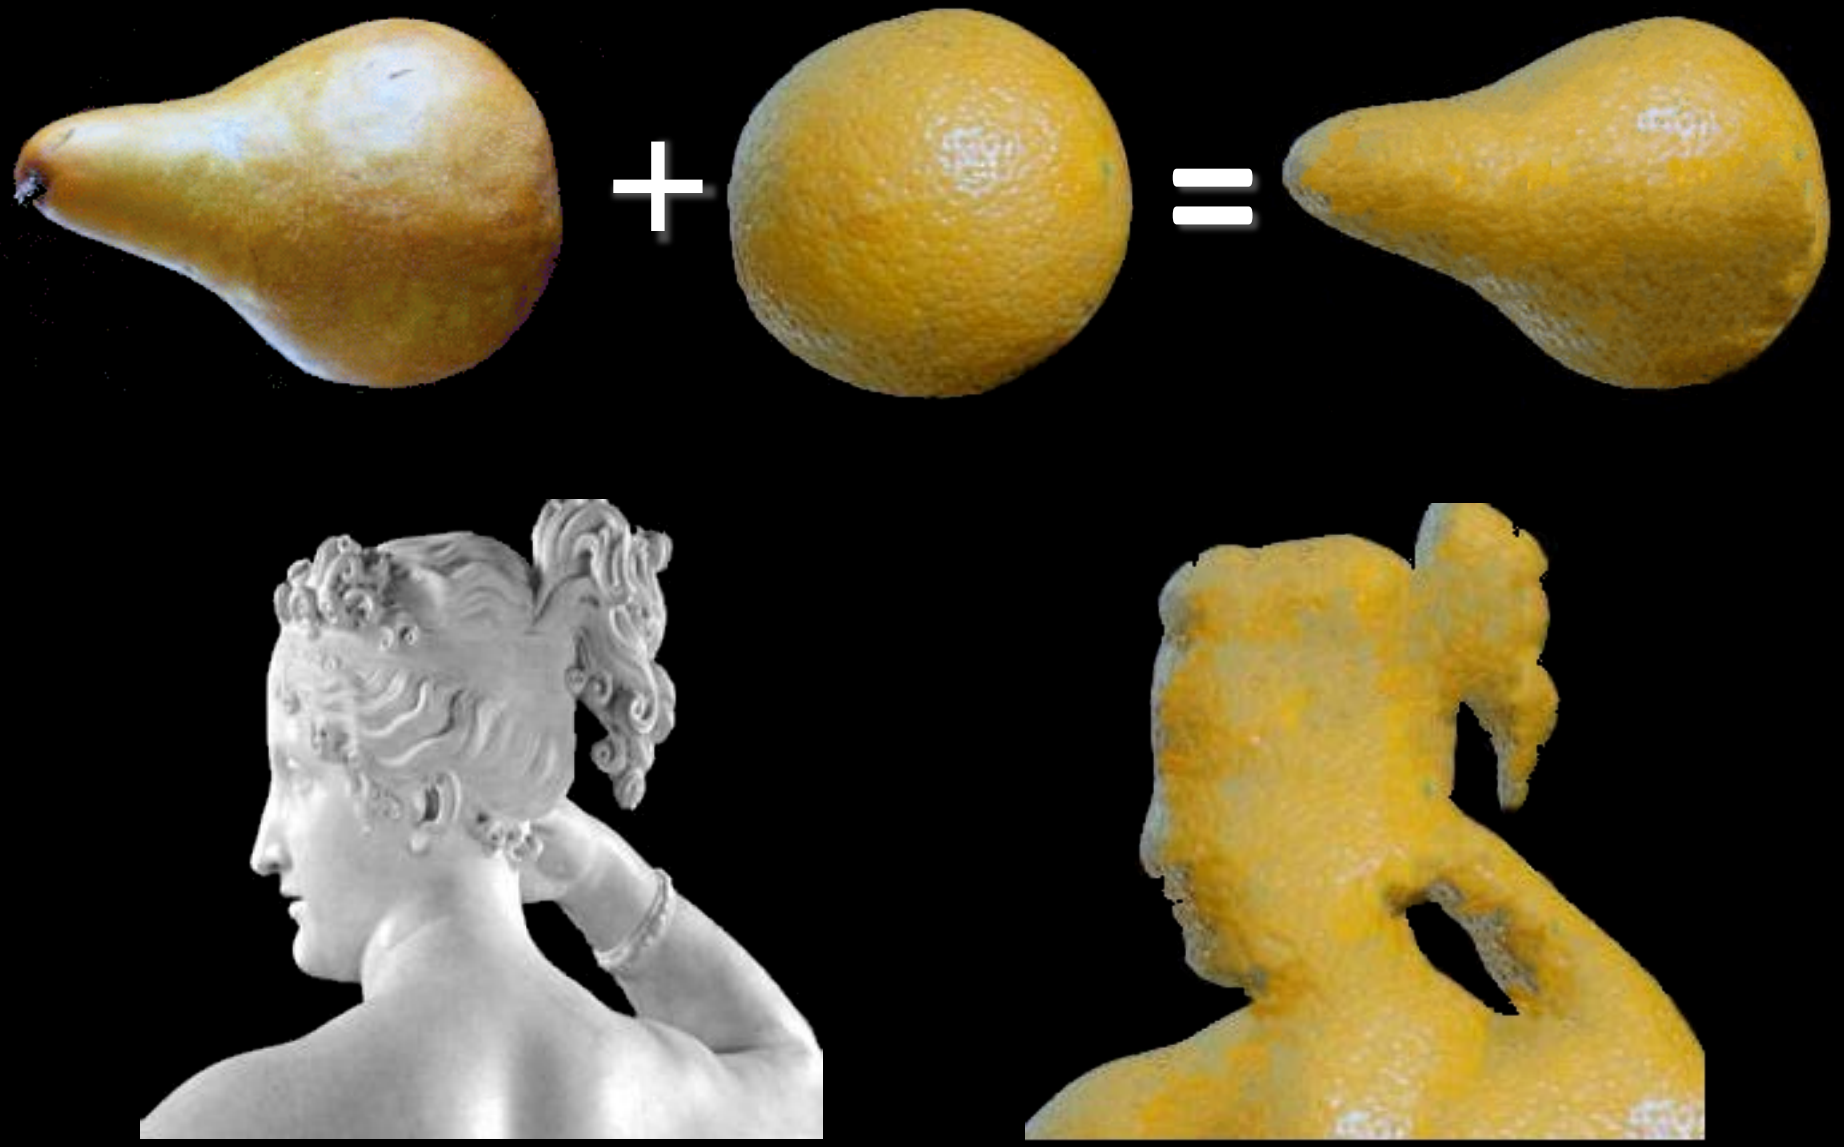

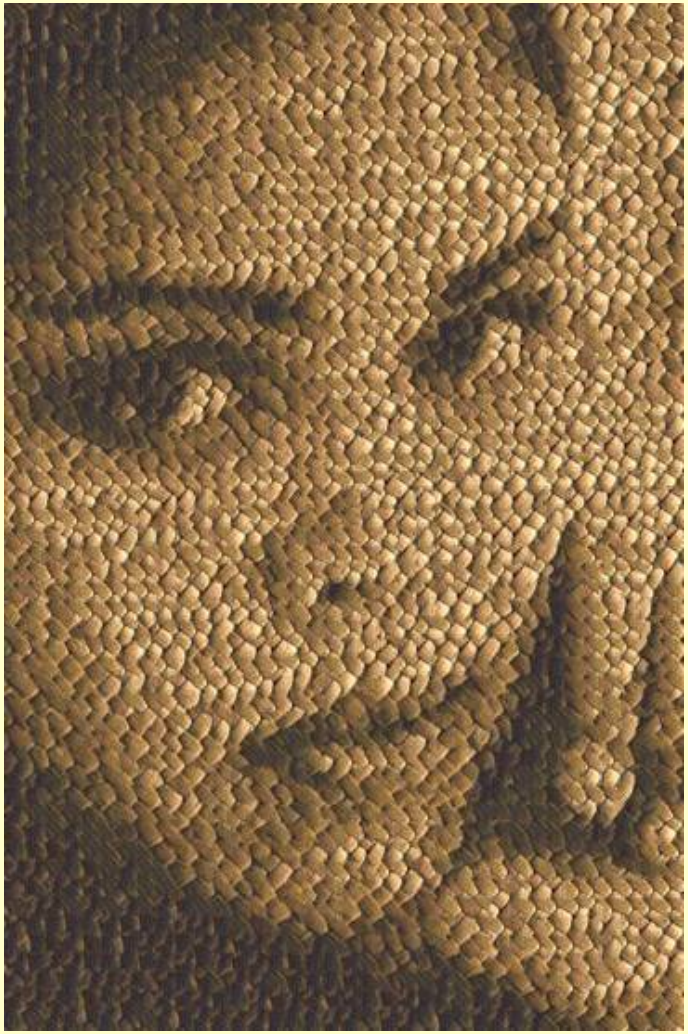

Image analogies (filter by example)

A to A' like B to B'

$A_1,\ldots,A_n : A_1',\ldots,A_n' :: B : B'$

texture segmentation

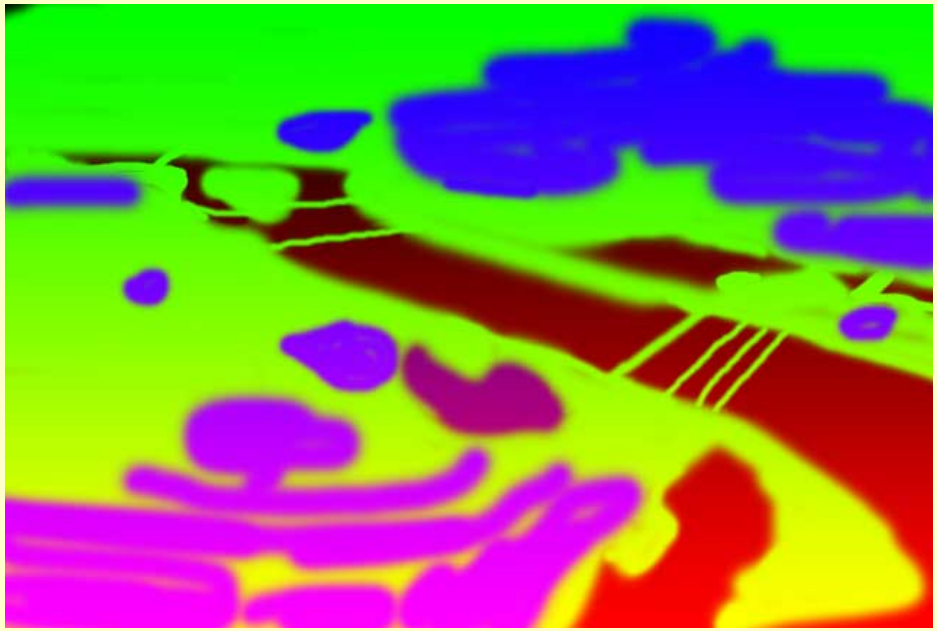

input

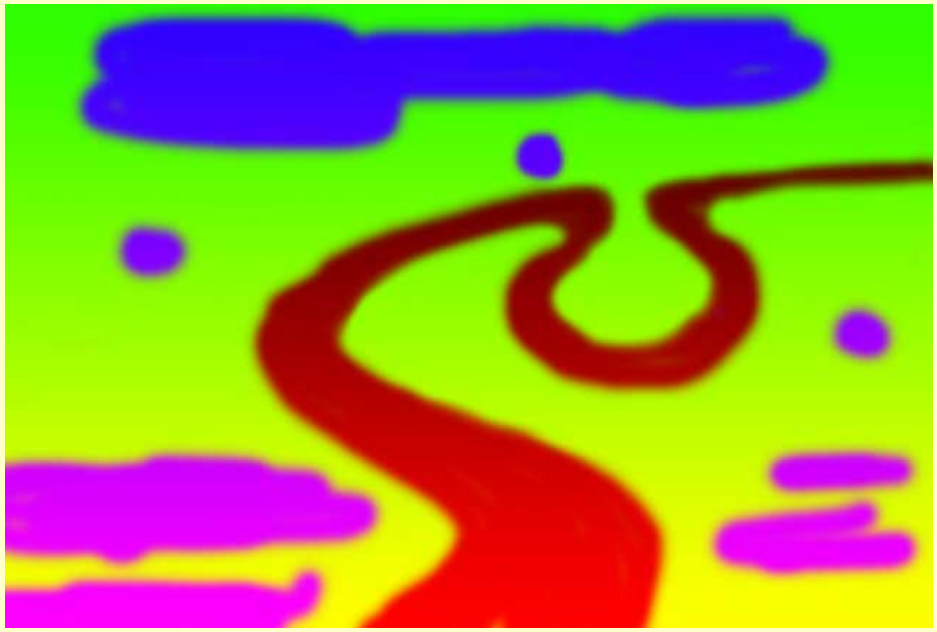

drawing with color coded textures

output

Applications – Artistic Filters (Cont.)

Target Pairs:

Source Pair:

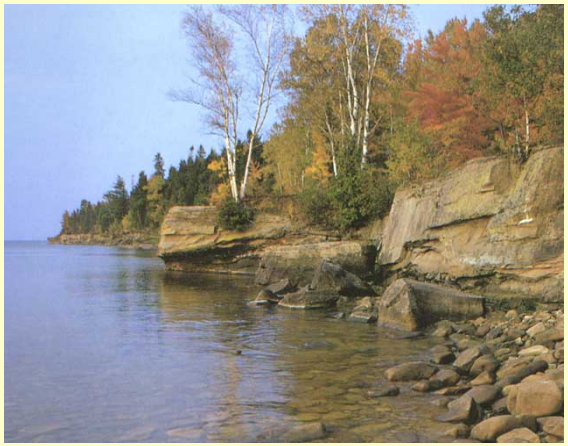

“Texture By Numbers”

- By color-labeling source image parts a realistic synthesized image can be created

Video

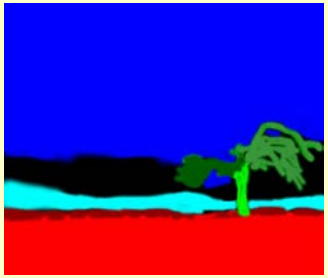

A

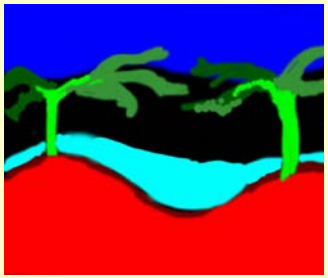

B

A`

B`

Fragment-based Image Completion (SIGGRAPH'03)

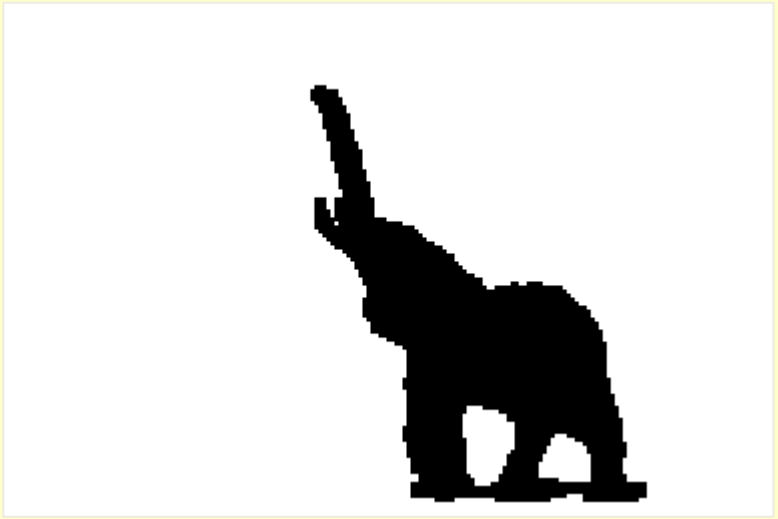

Fragment-based Image Completion (SIGGRAPH'03)

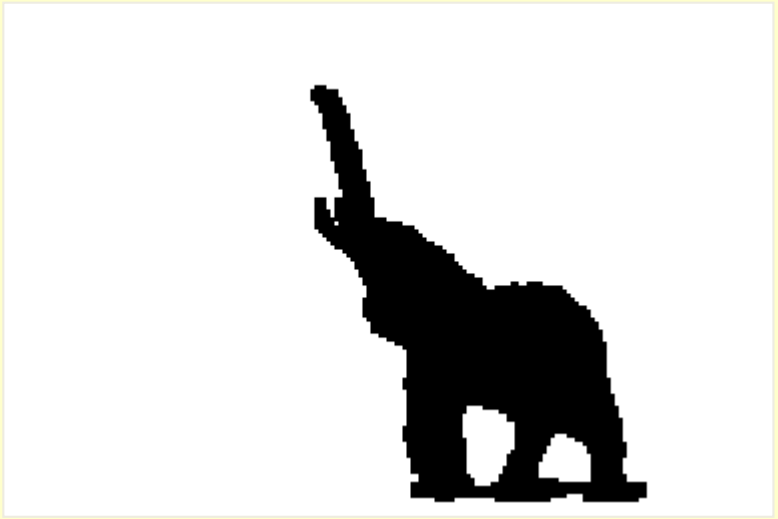

Completion process

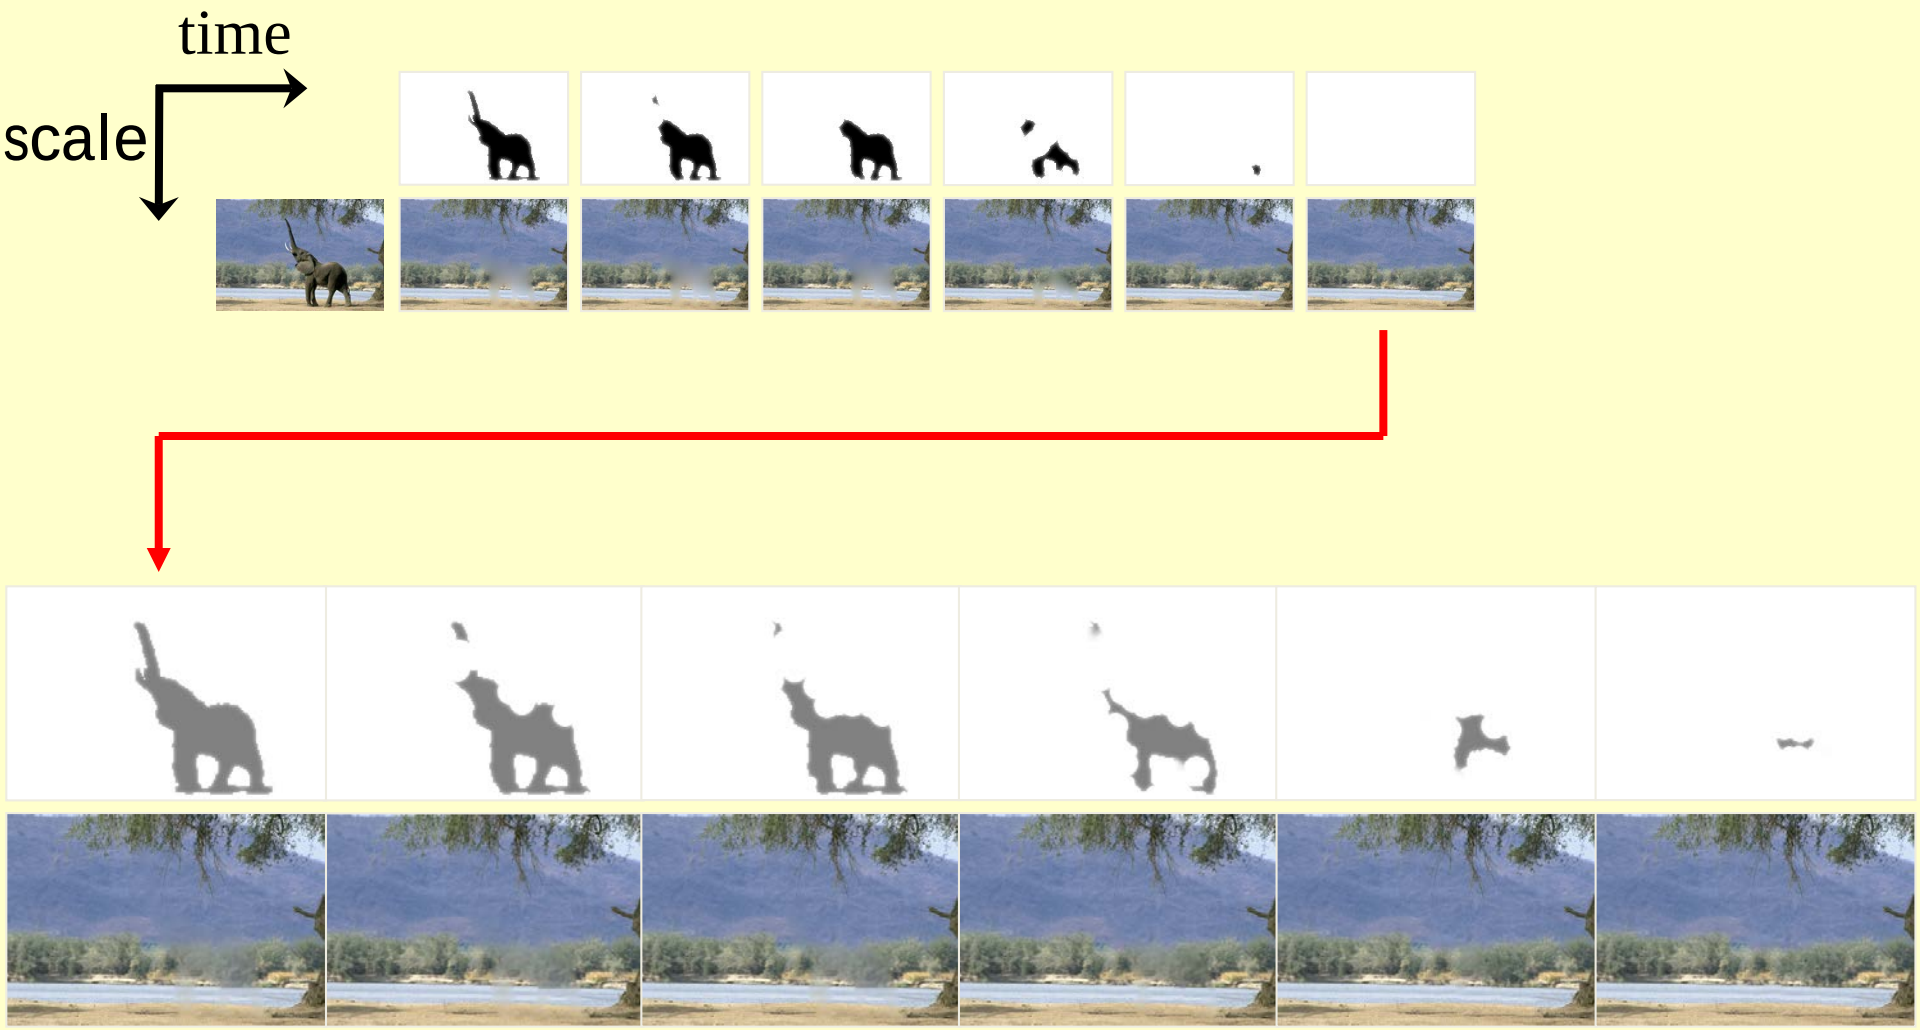

confidence and color at different time steps and scales

Results

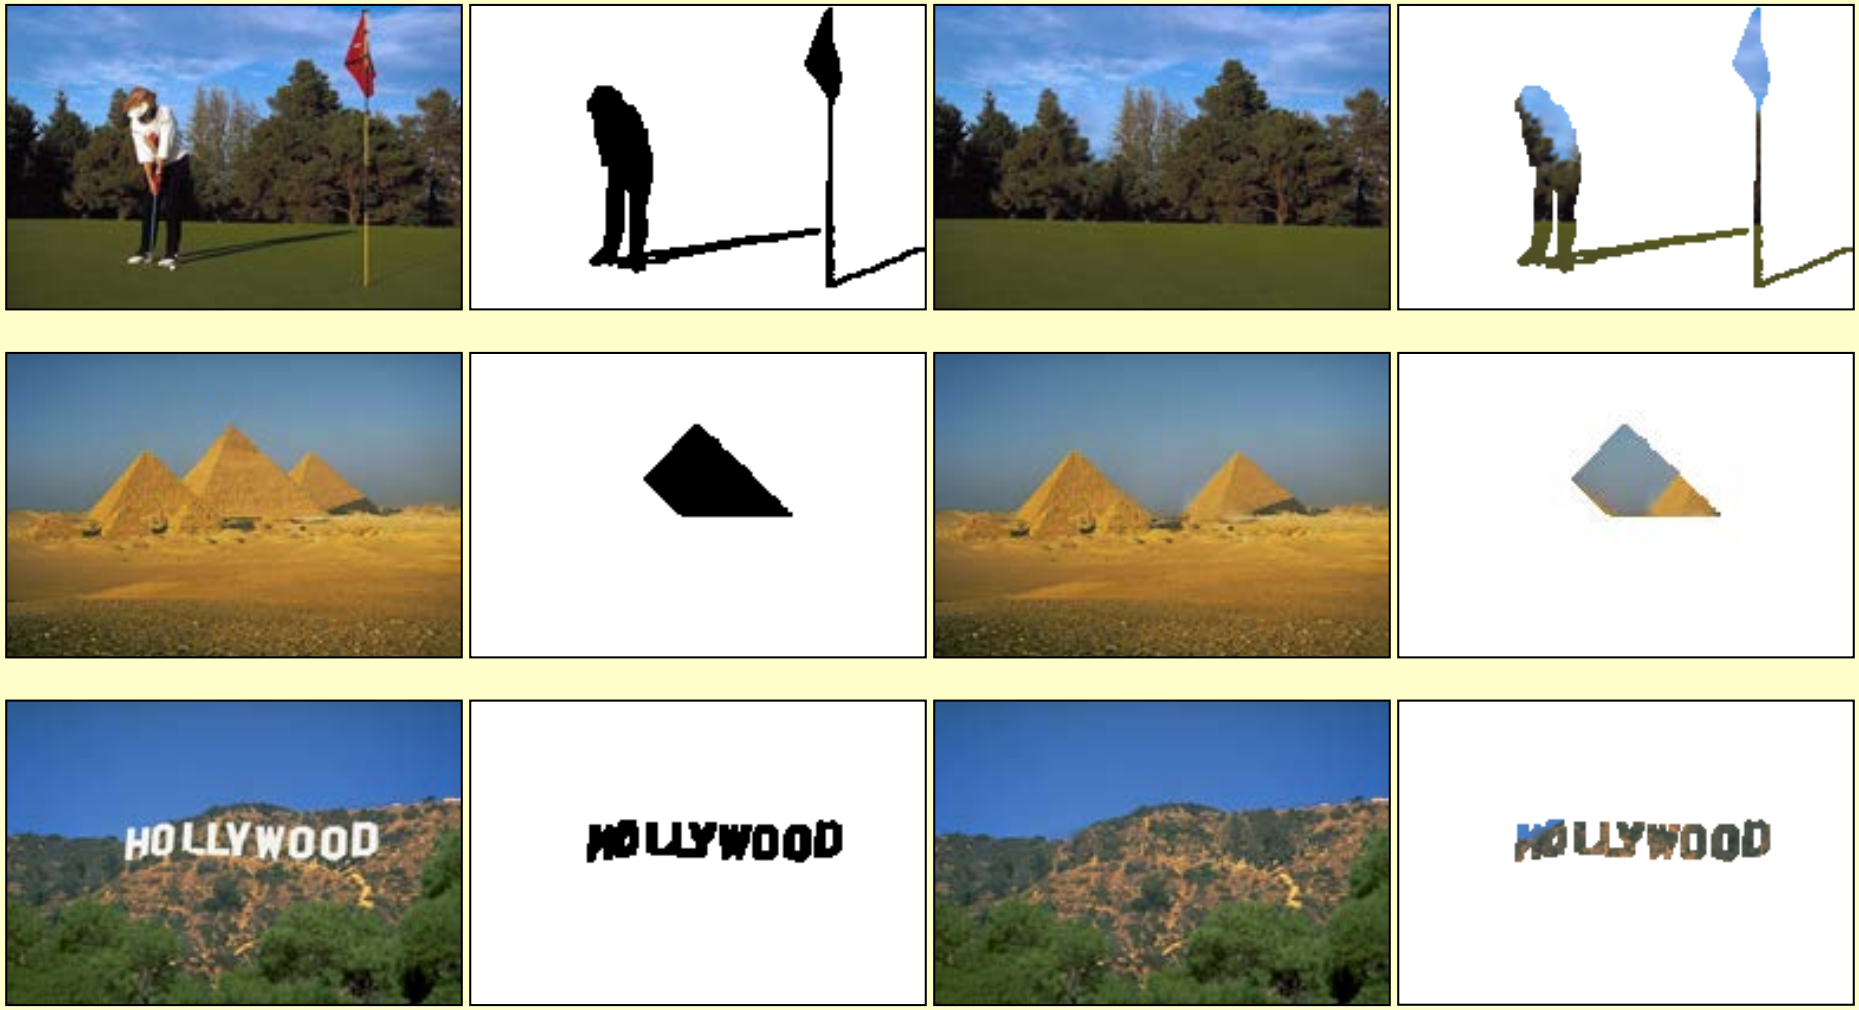

Results

input image

completion

Results

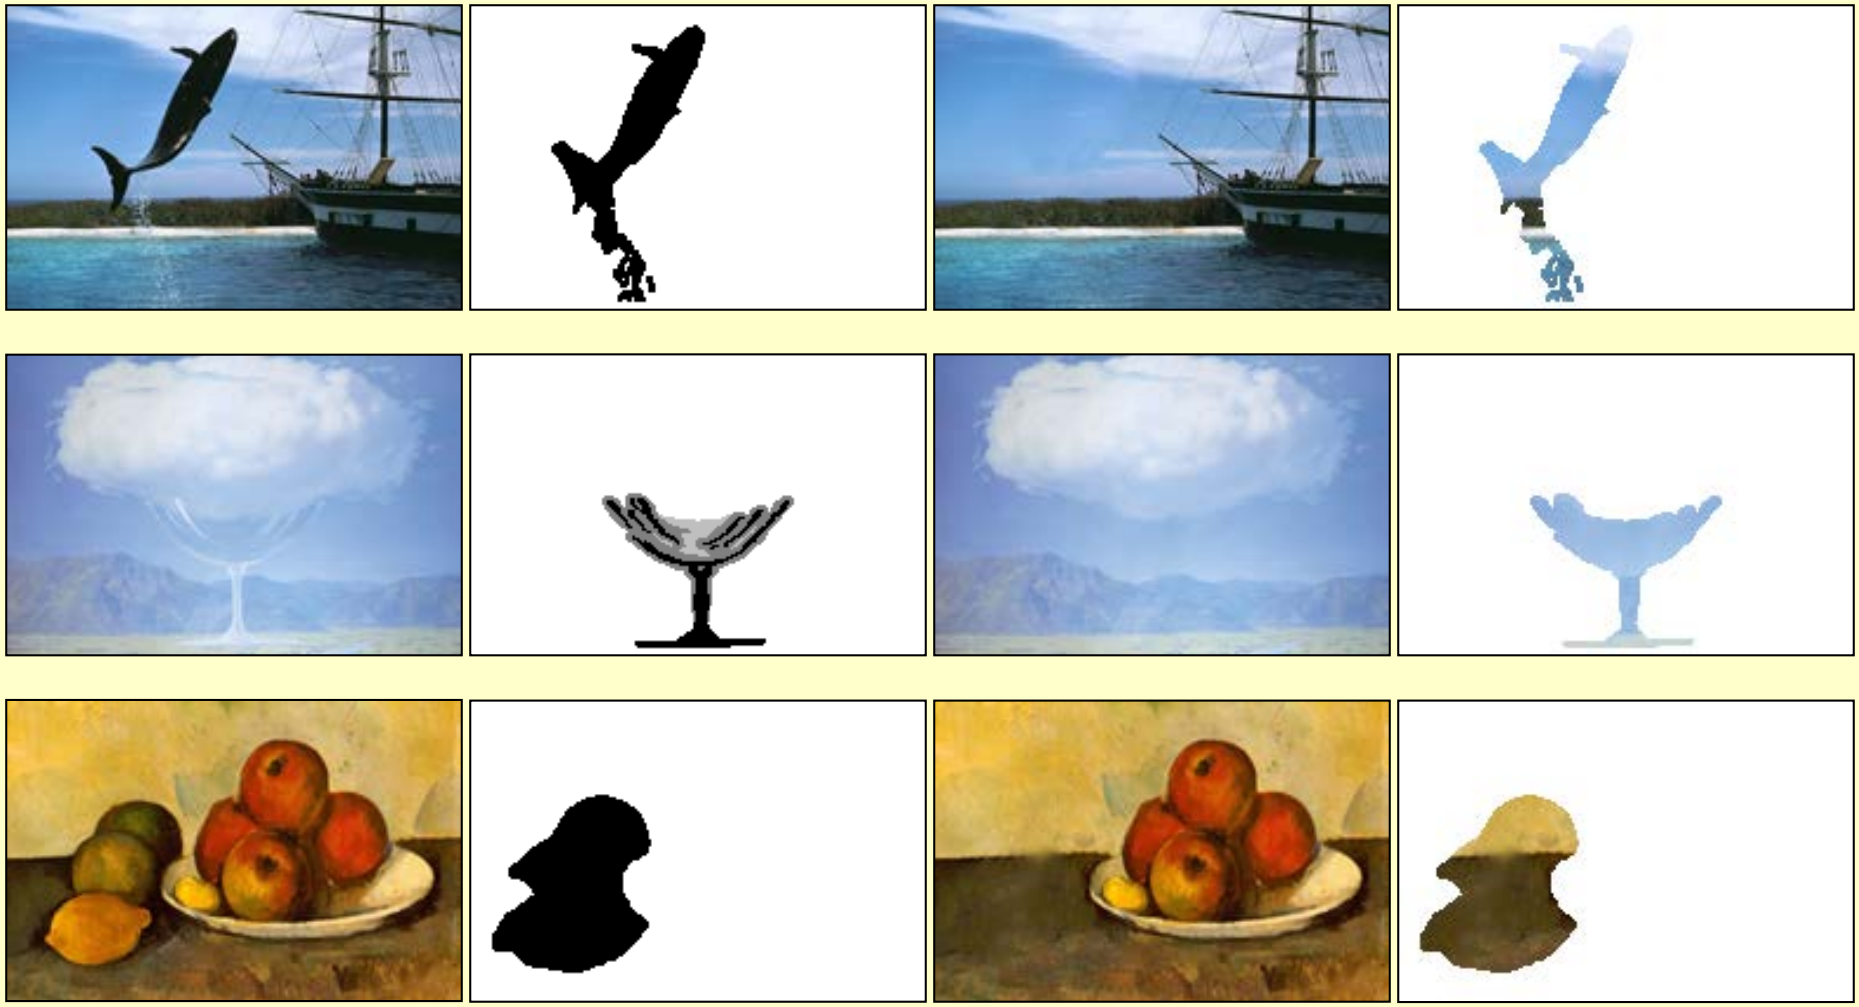

Video Completion

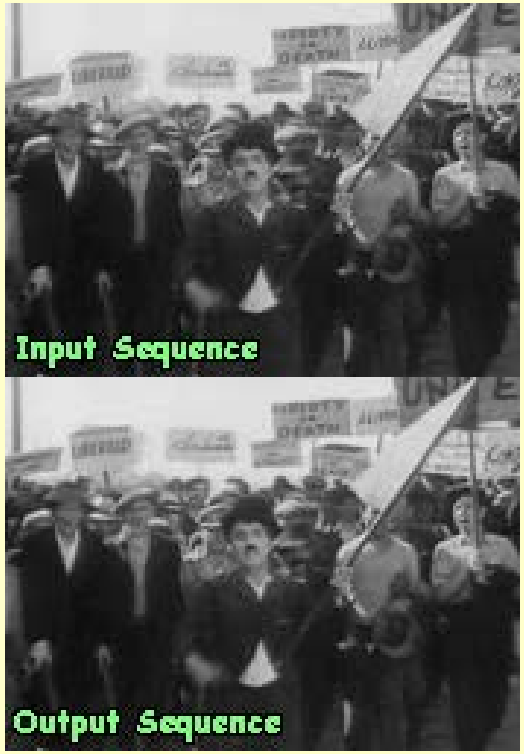

. Wexler E. Shechtman M. Irani; “Space-Time Video Completion"; CVPR’04.

Time-varying Weathering in Texture Space

Rachele Bellini, Yanir Kleiman, Daniel Cohen-Or

SIGGRAPH 2016

(Neural) Texture Synthesis

Given a sample patch of some texture, can we generate a bigger image of the same texture?

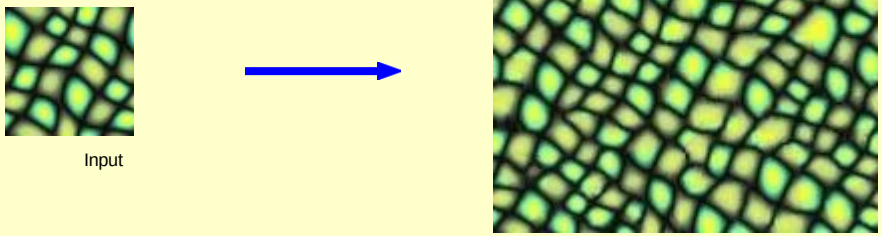

Input

Output

Neural Style Transfer

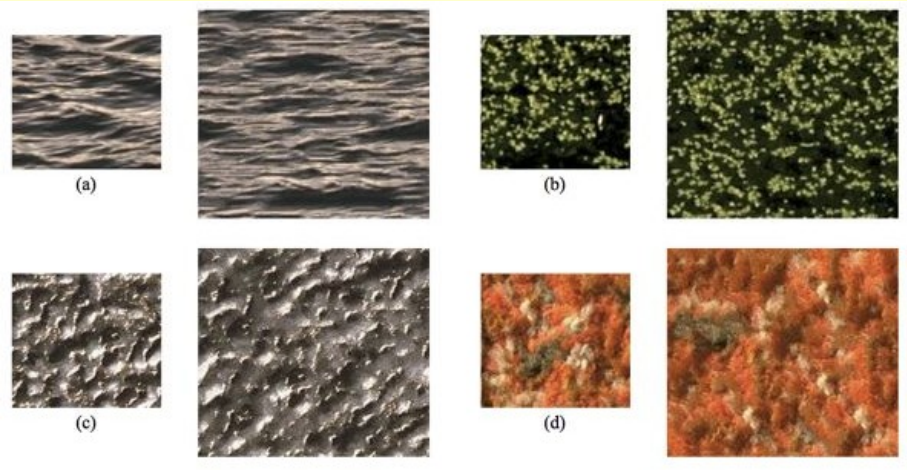

Wei and Levoy, “Fast Texture Synthesis using Tree-structured Vector Quantization”, SIGGRAPH 2000

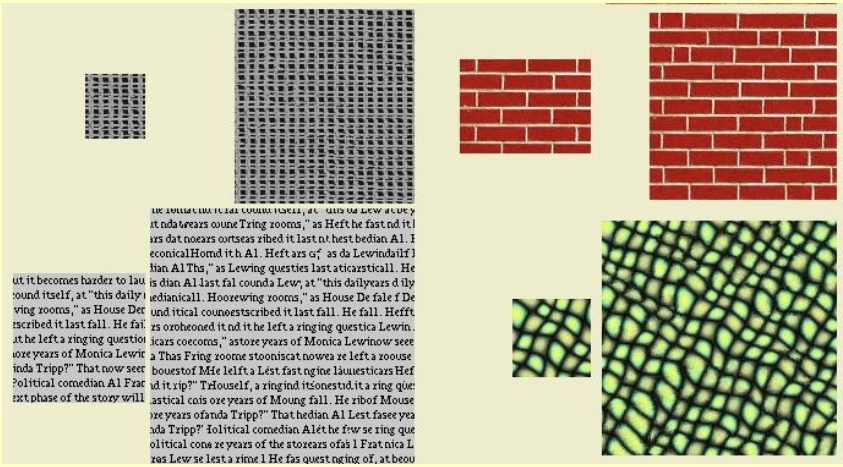

Efros and Leung, “Texture Synthesis by Non-parametric Sampling”, ICCV 1999

Neural Style Transfer

$$G^l_{ij} = \sum_k F^l_{ik} F^l_{jk}$$

1. Pretrain CNN on ImageNet (VGG-19)
2. Feed a texture, record activations
3. Compute Gram matrix for each layer:
4. Initialize image from random noise
5. Feed image through CNN, computing Gram matrices
6. Compute loss as the sum of L_2 distances between Gram matrices
7. Back-propagate to get a gradient for the input image
8. Update input image
9. Repeat 5-8 until convergence

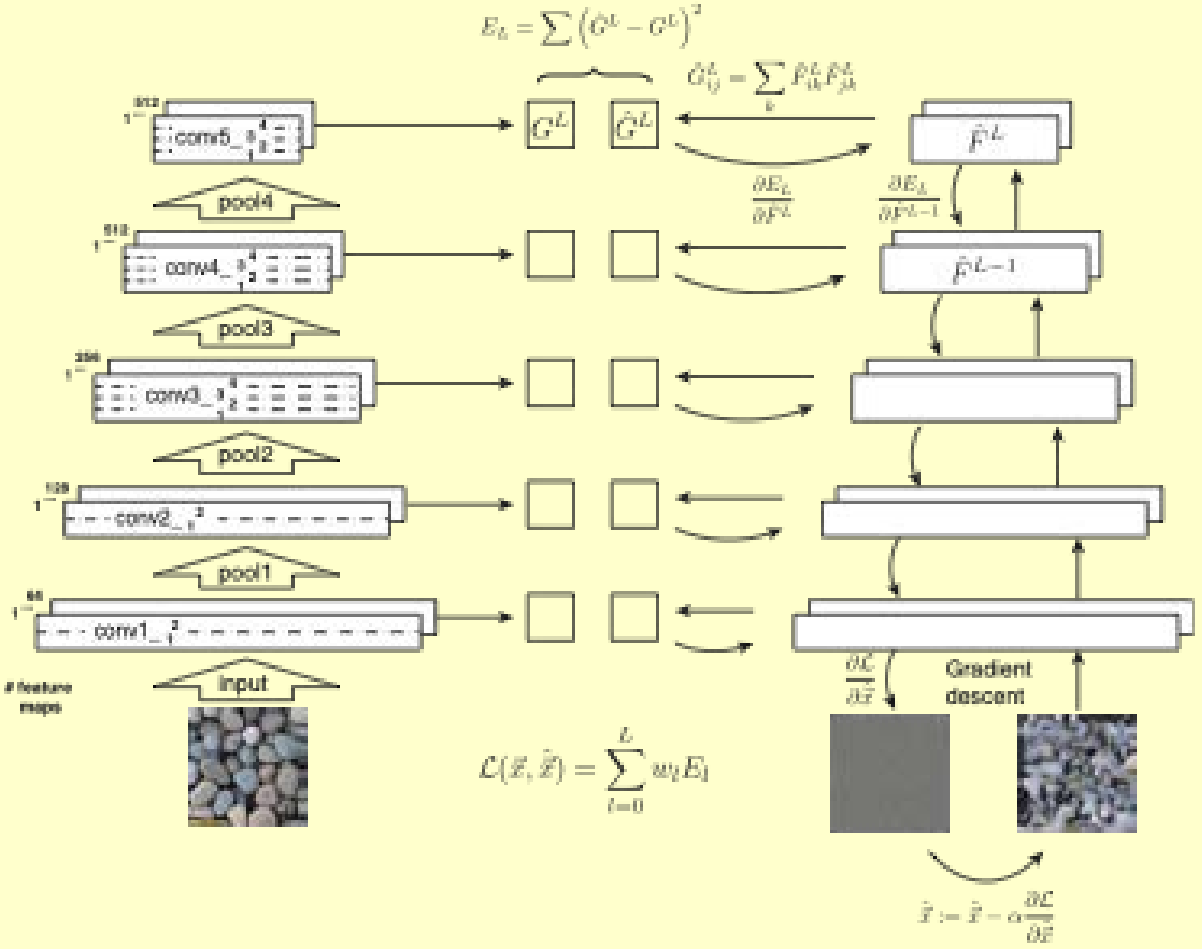

Style Loss

- Gram matrices $G^l, A^l \in R^{C_l \times C_l}$
  Features correlation matrix:

$$G^l_{\substack{i,j \\ 1<i,j<C^l}} = \frac{1}{C_l H_l W_l} \sum_{h=1}^{H_l} \sum_{w=1}^{W_l} F_{h,w,i} \cdot F_{h,w,j}$$

- Error at layer l:

$$E_l = \left\| G^l - A^l \right\|_F^2 = \sum_{i,j} \left(G^l_{i,j} - A^l_{i,j} \right)^2$$

- Weighted loss

$$L_{style}\left(x, G(z)\right) = \sum_{i=1}^{5} w_i E_i$$

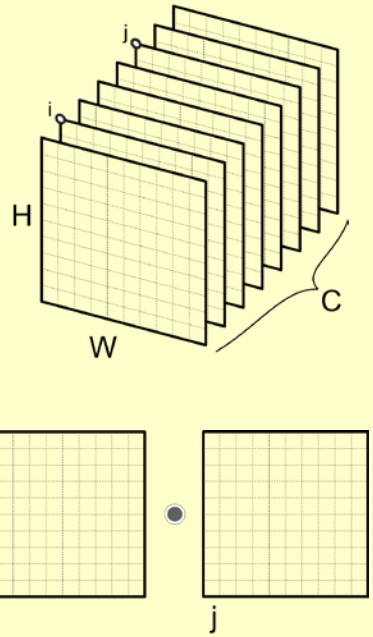

L.A. Gatys, A.S. Ecker, and M. Bethge. Texture Synthesis Using Convolutional Neural Networks. Advances in Neural Information Processing Systems 28 (May 2015)

Style Loss cont.

$$L_{style}(x, G(z)) = \sum_{i=1}^{5} w_i E_i$$

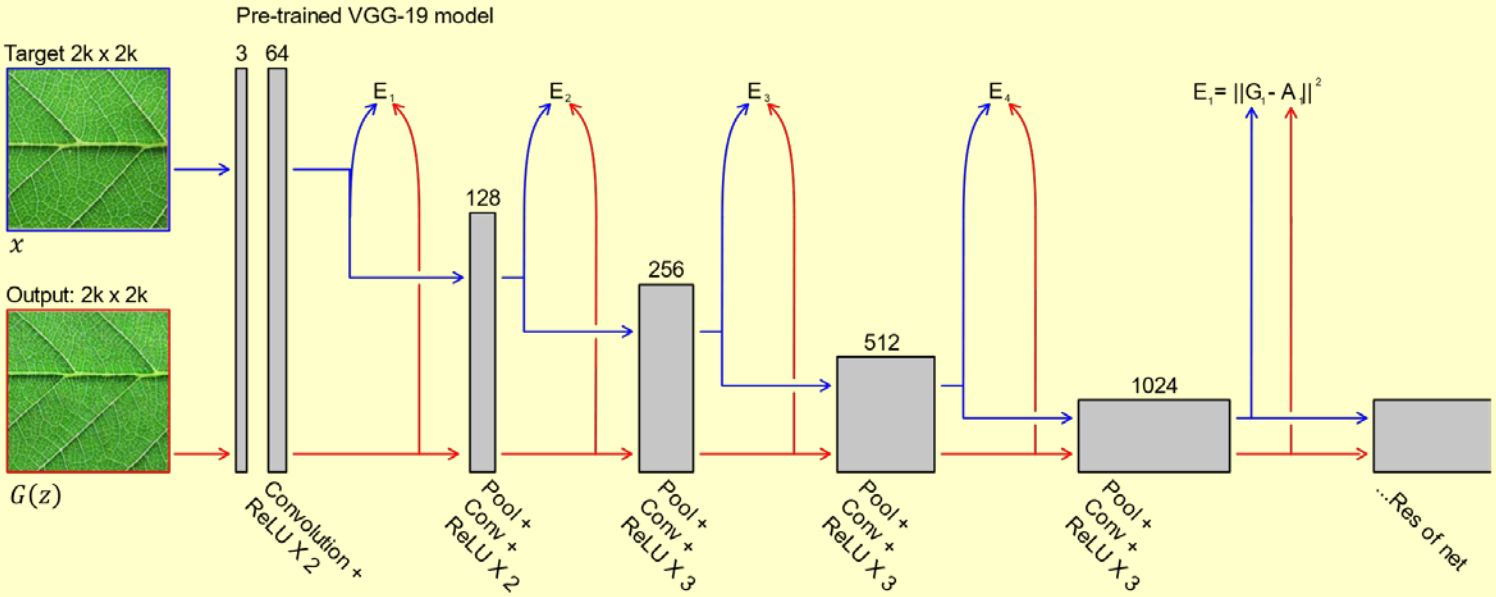

L.A. Gatys, A.S. Ecker, and M. Bethge. Texture Synthesis Using Convolutional Neural Networks. Advances in Neural Information Processing Systems 28 (May 2015)

Neural Style Transfer

Reconstructing from higher layers recovers larger features from the input texture

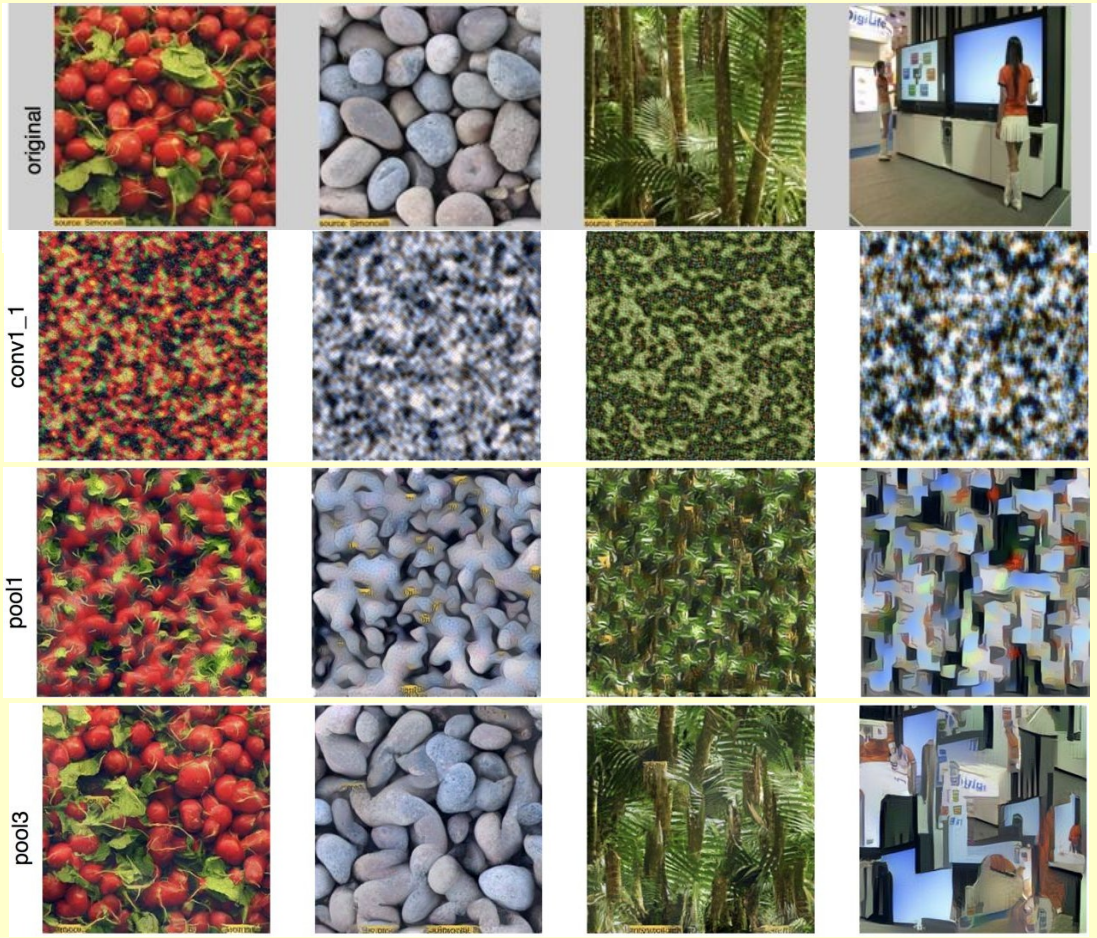

Gatys et al, “Texture Synthesis using Convolutional Neural Networks”, NIPS 2015

Style Transfer:
Feature Inversion + Texture Synthesis

Neural Style Transfer

Feature reconstruction

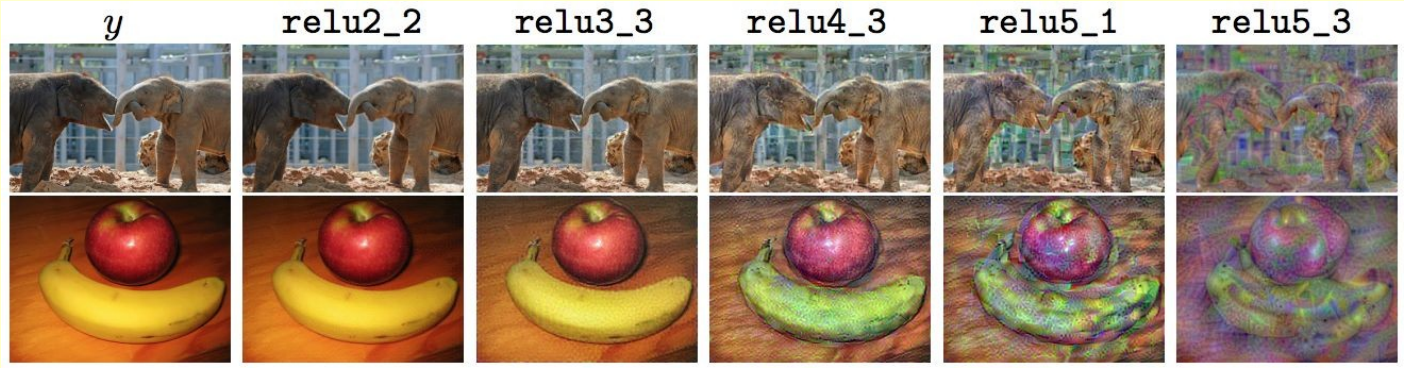

Texture synthesis (Gram reconstruction)

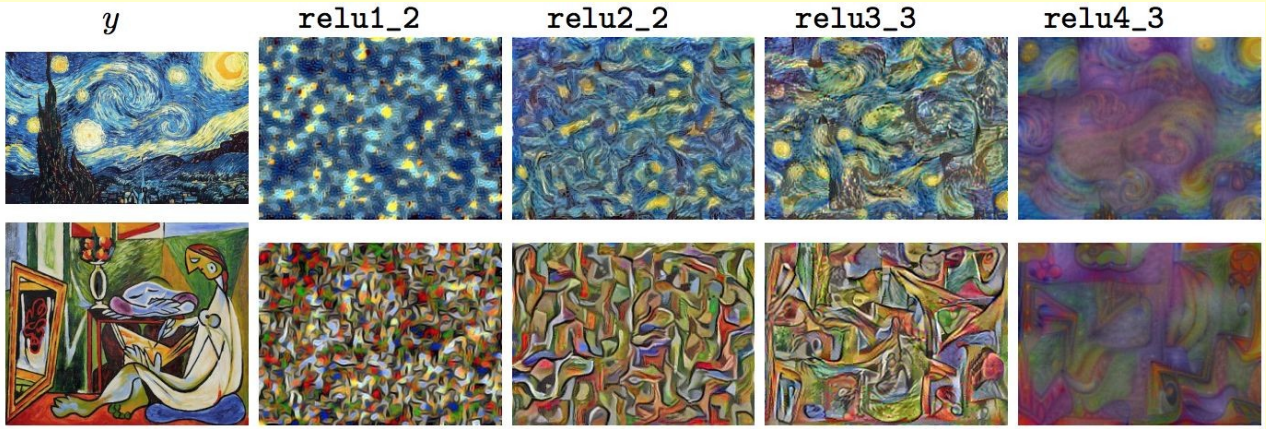

Figure credit: Johnson et al, “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”, ECCV 2016

Neural Style Transfer

Given a **content image** and a **style image**, find a new image that
- Matches the CNN features of the content image (feature reconstruction)
- Matches the Gram matrices of the style image (texture synthesis)

Combine feature reconstruction from Mahendran et al with Neural Texture Synthesis from Gatys et al, using the same CNN!

Content Image

+

Style Image

=

Stylized Result

Gatys et al, “A Neural Algorithm of Artistic Style”, arXiv 2015
Gatys et al, “Image Style Transfer using Convolutional Neural Networks”, CVPR 2016

Neural Style Transfer

1. Pretrain CNN
2. Compute features for content image
3. Compute Gram matrices for style image
4. Randomly initialize new image
5. Forward new image through CNN
6. Compute style loss (L2 distance between Gram matrices) and content loss (L2 distance between features)
7. Loss is weighted sum of style and content losses
8. Backprop to image
9. Take a gradient step
10. GOTO 5

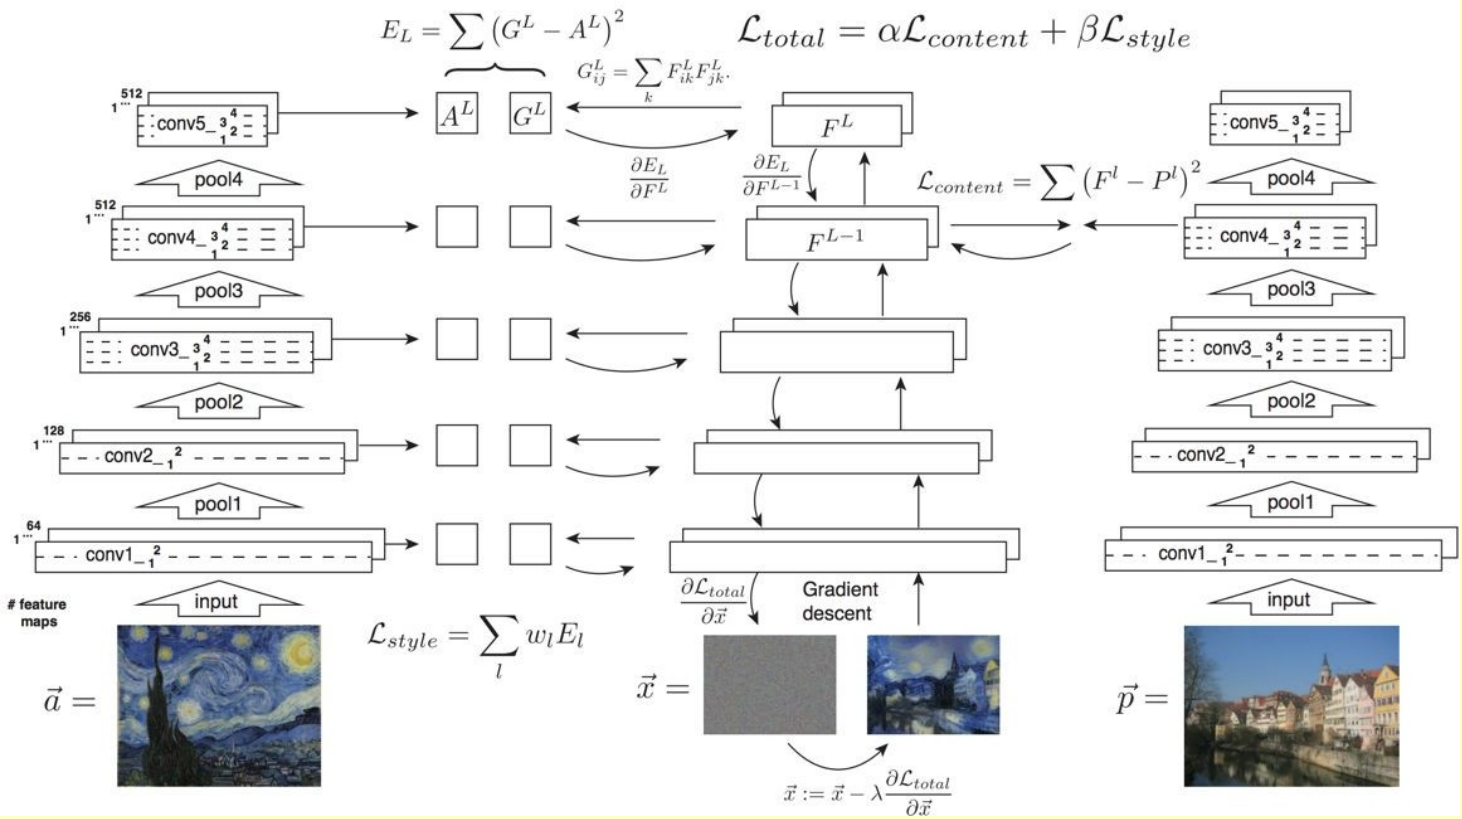

Gatys et al, “Image Style Transfer using Convolutional Neural Networks”, CVPR 2016

Neural Style Transfer

Gatys et al, “Image Style Transfer using Convolutional Neural Networks”, CVPR 2016

Neural Style Transfer

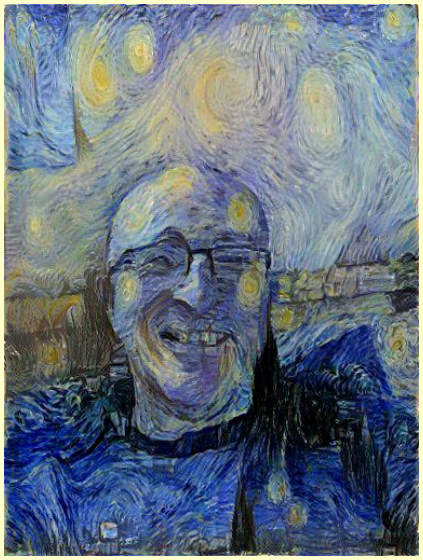
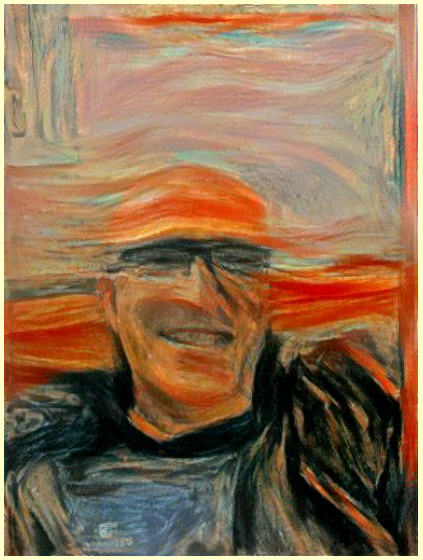
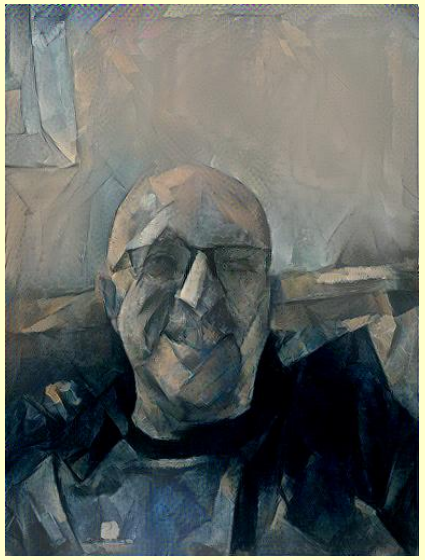
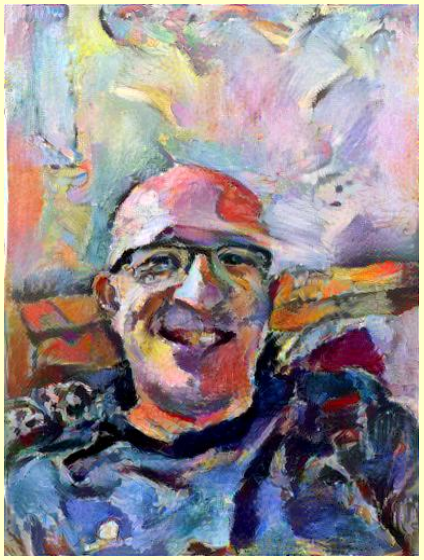

Neural Style Transfer: Style / Content Tradeoff

More weight to content loss

More weight to style loss

Justin Johnson, “neural-style”, https://github.com/jcjohnson/neural-style

Neural Style Transfer: Style Scale

Resizing style image before running style transfer algorithm can transfer different types of features

Larger style image ⟷ Smaller style image

Justin Johnson, “neural-style”, https://github.com/jcjohnson/neural-style

Neural Style Transfer: Multiple Style Images

Mix style from multiple images by taking a weighted average of Gram matrices

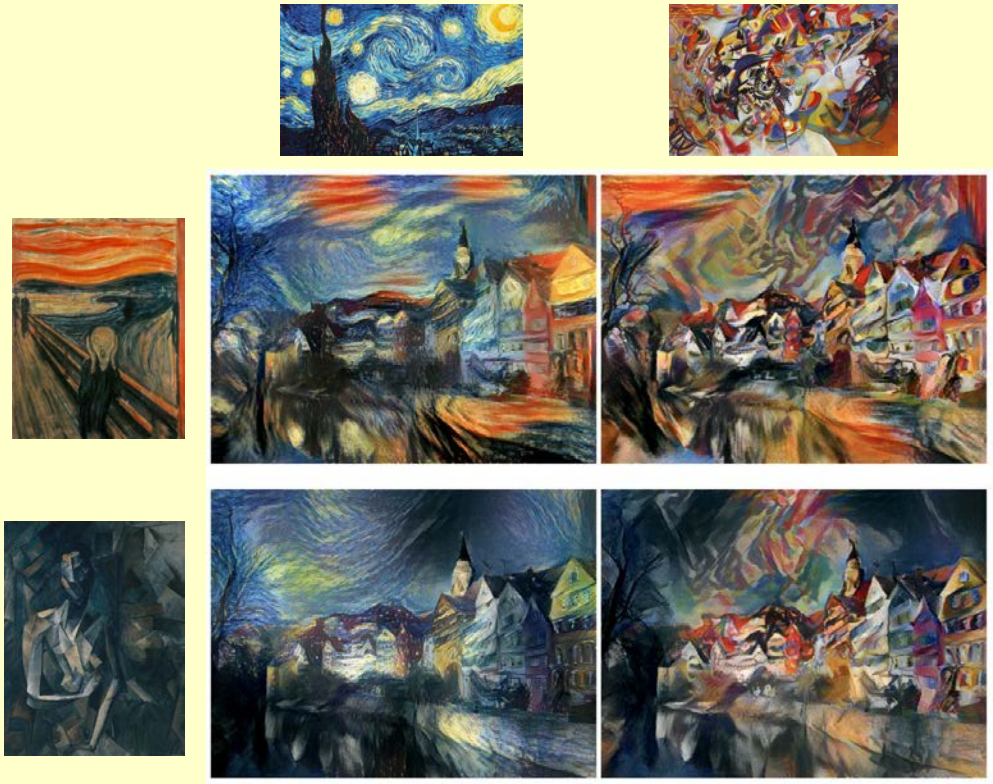

Justin Johnson, “neural-style”, https://github.com/jcjohnson/neural-style

Neural Style Transfer: Multiple Style Images

More “Scream” ⟷ More “Starry Night”

Justin Johnson, “neural-style”, https://github.com/jcjohnson/neural-style

Fast Style Transfer

Problem: Style transfer is slow; need hundreds of forward + backward passes of VGG

Solution: Train a feedforward network to perform style transfer!

Fast Style Transfer

(1) Train a feedforward network for each style
(2) Use pretrained CNN to compute same losses as before
(3) After training, stylize images using a single forward pass

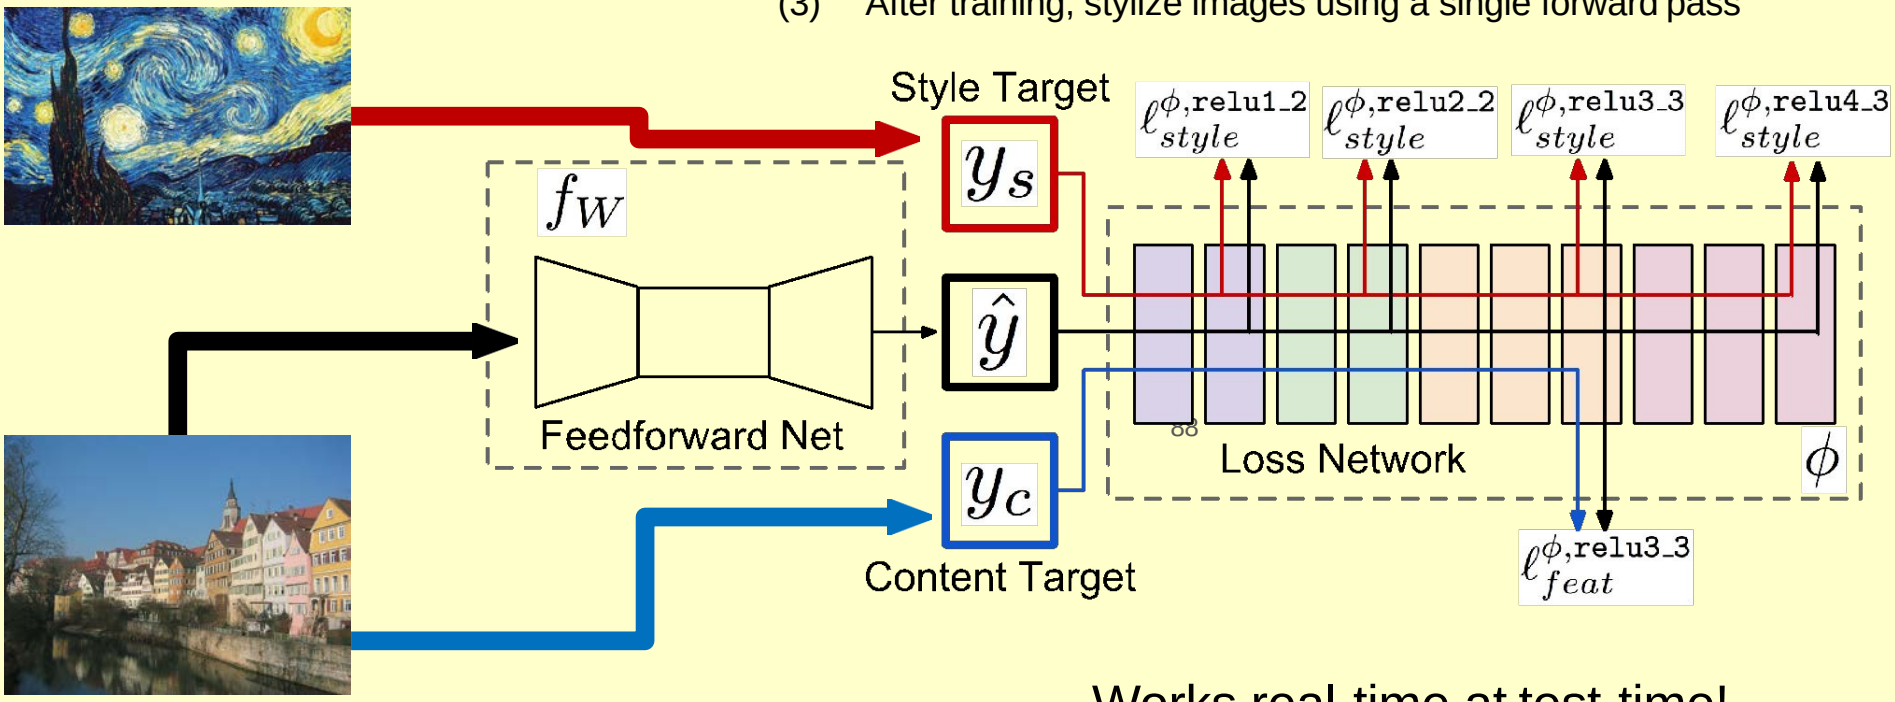

Works real-time at test-time!

Johnson et al, “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”, ECCV 2016

Fast Style Transfer

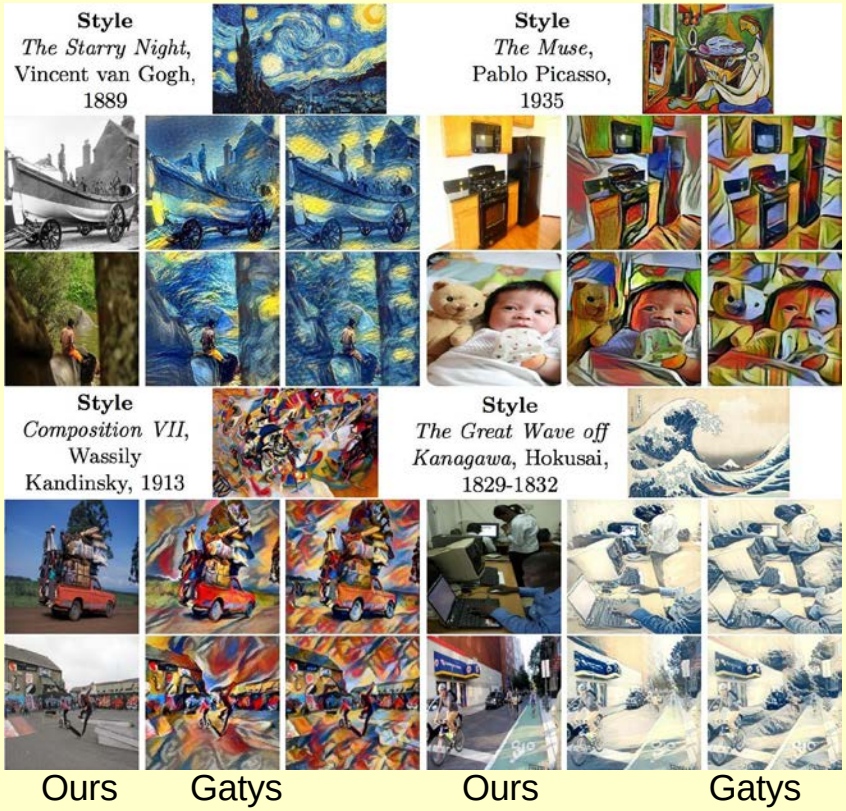

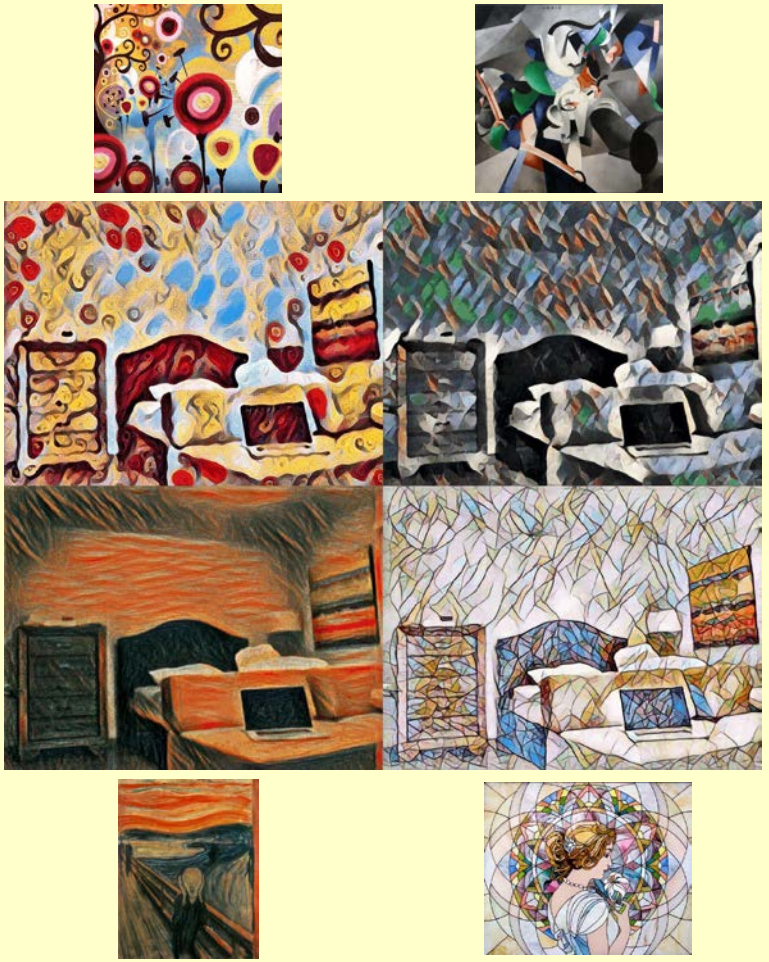

Johnson et al, “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”, ECCV 2016
https://github.com/jcjohnson/fast-neural-style

Works real-time on video!

Learning to transfer style

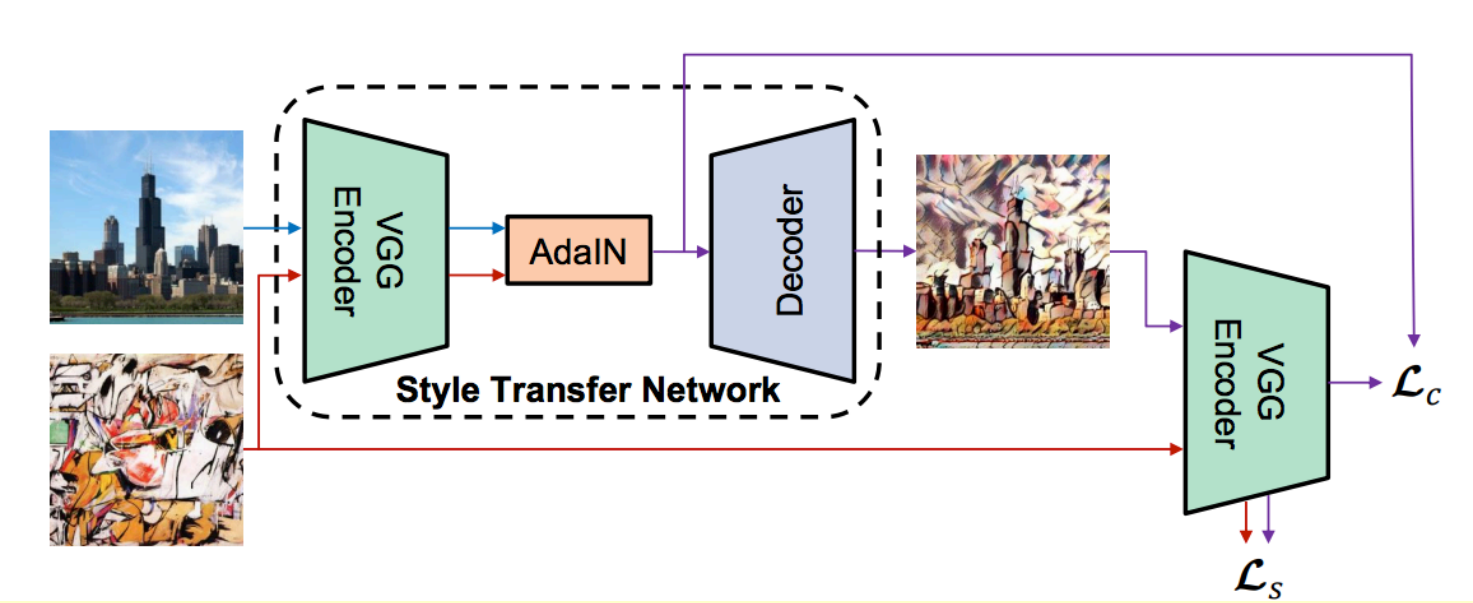

Huang, Xun; Belongie, Serge
Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization
International Conference on Computer Vision (ICCV), Venice, Italy, 2017, (Oral).