


Texture

CS 419

Slides by Zicheng Liao (Courtesy of Ali Farhadi)




Texture scandals






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.



What is Texture?



Texture spectrum

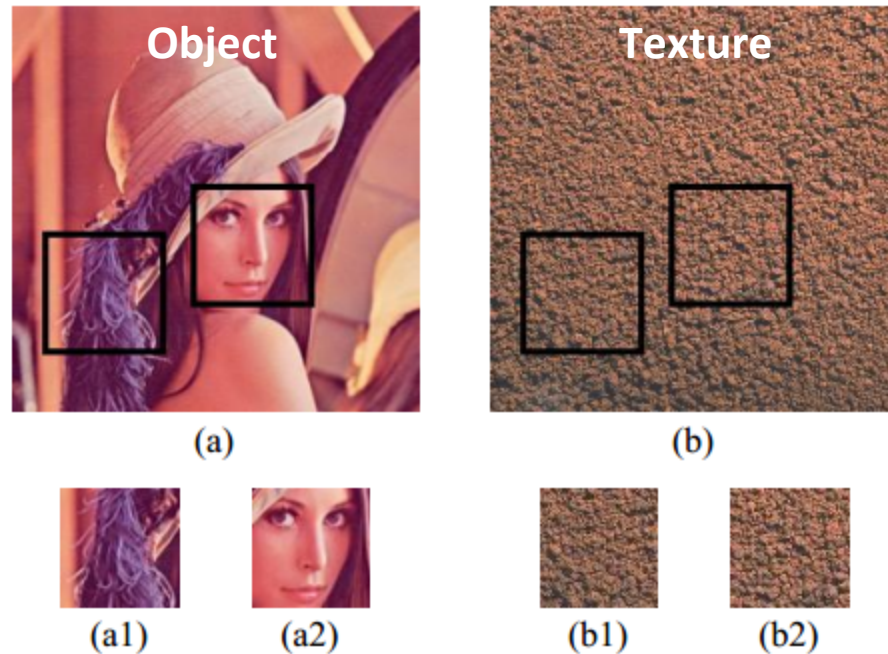
Diversity: color, scale, content, caused by geometry or color



Regularity


Randomness

How textures differ from objects




- **Stationary:** similar views as window moves around in (b), because the *statistical distribution is spatially invariant*.
- **Locality:** each pixel in (b) is only related to a small set of neighbors.

Texture synthesis



Input




Output

Two crucial algorithmic points

- Nearest neighbor search
- Dynamic programming

Algorithm I: NN + sampling



- How to determine a unknown pixel value?
 - Ask neighbors (locality)
- Find pixels with similar neighbors in the input, and then randomly take one



[Efros & Leung ICCV99]

NN patch match


- Find a few pixels with similar neighbors



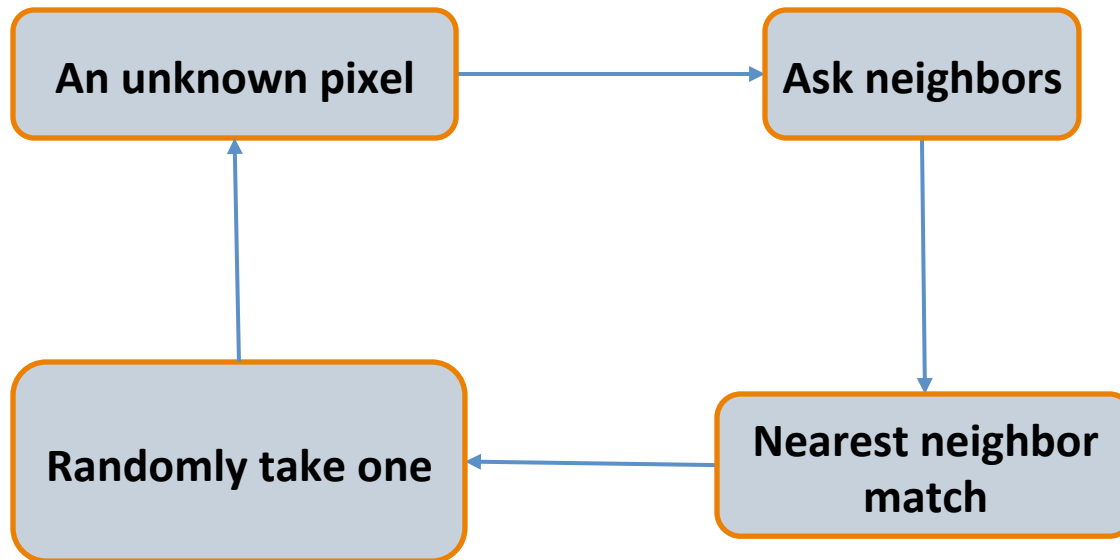
Input image

Randomly Sample

- Find a few pixels with similar neighbors
- Randomly take one matched pixel



The steps in a single iteration



Algorithm

1. Grow from the border of an input texture
2. For each unknown pixel p on the boundary
3. Gather the neighborhood centered at p : $N(p)$
4. Find patches $N(p')$ from input: $d(N(p), N(p')) < (1 + \epsilon)d(N(p), N_{best})$
5. Randomly take a p' to fill in p

\epsilon = 0.1

Discussion

1. Grow from the border of an input texture
2. For each unknown pixel p on the boundary
3. Gather the neighborhood centered at p : $N(p)$
4. Find patches $N(p')$ from input: $d(N(p), N(p')) < (1 + \epsilon)d(N(p), N_{best})$
5. Randomly take a p' to fill in p

$\epsilon = 0.1$

- How does the choice of ϵ effect the result?
 - Not much, some algorithm even set it to 0.
- Other than random sample?
 - Use the distance to bias the sampling, favor better matches

Discussion (2)


1. Grow from the border of an input texture
2. For each unknown pixel p on the boundary
3. Gather the neighborhood centered at p : $N(p)$
4. Find patches $N(p')$ from input: $d(N(p), N(p')) < (1 + \epsilon)d(N(p), N_{best})$
5. Randomly take a p' to fill in p

\epsilon = 0.1

- Distance metric
- Window size
- Order to synthesize

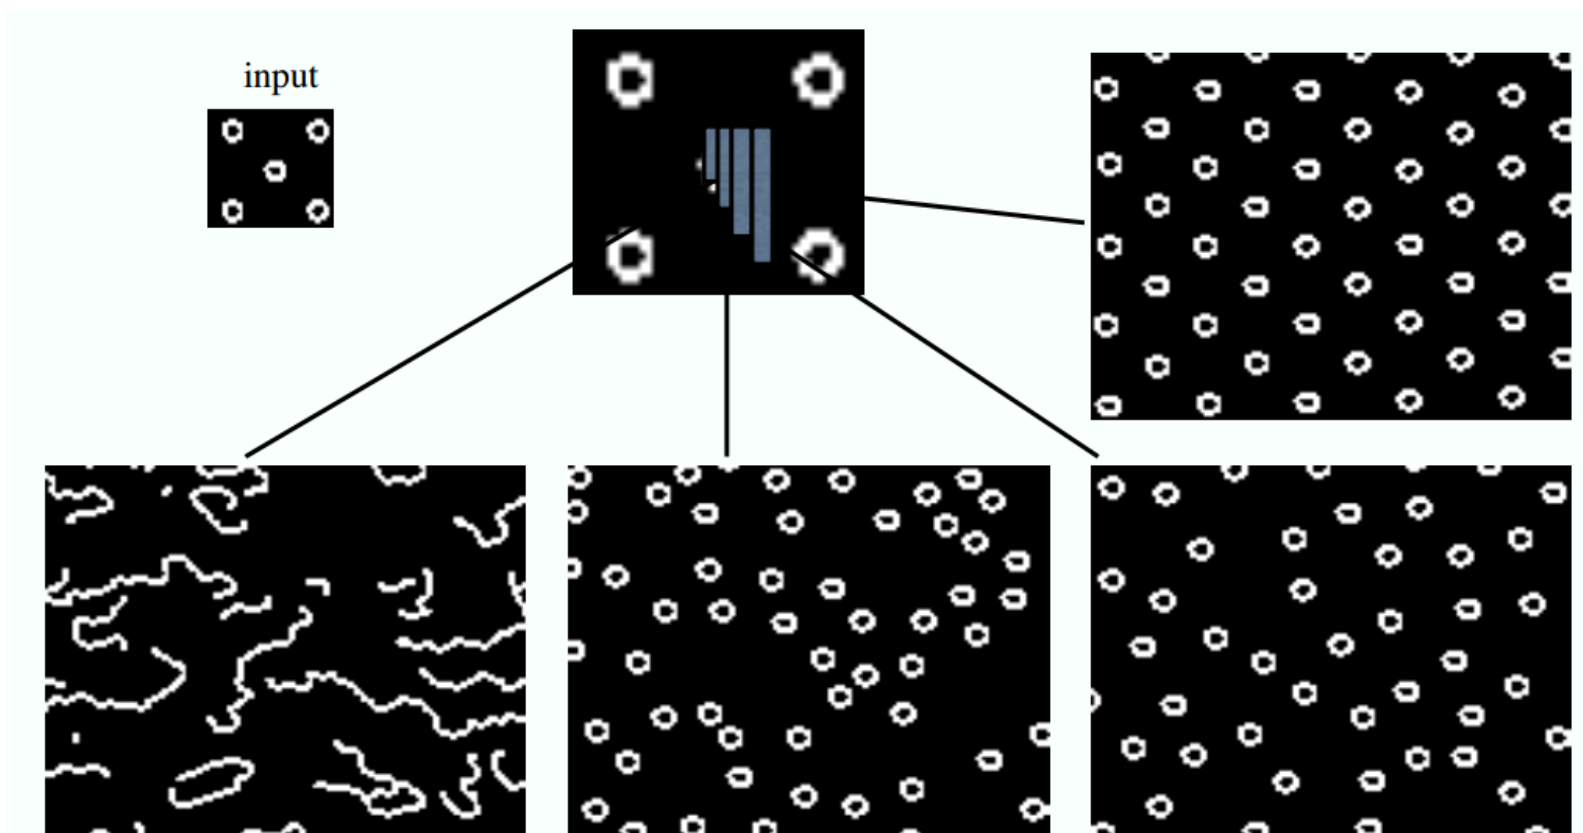
Patch distance metric

- Normalized sum of squared difference
- Further neighbors take less weight
 - Gaussian mask
 - Preserve local structure



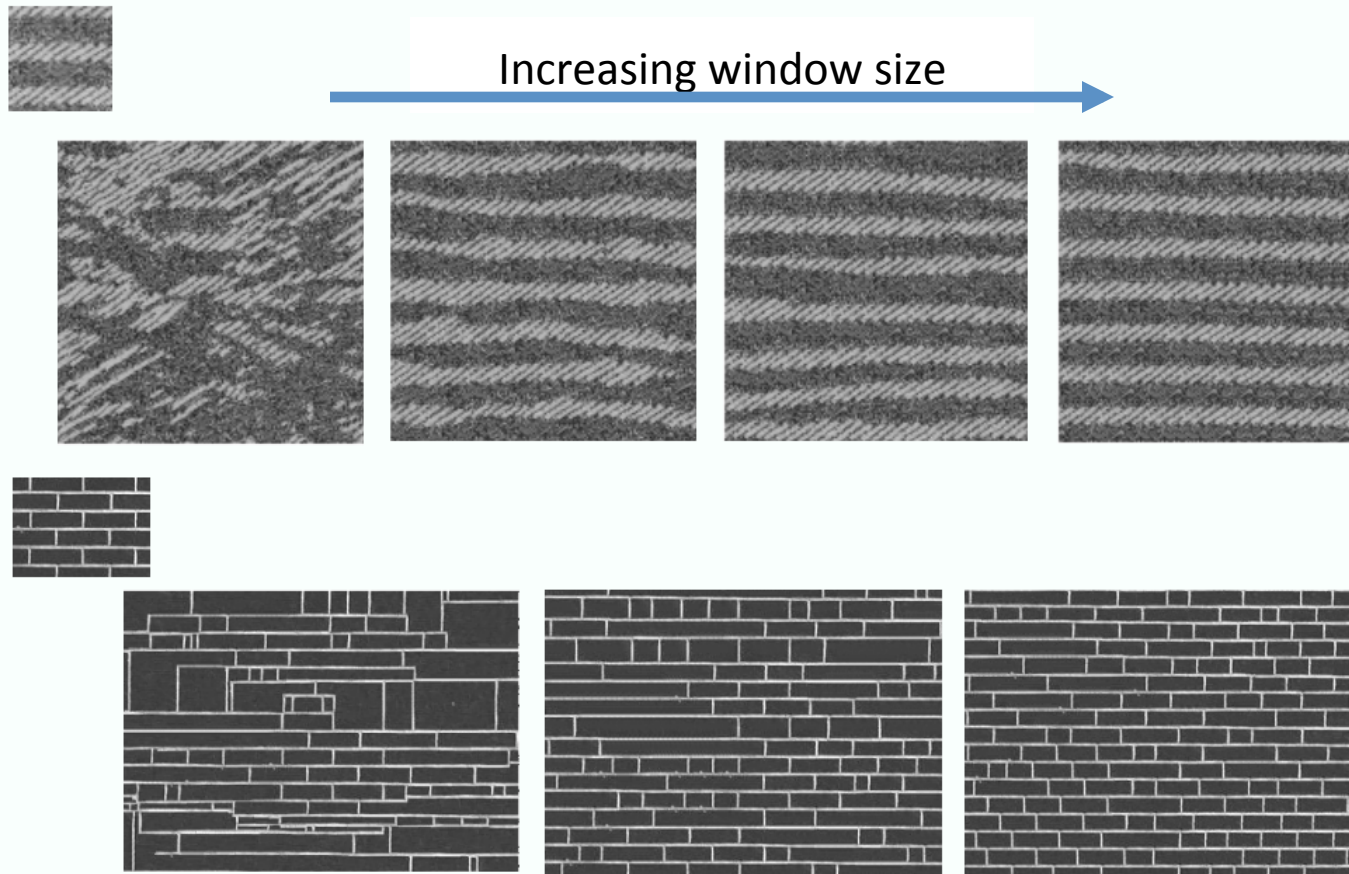
Window Size

- Effect of window size on the results




Window Size

- Control the degree of randomness






The order matters




[Efros & Leung ICCV99]

Some results




More results

ut it becomes harder to lau
ound itself, at "this daily
wing rooms," as House Der
scribed it last fall. He fai
at he left a ringing question
ore years of Monica Lewin
inda Tripp?" That now see
Political comedian Al Fra
ext phase of the story will





ut it becomes harder to lau
ound itself, at "this daily
wing rooms," as House Der
scribed it last fall. He fai
at he left a ringing question
ore years of Monica Lewin
inda Tripp?" That now see
Political comedian Al Fra
ext phase of the story will




More results

french canvas




rafia weave



Some results

wood




granite





Some results



white bread





brick wall





Hole Filling




Hole Filling



Extrapolation



Failure Cases



Growing garbage



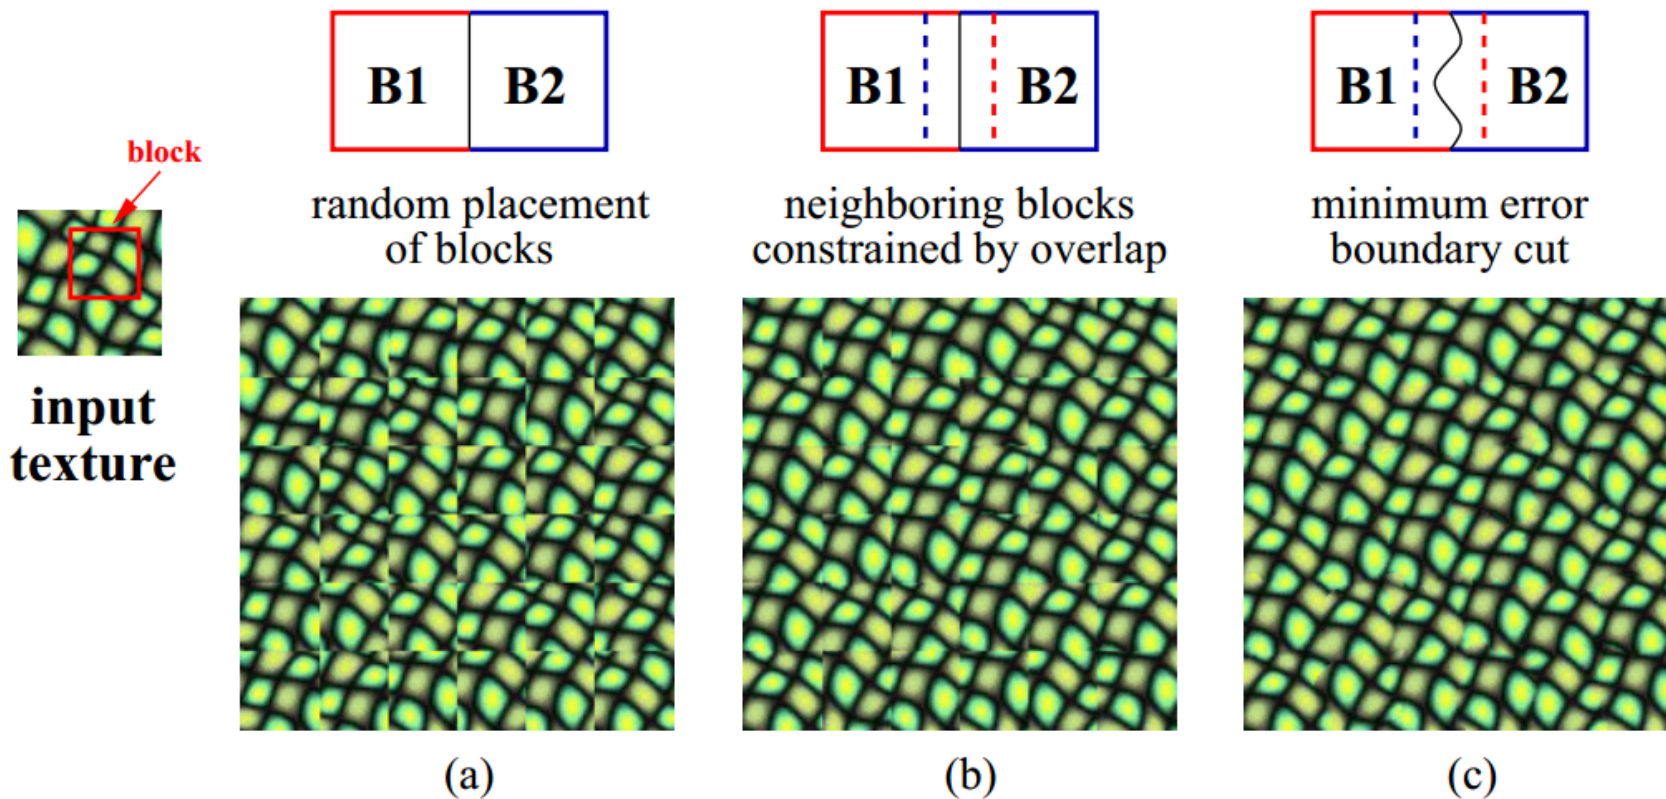
Verbatim copying

Pros and Cons

- Very simple
- Easy to implement: **32 lines of matlab code!**
- Works well for a variety of synthetic and real-world textures

- **VERY VERY slow!**
(A nearly identical idea was proposed in 1981 by Barber but discarded due to computational intractability)
- Idea
 - **A patch a time**, instead of a pixel

Image Quilting: Patch-based method

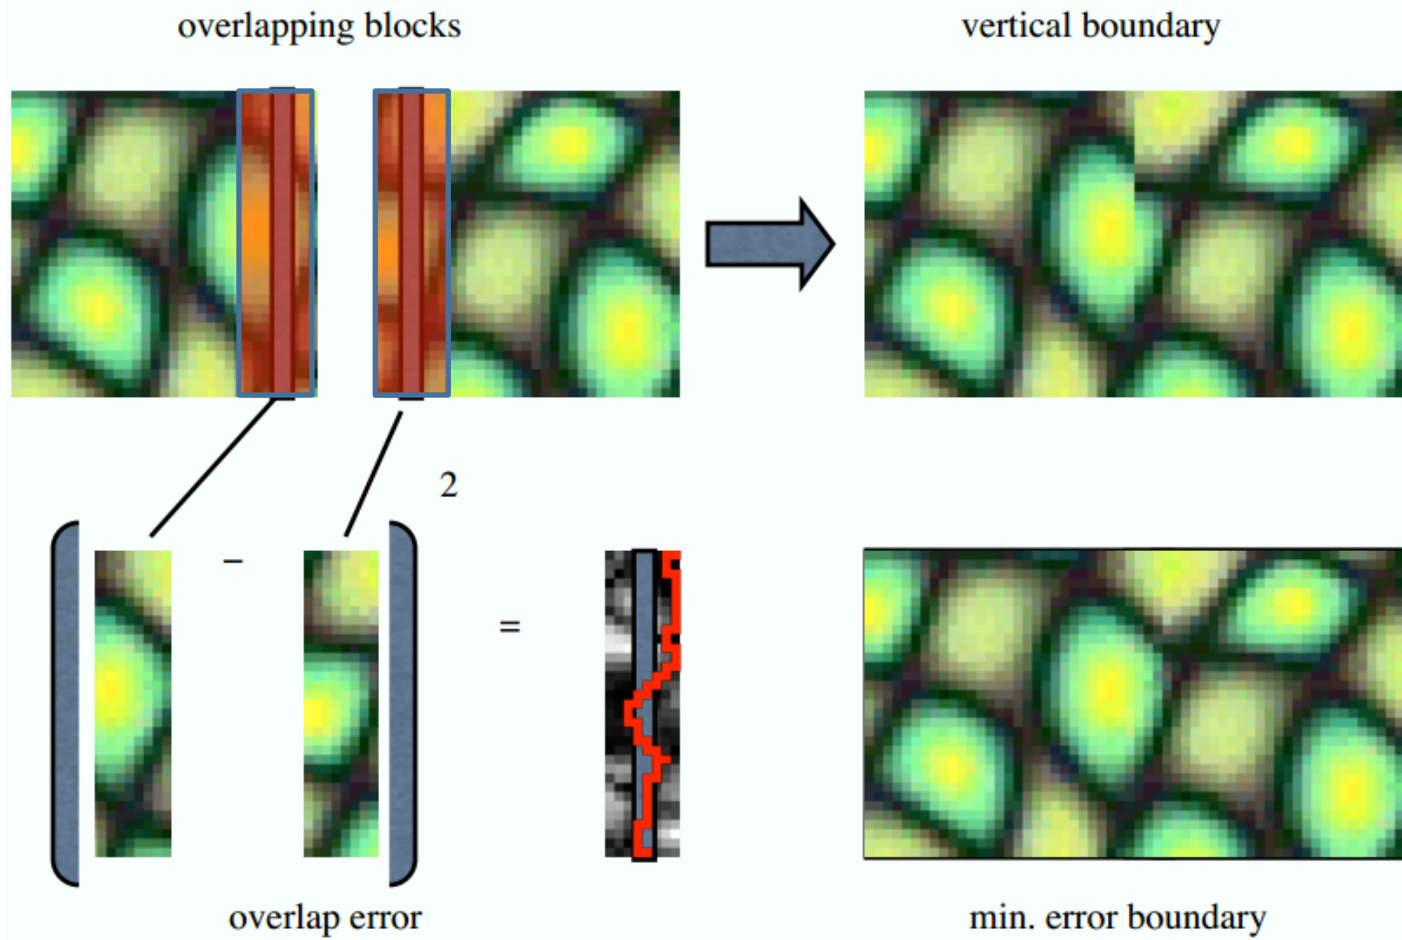


(a) random blocks concatenated together


(b) Blocks overlap, new block is chosen so that the overlap regions best agree

(c) A minimum cost (optimal) path is computed within the overlap

Curved path VS vertical path



How to find the optimal path?



$B1$



$B2$




$e = (B1-B2)^2$

- Brute force: exponential number of paths
- Greedy algorithm? No..
- Key observation: **every optimal sub-path is part of an optimal full path**
→ *Dynamic programming*

<http://community.topcoder.com/tc?module=Static&d1=tutorials&d2=dynProg>
(A nice dynamic programming tutorial)

Dynamic programming



$B1$



$B2$



e


```
Initialize:  $e_{i,j} = (B1_{ij} - B2_{ij})^2$   
for  $i = 2:h$ ; for  $j = 1:w$   
     $E_{i,j} = e_{i,j} + \min(E_{i-1,j-1}, E_{i-1,j}, E_{i-1,j+1})$   
     $k_{i,j} = \operatorname{argmin}(E_{i-1,j-1}, E_{i-1,j}, E_{i-1,j+1})$   
end; end
```

$e_{\{i,j\}}$: node cost at pixel (i,j)

$E_{\{i,j\}}$: optimal path cost up to node (i, j)




$k_{\{i,j\}}$: index to the optimal (next) sub-path

Dynamic programming





1. Compute path costs: start from the bottom, iteratively go up, end at the top
2. Get optimal path: compare the path cost of nodes on the top row, find the minimum cost node, and use $k_{\{i,j\}}$ to trace back for the optimal path down to the bottom

Results




[Efros & Freeman SIGGRAPH01]

More results




[Efros & Freeman SIGGRAPH01]

More results



[Efros & Freeman SIGGRAPH01]

More results



[Efron & Freeman SIGGRAPH01]

More results



[Efros & Freeman SIGGRAPH01]

More results




[Efros & Freeman SIGGRAPH01]

Failure cases



[Efros & Freeman SIGGRAPH01]

Texture Transfer




[Efros & Freeman SIGGRAPH01]

Texture Transfer



(u, v)

source texture




target image




(i, j)

texture transfer result



(u, v)



(i, j)

correspondence maps

$$e_{i,j} = \alpha(B1_{ij} - B2_{ij})^2 + (1 - \alpha)(Sc_{uv} - Tc_{i,j})^2$$

(u,v) : the coordinate of patch $B2_{i,j}$ in the source texture

Sc : source correspondence map

Tc : target correspondence map

More results



source texture




target images



texture transfer results

[Efros & Freeman SIGGRAPH01]

More results



[Efros & Freeman SIGGRAPH01]

More results

parmesan



+



=



rice



+



=

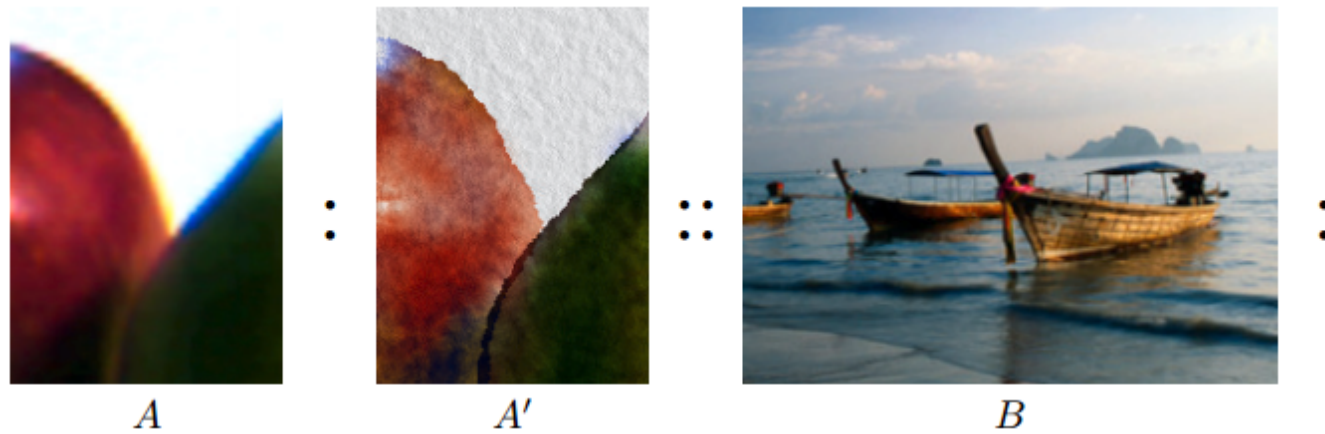


Pros and Cons

- Very simple
- Easy to implement
- Work well
- Fast!

- **Memoryless**
 - Cannot keep track of old solutions: A general problem for DP
- Better algorithm
 - Graph cut [*Graphcut texture, Kwatra SIGGRAPH03*]

Image Analogy




[Hertzman et al. SIGGRAPH01]

Image Analogy



A

•
•



A'

••
••



B

•
•



B'

[Hertzman et al. SIGGRAPH01]

Image Analogy



A

A'


B

B'

Problem(“IMAGE ANALOGY”): Given a pair of images A (*unfiltered source*) and filtered image A' (*filtered source*), Along with some additional unfiltered *target image* B, Synthesize a new *filtered target image* B' such that:

$$A : A' :: B : \mathbf{B'}$$

Basic idea



For q in B' , find index p in source images such that

$$A'(p) \sim B'(q) \text{ and}$$

$$A(p) \sim B(q)$$

Concatenate the windows of p in source pair, and that of q in the target pair, then match

Applications

- Learning a complicated filter from data
- Super-resolution
- Exemplar-based NPR
- Texture synthesis!

Training



Unfiltered source (A)




Filtered source (A')

[Hertzman et al. SIGGRAPH01]



Unfiltered target (B)



Learned filtered target (B')

[Hertzman et al. SIGGRAPH01]



Unfiltered target (B)



Learned filtered target (B')

Learn to Blur



Unfiltered source (A)



Filtered source (A')



Unfiltered target (B)



Filtered target (B')

[Hertzman et al. SIGGRAPH01]


Texture by Numbers



Unfiltered source (A)



Filtered source (A')



Unfiltered (B)



Filtered (B')

[Hertzman et al. SIGGRAPH01]

Colorization



Unfiltered source (A)



Filtered source (A')



Unfiltered target (B)




Filtered target (B')

Super-resolution



Unfiltered source




Filtered source

[Hertzman et al. SIGGRAPH01]

Super-resolution (result)





Unfiltered target



Filtered target

[Hertzman et al. SIGGRAPH01]

Super-resolution



Unfiltered source



Filtered source

Super-resolution (result)



[Hertzman et al. SIGGRAPH01]

Wrap-up: The two steps of texture synthesis

- **Modeling** → Visual fidelity
 - How to estimate the stochastic process from a given finite texture sample
 - Both used MRF model (locality and stationary)
- **Sampling** → Computational cost
 - How to develop an efficient sampling procedure based on the model
 - Pixel by pixel VS patch-based
 - Sampling under guidance: texture transfer, image analogy

Two problems remains

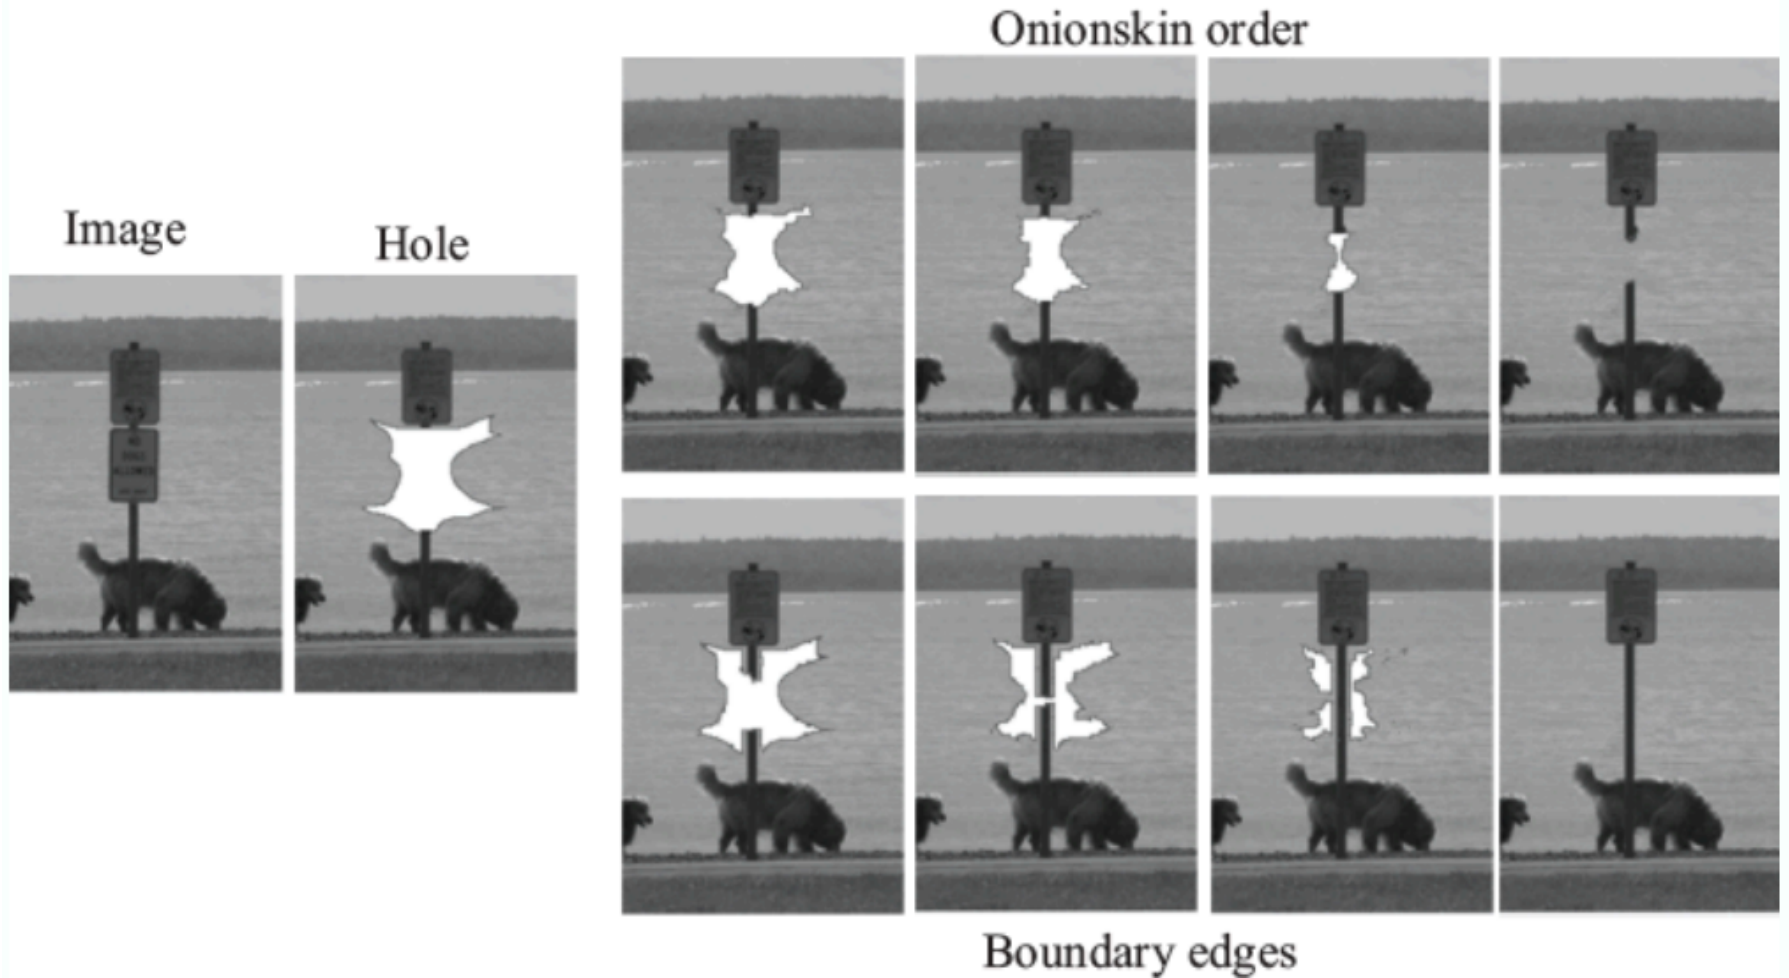
- Preserving scene structure
 - Prioritize synthesise order [*Criminisi et al. CVPR03*]
- Efficient patch matching
 - Fast approximate NN search [*Barnes et al. SIGGRAPH09*]

Inpainting



[Criminisi et al. CVPR03]

Synthesize Order Matters




Choose the order

- Confidence

- Favor unknown pixels with more (reliable) neighbor information

$$C(p) = \frac{\sum_{q \in \Psi_p \cap \Phi} C(q)}{|\Psi_p|}$$


Ψ_p : Neighborhood of p



Choose the order

- Data term
 - Favor starting from strong edges (indication of high saliency for structures)

$$D(p) = \frac{|\nabla I_p^\perp \cdot \mathbf{n}_p|}{\alpha}$$



Choose the order

- Priority


- Balance between confidence and data term

$$P(\mathbf{p}) = C(\mathbf{p}) - D(\mathbf{p})$$

- Algorithm

- Extract the manually selected initial front $\delta\Omega^0$.
- Repeat until done:
 - 1a. Identify the fill front $\delta\Omega^t$. If $\Omega^t = \emptyset$, exit.
 - 1b. Compute priorities $P(\mathbf{p}) \quad \forall \mathbf{p} \in \delta\Omega^t$.
 - 2a. Find the patch $\Psi_{\hat{\mathbf{p}}}$ with the maximum priority, *i.e.*, $\Psi_{\hat{\mathbf{p}}} \mid \hat{\mathbf{p}} = \arg \max_{\mathbf{p} \in \delta\Omega^t} P(\mathbf{p})$
 - 2b. Find the exemplar $\Psi_{\hat{\mathbf{q}}} \in \Phi$ that minimizes $d(\Psi_{\hat{\mathbf{p}}}, \Psi_{\hat{\mathbf{q}}})$.
 - 2c. Copy image data from $\Psi_{\hat{\mathbf{q}}}$ to $\Psi_{\hat{\mathbf{p}}}$.
 3. Update $C(\mathbf{p}) \quad \forall \mathbf{p} \mid \mathbf{p} \in \Psi_{\hat{\mathbf{p}}} \cap \Omega$

Results



More results



Input



Result

More results



Input



Masked target region



Result

More results



a



b



c



d



e



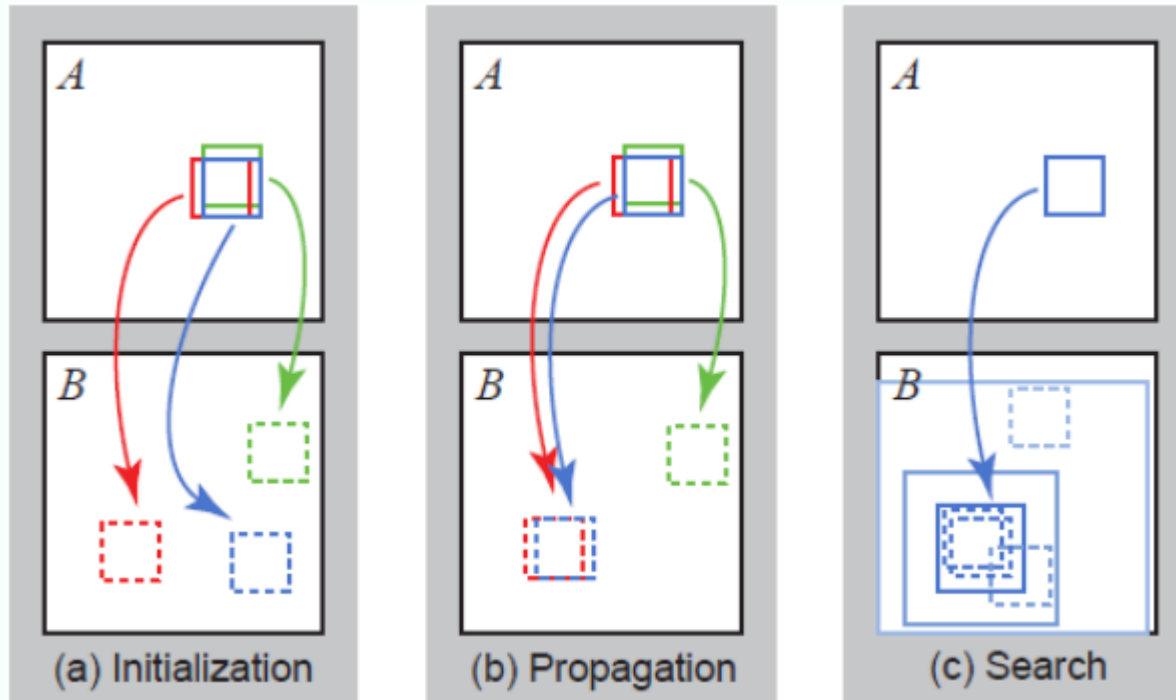
f

[Criminisi et al. CVPR03]

Fast Approximate NN search

- Brute force search
 - Sliding window
 - Exact answer , but too slow
- Better idea
 - **Local coherence**: nearby windows have high probability to match nearby windows in the source image

Fast ANN search



- (a) Random initialization,
- (b) Propagate good matches  **20-100x speedup**
- (c) Search nearby


Add User Constraints

- Mark an hole to fill in (standard process)
- Add extra label, partly inside the hole, partly outside
- Limit the search space for labeled pixels inside the hole to outside regions with the same label

Add User Constraints



(d) input




(e) hole



(f) completion (close up)



(g) same input



(h) hole and guides




(i) guided (close up)

More results



(a) input



(b) hole and guides



(c) completion result

More results



(a) original



(b) hole+constraints



(c) hole filled



(d) constraints



(e) constrained retarget



(f) reshuffle

Retargeting

- Make an image bigger or smaller in one direction
 - e.g. Change aspect ratio
- Traditional
 - Interpolation: content distortion
 - Cut off pixels: lose important content
- Seam Carving
 - Cut out a curve of pixels that “does not matter much”

Example



input



seam



scaling



cropping

[Avidan&Shamir. SIGGRAPH07]

Find a seam


- Define an optimal seam
 - Vertical seam, horizontal seam
 - Energy
 - **Gradient**
 - Entropy, HOG, Segmentation, Saliency, Corner detector, eye gaze

$$e_1(I) = \left| \frac{\partial}{\partial x} I \right| + \left| \frac{\partial}{\partial y} I \right|$$



- Optimal seam: A seam with minimum energy \rightarrow dynamic programming

Find a seam



Initialize: $e_{i,j} = \left| \frac{\partial}{\partial x} I_{ij} \right| + \left| \frac{\partial}{\partial y} I_{ij} \right|$

for $i = 2:h$; for $j = 1:w$

$E_{i,j} = e_{ij} + \min(E_{i-1,j-1}, E_{i-1,j}, E_{i-1,j+1})$

$k_{i,j} = \operatorname{argmin}(E_{i-1,j-1}, E_{i-1,j}, E_{i-1,j+1})$


end; end

Exactly the same DP, different energy term
(Compare with the algorithm of image quilting (P33)!!)


Visualization of seams



Image



Horizontal seams



Vertical seams

Blue: low cost seams
Red: high cost seams

Results



(a) Original



(b) Crop




(c) Column



(d) Seam

More results



input




output


Difference Energies make a difference




(a) Original




(b) e_1



(d) e_{HoG}




(c) $e_{Entropy}$




(e) Segmentation and L_1

Seam insertion




- Finding and inserting the optimum seam will most likely insert the same seam again and again, resulting in (b)
- Instead, inserting the seams in order of removal (c), achieves better result (d)

More results



input



widening by seam insertion

[Avidan&Shamir. SIGGRAPH07]

Seam carving and insertion



Input



Seam Carving result



Rescaling by interpolation

Failure case: need extra constraints



Figure 14: Retargeting the left image with e_1 alone (middle), and with a face detector (right).

Constrained Retargeting


Main idea: user supplied constraints (mask), high cost for seam crossing the masked region



Input + constraints




Retarget
without constraints




Retarget
with constraints


More results



Input + constraints




Retarget
without constraints



Retarget
with constraints

More results



Input + constraints



Retarget
without constraints



Retarget
with constraints

More results



Input + constraints




Retarget
without constraints



Retarget
with constraints


Local scale editing



(a) building marked by user



(b) scaled up, preserving texture



(c) bush marked by user



(d) scaled up, preserving texture.

Reshuffling




input

Reshuffled

[Barnes et al. SIGGRAPH09]

More Reshuffling



More Reshuffling



(a)

(b)


(c)



(d)

(e)

(f)



http://gfx.cs.princeton.edu/pubs/Barnes_2009_PAR/index.php

Summary

- Texture
 - Stationary
 - Locality
- Scene
 - Texture
 - Structure (edges, object contours, etc)
- Two crucial algorithmic points
 - Nearest neighbor search
 - Dynamic programming

Summary (2)

- Patch representation
 - RGB, luminance, CIELAB...
 - Gradient, HOG, Harris corner detector, steerable filter
 - Saliency
 - Entropy
- Patch distance
 - L1 norm, L2 norm, sum of square
 - Gaussian weighted sum of square
- Patch match algorithm
 - Brute force NN
 - Approximate NN: local coherence
 - Tree structure: kd-tree (used in the image analogy method)
 - Dimension reduction: PCA, vector quantization
- Multi-scale representation
 - Gaussian pyramid (used in the image analogy method)