



Texture

- What is texture?
 - Easy to recognize, hard to define
 - Deterministic textures (“thing-like”)
 - Stochastic textures (“stuff-like”)
- Tasks
 - Discrimination / Segmentation
 - Classification
 - Texture synthesis
 - Shape from texture
 - Texture transfer
 - Video textures


Texture Discrimination



Shape from Texture



Modeling Texture



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

Malik & Perona's Filters

Gabor filter kernels - product of symmetric Gaussian with oriented sinusoid


⇒ smoothed derivative kernels

DOG filters - difference of oriented Gaussians

+ DOG filters

⇒ spot and bar filters are many scales, orientations, and phases


$$\begin{cases} G_{\text{symmetric}}(x,y) = \cos(k_x x + k_y y) \exp\left\{-\frac{x^2 + y^2}{2\sigma^2}\right\} \\ G_{\text{antisym}}(x,y) = \sin(k_x x + k_y y) \exp\left\{-\frac{x^2 + y^2}{2\sigma^2}\right\} \end{cases}$$



Steerable (i.e., Oriented) Pyramids


- Laplacian pyramid not appropriate for texture analysis because it does not encode orientation info.

⇒ Oriented pyramids (Simoncelli)




Steerable Pyramids


- Multiresolution, multi-orientation image decomposition



Texture Synthesis [Efros & Leung, ICCV 99]



Synthesizing One Pixel



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

Markov Random Field


A Markov random field (MRF)

- generalization of Markov chains to two or more dimensions



First-order MRF:

- probability that pixel X takes a certain value given the values of neighbors $A, B, C,$ and D :

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




- Higher order MRF's have larger neighborhoods



Markov Chain

Markov Chain

- a *sequence* of random variables X_1, X_2, \dots, X_n
- 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**:

$$p(x_t | x_{t-1})$$


- The above is actually a *first-order* Markov chain
- An N 'th-order Markov chain: $p(x_t | x_{t-1}, \dots, x_{t-N})$

Really Synthesizing One Pixel




- An exact neighborhood 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

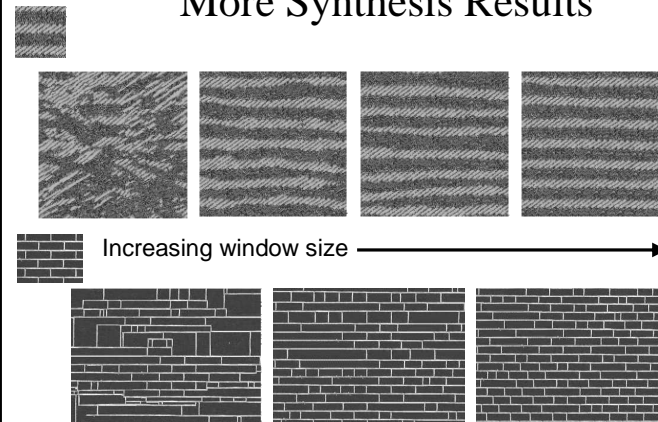


- Starting from the initial image, “grow” the texture one pixel at a time

Window Size Controls Regularity




More Synthesis Results




More Results


reptile skin




aluminum wire



Failure Cases

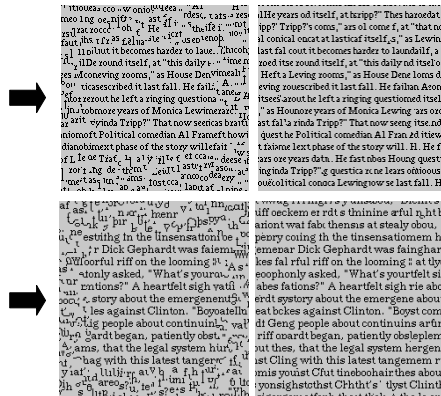
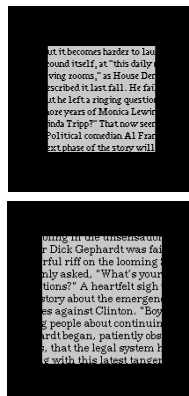


Growing garbage




Verbatim copying

Image-Based Text Synthesis




Efros & Leung '99 Extended




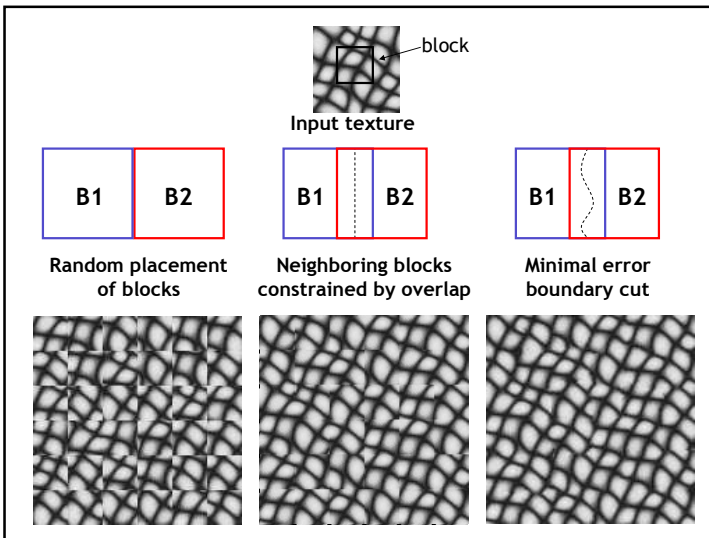
Synthesizing a block

non-parametric
sampling



Input image

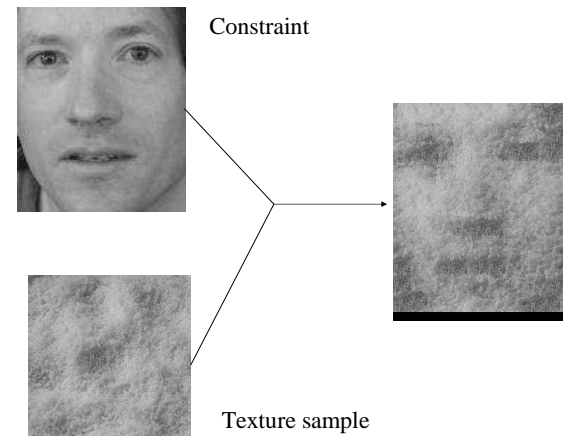
- Observation: neighbor pixels are highly correlated
- Idea: unit of synthesis = block
 - Exactly the same but now we want $P(B|N(B))$
 - Much faster: synthesize all pixels in a block at once



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 the only problem is connecting them together

Texture Transfer




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 luminance of underlying image at that spot

parmesan



rice



Shape from Texture

- Main Idea: Projection distorts texture geometry that depends on surface shape and geometry
- Witkin's Method
 - No model for texels
 - Assume natural textures do not mimic projection effects
 - Isotropy Assumption: All surface orientations equally likely and all edge orientations equally likely
- ⇒ 1) Model geometric distortion on edge orientations
- 2) Statistically model distribution of surface orientations and surface marking orientations

Imaging process distorts surface texture in 2 ways:

1. Distance

Farther objects appear smaller
Area subtended by solid angle Ω
 $= d^2 \Omega$ where d = distance to surface

2. Orientation



As slant α increases,
foreshortening makes point
appear closer together as
 $\frac{1}{\cos \alpha}$


- Features: Edge orientations of surface markings

- Geometric Model
Given curve $C(s)$ on surface \mathcal{S}
let $\beta(s)$ = direction of $C(s)$'s tangent at s



β = angle between $\beta(s)$
and x-axis is \mathcal{S}

What is projection of $\beta(s)$, β ,
 $C(s)$ into image I ?



Goal: Given image of a planar
surface, recover (r, t)

$$0 \leq r \leq \frac{\pi}{2}$$

$$-\frac{\pi}{2} \leq t \leq \frac{\pi}{2}$$


Assume orthographic projection
(\Rightarrow foreshortening effects only)

Let $C^*(s)$ = orthographic projection
of $C(s)$ in I

To compute $C^*(s)$:

1. Put $C(s)$ in I using I 's coords

$$\beta(s) = [\cos \beta, \sin \beta]$$



2. Rotate I by (r, t)
Initially, assume $t = 0$ (= tilt dir)

$$\Rightarrow (x, y, 0) \xrightarrow{R} (x \cos r, y, x \sin r)$$

3. Orthographic projection onto I plane

$$\Rightarrow (x, y, z) \rightarrow (x, y)$$

Combining 3 steps:

Edge orientation $[\cos \beta, \sin \beta]$ in \mathcal{S}
projects to $[\cos \beta \cos \sigma, \sin \beta]$ in \mathcal{I}

$$\Rightarrow \tan \beta^* = \frac{\sin \beta}{\cos \beta \cos \sigma} = \frac{\tan \beta}{\cos \sigma}$$


$$\Rightarrow \beta^* = \tan^{-1} \left(\frac{\tan \beta}{\cos \sigma} \right)$$

= angle b/w observed edge orient.
and tilt direc ($\pm \pi$ -axis)

More generally, if $\tau \neq 0$ then

$$\alpha^* = \tan^{-1} \left(\frac{\tan \beta}{\cos \sigma} \right) + \tau$$

where α^* = observed edge
orientation wrt
 \mathcal{I} 's x-axis



- Isotropy Assumption \Rightarrow all surface orientations equally likely ~~and~~ and all edge orientations equally likely

$$\Rightarrow p.d.f(\beta, \sigma, \tau) = \frac{1}{\pi} \cdot \frac{1}{\pi} \cdot \sin \sigma$$

\swarrow likelihood of $\beta = \beta_0$ in $[0, \pi]$ \swarrow likelihood of $\tau = \tau_0$ in $[0, \pi]$ \swarrow likelihood of $\sigma = \sigma_0$
 $= \frac{\sin \sigma}{\pi^2}$

Find $p.d.f(\alpha^*(\beta) | \sigma, \tau)$

$$= \frac{1}{\pi} \frac{\cos \sigma}{\cos^2(\alpha^* - \tau) + \sin^2(\alpha^* - \tau) \cos^2 \sigma}$$

i.e. says what histogram should look like for given (σ, τ) !

Assuming edge orientations in \mathcal{I} are independent \Rightarrow

$$p.d.f.(A^* = \{\alpha_1^*, \dots, \alpha_n^*\} | \sigma, \tau) = \prod_{i=1}^n p.d.f.(\alpha_i^* | \sigma, \tau)$$

By Bayes Rule:

$$p.d.f.(\sigma, \tau | A^*) = \frac{\prod_{i=1}^n \frac{\pi^{-2} \sin \sigma \cos \sigma}{\cos^2(\alpha_i^* - \tau) + \sin^2(\alpha_i^* - \tau) \cos^2 \sigma}}{\int \int \prod_{i=1}^n \frac{\pi^{-2} \sin \sigma \cos \sigma}{\cos^2(\alpha_i^* - \tau) + \sin^2(\alpha_i^* - \tau) \cos^2 \sigma} d\sigma d\tau}$$

Maximum likelihood estimator for $(\sigma, \tau) \Rightarrow$

1. Apply Edge operator
2. Compute edge orientation histogram A^*
3. For each (σ, τ) compute $p.d.f.(\sigma, \tau | A^*)$
4. Select (σ, τ) which is maximum.

Shape from Texture Using Texels

- Projective distortion changes size of texels due to distance, and shape of texels due to foreshortening

1. Detect Texels

- $\nabla^2 G_\sigma$ detectors find centers of "spots" of varying sizes

- Connected components analysis used to define texels

2. Estimate single planar surface that is maximally consistent with texels "painted" on surface (no micro-relief)

Heuristic: Texel-area gradient:

Area of texels decreases w/ distance and slant angle. Fastest in direction of tilt

- Approximate texel area using bounding box:



$$A_i \approx U_i F_i$$


- Relate image texel area, A_i , to physical texel area, scene plane orientation, (σ, τ) , and angle to texel from optical axis, θ

$$A_i = A_c (1 - \tan \theta \tan \sigma)^3$$


where A_c = area of texel at image center

$$\theta = \tan^{-1}((x \cos \tau + y \sin \tau)(r/f))$$

where (x, y) = texel center coords
 r = image radius
 f = focal length



- Relate to dimensions of texel's bounding box:



$$\begin{aligned} \theta &= \tan^{-1}(x_t(r/f)) & r &= \text{image radius} \\ &= \tan^{-1}((x \cos \tau + y \sin \tau)(r/f)) \\ \Rightarrow F_i &= F_c \frac{f}{d} \cos \theta (1 - \tan \theta \tan \sigma)^2 \\ &= F_c (1 - \tan \theta \tan \sigma)^2 \end{aligned}$$

where F_c = foreshortened dimension of texel at image center

Similarly,

$$U_i = U_c (1 - \tan \theta \tan \sigma)$$

$$\begin{aligned} A_i &= F_i U_i \\ &= A_c (1 - \tan \theta \tan \sigma)^3 \end{aligned}$$

- Search (A_c, σ, τ) space to maximize fit with observed texel areas, A_i 's.
- Discretize A_c, σ, τ values
- Use coarse-to-fine search

```

for  $A_c = \min A_c$  to  $\max A_c$  do
  for  $\theta = 0$  to  $\pi/2$  do
    for  $\tau = 0$  to  $\pi$  do
      fit := 0
      for  $t=1$  to num-textels do
        actual-area := area( $t$ )
        expected-area :=  $A_c (1 - \tan \theta \tan \tau)^3$ 
        region-fit :=  $\frac{\max(\text{expected}, \text{actual})}{\min(\text{expected}, \text{actual})}$ 

        fit := fit +
          (actual-area *
            |textl-contrast| *
             $\exp - \left\{ \frac{\text{region-fit}^2}{4} \right\}$ )
      od
      if fit > best-fit
      then best-fit best-fit := fit
    od
  od
od

```

Recovering Shape By Purposive Viewpoint Adjustment

Kiriakos N. Kutulakos Charles R. Dyer

Computer Sciences Department
University of Wisconsin
Madison, Wisconsin USA

Approaches to Recovering Shape

- Range sensors
Accuracy, distance and resolution limited
- Stereo
Surface texture required
- Shape from Shading
Surface reflectance characteristics required
- Shape from (Static) Contour
Ambiguous: many-to-1 mapping from shape to contour


Active Shape-Recovery

How can we recover surface shape using an observer able to move?


- Current approaches
 - Use a shape-from-motion module (e.g., [Cipolla & Blake; ICCV90])
 - Known viewer velocities and accelerations
 - Compute velocities and accelerations of image points
- Our approach
 - Control position relative to the surface
 - Maintain fixation
 - Measure relative viewing direction changes
 - Compute occluding contour curvatures

Motivation for the Approach

Some views provide more information than others about the shape of a surface:




Arbitrary view



Side view

- An active observer can use these views to recover shape information

Goal: Recover shape for points on the visible rim




- Recover
 1. Principal directions
 2. Principal curvatures
- Assume orthographic projection

Occluding Contour


Rim: Points where surface tangent plane contains the visual ray, $\vec{v} \cdot \vec{n} = 0$ (aka *Limb*, *Contour Generator*)


Visible Rim: Rim points and in view (i.e., not occluded)

Occluding Contour: Projection of the visible rim on image. Collection of open and closed smooth curves; endpoints are Cusps or T-junctions



Silhouette





Using the Occluding Contour

Occluding contour properties

- Dependence on viewpoint & shape well-understood
- Provides shape in absence of markings & surface reflectance information
- Can be efficiently tracked (Blake *et al.*, 1993)
- Recoverable from
 - stereo (Vaillant & Faugeras, 1992)
 - viewpoint control (Kutulakos & Dyer, 1994)

Properties of Occluding Contour

- Geometry is surface-dependent
- Projection of a limited set of surface points
- Geometry is viewpoint-dependent


Assumptions

- Smooth, opaque, stationary object (can be non-convex)
- Parallel projection
- Image features used: occluding contour only
- Observer moves on a sphere around object
- Angular changes in viewpoint known

Overview

- Basic steps.


1. Select a point on the visible rim (elliptic or hyperbolic)



2. Change position to recover the local surface shape at that point
3. Select a new point for shape-recovery

- Special case: Surfaces of revolution

Normal Sections



Principal curvatures = Normal curvatures along e_1, e_2



Elliptic



Hyperbolic



Parabolic




Planar

Shape from Occluding Contour

Relation between the occluding contour curvature and local surface shape [Blaschke]:

$$k_o^{-1} = k_1^{-1} \cos^2 \phi + k_2^{-1} \sin^2 \phi$$




Implications.

1. $k_o = k_1$ if the viewing direction is along e_2
2. If k_o, k_1, ϕ are known we can find k_2
3. $k_o(\phi)$ has only two maxima and two minima, along e_2 and e_1 respectively

The Shape-Recovery Algorithm


1. Compute k_o for the selected point at initial viewpoint
2. Compute point's tangent plane
3. Determine the direction of increasing k_o on point's tangent plane
4. Move in that direction until k_o is maximized
Now, $k_o = k_1$
5. Measure the angle ϕ between the initial and current viewing direction
6. Compute k_2 from ϕ, k_1 , and the initial value of k_o

Hyperbolic Points



Surfaces of Revolution

Local shape reveals global surface properties




- One of the principal directions corresponds to a side view

Results

Candlestick View 1 (arbitrary)


Degrees per Frame	Viewpoint (Radians)
1.000000	0.743038

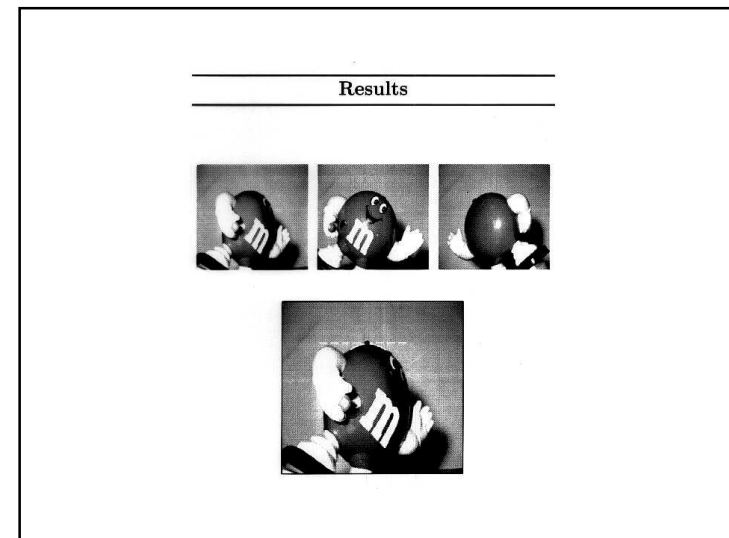





Results


Candlestick View 2 (curvature maximum)

Degrees per Frame	Viewpoint (Radians)
1.000000	1.59825







Results




Curvature variation with viewpoint:



Results



Curvature variation with viewpoint:



Successes and Contributions

Our active approach has a number of features:

- Recovers principal curvatures and principal directions
- Qualitative motion control
- Visual processing consists of curvature measurements on the occluding contour
- Recovers correct axis and generating curve of surfaces of revolution

What is the role of special views in an active context?

Long version of this paper available via ftp at
ftp.cs.wisc.edu (Technical Report #1035)