

Slide sources from A. Efros, S. Seitz, V. Vaish, M. Brown, D. Lowe, S. Lazebnik, P. Perez, R. Szeliski

Photo Collages



Photomosaics



Invented in 1993 by Joseph Francis, and patented by Robert Silvers in 2000





Text-to-Picture Communication

First the farmer gives hay to the goat. Then the farmer gets milk from the cow.





64

1











Reconstruction-based Super-Resolution

Reconstruct HR pixels by a linear combination of LR pixels so that when HR image is projected to LR, it is similar to all the LR images



Unknown Relationship b/w LR images

Photos taken by cameras with unknown translation, rotation, and scale (zoom)



Example-based Super-Resolution From a Single Image

> Daniel Glaser, Shai Bagon and Michal Irani



Natural images contain repetitive image content in small patches (e.g., 5 x 5)

Small patches in I are found "as is" in different locations and in other image scales of /

High-res "parent patches" (dashed squares) indicate what the high-res parents of patches in I might look like



Recurring patches in a single low-res image can be regarded as if extracted from multiple low-res images of the *same* high-res image

Employing Cross-Scale Patch Redundancy

- Build a cascade of decreasing resolution images from the input LR image
- For each patch in the LR image, search for its nearest neighbor in the even lower resolution images
- Take the found neighbor's parent in the original LR image and copy it to the HR image, providing a (linear) constraint on the output HR image

















Example-based Super Resolution

William T. Freeman, Thouis R. Jones and Egon C. Pasztor

Algorithm Overview

- Construct a DB of matching LR-HR patches based on a *separate set of natural images*
- Find the most similar patch assignment to generate high-res image



Goal: Given a static scene and a set of images (or video) of it, combine the images into a single "panoramic image" or "panoramic mosaic" that is optically correct





With the Cassini satellite's wide-angle camera aimed at Saturn, Cassini was able to capture 323 images in just over four hours in 2013. This final panorama used 141 of those images taken using red, green and blue spectral filters.

http://www.jpl.nasa.gov/spaceimages/details.php?id=pia17172

Mont Blanc Panorama

- 365 gigapixels, created from 70,000 images
- http://www.in2white.com/



Why Panoramas ?

• Virtual reality walkthroughs



Quicktime VR [Chen & Williams 95]





[Brown 2003]











Wisconsin Coastal Guide Panoramas









Panorama Capture Hardware





Panorama Stitching Algorithm

- 1. Capture Images: Capture a set of images of a static scene
- **2.** Alignment: Compute an image-to-image transformation that will map pixel coordinates in one image into corresponding pixel coordinates in a second image
- **3.** Warp: Warp each image using transform onto output compositing surface (e.g., plane, cylinder, sphere, cube)
- 4. Interpolate: Resample warped image
- 5. Composite: Blend images together so as to hide seams, exposure differences, lens distortion, scene motion, etc.





When can Two Images be Aligned?

Problems

- In general, warping function depends on the depth of the corresponding scene point since perspective projection defined by x' = fx/d, y' = fy/d
- Different views mean, in general, that parts that are visible in one image may be occluded in the other
- Special cases where the above problems can't occur
 - 1. Panorama: Camera rotates about its optical center, arbitrary 3D scene
 - No motion parallax as camera rotates, so depth unimportant
 - 2D projective transformation relates any 2 images (⇒ 8 unknowns)
 - 2. Planar mosaic: Arbitrary camera views of a planar scene
 - 2D projective transformation relates any 2 images

<section-header>









Image Reprojection

 Rather than thinking of this as a 3D reprojection, think of it as a 2D image warp from one image to another







Camera Transformations using Homogeneous Coordinates

- Computer vision and computer graphics usually represent points in Homogeneous coordinates instead of Cartesian coordinates
- Homogeneous coordinates are useful for representing perspective projection, camera projection, points at infinity, etc.
- Cartesian coordinates (x, y) represented as Homogeneous coordinates (wx, wy, w) for any scale factor w ≠ 0
- Given 3D homogeneous coordinates (x, y, w), the 2D Cartesian coordinates are (x/w, y/w)

Homogeneous Coordinates

















Perspective Projection

 Perspective projection with a pinhole camera is a matrix multiply using homogeneous coords!

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/f & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \begin{bmatrix} X \\ Y \\ Z/f \end{bmatrix} \Rightarrow \left(\frac{fX}{Z}, \frac{fY}{Z}\right)$$

This 3 x 4 matrix is called the camera perspective projection matrix





Projective Camera

- More general than Perspective Camera matrix, $\boldsymbol{\Pi}$
- Transformation matrix has only 11 DOFs since only the *ratios* of elements are important





Affine Camera

- Most general *linear* transformation
- 8 DOFs
- Reasonable assumption when scene objects are far away from camera (relative to focal length)







where does the pixel at coordinates (10, 5) in image 2 project to in image 1?



- 1. Convert point in image 2 to homogeneous coordinates: $p = (10, 5, 1)^T$
- **2**. Compute q = Hp:

1*10 + 2*5 + 3*1 = 23	7 [- 1	2	3	10	
4*10+1*5+0*1=45		4	1	0	5	
1*10+1*5+3*1=18		1	1	3	1	

3. Convert *q* to Cartesian coordinates:

 $q = (sx, sy, s)^{T} = (23, 45, 18)^{T}$ or, in Cartesian coords, (23/18, 45/18) = (1.28, 2.5) in image 1







Combining Images into Panoramas

- Theorem: Any 2 images of an arbitrary scene taken from 2 cameras with same camera center are related by p₂ ≈ K R K⁻¹ p₁ where p₁ and p₂ are homogeneous coords of 2 corresponding points, K is 3 x 3 camera calibration matrix, and R is 3 x 3 rotation matrix
- **K R K**⁻¹ is 3 x 3 matrix called the "homography induced by the plane at infinity"

Panorama Stitching Algorithm

- **1. Capture images:** Take a sequence of images , $I_1, ..., I_n$, from the same position by rotating the camera around its optical center
- Alignment: Compute an image-to-image transformation that will map pixel coordinates in one image into corresponding pixel coordinates in second image
- **3.** Warp: Warp each image using transform onto output compositing surface (e.g., plane, cylinder, sphere, cube)
- 4. Interpolate: Resample warped image
- 5. Composite: Blend images together so as to hide seams, exposure differences, lens distortion, scene motion, etc.

Finding the Homographies

How can we compute the homographies required for aligning a set of images?

Alignment

Direct methods

- Search over space of possible image warps and compare images by pixel intensity/color matching to find warp with minimum matching error
- Feature-based methods
 - Extract distinctive (point) features from each image and match these features to establish global correspondence, and then estimate geometric transformation

Alignment

- Direct method: Use image-based (intensity) correlation to determine best matching transformation
 - No correspondences needed
 - Statistically optimal (gives maximum likelihood estimate)
 - Useful for local image registration
- Feature-based method: Find feature point correspondences, and then solve for unknowns in "motion model"
 - Requires reliable detection of a sufficient number of corresponding features, at sub-pixel location accuracy

Direct Method for Computing Panoramic Mosaics

- Motion model is 2D projective transformation, so 8 parameters (DOFs)
- Assuming small displacement, minimize SSD error
- Use nonlinear minimization algorithm to solve

Panorama Stitching Algorithm

- **1.** Capture Images: Take a sequence of images , I₁, ..., I_n, from the same position by rotating the camera around its optical center
- Alignment: Compute an image-to-image transformation that will map pixel coordinates in one image into corresponding pixel coordinates in second image
- 3. Warp: Warp each image using transform onto output compositing surface (e.g., plane, cylinder, sphere, cube)
- 4. Interpolate: Resample warped image
- **5. Composite**: Blend images together so as to hide seams, exposure differences, lens distortion, scene motion, etc.



Finding the Homography





- By matching features across images
 - What features should we match?
 - How many features?

Finding the Homography

What features do we match across images ?

- Pixel values ?
- Edges ?
- Corners ?
- Lines ?
- Other features ?

Finding the Homography

What features do we match across images ?

- Pixel values
- Edges
- Corners
- Lines
- Feature points







Solving for Homography: 1 Pair of Corresponding Points





Alternatively, solve $\mathbf{Ah} = \mathbf{b}$ where \mathbf{A} is a $3n \ge 9$ matrix, \mathbf{b} is a $3n \ge 1$ matrix, \mathbf{h} is the same $9 \ge 1$ matrix to be solved for, and, for 1 pair of correspondences, $p1 = (x_1, x_2)$ and $p2 = (x_1, y_1)$:

 $\mathbf{b} = \begin{bmatrix} x_1' & y_1' & 1 \end{bmatrix}^{\mathrm{T}}$

 $\mathbf{A} = \begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & 0 & 0 \end{bmatrix}$





Warping Methods

• Forward, square-pixel based

- Consider pixel at (u,v) as a unit square in source image. Map square to a quadrilateral in destination image
- Assign (u,v)'s gray level to pixels that the quadrilateral overlaps



 Integrate source pixels' contributions to each output pixel.
 Destination pixel's gray level is weighted sum of intersecting source pixels' gray levels, where weight proportional to coverage of destination pixel

- Avoids holes, but not folds, and requires intersection test

Warping Methods

• Backward, point-based

- For each destination pixel at coordinates (x,y), apply backward mapping, U, V, to determine real-valued source coordinates (u,v)
- Interpolate gray level at (u,v) from neighboring pixels, and copy gray level to (x,y)



- Interpolation may cause artifacts such as aliasing, blockiness, and false contours
- Avoids holes and folds problems
- Method of choice

Backward Warping

```
    for x = xmin : xmax
    for y = ymin : ymax
    u = U(x, y);
    v = V(x, y);
    B(x, y) = A(u, v);
    end
```

end

- But (*u*, *v*) may not be at a pixel in A
- (*u*, *v*) may be out of A's domain
- If **U** and/or **V** are discontinuous, A may not be connected!

Pixel Interpolation

- Nearest-neighbor (0-order) interpolation
 - A(u, v) = gray level at **nearest** pixel (i.e., round (u, v) to nearest integers)
 - May introduce artifacts if image contains fine detail
- Bilinear (1st-order) interpolation
 - Given the 4 nearest neighbors, A(0, 0), A(0, 1), A(1, 0), A(1, 1), of a desired point A(u, v), where $0 \le u \le 1$ and $0 \le v \le 1$, compute gray level at A(u, v):
 - Interpolate linearly between A(0,0) and A(1,0) to obtain A(u,0)
 - Interpolate linearly between A(0,1) and A(1,1) to obtain A(u,1)
 - Interpolate linearly between A(u,0) and A(u,1) to obtain A(u,v)
 - Combining all three interpolation steps into one we get:
 - A(u,v) = (1-u)(1-v) A(0,0) + (1-u)v A(0,1) + u(1-v) A(1,0) + uv A(1,1)
- Bicubic spline interpolation



Example of Backward Warping

- Goal: Define a transformation that performs a scale change, which expands size of image by 2, i.e., U(x) = x/2
- A = 0 ... 0 2 2 2 0 ... 0
- 0-order interpolation, i.e., $u = \lfloor x/2 \rfloor$ B = 0 ... 0 2 2 2 2 2 2 0 ... 0
- Bilinear interpolation, i.e., u = x/2 and average 2 nearest pixels if u is not at a pixel

 $\mathsf{B}=0\,\dots\,0\,1\,2\,2\,2\,2\,2\,1\,0\,\dots\,0$

Panoramic Stitching Algorithm

Input: *N* images from camera rotating about center

- 1. Detect point features and their descriptors in all images
- 2. For adjacent images:
 - 1. Match features to get pairs of corresponding points
 - 2. [Optional: Eliminate bad matches]
 - 3. Solve for homography
- 3. Project images on common "image plane"
- 4. Blend overlapping images to obtain panorama

Panorama "Shape" Depends on Output Image Plane





Panorama Images

- Large field of view ⇒ should *not* map all images onto a plane
- Instead, map onto cylinder, sphere, or cube
- With a cylinder, first warp all images from rectilinear to cylindrical coordinates, then combine them
- "Undistort" (perspective correct) image from this representation prior to viewing





Beyond Panoramas: General Camera Motion



Affine model okay in practice when scene objects are far away from camera viewpoints relative to focal length, *f*, and baseline, *B*

$$\begin{bmatrix} x, y, 1 \end{bmatrix}^{T} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \qquad \begin{cases} x = a_{11}u + a_{12}v + a_{13} \\ y = a_{21}u + a_{22}v + a_{23} \end{cases}$$



Method 1: Averaging
• For each output pixel, compute average of
overlapping warped pixels, *x*:

$$C(x) = \sum_{k} w_{k}(x) \tilde{I}_{k}(x) / \sum_{k} w_{k}(x)$$
where $\tilde{I}_{k}(x)$ are the warped images and $w_{k}(x)$
is 1 at valid pixels and 0 elsewhere
• Weakness: Doesn't work well with exposure
differences, mis-registration, etc.

Method 2: Weighted Averaging

- Aka Feathering or Alpha Blending
- Weight pixels near center of each warped image more heavily than pixels near image border
- If image has holes, also down-weight values near border of hole
- Implement by computing a distance map = distance to nearest border pixel
- Weakness: blurs details such as edges















- Multi-resolution technique using image pyramid
- Hides seams but preserves sharp detail





Edge Detection in 2D

- Let *I*(x,y) be the image intensity function. It has derivatives in all directions
 - $\partial I(x, y)/\partial x = \lim I(x + \Delta x, y) I(x, y) / \Delta x \approx I(u+1, v) I(u, v)$
 - Gradient of *I*(x, y) is a vector ∇*I*(x, y) = [∂*I*/∂x, ∂*I*/∂y]^T
 specifying the direction of greatest rate of change in intensity (i.e., perpendicular to the edge's direction)
 - From gradient, we can determine the **direction** in which the first derivative is highest, and the **magnitude** of the first derivative in that direction
 - Magnitude = $[(\partial I/\partial x)^2 + (\partial I/\partial y)^2]^{1/2}$
 - Direction = $\tan^{-1} (\partial l / \partial y) / (\partial l / \partial x)$









Finite Differences and Noise

- Finite difference filters respond strongly to noise
 - obvious reason: image noise results in pixels that look very different from their neighbors
- Generally, the more noise, the stronger the response
- What can be done?
 - intuitively, most pixels in images look quite a lot like their neighbors
 - this is true even at an edge; along the edge they're similar, across the edge they're not
 - suggests that *smoothing* the image should help, by forcing pixels different from their neighbors (=noise pixels?) to look more like their neighbors

Finite Differences Responding to Noise



Increasing noise \rightarrow

(this is zero mean, additive, Gaussian noise)

Summary of Basic Edge Detection Steps

- 1. Smooth the image to reduce the effects of local intensity variations (i.e., noise)
- Differentiate the smoothed image using a digital gradient operator that assigns a magnitude and direction of the gradient at each pixel
 - Mathematically, we can apply the digital gradient operator to the digital smoothing filter, and then just convolve the differentiated smoothing filter to the image

Summary of Basic Edge Detection Steps

- 3. Threshold the gradient magnitude to eliminate low contrast edges
- 4. Apply a **non-maximum suppression** step to thin the edges to single pixel wide edges
 - Smoothing will produce an image in which the contrast at an edge is spread out in the neighborhood of the edge
 - Thresholding operation will produce thick edges

The Scale Space Problem

- Usually, any single choice of smoothing operator does not produce a good edge map
 - large amount of smoothing will produce edges from only the largest objects, and they will not accurately delineate the object because the smoothing reduces shape detail
 - small amount of smoothing will produce many edges and very jagged boundaries of many objects
- Scale-space approaches
 - detect edges at a range of scales [s₁, s₂]
 - combine the resulting edge maps









0 1 0

Laplacian Examples • l = ... 222888...• $\nabla^2 = 1 \cdot 21$ $\Rightarrow \nabla^2 l = ... 0006 \cdot 6000...$ • l = ... 2225888... $\Rightarrow \nabla^2 l = ... 00030 \cdot 3000...$

Laplacian Edge Detectors

- Laplacians are also combined with smoothing for better edge detection
 - Take the Laplacian of a Gaussian-smoothed image called the Laplacian-of-Gaussian (LoG), Mexican Hat operator, Differenceof-Gaussians (DoG), Marr-Hildreth, V²G operator
 - Locate the zero-crossings of the operator
 - these are pixels whose LoG is positive and which have at least one neighbor whose LoG is negative or zero
 - Often, include a final step that measures the gradient at these points to eliminate low contrast edges



LoG Properties

- Linear, shift invariant ⇒ convolution
- Circularly symmetric ⇒ isotropic
- Size of LoG operator approximately 6σ x 6σ
- LoG is separable
- LoG \approx G_{σ_1} G_{σ_2}, where σ_1 = 1.6 σ_2
- Analogous to spatial organization of receptive fields of retinal ganglion cells, with a central excitatory region and an inhibitory surround

Gaussian Pyramids



Gaussian Pyramid

- Multiresolution, low-pass filter
- Hierarchical convolution
 - G0 = input image
 - $G'_{k}(u, v) = \sum w(m, n) G_{k-1}(u-m, v-n)$; smooth
 - $G_k(u, v) = G'_k(2u, 2v), 0 < u, v < 2^{N-k}$; sub-sample
- w is a small (e.g., 5 x 5) separable generating kernel, e.g., 1/16
 [1 4 6 4 1]
- Cascading *w* is equivalent to applying one large kernel
 - Effective size of kernel at level k = 2M(2^k 1) + 1, where w has width 2M+1
 - Example: Let M=1. If k=1 then equivalent size = 5; k=2 then equivalent size = 13; k=3 then equivalent size = 27

Laplacian Pyramids

- Similar to results of edge detection
- Most pixels are 0
- Can be used for image compression

$$L_1 = g_1 - EXPAND[g_2]$$
$$L_2 = g_2 - EXPAND[g_3]$$
$$L_2 = g_2 - EXPAND[g_3]$$

$$L_3 = g_3 - EXPAND[g_4]$$













Image Compositing by Pyramid Blending

- **Given**: Two 2ⁿ x 2ⁿ images
- **Goal**: Create an image that contains left half of image A and right half of image B
- Algorithm
 - Compute Laplacian pyramids, LA and LB, from images A and B
 - Compute Laplacian pyramid LS by copying left half of LA and right half of LB. Pyramid nodes down the center line = average of corresponding LA and LB nodes ⇒ blend along center line
 - Expand and sum levels of LS to obtain output image S







Combining Apple & Orange using Laplacian Pyramids



Image Compositing from Arbitrary Regions

- Given: Two 2ⁿ x 2ⁿ images and one 2ⁿ x 2ⁿ binary mask
- Goal: Output image containing image A where mask=1, and image B where mask=0
- Algorithm:
 - Construct Laplacian pyramids LA and LB from images A and B
 - Construct Gaussian pyramid GR from mask R
 - Construct Laplacian pyramid LS:
 - $LS_{k}(u, v) = GR_{k}(u, v) LA_{k}(u, v) + (1 GR_{k}(u, v)) LB_{k}(u, v)$
 - Expand and sum levels of LS to obtain output image S





Method 4: Gradient Domain Blending

- Perez, Gangnet and Blake, Poisson Image Editing, *Proc. Siggraph*, 2003
- Aka Poisson Blending or Poisson Cloning
- Similar to Photoshop's "Healing Brush"



Simple Cut-and-Paste Result



Poisson Blending Result



Poisson Blending

- Key Idea: Allow the colors in foreground region to change, but preserve all the details (i.e., edges, corners)
- Blend should preserve the *gradients* of foreground region AND match background colors at seam, without changing the background
- Treat pixel colors as variables to be solved for
 - Minimize squared difference between gradients of foreground region and gradients of output region
 - Keep background pixels constant

 Rather than copying pixels, copy the gradients instead; then compute the pixel color values by solving a Poisson equation that matches the gradients while also satisfying fixed boundary conditions (i.e., pixel color values) at seam









<section-header>

source images

simple cloning

More Results





source images

Poisson blending

Poisson Blending: "Guiding" the Completion

- Use gradients from the source image (i.e., the foreground region) to *guide* the completion
- Find new pixels' values in output image's target region, *f*, so that their *gradients* are close to the gradients (vector field v) of the foreground image, *g*, while holding $f = f^*$ at the

boundary, δΩ



Poisson Blending

- Treat pixels as variables to be solved (with colors)
 - Minimize squared difference between gradients of foreground region and gradients of output pixels
 - Match background's boundary pixels

$\arg\min_{f} \iint_{\Omega} |\nabla f - v|^2 \quad s.t. \ f|_{\partial\Omega} = f^*|_{\partial\Omega}$

Equivalent to solving a Poisson equation, which can be formulated as a discrete quadratic optimization problem and solved using Gauss-Seidel or Jacobi methods

Perez et al. 2003



source images, g

Example Result



source images

Poisson blending result



Note: Target and source images must be (manually) aligned

Poisson Blending: Mixing Gradients

- There are situations where it is desirable to **combine** properties of *f** with those of foreground g, for example to add objects with holes, or partially transparent ones, on top of a textured or cluttered background
- Use gradients of source and/or destination



Figure 6: Inserting objects with holes. (a) The classic method, color-based selection and alpha masking might be time consuming and often leaves an undesirable halo; (b-c) seamless cloning, even averaged with the original image, is not effective; (d) mixed seam-less cloning based on a loose selection proves effective.



Mixing Gradients: Inserting Transparent Objects



Non-linear mixing of gradient fields picks out most salient structure at each pixel

Mixing Gradients: Inserting Objects
with HolesImage: Contract of the state of the st