6.098 Digital a... omputational Photography
6.8... ...ational Photography

Warning:
French Mathematicians inside

Gradient image processing

Bill Freeman
Frédo Durand
MIT - EECS

How was pset 2?

What have we learnt last time?

- **Log is good**
- **Luminance is different from chrominance**
- **Separate components:**
  – Low and high frequencies
- **Strong edges are important**

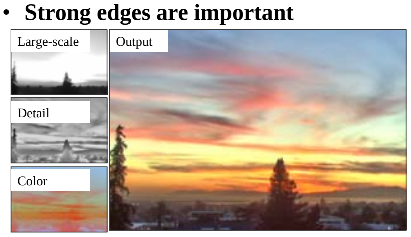

Homomorphic filtering

- **Oppenhein, in the sixties**
- **Images are the *product* of illumination and albedo**
  – Similarly, many sounds are the *product* of an envelope and a modulation
- **Illumination is usually slow-varying**
- **Perform albedo-illumination using low-pass filtering of the log image**

- http://www.cs.sfu.ca/~stella/papers/blairthesis/main/node33.html
- **See also Koenderink "Image processing done right"**
  http://www.springerlink.com/(l1bpumaapconcbjngteojwqv)/app/home/contribution.asp?referrer=parent&backto=issue,11,53;journal,1538,3333;linkingpublicationresults,1:105633,1

What's great about the bilateral filter

- **Separate image into two components**
- **Preserve strong edges**
- **Non-iterative**
  – More controllable, stable
- **Can be accelerated**

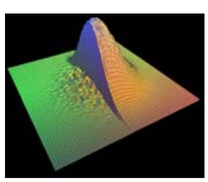

- **Lots of other applications**

Edit materials and lighting

- **With Oh, Chen and Dorsey**

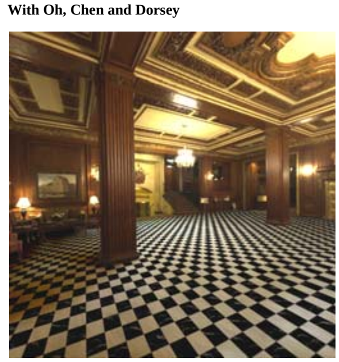

A Simple Relighting Example

- With Oh, Chen and Dorsey

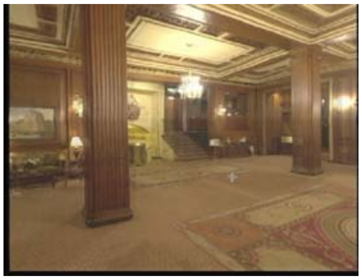

Flash Photography (Elmar Eisemann)

Bilateral filtering on meshes

- http://www.cs.tau.ac.il/~dcor/online_papers/papers/shachar03.pdf
- http://people.csail.mit.edu/thouis/JDD03.pdf

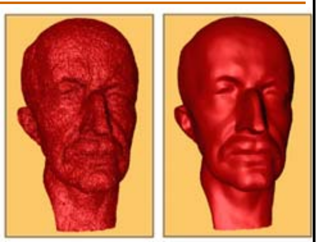

Figure 1: The dragon model (left) is artificially corrupted by Gaussian noise ($\sigma = 1/5$ of the mean edge length) (middle), then smoothed in a single pass by our method (right). Note that features such as sharp corners are preserved.

Questions?

Questions?

Today: Gradient manipulation

Idea:

- **Human visual system is very sensitive to gradient**
- **Gradient encode edges and local contrast quite well**

- **Do your editing in the gradient domain**
- **Reconstruct image from gradient**

∇

- **Various instances of this idea, I'll mostly follow Perez et al. Siggraph 2003 http://research.microsoft.com/vision/cambridge/papers/perez_siggraph03.pdf**

Problems with direct cloning

From Perez et al. 2003

Solution: clone gradient

Gradients and grayscale images

- **Grayscale image: n× n scalars**
- **Gradient: n× n *2D vectors***
- **Overcomplete!**
- **What's up with this?**
- **Not all vector fields are the gradient of an image!**
- **Only if they are curl-free (a.k.a. conservative)**
  - But it does not matter for us

Today message II

- **Variational approach**
  - Express your problem as an energy minimization over a space of functions

- **And we are going to spend our time going back and force between minimization and setting derivatives to zero. Your head will spin.**

Questions?

Seamless Poisson cloning

- **Given vector field *v* (pasted gradient), find the value of *f* in unknown region that optimize:**

Poisson equation with Dirichlet conditions

$$\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

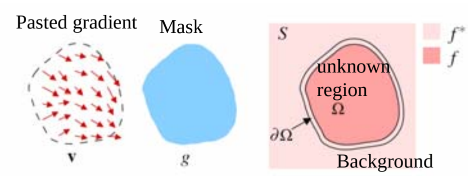

Figure 1: **Guided interpolation notations.** Unknown function f interpolates in domain Ω the destination function f^*, under guidance of vector field $\mathbf{v}$, which might be or not the gradient field of a source function g.

Warning:

- **What follows is not strictly necessary to implement Poisson image editing**
- **But**
  - It helps understand the properties of the equation
  - It helps to read the literature
  - It's cool math

Membrane interpolation

- **What if *v* is null?**
- **Laplace equation (a.k.a. membrane equation)**

$$\min_f \iint_\Omega |\nabla f|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

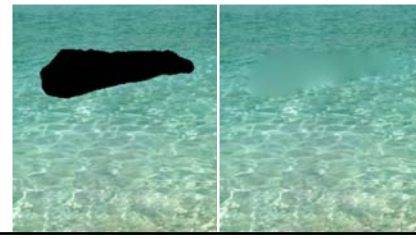

Membrane interpolation

- **What if *v* is null?**
- **Laplace equation (a.k.a. membrane equation)**

$$\min_f \iint_\Omega |\nabla f|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

- **Mathematicians will tell you there is an Associated Euler-Lagrange equation:**

$$\Delta f = 0 \text{ over } \Omega \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

- **Kind of the idea that we want a minimum, so we kind of derive and get a simpler equation**

Calculus

Simplified version:

- **Want to minimize *g(x)* over the space of real values *x***
- **Derive and set g'(x)=0**

- **Now we have a more complex equation: we want to minimize a *variational equation* over the space of functions *f***

$$\min_f \iint_\Omega |\nabla f|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

- **It's a complex business to derive wrt functions**
  - In general, derivatives are well defined only for functions over 1D domains

Derivative definition

- **1D derivative:**

$$\lim_{\epsilon\to 0} \; f'(x) = \frac{f(x+\epsilon) - (f(x)}{\epsilon}$$

- **multidimensional derivative:**
  - For a direction $\vec{w}$, directional derivative is

$$D_{\vec{w}} f(\vec{x}) = \lim_{\epsilon\to 0} \; \frac{f(\vec{x}+\epsilon\vec{w}) - (f(\vec{x})}{\epsilon}$$

- **For functionals ?**
  - Do something similar, replace vector by function

Calculus of variation – 1D

- **We want to minimize $\int_{x_1}^{x_2} f'(x)^2 dx$ with *f(x$_1$)=a, f(x$_2$)=b***
- **Assume we have a solution *f***

Try to define some notion of 1D derivative wrt to a 1D parameter ε in a given direction of functional space:

- **For a perturbation function η(x) that also respects the boundary condition (i.e. $\eta(x_1)=\eta(x_2)=0$) and scalar ε, the integral ∫ (f'(x)+ε η'(x))^2 dx should be bigger than for f' alone**

Calculus of variation – 1D

- $\int (f'(x)+\varepsilon\,\eta'(x))^2\,dx$ **should be bigger than for f' alone**
- $\int f'(x)^2 + 2\,\varepsilon\,\eta'(x)\,f'(x) + \varepsilon^2\eta'(x)^2\,dx$
- **The third term is always positive and is negligible when ε goes to zero**
- **Derive wrt ε and set to zero**
- $\int 2\,\eta'(x) f'(x)\,dx = 0$

Calculus of variation – 1D

$$\int_{x_1}^{x_2} \eta'(x)\; f'(x)\;\mathrm{dx}$$

- **How do we get rid of η ? And still include the knowledge that $\eta(x_1)=\eta(x_2)=0$**
- **When we have an integral of a product and we are playing with derivatives, look into integration by parts**
  - Now how do you remember integration by parts?
  - Integrate one, derive the other
  - It's about the derivative of a product in an integral

$$[gh]_{x_1}^{x_2} = \int_{x_1}^{x_2} \frac{d\,fg}{dx}\mathrm{dx} = \int_{x_1}^{x_2} f'(x)g(x) + f(x)g'(x)\mathrm{dx}$$

Calculus of variation – 1D

$$\int_{x_1}^{x_2} \eta'(x)\; f'(x)\;\mathrm{dx} = 0$$

- **Integrate by parts**

$$\int_{x_1}^{x_2} \eta'(x)\; f'(x)\;\mathrm{dx} = \left[\eta(x) f'(x)\right]_{x_1}^{x_2} - \int_{x_1}^{x_2} \eta(x)\; f''(x)\;\mathrm{dx}$$

- **We know that $\eta(x_1)=\eta(x_2)=0$**
- **We get** $\int_{x_1}^{x_2} \eta(x)\; f''(x)\;\mathrm{dx} = 0$
- **Must be true for any η**
- **Therefore, f''(x) must be zero everywhere**

Intuition

- **In 1D; just linear interpolation!**
  - The min of ∫ f' is the slope integrated over the interval
- **Locally, if the second derivative was not zero, this would mean that the first derivative is varying, which is bad since we want ∫ f' to be minimized**
- **Note that, in 1D: by setting f'', we leave two degrees of freedom. This is exactly what we need to control the boundary condition at x_1 and x_2**

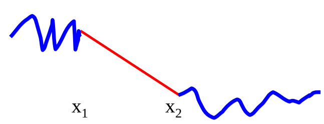

In 2D: membrane interpolation

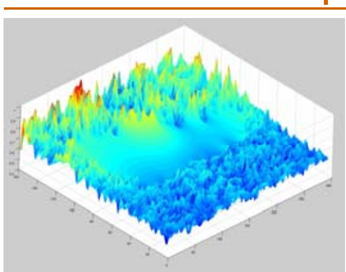

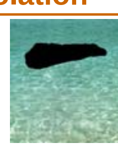

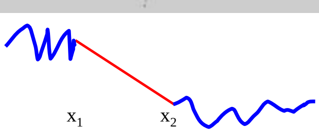

Recap

- **Variational minimization (integral of a functional) with boundary condition**

$$\min_f \iint_\Omega |\nabla f|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

- **Derive Euler-Lagrange equation:**
  - Use perturbation function
  - Calculus of variation. Set to zero. Integrate by parts.

$$\Delta f = 0 \text{ over } \Omega \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

Questions?

What if *v* is not null

Seamlessly paste  onto 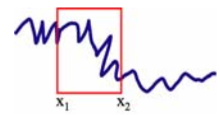

Just add a linear function so that the boundary condition is respected

What if *v* is not null

- **Variational minimization (integral of a functional) with boundary condition**

$$\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega},$$

- **Derive Euler-Lagrange equation:**

$$\Delta f = \operatorname{div}\mathbf{v} \text{ over } \Omega, \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

where $\operatorname{div}\mathbf{v} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}$ is the divergence of $\mathbf{v} = (u, v)$

In 2D, if *v* is conservative

- **If *v* is the gradient of an image g**
- **Correction function** $\hat{f}$ **so that** $f = g + \hat{f}$
- $\hat{f}$ **performs membrane interpolation over Ω:**

$$\Delta\tilde{f} = 0 \text{ over } \Omega, \; \tilde{f}|_{\partial\Omega} = (f^* - g)|_{\partial\Omega}$$

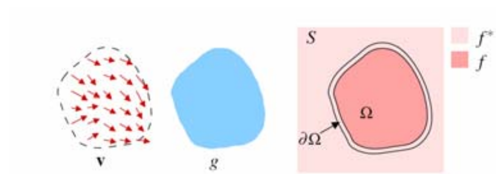

Questions?

Back to practical Poisson editing

sources/destinations

seamless cloning

Discrete Poisson solver

- **Two approaches:**
  - Minimize variational problem $\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$
  - Solve Euler-Lagrange equation $\Delta f = \text{div}\,\mathbf{v} \text{ over } \Omega, \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$

  In practice, variational is best
- **In both cases, need to discretize derivatives**
  - Finite differences over 4 pixel neighbors
  - We are going to work using pairs
    - Partial derivatives are easy on pairs
    - Same for the discretization of v

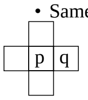

Discrete Poisson solver

- **Minimize variational problem** $\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$

$$\min_{f|_\Omega} \sum_{\langle p,q\rangle \cap \Omega \neq \emptyset} (f_p - f_q - v_{pq})^2, \text{ with } f_p = f_p^*, \text{for all } p \in \partial\Omega$$

Discretized gradient

(all pairs that are in Ω)

Discretized v: g(p)-g(q)

Boundary condition

- **Rearrange and call N_p the neighbors of p**

$$\text{for all } p \in \Omega, \quad |N_p| f_p - \sum_{q \in N_p \cap \Omega} f_q = \sum_{q \in N_p \cap \partial\Omega} f_q^* + \sum_{q \in N_p} v_{pq}$$

Only for boundary pixels

- **Big yet sparse linear system**

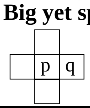

Result (eye candy)

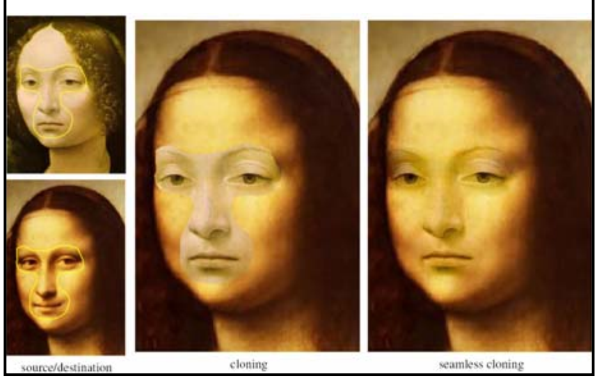

Questions?

Solving big matrix systems

- **Ax=b**
- **You can use Matlab's **
  - But not very scalable
- **In Pset 3, we ask you to implement conjugate gradient**
  - http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf
  - http://www.library.cornell.edu/nr/bookcpdf/c10-6.pdf

Conjugate gradient

An Introduction to
the Conjugate Gradient Method
Without the Agonizing Pain
Edition $1\frac{1}{4}$
Jonathan Richard Shewchuk
August 4, 1994

- "The Conjugate Gradient Method is the most prominent iterative method for solving sparse systems of linear equations. Unfortunately, many textbook treatments of the topic are written with neither illustrations nor intuition, and their victims can be found to this day babbling senselessly in the corners of dusty libraries. For this reason, a deep, geometric understanding of the method has been reserved for the elite brilliant few who have painstakingly decoded the mumblings of their forebears. Nevertheless, the Conjugate Gradient Method is a composite of simple, elegant ideas that almost anyone can understand. Of course, a reader as intelligent as yourself will learn them almost effortlessly."

Ax=b

- **A is square, symmetric and positive-definite**
- **When the A is dense, you're stuck, use backsubstitution**
- **When A is sparse, iterative techniques (such as Conjugate Gradient) are faster and more memory efficient**
- **Simple example:**

$$\begin{bmatrix} 3 & 2 \\ 2 & 6 \end{bmatrix} x = \begin{bmatrix} 2 \\ -8 \end{bmatrix}$$

(Yeah yeah, it's not sparse)

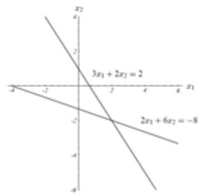

Turn Ax=b into a minimization problem

- **Minimization is more logical to analyze iteration (gradient ascent/descent)**
- **Quadratic form** $f(x) = \frac{1}{2}x^T A x - b^T x + c$
  - c can be ignored because we want to minimize
- **Intuition:**
  - the solution of a linear system is always the intersection of n hyperplanes
  - Take the square distance to them
  - A needs to be positive-definite so that we have a nice parabola

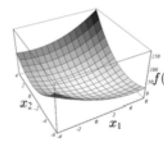

Graph of quadratic form $f(x) = \frac{1}{2}x^T A x - b^T x + c$. The minimum point of this surface is the solution to $Ax = b$.

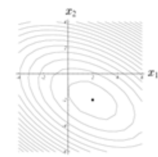

Contours of the quadratic form. Each ellipsoidal curve has constant $f(x)$.

Gradient of the quadratic form

$$f'(x) = \begin{bmatrix} \frac{\partial}{\partial x_1} f(x) \\ \frac{\partial}{\partial x_2} f(x) \\ \vdots \\ \frac{\partial}{\partial x_n} f(x) \end{bmatrix}$$

- Not our image gradient!
- Multidimensional gradient (as many dim as rows in matrix)

since $f(x) = \frac{1}{2}x^T A x - b^T x + c$

$$f'(x) = \frac{1}{2}A^T x + \frac{1}{2}Ax - b$$

And since A is symmetric

$$f'(x) = Ax - b$$

Not surprising: we turned Ax=b into the quadratic minimization

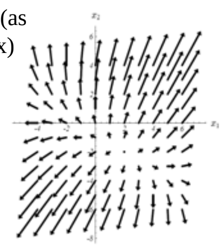

Gradient $f'(x)$ of the quadratic form. For every x, the gradient points in the direction of steepest increase of $f(x)$, and is orthogonal to the contour lines.

(if A is not symmetric, conjuagte gradient finds solution for $\frac{1}{2}(A^T + A)x = b$.

Steepest descent/ascent

- **Pick gradient direction**
- **Find optimum in this direction**

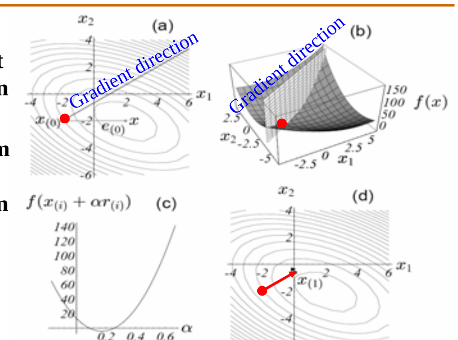

Energy along the gradient

The method of Steepest Descent.

Residual

- **At iteration i, we are at a point x(i)**
- **Residual r(i)=b-Ax(i)**
- **Cool property of quadratic form: residual = - gradient**

Behavior of gradient descent

- **Zigzag or goes straight depending if we're lucky**
  - Ends up doing multiple steps in the same direction

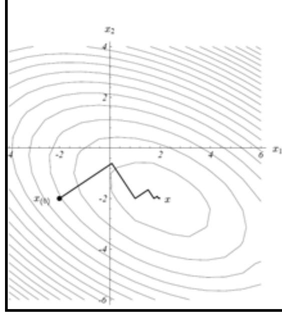

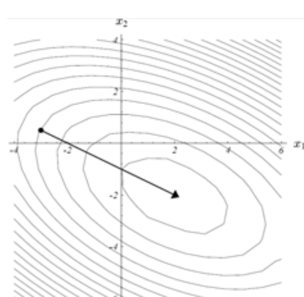

Conjugate gradient

- **Smarter choice of direction**
  - Ideally, step directions should be orthogonal to one another (no redundancy)
  - But tough to achieve
  - Next best thing: make them A-orthogonal (conjugate) That is, orthogonal when transformed by A: $d_{(i)}^T A d_{(j)} = 0$

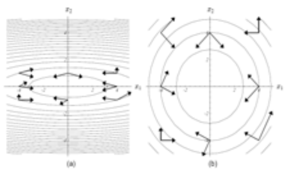

Figure 22: These pairs of vectors are A-orthogonal … because these pairs of vectors are orthogonal.

Conjugate gradient

- **For each step:**
  - Take the residual (gradient)
  - Make it A-orthogonal to the previous ones
  - Find minimum along this direction
- **Plus life is good:**
  - In practice, you only need the previous one
  - You can show that the new residual r(i+1) is already A-orthogonal to all previous directions p but p(i)

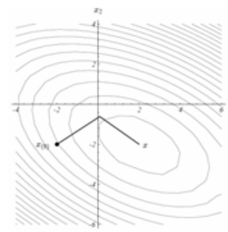

Figure 30: The method of Conjugate Gradients.

Recap

- **Poisson image cloning: paste gradient, enforce boundary condition**
- **Variational formulation** $\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2$ with $f|_{\partial\Omega} = f^*|_{\partial\Omega}$,
- **Also Euler-Lagrange formulation** $\Delta f = \operatorname{div}\mathbf{v}$ over Ω, with $f|_{\partial\Omega} = f^*|_{\partial\Omega}$
- **Discretize variational version, leads to big but sparse linear system**
- **Conjugate gradient is a smart iterative technique to solve it**

Questions?

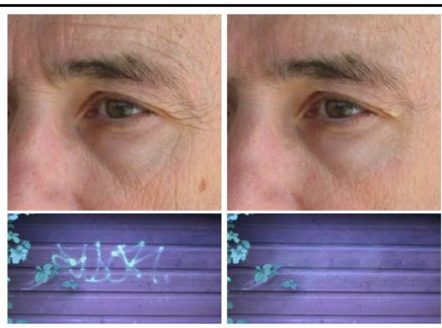

Figure 2: **Concealment**. By importing seamlessly a piece of the background, complete objects, parts of objects, and undesirable artifacts can easily be hidden. In both examples, multiple strokes (not shown) were used.

Manipulate the gradient

- **Mix gradients of g & f: take the max**

Figure 8: **Inserting one object close to another**. With seamless cloning, an object in the destination image touching the selected region Ω bleeds into it. Bleeding is inhibited by using mixed gradients as the guidance field.

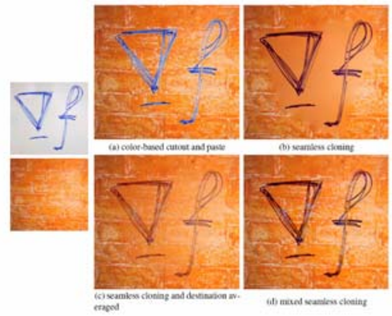

Figure 6: **Inserting objects with holes**. (a) The classic method, color-based selection and alpha masking might be time consuming and often leaves an undesirable halo; (b-c) seamless cloning, even averaged with the original image, is not effective; (d) mixed seamless cloning based on a loose selection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

Figure 7: **Inserting transparent objects**. Mixed seamless cloning facilitates the transfer of partly transparent objects, such as the rainbow in this example. The non-linear mixing of gradient fields picks out whichever of source or destination structure is the more salient at each location.

Reduce big gradients

- **Dynamic range compression**
- **See Fattal et al. 2002**

Figure 10: **Local illumination changes**. Applying an appropriate non-linear transformation to the gradient field inside the selection and then integrating back with a Poisson solver, modifies locally the apparent illumination of an image. This is useful to highlight under-exposed foreground objects or to reduce specular reflections.

Questions?

Fourier interpretation

- **Least square on gradient** $\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2$ with $f|_{\partial\Omega} = f^*|_{\partial\Omega}$
- **Parseval anybody?**
  - Integral of squared stuff is the same in Fourier and primal
- **What is the gradient/derivative in Fourier?**
  - Multiply coefficients by frequency
- **Seen in Fourier, Poisson editing does a weighted least square of the image where low frequencies have a small weight and high frequencies a big weight**

Issues with Poisson cloning

- **Colors**
- **Contrast**
- **The backgrounds in f & g should be similar**

Improvement: local contrast

- **Use the log**
- **Or use covariant derivatives (next slides)**

Covariant derivatives & Photoshop

- **Photoshop Healing brush**
- **Developed independently from Poisson editing by Todor Georgiev (Adobe)**

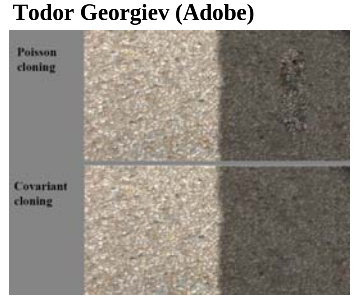

From Todor Georgiev's slides http://photo.csail.mit.edu/posters/todor_slides.pdf

Seamless Image Stitching in the Gradient Domain

- Anat Levin, Assaf Zomet**,** Shmuel Peleg, and Yair Weiss
  http://www.cs.huji.ac.il/~alevin/papers/eccv04-blending.pdf
  http://eprints.pascal-network.org/archive/00001062/01/tips05-blending.pdf
- **Various strategies (optimal cut, feathering)**

Fig. 1. Image stitching. On the left are the input images. ω is the overlap region. On top right is a simple pasting of the input images. On the bottom right is the result of the GIST1 algorithm.

Photomontage

- **http://grail.cs.washington.edu/projects/photomontage/photomontage.pdf**

Figure 6 We use a set of portraits (first row) to mix and match facial features, to either improve a portrait, or create entirely new people. The faces are first hand-aligned, for example, to place all the noses in the same location. In the first two images in the second row, we replace the closed eyes of a portrait with the open eyes of another. The user paints strokes with the *designated source* objective to specify desired features. Next, we create a fictional person by combining three source portraits. Gradient-domain fusion is used to smooth out skin tone differences. Finally, we show two additional mixed portraits.

Elder's edge representation

- http://elderlab.yorku.ca/~elder/publications/journals/ElderPAMI01.pdf

Gradient tone mapping

- **Fattal et al. Siggraph 2002**

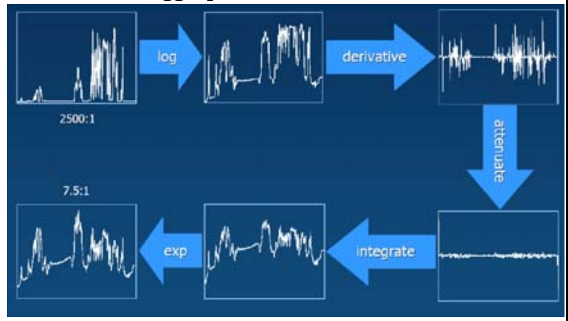

Slide from Siggraph 2005 by Raskar (Graphs by Fattal et al.)

Gradient attenuation

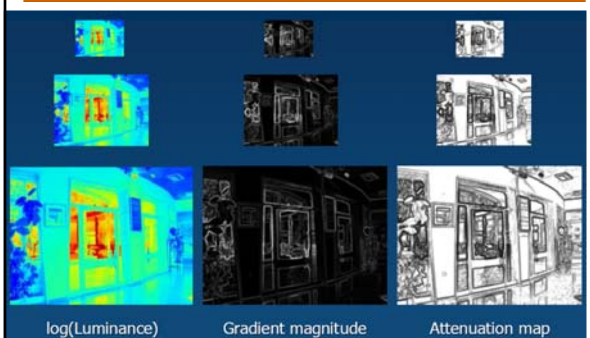

From Fattal et al.

Fattal et al. Gradient tone mapping

Gradient tone mapping

- **Socolinsky, D. *Dynamic Range Constraints in Image Fusion and Visualization*, in Proceedings of Signal and Image Processing 2000, Las Vegas, November 2000.**

(a) (b) (c) (d) (e)

Fig. 1. (a) Mediastinal window of thoracic CT scan. (b) Lung window of thoracic CT scan. (c) Clipped solution of equation (2) for the fusion of (a) and (b). (d) Linearly scaled solution of (2) for the fusion of (a) and (b). (e) Solution of equation (6) for the fusion of(a) and (b).

Gradient tone mapping

- **Socolinsky, D. *Dynamic Range Constraints in Image Fusion and Visualization*, in Proceedings of Signal and Image Processing 2000.**

Fig. 4. Left: average of images in figure 2. Middle: rendering of the sum of the images in figure 2 through adaptive histogram compression. Right: fusion of images in figure 2 using the obstacle method.

- **Socolinsky, D. and Wolff, L.B., *A new paradigm for multispectral image visualization and data fusion*, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Fort Collins, June 1999.**

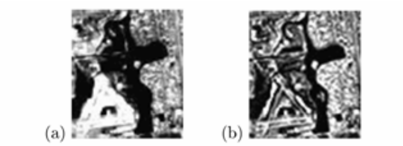

Figure 4: (a) Grayscale version of 9-band image computed through PCA. (b) Grayscale version of the same image computed through our algorithm.

Retinex

- **Land, Land and McCann (inventor/founder of polaroid)**
- **Theory of lightness perception (albedo vs. illumination)**
- **Strong gradients come from albedo, illumination is smooth**

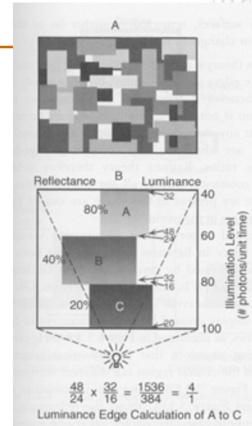

Questions?

Color2gray

- **Use Lab gradient to create grayscale images**

Color2Gray: Salience-Preserving Color Removal

Amy A. Gooch Sven C. Olsen Jack Tumblin Bruce Gooch

Northwestern University *

Figure 1: A color image (Left) often reveals important visual details missing from a luminance-only image (Middle). Our Color2Gray algorithm (Right) maps visible color changes to grayscale changes. *Image: Impressionist Sunrise by Claude Monet, courtesy of Artcyclopedia.com.*

Poisson Matting

- **Sun et al. Siggraph 2004**
- **Assume gradient of F & B is negligible**
- **Plus various image-editing tools to refine matte**

$$I = \alpha F + (1-\alpha)B$$

$$\nabla I = (F-B)\nabla\alpha + \alpha\nabla F + (1-\alpha)\nabla B$$

$$\nabla\alpha \approx \frac{1}{F-B}\nabla I$$

Figure 1: Pulling of matte from a complex scene. From left to right: a complex natural image for existing matting techniques where the color background is complex, a high quality matte generated by Poisson matting, a composite image with the extracted koala and a constant-color background, and a composite image with the extracted koala and a different background.

Gradient camera?

- **Tumblin et al. CVPR 2005** **http://www.cfar.umd.edu/~aagrawal/gradcam/gradcam.html**

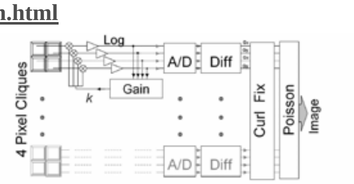

Figure 2. Log-gradient camera overview: intensity sensors organized into 4-pixel cliques share the same self-adjusting gain setting k, and send $\log(I_d)$ signals to A/D converter. Subtraction removes common-mode noise, and a linear 'curl fix' solver corrects saturated gradient values or 'dead' pixels, and a Poisson solver finds output values from gradients.

Poisson-ish mesh editing

- http://portal.acm.org/citation.cfm?id=1057432.1057456
- http://www.cad.zju.edu.cn/home/xudong/Projects/mesh_editing/main.htm
- http://people.csail.mit.edu/sumner/research/deftransfer/

Figure 1: An unknown mythical creature. Left: mesh components for merging and deformation (the arm). Right: final editing result.

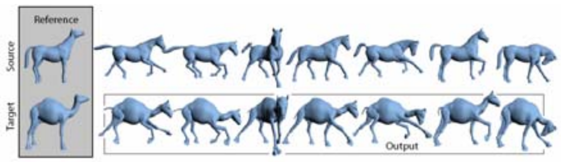

Figure 1: Deformation transfer copies the deformations exhibited by a source mesh onto a different target mesh. In this example, deformations of the reference horse mesh are transfered to the reference camel, generating seven new camel poses. Both gross skeletal changes as well as more subtle skin deformations are successfully reproduced.

Questions?

Alternative to membrane

- Thin plate:
  minimize *second* derivative

$$\min_f \int\int f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2 \mathrm{dxdy}$$

Data

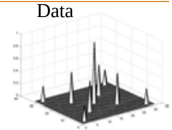

Membrane interpolation

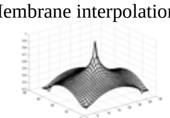

Thin-plate interpolation

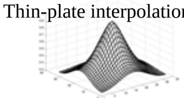

Inpainting

- More elaborate energy functional/PDEs
- http://www-mount.ee.umn.edu/~guille/inpainting.htm

Key references

- Socolinsky, D. *Dynamic Range Constraints in Image Fusion and Visualization* 2000. http://www.equinoxsensors.com/news.html
- Elder, Image editing in the contour domain, 2001 http://elderlab.yorku.ca/~elder/publications/journals/ElderPAMI01.pdf
- Fattal et al. 2002 Gradient Domain HDR Compression http://www.cs.huji.ac.il/%7Edanix/hdr/
- Poisson Image Editing Perez et al. http://research.microsoft.com/vision/cambridge/papers/perez_siggraph03.pdf
- Covariant Derivatives and Vision, Todor Georgiev (Adobe Systems) ECCV 2006

Poisson, Laplace, Lagrange, Fourier, Monge, Parseval

- Fourier studied under Lagrange, Laplace & Monge, and Legendre & Poisson were around
- They all raised serious objections about Fourier's work on Trigomometric series
- http://www.ece.umd.edu/~taylor/frame2.htm
- http://www.mathphysics.com/pde/history.html
- http://www-groups.dcs.st-and.ac.uk/~history/Mathematicians/Fourier.html
- http://www.memagazine.org/contents/current/webonly/wex80905.html
- http://www.shsu.edu/~icc_cmf/bio/fourier.html
- http://en.wikipedia.org/wiki/Simeon_Poisson
- http://en.wikipedia.org/wiki/Pierre-Simon_Laplace
- http://en.wikipedia.org/wiki/Jean_Baptiste_Joseph_Fourier
- http://www-groups.dcs.st-and.ac.uk/~history/Mathematicians/Parseval.html

Refs Laplace and Poisson

- http://www.ifm.liu.se/~boser/elma/Lect4.pdf
- http://farside.ph.utexas.edu/teaching/329/lectures/node74.html
- http://en.wikipedia.org/wiki/Poisson's_equation
- http://www.colorado.edu/engineering/CAS/courses.d/AFEM.d/AFEM.Ch03.d/AFEM.Ch03.pdf

Gradient image editing refs

- http://research.microsoft.com/vision/cambridge/papers/perez_siggraph03.pdf
- http://www.cs.huji.ac.il/~alevin/papers/eccv04-blending.pdf
- http://www.eg.org/EG/DL/WS/COMPAESTH/COMPAESTH05/075-081.pdf.abstract.pdf
- http://photo.csail.mit.edu/posters/Georgiev_Covariant.pdf
- **Covariant Derivatives and Vision, Todor Georgiev (Adobe Systems) ECCV 2006**
- http://www.mpi-sb.mpg.de/~hitoshi/research/image_restoration/index.shtml
- http://www.cs.tau.ac.il/~tommer/vidoegrad/
- http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=1467600
- http://grail.cs.washington.edu/projects/photomontage/
- http://www.cfar.umd.edu/~aagrawal/iccv05/surface_reconstruction.html
- http://www.merl.com/people/raskar/Flash05/
- http://research.microsoft.com/~carrot/new_page_1.htm
- http://www.idiom.com/~zilla/Work/scatteredInterpolation.pdf

PSet 3: write a review (6.882 only)

- **Choose a paper from the list**
  - Or suggest another paper
- **Write a review using the SIggraph form**

Peer review system (Siggraph biased)

- **Peer reviews, committees**
  - A paper chair forms a committee (~40 people)
  - Each paper is assigned to 2 committee members: a primary & a secondary
  - Each committee member assigns it to 1 or 2 external (a.k.a. tertiaries)
  - The committee meets and decides who gets accepted
- **Double blind process**
  - The authors don't know who reviews them
  - The tertiaries don't know who they review
  - In some fields, even the committee members don't know who they review.
  - Guessing who reviewed you?
    - A very bad idea. Too often wrong!

Other systems

- **Journals:**
  - No deadline, no committee meeting
  - Review cycle: reviewers critique, authors improve, until convergence
- **Non-blind system**
  - Some think that reviewer anonymity is bad:
    - Reviewers might not feel the need to do a good job since they're not cited
    - Competitors could slow down a paper to buy time

What to write in a review

- **Help committee with decision, assess work**
  - The score helps, but a concise discussion of the pros and cons, comparison to previous work is more important
- **Give feedback to authors, help them improve their work**
  - Technical points
  - Writing (most important)
  - As a reviewer, always a difficult balance between effort spent and doing a good job (sometimes you feel you should become a co-author for your contribution)

Reviewing

- **Ethical issues**
  - What if I work on the same subject?
  - Confidentiality
  - Conflicts
    - Advisor: lifetime conflict
    - Co-author (~ 2 to 3 years)
    - Co-principal investigator on a grant
    - Family
    - Same institution or could be perceived as same institution (e.g.CSAIL and Medialab, MSR Redmond and MSR Asia)
    - Anything that

Siggraph review form

- 1) Briefly describe the paper and its contribution to computer graphics and interactive techniques. Please give your assessment of the scope and magnitude of the paper's contribution.
- 2) Is the exposition clear? How could it be improved?
- 3) Are the references adequate? List any references that are needed.
- 4) Could the work be reproduced by one or more skilled graduate students? Are all important algorithmic or system details discussed adequately? Are the limitations and drawbacks of the work clear?
- 5) Please rate this paper on a continuous scale from 1 to 5, where: 1 = Reject, 2 = Doubtful, 3 = Possibly accept, 4 = Probably accept, 5 = Accept.
- 6) Please rate your expertise in the subject area of the paper on a continuous scale from 1 to 3, where: 1=Tyro, 2=Journeyman, 3=Expert.
- 7) Explain your rating by discussing the strengths and weaknesses of the submission. Include suggestions for improvement and publication alternatives, if appropriate. Be thorough -- your explanation will be of highest importance for any committee discussion of the paper and will be used by the authors to improve their work. Be fair -- the authors spent a lot of effort to prepare their submission, and your evaluation will be forwarded to them during the rebuttal period.
- 8) List here any questions that you want answered by the author(s) during the rebuttal period.
- 9) You may enter private comments for the papers committee here. These comments will not be sent to the paper author(s).

Importance of good writing

- **What is the use of creating the best innovative ideas if nobody else can understand them?**

- **See Fredo's slides "How to write a bad paper"** **http://people.csail.mit.edu/fredo/FredoBadWriting.pdf**
  useful links **http://people.csail.mit.edu/fredo/student.html**
- **Bill's slides and links:** **http://www.ai.mit.edu/courses/6.899/doneClasses.html** **(April 10)**

Kajiya on conference reviewing

"The reviewing process for SIGGRAPH is far from perfect, although most everyone is giving it their best effort.

The very nature of the process is such that many reviewers will not be able to spend nearly enough time weighing the nuances of your paper. This is something for which you must compensate in order to be successful."

Links

- **How to Get Your SIGGRAPH Paper Rejected, Jim Kajiya, SIGGRAPH 1993 Papers Chair, (link)**
- **Ted Adelson's Informal guidelines for writing a paper, 1991. (link)**
- **Notes on technical writing, Don Knuth, 1989. (pdf)**
- **What's wrong with these equations, David Mermin, Physics Today, Oct., 1989. (pdf)**
- **Ten Simple Rules for Mathematical Writing, Dimitri P. Bertsekas (link)**
- **Advice on Research and Writing (at CMU)**
- **How (and How Not) to Write a Good Systems Paper by Roy Levin and David D. Redell**
- **Things I Hope Not to See or Hear at SIGGRAPH by Jim Blinn**
- **How to have your abstract rejected**

Next time: how to take great pictures

Photos Steve McCurry