Hilbert and Fourier analysis
C10

M. Delbracio & G. Facciolo

Today's topics

- Distributions Theory
  - Summary from previous lecture
- Poisson Editor
  - Introduction
  - Numerical Solution
  - Seamless Cloning
  - Mixing Gradients
  - Selective Edition

The Variational Interpretation

Proposition 1: Let $\Omega = [0, 2\pi]^2$ and $f \in L^2_{\text{per}}(\mathbb{R}^2)$. Then there exists a unique function $u \in H^1_{\text{per}}(\mathbb{R}^2)$ minimizer of the energy

$$E(u) = \frac{1}{2} \int_\Omega |Du|^2(x, y)dxdy - \int_\Omega f(x, y)u(x, y)$$

and satisfying $u = 0$ in $\partial\Omega$.

This solution is the restriction to Ω of an odd function belonging to $H^2_{\text{per}}([-2\pi, 2\pi]^2)$ (4π-periodic functions).

This solution is the same as the one from the Poisson equation with Dirichlet conditions:

$$-\Delta u = f, \quad u = 0 \ \text{ in } \ \partial\Omega$$

Poisson Editor

Proposition 2: Let $\Omega = [0, 2\pi]^2$ and $V = (v_1, v_2) \in (L^2_{\text{per}}(\mathbb{R}^2))^2$ a vector field. Then there exists a function $u \in H^1_{\text{per}}(\mathbb{R}^2)$ minimizer of the energy

$$E(u) = \frac{1}{2}\int_\Omega |Du - V|^2(x,y)dxdy = \int_\Omega \left((u_x - v_1)^2 + (u_y - v_2)^2\right) dxdy$$

and satisfying $\frac{\partial u}{\partial n} = 0$ in $\partial\Omega$.

This solution is the restriction to Ω of an even (paire) function belonging to $H^1_{\text{per}}([-2\pi, 2\pi]^2)$ (4π-periodic functions).

This solution is the same as the one from the Poisson equation with Neumann conditions:

$$\Delta u = \text{div}(V) = \frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y}, \quad \frac{\partial u}{\partial n} = 0 \text{ in } \partial\Omega$$

Proof In order to naturally have the boundary condition $\frac{\partial u}{\partial n} = 0$ it is convenient to use the cosine basis decomposition for elements in $L^2([0, 2\pi]^2)$

$$\left(\cos\left(\tfrac{kx}{2}\right)\cos\left(\tfrac{ly}{2}\right)\right), \text{ with } k, l \in \mathbb{N}$$

Hence, we write

$$u(x,y) = \textstyle\sum_{k,l\in\mathbb{N}} c_{k,l}(u)\cos\left(\tfrac{kx}{2}\right)\cos\left(\tfrac{ly}{2}\right)$$

$$v_1(x,y) = \textstyle\sum_{k,l\in\mathbb{N}} c_{k,l}(v_1)\cos\left(\tfrac{kx}{2}\right)\cos\left(\tfrac{ly}{2}\right)$$

$$v_2(x,y) = \textstyle\sum_{k,l\in\mathbb{N}} c_{k,l}(v_2)\cos\left(\tfrac{kx}{2}\right)\cos\left(\tfrac{ly}{2}\right)$$

By using the respective Fourier expansions for u_x, u_y, (derivatives in the distribution sense) and v_1, v_2 the Poisson energy can be re-written as

$$E(u) = \frac{1}{2}\sum_{k,l\in\mathbb{N}}\left(\frac{k}{2}c_{k,l}(u) - c_{k,l}(v_1)\right)^2 + \left(\frac{l}{2}c_{k,l}(u) - c_{k,l}(v_2)\right)^2$$

Proof (cont. 1) Then, regrouping the energy $E(u)$:

$$E(u) = \frac{1}{2} \sum_{k,l \in \mathbb{N}^*} \frac{k^2+l^2}{4} c_{k,l}(u)^2 - \Big(k c_{k,l}(v_1) + l c_{k,l}(v_2)\Big) c_{k,l}(u) + c_{k,l}(v_1)^2 + c_{k,l}(v_2)^2$$

Each term depends on a different $c_{k,l}(u)$. We need to minimize

$$\left(\frac{k^2+l^2}{4}\right) c_{k,l}(u)^2 - \Big(k c_{k,l}(v_1) + l c_{k,l}(v_2)\Big) c_{k,l}(u)$$

so if we differentiate and $= 0$, we get

$$c_{k,l}(u) = 2\frac{k c_{k,l}(v_1) + l c_{k,l}(v_2)}{k^2+l^2}, \quad \forall k, l > 0$$

- u is in $H^1_{\text{per}}([-2\pi, 2\pi]^2)$ since $\frac{k}{2} c_{k,l}(u) \in l^2(\mathbb{Z}^2)$ (so $u_x \in L^2_{\text{per}}$).
- The solution is unique up to a constant (the term $c_{0,0}(u)$ is not set). If u is solution then $u + C$ for C constant is also solution. Thus,

$$u(x,y) = C + \sum_{k,l \in \mathbb{N}^*} 2\frac{k c_{k,l}(v_1) + l c_{k,l}(v_2)}{k^2+l^2} \cos\left(\frac{kx}{2}\right) \cos\left(\frac{ly}{2}\right)$$

- Same solution as the Poisson equation with the Neumann condition for

$$f = -\text{div}(V) = -(v_1)_x - (v_2)_y$$

□

Proposition 3: Let u be the solution of the Poisson editor for a gradient field $V = (v_1, v_2)$. Then u satisfies the equation

$$\Delta u = \mathrm{div}(V)$$

in the distribution sense.

Proof Consider a perturbation $t\varphi$, with $\varphi \in C_c^\infty(\Omega)$ and $t \in \mathbb{R}$. If u is a minimum then

$$E(u + t\varphi) = \int_\Omega (D(u + t\varphi) - V)^2 \geq \int_\Omega (Du - V)^2 = E(u)$$

Then,

$$\int_\Omega (Du - V)^2 + 2tD\varphi \cdot (Du - V) + t^2 Du \cdot Du \geq \int_\Omega (Du - V)^2$$

Since u is a minimum, the derivative in $t = 0$ should be zero

$$\int_\Omega D\varphi \cdot (Du - V) = 0$$

This is the same as

$$\int_\Omega \varphi_x(u_x - v_1) + \varphi_y(u_y - v_2) = 0$$

Since u_x, u_y, v_1 and v_2 are in L^2_{loc} they are also distributions:

$$< u_x - v_1, \varphi_x > + < u_y - v_2, \varphi_y >= 0$$

Next, from the definition of distribution derivative

$$- < u_{xx} - (v_1)_x, \varphi > - < u_{yy} - (v_2)_y, \varphi >= - < u_{xx} + u_{yy} - (v_1)_x - (v_2)_y, \varphi >= 0$$

and thus

$$\Delta u - \text{div}(V) = 0.$$

□

Periodic + Smooth decomposition

Proposition 4: Let u be a function such that $\partial^{\mathbf{i}} u \in L^2([0,2\pi]^2)$ for $|\mathbf{i}| \leq 2$. Then, there exists a function $v \in H^2_{\text{per}}([0,2\pi]^2)$ such that

$$\Delta v = \Delta u \quad \text{in } [0,2\pi]^2$$

This function is unique up to an additive constant. The difference $w = v - u$ verifies $\Delta w = 0$ and hence it is smooth in $[0,2\pi]^2$.

Periodic + Smooth decomposition

The P+S decomposition permits to visualize the Fourier spectrum getting rid of the boundary effects (e.g. discontinuities when periodizing). The periodic decomposition inherits all the image details since its laplacian is the same.

$$u \quad = \quad p \quad + \quad s$$

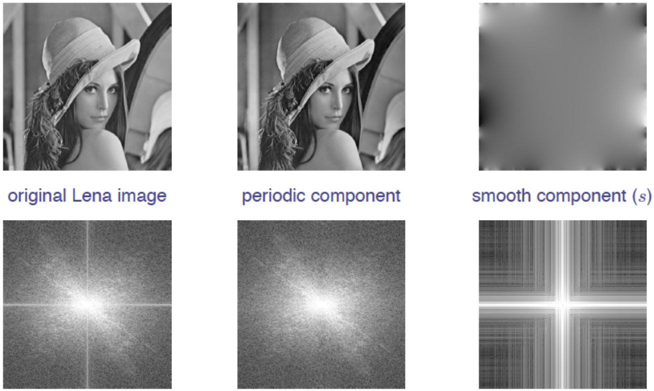

corresponding Discrete Fourier Transforms (log-modulus)

L. Moisan, *Periodic plus Smooth Image Decomposition*, Journal of Math. Imag. and Vision, 2011.

Theorem : *Unicity of the Fourier coefficients*

If u is a distribution with the Fourier expansion

$$u = \sum_{\mathbf{m}\in\mathbb{Z}^2} c_{\mathbf{m}} \mathrm{e}^{i\mathbf{m}\cdot\mathbf{x}} = 0$$

then $\forall \mathbf{m}$, $c_{\mathbf{m}} = 0$.

End of last lecture summary.

Image Edition

- **Goal:** Manipulate locally a digital image to change: color, replace detail, change some region, ..., in an imperceptible way.
- Our vision perceives the Laplacian of images more than the images themselves.
- A function in a compact domain can be reconstructed from its laplacian and the values in the boundary.

sources destinations cloning seamless cloning

Poisson Image Editing

Patrick Pérez* Michel Gangnet† Andrew Blake‡

Microsoft Research UK

Abstract

Using generic interpolation machinery based on solving Poisson equations, a variety of novel tools are introduced for seamless editing of image regions. The first set of tools permits the seamless importation of both opaque and transparent source image regions into a destination region. The second set is based on similar mathematical ideas and allows the user to modify the appearance of the image seamlessly, within a selected region. These changes can be arranged to affect the texture, the illumination, and the color of objects lying in the region, or to make tileable a rectangular selection.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques; I.4.3 [Image Processing and Computer Vision]: Enhancement—Filtering;

Keywords: Interactive image editing, image gradient, guided interpolation, Poisson equation, seamless cloning, selection editing

able effect. Conversely, the second-order variations extracted by the Laplacian operator are the most significant perceptually.

Secondly, a scalar function on a bounded domain is uniquely defined by its values on the boundary and its Laplacian in the interior. The Poisson equation therefore has a unique solution and this leads to a sound algorithm.

So, given methods for crafting the Laplacian of an unknown function over some domain, and its boundary conditions, the Poisson equation can be solved numerically to achieve seamless filling of that domain. This can be replicated independently in each of the channels of a color image. Solving the Poisson equation also has an alternative interpretation as a minimization problem: it computes the function whose gradient is the closest, in the L_2-norm, to some prescribed vector field — the *guidance* vector field — under given boundary conditions. In that way, the reconstructed function interpolates the boundary conditions inwards, while following the spatial variations of the guidance field as closely as possible. Section 2 details this guided interpolation.

[Pérez et al. 2003] P. Pérez, M. Gangnet and A. Blake, "Poisson Image Editing". *ACM Transactions on Graphics*. Proc. of ACM SIGGRAPH 2003, vol 22, Issue 3 Pages: 313 - 318 (July 2003).

Image Completion: *Simple Interpolation*

Suppose we want to interpolate a function $g(x)$ that we only know it in the contour $\partial\Omega$.

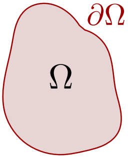

Proposed approach:

Minimize the following energy:

$$E(u) = \int_{\Omega} |Du|^2 d\mathbf{x} \quad \text{with} \quad u(\mathbf{x}) = g(\mathbf{x}) \quad \text{in } \partial\Omega.$$

This solution is the same as the one from the Laplace equation with Dirichlet boundary conditions:

$$\begin{cases} \Delta u = 0 & \text{if} \quad \mathbf{x} \in \Omega \\ u(\mathbf{x}) = g(\mathbf{x}) & \text{if} \quad \mathbf{x} \in \partial\Omega. \end{cases}$$

This solution is over blurred (too smooth since $\Delta u = 0$).

Image Completion: *Guided Interpolation*

- We want to interpolate a function $g(x)$ only known it in the contour $\partial\Omega$.
- We are given a *guidance field* $\mathbf{v} = (v_1, v_2)$.
- **Goal:** find a function u whose gradient is as similar as possible to $\mathbf{v}$ and compatible with the boundary condition.

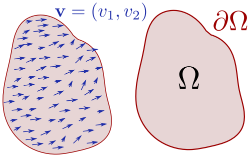

Minimize the following energy:

$$E(u) = \int_\Omega |Du - \mathbf{v}|^2 d\mathbf{x} \quad \text{with} \quad u(\mathbf{x}) = g(\mathbf{x}) \quad \text{in } \partial\Omega.$$

This solution is the same as the one from the Poisson equation with Dirichlet boundary conditions:

$$\begin{cases} \Delta u = \operatorname{div}(\mathbf{v}) & \text{if} \quad \mathbf{x} \in \Omega \\ u(\mathbf{x}) = g(\mathbf{x}) & \text{if} \quad \mathbf{x} \in \partial\Omega. \end{cases}$$

where we have called $\operatorname{div}(\mathbf{v}) = \partial {v_1}_x + \partial {v_2}_y$.

This equation is the cornerstore of the framework proposed in [Pérez et al. 2003].

Image Completion: *Guided Interpolation*

A simple case

- The *guidance field* is itself the gradient of a function w,

$$\mathbf{v} = Dw.$$

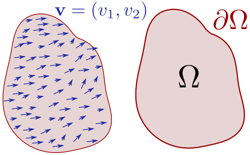

Then $\Delta u = \mathrm{div}(\mathbf{v}) = \mathrm{div}(Dw) = \Delta w$, and if we set $z = w - u$, we can rewrite the problem as

$$\begin{cases} \Delta z = 0 \quad \text{if} \quad \mathbf{x} \in \Omega \\ z(\mathbf{x}) = (g - w)(\mathbf{x}) \quad \text{if} \quad \mathbf{x} \in \partial\Omega. \end{cases}$$

Numerical Solution: Discrete Poisson Solver

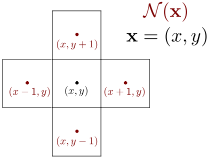

- Images are defined in a discrete grid Ω_d.
- We note $\partial\Omega_d$ the set of boundary pixels.
- For each pixel $\mathbf{x} = (x, y)$ we define $\mathcal{N}(\mathbf{x}) = \{(x-1, y), (x+1, y), (x, y-1), (x, y+1)\}$.
- A point $\mathbf{x} \in \partial\Omega_d$ iff $\mathbf{x} \notin \Omega_d$ and $\mathcal{N}(\mathbf{x}) \cap \Omega_d \neq \emptyset$.

  - If the domain Ω_d is not rectangular the Fourier technique doesn't work.
  - Directly solve the *discrete* variational (finite difference discretization)

Continuous energy:

$$\int_\Omega |Du - \mathbf{v}|^2 d\mathbf{x} \quad \text{with} \quad u(\mathbf{x}) = g(\mathbf{x}) \quad \text{in } \partial\Omega.$$

Discrete energy:

$$\sum_{\mathbf{x}\in\Omega_d} \sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})} \left(u(\mathbf{x}) - u(\mathbf{y}) - v_{i(\mathbf{x},\mathbf{y})}\left(\tfrac{\mathbf{x}+\mathbf{y}}{2}\right)\right)^2 \quad \text{with} \quad u(\mathbf{x}) = g(\mathbf{x}) \quad \text{in } \partial\Omega_d.$$

where $i(\mathbf{x}, \mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} = \mathbf{x} \pm (1,0), \\ 2 & \text{if } \mathbf{y} = \mathbf{x} \pm (0,1). \end{cases}$

Numerical Solution: Discrete Poisson Solver

Discrete energy:

$$\sum_{\mathbf{x}\in\Omega_d}\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})}\left(u(\mathbf{x})-u(\mathbf{y})-v_{i(\mathbf{x},\mathbf{y})}(\tfrac{\mathbf{x}+\mathbf{y}}{2})\right)^2 \quad \text{with} \quad u(\mathbf{x})=g(\mathbf{x}) \quad \text{in } \partial\Omega_d.$$

where $i(\mathbf{x},\mathbf{y})=\begin{cases}1 & \text{if } \mathbf{y}=\mathbf{x}\pm(1,0),\\ 2 & \text{if } \mathbf{y}=\mathbf{x}\pm(0,1).\end{cases}$ Its a quadratic optimization problem. Its solution satisfies the following simultaneous linear equations

$$\forall\mathbf{x}\in\Omega_d,\quad |\mathcal{N}(\mathbf{x})|u(\mathbf{x})-\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})\cap\Omega_d}u(\mathbf{y})=\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})\cap\partial\Omega_d}g(\mathbf{y})+\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})}\nu(\mathbf{x},\mathbf{y})$$

where $\nu(\mathbf{x},\mathbf{y})=v_{i(x,y)}(\frac{\mathbf{x}+\mathbf{y}}{2})$ and $|\mathcal{N}(\mathbf{x})|$ is the cardinal of $\mathcal{N}(\mathbf{x})$ (1 to 4) supposing that Ω_d is discrete connected (every point has at least one neighbor).

For the point in the interior of Ω_d there aren't boundary terms, then

$$\forall\mathbf{x}\in\text{interior}(\Omega_d),\quad |\mathcal{N}(\mathbf{x})|u(\mathbf{x})-\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})\cap\Omega_d}u(\mathbf{y})=\sum_{\mathbf{y}\in\mathcal{N}(\mathbf{x})}\nu(\mathbf{x},\mathbf{y})$$

This is a big linear system, with symmetric positive define *sparse* matrix ($|\Omega_d|$ unknown values). It can be solved by Gauss-Seidel iterative method.

Seamless Cloning

- Insert a part of image v_0 in image u_0.
- Copy and paste doesn't work.
- Guidance field $\mathbf{v}$ taken directly from the source image.

We set $\mathbf{v} = Dv_0$, then $\text{div}(\mathbf{v}) = \Delta v_0$. The problem becomes

$$\begin{cases} \Delta u(\mathbf{x}) = \Delta v_0(\mathbf{x}) & \text{if} \quad \mathbf{x} \in \Omega \\ u(\mathbf{x}) = u_0(\mathbf{x}) & \text{if} \quad \mathbf{x} \in \partial\Omega. \end{cases}$$

In the discrete world

$$\nu(\mathbf{x}, \mathbf{y}) = (v_0(\mathbf{x}) - v_0(\mathbf{y})).$$

Seamless Cloning

sources destinations cloning seamless cloning

Seamless Cloning

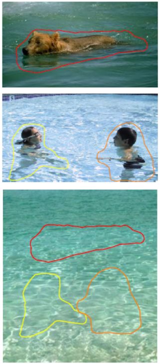
sources/destinations


cloning


seamless cloning

Seamless Cloning

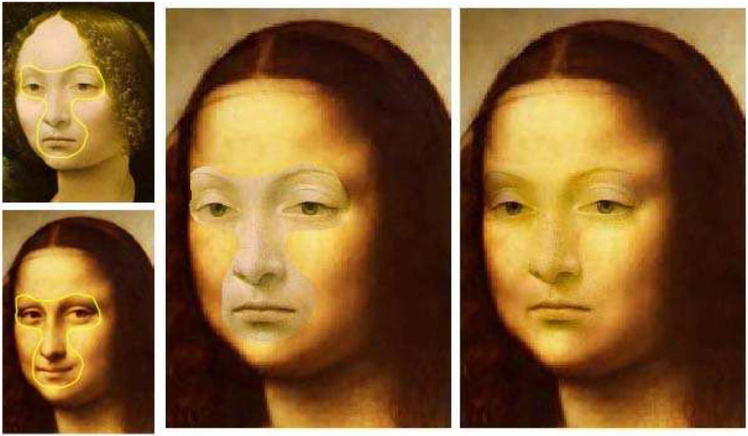

source/destination

cloning

seamless cloning

Mixing Gradients

- Insert a part of image v_0 in image u_0 but keep details of image u_0.
- See one image through the other.
- First idea: fill the domain Ω_d with $\frac{u_0+v_0}{2}$
- Problem: transitions in the boundary of Ω_d won't be continous.

Image details are coded in the parts with big gradient of the image.

We can define the guidance field $\mathbf{v}$ as

$$\mathbf{v}(\mathbf{x}) = \begin{cases} Dv_0(\mathbf{x}) & \text{if} \quad |Dv_0(\mathbf{x})| > |Du_0(\mathbf{x})|, \mathbf{x} \in \Omega \\ Du_0 & \text{otherwise.} \end{cases}$$

The discrete problem becomes

$$\nu(\mathbf{x}, \mathbf{y}) = \begin{cases} v_0(\mathbf{x}) - v_0(\mathbf{y}) & \text{if} \quad |v_0(\mathbf{x}) - v_0(\mathbf{y}))| > |u_0(\mathbf{x}) - u_0(\mathbf{y})|, \mathbf{x} \in \Omega_d \\ u_0(\mathbf{x}) - u_0(\mathbf{y}) & \text{otherwise.} \end{cases}$$

Mixing Gradients

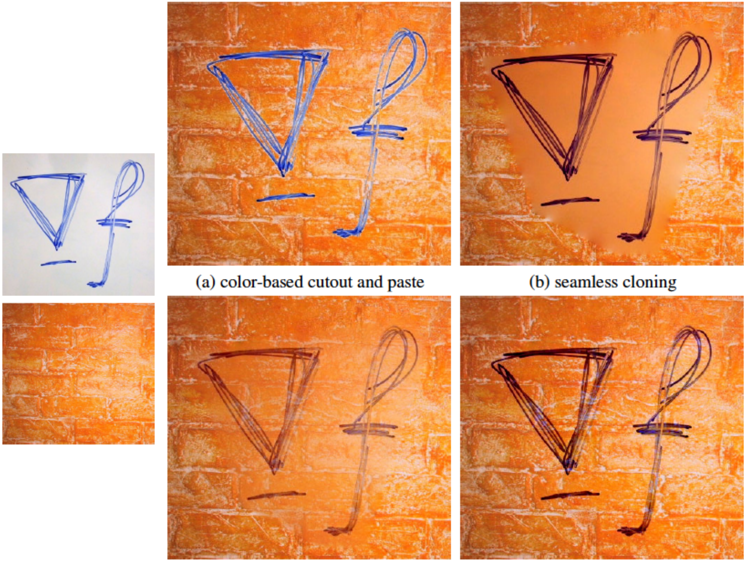

(a) color-based cutout and paste

(b) seamless cloning

(c) seamless cloning and destination averaged

(d) mixed seamless cloning

Mixing Gradients

source

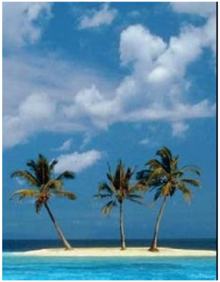

destination

Selective Edition

- Modify the image gradient field in order to keep only the values interested by the operator.

There are different ways of defining the guidance field **v** in a non linear way.

- Keep only the gradients bigger than certain threshold

$$\mathbf{v}(\mathbf{x}) = \begin{cases} Du_0(\mathbf{x}) & \text{if} \quad |Du_0(\mathbf{x})| > \delta, \mathbf{x} \in \Omega \\ 0 & \text{otherwise.} \end{cases}$$

- Enhance weak contrasts

$$\mathbf{v}(\mathbf{x}) = g(Du_0(\mathbf{x}))$$

with $g(v) = |v|^{-\alpha} v$ and $\alpha \in [0, 1]$.

Selective Edition

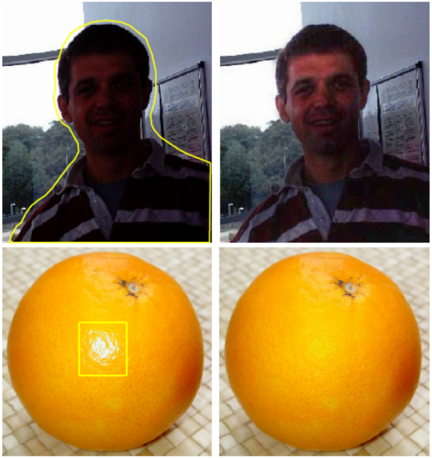

Figure 10: **Local illumination changes**. Applying an appropriate non-linear transformation to the gradient field inside the selection and then integrating back with a Poisson solver, modifies locally the apparent illumination of an image. This is useful to highlight under-exposed foreground objects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

Selective Edition

Figure 11: **Local color changes**. Left: original image showing selection Ω surrounding loosely an object of interest; center: background decolorization done by setting g to the original color image and f^* to the luminance of g; right: recoloring the object of interest by multiplying the RGB channels of the original image by 1.5, 0.5, and 0.5 respectively to form the source image.

Real-Time Gradient-Domain Painting

James McCann*
Carnegie Mellon University

Nancy S. Pollard†
Carnegie Mellon University

Figure 1: *Starting from a blank canvas, an artist draws and colors a portrait using gradient-domain brush strokes.* Additive *blend mode used for sketching,* directional *and* additive *for coloring. Artist time: 30 minutes.*

`http://graphics.cs.cmu.edu/projects/gradient-paint/`

[McCann and Pollard 2008] James McCann and Nancy S. Pollard, " Real-time Gradient-domain Painting". *ACM Transactions on Graphics (SIGGRAPH 2008)*. August 2008, Vol 27. No. 3.

the end.