Flash/No Flash and Dark Flash Photography

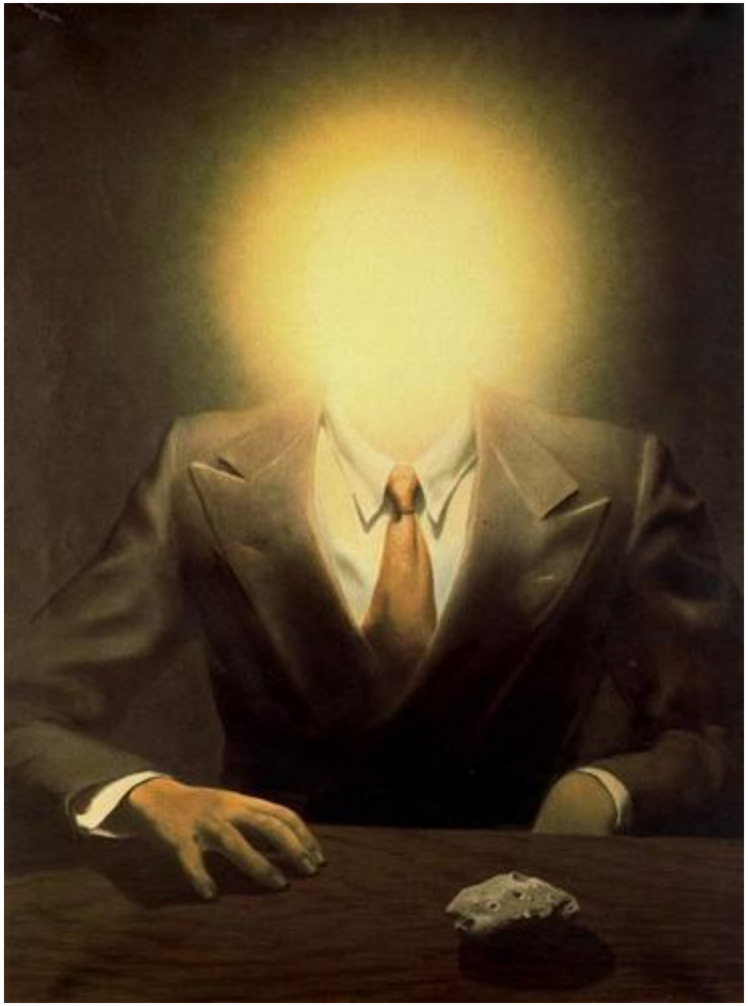

René Magritte, “The Pleasure Principle”

CSC320: Introduction to Visual Computing

Michael Guerzhoy

Slides from David Capel, Yael Amsterdamer

Digital Photography with Flash and No-Flash Image Pairs

Georg Petschnigg, Maneesh Agrawala, Hugues Hoppe
Richard Szeliski, Michael Cohen, Kentaro Toyama
Microsoft Corporation

Orig. (top) Detail Transfer (bottom) | Flash | No-Flash | Detail Transfer with Denoising

SIGGRAPH 2004

SIGGRAPH (**S**pecial **I**nterest **G**roup on **GRAPH**ics and Interactive Techniques) is the premier conference where new results on graphics and related fields are presented

To flash or not to flash?

- Natural lighting
- Low signal-to-noise ratio (SNR)
- Loss of details
- Longer exposure – motion blur

- Harsh, unnatural lighting
- High SNR
- More details
- May cause unwanted artifacts (red eye, shadows, specularities)

Slide by Yael Amsterdamer

Idea: combine both to get the best of both worlds

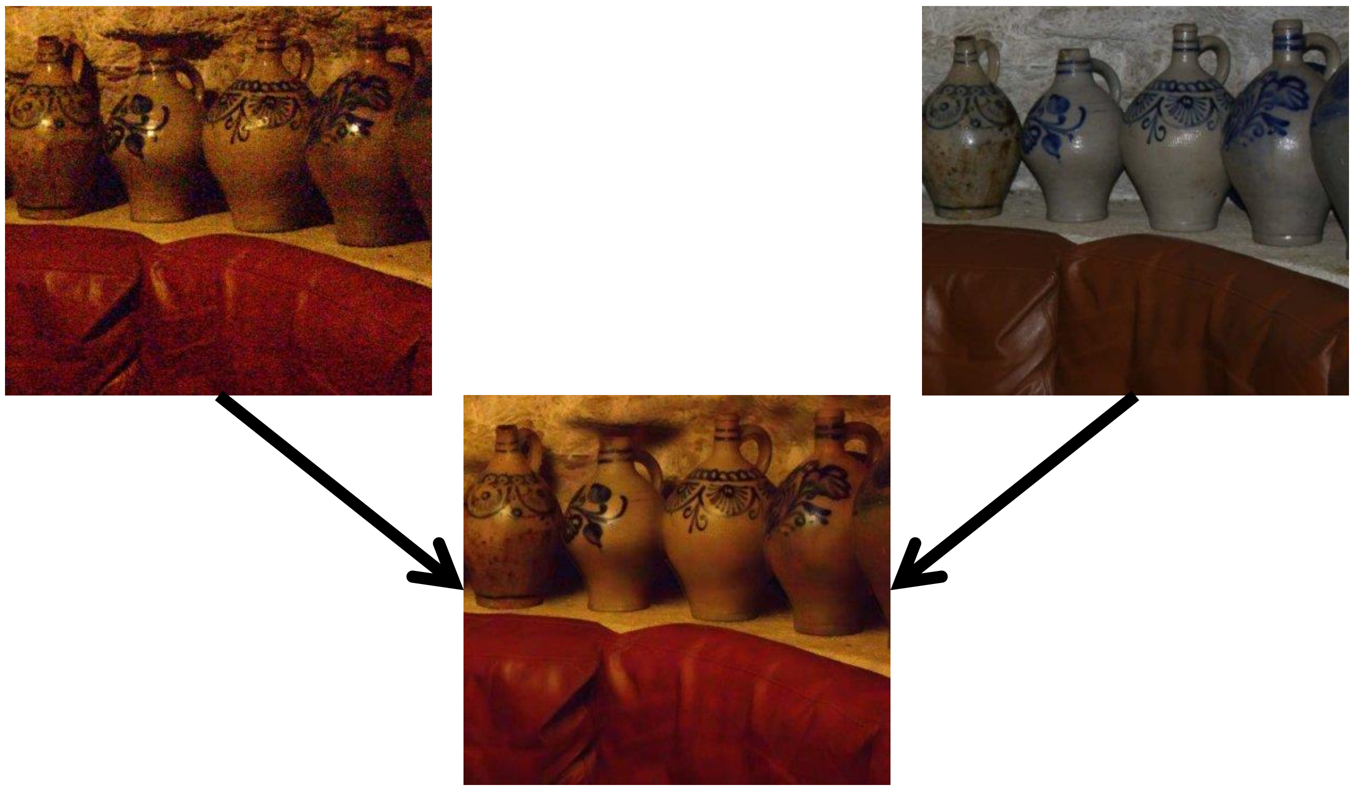

Slide by Yael Amsterdamer

Acquiring a flash/no flash photo pair

- Focus on the subject, lock camera settings
- Capture ambient image A
- Enable flash, capture flash image with minimal exposure time

Bilateral Filter

- Average pixels that are spatially close AND have similar intensities
- Implement edge-preserving smoothing

spatial kernel — intensity kernel

$$A_p^{Base} = \frac{1}{k(p)} \sum_{p' \in \Omega} g_d(p' - p)\, g_r(A_p - A_{p'})\, A_{p'}$$

(k(p) normalizes filter sum to 1)

spatial distance — intensity difference

- Away from edges, behaves like regular Gaussian smoothing
- Close to edges, the Gaussian kernel gets truncated
- Effect: this smoothing does not cross edges

spatial kernel → g_d; intensity kernel → g_r

$$A_p^{Base} = \frac{1}{k(p)} \sum_{p' \in \Omega} g_d(p' - p)\, g_r(A_p - A_{p'})\, A_{p'}$$

spatial distance → $p' - p$; intensity difference → $A_p - A_{p'}$

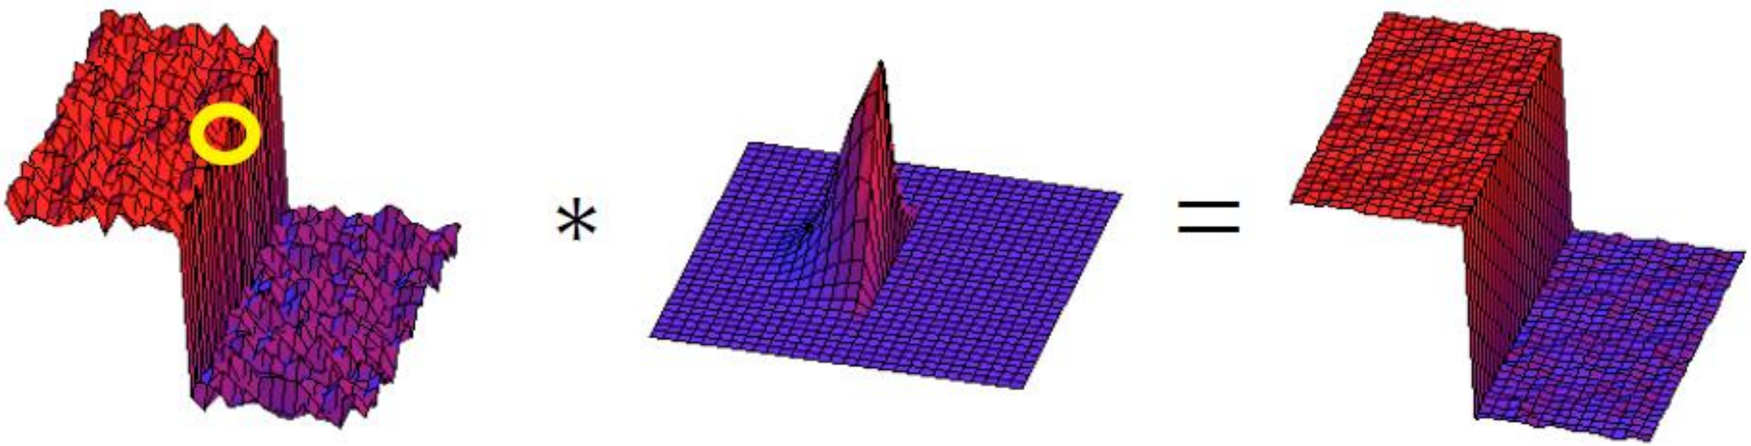

- Away from edges, it behaves like regular Gaussian smoothing
- Close to edges, the Gaussian kernel gets truncated
- Effect : Smoothing does not cross edges

Bilateral filtering example

Noisy image

Bilateral filtering

Joint Bilateral Filtering

- Flash image contains much better edge information
- **Idea: do smoothing in the ambient image, but use flash image to provide the edge-stopping**

spatial kernel — intensity kernel

$$A_p^{NR} = \frac{1}{k(p)} \sum_{p' \in \Omega} g_d(p' - p)\, g_r(F_p - F_{p'})\, A_{p'}$$

spatial distance (Ambient image) — intensity difference (in Flash image)

Flash/no flash Bilateral Filtering

(a) No-Flash Flash

(c) Denoised via Joint Bilateral Filter

Detail Transfer

- The flash image contains high-frequency content that may not be in the no-flash photo
- So transfer the details from the high-frequency image

Detail layer $F^{Detail} = \dfrac{F + \varepsilon}{F^{Base} + \varepsilon}$ $\varepsilon = 0.02$

- F^{Base} is the bilateral-filtered Flash image

Denoised image with detail transfer $A^{Final} = A^{NR} F^{Detail}$

(e) Detail Layer

(c) Denoised via Joint Bilateral Filter

(f) Detail Transfer

No-Flash | Detail Transfer with Denoising | Long Exposure Reference

Dark Flash Photography

Dilip Krishnan* Rob Fergus

Dept. of Computer Science, Courant Institute, New York University

Dark flash | Noisy ambient | Denoised | Ground truth

- Use "dark-flash" (UV+IR) to provide a bright, low-noise image
- Dark-flash is invisible to humans = unobtrusive
- Use flash *gradients* to guide reconstruction of ambient image

SIGGRAPH 2009

Dark-Flash Photography

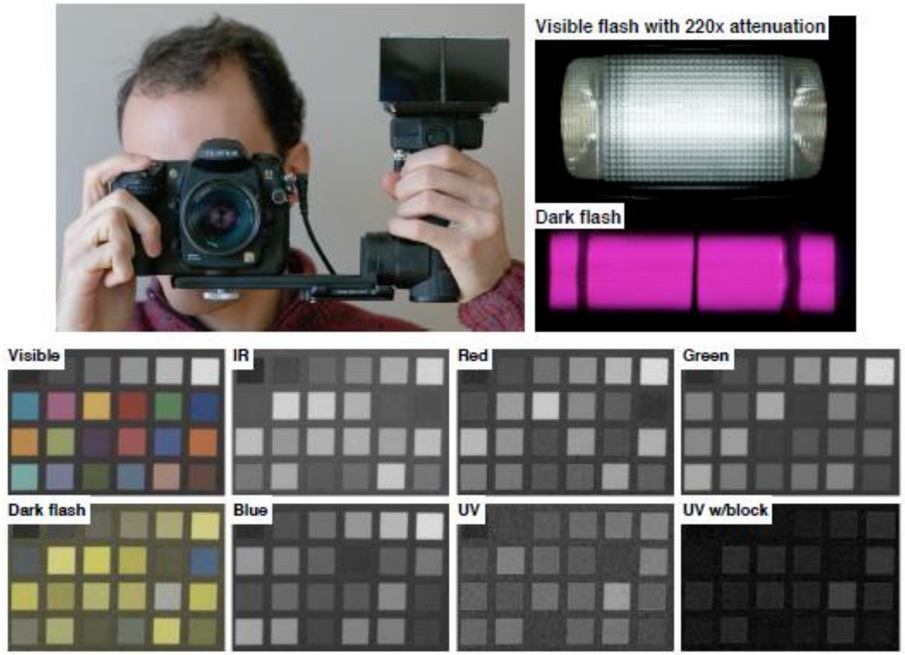

- Camera

  Fuji IS Pro (has no IR cut-filter)

- Flash

  Nikon SB-14 with UV absorbent coating removed from Xenon tube

  Hoya U360 filter to remove visible light (400nm to 700nm)

Spectral Response Curves

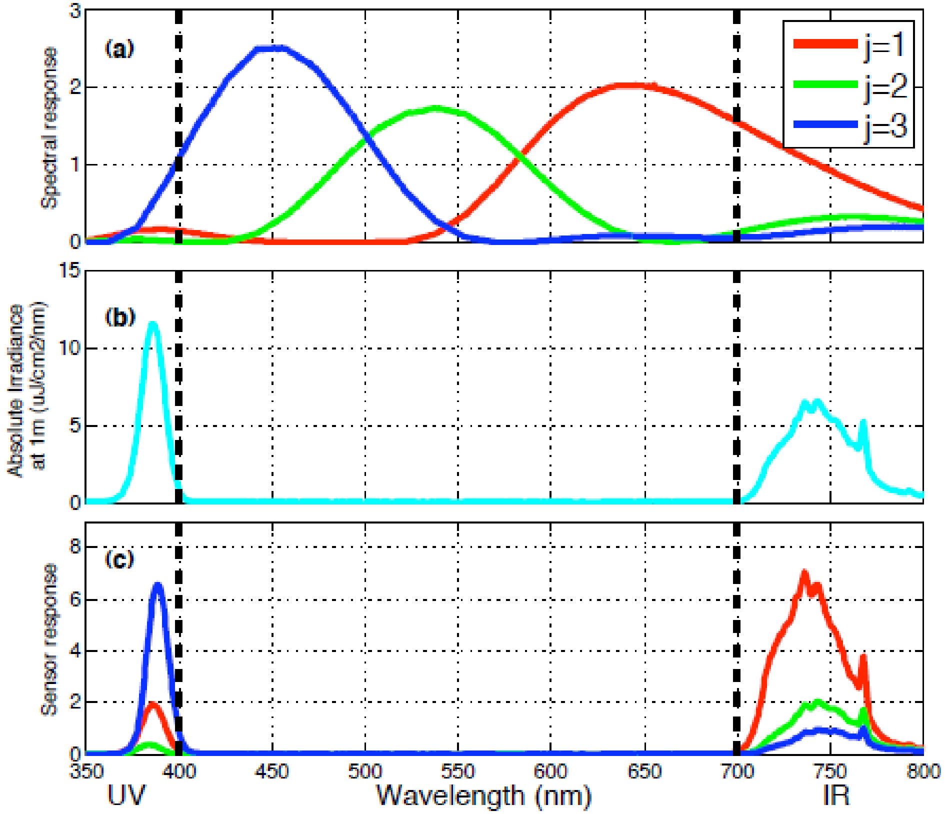

Camera response curves

Dark flash output

Camera response to dark flash

- Camera's response extends well into UV and IR regions
- Dark flash only emits light outside visible range
- UV appears mainly in camera Blue channel, IR in Red channel

Reconstructing the Image

- The dark-flash photo contains the details
- The ambient photo contains the right colours

Reconstructing the Image

$$\underset{R_j}{\operatorname{argmin}} \sum_p \Big[\underbrace{\mu_j \, m(p)(R_j(p) - A_j(p))^2}_{\text{Likelihood}} + \underbrace{\kappa \, m(p)|\nabla R_j(p)|^\alpha}_{\text{Spatial}} +$$

$$\underbrace{|\nabla R_j(p) - \nabla F_1(p)|^\alpha}_{\text{IR Spectral}} + \underbrace{|\nabla R_j(p) - \nabla F_3(p)|^\alpha}_{\text{UV Spectral}} \Big]$$

R is reconstructed image
A is ambient image
F_1,F_3 are IR,UV images

Reconstruction is guided by :

- Likelihood term = stay close to ambient image colors
- Spatial term = sparse smoothness prior, keep gradient small
- Spectral terms = keep gradients close to IR and UV gradients
- α=0.7 : sparse norm, encourages piecewise smoothness