# Nonphotorealistic rendering, and future cameras

## Computational Photography, 6.882

Bill Freeman
Fredo Durand
May 11, 2006

---

## Organization of NPR methods

- Automated methods
  - 2-d processing
  - 3-d processing
- Interactive methods
  - 2-d processing
  - 3-d processing

---

## Computer generated watercolor

Curtis has developed a method for automatic generation of watercolor effects using a water fluid simulation approach [Curtis97].

Figure 10: Computer Generated Watercolor, from [Curtis97]

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

---

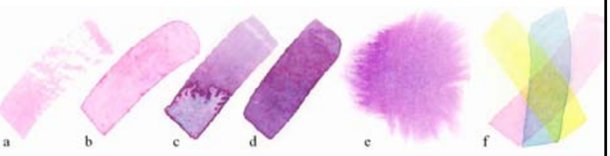

**Figure 1** Real watercolor effects: drybrush (a), edge darkening (b), backruns (c), granulation (d), flow effects (e), and glazing (f).

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

---

**Figure 2** Simulated watercolor effects created using our system.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

---

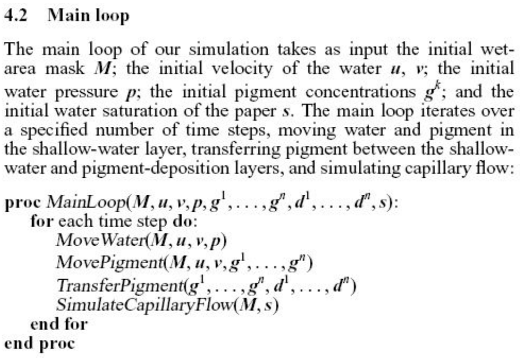

### 4.2 Main loop

The main loop of our simulation takes as input the initial wet-area mask $M$; the initial velocity of the water $u$, $v$; the initial water pressure $p$; the initial pigment concentrations $g^k$; and the initial water saturation of the paper $s$. The main loop iterates over a specified number of time steps, moving water and pigment in the shallow-water layer, transferring pigment between the shallow-water and pigment-deposition layers, and simulating capillary flow:

```
proc MainLoop($M, u, v, p, g^1, \ldots, g^n, d^1, \ldots, d^n, s$):
    for each time step do:
        MoveWater($M, u, v, p$)
        MovePigment($M, u, v, g^1, \ldots, g^n$)
        TransferPigment($g^1, \ldots, g^n, d^1, \ldots, d^n$)
        SimulateCapillaryFlow($M, s$)
    end for
end proc
```

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

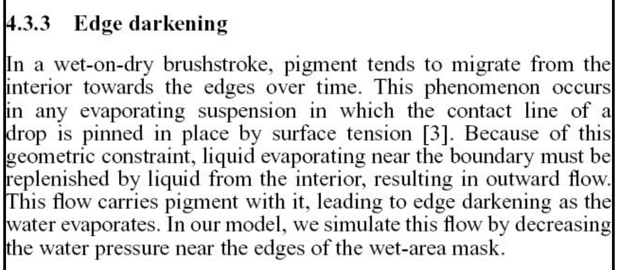

### 4.3.3 Edge darkening

In a wet-on-dry brushstroke, pigment tends to migrate from the interior towards the edges over time. This phenomenon occurs in any evaporating suspension in which the contact line of a drop is pinned in place by surface tension [3]. Because of this geometric constraint, liquid evaporating near the boundary must be replenished by liquid from the interior, resulting in outward flow. This flow carries pigment with it, leading to edge darkening as the water evaporates. In our model, we simulate this flow by decreasing the water pressure near the edges of the wet-area mask.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

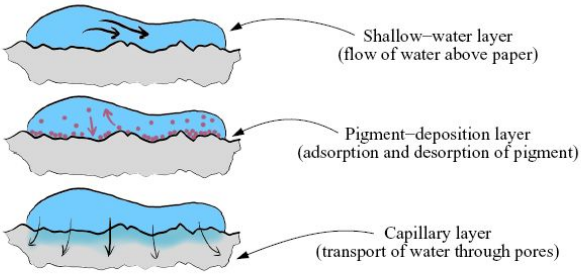

**Figure 3** The three-layer fluid model for a watercolor wash.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

**Figure 4** Example paper textures.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

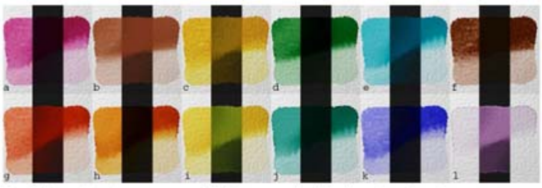

**Figure 5** Various synthetic pigments. The swatches were all created using identical initial conditions, with thicker pigment in the top half, and extra water in the upper left and lower right corners. The only changes from swatch to swatch are the pigments' optical and physical parameters, shown at right. The swatches are painted over a black stripe to distinguish the more opaque pigments such as "Indian Red" (b) from the more transparent ones such as "Brilliant Orange" (h).

|  | PIGMENT | $K_r$ | $K_g$ | $K_b$ | $S_r$ | $S_g$ | $S_b$ | $\rho$ | $\omega$ | $\gamma$ |
|---|---|---|---|---|---|---|---|---|---|---|
| a | "Quinacridone Rose" | 0.22 | 1.47 | 0.57 | 0.05 | 0.003 | 0.03 | 0.02 | 5.5 | 0.81 |
| b | "Indian Red" | 0.46 | 1.07 | 1.50 | 1.28 | 0.38 | 0.21 | 0.05 | 7.0 | 0.40 |
| c | "Cadmium Yellow" | 0.10 | 0.36 | 3.45 | 0.97 | 0.65 | 0.007 | 0.05 | 3.4 | 0.81 |
| d | "Hookers Green" | 1.62 | 0.61 | 1.64 | 0.01 | 0.012 | 0.003 | 0.09 | 1.0 | 0.41 |
| e | "Cerulean Blue" | 1.52 | 0.32 | 0.25 | 0.06 | 0.26 | 0.40 | 0.01 | 1.0 | 0.31 |
| f | "Burnt Umber" | 0.74 | 1.54 | 2.10 | 0.09 | 0.09 | 0.004 | 0.09 | 9.3 | 0.90 |
| g | "Cadmium Red" | 0.14 | 1.08 | 1.68 | 0.77 | 0.015 | 0.018 | 0.02 | 1.0 | 0.63 |
| h | "Brilliant Orange" | 0.13 | 0.81 | 3.45 | 0.005 | 0.009 | 0.007 | 0.01 | 1.0 | 0.14 |
| i | "Hansa Yellow" | 0.06 | 0.21 | 1.78 | 0.50 | 0.88 | 0.009 | 0.06 | 1.0 | 0.08 |
| j | "Phthalo Green" | 1.55 | 0.47 | 0.63 | 0.01 | 0.05 | 0.035 | 0.02 | 1.0 | 0.12 |
| k | "French Ultramarine" | 0.86 | 0.86 | 0.06 | 0.005 | 0.005 | 0.09 | 0.01 | 3.1 | 0.91 |
| l | "Interference Lilac" | 0.08 | 0.11 | 0.07 | 1.25 | 0.42 | 1.43 | 0.06 | 1.0 | 0.08 |

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

Interactive application. But too slow to let you paint in real-time.

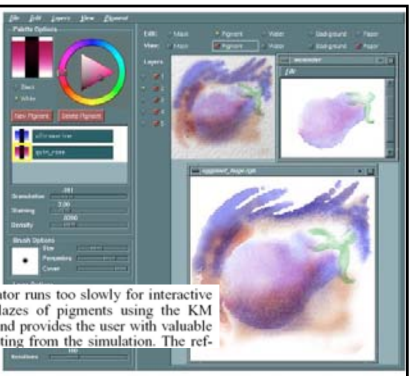

Although our watercolor simulator runs too slowly for interactive painting, compositing many glazes of pigments using the KM model is feasible in real-time, and provides the user with valuable feedback about the colors resulting from the simulation. The ref-

**Figure 7** An interactive painting application. At top center are the initial conditions painted by the user; at top right, a watercolor simulation in progress, showing two of the painting's five glazes. The large image is the finished painting.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

## Offline application

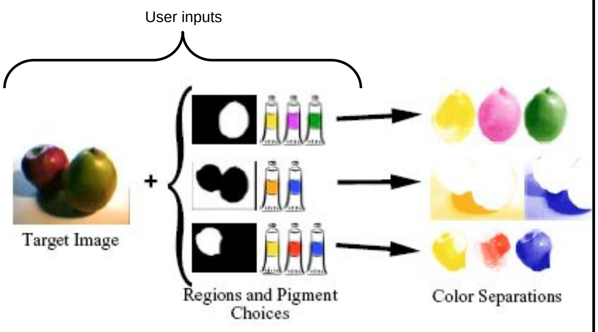

**Figure 8** Overview of the color separation process.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

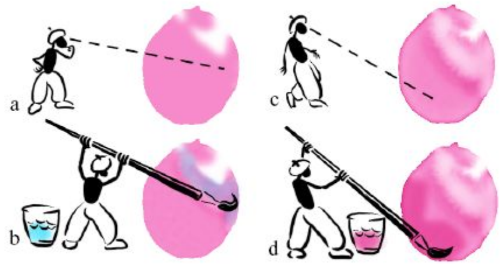

**Figure 9** Brushstroke planning. At given intervals, the planner identifies regions containing too much pigment (a) and thins them out by adding plain water (b). The planner can also compensate for a lack of pigment (c) by adding a pigmented wash (d).

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

**Figure 10** An automatic watercolorization (left) of a low resolution image captured using a poor-quality video camera (above). The finished painting consists of 11 glazes, using a total of 2750 iterations of the simulator, rendered at a resolution of 640 by 480 pixels in 7 hours on a 133 MHz SGI R4600 processor.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

Steps in the rendering

(user-supplied region mask not shown)

# Resulting watercolor

**Figure 10** An automatic watercolorization (left) of a low resolution image captured using a poor-quality video camera (above). The finished painting consists of 11 glazes, using a total of 2750 iterations of the simulator, rendered at a resolution of 640 by 480 pixels in 7 hours on a 133 MHz SGI R4600 processor.

http://citeseer.ist.psu.edu/cache/papers/cs/12745/http:zSzzSzcdserver.icemt.iastate.eduzSzcdzSzs97cpzSzcontentszSzpaperszSzcurtiszSzcurtis.pdf/curtis97computergenerated.pdf

# Source image for a 3d animation

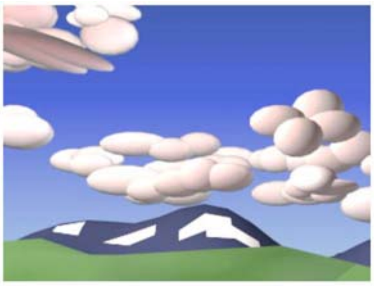

**Figure 12** The target image for Figure 14.

**Figure 14** Detail of one frame from Figure 13.

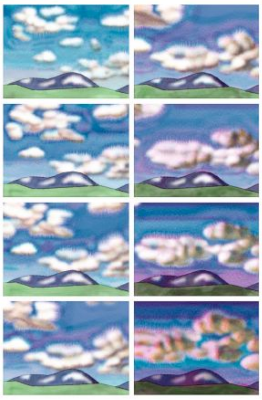

**Figure 13** Several frames from a non-photorealistic animation of moving clouds.

---

## Pen and ink illustration, exploiting 3d geometry

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

The work of Winkenbach provides emulation of pen-and-ink illustration by rendering 3D geometry in conjunction with stroke textures [Winkenbach94].

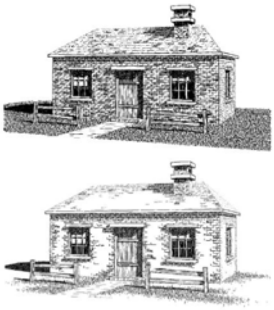

**Pen-and-ink renderings from [Winkenbach94]**

---

**Figure 5:** *Detail segments. The user interactively attaches "detail segments" to the surfaces to indicate roughly where details should appear.*

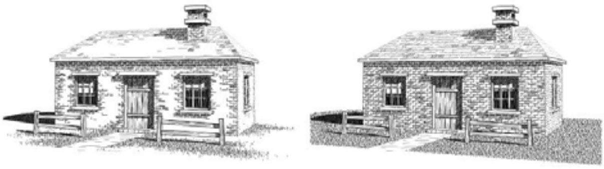

Figure 6: *Indicating texture. The left house is drawn using "indication"; the right house is not.*

---

**Non-Photorealistic Rendering - The Artist's Perspective**

***Simon Schofield***
***Slade School of Fine Art, University College London***

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

---

## The artist approved of this one…

**Non-Photorealistic Rendering - The Artist's Perspective**

***Simon Schofield***
***Slade School of Fine Art, University College London***

Winkenbach and Salesin's Pen and Ink work [19, 20] seem not to suffer as heavily from this problem as that of the "impressionist" systems. Their pen an ink renders are wholly algorithmic, yet posses a convincing hand drawn look without suffering from a seeming lack of authenticity. As discussed, we found that "blanket" or "blind" techniques, when applied overall to an image, often produced pleasing results, while techniques which simulated artists decisions, tended to fail. Etching, hatching and engraving methods can be, on a certain level, considered to be largely technical in their execution. Over smaller sections of an image, there may be very little consideration or decision given once a general approach to modelling form has been found. This does not mean that etchings or engravings are any less expressive than looser forms of painting, it simply means that the expressive decisions and virtuosity in these types of images lies elsewhere. In these cases, expression tends to lie in the overall composition, tonality and economy. Gustave Doré's genius lay not in his ability to engrave, which he often relied on others to do, but in his overall drawing, modelling and tone. It is still down to the user to provide these aspects when using Winkenbach and Salesin's systems.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

---

# Future cameras

Computational Photography,
6.882

Bill Freeman
Fredo Durand
May 11, 2006

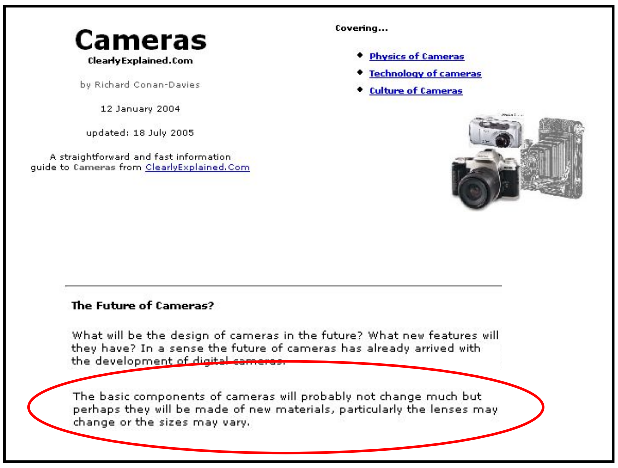

**Cameras**

ClearlyExplained.Com

by Richard Conan-Davies

12 January 2004

updated: 18 July 2005

A straightforward and fast information
guide to Cameras from ClearlyExplained.Com

**Covering...**

- Physics of Cameras
- Technology of cameras
- Culture of Cameras

**The Future of Cameras?**

What will be the design of cameras in the future? What new features will they have? In a sense the future of cameras has already arrived with the development of digital cameras.

The basic components of cameras will probably not change much but perhaps they will be made of new materials, particularly the lenses may change or the sizes may vary.

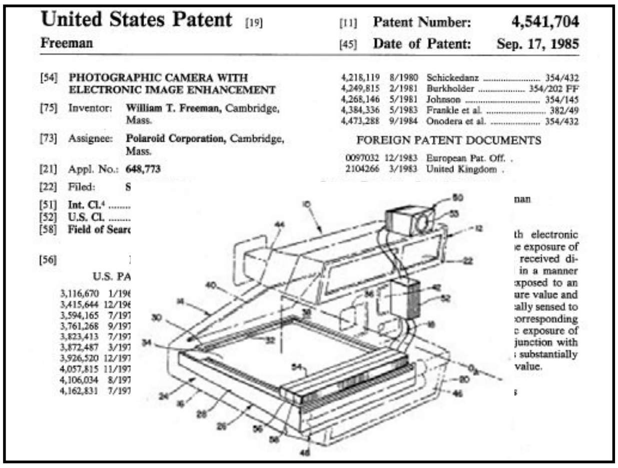

**United States Patent** [19]
**Freeman**

[11] **Patent Number:** **4,541,704**
[45] **Date of Patent:** **Sep. 17, 1985**

[54] **PHOTOGRAPHIC CAMERA WITH ELECTRONIC IMAGE ENHANCEMENT**

[75] Inventor: **William T. Freeman,** Cambridge, Mass.

[73] Assignee: **Polaroid Corporation,** Cambridge, Mass.

[21] Appl. No.: **648,773**

[22] Filed: S

[51] **Int. Cl.⁴** ........
[52] **U.S. Cl.** ........
[58] **Field of Searc**

[56]

U.S. PA

3,116,670 1/196
3,415,644 12/196
3,594,165 7/197
3,761,268 9/197
3,823,413 7/197
3,872,487 3/197
3,926,520 12/197
4,057,815 11/197
4,106,034 8/197
4,162,831 7/197

4,218,119 8/1980 Schickedanz .................... 354/432
4,249,815 2/1981 Burkholder .................. 354/202 FF
4,268,146 5/1981 Johnson ........................ 354/145
4,384,336 5/1983 Frankle et al. .................... 382/49
4,473,288 9/1984 Onodera et al. .................. 354/432

FOREIGN PATENT DOCUMENTS

0097032 12/1983 European Pat. Off. .
2104266 3/1983 United Kingdom .

nan

th electronic
e exposure of
received di-
in a manner
xposed to an
ure value and
ally sensed to
orresponding
c exposure of
junction with
substantially
value.

## What can be improved about current cameras?

(your list first…)

- Dynamic range
- Blurred photos
- Post-shot controllable depth of field
- Post-shot editable lighting, positions, etc.
- Size of camera

## What crazy other things?

- The previous list is all mostly with reference to the functionality of a film camera. Surely unexpected camera capabilities and uses, only possible with digital media, will come with future cameras.

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

## Color pixel mosaic

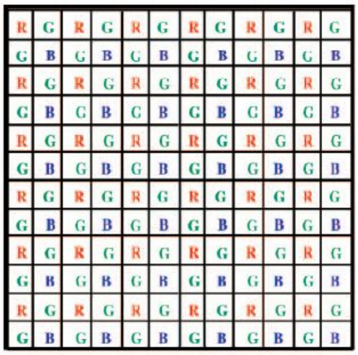

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

## Intensity attenuation mosaic

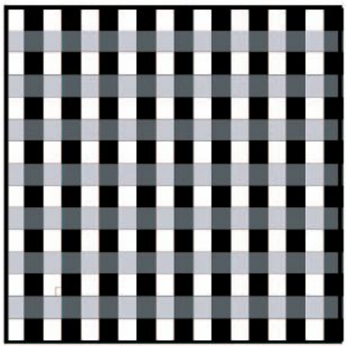

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

## Color and intensity mosaic

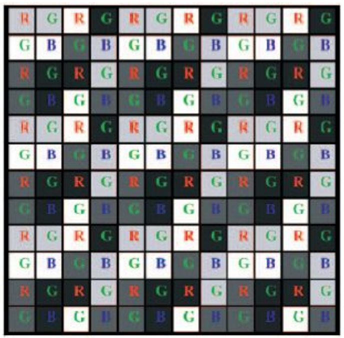

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

## Color and polarization mosaic

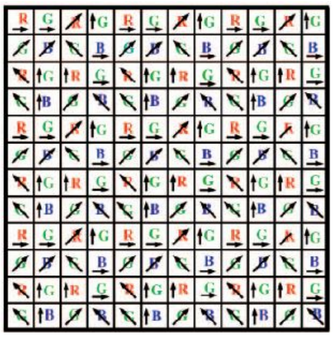

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

## Temporal sensitivity modulation

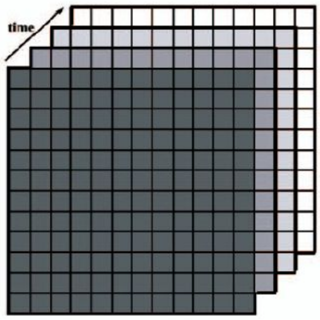

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

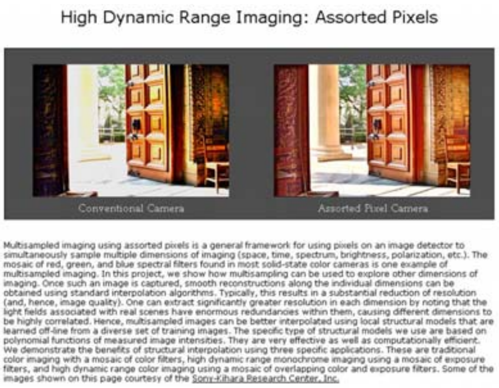

High Dynamic Range Imaging: Assorted Pixels

Multisampled imaging using assorted pixels is a general framework for using pixels on an image detector to simultaneously sample multiple dimensions of imaging (space, time, spectrum, brightness, polarization, etc.). The mosaic of red, green, and blue spectral filters found in most solid-state color cameras is one example of multisampled imaging. In this project, we show how multisampling can be used to explore other dimensions of imaging. Once such an image is captured, smooth reconstructions along the individual dimensions can be obtained using standard interpolation algorithms. Typically, this results in a substantial reduction of resolution (and, hence, image quality). One can extract significantly greater resolution in each dimension by noting that the light fields associated with real scenes have enormous redundancies within them, causing different dimensions to be highly correlated. Hence, multisampled images can be better interpolated using local structural models that are learned off-line from a diverse set of training images. The specific type of structural models we use are based on polynomial functions of measured image intensities. They are very effective as well as computationally efficient. We demonstrate the benefits of structural interpolation using three specific applications. These are traditional color imaging with a mosaic of color filters, high dynamic range monochrome imaging using a mosaic of exposure filters, and high dynamic range color imaging using a mosaic of overlapping color and exposure filters. Some of the images shown on this page courtesy of the Sony-Kihara Research Center, Inc.

Shree Nayar, Columbia University

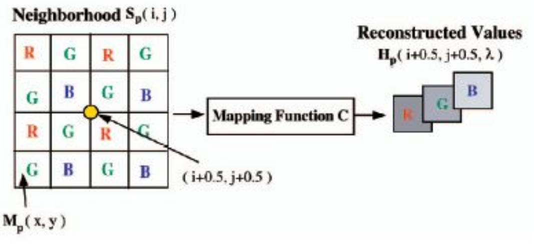

Fig. 8. The measured data $\mathbf{M_p}$ in the neighborhood $\mathbf{S_p}(i,j)$ around pixel $(i,j)$ are related to the high-quality color values $\mathbf{H_p}(i+0.5, j+0.5, \lambda)$ via a polynomial with coefficients $\mathbf{C_p}$.

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, VOL. 27, NO. 4, APRIL 2005

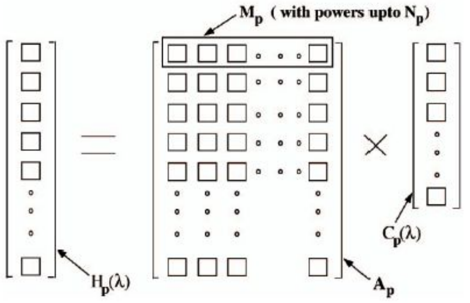

Fig. 9. The mapping function in (3) can be expressed as a linear system using matrices and vectors. For a given pattern $\mathbf{p}$ and color $\lambda$, $\mathbf{A_p}$ is the measurement matrix, $\mathbf{C_p}(\lambda)$ is the coefficient vector, and $\mathbf{H_p}(\lambda)$ is the reconstruction vector.

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, VOL. 27, NO. 4, APRIL 2005

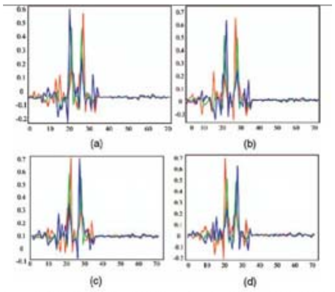

Fig. 10. Coefficients of the SVC structural model for the four different patterns. A second order polynomial is used and the neighborhood size is $6 \times 6$. The first 36 coefficients denote the coefficients corresponding to the first order terms and the remaining correspond to the second order terms in the SVC structural model. The first order terms capture most of the energy. (a) RGB coefficient for SVC pattern 1. (b) RGB coefficient for SVC pattern 2. (c) RGB coefficient for SVC pattern 3. (d) RGB coefficient for SVC pattern 4.

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

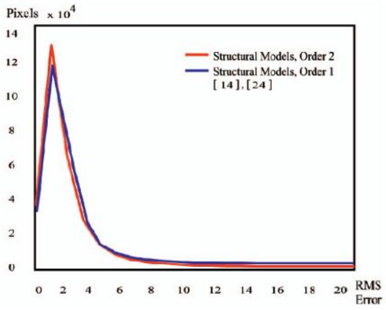

Fig. 17. Comparison of the histograms of luminance error (averaged over 20 test images) obtained by using structural models of orders 1 and 2, respectively. Note that the structural model with order 1 reduces to the works proposed in [14], [24]. Second order polynomials perform slightly better than the first order polynomials. Later, we will show that higher orders can have more significance in other multisampled images.

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

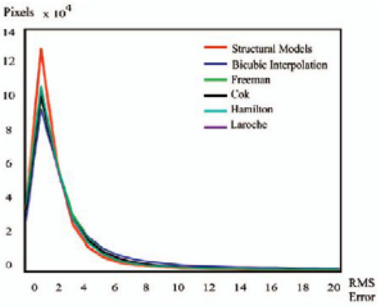

Fig. 16. Comparison of the histograms of luminance error (averaged over 20 test images) obtained by using five different methods. The average RMS error in luminance is 6.12 gray levels for bicubic interpolation and **3.27 gray levels for structural interpolation**. The RMS values for the methods—Freeman [11], Hamilton and Adams [13], Laroche and Prescott [15], and Cok [6]—are in between. Note that our method outperforms the other five methods. Structural: 3.7, Freeman 4.21, Hamilton 4.63, Laroche 4.78, Cok 4.96, and bicubic 6.12.

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

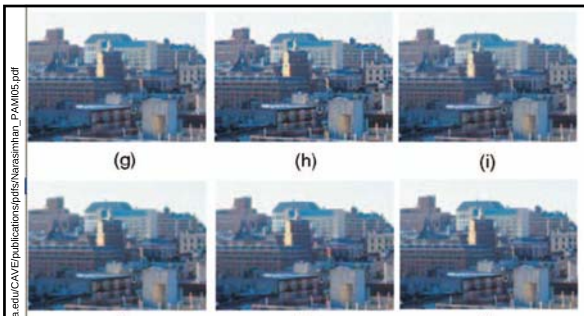

(j) (l)

Fig. 15. **(SVC Reconstruction)** Comparison of the reconstruction technique with bicubic interpolation and four other commonly used demosaicing techniques. Two regions of an urban scene are shown. As can be seen, the visual quality of our reconstruction is at least as good as the best demosaicing technique with the added advantage of avoiding spurious high frequencies. See Fig. 16 for a quantitative comparison. The methods used for demosaicing were obtained from four different papers [6], [11], [15], [13]. The implementations of these demosaicing methods (excluding bicubic interpolation) were obtained from Dr. Rajeev Ramanath [21]. (a) Original. (b) Structural Model. (c) Freeman. (d) Cok. (e) Laroche. (f) Hamilton. (g) Original. (h) Structural Model. (i) Freeman. (j) Cok. (k) Laroche. (l) Hamilton.

## Original (12 bits)


(a)

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

---

## As sampled (8 bits)

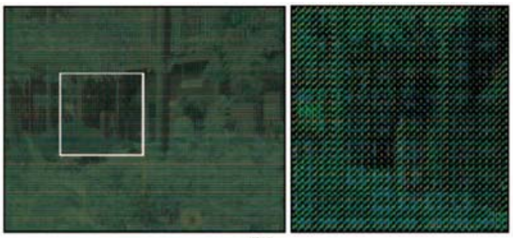
(b)

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

---

## Cubic spline interpolation to 12 bits


(c)

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

---

## Linear regression interpolation to 12 bits


(d)

http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.pdf

---

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

---

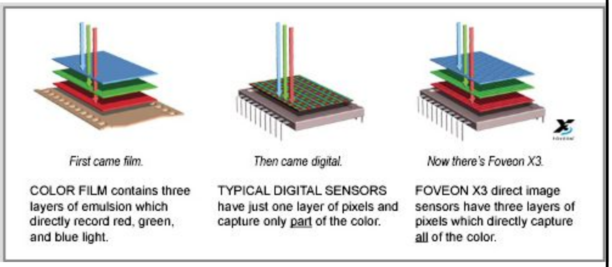
*First came film.*

COLOR FILM contains three layers of emulsion which directly record red, green, and blue light.

*Then came digital.*

TYPICAL DIGITAL SENSORS have just one layer of pixels and capture only part of the color.

*Now there's Foveon X3.*

FOVEON X3 direct image sensors have three layers of pixels which directly capture all of the color.

http://www.foveon.com/

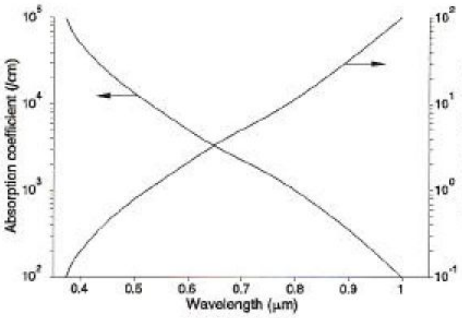

*Figure 1 Absorption coefficient and penetration depth in Silicon vs. wavelength.*

http://www.foveon.com/files/CIC13_Hubel_Final.pdf

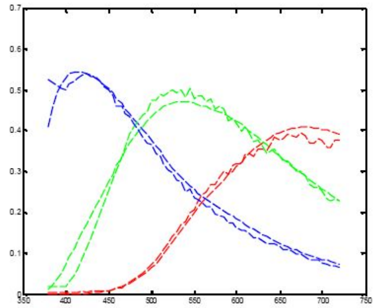

*Figure 2. Dashed curves show the color separation computed with a simple model of light absorption in silicon and the solid line show measurements measured from an early Foveon X3 image sensor.*

http://www.foveon.com/files/CIC13_Hubel_Final.pdf

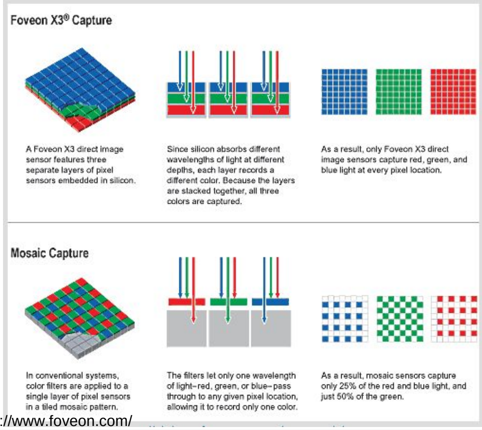

**Foveon X3® Capture**

A Foveon X3 direct image sensor features three separate layers of pixel sensors embedded in silicon.

Since silicon absorbs different wavelengths of light at different depths, each layer records a different color. Because the layers are stacked together, all three colors are captured.

As a result, only Foveon X3 direct image sensors capture red, green, and blue light at every pixel location.

**Mosaic Capture**

In conventional systems, color filters are applied to a single layer of pixel sensors in a tiled mosaic pattern.

The filters let only one wavelength of light–red, green, or blue–pass through to any given pixel location, allowing it to record only one color.

As a result, mosaic sensors capture only 25% of the red and blue light, and just 50% of the green.

http://www.foveon.com/

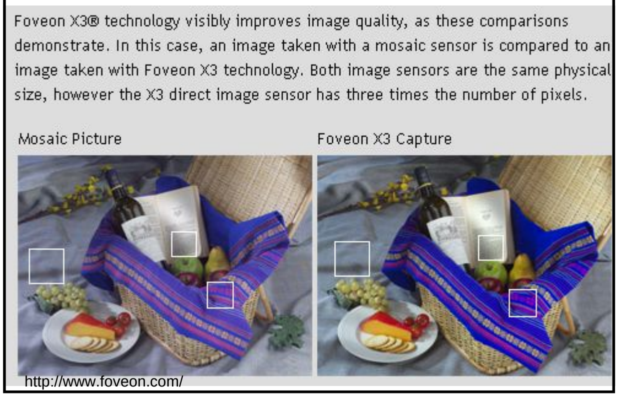

Foveon X3® technology visibly improves image quality, as these comparisons demonstrate. In this case, an image taken with a mosaic sensor is compared to an image taken with Foveon X3 technology. Both image sensors are the same physical size, however the X3 direct image sensor has three times the number of pixels.

Mosaic Picture

Foveon X3 Capture

http://www.foveon.com/

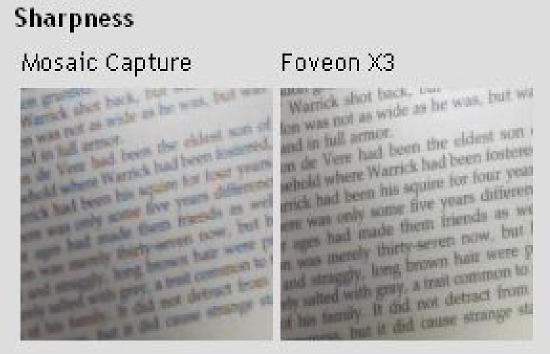

**Sharpness**

Mosaic Capture

Foveon X3

http://www.foveon.com/

**Color Detail**

Mosaic Capture

Foveon X3

http://www.foveon.com/

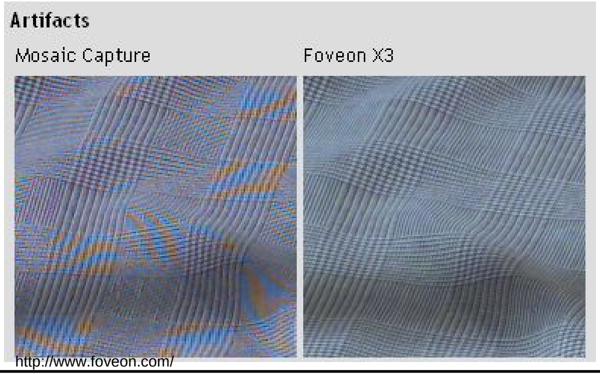

http://www.foveon.com/

## Foveon features

- Use the optical properties of silicon itself to separate colors.
  - Different wavelengths get absorbed at different depths of the silicon—blue, then green, then red.
- More efficient at capturing light—don't discard 2/3 of the spectrum at each pixel.
- Variable pixel size, depending on photo mode or video mode.

- 2002: "…destined to become the standard in image sensors for electronic cameras.", said Carver Mead, Foveon's founder. Status now…?

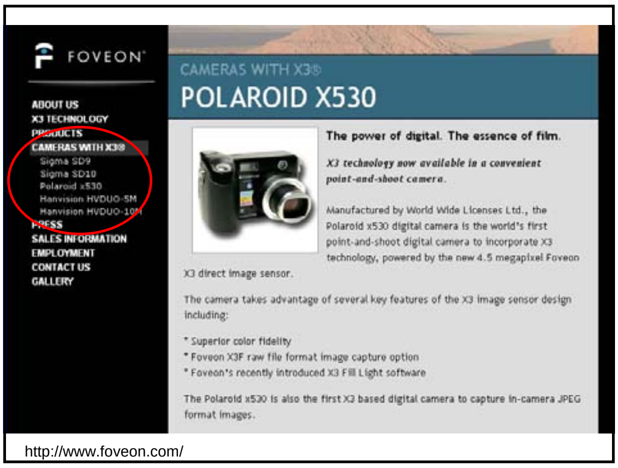

CAMERAS WITH X3®

POLAROID X530

**The power of digital. The essence of film.**

*X3 technology now available in a convenient point-and-shoot camera.*

Manufactured by World Wide Licenses Ltd., the Polaroid x530 digital camera is the world's first point-and-shoot digital camera to incorporate X3 technology, powered by the new 4.5 megapixel Foveon X3 direct image sensor.

The camera takes advantage of several key features of the X3 image sensor design including:

* Superior color fidelity
* Foveon X3F raw file format image capture option
* Foveon's recently introduced X3 Fill Light software

The Polaroid x530 is also the first X3 based digital camera to capture in-camera JPEG format images.

http://www.foveon.com/

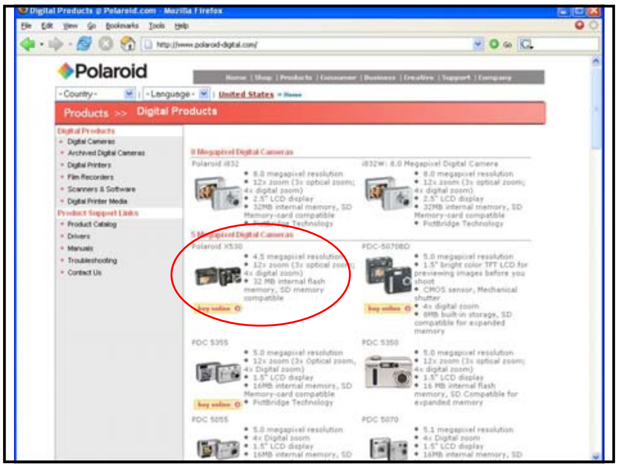

8 Megapixel Digital Cameras

Polaroid i832
- 8.0 megapixel resolution
- 12x zoom (3x optical zoom; 4x digital zoom)
- 2.5" LCD display
- 32MB internal memory, SD Memory-card compatible
- PictBridge Technology

i832W: 8.0 Megapixel Digital Camera
- 8.0 megapixel resolution
- 12x zoom (3x optical zoom; 4x digital zoom)
- 2.5" LCD display
- 32MB internal memory, SD Memory-card compatible
- PictBridge Technology

5 Megapixel Digital Cameras

Polaroid X530
- 4.5 megapixel resolution
- 12x zoom (3x optical zoom; 4x digital zoom)
- 32 MB internal flash memory, SD memory compatible

PDC-5070BD
- 5.0 megapixel resolution
- 1.5" bright color TFT LCD for previewing images before you shoot
- CMOS sensor, Mechanical shutter
- 4x digital zoom
- 8MB built-in storage, SD compatible for expanded memory

PDC 5355
- 5.0 megapixel resolution
- 12x zoom (3x Optical zoom, 4x Digital zoom)
- 1.5" LCD display
- 16MB internal memory, SD Memory-card compatible
- PictBridge Technology

PDC 5350
- 5.0 megapixel resolution
- 12x zoom (3x optical zoom; 4x digital zoom)
- 1.5" LCD display
- 16 MB internal flash memory, SD Compatible for expanded memory

PDC 5055
- 5.0 megapixel resolution
- 4x Digital zoom
- 1.5" LCD display
- 16MB internal memory, SD

PDC 5070
- 5.1 megapixel resolution
- 4x Digital zoom
- 1.5" LCD display
- 16MB internal memory, SD

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

**Coded Exposure Photography: Motion Deblurring using Fluttered Shutter**

Ramesh Raskar
Amit Agrawal
Mitsubishi Electric Research Laboratories (MERL)

Jack Tumblin
Northwestern University

Sponsored by ACM**SIGGRAPH**
2B0ST0N6

## Blurring convolution kernel

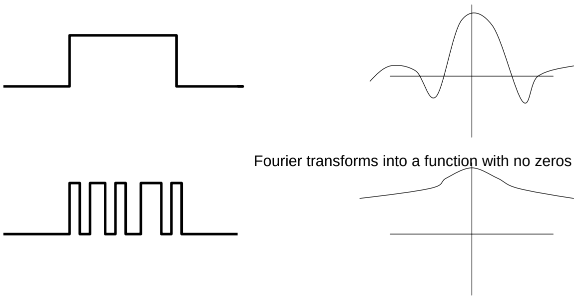

Fourier transforms into sinc

Fourier transforms into a function with no zeros

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

**Light Field Photography with a Hand-held Plenoptic Camera**

Ren Ng[*] Marc Levoy[*] Mathieu Brédif[*] Gene Duval[†] Mark Horowitz[*] Pat Hanrahan[*]
[*]Stanford University [†]Duval Design

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

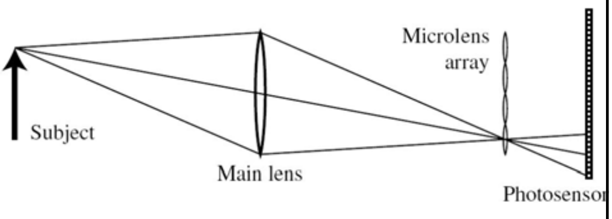

**Figure 1:** Conceptual schematic (not drawn to scale) of our camera, whic
is composed of a main lens, microlens array and a photosensor. The mai
lens focuses the subject onto the microlens array. The microlens array sep
arates the converging rays into an image on the photosensor behind it.

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

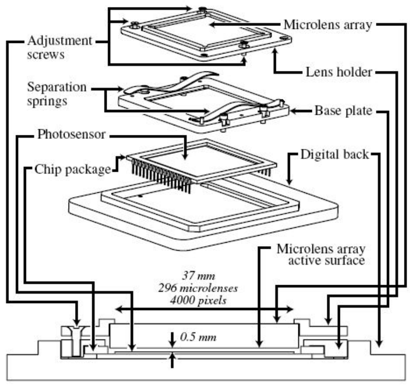

**Figure 8:** *Top:* Exploded view of assembly for attaching the microlens array to the digital back. *Bottom:* Cross-section through assembled parts.

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

**Figure 9:** Our light field camera in use.

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

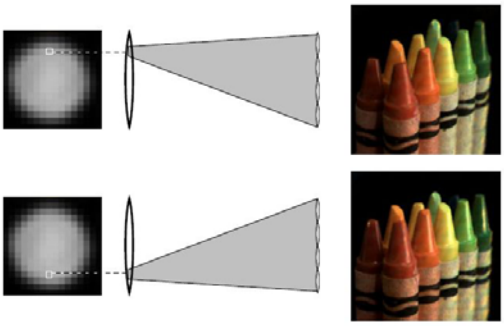

**Figure 4:** Two sub-aperture photographs obtained from a light field by extracting the shown pixel under each microlens (depicted on left). Note that the images are not the same, but exhibit vertical parallax.

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

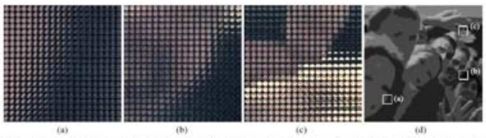

Figure 13: A complete light field captured by our prototype. Careful examination (zoom in on electronic version, or use magnifying glass in print) reveals 292 × 292 microlens images, each approximately 0.5 mm wide in print. Note the corner-to-corner clarity and detail of microlens images across the light field, which illustrates the quality of our microlens focusing. (a), (b) and (c) show magnified views of regions outlined on the key in (d). These close-ups are representative of three types of edges that can be found througout the image. (a) illustrates microlenses at depths closer than the focal plane. In these *right-side up* microlens images, the woman's cheek appears on the left, as it appears in the macroscopic image. In contrast, (b) illustrates microlenses at depths further than the focal plane. In these *inverted* microlens images, the man's cheek appears on the *right*, opposite the macroscopic world. This effect is due to inversion of the microlens' rays as they pass through the world focal plane before arriving at the main lens. Finally, (c) illustrates microlenses on edges *at* the focal plane (the fingers that are clasped together). The microlenses at this depth are constant in color because all the rays arriving at the microlens originate from the same point on the fingers, which reflect light diffusely.

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

**Figure 14:** Refocusing after a single exposure of the light field camera. Top is the photo that would have resulted from a conventional camera, focused on the clasped fingers. The remaining images are photographs refocused at different depths: middle row is focused on first and second figures; last row is focused on third and last figures. Compare especially middle left and bottom right for full effective depth of field.

---

http://graphics.stanford.edu/papers/lfcamera/lfcamera-150dpi.pdf

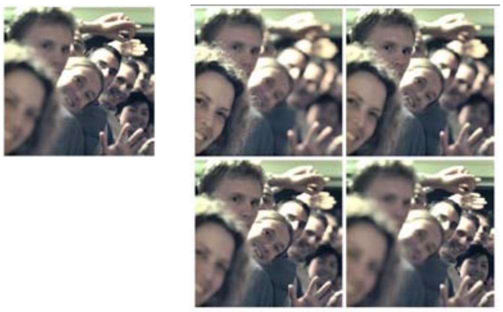

**Figure 14:** Refocusing after a single exposure of the light field camera. Top is the photo that would have resulted from a conventional camera, focused on the clasped fingers. The remaining images are photographs refocused at different depths: middle row is focused on first and second figures; last row is focused on third and last figures. Compare especially middle left and bottom right for full effective depth of field.

---

## Some possible future directions

- Assorted pixels
- Foveon imager
- Coded shutter flutter
- Light field camera
- Gradient camera

---

### Why I want a Gradient Camera

Jack Tumblin
Northwestern University
jet@cs.northwestern.edu

Amit Agrawal
University of Maryland
aagrawal@umd.edu

Ramesh Raskar
MERL
raskar@merl.com

#### Abstract

*We propose a camera that measures static gradients instead of static intensities. Quantizing sensed intensity differences between adjacent pixel values permits an ordinary A/D converter to measure detailed high contrast (HDR) scenes. We measure alternating 'cliques' of sensors (small groups) that locally determine their own best exposure, and reconstruct the image using a Poisson solver. This intrinsically differential design suppresses common-mode noise, hides and smoothes quantization, and can correct for its own saturated sensors. Simulations demonstrate these capabilities in side-by-side comparisons.*

#### 1 Introduction

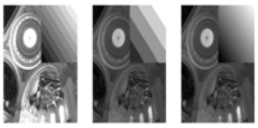

Figure 1. A log-gradient camera captures both large and small scene contrasts well: (a) an 8-bit intensity camera loses large contrasts ($> 10^{+5} : 1$) at A/D limits, but small contrasts (inset: 1.33 : 1 ramp, simulated) are finely quantized (11 levels); (b) an 8-bit log-intensity camera captures large contrasts well, but small contrasts are coarsely quantized (5 levels); (c) an 8 bit log-gradient camera preserves both, hiding errors smoothly everywhere.

---

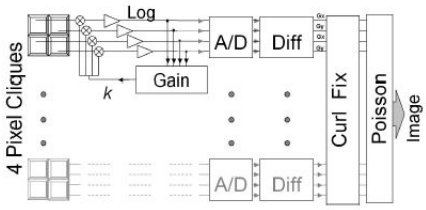

Figure 2. Log-gradient camera overview: intensity sensors organized into 4-pixel cliques share the same self-adjusting gain setting $k$, and send $\log(I_d)$ signals to A/D converter. Subtraction removes common-mode noise, and a linear 'curl fix' solver corrects saturated gradient values or 'dead' pixels,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

---

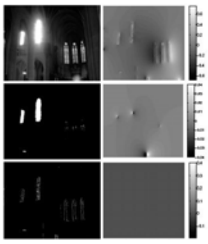

Figure 4. Sensor Error Correction: In the left column, an 8-bit intensity camera captures the Nave HDR scene [top], but loses 2.29% pixels to unmeasurable white; these pixels are marked white in [mid]. However, an 8-bit, 1% step log-gradient camera can measure all but 2788 or 0.41% of its gradients; these are marked white in [bott]. In the right column, [top] shows intensity error caused by reconstructing uncorrected gradients; the Poisson solver propagates these errors widely across the image. After curl correction, only 4 disjoint graphs remain, with a total of 28 unknown gradients, causing 4 small dimple-like errors in intensity [mid]. After disjoint graph correction [bot], error falls to 4 dots caused by underestimated offsets for the 4 graphs.

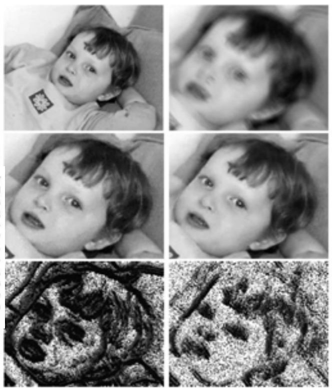

Figure 7. Gradient cameras reject motion and reduce noise better than intensity cameras when used as a gated estimator. A gated estimator finds the average of a stream of input values, locks its value permanently if given an outlier. [Top Row] source scene; time-average of source image + Gaussian noise moved in a 5-pixel-wide circle in 15 frame-times. [middle row]: Gated intensity camera result; Gated gradient camera result. [bottom row]:time-to-lock (from 0 to 15) for intensity camera; time-to-lock for gradient camera. Both pixel-gated cameras avoided motion blur well, but intensity camera lock time was much shorter and the results show much more noise. Similar adjacent gradients are much more common than similar adjacent intensities. Please use zoom with your PDF browser.

---

end