Where does volume and point data come from?

Marc Levoy

Computer Science Department
Stanford University

Three theses

Thesis #1: Many sciences lack good visualization tools.
Corollary: These are a good source for volume and point data.

Thesis #2: Computer scientists need to learn these sciences.
Corollary: Learning the science may lead to new visualizations.

Thesis #3: We also need to learn their data capture technologies.
Corollary: Visualizing the data capture process helps debug it.

Success story #1: volume rendering of medical data

Karl-Heinz Hoehne

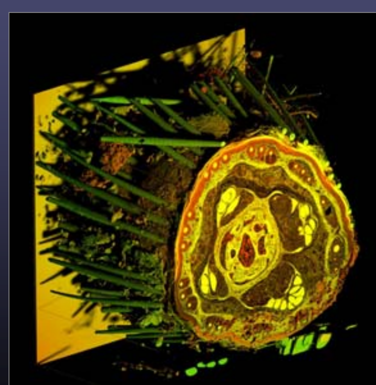

Resolution Sciences

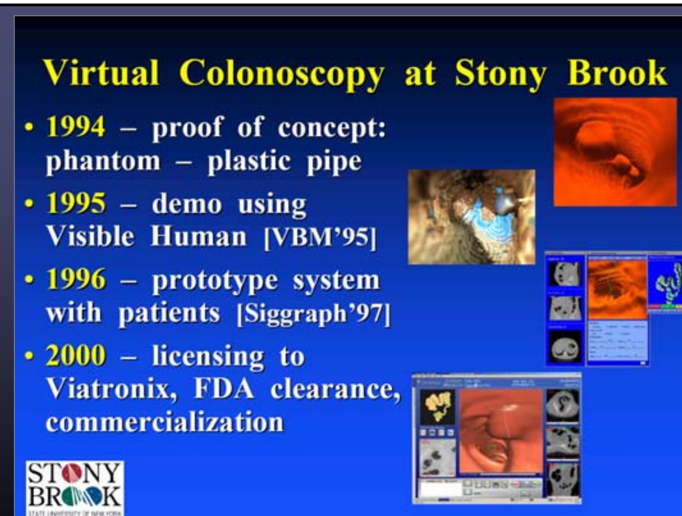

Arie Kaufman et al.

Success story #2: point rendering of dense polygon meshes

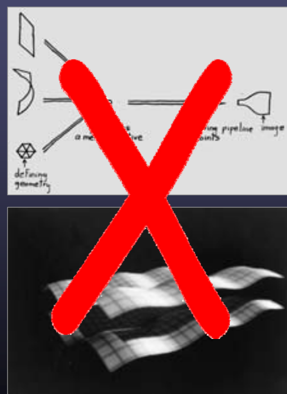

Levoy and Whitted (1985)

Szymon Rusinkiewicz's QSplat (2000)

Failure: volume rendering in the biological sciences

Bright	100
Gamma	1
Max Threshold	255
Min Threshold	0
Projection	Alpha
Storage Hint	Texture3D
View Aligned Slices	False
Volume Size	256

- *(a leading software package)*
  - limited control over opacity transfer functio
  - no control over surface appearance or lighting
  - no quantitative 3D probes
- Photoshop
  - converting 16-bit to 8-bit dithers the low-order bit
  - PhotoMerge (image mosaicing) performs poorly
  - no support for image stacks, volumes, n-D images

What's going on in the basic sciences?

- new instruments ⇒ scientific discoveries
- most important new instrument in the last 50 years: the digital computer
- computers + digital sensors = computational imaging

  Def: imaging methods in which computation is inherent in image formation.

  – B.K. Horn
- the revolution in medical imaging (CT, MR, PET, etc.) is now happening all across the basic sciences

(It's also a great source for volume and point data!)

Examples of computational imaging in the sciences

- medical imaging
  - rebinning ← *inspiration for light field rendering*
  - ✓ – transmission tomography
  - reflection tomography (for ultrasound)
- geophysics
  - ✓ – borehole tomography
  - seismic reflection surveying
- applied physics
  - ✓ – diffuse optical tomography
  - ✓ – diffraction tomography
  - ✓ – scattering and inverse scattering

✓ in this talk

- biology
  - ✓ – confocal microscopy ← *applicable at macro scale too*
  - ✓ – deconvolution microscopy
- astronomy
  - ✓ – coded-aperture imaging
  - – interferometric imaging
- airborne sensing
  - – multi-perspective panoramas
  - – synthetic aperture radar

- optics
  - – holography
  - – wavefront coding

Computational imaging technologies used in neuroscience

- Magnetic Resonance Imaging (MRI)
- Positron Emission Tomography (PET)
- Magnetoencephalography (MEG)
- Electroencephalography (EEG)
- Intrinsic Optical Signal (IOS)
- *In Vivo* Two-Photon (IVTP) Microscopy
- Microendoscopy
- Luminescence Tomography
- New Neuroanatomical Methods (3DEM, 3DLM)

The Fourier projection-slice theorem

(a.k.a. the central section theorem) [Bracewell 1956]

$P_\theta(t)$ → 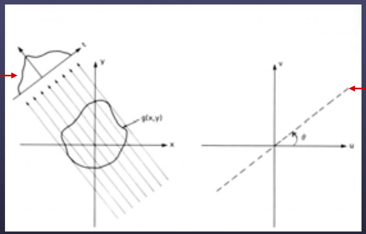 ← $G_\theta(\omega)$

(from Kak)

- $P_\theta(t)$ is the integral of g(x,y) in the direction θ
- G(u,v) is the 2D Fourier transform of g(x,y)
- $G_\theta(\omega)$ is a 1D slice of this transform taken at θ
- $\mathcal{F}^{-1}\ \{ G_\theta(\omega) \} = P_\theta(t)$!

Reconstruction of g(x,y) from its projections

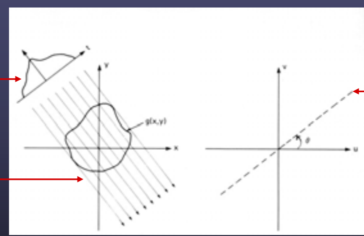

$$P_\theta(t) = \int_{-\infty}^{+\infty} P_\theta(t,s)\, ds$$

(from Kak)

- add slices $G_\theta(\omega)$ into u,v at all angles θ and inverse transform to yield g(x,y), or
- add 2D backprojections $P_\theta(t, s)$ into x,y at all angles θ

The need for filtering before (or after) backprojection

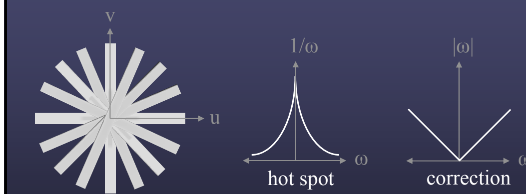

- sum of slices would create 1/ω hot spot at origin
- correct by multiplying each slice by |ω|, or
- convolve $P_\theta(t)$ by $\mathcal{F}^{-1}$ { |ω| } before backprojecting
- this is called filtered backprojection

Summing filtered backprojections

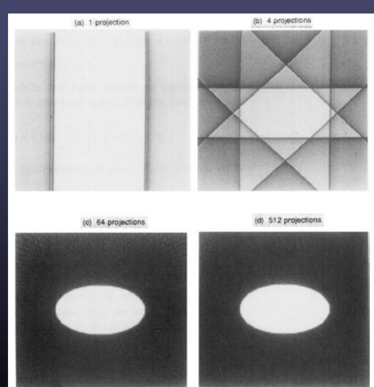

(from Kak)

Example of reconstruction by filtered backprojection

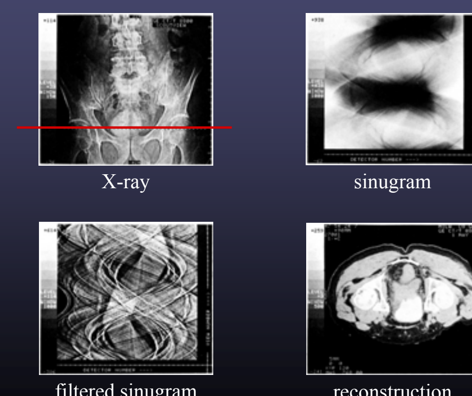

(from Herman)

More examples

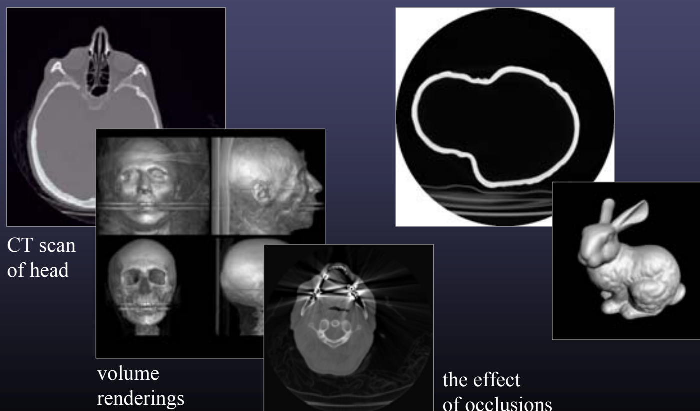

CT scan of head

volume renderings

the effect of occlusions

Limited-angle projections

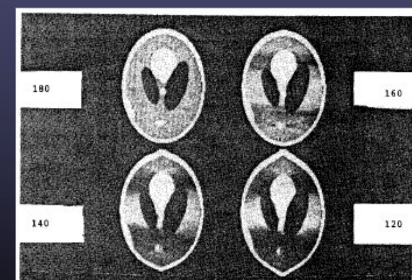

[Olson 1990]

Reconstruction using the Algebraic Reconstruction Technique (ART)

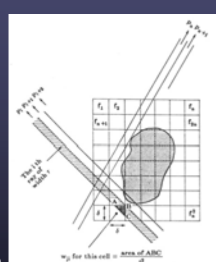

(from Kak)

$$p_i = \sum_{j=1}^{N} w_{ij} f_j, \quad i = 1, 2, \ldots, M$$

M projection rays
N image cells along a ray
p_i = projection along ray i
f_j = value of image cell j (n^2 cells)
w_{ij} = contribution by cell j to ray i
(a.k.a. resampling filter)

- applicable when projection angles are limited or non-uniformly distributed around the object
- can be under- or over-constrained, depending on N and M

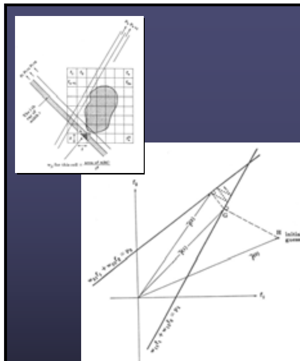

$$\bar{f}^{(k)} = \bar{f}^{(k-1)} - \frac{\bar{f}^{(k-1)} \bullet (\bar{w}_i - p_i)}{\bar{w}_i \bullet \bar{w}_i} \bar{w}_i$$

$\bar{f}^{(k)} = k^{\text{th}}$ estimate of all cells
$\bar{w}_i$ = weights $(w_{i1}, w_{i2}, \ldots, w_{iN})$ along ray i

Procedure

- make an initial guess, e.g. assign zeros to all cells
- project onto p_1 by increasing cells along ray 1 until $\Sigma = p_1$
- project onto p_2 by modifying cells along ray 2 until $\Sigma = p_2$, etc.
- to reduce noise, reduce by $\alpha \, \Delta \bar{f}^{(k)}$ for $\alpha < 1$

- linear system, but big, sparse, and noisy
- ART is solution by *method of projections* [Kaczmarz 1937]
- to increase angle between successive hyperplanes, jump by 90°
- SART modifies <u>all</u> cells using $f^{(k-1)}$, then increments k
- overdetermined if M > N, underdetermined if missing rays
- optional additional constraints:
  - $f > 0$ everywhere (positivity)
  - $f = 0$ outside a certain area

<u>Procedure</u>

- make an initial guess, e.g. assign zeros to all cells
- project onto p_1 by increasing cells along ray 1 until $\Sigma = p_1$
- project onto p_2 by modifying cells along ray 2 until $\Sigma = p_2$, etc.
- to reduce noise, reduce by $\alpha \Delta \bar{f}^{(k)}$ for $\alpha < 1$

- linear system, but big, sparse, and noisy
- ART is solution by *method of projections* [Kaczmarz 1937]
- to increase angle between successive hyperplanes, jump by 90°
- SART modifies <u>all</u> cells using $f^{(k-1)}$, then increments k
- overdetermined if M > N, underdetermined if missing rays
- optional additional constraints:
  - $f > 0$ everywhere (positivity)
  - $f = 0$ outside a certain area

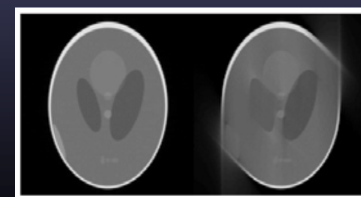

[Olson]

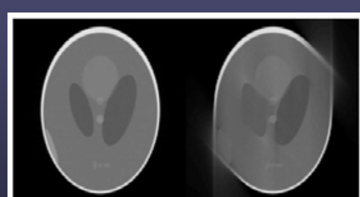

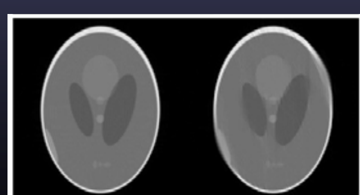

[Olson]

Borehole tomography

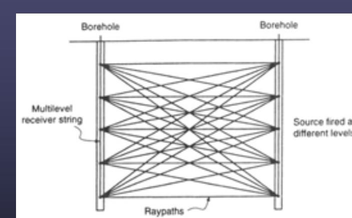

(from Reynolds)

- receivers measure end-to-end travel time
- reconstruct to find velocities in intervening cells
- must use limited-angle reconstruction methods (like ART)

Applications

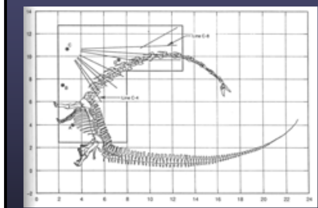

mapping a *seismosaurus* in sandstone using microphones in 4 boreholes and explosions along radial lines

mapping ancient Rome using explosions in the subways and microphones along the streets?

Optical diffraction tomography (ODT)

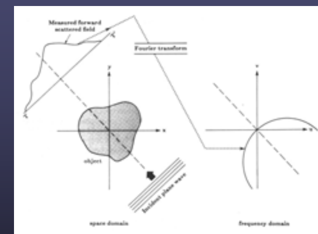

(from Kak)

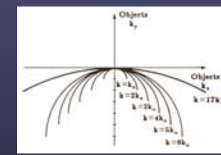

limit as $\lambda \rightarrow 0$ (relative to object size) is Fourier projection-slice theorem

- for weakly refractive media and coherent plane illumination
- if you record amplitude and phase of forward scattered field
- then the Fourier Diffraction Theorem says $\mathcal{F}$ {scattered field} = arc in $\mathcal{F}$ {object} as shown above, where radius of arc depends on wavelength λ
- repeat for multiple wavelengths, then take $\mathcal{F}^{-1}$ to create volume dataset
- equivalent to saying that a broadband hologram records 3D structure

[Devaney 2005]

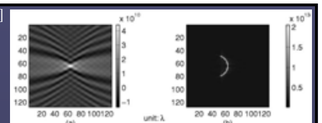

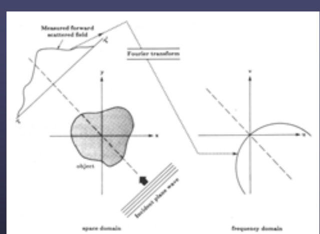

(from Kak)

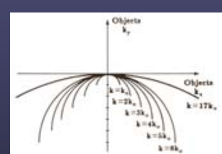

limit as $\lambda \rightarrow 0$ (relative to object size) is Fourier projection-slice theorem

- for weakly refractive media and coherent plane illumination
- if you record amplitude and phase of forward scattered field
- then the Fourier Diffraction Theorem says $\mathcal{F}$ {scattered field} = arc in $\mathcal{F}$ {object} as shown above, where radius of arc depends on wavelength λ
- repeat for multiple wavelengths, then take $\mathcal{F}^{-1}$ to create volume dataset
- equivalent to saying that a broadband hologram records 3D structure

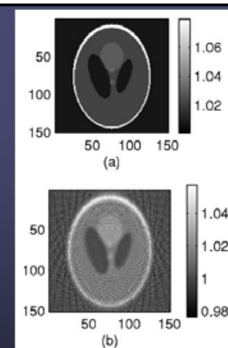

[Devaney 2005]

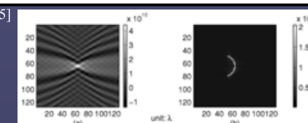

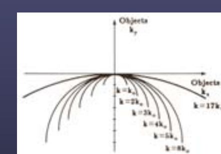

limit as $\lambda \rightarrow 0$ (relative to object size) is Fourier projection-slice theorem

- for weakly refractive media and coherent plane illumination
- if you record amplitude and phase of forward scattered field
- then the Fourier Diffraction Theorem says $\mathcal{F}$ {scattered field} = arc in $\mathcal{F}$ {object} as shown above, where radius of arc depends on wavelength λ
- repeat for multiple wavelengths, then take $\mathcal{F}^{-1}$ to create volume dataset
- equivalent to saying that a broadband hologram records 3D structure

Inversion by filtered backpropagation

backpropagation

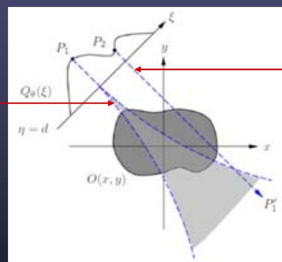

backprojection

[Jebali 2002]

- depth-variant filter, so more expensive than tomographic backprojection, also more expensive than Fourier method
- applications in medical imaging, geophysics, optics

Diffuse optical tomography (DOT)

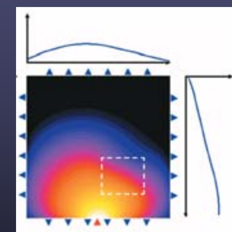

[Arridge 2003]

- assumes light propagation by multiple scattering
- model as diffusion process

Diffuse optical tomography

[Arridge 2003]

female breast with sources (red) and detectors (blue)

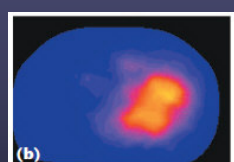

absorption (yellow is high)

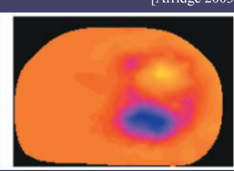

scattering (yellow is high)

- assumes light propagation by multiple scattering
- model as diffusion process
- inversion is non-linear and ill-posed
- solve using optimization with regularization (smoothing)

Computing vector light fields

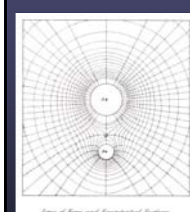

field theory (Maxwell 1873)

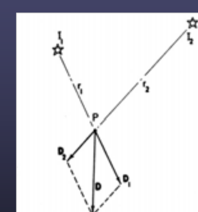

adding two light vectors (Gershun 1936)

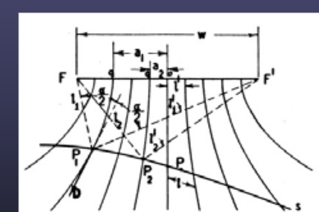

the vector light field produced by a luminous strip

Computing vector light fields

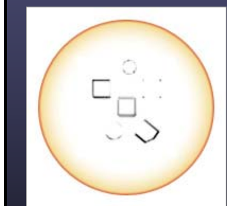

flatland scene with partially opaque blockers under uniform illumination

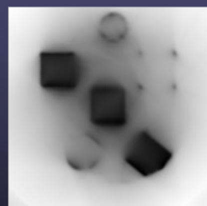

light field magnitude (a.k.a. irradiance)

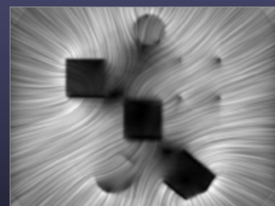

light field vector direction

From microscope light fields to volumes

- 4D light field → *digital refocusing* → 3D focal stack → *deconvolution microscopy* → 3D volume data

(DeltaVision)

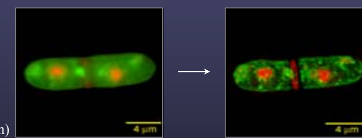

3D deconvolution

[McNally 1999]

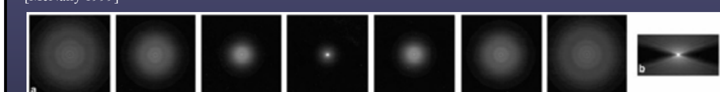

focus stack of a point in 3-space is the 3D PSF of that imaging system

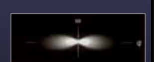

$\mathcal{F}$ {PSF}

- object * PSF → focus stack
- $\mathcal{F}$ {object} × $\mathcal{F}$ {PSF} → $\mathcal{F}$ {focus stack}
- $\mathcal{F}$ {focus stack} ÷ $\mathcal{F}$ {PSF} → $\mathcal{F}$ {object}
- spectrum contains zeros, due to missing rays
- imaging noise is amplified by division by ~zeros
- reduce by regularization (smoothing) or completion of spectrum
- improve convergence using constraints, e.g. object > 0

Silkworm mouth
(40x / 1.3NA oil immersion)

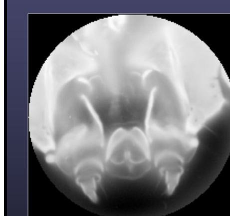

slice of focal stack

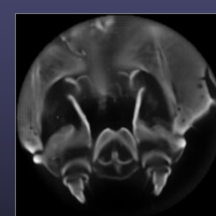

slice of volume

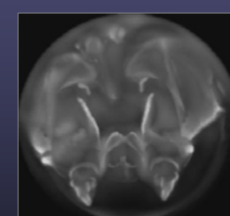

volume rendering

From microscope light fields to volumes

- 4D light field → *digital refocusing* → 3D focal stack → *deconvolution microscopy* → 3D volume data

(DeltaVision)

- 4D light field → *tomographic reconstruction* → 3D volume data

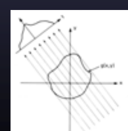

(from Kak)

Optical Projection Tomography (OPT)

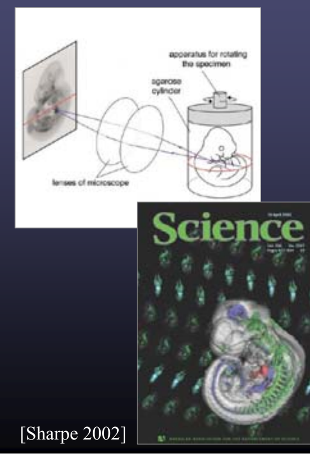

[Sharpe 2002]

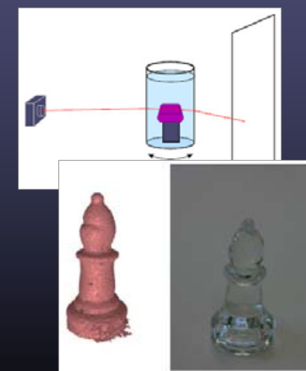

[Trifonov 2006]

Confocal scanning microscopy

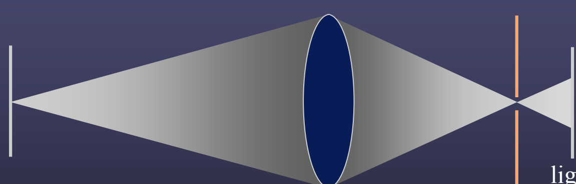

Confocal scanning microscopy

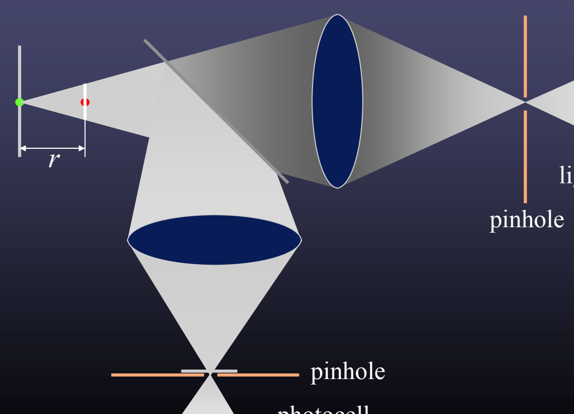

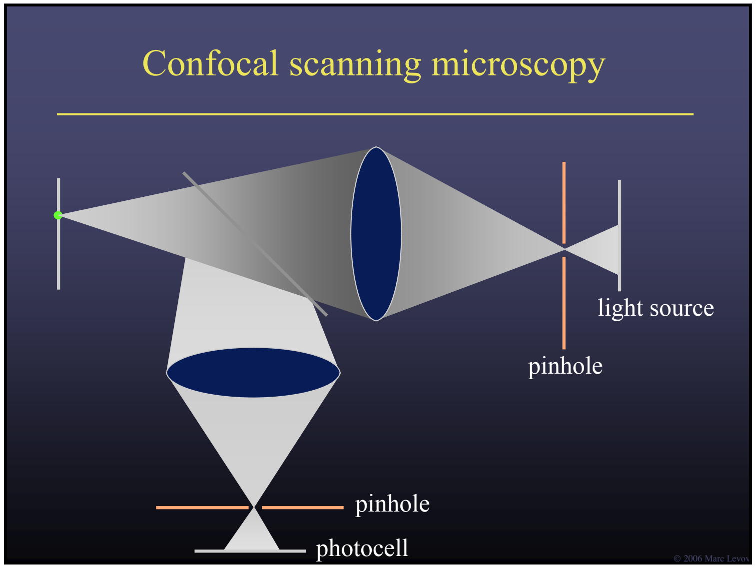
Confocal scanning microscopy
light source
pinhole
pinhole
photocell

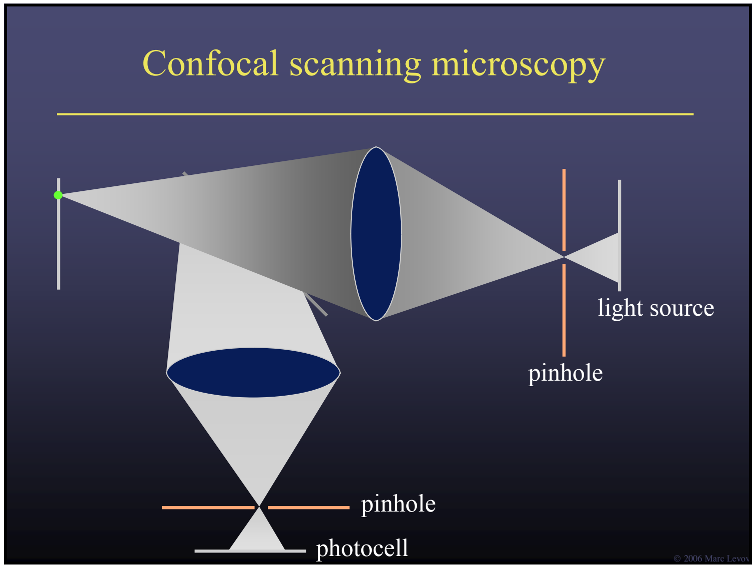
Confocal scanning microscopy
light source
pinhole
pinhole
photocell

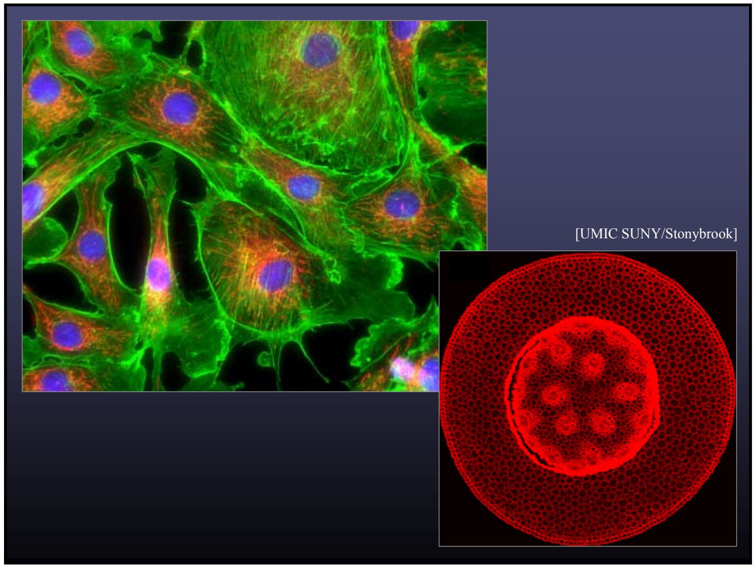
[UMIC SUNY/Stonybrook]

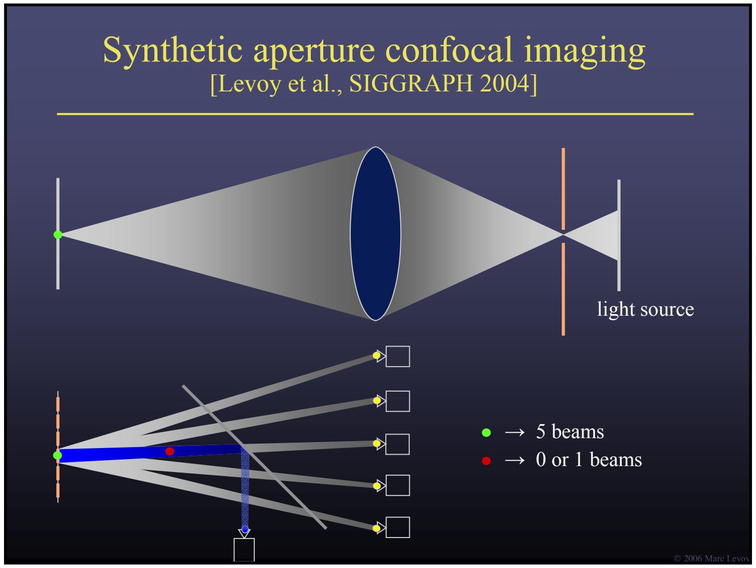
Synthetic aperture confocal imaging
[Levoy et al., SIGGRAPH 2004]
light source
→ 5 beams
→ 0 or 1 beams

Seeing through turbid water

Seeing through turbid water

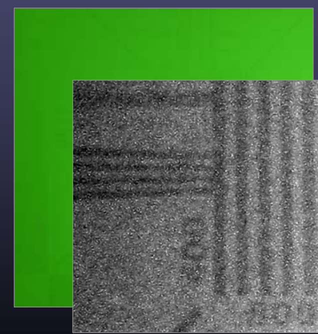

floodlit

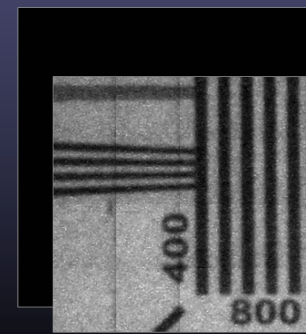

scanned tile

Coded aperture imaging

(from Zand)

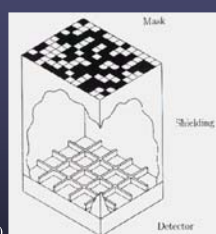

- optics cannot bend X-rays, so they cannot be focused
- pinhole imaging needs no optics, but collects too little light
- use multiple pinholes and a single sensor
- produces superimposed shifted copies of source

Reconstruction by backprojection

(from Zand)

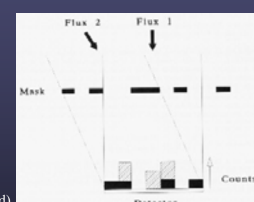

- backproject each detected pixel through each hole in mask
- superimposition of projections reconstructs source + a bias
- essentially a cross correlation of detected image with mask
- also works for non-infinite sources; use voxel grid
- assumes non-occluding source

Example using 2D images
(Paul Carlisle)

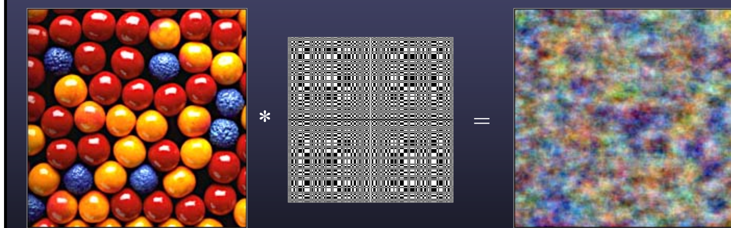

New sources for point data

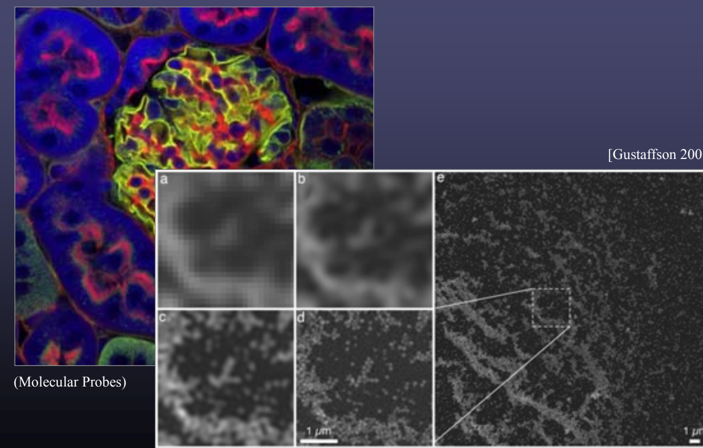

[Gustaffson 2005]

(Molecular Probes)

Three theses

<u>Thesis #1:</u> Many sciences lack good visualization tools.
Corollary: These are a good source for volume and point data.

<u>Thesis #2:</u> Computer scientists need to learn these sciences.
Corollary: Learning the science may lead to new visualizations.

<u>Thesis #3:</u> We also need to learn their data capture technologies.
Corollary: Visualizing the data capture process helps debug it.

The best visualizations are often created by domain scientists

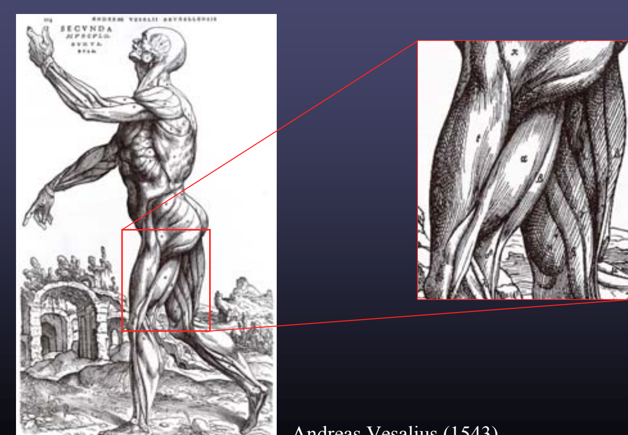

Andreas Vesalius (1543)

Three theses

<u>Thesis #1:</u> Many sciences lack good visualization tools.
Corollary: These are a good source for volume and point data.

<u>Thesis #2:</u> Computer scientists need to learn these sciences.
Corollary: Learning the science may lead to new visualizations.

<u>Thesis #3:</u> We also need to learn their data capture technologies.
Corollary: Visualizing the data capture process helps debug it.

Visualizing raw data helps debug the capture process

hollow fluorescent 15-micron sphere, manually captured Z-stack, 1-micron increments, 40×/1.3NA oil objective

X-Z cross-sectional slice of same stack

...or may force improvements in the capture technology

Shinya Inoué at his polarization microscope

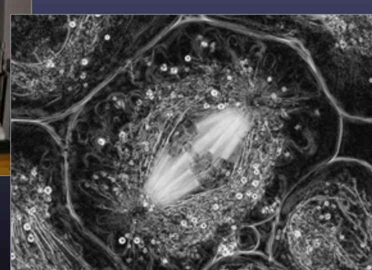

crane fly spermatocyte undergoing meiosis, image and video by Rudolf Oldenbourg

Final thought: the importance of building <u>useful</u> tools

"A toolmaker succeeds as, and only as, the users of his tool succeed with his aid. However shining the blade, however jeweled the hilt, however perfect the heft, a sword is tested only by cutting. That swordsmith is successful whose clients die of old age."

— Fred Brooks,
Computer Scientist as Toolsmith – II,
Proc. CACM 1996

Acknowledgements

- Fred Brooks (“Computer Scientist as Toolsmith”)
- Pat Hanrahan (“Self-Illustrating Phenomena”)
- Bill Lorensen (“The Death of Visualization”)
- Shinya Inoué (“History of Polarization Microscopy”)