Light Fields

Ivo Ihrke

Outline

- plenoptic function
- subsets of the plenoptic function
- light field:
  - concept
  - view synthesis
  - parametrization
- refresher:
  - signal processing
  - sampling
  - reconstruction
  - aliasing
- light field sampling analysis
- light field acquisition
- applications of light fields

Plenoptic Function

- plenoptic (latin *plenus:* full , *optic*: vision)
- plenoptic function [Adelson91] describes the radiance at
  - a position in space (3D)
  - in a certain direction (2D)
  - at a particular point in time (1D)
  - in a particular wavelength (1D)

  $L = P(x, y, z, \theta, \varphi, t, \lambda)$

  is a 7D function
- imagine a collection of dynamic environment maps covering the whole space

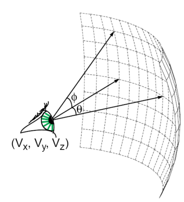

Grayscale snapshot

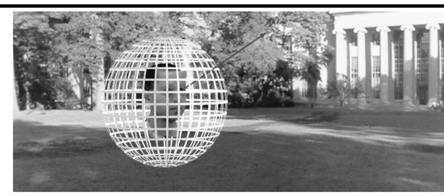

$P(\theta,\phi)$

is intensity of light

- Seen from a single view point
- At a single time
- Averaged over the wavelengths of the visible spectrum

(can also do *P(x,y),* but spherical coordinate are nicer)

Color snapshot

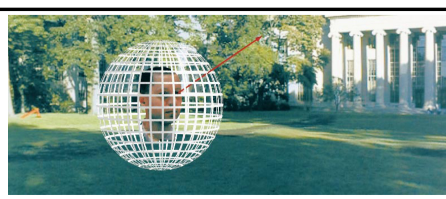

$P(\theta,\phi,\lambda)$

is intensity of light

- Seen from a single view point
- At a single time
- As a function of wavelength

A movie

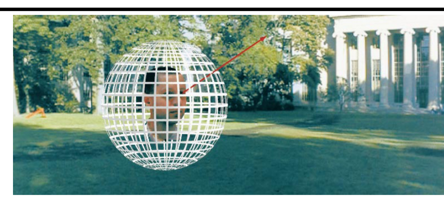

$P(\theta,\phi,\lambda,t)$

is intensity of light

- Seen from a single view point
- Over time
- As a function of wavelength

Holographic movie

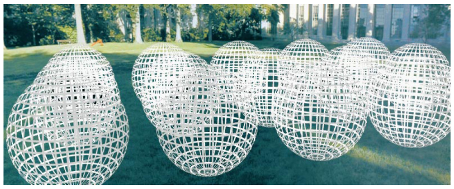

$$P(\theta, \phi, \lambda, t, V_X, V_Y, V_Z)$$

is intensity of light

- Seen from ANY viewpoint
- Over time
- As a function of wavelength

Plenoptic Function

- describes everything that can possibly be seen
  (and much more)
  - e.g. for example wavelength includes all electromagnetic radiation (not necessarily visible by human observer)
  - non-physical effects are covered
  - also time-varying and wave length-shifting effects like phosphorescence, etc.
- plenoptic function is unknown, what use does it have ?
  - conceptual tool to group imaging systems according to greater flexibility in view manipulation

Plenoptic Function

- imaging concepts using sub-sets of the plenoptic function
  - conventional photograph (2D sub-set of θ, φ)
  - panorama [Chen95] (2D – full range of θ, φ)
  - video sequence (3D sub-set of x, y, z, θ, φ, t)
  - light field [Levoy96, Gortler96]
    (4D sub-set of x, y, z, θ, φ)
  - dynamic light fields [Wilburn05]
    (5D sub-set of x, y, z, θ, φ, t)
- wavelength is usually discretely sampled in R,G,B
- in real imaging systems resulting radiance is limited in range
  - LDR for conventional cameras
  - HDR

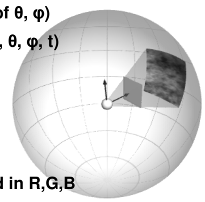

Plenoptic Function

- Drawbacks:
  - many scene parameters molded into time parameter
  - e.g.
    - dynamic scenes
    - illumination changes
    - light material interaction
- therefore: difficult to edit
- alternatives (not in this lecture):
  - plenoptic illumination function [Wong02]
  - reflectance fields [Debevec00]

Light Fields

- [McMillan95] use sampled 5D function (x, y, z, θ, φ) on a regular grid
- interpolate to generate new views

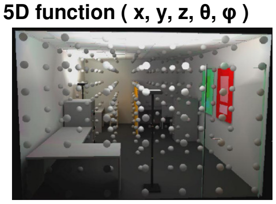

- light fields are only 4D
  - free space assumption
  - radiance is constant along a ray

Light Fields

space with occluders – 5D

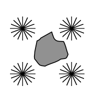

free space, radiance stays constant along the ray – 4D

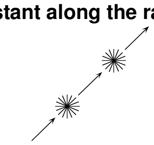

outside – in viewing

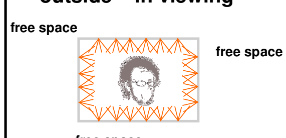

inside – out viewing

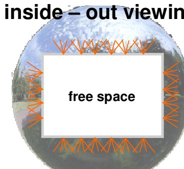

Light Fields – Principle of View Synthesis

- re-arrange ray samples to generate new views

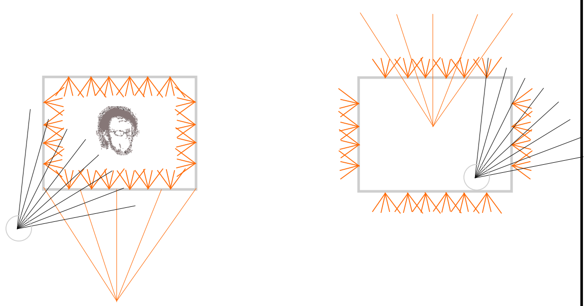

Light Fields - Properties

- Advantages
  - rendering complexity is independent of scene complexity
  - display algorithms are fast
  - complex view-dependent effects are simple
  - (no mathematical model required)
- Disadvantages
  - high storage requirements

    (although high correlation between images yields high compression ratios ~120:1 [Levoy96])
  - difficult to edit (no model)

Light Fields - Parametrization

- need a way to parametrize rays in space for simple sampling and retrieval
- should be adapted to sensor geometry
- new view synthesis should be fast
- Let's consider some candidate parametrizations

Light Fields - Parametrizations

- point on plane + direction L (u, v, θ, φ)
  - mixture between cartesian and trigonometric parameters
  - inefficient to evaluate
  - non-uniform sampling
  - directional interpolation difficult

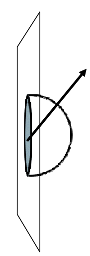

- alternatively arbitrary surface + direction,
  - should be convex to avoid duplicates

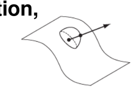

Light Fields - Parametrizations

- two points on sphere [Camahort98]
  - uniform sampling
  - needs a uniform subdivision of sphere into patches
  - needs a way to sample single rays
    - difficult for real scenes

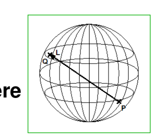

L(θ_1, φ_1, θ_2, φ_2)

- great circle + point on disk [Camahort98]
  - uniform sampling
  - needs orthographic projections to disk
    - less difficult than 2PS parametrization

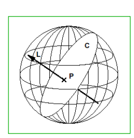

L(u, v, θ, φ)

Light Fields - Parametrizations

- two plane parametrization (light slab) [Levoy96]

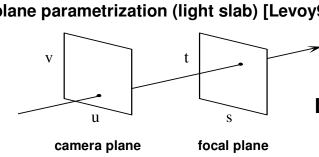

L (u, v, s, t)

- fast display algorithms (projective geometry)
- simple interpretation (array of images)
- most commonly used parametrization
- Drawback: only in one major direction
  - covering 360º requires at least 6 light slabs [Gortler96]
  - switching from one slab to the next introduces artifacts a.k.a. disparity problem

Light Fields - Parametrizations

- light field generation with two-plane parametrization

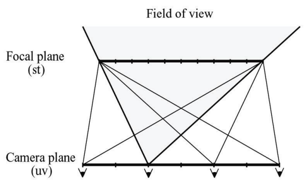

- off-axis perspective projections
- normal camera images need (simple) re-sampling

Light Fields – Parametrizations

- a two-plane parametrized light field is basically a collection of images

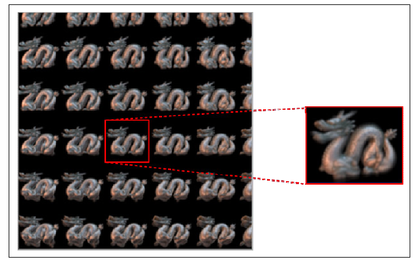

Light Fields - Parametrizations

- view generation from two-plane parametrization

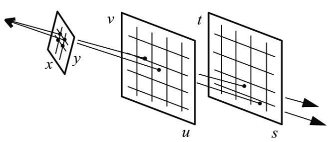

- at an observer position
  - project (u, v) and (s, t) parameter planes into virtual view (x, y)
  - for each pixel in virtual view use projected (u, v, s, t) to look up radiance L (u, v, s, t)
  - two perspective projections and one look-up determine virtual view → efficient rendering

Light Fields – Rendering 2D

involved samples

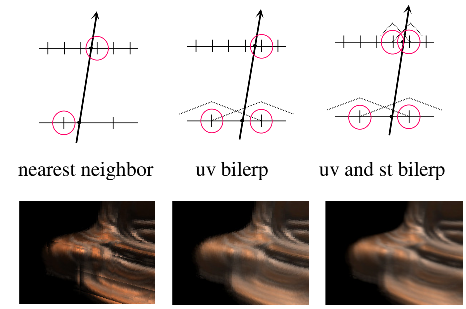

Light Field Rendering - Examples

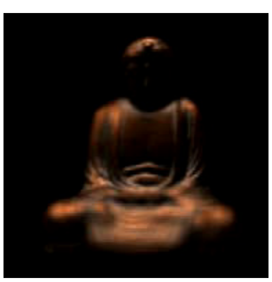

16x16 images

1 slab

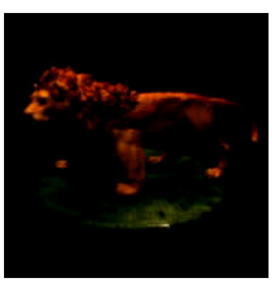

32 x 16 images

4 slabs

Depth Assisted Light Fields [Gortler96]

without depth knowledge

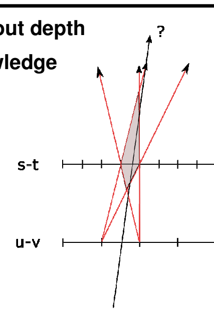

with depth knowledge

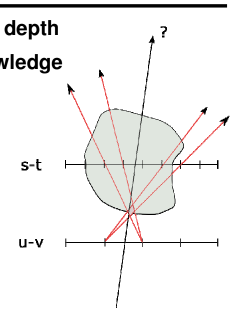

different pixels have to be interpolated !

Depth Assisted Light Fields

- Regions of uncertainty, depending on depth
  - closer objects have higher disparity
  - standard light field look-up as described previously yields poor results
  - need depth assisted warping
  - e.g. projective texture mapping [Debevec96]

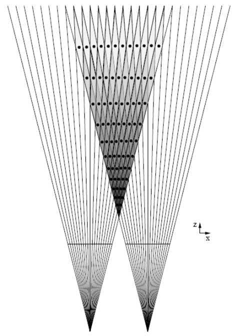

Depth Assisted Light Fields -Example

recorded images

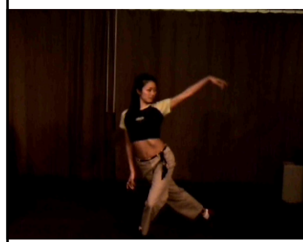

depth assisted view warping

Image-based vs. Model-based Rendering

- trade-off between image-based and model-based rendering approaches

Images Only

Mathematical Descriptions

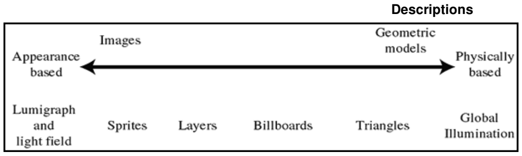

more data
less computation

less data
more computation

- Is there a way to find a good trade-off ?
- need some signal processing for analysis

Signal Processing Basics – Sampling

continuous function f

sampling function s (comb function/ shah function)

sampled function f_s

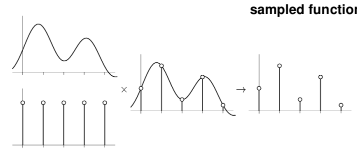

Sampling in the Space Domain

spatial domain: $f \cdot s = f_s$

Signal Processing Basics – Sampling

frequency spectrum of continuous function: F

frequency spectrum of sampling function: S

frequency spectrum of sampled function: F_s

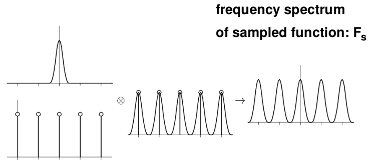

Sampling in the Frequency Domain

- convolution theorem

→ in frequency domain: $F \otimes S = F_s$

- frequency spectrum of original function is copied multiple times !

Signal Processing Basics – (Pre-)Aliasing

- if continuous function f has to high frequency content, aliasing occurs
- undersampling
- overlap of the copies in the fourier spectrum F_s

frequency domain:

F

S

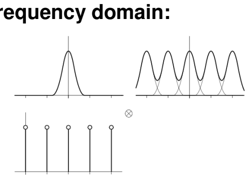

F_s

Signal Processing Basics - Reconstruction

- reconstruction in frequency domain is simple
  - get rid of the copies of the frequency spectrum
  - multiplication by box function
    (reconstruction filter)

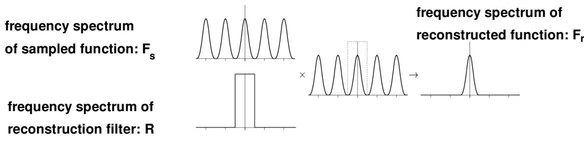

Reconstruction in the Frequency Domain

reconstruction in frequency domain: $F_s \cdot R = F_r$

Signal Processing Basics - Reconstruction

- reconstruction in spatial domain is a convolution
  convolution theorem: $f_s \otimes r = f_r$

- Fourier transform of box function is sinc function
  - sinc function has infinite support
  - → need all samples to reconstruct properly

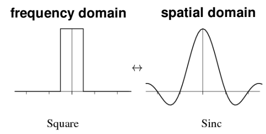

Signal Processing Basics - Reconstruction

- common reconstruction filters and their frequency spectrum

nearest neighbor

linear

cubic

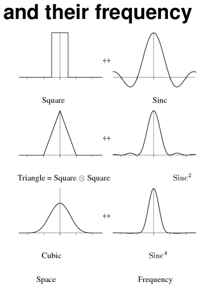

better approximation of sinc function in space yields smaller support in frequency domain

Signal Processing Basics – (Post-)Aliasing

- poor choice of reconstruction filter results in aliasing
- information from copies of the fourier spectrum gets included in the reconstruction
- example: nearest neighbor filter

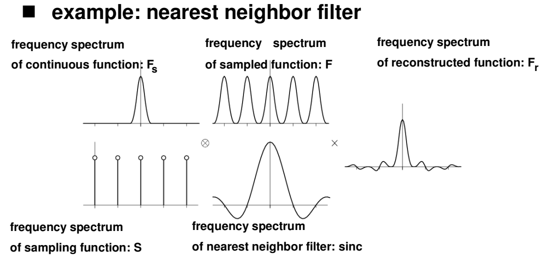

Signal Processing Basics – Sampling Rate

- increasing the sampling rate results in larger spacing of the frequency spectrum of the sampled function

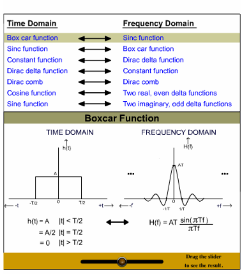

Signal Processing Basics - Lessons

- need band-limited function f
  - e.g. low-pass filtered (narrows support in frequency domain)
  - frequency spectrum with local support
- sampling rate must be sufficient (increases overlap-free area in the frequency domain)
- reconstruction filter should have local support in frequency domain to avoid post-aliasing

Plenoptic Sampling [Chai00]

- apply Fourier Analysis to light field rendering
- simplifying assumptions:
  - no occlusion
  - lambertian reflectance
- perform analysis in 2D
  - one spatial dimension
  - one directional dimension
- full 4D case analogous

Computational Photography   Hendrik Lensch, Summer 2007

Plenoptic Sampling – Epipolar Plane Images

- analyze epipolar plane image (EPI) and its frequency spectrum
  - main result: frequency spectrum of a light field is bounded by minimum and maximum scene depth
- EPI is a slice of the light field, e.g. (v, t)

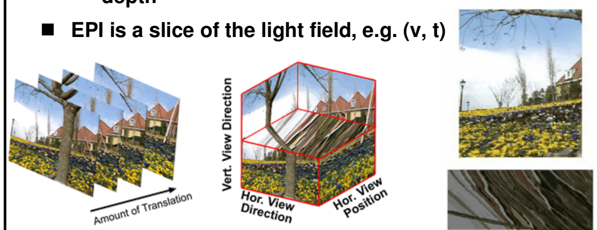

Computational Photography   Hendrik Lensch, Summer 2007

Multi-Dimensional Sampling

- 2D function → sampling function is "bed-of-nails" instead of comb function
- copies are spread in two dimensions

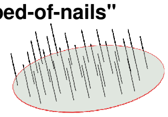

"continuous" image

sampled image

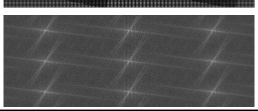

frequency spectrum with overlapping duplicates in two dimensions

Hendrik Lensch, Summer 2007

Plenoptic Sampling – Analysis Goal

- Analyze EPI images in frequency domain
- GOAL:
  - find minimum spacing of the copies of the frequency spectra without overlap
    → minimum sampling rate for anti-aliased rendering
  - analyze influence of known depth

Computational Photography   Hendrik Lensch, Summer 2007

Plenoptic Sampling - Analysis

- start with simple scene: textured plane parallel to camera translation plane

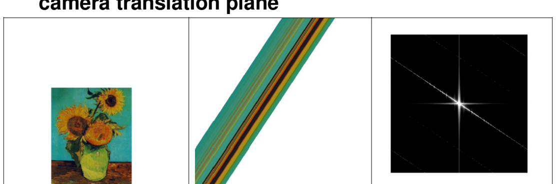

(a1) Scene image | (a2) EPI | (a3) Fourier transform of EPI

- ignore horizontal and vertical lines (artifacts of non-periodic nature of the function)
- frequency spectrum for parallel plane (constant depth) is a line !

Computational Photography   Hendrik Lensch, Summer 2007

Plenoptic Sampling - Analysis

- more complex: scene with two planes parallel to camera translation direction (two constant depths)

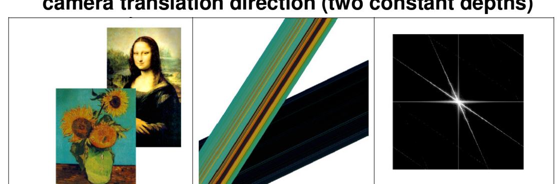

(b1)Scene image | (b2) EPI | (b3) Fourier transform of EPI

- a second line appears in the fourier spectrum !
- different slope
- What happens with non-parallel planes ?

Computational Photography   Hendrik Lensch, Summer 2007

Plenoptic Sampling - Analysis

- two parallel and one tilted plane

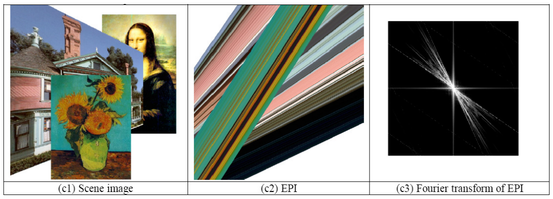

(c1) Scene image (c2) EPI (c3) Fourier transform of EPI

- frequency spectrum is still contained between the two lines corresponding to minimum and maximum depth !
- Is this also true for non-planar objects ?

Plenoptic Sampling - Analysis

- two parallel planes and a surface with complex depth variation

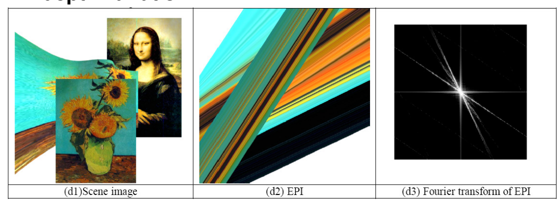

(d1)Scene image (d2) EPI (d3) Fourier transform of EPI

- the frequency spectrum is still bounded !
- boundedness (i.e. local support) is good, we can find a sampling rate that causes no overlap of the frequency spectra

Plenoptic Sampling – Sampling Rate

- determining the sampling rate

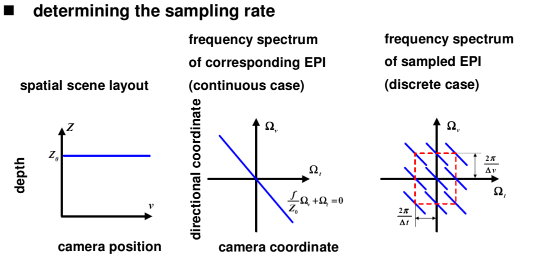

- Δt, Δv are the sampling rates on the camera and focal planes of the light field
- choose such that there is no overlap in the frequency domain

Plenoptic Sampling – Reconstruction Filter

- determining a suitable (depth assisted) reconstruction filter allows for denser packing in the frequency domain
  - → coarser sampling in the spatial domain
- depth assisted reconstruction filter is a tilted box filter

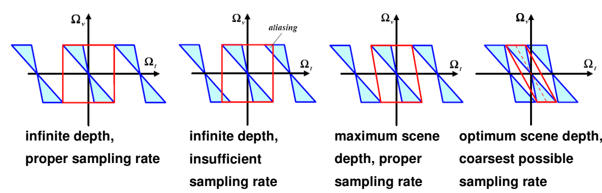

Plenoptic Sampling – Reconstruction Filter

- Is it possible to get away with even coarser sampling ?
- Yes, Fourier Transform is linear
  - → can decompose spectrum into a sum of spectra with different optimal depth reconstruction filters

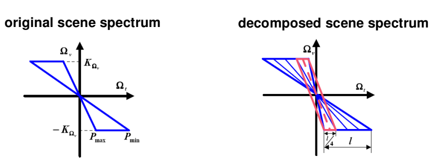

- multiple depth layers allow for denser packing in frequency space → coarser sampling in spatial domain

Plenoptic Sampling – Images and Geometry

- different combinations of number of images and number of available depth layers yield the same rendering result

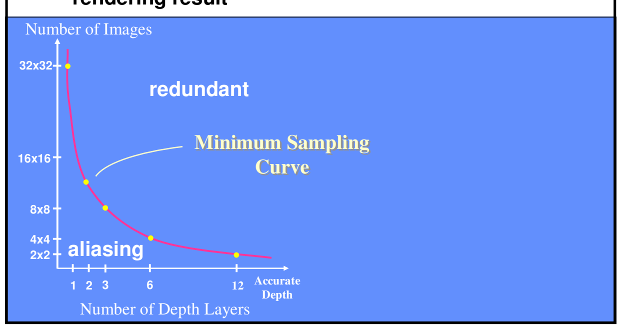

Plenoptic Sampling – Images and Geometry

- Example: 3 depth layers, 32x32 images, oversampled

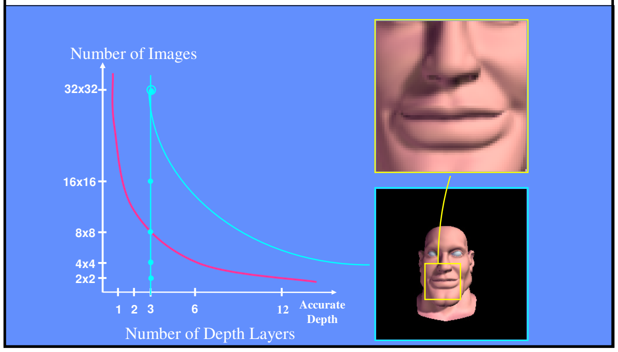

Plenoptic Sampling – Images and Geometry

- Example: 3 depth layers, 16x16 images, oversampled

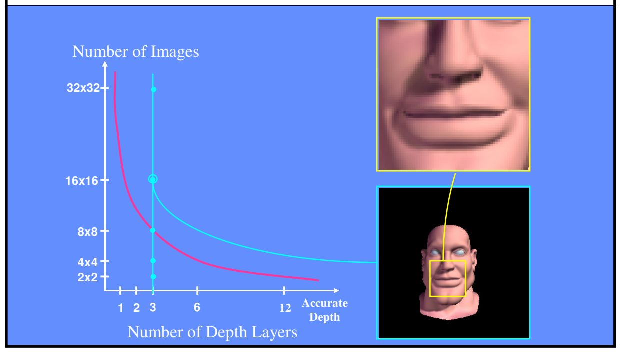

Plenoptic Sampling – Images and Geometry

- Example: 3 depth layers, 8x8 images, optimally sampled

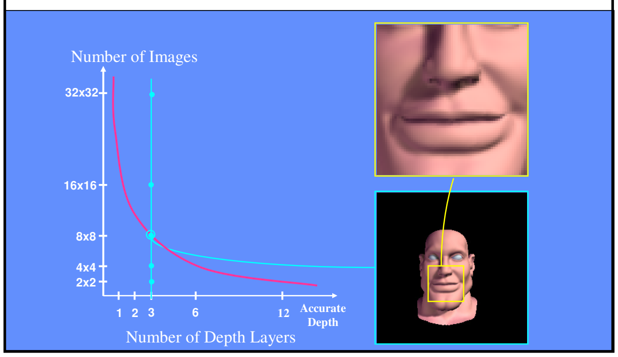

Plenoptic Sampling – Images and Geometry

- Example: 3 depth layers, 4x4 images, undersampled

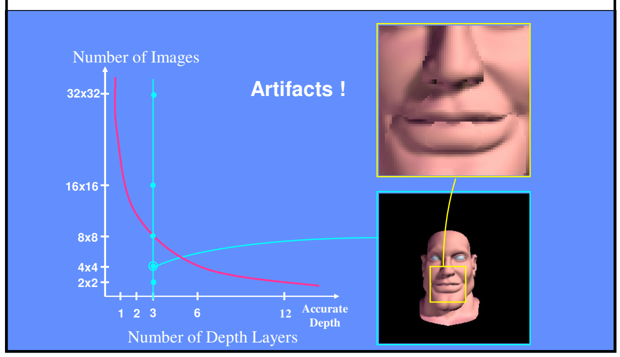

Plenoptic Sampling – Images and Geometry

- Example: 3 depth layers, 2x2 images, undersampled

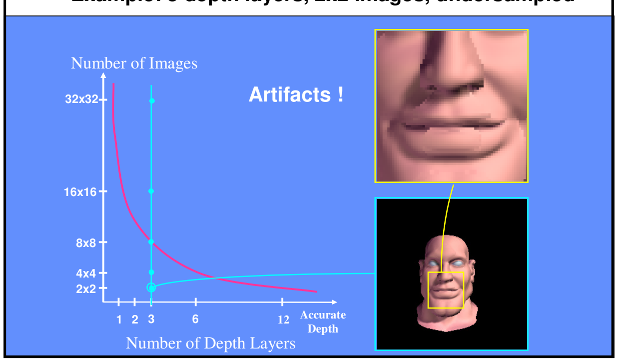

Plenoptic Sampling – Lessons

- Trade-off between geometric knowledge and image-based information
  - quantified by Minimum Sampling Curve
- for higher output resolution, Minimum Sampling Curve has to be shifted into redundant area
- this is also necessary for non-lambertian surfaces and heavy occlusion [Zhang03]

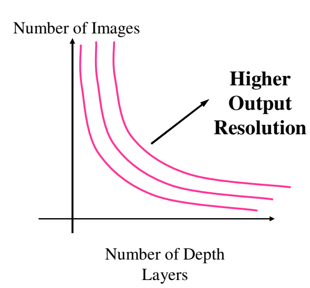

Light Fields – Acquisition & Applications

- spherical gantry [Levoy96]
  - 2 degrees of freedom for camera
  - 2 degrees of freedom for lamp
- Application:
  - acquisition of 360º light fields and reflectance fields

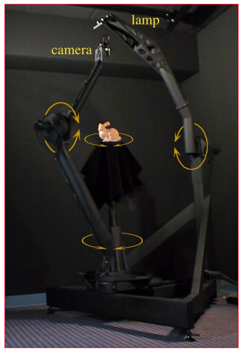

Light Fields – Acquisition & Applications

- Hand-Held Video Camera [Koch99]
  - structure-from-motion
  - 3D reconstruction (depth maps)

Light Fields – Acquisition & Applications

- Multi-Camera Array [Wilburn05]
  - 128 cameras (Stanford)
- Application:
  - dynamic light field acquisition
  - synthetic aperture imaging
  - spatio-temporal interpolation
  - HDR light field imaging

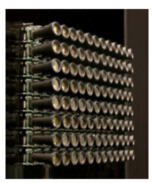

Light Fields – Acquisition & Applications

- Plenoptic Camera [Ng05]
  - conventional lens + microlens array
  - 4000x4000 pixels
  - 129x129 microlenses
  - =14x14 pixels per microlens
- Applications:
  - viewpoint shifts
  - perspective changes
  - digital refocusing

Kodak 16-megapixel sensor

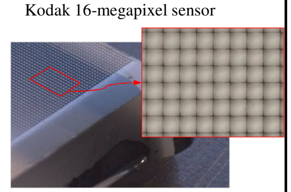

125µ square-sided microlenses

Light Fields – Acquisition & Applications

- principle of plenoptic camera

conventional camera

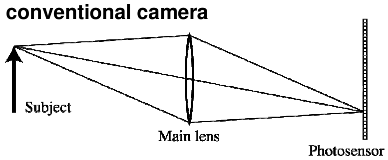

plenoptic camera

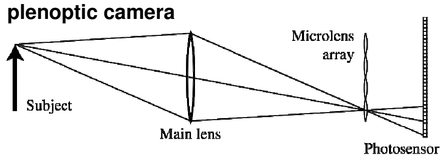

Light Fields – Acquisition & Applications

- refocusing example
- only one photograph taken
- refocus is performed computationally by light field manipulation

Light Fields – Acquisition & Applications

- light field microscope [Levoy06]
  - similar to plenoptic camera
  - occular + microlens array
- Applications:
  - perspective views
  - 3D reconstruction
  - refocusing

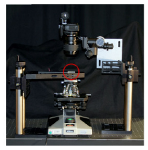

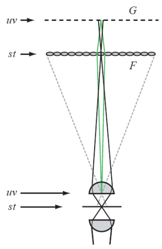

End

References

- [Adelson91] E. H. Adelson, J. R. Bergen, "The Plenoptic Function and the Elements of Early Vision", Computation Models of Visual Processing, 1991
- [Bracewell00] R. N. Bracewell, "The Fourier Transform and Its Applications", 3rd ed., McGraw-Hill 2000
- [Buehler01] C. Buehler, M. Bosse, L. McMillan, S. Gortler, M. Cohen, "Unstructured Lumigraph Rendering", SIGGRAPH 2001, pp. 425-432
- [Chai00] J.-X. Chai, X. Tong, S.-C. Chan, H.-Y. Shum, "Plenoptic Sampling", SIGGRAPH 2000, pp. 307-318
- [Camahort98] E. Camahort, A. Lerios, D. Fussell, "Uniformly Sampled Light Fields", EGWR 1998, pp. 117-130
- [Chen93] S. E. Chen, L. Williams, "View Interpolation for Image Synthesis", SIGGRAPH 1993, pp 279-288

References

- [Chen95] S. E. Chen, "Quicktime VR – An Image-Based Approach to Virtual Environment Navigation", SIGGRAPH 1995, pp. 29-38
- [Debevec96] P. E. Debevec, C. J. Taylor, J. Malik." Modeling and Rendering Architecture from Photographs", *SIGGRAPH 1996*, pp. 11-20
- [Debevec00] P. E. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, M. Sagar, "Acquiring the Reflectance Field of a Human Face", SIGGRAPH 2000, pp. 145-156
- [Gortler96] S. J. Gortler, R. Grzeszczuk, R. Szeliski, M. Cohen, "The Lumigraph", *SIGGRAPH 1996,* pp. 43-54.
- [Jones07] A. Jones, I. McDowall, H.Yamada, M. Bolas, P. Debevec, "Rendering for an Interactive 360º Light Field Display", SIGGRAPH 2007, to appear
- [Koch99] R. Koch, M. Pollefeys, B. Heigl, L. Van Gool, H. Niemann, "Calibration of Hand-held Camera Sequences for Plenoptic Modeling", ICCV 1999, pp. 585-591

References

- [Levoy96] M. Levoy, P. Hanrahan, "Light Field Rendering", SIGGRAPH 1996, pp. 31-42
- [Levoy04] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, M. Bolas, "Synthetic Aperture Confocal Imaging", SIGGRAPH 2004, pp. 825-834
- [Levoy06] M. Levoy, R. Ng, A. Adams, M. Footer, M. Horowitz, "Light Field Microscopy", SIGGRAPH 2006, pp. 924-934
- [McMillan95] L. McMillan, G. Bishop, "Plenoptic Modeling: An Image-Based Rendering System", SIGGRAPH 1995, pp. 39-46
- [Ng05a] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, P. Hanrahan, "Light Field Photography with a Hand-held Plenoptic Camera", Stanford Tech. Report CTSR 2005-02
- [Ng05b] R. Ng, "Fourier Slice Photography", SIGGRAPH 2005, pp. 735-744

References

- [Wilburn05] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, M. Levoy, "High Performance Imaging using Large Camera Arrays", SIGGRAPH 2005, pp. 765-776
- [Wong02] T.-T. Wong, C.-W. Fu, P.-A. Heng, C.-S. Leung, "The Plenoptic Illumination Function", IEEE Transactions on Multimedia 4(3), Sep. 2002, pp. 361-371
- [Zhang03] C. Zhang, T. Chen "Spectral Analysis for Sampling Image-Based Rendering Data", IEEE CSVT 2003, pp. 1038-1050

Acknowledgements

- Some slides by Marc Levoy, Alexei Efros
- Images for Signal Processing Basics by
  Pat Hanrahan
  lecture notes at
  http://www.cs.wisc.edu/graphics/Courses/559-s2001/notes/hanrahan.pdf
- Images from referenced papers