Visualization and Virtual Reality in Medicine

Prof. Dr.-Ing. Bernhard Preim, Dr.-Ing. Steffen Oeltze-Jafra
Lehrstuhl für Visualisierung
Universität Magdeburg

Tutorial Syllabus

Surface Visualization - Marching Cubes and its improvements - Smoothing of surface visualizations	(40 min.)
3D Vessel Visualization	(30 min.)
Labeling Medical Visualizations	(20 min.)
Break	(15 min.)
Direct Volume Visualization - Ray casting and texture-based approaches - Projection methods	(40 min.)
Multifield Medical Visualization	(30 min.)
Virtual Endoscopy	(20 min.)

Virtual Reality and Visualization

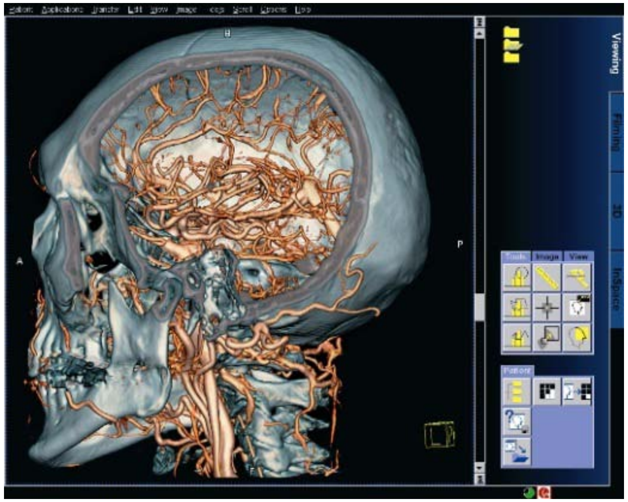

[Siemens Medical]

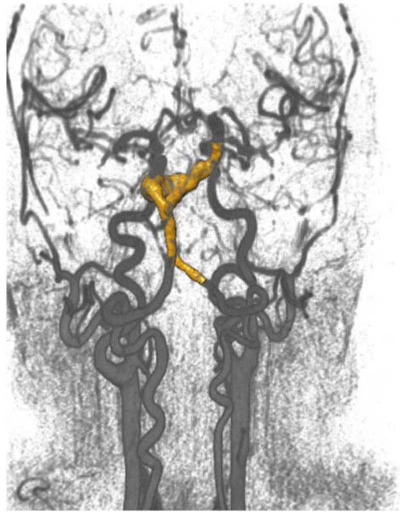

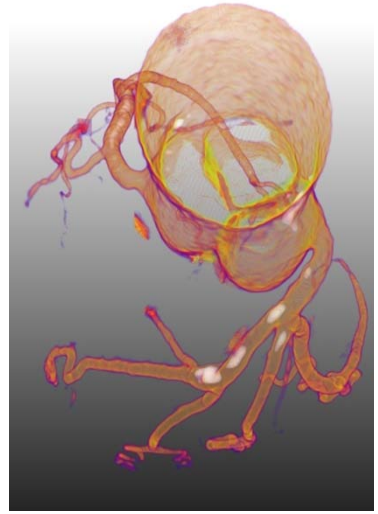

Volume visualization

[Vital Images]

Direct Volume Visualization: Outline

Direct Volume Visualization

- Introduction
- Image-based Volume Visualization
- Projection Methods
- Lightning
- Tagged Volume Rendering

Introduction

- Surface rendering is based on a binary decision which voxels should contribute to the visualization
- Direct volume rendering employs opacity transfer functions effectively leading to semi-transparent renderings
- Surface rendering is appropriate if there are clear boundaries, such as for skeletal structures
- Direct volume rendering is favorable in case of weak boundaries

Medical Data Sets

Examples:

- Inner ear with HR-CT:
  - 512x512x64,
    thickness: 1 mm,
    slice dist.: 0.5 mm
    resolution: 0.12 mm
- Intracranial vessels, CTA:
  - 512x512x856,
    resolution: 1 mm,
    thickness: 1 mm

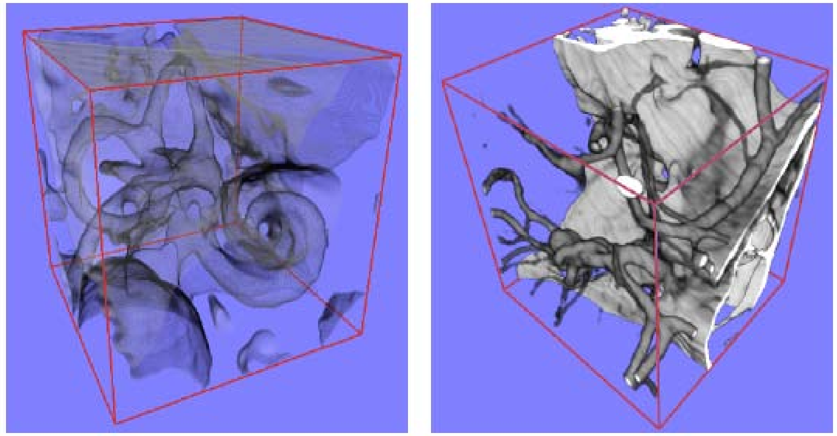

CT Inner Ear Detail
1 MB (128 x 128 x 64)

CTA Aneurysma Detail
2 MB (128 x 128 x 128)

Source: Rezk-Salama, 2002

Typical: many transparently or semi-transparently illustrated voxels

Specification: by an appropriate TF

Direct Volume Rendering: Introduction

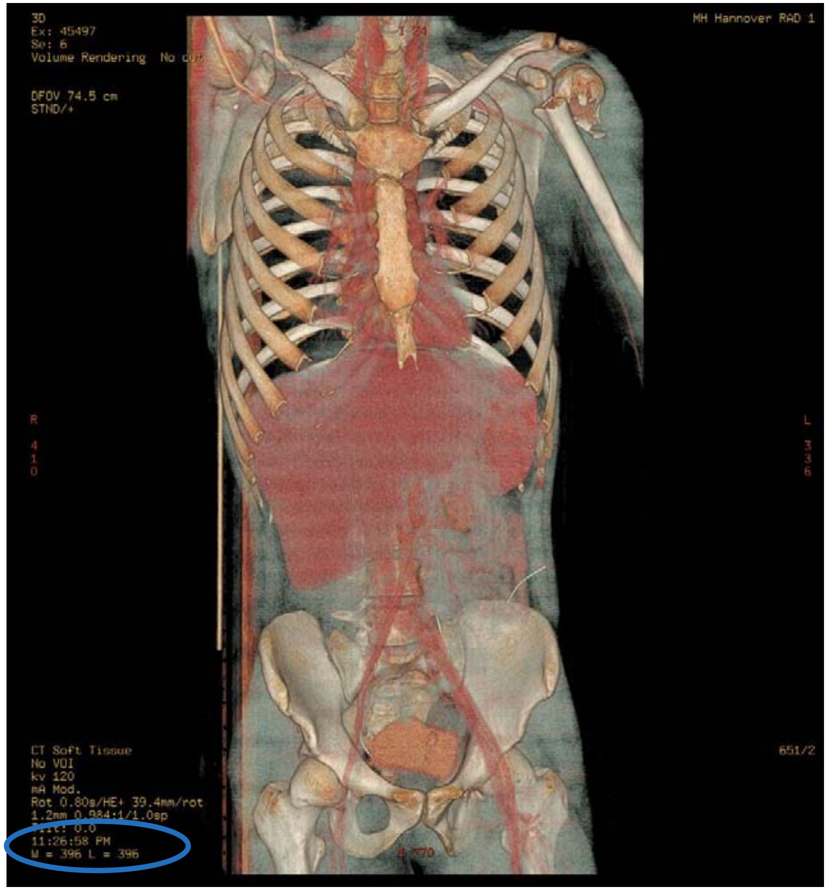

Modern volume rendering as part of clinical workstation
(Courtesy of Hoen-Oh Shin, Medical School Hannover, 2013)

Direct Volume Rendering: Introduction

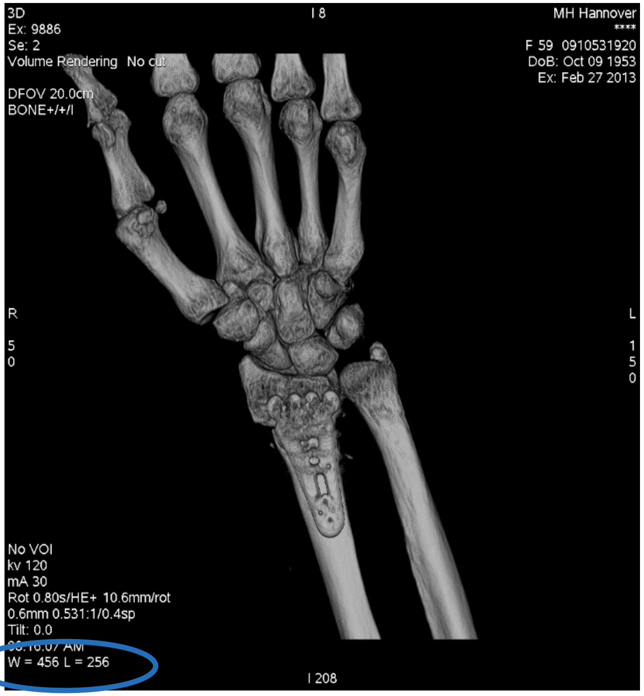

Modern volume rendering as part of clinical workstation

(Courtesy of Hoen-Oh Shin, Medical School Hannover, 2013)

Direct Volume Rendering: Introduction

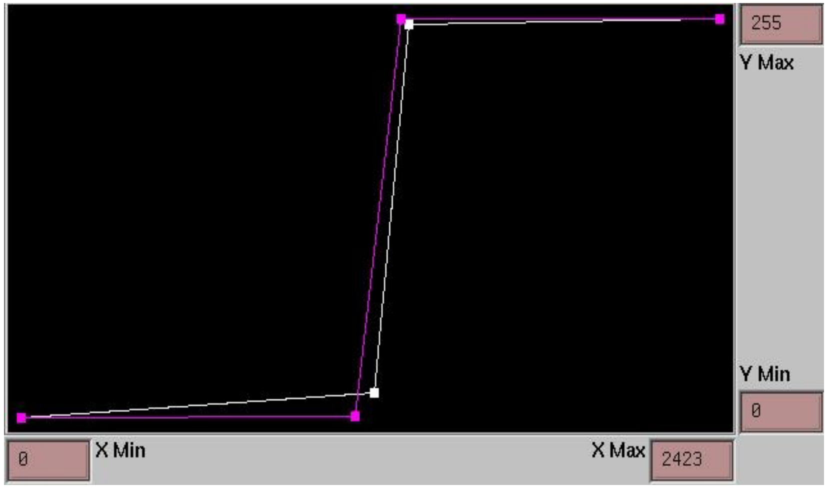

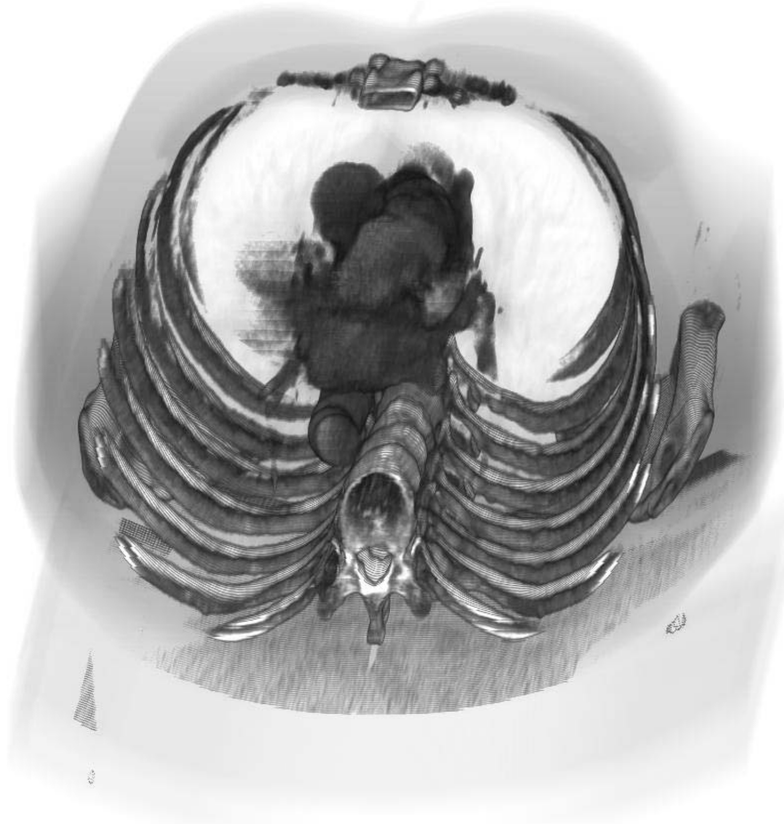

Transfer function and the resulting volume rendering of CT-thorax data

Direct Volume Rendering: Introduction

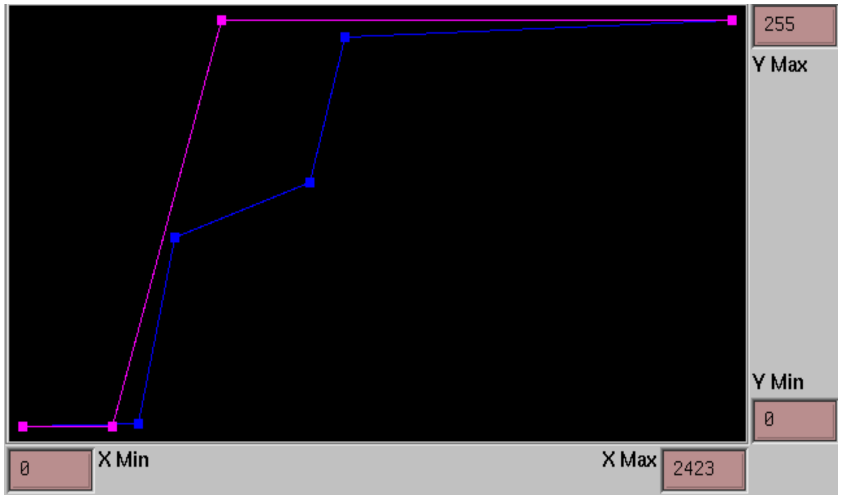

Different TF and the resulting volume rendering of (the same) CT-thorax data.
Simple piecewise linear 1D TFs. In practice, often presets are used.

- Setting of TFs for gray values and transparency (often linear functions).
- Histogram depicted as context in a graphical editor.

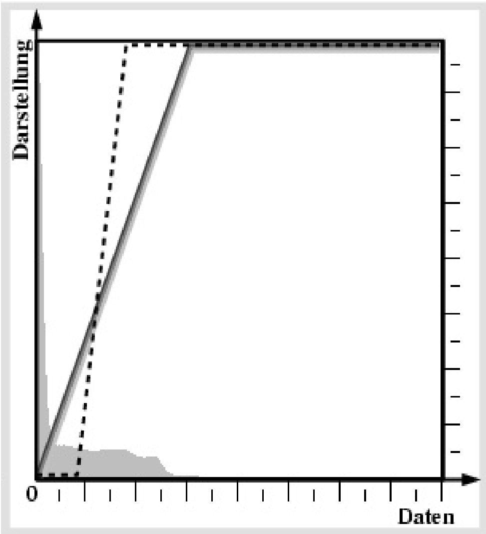

Source: Hastreiter, 1999

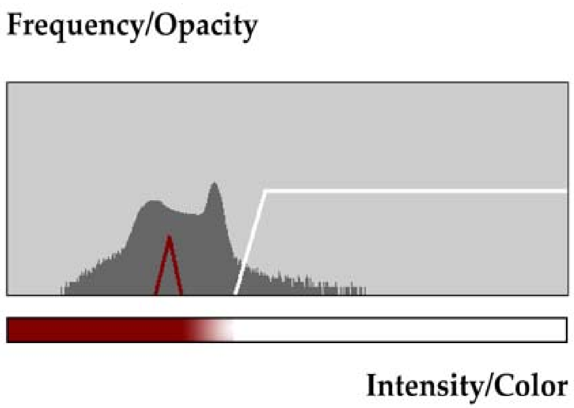

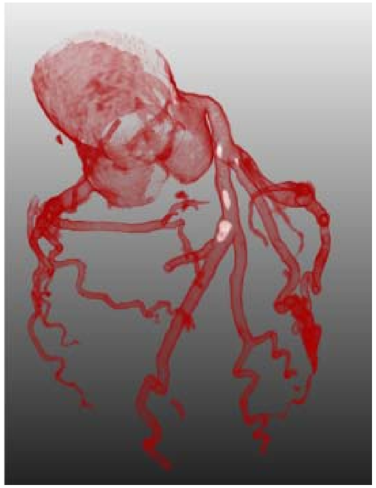

[Kubisch 2010A]

Classification

- Interpolation and application of TF (classification)
  - **Pre-Classification:** Application of TF to all edge points in the filter range (result: RGBA quadruple); afterwards: (tri)linear interpolation of this quadruple
  - **Post-Classification**: Interpolation of the intensity values from the data (e.g. Hounsfield Units); afterwards: application of transfer function to the interpolated result (pre-integrated for quality enhancement)
- Pre-classification problems:
  - Perception of the interpolated colors is non-linear
  - Imprecise classification

Pre-, post-, preintegrated classification [Engel06]

Classification

- Interpolation and application of the transfer function

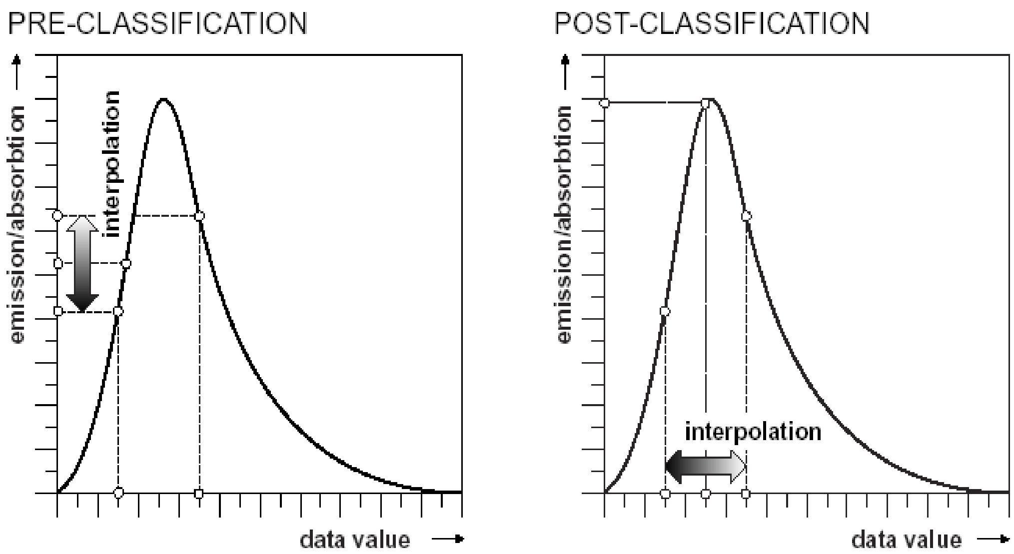

Late application of TF is more precise!

Source: Rezk-Salama, 2002

Classification

- Interpolation and application of the transfer function. Comparison of pre- and post-classification

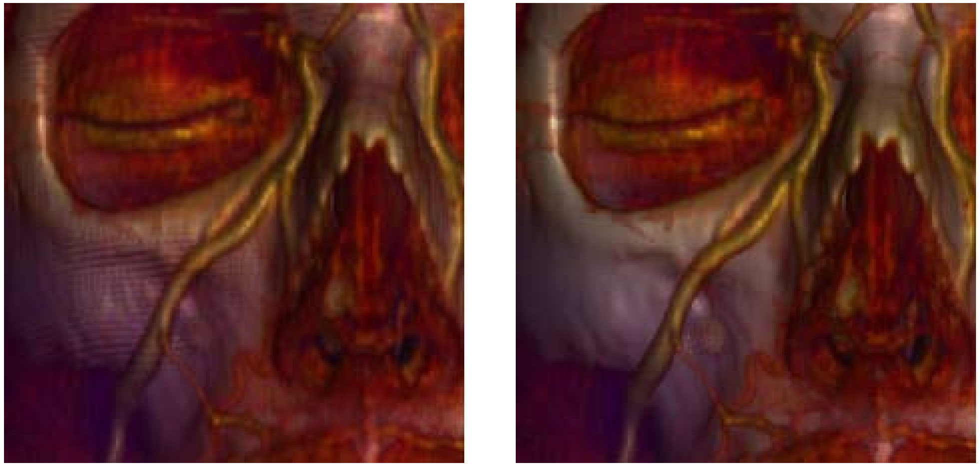

Source: Bruckner, 2004 (Master thesis)

Classification and Transfer Functions

Ambiguity of volume data: A certain intensity value in CT or MRI data may represent

- one tissue type or
- a mixture of tissue types in the boundary region (partial volume effect).

- The correct visualization requires to resolve this ambiguity (Drebin, 1988).
- Ideally, volume data is transformed in a *material percentage volume* representing for each voxel to which materials it belongs (with a certain percentage).

Prerequisite:

- A priori knowledge of material types and their intensity range

Classification and Transfer Functions

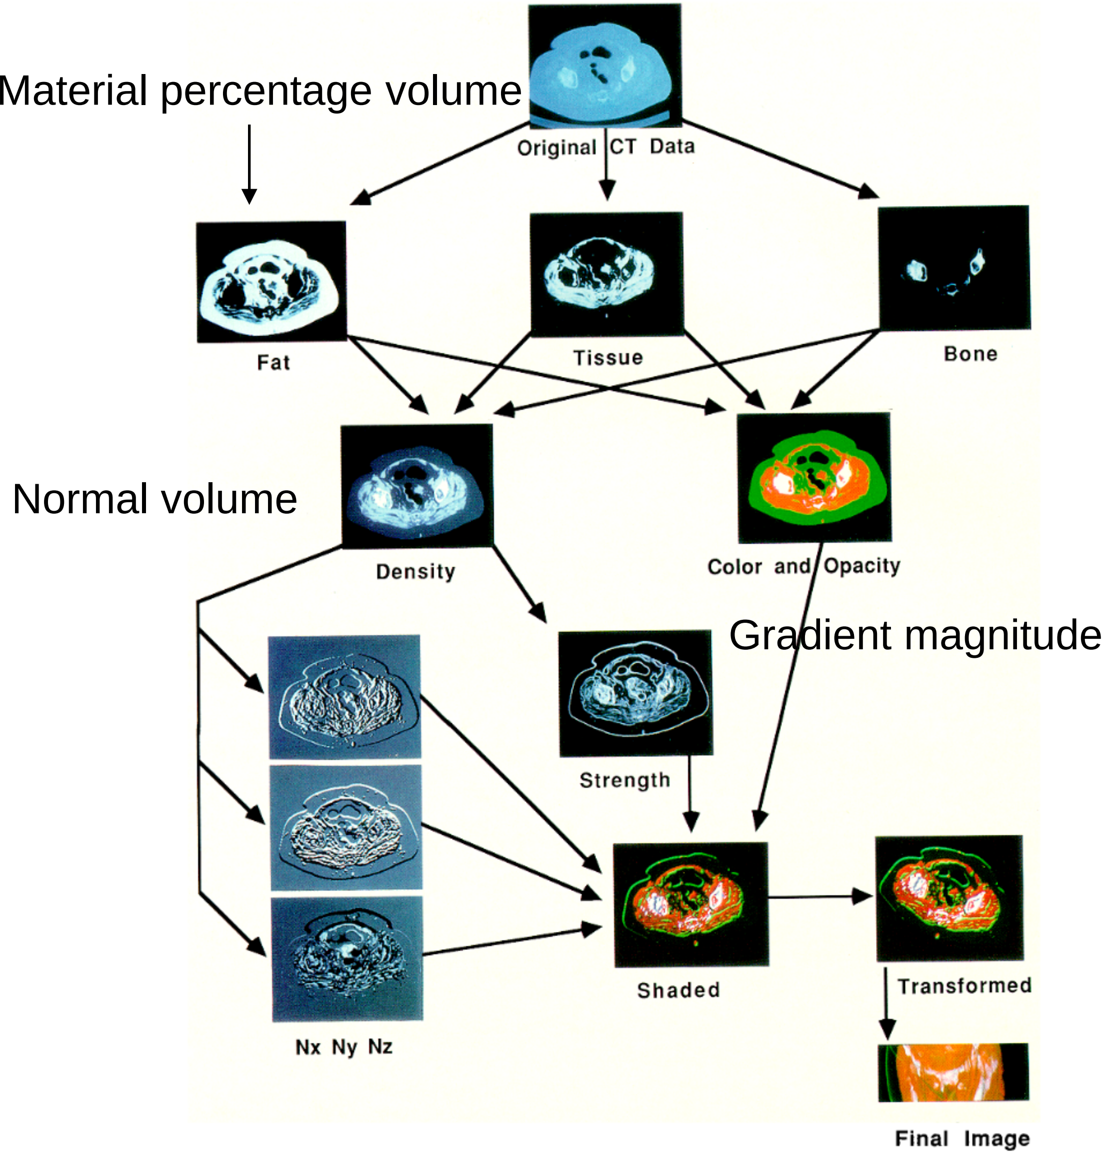

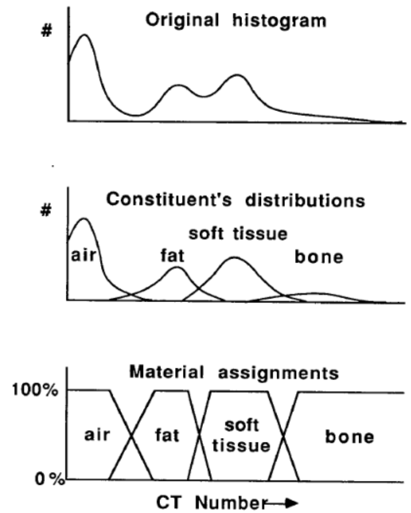

- Volume data modelled as a mixture model of four components.
- Percentages are determined linearly between two tissue types.
- The TF considers intensity values and material percentage data (Drebin, 1988).

Classification and Transfer Functions

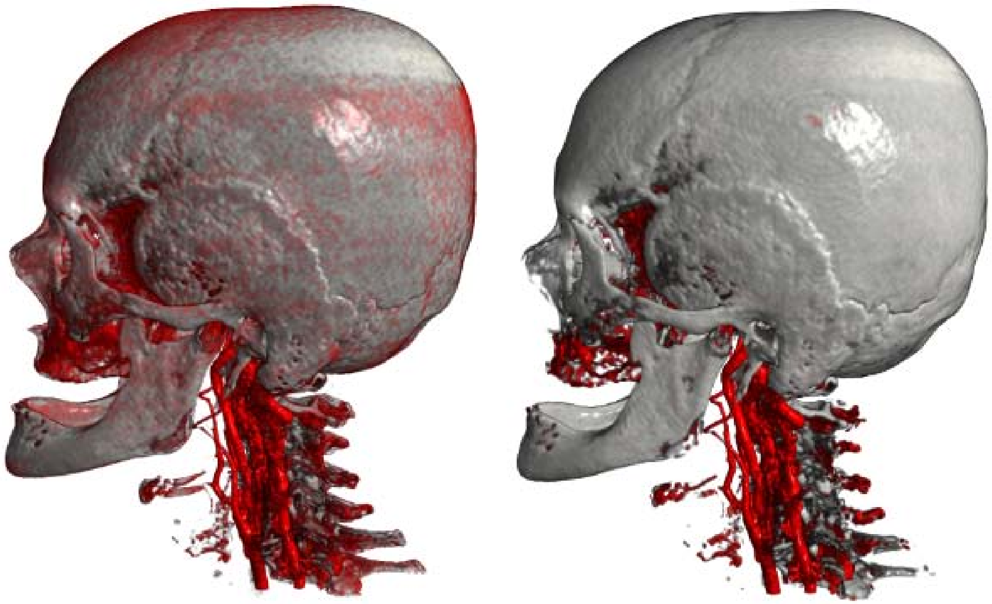

- **Boundary-aware volume reconstruction**. Without knowledge of materials and boundaries, the data is treated as continuous and reconstructed erroneously (left). Interpolation needs to be restricted to voxels within one tissue type and prevented between them (right) (Lindholm, 2014)

Classification and Transfer Functions

Simplifications in this model:

- Only transitions between two materials in one voxel are considered (similar to Marching Cubes!).
- Only transitions between materials where the intensity distributions overlap are considered.
  - In practice, there are transitions, e.g. between bones and air in the nasal cavity.
- Advanced models employ classification techniques, such as Maximum Likelihood to model voxels with more than two tissue types.

Alternative techniques for solving the ambiguity problem:

- Consider derived information (gradient magnitude, curvature)
- Tagged volume rendering (using segmentation)

Classification and Transfer Functions

Slight changes of a transfer function or isovalue may reveal a stenosis or hide it or indicate stenosis that is wrong (Lundström, 2007).

Animations that slightly change the TF would avoid this problem.

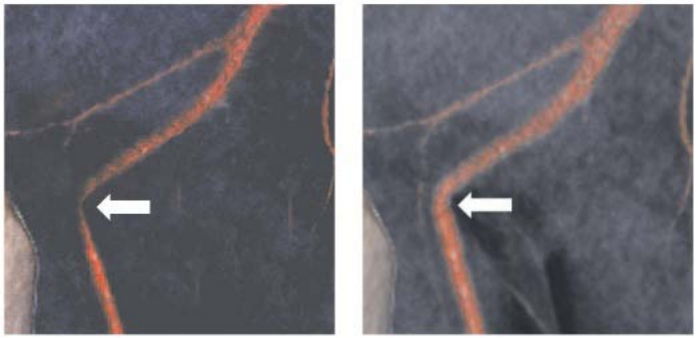

The default TF from a commercial programme indicates a stenosis (Lundström, 2007).

For a survey on medical problems and uncertainty visualizations (Ristovsky, 2014)

Volume Rendering Equation

Volume rendering equation (without shadows and diffusion):

$$I_\lambda(x,r) = \int_0^L c_\lambda(s) e^{-\int_0^s \alpha(t)dt} ds$$

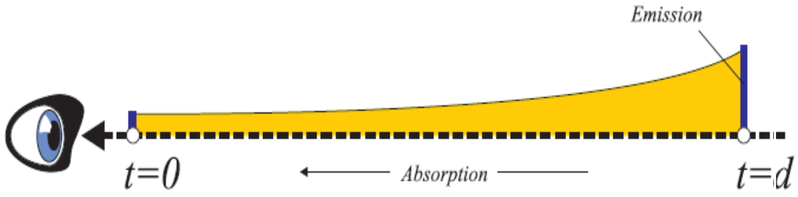

[SiggraphCourse2008]

I_λ Amount of light of wave length λ

r Direction of the viewed ray

x Position on the image plane

L Length of the light ray

s Current point on the light ray

C_λ Light of wavelength λ that is reflected or emitted from *s* in direction to *r*

α Density of the particles that are reflected to the viewer

Numerical Approximation

- Numerical approximation:

$$I_\lambda(x,r) = \int_0^L c_\lambda(s) e^{-\int_0^s \alpha(t)dt} ds$$

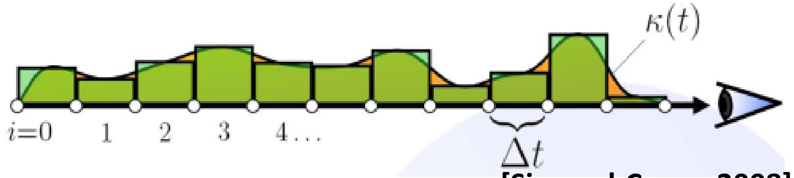

[SiggraphCourse2008]

$$I_\lambda(x,r) = \sum_{i=0}^{L/\Delta s} c_\lambda(s_i) * \prod_{j=0}^{i-1} (1-\alpha(s_j))$$

- $C_\lambda(s_i)$ Local color at position s_i
- $\alpha\ (s_i)$ Opacity at position s_i
- $C_\lambda(s_i)$ and α are the transfer functions
- Product describes the remaining visibility (after traversal of the light ray)

Resampling Filter

- Each volume rendering technique (image-based, object-based, and texture-based) involves resampling at sampling points in the volume
- Different filter kernels for resampling (here in 1D)
  - Nearest neighbor
  - Tent (linearly interpolated)
  - Sinc (ideal signal reconstruction, but too costly)

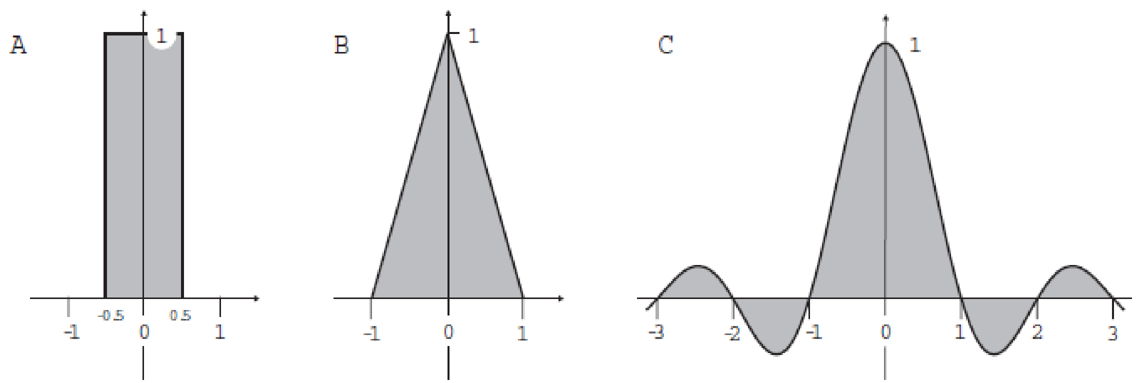

[Engel06]

Trilinear Interpolation

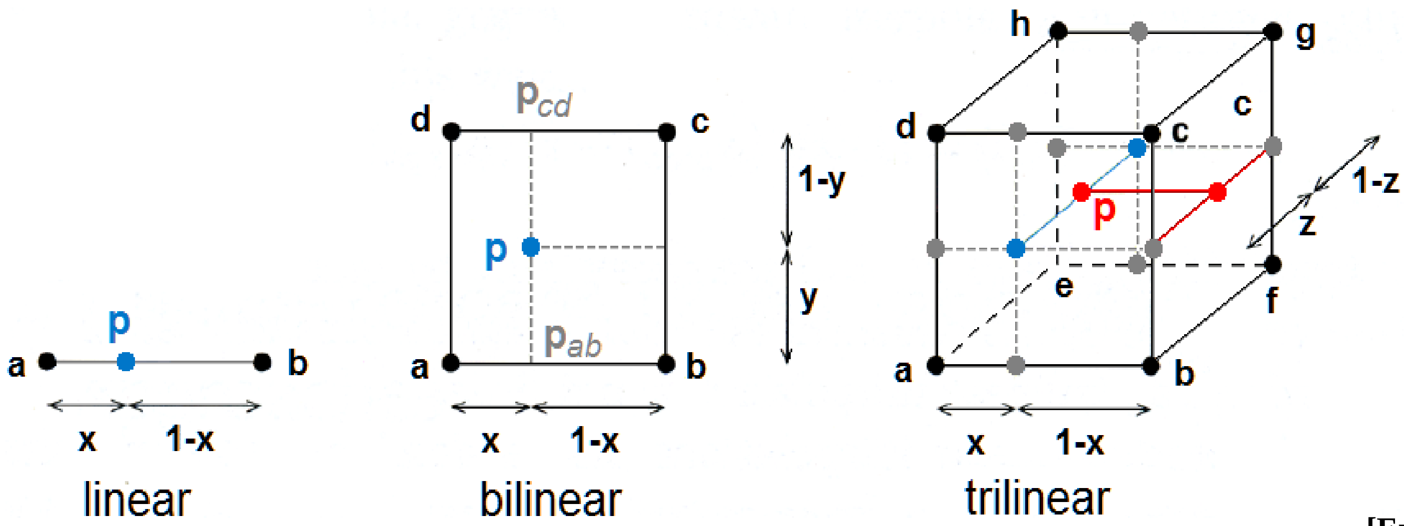

[Engel06]

- Trilinear interpolation (graphics hardware)
  - Stepwise 1D interpolation of extensions
    - Along an axis
    - Between the results of the first axis
    - … the second axis
  - Result: typically “rounded staircases”

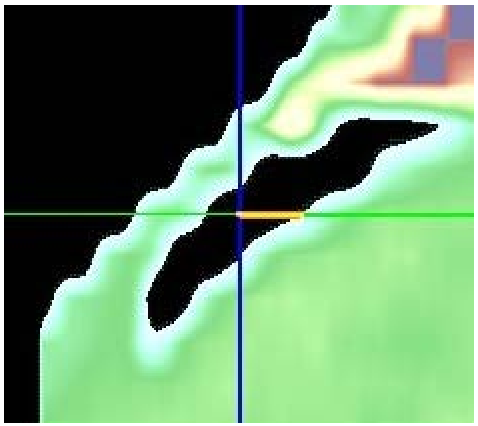

Compositing via α Blending

Set: I (0) = C(0), α(0) =1

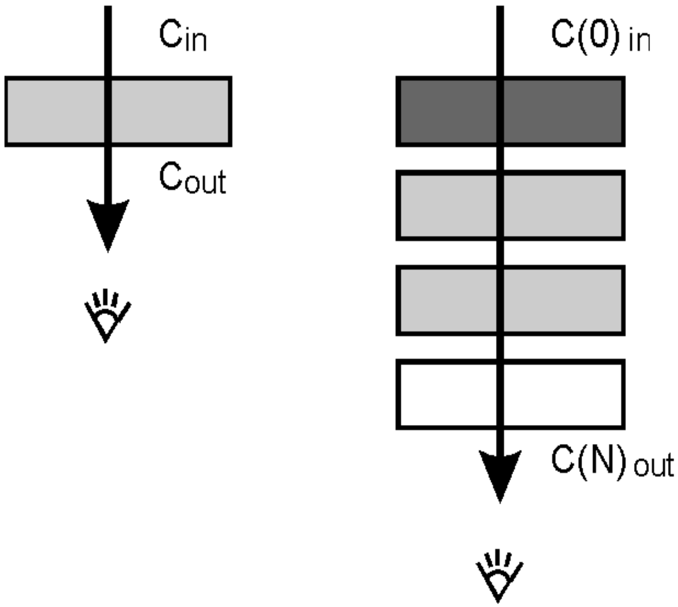

Color value and transparency of a pixels results from overlay of semi-transparent voxels (compositing).

$C_i := C_{i-1} * (1-\alpha_i) + \alpha_i C_i$ (color)

$\alpha_I := \alpha_{i-1} * (1-\alpha_i) + \alpha_i$ (opacity)

If 1-α_i becomes very small, a termination is possible.

Image-based Techniques

- Tracing of rays into the scene (Raycasting)
- Per scan point:
  - Rounding to the nearest voxel (nearest neighbor)
  - Resampling filter from the surrounding voxels
- Variations:
  - First-hit in case of iso-surfaces
  - Consideration of multiple reflexions (raytracing)

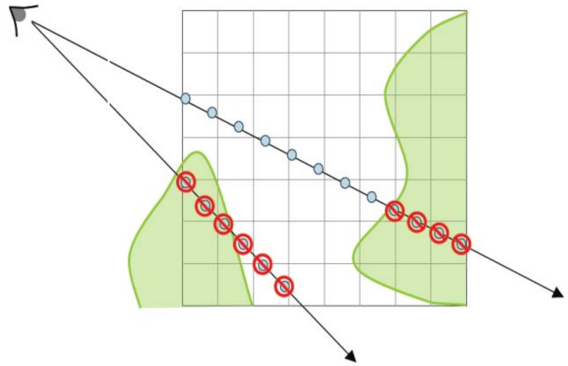

First Hit-Raycasting [Tariq07]

[Hadwiger05]

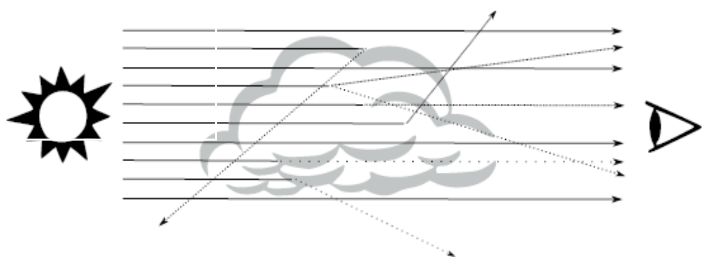

Light transfer in a variable density medium [Kratz06]

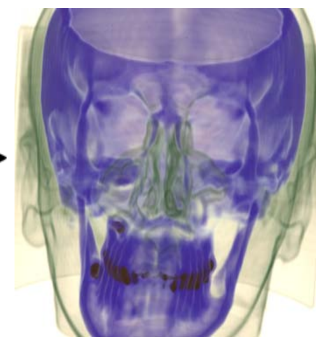

[Engel06]

Image-based Techniques

- Basic Algorithm Raycasting:

  for y_i = 1 **to** ImageHeight

  for x_i = 1 **to** ImageWidth

  for z_i = 1 **to** RayLength

  foreach x_0 in ResamplingFilter (x_i, y_i, z_i)

  foreach y_0 in ResamplingFilter (x_i, y_i, z_i)

  foreach z_0 in ResamplingFilter (x_i, y_i, z_i)

  add contribution of Voxel [x_0, y_0, z_0] to ImagePixel [x_i, y_i]

  Resampling filter usually corresponds to trilinear interpolation (2x2x2 values).
- **Problem**: Volume is not traversed in the order in which it lies in the storage. Usually, voxels are required that are not in the cache or in the main storage.

- **Problem**: uniform sampling in case of perspective projection (diverging rays)
- **Remedies for this problem**:
  – Splitting of rays
  – Rays integrate over a larger area in more distant areas

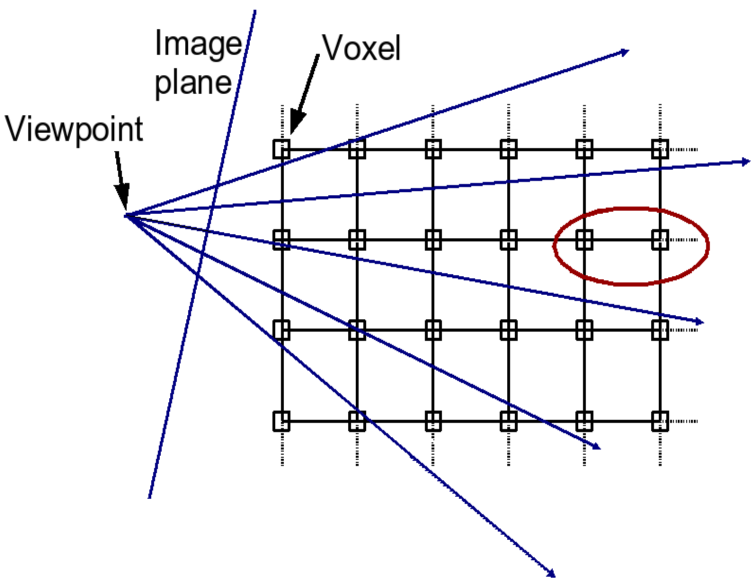

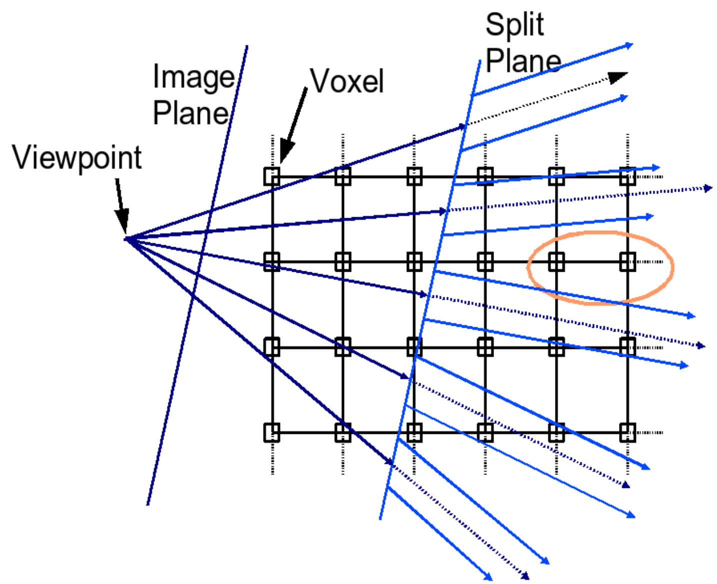

Image-based Techniques

- Raycasting was considered the slowest technique in 2000
- Acceleration Methods
  - **Early ray termination** (e.g. in case of 95% opacity) (Levoy [1990])
  - **Adaptive ray sampling** (sampling rate in strongly transparent areas, or increase in case of higher distance) (Danskin, Hanrahan [1992])
  - **Hitpoint refinement** (binary search of intersection point for iso-surfaces) (Sigg 2004)

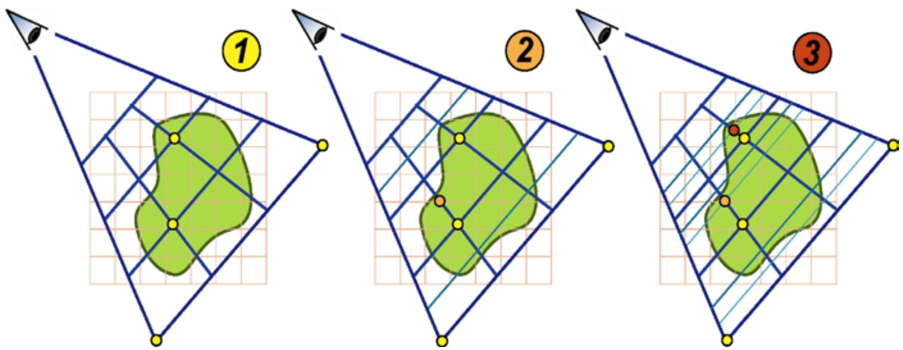

Binary search by dividing the step sizes [Sigg04]

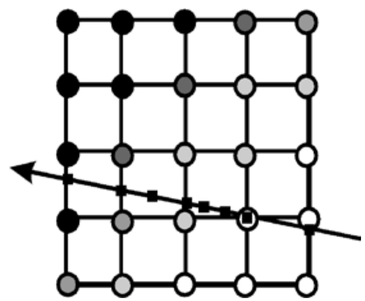

Adaptive ray sampling

Image-Based Techniques: GPU Raycasting

- Motivated by advances in graphics hardware (strongly enhanced flexibility with vertex and fragment shaders, multiple render targets and frame buffer objects since 2001)
- Creation of textures that encode the start and end positions of the rays (Krüger, 2003)
  - usually by means of rasterization of a simple hull geometry
- Integration along the ray
  - Simple programming model, very high performance
  - Transfer function and composition are part of the traversing for loop
- Compared to texture-based rendering:
  - More flexibility
  - Incorporation of acceleration strategies possible (early ray termination, empty space skipping)

- The volume is normalized in (0,1) range in all directions.
- Entry- and exit points are determined by rendering the cube twice (once with backface culling, once with frontface culling).
- A framebuffer object (a texture) is used as rendering target.
- The actual compositing is carried out in a third rendering

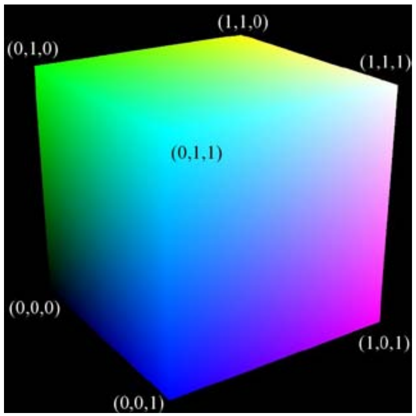

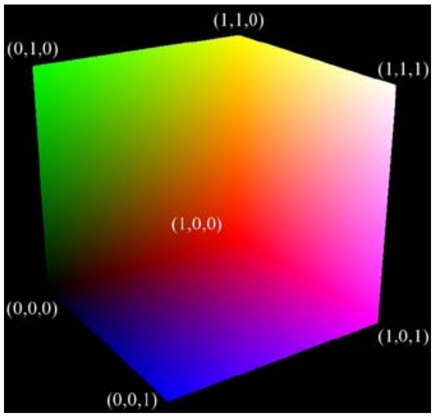

Courtesy of Cezary Bloch, Link

Image-Based Techniques: GPU Raycasting

- Again entry- and exit points are determined.
- The transfer function is evaluated to assess which parts of the date contribute significantly to the rendered image.
- Rays are shortened strongly to accelerate rendering.

Color encodes depth

[Glaßer, 2010]

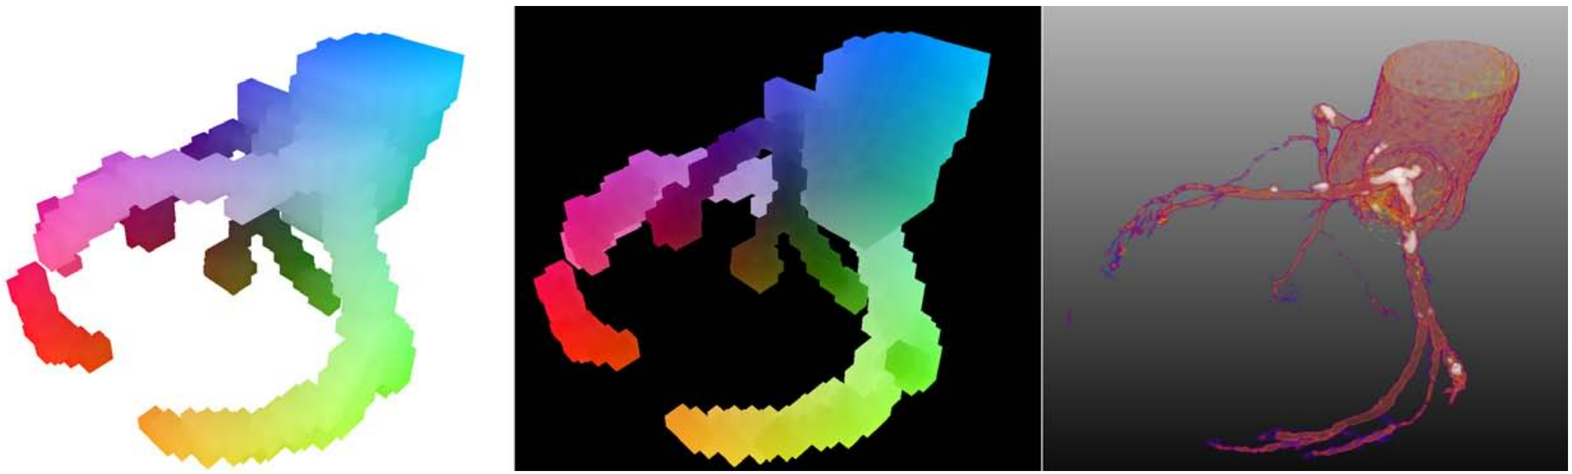

Direct Volume Visualization: Image-based Methods

- Influence of the sampling rate on alias effects (increment: 2.0 voxel, 1.0 voxel, 0.1 voxel), (© Schroeder et al. [1998])

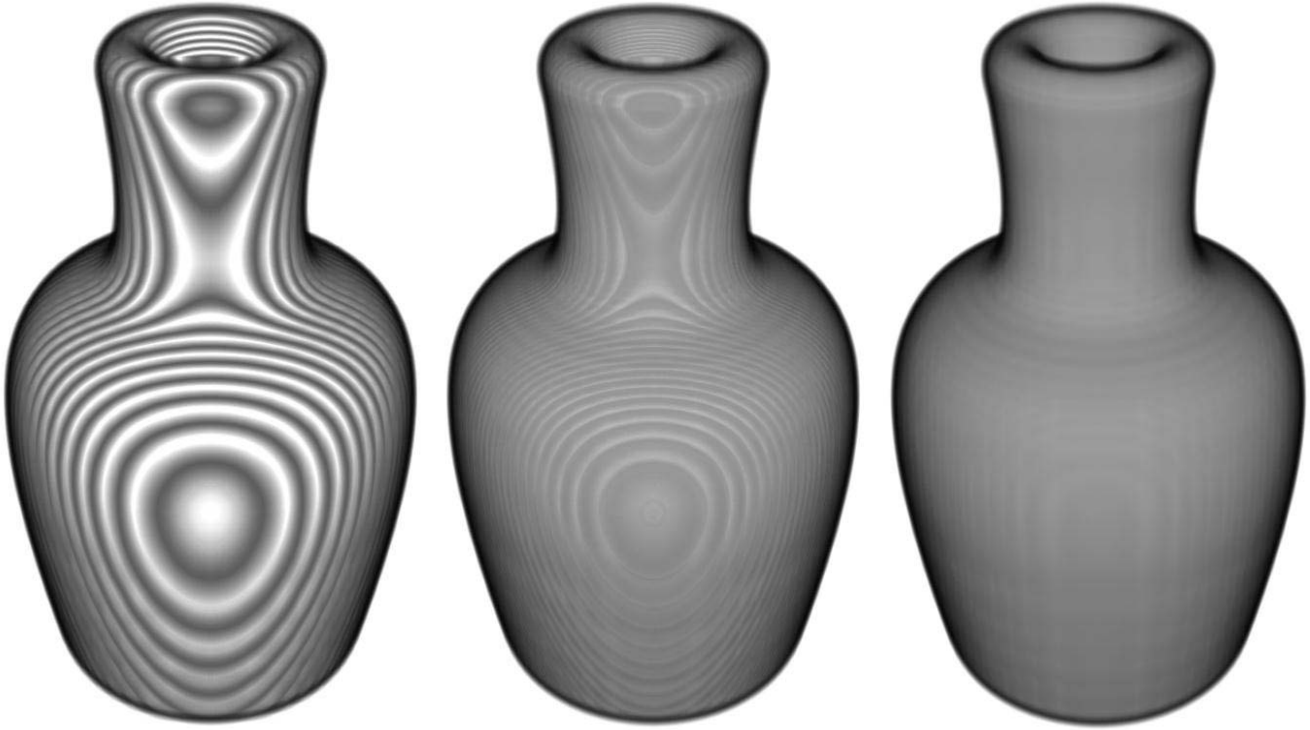

- Suggestion: increment < 0.5 voxels (according to the sampling theorem: sampling at least with the double frequency which is present in the discrete data).

Direct Volume Rendering: Image-based approach

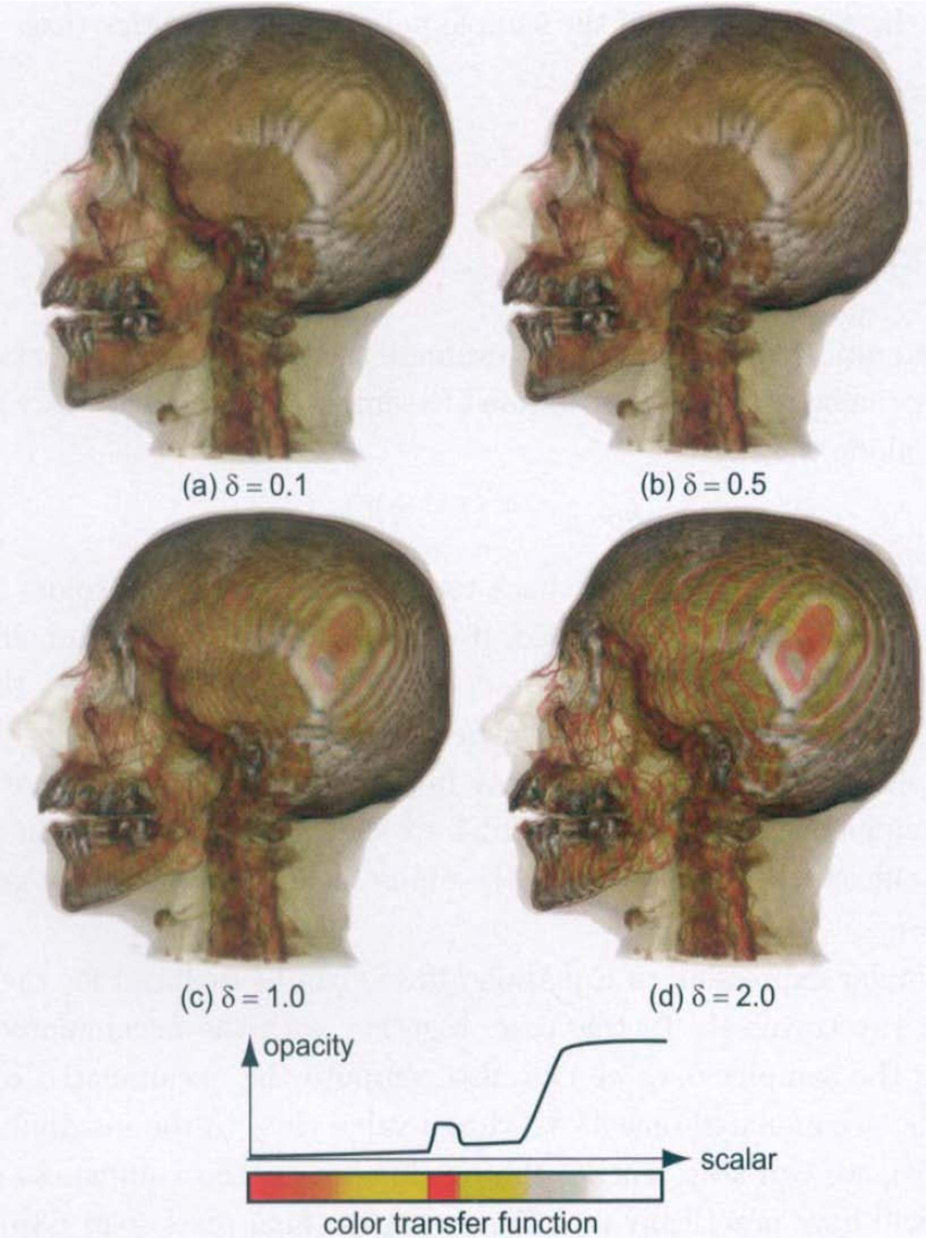

Comparison of image quality with four different values for the ray step size.

Source: Telea [2008]

Direct Volume Visualization: Projection Methods

Average Projection	Average of all hit voxels per ray	Simulation of x-ray projections
Maximum (minimum) Intensity Projection (M(m)IP)	Brightest and (darkest) voxel hit per ray	Illustration of vessels, noise-added data
Closest Vessel Projection (Zuiderveld [1995])	First hit voxel per ray above a threshold	Illustration of vessels

Direct Volume Visualization: Projection Methods

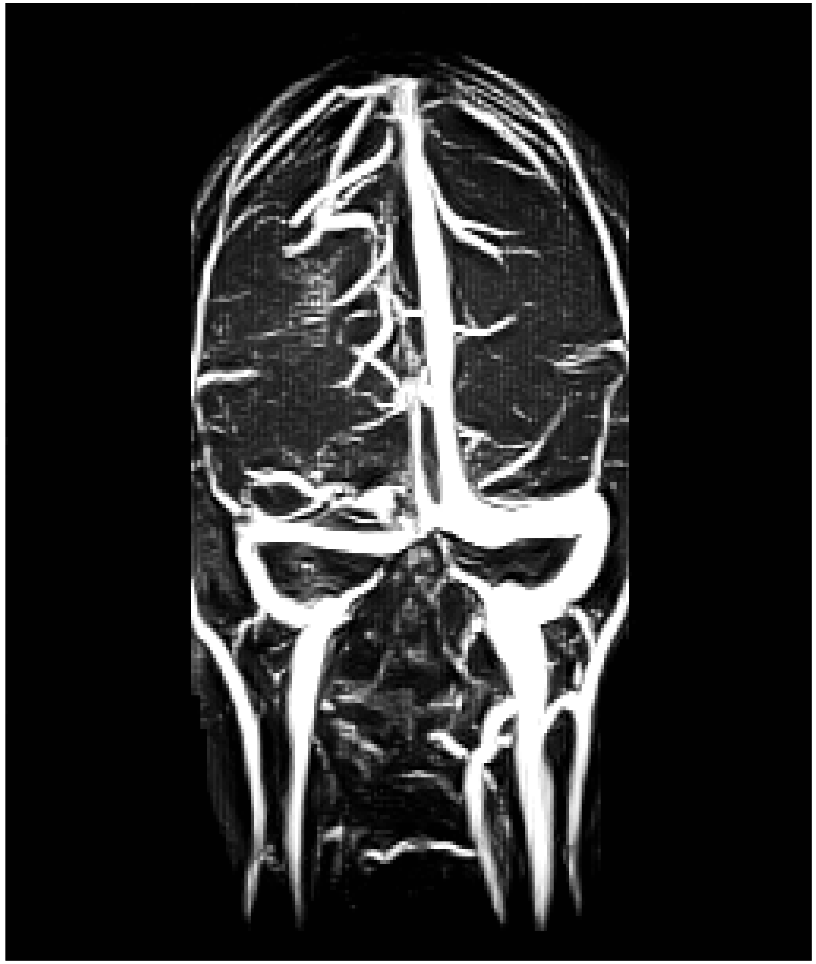

MIP (Data: MR angiography)

Direct Volume Visualization: Projection Methods

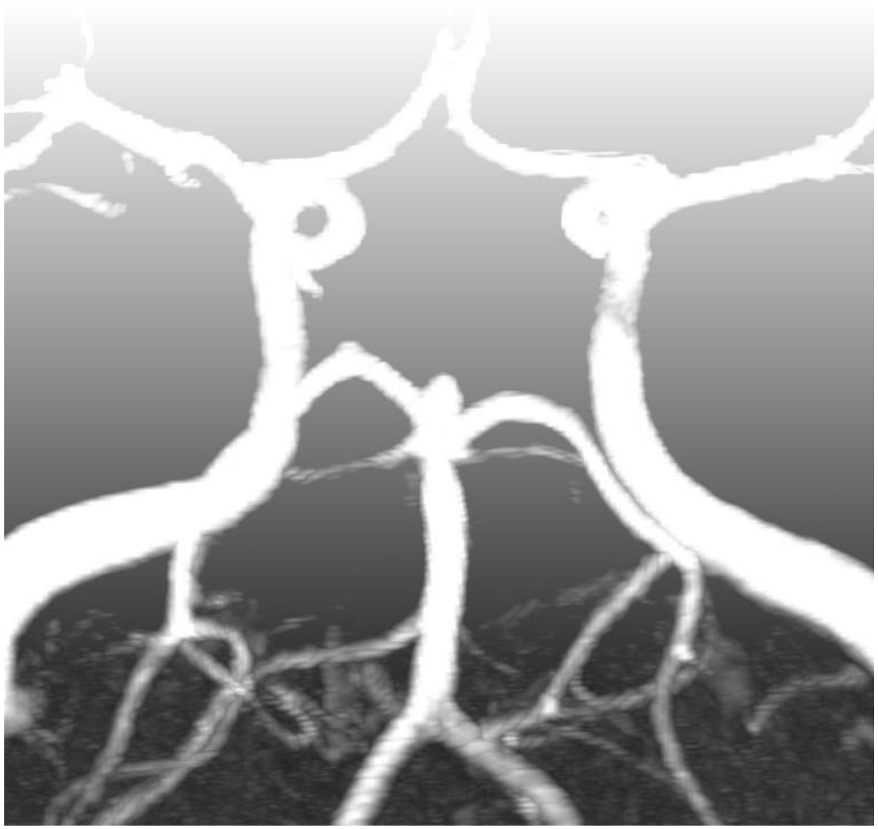

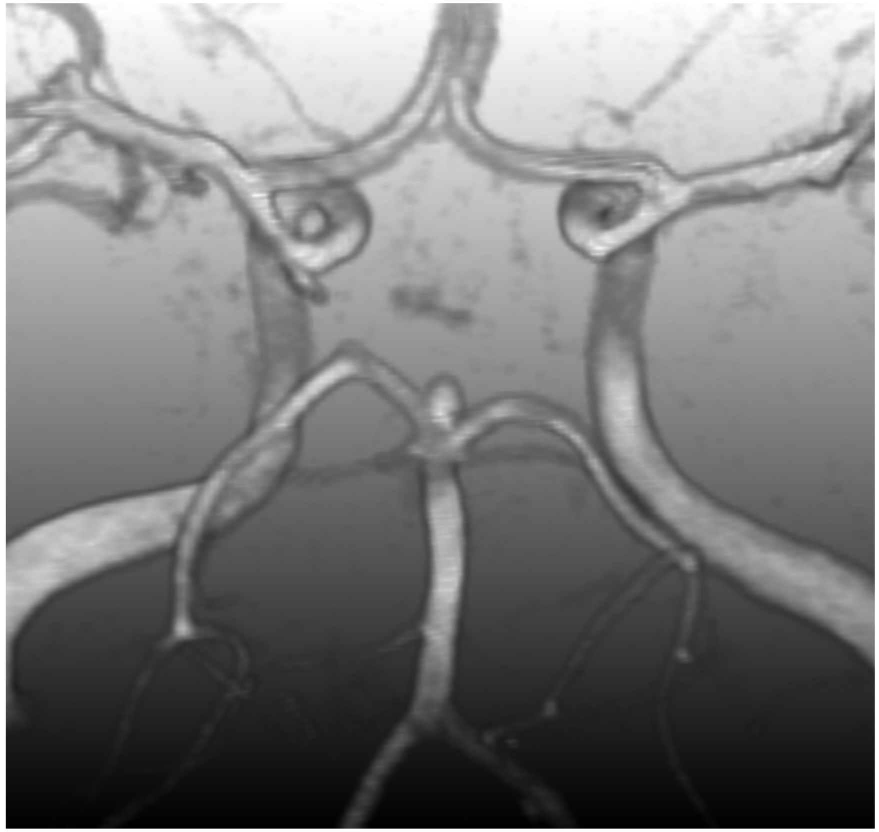

Comparison of MIP and DVR, cerebral vessels, purpose: diagnosis of aneurysms (Data: MR angiography, Prof. Terwey, Bremen)

Direct Volume Visualization: Projection Methods

Restriction of the data on which a MIP is applied:

- (1) Remove certain structures which disturb the MIP display.

  Example: Removal of bones (interactively byplacing a seed point and Region Growing).

  (2) Apply the MIP to a certain partial volume.

  Example: MIP illustration in a segmented organ for the selective evaluation of this organ

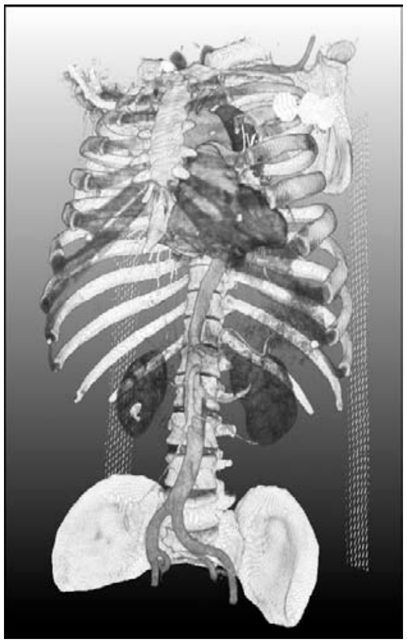

Before and after bone removal,

© Hans Drexl, Fraunhofer MEVIS

Direct Volume Visualization: Projection Methods

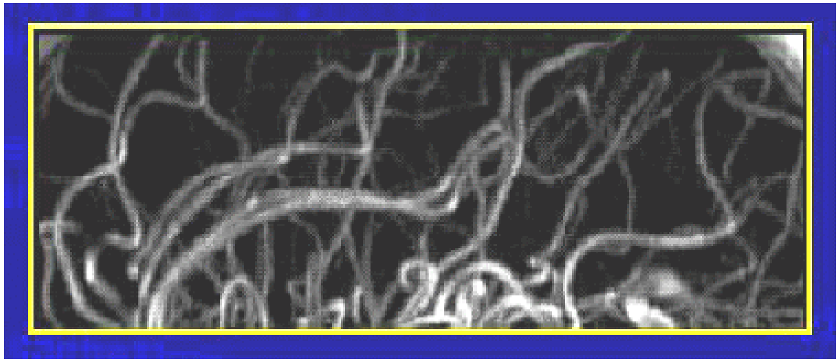

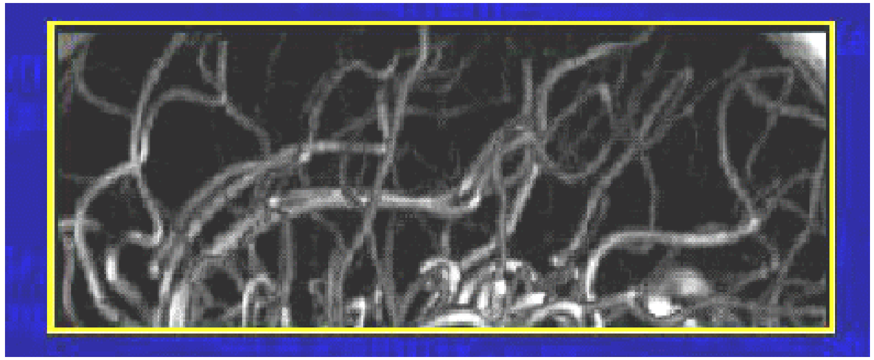

- MIP and CVP of brain vessels (© Karel Zuiderveld)
- To evaluate spatial relations, movies with rotations of MIP and CVP in a central perspective are often used.

Direct Volume Visualization: Shading

- Angle of incidence θ: angle between L and N (determines the diffuse reflection)
- Reflection angle r: angle between R and N.
- Angle Φ between V and R determines the intensity of the incident light.
- If V = R (respectively Φ =0), the light is reflected maximal to the viewer.

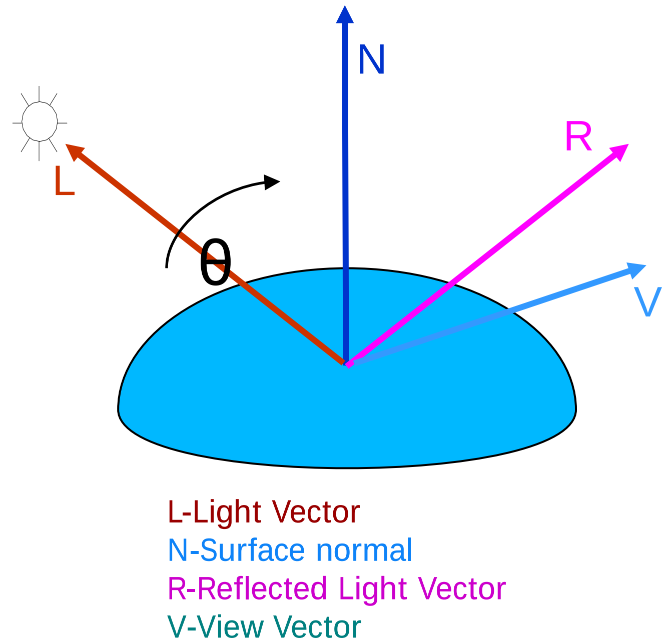

Direct Volume Visualization: Shading

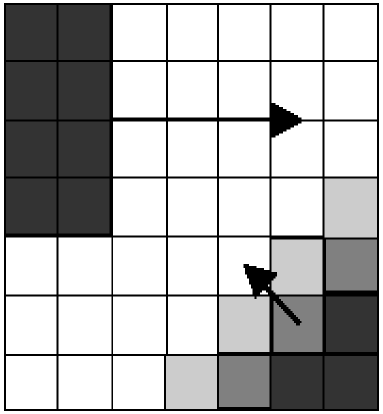

- Approximation of the surface normal by calculating the gradient (grey level gradient shading, [Höhne and Bernstein, 1986])
- Problem: Memory requirements: 4 Byte * 3 per voxel
- Indirect storage of the normals as indices in a field of normalized vectors (rounding)

  → Discretization of the normal in a gradient lookup table

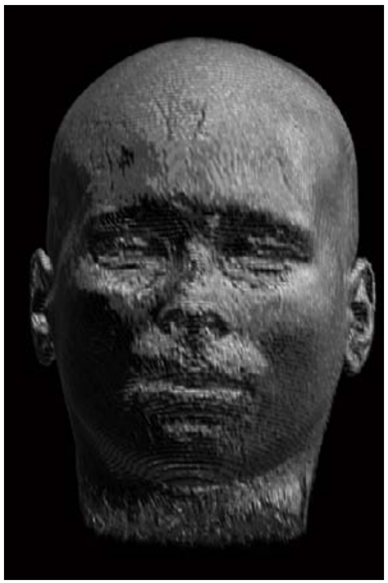

- Illuminated illustration of an MRI data set (high sampling rate and trilinear interpolation)

Problems:

- High noise sensibility (possibly smooth gradients) or ignore small gradients (use the threshold value)
- No consideration of the gradient strength

Direct Volume Visualization: Shading

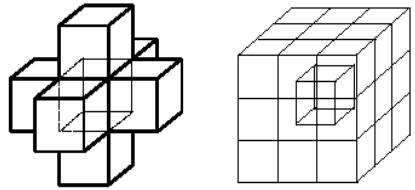

- Common variants of gradient estimation:

  (1) central differences (6 neighbors):

  $$\nabla V(X) = (\partial V/\partial x, \partial V/\partial y, \partial V/\partial z)$$

  $$\begin{aligned}\nabla V(xi, yj, zk) = (&\tfrac{1}{2}\,(V(xi+1, yj, zk) - (V(xi-1, yj, zk)),\\ (&\tfrac{1}{2}\,(V(xi, yj+1, zk) - (V(xi, yj-1, zk)),\\ (&\tfrac{1}{2}\,(V(xi, yj, zk+1) - (V(xi, yj, zk-1))\,)\end{aligned}$$

  (2) Gradient estimation of from the 26 neighbors (weighting according to the distance from the central voxel)

- The second variant is more complex than the first one, but qualitatively better.
- Problems: treatment of boundaries, line structures

Direct Volume Visualization: Shadow

- Shadow may further enhance depth perception
- Requires the definition of a light source and the analysis, how the voxels are oriented towards the light source.
- Method:
  - Two-Pass-Rendering: First Pass: illumination per voxel is computed and represented in the shadow-Buffer. Second pass: image generation based on the shadow buffer. (Levoy [1988])
  - Disadvantage: required size of a 3d-Shadow-Buffer
  - Recent refinements reduce memory consumption and increase performance:
    - Deep Shadow Maps (Kratz [2006]),
    - Adaptive Volumetric Shadow Maps (Salvi [2010])

Direct Volume Visualization: Shadow

(a) DVR (b) Shaded DVR (c) DVR with shadows

(a) DVR (b) Shaded DVR (c) DVR with Shadows

[Kratz2006]

Direct Volume Visualization: Ambient Occlusion

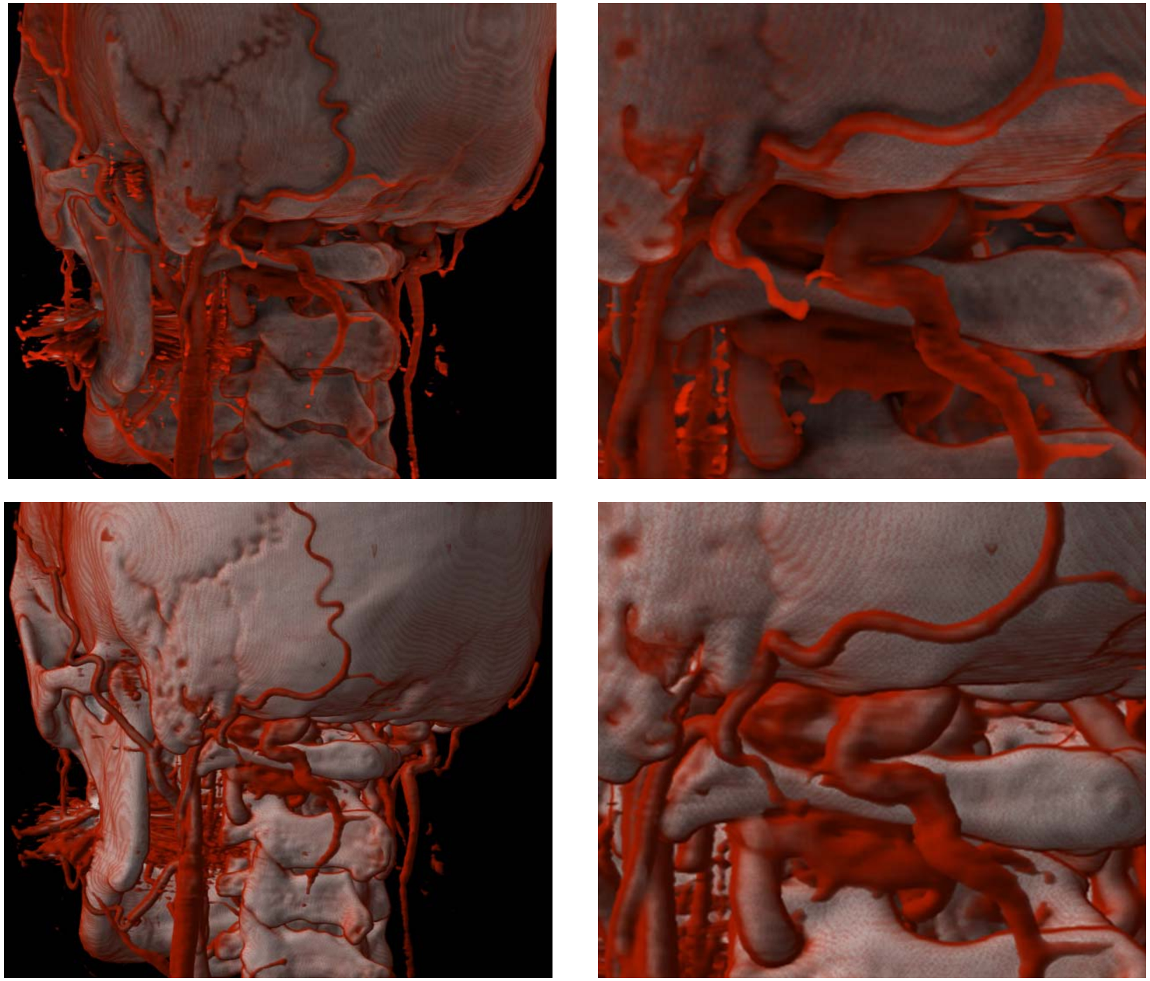

Top: Local ambient occlusion. Bottom: gradient-based shading (applied to CT angiography data). Understanding of vessel structures is improved (Hernell, 2007).

Direct Volume Visualization: Ambient Occlusion

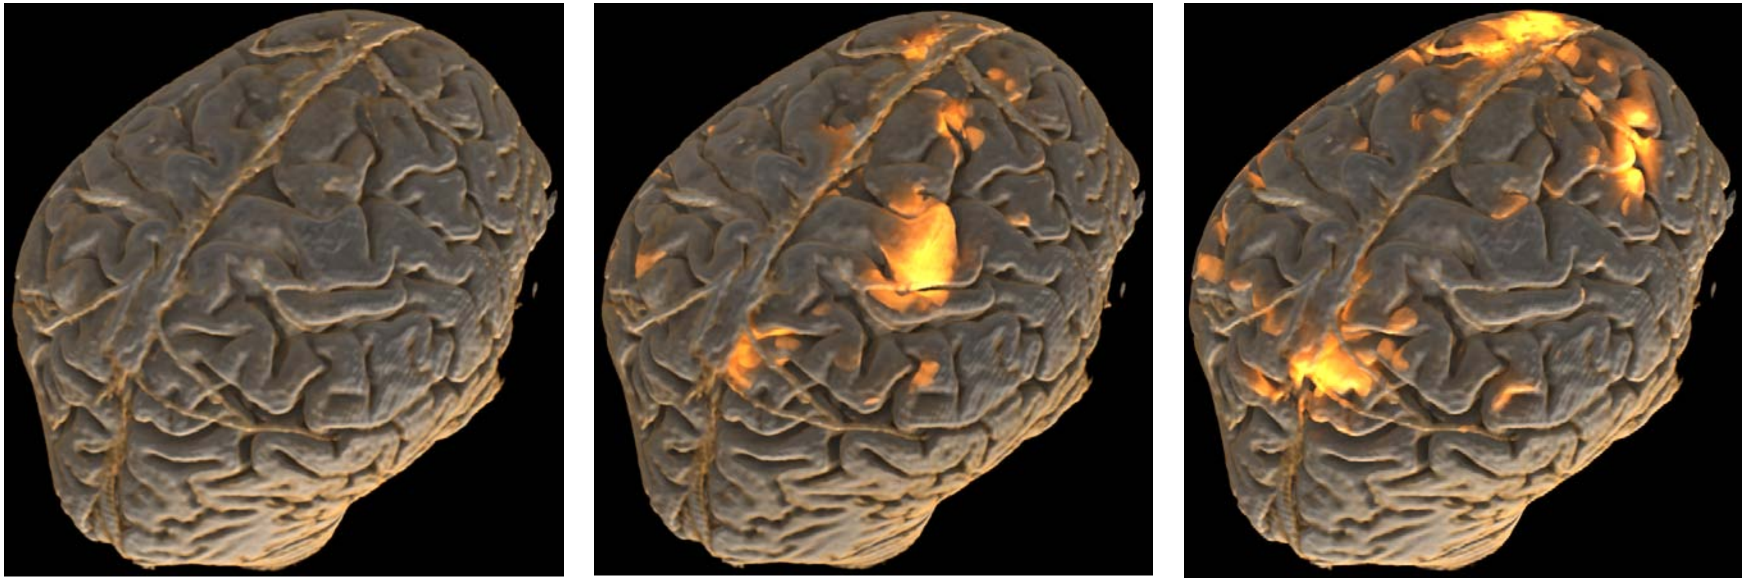

Left: ambient occlusion is used to enhance shape perception. **Middle and right**: emissive lightning and ambient occlusion are used to convey functional data (activity during math and language tasks) (Nguyen, 2010).

Direct Volume Visualization: Tagged VR

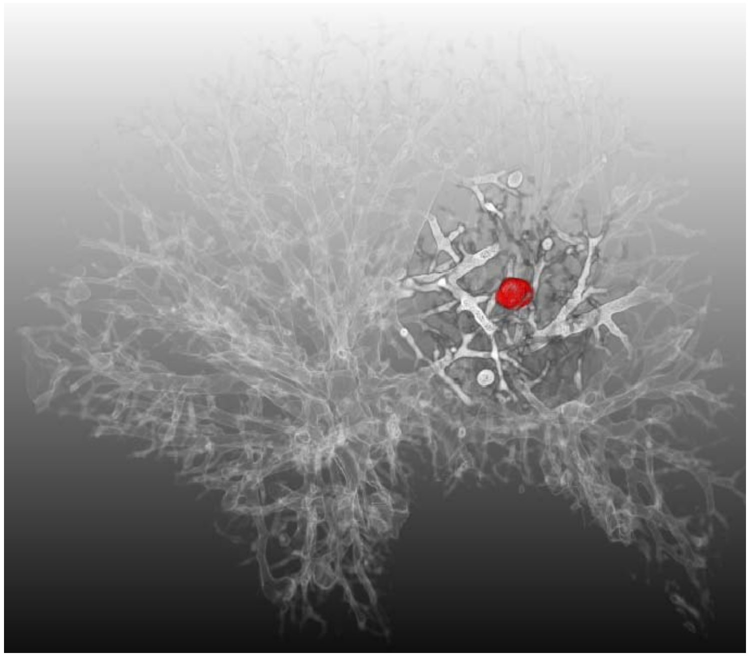

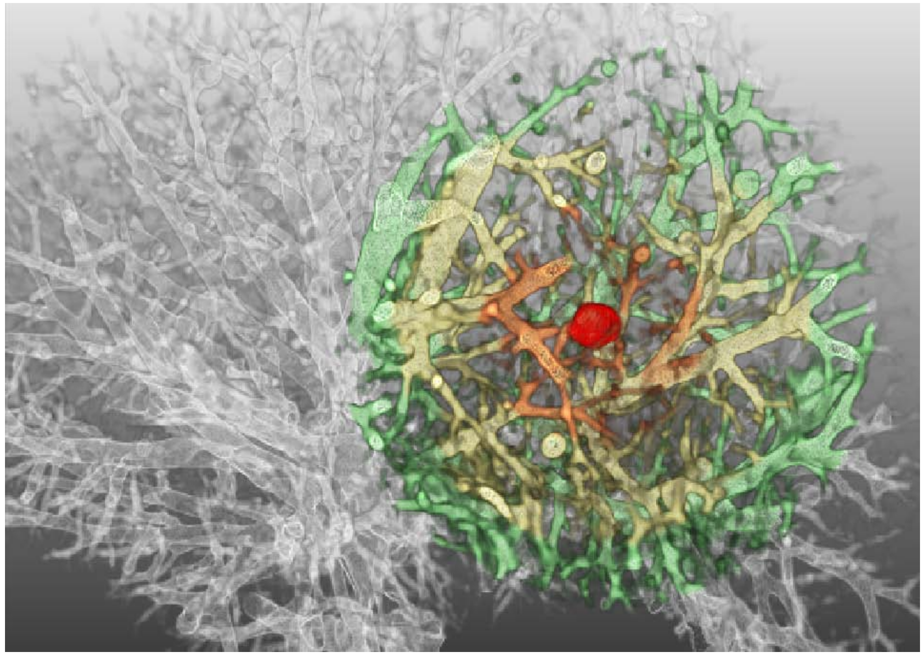

Tappenbeck [2006]

- Segmentation: Tumor
- Visualization: Distance-based TFs (distance to tumor mapped to opacity and color)

Direct Volume Visualization: Hierarchical Methods

- Goal: restrict rendering to visible portions and/or importance
- Typical data structure: Octree
- Node size, 16x16 …. 64x64
- Requires resampling, e.g. by means of a rank filter
- Overlap of the nodes for correct interpolation (1 voxel)
- Moderate additional memory load

Direct Volume Visualization: Hierarchical Methods

- Octree nodes are rendered back to front
- Order of nodes depends on the viewing direction
- Lower resolution may be used for interactive rendering

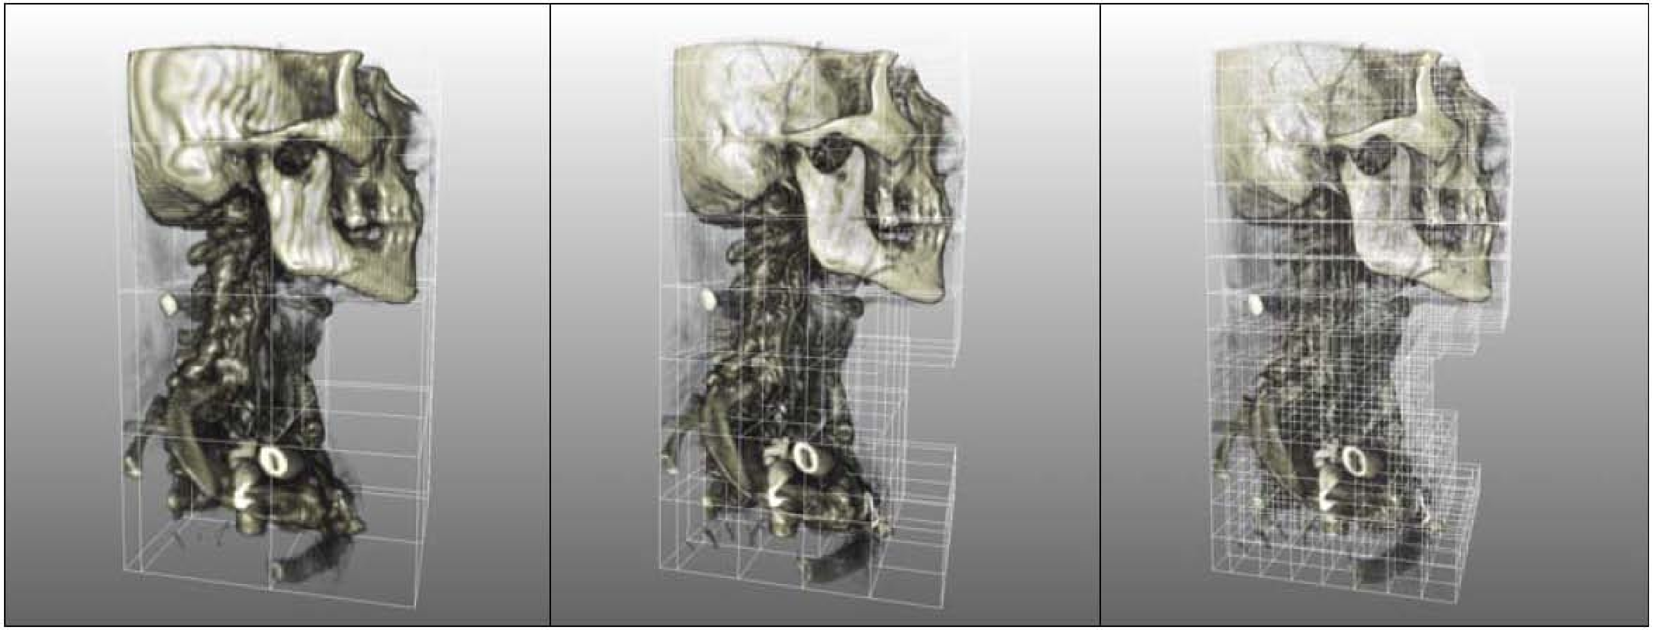

Link [2006]

Tagged Volume Rendering

- Additional volume saves "Tag ID"
  - usually 8 bit = 256 different structures are sufficient
- Problem: tag ID cannot be interpolated (discrete) and voxelwise classification (and pre-classification) is worse
  - » e.g. interpolation between ID 3 and 5 would yield 4. However, this ID usually does not exist in the neighborhood.
  - Use of trilinear weighting factors [0,1] to be able to use a threshold value
    - » BUT: impossible to distinguish between more than two IDs
    - » Segmentation must ensure this limit

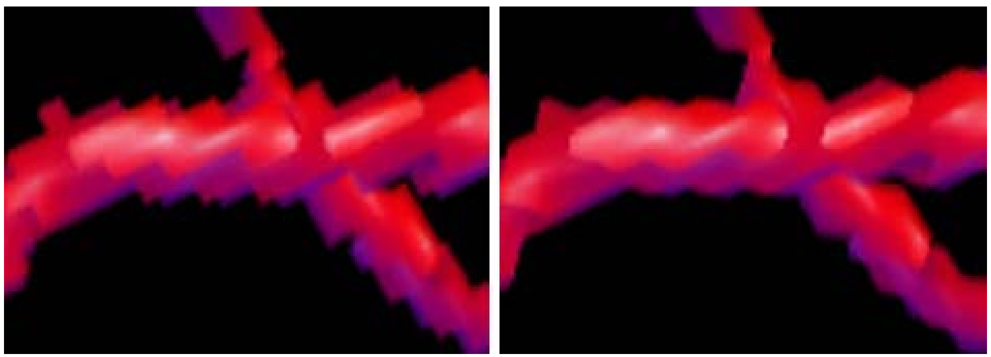

[Hadwiger, 2004]

Tagged Volume Rendering

- Additional shading variation is possible
  - Here, materials are encoded per "lit-spheres"

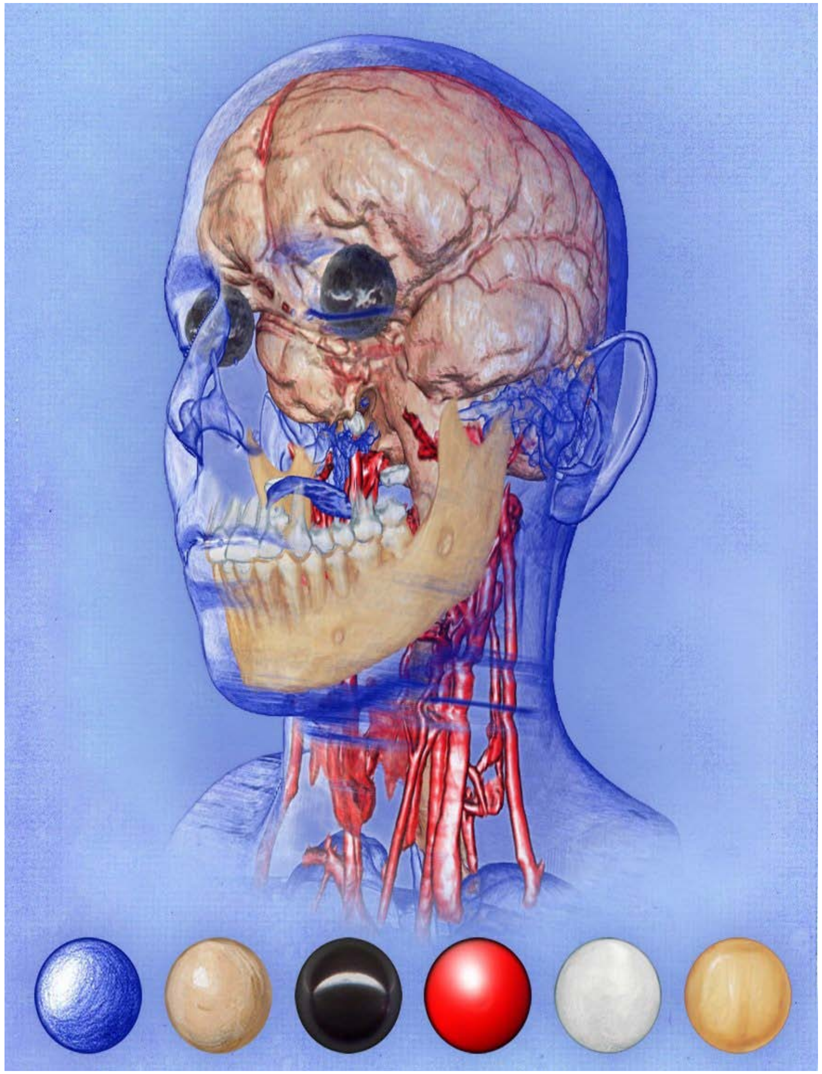

[Bruckner2007]

Summary

- Direct volume rendering does not require binary decisions which portions of the data are rendered.
- GPU-based raycasting as current standard.
- Transfer functions specify mappings to opacity and color.
- Gradient approximation and lightning improve spatial perception.
- Volume rendering benefits from segmentation (tagged volume rendering).

References

- S. Bruckner, E. Gröller: „Instant Volume Visualization using Maximum Intensity Difference Accumulation“, *Comput. Graph. Forum* 28(3): 775-782 (2009)
- B. Cabral, N. Cam, J. Foran. “Accelerated volume rendering and tomographic reconstruction using texture mapping hardware”. *Proc. of the 1994 symposium on Volume visualization*
- B. Csébfalvi, L. Mroz, H. Hauser, A. König, E. Gröller: Fast Visualization of Object Contours by Non-Photorealistic Volume Rendering. *Comput. Graph. Forum* 20(3): 452-460 (2001)
- C. Crassin, F. Neyret, S. Lefebvre, M. Sainz, E. Eisemann: „Beyond triangles: gigavoxels effects in video games“. *SIGGRAPH Posters 2009*
- R. A. Drebin, L. Carpenter, P. Hanrahan. 1988. Volume rendering. *SIGGRAPH Comput. Graph.* 22(4) (1988), 65-74
- V. Dicken, B. B. Wein, H. Schubert, J.-M. Kuhnigk, S. Krass, H.-O. Peitgen. „Projektionsansichten zur Vereinfachung der Diagnose von multiplen Lungenrundherden in CT-Thorax-Aufnahmen“. *Bildverarbeitung für die Medizin* 2003: 244-248
- K. Engel, M. Hadwiger, J. Kniss, C. Rezk-Salama, D. Weiskopf. *Real-time Volume Graphics*, AK Peters, 2006
- S. Glaßer, S. Oeltze, A. Hennemuth, C. Kubisch et al.: Automatic Transfer Function Specification for Visual Emphasis of Coronary Artery Plaque, *Comput. Graph. Forum* 29(1): 191-201 (2010)
- M. Hadwiger, C. Berger, H. Hauser. High-Quality Two-Level Volume Rendering of Segmented Data Sets on Consumer Graphics Hardware. *IEEE Visualization* 2003: 301-308
- Hadwiger, Markus; High-Quality Visualization and Filtering of Textures and Segmented Volume Data on Consumer Graphics Hardware; PhD Thesis, Vienna University of Technology; 2004
- H. Hauser, L. Mroz, G. I. Bischi, E. Gröller: Two-Level Volume Rendering. *IEEE Trans. Vis. Comput. Graph.* 7(3): 242-252 (2001)

References

- M. Hadwiger, C. Sigg, H. Scharsach, K. Bühler, M. H. Gross: Real-Time Ray-Casting and Advanced Shading of Discrete Isosurfaces. *Comput. Graph. Forum* 24(3): 303-312 (2005)
- A. Kratz, M. Hadwiger, R. Splechtna, A. Fuhrmann, K. Bühler. „GPU-Based High-Quality Volume Rendering For Virtual Environments*", Proc. of Internatl. Workshop on Augmented Environments for Medical Imaging and Computer Aided Surgery* (AMI-ARCS), 2006
- J. Krüger and R. Westermann. Acceleration Techniques for GPU-based Volume Rendering. In *Proc. of IEEE Visualization 2003*, pp. 38-45
- C. Kubisch, S. Glaßer, M. Neugebauer, B. Preim: Vessel Visualization with Volume Rendering. *Visualization in Medicine and Life Sciences* II 2012: 109-134
- M. Levoy: Efficient ray tracing of volume data. *ACM Trans. Graph.* 9(3): 245-261 (1990)
- S. Lindholm, D. Jönsson, C. Hansen, A. Ynnerman: Boundary Aware Reconstruction of Scalar Fields. *IEEE Trans. Vis. Comput. Graph.* 20(12): 2447-2455 (2014)
- C. Lundström, P. Ljung, A. Persson, A. Ynnerman: Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation. *IEEE Trans. Vis. Comput. Graph.* 13(6): 1648-1655, 2007
- C. Rezk-Salama. *Volume Rendering Techniques for General Purpose Graphics Hardware*, Phd Thesis, Computer Science Department, University of Erlangen-Nürnberg, 2002
- U. Tiede, T. Schiemann, K. H. Höhne. 1998. High quality rendering of attributed volume data. In *Proc. of IEEE Visualization '98*, pp. 255-262
- Wallis, J.W.; Miller, Tom R.; Lerner, C.A.; Kleerup, E.C., "Three-dimensional display in nuclear medicine," *IEEE Transactions on Medical Imaging*, Vol. 8(4), pp.297-, 1989
- Zuiderveld, Karel; *Visualization of Multimodality Medical Volume Data*; PhD Thesis, University Utrecht; 1995