

D²VR: High-Quality Volume Rendering of Projection-based Volumetric Data

Peter Rautek[†], Balázs Csébfalvi[‡], Sören Grimm[†], Stefan Bruckner[†], and Meister Eduard Gröller[†]

[†] Vienna University of Technology, Austria
[‡] Budapest University of Technology, Hungary




Figure 1: CT scan with insufficient resolution of a stag beetle as specimen. Figure (a) to (d) show DVR using different reconstruction schemes: (a) using nearest neighbor reconstruction, (b) using trilinear-, (c) using Catmull-Rom-, and (d) using polynomial interpolation of degree five. (e) shows D²VR of projection-based volumetric data. Figure (a) - (d) were rendered from grids with a resolution of 64^3 . Figure (e) was rendered from 64 projections each with a resolution of 64^2 .

Abstract

Volume rendering techniques are conventionally classified as either direct or indirect methods. Indirect methods require to transform the initial volumetric model into an intermediate geometrical model in order to efficiently visualize it. In contrast, direct volume rendering (DVR) methods can directly process the volumetric data. Modern CT scanners usually provide data as a set of samples on a rectilinear grid, which is computed from the measured projections by discrete tomographic reconstruction. Therefore the rectilinear grid can already be considered as an intermediate volume representation. In this paper we introduce direct direct volume rendering (D²VR). D²VR does not require a rectilinear grid, since it is based on an immediate processing of the measured projections. Arbitrary samples for ray casting are reconstructed from the projections by using the Filtered Back-Projection algorithm. Our method removes a lossy resampling step from the classical volume rendering pipeline. It provides much higher accuracy than traditional grid-based resampling techniques do. Furthermore we also present a novel high-quality gradient estimation scheme, which is also based on the Filtered Back-Projection algorithm.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism, I.3.3 [Computer Graphics]: Picture/Image Generation

1. Introduction

Modern 3D scanning technologies usually provide data values on rectilinear grid points. These data values are computed from measured projections by discrete tomographic

reconstruction [Rus92, KS88]. The set of the reconstructed data values (or samples) can be interpreted as a discrete representation of the underlying continuous phenomenon. In order to authentically visualize the original continuous signal, it has to be accurately reconstructed from the discrete samples. Note that such a signal reconstruction is differentiated from discrete tomographic reconstruction. From a signal-processing point of view, the original signal can be

[†] {rautek | grimm | bruckner | meister}@cg.tuwien.ac.at

[‡] cseb@iit.bme.hu

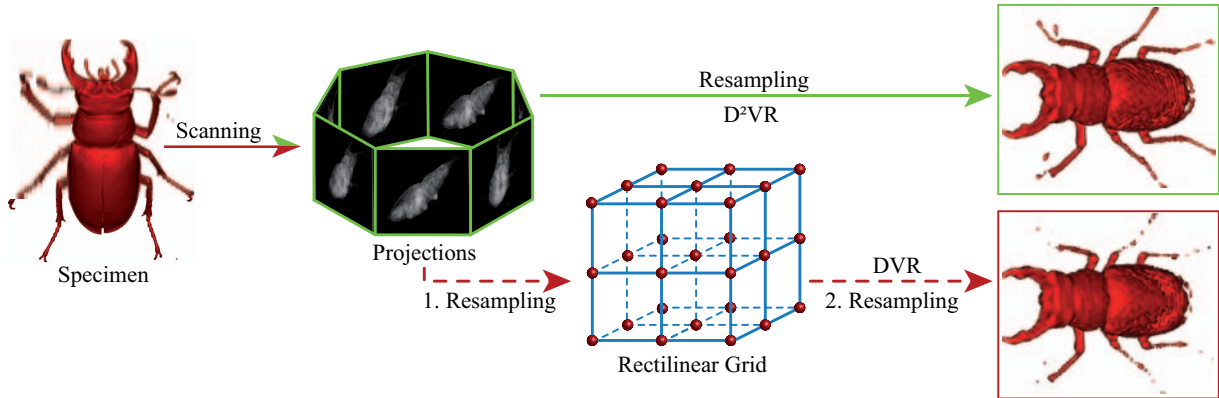


Figure 2: Data processing work flow of projection- and grid-based volume rendering. The dashed line corresponds to the traditional volume rendering pipeline. It requires two resampling steps in order to visualize the data. First an intermediate grid is resampled and then this grid is resampled again during ray traversal. The solid line corresponds to the projection-based volume rendering pipeline; one lossy resampling step is avoided.

perfectly reconstructed from discrete samples if it is band-limited and the sampling frequency is above the Nyquist limit [OS89]. Theoretically the perfect continuous reconstruction is obtained by convolving the discrete volume representation with the *sinc* function. The *sinc* function is considered to be the best reconstruction kernel, since it represents an ideal low-pass filter. In practice, however, it is difficult to convolve a discrete signal with the *sinc* kernel, because of its infinite support. Therefore practical reconstruction filters either approximate it or truncate it with an appropriate windowing function [ML94, TG00]. Moreover, real-world signals can hardly be considered band-limited. As a consequence, practical resampling results in a loss of information. Figure 2 shows the signal-processing approach of the traditional volume rendering pipeline (follow the dashed line). After the scanning process a rectilinear grid is computed from the measured projections. This first resampling step is done by discrete tomographic reconstruction. Although there exist different algorithms for tomographic reconstruction the most popular technique used is the Filtered Back-Projection algorithm [CCF94]. It first performs high-pass filtering on the measured projections. Afterwards the samples at rectilinear grid points are computed by back-projecting the filtered signals. As the projections are acquired by measuring accumulated attenuation by a limited number of sensors, they are actually available as discrete representations of continuous projection functions. Therefore high-pass filtering is performed in the discrete frequency domain, so the result is also a discrete function. In the back-projection phase, however, the rectilinear grid points are not necessarily projected exactly onto the discrete samples of the filtered projections. Therefore, resampling is necessary for back-projection, which results in the *first* loss of information in the pipeline. The obtained rectilinear volume can be visualized by different rendering tech-

niques. Using indirect methods, like the classical Marching Cubes algorithm [LC87], an intermediate geometrical model of an iso-surface is constructed from the volumetric model. This geometrical model is then interactively rendered by, for example, conventional graphics hardware. In contrast, Direct Volume Rendering approaches, like raycasting [Lev88] or splatting [Wes90, ZPvBG01] directly render the volumetric model without any intermediate representation. In both cases an interpolation technique is applied to define data values between the rectilinear grid points. In other words, a resampling of the discrete volume representation is required. This resampling results in the *second* loss of information in the traditional pipeline. In order to minimize the loss of information we propose to modify the traditional volume rendering pipeline by simply removing a resampling step (follow the solid line in Figure 2). To render the underlying continuous phenomenon, data samples at arbitrary sample points need to be defined, and for shading computation the corresponding gradients need to be determined. As it will be shown, both tasks can be solved using directly the filtered projections. This eventually leads to an alternative projection-based volume representation. Traditional direct volume rendering methods rely on an intermediate grid representation, so in this sense they are in fact indirect. In contrast, we present DVR *directly* from the measured raw data. To distinguish from the common DVR the novel approach is referred to as D^2VR [di: skweəd vi: a:]. With ideal reconstruction filters DVR and D^2VR would generate the same result. With practical but erroneous reconstruction filters the avoidance of one resampling step in D^2VR poses a clear quality advantage. In Section 2 we review previous work related to discrete tomographic reconstruction and volume resampling. In Section 3 our novel volume rendering approach is introduced. It is explained how to reconstruct data values and gradients directly from the projections by

using the Filtered Back-Projection algorithm. Section 4 reports the results. Finally in Section 5 the contribution of this paper is summarized and ideas for future work are given.

2. Related Work

In most of the practical volume rendering applications, especially in 3D medical imaging, the input data is usually generated from measured projections by using tomographic reconstruction [Rus92, KS88, NCN*97]. The set of projections is referred to as the Radon transform of the original signal. Therefore the tomographic reconstruction is, in fact, the inversion of the Radon transform. The inversion can be performed by using the classical Filtered Back-Projection [CCF94] algorithm, which is based on the Fourier projection-slice theorem [KS88, Mal93]. The output of tomographic reconstruction is a discrete (or sampled) representation of the underlying continuous phenomenon. Samples are conventionally generated on rectilinear grid points. The rectilinear grid has several advantages. For example, the sampled signal can be represented by 3D arrays, implicitly storing the locations of the samples. Furthermore, the neighborhood of a certain sample can be efficiently addressed, which is important for many volume processing or volume rendering algorithms. Nevertheless, in order to render the underlying continuous 3D function, data values need to be defined also between the rectilinear grid points. The *sinc* kernel as ideal reconstruction filter is impractical because of its infinite extent. In practice it is approximated by filters of finite support [ML94, TG00]. Generally, the wider the support of the reconstruction filter, the higher the quality of the reconstruction. On the other hand, the wider the support of the filter, the higher the computational cost of a spatial-domain convolution. Therefore several researchers analyzed different reconstruction filters, both in terms of accuracy and computational cost [MN88, ML94, MMMY97, MMK*98]. As the practical filters only approximate the ideal low-pass filter they result in either aliasing or smoothing [ML94], which can be interpreted as a loss of information. For frequent resampling tasks, like rotation, or upsampling, frequency-domain techniques can be alternatively applied [LM04, AMVG05, CCM99, CT99, TC99, UTY95]. In the frequency domain, it is exploited that a computationally expensive spatial-domain convolution is replaced by a simple multiplication. Although the frequency-domain resampling methods generally provide higher accuracy than spatial-domain methods, they assume that the new samples to be computed are also located at regular grid points. In order to avoid a lossy resampling step in the traditional volume rendering pipeline, we directly use the tomographic inversion in order to reconstruct the underlying function at arbitrary sample positions. Therefore we do not generate an intermediate rectilinear volume representation, but we directly process the filtered projections as an alternative volume representation. Using this projection-based volume rendering approach, the same accuracy can be ensured at arbitrary

sample positions. In contrast, using the traditional grid-based approach, accurate samples are available only at the grid points, while the accuracy of intermediate samples depends on the quality of the applied imperfect reconstruction filter.

3. D^2VR

We present D^2VR based on a raycasting approach. In order to perform raycasting the underlying 3D volumetric function needs to be reconstructed at arbitrary resampling locations. In case the data is given on a rectilinear grid the reconstructed function value is computed from a close neighborhood of samples as shown in Figure 3a. In contrast to that, raycasting based on the projections computes the reconstructed function value from filtered projections at the corresponding positions (see Figure 3b). Furthermore, gradients at these resample locations need to be determined in order to perform shading. Projection-based reconstruction of function values is described in Section 3.1. The estimation of gradients directly from the projections is described in Section 3.2.

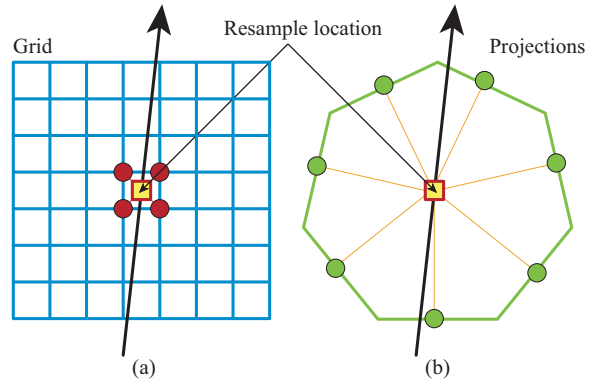


Figure 3: Difference of grid-based and projection-based resampling. (a) illustrates resampling along a ray on rectilinear volumetric data and (b) shows resampling along a ray directly from the filtered projections.

3.1. Data reconstruction

Data reconstruction from projection-based volumetric data is done using the Filtered Back-Projection algorithm. For simplicity we illustrate the Filtered Back-Projection in 2D based on a computed tomography scanning process using orthographic projection. Parallel projections are taken by measuring a set of parallel rays for a number of different angles. A projection is formed by combining a set of line integrals. The whole projection is a collection of parallel ray integrals as is given by $P_\theta(t)$ for a constant θ , see Figure 4. The line integrals are measured by moving an X-ray source and detector along parallel lines on opposite sides of the object. The

Filtered Back-Projection can be derived using the Fourier projection-slice theorem as follows:

The density function $f(x, y)$ can be expressed as:

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) e^{j2\pi(ux+vy)} dudv$$

where $F(u, v)$ denotes the two-dimensional Fourier transform of the density function $f(x, y)$. By moving from a Cartesian coordinate system in the frequency domain to a polar coordinate system, i.e., $u = w \cos \theta$, $v = w \sin \theta$, and $dudv = wdwd\theta$, we obtain:

$$f(x, y) = \int_0^{2\pi} \int_0^{\infty} F(w, \theta) e^{j2\pi w(x \cos \theta + y \sin \theta)} wdwd\theta$$

If we consider θ from 0 to π , the integral can be split as follows:

$$f(x, y) = \int_0^{\pi} \int_0^{\infty} F(w, \theta) e^{j2\pi w(x \cos \theta + y \sin \theta)} wdwd\theta + \int_0^{\pi} \int_0^{\infty} F(w, \theta + \pi) e^{j2\pi w(x \cos(\theta + \pi) + y \sin(\theta + \pi))} wdwd\theta$$

Since $F(w, \theta + \pi) = F(-w, \theta)$, the above expression can be written as:

$$f(x, y) = \int_0^{\pi} \left[\int_{-\infty}^{\infty} F(w, \theta) |w| e^{j2\pi wt} dw \right] d\theta$$

where $t = x \cos \theta + y \sin \theta$. By substituting $S_{\theta}(w)$ for the two-dimensional Fourier transform $F(w, \theta)$ the above integral can be expressed as:

$$f(x, y) = \int_0^{\pi} \int_{-\infty}^{\infty} S_{\theta}(w) |w| e^{j2\pi wt} dwd\theta$$

According to the Fourier projection-slice theorem $S_{\theta}(w)$ is the Fourier transform of $P_{\theta}(t)$. Let us define:

$$Q_{\theta}(t) = \int_{-\infty}^{\infty} S_{\theta}(w) |w| e^{j2\pi wt} dw \quad (1)$$

which is the inverse Fourier transform of $S_{\theta}(w) \cdot |w|$. As multiplication in the frequency domain corresponds to a convolution in the spatial domain, according to Equation 1, $Q_{\theta}(t)$ is obtained by high-pass filtering the measured projection $P_{\theta}(t)$. Other filters to reduce artifacts resulting from reconstruction can be applied, see [KS88].

In practice, the 2D density function $f(x, y)$ is discretely approximated by:

$$f(x, y) \approx \tilde{f}(x, y) = \frac{\pi}{K} \sum_{i=1}^K Q_{\theta_i}(x \cos \theta_i + y \sin \theta_i) \quad (2)$$

where Q_{θ_i} are the filtered projections. Thus, according to Equation 2 the density function can be reconstructed from a fixed number of projections. The Filtered Back-Projection algorithm is conventionally used for discrete tomographic reconstruction in order to obtain a rectilinear representation of the original density function. The formula in Equation 2 can also be considered as a resampling scheme to interpolate a density value at an arbitrary sample point. Previous

reconstruction techniques assume that accurate samples are available at the grid points. In order to maintain the same accuracy at any arbitrary sample location, we apply the Filtered Back-Projection to reconstruct the density value.




Figure 4: Parallel projection for a specific angle θ .

3.2. Derivative Estimation

In order to process or render volumetric data, often derivatives of the original density function are necessary. For example, for volume rendering the estimated gradients are used as surface normals to perform shading. In case of a grid based representation the straightforward way is to estimate the derivatives from a certain voxel neighborhood. To determine the gradient, common methods, such as intermediate difference gradient, central difference gradient, or higher order gradient estimation schemes are applied. In our case, computing the derivatives from a certain 3D neighborhood of samples requires to perform a large number of back-projections. Especially for higher order gradient estimation schemes, which need a large neighborhood of samples, the computational costs would be significantly high. However, the Filtered Back-Projection reconstruction scheme can also be exploited to compute derivatives.

For example the partial derivative \tilde{f}_x according to variable x can be expressed by using Newton's difference quotient:

$$\tilde{f}_x = \frac{\partial \tilde{f}(x, y)}{\partial x} = \lim_{\Delta x \rightarrow 0} \frac{1}{\Delta x} \left(\frac{\pi}{K} \left(\sum_{i=1}^K Q_{\theta_i}((x + \Delta x) \cos \theta_i + y \sin \theta_i) - \sum_{i=1}^K Q_{\theta_i}(x \cos \theta_i + y \sin \theta_i) \right) \right)$$

Substituting $t_i := x \cos \theta_i + y \sin \theta_i$ we obtain:

$$\begin{aligned} \tilde{f}_x &= \lim_{\Delta x \rightarrow 0} \frac{1}{\Delta x} \left(\frac{\pi}{K} \left(\sum_{i=1}^K Q_{\theta_i}(t_i + \Delta x \cos \theta_i) - \sum_{i=1}^K Q_{\theta_i}(t_i) \right) \right) \\ &= \lim_{\Delta x \rightarrow 0} \frac{1}{\Delta x} \left(\frac{\pi}{K} \sum_{i=1}^K Q_{\theta_i}(t_i + \Delta x \cos \theta_i) - Q_{\theta_i}(t_i) \right) \\ &= \frac{\pi}{K} \sum_{i=1}^K \lim_{\Delta x \rightarrow 0} \frac{1}{\Delta x} (Q_{\theta_i}(t_i + \Delta x \cos \theta_i) - Q_{\theta_i}(t_i)) \end{aligned}$$

The term

$$\lim_{\Delta x \rightarrow 0} \frac{1}{\Delta x} (Q_{\theta_i}(t_i + \Delta x \cos \theta_i) - Q_{\theta_i}(t_i)) \quad (3)$$

is the partial derivative of the projections Q_{θ_i} , but scaled with $\cos \theta_i$. We can therefore calculate the partial derivative \tilde{f}_x directly as sum of scaled derivatives of the projection data. Analogously, taking the difference quotient with respect to y we obtain:

$$\tilde{f}_y = \frac{\partial \tilde{f}(x,y)}{\partial y} = \frac{\pi}{K} \sum_{i=1}^K \lim_{\Delta y \rightarrow 0} \frac{1}{\Delta y} (Q_{\theta_i}(t_i + \Delta y \sin \theta_i) - Q_{\theta_i}(t_i))$$

It can be seen that applying Newton's difference quotient directly on the filtered projections is equivalent to applying Newton's difference quotient for the 2D density function $f(x,y)$. Moreover, any higher order derivative can be obtained by applying Newton's difference quotient multiple times.

Using Filtered Back-Projection for gradient estimation we expect higher accuracy than using the traditional gradient estimation schemes on the rectilinear grid. Consider central differences on the continuous reconstruction from a rectilinear representation. In order to calculate the gradient at an arbitrary sampling point six additional samples have to be interpolated. As interpolation usually causes loss of information, the introduced errors are accumulated in the estimated gradients. In contrast, using Filtered Back-Projection, the density values at the additional sample points are as accurate as the values at the grid points. Therefore, no interpolation error is introduced.

4. Results

In order to show the differences between grid-based and projection-based data reconstruction and gradient estimation we present results for different datasets. Projection-based datasets are acquired by a simulated Computed Tomography scanning process. A slice based scanning process produces a non-uniform error distribution along the z-axis. To avoid the non-uniform error distribution a simulation of a spiral scanning process was implemented. The three dimensional function used as input for the simulated scanning process is referred to as the gold standard. The volumetric function is defined by the projections and the Filtered Back-Projection algorithm. This projection-based volumetric function is assumed to be the initial point of the visualization pipeline. The difference between projection-based




Figure 5: Comparison of an iso-surface of the Marschner & Lobb function: (a) Analytically computed. (b) DVR rendering of a 64^3 grid. (c) DVR rendering of an eight times bigger grid (128^3). (d) D²VR from projection-based volumetric data (64 projections, each projection with a resolution of 64^2). Grids are reconstructed from 64 filtered projections, each projection with a resolution of 64^2 .

reconstruction (compare solid line in Figure 2) and grid-based reconstruction (compare dashed line in Figure 2) with different interpolation schemes was computed. To visually analyze the quality of D²VR a raycasting prototype for orthographic and perspective projection was implemented. For each pixel of the image plane, rays are cast through the volumetric space enclosed by the filtered projections. At each resample location the underlying 3D density function is reconstructed according to Equation 2. To perform shading gradient estimation is done directly on the projections. The central difference gradient is applied directly on the 2D projections and back-projected onto the resample location. The final color and opacity of the pixel are determined by the over-operator [PD84] in front-to-back order. For all experiments the stepsize along the ray was less than half the voxel size. For each experiment the stepsize used for DVR was equal to the stepsize used for D²VR. To evaluate our method we use the analytically defined Marschner & Lobb function as gold standard. The simulated scanning process was set up to compute 64 projections each with a resolution of 64^2 . The Filtered Back-Projection was used to reconstruct a regular rectilinear grid with a resolution of 64^3 . Using this grid different reconstruction schemes were used to reconstruct a rotated grid with the same resolution. The re-




Figure 6: CT scan of Carp (256x256x512): (a) DVR of original grid. (b) DVR of a 128x128x256 grid, reconstructed from 128 filtered projections, each projection with a resolution of 128x256. (c) D^2VR of projection-based volumetric data (128 projections, each projection with a resolution of 128x256). (d) close-up of DVR. (e) close-up of D^2VR .

sults were then compared to those of the projection-based reconstruction scheme. The differences between the grid-based reconstruction and the projection-based reconstruction decreases with more accurate grid-based reconstruction schemes. In Table 1 the difference of projection-based and grid-based reconstruction using nearest neighbor-, trilinear-, Catmull-Rom-, and polynomial interpolation is shown. We computed the average and the maximum of the absolute difference in percent of the range of values of the Marschner & Lobb function. To visualize the distribution of the error of projection-based reconstruction, we computed an iso-surface directly from the analytical Marschner & Lobb function. On the iso-surface a color coding was applied to visualize the differences between the analytical and the reconstructed value. Figure 8 shows the differences between the analytical values and the reconstructed values. Figure 8a

	NN	TRI	CAT	P3	P4	P5
AVG	2.05	1.15	0.88	0.64	0.59	0.59
MAX	76.9	37.1	42.1	36.2	38.8	35.8

Table 1: Average (AVG) and maximum (MAX) absolute difference of projection-based and grid-based reconstruction in percent of the data range, using different reconstruction schemes. Grid-based reconstruction schemes from left to right: nearest neighbor- (NN), trilinear- (TRI), Catmull-Rom- (CAT), and polynomial interpolation with polynomials of degree 3 (P3), 4 (P4) and 5 (P5).

shows the errors using trilinear interpolation for reconstruction on the grid (64^3) and Figure 8b shows the errors using Filtered Back-Projection for reconstruction directly from the projections. Green encodes low error, red encodes higher errors. In Figure 9 a similar color coding was used to visualize the angular error of the projection-based gradient estimation scheme. It is compared to grid-based gradient estimation based on central differences. Figure 9a shows the differences in degrees between the analytically computed gradients and the estimated gradients using central difference gradient estimation. Figure 9b shows the differences in degrees between the analytically computed gradients and the estimated gradients using our new projection-based gradient estimation method. Additionally, a grid with eight times more samples (128^3) was reconstructed from the projections and iso-surfaces were computed from the two different grids as well as directly from the analytical Marschner & Lobb function. The results can be seen in Figure 5. Figure 5a shows an analytical rendering. Figure 5b shows DVR of the 64^3 grid. Figure 5c shows DVR of the eight times bigger grid (128^3). And finally in Figure 5d D^2VR from 64 filtered projections, each projection with a resolution of 64^2 is shown.

Additionally we tested the behavior of the projection-based method with real world objects. Therefore, we use regular grids as gold standard. The resolution of the grids is high as compared to the resolution of the simulated scanner. The gold standard is undersampled to show the effect of using not sufficient resolution during the scanning process. Undersampling typically happens with real world objects. The

Stag beetle dataset with a resolution of $832 \times 832 \times 494$ voxels is used as the gold standard for the simulated scanning process. It is scanned resulting in 64 projections each with a resolution of 64^2 . From these projections a rectilinear grid is resampled with a resolution of 64^3 voxels. In Figure 1 a visual comparison of direct volume renderings of iso-surfaces using different reconstruction methods is given. Figure 1a to Figure 1d show different grid-based reconstruction methods while in Figure 1e the result of projection-based reconstruction can be seen. Note that the holes in the legs of the stag beetle get significantly smaller with D^2VR . In order to show the higher visual quality of D^2VR we also rendered the dataset of a carp with a semi-transparent transfer function. The rectilinear grid used as gold standard had a resolution of $256 \times 256 \times 512$. 128 projections were taken each with a resolution of 128×256 . From these projections a rectilinear grid was reconstructed with a resolution of $128 \times 128 \times 256$. In Figure 6a the high-resolution grid used as gold standard is shown. In Figure 6b a DVR of the grid-based data is shown using trilinear interpolation. Figure 6c shows a D^2VR of the projection-based data. In Figure 6d and Figure 6e close-ups of the grid-based DVR and D^2VR respectively are shown. Results in Figure 1 and Figure 6 clearly show that D^2VR produces higher quality than grid-based DVR using the same amount of pixels on all the projections as voxels on the regular grid. We also compare D^2VR with higher resolution grids. Rectilinear grids with twice the resolution in each direction as compared to the projections are computed, obtaining a grid with eight times more voxels than pixels on the projections. In Figure 7a an axial slice of the Carp dataset is

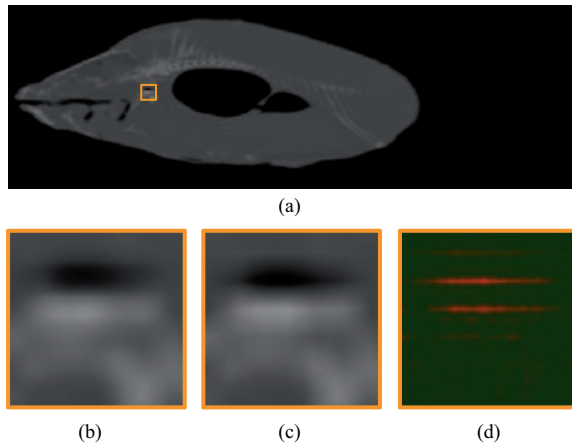


Figure 7: (a) Axial slice of the Carp dataset, (b) close-up using grid-based reconstruction with Catmull-Rom interpolation on a $256 \times 256 \times 512$ resolution grid, (c) close-up using projection-based reconstruction from 128 projections each with a resolution of 128×256 , (d) Absolute difference of data values of (b) and (c), dark green regions depicting zero difference, light red regions depicting a difference of up to 7.5% of the range of data values.

shown. In Figure 7b the grid-based reconstruction of a small region of the dataset is shown using Catmull-Rom interpolation on an eight times bigger grid. In Figure 7c the result of the projection-based reconstruction is shown. The absolute difference of the projection-based reconstruction and the grid-based reconstruction is shown in Figure 7d. While in most cases the used Catmull-Rom interpolation yields approximately the same results as the projection-based reconstruction, there are big differences when it comes to high frequency details. The maximum of the absolute difference of the depicted region is 7.5% of the range of the data values. A grid with twice the resolution as the projections using Catmull-Rom interpolation is not able to completely reconstruct the projection-based volumetric function.

To quantify the performance of D^2VR we measured the time, needed to reconstruct one million samples on 64 projections (each with a resolution of 64^2). Compared to trilinear interpolation on a regular grid (with a total of 64^3 voxels) projection-based reconstruction was approximately (but less than) 50 times slower. Compared to Catmull-Rom interpolation projection-based reconstruction was approximately (but less than) 10 times slower. To show the effect of increasing the resolution of the scanner, experiments with higher resolution datasets were performed. While the computation time of the projection-based approach increases linearly with the number of projections, the computation time of the grid-based approach increases slightly with the resolution of the grid. The experiments were repeated on different machines and with different numbers of samples showing no significant deviation.

5. Conclusion and Future Work

In this paper high-quality volume rendering of projection-based volumetric data has been introduced. It has been shown that volumetric raw data measured as a set of projections can be directly rendered without generating an intermediate grid-based volume representation. As our method avoids a lossy resampling step, it provides higher image quality than traditional direct volume rendering techniques do. Our novel projection-based gradient estimation scheme avoids the accumulation of interpolation errors. Traditional methods assume accurate samples at the grid points, while the accuracy of intermediate samples strongly depends on the quality of the applied interpolation method. In contrast, our approach provides accurate data values for arbitrary sample positions. The current performance of our prototype implementation is not comparable to that of well developed techniques. However, we believe that the increase in quality justifies further research in this direction. An interesting aspect of our work lies in the fact that scanners produce projection data. Investigating approaches that directly manipulate and process the projection data without any intermediate representation seems to be an interesting research direction in the future. Our work was inspired by the practical

tomographic reconstruction problem. Its theoretical significance is the demonstration of an alternative volume representation. In our future work we plan to explore other gridless volume representations, which are not necessarily related to the physical constraints of current scanning devices. For example, in order to achieve full uniformly distributed reconstruction quality, projection planes with uniformly distributed orientation might be applied. Although the adaptation of the Filtered Back-Projection algorithm to such a geometry requires further research, it would lead to a direction independent high-quality volume reconstruction scheme.

6. Acknowledgements

The work presented in this publication has been funded by the PVG project (FWF-P18547N13). The stag beetle dataset has been provided by G. Glaeser, University of Applied Arts Vienna. Data scanning is courtesy of J. Kastner, Wels College of Engineering. The carp dataset is courtesy of Michael Scheuring from Computer Graphics Group, University of Erlangen, Germany.

References

- [AMVG05] ARTNER M., MÖLLER T., VIOLA I., GRÖLLER M. E.: High-quality volume rendering with resampling in the frequency domain. In *Proceedings of Eurographics / IEEE VGTC Symposium on Visualization* (2005), pp. 85–92.
- [CCF94] CABRAL B., CAM N., FORAN J.: Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proceedings of IEEE Symposium on Volume Visualization* (1994), pp. 91–98.
- [CCM99] CHEN Q., CROWNOVER R., M. WEINHOUS: Subunity coordinate translation with Fourier transform to achieve efficient and quality three-dimensional medical image interpolation. *Med. Phys.* 26, 9 (1999), 1776–1782.
- [CT99] COX R., TONG R.: Two- and three-dimensional image rotation using the FFT. *IEEE Transactions on Image Processing* 8, 9 (1999), 1297–1299.
- [KS88] KAK A. C., SLANEY M.: *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics (Proceedings of SIGGRAPH '87)* (1987), 163–169.
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, Vol. 8, No. 3 (1988), 29–37.
- [LM04] LI A., MUELLER K.: Methods for efficient, high quality volume resampling in the frequency domain. In *Proceedings of IEEE Visualization* (2004), pp. 3–10.
- [Mal93] MALZBENDER T.: Fourier volume rendering. *ACM Transactions on Graphics*, Vol. 12, No. 3 (1993), 233–250.
- [ML94] MARSCHNER S., LOBB R.: An evaluation of reconstruction filters for volume rendering. In *Proceedings of IEEE Visualization* (1994), pp. 100–107.
- [MMK*98] MÖLLER T., MUELLER K., KURZION Y., MACHIRAJU R., YAGEL R.: Design of accurate and smooth filters for function and derivative reconstruction. In *Proceedings of IEEE Symposium on Volume Visualization* (1998), pp. 143–151.
- [MMMY97] MÖLLER T., MACHIRAJU R., MUELLER K., YAGEL R.: A comparison of normal estimation schemes. In *Proceedings of IEEE Visualization* (1997), pp. 19–26.
- [MN88] MITCHELL D., NETRAVALI A.: Reconstruction filters in computer graphics. In *Proceedings of SIGGRAPH* (1988), pp. 221–228.
- [NCN*97] NIKLASON L. T., CHRISTIAN B. T., NIKLASON L. E., KOPANS D. B., CASTLEBERRY D. E., OPSAHL-ONG B. H., C. E. LANDBERG P. J. S., GIARDINO A. A., MOORE R., ALBAGLI D., DEJULE M. C., FITZGERALD P. F., FOBARE D. F., GIAMBATISTA B. W., KWASNICK R. F., LIU J., LUBOWSKI S. J., POSSIN G. E., RICHOTTE J. F., WEI C. Y., WIRTH R. F.: Digital tomosynthesis in breast imaging. *Radiology* 205 (1997), 399–406.
- [OS89] OPPENHEIM A. V., SCHAFER R. W.: *Discrete-Time Signal Processing*. Prentice Hall Inc., Englewood Cliffs, 2nd edition, 1989.
- [PD84] PORTER T., DUFF T.: Compositing digital images. *Computer Graphics* 18, 3 (1984), 253–259.
- [Rus92] RUSS J.: *The Image Processing Handbook*. CRC Press, 1992.
- [TC99] TONG R., COX R.: Rotation of NMR images using the 2D chirp-z transform. *Magn. Reson. Med.* 41 (1999), 253–256.
- [TG00] THEUSSL T., GRÖLLER M. E.: Mastering windows: Improving reconstruction. In *Proceedings of IEEE Symposium on Volume Visualization* (2000), pp. 101–108.
- [UTY95] UNSER M., THEVENAZ P., YAROSLAVSKY L.: Convolution-based interpolation for fast, high-quality rotation of images. *IEEE Transactions on Image Processing* 4, 10 (1995).
- [Wes90] WESTOVER L.: Footprint evaluation for volume rendering. *Computer Graphics (Proceedings of SIGGRAPH '90)* (1990), 144–153.
- [ZPvBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: EWA volume splatting. In *Proceedings of IEEE Visualization 2001* (2001), pp. 29–36.




Figure 8: Color encoded differences between analytical value and (a) the reconstructed value using trilinear interpolation on the grid (64^3), (b) the reconstructed value using Filtered Back-Projection.




Figure 9: Color encoded differences in degrees between analytically computed gradients and (a) the estimated gradients using central difference gradient estimation on the grid, (b) the estimated gradients using our new projection-based gradient estimation method.