

Analysis of the HDAF for Interpolation and Noise Suppression in Volume Rendering

Kristoffer Andersson,^{1†} Ioannis A. Kakadiaris,¹ Manos Papadakis,² Donald J. Kouri,³ and David K. Hoffman⁴

¹ Visual Computing Lab, Department of Computer Science

² Department of Mathematics

³ Department of Physics, Mathematics, and Chemistry
University of Houston, Houston, TX 77204

⁴ Department of Chemistry, Iowa State University, Ames, Iowa 50011

Abstract

In this paper, we evaluate the HDAF (Hermite Distributed Approximating Functionals) family of interpolation and derivative functions, with respect to their accuracy for reliable volume rendering, and compare them with other interpolation and derivative estimation filters. We utilize several different evaluation methods, both analytical and experimental. The former includes the order of decay of the global error, the local spatial error, and the behavior of the filters in the frequency domain. In the experimental part, visualizations of both synthetic and medical data are produced and studied. We show that the HDAFs exhibit superior behavior if the volumetric data are distorted by high frequency noise, and perform well under noise free conditions. This due to their ability to adjust the range of recovered frequencies.

Categories and Subject Descriptors (according to ACM CCS): G.1.0 [General]: Error analysis, G.1.1 [Interpolation]: Interpolation formulas, G.1.2 [Approximation]: Special function approximations, I.3.7 [Three-Dimensional Graphics and Realism]: Ray-casting, shading, and I.4.5 [Reconstruction]: Series expansion methods.

1. Introduction

Data have been represented visually throughout human history because this format lends itself to intuitive interpretation and enables communication of information. Another task pertaining to the visualization field is to seek understanding and extract information from data. The evolution of medical imaging modalities like CT and MRI has made it possible to acquire three-dimensional data of various types for diagnosis purposes. The most common method of presenting this information is to view the data slice by slice. Oftentimes, better understanding can be attained if the data are visualized all at once, through *volume rendering*. One commonly used volume rendering method is *ray-casting*, and this method will be used for the volume visualizations throughout this

paper. The ray-casting algorithm constructs images by casting rays through the volume data, one for each pixel in the image, and retrieving values along the rays for classification and determination of the pixel color. If the known data points do not coincide with the needed sample points, interpolation of the values is required. Additionally, to be able to interpret the third dimension, which is lost after the projection onto the two-dimensional image, we need to add shading. The shading algorithm requires the derivative information, which has to be estimated as well, through derivative interpolation. The quality of the rendering depends on the interpolation and derivative estimation methods used. The accuracy of the latter has great importance as the shading of an object gives indications for its shape and orientation.

In medical applications, the accuracy of volume visualizations is extremely crucial. Removal of important information may result in false negative diagnosis, while addition of artifacts and distortions can make it difficult to extract useful

[†] This work was supported in part by the National Science Foundation Career Award CISE 9985482 and UH Texas Learning and Computation Center.

information and draw correct conclusions. Furthermore, the volumetric data might be corrupted by noise, which makes the extraction of accurate information even harder. The interpolation method should preferably discard the noise and provide a faithful reconstruction of the original signal.

The goal of this paper is to evaluate the Hermite Distributed Approximating Functionals, with respect to their accuracy for reliable volume rendering, and compare them with other interpolation filters. The comparisons are performed using different evaluation methods for the analysis of various properties. In Section 2, the mathematical foundations for interpolation and perfect reconstruction are presented. Section 3 surveys previous work in the field of visualization and filter analysis. The filters studied in this paper will be presented in Section 4 and analyzed in Section 5. Section 6 contains the experimental results. Finally, in Section 7, we summarize our findings.

2. Background

The interpolation operation can be written as a convolution of the discrete samples f_k with a continuous interpolation filter $w(x)$, or:

$$f_r(x) = \sum_{k=-\infty}^{\infty} f_k \cdot w\left(\frac{x}{\Delta} - k\right), \quad (1)$$

where Δ is the sampling distance. Similarly, convolution with a continuous derivative filter, which can be created by computing the derivative of the interpolation filter, under certain additional hypotheses produces the derivative of the function².

The interpolation filters usually satisfy the *interpolating constraint*, which requires that the interpolation operation does not alter the known values at the sample points, or:

$$f_r(k\Delta) = f_k, \forall k \in \mathcal{Z} \quad (2)$$

This translates into a filter that takes a unit value at the origin and vanishes at multiples of the unit distance from it. One problem with the interpolating constraint is that the same level of accuracy is not achieved both on and off the sample points. This leads to poor accuracy in the derivative at the sample points. Moreover, the acquired data are usually contaminated by noise, so there is no need to preserve them exactly. By regarding the interpolating constraint as a soft constraint, all points can be treated the same, which has been called the *well-tempered property*²¹.

According to the sampling theorem by Shannon¹¹, any bandlimited function can be recovered uniquely from its samples if it was sampled at or above the *Nyquist rate*, which is twice the highest frequency content of the function. The ideal interpolation filter that produces the perfect reconstruction is the sinc function. However, the sinc function is never used in practise since it has a very slow decay at infinity. While truncation of the function will make it feasible to use,

it will also result in significant reconstruction errors^{2, 14}. The Fourier transform of a truncated sinc contains large lobes, which causes aliasing and ringing effects. One solution is to approximate the sinc, but minimize the artifacts and other reconstruction errors.

3. Previous Work

Strang and Fix¹² showed a number of equivalent ways to retrieve the *approximation order*, which describes the decay rate of the approximation error as the sampling step becomes smaller. Thévenaz et al.¹⁴ presented the approximation order of various interpolation filters. After performing several experiments, they concluded that it was the most important index of quality. Möller et al.⁸ studied the local spatial behavior of the error through the Taylor series expansion of the reconstruction process and used it to classify different cubic interpolation and derivative filters. The analysis was later expanded to produce a spatial filter design scheme⁹. Theußl et al.¹³ made use of the same analysis method in their study of various windowing functions that can be used to bound the sinc/cos filters. Bentum² compared the frequency response of various interpolation and derivative methods in his research, as deviations from the ideal filter lead to distortions in the volume rendering. Marschner and Lobb⁷ proposed a number of metrics to measure the quality of interpolation filters from their frequency representation.

Experimental comparisons involve reconstructing functions from sampled data³, enhancing the resolution of an image^{2, 15}, and performing volume rendering on analytical or medical volumetric data^{1, 2, 7, 8, 9, 10, 13, 18}. The inspection of the rendered images is subjective and conclusions can be drawn only if the differences are very evident, like in the comparison performed by Neumann et al.¹⁰ between their new gradient estimation method and the simple and poor central difference. Their method, which is based on linear regression, can also be used for smoothing purposes, and therefore produces much better results. To make the comparison of visualizations more objective, Volpi¹⁸ used five image level metrics in his evaluation. Kim et al.⁶ try to overcome the limitations of image level comparisons with a data level approach using intermediate rendering information. Möller et al.⁹ computed and visualized the angular error in the reconstruction of an analytical function to compare different filters and filter settings.

4. Interpolation Techniques

In this section, the interpolation filters that we analyzed and used in the experiments are presented. The filters are all *separable*, which means that the three-dimensional interpolation filters are constructed from three one-dimensional functions as follows:

$$w(x, y, z) = w(x)w(y)w(z) \quad (3)$$

The partial derivatives of separable multidimensional filters are constructed the same way, using two one-dimensional interpolation filters and one one-dimensional derivative filter².

4.1. HDAF Approximation

Hoffman and Kouri^{4,5} introduced the *Distributed Approximating Functionals* (DAFs) as approximations of the *Dirac delta function* in the time domain, for the purpose of approximating continuously differentiable functions from discrete samples that had possibly been corrupted by noise. Both interpolative and non-interpolative variants of the DAF exist. *Hermite Distributed Approximating Functional* (HDAF), the first DAF developed, is not interpolative and will not preserve the input data. Instead, it can achieve similar order of accuracy both at and between sample points, also referred to as the *well-tempered property*. This has been proven to be useful in padding, periodic extension, and non-periodic extrapolation²⁰. The family of HDAFs $\delta_N(x)$, all having a regularity of C^∞ , is given by²¹:

$$\delta_N(x - x' | \sigma) = \frac{e^{-z^2}}{\sqrt{2\pi\sigma}} \sum_{n=0}^N \left(-\frac{1}{4}\right)^n \frac{1}{n!} H_{2n}(z) \quad (4)$$

and the Fourier transform is given by:¹⁹

$$\hat{\delta}_N(\xi | \sigma) = e^{-\frac{\xi^2 \sigma^2}{2}} \sum_{n=0}^N \frac{(\xi^2 \sigma^2)^n}{2^n n!}, \quad (5)$$

where $z = (x - x')/\sigma\sqrt{2}$, $H_n(x)$ is a *Hermite polynomial* of degree n , and N and σ control the behavior of the HDAF. The Gaussian factor e^{-z^2} ensures that the function decays fast enough in the spatial domain. This makes it possible to truncate it at a point where the loss will not alter the approximation too much. Therefore, the σ controls the effective width of the function. The HDAF approximation is given by:

$$f(x) \approx f_{HDAF}(x) = \Delta \sum_j f(x_j) \delta_N(x - x_j) \quad (6)$$

$$f^{(l)}(x) \approx f_{HDAF}^{(l)}(x) = \Delta \sum_j f(x_j) \delta_N^{(l)}(x - x_j), \quad (7)$$

With a suitable choice of the parameters N and σ , it is possible to estimate the function and a finite set of derivatives to a desired level of accuracy²¹.

The HDAFs chosen for this evaluation are the HDAF₁, HDAF₄, and HDAF₆ (see Figure 1), together with their derivatives for the gradient estimation. The σ parameter is initially assigned a value at which the HDAFs are close to interpolative when $\Delta = 1$. These settings ($\sigma = 0.598413420$, 0.981772018 , 1.169944988 , respectively) will be changed later to study the effects on the reconstruction when the HDAFs are more approximating.

4.2. Piecewise Polynomial Interpolation

Traditionally, ray-casting implementations have used simple interpolation methods, like linear interpolation and finite-

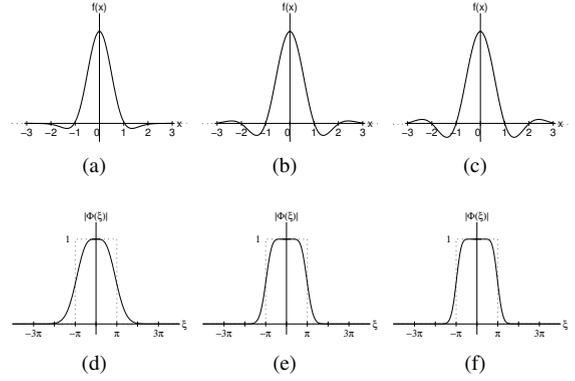


Figure 1: (a) HDAF₁ ($\sigma = 0.598413420$), (b) HDAF₄ ($\sigma = 0.981772018$), and (c) HDAF₆ ($\sigma = 1.169944988$) interpolation filters and their Fourier transform (d-f).

difference gradient estimation, due to the computational effort required in volume rendering and the limitations in computer speed. With special purpose hardware, interactive speeds can be accomplished even when using larger filters, like the *Catmull-Rom* filter and its derivative². This piecewise cubic polynomial filter is given by⁹:

$$w_{C-R}(x) \begin{cases} \frac{3}{2}|x|^3 - \frac{5}{2}|x|^2 + 1, & 0 \leq |x| < 1 \\ -\frac{1}{2}|x|^3 + \frac{5}{2}|x|^2 - 4|x| + 2, & 1 \leq |x| < 2 \\ 0, & 2 \leq |x|. \end{cases} \quad (8)$$

It belongs to the class of C^1 functions and has a support of $[-2, 2]$.

Thévenaz et al.¹⁴ found that the *Moms functions* (Maximal Order and Minimal Support) gave the best combination of quality and speed. Two of the presented Moms functions have been chosen for this comparison: the *cardinal cubic B-spline* and the *cardinal cubic o-Moms*. The filters are interpolative and have infinite support, but the same results can be achieved by prefiltering the data and applying the non-interpolative cubic B-spline and cubic o-Moms with compact support. The prefiltering converts the samples to the correct coefficients and can be implemented efficiently through recursion^{17, 16}. The cubic B-spline is given by⁹:

$$\beta^3(x) \begin{cases} \frac{1}{2}|x|^3 - |x|^2 + \frac{2}{3}, & 0 \leq |x| < 1 \\ -\frac{1}{6}|x|^3 + |x|^2 - 2|x| + \frac{4}{3}, & 1 \leq |x| < 2 \\ 0, & 2 \leq |x| \end{cases} \quad (9)$$

and the cubic o-Moms can be constructed using a cubic B-spline and its second derivative as follows¹⁴:

$$\text{o-Moms}^3(x) = \beta^3(x) + \frac{1}{42} \frac{d^2}{dx^2} \beta^3(x) \quad (10)$$

5. Analysis

Three analytical comparison methods were utilized in the analysis of the filters: the approximation order, the local spatial error, and the frequency response of the filters.

5.1. Approximation Order

The rate at which the reconstructed function converges to the true function as the sampling distance decreases determines how good the interpolation operator is. When the sampling step becomes sufficiently small, the behavior of the error can be described by two variables, the approximation order L and the approximation constant C , as shown below¹⁴:

$$\epsilon(\Delta) = \|f - f_\Delta\|_{L_2} = C\Delta^L \|f^{(L)}\|_{L_2}, \text{ as } \Delta \rightarrow 0. \quad (11)$$

The approximation order estimates the global decay of the error, where a higher value of L gives a faster decay, and offers a first classification of the interpolation filters in the form of a single value. It can be retrieved from one of the Strang-Fix conditions¹², as follows:

$$\begin{cases} \Phi(0) = 1 \\ \Phi^{(n)}(2\pi k) = 0, \quad k \in \mathcal{Z} \text{ and } n \in [0, L-1], \end{cases} \quad (12)$$

where $\Phi(\xi)$ is the interpolation function in the frequency domain (see Table 1). Similar computations can be made for the derivative filters using the following condition:

$$\begin{cases} \Phi^{(1)}(0) = 1 \\ \Phi^{(n)}(2\pi k) = 0, \quad k \in \mathcal{Z} \text{ and } n \in [0, L] \end{cases} \quad (13)$$

The approximation order for the interpolation and derivative filters under study are given in Table 2.

Table 1: Interpolation filters in the frequency domain.

Name	$\Phi(\xi)$
Catmull-Rom	$[3\text{sinc}(\frac{\xi}{2\pi}) - 2\cos(\frac{\xi}{2})][\text{sinc}(\frac{\xi}{2\pi})]^3$
Cardinal cubic B-spline	$\frac{[\text{sinc}(\frac{\xi}{2\pi})]^4}{\frac{2}{3} + \frac{2}{6}\cos(\xi)}$
Cardinal cubic o-Moms	$\frac{(1 - \frac{1}{42}\xi^2)[\text{sinc}(\frac{\xi}{2\pi})]^4}{\frac{13}{21} + \frac{8}{21}\cos(\xi)}$
HDAF _N	$e^{-\frac{\xi^2\sigma^2}{2}} \sum_{n=0}^N \frac{(\xi^2\sigma^2)^n}{2^n n!}$

[†] Neither the Fourier transform of the HDAF (or its derivative) nor any derivatives of the Fourier transform is zero at multiples of 2π . However, the condition can be met arbitrarily close by changing the σ parameter.

Table 2: Approximation order of the interpolation and derivative filters.

Name	L_{int}	L_{deriv}
Catmull-Rom	3	2
Cardinal cubic B-spline	4	3
Cardinal cubic o-Moms	4	3
HDAF _N	$2N+2^\dagger$	$2N+2^\dagger$

5.2. Local Spatial Error Analysis

Möller et al.⁸ presented a method for classification and evaluation of interpolation and derivative filters. The method studies local spatial behavior of the approximation error, and the analysis is based on the Taylor series expansion of the reconstruction process. By substituting the series for the sample points into Equation 1, we can write the reconstructed function $f_r(x)$ as a weighted sum of the original function itself $f^{(0)}(x)$ and its derivatives $f^{(n)}(x)$, $n > 0$, as follows⁸:

$$f_r(x) = \sum_{n=0}^N a_n(\tau) f^{(n)}(x) + r_{N,i}(\tau)$$

$$a_n(\tau) = \frac{\Delta^n}{n!} \sum_{k=-\infty}^{\infty} (k-\tau)^n w(\tau-k) \quad (14)$$

$$r_{N,i}(\tau) = \frac{\Delta^{(N+1)}}{(N+1)!} \sum_{k=-\infty}^{\infty} f^{(N+1)}(\xi_{k,i})(k-\tau)^{(N+1)} w(\tau-k)$$

where a_n are called approximation coefficients, $r_{N,i}$ is the error term, with $\xi_{k,i} \in [x, (k+i)\Delta]$, and τ is the offset from the sample points, with $0 \leq \tau < 1$. The approximation coefficients only depend on the interpolation filter, and can therefore be used to classify it. If N in Equation 14 is the largest number for which the following is true:

$$a_n(\tau) \begin{cases} 1 & n=0 \\ 0 & 1 \leq n \leq N \end{cases} \quad (15)$$

then, the interpolation filter is classified as having an approximation order of $N+1$. Similar observations can be made for the derivative interpolators. If N is the largest number for which the corresponding approximation coefficients can be written as:

$$a_n(\tau) \begin{cases} 0 & n=0 \\ 1 & n=1 \\ 0 & 2 \leq n \leq N \end{cases} \quad (16)$$

then, the derivative filter has an approximation order of N . When the sampling step is small, the second non-zero coefficient has superior control over the behavior of the error. By looking at its placement and how it varies with τ , we can analyze the local spatial approximation error of the filters.

In our analysis the sampling distance was set to $\Delta = 1$ and

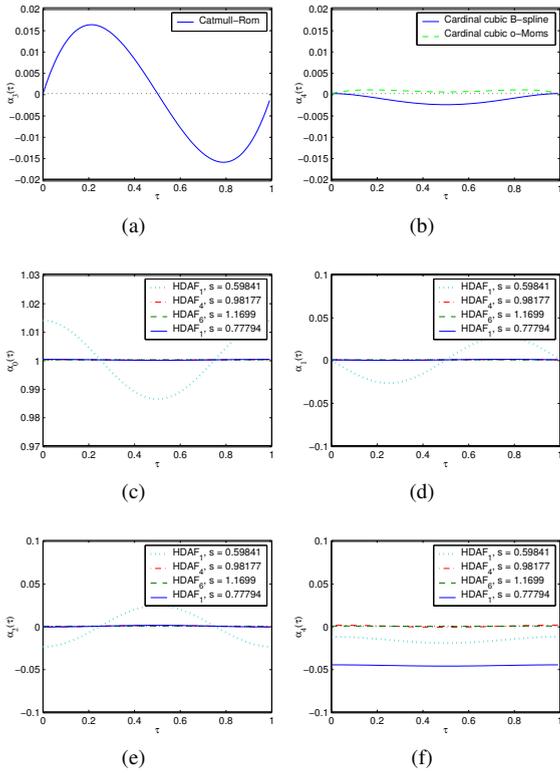


Figure 2: Coefficients of the interpolation filters.

the extent of the cardinal cubic B-spline, the cardinal cubic o-Moms, and the HDAFs was set to $[-20, 20]$, as a larger extent did not produce any significant change. From the graph in Figure 2(a), we observe that the Catmull-Rom filter has an approximation order equal to 4 at $\tau = 0.5$. This is important information if we are interested in interpolating in the middle of two sample points only. The largest error for the Catmull-Rom filter is found at $\tau = 0.25$ and $\tau = 0.75$. Figures 2(c-f) depict the approximation coefficients for the HDAF interpolation filters. Here we observe how the HDAF₁ converges to the approximation order given in Table 2. From Figure 2(c), it is clear that the HDAF₁, with $\sigma = 0.598413420$, is not appropriate even to reconstruct constant functions, although the approximation order for HDAF₁ was stated to be 4. After increasing the σ parameter by 30% to the value 0.777937446, better results are obtained. The side-effect of the increase can be seen in Figure 2(f) where the error term has become larger. For higher values of N , the frequency response of the HDAF has a faster decay, which is why the approximation coefficients for HDAF₄ and HDAF₆ are better.

The local spatial error can also be analyzed for the derivative filters. The Catmull-Rom derivative filter exhibits the opposite behavior to the interpolation filter (see Figure 3(a)). Where the latter had a higher approximation order, the for-

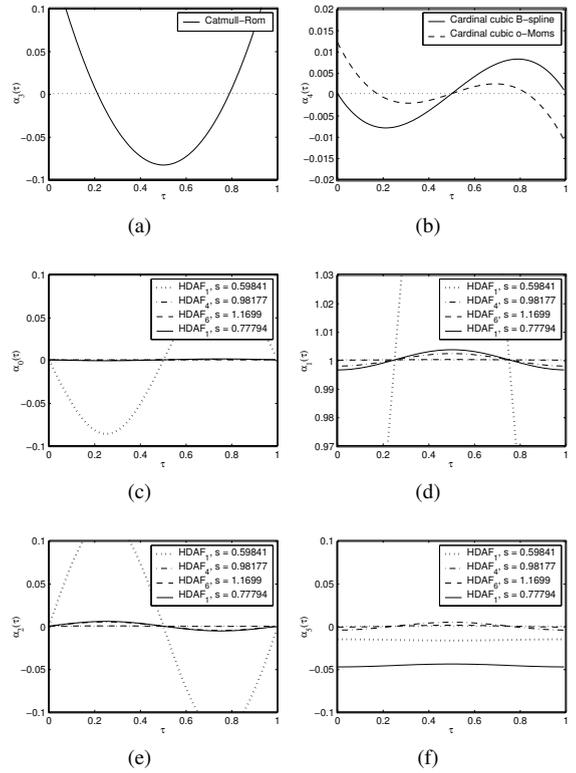


Figure 3: Coefficients of the derivative filters.

mer exhibits larger errors, and vice versa. This complies with the idea that better estimations of the function value will lead to worse estimations of the derivative. Figures 3(c-f) are the graphs of the approximation coefficients for the HDAF derivative filters. The problem with the HDAF₁, where $\sigma = 0.598413420$, can be observed in this set of plots too. The approximation order does converge to 4, as stated in Section 5.1, but only as σ increases. For higher values of N , the frequency response of the HDAF approximates the frequency response of the ideal filter and therefore its values are closer to zero at multiples of 2π other than the origin. This explains the better behavior of the approximation coefficients for HDAF₄ and HDAF₆.

5.3. Behavior of the Frequency Response

The reconstruction process can be studied in the frequency domain as well. This extraction is carried out by applying a lowpass filter. The goal is to approximate the ideal lowpass filter (i.e., the box filter) without its disadvantages. To observe how well the interpolation filters in this study approximate the ideal filter, we can compare their frequency response with the box filter (see Figure 4). The overlap percentage is found in Table 4 and the area outside of the box filter in Table 3. The Catmull-Rom and HDAF₁ with

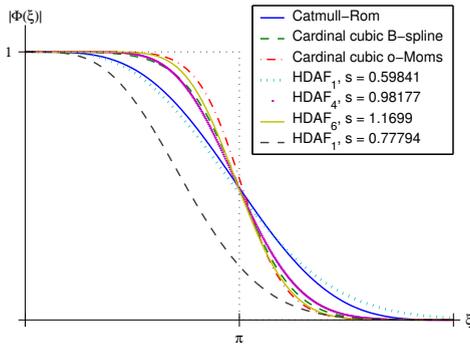
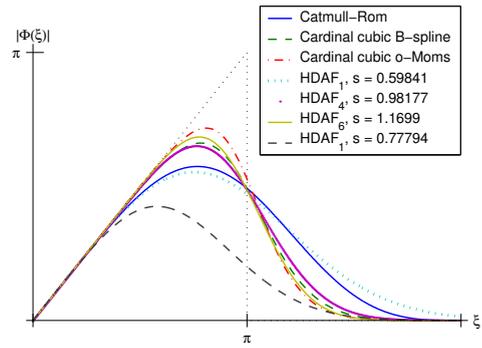

Figure 4: Frequency response of the interpolation filters.

Figure 5: Frequency response of the derivative filters.

Table 3: Area of the interpolators outside the box filter.

Name	Area: $\pi - 2\pi$	$2\pi - 3\pi$	$3\pi - 4\pi$
Catmull-Rom	0.3987	0.0088	0.0122
Cardinal cubic B-spline	0.2515	0.0078	0.0044
Cardinal cubic o-Moms	0.2312	0.0082	0.0078
HDAF ₁ , $\sigma = 0.598413420$	0.4467	0.0031	0.0000
HDAF ₁ , $\sigma = 0.777937446$	0.1112	0.0000	0.0000
HDAF ₄ , $\sigma = 0.981772018$	0.2766	0.0000	0.0000
HDAF ₆ , $\sigma = 1.169944988$	0.2316	0.0000	0.0000

$\sigma = 0.598413420$ form a group with similar frequency responses. They exhibit the same behavior close to the origin, but the former has a faster decay and vanishes to zero at 2π , where the frequency response of the latter is about 0.0069. This will lead to introduction of more aliasing artifacts for the HDAF₁. The result of increasing the σ parameter is better behavior outside the box filter, but also a smoothing effect, while it covers less of the box area. The rest of the filters form a second group, where the cardinal cubic o-Moms exhibits the best behavior in Figure 4, followed by the HDAF₆, the cardinal cubic B-spline, and HDAF₄. Better coverage of the box filter area is accomplished with these filters, and the transition is sharper, which results in less aliasing.

Similar analysis can be performed for the derivative re-

Table 4: Overlap percentage of (a) the box filter area for the interpolation filters, and (b) the ramp filter area for the derivative filters.

Name	Overlap: box filter	ramp filter
Catmull-Rom	86.5%	79.2%
Cardinal cubic B-spline	91.0%	85.5%
Cardinal cubic o-Moms	93.7%	89.6%
HDAF ₁ , $\sigma = 0.598413420$	85.2%	77.2%
HDAF ₁ , $\sigma = 0.777937446$	72.9%	58.7%
HDAF ₄ , $\sigma = 0.981772018$	90.7%	84.8%
HDAF ₆ , $\sigma = 1.169944988$	92.1%	87.0%

Table 5: Area of the derivative filters outside the ramp filter.

Name	Area: $\pi - 2\pi$	$2\pi - 3\pi$	$3\pi - 4\pi$
Catmull-Rom	1.5061	0.0713	0.1258
Cardinal cubic B-spline	0.9013	0.0679	0.0443
Cardinal cubic o-Moms	0.8157	0.0727	0.0781
HDAF ₁ , $\sigma = 0.598413420$	1.7528	0.0209	0.0000
HDAF ₁ , $\sigma = 0.777937446$	0.4058	0.0001	0.0000
HDAF ₄ , $\sigma = 0.981772018$	1.0031	0.0000	0.0000
HDAF ₆ , $\sigma = 1.169944988$	0.8221	0.0000	0.0000

construction filters. The ideal derivative filter in the frequency domain is the ramp filter. It is depicted together with the frequency response of the derivative filters in Figure 5. The coverage of the ramp filter is listed in Table 4 and their area outside the ramp can be found in Table 5. From the graph and the tables we observe that the HDAF₁ with $\sigma = 0.777937446$ does not approximate the ideal filter very well. The Catmull-Rom and the HDAF₁ with $\sigma = 0.598413420$ offer better coverage, but also extend further outside the ramp – especially the latter. The HDAF₁ filter has a value of about 0.043 at 2π , which will lead to more aliasing. The derivative of the cardinal cubic o-Moms gives the best approximation in Figure 5, followed by the HDAF₆, the cardinal cubic B-spline, and HDAF₄.

6. Experimental Results

The previous section concentrated on analytical studies of the filters and provided theoretical results for the filter properties. For our volume rendering experiments we used two data sets: an MRI (Magnetic Resonance Imaging) data set and a synthetic data set, constructed by sampling an analytic function. In the first subsection, the volume rendering of the MRI data for visual comparisons is addressed. The second subsection presents the reconstruction of the synthetic data, along with the reconstruction errors. The sampling distance Δ was assumed to be 1 for all data sets to simplify the computations and avoid the need for normalization. This tran-

scribes into a highest allowable frequency content of $\xi_l = \pi$ in the function. The recursive prefiltering step for the cardinal cubic B-spline and the cardinal cubic o-Moms uses mirroring of the data at the boundaries^{15,17}. Mirroring of the data was the solution to the boundary condition in the interpolation step as well.

The HDAFs have infinite support and have to be truncated at some point. The extent for the HDAF interpolation and derivative filters can be found in Table 6. A larger support did not produce any significant change in the volume rendering results. Because the extent has the greatest impact on the computational cost of a filter¹⁴, the figures in Table 6 indicate that the HDAFs require more computational effort than the other filters. Even when look-up tables with pre-computed filter values are used, the cost of the interpolation operation in volumes is $O(W^3)$, where W is the filter extent. The cost of evaluating the piecewise polynomial or the HDAFs on request is $O(Wn)$, where n is the polynomial degree of the filter. The disadvantage with the HDAFs are is larger size and higher polynomial degree.

Table 6: Extents for the HDAF interpolation and derivative filters.

Name	σ	Extent: $\delta_N(x \sigma)$	$\delta_N^{(1)}(x \sigma)$
HDAF ₁	0.598413420	[-3, 3]	[-4, 4]
HDAF ₁	0.777937446	[-4, 4]	[-4, 4]
HDAF ₁	0.897620130	[-5, 5]	[-5, 5]
HDAF ₄	0.981772018	[-4, 4]	[-5, 5]
HDAF ₄	1.472658027	[-6, 6]	[-7, 7]
HDAF ₆	1.169944988	[-5, 5]	[-6, 6]
HDAF ₆	1.754917482	[-7, 7]	[-8, 8]

6.1. MRI Data

The MRI data set is the *Volume Rendering Test Dataset Volume II* from the University of North Carolina, also used by Bentum² and Möller et al.^{8,9}. It depicts a head with the top part of the skull removed so that the surface of the brain is visible. The images of the MRI data were generated using the ray-casting algorithm, with the step size along the rays set to 0.05 times the voxel length. At each step the value was interpolated using the filters from Section 4, and an iso-surface classification function was used to produce a surface with 600 as the isovalue. The rays were terminated after this single threshold and the normal was computed with the corresponding derivative filters for the shading.

Figure 6 contains the visualization results using the different filters. The HDAFs produce smooth results in Figures 6(g-i), while the rest of the images contain more details (or noise). The HDAFs clearly offer a wide range of possibilities, from close to interpolative reconstructions to coarse scale approximations.

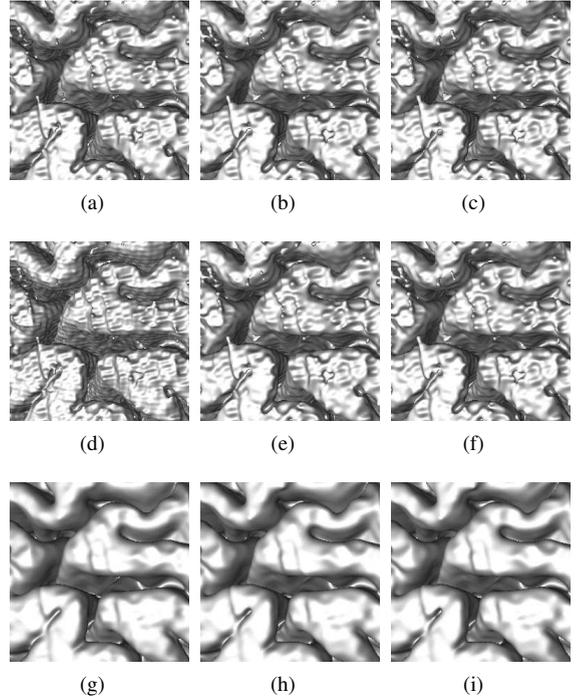


Figure 6: Volume visualization of the MRI data set using (a) Catmull-Rom, (b) cardinal cubic B-spline, (c) cardinal cubic o-Moms, (d) HDAF₁ ($\sigma = 0.598413420$), (e) HDAF₄ ($\sigma = 0.981772018$), (f) HDAF₆ ($\sigma = 1.169944988$), (g) HDAF₁ ($\sigma = 0.897620130$), (h) HDAF₄ ($\sigma = 1.472658027$), and (i) HDAF₆ ($\sigma = 1.754917482$).

6.2. Volume Rendering of Synthetic Data

The synthetic data set was constructed by sampling an analytic function, introduced by Marschner and Lobb⁷, with a sampling rate of 20 samples per unit distance over the range $-1 \leq x, y, z \leq 1$, and scale the values by 255. This produces a $41 \times 41 \times 41$ volume data set. In the reconstruction, however, the sampling distance Δ is assumed to be 1 to simplify the computations, as mentioned before. This is equivalent to dilating the function and increasing the sampling range by 1:20. The Marschner-Lobb function is given by:

$$f(x, y, z) = \frac{1}{2} - \frac{2}{5} \sin\left(\frac{\pi}{2}z\right) + \frac{1}{10} \cos\left(12\pi \cos\left(\frac{\pi}{2}\sqrt{x^2 + y^2}\right)\right)$$

Marschner and Lobb claimed that this function is bandlimited and that the sampling rate given above is just above the Nyquist rate. This produces a very interesting function to reconstruct, while the periodic repetitions of the Fourier transform of the function are only just separated and a considerable amount of frequency content lies at the border between the repetitions.

An isosurface was produced using the ray-casting algorithm, with half of the maximum function value as the iso-

value. The step size along the rays was set to 0.005 times the voxel length. In Figure 7(a) the goal surface to be reproduced is depicted, and Figure 7(b) gives a close up view of its central part. Since it sometimes can be hard to assess the performance of the filters from the volume rendering results, the reconstruction errors have been computed and visualized as well. The first error is the *interpolation error*, computed as the absolute difference between the interpolated value and the true function value. The second error is the *angular error*, given by the angular difference between the interpolated surface normal and the true normal. In the error plots, black represents perfect reconstruction, while white depicts an interpolation error greater or equal to 4 or an angular error of at least 15 degrees, respectively. Values in between produce a grey level proportional to the error.

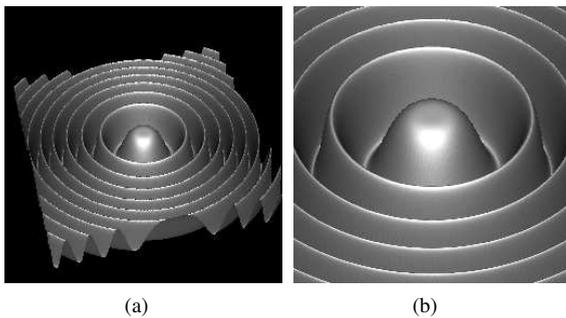


Figure 7: Isosurface rendering of the Marschner-Lobb function. (a) Overview and (b) central part.

6.2.1. Without Noise

Figures 8(a-c) depict the results of the volume renderings of the synthetic data using the different filters, along with the plots of the interpolation error and angular error. Statistical information about the errors can be found in Table 7. The HDAF₁ with $\sigma = 0.598413420$ introduces a number of aliasing artifacts and gives the poorest reconstruction. With $\sigma = 0.777937446$, the HDAF₁ smooths out the function and clearly produces the largest interpolation error. The shape of the reconstructed surface can still be recognized, but it does not have as strong an upward and downward inclination or height and depth as the goal surface, which leads to the angular errors in Figure 8(c). The cardinal cubic B-spline, the cardinal cubic o-Moms, HDAF₄, and HDAF₆ all seem to have similar reconstruction results.

Overall, the interpolation error images depict a zebra pattern with areas containing good estimations alternating with patches of large errors. The darker areas lie along the slices of sample points, as the interpolating constraint forces the reconstruction to be close to the real function there. In between the slices the error is larger, and thereof the zebra pattern. Stripe patterns are seen in the angular error images too.

6.2.2. With Noise

While most interpolation filters were designed to give highest accuracy under the constraint of preserving the sample points, the HDAFs are used under the assumption that the data set might be corrupted by noise. The requirement of preserving the sample points is then no longer justified. If the noise satisfies two conditions given below, it can be filtered by adjusting the width of the frequency response of the HDAFs. The first condition is that the noise be random. It will then contain mostly high frequency components, while the real objects are smooth and contain mostly lower frequencies. The second is that the noise be *additive* and consequently totally independent of the real data. The width of the frequency response is adjusted by changing the value of σ , where a higher value leads to attenuation of high frequencies and thereby removes the noise.

For this reason, random noise with a gaussian probability distribution (mean set to 0 and standard deviation equal to 2) was added to the sample points. Since the sampling rate used before was just above the Nyquist rate, the frequency content of the function and the noise would blend. To get a better separation of their frequency content, the sampling rate was increased by 50% to 30 samples per unit distance, resulting in a 61x61x61 volume data set. The σ parameter for the HDAFs was also increased by 50% from the original settings, to filter away the noise but still reconstruct the function. Making the same changes to the width of the frequency response of the other filters would lead to aliasing artifacts, due to the side lobes found in their frequency response¹.

Figure 8(d) depicts the performance of the different filters in the recovery of the function, while the interpolation error and angular error plots can be found in Figure 8(e) and Figure 8(f), respectively. The Catmull-Rom, cardinal cubic B-spline and cardinal cubic o-Moms are all interpolative and are forced to preserve the noisy data points exactly. The errors for these filters are therefore concentrated in a spotty pattern. The HDAFs, on the other hand, suppress the high frequency noise due to the increase of the σ parameter. Since there are no special points anymore, the errors are more evenly distributed.

7. Conclusions

Both the analytical and the experimental analyses suggest that the approximation order of a filter provides a reasonable first indication of the accuracy of the reconstruction. For further analysis and comparison, the behavior of the local spatial error and the frequency response must be studied. For example, deeper analysis was vital for the selection of the σ parameter for the HDAFs₁.

The Catmull-Rom filter provides an acceptable tradeoff between the quality of the volume renderings and the computational effort. It has small support, which makes the visualization process fast. However, better results can be ac-

complished with the cubic B-spline and the cubic o-Moms, where the only additional computational cost comes from a simple prefiltering step. Other interpolation methods (e.g. windowed sinc/cosc) should be included in future studies to get a more comprehensive comparison.

The HDAF family of functions offers a set of parameterized interpolation and derivative filters, with high control over the reconstruction process and the type of functions that can be recovered. In noise-free conditions, they can produce results with the same quality as the cardinal cubic B-Spline and cardinal cubic o-Moms. Although the extent is larger, implying greater computational effort and time, their excellent noise suppression qualities makes them the superior choice when additive noise is present. In the future, we plan to explore how the HDAFs perform when the data is corrupted by multiplicative noise, compared to other existing noise filtering methods.

References

1. K. Andersson. Interpolation of Volumetric Data for Volume Visualization. Master's thesis, University of Houston, Houston TX, USA, August 2002.
2. M. Bentum. *Interactive Visualization of Volume Data*. PhD thesis, University of Twente, Enschede, the Netherlands, June 1996.
3. H. Choi and R. Baraniuk. Interpolation and Denoising of Nonuniformly Sampled Data Using Wavelet-Domain Processing. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, Phoenix AZ, USA, March 1999.
4. D. K. Hoffman and D. J. Kouri. Distributed Approximating Functional Theory for an Arbitrary Number of Particles in a Coordinate System Independent Formalism. *Journal of Physical Chemistry*, 97:4984–4988, 1993.
5. D. K. Hoffman, N. Nayar, O. Sharafeddin, and D. J. Kouri. On an Analytic Banded Approximation for the Discretized Free Propagator. *Journal of Physical Chemistry*, 95:8299, 1991.
6. K. Kim, C. M. Witterbrink, and A. Pang. Extended Specifications and Test Data Sets for Data Level Comparisons of Direct Volume Rendering Algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 7(4):299–317, 2001.
7. S. R. Marschner and R. J. Lobb. An evaluation of Reconstruction Filters for Volume Rendering. In *Proceedings of Visualization '94*, pages 100–107. IEEE CS Press, October 1994.
8. T. Möller, R. Machiraju, K. Mueller, and R. Yagel. Classification and Local Error Estimation of Interpolation and Derivative Filters for Volume Rendering. In *Proceedings 1996 Symposium on Volume Visualization*, pages 71–78, 1996.
9. T. Möller, K. Mueller, Y. Kurzion, R. Machiraju, and R. Yagel. Design of Accurate and Smooth Filters for Function and Derivative Reconstruction. In *Proceedings 1998 Symposium on Volume Visualization*, pages 143–151, 1998.
10. L. Neumann, B. Csébfalvi, A. König, and M. E. Gröller. Gradient Estimation in Volume Data using 4D Linear Regression. In *Computer Graphics Forum (Proceedings of Eurographics 2000)*, volume 19, pages 351–358, 2000.
11. C. Shannon. Communication in the Presence of Noise. In *Proceedings of the IRE*, volume 37, pages 10–21, January 1949.
12. G. Strang and G. Fix. A Fourier Analysis of the Finite Element Variational Method. In G. Geymonat, editor, *Constructive Aspects of Functional Analysis*, pages 795–840, Rome, Italy, 1973. Edizioni Cremonese.
13. T. Theußl, H. Hauser, and M. E. Gröller. Mastering Windows: Improving Reconstruction. In *Proceedings 2000 Symposium on Volume Visualization*, pages 101–108, October 2000.
14. P. Thévenaz, T. Blu, and M. Unser. Interpolation Revisited. In *IEEE Transactions on Medical Imaging*, volume 19, pages 739–758, July 2000.
15. M. Unser, A. Aldroubi, and M. Eden. Fast B-Spline Transforms for Continuous Image Representation and Interpolation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(3):277–285, March 1991.
16. M. Unser, A. Aldroubi, and M. Eden. B-Spline Signal Processing: Part I - Theory. *IEEE Transactions on Signal Processing*, 41(2):821–833, February 1993.
17. M. Unser, A. Aldroubi, and M. Eden. B-Spline Signal Processing: Part II - Efficient Design and Applications. *IEEE Transactions on Signal Processing*, 41(2):834–848, February 1993.
18. W. H. Volpi. Comparison of Interpolation Methods for Volume Rendering. Master's thesis, University of Houston, Houston TX, USA, May 2001.
19. G. Wei, H. Wang, D. J. Kouri, M. Papadakis, I. A. Kakadiaris, and D. K. Hoffman. On the Mathematical Properties of Distributed Approximating Functionals. *Journal of Mathematical Chemistry*, 30:83–107, 2001.
20. D. S. Zhang. *Distributed Approximating Functionals, Partial Differential Equations, and Signal Processing*. PhD thesis, University of Houston, Houston TX, USA, August 1999.
21. D. S. Zhang, G. W. Wei, D. J. Kouri, D. K. Hoffman, M. Gorman, A. Palacios, and G. H. Gunaratne. Integrating the Kuramoto-Sivashinsky Equation in Polar Coordinates: Application of the Distributed Approximating Functional Approach. *Physical Review E*, 60(3):3353–3360, 1999.

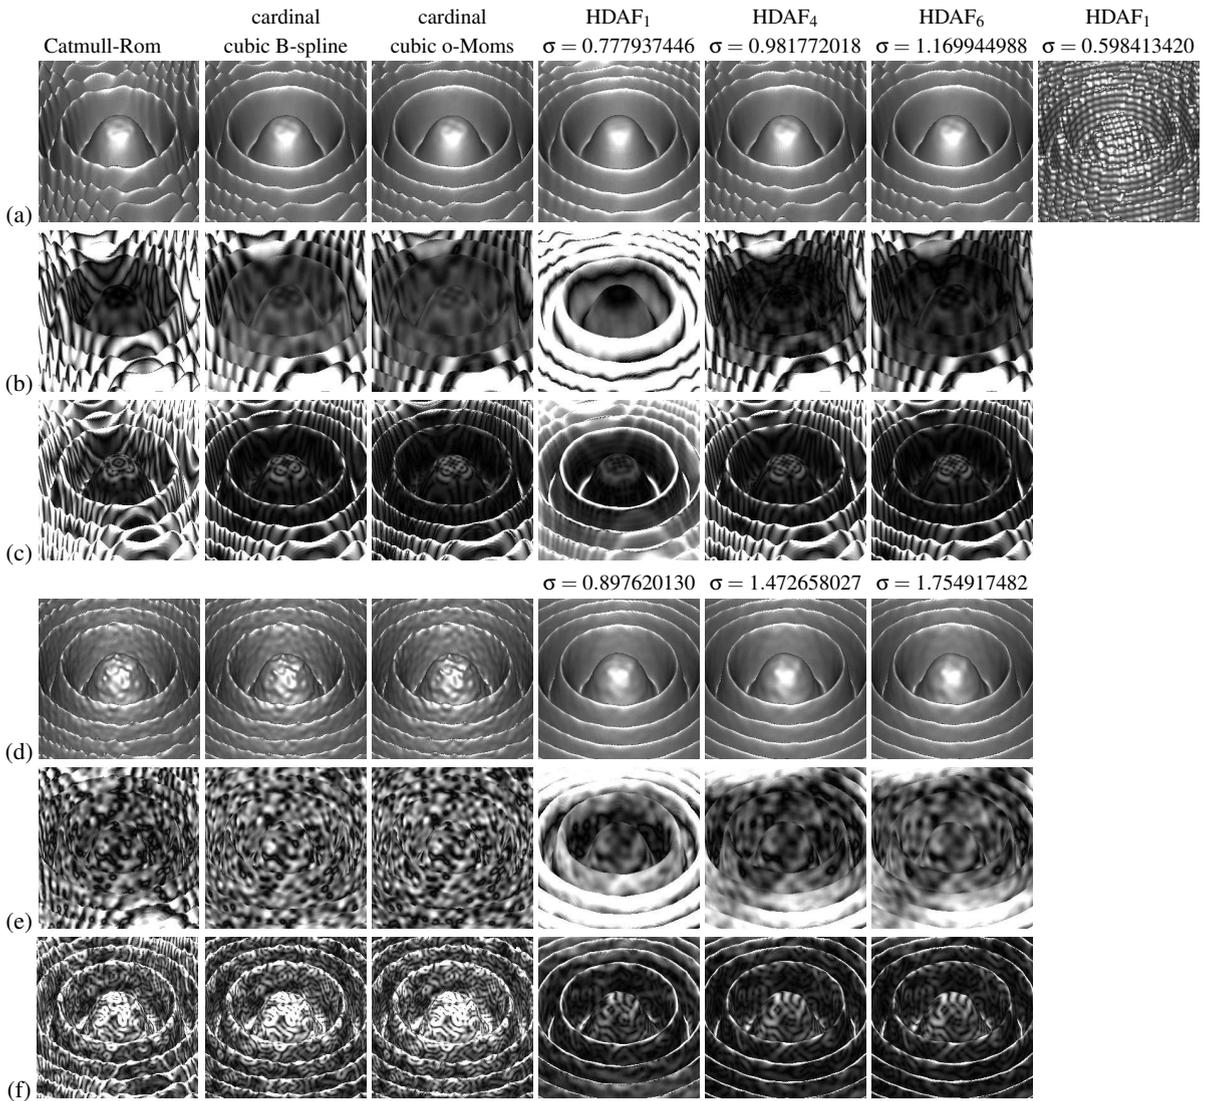


Figure 8: Volume rendering of reconstructed function, interpolation error, and angular error using noise free Marschner-Lobb data (a, b, and c, respectively) and corrupted data (e, d, and f, respectively).

Table 7: Average interpolation and angular error (in degrees) in the reconstruction, along with the standard deviation.

Name	Without noise				With noise			
	Interpolation error		Angular error		Interpolation error		Angular error	
	Mean	Std dev.	Mean	Std dev.	Mean	Std dev.	Mean	Std dev.
Catmull-Rom	2.9483	3.1530	12.7876	18.1015	1.6317	1.1630	8.0124	8.8786
Cardinal cubic B-spline	2.0188	2.1189	5.9674	11.8388	1.6432	0.8880	6.5657	6.3843
Cardinal cubic o-Moms	1.5604	1.6928	4.6545	9.8680	1.5444	0.8958	7.0196	6.6821
HDAF ₁	4.7078	4.0845	8.2500	9.2905	3.2231	2.3672	3.9731	3.9458
HDAF ₄	1.6710	2.2278	6.8339	12.3800	1.9119	1.5868	3.2057	3.8532
HDAF ₆	1.3588	1.8725	4.9280	10.4299	1.9073	1.3816	3.1245	3.8905