Extensions of the Zwart-Powell Box Spline for Volumetric Data Reconstruction on the Cartesian Lattice

Alireza Entezari, and Torsten Möller, *Member, IEEE*

Abstract—In this article we propose a box spline and its variants for reconstructing volumetric data sampled on the Cartesian lattice. In particular we present a tri-variate box spline reconstruction kernel that is superior to tensor product reconstruction schemes in terms of recovering the proper Cartesian spectrum of the underlying function. This box spline produces a C^2 reconstruction that can be considered as a three dimensional extension of the well known Zwart-Powell element in 2D. While its smoothness and approximation power are equivalent to those of the tri-cubic B-spline, we illustrate the superiority of this reconstruction on functions sampled on the Cartesian lattice and contrast it to tensor product B-splines. Our construction is validated through a Fourier domain analysis of the reconstruction behavior of this box spline. Moreover, we present a stable method for evaluation of this box spline by means of a decomposition. Through a convolution, this decomposition reduces the problem to evaluation of a four directional box spline that we previously published in its explicit closed form [8].

Index Terms—Volumetric data interpolation, reconstruction, box splines.

---◆---

1 INTRODUCTION

It is a well known fact that optimal lattice sampling – in the sense of capturing the maximum amount of information using the sparsest possible sampling pattern – demands non-Cartesian lattices. Such sampling strategies, however, have not been widely used. The reason, in part, is that there are no adequate tools to deal with such lattices. A number of recent publications have been trying to address this problem with great success [8, 23, 9, 26, 25]. However, the Cartesian lattice will remain attractive due to its simple structure that allows one to apply a tensor product of one dimensional analysis and processing tools. The power of the dimensionality reduction will remain the major reason that the Cartesian lattice will be the preferred tool in numerical algorithms.

Tensor product reconstruction schemes on the volumetric data sampled on the Cartesian lattice are popular and easy to understand. However, in this paper we will build the case, that one should consider applying more generic reconstruction kernels for the reconstruction and rendering of data sampled on the Cartesian lattice.

For reconstructing in the space of "bandlimited functions", the design of reconstruction methods aims at constructing kernels that are as close as possible to the ideal reconstruction function, which is represented in the Fourier domain by a box function in 1D. In the space domain this is the sinc function which has infinite support. Typically we only approximate this non-compactly supported kernel with a compact kernel in the spatial domain. This approximation tends to show a smooth drop-off from a DC-value of one to a value of zero at the Nyquist frequency. This is of course far from ideal, but the best one can do. In computer graphics and visualization we seek reconstruction kernels with spatial properties such as smoothness and reconstruction accuracy.

Filter design typically is done in 1D and extended through a tensor product. As expected, a tensor product extension of this 1D box function is the ideal reconstruction kernel for higher dimensional Cartesian lattices. However, for non-ideal reconstruction kernels, their tensor-product extensions show a much more pronounced attenuation of frequencies near the Nyquist region at the corners of the tensor-product box compared to the axis aligned Nyquist region. On the flip side, aliased spectra in the diagonal directions disappear quicker, than the

- *School of Computing Science, Simon Fraser University, E-mail: {aentezar,torsten}@cs.sfu.ca.*

aliased spectra aligned along the axis directions. Hence reconstruction along the axis receives extra treatment.

(a) bi-linear (b) ZP element
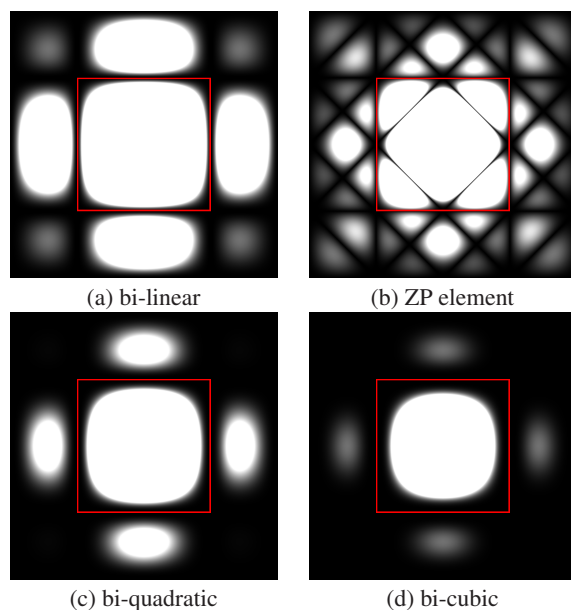
(c) bi-quadratic (d) bi-cubic

Fig. 1. Comparison of the frequency response of (a) a tensor product kernel (the Hat function that is the *bi-linear* B-spline); (b) the ZP element; (c) the tensor-product quadratic B-spline as well as the (d) tensor-product cubic B-spline. For convenience, responses higher than 0.005 are set to white.

This phenomenon is illustrated in Figure 1, where all magnitudes larger than 0.005 have been set to white. Figure 1a shows the simple bi-linear interpolant (which is a first order B-spline). While only frequencies fairly close to the Nyquist region (drawn as a red box) have been almost completely diminished (the stop-band), the preservation of large parts of the aliased spectra is typically unacceptable. The sensible idea in tensor product scheme is to convolve the bi-linear interpolant with a box function, which yields the bi-quadratic reconstruction kernel (or the second order B-spline) which is shown in Figure 1c. We notice that the influence of the aliased spectra is greatly diminished (especially in the diagonal regions), but the stop-band region has been enlarged as well - more so at the corners of the primary spectrum, and less so at the axis aligned edge regions. This behavior

is even more pronounced for the third order (bi-cubic) B-spline shown in Figure 1d.

This situation can be improved only by designing non-separable schemes. The idea is simple – we enhance the corner regions by simply convolving the original axis-aligned bi-linear interpolant with a box function that is rotated by 45 degrees and scaled properly, such that its zero crossings are at the center of the aliased spectra. According to Strang-Fix [24] this will improve the numerical accuracy of the reconstruction. As can be seen in Figure 1b this yields a reconstruction spectrum that enhances the pass band. The aliased spectra in the axis-aligned directions are now greatly reduced as well under the trade-off, that the aliased spectra in the diagonal directions are somewhat enhanced. The net effect is that the reconstruction of the underlying function will show much fewer grid-aligned artifacts than before – a desirable property. The particular kernel shown in Figure 1b is also known as the Zwart-Powell element.

In other words, there is something to be gained by designing reconstruction schemes in a non-separable fashion. The impulse for this research came from design of reconstructions for general lattices, especially optimal lattices, such as BCC and FCC lattices and is based on box splines. This paper will design box splines, based on the 2D Zwart-Powell element and devise a non-separable kernel for the Cartesian lattice. This kernel captures the isotropic nature of the underlying function much better, without having to resample the underlying function on any new lattice. The advantage of our proposed reconstruction scheme can be demonstrated through the representation of planar surfaces within a Cartesian lattice as shown in Figure 6 as well as in Figure 13.

After a review of previous work in Section 2, we will give a brief introduction to box splines in Section 3 as well as a derivation of the 3D Zwart-Powell element and its variations. Section 4 will discuss implementation details which is followed by results in Section 5.

2 RELATED RESEARCH

The design of reconstruction filters is a very rich area of signal processing and approximation theory. The literature in the domain of signal processing typically uses constraints in the Fourier Domain in order to guide the filter design process, which often results in discrete filter weights (or "taps"). The goal of these methods is to design a "good" approximation to the ideal filter (a box in 1D in the frequency domain) which is compact (with a finite support) in the spatial domain. Sophisticated algorithms have been developed to enhance the behavior in what is known as the pass band, the stop band and the transition band (see the Parks-McClellan as well as Butterworth filter designs [20, 4] as well as the max-flat filter design of Dutta Roy [7]).

These filters resample a given regular sampling pattern on a new regular sampling pattern. While these approaches are important in many applications (such as compression, multi-resolution analysis and others), graphics and rendering algorithms often require continuous reconstruction filters, that allow the computation of the underlying function values at arbitrary offsets. Other features of the reconstruction process often include smoothness and accuracy criteria – properties best specified in the spatial domain. Approaches from approximation theory often use constraints in the spatial domain in order to design continuous kernels, that allow the reconstruction of samples with arbitrary offsets from the given sampling lattice [11, 16, 18].

While all these design approaches yield one-dimensional filters, in the areas of image processing and volume rendering we need to reconstruct higher-dimensional functions. The common approach is to design the filter in one dimension and then to extend the filter into higher dimensions through a separable extension (tensor product) or through a spherical extension. While separable extensions are justified due to the separability of the sampling lattice, spherical extensions often suffer from the fact that it is difficult to guarantee zero-crossings of the frequency response at all of the replicas of the spectrum. These zero-crossings are crucial in order to guarantee accuracy of the reconstruction process [24]. These problems of spherical extensions remain in all sampling lattices.

In the field of approximation theory there has been a great deal of effort on approximating multidimensional functions. In particular, Birkhoff and de Boor discuss multivariate splines that are tensor products of one dimensional splines [3, 2, 1]. When approximating a function over $\mathbb{R}^n$ with a spline of *degree* s, these splines are constructed as a tensor product of splines of a maximum degree s in any of the variables and the tensor product is of total degree ns.

Zwart [27] considered a different family of multivariate splines which aren't necessarily a tensor product, but rather they are piecewise polynomial functions where each piece is a polynomial of total degree s. The well known Zwart-Powell (ZP) element was one of the example splines where the reconstruction kernel is not a tensor product of univariate splines. Rather it has non-degenerate partitions and on each partition it is a polynomial of second degree. This element gained popularity as an example of a nontrivial bivariate box spline [6]. Powell also devised the corresponding subdivision scheme to reconstruct functions with continuous first derivatives that was widely used for contour drawing programs [22]. The significance of the ZP element is that it achieves a C^1 smoothness with only a polynomial of degree two in 2D. A tensor product solution can only do the job with a polynomial of degree four. Hence the ZP element stands out among the low degree splines that achieves the maximal smoothness.

Smooth piece-wise polynomial reconstruction of functions are used in the numerical solution of partial differential equations through construction of Finite Elements and in modeling of smooth surfaces in Computer Aided Geometric Design. One and two dimensional elements are widely studied and implemented throughout the literature that provide convincing solutions in terms of numerical accuracy and smoothness for reconstruction.

3 BOX SPLINES

Box splines offer a mathematically elegant framework for constructing a class of elements with flexible shape and support suited for various reconstruction purposes. The general topic of box splines is rather intricate and a general survey of results on the topic have been gathered in [6]. We begin by briefly introducing box splines and state their properties that will be useful in our exposition.

3.1 Definitions and Properties of Box Splines

A box spline is specified by a set of direction vectors that determine the shape of the support of the box spline and also the smoothness and accuracy of the functions spanned by that spline. These direction vectors are usually gathered in a matrix that specifies the box spline. A box spline in $\mathbb{R}^s$ is specified by $n \geq s$ vectors in $\mathbb{R}^s$ that are columns of the box spline matrix $\mathbf{\Xi} = [\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \ldots, \boldsymbol{\xi}_n]$. The support of the box spline is all points $\boldsymbol{x} \in \mathbb{R}^s$ such that $\boldsymbol{x} = \mathbf{\Xi}\boldsymbol{t}$ where $\boldsymbol{t} \in \mathbb{R}^n$ and $0 \leq t_i < 1$ for $1 \leq i \leq n$. In other words, the support of the box spline is the convex combination of these direction vectors.

The simplest box spline is constructed by $n = s$ vectors. This box spline is the characteristic function of its support area.

$$M_{\mathbf{\Xi}}(\boldsymbol{x}) = \begin{cases} \frac{1}{|\det \mathbf{\Xi}|} & \text{where } \boldsymbol{x} = \mathbf{\Xi}\boldsymbol{t} \text{ and } \boldsymbol{t} \in [0,1)^n \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

It is clear from Equation 1 that this box spline is discontinuous at the boundary of its support. The simplest one dimensional box spline amounts to the *boxcarr* function which is the indicator function for the interval $[0, 1)$.

For the general case where there are more direction vectors $n > s$, the box splines are defined recursively:

$$M_{[\mathbf{\Xi},\boldsymbol{\xi}]}(\boldsymbol{x}) = \int_0^1 M_{\mathbf{\Xi}}(\boldsymbol{x} - t\boldsymbol{\xi})dt \quad (2)$$

This inductive definition implies that starting from the base case as in Equation 1 the indicator function is smeared *along* the new direction vector. Hence, the convolution of two box splines is yet another box spline:

$$M_{[\mathbf{\Xi}_1,\mathbf{\Xi}_2]}(\boldsymbol{x}) = (M_{\mathbf{\Xi}_1} * M_{\mathbf{\Xi}_2})(\boldsymbol{x}) \quad (3)$$

A box spline is a piece-wise polynomial of degree at most $n - s$. Moreover, let ρ be the minimal number of vectors that if removed from Ξ the remaining vectors would *not* span $\mathbb{R}^s$. Then $M_\Xi \in C^{\rho-2}$, where C^m is the space of m-times differentiable smooth functions. The Fourier transform of a box spline can be shown to be:

$$\hat{M}_\Xi(\boldsymbol{\omega}) = \prod_{\xi \in \Xi} \frac{1 - \exp(-i\boldsymbol{\xi} \cdot \boldsymbol{\omega})}{i\boldsymbol{\xi} \cdot \boldsymbol{\omega}} \qquad (4)$$

As an example in 2D, the simplest box spline is specified by:

$$\Xi_0 = \begin{bmatrix} \boldsymbol{\xi}_1 & \boldsymbol{\xi}_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

which is the indicator function of the unit square $[0, 1)^2$. This box spline generates a discontinuous spline space over $\mathbb{R}^2$.

3.2 The Zwart-Powell Element

Adding a direction vector of $\boldsymbol{\xi}_3 = \begin{bmatrix} 1 & 1 \end{bmatrix}^\top$ to Ξ_0 smears the unit square across its diagonal. This is illustrated in Figure 2(a-b). As the basic box spline is a constant function on the unit square, the result of smearing it along the diagonal produces a linear order box spline that is represented by $[\Xi_0, \boldsymbol{\xi}_3]$. The support of this box spline is illustrated in Figure 2(b). This box spline is a two dimensional piece-wise polynomial of degree one over its support. This box spline generates a C^0 spline function space, as $\rho = 2$.

Adding one more direction vector to the above box spline produces a quadratic order box spline. The choice of $\boldsymbol{\xi}_4 = \begin{bmatrix} -1 & 1 \end{bmatrix}^\top$ produces a symmetric octagonal shape of this box spline which is known as the Zwart-Powell element. The process of this convolution is illustrated in Figure 2(c-d). The ZP element contains two
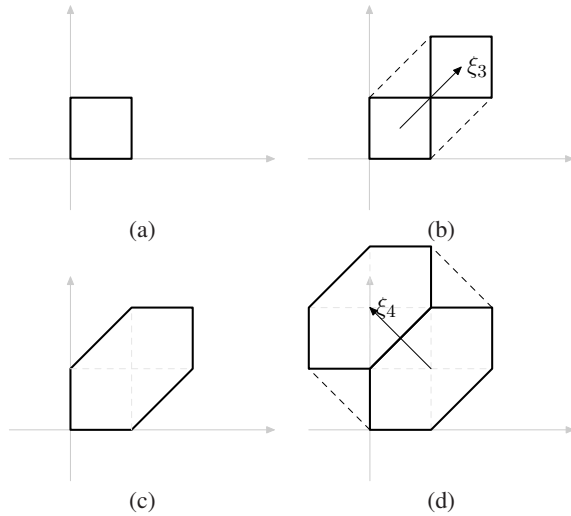
Fig. 2. Construction of the Zwart-Powell element from the linear box spline

dimensional polynomials that are only of second degree and yet they achieve a C^1 reconstruction in 2D ($\rho = 3$). As has been argued in the introduction section and Figure 1, this element exhibits more isotropy in the stop-band and a better Cartesian behavior in the pass-band than the elements that are constructed using tensor product extensions.

3.3 The Seven Directional Box Spline

Our choice of the box spline for 3D volume rendering is motivated by the improved frequency behavior that the ZP element exhibits in 2D. Since the ZP element is obtained by direction vectors that are forming the square (which is the Voronoi cell of the 2D Cartesian lattice) and its diagonals, our 3D construction is also a box spline whose direction vectors form a cube (that is the Voronoi cell of the 3D Cartesian lattice) and its four diagonals, see Figure 3. While one could extend the

2D ZP element by adding other directions in 3D, we have chosen the diagonal directions specifically in order to get a good approximation to a spherically symmetric reconstruction kernel. Our choice of box spline directions is represented by the matrix:

$$\Xi = \begin{bmatrix} 1 & 0 & 0 & 1 & -1 & -1 & 1 \\ 0 & 1 & 0 & -1 & 1 & -1 & 1 \\ 0 & 0 & 1 & -1 & -1 & 1 & 1 \end{bmatrix} \qquad (5)$$
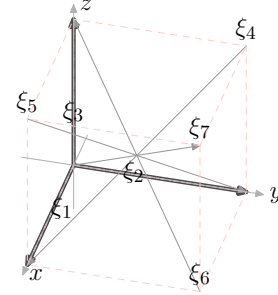
Fig. 3. The seven directions of the box spline. The first three directions are axis aligned and form a cube. The last four are the antipodal diagonal directions.
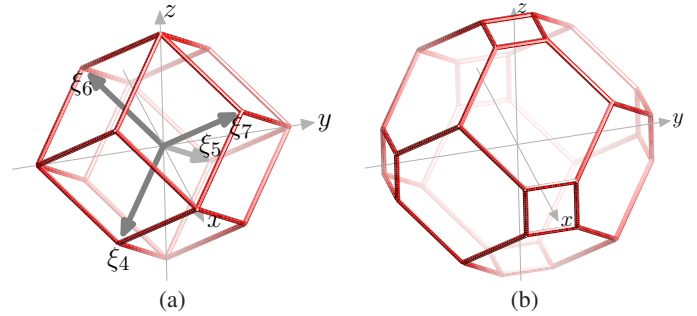
Fig. 4. (a) The four diagonal directions constitute a box spline whose support is a rhombic dodecahedron. (b) The support of the seven directional box spline is a truncated rhombic dodecahedron which is the convolution of the indicator function of a cube and rhombic dodecahedron.

3.3.1 Space Domain Analysis

The seven directional box spline offers a C^2 reconstruction since $\boldsymbol{\xi}_1, \boldsymbol{\xi}_4$ and $\boldsymbol{\xi}_7$ are co-planar and $\rho = 7 - 3 = 4$. This smoothness parallels that of the tri-cubic B-spline. However the seven directional box spline only consists of polynomials of degree four, but the tri-cubic B-spline consists of degree nine polynomials. Peters exploits the advantage of this box spline for constructing a curvature continuous surface through tracing its zero sets [21]. He also describes a subdivision scheme that can be used to approximate this box spline at a particular point.

We note that the first three direction vectors of Ξ are along the coordinate axis, and the last four direction vectors are the ones on the diagonals of a cube. We can decompose this box spline matrix into two sub-matrices where the axis aligned vectors are separated from the diagonal ones:

$$\Xi_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad \Xi_2 = \begin{bmatrix} 1 & -1 & -1 & 1 \\ -1 & 1 & -1 & 1 \\ -1 & -1 & 1 & 1 \end{bmatrix}$$

Due to the convolution property of box splines as in Equation 3 the seven-directional box spline can be written as the convolution of the

box splines induced by these matrices:

$$M_{\Xi}(\boldsymbol{x}) = (M_{\Xi_1} * M_{\Xi_2})(\boldsymbol{x}) \qquad (6)$$

We can recognize M_{Ξ_1} as the three dimensional box function or the indicator function of the interval $[0,1)^3$. Moreover M_{Ξ_2} is the linear box spline that was introduced for reconstruction of the Body Centered Cubic lattice in [8]. The support of this box spline (when centered at the origin) is a rhombic dodecahedron that is contained in the interval $[-2,2)^3$, see Figure 4(a). To determine the support of M_{Ξ}, we note that it is the convolution of M_{Ξ_1} with M_{Ξ_2}. Therefore, the result of the convolution of these two box splines, results in a support that is contained in $[-5/2, 5/2)^3$.

In [21] the author claims that the support of the seven directional box spline is an octahedron. However, since the support of M_{Ξ_2} is a rhombic dodecahedron as illustrated in Figure 4(a) and the support of M_{Ξ_1} is a cube, the support of M_{Ξ} is a convolution of the indicator function of a rhombic dodecahedron and a cube which is a *truncated rhombic dodecahedron* as illustrated in Figure 4(b).

3.3.2 Fourier Domain Analysis

The center of this box spline is $\frac{1}{2}\sum_{\xi \in \Xi}\xi$. To center this box spline at the origin, we can introduce a shift in the space domain (which amounts to a phase shift in the Fourier domain). Then the Fourier transform of this centered box spline is:

$$\hat{M}_{\Xi}^c(\boldsymbol{\omega}) = \prod_{\xi \in \Xi} \text{sinc}\,(\xi \cdot \boldsymbol{\omega}) \qquad (7)$$

where $\text{sinc}\,(t) = \sin(t/2)/(t/2)$. In order to avoid any signal shifts, we use this centered version of the box spline during reconstruction.

The ideal reconstruction kernel in the space of band-limited functions is the sinc function whose Fourier transform is the indicator function of the unit cube in the Fourier domain. Reconstruction with the sinc function allows the perfect recovery of every wave mode contained in the spectrum of the sampled signal and at the same time it perfectly eliminates the aliasing wave-numbers outside the primary spectrum.

One measure of performance for compactly supported reconstruction kernels is the number of *vanishing moments*, which is the (minimum) order of zero crossings of the kernel at the center of each aliasing wave number in the Fourier domain. The vanishing moments determine the approximation power of reconstruction by the result from Strang and Fix [24]. Further, the more a reconstruction kernel diminishes the wave-numbers inside the primary spectrum, the more errors occur during the reconstruction. Marschner and Lobb introduced a numerical measure for over-smoothing and post-aliasing errors [13].

Using Equation 7, we can describe the Fourier domain behavior of this box spline and compare it with the tri-cubic B-spline recalling that the Fourier transform of the tri-cubic B-spline is $\hat{\beta}_3(\omega_x, \omega_y, \omega_z) = \text{sinc}^4(\omega_x)\,\text{sinc}^4(\omega_y)\,\text{sinc}^4(\omega_z)$. From careful examination of the Fourier transform of this box spline and the tri-cubic B-spline, we realize that they both attain a minimum of *four* vanishing moments at the center of every aliasing frequency; hence both of them have a fourth order approximation power.

From considering the replica in the Fourier domain along each coordinate axis at $(2\pi, 0, 0)$, $(0, 2\pi, 0)$ and $(0, 0, 2\pi)$ we note that the seven-directional box spline has *five* vanishing moments and the tri-cubic B-spline has *four*. Due to greater number of zero crossings, the seven-directional box spline introduces slightly (i.e. 25%) more diminishing power along the axis. For the replica located on the sub-diagonals at $(2\pi, 2\pi, 0)$, $(2\pi, 0, 2\pi)$ and $(0, 2\pi, 2\pi)$ the seven-directional box spline has *four* vanishing moments while the tri-cubic B-spline has *eight*. Here the tri-cubic B-spline introduces significantly (200%) more diminishing power. Hence frequencies inside the primary spectrum along the sub-diagonals suffer from sever diminishing. Further, on the diagonal direction $(2\pi, 2\pi, 2\pi)$ the seven-directional box spline introduces *seven* vanishing moments while the tri-cubic B-spline has *twelve*. Therefore

the tensor product again introduces a significant diminishing of frequencies along the diagonal directions on the primary spectrum of the signal.

To illustrate the differences between the seven-directional box spline and the tensor product splines, we plotted a 2D slice of the frequency response on the XY plane in Figure 5(a-d) as well as a slice spanned by two antipodal diagonals of the cubic support of the primary spectrum in Figure 5(e-h). These figures display the wave numbers that are being recovered (through white regions) during the reconstruction and compare tri-linear, tri-quadratic and tri-cubic to the seven directional box spline. In these figures the thin red band indicates the boundary of the primary spectrum. The goal of reconstruction is to recover frequencies inside the primary spectrum and eliminate the frequencies outside. By diminishing the frequencies inside the main spectrum, we introduce over-smoothing while the remaining frequencies outside the primary spectrum lead to aliasing. Hence, it is apparent from these figures that the seven-directional box spline strikes the best balance between over-smoothing and aliasing among tri-linear, tri-quadratic and tri-cubic B-splines.

(a) tri-linear, XY plane   (b) box spline Ξ, XY plane

(c) tri-quadratic in XY plane   (d) tri-cubic, XY plane

(e) tri-linear, diagonal plane   (f) box spline Ξ, diagonal plane

(g) tri-quadratic, diagonal plane   (h) tri-cubic, diagonal plane
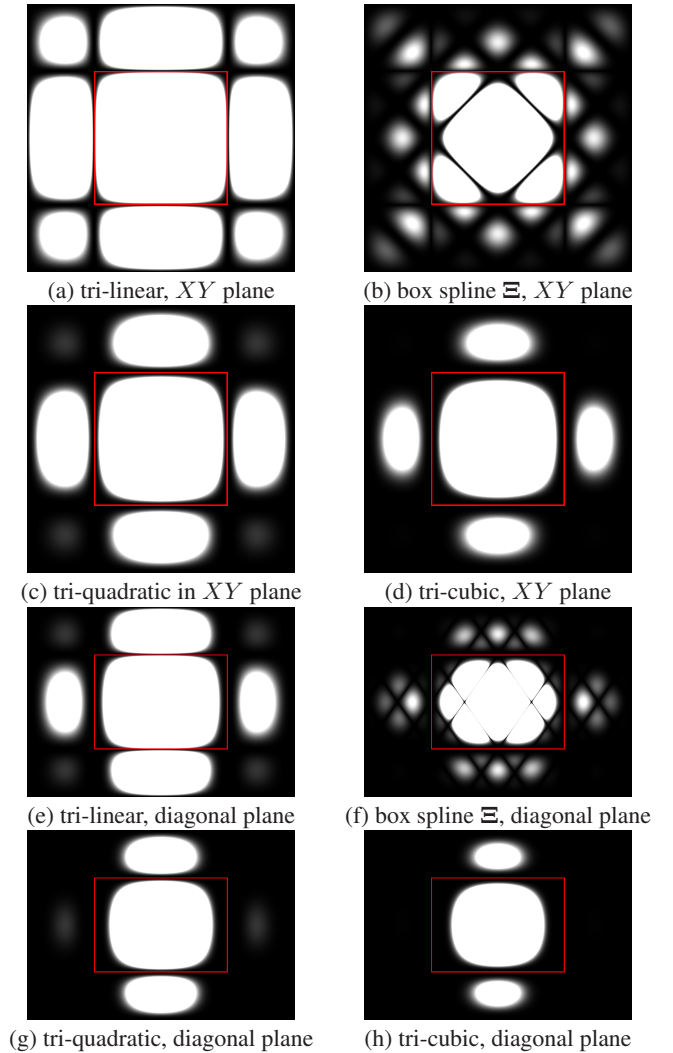
Fig. 5. The frequency responses of various reconstruction kernels on the XY plane (a-d) and on a plane spanned by 2 diagonal directions (e-h). Comparison of the frequency response of (a,e) the tri-linear tensor product kernel (Hat function or *linear* B-spline); (b,f) the seven-directional box spline; (c,g) the tri-quadratic B-spline; and (d,h) the tri-cubic B-spline. In terms of smoothness and accuracy, tri-cubic matches the box spline. However, the box spline allows for a fuller reconstruction of the primary spectrum. For emphasis, responses higher than 0.0005 are set to white.

3.3.3 Variants

If we consider the decomposition of the direction vectors as above, we realize that reconstruction by Ξ is a 3D nearest neighbor Ξ_1 convolution followed by a convolution with the linear order Ξ_2. Reconstruction with the box spline of Ξ_1 is the nearest neighbor interpolation; this kernel provides a discontinuous reconstruction with first order approximation power. The box spline of Ξ_2 by itself is a linear order reconstruction on the rhombic dodecahedral support of this box spline. When centered on a Cartesian lattice point, the support of this rhombic dodecahedron extends to the two neighbors on the left and the right on each axis. On the diagonal directions its support is only to the nearest neighbors, hence we concluded earlier that the support is contained in $[-2, 2)^3$. For this box spline $\rho = 2$ since it provides a C^0 reconstruction. Examining vanishing moments, one can verify that it is only of second order in terms of approximation power.

One can shrink the support of this box spline by considering the box spline induced by $\Xi'_2 = \frac{1}{2}\Xi_2$. The support of this box spline is a rhombic dodecahedron whose extent goes to the closest neighbors on each axis and half way to the nearest neighbor on the diagonal direction. Therefore, the support of this box spline is contained in $[-1, 1)^3$. Putting things back together the box spline induced from $[\Xi_1, \Xi'_2]$ offers the same C^2 smoothness in reconstruction since the directions of the box spline didn't change hence ρ is the same as Ξ. The support of this box spline is contained in $[-3/2, 3/2)^3$ and examining its vanishing moments one can show that its approximation power is decreased to only *one*. This box spline would be a suitable reconstruction filter only when smoothness of the reconstruction is important and the approximation order is less important. Since it has a smaller support it offers a faster reconstruction.

Another possibility is to use a tri-linear interpolant instead of Ξ_1. The tri-linear interpolant is induced from $[\Xi_1, \Xi_1]$. The box spline that is induced from $[\Xi_1, \Xi_1, \Xi'_2]$ is also of interest. The support of the tri-linear is $[-1, 1)^3$, hence the support of this box spline is contained in $[-2, 2)^3$. For this box spline $\rho = 6$. Hence it provides with a C^4 reconstruction while its vanishing moments reveal that it is of *second* order approximation power. This box spline obviously favors the axis directions more strongly due to its tri-linear component.

These variants offer a variety of approximation order and smoothness combinations that are surprising yet interesting.

4 IMPLEMENTATION

The most efficient method to evaluate a generic box spline is the evaluation using de Boor's recurrence relation:

$$M_\Xi(\boldsymbol{x}) = \sum_{\boldsymbol{\xi} \in \Xi} t_{\boldsymbol{\xi}} M_{\Xi \backslash \boldsymbol{\xi}}(\boldsymbol{x}) + (1 - t_{\boldsymbol{\xi}}) M_{\Xi \backslash \boldsymbol{\xi}}(\boldsymbol{x} - \boldsymbol{\xi}) \quad (8)$$

for $\boldsymbol{x} = \Xi \boldsymbol{t}$. Evaluation of box splines is an active research topic since they are prone to significant numerical instabilities [5] and are inefficient computationally. For a generic box spline the recursive evaluation of box splines is the most favorable method computationally, but the numerical instabilities that makes the recursive formula impractical. Kobbelt addresses the issue of numerical instabilities by delaying the evaluation of the discontinuous Heaviside function to the latest stages of recursion [12]. Even though the numerical inaccuracies of the recursive algorithm can be minimized, to make box splines practical in volume rendering their computational complexity demands a more explicit derivation that is faster.

It is not uncommon to use lookup tables for the purpose of rendering (e.g. most hardware based rendering algorithms [19, 10]). Hence, we can evaluate our box splines on a fine resolution volume and use this volume as a look up table of weights during the reconstruction phase in volume rendering.

Since box splines have a simple Fourier domain representation Mc-Cool [14, 15] suggests a Fourier domain sampling of the box spline, which can yield the space domain sampling of the box splines using an inverse FFT. However, since truncating sinc functions in Equation 7 leads to ringing artifacts explained by the Gibbs phenomenon, there needs to be a windowing function applied in the Fourier domain

to weaken these artifacts. However, the windowing function has a smoothing effect which changes the reconstruction quality of the box spline if this approach is used.

Instead of introducing smoothing we used the structure of the seven directional box spline – the decomposition in Equation 6. The box spline of M_{Ξ_1}, when centered, is simply a constant function over $[-\frac{1}{2}, \frac{1}{2})^3$. In [8] we arrived at the following explicit form of M_{Ξ_2}:

$$M_{\Xi_2}(x, y, z) = 2 \max(0, 1 - \max(|x| + |y|, |x| + |z|, |y| + |z|)). \quad (9)$$

Using this decomposition, we obtain a fine sampling of M_{Ξ_1} and M_{Ξ_2} on two volumes. In our implementation, we used a sampling resolution of 100^3 for sampling the interval $[-5/2, 5/2)^3$ which contains the support of M_Ξ. The first box spline has a support of only $[-\frac{1}{2}, \frac{1}{2})^3$ and the support of the second box spline is contained in $[-2, 2)^3$. Hence, a 100^3 sampling of each box spline on the interval of $[-5/2, 5/2)^3$ was obtained; therefore, an in place convolution allows for the larger support of the box spline. To obtain a volume that is a sampling of M_Ξ, we need to perform a discrete convolution of the original two volumes representing M_{Ξ_1} and M_{Ξ_2}. This convolution can be performed as a multiplication in the Fourier domain. The resulting representation for M_Ξ was stable and free from discretization and sampling artifacts as increasing the sampling resolution from 100 did not improve the quality of reconstructed images. Also, no ringing artifacts existed since we avoid truncating the sinc function in the Fourier domain by having an explicit representation of M_{Ξ_1} and M_{Ξ_2} in the space domain. Once this volume is obtained, it can be used as a lookup table during volume rendering.

5 RESULTS

In order to validate our theoretical expectations on the decrease of directional artifacts during the reconstruction, we have used several datasets to compare the seven directional box spline against the popular tri-cubic B-spline reconstruction. Since both of these filters produce C^2 reconstructions and they both have a fourth order approximation power, they are comparable filters. Since the neighborhood of the tri-cubic B-spline has a support of $4 \times 4 \times 4$ and the box spline has a support of $5 \times 5 \times 5$ in our experiments, we have also chosen filters of wider support such as $d0\_c2\_4ef$ and $d0\_c3\_4ef$ from [17, 18]. In our experiments these more powerful filters (with support of six taps) were not able to reduce any of the artifacts that are observed in the tri-cubic B-spline case. The dominance of the box spline reconstruction over these tensor product reconstructions was preserved. For convenience of comparisons, we rendered the datasets with an opaque transfer function so that the renderer produces a crisp opaque surface where we can pick out the visible differences between various reconstructions. The gradients for shading purposes was estimated by central differencing of step size 1; also, all of the images were generated with a sampling resolution of 0.1 along the ray.

Our first case study is on reconstructing sharp planar surfaces with varying inclination. This experiment is illustrated in Figure 6. We designed a set of voxelized planar surface datasets that are essentially half spaces angled from zero to forty five degrees at increments of five degrees sampled at $21 \times 21 \times 21$ resolution on the Cartesian lattice. The zero degree surface represents a plane that is aligned with a coordinate axis of the Cartesian lattice, while the forty five degrees planar surface makes a forty five degrees angle with the Z-axis in the XZ plane. While the box spline reconstruction performed equally well when compared with tri-cubic B-spline for planes between $0 - 30$ degrees, sharper angled planes showed a big difference in the two reconstructions. This is as expected, since for small angels the main features of the planar surface are indeed aligned on the coordinate axis. The isotropy of the box spline reconstruction allowed the extremely smooth reconstruction of the planar surface, while the tri-cubic B-spline solution exhibited grid-aligned stripes on the surface known as stair-casing artifacts. In our experiments none of the tensor product reconstructions (even with wider support) in [17, 18] could decrease these stair-casing artifacts. We have also experimented with a lot of different inclination

angles and the behavior of the box spline reconstruction was consistent with what we reported above. This is due to the shape of the support and continuity of the box spline that is guaranteed to produce a C^2 reconstruction.
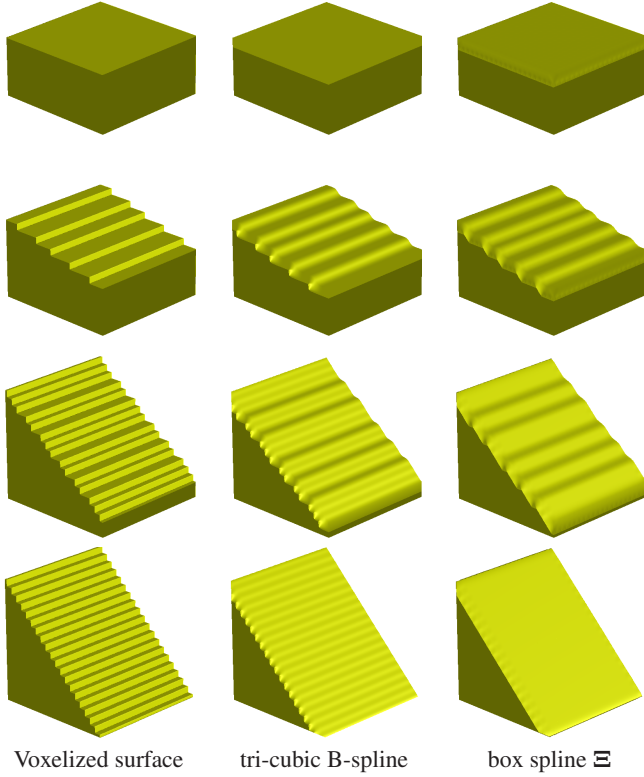

Voxelized surface   tri-cubic B-spline   box spline Ξ

Fig. 6. Planar surfaces with varying inclination sampled at the resolution of $21 \times 21 \times 21$. The first row shows the plane at 0 degrees, the second row is at 15 degrees, the third row is at 35 degrees and the last row at 45 degrees. The left column is the voxelized surface rendered with nearest neighbor reconstruction to reflect the discretization of the planar surface. The middle column uses the tri-cubic B-spline and the right column uses the box spline Ξ. The wider stripes are due to the aliasing *existing* in the volume dataset. The stair-casing artifacts appear as the thin stripes which are dominant in the left and the middle column are due to aliasing in the reconstruction.

We also experimented with the dataset proposed by Marschner and Lobb in [13] as a benchmark dataset for evaluating the power of reconstruction kernels. The original analytical function is depicted in Figure 7. The result of the reconstruction and the corresponding error images appear in Figure 8. The error images in the second row convey the amount of angular error occurred when estimating the gradients on the surface. A maximum angular error of 1 radian was mapped to white. The first column demonstrates the tri-cubic B-spline reconstruction. The second column in Figure 8 demonstrates the reconstruction achieved by the variant suggested in

Fig. 7. The rendition of the explicit function introduced by Marschner-Lobb.

Section 3. This box spline is represented by $[\Xi_1, \Xi_1, 1/2\Xi_2]$. The third column demonstrates the seven directional box spline reconstruction by Ξ. As discussed before, the variant box spline (in the second column) favors the axis aligned frequencies it resembles more like the tri-cubic reconstructions. Although the aliasing artifacts are clearly visible in this reconstruction due to its poor second order approxima-

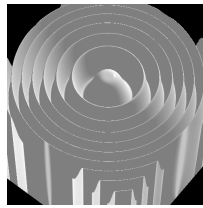tion power, yet the reconstruction is C^4 smooth. This reconstruction achieves the lowest angular error among the three reconstructions, as its error image is darker among the three. A close comparison of the reconstruction by the box spline of Ξ with the tri-cubic B-spline reconstruction reveals that the first, second and the third rings (from the center) have preserved their circular shape in the box spline reconstruction while they have been distorted in the tensor-product reconstruction. Also, the amount of distortion on the outer rings is smaller for the box spline reconstruction. Note that, in the areas close to the corners, the outer rings have sharp spikes in the tri-cubic B-spline case while these sharp spikes are nicely removed in the box spline reconstruction. These images support the isotropic property of the seven-directional box spline reconstruction as the circular rings of the dataset are preserved more accurately while with the tri-cubic B-spline the aliasing artifacts destroyed the circular structure of the middle rings of the dataset. However, close examination of the angular errors when estimating gradients shows a slight increase of the gradient angular errors on the box spline reconstruction. This is expected since the box spline is only a fourth order polynomial approximation, while the tri-cubic B-spline is a ninth degree polynomial approximation.
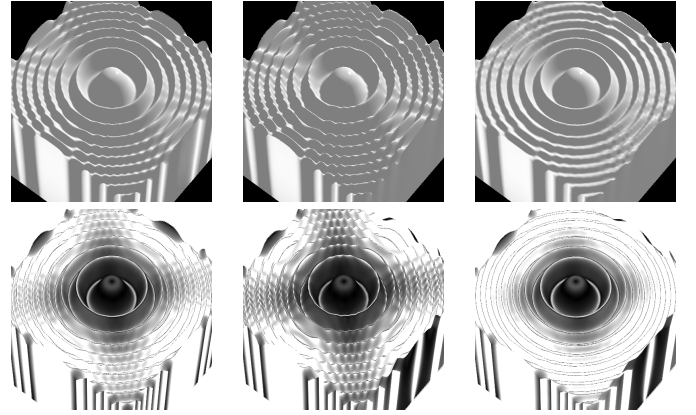

Fig. 8. The Marschner-Lobb dataset sampled at the resolution of $40 \times 40 \times 40$. The first column is reconstructed using tri-cubic B-spline. The second column is reconstructed using the variant represented by $[\Xi_1, \Xi_1, 1/2\Xi_2]$. The third column shows the reconstruction using the box spline of Ξ. The second row depicts the angular errors in estimating the gradient where white denotes the error of 1 radian.

Furthermore we experimented with a discretization of a solid sphere in Figure 9. As the grid-aligned artifacts in the tensor product reconstruction are apparent in the left images, the right images demonstrate the superior reconstruction using the proposed seven-directional box spline. However, a nicer reconstruction is achieved by the tri-cubic B-spline on faces along axis.

We also experimented with voxelized surfaces. The Stanford bunny dataset was voxelized on the Cartesian lattice with an 18 connectivity neighborhood. The resulting volume was rendered using the tri-linear and tri-cubic B-spline and the box spline Ξ in Figure 10. The resulting images also confirm the predicted behavior of the box spline reconstruction since the surface is reconstructed with a lot less artifacts.

For a real dataset, we experimented with the hydrogen atom dataset. Figure 11 illustrates the more isotropic reconstruction achieved by the box spline solution as there are less artifacts present in the box spline reconstruction. The artifacts present in the dataset are reduced in the box spline reconstruction in the right image and are more pronounced in the tensor product reconstruction in the left image.

We also experimented with the UNC brain dataset. The resulting reconstruction images are tabulated in Figure 12 and Figure 13. These images demonstrate that box spline reconstructions significantly removed the artifacts that are present in the tri-cubic B-spline reconstructions. A second view of the same dataset shows similar improvements in the second and third row of Figure 12.

| tri-linear B-spline | tri-cubic B-spline | box spline Ξ |

Fig. 10. The Stanford Bunny dataset voxelized with an 18 connectivity neighborhood at the sampling resolution of $182 \times 182 \times 182$.

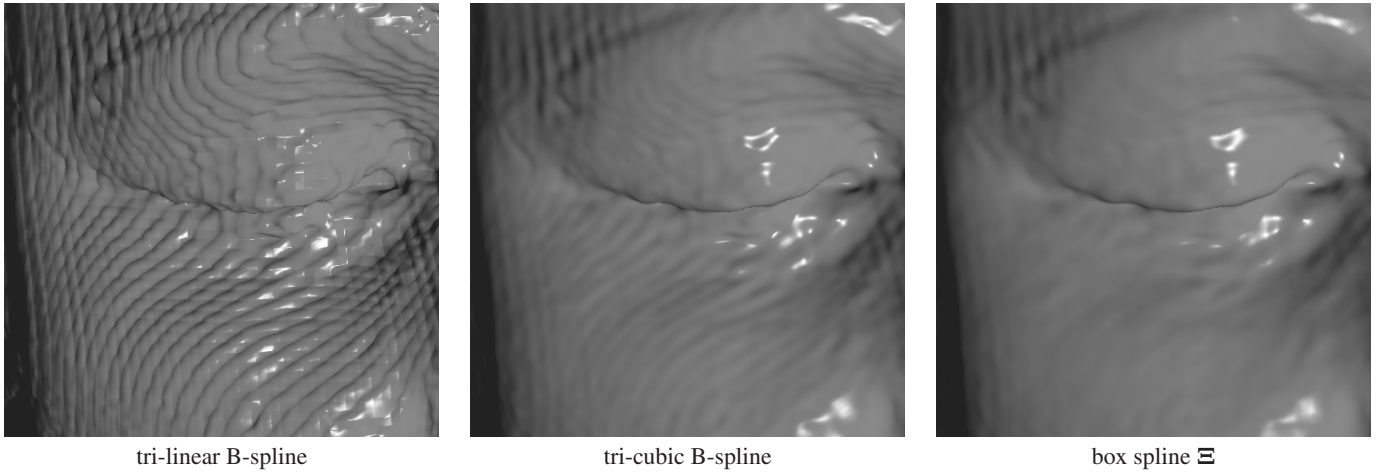| tri-linear B-spline | tri-cubic B-spline | box spline Ξ |

Fig. 13. Reconstruction of the eye area in the UNC brain dataset. The grid aligned artifacts are significantly reduced in the proposed reconstruction.

In terms of rendering time performance, since the reconstruction kernels were pre-computed in form of a look-up table, there was no penalty for evaluating the reconstruction kernels. However, since the box spline Ξ has a support of $5 \times 5 \times 5$, its rendering time is slightly more than tri-cubic B-spline whose support is $4 \times 4 \times 4$. On the other hand, box spline reconstruction was slightly faster than such six tap filters as $d0\_c2\_4ef$ and $d0\_c3\_4ef$ from [17, 18].

6 CONCLUSION

In this paper we motivated the use and purpose of non-separable reconstruction for the Cartesian lattice. Unlike tensor product extensions, non-separable reconstruction allows one to adapt the reconstruction behavior of the applied kernel in a much more isotropic way. This dramatically reduces the directional bias of reconstruction of functions sampled on the Cartesian lattice. We have adapted our reconstruction kernel in seven directions, as opposed to the three directions one is typically constrained with using tensor product extensions.

We have demonstrated our theory on the specific case of a seven-directional box spline, that can be seen as the 3D extension of the well known Zwart-Powell element in 2D. We proposed the use of this box spline for rendering purposes and we offered a thorough Fourier domain analysis of its advantages when used in reconstructing volumetric data. We were also able to show the correct shape of this box spline, which has been miss-reported in the previous literature. Further, this box spline merely serves as an example for the potential of non-separable reconstruction in graphics and visualization applications, which as far as we know is novel.

We have also demonstrated that the use of this reconstruction is rather efficient and can be implemented using a look-up table. However, for more esoteric applications that demand analytical evaluations, we plan to derive explicit polynomial partitions of this box spline.
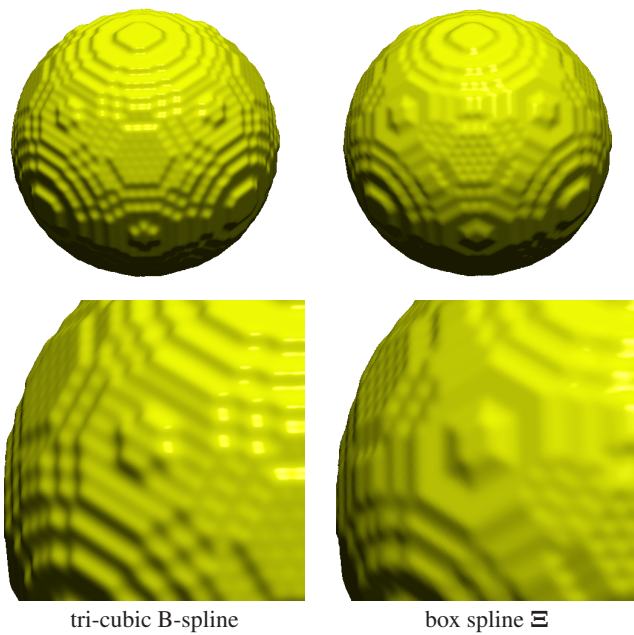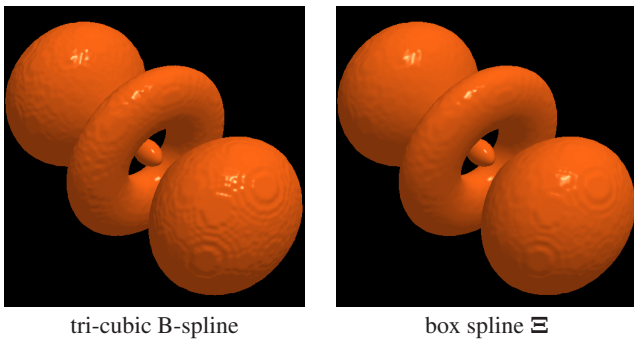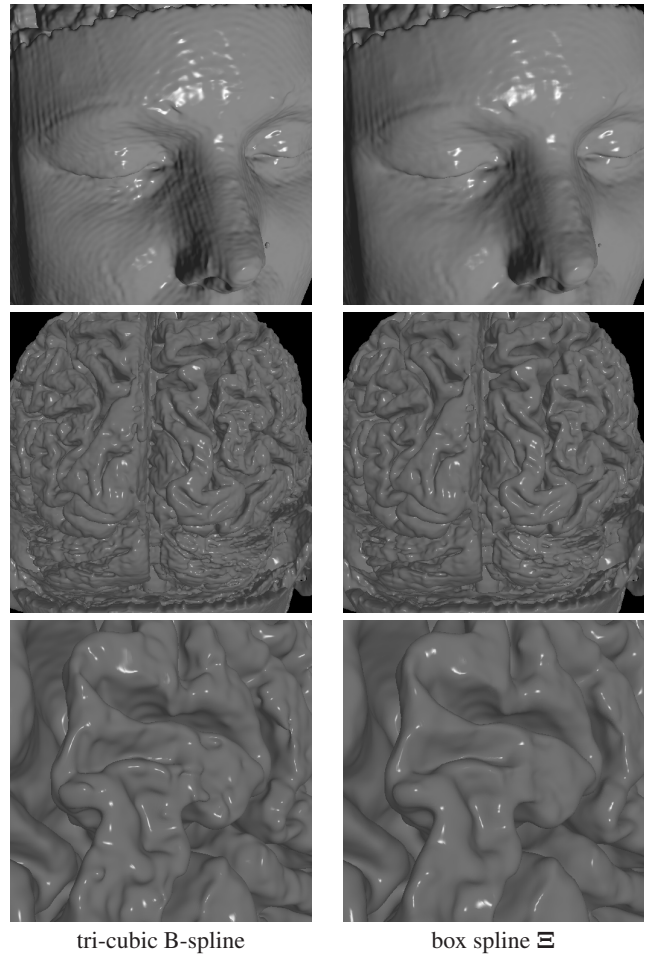
REFERENCES

[1] G. Birkhoff. Piecewise Bicubic Interpolation and Approximation in Polygons. *Approximation with Special Emphasis on Spline Functions V*, 1969.

[2] G. Birkhoff and C. de Boor. Piecewise Polynomial Interpolation and Approximation. *Approximation of Functions*, 1965.

[3] G. Birkhoff and H. L. Garabedian. Smooth Surface Interpolation. *J. Math. and Phys.*, 39:258–268, 1960.

[4] I. Carlbom. Optimal Filter Design for Volume Reconstruction and Visualization. In *Proceedings of the IEEE Conference on Visualization*, pages 54–61, Oct. 1993.

[5] C. de Boor. On the Evaluation of Box Splines. *Numerical Algorithms*, 5(1–4):5–23, 1993.

[6] C. de Boor, K. Höllig, and S. Riemenschneider. *Box Splines*. Springer Verlag, 1993.

tri-cubic B-spline   box spline Ξ

Fig. 9. The sphere dataset sampled at the resolution of $64 \times 64 \times 64$.

tri-cubic B-spline   box spline Ξ

Fig. 11. The hydrogen atom dataset at the resolution of $128 \times 128 \times 128$.

tri-cubic B-spline   box spline Ξ

Fig. 12. The UNC brain dataset at the resolution of $256 \times 256 \times 145$.

ume 22, pages 221–228, Aug. 1988.

[7] S. C. Dutta Roy and B. Kumar. *Handbook of Statistics*, volume 10, chapter Digital Differentiators, pages 159–205. Elsevier Science Publishers B. V., North Holland, 1993.

[8] A. Entezari, R. Dyer, and T. Möller. Linear and Cubic Box Splines for the Body Centered Cubic Lattice. In *Proceedings of the IEEE Conference on Visualization*, pages 11–18, Oct. 2004.

[9] A. Entezari, T. Meng, S. Bergner, and T. Möller. A Granular Three Dimensional Multiresolution Transform. In *Proceedings of the Eurographics/IEEE-VGTC Symposium on Visualization*, pages 267–274, May 2006.

[10] M. Hadwiger, T. Theußl, H. Hauser, and E. Gröller. Hardware-Accelerated High-Quality Filtering on PC Hardware. In *Workshop on Vision, Modeling and Visualization*, pages 105–112, 2001.

[11] R. G. Keys. Cubic Convolution Interpolation for Digital Image Processing. *IEEE Tran. Acoustics, Speech, and Signal Processing*, 29(6):1153–1160, Dec. 1981.

[12] L. Kobbelt. Stable Evaluation of Box Splines. *Numerical Algorithms*, 14(4):377–382, 1997.

[13] S. R. Marschner and R. J. Lobb. An Evaluation of Reconstruction Filters for Volume Rendering. In *Proceedings of the IEEE Conference on Visualization*, pages 100–107, 1994.

[14] M. D. McCool. Analytic Antialiasing With Prism Splines. *SIGGRAPH Conference Proceedings*, pages 429–436, 1995.

[15] M. D. McCool. Accelerated Evaluation of Box Splines via a Parallel Inverse FFT. *Computer Graphics Forum*, 15(1):35–45, 1996.

[16] D. P. Mitchell and A. N. Netravali. Reconstruction Filters in Computer Graphics. In *Computer Graphics (Proceedings of SIGGRAPH 88)*, vol-

[17] T. Möller, R. Machiraju, K. Mueller, and R. Yagel. A Comparison of Normal Estimation Schemes. In *Proceedings of the IEEE Conference on Visualization*, pages 19–26, Oct. 1997.

[18] T. Möller, K. Mueller, Y. Kurzion, R. Machiraju, and R. Yagel. Design of Accurate and Smooth Filters for Function and Derivative Reconstruction. *Proceedings of the Symposium on Volume Visualization*, pages 143–151, Oct 1998.

[19] K. Mueller, T. Möller, J. E. Swan II, R. Crawfis, N. Shareef, and R. Yagel. Splatting Errors and Antialiasing. *IEEE Transactions on Visualization and Computer Graphics*, 4(2):178–191, 1998.

[20] A. Oppenheim and R. Schafer. *Discrete-Time Signal Processing*. Prentice Hall Inc., Englewoods Cliffs, NJ, 1989.

[21] J. Peters. C^2 Surfaces Built from Zero Sets of the 7-direction Box Spline. In *The Mathematics of Surfaces VI*, pages 463–474, 1996.

[22] M. J. D. Powell and M. A. Sabin. Piecewise Quadratic Approximations on Triangles. *ACM Trans. Math. Softw.*, 3(4):316–325, 1977.

[23] W. Qiao, D. Ebert, A. Entezari, M. Korkusinski, and G. Klimeck. VolQD: Direct Volume Rendering of Multi-million Atom Quantum Dot Simulations. In *Proceedings of the IEEE Conference on Visualization*, pages 319–326, Oct. 2005.

[24] G. Strang and G. J. Fix. A Fourier Analysis of the Finite Element Variational Method. *Construct. Aspects of Funct. Anal.*, pages 796–830, 1971.

[25] D. Van De Ville, T. Blu, and M. Unser. On the Multidimensional Extension of the Quincunx Subsampling Matrix. *IEEE Signal Processing Letters*, 12(2):112–115, February 2005.

[26] D. Van De Ville, T. Blu, M. Unser, W. Philips, I. Lemahieu, and R. Van de Walle. Hex-Splines: A Novel Spline Family for Hexagonal Lattices. *IEEE Transactions on Image Processing*, 13(6):758–772, June 2004.

[27] P. B. Zwart. Multivariate Splines with Nondegenerate Partitions. *SIAM Journal on Numerical Analysis*, 10(4):665–673, Sept. 1973.