The Visual Computer manuscript No. 206

Erald Vuçini · Torsten Möller · M. Eduard Gröller

Efficient Reconstruction from Non-uniform Point Sets

Abstract We propose a method for non-uniform reconstruction of 3D scalar data. Typically, radial basis functions, trigonometric polynomials or shift-invariant functions are used in the functional approximation of 3D data. We adopt a variational approach for the reconstruction and rendering of 3D data. The principle idea is based on data fitting via thin-plate splines. An approximation by B-splines offers more compact support for fast reconstruction. We adopt this method for large datasets by introducing a block-based reconstruction approach. This makes the method practical for large data sets. Our reconstruction will be smooth across blocks. We give reconstruction measurements as error estimations based on different parameter settings and also an insight on the computational effort. We show that the block size used in reconstruction has a negligible effect on the reconstruction error. Finally we show rendering results to emphasize the quality of this 3D reconstruction technique.

Keywords Non-uniform Reconstruction · Variational Approximation · B-splines · 3D Object Modeling

1 Introduction

Regular (Cartesian) grids represent one of the most common data types in volumetric rendering. Although Cartesian grids are the standard for a wide variety of situations, many applications, including fluid dynamics, weather modeling, Doppler measurements etc. use non-uniform (irregular) grids. In such cases the use of non-uniform grids is feasible because we can represent regions of space having little variation with fewer voxels, and regions with rapid changes with more voxels.

Erald Vuçini
Vienna University of Technology
E-mail: vucini@cg.tuwien.ac.at

Torsten Möller
Simon Fraser University
E-mail: torsten@sfu.ca

M. Eduard Gröller
Vienna University of Technology
E-mail: groeller@cg.tuwien.ac.at

While the reconstruction of Cartesian grids is well-understood and can be solved via a convolution of tensor-product kernels, the reconstruction of non-uniform data is more expensive. In this paper we pursue the strategy to build a regular Cartesian grid out of the non-uniform grid before rendering the data. Therefore we can apply standard, fast volume rendering algorithms for the display of the data. A list of previous work is provided in Section 2.

We build our technique on previous work of Arigovindan et al. [4], [5]. An overview of this approach is given in Section 3. Although their results are very good, the straightforward implementation of their technique for 3D non-uniform data is often not feasible and introduces problems regarding reconstruction errors and memory requirements. In our work we propose to break-up the reconstruction data into axis-aligned blocks to use this method even for large data sizes.

The main contribution of this paper is the application of the variational reconstruction approach of Arigovindan et al. to volumetric data. In order to make this feasible, we divide the non-uniform point cloud into an assembly of blocks, so that each block can be reconstructed separately. In order to deal with discontinuity problems we overlap the blocks by a 2 voxel margin. The details are found in Section 4. Therefore, we can work on datasets of different sizes without memory constraints. A number of results for various data sets are detailed in Section 5. Since this paper presents a feasibility study, we have not yet applied our work to actual non-uniform data. It would be difficult to measure the error with respect to the "ground truth" for such data. Instead we have been following the approach of Arigovindan et al. and produced non-uniformly sampled test cases given a Cartesian representation of the data set. This is achieved by thresholding the gradient data after applying a Laplacian filter.

In Section 6 our work is summarized and ideas for future research are given.

2 Related Work

Recent technological advances have improved the process of 3D data acquisition. The amount and size of such data is increasing steadily. This brings the necessity of sampling such data adaptively, requiring more sophisticated reconstruction and processing algorithms. Various research in this area can be divided into several groups.

A common approach for dealing with non-uniformly sampled data is (a finite element style approach) by first polyhedralizing the given set of points and then rendering the polyhedra through specific volume rendering techniques. A seminal work in this direction was introduced by Garrity [9], who used raytracing as a technique for rendering the convex cells structured as tetrahedra. A problem that arises is the proper visibility sorting. This was addressed by a number of researchers, including Silva et al. [21] and Weiler et al. [25]. However, visibility sorting tends to be computationally expensive in terms of time and memory.

Radial Basis Functions (RBFs) are another widely used method for approximating non-uniform data by applying an implicit representation. Several RBF-based approaches have been proposed for applications related to surface fitting, geometric modeling and visualisation of 3D scattered data [10], [12]. However these methods suffer from high computational costs and also from smoothing [10], [12]. Jang et al. [12] introduce an RBF approximation technique for non-uniformly sampled data. The centers and weights of the RBFs are found through an iterative PCA-based clustering technique and truncated Gaussians are used as basis functions. In their work it is stated that the method has very high fitting times and has best results for spherical structures. In order to overcome these problems researchers have proposed the use of ellipsoids instead of RBFs for sparse data fitting [11], [18]. Jang et al. [11] improve on their previous work by introducing non-uniformly sampled data encoding via axis-aligned or arbitrary-aligned ellipsoids. Although this method improves the results and the fitting times, these issues are still not fully solved. Juba and Varshney [13] propose a hierarchical RBF-based method for uniform data fitting. Their technique is based on Maximum Likelihood Estimation and since it is non-iterative, it does not suffer from high fitting times. Their method works on regular grids and its extension to non-uniformly sampled data is non-trivial. Most of the above mentioned approaches use truncated Gaussian functions as basis functions, leading to artifacts in the regions where the truncations occur. In our work we use compact basis functions such as B-splines instead.

Several approaches using point-based approximation techniques have been introduced in recent years. One motivation behind this recent trend is that a point primitive offers greater simplicity and detail than a small triangle [27]. Most papers in this domain are related to surface approximation and surface rendering [3], [8]. Alexa et al. [3] introduce a "lighter" point set, which is a subset of the original data, obtained by a downsampling technique based on moving least squares approximation. They use this "new" representation for accelerating and improving their surface rendering algorithm. Our technique is similar to their approach with regard to the fact that we initially also find a point subset from the original data. As it will be shown next, this point subset gives very good results with our approximation technique.

Welsh and Mueller [26] introduce a frequency-sensitive point hierarchy where high frequency points are found by means of a frequency-space analysis based on Gabor wavelets. Our work is similar in the sense that we also downsample the dataset by selecting high frequency points, but for this we use a 3D Laplacian filter.

Lee et al. [15] use multilevel B-splines to compute a C^2 continuous surface from a set of irregularly sampled data. Their algorithm is comparable to the one presented in this paper with regard to the use of a coarse-to-fine two-scale relation of odd degree B-splines. Since they do surface reconstruction they use bicubic B-spline basis functions, opposite to the tricubic functions used in this paper. Rössl et al. [20] introduce volumetric data approximation using piecewise cubic polynomials. They reconstruct uniform type-6 tetrahedra partitions of volumetric data. While their method is applicable only after tetrahedra partitioning of the volumes, they report low fitting times and their focus is mainly on iso-surface rendering.

The decomposition of volumes into blocks and processing them independently is a well-known technique in signal processing, visualisation and computer graphics. Tuncer [23] proposes a block-based signal reconstruction from band-limited non-uniform samples using an interpolation based on the Discrete Fourier Transform. Ljung et al. [16] propose an interblock interpolation technique that enables direct volume rendering of block-based multiresolution volumes. Several approaches use volume subdivison into hierarchical blocks to accelerate volume rendering [6] or to improve rendering quality by blending renderings of blocks with different representations [14], [17].

3 Variational Reconstruction - Overview

In this section we give a short introduction to the variational approximation approach via B-splines. For a deeper insight into the method we refer the reader to Arigovindan [4].

Non-uniform sampling and reconstruction techniques have received special attention in recent years especially for two dimensional image and signal data. Aldroubi and

Gröchenig [2] and Aldroubi and Feichtinger [1] introduce several mathematical concepts related to reconstruction from non-uniform data in shift-invariant function bases. Generally, if there is no restriction on the distribution of the samples, the reconstruction is not uniquely defined and hence is ill-posed. In such cases a variational approach is used and the reconstruction routine is formulated as an minimization of two terms: *a*) the sum of squared errors, and *b*) the regularization term that controls the smoothness of the solution. The first part guarantees that the solution is close to the sample points, while the second part ensures that there are no discontinuities in the reconstruction. In variational theory the best results with regard to approximation accuracy are given by RBFs, and particularly by a specific class of basis functions known as thin plate splines [7]. While thin-plate splines are one of the preferred approaches to deal with multi-dimensional non-uniform data, they tend to be computationally expensive when the number of points increases significantly. To overcome this problem it is proposed to discretize the thin-plate splines using uniform B-splines attached to the reconstruction grid. While this theory holds mathematically for one dimensional signal reconstruction, for higher dimensions there are no compactly supported B-splines that span the same space as the thin-plate splines. However cubic B-splines are very good candidates for the reconstruction process.

As described by Thévenaz et al. [22], B-splines have several properties which make them very suitable for signal approximation. We mention properties such as easy analytical manipulation, several recursion relations, the m-scale relation, minimal curvature, easy extension to quasi interpolation, simplicity of their parametrization etc.. One basic feature, which makes B-splines very suitable in applications related to signal approximation, is that they enjoy the maximal order of approximation for a given integer support, providing the best quality for a given computational cost.

Given a set of sample points, $p_i = (x_i, y_i, z_i)$, i = 1,2,...,M, let f_i be the scalar value associated with p_i. We define the B-spline approximation through the form

$$S(x,y,z) = \sum_{k=0}^{N-1}\sum_{l=0}^{N-1}\sum_{m=0}^{N-1} c_{k,l,m}\beta^3(x-k)\beta^3(y-l)\beta^3(z-m) \quad (1)$$

where $\beta^3(x)$ is the cubic B-spline basis function. In order to find the coefficients $c_{k,l,m}$ the following cost function is minimized

$$C(S) = \sum_i \|S(x_i,y_i,z_i) - f_i\|^2 + \lambda \int\int\int \|D^p S\|^2 dxdydz \quad (2)$$

where λ is a parameter that together with the regularization term $D^p S$ controls the smoothness of the solution. The crucial part of the variational technique is to express the second term in equation 2 by means of the first term. This can be achieved using Duchon's semi-norms which are represented by the second term in Eq. 2. These norms are a combination of the sum of partial derivatives of a degree chosen respectively to the reconstruction technique and in this case also to the spline degree. We can express Eq. 2 with a simpler formulation using the following representations:

$$\begin{cases} c = [c_{0,0,0}\ldots, c_{N-1,0,0}, \ldots, c_{N-1,N-1,N-1}] \\ f = [\ldots f_i \ldots] \\ S_{i,N^2m+Nl+k} = \beta^3(x_i-k)\beta^3(y_i-l)\beta^3(z_i-m) \end{cases} \quad (3)$$

where $c_{k,l,m}$ are the B-spline coefficients and N is the size of the dataset in each dimension. The cost function now can be rewritten as:

$$C(S) = \|f - Sc\|^2 + \lambda c^T Rc \quad (4)$$

where R is a block-circulant filter that corresponds to a regularization filter which is derived from the Duchon's semi-norm. By applying the Euler-Langrange functional equation for variable c we have

$$[S^T S + \lambda R]c = S^T f \quad (5)$$

We denote $A = S^T S + \lambda R$ and $b = S^T f$ for the sake of simplicity. Equation 5 is well-posed and can be solved through different numerical analysis methods. However, being based on the basic feature of the two-scale relation of odd degree B-splines Arigovindan et al. [5] proposed a multigrid iteration algorithm for finding the solution to the cost minimization problem. Considering the reconstruction at different scales, we specify 2^j as the scale size and we have

$$S_j(x,y,z) = \sum_{k=0}^{\frac{N-1}{2^j}}\sum_{l=0}^{\frac{N-1}{2^j}}\sum_{m=0}^{\frac{N-1}{2^j}} c_{k,l,m}^{(j)}\beta^3(\frac{x}{2^j}-k)\beta^3(\frac{y}{2^j}-l)\beta^3(\frac{z}{2^j}-m) \quad (6)$$

For $j = 0$ the reconstruction is at its full dimensionality $(N \times N \times N)$ and for $j = 1$ each dimension is divided by two. Once we specify the desired resolution level, we can make use of the downsampling and upsampling of the signal related to the two-scale relation of B-splines. The idea is to downsample the signal to a coarser resolution, then solve equation 6 iteratively and then upsample the signal for getting a finer resolution. The upsampled signal will serve as initialization for the B-spline coefficients at a finer level of resolution. At each level of resolution an error refinement scheme is applied. The multigrid scheme ensures the fast convergence of Eq. 2 to its solution in each dimension. At the end this scheme of course will give our desired reconstructed signal. The resolution coarsening can be defined through the following equations:

$$\begin{cases} A_{j+1} = U_j^T A_j U_j \\ R_{j+1} = U_j^T R_j U_j \\ b_{j+1} = U_j^T b_j U_j \end{cases} \quad (7)$$

where U is a matrix representing the downsampling operation which is achieved by convolving the signal with a circulant matrix corresponding to the filter kernel of the B-spline two-scale relation formula ([19] and [24]). Once

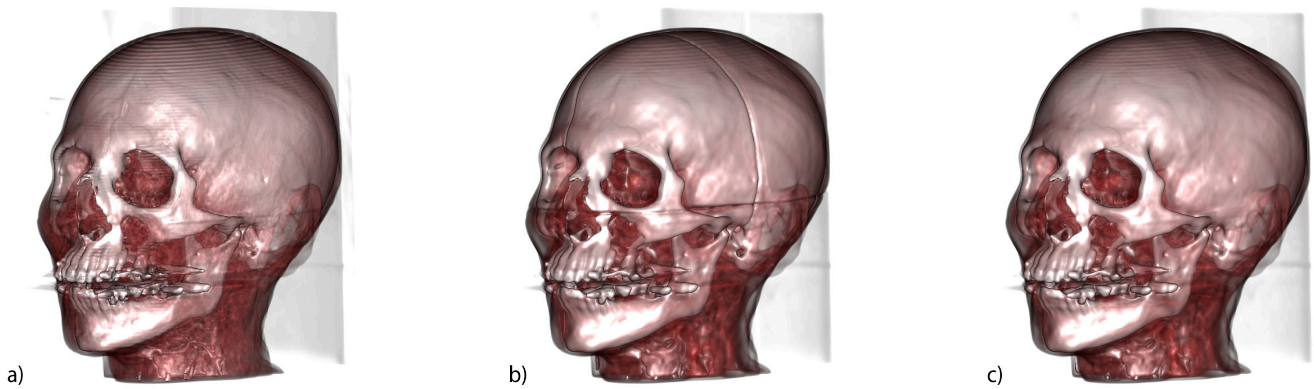

Fig. 1 Rendering of the CT-Head dataset: a) original dataset, b) reconstructed dataset with no overlap between blocks, c) reconstructed dataset with 4-voxel overlap between blocks. 20% of the original dataset points were used in b) and c). The block size used is 128x128x128.

a coarser resolution signal is obtained the equation 5 is solved through a Gauss-Seidel iterator. The advantage of this multigrid interscale technique is that the solution in the lower resolution is more efficient and faster. At the end the upsampling operation is applied for getting the target resolution.

4 Block Structure

A custom solution to the variational method with thin-plate splines as basis functions would require the calculation of the weights as well as the centers of the basis functions. The linear system to find such a solution is mostly ill-posed and has a poor numerical behavior. For solving such a system, when M non-uniformly sampled points are given, we have to deal with an $\mathcal{O}(M)$ complexity. Once the weights are specified, the next step would be to resample the thin-plate splines on a regular grid, which will require an additional $\mathcal{O}(MN^3)$ operations where N is the dataset resolution [4].

The method that we present in this paper has several advantages over the thin-plate spline solution. Since B-splines have a compact support the system is better conditioned. Thanks to the multigrid interscale relation the solution of the system is very efficient and the complexity is reduced to $\mathcal{O}(N^3)$. Furthermore, there is no need for a resampling step since the samples at the grid positions can be obtained by a simple filtering of the B-spline coefficients. Hence, the reconstruction time is dependent not on the number of non-uniform points but on the size of the dataset.

One of the main problems of the variational method is memory requirements. For each grid position we estimate the B-spline basis functions $(\beta^3(x), \beta^3(y), \beta^3(z))$, that vary in accordance with the point coordinates. Each grid point is represented by four values (estimated from the basis functions) for each dimension, hence we are dealing with $4N \times 4N \times 4N$ data. Assuming floating point numbers, for a dataset of size $256 \times 256 \times 256$ we will need 4GB of memory.

This bottleneck brought us to the idea of reconstructing the point set in blockwise fashion. One important issue we faced in the straightforward implementation of block-based reconstruction was the discontinuity problem between neighboring blocks. To overcome this problem we decided to extend the blocks in each direction by a certain number of voxels. Taking into consideration the local support of a cubic B-spline and also the reconstruction results, we extended each block by two voxels in each direction, having thus a 4-voxel overlap between blocks. In Fig. 1 we show the rendering of the CT-Head dataset with and without block-overlap. No visual discontinuities are present when we apply a 4-voxel block overlap.

In order to improve performance, the implementation of the variational method is based on reconstruction of blocks with sizes that are a power of two. The size of a block along each dimension for which the lowest reconstruction time is required can be found through the following reasoning. If we denote with N_x one of the dataset dimensions, e.g. its width, and 2^Q is the maximum block size dimension due to memory constraints, then the optimal block-size is 2^{Q-k} where k minimizes the following function:

$$f(k) = \left\lceil \frac{N_x - L}{2^{Q-k} - L} \right\rceil * 2^{-k} \tag{8}$$

L is the overlap between blocks, $\lceil x \rceil$ is the smallest integer greater or equal to x and $(2^{Q-k} - L) > 1$.

5 Results

In order to create non-uniformly sampled data, we used Cartesian data sets and adaptively sampled them, similar to the approach of Arigovindan et al. [5]. For the adaptive sampling of the data we used a 3D Laplacian kernel (see Equation 9). After convolving the data with this 3D filter we sorted the values according to their magnitudes and retain only the ones that have the biggest absolute values (i.e., 20% of all points in our experiments).

Table 1 Reconstruction timings (in minutes) and RMS(%) errors for different block sizes applied to several datasets.

Dataset		Timings(min) and RMS(%)			
Name	Size	16x16x16	32x32x32	64x64x64	128x128x128
Engine	256x256x128	4.23 — 2.20	2.40 — 2.26	2.88 — 2.24	4.97 — 2.24
Tooth	256x256x160	5.57 — 0.24	3.32 — 0.23	2.95 — 0.23	5.05 — 0.23
CT-Head	256x256x224	7.50 — 3.04	4.30 — 2.92	3.95 — 2.93	5.16 — 2.93
Carp	256x256x512	17.08 — 0.57	9.60 — 0.55	8.83 — 0.50	12.88 — 0.55
CT-Chest	394x394x240	17.78 — 1.33	10.73 — 1.31	9.65 — 1.31	9.83 — 1.32
Christmas-Tree	512x499x512	65.78 — 0.50	38.66 — 0.50	29.35 — 0.50	37.32 — 0.50
Stag-Beetle	832x832x494	177.15 — 0.32	91.36 — 0.31	79.23 — 0.31	95.56 — 0.31

This is equivalent to keeping the points on both sides of boundary regions. Other filters could have been used, but since the idea of non-uniformely sampled datasets is to represent higher frequency regions with more points, convolution with a Laplacian filter would result in a similar effect.

$$L(x, y, z) = \frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} + \frac{\partial^2 V}{\partial z^2} \tag{9}$$

We have tested our method with a variety of datasets and with different parameter settings. Since we obtained the non-uniform representation from a Cartesian (high-resolution) representation, we were able to measure the error of the variational reconstruction. There are several possible error measures. In our work we have chosen the Root Mean Square error (RMS(%)) calculated as follows:

$$RMS(\%) = \sqrt{\frac{\|V_o - V_r\|}{N_x \times N_y \times N_z}} \times \frac{100}{MAX_V} \tag{10}$$

where V_o and V_r are the original and the reconstructed volumes, N_x, N_y and N_z are the dimensions of the volume and MAX_V is the maximum value measured in the original volume.

First we tested various block sizes in order to find the optimal one. Here, optimal refers to minimal reconstruction error and best timing performance. In Table 1 we show the reconstruction timings and errors of several datasets for four different block sizes. As we found in our results, the variation of block size has a negligible effect on the reconstruction errors. However, timings are strongly dependent on the block size. According to these results and the mathematical concept introduced in the previous section, the optimal block size for most of the datasets is 64x64x64. As we reduce the block size, the overlap portion becomes decisive in the timing performance. When $16 \times 16 \times 16$ blocks are used the reconstruction timings are almost twice the timings of the $64 \times 64 \times 64$ block-size cases.

The calculated errors for some well-known and widely used datasets are given in Table 2. For each dataset we take only 20% of the points from a Cartesian dataset after applying a Laplacian filter. Then we show the reconstruction error and the timing (in minutes) required for reconstructing the whole dataset from the non-uniform point set. All the reported errors were computed with a block-based reconstruction, except for the Neghip and Hydrogen dataset, which have dimensions that allow non-block-based reconstructions. The testings were done with a Dual Core AMD Opteron 2.41GHz processor machine with 8GB of RAM. Since our program is single threaded we are using only one dedicated processor during the reconstruction process.

Table 2 Reconstruction errors and timings (in minutes) for the variational method. Each reconstruction is based on 20% of the points of the original dataset. Optimal block size is used in the reconstruction process.

Dataset	Block size	RMS(%)	Timing
Neghip (64x64x64)	64x64x64	2.14	0.03
Hydrogen (128x128x128)	128x128x128	0.17	0.32
Lobster (301x324x56)	31x128x64	1.21	1.36
Statue Leg (341x341x93)	128x128x128	0.95	2.13
Engine (256x256x128)	64x64x32	2.24	2.38
Tooth (256x256x160)	64x64x64	0.23	2.95
CT-Head (256x256x224)	64x64x128	2.93	3.95
Foot (256x256x256)	64x64x64	2.16	4.80
Carp (256x256x512)	64x64x64	0.50	8.83
CT-Chest (394x394x240)	64x64x128	1.31	7.95
Christmas-Tree (512x499x512)	64x64x64	0.50	29.35
Stag-Beetle (832x832x494)	64x64x64	0.31	79.23

For the rendering of the datasets we have used VolumeShop [6] which is an open source volume rendering platform. The volumes are rendered with a GPU-based raycaster with a sampling step of 0.25. For some of the datasets the rendered images are shown in Fig. 2, 3 and 4.

In non-uniform reconstruction approaches which apply exact interpolation techniques, the number of points used for reconstruction highly affects the reconstruction error. In quasi-interpolation approaches like the varia-

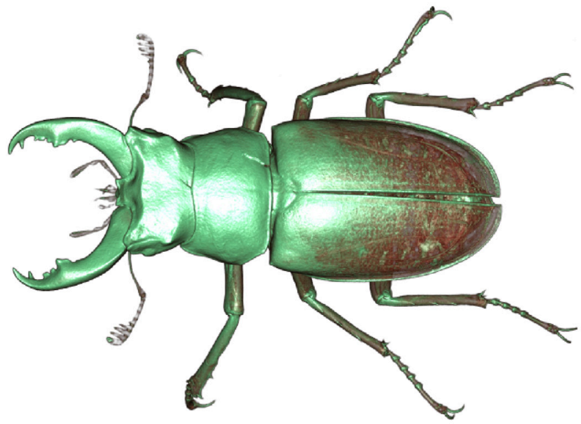

Fig. 2 Rendering of the Stag Beetle dataset (832x832x494). Original dataset (left) and reconstructed one (right) using 20% of points. The RMS(%) error is 0.31.

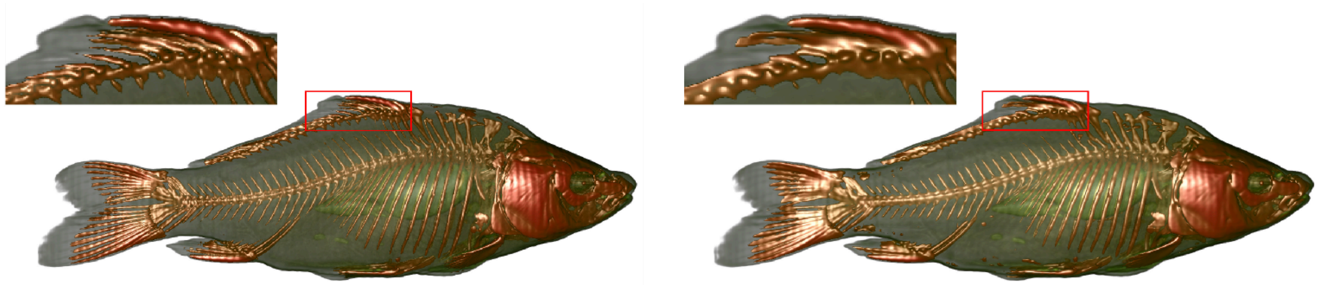

Fig. 3 Rendering of the Carp dataset (256x256x512). Original dataset (left) and reconstructed one (right) using 20% of points. The RMS(%) error is 0.5.

tional method that we have presented here there is always a certain limit where even if we increase the number of points the reconstruction error will remain stable. This is strongly connected to the regularization parameter which controls the smoothing. In our experiments we concluded that we can achieve a stable reconstruction rate when using 15%-20% of the points.

Smoothing is another factor that affects the reconstruction error. Smoothing lowers the noise levels but it also eliminates details in the data. A compromise is required between accuracy and smoothness. In Fig. 5 we display the CT-Chest dataset for different levels of smoothing. In the first reconstructed image there is too much visual noise due to low smoothing. In the rightmost image the high frequencies are removed due to the high smoothing operator. Although we do not aim for a compression technique, our method achieves a reduction of up to 60% of the original dataset size when 20% of points are kept for reconstruction. We do not apply any compression technique, but just save the coordinates and values in a slice/row basis. The reconstruction error varies with the dataset and the worst case scenarios were in the range of 2-3%, which is comparable to other techniques [11], [13], [26]. The reconstruction timings are difficult to compare since previous papers have not reported timings for specific datasets. However, our method is several orders of magnitude faster than other RBF-based reconstruction techniques introduced by Jang et al. [11], [12]. Compared to the work of Juba and Amitabh [13] our encoding timings are in the same range.

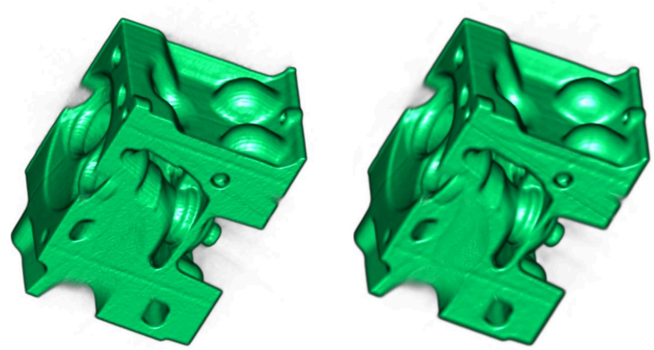

Fig. 4 Rendering of the Engine dataset (256x256x128). Original dataset (left) and reconstructed one (right) using 20% of points. The RMS(%) error is 2.24.

6 Conclusion and Future Work

In this paper we introduce a method for 3D non-uniform data reconstruction based on a variational approach. We present how to reconstruct on a uniform grid, datasets that were given as point sets consisting of 20% of the original number of points.

The algorithm is based on B-spline approximation. B-spline coefficients are efficiently determined by minimizing the approximation error at each sample point. A tradeoff exists between the smoothness and the approximation accuracy depending on the number of points kept for reconstruction. We also improved the variational reconstruction to overcome the memory constraints by introducing a block-based approach and a simple control over the continuity at block-boundary regions.

The next step in our research is to apply the presented technique to data originally given on a non-uniform grid. Further research is needed to optimize and automatize the parameter settings controlling the smoothness and reconstruction for different resolutions. A GPU-implementation for fast reconstruction of non-uniform data is another important issue we want to deal with in our future research. All in all, we believe that this approach opens new possible directions in the research area of 3D non-uniform signal reconstruction and visualization.

Acknowledgements The work presented in this publication is carried out as part of the Point-based Volume Graphics project (www.cg.tuwien.ac.at/research/vis/pvg) supported by the Austrian Science Fund (FWF) grant no. P18547, and partially funded by the National Science and Engineering Research Council of Canada. Special thanks to Muthuvel Arigovindan for providing the basic codes of the variational reconstruction method. We would also like to thank Michael Unser for helpful discussions.

References

1. Aldroubi, A., Feichtinger, H.: Non-uniform sampling: Exact reconstruction from non-uniformly distributed weighted-averages. In: Proceedings of the international conference of computational harmonic analysis, vol. 1, pp. 1–8 (2002)
2. Aldroubi, A., Gröchenig, K.: Nonuniform sampling and reconstruction in shift-invariant spaces. SIAM Rev. **43**(4), 585–620 (2001)
3. Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., Silva, C.T.: Point set surfaces. In: Proceedings of IEEE Visualization, pp. 21–28 (2001)
4. Arigovindan, M.: Variational reconstruction of vector and scalar images from non-uniform samples. Ph.D. thesis, Ecole Polytechnique Federale de Lausanne (2005)
5. Arigovindan, M., Suhling, M., Hunziker, P., Unser, M.: Variational image reconstruction from arbitrarily spaced samples: A fast multiresolution spline solution. In: IEEE Transactions on Image Processing, vol. 14, pp. 450–460 (2005)
6. Bruckner, S., Gröller, M.E.: VolumeShop: An interactive system for direct volume illustration. In: IEEE Visualization 2005, pp. 671–678 (2005)
7. Buhmann, M.D.: Radial Basis Functions: Theory and Implementations. Cambridge University Press (2003)
8. Chen, B., Nguyen, M.X.: Pop: A hybrid point and polygon rendering system for large data. In: IEEE Visualization, pp. 45–52 (2001)
9. Garrity, M.: Ray tracing irregular volume data. Workshop on Volume visualization **24**, 35:50 (1990)
10. Haber, J., Zeilfelder, F., Davydov, O., Seidel, H.P.: Smooth approximation and rendering of large scattered datasets. In: IEEE Visualisation, pp. 341–347 (2001)
11. Jang, Y., Botchen, R.P., Lauser, A., Ebert, D.S., Gaither, K.P., Ertl, T.: Enhancing the interactive visualization of procedurally encoded multifield data with ellipsoidal basis functions. In: Computer Graphics Forum, vol. 25, pp. 587–596 (2006)
12. Jang, Y., Weiler, M., Hopf, M., Huang, J., Ebert, D.S., Gaither, K.P., Ertl, T.: Interactively visualizing procedurally encoded scalar fields. In: Proceedings of Joint Eurographics-IEEE TCVG Symposium on Visualization, pp. 35–44 (2004)
13. Juba, D., Varshney, A.: Modelling and rendering large volume data with Gaussian radial basis functions. Tech. rep., University of Maryland (2007)
14. Kähler, R., Abel, T., Hege, H.C.: Simultaneous gpu-assisted raycasting of unstructured point sets and volumetric grid data. In: Proceedings of IEEE/EG International Symposium on Volume Graphics, pp. 49–56 (2007)
15. Lee, S., Wolberg, G., Shin, S.Y.: Scattered data interpolation with multilevel B-splines. In: IEEE Transactions on Visualization and Computer Graphics, vol. 3, pp. 228–244 (1997)
16. Ljung, P., Lundström, C., Ynnerman, A.: Multiresolution interblock interpolation in direct volume rendering. In: Proceedings of Joint Eurographics-IEEE TCVG Symposium of Visualization, pp. 259–266 (2006)
17. Muigg, P., Hadwiger, M., Doleisch, H., Hauser, H.: Scalable hybrid unstructured and structured grid raycasting. In: IEEE Transactions on Visualization and Computer Graphics, vol. 13, pp. 1592–1599 (2007)
18. Neophytou, N., Mueller, K., McDonnell, K.T., Hong, W., Guan, X., Qin, H., Kaufman, A.E.: GPU-accelerated volume splatting with elliptical RBFs. In: Proceedings of Joint Eurographics-IEEE TCVG Symposium on Visualization, pp. 13–20 (2006)
19. Petrosian, A.A., Meyer, F.G.: Wavelets in Signal and Image Analysis: From Theory to Practice. Springer (2001)
20. Rössl, C., Zeilfelder, F., Nürnberger, G., Seidel, H.P.: Spline approximation of general volumetric data. In: G. Elber, N. Patrikalakis, P. Brunet (eds.) Proceedings of the 9th ACM Symposium on Solid Modeling and Applications, pp. 71–82 (2004)
21. Silva, C., Mitchell, J., Kaufman, A.: Fast rendering of irregular grids. In: ACM/IEEE Volume Visualization Symposium, pp. 15–22 (1996)
22. Thévenaz, P., Blu, T., Unser, M.: Image interpolation and resampling. In: Handbook of Medical Imaging, Processing and Analysis. Academic Press (2000)
23. Tuncer, T.E.: Block-based methods for the reconstruction of finite-length signals from nonuniform samples. In: IEEE Transactions on Signal Processing, vol. 55, pp. 530–541 (2007)
24. Vetterli, M., Kovacevic, J.: Wavelets and subband coding. Prentice Hall, Englewood Cliffs, NJ, USA (1995)
25. Weiler, M., Kraus, M., Merz, M., Ertl, T.: Hardware-based raycasting for tetrahedral meshes. In: Proceedings of IEEE Visualisation, pp. 71–78 (2004)
26. Welsh, T., Mueller, K.: A frequency-sensitive point hierarchy for images and volumes. In: IEEE Visualization, pp. 425–432 (2003)
27. Zwicker, M., Pfister, H., Baar, J.V., Gross, M.: Surface splatting. In: Proceedings of ACM SIGGRAPH, pp. 371–378 (2001)

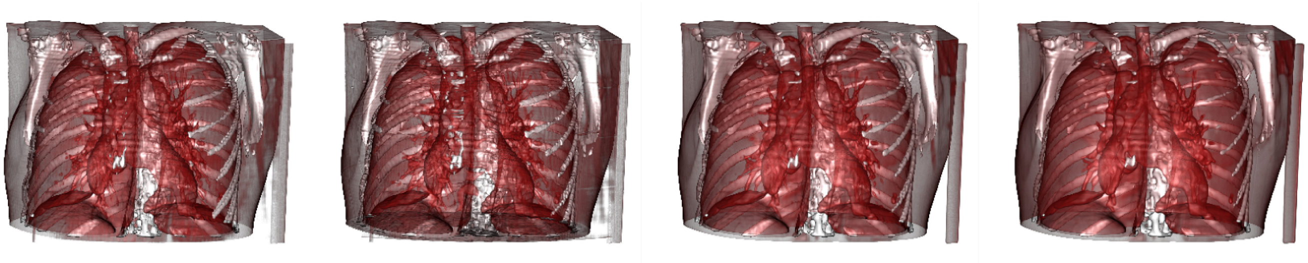

Fig. 5 Examples of reconstructions using different levels of smoothing of the CT-Chest dataset (394x394x240). From left to right: original dataset, data reconstructed with lower to higher levels of smoothing. The reconstruction errors (RMS(%) are, respectively, 2.43, 1.34 and 1.76. 20% of the original data points were used.

Erald Vuçini is a PhD reseacher at the Institute of Computer Graphics and Algorithms, Vienna University of Technology. He received a BSc degree with high honours in Computer Engineering from Yildiz Technical University (2004), and a MSc degree in Computer Engineering from Istanbul Technical University (2006). He won for two consecutive years the Albanian National Mathematical Olympiad for high schools (1998, 1999), won the Bronze Medal in the Balkan Mathematical Olympiad (Ohrid, Macedonia 1999) and was awarded the Gold Medal of Albania for his excellent results in the end of the high school studies (1999). His research interests include volume visualization, image processing and reconstruction techniques on non-uniform point sets.

Torsten Möller is an Associate Professor at the School of Computing Science at Simon Fraser University. His research interests include the fields of Visualization and Computer Graphics, especially the mathematical foundations thereof. He is the director of Vivarium, co-director of the Graphics, Usability and Visualization Lab and serves on the Board of Advisors for the Centre for Scientific Computing at Simon Fraser University. Dr. Möller has been appointed Vice Chair for Publications of the IEEE Visualization and Graphics Technical Committee. He is currently serving on the steering committee of the Symposium on Volume Graphics as well as the paper co-chair for the 2008 Eurographics/IEEE Symposium on Visualization. Further, he is an associate editor for the IEEE Transactions on Visualization and Computer Graphics as well as the Computer Graphics Forum. He received his PhD in Computer and Information Science from Ohio State University in 1999 and a Vordiplom (BSc) in mathematical computer science from Humboldt University of Berlin, Germany. Dr. Möller has been papers co-chair for the 2003 Graphics Interface conference, the 2006 Workshop on Volume Graphics (VG'06), the 2007 Eurographics/IEEE Symposium on Visualization, as well as the Visualization track of the 2007 International Symposium on Visual Computing. He has been the general chair for the 2007 Workshop on Volume Graphics. He has also co-organized the 2004 Workshop on Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration at the Banff International Research Station.

M. Eduard Gröller is Associate Professor at the Institute of Computer Graphics and Algorithms (ICGA), Vienna University of Technology. In 1993 he received his PhD from the same university. His research interests include computer graphics, flow visualization, volume visualization, medical visualization, and information visualization. He is heading the visualization group at ICGA. The group performs basic and applied research projects in the area of scientific visualization (http://www.cg.tuwien.ac.at/research/vis/). Dr. Gröller is a scientific proponent and member of the Scientific Review Committee of the VRVis Kplus center of excellence (http://www.vrvis.at/) The center performs applied research in virtual reality and visualization. Dr. Gröller is adjunct professor of computer science at the University of Bergen, Norway (since 2005). He is head of the working group on computer graphics of the Austrian Computer Society and member of IEEE Computer Society, ACM (Association of Computing Machinery), GI (Gesellschaft für Informatik), OCG (Austrian Computer Society).