

On Visualization and Reconstruction from Non-Uniform Point Sets using B-splines

Erald Vućini¹ and Torsten Möller² and M. Eduard Gröller¹

¹Vienna University of Technology, Austria

²Simon Fraser University, BC, Canada

Abstract

In this paper we present a novel framework for the visualization and reconstruction from non-uniform point sets. We adopt a variational method for the reconstruction of 3D non-uniform data to a uniform grid of chosen resolution. We will extend this reconstruction to an efficient multi-resolution uniform representation of the underlying data. Our multi-resolution representation includes a traditional bottom-up approach and a novel top-down hierarchy for adaptive hierarchical reconstruction. Using a hybrid regularization functional we can improve the reconstruction results. Finally, we discuss further application scenarios and show rendering results to emphasize the effectiveness and quality of our proposed framework. By means of qualitative results and error comparisons we demonstrate superiority of our method compared to competing methods.

Categories and Subject Descriptors (according to ACM CCS): Numerical Analysis [G.1.2]: Spline and piecewise polynomial approximation.—Image Processing and Computer Vision [I.4.5]: Reconstruction.—Computer Graphics [I.3.5]: Computational Geometry and Object Modeling.—

1. Introduction

The traditional sources of volumetric data are simulations as well as data acquisition devices on uniform (Cartesian) lattices. In an effort to study larger and more complex problems, there has been a move toward non-uniform data representations, since they offer a way of adapting the measure location (or sample points) according to the importance (variance) of the data. Examples include a) simple data loss during data communication in sensor networks, b) Doppler measurements or other novel acquisition models (polar or spiral) for tomography and magnetic resonance imaging, c) adaptive and moving mesh approaches in mathematical simulations in the physical sciences, and d) particle simulations.

While the acquisition of data on non-uniform grids has become wide-spread, the available tools for processing, filtering, analysis, and rendering of data are most efficient on uniform representations. We do not make use of the explicit neighborhood information in non-uniform grids in this work and hence they are used like point-sets. There are two competing efforts to deal with non-uniform data: a) create novel and efficient tools that directly work on them, or b) convert the non-uniform representation into an efficient intermediate uniform representation and apply standard tools. Both ap-

proaches have advantages and disadvantages. In this paper we make a contribution towards the latter approach. Among other things, this will allow us to better exploit the capabilities of modern GPUs. In Section 2 we will contrast these two approaches further and review alternative works.

In order to find the best way to transform the non-uniform data onto a uniform grid, we first need to analyze the nature of the given data. One reason for non-uniformity is the ability to capture different scales of information density (e.g. mathematical simulation of shock waves). Another reason for non-uniform data representations could come from imprecise measurement devices (e.g. ultrasound) or sparse representations (e.g. compressive sensing). While in the former case multi-resolution representations might be most suitable, in the latter case a single resolution representation might be all what is needed. Therefore, we develop some heuristics, based on a statistical analysis, to adapt to each scenario.

In this paper we propose a uniform representation consisting of B-spline coefficients which define a C^2 continuous function across the whole volume. Our main contributions are: a) a statistical approach for selecting the resolution of reconstruction for non-uniform datasets (Section 4.1), b) a bottom-up multi-resolution pyramid (Section 4.2), c) a novel

top-down adaptive multi-resolution scheme (Section 4.3), and d) a novel hybrid regularization functional for the variational approach leading to improved accuracy (Section 5).

We compare our approaches with similar approaches, that perform a resampling of the data domain on a variety of different data sets (Section 6). Conclusions and ideas for future work are summarized in Section 7.

2. Related Work

There is a considerable body of literature on the rendering of non-uniform data without any resampling steps (see e.g. [KSW01, WKME04]). While there are very good reasons to adapt such an approach for rendering, we postulate here that an intermediate transformation onto a regular data structure opens up the possibilities for much more sophisticated data processing in general and henceforth focus on such a pipeline.

A number of approaches have been proposed for the reconstruction of non-uniformly sampled data, especially for one- and two-dimensional signals. Most of the methods are based on the reconstruction of the data by solving large systems of equations [FGS95, GS04]. However, their use of basis functions with infinite support makes them impractical to use for real-time visualization applications, where finite support reconstruction kernels are desired. Park et al. [PLKO06] presented an efficient discretization of Sibson's natural-neighbor interpolation for 2D and 3D data fitting. They ensure C^1 continuity, but they fail to report experiments for real non-uniform point sets. Nielson [Nie93] presented an overview of several approximation techniques for non-uniform point sets. While each technique performs best only in particular cases, the use of local compact operators is considered the fastest approach.

The necessity of basis functions with infinite support comes from a desire of reconstructing bandlimited signals. Unser [Uns00] suggests to replace the concept of bandlimitedness by minimum-error projection on a space of shift-invariant functions. A more general overview on modern non-uniform reconstruction techniques in shift-invariant spaces has been summarized by Aldroubi and Gröchenig [AG01]. Perhaps the most popular shift-invariant spaces are based on Radial Basis Functions (RBFs). They have been used for surface [ABCO*01, OBS04, CBC*01] as well as volumetric [JWH*04, JBL*06, WM03] approximation and reconstruction techniques. Our focus is on volumetric data.

Jang et al.'s [JWH*04] method is formulated as an iterative algorithm for finding the centers and weights of the RBFs using a PCA-based clustering technique by applying truncated Gaussians as basis functions. This technique suffers from high-encoding times and is best suited for locally isotropic structures. Later they [JBL*06] adapt their technique to ellipsoidal basis functions (EBFs). The high computational cost is still the main bottleneck of this approach.

Our approach uses B-splines as basis functions. B-splines, with their smoothness and compact support, offer optimal conditions for fast and accurate reconstruction results. They are related to RBF-based approaches since B-splines are very good approximators of thin-plate splines, which in turn

are widely used RBFs in approximation theory. Further, our B-spline basis will be anchored on a regular grid, preventing the need to store the grid geometry explicitly and opening the door for efficient multi-resolution representations.

Arigovindan et al. [ASHU05] proposed to use B-splines in a multi-grid framework for the reconstruction of non-uniform 2D data. Vučini et al. [VMG08] extended these ideas for 3D volumes and proposed a block-based variational reconstruction technique for large datasets. In this paper we build on these ideas.

Multi-resolution approaches have been introduced to improve the rendering speed as well as the quality of the data representation adaptively while minimizing the memory overhead [Sam05]. Based on the data structure, multi-resolution schemes can be divided into regular schemes [LHJ99] (e.g., octrees) and irregular schemes [KSH03] (e.g., adaptively refined meshes). Our scheme adopts concepts from both classes. It is a multi-level hierarchy where the first level represents the coarse resolution and has a regular representation. Additional levels encode the errors and are refined adaptively. The structure of the refined cells is again regular.

The proper continuous interpolation between different octree levels has remained a challenge in multi-resolution volume rendering. Several proposed approaches ensure only a C^0 continuity in their rendering algorithms. The multi-resolution function in our approach is a hierarchical sum of C^2 continuous functions, ensuring the C^2 continuity over the entire domain.

3. Variational Reconstruction Basics

Variational reconstruction is a well-known technique applied to solving ill-posed problems such as the reconstruction from non-uniform point sets. The variational functional is formulated as the sum 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.

Given a set of sample points, $\mathbf{p}_i = (x_i, y_i, z_i)$, $i = 1, 2, \dots, M$, let f_i be the scalar values associated with \mathbf{p}_i . We define the B-spline approximation through the form:

$$F(\mathbf{p}) = \sum_{\mathbf{k}=0}^{\mathbf{N}-1} c_{\mathbf{k}} \beta^3(\mathbf{p} - \mathbf{k}) \quad (1)$$

where $\beta^3()$ is the tensor product of cubic B-spline basis functions and where we denote with $\mathbf{N} = (N_x, N_y, N_z)$ the resolution of the axis-aligned bounding box of our non-uniform data set. Although cubic B-splines do not enjoy the interpolation property, they have the maximal order of approximation for a given integer support, providing the best quality for a given computational cost [TBU00, Uns99]. In order to find the coefficients $c_{\mathbf{k}}$ the following cost function is minimized:

$$C(F) = \sum_{i=1}^M \|F(\mathbf{p}_i) - f_i\|^2 + \lambda \iiint \|D^p F\|^2 dx dy dz \quad (2)$$

where λ is a parameter that controls the smoothness and

the second term is the regularization functional that uses Duchon's seminorms $D^p F$ [Duc79].

The key idea of the variational reconstruction is to express the second term of Eq. 2 by means of the first term and then minimize the error $C(F)$ with regard to the B-spline coefficients c_k in Eq. 1. Once we solve the equation, we have $N_x \times N_y \times N_z$ B-spline coefficients defined at the reconstruction-grid positions. We can compute $F(\mathbf{p})$ (a C^2 -continuous function) at any position inside the volume, from a 64 point neighborhood using Eq. 1. For a deeper insight into the method we refer the reader to Arigovindan [ASHU05].

4. Single and Multi-Resolution Reconstruction

The key issue in resampling a non-uniform point set into a uniform representation is the selection of the proper resolution. This will be the central question we are trying to answer in this section. We first assume that we can only afford a single resolution and we make suggestions on how this resolution can be best obtained. This is applicable for non-uniform data, where the distribution of samples is even (in the sense of a discrepancy measure), e.g., ultrasound data or seismic data.

4.1. Single Resolution

Increasing the resolution results in a decrease of the error, since the oscillations in the data can be captured with more precision. Finding the optimal resolution ($N_x \times N_y \times N_z$) would therefore require estimating the error. We propose to do so by simply looking at the error within a single grid cell. If there are many non-uniform points crammed into a cell, and their standard deviation σ_k is large, the cell might be too large. Therefore, we propose to approximate the error by using the average standard deviation, defined as:

$$\sigma_{avg} = \frac{\sum_{k=0}^{N-1} \sigma_k}{N_x \cdot N_y \cdot N_z} \quad (3)$$

as an indicator for the proper uniform grid resolution. Empty cells are considered as cells with zero standard deviation.

In Section 6.1 we analyze a number of data sets in order to arrive at a reasonable threshold. Our idea is motivated by the strong correlation observed between the reconstruction error and the average standard deviation of point values.

4.2. Bottom-up Multi-Resolution Pyramid (BMRP)

There are many scenarios where we observe a large variance in the density of the data points. Hence, finding a single resolution to minimize the error in a uniform representation leads to very large data sets with lots of redundancy. In such a case, it is typical to encode the data in a multi-resolution pyramid. One usually starts with the highest resolution and gradually finds coarser representations. To tackle this problem we propose a multi-resolution scheme based on the interscale relation of the B-splines of odd degree:

$$\beta^n\left(\frac{x}{2^j}\right) = \sum h(k) \beta^n\left(\frac{x}{2^{j-1}} - k\right) \quad (4)$$

where $h(k)$ is the binomial filter [AL96].

We consider a 3D signal being represented by a set of coefficients $c^{(j)}$ at scale j :

$$F^{(j)}(\mathbf{p}) = \sum_{k=0}^{\frac{N-1}{2^j}} c_k^{(j)} \beta^3\left(\frac{\mathbf{p}}{2^j} - \mathbf{k}\right) \quad (5)$$

Using results from multi-resolution analysis, the same signal can be represented at a finer scale ($j-1$) by the coefficients $c^{(j-1)}$, which are obtained by first upsampling $c^{(j)}$ and then filtering with $h(k)$. In the same fashion, by using the inverse transform of Eq. 4 we can filter and downsample the $c^{(j)}$ to get a projection of these coefficients to the space spanned by the coarser coefficients $c^{(j+1)}$. For a specific scale j we denote the upsampling and downsampling process by U_j and D_j respectively [ASHU05].

We initially estimate the coefficients at the finest resolution and then process them to create a top-down hierarchy of coarser resolutions. We obtain the coefficients of the coarser resolution ($j+1$) by downsampling from the finer resolution (j): $c^{(j+1)} = D_j c^{(j)}$. Ignoring the finer resolution completely, would create an error at scale (j), i.e., $e^{(j)} = c^{(j)} - U_j c^{(j+1)}$. By saving the coarser scale coefficients and also part of the error volumes (where the error is high) we can reconstruct the data at a finer resolution with little or no error.

For example, for a signal which we want to reconstruct with the finest resolution of $N_x \times N_y \times N_z$, using three levels of hierarchy, we first estimate the finest coefficients $c^{(0)}$ by minimizing Eq. 2. Then, by using the interscale relation we estimate $c^{(1)}$ and $c^{(2)}$, as well as $e^{(0)}$ and $e^{(1)}$. In our scheme we save only $c^{(2)}$ and parts of $e^{(0)}$ and $e^{(1)}$, which we denote by $e_p^{(0)}$ and $e_p^{(1)}$. When visualizing the data, we can either use the coefficients $c^{(2)}$ for a coarse resolution representation, or the approximations $\tilde{c}^{(1)} = U_1 c^{(2)} + e_p^{(1)}$ or $\tilde{c}^{(0)} = U_0 \tilde{c}^{(1)} + e_p^{(0)}$ for a finer resolution representation.

This approach requires an explicit intermediate representation of the finest resolution, which might not be feasible computationally. Hence, we propose a novel algorithm to build an adaptive multi-resolution data structure.

4.3. Adaptive Multi-Resolution (AMR)

Whenever our σ_{avg} demands a resolution that is too large to handle directly, we decide to create a multi-resolution representation starting from a coarse resolution first. This prevents us from having to compute the highest resolution explicitly. Estimating a reasonable coarse resolution is typically tied to hardware constraints. One should not choose a very high resolution, such that it compromises real-time rendering or analysis performance, yet, it should not be too coarse to avoid storing too many levels in the hierarchy. We call this maximum resolution N_{max} . Next, we determine whether each cell of the coarse resolution should be subdivided or not, i.e., whether it is composite or not. This is done based on an error criterion. These steps applied recursively will create a multi-resolution hierarchy, that adapts to

the variance in the data. What follows is pseudo-code outlining this algorithm as well as the procedure to use the multi-resolution hierarchy to determine the value of the function. We will use the notation introduced in Section 3.

Procedure 1 $c_V^{(j)}$ = AMR(Volume V, point set P_V with values f_V , level j)

- 1: determine the resolution N for volume V ($\leq N_{\max}$)
- 2: determine the B-spline representation F_V with coefficients c_V
- 3: **for all** cells U of grid V **do**
- 4: estimate reconstruction error $\epsilon_U = F_V - f_V$ of all points P_U inside the cell U
- 5: **if** (ϵ_U is too large) AND ($|P_U| > \tilde{M}$) **then**
- 6: $c_U^{(j+1)}$ = AMR(U , P_U with values ϵ_U , level $j+1$)
- 7: **end if**
- 8: **end for**

Procedure 2 Evaluate Function F_V at (x, y, z) for Volume V

- 1: evaluate $F_V(x, y, z)$ by using coefficients c_V
- 2: **if** (x, y, z) is in composite cell U **then**
- 3: Evaluate Function F_U at (x, y, z) for Volume U
- 4: return $F_V(x, y, z) + F_U(x, y, z)$
- 5: **else**
- 6: return $F_V(x, y, z)$
- 7: **end if**

Procedure 1 starts by determining the resolution of the volume V (line 1). This is done based on the σ_{avg} threshold. In order to create a balanced tree, the chosen resolution cannot exceed a maximum resolution N_{\max} . Given a resolution we can then determine the B-spline coefficients c_V for this resolution using Equation 2 (line 2). For each cell of this resolution, we will determine whether we should recursively subdivide (line 5). This is based on the cell reconstruction error (computed in line 4) as well as whether the number of points in the cell is above a threshold \tilde{M} . \tilde{M} is used to prevent the algorithm from subdividing cells with a low number of points inside. Once we determine that we should subdivide the given cell, we only reconstruct the error function ($F_V - f_V$) in line 6.

Procedure 2 is used during the raycasting process. It chooses the B-spline coefficients to use in Eq. 1 accordingly. If the point is in a composite cell, it recursively adds the error estimation of each level of the hierarchy.

5. Improving Regularization

Regularization provides a framework for converting ill-posed problems into well-posed ones by restricting the domain of possible solutions via smoothing constraints. Arigovindan et al. [ASHU05] suggest using Duchon's seminorms ($D^p F$) for regularization. For $p = 1$ and $p = 2$ this norm yields a minimization of an energy functional associated with a membrane and a plate model respectively [Duc79]. Here, we propose a hybrid regularization functional in order to reduce reconstruction errors for anisotropic signals. The main motivation for this idea lies in the fact that cubic B-splines have a better response to high frequencies,

that can be detected and preserved by convolving the signal with a Laplacian regularization kernel [Uns00]. The regularization functional consists of the sum of second degree derivatives if cubic B-splines are used as a basis function for reconstruction:

$$R_2(F, \lambda) = \lambda \iiint (d_x^2 F)^2 + (d_y^2 F)^2 + (d_z^2 F)^2 dx dy dz \quad (6)$$

where $(d_x^2 F)^2$, $(d_y^2 F)^2$ and $(d_z^2 F)^2$ are the directional second degree derivatives of F . In order to deal with anisotropic characteristics we extend Eq. 6, so that we can achieve a different regularization in each direction.

$$R_2(F, \lambda_x, \lambda_y, \lambda_z) = \iiint \lambda_x (d_x^2 F)^2 + \lambda_y (d_y^2 F)^2 + \lambda_z (d_z^2 F)^2 dx dy dz \quad (7)$$

Eq. 7 provides a very good application scenario in cases when we have apriori knowledge of the directional variance of the data we are reconstructing. In other cases, we suggest a pre-estimation of the variance of directional gradients and hence setting λ_x , λ_y , λ_z accordingly to an inversely proportional formula. A high variance in the x-direction for example, means we should set a lower λ_x and vice-versa. As opposed to Duchon's seminorms, the regularization terms introduced in Eq. 6 and Eq. 7 do not enjoy the rotational invariance property.

6. Implementation and Results

Our test platform is an Intel Dual Core 2.70 GHz processor machine with 6GB of RAM. We tested our framework on several data either given from originally non-uniform data, or obtained by taking Laplacian points from a uniform dataset [VMG08]. Unless stated otherwise our experiments are using Eq. 6 where we set $\lambda = 0.3$. In order to evaluate the quality of our reconstruction, we use the Root Mean Square error (RMS) defined as follows:

$$RMS = \sqrt{\frac{\sum_i^M (F(x_i, y_i, z_i) - f_i)^2}{M}} \times \frac{100}{Max_{value}} \quad (8)$$

where Max_{value} is the maximum value in the given point set.

6.1. Determining the threshold σ_{avg}

In order to determine an appropriate resolution for a uniform grid representation of our non-uniform data points, we would ideally vary the value of N_x , reconstruct using this resolution and measure the error. N_y and N_z are determined by the proper aspect ratio of our underlying axis-aligned bounding box enclosing the given non-uniform data points. Unfortunately, this is computationally infeasible. Hence, in the search for a good heuristic, we did indeed reconstruct a number of test data sets under various resolutions. Then we measured the reconstruction RMS of the point set as well as the average variance of point values (as opposed to the reconstruction error) in each cell, according to Equation 3. The resulting relationship for the Bypass data set can be seen in Fig. 1. We found a similar relationship in all test data sets (a complete listing can be found in the supplemental material).

Based on this analysis, we suggest, that a value of $\sigma_{avg} = 0.05$ yields a low RMS. In order to use this effectively for an unknown data set, we start from a low value of N_x and increase it until the value of σ_{avg} falls below the desired value. σ_{avg} monotonically decreases with the increase of resolution. Typically we double N_x in each step, using $N_x = 8$ as a starting resolution. Once σ_{avg} has a value lower than the specified threshold we refine the exact value of N_x with a binary-search method.

Arigovindan et al. [ASHU05] suggested a heuristic such that the number of uniform points is 4-5 times the number of non-uniform points (i.e., $N_x \cdot N_y \cdot N_z \approx 5 \cdot M$). With our heuristic, on the other hand, we sometimes find $\approx 4 \cdot M$ to be good enough.




Figure 1: Graph showing relation of RMS to σ_{avg} for the Bypass dataset. A hairline shows the suggested threshold of σ_{avg}

6.2. Single-Resolution and Improving Regularization

A GPU-based raycaster is employed for single resolution rendering. The renderer is developed inside the VolumeShop platform [BG05]. The rendering integral is evaluated at each point along the ray by using Eq. 1. On the fly gradient estimation is used by taking partial derivatives of the function defined in Eq. 1 and applying the fact that the derivative of a B-spline of degree n is a B-spline of degree $n - 1$ [Uns99]. It can be defined as follows:

$$\frac{\partial \beta^n(x)}{\partial x} = \beta^{n-1}(x + \frac{1}{2}) - \beta^{n-1}(x - \frac{1}{2}) \quad (9)$$

Table 1 compares our hybrid regularization to our previous work [VMG08]. We used uniform data, computed and thresholded their Laplacian to keep 20% of the original points and reconstructed the complete uniform data set from this sparse representation.

While computational times remain the same, we observe a 20%-60% improvement in the reconstruction error compared to our previous results [VMG08].

Table 2 compares our method to the work presented by Jang et al. [JWH*04, JBL*06]. Our method has lower reconstruction errors and improves computation time by several orders of magnitude.

So far we computed the error only at the points used for the reconstruction (the input-points). However, an important measure of the quality of reconstruction is the quality of the

Table 1: RMS errors and computation times (in minutes) for different non-uniform datasets created by taking 20% of Laplacian points from their original uniform representation. Here we use our hybrid regularization (Eq. 6) with $\lambda = 0.3$, while previously [VMG08] the results are obtained by using Duchon's regularization and $\lambda = 1.0$.

Dataset		RMS and Times (min)	
Name	Size	Our method	[VMG08]
Engine	256x256x128	0.94 1.28	2.24 1.28
Tooth	256x256x160	0.18 1.88	0.23 1.88
CT-Head	256x256x224	1.17 2.60	2.93 2.60
CT-Chest	394x394x240	0.60 5.08	1.31 5.08
Carp	256x256x512	0.25 5.73	0.50 5.73

reconstruction at other locations. Since the ground truth is not given, this is typically hard to evaluate. Instead we took a dataset given on a uniform grid, created a non-uniform version by removing 80% of its values (according to a Laplacian threshold) and reconstructed it on the original grid. In Fig. 2 we show such a scenario using the Laplacian points from the Tooth dataset. We compute the errors at the non-uniform points used for the reconstruction as well as all the original uniform data points. While our approach has the same error rates in both cases, Jang et al. [JWH*04, JBL*06] show a significant increase in the reconstruction error at the non-input points, which is quite visible.

The Synthetic Chirp is a synthetic radial sinusoidal wave with a spatial frequency that decreases from the center to the edges. We create a non-uniform point set by evaluating the Chirp function for 75,000 random points. The function is changing very fast in the xy plane (the screen plane), while it is changing very slowly along the z axis. In order to reduce the reconstruction error a lower smoothness control along the xy plane is required. In Fig. 3 we show the original dataset, our reconstruction with a regularization term as defined in Eq. 6 and our reconstruction with a directional regularization term as defined in Eq. 7. All three cases were reconstructed on a $64 \times 64 \times 64$ grid, selected based on the $\sigma_{avg} = 0.05$ threshold. There is a clear improvement in the visual quality when directional regularization is used; the error is reduced by 54%.

6.3. Bottom-up Multi-Resolution Pyramid

In order to implement our BMRP scheme we need to find an error threshold, that determines which detail coefficients to keep. In our experiments we found that keeping 20% of the coefficients with the highest error in each level is a good trade-off between storage overhead and accuracy. Although these error coefficients can be anywhere in the volume, they are still located on a uniform grid. Hence, using a run-length encoding data-structure, we found that for 20% of the points of a uniform dataset we need approximately 40% of the storage required for the entire uniform dataset.

The Bypass dataset consists of 421 timesteps of a simulation from a laminar-turbulent transition in a boundary layer that is subject to free stream turbulence. Each timestep is

Table 2: RMS errors and computation times (in minutes) for different non-uniform datasets for our approach, the method in [JWH*04] and in [JBL*06]. Size (MB) in the third and fifth column shows the storage (in megabyte) required for the non-uniform dataset and the reconstructed dataset, respectively. The resolutions in our approach are selected based on the $\sigma_{avg} = 0.05$ threshold.

Dataset					RMS and Times (min)		
Name	Points	Size(MB)	Resolution	Size(MB)	Our method	[JWH*04]	[JBL*06]
Oil	29,094	0.44	38x40x38	0.22	0.19 0.07	1.02 1.10	1.08 0.21
Natural Convection	68,921	1.05	61x61x61	0.87	0.63 0.07	1.51 6.95	1.41 4.16
Synthetic Chirp	75,000	1.14	64x64x64	1.00	1.12 0.08	3.06 229	1.37 36.4
Bypass	7,929,856	121.00	766x92x192	51.62	0.61 6.40	3.38 3987	3.33 3889
Blunt-Fin	40960	0.63	93x36x25	0.32	1.14 0.12	1.58 6.83	1.41 5.38




Figure 2: Renderings of the Tooth dataset: a) original uniform dataset, b) reconstruction from 2,110,259 non-uniform points using our method. The resolution of reconstruction is selected to be the same as in the original 256x256x160 dataset. The RMS error is 0.19 at the input points, and 0.18 at the entire uniform volume, c) reconstruction from the same set of input points using RBFs proposed in [JWH*04]. The RMS error is 1.26 at input points, and 2.87 for the entire volume, and d) reconstruction from the same set of input points using EBFs as proposed in [JBL*06]. The RMS error is 0.76 at input points, and 2.45 at the entire uniform volume.

represented by 7,929,856 non-uniform points in a curvilinear grid with uniform spacing across the x- and z-axes and with non-uniform spacing along the y-axis. The visualization of this simulation is of great importance to better analyze how the "bypass" of the Tollmien-Schlichting (TS) waves develops. In Fig. 4 we show timestep 360 from this dataset (focusing on the "bypass" process, i.e., the creation of vortex-shape structures) reconstructed with our BMRP approach. There is a visible difference in the level of detail in the different resolutions. The file size for saving $c^{(2)}$, $e_p^{(0)}$, and $e_p^{(1)}$ altogether is 49% of the size of the non-uniform dataset. Analyzing the plot results of the relation of the RMS error to the percentages of $e^{(0)}$ and $e^{(1)}$, we observe a drastic change in the errors in the 20% region. Hence, we decided to keep only 20% of the error coefficients in each level.

6.4. Adaptive Multi-Resolution

To visualize the multi-resolution hierarchy we have adapted our CPU-based raycaster to implement Procedure 2, which takes all resolution levels into account during rendering.

One of our main concerns is the continuity or smoothness preservation through different levels of resolution. However, since each level of the hierarchy is C^2 continuous and we are simply adding these levels, the final result remains a C^2 continuous function. In order to avoid any discontinuity at the

boundaries, we extend the borders of the cells in each direction by a specific number of voxels of value zero (here the voxel size depends on the resolution of the cell). We take into consideration the finite support of cubic B-splines. Extending by two voxels in each direction ensures that the function representing the cell smoothly goes to zero as it approaches these extended borders and is zero-valued everywhere beyond them.

Taking into consideration rendering performance a suitable choice of N_{max} could be 8, 16, or 32. In fact, in our experiments we chose 32 for the initial level, but experimented with different N_{max} for the subsequent levels. The decision whether a cell has to be refined is based on the error of reconstruction of that cell (see line 5 in Procedure 1). The error threshold is always set to 1.0. In order to prevent the subdivision of cells with only few non-uniform points we set $\bar{M} = 100$.

The X38 dataset consists of 323,192 non-uniform points emulating the X38 Crew Return Vehicle. It is a typical non-uniform dataset where 99% of its points are concentrated in about 5% of the volume. In Fig. 5 we show the dataset reconstructed with our multi-resolution scheme consisting of two levels. Due to the aspect ratio of the axis-aligned bounding box, the coarse resolution is 32x23x17 with an RMS of 6.39.




Figure 3: Renderings of the Chirp dataset: a) the original uniform data, b) reconstruction from 75,000 non-uniform points using regularization as defined in Eq. 6 ($\lambda = 0.3$), RMS is 1.12 with a reconstruction time of 0.08 min, and c) reconstruction from 75,000 non-uniform points using regularization as defined in Eq. 7 ($\lambda_x = \lambda_y = 0.3, \lambda_z = 1.0$), RMS is 0.51 with a reconstruction time of 0.08 min.




Figure 4: Renderings of the Bypass dataset consisting of 7,929,856 non-uniform points, reconstructed using our BMRP scheme (finest resolution 1024x120x256): a) coarse representation reconstructed from $c^{(2)}$ coefficients, RMS is 4.55, b) finer representation reconstructed from $\tilde{c}^{(1)} = U_j c^{(2)} + e_p^{(1)}$, where we used 20% of the points from the $e^{(1)}$ error volume, RMS is 2.69, c) finest representation reconstructed from $\tilde{c}^{(0)} = U_j \tilde{c}^{(1)} + e_p^{(0)}$, where we used 20% of the points from the error volume $e^{(0)}$, RMS is 0.6, and d) finest representation reconstructed where we used 100% of the points from the error volumes $e^{(0)}$ and $e^{(1)}$, RMS is 0.4.




Figure 5: The X38 aircraft dataset consisting of 323,192 non-uniform points rendered with our multi-resolution scheme using two levels of hierarchy.

Table 3 summarizes several scenarios we tested to analyze the behavior and performance of our method. In all cases we specified the coarse resolution to be 32, while for the composite cells we either selected 8, 16 or an adaptive resolution by using the $\sigma_{avg} = 0.05$ threshold. Using a resolution of 8 requires more levels in order to capture the data accurately, however, the storage per level is reduced. By using a resolution of 16 or an adaptive resolution we increase the storage requirement per level, but achieve a better approximation of the data. When using the adaptive resolution the impact of the third or higher levels in the reconstruction error is very

small. For the Bypass dataset there is no refinement in level three since no cell has an error higher than 1.0.

In addition we also analyzed the impact of the threshold \tilde{M} . Lowering this threshold, lowers the error, but increases the storage requirements drastically. Since the minimal refinement of the composite cells is $8^3 = 512$, having a threshold \tilde{M} set to 100 is sensible.

7. Conclusion and Future work

In this work we presented a framework for reconstruction and visualization from non-uniform point sets on uniform grids using B-spline basis functions. We improve our reconstruction results by introducing a new regularization functional. We introduce a new link between non-uniform representations and the interscale B-spline relation in a multi-resolution context. We show the performance and quality of our technique when compared to other competing techniques.

In the future we plan to improve on our work by: a) analyzing the possibility of an automatic setting of the λ parameters, b) extending our framework to vector and time-varying data, c) providing a GPU implementation of the CPU-based multi-resolution raycaster, and d) investigating alternative lattice structures.

Table 3: AMR results for different datasets and settings. Size shows the storage requirements in MB, Sub shows the resolution of reconstruction of the composite cells (i.e., either fixed or selected adaptively (adp) based on the $\sigma_{avg} = 0.05$ threshold), Thresh is the minimum number of points for which a cell can be subdivided, Cells is the number of composite cells in each level of the hierarchy and Times are in minutes.

Dataset				Level 1			Level 2				Level 3			
Name	Size	Sub	Thresh	RMS	Times	Size	Cells	RMS	Times	Size	Cells	RMS	Times	Size
X38	5.11	8 ³	100	6.39	0.03	0.05	136	2.66	0.07	0.27	1765	1.99	0.27	3.46
X38	5.11	16 ³	100	6.39	0.03	0.05	136	1.67	0.17	2.13	651	1.60	0.58	10.18
X38	5.11	adp	100	6.39	0.03	0.05	136	1.60	0.52	3.14	33	1.58	0.03	0.18
Bypass	121	8 ³	100	4.33	0.52	0.01	938	2.41	1.22	1.84	13270	1.54	1.95	25.99
Bypass	121	16 ³	100	4.33	0.52	0.01	938	0.79	1.72	14.66	-	-	-	-
Bypass	121	adp	100	4.33	0.52	0.01	938	0.34	7.53	58.91	-	-	-	-
X38	5.11	adp	50	6.39	0.03	0.05	183	1.59	0.55	3.32	292	1.55	0.25	1.37
X38	5.11	adp	10	6.39	0.03	0.05	427	1.33	0.58	3.58	7535	0.97	5.03	28.63
X38	5.11	adp	1	6.39	0.03	0.05	787	1.30	0.63	3.88	17948	0.75	12.43	81.60

Acknowledgments

We thank Muthuvel Arigovindan and Michael Unser for helpful discussions about the variational approach, Philippe Thévenaz for the inspiration of the chirp data set, and Yun Jang and Anders Ynnerman for the discussions and providing the datasets. This work is supported by the Austrian Science Fund (FWF) grant no. P18547, and is partially funded by the National Science and Engineering Research Council of Canada.

References

- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Point set surfaces. In *Proceedings of IEEE Visualization* (2001), pp. 21–28.
- [AG01] ALDROUBI A., GRÖCHENIG K.: Nonuniform sampling and reconstruction in shift-invariant spaces. *SIAM Rev.* 43, 4 (2001), 585–620.
- [AL96] AUBURY M., LUK W.: Binomial filters. *Journal of VLSI Signal Processing* 12 (1996), 35–50.
- [ASHU05] 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* (2005), vol. 14, pp. 450–460.
- [BG05] BRUCKNER S., GRÖLLER M. E.: VolumeShop: An interactive system for direct volume illustration. In *Proceedings of IEEE Visualization* (2005), pp. 671–678.
- [CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3D objects with radial basis functions. In *Proc. of SIGGRAPH* (2001), pp. 67–76.
- [Duc79] DUCHON J.: Splines minimizing rotation-invariant seminorms in Sobolev spaces. In *Multivariate Approx. Theory, W. Schempp and K. Zeller, Eds. Birkhäuser-Verlag* (1979), 85–100.
- [FGS95] FEICHTINGER H., GRÖCHENIG K., STROHMER T.: Efficient numerical methods in non-uniform sampling theory. *Numerische Mathematik* 69 (1995), 423–440.
- [GS04] GRISHIN D., STROHMER T.: Fast multi-dimensional scattered data approximation with Neumann boundary conditions. *Linear Algebra Applications* 391 (2004), 99–123.
- [JBL*06] 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 *Comp. Graph. Forum* (2006), vol. 25, pp. 587–596.
- [JWH*04] JANG Y., WEILER M., HOPF M., HUANG J., EBERT D. S., GAITHER K. P., ERTL T.: Interactively visualizing procedurally encoded scalar fields. In *Proc. of Joint EG-IEEE TCVG Symp. on Visualization* (2004), pp. 35–44.
- [KSH03] KÄHLER R., SIMON M., HEGE H.-C.: Interactive volume rendering of large sparse data sets using adaptive mesh refinement hierarchies. *IEEE Transactions on Visualization and Computer Graphics* 9, 3 (2003), 341–351.
- [KSW01] KRISHNAN S., SILVA C. T., WEI B.: A hardware-assisted visibility ordering algorithm with applications to volume rendering of unstructured grids. In *Proceedings of EG/IEEE TCVG Symposium on Visualisation* (2001), pp. 27–34.
- [LHJ99] LAMAR E., HAMANN B., JOY K. I.: Multiresolution techniques for interactive texture-based volume visualization. In *Proceedings of IEEE Visualization* (1999), pp. 355–361.
- [Nie93] NIELSON G. M.: Scattered data modeling. *IEEE Comput. Graph. Appl.* 13 (1993), 60–70.
- [OBS04] OHTAKE Y., BELYAEV A. G., SEIDEL H.-P.: 3D scattered data approximation with adaptive compactly supported radial basis functions. In *International Conference on Shape Modeling and Applications* (2004), pp. 31–39.
- [PLKO06] PARK S. W., LINSSEN L., KREYLOS O., OWENS J. D.: Discrete Sibson interpolation. *IEEE Transactions on Visualization and Computer Graphics* 12, 2 (2006), 243–253.
- [Sam05] SAMET H.: *Foundations of Multidimensional and Metric Data Structures*. Morgan Kaufmann Publ. Inc., USA, 2005.
- [TBU00] THÉVENAZ P., BLU T., UNSER M.: Image interpolation and resampling. In *Handbook of Medical Imaging, Processing and Analysis*. Academic Press, 2000, pp. 393–420.
- [Uns99] UNSER M.: Splines: A perfect fit for signal and image processing. *IEEE Signal Proc. Magazine* 16, 6 (1999), 22–38.
- [Uns00] UNSER M.: Sampling—50 Years after Shannon. *Proceedings of the IEEE* 88, 4 (2000), 569–587.
- [VMG08] VUČINI E., MÖLLER T., GRÖLLER M. E.: Efficient reconstruction from non-uniform point sets. *The Visual Computer, Springer Berlin / Heidelberg* 24, 7–9 (2008), 555–563.
- [WKME04] WEILER M., KRAUS M., MERZ M., ERTL T.: Hardware-based raycasting for tetrahedral meshes. In *Proceedings of IEEE Visualisation* (2004), pp. 71–78.
- [WM03] WELSH T., MUELLER K.: A frequency-sensitive point hierarchy for images and volumes. In *Proceedings of IEEE Visualization* (2003), pp. 425–432.

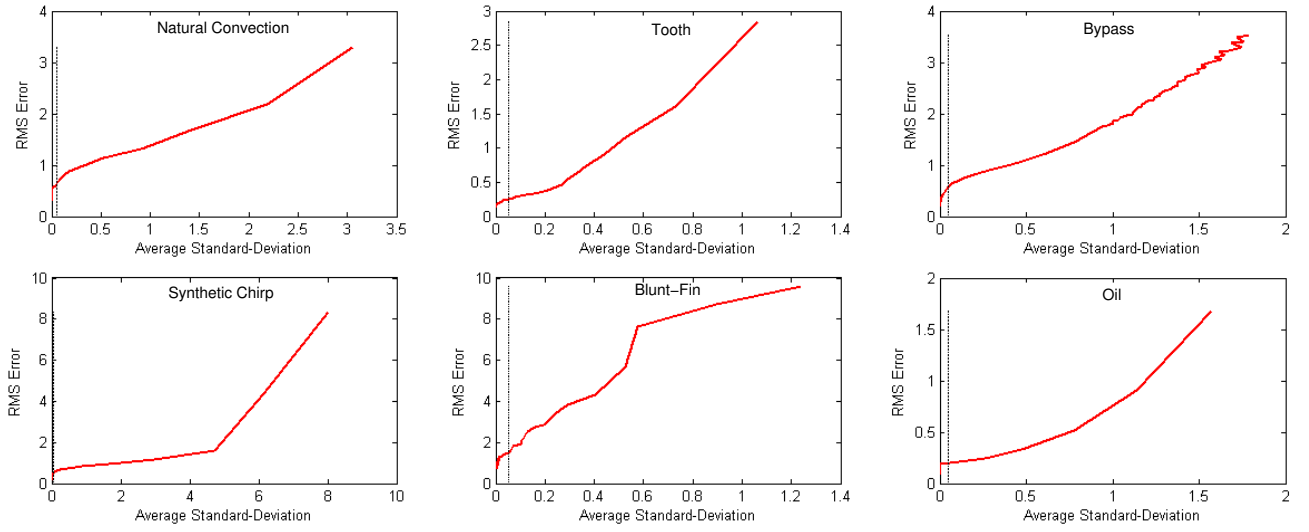


Figure 6: The relation of RMS as computed in Equation 8 vs. the average standard deviation σ_{avg} as computed in Equation 3 for several non-uniform datasets. A threshold of $\sigma_{avg} = 0.05$ is shown with a vertical line.

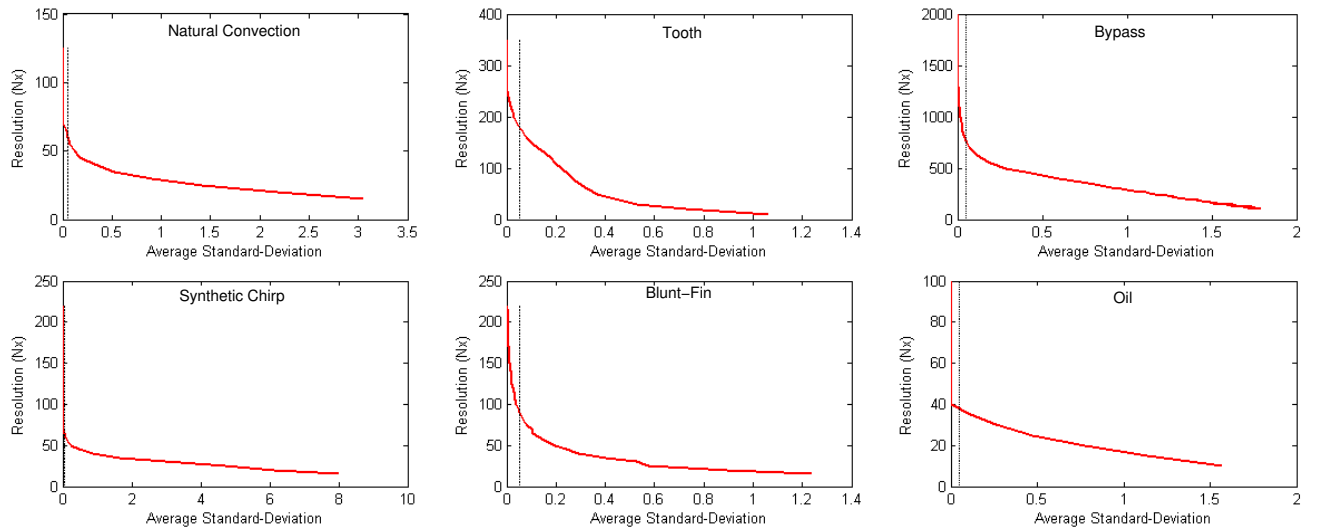


Figure 7: The relation of resolution of reconstruction vs. average standard deviation σ_{avg} as computed in Equation 3 for several non-uniform datasets. A threshold of $\sigma_{avg} = 0.05$ is shown with a vertical line.




Figure 8: Renderings of the Natural Convection dataset: a) reconstruction using our method, RMS is 0.63, b) reconstruction using RBFs proposed in [JWH*04], RMS is 1.51, and c) reconstruction using EBFs proposed in [JBL*06], RMS is 1.41.




Figure 9: Renderings of the Oil dataset: a) reconstruction using our method, RMS is 0.19, b) reconstruction using RBFs proposed in [JWH*04], RMS is 1.02, and c) reconstruction using EBFs proposed in [JBL*06], RMS is 1.08.




Figure 10: Renderings of the Synthetic Chirp dataset: a) reconstruction using our method, RMS is 1.12, b) reconstruction using RBFs proposed in [JWH*04], RMS is 3.06, and c) reconstruction using EBFs proposed in [JBL*06], RMS is 1.37.