EG UK Theory and Practice of Computer Graphics (2007)
Ik Soo Lim, David Duce (Editors)

A case-study of inconsistent surface reconstruction in recent literature resulting from Octree rotation-variance

David Meirion Hughes and Ik Soo Lim

School of Computer Science,
University of Wales, Bangor,
United Kingdom
{meirion, iksoolim}@informatics.bangor.ac.uk

Abstract

We review the use of octree and PCA (Principle Component Analysis) in current literature and explore a leading piece of research, as a case study, to highlight how overlooking octree rotation-variance has resulted in inconsistent results. We provide a simple method using PCA to re-orientate an octree to the intrinsic-orientation defined by data variance. In our case-study we explore and provide a method for consistency within multi-scale CSRBFs (Compactly Supported Radial Basis Functions). By utilizing PCA we provide rotation-invariant multi-scale surface reconstructions. We show, by curvature analysis, that the original surface reconstruction method is variant to data orientation and that our approach results in rotation-invariant reconstructions. In addition we also provide a technique for more flexibility when choosing a CSRBF for use with multi-scale surface reconstruction.

Categories and Subject Descriptors (according to ACM CCS): surface fitting, reconstruction, reliability, performance

1. Introduction

Consistent results are a fundamental aspect of computer science. Indeed without a high standard for consistency, research may not be entirely reproducible in a consistent and reliable manner. For many years octrees have been employed for a magnitude of tasks, ranging from space partitioning to data down-sampling. For such an important tool it is vital that its use is consistent and reliable. However, it is apparent that even recent research can be plagued by inconsistencies due to the use of an octree.

By design the octree is used to partition space into octants, which can further be subdivided. Typically, when applied to a space with some form of spatial data, the octree is centered and bound to that data; this ensures efficiency, and is widely practiced. While the sampling of an octree that has been centered and bound to the data is both position- and scale-invariant, it is not rotation-invariant. This is to say that if the data was scaled or translated, the octree results would be identical but not if the data was rotated.

In this paper we explore the important field of surface reconstruction and highlight recent research where the rotation-variance of octrees has not been addressed. We provide a method utilizing PCA (Principle Component Analysis) by which rotation-invariance can be achieved for an octree employed for surface reconstruction. In addition we show the inconsistency of the previous method and the consistency of our approach using a curvature analysis.

1.1. Previous Work involving Octree and PCA

Octrees are have been used extensively in surface reconstruction. We underline this fact by highlighting recent uses of octrees. While we only explore and prove one example of inconsistency resulting from the use of an octree, it is possible -but in no way asserted or proven here- that other methods employing an octree may be rotation-variant and subsequently inconsistent.

Park [PGSQ06] utilizes octrees in an automatic and interactive system to repair both the shape and appearance of defective point sets. Ohtake [OBMA03] employs octree-based subdivision of very large sets of points to reconstruct surface models using multi-level partition of unity implicit surfaces. Xie [XMQ04] organizes sample points using an octree for surface reconstruction, which is able to recover high-quality surfaces from noisy and defective data sets without any normal or orientation information. Dalmasso [DN04] describes a volumetric approach to surface reconstruction from nonuniform data. The data volume is split and classified at different scales of spatial resolution into surface, internal and external voxels and hierarchically described using an octree. Kindlmann [KRTM03] uses efficient octree algorithm to permit handling of large data sets for CSRBFs surface reconstruction. Hornung [HK06] uses a octree in a new volumetric method for reconstructing watertight triangle meshes from arbitrary, un-oriented point clouds. Tobor [TRS06] reconstructs multi-scale implicit surfaces with attributes, given discrete point sets with attributes. Ohtake [OBS05] uses an an octree to build a hierarchy of points-sets for use in multi-scale RBF surface reconstruction technique (we further explore this method as our case study). Kalaiah [KA03] pre-processes input data using an octree. Wilhelms [WG92] used an octree for faster iso-surface generation. Knoll [KWSC06] used a lossless-compression octree representation to store compressed volume data for fast iso-surfacing. Hadwiger [HSSB] describes a two-level hierarchical representation utilizing a form of octree to allow object-order and image-order empty space skipping for real-time ray-casting of discrete isosurfaces.

PCA has been widely used for the representation of shape, appearance and motion in the computer graphics field. It is worthwhile to highlight some research utilizing PCA. Kalaiah [KA03] uses PCA analysis on a group of points *inside* octants of an octree to estimate of their local orientation frame, the mean, and variance. Sloan [SHHS03] compresses the storage and accelerates performance of pre-computed radiance transfer (PRT), which captures the way an object shadows, scatters, and reflects light, using clustered principal component analysis. Kristensen [KAMJ05] uses PCA to form a compressed representation for real-time relighting of scenes illuminated by local light sources. Feng [FYD00] used PCA for Human face recognition, while Gouaillier [GGA97] explored PCA as a means for ship silhouette recognition. Torre [TB01] provided robust PCA (RPCA) for computer vision for improved PCA representation of shape, appearance, and motion.

2. Surface reconstruction

Surface reconstruction from unorganized point clouds is an important problem, especially for the recreation of real world objects that have been digitally scanned. Most object scanning technologies do not, by design, provide a surface model to be used instantly, but rather supply data by which a surface or an object can be recreated. There are a variety of sources from which data is obtained. Contour slices, where an object has been scanned using a CT scanner and an iso-surface has been defined, it a typical source. Another source is interactive tools, where data is created in real-time by a user. However, most prominent is range-data, where an object has been scanned using a laser to measure distances to the areas of an object.

There have been numerous solutions to the surface reconstruction problem. Hoppe [HTDM92] used an implicit surface model were surface reconstruction was defined as the zero set of an estimated signed distance function. Bernardini [BMRS99] used a the rolling ball technique, Curless [CL96] used a volumetric approach and Carr [CB01] used Radial Basis Functions to solve a scattered data interpolation problem and reconstruct surfaces. Nina Amenta [ABK98, AB99] used Voronoi vertices and Delaunay triangulation to create a piece-wise linear approximation of a smooth surface with better noise reduction.

Typically the problem of reconstructing a surface requires that the input data be converted to an unorganized point cloud in three-dimensional space. In this paper we are exploring an approach whereby the surface reconstruction task is cast as a scattered data interpolation problem and the reconstruction is defined as an implicit surface.

A surface that is not explicitly defined, but rather is embedded within another property, is called an implicit surface. A distance field is an example of an implicit surface. The surface of a distance field is typically defined as zero and all space exterior or interior to the surface is non-zero. Surface reconstructions based on implicit surfaces is a popular approach due to a number of advantages it has over other representations. A particularly noteworthy advantage is the ability of implicit surfaces to easily represent models of complex topology.

A well known approach for solving the scattered data interpolation problem is a RBF (Radial Basis Function) Network.

2.1. Radial Basis Functions

For the scattered data point interpolation, a RBF network is defined as [Pow87, Hay99]

$$f(\mathbf{x}) = \sum_{i=1}^{N} w_i \phi(\|\mathbf{x} - \mathbf{x}_i\|) \tag{1}$$

which satisfies the interpolation conditions $f(\mathbf{x}_i) = y_i$ where $\mathbf{x}_i \in \mathbb{R}^3$ are data points, and $y_i \in \mathbb{R}$ are function values. Unlike 'height' function interpolation, a surface embedded in three-dimensional space is often defined as a zero-level set $f(\mathbf{x}) = 0$. To avoid the trivial solution that f is zero everywhere, off-surface points are typically appended to the input data and are given non-zero values $y_i \neq 0$ whilst the on-surface points are defined with $y_i = 0$ [DTS02, TO02]. The coefficients w_i are determined by solving a linear system $\mathbf{Gw} = \mathbf{y}$ which is obtained by inserting the interpolation conditions into Eq. 1. If the matrix $\mathbf{G}$ is full, however, this approach is limited to a small data set; approximately a thousand points or so. Given a large data set, a naive approach is to use a small subset of it and discard the remaining data points [DTS02]. A better approach is to use CSRBFs (Compactly Supported RBFs) since their compact supports lead to a sparse linear system suitable for a large data set [BTD01]. However, it is sensitive to the density of scattered data and, therefore, a careful selection of the support size for CSRBFs is required in surface reconstruction [BTD01].

2.2. Multi-Layer Radial Basis Functions

To get around the CSRBFs problems, multi-level interpolation with a point hierarchy was proposed [OBS05]. Given a set of points $\mathcal{P} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ sampled from a smooth surface, the multi-scale hierarchy of point sets $\{\mathcal{P}^1, \mathcal{P}^2, \ldots, \mathcal{P}^M = \mathcal{P}\}$ is first constructed by spatial down sampling. Then the multi-level interpolation procedure proceeds in a coarse-to-fine way with decreasing support sizes. It recursively determines the set of interpolating functions $f^k(\mathbf{x}) = f^{k-1}(\mathbf{x}) + o^k(\mathbf{x})$ such that $f^k(\mathbf{x}) = 0$ interpolates $\mathcal{P}^k$ for $k = 1, 2, \ldots, M$ and $f^0(\mathbf{x}) = -1$. The offsetting function $o^k(\mathbf{x})$ has the form of

$$\sum_{\mathbf{p}_i \in \mathcal{P}^k} [g_i(\mathbf{x}) + w_i]\, \phi_\sigma(\|\mathbf{x} - \mathbf{p}_i\|) \tag{2}$$

where $g_i(\mathbf{x})$ are local polynomial approximations determined via least square fitting to $\mathcal{P}^k$ and $\phi_\sigma(\|\mathbf{x} - \mathbf{p}_i\|)$ are CSRBFs. The coefficients w_i are found by solving a linear system

$$\Phi w = -(\mathbf{f} + \mathbf{g}) \tag{3}$$

obtained by the interpolation conditions $f^k(\mathbf{p}_j) = 0$ for every point $\mathbf{p}_j \in \mathcal{P}^k$. The point hierarchy is created using octree-based subdivision. It starts with an axis-aligned box that encompasses the point set $\mathcal{P}$, and is followed by recursive subdivision of the space and points into eight octants or cells. $\mathcal{P}$ is clustered with respect to the cells by computing centroids of the points in each cell. Depending on the coordinate system used to represent the points, however, it can lead to inconsistent surface reconstruction and geometry. For example, surface curvatures are important for matching and registration tasks and can result in different values even with the same point set if represented in different coordinate systems. An actual research example, Hadwiger [HSSB], not only uses an octree for hierarchical representation of a volume, but also explicitly extracts curvature for visualisation, and it is possible the variance problem affects it too. The variance, attributed to the coordinate system, is due to the octree subdivision such that each side of the cells is parallel to an axis of the coordinate system: rotation is especially problematic whilst the octree is invariant to other coordinate transforms such as translation, scaling and flipping.

3. Consistent Surface Reconstruction using PCA

Ohtake's [OBS05] multi-layer approach does not address the rotation-variance inherent to octree down-sampling. Indeed we prove later that an arbitrarily aligned octree results in inconsistencies in the surface reconstruction. Such inconsistencies could result in difficulty during object mark-up and reconstruction comparison, where consistency is vital.

To solve the problem of rotation-variance we turn to PCA (Principle Component Analysis). PCA involves a linear transformation of a data-set, such that the first principle component is the data-set projection with the greatest variance, the second principle component is the second greatest variance and so on. The first three principle components can be viewed as the cartesian axis defining the intrinsic orientation of the data. The appeal of this method is that any coordinate transformations applied to the data set will also effect its intrinsic orientation.

We arrive at an orthogonal coordinate system (the intrinsic-orientation of the data) from calculating the eigenvectors of the covariance matrix $\mathbf{C} = \mathbf{D}\mathbf{D}^T$ where $\mathbf{D} = [\mathbf{x}_1 - \bar{\mathbf{x}}, \ldots, \mathbf{x}_N - \bar{\mathbf{x}}]$ and $\bar{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}_i$. We then orientate the octree to the data using this coordinate system.

This approach is different than that of Kalaiah [KA03] where PCA is used on a group of points structured in an octree and used to determine the *local* orientation frame of the group, as we are re-orientating the *entire* octree *prior* to space partitioning and centroid calculation. Kalaiah [KA03] does not address the rotation-variance of octrees.

4. Flexible Basis Functions

For a wider choice of basis functions, we can use an approximation scheme. In the aforementioned interpolation approach, the down-sampled points $\mathbf{p}_j \in \mathcal{P}^k$ serve as both the basis centres and the data points as in Eq. 2 and 3, and only a few types of functions can make the linear system in Eq. 3 solvable. In the proposed approximation approach, we use $\mathcal{P}^k$ only for the basis centres and $\mathcal{P} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ for the data points. In this approximation scheme, we may obtain a linear system equivalent to Eq. 3 from the conditions $f^k(\mathbf{x}_i) = 0$. However, it is over-determined since the number of the basis functions is less than that of the data points, *i.e.*, $|\mathcal{P}^k| < |\mathcal{P}|$. As the least square solution, instead, we can determine the weights w_i by solving the following linear system

$$\Phi^T \Phi \mathbf{w} = -\Phi^T (\mathbf{f} + \mathbf{g}). \tag{4}$$

In choosing the basis functions, this approximation scheme provides more flexibility than the interpolation does since there are more functions making Eq. 4 solvable than those for Eq. 3.

5. Experiments

In order to show the consistency of our method we analyzed the curvature of the reconstructed models. We used the curvature calculation method provided by Kindlmann [KRTM03] to calculate the mean curvature at each center.

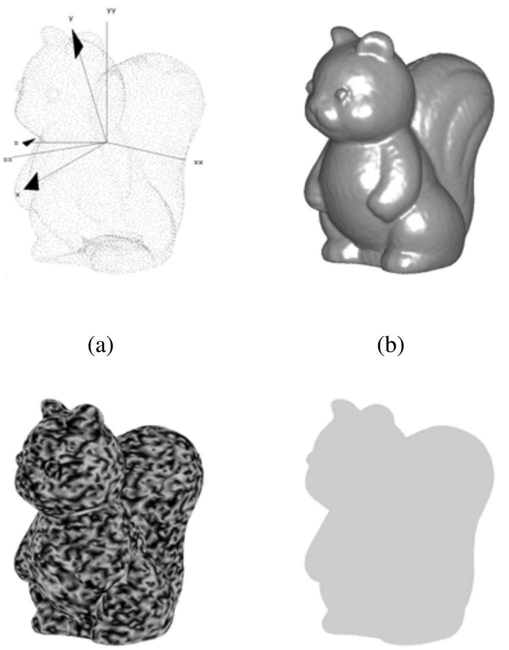

(a) (b) (c) (d)

Figure 1: *Surface reconstruction invariant to coordinate transforms. (a) the input data set represented by two different coordinate systems; (b) an example of surface reconstruction; (c) discrepancies in mean curvatures between the reconstructed surfaces with the octrees aligned to the input coordinate systems (the darker, the wider discrepancy); (d) discrepancies when aligned to the PCA coordinate system.*

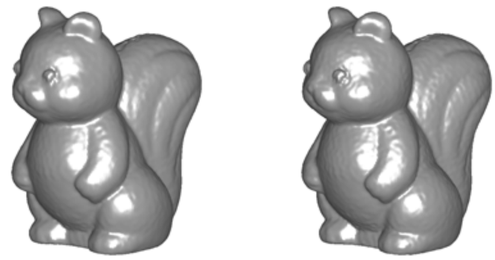

(a) (b)

Figure 2: *Interpolation vs. approximation. Surface reconstructions using (a) the interpolation and (b) the approximation scheme both using* $\phi(r) = \left(1 - \frac{r}{\sigma}\right)_+^4 \left(4\frac{r}{\sigma} + 1\right)$*; (c) the approximation using* $\phi(r) = \left(1 - \frac{r^2}{\sigma^2}\right)_+^4$

First we formed two copies a 54K point-set, sampled from squirrel a data-set, and applied a rotation of 70 degrees around the y-axis to one of them. We then applied the PCA-Octree reorientation step on both data sets. After fitting two multi-scale CSRBFs, we calculated the mean curvature for each data-set. Figure 1 shows the curvature residual of two experiments: one with the PCA-Octree orientation step and one without. When the octrees are aligned with the input coordinate systems as in [OBS05], they show noticeable discrepancies in these curvature values: 8.2 on average. When aligned with the PCA coordinate system as proposed, they show virtually no discrepancies: $1.0 * 10^{-4}$, of which attributed to numerical inaccuracy.;

We also experiment with the approximation scheme for the reconstruction and compare its results with that of the interpolation (Fig. 2). In addition to the aforementioned CSRBFs, we use simpler basis functions $\phi(r) = \left(1 - \frac{r^2}{\sigma^2}\right)_+^4$ which would be not usable in the interpolation scheme. These basis functions only need to compute r^2, but not r which would involve expensive computation of *sqrt*() as in the CSRBFs whilst the reconstruction qualities are comparable.

6. Conclusion

We have reviewed the use of octrees and PCA in current research. We have explored, using a surface reconstruction case-study, how leading research does not account for the rotation-variance inherent to octrees when partition a space and spatial data. Indeed we show that not accounting for the problem produces inconsistent results within derived techniques. As a case-study we explored the surface reconstruction approach utilizing an octree and CSRBFs. We highlight the need for researchers to consider the rotation-variance of octrees when applied to any form of spatial data.

We examined Ohtake's [OBS05] multi-scale CSRBFs surface reconstruction technique and have presented a method for rotation-invariance by utilizing PCA. PCA was used to re-orientate the scattered data points prior to fitting; in-effect orientating the octree to the intrinsic-orientation of the data-set defined by the data variance. We have shown that employing the PCA-Octree method produces consistent reconstructions and consistent analysis results, such as curvature analysis, of arbitrarily-orientated data.

In addition we have also introduced more flexibility to multi-scale CSRBFs by employing RBF approximation. We apply RBF approximation and provide an example of a computationally inexpensive (relatively) compact support radial basis function, which is shown to produce similar results as the standard method.

References

[AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discrete and Computational Geometry 22* (1999), 481–504.

[ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *SIGGRAPH* (1998), pp. 415–421.

[BMRS99] BERNARDINI F., MITTLEMAN J., RUSHMEIR H., SILVA C.: The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Vision and Computer Graphics 5*, 4 (1999).

[BTD01] B.MORSE, T.YOO, D.CHEN: Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions. In *Proceedings of the International Conference on Shape Modeling & Applications* (2001), pp. 89–98.

[CB01] CARR J. C., BEATSON R. K.: Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of conference on Computer graphics and interactive techniques* (2001), pp. 67–76.

[CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *In Proceedings of ACM SIGGRAPH (1996)* (1996), pp. 303 – 312.

[DN04] DALMASSO P., NERINO R.: Hierarchical 3d surface reconstruction based on radial basis functions. In *3D Data Processing, Visualization and Transmission* (September 2004), pp. 574–579.

[DTS02] DINH H. Q., TURK G., SLABAUGH G.: Reconstructing surfaces by volumetric regularization using radial basis functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24*, 10 (October 2002), 1358–1371.

[FYD00] FENG G. C., YUEN P. C., DAI D. Q.: Human face recognition using pca on wavelet subband. *Journal of Electronic Imaging 9*, 2 (2000), 362.

[GGA97] GOUAILLIER V., GAGNON L., ANDREW G. T.: Ship silhouette recognition using principal components analysis. In *Applications of digital image processing, Conference* (1997), vol. 3164, pp. 59–69.

[Hay99] HAYKIN. S.: *Neural networks:a comprehensive foundation*, 2nd. edition ed. Prentice Hall inc., 1999, ch. 5, p. 280Ű284.

[HK06] HORNUNG A., KOBBELT L.: Robust reconstruction of watertight 3d models from non-uniformly sampled point clouds without normal information. In *Eurographics Symposium on Geometry Processing* (2006)), pp. 41–50.

[HSSB] HADWIGER M., SIGG C., SCHARSACH H., BÜHLER K.: Real-time ray-casting and advanced shading of discrete isosurfaces.

[HTDM92] HOPPE H., T.DEROSE, DUCHAMP T., MCDONALD J.: Surface reconstruction from unorganized points. In *SIGGRAPH* (1992), pp. 71 – 78.

[KA03] KALAIAH A., A V.: Statistical point geometry. In *Eurographics Symposium on Geometry Processing* (2003), pp. 107 – 115.

[KAMJ05] KRISTENSEN A. W., AKENINE-MÖLLER T., JENSEN H. W.: Precomputed local radiance transfer for real-time lighting design. In *In Proceedings of ACM SIGGRAPH (2005)* (2005), pp. 1208 – 1215.

[KRTM03] KINDLMANN G., R.WHITAKER, TASDIZEN T., MOLLER T.: Curvature-based transfer functions for direct volume rendering: methods and applications. *IEEE Visualization 5*, 4 (2003), 513– 520.

[KWSC06] KNOLL A., WALD I., S.PARKER, C.HANSEN: Interactive isosurface ray tracing of large octree volumes. In *IEEE Symposium on Interactive Ray Tracing (2006)* (2006), pp. 115–124.

[OBMA03] OHTAKE Y., BELYAEV A., M. ALEXA G. T.: Multi-level partition of unity implicits. In *Proceedings of the ACM SIGGRAPH (2003)* (2003), pp. 463 – 470.

[OBS05] OHTAKE Y., BELYAEV A., SEIDEL H.: 3d scattered data interpolation and approximation with multilevel compactly supported RBFs. *Graphical Models 67* (2005), 150Ű–165.

[PGSQ06] PARK S., GUO X., SHIN H., QIN H.: Surface completion for shape and appearance. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22* (March 2006), 168 – 180.

[Pow87] POWELL M.: Radial basis functions for multivariable interpolation: a review. In *Clarendon Press Institute Of Mathematics And Its Applications* (1987), pp. 143 – 167.

[SHHS03] SLOAN P.-P., HALL J., HART J., SNYDER J.: Clustered principal components for precomputed radiance transfer. In *In Proceedings of ACM SIGGRAPH (2003)* (2003), pp. 382 – 391.

[TB01] TORRE F., BLACK M. J.: Robust principal component analysis for computer vision. In *Eighth International Conference on Computer Vision (ICCV'01)* (2001), vol. 1, pp. 59–69.

[TO02] TURK G., O'BRIEN J. F.: Modelling with implicit surfaces that interpolate. *ACM Trans. Graph. 21*, 4 (2002), 855–873.

[TRS06] TOBOR I., REUTER P., SCHLICK C.: Reconstructing multi-scale variational partition of unity implicit surfaces with attributes. *Graphical Models 68*, 1 (2006), 25 – 41.

[WG92] WILHELMS J., GELDER A. V.: Octrees for faster isosurface generation. *ACM Transactions on Graphics 11*, 3 (1992), 201 – 227.

[XMQ04] XIE H., MCDONNELL K. T., QIN H.: Surface reconstruction of noisy and defective data sets. In *IEEE Visualization* (2004), pp. 259 – 266.