Computers & Graphics 49 (2015) 58–68

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on Processing Large Geospatial Data

Locally refined spline surfaces for representation of terrain data

Vibeke Skytt*, Oliver Barrowclough, Tor Dokken

SINTEF ICT, Forskningsveien 1, PO Box 124 Blindern, 0314 Oslo, Norway

ARTICLE INFO

Article history:
Received 5 September 2014
Received in revised form
12 March 2015
Accepted 12 March 2015
Available online 25 March 2015

Keywords:
Terrain data
Point clouds
LR B-spline surface
Approximation methods

ABSTRACT

In this paper we describe the use of a novel representation, LR B-spline surfaces, and apply this representation in the treatment of geographical data. These data sets are typically very large and LR B-spline surfaces offer a compact representation of overall smooth data with local details. We briefly describe the properties of the LR B-spline representation, and discuss the details of two approximation methods adapted for LR B-splines: least squares approximation and multilevel B-spline approximation (MBA). The described techniques are demonstrated on several examples of terrain data in the form of point clouds.

1. Introduction

Unprocessed terrain data often takes the form of huge amounts of unstructured points. To make the data appropriate for use, the data must be processed and represented in a more structured format. Furthermore, the data size must be reduced with a minimum loss of the information carried by the point cloud. We will, in this section, look into different methods for processing sets of scattered data before we concentrate on the topic of locally refined splines and the use of this technique for terrain data. It is assumed that the input data is already tiled such that the input data is suitable for being represented by one raster or one surface.

The digital elevation model (DEM) is the most common format for processed terrain data. DEM uses a raster format for storage. The measured area is divided into uniformly spaced cells and the elevation is represented by one point in each cell. This point is often placed in the cell centre and existing height values are used to estimate the elevation in this point. This procedure is referred to as interpolation. DEM is in essence a regular format although the elevation value is not necessarily defined in all cells. A continuous surface is represented by a discrete set of points. To fetch the value of a point not lying at a grid intersection leads to another height value estimation. Thus, DEM is an approximate representation of the terrain where the accuracy depends on the interpolation method and the grid density. A globally smooth terrain is necessarily more accurately represented than a terrain with high variation in shape.

A number of different methods have been applied to estimate the elevation in grid cells. An overview can be found in [22]. Most GIS systems offer the possibility of the inverse weighted interpolation method originally proposed by Shepard, see [34]. The method is simple, but tends to create the artifact of flat spots at the data points. Other methods include Kriging [25], natural neighbour [2] and splines. Kriging is an advanced technique using a Geo-statistical approach and is mostly relevant when the estimation errors are of interest. Natural neighbour interpolation is originally based on Voronoi tessellation and tends to give a more smooth result than the inverse weighted interpolation. The term ‘spline’ in the context of raster interpolation refers to splines with tension and regularized splines, see [22], which is different from splines as piecewise polynomials. However, the variational approach utilized in this context is related to one of the approximation methods we will describe in Section 3. Raster interpolation using splines can produce estimated values that are above or below the given sample data. This may be a wanted feature, but is less adequate if the sample points are close together and have extreme differences in value.

The method of radial basis functions (RBF) is central in scattered data interpolation and there is a choice of several types of basis functions. Franke published a test on some types in [15]. The quality of the interpolated surface varies considerably with the choice of basis functions. Initially RBF was a global method which resulted in an equation system of the same magnitude as the number of data points, but also local methods exist, see for instance [11]. A combination of moving least squares (MLS), see [20] and RBF is presented in [26].

The main field for wavelets [4] and wavelet transforms is signal processing, but it is also used for data compression and noise reduction. The latter is achieved by keeping only those terms in the wavelet decomposition which have a coefficient larger than a

* Corresponding author.
E-mail addresses: Vibeke.Skytt@sintef.no (V. Skytt), oliver.barrowclough@sintef.no (O. Barrowclough), tor.dokken@sintef.no (T. Dokken).

http://dx.doi.org/10.1016/j.cag.2015.03.006
0097-8493/

given threshold. Unser considers splines and spline wavelets for signal and image processing in [35].

Subdivision surfaces are flexible tools for shape modelling [27], but to little extent used for approximation of scattered data. However, a few approaches can be found and Ref. [32] focuses on the approximation of large data sets primarily for visualization.

Tensor-product spline surfaces (piecewise polynomials and NURBS) are a well established representation in fields such as computer aided design (CAD), where they are used primarily to model smooth shapes. Their use has also previously been explored in representing topographic data sets. For example, spline approximations of data from the Shuttle Radar Topography Mission (SRTM) have been used, particularly for the application of filling data voids [9]. A spline surface representing a terrain can be used to effectively generate rasters of different density on demand.

Both raster representations and tensor-product splines approximate the original measured data and reproduces the terrain best if it is globally smooth. Triangulated surfaces, on the other hand, are able to interpolate unorganized data. For huge data sets, a data reduction is nevertheless required. Thinning of the initial point cloud and hierarchical triangulation [12] are relevant methods in this context. A triangulated surface is appropriate for areas of high variation and sharp features, but gives a less efficient mean of storage for smooth areas. To get an indication on how a linear representation such as a triangulation approximates a smooth function compared to a higher order continuous representation, we can look at the corresponding figures for splines. When approximating a C^k-continuous function with a C^d-continuous spline of degree d where $d<k$ the error term is $O(h^{d+1})$. h is the size of the polynomial segments. So splitting all the polynomial segments of the spline in two will typically reduce the error to $1/4$ for the linear spline, while the error for the cubic spline will be reduced to $1/16$. Consequently, for the spline approximation of smooth shapes increasing the polynomial degree has the potential of representing the shape with less data.

For data which exhibits local detail, tensor-product spline surfaces do not possess sufficient flexibility. Non-uniform refinement is available, but it is global for each parameter direction resulting in large amounts of data. This means that such representations are limited to globally smooth data sets. Locally refined spline surfaces have, in contrast to tensor-product spline surfaces, the ability to represent local variations in shape without globally increasing the data size of the surface. In the past decades, a number of approaches to local refinement of splines have been pursued, including hierarchical splines [13], truncated hierarchical splines [17], T-splines [33] and LR B-splines [7]. In this paper, we focus our attention on LR B-splines. We will describe the properties of LR B-splines and discuss two methods for approximation of point clouds: least squares approximation with smoothing and multilevel B-spline approximation (LR-MBA).

Section 2 presents the spline representation format with emphasis on the LR B-spline surfaces, Section 3 focuses on approximation methods, and the two methods are illustrated with a number of examples in Section 4. Section 5 presents a conclusion and some considerations about further work.

2. Spline representations

2.1. The polynomial spline representation

The spline format is a well established representation and described by several authors, see for instance [8]. In the following, we will present a short summary to show the context where LR B-spline surfaces belong to. A spline curve is a piecewise polynomial curve where the polynomial pieces are joined together at specified values known as knots. The curve $c(t)$ is expressed as $c(t)=\sum_{i=1}^{n} c_i N_{i,\theta}(t)^d$ where c_i, $i=1\ldots n$ are the spline coefficients, n is the number of coefficients while N_i, $i=1\ldots n$ are basis functions or B-splines. The B-splines are piecewise polynomial functions with joints at the knots. The curve has polynomial degree d. It is parameterized on an interval $t\in[\theta_1,\theta_{n+d+1}]$. θ is the knot vector, the vector of parameter values of the joints between the polynomial pieces. For all entries in the knot vector $\theta_i \le \theta_{i+1}$. The curve will initially be C^{d-1} continuous, but the continuity may be reduced by allowing for multiple knots, i.e. $\theta_i=\theta_{i+1}$.

The B-splines, N_i, have limited support, that is N_i is non-zero only on the interval $[\theta_i,\theta_{i+d+1}]$ The B-splines are non-negative, linearly independent and add up to 1 for all t. This is denoted partition of unity. The curve can be refined by adding new knots and the polynomial degree of the curve can be increased without changing the curve.

The B-splines are a stable basis for the spline space. For a spline function, there exist constants which bound the function from above and below by the spline coefficients. This result is independent of the dimension of the spline space and the knot vector. In [30], Quak establishes a relation between this result and the Riesz basis in multi-resolution analysis in the context of spline wavelets.

A tensor-product spline surface is constructed by taking the tensor-product between two spline curves, i.e. $S(u,v)=\sum_{i=1}^{n}\sum_{j=1}^{m} c_{i,j}N_{i,\theta}(u)^{d_1}N_{j,\gamma}(v)^{d_2}$ where d_1 and d_2 are the polynomial degree in the first and the second parameter direction, respectively, and θ and γ are the two knot vectors. The surface is defined on the area $\Omega=[\theta_1,\theta_{n+d_1+1}]\times[\gamma_1,\gamma_{m+d_2+1}]$. This construction gives a surface possessing the same properties as the curve. The surface is defined on a regular grid and knot insertion will lead to a global refinement of the surface in one parameter direction. Local refinement is not possible.

The polynomial spline curves and surfaces can be extended with a rational term to be able to represent algebraic curves and surface such as circles, spheres and cylinders together with free form curves and surfaces in a uniform way, see [28]. The curve expression now becomes $c(t)=\sum_{i=1}^{n} h_i c_i N_{i,\theta}(t)^d / \sum_{i=1}^{n} h_i N_{i,\theta}(t)^d$ where $h_i\in\mathbb{R}$ are positive weights. The weights bring additional flexibility to the construction, but are seldom used in approximation as that would lead to a non-linear optimization problem. In our context, we will restrict ourselves to polynomial splines.

2.2. Locally refined splines with emphasis on LR B-spline surfaces

LR B-splines is a new approach for local refinement of spline spaces published in 2013, see [7]. Other approaches addressing this topic are hierarchical splines and T-splines. LR B-splines and (truncated) hierarchical B-splines build a sequence of nested spline spaces starting from a tensor product B-spline space. While the spline space of hierarchical B-splines is spanned by uniform B-splines at different levels of refinement, LR B-splines are based on non-uniform B-splines and allow all refinements where a B-spline is split. The splines space of hierarchical B-splines satisfying the strong condition (the knot lines defines a T-mesh) is included in the spline space spanned by LR B-splines over the T-mesh. Truncated hierarchical B-splines introduce in the transition zone between refinement levels special basis functions built from B-splines of the finer levels to reduce the size of the support of the basis functions. For LR B-splines successive refinement is performed when necessary to ensure that all B-splines have minimal support. To ensure positive partition of unity LR B-splines scale the B-splines used if needed, while the construction of the special basis functions of truncated B-splines is ensured by proper scaling of the B-splines from which they are composed. While (truncated) hierarchical B-splines are always linearly independent, linear independence for LR B-splines is ensured by supervising the