

Data-Dependent MLS for Faithful Surface Approximation

Yaron Lipman Daniel Cohen-Or David Levin

Tel-Aviv University

Abstract

In this paper we present a high-fidelity surface approximation technique that aims at a faithful reconstruction of piecewise-smooth surfaces from a scattered point set. The presented method builds on the Moving Least-Squares (MLS) projection methodology, but introduces a fundamental modification: While the classical MLS uses a fixed approximation space, i.e., polynomials of a certain degree, the new method is data-dependent. For each projected point, it finds a proper local approximation space of piecewise polynomials (splines). The locally constructed spline encapsulates the local singularities which may exist in the data. The optional singularity for this local approximation space is modeled via a Singularity Indicator Field (SIF) which is computed over the input data points. We demonstrate the effectiveness of the method by reconstructing surfaces from real scanned 3D data, while being faithful to their most delicate features.

1. Introduction

Approximating a surface from scattered data is a fundamental problem with various applications in computer graphics. In particular, reconstructing or upsampling an unorganized point set is an important problem in surface reconstruction [HDD*92, CBC*01, OBA*03]. The main challenge is to *faithfully* reconstruct the unknown surface from a scattered set of samples. A faithful reconstruction adheres to the data, and aims to respect even its most delicate singularities or sharp features. As such, the reconstruction should necessarily be piecewise-smooth.

Common techniques for surface reconstruction assume that the unknown surface is smooth everywhere, and hence make use of approximation spaces of a smooth functions to reconstruct the surface. A powerful method to approximate an unknown surface from an unorganized noisy sample set, is to fit local polynomials by the moving least-squares (MLS) method [ABCO*01]. The functional MLS approximation procedure [LS81, Lev98, Wen01] defines a smooth approximant/interpolant for the unknown function, based upon function samples at unstructured data sites. Combining the MLS functional approximation scheme with a projection operator [Lev03, AK04] then yields an efficient method for smoothly approximating surfaces from 3D data sets. This furnishes a highly generic and versatile tool for defining, manipulating and reconstructing surfaces based on irregu-

lar point samples such as those obtained from 3D scanners [PKKG03].

However, the MLS method is designed to reconstruct smooth surfaces from data which is sampled from smooth surfaces. This leads to an erroneous surface reconstruction at regions where the smoothness assumption is invalid. One noticeable phenomenon of forcing smoothness is known as the Gibbs phenomenon (see Figure 4 and 6(c)). In other cases, the MLS might smooth out features. In both cases the reconstructed surface is clearly unfaithful to the original geometry (see Figure 1). A main flaw of the MLS reconstruction stems from the fact that its approximation space (i.e., polynomials) does not respect singularities in the data.

In this paper, we present a data-dependent moving least squares (DDMLS) method for surface reconstruction. It is data-dependent in the sense that it is faithful to the data singularities. Hereafter, the terms sharp features or singularities refer to discontinuities of arbitrary magnitudes in the function or its derivatives.

The technique we present here retains the simplicity of the MLS machinery: For each point x , it fits a function $p \in \mathcal{F}$, from a class of functions \mathcal{F} . Then, $p(x)$ is the desired local approximation. However, here we use a varying, locally fitted, piecewise polynomial (splines) space as \mathcal{F} . The main challenge is to define the location (locus) of the spline singularity, so that the spline approximation space




Figure 1: Top row: original surface; Middle: MLS reconstruction; Bottom: DDMLS reconstruction. Note the faithful reconstruction by the DDMLS at singularities as expressed by the isophote lines (left column).

furnishes a local faithful approximant. Our approach to this problem is based on the following rationale. The problem of reconstructing a surface from discrete samples is an ill-posed problem since there is no unique surface that fits the data. In particular, it is impossible to distinguish between sharp and non-sharp data points (see Figure 3). Thus, for each data point we define a *continuous* value that measures its singularity potential. This defines a singularity indicator field (SIF), which is used to construct the spline space \mathcal{F} , from which the local approximation p is derived.

An important feature of the new method is that it automatically approximates the singularities in a *threshold-free* manner. This enables the method to be faithful even to the most delicate singularities, even in vicinity of abrupt singularities. This is demonstrated in Figure 4, where the height function model (b) contains two types of circular features. Across the outer circle there is an abrupt discontinuity, and across the inner circle there is a delicate derivative discontinuity. Another example is shown in Figure 1.

We show that the DDMLS method allows the reconstruction of surfaces with high fidelity. All the examples that we use exhibit some sharp features and fine details like the little scratch on the soap in Figure 8. We also show that this technique is applicable to the enhancing of images.




Figure 2: A closeup of projection of the black point. The point cloud is colored from blue to red proportional to the SIF at each data point. The green surface patch is the best local approximant p from the data-dependent space \mathcal{F} . The yellow ball is the result of the projection. Note that although the point cloud is sparse the local singularity is well estimated.

1.1. Background and related works

In recent years many researchers have dealt with defining and manipulating point set surfaces following the MLS methodology in [Lev03]. Alexa et al. [ABCO*03] defined a surface approximating a given 3D point set, and introduced algorithms for representation and up-sampling. Amenta et al. [AK04] have further deepened the understanding of the MLS projection operator by expressing it as an extremal surface and presenting an efficient algorithm to project a point on the MLS surface and its variants. Pauly et al. [PKKG03] have constructed a shape modeling framework based on point set geometry.

Fleishman et al. [FCOS05] use the MLS methodology for reconstructing surfaces with sharp features. They use a trial-and-error process to define the neighborhoods of points in the approximation. That is, for each point to be approximated, an initial group of nearby points is chosen by an exhaustive search. Then, other nearby points are added to the initial group as long as a max-norm threshold is not violated. This method, although robust to outliers, does not necessar-



Figure 3: Surface reconstruction from a discrete set of points is an ill-posed problem. The figure demonstrates three possible reconstruction procedures. One which smooths the data, one interpolating the data, and one which allows sharp features.




Figure 4: The height function in (b) contains two types of circular features. Across the outer circle there is a function discontinuity, and across the inner circle there is a delicate derivative discontinuity. The data set consists of the irregular point samples (a). (d) and (f) demonstrate reconstruction by MLS and DDMLS, respectively. To better visualize the Gibbs phenomenon and the over-smoothing, the corresponding isophote lines are presented at the bottom row. Note the oscillations near the boundaries and the smoothing effect of the delicate inner circular feature caused by the MLS, compared with the faithfulness of the DDMLS. To further illustrate the effect, 1D slices of the original, MLS, and DDMLS are displayed in (h) from top to bottom.

ily reconstruct a continuous surface near the sharp edges (see Figure 13).

Reuter et al. [RJT*05] have introduced an alternative projection operator to the MLS that respects sharp features. However, in their work the user has to manually tag the sharp features in the point cloud as an input to the algorithm.

Lipman et al. [LCOL06] looked for radial neighborhoods which minimize a bound on the error of the approximation. In their work, the authors dealt with smooth surfaces only and did not relate to the problem of surfaces with singularities and sharp features. However, in that paper the authors provided some error estimates for the MLS which will be of use in this paper.

The technique presented here is most related to the work on essentially non-oscillatory (ENO) interpolation [HEOC87]. ENO schemes originate and are in use in the field of numerical methods for PDEs [SO89, JS96, JP00]. In rough terms, for each point x , the ENO reconstruction scheme chooses a different stencil of samples for the reconstruction. In particular, ENO uses samples from the neighborhood of x which minimize the divided difference. This yields an approximant which interpolates the data samples and does not cross singularities, and hence does not suffer from the Gibbs phenomenon and does not over smooth the data.

These properties are most relevant to the reconstruction of surfaces from irregular point samples. However, there are two fundamental problems: First, ENO is an interpola-

tory scheme and thus cannot deal with noisy data. Second, the ENO concept is defined for 1D data and does not seem to have a straightforward generalization to irregular data in higher dimensions. This is mainly due to the difficulty in recognizing well-posed sets for multivariate interpolation, especially in the case of irregular data points. In particular, there is no natural generalization of divided difference to the multivariate case.

2. Overview

The core element of the present work is a new multivariate substitute for the divided difference check of the ENO scheme. We compute a Singularity Indicator Field (SIF), which intuitively, assigns to each point an estimate of its proximity to a singularity. This serves as a basis for features detection and for the definition of the appropriate local approximating spline space. Singularities of surfaces are typically 1-manifold, therefore, in order to reconstruct and represent them, it is sensible to employ manifold approximation techniques. In this work we use the MLS machinery again to approximate these 1-manifolds. The non-trivial part is defining an indicator field over the data, on which the MLS is applied. The benefit would be a consistent approximation to the singularity manifold.

The SIF is illustrated in Figures 5 and 6(b) with pseudo colors. In Section 4 we rigorously develop and define the SIF and describe its construction. In Section 5 we describe the local spline fitting, namely, using a piecewise polynomial




Figure 5: Visualization of the SIF of the surface from Figure 4 by a logarithmic scale where colors range from dark blue to dark red. (b) shows a close-up view of the upper-right corner. Note that the closer a point is to the singularity area, the larger is its SIF relative to other points in its vicinity.

space that respects the (possible) local 1-manifold singularity. The surface is then reconstructed by projecting points over these local (moving) splines, as illustrated in Figure 2. Next we outline a pseudo-code of the method:

1. Preprocess. For each data point:
 - a. Find a local reference plane.
 - b. Calculate the SIF value for that point.
2. Projecting a point s .
 - a. Find a local reference plane π . Denote by x the projection of s onto it.
 - b. Project the local data points on π to create a functional settings.
 - c. Apply 2D univariate MLS to find the curve defined by the SIF.
 - d. Construct the local spline space S_m such that its singularity line passes through and tangent to the univariate MLS projection.
 - e. Find the best least-squares fit $p \in S_m$ to the functional data (S_m includes also the smooth polynomials).
 - f. $p(x)$ is the desired DDMLS projection.

In the following section we review the basic MLS mechanism, which is employed here triply: first to define the SIF, then to define the local 1-manifold singularity curve, and finally to reconstruct the 2-manifold in a data-dependent manner.

3. Moving Least Squares

The key element in the MLS procedure for surfaces is the MLS approximation of functions, which we describe below. We are given data, sampled from a function f , at some irregular set of data sites: $(X, f(X)) \subset \Omega \times \mathbf{R}$, where Ω is a domain in \mathbf{R}^d , and $X = \{x_i : i \in I\}$, $I = \{1, 2, \dots, N\}$. To define the approximation at an arbitrary point $x \in \Omega$ the following quadratic minimization problem is solved:

$$p := \operatorname{argmin}_{p \in \mathcal{F}} \left\{ \sum_{i \in I} |p(x_i) - f(x_i)|^2 \phi(\|x_i - x\|) \right\}, \quad (1)$$




Figure 6: Surface reconstruction from a sparse point cloud (a). Note for example the low density of points in the region marked by a black circle. (b) presents the SIF values colored from blue to red. In (c) we exhibit that the MLS reconstruction smooths out the corners and suffers from the Gibbs effect. In (d) we see that a DDMLS reconstruction is more faithful to the sharp features of the input data.

where $\mathcal{F} = \Pi_m$, the subspace of polynomials of total degree m , $\phi(r)$ is a fast decreasing smooth function of finite support size h , and h is a data parameter. In this paper we have used the approximation $\phi(r) = e^{-\frac{r^2}{(h/4)^2}}$. Then the MLS approximation at x is defined by

$$\mathcal{M}_f(x) = p(x).$$

An equivalent definition of the same approximation is given in terms of a polynomial reproduction property [Lev98]:

$$\mathcal{M}_f(x) = \sum_{i \in I} f(x_i) L_i^{X,m}(x),$$

where the “shape functions” $\{L_i^{X,m}\}_{i \in I}$ minimize the quadratic form

$$Q = \sum_{i \in I} |L_i^{X,m}(x)|^2 \phi(\|x_i - x\|) \quad (2)$$

subject to the linear constraints of polynomial reproduction

$$\sum_{i \in I} L_i^{X,m}(x) b_j(x_i) = b_j(x), \quad j = 1, \dots, K \quad (3)$$

where $b_1(x), \dots, b_K(x)$ is a basis for Π_m . We shall make use of both presentations of the MLS approximation in the development of the new tool presented here.

To apply the above functional MLS approximation to surface-sampled data $\{p_i\}$, a local reference plane is defined for each projected point [Lev03]. In order to define a surface, this plane should be chosen with care. In this work we define the reference plane using a local weighted PCA, where




Figure 7: Scanned CAD model (top left, photograph). On the top right is the DDMLS reconstruction. The two bottom rows show enlarged parts, where the left most are photographs of the scanned model (ground truth), the middle column is the DDMLS reconstruction and the right-most column shows an MLS reconstruction. Note the delicate feature marked by a small arrow, which is smoothed in the MLS reconstruction.

the weights are defined to be $w_i = \phi(\|p_i - s\|)$, where here s is the projected point. Although this fails to define a projection operator in all cases [AK04], we found it practically convenient and robust.

4. Singularity indicator field

In this section we define the singularity indicator field (SIF), and describe its construction. Singularities are related to unbounded high derivatives of the unknown function f , and we would like to assign to each data point a value which indicates its proximity to a singular point. Given a discrete data set sampled from a function f , no *upper bound* to the derivatives of f can be found, since, as mentioned above, the reconstruction problem has no unique solution (see Figure 3). As we show below, it is possible to find a good local lower bound to the absolute value of the derivatives of the unknown function f . These lower bounds will be used to locate the points or curves of singularities of the surface.

To construct a lower bound of the derivative, we use the error expression of the MLS approximation [LCOL06]. Denote by p the polynomial that is fitted to the data X , as defined in (1), then the error in the MLS approximation of x is

$$f(x) - p(x) = - \sum_{i,v} \frac{D^v f(\eta_i(x_i - x) + x)}{v!} (x_i - x)^v L_i^{X,m}(x), \quad (4)$$

where $\sum_{i,v}$ stands for $\sum_{|v|=m+1} \sum_{i \in I_h(x)}$, $I_h(x) = \{i \mid x_i \in$

$X \cap B(x, h)\}$, where $B(x, h)$ is a ball of radius h centered at x , and $0 \leq \eta_i \leq 1$. From this equation we get

$$|f(x) - p(x)| \leq \max_{|v|=m+1, y \in \langle X_h \cup x \rangle} |D^v f(y)| \sum_{i,v} \frac{|x_i - x|^v}{v!} |L_i^{X,m}(x)|, \quad (5)$$

where $\langle X_h \cup x \rangle$ denotes the convex hull of the set of points $X_h \cup x$, where $X_h = \{x_i\}_{i \in I_h(x)}$.

After rearranging, we get

$$\max_{|v|=m+1, y \in \langle X_h \cup x \rangle} |D^v f(y)| \geq \frac{|f(x) - p(x)|}{\sum_{i,v} \frac{|x_i - x|^v}{v!} |L_i^{X,m}(x)|}. \quad (6)$$

Taking x to be one of the data points, all the expressions on the r.h.s. of the above inequality can be computed, and we get a lower bound for the absolute values of the $(m+1)$ th order derivatives of the unknown function f near a data point. Motivated by this inequality we define the singularity indicator Λ_j at point $x_j \in X$ by

$$\Lambda_j = \frac{|f(x_j) - p(x_j)|}{\sum_{i,v} \frac{|x_i - x_j|^v}{v!} |L_i^{X,m}(x_j)|}, \quad (7)$$

and we have an immediate property of Λ_j , derived from (6):

$$\Lambda_j \leq \max_{|v|=m+1, y \in \langle X_h \cup x \rangle} |D^v f(y)|.$$

The values Λ_j can be computed: The $f(x_j)$ are known, and the shape functions of the MLS $|L_i^{X,m}(x)|$ are obtained via the minimization of the quadratic form (2) subject to the linear constraints (3). The minimizer is computed by solving a linear system derived by Lagrange multipliers [Lev98]. Note, that the cost of computing each Λ_j is the same as a single MLS fit.

Let us assume that the function f is piecewise C^{m+1} , with a jump-discontinuity in one of its $(k+1)$ th order derivatives, $k \leq m+1$, near the point x_j . Then, using a local Taylor series approximation of total degree k , and using the polynomial reproduction property of the MLS, it can be shown that $\Lambda_j = O(h^{k-m-1})$. For a non-degenerate distribution of data points, we observe that $\Lambda_j = \Theta(h^{k-m-1})$ for $k = -1, 0$ (the examples in the paper). This observation implies that Λ_j can be used as a local indicator to the presence of jump-discontinuities, or large *magnitude* of the derivatives of the unknown function f at the vicinity of x_j . Therefore, we define the Singularity Indicator Field (SIF) of the given data as $\{x_i, \Lambda_i\}_{i \in I}$. Figure 5 shows an example of the SIF.

The above SIF should be computed for each projected data point. However, to reduce computational cost, we approximate the SIF by precomputing the singularity indicator of each data point once with respect to its own reference plane. Since changing slightly the reference plane does not change the relative magnitudes of the local derivatives this approximation is good enough for our purposes. Figure 6(b) shows the precomputed SIF for a specific example.




Figure 8: A scanned soap. (a) shows the ground truth (photograph). In (b) pct is taken as the 100th percentile, i.e., regular MLS. In (c) and (d) pct is taken as the 90th and the 75th percentile, respectively. Note how the fine scratch on the soap (marked by small arrow) is reconstructed by the DDMLS and smoothed out by the MLS. The small upper window at (a) shows a portion of the scanned point cloud.

5. Local data-dependent spline space

Instead of using the polynomial space Π_m in the approximation for x as is done in the standard MLS (1), we define the space from which the minimizer is sought in a data-dependent manner. That is, we define the modified MLS operator, first, by solving the minimization problem (1), where \mathcal{F} is now taken as the space of piecewise polynomials fitted to the point x . I.e., we let $\mathcal{F} = S_m$, the space of splines of total degree $\leq m$, whose singularity locus is an approximation to the singularity in the data, near the point x , as derived using the SIF. Thus we solve for

$$p := \operatorname{argmin}_{p \in S_m} \left\{ \sum_{i \in I} |p(x_i) - f(x_i)|^2 \phi(\|x_i - x\|) \right\}, \quad (8)$$

Then, the modified MLS operator $\widehat{\mathcal{M}}_f$ at x is defined by

$$\widehat{\mathcal{M}}_f(x) = p(x). \quad (9)$$

For the definition of the local approximation spline space \mathcal{F} to be used at a given point x , we apply the MLS framework on the SIF data, and construct a local univariate (curve) approximation to the singularities. I.e., we look for a curve which is close, in a least-squares sense, to data points with high SIF value. We explain this part of the procedure through an example: In Figure 9 (a) the SIF at each data point is drawn by a red circle whose radius is proportional to the SIF value at that corresponding data point. The green point denotes the point of approximation x . We then use the MLS operator to find, now in a univariate setting, a univariate polynomial which fits high valued SIF points locally, e.g., the green curve at Figure 9(a). More precisely, we first choose a *reference line* $\ell(x)$ (the thin black line in (a)). This can be done in several ways, as mentioned before. Here, we have used again a weighted PCA, where the weights are taken as the singularity indicator values times the smooth radial weights, i.e., $\phi(\|x_i - x\|)\Lambda_i$. Denote the projection of x on $\ell(x)$ by x_p , and by $v(x), n(x)$ the tangent unit vector and the normal to $\ell(x)$, respectively. Next, we project the data points $x_i \in X_h$ on the reference line to represent the data in the local

frame $v(x), n(x)$, where x_p serves as the origin of axes:

$$(t_i, y_i) = (\langle x_i - x_p, v(x) \rangle, \langle x_i - x_p, n(x) \rangle).$$

Also we denote x in that coordinate system by $(0, y)$. The local approximation to the singularity is then defined by

$$q := \operatorname{argmin}_{q \in V} \left\{ \sum_{i \in I} |q(t_i) - y_i|^2 \psi(\|x_i - x\|) \Lambda_i \right\}, \quad (10)$$

where V is a univariate polynomials space $\Pi_m(\mathbf{R})$, and $\psi(r)$ is a fast decreasing weight function. We now use the curve




Figure 9: Local data-dependent spline space for MLS the projection. In (a) we demonstrate the construction by a discontinuous spline space, and in (b) by a continuous spline space.

outlined by q as the locus of the local singularity of the space S_m to be used in Eq. (8). We offer two possible different constructions of the space S_m , one discontinuous and one continuous:

Type-0 singularity. In this case the functions in S_m are chosen to be discontinuous along the curve $q(t)$ (illustrated by the green line in Figure 9(a)). Thus, there are two independent polynomials each defined on each side of the curve q . In this case the approximation defined by Eq. 9 is equivalent to solving Eq. (8) only for points x_i which are on the same side as x with respect to the approximated singularity curve q . More formally, we use x_i if and only if

$$\operatorname{sign}(q(t_i) - y_i) = \operatorname{sign}(q(0) - y).$$

This is illustrated in Figure 9(a), where only the orange points participate in the approximation at x (the green point). Since we only take an approximation to the singularity, in order to better ensure we do not take points from the other side of the singularity, we can take points which also satisfy $|q(t_i) - y_i| > \frac{1}{4}|q(0) - y|$.

Type-1 singularity. In case of surface reconstruction it is usually required to force continuity. To that end, we would use continuous spline functions as the approximation space. Thus S_m is defined as the space of continuous piecewise polynomials, of total degree $\leq m$, with possible discontinuity in the derivatives across the line

$$\gamma(t) := (0, q(0)) + t(1, q'(0)), t \in \mathbf{R}.$$

γ is illustrated by the bold black line in Figure 9(b). For calculating the minimizer $p \in \mathcal{F}$ of Eq. (8) in this case, we augment the standard basis of Π_m by the piecewise polynomials. This can be done by first aligning the x -axis to γ and then adding as basis functions $x^\alpha(y)_+^\beta$, where $\alpha \geq 0$, $\beta \geq 1$ and $\alpha + \beta \leq m$.

If x is far from a singularity the system might turn out to be ill-conditioned. In such a case it is natural to turn back to the old approximation space Π_m : Practically, switching to $\mathcal{F} = \Pi_m$ for bad-conditioned system worked well for us.

Go recursively: singularities of singularities The same argument as above can be applied in the univariate setting. Instead of approximating the singularity using a smooth polynomial, we can use a univariate spline with a *point* singularity. Now, the SIF is to be considered as a point cloud with confidence values. The point of singularity can be deduced from a SIF of the SIF. In turn, the SIF of the singularity can be defined as described in section 4. See for example Figure 12.

6. Noisy Data

In the presence of noise, one needs to compromise since in the lack of any priors, one cannot distinguish noise from features. To handle noise we can clamp the SIF based on the following observation: The singularities in a surface (2-manifold) are (usually) a lower dimension manifolds, that is, curves (1-manifold) or points (0-manifolds). Hence, one can assume that most of the data points are not near a singular point. Define the modified SIF of the data Λ by:

$$\Lambda_j^{new} = H(\Lambda_j - pct), \quad (11)$$

where *pct* is, for instance, 90th percentile of the Λ , and H is the Heaviside function.

The practical effect of Eq. (11), is applying regular MLS in areas with SIF values smaller than *pct*, which are generally the smooth areas of the surface. Areas with SIF values larger than *pct* are approximated by a data-dependent MLS, taking into account the (now non-zero) SIF. Hence, reconstructing features at areas with large SIF values, relative to




Figure 10: (a) A photograph of a Lilly flower. In (b) we see the irregular sampled data (12% of the points), and in (c) a Lucy-Richardson deblurring applied to the MLS reconstruction. (d) depicts the DDMLS reconstruction. Note the faithful reconstruction of the silhouette edges and delicate lines in the background.




Figure 11: Noise direction analysis. (a) exhibits the data points (10K points) colored by the SIF values, and (b) the original data points rendered as a mesh to show the high noise level. Note the delicacy of the sharp feature w.r.t. the noise level. In (c) we see the reconstruction using DDMLS with a directionality analysis to smooth noise and preserve the features.

the rest of the surface. The effect of applying different *pct* values is demonstrated in Figure 8.

Another way of dealing with noise is based on the observation that noise does not have a structured shape like, for example, a direction, whereas features typically have. In the simplest form, this can be integrated into the method during weighted PCA construction of the local reference line. Let us define the directionality score of the local SIF as the ratio of the singular values in the weighted PCA: $\mu(x) := |s_1/s_2|$, $i \in I$, where $|s_1| > |s_2|$ are the singular values in the PCA analysis when considering point x . Then, fix a blending interval $[\mu_1, \mu_2]$, and a smooth blending function $\psi(t)$ where




Figure 12: Singularities of singularity curves. (a) shows the SIF on the jittered grid. In (b) the original function is drawn as an image, where the singularity curve has a point singularity. (c) exhibits reconstruction using a polynomial approximation to the singularity curve, and (d) reconstruction using a piecewise-polynomial to approximate the singularity curve.

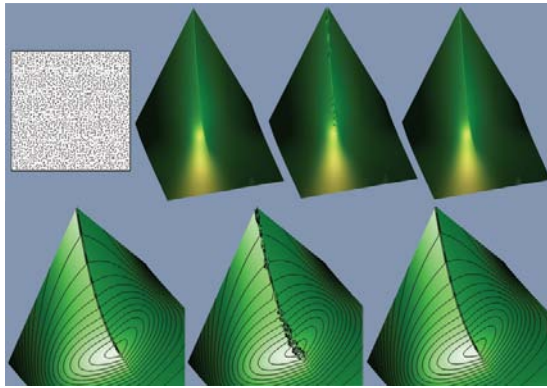


Figure 13: Comparison with Fleishman et al. [2005]. The original surface (left) has a sharp feature with a vanishing magnitude. The detection of a sharp feature using a threshold will necessarily fail. The surface is sparsely sampled (top left) with a jittered grid of 2.5K samples. The reconstruction with Fleishman et al. [2005] is shown in the middle, and with the DDMLS on the right. The faithfulness of the reconstruction is demonstrated with the isophotes (bottom row).

$\psi(t) = 0$ for $t < \mu_1$, and $\psi(t) = 1$ for $t > \mu_2$. Now define the operator as

$$\mathcal{M}_f^B(x) = \psi(\mu(x))\mathcal{M}_f(x) + (1 - \psi(\mu(x)))\widehat{\mathcal{M}}_f(x).$$

Figure 11 demonstrates this technique for resampling a relatively sparse noisy data with a delicate sharp feature. In this context, we note that using an enhanced SIF field can be useful, for example we have used $(\Lambda_i)^4$.

7. Discussion

The main contribution of this work is the use of an adaptive spline space in the MLS approximation via a novel feature detection method in a scattered data setting. An interesting observation is that we use the inability of the MLS to reconstruct discontinuities in order to detect them, as expressed by the SIF. We have applied DDMLS to real scanned models,

synthetic surfaces and images; for example Figures 1, 4, 6, 7, 8, 10. We have used mostly cubic piecewise polynomials ($m = 3$).

Our method has the important property of being threshold free, which is important for faithful surface reconstruction as well as for processing and representing piecewise smooth geometry. The local singularity structure is realized through a spline approximation space used in the MLS method, where the piecewise polynomial singularity locations are defined in a relative, local, threshold free manner. Note however, that as we described in Section 6, there is a threshold parameter that is used in cases of noisy data, in order to distinguish features from noise.

In Figure 13 we compare our method to the method of Fleishman et al. [2005]. We applied the two methods to a surface with a sharp feature of a magnitude decreasing to zero. The detection of such a sharp feature using a threshold based method will necessarily fail. The surface is sparsely sampled with a jittered grid of 2.5K samples. The faithfulness of the reconstruction is demonstrated by isophotes curves. Moreover, note that the iterative construction of Fleishman et al. [FCOS05] does not consider the singularity as manifold. Hence, it cannot guarantee a consistent approximation, nor a continuous reconstruction across sharp edges.

Another advantage of the new method is the ability to reconstruct sharp features using relatively sparse point set. For example, the models in Figures 4, 6, 13 consist of only 5K, 12K, 2.5K input data points, respectively. Generally, we retain the computational complexity of the standard MLS since we perform two MLS projections for each point. In our MATLAB implementation the method projects around 200 points per second, and the preprocess calculates 500 Λ values per second. This is two orders of magnitude faster than the method of Fleishman et al. [FCOS05]. For example, in Figure 13, Fleishman's method is 70 times slower than DDMLS.

Currently, the limitation of the technique is that it requires careful tuning in case of noisy data. Otherwise, it may amplify the noise or suppress delicate features. Another limitation is that it does not deal well with non-manifold singularity structure, like the corner of a cube. An interesting research direction is to extend the method to construct such non-manifold singularities. Note that the presented method can be applied to data of arbitrary dimension. In our setting, it is applied to a 2D surface while respecting 1D singularity. However, it can also be applied to volumetric data with discontinuities. Another interesting direction, which is demonstrated by a simple example in Figure 12, is the detection of breakpoints in curves of sharp features and using this to further adapt the local spline space.

8. Acknowledgements

This work was supported by the Israel Science Foundation.

References

- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C.: Point set surfaces. *Proceedings of the conference on Visualization'01* (2001), 21–28.
- [ABCO*03] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Computing and rendering point set surfaces. *IEEE Transactions on Visualization and Computer Graphics* 9, 1 (2003), 3–15.
- [AK04] AMENTA N., KIL Y. J.: Defining point-set surfaces. *ACM Trans. Graph.* 23, 3 (2004), 264–270.
- [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 *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 67–76.
- [FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. *ACM Trans. Graph.* 24, 3 (2005), 544–552.
- [HDD*92] HOPPE H., DE ROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. *Computer Graphics* 26, 2 (1992), 71–78.
- [HEOC87] HARTEN A., ENGQUIST B., OSHER S., CHAKRAVARTHY S. R.: Uniformly high order accurate essentially non-oscillatory schemes, 111. *J. Comput. Phys.* 71, 2 (1987).
- [JP00] JIANG G.-S., PENG D.: Weighted eno schemes for hamilton-jacobi equations. *SIAM Journal on Scientific Computing* 21, 6 (2000), 2126–2143.
- [JS96] JIANG G.-S., SHU C.-W.: Efficient implementation of weighted eno schemes. *J. Comput. Phys.* 126, 1 (1996), 202–228.
- [LCOL06] LIPMAN Y., COHEN-OR D., LEVIN D.: Error bounds and optimal neighborhoods for mls approximation. *Symposium on Geometry Processing* (2006).
- [Lev98] LEVIN D.: The approximation power of moving least-squares. *Math. Comput.* 67, 224 (1998), 1517–1531.
- [Lev03] LEVIN D.: Mesh-independent surface interpolation. *Geometric Modeling for Scientific Visualization* (2003), 37–49.
- [LS81] LANCASTER P., SALKAUSKAS K.: Surfaces generated by moving least squares methods. *Math. Comp.* 37, 155 (1981), 141–158.
- [OBA*03] OHTAKE Y., BELYAEV A., ALEXA M., TURK G., SEIDEL H.-P.: Multi-level partition of unity implicit. *ACM Trans. Graph.* 22, 3 (2003), 463–470.
- [PKKG03] PAULY M., KEISER R., KOBBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. *ACM Trans. Graph.* 22, 3 (2003), 641–650.
- [RJT*05] REUTER P., JOYOT P., TRUNZLER J., BOUBEKEUR T., SCHLICK C.: Surface reconstruction with enriched reproducing kernel particle approximation. In *Proceedings of the IEEE/Eurographics Symposium on Point-Based Graphics* (2005), Eurographics/IEEE Computer Society, pp. 79–87.
- [SO89] SHU C.-W., OSHER S.: Efficient implementation of essentially non-oscillatory shock-capturing schemes,ii. *J. Comput. Phys.* 83, 1 (1989), 32–78.
- [Wen01] WENDLAND H.: Local polynomial reproduction and moving least squares approximation. *IMA Journal of Numerical Analysis* 21, 1 (2001), 285–300.