SCI INSTITUTE

TECHNICAL REPORT

A Unified Projection Operator for Moving Least Squares Surfaces

Tilo Ochotta, Carlos Scheidegger, John Schreiner, Yuri Lima, Robert M. Kirby and Claudio Silva

UUSCI-2007-006

Scientific Computing and Imaging Institute
University of Utah
Salt Lake City, UT 84112 USA

April 25, 2007

Abstract:

Moving-Least Squares (MLS) surfaces are a popular way to define a smooth manifold surface from a set of unorganized points without normals. In this paper we present a formulation of MLS surfaces that sheds light on shortcomings of the original technique proposed by Levin, which is based on a two-step minimization procedure. We show that there are cases intrinsic to the geometry of the underlying surface from which the points are sampled where Levins projection fails to find an adequate fit. These shortcomings occur regardless of sampling density or the amount of noise. Our formulation solves this problem by directly fitting a local approximating function to the surface using a unified minimization scheme. We present a modification of Levins original proof that can be directly adapted to our unified approach. Consequently, this suggests our method can be used to create different families of MLS surfaces, depending on the function space used for the fit. This allows specific priors to be used in the approximation, leading to better reconstructions. We present experimental results that show our technique performs adequately in a wide range of conditions.

THE UNIVERSITY OF UTAH

A Unified Projection Operator for Moving Least Squares Surfaces SCI Institute Tech Report UUSCI-2007-006

Tilo Ochotta Carlos Scheidegger John Schreiner
Yuri Lima Robert M. Kirby Claudio Silva

April 25, 2007

Abstract

Moving-Least Squares (MLS) surfaces are a popular way to define a smooth manifold surface from a set of unorganized points without normals. In this paper we present a formulation of MLS surfaces that sheds light on shortcomings of the original technique proposed by Levin, which is based on a two-step minimization procedure. We show that there are cases intrinsic to the geometry of the underlying surface from which the points are sampled where Levin's projection fails to find an adequate fit. These shortcomings occur regardless of sampling density or the amount of noise. Our formulation solves this problem by directly fitting a local approximating function to the surface using a unified minimization scheme. We present a modification of Levin's original proof that can be directly adapted to our unified approach. Consequently, this suggests our method can be used to create different families of MLS surfaces, depending on the function space used for the fit. This allows specific priors to be used in the approximation, leading to better reconstructions. We present experimental results that show our technique performs adequately in a wide range of conditions.

1 Introduction

Recently, there has been substantial interest in the area of surface reconstruction from point-sampled data. This work is driven by a set of important applications where the ability to define surfaces out of a set of discrete samples is necessary. For instance, devices capable of acquiring high-resolution 3D models have become affordable and commercially available, and such reconstruction techniques are required for the effective use of these devices. A particularly powerful approach has been the use of the moving least-squares (MLS) technique of Levin [Lev03] for modeling point-set surfaces [ABCO*01, ABCO*03, AK04a]. Variants of this framework have become the basis for much of the current point-based modeling work in the graphics, visualization, and computational geometry communities.

The key idea in Levin's original formulation is to define the surface in terms of a projection operator, the fixed points of which are defined as the surface. The idea

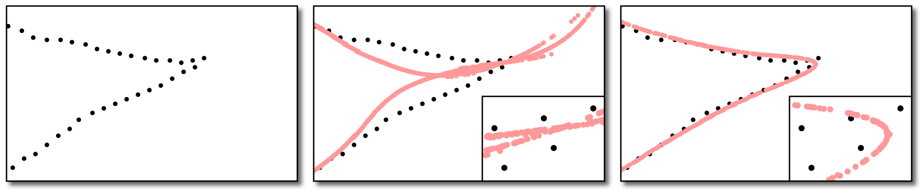

(a) (b) (c)

Figure 1: When Levin's original MLS projection [Lev03] is applied to a wedge shaped set of input points (shown in black) (a), the resulting surface (shown in red) can pinch in an undesirable way as seen in (b). The surface pinches because the local reference frame bisects the input points (see Figure 4), which does not allow a good functional approximation to the points. Our new unified projection operator avoids this pinching artifact and rounds the surface as one might expect (c).

is similar to the seminal work of Lancaster and Salkauskas [LS81] for the interpolation and approximation of functions. Levin generalizes this previous work in function approximation theory to accommodate manifolds. Levin's operator involves a non-linear optimization for each point projection, but unlike simpler, subsequent definitions [AA03], it does not require as input normal information at the points. The operator is defined as a two-step optimization procedure. The first step computes a reference frame for a local neighborhood of the point being projected by a non-linear weighted least squares fit. (Note that this is the main difference from [LS81].) The second (linear) step finds a best-approximating function in the reference frame computed during the first step. Typically a tensor-product quadratic is fit to the input points, from which differential-geometric properties of the MLS surface can be approximated. It is possible to skip the second step, which is equivalent to *fitting* a zero degree function to the point set, and this is advocated in many works (e.g., [AK04a]). This second step is critical for our reformulation of the projection procedure. We show that assuming a constant function will make the original MLS definition fail to produce a suitable surface reconstruction for certain geometric configurations, regardless of neighborhood size or sampling density.

It can be demonstrated that cases exist when during the finding of a non-constant function f in the second stage, the reference frame computed by the first stage will generate a poor fit to the input points. The motivating insight for our reformulation lies in the observation that for the same input configuration *there still exists another reference frame that would allow an accurate fit* (see Figure 4). We address this issue by incorporating the function fitting into the non-linear optimization, thereby unifying the projection procedure into a single fitting step. Furthermore, we show that it is possible to tailor this new MLS formulation for a variety of geometric processing tasks by changing the function space from which the function f is selected.

2 Related work

The problem of defining surfaces out of point samples has been actively approached by researchers for many years. Pioneering work has been done by Hoppe et al. [HDD*92] and others [CL96, TL94, WSI98] in the context of surface reconstruction, where the primary focus was on building triangle meshes of the sampled surface. Early important work in the area was also done in the computational geometry community, particularly in the area of connectivity reconstruction using techniques based on the Delaunay triangulation [ABK98].

During the development of *Point-Based Graphics* as an independent sub-area of interest [PZvBG00, RL00], it was natural to consider the more general problem of defining surfaces directly from point sets. Alexa et al. [ABCO*01]'s Point Set Surfaces was one of the first papers in this area. It showed how a simple and effective representation could be achieved by the use of a moving least-squares (MLS) technique [Lev03], even in the case of noisy sets of points. Other related formulations followed, e.g., [AK04a, FCOAS03, MVF03, PKKG03, XWH*03].

At this point, there are many variations and extensions of the original approach to defining point set surfaces. A set of popular techniques are based on defining the point projection in terms of a combination of weighted centroids and a normal field [AA07, AA04, AA03, AK04a]. This type of projection defines the surface in terms of a level set of a scalar function. This function can be evaluated very efficiently, but an iterative method is required to project points onto the surface. Also, their simplicity makes them more suitable for analysis that give strong theoretical guarantees [Kol05, AA07, DS05, AK04a, BH05]. We note that some of these linear techniques require normal information, which may be unavailable or unreliable.

There is also continued interest in further analysis of the original approach using Levin's MLS [Lev03] which involve a non-linear projection, e.g., [LCOL06, FCOS05, SSFS06, ABCO*03]. The work of Amenta and Kil [AK04b] raised many important practical and theoretical issues, including what happens to the projection near edges and corners; they show that the original projection sometimes has undesirable behavior in those locations. Fleishman et al. [FCOS05] propose to use robust statistics and a modified projection scheme for recovering sharp features as well as increasing stability of the projection operator near features. This approach is sound and produces good results, but it introduces extra processing steps which break the natural elegance of the original formulation.

In our work, we follow on the footsteps of Amenta and Kil, and show that the failures near corners and edges are unrelated to sampling conditions – they are in fact intrinsic to the geometry of the surface from which the points are sampled. Our work also proposes an alternative way to define sharp features. Instead of using multiple steps, our technique is based on a conceptually simple modification (extension) of the original non-linear projection by Levin.

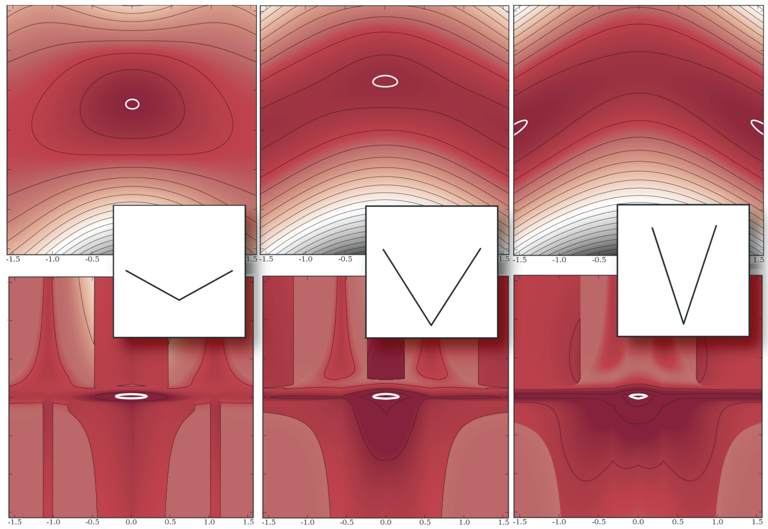

Figure 2: Energy fields for several input point configurations are visualized for Levin's projection (Equation (1), top row), and our unified approach (Equation (4), bottom row). The horizontal and vertical axis correspond to a parameterizaton of all possible reference frames, whose center is at an angle θ from the vertical wedge direction, and a distance ρ from the center point of the wedge, respectively. Notice that for the sharpest wedge, the minimum (always encircled in a white contour) appears in the wrong place for Levin's formulation, while ours has a minimum with angle zero in all situations.

3 The Unified Projection Operator

We wish to project a point r in a domain $D \subseteq \mathbb{R}^3$ to a point on the MLS surface, defined by a set of input points p_i. This is done by defining a function $\mathcal{P} : D \to \mathcal{S}$. We begin by reviewing Levin's original MLS projection. We then show how our new projection relates to Levin's formulation, and give a proof to show that it is indeed a projection.

3.1 Levin's Projection

Levin's projection [Lev03], $\mathcal{P}_L$, has two stages. It first looks for a local reference frame defined by a point q relative to r. This is the plane that passes through q, with normal $(q-r)/|q-r|$, and is denoted by $H_{r,q}$. The orthogonal projection onto the plane is denoted by $h_{r,q} : \mathbb{R}^3 \to H_{r,q}$. The plane $H_{r,q}$ is found by

$$\underset{q}{\operatorname{argmin}} \frac{\sum_i |h_{r,q}(p_i) - p_i|^2 \omega(|p_i - q|)}{\sum_i \omega(|p_i - q|)}, \tag{1}$$

where ω is a weighting function (often a Gaussian). Note that a non-linear minimization is required to find this weighted least-squares fit due to the weighting function's

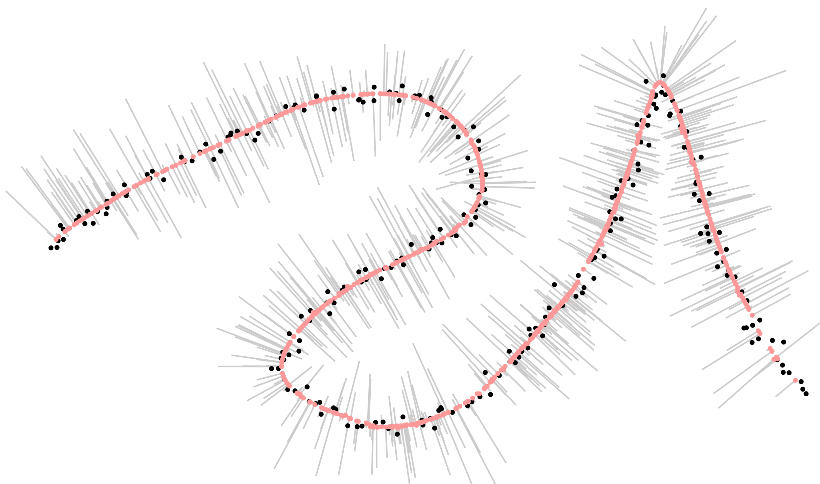

Figure 3: Like the original MLS, our unified projection is quite resilient to noisy and irregularly sampled inputs. Our approach has the extra benefit of properly rounding sharp corners. The grey lines show the paths of the point projections.

dependence on q, and is typically found using a Powell method [PTVF92, Chapter 10.5]. Once this local reference frame has been found, a second step is performed where a function $f : H_{r,q} \to \mathbb{R}$, from a function space $\mathcal{F}$, is fit to the points p_i, in a weighted least-squares sense. Let $g_{r,q} : (H_{r,q} \times \mathbb{R}) \to \mathbb{R}^3$ be the linear function mapping the plane and a scalar value back into the global $\mathbb{R}^3$ domain (similar to h^{-1}). With r,q subscripts omitted for clarity, the function f can then be found by

$$\operatorname*{argmin}_f \frac{\sum_i |g(h(p_i), f(h(p_i))) - p_i|^2 \omega(|q - p_i|)}{\sum_i \omega(|q - p_i|)}. \tag{2}$$

Intuitively, this can be thought of as finding a weighted least squares function approximation to the points p_i over the $\mathbb{R}^2$ local domain formed by H. Since the weights are constant, the function f can be computed directly by solving the resulting linear system, when $\mathcal{F}$ is a finite-dimensional linear function space. After the function f is found, Levin's projection of r onto the MLS surface is simply defined to be $\mathcal{P}_L(r) = g(h(r), f(h(r)))$, i.e. the projection of r onto f in the direction $(q-r)/|q-r|$.

This two-step formulation works well when applied to most point configurations, but can sometimes break down. The intuition for the two-step process is that the first step is trying to find a local reference frame for the surface, analogous to the traditional differential-geometric reference frame, while the second step actually computes the function.

Unfortunately, the first step of the process fails to produce the expected reference frame when the points p_i are in certain, but common, geometric configurations. Specifically, we consider the covariance matrix of the neighborhood of points p_i around r, represented by the covariance matrix of the points. The reference frame found by the optimization process will have its normal aligned with the minor eigenvector of the matrix. In some cases, this is not the normal of the natural reference frame for the local surface (see Figure 4(left)). When this happens, the MLS projection operator will pinch and extend away from the points p_i (as demonstrated in Figure 1(b)).

Even if the function space used in the second step of the projection procedure is capable of accurately representing the local surface, it needs a reasonable reference frame from which to be applied. It is natural and probably intuitive to target the sampling density as the source of the problem. In particular, we know from differential geometry that an appropriate plane must exist, but the original MLS definition is simply not finding it. The reason for the original MLS definition's deficiency is that it is performing the optimization in two different spaces. The first stage of the projection is fundamentally limited to planar approximations, and there is no chance for the second stage to correct (or adequately compensate for) the problems generated by the first stage.

3.2 Unifying the Projection Operator

We now describe a generalization of the Levin projection that unifies the search for H and f. Since the search for H is the aspect that leads to poor results, we take the approach of trying to minimize its influence on the projection procedure. We begin by assuming that the function space $\mathcal{F}$ is the set of constant functions. Then when H is a minimizer of Equation (1), Equation (2) will be minimized by $f = 0$. Since $g(h(p_i), f(h(p_i))) = h(p_i)$, we can then rewrite Equation (1) as

$$\operatorname*{argmin}_{q,f} \frac{\sum_i |g(h(p_i), f(h(p_i))) - p_i|^2 \omega(|q - p_i|)}{\sum_i \omega(|q - p_i|)}. \tag{3}$$

In this simplified setting, with constant f, it is clear that equation (3) is equivalent to equation (1). We further generalize the projection by writing

$$\operatorname*{argmin}_{q,f} \frac{\sum_i |C_{H,f}(p_i) - p_i|^2 \omega(|C_{H,f}(r) - p_i|)}{\sum_i \omega(|C_{H,f}(r) - p_i|)}, \tag{4}$$

where $C_{H,f}(x)$ is the closest point projection of x onto the function f over domain H, and the result of the unified projection is defined to be $\mathcal{P}_U(r) = C_{H,f}(r)$. Again, when f is constant, this is equivalent to equation (1). For more complex function spaces $\mathcal{F}$, however, f is not constant, and the unified projection is quite different than Levin's version. The first difference is that the weights computed by ω are in reference to the final result of the projection, rather than the projection onto the plane H. This is one step in removing the dependence on the local frame on which f is defined. A second difference is in the way the error residuals are measured when fitting the function f. The original projection measured the residuals in the direction orthogonal to H. Our unified approach measures them orthogonally to the function f itself. The encapsulation of H into the closest point projection function C effectively removes any major dependence on the local reference frame. H can now simply be thought of as a means of parameterizing all possible functions in $\mathcal{F}$, over all orientations and translations.

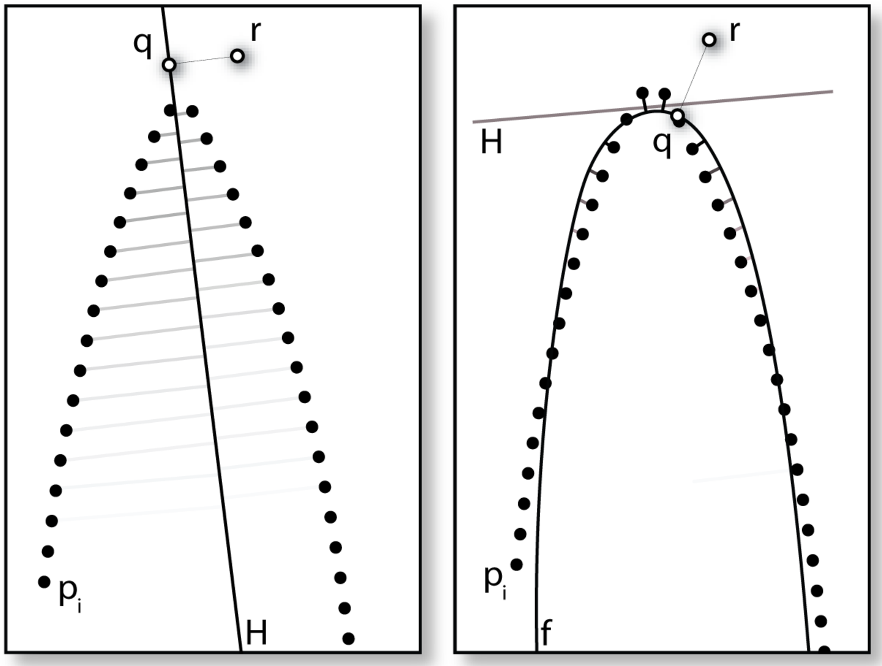

Figure 4: When projecting a point r onto the surface, Levin's MLS (left) may find a reference frame H that does not naturally allow the input points to be approximated by a function. Our approach (right) searches for H and the approximating function f simultaneously. Errors are measured differently. In Levin's, the error residuals are measured perpendicular to H, and weighted by the distance to q. In ours, the residuals are measured by closest point distances, weighted by the closest point projection of r onto f.

Another significant change is in the way that the point r is projected onto the function f approximating the surface. Levin's projection always projects in the direction $(q-r)/|q-r|$. The unified projection is always orthogonal to f. This allows f to be found in a way that best approximates the surface, followed by a projection onto that approximation.

In essence, our projection operator is aware of the local function approximations being used, rather than searching for a reference frame with the hope that a function fit over that domain will be adequate. Levin's projection essentially searches the reference frame whose normal agrees with the minor eigenvector of the inertia tensor. However, this might be a reference frame in which the function fit and resulting surface approximation to p_i will be poor. With the new operator, finding the local reference frame on which f is defined is an effect of finding a good surface approximation.

3.3 Proof of Projection

An important question that must be answered is whether or not our new operator, $\mathcal{P}_U$, is still a projection (for the reasons outlined in Levin's original work). We start by showing an alternative proof that Levin's operator is a projection. We refer the reader to step 1 of Section 3 in [Lev03]. Assume Levin's operator was applied to a point r such that $P_L(r) = q$, and $r \neq q$. We will now examine the effect of P_L in a one-dimensional neighborhood of r along $(q-r)/|(q-r)|$ (Levin calls this direction $a = a(q)$). Let $r' = r + u(r-q), u \in (-\varepsilon, \varepsilon)$, with ε such that the stated assumptions in Levin's work hold. It is clear that $a' = (q-r')/|q-r'| = a$, so using q and a as directions for projecting r' is admissible. Then, the first and third conditions in [Lev03] hold for r', and the second is satisfied by construction. The consequence of these conditions is that q is the result of the first stage of Levin's projection for all u. Since the second stage of the projection procedure only moves q in the direction a and does not depend on r, the final result is

$$\nabla_a(P_L(r)) = \mathbf{0}. \tag{5}$$

This means that the null space of the Jacobian $J(P_L(r))$ is non-trivial, and therefore rank-deficient. Therefore it follows that P_L is a projection.

In order to prove that our operator is indeed a projection we will similarly show that the Jacobian of $P_U(r)$ is rank-deficient. We use notation as used previously: $P_U(r) = q, r \neq q, a = (q-r)/|(q-r)|$, and so forth. If q minimizes Equation (4) for r, then there exists an ε for which any $r' = r + u(r-q)$, $u \in (-\varepsilon, \varepsilon)$ has the same minimizer. First note that $C_{H,f}(r') = q$ for all u, and therefore $\nabla_a(C_{H,f}(r)) = \mathbf{0}$. This is because the closest point projection follows a distance field whose gradient agrees with a. Now, since the only dependence of Equation (4) on r is through $C_{H,f}(r)$, applying the chain rule results with

$$\nabla_a(P_U(r)) = \mathbf{0}. \tag{6}$$

Again, since the null space of the Jacobian is non-trivial, P_U is a projection. In the case of $r = q$, we simply set a to be in the direction of the normal of f at q, and the same argument applies. This result is quite general as any function space used in P_U will produce a projection operator.

4 Results and Discussion

Our proposed operator has several desirable features, compared to the original MLS definition and subsequent work. It directly fits local functions to the input, and so is less susceptible to incorrect minima arising from an inadequate reference frame. Therefore, some feature from the sampled model might be "pinched" in these definitions. In our unified projection, these corners are "rounded" — the behavior is more similar to a surface under low-pass filtering. This can be seen in Figures 1, 3, and 8.

At the same time, our formulation can also be used to exactly reconstruct C^0 features of the surface. If the functions we are using for the optimization are only C^0, there will be 3-dimensional balls in R^3 that project down to a single point, as illustrated in Figure 5. This runs counter to the intuition that our surfaces should be as continuous as

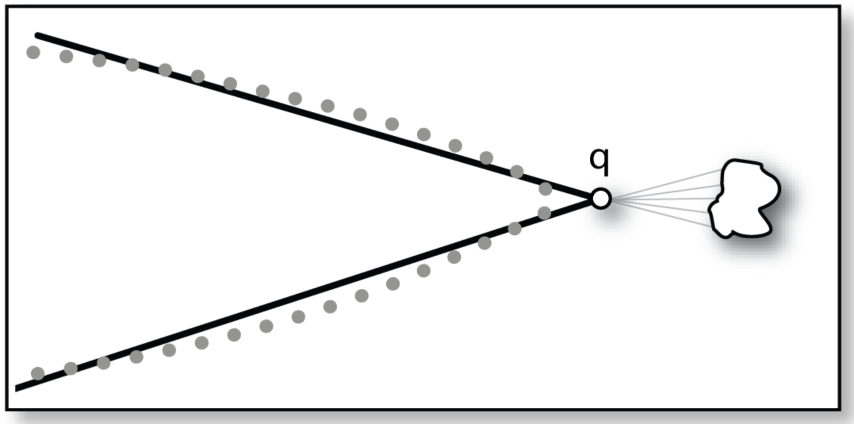

Figure 5: The unified operator can reconstruct C^0 surfaces when the function space $\mathcal{F}$ consists of wedge shaped functions. A d-dimensional ball R collapses to a point q at C^0 features under the closest point projection. Since this is the point that determines the weights, if q is a minimum for some $r \in R$, then it is a minimum for all of them. Hence the surface retains the C^0 feature.

the weighting functions, as Levin conjectures in [Lev03]. To the best of our knowledge, this is the first technique to directly incorporate sharp features, without formulations that use intersections of smooth surfaces or explicit tagging [FCOS05,RJT*05]. Figure 6 illustrates the effect using $f(x) = |x|$ as the single basis vector for the function space, which is then searched over all origins and orientations.

It is interesting to consider the conditions that cause the breakdown of the original MLS operator. This happens when the minor eigenvector of the inertia tensor switches from being aligned with the normal of the expected reference frame to being tangent to the frame. We believe that this is related to the uniqueness of the minor eigenvector of the inertial tensor and how it changes over the surface. Bremer and Hart [BH05] have recently used a related technique for proving sampling conditions under which an MLS formulation properly reconstructs a surface.

The flexibility that the unified projection operator affords can be exploited for unusual applications. One can define the function space $\mathcal{F}$ used by the unified operator to contain the expected local surfaces for the model being reconstructed. By biasing the function space in this way, we can interpret this unified formulation in terms of providing *surface priors* for the reconstruction. One interesting application of this technique is for automatic completion. In Figure 7, we show two examples of reconstructions from an incomplete set of input points. On the right, the function space $\mathcal{F}$ only contains smooth functions, so a smooth surface is reconstructed. In contrast, the left shows a a case where the $\mathcal{F}$ includes functions with sharp features, so a sharp point is reconstructed to fill the gap.

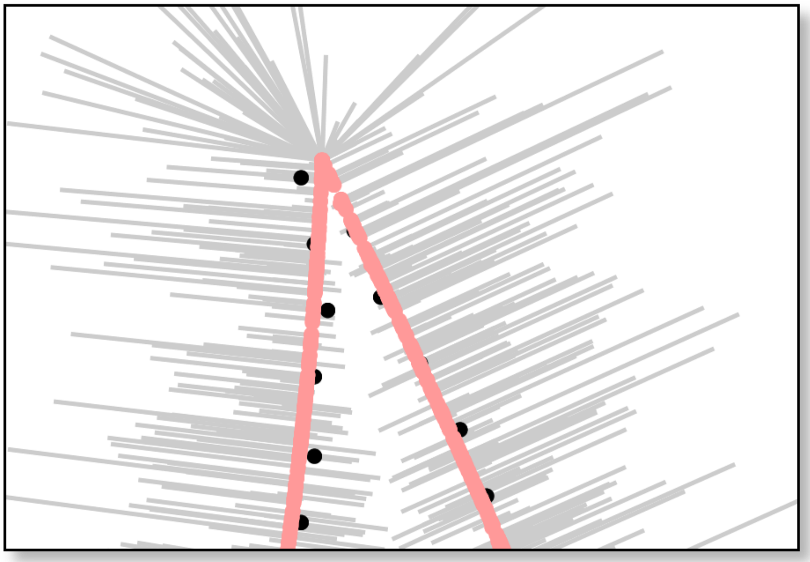

Figure 6: Using C^0 functions in the function space leads to a C^0 reconstruction of the surface.

5 Conclusions and Future Work

In this paper, we have proposed a unified MLS projection operator. Our operator is close in spirit to [Lev03, LS81], and it allows for the use of rich function spaces that contain functions relevant to the surface being reconstructed. It is a natural extension to Levin's approach, and addresses a key shortcoming of the technique. We have provided an alternate proof that Levin's formulation is a projection, which naturally extends to our more general operator.

We believe that our approach has many applications. Our unified projection allows manifolds to be defined on embedding spaces richer than R^3. For example, it should be possible to reconstruct color information by operating on the product space of color and geometric information. The same technique should be applicable to texture coordinates, and other continuous fields associated to the manifold.

An exciting avenue for future work is to automatically determine the most likely priors within the point set and use those to iteratively improve the point set surface definition, computing an intrinsically best point set surface for the particular scan. Similarly, it may be possible to use parts of a surface as priors for the reconstruction of other, more poorly sampled areas.

Acknowledgments

Removed for anonymity.

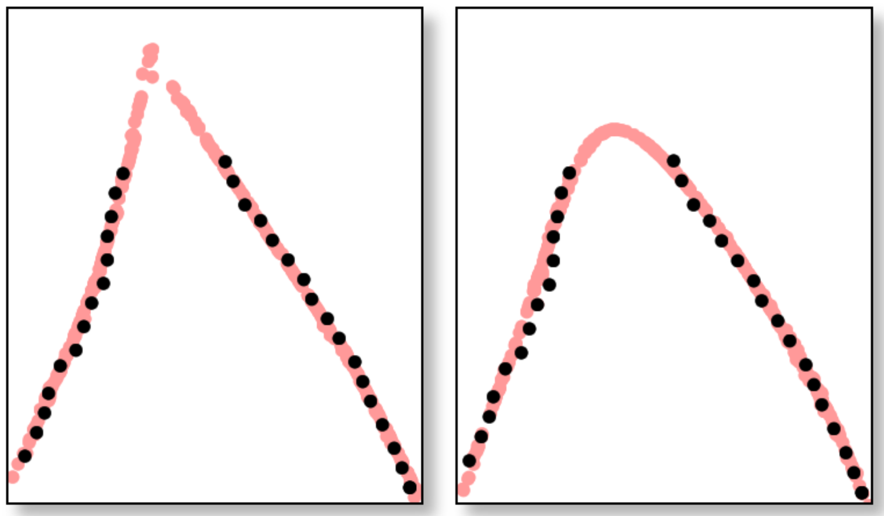

Figure 7: A function space $\mathcal{F}$ that contains predetermined priors that fit the original surface can be used to complete missing details. On the left, $\mathcal{F}$ contains wedge shaped functions, so a sharp point is reconstructed to fill the gap. On the right, $\mathcal{F}$ only contains smooth functions, so the gap is filled with a smooth reconstruction.

References

[AA03] ADAMSON A., ALEXA M.: Approximating and intersecting surfaces from points. In *Proc. Symposium on Geometry Processing* (2003), pp. 230–239.

[AA04] ALEXA M., ADAMSON A.: On normals and projection operators for surfaces defined by point sets. In *Proc. Symposium on Point-Based Graphics 2004* (2004), pp. 149–155.

[AA07] ALEXA M., ADAMSON A.: *Interpolatory Point Set Surfaces - Convexity and Hermite Data*. Tech. rep., T.U. Berlin, 2007.

[ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Point set surfaces. In *Proc. IEEE Visualization* (2001), pp. 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.

[ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. *Proc. ACM SIGGRAPH* (1998), 415–422.

[AK04a] AMENTA N., KIL Y. J.: Defining point-set surfaces. *ACM Transactions on Graphics 23*, 3 (2004), 264–270.

[AK04b] AMENTA N., KIL Y. J.: The domain of a point set surface. In *Proc. Symposium on Point-Based Graphics* (2004), pp. 139–147.

[BH05] BREMER P.-T., HART J. C.: A sampling theorem for mls surfaces. In *Proc. Symposium on Point-Based Graphics* (2005), pp. 47–54.

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

[DS05] DEY T. K., SUN J.: An adaptive mls surface for reconstruction with guarantees. In *Proc. Symposium on Geometry Processing* (2005), pp. 43–52.

[FCOAS03] FLEISHMAN S., COHEN-OR D., ALEXA M., SILVA C. T.: Progressive point set surfaces. *ACM Transactions on Graphics 22*, 4 (2003), 997–1011.

[FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. *ACM Transactions on Graphics 24*, 3 (2005), 544–552.

[HDD*92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proc. ACM SIGGRAPH* (1992), pp. 71–78.

[Kol05] KOLLURI R.: Provably good moving least squares. In *Proc. ACM-SIAM Symposium on Discrete Algorithms* (2005), pp. 1008–1017.

[LCOL06] LIPMAN Y., COHEN-OR D., LEVIN D.: Error bounds and optimal neighborhoods for mls approximation. In *Proc. Symposium on Geometry Processing* (2006), pp. 71–80.

[Lev03] LEVIN D.: *Geometric Modeling for Scientific Visualization*. Springer-Verlag, 2003, ch. Mesh-independent surface interpolation, pp. 37–49.

[LS81] LANCASTER P., SALKAUSKAS K.: Surfaces generated by moving least squares methods. *Mathematics of Computation 37*, 155 (1981), 141–158.

[MVF03] MEDEROS B., VELHO L., FIGUEIREDO L.: Moving least squares multiresolution surface approximation. In *XVI Brazilian Symposium on Computer Graphics and Image* (2003), pp. 19–26.

[PKKG03] PAULY M., KEISER R., KOBBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. *ACM Transactions on Graphics 22*, 3 (2003), 641–650.

[PTVF92] PRESS W., TEUKOLSKY S., VETTERLING W., FLANNERY B.: *Numerical Recipes in C*. Cambridge, 1992.

[PZvBG00] PFISTER H., ZWICKER M., VAN BAAR J., GROSS M.: Surfels: Surface elements as rendering primitives. In *Proc. ACM SIGGRAPH* (2000), pp. 335–342.

[RJT*05] REUTER P., JOYOT P., TRUNZLER J., BOUBEKEUR T., SCHLICK C.: *Point Set Surfaces with Sharp Features*. Tech. rep., LaBRI, 2005.

[RL00] RUSINKIEWICZ S., LEVOY M.: Qsplat: A multiresolution point rendering system for large meshes. In *Proc. of ACM SIGGRAPH* (2000), pp. 343–352.

[SSFS06] SCHREINER J., SCHEIDEGGER C. E., FLEISHMAN S., SILVA C. T.: Direct (re)meshing for efficient surface processing. *Computer Graphics Forum 25*, 3 (2006), 527–536.

[TL94] TURK G., LEVOY M.: Zippered polygon meshes from range images. In *Proc. ACM SIGGRAPH* (1994), pp. 311–318.

[WSI98] WHEELER M. D., SATO Y., IKEUCHI K.: Consensus surfaces for modeling 3D objects from multiple range images. In *ICCV '98* (1998), pp. 917–924.

[XWH*03] XIE H., WANG J., HUA J., QIN H., KAUFMAN A.: Piecewise g^1 continuous surface reconstruction from noisy point cloud via local implicit quadric regression. In *Proc. IEEE Visualization* (2003), pp. 259–266.

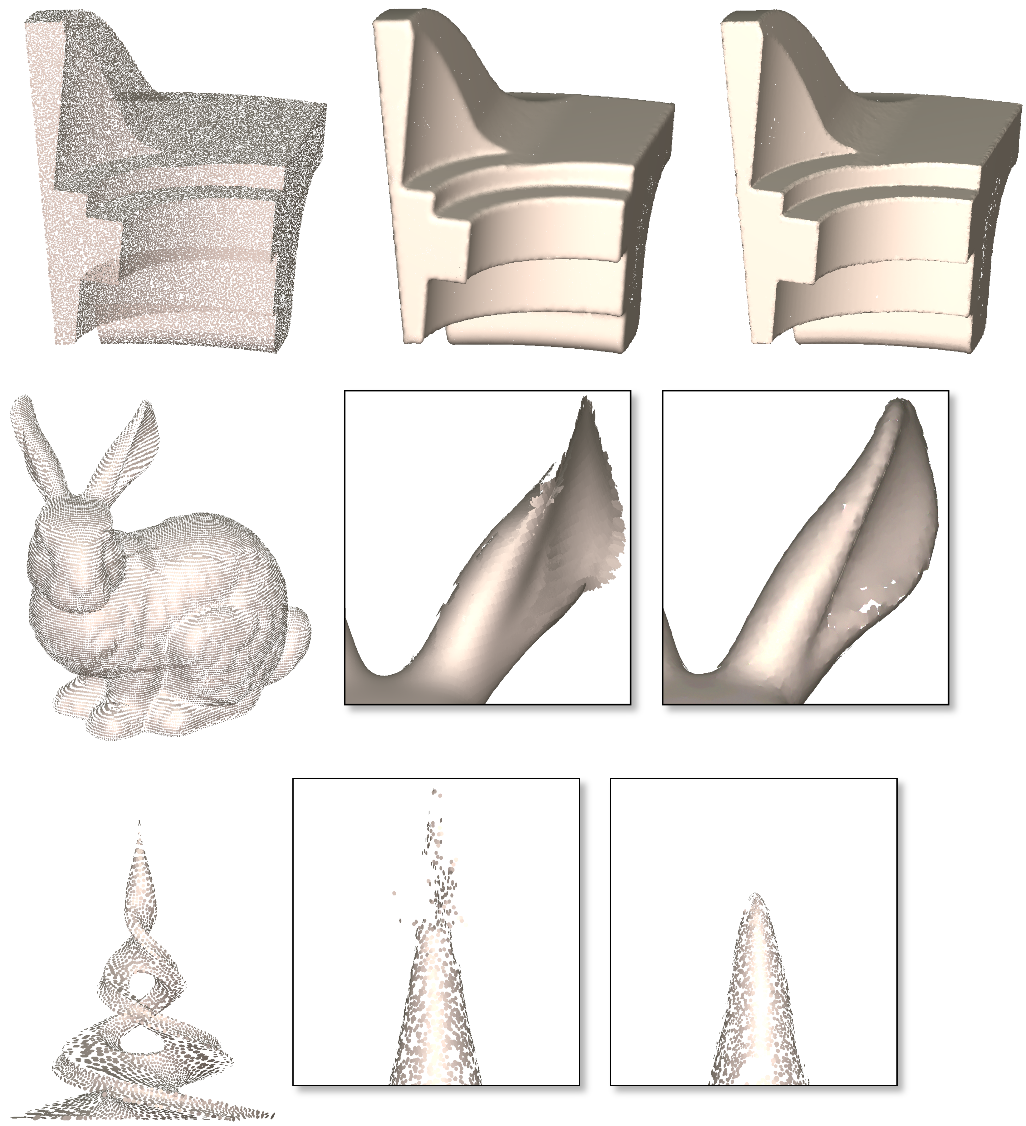

Figure 8: On reasonably well-behaved models like the fandisk (first row), our formulation (right columns) performs similarly to Levin's (middle column). On an undersampled point set of the Stanford bunny, or the Twirl model, Levin's projection produces degeneracies, while our direct fitting generates an appropriate surface.