PhD Showcase: Automated artifact-free seafloor surface reconstruction with two-step ODETLAP

PhD Student: Tsz-Yam Lau
Rensselaer Polytechnic Institute
110, Eighth Street
Troy, NY, USA
laut@cs.rpi.edu

PhD supervisor: W. Randolph Franklin
Rensselaer Polytechnic Institute
110, Eighth Street
Troy, NY, USA
mail@wrfranklin.org

ABSTRACT

We present an updated artifact-free seafloor surface reconstruction scheme which preserves more terrain features than our previous attempt using overdetermined Laplacian Partial Differential Equation (ODETLAP) and automates the adjustment of smoothing parameter. The high resolution version of such a surface fitting problem remains a challenge since we are still confined to extremely unevenly distributed depth samples collected along and near the ships, in which case numerous generic reconstruction algorithms generate unacceptable surfaces featuring abnormal depth fluctuations which are correlated with the trackline locations. Previously we reported the use a modified ODETLAP scheme, which integrates data-density-dependent smoothing into the reconstruction process, to generate surfaces which are free from such acquisition footprint. However, that scheme still suffers from terrain feature loss due to smoothing, and the reliance of human to decide appropriate smoothing factor. This paper aims to fix these two problems with a two-step ODETLAP procedure. The procedure first applies an accuracy-biased ODETLAP to complete the missing depth data from the given samples. After that, the vigorous depth fluctuations along the tracklines are removed by applying a smoothing-biased ODETLAP on the completed depth grid. To decide the optimal smoothing factor automatically, the procedure computes the areas of the individual bumps on the reconstructed surface. A surface suffering heavily from the artifacts has many small bumps but few big ones. Smoothing reduces such skewness. We find that for many datasets, the artifact is mostly gone when the coefficient of variation of the areas drops to around 1.3. Using that value to gauge the smoothing factor, the automated scheme successfully generates artifact-free seafloor surfaces within a limited error budget.

Categories and Subject Descriptors

I.3.5 [**Computing Methodologies**]: Computer Graphics Computational Geometry and Object Modeling

The SIGSPATIAL Special, Volume 4, Number 3, November 2012.
Copyright is held by the author/owner(s).

Keywords

GIS, ODETLAP, bathymetry, surface reconstruction, sparse height grid

1. INTRODUCTION

A bathymetric chart, the underwater equivalent of a topographic map, represents the depth and features of the ocean floor. This data helps solve not only applications such as tsunami hazard assessment, communications cable and pipeline route planning, resource exploration, habitat management, and territorial claims under the Law of the Sea, but also fundamental Earth science questions, such as what controls seafloor shape and how seafloor shape influences global climate [12].

While wide-area, high-resolution ground heights can now be measured quickly with aerial electromagnetic survey technologies such as standard photogrammetry and LIDAR [10], the same is not true with seafloor depths. Good discussions on this issue are available at [12, 13]. In short, the problem lies in the 3000-5000 meters of salty water which masks the penetration of electromagnetic waves.

The altimeter method is an attempt for wide-area coverage. The method exploits the fact that the ocean surface has broad bumps and dips that mimic the topography of the ocean floor. By surveying the shape of the water surface instead, we avoid the need of shooting electromagnetic waves into the water, yet allowing us to deduce how the underlying seafloor looks like. However, features on the ocean floor that are narrower than the average ocean depth of 3-5 kilometers do not produce measurable bumps on the ocean surface.

To obtain high-resolution data, we have no choice but to sail across the ocean, and on the way send out acoustic pulses to the seafloor. From the time it takes the pulses to leave the ship, be reflected by the seafloor and eventually get back the ship again, the ocean depths can be estimated. However, since ships travel slowly, the ocean remains largely uncharted. It has been estimated that 300 years are needed to cover the whole ocean area. With the most popular multibeam bathymetry technique [15], we can collect many data points with 10m resolution in a swath up to 10km wide along a ship's trackline. However, between the tracklines there is no data. In the southern oceans, these survey lines can sometimes be close together, but more often they are hundreds of kilometers apart [14]. Figure 1 demonstrates the spatial distribution of such shipboard data samples in a few 10×10 degree regions.

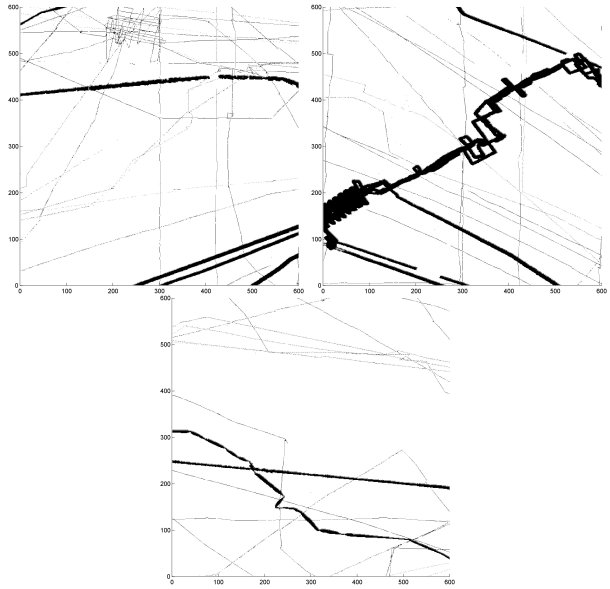

Figure 1: Locations with available depth samples. (Top left) Region 1: 30S-40S, 90E-100E. (Top right) Region 2: 40S-50S, 80E-90E. (Bottom) Region 3: 40S-50S, 90E-100E.

Since a full data grid is assumed for many terrain analyses, we need a *surface reconstruction*. The problem is defined over a spatial domain of dimensions $n \times n$. Available are the measured depth values of $k \ll n^2$ positions (x_1, y_1), (x_2, y_2), $\ldots$, (x_k, y_k), denoted as h_{x_1,y_1}, h_{x_2,y_2}, $\ldots$, h_{x_k,y_k}. The task is to predict the depths for all the $n \times n$ positions in the domain, including those of the k known positions and those of the remaining $n^2 - k$ unknown positions. This assumes for each possible location (i, j), the corresponding predicted depth $z_{i,j}$ is single-valued; caves or overhangs are not allowed.

Acquisition footprint is the major problem of using general reconstruction schemes on those extremely unevenly distributed depth data, even with a few current bathymetry charts such as the one published by National Oceanic and Atmospheric Administration (NOAA) [9]. It refers to the artifact which makes the tracklines visible. The artifact is especially obvious under *shaded relief*, a graphic technique which is often used to highlight terrain surface variations [17]. Figure 2 shows how the reconstructed seafloor surfaces look like under a few such schemes such as Kriging [4] and natural neighbor [16]. While both come up with a surface of a similar general shape, we observe abnormal high-frequency depth fluctuations which are correlated with the few trackline locations. We aim at a surface from which we cannot deduce the trackline locations. Meanwhile, its general shape as portrayed by relatively low frequencies should be preserved.

To remove the artifact, a typical routine is to smooth the reconstructed surface. The difficulty lies in the fact that it is not a simple high-frequency removal problem. Each local region has its unique smoothing requirement. We may not need to smooth certain noisy trackline locations, as long as their neighborhood is equally noisy. In contrast, we cannot accept even small fluctuations along the tracklines if they are surrounded by relatively calm surfaces which make the tracklines stand out. One known attempt, CleanTOPO2 [11], involves post-processing, manually removing the artifacts with repeated applications of a morphological averaging filter. Though easy for humans to decide where filtering should be done and when to stop, the process is subjective and time-consuming. As we are obtaining new trackline depth data from time to time, it may become impractical to do such a manual process every time the database expands. This motivates the automation of artifact-free surface reconstruction.

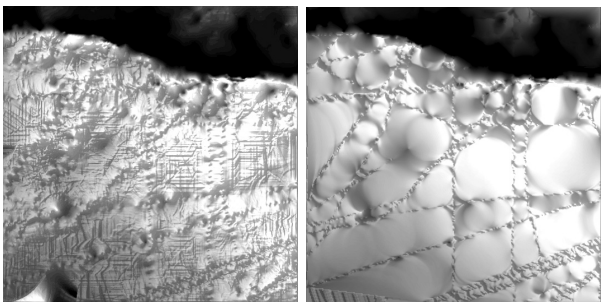

Figure 2: Region 1 reconstructions with (Left) Kriging, (Right) natural neighbor.

Our previous work proposed Overdetermined Laplacian Partial Differential Equation (ODETLAP) as a partial solution. Its original version was designed with smoothing considered right in the reconstruction process, and yields much better reconstruction accuracy than conventional algorithm. At that time, we further improved its accuracy by allowing the smoothing factor to vary according to local smoothing need. Section 2 will give information about this solution.

This paper aims to improve that ODETLAP implementation in two ways. First, we find that quite a few terrain features have been lost in the above variable-smoothing ODETLAP. We fix it by a two-step approach which reconstructs a preliminary surface with ODETLAP of a high-accuracy setting, and then applying a smoothing-biased ODETLAP over that preliminary surface. Second, our previously-reported implementation relied on human to set smoothing parameters. We automate the process based on the bump area distribution. Details will be given in Section 3, before we conclude the paper in Section 4.

2. ODETLAP

We first presented Overdetermined Laplacian Partial Differential Equation (ODETLAP) as a superior above-ground terrain reconstruction and lossy terrain compression [2, 20, 21], since it predicts neighboring terrain heights better than other conventional prediction schemes. Besides, it can work with contour lines (continuous or intermittently broken), infer mountain tops inside a ring of contours, and enforce continuity of slope across contours. All these are favorable features of natural-looking terrains.

Its formulation sets up an overdetermined system $\mathbf{Az} = \mathbf{b}$, as shown in Figure 3, to solve for the depths of the whole seafloor depth grid $\mathbf{z}$ in the bathymetry case presented here. The system includes an exact equation for each of the k known-depth positions. That equation aims to set the depth value

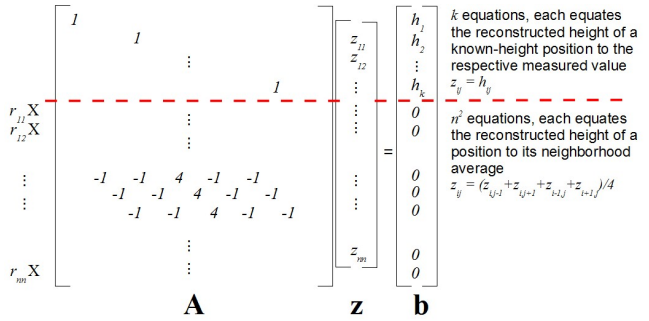

Figure 3: ODETLAP.

of the respective position to its known value. The system also contains an averaging equation for all n^2 positions. That equation attempts to regularize the respective depth to the average of its immediate four neighbors. Through adjusting the weights $r_{i,j}$ of averaging equations, we can change how the errors are distributed over the equations and hence obtain terrain surfaces of the desired accuracy-smoothness tradeoff.

If the r values are low (e.g. 1), the system is *accuracy-centric*. The reconstructed values of the known-depth locations will be close to the measured values. However, acquisition footprint appears extensively as vigorous fluctuations along the tracklines, as shown in Figure 4, top left. Note that in areas with no data, the surface is relatively plain. This is superior to the surfaces done with a few conventional reconstruction schemes mentioned earlier, in which the artifact also affects those regions. To alleviate such artifact, we can use the *smoothing-centric* version in which the r values are high (e.g. 50). Accuracy at known-depth locations (especially those with exceptional values) are sacrificed for a smooth surface as implied by the averaging equations, as shown in Figure 4, middle left.

In the original version of ODETLAP, we set all $r_{i,j}$ to the same value. In our previous paper [5], we demonstrated how unnecessary error budget could be saved by adjusting the $r_{i,j}$ values for individual locations. We observed that the trackline locations were usually of high data density and hence required relatively higher smoothing. We asked the users to specify the smoothing factors for locations with lowest local sample density and locations with highest sample density, and then allowed locations with intermediate data density to vary between these two values. As a result, we achieved a surface of similar smoothness with a smaller error budget. Figure 4, bottom left, shows a typical reconstruction results with such a variable-smoothing ODETLAP system.

The time complexity of ODETLAP is $O(n^3 + k)$. In practice, we transform the system to $\mathbf{A^T A z} = \mathbf{A^T b}$ before solving for $\mathbf{z}$. In this equivalent system $\mathbf{A'z} = \mathbf{b'}$ where $\mathbf{A'} = \mathbf{A^T A}$ and $\mathbf{b'} = \mathbf{A^T b}$, $\mathbf{A'}$ is *symmetric positive definite*. We can then take advantage of the fast Cholesky factorization to keep the actual solving time to within seconds even for large datasets [7]. Note that the matrix A is indeed a $n^2 \times n^2$ sparse matrix because the number of non-zero entries in each row is upper bounded by the number of possible immediate neighbors, which is 4. With the recent advances of graphical display units (GPU) in solving sparse linear systems [1, 6, 8], this

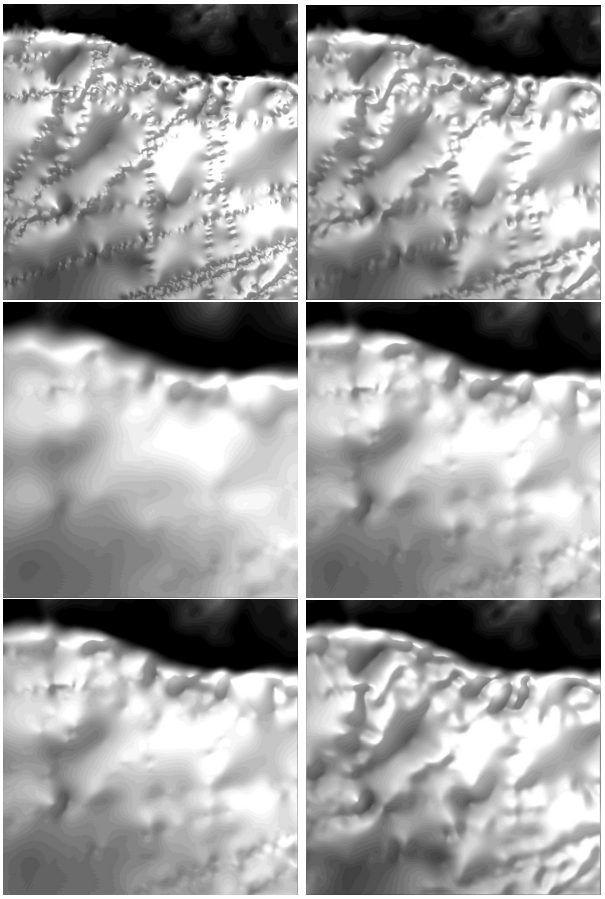

Figure 4: [Left] Region 1 reconstructions with ODETLAP: (Top) accuracy-centric. (Middle) smoothing-centric. (Bottom) variable smoothness factors. [Right] Region 1 reconstructed seafloor surfaces with mean error around 130m. (Top) a mean filter with a square structure element of length 10, (Middle) variable-smoothness ODETLAP, (Bottom) two-step ODETLAP.

approach may offer even faster and more efficient solutions to large data grids.

3. AUTOMATED TWO-STEP ODETLAP

Figure 5 gives the flow diagram of the two-step ODETLAP as our solution of the automated artifact-free seafloor surface reconstruction problem. Below we will describe the algorithm in terms of its two feature characteristics: *terrain feature preservation* and *automated smooth factor determination*.

3.1 Terrain feature preservation

Our new algorithm first reconstructs a highly-accurate preliminary surface from the available depth data, and then smoothes that preliminary surface.

In the first step which reconstructs a preliminary seafloor surface, we choose ODETLAP since this reconstruction scheme works best in deducing the missing seafloor depth data, as described in Section 2. In case the resultant surface needs no further smoothing, we are likely to have obtained the most

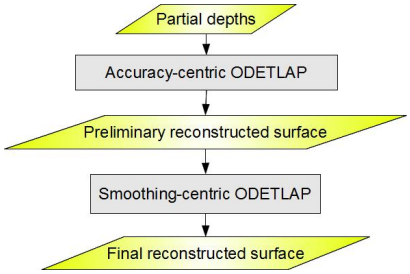

Figure 5: Two-step ODETLAP workflow.

SRTM1	Natural neighbor	ODETLAP
W121N38 1201:1600,1201:1600	31m (413m)	24m (292m)
W121N38 2800:3200, 801:1200	20m (309m)	18m (242m)
W121N38 3201:3600, 401:800	17m (299m)	15m (164m)
W111N31 401:800, 1:400	4m (93m)	3m (79m)
W111N31 401:800, 401:800	14m (173m)	11m (120m)
W111N31 401:800, 801:1200	5m (134m)	4m (134m)

Table 1: Mean errors and maximum errors (in parentheses) of the terrain surfaces reconstructed using height samples on the tracklines of Region 1.

accurate surface. We set the weighting between the exact equations and the averaging equations to 1:1. Beyond that point, increasing the weightings of the exact equations does not change the accuracy of the preliminary terrain too much.

To investigate the effect of switching the terrain reconstruction scheme in the first step from ODETLAP to the others, we first remove the height values of a few full terrains except those falling on the tracklines shown in Figure 1. (Note we do not use seafloor surfaces as no ground truth is available for error comparison.) Then we reconstruct the terrain with different techniques and compute the errors with respect to the ground truth. ODETLAP does the best in guessing the missing heights, as reflected by the generally lower mean errors and maximum errors shown in Table 1. This result supports our choice of ODETLAP even if we are now working on extremely unevenly distributed depth samples.

In the second step which smoothes the preliminary surface, we once again pick ODETLAP since this scheme provides better smoothing of the artifacts than others under the same error budget. Figure 4, right, compares the smoothing results with a mean filter (similar to the one used in CleanTOPO2 mentioned in Section 2) and ODETLAP under similar error budget. While the acquisition artifact is almost gone with ODETLAP smoothing, it is not the case with the mean filter. Table 2 shows the respective mean error budgets needed by the sample seafloors to achieve the respective optimal smoothing levels (using the metric defined in the next subsection). Our ODETLAP smoothing scheme just needs half of the error budget as the average filter counterpart. The results above demonstrate the capability of our scheme in distributing the limited error budget to smoothing locations.

Region	ODETLAP	Mean filter
Region 1	130m	297m
Region 2	75m	178m
Region 3	214m	360m

Table 2: Mean error budgets to reach optimal smoothing levels.

The only variable of this algorithm is the smoothing factor in the second step. Figures 6–8, left, show how the surface varies as smoothing increases. A higher smoothing factor means a higher mean error but at the same time better smoothing-out of the small bumps along the tracklines. When compared with our original implementation which requires the specification of the lower and upper smoothing factors, we now have one fewer degree of freedom, make it easier to adjust. Also note that on raising the error budget, the high-frequency acquisition footprint is gone before those lower-frequency terrain features which account for the general terrain shape. Such a smoothing priority makes the scheme superior over our previously-proposed variable-smoothing ODETLAP. In the surface reconstructed with variable-smoothing ODETLAP such as the one in Figure 4 middle right, even though the acquisition footprint is also almost gone, we also lose quite a few terrain features.

3.2 Automated smoothing determination

The above two-step procedure features a single parameter controlling the smoothing of the final reconstructed seafloor surface. To allow automated determination of the optimal smoothing level, first we need to convert the acquisition footprint to a form that is recognizable by computers. After several experiments, we find that the graph of Gaussian curvature may be used. As shown on the right hand side of Figures 6–8, the artifact appears as relatively small patches of positive Gaussian curvature concentrated at the trackline locations. In fact, patches of Gaussian curvature is regarded as a view-independent indicator of regions with potential shaded relief in a few other research [3].

As observed from the same set of figures, smoothing helps enlarge those patches or simply remove them. With a small smoothing factor, along the tracklines we have a huge number of small such patches. This is in contrast with regions with no data where there are few, much bigger patches. On increasing smoothing, mean error at known-height location increases. Meanwhile, the bumps along the tracklines become fewer and bigger, while those in no-data regions hardly change in terms of both size and quantity. This leads to a drop of the variations among the patch areas.

To utilize the above phenomenon in automated smoothing factor determination, we first apply a morphological erosion [19] with a 3-pixel-width square component on the positive Gaussian curvature graph to help discriminate the patches. We then compute the coefficient of variation c_v, which is a normalized measure of dispersion of a probability distribution [18], of the patch areas. The acquisition footprint is found to be almost gone when that coefficient drops to around 1.3. Figures 6–8, bottom, correspond to a smoothing level with

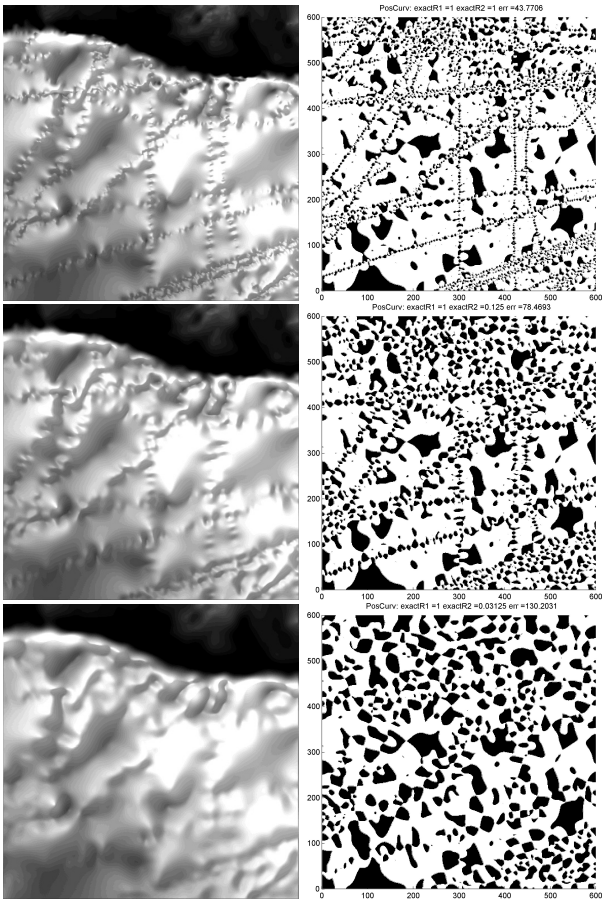

Figure 6: (Left) Region 1 reconstructed seafloor surfaces. (Right) Respective positive Gaussian curvature locations. The mean errors, from top to bottom, are around 44m, 78m and 130m respectively.

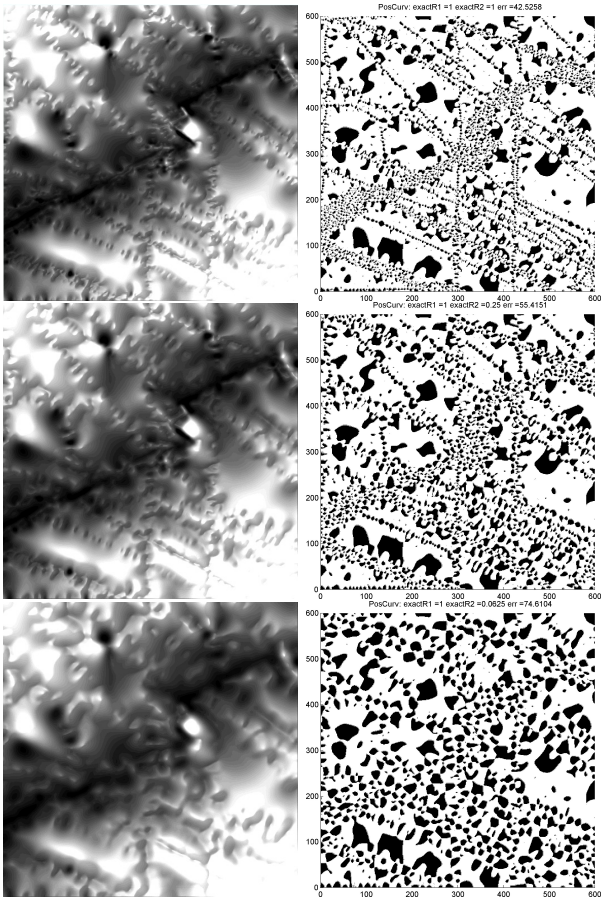

Figure 7: (Left) Region 2 reconstructed seafloor surfaces. (Right) Respective positive Gaussian curvature locations. The mean errors, from top to bottom, are around 43m, 55m and 75m respectively.

around that coefficient value. Note that different datasets may need different error budgets to remove the artifacts. For example, while Region 3 requires a mean error budget as high as 214m to have the artifact removed, Region 2 needs 75m only. Using that coefficient as a gauge helps reduce unnecessary smoothing and hence reduce the errors needed to achieve an artifact-free surface.

4. CONCLUSION

We have presented an improved ODETLAP procedure for the automated reconstruction of artifact-free seafloor surfaces within a limited error budget. It has the smoothing factor as its only parameter, making it easier to adjust than our previous attempt which requires two parameter inputs. By smoothing an accurate reconstruction with ODETLAP, we allow terrain features to be better preserved than our previously-reported scheme that reconstructs from the given measured values directly. To automate the adjustment of the smoothness parameter, we analyze the Gaussian curvature of the reconstructed surfaces. Areas of positive Gaussian curvatures highly resemble the locations of the bumps that we observe on the shaded reconstructed surface. Increasing smoothing enlarges the small bumps and reduces their numbers along the tracklines, thus alleviating the acquisition footprint. When the coefficient of variation of such areas is around 1.3, the artifact is almost gone. We use this observation to determine the minimum smoothing needed for the artifact-free surface.

In the future, we will look into the automatic stopping criterion further. More tests will be done on a variety of trackline depth samples. Even more accurate stopping criteria will be investigated. Our current work embraces data from the tracklines but not the altimeter. It is interesting to see how data from different sources could combine.

This research was partially supported by NSF grants CMMI-0835762 and IIS-1117277.

5. REFERENCES

[1] N. Bell and M. Garland. Implementing sparse matrix-vector multiplication on throughput-oriented processors. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC '09, pages 18:1–18:11, New York, NY, USA, 2009. ACM.

[2] M. B. Gousie and W. R. Franklin. Augmenting grid-based contours to improve thin plate DEM

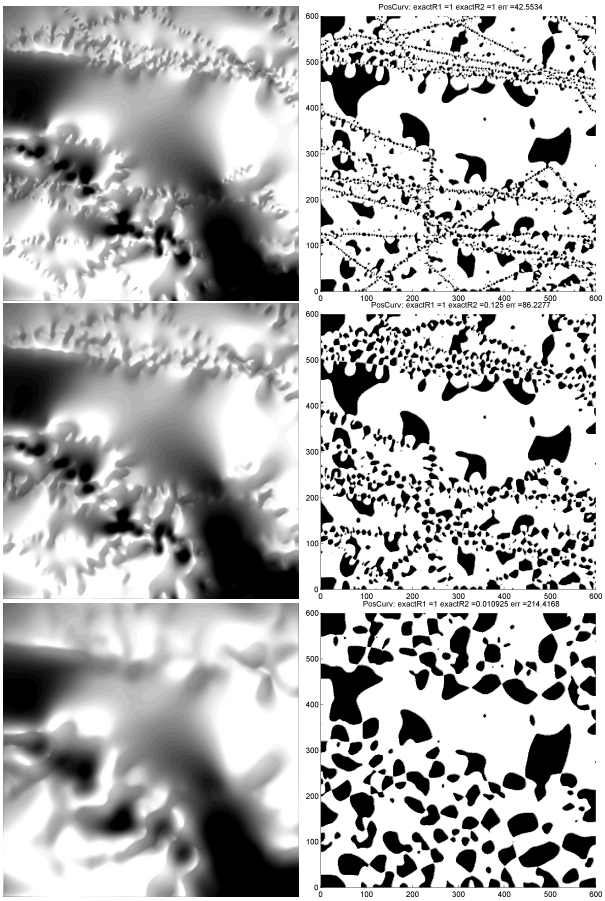

Figure 8: (Left) Region 3 reconstructed seafloor surfaces. (Right) Respective positive Gaussian curvature locations. The mean errors, from top to bottom, are around 43m, 86m and 214m respectively.

generation. *Photogrammetric Engineering & Remote Sensing*, 71(1):69–79, 2005.

[3] Y. Iwahori, S. Fukui, C. Fujitani, Y. Adachi, and R. J. Woodham. Relative magnitude of gaussian curvature from shading images using neural network. In *Proceedings of the 9th international conference on Knowledge-Based Intelligent Information and Engineering Systems - Volume Part I*, KES'05, pages 813–819, Berlin, Heidelberg, 2005. Springer-Verlag.

[4] D. G. Krige. A statistical approach to some mine valuations and allied problems at the Witwatersrand. Master's thesis, University of Witwatersrand, 1951.

[5] T.-Y. Lau, Y. Li, Z. Xie, and W. R. Franklin. Sea floor bathymetry trackline surface fitting without visible artifacts using ODETLAP. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '09, pages 508–511, New York, NY, USA, 2009. ACM.

[6] Y. Li. *CUDA-accelerated HD-ODETLAP: a High Dimensional Geospatial Data Compression Framework*. PhD thesis, Rensselaer Polytechnic Institute, 2011.

[7] Y. Li, T. Y. Lau, C. S. Stuetzle, P. Fox, and W. R. Franklin. 3D oceanographic data compression using 3D-ODETLAP. In *18th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2010.

[8] D. Luebke, M. Harris, N. Govindaraju, A. Lefohn, M. Houston, J. Owens, M. Segal, M. Papakipos, and I. Buck. GPGPU: general-purpose computation on graphics hardware. In *Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, SC '06, New York, NY, USA, 2006. ACM.

[9] National Oceanic and Atmospheric Administration. Bathymetry data viewer. `http://maps.ngdc.noaa.gov/viewers/bathymetry/`, (retrieved 7/2/2012), 2012.

[10] National Oceanic and Atmospheric Administration Coastal Services Center. LiDAR 101: An introduction to LiDAR technology, data, and applications. `http://www.csc.noaa.gov/digitalcoast/data/coastallidar/_pdf/What_is_Lidar.pdf`, (retrieved 2/18/2011), 2008.

[11] T. Patterson. CleanTOPO2: Edited SRTM30 plus World Elevation Data. `http://www.shadedrelief.com/cleantopo2/`, (retrieved 7/2/2012), 2007.

[12] D. T. Sandwell. Ocean Bumps and Dips. *The World & I*, 3:252–255, 1992.

[13] W. H. Smith and D. T. Sandwell. Global sea floor topography from satellite altimetry and ship depth soundings. *Science Magazine*, 277, issue 5334, 1997.

[14] W. H. Smith and D. T. Sandwell. Predicted seafloor topography: NGDC data announcement number: 94-MGG-04. `http://www.ngdc.noaa.gov/mgg/fliers/94mgg04.html`, (retrieved 6/8/2009), July 2008.

[15] University of Rhode Island, Office of Marine Programs. Discovery of sound in the sea: Echo sounder-multibeam. `http://www.dosits.org/gallery/tech/osf/esm1.html`, (retrieved 6/8/2009), Mar. 2008.

[16] Wikipedia. Natural neighbor. `http://en.wikipedia.org/wiki/Natural_neighbor`, (retrieved Dec 1, 2010), 2010.

[17] Wikipedia. Cartographic relief depiction. `http://en.wikipedia.org/wiki/Cartographic_relief_depiction`, (retrieved 7/2/2012), 2012.

[18] Wikipedia. Coefficient of variation. `http://en.wikipedia.org/wiki/Coefficient_of_variation`, (retrieved 7/2/2012), 2012.

[19] Wikipedia. Erosion (morphology). `http://en.wikipedia.org/wiki/Erosion_%28morphology%29`, (retrieved 7/2/2012), 2012.

[20] Z. Xie, M. A. Andrade, W. R. Franklin, B. Cutler, M. Inanc, D. M. Tracy, and J. Muckell. Approximating terrain with over-determined Laplacian PDEs. In *17th Fall Workshop on Computational Geometry*, 2–3 Nov 2007. poster session, no formal proceedings.

[21] Z. Xie, W. R. Franklin, B. Cutler, M. A. Andrade, M. Inanc, and D. M. Tracy. Surface compression using over-determined Laplacian approximation. In *Proceedings of SPIE Vol. 6697 Advanced Signal Processing Algorithms, Architectures, and Implementations XVII*, San Diego CA, 27 August 2007. International Society for Optical Engineering. paper 6697-15.