

Joint Bilateral Mesh Denoising using Color Information and Local Anti-Shrinking

Oliver Wasenmüller
oliver.wasenmueller@dfki.de

Gabriele Bleser
gabriele.bleser@dfki.de

Didier Stricker
didier.stricker@dfki.de

DFKI GmbH - German Research Center for Artificial Intelligence
Augmented Vision Department
Trippstadter Str. 122
67663 Kaiserslautern, Germany

ABSTRACT

Recent 3D reconstruction algorithms are able to generate colored meshes with high resolution details of given objects. However, due to several reasons the reconstructions still contain some noise. In this paper we propose the new Joint Bilateral Mesh Denoising (JBMD), which is an anisotropic filter for highly precise and smooth mesh denoising. Compared to state of the art algorithms it uses color information as an additional constraint for denoising; following the observation that geometry and color changes often coincide. We face the well-known mesh shrinking problem by a new local anti-shrinking, leading to precise edge preservation. In addition we use an increasing smoothing sensitivity for higher numbers of iterations. We show in our evaluation with three different categories of testdata that our contributions lead to high precision results, which outperform competing algorithms. Furthermore, our JBMD algorithm converges on a minimal error level for higher numbers of iterations.

Keywords

Mesh Denoising, Smoothing, Fairing, Joint Bilateral Filter, Local Anti-Shrinking, Color Information.

1 INTRODUCTION

Many applications, such as urban planning, industrial measurement or human anthropometry, require reconstructed 3D models of the respective objects with very high precision. The traditional approach is to acquire these models by laser scanners, since they promise high quality results. However, they are expensive, impractical to use and contain still some noise. Meanwhile 3D reconstruction algorithms, such as e.g. [Agi, FP10, NIH⁺11, SSC14], are able to generate colored mesh models from devices like standard color cameras and/or depth cameras, which are widely spread, cheap and easy to use. These reconstructions contain color information together with high-resolution details, but also suffer from noise in their 3D geometry. To get rid of this noise, several (mostly iterative) methods for mesh denoising were proposed in the literature. Sometimes they are also referred to as smoothing, filtering or fairing methods. They use directly the 3D geometry or derived measures, like distances or normals of

the mesh, to estimate new vertex positions. However, none of them explicitly uses the color information provided by e.g. one of the above mentioned algorithms. Thus, we present in this paper a new anisotropic method called Joint Bilateral Mesh Denoising (JBMD), which uses - besides geometric information - the color information as an additional constraint for edge preserving denoising.

Another well-known problem of mesh denoising are shrinking effects. They occur mostly in curves regions of the mesh and are caused by homogeneous shifts of vertices in a neighborhood into one major direction. Current approaches try to compensate that ef-




Figure 1: Comparison of different mesh denoising algorithms for the *fandisk* mesh. Top row: meshes. Bottom row: color-coded error distribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

fect, whereas they focus more on visually appealing results than on precision. Thus, in our JBMD algorithm we propose a new approach to avoid this effect by a precise local anti-shrinking. Furthermore, many current algorithms suffer from their high dependence on the number of iterations. Therefore our new algorithm increases the denoising sensitivity per iteration leading to constantly low errors. A short overview of our algorithm is also given in the following video: <https://youtu.be/odm8kr2rKPA>

Summarizing, the main contributions of our proposed JBMD algorithm are:

- explicit usage of color information as an addition constraint for denoising,
- precise local anti-shrinking and
- increasing denoising sensitivity.

The remainder of this paper is organized as follows: Section 2 gives an overview of existing methods for mesh denoising. The proposed JBMD algorithm is motivated and explained in detail in Section 3, while it is evaluated regarding precision in Section 4. The work is concluded in Section 5.

2 RELATED WORK

Image and mesh denoising is an ongoing research topic in the image processing and computer vision community.

Often, mesh denoising methods are related to image denoising approaches. State of the art approaches in image denoising include methods such as anisotropic diffusion [PM90], total variation [ROF92], wavelet denoising [Don95], robust diffusion [BSMH98], bilateral filter [TM98] and joint bilateral filter [KCLU07, HSJS08]. In particular the joint bilateral filter uses, similar to our new approach, the color information as an additional constraint.

State of the art algorithms for mesh denoising are amongst others Laplacian [Fie88, Tau95, VMM99] and bilateral [FDCO03, JDD03, ZFAT11] mesh denoising (BMD) methods. Laplacian mesh denoising is an iterative isotropic procedure, where the new vertex positions are directly calculated from the positions of the neighboring vertices. In contrast, bilateral mesh denoising is an iterative edge preserving anisotropic approach. New vertex positions are estimated from the vertex's neighborhood, where the influence of neighboring vertices depends on their distance and on their offset to the tangent plane. Parts of this approach are also used for our new algorithm.

A general and well-known problem of mesh denoising is that the mesh shrinks in convex regions with each application of the particular algorithm, which is a huge

problem especially for iterative approaches. [Tau95] solves this problem by alternating shrinking and expansion steps. Admittedly the precision of this approach depends heavily on the geometry of the particular mesh [DMSB99]. Another common approach, which is e.g. used in the Bilateral Mesh Denoising [FDCO03], is to preserve the volume of the mesh by a global correction step as proposed in [DMSB99]. The algorithm estimates the volume V^n of a mesh after the n-th iteration by the sum of volumes of all ordered pyramids centered at the origin and with a triangle of the mesh as base. Each vertex of the mesh is then scaled by the factor β , which is defined by

$$\beta = \left(\frac{V^0}{V^n} \right)^{\frac{1}{3}}, \quad (1)$$

to achieve the original volume V^0 . However, as mesh shrinkage occurs only in convex regions contrary to flat regions, a global correction has indeed appealing effects but is not precise. Thus, we propose in this paper a precise local shrinkage correction.

3 METHOD

In this paper we propose the Joint Bilateral Mesh Denoising (JBMD), which is a filtering method for meshes using local neighborhoods. The method can be subdivided into two parts: the denoising itself and the subsequent local anti-shrinking.

The main idea of our new denoising algorithm is related to image processing, namely motivated through the Joint Bilateral Filter (JBF) [KCLU07]. This anisotropic edge-preserving filter is often used to denoise depth images by using color images as additional constraints. The main idea is to compute a new depth value as a weighted average of surrounding depth values, where the weights depend on their deviation in position (space) and color value (range). The assumption of JBF are coherent depth and color discontinuity, meaning that edges in the color image coincide with edges in the depth image and vice versa. This coherence assumption was validated in many image processing publications [KCLU07, WBS15] and we show in Section 4 that it also holds for meshes.

The intention of our new local shrinkage correction is - contrary to alternating [Tau95] or global [DMSB99] correction - to adjust vertex positions only where shrinkage effects occurred. This effect arises only in convex regions, whereas flat regions are not affected. We observed that the weighted mean signed shift of vertices in the neighborhood, which were estimated by the denoising in the first step, equalize in noisy flat regions, whereas in convex (and thus shrunk) regions it is a precise local measure for a shrinkage correction.

Like in many other mesh denoising algorithms [FDCO03, ZFAT11] we estimate in our JBMD new

vertex positions v'' in a mesh by shifting along the normal direction n . This has the positive effect that irregularities in the resulting mesh are avoided. Our two-step algorithm can be described by

$$v'' = v + (x' \cdot n) - (x'' \cdot n), \quad (2)$$

where $(x' \cdot n)$ is the denoising part, $(x'' \cdot n)$ the correction part and $x' | x''$ refer to the magnitude of the shift. The algorithm is illustrated in detail in Figure 2 and defined in the following:

Algorithm 1 Joint Bilateral Mesh Denoising (JBMD) for the m -th iteration

Require: Vertex v , Normal n

```

1:  $\{q_i\} = \text{neighborhood}(v)$ 
2:  $sum, norm, sum', norm' = 0$ 
3: for all  $i$  do
4:    $d_i = \|v - q_i\|$ 
5:    $o_i = \langle n, v - q_i \rangle$ 
6:    $c_i = (v.r - q_i.r)^2 + (v.g - q_i.g)^2 + (v.b - q_i.b)^2$ 
7:    $w_i^d = \exp(-d_i^2 / 2\sigma_d^2)$ 
8:    $w_i^o = \exp(-o_i^2 / 2(\sigma_o \cdot \lambda^m)^2)$ 
9:    $w_i^c = \exp(-c_i / 2\sigma_c^2)$ 
10:   $sum += (w_i^d \cdot w_i^o \cdot w_i^c) \cdot o_i$ 
11:   $norm += w_i^d \cdot w_i^o \cdot w_i^c$ 
12: end for
13:  $x' = sum / norm$ 
14:  $v' = v + x' \cdot n$ 
15: for all  $i$  do
16:    $d'_i = \|v' - q'_i\|$ 
17:    $w_i^{d'} = \exp(-(d'_i)^2 / 2\sigma_d^2)$ 
18:    $sum' += w_i^{d'} \cdot x'_i$ 
19:    $norm' += w_i^{d'}$ 
20: end for
21:  $x'' = sum' / norm'$ 
22:  $v'' = v' - x'' \cdot n$ 
23: return  $v''$ 

```

In the first step of our algorithm (line 3-14) the new position v' of a vertex v is estimated as a weighted average of neighboring vertex position q_i , where the weights depend on three influencing factors: distance (w_i^d), offset (w_i^o) and color difference (w_i^c). For computing the distance d_i between a vertex v and neighboring vertex q_i , the geodesic distance on the smooth surface would be the correct measure. However, for efficiency reasons we approximate d_i using the Euclidean distance in line 4, since [FDCO03] demonstrated already a sufficient impact. The offset o_i is defined as the distance of vertex q_i to the tangent plane of vertex v . The intention of using this offset o_i is that neighboring points in flat regions should have a higher influence than in convex or edge regions. As described in line 5, o_i can be easily estimated using the dot product. The last influence factor is the color difference c_i between a vertex q_i and




Figure 2: Joint Bilateral Mesh Denoising (JBMD) applied to vertex v : (a) Denoising step ($v \rightarrow v'$). (b) Shrinking correction step ($v' \rightarrow v''$). (c) Final result.

vertex v , which is estimated in line 6. To map the influence factors d_i , o_i and c_i to weights w_i^d , w_i^o and w_i^c , we use the Gaussians of lines 7-9. The final shift x' of vertex v is the normalized weighted sum of offsets o_i of neighboring vertices q_i .

In the second step of our algorithm (line 15-22) we correct the position of a vertex v' due to possible shrinking effects. As already mentioned before, we observed that the weighted mean signed shift x'' of vertices in the neighborhood, which were estimated by the denoising in the first step, is a good local measure for a shrinkage correction. For the estimation of the weights for x'' we use the distance d'_i between a vertex q'_i and v' together with a mapping function (line 16-17) similar to the first step of the algorithm. The weighted mean signed shift x'' is calculated by summing up the weighted signed shifts x'_i (line 18) and normalizing afterwards (line 21).




Figure 3: Mean errors of the *fandisk* mesh for different numbers of iterations. Quantitative comparison of different features of our new JBMD algorithm and comparison against BMD.

Dark blue: BMD [FDCO03]. Red: BMD without volume preservation. Green: Our new JBMD. Purple: JBMD without increasing smoothing sensitivity. Light blue: JBMD without local anti-shrinking. Orange: JBMD without increasing smoothing sensitivity and without local anti-shrinking.

The described JBMD algorithm is applied locally for each vertex of the mesh. However, vertices at a boundary of a mesh do not have a well defined neighborhood (line 1). In our algorithm we define the size of the neighborhood as a fixed number k . The neighborhood of a vertex v is then defined by the k closest vertices q_i . Obviously, the shape of our neighborhood changes from vertex to vertex, but since the distance d_i between vertices v and q_i is an influencing parameter, this artifact has negligible influence.

Our JBMD algorithm is - like many other [FDCO03, ZFAT11] - an iterative approach. In the first iteration major noise is eliminated, whereas with higher number of iterations the overall level of noise decreases. Thus, we consider this aspect by an increasing smoothing sensitivity. In our JBMD algorithm the noise influences the result via the offset o_i , whereas the corresponding mapping function depends on σ_o . Therefore, we decrease the parameter σ_o by a constant factor λ with each iteration; leading to constantly low error.

4 EVALUATION

In this section we benchmark our JBMD algorithm by comparing it to competing algorithms, namely Laplacian Denoising and Bilateral Mesh Denoising (BMD). These algorithms are described in more detail in Section 2. All methods - including ours - depend on some parameters. For our JBMD algorithm these are σ_d , σ_o , σ_c and λ . Thus, we run each algorithm with a huge number of possible parameter combinations to detect the optimal setting. All results (Figures, diagrams, etc.) shown in this paper are generated with optimal parameter settings and numbers of iterations. For our JBMD algorithm we used the parameter settings of Table 1 for the given datasets. Note, these parameters depend highly on the mean vertex distance (MVD) of the given mesh. According to our experiments $MVD \approx 2\sigma_d \approx 4\sigma_o$ can be used as a rough guideline for setting the parameters.

Unfortunately, a groundtruth comparison on real world data is very difficult, since no datasets are available, which provide both real noisy data and real denoised data. Thus, in the recent literature it is common to use precise models of an object and generate the noise on it synthetically. For this paper we decided to use three categories of testdata in our evaluation.

The first category are colored synthetic meshes with sharp edges together with an artificially noisy version of this mesh. We use here the well-known *fandisk* mesh, where each part of the surface has another color. Furthermore, we add a Gaussian noise, where the standard deviation is roughly half of the vertex distance. The second category of testdata are highly precise reconstructions of real objects acquired by a camera-projector-system [KNRS13]. We also added here a Gaussian noise with a standard deviation of approximately half vertex distance. The meshes used in this paper are the *lion* and *allegorie* reconstructions. The third category of testdata are reconstructed meshes, which are generated by standard cameras and Agisoft PhotoScan [Agi]. We use in this paper the heads of two persons: *person 1* and *person 2*. These reconstructions include partially strong noise due to the lack of characteristic features. Obviously, for these reconstructions no ground truth is available, but they are a real world scenario, where the application of a mesh denoising algorithm is required.

	MVD	σ_d	σ_o	σ_c	λ
<i>fandisk</i>	0.1897	0.1	0.05	30	0.8
<i>lion</i>	0.4401	1.0	0.5	30	0.6
<i>allegorie</i>	0.0003	0.0008	0.0004	35	0.7
<i>person 1</i>	0.0030	0.01	0.005	20	0.7
<i>person 2</i>	0.0041	0.01	0.005	20	0.7

Table 1: Parameter settings in the evaluation of our new JBMD algorithm for the given datasets with specified mean vertex distance (MVD).




Figure 4: Comparison of different mesh denoising algorithms for the *lion* mesh. Top row: meshes. Bottom row: color-coded error distribution.




Figure 5: Mean errors of the *lion* mesh for different numbers of iterations. Quantitative comparison of different features of our new JBMD algorithm and comparison against BMD.

As a quality measure for denoising we use the mean error (ME), which is defined for a denoised mesh X and a ground truth mesh G by

$$ME(X, G) = \frac{1}{n} \sum_{i=1}^n \|x_i - g_p\| \quad (3)$$

$$x_i \in X; g_p \in G; \forall i p = \arg \min_p \|x_i - g_p\|,$$

where n is the number of vertices in the mesh X . More intuitive, it is defined by the mean distance of each vertex in the mesh X to the respective closest vertex in the ground truth mesh G .

First of all we evaluate our algorithm with the *fandisk* mesh in Figure 1, which has ideal preconditions for our JBMD, since all sharp edges coincide with color changes. From a visual point of view, all three mesh denoising algorithms provide smooth results without visible noise. However, they differ strongly in their precision, as visible in the color-coded error distribution in the bottom row. The blue color indicates low errors, whereas red represents high errors. Both BMD and Laplacian denoising have imprecise vertices at the

edges of the mesh, whereas our JBMD has only some minor inaccuracy. Figure 3 depicts the mean error of the *fandisk* mesh depending on the number of iterations. Our JBMD has the lowest error and converges in particular on this low error level. If our JBMD is used without the increasing smoothing sensitivity, the mean error increases again from the fourth iteration on. If we switch off our shrinking correction, we achieve better results for a small number of iterations. This is caused by the inhibiting effect of the shrinking correction, since it reverts the denoising to some extent. However, for larger numbers of iterations superior results can be achieved with our new local anti-shrinking. Looking at the effects of using the color information as an additional parameter, we see that our JBMD has - even without increasing smoothing sensitivity and without local anti-shrinking (orange line) - always a lower mean error than the BMD.

The *lion* and *allegorie* meshes, which correspond to the testdata category of real reconstructions with synthetic noise, are depicted in Figure 4 and 7 respectively. Again, from a visual point of view all three mesh de-




Figure 6: Comparison of different mesh denoising algorithms for the *person 1* mesh. Top row: meshes. Bottom row: color-coded comparison against the input mesh.

noising algorithms deliver smooth results. However, for both datasets at the edges of the mesh BMD and Laplacian denoising are less precise than our JBMD. Figure 5 and Figure 8 depict the respective mean error of the *lion* and *allegorie* meshes depending on the number of iterations. Again, our JBMD outperforms the competing algorithms and converges at the lowest error level.

The *person 1* and *person 2* meshes correspond to the testdata category of reconstructions with real world noise. Since no groundtruth data is available for these datasets, we compare the denoised meshes against the original mesh in Figure 6 and 9 respectively. The BMD algorithms results in the biggest differences to the original mesh. Especially the nose, but also the eyebrows and mouth, have a huge deviation and are not precise. The Laplacian smoothing shows less deviation, but is by far not as close to the original mesh as our JBMD. Of course, smaller deviations to the original do not mandatory result in a better quality, but from a visual point of view all results are similarly smooth. Thus, also for this category of testdata our JBMD outperforms the competing algorithms.

Summarizing the evaluation results, we found out that our JBMD algorithm outperforms competing algorithms for all tested datasets in terms of precision while creating smooth results. Notably is in particular that our JBMD converges on the lowest error level for higher numbers of iterations. This is the achievement of all three main contributions of our paper: Using color information as additional constraint, correcting

shrinking effects locally and increasing the smoothing sensitivity with each iteration. Like illustrated in Figure 3, 5 and 8 this is only possible with the combination of all these three contributions. As long as at least one of them is not activated, the mean error is not minimal and does not converge. With the local anti-shrinking it is possible to denoise especially edges very precisely. Furthermore, we verified with our convincing result that the coherence assumption of coinciding geometry and color changes holds also for meshes.

5 CONCLUSION

In this paper we proposed the new Joint Bilateral Mesh Denoising (JBMD), which is an anisotropic filter for highly precise and smooth mesh denoising. Under the assumption of coinciding geometry and color changes it uses color information as an additional constraint for denoising. This assumption is adapted from the Joint Bilateral Filter (JBF) of the recent image processing research and we showed in this paper that this coherence assumption also holds for meshes. Furthermore, we proposed a precise local anti-shrinking, which leads to precision improvements especially at the edges of the mesh. Our third contribution increases the smoothing sensitivity for higher numbers of iterations. In our evaluation we compared our new JBMD algorithm against competing algorithm based on three categories of test data. We showed that our contributions lead to high precision results with lowest errors. In addition our algorithm converges to the minimum error level for higher numbers of iterations.




Figure 7: Comparison of different mesh denoising algorithms for the *allegorie* mesh. Top row: meshes. Bottom row: color-coded error distribution.




Figure 8: Mean errors of the *allegorie* mesh for different numbers of iterations. Quantitative comparison of different features of our new JBMD algorithm and comparison against BMD.




Figure 9: Comparison of different mesh denoising algorithms for the *person 2* mesh. Top row: meshes. Bottom row: color-coded comparison against the input mesh.

ACKNOWLEDGEMENTS

This work was partially funded by the Federal Ministry of Education and Research (Germany) in the context of the project DENSITY (01IW12001).

6 REFERENCES

- [Agi] Agisoft. PhotoScan. <http://www.agisoft.com/>.
- [BSMH98] Michael J Black, Guillermo Sapiro, David H Marimont, and David Heeger. Robust anisotropic diffusion. *Image Processing, IEEE Transactions on*, 7(3):421–432, 1998.
- [DMSB99] Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 317–324. ACM Press/Addison-Wesley Publishing Co., 1999.
- [Don95] David L Donoho. De-noising by soft-thresholding. *Information Theory, IEEE Transactions on*, 41(3):613–627, 1995.
- [FDCO03] Shachar Fleishman, Iddo Drori, and Daniel Cohen-Or. Bilateral mesh denoising. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 950–953. ACM, 2003.
- [Fie88] David A Field. Laplacian smoothing and delaunay triangulations. *Communications in applied numerical methods*, 4(6):709–712, 1988.
- [FP10] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(8):1362–1376, 2010.
- [HSJS08] Benjamin Huhle, Timo Schairer, Philipp Jenke, and Wolfgang Straßer. Robust non-local denoising of colored depth data. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–7. IEEE, 2008.
- [JDD03] Thouis R Jones, Frédo Durand, and Mathieu Desbrun. Non-iterative, feature-preserving mesh smoothing. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 943–949. ACM, 2003.
- [KCLU07] Johannes Kopf, Michael F Cohen, Dani Lischinski, and Matt Uyttendaele. Joint bilateral upsampling. In *ACM Transactions on Graphics (TOG)*, volume 26, page 96. ACM, 2007.
- [KNRS13] Johannes Kohler, Tobias Noll, Gerd Reis, and Didier Stricker. A full-spherical device for simultaneous geometry and reflectance acquisition. In *Applications of Computer Vision (WACV), 2013 IEEE Workshop on*, pages 355–362. IEEE, 2013.
- [NIH⁺11] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on*, pages 127–136. IEEE, 2011.
- [PM90] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(7):629–639, 1990.
- [ROF92] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [SSC14] Frank Steinbrucker, Jurgen Sturm, and Daniel Cremers. Volumetric 3d mapping in real-time on a cpu. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2021–2028. IEEE, 2014.
- [Tau95] Gabriel Taubin. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 351–358. ACM, 1995.
- [TM98] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Computer Vision, 1998. Sixth International Conference on*, pages 839–846. IEEE, 1998.
- [VMM99] Jörg Vollmer, Robert Mencl, and Heinrich Mueller. Improved laplacian smoothing of noisy surface meshes. In *Computer Graphics Forum*, volume 18, pages 131–138. Wiley Online Library, 1999.
- [WBS15] Oliver Wasenmüller, Gabriele Bleser, and Didier Stricker. Combined bilateral filter for enhanced real-time upsampling of depth images. *International Conference on Computer Vision Theory and Applications*, 2015.
- [ZFAT11] Youyi Zheng, Hongbo Fu, OK-C Au, and Chiew-Lan Tai. Bilateral normal filtering for mesh denoising. *Visualization and Computer Graphics, IEEE Transactions on*, 17(10):1521–1530, 2011.