

Seamless Mosaicing of Image-Based Texture Maps

Victor Lempitsky

Department of Mathematics and Mechanics
Moscow State University
victorlempitsky@gmail.com

Denis Ivanov

Department of Mathematics and Mechanics
Moscow State University
dvi@fit.com.ru

Abstract

Image-based object modeling has emerged as an important computer vision application. Typically, the process starts with the acquisition of the image views of an object. These views are registered within the global coordinate system using structure-and-motion techniques, while on the next step the geometric shape of an object is recovered using stereo and/or silhouette cues. This paper considers the final step, which creates the texture map for the recovered geometry model.

The approach proposed in the paper naturally starts by backprojecting original views onto the obtained surface. A texture is then mosaiced from these backprojections, whereas the quality of the mosaic is maximized within the process of Markov Random Field energy optimization. Finally, the residual seams between the mosaic components are removed via seam levelling procedure, which is similar to gradient-domain stitching techniques recently proposed for image editing.

Unlike previous approaches to the same problem, intensity blending as well as image resampling are avoided on all stages of the process, which ensures that the resolution of the produced texture is essentially the same as that of the original views. Importantly, due to restriction to non-greedy energy optimization techniques, good results are produced even in the presence of significant errors on image registration and geometric estimation steps.

1. Introduction

The problem of 3D object modeling from registered photographs has received a lot of attention within computer vision community. The emphasis was however put on the geometry recovery [21, 7], while the problem of texture map creation was mainly overlooked. However, a texture map is an important component of a geometric model, and the texture quality and resolution have a key impact on the model realism. In this paper, we therefore investigate how high-resolution texture map can be obtained for a geometric model obtained with some image-based 3D reconstruction method.




Figure 1. An image view (left) is backprojected onto a surface model (middle) yielding a texture fragment (right).

Thus, we consider the creation of a texture map for a geometric model of an object given in a form of triangular mesh. We assume that the image views of the object and the recovered model are registered within the global coordinate system. Consequently, for each point on the surface of the model it is possible to find its projection onto each of the images.

Given such registration, creating image-based texture maps is essentially the problem of the combination of *texture fragments*. Consider an image view V of some object and its surface model S that is defined by a triangular mesh. Once the camera parameters for an image view are known, they establish a one-to-one correspondence between the projection of S onto V and the part of S visible in V . In other words, the image can be backprojected onto the surface. The result of such backprojection is a *texture fragment* (Fig. 1). Multiple views produce multiple texture fragments. The domains of these fragments are different parts of the surface model. Overlapping fragments may differ *photometrically* due to different lighting, camera settings, or non-lambertian object surface and *geometrically* due to model imprecision and/or imperfect registration.

To overcome these misalignments, current solutions to the considered problem employ some kind of blending between the fragments. Thus, some works perform weighted averaging of fragments over the whole mesh [3, 6, 15, 19]. Others employ mosaicing strategy but use feathering to

mask seams after the mosaicing process [2, 11, 14, 20]. These strategies based on intensity blending have the following drawbacks. Firstly, in the presence of geometric misalignments, blending causes ghosting (double imaging) and blurring. Secondly, to blend the fragments, one needs to resample them to some common reference frame (unified parametrization). Constructing good parametrizations for objects with complex topology and geometry is difficult, while inevitable distortions in the computed parametrizations cause image quality degradation during resampling.

At the same time, the problem of combination of texture fragments shares a lot in common with planar image stitching and texture synthesis, where recent introduction of graph cut optimization and *gradient-domain* techniques improved the performance over intensity blending significantly [9, 18, 1, 12]. Therefore, the contribution of this paper is the extension of these techniques to the problem of texture fragment combination on a 3D surface, which to the best of our knowledge has not been done before.

The approach proposed in the paper accomplishes fragment combination in two-steps. On the first step the fragments are *mosaiced* into the texture, whereas the optimal mosaic is sought within Markov Random Field energy optimization framework. On the second step, the mosaic seams are removed via *seam levelling* procedure, which is equivalent to gradient domain based stitching techniques. Importantly, with this approach the fragments need not be resampled to unified texture coordinates. Instead, our method uses what we call a *natural* texture mapping. This texture mapping is defined by the geometric projection from the model surface to the original image which is used to texture the particular surface part. The texture map in this case can be obtained by a simple stacking of the initial views (Fig. 2 middle) and does not incur any resampling. To make the texture more compact, one can remove the unused parts of the texture (Fig. 2 right) and pack the used parts more tightly on a plane. Note, that all such operations do not require resampling and hence do not degrade the quality of the texture.

2. MRF-based mosaicing

2.1. Energy formulation

The first step of our texture mosaicing approach can be regarded as a labelling procedure. Consider a mesh with the faces F_1, F_2, \dots, F_K , and the set of texture fragments V^1, V^2, \dots, V^N , each corresponding to one of the initial views. Then the texture mosaic is defined by a labelling vector $\mathbf{M} = \{m_1, m_2, \dots, m_K\} \in \{0..N\}^K$, prescribing to texture the face F_i from the fragment V^{m_i} .

Obviously, not all fragments containing a face are equally suitable for texturing it. We assume that the quality of a face F_i in the fragment V^j is defined by the cost




Figure 2. Texture maps created with our approach employ a *natural* texture mapping defined by projection operators (creating unified parametrization is not required). Left – a model from Fig. 1 with the mosaiced texture map. Middle – a *natural* texture map obtained by stacking the initial views. Right – the same, but the unused parts are removed.

value w_i^j , which is smaller for the fragments with better quality. If at least part of the face does not fit into the fragment then w_i^j is set to ∞ . For other fragments, different strategies can be used to compute cost values. Thus w_i^j may be computed based on the angle between the local viewing direction of the corresponding view and the face normal (e.g. $w_i^j = \sin^2 \phi + \alpha$), giving preference to less oblique backprojections. We have also tried some more sophisticated variants, which involved surface distances to surface boundaries and the differences between the fragment color and the median local color, and found that they tended to give similar results. A number of previous approaches chose the fragment for each face based on such quality measure, which correspond to the labelling vector $\mathbf{M} = \{m_i | m_i = \arg \min_j w_i^j\}$.

However, such “*best fragment*” approach to mosaicing does not take into account the visibility of the seams that appear when two adjacent faces are textured from different fragments (Fig. 2 left). Therefore, we propose to minimize the visibility of such seams simultaneously with the maximization of fragment quality, thus measuring the quality of each possible mosaic \mathbf{M} with a two-termed energy:

$$E(\mathbf{M}) = E_Q(\mathbf{M}) + \lambda E_S(\mathbf{M}). \quad (1)$$

Here, the **first term** $E_Q(\mathbf{M}) = \sum_{i=1}^K w_i^{m_i}$ is the energy term corresponding to the quality of fragments parts used in the mosaic, while the **second term** $E_S(\mathbf{M})$ in (1) corresponds to the distinguishability of the seams. Assuming that two adjacent faces F_i and F_j share an edge E_{ij} , the visibility of a seam between F_i and F_j is measured by the integral difference of the colors of this edge in corresponding fragments:

$$w_{i,j}^{m_i, m_j} = \int_{E_{ij}} d(Pr_{m_i}(X), Pr_{m_j}(X)) dX. \quad (2)$$

Here, Pr_r is a projection operator for the view r , and $d(\cdot, \cdot)$ is some metric on colors or intensities (we used a Euclidean

distance in RGB color space). Naturally, if $m_i = m_j$ then $w_{i,j}^{m_i,m_j}$ equals zero (texturing from the same view does not produce seams).

If \mathcal{N} denotes a set of pairs of adjacent faces, then the second term in (1) can be written as:

$$E_S(\mathbf{M}) = \sum_{\{i,j\} \in \mathcal{N}} w_{i,j}^{m_i,m_j}, \quad (3)$$

and the overall energy (1) can be rewritten as:

$$E(\mathbf{M}) = \sum_{i=1}^K w_i^{m_i} + \sum_{\{i,j\} \in \mathcal{N}} w_{i,j}^{m_i,m_j}. \quad (4)$$

2.2. Energy optimization

The functionals similar to (4) are common in computer vision, as they are associated with probability distributions of pairwise Markov Random Fields [13]. Such Markov Random Fields (MRF) have been recently applied for mosaicing in texture synthesis and image stitching [9, 1]. For those problems, however, the structure of MRF is typically defined by the pixel or voxel grid, the nodes are associated with pixels or voxels, and the node interactions are added for adjacent pixels and voxels. In our case, the MRF is mesh-based, its nodes correspond to mesh faces, and the node interactions are defined by faces adjacency.

The wide applicability of pairwise discrete MRFs greatly promoted the search for the efficient optimization methods in recent years. For the minimization of our energy, we considered three state-of-the-art algorithms: max-product Belief Propagation (BP) [16], α -expansion Graph Cuts [4], and Convergent Tree Reweighted Message Passing (TRW-S) [8] and found out α -expansion to demonstrate the best performance. It converged to a lowest value of energy using the smallest amount of time (just several seconds for a mesh with hundreds of thousand faces). Therefore all further experimental results were produced using this algorithm. As can be seen from fig. 3, the optimized mosaics gives textures, which are visually superior to “best fragment” approach. Not only the overall length of the seams is reduced in the optimized mosaic, but also the seams there pass through the regions where misalignments between fragments are lower.

3. Seam levelling

MRF-mosaicing decreases the distinguishability of seams significantly. However, in the presence of large photometric differences between original images it can not make seams completely unnoticeable. Therefore, for truly seamless mosaicing, the remaining intensity discontinuities on seams require further processing. A number of gradient-domain (or equivalent) methods for seamless

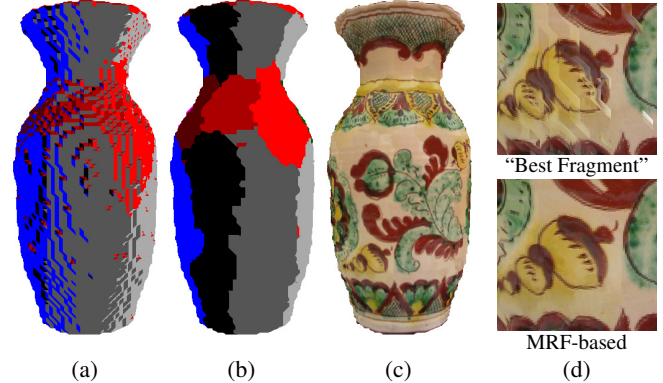


Figure 3. MRF-based mosaicing results in better quality than “best fragment” approach. (a) and (b) – “best fragment” and optimized mosaics; face color corresponds to the number of fragment it is textured with. (c) – a model with an optimized texture map. (d) – closeup comparison of “best fragment” and MRF-optimized mosaics.

stitching, which avoid intensity blending, were proposed in [17, 18, 12] for planar images. Below, we present the seam levelling procedure, which is an adaptation of these methods to the problem of fragment combination on arbitrary manifolds, e.g. meshes.

3.1. Seam levelling on manifolds




Figure 4. Seam levelling on a circumference (function values are shown as elevations above the circumference). By summing the original discontinuous function (left) with the specially computed levelling function (middle), the continuous seam-levelled function (right) that preserves the higher frequencies of the original function is obtained. The discontinuity points of the original function are indicated with radial sticks.

In this subsection, we propose a continuous formulation of seam levelling on an arbitrary smooth oriented manifold M (e.g. a circumference on Fig. 4). Let f be a piecewise-continuous function on M and let S be a seam submanifold of codimension one formed by the discontinuity points of f .

The levelling function g is a piecewise-continuous function with the same discontinuity set as f , defined as follows:

$$g = \arg \min \int_{M \setminus S} (\nabla h)^2 dX \quad (5)$$

$$s.t. [g]_X = -[f]_X, \forall X \in S. \quad (6)$$

Here, the first condition (5) ensures the minimality of gradient magnitude for the levelling function, while the second condition (6) demands that the jump of g at each discontinuity point must be equal to the minus jump of f . As a consequence, the sum $f + g$ (Fig. 4-right) is an everywhere continuous function on M .

Note that due to the condition (5), the levelling function is as smooth as possible within each continuity region. Therefore, this seam levelling procedure preserves higher frequencies of the initial function (Fig. 4).

Note, that the two conditions (5)(6) define g up to an additive constant, which can be chosen arbitrarily, e.g. by equating the mean value of g to zero (for an unconnected manifold, this constant may be chosen independently in every connected component).

3.2. Implementation for meshes

Seam levelling can be applied to remove the seams on the fragments mosaic. Consider some given mosaic M . This mosaic determines some texture map on the surface, i.e. it defines a texture intensity function f on the mesh surface ($f : S \rightarrow \mathbf{R}$). For color textures, we consider three separate functions corresponding to RGB components.

Let C_1, C_2, \dots, C_T be connected components of the mosaic M , i.e. connected sets of triangles textured from the same texture fragment. Then we treat the texture intensity function f as continuous on each C_j and having discontinuities on the boundaries between different C_j . For this f , we seek for the levelling function g .

For the approximate computation of the levelling function g , we make use of the mesh structure. We aim at computing function g in each vertex of the mesh. The values of g between vertices are calculated by interpolation. More specifically, we consider the pair set \mathcal{M} containing all (i, j) -pairs such as at least one triangle adjacent to the vertex V_i belongs to the connected component C_j . For each $(i, j) \in \mathcal{M}$ we compute the value g_i^j , which will be later used to interpolate the value of g over the triangles adjacent to V_i in the fragment C_j .

Let us denote with f_i^j the value of the original texture function on the fragment C_j in vertex V_i (note, that these values will be different for the same V_i but different C_j since f is discontinuous in such points). Also let \mathcal{L} be the set of pairs of adjacent vertices on the mesh. Then the following least-squares energy over the variables g_i^j approximates the conditions (5), (6):

$$\sum_{\substack{(i_1, j) \in \mathcal{M} \\ (i_2, j) \in \mathcal{M} \\ (i_1, i_2) \in \mathcal{L}}} (g_{i_1}^j - g_{i_2}^j)^2 + \lambda \sum_{\substack{(i, j_1) \in \mathcal{M} \\ (i, j_2) \in \mathcal{M}}} (g_i^{j_1} - g_i^{j_2} - (f_i^{j_2} - f_i^{j_1}))^2$$

Indeed, the first term ensures the small absolute value of ∇g and therefore approximates the condition (5), while the



Bust Rabbit Matreshkas
Figure 6. Samples of the source photographs for the experiments (“bust” sequence downloaded from Long Quan’s homepage). Note exposure and white balance variations, as well as surface specularities and varying lighting (flash) in case of “matreshkas”.

second term approximates the condition (6). Since (6) is a hard constraint, we set the parameter λ to a large value (e.g. 100). The values g_i^j are obtained after solving least squares problem with sparse solver and are used to interpolate the values of g over the whole mesh. The resulting levelling function is added pixel-wise to f yielding the final texture (Fig. 5).

4. Experimental Validation


In this section, the ability of our method to generate seamless mosaiced texture maps under a large variety of conditions is validated in a series of experiments.

Experiment with synthetic data. An important question is the performance of the method in the presence of significant errors in the supplied geometry. To control the degree of geometry inaccuracy, we tested the method on a set of 20 computer-generated images of a unit sphere textured with an Earth map along with a set of concentric spheres of different radii. While the precise geometry (a unit sphere) produced the exact copy of the original scene, our method was also able to produce visually pleasant seamless textures for the spheres with other radii (Fig. 7), thus demonstrating the ability to cope with significant errors in geometry.

Experiments with real data. Our method was also tested on several real sequences, containing 16-25 registered photos. In each case, the geometric model was obtained using silhouette intersection on a volumetric grid. Resulting *visual hulls* of the models [10] constituted coarse approximations to the actual object shape, which caused significant misalignments between different fragments and made seamless mosaicing challenging. Arguably, such sil-



Figure 5. Seam levelling for fragments mosaicing. From left to right: the geometric model, the model with MRF-mosaiced texture (arrows point to residual seams), computed levelling function (contrast enhanced), the model with the seam-levelled texture. Seam levelling successfully removed seams preserving the original detailization.



One of source images +10% Correct radius -10% -20% -33% -50%
 Figure 7. "Synthetic Earth" experiment: spheres of different radii are textured based on the images of a unit sphere. Visually-consistent seamless textures are produced for radii significantly deviating from 1 demonstrating the ability of our method to cope with errors on geometry estimation step.




Figure 8. Our method applied to rabbit sequence. Note a coarseness of geometry. Left – the geometric model. Middle – the mosaic structure and the levelling function. Right – the model with the texture produced with our method.

houette intersection is the simplest and most robust yet imprecise way to capture the model geometry.

The following table summarizes the properties of the considered real sequences ordering them from "vase"(simplest) to "matreshkas"(most challenging):

Setup	Exposure & White Bal.	Lighting	Surface	Result
Vase	fixed	fixed	non-lambert.	Fig. 3
Bust	varied	fixed	lambert.	Fig. 5
Rabbit	varied	fixed	lambert.	Fig. 8
Matreshkas	varied	varied	non-lambert.	Fig. 9

Comparative evaluation. Figure 10 presents the close-ups of two models with texture maps obtained using different fragment combination techniques, including traditional fragments blending and unoptimized "best fragment" mosaics. For blending, the weights were chosen according to the angle between the viewing direction and the face normal. After that they were smoothly truncated to zero in the ramp near the border of each fragment.

While the results of blending and "best fragment" approaches exhibit significant artifacts, the addition of seam levelling or MRF optimization improved the texture quality significantly. Importantly, the combination of MRF optimization and seam levelling yielded the results superior to each of the techniques on its own.

5. Conclusion

In this paper, we addressed the problem of creation of seamless texture maps for image-based modeling. Such maps are composed from the backprojected texture fragments using *MRF-based mosaicing* with subsequent *seam levelling* derived from the analogous techniques recently proposed for planar image stitching. Unlike previous approaches to the problem, intensity blending between several fragments and resampling are avoided. As a result, the produced maps have essentially the same resolution as the original photographs. Experiments demonstrate the ability of our method to produce seamless, detail-rich texture maps under a large variety of conditions.



Figure 9. Results for the matreshkas sequence: the geometry, the mosaic, and the views of the textured model. Despite strong specularities on varnished surfaces and varying lighting (built-in camera flash) visually-consistent seamless texture is obtained.




Figure 10. Close-up comparison of different fragment combination techniques for “synthetic Earth” and “bust” setups. For “synthetic Earth” the radius was decreased by 20%. While intensity-blending and “best fragment” mosaic produces inconsistent results, the best quality is attained by the combination of MRF mosaicing and seam levelling proposed in this paper.

6. Acknowledgements

The authors wish to thank Yuri Boykov, Vladimir Kolmogorov, and Olga Veksler for the code of MRF-optimization algorithms [4, 5], and Long Quan for placing online “bust” image sequence.

References

- [1] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive Digital Photomontage. In SIGGRAPH, 2004. 2, 3
- [2] A. Baumberg. Blending Images for Texturing 3D Models. In BMVC, 2002. 2
- [3] F. Bernardini, I. Martin, and H. Rushmeier. High Quality Texture Reconstruction from Multiple Scans. In TVCG, 7(4), 2001. 1
- [4] Y. Boykov, O. Veksler, and R. Zabih. Efficient Approximate Energy Minimization via Graph Cuts. In TPAMI, 20(12), 2001. 3, 6
- [5] Y. Boykov and V. Kolmogorov. An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision. In TPAMI, 26(9), 2004. 6
- [6] P. Debevec, C. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry and image-based approach. In SIGGRAPH, 1996. 1
- [7] C. Dyer. Volumetric Scene Reconstruction from Multiple Views. In *Foundations of Image Understanding*, Kluwer, Boston, 2001. 1
- [8] V. Kolmogorov. Convergent Tree-Reweighted Message Passing for Energy Minimization. In TPAMI, 28(10), 2006. 3
- [9] V. Kwatra, A. Schodl, I. Essa, G. Turk and A. Bobick Graphcut Textures: Image and Video Synthesis Using Graph Cuts. In SIGGRAPH, 2003. 2, 3
- [10] A. Laurentini. The visual hull concept for silhouette-based image understanding. In TPAMI, 16(2), 1994. 4
- [11] H. Lensch, W. Heidrich, H. Seidel. A Silhouette-Based Algorithm for Texture Registration and Stitching. *Graphical Models*, 63(4), 2001. 2
- [12] A. Levin, A. Zomet, S. Peleg, and Y. Weiss. Seamless Image Stitching in the Gradient Domain. In ECCV, 2004. 2, 3
- [13] S. Z. Li, Markov random field modeling in computer vision, Springer-Verlag, 1995. 3
- [14] W. Niem. Automatic reconstruction of 3D objects using a mobile camera. In IVC, vol.17, 1999. 2
- [15] E. Ofek, E. Shilat, M. Werman. Multiresolution Textures from Image Sequences. In *Computer Graphics & Applications*, 17(2), 1997. 1
- [16] J. Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of plausible Inference. CA: Morgan Kaufmann, San Francisco, 1988. 3
- [17] S. Peleg. Elimination of Seams for Photomosaics, In Conf. on Pattern Recognition and Image Processing, 1981. 3
- [18] P. Perez, M. Gangnet, A. Blake. Poisson image editing. In SIGGRAPH, 2003. 2, 3
- [19] F. Pighin, J. Hecker, D. Lischinski, R. Szeliski, D. Salesin. Synthesizing realistic facial expressions from photographs. In Conf. on Comp. Graphics and Interactive Techniques, 1998. 1
- [20] C. Rocchini, P. Cignoni, C. Montani. Multiple textures stitching and blending on 3D objects. In EuroGr. Workshop on Rendering, 1999. 2
- [21] S. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms, In CVPR, 2006. 1