Regularization of Voxel Art

David Coeurjolly
CNRS, Université de Lyon

Pierre Gueth
Arskan

Jacques-Olivier Lachaud
Université Savoie Mont-Blanc

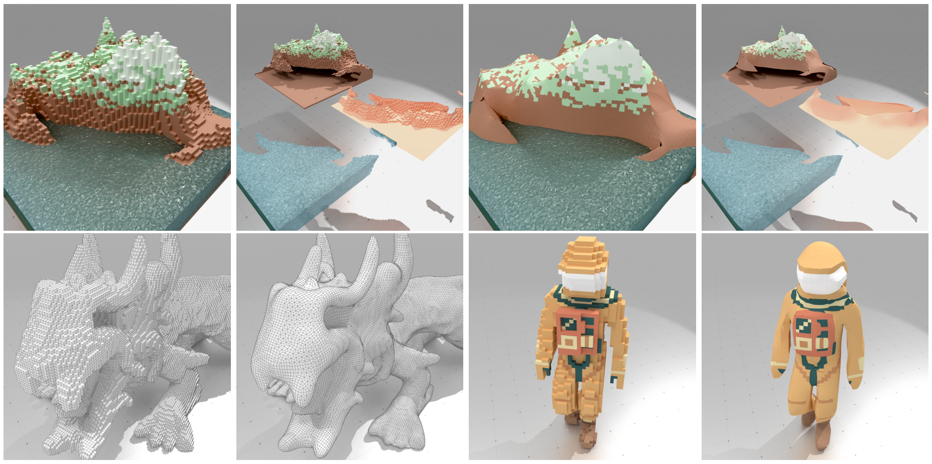

Figure 1: First row, from left to right: input voxel art decomposed into two labels (sea, island, $64 \times 68 \times 40$) and three interfaces, piecewise smooth regularization with consistent interfaces between labels. Second row: high resolution regularization and low resolution voxel art reconstruction.

ABSTRACT

Voxel based modeling is an attractive way to represent complex multi-material objects. Multi-labeled voxel models are also ubiquitous in material sciences, medical imaging or numerical simulations. We present here a variational method to regularize interfaces in multi-labeled digital images. It builds piecewise smooth quadrangulated surfaces efficiently, with theoretical guarantees of stability. Complex topological events when several materials are considered are handled naturally . We show its usefulness for digital surface approximation, for voxel art regularization by transferring colorimetric information, and for upscaling voxel models to speed up coarse-to-fine modeling.

CCS CONCEPTS

• **Computing methodologies → Mesh geometry models**; **Volumetric models**; **Shape analysis**;

KEYWORDS

Digital Geometry, Meshing, Voxel Art, Variational Approach

ACM Reference Format:

David Coeurjolly, Pierre Gueth, and Jacques-Olivier Lachaud. 2018. Regularization of Voxel Art. In *Proceedings of SIGGRAPH '18 Talks.* ACM, New York, NY, USA, Article 37, 2 pages. https://doi.org/10.1145/3214745.3214748

SIGGRAPH '18 Talks, August 12-16, 2018, Vancouver, BC, Canada

ACM ISBN 978-1-4503-5820-0/18/08.
https://doi.org/10.1145/3214745.3214748

INTRODUCTION

In Computer Graphics, the representation of complex objects as data attached to voxels or elements of a regular grid becomes more and more familiar. A key feature of voxel based modeling is that many geometry and topology processing tasks (collision detection, constructive solid geometry, ray-object intersection, homotopic transforms…) are trivially or efficiently defined within this model. Together with very efficient data structures [Kämpe et al. 2013], voxel grids are widely used to represent or simulate complex phenomena that are volumetric, multi-material and/or dynamic (*e.g.*[Wojtan et al. 2007]). Voxel based techniques allow efficient rendering by global illumination [Crassin et al. 2009]. Voxel art and low resolution voxel modeling are also present in the game industry thanks to some dedicated modeling tools [@ephtracy [n. d.]]. In the spirit of [Kopf and Lischinski 2011] which addresses the issue of depixelizing pixel art, our method transforms the crispy representation of voxel surfaces into a smooth reconstruction even for low resolution objects. In this work, we present a new variational formulation to efficiently regularize interfaces of voxel sets and partitions in multi-labeled 3D images. Starting from digital surfaces defined as boundaries of voxel sets, we construct a piecewise smooth regular quadrangulation that captures the singularities of the input voxel geometry. As we maintain a mapping between original digital surface quads and regularized ones, we can transfer information (*e.g.* color, embossing or bump map) from the multi-labeled image to the regularized interfaces. Such smooth surfaces can either be used for visualization purposes, or as an intermediate representation for later geometry processing tasks.

In related works, the popular Marching-Cubes (MC) extract isosurfaces from volumetric data by gluing triangles defined by local configurations [Lorensen and Cline 1987]. When Hermite data are available (points and exact normal vector field), Dual-Contouring (DC) methods solve a quadratic problem in each MC local configuration to construct a smoother dual iso-surface [Ju et al. 2002]. Despite many extensions (adaptiveness, manifoldness…), the overall reconstruction is highly sensitive to normal and position errors, as illustrated in the supplementary. On binary grids, [Chica et al. 2008] proposes a complex two-step algorithm combining a decomposition of the surface into maximal flat patches and a specific smoothing step. We propose a GPU friendly simpler yet robust formulation. Related works are also found in volumetric meshing (*e.g.* [Faraj et al. 2016]) but mainly focus on the interior regularity of 3D mesh rather than on the interfaces. We process here only multi-label interfaces and thus obtain better piecewise smooth surfaces even on low resolution data. As we regularize surfaces, variational remeshing approaches can also be considered (*e.g.* [Zhang et al. 2015], see the supplementary). We provide here an energy formulation dedicated to digital surfaces which is simple, robust and easy to implement efficiently. Finally, for voxel art reconstruction, the work of Muniz *et al.* [Muniz et al. 2013] is mainly focused on colorimetric transfers between voxel objects and MC-like surfaces. Our approach follows a similar idea for color mapping but uses a robust regularization of surfaces geometry.

METHOD OVERVIEW

Let us consider a single voxel object and its boundary defined as isothetic quads. The core of the method relies on a quadratic energy minimization problem. If P denotes the vertices of the input digital surface, F the set of (quadrilateral) faces and $\mathbf{n}_f$ an estimated normal vector on the face f, we want the quad surface vertex positions P^* that minimizes the following energy function:

$$\mathcal{E}(\hat{P}) := \alpha \sum_{i=1}^{n} \|\mathbf{p}_i - \hat{\mathbf{p}}_i\|^2 + \beta \sum_{f \in F} \sum_{\mathbf{e}_j \in \partial f} (\mathbf{e}_j \cdot \mathbf{n}_f)^2 + \gamma \sum_{i=1}^{n} \|\hat{\mathbf{p}}_i - \hat{\mathbf{b}}_i\|^2 . \quad (1)$$

where " $\cdot$ " is the standard $\mathbb{R}^3$ scalar product, $\mathbf{e}_j \in \partial f$ is an edge of the face f (and is equal to some $\mathbf{p}_k - \mathbf{p}_l$) and $\hat{\mathbf{b}}_i$ is the barycenter of the vertices adjacent to $\hat{\mathbf{p}}_i$. The first term is a data attachment term so that the regularized surface stays relatively close to the original one. The second one forces the regularized quad to be perpendicular to the input normal vector field. The last one is a fairness term that enhances the quad aspect ratio and avoids degenerate quads. Each term being quadratic, a unique minimizer P^* exists. Furthermore, closed forms of the energy gradients can be expressed as linear operators. The consequence is twofold. First, one can obtain P^* by solving a linear system (*e.g.* using robust Cholesky decomposition). Second, we can approximate solutions converging to P^* by a simple iterative gradient descent, which allows interactive regularization with a GPU implementation (less than 3ms per gradient step for the Lake scene in the teaser, visually acceptable solution in less than 10 steps). Note that some terms of the energy function in (1) appeared in various forms in previous works on triangular or quad mesh processing. The originality of the contribution consists in an efficient minimization of (1) and its versatility for voxel art regularization as detailed below.

In a mathematical framework where the digital object is the voxelization of a continuous manifold X, we can show that if the normal vector field $\{\mathbf{n}_f\}$ converges to the normal vector field on ∂X as the grid step tends to zero (*e.g.* using [Coeurjolly et al. 2014]), then the regularized surface P^* is close to ∂X in the l_2 sense. The method is stable and robust both to position and normal vector perturbations. Experimentally, we show in Supplementary that our approach outperforms Marching-Cubes or Dual-Contouring approaches when the normal vector field is noisy, and is way more stable to perturbations of positions.

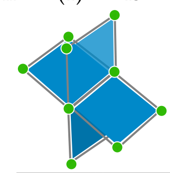

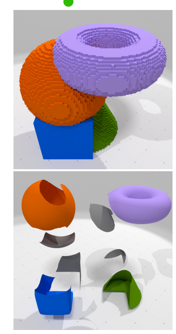

Even if the variational formulation (1) is a global quadratic energy, its evaluation and its gradient are local and can be extended to regularize complex multi-labeled interfaces. On multi-labeled images, we rely on the limited number of topological configurations of digital interfaces: each original quad is shared by two voxels with different labels and each edge is a junction of at most 4 interfaces. In the energy formulation, an edge with more that two adjacent quads contributes more than twice to the energy function (1) (in the second and third term) but the overall formulation is perfectly valid and the same minimization tool is used. Thanks to the robustness of the normal vector field estimated per interface, optimal solution for (1) defines regularized patches that perfectly coincide along topological junctions.

As we maintain a mapping between input and regularized quads, colorimetric information on the input voxel set can be pushed to its smooth interface. In supplementary material, we provide more results of voxel art regularization.

Finally, as we reconstruct piecewise smooth models from voxel shapes, we can build refined voxelizations of initial voxel shapes, just by binarizing our piecewise smooth model. This application is extremely useful for interactive modeling with initial coarse sketches. Please note that the release code of our regularization tool will be publicly available under an open source license.

REFERENCES

A. Chica, J. Williams, C. Andujar, P. Brunet, I. Navazo, J. Rossignac, and A. Vinacua. 2008. Pressing: Smooth isosurfaces with flats from binary grids. *Computer Graphics Forum* 27, 1 (2008), 36–46.

D. Coeurjolly, J.-O. Lachaud, and J. Levallois. 2014. Multigrid Convergent Principal Curvature Estimators in Digital Geometry. *Computer Vision and Image Understanding* 129, 1 (June 2014), 27–41.

C. Crassin, F. Neyret, S. Lefebvre, and E. Eisemann. 2009. GigaVoxels: ray-guided streaming for efficient and detailed voxel rendering. In *SI3D*. ACM, 15–22.

@ephtracy. [n. d.]. MagicaVoxel. ([n. d.]). https://ephtracy.github.io

N. Faraj, J.-M. Thiery, and T. Boubekeur. 2016. Multi-Material Adaptive Volume Remesher. *Compurer and Graphics Journal (proc. Shape Modeling International 2016)* (2016).

T. Ju, F. Losasso, S. Schaefer, and J. Warren. 2002. Dual contouring of hermite data. In *ACM transactions on graphics (TOG)*, Vol. 21. ACM, 339–346.

V. Kämpe, E. Sintorn, and U. Assarsson. 2013. High Resolution Sparse Voxel DAGs. *ACM Transactions on Graphics* 32, 4 (2013). SIGGRAPH 2013.

J. Kopf and D. Lischinski. 2011. Depixelizing pixel art. In *ACM Transactions on graphics (TOG)*, Vol. 30. ACM, 99.

W. E. Lorensen and H. E. Cline. 1987. Marching cubes: A high resolution 3D surface construction algorithm. In *ACM siggraph computer graphics*, Vol. 21. ACM, 163–169.

C. E. V. Muniz, A. Montenegro, M. Lage, and C. N. Vasconcelos. 2013. Polygonal Mesh Extraction From Digital Voxel Art. In *Graphics, Patterns and Images (SIBGRAPI), 2013 26th SIBGRAPI-Conference on*. IEEE, 171–178.

C. Wojtan, M. Carlson, P. J. Mucha, and G. Turk. 2007. Animating Corrosion and Erosion.. In *NPH*. 15–22.

W. Zhang, B. Deng, J. Zhang, S. Bouaziz, and L. Liu. 2015. Guided mesh normal filtering. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 23–34.