Volume xx (200y), Number z, pp. 1–4

Symmetry-Aware Mesh Segmentation into Uniform Overlapping Patches (Supplementary Material)

A. Dessein[1], W. A. P. Smith[2], R. C. Wilson[2], and E. R. Hancock[2]

[1]IMB / LaBRI, Université de Bordeaux, France [2]Department of Computer Science, University of York, UK

Abstract

In this supplementary material, we provide a proof to bound the asymptotic number of patch neighbours, some additional results on mesh segmentation and patch growing experiments as well as some examples of failure cases, and further explanations on our Poisson blending framework.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Bounded Asymptotic Number of Patch Neighbours

For a sufficiently fine segmentation, we may assume that the Delaunay triangulation of the samples exists and has the same topology as $\mathcal{M}$ in terms of genus g and number of boundary components b, which at least holds in the limit when $M = N$ vertices are sampled. We note V, E, F the number of vertices, edges, faces of the triangulation, and show that the average degree of vertices δ (i.e. the expected number of patch neighbours) is asymptotically bounded.

First, each edge defines 2 neighbours so that we have a total of $N = 2E$ neighbours. Second, each face may either share 1, 2, or 3 edges with adjacent faces. An extreme case is that of a narrow strip where faces are assembled iteratively so as to share an edge only with the previous face. Hence, all vertices are on the boundary and $E = 2F + 1$. The opposite extreme is that of a closed mesh where all faces share their 3 edges with adjacent faces. Hence, there is no boundary and $E = 3F/2$. We thus have $3F/2 \leq E \leq 2F + 1$. Introducing the Euler-Poincaré characteristic $\chi = V - E + F$, we can bound δ as $4(1 - \chi/V - 1/2V) \leq \delta \leq 6(1 - \chi/V)$. We know that $\chi = 2 - 2g - b$, or $\chi = 2 - g - b$, for orientable and non-orientable surfaces, respectively. Since both g and b are asymptotically independent of V, χ/V is asymptotically null. This proves that δ is asymptotically bounded between 4 and 6 (and that these bounds are tight since they are reached by the two above extreme cases).

2. Additional Mesh Segmentation Results

To illustrate our mesh segmentation approach with no symmetry cues, we use the Happy Buddha dataset from the Stanford 3D Scanning Repository. This 3D mesh features 543,652 vertices, no boundary components, but a large genus resulting from 5 visible handles and numerous small bridges due to space carving.

We sample an increasing number of vertices for patch centres, and compute the associated segmentations (Fig. 1). It shows that controlling the number of patches allows the patch size to become more regular, thus producing a uniform segmentation as the granularity gets finer. This example, simply used to discuss uniformity, involves a basic version of our algorithm which is actually equivalent to classical furthest-point sampling [PC06]. Improvements that we have introduced, namely seeding, symmetry and overlap, are showcased with subsequent examples.

We now use a scan from the 10 out-of-sample faces of the Basel Face Model database [PKA*09]. The mesh consists of 53,490 vertices, and possesses 1 boundary component but a null genus. Compared to the previous example, we further introduce some semantic constraints via keypoint specification and symmetry cues. Symmetry detection and symmetric remeshing in this example are processed by template fitting.

We compare 4 variants with asymmetric vs symmetric sampling, and seeding vs no seeding (Fig. 2). For a fair com-

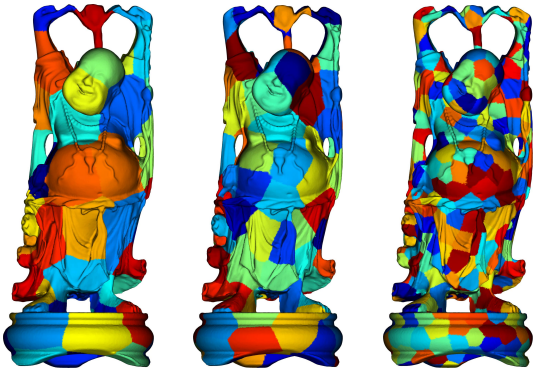

(a) 50 patches. (b) 100 patches. (c) 500 patches.

Figure 1: *Visualisation of mesh segmentation examples. The Happy Buddha dataset is segmented into an increasing number of patches. The resulting patches are displayed by assigning different colours. The patch size gets more regular and leads to a uniform segmentation as the number of patches augments.*

parison, we segment 200 patches in each version. We use 70 hand-labelled seeds forming 41 symmetric pairs according to Farkas' anthropometric set of fiducial points [Far94]. The results prove that seeding and symmetry contribute to providing some semantically meaningful patches. In the asymmetric version with random seeding, some patches split sensitive regions, such as the eyes, where we might want to constrain solutions. As a result, it is likely that this segmentation will not preserve a uniform colour over the whole iris when used for texturing. Moreover, the patches do not preserve the original mesh symmetry. This is addressed in the symmetric segmentation. However, the eyes are still split between different patches. In contrast, using an informed seeding addresses the issues due to splitting sensitive regions. The seeds on the eye centres now allow the irises to be contained within single patches, at least as long as the patch size is not chosen too small. In the end, the most relevant segmentation seems to be obtained for both informed seeding and symmetric sampling. We notice, however, that due to using many seeds compared to the number of patches, the uniformity is hindered in regions where some fiducial points are too close to each other, which is markedly visible above the mouth in particular. The eyes seem oversampled too, where 5 fiducial points are used for each (eye centre, inner and outer eye corners, midpoints of the upper and lower eyelids). Hence, it might be relevant for applications to reduce the seeding set to encourage overall uniformity. For example, in our application to face texture synthesis, we just use the symmetric eye centres for seeding to remove stochastic indeterminacy, and more importantly to ensure that single patches span the eyes so that a single colour is assigned to the irises.

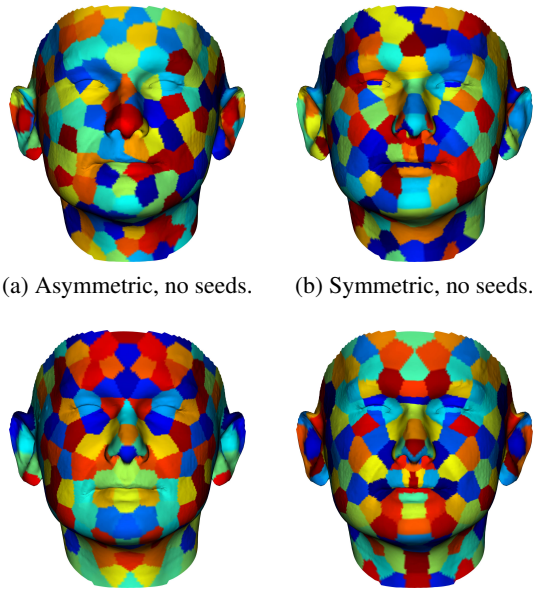

(a) Asymmetric, no seeds. (b) Symmetric, no seeds.

(c) Asymmetric, seeds. (d) Symmetric, seeds.

Figure 2: *Visualisation of mesh segmentation examples. A scan from the Basel Face Model is segmented with or without seeds and symmetry. Coupling seeds and symmetry enables a relevant segmentation of sensitive regions, though overusing seeds should be avoided to maintain uniformity.*

3. Additional Patch Growing Results

To illustrate our patch growing approach with no symmetry cues, we use the Bunny dataset from the Stanford 3D Scanning Repository. The mesh has 35,947 vertices, a null genus, but 5 boundary components due to holes in the bottom.

We segment the mesh into 100 patches and grow them within neighbours using an increasing overlap ratio (Fig. 3). It shows that controlling the degree of overlap allows patches to grow more or less within their neighbours, so that we can modify the influence of neighbouring constraints.

To assess the effect of symmetry cues on patch growing, we now consider the Dino dataset with symmetry detection and symmetric remeshing as performed in [SKS06]. The original dataset features 3,322 vertices, and a simple topology with a null genus and no boundary components. In this example, we upsample the mesh to 19,581 vertices by preserving a symmetric connectivity to obtain finer results, and rig the mesh to deform it and simulate a walking gait with significant geometrical asymmetry between sides.

We compare the asymmetric and symmetric variants by segmenting 100 patches with an overlap ratio of 0.2 (Fig. 4). The asymmetric version does not preserve the mesh symmetries, whereas the symmetric version is robust against geometrical asymmetries and preserves the original symmetries.

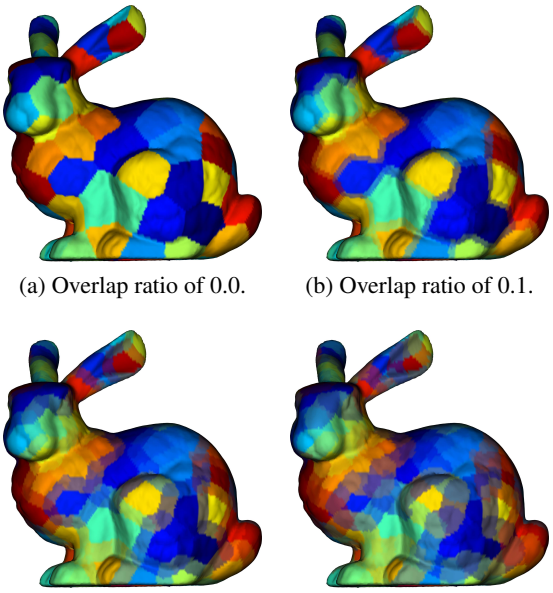

(a) Overlap ratio of 0.0. (b) Overlap ratio of 0.1.

(c) Overlap ratio of 0.2. (d) Overlap ratio of 0.3.

Figure 3: *Visualisation of patch growing examples. The Bunny dataset is segmented and the obtained patches are grown with increasing overlap ratios. The resulting patches are displayed by blending colours in overlap regions. These regions, and thus the influence of neighbouring constraints, spread as the overlap ratio augments.*

4. Failure Cases

We now show some examples of what could be considered as failure cases for our segmentation algorithm (Fig. 5). In some sense, the algorithm never really fails. It always outputs a segmentation, and it is always symmetric if symmetry information is provided. Hence, what we investigate here is the worst performance obtained over the TOSCA dataset. For each object, we consider all poses and patch numbers (M), and select the less uniform result (highest SU value).

From the results, it is clear that asymmetry is not the determining factor in poor performance. In two of the cases, the meshes are almost extrinsically symmetric, whereas they are highly asymmetric in the other two. The worst performance occurs when the number of patches is the smallest. In both the visualisation and histogram, it is evident that this is due to a peak of atypically small patches. For all four shapes, features such as feet and hands are segmented into much smaller patches than less curved regions on the body. As the number of patches increases, this effect is reduced. This is a property of using furthest-point sampling rather than a property of our symmetric enhancement. It is caused by the fact that keypoint placement seeks to maximise the minimum distance to previous keypoints, which is not the same as selecting patch centres that maximise directly the patch size uniformity.

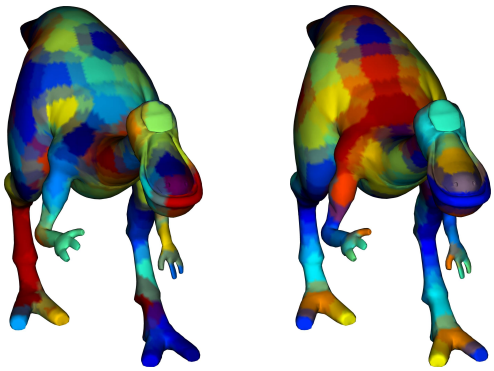

(a) Asymmetric. (b) Symmetric.

Figure 4: *Visualisation of patch growing examples. The Dino dataset is segmented and the obtained patches are grown with or without accounting for symmetry cues. The asymmetric version is not robust against geometrical irregularities, whereas the symmetric variant provides a symmetry-preserving structure of uniform overlapping patches.*

5. Poisson Blending Framework

In order to blend the overlapping textures in our segmentation applications, we adapt a principled approach to Poisson's equation on a mesh with Dirichlet boundary conditions [YZX*04] so as to incorporate a screening term instead. A discrete vector field V is a piecewise constant vector function defined for each triangle T_l by a coplanar vector v_l. A discrete potential field is a piecewise linear function $\phi(s) = \sum_{i\in\mathcal{K}} \phi_i B_i(s)$ on the mesh surface, where B_i is the piecewise linear basis function valued 1 at vertex i and 0 at other vertices, and ϕ_i specifies the value of ϕ at vertex i. The discrete gradient of ϕ for triangle T_l is $\nabla\phi_l = \sum_{i\in\mathcal{K}} \phi_i \nabla B_{il}$, where ∇B_{il} is the gradient of B_i within T_l. The divergence of V at vertex i is $\mathrm{div}\,V(i) = \sum_{T_l\in\mathcal{K}_i} |T_l|\,\nabla B_{il}^\top\,\mathbf{v}_l$, where $\mathcal{K}_i$ is the set of triangles sharing vertex i and $|T_l|$ is the area of triangle T_l. Writing Poisson's equation $\mathrm{div}\,\nabla\phi = \mathrm{div}\,V$ in this framework leads to a linear system of equations $\mathbf{Ax} = \mathbf{y}$ for the unknown potential values $x_i = \phi_i$, where:

$$a_{ij} = \sum_{T_l\in\mathcal{K}_i} |T_l|\,\nabla B_{il}^\top \nabla B_{jl} \quad , \quad y_i = \sum_{T_l\in\mathcal{K}_i} |T_l|\,\nabla B_{il}^\top \mathrm{v}_l \quad . \quad (1)$$

This equation can be interpreted as seeking for a potential field ϕ whose gradient $\nabla\phi$ matches the guide vector field V. If V is conservative, i.e., it is the gradient of an existing potential field ϕ, then ϕ is the exact solution. Otherwise, a more general minimizer can still be obtained by least squares but its gradient differs from V. In addition, we regularize the minimization via screening:

$$\min_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{Ax}-\mathbf{y}\|_2^2 + \lambda\left\|\mathbf{x}-\mathbf{x}'\right\|_2^2 \quad , \quad (2)$$

where $\lambda > 0$ and $\mathbf{x}'$ is a guide potential field ϕ' with $\phi'_i = x'_i$.

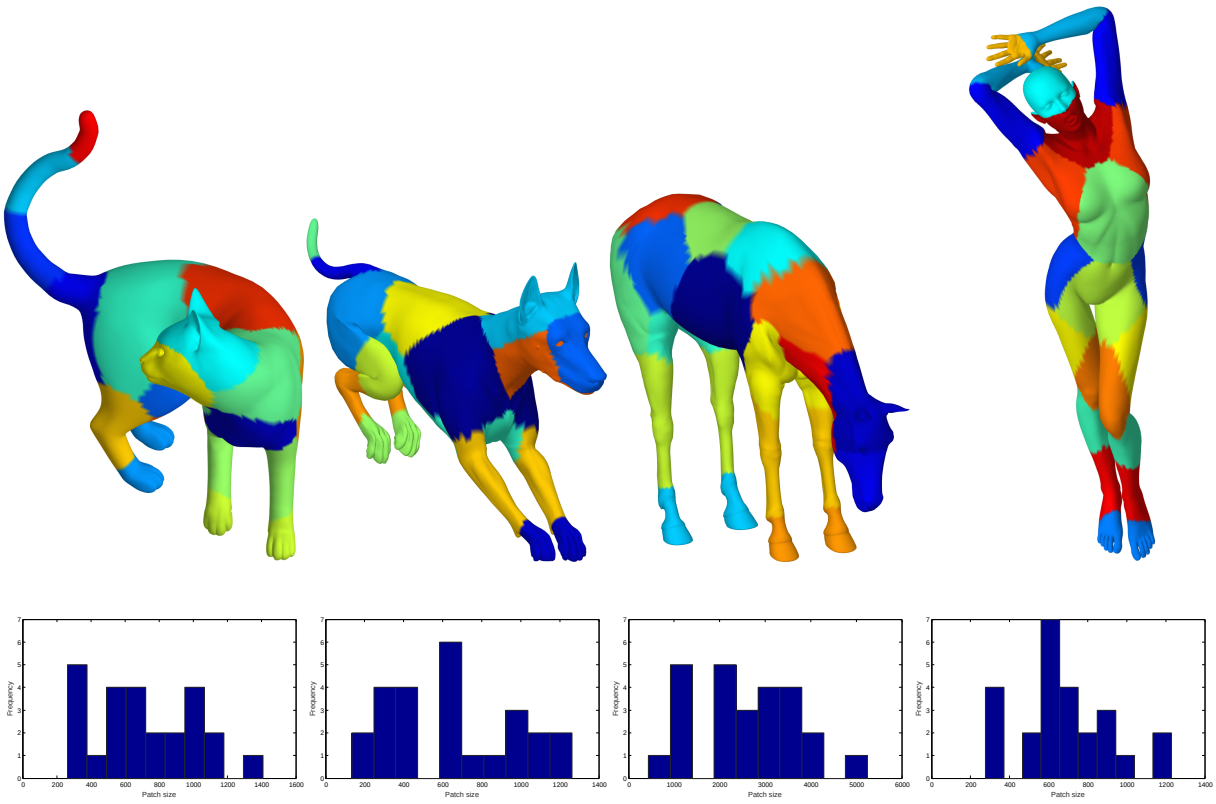

(e) $M=25$, SU$=8.6\%$ (f) $M=25$, SU$=10.4\%$ (g) $M=25$, SU$=9.1\%$ (h) $M=25$, SU$=7.3\%$

Figure 5: *Visualisation and patch size histogram for failure cases. For each object included in our experiments on the TOSCA dataset, we present the worst segmentation result over all patch numbers (M) in terms of size uniformity (*SU*). In the top row, we show a visualisation of the segmentation. In the bottom, row we draw the histogram for the distribution of patch sizes.*

References

[Far94] FARKAS L. G.: *Anthropometry of the Head and Face*. Raven Press, New York, 1994. 2

[PC06] PEYRÉ G., COHEN L. D.: Geodesic remeshing using front propagation. *Int. J. Comput. Vis. 69*, 1 (2006), 145–156. 1

[PKA*09] PAYSAN P., KNOTHE R., AMBERG B., ROMDHANI S., VETTER T.: A 3D face model for pose and illumination invariant face recognition. In *IEEE Int. Conf. on Advanced Video and Signal Based Surveillance* (2009), pp. 296–301. 1

[SKS06] SIMARI P., KALOGERAKIS E., SINGH K.: Folding meshes: Hierarchical mesh segmentation based on planar symmetry. In *Eurographics Symp. on Geometry Processing* (2006), pp. 111–119. 2

[YZX*04] YU Y., ZHOU K., XU D., SHI X., BAO H., GUO B., SHUM H.-Y.: Mesh editing with Poisson-based gradient field manipulation. *ACM Trans. Graph. 23*, 3 (2004), 644–651. 3