

Persistent Homology for 3D Reconstruction Evaluation

Antonio Gutierrez¹, David Monaghan²,
María José Jiménez¹, and Noel E. O'Connor²

¹ Applied Math Department, School of Computer Engineering,
University of Seville,
Campus Reina Mercedes, 41012 Sevilla, Spain
`majiro@us.es`

² CLARITY: Centre for Sensor Web Technologies,
Dublin City University, Ireland
`{david.monaghan,noel.oconnor}@dcu.ie`

Abstract. Space or voxel carving is a non-invasive technique that is used to produce a 3D volume and can be used in particular for the reconstruction of a 3D human model from images captured from a set of cameras placed around the subject. In [1], the authors present a technique to quantitatively evaluate spatially carved volumetric representations of humans using a synthetic dataset of typical sports motion in a tennis court scenario, with regard to the number of cameras used. In this paper, we compute persistent homology over the sequence of chain complexes obtained from the 3D outcomes with increasing number of cameras. This allows us to analyze the topological evolution of the reconstruction process, something which as far as we are aware has not been investigated to date.

Keywords: voxel carving, volume reconstruction, persistent homology, evaluation.

1 Introduction

Homology is topologically invariant, meaning it is a property of an object that does not change under continuous (elastic) transformations of the object. Roughly speaking, homology characterizes “holes” in any dimension (e.g. connected components, tunnels and cavities in a 3D space). Homology computation can be carried out over a combinatorial structure called *cell complex*, which is built up by basic elements (*cells*) of different dimensions (vertices, edges, faces, etc.). One can take advantage of the combinatorial nature of a digital image (as a set of voxels) to compute homology by taking as input the (algebraic) cubical complex associated to the image. *Persistent homology* studies homology classes and their life-times (persistence) in the belief that significant topological attributes must have a long life-time in a filtration (an increasing nested sequence of subcomplexes). In this paper, we compute persistent homology via the Incremental and Decremental Algorithms for computing AT-models (see [5]), which allow to combine an incremental with a decremental technique in the case of a non-increasing

filtration, that is, a sequence of subcomplexes. In the following Section, we describe the context in which we apply persistent homology computation. Section 3 is devoted to recall basic tools used in our computations. Section 4 describes the application of persistent homology to the evaluation of the voxel carving process. We draw some conclusions and ideas for future work in the last Section.

2 Voxel Carving Approach

Space carving is a well-known method for constructing three-dimensional models of objects from a set of images. The process involves capturing a series of images of an object, and, by analysis of these images, deriving a description of the shape of the object. In particular, space (or voxel) carving approaches [2,3,9,11] are non-invasive techniques that allow the reconstruction of a 3D human model from the images captured from a set of cameras placed around the subject. In each image, firstly, the region of interest (subject silhouette) is segmented from the background by an autonomous adaptive “approximate median” background modelling algorithm; then a 3D bounding box is drawn around the subject’s approximate position in 3D space. By using extracted silhouettes from each image, inconsistent voxels are eliminated from the defined volume, iterating through each of the cameras [9]. In [1], the authors present a technique to quantitatively evaluate spatially carved volumetric representations of humans using a synthetic dataset of typical sports motion in a tennis court scenario. Such a quantification is based on the computation of Normalised Mean Square Error (NMSE) of a groundtruth volumetric reconstruction (which has been considered at 50 cameras, based on experimental observation) against any reconstruction from an inferior camera setup (with less cameras than the setup used to carve the ground truth). The aim of such an evaluation is to somehow quantify the accuracy of the 3D volume produced by the voxel carving process with regard to the number of cameras used. This investigation was motivated by the fact that very little work has been done to date on evaluating the quality of space carving results. In this paper, we intend to give a different insight into the voxel carving work by homologically characterising the sequence of reconstruction volumes. This may be interesting as the surfaces produced with a few cameras are quite noisy with many holes, which are irrelevant topological information that can be discarded by using persistent homology. Given the nature of the carvings, we believe that a homology-based approach is a more appropriate quantification than the relatively simple NMSE-based approach used previously.

3 Homology Computations on a Set of Voxels

A *cell complex* is a general topological structure by which a space is decomposed into basic elements (*cells*) of different dimensions, which are *glued* together by their boundaries (see a formal definition of CW-complex in [8]). Due to the nature of our input data, we focus on a special type of cell complex: *cubical complex*. A cubical complex Q in \mathbf{R}^3 , is given by a finite collection of p -cubes




Fig. 1. Voxel carving approach for 3D reconstruction. Process with 4 cameras around the subject and an overhead camera.

such that a 0-cube is a vertex, a 1-cube is an edge, a 2-cube is a filled square (we call it, simply, a square) and a 3-cube is a filled cube (resp. a cube); together with all their faces and such that the intersection between two of them is either empty or a face of each of them.

We consider $\mathbf{Z}/2$ as the ground ring for algebraic computations, since we do not need to deal with torsion. The *cubical chain complex* associated to the cubical complex Q is the collection $\mathcal{C}(Q) = \{C_p(Q), \partial_p\}_p$ where:

- (a) each $C_p(Q)$ is the corresponding chain group generated by the p -cubes of Q , over $\mathbf{Z}/2$;
- (b) the boundary operator $\partial_p : C_p(Q) \rightarrow C_{p-1}(Q)$ connects two immediate dimensions. The boundary of a p -cube is the formal sum (mod 2) of all its facets (proper faces of maximal dimension). It is extended to p -chains by linearity.

Roughly speaking, the homology groups of a cubical chain complex will be a chain group whose elements are equivalence classes of *cycles*, such that if one cycle can be obtained from another by continuous deformation through the object, then they are *homologous* (or equivalent). For example, two vertices are homologous if there exists a path through the object between them. Formally, a p -cycle is a p -chain a such that $\partial_p(a) = 0$. If $a = \partial_{p+1}b$ for some $p + 1$ -chain b then a is called a p -boundary. We say that two p -cycles a and b are *homologous* if there exists a $(p + 1)$ -chain c such that $a = b + \partial_{p+1}c$. Define the p -th homology group to be the quotient group of p -cycles mod p -boundaries denoted by $H_p(Q)$. Each element $[a]$ of $H_p(Q)$ is a quotient class obtained by adding each p -boundary to a given p -cycle a called a *representative cycle* of the homology class $[a]$. The *homology* of Q is the chain group $\mathcal{H}(Q) = \{H_p(Q)\}_p$. See [10] for further details.

3.1 Incremental-Decremental Algorithms for Computing Persistent Homology

We focus on homology computation methods based on the concept of *AT-model* [7]. Given a cell complex, Incremental Algorithm for computing AT-models [7]

computes homology information of the cell complex by an incremental technique, considering the addition of a cell each time. Once homology of an object has been computed, the same algorithm can be used again to update homology information if new cells are added to the existing complex; Decremental Algorithm for computing AT-models [6] can be used for the same aim, in the case that some cells are deleted.

Given a cubical complex Q , an algebraic-topological model (*AT-model* [7]) for Q is a set of data (Q, H, f, g, ϕ) , such that:

- Q is the cubical complex itself.
- H is a subset of Q that characterizes the homology of Q by containing a p -cube for each p -homology class, for all p . In 3D, H can only have points, edges and squares: each point of H represents a connected component of Q , each edge represents a “tunnel” and each square represents a “void” (i.e. a connected component of the background inaccessible from the outside).
- f is a chain map from $\mathcal{C}(Q)$ to $\mathcal{C}(H)$. This map provides the equivalence relation between cycles (that is, if two cycles, a and b , are equivalent, then $f(a) = f(b)$). Moreover, $f g(c) = c$ for any $c \in H$.
- g is a chain map from $\mathcal{C}(H)$ to $\mathcal{C}(Q)$. For each cube c in H , $g(c)$ is a representative cycle of a homology class.
- ϕ is a map from $\mathcal{C}(Q)$ to $\mathcal{C}(Q)$ that is a chain homotopy (see [10]) from $g f$ to the identity homomorphism on $\mathcal{C}(Q)$. This map can be seen as a kind of boundary inverse. For example, if c is a vertex, then $\phi(c)$ is the path from c to the vertex $v \in H$ homologous to c .




Fig. 2. A simple example of execution of Incremental Algorithm for computing AT-models. a) The input cubical complex, a filled square with all its faces (only the labels of the vertices are shown). b) The elements in H . c) The table with the information related to f , g and ϕ . Read, for instance, $f(16) = 16$, $g(16) = 16$, $\phi(16) = 0$, $\phi(17) = 16 - 17$ (edge from 16 to 17).

In [5], the authors revisit the algorithm for computing AT-models using an incremental technique that appears in [7] (we will refer to it as the Incremental Algorithm) with the aim of setting its equivalence with persistent homology computation algorithm [4,12]. Given a cubical complex Q associated to a 3D digital image, consider a full ordering of its cubes $\{c^1, \dots, c^n\}$ such that if c^i is a

face of c^j , then $i < j$; take a nested sequence of subcomplexes $\emptyset = Q^0 \subseteq Q^1 \dots \subseteq Q^n$ (a filtration over Q) such that $Q^i = \{c^1, \dots, c^i\}$ (notice that all the proper faces of c^i are in Q^{i-1}). Under these conditions, Incremental Algorithm can be applied to compute persistent homology over the filtration.

See Fig. 2 as a simple example of execution of the Incremental Algorithm for computing AT-models.




Fig. 3. A simple example of execution of Decremental Algorithm for computing an AT-model (Q', H', f', g', ϕ') after removing a 2-cube (the square 16 – 17 – 24 – 23) from Fig. 2.a. a) The output cubical complex (only the labels of the vertices are shown) after deleting the square. b) The cubes in H' . c) The table with the information related to f' , g' and ϕ' .

Now, let (Q, H, f, g, ϕ) be an AT-model for a cubical complex Q computed by the Incremental Algorithm. Let c^m be a maximal cube of Q . Then an AT-model for $Q' = Q \setminus \{c^m\}$, (Q', H', f', g', ϕ') , can be constructed by the Decremental Algorithm given in [5], where it was redefined (with respect to the one of [6]) with the aim of extending the concept of persistent homology for objects with a filtration that is not necessarily increasing.

See Fig. 3 as an example of execution of Decremental Algorithm for computing AT-models. Notice that by removing the 2-cube from the initial cubical complex on Fig. 2, a new homology class is created. The output of the algorithm is the set (Q', H', f', g', ϕ') represented in a table form in Fig. 3.c).

Now, let $\emptyset = Q^0 \leftrightarrow Q^1 \leftrightarrow \dots \leftrightarrow Q^n$ be a *zig-zag filtration*, that is, a sequence of cell complexes such that every two consecutive complexes differ by a single cell c , i.e. either $Q^i = Q^{i-1} \cup \{c\}$ or $Q^i = Q^{i-1} \setminus \{c\}$. Then, one can compute persistent homology over the filtration by combining the application of Incremental and Decremental Algorithms depending on whether a cell c is added or deleted each time.

4 Persistent Homology for 3D Reconstruction Evaluation

We are concerned with the application of persistent homology computation to provide topological evaluation of the 3D reconstruction process by the voxel carving technique. The new insight could significantly enrich the evaluation made

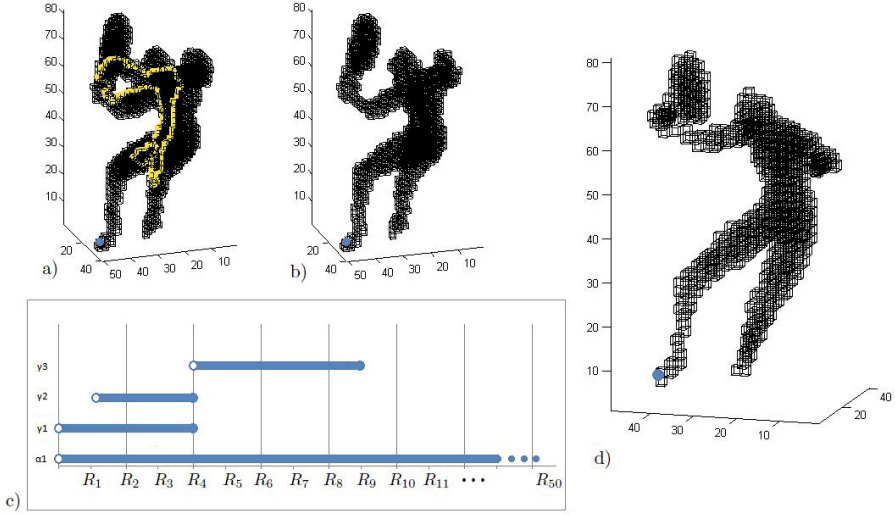


Fig. 4. 3D Reconstructions using a) 4 cameras and b) 10 cameras. Representative cycles of homology are highlighted in both cases. c) Barcode associated to the whole sequence of 3D reconstructions with increasing number of cameras (from 1 to 50). d) 3D reconstruction using 50 cameras, what is considered the groundtruth model.

in [1] by means of NMSE quantification. For this aim, we must consider the sequence of different 3D models, obtained by voxel carving under increasing number of cameras, as a whole object on which we have to set up a filtration over which to compute persistent homology. This way, in particular, we can get an estimation of the minimum number of cameras needed in order to obtain a topologically correct 3D model (which in general has only one connected component and no tunnels or voids).

We denote by R_k the cubical complex associated to the 3D reconstruction obtained using k cameras (which are randomly chosen). Starting from the first reconstruction R_1 (obtained by “one carving” of the initial 3D bounding box), we can use Incremental Algorithm to compute its homology. Notice that R_{k+1} may be obtained from R_k by removing some voxels (cubes, together with all their faces in the cubical complex). This fact makes this context good for making use of the Decremental Algorithm for getting homology computations through increasing number of cameras. Both, Incremental and Decremental Algorithms provide all the pairs of cells responsible for the creation/destruction of homology classes along the process, what allows to follow the evolution of these classes with respect to time, that is, the number of cameras used. Actually, to compute persistent homology of the whole sequence of 3D models, $\{R_k\}_k$, the zigzag filtration is given by the sequence $\{R_k\}_k$ itself with the inclusion, between R_k and R_{k+1} , of a sequence of complexes $\{R_k^{i_k}\}_{i_k=1\dots n_k}$ given by the addition or deletion of a cell, each time. Compute, then, a big barcode for visualizing the hole computation in

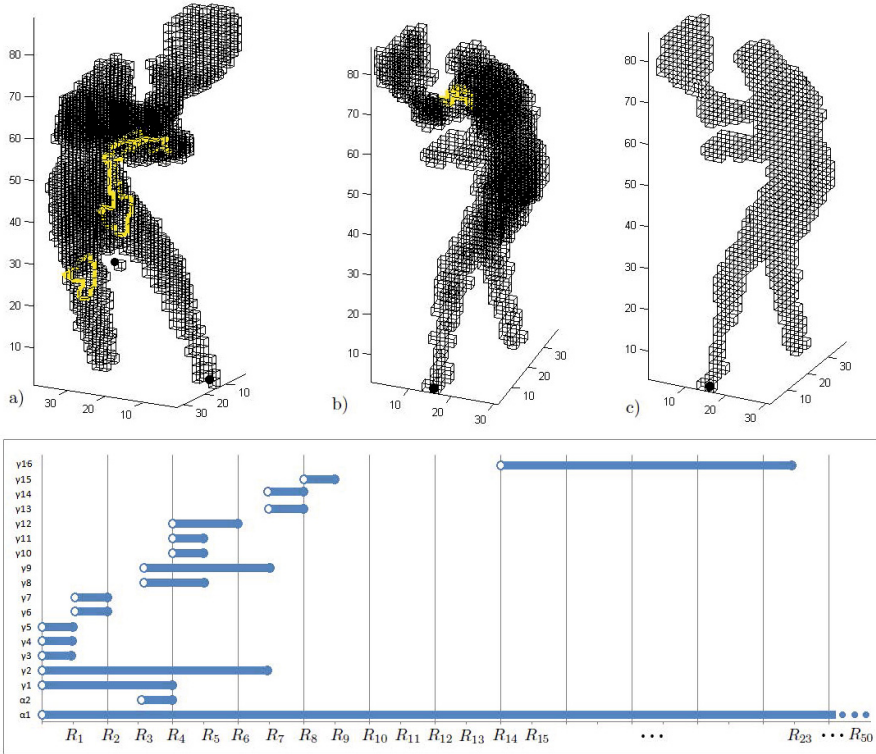


Fig. 5. 3D Reconstructions (viewed from different angles) using different number of cameras: a) 4 cameras, b) 15 cameras and c) 24 cameras, which is similar to the one obtained with 50 cameras (groundtruth model). Representative cycles of homology are highlighted. Below, barcode associated to the whole sequence of 3D reconstructions from 1 to 50 cameras.

order to easily analyze the stability of the elements of homology. We want also to remark that, due to the nature of the voxel carving process, only voxels on the surface of the object are removed each time, so different connected components and tunnels (but no cavities) may arise.

We have used for computation five different frames extracted from a 3D video sequence with a voxel resolution of 4 cm, that is, the spacing between each voxel is 4 cm in the OX , OY and OZ directions. This means 15,625 voxels per cubic metre. We have appreciated, as it was expected, that simpler poses of the subject produce simpler barcodes while more complex poses give place to more interesting homological information. Fig. 4 shows that the carving process, in a case of simple pose, stabilizes at 10 cameras (with a unique connected component), while below that point, 3 different tunnels have been living for some time. That means that, in order to produce a topologically correct model,

at least 10 cameras are needed. Fig. 5 reflects a more complex pose, though it also corresponds to a 3D object with one connected component and no tunnels or voids. Notice the more complex barcode associated (in which 2 connected components and 16 tunnels are represented) and, especially, the fact that a 1-homology class is created at time $k = 15$, that persists until $k = 23$. So stabilization of one connected component as final state, occurs much later than in the former case.

We are working also on other approaches:

- To compute persistent homology of the sequence of 3D difference complexes $\{D_k\}_k$ with respect to the groundtruth model (R_∞), where $D_k = R_k \setminus R_\infty$. Now the barcode for the whole sequence will provide different information about the whole process that might complement the one given by the reconstructions themselves.
- To compute persistent homology of the sequence of 3D complexes generated by the convex deficiencies of each 3D reconstruction, that is, the complexes obtained by the subtraction of each 3D reconstruction to its 3D convex hull.

5 Conclusions and Future Work

Persistent homology computation provides an interesting new insight into the 3D model reconstruction process explained in this paper. There are lots of ideas and experimentation still to be investigated. An important point is to study the dependence of the observations on the resolution of the input data. An alternative approach could be to compute some homology-based features extracted from each reconstruction R_k and to compare them against a groundtruth model. These features should be measurable so that a distance with respect to the groundtruth model could be computed. These parameters could be extracted from the comparison of weighted histograms of connected components (for 0-homology study) or minimal-length (in some sense) representative cycles of 1-homology. Another interesting question arises by fixing a certain number of cameras and considering the sequence of 3D reconstructions along time. Holes can be produced along time by different movements of the body that should not be stable along the complete scene, so the homological analysis of voxel carving performance of video sequences could shed some light on the classification of these movements.

References

1. Monaghan, D., Kelly, P., OConnor, N.E.: Quantifying Human Reconstruction Accuracy for Voxel Carving in a Sporting Environment. In: ACM MM, Scottsdale, AZ, November 28 - December 1 (2011)
2. Broadhurst, A., Drummond, T., Cipolla, R.: A probabilistic framework for space carving. In: Conf. on Computer Vision, vol. 1, p. 388 (2001)
3. Culbertson, W.B., Malzbender, T., Slabaugh, G.: Generalized voxel coloring. In: Intern. Workshop on Vision Algorithms: Theory and Practice, pp. 100–115 (1999)

4. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. In: FOCS 2000, pp. 454–463. IEEE Computer Society (2000)
5. Gonzalez-Díaz, R., Ion, A., Jimenez, M.J., Poyatos, R.: Incremental-Decremental Algorithm for Computing AT-Models and Persistent Homology. In: Real, P., Diaz-Pernil, D., Molina-Abril, H., Berciano, A., Kropatsch, W. (eds.) CAIP 2011, Part I. LNCS, vol. 6854, pp. 286–293. Springer, Heidelberg (2011)
6. Gonzalez-Díaz, R., Medrano, B., Sánchez-Peláez, J., Real, P.: Simplicial Perturbation Techniques and Effective Homology. In: Ganzha, V.G., Mayr, E.W., Vorozhtsov, E.V. (eds.) CASC 2006. LNCS, vol. 4194, pp. 166–177. Springer, Heidelberg (2006)
7. Gonzalez-Díaz, R., Real, P.: On the cohomology of 3D digital images. *Discrete Applied Math.* 147(2-3), 245–263 (2005)
8. Hatcher, A.: *Algebraic Topology*. Cambridge University Press (2002)
9. Kutulakos, K.N., Seitz, S.M.: A theory of shape by space carving. *Intern. Journal of Computer Vision* 38, 199–218 (2000)
10. Munkres, J.: *Elements of Algebraic Topology*. Addison-Wesley Co. (1984)
11. Seitz, S.M., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 519–528 (2006)
12. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete and Computational Geometry* 33(2), 249–274 (2005)