

Opening Holes in Discrete Objects with Digital Homotopy

Aldo Gonzalez-Lorenzo^(✉), Alexandra Bac, and Jean-Luc Mari

Aix-Marseille Université, CNRS, LSIS UMR 7296, Marseille, France
aldo.gonzalez-lorenzo@univ-amu.fr

Abstract. Discrete objects are sets of pixels, voxels or their analog in higher dimension. A three-dimensional discrete object can contain holes such as tunnels, handles or cavities. Opening the holes of an object consists in erasing all its holes by removing some parts of it. The main idea is to take a point of the object and to dilate it inside the object without changing its homotopy type: the remaining points in the object are those which have to be removed. This process does not require the computation of the homology groups of the object and is only based on the identification of simple points.

In this experimental paper we propose two algorithms for opening the holes of a discrete object endowed with any adjacency relation in arbitrary dimension. Both algorithms are based on the distance transform of the object and differ in how the dilation is performed, favoring either time complexity or the quality of the output. Moreover, these algorithms contain a parameter that controls the thickness of the removed parts.

1 Introduction

Topological features such as connected components, tunnels, handles or cavities allow us to understand the essential structure of an object, which is invariant under a continuous deformation. They are defined through algebraic topology. Roughly speaking, homology theory defines a q -dimensional hole as a q -dimensional sphere which is not the boundary of a $(q + 1)$ -dimensional ball. Hence, connected components are 0-dimensional holes; tunnels and handles are 1-dimensional holes and cavities are 2-dimensional holes. On the other hand, homotopy theory detects q -dimensional holes with q -dimensional spheres that cannot be continuously deformed into a point.

In order to erase a hole, one can add matter to the object so it disappears. This is called *closing* a hole. For instance, we can remove a connected component by adding a *bridge* to another connected component, a handle by adding a patch or a cavity by filling its interior. Intuitively, we can erase a q -dimensional hole by adding a $(q + 1)$ -dimensional ball.

Another way to erase a hole is by removing matter from the object. We call this *opening* a hole. As an illustration, we can open a 0-dimensional hole by removing the whole connected component, a handle by cutting a slice or a

cavity by digging a well. In this case, if the object is embedded in the three-dimensional space, we say that we can open a q -dimensional hole by removing a $(3 - q)$ -dimensional ball.

The problem of closing holes has been addressed in [1, 7, 8]. We recall that an object is *contractible* if it has the homotopy type of a point [6]. Hence, a contractible object has no holes at all. A trivial way of filling the holes of an object is to consider a contractible object that contains it: the difference between them is the matter that we have to add to close its holes. In order to optimize the quantity of matter that we add, we can shrink it while preserving its contractibility property and keeping the object inside.

In this experimental article we study the problem of opening the holes in an object. The main idea is to find a contractible subset of the object by taking a point in its interior and expanding it without changing its homotopy type. This has been shortly investigated in [10] for segmenting the brain cortex in magnetic resonance images. In this paper we develop farther this problem and design two algorithms for opening the holes of a discrete object using additional geometric conditions to obtain a visually pleasant output.

2 Preliminaries

2.1 Discrete Object

A d -dimensional *discrete object* is a (finite) subset of \mathbb{Z}^d . Its elements are called pixels when $d = 2$, voxels when $d = 3$ or points in general.

We endow a discrete object with an *adjacency relation*. Let us mention two of them: two points $x, y \in \mathbb{Z}^d$ are $(2d)$ -adjacent (resp. $(3^d - 1)$ -adjacent) if $\|x - y\|_1 \leq 1$ (resp. $\|x - y\|_\infty \leq 1$). Let $\alpha \in \{2d, 3^d - 1\}$, the α -neighborhood of a point x , denoted $N_\alpha(x)$, is the set of its α -adjacent points. We also denote $N_\alpha^*(x) := N_\alpha(x) - \{x\}$. The *outer α -boundary* of a discrete object X is $N_\alpha^*(X) := (\bigcup_{x \in X} N_\alpha(x)) - X$, the set of points in $\mathbb{Z}^d - X$ having an α -adjacent point in X .

2.2 Simple Points

Roughly speaking, a point is *simple* for a discrete object X if its addition or removal from X does not change the homotopy type of X . As a discrete object is just a set of isolated points, this only makes sense if we endow the discrete object with a topological space. This is usually done in terms of the adjacency relation considered for the object, so there are different characterizations of simple points. See [2, 3, 5, 11] to cite a few.

We consider in this paper the $(3^d - 1)$ -adjacency relation and its associated notion of simple point as described in [5], though any other definition of simple point can be used.

3 Homotopic Opening

A topological space is *contractible* if it has the homotopy type of a point. Intuitively, this means that we can shrink it to a point. We can translate this concept into the discrete context by using the notion of simple point. We recall that a point in a discrete object is simple if its removal does not change the homotopy type of the object. Consequently, a discrete object is contractible if we can reduce it to a point by a sequence of simple points deletions.

A contractible discrete object has no holes. Given a discrete object X , we say that $Y \subset X$ is a *homotopic opening* if $X - Y$ is contractible. Hence, the homotopic opening consists of a set of *cuts* that remove the holes from the object. Note that, if the object is not connected, then all the connected components except one will be in the homotopic opening, so it may be more interesting to consider the homotopic opening of each connected component separately.

The homotopic opening is clearly not unique. We could be interested in obtaining a minimal (in the number of points) homotopic opening, but this is not useful in the discrete context (see Fig. 1). Also, it can be useful to have thick cuts for better representing the holes. For this reason, we evaluate the homotopic opening by visual inspection and we do not try to define a unique or optimal homotopic opening.



Fig. 1. Two homotopic openings (black) for the same object (gray) with the same size. This example shows why we cannot judge a homotopic opening by its size.

4 Computing Homotopic Openings

A simple algorithm for computing a homotopic opening, which is the base for the more elaborate algorithms that we introduce later, is the following. Let X be a discrete object, choose some point $x \in X$ and set $C = \{x\}$. Then find a point in $X - C$ which is simple for C and add it to C . Repeat this operation while there are such points. At the end C is obviously contractible and thus $X - C$ is a homotopic opening for X . Note that, according to [10], C is a homotopic dilation of $\{x\}$ constrained by X .

Figure 2 illustrates the output of such algorithm. While the cuts are thin, they are far too long. A more elaborate algorithm consists in using the distance transform.




Fig. 2. A 2D discrete object (gray) and a homotopic opening (black) obtained by choosing simple points at random.

The (Euclidean) distance transform of a discrete object X is the map $dt_X : X \rightarrow \mathbb{R}$ that assigns to each point of X its distance to $\mathbb{Z}^d - X$ or, in other words, its depth in the object. Hence, at the beginning we can set $C = \{x\}$ where x is a maximum for dt_X and then we give priority to points with higher dt_X value. By using the distance transform, we lead the propagation in order to obtain the cuts in the thinner parts of the object. In the following we describe two ways of doing this.

4.1 Random Propagation

A main concern when using the distance transform in the propagation of the contractible subset is how to deal with points with equal value. A simple idea for having an isotropic homotopic opening is to randomly pick a point among the points with equal distance transform value. Algorithm 1 describes this approach.

Algorithm 1. Homotopic opening with random propagation

Input: $X \subset \mathbb{Z}^d$
Output: Homotopic opening for X
 $C \leftarrow \{x\}$ for a random $x \in X$ such that $dt_X(x)$ is maximal;
 $S \leftarrow N_\alpha^*(x) \cap (X - C)$;
while $S \neq \emptyset$ **do**
 $x \leftarrow$ random point in S such that $dt_X(x)$ is maximal;
 $S \leftarrow S - \{x\}$;
 if x is simple for C **then**
 $C \leftarrow C \cup \{x\}$;
 $S \leftarrow S \cup (N_\alpha^*(x) \cap (X - C))$;
return $X - C$;

Complexity of Algorithm 1. Let n denote the number of points in the bounding box of the discrete object. The distance transform of X can be computed in $\mathcal{O}(n)$ [4,9]. Since every point is inserted into S each time a neighbor is added to C , then it is added at most $(3^d - 1)$ times. Thus, the *while* loop is executed at most $3^d n$ times. The set S can be implemented as a priority queue using the distance transform as key. Hence, insertion and deletion can be performed in logarithmic time while finding the deepest point is done in constant time. Consequently, Algorithm 1 has complexity $\mathcal{O}(3^d n(\log n + f(d)))$, where $f(d)$ denotes the complexity of checking if a point is simple in \mathbb{Z}^d .

For dimension $d \leq 3$, we can use a look-up table for recognizing simple points, and hence the complexity of Algorithm 1 is $\mathcal{O}(n \log n)$.

Figure 3 illustrates some results of Algorithm 1. The holey disk (top-left) presents short cuts, though they do not seem *straight* segments, specially the bottom-left cut. This kind of issue is treated in Sect. 4.3. The hollow ball (top-center) and the sculpture (top-right) present well-shaped cuts. Regarding the torus (bottom), we can appreciate two differentiated cuts. One is around the central hole, which is well-shaped, while the other one is far from looking like a ring around the cavity.




Fig. 3. Homotopic openings obtained with Algorithm 1.

This kind of problem is due to the propagation of the contractible subset. If we randomly propagate a contractible object inside a plane, we do not obtain a uniform disk, but a tree-shaped object. A 3D discrete object can contain a two-manifold-like set of points which are equidistant to the complement of the object, so the propagation has a similar behavior when it reaches their depth.

A concrete example is shown in Fig. 4 (top), which depicts a thickened torus—so its set of deepest points is a torus—and three steps in the propagation of the contractible subset. Hence, it is natural that Algorithm 1 produces strange homotopic openings on such objects. Note that this problem does not happen in 2D discrete objects.

4.2 Propagation by Layers

The previous example on the torus motivates the following algorithm. In this case, we treat the object by layers. The idea is that at each step we traverse the simple points in the outer boundary of the contractible subset to find the maximum distance transform value. Then we mark the points in the outer boundary whose distance transform value is close to that maximum and we try to add them to the object. This is described in Algorithm 2. This alternative strategy produces a more uniform propagation of the contractible subset. Figure 4 illustrates the difference between the propagations performed by Algorithms 1 and 2.




Fig. 4. Different steps of the propagation performed in Algorithms 1 (top) and 2 (bottom) for the same object.

Complexity of Algorithm 2. An upper bound for the number of executions of the *repeat* loop is n , since at each execution (except the last one), at least one point of $X - C$ is added to C . At each iteration we must find the outer boundary of the contractible subset and check if the points are simple, thus needing $\mathcal{O}(3^d n f(d))$ operations. The later traversal of the outer boundary does not affect the complexity of the iteration. Hence, the complexity of Algorithm 2 is $\mathcal{O}(3^d n^2 f(d))$.

Again, the complexity of Algorithm 2 is $\mathcal{O}(n^2)$ for dimension $d \leq 3$.

At each step, we put in the list L not only the deepest simple points in the outer boundary of the contractible subset, but all the points with distance transform value close to the maximum. Concretely, we consider the points whose

Algorithm 2. Homotopic opening with propagation by layers

Input: $X \subset \mathbb{Z}^d$
Output: Homotopic opening for X
 $C \leftarrow \{x\}$ for a random $x \in X$ with highest dt_X value;
repeat
 $m \leftarrow \max \{dt_X(x) \mid x \in N_\alpha^*(C) \cap X, x \text{ simple for } C\};$
 $L \leftarrow dt_X^{-1}([m - 1, m]) \cap N_\alpha^*(C);$
 foreach $x \in L$ **do**
 if x is simple for C **then**
 $C \leftarrow C \cup \{x\};$
until idempotency;
return $X - C;$

distance transform value is in the interval $[m - a, m]$, for $a = 1$. We have chosen this value of a after visual inspection of several examples, but a different value can also be considered.

Some results of Algorithm 2 are depicted in Fig. 5. The cuts of the holey disk (top-left) look slightly better. There is no remarkable improvement for the hollow ball (top-center) nor the sculpture (top-right). However, the homotopic opening of the torus (bottom) is definitely better.

In conclusion, Algorithm 2 can be better than Algorithm 1, but it has worse complexity.




Fig. 5. Homotopic openings obtained with Algorithm 2.

4.3 More Than Simple Points

Both previous algorithms add as many simple points as possible to the contractible subset, which produces thin cuts. However, leaving thicker cuts can be interesting for two reasons:

1. Cuts usually look like polygonal lines instead of discretized straight lines. When two wavefronts collide, the next added simple points depend strongly on their neighborhood. This produces a chain reaction that makes the cut be far from what we would expect.
2. One-dimensional thin cuts in big objects can be very difficult to see.

We propose a simple method for obtaining thick cuts. Instead of adding simple points, we check if we can add a discrete ball without changing the homotopy type of the contractible subset, and then we just add its center point (which must be a simple point). Given $r \in \mathbb{N}$, we consider $B_r = \{x \in \mathbb{Z}^d \mid \|x\|_2 \leq r\}$, the discrete ball of radius r . Hence, given a discrete object X , a contractible subset C , a point p and a radius r , we check if we can reduce $(C \cup (p + B_r)) \cap X$ to C via a sequence of simple points deletions. A simple way of doing this is to put the $\mathcal{O}(r^d)$ points of $((p + B_r) \cap X) - C$ in a list and repeatedly traverse it and remove simple points until stability. Hence, the complexity of checking if a ball is simple is $\mathcal{O}(r^{2d} f(d))$.

By using *simple balls* instead of simple points, Algorithms 1 and 2 are enriched with a parameter r which is the radius of the discrete ball. Thus, their complexities are $\mathcal{O}(3^d n (\log n + r^{2d} f(d)))$ and $\mathcal{O}(3^d n^2 r^{2d} f(d))$ respectively.

Intuitively, using simple balls instead of simple points should be sufficient to obtain thick cuts, but using a big radius can lead to the *tangency problem* illustrated in Fig. 6.

The discrete ball (red) is not simple for the contractible object (black) since it creates a one-pixel hole. Consequently, the center point (green) is not added. This kind of phenomenon is quite frequent when we use a big ball. It prevents a correct propagation of the contractible subset which can eventually produce a



Fig. 6. Tangency problem. The discrete ball (red) is not simple for the object (black). (Color figure online)




Fig. 7. Homotopic openings computed by Algorithms 1 (top) and 2 (bottom) with radiuses 10 (left), 20 (center) and 50 (right).

strange homotopic opening. This is illustrated in Fig. 7, which depicts the output of Algorithms 1 (top) and 2 (bottom) with radiuses 10, 20 and 50.

An attempt to avoid this kind of configurations is to center the ball not only in the point, but also in its $(2d)$ -neighborhood. If one of the $2d + 1$ balls is simple for the contractible subset, then we add the point. We call this the *relaxed test for simple balls*. Figure 8 shows how the homotopic openings of Fig. 7 are when we use a relaxed test for simple balls. We can appreciate that, after using this approach, the outputs of Algorithms 1 (top) and 2 (bottom) are almost indistinguishable. Sadly, this idea does not seem to work for 3D objects. Figure 9 depicts the homotopic openings computed by Algorithms 1 (top) and 2 (bottom) using a ball of radius 5. The relaxed test for simple balls yet produces strange results, but we can appreciate that Algorithm 2 behaves notably better.

We observe that using a positive radius prevents fronts from approaching too much, so we obtain quite smooth cuts. We can then continue to propagate the contractible subset with balls of smaller radiuses in order to obtain thin and smooth cuts. We have chosen to divide the radius by two at each step. Figure 10 shows the homotopic openings computed by Algorithm 1 when using initial radiuses 0, 2, 4 and 8. We can appreciate how the cuts look more and more well-shaped as we increase the initial radius.

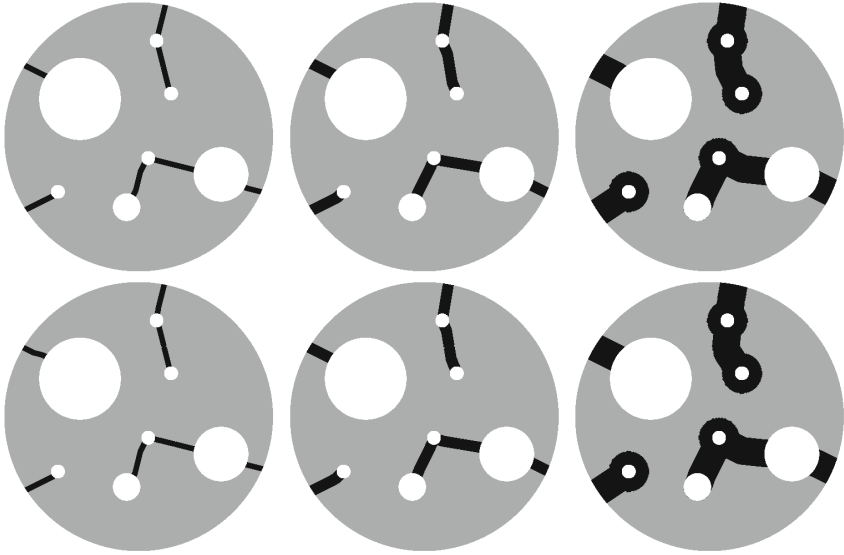


Fig. 8. Homotopic openings computed by Algorithms 1 (top) and 2 (bottom) with radiuses 10 (left), 20 (center) and 50 (right) and relaxed test for simple balls.




Fig. 9. Homotopic openings computed by Algorithms 1 (top) and 2 (bottom) with radius 5 and relaxed test for simple balls.




Fig. 10. Homotopic openings computed by Algorithm 1 with initial radiuses 0, 2, 4 and 8 and successive reductions. The cuts have been thickened for better visibility.

5 Conclusion and Future Work

This paper addresses the problem of erasing the holes of a discrete object by removing parts of it. We have defined the homotopic opening and we have presented two algorithms for computing it. The outputs of these algorithms are evaluated by visual inspection. While the second algorithm has worse complexity than the first one, it seems to produce better results. Then we have enriched these algorithms with a parameter corresponding to the radius of the balls used in the propagation of the contractible subset.

We do not obtain satisfactory results for big radiuses in 3D discrete objects, due to the tangency problem. We think that this problem deserves more attention. Our algorithms can also be generalized by considering other metrics and other structural elements for the propagation. Moreover, we suspect that the complexity of Algorithm 2 can be improved by smartly managing the layer update.

References

1. Aktouf, Z., Bertrand, G., Perrotton, L.: A three-dimensional holes closing algorithm. *Pattern Recogn. Lett.* **23**(5), 523–531 (2002)
2. Bertrand, G.: New notions for discrete topology. In: *Proceedings of the 8th International Conference Discrete Geometry for Computer Imagery, DCGI 1999, Marne-la-Vallée, France, 17–19 March 1999*, pp. 218–228 (1999)
3. Bertrand, G., Malandain, G.: A new characterization of three-dimensional simple points. *Pattern Recogn. Lett.* **15**(2), 169–175 (1994)
4. Coeurjolly, D., Montanvert, A.: Optimal separable algorithms to compute the reverse Euclidean distance transformation and discrete medial axis in arbitrary dimension. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(3), 437–448 (2007)
5. Couprie, M., Bertrand, G.: New characterizations of simple points in 2D, 3D, and 4D discrete spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(4), 637–648 (2009)
6. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
7. Janaszewski, M., Couprie, M., Babout, L.: Geometric approach to hole segmentation and hole closing in 3D volumetric objects. In: Bayro-Corrochano, E., Eklundh, J.-O. (eds.) *CIARP 2009*. LNCS, vol. 5856, pp. 255–262. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10268-4_30](https://doi.org/10.1007/978-3-642-10268-4_30)

8. Janaszewski, M., Couprie, M., Babout, L.: Hole filling in 3D volumetric objects. *Pattern Recogn.* **43**(10), 3548–3559 (2010)
9. Maurer Jr., C.R., Qi, R., Raghavan, V.: A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(2), 265–270 (2003)
10. Rueda, A., Acosta, O., Couprie, M., Bourgeat, P., Fripp, J., Dowson, N., Romero, E., Salvado, O.: Topology-corrected segmentation and local intensity estimates for improved partial volume classification of brain cortex in MRI. *J. Neurosci. Methods* **188**(2), 305–315 (2010)
11. Saha, P.K., Chaudhuri, B.B., Chanda, B., Majumder, D.D.: Topology preservation in 3D digital space. *Pattern Recogn.* **27**(2), 295–300 (1994)