ESTIMATING THE DISCRETE GEOMETRIC LUSTERNIK–SCHNIRELMANN CATEGORY

Brian Green — Nicholas A. Scoville — Mimi Tsuruga

————————————

ABSTRACT. Let K be a simplicial complex and suppose that K collapses onto L. Define n to be 1 less than the minimum number of collapsible sets it takes to cover L. Then the discrete geometric Lusternik–Schnirelmann category of K is the smallest n taken over all such L. In this paper, we give an algorithm which yields an upper bound for the discrete geometric category. We show our algorithm is correct and give several bounds for the discrete geometric category of well-known simplicial complexes. We show that the discrete geometric category of the dunce cap is 2, implying that the dunce cap is "further" from being collapsible than Bing's house whose discrete geometric category is 1.

1. Introduction

The goal of this paper is to introduce a computational algorithm to give bounds on the discrete geometric Lusternik–Schnirelmann (LS) category or just discrete geometric category of a simplicial complex. The discrete geometric category was introduced in [1] as a discrete analogue of the classical LS category for topological spaces [9]. Its main properties are summarized in Section 2. Our notion of the discrete geometric category is based on collapsibility, which has been used to study simple homotopy type [7]. See Definition 2.1 for the definition of

collapsibility. Let $f\colon K \to \mathbb{R}$ be a discrete Morse function in the sense of R. Forman [12], [13]. It was shown in [1] that the discrete geometric category of K bounds from below the number of critical points of f. There is much interest in the relationship between discrete Morse theory and collapsibility. For example, R. Ayala et al. [3] have used the collapse number of a 2-dimensional complex to study certain classes of discrete Morse functions. In addition, B. Benedetti and F.H. Lutz recently introduced so-called random discrete Morse theory [5]. They propose obtaining a discrete Morse vector by collapsing a complex until it contains no more free faces, removing a top dimensional face, and repeating. While this approach uses discrete Morse functions to provide an indicator of the complexity of a simplicial complex, we provide an alternative measure of the complexity of a simplicial complex by considering the minimum number of collapsible subcomplexes it takes to cover the complex. After introducing the discrete geometric category and reviewing its basic properties, we propose an algorithm to determine an upper bound for any finite type simplicial complex. We show that the algorithm yields an upper bound by constructing a collapsible cover and discuss some experiments for several well-known simplicial complexes. It immediately follows that our algorithm provides computational upper bounds for the topological complexity of a space. Recently there has been a renewed interest in Lusternik–Schnirelmann type invariants because of the work on M. Farber and others in topological complexity and robot motion planning [10], [11]. Since the discrete geometric category bounds above the (classical) category of the geometric realization of a space, certain upper bounds for topological complexity easily follow. These are summarized in Corollary 4.3.

Recall that a topological space is contractible if it has the homotopy type of a point. A simplicial complex which is collapsible always has a contractible geometric realization (Proposition 2.3) but the converse is not true. It is well known that Bing's house with two rooms [6] and the dunce cap [22] provide examples of complexes with contractible geometric realization, but which are not collapsible. In Proposition 4.6, we show that the discrete geometric category of the dunce cap is 2, while the discrete geometric category of Bing's house is only 1, and hence the dunce cap is in a certain sense further from being collapsible than Bing's house. This raises the question as to the existence of contractible simplicial complexes with arbitrarily large discrete geometric category.

2. Simplicial complexes and discrete geometric category

We begin by reviewing the basic terms used throughout this paper. We work with simplicial complexes because of the relative ease of implementing our algorithm, although there seems to be no reason in theory why our results will not carry over to the setting of a regular CW complex. All simplicial complexes

are assumed to be connected. Let $[n] = \{1, \ldots, n\}$. An *abstract (finite type) simplicial complex* K on $[n]$ is a collection of subsets of $[n]$ such that

(1) If $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$.

(2) $\{i\} \in K$ for every $i \in [n]$.

An element $\sigma \in K$ of cardinality $i+1$ is called an *i-dimensional face* or an *i-face* of K. The *dimension* of K, denoted $\dim(K)$, is the maximum dimension over all its faces. If $\sigma, \tau \in K$ with $\tau \subseteq \sigma$, then τ is a *face* of σ and σ is a *coface* of τ. We also say that τ is a *proper face* of σ if $\dim(\sigma) = \dim(\tau) + 1$. If $\tau \subseteq \sigma$ and $n = \dim(\sigma) - \dim(\tau)$, we say that τ is of *codimension n* with respect to σ. A face of K that is not contained in any other face is called a *facet* of K. A *(closed)subcomplex* L of K, denoted $L \subseteq K$, is a subset L of K such that L is also a simplicial complex. Denote by $\overline{\sigma}$ the smallest simplicial complex containing σ. For any $a_0, a_1, \ldots, a_i \in [n]$, define $a_0 a_1 \ldots a_i = \overline{\{a_0, a_1, \ldots, a_i\}}$. We are careful to use the term simplex for an element σ of K and the term complex for a subcomplex L of K. The *boundary* of σ, denoted $\mathrm{bd}(\sigma)$, is the collection of all its faces. Define $\mathrm{cbd}(\sigma) = \{\tau \in \mathrm{bd}(\sigma) : \tau \text{ is a proper face of } \sigma\}$. Clearly if $\dim(\sigma) = n$, then $|\mathrm{cbd}(\sigma)| = n + 1$.

DEFINITION 2.1. If K contains a pair of simplices σ, τ such that τ is a proper face of σ and τ has no other cofaces, then $K - \{\sigma, \tau\}$ is a simplicial complex called an *elementary (simplicial) collapse* of K. The simplicial complex K is said to *collapse* onto L if L can be obtained from K through a finite series of elementary collapses, denoted $K \searrow L$. If K collapses onto L, we also say that L *expands* to K, denoted $L \nearrow K$. The pair $\{\sigma, \tau\}$ is said to be a *free pair*, a term we will use to denote a pair that can either be collapsed or expanded with respect to K. In the case where $L = \{v\}$ is a single point, we say that K is *collapsible*.

The following is immediate from the definition of a free pair. It will be used to prove our algorithm is correct.

LEMMA 2.2. *Let L be a simplicial complex with σ a simplex of L and $\tau \in \mathrm{cbd}(\sigma)$. Then $K := L - \{\sigma, \tau\} \nearrow L$ if and only if $\mathrm{cbd}(\sigma) - \tau \subseteq K$.*

If K is a simplicial complex, let $|K|$ denote its geometric realization. Since an elementary collapse corresponds to a deformation retraction, we have the following proposition.

PROPOSITION 2.3 ([16, Proposition 6.14]). *If K and L have the same simple homotopy type, then $|K|$ and $|L|$ have the same homotopy type. In particular, if K is collapsible, then $|K|$ is contractible.*

DEFINITION 2.4. Let $L \subseteq K$ be a subcomplex. We say that *L has discrete geometric pre-category less than or equal to n in K*, denoted $\widetilde{\mathrm{dgcat}}_K(L) \leq n$,

if there exists $n + 1$ closed subcomplexes $\{U_0, \ldots, U_n\}$, $U_i \subseteq K$ for $0 \le i \le n$, each of which is collapsible such that $L \subseteq \bigcup_{i=0}^{n} U_i$. If $\widetilde{\mathrm{dgcat}}_K(L) \not< n$, then $\widetilde{\mathrm{dgcat}}_K(L) := n$.

DEFINITION 2.5. The *discrete geometric category* of L in K is defined by
$$\mathrm{dgcat}_K(L) := \min\{\widetilde{\mathrm{dgcat}}_K(L') : L \text{ collapses to } L'\}.$$
We write $\mathrm{dgcat}(K) := \mathrm{dgcat}_K(K)$.

REMARK 2.6. A word is in order concerning our definition. The need to define the pre-category of a complex is in order to guarantee that an elementary collapse does not decrease the discrete category.

EXAMPLE 2.7. Let $K = K_6$, the complete graph on 6 vertices. In other words, the 0-dimensional simplices of K consist of 6 vertices and the 1-dimensional simplices of K consist of all possible combinations of pairs of distinct 0-simplices. A collapsible cover of K is shown below using the three colors red, green, and blue in Figure 1. Hence $\widetilde{\mathrm{dgcat}}_K(L) \le 2$. By Proposition 3.1 (see below), we have that $\lceil 15/(6-1) \rceil - 1 = 2 \le \mathrm{dgcat}(K)$. We conclude that $\mathrm{dgcat}(K) = 2$.
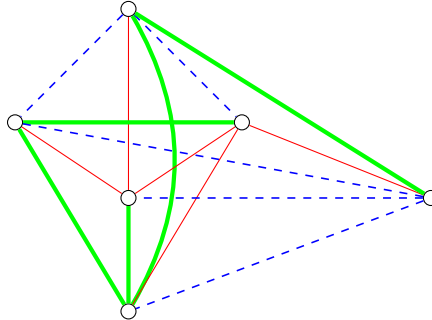

FIGURE 1. 3 collapsible sets distinguished by line type which cover K_6

REMARK 2.8. In general, if G is any 1-dimensional complex or graph, then the discrete geometric category coincides with the arboricity [14] of G. Nash–Williams has computed this invariant for all graphs [18]. In particular, if $G = K_n$, the complete graph on n nodes, then $\mathrm{dgcat}(K_n) = \lceil n/2 \rceil - 1$. We will make note of this fact in Table 1.

We note the relationship between the classical LS category and discrete geometric category.

PROPOSITION 2.9 ([1, Corollary 12]). *Let K be a simplicial complex. Then*
$$\mathrm{cat}(|K|) \le \mathrm{dgcat}(K).$$

Note that Example 2.7 provides an example of strict inequality for the above Proposition, as $\text{cat}(|K_6|) = 1 < \text{dgcat}(K_6) = 2$.

3. Combinatorial lower bound

Let K be a simplicial complex of dimension n or $n+1$, and $H_i(K)$ the i^{th} (unreduced) simplicial homology group of K. Let c_i^K denote the number of simplices of K of dimension i. Recall that the *Euler characteristic* of K is defined by $\mathcal{X}(K) = \sum_i (-1)^i c_i^K$. If β_i^K denotes the i^{th} Betti number of K, it is easy to show that $\mathcal{X}(K) = \sum_i (-1)^i \beta_i^K$ [19, Theorem 1.31]. Let $E(c^K) := c_0^K + c_2^K + \ldots + c_n^K$ and $O(c^K) := c_1^K + c_3^K + \ldots + c_{n\pm1}^K$.

PROPOSITION 3.1. *Let K be a simplicial complex of dimension n with c_i the number of simplices of K of dimension i, $0 \le i \le n$. If $E(c^K) - 1 \ge O(c^K)$, then $\lceil (E(c^K) - 1)/O(c^K) \rceil - 1 \le \text{dgcat}(K)$. If $E(c^K) - 1 \le O(c^K)$, then $\lceil O(c^K)/(E(c^K) - 1) \rceil - 1 \le \text{dgcat}(K)$.*

PROOF. We only show the first inequality, as the other one is similar. Let U be a collapsible subcomplex of K. Since $H_i(U) = 0$ for all $i \ge 1$ whenever U is collapsible [15, Corollary 2.3.5], it follows that any collapsible set must satisfy

$$\mathcal{X}(U) = 1 = \sum_i (-1)^i \beta_i^U = \sum_i (-1)^i c_i^U.$$

Rearranging this equation yields $O(c^U) = E(c^U) - 1$.

Now since $E(c^K) - 1 \ge O(c^K)$, the largest amount of odd-dimensional simplices in any collapsible cover is at most $E(c^U) - 1 = O(c^K)$. Thus in order to satisfy this equation, we need at least $\lceil (E(c^K) - 1)/O(c^K) \rceil$ collapsible sets, and $\lceil (E(c^K) - 1)/O(c^K) \rceil - 1 \le \widetilde{\text{dgcat}}(K)$.

If $K \searrow K'$ is any elementary collapse, then

$$\left\lceil \frac{E(c^K) - 1}{O(c^K)} \right\rceil - 1 \le \left\lceil \frac{E(c^K) - 1 - 1}{O(c^K) - 1} \right\rceil - 1 \le \widetilde{\text{dgcat}}_K(K').$$

Thus $\lceil (E(c^K) - 1)/O(c^K) \rceil - 1 \le \text{dgcat}(K)$. □

4. Algorithm

Let K be a simplicial complex. Let $\mathcal{H}$ be a graph encoding the incidence relations of the simplices of K; every node of $\mathcal{H}$ is a simplex of K and there is an edge between two simplices σ, τ whenever τ is a proper face of σ. This graph $\mathcal{H}$ is called the *Hasse diagram* of K [21]. By abuse of language, we will not distinguish between a simplex and a node of $\mathcal{H}$ representing the simplex. Let $\mathcal{H}(i)$ be the nodes of $\mathcal{H}$ corresponding to the i-simplices of K. We refer to $\mathcal{H}(i)$ as *level i*. Each node of $\mathcal{H}$ is equipped with an on/off switch consisting of three colors: red, green, and black. A node colored red means that it is not in the

cover $\mathcal{U}$ nor the current collapsing set U, a node colored green means that it is in the current collapsible set U, and a node colored black means that it is in the cover $\mathcal{U}$. Note that if a node is colored green or black, its red switch must be off. This fact will be used but not stated below. A node can be both black and green. Let $\mathcal{H}_r$ denote the nodes of $\mathcal{H}$ colored red. If $v \in \mathcal{H}(i+1)$, let $N^i(v)$ be the set of all green nodes on level i connected to v by an edge in $\mathcal{H}$ (i.e. a neighbour of v); then $N^i(v) = \{u \in \mathcal{H}(i) : u \text{ is green}, u \text{ is a proper face of } v\}$. Define the *expansion set* in row $i+1$ by $E(i+1) = \{v \in \mathcal{H}(i+1) : |N^i(v)| = i+1\}$; $E(i+1)$ collects all $(i+1)$-simplices σ such that all but one of its proper faces cbd(σ) are colored green. The *critical expansion set* in row $i+1$ is defined by $CE(i+1) = \{v \in E(i+1) : v \text{ is red}\}$. The paper [5] by Benedetti and Lutz serves as inspiration for our algorithm.

ALGORITHM 4.1. Discrete geometric category upper bound.
Input: A non-empty connected simplicial complex K.
Output: A collapsible cover $\mathcal{U}$ of K.

(1) Set $\mathcal{U} = \emptyset$ and build the Hasse diagram $\mathcal{H}$ of K. Color all nodes red.
(2) Set $U = \emptyset$.
(3) Pick a random red facet σ such that σ has maximum dimension over all red facets. For every $\tau \subseteq \sigma$ of any dimension, color τ green.
(4) Initialize $i = 0$.
(5) If $E(i+1) = \emptyset$, go to step 6. If $CE(i+1) = \emptyset$, choose a random $v \in E(i+1)$. Otherwise, choose a random $v \in CE(i+1)$. Color v (and all $\tau \subseteq v$) and its unique non-green proper face u on level i (and all $\tau \subseteq u$) green. Repeat step 5.
(6) Increment $i = i+1$. If $i = \dim(K)$, go to step 7. Otherwise go to step 5.
(7) Add all green nodes to U. Color every node in U black and turn off green. Add U to $\mathcal{U}$. If $\mathcal{H}_r = \emptyset$, then terminate algorithm. Otherwise, go to step 2.

The set $\mathcal{U}$ obtained in the above algorithm is a collapsible cover of K so that dgcat$(K) \leq |\mathcal{U}| - 1$. Since the complex induced by a facet of U is added in step 3, it follows that it will take at most the number of facets of K iterations of the algorithm to find a collapsible cover of K, and thus the algorithm will terminate.

The idea behind the algorithm is to determine whether or not an expansion is possible from the information provided by the Hasse diagram. The algorithm begins by picking a random top dimensional subcomplex, and begins to perform elementary expansions by expanding along as many 0-simplices as possible, as many 1-simplices as possible, etc. If all of the level i neighbours of a node on level $i+1$ are colored green except one neighbour, this means that all the boundary elements of codimension 1 except one of the corresponding simplex are in the set U, and hence we may perform an elementary expansion. A node with color

red has not been added to the cover yet, so preference is given to expanding along those nodes on level $i+1$ which are colored red. Since performing a finite number of elementary expansions can be undone by performing elementary collapses in reverse order, the subcomplex we obtain at the end of one full iteration of the algorithm is collapsible. Formally, we have the following.

PROPOSITION 4.2. *Algorithm* 4.1 *returns a set of subcomplexes* $U_0, \ldots, U_n$ *of a complex* K *such that each* U_i *is collapsible and* $\bigcup U = K$.

PROOF. We first show that any U obtained from Algorithm 4.1 is collapsible by induction. According to step 3, $U = \overline{\sigma}$ which is clearly collapsible. Assume that U is collapsible going into step 5. If $E(i + 1) = \emptyset$ and we end up in step 7, then we are done. Otherwise, we end up back in step 5 so assume that $E(i + 1) \neq \emptyset$ and choose a random $u \in E(i + 1)$ or $CE(i + 1)$. By definition of these sets, $|N^i(u)| = i + 1$ so that $i + 1$ boundary simplices of v of dimension i are in U and the u found in step 5 is not in U. In other words, $\mathrm{cbd}(v) - u \subseteq U$. By Lemma 2.2, $\{u, v\}$ is a free pair of U so that $U \nearrow U \cup \{u, v\}$ is an elementary expansion. Thus U is collapsible. Now let $\sigma \in K$. Since σ is in a collapsible U if and only if $\sigma \notin \mathcal{H}_r$ and the algorithm terminates only when $\mathcal{H}_r = \emptyset$, it follows that there exists $U \in \mathcal{U}$ such that $\sigma \in U$. Hence $\bigcup U = K$. □

Because the above algorithm gives an upper bound for the discrete geometric category, Proposition 2.9 along with the well-known relations between classical LS category and topological complexity yields the following:

COROLLARY 4.3. *Let* $\mathrm{TC}(X)$ *denote the topological complexity of space* X. *If* K *is a simplicial complex then*:
  (a) $\mathrm{TC}(|K|) \leq 2 \cdot \mathrm{dgcat}(K) - 1$ [10, Theorem 5],
  (b) $\mathrm{TC}(|K|) \leq \mathrm{dgcat}(K \times K)$ [11, Lemma 9.2].

4.1. Analysis of 1-dimensional case. This section is devoted to analyzing Algorithm 4.1 in the special case where $K = G$ is a 1-dimensional connected simplicial complex (i.e. a graph). We first note that each finding of a collapsible set in Algorithm 4.1 is equivalent to an implementation of Prim's algorithm to find a minimum weighted spanning tree [2, p. 125]. Indeed, label any red edge -1 and any black edge 0. Then Algorithm 4.1 finds a spanning tree by choosing an edge with minimum weight at each iteration of step 5, which is precisely Prim's algorithm. This guarantees that each collapsible cover adds the most uncovered edges possible at that iteration. This fact will be used below. Let $E(G)$ denote the set of edges (facets) of G. Since a graph is collapsible if and only if it is a tree, we use T_i to denote a collapsible set. For $n \in \mathbb{Z}^{>0}$, let $H(n)$ denote the nth harmonic number i.e. $H(n) := \sum_{i=1}^{n} \frac{1}{n}$ and set $H(0) = 0$. The following proof is nearly identical to [8, Theorem 35.4], but we include it here for completeness.

PROPOSITION 4.4. *Let G be a 1-dimensional connected simplicial complex. Then Algorithm 4.1 is an $H(|E(G)|)$-approximation algorithm for $\mathrm{dgcat}(G)$. In other words, if $\mathcal{U}$ is any set obtained in Algorithm 4.1, then*

$$|\mathcal{U}| \leq H(|E(G)|) \cdot \mathrm{dgcat}(G).$$

PROOF. Let $\mathcal{U} = \{T_0, \ldots, T_n\}$ be a collapsible cover of G obtained in Algorithm 4.1. If an edge $e \in E(T_i)$ appears for the first time in T_i so that $e \notin E(T_j), 0 \leq j < i$, define the *weight* of e by

$$\omega(e) = \frac{1}{|E(T_i) - (E(T_0) \cup \ldots \cup E(T_{i-1}))|}.$$

Then $|\mathcal{U}| = \sum\limits_{e \in E(G)} \omega(e)$. For an optimal cover $\mathcal{U}^*$ of G, we have

$$\sum_{T \in \mathcal{U}^*} \sum_{e \in E(T)} \omega(e) \geq \sum_{e \in E(G)} \omega(e)$$

and hence

$$|\mathcal{U}| \leq \sum_{T \in \mathcal{U}^*} \sum_{e \in E(T)} \omega(e).$$

Now let $T \subseteq G$ be any tree. Let $u_i := |E(T) - (E(T_0) \cup E(T_1) \cup \ldots \cup E(T_i))|$ with $u_{-1} = |E(T)|$ and k the least index such that $u_k = 0$. Clearly $u_{i-1} \geq u_i$ and $u_{i-1} - u_i$ edges of T are covered for the first time by T_i. Thus

$$\sum_{e \in E(G)} \omega(e) = \sum_{i=0}^{k} (u_{i-1} - u_i) \cdot \frac{1}{|E(T_i) - (E(T_0) \cup E(T_1) \cup \ldots \cup E(T_{i-1}))|}.$$

Since the maximum amount of non-covered edges are covered by $E(T_i)$ as noted above, we have that

$$|E(T_i) - (E(T_0) \cup E(T_1) \cup \ldots \cup E(T_i))|$$
$$\geq |E(T) - (E(T_0) \cup E(T_1) \cup \ldots \cup E(T_{i-1}))| = u_{i-1}.$$

We have

$$\sum_{e \in E(G)} \omega(e) = \sum_{i=0}^{k} (u_{i-1} - u_i) \cdot \frac{1}{|E(T_i) - (E(T_0) \cup E(T_1) \cup \ldots \cup E(T_{i-1}))|}$$

$$\leq \sum_{i=0}^{k} (u_{i-1} - u_i) \cdot \frac{1}{u_{i-1}} = \sum_{i=0}^{k} \sum_{j=u_i+1}^{u_{i-1}} \frac{1}{u_{i-1}}$$

$$\leq \sum_{i=0}^{k} \sum_{j=u_i+1}^{u_{i-1}} \frac{1}{j} = \sum_{i=0}^{k} \left(\sum_{j=1}^{u_{i-1}} \frac{1}{j} - \sum_{j=1}^{u_i} \frac{1}{j} \right)$$

$$= \sum_{i=0}^{k} H(u_{i-1}) - H(u_i) = H(u_{-1}) - H(u_k) = H(|E(T)|).$$

To complete the proof, we see that

$$|\mathcal{U}| \leq \sum_{T \in \mathcal{U}^*} \sum_{e \in E(T)} \omega(e) \leq \sum_{T \in \mathcal{U}^*} H(|E(t)|)$$

$$\leq |\mathcal{U}^*| \cdot H(|E(t)|) = \mathrm{dgcat}(G) \cdot H(|E(t)|). \qquad \square$$

Remark 4.5. It is not always the case that Algorithm 4.1 adds the maximum amount of facets in each element of the cover. Indeed, consider the simplicial complex K on the set $\{1, 2, 3, 4\}$ with all six 1-dimensional simplices and 2-simplices given by $\{1, 2, 4\}$ and $\{1, 3, 4\}$. Then step 3 of Algorithm 4.1 could begin by adding facet $\{1, 2, 4\}$ followed by 1-simplex $\{2, 3\}$ in the very first implementation of step 5. At this point, there are no free 1-dimensional faces, so the Algorithm moves on to step 6. Now there are also no free 2-dimensional faces, and we obtain one element in a cover of K, and this element contains only one facet of K. However, had the algorithm chosen the 1-simplex $\{1, 3\}$ instead of $\{2, 3\}$ in step 3, then $\{1, 3, 4\}$ would have been a free face and thus added to the cover when the Algorithm was searching for free 2-simplices. Hence, the algorithm does not always produce the maximum number of facets in each element of a cover.

4.2. Computations. In this section we present and discuss some experiments we performed on `polymake` [20] using Algorithm 4.1 to estimate the discrete geometric category of several well known simplicial complexes. Our experiments were run on a quad-core Intel® Xeon® X3460, 2.8 GHz, with 16GB of RAM. A downloadable version of our program (^{1}) is available and will be available as part of the `polymake` distribution from version 2.14.

A summary of our results is listed below in Table 1. The lower bound in each row is obtained by using a combination of Proposition 3.1 and other arguments discussed below. The "best found" column lists the smallest value obtained by running the `polymake` implementation of Algorithm 4.1.

Because the discrete geometric category depends on the particular simplicial structure chosen, the list of facets for these complexes may be found in the Simplicial Complex Library website (^{2}). In particular, it should be noted that we do not necessarily expect a complex with the topology of a sphere, such as the 3-sphere with a knotted triangle, to have discrete geometric category of 1, nor do we necessarily expect a complex with the topology of a ball, such as the 3-ball with a knotted hole, to have discrete geometric category of 0.

Our algorithm is accurate for many complexes. The problem of whether or not a complex is collapsible is undecidable [17], so one does not expect to be able

---

(^{1}) `http://webpages.ursinus.edu/nscoville/research-papers.html`

(^{2}) `http://infoshako.sk.tsukuba.ac.jp/∼hachi/math/library/index_eng.html`, accessed on February 12, 2015.

Complex K	theor lower bound	best found	dgcat(K)
Bing's house with 2 rooms	1	1	1
Dunce hat	2	2	2
Lockeberg's 4-polytope	1	1	1
Mani and Walkup's 3-sphere	1	1	1
3-sphere with knotted triangle	1	1	1
Projective plane ($\mathbb{R}P^2$)	2	2	2
Rudin's 3-ball	0	0	0
3-ball with a knotted hole	0	2	?
Non-PL 5-sphere	1	3	?
Poincaré sphere	1	8	?
K_5	2	2	2
K_{10}	4	4	4
K_{20}	9	9	9
K_{50}	24	24	25
K_{100}	49	49	50

TABLE 1. Summary of experiments. The second column refers to the best, i.e. smallest, value obtained in our experiments over many runs.

to compute the discrete geometric category in general. Our algorithm avoids the pitfalls of the undecidability problem by simply claiming to provide an upper bound.
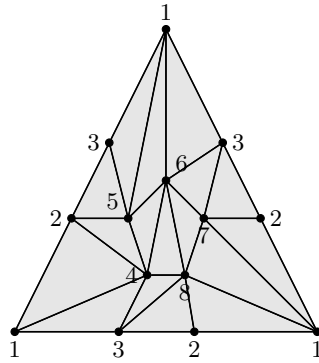

FIGURE 2. A triangulation of the dunce cap

Bing's house with two rooms [6] and the dunce cap have no free faces (and hence are not collapsible) but have contractible geometric realization. Because

these two complexes have no free faces, their discrete geometric category is bounded below by 1. In addition, Proposition 2.3 implies that $\text{cat}(|K|) \leq \text{dgcat}(K)$, where cat is the classical LS category. Hence $\text{cat}(|\mathbb{R}P^2|) = 2 \leq \text{dgcat}(\mathbb{R}P^2)$. As mentioned in Remark 2.8, the actual discrete geometric category for graphs was computed by Nash–Williams.

We now show that $2 \leq \text{dgcat}(D)$ where D is the triangulation of the dunce cap given below.

Call any edge of the edges $1\,2, 1\,3$, or $2\,3$ *formal* and any facet containing a formal edge formal.

PROPOSITION 4.6. *Let D be the dunce cap given by the triangulation above. Then $\text{dgcat}(D) = 2$.*
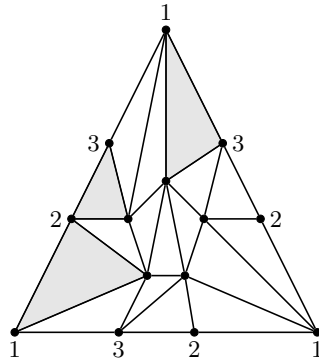

FIGURE 3. If U contains all three formal edges, then U cannot be collapsible.

PROOF. By the table above, computational results have shown that $\text{dgcat}(D) \leq 2$. Using the above labeling, we show by contradiction that $\text{dgcat}(D) > 1$, which yields the result. Suppose that $D = U \cup V$ with U, V collapsible subcomplexes of D. We will utilize the fact discussed leading up to Proposition 3.1 that a necessary condition for collapsibility is that a complex satisfy Euler's formula $v + f - 1 = e$. Since D is composed of 9 formal facets, at least one of U, V must contain 5 such formal facets, say U. We first claim that if U contains at least three formal facets with $1\,2, 1\,3$, and $2\,3$ in their boundary, then either U has non-trivial homology or $U = D$.

The configuration satisfies $6 + 3 - 1 \leq 9$, which implies that the complex is not contractible and hence not collapsible. In order to satisfy Euler's equation (again, a necessary condition for U to be collapsible), we must add at some point add either a node or a face without adding an edge. Since a collapsible set needs to be connected, the addition of any node will come with the addition of an edge. Every addition of a face will add either a face and an edge, a face and 2 edges, or
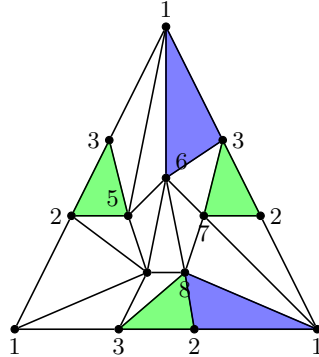
FIGURE 4. The minimum collection of facets of U are colored blue and the minimum collection of facets of V are colored green. The addition of $6\,8\,7$ to U or V containing the above blue and green facets, respectively, yields a cycle which cannot be killed.

a face along with 2 edges and a node. In any case, at least one edge is added for every other simplex so that $v + f - 1 = e$ will never be satisfied unless $U = D$, which is not collapsible.

Hence assume that none of the formal facets in U contain the edge $2\,3$ as the other two cases are similar. Since U does not contain $2\,3$, V contains $2\,3\,8, 2\,3\,7$, and $2\,3\,5$.

Furthermore, suppose that U contains $1\,3\,6$, as a similar analysis shows that if V contains $1\,3\,6$, then one of U, V cannot be collapsible. If U also contains $1\,2\,8$, then it is easy to see that with U and V containing at least the facets mentioned above, the facet $6\,7\,8$ will always create a cycle in U and in V. An argument as above considering the need to satisfy Euler's equation then shows that the cycle cannot be killed without creating all of D.

Otherwise, V must contain $1\,2\,8$. Then U contains $1\,7\,8$ for otherwise the addition of $7\,8$ would create a an unkillable cycle in V. So we must place $1\,7\,8$ in V. But then facet $6\,8\,7$ will create a cycle in both U and V, and again an Euler equation argument shows this cannot be killed without creating D. This homologically nontrivial cycle can easily been seen from Figure 3. Thus D cannot be written as the union of two collapsible sets and $1 < \mathrm{dgcat}(D)$. $\qquad\square$

REMARK 4.7. As an alternative to show that $\mathrm{dgcat}(D) \leq 2$, one could use the discrete Lusternik–Schnirelmann theorem [1] which says that if $f\colon K \to \mathbb{R}$ is a discrete Morse function with m critical values, then $\mathrm{dgcat}(K) + 1 \leq m$. There is a discrete Morse function $g\colon D \to \mathbb{R}$ with exactly 3 critical values [4] so that by the discrete LS theorem, $\mathrm{dgcat}(D) + 1 \leq 3$. As computed above, $\mathrm{dgcat}(D) = 2$ so this provides an example of equality in the discrete LS theorem.

Although the above evidence does not warrant a conjecture, Proposition 4.6 suggests that the existence of a contractible complex with discrete geometric category any positive integer is worth investigating; that is, given a positive integer n, does there exist a contractible simplicial complex $A(n)$ such that $\mathrm{dgcat}(A(n)) = n$? Bing's house with two rooms and the dunce cap answer the question in the affirmative for $n = 1$ and $n = 2$, respectively.

References

[1] S. Aaronson and N. A. Scoville, *Lusternik–Schnirelmann category for simplicial complexes*, Illinois J. Math. **57** (2013), no. 3, 743–753.

[2] G. Agnarsson and R. Greenlaw, *Graph theory: modeling, applications, and algorithms*, Pearson Prentice Hall, Upper Saddle River, NJ, 2007.

[3] R. Ayala, D. Fernández-Ternero and J.A. Vilches, *Perfect discrete Morse functions on triangulated 3-manifolds*, Computational Topology in Image Context, Lecture Notes in Computer Science, vol. 7309, 2012, 11–19.

[4] ———, *Perfect discrete Morse functions on 2-complexes*, Computational Topology in Image Context **33** (2012), 1495–1500.

[5] B. Benedetti and F.H. Lutz, *Random discrete Morse theory and a new library of triangulations*, Exp. Math. (to appear).

[6] R.H. Bing, *Some aspects of the topology of 3-manifolds related to the Poincaré conjecture*, Lectures on Modern Mathematics, Vol. II, Wiley, New York, 1964, 93–128.

[7] M.M. Cohen, *A course in Simple-Homotopy Theory*, Grad. Texts in Math., vol. 10, Springer-Verlag, New York, 1973.

[8] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein, *Introduction to algorithms*, third ed., MIT Press, Cambridge, MA, 2009.

[9] O. Cornea, G. Lupton, J. Oprea and D. Tanré, *Lusternik–Schnirelmann Category*, Mathematical Surveys and Monographs, vol. 103, Amer. Math. Soc., Providence, RI, 2003.

[10] M. Farber, *Topological complexity of motion planning*, Discrete Comput. Geom. **29** (2003), 211–221.

[11] ———, *Topology of robot motion planning*, Morse Theoretic Methods in Nonlinear Analysis and in Symplectic Topology, NATO Sci. Ser. II Math. Phys. Chem., vol. 217, Springer, Dordrecht, 2006, 185–230.

[12] R. Forman, *Morse theory for cell complexes*, Adv. Math. **134** (1998), no. 1, 90–145.

[13] ———, *A user's guide to discrete Morse theory*, Sém. Lothar. Combin., **48** (2002).

[14] F. Harary, *Graph Theory*, Addison-Wesley Publishing Co., Reading, Mass.-Menlo Park, Calif.-London, 1969.

[15] J. Jonsson, *Introduction to simplicial homology*, 2011, Available at the website `http://www.math.kth.se/~jakobj/doc/homology/homology.pdf` accessed on February 2, 2015.

[16] D. Kozlov, *Combinatorial Algebraic Topology*, Algorithms and Computation in Mathematics, vol. 21, Springer, Berlin, 2008.

[17] A. Markov, *The insolubility of the problem of homeomorphy*, Dokl. Akad. Nauk SSSR, **121** (1958), 218–220.

[18] C.St.J.A. Nash-Williams, *Edge-disjoint spanning trees of finite graphs*, J. London Math.
     Soc. **36** (1961), 445–450.
[19] V.V. Prasolov, *Elements of homology theory*, Graduate Studies in Mathematics, vol. 81,
     American Mathematical Society, Providence, RI, 2007, Translated from the 2005 Russian
     original by Olga Sipacheva.
[20] E. Gawrilow and M. Joswig, *polymake: a framework for analyzing convex polytopes*,
     Polytopes – Combinatorics and Computation (Gil Kalai and Günter M. Ziegler, eds.),
     Birkhäuser, 2000, pp. 43–74.
[21] H. Vogt, *Leçons sur la résolution algébrique des équations*, 1895.
[22] E.C. Zeeman, *On the dunce hat*, Topology **2** (1964), 341–358.

Brian Green and Nicholas A. Scoville
Department of Mathematics and Computer Science
601 E. Main Street
Collegeville, PA 19426, USA
    E-mail address: nscoville@ursinus.edu

Mimi Tsuruga
Straße des 17.Juni 136, Sekr. MA 8-1
13351 Berlin, GERMANY
    E-mail address: tsuruga@math.tu-berlin.de