A Topological Method for Three-Dimensional Symmetry Detection in Images

Jonah Brown-Cohen

Introduction

Recent work in computational topology has made it possible to compute the homology of geometric objects given only a description in the form of point cloud data–a finite collection of sample points lying on or near the space being sampled. The main technique is that of persistent homology developed in [8]. Additionally, a program for the computation of persistent homology known as PLEX has been developed by the computational topology research group at Stanford.

A portion of the previous applications of persistent homology to data analysis have been in the domain of image and shape processing. In [4], persistent homology techniques were used to determine that the set of high-contrast patches in natural images has the topology of a klein bottle. Other applications have focused on using persistent homology as an invariant or signature to distinguish shapes. In [5] the authors develop a hybrid method using techniques from differential geometry and persistent homology to classify shapes.

In this paper, we present a new application of persistent homology to approximate symmetry detection in sets of images. The motivation for our work is the desire to identify the symmetries of a three-dimensional object given a collection of two-dimensional images (e.g. photographs) of the object under different rotations. In this situation if two images of an object under two different rotations appear the same, then it seems likely that the object has some rotational symmetry. The question then is how to extract the symmetry information encoded by a set of images of an object.

To answer this question, we define the symmetrized rotation space of an object to be the set of all possible three-dimensional rotations of the object.

A set of 2D-images of an object under different rotations is therefore a set of samples from the objects symmetrized rotation space with one caveat: two-dimensional images or photographs are obtained by perspective projection of an object in three dimensions. So, the perspective projection must be taken into account when regarding images as samples.

Our method for detecting symmetries is thus to first perform theoretical calculations for the homology of the symmetrized rotation space of several model objects exhibiting different symmetry types. We then determine the effect of perspective projection on the homology of these spaces. Finally, given a set of images of an object we treat the images as a point-cloud sampled from the object's symmetrized rotation space and compute their persistent homology. By finding a match between the results of the computation and the theoretical results for some model symmetric object we can then determine the object's approximate symmetries.

In Section 1 we give background on the computation of persistent homology. In Section 2 we introduce the unit quaternions and the generalized quaternion groups. We then show that the unit quaternions are a double cover of $SO(3)$. In Section 3 we relate the topology of the symmetrized rotation space of an object to its group of symmetries and define a model for perspective projection. We then perform theoretical calculations for some model objects to determine the homology of their symmetrized rotation spaces both with and without perspective projection. In Section 4 we use PLEX to perform persistent homology computations on sets of images of the model objects and we compare the results to the theoretical calculations. We also present the unexpected discovery that perspective projection tends to eliminate the approximate topological information related to a certain flip rotational symmetry of some objects. This final result was first observed from computational results and later a theoretical explanation was derived, demonstrating the power of persistent homology in discovering entirely non-obvious topological relationships in data.

1 Persistent Homology

The motivation for the development of persistent homology is the question: how can the homology of point-cloud data be computed? More specifically, how can a simplicial complex be built from a discrete set of points so that the homology of the complex reflects the homology of the underlying space

from which the points were sampled. A natural way to construct a simplicial complex is given by the *Vietoris-Rips complex* described in [3]. The idea is to first decide on some closeness parameter ϵ and then to build a complex where the simplicies are simply the sets of points which are all pairwise close.

Definition 1.1. Let Z be a discrete set of points and let d be the metric on these points. The *Vietoris-Rips complex* with parameter ϵ, denoted by $W(Z, \epsilon)$, is the simplicial complex given by

1. The vertex set of $VR(Z, \epsilon)$ is the set of points in Z.
2. A k-simplex $(v_1, v_2, \ldots, v_k)$ is in $VR(Z, \epsilon)$ if and only if $d(v_i, v_j) \leq \epsilon$ for all $0 \leq i, j \leq k$.

The issue then becomes the question of how to choose epsilon to correctly reflect the topology of the underlying space. If ϵ is chosen to be too small, there may be cycles in the homology for which not all the vertices are within ϵ. So the complex would not capture these cycles. On the other hand, if ϵ is chosen to be too large, then a small cycle representing a non-trivial element of homology may have all of its vertices within ϵ of each other. So, the vertices would define a simplex and the cycle would appear to be bounding. Thus, both types of errors lose homological information and there is often no apriori obvious choice for ϵ given a data set.

The solution presented in [8] to the problem of choosing ϵ is to build a structure which simultaneously considers all possible values of ϵ. The idea is that as ϵ increases, non-bounding cycles in the homology of the complex $VR(Z, \epsilon)$ are created and then eventually filled in and destroyed. We would like to consider the cycles that persist across a wide range of values of ϵ to be the true cycles in the underlying topological space. Cycles which appear and then disappear rapdily are regarded as noise resulting from the sampling. This is the intuition behind persistent homology.

The first step toward computing persistent homology is to construct an increasing sequence of simplicial complexes $K^0 \subset K^1 \subset \ldots \subset K^n = K$ called a *filtered simplicial complex* or a *filtration*. Observe that for $\epsilon_1 < \epsilon_2$ we have the containment $VR(Z, \epsilon_1) \subset VR(Z, \epsilon_2)$. So if we simply take a sequence of increasing values of ϵ over some range the corresponding Vietoris-Rips complexes form a filtered complex. For each complex K^i in a filtration we have the standard chain complex of simplicial homology which we will denote by C_k^i with boundary maps ∂_k^i.

$$\cdots \longrightarrow C^i_{k+1} \xrightarrow{\partial^i_{k+1}} C^i_k \xrightarrow{\partial^i_k} C^i_{k-1} \xrightarrow{\partial^i_{k-1}} \cdots$$

The containments $K^i \subset K^{i+1}$ clearly induce chain maps $f^i : C^i_* \to C^{i+1}_*$ on these chain complexes. This allows for the combination of all the chain complexes for each simplicial complex in the filtration into one structure. In [8], the combined structure is called a *persistence complex* and is given by the following definition:

Definition 1.2. A *persistence complex* is a collection of chain complexes $\{C^i_*\}_{i\geq 0}$ over a ring R with maps $f^i : C^i_* \to C^{i+1}_*$ giving rise to the diagram:

$$0 \longrightarrow C^0_* \xrightarrow{f^0} C^1_* \xrightarrow{f^1} C^2_* \xrightarrow{f^2} \cdots$$

This is simply the natural way of attaching the simplicial chain complexes of a filtration. Now we turn to the definition of the persistent homology groups of a persistence complex. First, let $Z^i_k = \mathrm{Ker}(\partial^i_k)$ and let $B^i_k = \mathrm{Im}(\partial^{i+1}_k)$. These are the standard cycles and boundaries of the simplicial chain complexes C^i_*.

Definition 1.3. Let K be a filtered complex. The *p-persistent kth homology group* of K^i is

$$H^{i,p}_k = Z^i_k/(B^{i+p}_k \cap Z^i_k)$$

The definition makes sense because the chain complexes of K form a persistence complex, and the inclusion maps f^i mean that Z^i_k is a subgroup of C^{i+p}_k. Thus, the intersection in the denominator, as well as the quotient itself, are well-defined. If the filtered complex K is given by a sequence of Vietoris-Rips complexes $VR(Z, \epsilon_i)$ for increasing ϵ_i, then the above homology group corresponds to the set of cycles from $VR(Z, \epsilon_i)$ that persist for at least p more values of ϵ in the sequence. This precisely matches the intuition that we are interested in the cycles in homology that persist over a large range of values of ϵ. However, computation of the p-persistent kth homology for a filtered complex requires a compatible basis for Z^i_k and B^{i+p}_k.

We follow the method given in [8] for the construction of compatible bases. First define a *persistence module* to be a collection of R-modules M^i with homomorphisms $\varphi^i : M^i \to M^{i+1}$. Clearly the homology groups in

each dimension of a persistence complex form a persistence module. Now we consider the direct sum

$$\bigoplus_{i=0}^{\infty} M^i.$$

This sum has the structure of a graded module over the polynomial ring $R[t]$, where the action of t is given by

$$t \cdot (m^0, m^1, \ldots) = (0, \varphi^0(m^0), \varphi^1(m^1), \ldots)$$

That is, t moves elements up the filtration. Intuitively, the exponent of t in this $R[t]$ module tells us the time at which a given simplex σ first appeared. So if $t^j \cdot \sigma$ is in the ith gradation M^i. Then σ first appeared in the $i - j$th filtration. Thus, the graded $R[t]$ module defined in the sum above is a single module encoding all the information of the homology of the filtration.

The question of finding compatible bases for homology in a persistence module thus becomes a question about the structure of $R[t]$ modules. When R is not a field, for example $R = \mathbb{Z}$, then there is no known classification theorem for $R[t]$ modules. However, when $R = F$ for a field F, then $F[t]$ is a principal ideal domain, and a simple classification exists. In [8] the classification theorem for modules over a PID is used to show that the isomorphism classes of finitely generated graded $F[t]$-modules are in bijection with finite collections of intervals (i, j). Here we require $0 \leq i < j \in \mathbb{Z} \cup \infty$. In particular, the construction of this bijection means that an interval (i, j) corresponds to a basis element (i.e. a non-bounding cycle) of the homology groups of K that arises at time i and persists until time j when the cycle becomes a boundary. Thus, the set of intervals for a persistence module completely describe the persistent homology of the filtered complex K.

The algorithm given in [8] for computing these intervals simply modifies the standard reduction algorithm for homology to compute homology for the persistence module in each dimension. The output of the algorithm is, for each dimension k, the set of intervals (i, j) described above. This collection of intervals can be graphed to determine which intervals persist across a large range of the filtration. Such a graph is called a barcode. The interpretation of a barcode is that long intervals correspond to true generators of the homology of the underlying space, and short intervals correspond to noise resulting from sampling. The program JPLEX computes the barcode in each dimension given a filtered simplicial complex.

Though we have used the Vietoris-Rips complex as an example throughout this section because it is simple to understand, in practice it creates a filtered complex with two many simplicies for efficient computation. In order to speed up actual computations a different method described in [3] for constructing a complex from point-cloud data is often used. This method relies on selecting a subset of *landmark* points which will serve as the vertices of the complex, but using all the points in the cloud to determine simplices among the landmarks.

Definition 1.4. Let Z be a set of points and L be a subset of *landmark* points. For $z \in Z$ let $m_k(z)$ be the distance from z to the $(k+1)$st closest landmark in L. The witness complex with landmark set L and parameter ϵ, denoted by $W(Z, L, \epsilon)$, is the simplicial complex given by

1. The vertex set of $W(Z, L, \epsilon)$ is the set of points in L.
2. A k-simplex $(v_1, v_2, \ldots, v_k)$ is in $W(Z, L, \epsilon)$ if and only if all of its faces are in $W(Z, L, \epsilon)$ and there exists a *witness* $z \in Z$ such that $d(v_i, z) \leq m_k(z) + \epsilon$ for all $0 \leq i \leq k$.

The idea is to reduce the number of simplices by choosing a small landmark set, but to retain the information of all the points in the cloud by letting any point serve as a witness. Again we have the inclusion $W(Z, L, \epsilon_1) \subset W(Z, L, \epsilon_2)$ for $\epsilon_1 < \epsilon_2$, so we may construct a filtered simplicial complex from a set of witness complexes corresponding to a sequence of increasing values of ϵ. For our computations in Section 4 we will use witness complexes to triangulate point-clouds.

2 The Unit Quaternions and SO(3)

In this section we give some mathematical background which will be useful for the computations in Section 3. We begin by giving a surjective homomorphism from the unit quaternions to $SO(3)$ which has kernel $\{\pm 1\}$. We closely follow the construction given in [1]. First we define the quaternions, which we will denote by $\mathbb{H}$, to be a four dimensional vector space over $\mathbb{R}$ with basis $\{1, i, j, k\}$. That is, for $a, b, c, d \in \mathbb{R}$ an element $\omega \in \mathbb{H}$ is given by

$$\omega = a + bi + cj + dk$$

To define the group multiplication in $\mathbb{H}$ we define multiplication of the basis vectors and extend it by distributivity to all of $\mathbb{H}$. The basis vector 1 is the multiplicative identity and for i, j and k we define the relations

$$\begin{aligned} i^2 = j^2 = k^2 &= -1 \\ ij &= -ji = k \\ jk &= -kj = i \\ ki &= -ik = j \end{aligned}$$

It is a straightforward calculation to check that these relations make $\mathbb{H} - \{0\}$ into a group. Now we define quaternion conjugation in $\mathbb{H}$ to be a natural extension of complex conjugation.

Definition 2.1. Let $\omega = a + bi + cj + dk$ then the quaternion conjugate of ω is

$$\omega^* = a - bi - cj - dk$$

It is easy to check that $(\alpha\beta)^* = \beta^*\alpha^*$. To further extend the analogy between conjugation of quaternions and complex conjugation we note that

$$\omega^*\omega = a^2 + b^2 + c^2 + d^2$$

for $\omega = a+bi+cj+dk$. In particular, multiplication by the complex conjugate gives a non-negative real number which is zero only if ω is zero. Therefore, as in the case of the complex numbers, multiplication by the complex conjugate defines a norm on $\mathbb{H}$. We define the unit quaternions to be those $\omega \in \mathbb{H}$ with unit norm i.e. $\omega^*\omega = 1$. Thus the unit quaternions can be identified with S^3, the unit sphere in $\mathbb{R}^4$, by the map

$$\omega = a + bi + cj + dk \mapsto (a, b, c, d)$$

because $\omega^*\omega = 1$ implies that $a^2+b^2+c^2+d^2 = 1$. In light of this identification we will use S^3 to denote the unit quaternion group.

Now we relate S^3 to $SO(3)$ by giving an action of S^3 on $\mathbb{R}^3$. Let $V \subset \mathbb{H}$ be the pure imaginary quaternions. That is V is the subspace of vectors $v = bi + cj + dk$ with no component in the direction of the basis element 1. Equivalently, V is the set of $\omega \in \mathbb{H}$ with $\omega^* = -\omega$. Then we have the following proposition.

Proposition 2.2. *Let V be the pure imaginary quaternions and $O(3)$ the group of three-dimensional orthogonal transformations. Then the subspace $V \subset \mathbb{H}$ is isomorphic to $\mathbb{R}^3$ and the action of S^3 on V by conjugation defines a homomorphism $\phi : S^3 \to O(3)$.*

Proof. Clearly V is isomorphic as a vector space to $\mathbb{R}^3$ with the isomorphism given by

$$v = bi + cj + dk \mapsto (b, c, d)$$

This identification also equates the standard Euclidean norm on $\mathbb{R}^3$ with the norm given by quaternion conjugation in V since $v^*v = b^2+c^2+d^2$. Consider the action on V given by conjugation by an element of S^3 that is

$$\omega \cdot v = \omega v \omega^*$$

where multiplication on the right hand side is in $\mathbb{H}$. Observe that

$$(\omega v \omega^*)^* = \omega^{**} v^* \omega^* = \omega(-v)\omega^* = -\omega v \omega^*$$

Thus $\omega \cdot v \in V$, so conjugation by an element of S^3 is a group action on V. Futher for $a \in \mathbb{R}$ and $u, v \in V$ we have

$$\begin{aligned} \omega \cdot (au + v) &= \omega(au + v)\omega^* \\ &= a\omega u\omega^* + \omega v \omega^* \\ &= a(\omega \cdot u) + \omega \cdot v \end{aligned}$$

So conjugation by an element of S^3 is a linear transformation on V. Taking the norm of $\omega \cdot v$ yields

$$\begin{aligned} (\omega \cdot v)^*(\omega \cdot v) &= -(\omega v \omega^*)(\omega v \omega^*) \\ &= -\omega v^2 \omega^* \\ &= \omega(-v^2)\omega^* \\ &= \omega(v^*v)\omega^* \\ &= \omega\omega^*(v^*v) \\ &= v^*v \end{aligned}$$

So $\omega \cdot v$ has the same norm as v. Since we saw above that this corresponds to the Euclidean norm when V is identified with $\mathbb{R}^3$, we have that the action of S^3 on V preserves the Euclidean norm. Therefore conjugation by an element

of S^3 is an orthogonal linear transformation of $\mathbb{R}^3$. In other words, the action of S^3 on V defines a map $\phi : S^3 \to O(3)$.

To complete the proof we show that this map is a homomorphism. Let $\alpha, \beta \in S^3$ and $v \in V$.

$$\begin{aligned} \alpha\beta \cdot v &= \alpha\beta v\beta^*\alpha^* \\ &= \alpha \cdot \beta v\beta^* \\ &= \alpha \cdot (\beta \cdot v) \end{aligned}$$

It follows that ϕ is a homomorphism. □

With a little more work it can be shown that ϕ maps surjectively onto $SO(3) \subset O(3)$. We state this fact but we omit the proof here. Full details of the proof can be found in [1].

Proposition 2.3. *The homomorphism $\phi : S^3 \to O(3)$ is a surjection onto $SO(3) \subset O(3)$ with kernel $\{\pm 1\}$.*

We now introduce a class of finite subgroups of S^3 known as the *generalized quaternion groups* that will be useful in the theoretical calculations of the next section. Let n be an integer and let $\alpha, \beta \in S^3$ be defined by $\alpha = e^{i\theta}$ for $\theta = \pi/n$ and $\beta = j$.

Definition 2.4. The *generalized quaternion group* Q_{4n} is the subgroup of S^3 generated by α and β.

First we note that α and β have order $2n$ and 4 respectively. Observe that

$$\begin{aligned} \alpha^n &= e^{\pi i} = -1 \\ \beta^2 &= -1 \end{aligned}$$

It can also be shown that the only other relation between the two generators is $\beta^{-1}\alpha\beta = \alpha^{-1}$. These relations completely define Q_{4n}. This concludes our discussion of $SO(3)$ and the unit quaternions. In the next section we give theoretical calculations for the homology of symmetrized rotation spaces.

3 Homology of Projected Symmetrized Rotation Spaces

In this section we make the basic definitions and prove a simple proposition necessary to compute the homology of projected symmetrized rotation spaces. We then carry out the computation for two model symmetric objects: a pair of colored points, and a regular polygon with n sides.

The following definition formalizes the idea of the space of all possible rotations of an object in three dimensions.

Definition 3.1 (Symmetrized Rotation Space)**.** Let $SO(3)$ act by rotations on a set $S \subset \mathbb{R}^3$. The set $\{A(S)|A \in SO(3)\}$ of all possible rotated images of S is the *symmetrized rotation space* of S, denoted by X_S.

We will often refer to symmetrized rotation spaces as SR-spaces for brevity. Note that a single point in the space X_S is an entire rotated image $A(S)$ of the subset S. In general, we will use the Hausdorff metric for subsets of $\mathbb{R}^3$ as the metric on X_S. That is, the distance between two rotated copies of S in X_S is just the Hausdorff distance between the two as subsets of R^3. However any metric will do for which the action of $SO(3)$ on elements of X_S is continuous. This gives a topology on X_S.

Note that if a subset S has no rotational symmetries, then every element of $SO(3)$ will rotate S to a distinct subset of $\mathbb{R}^3$, and X_S will be homeomorphic to $SO(3)$. At the other extreme, if S is mapped to itself by all rotations, for example if S is a sphere centered at the origin, then X_S consists of a single point. The more interesting cases occur when S exhibits only some rotational symmetry, i.e. when the stabilizer of S in $SO(3)$ is a nontrivial, strict subgroup.

Images of a three dimensional object are perspective projections of that object onto some two dimensional plane in front of the viewer. For simplicty we will set the viewer to be at the point $(0, 0, 2) \in \mathbb{R}^3$ looking towards the origin and set the viewing plane to be defined by the equation $z = 1$. These two choices are sufficient to describe the perspective projection map P. A point $x \in \mathbb{R}^3$ is projected by drawing the line from x to the viewer at $(0, 0, 2)$, and finding the intersection, $P(x)$, with the viewing plane. The coordinates of this intersection can be computed using similar triangles, yielding:

$$P(x, y, z) = \left(\frac{x}{2 - z}, \frac{y}{2 - z}\right)$$

For symmetrized rotation spaces X_S, we will write $P(X_S)$ to denote the space gotten by using P to project each rotated image of the subset S in X_S to the viewing plane. Formally, $P(X_S) = \{P(A(S)) \mid A \in SO(3)\}$.

In order to compute the homology of $P(X_S)$ for a subset S it is useful to break up the computation into two steps. First, compute the homology of X_S and second, determine the effect of the map P on the space X_S to find the homology of $P(X_S)$. In particular, in the second step one should find the identifications–if any–that the map P makes. This division into two steps allows us to understand the distortions of the symmetries of an object caused by perspective projection. For the first step it seems intuitively that X_S should be some quotient of $SO(3)$ depending on the symmetries of S. The following lemma formalizes this intuition:

Lemma 3.2. *Let S be a compact subset of $\mathbb{R}^3$ and X_S be the corresponding symmetrized rotation space. Let G_S be the stabilizer of S in $SO(3)$. Then X_S is homeomorphic to $SO(3)/G_S$.*

Proof. By the orbit-stabilizer theorem of elementary group theory there is a bijection f from the quotient by the stabilizer, $SO(3)/G_S$, to the orbit of S under the action of $SO(3)$. In this case the space X_S is by definition the orbit of S. So the map f is given by

$$\begin{aligned} f : SO(3)/G_S &\to X_S \\ [A] &\mapsto A(S) \end{aligned}$$

where $[A] = AG_S$ is the left coset of G_S represented by A. Since $SO(3)$ is compact and the quotient map by G_S is continuous, $SO(3)/G_S$ is also compact. Further since the action of $SO(3)$ on X_S is continuous, the map f is continuous. Thus X_S is the continuous image of the compact space $SO(3)/G_S$, which implies that X_S is also compact. Now f is a continuous bijection between compact spaces, so f is a homeomorphism. □

Lemma 3.2 provides a general description of the topology of a symmetrized rotation space X_S. The second step of computing the identifications introduced by the projection to $P(X_S)$ must be handled on a case-by-case basis.

3.1 The Symmetrized Rotation Space of Two Colored Points

We now apply the technique described above to calculate the homology of the SR space of two colored points, as well as that of the corresponding projected SR space. Let $C = \{(1,0,0), (-1,0,0) \in \mathbb{R}^3\}$. Further, let us color the two points with different colors, say $r = (1,0,0)$ is red and $g = (-1,0,0)$ is green. Here coloring is simply a way of making the points distinguishable. In particular, it means that any of the 180 degree rotations that exchange r and g, are not elements of the stabilizer of S. In this case, we may think of $A(S) \in X_C$ as the ordered pair $(A(r), A(g)) \in \mathbb{R}^6$, and take our metric to be the standard Euclidean distance.

Proposition 3.3. *The symmetrized rotation space X_C is homeomorphic to S^2.*

Proof. Note the stabilizer of C is the set of all rotations around the x-axis. This is isomorphic to the two-dimensional rotation group $SO(2)$. By Lemma 3.2 the symmetrized rotation space X_C is homeomorphic to $SO(3)/SO(2)$. Note that any rotation in $SO(3)$ can be described by an axis v, and an angle θ of rotation around that axis. Since each rotation in $SO(2)$ is totally described by one angle θ, quotienting $SO(3)$ by $SO(2)$ kills the rotation angle θ leaving only the axis v. The set of all possible axes of rotation v is simply the unit sphere S^2. Thus X_C is homeomorphic to S^2. □

If we think of S^2 as the unit sphere in R^3 then the homeomorphism $X_C \to S^2$ can in fact be written explicitly as the map taking $\{r, g\} \mapsto r$. The proposition immediately implies that the homology of X_C is just the homology of S^2, that is:

$$H_n(X_C) \cong \begin{cases} \mathbb{Z} & \text{for } n = 0, 2 \\ 0 & \text{otherwise} \end{cases}$$

The next step is to determine the effect of the projection map P on X_C. We are specifically interested in the additional identifications that P makes on the space X_C, as this will in general determine the topology. The following result shows that if X_C is thought of as the sphere S^2, then P identifies the north pole with the south pole.

Proposition 3.4. *Let $\{r, g\} \in X_C$ for $r = (r_x, r_y, r_z)$ and $g = (g_x, g_y, g_z)$. Then the map $P : X_C \to P(X_C)$ identifies the points $\{(0,0,1),(0,0,-1)\}$ and $\{(0,0,-1),\ (0,0,1)\}$ and is injective for all other $\{r, g\}$.*

Proof. Suppose first that $r_x = r_y = 0$. Since the vector r is a rotation of $(1,0,0)$ it must have length 1, so $r_z = \pm 1$. Note both $\{(0,0,1),(0,0,-1)\}$ and $\{(0,0,-1),(0,0,1)\}$ project to $\{(0,0),(0,0)\}$ under P. That is P identifies the two points where $r_x = r_y = 0$.

Now consider the case where r_x and r_y are not both zero. Without loss of generality, assume $r_x \neq 0$. Suppose $\{P(r), P(g)\} = \{P(r'), P(g')\}$ for $\{r, g\}, \{r', g'\} \in X_C$. Then we have

$$\frac{r_x}{2 - r_z} = \frac{r'_x}{2 - r'_z} \tag{1}$$

$$\frac{r_y}{2 - r_z} = \frac{r'_y}{2 - r'_z} \tag{2}$$

$$\frac{g_x}{2 - g_z} = \frac{g'_x}{2 - g'_z} \tag{3}$$

$$\frac{g_y}{2 - g_z} = \frac{g'_y}{2 - g'_z} \tag{4}$$

Note that for $\{r, g\} \in X_C$ we have $r = -g$. This is because each (r, g) is a rotated copy of the pair $(1,0,0)$ and $(-1,0,0)$, and rotations preserve the length and dot product of two vectors. Now since $r = -g$ and $r' = -g'$, Equation (3) becomes

$$\frac{r_x}{2 + r_z} = \frac{r'_x}{2 + r'_z}$$

Solving the above equation for r'_x and plugging the result into (1) then gives

$$r_x \left(\frac{2 + r'_z}{2 + r_z} \right) = r_x \left(\frac{2 - r'_z}{2 - r_z} \right)$$

By assumption $r_x \neq 0$ so we may cancel the r_x from both sides. Simplifying then shows that $r_z - r'_z = r'_z - r_z$, which implies that $r_z = r'_z$. Substituting this into (1) and (2) cancels the denominators yielding $r_x = r'_x$ and $r_y = r'_y$ respectively. Thus $r = r'$. Since $g = -r$ and $g' = -r'$, we additionally have that $g = g'$, which implies that P is injective in this case.

□

Under the homeomorphism $X_C \to S^2 \subset \mathbb{R}^3$ which maps $\{r, g\} \mapsto r$, the two points identified by P map to $(1, 0, 0)$ and $(-1, 0, 0)$. That is, the proposition implies that the north and south pole of X_C thought of as the sphere S^2 are identified. The integer homology of $P(X_C)$ is then given by

$$H_n(P(X_C)) \cong \begin{cases} \mathbb{Z} & \text{for } n = 0, 1, 2 \\ 0 & \text{otherwise} \end{cases}$$

This completes the calculation of the homology of the projected SR space X_C. Although in this case it was straightforward to explicitly construct the topological space $P(X_C)$ as a 2-sphere with north and south pole identified, this is not generally true. The next example will show that even in relatively simple cases, such an explicit contruction is difficult, and the algebraic tools of homology and cohomology become much more useful.

3.2 The SR-Space of a Regular Polygon

Consider a regular polygon T with n vertices centered at the origin in $\mathbb{R}^3$. For convenience we also require that one vertex of T is at the point $(1, 0, 0)$. As usual the symmetrized rotation space of T will be written as X_T, and we will take the elements $x \in X_T$ to be n-tuples of vertices $(x^1, x^2, \ldots x^n) \in \mathbb{R}^{3n}$ representing rotated copies of T. For this object we will consider the vertices to be indistinguishable–a fact which must be reflected in the choice of metric.

In particular, treating elements as points in $\mathbb{R}^{3n}$ with the standard Euclid-den metric would be incorrect. For example the rotation A which cyclically permutes the vertices of the polygon clearly fixes T as a set, and so $A(T) = T$ in X_T. However the corresponding vertices between T and $A(T)$ are now off-set by one in the n-tuples of vertices given as coordinates. Thus the Euclidean distance between T and $A(T)$ thought of as points in $\mathbb{R}^{3n}$ will be nonzero even though $T = A(T)$. To remedy this, we define a polygon rotation metric d_T.

Definition 3.5. Let Σ_n, the permutation group on n elements, act on $x \in X_T$ by permuting the n vertices of the polygon x. Then for $x, y \in X_T$

$$d_T(x, y) = \min_{\sigma \in \Sigma_n} ||x - \sigma(y)||$$

where $|| \cdot ||$ is the Euclidean norm on $\mathbb{R}^{3n}$.

The intuition is that two polygons x and y should be near each other in X_T if each vertex of x is near a vertex of y. However, we might have that x_1 is near y_2, x_2 is near y_3 and so on, but that x_i is far from y_i for $i = 1$ to n. In this case the Euclidean distance between x and y as points in $\mathbb{R}^{3n}$ would be large. So we allow reordering of the vertices of the triangles by permutations in Σ_n to minimize the Euclidean distance. The following proposition shows that d_T is a metric.

Proposition 3.6. *The function $d_T : X_T \to \mathbb{R}$ is a metric.*

Proof. Evidently, $d_T(x, y) \geq 0$. Further, $d_T(x, y) = 0$ if and only if there is some permutation σ of the vertices of y such that for $i = 1$ to n the equality $\sigma(y_i) = x_i$ for holds (i.e. the vertices of x and y are identical as sets). This implies, $d_T(x, y) = 0$ if and only if $x = y$. Further note that $||x - \sigma(y)|| = ||\sigma^{-1}(x) - y||$ so d_T is symmetric, for

$$\begin{aligned} d_T(x, y) &= \min_{\sigma \in \Sigma_n} ||x - \sigma(y)|| \\ &= \min_{\sigma \in \Sigma_n} ||\sigma^{-1}(x) - y)|| \\ &= \min_{\sigma \in \Sigma_n} ||y - \sigma(x)|| \\ &= d_T(y, x) \end{aligned}$$

Finally, we show that the triangle inequality holds. For $x, y \in X_T$ let $\sigma_{x,y} \in \Sigma_n$ be the permutation which minimizes $||x - \sigma(y)||$. Then we have

$$\begin{aligned} d_T(x, y) + d_T(y, z) &= d_T(y, x) + d_T(y, z) \\ &= \min_{\sigma \in \Sigma_n} ||y - \sigma(x)|| + \min_{\sigma \in \Sigma_n} ||y - \sigma(z)|| \\ &= ||y - \sigma_{y,x}(x)|| + ||y - \sigma_{y,z}(z)|| \\ &\geq ||\sigma_{y,x}(x) - \sigma_{y,z}(z)|| \\ &= ||x - \sigma_{y,x}^{-1} \circ \sigma_{y,z}(z)|| \\ &\geq \min_{\sigma \in \Sigma_n} ||x - \sigma(z)|| \\ &= d_T(x, z) \end{aligned}$$

where the first inequaliy follows from the triangle inequality for the Euclidean metric on $\mathbb{R}^{3n}$. Thus d_T is a metric on X_T. □

Clearly the action of $SO(3)$ on X_T is continuous, because the distance between corresponding vertices of x and $A(x)$ for $A \in SO(3)$ can be made arbitrarily close to zero by choosing A to be a rotation by a sufficiently small angle. With such a metric d_T in hand, we now have a complete description of the SR-space X_T and can begin to compute it's homology.

Proposition 3.7. *Let D_{2n} be the dihedral group of order $2n$. Then the SR-space X_T is homeomorphic to $SO(3)/D_{2n}$.*

Proof. The polygon T has two distinct types of rotational symmetries. First there are the $n-1$ rotations by $2\pi/n$ about the z-axis which cyclically permute the n vertices. These are generated by the rotation

$$a = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

for $\theta = 2\pi/n$. The second class of symmetries are the rotations by π about an axis through one vertex of the polygon. Such rotations permute the other $n-1$ vertices if n is odd, and $n-2$ vertices if n is even. Let

$$b = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}$$

the rotation about the x-axis by π. Then the rotations $b, ab, a^2b, \ldots, a^{n-1}b$ correspond to rotations by π about the axes going through each of the n vertices.

Thus the two rotations a and b generate G_T, the stabilizer of T in $SO(3)$. Further, we have that $a^3 = b^2 = 1$ and it is easy to check that $bab^{-1} = a^{-1}$. This shows that G_T is simply D_{2n}, the dihedral group of order $2n$. By Lemma 3.2 we then have that the SR space X_T is homeomorphic to $SO(3)/D_{2n}$. □

What remains now is to calculate the homology $SO(3)/D_{2n}$ and to determine the effect of the projection map on the space X_T. To compute the homology, it is useful to first pass to S^3, the universal cover of $SO(3)$. As in Section 2, we will give S^3 the group structure of the unit quaternions. To this end we prove the following lemma.

Lemma 3.8. *Let Q_{4n} be the generalized quaternion group of order $4n$. Then the space $SO(3)/D_{2n}$ is homeomorphic to S^3/Q_{4n}.*

Proof. By Proposition 2.3 $SO(3)$ is a quotient of S^3, with quotient homomorphism $\phi : S^3 \to SO(3)$ having kernel $\{\pm 1\}$. Further, Q_{4n} is a subgroup generated by α and β satisfying $\alpha^n = \beta^2 = -1$ and $\beta\alpha\beta^{-1} = \alpha^{-1}$. So, the image of Q_{4n} under ϕ is simply D_{2n}.

Now let $\psi : SO(3) \to SO(3)/D_{2n}$ be the natural quotient map. Then, the quotient map $S^3 \to S^3/Q_{4n}$ factors through $SO(3)$ and is in fact simply the composition $\psi \circ \phi$. In particular, this gives a bijection

$$\begin{aligned} f : S^3/Q_{4n} &\to SO(3)/D_{2n} \\ [\omega] &\mapsto [\phi(\omega)] \end{aligned}$$

Since each of the maps ψ and ϕ are continuous and S^3 and $SO(3)$ are compact, f is a continuous bijection between compact spaces. So f is a homeomorphism. □

The advantage of viewing this space as S^3/Q_{4n} is that S^3 is simply connected, and since Q_{4n} is a finite group acting freely on S^3, the action is properly discontinuous. So the fundamental group of S^3/Q_{4n} is isomorphic to Q_{4n}. This immediately allows us to calculate the first dimensional homology of X_T. Further, the fact that S^3 is sphere allows us to use techniques from group cohomology–which relate the cohomology of groups to their action on contractible spaces such as infinite spheres–to compute the higher dimensional homology groups.

The following argument makes use of a variety of facts from both algebraic topology and group cohomology. However a full discussion of these facts is beyond the scope of this paper. Chapters one through three of [6] are a good reference for the theorems from algebraic topology, and chapter two of [1] contains all the requisite facts in group cohomology.

Proposition 3.9. *Let T be a regular polygon with n vertices centered at the origin. Then the homology of X_T is given by*

$$H_k(X_T)) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 3 \\ \mathbb{Z}/(4) & \text{for } k = 1,\ n \text{ odd} \\ \mathbb{Z}/(2) \times \mathbb{Z}/(2) & \text{for } k = 1,\ n \text{ even} \\ 0 & \text{otherwise} \end{cases}$$

Proof. Proposition 3.7 and Lemma 3.8 imply that X_T is homeomorphic to S^3/Q_{4n}. First we compute the first dimensional homology from the fundamental group. As we saw above, the action of Q_{4n} on S^3 is properly discontinuous. So since S^3 is simply connected, $\pi_1(X_T) \cong Q_{4n}$. So $H_1(X_T)$ is the abelianization of Q_{4n}. Quotienting by the commutator of the generators $\alpha, \beta \in Q_{4n}$ has the following effect on the second relation

$$\begin{aligned}\alpha^{-1} &= \beta^{-1}\alpha\beta \\ &= \beta^{-1}\beta\alpha \\ &= \alpha\end{aligned}$$

which implies that $\alpha^2 = 1$. Combining this with the first relation gives

$$\beta^2 = \alpha^n = \alpha^k$$

for $k = n \pmod 2$. If n is odd, $k = 1$ so $\alpha = \beta^2$ and the whole group is generated by β which has order 4. If n is even, $k = 0$ so $\beta^2 = 1$. This gives that the abelianization is generated by a pair of commuting elements each of order 2. In summary, we have $H_1(X_T) \cong \mathbb{Z}/(4)$ for n odd and $H_1(X_T) \cong \mathbb{Z}/(2) \times \mathbb{Z}/(2)$ for n even, as desired.

To compute the two higher homology groups H_2 and H_3, a different approach using an application of group cohomology is needed. First let $E = S^3 * S^3 * \dots$ be the infinite join of copies of S^3. Note that since the join of spheres $S^n * S^m$ is S^{n+m+1} the space E is simply S^∞. In particular we have that E is contractible. So the homology $H_*(E)$ has a single $\mathbb{Z}$ in dimension 0, and is trivial in every other dimension. Now consider the short exact sequence of chain complexes:

$$0 \longrightarrow C_k(S^3) \xrightarrow{i} C_k(E) \longrightarrow C_k(E, S^3) \longrightarrow 0$$

The corresponding long exact sequence of homology shows that $H_4(E, S^3) \cong \mathbb{Z}$ and every other relative homology gruop $H^k(E, S^3)$ is zero. Thus, the relative homology $H_k(E, S^3)$ is simply that of $H_k(E)$ suspended up four dimensions. That is, $H_k(E, S^3) \cong H_{k-4}(E)$. Now we note that the group $G = Q_{4n}$ acts freely on E by the action of G on the copy of S^3 in each dimension. So, since E is contractible, E/G is the classifying space of G.

Thus we may compute the cohomology $H^k(E/G)$ simply by computing the group cohomology $H^k(G)$, given algebraically by $\mathrm{Ext}^k_{\mathbb{Z}[G]}(\mathbb{Z}, \mathbb{Z})$. Further,

because $C_k(E, S^3)$ and $C_k(E)$ are both acted on freely by G and bounded below, it is a general fact in homological algebra that the isomorphism $H_k(E, S^3) \cong H_{k-4}(E)$ on homology induces an isomorphism:

$$H^k(E/G, S^3/G) \cong H^{k-4}(E/G)$$

In addition, the long exact sequence of cohomology for the pair $(E/G, S^3/G)$ gives

$$\cdots \longrightarrow H^k(E/G) \xrightarrow{i^*} H^k(S^3/G) \xrightarrow{\delta} H^{k+1}(E/G, S^3/G) \longrightarrow \cdots \qquad (5)$$

So given the group cohomology of G, we will know the first and last group in the above diagram, and will thus be able to determine $H^k(S^3/G)$ by exactness of the sequence.

We now compute the group cohomology of G. Let R be the ring $\mathbb{Z}[G]$. To compute $\mathrm{Ext}^k_R(\mathbb{Z}, \mathbb{Z})$ we use the following R-free resolution of $\mathbb{Z}$ which can be found in [7]:

$$\cdots \longrightarrow R \xrightarrow{d_4} R \xrightarrow{d_3} R^2 \xrightarrow{d_2} R^2 \xrightarrow{d_1} R \xrightarrow{\epsilon} \mathbb{Z} \longrightarrow 0$$

Let σ, σ_1 and σ_2 be elements of $\mathbb{Z}[G]$. Then the four maps for the resolution are given by:

$$\begin{aligned}
&d_1(\sigma_1, \sigma_2) = \sigma_1(\alpha - 1) + \sigma_2(\beta - 1) \\
&d_2(\sigma_1, \sigma_2) = (\sigma_1(1 + \alpha + \ldots + \alpha^{n-1}) + \sigma_2(\alpha\beta + 1), -\sigma_1(\beta + 1) + \sigma_2(\alpha - 1)) \\
&d_3(\sigma) = (\sigma(\alpha - 1), -\sigma(\alpha\beta - 1)) \\
&d_4(\sigma) = \sigma \sum_{g \in G} g
\end{aligned}$$

Now we apply the functor $\mathrm{Hom}_{\mathbb{Z}[G]}(\cdot, \mathbb{Z})$ and take the homology of the resulting complex to yield the group cohomology. The final result for dimensions zero through three is

$$H^k(E/G) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0 \\ 0 & \text{for } k = 1, 3 \\ \mathbb{Z}/(4) & \text{for } k = 2,\ n \text{ odd} \\ \mathbb{Z}/(2) \times \mathbb{Z}/(2) & \text{for } k = 2,\ n \text{ even} \end{cases}$$

Note that since $H^k(E/G, S^3/G)$ is just the above cohomology, suspended four dimensions up, we have that $H^k(E/G, S^3/G)$ is trivial for $k < 4$. Thus, the $k = 2$ portion of the long exact sequence (5) becomes

$$0 \longrightarrow H^2(E/G) \xrightarrow{i^*} H^2(S^3/G) \xrightarrow{\delta} 0$$

which implies that $H^2(S^3/G) \cong H^2(E/G)$.

Note $H^3(E/G) = 0$, and the relative homology group $H^4(E/G, S^3/G) \cong H^0(E/G) \cong \mathbb{Z}$. So, the $k = 3$ portion of (5) becomes

$$0 \xrightarrow{i^*} H^3(S^3/G) \xrightarrow{\delta} \mathbb{Z} \longrightarrow \cdots$$

Thus, δ is an injection into $\mathbb{Z}$. This implies that $H^3(S^3/G)$ is isomorphic to some ideal in $\mathbb{Z}$. However all ideals of $\mathbb{Z}$ are isomorphic to $\mathbb{Z}$ itself, so $H^3(S^3/G) \cong \mathbb{Z}$.

Now we apply the universal coefficient theorem of cohomology to convert $H^k(S^3/G)$ to the corresponding homology groups. The theorem implies that the torsion subgroups of cohomology are one dimension higher than those of homology. The only torsion present in $H^k(S^3/G)$ occurs when $k = 2$. Thus, we get $H_2(S^3/G) = 0$ and

$$H_1(S^3/G) \cong H^2(S^3/G) \cong \begin{cases} \mathbb{Z}/(4) & \text{for } n \text{ odd} \\ \mathbb{Z}/(2) \times \mathbb{Z}/(2) & \text{for } n \text{ even} \end{cases}$$

For the other dimensions $k = 0, 3$, the homology groups are the same as the cohomology groups. This completes the calculation. □

It remains to determine the effect of the projection map on the topology of X_T. The next proposition shows that in this case the map P is injective, and so $P(X_T)$ is homeomorphic to X_T.

Proposition 3.10. *The projection map $P : X_T \to P(X_T)$ is injective.*

Proof. Let $x^1, x^2, ..., x^n$ be the n vertices of $x \in X_T$, and let x^i_j be the jth coordinate of x^i in $\mathbb{R}^3$. Let $x, y \in X_T$ with $P(x) = P(y)$. We will show that $x = y$, which will then imply that P is injective.

Without loss of generality we may assume that the vertices of y have been reordered such that $P(x^i) = P(y^i)$ for all i. This gives us the equations

$$\frac{x^i_j}{2 - x^i_3} = \frac{y^i_j}{2 - y^i_3}$$

for $i = 1$ to n and $j = 1, 2$. Solving for x_j^i yields the following equivalent form

$$x_j^i = \left(\frac{2 - x_3^i}{2 - y_3^i}\right) y_j^i \tag{6}$$

Now note that rotations preserve the inner product between pairs of rotated points. Since y is simply the image of x under some rotation, we have for each k

$$\left\langle x_j^i, x_j^k \right\rangle = \left\langle y_j^i, y_j^k \right\rangle \tag{7}$$

Let $\lambda_i = \left(\frac{2-x_3^i}{2-y_3^i}\right)$. Now using Equation (6) to substitute for x_j^i yields

$$\begin{aligned} \left\langle y_j^i, y_j^k \right\rangle &= \left\langle x_j^i, x_j^k \right\rangle \\ &= \left\langle \lambda_i y_j^i, \lambda_k y_j^k \right\rangle \\ &= \lambda_i \lambda_k \left\langle y_j^i, y_j^k \right\rangle \end{aligned}$$

which implies that $\lambda_i \lambda_k = 1$. Now let us apply this fact to three indices in the polygons, say $1, 2$ and 3. We have

$$\lambda_1 \lambda_2 = \lambda_2 \lambda_3 = \lambda_1 \lambda_3 = 1$$

Dividing through by λ_1 gives $\lambda_2 = \lambda_3$ which implies that $\lambda_2 = \lambda_3 = 1$. Similar results are obtained by dividing by each of λ_2 and λ_3. In summary we have,

$$\lambda_1 = \lambda_2 = \lambda_3 = 1$$

This argument can be repeated for each triple λ_{i-1}, λ_i and λ_{i+1} giving that $\lambda_i = 1$ for all i from 1 to n. Now by definition

$$\left(\frac{2 - x_3^i}{2 - y_3^i}\right) = \lambda_i = 1$$

which implies that $x_3^i = y_3^i$ for all i. In other words, all vertices of both x and y must have the same z-coordinate. But two points with the same z-coordinate will have the same projection under P only if their first two coordinates are also the same. Thus, $x^i = y^i$ for all i which implies that $x = y$. □

The projection map P now forms a continuous bijection between X_T and $P(X_T)$. As usual, since both of these spaces are compact, we have that

P is homeomorphism. In particular this means that $P(X_T)$ has the same homology as X_T. This concludes the theoretical calculations of homology for projected SR-spaces. In the next section we will present the results of persistent homology computations performed on point-cloud data sets corresponding to samples from the spaces described above.

4 Persistent Homology Computations

In this section we perform persistent homology computations for point-clouds S_C and S_T sampled from the the spaces X_C and X_T respectively. We compute the persistent homology both for the sampled spaces before perspective projection, and for the projected SR-spaces $P(S_C)$ and $P(S_T)$. The point of these computations is to demonstrate that the persistent homology barcode of the SR-space of an object can be a good way of determining the object's approximate symmetries. In particular the results show that, at least for the spaces X_T and X_C, the theoretical calculation almost exactly matches the approximate computation. Additionally for the one case where there is a mismatch, we show that the approximate computation actually reveals something interesting (and non-obvious from the theory) about the effect of the projection map P on the toplogy of SR-spaces.

For these computations we chose T to be the regular polygon with three vertices i.e. an equilateral triangle. The program used for these computations is JPLEX, a persistent homology program developed by Stanford's CompTop project. The program is written in Java, and can be run through matlab's Java plugin. All of our programs for sampling or metric generation are written either as matlab scripts or in Java. We begin with a description of the method of sampling from the spaces.

Both S_C and S_T consist of $50,000$ points sampled uniformly at random from X_C and X_T. The sampling was performed by starting with the subset of $\mathbb{R}^3$ corresponding to each space, and applying $50,000$ rotations chosen uniformly at random from $SO(3)$. To choose uniform random rotations from $SO(3)$ we implemented, in matlab code, the algorithm described in [2]. To generate S_C we took the pair of points $r = (1, 0, 0)$ and $g = (-1, 0, 0)$ and stored pairs of images (Ar, Ag) for 50,000 matrices A chosen uniformly at random from $SO(3)$. Similarly for S_T we started with the equilateral triangle

given by the vertices

$$x_1 = (1, 0, 0)$$
$$x_2 = \left(\frac{-1}{2}, \frac{\sqrt{3}}{2}, 0\right)$$
$$x_3 = \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}, 0\right)$$

and stored 50,000 triples of points (Ax_1, Ax_2, Ax_3).

In both cases we used witness complexes with 60 landmark points to generate the filtered simplicial complexes in JPLEX. We chose our landmarks using the MaxMin greedy algorithm. This algorithm starts with a random point in the data set. For each subsequent step, MaxMin chooses as the next landmark the point x which maximizes the minimum distance between x and a point in the set of landmarks already chosen. This algorithm generally gives good coverage of the dataset, even for a relatively small number of landmarks.

Finally it is important to note that persistent homology computations only work for coefficients in a finite field $\mathbb{Z}/(p)$ for some prime p. Since all the theoretical calculations from Section 3 are done for homology with integer coefficients, we must use the universal coefficient theorem of homology to convert to coefficients in a finite field. Such a conversion can in some cases– for example when the integer homology has some torsion–lose information. In cases where such a loss may occur, we do the persistent homology computation multiple times with coefficients in multiple finite fields in order to give a more accurate picture of the integer homology. We now present the results of the computations for each of S_C and S_T.

4.1 Persistent Homology of S_C

Recall that for X_C thought of as pairs of points in $\mathbb{R}^3$, the regular Euclidean metric on $\mathbb{R}^6$ is continuous with respect to the action of $SO(3)$. So it is reasonable when building witness complexes for S_C to use the standard Euclidean distance as the metric. JPLEX has built in classes for handling point-cloud data with a Euclidean metric, so we used this built in functionality for all computations involving S_C. We break the results of the persistent homology

calculations for S_C into two parts. First, Figure 1 shows the persistent homology with coefficients in $\mathbb{Z}/(2)$ of S_C before perspective projection. Only

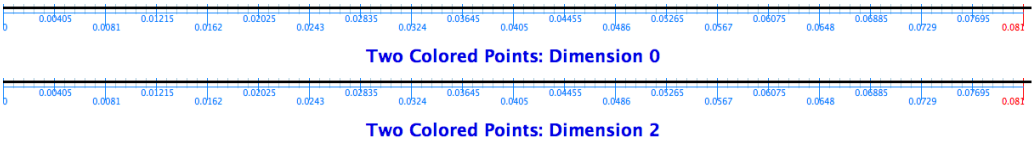

Figure 1: Barcode for Persistent Homology of S_C

dimensions zero and two are included because there were no persistence intervals in any other dimensions. In each of dimension zero and two, there is one interval that persists at all times. So, we conclude that the $\mathbb{Z}/(2)$ homology of S_C has one copy of $\mathbb{Z}/(2)$ in dimensions zero and two. Proposition 3.3 states that the $\mathbb{Z}$-homology of X_C is that of the sphere S^2, which has one copy of $\mathbb{Z}$ in dimensions zero and two. Since there is no torsion present in the $\mathbb{Z}$-homology, the universal coefficient theorem implies that the betti numbers for $\mathbb{Z}/(2)$-homology should be the same as for $\mathbb{Z}$-homology. This is exactly what we get from our persistence computation, $betti_0 = betti_2 = 1$ for $\mathbb{Z}/(2)$-homology. So the persistent homology of the sampled space S_C exactly matches the theoretical homology of the actual space X_C.

Now we move on to the projected sampled space $P(S_C)$. Recall from Section 3.1 that the integer homology of $P(X_C)$ had a $\mathbb{Z}$ in dimensions $0, 1, 2$ and 0 in all other dimensions. So the betti numbers are $betti_i = 1$ for $i = 0, 1, 2$ and $betti_i = 0$ otherwise. Again since there is no torsion, the betti numbers for the $\mathbb{Z}/(2)$-homology of $P(X_C)$ should be the same. Figure 2 shows that this holds for the projected sampled space $P(S_C)$. So the approximate computation again matches the theoretical result.

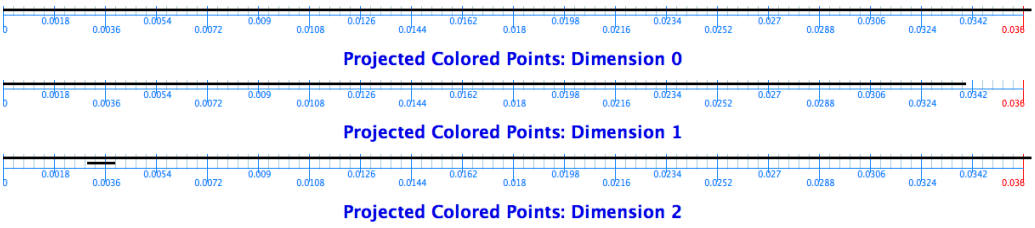

Figure 2: Barcode for Persistent Homology of $P(S_C)$

It is useful to note that there is a little more noise in the results of this

computation, such as the very short lived cycle represented by the short interval in dimension two. This noise probably resulted from the distortion of the metric caused by the projection map P. In particular, though the exact identifications resulting from the projection map are correctly given by the theoretical calculations in Section 3, there may be approximate identifications resulting from the flattening of the SR-space under P.

In particular, points that were previously far apart might be projected relatively near to each other under P. This could cause chains in the simplicial complex that were once not cycles to appear as approximate cycles for some times in the filtration, resulting in noise in the form of short persistence intervals. Also interesting is the case where P projects all the points on some cycle to a small enough area that the cycle disappears in the approximate computation. As we will see in the next section, this second case plays an important role in the persistent homology of S_T.

4.2 Persistent Homology of S_T

Recall from Section 3.2 that we defined a metric d_T on X_T that was required to correctly capture the symmetries of the polygon. The sampled data for S_T is represented as a length 50,000 array of points in $\mathbb{R}^9$, so in order to correctly construct witness complexes for S_T we had to add capability to measure distance using the metric d_T to JPLEX. We extended JPLEX with a class for triangle data–points in $\mathbb{R}^9$ thought of as triples of vertices of triangles–which implements the metric d_T for use in construction of simplicial complexes.

Proposition 3.9 shows that the homology of X_T contains torsion. In particular, $H_1(X_T) \cong \mathbb{Z}/(4)$ when T is a triangle (and so n is odd). The universal coefficient theorem for homology then implies that for $\mathbb{Z}/(p)$-homology with p dividing 4, we will have $\mathbb{Z}/(p)$ in dimensions 1 and 2. For p not dividing 4, we will have 0 in dimensions 1 and 2. In summary, the $\mathbb{Z}/(p)$-homology of X_T is given by

$$H_k(X_T; \mathbb{Z}/(p)) \cong \begin{cases} \mathbb{Z}/(p) & \text{for } k = 0, 3 \\ \mathbb{Z}/(p) & \text{for } k = 1, 2 \text{ and } p \mid 4 \\ 0 & \text{for } k = 1, 2 \text{ and } p \nmid 4 \end{cases}$$

Thus, to attempt to retain as much information as possible from the integer homology, we compute both $\mathbb{Z}/(2)$ and $\mathbb{Z}/(3)$ persistent homology of S_T.

Figures 3 and 4 give the results of these computations.

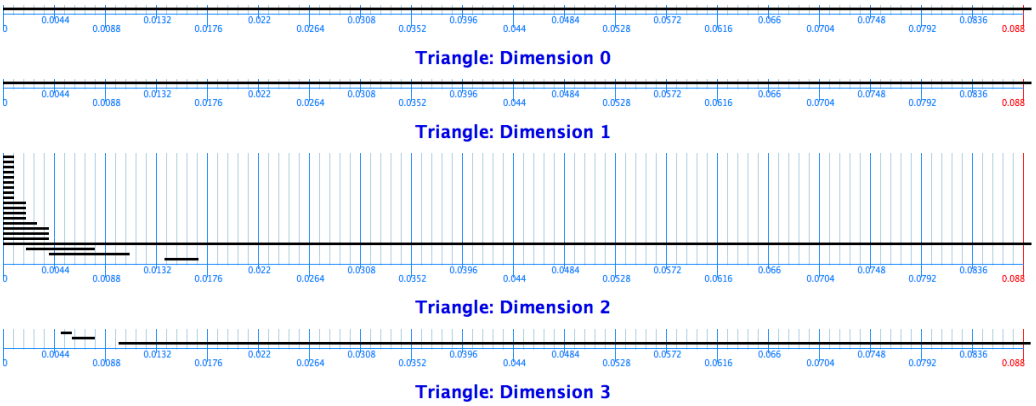

Figure 3: Barcode for Persistent Homology of S_T with $\mathbb{Z}/(2)$ Coefficients

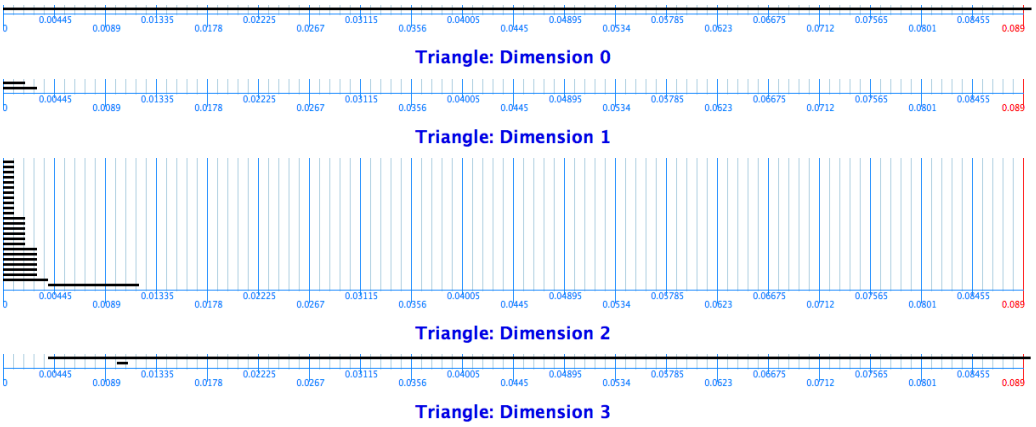

Figure 4: Barcode for Persistent Homology of S_T with $\mathbb{Z}/(3)$ Coefficients

The intervals that persist across the whole barcodes correspond precisely, in both the $\mathbb{Z}/(2)$ and $\mathbb{Z}/(3)$ case, to the theoretical predictions for the homology of X_T. The next step is to compute the persistent homology of $P(S_T)$. The result of the computation for $\mathbb{Z}/(2)$ coefficients is given in Figure 5.

For this barcode there are no intervals in dimension one and all the intervals in dimension two are too short to indicate any sort of cycle that actually persists through a large number of filtration times. The homology groups

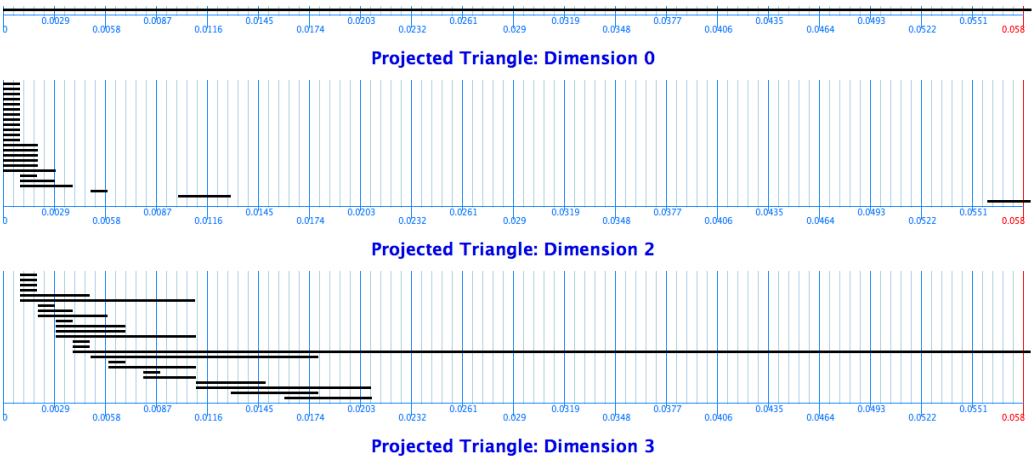

Figure 5: Barcode for Persistent Homology of $P(S_T)$ with $\mathbb{Z}/(2)$ Coefficients

implied by the barcode are then simply

$$H_k(P(S_T); \mathbb{Z}/(p)) \cong \begin{cases} \mathbb{Z}/(p) & \text{for } k = 0, 3 \\ 0 & \text{otherwise} \end{cases}$$

Since Proposition 3.10 shows that $P(X_T)$ is homeomorphic to X_T, the persistent homology computed for $P(S_T)$ does not match the theoretical prediction. However, a more careful examination of why this inconsistency occurs reveals an interesting fact about the effect of the perspective projection on the *approximate* symmetries of an object.

Recall that the two $\mathbb{Z}/(2)$s in dimensions one and two of $H_*(X_T; \mathbb{Z}/(2))$ both arise from the torsion in $H_1(X_T; \mathbb{Z})$ via the universal coefficient theorem. Thus, the trivial homology in dimensions one and two of the barcode shown in Figure 5 indicates that the projection map P is killing the first dimensional $\mathbb{Z}$-homology of X_T. Since $H_1(X_T)$ is simply the abelianization of the fundamental group $\pi_1(X_T)$, determining the effect of P on π_1 will explain its effect on H_1.

To this end, consider a loop γ_β representing the homotopy class of the generator $\beta \in \pi_1(X_T) \cong Q_{4n}$. Recall from Lemma 3.8 that the generator $\beta \in Q_{4n}$ maps to the generator $b \in D_{2n}$ under the quotient map $S^3 \to SO(3)$. We also saw in Proposition 3.7 that b is the rotation by π about an axis from the origin through one vertex of the polygon. Thus a loop in X_T representing

β is given by

$$\gamma_\beta(t) = B_t(x_0)$$

where x_0 is the basepoint in X_T and B_t is a rotation by πt about an axis going through one vertex of the polygon x_0 in X_T.

Now let us return to the specific case of S_T where the polygon T is a triangle. Let the basepoint x_0 be the equilateral triangle given in coordinates that we used to generate our $50,000$ random rotation samples. In particular x_0 lies in the xy-plane, centered at the origin, with one vertex at the point $(1, 0, 0)$. Let B_t be the rotation by πt about the x-axis. Then the loop γ_β is simply a continuous sequence of triangles rotating over a total angle of π through the bisector of one vertex. This loop represents the non-trivial homotopy class $\beta \in \pi_1(S_T)$ and so is not homotopic to a point.

In coordinates, the triangles that comprise this loop are given by

$$\gamma_\beta(t) = \begin{pmatrix} 1 & 0 & 0 \\ -\frac{1}{2} & \frac{\sqrt{3}\cos(t)}{2} & \sin(t) \\ -\frac{1}{2} & -\frac{\sqrt{3}\cos(t)}{2} & -\sin(t) \end{pmatrix}$$

where the ith row of the above matrix gives the coordinates of the ith vertex of the triangle. For $t \in [0, \frac{\pi}{2}]$ the second vertex rotates toward the viewer, and the third rotates away from the viewer until they both lie on the yz-plane at $t = \frac{\pi}{2}$. Then for $t \in [\frac{\pi}{2}, \pi]$ the second vertex begins to rotate downward away from the viewer, and the third vertex begins rotating upward until at $t = \pi$ they both lie flat again on the xy-plane. Now we apply the projection map to the loop γ_β.

$$P(\gamma_\beta(t)) = \begin{pmatrix} \frac{1}{2} & 0 \\ -\frac{1}{4-2\sin(t)} & \frac{\sqrt{3}\cos(t)}{4-2\sin(t)} \\ -\frac{1}{4+2\sin(t)} & -\frac{\sqrt{3}\cos(t)}{4+2\sin(t)} \end{pmatrix}$$

Under projection, the interval $t \in [0, \frac{\pi}{2}]$ corresponds to a triangle where both vertices move inward unevenly until at $t = \frac{\pi}{2}$ the triangle becomes a straight line. For the second half of the loop when $t \in [\frac{\pi}{2}, \pi]$, the two vertices separate and move back outward until they return to their original places. Though the paths taken in these two halves of the loop γ_β do not match exactly, they are very close. Intuitively this means that γ_β is a long and skinny loop that is unlikely to be detected by sampling a few points from it.

That is, a pair of points across the loop from each other may be closer together than they are to the next point sampled along the actual path of the loop. Such a loop would have its interior triangulated in a filtered persistence complex before all the edges that comprise the loop are added. Thus, it would appear trivial in the persistent homology computation. The following lemma shows that if $\beta \in \pi_1(S_T)$ is trivial, then the first integer homology of S_T is trivial.

Lemma 4.1. *Let n be odd and $\beta \in \pi_1(S_T) \cong Q_{4n}$ be trivial. Then the first integer homology group $H_1(S_T; \mathbb{Z})$ is trivial.*

Proof. Recall the relation in Q_{4n} given by $\beta^{-1}\alpha\beta = \alpha^{-1}$. If β is trivial this implies that $\alpha^2 = 1$. Further, we have $\alpha^n = \beta^2 = 1$ for trivial β. For odd n having both $\alpha^2 = 1$ and $\alpha^n = 1$ implies that $\alpha = 1$. So, $\pi_1(S_T)$ is trivial which implies that $H_1(S_T; \mathbb{Z})$ is also trivial as it is the abelianization of a trivial group. □

As was mentioned before, if $H_1(S_T, \mathbb{Z})$ is killed by the projection map, then the homology from the barcode in Figure 5 is correct for $P(S_T)$. In summary, what we have shown is that the map P distorts the SR-space S_T in such a way as to eliminate the first dimensional homology. Viewed from another perspective, this means that the symmetry of the triangle rotating by π about an axis through one vertex is somehow lost in the perspective projection.

It is important to observe that there is really nothing special about the triangle in the arguments above. The loss of the symmetry corresponding to rotation by π is actually more of a property of the perspective projection map P. This observation can therefore be extended to a large class of objects that have a symmetry of a rotation by π. For example, any of the other polygons with n sides will behave just as the triangle did and lose the generator β of the fundamental group. For other objects, the effect on homology may be different, but the loss of some generator for the fundamental group will remain, and must therefore be taken into account when computing persistent homology of SR-spaces.

5 Conclusion and Future Work

In summary, we first introduced the concepts of symmetrized rotation spaces and perspective projection for understanding the symmetries of an object

from two-dimensional images of the object. We performed theoretical calculations of the homology of the SR-spaces of two model objects: a pair of colored points, and a regular polygon with n sides. We then generated 50,000 images of each model object and used JPLEX to compute the persistent homology of these sampled SR-spaces.

In all cases but one, the persistent homology agreed with the theoretical calculation. The one exception was for the perspective projection of the SR-space of the regular polygon on n sides. In this case we found that the perspective projection distorted the metric in such a way as to approximately destroy the homological information about one of the polygon's symmetries. We observed that this is actually a quite general feature of the perspective projection map–a feature that would have been difficult to discover without the computational tool of persistent homology.

Future work on topological detection of symmetries in images could be in two possible directions. First, it would be useful to compute the homology by hand of a larger collection of model symmetric objects. Second, the persistent homology computations performed for this paper were done with a very idealized method for sampling two-dimensional images of an object. A natural next step would be to use computer graphics renderings of an object under various rotations to see if it was possible to achieve the same results. In either case, our work demonstrates the power of computational topology in detecting symmetries in images, and should allow interesting applications of the techniques we have described to problems in image processing.

6 Acknowledgements

I would like to thank my advisor Professor Gunnar Carlsson for all of his help with this honors thesis over the past year. I would also like to thank my fellow student Peter Smillie for helping with the homology calculations, and in general for doing math with me for as long as I can remember.

References

[1] Adem, A., Milgram, R. J. *Cohomology of Finite Groups.* Springer-Verlag: Berlin. 1994.

[2] Arvo, J. (1992), "Fast random rotation matrices", in David Kirk, Graphics Gems III, San Diego: Academic Press Professional, pp. 117–120, ISBN 978-0-12-409671-4

[3] Carlsson, G., "Topology and Data" , Bulletin of the American Mathematical Society , vol 46, Number 2, pp.255-308 . Jan 2009.

[4] Carlsson, G., Ishkhanov, T., de Silva, V., and Zomorodian, A. "On the local behavior of spaces of natural images." International Journal of Computer Vision 76, 1 (2008), 1–12.

[5] Carlsson, G., Zomorodian, A., Collins, A. and Guibas, L., "Persistence barcodes for shapes" , International Journal of Shape Modeling , 11 (2005), pp. 149–187. Jan 2005

[6] Hatcher, A. *Algebraic Topology.* Cambridge University Press. 2001.

[7] Hayami, Takao and Sanada, Katsunori(2002) "Cohomology ring of the generalized quaternion group with coefficients in an order", Communications in Algebra, 30: 8, 3611 – 3628

[8] Zomorodian, A. and Carlsson, G. "Computing persistent homology," Discrete and Computational Geometry, 33 (2), pp. 247–274.