

On the Local Behavior of Spaces of Natural Images

Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, Afra Zomorodian*

Abstract

In this study we concentrate on qualitative topological analysis of the local behavior of the space of natural images. To this end, we use a space of 3 by 3 high-contrast patches \mathcal{M} studied by Mumford et al. We develop a theoretical model for the high-density 2-dimensional submanifold of \mathcal{M} showing that it has the topology of the Klein bottle. Using our topological software package *PLEX* we experimentally verify our theoretical conclusions. We use polynomial representation to give coordinatization to various subspaces of \mathcal{M} . We find the best-fitting embedding of the Klein bottle into the ambient space of \mathcal{M} . Our results are currently being used in developing a compression algorithm based on a Klein bottle dictionary.

Keywords : topology; natural images; manifold; filtration; Klein bottle; persistent homology

Introduction

A natural image, such as a digital photograph, may be viewed as a vector in a very high-dimensional space \mathcal{P} . The dimension of \mathcal{P} is the number of pixels in the format used by the camera and the image is associated to the vector whose coordinates are grey scale values of the pixels. There has been a considerable amount of interest in studying families of images viewed as an approximation to a ‘submanifold’ within \mathcal{P} and this point of view has provided some interesting parametrizations of ‘phase spaces’ such as the space of all positions of a certain object in space or the hand written copies of a particular digit. Alternatively, David Mumford has asked what can be said about the set $\mathcal{I} \subseteq \mathcal{P}$ of all images taken outdoors. Intuitively, one feels that although \mathcal{I} is very high-dimensional, it should also in some sense be of very high codimension within the ambient space \mathcal{P} . The intuition behind this is that one believes that a vector chosen at random from \mathcal{P} will almost never be close to an actual image vector. However, this kind of global study is difficult since even if we believe that \mathcal{I} in some sense approximates a submanifold, we are confronted with the fact that both its dimension and codimension are very high, so we are unlikely to obtain much information directly. An approach suggested by Mumford and carried out in [3] is instead to analyze the local behavior, modeled by spaces of small ‘patches’ of pixels within images. Specifically, Mumford et al. study a large space of 3 by 3 patches of pixels obtained from a database of images constructed by H. van Hateren ([2]), and make a number of extremely interesting observations about the resulting space of patches.

The goal of this paper is to analyze a similar space of patches, which will be denoted \mathcal{M} throughout the paper, in a systematic way. Our results are qualitative in nature in the sense that we describe the global topological structure of various spaces within \mathcal{M} rather than their local geometric behavior. In particular, we show that there is a large 2-dimensional subset \mathcal{K} of \mathcal{M} , whose homology is that of a Klein bottle. We also provide theoretical explanations for why in fact it is the case that \mathcal{K} is topologically equivalent, i.e.

*e-mail addresses: gunnar@math.stanford.edu, tigran@math.stanford.edu, Vin.deSilva@pomona.edu, afra@cs.dartmouth.edu

homeomorphic, to the Klein bottle. To this end, we use a parametrization of \mathcal{M} by a space of polynomials in two variables. The Klein bottle model can be regarded as a 2-dimensional skeleton of the space of 3 by 3 patches thought of as a topological CW -complex ([4]). Finally, we show how to flow the initial (theoretical) embedding of the Klein bottle in the ambient space of \mathcal{M} into the one which best approximates the high-density 2-dimensional subspace of \mathcal{M} .

In order to obtain qualitative information about various subspaces of \mathcal{M} we use a mathematical formalism called algebraic topology, and specifically the newly developed portion of that formalism called persistent homology ([1], [4]).

The work described in this paper has interesting connections with research in neuroscience. It is possible to analyze the topology of the space of images presented to the visual cortex area V1 by measuring the neuron responses recorded from some subset of neurons of V1 ([10]). We also believe the research initiated in this paper will prove useful in the area of image compression. One approach which is currently being explored is to create a Huffman type code for encoding patches from a high-density manifold sitting inside of the space of high-contrast patches. This code can be used to map high-contrast patches extracted from a particular image to the points in the codebook which are the closest to them.

The main advantage of the topological approach taken in this paper is that it allows one to find nonlinear structures within the data which would be very hard to find using statistical techniques. More concretely, in addition to allowing us to determine whether the data splits into clusters, our methods allow us to discover the topological structure of each of the clusters.

1 Topological approach

1.1 Homology

Bearing in mind our goal of analyzing the space of images qualitatively we turn to algebraic topology, which provides a rich set of tools aimed at understanding the global properties of the space. Ideally, we would like to get the homeomorphism type of the space at hand, i.e. understand its geometry up to stretching and bending but not tearing or gluing. For example, despite having different geometries, the circle and the ellipse have the same homeomorphism type since one can be deformed into the other without tearing or gluing. However, even in the classical setting when the space is given in closed form, for instance cut out by a system of equations, it is usually hard to obtain its homeomorphism type and so one has to resort to finding some partial invariants, e.g. homotopy, homology etc. In our case there are also serious computational limitations since the space is given as point cloud data. To deal with this situation and calculate space invariants numerically we would like to represent our data combinatorially. One way to do this is to approximate our space by a simplicial complex and compute its simplicial homology groups. Instead of developing the homology formalism from scratch, here we state the main properties that we are interested in. For a more in-depth treatment of homology please refer to [4].

- To every space and non-negative integer k one associates a vector space $H_k(X)$ over a field F , which is specified once and for all. We will deal with the field with two elements 0 and 1, but the real or rational numbers could also be chosen. The dimension of $H_k(X)$ is called k -th Betti number of X .
- The 0-th Betti number of X is always the number of connected components of X . In other words, the number of blocks in a decomposition of X into pieces which can be separated from each other.
- The construction of $H_k(X)$ is functorial in X , meaning that any continuous map of spaces $f : X \rightarrow Y$ induces a linear transformation of vector spaces $H_k(f) : H_k(X) \rightarrow H_k(Y)$.

Roughly speaking the k -th homology group measures the number of k -dimensional subspaces of X that

have no boundary in X and themselves are not a boundary of any $k + 1$ -dimensional subspace.

Given this, the next question is, “how do we compute the Betti numbers?” A simple answer is that the Betti numbers can be computed by hand using Gaussian elimination of certain matrices over the field F , provided one has complete information about one’s space. However, when complete information is unavailable (as in our case), one has to develop methods which make up for this lack of information. This leads to the idea of persistent homology discussed in the next section.

1.2 Persistent homology

In most applications one deals with the data in the form of a point cloud, i.e. a large but finite set of points in Euclidean space sampled from some underlying geometric object. Persistent homology is in a sense a multiscale approach. Its utility is best illustrated by considering the following simple geometric situation. Suppose we are given a set consisting of two blobs of points separated from each other by some “small” distance ϵ . Because our invariants highly depend on the connectivity of the spaces that we are trying to describe, we need to specify how small ϵ should be for the two blobs to qualify as one connected component. To avoid making such choices one utilizes the idea of persistence, which we sketch below. The first step is to construct a simplicial complex which approximates the homeomorphism type of the underlying space. One way to do this is via a “Rips complex”. Consider a point cloud data set X (which we will refer to as $PCDX$) in Euclidean space and a number $\epsilon > 0$. We let $R(X, \epsilon)$ denote the simplicial complex whose vertex set is X and where we declare that the set $\{x_0, x_1, \dots, x_k\}$ spans a k -simplex if and only if $d(x_i, x_j) \leq \epsilon$.

Figure 1 shows two contrasting situations. The Rips complex on the left is built with ϵ equal to the longest side of the polygon (vertices of the complex are vertices of the polygon); in the Rips complex on the right ϵ is equal to the longest diagonal ab which results in all possible simplices being present. The homology groups are different in these two cases.

If we have no information beyond the raw point cloud data, such as curvature or some statistical data about the space, it is hard to choose a value for ϵ that produces a simplicial complex whose homology agrees with our intuitive understanding of the space. Here the idea of persistence comes into play. We also use the functoriality of homology described above. Whenever $\epsilon \leq \epsilon'$ there is an inclusion $R(X, \epsilon) \hookrightarrow R(X, \epsilon')$ and therefore by functoriality, there is a linear transformation $H_k(R(X, \epsilon)) \rightarrow H_k(R(X, \epsilon'))$ for each $k \geq 0$. This means that instead of obtaining just one vector space, we obtain a directed system of vector spaces $\{V_\epsilon\}_{\epsilon \geq 0}$, i.e. a family of vector spaces V_ϵ for each $\epsilon \geq 0$ together with linear transformations $V_\epsilon \rightarrow V_{\epsilon'}$ whenever $\epsilon \leq \epsilon'$. This observation is extremely useful since there is a simple classification of directed systems of vector spaces up to isomorphism.

By definition, a **barcode** is a finite set of intervals on the real line, where each interval is required to be bounded below but not necessarily above. The intervals are drawn one above the other like in figure 2, which is a barcode for the Betti numbers of a space comprised of two circles meeting at a single point (this




Figure 1: left: Rips complex with a smaller ϵ , right: larger ϵ , no 1-dimensional homology




Figure 2: The barcode for the wedge of two circles

space is known as the wedge of two circles). It is a theorem that barcodes are in bijective correspondence with isomorphism classes of directed systems of vector spaces ([1]).

We can think of homology as a measure of the number of holes of various dimensions in our space. Thus an interval starting at time $t = t_0$ and ending at $t = t_1$ can be interpreted as a hole that appears in the Rips complex at $\epsilon = t_0$ and gets filled in at $\epsilon = t_1$. Long intervals correspond to holes persisting through a large range of values of the persistence parameter ϵ while short intervals correspond to holes that get filled in quickly. The natural interpretation is that long intervals give true homological information about the underlying object while short ones are “topological noise” coming from irregularities in the point sampling rather than reflecting the topology of the space itself. In figure 2 we see one long line in $H_0(X)$ and two in $H_1(X)$ which reflects the fact that a wedge of two circles has one connected component and two one-dimensional holes.

For the computations done in this paper we use a variation of the Rips complex called the witness complex ([5]), which considerably speeds up the construction of a simplicial complex approximating the space X . The main reason for this is that the witness complex uses only a small subset of points, so-called landmarks, of the original $PCD X$ as vertices of the simplicial complex. The rest of the points are used as witnesses to inclusion of simplices into the simplicial complex. The conditions determining when the simplex enters the simplicial complex are closely related to the ones used to construct the Delaunay complex for a collection of points in Euclidean space ([7]).

The directed system of complexes obtained in this way turns out to be an excellent approximation to the underlying spaces in the sense that its persistent homology often effectively computes the actual homology, and the amount of noise is much smaller than in the case of Rips complexes. The barcodes are computable from point cloud data using the *PLEX* software package developed by G. Carlsson, V. de Silva, A. Zomorodian and P. Perry. The *PLEX* software is our main tool in getting the qualitative (homological) information about various spaces in this paper.

1.3 The density filtration

The subject of density estimation is a highly developed one in statistics ([9]). In this paper we use the ‘nearest neighbor’ estimation of the local density of the space at a point x . In other words, fix $x \in X$ and $k > 0$, let $\rho_k(x)$ denote the distance from x to the k -th nearest neighbor of x . The function ρ_k is inversely proportional to the density at x since small values of ρ_k indicate that there are k points of X already at this distance from x . Different values of k give rise to different density estimations at x , with larger values giving more global estimations. Once we fix a value of k and order the points of X by descending density we need to choose a ‘cut’ parameter p , i.e. a percentage of densest points we extract for our topological computations. In most data sets the points are concentrated around some ‘core’ subset. Finding this core set may give a clue about some phenomenon occurring within the data which otherwise would be very hard to predict; this is exactly what happened in our case. To this end, the study of the set of points falling into the top p percent (denoted $X(k, p)$ due to its dependence on both k and p parameters) may provide important topological information about the ‘core’ set, which could be lost when considering all the points of X .

1.4 Denoising

Another technique we employed in the study of the image space is denoising ([9]). The idea is best illustrated by a simple example. Suppose we have a space consisting of a unit circle in \mathbb{R}^2 together with some extra points very near the circle. The homology of this space is the same as that of a circle but in the construction of an approximating simplicial complex there may appear some extra simplices which produce noise and may considerably reduce the computational efficacy. To avoid this problem we employ the denoising algorithm, which is also based on the ‘nearest neighbor’ method. We start by replacing each point in the original space by an average of points in its neighborhood. This produces a new space, and we iterate this procedure. Two iterations is usually enough to give much cleaner and faster results for homology calculations. The method requires one parameter - the neighborhood size in terms of number of points. We set it experimentally depending on the density variation within the space.

2 The space of patches

Our main space \mathcal{M} is a collection of $4 \cdot 10^6$ ‘3 by 3’ patches of high contrast obtained from the collection of still images gathered by H. van Hateren and A. van der Schaaf. \mathcal{M} is a subset of a larger set $\tilde{\mathcal{M}}$, provided to us by K. Pedersen. The size of $\tilde{\mathcal{M}}$ is about $8 \cdot 10^6$. Below is a series of steps performed to obtain a set of high-contrast patches from a particular image. (See [3] for more details).

1. Select an image from the still image collection.
2. Extract at random 5000 3 by 3 patches from the image. Regard each patch as a vector in 9-dimensional space.
3. Next, for each patch do the following.
 - (a) Compute the logarithm of intensity at each pixel.
 - (b) Subtract an average of all coordinates from each coordinate. This produces a new 9-vector.
 - (c) For this vector of logarithms compute the contrast or “ D -norm” of the vector. The D -norm of a vector x is defined as $\sqrt{x^T D x}$, where D is a certain positive definite symmetric 9×9 matrix.

- (d) Keep this patch if its D -norm is among the top 20 percent of all patches taken from the image.
4. Normalize each of the selected vectors by dividing by their respective D -norms. This places them on the surface of a 7-dimensional ellipsoid.
 5. Perform a change of coordinates so that the resulting set lies on the actual 7-dimensional sphere in \mathbb{R}^8 .

The space $\tilde{\mathcal{M}}$ is obtained by applying the above procedure to a subset of images from the still image collection.

In [3] an extremely convenient basis for the space of patches was also constructed. We give it here for reader's convenience and refer to it later as the DCT basis.

$$\begin{aligned}
\mathbf{e}_1 &= (1/\sqrt{6}) \begin{pmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{pmatrix}, & \mathbf{e}_2 &= (1/\sqrt{6}) \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{pmatrix}, & \mathbf{e}_3 &= (1/\sqrt{54}) \begin{pmatrix} 1 & -2 & 1 \\ 1 & -2 & 1 \\ 1 & -2 & 1 \end{pmatrix}, \\
\mathbf{e}_4 &= (1/\sqrt{54}) \begin{pmatrix} 1 & 1 & 1 \\ -2 & -2 & -2 \\ 1 & 1 & 1 \end{pmatrix}, & \mathbf{e}_5 &= (1/\sqrt{8}) \begin{pmatrix} 1 & 0 & -1 \\ -2 & 0 & 2 \\ 1 & 0 & -1 \end{pmatrix}, & \mathbf{e}_6 &= (1/\sqrt{48}) \begin{pmatrix} 1 & 0 & -1 \\ -2 & 0 & 2 \\ 1 & 0 & -1 \end{pmatrix}, \\
\mathbf{e}_7 &= (1/\sqrt{48}) \begin{pmatrix} 1 & -2 & 1 \\ 0 & 0 & 0 \\ -1 & 2 & 1 \end{pmatrix}, & \mathbf{e}_8 &= (1/\sqrt{216}) \begin{pmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{pmatrix}
\end{aligned}$$

The basis vectors are normalized so as to have D -norm equal to one.

3 The space of polynomials in two variables

In what follows it will be important to identify various subspaces of \mathcal{M} in a more common closed form, i.e. as a family of polynomials depending on one or more parameters. In order to obtain such a representation we note that each individual patch can be thought of as belonging to an xy -plane with each of the nine pixels having coordinates (x_0, y_0) , $x_0 \in \{-1, 0, 1\}$, $y_0 \in \{-1, 0, 1\}$. We denote the grid made by these nine points in a plane by H . Let us denote the space of all polynomials in two variables by P . Evaluating a polynomial $p(x, y)$ at each of the points of H produces a 9-vector. Subtracting the mean from each coordinate and dividing the result by the Euclidean norm produces a point on the surface of the 7-sphere S^7 in \mathbb{R}^8 . By making $p(x, y)$ depend on parameters we obtain a certain topological space $T \subseteq P$, where each point is a polynomial corresponding to a particular choice of these parameters. Let us define a map $q : P \rightarrow S^7$ as a composite of evaluating the polynomial on H , subtracting the mean and normalizing. In the instance where we know the topology of T and the behavior of the map $q|_T$ we can deduce the topology of the space $im(q|_T)$ in S^7 . We will make use of this idea later.

4 Klein bottle

In this section we describe the space which plays a key role in our study of \mathcal{M} , the Klein bottle. To visualize this surface within \mathbb{R}^3 one can picture it as a tube which is allowed to move through itself to connect up with the opposite end. Figure 3 shows an immersion of the Klein bottle in \mathbb{R}^3 .



Figure 3: Klein bottle immersion in \mathbb{R}^3

In order to make a precise definition of the Klein bottle, we will need to establish the concept of identifying points in a topological space. Given a topological space, we can always construct a new topological space by specifying sets of points to be identified. For example, starting with the closed interval $[0, 1]$ and identifying its endpoints produces a space homeomorphic to a circle. In the case of Klein bottle one starts with a unit square $[-1, 1] \times [-1, 1]$ in the plane and makes the following identifications on the boundary: $(-1, t)$ with $(1, -t)$ and $(s, -1)$ with $(s, 1)$ as in figure 4.




Figure 4: Klein bottle representation as a rectangle with opposite edges identified

Note that even though this space is two-dimensional it cannot be embedded in \mathbb{R}^3 without self-intersections. It can, however, be embedded as a subspace in \mathbb{R}^4 . Another interesting feature of the Klein bottle is that it is non-orientable, i.e. it essentially has only one side unlike surfaces such as the sphere or the torus.

5 Results for $X(k, p)$ spaces

We would like to analyze the structure of the dense subspaces within the spaces of patches.

In [8] it was observed that if we take a space of patches of size $5 \cdot 10^4$ (note that this is much smaller than \mathcal{M}) and extract 30 percent of the densest points using the density estimator $k = 15$, we obtain a space whose first Betti number is equal to 5. One can construct many different spaces that satisfy this condition on the Betti numbers. However, if we take into account the nature of the space, one reasonable candidate is

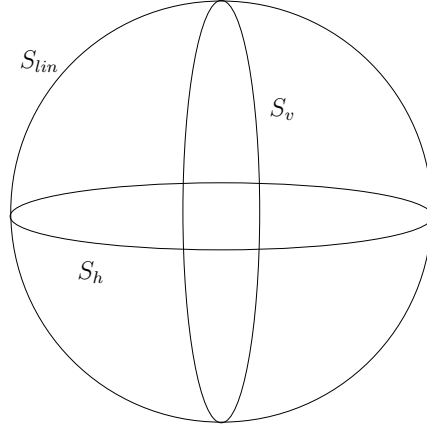


Figure 5: The ‘three circle’ space




Figure 6: 3 by 3 patches parametrized by the Klein bottle

the three circle model (figure 5), denoted C_3 . In the three circle space, two of the circles (labeled S_v and S_h in the picture) intersect the third one (S_{lin}) in exactly two points, while they themselves do not intersect (the two ellipsoidal circles in the figure do not intersect inside S^7). Note that the space C_3 lives in \mathbb{R}^8 where this geometric situation is possible, but when drawn in the plane the extra intersections cannot be avoided without distorting the actual pattern. Figure 6 shows the same three circles sitting in the Klein bottle parametrization space (cf. figure 4). The red line in the figure which appears to be two connected components is actually a single circle (corresponding to S_{lin}) on the surface of the Klein bottle due to identifications made on the vertical edges of the square. The green and black lines correspond to S_v and S_h respectively.

In section 6 below we provide justification to why each of the three components of C_3 is topologically a circle. This model was verified by inspection and suggested in [8]. In terms of polynomial representation of patches (sections 3 and 6) the circle S_{lin} (primary circle) corresponds to linear gradients while the other two (secondary circles) to quadratic polynomials in a single variable x (or y). To verify that the first Betti number of this space is equal to 5 we note that the space C_3 can be thought of as a graph with 8 arcs, 4 vertices and one connected component. Therefore, using the standard formula we arrive at $b_1 = \#(arcs) -$




Figure 7: *PLEX* results for $X(15, 30)$

$\#(\text{vertices}) + \#(\text{connected components}) = 5$. The *PLEX* result from [8] for $X(15, 30)$ is shown in figure 7.

The above result for $X(15, 30)$ suggests that there are two competing preferences in the natural images. One is a preference for linear intensity functions on small patches and the other is a preference for vertical and horizontal directions as opposed to intermediate directions. Another explanation for the presence of secondary circles lies in our choice of patches as squares with edges in vertical and horizontal directions. When we analyzed the set of patches coming from the images taken while holding the camera at a 45° angle we found that the set of densest points projected onto the first two coordinates (corresponding to linear gradients) of the DCT basis is also rotated by the same angle as compared with a similar projection of the set of patches coming from \mathcal{M} . However, a strong bias in favor of vertical and horizontal directions (for purely quadratic gradients) remained, thereby indicating that the choice of patch shape also affects the density distribution.

In [8] it was also shown that for higher values of the parameter k (i.e. less localized density estimation) with the fixed cut parameter p the space $X(k, p)$ loses both secondary circles. Figure 8 shows the results for $X(300, 30)$. The only interval in the Betti one barcode corresponds to a primary circle of linear intensity gradients.

These results raise the question of what happens when we decrease k below 15. Can we detect the intermediate directions as well? More precisely, can we detect whether the space $X(k, p)$ spans a two-manifold for some cut parameter p with $k < 15$?

Since reducing the value of k below 15 seems unreasonable due to noise issues we decided to examine the situation for a much bigger space of patches (recall that the size of \mathcal{M} is $4 \cdot 10^6$). In this new situation the value of $k = 15$ for a space of size $5 \cdot 10^4$ would correspond to a value $k = 15 \times (4 \cdot 10^6 / 5 \cdot 10^4) = 1200$ for our space \mathcal{M} . Figure 9 shows the result for the space $X(100, 10) \cup Q \subseteq \mathcal{M}$, here Q is the set of size $|Q| = 30$, which is less than 0.01% of the size of $X(100, 10)$. Note that $k = 100$ corresponds to $k = 1.25$ for the space of size $5 \cdot 10^4$.

The size of Q can be thought of as a measure of the failure of the density function to cut out the 2-dimensional manifold. To obtain the set Q , we first compute distances from points of \mathcal{M} to 30 points sampled from \overline{Q} , where \overline{Q} is the set of patches with purely quadratic gradients, which in polynomial rep-




Figure 8: *PLEX* results for $X(300, 30)$

resentation (cf. section 6) correspond to either $(ax + by)^2$ or $-(ax + by)^2$ with neither a nor b equal to 0. Geometrically \overline{Q} is a pair of circles with neighborhoods of two points removed from each. These two points correspond to vertical and horizontal gradient directions. Second, for each of the sampled points, we select the closest point of \mathcal{M} .

The reason that the points of Q have a much lower density than the rest of the 2-manifold is twofold. On the one hand, vertical and horizontal directions of intensity gradient occur more frequently than other directions in natural images; on the other hand the patch shape was chosen to be a square with vertical and horizontal sides.

For homological computations, we randomly select a subset S of size $|S| = 10,000$ from $X(100, 10)$ and add points of Q to it. We repeat this many times (each time selecting different S) to make sure the results are stable and capture the actual topology of the space $X(100, 10) \cup Q$.

The barcode (figure 9) reflects the situation where there are two essential one-dimensional cycles and one two-dimensional cycle. The presence of an interval in Betti two implies that the underlying space is no longer one-dimensional. Moreover, there is a general result from algebraic topology that any two-manifold has $b_2 = 1$ when homology is taken with coefficients in \mathbb{Z}_2 . The next section provides explanations for the result shown in figure 9.

Combined together, the results of this section suggest strong evidence that as the density estimation parameter decreases, the space of densest points with an appropriate value of the cut parameter fills out a two-manifold: initially, for large k it consists only of a primary circle of linear gradients; then it acquires two additional circles corresponding to quadratic gradients in the ‘preferred’ vertical and horizontal directions; finally it admits all the intermediate directions with the exception of non-vertical/non-horizontal purely quadratic gradients.

6 Klein bottle and a family of degree 2 polynomials

There is a pair of two-manifolds whose homology groups with \mathbb{Z}_2 coefficients agree with the result at the end of the previous section (figure 9). These are the torus and the Klein bottle. There is experimental and theoretical evidence in favor of the latter space.




Figure 9: *PLEX* results for $X(100, 10) \cup Q$

The experimental argument goes as follows. First, generate an ideal Klein bottle \mathcal{K} (we did this using the method presented later in this section, see also section 7). Next, consider a filtered simplicial complex for the union $\mathcal{K} \cup (X(100, 10) \cup Q)$ with the following special filtration: all simplices, witnesses for which are points of the ideal Klein bottle, have filtration value 0; all other simplices have filtration values starting at some value $t = t_0 > 0$. Since we know the barcode for \mathcal{K} , we know if the barcode between filtration values 0 and t_0 represents the homology of the Klein bottle. In the case that it does, we only need to observe what happens at t_0 and $t > t_0$. Figure 10 shows the results (with $t_0 = 0.1$). Long lines in the barcode corresponding to generators in H_1 and H_2 start at 0, continue through t_0 , and no other lines appear. This indicates that $X(100, 10) \cup Q$ does not contribute any new generators and none of the existing generators in homology get killed. Therefore, all the non-trivial homology comes from the ideal Klein bottle \mathcal{K} .




Figure 10: *PLEX* results for $\mathcal{K} \cup (X(100, 10) \cup Q)$, $t_0 = 0.1$




Figure 11: Klein bottle, the image of the map g

The following argument gives a theoretical reason to why it is plausible that we are seeing a Klein bottle. It also allows us to generate an ideal Klein bottle inside S^7 .

Consider the space K of all degree two polynomials of the form $c(ax + by)^2 + d(ax + by)$, where a, b, c, d are real parameters such that $(a, b) \in S^1$ and $(c, d) \in S^1$, where S^1 denotes the unit circle in the plane. Varying the 4-tuple (a, b, c, d) on the surface of the torus $S^1 \times S^1$ and evaluating the resulting polynomials at 9 points of H , followed by the subtraction of the mean and division by the norm, we get patches with either linear or quadratic gradients. Since we are dealing with the high-contrast parts of the image we expect that the patches with the above properties occur with high frequency in our space \mathcal{M} .

Define the map $g : S^1 \times S^1 \rightarrow K$ by $(a, b, c, d) \mapsto c(ax + by)^2 + d(ax + by)$. Map g is onto, but not 1-1, since the points (a, b, c, d) and $(-a, -b, c, -d)$ are mapped to the same polynomial. Let us denote the pairs (a, b) and (c, d) by θ and ϕ respectively. Both θ and ϕ vary in the interval $[0, 2\pi]$. In these terms the relation $(a, b, c, d) \sim (-a, -b, c, -d)$ becomes $(\theta, \phi) \sim (\theta + \pi, 2\pi - \phi)$. Since there are no other identifications produced by g , the space $K = im(g)$ is homeomorphic to $S^1 \times S^1 / (\theta, \phi) \sim (\theta + \pi, 2\pi - \phi)$.

The torus has a representation similar to the representation of the Klein bottle as a square with the opposite edges identified, but without the orientation reversal. Figure 11 shows the effect of the map g on the torus. The right half, denoted R in the picture, gets identified by g with the left half, denoted L , with the identifications on the boundary as shown. Each half is the standard representation of the Klein bottle (figure 4, section 4) thus the image of g is homeomorphic to the Klein bottle and hence so is K .

The space K also incorporates the ‘three circle’ space C_3 , described in section 5 above. The primary circle of C_3 is the subspace obtained by setting $(c, d) = (0, 1)$ and letting (a, b) vary on S^1 , while the secondary circles are obtained by setting $\{a = 1, b = 0\}$ on the one hand or $\{a = 0, b = 1\}$ on the other. Denoting the primary circle by C_{lin} and the secondary ones by $C_{1,0}$ and $C_{0,1}$ respectively, we note that their intersection pattern is the one described in section 5, namely $C_{lin} \cap C_{1,0} = \{x, -x\}$, $C_{lin} \cap C_{0,1} = \{y, -y\}$ and $C_{1,0} \cap C_{0,1} = \emptyset$.

Recall the map $q : P \rightarrow S^7$ from section 3. We will need the following result.

Proposition. The restriction of q to the subspace K of P is one-to-one.

The proposition together with a result from general topology that a continuous one-to-one map on a compact space is a homeomorphism onto its image implies that $im(q|_K)$ in S^7 is homeomorphic to the Klein bottle.

Proof.

Let $p = c(ax + by)^2 + d(ax + by)$ and $p' = c'(a'x + b'y)^2 + d'(a'x + b'y)$ be two polynomials such

that their images under q agree, i.e. $q(p) = q(p')$. We show that $p = p'$.

Let us introduce some notation. Denote by m and m' the mean values of the 9-vectors v and v' , obtained by evaluating the polynomials p and p' on the grid H and by k and k' the Euclidean norms of the vectors $v - m \cdot e$ and $v' - m' \cdot e$, where e is a 9-vector with all components equal to one.

Since $q(p) = q(p')$ we have the following two equations in coordinates corresponding to $(1, 0) \in H$ and $(-1, 0) \in H$:

$$(ca^2 + da - m)/k = (c'a'^2 + d'a' - m')/k' \quad (1)$$

$$(ca^2 - da - m)/k = (c'a'^2 - d'a' - m')/k' \quad (2)$$

Subtracting the second equation from the first we get $kda = k'd'a'$. Doing the same for the coordinates corresponding to $(0, 1)$ and $(0, -1)$ we similarly obtain $kdb = k'd'b'$. From these last two equations it follows that either (1) $a/b = a'/b'$ or (2) $d = d' = 0$. Note that $m = 2(a^2 + b^2)/3$, $m' = 2(a'^2 + b'^2)/3$ and $k, k' > 0$.

Consider case (1). Since $(a, b) \in S^1$ and $(a', b') \in S^1$ there are only two possibilities. Either (a) $(a, b) = (a', b')$ and hence $d = d'$, in which case either $c = c'$ giving $p = p'$ or $c = -c'$. In this last case $m = -m'$ and looking at the coordinate corresponding to $(0, 0)$ we get that $m = m' = 0$ ($-mk = -m'k'$ with $k, k' > 0$) and hence $c = -c' = 0$ implying $p = p'$. Or (b) $(a, b) = -(a', b')$, in which case $d = -d'$ and hence either $p = p'$ or $c = c'$, and the last situation is treated the same way as in the case (a), i.e. appealing to the $(0, 0)$ -coordinate.

In case (2) we have two essentially distinct subcases: (a) $c = 1, c' = -1$ and (b) $c = c' = 1$.

In case (a) we get $m > 0, m' < 0$, and equality of vectors at the $(0, 0)$ -coordinate implies that $m = m' = 0$ hence $c = c' = 0$, a contradiction. It remains to consider case (b). We have $c = c' = 1, d = d' = 0$. Looking at the $(1, 1)$ and $(-1, -1)$ -coordinates gives the following pair of equations:

$$(a^2 + b^2 + 2ab - m)/k = (a'^2 + b'^2 + 2a'b' - m')/k' \quad (3)$$

$$(a^2 + b^2 - 2ab - m)/k = (a'^2 + b'^2 - 2a'b' - m')/k' \quad (4)$$

Subtracting the second equation from the first yields $ab/k = a'b'/k'$ or $ab = (k/k')a'b'$. At the same time the $(1, 0)$ and $(0, 1)$ -coordinates give the following two equations:

$$(a^2 - m)/k = (a'^2 - m')/k' \quad (5)$$

$$(b^2 - m)/k = (b'^2 - m')/k' \quad (6)$$

Subtracting (6) from (5) yields $(a^2 - b^2)/k = (a'^2 - b'^2)/k'$ or $(a^2 - b^2) = (k/k')(a'^2 - b'^2)$. Rewriting $(a, b) = (\cos(\alpha), \sin(\alpha))$ and $(a', b') = (\cos(\beta), \sin(\beta))$ and recalling the trigonometric identities $\sin(2x) = 2 \sin(x) \cos(x)$ and $\cos(2x) = \cos(x)^2 - \sin(x)^2$ the previous two equations

$$ab = (k/k')a'b' \quad (7)$$

$$a^2 - b^2 = (k/k')(a'^2 - b'^2) \quad (8)$$

become

$$\sin(2\alpha) = (k/k') \sin(2\beta) \quad (9)$$

$$\cos(2\alpha) = (k/k') \cos(2\beta) \quad (10)$$

There are two solutions $\alpha = \beta$, i.e. $(a, b) = (a', b')$ and $\alpha = -\beta$, i.e. $(a, b) = -(a', b')$. In both cases $p = p'$.

Hence the map q is one-to-one as claimed. \square

We use the proposition to show that we obtain a ‘three circle’ space as a subset of $im(q|_K)$ inside S^7 .

Since $q|_K$ is one-to-one, the image of $C_3 \subseteq K$ in S^7 is homeomorphic to C_3 , since C_3 is compact. It remains to show that any patch with linear gradient or quadratic gradient in the vertical or horizontal direction is in the set $im(q|_{C_3})$, the image of q restricted to C_3 . This is obvious in the linear gradient case. Since the vertical and horizontal cases are identical we concentrate on the latter.

Any patch with quadratic gradient in the horizontal direction can be obtained by evaluating on H a polynomial $p = cx^2 + bx + a$, whose mean value $m(p)$ on H is equal to zero and whose norm $k(p) = \sum_H p(x)^2$ is one. The coefficients of polynomials with these two properties satisfy the following pair of equations:

$$3a + 2c = 0 \tag{11}$$

$$6(a + c)^2 + 3a^2 + 6b^2 = 1 \tag{12}$$

The first equation restricts the possible values of (a, b, c) to a plane in (a, b, c) -space, and the second defines an ellipsoid. The intersection of the two will be an ellipse, which is topologically equivalent (homeomorphic) to a circle. Therefore, the space of all patches with quadratic gradient in the horizontal direction (denoted S_h) is also topologically a circle in S^7 (quadratic function in one variable is uniquely determined by its values at three distinct points, here these are $x = \{-1, 0, 1\}$). On the other hand the image of $C_{1,0}$ under q lies inside S_h and is also homeomorphic to a circle. Hence $S_h = im(q|_{C_{1,0}})$, since if $im(q|_{C_{1,0}})$ missed at least one point of S_h it would be contractible. The same argument applies to the case of patches with quadratic gradient in the vertical direction, yielding the ‘three circle’ space sitting inside the set $im(q|_K)$.

Taking the point of view that the space \mathcal{M} should be homeomorphic to some CW -space \mathcal{M}_{cw} (cf. [4]), we can think of the spaces $im(q|_{C_3})$ and $im(q|_K)$ as being the first and the second skeleton respectively of the cell structure of \mathcal{M}_{cw} .

7 Embedding of the Klein bottle into S^7

The homology of spaces $X(100, p)$ with $p > 10$ no longer agrees with the homology of the Klein bottle (cf. figure 9). Probably the sporadic high-density regions due to sampling within the space \mathcal{M} start to take effect as the cut parameter increases beyond 10.

We try a different approach. The idea is to first embed the idealized Klein bottle into the 7-sphere where the set \mathcal{M} lives and then move it towards higher density regions making sure that its topology stays intact along the way, i.e. there is no tearing or gluing of the surface as it moves within \mathbb{R}^8 .

As an initial step for the embedding of the Klein bottle into S^7 we select 1600 points from the torus $S^1 \times S^1$ by taking the set of uniformly spaced 40 points on each of the circles (denoted $\{a_1, \dots, a_{40}\}$ and $\{b_1, \dots, b_{40}\}$ respectively) and forming all possible tuples (a_i, b_j) . Next, to each of the 1600 points we apply the map $h \circ g$, where g is defined in section 6 and h is the composite of subtraction of the mean from each coordinate, normalization by the D -norm, and transformation into the DCT basis. This gives a point cloud data in S^7 .

Due to the identification $(\theta, \phi) \sim (\theta + \pi, 2\pi - \phi)$ produced by g , not all 1600 points are going to be distinct after applying the above procedure. In fact, we obtain a total of 939 distinct points as our point cloud data for the embedded Klein bottle. Let us denote this set by \mathcal{K} . Figure 12 shows the *PLEX* result for the homology of the space \mathcal{K} . The barcode supports the theoretical conclusion of the previous section about the topology of the embedded space giving \mathbb{Z}_2 -homology of the Klein bottle. By construction the points of \mathcal{K} are patches with linear or quadratic gradients.




Figure 12: *PLEX* results for \mathcal{K}

8 Moving the initial embedding

We would like to understand just how much of \mathcal{M} is captured by the Klein bottle model, i.e. what is the largest subspace in \mathcal{M} that still has the homology of the Klein bottle. Recall in this context the result from figure 9 which suggests that the space consisting of 10 percent of the densest points of \mathcal{M} (with density estimation parameter equal to 100) does have the Klein bottle’s homology. Going beyond 10 percent for the value of the cut parameter, however, changes the topological type of the space $X(100, p)$. As was mentioned in the beginning of the previous section, what we try to do instead is to make the local density of \mathcal{M} guide our initial Klein bottle point cloud data toward the region where it best fits the actual data.

Let us rename the initial embedding \mathcal{K} by \mathcal{K}_0 and denote by $\mathcal{A}_{0,p}$ the subspace of \mathcal{M} obtained by taking p percent of the closest points to \mathcal{K}_0 . More specifically, for each point of \mathcal{M} we find the Euclidean distance from this point to the closest point of \mathcal{K}_0 , then take the top p percent of the points for which this distance is smallest. As p gets larger we expect the space $\mathcal{A}_{0,p}$ to look less and less like a Klein bottle. Indeed, at some value of p slightly above 50 the space undergoes a topological change. Figure 13 shows the result for the space $\mathcal{A}_{0,50}$, which is the largest subspace \mathcal{M} we found giving a true topological approximation to \mathcal{K}_0 .

To move the initial embedding \mathcal{K}_0 towards higher density region we use a variant of a self-organizing map. For each point x_0 of \mathcal{K}_0 we find the densest point in its Voronoi cell V_{x_0} , where V_{x_0} consists of all points y of \mathcal{M} such that $d(x_0, y) \leq d(x, y)$ for all $x \in \mathcal{K}_0$. We move each point x of \mathcal{K}_0 towards the densest point in its Voronoi cell by an amount a , which we set experimentally (the value for a used in obtaining the results of this section is 0.3). This produces a new space \mathcal{K}_1 , and we iterate this procedure. At each step of the iteration we compute the homology of the space \mathcal{K}_i to make sure that it still has the original topology of \mathcal{K}_0 . We also compute the homology of spaces $\mathcal{A}_{i,p}$, starting with $p = 5\%$ and using increments of 5%.

Figures 14 and 15 show homological results for \mathcal{K}_2 and $\mathcal{A}_{2,60}$ respectively. The latter space is the largest for which we still recover the homology of the Klein bottle though only for a limited range. In other words, the topology of the space consisting of 60 percent of the points of \mathcal{M} can be sufficiently well approximated by the Klein bottle model.

In obtaining the experimental homological results of this section we used the following procedure. First, select at random a subset S_0 with $|S_0| = 10,000$ from $\mathcal{A}_{i,p}$. Apply the denoising algorithm (two iterations)




Figure 13: *PLEX* results for $\mathcal{A}_{0,50}$

to S_0 with the neighborhood size parameter set to 20, obtaining a new space S_2 (see section 1.4). Compute the homology of the space S_2 . Repeat many times, each time selecting a new set S_0 to make sure the results are stable and convey the actual homology of $\mathcal{A}_{i,p}$.

To get a better understanding of the space \mathcal{K}_2 we obtain a histogram of distances between this perturbed ideal Klein bottle and points of the space $\mathcal{A}_{2,60}$ (figure 16). Using Monte Carlo method (10^4 iterations) we found that while the volume of the entire dataset \mathcal{M} is 84% of the volume of S^7 , the volume of $\mathcal{A}_{2,60}$ is only 21%, indicating that the neighborhood of the perturbed Klein bottle has a very high concentration of points of \mathcal{M} .

Summary and concluding remarks

In this paper we continued the qualitative (topological) approach to the study of the space of 3 by 3 patches coming from natural images initiated in [8]. The key advantage of this approach is that it allows one to find highly non-linear yet extremely important subsets within the data which would be very hard to discover using statistical techniques.

We established a correspondence between the space of patches and the space of polynomials in two variables and used it as an initial step in finding that there is a large portion of \mathcal{M} which is topologically equivalent to a well-known two-manifold, the Klein bottle. We were using *PLEX* topological software to experimentally verify our theoretical conjectures about the qualitative structure of various subspaces within \mathcal{M} .

If we think of the 2-manifold that we found as representing the two-skeleton of \mathcal{M} further development in this direction may concentrate on trying to find the three- and higher-dimensional skeletons of the space \mathcal{M} , thereby giving it the structure of a *CW*-complex. Another interesting way to proceed which is currently being explored is to use the Klein bottle embedded into S^7 for the purposes of image compression, i.e. create an efficient and reliable algorithm which replaces the high-contrast patches of the given image by the points on the surface of the embedded Klein bottle that best approximate these high-contrast patches. This paves the way for very useful practical applications.




Figure 14: *PLEX* results for \mathcal{K}_2

Acknowledgments

We gratefully acknowledge the support of DARPA through grants 1092228 and HR0011-06-1-0038 and by NSF through DMS 0354543. The work was carried out at the Department of Mathematics at Stanford University.

References

- [1] Zomorodian A., Carlsson G. *Computing persistent homology*. 20th ACM Symposium on Computational Geometry (2004)
- [2] van Hateren J.H., van der Schaaf A. *Independent component filters of natural images compared with simple cells in primary visual cortex*. Proc.R.Soc.Lond. B 265(1998), 359-366
- [3] Lee A.B., Pedersen K.S., Mumford D. *The non-linear statistics of high-contrast patches in natural images*. International Journal of Computer Vision 54, 1-3(2003), 83-103
- [4] Hatcher A. *Algebraic topology*. Cambridge University Press (2001)
- [5] de Silva V. *A weak definition of Delaunay triangulation*. (2003)
- [6] Edelsbrunner H., Letscher D., Zomorodian A. *Topological persistence and simplification*. IEEE Symposium on Foundations of Computer Science (2000)
- [7] de Berg M., van Kreveld M., Overmars M., Schwarzkopf O. *Computational Geometry: Algorithms and Applications*. Springer, 2nd edition (2000)
- [8] Carlsson G., de Silva V. *Topological estimation using witness complexes*. Symposium on Point-Based Graphics (2004)
- [9] Silverman B.W. *Density Estimation for statistics and data analysis*. Chapman & Hall/CRC (1986)




Figure 15: *PLEX* results for $\mathcal{A}_{2,60}$

- [10] Carlsson G., Ishkhanov T., Mémoli F., Ringach D., Sapiro G. with Singh G. *Topological analysis of the Space of Responses of Neurons in V1*. [in preparation]




Figure 16: Histogram of distances from $\mathcal{A}_{2,60}$ to \mathcal{K}_2