Lecture 4: Application to Natural Images

May 14, 2019

Gunnar Carlsson

Stanford University

Application to Natural Image Statistics

With V. de Silva, T. Ishkanov, A. Zomorodian

Hubel-Wiesel Experiment

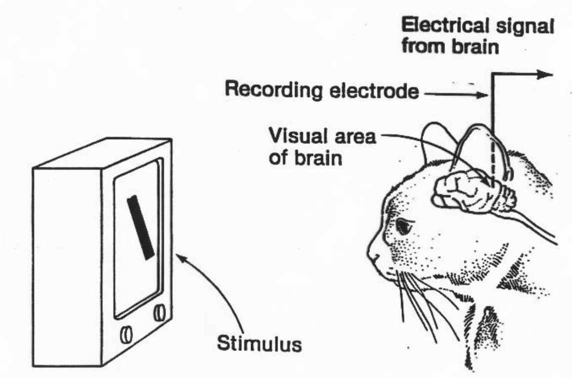

van Hateren-van der Schaaf Data Set

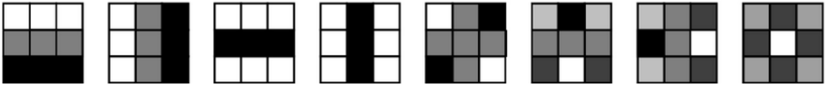

Natural Images

An image taken by black and white digital camera can be viewed as a vector, with one coordinate for each pixel

Natural Images

An image taken by black and white digital camera can be viewed as a vector, with one coordinate for each pixel

Each pixel has a “gray scale” value, can be thought of as a real number (in reality, takes one of 255 values)

Natural Images

An image taken by black and white digital camera can be viewed as a vector, with one coordinate for each pixel

Each pixel has a "gray scale" value, can be thought of as a real number (in reality, takes one of 255 values)

Typical camera uses tens of thousands of pixels, so images lie in a very high dimensional space, call it *pixel space*, $\mathcal{P}$

Natural Images

D. Mumford: What can be said about the set of images $\mathcal{I} \subseteq \mathcal{P}$ one obtains when one takes many images with a digital camera?

Natural Images

1. $\mathcal{I}$ is very high dimensional, because images are so *expressive*

Natural Images

1. $\mathcal{I}$ is very high dimensional, because images are so *expressive*

2. $\mathcal{I}$ is very sparse (of high codimension) in $\mathcal{P}$, because a random choice of values for each pixel will not give anything close to an image

Natural Images

1. $\mathcal{I}$ is very high dimensional, because images are so *expressive*

2. $\mathcal{I}$ is very sparse (of high codimension) in $\mathcal{P}$, because a random choice of values for each pixel will not give anything close to an image

3. The direct study of $\mathcal{I}$ will be very difficult, unless one severely restricts the range of subjects for the images

Natural Images

Solution (Lee, Mumford, Pedersen): Study *local* structure of images statistically, where there is less variation

Natural Images

Solution (Lee, Mumford, Pedersen): Study *local* structure of images statistically, where there is less variation

Specifically, study 3×3 patches in the image.

Patches

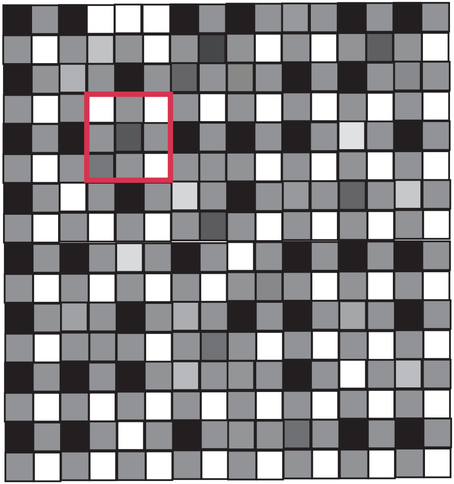

Patches

Observations:

Patches

Observations:

1. Each patch gives a vector in $\mathbb{R}^9$

Patches

Observations:

1. Each patch gives a vector in $\mathbb{R}^9$
2. Most patches will be nearly constant, or *low contrast*, because of the presence of regions of solid shading in most images

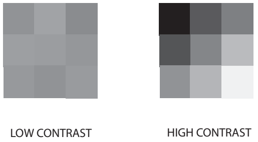

Patches

Observations:

1. Each patch gives a vector in $\mathbb{R}^9$
2. Most patches will be nearly constant, or *low contrast*, because of the presence of regions of solid shading in most images

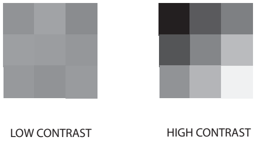

3. Low contrast will dominate statistics, not interesting

Patches

- Lee-Mumford-Pedersen [LMP] study only high contrast patches

Patches

- Lee-Mumford-Pedersen [LMP] study only high contrast patches
- Contrast measured by quadratic form on space of patches, so-called *D-norm*

Patches

- Lee-Mumford-Pedersen [LMP] study only high contrast patches
- Contrast measured by quadratic form on space of patches, so-called *D-norm*
- Restricts study to parts of the images where transitions in intensity occur, i.e. edges, lines, etc.

Patches

- Lee-Mumford-Pedersen [LMP] study only high contrast patches
- Contrast measured by quadratic form on space of patches, so-called *D-norm*
- Restricts study to parts of the images where transitions in intensity occur, i.e. edges, lines, etc.
- How do they study the high contrast patches?

Patches

- Collect c:a 4.5×10^6 high contrast patches from the collection of images obtained by van Hateren and van der Schaaf

Patches

- Collect c:a 4.5×10^6 high contrast patches from the collection of images obtained by van Hateren and van der Schaaf
- Normalize mean intensity by subtracting mean from each pixel value to obtain patches with mean intensity $= 0$

Patches

- Collect c:a 4.5×10^6 high contrast patches from the collection of images obtained by van Hateren and van der Schaaf
- Normalize mean intensity by subtracting mean from each pixel value to obtain patches with mean intensity $= 0$
- Puts data on an 8-dimensional hyperplane, $\cong \mathbb{R}^8$

Patches

- Collect c:a 4.5×10^6 high contrast patches from the collection of images obtained by van Hateren and van der Schaaf
- Normalize mean intensity by subtracting mean from each pixel value to obtain patches with mean intensity $= 0$
- Puts data on an 8-dimensional hyperplane, $\cong \mathbb{R}^8$
- Means that we will consider as equivalent patches which can be obtained from each other by turning the intensity knob

Patches

- Normalize contrast by dividing by the D-norm, so obtain patches with D-norm $= 1$

Patches

- Normalize contrast by dividing by the D-norm, so obtain patches with D-norm $= 1$
- Means that data now lies on a 7-D ellipsoid, $\cong S^7$

Patches

- Normalize contrast by dividing by the D-norm, so obtain patches with D-norm $= 1$
- Means that data now lies on a 7-D ellipsoid, $\cong S^7$
- Normalization means that we will consider patches which can be obtained from each other by turning contrast knob to be the same

Patches

Patches

Result: Point cloud data $\mathcal{M}$ lying on a sphere in $\mathbb{R}^8$

Patches

Result: Point cloud data $\mathcal{M}$ lying on a sphere in $\mathbb{R}^8$

We wish to analyze it with persistent homology to understand it *qualitatively*

Analysis

First Observation: The points fill out S^7 in the sense that every point in S^7 is "close" to a point in $\mathcal{M}$

Analysis

First Observation: The points fill out S^7 in the sense that every point in S^7 is "close" to a point in $\mathcal{M}$

Initially disappointing, since it means that nothing special can be said about the actual patches different from patches chosen at random

Analysis

First Observation: The points fill out S^7 in the sense that every point in S^7 is “close” to a point in $\mathcal{M}$

Initially disappointing, since it means that nothing special can be said about the actual patches different from patches chosen at random

However, density of points varies a great deal from region to region

Analysis

First Observation: The points fill out S^7 in the sense that every point in S^7 is “close” to a point in $\mathcal{M}$

Initially disappointing, since it means that nothing special can be said about the actual patches different from patches chosen at random

However, density of points varies a great deal from region to region

Study the subsets of high density, as measured by different density estimators.

Primary Circle

$$5 \times 10^4 \text{ points}, \ k = 300, \ T = 25$$

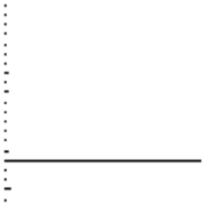

One-dimensional barcode, suggests $\beta_1 = 1$

Primary Circle

$$5 \times 10^4 \text{ points, } k = 300, T = 25$$

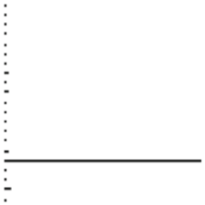

One-dimensional barcode, suggests $\beta_1 = 1$

Is the set clustered around a circle?

Primary Circle

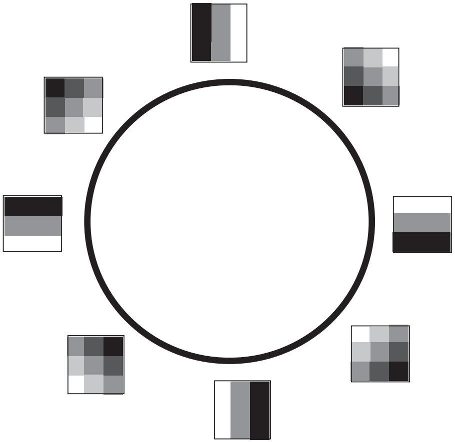

PRIMARY CIRCLE

Three Circle Model

$$5 \times 10^4 \text{ points}, \ k = 15, \ T = 25$$

One-dimensional barcode, suggests $\beta_1 = 5$

Three Circle Model

5×10^4 points, $k = 15$, $T = 25$

One-dimensional barcode, suggests $\beta_1 = 5$

What's the explanation for this?

Three Circle Model

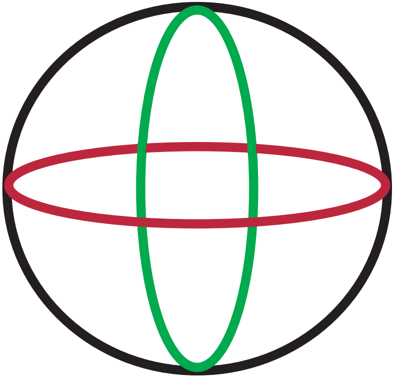

Three Circle Model

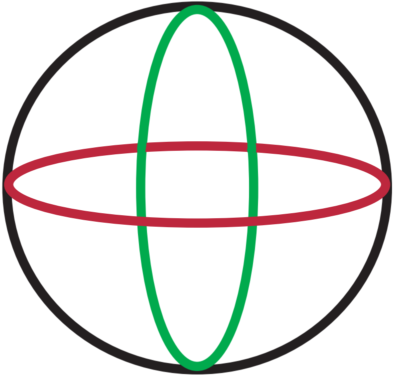

THREE CIRCLE MODEL

Three Circle Model

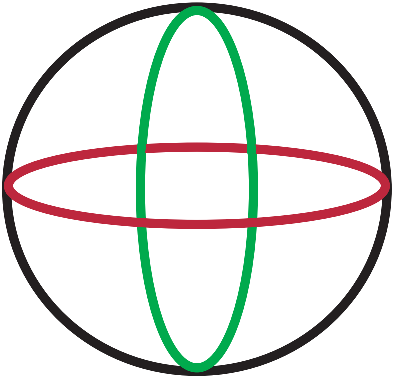

Red and green circles do not touch, each touches black circle

Three Circle Model

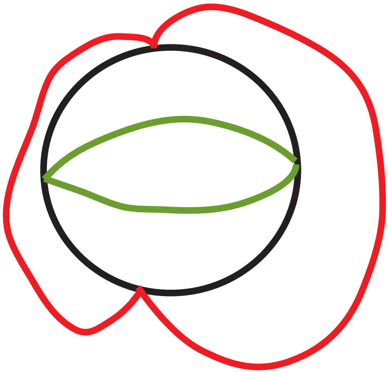

Three Circle Model

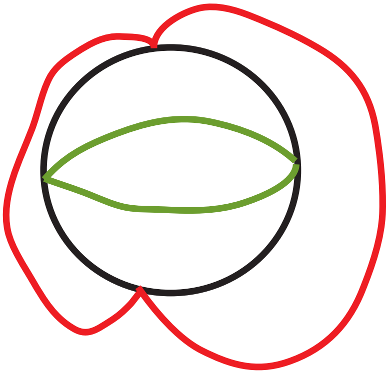

$$\beta_1 = 5$$

Three Circle Model

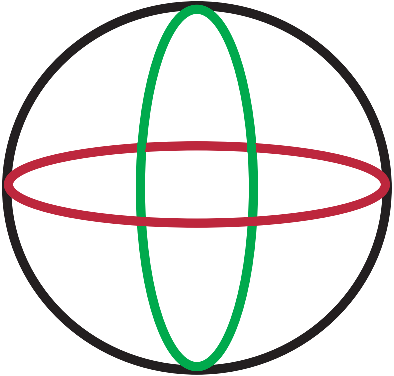

Does the data fit with this model?

Three Circle Model

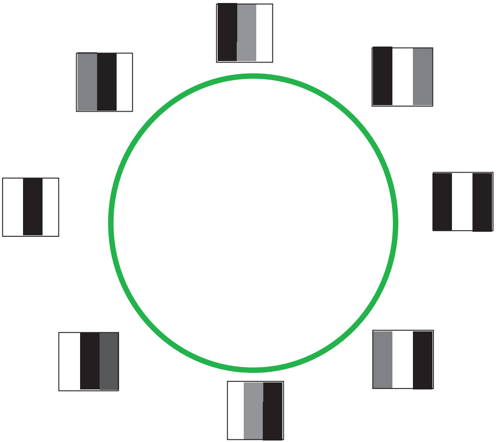

SECONDARY CIRCLE

Three Circle Model

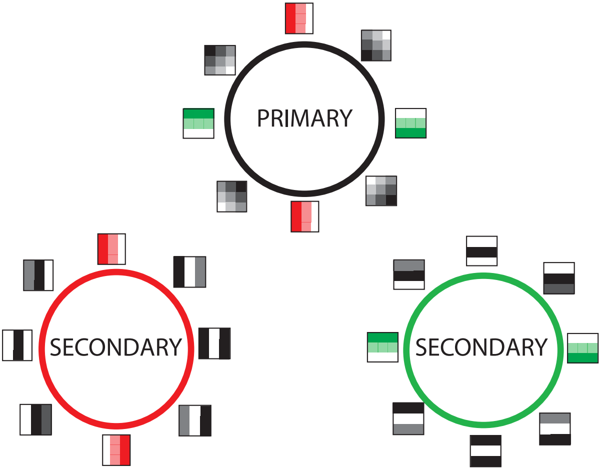

Database

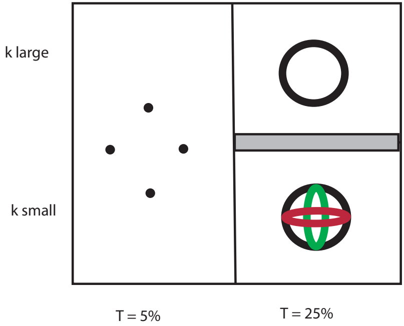

Three Circle Model

IS THERE A TWO DIMENSIONAL SURFACE IN WHICH THIS PICTURE FITS?

More Relaxed Threshold

4.5×10^6 points, $k = 100$, $T = 10$

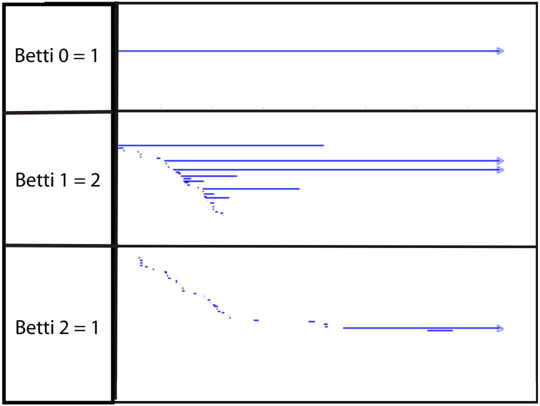

Surfaces

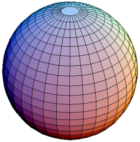

Sphere

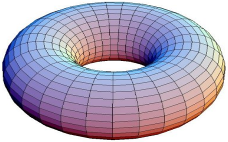

Torus

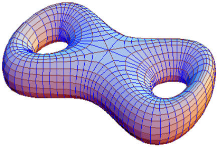

Two-holed Torus

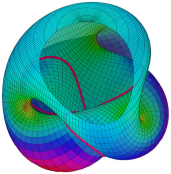

Projective Plane = P^2

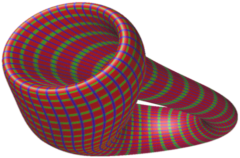

Klein Bottle = K^2

. . .

Betti Numbers

	S^2	Torus	2-holed Torus	P^2	K^2
mod 2 β_0	1	1	1	1	1
mod 2 β_1	0	2	4	1	2
mod 2 β_2	1	1	1	1	1
mod 3 β_0	1	1	1	1	1
mod 3 β_1	0	2	4	0	1
mod 3 β_2	1	1	1	0	0

Betti Numbers

- How are these numbers obtained?

Betti Numbers

- How are these numbers obtained?
- Uses methods from Lecture 2

Betti Numbers

- How are these numbers obtained?
- Uses methods from Lecture 2
- Decompose the spaces into pieces

Betti Numbers of Torus

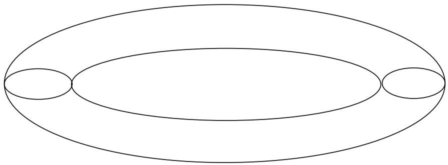

Torus

Betti Numbers of Torus

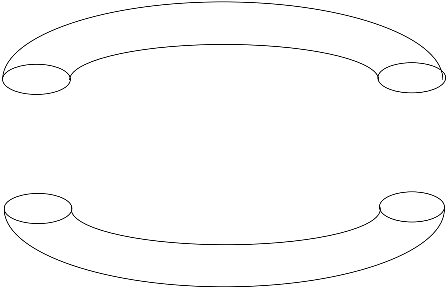

Torus decomposed into two pieces U and V

Betti Numbers of Torus

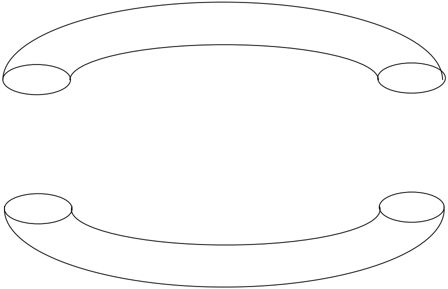

U and V are both homotopy equivalent to circle S^1

Betti Numbers of Torus

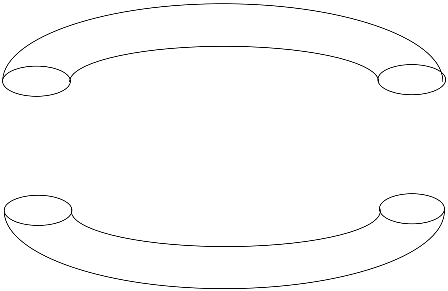

$$\beta_0(U) = \beta_0(V) = 1 \text{ and } \beta_1(U) = \beta_1(V) = 1$$

Betti Numbers of Torus

$U \cap V$

Betti Numbers of Torus

$$\beta_0(U \cap V) = 2 \text{ and } \beta_1(U \cap V) = 2$$

Betti Numbers of Torus

$$\to H_2(U) \oplus H_2(V) \to H_2(T) \to H_1(U \cap V) \to H_1(U) \oplus H_1(V) \to$$

$$\to H_1(T) \to H_0(U \cap V) \to H_0(U) \oplus H_0(V) \to H_0(T)$$

Mayer-Vietoris exact sequence

Betti Numbers of Torus

$$\to 0 \oplus 0 \to H_2(T) \to k \oplus k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(T) \to k \oplus k \xrightarrow{\beta} k \oplus k \to H_0(T)$$

Mayer-Vietoris exact sequence

Betti Numbers of Torus

$$\to 0 \oplus 0 \to H_2(T) \to k \oplus k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(T) \to k \oplus k \xrightarrow{\beta} k \oplus k \to H_0(T)$$

The maps α and β both have matrix $\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$, so have rank $=1$

Betti Numbers of Torus

Gives $\beta_0(T) = 1$, $\beta_1(T) = 2$, and $\beta_2(T) = 1$. Holds true for any coefficient field.

Klein Bottle

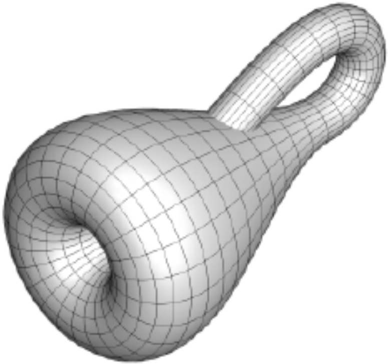

K^2 - Klein Bottle

Betti Numbers of Klein Bottle

Klein Bottle - Identification Space Representation

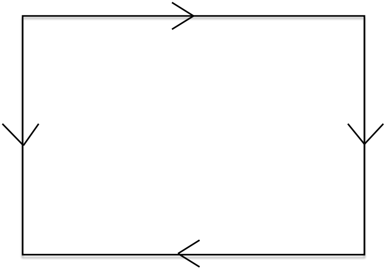

Betti Numbers of Klein Bottle

Klein Bottle - Identification Space Representation

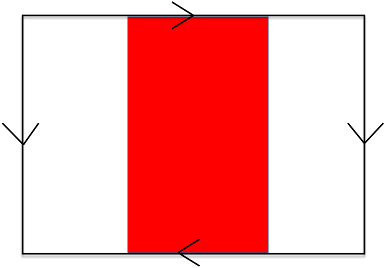

U

Betti Numbers of Klein Bottle

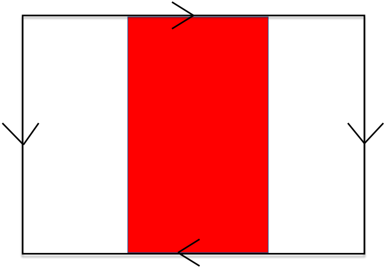

Möbius Band

Betti Numbers of Klein Bottle

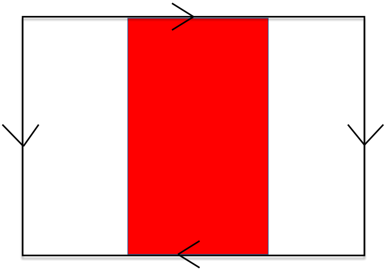

Homotopy Equivalent to Circle

Betti Numbers of Klein Bottle

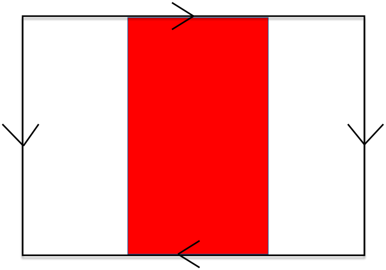

$$\beta_0(U) = 1,\ \beta_1(U) = 1$$

Betti Numbers of Klein Bottle

Klein Bottle - Identification Space Representation

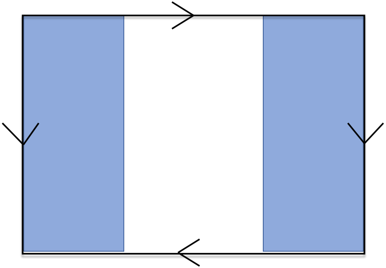

V

Betti Numbers of Klein Bottle

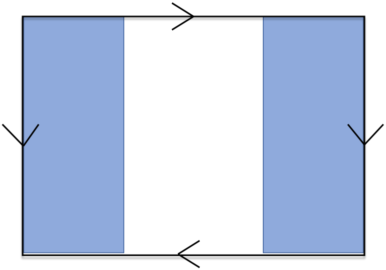

Möbius Band

Betti Numbers of Klein Bottle

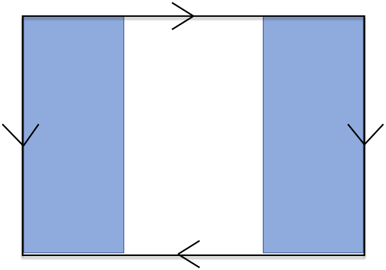

$$\beta_0(V) = 1,\ \beta_1(V) = 1$$

Betti Numbers of Klein Bottle

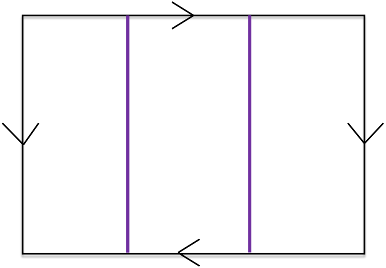

$U \cap V$

Betti Numbers of Klein Bottle

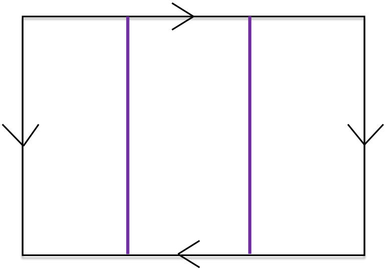

Circle

Betti Numbers of Klein Bottle

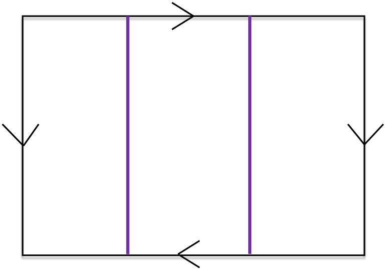

$$\beta_0(U \cap V) = 1, \beta_1(U \cap V) = 1$$

Betti Numbers of Klein Bottle

$$\to H_2(U) \oplus H_2(V) \to H_2(K) \to H_1(U \cap V) \to H_1(U) \oplus H_1(V) \to$$

$$\to H_1(K) \to H_0(U \cap V) \to H_0(U) \oplus H_0(V) \to H_0(K)$$

Mayer-Vietoris exact sequence

Betti Numbers of Klein Bottle

$$\to 0 \oplus 0 \to H_2(K) \to k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(K) \to k \xrightarrow{\beta} k \oplus k \to H_0(K)$$

Mayer-Vietoris exact sequence

Betti Numbers of Klein Bottle

$$\to 0 \oplus 0 \to H_2(K) \to k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(K) \to k \xrightarrow{\beta} k \oplus k \to H_0(K)$$

The map α has matrix $\begin{bmatrix} 2 \\ 2 \end{bmatrix}$ and β has matrix $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$

Betti Numbers of Klein Bottle

$$\to 0 \oplus 0 \to H_2(K) \to k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(K) \to k \xrightarrow{\beta} k \oplus k \to H_0(K)$$

Means behavior for $\mathbb{F}_2$ is different from behavior for $\mathbb{F}_p$ (p odd) and $\mathbb{Q}$

Betti Numbers of Klein Bottle

$$\to 0 \oplus 0 \to H_2(K) \to k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(K) \to k \xrightarrow{\beta} k \oplus k \to H_0(K)$$

For $\mathbb{F}_2$, get $\beta_0(K) = 1$, $\beta_1(K) = 2$, and $\beta_2(K) = 1$

Betti Numbers of Klein Bottle

$$\to 0 \oplus 0 \to H_2(K) \to k \xrightarrow{\alpha} k \oplus k \to$$

$$\to H_1(K) \to k \xrightarrow{\beta} k \oplus k \to H_0(K)$$

For $\mathbb{F}_p$ and $\mathbb{Q}$, get $\beta_0(K) = 1$, $\beta_1(K) = 1$, and $\beta_2(K) = 0$

Do the three circles fit naturally inside $\mathcal{K}$?

Klein Bottle

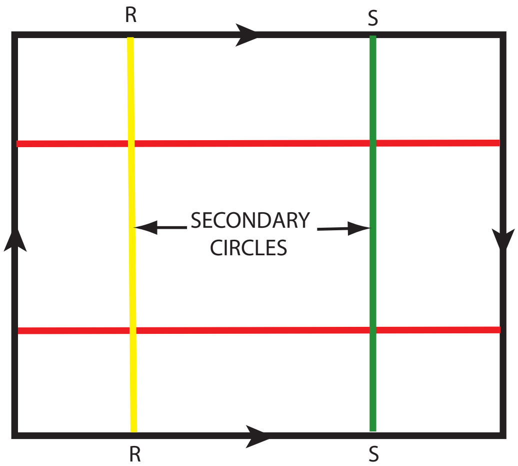

Mapping Patches

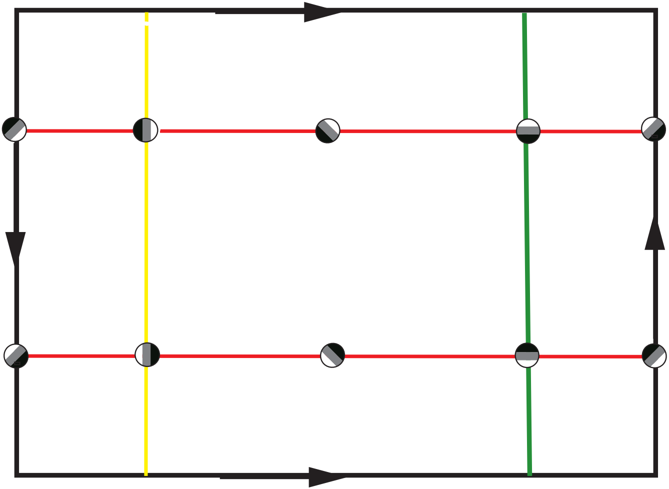

Mapping Patches

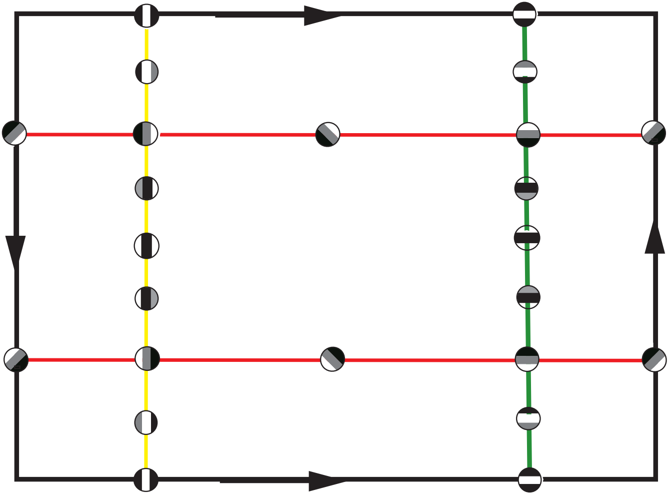

Mapping Patches

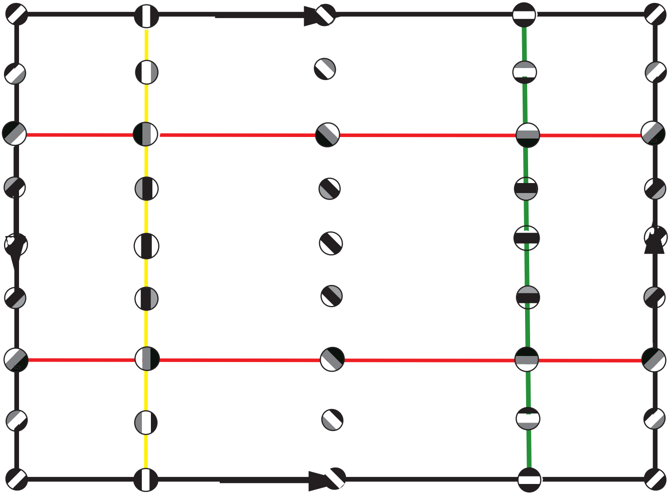

Natural Image Statistics

Klein bottle makes sense as the space K of quadratic polynomials in two variables on the unit disc D, as the space of quadratic polynomials of the form

$$f = q(\lambda(x))$$

satisfying

1. q is single variable quadratic
2. λ is a linear functional
3. $\int_D f = 0$
4. $\int_D f^2 = 1$

Natural Image Statistics

Klein bottle makes sense as the space K of quadratic polynomials in two variables on the unit disc D, as the space of quadratic polynomials of the form

$$f = q(\lambda(x))$$

satisfying

1. q is single variable quadratic
2. λ is a linear functional
3. $\int_D f = 0$
4. $\int_D f^2 = 1$

Why is this space homeomorphic to a Klein bottle?

Natural Image Statistics

- Consider first the space Q_2 of one variable quadratic functions on $[-1, 1]$, with mean value $= 0$ and with L^2-norm equal to 1

Natural Image Statistics

- Consider first the space Q_2 of one variable quadratic functions on $[-1, 1]$, with mean value $= 0$ and with L^2-norm equal to 1
- Orthonormal basis for the space of degree two polynomials on $[-1, 1]$ is given by the first three Legendre polynomials $P_0 = 1, P_1 = x$, and $P_2 = \frac{1}{2}(3x^2 - 1)$.

Natural Image Statistics

- Consider first the space Q_2 of one variable quadratic functions on $[-1, 1]$, with mean value $= 0$ and with L^2-norm equal to 1
- Orthonormal basis for the space of degree two polynomials on $[-1, 1]$ is given by the first three Legendre polynomials $P_0 = 1, P_1 = x$, and $P_2 = \frac{1}{2}(3x^2 - 1)$.
- A quadratic polynomial $f = c_0 P_0 + c_1 P_1 + c_2 P_2$ satisfies the conditions described above if (a) $c_0 = 0$ and (b) $c_1^2 + c_2^2 = 1$

Natural Image Statistics

- Consider first the space Q_2 of one variable quadratic functions on $[-1,1]$, with mean value $= 0$ and with L^2-norm equal to 1
- Orthonormal basis for the space of degree two polynomials on $[-1,1]$ is given by the first three Legendre polynomials $P_0 = 1, P_1 = x$, and $P_2 = \frac{1}{2}(3x^2 - 1)$.
- A quadratic polynomial $f = c_0P_0 + c_1P_1 + c_2P_2$ satisfies the conditions described above if (a) $c_0 = 0$ and (b) $c_1^2 + c_2^2 = 1$
- Consequently Q_2 is homeomorphic to a circle

Natural Image Statistics

- Consider first the space Q_2 of one variable quadratic functions on $[-1, 1]$, with mean value $= 0$ and with L^2-norm equal to 1
- Orthonormal basis for the space of degree two polynomials on $[-1, 1]$ is given by the first three Legendre polynomials $P_0 = 1, P_1 = x$, and $P_2 = \frac{1}{2}(3x^2 - 1)$.
- A quadratic polynomial $f = c_0P_0 + c_1P_1 + c_2P_2$ satisfies the conditions described above if (a) $c_0 = 0$ and (b) $c_1^2 + c_2^2 = 1$
- Consequently Q_2 is homeomorphic to a circle
- Parametrize it by an angle θ_1, so $c_1 = sin\theta_1$, $c_2 = cos\theta_1$.

Natural Image Statistics

- For any angle θ_2, let $\varphi_{\theta_2}(x, y) = x cos\theta_2 + y sin\theta_2$. Let L denote the space of all functions φ_θ on the unit disc in $\mathbb{R}^2$.

Natural Image Statistics

- For any angle θ_2, let $\varphi_{\theta_2}(x, y) = xcos\theta_2 + ysin\theta_2$. Let L denote the space of all functions φ_θ on the unit disc in $\mathbb{R}^2$.
- The space L is homeomorphic to the circle.

Natural Image Statistics

- For any angle θ_2, let $\varphi_{\theta_2}(x, y) = xcos\theta_2 + ysin\theta_2$. Let L denote the space of all functions φ_θ on the unit disc in $\mathbb{R}^2$.
- The space L is homeomorphic to the circle.
- The formula $\Psi(\theta_1, \theta_2)(x, y) = f_{\theta_1}(\varphi_{\theta_2}(x, y))$ gives a continuous map $Q_2 \times S^1 \longrightarrow K$.

Natural Image Statistics

- For any angle θ_2, let $\varphi_{\theta_2}(x, y) = xcos\theta_2 + ysin\theta_2$. Let L denote the space of all functions φ_θ on the unit disc in $\mathbb{R}^2$.
- The space L is homeomorphic to the circle.
- The formula $\Psi(\theta_1, \theta_2)(x, y) = f_{\theta_1}(\varphi_{\theta_2}(x, y))$ gives a continuous map $Q_2 \times S^1 \longrightarrow K$.
- The map ψ is surjective

Natural Image Statistics

- For any angle θ_2, let $\varphi_{\theta_2}(x, y) = xcos\theta_2 + ysin\theta_2$. Let L denote the space of all functions φ_θ on the unit disc in $\mathbb{R}^2$.
- The space L is homeomorphic to the circle.
- The formula $\Psi(\theta_1, \theta_2)(x, y) = f_{\theta_1}(\varphi_{\theta_2}(x, y))$ gives a continuous map $Q_2 \times S^1 \longrightarrow K$.
- The map ψ is surjective
- The map ψ satisfies the condition $\psi(\theta_1, \theta_2 + \pi) = \psi(-\theta_1, \theta_2)$

Natural Image Statistics

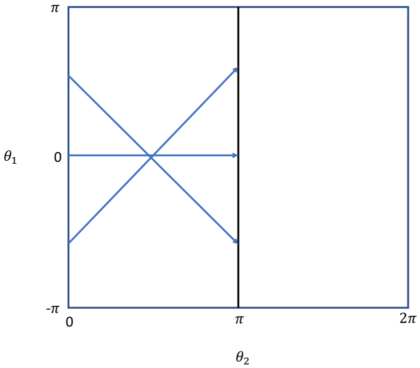

Range Patches

- Range camera records distance to image rather than optical intensity

Range Patches

- Range camera records distance to image rather than optical intensity

- Mumford et al also constructed a database of high contrast range images

Range Patches

- Range camera records distance to image rather than optical intensity

- Mumford et al also constructed a database of high contrast range images
- Initial analysis of 3×3 patches revealed interesting structures in the form of clusters, but no systematic overview

Range Patches

- Persistent homology analysis of range patches joint with H. Adams

Range Patches

- Persistent homology analysis of range patches joint with H. Adams
- Construct spaces of 3×3 patches $X(k, T)$, with density estimator associated with k, and T a percentage threshhold.

Range Patches

- Persistent homology analysis of range patches joint with H. Adams
- Construct spaces of 3×3 patches $X(k, T)$, with density estimator associated with k, and T a percentage threshhold.
- No higher dimensional structure, but β_0 barcodes with long lines, indicating clustering.

Range Patches

- Persistent homology analysis of range patches joint with H. Adams
- Construct spaces of 3×3 patches $X(k, T)$, with density estimator associated with k, and T a percentage threshhold.
- No higher dimensional structure, but β_0 barcodes with long lines, indicating clustering.
- Patches occurring in 30% region are like binary patches.

Range Patches

- Extend to 7×7 patches

Range Patches

- Extend to 7×7 patches
- Obtain non-trivial β_1 barcodes

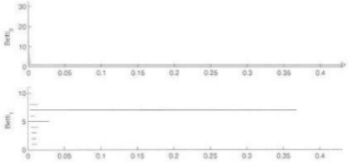

FIGURE 6. Barcodes for $X^7(100, 30)$

Range Patches

- Suggests the possibility of version of primary circle

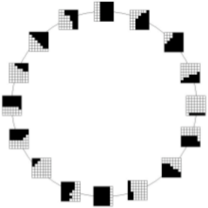

Range Patches

- Suggests the possibility of version of primary circle

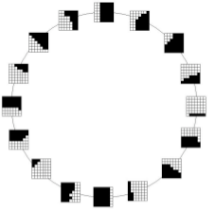

- 3×3 binary patches cannot encode primary circle, but 7×7 ones can

Kleinlet Compression

- This understanding of density can be applied to develop compression schemes

Kleinlet Compression

- This understanding of density can be applied to develop compression schemes
- Earlier work, based on primary circle, called “Wedgelets”, done by Baraniuk, Donoho, et al.

Kleinlet Compression

- This understanding of density can be applied to develop compression schemes
- Earlier work, based on primary circle, called “Wedgelets”, done by Baraniuk, Donoho, et al.
- Extension to Klein bottle dictionary of patches natural

Kleinlet Compression

A Picture is worth 1,000 words

The evidence for Kleinlets over Wedglets

Original

Coded by Kleinlet at .71bpp
PSNR= 29dB

Coded by Wedgelet at .8bpp
PSNR= 27.7dB

Kleinlet Wedgelet

Kleinlet

Wedgelet

Kleinlet Compression

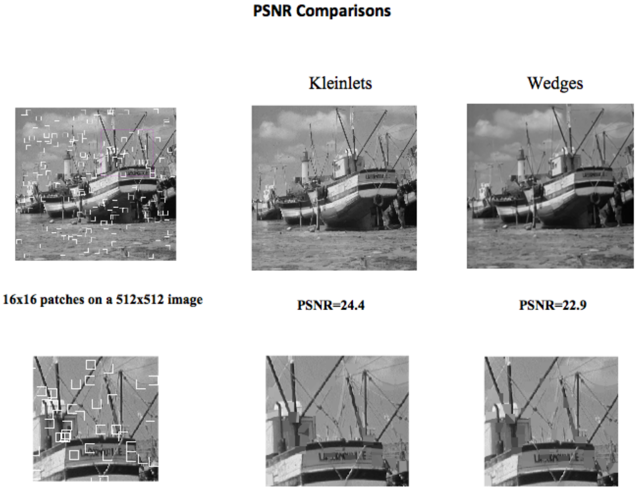

Kleinlet Compression

Compression comparison between kleinlets and wedgelets

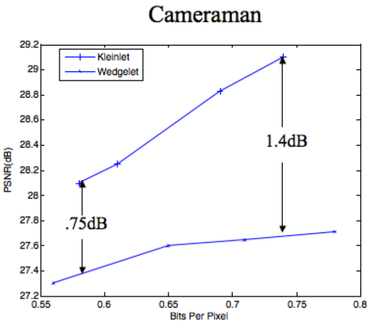

Texture Recognition

- Texture patches can be sampled for high contrast patches

Texture Recognition

- Texture patches can be sampled for high contrast patches
- Yields distribution on Klein bottle

Texture Recognition

- Klein bottle has a natural geometry, and supports its own Fourier Analysis

Texture Recognition

- Klein bottle has a natural geometry, and supports its own Fourier Analysis
- Textures provide distributions on the Klein bottle

Texture Recognition

- Klein bottle has a natural geometry, and supports its own Fourier Analysis
- Textures provide distributions on the Klein bottle
- Pdf's can be given Fourier expansions, gives coordinates for texture patches (Jose Perea)

Texture Recognition

- Klein bottle has a natural geometry, and supports its own Fourier Analysis
- Textures provide distributions on the Klein bottle
- Pdf's can be given Fourier expansions, gives coordinates for texture patches (Jose Perea)
- Gives methods comparable to state of the art in performance, but in which effect of transformations such as rotation is predictable

Texture Recognition

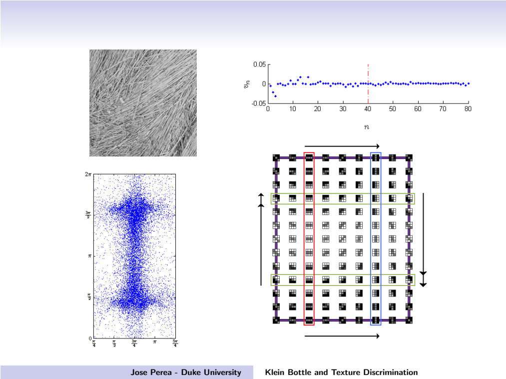

Summary

- Compression and texture recognition often obtained by using finite dictionaries

Summary

- Compression and texture recognition often obtained by using finite dictionaries
- Geometry gives alternate notions of "finiteness", i.e finite geometric descriptions of finite sets

Summary

- Compression and texture recognition often obtained by using finite dictionaries
- Geometry gives alternate notions of "finiteness", i.e finite geometric descriptions of finite sets
- Permits analysis using more mathematics, in particular coordinate changes