The Least Squares Klein Bottle for Image Patches

Bradley Nelson[*1], Gunnar Carlsson[†]

[*]Institute for Computational and Mathematical Engineering
[†]Department of Mathematics
Stanford University

Applied Algebraic Topology
Hokkaido University, Sapporo, Japan
August 11, 2017

[1]Supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program

Outline

What is an Image Patch?

Image from van Hateren Natural Images Database [3]

- A $d \times d$ pixel block of a natural image
- This talk: black & white images
  - Pixel values measure light intensity
- We're interested in "high contrast" image patches

Why Study Image Patches?

Natural images are complex. Patches are (relatively) simple.

A sparse, rough and incomplete picture:

- Image Compression, Harmonic Analysis
  - Ridgelets, Curvelets, etc. [1]
- Image denoising/interpolation [2]
- Vision - understand the visual cortex [3, 4]

Image Patches in TDA

Circles and Klein Bottles

- Carlsson & de Silva (2004) use as an example in witness complex construction. [5]
  - 3 circle model for a dense region of 3×3 patches
- Carlsson et al. (2008) formulate Klein bottle model. [6]
  - Showed that a dense region of 3×3 patches lies near a Klein bottle
- Adams & Carlsson (2009) - range images primary circle [7]

Applications:

- Compression scheme by Maleki, Shahram, Carlsson (2008): BiWedgelets / Kleinlets [8]
- Perea & Carlsson (2014) - rotation-invariant pattern recognition using Klein bottle [9]

Our Contributions

We introduce a general Klein bottle model for image patches

- Generalizes model used in Carlsson et al (2008)
- Works for general patch sizes
- Parametric model for the image patch Klein bottle
- We consider the problem of finding a ‘least squares fit’ Klein bottle for high contrast image patches

Outline

Whence the Klein Bottle?

???

Klein Bottle artist: Ron Estrin

Our Viewpoint

- Image patches are functions (intensity) on a $d \times d$ grid
- Image patches are a discretization of a continuous phenomenon

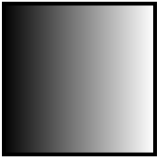

$f(\mathbf{x})$

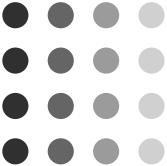

$f(\mathbf{x_i})$

Obtaining Patches

1. Randomly select 5000 patches from image
2. Take log intensities of patches
3. mean-center patches
4. take top 20% by contrast norm
5. normalize by contrast norm

Procedure from Lee, Pedersen, Mumford (2003) [4] applied to van Hateren dataset [3]

Edges

Our model for "high-contrast" image patches is based on the observation that patches that look like edges are common.

Example image patches from the van Hateren data set [3]

Sources of Variation

Two sources of variation in edges:

- Orientation of edges (primary circle)
- Type of edge (secondary circles)

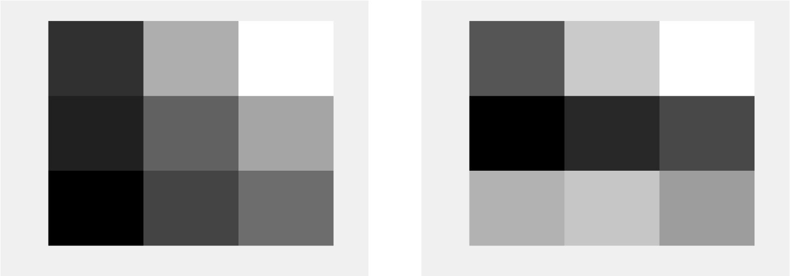

Example image patches from the van Hateren data set [3]

Odd Functions

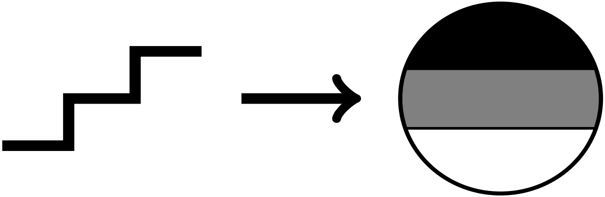

Odd functions capture the behavior of transitions from one intensity to another

Even Functions

Even functions capture the behavior of lines

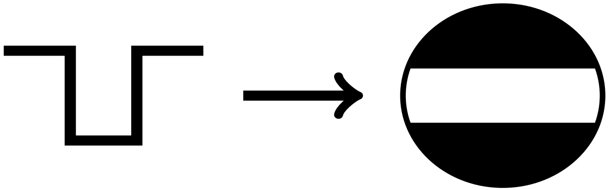

Combined Functions

Edges can also be a combination of even and odd functions

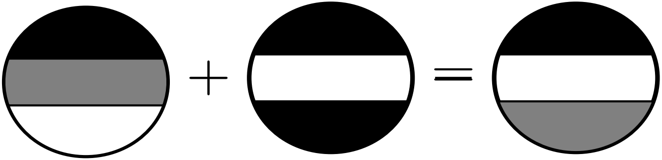

Outline

Klein Bottle Elements

We'll show how to construct Klein bottles given the following ingredients:

- A symmetric domain $\mathbf{X}$
- A space of functions $\mathcal{F}(\mathbf{X})$ mapping $\mathbf{X}$ to $\mathbb{R}$
- A metric on $\mathcal{F}(\mathbf{X})$
- Circles of even and odd functions in $\mathcal{F}(\mathbf{X})$
- A "phase" stitching together the circles

Symmetric Domain

We say a space $\mathbf{X}$ is symmetric if for all $x \in \mathbf{X}$, $-x \in \mathbf{X}$.

We'll generally think of compact subsets of $\mathbb{R}^n$.

Examples:

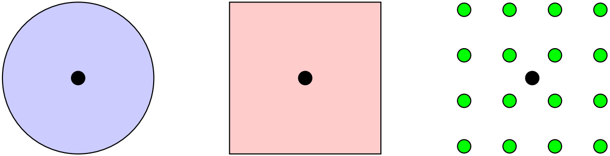

Space of Functions

Ingredient 2: A space of functions $\mathcal{F}(\mathbf{X})$ mapping $\mathbf{X}$ to $\mathbb{R}$:
Examples:

- L_2 functions on $\mathbf{X}$
- Continuous bounded functions on $\mathbf{X}$
- degree ℓ polynomials on $\mathbf{X}$

A metric

$\mathcal{F}(\mathbf{X})$ may come equipped with a metric, but sometimes we may be able to choose one.

In general we will consider the L_2 metric. For continuous $\mathbf{X}$:

$$d(f,g) = \left(\int_{\mathbf{X}} (f(x) - g(x))^2 \right)^{1/2}$$

For discrete $\mathbf{X}$:

$$d(f,g) = \left(\sum_{\mathbf{X}} (f(x) - g(x))^2 \right)^{1/2}$$

Even and Odd

Let $f \in \mathcal{F}(\mathbf{X})$.

We say f is odd if $f(-x) = -f(x)$ for all $x \in \mathbf{X}$.

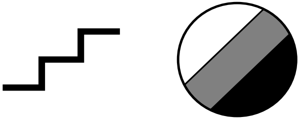

We say f is even if $f(-x) = f(x)$ for all $x \in \mathbf{X}$.

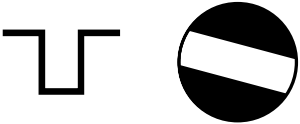

Circle of Odd Functions

Let $\mathcal{F}_o \subset \mathcal{F}(\mathbf{X})$ be a set of odd functions homeomorphic to S^1. We can think of $\mathcal{F}_o$ as being parameterized by $\phi \in [0, 2\pi]$, and denote $f_o \in \mathcal{F}_o$ as $f_o(x; \phi)$. Additionally, we require that $f_o(x; \phi) = -f_o(x; \phi + \pi)$.

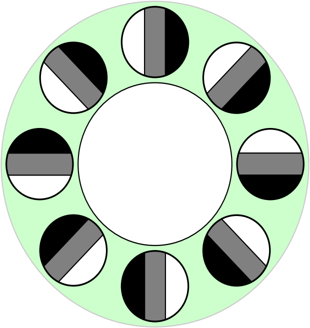

(This is the "Primary Circle" [5])

Circle of Even Functions

Let $\mathcal{F}_e \subset \mathcal{F}(\mathbf{X})$ be a set of even functions homeomorphic to S^1. We can think of $\mathcal{F}_e$ as being parameterized by $\phi \in [0, 2\pi]$, and denote $f_e \in \mathcal{F}_o$ as $f_e(x;\phi)$. Additionally, we require $f_e(x;\phi_1) \neq -f_e(x;\phi_2)$, $\phi_1, \phi_2 \in [0, 2\pi]$

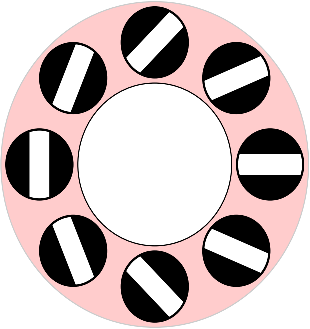

Note that one rotation in this circle is half a rotation in angle.

Mixing Even and Odd

We introduce a "phase" θ to combine even and odd functions. Let $\mathcal{F}_k \subset \mathcal{F}(\mathbf{X})$ be defined as

$$\mathcal{F}_k = \{\cos(\theta) f_e(x; 2\phi) + \sin(\theta) f_o(x; \phi) \mid \theta \in [0, 2\pi], \phi \in [0, \pi]\}$$

We will denote an element of $\mathcal{F}_k$ as $f_k(x; \theta, \phi)$.

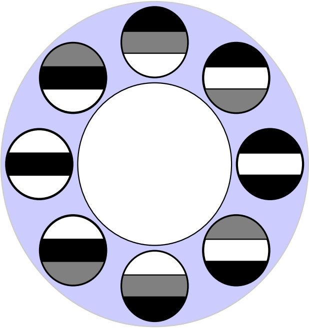

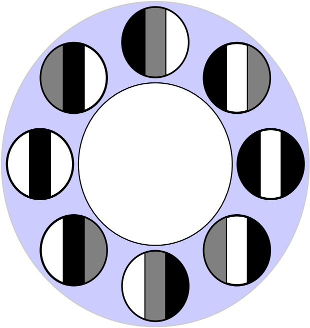

(These are "Secondary Circles" [5])

Identification

$\mathcal{F}_k$ is a Klein bottle, as seen through an identification on the torus.
Recall: $f_o(x;\phi) = -f_o(x;\phi+\pi)$

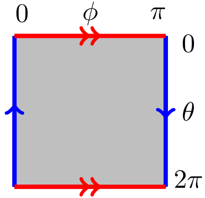

$$
\begin{aligned}
&f_k(x;2\pi-\theta,\phi+\pi)\\
&= \cos(-\theta)f_e(x;2\phi+2\pi)+\sin(-\theta)f_o(x;\phi+\pi)\\
&= \cos(\theta)f_e(x;2\phi+2\pi)-\sin(\theta)f_o(x;\phi+\pi)\\
&= \cos(\theta)f_e(x;2\phi)+\sin(\theta)f_o(x;\phi)\\
&= f_k(x;\theta,\phi)
\end{aligned}
$$

Red: $f_k(x;\theta=0,\phi) = f_k(x;\theta=2\pi,\phi)$ (full phase)
Blue: $f_k(x;\theta,\phi=0) = f_k(x;2\pi-\theta,\phi=\pi)$ (even/odd, sin/cos)

Additional identifications do not occur due to the condition
$f_e(x;\phi_1) \neq -f_e(x;\phi_2)$, for any $\phi_1,\phi_2 \in [0,2\pi]$.

Klein Bottle Identification

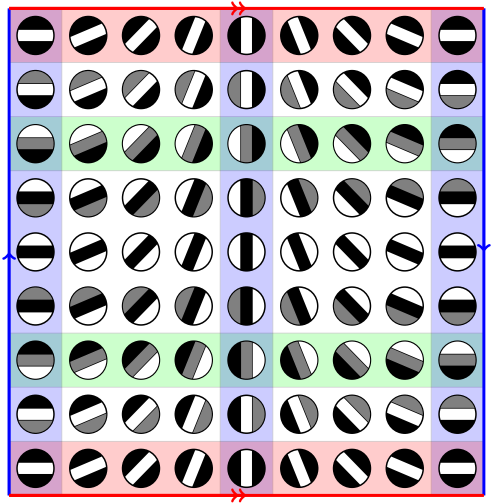

Green: primary circle, Blue: secondary circles

Comments

- We've made few assumptions about $\mathbf{X}$
  - For patches, no assumptions about size of pixel grid [2]
  - Can consider more general domains
- We've said little about $\mathcal{F}_e$ and $\mathcal{F}_o$
  - Don't need to look anything like edges
  - Also, many possible candidates for edge-like functions

[2] We need to be able to satisfy conditions on $\mathcal{F}_e$ and $\mathcal{F}_o$, which rules out 2×2 grids and smaller if normalizing patches

Outline

Linear Combinations

We can construct circles $\mathcal{F}_e$ and $\mathcal{F}_o$ explicitly using linear combinations of several even and odd functions.
For example, let f_{o1}, f_{o2} be odd functions and $f_{e1}, f_{e2}, f_{e3}, f_{e4}$ be even functions.

$$\mathcal{F}_o = \{\sin(\phi) f_{o1} + \cos(\phi) f_{o2} \mid \phi \in [0, 2\pi]\}$$

satisfies $f_o(x; \phi + \pi) = -f_o(x; \phi)$.

$$\mathcal{F}_e = \left\{ \begin{cases} \sin(\phi) f_{e3} + \cos(\phi) f_{e1} & 0 \leq \phi < \pi/2 \\ \sin(\phi) f_{e3} - \cos(\phi) f_{e2} & \pi/2 \leq \phi < \pi \\ -\sin(\phi) f_{e4} - \cos(\phi) f_{e2} & \pi \leq \phi < 3\pi/2 \\ -\sin(\phi) f_{e4} + \cos(\phi) f_{e1} & 3\pi/2 \leq \phi < 2\pi \end{cases} \quad \phi \in [0, 2\pi] \right\}$$

satisfies $f_e(x; \phi_1) \neq -f_e(x; \phi_2)$, for all $\phi_1, \phi_2 \in [0, 2\pi]$

A DCT Klein Bottle

We use the lowest contrast modes of the discrete cosine transform as our even and odd functions.

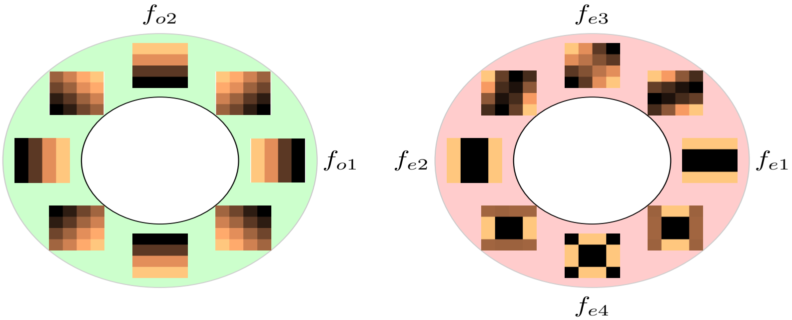

A DCT Klein Bottle

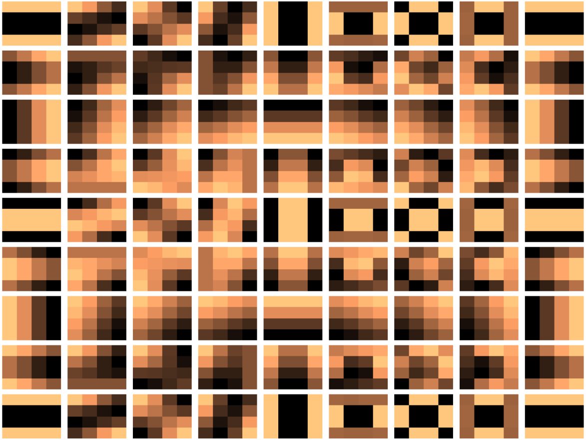

Klein bottle generated using DCT basis functions on 4×4 patches.

Polynomial Patches

We can also turn a 1-dimensional function into a 2-dimensional function via projection before evaluation.

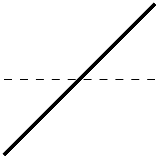

$f_o(x)$

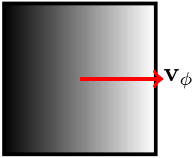

$f_o(\mathbf{x};\phi) = f_o(\mathbf{x}^T\mathbf{v}_\phi)$

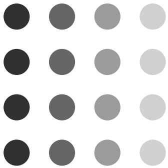

$f_o(\mathbf{x_i};\phi)$

$$\mathbf{v}_\phi = \begin{bmatrix} \cos(\phi) \\ \sin(\phi) \end{bmatrix}$$

Original Klein Bottle Model

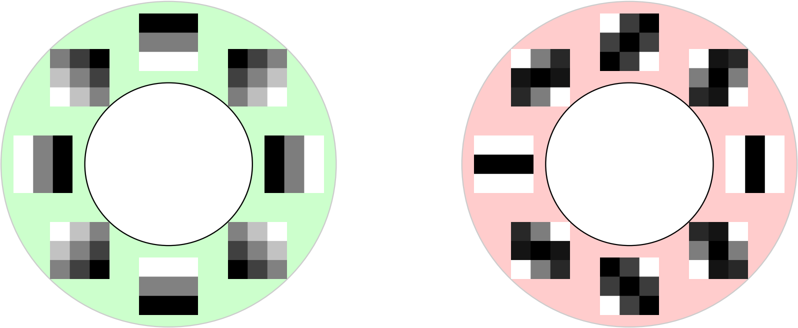

The original image patch Klein bottle of Carlsson et al. [6]

$$f_o(x) = x \qquad f_e(x) = x^2$$

Original Klein Bottle

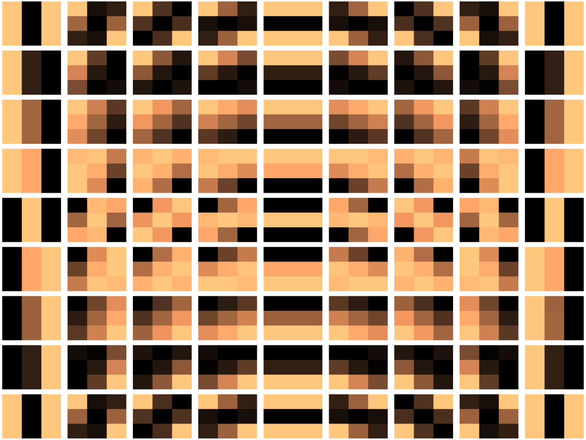

The original image patch Klein bottle of Carlsson et al.

$$f_o(x) = x \qquad f_e(x) = x^2$$

Outline

Least Squares Formulation

Consider a family of Klein bottle models parameterized by γ, with Klein bottle K_γ, and projection operation

$$P_\gamma y = \arg \inf_{f_k \in K_\gamma} \|y - f_k\|$$

We would like to find the parameter γ that solves

$$\underset{\gamma}{\text{minimize}} \sum_i \|y^i - P_\gamma y^i\|^2$$

For a collection of patches $\{y^i\}$.

We will consider the case where $\mathcal{F}_e$ and $\mathcal{F}_o$ are determined by low degree polynomials ($\gamma = (\alpha, \beta)$):

$$f_o(\mathbf{x}, \phi) = \Big(\sum_{j=0}^{\ell} \alpha_j (\mathbf{x}^T \mathbf{v}_\phi)^{2j+1} \Big)$$

$$f_e(\mathbf{x}, \phi) = \Big(\sum_{j=1}^{\ell} \beta_j (\mathbf{x}^T \mathbf{v}_\phi)^{2j} \Big)$$

Decomposing a Patch

To determine the best choice of $\mathcal{F}_o(\alpha)$ and $\mathcal{F}_e(\beta)$ for the Klein bottle, we want to determine the orientation ϕ and phase θ independently. For each patch $y^i \in \{y^i\}$

- ϕ^i is estimated using the Harris Edge detector [10]. This gives us the coordinates $(x_j^i)^T \mathbf{v}_\phi^i$ to use for function evaluation
- We decompose the patch into odd and even parts

$$y_e^i(x) = \big(y^i(x) + y^i(-x)\big)/2$$
$$y_o^i(x) = \big(y^i(x) - y^i(-x)\big)/2$$

- θ^i is estimated using the norms of y_e^i and y_o^i

Least Squares System

We can set up a linear least squares equation for the odd term coefficients α,

$$X\alpha = y$$

where α is the coefficient vector, and X and y have vertical blocks of the form

$$X_i = \sin(\theta^i) \begin{bmatrix} (x_1^i)^T \mathbf{v}_\phi^i & ((x_1^i)^T \mathbf{v}_\phi^i)^3 & \dots & ((x_1^i)^T \mathbf{v}_\phi^i)^{2\ell+1} \\ \vdots & \vdots & & \vdots \\ (x_{d^2}^i)^T \mathbf{v}_\phi^i & ((x_{d^2}^i)^T \mathbf{v}_\phi^i)^3 & \dots & ((x_{d^2}^i)^T \mathbf{v}_\phi^i)^{2\ell+1} \end{bmatrix}$$

$$y_i = \begin{bmatrix} y_{o,1}^i \\ \vdots \\ y_{o,d^2}^i \end{bmatrix}$$

A similar system can be derived for the even coefficients β.

Empirical Results


Example images from the Van Hateren dataset [3]

We use distance to a patch's kth nearest neighbor as a proxy for density.
We'll look at the top p percent "densest" points for different values of p.

Empirical Results

Van Hateren Data set. Numbers reported are the relative 2-norm approximation error.
$k = 50, p = 0.2, \ell = 9$

d	PCA2	PCA3	PCA4	KB(DCT)	KB(old)	KB(LS)
3	0.2446	0.1897	0.1328	0.2630	0.1268	**0.1172**
4	0.3434	0.2792	0.2132	0.3718	0.2240	**0.1997**
5	0.4402	0.3558	0.2730	0.4690	0.3014	**0.2700**
6	0.5185	0.4215	**0.3290**	0.5630	0.3940	0.3524

Empirical Results

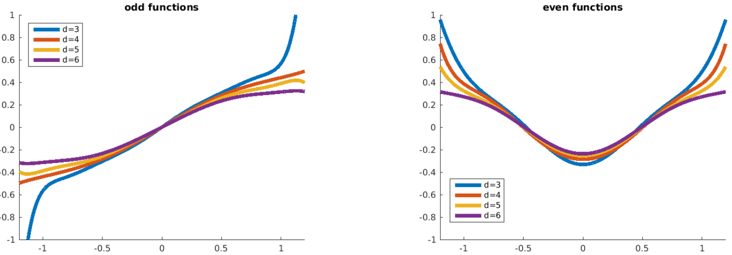

Least Squares Polynomials for $k = 50, p = 0.2, \ell = 9$.

Empirical Results

$k = 50, p = 0.2, d = 5$. Witness Complex (70 landmarks) computed on data augmented with least squares Klein bottle (100×100 sample on (θ, ϕ) grid).

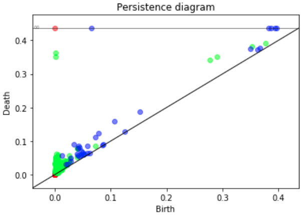

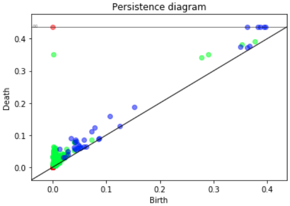

Left: persistence diagram computed over $\mathbb{Z}_2$, right, persistence diagram computed over $\mathbb{Z}_3$. Computations performed with Gudhi.

Empirical Results

Van Hateren Data set. Numbers reported are the relative 2-norm approximation error.
$k = 50, p = 0.5, \ell = 9$

d	PCA2	PCA3	PCA4	KB(DCT)	KB(old)	KB(LS)
3	0.3128	0.2550	0.1924	0.3159	0.1976	**0.1895**
4	0.4452	0.3646	**0.2927**	0.4419	0.3170	**0.2944**
5	0.5250	0.4368	**0.3613**	0.5409	0.4285	0.3895
6	0.5875	0.4993	**0.4223**	0.6370	0.5378	0.4877

Empirical Results

Van Hateren Data set. Numbers reported are the relative 2-norm approximation error.
$k = 50, p = 1.0, \ell = 9$ (all high-contrast patches)

d	PCA2	PCA3	PCA4	KB(DCT)	KB(old)	KB(LS)
3	0.4625	0.3706	**0.2842**	0.4175	0.3415	**0.3331**
4	0.5486	0.4674	**0.3838**	0.5406	0.4838	0.4610
5	0.6045	0.5278	**0.4497**	0.6372	0.5964	0.5636
6	0.6472	0.5738	**0.5009**	0.7240	0.6952	0.6577

Discussion

- The least squares Klein bottle model does well on the high density portions of the patch data
- As dimension increases, performance suffers as edges are not the only thing that appears
- The least squares Klein bottle does better than 2-D PCA

Outline

Conclusion

We've shown

- There are many ways to come up with Klein bottles of functions on symmetric spaces
- Edge models for image patches are a natural way to arrive at the Klein bottle
  - This doesn't depend on patch size
- Klein bottle models empirically work well for high contrast patch compression

Ongoing/Future Work

- Different (better?) bases for even and odd functions
- Bases for individual images/textures
  - Functions learned for textures may be interesting
- Improvements in optimization process/understanding
  - Currently, obtaining ϕ is somewhat heuristic
  - Statistical guarantees?

Questions?

(Some starters)

- Where else might elements of these Klein bottle models appear?
- What image patch features aren't captured by the edge model?
- Are there reasonable models that can be built on top of the Klein bottle?

References I

- D.L. Donoho, A.G. Flesia, Can recent innovations in harmonic analysis explain key findings in natural image statistics?. Network: computation in neural systems, 12(3), pp.371-393. 2001.
- S. Osher, Z. Shi, and W. Zhu. Low Dimensional Manifold Model for Image Processing. Technical Report UCLA CAM Report 16-04, 2016.
- J.H. van Hateren and A. van der Schaaf. Independent component filters of natural images compared with simple cells in primary visual cortex. Proc. R. Soc. Lond., B265:359366, 1998.
- A.B. Lee, K.S. Pedersen, and D.B. Mumford. The nonlinear statistics of high-contrast patches in natural images. International Journal of Computer Vision, 54:83103, 2003.
- V. de Silva, G. Carlsson. Topological estimation using witness complexes. Eurographics Symposium on Point-Based Graphics. 2004.
- G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. International Journal of Computer Vision, 76(1):112, 2008.

References II

- H. Adams, G. Carlsson. On the nonlinear statistics of range image patches. SIAM J. Imaging Sciences, 2(1):110117, 2009.
- A. Maleki, M. Shahram, G. Carlsson A Near Optimal Coder For Image Geometry With Adaptive Partitioning. 15th IEEE International Conference on Image Processing, 2008.
- J. Perea, G. Carlsson. A klein-bottle-based dictionary for texture representation. International Journal of Computer Vision, 107(1):7597, 2014.
- C. Harris, M. Stephens. A combined corner and edge detector. 1988.

Texture Recognition

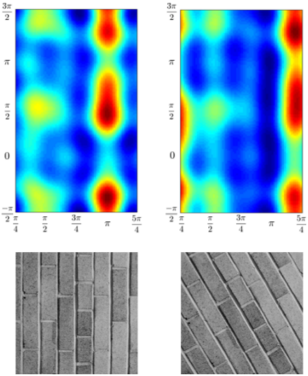

J. Perea, G. Carlsson "A Klein-Bottle-Based Dictionary for Texture Representation" (2014)

Compression

A. Maleki, M. Shahram, G. Carlsson "A Near Optimal Coder For Image Geometry With Adaptive Partitioning" (2008)