Persistence Diagrams of Cortical Surface Data

Moo K. Chung[1,2], Peter Bubenik[4], Peter T. Kim[5]
Kim M. Dalton[2], Richard J. Davidson[2,3]

[1]Department of Biostatistics and Medical Informatics
[2]Waisman Laboratory for Brain Imaging and Behavior
University of Wisconsin, Madison, WI 53706, USA
[3]Department of Psychology and Psychiatry
[4]Department of Mathematics
Cleveland State University, Cleveland, Ohio 44115-2214, USA
[5] Department of Mathematics and Statistics
University of Guelph, Guelph, Ontario N1G 2W1, Canada

mkchung@wisc.edu

Abstract. We present a novel framework for characterizing signals in images using techniques from computational algebraic topology. This technique is general enough for dealing with noisy multivariate data including geometric noise. The main tool is persistent homology which can be encoded in persistence diagrams. These are scatter plots of paired local critical values of the signal. One of these diagrams visually shows how the number of connected components of the sublevel sets of the signal changes. The use of local critical values of a function differs from the usual statistical parametric mapping framework, which mainly uses the mean signal in quantifying imaging data. Our proposed method uses all the local critical values in characterizing the signal and by doing so offers a completely new data reduction and analysis framework for quantifying the signal. As an illustration, we apply this method to a 1D simulated signal and 2D cortical thickness data.

1 Introduction

In neuroimaging, it is usually assumed that measurements f in images follow the familiar signal plus noise framework

$$f(x) = \mu(x) + \epsilon(x),\ x \in \mathbb{M} \subset \mathbb{R}^d, \tag{1}$$

where μ is the unknown mean signal, to be estimated, and ϵ is noise [5] [15] [19] [20] [26] [37]. The unknown signal is usually estimated by various spatial image smoothing techniques over $\mathbb{M}$. The most widely used smoothing technique is kernel smoothing and its variants because of their simplicity, and because they provide the theoretical context for scale spaces and Gaussian random field theory [33] [37].

In the usual statistical parametric mapping framework [15] [20] [37], inference on the model (1) proceeds as follows. If we denote an estimate of the signal by

$\widehat{\mu}$, the residual $f - \widehat{\mu}$ gives an estimate of the noise. One then constructs a test statistic $T(x)$, corresponding to a given hypothesis about the signal. As a way to account for spatial correlation of the statistic $T(x)$, the global maximum of the test statistic over the search space $\mathbb{M}$ is taken as the subsequent test statistic. Hence a great deal of the neuroimaging and statistical literature, have been devoted to determining the distribution of $\sup_{x \in \mathbb{M}} T(x)$ using random field theory [35] [37], permutation tests [31] and the Hotelling–Weyl volume of tubes calculation [30].

The use of the mean signal is one way of performing data reduction, however, this may not necessarily be the best way to characterize complex multivariate imaging data. Thus instead of using the mean signal, in this paper we propose to use what is known as persistent homology, which pairs the local critical values [12] [13] [40]. It is intuitive that the local critical values of $\widehat{\mu}$ approximately characterizes the shape of the continuous signal μ using only a finite number of scalar values. By pairing these local critical values in a nonlinear fashion and plotting them, one constructs the persistence diagram [7] [12] [29] [39].

Although persistent homology is popular in computational algebraic topology with applications in protein structure analysis [32], gene expression [11], and sensor networks [9], as far as the authors are aware, there is no such application in medical image analysis. This is the first paper that applies the concept of persistent homology to medical imaging data. The proposed method is illustrated using both simulated and real neuroimaging data. For the simulation, we use 1D Gaussian noise in (1). The 2D neuroimaging data comes from an MRI autism study where the interest is in quantifying the abnormal cortical thickness pattern in autistic subjects if there is any.

2 Persistence Diagrams

A function is called a Morse function if all critical values are unique and non-degenerate, i.e. the Hessian does not vanish [28]. We note that for integer valued digital images, critical values of intensity may not be all unique; however, the underlying continuous signal μ in (1) is likely and assumed to be a Morse function. We estimate the signal using a kernel function and obtain a smooth estimate.

For illustrative purposes, we will show how to construct the persistence diagram for a 1D Morse function. Assuming μ is a Morse function with a finite number of critical values, define a sublevel set $R(y) = \mu^{-1}(-\infty, y]$. The sublevel set is the subset of $\mathbb{R}$ that satisfies $\mu(x) \leq y$. The sublevel set can have many disjoint components. Let $\#R(y)$ be the number of connected components in the sublevel set. Let us denote the local minimums as $g_1, \cdots, g_m$ and the local maximums as $h_1, \cdots, h_n$. Since the critical values of the Morse function are all unique, we can strictly order the local minimums from the smallest to the largest as

$$g_{(1)} < g_{(2)} < \cdots < g_{(m)}$$

and similarly for the local maximums as

$$h_{(1)} < h_{(2)} < \cdots < h_{(n)}.$$

We further collect all the critical values,

$$z_1 = g_1, \ldots, z_m = g_m, z_{m+1} = h_1, \ldots, z_{m+n} = h_n$$

and order them as

$$z_{(1)} < z_{(2)} < \cdots < z_{(m+n)}.$$

At each minimum, we have the birth of a new component, i.e.

$$\#R(g_i) = \#R(g_i - \varepsilon) + 1$$

for sufficiently small ε. The new component is identified with the local minimum g_i. Similarly at each maximum, we have the death of a component, i.e.

$$\#R(h_i) = \#R(h_i - \varepsilon) - 1,$$

and two components will merge as one. The number of connected components will only change if we pass through critical points and we can iteratively compute $\#R$ at each critical value as

$$\#R(z_{(i+1)}) = \#R(z_{(i)}) \pm 1.$$

The sign depends on whether $z_{(i+1)}$ is a maximum (-1) or a minimum $(+1)$. This is the basis of Morse theory [28] that says the topological characteristics of a topological spaces are characterized by the local behavior at critical points of a Morse function on that space. Persistent homology produces pairs (g_i, h_j) of critical values so that a component is born at g_i and dies at h_j. Of course these are the (topological) parameters of interest which are unknown and to be statistically estimated with data generated according to (1).

As an example, the birth and death processes are illustrated in Figure 1, where the gray dots are simulated with Gaussian noise with mean 0 and variance 0.2^2 as

$$f(x) = \mu(x) + N(0, 0.2^2)$$

with signal $\mu(t) = 10(t - 1/2)^2 + \cos(7\pi t)/2$. The signal μ is estimated using heat kernel smoothing [5] and plotted as the red line. Now we increase y from $-\infty$ to ∞. When we hit the first critical value $y = a$, the sublevel set consists of a single point, i.e. $\widehat{\#R(a)} = 1$. When we hit the minimum at $y = b$, we have the birth of a new component at b, i.e. $\widehat{\#R(b)} = 2$. When we hit the maximum at $y = c$, the two components identified by a and c are merged together to form a single component, i.e. $\widehat{\#R(c)} = 1$.

When we pass through a maximum and merge two components, we pair the maximum with the higher of the two minimums of the two components [12]. Doing so we are pairing the birth of a component to its death. Obviously the paired extremes do not have to be adjacent to each other. If there is a boundary, the function value evaluated at the boundary is treated as a critical value. In our simulated example, we need to pair (b, c) and (d, e). Other critical values are paired similarly. The reduced persistence diagram is then the scatter plot of

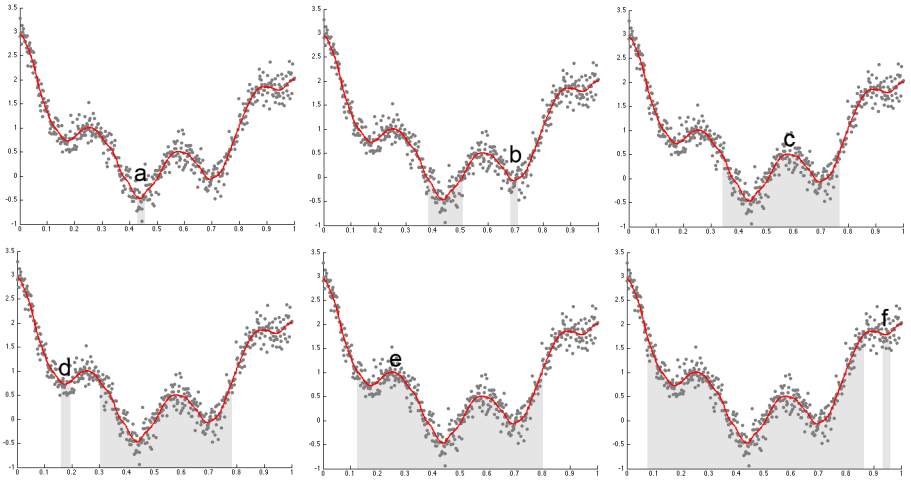

Fig. 1. The births and deaths of components in sublevel sets. We have critical values a, b, c, d, e, f, where $a < b < d < f$ are minimums and $c < e$ are maximums. At $y = a$, we have a single component marked by a single gray area. When we increase the level to $y = b$, we have the birth of a new component in addition to the existing component born at a. At the maximum $y = c$, the two components merge together to form a single component. Following the pairing rule [12], we pair (b, c) and (d, e). Other critical values are paired similarly.

these pairings. For technical reasons, the persistence diagram also include all of the points (t, t), where $t \in \mathbb{R}$.

NOTATION: PAIRING IS DENOTED BY (t, t), (a, b), (x, y) IN VARIOUS PLACES. CAN WE HAVE UNIFORM SYMBOLS TO DENOTE PAIRING.

2.1 Persistence Diagram for Cortical Data

For a 2D Morse function defined on a cortical manifold $\mathbb{M} \subset \mathbb{R}^3$, we need to also consider saddle points so the situation is more complicated. At a saddle point, we can have two possible pairings corresponding to either birth or death. A saddle point may join two components. This case is analogous to the local maximum in the 1D case. In this case, persistent homology pairs the value of the saddle point with the larger of the minimums of the two components. This pair is recorded in the persistence diagram of degree 0 (Figure 4). If the saddle point does not join to disconnected components, then a hole is born in the sublevel sets. Persistent homology pairs the value at this saddle point with the value of the local maximum where this hole disappears. A precise definition is given in Section 4.

DEFINITION OF DEGREE-1 PD IS MISSING HERE.

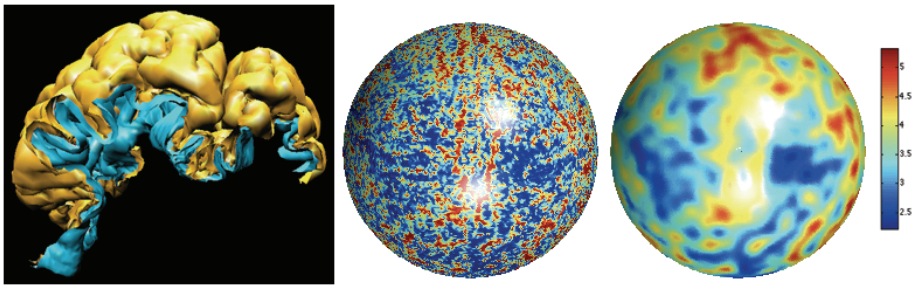

Fig. 2. Cortical thickness is computed as the distance between the outer (yellow) and the inner cortical (blue) surfaces. The cortical thickness is mapped onto a unit sphere and goes through heat kernel smoothing [5]

Among various cortical measures, in this paper we consider cortical thickness, which has been used in characterizing various clinical populations [6] [14] [23] [24] [27] [38]. High resolution magnetic resonance images of age-matched right-handed males (6 high functioning autistic and 11 normal controls) were obtained using a 3-Tesla GE SIGNA scanner. The collected images went through intensity nonuniformity correction [34] and were spatially normalized into the MNI stereotaxic space *via* a global affine transformation [8]. Subsequently a supervised neural network classifier was used for tissue segmentation [22]. Brain substructures such as the brain stem were removed to make both the outer and the inner surfaces to be topologically equivalent to a sphere. A deformable surface algorithm [25] was used to obtain the inner cortical surface by deforming from a spherical mesh (Figure 2). Then the outer surface was obtained by deforming the inner surface. The deformation process establishes the structural correspondence between the two surfaces. The cortical thickness f is then defined as the distance between the corresponding vertices along the cortical mesh $\mathbb{M}$. **DON'T FIX ANYTHING HERE. WE USE PAST TENSE FOR DESCRIBING EXPERIMENTS, IMAGE PROCESSING DONE TO THE DATA WE ARE USING IN THIS PAPER.**

Since the deformable surface algorithm starts with a spherical mesh, there is no need to use other available surface flattening algorithms [1] [2] [16] [17] [36] for mapping thickness to the unit 2–sphere S^2. Let $\zeta : \mathbb{M} \to S^2$ be a sufficiently smooth surface flattening obtained from the deformable surface algorithm. Then the pullback $(\zeta^{-1})^* \widehat{\mu} = \widehat{\mu} \circ \zeta^{-1}$ projects the cortical thickness from the cortical surface $\mathbb{M}$ to the unit sphere. Figure 2 shows the pull back and the corresponding heat kernel smoothing on S^2. Note that in the process of flattening, the critical values do not change so the persistence diagram should be identical for $\widehat{\mu}$ and its pullback $(\zeta^{-1})^* \widehat{\mu}$. Therefore, we will construct the persistence diagram on the unit 2–sphere by projecting the cortical data to the sphere.

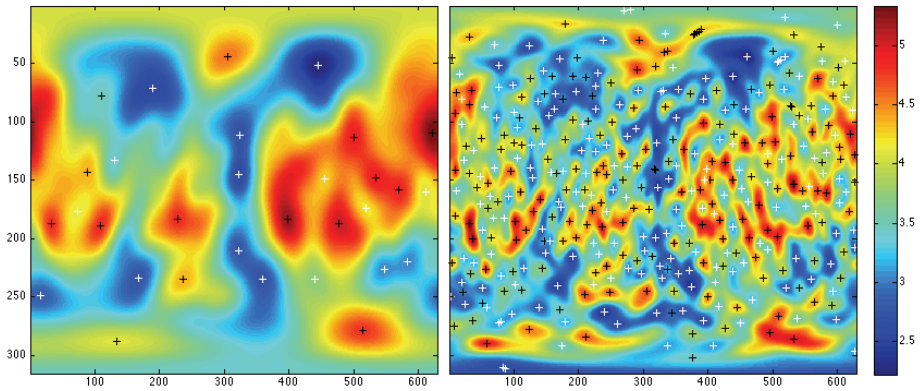

Fig. 3. The flat maps of cortical thickness at different smoothing scales. The maximums and minimums are denoted with black and white crosses respectively. The smoothing is done along the unit sphere and flattened using the Eulger angles θ (vertical axis) and φ (horizontal axis) associated with the 2-sphere. Smoother thickness produces less number of critical points and, in turn, less number of pairings.

3 Kernel Smoothing

As described in Section 2.1, after the application of a deformable surface algorithm, our data is on the unit 2– sphere, S^2. So our measurement, $f : S^2 \to \mathbb{R}$ is given by the nonparametric regression formula (1), where μ is the unknown signal and ϵ is the noise. In this section, we estimate the persistent homology of the sublevel sets of $\widehat{\mu}$, an estimator of μ.

We begin by smoothing the data using the kernel,

$$K_{x_0}(x) = \max(1 - \kappa \arccos(x_0' x), 0),$$

where κ is given in [21] and $\arccos(x'y)$ gives the geodesic distance between x and y on the unit sphere. We smooth the data using the usual kernel function estimator

$$\widehat{\mu}(x) = \frac{\sum_i f(x_i) K_{x_i}(x)}{\sum_i K_{x_i}(x)}. \tag{2}$$

To implement this we need to choose the corresponding design points which we do in the following way. We start by choosing a triangulation, $\mathcal{T}$, of the sphere whose number of vertices satisfies the conditions in [3]. For our data, we start with an icosahedron and iteratively subdivide it three times, obtaining a triangulation with 1280 faces and 642 vertices.

For a sample of size n, define the estimator $\widehat{\mu}_n$ in the following way. For each vertex v in our triangulation, we define $\widehat{\mu}_n(v) = \widehat{\mu}(v)$ according to (2). For each face in our triangulation, we define $\widehat{\mu}_n$ on the face by affine interpolation from the values on the vertices. This construction is well defined on the edges, and defines a function on the sphere.

3.1 The persistence diagrams of $\widehat{\mu}_n$

It remains to calculate the persistence diagrams of the sublevel sets of $\widehat{\mu}_n$. We will see that because of the way $\widehat{\mu}_n$ is constructed, we can calculate its persistence diagrams using our triangulation, $\mathcal{T}$.

We filter $\mathcal{T}$ using $\widehat{\mu}_n$ as follows. Let $r_1 \leq r_2 \leq \ldots \leq r_m$ be the ordered list of values of $\widehat{\mu}_n$ on the vertices of the triangulation. For $1 \leq i \leq m$, let $\mathcal{T}_i$ be the subcomplex of $\mathcal{T}$ containing all vertices v with $\widehat{\mu}_n(v) \leq r_i$ and all edges whose boundaries are in $\mathcal{T}_i$ and all faces whose boundaries are in $\mathcal{T}_i$. This construction is called a Vietoris–Rips complex. We obtain the following filtration of $\mathcal{T}$,

$$\phi = \mathcal{T}_0 \subset \mathcal{T}_1 \subset T_2 \subset \cdots \subset \mathcal{T}_m = \mathcal{T}.$$

The end result is that the topological properties of the sublevel sets of $\widehat{\mu}_n$ will equal the topological properties of the above filtration of $\mathcal{T}$.

Using the software Plex, [10], we calculate the persistent homology, in degrees 0, 1 and 2 of the triangulation $\mathcal{T}$ filtered according to the estimator for each of the 27 subjects. Since the data is two–dimensional, we do not expect any interesting homology in higher degrees. In degree two, the persistent homology consists of a single persistence pair (a, ∞), where a is the maximum of $\widehat{\mu}_n$.

To compare the autistic subjects and control subjects, we take the union of the persistence diagrams of the subjects (Figure 4).

4 Statistical Properties of Persistence Diagram

In this section we will make more precise definition of a persistence diagram [7] and present results that compare the topological parameters and their estimators [3] [4].

The persistent homology of the signal, μ, is encoded in its reduced persistence diagram, $\bar{D}(\mu)$, which is a multiset of points each corresponding to the persistence of one topological feature, as in the examples above. In order to define a metric for such diagrams, it is convenient to add the ordered pairs (a, a) for all $a \in \mathbb{R}$, each with infinite multiplicity. Call this multiset the *persistence diagram* of μ, denoted $D(\mu)$. We now give the precise definition.

Let k be a nonnegative integer. Given $\mu : S^2 \to \mathbb{R}$ and $a \leq b \in \mathbb{R}$ the inclusion of sublevel sets $i_a^b : S^2_{\mu \leq a} \hookrightarrow S^2_{\mu \leq b}$ induces a map on homology

$$H_k(i_a^b) : H_k(S^2_{f \leq a}) \to H_k(S^2_{f \leq b}).$$

The image of $H_k(i_a^b)$ is the persistent homology group from a to b. Let β_a^b be its dimension. This counts the independent homology classes which are born by time a and die after time b.

Call a real number a a *homological critical value* of μ if for all sufficiently small $\varepsilon > 0$ the map $H_k(i_{a-\varepsilon}^{a+\varepsilon})$ is not an isomorphism. Call μ *tame* if it has finitely many homological critical values, and for each $a \in \mathbb{R}$, $H_k(S^2_{\mu \leq a})$ is finite dimensional. In particular, any Morse function on a compact manifold is tame.

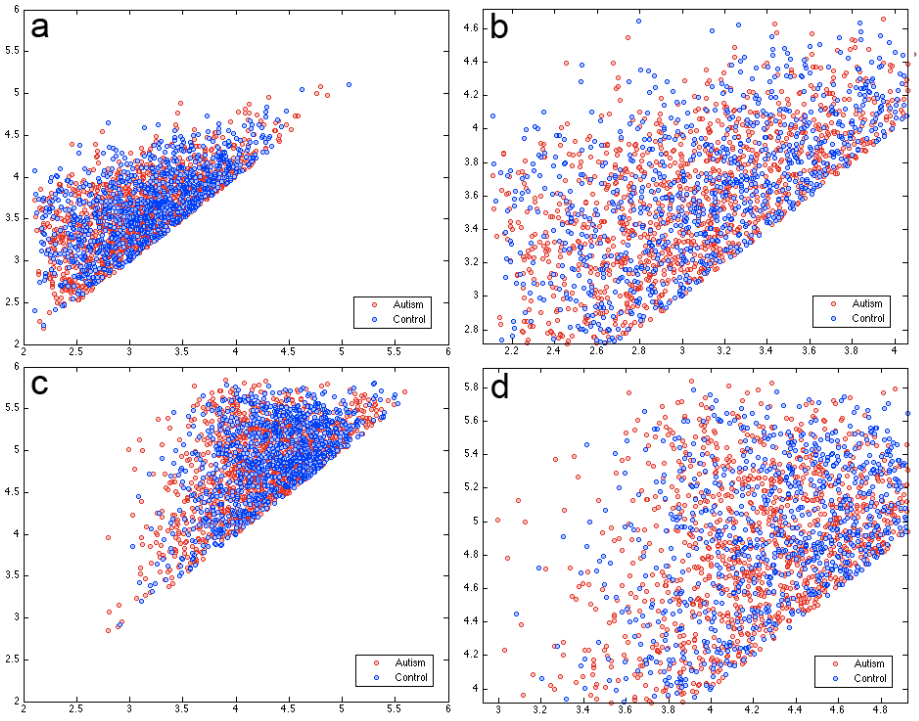

Fig. 4. The persistence diagrams for 11 control (blue) and 16 autistic (red) subjects. One notices an additional layer of structure in the autistic group in both persistence diagrams. The figure clearly demonstrates the feasibility of using the persistence diagram for discriminating populations. (a) (b) 0-degree persistence diagram (c) (d) 1-degree persistence diagram.

Assume that μ is tame. Choose ε smaller than the distance between any two homological critical values. For each pair of homological critical values $a < b$, we define their *multiplicity* μ_a^b which we interpret as the number of independent homology classes that are born at a and die at b. We count the homology classes born by time $a + \varepsilon$ that die after time $b - \varepsilon$. Among these subtract those born by $a - \varepsilon$ and subtract those that die after $b + \varepsilon$. This double counts those born by $a - \varepsilon$ that die after $b + \varepsilon$, so we add them back. That is,

$$\mu_a^b = \beta_{a+\varepsilon}^{b-\varepsilon} - \beta_{a-\varepsilon}^{b-\varepsilon} - \beta_{a+\varepsilon}^{b+\varepsilon} + \beta_{a-\varepsilon}^{b+\varepsilon}.$$

The *reduced persistence diagram* of μ, $\bar{D}(\mu)$, is the multiset of pairs (a, b) together with their multiplicities μ_a^b. We call this a diagram since it is convenient to plot these points on the plane. We will see that it is useful to add homology classes which are born and die at the same time. Let the *persistence diagram* of μ, $D(\mu)$, be given by the union of $\bar{D}(\mu)$ and $\{(a, a)\}_{a \in \mathbb{R}}$ where each (a, a) has infinite multiplicity.

A metric on the space of persistence diagrams is the bottleneck distance which bounds the Hausdorff distance [7]. It is given by

$$d_B(D(\mu), D(\nu)) = \inf_{\gamma} \sup_{p \in D(\mu)} \|p - \gamma(p)\|_\infty, \tag{3}$$

where the infimum is taken over all bijections $\gamma : D(\mu) \to D(\nu)$.

In [7], the following result is proven:

$$d_B(D(\mu), D(\nu)) \leq \|\mu - \nu\|_\infty \tag{4}$$

where $\mu, \nu : \mathbb{M} \to \mathbb{R}$ are tame functions. As an immediate consequence of (4), we can apply it to the model (1). Let $\Lambda_t(\beta, L)$ denote the subset of tame functions in $\Lambda(\beta, L)$ the class of Hölder functions

$$\Lambda(\beta, L) = \{f : S^2 \to \mathbb{R} \mid |f(x) - f(z)| \leq L(\arccos(x'y))^\beta, x, z \in S^2\}, \tag{5}$$

where $0 < \beta \leq 1$ and $L > 0$.

If we assume $\mu \in \Lambda_t(\beta, L)$ for the model (1) ϵ is $N(0, \sigma^2)$, for the estimator $\widehat{\mu}_n$ with $0 < \beta \leq 1$ and $L > 0$,

$$\sup_{\mu \in \Lambda_t(\beta, L)} \mathbb{E} d_B\left(D(\hat{\mu}_n), D(\mu)\right) \leq L^{2/(2\beta+2)} \left(\frac{\sigma^2 \ (\beta+2) 2^3}{\beta^2} \ \frac{\log n}{n} \right)^{\beta/(2\beta+2)}$$

as $n \to \infty$ [3].

SO HOW MUCH BOUND ARE WE EXPECTING IN NUMBER FOR OUR DATA? THIS IS A NICE RESULT BUT IMAGING PEOPLE WOULD WANT TO KNOW HOW MUCH n SHOULD BE TO DERIVE SOME VALID INFERENCE ON PD. GIVE SOME GUIDANCE. n SHOULD BE LARGER THAN 100? or 1000?

Consequently, because of the large n in medical image data, Figure 4 is an accurate description of the population parameters and therefore the extra layer of homology classes for the autistic group is most likely significant.

5 Discussion

We have presented the concept of persistence diagrams and described the filtration based algorithm for constructing the persistence diagrams. Since cortical thickness is highly noise, kernel smoothing is applied to remove high frequency spatial noise before the filtration. The constructed degree 0 and 1 persistence diagrams seem to show extra layer of homology classes for the autistic group (Figure 4). It seems the autistic subjects introduce more pairings. For the degree-1 persistence diagram, there are 907 and 1347 pairings respectively for controls and autistic subjects. For the degree-0 persistence diagram, there are 930 and 1349 pairings respectively for controls and autistic subjects. Cortical thickness has been shown to be a discriminating anatomical feature for autism although the underlying biomedical mechanism is still unknown [6]. From the previous

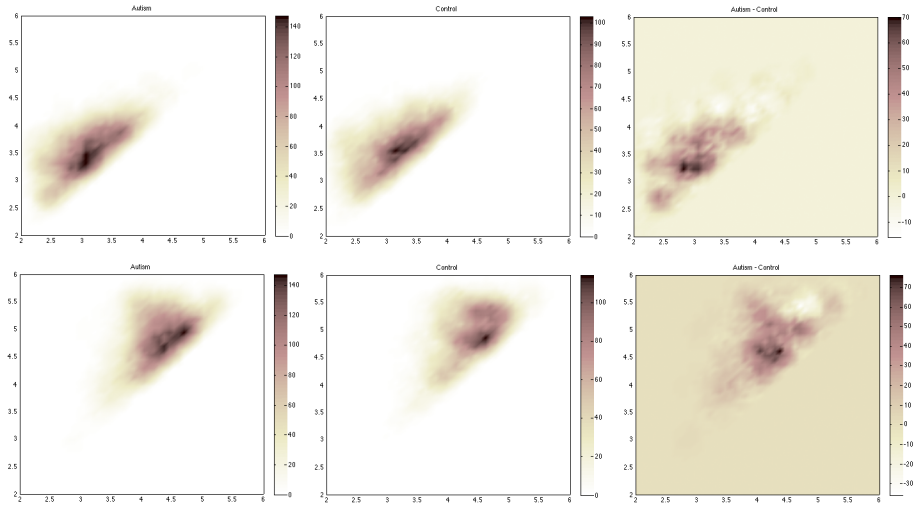

Fig. 5. Pairing concentration $P(x)$ is computed by counting the number of pairings within a circle of fixed radius $r = 0.2$ at the point $x \in [2, 6]^2$. The first row is for degree 0 persistence and the second row is for degree 1 persistence. The first (second) column is the concentration map of autistic (control). The concentration difference (autism - control) is given in the last column which shows dominating positive concentration in almost all points.

morphometric studies, it is found that autistic subjects exhibits more structural variability than normal controls. It is likely that autistics subjects exhibit more complex underlying geometric and topological pattern of brain.

It is unclear how one determines the statistical significance level of seemingly extra layer of homology classes for the lack of statistical inference frameworks for persistence diagrams. One may tempted to go for hypothesis-free classification frameworks for inference. However, figure 4 shows the classification based on possibly discriminating spatial pattern might be difficult if not impossible. Note that the autistic scatter plots basically encompass the control scatter plots for the degree 0 and degree 1 persistence diagrams. Since there are so much overlap, it may be difficult to directly apply various statistical and machine learning techniques. One the other hand, there seems to be spatial concentration difference in the pairings. To visualize if there is any, we have computed the pairing concentration $P(x)$, which is the number of parings within a circle of radius $r = 0.2$ at the point $x \in [2, 6]^2$. The pairing concentration maps are displayed in Figure 5. From the figure, we can see huge concentration difference for both the degree 0 and 1 persistence diagrams. The concentration maps for autism is lager in most points.

The inference on the concentration map is heuristically given as follows. The number of pairings in a circle will likely follow a Poisson distribution. Suppose

the concentration maps of autistic (P_a) and control (P_c) groups are Poisson with parameters λ_a and λ_c respectively. Under the null hypothesis of identical distributions across the groups, the concentration difference $P_a - P_c$ follows the Skellam distribution of the form $e^{-\lambda_a-\lambda_c}J_x(2)$, where J_x is the modified Bessel function of the first kind [18]. Since the Skellam distribution is cumbersome to manipulate, we approximate it with a Gaussian distribution $N(0, \lambda_a + \lambda_c)$ by matching the first moments and the variances. The variance $\lambda_a + \lambda_c$ is estimated as the pooled sample variance of the concentration maps, and is 32.21^2 and 31.86^2 for the degree 0 and 1 persistence diagrams respectively. The maximum concentration difference is 70 and 76 for the degree 0 and 1 persistence diagrams respectively. Then the pvalue of the one sided test of autism concentration larger than the control is 0.0149 and 0.0085 for the degree 0 and 1 persistence diagrams respectively. Since we took an approximation, we do not claim this to be exact but rather an approximation in the order of magnitude. A better inferencing might be achieved using nonparametric test procedures such as permutation tests.

Based on the above heuristic argument we have shown the significant concentration difference confirming the conjecture of autistic brain to be more likely to exhibit complex geometric and topological pattern than controls. It is hoped that this paper offers a spring board for further investigation of the persistence diagram based characterization of medical images. Due to the space limitation, there are a lot of methodological issues we have not discussed in the paper such as a more rigorous inferential procedure or the estimation of confidence circles around paired points in various ways including the bootstrap. These we hope to do in future works.

Acknowledgment

Authors wish to thank Vikas Singh of the University of Wisconsin-Madison for discussion on the persistence diagram.

References

1. S. Angenent, S. Hacker, A. Tannenbaum, and R. Kikinis. On the laplace-beltrami operator and brain surface flattening. *IEEE Transactions on Medical Imaging*, 18:700–711, 1999.
2. C. Brechbuhler, G. Gerig, and O. Kubler. Parametrization of closed surfaces for 3d shape description. *Computer Vision and Image Understanding*, 61:154–170, 1995.
3. P. Bubenik, G. Carlsson, P.T. Kim, and Z. Luo. Asymptotic minimax sup-norm risk on manifolds with application to topology. *Preprint*, 2008.
4. P. Bubenik and P.T. Kim. Aa statistical approach to persistent homology. *Homology Homotopy and Applications*, 9:337–362, 2007.
5. M.K. Chung, L. Shen Dalton, K.M., A.C. Evans, and R.J. Davidson. Weighted fourier representation and its application to quantifying the amount of gray matter. *IEEE Transactions on Medical Imaging*, 26:566–581, 2007.

6. M.K. Chung, S. Robbins, Davidson R.J. Alexander A.L. Dalton, K.M., and A.C. Evans. Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage*, 25:1256–1265, 2005.
7. D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37, 2007.
8. D.L. Collins, P. Neelin, T.M. Peters, and A.C. Evans. Automatic 3d intersubject registration of mr volumetric data in standardized talairach space. *J. Comput. Assisted Tomogr.*, 18:192–205, 1994.
9. Ghrist R. de Silva, V. Homological sensor networks. *Notic Amer Math Soc*, 54:10–17, 2007.
10. V. de Silva and P. Perry. Plex version 2.5. *http://math.stanford.edu/comptop/programs/plex*, 2008.
11. Dequent M-L. Mileyko Y. Edelsbrunner, H. and O. Pourquie. Assessing periodicity in gene expression as measured by microarray data. *Preprint*, 2008.
12. H. Edelsbrunner and J. Harer. Persistent homology - a survey. In *Twenty Years After, American Mathematical Society*, page in press, 2008.
13. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
14. B. Fischl and A.M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *PNAS*, 97:11050–11055, 2000.
15. K.J. Friston. A short history of statistical parametric mapping in functional neuroimaging. Technical Report Technical report, Wellcome Department of Imaging Neuroscience, ION, UCL., London, UK., 2002.
16. X. Gu, Y.L. Wang, T.F. Chan, T.M. Thompson, and S.T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23:1–10, 2004.
17. M. K. Hurdal and K. Stephenson. Cortical cartography using the discrete conformal approach of circle packings. *NeuroImage*, 23:S119S128, 2004.
18. J. O. Irwin. The frequency distribution of the difference between two independent variates following the same poisson distribution. *Journal of the Royal Statistical Society: Series A*, 100:415416, 1937.
19. S.C. Joshi. Large deformation diffeomorphisms and gaussian random fields for statistical characterization of brain sub-manifolds. 1998.
20. S.J. Kiebel, J.-P. Poline, K.J. Friston, A.P. Holmes, and K.J. Worsley. Robust smoothness estimation in statistical parametric maps using standardized residuals from the general linear model. *NeuroImage*, 10:756–766, 1999.
21. J. Klemelä. Asymptotic minimax risk for the white noise model on the sphere. *Scand J Statist*, 26:465–473, 1999.
22. K. Kollakian. Performance analysis of automatic techniques for tissue classification in magnetic resonance images of the human brain. Technical Report Master's thesis, Concordia University, Montreal, Quebec, Canada, 1996.
23. J. P. Lerch and A.C. Evans. Cortical thickness analysis examined through power analysis and a population simulation. *NeuroImage*, 24:163–173, 2005.
24. E. Luders, K.L. Narr, P.M. Thompson, D.E. Rex, R.P. Woods, and Jancke L. andToga A.W. DeLuca, H. Gender effects on cortical thickness and the influence of scaling. *Human Brain Mapping*, 27:314–324, 2006.
25. J.D. MacDonald, N. Kabani, D. Avis, and A.C. Evans. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from mri. *NeuroImage*, 12:340–356, 2000.

26. M.I. Miller, A. Banerjee, G.E. Christensen, S.C. Joshi, N. Khaneja, U. Grenander, and L. Matejic. Statistical methods in computational anatomy. *Statistical Methods in Medical Research*, 6:267–299, 1997.
27. M.I. Miller, A.B. Massie, J.T. Ratnanather, K.N. Botteron, and J.G. Csernansky. Bayesian construction of geometrically based cortical thickness metrics. *NeuroImage*, 12:676–687, 2000.
28. J. Milnor. *Morse Theory*. Princeton University Press, 1973.
29. D. Morozov. *Homological Illusions of Persistence and Stability*. Ph.D. thesis, Duke University, 2008.
30. D.Q. Naiman. Volumes for tubular neighborhoods of spherical polyhedra and statistical inference. *Ann. Statist.*, 18:685–716, 1990.
31. T. Nichols and S. Hayasaka. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Stat Methods Med. Res.*, 12:419–446, 2003.
32. Ozturk O. Ferhatosmanoglu H. Sacan, A. and Y. Wang.
33. D.O. Siegmund and K.J. Worsley. Testing for a signal with unknown location and scale in a stationary gaussian random field. *Annals of Statistics*, 23:608–639, 1996.
34. J.G. Sled, A.P. Zijdenbos, and A.C. Evans. A nonparametric method for automatic correction of intensity nonuniformity in mri data. *IEEE Transactions on Medical Imaging*, 17:87–97, 1988.
35. J.E. Taylor and K.J. Worsley. Random fields of multivariate test statistics, with applications to shape analysis. *Annals of Statistics*, page in press, 2008.
36. B. Timsari and R. Leahy. An optimization method for creating semi-isometric flat maps of the cerebral cortex. In *The Proceedings of SPIE, Medical Imaging*, 2000.
37. K.J. Worsley, S. Marrett, P. Neelin, A.C. Vandal, K.J. Friston, and A.C. Evans. A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping*, 4:58–73, 1996.
38. A. Yezzi and J.L. Prince. An eulerian pde approach for computing tissue thickness. *IEEE Transactions on Medical Imaging*, 22:1332–1339, 2003.
39. A.J. Zomorodian. *Computing and Comprehending Topology: Persistence and Hierarchical Morse Complexes*. Ph.D. Thesis, University of Illinois, Urbana-Champaign, 2001.
40. A.J. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33:249–274, 2005.