Manifold Learning on Fibre Bundles

Tingran Gao

Committee on Computational and Applied Mathematics Department of Statistics The University of Chicago

Computational and Applied Mathematics Colloquium The University of Chicago, Chicago IL

Thursday May 16, 2019

・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ うへで

1/76

Outline

Background

Spectral Geometry and Data Analysis

Manifold Learning on Fibre Bundles

- Motivation: Comparative Biology
- Horizontal Diffusion Maps

From Geometry to Learning: Synchronization Problems

Manifold Learning



Manifold Learning

Motivation

- Curse of Dimensionality: For D-dimensional data, need $\Omega\left(\epsilon^{-1}\exp\left(\frac{D}{\epsilon}\log\frac{1}{\epsilon}\right)\right)$ samples to ensure estimation error $\leq \epsilon$
- Manifold Assumption: Data lie approximately on a d-dimensional submanifold of ℝ^D, with d ≪ D; expecting sampling complexity Ω (ε⁻¹ exp (d log 1/ε)) instead



/76

Manifold Learning and Discrete Geometry – Various Interpretations

- Nonparametric Statistics: Regression and density estimation under manifold assumptions (L. Wasserman, P. Bickel et al.)
- Gaussian Processes: Gaussian process latent variable models (N. Lawrence et al.)
- Coarse Geometry: Geometric Whitney Problem Metric space approximation under the Gromov–Hausdorff distance (C. Fefferman et al.)
- Finite Elements: Discrete exterior calculus (P. Schröder, D. Arnold et al.)

Manifold Learning and Discrete Geometry – Various Interpretations

- Nonparametric Statistics: Regression and density estimation under manifold assumptions (L. Wasserman, P. Bickel et al.)
- Gaussian Processes: Gaussian process latent variable models (N. Lawrence et al.)
- Coarse Geometry: Geometric Whitney Problem Metric space approximation under the Gromov–Hausdorff distance (C. Fefferman et al.)
- Finite Elements: Discrete exterior calculus (P. Schröder, D. Arnold et al.)
- Spectral Geometry: Laplacian Eigenmaps (M. Belkin & P. Niyogi, S. Lafon & R. Coifman et al.)

CAN ONE HEAR THE SHAPE OF A DRUM?

MARK KAC, The Rockefeller University, New York

To George Eugene Uhlenbeck on the occasion of his sixty-fifth birthday

"La Physique ne nous donne pas seulement l'occasion de résoudre des problèmes . . . , elle nous fait presentir la solution." H. POINCARÉ.

$$\Delta_M u_n = -\lambda_n u_n, \quad n = 0, 1, 2, \dots$$
$$0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \dots \nearrow \infty$$

Does knowing all λ_n's determine M up to isometry?
 (Spoiler Alert: No)

$$\Delta_M u_n = -\lambda_n u_n, \quad n = 0, 1, 2, \dots$$
$$0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \dots \nearrow \infty$$

- ► Does knowing all λ_n's determine M up to isometry? (Spoiler Alert: No)
- Isospectral but non-isomstric:
 - Flat tori (Milnor 1964)
 - Riemann surfaces with constant negative curvature (Fignéras 1980)
 - Lens spaces with constant curvature (lkeda 1983)
 - Riemannian covering spaces (Sunada 1985)

$$\Delta_M u_n = -\lambda_n u_n, \quad n = 0, 1, 2, \dots$$
$$0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \dots \nearrow \infty$$

- ► Does knowing all λ_n's determine M up to isometry? (Spoiler Alert: No)
- Isospectral but non-isomstric:
 - Flat tori (Milnor 1964)
 - Riemann surfaces with constant negative curvature (Fignéras 1980)
 - Lens spaces with constant curvature (lkeda 1983)
 - Riemannian covering spaces (Sunada 1985)
- Cospectral graphs

$$\Delta_M u_n = -\lambda_n u_n, \quad n = 0, 1, 2, \dots$$
$$0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \dots \nearrow \infty$$

► Does knowing all \u03c6_n's and all \u03c6_{un}\u03c6's determine M up to isometry?

$$\Delta_M u_n = -\lambda_n u_n, \quad n = 0, 1, 2, \dots$$
$$0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \dots \nearrow \infty$$

► Does knowing all \u03c6_n's and all {u_n}'s determine M up to isometry?

Yes – via the *heat kernel*: $\forall x, y \in M, t \in [0, \infty)$

$$k_M(x,y;t) = \sum_{n=0}^{\infty} e^{-\lambda_n t} u_n(x) u_n(y)$$

Heat kernel completely determines the metric up to isometry

Spectral Embedding

(Bérard-Besson-Gallot 1994) Any closed Riemannian manifold in

$$\mathcal{M}_{n,k,D} = \left\{ \left(M, g \right) \, \middle| \, \dim \left(M \right) = n, \operatorname{Ric} \left(g \right) \ge \left(n - 1 \right) k g, \\ \operatorname{diam} \left(M \right) \le D \right\}$$

can be embedded into the infinite dimensional Hilbert space ℓ^2 using the *heat kernel map*

$$M \ni x \longmapsto \Phi_t(x) := \left(e^{-\lambda_0 t/2} u_0(x), e^{-\lambda_1 t/2} u_1(x), \dots\right) \in \ell^2$$

$$\blacktriangleright \langle \Phi_t(x), \Phi_t(y) \rangle_{\ell^2} = k_M(x, y; t) \text{ RKHS}$$

• Estimates for $k_M \leftrightarrow$ (Gromov) Precompactness of $\mathcal{M}_{n,k,D}$

<ロ > < 母 > < 三 > < 三 > < 三 > < 三 > < 9/76

Spectral Embedding: Beyond BBG'94

- (Jones-Maggioni-Schul 2010) Local, finitely many eigenfunctions used to make bi-Lipschitz coordinate charts
- ► (Bates 2014) Make (Jones-Maggioni-Schul 2010) global
- (Wang-Zhu 2015) Heat kernel maps can be made isometric using Nasher-Moser
- (Portegies 2016) Global, finite, and almost isomstric embedding using harmonic radius arguments
- (Wu 2017) BGG'94-type embedding, but with the heat kernels of the connection Laplacian (rough Laplacian)
- ▶ (Lin-Wu 2018) Embedding in (Wu 2017) can be made finite

Data Analysis: Discrete/Combinatorial/Probabilistic

Underpinning methodology (Lim 2015):

Graphs are discrete Riemannian manifolds

Riemannian Manifold	Graph
tangent vectors	edges
Laplacian	graph Laplacian
heat kernel	heat kernel
diffusion process	random walk

L.-K. Lim, "Hodge Laplacians on Graphs," arxiv:1507.05319

M. Belkin, P. Niyogi, "Laplacian Eigenmaps for Dimensionality Reduction and Data Representation," Neural Computation 15 (6), 1373-1396

R. Coifman, S. Lafon, "Diffusion Maps," Applied and Computational Harmonic Analysis 21 (2006), no. 1, 5-30

Random Walks "Knit Together" Local Geometry



◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□▶ ◆□▶

Diffusion Maps

- ▶ Data $\mathscr{X} := \{x_1, \dots, x_n\} \subset \mathbb{R}^D$, point cloud
- ▶ Build weighted nearest-neighbor graph G = (V, E, w) for \mathscr{X}
- ▶ Define weighted adjacency matrix W for G, and diagonal matrix $D \in \mathbb{R}^{n \times n}$ with

$$D_{ii} = \sum_{k=1}^{n} w_{ij}$$

▶ Build graph random walk Laplacian $L = I_n - D^{-1}W$, and perform eigen-decomposition

$$Lu_i = \lambda_i u_i, \quad i = 1, \ldots, n$$

• For any $1 \leq d \leq n$, embed $\mathscr X$ into $\mathbb R^d$ by

$$x_{k}\mapsto\left(\lambda_{1}^{1/2}u_{1}\left(k
ight),\ldots,\lambda_{d}^{1/2}u_{d}\left(k
ight)
ight)$$

▲□▶ ▲□▶ ▲□▶ ▲□▶ □□ − つへぐ

13/76

Graph Random Walk Laplacian

- $D^{-1}W$ defines a random walk on the graph G
- ► Typical choice for weights: For some *bandwidth* parameter ϵ > 0,

$$W_{ij} = \exp\left(-\left\|x_i - x_j\right\|^2 / \epsilon\right)$$

- The choice of ϵ reflects our belief for the "locality" of the data
- ► If we eigen-decompose W instead of D⁻¹W, the algorithm is known as Laplacian eigenmaps

Continuous Limit

• Explicitly, $D^{-1}W$ is a Markov operator: for any $v \in \mathbb{R}^n$,

$$D^{-1}Wv = \frac{\sum_{j=1}^{n} w_{ij}v_j}{\sum_{j=1}^{n} w_{ij}}$$

▶ As $n \to \infty$, $D^{-1}W$ converges weakly to the integral operator

$$P_{\epsilon}f(x) = \frac{\int_{M} \exp\left(-\|x-y\|^{2}/\epsilon\right) f(y) \operatorname{dvol}(y)}{\int_{M} \exp\left(-\|x-y\|^{2}/\epsilon\right) \operatorname{dvol}(y)}$$

for all $f \in L^1(M)$.

<□ > < □ > < □ > < Ξ > < Ξ > Ξ < つ < ○ 15/76</p>

Diffusion Maps: Asymptotic Theory

Theorem (Belkin-Niyogi, 2005). Let data points x_1, \dots, x_n be sampled from a **uniform** distribution on M. Under mild technical assumptions, there exist a sequence of real numbers $t_n \to 0$ and a constant C such that for any $f \in C^{\infty}(M)$

$$\lim_{n\to\infty}C\frac{\left(4\pi t_{n}\right)^{-\frac{k+2}{2}}}{n}\frac{P_{t_{n}}-I}{t_{n}}f\left(x\right)=\Delta_{M}f\left(x\right),\quad\forall x\in M.$$

Theorem (Coifman-Lafon, 2006). As $\epsilon \to 0$, for any $f \in C^{\infty}(M)$ and $x \in M$, if $\{x_i\}_{i=1}^n \sim p(x) \operatorname{dvol}_M(x)$, then w.h.p.

$$\begin{aligned} & P_{\epsilon}^{(\alpha)} f(x) \\ &= f(x) + \epsilon \frac{m_2}{2m_0} \left[\frac{\Delta_M \left[f p^{1-\alpha} \right](x)}{p^{1-\alpha}(x)} - f(x) \frac{\Delta_M p^{1-\alpha}(x)}{p^{1-\alpha}(x)} \right] + O(\epsilon^2) \,. \end{aligned}$$

<ロ> < □ > < □ > < 三 > < 三 > < 三 > ○ < ○ 16/76

Decoupling Geometry from Probability



・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・

The Convergence Rate: Diffusion Maps

Theorem (Singer, 2006). Suppose *N* points are i.i.d. uniformly sampled from a *d*-dimensional Riemannian manifold *M*. The graph diffusion operator $P_{\epsilon,\alpha}$ converges to its smooth limit at rate

$$O\left(N^{-\frac{1}{2}}\epsilon^{\frac{1}{2}-\frac{d}{4}}\right).$$

Corollary. Under the same assumption, non-uniform sampling has convergence rate

$$O\left(\mathbf{N}^{-\frac{1}{2}}\epsilon^{-\frac{d}{4}}\right).$$

・ロト ・ 日 ・ モ ・ モ ・ モ ・ つくぐ

18/76

1

This Talk: Horizontal Diffusion Maps on Fibre Bundles



<ロト < 部 > < 注 > < 注 > 注 の < で 19/76

Outline

Background

Spectral Geometry and Data Analysis

Manifold Learning on Fibre Bundles

- Motivation: Comparative Biology
- Horizontal Diffusion Maps

From Geometry to Learning: Synchronization Problems

A Motivating Case Study: Morphology and Classification



Systema Naturae, 1735



Carl Linnaeus (1707-1778)

<ロト</th>
・< 目 > < 目 > < 目 > < 目 > < 21/76</th>

Morphometrics

5.4

4.6

5.0

3.9

3.4

3.4

1.7

1.4

1.5

0.4

0.3

0.2



R.A. Fisher, "The Use of Multiple Measurements in Taxonomic Problems," Annals of Eugenics 7.2 (1936):

4.5

4.7

3.3

1.3

1.6

1.0

7.6

4.9

7.3

2.8

3.3

2.4

5.7

6.3

4.9

179-188.

3.0

2.5

2.9

2.1

1.7

1.8

6.6

4·5 6·3

Geometric Morphometrics



▲□▶ ▲圖▶ ▲圖▶ ▲圖▶ = ■ - ∽��?

23/76

Landmark-Based Geometric Morphometrics











Boyer et al. "Algorithms to Automatically Quantify the Geometric Similarity of Anatomical Surfaces," *Proceedings of the National Academy of Sciences* 108.45 (2011): 18221-18226.

Data Acquisition: microCT (High-Resolution X-ray CT)



Surface reconstructed from μ CT-scanned voxel data

Data Acquisition: Morphosource.org



Sharing the Bones

Duke researchers bring digital tools to the Stone Age findings of the Rising Star cave expedition

Writer: Louise Flynn December 11, 2015



Successful and the second seco



"Big Data" for Biologists: Impossible to landmark them all!!

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三日 • のへで

A Zoo of (Landmark-Free) Shape Distances...

 $\begin{array}{ll} d_{\rm cWn}\left(S_1,S_2\right): & {\rm Conformal \ Wasserstein \ Distance \ (CWD)} \\ d_{\rm cP}\left(S_1,S_2\right): & {\rm Continuous \ Procrustes \ Distance \ (CPD)} \\ d_{\rm cKP}\left(S_1,S_2\right): & {\rm Continuous \ Kantorovich-Procrustes \ Distance \ (CKPD)} \end{array}$

$$d_{\mathrm{cP}}\left(S_{1},S_{2}\right) = \inf_{\mathcal{C} \in \mathcal{A}\left(S_{1},S_{2}\right)} \inf_{R \in \mathbb{E}\left(3\right)} \left(\int_{S_{1}} \left\| R\left(x\right) - \mathcal{C}\left(x\right) \right\|^{2} d\mathrm{vol}_{S_{1}}\left(x\right) \right)^{\frac{1}{2}}$$







◆□ ▶ ◆□ ▶ ◆ 三 ▶ ◆ 三 ● ● ● ● ●

$$d_{\mathrm{cP}}\left(S_{i},S_{j}\right) = \inf_{\mathcal{C}\in\mathcal{A}\left(S_{i},S_{j}\right)} \inf_{R\in\mathbb{E}(3)} \left(\int_{S_{i}} \|R(x)-\mathcal{C}(x)\|^{2} d\mathrm{vol}_{S_{i}}(x)\right)^{\frac{1}{2}}$$



Si

▲□▶ ▲圖▶ ▲≧▶ ▲≧▶ 三三 のへで

28/76



29/76

▲□▶ ▲圖▶ ▲≣▶ ▲≣▶ ▲国 • のへぐ

Interpretability Issue



Even mistakes made by CPD were similar to biologists' mistakes!

30/76

Trust Small Distances — Let Diffusion Maps Do it for You!



"Correct" like a biologist, but *automatically*?

small distances between $S_1, S_2 \longrightarrow OK$ maps larger distances \longrightarrow not OK

Gao et al. (2018) "Development and Assessment of Fully Automated and Globally Transitive Geometric Morphometric Methods, with Application to a Biological Comparative Dataset with High Interspecific Variation," *The Anatomical Record* 301 (4), 636-658 (2017)

▲□▶ ▲□▶ ★ 国▶ ★ 国▶ → 国 → のへで →

Trust Only Small Distances: Geodesics in Shape Space



▲□▶ ▲圖▶ ▲圖▶ ▲圖▶ - 圖 - 釣ぬで

32/76
Geometric Model — Parallel Transport on Fibre Bundles



< □ ▶ < □ ▶ < Ξ ▶ < Ξ ▶ < Ξ > Ξ < つ < () 33/76

Geometric Model — Parallel Transport on Fibre Bundles



◆□▶ ◆□▶ ◆ □▶ ◆ □▶ ● □ ● ● ● ●

Shape Space is **NOT** a Trivial Fibre Bundle



▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへで

Ideal, Noise-Free Case: Trivial Bundle/No Holonomy



▲□▶ ▲圖▶ ▲≣▶ ▲≣▶ ▲国 ● ● ●

Cruel, Real Life: Non-trivial Bundle/Holonomy



◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 の�?

Two Possible Sources of the Holonomy

Random Noise/Discretization Error: Variance (G.,2015; G., 2016)

 The fibre bundle is truly non-trivial: Bias (G. et al., 2016; Bajaj et al., 2018)

• Tingran Gao, "Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics," PhD Thesis, Duke University (2015)

• Tingran Gao. "The Diffusion Geometry of Fibre Bundles: Horizontal Diffusion Maps." arXiv:1602.02330 (2016)

• Tingran Gao, Jacek Brodzki, Sayan Mukherjee. "The Geometry of Synchronization Problems and Learning Group Actions." Discrete & Computational Geometry, to appear (2019)

 Chandrajit Bajaj, Tingran Gao, Zihang He, Qixing Huang, and Zhenxiao Liang, "SMAC: Simultaneous Mapping and Clustering Using Spectral Decompositions," *Proceedings of the 35th International Conference on Machine Learning, PMLR* 80:324-333 (2018)

▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへの

Horizontal Random Walk on a Fibre Bundle



▲□▶ ▲□▶ ★ □▶ ★ □▶ □ のへで

Horizontal Diffusion Process in Stochastic Geometry

- K.D. Elworthy, W.S. Kendall. "Factorization of Harmonic Maps and Brownian Motions." University of Warwick, 1985.
- M. Liao, "Factorization of Diffusions on Fibre Bundles." Transactions of the American Mathematical Society. 311.2 (1989): 813-827.
- M. Arnaudon, A. Thalmaier. "Horizontal Martingales in Vector Bundles." Séminaire de Probabilits de Strasbourg. 36 (2002): 419-456.
- K.D. Elworthy, Y. Le Jan, and X. Li. "The Geometry of Filtering." Springer Basel, 2010. 33-59.
- F. Baudoin. "An Introduction to the Geometry of Stochastic Flows." London: Imperial College Press, 2004.



◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 三臣 - のへで







◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 三臣 - のへで - -



▲□▶ ▲□▶ ▲注▶ ▲注▶ 注目 めんの



▲□▶ ▲□▶ ▲ 臣▶ ▲ 臣▶ 臣 - ∽ Q (~ 44/76

Assume there are *n* shapes, S_1, \ldots, S_n , and each S_i has ℓ_i vertices.

▶ Compute $f_{ij}: S_i \to S_j$ for each pair of $1 \le i \ne j \le n$

•
$$f_{ij} \longleftrightarrow \rho_{ij}^{\delta} \in \mathbb{R}^{\ell_j \times \ell}$$

• Let \mathcal{W} be an *n*-by-*n* **block** matrix, with (i, j)th block

$$\exp\left(-\frac{d_{\mathrm{cP}}^2\left(\mathit{S}_i, \mathit{S}_j\right)}{\epsilon}\right) \cdot \rho_{ij}^{\delta} \in \mathbb{R}^{\ell_j \times \ell_i}$$

- ϵ : horizontal bandwidth parameter
- δ : vertical bandwidth parameter
- Let D be a diagonal matrix, with kth diagonal entry equaling the kth row sum of W
- Horizontal Random Walk Laplacian: $I D^{-1}W$

Remark. \mathcal{W} can be viewed as a flattening of a 4-tensor



HDM: Continuous Limit

M – base manifold; F – template fibre; E – total manifold; F_x – fibre over $x \in M$; $P_{yx} : F_x \to F_y$ – parallel transport

As n, ℓ_i → ∞ for all 1 ≤ ℓ ≤ n, D⁻¹W converges weakly to the integral operator H_{ε,δ} : L¹(E) → L¹(E)

$$=\frac{\int_{M}\int_{F_{y}}K_{\epsilon,\delta}(x,v;y,w)f(y,w)p(y,w)d\mathrm{vol}_{F_{y}}(w)d\mathrm{vol}_{M}(y)}{\int_{M}\int_{F_{y}}K_{\epsilon,\delta}(x,v;y,w)p(y,w)d\mathrm{vol}_{F_{y}}(w)d\mathrm{vol}_{M}(y)}$$

for all $f \in L^{1}(E)$ and $v \in F_{x}$, $w \in F_{y}$, where

$$K_{\epsilon,\delta}(x,v;y,w) = \exp\left(-\frac{d_M^2(x,y)}{\epsilon} - \frac{d_{F_y}^2(P_{yx}v,w)}{\delta}\right)$$

◆□ ▶ ◆□ ▶ ◆ 臣 ▶ ◆ 臣 ◆ ○ ♀ ○ 47/76

Horizontal Diffusion Maps: Embedding the Entire Bundle



Horizontal Diffusion Maps: Embedding the Entire Bundle



▲ロト ▲御 → ▲ 臣 → ▲ 臣 → 今久で -----

Horizontal Diffusion Maps



▲□▶ ▲□▶ ▲三▶ ▲三▶ 三三 - シタペー

Automatic Landmarking — Interpretability





▲口▶ ▲圖▶ ▲圖▶ ▲圖▶ 三国 - 釣A@ -



<ロ> < 理> < 理> < 理> < 理> < 理> のQ()



▲口▶ ▲御▶ ▲注▶ ▲注▶ 「注」 釣A@



◆□▶ ◆□▶ ◆ ≧▶ ◆ ≧ ▶ ○ ② ○ ○ 55/76



<ロト < @ ト < 臣 > < 臣 > ○ ○ ○ 5

Species Clustering



Horizontal Base Diffusion Distance (with Maps)



Diffusion Distance (without Maps)

・ロ ・ ・ 「日 ・ ・ 三 ・ ・ 三 ・ つ へ (や 57/76)

Species Clustering



Horizontal Base Diffusion Distance (with Maps)





▲□▶ ▲□▶ ▲□▶ ▲□▶ □ の�?

HDM: Continuous Limit

$$=\frac{\int_{M}\int_{F_{y}}K_{\epsilon,\delta}(x,v;y,w)f(y,w)p(y,w)d\mathrm{vol}_{F_{y}}(w)d\mathrm{vol}_{M}(y)}{\int_{M}\int_{F_{y}}K_{\epsilon,\delta}(x,v;y,w)p(y,w)d\mathrm{vol}_{F_{y}}(w)d\mathrm{vol}_{M}(y)}$$

for all $f \in L^{1}(E)$ and $v \in F_{x}$, $w \in F_{y}$, where

$$K_{\epsilon,\delta}(x,v;y,w) = \exp\left(-\frac{d_M^2(x,y)}{\epsilon} - \frac{d_{F_y}^2(P_{yx}v,w)}{\delta}\right)$$

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三回 のへの

Theorem (G., 2016). If
$$\delta = O(\epsilon)$$
 as $\epsilon \to 0$, then for any
 $f \in C^{\infty}(E)$ and $(x, v) \in E$, as $\epsilon \to 0$,

$$H_{\epsilon,\delta}^{(\alpha)} f(x, v)$$

$$= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H (fp^{1-\alpha})(x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_H p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right]$$

$$+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_E^V (fp^{1-\alpha})(x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_E^V p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right]$$

$$+ O(\epsilon^2 + \epsilon \delta + \delta^2).$$

Tingran Gao, "The Diffusion Geometry of Fibre Bundles: Horizontal Diffusion Maps," arXiv:1602.02330

<□▶ <□▶ < 글▶ < 글▶ < 글 < ⊃ < 60/76</p>

Theorem (G., 2016). If
$$\delta = O(\epsilon)$$
 as $\epsilon \to 0$, then for any
 $f \in C^{\infty}(E)$ and $(x, v) \in E$, as $\epsilon \to 0$,

$$H_{\epsilon,\delta}^{(\alpha)}f(x, v)$$

$$= f(x, v) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H(fp^{1-\alpha})(x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_H p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right]$$

$$+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_E^V(fp^{1-\alpha})(x, v)}{p^{1-\alpha}(x, v)} - f(x, v) \frac{\Delta_E^V p^{1-\alpha}(x, v)}{p^{1-\alpha}(x, v)} \right]$$

$$+ O(\epsilon^2 + \epsilon \delta + \delta^2).$$

$$\blacktriangleright \Delta_E^V$$
 is the vertical Laplacian on E

< □ > < □ > < □ > < Ξ > < Ξ > Ξ の < ⊙ 61/76

Theorem (G., 2016). If $\delta = O(\epsilon)$ as $\epsilon \to 0$, then for any $f \in C^{\infty}(E)$ and $(x, v) \in E$, as $\epsilon \to 0$,

$$\begin{aligned} H_{\epsilon,\delta}^{(\alpha)}f(x,v) &= f(x,v) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H \left(fp^{1-\alpha} \right) (x,v)}{p^{1-\alpha} (x,v)} - f(x,v) \frac{\Delta_H p^{1-\alpha} (x,v)}{p^{1-\alpha} (x,v)} \right] \\ &+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_E^V \left(fp^{1-\alpha} \right) (x,v)}{p^{1-\alpha} (x,v)} - f(x,v) \frac{\Delta_E^V p^{1-\alpha} (x,v)}{p^{1-\alpha} (x,v)} \right] \\ &+ O \left(\epsilon^2 + \epsilon \delta + \delta^2 \right). \end{aligned}$$

・ロト ・ 同 ト ・ ヨ ト ・ ヨ ・ りへぐ

61/76

• Δ_E^V is the vertical Laplacian on E

 $\blacktriangleright \Delta_H$ is the Bochner horizontal Laplacian on E

Theorem (G., 2016). If $\delta = O(\epsilon)$ as $\epsilon \to 0$, then for any $f \in C^{\infty}(E)$ and $(x, v) \in E$, as $\epsilon \to 0$,

$$\begin{split} H_{\epsilon,\delta}^{(\alpha)}f\left(x,v\right) &= f\left(x,v\right) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_H \left(fp^{1-\alpha}\right)\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} - f\left(x,v\right) \frac{\Delta_H p^{1-\alpha}\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} \right] \\ &+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_E^V \left(fp^{1-\alpha}\right)\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} - f\left(x,v\right) \frac{\Delta_E^V p^{1-\alpha}\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} \right] \\ &+ O\left(\epsilon^2 + \epsilon\delta + \delta^2\right). \end{split}$$

• Δ_E^V is the vertical Laplacian on E

- $\blacktriangleright \Delta_H$ is the Bochner horizontal Laplacian on E
- ▶ In general $\Delta_H + \Delta_E^V \neq \Delta_E$, true if and only if π is harmonic

Theorem (G., 2016). If $\delta = O(\epsilon)$ as $\epsilon \to 0$, then for any $f \in C^{\infty}(E)$ and $(x, v) \in E$, as $\epsilon \to 0$,

$$\begin{split} H_{\epsilon,\delta}^{(\alpha)}f\left(x,v\right) &= f\left(x,v\right) + \epsilon \frac{m_{21}}{2m_0} \left[\frac{\Delta_{H}\left(fp^{1-\alpha}\right)\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} - f\left(x,v\right) \frac{\Delta_{H}p^{1-\alpha}\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} \right] \\ &+ \delta \frac{m_{22}}{2m_0} \left[\frac{\Delta_{E}^{V}\left(fp^{1-\alpha}\right)\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} - f\left(x,v\right) \frac{\Delta_{E}^{V}p^{1-\alpha}\left(x,v\right)}{p^{1-\alpha}\left(x,v\right)} \right] \\ &+ O\left(\epsilon^{2} + \epsilon\delta + \delta^{2}\right). \end{split}$$

• Δ_E^V is the vertical Laplacian on E

- Δ_H is the Bochner horizontal Laplacian on E
- ▶ In general $\Delta_H + \Delta_E^V \neq \Delta_E$, true if and only if π is harmonic
- ► ⇒ HDM is not applying diffusion maps to the total space!

The Convergence Rate: Diffusion Maps

Theorem (Singer, 2006). Suppose *N* points are i.i.d. uniformly sampled from a *d*-dimensional Riemannian manifold *M*. The graph diffusion operator $P_{\epsilon,\alpha}$ converges to its smooth limit at rate

$$O\left(N^{-\frac{1}{2}}\epsilon^{\frac{1}{2}-\frac{d}{4}}\right).$$

Corollary. Under the same assumption, non-uniform sampling has convergence rate

$$O\left(N^{-\frac{1}{2}}\epsilon^{-\frac{d}{4}}\right).$$

・ロト ・ 日 ・ モ ・ モ ・ モ ・ つくぐ

The Convergence Rate: HDM on Unit Tangent Bundles

Theorem (G., 2015). Suppose N_B points are i.i.d. sampled from a *d*-dimensional Riemannian manifold M, and N_F unit tangent vectors are i.i.d. sampled at each of the N_B samples. The graph horizontal diffusion operator $H_{\epsilon,\delta}^{(\alpha)}$ converges to its smooth limit at rate

$$O\left(\theta_*^{-1} N_B^{-\frac{1}{2}} \epsilon^{-\frac{d}{4}}\right),$$

where

$$heta_* = 1 - rac{1}{1 + \epsilon^{rac{d}{4}} \delta^{rac{d-1}{4}} \sqrt{rac{N_F}{N_B}}}.$$

Remark. Diffusion maps on the total manifold in this setting has convergence rate of $O\left(N_F^{-\frac{1}{2}}N_B^{-\frac{1}{2}}\epsilon^{-\frac{d}{4}}\right)$, by (Singer, 2006)

Tingran Gao, "Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics," PhD Thesis, Duke University (2015)

・ロ ・ ・ @ ・ ・ 注 ・ 注 ・ う へ (?)
63/76

Outline

Background

Spectral Geometry and Data Analysis

Manifold Learning on Fibre Bundles

- Motivation: Comparative Biology
- Horizontal Diffusion Maps

From Geometry to Learning: Synchronization Problems
Synchronization Problems

$$y_i = R_i x + \xi_i$$

 $R_i \in O(d), \quad \xi_i \sim \text{i.i.d. noise}$



Afonso S. Bandeira. "Ten Lectures and Forty-Two Open Problems in the Mathematics of Data Science." (2015).

うっつ 叫 へ叫 く 山 マ ふし く 自 マ

65/76

Average Out the Variance, but Respect the Bias



Tingran Gao, Jacek Brodzki, Sayan Mukherjee. "The Geometry of Synchronization Problems and Learning Group Actions." Discrete & Computational Geometry, to appear (2019)

<ロ> < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

Algorithm: SynCut

Input: $\Gamma = (V, E, w), \rho \in C^1(\Gamma; G)$, number of partitions KOutput: Partitions $\{S_1, \dots, S_K\}$

- 1. Solve synchronization problem over Γ for ρ , obtain $f \in C^0(\Gamma; G)$
- 2. Compute $d_{ij} = \exp\left(-w_{ij} \|f_i \rho_{ij}f_j\|\right)$ on all edges $(i, j) \in E$
- 3. Spectral clustering on weighted graph (V, E, d) to get $\{S_1, \dots, S_k\}$
- 4. Solve synchronization problem within each partition S_j , "glue up" the local solutions to obtain $f_* \in C^0(\Gamma; G)$
- 5. $f \leftarrow f_*$, repeat from Step 2

Tingran Gao, Jacek Brodzki, Sayan Mukherjee. "The Geometry of Synchronization Problems and Learning Group Actions." Discrete & Computational Geometry, to appear (2019)

Compare with Graph-based Spectral Clustering



- ► Consider graph Γ = (V, E), where V = (v₁, · · · , v_n) are separated into 2 communities
- ► Fix orthogonal matrices R₁, · · · , R_n ∈ O(d), one at each corresponding vertex
- Fix $0 \le q . Consider the following <math>\rho \in C^1(\Gamma; O(d))$:
 - If v_i, v_j belong to the same community,

$$ho_{ij} = egin{cases} R_i^{ op} R_j & ext{with probability } p \ ext{Unif}\left(O\left(d
ight)
ight) & ext{with probability } 1-p \end{cases}$$

If v_i, v_j belong to different communities

$$ho_{ij} = egin{cases} R_i^{ op} R_j & ext{with probability } q \ ext{Unif}\left(O\left(d
ight)
ight) & ext{with probability } 1-q \end{cases}$$

<ロ> < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < 0 < 0 < 69/76

Fix $0 \le q . Consider the following <math>\rho \in C^1(\Gamma; O(d))$:

If v_i, v_j belong to the same community,

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } p \\ \text{Unif}(O(d)) & \text{with probability } 1 - p \end{cases}$$

If v_i, v_j belong to different communities

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } q \\ \text{Unif}(O(d)) & \text{with probability } 1 - q \end{cases}$$

▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへの

70/76

• "Clean" case (p = q = 0): synchronization solution

$$\tilde{\rho}_{ij} = R_i^\top R_j$$

- Fix $0 \le q . Consider the following <math>\rho \in C^1(\Gamma; O(d))$:
 - If v_i , v_j belong to the same community,

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } p \\ \text{Unif}(O(d)) & \text{with probability } 1 - p \end{cases}$$

If v_i, v_j belong to different communities

$$ho_{ij} = egin{cases} R_i^ op R_j & ext{with probability } q \ ext{Unif}\left(O\left(d
ight)
ight) & ext{with probability } 1-q \end{cases}$$

"Perturbed case (q > p > 0)": synchronization solution

$$\tilde{\rho}_{ij} \approx R_i^\top R_j$$

<□ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

Fix $0 \le q . Consider the following <math>\rho \in C^1(\Gamma; O(d))$:

If v_i, v_j belong to the same community,

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } p \\ \text{Unif}(O(d)) & \text{with probability } 1 - p \end{cases}$$

► If v_i, v_j belong to different communities

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } q \\ \text{Unif}(O(d)) & \text{with probability } 1 - q \end{cases}$$

"Perturbed case (q > p > 0)": synchronization solution

$$\tilde{\rho}_{ij} \approx R_i^{\top} R_j$$

▶ But || p̃_{ij} − p_{ij} || deviates from 0 more frequently on edges connecting different communities!

- Fix $0 \le q . Consider the following <math>\rho \in C^1(\Gamma; O(d))$:
 - If v_i , v_j belong to the same community,

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } p \\ \text{Unif}(O(d)) & \text{with probability } 1 - p \end{cases}$$

► If v_i, v_j belong to different communities

$$\rho_{ij} = \begin{cases} R_i^\top R_j & \text{with probability } q \\ \text{Unif}(O(d)) & \text{with probability } 1 - q \end{cases}$$

"Perturbed case (q > p > 0)": synchronization solution

$$\tilde{\rho}_{ij} \approx R_i^\top R_j$$

- ▶ But || p̃_{ij} p_{ij} || deviates from 0 more frequently on edges connecting different communities!
- ▶ suggesting normalized graph cut with new weights $\|\tilde{\rho}_{ij} \rho_{ij}\|$

71/76

Simultaneous Matching and Clustering (SMAC)

- Established results for permutation groups in the context of collection shape matching
- Stochastic Block + Random Corruption for the maps, Erdős-Rényi for the observation graph, K ≥ 2 balanced clusters, synchronization with spectral methods
- ► The recovery rate for the underlying clusters matches the information theoretic lower bound in [Chen et al. 2017]

• Yuxin Chen, Changho Suh, and Andrea J. Goldsmith. "Information Recovery from Pairwise Measurements." *IEEE Transactions on Information Theory* 62, no. 10 (2016): 5881-5905.

• Chandrajit Bajaj, **Tingran Gao**, Zihang He, Qixing Huang, and Zhenxiao Liang, "SMAC: Simultaneous Mapping and Clustering Using Spectral Decompositions," *Proceedings of the 35th International Conference on Machine Learning, PMLR* 80:324-333, 2018 Algorithm 1 PermSMAC: Simultaneously mapping and clustering

- **Input:** Observation graph $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ and initial pairwise maps $X_{ij}^{in}, (i, j) \in \mathcal{E}$
- **Output:** Underlying clusters $S = c_1 \cup \cdots \cup c_k$ and optimized pairwise maps $X_{ij}, 1 \le i, j \le n$
 - 1: {**Step 1**} Simultaneously compute the intra-cluster maps and extract the underlying clusters:
 - 2: {**Step 1.1**} Form data matrix based on (1).
 - 3: {Step 1.2} Compute the critical value $r = \underset{2 \le i \le nm}{\operatorname{argmax}} \frac{\lambda_i \lambda_{i+1}}{\lambda_i + \lambda_{i+1}}.$
 - 4: {Step 1.3} Let $U \in \mathbb{R}^{nm \times r}$ store the leading r eigenvectors of X. Compute pair-wise maps X_{ij}^{\star} by solving (2)
 - 5: {Step 1.4} Use $f_{ij}(X_{ij}^{\star})$ as the affinity score and apply single-linkage clustering to obtain the underlying clusters
- 6: {**Step 2**} compute the inter-cluster maps by solving (6)

SMAC: Exact Recovery Conditions

Theorem (Bajaj et al. 2018). Under mild assumptions, PermSMAC recovers the underlying clusters and intra-cluster maps with high probability provided

$$p-q \ge ck\sqrt{rac{\log n}{nt}}$$

for some constant c > 0, where $t = \Omega (\log (n) / n)$. The inter-cluster maps are recovered with high probability provided

$$q \ge c'k\sqrt{rac{\log n}{n^2t}}$$

for some constant c' > 0.

 Chandrajit Bajaj, Tingran Gao, Zihang He, Qixing Huang, and Zhenxiao Liang, "SMAC: Simultaneous Mapping and Clustering Using Spectral Decompositions," Proceedings of the 35th International Conference on Machine Learning, PMLR 80:324-333, 2018

<ロト
・ロト
・日
・<</p>
・<</p>
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・
・</

Collaborators







Jacek Brodzki Southamption

Ingrid Daubechies Duke



Qixing Huang UT Austin



Sayan Mukherjee Duke



Chandrajit Bajaj UT Austin

75/76

Thank You!



• Tingran Gao, "Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics," PhD Thesis, Duke University (2015)

• Tingran Gao, "The Diffusion Geometry of Fibre Bundles: Horizontal Diffusion Maps." arXiv:1602.02330

• Tingran Gao, Jacek Brodzki, Sayan Mukherjee. "The Geometry of Synchronization Problems and Learning Group Actions." Discrete & Computational Geometry, to appear (2019)

 Chandrajit Bajaj, Tingran Gao, Zihang He, Qixing Huang, and Zhenxiao Liang, "SMAC: Simultaneous Mapping and Clustering Using Spectral Decompositions," Proceedings of the 35th International Conference on Machine Learning, PMLR 80:324-333, 2018