

Solution manifold and its statistical applications*

Yen-Chi Chen

*Department of Statistics,
University of Washington, Seattle
e-mail: yenchi@uw.edu*

Abstract: A solution manifold is the collection of points in a d -dimensional space satisfying a system of s equations with $s < d$. Solution manifolds occur in several statistical problems including missing data, algorithmic fairness, hypothesis testing, partial identifications, and nonparametric set estimation. We theoretically and algorithmically analyze solution manifolds. In terms of theory, we derive five useful results: smoothness theorem, stability theorem (which implies the consistency of a plug-in estimator), convergence of a gradient flow, local center manifold theorem and convergence of the gradient descent algorithm. We propose a Monte Carlo gradient descent algorithm to numerically approximate a solution manifold. In the case of the likelihood inference, we design a manifold constraint maximization procedure to find the maximum likelihood estimator on the manifold.

Keywords and phrases: Set estimation, gradient descent, nonconvex optimization, manifold.

Received September 2021.

1. Introduction

A solution manifold [64] is the collection of points in d -dimensional space that solves a system of s equations where $s < d$. Namely, feasible set is a collection of points in an under-constrained system. Under smoothness conditions, the feasible set forms a manifold known as a solution manifold.

The solution manifold occurs in many problems in statistics such as missing data (Example 1), algorithmic fairness (Example 2), constrained likelihood space (Example 3), and density ridges/level sets (Example 4). In the regular case that $s = d$, the solution manifold reduces to the usual problems such as the Z-estimators [77] or estimating equations [45]. While there has been a tremendous amount of literature on the analysis of regular cases ($s = d$), little is known when $s < d$. This study aims to analyze the problem when $s < d$ and design a practical algorithm to find the manifold.

Formally, let $\Psi : \mathbb{R}^d \mapsto \mathbb{R}^s$ be a vector-valued function with $s < d$. The solution set of Ψ

$$M = \{x : \Psi(x) = 0\}$$

*Supported by NSF grant DMS 2112907 and DMS 195278 and NIH grant U24 AG072122.




FIG 1. An example of a solution manifold formed by the parameter (μ, σ) of a Gaussian with a tail probability bound $P(-5 < Y < 2) = 0.5$. The left panel shows 1000 random initializations (uniformly distributed within $[1, 3] \times [2, 4]$). We keep applying the gradient descent algorithm until convergence (right panel). The black dashed line indicates the actual location of the solution manifold.

is called the solution manifold and we call Ψ the generator (function) of M . Note that in some applications, x represents the parameter in a model; thus, sometimes we write $M = \{\theta : \Psi(\theta) = 0\}$. Here we provide examples of solution manifolds from various statistical problems.

Example 1 (Missing data). Consider a simple missing data problem where we have a binary response variable Y and a binary covariate X . The response variable is subject to missing. We use a binary variable R to indicate the response pattern of Y (i.e., Y is observed if $R = 1$). Depending on the value of R , we may observe $(X, Y, R = 1)$ or $(X, R = 0)$. In this case, the entire distribution is characterized by the following parameters:

$$\zeta_{x,y} = P(R = 1 | X = x, Y = y), \quad \mu_x = P(Y = 1 | X = x), \quad \xi = P(X = x)$$

for $x, y \in \{0, 1\}$. Parameter $\zeta_{x,y}$ is called missing data mechanism [47]. μ_x is the regression function, and ξ is the marginal mean of X . Thus, this problem has seven parameters. From the observed data (IID elements in the form of $(X, Y, R = 1)$ or $(X, R = 0)$), we can identify $P(x, y, R = 1)$ and $P(x, R = 0)$ for $x, y \in \{0, 1\}$, which leads to six constraints (note that $P(x, y, r) = P(X = x, Y = y, R = r)$):

$$\begin{aligned} P(1, 1, 1) &= \zeta_{11}\mu_1\xi, & P(1, 0, 1) &= \zeta_{10}(1 - \mu_1)\xi, \\ P(0, 1, 1) &= \zeta_{01}\mu_0(1 - \xi), & P(0, 0, 1) &= \zeta_{00}(1 - \mu_0)(1 - \xi) \\ P(X = 0, R = 0) &= (1 - \zeta_{01})\mu_0(1 - \xi) + (1 - \zeta_{00})(1 - \mu_0)(1 - \xi) \\ P(X = 1, R = 0) &= (1 - \zeta_{11})\mu_1\xi + (1 - \zeta_{10})(1 - \mu_1)\xi. \end{aligned} \quad (1)$$

Thus, the feasible values of the parameters $(\zeta_{x,y}, \mu_x, \xi)$ will form a solution manifold and the above constraints describe the generator Ψ . Note that the resulting solution manifold is related to the nonparametric bound of the parameters [49, 50, 17]; see Section 4.1 for more discussions.

Example 2 (Algorithmic fairness). The algorithmic fairness is a trending topic in modern machine learning research [40, 26, 27]. We consider a post-processing

method in the algorithmic fairness study. Suppose we have a binary response $Y \in \{0, 1\}$, a sensitive binary variable $A \in \{0, 1\}$ that we wish to protect, and an output from a trained classifier $W \in \{0, 1\}$ (one can view it as $W = c(X, A)$, where X is the covariate/feature and c is a trained classifier). The sensitive variable is often the race or gender indicator. The classification result based on W may discriminate against the sensitive variable A ; that is, it is likely that $A = 1$ and $W = 1$ occur at the same time. We want to construct a new ‘fair’ classifier $Q \in \{0, 1\}$ such that Q is a random variable whose distribution depends only on A and W and the output of Q will not discriminate against A . In other words, we design a new random variable Q such that $Q \perp Y|A, W$. While there are many principles of algorithmic fairness, we consider the test fairness [26]: We want to construct Q such that

$$P(Y = 1|Q = s, A = 0) = P(Y = 1|Q = s, A = 1), \quad (2)$$

for each $s = 0, 1$. To construct Q that satisfies the above constraint, we generate Q based on A, W such that its distribution is determined by parameter $q_{w,a} = P(Q = 1|W = w, A = a)$. As long as we can properly choose $q_{w,a}$, the resulting Q will satisfy equation (2). In this case, any $q_{w,a}$ solving the following two equations will satisfy equation (2) (see Appendix B):

$$\begin{aligned} \frac{\sum_w q_{w,0} P(W = w, Y = 1|A = 0)}{\sum_{w'} q_{w',0} P(W = w'|A = 0)} &= \frac{\sum_w q_{w,1} P(W = w, Y = 1|A = 1)}{\sum_{w'} q_{w',1} P(W = w'|A = 1)}. \\ \frac{\sum_w (1 - q_{w,0}) P(W = w, Y = 1|A = 0)}{\sum_{w'} (1 - q_{w',0}) P(W = w'|A = 0)} &= \frac{\sum_w (1 - q_{w,1}) P(W = w, Y = 1|A = 1)}{\sum_{w'} (1 - q_{w',1}) P(W = w'|A = 1)}. \end{aligned} \quad (3)$$

Note that $P(W = w, Y = y, A = a)$ is identifiable from the data. The original parameter $\{q_{w,a} : w, a = 0, 1\}$ is in four-dimensional space and we have two constraints, leading to a solution manifold of two dimensions.

Example 3 (Constrained likelihood space). Consider a random variable Y from an unknown distribution. We place a parametric model $p(y; \theta)$ of the underlying PDF of Y where $\theta \in \mathbb{R}^d$ is the parameter vector. Suppose that we have a set of constraints on the model such that a feasible parameter must satisfy

$$f_1(\theta) = f_2(\theta) = \dots = f_s(\theta) = 0$$

for some given functions f_1, \dots, f_s . These functions may be from independence assumptions or moment constraints $\mathbb{E}(g_j(Y)) = 0$ for some functions g_1, \dots, g_s . The set of parameters that satisfies these constraints is

$$\Theta_0 = \{\theta : f_\ell(\theta) = 0, \ell = 1, \dots, s\} = \{\theta : \Psi(\theta) = 0\}, \quad (4)$$

which is a solution manifold with $\Psi_\ell(\theta) = \int f_\ell(y) p(y; \theta) dy$. The above model is used in algebraic statistics [32, 55], partial identification problems with equality constraints [38, 25], and mixture models with moment constraints [46, 14]. We will return to this problem in Section 4.3.2. Figure 1 shows an example of a solution manifold formed by the tail probability constraint $P(-5 < Y < 2) = 0.5$

where $Y \sim N(\mu, \sigma^2)$. We have two parameters (μ, σ^2) and one constraint; thus, the resulting solution set is a one-dimensional manifold. From the left to the right panels, we show that we can recover the underlying manifold by random initializations with a suitable gradient descent process (Algorithm 1).

Example 4 (Density ridges). A k -ridge [35] of a density function $p(x)$ is defined as the collection of points satisfying

$$\{x : V_k(x)^T \nabla p(x) = 0, \lambda_k(x) < 0\},$$

where $V_k(x) = [v_k(x), \dots, v_d(x)] \in \mathbb{R}^{(k-d)}$ is the collection of eigenvectors of $H(x) = \nabla \nabla p(x)$, and $\lambda_k(x)$ is the k -th eigenvalue. The eigen-pairs are ordered as $\lambda_1(x) \geq \lambda_2(x) \geq \dots \geq \lambda_d(x)$. In this case $\Psi(x) = V_k(x)^T \nabla p(x)$; hence, ridges are also solution manifolds. In addition to ridges, the level sets and critical points of a function [79, 48, 13] are examples of solution manifolds. We will discuss this in Section 4.4

Although the aforementioned examples are from different statistical problems, they all share a similar structure that the feasible set forms a solution manifold. Thus, we study properties of solution manifolds in this paper, and our results will be applicable to all of these cases.

Main results. Our main results include theoretical developments and algorithmic innovations. In the theoretical analysis, we show that under a similar sets of assumptions, we have the following:

1. **Smoothness theorem.** The solution manifold is a $(d - s)$ -dimensional manifold with a positive reach (Lemma 1 and Theorem 3).
2. **Stability theorem.** As long as $\hat{\Psi}$ and Ψ and their derivatives are sufficiently close, $\hat{M} = \{x : \hat{\Psi}(x) = 0\}$ converges to M under the Hausdorff distance (Theorem 5).
3. **Convergence of a gradient flow.** For the gradient descent flow of $\|\Psi(x)\|^2$, the flow converges (in the normal direction of M) to a point on M when the starting point is sufficiently close to M (Theorem 7).
4. **Local center manifold theorem.** The collection of points converging to the same location $z \in M$ forms an s -dimensional manifold (Theorem 8).
5. **Convergence of a gradient descent algorithm.** With a good initialization, the gradient descent algorithm of $\|\Psi(x)\|^2$ converges linearly to a point in M (Theorem 9) when the step size is sufficiently small.

We propose three algorithms to numerically find solution manifolds and use them to handle statistical problems:

1. **Monte Carlo gradient descent algorithm:** an algorithm generating points over M that requires only the access to Ψ and its gradient (Section 3 and Algorithm 1).
2. **Manifold-constraint maximizing algorithm:** an algorithm that finds the MLE on the solution manifold (Section 4.3.1 and Algorithm 2).
3. **Approximated manifold posterior algorithm:** a Bayesian procedure that approximates the posterior distribution on a manifold (Appendix A

and Algorithm 3).

We would like to emphasize that while some of the theoretical results have appeared in the regular case ($s = d$, i.e., the solution manifold is a collection of points or just a single point), generalizing these results to the manifold cases ($s < d$) requires non-trivial extensions of existing techniques. The major challenge comes from the fact that the set M contains an infinite number of points and the geometry of M will pose technical issues during the theoretical analysis. Also, although the stability theorem has appears for specific examples such as level set and ridge estimation [12, 66, 35], there is no such result for the general class of solution manifolds. This paper provides a unified framework of analyzing solution manifolds.

The impact of this paper is beyond statistics. Our result provides a new analysis of the partial identification problem in econometrics [38, 25]. The local center manifold theorem offers a new class of statistical problems where the dynamical system interacts with statistics [61]. The Monte Carlo approximation of a solution manifold leads to a point cloud over the manifold, which is a common scenario in computational geometry [24, 23, 30]. The algorithmic convergence of the gradient descent demonstrates a new class of non-convex functions for that we still obtain the linear convergence [10, 58].

Outline. The remainder of this paper is organized as follows. Section 2 provides a formal definition of a solution manifold and studies the smoothness and stability of the manifold. Section 3 presents an algorithm for approximating the solution manifold and an analysis of its properties. Section 4 discusses several statistical applications of solution manifolds. Section 5 provides future directions, connections with other fields, and some manifolds in statistics that are not in a solution form.

Notations. Let $v \in \mathbb{R}^d$ be a vector and $V \in \mathbb{R}^{n \times m}$ be a matrix. $\|v\|_2$ is the L_2 norm (Euclidean norm) of v , and $\|v\|_{\max} = \max\{|v_1|, \dots, |v_d|\}$ is the vector max norm. For matrices, we use $\|V\| = \|V\|_2 = \max_{\|u\|=1, u \in \mathbb{R}^m} \frac{\|Vu\|_2}{\|u\|_2}$ as the L_2 norm and $\|V\|_{\max} = \max_{i,j} \|V_{ij}\|$ as the max norm. For a squared matrix A , we define $\lambda_{\min}(A)$, $\lambda_{\max}(A)$ to be the minimal and maximal eigenvalue respectively, and $\lambda_{\min, >0}^2(A)$ as the smallest non-zero eigenvalue. For a vector value function Ψ , we define a maximal norm of derivatives as

$$\|\Psi\|_{\infty}^{(J)} = \sup_x \max_i \max_{j_1} \cdots \max_{j_J} \left| \frac{\partial^J}{\partial x_{j_1} \cdots \partial x_{j_J}} \Psi_i(x) \right|$$

for $J = 0, 1, 2, 3$. When Ψ is a scalar function, this is reduced to

$$\|\Psi\|_{\infty}^{(1)} = \sup_x \|\nabla \Psi(x)\|_{\max}, \quad \|\Psi\|_{\infty}^{(2)} = \sup_x \|\nabla \nabla \Psi(x)\|_{\max},$$

which are the usual maximal norm of the gradient vector and Hessian matrix over all x . We also define

$$\|\Psi\|_{\infty, J}^* = \max_{j=0, \dots, J} \|\Psi\|_{\infty}^{(j)}$$

as a norm that measures distance using upto the J -th derivative. The Jacobian (gradient) of $\Psi(x)$ is an $s \times d$ matrix

$$G_{\Psi}(x) = \nabla\Psi(x) = \begin{bmatrix} \nabla\Psi_1(x)^T \\ \nabla\Psi_2(x)^T \\ \dots \\ \nabla\Psi_s(x)^T \end{bmatrix} \in \mathbb{R}^{s \times d}$$

and the Hessian of $\Psi(x)$ will be an $s \times d \times d$ array

$$H_{\Psi}(x) = \nabla\nabla\Psi(x) \in \mathbb{R}^{s \times d \times d}, \quad [H_{\Psi}(x)]_{ijk} = \frac{\partial^2}{\partial x_j \partial x_k} \Psi_i(x)$$

and third derivative of Ψ will be an array

$$\nabla\nabla\nabla\Psi(x) \in \mathbb{R}^{s \times d \times d \times d}, \quad [\nabla\nabla\nabla\Psi(x)]_{ijkl} = \frac{\partial^3}{\partial x_j \partial x_k \partial x_l} \Psi_i(x).$$

Let A be a set and x be a point. We then define

$$d(x, A) = \inf\{\|x - y\| : y \in A\}$$

as the projected distance from x to A . For a set A and a positive number r , we denote $A \oplus r = \{x : d(x, A) \leq r\}$.

2. Solution manifold and its geometry

Let $\Psi : \mathbb{R}^d \mapsto \mathbb{R}^s$ be a vector-valued function and $M = \{x : \Psi(x) = 0\}$ be the solution set/manifold. When the Jacobian matrix $G_{\Psi}(x) = \nabla\Psi(x)$ has rank s at every $x \in M$, the set M is an $(d - s)$ -dimensional manifold locally at every point x due to the implicit function theorem [72]. In algebraic statistics, the parameters in the solution set $\{x : \Psi(x) = 0\}$ are called an implicit (statistical) algebraic model [36].

For a solution manifold, its normal space can be characterized using the following lemma.

Lemma 1. *For every point $x \in M$, the row space of $G_{\Psi}(x) \in \mathbb{R}^{s \times d}$ spans the normal space of M at x .*

Lemma 1 is an elementary result from geometry (see Section 6.5.1 of [34]); hence, we omit its proof. Lemma 1 states that the Jacobian/gradient of Ψ is normal to the solution manifold. This is a natural result because the gradient of a function is always normal to the level set and the solution manifold can be viewed as the intersection of level sets of different functions.

2.1. Assumptions

We will make the following two major assumptions in this paper. All the theoretical results rely on these two assumptions.

(D-k) $\Psi(x)$ is bounded k -times differentiable.

(F) There exists $\lambda_0, \delta_0, c_0 > 0$, such that

A. $\lambda_{\min}(G_{\Psi}(x)G_{\Psi}(x)^T) \equiv \lambda_{\min, > 0}(G_{\Psi}(x)^T G_{\Psi}(x)) > \lambda_0^2$ for all $x \in M \oplus \delta_0$ and

B. $\|\Psi(x)\|_{\max} > c_0$ for all $x \notin M \oplus \delta_0$.

Assumption (D-k) is an ordinary smoothness of the generator function. It may be relaxed by a Hölder type condition on $\Psi(x)$. Assumption (D-k) is weaker for a smaller integer k . In the stability analysis, we need (D-1) (i.e., bounded first order derivative) condition. If we want to control the smoothness of the manifold, we need (D-2) or a higher-order condition. The stability of a gradient flow requires a (D-3) condition. Assumption (F) is a curvature assumption of Ψ around the solution manifold. Lemma 1 implies that the normal space of M at each point is well-defined. This assumption will reduce to commonly assumed conditions in the literature. For instance, in the case of mode estimation (finding the local modes of a PDF $p(x)$), (F) reduces to the assumption that the local modes are well-defined as separated [68, 69]. This is similar to the assumption that the PDF $p(x)$ is a Morse function [19, 43]. In the MLE theory, (F) refers to the Fisher's information matrix being positive definite at the MLE (and other local maxima), which is often viewed as a classical condition in the MLE theory (see, e.g., Chapter 5 of [77]). In the problem of finding the density level sets (finding the set $\{x : p(x) = \lambda\}$), this assumption is equivalent to assuming that $p(x)$ has a non-zero gradient around the level set [56, 76, 57, 12, 48, 44]. The assumption (F) is critical to the that the set M forms a manifold. When there is no lower bound λ_0 , i.e., the gradient $\lambda_{\min}(G_{\Psi}(x)G_{\Psi}(x)^T)$ attains 0, the set M may not form a manifold. One scenario that this occurs is the density level set $\{x : p(x) = \lambda\}$ such that the density function $p(x)$ is flat at the level λ .

The constants in (F) can be further characterized by the following lemma.

Lemma 2. *Assume (D-2) and that there exists $\lambda_M > 0$ such that*

$$(F') \quad \inf_{x \in M} \lambda_{\min}(G_{\Psi}(x)G_{\Psi}(x)^T) \geq \lambda_M^2.$$

Then the constants in (F) can be chosen as

$$\lambda_0 = \frac{1}{2}\lambda_M, \quad \delta_0 = \frac{3\lambda_M^2}{8\|\Psi_{\infty,1}^*\|\|\Psi_{\infty,2}^*\|}, \quad c_0 = \inf_{x \notin M \oplus \delta_0} \|\Psi\|_{\max}.$$

Lemma 2 only places assumptions on the behavior of Ψ and its derivatives on the manifold M . (F') is the eigengap conditions for the row space of $G_{\Psi}(x)$. The assumption in Lemma 2 is very mild. When estimating the local modes of a function, the assumption is the same as requiring that the Hessian matrix at local modes have all eigenvalues being positive. The requirement $\inf_{x \in M} \lambda_{\min}(G_{\Psi}(x)G_{\Psi}(x)^T) \geq \lambda_M^2$ implies that rows of $G_{\Psi}(x)$ are linearly independent for all x .




FIG 2. An illustration for reach. Reach is the largest radius for a ball that can freely move along the manifold M without penetrating any part of M . In (a), the radius of the pink ball is the same to the reach. In (b), the radius is too large; hence, it cannot pass the small gap on M .

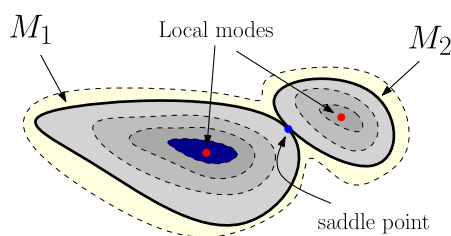


FIG 3. An example of a smooth generator Ψ but the resulting solution manifold $M = M_1 \cup M_2$ may have 0 reach. The dashed line is the contour lines for different levels.

2.2. Smoothness of solution manifolds

We introduce the concept of *reach* [33, 28, 1, 2] (also known as condition number in [59] and minimal feature size in [15]) to describe the smoothness of a manifold. The reach is the longest distance away from M , in which every point within this distance to M has a unique projection onto M , that is,

$$\text{reach}(M) = \sup\{r > 0 : \forall x \in M \oplus r, x \text{ has a unique projection onto } M\}. \quad (5)$$

The reach can be viewed as the largest radius of a ball that can freely move along the manifold M . See figure 2 for an example. The reach has been used in nonparametric set estimation as a condition to guarantee the stability of a set estimator [18, 20, 28].

The smoothness of Ψ does not suffice to guarantee the smoothness of a solution manifold. Consider the example of a density level set $\{x : p(x) = \lambda\}$ with a smooth density $p(x)$. By construction, this level set is a solution manifold with $\Psi(x) = p(x) - \lambda$, a smooth function. Suppose that $p(x)$ has two modes and a saddle point c . If we choose $\lambda = p(c)$, the level set does not have a positive reach. See Figure 3 for an example.

Although the smoothness of Ψ is not enough to guarantee a smooth M , with an additional condition (F), the solution manifold will be a smooth one as will be described in the following theorem.

Theorem 3 (Smoothness Theorem). *Conditions (D-2) and (F) imply that the reach of M has a lower bound*

$$\text{reach}(M) \geq \min \left\{ \frac{\delta_0}{2}, \frac{\lambda_0}{\|\Psi\|_{\infty,2}^*} \right\}.$$

Theorem 3 shows a lower bound on the reach of a solution manifold. Essentially, it shows that as long as the generator function is not flat around the solution manifold (assumption (F)), the resulting manifold will be smooth. Note that for Theorem 3, condition (D-2) can be relaxed to a 2-Hölder condition and the quantity $\|\Psi\|_{\infty,2}^*$ can be replaced by the corresponding Hölder's constant.

Remark 1. *Reach is related to the curvature and a quantity called folding [65]. Folding is defined as the smallest radius r such that $B(x,r) \cap M$ is connected for each $x \in M$. The first quantity $\frac{\delta_0}{2}$ is related to the folding. The second quantity $\frac{\lambda_0}{K}$ is related to the curvature. When $s = 1$, a similar result to Theorem 3 appeared in [79]. Note that the reach is also related to the ‘rolling properties’ and ‘ α -convexity’; see [28] and Appendix A of [60].*

2.3. Stability of solution manifolds

In this section, we show that when two generator functions are sufficiently close, the associated solution manifolds will be similar as well. We first start with a partial stability theorem, which holds under a weaker smoothness condition.

Proposition 4. *Let $\Psi, \tilde{\Psi} : \mathbb{R}^d \mapsto \mathbb{R}^s$ be two vector-valued functions and let*

$$M = \{x : \Psi(x) = 0\}, \quad \tilde{M} = \{x : \tilde{\Psi}(x) = 0\}$$

be the corresponding solution manifolds. Moreover, let δ_0 and c_0 be the constants in (F). Assume that (D-1) holds within $M \oplus \delta_0$ and (F). When $\|\tilde{\Psi} - \Psi\|_{\infty,0}^ < c_0$, we have*

$$\sup_{y \in \tilde{M}} d(y, M) = O \left(\|\tilde{\Psi} - \Psi\|_{\infty}^{(0)} \right).$$

In the case of $s = d$, the conditions in Proposition 4 show many connections to the existing work. For instance, the convergence rate of estimating a mode (or a local mode) is often based on a similar condition to (D-1); see, e.g., Theorem 2 of [78]. In the MLE and Z-estimation theory (see, e.g., Section 5.6 and 5.7 of [77]), classical conditions often require the first order derivative of the score equation or estimating equations to be uniformly bounded, which are similar to conditions of (D-1). Moreover, in the MLE theory, we often need the derivative of the score function (or Fisher's information matrix under appropriate conditions) to be non-singular within a small neighborhood of the population maximizer. This is exactly the same as condition (F). The constant λ_0 in (F) is the smallest absolute eigenvalue of the derivative of the score function in the case of MLE problems.

Note that in the case of $s = d$, Proposition 4 is often enough for statistical consistency because M is a collection of disjoint points; thus, there is no need to consider the smoothness of M . However, when $s < d$, the set M contains infinite amount of points and the smoothness of M plays a role in terms of analyzing its stability. Therefore, we will need additional derivatives.

Before we formally state the stability theorem, we first introduce the concept of the Hausdorff distance, a common metric of sets. The *Hausdorff distance* is defined as

$$\text{Haus}(A, B) = \max \left\{ \sup_{x \in B} d(x, A), \sup_{x \in A} d(x, B) \right\}.$$

The Hausdorff distance is a distance between two sets and can be viewed as an L_∞ type distance between sets.

Theorem 5 (Stability Theorem). *Let $\Psi, \tilde{\Psi} : \mathbb{R}^d \mapsto \mathbb{R}^s$ be two vector-valued functions and let*

$$M = \{x : \Psi(x) = 0\}, \quad \tilde{M} = \{x : \tilde{\Psi}(x) = 0\}$$

be the corresponding solution manifolds. Assume (D-2) and (F) and that $\tilde{\Psi}$ is bounded two-times differentiable. When $\|\tilde{\Psi} - \Psi\|_{\infty,2}^$ is sufficiently small,*

1. (F) holds for $\tilde{\Psi}$.
2. $\text{Haus}(M, \tilde{M}) = O\left(\|\tilde{\Psi} - \Psi\|_{\infty}^{(0)}\right)$.
3. $\text{reach}(\tilde{M}) \geq \min\left\{\frac{\delta_0}{2}, \frac{\lambda_0}{\|\Psi\|_{\infty,2}^*}\right\} + O\left(\|\tilde{\Psi} - \Psi\|_{\infty,2}^*\right)$.

Theorem 5 shows that two similar generator functions have similar solution manifolds. Claim 2 is a geometric convergence property indicating that a consistent generator function estimator implies a consistent manifold estimator. The need of a second-order derivative comes from the constants in (F). These constants are associated with the second-order derivatives through Lemma 2. Claim 3 is the convergence in smoothness, which implies that \tilde{M} cannot be very wiggly when $\tilde{\Psi}$ is sufficiently closed to Ψ .

Example 5 (Missing data). *The stability theorem (Theorem 5) provides a simple approach for obtaining the convergence rate of an estimator. Consider the missing data example in Example 1. The ‘population’ solution manifold is the parameters $\theta = (\zeta_{x,y}, \mu_x, \xi)$*

$$\Theta = \{\theta : \Psi(\theta) = 0\} \subset \mathbb{R}^7$$

such that $\Psi(\theta) \in \mathbb{R}^6$ is based on equation (1). When we observed random samples of size n in the form of $(X_i, Y_i, R_i = 1)$ or $(X_i, R_i = 0)$, we derived estimators $\hat{P}(X = x, Y = y, R = 1)$ and $\hat{P}(X = x, R = 0)$ and constructed an estimated version $\hat{\Psi}_n(\theta)$ by replacing $P(x, y, r = 1)$ and $P(x, r = 0)$ in equation (1) with the estimated versions. This leads to an estimated solution manifold

$$\hat{\Theta}_n = \{\theta : \hat{\Psi}_n(\theta) = 0\} \subset \mathbb{R}^d.$$

Algorithm 1 Monte Carlo gradient descent algorithm

1. Randomly choose an initial point $x_0 \sim Q$, where Q is a distribution over the region of interest \mathbb{K} .

2. Iterates

$$x_{t+1} \leftarrow x_t - \gamma \nabla f(x_t) \tag{7}$$

until convergence. Let x_∞ be the convergent point.

3. If $\Psi(x_\infty) = 0$ (or sufficiently small), we keep x_∞ ; otherwise, discard x_∞ .

4. Repeat the above procedure until we obtain enough points for approximating M .

The stability theorem (Theorem 5) bounds the distance between $\hat{\Theta}_n$ and Θ via the difference

$$\max\{|\hat{P}(x, y, r = 1) - P(x, y, r = 1)|, |\hat{P}(x, r = 0) - P(x, r = 0)| : x, y = 0, 1\}.$$

3. Monte Carlo approximation to solution manifolds

Given Ψ or its estimator/approximation $\hat{\Psi}$, numerically finding the solution manifold M is a non-trivial task. In this section, we propose a simple gradient descent procedure to find a point on M . Note that though we describe the algorithm using Ψ , we will apply the algorithm to $\hat{\Psi}$ in practice. M is the solution set of Ψ ; thus, we may rewrite it as

$$\begin{aligned} M &= \{x : \Psi(x) = 0\} = \{x : f(x) = 0\}, \\ f(x) &= \Psi(x)^T \Lambda \Psi(x), \end{aligned} \tag{6}$$

where Λ is an $s \times s$ positive definite matrix. Let x be an initial point. Consider the gradient flow $\pi_x(t)$:

$$\pi_x(0) = x, \quad \pi_x'(t) = -\nabla f(\pi_x(t)).$$

Points in M are stationary points of the gradient system; moreover, they are the minima of the function $f(x)$. Thus, we can use a gradient descent approach to find points on M . Algorithm 1 summarizes the gradient descent procedure for approximating M . Note that we may choose $\Lambda = \mathbb{I}$ to be the identity matrix. In this case, $f(x) = \|\Psi(x)\|^2$, so we will be investigating the gradient descent flow of $\|\Psi(x)\|^2$. For the case of $d = s$, this is a common method in numerical analysis to find the solution set of non-linear equations; see, e.g., Section 6.5 of [29].

Algorithm 1 consists of three steps: a random initialization step, a gradient descent step, and a rejection step. The random initialization step allows us to explore different parts of the manifold. The gradient descent step moves the initial points to possible candidates on M by iterating the gradient descent. The rejection step ensures that points being kept are indeed on the solution manifold.

Note that the random initialization and rejection steps are popular strategies in numerical analysis. They serve as a remedy to resolve the problem of local

minimizers of f that is not in the solution manifold M . See the discussion in page 152 of [29].

Figure 1 shows an example of finding the solution manifold

$$\{(\mu, \sigma) : P(-5 < Y < 2) = 0.5, Y \sim N(\mu, \sigma^2)\}$$

using random initializations (from a uniform distribution over $[1, 3] \times [2, 4]$) and the gradient descent (we will provide more details on the implementations later in Example 6). We recover the underlying 1-dimensional manifold structure using Algorithm 1.

3.1. Analysis of the gradient flow

When an initial point is given, we perform gradient descent to find a minimum. We analyze this process by starting with an analysis of the (continuous-time) gradient flow $\pi_x(t)$. The gradient descent algorithm can be viewed as a discrete-time approximation to the continuous-time gradient flow. To analyze the convergence of the flow and the algorithm, we denote Λ_{\max} and Λ_{\min} as the largest and smallest eigenvalues of Λ , respectively.

Lemma 6. *Assume (D-2) and (F). Let $G_f(x) = \nabla f(x)$ and $H_f(x) = \nabla \nabla f(x)$ and $G_\psi(x) = \nabla \Psi(x) \in \mathbb{R}^{s \times d}$. Then we have the following properties:*

1. For each $x \in M$,
 - (a) the non-zero eigenvectors of $H_f(x)$ span the normal space of M at x .
 - (b) the minimal non-zero eigenvalue

$$\lambda_{\min, > 0}(H_f(x)) \geq \psi_{\min}^2(x) \equiv \lambda_{\min, > 0}^2(G_\psi(x)^T G_\psi(x)) \Lambda_{\min} \geq 2\lambda_0^2 \Lambda_{\min}.$$

- (c) the minimal eigenvalue in the normal space of M at x $\lambda_{\min, \perp}(H_f(x)) = \lambda_{\min, > 0}(H_f(x))$.

2. Suppose that x has a unique projection $x_M \in M$ and let $N_M(x)$ be the normal space of M at x_M . If $d(x, M) < \delta_c = \min \left\{ \delta_0, \frac{\Lambda_{\min}}{8d\Lambda_{\max}} \frac{\lambda_0^2}{\|\Psi\|_{\infty, 2}^* \|\Psi\|_{\infty, 3}^*} \right\}$, then

$$\lambda_{\min, \perp, M}(H_f(x)) \equiv \min_{v \in N_M(x_M)} \frac{v^T H_f(x) v}{\|v\|^2} = \min_{v \in N_M(x_M)} \frac{\|H_f(x) v\|}{\|v\|} \geq \lambda_0^2 \Lambda_{\min}.$$

Property 1 in Lemma 6 describes the behavior of the Hessian $H_f(x)$ on the manifold. The eigenspace (corresponds to non-zero eigenvalues) is the same as the normal space of the manifold. With this insight, it is easy to understand property 1-(c) that the minimal eigenvalue in the normal space is the same as the minimal non-zero eigenvalue. Property 2 is about the behavior of $H_f(x)$ around the manifold. The Hessian $H_f(x)$ is well-behaved as long as we are sufficiently close to M . The following theorem characterizes several important properties of the gradient flow.

Theorem 7 (Convergence of gradient flows). *Assume (D-3) and (F). Let δ_c be defined in Lemma 6. Define $\pi_x(\infty) = \lim_{t \rightarrow \infty} \pi_x(t)$. The gradient flow $\pi_x(t)$ satisfies the following properties:*

1. (Convergence radius) For all $x \in M \oplus \delta_c$, $\pi_x(\infty) \in M$.
2. (Terminal flow orientation) Let $v_x(t) = \frac{\pi'_x(t)}{\|\pi'_x(t)\|}$ be the orientation of the gradient flow. If $\pi_x(\infty) \in M$, then $v_x(\infty) = \lim_{t \rightarrow \infty} v_x(t) \perp M$ at $\pi_x(\infty)$.

The first result of Theorem 7 defines the convergence radius of the gradient flow. The flow converges to the manifold as long as the gradient flow starts within δ_c distance to the manifold. The second statement of the theorem characterizes the orientation of the gradient flow. The flow intersects the manifold from the normal space of the manifold. Namely, the flow hits the manifold orthogonally. If we choose $\Lambda = \mathbb{I}$ to be the identity matrix ($\Lambda_{\min} > 0$ in this case), Theorem 7 implies the convergence of the gradient flow of $\|\Psi\|^2$.

Note that Theorem 7 requires one additional derivative (D-3) because we need to perform a Taylor expansion of the Jacobian of Ψ around M to ensure that the gradient flow converges from a normal direction to a manifold. We need the third-order derivatives to ensure that the remainders are small.

Suppose that the initial point x is drawn from the distribution Q , which has a PDF q , the convergent point $\pi_x(\infty)$ can be viewed as a random draw from a distribution Q_M defined over the manifold M . The distribution Q and the distribution Q_M are associated via the mapping induced by the gradient descent process; thus, Q_M is a pushforward measure of Q [7]. We now investigate how Q and Q_M are associated.

For every point $z \in M$, we define its basin of attraction [13, 19] as

$$A(z) = \{x : \pi_x(\infty) = z\}.$$

Namely, $A(z)$ is the collection of initial points that the gradient flow converges to $z \in M$. Let $\mathcal{A}_M = \cup_{z \in M} A(z)$ be the union of all basins of attraction. The set \mathcal{A}_M characterizes the regions where the initialization leads to an accepted point in Algorithm 1. Thus, the acceptance probability of the rejection step of Algorithm 1 is $Q(\mathcal{A}_M) = \int_{\mathcal{A}_M} Q(dx)$.

Basin $A(z)$ has an interesting geometric property of forming an s -dimensional manifold under smoothness assumption. This result is similar to the stable manifold theorem in dynamical systems literature [53, 54, 5]. In fact, it is more relevant to the local center manifold theorem (see, e.g., Section 2.12 of [61]).

Theorem 8 (Local center manifold theorem). *Assume (D-3) and (F). The basin of attraction $A(z)$ forms an s -dimensional manifold at each $z \in M$.*

An outcome from Theorem 8 is that the pushforward measure Q_M has an s -dimensional Hausdorff density function [52, 62] if Q has a regular PDF q . Note that an s -dimensional Hausdorff density at point x is defined through

$$\lim_{r \rightarrow 0} \frac{Q_M(B(x, r))}{C_s r^s},$$

where C_s is the s -dimensional volume of a unit ball. If the limit of the above equation exists, Q_M has an s -dimensional Hausdorff density at point x .

Thus, if we obtain $Z_1, \dots, Z_N \in M$ from applying Algorithm 1, we may view them as IID observations from a distribution Q_M defined over the manifold M , and this distribution Q_M has an s -dimensional Hausdorff density function. The model that we observe IID Z_1, \dots, Z_N from a distribution supported over a lower-dimensional manifold is common in computational geometry [24, 30, 31, 16]. Hence, Theorem 8 implies that Algorithm 1 provides a new statistical example for this model.

Note that [3] proved the stability of a gradient ascent flow when the target is to find the local modes of density function. The stability of basins of attraction was studied in [21] in a similar scenario. These results may be generalized to solution manifolds. The major technical issue that we need to solve is that the convergent points of flows form a collection of infinite number of points. Therefore, the analysis is much more complicated. We leave this as a future work.

3.2. Analysis of the gradient descent algorithm

In Algorithm 1, we did not perform the gradient descent using the flow π_x ; instead, we used an iterative gradient descent approach that creates a sequence of discrete points $x_0, x_1, \dots, x_\infty$ such that

$$x_{t+1} = x_t - \gamma \nabla f(x_t), \quad x_0 = x, \tag{8}$$

where $\gamma > 0$ is the step size.

The gradient descent algorithm will diverge if the step size γ is chosen incorrectly. Thus, it is crucial to investigate the range of γ leading to a convergent point x_∞ and how fast the sequence $\{x_t : t = 0, 1, \dots\}$ converges to a point on M . The following theorem characterizes the algorithmic convergence along with a feasible range of the step size γ .

Theorem 9 (Linear convergence). *Assume (D-2) and (F). When the initial point x_0 and step size γ satisfy*

$$d(x_0, M) < \delta_c = \min \left\{ \delta_0, \frac{\Lambda_{\min} \lambda_0^2}{8d\Lambda_{\max} \|\Psi\|_{\infty,2}^* \|\Psi\|_{\infty,3}^*} \right\},$$

$$\gamma < \min \left\{ \frac{1}{\Lambda_{\max} \|\Psi\|_{\infty,2}^*}, \frac{\Lambda_{\max} \|\Psi\|_{\infty,2}^*}{4\lambda_0^4 \Lambda_{\min}^2}, \delta_c \right\},$$

we have the following properties for $t = 0, 1, 2, \dots$:

$$f(x_t) \leq f(x_0) \cdot \left(1 - \gamma \frac{\lambda_0^4 \Lambda_{\min}^2}{\Lambda_{\max} \|\Psi\|_{\infty,2}^*} \right)^t,$$

$$d(x_t, M) \leq d(x_0, M) \left(1 - \gamma \lambda_0^2 \Lambda_{\min} \right)^{\frac{t}{2}}.$$

The convergence radius δ_c is the same as in Theorem 7. Theorem 9 shows that when the initial point is within the convergence radius of the gradient flow and the step size is sufficiently small, the gradient descent algorithm converges linearly to a point on the manifold. An equivalent statement of Theorem 9 is that the algorithm takes only $O(\log(1/\epsilon))$ iterations to converge to ϵ -error to the minimum.

The key step in the derivation of Theorem 9 is to investigate the minimal eigenvalue of the normal space $\lambda_{\min,\perp}(H_f(x))$ for each $x \in M$. This quantity (appears in the theorem through the lower bound $\lambda_0^2 \Lambda_{\min}$) controls the flattest direction of $f(x)$ in the normal space. The three requirements on the step sizes are due to different reasons. The first requirement ($\gamma < \frac{1}{\Lambda_{\max} \|\Psi\|_{\infty,2}^*}$) ensures that the objective function is decreasing. The second requirement ($\gamma < \frac{\Lambda_{\max} \|\Psi\|_{\infty,2}^*}{\lambda_0^2 \Lambda_{\min}}$) establishes the convergence rate. The third requirement ($\gamma < \delta_c$) guarantees that the Hessian matrix behaves of f is well-behaved when applying the gradient descent algorithm. The first and third requirements together are enough for the convergence of the gradient descent algorithm but will not lead to the convergence rate. We need the additional second requirement to obtain the convergence rate.

Theorem 9 is a very interesting result. The function $f(x)$ is non-convex within any neighborhood of M (i.e., not locally convex), but the gradient descent algorithm still converges linearly to a stationary point. An intuitive explanation of this result is that the function $f(x)$ is a ‘directionally’ convex function in the normal subspace of M (Property 2 in Lemma 6). Note that similar to Theorem 7, Theorem 9 applies to the gradient descent algorithm with $\Lambda = \mathbb{I}$.

4. Statistical applications

In this section, we show that the idea of solution manifolds can be applied to various statistical problems. We also include a Bayesian approach that finds a credible region on a solution manifold in Appendix A.

4.1. Missing data

The solution manifold framework we developed can be used to analyze the nonparametric bound in the missing data problem [49, 50, 17]. We use Example 1 to illustrate the idea, but our analysis can be generalized to a complex missing data scenario. The nonparametric bound refers to the feasible range of parameters $\theta = (\zeta_{x,y}, \mu_x, \xi : x, y = 0, 1) \in \mathbb{R}^7$ without any additional assumptions. Hence, the only constraint for these seven parameters is the six equations in equation (1). Thus, we know that the resulting parameter space is a one-dimensional manifold.

This manifold will be a smooth one due to Theorem 3. The stability theorem informs us that when we estimate the constraints by sample analogues, the estimated manifold (can be viewed as an estimated nonparametric bound) will

be at $O_P(1/\sqrt{n})$ distance to the population manifold by the stability theorem (Theorem 5).

Algorithm 1 provides a simple approach for numerically finding points on this solution manifold. We can obtain a point cloud approximation of the 1D manifold characterizing the nonparametric bound of all the parameters with multiple random initializations.

4.2. Algorithmic fairness

We now revisit the algorithmic fairness problem in Example 2. We have shown that a simple method of generating a test fair classifier Q from W, A is to sample from a Bernoulli random variable with a parameter $q_{W,A} = P(Q = 1|W, A)$ that satisfies equation (3). This leads to a solution manifold $\Theta = \{\theta = (q_{w,a} : w, a = 0, 1) : \Psi(\theta) = 0\}$, where $\Psi(\theta)$ is described in equation (3). By Theorem 3, the collection of feasible parameters will be a smooth manifold. When we estimate the underlying constraint by a random sample, the convergence rate (of manifolds) is described by Theorem 5.

In practice, finding Θ is often not the ultimate goal. Our goal is to find a classifier that is test fair and has a good classification error [40, 27]. A conventional approach of measuring classification accuracy is via a loss function $L(y, y')$ and we want to find the optimal $q_{w,a}^* \in \Theta$ such that

$$q^* = \operatorname{argmin}_{q \in \Theta} R(q) = \mathbb{E}_{Q \sim q}(L(Y, Q)).$$

This is essentially a manifold constraint maximization/minimization problem. This problem also occurs in the constraint likelihood inference (see next section) where we want to find the MLE under a solution manifold constraint. We will discuss a unified treatment of this manifold-constraint optimization problem in the next section and propose an algorithm for it (Algorithm 2). While the algorithm is written in the form of finding the MLE, one can easily adapt it to the test fairness problem.

4.3. Parametric model

One scenario that the solution manifolds will be useful is analyzing parametric models. We provide two different examples showing how solution manifolds can be used in parametric modeling. Suppose that we observe IID observations X_1, \dots, X_n from some distribution P , and we model the distribution using a parametric model P_θ and $\theta \in \Theta$. Let p_θ be the PDF/PMF of P_θ and let

$$\ell(\theta|X_1, \dots, X_n) = \frac{1}{n} \sum_{i=1}^n \log p_\theta(X_i)$$

be the log-likelihood function. Note that in Appendix A, we also provide a Bayesian procedure that approximates the posterior distribution on a manifold (Algorithm 3).

4.3.1. Constrained MLE

In the likelihood inference, we may need to compute the MLE under some constraints. One example is the likelihood ratio test when the parametric space Θ_0 under H_0 is generated by equality constraints. Namely,

$$\Theta_0 = \{\theta \in \Theta : \Psi(\theta) = 0\}.$$

This problem occurs in algebraic statistic and asymptotic theories can be found in [55] and Section 5 of [32].

To use the likelihood ratio test, we need to find the MLEs under both Θ_0 and Θ . Finding the MLE under Θ is a regular statistical problem. However, finding the MLE under Θ_0 may not be easy because of the constraint $\Psi(\theta) = 0$. We propose to a procedure combining the gradient ascent of the likelihood function and the gradient descent to the manifold to compute the constrained MLE. Algorithm 2 describes the procedure, and Figure 4 provides a graphical illustration. This algorithm consists of a one-step gradient ascent of the likelihood function (Step 3) and a gradient descent to manifold (Algorithm 1; steps 4 and 5).

The stopping criterion (Step 6) is that $\nabla \ell(\theta_\infty^{(m)} | X_1, \dots, X_n)$ belongs to the row space of $\nabla \Psi(\theta_\infty^{(m)})$. Due to Lemma 1, the row space of $\nabla \Psi(\theta_\infty^{(m)})$ is the normal space of M at $\theta_\infty^{(m)}$. It is easy to see that any critical points of the log-likelihood function on the manifold satisfy the condition that the likelihood gradient belongs to the row space of $\nabla \Psi$; thus, the constrained MLE is a stationary point in Algorithm 2. As a result, we stop the algorithm when the stopping criterion occurs. However, other local modes and saddle points and local minima are also the stationary points. Hence, in practice, we need to run the algorithm with multiple initial points to increase the chance of finding the true MLE.

Note that one may replace the gradient ascent process by the EM algorithm. However, the EM algorithm is not identical to a gradient ascent, so it is unclear if the movement $\theta_0^{(m+1)} - \theta_\infty^{(m)}$ will be normal to the manifold Θ_0 or not.

Example 6 (Testing a tail probability in a Gaussian model). *To illustrate the idea, suppose that $X_i \in \mathbb{R}$ is from an unknown distribution that we place a parametric model on it. We further assume that the parametric distribution is a Gaussian $N(\mu, \sigma^2)$ with unknown mean and variance. Consider the null hypothesis*

$$H_0 : P(r_0 \leq Y \leq r_1) = s_0$$

for some given $s_0 > 0$ and r_0, r_1 (note that this example also appears in Figure 1). Let $\Phi(y) = P(Z \leq y)$ denote the CDF of a standard normal. The null hypothesis H_0 forms the following constraint on (μ, σ^2) :

$$s_0 = \Phi\left(\frac{r_1 - \mu}{\sigma}\right) - \Phi\left(\frac{r_0 - \mu}{\sigma}\right).$$

Thus,

$$\Psi(\mu, \sigma) = \Phi\left(\frac{r_1 - \mu}{\sigma}\right) - \Phi\left(\frac{r_0 - \mu}{\sigma}\right) - s_0 \in \mathbb{R}.$$

Algorithm 2 Manifold-constraint maximizing algorithm

1. Randomly choose an initial point $\theta_0^{(0)} = \theta_\infty^{(0)} \in \Theta$.
2. For $m = 1, 2, \dots$, do step 3-6:
3. **Ascent of likelihood.** Update

$$\theta_0^{(m)} = \theta_\infty^{(m-1)} + \alpha \nabla \ell(\theta_\infty^{(m-1)} | X_1, \dots, X_n), \quad (9)$$

where $\alpha > 0$ is the step size of the gradient ascent over likelihood function and $\ell(\theta | X_1, \dots, X_n)$ is the log-likelihood function.

4. **Descent to manifold.** For each $t = 0, 1, 2, \dots$ iterates

$$\theta_{t+1}^{(m)} \leftarrow \theta_t^{(m)} - \gamma \nabla f(\theta_t^{(m)})$$

until convergence. Let $\theta_\infty^{(m)}$ be the convergent point.

5. If $\Psi(\theta_\infty^{(m)}) = 0$ (or sufficiently small), we keep $\theta_\infty^{(m)}$; otherwise, discard $\theta_\infty^{(m)}$ and return to step 1.
6. If $\nabla \ell(\theta_\infty^{(m)} | X_1, \dots, X_n)$ belongs to the row space of $\nabla \Psi(\theta_\infty^{(m)})$, we stop and output $\theta_\infty^{(m)}$.

The feasible set of (μ, σ^2) forms a 1D solution manifold in \mathbb{R}^2 . It is difficult to find the analytical form of the MLE under H_0 , but we may use the method in Algorithm 2 to obtain a numerical approximation. The derivative of $\Psi(\mu, \sigma)$ with respect to μ and σ has the following closed-form:

$$\begin{aligned} \frac{\partial}{\partial \mu} \Psi(\mu, \sigma) &= -\frac{1}{\sigma} \phi\left(\frac{r_1 - \mu}{\sigma}\right) + \frac{1}{\sigma} \phi\left(\frac{r_0 - \mu}{\sigma}\right) \\ \frac{\partial}{\partial \sigma} \Psi(\mu, \sigma) &= -\frac{r_1 - \mu}{\sigma^2} \phi\left(\frac{r_1 - \mu}{\sigma}\right) + \frac{r_0 - \mu}{\sigma^2} \phi\left(\frac{r_0 - \mu}{\sigma}\right), \end{aligned}$$

where $\phi(y) = \frac{1}{\sqrt{2\pi}} e^{-y^2/2}$ is the PDF of the standard normal. Algorithm 2 can easily be implemented with the above derivatives. Figure 4 shows an example of applying Algorithm 2 to this example with $r_1 = -5, r_2 = 2$ and $s_0 = 0.5$ (and 1000 random numbers generated from $N(1.5, 3^2)$). All five random initial points converge to the maximum on the manifold.

4.3.2. Partial identification and generalized method of moments

The solution manifolds appear in the partial identification problem [51] in Econometrics. One example is the moment constraint problem [25], also known as the generalized method of moments [38, 39]. In this case, we want to estimate parameter $\theta \in \mathbb{R}^d$ that solves the moment equation $\mathbb{E}(g(Y; \theta)) = 0$, where $g(y; \theta) \in \mathbb{R}^s$ is a vector-valued function, and X is a random variable denoting the observed data. When $s < d$, the solution set (also called an identified set in [25]) $M = \{\theta : \mathbb{E}(f(Y; \theta)) = 0\}$ forms a solution manifold.

Thus, the smoothness theorem (Theorem 3) and the stability theorem (Theorem 5) can be applied to this case. When the estimator is obtained by the empirical moment equation $\widehat{M}_n = \{\theta : \frac{1}{n} \sum_{i=1}^n f(Y_i; \theta) = 0\}$, Theorem 5 implies $\text{Haus}(\widehat{M}_n, M) \xrightarrow{P} 0$ when the empirical moments $\frac{1}{n} \sum_{i=1}^n f(Y_i; \theta)$ and its

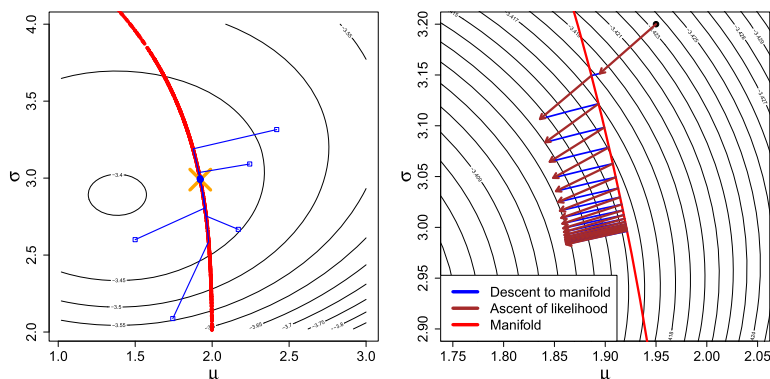


FIG 4. An example illustrating how Algorithm 2 works. We consider the example of estimating the tail probability in a Gaussian model $N(\mu, \sigma^2)$ with the constraint $P(-5 \leq X \leq 2) = 0.5$. We generate $n = 1000$ points from $N(1.5, 3^2)$ and display the log-likelihood function in the two panels (contours are the log-likelihood surface). **Left:** We initialize Algorithm 2 with five random points (blue boxes) and the algorithm creates an ascending path (blue lines) to the maximum point (orange cross). **Right:** We illustrate Algorithm 2 by showing points in each iteration in the algorithm in a zoom-in area relative to the left panel. Starting from the solid black point, we first perform a gradient ascent with respect to the log-likelihood function (brown arrow) then apply Algorithm 1 to descent to the solution manifold. We keep repeating this process until it converges.

derivatives with respect to θ converge to the population moments $\mathbb{E}(g(X; \theta))$ and its derivatives, respectively.

In generalized method of moments, a common approach of finding a solution to $\mathbb{E}(g(Y; \theta)) = 0$ is by minimizing a criterion function $Q(\theta) = \mathbb{E}(g(Y; \theta))^T \Lambda \mathbb{E}(g(Y; \theta))$, where Λ is a positive definite matrix [39]. This is identical to the function f defined in equation (6). Thus, the analysis in Section 3 can be used to study the minimization problem in the generalized method of moments.

Remark 2. In econometrics, a similar problem to the solution manifold is the inequality constraint problem, which occurs when we replace the equality constraints with inequality constraints [75, 70], i.e., $\mathbb{E}(g_\ell(Y; \theta)) \leq 0$ for $\ell = 1, \dots, s$. The goal is to find θ satisfying the above inequality constraint. A common approach to finding the feasible set is by defining an objective function

$$Q(\theta) = \left| \sum_{\ell=1}^s [\mathbb{E}(g_\ell(Y; \theta))]_+ \right|^2, \quad [y]_+ = \max(y, 0)$$

such that the feasible set is $\{\theta : Q(\theta) = 0\}$. The inequality constraint implies that $\{\theta : Q(\theta) = 0\}$ may not form a lower-dimensional manifold but a subset of the original parameter space.

A common estimator of $\{\theta : Q(\theta) = 0\}$ is

$$\left\{ \theta : \widehat{Q}_n(\theta) \leq c_n \right\}, \quad \widehat{Q}_n(\theta) = \left| \sum_{\ell=1}^s \left[\frac{1}{n} \sum_{i=1}^n g_\ell(Y_i; \theta) \right]_+ \right|^2$$




FIG 5. An example of approximating a density level set. This is the Graft-versus-Host Disease (GvHD) data that we borrowed from `mclust` package in R. We use the control group and choose three variables: CD3, CD4, and CD8b. We apply a Gaussian kernel density estimator with bandwidth $h = 20$ (same in all coordinates). The level set of interest is the density level corresponding to 25% quantile of densities at all observations. The three panels display the level set under three different angles.

for some sequence $c_n \rightarrow 0$. Note that by properly choosing c_n , we may construct both an estimator and a confidence region; see [25] and [70] for more details.

4.4. Nonparametric set estimation

Solution manifolds occur in many scenarios of nonparametric set estimation problems. One famous example is the density level set problem in which the parameter of interest is the (density) level set $\{x : p(x) = \lambda\}$, and p is the PDF that generates our data, and λ is a pre-specified level. In this case, the smoothness theorem (Theorem 3) yields the same result as [20]. Moreover, the stability theorem (Theorem 5) suggests that the convergence rate under the Hausdorff distance will be the rate of estimating the density function, which is consistent with several existing works [12, 67, 66].

The methods developed in Section 3 can be used to find points on the level set. As an illustration, Figure 5 shows an example of approximating the level set using Algorithm 1 with the Graft-versus-Host Disease (GvHD) data [11] from `mclust` package in R [6]. We use variables CD3, CD4, and CD8b and focus on the control group. The density is computed using a Gaussian kernel density estimator with an equal bandwidth $h = 20$ in all coordinates. We choose the level as the 25% quantile of all observations' densities. The three panels display the approximated level sets from three angles. The three surfaces in the data indicate three connected components in the regions where the density is above this threshold.

In addition to the level set problem, the density ridges [18, 35] are also examples of solution manifolds. The density k -ridges are the collection $\{x : V_k(x)^T \nabla p(x) = 0, \lambda_{k+1}(x) < 0\}$, where $V_k(x) = [v_{k+1}(x), \dots, v_d(x)]$ denotes the matrix of $d - k$ eigenvectors of $\nabla \nabla p(x)$ corresponding to the smallest eigenvalues and $\lambda_k(x)$ is the k -th largest eigenvalue. If we only pick the lowest $d - k$ eigenvectors, we obtain a system of equations with $d - k$ equations, leading to a k -dimensional manifold (under smoothness conditions). The stability theorem

(Theorem 5) states that the convergence rate of a density ridge estimator will be at the rate of estimating the Hessian matrix, which is consistent with the findings of [35].

Note that the density local modes can be viewed as 0-dimensional ridges. In this special case, the matrix $V_1(x)$ is full rank under assumption (F); thus,

$$\{x : V_0(x)^T \nabla p(x) = 0, \lambda_1(x) < 0\} = \{x : \nabla p(x) = 0, \lambda_1(x) < 0\}.$$

As a result, the function Ψ that generates the local modes does not involve the Hessian matrix, leading to the convergence rate of the mode estimator to be the same as the gradient estimation rate rather than the rate of estimating the Hessian.

5. Discussion

In this paper, we investigate both geometric and algorithmic properties of solution manifolds. While the solution manifolds may seem to be abstract, we showed that they appear in various statistical problems including missing data, algorithmic fairness, likelihood inference, and nonparametric set estimation. Hence, the methodologies and theories developed in this paper provide a generic framework for analyzing all these problems. This framework may inform us of the hidden relation among all these seemingly different statistical problems. In what follows, we discuss some relevant topics to the solution manifolds.

5.1. Smoothness, stability, and convergence of gradient flow

We developed 5 major theoretical results: smoothness theorem (Theorem 3), stability theorem (Theorem 5), gradient flow theorem (Theorem 7), local center manifold theorem (Theorem 8), and algorithmic convergence theorem (Theorem 9). These results characterize different properties of solution manifolds and are often studied in various fields. In our work, we showed that they all rely on a similar set of assumptions: (D), the smoothness of Ψ , and (F), the curvature assumption of Ψ around M .

Assumption (D) is more than enough for some theoretical results. The smoothness theorem can be relaxed to assuming that Ψ satisfies β -Hölder condition with various β . For instance, the (D-1) condition in the weak stability result (Proposition 4) can be relaxed to the 1-Hölder condition. The condition (D-2) in the smoothness theorem (Theorem 3) can be replaced by the 2-Hölder condition.

Moreover, we observe a hierarchy of smoothness corresponding to different theoretical results. If we only have (D-1), the first order derivatives, we have a weak stability result from Proposition 4. If we make a further assumption (D-2), we have a stability theorem (Theorem 5), a characterization of smoothness (Theorem 3), and an algorithmic convergence (Theorem 9). Under an additional assumption (D-3), we can derive an even stronger result of the corresponding gradient flow (Theorem 7 and 8)

5.2. Connections to other fields

We would like to point out that the results of this paper have several connections to other fields.

- **Econometrics.** Solution manifolds occur in the partial identification problem (Section 4.3.2); hence, our analysis provides some insights into the moment equality constraint problem [25]. Our analysis on the gradient descent (e.g., Theorem 7) can be applied to investigate the property of the minimization problem in the generalized method of moments approach [38, 39].
- **Dynamical systems.** As mentioned before, Theorem 8 is related to the stable manifold theorem and the local center manifold theorem in dynamical systems [53, 54, 5, 61]. Our analysis provides statistical examples that these theorems may be useful in data analysis.
- **Computational geometry.** If we stop the gradient descent process early, we do not obtain points that are on the manifold. The resulting points Z_1, \dots, Z_n may be viewed as from $Z_i = X_i + \epsilon_i$, where $X_i \in M$ is from a distribution over the manifold and ϵ_i is some additive noise. This model is a common additive noise model in the computational geometry literature [24, 23, 30, 31, 16, 8]. Our proposed method provides another concrete example of the manifold additive noise model.
- **Optimization.** In general, a gradient descent method has a linear convergence rate when the objective function is strongly convex and has a smooth gradient [10, 58]. However, in our setting, the objective function $f(x)$ is non-convex (and is not locally convex), but the gradient descent algorithm still obtains a linear (algorithmic) convergence rate (Theorem 7). This reveals a class of non-convex objective functions that can be minimized quickly using a gradient descent algorithm.

5.3. Future work

The framework developed in this paper has many potentials in other problems. We provide some possible directions that we plan to pursue in the future.

- **Log-linear model.** The log-linear model of categorical variables is an interesting example in the sense that it can be expressed as a solution manifold when there are constraints like conditional independence but it may be unnecessary to use the developed techniques. Consider a d -dimensional categorical random vector X that takes values in $\{0, 1, 2, \dots, J-1\}^d$. The joint PMF of X $p(x_1, \dots, x_d)$ has J^d entries with the constraint that $\sum_x p(x_1, \dots, x_d) = 1$, so it has $J^d - 1$ degrees of freedom. In the log-linear model, we reparametrize the PMF using the log-linear expansion: $\log p(x) = \sum_A \psi_A(x_A)$, where A is any non-empty subset of $\{1, 2, \dots, J\}$ and $x_A = (x_j : j \in A)$ with the constraint that $\psi_A(x_A) = 0$ if any $x_j = 0$ for $j \in A$. Under the log-linear model, we reparametrize the joint PMF

using the parameters $\Theta_{\text{LL}} = \{\psi_A(x_A) : A \subset \{1, 2, \dots, J\}, x_A \in \{0, 1\}^{|A|}\}$, where $|A|$ is the cardinality of A . The feasible parameters in Θ_{LL} forms a solution manifold due to the aforementioned constraints. However, common constraints in the log-linear model are that interaction terms $\psi_A = 0$ for some A . This leads to a flat manifold; hence, there is no need to use the developed technique. We may need to use techniques from the solution manifold when the constraint is placed on the PMF $p(x_1, \dots, x_d)$ rather than the log-linear models because the constraints on the PMF lead to an implicit constraint on Θ_{LL} . We leave this as future work.

- **Confidence regions of solution manifolds.** Another future direction is to develop a method for constructing the confidence regions of solution manifolds. There are two common approaches to construct a confidence region of a set. The first one is based on the “vertical uncertainty”, which is the uncertainty due to $\widehat{\Psi} - \Psi$. This idea has been applied in generalized method of moment problems [25, 70, 17] and level set estimations [48, 63, 22]. The other approach is based on the “horizontal uncertainty”, which is the uncertainty due to $\text{Haus}(\widehat{M}, M)$. This technique has been used in constructing confidence sets of density ridges and level sets [18, 20]. Based on these results, we believe that it is possible to develop a procedure for constructing the confidence regions of solution manifolds. We leave this as future work.
- **A new class of non-convex problems.** We observe an interesting phenomenon in Theorem 9. Although the objective function $f(x) = \Psi^T \Lambda \Psi(x)$ is non-convex, we still obtain a linear (algorithmic) convergence. Note that for a non-convex but locally convex around the minimizer, the linear convergence can be established via assuming a local strong convexity of the objective function [4], i.e., $f(x)$ is strongly convex within $B(x^*, r)$ for some radius $r > 0$ and x^* is the global minimizer. However, our problem is more complicated in the sense that $f(x)$ is flat along M , so it is not locally strongly convex. The key element in our result is assumption (F) stating that $f(x)$ behaves like being “locally strongly-convex” in the normal direction of M . Thus, with some additional structure on the non-convex function, we may still obtain a fast convergence. We will investigate how this may be useful in other non-convex optimization problems. In addition, the analysis may be applied to other forms of $f(x)$ that are not limited to a “squared”-type transformation of $\Psi(x)$ ($f(x)$ behaves like the square of Ψ), which may further improve the convergence rate. For instance, the gradient descent over $f_1(x) = \|\Psi(x)\|_1$ may also converge faster than over the function $f(x)$. We will investigate this in the future.

Acknowledgements

We thank the reviewers and M. Drton for insightful comments.

Appendix A: Bayesian inference

The techniques we developed for solution manifolds can be used for the Bayesian inference after some modifications. One example is the univariate Gaussian with unknown mean μ and variance σ^2 , and a second moment constraint. The parameter space is $\Theta(s_0) = \{(\mu, \sigma^2) : \mathbb{E}(Y^2) = \mu^2 + \sigma^2 = s_0^2\}$. We place a prior $\pi(\theta)$ over $\Theta(s_0)$ that reflects our prior belief about the parameter $\theta = (\mu, \sigma)$. However, how to sample from π (and the posterior) is a non-trivial task because π is supported on a manifold.

The Monte Carlo approximation method in Section 3 offers a solution to sampling from π . With a little modification of Algorithm 1, we can approximate the posterior distribution defined on the solution manifold. Let π be a prior PDF defined over the solution set $M = \{\theta : \Psi(\theta) = 0\}$ where $\Psi : \Theta \mapsto \mathbb{R}^k$. We observe IID observations X_1, \dots, X_n that are assumed to be from a parametric model $p(x|\theta)$. The posterior distribution of θ will be

$$\pi(\theta|X_1, \dots, X_n) \propto \begin{cases} \pi(\theta) \prod_{i=1}^n p(X_i|\theta), & \text{if } \theta \in M; \\ 0, & \text{if } \theta \notin M. \end{cases}$$

We propose a method that approximates the posterior distribution using a weighted point cloud. Our approach is formally described in Algorithm 3. Note that the algorithm we develop only requires the ability to evaluate a function $\rho(\theta) \propto \pi(\theta)$. We do not need the exact value of the prior density.

Example 7 (Bayesian analysis of Example 6). *Figure 6 shows an example of 90% credible intervals and the MAPs under three scenarios: prior distribution only (left panel), posterior distribution with $n = 100$ (middle panel), and posterior distribution with $n = 1000$ (right panel). This is the same setting in Example 6 and Figure 4, where the manifold is formed by the constraint $P(-5 < X < 2) = 0.5$ with $X \sim N(\mu, \sigma)$. We choose the prior distribution (density) as*

$$\pi(\mu, \sigma) \propto \phi(\mu; 2, 0.2)\phi(\sigma; 2.5, 0.2)I((\mu, \sigma) \in M),$$

where $\phi(x; a, b)$ is the density of $N(a, b^2)$. In the left panel, the credible interval is completely determined by the prior distribution and the MAP is the mode of the prior. In the middle and right panels, the data are incorporated into the posterior distributions. Both the credible intervals and MAPs are changing because of the influence of the data. Our method (Algorithm 3 and equation (12)) provides a simple and elegant approach of approximating the credible intervals on the manifold.

To see why the outputs from Algorithm 3 are a valid approximation to the posterior density, note that the density score $\widehat{\rho}_{i,N}$ is proportional to the underlying density of Z_1, \dots, Z_M defined over M . Hence, the weighted point cloud $(Z_1, \widehat{\rho}_{i,N}^{-1}), \dots, (Z_N, \widehat{\rho}_{N,N}^{-1})$ behaves like a uniform sample over M . Thus, to account for the unweighted point cloud density, we have to rescale the posterior

Algorithm 3 Approximated manifold posterior algorithm

1. Apply Algorithm 1 to generate many points $Z_1, \dots, Z_N \in M$.
2. Estimate a density score of Z_i using

$$\hat{\rho}_{i,N} = \frac{1}{N} \sum_{j=1}^N K\left(\frac{\|Z_i - Z_j\|}{h}\right),$$

where $h > 0$ is a tuning parameter and K is a smooth function such as a Gaussian.

3. Compute the posterior density score of Z_i as

$$\hat{\omega}_{i,N} = \frac{1}{\hat{\rho}_{i,N}} \cdot \hat{\pi}_{i,N}, \quad \hat{\pi}_{i,N} = \pi(Z_i) \cdot \prod_{j=1}^n p(X_j|Z_i), \quad (10)$$

return Weighted point clouds $(Z_1, \hat{\omega}_{1,N}), \dots, (Z_N, \hat{\omega}_{N,N})$.

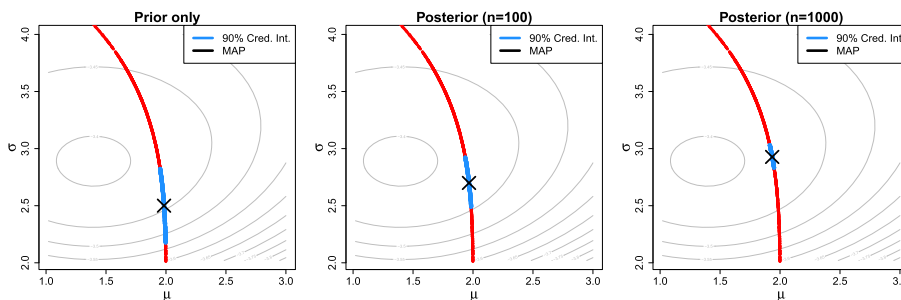


FIG 6. An example showing the credible interval (credible region) and MAP on the manifold. We use the same example in Figure 4 and the prior distribution $\pi(\mu, \sigma) \propto \phi(\mu; 2, 0.2)\phi(\sigma; 2.5, 0.2)I((\mu, \sigma) \in M)$, where $\phi(x; a, b)$ is the density of $N(a, b^2)$. **Left:** 90% credible interval along with the MAP using only the prior distribution. **Middle:** we randomly generate $n = 100$ observations from $N(1.5, 3^2)$ and compute 90% credible interval and MAP from the posterior distribution. **Right:** the same analysis as the middle panel but now we use a sample of size $n = 1000$. Note that the background gray contours show the log-likelihood function (as an indication of how the likelihood function will influence posterior).

score of Z_i in equation (10) by the factor $\hat{\rho}_{i,N}^{-1}$. Note that the value $\hat{\pi}_{i,N}$ is proportional to the posterior density $\pi(Z_i|X_1, \dots, X_n)$ evaluated at point $\theta = Z_i$. The quantity h is the smoothing bandwidth in the kernel density estimation. Because this is a density estimation problem, we would recommend to choose it using the Silverman's rule of thumb [74] or other popular approaches such as least square cross-validation [71, 9]; see the review paper of [73] for a list of reliable methods.

With the output from Algorithm 3, the posterior density $\pi(\theta|X_1, \dots, X_n)$ is represented by the collection of points Z_1, \dots, Z_N along with the corresponding weights $\hat{\omega}_{1,N}, \dots, \hat{\omega}_{N,N}$. The posterior mean can be approximated using

$$\hat{\theta}_{\text{Pmean}} = \frac{\sum_{i=1}^n \hat{\omega}_{i,N} Z_i}{\sum_{i=1}^n \hat{\omega}_{i,N}}.$$

This estimator is essentially the importance sampling estimator. The posterior

mode (MAP: maximum a posteriori) can be approximated using

$$\hat{\theta}_{\text{MAP}} = Z_{i^*}, \quad i^* = \operatorname{argmax}_{i \in \{1, \dots, N\}} \hat{\pi}_{i,N}.$$

The weighted point cloud also leads to an approximated credible region. Let $1 - \alpha$ be the credible level and $Z_{(1)}, \dots, Z_{(N)}$ be the ordered points such that

$$\hat{\pi}_{(1),N} \geq \hat{\pi}_{(2),N} \geq \dots \geq \hat{\pi}_{(N),N}.$$

Define

$$i(\alpha) = \operatorname{argmin} \left\{ i : \frac{\sum_{j=1}^i \hat{\omega}_{(j),N}}{\sum_{\ell=1}^N \hat{\omega}_{(\ell),N}} \geq 1 - \alpha \right\}. \quad (11)$$

Then we may use the collection of points

$$\{Z_{(1)}, \dots, Z_{(i(\alpha))}\} \quad (12)$$

as an approximation of a $1 - \alpha$ credible region. Alternatively, one may use the set

$$\{\theta \in M : \pi(\theta|X_1, \dots, X_n) \geq \pi(Z_{i(\alpha)}|X_1, \dots, X_n)\}$$

as another approximation of a $1 - \alpha$ credible region.

Here is an explanation of the choice in equation (11). $\hat{\pi}_{i,N}$ is proportional to the posterior value at Z_i ; hence,

$$\pi(Z_{(1)}|X_1, \dots, X_n) \geq \pi(Z_{(2)}|X_1, \dots, X_n) \geq \dots \geq \pi(Z_{(N)}|X_1, \dots, X_n).$$

Define the upper-level set of level λ of the posterior distribution as

$$L(\lambda) = \{\theta : \pi(\theta|X_1, \dots, X_n) \geq \lambda\}.$$

The posterior probability within $L(\lambda)$ is

$$\pi(L(\lambda)|X_1, \dots, X_n) = \int I(\theta \in L(\lambda)) \pi(\theta|X_1, \dots, X_n) d\theta.$$

A $1 - \alpha$ credible region can be constructed by choosing the minimal value λ_α such that

$$\lambda_\alpha = \inf \{\lambda : \pi(L(\lambda)|X_1, \dots, X_n) \geq 1 - \alpha\}. \quad (13)$$

With the weights $\hat{\omega}_{1,N}, \dots, \hat{\omega}_{N,N}$, an approximation to $\pi(L(\lambda)|X_1, \dots, X_n)$ is

$$\hat{\pi}(L(\lambda)|X_1, \dots, X_n) = \frac{\sum_{j=1}^n \hat{\omega}_{j,N} I(Z_j \in L(\lambda))}{\sum_{\ell=1}^n \hat{\omega}_{\ell,N}}.$$

The posterior levels $\hat{\pi}_{1,N}, \dots, \hat{\pi}_{N,N}$ form a discrete approximation of all levels of λ . Thus, an approximation to equation (13) is

$$\hat{\lambda}_\alpha = \min \left\{ \hat{\pi}_{i,N} : \frac{\sum_{j=1}^n \hat{\omega}_{j,N} I(Z_j \in L(\hat{\pi}_{i,N}))}{\sum_{\ell=1}^n \hat{\omega}_{\ell,N}} \geq 1 - \alpha \right\}$$

$$\begin{aligned}
&= \min \left\{ \hat{\pi}_{i,N} : \frac{\sum_{j=1}^n \hat{\omega}_{j,N} I(\hat{\pi}_{j,N} \geq \hat{\pi}_{i,N})}{\sum_{\ell=1}^n \hat{\omega}_{\ell,N}} \geq 1 - \alpha \right\} \\
&= \min \left\{ \hat{\pi}_{(i),N} : \frac{\sum_{j=1}^i \hat{\omega}_{(j),N}}{\sum_{\ell=1}^n \hat{\omega}_{(\ell),N}} \geq 1 - \alpha \right\} \\
&= \hat{\pi}_{(i(\alpha)),N}, \quad i(\alpha) = \operatorname{argmin} \left\{ i : \frac{\sum_{j=1}^i \hat{\omega}_{(j),N}}{\sum_{\ell=1}^n \hat{\omega}_{(\ell),N}} \geq 1 - \alpha \right\}.
\end{aligned}$$

The choice in equation (11) is from the above approximation to the level λ_α .

Remark 3. Note that the posterior mean may not be on the manifold. One may replace it by the posterior Fréchet mean [37] defined as

$$\hat{\theta}_{\text{PFmean}} = Z_{i^\dagger}, \quad i^\dagger = \operatorname{argmin}_{i \in \{1, \dots, N\}} \sum_{j=1}^N \hat{\omega}_{j,N} (Z_i - Z_j)^2.$$

The Fréchet mean defines a mean of a random variable X using the minimization problem $\operatorname{argmin}_\mu E(X - \mu)^2$ and constraints the minimizer to be in the manifold. Here, we use the weighted point approximation to this minimization.

Appendix B: Derivation of equation (2)

To derive the constraint in equation (3) from equation (2), we expand the first term and consider $s = 1$:

$$\begin{aligned}
P(Y=1|Q=1, A=0) &= \frac{P(Y=1, Q=1|A=0)}{P(Q=1|A=0)} \\
&= \frac{\sum_w P(Y=1, Q=1, W=w|A=0)}{\sum_{w'} P(Q=1, W=w'|A=0)} \\
&\stackrel{(*)}{=} \frac{\sum_w P(Q=1|W=w, A=0)P(W=w, Y=1|A=0)}{\sum_{w'} P(Q=1|W=w', A=0)P(W=w'|A=0)} \\
&= \frac{\sum_w q_{w,0} P(W=w, Y=1|A=0)}{\sum_{w'} q_{w',0} P(W=w'|A=0)}.
\end{aligned}$$

Note that the equality labeled with (*) is due to $Q \perp Y|A, W$. The two probabilities $P(W=w, Y=1|A=0)$ and $P(W=w'|A=0)$ are identifiable from the data. A similar calculation shows that

$$P(Y=1|Q=1, A=1) = \frac{\sum_w q_{w,1} P(W=w, Y=1|A=1)}{\sum_{w'} q_{w',1} P(W=w'|A=1)}.$$

So the test fairness constraint in equation (3) requires

$$\frac{\sum_w q_{w,0} P(W=w, Y=1|A=0)}{\sum_{w'} q_{w',0} P(W=w'|A=0)} = \frac{\sum_w q_{w,1} P(W=w, Y=1|A=1)}{\sum_{w'} q_{w',1} P(W=w'|A=1)}.$$

Also, for the case of $s = 0$, the above constraint becomes

$$\frac{\sum_w (1 - q_{w,0}) P(W=w, Y=1|A=0)}{\sum_{w'} (1 - q_{w',0}) P(W=w'|A=0)} = \frac{\sum_w (1 - q_{w,1}) P(W=w, Y=1|A=1)}{\sum_{w'} (1 - q_{w',1}) P(W=w'|A=1)}.$$

The above two equations are what equation (3) refers to as.

Appendix C: Proofs

Proof of Lemma 2. Essentially, we only need to show that when $d(x, M) \leq \frac{3\lambda_M^2}{8\|\Psi\|_{\infty,1}^* \|\Psi\|_{\infty,2}^*}$, the minimal eigenvalue $\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) \geq \frac{1}{4}\lambda_M^2$.

For any point x with $d(x, M) \leq \frac{3\lambda_M^2}{8\|\Psi\|_{\infty,1}^* \|\Psi\|_{\infty,2}^*}$, let x_M be the projection on M . The minimal eigenvalue

$$\begin{aligned} \lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) &= \lambda_{\min}(G_\Psi(x_M)G_\Psi(x_M)^T) \\ &\quad + (\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) - \lambda_{\min}(G_\Psi(x_M)G_\Psi(x_M)^T)) \\ &\geq \lambda_M^2 - |\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) - \lambda_{\min}(G_\Psi(x_M)G_\Psi(x_M)^T)|. \end{aligned} \tag{14}$$

The Weyl's theorem (see, e.g., Theorem 4.3.1 of [41]) shows that the eigenvalue difference can be bounded via

$$\begin{aligned} &|\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) - \lambda_{\min}(G_\Psi(x_M)G_\Psi(x_M)^T)| \\ &\leq \|G_\Psi(x)G_\Psi(x)^T - G_\Psi(x_M)G_\Psi(x_M)^T\|_{\max} \\ &\leq 2\|\Psi\|_{\infty,1}^* \|\Psi\|_{\infty,2}^* \|x - x_M\| \quad (\text{by Taylor's theorem}) \\ &= 2\|\Psi\|_{\infty,1}^* \|\Psi\|_{\infty,2}^* d(x, M). \end{aligned}$$

Thus, as long as

$$2\|\Psi\|_{\infty,1}^* \|\Psi\|_{\infty,2}^* d(x, M) \leq \frac{3}{4}\lambda_M^2$$

we have $\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) \geq \frac{1}{4}\lambda_M^2$, which completes the proof. \square

Proof of Theorem 3. The proof is modified from the proof of Lemma 4.11 and Theorem 4.12 in [33].

Let $r_0 = \min\{\frac{\delta_0}{2}, \frac{\lambda_0}{\|\Psi\|_{\infty,2}^*}\}$. We will show that r_0 is a lower bound of reach (M). We proceed by the proof of contradiction.

Suppose that the conclusion is incorrect that the reach is less than r_0 . Then there exists a point x such that $d(x, M) < r_0$ and x has two projections onto M , denoted as $b, c \in M$.

Since $b, c \in M$, $\Psi(b) = \Psi(c) = 0$ and by Taylor's remainder theorem and condition (D-2) and (F),

$$\begin{aligned} \|G_\Psi(b)(b-c)\|_2 &= \|f(b) - f(c) - G_\Psi(b)(b-c)\|_2 \\ &\leq \frac{1}{2}\|b-c\|_2^2 \|\Psi\|_{\infty,2}^*. \end{aligned} \tag{15}$$

By the nature of projection, we can find a vector $t_b \in \mathbb{R}^s$ such that $x - b = t_b^T G_\Psi(b)$ because the normal space is spanned by the row space of $G_\Psi(b)$ (Lemma 1). Together with (15), this implies

$$\begin{aligned} 2|(x - b)^T(b - c)| &= 2|t_b^T G_\Psi(b)(b - c)| \\ &\leq \|G_\Psi(b)(b - c)\|_2 \|t_b\|_2 \\ &\leq \|\Psi\|_{\infty,2}^* \|b - c\|_2^2 \|t_b\|_2. \end{aligned} \quad (16)$$

Since b, c are projections of x onto M ,

$$\|x - b\|_2 = \|x - c\|_2.$$

As a result,

$$\begin{aligned} 0 &= \|x - c\|_2^2 - \|x - b\|_2^2 \\ &= \|b - c\|_2^2 + 2(b - c)^T(x - b) \\ &\geq \|b - c\|_2^2 - \|\Psi\|_{\infty,2}^* \|b - c\|_2^2 \|t_b\|_2 \quad (16) \\ &= \|b - c\|_2^2 (1 - \|\Psi\|_{\infty,2}^* \|t_b\|_2). \end{aligned} \quad (17)$$

However, starting from the definition of r_0 , we have

$$\begin{aligned} \frac{\lambda_0}{\|\Psi\|_{\infty,2}^*} > r_0 \geq \|x - b\|_2 = \|t_b^T G_\Psi(b)\|_2 \\ \geq \underbrace{\lambda_0}_{(F)} \|t_b\|_2. \end{aligned} \quad (18)$$

As a result,

$$\|\Psi\|_{\infty,2}^* \|t_b\|_2 < 1$$

so $\|b - c\|_2^2 = 0$ by (17), which implies $b = c$, a contradiction. Accordingly, x must have a unique projection so the reach has a lower bound r_0 . \square

Proof of Proposition 4. Consider a point $x \in \tilde{M}$. By the condition $\|\tilde{\Psi} - \Psi\|_{\infty,0}^* < c_0$ and assumption (F), we know that $d(x, M) \leq \delta_0$, where c_0 and δ_0 are the constants in (F).

Define

$$h(x) = \|\Psi(x)\|_2 = \sqrt{\Psi(x)^T \Psi(x)} \quad (19)$$

to be the L_2 norm for Ψ . The derivative of $h(x)$

$$\nabla h(x) = \frac{\Psi(x)^T G_\Psi(x)}{\|\Psi(x)\|_2} \quad (20)$$

is a vector of \mathbb{R}^d . Note that $G_\Psi(x) = \nabla \Psi(x) \in \mathbb{R}^{s \times d}$ is the Jacobian.

For any point $x \in \tilde{M}$, we define a flow

$$\phi_x : \mathbb{R} \mapsto \mathbb{R}^d \quad (21)$$

such that

$$\phi_x(0) = x, \quad \frac{\partial}{\partial t} \phi_x(t) = -\nabla h(\phi(t)). \quad (22)$$

Later we will prove in Theorem 7 that $\phi_x(\infty) \in M$ when $x \in M \oplus \delta_c$, where δ_c is defined in Theorem 7.

By Theorem 3.39 in [42], $\phi_x(t)$ is uniquely defined since the gradient $\nabla h(x)$ is well-defined for all $x \notin M$. We define an arc-length flow (i.e., a constant velocity flow) based on ϕ_x :

$$\gamma_x(0) = x, \quad \frac{\partial}{\partial t} \gamma_x(t) = -\frac{\nabla h(\gamma_x(t))}{\|\nabla h(\gamma_x(t))\|_2}. \quad (23)$$

The time traveled in this flow is the same as the distance traveled (due to the velocity being a unit vector). Let $T_x = \inf\{t > 0 : \gamma_x(t) \in M\}$ be the terminal time point and let $\gamma_x(T_x) \in M$ as the endpoint of the flow. This means that T_x is the length of the flow from x to the destination on M . The goal is to bound T_x since the length must be greater or equal to the projection distance for $x \in \tilde{M}$.

We define $\xi_x(t) = h(\gamma_x(t)) - h(\gamma_x(T_x)) = h(\gamma_x(t))$. Differentiating $\xi_x(t)$ with respect to t leads to

$$\begin{aligned} \xi'_x(t) &= -\frac{d}{dt} h(\gamma_x(t)) \\ &= -[\nabla h(\gamma_x(t))]^T \frac{d}{dt} \gamma_x(t) \\ &= -\|\nabla h(\gamma_x(t))\| \\ &= -\frac{\|\Psi(\gamma_x(t))^T G_\Psi(\gamma_x(t))\|_2}{\|\Psi(\gamma_x(t))\|_2} \\ &\leq -\lambda_{\min}(G_\Psi(\gamma_x(t))G_\Psi(\gamma_x(t))^T) \\ &\leq -\lambda_0 \end{aligned} \quad (24)$$

because $\gamma_x(t) \in M \oplus \delta_0$ for all t .

Let $\epsilon_0 = \|\Psi - \tilde{\Psi}\|_{\infty,0}^* = \sup_x \|\Psi(x) - \tilde{\Psi}(x)\|_{\max}$ and recall that $x \in \tilde{M}$ so $\tilde{\Psi}(x) = 0$. Then by the fact that $\|v\|_2 \leq \sqrt{d} \times \|v\|_{\max}$ for vector v ,

$$\begin{aligned} \sqrt{d} \cdot \epsilon_0 &= \sqrt{d} \sup_x \|\Psi(x) - \tilde{\Psi}(x)\|_{\max} \\ &\geq \sup_x \|\Psi(x) - \tilde{\Psi}(x)\| \\ &\geq \|\Psi(x) - \tilde{\Psi}(x)\| \\ &\geq h(x) \\ &= h(\gamma_x(0)) - h(\gamma_x(T_x)) \quad (\text{since } h(\gamma_x(T_x)) = 0) \\ &= \xi(0) - \xi(T_x) \quad (\xi(T_x) = 0 \text{ and } \xi(0) = h(0)) \\ &= -T_x \xi'(T_x^*) \quad (\text{mean value Theorem}) \\ &\geq T_x \lambda_0 \quad \text{by equation (24)}. \end{aligned}$$

Hence, $T_x \leq \frac{\sqrt{d}}{\lambda_0} \epsilon_0 = O(\epsilon_0)$ which is independent of x . This implies that

$$\sup_{x \in \tilde{M}} d(x, M) \leq \frac{\sqrt{d}}{\lambda_0} \epsilon_0 = O(\|\tilde{\Psi} - \Psi\|_{\infty,0}^*). \quad \square$$

Proof of Theorem 5. 1. Since condition (F) involves only Ψ and its derivative, when $\|\Psi - \tilde{\Psi}\|_{\infty,2}^*$ is sufficiently small, (F) holds for $\tilde{\Psi}$.

2. By the first assertion, condition (F) holds for $\tilde{\Psi}$.

Thus, we can exchange \tilde{M} and M and repeat the proof of Proposition 4, which leads to

$$\sup_{x \in M} d(x, \tilde{M}) \leq \frac{\sqrt{d}}{\lambda_0} \epsilon_0.$$

As a result, we conclude that $\text{Haus}(\tilde{M}, M) \leq \frac{\sqrt{d}}{\lambda_0} \epsilon_0 = O(\epsilon_0)$.

3. By Theorem 3, the reach of M has lower bound $\min\{\delta_0/2, \lambda_0/\|\Psi\|_{\infty,2}^*\}$.

Note that δ_0, λ_0 depends on the first derivative of Ψ . Hence, the lower bound for reach of M and \tilde{M} will be bounded at rate $O(\|\Psi - \tilde{\Psi}\|_{\infty,2}^*)$. \square

Before moving forward, we would like to note that the Jacobian and Hessian of f can be expressed as

$$G_f(x) = \nabla f(x) = 2\Psi(x)^T \Lambda[\nabla\Psi(x)] \quad (25)$$

$$H_f(x) = \nabla\nabla f(x) = 2[\nabla\Psi(x)]^T \Lambda[\nabla\Psi(x)] + 2\Psi(x)\Lambda[\nabla\nabla\Psi(x)], \quad (26)$$

$$\nabla\nabla\nabla f(x) = 6[\nabla\Psi(x)]^T \Lambda[\nabla\nabla\Psi(x)] + 2\Psi(x)\Lambda[\nabla\nabla\nabla\Psi(x)], \quad (27)$$

where $\nabla\Psi(x) \in \mathbb{R}^{s \times d}$ and $\nabla\nabla\Psi(x) \in \mathbb{R}^{s \times d \times d}$.

Proof of Lemma 6. Property 1 (For each $x \in M$).

1-(a). By equation (26) and the fact that $\Psi(x) = 0$ whenever $x \in M$, we obtain

$$H_f(x) = 2[\nabla\Psi(x)]^T \Lambda[\nabla\Psi(x)].$$

Because Λ is positive definite, it can be decomposed into $\Lambda = UDU^T$ where D is a diagonal matrix so the eigenvectors corresponding to non-zero eigenvalues of H_f will be the the rows of $U^T[\nabla\Psi(x)]$, which spans the same subspace as the row space of $[\nabla\Psi(x)]$ so by Lemma 1, the non-zero eigenvectors spans the normal space of M at x .

1-(b). Because $H_f(x) = 2[\nabla\Psi(x)]^T \Lambda[\nabla\Psi(x)]$ when $x \in M$, the minimal non-zero eigenvalue

$$\lambda_{\min,>0}(H_f(x)) = 2\lambda_{\min,>0}(G_\Psi(x)^T \Lambda G_\Psi(x)).$$

Since Λ is positive definite and symmetric, we can decompose

$$G_\Psi(x)^T \Lambda G_\Psi(x) = G_\Psi(x)^T \Lambda^{1/2} \Lambda^{1/2} G_\Psi(x)$$

so we obtain

$$\lambda_{\min,>0}(H_f(x)) = 2\lambda_{\min,>0}(G_\Psi(x)^T \Lambda G_\Psi(x))$$

$$\begin{aligned}
&= 2\lambda_{\min}(\Lambda^{1/2}G_{\Psi}(x)G_{\Psi}(x)^T\Lambda^{1/2}) \\
&\geq 2\Lambda_{\min}\lambda_{\min}(G_{\Psi}(x)G_{\Psi}(x)^T) \\
&\geq 2\Lambda_{\min}\lambda_0^2.
\end{aligned}$$

1-(c). Because the normal space of M at x is spanned by the rows of $G_{\Psi}(x) = \nabla\Psi(x)$, which by 1-(a) is spanned by the non-zero eigenvectors of $H_f(x)$, the result follows.

Property 2. Because $d(x, M) < \delta_c$ so x is within the reach of M and thus, x_M , the projection from x onto M , is unique. As a result, the normal space of x_M , $N_M(x)$, is well-defined.

We can decompose

$$\begin{aligned}
\lambda_{\min, \perp, M}(H_f(x)) &= \min_{v \in N_M(x)} \frac{v^T H_f(x) v}{\|v\|^2} \\
&= \min_{v \in N_M(x)} \frac{v^T (H_f(x_M) + H_f(x) - H_f(x_M)) v}{\|v\|^2} \\
&\geq \min_{v \in N_M(x)} \frac{v^T H_f(x_M) v}{\|v\|^2} - \max_{v \in N_M(x)} \frac{v^T (H_f(x) - H_f(x_M)) v}{\|v\|^2} \\
&\geq \min_{v \in N_M(x)} \frac{v^T H_f(x_M) v}{\|v\|^2} - d \|H_f(x) - H_f(x_M)\|_{\max} \\
&\geq 2\lambda_0^2 \Lambda_{\min} - d \|H_f(x) - H_f(x_M)\|_{\max}
\end{aligned}$$

By equation (26),

$$\begin{aligned}
\|H_f(x) - H_f(x_M)\|_{\max} &\leq 2\|G_{\Psi}(x)^T \Lambda G_{\Psi}(x) - G_{\Psi}(x_M)^T \Lambda G_{\Psi}(x_M)\|_{\max} \\
&\quad + 2\|\Psi(x) \Lambda H_{\Psi}(x) - \Psi(x_M) \Lambda H_{\Psi}(x_M)\|_{\max} \\
&\leq 4\|\Psi\|_{\infty, 1}^* \Lambda_{\max} \|\Psi\|_{\infty, 2}^* \|x - x_M\| \\
&\quad + 4\|\Psi\|_{\infty, 2}^* \Lambda_{\max} \|\Psi\|_{\infty, 3}^* \|x - x_M\| \\
&\leq 8\|\Psi\|_{\infty, 2}^* \Lambda_{\max} \|\Psi\|_{\infty, 3}^* \|x - x_M\|.
\end{aligned}$$

Thus, as long as

$$\|x - x_M\| = d(x, M) \leq \frac{\lambda_0^2 \Lambda_{\min}}{8d \Lambda_{\max} \|\Psi\|_{\infty, 2}^* \|\Psi\|_{\infty, 3}^*},$$

we have

$$\lambda_{\min, \perp, M}(H_f(x)) \geq \lambda_0^2 \Lambda_{\min},$$

which completes the proof. \square

Proof of Theorem 7. 1. Convergence radius. We prove this by showing that for any $x \in M \oplus \delta_c$ and $x \notin M$, the destination $\pi_x(\infty) \in M$. The idea of the proof relies on two properties:

(P1) Any stationary point of f inside $M \oplus \delta_c$ must be a point in M .

(P2) Let x_M be a point on M that is closest to x . For any point $x \in M \oplus \delta_c$, $(x - x_M)^T \nabla f(x) > 0$. Namely, the gradient flow only moves $\pi_x(t)$ closer toward M .

With the above two properties, it is easy to see that if we start a gradient flow π_x from $x \in M \oplus \delta_c$, then by (P2) this flow must stay within $M \oplus \delta_c$. Because stationary points within $M \oplus \delta_c$ are all in M by (P1) and the destination of a gradient flow must be a stationary point, we conclude that $\pi_x(\infty) \in M$, which completes the proof of convergence radius. In what follows, we show the two properties.

Property P1: Any stationary point inside $M \oplus \delta_c$ must be a point in M . Because $\nabla f(x) = \Psi(x) \Lambda G_\Psi(x)$ and Λ is positive definite, there are only two cases that $\nabla f(x) = 0$: 1. $\Psi(x) = 0$ and 2., row space of $G_\Psi(x)$ has a dimension less than s (in fact, if $\Psi(x) \neq 0$, then the second case is a necessary condition). The first case is the solution manifold M so we only need to focus on showing that the second case will not happen for $x \in M \oplus \delta_c$.

The row space of $G_\Psi(x)$ has a dimension less than s when there exists a singular value of $G_\Psi(x)$ being 0; or equivalently, $\lambda_{\min}(G_\Psi(x)G_\Psi(x)^T) = 0$. However, assumption (F) already requires that this will not happen within $M \oplus \delta_c$. Thus, this property holds.

Property P2: For any $x \in M \oplus \delta_c$, the directional gradient $(x - x_M)^T \nabla f(x) > 0$. By Taylor expansion and property 2 of Lemma 6,

$$\begin{aligned} (x - x_M)^T \nabla f(x) &= (x - x_M)^T (\nabla f(x) - \underbrace{\nabla f(x_M)}_{=0}) \\ &= (x - x_M)^T \int_{\epsilon=0}^{\epsilon=1} H_f(x_M + \epsilon(x - x_M))(x - x_M) d\epsilon \quad (28) \\ &\geq \|x - x_M\|^2 \inf_{y \in M \oplus \delta_c} \lambda_{\min, \perp, M}(H_f(y)) \\ &\geq d(x, M)^2 \lambda_0^2 \Lambda_{\min} > 0 \end{aligned}$$

2. Terminal flow orientation. To study the gradient flow close to M , it suffices to analyze the behavior of gradient close to M . Let $x \in M$ and define u to be a unit vector in the normal space of M at x . By Lemma 1, u belongs to the row space of $\nabla \Psi(x) = G_\Psi(x)$.

Now we consider the gradient at $x + \epsilon u$ when $\epsilon \rightarrow 0$. By Taylor's theorem and the fact that f has bounded third derivatives (from (D-3)),

$$G_f(x + \epsilon u) \equiv \nabla f(x + \epsilon u) = \nabla f(x + \epsilon u) - \nabla f(x) = \epsilon H_f(x)u + O(\epsilon^2).$$

Thus,

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} G_f(x + \epsilon u) = H_f(x)u.$$

By equation (26),

$$H_f(x) = 2G_\Psi(x)^T \Lambda G_\Psi(x) + 2\Psi(x) \Lambda H_\Psi(x) = 2G_\Psi(x)^T \Lambda G_\Psi(x)$$

because $\Psi(x) = 0$ when $x \in M$. Using the fact that $G_\Psi(x)^T = [\nabla\Psi_1(x), \dots, \nabla\Psi_s(x)]$, it is easy to see that

$$H_f(x)u = \sum_{\ell=1}^s a_\ell \nabla\Psi_\ell(x),$$

where $a_\ell = e_\ell^T \Lambda G_\Psi(x)u$ with $e_\ell = (0, 0, \dots, 0, 1, 0, \dots, 0)^T \in \mathbb{R}^s$ is the coordinate vector pointing toward ℓ -th coordinate. Thus, by Lemma 1 $\nabla\Psi_\ell(x)u$ belongs to the normal space of M at x , which completes the proof of terminal orientation. \square

Proof of Theorem 8. We prove this result using the idea of the Lyapunov-Perron method [61]. Recall that $A(z) = \{x : \pi_x(\infty) = z\}$ for $z \in M$ is the basin of attraction of point z . Consider a ball $B(z, r)$ such that any gradient flow $\pi_x(t)$ that converges to $z = \pi_x(\infty)$ intersects one and only one point at the boundary $\partial B(z, r) = \{y : \|y - z\| = r\}$. This occurs when $r < \delta_c$ due to property (P2) in the proof of Theorem 7.

Consider the gradient flow $\pi_x(t)$ with $x \in \partial B(z, r)$ and $\pi_x(\infty) = z$. By Taylor's theorem, this flow solves the following equation

$$\begin{aligned} \pi'_x(t) &= -G_f(\pi_x(t)) = -G_f(\pi_x(t)) + \underbrace{G_f(\pi_x(\infty))}_{=0} \\ &= -H_f(\pi_x(\infty))(\pi_x(t) - \pi_x(\infty)) + \epsilon(\pi_x(t)), \end{aligned} \quad (29)$$

where $\|\epsilon(\pi_x(t))\| \leq C_0 \|\pi_x(t) - \pi_x(\infty)\| \leq C_0 r$ for some finite constant C_0 due to Assumption (D-3). Equation (29) is a perturbed ODE with a fixed point $\pi_x(\infty)$ and by the variation of parameters, its solution can be written as

$$\pi_x(t) - \pi_x(\infty) = e^{-tH_f(\pi_x(\infty))}(\pi_x(0) - \pi_x(\infty)) + \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon(\pi_x(s)) ds.$$

Denoting $v_x = \pi_x(0) - \pi_x(\infty)$, we can rewrite the flow as

$$\pi_x(t) - \pi_x(\infty) = e^{-tH_f(\pi_x(\infty))} v_x + \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon(\pi_x(s)) ds.$$

By Lemma 1, the normal space of M at $z = \pi_x(\infty)$ is the row space of $G_\Psi(z) = \nabla\Psi(z)$, which will also be the space spanned by the eigenvectors of $H_f(\pi_x(\infty))$ that corresponds to non-zero eigenvalues (Lemma 6 1-(a)). The spectral decomposition shows $H_f(\pi_x(\infty)) = \sum_{\ell=1}^s \lambda_\ell u_\ell u_\ell^T$ and we define the projection matrix onto the normal space of M as $\Pi_N = \sum_{\ell=1}^s u_\ell u_\ell^T$ and the projection matrix onto the tangent space of M as $\Pi_T = I_d - \Pi_N$. By construction, $\Pi_N H_f(\pi_x(\infty)) = H_f(\pi_x(\infty))$ and $\Pi_T H_f(\pi_x(\infty)) = 0$ so $\Pi_N e^{-tH_f(\pi_x(\infty))} = e^{-tH_f(\pi_x(\infty))}$ and $\Pi_T e^{-tH_f(\pi_x(\infty))} = \Pi_T$.

We decompose

$$\begin{aligned}
\pi_x(t) - \pi_x(\infty) &= \Pi_T(\pi_x(t) - \pi_x(\infty)) + \Pi_N(\pi_x(t) - \pi_x(\infty)) \\
&= \Pi_T e^{-tH_f(\pi_x(\infty))} v_x + \Pi_T \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon(\pi_x(s)) ds \\
&\quad + \Pi_N e^{-tH_f(\pi_x(\infty))} v_x + \Pi_N \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon(\pi_x(s)) ds \\
&= v_{x,T} + \int_{s=0}^{s=t} \epsilon_T(\pi_x(s)) ds + e^{-tH_f(\pi_x(\infty))} v_{x,N} \\
&\quad + \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon_N(\pi_x(s)) ds,
\end{aligned} \tag{30}$$

where

$$v_{x,T} = \Pi_T v_x, \quad v_{x,N} = \Pi_N v_x, \quad \epsilon_T(\pi_x(s)) = \Pi_T \epsilon(\pi_x(s)), \quad \epsilon_N(\pi_x(s)) = \Pi_N \epsilon(\pi_x(s)).$$

In the tangent direction, when $t \rightarrow \infty$

$$\begin{aligned}
0 &= \lim_{t \rightarrow \infty} \Pi_T(\pi_x(t) - \pi_x(\infty)) \\
&= \lim_{t \rightarrow \infty} \Pi_T e^{-tH_f(\pi_x(\infty))} v_x + \lim_{t \rightarrow \infty} \Pi_T \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon(\pi_x(s)) ds \\
&= \Pi_T v_x + \int_{s=0}^{s=\infty} \Pi_T \epsilon(\pi_x(s)) ds \\
&= v_{x,T} + \int_{s=0}^{s=\infty} \epsilon_T(\pi_x(s)) ds.
\end{aligned}$$

Thus,

$$v_{x,T} = - \int_{s=0}^{s=\infty} \epsilon_T(\pi_x(s)) ds \tag{31}$$

and equation (30) can be rewritten as

$$\begin{aligned}
\pi_x(t) - \pi_x(\infty) &= - \int_{s=0}^{s=\infty} \epsilon_T(\pi_x(s)) ds + \int_{s=0}^{s=t} \epsilon_T(\pi_x(s)) ds \\
&\quad + e^{-tH_f(\pi_x(\infty))} v_{x,N} + \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon_N(\pi_x(s)) ds \\
&= e^{-tH_f(\pi_x(\infty))} v_{x,N} + \int_{s=0}^{s=t} e^{-(t-s)H_f(\pi_x(\infty))} \epsilon_N(\pi_x(s)) ds \\
&\quad - \int_{s=t}^{s=\infty} \epsilon_T(\pi_x(s)) ds.
\end{aligned} \tag{32}$$

The latter two terms involving integral are determined entirely by the Taylor remainder terms $\epsilon(\pi_x(t))$. Thus, to uniquely determine a point on the gradient flow $\pi_x(t)$ that converges to z (and is inside $B(z, r)$), we only need to specify

the time t and the vector $v_{x,N}$ that belongs to the normal space of M at z with $\|v_{x,N}\| = r$. Namely, there exists a mapping (due to equation (32)) Ω such that

$$\pi_x(t) = \Omega(t, v_{x,N})$$

for all $\pi_x(t)$ with $\|x - z\| = r$. Note that equation (32) implies that the mapping Ω has bounded derivative with respect to both t and $v_{x,N}$. Therefore, the set

$$A(z) \cap B(z, r) = \left\{ \pi_x(t) = \Omega(t, v_{x,N}) : t \in [0, \infty), v_{x,N} = \sum_{\ell=1}^s a_\ell u_\ell, \sum_{\ell=1}^s a_\ell^2 = r^2 \right\}$$

is parameterized by (t, a_1, \dots, a_s) with a constraint $\sum_{\ell=1}^s a_\ell^2 = r^2$ so it is an s -dimensional manifold.

To generalize this to the entire set $A(z)$, note that every gradient flow ending at z must pass the boundary $\partial B(z, r)$ so allowing the gradient $\pi_x(t)$ to move toward $t \rightarrow -\infty$ covers the entire basin, i.e.,

$$A(z) = \left\{ \pi_x(t) = \Omega(t, v_{x,N}) : t \in \mathbb{R}, v_{x,N} = \sum_{\ell=1}^s a_\ell u_\ell, \sum_{\ell=1}^s a_\ell^2 = r^2 \right\}.$$

This implies that $A(z)$ is parametrized by (t, a_1, \dots, a_s) with a constraint $\sum_{\ell=1}^s a_\ell^2 = r^2$ so again it is an s -dimensional manifold. \square

Proof of Theorem 9. Convergence of $f(x_t)$. Because $x_{t+1} = x_t - \gamma G_f(x_t)$, simple Taylor expansion shows that

$$\begin{aligned} f(x_t) - f(x_{t+1}) &= f(x_t) - f(x_t - \gamma G_f(x_t)) \\ &= \gamma \|G_f(x_t)\|^2 - \frac{1}{2} \gamma^2 \int_{\epsilon=0}^{\epsilon=1} G_f(x_t) H_f(x_t - \epsilon \gamma G_f(x_t)) G_f(x_t) d\epsilon \\ &\geq \gamma \|G_f(x_t)\|^2 - \frac{1}{2} \gamma^2 \|G_f(x_t)\|^2 \sup_z \|H_f(z)\|_2. \end{aligned}$$

Note that one can also use the fact that the gradient G_f is Lipschitz to obtain a similar bound.

Thus, when $\gamma < \frac{2}{\sup_z \|H_f(z)\|_2}$, we obtain

$$f(x_t) - f(x_{t+1}) > 0$$

which implies that $f(x_{t+1}) < f(x_t)$, i.e., the objective function is decreasing. We can summarize the result as

$$f(x_{t+1}) \leq f(x_t) - \gamma \|G_f(x_t)\|^2 \left(1 - \frac{1}{2} \gamma \sup_z \|H_f(z)\|_2 \right). \quad (33)$$

To obtain the algorithmic convergence rate, we need to associate the objective function $f(x)$ and the squared gradient $\|G_f(x)\|^2$. We focus on the case of $t = 0$, and investigate

$$f(x_1) \leq f(x_0) - \gamma \|G_f(x_0)\|^2 \left(1 - \frac{1}{2} \gamma \sup_z \|H_f(z)\|_2 \right). \quad (34)$$

Because $d(x_0, M) \leq \delta_c \leq \text{reach}(M)$, there is a unique projection $x_M \in M$ from x_0 . Note that $d(x_0, M) = \|x_0 - x_M\|$. The gradient has a lower bound from the following Taylor expansion:

$$\begin{aligned} \|G_f(x_0)\| &= \|G_f(x_0) - \underbrace{G_f(x_M)}_{=0}\| \\ &= \left\| \int_{\epsilon=0}^{\epsilon=1} H_f(x_M + \epsilon(x_0 - x_M))(x_0 - x_M) d\epsilon \right\| \\ &\geq \|x_0 - x_M\| \inf_{\epsilon \in [0,1]} \lambda_{\min, \perp}(H_f(x_M + \epsilon(x_0 - x_M))) \\ &\geq \|x_0 - x_M\| \lambda_0^2 \Lambda_{\min} \\ &\geq d(x_0, M) \lambda_0^2 \Lambda_{\min}, \end{aligned}$$

where the second to the last inequality is due to property 2 in Lemma 6. Thus,

$$\|G_f(x_0)\|^2 \geq d(x_0, M)^2 \lambda_0^4 \Lambda_{\min}^2. \quad (35)$$

The distance $d(x_0, M)$ and the objective function $f(x_0)$ can also be associated using another Taylor expansion:

$$\begin{aligned} f(x_0) &= f(x_0) - f(x_M) \\ &= (x_0 - x_M)^T \underbrace{G_f(x_M)}_{=0} \\ &\quad + \frac{1}{2} (x_0 - x_M)^T \int_{\epsilon=0}^{\epsilon=1} H_f(x_M + \epsilon(x_0 - x_M)) d\epsilon (x_0 - x_M) \\ &\leq \frac{1}{2} d^2(x_0, M) \sup_z \|H_f(z)\|_2. \end{aligned}$$

Thus,

$$d^2(x_0, M) \geq \frac{2f(x_0)}{\sup_z \|H_f(z)\|_2}$$

which implies an improved bound on equation (35) as

$$\|G_f(x_0)\|^2 \geq d(x_0, M)^2 \lambda_0^4 \Lambda_{\min}^2 \geq \frac{\lambda_0^4 \Lambda_{\min}^2}{\sup_z \|H_f(z)\|_2} 2f(x_0). \quad (36)$$

By inserting equation (36) into equation (34), we obtain

$$\begin{aligned} f(x_1) &\leq f(x_0) - \gamma \|G_f(x_0)\|^2 \left(1 - \frac{1}{2} \gamma \sup_z \|H_f(z)\|_2\right) \\ &\leq f(x_0) - \gamma \left(1 - \frac{1}{2} \gamma \sup_z \|H_f(z)\|_2\right) \frac{\lambda_0^4 \Lambda_{\min}^2}{\sup_z \|H_f(z)\|_2} 2f(x_0) \\ &= f(x_0) \left(1 - 2\gamma \left(1 - \frac{1}{2} \gamma \sup_z \|H_f(z)\|_2\right) \frac{\lambda_0^4 \Lambda_{\min}^2}{\sup_z \|H_f(z)\|_2}\right) \end{aligned}$$

When $\gamma < \frac{1}{\sup_z \|H_f(z)\|_2}$, the above inequality can be simplified as

$$f(x_1) \leq f(x_0) \cdot \left(1 - \gamma \frac{\lambda_0^4 \Lambda_{\min}^2}{\sup_z \|H_f(z)\|_2}\right).$$

Thus, we have proved the result for $t = 0$. The same derivation works for other t (by treating x_t as x_0). By telescoping, we conclude that

$$f(x_t) \leq f(x_0) \cdot \left(1 - \gamma \frac{\lambda_0^4 \Lambda_{\min}^2}{\sup_z \|H_f(z)\|_2}\right)^t.$$

Finally, using the fact that $\sup_z \|H_f(z)\|_2 \leq \Lambda_{\max} \|\Psi\|_{\infty,2}^*$, we obtain the desired bound.

Convergence of $d(x_t, M)$. Let $x_{t,M} \in M$ be the point on the manifold that is closest to x_t ; again, due to the reach condition this projection is unique. The Taylor expansion along with property 2 in Lemma 6 shows that

$$\begin{aligned} -f(x_t) &= f(x_{t,M}) - f(x_t) \\ &= (x_{t,M} - x_t)^T G_f(x_t) \\ &\quad + \frac{1}{2} (x_{t,M} - x_t)^T \int_{\epsilon=0}^{\epsilon=1} H_f(x_t + \epsilon(x_{t,M} - x_t)) (x_t - x_{t,M}) d\epsilon \\ &\geq (x_{t,M} - x_t)^T G_f(x_t) + \frac{1}{2} \|x_t - x_{t,M}\|^2 \lambda_0^2 \Lambda_{\min}. \end{aligned}$$

Thus,

$$-f(x_t) - \frac{1}{2} \|x_{t,M} - x_t\|^2 \lambda_0^2 \Lambda_{\min} \geq -(x_t - x_{t,M})^T G_f(x_t). \quad (37)$$

Because of equation (33) and that $\sup_z \|H_f(z)\|_2 \leq d \|\Psi\|_{\infty,2}^*$, we have

$$f\left(x - \frac{1}{d \Lambda_{\max} \|\Psi\|_{\infty,2}^*} G_f(x)\right) - f(x) \leq -\frac{1}{2d \Lambda_{\max} \|\Psi\|_{\infty,2}^*} \|G_f(x)\|^2.$$

Using the fact that $f\left(x - \frac{1}{d \Lambda_{\max} \|\Psi\|_{\infty,2}^*} G_f(x)\right) \geq 0$, we conclude that

$$\frac{1}{2d \Lambda_{\max} \|\Psi\|_{\infty,2}^*} \|G_f(x)\|^2 \leq f(x),$$

which implies

$$\|G_f(x)\|^2 \leq 2d \Lambda_{\max} \|\Psi\|_{\infty,2}^* f(x). \quad (38)$$

For any t , we have

$$\begin{aligned} d(x_{t+1}, M)^2 &\leq \|x_{t+1} - x_{t,M}\|^2 \\ &= \|x_t - x_{t,M} - \gamma G_f(x_t)\|^2 \\ &= \|x_t - x_{t,M}\|^2 - 2(x_t - x_{t,M})^T G_f(x_t) + \gamma^2 \|G_f(x_t)\|^2 \end{aligned}$$

$$\begin{aligned}
& \stackrel{(37)}{\leq} \|x_t - x_{t,M}\|^2(1 - \gamma\lambda_0^2\Lambda_{\min}) - 2\gamma f(x_t) + \gamma^2\|G_f(x_t)\|^2 \\
& \stackrel{(38)}{\leq} \|x_t - x_{t,M}\|^2(1 - \gamma\lambda_0^2\Lambda_{\min}) - \underbrace{2\gamma f(x_t)(1 - d\gamma\Lambda_{\max}\|\Psi\|_{\infty,2}^*)}_{\geq 0} \\
& \leq \|x_t - x_{t,M}\|^2(1 - \gamma\lambda_0^2\Lambda_{\min}) \\
& = d(x_t, M)^2(1 - \gamma\lambda_0^2\Lambda_{\min})
\end{aligned}$$

whenever $\gamma < \frac{1}{\Lambda_{\max}\|\Psi\|_{\infty,2}^*}$. By telescoping, the result follows. \square

References

- [1] AAMARI, E., KIM, J., CHAZAL, F., MICHEL, B., RINALDO, A. and WASSERMAN, L. (2019). Estimating the Reach of a Manifold. *Electronic Journal of Statistics* **13** 1359–1399. [MR3938326](#)
- [2] AAMARI, E. and LEVRARD, C. (2019). Nonasymptotic rates for manifold, tangent space and curvature estimation. *The Annals of Statistics* **47** 177–204. [MR3909931](#)
- [3] ARIAS-CASTRO, E., MASON, D. and PELLETIER, B. (2016). On the Estimation of the Gradient Lines of a Density and the Consistency of the Mean-Shift Algorithm. *Journal of Machine Learning Research* **17** 1–28. [MR3491137](#)
- [4] BALAKRISHNAN, S., WAINWRIGHT, M. J. and YU, B. (2017). Statistical guarantees for the EM algorithm: From population to sample-based analysis. *The Annals of Statistics* **45** 77–120. [MR3611487](#)
- [5] BANYAGA, A. and HURTUBISE, D. (2013). *Lectures on Morse homology* **29**. Springer Science & Business Media. [MR2145196](#)
- [6] BAUDRY, J.-P., RAFTERY, A. E., CELEUX, G., LO, K. and GOTTARDO, R. (2010). Combining mixture components for clustering. *Journal of computational and graphical statistics* **19** 332–353. [MR2758307](#)
- [7] BOGACHEV, V. I. (2007). *Measure theory* **1**. Springer Science & Business Media. [MR2267655](#)
- [8] BOISSONNAT, J.-D. and GHOSH, A. (2014). Manifold reconstruction using tangential Delaunay complexes. *Discrete & Computational Geometry* **51** 221–267. [MR3148657](#)
- [9] BOWMAN, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika* **71** 353–360. [MR0767163](#)
- [10] BOYD, S. and VANDENBERGHE, L. (2004). *Convex optimization*. Cambridge university press. [MR2061575](#)
- [11] BRINKMAN, R. R., GASPARETTO, M., LEE, S.-J. J., RIBICKAS, A. J., PERKINS, J., JANSSEN, W., SMILEY, R. and SMITH, C. (2007). High-content flow cytometry and temporal data analysis for defining a cellular signature of graft-versus-host disease. *Biology of Blood and Marrow Transplantation* **13** 691–700.

- [12] CADRE, B. (2006). Kernel estimation of density level sets. *Journal of multivariate analysis* **97** 999–1023. [MR2256570](#)
- [13] CHACÓN, J. E. (2015). A population background for nonparametric density-based clustering. *Statistical Science* **30** 518–532. [MR3432839](#)
- [14] CHAUVEAU, D. and HUNTER, D. (2013). ECM and MM algorithms for normal mixtures with constrained parameters.
- [15] CHAZAL, F. and LIEUTIER, A. (2005). The “ λ -medial axis”. *Graphical Models* **67** 304–331.
- [16] CHAZAL, F. and LIEUTIER, A. (2008). Smooth manifold reconstruction from noisy and non-uniform approximation with guarantees. *Computational Geometry* **40** 156–170. [MR2400541](#)
- [17] CHEN, X., CHRISTENSEN, T. M. and TAMER, E. (2018). Monte Carlo confidence sets for identified sets. *Econometrica* **86** 1965–2018. [MR3901244](#)
- [18] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2015). Asymptotic theory for density ridges. *The Annals of Statistics* **43** 1896–1928. [MR3375871](#)
- [19] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2016). A comprehensive approach to mode clustering. *Electronic Journal of Statistics* **10** 210–241. [MR3466181](#)
- [20] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2017). Density level sets: Asymptotics, inference, and visualization. *Journal of the American Statistical Association* **112** 1684–1696. [MR3750891](#)
- [21] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2017). Statistical inference using the Morse-Smale complex. *Electronic Journal of Statistics* **11** 1390–1433. [MR3635917](#)
- [22] CHENG, G., CHEN, Y.-C. et al. (2019). Nonparametric inference via bootstrapping the debiased estimator. *Electronic Journal of Statistics* **13** 2194–2256. [MR3980957](#)
- [23] CHENG, S.-W., DEY, T. K. and RAMOS, E. A. (2005). Manifold reconstruction from point samples. In *SODA* **5** 1018–1027. [MR2298361](#)
- [24] CHENG, S. W., FUNKE, S., GOLIN, M., KUMAR, P., POON, S. H. and RAMOS, E. (2005). Curve reconstruction from noisy samples In *Computational Geometry* **31**.
- [25] CHERNOZHUKOV, V., HONG, H. and TAMER, E. (2007). Estimation and confidence regions for parameter sets in econometric models 1. *Econometrica* **75** 1243–1284. [MR2347346](#)
- [26] CHOULDECHOVA, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* **5** 153–163.
- [27] CORBETT-DAVIES, S., PIERSON, E., FELLER, A., GOEL, S. and HUQ, A. (2017). Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining* 797–806.
- [28] CUEVAS, A. (2009). Set estimation: Another bridge between statistics and geometry. *Bol. Estad. Investig. Oper* **25** 71–85. [MR2750781](#)
- [29] DENNIS JR, J. E. and SCHNABEL, R. B. (1996). *Numerical methods for unconstrained optimization and nonlinear equations*. SIAM. [MR1376139](#)

- [30] DEY, T. K. (2006). *Curve and surface reconstruction: algorithms with mathematical analysis* **23**. Cambridge University Press. [MR2267420](#)
- [31] DEY, T. K. and GOSWAMI, S. (2006). Provable surface reconstruction from noisy samples. *Computational Geometry* **35** 124–141. [MR2231219](#)
- [32] DRTON, M. and SULLIVANT, S. (2007). Algebraic statistical models. *Statistica Sinica* 1273–1297. [MR2398596](#)
- [33] FEDERER, H. (1959). Curvature measures. *Transactions of the American Mathematical Society* **93** 418–491. [MR0110078](#)
- [34] GARRITY, T. A. (2001). All the mathematics you missed: but need to know for graduate school. [MR1889610](#)
- [35] GENOVESE, C. R., PERONE-PACIFICO, M., VERDINELLI, I. and WASSERMAN, L. (2014). Nonparametric ridge estimation. *The Annals of Statistics* **42** 1511–1545. [MR3262459](#)
- [36] GIBILISCO, P., RICCOMAGNO, E., ROGANTIN, M. P. and WYNN, H. P. (2010). *Algebraic and geometric methods in statistics*. Cambridge University Press. [MR2642656](#)
- [37] GROVE, K. and KARCHER, H. (1973). How to conjugate π_1 -close group actions. *Mathematische Zeitschrift* **132** 11–20. [MR0356104](#)
- [38] HANSEN, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society* 1029–1054. [MR0666123](#)
- [39] HANSEN, L. P. and SINGLETON, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society* 1269–1286. [MR0673689](#)
- [40] HARDT, M., PRICE, E. and SREBRO, N. (2016). Equality of opportunity in supervised learning. In *Advances in neural information processing systems* 3315–3323.
- [41] HORN, R. A. and JOHNSON, C. R. (2012). *Matrix analysis*. Cambridge university press. [MR2978290](#)
- [42] IRWIN, M. (1980). *Smooth dynamical systems*. Academic Press. [MR0586942](#)
- [43] JISU, K., CHEN, Y.-C., BALAKRISHNAN, S., RINALDO, A. and WASSERMAN, L. (2016). Statistical inference for cluster trees. In *Advances in Neural Information Processing Systems* 1839–1847.
- [44] LALOE, T. and SERVIEN, R. (2013). Nonparametric estimation of regression level sets. *Journal of the Korean Statistical Society*.
- [45] LIANG, K.-Y. and ZEGER, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika* **73** 13–22. [MR0836430](#)
- [46] LINDSAY, B. G. (1995). Mixture models: theory, geometry and applications. In *NSF-CBMS regional conference series in probability and statistics* i–163. JSTOR.
- [47] LITTLE, R. J. and RUBIN, D. B. (2019). *Statistical analysis with missing data* **793**. John Wiley & Sons. [MR1925014](#)
- [48] MAMMEN, E. and POLONIK, W. (2013). Confidence regions for level sets. *Journal of Multivariate Analysis* **122** 202–214. [MR3189318](#)
- [49] MANSKI, C. F. (1990). Nonparametric Bounds On Treatment Effects. *The*

- American Economic Review* **80** 319.
- [50] MANSKI, C. F. (1999). *Identification problems in the social sciences*. Harvard University Press.
- [51] MANSKI, C. F. (2003). *Partial identification of probability distributions*. Springer Science & Business Media. [MR2151380](#)
- [52] MATTILA, P. (1999). *Geometry of sets and measures in Euclidean spaces: fractals and rectifiability* **44**. Cambridge university press. [MR1333890](#)
- [53] MCGEHEE, R. (1973). A stable manifold theorem for degenerate fixed points with applications to celestial mechanics. *Journal of Differential Equations* **14** 70–88. [MR0362403](#)
- [54] MCGEHEE, R. and SANDER, E. (1996). A new proof of the stable manifold theorem. *Zeitschrift für angewandte Mathematik und Physik ZAMP* **47** 497–513. [MR1407968](#)
- [55] MICHAŁEK, M., STURMFELS, B., UHLER, C. and ZWIERNIK, P. (2016). Exponential varieties. *Proceedings of the London Mathematical Society* **112** 27–56. [MR3458144](#)
- [56] MOLCHANOV, I. (1991). Empirical estimation of distribution quantiles of random closed sets. *Theory of Probability & Its Applications* **35** 594–600. [MR1091220](#)
- [57] MOLCHANOV, I. S. (1998). A limit theorem for solutions of inequalities. *Scandinavian Journal of Statistics* **25** 235–242. [MR1614288](#)
- [58] NESTEROV, Y. (2018). *Lectures on convex optimization* **137**. Springer. [MR3839649](#)
- [59] NIYOGI, P., SMALE, S. and WEINBERGER, S. (2008). Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry* **39** 419–441. [MR2383768](#)
- [60] PATEIRO LÓPEZ, B. (2008). *Set estimation under convexity type restrictions*. Univ Santiago de Compostela.
- [61] PERKO, L. (2001). *Differential equations and dynamical systems* **7**. Springer Science & Business Media. [MR1801796](#)
- [62] PREISS, D. (1987). Geometry of measures in \mathbb{R}^n : distribution, rectifiability, and densities. *Annals of Mathematics* **125** 537–643. [MR0890162](#)
- [63] QIAO, W. and POLONIK, W. (2019). Nonparametric confidence regions for level sets: Statistical properties and geometry. *Electronic Journal of Statistics* **13** 985–1030. [MR3934621](#)
- [64] RHEINBOLDT, W. C. (1988). On the Computation of Multi-Dimensional Solution Manifolds of Parametrized Equations. *Numerische Mathematik*. [MR0946374](#)
- [65] RICE, J. R. (1967). Nonlinear Approximation. II Curvature in Minkowski Geometry and Local Uniqueness. *Transactions of the American Mathematical Society* **128** 437–459. [MR0226278](#)
- [66] RINALDO, A., SINGH, A., NUGENT, R. and WASSERMAN, L. (2012). Stability of density-based clustering. *The Journal of Machine Learning Research* **13** 905–948. [MR2930628](#)
- [67] RINALDO, A. and WASSERMAN, L. (2010). Generalized density clustering. *The Annals of Statistics* **38** 2678–2722. [MR2722453](#)

- [68] ROMANO, J. P. (1988). Bootstrapping the mode. *Annals of the Institute of Statistical Mathematics* **40** 565–586. [MR0964293](#)
- [69] ROMANO, J. P. (1988). On weak convergence and optimality of kernel density estimates of the mode. *The Annals of Statistics* 629–647. [MR0947566](#)
- [70] ROMANO, J. P. and SHAIKH, A. M. (2010). Inference for the identified set in partially identified econometric models. *Econometrica* **78** 169–211. [MR2642860](#)
- [71] RUDEMO, M. (1982). Empirical choice of histograms and kernel density estimators. *Scandinavian Journal of Statistics* 65–78. [MR0668683](#)
- [72] RUDIN, W. (1964). *Principles of mathematical analysis* **3**. McGraw-hill New York. [MR0166310](#)
- [73] SHEATHER, S. J. (2004). Density estimation. *Statistical Science* **19** 588–597. [MR2185580](#)
- [74] SILVERMAN, B. W. (1986). *Density estimation for statistics and data analysis* **26**. CRC press. [MR0848134](#)
- [75] TAMER, E. (2010). Partial identification in econometrics. *Annu. Rev. Econ.* **2** 167–195.
- [76] TSYBAKOV, A. B. (1997). On nonparametric estimation of density level sets. *The Annals of Statistics* **25** 948–969. [MR1447735](#)
- [77] VAN DER VAART, A. (1998). *Asymptotic Statistics*. Cambridge University Press, Cambridge. [MR1652247](#)
- [78] VIEU, P. (1996). A note on density mode estimation. *Statistics & probability letters* **26** 297–307. [MR1393913](#)
- [79] WALTHER, G. (1997). Granulometric smoothing. *The Annals of Statistics* **25** 2273–2299. [MR1604445](#)