
The Expected Jacobian Outerproduct: Theory and Empirics

S. Trivedi
CS & AI Lab
MIT
Cambridge, MA 02139

J. Wang
Dept. of Computer Science
The University of Chicago
Chicago, IL 60637

Abstract

The expected gradient outerproduct (EGOP) of an unknown regression function is an operator that arises in the theory of multi-index regression, and is known to recover those directions that are most relevant to predicting the output. However, work on the EGOP, including that on its cheap estimators, is restricted to the regression setting. In this work, we adapt this operator to the multi-class setting, which we dub the expected Jacobian outerproduct (EJOP). Moreover, we propose a simple rough estimator of the EJOP and show that somewhat surprisingly, it remains statistically consistent under mild assumptions. Furthermore, we show that the eigenvalues and eigenspaces also remain consistent. Finally, we show that the estimated EJOP can be used as a metric to yield improvements in real-world non-parametric classification tasks: both by its use as a metric, and also as cheap initialization in metric learning tasks.

1 Introduction

In high-dimensional classification and regression problems, the task is to infer the unknown function f with $Y \approx f(X)$, given a set of observations $(\mathbf{x}, \mathbf{y})_i, i = 1, 2, \dots, n$, with $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d$ and labels \mathbf{y}_i being noisy versions of the function values $f(\mathbf{x}_i)$. We are interested in distance based (non-parametric) regression, which provides our function estimate $f_n(\mathbf{x}) = \sum_{i=1}^n w(\mathbf{x}, \mathbf{x}_i)\mathbf{y}_i$, where $w(\mathbf{x}, \mathbf{x}_i)$ depends on the distance ρ , defined as $\rho(\mathbf{x}, \mathbf{x}') = \sqrt{(\mathbf{x} - \mathbf{x}')\mathbf{W}(\mathbf{x} - \mathbf{x}')}}$ and $\mathbf{W} \succeq 0$. The problem of estimating unknown f becomes significantly harder as d increases due to the *curse of dimensionality*. To remedy this situation, several pre-processing techniques are used, each of which rely on a suitable assumption about the data and/or about f . For instance, a conceptually simple, yet often reasonable assumption that can be made is that f might not vary equally along all coordinates of \mathbf{x} . Letting $f'_i = \nabla f^T e_i$ denote the derivative along coordinate i , and $\|f'_i\|_{1,\mu} \equiv \mathbb{E}_{\mathbf{x} \sim \mu} f'_i(\mathbf{x})$, we can use the above distance based estimator by setting ρ such that $\mathbf{W}_{i,j} = \|f'_i\|_{1,\mu}$ when $i = j$ and 0 otherwise. This *gradient weighting* rescales the space such the ball \mathcal{B}_ρ contains more points relative to the Euclidean ball \mathcal{B} . This is the intuition pursued in works such as [6], [7], with an emphasis on deriving an efficient, yet consistent estimator for the gradient. Using gradient weights for coordinate scaling in this manner has strong theoretical grounding: it has the effect of reducing the regression variance, while keeping the bias in control. The same assumption that f may not equally vary along all coordinates is also the motivation for a plethora of variable selection methods. In the simplest setting, we use $f(X) = g(PX)$, where $P \in \{0, 1\}^{k \times d}$ projects X down to $k < d$ coordinates that are most relevant to predicting the output. This idea is relaxed further in *multi-index regression* e.g. [8, 13, 2, 23]), by letting $P \in \mathbb{R}^{k \times d}$, which projects X down to a k -dimensional subspace of \mathbb{R}^d . The motivation for multi-index regression is that while that while f might vary along all coordinates, it actually may only depend on an unknown k -dimensional subspace, called a *relevant* subspace. The task then becomes finding the said relevant subspace rather than chopping coordinates since

they all might be relevant in predicting the output y . Work to recover this relevant subspace (which is sometimes also referred to in the literature as *effective dimension reduction* [8]) gives rise to the expected gradient outerproduct (EGOP) [17]: $\mathbb{E}_X G(X) \triangleq \mathbb{E}_X (\nabla f(X) \cdot \nabla f(X)^\top)$. Notice that if f does not vary along some direction $\mathbf{v} \in \mathbb{R}^d$, then \mathbf{v} must lie in the nullspace of the operator. Furthermore, it is easy to see that under mild assumptions on f , the column space of the EGOP is exactly the relevant subspace described above. Since the EGOP recovers the average variation of f in all directions, it is also useful beyond the multi-index motivation. That is, even in the case when there isn't a clear relevant subspace, it is still reasonable to assume that f may not vary equally in all directions. Thus, the EGOP can be used to weigh any direction $v \in \mathbb{R}^d$ according to its relevance as captured by the average variation of f along v . Letting VDV^\top be the spectral decomposition of the estimated EGOP, we can use it to transform the input \mathbf{x} as $D^{1/2}V^\top \mathbf{x}$, which can be used as a distance function in non-parametric regression tasks.

However, all prior work on multi-index regression and cognate topics only revolves around regression (and binary classification), in part due to the complexity of analysis. In this work we attack the more general multi-class case, which is treated as a multinomial regression problem with c outputs, with the unknown function denoted as $f : \mathbb{R}^d \rightarrow \mathbb{S}^c$ where $\mathbb{S}^c = \{\mathbf{y} \in \mathbb{R}^c | \mathbf{y} \geq 0, \mathbf{y}^\top \mathbf{1} = 1\}$. We are led to an operator similar to the EGOP based on computing the Jacobian of f , which we call the Expected Jacobian Outer Product (EJOP) $\mathbb{E}_X G(X) \triangleq \mathbb{E}_{\mathbf{x}} (\mathbf{J}_f(X) \cdot \mathbf{J}_f(X)^\top)$. For constructing the EJOP, we need to compute gradient estimates, for which optimal estimators can be expensive in practice. We propose a simple, efficient, difference based estimator and show that despite its crudity it remains statistically consistent under mild assumptions. This approach is also online and cheap: we only require $2d$ estimates of the function f at \mathbf{x} . We also show that the EJOP can be used for metric weighing for distance-based non-parametric classification, as well as used as a pre-processing metric for standard metric learning tasks.

2 The Expected Jacobian Outerproduct

Recall that in high dimensional classification problems over \mathbb{R}^d , the unknown (multinomial regression) function f is a vector-valued function mapping to a probability simplex $\mathbb{S}^c = \{\mathbf{y} \in \mathbb{R}^c | \forall i y_i \geq 0, \mathbf{y}^\top \mathbf{1} = 1\}$, where c is the number of outputs. For classification, the prediction for some point \mathbf{x} is then given by: $y = \text{argmax}_{i=1,\dots,c} f_i(\mathbf{x})$. For f , at point \mathbf{x} , denoting the Jacobian as \mathbf{J}_f , we are interested in the quantity $\mathbf{J}_f(\mathbf{x})\mathbf{J}_f(\mathbf{x})^\top$. Let f_n be an initial estimate of f , for which we use a kernel estimate, then for the $(i, j)^{\text{th}}$ element of $\mathbf{J}_f(\mathbf{x})$, we can use the following rough estimate:

$$\Delta_{t,i,j} f_n(\mathbf{x}) = \frac{f_{n,i}(\mathbf{x} + t\mathbf{e}_j) - f_{n,i}(\mathbf{x} - t\mathbf{e}_j)}{2t}, t > 0$$

Let $G_n(\mathbf{x})$ be the outerproduct using the estimated Jacobian, the EJOP is estimated as $\mathbb{E}_n G_n(X)$.

2.1 Function Estimate

First, we need to specify the function estimate, that is used both for the theoretical analysis, and the experiments reported. We denote the vector-valued function estimate by $\bar{f}_{n,h}(\mathbf{x}) \in \mathbb{S}^c$, for which we employ the estimate using an admissible kernel K : $\bar{f}_{n,h,c}(\mathbf{x}) = \sum_i w_i \mathbb{1}\{Y_i = c\}$, where:

$$w_i(\mathbf{x}) = \frac{K(\|\mathbf{x} - \mathbf{x}_i\|/h)}{\sum_j K(\|\mathbf{x} - \mathbf{x}_j\|/h)} \text{ if } \mathcal{B}(\mathbf{x}, h) \cap \mathbf{x} \neq \emptyset,$$

$$w_i(\mathbf{x}) = \frac{1}{n} \text{ otherwise}$$

Note that for k -NN, $w_i(\mathbf{x}) = \frac{1}{k}$ and for h -NN, $w_i(\mathbf{x}) = \frac{1}{|\mathcal{B}(\mathbf{x}, h)|}$. While estimating gradients, we actually work with the softmaxed output

$$\bar{f}_{n,h,i}(\mathbf{x}) = \frac{\exp(\bar{f}_{n,h,i}(\mathbf{x}))}{\sum_j \exp(\bar{f}_{n,h,j}(\mathbf{x}))}$$

3 Notation and Setup

For a vector $x \in \mathbb{R}^d$, we denote the euclidean norm as $\|x\|$. For a matrix, we denote the spectral norm, which is the largest singular value of the matrix $\sigma_{\max}(A)$ as $\|A\|_2$. The column space of a matrix $A \in R^{n \times m}$ is denoted as $\text{im}(A)$ where $\text{im}(A) = \{\mathbf{Y} \in \mathbb{R}^n | \mathbf{Y} = Ax \text{ for some } \mathbf{x} \in \mathbb{R}^m\}$, and $\ker(A)$ is used to denote the null space of matrix $A \in R^{n \times m}$: $\ker(A) = \{\mathbf{x} \in \mathbb{R}^m | Ax = 0\}$. We use $A \circ B$ to denote the Hadamard product of matrices A and B . Let the estimated nonparametric function be $f_{n,h,c}(\mathbf{x}) = \sum_i \omega_i(\mathbf{x}) \mathbb{1}\{y_i = c\}$, and $\tilde{f}_{n,h,c}(\mathbf{x}) = \sum_i \omega_i(\mathbf{x}) \mathbb{P}(y_i = c | x_i)$. Our estimated gradient at dimension i is given as

$$\Delta_{t,i} f_{n,h,c}(\mathbf{x}) = \frac{f_{n,h,c}(\mathbf{x} + te_i) - f_{n,h,c}(\mathbf{x} - te_i)}{2t},$$

and the estimated and true gradients for class c are given as:

$$\hat{\nabla} f_{n,h,c}(\mathbf{x}) = \begin{bmatrix} \Delta_{t,1} f_{n,h,c}(\mathbf{x}) \cdot \mathbb{1}_{A_{n,1}(\mathbf{x})} \\ \Delta_{t,2} f_{n,h,c}(\mathbf{x}) \cdot \mathbb{1}_{A_{n,2}(\mathbf{x})} \\ \dots \\ \Delta_{t,d} f_{n,h,c}(\mathbf{x}) \cdot \mathbb{1}_{A_{n,d}(\mathbf{x})} \end{bmatrix}, \hat{\nabla} f_c(\mathbf{x}) = \begin{bmatrix} \Delta_{t,1} f_c(\mathbf{x}) \cdot \mathbb{1}_{A_{n,1}(\mathbf{x})} \\ \Delta_{t,2} f_c(\mathbf{x}) \cdot \mathbb{1}_{A_{n,2}(\mathbf{x})} \\ \dots \\ \Delta_{t,d} f_c(\mathbf{x}) \cdot \mathbb{1}_{A_{n,d}(\mathbf{x})} \end{bmatrix}$$

Where $A_{n,i}(X)$ is the event that enough samples contribute to the estimate $\Delta_{t,i} f_{n,h}(X)$:

$$A_{n,i}(X) \equiv \min_{\{t,-t\}} \mu_n(B(X + se_i, h/2)) \geq \frac{2d \ln 2n + \ln(4/\delta)}{n}$$

$$A_i(X) \equiv \min_{\{t,-t\}} \mu(B(X + se_i, h/2)) \geq 3 \cdot \frac{2d \ln 2n + \ln(4/\delta)}{n}$$

μ_n, μ are empirical mass and mass of a ball, respectively. We denote indicators for events $A_{n,i}(X)$ and $A_i(X)$ as:

$$\mathbb{I}_n(x) = \begin{bmatrix} \mathbb{1}_{A_{n,1}(\mathbf{x})} \\ \mathbb{1}_{A_{n,2}(\mathbf{x})} \\ \dots \\ \mathbb{1}_{A_{n,d}(\mathbf{x})} \end{bmatrix}, \overline{\mathbb{I}_n(x)} = \begin{bmatrix} \mathbb{1}_{\bar{A}_{n,1}(\mathbf{x})} \\ \mathbb{1}_{\bar{A}_{n,2}(\mathbf{x})} \\ \dots \\ \mathbb{1}_{\bar{A}_{n,d}(\mathbf{x})} \end{bmatrix}, \mathbb{I}(x) = \begin{bmatrix} \mathbb{1}_{A_1(\mathbf{x})} \\ \mathbb{1}_{A_2(\mathbf{x})} \\ \dots \\ \mathbb{1}_{A_d(\mathbf{x})} \end{bmatrix}, \overline{\mathbb{I}(x)} = \begin{bmatrix} \mathbb{1}_{\bar{A}_1(\mathbf{x})} \\ \mathbb{1}_{\bar{A}_2(\mathbf{x})} \\ \dots \\ \mathbb{1}_{\bar{A}_d(\mathbf{x})} \end{bmatrix}.$$

The Jacobian outer product matrix is $G(\mathbf{x}) = \mathbf{J}_f(\mathbf{x}) \mathbf{J}_f(\mathbf{x})^T$, with the estimated Jacobian matrix being:

$$\hat{\mathbf{J}}_f(\mathbf{x}) = [\hat{\nabla} f_{n,h,1}(\mathbf{x}) \quad \hat{\nabla} f_{n,h,2}(\mathbf{x}) \quad \dots \quad \hat{\nabla} f_{n,h,k}(\mathbf{x})]$$

the EJOP is denoted $\hat{G}(\mathbf{x}) = \hat{\mathbf{J}}_f(\mathbf{x}) \hat{\mathbf{J}}_f(\mathbf{x})^T$.

3.1 Assumptions

We make minimalistic assumptions: Simply that f is continuously differentiable and μ has lower-bounded density on a compact support \mathcal{X} . All of the detailed assumptions below will then hold:

- Noise: Let $\eta(X) \triangleq Y - f(X)$. We assume $\forall \delta > 0$ there exists $p > 0$ such that $\sup_{x \in \mathbf{X}} \mathbb{P}_{Y|X=x}(|\eta(x)| > p) \leq \delta$. The infimum over all p is denoted by $C_Y(\delta)$. Moreover, the variance of $(Y|X=x)$ is upper-bounded by a constant σ_Y^2 uniformly over $x \in \mathbf{X}$.
- Bounded Gradient: Defining the τ -envelope of \mathcal{X} as $\mathcal{X} + B(0, \tau) \triangleq \{z \in B(x, \tau), x \in \mathcal{X}\}$. We assume there exists τ such that f is continuously differentiable on the τ -envelope $\mathcal{X} + B(0, \tau)$. Furthermore, for all $x \in \mathcal{X} + B(0, \tau)$, $k \in [c]$, we have $\|\nabla f_k(x)\| \leq R$ for some $R > 0$, and ∇f is uniformly continuous on $\mathcal{X} + B(0, \tau)$.
- Modulus of continuity of ∇f_k : Let $\epsilon_{t,k,i} = \sup_{\mathbf{x} \in \mathcal{X}, s \in [-t, t]} \left| \frac{\partial f_k(\mathbf{x})}{\partial x_i} - \frac{\partial f_k(\mathbf{x} + se_i)}{\partial x_i} \right|$ and $\epsilon_{t,i} = \max_k \epsilon_{t,c,i}$, define the (t, i) -boundary of \mathcal{X} as $\partial_{t,i}(\mathcal{X}) = \{\mathbf{x} : \{\mathbf{x} + te_i, \mathbf{x} - te_i\} \not\subseteq \mathcal{X}\}$. When μ has continues density on \mathcal{X} and ∇f_k is uniformly continuous on $\mathcal{X} + B(0, \tau)$, we have $\mu(\partial_{t,i}(\mathcal{X})) \xrightarrow{t \rightarrow 0} 0$ and $\epsilon_{t,k,i} \xrightarrow{t \rightarrow 0} 0$.

4 Consistency of Estimator $\mathbb{E}_n \hat{G}(X)$ of the Jacobian Outerproduct $\mathbb{E}_X G(X)$

To show that the estimator $\mathbb{E}_n \hat{G}(X)$ is consistent, we proceed to bound $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2$ for finite n , which is encapsulated in the theorem that follows. There are two main difficulties in the proof, which are addressed by a sequence of lemmas. One has to do with the fact that the gradient estimate at any point depends on all other points, and second, having gradient estimates for c classes.

MAIN RESULT

Theorem 1. *Let $t + h \leq \tau$, and let $0 \leq \delta \leq 1$. There exist $C = C(\mu, K(\cdot))$ and $N = N(\mu)$ such that the following holds with probability at least $1 - 2\delta$. Define $A(n) = \sqrt{Cd \cdot \log(kn/\delta)} \cdot 0.25/\log^2(n/\delta)$. Let $n \geq N$, we have:*

$$\begin{aligned} \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2 &\leq \frac{6R^2}{\sqrt{n}} \left(\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}} \right) + k \left(3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR+1}{t} \right) \right). \\ &\left[\frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^4} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} \right] \end{aligned}$$

Proof. We begin with the following decomposition:

$$\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2 \leq \|\mathbb{E}_n G(X) - \mathbb{E}_X G(X)\|_2 + \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2$$

The first term on the right hand side i.e. $\|\mathbb{E}_n G(X) - \mathbb{E}_X G(X)\|_2$ is bounded using Lemma 2; by using Lemma 3 we bound the second term $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2$, this is done with respect to $\sum_{k \in [c]} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$; therefore we need to bound $\sum_{k \in [c]} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$, which is done by employing Theorem 2 which concludes the proof. \square

Note that consistency is implied for $t \xrightarrow{n \rightarrow \infty} 0$, $h \xrightarrow{n \rightarrow \infty} 0$, $h/t^2 \xrightarrow{n \rightarrow \infty} 0$, and $(n/\log n)h^dt^4 \xrightarrow{n \rightarrow \infty} \infty$, this is satisfied for many settings, for example $t \propto h^{1/4}$, $h \propto \frac{1}{\ln n}$.

4.1 Bounding $\|\mathbb{E}_n G(X) - \mathbb{E}_X G(X)\|_2$

To bound this term, we use the following random matrix concentration result.

Lemma 1. [18, 3]. *For the random matrix $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$ with bounded spectral norm $\|\mathbf{X}\|_2 \leq M$, let $d = \min\{d_1, d_2\}$, and $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ are i.i.d. samples, with probability at least $1 - \delta$, we have*

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i - \mathbb{E} \mathbf{X} \right\|_2 \leq \frac{6M}{\sqrt{n}} \left(\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}} \right)$$

Using the bounded gradient assumption we can apply the above lemma to i.i.d matrices $G(X), X \in \mathbf{X}$, yielding the following lemma.

Lemma 2. *With probability at least $1 - \delta$ over i.i.d sample X*

$$\|\mathbb{E}_n G(X) - \mathbb{E}_X G(X)\|_2 \leq \frac{6R^2}{\sqrt{n}} \left(\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}} \right)$$

Next we proceed to bound the second term in the decomposition mentioned in the proof of theorem 1.

4.2 Bounding $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2$

We derive a first bound by the following lemma:

Lemma 3. *There exists a constant c , such that with probability at least $1 - \delta$:*

$$\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2 \leq \sum_{k \in [c]} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \cdot \max_{x \in \mathbf{X}} \|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2$$

Proof. The term on the l.h.s can be written in terms of the gradients for each class:

$$\begin{aligned} \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2 &= \|\mathbb{E}_n [\hat{G}(X) - G(X)]\|_2 \\ &= \left\| \sum_{k \in [c]} \mathbb{E}_n [\nabla f_k(X) \cdot \nabla f_k(X)^T - \hat{\nabla} f_{n,h,k}(X) \cdot \hat{\nabla} f_{n,h,k}(X)^T] \right\|_2 \\ &\leq \sum_{k \in [c]} \left\| \mathbb{E}_n [\nabla f_k(X) \cdot \nabla f_k(X)^T - \hat{\nabla} f_{n,h,k}(X) \cdot \hat{\nabla} f_{n,h,k}(X)^T] \right\|_2 \end{aligned}$$

Note that $\nabla f_k(\mathbf{x}) \cdot \nabla f_k(\mathbf{x})^T - \hat{\nabla} f_{n,h,k}(\mathbf{x}) \cdot \hat{\nabla} f_{n,h,k}(\mathbf{x})^T$ may be rewritten as:

$$\begin{aligned} \nabla f_k(\mathbf{x}) \cdot \nabla f_k(\mathbf{x})^T - \hat{\nabla} f_{n,h,k}(\mathbf{x}) \cdot \hat{\nabla} f_{n,h,k}(\mathbf{x})^T &= \frac{1}{2} \cdot (\nabla f_k(\mathbf{x}) + \hat{\nabla} f_{n,h,k}(\mathbf{x})) \cdot (\nabla f_k(\mathbf{x}) - \hat{\nabla} f_{n,h,k}(\mathbf{x}))^T \\ &+ \frac{1}{2} \cdot (\nabla f_k(\mathbf{x}) - \hat{\nabla} f_{n,h,k}(\mathbf{x})) \cdot (\nabla f_k(\mathbf{x}) + \hat{\nabla} f_{n,h,k}(\mathbf{x}))^T \end{aligned}$$

Using this yields:

$$\begin{aligned} \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2 &\leq \frac{1}{2} \sum_{k \in [c]} \|\mathbb{E}_n [(\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X))^T]\|_2 \\ &+ \frac{1}{2} \sum_{k \in [c]} \|\mathbb{E}_n [(\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X))^T]\|_2 \\ &= \sum_{k \in [c]} \|\mathbb{E}_n [(\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X))^T]\|_2 \end{aligned}$$

Employing Jensen's inequality:

$$\begin{aligned} \mathbb{E}_n [(\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X))^T] &\|_2 \\ &\leq \mathbb{E}_n \|(\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X))^T\|_2 \end{aligned}$$

combining the above, gives us the following bound on $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2$

$$\begin{aligned} \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_n G(X)\|_2 &\leq \sum_{k \in [c]} \mathbb{E}_n \|(\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)) \cdot (\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X))^T\|_2 \\ &= \sum_{k \in [c]} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \cdot \|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2. \\ &\leq \sum_{k \in [c]} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \cdot \max_{X \in \mathbf{X}} \|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2. \end{aligned}$$

□

The above bound has a dependence on $\|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2$, which we now proceed to bound below:

4.3 Bounding $\|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2$

We first bound the max term, by the following lemma:

Lemma 4. $\forall c \in [k]$, we have

$$\max_{X \in \mathbf{X}} \|\nabla f_k(X) + \hat{\nabla} f_{n,h,k}(X)\|_2 \leq 3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR + 1}{t} \right)$$

Proof. $\forall x \in \mathbf{X}$, we have

$$\begin{aligned} \|\nabla f_k(\mathbf{x}) + \hat{\nabla} f_{n,h,k}(\mathbf{x})\|_2 &\leq \|\nabla f_k(\mathbf{x})\|_2 + \|\hat{\nabla} f_{n,h,k}(\mathbf{x})\|_2 \leq 2\|\nabla f_k(\mathbf{x})\|_2 + \|\nabla f_k(\mathbf{x}) - \hat{\nabla} f_{n,h,k}(\mathbf{x})\|_2 \\ &\leq 2R + \|\nabla f_k(\mathbf{x}) - \hat{\nabla} f_k(\mathbf{x})\|_2 + \|\hat{\nabla} f_k(\mathbf{x}) - \hat{\nabla} f_{n,h,k}(\mathbf{x})\|_2 \end{aligned}$$

Next, we adopt the steps as in the proof for Lemma 9, and get the following bound:

$$\|\hat{\nabla} f_k(\mathbf{x}) - \hat{\nabla} f_{n,h,k}(\mathbf{x})\|_2 \leq \sqrt{\sum_{i \in [d]} (|\Delta_{t,i} f_{n,h,k}(\mathbf{x}) - \Delta_{t,i} f_k(\mathbf{x})| \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})})^2},$$

this is because

$$\begin{aligned} |\Delta_{t,i} f_{n,h,k}(\mathbf{x}) - \Delta_{t,i} f_k(\mathbf{x})| \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})} &\leq \frac{1}{t} \max_{s \in \{-t,t\}} |\tilde{f}_{n,h,k}(\mathbf{x} + s e_i) - f_k(\mathbf{x} + s e_i)| \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})} \\ &+ \frac{1}{t} \max_{s \in \{-t,t\}} |\tilde{f}_{n,h,k}(\mathbf{x} + s e_i) - f_{n,h,k}(\mathbf{x} + s e_i)| \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})}, \end{aligned}$$

we also know that

$$\max_{s \in \{-t,t\}} |\tilde{f}_{n,h,k}(\mathbf{x} + s e_i) - f_{n,h,k}(\mathbf{x} + s e_i)| \leq 1.$$

Thus we obtain the following bound:

$$\|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \sqrt{d} \left(\frac{hR + 1}{t} \right)$$

While, we also have that

$$\|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 \leq \|\nabla f_k(X) \circ \mathbf{I}_n(X) - \hat{\nabla} f_k(X)\|_2 + \|\nabla f_k(X) \circ \overline{\mathbf{I}_n(X)}\|_2 \leq R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2}$$

Combining the above completes the proof \square

Next we need to bound $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$, which we do so in the next subsection:

4.4 Bound on $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$

We first decompose $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$ as:

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 + \mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$$

the first term in the r.h.s of the above i.e. $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2$ can in turn be decomposed as:

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 \leq \mathbb{E}_n \|\nabla f_k(X) \circ \mathbf{I}_n(X) - \hat{\nabla} f_k(X)\|_2 + \mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbf{I}_n(X)}\|_2$$

We need to bound both terms that appear on the r.h.s of the above, which we do so in the next two subsections, starting with the second term.

4.4.1 Bounding $\mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2$

Lemma 5. *With probability at least $1 - \delta$ over the choice of X :*

$$\mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2 \leq R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right)$$

Proof. We begin by recalling the bounded gradient assumption: $\|\nabla f(\mathbf{x})\|_2 \leq R$, using which we get

$$\mathbb{E}_n \|\nabla f(X) \circ \overline{\mathbb{I}_n(X)}\|_2 \leq R \mathbb{E}_n \|\overline{\mathbb{I}_n(X)}\|_2$$

By relative VC bounds [19], if we set $\alpha_n = \frac{2d \ln 2n + \ln(4/\delta)}{n}$, then with probability at least $1 - \delta$ over the choice of X , for all balls $B \in R^d$ we have $\mu(B) \leq \mu_n(B) + \sqrt{\mu_n(B)\alpha_n} + \alpha_n$. Thus, with probability at least $1 - \delta$, $\forall i \in [d]$, $\bar{A}_{n,i}(X) \Rightarrow \bar{A}_i(X)$. Moreover, since $\|\overline{\mathbb{I}(X)}\|_2 \leq \sqrt{d}$, then by Hoeffding's inequality,

$$\mathbb{P}(\mathbb{E}_n \|\overline{\mathbb{I}(X)}\|_2 - \mathbb{E}_X \|\overline{\mathbb{I}(X)}\|_2 \geq \epsilon) \leq e^{-\frac{2n\epsilon^2}{d}}$$

applying the union bound, we have the following with probability at least $1 - \delta$

$$\mathbb{E}_n \|\overline{\mathbb{I}_n(X)}\|_2 \leq \mathbb{E}_n \|\overline{\mathbb{I}(X)}\|_2 \leq \mathbb{E}_X \|\overline{\mathbb{I}_n(X)}\|_2 + \sqrt{\frac{d \ln \frac{d}{\delta}}{2n}}$$

But note that we have: $\mathbb{E}_X \mathbb{1}_{\bar{A}_i(X)} \leq \mathbb{E}_X [\mathbb{1}_{\bar{A}_i(X)} | X \in \mathcal{X} \setminus \partial_{t,i}(\mathcal{X})] + \mu(\partial_{t,i}(\mathcal{X}))$ to see why this is true observe that $\mathbb{E}_X [\mathbb{1}_{\bar{A}_i(X)} | X \in \mathcal{X} \setminus \partial_{t,i}(\mathcal{X})] = 0$ because $\mu(B(x + se_i, h/2)) \geq C_\mu(h/2)^d \geq 3\alpha$ when we set $h \geq (\log^2(n/\delta)/n)^{1/d}$. So, we have:

$$\mathbb{E}_X \|\overline{\mathbb{I}_n(X)}\|_2 \leq \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))}$$

Thus with probability at least $1 - \delta$, we obtain the following:

$$\mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2 \leq R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right)$$

□

Next we need to bound the first term that appeared on the r.h.s. of the decomposition of $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2$, reproduced below for ease of exposition:

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 \leq \mathbb{E}_n \|\nabla f_k(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f_k(X)\|_2 + \mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2$$

4.4.2 Bounding $\mathbb{E}_n \|\nabla f_k(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f_k(X)\|_2$

This bound is encapsulated in the following lemma

Lemma 6.

$$\mathbb{E}_n \|\nabla f_k(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f_k(X)\|_2 \leq \sqrt{\sum_{i \in [d]} \epsilon_{t,c,i}^2}$$

Proof. We start with the simple observation regarding the envelope:

$$f_k(\mathbf{x} + te_i) - f_k(\mathbf{x} - te_i) = \int_{-t}^t \frac{\partial f_k(\mathbf{x} + se_i)}{\partial x_i} ds$$

using this we have

$$2t \left(\frac{\partial f'_k(\mathbf{x})}{\partial x_i} - \epsilon_{t,k,i} \right) \leq f_k(\mathbf{x} + te_i) - f_k(\mathbf{x} - te_i) \leq 2t \left(\frac{\partial f'_k(\mathbf{x})}{\partial x_i} + \epsilon_{t,k,i} \right)$$

Thus we have

$$\left| \frac{1}{2t} (f_c(\mathbf{x} + te_i) - f_c(\mathbf{x} - te_i)) - \frac{\partial f'_c(\mathbf{x})}{\partial x_i} \right| \leq \epsilon_{t,c,i}$$

using which we obtain the following

$$\begin{aligned} \|\nabla f_k(\mathbf{x}) \circ \mathbb{I}_n(x) - \hat{\nabla} f_k(\mathbf{x})\|_2 &= \sqrt{\sum_{i=1}^d \left| \frac{\partial f'_k(\mathbf{x})}{\partial x_i} \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})} - \Delta_{t,i} f_k(\mathbf{x}) \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})} \right|^2} \\ &\leq \sqrt{\sum_{i=1}^d \left| \frac{1}{2t} (f_k(\mathbf{x} + te_i) - f_k(\mathbf{x} - te_i)) - \frac{\partial f'_k(\mathbf{x})}{\partial x_i} \right|^2} \leq \sqrt{\sum_{i \in [d]} \epsilon_{t,k,i}^2} \end{aligned}$$

Taking empirical expectation on both sides finishes the proof. \square

Taking a step back, recall again the decomposition of $\mathbb{E}_n \|\nabla f_c(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$:

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 + \mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$$

the first term in the r.h.s of the above i.e. $\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2$ was in turn decomposed as:

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_k(X)\|_2 \leq \mathbb{E}_n \|\nabla f_k(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f_k(X)\|_2 + \mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2$$

The analysis in the previous subsection was bounding these two terms individually. Now we turn our attention towards bounding $\mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2$

4.4.3 Bounding $\mathbb{E}_n \|\hat{\nabla} f(X) - \hat{\nabla} f_{n,h}(X)\|_2$

First we introduce a lemma which is a modification of Lemma 6 appearing in [6]

Lemma 7. *Let $t + h \leq \tau$. We have for all $i \in [d]$, and all $s \in \{-t, t\}$: $|\tilde{f}_{n,h,c}(\mathbf{x} + se_i) - f_c(\mathbf{x} + se_i)| \cdot \mathbb{1}_{A_{n,i}(\mathbf{x})} \leq hR$*

Proof. The proof follows the same logic as in [6], with the last step modified appropriately. To be more specific, let $x = X + se_i$, let $v_i = \frac{X_i - x}{\|X_i - x\|_2}$, then we have

$$\begin{aligned} |\tilde{f}_{n,h,c}(\mathbf{x} + se_i) - f_c(\mathbf{x} + se_i)| &\leq \sum_{i \in [d]} w_i(x) |f(X_i) - f(x)| \\ &= \sum_{i \in [d]} w_i(x) \left| \int_0^{\|X_i - x\|_2} v_i^T \nabla f(x + tv_i) dt \right| \leq \sum_{i \in [d]} w_i(x) \|X_i - x\|_2 \cdot \max_{x' \in \mathcal{X} + B(0, \tau)} \|v_i^T \nabla f(x)\|_2 \\ &\leq \sum_{i \in [d]} w_i(x) \|X_i - x\|_2 \leq hR \end{aligned}$$

\square

Lemma 8. *There exist a constant $C = C(\mu, K(\cdot))$, such that the following holds with probability at least $1 - 2\delta$ over the choice of X . Define $A(n) = 0.25 \cdot \sqrt{Cd \cdot \ln(kn/\delta)}$, for all $i \in [d], k \in [c]$, and all $s \in \{-t, t\}$:*

$$\mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_{n,h,k}(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}(X)} \leq \frac{A(n)}{nh^d}$$

Proof. The proof follows a similar line of argument as made for the proof of Lemma 7 in [6]. First fix any $k \in [c]$, Assume $A_{n,i}(X)$ is true, and fix $x = X + se_i$. Taking conditional expectation on $\mathbf{Y}^n = Y_1, \dots, Y_n$ given $\mathbf{X}^n = X_1, \dots, X_n$, we have

$$\mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n} |f_{n,h,k}(x) - \tilde{f}_{n,h,k}(x)|^2 \leq 0.25 \cdot \sum_{i \in [n]} (w_i(x))^2 \leq 0.25 \cdot \max_{i \in [n]} w_i(x)$$

Use \mathbf{Y}_x^n to denote corresponding Y_i of samples $X_i \in B(x, h)$.

Next, we consider the random variable

$$\psi(\mathbf{Y}_x^n) = |f_{n,h,k}(x) - \tilde{f}_{n,h,k}(x)|^2$$

Let \mathcal{Y}_δ denote the event that for all $Y_i \in \mathbf{Y}^n$, $|Y_i - f(X_i)|^2 \leq 0.25$. We know \mathcal{Y}_δ happens with probability at least $1/2$. Thus

$$\begin{aligned} \mathbb{P}_{\mathbf{Y}^n | \mathbf{X}^n} (\psi(\mathbf{Y}_x^n) > 2\mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n} \psi(\mathbf{Y}_x^n) + \epsilon) &\leq \mathbb{P}_{\mathbf{Y}^n | \mathbf{X}^n} (\psi(\mathbf{Y}_x^n) > \mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n, \mathcal{Y}_\delta} \psi(\mathbf{Y}_x^n) + \epsilon) \\ &\leq \mathbb{P}_{\mathbf{Y}^n | \mathbf{X}^n, \mathcal{Y}_\delta} (\psi(\mathbf{Y}_x^n) > \mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n, \mathcal{Y}_\delta} \psi(\mathbf{Y}_x^n) + \epsilon) + \delta/2 \end{aligned}$$

By McDiarmid's inequality, we have

$$\mathbb{P}_{\mathbf{Y}^n | \mathbf{X}^n, \mathcal{Y}_\delta} (\psi(\mathbf{Y}_x^n) > \mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n, \mathcal{Y}_\delta} \psi(\mathbf{Y}_x^n) + \epsilon) \leq \exp \left\{ -2\epsilon^2 \cdot \delta_Y^4 \sum_{i \in [n]} w_i^4(x) \right\}$$

The number of possible sets \mathbf{Y}_x^n (over $x \in \mathcal{X}$) is at most the n -shattering number of balls in \mathbb{R}^d , using Sauer's lemma we get the number is bounded by $(2n)^{d+2}$. By union bound, with probability at least $1 - \delta$, for all $x \in \mathcal{X}$ satisfying $B(x, h/2) \cap \mathbf{X}^n \neq \emptyset$,

$$\begin{aligned} \psi(\mathbf{Y}_x^n) &\leq 2\mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n} \psi(\mathbf{Y}_x^n) + \sqrt{0.25(d+2) \log(n/\delta) \sum_{i \in [n]} w_i^4(x)} \leq 2\sqrt{\mathbb{E}_{\mathbf{Y}^n | \mathbf{X}^n} \psi^2(\mathbf{Y}_x^n)} \\ &+ \sqrt{0.25(d+2) \log(n/\delta) \delta_Y^4 \max_{i \in [n]} w_i^2(x)} \leq \sqrt{Cd \cdot \log(n/\delta) \cdot 0.25/n^2 \mu_n^2(B(x, h/2))} \end{aligned}$$

Take a union bound over $k \in [c]$, and take empirical expectation, we get $\forall k \in [c]$

$$\mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_{n,h,k}(X + se_i)|^2 \leq \frac{0.25 \cdot \sqrt{Cd \cdot \ln(cn/\delta)}}{n} \sum_{i \in [n]} \frac{1}{n(x_i, h/2)}$$

where $n(x_i, h/2) = n\mu_n(B(x_i, h/2))$ is the number of points in Ball $B(x_i, h/2)$.

Let \mathcal{Z} denote the minimum $h/4$ cover of $\{x_1, \dots, x_n\}$, which means for any x_i , there is a $z \in \mathcal{Z}$, such that x_i is contained in the ball $B(z, h/4)$. Since $x_i \in B(z, h/4)$, we have $B(z, h/4) \in B(x_i, h/2)$. We also assume every x_i is assigned to the closest $z \in \mathcal{Z}$, and write $x_i \rightarrow z$ to denote such x_i . Then we have:

$$\begin{aligned} \sum_{i \in [n]} \frac{1}{n(x_i, h/2)} &= \sum_{z \in \mathcal{Z}} \sum_{x_i \rightarrow z} \frac{1}{n(x_i, h/2)} \leq \sum_{z \in \mathcal{Z}} \sum_{x_i \rightarrow z} \frac{1}{n(z, h/4)} \\ &\leq \sum_{z \in \mathcal{Z}} \frac{n(z, h/4)}{n(z, h/4)} = |\mathcal{Z}| \leq C_\mu (h/4)^{-d} \end{aligned}$$

Combining above analysis finishes the proof. \square

Lemma 9. *There exists a constant $C = C(\mu, K(\cdot))$, such that the following holds with probability at least $1 - 2\delta$. Define $A(n) = 0.25 \cdot \sqrt{Cd \cdot \ln(kn/\delta)}$, $\forall k \in [c]$:*

$$\mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2}$$

Proof. First we can write the following bound for the l.h.s:

$$\begin{aligned} \mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 &\leq \mathbb{E}_n \sqrt{\sum_{i \in [d]} |\Delta_{t,i} f_{n,h,k}(X) - \Delta_{t,i} f_k(X)|^2 \cdot \mathbb{1}_{A_{n,i}}(X)} \\ &\leq \sqrt{\sum_{i \in [d]} \mathbb{E}_n |\Delta_{t,i} f_{n,h,k}(X) - \Delta_{t,i} f_k(X)|^2 \cdot \mathbb{1}_{A_{n,i}}(X)} \\ &\leq \sqrt{\sum_{i \in [d]} \frac{1}{t^2} \max_{s \in \{-t, t\}} \mathbb{E}_n |f_{n,h,k}(X + se_i) - f_k(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X)} \end{aligned}$$

First observe that:

$$\begin{aligned} \mathbb{E}_n |f_{n,h,k}(X + se_i) - f_k(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X) &\leq \mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_k(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X) \\ &+ \mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_{n,h,k}(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X) \end{aligned}$$

Also notice that: $\mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_k(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X)$ and $\mathbb{E}_n |\tilde{f}_{n,h,k}(X + se_i) - f_{n,h,k}(X + se_i)|^2 \cdot \mathbb{1}_{A_{n,i}}(X)$ can be respectively bounded by two lemmas from above, thus we get with probability at least $1 - 2\delta$

$$\mathbb{E}_n |f_{n,h,k}(X + se_i) - f_k(X + se_i)|^2 \leq h^2 R^2 + \sqrt{\frac{A(n)}{nh^d}}$$

Combining above we get with probability at least $1 - 2\delta$, $\forall k \in [c]$

$$\|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2}$$

□

The following theorem provides a bound on $\mathbb{E}_n \|\nabla f(X) - \hat{\nabla} f_{n,h}(X)\|_2$:

Theorem 2. *With probability at least $1 - 2\delta$ over the choice of X , we have $\forall k \in [c]$:*

$$\mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 \leq \frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2}$$

Proof. We start with the now familiar decomposition:

$$\begin{aligned} \mathbb{E}_n \|\nabla f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 &\leq \mathbb{E}_n \|\hat{\nabla} f_k(X) - \hat{\nabla} f_{n,h,k}(X)\|_2 + \mathbb{E}_n \|\nabla f_k(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f_k(X)\|_2 \\ &+ \mathbb{E}_n \|\nabla f_k(X) \circ \overline{\mathbb{I}_n(X)}\|_2 \end{aligned}$$

By Lemma 5 we bound $\mathbb{E}_n \|\nabla f(X) \circ \overline{\mathbb{I}_n(X)}\|_2$; by Lemma 6 we bound $\mathbb{E}_n \|\nabla f(X) \circ \mathbb{I}_n(X) - \hat{\nabla} f(X)\|_2$; by Lemma 9 we bound $\mathbb{E}_n \|\hat{\nabla} f(X) - \hat{\nabla} f_{n,h}(X)\|_2$. Combining these results concludes the proof. □

5 Bounds on Eigenvalues and Eigenspace variations

In the above section, we established that $\mathbb{E}_n \hat{G}(X)$ is a consistent estimator of $\mathbb{E}_X G(X)$. In this section, we also establish consistency of its eigenvalues and eigenspaces,. The analysis here is based upon results from matrix perturbation theory [10, 11].

5.1 Eigenvalues variation

We consider the following lemma for eigenvalues variation from matrix perturbation theory:

Lemma 10. [10] Suppose both G and \hat{G} are Hermitian matrices of size $d \times d$, and admit the following eigen-decompositions:

$$G = X\Lambda X^{-1} \quad \text{and} \quad \hat{G} = \hat{X}\hat{\Lambda}\hat{X}^{-1}$$

where X and \hat{X} are nonsingular and

$$\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_d) \quad \text{and} \quad \hat{\Lambda} = \text{diag}(\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_d)$$

and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$, $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$. Thus for any unitary invariant norm $\|\cdot\|$, we have

$$\|\text{diag}(\lambda_1 - \hat{\lambda}_1, \lambda_2 - \hat{\lambda}_2, \dots, \lambda_d - \hat{\lambda}_d)\| \leq \|G - \hat{G}\|$$

More specifically, when considering the spectral norm, we have

$$\max_{i \in [d]} |\lambda_i - \hat{\lambda}_i| \leq \|G - \hat{G}\|_2$$

and when considering the Frobenius norm, we have

$$\sqrt{\sum_{i \in [d]} |\lambda_i - \hat{\lambda}_i|^2} \leq \|G - \hat{G}\|_F$$

Using the above lemma, we obtain the following theorem that bounds the eigenvalue variation:

EIGENVALUE VARIATION BOUND

Theorem 3. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ be the eigen-values of $\mathbb{E}_X G(X)$, let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ be the eigen-values of $\mathbb{E}_n \hat{G}(X)$. There exist $C = C(\mu, K(\cdot))$ and $N = N(\mu)$ such that the following holds with probability at least $1 - 2\delta$. Define $A(n) = \sqrt{Cd \cdot \log(n/\delta)} \cdot C_Y^2(\delta/2n) \cdot \sigma_Y^2 / \log^2(n/\delta)$. Let $n \geq N$, we have:

$$\begin{aligned} \max_{i \in [d]} |\lambda_i - \hat{\lambda}_i| &\leq \frac{6R^2}{\sqrt{n}} (\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}}) + \left(3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR + C_Y(\delta)}{t} \right) \right). \\ &\left[\frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} \right] \end{aligned}$$

Proof. By Lemma 10, we bound $\max_{i \in [d]} |\lambda_i - \hat{\lambda}_i|$ with respect to $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2$; by Theorem 1 we bound $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2$. \square

5.2 Eigenspace variation

First we introduce the following definition:

Definition 1. (Angles between two subspaces) Let $X, \hat{X} \in \mathbb{R}^{d \times k}$ have full column rank k . The angle matrix $\Theta(X, \hat{X})$ between X and \hat{X} is defined as:

$$\arccos((X^T X)^{-\frac{1}{2}} X^T \hat{X} (\hat{X}^T \hat{X})^{-1} \hat{X}^T X (X^T X)^{-\frac{1}{2}})^{\frac{1}{2}}$$

More specifically, when $k = 1$, it reduces to the angle between two vectors: $\Theta(\mathbf{x}, \hat{\mathbf{x}}) = \arccos \frac{|\mathbf{x}^T \hat{\mathbf{x}}|}{\|\mathbf{x}\|_2 \|\hat{\mathbf{x}}\|_2}$

Armed with this definition, we consider the following lemma on eigenspace variation:

Lemma 11. [11] Suppose both G and \hat{G} are Hermitian matrices of size $d \times d$, and admit the following eigen-decompositions:

$$G = [X_1 \quad X_2] \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix} \begin{bmatrix} X_1^{-1} \\ X_2^{-1} \end{bmatrix} \quad \text{and } \hat{G} = [\hat{X}_1 \quad \hat{X}_2] \begin{bmatrix} \hat{\Lambda}_1 & 0 \\ 0 & \hat{\Lambda}_2 \end{bmatrix} \begin{bmatrix} \hat{X}_1^{-1} \\ \hat{X}_2^{-1} \end{bmatrix}$$

where $X = [X_1 \quad X_2]$ and $\hat{X} = [\hat{X}_1 \quad \hat{X}_2]$ are unitary. We have

$$\|\sin \Theta(X_1, \hat{X}_1)\|_2 \leq \frac{\|(\hat{G} - G)X_1\|_2}{\min_{\lambda \in \lambda(\Lambda_1), \hat{\lambda} \in \lambda(\Lambda_2)} |\lambda - \hat{\lambda}|}$$

Using the above lemma, we get the following theorem for eigenspaces variation:

EIGENSPACE VARIATION

Theorem 4. Write the eigen-decompositions of $\mathbb{E}_X G(X)$ and $\mathbb{E}_n \hat{G}(X)$ as

$$\mathbb{E}_X G(X) = [X_1 \quad X_2] \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix} \begin{bmatrix} X_1^{-1} \\ X_2^{-1} \end{bmatrix}, \quad \mathbb{E}_n \hat{G}(X) = [\hat{X}_1 \quad \hat{X}_2] \begin{bmatrix} \hat{\Lambda}_1 & 0 \\ 0 & \hat{\Lambda}_2 \end{bmatrix} \begin{bmatrix} \hat{X}_1^{-1} \\ \hat{X}_2^{-1} \end{bmatrix}$$

There exist constants $C = C(\mu, K(\cdot))$ and $N = N(\mu)$ such that the following holds with probability at least $1 - 2\delta$. Define $A(n) = \sqrt{Cd \cdot \log(n/\delta)} \cdot C_Y^2(\delta/2n) \cdot \sigma_Y^2 / \log^2(n/\delta)$. Let $n \geq N$:

$$\begin{aligned} \|\sin \Theta(X_1, \hat{X}_1)\|_2 &\leq \\ &\left(\frac{\|X_1\|_2}{\min_{\lambda \in \lambda(\Lambda_1), \hat{\lambda} \in \lambda(\Lambda_2)} |\lambda - \hat{\lambda}|} \right) \cdot \left(\frac{6R^2}{\sqrt{n}} (\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}}) + \left(3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR + C_Y(\delta)}{t} \right) \right) \cdot \right. \\ &\left. \left[\frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} \right] \right) \end{aligned}$$

Proof. By Lemma 11, we bound $\|\sin \Theta(X_1, \hat{X}_1)\|_2$ with respect to $\|X_1(\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X))\|_2$, since $\|X_1(\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X))\|_2 \leq \|X_1\|_2 \cdot \|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2$, and by Theorem 1 we bound $\|\mathbb{E}_n \hat{G}(X) - \mathbb{E}_X G(X)\|_2$. Combining these concludes the proof. \square

6 Recovery of projected semiparametric regression model

In this section, we return to the multi-index motivation of the EGOP and EJOP discussed in the introduction to this chapter. For ease of exposition, we restrict our discussion to the EGOP, but the same argument also works for the EJOP.

Consider the following projected semiparametric regression model: $f(\mathbf{x}) = g(V^T \mathbf{x})$ where $V \in \mathbb{R}^{d \times r}$, $r \ll d$ is a dimension-reduction projection matrix, and g is a nonparametric function. Without loss of generality, we assume $V = [v_1, v_2, \dots, v_r]$, where $v_i \in \mathbb{R}^d$, $i \in [r]$ is a set of orthonormal vectors, and the gradient outer product (GOP) matrix of g : $\mathbb{E}_X [\nabla g(V^T X) \cdot \nabla g(V^T X)^T]$ is non-singular. The following proposition gives the eigen-decomposition of gradient outer product (GOP) matrix of f : $\mathbb{E}_X G(\mathbf{x})$.

Proposition 1. Suppose the eigen-decomposition of $\mathbb{E}_X [\nabla g(V^T X) \cdot \nabla g(V^T X)^T]$ is given by:

$$\mathbb{E}_X [\nabla g(V^T X) \cdot \nabla g(V^T X)^T] = Z \Lambda Z^{-1}$$

then we have the following eigen-decomposition of $\mathbb{E}_X G(X)$:

$$\mathbb{E}_X G(X) = [VZ \quad U] \begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} Z^{-1}V^T \\ U^T \end{bmatrix}$$

where $U = [u_1, u_2, \dots, u_{d-r}]$, $u_i \in [d-r]$ is a set of orthonormal vectors in $\ker(V^T)$.

Proof. Since $f(\mathbf{x}) = g(V^T \mathbf{x})$, we have $\nabla f(\mathbf{x}) = V \nabla g(V^T \mathbf{x})$. Thus we get:

$$\mathbb{E}_X G(X) = V \mathbb{E}_X [\nabla g(V^T X) \cdot \nabla g(V^T X)^T] V^T = V Z \Lambda Z^{-1} V^T$$

When we check the eigen-decomposition given in the proposition, the above equation is satisfied. Moreover, since

$$[VZ \quad U] \begin{bmatrix} Z^{-1} V^T \\ U^T \end{bmatrix} = \begin{bmatrix} Z^{-1} V^T \\ U^T \end{bmatrix} [VZ \quad U] = I$$

concludes the proof. \square

Since Z in the above proposition is nonsingular, we get that $\text{im}(V) = \text{im}(VZ)$, which means that the column space of projection matrix V is exactly the subspace spanned by the top- r eigenvectors of the GOP matrix $\mathbb{E}_X G(X)$. This point has also been noticed by [9, 22, 21].

Lastly, we need to show that the projection matrix V can be recovered using the estimated GOP matrix. This is captured in the following:

RECOVERY OF SEMI-PARAMETRIC MODEL

Theorem 5. Suppose the function f we want to estimate has the form $f(\mathbf{x}) = g(V^T \mathbf{x})$, and $\tilde{V} \in \mathbb{R}^{d \times r}$ is the matrix composed by the top- r eigenvectors of $\mathbb{E}_n \hat{G}(X)$, then with probability at least $1 - 2\delta$:

$$\|\sin \Theta(V, \tilde{V})\|_2 \leq \frac{1}{\lambda_{\min}} \left(\frac{6R^2}{\sqrt{n}} (\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}}) + \left(3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR + C_Y(\delta)}{t} \right) \right) \right. \\ \left. \left[\frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} \right] \right)$$

where λ_{\min} is the smallest eigenvalue of $\mathbb{E}_X [\nabla g(V^T X) \cdot \nabla g(V^T X)^T]$. Suppose $\lambda_1, \lambda_2, \dots, \lambda_{d-r}$ are the lowest $d - r$ eigenvalues of $\mathbb{E}_n \hat{G}(X)$, and with probability at least $1 - 2\delta$:

$$\max_{i \in [d-r]} |\lambda_i| \leq \left(\frac{6R^2}{\sqrt{n}} (\sqrt{\ln d} + \sqrt{\ln \frac{1}{\delta}}) + \left(3R + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} + \sqrt{d} \left(\frac{hR + C_Y(\delta)}{t} \right) \right) \right. \\ \left. \left[\frac{\sqrt{d}}{t} \sqrt{\frac{A(n)}{nh^d} + h^2 R^2} + R \left(\sqrt{\frac{d \ln \frac{d}{\delta}}{2n}} + \sqrt{\sum_{i \in [d]} \mu^2(\partial_{t,i}(\mathcal{X}))} \right) + \sqrt{\sum_{i \in [d]} \epsilon_{t,i}^2} \right] \right)$$

Proof. We only sketch the proof. First of all, notice that V is a semi-orthogonal matrix, therefore $\|V\|_2 = 1$. When this observation is combined with above proposition and Theorem 4, we get a proof of the first part of the theorem. For proving the second part of the theorem, first observe that by proposition 1, the lowest $d - r$ eigenvalues of $\mathbb{E}_X G(X)$ are all zeros. This observation when combined with lemma 10 finishes the proof. \square

7 Classification Experiments

In this section, we examine the utility of the EJOP as a technique for metric estimation, when used in the setting of non-parametric classification. We consider non-parametric classifiers that rely on the notion of distance, parameterized by a matrix $\mathbf{M} \succeq 0$, with the squared distance computed as $(\mathbf{x} - \mathbf{x}')^T \mathbf{M} (\mathbf{x} - \mathbf{x}')$. In the experiments reported in this section, we consider three different choices for \mathbf{M} : The first is $\mathbf{M} = \mathbf{I}$, which corresponds to the Euclidean distance. Second we consider the

Method	Error %
Euclidean	4.93
ReliefF	4.11
EJOP (k NN)	2.08
EJOP (h NN)	2.17

Table 1: Error rates on MNIST using EJOP as the underlying metric

case when $\mathbf{M} = \mathbf{D}$, where \mathbf{D} is a diagonal matrix, the notion of distance in this case corresponds to a scaled Euclidean distance. In particular, in the absence of a gradients weights [6], [7] like approach for the multiclass case, we instead obtain weights by using the ReliefF procedure [5], which estimates weights for the multiclass case by a series of one versus all binary classifications. We use this as a baseline. Finally, we consider $\mathbf{M} = \mathbb{E}_n G_n(X)$, where $\mathbb{E}_n G_n(X)$ is the estimated EJOP matrix. In particular, letting $V D V^\top$ denote the spectral decomposition of \mathbf{M} , we use it to transform the input \mathbf{x} as $D^{1/2} V^\top \mathbf{x}$ for the distance computation. Next, for a fixed choice of \mathbf{M} , we can define nearest neighbors of a query point \mathbf{x} in various ways. We consider the following two cases: First, k nearest neighbors (denoted henceforth as k NN) for fixed k and second, neighbors that have distance $\leq h$ for fixed h from the query. We denote this as h NN. This corresponds to nonparametric classification using a boxcar kernel.

7.1 A First Experiment on MNIST

We first consider the MNIST dataset to test the quality of the EJOP metric. We set aside 10,000 points as a validation set, which is used to obtain the ReliefF weights, as well as for tuning the parameter t_i for $i = 1, \dots, 784$ in the EJOP estimation, as well as for tuning the parameter h . For the k NN case, we fix $k = 7$. While the t_i can be tuned separately for each class, we observed that it doesn't afford significant advantages over using a single t . Note that no preprocessing is applied on the images, and the metric estimation, as well as classification is done using the raw images. The results on the test set are illustrated in table 1. While MNIST is a considerably easy task, the improvement given by the use of the EJOP as the distance metric over the plain Euclidean distance is substantial.

7.2 Classification Experiments

Next, we consider the datasets considered in [16] and [4]. First we report experiments using plain Euclidean distance, h NN and k NN when the EJOP is used as the metric. The train/test splits are reported in the table. We split 20 % of the training portion to tune for h , k and t_i , the results reported are over 10 random runs.

Dataset	d	N	train/test	Euclidean	h-NN	k -NN
Isolet	172	7797	4000/2000	14.17 ± 0.7	10.14 ± 0.9	8.67 ± 0.6
USPS	256	9298	4000/2000	7.87 ± 0.2	7.14 ± 0.3	6.67 ± 0.4
Letters	16	20000	4000/2000	7.65 ± 0.3	5.12 ± 0.7	4.37 ± 0.4
DSLR	800	157	100/50	84.85 ± 4.8	41.13 ± 2.1	35.01 ± 1.4
Amazon	800	958	450/450	66.17 ± 2.8	41.07 ± 2.3	39.85 ± 1.5
Webcam	800	295	145/145	61.43 ± 1.7	24.86 ± 1.2	23.71 ± 2.1
Caltech	800	1123	550/500	85.41 ± 3.5	54.65 ± 2.6	52.86 ± 3.1

Table 2: Classification error rates on the datasets used in [4] using Euclidean distance, h NN and k NN while using the EJOP as the metric

Next, we report results obtained on the same folds using three popular metric learning methods. In particular, we consider Large Margin Nearest Neighbors (LMNN) [20], Information Theoretic Metric Learning (ITML) [1] and Metric Learning to Rank (MLR) [12]. Since these methods explicitly optimize for the metric over a space of possible metrics, the comparison is manifestly unfair, since in the case of the EJOP, there is only one metric, which is estimated from the training samples. The setup is the same as discussed above, with the following addition for the metric learning methods: We learn the metric for $k = 5$, and test is using whatever k that was returned while tuning for the EJOP. We observe that despite its simplicity, EJOP yields performance comparable to the metric learning methods, in some cases returning error rates comparable to those returned by MLR and ITML.

Dataset	h-NN	k -NN	ITML	LMNN	MLR
Isolet	10.14 ± 0.9	8.67 ± 0.6	8.43 ± 0.3	5.3 ± 0.4	6.59 ± 0.3
USPS	7.14 ± 0.3	6.67 ± 0.4	6.57 ± 0.2	6.23 ± 0.5	6.76 ± 0.3
Letters	5.12 ± 0.7	4.37 ± 0.4	5 ± 0.7	4.1 ± 0.4	17.81 ± 5.1
DSLR	41.13 ± 2.1	35.01 ± 1.4	21.65 ± 3.1	29.65 ± 3.7	41.54 ± 2.3
Amazon	41.07 ± 2.3	39.85 ± 1.5	39.83 ± 3.5	33.08 ± 4.2	29.65 ± 2.6
Webcam	24.86 ± 1.2	23.71 ± 2.1	15.31 ± 4.3	19.78 ± 1.5	27.54 ± 3.9
Caltech	54.65 ± 2.6	52.86 ± 3.1	52.37 ± 4.2	52.15 ± 3.2	51.34 ± 4.5

Table 3: Classification error rates given by the EJOP, and three popular metric learning methods

Finally, table 4 presents results comparing classification error rates obtained by using LMNN and MLR in the plain vanilla case and when the estimated EJOP is used as an initialization metric, showing that it improves performance. We also observed that the convergence of the two methods was also much faster when EJOP was used as an initialization, indicating that the EJOP serves as a good prior.

Dataset	LMNN	LMNN (EJOP Init.)	MLR	MLR (EJOP init.)
Isolet	5.3 ± 0.4	4.9 ± 0.3	6.59 ± 0.3	6.11 ± 0.2
USPS	6.23 ± 0.5	5.96 ± 0.3	6.76 ± 0.3	6.21 ± 0.3
Letters	4.1 ± 0.4	4 ± 0.3	17.81 ± 5.1	15.17 ± 4.6
DSLR	29.65 ± 3.7	26.13 ± 3.9	41.54 ± 2.3	38.61 ± 2.9
Amazon	33.08 ± 4.2	31.17 ± 3.7	29.65 ± 2.6	26.19 ± 2.9
Webcam	19.78 ± 1.5	19.5 ± 1.7	27.54 ± 3.9	24.11 ± 3.5
Caltech	52.15 ± 3.2	50.35 ± 3.4	51.34 ± 4.5	50.1 ± 3.7

Table 4: Classification error rate improvements when EJOP is used as an initialization

8 Conclusion

In this article we adapted the expected gradient outerproduct to the multi-class setting. We studied its theoretical properties, proposed a simple estimator and showed that it affords improvements in non-parametric classification tasks.

References

- [1] Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pages 209–216. ACM, 2007.
- [2] Wolfgang Hardle, Peter Hall, Hidehiko Ichimura, et al. Optimal smoothing in single-index models. *The annals of Statistics*, 21(1):157–178, 1993.
- [3] Sham Kakade. Lecture notes on multivariate analysis, dimensionality reduction, and spectral methods. *STAT 991, Spring*, 2010.
- [4] Dor Kedem, Stephen Tyree, Kilian Weinberger, Fei Sha, and Gert Lanckriet. Non-linear metric learning. In *Advances in Neural Information Processing Systems 25*, pages 2582–2590, 2012.
- [5] Kenji Kira and Larry M. Rendell. The feature selection problem: Traditional methods and a new algorithm. In *Proceedings of AAAI*, pages 129–134. AAAI, 1992.
- [6] Samory Kpotufe and Abdeslam Boularias. Gradient weights help nonparametric regressors. In *NIPS*, pages 2870–2878, 2012.
- [7] Samory Kpotufe, Abdeslam Boularias, Thomas Schultz, and Kyoungok Kim. Gradient weights improve regression and classification. *Journal of Machine Learning Research*, 17:1–34, 2016.
- [8] Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- [9] Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- [10] Ren-Cang Li. Relative perturbation theory i: eigenvalue and singular value variations. *SIAM Journal on Matrix Analysis and Applications*, 19:956–982, 1998.

- [11] Ren-Cang Li. Relative perturbation theory ii: eigenspace and singular space variations. *SIAM Journal on Matrix Analysis and Applications*, 20:471–492, 1999.
- [12] B. McFee and G. R. G. Lanckriet. Metric learning to rank. In *Proceedings of the 27th International Conference on Machine Learning (ICML'10)*, 2010.
- [13] James L Powell, James H Stock, and Thomas M Stoker. Semiparametric estimation of index coefficients. *Econometrica: Journal of the Econometric Society*, pages 1403–1430, 1989.
- [14] Shubhendu Trivedi. Notes on asymmetric metric learning for k-nn classification. *Unpublished notes*, 2015.
- [15] Shubhendu Trivedi. Discriminative learning of similarity and group equivariant representations. *arXiv Preprint, arXiv:1808.10078*, 2018.
- [16] Shubhendu Trivedi, David Mcallester, and Gregory Shakhnarovich. Discriminative metric learning by neighborhood gerrymandering. In *Advances in Neural Information Processing Systems*, pages 3392–3400, 2014.
- [17] Shubhendu Trivedi, Jialei Wang, Samory Kpotufe, and Gregory Shakhnarovich. A consistent estimator of the expected gradient outerproduct. In *Proceedings of the 30th International Conference on Uncertainty in Artificial Intelligence*, pages 819–828. AUAI, 2014.
- [18] Joel. A. Tropp. User-friendly tools for random matrices: An introduction. *Tutorial at NIPS*, 2012.
- [19] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, 16(2):264–280, 1971.
- [20] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 10:207–244, 2009.
- [21] Qiang Wu, Justin Guinney, Mauro Maggioni, and Sayan Mukherjee. Learning gradients: Predictive models that infer geometry and statistical dependence. *Journal of Machine Learning Research*, 11:2175–2198, 2010.
- [22] Yingcun Xia, Howell Tong, W. K. Li, and Li-Xing Zhu. An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):363–410, 2002.
- [23] Yingcun Xia, Howell Tong, WK Li, and Li-Xing Zhu. An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):363–410, 2002.