



**HAL**  
open science

# Local logistic regression for dimension reduction in classification

Touqeer Ahmad, François Portier, Gilles Stupfler

► **To cite this version:**

Touqeer Ahmad, François Portier, Gilles Stupfler. Local logistic regression for dimension reduction in classification. 2024. hal-04647345

**HAL Id: hal-04647345**

**<https://hal.science/hal-04647345>**

Preprint submitted on 14 Jul 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Local logistic regression for dimension reduction in classification

Touqeer Ahmad<sup>a</sup>, François Portier<sup>a</sup> & Gilles Stupfler<sup>b</sup>

<sup>a</sup> CREST, ENSAI, University of Rennes, France

<sup>b</sup> Univ Angers, CNRS, LAREMA, SFR MATHSTIC, F-49000 Angers, France

**Abstract.** Sufficient dimension reduction has received much interest over the past 30 years. Most existing approaches focus on statistical models linking the response to the covariate through a regression equation, and as such are not adapted to binary classification problems. We address the question of dimension reduction for binary classification by fitting a localized nearest-neighbor logistic model with  $\ell_1$ -penalty in order to estimate the gradient of the conditional probability of interest. Our theoretical analysis shows that the pointwise convergence rate of the gradient estimator is optimal under very mild conditions. The dimension reduction subspace is estimated using an outer product of such gradient estimates at several points in the covariate space. Our implementation uses cross-validation on the misclassification rate to estimate the dimension of this subspace. We find that the proposed approach outperforms existing competitors in synthetic and real data applications.

**AMS Subject Classifications:** 62G08, 62G20, 62J07, 62J12

**Keywords:** Binary classification, Cross-entropy, Dimension reduction, Local linear estimation, Nearest neighbors, Penalization, Weak convergence.

## 1 INTRODUCTION

A number of statistical problems require predicting whether an event of interest  $A$  happens given auxiliary information in the form of a random covariate  $X \in \mathbb{R}^p$ . In finance, for example, it is of interest to predict whether a given customer, characterized by information such as age, occupation and net income, is going to default on a loan. In hydrology, one may have to predict whether flood defenses are going to be breached during a storm based on weather variables and local topographical information. Marketing studies often try to assess customer churn within slices of a customer database, which amounts to predicting whether a customer is going to stop using a product given their past behavior and responses to marketing communications. In each of those problems, a natural approach is to set  $Y = \mathbb{1}_A$ , the indicator function of  $A$ , and to model the conditional probability  $\mathbb{P}(Y = 1 | X)$ . A first, naive solution is then to model this conditional probability as  $\mathbb{P}(Y = 1 | X) = \exp(\beta^T X) / (1 + \exp(\beta^T X))$  and to estimate the direction  $\beta \in \mathbb{R}^p$  with logistic regression. This is a well-understood methodology, part of generalized linear modelling, for which an excellent reference is [Agresti \(2013\)](#). Logistic

regression, and its alternatives such as probit regression or log-log regression, is sensitive to misspecification like any other parametric technique. Moreover, it relies on all the information brought by  $X$  about  $Y$  to lie in a single projection of  $X$ , meaning that up to rotation within the covariate space, only a single component of  $X$  should contain information about the event of interest. These strong restrictions have motivated the construction and study of several nonparametric and semiparametric regression methods over the past 30 years.

Many such methods are partly built on a class of techniques aimed at estimating so-called sufficient dimension reduction subspaces. A dimension reduction subspace  $E$  for the distribution of a random pair  $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$  is any linear subspace of  $\mathbb{R}^p$  spanned by the  $(p-d)$ -dimensional columns of a matrix  $B \in \mathbb{R}^{p \times d}$  with  $p$  rows and  $d$  columns such that  $Y$  and  $X$  are conditionally independent in the sense of Dawid (1979) given the vector  $B^T X \in \mathbb{R}^d$ ; the notion of conditional independence, which is linked to fundamental statistical notions such as sufficiency and ancillarity, is also discussed from a different viewpoint in Basu and Pereira (1983). Of course, dimension reduction subspaces always exist and are not unique in general, but the central subspace, that is, the intersection of all dimension reduction subspaces, is identifiable and is in many cases a dimension reduction subspace as well (Cook, 1994, 1996; Portier and Delyon, 2013). An interesting case occurs when the central subspace has dimension  $d = 1$ , which is the so-called single-index regression situation (Hall, 1989) where the distribution of  $Y$  given  $X$  is an unknown function of a single, unknown projection  $\beta^T X$  to be estimated. This model extends traditional parametric regression models by allowing the link function  $g$  satisfying  $\mathbb{E}(Y|X) = g(\beta^T X)$  to be unknown. More generally, whenever a matrix  $B$  forms a dimension reduction subspace,  $Y$  can in fact be written as a measurable function of  $B^T X$  and  $\varepsilon$ , where  $\varepsilon$  is independent from  $X$ ; this makes the model  $Y = g(B^T X) + \varepsilon$ , with  $g : \mathbb{R}^d \rightarrow [0, 1]$  unknown, a particularly appealing model for dimension reduction.

Among dimension reduction methods, *Sliced Inverse Regression* (SIR), pioneered by Li (1991), seeks to estimate the central subspace by recovering the conditional expectation  $\mathbb{E}(X|Y)$  from a family of moments  $\mathbb{E}(\psi(Y)X)$  depending on a class of (indicator) functions  $\psi$  encoding whether  $Y$  belongs to a particular slice of the data. Several extensions have been proposed, including *Sliced Average Variance Estimation* (SAVE) by Cook and Weisberg (1991). These methods rely on the strong assumption that  $X$  has an elliptically symmetric distribution (Cook and Weisberg, 1991). Directional Regression (DR) approaches as in Li and Wang (2007) make assumptions that are weaker but related to elliptical symmetry, and Likelihood Acquired Directions (LAD) as in Cook and Forzani (2009) assume that  $X$  given  $Y$  is Gaussian. Differently from these techniques, the *Principal Hessian Direction* (PHD) approach developed by e.g. Li (1992) for Gaussian multivariate data, and further investigated by Cook (1998) under conditions akin to elliptical symmetry again, directly work on the conditional expectation  $\mathbb{E}(Y|X)$  in order to recover the central subspace. Li (2007) modifies these methods to obtain sparse (with respect to the covariate index) dimension reduction subspaces. Adaptations of the SIR, SAVE, and PHD methods to rare event classification were studied by Aghbalou et al. (2024).

*Gradient-based estimation methods* do not assume elliptical symmetry and are based on the estimation of the gradient of a regression function in order to recover vectors in the space of interest. Gradient estimation methods have been introduced by Härdle and Stoker (1989) and Powell et al. (1989); the works by Klein and Spady (1993) and Huh and Park (2002) focus specifically on binary regression but only allows for a one-dimensional dimension reduction subspace. These ideas have been extended in various ways using local linear

estimation, see for example [Hristache et al. \(2001\)](#), which improves the average derivative estimator by focusing on the hyperplane orthogonal to the target direction in a single-index model, the Minimum Average Variance Estimation (MAVE) procedure of [Xia et al. \(2002\)](#) that simultaneously estimates the conditional variance of the observations given directions in the data and minimizes this quantity with respect to these directions, the method of [Dalalyan et al. \(2008\)](#) combining local linear estimation with an  $\ell_\infty$ -type Principal Component Analysis technique, and nearest-neighbor local linear estimators combined with a LASSO penalization devised by [Ausset et al. \(2021\)](#) in order to handle the high-dimensional setting. Local linear estimation, for which a comprehensive reference is [Fan and Gijbels \(1996\)](#), provides an efficient way to estimate the derivatives of the regression function and as such, it has an essential place in many dimension reduction methods such as those we have just highlighted. Somewhat tangentially, [Delecroix et al. \(2006\)](#) also build estimators based on minimizing general loss functions. With the exception of the work by [Klein and Spady \(1993\)](#) and [Delecroix et al. \(2006\)](#), the gradient-based methods outlined here are designed for dimension reduction in location regression models of the form  $Y = g(\beta^T X) + \varepsilon$ , which excludes the binary regression and classification frameworks; however, [Klein and Spady \(1993\)](#) and [Delecroix et al. \(2006\)](#) only handle single-index classification.

A more recent line of work has focused on using powerful modern tools of mathematical statistics such as *Reproducing Kernel Hilbert Spaces* (RKHS). The use of RKHS in dimension reduction appears to have been pioneered by [Fukumizu et al. \(2004\)](#), who use the covariance operator to characterize the central subspace when the response and covariate are finite-dimensional. Later, [Hsing and Ren \(2009\)](#) paired the use of RKHS with SIR to handle dimension reduction when the covariate is a stochastic process. Subsequent work by [Li et al. \(2011\)](#) unified the construction of linear and nonlinear dimension reduction subspaces again using support vector machine approaches that can handle binary regression models; this technique is closely linked to the Principal Optimal Transport Direction (POTD) method investigated in [Meng et al. \(2020\)](#). The recent solution put forward by [Kang and Shin \(2022\)](#), and building on earlier work by [Mukherjee and Wu \(2006\)](#), makes use of local estimation techniques based on an RKHS representation theorem to fit the function as well as its gradient. Their theory requires the true functions and gradient to lie within the RKHS space of interest, and to specify the RKHS kernel correctly. Computationally speaking, their optimization procedure relies on a risk function that is a second-order  $U$ -statistic involving  $n^2$  terms. Other machine learning techniques, such as ensemble learning, have also been paired with local estimation (see for example [Cai et al., 2023](#)) although they tend to rely on local least squares and thus not on loss functions better suited to binary classification.

This article proposes a new dimension reduction framework tailored to classification, handling dimension reduction subspaces not limited to a single projection, and able to perform variable selection while dealing with the curse of dimensionality. Assume that  $Y$  satisfies

$$\pi(X) := \mathbb{P}(Y = 1|X) = g(\beta^T X)$$

for some  $\beta \in \mathbb{R}^{p \times d}$ ,  $d \leq p$  and a measurable function  $g : \mathbb{R}^d \rightarrow [0, 1]$ . The function  $g$ , treated as a nuisance parameter, and the matrix  $\beta$  are both unknown. If we had  $d = 1$  and  $g(t) = \text{expit}(t) := \exp(t)/(1 + \exp(t))$ , then we would have  $\text{logit}(\pi(X)) := \log(\pi(X)/(1 - \pi(X))) = \beta^T X$ , and  $\beta \in \mathbb{R}^p$  would be exactly the gradient of  $\ell := \text{logit}(\pi)$ . Our main idea is that, if the conditional probability  $\pi$  is sufficiently regular at a given point  $x$  then  $\ell(X)$  can be approximated by a linear function of  $X - x$  in a neighborhood of  $x$ , whose intercept

$a = a(x)$  is just  $\ell(x)$  and whose gradient  $b = b(x) = \nabla\ell(x)$  belongs to the vector space  $\text{span}(\beta)$  that is spanned by the columns of  $\beta$ . We then construct, at each point  $x$ , a nearest-neighbor, penalized, local logistic (and hence convex) loss function depending on two arguments  $a \in \mathbb{R}$  and  $b \in \mathbb{R}^p$ . The weak convergence of the estimator follows, under very mild conditions, from a new general empirical process theory result on a family of nearest-neighbor estimators that is of independent interest. The gradient estimates at several points are then aggregated using the standard outer product of gradients of *e.g.* Section 3.1 in [Xia et al. \(2002\)](#), and  $\text{span}(\beta)$  is recovered by finding the eigenspaces of this matrix corresponding to its  $d$  largest eigenvalues. Using nearest neighbors (see [Biau and Devroye \(2015\)](#) for a recent textbook), instead of the Nadaraya-Watson weights as in [Fan et al. \(1995\)](#), ensures an adaptive bandwidth choice while keeping computational time reasonable thanks to algorithms such as  $k - d$  tree search. Our use of a penalized loss function ensures sparsity of the estimated vectors in a similar spirit as in [Li \(2007\)](#) but without requiring a first step non-sparse estimation and, unlike earlier local generalized linear techniques such as that of [Fan et al. \(1995\)](#), makes it possible to deal with potentially large dimensions (of the order of several dozens when the sample size is  $n = 1000$ ). While most of the existing dimension reduction methods select the number  $d$  of components in the reduction subspace by using rank testing procedures ([Bura and Yang, 2011](#); [Portier and Delyon, 2014](#)), we here propose to rely on the underlying classification context by developing a simple cross-validation approach comparing the different sets of ordered eigenvectors given by the outer product of gradients. Surprisingly, this is helpful in improving the final accuracy of the classification method even when the true dimension of the reduction subspace is known.

The paper is organized as follows. In Section 2.1, we describe our statistical framework in detail and we then construct our nearest-neighbor, penalized local logistic loss function, whose maximization gives rise to the proposed gradient estimator. In Section 2.2, we establish the pointwise weak convergence of this estimator and in Section 2.3 we employ an outer-product of gradients to estimate the central subspace. In Section 3, we discuss several practical aspects of the proposed algorithm, we construct our classification method, and we compare our approach to several competitors on synthetic and real data. The mathematical proofs and additional finite-sample results are postponed to an online Supplementary Material document.

## 2 STATISTICAL FRAMEWORK AND MAIN RESULTS

### 2.1 BACKGROUND

Let  $Y \in \{0, 1\}$  be a binary response variable, with random covariate  $X \in \mathbb{R}^p$ . The data at hand is supposed to be made of independent copies  $(Y_i, X_i)$  of the random pair  $(Y, X)$ , for  $1 \leq i \leq n$ . Suppose that for an  $x$  in the support  $S_X$  of  $X$  (assumed to have nonempty interior),

$$\pi(x) = \mathbb{P}(Y = 1|X = x) = g(\beta^T x),$$

for some unknown  $\beta \in \mathbb{R}^{p \times d}$ ,  $d \leq p$  and an unknown measurable function  $g : \mathbb{R}^d \rightarrow [0, 1]$ . The Kullback-Leibler divergence of the Bernoulli distribution with parameter  $\pi \in (0, 1)$  from the Bernoulli distribution with parameter  $q \in (0, 1)$  is

$$D_{\text{KL}}(\pi||q) = \pi \log \left( \frac{\pi}{q} \right) + (1 - \pi) \log \left( \frac{1 - \pi}{1 - q} \right).$$

Gibbs' theorem entails  $D_{\text{KL}}(\pi||q) \geq 0$ , with equality if and only if  $\pi = q$ . Equivalently, the cross-entropy function

$$H(\pi, q) = D_{\text{KL}}(\pi||q) - \{\pi \log(\pi) + (1 - \pi) \log(1 - \pi)\} = -\pi \log(q) - (1 - \pi) \log(1 - q),$$

obtained by adding to  $D_{\text{KL}}(\pi||q)$  the entropy of the Bernoulli distribution with parameter  $\pi$ , is minimal if and only if  $\pi = q$ . In the classification setup, where the random covariate  $X$  has distribution  $P_X$ , the natural loss function to consider is the integrated cross-entropy

$$\int_{x \in \mathcal{S}_X} H(\pi(x), q(x)) P_X(dx) = -\mathbb{E}[Y \log(q(X)) + (1 - Y) \log(1 - q(X))]$$

viewed as a function of the map  $x \mapsto q(x)$ . When  $g = g_0$  is known, one should search for a function  $q$  of the form  $x \mapsto g_0(b^T X)$ , for some  $b \in \mathbb{R}^{p \times d}$ , meaning that one should minimize  $-\mathbb{E}[Y \log(g_0(b^T X)) + (1 - Y) \log(1 - g_0(b^T X))]$  as a function of  $b$ . This suggests the estimator

$$\arg \max_{b \in \mathbb{R}^{p \times d}} \sum_{i=1}^n Y_i \log(g_0(b^T X_i)) + (1 - Y_i) \log(1 - g_0(b^T X_i)).$$

For  $d = 1$  and  $g_0 = \text{expit}$ , this is the logistic regression estimator, with the empirical integrated cross-entropy loss being just the negative conditional log-likelihood in the logistic model.

When  $g$  is unknown, the logistic regression estimator is not guaranteed to be a consistent estimator of  $\beta$ . However, if  $\pi$  is differentiable at  $x$ , then the gradient  $\nabla \pi(x)$  of  $\pi$  at  $x$  satisfies

$$\nabla \pi(x) = \beta \nabla g(\beta^T x).$$

As such,  $\nabla \pi(x) \in \text{span}(\beta)$ . If moreover  $0 < \pi(x) < 1$ , then likewise

$$\nabla \ell(x) = \frac{1}{\pi(x)(1 - \pi(x))} \nabla \pi(x) \in \text{span}(\beta), \text{ where } \ell = \text{logit } \pi.$$

To recover  $\text{span}(\beta)$ , it then suffices to estimate enough gradients of the form  $\nabla \ell(x_j)$ . Now a Taylor expansion of  $\ell$  around  $x$  yields  $\pi(y) \approx \text{expit}(a + b^T(y - x))$ , with  $a = \text{logit}(\pi(x)) = \ell(x)$  and  $b = \nabla \text{logit}(\pi(x)) = \nabla \ell(x)$ , for  $y$  close to  $x$ . This suggests that the logistic regression estimator can still be used to produce an estimator of  $\nabla \ell(x)$  if it is restricted to data points close enough to  $x$ . Namely, one should expect the local integrated cross-entropy

$$-\mathbb{E}[Y \log(\text{expit}(a + b^T(X - x))) + (1 - Y) \log(1 - \text{expit}(a + b^T(X - x))) \mid X \in B(x, \varepsilon)]$$

to be minimized at  $(a_\varepsilon(x), b_\varepsilon(x)) \approx (\ell(x), \nabla \ell(x))$  as  $\varepsilon \downarrow 0$ , where  $B(x, \varepsilon)$  is the closed Euclidean ball with center  $x$  and radius  $\varepsilon$  in  $\mathbb{R}^p$ . This shall be the cornerstone of our methodology.

## 2.2 NEAREST-NEIGHBOR PENALIZED LOCAL LOGISTIC ESTIMATOR

The localization induced by conditioning upon covariate values belonging to  $B(x, \varepsilon)$  is reproduced empirically using a nearest-neighbor procedure. Fix  $x \in \mathbb{R}^p$  and let  $N_k(x) \subset \{1, \dots, n\}$  be the set gathering the indices of the  $k$ -nearest neighbors  $X_i$  of the point  $x$ ; we shall assume in the theory below that the distribution of  $X$  has a density w.r.t. Lebesgue measure, so that

ties will not happen with probability 1 and  $N_k(x)$  is well-defined. The empirical counterpart of the local integrated cross-entropy is then

$$-\frac{1}{k} \sum_{i \in N_k(x)} Y_i \log(\text{expit}(a + b^T(X_i - x))) + (1 - Y_i) \log(1 - \text{expit}(a + b^T(X_i - x))).$$

Introducing the LASSO penalty  $\lambda \|b\|_1$ , where  $\|\cdot\|_1$  is the  $\ell_1$ -norm on  $\mathbb{R}^p$ , and rescaling, we naturally obtain a nearest-neighbor, penalized, local logistic estimator as

$$(\hat{a}_n(x), \hat{b}_n(x)) = \arg \max_{(a,b) \in \mathbb{R} \times \mathbb{R}^p} \{L_n(a,b) - \lambda \|b\|_1\} \quad (1)$$

with

$$L_n(a,b) = \sum_{i \in N_k(x)} Y_i \log(\text{expit}(a + b^T(X_i - x))) + (1 - Y_i) \log(1 - \text{expit}(a + b^T(X_i - x))).$$

Our main theoretical result is that one can obtain the asymptotic distribution of this pair of estimators under very weak assumptions on the distribution of  $X$  and the conditional distribution of  $Y|X = x$ . We spell out these conditions and their interpretation below.

- (A1) The distribution of  $X$  has a continuous density  $f_X$  with respect to the Lebesgue measure on  $\mathbb{R}^p$  and  $f_X(x) > 0$ .
- (A2) The function  $\pi : \mathbb{R}^p \rightarrow [0,1]$  is twice differentiable with continuous second order derivatives at  $x$  and such that  $\pi(x) \in (0,1)$ .

Assumption (A1) ensures that there are enough points around  $x$  for the nearest-neighbor procedure to work. It also ensures the good probabilistic behavior of the bandwidth

$$\hat{\tau}_{n,k}(x) := \inf \left\{ \tau \geq 0 : \sum_{i=1}^n \mathbb{1}_{B(x,\tau)}(X_i) \geq k \right\}$$

corresponding to the smallest radius  $\tau \geq 0$  such that the ball  $B(x, \tau)$  contains at least  $k$  points from the sample. Actually, the fact that  $X$  has a continuous distribution w.r.t. Lebesgue measure yields

$$\begin{aligned} L_n(a,b) &= \sum_{i=1}^n [Y_i \log(\text{expit}(a + b^T(X_i - x))) + (1 - Y_i) \log(1 - \text{expit}(a + b^T(X_i - x)))] \mathbb{1}_{B(x, \hat{\tau}_{n,k}(x))}(X_i). \end{aligned}$$

It also turns out that if  $k = k_n \rightarrow \infty$  with  $k/n \rightarrow 0$ , then under (A1),  $\hat{\tau}_{n,k}(x)/\tau_{n,k}(x) \rightarrow 1$  in probability, where, if  $V_p$  denotes the volume of the Euclidean unit ball in  $\mathbb{R}^p$ ,

$$\tau_{n,k}(x) = \left( \frac{k}{n f_X(x) V_p} \right)^{1/p}.$$

See Lemma 1 in [Portier \(2021\)](#). It is then reasonable to write, for  $n$  large enough,

$$\begin{aligned} L_n(a,b) &\approx \bar{L}_n(a,b) \\ &= \sum_{i=1}^n [Y_i \log(\text{expit}(a + b^T(X_i - x))) + (1 - Y_i) \log(1 - \text{expit}(a + b^T(X_i - x)))] \mathbb{1}_{B(x, \tau_{n,k}(x))}(X_i). \end{aligned}$$

The asymptotic behavior of  $\bar{L}_n(a, b)$  is much easier to study than that of  $L_n(a, b)$ , since it is a sum of independent and identically distributed random variables; like  $L_n(a, b)$ , it defines a concave objective function and therefore one should expect that (up to technical details) the asymptotic behavior of  $(\hat{a}_n(x), \hat{b}_n(x))$  will follow from the pointwise convergence of  $\bar{L}_n(a, b)$ . The key result in order to make this intuition rigorous is of independent interest and is the first main result of this paper. Here and throughout  $\|\cdot\|_2$  denotes the standard Euclidean norm.

**Theorem 1** (Central limit theorem for nearest-neighbor estimators). *Suppose that (A1) is fulfilled. Assume that  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ . Let  $\Psi_n : \{0, 1\} \times \mathbb{R}^p \rightarrow \mathbb{R}^q$  be a sequence of measurable vector-valued functions and suppose that there is a positive integer  $n_0$  such that*

$$\sup_{n \geq n_0} \sup_{z \in B(x, (3/2)^{1/p} \tau_{n,k}(x))} \max(\|\Psi_n(0, z)\|_2, \|\Psi_n(1, z)\|_2) < \infty.$$

Define a centered and càdlàg stochastic process  $(Z_n(\tau))_{\tau > 0}$  by

$$Z_n(\tau) = \frac{1}{\sqrt{k}} \sum_{i=1}^n \left\{ \Psi_n(Y_i, X_i) \mathbb{1}_{B(x, \tau)}(X_i) - \mathbb{E}[\Psi_n(Y, X) \mathbb{1}_{B(x, \tau)}(X)] \right\}.$$

(i) Then  $Z_n(\hat{\tau}_{n,k}(x)) = O_{\mathbb{P}}(1)$ .

(ii) Let  $\Sigma_n^2(X) = \mathbb{E}[\Psi_n(Y, X) \Psi_n(Y, X)^T | X]$ . If moreover there is a (positive semidefinite) matrix-valued function  $t \mapsto \Sigma^2(t, x)$  such that

$$\forall t \in [1/2, 3/2], \int_{B(0,1)} \Sigma_n^2(x + \tau_{n,k}(x) t^{1/p} v) dv \rightarrow V_p \Sigma^2(t, x),$$

then

$$Z_n(\hat{\tau}_{n,k}(x)) = Z_n(\tau_{n,k}(x)) + o_{\mathbb{P}}(1) \xrightarrow{d} \mathcal{N}(0, \Sigma^2(1, x)).$$

Theorem 1(ii) is obtained by, first, showing that the empirical process  $(Z_n(\tau))$  is tight and converges weakly to a continuous Gaussian process within a suitable functional space when the radius  $\tau$  is restricted to being of similar order as the nearest neighbor radius  $\tau_{n,k}(x)$ , and then by combining that with the fact that  $\hat{\tau}_{n,k}(x)/\tau_{n,k}(x) \rightarrow 1$  in probability, see Lemma 2.

Because of the LASSO penalty term in (1), the asymptotic distribution of  $\hat{b}_n(x) - \nabla \ell(x)$  will depend on the so-called local active set associated to  $\nabla \ell(x)$ , defined as the set of indices  $j$  such that  $\nabla \ell_j(x) \neq 0$ . This is in line with the asymptotic distribution obtained for the standard least squares LASSO regression estimator, see Knight and Fu (2000) and Zou (2006). Let, for any real number  $t$ , the quantity  $\text{sgn}(t) = \mathbb{1}_{(0, \infty)}(t) - \mathbb{1}_{(-\infty, 0)}(t)$  be the sign of  $t$ , that is,  $\text{sgn}(t) = 1$  when  $t \geq 0$  and  $-1$  otherwise. Finally, define

$$\Gamma(x) = \pi(x)(1 - \pi(x)) \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{1}{p+2} I_p \end{pmatrix}$$

where  $0_p$  is the zero vector in  $\mathbb{R}^p$  and  $I_p$  denotes the identity matrix of order  $p$ . Our first main result provides the limiting distribution of the pair  $(\hat{a}_n(x), \hat{b}_n(x))$ . In this result, assumption (A2) ensures that the gradient of  $\ell = \text{logit } \pi$  is well-defined, with the second order derivatives of  $\pi$  and hence of  $\ell$  coming into play when evaluating the rate of convergence of the bias term incurred when localizing. Denote by  $\Delta \pi(x)$  the Laplacian of  $\pi$  at  $x$ .



**Theorem 2** (Convergence of nearest-neighbor penalized local logistic regression estimators). *Suppose that (A1) and (A2) are fulfilled. If  $k := k_n \rightarrow \infty$  and  $\lambda := \lambda_n$  are such that  $k^{1+p/2}/n \rightarrow \infty$ ,  $k^{1+p/4}/n$  is bounded and  $n\lambda^p/k^{1+p/2} \rightarrow c \in [0, \infty)$  then we have*

$$\begin{aligned} & \begin{pmatrix} \sqrt{k}(\widehat{a}_n(x) - \ell(x)) \\ \tau_{n,k}(x)\sqrt{k}(\widehat{b}_n(x) - \nabla\ell(x)) \end{pmatrix} \\ & \stackrel{d}{=} \arg \max_{u=(u_0, u_1, \dots, u_p)^T \in \mathbb{R}^{p+1}} \left\{ u^T(W_n(x) + T_n(x)) - \frac{1}{2}u^T\Gamma(x)u \right. \\ & \quad \left. - (cf_X(x)V_p)^{1/p} \left( \sum_{j=1}^p \operatorname{sgn}(\nabla\ell_j(x))u_j\mathbb{1}_{\{\nabla\ell_j(x) \neq 0\}} + |u_j|\mathbb{1}_{\{\nabla\ell_j(x) = 0\}} \right) \right\} + o_{\mathbb{P}}(1) \end{aligned}$$

with  $W_n(x) \xrightarrow{d} \mathcal{N}(0, \Gamma(x))$  and

$$T_n(x) = \tau_{n,k}^2(x)\sqrt{k} \left( \frac{1}{2(p+2)} \left( \Delta\pi(x) - \frac{1-2\pi(x)}{\pi(x)(1-\pi(x))} \|\nabla\pi(x)\|_2^2 \right) \begin{pmatrix} 1 \\ 0_p \end{pmatrix} + o_{\mathbb{P}}(1) \right).$$

The conditions on  $k$  and  $\lambda$  in Theorem 2 are equivalent to assuming that  $\tau_{n,k}(x)\sqrt{k} \rightarrow \infty$ ,  $\tau_{n,k}^2(x)\sqrt{k}$  is bounded and  $\lambda/(\tau_{n,k}(x)\sqrt{k})$  converges to the finite constant  $(cf_X(x)V_p)^{1/p}$ . Condition  $\tau_{n,k}(x)\sqrt{k} \rightarrow \infty$  is necessary in order to be able to write a Taylor expansion of the penalized component in the loss function (although it is, strictly speaking, not needed for the analysis of the non-penalized local logistic maximum likelihood function, see Lemma 5). In pointwise results on local linear kernel quasi-maximum likelihood estimation with one-dimensional covariates, Theorem 1a in Fan et al. (1995) requires  $\sqrt{nh^3} = h\sqrt{nh} \rightarrow \infty$ , where  $h$  is the kernel bandwidth; note that  $h$  and  $\tau_{n,k}(x)$  play the same role and that, for  $p = 1$ , the kernel regression analogue of  $k$  is a quantity proportional to  $nh$ , so that conditions  $h\sqrt{nh} \rightarrow \infty$  and  $\tau_{n,k}(x)\sqrt{k} \rightarrow \infty$  are indeed analogous. Of course, condition  $\tau_{n,k}(x)\sqrt{k} \rightarrow \infty$  is automatically satisfied if  $\tau_{n,k}^2(x)\sqrt{k}$  converges to a finite positive limit, namely, when the nonparametric bias-variance tradeoff is achieved. We also highlight that condition  $k/n \rightarrow 0$ , which is standard in nearest-neighbor estimation, follows from assuming that  $k \rightarrow \infty$  and  $\tau_{n,k}^2(x)\sqrt{k}$  is bounded.

It follows that, under conditions (A1) and (A2) plus classical conditions linking the number of nearest neighbors and the penalizing constant  $\lambda$ , the two estimators  $\widehat{a}_n(x)$  and  $\widehat{b}_n(x)$  converge respectively at the rate  $1/\sqrt{k}$  and  $1/(\tau_{n,k}(x)\sqrt{k})$ . Note that for non-penalized local logistic regression ( $\lambda = 0$  and thus  $c = 0$ ), the result of Theorem 2 is just

$$\begin{pmatrix} \sqrt{k}(\widehat{a}_n(x) - \ell(x)) \\ \tau_{n,k}(x)\sqrt{k}(\widehat{b}_n(x) - \nabla\ell(x)) \end{pmatrix} \stackrel{d}{=} \mathcal{N}(0, \Gamma^{-1}(x)) + O_{\mathbb{P}}(\tau_{n,k}^2(x)\sqrt{k}) + o_{\mathbb{P}}(1).$$

In particular, if  $k\tau_{n,k}^4(x) \rightarrow 0$ , a straightforward application of the delta-method yields

$$\sqrt{k}(\operatorname{expit}(\widehat{a}_n(x)) - \pi(x)) \xrightarrow{d} \mathcal{N}(0, \pi(x)(1-\pi(x)))$$

as expected from standard maximum likelihood theory when the logistic regression model is valid. Condition  $k\tau_{n,k}^4(x) \rightarrow 0$ , which makes the bias term  $T_n(x)$  vanish asymptotically, again has a straightforward analogue in local linear kernel quasi-maximum likelihood estimation; for  $p = 1$ , the corresponding condition is  $nh^5 \rightarrow 0$ , which is exactly the bias condition necessary to eliminate the smoothing bias term in Theorem 1a in Fan et al. (1995).

Another immediate corollary of Theorem 2 can also be given on the estimation rate of the gradient  $\nabla\ell(x)$  by  $\widehat{b}_n(x)$ .

**Corollary 3** (Rate of convergence of the gradient estimator). *Under the conditions of Theorem 2,  $\widehat{b}_n(x)$  is a consistent estimator of  $\nabla\ell(x)$ , and*

$$\widehat{b}_n(x) - \nabla\ell(x) = O_{\mathbb{P}}\left(\frac{1}{\tau_{n,k}(x)\sqrt{k}}\right) + O_{\mathbb{P}}(\tau_{n,k}(x)).$$

These results provide statistical guarantees on the estimation of directions belonging to the dimension reduction subspace, and will therefore be useful in order to find such guarantees on the dimension reduction procedure. The rate obtained in Corollary 3 has also been obtained as part of non-asymptotic bounds on nearest neighbor estimation of the gradient of the regression function, although this of course requires strong assumptions on the tail behavior of the noise, see e.g. Theorem 1 in Ausset et al. (2021). The best achievable rate of convergence, obtained for  $k = n^{4/(p+4)}$ , is  $n^{-1/(p+4)}$  for the estimation of the gradient vector. This rate is optimal, see Stone (1982), is the same as the one obtained in Ausset et al. (2021), and is better than the rate  $n^{-1/(3(p+4))}$  obtained in Kang and Shin (2022).

The above results only require  $X$  to have a continuous density that is positive around  $x$ , and the conditional probability to be twice continuously differentiable at  $x$ . As a consequence the best that can be hoped for is, of course, a pointwise asymptotic convergence result in the spirit of Theorem 2. Stronger results, such as uniform convergence results for  $\widehat{b}_n(x)$  over compact subsets of the support of  $X$ , could be obtained under much more restrictive conditions.

### 2.3 AGGREGATING THE DIRECTIONS

The optimization procedure (1) allows to estimate  $\nabla\ell(x)$  which, as explained in Section 2.1, belongs to the central subspace for each  $x \in S_X$ . This subspace can be recovered by estimating several such directions at different points  $x$ . To this aim, given a probability measure  $\mu$  supported on  $S_X$ , let

$$M = \int_{\mathbb{R}^p} \nabla\ell(x)\nabla\ell(x)^T \mu(dx) = \mathbb{E}_{X^* \sim \mu}[\nabla\ell(X^*)\nabla\ell(X^*)^T].$$

To estimate the matrix  $M$ , generate  $X_i^* \sim \mu$ ,  $i = 1, \dots, m$  independently and compute

$$\widehat{M} = \frac{1}{m} \sum_{i=1}^m \widehat{b}_n(X_i^*)\widehat{b}_n(X_i^*)^T. \quad (2)$$

One can then define  $(\widehat{\beta}_1, \dots, \widehat{\beta}_p)$  as the set of orthogonal eigenvectors of  $\widehat{M}$ , ordered according to their eigenvalues (in decreasing order). Finally, given  $d \in \{1, \dots, p\}$  (which is chosen in practice by cross-validation, as we shall explain next), the projection matrix

$$\widehat{P}_{\widehat{\beta}} = \widehat{P}_{\widehat{\beta},d} = \sum_{k=1}^d \widehat{\beta}_k \widehat{\beta}_k^T$$

defines an estimator of the projection on the central subspace of interest. A corollary from our main results can now be stated on the estimation of  $M$ .

**Corollary 4.** Suppose that  $\mu$  is a finitely supported measure on  $S_X$  and that for each  $x \in \text{supp}(\mu)$ , assumptions (A1) and (A2) are satisfied. If  $k := k_n \rightarrow \infty$  and  $\lambda := \lambda_n$  are such that  $k^{1+p/2}/n \rightarrow \infty$ ,  $k^{1+p/4}/n$  is bounded and  $n\lambda^p/k^{1+p/2}$  converges to a finite constant, we have

$$\widehat{M} - M = O_{\mathbb{P}}\left(\frac{1}{\tau_{n,k}(x)\sqrt{k}}\right) + O_{\mathbb{P}}(\tau_{n,k}(x)) + O_{\mathbb{P}}\left(\frac{1}{\sqrt{m}}\right).$$

The above rate of convergence can be extended to the eigenprojector  $\widehat{P}_{\widehat{\beta}}$  as indicated by Lemma 4.1 in Tyler (1981). Also note that a natural choice for  $\mu$ , which is also the one made in our numerical experiments, is the empirical measure of  $X_1, \dots, X_n$ . This is, of course, a random measure whose number of atoms is not bounded with respect to  $n$ . From a theoretical perspective, we conjecture, following results given in Hristache et al. (2001) and Dalalyan et al. (2008), that such a choice would lead to a different rate of convergence compared to the one given in the above corollary, under substantially stronger regularity assumptions than ours.

### 3 NUMERICAL EXPERIMENTS

This section first explains how the algorithms for the proposed methods work to estimate the dimension reduction matrix  $M$ . We then describe the different competitors and evaluation metrics that we shall use. We next analyze synthetic data examples and finally consider real data examples in order to showcase the benefits of our methodology in classification tasks. The code that may be used to reproduce our experiments is publicly available on GitHub<sup>1</sup>.

#### 3.1 THE ALGORITHM

##### 3.1.1 ESTIMATION OF THE DIMENSION REDUCTION MATRIX

For a given choice of  $m$  (the number of  $b$ 's used to estimate  $M$ ), a given value of  $\lambda$ , the penalization parameter in the optimization problem, and  $k$  the number of neighbors, the computation of the dimension reduction matrix is straightforward, see Algorithm 1 below (the optimization is carried out using the R function `glmnet` from the package of the same name).

---

#### Algorithm 1 Estimation of $M$

---

- 1: **Input:**  $(X_1, Y_1), \dots, (X_n, Y_n) \in \mathbb{R}^p \times \{0, 1\}$ ,  $\lambda > 0$ ,  $k \in \{1, \dots, n\}$  and  $m \in \{1, \dots, n\}$
  - 2: **Output:** Dimension reduction matrix  $\widehat{M}$
  - 3: Draw uniformly a list  $\mathcal{X}_m$  of  $m$  observations among  $\mathcal{X} = (X_1, \dots, X_n)$  without replacement
  - 4: **for** each  $x \in \mathcal{X}_m$  **do**
  - 5:     Compute  $N_k(x)$ , the index of the  $k$  nearest neighbors to  $x$  among  $\mathcal{X}$
  - 6:     Compute  $\widehat{a}_n(x)$  and  $\widehat{b}_n(x)$  according to (1) using gradient descent
  - 7: **end for**
  - 8: Return  $\widehat{M} = \frac{1}{m} \sum_{x \in \mathcal{X}_m} \widehat{b}_n(x) \widehat{b}_n(x)^T$
- 

We now discuss the choice of the hyperparameters. Ideally, one should set  $m = n$  so that the gradient is estimated at each data point; we set here  $m = n/4$  to save computing time, as in our experience this choice does not adversely affect finite-sample results substantially.

<sup>1</sup>[https://github.com/touqeerahmadunipd/LLO\\_regression](https://github.com/touqeerahmadunipd/LLO_regression)

We also set  $k = \lfloor \sqrt{n} \rfloor$ . Furthermore, to mitigate biases that may arise from imbalanced class distributions, we exclude samples where after finding the closest neighbors, either class 0 or class 1 is rare, defined here as having fewer than 5 points within one of the two classes.

To estimate the directions featured in (2), we use either the pure nearest-neighbor logistic log-likelihood without penalization or its penalized version. For the latter, we require an optimal choice of  $\lambda$  which balances the regularization and goodness-of-fit. We propose to select the same parameter  $\lambda$  for all  $x \in \mathcal{X}_m$  in order to alleviate the computational burden. We therefore select  $\lambda$  by 10-fold cross-validation at the average point  $x = \bar{x} = \frac{1}{n} \sum_{i=1}^n X_i$ . In other words, we divide the data into 10 randomly selected subsets of equal size, each subset serving as a validation set while the remaining subsets are used for model fitting. The fitted model is assessed on the validation set using the misclassification error, that is, the proportion of observations whose label is not correctly predicted, which is nothing but the empirical counterpart of the misclassification risk  $\mathcal{R}(g) = \mathbb{P}(g(X) \neq Y)$  for a given classifier  $g$ ; in this cross-validation procedure, an observation is labeled as 1 if and only if its predicted probability of success by the nearest-neighbor logistic log-likelihood estimator  $\hat{\pi}_n(\bar{x}) = \text{expit}(\hat{a}_n(\bar{x}))$  exceeds 1/2. This evaluation of the quality of the fitted model is then done across a sequence of  $\lambda$  values: more precisely, we use the R function `cv.glmnet` from the `glmnet` package with `type.measure=class`, which automatically computes the average misclassification error across all validation sets, and the regularization parameter  $\lambda$  selected is the one minimizing this error.

### 3.1.2 SELECTING THE DIMENSION OF THE REDUCTION SUBSPACE

Unlike some dimension reduction approaches such as that of the seminal paper by Li (1991), we estimate the dimension of the reduction subspace using cross-validation on the underlying prediction problem, and not from a test on the eigenvalues of the outer product of gradients  $\hat{M}$ . We refer to Bura and Yang (2011) and Portier and Delyon (2014) for more details on the rank testing procedures used to estimate the dimension of the reduction subspace. Our procedure first divides the data into a training set  $(X_{\text{train}}, Y_{\text{train}})$  and a testing set  $(X_{\text{test}}, Y_{\text{test}})$ . The matrix  $\hat{M}$  and its  $p$  orthogonal eigenvectors  $\hat{\beta}_1, \dots, \hat{\beta}_p$ , in decreasing order according to their eigenvalues, are estimated from the training set following the procedure described in Algorithm 1. For every  $d \in \{1, \dots, p\}$ , the first  $d$  eigenvectors are then gathered in a matrix  $\hat{\beta}_{(1:d)} \in \mathbb{R}^{p \times d}$ , the sets of covariates  $X_{\text{train}}$  and  $X_{\text{test}}$  are projected onto the sets of lower-dimensional covariates  $X_{\text{train}}\hat{\beta}_{(1:d)}$  and  $X_{\text{test}}\hat{\beta}_{(1:d)}$ , and a classifier is learned based on the training set  $(X_{\text{train}}\hat{\beta}_{(1:d)}, Y_{\text{train}})$ . We use here mainly the knn nearest-neighbor classifier, that is, at a given point  $x$ , the result of the majority vote among the nearest neighbors of  $x$  within the space of projections of the covariates in the training set (with ties broken at random). For this classifier, the training step merely consists of storing the covariates in the training set along with their labels, which will form the basis for the vote at each point. In the real data analysis, we shall also compare our results with the Random Forest classifier, whose training step is nontrivial. The performance of the chosen classifier is then evaluated on the test set  $(X_{\text{test}}\hat{\beta}_{(1:d)}, Y_{\text{test}})$ , and the dimension retained is the one for which the classifier has the lowest misclassification risk. This forward iterative procedure, summarized in Algorithm 2, makes it possible to strike a balance between dimension reduction and the preservation of relevant information within the dataset. We shall compare the results obtained with the situation where the correct dimension of the reduction subspace is known in order to assess the influence of the dimension selection step.

---

**Algorithm 2** Estimation of the dimension  $d$ 

---

- Input:** Dataset  $(X, Y)$  with  $X \in \mathbb{R}^{n \times p}$  and  $Y \in \{0, 1\}^n$ , classification algorithm  $g$  (kNN, random forest...), and parameters  $\lambda > 0$ ,  $k \in \{1, \dots, n\}$ ,  $m \in \{1, \dots, n\}$  and  $K \geq 2$
- 2: **Output:** Dimension of reduction subspace  
Estimate  $\widehat{M}$  using algorithm 1 and compute the eigenvectors  $\widehat{\beta}_1, \dots, \widehat{\beta}_p$  of  $\widehat{M}$
- 4: Split  $(X, Y)$  into  $K$  folds  $(X_{(j)}, Y_{(j)})_{j=1, \dots, K}$   
**for** each  $d \in \{1, \dots, p\}$  **do**
- 6: Define  $\widehat{\beta}_{(1:d)} = [\widehat{\beta}_1 \cdots \widehat{\beta}_d]$   
**for** each  $j \in \{1, \dots, K\}$  **do**
- 8: Define  $(X_{\text{train}}, Y_{\text{train}}) = (X, Y) \setminus (X_{(j)}, Y_{(j)})$   
Train the classification rule  $g$  on data  $(X_{\text{train}} \widehat{\beta}_{(1:d)}, Y_{\text{train}})$
- 10: Evaluate its misclassification risk  $R_{j,d}$  using  $(X_{(j)} \widehat{\beta}_{(1:d)}, Y_{(j)})$   
**end for**
- 12: Compute  $R_d = \frac{1}{K} \sum_{j=1}^K R_{j,d}$   
**end for**
- 14: Return:  $d$  minimizing  $R_d$
- 

### 3.2 SIMULATION STUDY

In all the examples presented below,  $(X_i, Y_i)_{i=1, \dots, n}$  is a collection of independent and identically distributed random copies of the pair  $(X, Y)$ , and  $X$  is a vector of independent centered and unit Gaussian covariates with dimension  $p \geq 8$ . We consider the following four examples:

**Example 1.** The binary response  $Y$  has a Bernoulli distribution with parameter  $\text{expit}(X_1)$ .

**Example 2.** The binary response is  $Y = \text{sign}\{\sin(X_1) + X_2^2 + 0.2\epsilon\}$ , where  $\epsilon \sim \mathcal{N}(0, 1)$ .

**Example 3.** The binary response is  $Y = \text{sign}\{(X_1 + 0.5)(X_2 - 0.5)^2 + 0.2\epsilon\}$ , where  $\epsilon \sim \mathcal{N}(0, 1)$ .

**Example 4.** The binary response is  $Y = \text{sign}\{\log(X_1^2)(X_2^2 + X_3) + 0.2\epsilon\}$ , where  $\epsilon \sim \mathcal{N}(0, 1)$ .

Example 1 is the simplest instance of a logistic regression model with a single relevant feature, while Examples 2, 3, and 4 are closely related to examples considered by [Meng et al. \(2020\)](#). We compare the proposed approach, with penalization (denoted by  $\text{LLO}(\lambda > 0)$ ) and without penalization (called  $\text{LLO}(\lambda = 0)$ ) with three existing competitors: SAVE and PHD, which are inverse regression techniques ([Cook and Weisberg, 1991](#); [Li, 1992](#)) (both using the  $\text{dr}$  function from the R package of the same name) and POTD, relying on optimal transport ([Meng et al., 2020](#)). These methods are compared using, on the one hand, the distance between the estimated central subspace and the true central subspace using the Frobenius distance between the projection matrices on these subspaces, that is,  $d(\mathcal{S}(\widehat{\beta}), \mathcal{S}(\beta)) = \|P_{\widehat{\beta}} - P_{\beta}\|_F$  with  $P_{\beta} = \beta\beta^T$  for any orthogonal matrix  $\beta$  of  $p$ -dimensional vectors. In each example, the central subspace is explicit: in Example 1, the central subspace is spanned by the first vector of the canonical basis in  $\mathbb{R}^p$ , it is spanned by the first two vectors of this basis in Examples 2 and 3, and it is spanned by the first three vectors of this basis in Example 4. As such, Examples 1 to 4 are ranked in order of complexity of the dimension reduction problem.

On the other hand, these dimension reduction methods naturally give rise to classification procedures, in the following way: if  $P$  denotes the projection matrix on an estimated central subspace of dimension  $d$ , then a nearest-neighbor classifier at  $X = x$  is defined by the result of the majority vote among the nearest neighbors of  $Px$ . We therefore compare the nearest-neighbor classifiers obtained in this way using each of the dimension reduction procedures

we consider; for this, and in each example, we generate independently a test data set on which all the nearest-neighbor classifiers are run (for this classification task we take  $k = 10$ , as in Section 5 of Meng et al. (2020)), and their misclassification rates on this test data set are stored. It is worth noting that an easy dimension reduction problem may not automatically translate into an easy classification problem: in Example 1, for instance, the conditional probability that  $Y = 1$  is often close to  $1/2$ , meaning that the misclassification risk is bound to be high even if the central subspace is correctly identified. The distance to the central subspace and the misclassification risk thus give two different pieces of information about the accuracy of each method and the difficulty of each problem. We consider sample sizes  $n = 500, 1000, 2000, 3000, 4000$ , and in each case the misclassification risk and the distance to the central subspace are averaged over  $N = 1000$  independent replications.

Results are represented in Figures 1 and 2, first in the case when the dimension  $d$  of the reduction subspace is assumed to be known and correctly specified in each example, and the dimension of the full covariate space is  $p = 8$ . We observe that the penalized nearest-neighbor local logistic method  $\text{LLO}(\lambda > 0)$  performs best in terms of estimation of the central subspace, especially for low sample sizes. As expected, the non-penalized version of this method is substantially less accurate, which highlights the importance of penalization. The nearest-neighbor majority vote classifier fed with the estimated central subspace obtained using the  $\text{LLO}(\lambda > 0)$  gradient estimates is also best among the tested methods as far as misclassification risk is concerned, and performs almost as well as if the correct central subspace was used. For larger sample sizes, the closest competitor overall seems to be the POTD approach; for lower sample sizes, the worst two approaches are those based on the non-penalized local logistic estimation of gradients and the PHD technique. In Figure 3, we moreover examine how the methods perform when the correct dimension reduction subspace and sample size are fixed but the dimensions of the ambient space and estimated central subspace vary, *i.e.* we consider Example 4 for  $n = 1000$  with  $X$  a vector of standard Gaussian random variables having dimension  $p \in \{8, 16, 32, 64\}$ , and various dimensions  $1 \leq d \leq 6$  for the estimated central subspace. The  $\text{LLO}(\lambda > 0)$  method again performs best overall, and markedly improves over competitors when  $p \leq 16$ . It is interesting to note that the best results for  $\text{LLO}(\lambda > 0)$  are found when the dimension  $d$  is correctly specified if  $p \leq 16$ ; for higher dimensions this is not necessarily the case, although a dimension close to the dimension of the correct central subspace will tend to yield better results. The results related to Examples 1 to 3 are similar and are deferred to Figures S.1–S.2 for the sake of brevity.

We finally assess the performance of our complete workflow, namely, dimension reduction and selection through Algorithms 1 and 2 and then classification using the nearest-neighbor classifier. Again, we focus on Example 4 for the sake of brevity, in the case  $n = 1000$  with  $X$  a vector of independent centered and unit Gaussian random variables covariates having dimension  $p \in \{8, 16, 32, 64\}$ . We first ran a series of experiments to check the accuracy of our cross-validation dimension selection procedure outlined in Algorithm 2 under the same experimental design. It is seen in the left panel of Figure 4 that in the majority of cases, the dimension selected is within one unit of the correct dimension, and that dimension selection is typically quite accurate for  $p \leq 16$ . In the right panel of Figure 4, we represent the misclassification risks of the  $\text{LLO}(\lambda = 0)$  and  $\text{LLO}(\lambda > 0)$  methods following the full workflow, and we compare their misclassification rates with the misclassification probability of their versions when  $d$  is chosen as the dimension of the correct central subspace. It is

striking that the rate of misclassification tends to be lower when following the full workflow (and therefore when estimating the dimension by cross-validation on the misclassification probability) than when using the correct dimension. The interpretation is that while using the correct dimension is optimal as far as the estimation of the dimension reduction subspace is concerned, it is not necessarily so regarding the misclassification rate, and for the latter, it is better to get through the selection of  $d$  based on optimizing the misclassification risk. It is also apparent that the performance of the full workflow tends to be slightly more robust to an increase in the dimension of the ambient covariate space. The results in Examples 1 to 3 are broadly similar and can be found in Figures S.3–S.4.

### 3.3 REAL DATA ANALYSES

We apply the proposed methodology to three real data sets, all freely available from the UCI repository and on file with the authors:

- The Hill-Valley (HV) dataset<sup>2</sup>. Each data point is made of 100 real numbers  $x_i = (x_{i,j})_{1 \leq j \leq 100}$  which, when plotted in order, create a curve in the two-dimensional plane that features a hill (a “bump” in the curve) or a valley (a “dip” in the curve). The data consists of the  $n = 1212$  pairs  $(Y_i, x_i) \in \{0, 1\} \times \mathbb{R}^{100}$ , where  $Y_i = 0$  if the  $i$ th curve features a valley and 1 if it features a hill.
- The Mice Protein Expression (MPE) dataset<sup>3</sup>. After data cleaning, the dataset contains  $n = 1047$  observations with  $p = 71$  attributes, consisting of healthy mice and mice diagnosed with Down’s syndrome.
- The Wisconsin Diagnostic Breast Cancer (WDBC) dataset<sup>4</sup> used in Shin et al. (2014). A total of  $n = 569$  subjects are diagnosed with breast tumors, either benign or malignant. Ten features of breast cell nuclei are measured for each subject, with the mean, standard error, and largest values recorded for each feature, leading to  $p = 30$  predictors in total.

Each dataset is divided at random into a training set and a testing set, approximately made of 70% and 30% of the original data, respectively. Table 1 compares the performance of the complete workflow (dimension reduction and selection and then classification) using the full space of covariates, the covariates projected on the dimension reduction subspace provided by the proposed LLO( $\lambda > 0$ ) method, and its version obtained using the non-penalized version LLO( $\lambda = 0$ ), paired with either the knn or the RandomForest classifier, when applied to the testing set. The classification procedure using LLO( $\lambda > 0$ ) generally has a comparable or lower misclassification risk, a comparable or higher AUC (this can also be seen by comparing the ROC curves of each classification procedure, see Figure S.6) with comparable or lower computing time with respect to the non-penalized version, and always improves substantially upon the classifier not featuring dimension reduction. It is interesting to note that, as expected, the cross-validation procedure for the dimension of the reduction subspace involving the estimation of the matrix  $M$  via LLO( $\lambda > 0$ ) tends to select fewer components than its analogue using LLO( $\lambda = 0$ ), with comparable or higher accuracy, as shown in the top panels in Figure 5

<sup>2</sup><https://archive.ics.uci.edu/dataset/166/hill+valley>

<sup>3</sup><https://archive.ics.uci.edu/dataset/342/mice+protein+expression>

<sup>4</sup><https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic>

for the nearest neighbor classifier (see Figure S.5 for similar results with the Random Forest classifier). It is noted that in the WBDC real data analysis, all the eigenvalues of the empirical outer product  $\hat{M}$  were found to be 0 from dimension  $d = 15$  and  $d = 22$  onwards when using the LLO( $\lambda > 0$ ) and LLO( $\lambda = 0$ ) method, respectively.

Table 1: Misclassification risk, the area under the ROC curve, and computing time for each classifier applied to the three real datasets. In each case, the prediction exercise is carried out on the selected testing set.

Classifier	Random Forest			knn		
	Miscl. risk	AUC	Est. time (secs)	Miscl. risk	AUC	Est. time (secs)
Hill-Valley (HV)						
No dimension reduction	0.437	0.563	3.95	0.481	0.516	1.48
LLO( $\lambda = 0$ )	0.212	0.855	3.16	0.429	0.597	<b>0.73</b>
LLO( $\lambda > 0$ )	<b>0.115</b>	<b>0.952</b>	<b>1.44</b>	<b>0.126</b>	<b>0.953</b>	<b>0.73</b>
Mice Protein Expression (MPE)						
No dimension reduction	0.019	0.998	1.78	0.105	0.962	1.04
LLO( $\lambda = 0$ )	<b>0.013</b>	<b>0.999</b>	1.35	<b>0.07</b>	0.98	0.94
LLO( $\lambda > 0$ )	<b>0.013</b>	<b>0.999</b>	<b>0.62</b>	0.08	<b>0.985</b>	<b>0.41</b>
Wisconsin Diagnostic Breast Cancer (WBDC)						
No dimension reduction	0.064	0.979	0.88	0.058	0.972	0.41
LLO( $\lambda = 0$ )	0.053	0.987	0.8	0.058	<b>0.985</b>	<b>0.28</b>
LLO( $\lambda > 0$ )	<b>0.035</b>	<b>0.99</b>	<b>0.69</b>	<b>0.029</b>	0.982	0.32

We also compare these classifiers to those obtained following dimension reduction via the SAVE, PHD, and POTD methods. The misclassification risk is compared for estimated central subspaces of dimension  $d \in \{2, 3, 4, 5, 10\}$ . It can be seen in the bottom panels of Figure 5 that the classifier obtained after dimension reduction through the proposed LLO( $\lambda > 0$ ) approach is consistently superior to its competitors. The classifier based on LLO( $\lambda = 0$ ) dimension reduction performs worse but its performance gets closer to that of the LLO( $\lambda > 0$ )-based classifier as the dimension of the reduction subspace increases. Similar conclusions can be reached by considering the ROC curves of the different classifiers; see Figures S.7–S.9.

## ACKNOWLEDGMENTS

T. Ahmad acknowledges support from the Région Bretagne through project SAD-2021-MaEva. G. Stupfler acknowledges support from grants ANR-19-CE40-0013 (ExtremReg project), ANR-23-CE40-0009 (EXSTA project) and ANR-11-LABX-0020-01 (Centre Henri Lebesgue), the TSE-HEC ACPR Chair “Regulation and systemic risks”, and the Chair Stress Test, RISK Management and Financial Steering of the Foundation Ecole Polytechnique.



## REFERENCES

- Aghbalou, A., F. Portier, A. Sabourin, and C. Zhou (2024). Tail inverse regression: dimension reduction for prediction of extremes. *Bernoulli* 30(1), 503–533.
- Agresti, A. (2013). *Categorical Data Analysis (Third Edition)*. Wiley.
- Ausset, G., S. Cléménçon, and F. Portier (2021). Nearest neighbour based estimates of gradients: Sharp nonasymptotic bounds and applications. In *International Conference on Artificial Intelligence and Statistics*, pp. 532–540. PMLR.
- Basu, D. and C. A. B. Pereira (1983). Conditional independence in statistics. *Sankhyā: The Indian Journal of Statistics* 45(3), 324–337.
- Biau, G. and L. Devroye (2015). *Lectures on the nearest neighbor method*. Springer.
- Billingsley, P. (1999). *Convergence of Probability Measures (second edition)*. John Wiley & Sons, New York.
- Bura, E. and J. Yang (2011). Dimension estimation in sufficient dimension reduction: a unifying approach. *Journal of Multivariate analysis* 102(1), 130–142.
- Cai, Z., Y. Xia, and W. Hang (2023). An outer-product-of-gradient approach to dimension reduction and its application to classification in high dimensional space. *Journal of the American Statistical Association* 118(543), 1671–1681.
- Cook, R. D. (1994). On the interpretation of regression plots. *Journal of the American Statistical Association* 89(425), 177–189.
- Cook, R. D. (1996). Graphics for regressions with a binary response. *Journal of the American Statistical Association* 91(435), 983–992.
- Cook, R. D. (1998). Principal Hessian directions revisited. *Journal of the American Statistical Association* 93(441), 84–94.
- Cook, R. D. and L. Forzani (2009). Likelihood-based sufficient dimension reduction. *Journal of the American Statistical Association* 104(485), 197–208.
- Cook, R. D. and S. Weisberg (1991). Sliced inverse regression for dimension reduction: Comment. *Journal of the American Statistical Association* 86(414), 328–332.
- Dalalyan, A. S., A. Juditsky, and V. Spokoiny (2008). A new algorithm for estimating the effective dimension-reduction subspace. *Journal of Machine Learning Research* 9(53), 1648–1678.
- Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B* 41(1), 1–31.
- Delecroix, M., M. Hristache, and V. Patilea (2006). On semiparametric  $M$ -estimation in single-index regression. *Journal of Statistical Planning and Inference* 136(3), 730–769.
- Fan, J. and I. Gijbels (1996). *Local Polynomial Modelling and its Applications*, Volume 66 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.

- Fan, J., N. E. Heckman, and M. P. Wand (1995). Local polynomial kernel regression for generalized linear models and quasi-likelihood functions. *Journal of the American Statistical Association* 90(429), 141–150.
- Fukumizu, K., F. R. Bach, and M. I. Jordan (2004). Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research* 5(Jan), 73–99.
- Hall, P. (1989). On projection pursuit regression. *The Annals of Statistics* 17(2), 573–588.
- Härdle, W. and T. M. Stoker (1989). Investigating smooth multiple regression by the method of average derivatives. *Journal of the American Statistical Association* 84(408), 986–995.
- Hjort, N. L. and D. Pollard (1993). Asymptotics for minimisers of convex processes. Statistical Research Report, University of Oslo.
- Hristache, M., A. Juditsky, and V. Spokoiny (2001). Direct estimation of the index coefficient in a single-index model. *Annals of Statistics* 29(3), 595–623.
- Hsing, T. and H. Ren (2009). An RKHS formulation of the inverse regression dimension-reduction problem. *Annals of Statistics* 37(2), 726–755.
- Huh, J. and B. Park (2002). Likelihood-based local polynomial fitting for single-index models. *Journal of multivariate analysis* 80(2), 302–321.
- Kang, J. and S. J. Shin (2022). A forward approach for sufficient dimension reduction in binary classification. *Journal of Machine Learning Research* 23(199), 1–31.
- Klein, R. W. and R. H. Spady (1993). An efficient semiparametric estimator for binary response models. *Econometrica* 61(2), 387–421.
- Knight, K. and W. Fu (2000). Asymptotics for Lasso-type estimators. *Annals of Statistics* 28(5), 1356–1378.
- Li, B., A. Artemiou, and L. Li (2011). Principal support vector machines for linear and nonlinear sufficient dimension reduction. *Annals of Statistics* 39(6), 3182–3210.
- Li, B. and S. Wang (2007). On directional regression for dimension reduction. *Journal of the American Statistical Association* 102(479), 997–1008.
- Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* 86(414), 316–327.
- Li, K.-C. (1992). On principal Hessian directions for data visualization and dimension reduction: Another application of Stein’s lemma. *Journal of the American Statistical Association* 87(420), 1025–1039.
- Li, L. (2007). Sparse sufficient dimension reduction. *Biometrika* 94(3), 603–613.
- Meng, C., J. Yu, J. Zhang, P. Ma, and W. Zhong (2020). Sufficient dimension reduction for classification using principal optimal transport direction. *Advances in Neural Information Processing Systems* 33, 4015–4028.

- Mukherjee, S. and Q. Wu (2006). Estimation of gradients and coordinate covariation in classification. *Journal of Machine Learning Research* 7(88), 2481–2514.
- Niemiro, W. (1992). Asymptotics for  $M$ -estimators defined by convex minimization. *Annals of Statistics* 20(3), 1514–1533.
- Portier, F. (2021). Nearest neighbor empirical processes. *Bernoulli*, to appear. arXiv:2110.15083.
- Portier, F. and B. Delyon (2013). Optimal transformation: A new approach for covering the central subspace. *Journal of Multivariate Analysis* 115, 84–107.
- Portier, F. and B. Delyon (2014). Bootstrap testing of the rank of a matrix via least-squared constrained estimation. *Journal of the American Statistical Association* 109(505), 160–172.
- Powell, J. L., J. H. Stock, and T. M. Stoker (1989). Semiparametric estimation of index coefficients. *Econometrica* 57(6), 1403–1430.
- Shin, S. J., Y. Wu, H. H. Zhang, and Y. Liu (2014). Probability-enhanced sufficient dimension reduction for binary classification. *Biometrics* 70(3), 546–555.
- Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics* 10(4), 1040–1053.
- Tyler, D. E. (1981). Asymptotic inference for eigenvectors. *The Annals of Statistics* 9(4), 725–736.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*, Volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge: Cambridge University Press.
- van der Vaart, A. W. and J. A. Wellner (1996). *Weak Convergence and Empirical Processes. With Applications to Statistics*. Springer Series in Statistics. New York: Springer-Verlag.
- Xia, Y., H. Tong, W. K. Li, and L.-X. Zhu (2002). An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 64(3), 363–410.
- Zou, H. (2006). The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association* 101(476), 1418–1429.

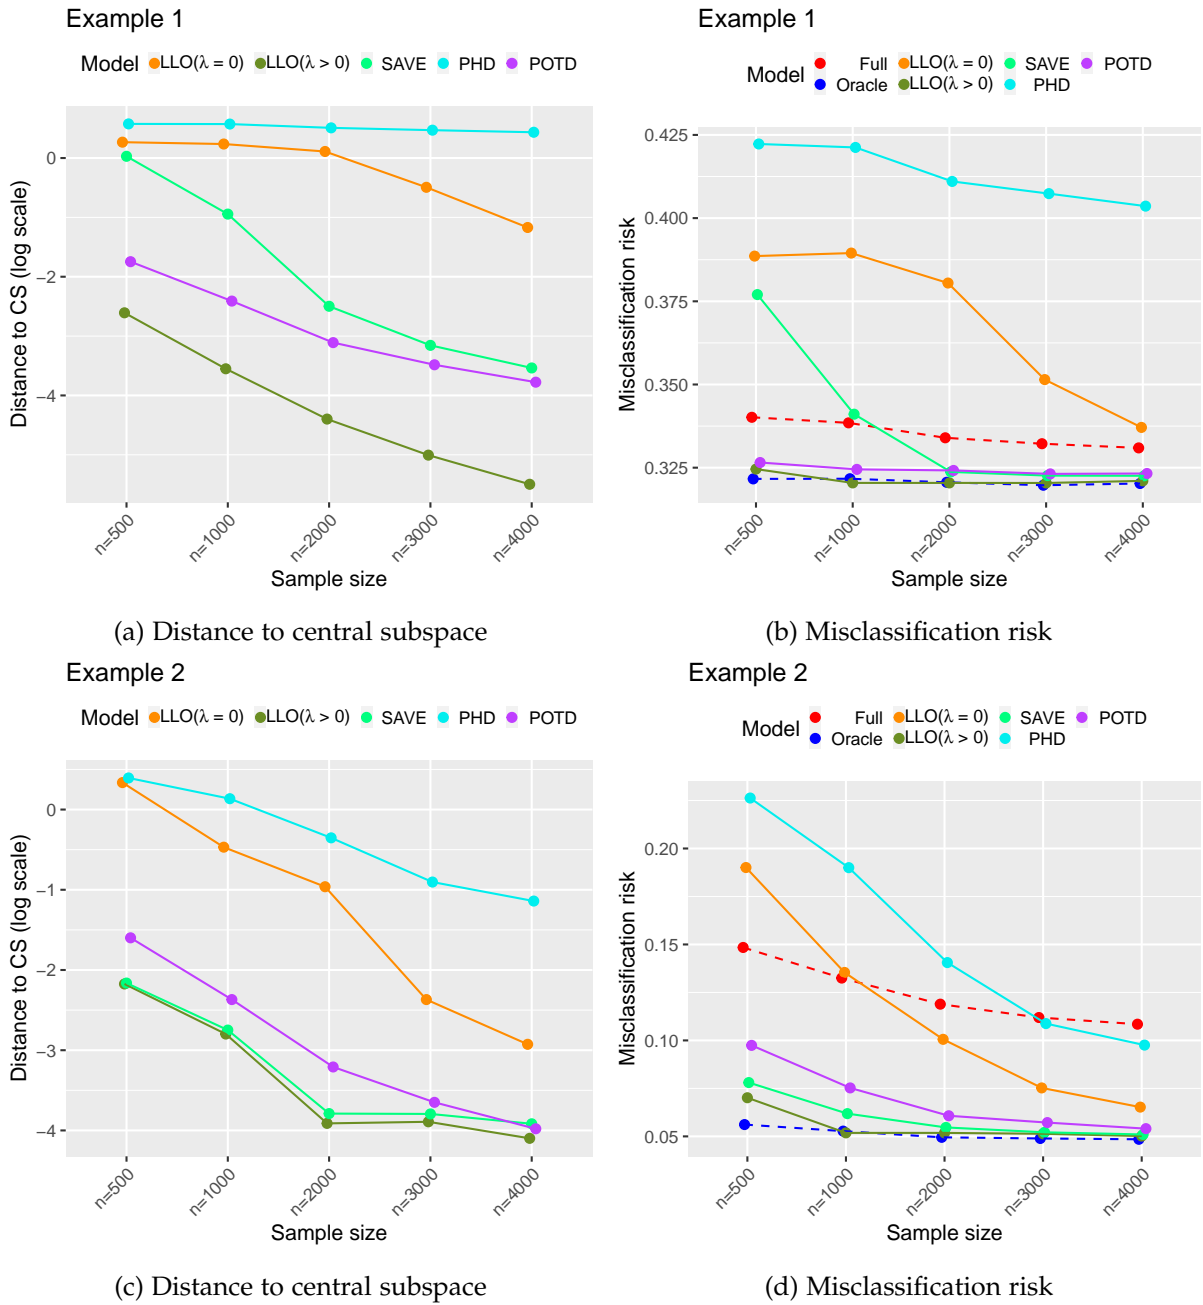


Figure 1: Simulation study – Distance to the central subspace (left, on the log scale) and misclassification risk (right) in Example 1 (top) and Example 2 (bottom), averaged over  $N = 1000$  replications in each situation. The covariate has dimension  $p = 8$  and the dimension  $d$  is chosen as the dimension of the correct population central subspace (*i.e.*  $d = 1$  in Example 1 and  $d = 2$  in Example 2). In the right-hand panels, the misclassification risk represented relates to the nearest neighbor majority vote classifier using the set of projected covariates on the estimated subspace produced by each method; in addition, “Full” denotes this classifier on the full, non-projected set of covariates, and “Oracle” denotes this classifier using the covariates projected on the correct population central subspace.




Figure 2: Simulation study – Distance to the central subspace (left) and misclassification risk (right) in Example 3 (top) and Example 4 (bottom), averaged over  $N = 1000$  replications in each situation. The covariate has dimension  $p = 8$  and the dimension  $d$  is chosen as the dimension of the correct population central subspace (*i.e.*  $d = 2$  in Example 3 and  $d = 3$  in Example 4). In the right-hand panels, the misclassification risk represented relates to the nearest neighbor majority vote classifier using the set of projected covariates on the estimated subspace produced by each method; in addition, “Full” denotes this classifier on the full, non-projected set of covariates, and “Oracle” denotes this classifier using the covariates projected on the correct population central subspace.

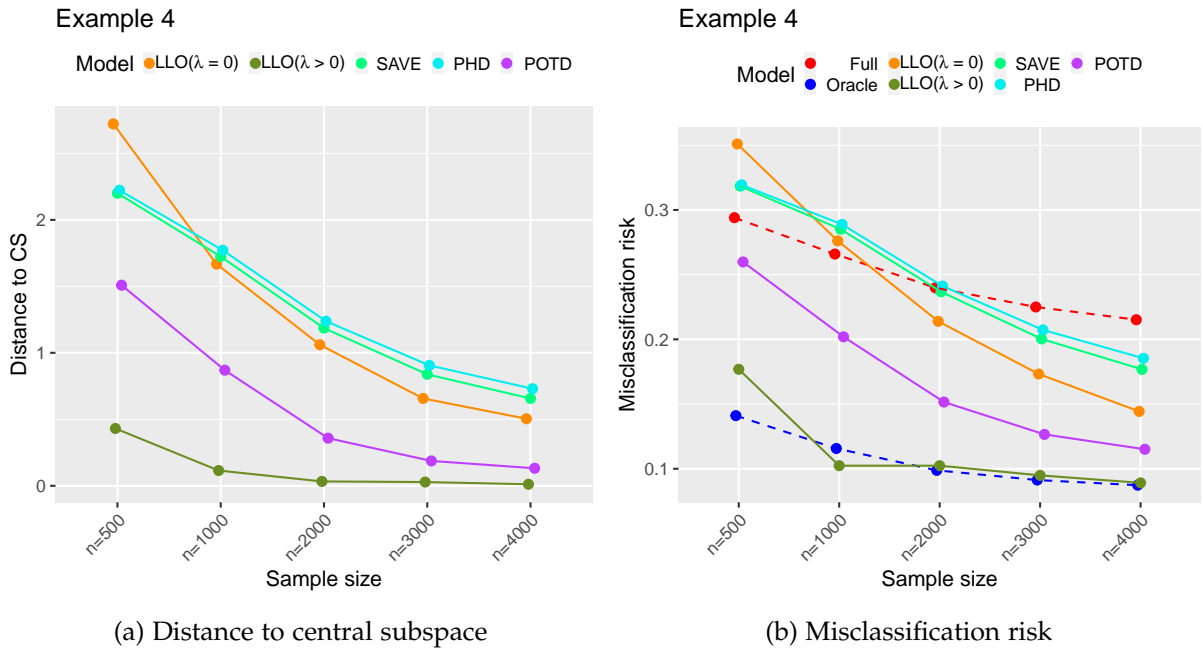


Figure 3: Simulation study – Distance to the central subspace (left) and misclassification risk (right) in Example 4, averaged over  $N = 1000$  replications of a sample of size  $n = 1000$ , as a function of the dimension  $d \in \{1, \dots, 6\}$  of the estimated central subspace and  $p \in \{8, 16, 32, 64\}$  of the full covariate space. In the right-hand panels, the red dashed line corresponds to the nearest-neighbor classifier with  $d = p$ , and the blue dashed line corresponds to this classifier using the covariates projected on the correct population central subspace.

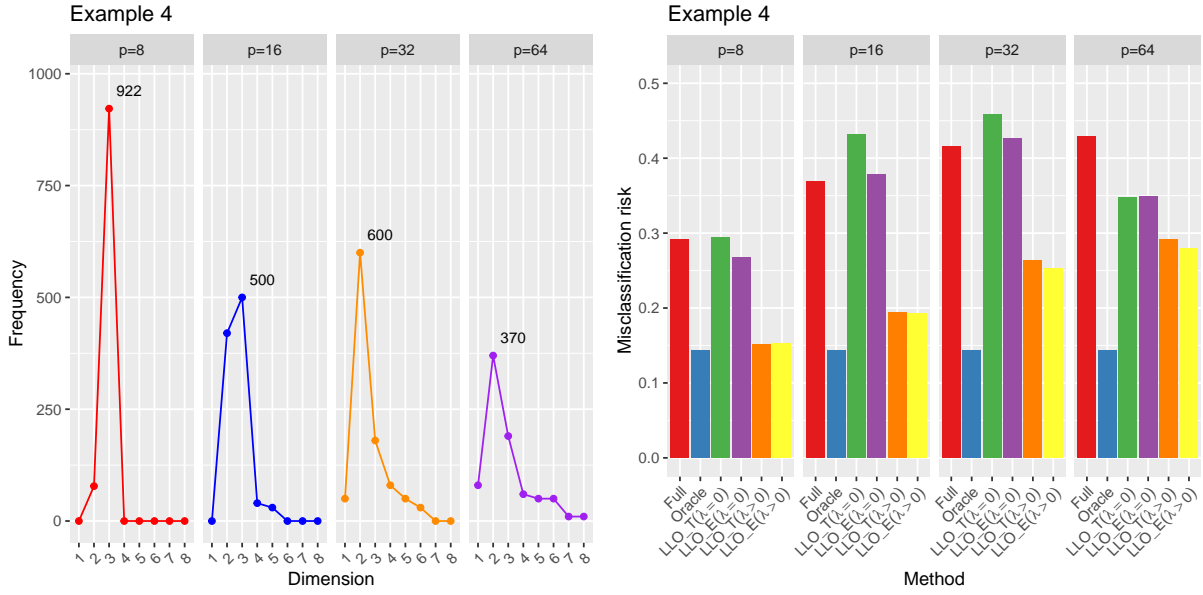


Figure 4: Simulation study – Left panel: Dimension selection through Algorithm 2, where the number indicated above the curve gives the number of times the dimension selected in the (absolute or relative) majority of cases was chosen. Right panel: Misclassification risk of the nearest-neighbor classifier with, from left to right,  $d = p$  (red bar), the covariates projected on the correct population central subspace (blue bar), the central subspace estimated using the non-penalized LLO( $\lambda = 0$ ) method under correct specification of the dimension (green bar) and with the dimension estimated by cross-validation (purple bar), and the central subspace estimated using the penalized LLO( $\lambda > 0$ ) method under correct specification of the dimension (orange bar) and with the dimension estimated by cross-validation (yellow bar). Both panels are produced using  $N = 1000$  independent replications of a sample of size  $n = 1000$  and dimensions  $p \in \{8, 16, 32, 64\}$  of the full covariate space are considered.

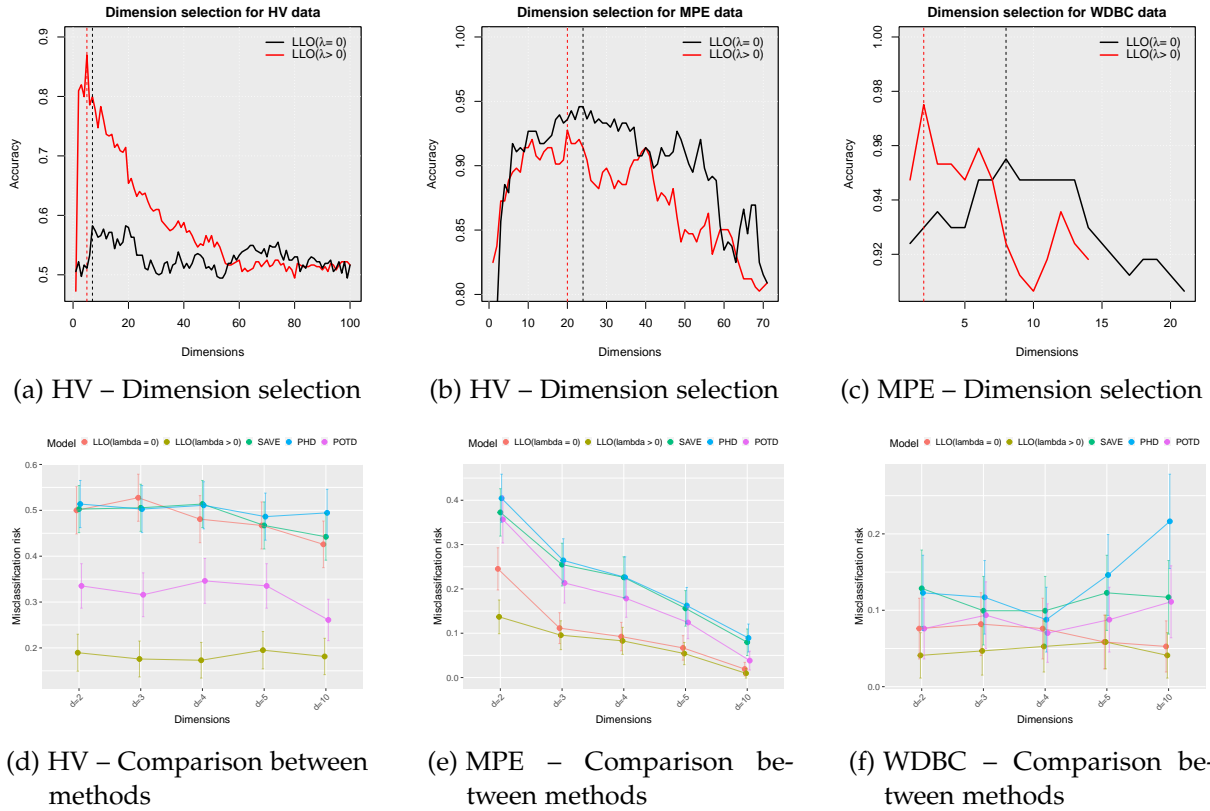


Figure 5: Real data analysis - Top panels: Dimension selection through cross-validation for  $LLO(\lambda = 0)$  and  $LLO(\lambda > 0)$ , using the nearest neighbor classifier. Bottom panels: Estimated misclassification risk related to the nearest neighbor classifier on the set of projected covariates on the estimated subspace produced by each method for a dimension  $d$  of the central subspace in  $\{2, 3, 4, 5, 10\}$ , with 95% asymptotic Gaussian Wald-type confidence intervals (the sample size used is the size of the testing set).



SUPPLEMENTARY MATERIAL FOR  
 “LOCAL LOGISTIC REGRESSION FOR DIMENSION REDUCTION IN  
 CLASSIFICATION”

S.1 MATHEMATICAL PROOFS

S.1.1 WEAK CONVERGENCE OF ARGMINs OF CONVEX FUNCTIONS

We first adapt a result from [Hjort and Pollard \(1993\)](#) to obtain a convergence result on the minima of our loss function, which has a very specific structure due to the LASSO penalty. Recall that a function  $F : \mathbb{R}^q \rightarrow \mathbb{R}$  is said to be  $\mu$ -strongly convex, for a given  $\mu > 0$ , whenever  $u \mapsto F(u) - \frac{1}{2}\mu\|u\|_2^2$  is convex. In particular:

- Strongly convex functions on  $\mathbb{R}^q$  have a unique global minimum on  $\mathbb{R}^q$ : this follows from the fact that (i) strongly convex functions on  $\mathbb{R}^q$  are strictly convex and continuous, and (ii) an equivalent definition of strong convexity of  $F$  is that, for any  $x, x_0 \in \mathbb{R}^q$  and any subgradient  $v$  of  $F$  at  $x_0$ ,

$$F(x) \geq F(x_0) + v^T(x - x_0) + \frac{\mu}{2}\|x - x_0\|_2^2, \quad (\text{S.1})$$

and therefore strongly convex functions are bounded from below by strictly convex polynomials of degree 2, meaning that they tend to  $+\infty$  as  $\|x\|_2 \rightarrow +\infty$ .

- If  $\Gamma \in \mathbb{R}^{q \times q}$  is symmetric and positive definite, and  $f : \mathbb{R}^q \rightarrow \mathbb{R}$  is convex, then  $F : u \mapsto u^T \Gamma u + f(u)$  is  $\mu$ -strongly convex, with  $\mu/2$  being the smallest eigenvalue of  $\Gamma$ .

We may now state our first result, on the convergence of minimizers of random convex functions approximated pointwise by random strongly convex functions.

**Lemma 1.** *Let  $F : \mathbb{R}^q \rightarrow \mathbb{R}$  be  $\mu$ -strongly convex,  $(S_n)$  be a random sequence and  $B_n : u \in \mathbb{R}^q \mapsto u^T S_n + F(u) \in \mathbb{R}$ . Let  $\beta_n = \arg \min_{u \in \mathbb{R}^q} B_n(u) = O_{\mathbb{P}}(1)$ .*

(i) *If  $S_n = O_{\mathbb{P}}(1)$ , then  $\beta_n = O_{\mathbb{P}}(1)$ .*

(ii) *Let moreover  $A_n : \mathbb{R}^q \rightarrow \mathbb{R}$ ,  $n \geq 1$ , be a sequence of random convex functions. If, pointwise in  $u \in \mathbb{R}^q$ ,*

$$A_n(u) - B_n(u) \xrightarrow{\mathbb{P}} 0,$$

*then  $\arg \min_{u \in \mathbb{R}^q} A_n(u)$  is nonempty for  $n$  large enough and one can construct a measurable sequence  $\alpha_n \in \mathbb{R}^q \cup \{\infty\}$ ,  $n \geq 1$ , such that  $\alpha_n \in \arg \min_{u \in \mathbb{R}^q} A_n(u)$  and  $\|\alpha_n - \beta_n\|_2 = o_{\mathbb{P}}(1)$ .*

*Proof of Lemma 1.* (i) For an arbitrary  $s \in \mathbb{R}^q$ , let  $B_s : u \in \mathbb{R}^q \mapsto u^T s + F(u) \in \mathbb{R}$ , so that  $B_n \equiv B_{S_n}$ . Then  $B_s$  is  $\mu$ -strongly convex and has a unique minimizer  $u(s) \in \mathbb{R}^q$  for any  $s$ . Let  $s, t \in \mathbb{R}^q$ . The characterization (S.1) of strong convexity and the fact that  $u(s)$  (resp.  $u(t)$ ) minimizes  $B_s$  (resp.  $B_t$ ) together imply that

$$B_s(u(t)) - B_s(u(s)) \geq \frac{\mu}{2}\|u(t) - u(s)\|_2^2 \quad \text{and} \quad B_t(u(s)) - B_t(u(t)) \geq \frac{\mu}{2}\|u(t) - u(s)\|_2^2.$$

Consequently

$$\begin{aligned} \mu \|u(t) - u(s)\|_2^2 &\leq B_s(u(t)) - B_t(u(t)) + B_t(u(s)) - B_s(u(s)) \\ &= (u(t) - u(s))^T (s - t) \\ &\leq \|u(t) - u(s)\|_2 \|t - s\|_2. \end{aligned}$$

It follows that  $u$  is Lipschitz continuous. Let now  $\varepsilon > 0$ . Since  $S_n = O_{\mathbb{P}}(1)$  and  $\beta_n = u(S_n)$ , there is a compact set  $K \subset \mathbb{R}^q$  such that  $\liminf_{n \rightarrow \infty} P(\beta_n \in u(K)) \geq \liminf_{n \rightarrow \infty} P(S_n \in K) \geq 1 - \varepsilon$ . By continuity of  $u$ , the image  $u(K)$  of  $K$  by  $u$  is compact, so that indeed  $\beta_n = O_{\mathbb{P}}(1)$ .

(ii) Lemma 2 in [Hjort and Pollard \(1993\)](#) and the discussion right before it (drawn from the appendix in [Niemiro, 1992](#)) yield the existence of a measurable sequence of minima  $\alpha_n$  of  $A_n$ , satisfying

$$\mathbb{P}(|\alpha_n - \beta_n| \geq \delta) \leq \mathbb{P}\left(\sup_{\|u - \beta_n\|_2 \leq \delta} |A_n(u) - B_n(u)| \geq \frac{1}{2} \inf_{\|u - \beta_n\|_2 = \delta} (B_n(u) - B_n(\beta_n))\right).$$

By strong convexity of  $B_n$ ,

$$\mathbb{P}(|\alpha_n - \beta_n| \geq \delta) \leq \mathbb{P}\left(\sup_{\|u - \beta_n\|_2 \leq \delta} |A_n(u) - B_n(u)| \geq \frac{\delta^2 \mu}{4}\right).$$

Consequently, it is enough to show that

$$\sup_{\|u - \beta_n\|_2 \leq \delta} |A_n(u) - B_n(u)| = o_{\mathbb{P}}(1).$$

Let  $\eta > 0$  and, for any compact set  $K \subset \mathbb{R}^q$ , let  $K_\delta$  be the (compact) set of those points  $x \in \mathbb{R}^q$  whose distance to  $K$  is not greater than  $\delta$ . Then

$$\mathbb{P}\left(\sup_{\|u - \beta_n\|_2 \leq \delta} |A_n(u) - B_n(u)| > \eta\right) \leq \mathbb{P}\left(\sup_{u \in K_\delta} |A_n(u) - B_n(u)| > \eta\right) + \mathbb{P}(\beta_n \notin K).$$

Combine Lemma 1 in [Hjort and Pollard \(1993\)](#) and (i) of the present Lemma to obtain that the left-hand side converges to 0, which is the required result.  $\square$

### S.1.2 FUNCTIONAL WEAK CONVERGENCE OF CERTAIN NEAREST-NEIGHBOR EMPIRICAL PROCESSES (PROOF OF THEOREM 1)

We first recall the following lemma on the convergence of the  $k$ -NN bandwidth; see Lemma 1 in [Portier \(2021\)](#).

**Lemma 2** ([Portier \(2021\)](#)). *Suppose that (A1) is fulfilled. Assume that  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ . Then  $\hat{\tau}_{n,k}(x)/\tau_{n,k}(x) \xrightarrow{\mathbb{P}} 1$ .*

Recall the notation

$$Z_n(\tau) = \frac{1}{\sqrt{k}} \sum_{i=1}^n \left\{ \Psi_n(Y_i, X_i) \mathbb{1}_{B(x,\tau)}(X_i) - \mathbb{E}[\Psi_n(Y, X) \mathbb{1}_{B(x,\tau)}(X)] \right\}.$$

We then state a result on the weak convergence of  $Z_n(\tau)$  as a stochastic process, which is the key to the proof of Theorem 1. Set  $s \wedge t = \min(s, t)$  and  $s \vee t = \max(s, t)$  for  $s, t \in \mathbb{R}$ , and recall the notation  $V_p$  for the volume of the unit Euclidean ball in  $\mathbb{R}^p$ . Let also  $\ell^\infty([1/2, 3/2])$  denote the space of uniformly bounded vector-valued functions defined on  $[1/2, 3/2]$  (we do not emphasize the dimension of the image space for the sake of notational convenience). This is a metric space with respect to the uniform metric  $d(f, g) = \sup_{t \in [1/2, 3/2]} \|f(t) - g(t)\|_2$ .

**Lemma 3** (Tightness and weak convergence of  $Z_n$ ). *Suppose that (A1) is fulfilled. Let  $\Psi_n : \{0, 1\} \times \mathbb{R}^p \rightarrow \mathbb{R}^q$  be a sequence of measurable vector-valued functions and suppose that there is a positive integer  $n_0$  such that*

$$\Psi_\infty := \sup_{n \geq n_0} \sup_{z \in B(x, (3/2)^{1/p} \tau_{n,k}(x))} \max(\|\Psi_n(0, z)\|_2, \|\Psi_n(1, z)\|_2) < \infty.$$

(i) *If  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ , then the stochastic process*

$$\left\{ Z_n(t^{1/p} \tau_{n,k}(x)) \right\}_{t \in [1/2, 3/2]}$$

*is tight in  $\ell^\infty([1/2, 3/2])$ .*

(ii) *Let  $\Sigma_n^2(X) = \mathbb{E}[\Psi_n(Y, X)\Psi_n(Y, X)^T | X]$ . If moreover there is a (positive semidefinite) matrix-valued function  $t \mapsto \Sigma^2(t, x)$  such that*

$$\forall t \in [1/2, 3/2], \int_{B(0,1)} \Sigma_n^2(x + \tau_{n,k}(x)t^{1/p}v) dv \rightarrow V_p \Sigma^2(t, x),$$

*then this same stochastic process converges weakly in  $\ell^\infty([1/2, 3/2])$  to a continuous Gaussian process with covariance function  $(s, t) \mapsto (s \wedge t)\Sigma^2(s \wedge t, x)$ .*

We shall prove (i) by showing tightness of any real-valued projection of the stochastic process of interest with respect to the uniform metric on  $\ell^\infty([1/2, 3/2])$  using general empirical process theory from [van der Vaart and Wellner \(1996\)](#) and [van der Vaart \(1998\)](#) linking weak convergence on  $\ell^\infty([1/2, 3/2])$  to asymptotic uniform equicontinuity on  $[1/2, 3/2]$ , equipped with the standard distance between real numbers. Continuity of the limiting process in (ii) on  $[1/2, 3/2]$  will then follow from Theorem 1.5.7 and Addendum 1.5.8 p.37 in [van der Vaart and Wellner \(1996\)](#). As a side note, let us highlight that in doing so we shall in fact check the usual tightness conditions of the space of continuous functions on  $[1/2, 3/2]$ , see Theorem 7.3 p.82 in [Billingsley \(1999\)](#); since the stochastic process of interest actually lives in the space of càdlàg functions on  $[1/2, 3/2]$ , Theorem 13.4 p.142 in [Billingsley \(1999\)](#) and its Corollary provide an alternative route to the proof of weak convergence and continuity of the limiting process.

We first recall a few definitions. Given a probability measure  $Q$  on a measurable space  $(S, \mathcal{S})$ , the metric space of square-integrable, Borel measurable real-valued functions on  $S$  with respect to  $Q$  is defined as

$$L_2(Q) = \left\{ g : (S, \mathcal{S}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R})) \text{ such that } \|g\|_{L_2(Q)}^2 := Q(g^2) := \int_S g^2 dQ < \infty \right\}.$$

For two functions  $\underline{f}, \bar{f} \in L_2(Q)$ , the *bracket*  $[\underline{f}, \bar{f}]$  is the set of all functions  $g$  in  $L_2(Q)$  such that  $\underline{f} \leq g \leq \bar{f}$  on  $S$ . A bracket  $[\underline{f}, \bar{f}]$  such that  $\|\underline{f} - \bar{f}\|_{L_2(Q)} \leq \varepsilon$  is called an  $\varepsilon$ -bracket, and for any  $\mathcal{G} \subset L_2(Q)$ , the  $\varepsilon$ -bracketing number, denoted by  $\mathcal{N}_{[]}(\mathcal{G}, L_2(Q), \varepsilon)$ , is defined as the smallest number of  $\varepsilon$ -brackets needed to cover  $\mathcal{G}$ . Finally, call an *envelope function* for  $\mathcal{G}$  any function  $G : S \rightarrow \mathbb{R}$  such that  $|g| \leq G$  on  $S$  for any  $g \in \mathcal{G}$ .

*Proof of Lemma 3.* (i) Work on the space  $T = [1/2, 3/2]$  equipped with the usual distance between real numbers. Fix  $u \in \mathbb{R}^q \setminus \{0\}$ . Let  $P$  (resp.  $P_n$ ) denote the probability measure of  $(Y, X)$  (resp. the empirical probability measure on the set of pairs  $(X_i, Y_i)$ ,  $1 \leq i \leq n$ ), so that  $u^T Z_n(t^{1/p} \tau_{n,k}(x)) = \sqrt{n}(P_n - P)(f_{n,t})$  with

$$f_{n,t}(y, z) = \sqrt{\frac{n}{k}} u^T \Psi_n(y, z) \mathbb{1}_{B(x, t^{1/p} \tau_{n,k}(x))}(z).$$

An obvious envelope function for the set of measurable functions  $\mathcal{F}_n = \{f_{n,t} : t \in T\}$  is

$$F_n(y, z) = \sqrt{\frac{n}{k}} \|u\|_2 \Psi_\infty \mathbb{1}_{B(x, (3/2)^{1/p} \tau_{n,k}(x))}(z).$$

Notice that when  $n$  is large enough,  $f_X(z) \leq 2f_X(x)$  for any  $z \in B(x, (3/2)^{1/p} \tau_{n,k}(x))$ . Therefore, for  $n$  large enough,

$$P(F_n^2) \leq 2f_X(x) \|u\|_2^2 \Psi_\infty^2 \times \frac{n}{k} \int_{B(x, (3/2)^{1/p} \tau_{n,k}(x))} dz = 3\Psi_\infty^2 < \infty.$$

According to Theorem 2.11.23 p.221 in [van der Vaart and Wellner \(1996\)](#) (see also Theorem 19.28 on p.282 in [van der Vaart, 1998](#)), it is then sufficient to show that

$$\forall \eta > 0, \lim_{n \rightarrow \infty} P(F_n^2 \mathbb{1}_{\{F_n > \eta \sqrt{n}\}}) = 0, \quad (\text{S.2})$$

$$\text{and for every } \delta_n \downarrow 0, \lim_{n \rightarrow \infty} \sup_{|t-s| \leq \delta_n} P[(f_{n,t} - f_{n,s})^2] = 0 \quad (\text{S.3})$$

$$\text{and } \lim_{n \rightarrow \infty} \int_0^{\delta_n} \sqrt{\log \mathcal{N}_{[]}(\mathcal{F}_n, L_2(P), \varepsilon \|F_n\|_{L_2(P)})} d\varepsilon = 0. \quad (\text{S.4})$$

Clearly

$$\forall \eta > 0, P(F_n^2 \mathbb{1}_{\{F_n > \eta \sqrt{n}\}}) \leq \frac{n}{k} \|u\|_2^2 \Psi_\infty^2 \mathbb{1}_{\{\|u\|_2 \Psi_\infty > \eta \sqrt{k}\}} = 0 \text{ for } n \text{ large enough.}$$

This shows (S.2). To prove (S.3), pick  $s, t \in [1/2, 3/2]$  and write, for  $n$  large enough,

$$\begin{aligned} P[(f_{n,t} - f_{n,s})^2] &\leq 2f_X(x) \|u\|_2^2 \Psi_\infty^2 \times \frac{n}{k} \int_{\mathbb{R}^p} \left( \mathbb{1}_{B(x, (s \vee t)^{1/p} \tau_{n,k}(x))}(z) - \mathbb{1}_{B(x, (s \wedge t)^{1/p} \tau_{n,k}(x))}(z) \right) dz \\ &= 2\|u\|_2^2 \Psi_\infty^2 |t - s|. \end{aligned}$$

This proves (S.3). Convergence (S.4) of the sequence of bracketing integrals is easily obtained by following the proof of Lemma 2 in [Portier \(2021\)](#): brackets for  $\sqrt{\frac{n}{k}} \mathbb{1}_{B(x, t^{1/p} \tau_{n,k}(x))}$ ,  $t \in [1/2, 3/2]$ , can be constructed as

$$\left[ \sqrt{\frac{n}{k}} \mathbb{1}_{B(x, t_j^{1/p} \tau_{n,k}(x))}, \sqrt{\frac{n}{k}} \mathbb{1}_{B(x, t_{j+1}^{1/p} \tau_{n,k}(x))} \right]$$

where the  $t_j$  make up an  $\varepsilon$ -spaced set of increasing points in  $[1/2, 3/2]$ . This concludes the proof of (i).

(ii) Fix again  $u \in \mathbb{R}^q \setminus \{0\}$  and pick  $s, t \in [1/2, 3/2]$ . By the Cramér-Wold device, and according to Theorem 2.11.23 p.221 in [van der Vaart and Wellner \(1996\)](#), weak convergence to a (tight) Gaussian process will be guaranteed if we can show that, with the notation of (i),  $P(f_{n,t} f_{n,s}) - P(f_{n,t})P(f_{n,s})$  converges to  $(s \wedge t) u^T \Sigma^2(s \wedge t, x) u$ . This Gaussian process will

then necessarily be centered because  $Z_n(\tau)$  is so, and each of its univariate projections will be continuous on  $[1/2, 3/2]$  by Theorem 1.5.7 p.37 in [van der Vaart and Wellner \(1996\)](#) and its Addendum 1.5.8.

First of all,  $|P(f_{n,t})| \leq P(F_n)$  and, for  $n$  large enough,

$$P(F_n) \leq 2f_X(x)\|u\|_2\Psi_\infty \times \sqrt{\frac{n}{k}} \int_{B(x,(3/2)^{1/p}\tau_{n,k}(x))} dz = O\left(\sqrt{\frac{k}{n}}\right) \rightarrow 0.$$

It then suffices to prove that  $P(f_{n,t}f_{n,s}) \rightarrow (s \wedge t)u^T\Sigma^2(s \wedge t, x)u$ . Now

$$P(f_{n,t}f_{n,s}) = \frac{n}{k}u^T\mathbb{E}[\Psi_n(Y, X)\Psi_n(Y, X)^T\mathbb{1}_{B(x,(s \wedge t)^{1/p}\tau_{n,k}(x))}(X)]u.$$

A change of variables gives

$$P(f_{n,t}f_{n,s}) = \frac{n}{k}(s \wedge t)\tau_{n,k}^p(x) \times u^T \left( \int_{B(0,1)} (f_X\Sigma_n^2)(x + \tau_{n,k}(x)(s \wedge t)^{1/p}v)dv \right) u.$$

Finally

$$\begin{aligned} & \left| u^T \left( \int_{B(0,1)} (f_X(x + \tau_{n,k}(x)(s \wedge t)^{1/p}v) - f_X(x))\Sigma_n^2(x + \tau_{n,k}(x)(s \wedge t)^{1/p}v)dv \right) u \right| \\ & \leq \sup_{v \in B(0,1)} |f_X(x + \tau_{n,k}(x)(s \wedge t)^{1/p}v) - f_X(x)| \times u^T \left( \int_{B(0,1)} \Sigma_n^2(x + \tau_{n,k}(x)(s \wedge t)^{1/p}v)dv \right) u \\ & \rightarrow 0 \end{aligned}$$

and then  $P(f_{n,t}f_{n,s}) \rightarrow (s \wedge t)u^T\Sigma^2(s \wedge t, x)u$ . The proof is complete.  $\square$

We can now combine Lemmas 2 and 3 to write a proof of Theorem 1.

*Proof of Theorem 1.* Write

$$Z_n(\widehat{\tau}_{n,k}(x)) = Z_n(t_n^{1/p}\tau_{n,k}(x)) = Z_n(\tau_{n,k}(x)) + \{Z_n(t_n^{1/p}\tau_{n,k}(x)) - Z_n(\tau_{n,k}(x))\}$$

with  $t_n = (\widehat{\tau}_{n,k}(x)/\tau_{n,k}(x))^p$ .

(i) By Lemma 2,  $t_n \xrightarrow{\mathbb{P}} 1$ , so that by tightness of the stochastic process  $\{Z_n(t^{1/p}\tau_{n,k}(x))\}_{t \in [1/2, 3/2]}$  in  $\ell^\infty([1/2, 3/2])$  following from Lemma 3(i), one has  $Z_n(\widehat{\tau}_{n,k}(x)) = O_{\mathbb{P}}(1)$ .

(ii) By Lemma 3(ii),  $Z_n(\tau_{n,k}(x))$  converges weakly to a Gaussian distribution with mean 0 and covariance matrix  $\Sigma^2(1, x)$ , and the stochastic process  $\{Z_n(t^{1/p}\tau_{n,k}(x))\}_{t \in [1/2, 3/2]}$  converges weakly to a continuous Gaussian process in  $\ell^\infty([1/2, 3/2])$ . It follows that  $|Z_n(t_n^{1/p}\tau_{n,k}(x)) - Z_n(\tau_{n,k}(x))| \rightarrow 0$  in probability, so that the desired weak convergence property of  $Z_n(\widehat{\tau}_{n,k}(x))$  holds.  $\square$

Besides being the crucial tool in the asymptotic analysis of the nearest-neighbor local logistic log-likelihood, Lemma 3 and Theorem 1 make it possible to obtain laws of large numbers for certain weighted averages of nearest neighbors and nearest-neighbor weighted empirical Gram matrices. These laws of large numbers will be used several times subsequently.

**Proposition 5** (Laws of large numbers for local linear nearest-neighbor estimators). *Suppose that (A1) is fulfilled. Let  $\varphi : \mathbb{R}^p \rightarrow \mathbb{R}$  be measurable, and continuous at the point  $x$ . If  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ , then*

$$\frac{1}{k} \sum_{i \in N_k(x)} \varphi(X_i) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T \xrightarrow{\mathbb{P}} \varphi(x) \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{1}{p+2} I_p \end{pmatrix}.$$

Moreover, for any symmetric matrix  $M \in \mathbb{R}^{p \times p}$ ,

$$\frac{1}{k} \sum_{i \in N_k(x)} \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix}^T M \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \xrightarrow{\mathbb{P}} \frac{\text{tr}(M)}{p+2} \begin{pmatrix} 1 \\ 0_p \end{pmatrix}.$$

*Proof of Proposition 5.* Identify in this proof the vector space of square matrices  $(p+1) \times (p+1)$  having real coefficients with  $\mathbb{R}^{(p+1)^2}$  equipped with its standard Euclidean norm  $\|\cdot\|_2$  and consider the stochastic process

$$Z_n(\tau) = \frac{1}{\sqrt{k}} \sum_{i=1}^n \left\{ \Psi_n(Y_i, X_i) \mathbb{1}_{B(x,\tau)}(X_i) - \mathbb{E}[\Psi_n(Y, X) \mathbb{1}_{B(x,\tau)}(X)] \right\}$$

with

$$\Psi_n(y, z) = \varphi(z) \begin{pmatrix} 1 \\ \frac{z-x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{z-x}{\tau_{n,k}(x)} \end{pmatrix}^T.$$

Let  $|\varphi|_\infty$  be a finite upper bound for  $|\varphi|$  in a sufficiently small neighborhood  $U$  of  $x$ , and let  $n_0$  be such that for  $n \geq n_0$ ,  $B(x, (3/2)^{1/p} \tau_{n,k}(x)) \subset U$ . Then

$$\sup_{n \geq n_0} \sup_{z \in B(x, (3/2)^{1/p} \tau_{n,k}(x))} \max(\|\Psi_n(0, z)\|_2, \|\Psi_n(1, z)\|_2) \leq |\varphi|_\infty \left( 1 + \left( \frac{3}{2} \right)^{2/p} \right).$$

Conclude, by Lemma 3(i), that the stochastic process  $\{Z_n(t^{1/p} \tau_{n,k}(x))\}_{t \in [1/2, 3/2]}$  is tight in  $\ell^\infty([1/2, 3/2])$ . It follows that, first of all,

$$\begin{aligned} \sup_{t \in [1/2, 3/2]} \left\| \frac{1}{k} \sum_{i=1}^n \varphi(X_i) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T \mathbb{1}_{B(x, t^{1/p} \tau_{n,k}(x))}(X_i) \right. \\ \left. - \frac{n}{k} \mathbb{E} \left[ \varphi(X) \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix}^T \mathbb{1}_{B(x, t^{1/p} \tau_{n,k}(x))}(X) \right] \right\|_2 \xrightarrow{\mathbb{P}} 0. \end{aligned}$$

Obviously  $\int_{B(0,1)} dz = V_p$  and  $\int_{B(0,1)} z dz = 0$ ; moreover  $\int_{B(0,1)} z_i z_j dz = 0$  when  $i \neq j$ , so that by rotational symmetry and a change to polar coordinates,

$$\int_{B(0,1)} z z^T dz = \left( \int_{B(0,1)} z_1^2 dz \right) I_p = \frac{1}{p} \left( \int_{B(0,1)} \|z\|_2^2 dz \right) I_p = \frac{V_p}{p+2} I_p.$$

Using the continuity of  $\varphi$  and  $f_X$  at  $x$  and a linear change of variables, one finds

$$\begin{aligned} \sup_{t \in [1/2, 3/2]} \left\| \frac{n}{k} \mathbb{E} \left[ \varphi(X) \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix}^T \mathbb{1}_{B(x, t^{1/p} \tau_{n,k}(x))}(X) \right] - t \varphi(x) \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{t^{2/p}}{p+2} I_p \end{pmatrix} \right\|_2 \\ \leq \frac{3}{2} \times \frac{n}{k} \tau_{n,k}^p(x) \sup_{t \in [1/2, 3/2]} \left\| \int_{B(0,1)} ((\varphi f_X)(x + t^{1/p} \tau_{n,k}(x) z) - (\varphi f_X)(x)) \begin{pmatrix} 1 & t^{1/p} z^T \\ t^{1/p} z & t^{2/p} z z^T \end{pmatrix} dz \right\|_2 \\ \rightarrow 0. \end{aligned}$$

Taking  $t = t_n = (\hat{\tau}_{n,k}(x)/\tau_{n,k}(x))^p$ , which converges to 1 in probability, then yields

$$\begin{aligned} \frac{1}{k} \sum_{i \in N_k(x)} \varphi(X_i) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T &= \frac{1}{k} \sum_{i=1}^n \varphi(X_i) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T \mathbb{1}_{B(x, \hat{\tau}_{n,k}(x))}(X_i) \\ &= \varphi(x) \begin{pmatrix} 1 & \mathbf{0}_p^T \\ \mathbf{0}_p & \frac{1}{p+2} I_p \end{pmatrix} + o_{\mathbb{P}}(1). \end{aligned}$$

This proves the first convergence. To show the second one, note it is an immediate consequence of the first convergence that

$$\frac{1}{k} \sum_{i \in N_k(x)} \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix}^T M \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix} \xrightarrow{\mathbb{P}} \frac{\text{tr}(M)}{p+2}.$$

Finally, to show that

$$\frac{1}{k} \sum_{i \in N_k(x)} \left\{ \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix}^T M \begin{pmatrix} X_i - x \\ \tau_{n,k}(x) \end{pmatrix} \right\} \frac{X_i - x}{\tau_{n,k}(x)} \xrightarrow{\mathbb{P}} \mathbf{0}_p,$$

repeat the proof of the first statement with the function

$$\Psi_n(y, z) = \left\{ \begin{pmatrix} z - x \\ \tau_{n,k}(x) \end{pmatrix}^T M \begin{pmatrix} z - x \\ \tau_{n,k}(x) \end{pmatrix} \right\} \frac{z - x}{\tau_{n,k}(x)}$$

and note that for any  $i, j, k \in \{1, \dots, p\}$ ,  $\int_{B(0,1)} z_i z_j z_k \mathbf{d}z = 0$ . □

### S.1.3 CONVERGENCE PROPERTIES OF THE NEAREST-NEIGHBOR LOCAL LOGISTIC LOG-LIKELIHOOD

It follows from (1) that proving Theorem 2 is equivalent to obtaining the convergence of

$$\begin{aligned} &(\sqrt{k}(\hat{a}_n(x) - \ell(x)), \tau_{n,k}(x)\sqrt{k}(\hat{b}_n(x) - \nabla \ell(x))) \\ &= \arg \max_{(a,b) \in \mathbb{R} \times \mathbb{R}^p} \left\{ L_n \left( \ell(x) + \frac{a}{\sqrt{k}}, \nabla \ell(x) + \frac{b}{\tau_{n,k}(x)\sqrt{k}} \right) - \lambda \left\| \nabla \ell(x) + \frac{b}{\tau_{n,k}(x)\sqrt{k}} \right\|_1 \right\}. \end{aligned}$$

The objective function is concave, so Lemma 1 suggests that it is enough to consider its convergence properties in order to recover the convergence of its minimizer. The quantity  $L_n(a, b)$  is a log-likelihood, so we analyze the convergence of the first term in the above objective function using a Taylor expansion of order 2. This requires obtaining the asymptotic behavior of the corresponding score function at  $(\ell(x), \nabla \ell(x))$ , that is

$$S_n(x) = \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} (Y_i - \text{expit}(\ell(x) + \nabla \ell(x)^T (X_i - x))) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix},$$

and of its Hessian matrix at  $(\ell(x), \nabla \ell(x))$ , namely

$$H_n(x) = -\frac{1}{k} \sum_{i \in N_k(x)} \text{expit}'(\ell(x) + \nabla \ell(x)^T (X_i - x)) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T$$

where  $\text{expit}' : s \mapsto \text{expit}(s)(1 - \text{expit}(s))$  denotes the derivative of  $\text{expit}$ . Recall the notation

$$\Gamma(x) = \pi(x)(1 - \pi(x)) \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{1}{p+2} I_p \end{pmatrix}.$$

Write  $\Delta\pi(x)$  for the Laplacian of  $\pi$  at  $x$ , *i.e.* the trace of its Hessian matrix.

**Lemma 4** (Convergence of the score function and Hessian of the local logistic log-likelihood). *Suppose that (A1) and (A2) are fulfilled. If  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ , it holds that  $S_n(x) = W_n(x) + T_n(x)$ , where*

$$W_n(x) = \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} (Y_i - \pi(X_i)) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \xrightarrow{d} \mathcal{N}(0, \Gamma(x))$$

and  $T_n(x) = \tau_{n,k}^2(x) \sqrt{k} \left( \frac{1}{2(p+2)} \left( \Delta\pi(x) - \frac{1 - 2\pi(x)}{\pi(x)(1 - \pi(x))} \|\nabla\pi(x)\|_2^2 \right) \begin{pmatrix} 1 \\ 0_p \end{pmatrix} + o_{\mathbb{P}}(1) \right).$

Moreover,  $H_n(x) \xrightarrow{\mathbb{P}} -\Gamma(x)$ .

The control of  $T_n(x)$  uses the following consequence of the Taylor formula with integral form of the remainder: let, for any function  $F : \mathbb{R}^q \rightarrow \mathbb{R}$  that is twice continuously differentiable at a point  $z$ ,  $H_F(z')$  denote its Hessian matrix at the point  $z'$  when  $z'$  is close enough to  $z$ . Then

$$\lim_{\eta \rightarrow 0} \sup_{\substack{z' \in B(z, \eta) \\ z' \neq z}} \frac{1}{\|z' - z\|_2^2} \left| F(z') - F(z) - (z' - z)^T \nabla F(z) - \frac{1}{2} (z' - z)^T H_F(z) (z' - z) \right| = 0. \quad (\text{S.5})$$

More precisely, for  $z' \neq z$  but close enough to  $z$ ,

$$\begin{aligned} & \frac{1}{\|z' - z\|_2^2} \left| F(z') - F(z) - (z' - z)^T \nabla F(z) - \frac{1}{2} (z' - z)^T H_F(z) (z' - z) \right| \\ & \leq \frac{1}{2} \|H_F(z') - H_F(z)\|_2 \end{aligned} \quad (\text{S.6})$$

where in the upper bound  $\|\cdot\|_2$  is the operator norm induced by the Euclidean norm on  $\mathbb{R}^q$ .

*Proof of Lemma 4.* Since  $\text{expit}' = \text{expit}(1 - \text{expit})$ , the convergence in probability of  $H_n(x)$  to  $-\Gamma(x)$  is an obvious consequence of Proposition 5. We concentrate on the convergence of  $S_n(x)$ . Write

$$\begin{aligned} S_n(x) &= \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} (Y_i - \pi(X_i)) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \\ &\quad - \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} (\text{expit}(\ell(x) + \nabla\ell(x)^T(X_i - x)) - \pi(X_i)) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} =: W_n(x) + T_n(x). \end{aligned}$$

We study first the quantity  $W_n(x)$  and we then examine  $T_n(x)$ .

**Convergence of  $W_n(x)$ :** Consider the stochastic process

$$Z_n(\tau) = \frac{1}{\sqrt{k}} \sum_{i=1}^n \left\{ \Psi_n(Y_i, X_i) \mathbb{1}_{B(x, \tau)}(X_i) - \mathbb{E}[\Psi_n(Y, X) \mathbb{1}_{B(x, \tau)}(X)] \right\}$$



with

$$\Psi_n(y, z) = (y - \pi(z)) \begin{pmatrix} 1 \\ \frac{z-x}{\tau_{n,k}(x)} \end{pmatrix}.$$

Then

$$\sup_{n \geq 1} \sup_{z \in B(x, (3/2)^{1/p} \tau_{n,k}(x))} \max(\|\Psi_n(0, z)\|_2, \|\Psi_n(1, z)\|_2) \leq \sqrt{1 + \left(\frac{3}{2}\right)^{2/p}} < \infty.$$

Obviously  $\mathbb{E}[\Psi_n(Y, X) \mathbb{1}_{B(x, \tau)}(X)] = 0$  since  $\mathbb{E}[Y - \pi(X) | X] = 0$ , and

$$\begin{aligned} \Sigma_n^2(X) &:= \mathbb{E} \left[ (Y - \pi(X))^2 \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix}^T \middle| X \right] \\ &= \pi(X)(1 - \pi(X)) \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X-x}{\tau_{n,k}(x)} \end{pmatrix}^T. \end{aligned}$$

Recall also from the proof of Proposition 5 that

$$\int_{B(0,1)} zz^T dz = \frac{V_p}{p+2} I_p.$$

Assumption (A2) then yields, for any  $t > 0$ ,

$$\int_{B(0,1)} \Sigma_n^2(x + \tau_{n,k}(x)t^{1/p}v) dv \rightarrow V_p \pi(x)(1 - \pi(x)) \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{t^{2/p}}{p+2} I_p \end{pmatrix}.$$

For  $t = 1$  the right-hand side is exactly  $V_p \Gamma(x)$ . Conclude, using Theorem 1(ii), that  $Z_n(\hat{\tau}_{n,k}(x)) \xrightarrow{d} \mathcal{N}(0, \Gamma(x))$ , that is,

$$\begin{aligned} \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} (Y_i - \pi(X_i)) \begin{pmatrix} 1 \\ \frac{X_i-x}{\tau_{n,k}(x)} \end{pmatrix} &= \frac{1}{\sqrt{k}} \sum_{i=1}^n (Y_i - \pi(X_i)) \begin{pmatrix} 1 \\ \frac{X_i-x}{\tau_{n,k}(x)} \end{pmatrix} \mathbb{1}_{B(x, \hat{\tau}_{n,k}(x))}(X_i) \\ &\xrightarrow{d} \mathcal{N}(0, \Gamma(x)) \end{aligned}$$

as announced.

**Convergence of  $T_n(x)$ :** Fix  $\varepsilon > 0$ . Write, for any  $z$  close enough to  $x$ ,

$$\begin{aligned} \text{expit}(\ell(x) + \nabla \ell(x)^T(z-x)) - \pi(z) &= \text{expit}(\ell(x) + \nabla \ell(x)^T(z-x)) - \pi(x) - \nabla \pi(x)^T(z-x) \\ &\quad - (\pi(z) - \pi(x) - \nabla \pi(x)^T(z-x)). \end{aligned}$$

Recall that  $\pi \mapsto \text{logit}(\pi)$  has derivative  $1/(\pi(1-\pi))$  on  $(0, 1)$ . As a consequence,  $\text{expit}(\ell(x))(1 - \text{expit}(\ell(x))) \nabla \ell(x) = \nabla \pi(x)$ . Moreover,  $s \mapsto \text{expit}(s)$  has second derivative  $s \mapsto \text{expit}(s)(1 - \text{expit}(s))(1 - 2 \text{expit}(s))$ . Applying (S.5), first to the function  $\text{expit}$  and then to the function  $\pi$ , which is twice continuously differentiable at  $x$  by (A2), leads to the existence of  $\eta > 0$  such that for all  $z \in B(x, \eta)$ ,

$$\begin{aligned} &\left| \text{expit}(\ell(x) + \nabla \ell(x)^T(z-x)) - \pi(x) - \nabla \pi(x)^T(z-x) - \frac{1}{2} \frac{1 - 2\pi(x)}{\pi(x)(1 - \pi(x))} (\nabla \pi(x)^T(z-x))^2 \right| \\ &\leq \frac{\varepsilon}{2} \|z - x\|_2^2 \end{aligned}$$

and

$$\left| \pi(z) - \pi(x) - \nabla \pi(x)^T (z - x) - \frac{1}{2} (z - x)^T H_\pi(x) (z - x) \right| \leq \frac{\varepsilon}{2} \|z - x\|_2^2.$$

Since by Lemma 2 we have  $\widehat{\tau}_{n,k}(x)/\tau_{n,k}(x) \rightarrow 1$  in probability and  $\tau_{n,k}(x) \rightarrow 0$ , one may conclude that

$$T_n(x) = \frac{1}{2\sqrt{k}} \sum_{i \in N_k(x)} \left( (X_i - x)^T H_\pi(x) (X_i - x) - \frac{1 - 2\pi(x)}{\pi(x)(1 - \pi(x))} (\nabla \pi(x)^T (X_i - x))^2 \right) \left( \frac{1}{\tau_{n,k}(x)} \right) + o_{\mathbb{P}}(\tau_{n,k}^2(x)\sqrt{k}).$$

Now

$$\begin{aligned} & \frac{1}{\sqrt{k}} \sum_{i \in N_k(x)} \left( (X_i - x)^T H_\pi(x) (X_i - x) - \frac{1 - 2\pi(x)}{\pi(x)(1 - \pi(x))} (\nabla \pi(x)^T (X_i - x))^2 \right) \left( \frac{1}{\tau_{n,k}(x)} \right) \\ &= \tau_{n,k}^2(x)\sqrt{k} \left( \frac{1}{k} \sum_{i \in N_k(x)} \left( \frac{X_i - x}{\tau_{n,k}(x)} \right)^T H_\pi(x) \left( \frac{X_i - x}{\tau_{n,k}(x)} \right) \left( \frac{1}{\tau_{n,k}(x)} \right) \right. \\ & \quad \left. - \frac{1}{k} \sum_{i \in N_k(x)} \frac{1 - 2\pi(x)}{\pi(x)(1 - \pi(x))} \left( \frac{X_i - x}{\tau_{n,k}(x)} \right)^T \nabla \pi(x) \nabla \pi(x)^T \left( \frac{X_i - x}{\tau_{n,k}(x)} \right) \left( \frac{1}{\tau_{n,k}(x)} \right) \right) \end{aligned}$$

and the second convergence of Proposition 5 applies. The proof is complete.  $\square$

#### S.1.4 PROOF OF THEOREM 2

Note that  $k/n \rightarrow 0$  because  $\tau_{n,k}^2(x)\sqrt{k}$  is bounded. Let the sequence of rescaling matrices  $D_n$  be defined as

$$D_n = \frac{1}{\sqrt{k}} \begin{pmatrix} 1 & 0_p^T \\ 0_p & \frac{1}{\tau_{n,k}(x)} I_p \end{pmatrix}$$

so that

$$\forall (a, b) \in \mathbb{R} \times \mathbb{R}^p, \frac{1}{\sqrt{k}} \begin{pmatrix} a \\ b \end{pmatrix} = D_n \begin{pmatrix} a \\ b \end{pmatrix}.$$

We seek to apply Lemma 1(ii) with  $A_n(u) = -A_{n,1}(u) + A_{n,2}(u)$ , where, for given  $u = (u_0, u_1, \dots, u_p) \in \mathbb{R}^{p+1}$ , we let

$$\begin{aligned} A_{n,1}(u) &= L_n \left( \begin{pmatrix} \ell(x) \\ \nabla \ell(x) \end{pmatrix} + D_n u \right) - L_n \begin{pmatrix} \ell(x) \\ \nabla \ell(x) \end{pmatrix}, \\ A_{n,2}(u) &= \lambda \left\{ \left\| \nabla \ell(x) + \frac{1}{\tau_{n,k}(x)\sqrt{k}} u_{(1:p)} \right\|_1 - \|\nabla \ell(x)\|_1 \right\}, \text{ with } u_{(1:p)} = \begin{pmatrix} u_1 \\ \vdots \\ u_p \end{pmatrix}, \end{aligned}$$

and, if  $S_n = S_n(x)$  denotes the score function of Lemma 4,

$$\begin{aligned} B_n(u) &= -u^T S_n(x) + \frac{1}{2} u^T \Gamma(x) u \\ & \quad + (cf_X(x) V_p)^{1/p} \left( \sum_{j=1}^p \text{sgn}(\nabla \ell_j(x)) u_j \mathbb{1}_{\{\nabla \ell_j(x) \neq 0\}} + |u_j| \mathbb{1}_{\{\nabla \ell_j(x) = 0\}} \right). \end{aligned}$$

Since  $\Gamma(x)$  is a positive definite matrix and  $S_n(x)$  is bounded in probability by Lemma 4, the function  $B_n$  satisfies the assumptions of Lemma 1. Noting that  $L_n$  is concave, it follows that  $A_n$  is convex, so that it is enough to show that  $A_n - B_n$  converges pointwise to 0 in probability in order to apply Lemma 1(ii). Given that  $H_n(x)$  converges to  $-\Gamma(x)$  in probability by Lemma 4 again, it is sufficient to prove that

$$R_n(u) = A_{n,1}(u) - u^T S_n(x) - \frac{1}{2} u^T H_n(x) u \xrightarrow{\mathbb{P}} 0$$

and  $A_{n,2}(u) \xrightarrow{\mathbb{P}} (cf_X(x) V_p)^{1/p} \left\{ \sum_{j=1}^p \text{sgn}(\nabla \ell_j(x)) u_j \mathbb{1}_{\{\nabla \ell_j(x) \neq 0\}} + |u_j| \mathbb{1}_{\{\nabla \ell_j(x) = 0\}} \right\}$ .

The remainder term  $R_n(u)$  is dealt with using the following lemma.

**Lemma 5** (Pointwise approximation of the nearest-neighbor local logistic log-likelihood). *Suppose that (A1) and (A2) are fulfilled. If  $k := k_n \rightarrow \infty$  is such that  $k/n \rightarrow 0$ , we have*

$$\forall u \in \mathbb{R}^{p+1}, |R_n(u)| \leq \frac{\|u\|_2^3}{2\sqrt{2}} \times \frac{1}{\sqrt{k}} \left( \frac{\widehat{\tau}_{n,k}(x)}{\tau_{n,k}(x)} \vee 1 \right)^3,$$

and in particular  $|R_n(u)| = O_{\mathbb{P}}(1/\sqrt{k})$ .

*Proof of Lemma 5.* From (S.6), we find that if  $F : \mathbb{R}^{p+1} \rightarrow \mathbb{R}$  is a twice continuously differentiable function having a Lipschitz continuous Hessian matrix  $u \mapsto H_F(u)$ , that is, there is  $C > 0$  with  $\|H_F(v) - H_F(u)\|_2 \leq C\|v - u\|_2$  for any  $u, v \in \mathbb{R}^{p+1}$ , then

$$\left| F(v) - F(u) - (v - u)^T \nabla F(u) - \frac{1}{2} (v - u)^T H_F(u) (v - u) \right| \leq \frac{C}{2} \|v - u\|_2^3. \quad (\text{S.7})$$

We apply (S.7) with  $F : u \mapsto L_n(z + D_n u)$  where  $z = (\ell(x), \nabla \ell(x)^T)^T$ . We know that the gradient of  $F$  at  $0_{p+1}$  is  $S_n(x)$  and its Hessian matrix at  $0_{p+1}$  is  $H_n(x)$ ; more generally

$$H_F(u) = -\frac{1}{k} \sum_{i \in N_k(x)} \text{expit}' \left( (z + D_n u)^T \begin{pmatrix} 1 \\ X_i - x \end{pmatrix} \right) \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix}^T.$$

Now  $s \mapsto \text{expit}''(s) = \text{expit}(s)(1 - \text{expit}(s))(1 - 2 \text{expit}(s))$  is bounded by  $1/4$ , so by the mean value theorem,

$$\left| \text{expit}' \left( (z + D_n u)^T \begin{pmatrix} 1 \\ X_i - x \end{pmatrix} \right) - \text{expit}' \left( z^T \begin{pmatrix} 1 \\ X_i - x \end{pmatrix} \right) \right| \leq \frac{1}{\sqrt{k}} \times \frac{1}{4} \|u\|_2 \left\| \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \right\|_2.$$

Therefore

$$\|H_F(u) - H_F(0_{p+1})\|_2 \leq \frac{1}{\sqrt{k}} \times \frac{1}{4} \|u\|_2 \times \frac{1}{k} \sum_{i \in N_k(x)} \left\| \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \right\|_2^3.$$

Inequality (S.7) applies and yields, for any  $u \in \mathbb{R}^{p+1}$ ,

$$|R_n(u)| \leq \frac{1}{\sqrt{k}} \times \frac{1}{8} \|u\|_2^3 \times \frac{1}{k} \sum_{i \in N_k(x)} \left\| \begin{pmatrix} 1 \\ \frac{X_i - x}{\tau_{n,k}(x)} \end{pmatrix} \right\|_2^3 \leq \frac{\|u\|_2^3}{2\sqrt{2}} \times \frac{1}{\sqrt{k}} \left( \frac{\widehat{\tau}_{n,k}(x)}{\tau_{n,k}(x)} \vee 1 \right)^3$$

as required. The conclusion on the rate of convergence of  $R_n(u)$  to 0 follows from Lemma 2.  $\square$

To show the convergence of  $A_{n,2}(u)$ , note that for each  $z \neq 0$  and  $u \in \mathbb{R}$ , there is  $t > 0$  small enough such that  $|z + tu| - |z| = tu \operatorname{sgn}(z)$ ; besides, if  $z = 0$ , then  $|z + tu| - |z| = t|u|$ . Working componentwise then immediately yields the below lemma, from which the pointwise limit of  $A_{n,2}$  follows.

**Lemma 6.** *For any  $u, v \in \mathbb{R}^p$ , there exists  $t_0 > 0$  such that for all  $t \in [0, t_0]$ ,*

$$\|v + tu\|_1 - \|v\|_1 = t \sum_{j=1}^p \operatorname{sgn}(v_j) u_j \mathbb{1}_{\{v_j \neq 0\}} + |u_j| \mathbb{1}_{\{v_j = 0\}}.$$

### S.1.5 PROOF OF COROLLARY 4

The proof follows from the decomposition

$$\begin{aligned} \|\widehat{M} - M\|_F &\leq \left\| \frac{1}{m} \sum_{i=1}^m \{\widehat{b}_n(X_i^*) \widehat{b}_n(X_i^*)^T - \nabla \ell(X_i^*) \nabla \ell(X_i^*)^T\} \right\|_F \\ &\quad + \left\| \frac{1}{m} \sum_{i=1}^m \nabla \ell(X_i^*) \nabla \ell(X_i^*)^T - \int_{\mathbb{R}^p} \nabla \ell(x) \nabla \ell(x)^T d\mu(x) \right\|_F, \end{aligned}$$

where  $\|\cdot\|_F$  is the Frobenius norm. The first term in the above upper bound can be treated using Theorem 2:

$$\begin{aligned} \left\| \frac{1}{m} \sum_{i=1}^m \{\widehat{b}_n(X_i^*) \widehat{b}_n(X_i^*)^T - \nabla \ell(X_i^*) \nabla \ell(X_i^*)^T\} \right\|_F &\leq \max_{x \in \operatorname{supp}(\mu)} \|\widehat{b}_n(x) \widehat{b}_n(x)^T - \nabla \ell(x) \nabla \ell(x)^T\|_F \\ &= O_{\mathbb{P}} \left( \frac{1}{\tau_{n,k}(x) \sqrt{k}} \right) + O_{\mathbb{P}}(\tau_{n,k}(x)). \end{aligned}$$

The other term is controlled by computing its second moment:

$$\begin{aligned} &\mathbb{E} \left( \left\| \frac{1}{m} \sum_{i=1}^m \nabla \ell(X_i^*) \nabla \ell(X_i^*)^T - \int_{\mathbb{R}^p} \nabla \ell(x) \nabla \ell(x)^T d\mu(x) \right\|_F^2 \right) \\ &= \frac{1}{m} \mathbb{E} \left( \left\| \nabla \ell(X^*) \nabla \ell(X^*)^T - \int_{\mathbb{R}^p} \nabla \ell(x) \nabla \ell(x)^T d\mu(x) \right\|_F^2 \right) \\ &\leq \frac{1}{m} \mathbb{E} \left( \left\| \nabla \ell(X^*) \nabla \ell(X^*)^T \right\|_F^2 \right) = \frac{1}{m} \int_{\mathbb{R}^p} \|\nabla \ell(x)\|_2^4 d\mu(x) \leq \frac{1}{m} \max_{x \in \operatorname{supp}(\mu)} \|\nabla \ell(x)\|_2^4. \end{aligned}$$

The upper bound is finite by assumption.  $\square$

## S.2 ADDITIONAL NUMERICAL RESULTS

### S.2.1 SIMULATION STUDY

Similarly to Figure 3, we represent in Figures S.1 and S.2 the distance to the central subspace and the misclassification risk in Examples 1, 2 and 3 when the correct dimension reduction subspace and sample size are fixed but the dimensions of the ambient space and estimated

central subspace vary, *i.e.* we fix  $n = 1000$ , simulate a vector  $X$  of independent centered and unit Gaussian random variables covariates having dimension  $p \in \{8, 16, 32, 64\}$ , and the estimated central subspace has dimension  $1 \leq d \leq 6$ . Again, the proposed LLO( $\lambda > 0$ ) method consistently yields superior results across all evaluated settings. As  $p$  increases, the performances of LLO( $\lambda = 0$ ) and LLO( $\lambda > 0$ ) get closer, while the performances of the SAVE and POTD methods tend to substantially deteriorate. Similarly to Figure 4, we also provide in Figures S.3 and S.4 results related to the dimension  $d$  selected through the full workflow summarized in Algorithms 1 and 2 and the misclassification risk of the hence obtained nearest-neighbor classifier in Examples 1, 2 and 3.

## S.2.2 REAL DATA ANALYSES

Figures S.5–S.9 contain extra results about the number of selected components by the proposed method with the random forest classifier, and further elements about the predictive quality of the nearest neighbor and random forest classifiers when paired with one of the dimension reduction methods we consider.




Figure S.1: Simulation study – Distance to the central subspace and misclassification risk in Examples 1, 2 and 3, averaged over  $N = 1000$  replications of a sample of size  $n = 1000$ , as a function of the dimension  $d \in \{1, \dots, 6\}$  of the estimated central subspace and  $p \in \{8, 16, 32, 64\}$  of the full covariate space.

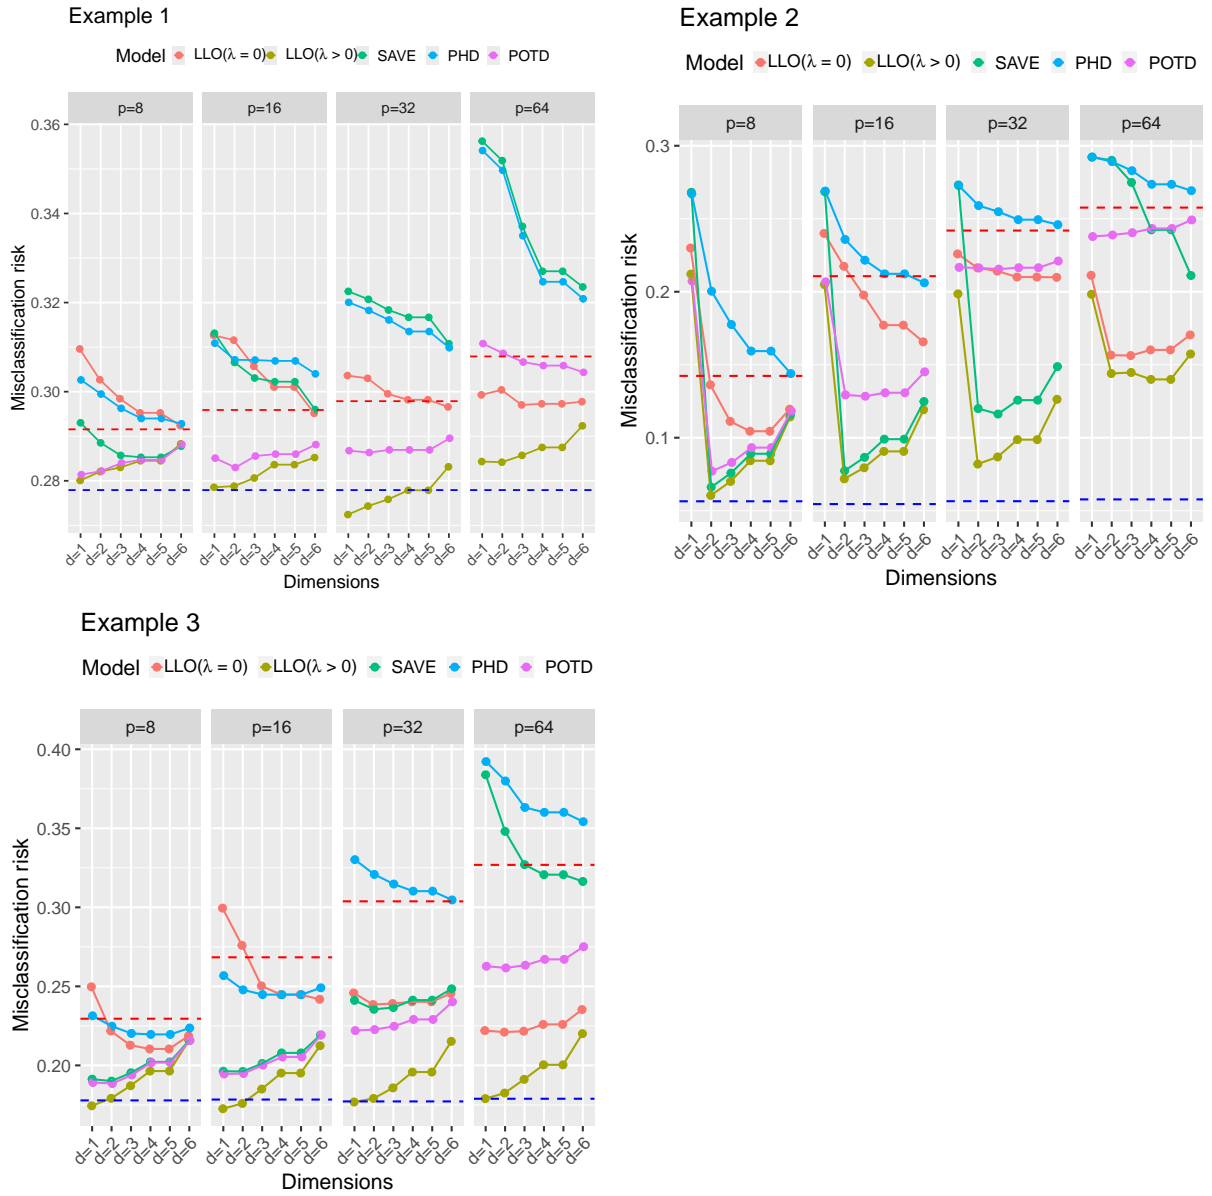


Figure S.2: Simulation study – Misclassification risk in Examples 1, 2 and 3, averaged over  $N = 1000$  replications of a sample of size  $n = 1000$ , as a function of the dimension  $d \in \{1, \dots, 6\}$  of the estimated central subspace and  $p \in \{8, 16, 32, 64\}$  of the full covariate space. In the right-hand panels, the red dashed line corresponds to the nearest-neighbor classifier with  $d = p$ , and the blue dashed line corresponds to this classifier using the covariates projected on the correct population central subspace.




Figure S.3: Simulation study – Dimension selection through Algorithm 2 over  $N = 1000$  independent replications of a sample of size  $n = 1000$ , as a function of the dimension  $p \in \{8, 16, 32, 64\}$  of the full covariate space. In each panel the number indicated above the curve gives the number of times the dimension selected in the (absolute or relative) majority of cases was chosen.






Figure S.4: Simulation study – Misclassification risk of the nearest-neighbor classifier with, from left to right,  $d = p$  (red bar), the covariates projected on the correct population central subspace (blue bar), the central subspace estimated using the non-penalized LLO( $\lambda = 0$ ) method under correct specification of the dimension (green bar) and with the dimension estimated by cross-validation (purple bar), and the central subspace estimated using the penalized LLO( $\lambda > 0$ ) method under correct specification of the dimension (orange bar) and with the dimension estimated by cross-validation (yellow bar). All panels are produced using  $N = 1000$  independent replications of a sample of size  $n = 1000$  and considering dimensions  $p \in \{8, 16, 32, 64\}$  of the full covariate space.







Figure S.5: Real data analysis – Dimension selection through cross-validation for  $LLO(\lambda = 0)$  and  $LLO(\lambda > 0)$ , using the random forest classifier. In the WDBC real data analysis, all the eigenvalues of the empirical outer product  $\hat{M}$  were found to be 0 from dimension  $d = 15$  and  $d = 22$  onwards when using the  $LLO(\lambda > 0)$  and  $LLO(\lambda = 0)$ , respectively.




(a) HV




(b) HV




(c) MPE



(d) MPE



(e) WDBC



(f) WDBC

Figure S.6: Real data analysis – ROC curve of the nearest neighbor classifier (left) and random forest classifier (right) with no dimension reduction (black curve), dimension reduction following the LLO( $\lambda = 0$ ) procedure (red curve) and dimension reduction following the LLO( $\lambda > 0$ ) procedure (blue curve). In each case, the prediction exercise is carried out on the selected testing set.




Figure S.7: Real data analysis, HV dataset – ROC curve of the nearest neighbor classifier following the LLO( $\lambda = 0$ ), LLO( $\lambda > 0$ ), SAVE, PHD and POTD dimension reduction procedures, for a dimension  $d$  of the dimension reduction subspace in  $\{2, 3, 4, 5, 10\}$ . In each case, the prediction exercise is carried out on the selected testing set.

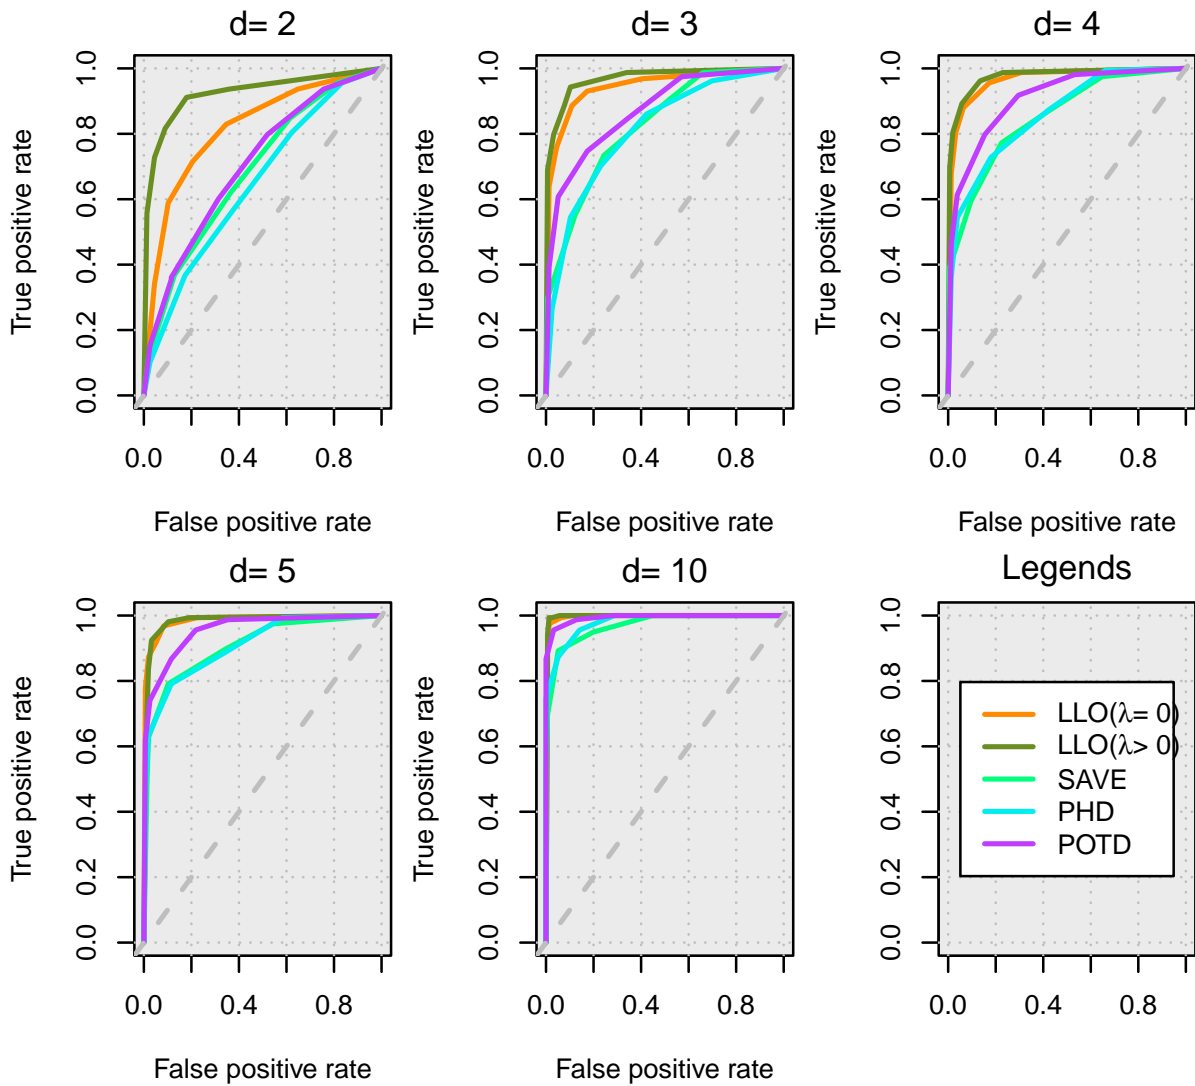


Figure S.8: Real data analysis, MPE dataset – ROC curve of the nearest neighbor classifier following the LLO( $\lambda = 0$ ), LLO( $\lambda > 0$ ), SAVE, PHD and POTD dimension reduction procedures, for a dimension  $d$  of the dimension reduction subspace in  $\{2, 3, 4, 5, 10\}$ . In each case, the prediction exercise is carried out on the selected testing set.

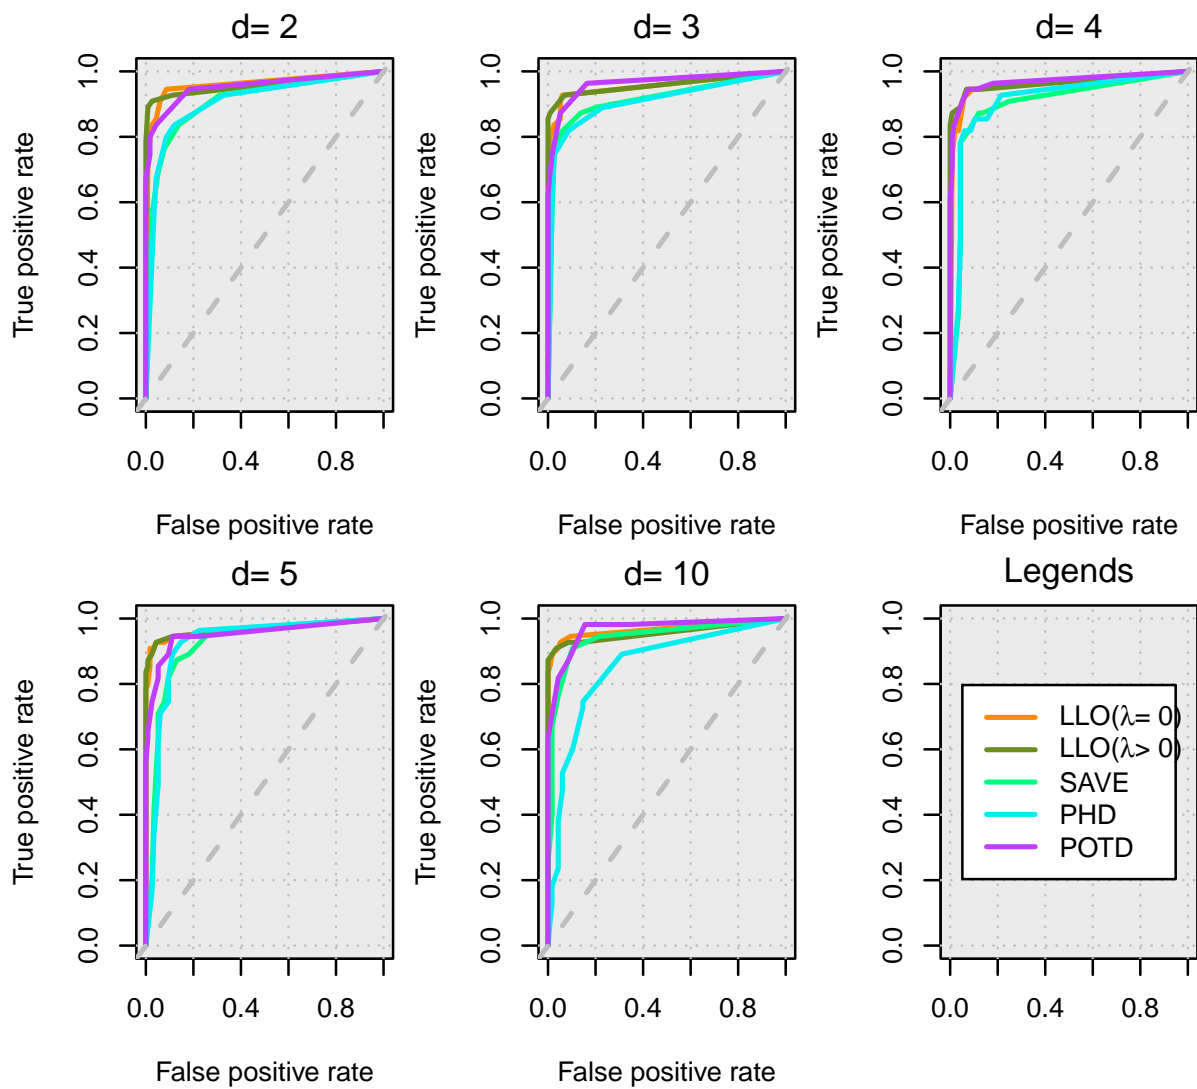


Figure S.9: Real data analysis, WDBC dataset – ROC curve of the nearest neighbor classifier following the  $LLO(\lambda = 0)$ ,  $LLO(\lambda > 0)$ , SAVE, PHD and POTD dimension reduction procedures, for a dimension  $d$  of the dimension reduction subspace in  $\{2, 3, 4, 5, 10\}$ . In each case, the prediction exercise is carried out on the selected testing set.