

# Boosting Nearest Neighbors for the Efficient Estimation of Posteriors

Roberto D’Ambrosio<sup>1,3</sup>, Richard Nock<sup>2</sup>, Wafa Bel Haj Ali<sup>3</sup>, Frank Nielsen<sup>4</sup>,  
and Michel Barlaud<sup>3,5</sup>

<sup>1</sup> University Campus Bio-Medico of Rome, Rome, Italy  
`r.dambrosio@unicampus.it`

<sup>2</sup> CEREGMIA - Université Antilles-Guyane, Martinique, France  
`rnock@martinique.univ-ag.fr`

<sup>3</sup> CNRS - U. Nice, France  
`{belhajal,barlaud}@i3s.unice.fr`

<sup>4</sup> Sony Computer Science Laboratories, Inc., Tokyo, Japan  
`Frank.Nielsen@acm.org`

<sup>5</sup> Institut Universitaire de France

**Abstract.** It is an admitted fact that mainstream boosting algorithms like AdaBoost do not perform well to estimate class conditional probabilities. In this paper, we analyze, in the light of this problem, a recent algorithm, UNN, which leverages nearest neighbors while minimizing a convex loss. Our contribution is threefold. First, we show that there exists a subclass of surrogate losses, elsewhere called balanced, whose minimization brings simple and statistically efficient estimators for Bayes posteriors. Second, we show *explicit* convergence rates towards these estimators for UNN, for any such surrogate loss, under a Weak Learning Assumption which parallels that of classical boosting results. Third and last, we provide experiments and comparisons on synthetic and real datasets, including the challenging SUN computer vision database. Results clearly display that boosting nearest neighbors may provide highly accurate estimators, sometimes more than a hundred times more accurate than those of other contenders like support vector machines.

## 1 Introduction

*Boosting* refers to the iterative combination of classifiers which produces a classifier with reduced true risk (with high probability), while the base classifiers may be weakly accurate [1]. The final, *strong* classifier  $h$ , satisfies  $\text{im}(h) \subseteq \mathbb{R}$ . Such an output carries out two levels of information. The simplest one is the sign of the output. This discrete value is sufficient to classify an unknown observation  $\mathbf{x}$ :  $h(\mathbf{x})$  predicts that  $\mathbf{x}$  belongs to a class of interest iff it is positive. The most popular boosting results typically rely on this sole information [1–3] (and many others). The second level is the real value itself, which carries out as additional information a magnitude which can be interpreted as a “confidence” in the classification. This continuous information may be fit into a link function

$f : \mathbb{R} \rightarrow [0, 1]$  to estimate conditional class probabilities, thus lifting the scope of boosting to that of Bayes decision rule [4]:

$$\hat{\mathbf{Pr}}[y = 1 | \mathbf{x}] = f(h(\mathbf{x})) . \quad (1)$$

To date, estimating posteriors with boosting has not met the same success as predicting (discrete) labels. It is widely believed that boosting and conditional class probability estimation are, up to a large extent, in conflict with each other, as boosting iteratively improves classification at the price of progressively overfitting posteriors [4, 5]. Experimentally, limiting overfitting is usually obtained by tuning the algorithms towards early stopping [6]. Very recently, a new algorithm was proposed to leverage the famed nearest neighbor (NN) rules [7, 8]. This algorithm, UNN, fits real-valued coefficients for examples in order to minimize a surrogate risk [9, 2]. These leveraging coefficients are used to balance the votes in the final  $k$ -NN rule. It is proven that, as the number of iterations  $T \rightarrow \infty$ , UNN achieves the global optimum of the surrogate risk at hand for a wide class of surrogates called strictly convex surrogates [10, 2]. An explicit convergence rate is obtained for the specific case of the exponential loss, under a so-called “weak index assumption” [8], generalized in [7]. Our contribution is threefold. First, we show that there exists a subclass of surrogate losses, elsewhere called *balanced*, whose minimization brings simple and efficient estimators for Bayes posteriors (1). Second, we show explicit convergence rates for UNN for *any* such surrogate loss under a Weak Learning Assumption which parallels that of classical boosting results [3]. Third and last, we provide experiments on simulated and real domains, displaying that boosting nearest neighbors brings very good results from the conditional class probabilities estimation standpoint, *without* the overfitting problem of classical boosting approaches. A serious challenger to the popular logistic estimator for posteriors estimation also emerges. We end up with the conclusion that learning posteriors with boosting nearest neighbors benefits from two advantages. First, the weak classifiers being simple examples, they naturally limit the risk of overfitting compared to more complex weak learners. Second, we end up learning posteriors using a *natural, fixed* topology of data, and not an *ad hoc* topology relying on an induced classifier.

The remaining of the paper is structured as follows: the next Section presents definitions, followed by a Section on convex losses and the estimation of posteriors. Then, a Section presents algorithms and results on boosting nearest neighbors. The two last Sections present experiments with discussions, and conclude.

## 2 Definitions

### 2.1 Estimation

Our setting is that of multiclass multilabel classification (See *e.g.* [3]). We have access to an input set of  $m$  examples, also called prototypes,  $\mathcal{S} \doteq \{(\mathbf{x}_i, \mathbf{y}_i), i = 1, 2, \dots, m\}$ . Vector  $\mathbf{y}_i \in \{-1, 1\}^C$  encodes class memberships, assuming  $y_{ic} = 1$  means that observation  $\mathbf{x}_i$  belongs to class  $c$ .  $\mathcal{S}$  is sampled i.i.d. according to an

unknown distribution  $\mathcal{D}$ . Given an observation  $\mathbf{x} \in \mathcal{O}$ , we wish to estimate the conditional class probabilities for each class  $c$ , also called (estimated) posteriors:

$$\hat{p}_c(\mathbf{x}) \doteq \hat{\mathbf{Pr}}[y_c = 1|\mathbf{x}] . \quad (2)$$

We note  $p_c(\mathbf{x}) \doteq \mathbf{Pr}_{\mathcal{D}}[y_c = 1|\mathbf{x}]$  the corresponding Bayes (true) posteriors.

## 2.2 Surrogates, losses and risks

Perhaps the simplest road towards computing these estimators consists in first crafting  $C$  separate classification problems, each of which leads to estimators for one class (2). Normalizing estimators to 1 over the  $C$  classes yields the values in (2). Each of these  $C$  problems is a one-versus-all classification task, say for class  $c$ , with corresponding sample  $\mathcal{S}^{(c)} = \{(\mathbf{x}_i, y_{ic}), i = 1, 2, \dots, m\}$ . For each of these problems, we learn from  $\mathcal{S}$  a classifier  $h : \mathcal{O} \rightarrow \mathbb{R}$  out of which we may accurately compute (2), typically with  $\hat{p}_c(\mathbf{x}) = f(h(\mathbf{x}))$  for some relevant function  $f$ . More sophisticated approaches exist that reduce the number of classifiers by folding classes in observation variables [4, 3]. Each of them equivalently learn on a sample of  $\Omega(mC)$  examples, and it is an easy task to craft from their output a set of  $C$  classifiers that fit into the framework we consider.

There exists a convenient approach to carry out this path as a whole, for each class  $c = 1, 2, \dots, C$ : learn  $h$  by minimizing a *surrogate risk* over  $\mathcal{S}$  [9, 2, 10]. A surrogate risk has general expression:

$$\varepsilon_{\mathcal{S}}^{\psi}(h, c) \doteq \frac{1}{m} \sum_{i=1}^m \psi(y_{ic}h(\mathbf{x}_i)) , \quad (3)$$

for some function  $\psi$  that we call a *surrogate loss*. Quantity  $y_{ic}h(\mathbf{x}) \in \mathbb{R}$  is called the *edge* of classifier  $h$  on example  $(\mathbf{x}_i, y_i)$  for class  $c$ . The surrogate risk is an estimator of the *true surrogate risk* computed over  $\mathcal{D}$ :

$$\varepsilon_{\mathcal{D}}^{\psi}(h, c) \doteq \mathbf{E}_{\mathcal{D}}[\psi(y_{ic}h(\mathbf{x}))] . \quad (4)$$

Any surrogate loss relevant to classification [9] has to meet  $\text{sign}(h_{\text{opt}}(\mathbf{x}^*)) = \text{sign}(2\mathbf{Pr}_{\mathcal{D}}[y_c = 1|\mathbf{x} = \mathbf{x}^*] - 1)$ , where  $h_{\text{opt}}$  minimizes  $\mathbf{E}_{\mathcal{D}}[\psi(y_c h(\mathbf{x}))|\mathbf{x} = \mathbf{x}^*]$ . Hence, the sign of the optimal classifier  $h_{\text{opt}}$  is as accurate to predict class membership as Bayes decision rule. This Fisher consistency requirement for  $\psi$  is called *classification calibration* [9]. We focus in this paper on the subclass of classification calibrated surrogates that are strictly convex and differentiable.

**Definition 1.** [2] A *strictly convex loss* is a strictly convex function  $\psi$  differentiable on  $\text{int}(\text{dom}(\psi))$  satisfying (i)  $\text{im}(\psi) \subseteq \mathbb{R}^+$ , (ii)  $\text{dom}(\psi)$  symmetric around 0, (iii)  $\nabla_{\psi}(0) < 0$ .

Definition 1 is extremely general: should we have removed conditions (i) and (ii), Theorem 6 in [9] brings that it would have encompassed the intersection between strictly convex differentiable functions and classification calibrated functions.

Conditions (i) and (ii) are mainly conveniences for classification: in particular, it is not hard to see that modulo scaling by a positive constant, the surrogate risk (3) is an upperbound of the empirical risk for any strictly convex loss. Minimizing the surrogate risk amounts thus to minimize the empirical risk up to some extent. We define the Legendre conjugate of any strictly convex loss  $\psi$  as  $\psi^*(x) \doteq x\nabla_{\psi}^{-1}(x) - \psi(\nabla_{\psi}^{-1}(x))$ . There exists a particular subset of strictly convex losses of independent interest [2]. A function  $\phi : [0, 1] \rightarrow \mathbb{R}^+$  is called *permissible* iff it is differentiable on  $(0, 1)$ , strictly concave and symmetric around  $x = 1/2$  [11, 2]. We adopt the notation  $\bar{\phi} = -\phi$  [2].

**Definition 2.** [2] *Given some permissible  $\phi$ , we let  $\psi_{\phi}$  denote the **balanced convex loss** with signature  $\phi$  as:*

$$\psi_{\phi}(x) \doteq \frac{\bar{\phi}^*(-x) - \phi(0)}{\phi(1/2) - \phi(0)} . \quad (5)$$

Balanced convex losses have an important rationale: up to differentiability constraints, they match the set of symmetric lower-bounded losses defining proper scoring rules [2], that is, basically, the set of losses that fit to classification problems without class-dependent misclassification costs. Table 1 provides examples of surrogate losses, most of which are strictly convex surrogates, some of which are balanced convex surrogates. We have derived Amari’s  $\alpha$ -loss from Amari’s famed  $\alpha$  divergences [12] (proof omitted). The linear Hinge loss is *not* a balanced convex loss, yet it figures the limit behavior of balanced convex losses [2]. Remark that all signatures  $\phi$  are well-known in the domain of decision-tree induction : from the top-most to the bottom-most, one may recognize Gini criterion, the entropy (two expressions), Matsushita’s criterion and the empirical risk [11, 10].

### 2.3 One dimensional exponential families and posteriors estimation

A (regular) one dimensional *exponential family* [12] is a set of probability density functions whose elements admit the following canonical form:

$$p[x|\theta] \doteq \exp(x\theta - \psi(\theta)) p_0(x) , \quad (6)$$

where  $p_0(x)$  normalizes the density,  $\psi$  is a strictly convex differentiable function that we call the *signature* of the family, and  $\theta$  is the density’s natural parameter. It was shown in [2] that the efficient minimization of any balanced convex surrogate risk — *i.e.* a surrogate risk with a balanced convex loss — amounts to a maximum likelihood estimation  $\hat{\theta} = H(\mathbf{x})$  at some  $\mathbf{x}$  for an exponential family whose signature depends solely on the permissible function  $\phi$ . [2] suggest to use the corresponding *expected* parameter of the exponential family as the posterior:

$$\hat{\Pr}[y = 1|\mathbf{x}] = \hat{\Pr}_{\phi}[y = 1|\mathbf{x}; H] \doteq \nabla_{\bar{\phi}}^{-1}(H(\mathbf{x})) \in [0, 1] . \quad (7)$$

$\nabla_{\bar{\phi}}^{-1}$  plays the role of the link function (1). The quality of such an estimator shall be addressed in the following Section.

|   | $\psi$  | $\hat{p}_c(\mathbf{x})$   | $\phi$                                   |
|---|---|---|--|
| A | $(1-x)^2$   | $\frac{1}{2}(1+x)$  | $x(1-x)$                                 |
| B | $\log_2(1+\exp(-x))$  | $[1+\exp(-x)]^{-1}$   | $\frac{-x \ln x}{-(1-x) \ln(1-x)}$       |
| C | $\log_2(1+2^{-x})$  | $[1+2^{-x}]^{-1}$   | $\frac{-x \log_2 x}{-(1-x) \log_2(1-x)}$ |
| D | $-x + \sqrt{1+x^2}$   | $\frac{1}{2} \left( 1 + \frac{x}{\sqrt{1+x^2}} \right)$   | $\sqrt{x(1-x)}$                          |
| E | $\frac{1}{2}x(\text{sign}(x) - 1)$                                    | $\begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x < 0 \end{cases}$                          | $2 \min\{x, 1-x\}$                       |
| F | $\exp(-x)$  | $[1+\exp(-2x)]^{-1}$  | N/A                                      |
| G | $\left(1 + \frac{1-\alpha^2}{4}x\right)^{-\frac{1+\alpha}{1-\alpha}}$ | $\left[1 + \left(\frac{4-(1-\alpha^2)x}{4+(1-\alpha^2)x}\right)^{\frac{2}{1-\alpha}}\right]^{-1}$ | N/A                                      |

Table 1: Examples of surrogates  $\psi$  (Throughout the paper, we let  $\ln$  denote the base- $e$  logarithm, and  $\log_z(x) \doteq \ln(x)/\ln(z)$  denote the base- $z$  logarithm). From top to bottom, the losses are known as: squared loss, (normalized) logistic loss, binary logistic loss, Matsushita loss [10, 2], linear Hinge loss, exponential loss, Amari’s  $\alpha$ -loss, for  $\alpha \in (-1, 1)$  [2]. Strictly convex losses are A, B, C, D, F, G. Balanced convex losses are A, B, C, D (E corresponds to a limit behavior of balanced convex losses [2]). For each  $\psi$ , we give the corresponding estimators  $\hat{p}_c(\mathbf{x})$  (Theorem 1 and Eqs (9, 11) below: replace  $x$  by  $h_{\text{opt}}(\mathbf{x})$ ), and if they are balanced convex losses, the corresponding concave signature  $\phi$  (See text for details).

### 3 Strictly convex losses and the efficient estimation of posteriors

There is a rationale to use (7) as the posterior: the duality between natural and expectation parameters of exponential families, via Legendre duality [9, 2], and the fact that the domain of the expectation parameter of one dimensional exponential families whose signature is (minus) a permissible function is the interval  $[0, 1]$  [2]. We improve below this rationale, with the proof that *Bayes posteriors* satisfy (7) for the classifier which is the population minimizer of (7).

**Theorem 1.** *Suppose  $\psi$  strictly convex differentiable. The true surrogate risk  $\mathbf{E}_{\mathcal{D}}[\psi(y_{ic}h(\mathbf{x}))]$  is minimized at the unique  $h_{\text{opt}}(\mathbf{x})$  satisfying:*

$$\frac{\nabla_{\psi}(-h_{\text{opt}}(\mathbf{x}))}{\nabla_{\psi}(h_{\text{opt}}(\mathbf{x}))} = \frac{p_c(\mathbf{x})}{1-p_c(\mathbf{x})}. \quad (8)$$

Furthermore, if  $\psi$  is a balanced convex loss, then the population minimizer  $h_{\text{opt}}$  of  $\mathbf{E}_{\mathcal{D}}[\psi_{\phi}(y_{ic}h(\mathbf{x}))]$  satisfies:

$$p_c(\mathbf{x}) = \nabla_{\phi}^{-1}(h_{\text{opt}}(\mathbf{x})), \quad (9)$$

for which

$$\mathbf{E}_{\mathcal{D}}[\psi_{\phi}(y_{ic}h_{\text{opt}}(\mathbf{x}))] = \frac{\phi(p_c(\mathbf{x})) - \phi(0)}{\phi(1/2) - \phi(0)}. \quad (10)$$

(Proof omitted) Table 1 provides examples of expressions for  $p_c(\mathbf{x})$  as in (9). Eq. (8) in Theorem (1) brings that we may compute an estimator  $\hat{p}_c(\mathbf{x})$  as:

$$\hat{p}_c(\mathbf{x}) = \frac{\nabla_{\psi}(-h(\mathbf{x}))}{\nabla_{\psi}(h(\mathbf{x})) + \nabla_{\psi}(-h(\mathbf{x}))} . \quad (11)$$

This simple expression is folklore, at least for the logistic and exponential losses [6, 4]. The essential contribution of Theorem 1 relies on bringing a strong rationale to the use of (7), as the estimators converge to Bayes posteriors in the infinite sample case. Let us give some finite sample properties for the estimation (7). We show that the sample-wise estimators of (9) are efficient estimators of (9); this is not a surprise, but comes from properties of exponential families [13]. What is perhaps more surprising is that the corresponding aggregation of classifiers is not a linear combination of all estimating classifiers, but a generalized  $\nabla_{\phi}^{-1}$ -mean.

**Theorem 2.** *Suppose we sample  $n$  datasets  $\mathcal{S}_j^{(c)}, j = 1, 2, \dots, n$ . Denote  $\hat{h}_{\text{opt},j}$  the population minimizer for  $\mathbf{E}_{\mathcal{S}_j^{(c)}}[\psi_{\phi}(y_{ic}h(\mathbf{x}))]$ . Then each  $\hat{p}_{c,j}(\mathbf{x}) \doteq \nabla_{\phi}^{-1}(\hat{h}_{\text{opt},j}(\mathbf{x}))$  is the only efficient estimator for  $p_c(\mathbf{x})$ . The corresponding classifier  $\hat{h}_{\text{opt}}$  aggregating all  $\hat{h}_{\text{opt},j}$ , is:  $\hat{h}_{\text{opt}}(\mathbf{x}) \doteq \nabla_{\phi}^{-1}\left(\frac{1}{n_{\mathbf{x}}} \sum_{j:(\mathbf{x}, \cdot) \in \mathcal{S}_j^{(c)}} \nabla_{\phi}^{-1}(\hat{h}_{\text{opt},j}(\mathbf{x}))\right), \forall \mathbf{x} \in \cup_j \mathcal{S}_j$ , where  $1 \leq n_{\mathbf{x}} \leq n$  is the number of subsets containing  $\mathbf{x}$ .*

(Proof omitted)

## 4 Leveraging and boosting Nearest Neighbors

The nearest neighbor rule belongs to the oldest, simplest and most widely studied classification algorithms [14, 15]. We denote by  $\text{NN}_k(\mathbf{x})$  the set of the  $k$ -nearest neighbors (with integer constant  $k > 0$ ) of an example  $(\mathbf{x}, \mathbf{y})$  in set  $\mathcal{S}$  with respect to a non-negative real-valued “distance” function. This function is defined on domain  $\mathcal{O}$  and measures how much two observations differ from each other. This dissimilarity function thus may not necessarily satisfy the triangle inequality of metrics. For the sake of readability, we let  $j \sim_k \mathbf{x}$  denote the assertion that example  $(\mathbf{x}_j, \mathbf{y}_j)$  belongs to  $\text{NN}_k(\mathbf{x})$ . We shall abbreviate  $j \sim_k \mathbf{x}_i$  by  $j \sim_k i$ . To classify an observation  $\mathbf{x} \in \mathcal{O}$ , the  $k$ -NN rule  $\mathcal{H}$  over  $\mathcal{S}$  computes the sum of class vectors of its nearest neighbors, that is:  $\mathcal{H}(\mathbf{x}) = \sum_{j \sim_k \mathbf{x}} \mathbf{1} \circ \mathbf{y}_j$ , where  $\circ$  is the Hadamard product.  $\mathcal{H}$  predicts that  $\mathbf{x}$  belongs to each class whose corresponding coordinate in the final vector is positive. A *leveraged*  $k$ -NN rule is a generalization of this to:

$$\mathcal{H}(\mathbf{x}) = \sum_{j \sim_k \mathbf{x}} \boldsymbol{\alpha}_j \circ \mathbf{y}_j , \quad (14)$$

where  $\boldsymbol{\alpha}_j \in \mathbb{R}^C$  is a leveraging vector for the classes in  $\mathbf{y}_j$ . Leveraging approaches to nearest neighbors are not new [16, 17], yet to the best of our knowledge no

---

**Algorithm 1** Algorithm UNIVERSAL NEAREST NEIGHBORS,  
UNN( $\mathcal{S}, \psi, k$ )

---

**Input:**  $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i), i = 1, 2, \dots, m, \mathbf{x}_i \in \mathcal{O}, \mathbf{y}_i \in \{-1, 1\}^C\}$ ,  $\psi$  strictly convex loss (Definition 1),  $k \in \mathbb{N}_*$ ;

Let  $\boldsymbol{\alpha}_j \leftarrow \mathbf{0}, \forall j = 1, 2, \dots, m$ ;

**for**  $c = 1, 2, \dots, C$  **do**

    Let  $\mathbf{w} \leftarrow -\nabla_{\psi}(0)\mathbf{1}$ ;

**for**  $t = 1, 2, \dots, T$  **do**

**[I.0]** Let  $j \leftarrow \text{WIC}(\mathcal{S}, \mathbf{w})$ ;

**[I.1]** Let  $\delta_j \in \mathbb{R}$  solution of:

$$\sum_{i: j \sim_k i} y_{ic} y_{jc} \nabla_{\psi} \left( \delta_j y_{ic} y_{jc} + \nabla_{\psi}^{-1}(-w_i) \right) = 0 ; \quad (12)$$

**[I.2]**  $\forall i : j \sim_k i$ , let

$$w_i \leftarrow -\nabla_{\psi} \left( \delta_j y_{ic} y_{jc} + \nabla_{\psi}^{-1}(-w_i) \right) , \quad (13)$$

**[I.3]** Let  $\alpha_{jc} \leftarrow \alpha_{jc} + \delta_j$ ;

**Output:**  $\mathcal{H}(\mathbf{x}) \doteq \sum_{j \sim_k \mathbf{x}} \boldsymbol{\alpha}_j \circ \mathbf{y}_j$

---

convergence results or rates were known, at least until the algorithm UNN [7, 8]. Algorithm 1 gives a simplified version of the UNN algorithm of [7, 8] which learns a leveraged  $k$ -NN. Oracle  $\text{WIC}(\mathcal{S}, \mathbf{w})$  is the analogous for NN of the classical weak learners for boosting: it takes learning sample  $\mathcal{S}$  and weights  $\mathbf{w}$  over  $\mathcal{S}$ , and returns the index of some example in  $\mathcal{S}$  which is to be leveraged. [8] prove that for any strictly convex loss  $\psi$ , UNN converges to the global optimum of the surrogate risk at hand. However, they prove boosting-compliant convergence rates only for the exponential loss. For all other strictly convex losses, there is no insight on the rates with which UNN may converge towards the optimum of the surrogate risk at hand. We now provide such explicit convergence rates under the following *Weak Learning Assumption*:

**WLA:** There exist some  $\vartheta > 0, \varrho > 0$  such that, given any  $k \in \mathbb{N}_*$ ,  $c = 1, 2, \dots, C$  and any distribution  $\mathbf{w}$  over  $\mathcal{S}$ , the weak index chooser oracle  $\text{WIC}$  returns an index  $j$  such that the following two statements hold:

(i)  $\Pr_{\mathbf{w}}[j \sim_k i] \geq \varrho$ ;

(ii)  $\Pr_{\mathbf{w}}[y_{jc} \neq y_{ic} | j \sim_k i] \leq 1/2 - \vartheta$  or  $\Pr_{\mathbf{w}}[y_{jc} \neq y_{ic} | j \sim_k i] \geq 1/2 + \vartheta$ .

Requirement (i) is a weak *coverage* requirement, which “encourages”  $\text{WIC}$  to choose indexes in dense regions of  $\mathcal{S}$ . Before studying the boosting abilities of UNN, we focus again on surrogate risks. So far, the surrogate risk (3) has been evaluated with respect to a single class. In a multiclass multilabel setting, we may compute the *total* surrogate risk over all classes as:

$$\varepsilon_{\mathcal{S}}^{\psi}(\mathcal{H}) \doteq \frac{1}{C} \sum_{c=1}^C \varepsilon_{\mathcal{S}}^{\psi}(h_c, c) , \quad (15)$$

where  $\mathcal{H}$  is the set of all  $C$  classifiers  $h_1, h_2, \dots, h_C$  that have been trained to minimize each  $\varepsilon_S^\psi(\cdot, c), c = 1, 2, \dots, C$ . We split classifiers just for convenience in the analysis: if one trains a single classifier  $H : \mathcal{O} \times \{1, 2, \dots, C\} \rightarrow \mathbb{R}$  like for example [3], then we define  $h_c$  to be  $H$  in which the second input coordinate is fixed to be  $c$ . Minimizing the total surrogate risk is not only efficient to estimate posteriors (Section 3): it is also useful to reduce the error in label prediction, as the total surrogate risk is an upperbound for the *Hamming risk* [3]:  $\varepsilon_S^H(\mathcal{H}) \doteq (1/(mC)) \sum_{c=1}^C \sum_{i=1}^m \mathbb{I}[y_{ic} h_c(\mathbf{x}_i) < 0]$ , where  $\mathbb{I}[\cdot]$  denotes the indicator variable. It is indeed not hard to check that for any strictly convex surrogate loss  $\psi$ , we have  $\varepsilon_S^H(\mathcal{H}) \leq (1/\psi(0)) \times \varepsilon_S^\psi(\mathcal{H})$ . We are left with the following question about UNN:

“are there sufficient conditions on the surrogate loss  $\psi$  that guarantee, under the sole **WLA**, a *convergence rate* towards the optimum of (15) with UNN?”

We give a positive answer to this question when the surrogate loss meets the following smoothness requirement.

**Definition 3.** [18]  $\psi$  is said to be  $\omega$  strongly smooth iff there exists some  $\omega > 0$  such that, for all  $x, x' \in \text{int}(\text{dom}(\psi))$ ,  $D_\psi(x' \| x) \leq \frac{\omega}{2}(x' - x)^2$ , where

$$D_\psi(x' \| x) \doteq \psi(x') - \psi(x) - (x' - x)\nabla_\psi(x) \quad (16)$$

denotes the Bregman divergence with generator  $\psi$  [2].

Denote  $n_j \doteq |\{i : j \sim_k i\}|$  the number of examples in  $\mathcal{S}$  of which  $(\mathbf{x}_j, \mathbf{y}_j)$  is a nearest neighbor, and  $n_* \doteq \max_j n_j$ . Denote also  $\mathcal{H}_{\text{opt}}$  the leveraged  $k$ -NN which minimizes  $\varepsilon_S^\psi(\mathcal{H})$ ; it corresponds to the set of classifiers  $\hat{h}_{\text{opt}}$  of Section 3 that would minimize (3) over each class. We are now ready to state our main result (remark that  $\varepsilon_S^\psi(\mathcal{H}_{\text{opt}}) \leq \psi(0)$ ).

**Theorem 3.** Suppose (**WLA**) holds and choose as  $\psi$  is any  $\omega$  strongly smooth, strictly convex loss. Then for any fixed  $\tau \in [\varepsilon_S^\psi(\mathcal{H}_{\text{opt}}), \psi(0)]$ , UNN has fit a leveraged  $k$ -NN classifier  $\mathcal{H}$  satisfying  $\varepsilon_S^\psi(\mathcal{H}) \leq \tau$  provided the number of boosting iterations  $T$  in the inner loop satisfies:

$$T \geq \frac{(\psi(0) - \tau)\omega m n_*}{2\vartheta^2 \rho^2}. \quad (17)$$

**Proof sketch:** To fit UNN to the notations of (15), we let  $h_c$  represent the leveraged  $k$ -NN in which each  $\alpha_j$  is restricted to  $\alpha_{jc}$ . We first analyze  $\varepsilon_S^\psi(h_c, c)$  for some fixed  $c$  in the outer loop of Algorithm 1, after all  $\alpha_{jc}$  have been computed in the inner loop. We adopt the following notations in this proof: we plug in the weight notation the iteration  $t$  and class  $c$ , so that  $w_{ii}^{(c)}$  denotes the weight of example  $\mathbf{x}_i$  at the beginning of the “for  $c$ ” loop of Algorithm 1.

$\psi$  is  $\omega$  strongly smooth is equivalent to  $\tilde{\psi}$  being strongly convex with parameter  $\omega^{-1}$  [18], that is,

$$\tilde{\psi}(w) - \frac{1}{2\omega}w^2 \text{ is convex,} \quad (18)$$



where we use notation  $\tilde{\psi}(x) \doteq \psi^*(-x)$ . Any convex function  $h$  satisfies  $h(w') \geq h(w) + \nabla_h(w)(w' - w)$ . We apply this inequality taking as  $h$  the function in (18). We obtain,  $\forall t = 1, 2, \dots, T, \forall i = 1, 2, \dots, m, \forall c = 1, 2, \dots, C$ :

$$D_{\tilde{\psi}}\left(w_{(t+1)i}^{(c)} \| w_{ti}^{(c)}\right) \geq \frac{1}{2\omega} \left(w_{(t+1)i}^{(c)} - w_{ti}^{(c)}\right)^2. \quad (19)$$

On the other hand, Cauchy-Schwartz inequality and (12) yield:

$$\forall j \in \mathcal{S}, \sum_{i:j \sim_k i} \left(r_{ij}^{(c)}\right)^2 \sum_{i:j \sim_k i} \left(w_{(t+1)i}^{(c)} - w_{ti}^{(c)}\right)^2 \geq \left(\sum_{i:j \sim_k i} r_{ij}^{(c)} w_{ti}^{(c)}\right)^2. \quad (20)$$

**Lemma 1.** *Under the WLA, index  $j$  returned by WIC at iteration  $t$  satisfies  $\left|\sum_{i:j \sim_k i} w_{ti}^{(c)} r_{ij}^{(c)}\right| \geq 2\vartheta \varrho$ .*

(proof omitted) Letting  $e(t) \in \{1, 2, \dots, m\}$  denote the index of the example returned at iteration  $t$  by WIC in Algorithm 1, and using (19), (20), Lemma 1 in this order, we arrive after few derivations (skipped because of the lack of space):

$$\frac{1}{m} \sum_{i=1}^m D_{\tilde{\psi}}\left(w_{(t+1)i}^{(c)} \| w_{ti}^{(c)}\right) \geq \frac{1}{2\omega m} \sum_{i:e(t) \sim_k i} \left(w_{(t+1)i}^{(c)} - w_{ti}^{(c)}\right)^2 \geq \frac{2\vartheta^2 \varrho^2}{\omega m n_{e(t)}} \geq \frac{2\vartheta^2 \varrho^2}{\omega m n_*}.$$

Summing these inequalities for  $t = 1, 2, \dots, T$  yields:

$$\sum_{t=1}^T \frac{1}{m} \sum_{i=1}^m D_{\tilde{\psi}}\left(w_{(t+1)i}^{(c)} \| w_{ti}^{(c)}\right) \geq \frac{2T\vartheta^2 \varrho^2}{\omega m n_*}. \quad (21)$$

Now, UNN meets the following property ([8], A.2):

$$\varepsilon_{\mathcal{S}}^{\psi}(h_{(t+1)c}, c) - \varepsilon_{\mathcal{S}}^{\psi}(h_{tc}, c) = -\frac{1}{m} \sum_{i=1}^m D_{\tilde{\psi}}\left(w_{(t+1)i}^{(c)} \| w_{ti}^{(c)}\right), \quad (22)$$

where  $h_{(t+1)c}$  denotes  $h_c$  after the  $t^{\text{th}}$  iteration in the inner loop of Algorithm 1. We unravel (22), using the fact that all  $\alpha$  are initialized to the null vector, and obtain that at the end of the inner loop,  $h_c$  satisfies:

$$\varepsilon_{\mathcal{S}}^{\psi}(h_c, c) = \psi(0) - \sum_{t=1}^T \frac{1}{m} \sum_{i=1}^m D_{\tilde{\psi}}\left(w_{(t+1)i}^{(c)} \| w_{ti}^{(c)}\right) \leq \psi(0) - \frac{2T\vartheta^2 \varrho^2}{\omega m n_*}, \quad (23)$$

from (21). There remains to compute the minimal value of  $T$  for which the right hand side of (23) becomes no greater than some user-fixed  $\tau \in [0, 1]$  to obtain that  $\varepsilon_{\mathcal{S}}^{\psi}(h_c, c) \leq \tau$ . The aggregation of the bounds for each  $c = 1, 2, \dots, C$  in  $\varepsilon_{\mathcal{S}}^{\psi}(\mathcal{H})$  is immediate as it is an average of  $\varepsilon_{\mathcal{S}}^{\psi}(h_c, c)$  over all classes. Hence, this minimal value of  $T$ , used for each  $c = 1, 2, \dots, C$ , also yields  $\varepsilon_{\mathcal{S}}^{\psi}(\mathcal{H}) \leq \tau$ . This ends the proof of Theorem 3.  $\square$

Section 3 has underlined the importance of balanced convex losses in obtaining simple efficient estimators for conditional class probabilities. Coupled with Theorem 3, we now show that UNN may be a fast approach to obtain such estimators.

|   | $\delta_{jc}$ , see (26)  | $g : w_i \leftarrow g(w_i)$  |
|---|---|--|
| A | $\frac{2W_{jc}-1}{2W_{jc}}$   | $w_i - 2\delta_{jc}y_{ic}y_{jc}$   |
| B | $\ln \frac{W_{jc}}{1-W_{jc}}$   | $\frac{w_i}{w_i \ln 2 + (1-w_i) \ln 2} \times \exp(\delta_{jc}y_{ic}y_{jc})$   |
| C | $\log_2 \frac{W_{jc}}{1-W_{jc}}$  | $\frac{w_i}{w_i + (1-w_i) \times 2^{\delta_{jc}y_{ic}y_{jc}}}$   |
| D | $\frac{2W_{jc}-1}{2\sqrt{W_{jc}(1-W_{jc})}}$  | $1 - \frac{1-w_i + \sqrt{w_i(2-w_i)}\delta_{jc}y_{ic}y_{jc}}{\sqrt{1+\delta_{jc}^2 w_i(2-w_i)+2(1-w_i)\sqrt{w_i(2-w_i)}\delta_{jc}y_{ic}y_{jc}}}$                        |
| E | N/A   | N/A  |
| F | $\frac{1}{2} \ln \frac{W_{jc}}{1-W_{jc}}$   | $\exp(-\delta_{jc}y_{ic}y_{jc})$   |
| G | $\frac{4}{1-\alpha^2} \left( \frac{(W_{jc})^{\frac{2}{1-\alpha}} - (1-W_{jc})^{\frac{2}{1-\alpha}}}{(W_{jc})^{\frac{2}{1-\alpha}} + (1-W_{jc})^{\frac{2}{1-\alpha}}} \right)$ | $\frac{4}{1-\alpha^2} \times \left( \frac{1-\alpha^2}{4} \delta_{jc}y_{ic}y_{jc} + \left( \frac{1+\alpha}{2\sqrt{w_i}} \right)^{1-\alpha} \right)^{-\frac{2}{1-\alpha}}$ |

Table 2: Computation of  $\delta_{jc}$  and the weight update rule of our implementation of UNN, for the strictly convex losses of Table 1. UNN leverages example  $j$  for class  $c$ , and the weight update is that of example  $i$  (See text for details and notations).

**Corollary 1.** *Consider any permissible  $\phi$  that has been scaled without loss of generality so that  $\phi(1/2) = 1$ ,  $\phi(0) = \phi(1) = 0$ . Then for the corresponding balanced convex loss  $\psi = \psi_\phi$  and under the **WLA**, picking*

$$T > \frac{mn_*}{2\vartheta^2 \varrho^2 \min_{x \in (0,1)} \left| \frac{\partial^2 \phi}{\partial x^2} \right|} \quad (24)$$

in the inner loop of UNN, for each  $c = 1, 2, \dots, C$ , guarantees to yield an optimal leveraged  $k$ -NN  $\mathcal{H}$ , satisfying  $\varepsilon_S^\psi(\mathcal{H}) = \varepsilon_S^\psi(\mathcal{H}_{\text{opt}})$ . This leveraged  $k$ -NN yields efficient estimators for conditional class probabilities, for each class, by computing:

$$\hat{p}_c(\mathbf{x}) = \nabla_\phi^{-1}(h_c(\mathbf{x})) . \quad (25)$$

(Proof omitted) For the most popular permissible functions (Table 1), quantity  $\min_{x \in (0,1)} \left| \frac{\partial^2 \phi}{\partial x^2} \right|$  does not take too small value: its values are respectively 8,  $4/\ln 2$ , 4 for the permissible functions corresponding to the squared loss, logistic loss, Matsushita loss. Hence, in these cases, the bound for  $T$  in (24) is not significantly affected by this term.

## 5 Experiments

### 5.1 Computing leveraging coefficients and weights update

Fix for short  $\mathcal{S}_{jb}^{(c)} \doteq \{i : j \sim_k i \wedge y_{ic} = by_{jc}\}$  for  $b \in \{+, -\}$ . (12) may be simplified as  $\sum_{i \in \mathcal{S}_{j+}^{(c)}} \nabla_\psi \left( \delta + \nabla_\psi^{-1}(-w_i) \right) = \sum_{i \in \mathcal{S}_{j-}^{(c)}} \nabla_\psi \left( -\delta + \nabla_\psi^{-1}(-w_i) \right)$ . There is no closed form solution to this equation in the general case. While it can be simply approximated with dichotomic search, it buys significant computation time, as this approximation has to be performed for each couple  $(c, t)$ . We

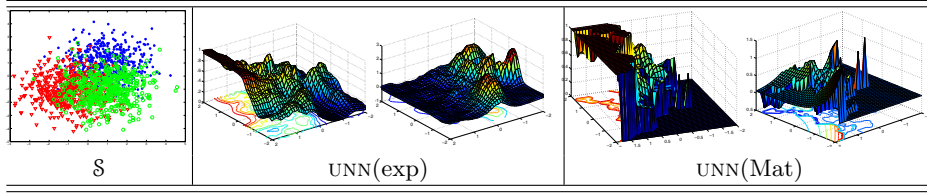


Fig. 1: From left to right: example of simulated dataset with  $\sigma = 1.1$ ; the estimated posterior for class 1 obtained by UNN(exp); the corresponding gridwise KL divergence for class 1; the estimated posterior for class 1 obtained by UNN(Mat); the corresponding gridwise KL divergence for class 1 (see (28) and text for details).

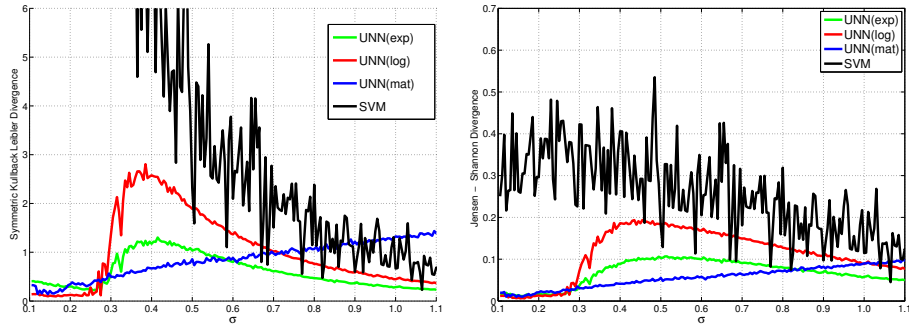


Fig. 2: Average Symmetric KL-divergence (left) and JensenShannon divergence (right) as a function of  $\sigma$  on simulated datasets, for UNN(exp), UNN(log), UNN(Mat) (left,  $k = 10$ ) and SVM .

tested a much faster alternative which produces results that are in general experimentally quite competitive, consisting in solving instead:  $\sum_{i \in \mathcal{S}_{j+}^{(c)}} w_i \nabla_{\psi}(\delta) = \sum_{i \in \mathcal{S}_{j-}^{(c)}} w_i \nabla_{\psi}(-\delta)$ . We get equivalently that  $\delta$  satisfies:

$$\frac{\nabla_{\psi}(-\delta)}{\nabla_{\psi}(\delta)} = \frac{W_{jc}}{1 - W_{jc}}, \quad (26)$$

with  $W_{jc} \doteq (\sum_{i \in \mathcal{S}_{j+}^{(c)}} w_i) / (\sum_{i \in \mathcal{S}_{j+}^{(c)}} w_i + \sum_{i \in \mathcal{S}_{j-}^{(c)}} w_i)$ . Remark the similarity with (8). Table 2 gives the corresponding expressions for  $\delta$  and the weight updates.

## 5.2 General experimental settings

We have tested three flavors of UNN: with the exponential loss (F in Table 1), the logistic loss (B in Table 1) and Matsushita's loss (D in Table 1). All three are respectively referred to as UNN(exp), UNN(log) and UNN(Mat). It is the first time this last flavor is tested, even from the classification standpoint. We chose support vector machines (SVM) as the contender against which to compare UNN:

SVM are large margin classifiers with convenient methods to obtain estimators for the posteriors [19]. For all these algorithms, we compute the estimation of posteriors as follows: we use (11) for UNN(exp), (25) for UNN(log) and UNN(Mat). For SVM, we use the method of [19], which, given a SVM  $f$  for class  $c$ , forms the posterior:

$$\hat{p}_c(\mathbf{x}) \doteq \frac{1}{1 + \exp(af(\mathbf{x}) + b)} , \quad (27)$$

where  $a$  and  $b$  are estimated by maximizing the log-likelihood of the training sample with a five-fold cross validation. On synthetic datasets SVM performs equally using both linear and radial basis function kernel. Therefore, in the following we indicate with *SVM* the linear Support Vector Machine. We use three metrics to evaluate the algorithms: Two computed to evaluate performance on synthetic data and one to evaluate performance on real datasets. On simulated data, we compute first Kullback-Leibler (KL) divergences between the true and estimated posterior and after their mean obtaining the Symmetric Kullback-Leibler divergences:

$$D_{\text{KL}}(\hat{p}||p) \doteq \sum_c \Pr[c] \int \Pr[\mathbf{x}] \hat{p}_c(\mathbf{x}) \ln \frac{\hat{p}_c(\mathbf{x})}{p_c(\mathbf{x})} d\mu , \quad D_{\text{KL}}(p||\hat{p}) \doteq \sum_c \Pr[c] \int \Pr[\mathbf{x}] p_c(\mathbf{x}) \ln \frac{p_c(\mathbf{x})}{\hat{p}_c(\mathbf{x})} d\mu \quad (28)$$

$$\text{Symm}D_{\text{KL}} \doteq \frac{1}{2}(D_{\text{KL}}(\hat{p}||p) + D_{\text{KL}}(p||\hat{p})) \quad (29)$$

and also we compute JensenShannon (JS) divergence:

$$D_{\text{JS}} \doteq \frac{1}{2}(D_{\text{KL}}(\hat{p}||q) + D_{\text{KL}}(p||q)) \quad (30)$$

where  $q$  is the average of the two distribution. Our estimate,  $\text{Symm}\hat{D}_{\text{KL}}$  and  $\hat{D}_{\text{JS}}$  rely on a simple fine-grained grid approximation of the integral over the subsets of  $\mathcal{O}$  of sufficient mass according to  $\mu$ . On real data, we compute a couple of metrics. First, we compute the F-measure of the classifiers (the harmonic average of precision and recall), based on thresholding the probabilistic output and deciding that  $\mathbf{x}$  belong to class  $c$  iff  $\hat{p}_c(\mathbf{x}) \geq \kappa$ , for varying  $\kappa \in (1/2, 1)$ . Second, we compute the rejection rate, that is, the proportion of observations for which  $\hat{p}_c(\mathbf{x}) < \kappa$ . Either we plot couples of curves for the F-measure and rejection rates, or we summarize both metrics by their average values as  $\kappa$  ranges through  $(1/2, 1)$ , which amounts to compute the area under the corresponding curves.

### 5.3 Results on simulated data

We evaluated the goodness-of-fit of the estimates on simulated datasets with the following experiments. We crafted a general domain consisting of  $C = 3$  equiprobable classes, each of which follows a Gaussian  $\mathcal{N}(\boldsymbol{\mu}, \sigma\mathbf{I})$ , for  $\sigma \in [0.1, 1.1]$  with steps of 0.005, and  $\boldsymbol{\mu}$  remains the same. For each value of  $\sigma$ , we compute

|                    | $k$ | UNN(exp) | UNN(log) | UNN(Mat) | SVM   |
|--------------------|-----|----------|----------|----------|-------|
| $Symm\hat{D}_{KL}$ | 10  | 0.599    | 1.029    | 0.848    | 3.533 |
|                    | 20  | 0.372    | 0.760    | 0.687    |       |
|                    | 30  | 0.293    | 0.610    | 0.646    |       |
|                    | 40  | 0.254    | 0.534    | 0.632    |       |
| $\hat{D}_{JS}$     | 10  | 0.067    | 0.113    | 0.0562   | 0.256 |
|                    | 20  | 0.045    | 0.086    | 0.045    |       |
|                    | 30  | 0.036    | 0.072    | 0.043    |       |
|                    | 40  | 0.032    | 0.065    | 0.043    |       |
| F-measure          | 10  | 90.32    | 89.59    | 90.58    | 91.02 |
|                    | 20  | 90.62    | 89.53    | 90.81    |       |
|                    | 30  | 90.70    | 89.26    | 90.84    |       |
|                    | 40  | 90.72    | 88.82    | 90.88    |       |

Table 3: Average results over simulated data, for UNN(exp), UNN(log), UNN(Mat) with four different values of  $k$ , and for support vector machines with linear (SVM ).

the average over ten simulations, each of which consists of 1500 training examples and 4500 testing examples. We get overall several thousands datasets, on which all algorithms are tested. Figure 1 presents an example of such datasets, along with results obtained by UNN(exp) and UNN(Mat) from the standpoints of the posterior estimates and KL-divergence on the same class. The estimators are rather good, with the largest mismatches (KL-divergence) located near the frontiers of classes. Also, UNN(Mat) tends to outperform UNN(exp).

Figure 2 synthesizes the results from the KL and JS divergence standpoints. Two clear conclusions can be drawn from these results. First, UNN is the clear winner over SVM for the posteriors estimation task. The results of each flavor of UNN is indeed better than those of SVM by orders of magnitude. This is all the more important as the kernels we used are the theoretical kernels of choice given the way we have simulated data. The second conclusion is that UNN(Mat) is the best of all flavors of UNN, a fact also confirmed by the synthetic results of Table 3. Its behavior (Figure 2) is also monotonous: it is predictable that it increases with the degree of overlap between classes, that is, with  $\sigma$ . From the *classification* standpoint, the average F-measure metrics display a very slight advantage to SVM.

The most important conclusion that can be drawn from the simulated data is shown in Figure 3: as the number of boosting iterations  $T$  increase, UNN does *not* overfit posteriors in general. The only hitch — not statistically significant — is the case  $\sigma > 0.7$  for UNN(Mat), but the differences are of very small order compared to the standard deviations of the KL-divergence.

#### 5.4 Results on the SUN database domains

We have crafted, out of the challenging SUN computer vision database [20], three datasets, consisting in taking all pictures from the first ten (SUN 10), twenty (SUN 20) or thirty (SUN 30) classes. We have compared UNN(exp), UNN(log), UNN(Mat) and SVM on each dataset, by computing the average values, over the threshold  $\kappa$ , of the F-measure and the rejection rate. Table 4 summarizes

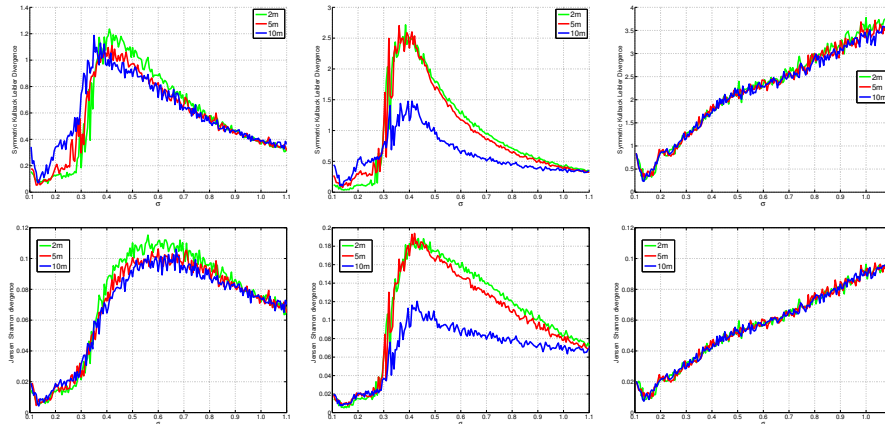


Fig. 3: Average symmetric KL-divergence (top) and JensenShannon divergence (bottom) as a function of  $\sigma$  on simulated datasets, for UNN(exp) (left), UNN(log) (center), UNN(Mat) (right), when the number of boosting iterations  $T$  varies in  $\{2m, 5m, 10m\}$ . The color code in the same on each plot.

|        | UNN(exp)     |       | UNN(log) |       | UNN(Mat) |             | SVM <sub>l</sub> |       |
|--------|--------------|-------|----------|-------|----------|-------------|------------------|-------|
|        | F            | R     | F        | R     | F        | R           | F                | R     |
| SUN 10 | <b>89.91</b> | 21.35 | 84.46    | 5.18  | 72.47    | <b>3.39</b> | 87.99            | 22.32 |
| SUN 20 | <b>82.82</b> | 36.64 | 72.34    | 8.51  | 55.46    | <b>2.51</b> | 74.60            | 33.25 |
| SUN 30 | <b>73.39</b> | 49.92 | 61.02    | 14.99 | 40.83    | <b>5.99</b> | 62.81            | 39.95 |

Table 4: Area under the (F)-measure (in percentage) and (R)ejection rate on the SUN databases. For each database, the best F and R are written in **bold faces**.

the results obtained. This table somehow confirms that classification and posterior estimation may be conflicting goals when it comes to boosting [4, 5], as UNN(Mat) achieves very poor results compared to the other algorithms. Furthermore, UNN(exp) appears to be the clear winner over all algorithms for this classification task. These results have to be appreciated in the light of the rejection rates: in comparison with the other algorithms, UNN(Mat) rejects a very small proportion of the examples, this indicating a high recall for the algorithm. Figure 4 completes the picture by detailing F-measure and rejection rates plots. The F-measure plots clearly display the better performances of UNN(exp) compared to the other algorithms, and the fact that UNN(Mat) displays very stable performances. The rejection rates plots show that UNN(Mat) indeed rejects a very small proportion of examples, even for large values of  $\kappa$ .

## 6 Conclusion

Boosting algorithms are remarkably simple and efficient from the classification standpoint, and are being used in a rapidly increasing number of domains and

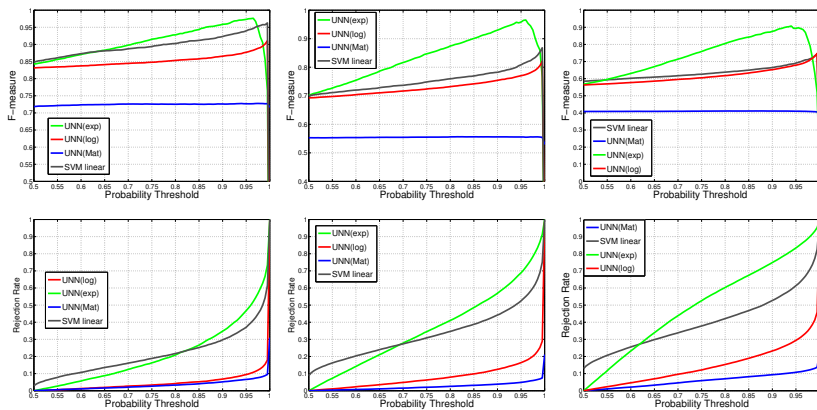


Fig. 4: F-measure (top row) and rejection rates (bottom row) on the SUN domains, with  $C = 10$  (left),  $C = 20$  (center) and  $C = 30$  (right, see Table 3 for notations).

problems [6]. In some sense, it would be too bad that such successes be impeded when it comes to posterior estimation [5]. Experimental results display that this estimation is possible, but it necessitates a very fine tuning of the algorithms [6]. The point of our paper is that estimating class conditional probabilities may be possible, without such tedious tunings, and sometimes even *without overfitting*, if we boost topological approaches to learning like nearest neighbors [8]; such approaches offer very promising results in the field of computer vision [7]. There is a simple explanation to this fact. For any classifier, the conditional class probability estimation for some  $\mathbf{x}$  in (7) is the same as for any other observation in the vicinity of  $\mathbf{x}$ , where the “vicinity” is to be understood from the *classifier* standpoint. When boosting decision trees, the vicinity of  $\mathbf{x}$  corresponds to observations classified by the same leaf as  $\mathbf{x}$ . As the number of leaves of the tree increases, the vicinity gets narrowed, which weakens the estimation in (7) and thus overfits the corresponding estimated density. Ultimately, linear combinations of such trees, such as those performed in AdaBoost, make such a fine-grained approximation of the local topology of data that the estimators get irreparably confined to the borders of the interval  $[0, 1]$  [5]. Nearest neighbors do not have such a drawback, as the set of  $k$ -nearest neighbors in  $\mathcal{S}$  of some observation  $\mathbf{x}$  spans a region of  $\mathcal{O}$  which does not change throughout the iterations. Furthermore, nearest neighbor rules exploit a topology of data which, under regularity conditions about the true posteriors, also carries out information about these posteriors. For these reasons, nearest neighbors might be a key entry for a reliable estimation of posteriors with boosting.

## 7 Acknowledgments

R. Nock acknowledges a visiting grant from Institut Universitaire de France / Université de Nice.

## References

1. Schapire, R.E., Freund, Y., Bartlett, P., Lee, W.S.: Boosting the margin : a new explanation for the effectiveness of voting methods. *Annals of statistics* **26** (1998) 1651–1686
2. Nock, R., Nielsen, F.: On the efficient minimization of classification-calibrated surrogates. In: *NIPS\*21*. (2008) 1201–1208
3. Schapire, R.E., Singer, Y.: Improved boosting algorithms using confidence-rated predictions. *Machine Learning Journal* **37** (1999) 297–336
4. Friedman, J., Hastie, T., Tibshirani, R.: Additive Logistic Regression : a Statistical View of Boosting. *Annals of Statistics* **28** (2000) 337–374
5. Buja, A., Mease, D., Wyner, A.J.: Comment: Boosting algorithms: regularization, prediction and model fitting. *Statistical Science* **22** (2007) 506–512
6. Bühlmann, P., Hothorn, T.: Boosting algorithms: regularization, prediction and model fitting. *Statistical Science* **22** (2007) 477–505
7. Nock, R., Piro, P., Nielsen, F., Bel Haj Ali, W., Barlaud, M.: Boosting  $k$ -NN for categorization of natural scenes. *International Journal of Computer Vision* (2012, to appear)
8. Piro, P., Nock, R., Nielsen, F., Barlaud, M.: Leveraging  $k$ -NN for generic classification boosting. *Neurocomputing* **80** (2012) 3–9
9. Bartlett, P., Jordan, M., McAuliffe, J.D.: Convexity, classification, and risk bounds. *Journal of the Am. Stat. Assoc.* **101** (2006) 138–156
10. Nock, R., Nielsen, F.: Bregman divergences and surrogates for learning. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **31**(11) (2009) 2048–2059
11. Kearns, M., Mansour, Y.: On the boosting ability of top-down decision tree learning algorithms. *Journal of Comp. Syst. Sci.* **58** (1999) 109–128
12. Amari, S.I., Nagaoka, H.: *Methods of Information Geometry*. Oxford University Press (2000)
13. Müller-Funk, U., Pukelsheim, F., Witting, H.: On the attainment of the Cramér-Rao bound in  $L_r$ -differentiable families of distributions. *Annals of Statistics* (1989) 1742–1748
14. Cover, T.M., Hart, P.: Nearest neighbor pattern classification. *IEEE Trans. on Information Theory* **13** (1967) 21–27
15. Devroye, L., Györfi, L., Lugosi, G.: *A Probabilistic Theory of Pattern Recognition*. Springer (1996)
16. Sebban, M., Nock, R., Lallich, S.: Boosting Neighborhood-Based Classifiers. In: *Proc. of the 18<sup>th</sup> International Conference on Machine Learning*, Morgan Kaufman (2001) 505–512
17. Sebban, M., Nock, R., Lallich, S.: Stopping criterion for boosting-based data reduction techniques: from binary to multiclass problems. *J. of Mach. Learn. Res.* **3** (2003) 863–885
18. Kakade, S., Shalev-Shwartz, S., Tewari, A.: Applications of strong convexity–strong smoothness duality to learning with matrices. Technical Report CoRR abs/0910.0610, Computing Res. Repository (2009)
19. Platt, J.C.: Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In: *Advances in Large Margin Classifiers*, MIT Press (1999) 61–74
20. Xiao, J., Hays, J., Ehringer, K.A., Oliva, A., Torralba, A.: SUN database: Large-scale scene recognition from abbey to zoo. In: *Proc. of IEEE International Conference on Computer Vision and Pattern Recognition*. (2010) 3485–3492