Rates of convergence for nearest neighbor classification

Kamalika Chaudhuri
Computer Science and Engineering
University of California, San Diego
kamalika@cs.ucsd.edu

Sanjoy Dasgupta
Computer Science and Engineering
University of California, San Diego
dasgupta@cs.ucsd.edu

Abstract

We analyze the behavior of nearest neighbor classification in metric spaces and provide finite-sample, distribution-dependent rates of convergence under minimal assumptions. These are more general than existing bounds, and enable us, as a by-product, to establish the universal consistency of nearest neighbor in a broader range of data spaces than was previously known. We illustrate our upper and lower bounds by introducing a new smoothness class customized for nearest neighbor classification. We find, for instance, that under the Tsybakov margin condition the convergence rate of nearest neighbor matches recently established lower bounds for nonparametric classification.

1   Introduction

In this paper, we deal with binary prediction in metric spaces. A classification problem is defined by a metric space $(\mathcal{X}, \rho)$ from which instances are drawn, a space of possible labels $\mathcal{Y} = \{0, 1\}$, and a distribution $\mathbb{P}$ over $\mathcal{X} \times \mathcal{Y}$. The goal is to find a function $h : \mathcal{X} \to \mathcal{Y}$ that minimizes the probability of error on pairs (X, Y) drawn from $\mathbb{P}$; this error rate is the *risk* $R(h) = \mathbb{P}(h(X) \neq Y)$. The best such function is easy to specify: if we let μ denote the marginal distribution of X and η the conditional probability $\eta(x) = \mathbb{P}(Y = 1 | X = x)$, then the predictor $1(\eta(x) \geq 1/2)$ achieves the minimum possible risk, $R^* = \mathbb{E}_X[\min(\eta(X), 1 - \eta(X))]$. The trouble is that $\mathbb{P}$ is unknown and thus a prediction rule must instead be based only on a finite sample of points $(X_1, Y_1), \ldots, (X_n, Y_n)$ drawn independently at random from $\mathbb{P}$.

Nearest neighbor (NN) classifiers are among the simplest prediction rules. The *1-NN classifier* assigns each point $x \in \mathcal{X}$ the label Y_i of the closest point in $X_1, \ldots, X_n$ (breaking ties arbitrarily, say). For a positive integer k, the *k-NN classifier* assigns x the majority label of the k closest points in $X_1, \ldots, X_n$. In the latter case, it is common to let k grow with n, in which case the sequence $(k_n : n \geq 1)$ defines a k_n-*NN classifier*.

The asymptotic consistency of nearest neighbor classification has been studied in detail, starting with the work of Fix and Hodges [7]. The risk of the NN classifier, henceforth denoted R_n, is a random variable that depends on the data set $(X_1, Y_1), \ldots, (X_n, Y_n)$; the usual order of business is to first determine the limiting behavior of the expected value $\mathbb{E}R_n$ and to then study stronger modes of convergence of R_n. Cover and Hart [2] studied the asymptotics of $\mathbb{E}R_n$ in general metric spaces, under the assumption that every x in the support of μ is either a continuity point of η or has $\mu(\{x\}) > 0$. For the 1-NN classifier, they found that $\mathbb{E}R_n \to \mathbb{E}_X[2\eta(X)(1 - \eta(X))] \leq 2R^*(1 - R^*)$; for k_n-NN with $k_n \uparrow \infty$ and $k_n/n \downarrow 0$, they found $\mathbb{E}R_n \to R^*$. For points in Euclidean space, a series of results starting with Stone [15] established consistency without any distributional assumptions. For k_n-NN in particular, $R_n \to R^*$ almost surely [5].

These consistency results place nearest neighbor methods in a favored category of nonparametric estimators. But for a fuller understanding it is important to also have rates of convergence. For

1

instance, part of the beauty of nearest neighbor is that it appears to adapt automatically to different distance scales in different regions of space. It would be helpful to have bounds that encapsulate this property.

Rates of convergence are also important in extending nearest neighbor classification to settings such as active learning, semisupervised learning, and domain adaptation, in which the training data is not a fully-labeled data set obtained by i.i.d. sampling from the future test distribution. For instance, in active learning, the starting point is a set of unlabeled points $X_1, \ldots, X_n$, and the learner requests the labels of just a few of these, chosen adaptively to be as informative as possible about η. There are many natural schemes for deciding which points to label: for instance, one could repeatedly pick the point furthest away from the labeled points so far, or one could pick the point whose k nearest labeled neighbors have the largest disagreement among their labels. The asymptotics of such selective sampling schemes have been considered in earlier work [4], but ultimately the choice of scheme must depend upon finite-sample behavior. The starting point for understanding this behavior is to first obtain a characterization in the non-active setting.

1.1 Previous work on rates of convergence

There is a large body of work on convergence rates of nearest neighbor estimators. Here we outline some of the types of results that have been obtained, and give representative sources for each.

The earliest rates of convergence for nearest neighbor were *distribution-free*. Cover [3] studied the 1-NN classifier in the case $\mathcal{X} = \mathbb{R}$, under the assumption of class-conditional densities with uniformly-bounded third derivatives. He showed that $\mathbb{E}R_n$ converges at a rate of $O(1/n^2)$. Wagner [18] and later Fritz [8] also looked at 1-NN, but in higher dimension $\mathcal{X} = \mathbb{R}^d$. The latter obtained an asymptotic rate of convergence for R_n under the milder assumption of non-atomic μ and lower semi-continuous class-conditional densities.

Distribution-free results are valuable, but do not characterize which properties of a distribution most influence the performance of nearest neighbor classification. More recent work has investigated different approaches to obtaining distribution-dependent bounds, in terms of the smoothness of the distribution.

A simple and popular smoothness parameter is the Holder constant. Kulkarni and Posner [12] obtained a fairly general result of this kind for 1-NN and k_n-NN. They assumed that for some constants K and α, and for all $x_1, x_2 \in \mathcal{X}$,

$$|\eta(x_1) - \eta(x_2)| \leq K\rho(x_1, x_2)^{2\alpha}.$$

They then gave bounds in terms of the Holder parameter α as well as covering numbers for the marginal distribution μ. Gyorfi [9] looked at the case $\mathcal{X} = \mathbb{R}^d$, under the weaker assumption that for some function $K : \mathbb{R}^d \to \mathbb{R}$ and some α, and for all $z \in \mathbb{R}^d$ and all $r > 0$,

$$\left| \eta(z) - \frac{1}{\mu(B(z,r))} \int_{B(z,r)} \eta(x)\mu(dx) \right| \leq K(z)r^\alpha.$$

The integral denotes the average η value in a ball of radius r centered at z; hence, this α is similar in spirit to the earlier Holder parameter, but does not require η to be continuous. Gyorfi obtained asymptotic rates in terms of α. Another generalization of standard smoothness conditions was proposed recently [17] in a "probabilistic Lipschitz" assumption, and in this setting rates were obtained for NN classification in bounded spaces $\mathcal{X} \subset \mathbb{R}^d$.

The literature leaves open several basic questions that have motivated the present paper. (1) Is it possible to give tight finite-sample bounds for NN classification in metric spaces, without any smoothness assumptions? What aspects of the distribution must be captured in such bounds? (2) Are there simple notions of smoothness that are especially well-suited to nearest neighbor? Roughly speaking, we consider a notion suitable if it is possible to sharply characterize the convergence rate of nearest neighbor for all distributions satisfying this notion. As we discuss further below, the Holder constant is lacking in this regard. (3) A recent trend in nonparametric classification has been to study rates of convergence under "margin conditions" such as that of Tsybakov. The best achievable rates under these conditions are now known: does nearest neighbor achieve these rates?

2
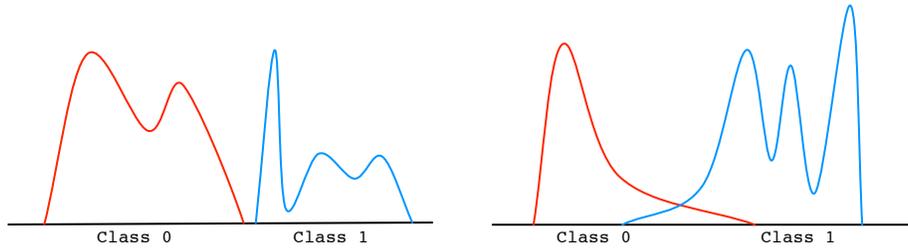
Figure 1: One-dimensional distributions. In each case, the class-conditional densities are shown.

1.2 Some illustrative examples

We now look at a couple of examples to get a sense of what properties of a distribution most critically affect the convergence rate of nearest neighbor. In each case, we study the k-NN classifier.

To start with, consider a distribution over $\mathcal{X} = \mathbb{R}$ in which the two classes ($Y = 0, 1$) have class-conditional densities μ_0 and μ_1. Assume that these two distributions have disjoint support, as on the left side of Figure 1. The k-NN classifier will make a mistake on a specific query x only if x is near the boundary between the two classes. To be precise, consider an interval around x of probability mass k/n, that is, an interval $B = [x-r, x+r]$ with $\mu(B) = k/n$. Then the k nearest neighbors will lie roughly in this interval, and there will likely be an error only if the interval contains a substantial portion of the wrong class. Whether or not η is smooth, or the μ_i are smooth, is irrelevant.

In a general metric space, the k nearest neighbors of any query point x are likely to lie in a ball centered at x of probability mass roughly k/n. Thus the central objects in analyzing k-NN are balls of mass $\approx k/n$ near the decision boundary, and it should be possible to give rates of convergence solely in terms of these.

Now let's turn to notions of smoothness. Figure 1, right, shows a variant of the previous example in which it is no longer the case that $\eta \in \{0, 1\}$. Although one of the class-conditional densities in the figure is highly non-smooth, this erratic behavior occurs far from the decision boundary and thus does not affect nearest neighbor performance. And in the vicinity of the boundary, what matters is not how much η varies within intervals of any given radius r, but rather within intervals of probability mass k/n. Smoothness notions such as Lipschitz and Holder constants, which measure changes in η with respect to x, are therefore not entirely suitable: what we need to measure are changes in η with respect to the underlying marginal μ on $\mathcal{X}$.

1.3 Results of this paper

Let us return to our earlier setting of pairs (X, Y), where X takes values in a metric space $(\mathcal{X}, \rho)$ and has distribution μ, while $Y \in \{0, 1\}$ has conditional probability function $\eta(x) = \Pr(Y = 1 | X = x)$. We obtain rates of convergence for k-NN by attempting to make precise the intuitions discussed above. This leads to a somewhat different style of analysis than has been used in earlier work.

Our main result is an upper bound on the misclassification rate of k-NN that holds for any sample size n and for any metric space, with no distributional assumptions. The bound depends on a novel notion of the *effective boundary* for k-NN: for the moment, denote this set by $A_{n,k} \subset \mathcal{X}$.

- We show that with high probability over the training data, the misclassification rate of the k-NN classifier (with respect to the Bayes-optimal classifer) is bounded above by $\mu(A_{n,k})$ plus a small additional term that can be made arbitrarily small (Theorem 5).

- We lower-bound the misclassification rate using a related notion of effective boundary (Theorem 6).

- We identify a general condition under which, as n and k grow, $A_{n,k}$ approaches the actual decision boundary $\{x \mid \eta(x) = 1/2\}$. This yields universal consistency in a wider range of metric spaces than just $\mathbb{R}^d$ (Theorem 1), thus broadening our understanding of the asymptotics of nearest neighbor.

3

We then specialize our generalization bounds to smooth distributions.

- We introduce a novel smoothness condition that is tailored to nearest neighbor. We compare our upper and lower bounds under this kind of smoothness (Theorem 3).
- We obtain risk bounds under the margin condition of Tsybakov that match the best known results for nonparametric classification (Theorem 4).
- We look at additional specific cases of interest: when η is bounded away from $1/2$, and the even more extreme scenario where $\eta \in \{0, 1\}$ (zero Bayes risk).

2 Definitions and results

Let $(\mathcal{X}, \rho)$ be any separable metric space. For any $x \in \mathcal{X}$, let
$$B^o(x, r) = \{x' \in \mathcal{X} \mid \rho(x, x') < r\} \text{ and } B(x, r) = \{x' \in \mathcal{X} \mid \rho(x, x') \leq r\}$$
denote the open and closed balls, respectively, of radius r centered at x.

Let μ be a Borel regular probability measure on this space (that is, open sets are measurable, and every set is contained in a Borel set of the same measure) from which *instances* X are drawn. The *label* of an instance $X = x$ is $Y \in \{0, 1\}$ and is distributed according to the measurable conditional probability function $\eta : \mathcal{X} \to [0, 1]$ as follows: $\Pr(Y = 1 | X = x) = \eta(x)$.

Given a data set $S = ((X_1, Y_1), \ldots, (X_n, Y_n))$ and a query point $x \in \mathcal{X}$, we use the notation $X^{(i)}(x)$ to denote the i-th nearest neighbor of x in the data set, and $Y^{(i)}(x)$ to denote its label. Distances are calculated with respect to the given metric ρ, and ties are broken by preferring points earlier in the sequence. The k-NN classifier is defined by

$$g_{n,k}(x) = \begin{cases} 1 & \text{if } Y^{(1)}(x) + \cdots + Y^{(k)}(x) \geq k/2 \\ 0 & \text{otherwise} \end{cases}$$

We analyze the performance of $g_{n,k}$ by comparing it with $g(x) = 1(\eta(x) \geq 1/2)$, the omniscient Bayes-optimal classifier. Specifically, we obtain bounds on $\Pr_X(g_{n,k}(X) \neq g(X))$ that hold with high probability over the choice of data S, for any n. It is worth noting that convergence results for nearest neighbor have traditionally studied the excess risk $R_{n,k} - R^*$, where $R_{n,k} = \Pr(Y \neq g_{n,k}(X))$. If we define the pointwise quantities

$$R_{n,k}(x) = \Pr(Y \neq g_{n,k}(x) | X = x)$$
$$R^*(x) = \min(\eta(x), 1 - \eta(x)),$$

for all $x \in \mathcal{X}$, we see that

$$R_{n,k}(x) - R^*(x) = |1 - 2\eta(x)| 1(g_{n,k}(x) \neq g(x)). \tag{1}$$

Taking expectation over X, we then have $R_{n,k} - R^* \leq \Pr_X(g_{n,k}(X) \neq g(X))$, and so we also obtain upper bounds on the excess risk.

The technical core of this paper is the finite-sample generalization bound of Theorem 5. We begin, however, by discussing some of its implications since these relate directly to common lines of inquiry in the statistical literature. All proofs appear in the appendix.

2.1 Universal consistency

A series of results, starting with [15], has shown that k_n-NN is strongly consistent ($R_n = R_{n,k_n} \to R^*$ almost surely) when $\mathcal{X}$ is a finite-dimensional Euclidean space and μ is a Borel measure. A consequence of the bounds we obtain in Theorem 5 is that this phenomenon holds quite a bit more generally. In fact, strong consistency holds in any metric measure space $(\mathcal{X}, \rho, \mu)$ for which the Lebesgue differentiation theorem is true: that is, spaces in which, for any bounded measurable f,

$$\lim_{r \downarrow 0} \frac{1}{\mu(B(x, r))} \int_{B(x,r)} f \, d\mu = f(x) \tag{2}$$

for almost all (μ-a.e.) $x \in \mathcal{X}$.

For more details on this differentiation property, see [6, 2.9.8] and [10, 1.13]. It holds, for instance:

4

- When $(\mathcal{X}, \rho)$ is a finite-dimensional normed space [10, 1.15(a)].
- When $(\mathcal{X}, \rho, \mu)$ is *doubling* [10, 1.8], that is, when there exists a constant $C(\mu)$ such that $\mu(B(x, 2r)) \leq C(\mu)\mu(B(x, r))$ for every ball $B(x, r)$.
- When μ is an atomic measure on $\mathcal{X}$.

For the following theorem, recall that the risk of the k_n-NN classifier, $R_n = R_{n,k_n}$, is a function of the data set $(X_1, Y_1), \ldots, (X_n, Y_n)$.

Theorem 1. *Suppose metric measure space $(\mathcal{X}, \rho, \mu)$ satisfies differentiation condition (2). Pick a sequence of positive integers (k_n), and for each n, let $R_n = R_{n,k_n}$ be the risk of the k_n-NN classifier g_{n,k_n}.*

1. *If $k_n \to \infty$ and $k_n/n \to 0$, then for all $\epsilon > 0$,*

$$\lim_{n \to \infty} \mathrm{Pr}_n(R_n - R^* > \epsilon) = 0.$$

   Here Pr_n denotes probability over the data set $(X_1, Y_1), \ldots, (X_n, Y_n)$.

2. *If in addition $k_n/(\log n) \to \infty$, then $R_n \to R^*$ almost surely.*

2.2 Smooth measures

Before stating our finite-sample bounds in full generality, we provide a glimpse of them under smooth probability distributions. We begin with a few definitions.

The support of μ. The support of distribution μ is defined as

$$\mathrm{supp}(\mu) = \{x \in \mathcal{X} \mid \mu(B(x, r)) > 0 \text{ for all } r > 0\}.$$

It was shown by [2] that in separable metric spaces, $\mu(\mathrm{supp}(\mu)) = 1$. For the interested reader, we reproduce their brief proof in the appendix (Lemma 24).

The conditional probability function for a set. The conditional probability function η is defined for points $x \in \mathcal{X}$, and can be extended to measurable sets $A \subset \mathcal{X}$ with $\mu(A) > 0$ as follows:

$$\eta(A) = \frac{1}{\mu(A)} \int_A \eta \, d\mu. \tag{3}$$

This is the probability that $Y = 1$ for a point X chosen at random from the distribution μ restricted to set A. We exclusively consider sets A of the form $B(x, r)$, in which case η is defined whenever $x \in \mathrm{supp}(\mu)$.

2.2.1 Smoothness with respect to the marginal distribution

For the purposes of nearest neighbor, it makes sense to define a notion of smoothness with respect to the marginal distribution on instances. For $\alpha, L > 0$, we say the conditional probability function η is (α, L)-*smooth* in metric measure space $(\mathcal{X}, \rho, \mu)$ if for all $x \in \mathrm{supp}(\mu)$ and all $r > 0$,

$$|\eta(B(x, r)) - \eta(x)| \leq L \, \mu(B^o(x, r))^\alpha.$$

(As might be expected, we only need to apply this condition locally, so it is enough to restrict attention to balls of probability mass upto some constant p_o.) One feature of this notion is that it is scale-invariant: multiplying all distances by a fixed amount leaves α and L unchanged. Likewise, if the distribution has several well-separated clusters, smoothness is unaffected by the distance-scales of the individual clusters.

It is common to analyze nonparametric classifiers under the assumption that $\mathcal{X} = \mathbb{R}^d$ and that η is α_H-*Holder continuous* for some $\alpha > 0$, that is,

$$|\eta(x) - \eta(x')| \leq L\|x - x'\|^{\alpha_H}$$

for some constant L. These bounds typically also require μ to have a density that is uniformly bounded (above and/or below). We now relate these standard assumptions to our notion of smoothness.

5

Lemma 2. *Suppose that $\mathcal{X} \subset \mathbb{R}^d$, and η is α_H-Holder continuous, and μ has a density with respect to Lebesgue measure that is $\geq \mu_{\min}$ on $\mathcal{X}$. Then there is a constant L such that for any $x \in supp(\mu)$ and $r > 0$ with $B(x,r) \subset \mathcal{X}$, we have $|\eta(x) - \eta(B(x,r))| \leq L\mu(B^o(x,r))^{\alpha_H/d}$.*

(To remove the requirement that $B(x,r) \subset \mathcal{X}$, we would need the boundary of $\mathcal{X}$ to be well-behaved, for instance by requiring that $\mathcal{X}$ contains a constant fraction of every ball centered in it. This is a familiar assumption in nonparametric classification, including the seminal work of [1] that we discuss shortly.)

Our smoothness condition for nearest neighbor problems can thus be seen as a generalization of the usual Holder conditions. It applies in broader range of settings, for example for discrete μ.

2.2.2 Generalization bounds for smooth measures

Under smoothness, our general finite-sample convergence rates (Theorems 5 and 6) take on an easily interpretable form. Recall that $g_{n,k}(x)$ is the k-NN classifier, while $g(x)$ is the Bayes-optimal prediction.

Theorem 3. *Suppose η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$. The following hold for any n and k.*

(Upper bound on misclassification rate.) Pick any $\delta > 0$ and suppose that $k \geq 16 \ln(2/\delta)$. Then

$$\Pr_X(g_{n,k}(X) \neq g(X)) \leq \delta + \mu\left(\left\{x \in \mathcal{X} \;\middle|\; |\eta(x) - \frac{1}{2}| \leq \sqrt{\frac{1}{k}\ln\frac{2}{\delta}} + L\left(\frac{k}{2n}\right)^\alpha\right\}\right).$$

(Lower bound on misclassification rate.) Conversely, there is an absolute constant c_o such that

$$\mathbb{E}_n \Pr_X(g_{n,k}(X) \neq g(X)) \geq c_o\mu\left(\left\{x \in \mathcal{X} \;\middle|\; \eta(x) \neq \frac{1}{2}, \; |\eta(x) - \frac{1}{2}| \leq \frac{1}{\sqrt{k}} - L\left(\frac{2k}{n}\right)^\alpha\right\}\right).$$

Here $\mathbb{E}_n$ is expectation over the data set.

The optimal choice of k is $\sim n^{2\alpha/(2\alpha+1)}$, and with this setting the upper and lower bounds are directly comparable: they are both of the form $\mu(\{x : |\eta(x) - 1/2| \leq \tilde{O}(k^{-1/2})\})$, the probability mass of a band of points around the decision boundary $\eta = 1/2$.

It is noteworthy that these upper and lower bounds have a pleasing resemblance for *every* distribution in the smoothness class. This is in contrast to the usual minimax style of analysis, in which a bound on an estimator's risk is described as "optimal" for a class of distributions if there exists even a single distribution in that class for which it is tight.

2.2.3 Margin bounds

An achievement of statistical theory in the past two decades has been *margin bounds*, which give fast rates of convergence for many classifiers when the underlying data distribution $\mathbb{P}$ (given by μ and η) satisfies a *large margin condition* stipulating, roughly, that η moves gracefully away from $1/2$ near the decision boundary.

Following [13, 16, 1], for any $\beta \geq 0$, we say $\mathbb{P}$ satisfies the β-*margin condition* if there exists a constant $C > 0$ such that

$$\mu\left(\left\{x \;\middle|\; |\eta(x) - \frac{1}{2}| \leq t\right\}\right) \leq Ct^\beta.$$

Larger β implies a larger margin. We now obtain bounds for the misclassification rate and the excess risk of k-NN under smoothness and margin conditions.

Theorem 4. *Suppose η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$ and satisfies the β-margin condition (with constant C), for some $\alpha, \beta, L, C \geq 0$. In each of the two following statements, k_o and C_o are constants depending on α, β, L, C.*

(a) *For any $0 < \delta < 1$, set $k = k_o n^{2\alpha/(2\alpha+1)}(\log(1/\delta))^{1/(2\alpha+1)}$. With probability at least $1 - \delta$ over the choice of training data,*

$$\Pr_X(g_{n,k}(X) \neq g(X)) \leq \delta + C_o\left(\frac{\log(1/\delta)}{n}\right)^{\alpha\beta/(2\alpha+1)}.$$

6

(b) Set $k = k_o n^{2\alpha/(2\alpha+1)}$. Then $\mathbb{E}_n R_{n,k} - R^* \leq C_o n^{-\alpha(\beta+1)/(2\alpha+1)}$.

It is instructive to compare these bounds with the best known rates for nonparametric classification under the margin assumption. The work of Audibert and Tsybakov [1] (Theorems 3.3 and 3.5) shows that when $(\mathcal{X}, \rho) = (\mathbb{R}^d, \|\cdot\|)$, and η is α_H-Holder continuous, and μ lies in the range $[\mu_{\min}, \mu_{\max}]$ for some $\mu_{\max} > \mu_{\min} > 0$, and the β-margin condition holds (along with some other assumptions), an excess risk of $n^{-\alpha_H(\beta+1)/(2\alpha_H+d)}$ is achievable and is also the best possible. This is exactly the rate we obtain for nearest neighbor classification, once we translate between the different notions of smoothness as per Lemma 2.

We discuss other interesting scenarios in Section C.4 in the appendix.

2.3 A general upper bound on the misclassification error

We now get to our most general finite-sample bound. It requires no assumptions beyond the basic measurability conditions stated at the beginning of Section 2, and it is the basis of the all the results described so far. We begin with some key definitions.

The radius and probability-radius of a ball. When dealing with balls, we will primarily be interested in their probability mass. To this end, for any $x \in \mathcal{X}$ and any $0 \leq p \leq 1$, define

$$r_p(x) = \inf\{r \mid \mu(B(x,r)) \geq p\}.$$

Thus $\mu(B(x, r_p(x))) \geq p$ (Lemma 23), and $r_p(x)$ is the smallest radius for which this holds.

The effective interiors of the two classes, and the effective boundary. When asked to make a prediction at point x, the k-NN classifier finds the k nearest neighbors, which can be expected to lie in $B(x, r_p(x))$ for $p \approx k/n$. It then takes an average over these k labels, which has a standard deviation of $\Delta \approx 1/\sqrt{k}$. With this in mind, there is a natural definition for the *effective interior* of the $Y = 1$ region: the points x with $\eta(x) > 1/2$ on which the k-NN classifier is likely to be correct:

$$\mathcal{X}_{p,\Delta}^+ = \left\{ x \in \text{supp}(\mu) \mid \eta(x) > \frac{1}{2}, \ \eta(B(x,r)) \geq \frac{1}{2} + \Delta \ \text{ for all } r \leq r_p(x) \right\}.$$

The corresponding definition for the $Y = 0$ region is

$$\mathcal{X}_{p,\Delta}^- = \left\{ x \in \text{supp}(\mu) \mid \eta(x) < \frac{1}{2}, \ \eta(B(x,r)) \leq \frac{1}{2} - \Delta \ \text{ for all } r \leq r_p(x) \right\}.$$

The remainder of $\mathcal{X}$ is the effective boundary,

$$\partial_{p,\Delta} = \mathcal{X} \setminus (\mathcal{X}_{p,\Delta}^+ \cup \mathcal{X}_{p,\Delta}^-).$$

Observe that $\partial_{p',\Delta'} \subset \partial_{p,\Delta}$ whenever $p' \leq p$ and $\Delta' \leq \Delta$. Under mild conditions, as p and Δ tend to zero, the effective boundary tends to the actual decision boundary $\{x \mid \eta(x) = 1/2\}$ (Lemma 14), which we shall denote ∂_o.

The misclassification rate of the k-NN classifier can be bounded by the probability mass of the effective boundary:

Theorem 5. *Pick any $0 < \delta < 1$ and positive integers $k < n$. Let $g_{n,k}$ denote the k-NN classifier based on n training points, and $g(x)$ the Bayes-optimal classifier. With probability at least $1 - \delta$ over the choice of training data,*

$$\Pr_X(g_{n,k}(X) \neq g(X)) \ \leq \ \delta + \mu(\partial_{p,\Delta}),$$

where

$$p = \frac{k}{n} \cdot \frac{1}{1 - \sqrt{(4/k)\ln(2/\delta)}}, \quad and \quad \Delta = \min\left(\frac{1}{2}, \sqrt{\frac{1}{k}\ln\frac{2}{\delta}} \right).$$

7

2.4 A general lower bound on the misclassification error

Finally, we give a counterpart to Theorem 5 that lower-bounds the expected probability of error of $g_{n,k}$. For any positive integers $k < n$, we identify a region close to the decision boundary in which a k-NN classifier has a constant probability of making a mistake. This high-error set is $\mathcal{E}_{n,k} = \mathcal{E}_{n,k}^{+} \cup \mathcal{E}_{n,k}^{-}$, where

$$\mathcal{E}_{n,k}^{+} = \left\{ x \in \mathrm{supp}(\mu) \mid \eta(x) > \frac{1}{2}, \ \ \eta(B(x,r)) \leq \frac{1}{2} + \frac{1}{\sqrt{k}} \text{ for all } r_{k/n}(x) \leq r \leq r_{(k+\sqrt{k}+1)/n}(x) \right\}$$

$$\mathcal{E}_{n,k}^{-} = \left\{ x \in \mathrm{supp}(\mu) \mid \eta(x) < \frac{1}{2}, \ \ \eta(B(x,r)) \geq \frac{1}{2} - \frac{1}{\sqrt{k}} \text{ for all } r_{k/n}(x) \leq r \leq r_{(k+\sqrt{k}+1)/n}(x) \right\}.$$

(Recall the definition (3) of $\eta(A)$ for sets A.) For smooth η this region turns out to be comparable to the effective decision boundary $\partial_{k/n,1/\sqrt{k}}$. Meanwhile, here is a lower bound that applies to any $(\mathcal{X}, \rho, \mu)$.

Theorem 6. *For any positive integers $k < n$, let $g_{n,k}$ denote the k-NN classifier based on n training points. There is an absolute constant c_o such that the expected misclassification rate satisfies*

$$\mathbb{E}_n \mathrm{Pr}_X(g_{n,k}(X) \neq g(X)) \geq c_o\, \mu(\mathcal{E}_{n,k}),$$

where $\mathbb{E}_n$ is expectation over the choice of training set.

Acknowledgements

The authors are grateful to the National Science Foundation for support under grant IIS-1162581.

8

References

[1] J.-Y. Audibert and A.B. Tsybakov. Fast learning rates for plug-in classifiers. *Annals of Statistics*, 35(2):608–633, 2007.

[2] T. Cover and P.E. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13:21–27, 1967.

[3] T.M. Cover. Rates of convergence for nearest neighbor procedures. In *Proceedings of The Hawaii International Conference on System Sciences*, 1968.

[4] S. Dasgupta. Consistency of nearest neighbor classification under selective sampling. In *Twenty-Fifth Conference on Learning Theory*, 2012.

[5] L. Devroye, L. Gyorfi, A. Krzyzak, and G. Lugosi. On the strong universal consistency of nearest neighbor regression function estimates. *Annals of Statistics*, 22:1371–1385, 1994.

[6] H. Federer. *Geometric Measure Theory*. Springer, 1969.

[7] E. Fix and J. Hodges. Discriminatory analysis, nonparametric discrimination. *USAF School of Aviation Medicine, Randolph Field, Texas, Project 21-49-004, Report 4, Contract AD41(128)-31*, 1951.

[8] J. Fritz. Distribution-free exponential error bound for nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 21(5):552–557, 1975.

[9] L. Gyorfi. The rate of convergence of k_n-nn regression estimates and classification rules. *IEEE Transactions on Information Theory*, 27(3):362–364, 1981.

[10] J. Heinonen. *Lectures on Analysis on Metric Spaces*. Springer, 2001.

[11] R. Kaas and J.M. Buhrman. Mean, median and mode in binomial distributions. *Statistica Neerlandica*, 34(1):13–18, 1980.

[12] S. Kulkarni and S. Posner. Rates of convergence of nearest neighbor estimation under arbitrary sampling. *IEEE Transactions on Information Theory*, 41(4):1028–1039, 1995.

[13] E. Mammen and A.B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.

[14] E. Slud. Distribution inequalities for the binomial law. *Annals of Probability*, 5:404–412, 1977.

[15] C. Stone. Consistent nonparametric regression. *Annals of Statistics*, 5:595–645, 1977.

[16] A.B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.

[17] R. Urner, S. Ben-David, and S. Shalev-Shwartz. Access to unlabeled data can speed up prediction time. In *International Conference on Machine Learning*, 2011.

[18] T.J. Wagner. Convergence of the nearest neighbor rule. *IEEE Transactions on Information Theory*, 17(5):566–571, 1971.

9

A Proofs of the general upper and lower bounds (Theorems 5 and 6)

A.1 A tie-breaking mechanism

In some situations, such as discrete instance spaces, there is a non-zero probability that two or more of the training points will be equidistant from the query point. In practice, we break ties by a simple rule such as preferring points that appear earlier in the sequence. To accurately reflect this in the analysis, we adopt the following mechanism: for each training point X, we also draw a value Z independently and uniformly at random from $[0, 1]$. When breaking ties, points with lower Z value are preferred. We use the notation $X' = (X, Z) \in \mathcal{X} \times [0, 1]$ to refer to the *augmented* training instances, drawn from the product measure $\mu \times \nu$, where ν is Lebesgue measure on $[0, 1]$.

Given a query point $x \in \mathcal{X}$ and training points $X'_1, \ldots, X'_n \in \mathcal{X} \times [0, 1]$, let $X'_{(1)}(x), \ldots, X'_{(n)}(x)$ denote a reordering of these points by increasing distance from x, where each $X'_{(i)}$ is of the form $(X_{(i)}, Z_{(i)})$. With probability 1, this ordering is unambiguous. Also, let $Y_{(1)}(x), \ldots, Y_{(n)}(x)$ be the corresponding labels.

We will need to consider balls in the augmented space. For $x_o \in \mathcal{X}$, $r_o \geq 0$, and $z_o \in [0, 1]$, define

$$B'(x_o, r_o, z_o) = \{(x, z) \in \mathcal{X} \times [0, 1] \mid \text{either } \rho(x_o, x) < r_o \text{ or } (\rho(x_o, x) = r_o \text{ and } z < z_o)\}$$

$$= \left(B^o(x_o, r_o) \times [0, 1]\right) \bigcup \left((B(x_o, r_o) \setminus B^o(x_o, r_o)) \times [0, z_o)\right).$$

Given a set of training points (X'_i, Y_i) and an augmented ball $B' = B'(x_o, r_o, z_o)$, let $\widehat{Y}(B')$ denote the mean of the Y_i for points $X'_i \in B'$; if there is no $X'_i \in B'$, then this is undefined.

Let $\eta(B')$ denote the mean probability that $Y = 1$ for points $(x, z) \in B'$; formally, it is given by

$$\eta(B') = \frac{1}{(\mu \times \nu)(B')} \int_{B'} \eta \, d(\mu \times \nu)$$

whenever $(\mu \times \nu)(B') > 0$. Here $\eta(x, z)$ is defined to be $\eta(x)$.

The ball $B' = B(x_o, r_o, z_o)$ in the augmented space can be thought of as lying between the open ball $B^o = B^o(x_o, r_o)$ and the closed ball $B = B(x_o, r_o)$ in the original space; and indeed $\eta(B')$ is a convex combination of $\eta(B)$ and $\eta(B^o)$ (Lemma 25).

A.2 Proof of Theorem 5

Theorem 5 rests on the following basic observation.

Lemma 7. *Let $g_{n,k}$ denote the k-NN classifier based on training data $(X'_1, Y_1), \ldots, (X'_n, Y_n)$. Pick any $x_o \in \mathcal{X}$ and any $0 \leq p \leq 1$, $0 \leq \Delta \leq 1/2$. Let $B' = B'(x_o, \rho(x_o, X_{(k+1)}(x_o)), Z_{(k+1)})$. Then*

$$
\begin{aligned}
1(g_{n,k}(x_o) \neq g(x_o)) \quad \leq \quad & 1(x_o \in \partial_{p,\Delta}) + \\
& 1(\rho(x_o, X_{(k+1)}(x_o)) > r_p(x_o)) + \\
& 1(|\widehat{Y}(B') - \eta(B')| \geq \Delta).
\end{aligned}
$$

Proof. Suppose $x_o \notin \partial_{p,\Delta}$. Then, without loss of generality, x_o lies in $\mathcal{X}^+_{p,\Delta}$, whereupon $\eta(B(x_o, r)) \geq 1/2 + \Delta$ for all $r \leq r_p(x_o)$.

Next, suppose $r = \rho(x_o, X_{(k+1)}(x_o)) \leq r_p(x_o)$. Then $\eta(B(x_o, r))$ and $\eta(B^o(x_o, r))$ are both $\geq 1/2 + \Delta$ (Lemma 26). By Lemma 25, $\eta(B')$ is a convex combination of these and is thus also $\geq 1/2 + \Delta$.

The prediction $g_{n,k}(x_o)$ is based on the average of the Y_i values of the k points closest to x_o, in other words, $\widehat{Y}(B')$. If this average differs from $\eta(B')$ by less than Δ, then it is $> 1/2$, whereupon the prediction is correct. $\qquad \square$

When we take expectations in the inequality of Lemma 7, we see that there are three probabilities to be bounded. The second of these, the probability that $\rho(x_o, X_{(k+1)}(x_o)) > r_p(x_o)$, can easily be controlled when p is sufficiently large.

10

Lemma 8. *Fix any $x_o \in \mathcal{X}$ and $0 \leq p, \gamma \leq 1$. Pick any positive integer $k \leq (1 - \gamma)np$. Let $X_1, \ldots, X_n$ be chosen uniformly at random from μ. Then*

$$\Pr_n(\rho(x_o, X_{(k+1)}(x_o)) > r_p(x_o)) \leq e^{-np\gamma^2/2} \leq e^{-k\gamma^2/2}.$$

Proof. The probability that any given X_i falls in $B(x_o, r_p(x_o))$ is at least p (Lemma 23). The probability that $\leq k \leq (1 - \gamma)np$ of them land in this ball is, by the multiplicative Chernoff bound, at most $e^{-np\gamma^2/2}$. $\square$

To bound the probability that $\widehat{Y}(B')$ differs substantially from $\eta(B')$, a slightly more careful argument is needed.

Lemma 9. *Fix any $x_o \in \mathcal{X}$ and any $0 \leq \Delta \leq 1/2$. Draw $(X_1, Z_1, Y_1), \ldots, (X_n, Z_n, Y_n)$ independently at random and let $B' = B'(x_o, \rho(x_o, X_{(k+1)}(x_o)), Z_{(k+1)}) \subset \mathcal{X} \times [0, 1]$. Then*

$$\Pr_n(|\widehat{Y}(B') - \eta(B')| \geq \Delta) \leq 2e^{-2k\Delta^2}.$$

Moreover, if $\eta(B') \in \{0, 1\}$ then $\widehat{Y}(B') = \eta(B')$ with probability one.

Proof. We will pick the points $X'_i = (X_i, Z_i)$ and their labels Y_i in the following manner:

1. First pick a point $(X_1, Z_1) \in \mathcal{X} \times [0, 1]$ according to the marginal distribution of the $(k + 1)$st nearest neighbor of x_o.

2. Pick k points uniformly at random from the distribution $\mu \times \nu$ restricted to $B' = B'(x_o, \rho(x_o, X_1), Z_1)$.

3. Pick $n - k - 1$ points uniformly at random from the distribution $\mu \times \nu$ restricted to $(\mathcal{X} \times [0, 1]) \setminus B'$.

4. Randomly permute the n points obtained in this way.

5. For each (X_i, Z_i) in the permuted order, pick a label Y_i from the conditional distribution $\eta(X_i)$.

The k nearest neighbors of x_o are the points picked in step 2. Their Y values are independent and identically distributed with expectation $\eta(B')$. The main bound in the lemma now follows from a direct application of Hoeffding's inequality.

The final statement of the lemma is trivial and is needed to cover situations in which $\Delta = 1/2$. $\square$

We now complete the proof of Theorem 5. Adopt the settings of p and Δ from the theorem statement, and define the central bad event to be

$$\text{BAD}(X_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n) = \mathbb{1}(\rho(X_o, X_{(k+1)}(X_o)) > r_p(X_o)) + \mathbb{1}(|\widehat{Y}(B') - \eta(B')| \geq \Delta),$$

where B' is a shorthand for $B'(X_o, \rho(X_o, X_{(k+1)}(X_o)), Z_{(k+1)})$, as before. Fix any $x_o \in \mathcal{X}$. If $\Delta < 1/2$, then by Lemmas 8 and 9,

$$\mathbb{E}_n \text{BAD}(x_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n) \leq \exp(-k\gamma^2/2) + 2\exp(-2k\Delta^2) \leq \delta^2,$$

where $\gamma = 1 - (k/np) = \sqrt{(4/k)\ln(2/\delta)}$ and $\mathbb{E}_n$ is expectation over the choice of training data. If $\Delta = 1/2$ then $\eta(B') \in \{0, 1\}$ and we have

$$\mathbb{E}_n \text{BAD}(x_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n) \leq \exp(-k\gamma^2/2) \leq \delta^2.$$

Taking expectation over X_o,

$$\mathbb{E}_{X_o} \mathbb{E}_n \text{BAD}(X_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n) \leq \delta^2,$$

from which, by switching expectations and applying Markov's inequality, we have

$$\Pr_n(\mathbb{E}_{X_o} \text{BAD}(X_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n) \geq \delta) \leq \delta.$$

The theorem then follows by writing the result of Lemma 7 as

$$\Pr_{X_o}(g_{n,k}(X_o) \neq g(X_o)) \leq \mu(\partial_{p,\Delta}) + \mathbb{E}_{X_o} \text{BAD}(X_o, X'_1, \ldots, X'_n, Y_1, \ldots, Y_n).$$

11

A.3 Proof of Theorem 6

For positive integer n and $0 \leq p \leq 1$, let $\text{bin}(n, p)$ denote the (binomial) distribution of the sum of n independent Bernoulli(p) random variables. We will use $\text{bin}(n, p; \geq k)$ to denote the probability that this sum is $\geq k$; and likewise $\text{bin}(n, p; \leq k)$.

It is well-known that the binomial distribution can be approximated by a normal distribution, suitably scaled. Slud [14] has finite-sample results of this form that will be useful to us.

Lemma 10. *Pick any $0 < p \leq 1/2$ and any nonnegative integer ℓ.*

    (a) [14, p. 404, item (v)] If $\ell \leq np$, then $\text{bin}(n, p; \geq \ell) \geq 1 - \Phi((\ell - np)/\sqrt{np})$.

    (b) [14, Thm 2.1] If $np \leq \ell \leq n(1 - p)$, then $\text{bin}(n, p; \geq \ell) \geq 1 - \Phi((\ell - np)/\sqrt{np(1 - p)})$.

Here $\Phi(a) = (2\pi)^{-1/2} \int_{-\infty}^{a} \exp(-t^2/2)dt$ is the cumulative distribution function of the standard normal.

Now we begin the proof of Theorem 6. Fix any integers $k < n$, and any $x_o \in \mathcal{E}_{n,k}$. Without loss of generality, $\eta(x_o) < 1/2$.

Pick $X_1, \ldots, X_n$ and $Z_1, \ldots, Z_n$ (recall the discussion on tie-breaking in Section A.1) in the following manner:

1. First pick a point $(X_1, Z_1) \in \mathcal{X} \times [0, 1]$ according to the marginal distribution of the $(k + 1)$st nearest neighbor of x_o.

2. Pick k points uniformly at random from the distribution $\mu \times \nu$ restricted to $B' = B'(x_o, \rho(x_o, X_1), Z_1)$; recall the earlier definition of the augmented space $\mathcal{X} \times [0, 1]$ and augmented balls within this space.

3. Pick $n - k - 1$ points uniformly at random from the distribution $\mu \times \nu$ restricted to $(\mathcal{X} \times [0, 1]) \setminus B'$.

4. Randomly permute the n points obtained in this way.

The $(k + 1)$st nearest neighbor of x_o, denoted $X_{(k+1)}(x_o)$, is the point chosen in the first step. With constant probability, it lies within a ball of probability mass $(k + \sqrt{k} + 1)/n$ centered at x_o, but not within a ball of probability mass k/n. Call this event G_1:

$$G_1: \quad r_{k/n}(x_o) \leq \rho(x_o, X_{(k+1)}(x_o)) \leq r_{(k+\sqrt{k}+1)/n}(x_o)$$

Lemma 11. *There is an absolute constant $c_1 > 0$ such that $\Pr(G_1) \geq c_1$.*

Proof. The expected number of points X_i that fall in $B(x_o, r_{(k+\sqrt{k}+1)/n}(x_o))$ is $\geq k + \sqrt{k} + 1$; the probability that the actual number is $\leq k$ is at most $\text{bin}(n, (k + \sqrt{k} + 1)/n; \leq k)$. Likewise, the expected number of points that fall in $B^o(x_o, r_{k/n}(x_o))$ is $\leq k$, and the probability that the actual number is $\geq k + 1$ is at most $\text{bin}(n, k/n; \geq k + 1)$. If neither of these bad events occurs, then G_1 holds. Therefore,

$$\Pr(G_1) \geq 1 - \text{bin}\left(n, \frac{k + \sqrt{k} + 1}{n}; \leq k\right) - \text{bin}\left(n, \frac{k}{n}; \geq k + 1\right).$$

The last term is easy to bound: it is $\leq 1/2$ since k is the median of $\text{bin}(n, k/n)$ [11]. To bound the first term, we use Lemma 10(a):

$$\text{bin}\left(n, \frac{k + \sqrt{k} + 1}{n}; \leq k\right) = 1 - \text{bin}\left(n, \frac{k + \sqrt{k} + 1}{n}; \geq k + 1\right)$$

$$\leq \Phi\left(\frac{(k + 1) - (k + \sqrt{k} + 1)}{\sqrt{k + \sqrt{k} + 1}}\right) \leq \Phi(-1/\sqrt{3}),$$

which is $1/2 - c_1$ for some constant $c_1 > 0$. Thus $\Pr(G_1) \geq c_1$. $\qquad\square$

12

Next, we lower-bound the probability that (conditional on event G_1), the k nearest neighbors of x_o have an average Y value with the wrong sign. Recalling that $\eta(x_o) < 1/2$, define the event

$$G_2 : \quad \widehat{Y}(B') > 1/2$$

where as before, B' denotes the ball $B'(x_o, X_{(k+1)}(x_o), Z_{k+1})$ in the augmented space.

Lemma 12. *There is an absolute constant $c_2 > 0$ such that $\Pr(G_2|G_1) \geq c_2$.*

Proof. Event G_1 depends only on step 1 of the sampling process. Assuming this event occurs, step 2 consists in drawing k points from the distribution $\mu \times \nu$ restricted to B'. Since $x_o \in \mathcal{E}_{n,k}$, we have (by an application of Lemmas 25 and 26) that $\eta(B') \geq 1/2 - 1/\sqrt{k}$. Now, $\widehat{Y}(B')$ follows a $\text{bin}(k, \eta(B'))$ distribution, and hence, by Lemma 10(b),

$$\Pr\left(\widehat{Y}(B') > \frac{k}{2}\right) = \Pr\left(\widehat{Y}(B') \geq \left\lceil \frac{k+1}{2} \right\rceil\right) \geq \Pr\left(Z \geq \frac{2\sqrt{k}+2}{\sqrt{k}}\right)$$

where Z is a standard normal. The last tail probability is at least some constant c_2. $\qquad\square$

In summary, for $x_o \in \mathcal{E}_{n,k}$,

$$\Pr_n(g_{n,k}(x_o) \neq g(x_o)) \geq \Pr(G_1 \wedge G_2) \geq c_1 c_2.$$

Taking expectation over x_o, we then get

$$\mathbb{E}_n \Pr_X(g_{n,k}(x) \neq g(x)) \geq c_1 c_2 \mu(\mathcal{E}_{n,k}),$$

as claimed.

B Proof of universal consistency (Theorem 1)

Recall that we define $R_n = \Pr_X(g_{n,k_n}(X) \neq Y)$. From equation (1), we have:

$$R_n - R^* \leq \Pr_X(\eta(X) \neq 1/2 \text{ and } g_{n,k_n}(X) \neq g(X)).$$

Defining $\partial_o = \{x \in \mathcal{X} \mid \eta(x) = 1/2\}$ to be the decision boundary, we then have the following corollary of Theorem 5.

Corollary 13. *Let (δ_n) be any sequence of positive reals, and (k_n) any sequence of positive integers. For each n, define (p_n) and (Δ_n) as in Theorem 5. Then*

$$\Pr_n\left(R_n - R^* > \delta_n + \mu\left(\partial_{p_n, \Delta_n} \setminus \partial_o\right)\right) \leq \delta_n,$$

where $\Pr_n$ is probability over the choice of training data.

For the rest of the proof, assume that $(\mathcal{X}, \rho, \mu)$ satisfies Lebesgue's differentiation theorem: that is, for any bounded measurable $f : \mathcal{X} \to \mathbb{R}$,

$$\lim_{r \downarrow 0} \frac{1}{\mu(B(x,r))} \int_{B(x,r)} f \, d\mu = f(x)$$

for almost all (μ-a.e.) $x \in \mathcal{X}$. We'll see that, as a result, $\mu(\partial_{p_n, \Delta_n} \setminus \partial_o) \to 0$.

Lemma 14. *There exists $\mathcal{X}_o \subset \mathcal{X}$ with $\mu(\mathcal{X}_o) = 0$, such that any $x \in \mathcal{X} \setminus \mathcal{X}_o$ with $\eta(x) \neq 1/2$ lies in $\mathcal{X}_{p,\Delta}^+ \cup \mathcal{X}_{p,\Delta}^-$ for some $p, \Delta > 0$.*

Proof. As a result of the differentiation condition,

$$\lim_{r \downarrow 0} \eta(B(x,r)) = \lim_{r \downarrow 0} \frac{1}{\mu(B(x,r))} \int_{B(x,r)} \eta \, d\mu = \eta(x) \qquad (4)$$

for almost all (μ-a.e.) $x \in \mathcal{X}$. Let $\mathcal{X}_o$ denote the set of x's for which (4) fails to hold or that are outside $\text{supp}(\mu)$. Then, $\mu(\mathcal{X}_o) = 0$.

Now pick any $x \notin \mathcal{X}_o$ such that $\eta(x) \neq 1/2$. Without loss of generality, $\eta(x) > 1/2$. Set $\Delta = (\eta(x) - 1/2)/2 > 0$. By (4), there is some $r_o > 0$ such that $\eta(B(x,r)) \geq 1/2 + \Delta$ whenever $0 \leq r \leq r_o$. Define $p = \mu(B(x, r_o)) > 0$. Then $r_p(x) \leq r_o$ and $x \in \mathcal{X}_{p,\Delta}^+$. $\qquad\square$

13

Lemma 15. *If $p_n, \Delta_n \downarrow 0$, then*
$$\lim_{n \to \infty} \mu\big(\partial_{p_n, \Delta_n} \setminus \partial_o\big) = 0.$$

Proof. Let $A_n = \partial_{p_n, \Delta_n} \setminus \partial_o$. Then $A_1 \supset A_2 \supset A_3 \supset \cdots$. We've seen earlier that for any $x \in \mathcal{X} \setminus (\mathcal{X}_o \cup \partial_o)$ (where $\mathcal{X}_o$ is defined in Lemma 14), there exist $p, \Delta > 0$ such that $x \notin \partial_{p, \Delta}$. Therefore,
$$\bigcap_{n \geq 1} A_n \subset \mathcal{X}_o,$$
whereupon, by continuity from above, $\mu(A_n) \to 0$. $\qquad\square$

Convergence in probability follows immediately.

Lemma 16. *If $k_n \to \infty$ and $k_n/n \to 0$, then for any $\epsilon > 0$,*
$$\lim_{n \to \infty} \Pr_n(R_n - R^* > \epsilon) = 0.$$

Proof. First define $\delta_n = \exp(-k_n^{1/2})$, and define the corresponding p_n, Δ_n as in Theorem 5. It is easily checked that the three sequences δ_n, p_n, Δ_n all go to zero.

Pick any $\epsilon > 0$. By Lemma 15, we can choose a positive integer N so that $\delta_n \leq \epsilon/2$ and $\mu(\partial_{p_n, \Delta_n} \setminus \partial_o) \leq \epsilon/2$ whenever $n \geq N$. Then by Corollary 13, for $n \geq N$,
$$\Pr_n(R_n - R^* > \epsilon) \leq \delta_n.$$
Now take $n \to \infty$. $\qquad\square$

We finish with almost sure convergence.

Lemma 17. *Suppose that in addition to the conditions of Lemma 16, we have $k_n/(\log n) \to \infty$. Then $R_n \to R^*$ almost surely.*

Proof. Choose $\delta_n = 1/n^2$, and for each n set p_n, Δ_n as in Theorem 5. It can be checked that the resulting sequences (p_n) and (Δ_n) both go to zero.

Pick any $\epsilon > 0$. Choose N so that $\sum_{n \geq N} \delta_n \leq \epsilon$. Letting ω denote a realization of an infinite training sequence $(X_1, Y_1), (X_2, Y_2), \ldots$, we have from Corollary 13 that
$$\Pr\big\{\omega \mid \exists n \geq N : R_n(\omega) - R^* > \delta_n + \mu\big(\partial_{p_n, \Delta_n} \setminus \partial_o\big)\big\} \leq \sum_{n \geq N} \delta_n \leq \epsilon.$$

Therefore, with probability at least $1 - \epsilon$ over ω, we have
$$R_n(\omega) - R^* \leq \delta_n + \mu\big(\partial_{p_n, \Delta_n} \setminus \partial_o\big)$$
for all $n \geq N$, whereupon, by Lemma 15, $R_n(\omega) \to R^*$. The result follows since this is true for any $\epsilon > 0$. $\qquad\square$

C Generalization bounds under smoothness conditions

C.1 Proof of Lemma 2

Suppose that η satisfies the α-Holder condition so that for some constant $C > 0$,
$$|\eta(x) - \eta(x')| \leq C\|x - x'\|^{\alpha_H}$$
whenever $x, x' \in \mathcal{X}$. For any $x \in \text{supp}(\mu)$ and $r > 0$, we then have
$$|\eta(x) - \eta(B(x, r))| \leq Cr^{\alpha_H}.$$
If μ has a density that is lower-bounded by $\mu_{\min}$, and $B(x, r) \subset \mathcal{X}$, we also have
$$\mu(B^o(x, r)) \geq \mu_{\min} v_d r^d,$$
where v_d is the volume of the unit ball in $\mathbb{R}^d$. The lemma follows by combining these two inequalities.

14

C.2 Proof of Theorem 3

Recall that the key term in the upper bound (Theorem 5) is $\mu(\partial_{p,\Delta})$, for $p \approx k/n$ and $\Delta \approx 1/\sqrt{k}$.

Lemma 18. *If η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$, then for any $p, \Delta \geq 0$,*

$$\partial_{p,\Delta} \cap \operatorname{supp}(\mu) \ \subset \ \left\{ x \in \mathcal{X} \ \middle| \ \left| \eta(x) - \frac{1}{2} \right| \leq \Delta + Lp^{\alpha} \right\}.$$

Proof. Pick any $x \in \operatorname{supp}(\mu)$ and any $p \geq 0$. For $r \leq r_p(x)$, we have $\mu(B^o(x,r)) \leq p$ and thus, by the definition of (α, L)-smoothness,

$$|\eta(B(x,r)) - \eta(x)| \ \leq \ Lp^{\alpha}.$$

As a result, if $\eta(x) > 1/2 + \Delta + Lp^{\alpha}$ then $\eta(B(x,r)) > 1/2 + \Delta$ whenever $r \leq r_p(x)$. Therefore, such an x lies in the effective interior $\mathcal{X}^+_{p,\Delta}$. A similar result applies to x with $\eta(x) < 1/2 - \Delta - Lp^{\alpha}$. Therefore, the boundary region $\partial_{p,\Delta}$ can only contain points x for which $|\eta(x) - 1/2| \leq \Delta + Lp^{\alpha}$, as claimed. $\square$

This yields a bound on $\operatorname{Pr}_X(g_{n,k}(X) \neq g(X))$ that is roughly of the form $\mu(\{x \mid |\eta(x) - 1/2| \leq k^{-1/2} + L(k/n)^{\alpha}\})$. The optimal setting of k is then $\sim n^{2\alpha/(2\alpha+1)}$.

The key term in the lower bound of Theorem 6 is $\mu(\mathcal{E}_{n,k})$. Under the smoothness condition, it becomes directly comparable to the upper bound.

Lemma 19. *If η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$, then for any k, n,*

$$\mathcal{E}_{n,k} \supset \left\{ x \in \operatorname{supp}(\mu) \ \middle| \ \eta(x) \neq \frac{1}{2}, \ \left| \eta(x) - \frac{1}{2} \right| \leq \frac{1}{\sqrt{k}} - L\left(\frac{k + \sqrt{k} + 1}{n} \right)^{\alpha} \right\}.$$

Proof. The proof is similar to that of the upper bound. Any point $x \in \operatorname{supp}(\mu)$ with

$$\frac{1}{2} < \eta(x) \leq \frac{1}{2} + \frac{1}{\sqrt{k}} - L\left(\frac{k + \sqrt{k} + 1}{n} \right)^{\alpha}$$

has $\eta(B(x,r)) \leq 1/2 + 1/\sqrt{k}$ for all $r \leq r_{(k+\sqrt{k}+1)/n}(x)$, and therefore lies in $\mathcal{E}^+_{n,k}$. Likewise for $\mathcal{E}^-_{n,k}$. $\square$

C.3 Proof of Theorem 4

Assume that η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$, that is,

$$|\eta(B(x,r)) - \eta(x)| \leq L\mu(B^o(x,r))^{\alpha} \tag{5}$$

for all $x \in \operatorname{supp}(\mu)$ and all $r > 0$, and also that it satisfies the β-margin condition (with constant C), under which, for any $t \geq 0$,

$$\mu\left(\left\{ x \ \middle| \ \left| \eta(x) - \frac{1}{2} \right| \leq t \right\} \right) \ \leq \ Ct^{\beta}. \tag{6}$$

Proof of Theorem 4(a)

Set p, Δ as specified in Theorem 5. It follows from that theorem and from Lemma 18 that under (5) and (6), for any $\delta > 0$, with probability at least $1 - \delta$ over the choice of training data,

$$\operatorname{Pr}_X(g_{n,k}(X) \neq g(X)) \ \leq \ \delta + \mu(\partial_{p,\Delta}) \ \leq \ \delta + C(\Delta + Lp^{\alpha})^{\beta}.$$

Expanding p, Δ in terms of k and n, this becomes

$$\operatorname{Pr}_X(g_{n,k}(X) \neq g(X)) \ \leq \ \delta + C\left(\left(\frac{\ln(2/\delta)}{k} \right)^{1/2} + L\left(\frac{2k}{n} \right)^{\alpha} \right)^{\beta},$$

provided $k \geq 16 \ln(2/\delta)$. The result follows by setting $k \propto n^{2\alpha/(2\alpha+1)}(\log(1/\delta))^{1/(2\alpha+1)}$.

15

Proof of Theorem 4(b)

Theorem 4(b) is an immediate consequence of Lemma 21 below. We begin, however, with an intermediate result about the pointwise expected risk.

Fix any n and any $k < n$, and set $p = 2k/n$. Define

$$\Delta(x) = |\eta(x) - 1/2|$$
$$\Delta_o = Lp^\alpha$$

Recall that the Bayes classifier $g(x)$ has risk $R^*(x) = \min(\eta(x), 1 - \eta(x))$ at x. The pointwise risk of the k-NN classifier $g_{n,k}$ is denoted $R_{n,k}(x)$.

Lemma 20. *Pick any $x \in \operatorname{supp}(\mu)$ with $\Delta(x) > \Delta_o$. Under (5),*

$$\mathbb{E}_n R_{n,k}(x) - R^*(x) \leq \exp(-k/8) + 4\Delta(x) \exp(-2k(\Delta(x) - \Delta_o)^2).$$

Proof. Assume without loss of generality that $\eta(x) > 1/2$. By (5), for any $0 \leq r \leq r_p(x)$, we have

$$\eta(B(x,r)) \geq \eta(x) - Lp^\alpha = \eta(x) - \Delta_o = \frac{1}{2} + (\Delta(x) - \Delta_o),$$

whereby $x \in \mathcal{X}_{p,\Delta(x)-\Delta_o}^+$ (and thus $x \notin \partial_{p,\Delta(x)-\Delta_o}$).

Next, recalling (1), and then applying Lemma 7,

$$R_{n,k}(x) - R^*(x) = 2\Delta(x)1(g_{n,k}(x) \neq g(x))$$
$$\leq 2\Delta(x)\left(1(\rho(x, X_{(k+1)}(x)) > r_p(x)) + 1(|\widehat{Y}(B') - \eta(B')| \geq \Delta(x) - \Delta_o)\right),$$

where B' is as defined in that lemma statement. We can now take expectation over the training data and invoke Lemmas 8 and 9 to conclude

$$\mathbb{E}_n R_{n,k}(x) - R^*(x) \leq 2\Delta(x)\left(\Pr{}_n(\rho(x, X_{(k+1)}(x)) > r_p(x)) + \Pr{}_n(|\widehat{Y}(B') - \eta(B')| \geq \Delta(x) - \Delta_o)\right)$$
$$\leq 2\Delta(x)\left(\exp\left(-\frac{k}{2}\left(1 - \frac{k}{np}\right)^2\right) + 2\exp\left(-2k(\Delta(x) - \Delta_o)^2\right)\right),$$

from which the lemma follows by substituting $p = 2k/n$ and observing $\Delta(x) \leq 1/2$. □

Lemma 21. *Under (5) and (6),*

$$\mathbb{E}_n R_{n,k} - R^* \leq \exp(-k/8) + 6C\max\left(2L\left(\frac{2k}{n}\right)^\alpha, \sqrt{\frac{8(\beta + 2)}{k}}\right)^{\beta+1}.$$

Proof. Recall the definitions of $p(= 2k/n)$ and $\Delta_o, \Delta(x)$ above. Further, for each integer $i \geq 1$, define $\Delta_i = \Delta_o \cdot 2^i$. Fix any $i_o \geq 1$.

Lemma 20 bounds the expected pointwise risk for any x with $\Delta(x) > \Delta_o$. We will apply it to points with $\Delta(x) > \Delta_{i_o}$. For all remaining x, we have $\mathbb{E}_n R_{n,k}(x) - R^*(x) \leq 2\Delta_{i_o}$. Taking expectation over X,

$$\mathbb{E}_n R_n - R^*$$
$$\leq \mathbb{E}_X\left[2\Delta_{i_o}1(\Delta(X) \leq \Delta_{i_o}) + \exp(-k/8) + 4\Delta(X)\exp(-2k(\Delta(X) - \Delta_o)^2)1(\Delta(X) > \Delta_{i_o})\right]$$
$$\leq 2\Delta_{i_o}\Pr{}_X(\Delta(X) \leq \Delta_{i_o}) + \exp(-k/8) + 4\mathbb{E}_X\left[\Delta(X)\exp(-2k(\Delta(X) - \Delta_o)^2)1(\Delta(X) > \Delta_{i_o})\right].$$

By the margin condition (6), we have $\Pr_X(\Delta(X) \leq t) \leq Ct^\beta$. Thus only the last expectation remains to be bounded. We do so by considering each interval $\Delta_i < \Delta(X) \leq \Delta_{i+1}$ separately:

$$\mathbb{E}_X\left[\Delta(X)\exp(-2k(\Delta(X) - \Delta_o)^2)1(\Delta_i < \Delta(X) \leq \Delta_{i+1})\right]$$
$$\leq \Delta_{i+1}\exp(-2k(\Delta_i - \Delta_o)^2)\Pr{}_X(\Delta(X) \leq \Delta_{i+1})$$
$$\leq C\Delta_{i+1}^{\beta+1}\exp(-2k(\Delta_i - \Delta_o)^2). \tag{7}$$

16

If we set

$$i_o \;=\; \max\left(1, \left\lceil \log_2 \sqrt{\frac{2(\beta+2)}{k\Delta_o^2}} \right\rceil \right),$$

then for $i \geq i_o$, the terms (7) are upper-bounded by a geometric series with ratio $1/2$. This is because the ratio of two successive terms can be bounded as

$$
\begin{aligned}
\frac{C\Delta_{i+1}^{\beta+1}\exp(-2k(\Delta_i - \Delta_o)^2)}{C\Delta_i^{\beta+1}\exp(-2k(\Delta_{i-1}-\Delta_o)^2)} &= 2^{\beta+1}\exp(-2k((2^i\Delta_o - \Delta_o)^2 - (2^{i-1}\Delta_o - \Delta_o)^2)) \\
&= 2^{\beta+1}\exp(-2k\Delta_o^2((2^i-1)^2 - (2^{i-1}-1)^2)) \\
&\leq 2^{\beta+1}\exp(-2^{2i-1}k\Delta_o^2) \\
&\leq 2^{\beta+1}\exp(-(\beta+2)) \;\leq\; 1/2.
\end{aligned}
$$

Therefore

$$
\begin{aligned}
\mathbb{E}_X &\left[\Delta(X)\exp(-2k(\Delta(X)-\Delta_o)^2)1(\Delta(X) > \Delta_{i_o})\right] \\
&= \sum_{i \geq i_o}\mathbb{E}_X\left[\Delta(X)\exp(-2k(\Delta(X)-\Delta_o)^2)1(\Delta_i < \Delta(X) \leq \Delta_{i+1})\right] \\
&\leq \sum_{i \geq i_o} C\Delta_{i+1}^{\beta+1}\exp(-2k(\Delta_i - \Delta_o)^2) \\
&\leq C\Delta_{i_o}^{\beta+1}\exp(-2k(\Delta_{i_o-1}-\Delta_o)^2) \;\leq\; C\Delta_{i_o}^{\beta+1}.
\end{aligned}
$$

Putting these together, we have $\mathbb{E}_n R_{n,k} - R^* \;\leq\; 6C\Delta_{i_o}^{\beta+1} + e^{-k/8}$. We finish by substituting $\Delta_{i_o} = 2^{i_o}\Delta_o$. $\qquad\square$

C.4 Margin bounds under various scenarios of interest

One interesting scenario is when η is bounded away from $1/2$, that is, there exists some Δ^* for which

$$\mu\left(\left\{x \;\Big|\; |\eta(x) - \tfrac{1}{2}| \leq \Delta^*\right\}\right) = 0$$

It follows from Lemma 18 that if η is (α, L)-smooth in $(\mathcal{X}, \rho, \mu)$, then $\mu(\partial_{p,\Delta}) = 0$ whenever $\Delta + Lp^\alpha \leq \Delta^*$. Invoking Theorem 5 with

$$k = \frac{n}{2}\left(\frac{\Delta^*}{2L}\right)^{1/\alpha}, \quad \delta = 2e^{-k(\Delta^*)^2/4},$$

yields an exponentially fast rate of convergence: $\Pr(g_n(X) \neq g(X)) \leq 2e^{-C_o n}$, where $C_o = (\Delta^*)^{2+1/\alpha}/(8(2L)^{1/\alpha})$.

A final case of interest is when $\eta \in \{0, 1\}$ almost everywhere, so that the Bayes risk R^* is zero: that is, there is no inherent uncertainty in the conditional probability distribution $p(y|x)$.

A useful quantity to consider in this case is the *effective interiors* of the classes as a whole:

$$
\begin{aligned}
\mathcal{X}_p^+ &= \{x \in \text{supp}(\mu) \mid \eta(x) = 1, \eta(B(x,r)) = 1 \text{ for all } r \leq r_p(x)\}. \\
\mathcal{X}_p^- &= \{x \in \text{supp}(\mu) \mid \eta(x) = 0, \eta(B(x,r)) = 0 \text{ for all } r \leq r_p(x)\}.
\end{aligned}
$$

Thus, $\mathcal{X}_p^+ = \mathcal{X}_{p,1/2}^+$, and $\mathcal{X}_p^- = \mathcal{X}_{p,1/2}^-$. The rest of $\mathcal{X}$ is the *effective boundary* between the two classes:

$$\partial_p = \mathcal{X} \setminus (\mathcal{X}_p^+ \cup \mathcal{X}_p^-).$$

Incorporating these two quantities into Theorem 5 yields a bound of the following form.

Lemma 22. *Let δ be any positive real and let $k < n$ be positive integers. With probability $\geq 1 - \delta$ over the choice of the training data, the error of the k-nearest neighbor classifier $g_{n,k}$ is bounded as:*

$$\Pr_X(g_{n,k}(X) \neq g(X)) \leq \delta + \mu(\partial_p),$$

17

where

$$p = \frac{k}{n} + \frac{2\log(2/\delta)}{n}\left(1 + \sqrt{1 + \frac{k}{\log(2/\delta)}}\right)$$

Proof. The proof is the same as that of Theorem 5, except that the probability of the central bad event is different. We will therefore bound the probabilities of these events.

Observe that under the conditions of the lemma, for any p,

$$\partial_{p,\frac{1}{2}} = \partial_p$$

Moreover, for any $x_0 \notin \partial_p$, $\eta(B'(x_0, r_p(x_0)))$ is either 0 or 1; this implies that for all $x \in B'(x_0, r_p(x_0))$ except those in a measure zero subset, $\eta(x)$ is either 0 or 1. Therefore, the probability $\Pr(\hat{Y}(B') \neq \eta(B'))$ is zero.

The rest of the lemma follows from plugging this fact in to the proof of Theorem 5 and some simple algebra. $\square$

In particular, observe that since p increases with increasing k, and the dependence on Δ is removed, the best bounds are achieved at $k = 1$ for:

$$p = \frac{1}{n} + \frac{2(1 + \sqrt{2})\log(2/\delta)}{n}$$

This corroborates the admissibility results of [2], which essentially state that there is no $k > 1$ such that the k-nearest neighbor algorithm has equal or better error than the 1-nearest neighbor algorithm against all distributions.

D Assorted technical lemmas

Lemma 23. *For any $x \in \mathcal{X}$ and $0 \leq p \leq 1$, we have $\mu(B(x, r_p(x))) \geq p$.*

Proof. Let $r^* = r_p(x) = \inf\{r \mid \mu(B(x,r)) \geq p\}$. For any $n \geq 1$, let $B_n = B(x, r^* + 1/n)$. Thus $B_1 \supset B_2 \supset B_3 \supset \cdots$, with $\mu(B_n) \geq p$. Since $B(x, r^*) = \bigcap_n B_n$, it follows by continuity from above of probability measures that $\mu(B_n) \to \mu(B(x, r^*))$, so this latter quantity is $\geq p$. $\square$

Lemma 24 (Cover-Hart). *$\mu(\mathrm{supp}(\mu)) = 1$.*

Proof. Let $\mathcal{X}_o$ denote a countable dense subset of $\mathcal{X}$. Now, pick any point $x \notin \mathrm{supp}(\mu)$; then there is some $r > 0$ such that $\mu(B(x, r)) = 0$. It is therefore possible to choose a ball B_x centered in $\mathcal{X}_o$, with rational radius, such that $x \in B_x$ and $\mu(B_x) = 0$. Since there are only countably many balls of this sort,

$$\mu(\mathcal{X} \setminus \mathrm{supp}(\mu)) \leq \mu\Big(\bigcup_{x \notin \mathrm{supp}(\mu)} B_x\Big) = 0.$$

$\square$

Lemma 25. *Pick any $x_o \in \mathrm{supp}(\mu)$, $r_o > 0$, and any Borel set $I \subset [0, 1]$. Define $B^o = B^o(x_o, r_o)$ and $B = B(x_o, r_o)$ to be open and closed balls centered at x_o, and let $A \subset \mathcal{X} \times [0, 1]$ be given by $A = (B^o \times [0,1]) \bigcup ((B \setminus B^o) \times I)$. Then*

$$\eta(A) = \frac{\mu(B)\nu(I)}{\mu(B)\nu(I) + \mu(B^o)(1 - \nu(I))}\eta(B) + \frac{\mu(B^o)(1 - \nu(I))}{\mu(B)\nu(I) + \mu(B^o)(1 - \nu(I))}\eta(B^o).$$

18

Proof. Since x_o lies in the support of μ, we have $(\mu \times \nu)(A) \geq \mu(B^o) > 0$; hence $\eta(A)$ is well-defined.

$$
\begin{aligned}
\eta(A) &= \frac{1}{(\mu \times \nu)(A)} \int_A \eta \, d(\mu \times \nu) \\
&= \frac{1}{\mu(B^o) + \mu(B \setminus B^o)\nu(I)} \left(\int_{B^o} \eta \, d\mu + \int_{B \setminus B^o} \nu(I)\eta \, d\mu \right) \\
&= \frac{1}{\mu(B^o) + (\mu(B) - \mu(B^o))\nu(I)} \left(\int_{B^o} \eta \, d\mu + \nu(I) \left(\int_B \eta \, d\mu - \int_{B^o} \eta \, d\mu \right) \right) \\
&= \frac{\mu(B^o)\eta(B^o) + \nu(I)(\mu(B)\eta(B) - \mu(B^o)\eta(B^o))}{\mu(B^o)(1 - \nu(I)) + \mu(B)\nu(I)},
\end{aligned}
$$

as claimed. □

Lemma 26. *Suppose that for some $x_o \in \operatorname{supp}(\mu)$ and $r_o > 0$ and $q > 0$, it is the case that $\eta(B(x_o, r)) \geq q$ whenever $r \leq r_o$. Then $\eta(B^o(x_o, r_o)) \geq q$ as well.*

Proof. Let $B^o = B^o(x_o, r_o)$. Since

$$
B^o = \bigcup_{r < r_o} B(x_o, r),
$$

it follows from the continuity from below of probability measures that

$$
\lim_{r \uparrow r_o} \mu(B(x_o, r)) = \mu(B^o)
$$

and by dominated convergence that

$$
\lim_{r \uparrow r_o} \int_{B(x_o, r)} \eta \, d\mu = \lim_{r \uparrow r_o} \int_{B^o} 1(x \in B(x_o, r))\eta(x)\mu(dx) = \int_{B^o} \eta \, d\mu.
$$

For any $r \leq r_o$, we have $\eta(B(x_o, r)) \geq q$, which can be rewritten as

$$
\int_{B(x_o, r)} \eta \, d\mu - q \, \mu(B(x_o, r)) \geq 0.
$$

Taking the limit $r \uparrow r_o$, we then get the desired statement. □