

UNIVERSAL CONSISTENCY OF THE k -NN RULE IN METRIC SPACES AND NAGATA DIMENSION^{*,**}

BENOÎT COLLINS¹, SUSHMA KUMARI² AND VLADIMIR G. PESTOV^{3,4,***}

Abstract. The k nearest neighbour learning rule (under the uniform distance tie breaking) is universally consistent in every metric space X that is sigma-finite dimensional in the sense of Nagata. This was pointed out by Cérou and Guyader (2006) as a consequence of the main result by those authors, combined with a theorem in real analysis sketched by D. Preiss (1971) (and elaborated in detail by Assouad and Quentin de Gromard (2006)). We show that it is possible to give a direct proof along the same lines as the original theorem of Charles J. Stone (1977) about the universal consistency of the k -NN classifier in the finite dimensional Euclidean space. The generalization is non-trivial because of the distance ties being more prevalent in the non-Euclidean setting, and on the way we investigate the relevant geometric properties of the metrics and the limitations of the Stone argument, by constructing various examples.

Mathematics Subject Classification. 62H30, 54F45.

Received February 28, 2020. Accepted June 22, 2020.

1. INTRODUCTION

The k -nearest neighbour classifier, in spite of being arguably the oldest supervised learning algorithm in existence, still retains his importance, both practical and theoretical. In particular, it was the first classification learning rule whose (weak) universal consistency (in the finite-dimensional Euclidean space $\mathbb{R}^d = \ell^2(d)$) was established, by Charles J. Stone in [20].

Stone's result is easily extended to all finite-dimensional normed spaces, see, *e.g.*, [6]. However, the k -NN classifier is no longer universally consistent already in the infinite-dimensional Hilbert space ℓ^2 . A series of examples of this kind, obtained in the setting of real analysis, belongs to Preiss, and the first of them [17] is

* S.K. would like to thank JICA-FRIENDSHIP scholarship (fellowship D-1590283/ J-1593019) for supporting her doctoral study at Kyoto University and Department of Mathematics, Kyoto University for supporting the Brazil research trip under the KTGU project.

** V.G.P. wants to acknowledge support from CNPq (bolsa Pesquisador Visitante, processo 310012/2016) and CAPES (bolsa Professor Visitante Estrangeiro Sênior, processo 88881.117018/2016-01).

Keywords and phrases: k -NN classifier, universal consistency, geometric Stone lemma, distance ties, Nagata dimension, sigma-finite dimensional metric spaces.

¹ Department of Mathematics, Kyoto University, Kitashirakawa Oiwake-cho, Sakyo-ku 606-8502, Japan.

² Department of Mathematical Engineering, Musashino University, 1 Chome-1-20 Shinmachi, Nishitokyo, Tokyo 202-8585, Japan.

³ Instituto de Matemática e Estatística, Universidade Federal da Bahia, Ondina, Salvador-BA, 40.170-115, Brazil.

⁴ Department of Mathematics and Statistics, University of Ottawa, Ottawa, ON K1N 6N5, Canada.

*** Corresponding author: vpest283@uottawa.ca

so simple that it can be described in a few lines. We will reproduce it in the article, since the example remains virtually unknown in the statistical machine learning community.

There is sufficient empirical evidence to support the view that the performance of the k -NN classifier greatly depends on the chosen metric on the domain (see *e.g.*, [9]). There is a supervised learning algorithm, Large Margin Nearest Neighbour Classifier (LMNN), based on the idea of optimizing the k -NN performance over all Euclidean metrics on a finite dimensional vector space [21]. At the same time, it appears that a theoretical foundation for such an optimization over a set of distances is still lacking. The first question to address in this connection, is of course to characterize those metrics (generating the original Borel structure of the domain) for which the k -NN classifier is (weakly) universally consistent.

While the problem in this generality remains still open, a great advance in this direction was made by Cérou and Guyader in [2]. They have shown that the k -NN classifier is consistent under the assumption that the regression function $\eta(x)$ satisfies the weak Lebesgue–Besicovitch differentiation property:

$$\frac{1}{\mu(B_r(x))} \int_{B_r(x)} \eta(x) \, d\mu(x) \rightarrow \eta(x), \tag{1.1}$$

where the convergence is in measure, that is, for each $\epsilon > 0$,

$$\mu \left\{ x \in \Omega : \left| \frac{1}{\mu(B_r(x))} \int_{B_r(x)} \eta(x) \, d\mu(x) - \eta(x) \right| > \epsilon \right\} \rightarrow 0 \text{ when } r \downarrow 0.$$

The probability measure μ above is the sample distribution law. The proof extended the ideas of the paper [4], in which it was previously observed that Stone’s universal consistency can be deduced from the classical Lebesgue–Besicovitch differentiation theorem: every $L^1(\mu)$ -function f on \mathbb{R}^d satisfies equation (1.1), even in the strong sense (convergence almost everywhere). See also [8].

Those separable metric spaces in which the weak Lebesgue–Besicovitch differentiation property holds for every Borel probability measure (equivalently, for every sigma-finite locally finite Borel measure) have not yet been characterized. But the complete separable metric spaces in which the *strong* Lebesgue–Besicovitch differentiation property holds for every such measure as above have been described by Preiss [18]: they are exactly those spaces that are *sigma-finite dimensional* in the sense of Nagata [13, 16]. (For finite dimensional spaces in the sense of Nagata, the sketch of a proof by Preiss, in the sufficiency direction, was elaborated by Assouad and Quentin de Gromard in [1]. The completeness assumption on the metric space is only essential for the necessity part of the result.) In particular, it follows that every sigma-finite dimensional separable metric space satisfies the weak Lebesgue–Besicovitch differentiation property for every probability measure.

Combining the result of Preiss with that of Cérou–Guyader, one concludes that the k -NN classifier is universally consistent in every sigma-finite dimensional separable metric space, as was noted in [2].

The authors of [2] mention in their paper that “[Stone’s theorem] is based on a geometrical result, known as Stone’s Lemma. This powerful and elegant argument can unfortunately not be generalized to infinite dimension.” The aim of this article is to show that at least Stone’s original proof, including Stone’s geometric lemma as its main tool, can be extended from the Euclidean case to the sigma-finite dimensional metric spaces. In fact, as we will show, the geometry behind Stone’s lemma, even if it appears to be essentially based on the Euclidean structure of the space, is captured by the notion of Nagata dimension, which is a purely metric concept. In this way, Stone’s geometric lemma and indeed the original Stone’s proof of the universal consistency of the k -NN classifier, become applicable to a wide range of metric spaces.

In the absence of distance ties (that is, in case where every sphere is a μ -negligible set with regard to the underlying measure μ), the extension is quite straightforward, indeed almost literal. However, this is not so in the presence of distance ties: an example shows that the conclusion of Stone’s geometric lemma may not hold. Another example shows that even in the compact metric spaces of Nagata dimension zero, the distance ties may be everpresent. We also show that an attempt to reduce the case to the situation without distance ties

by learning in the product of Ω with the unit interval (an additional random variable used for tie-breaking) cannot work, because already the product of a zero-dimensional space in the sense of Nagata with the interval (which has dimension one) can have an infinite Nagata dimension. Stone's geometric lemma has to be modified, to parallel the Hardy–Littlewood inequality in the geometric measure theory.

We do not touch upon the subject of strong universal consistency in general metric spaces. The main open question left is whether every metric space in which the k -NN classifier is universally consistent is necessarily sigma-finite dimensional. A positive answer, modulo the work of [2] and [18], would also answer in the affirmative an open question in real analysis going back to Preiss: suppose a metric space X satisfies the weak Lebesgue–Besicovitch differentiation property for every sigma-finite locally finite Borel measure, will it satisfy the strong Lebesgue–Besicovitch differentiation property for every such measure?

2. SETTING FOR STATISTICAL LEARNING

Here we will recall the standard probabilistic model for statistical learning theory. The *domain*, Ω , means a standard Borel space, that is, a set equipped with a sigma-algebra which coincides with the sigma-algebra of Borel sets generated by a suitable separable complete metric. (Recall that the Borel structure generated by a metric ρ on a set Ω is the smallest sigma-algebra containing all open subsets of the metric space (Ω, ρ) .) The distribution laws for datapoints, both unlabelled and labelled, are Borel probability measures defined on the corresponding Borel sigma-algebra.

Since we will be dealing with the k -NN classifier, the domain, Ω , will actually be a metric space, which we also assume to be separable.

Labelled data pairs (x, y) , where $x \in \Omega$ and $y \in \{0, 1\}$, will follow an unknown probability distribution $\tilde{\mu}$, that is, a Borel probability measure on $\Omega \times \{0, 1\}$. We denote the corresponding random element $(X, Y) \sim \tilde{\mu}$. Define two Borel measures on Ω , μ_i , $i = 0, 1$, by $\mu_i(A) = \tilde{\mu}(A \times \{i\})$. In this way, μ_0 is governing the distribution of the elements labelled 0, and similarly for μ_1 . The sum $\mu = \mu_0 + \mu_1$ (the direct image of $\tilde{\mu}$ under the projection from $\Omega \times \{0, 1\}$ onto Ω) is a Borel probability measure on Ω , the distribution law of unlabelled data points (See Fig. 1). Clearly, μ_i is absolutely continuous with regard to μ , that is, if $\mu(A) = 0$, then $\mu_i(A) = 0$ for $i = 0, 1$. The corresponding Radon-Nikodým derivative in the case $i = 1$ is just the conditional probability for a point x to be labeled 1:

$$\eta(x) = \frac{d\mu_1}{d\mu}(x) = P[Y = 1|X = x].$$

In statistical terminology, η is the *regression function*.

Together with the Borel probability measure μ on Ω , the regression function allows for an alternative, and often more convenient, description of the joint law $\tilde{\mu}$. Namely, given $A \subseteq \Omega$,

$$\mu_1(A) = \int_A \eta(x) \, d\mu,$$

and

$$\mu_0(A) = \int_A (1 - \eta(x)) \, d\mu,$$

which allows to reconstruct the measure $\tilde{\mu}$ on $\Omega \times \{0, 1\}$.

Let $\mathcal{B}(\Omega, \{0, 1\})$ denote the collection of all Borel measurable binary functions on the domain, that is, essentially, the family of all Borel subsets of Ω . Given such a function $f: \Omega \rightarrow \{0, 1\}$ (a *classifier*), the *misclassification error* is defined by

$$\text{err}_{\tilde{\mu}}(f) = \tilde{\mu}\{(x, y): f(x) \neq y\} = P[f(X) \neq Y].$$

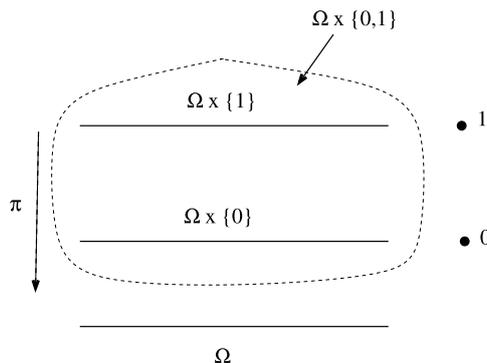


FIGURE 1. Labeled domain and the projection $\pi : \Omega \times \{0, 1\} \rightarrow \Omega$.

The *Bayes error* is the infimal misclassification error taken over all possible classifiers:

$$\ell^* = \ell^*(\tilde{\mu}) = \inf_f \text{err}_{\tilde{\mu}}(f).$$

It is a simple exercise to verify that the Bayes error is achieved on some classifier (and thus is the minimum), which is called a *Bayes classifier*. For instance, every classifier satisfying

$$T_{\text{bayes}}(x) = \begin{cases} 1, & \text{if } \eta(x) > \frac{1}{2}, \\ 0, & \text{if } \eta(x) < \frac{1}{2}, \end{cases}$$

is a Bayes classifier.

The Bayes error is zero if and only if the learning problem is deterministic, that is, the regression function η is equal almost everywhere to the indicator function, χ_C , of a *concept* $C \subseteq \Omega$, a Borel subset of the domain.

A *learning rule* is a family of mappings $\mathcal{L} = (\mathcal{L}_n)_{n=1}^\infty$, where

$$\mathcal{L}_n : \Omega^n \times \{0, 1\}^n \rightarrow \mathcal{B}(\Omega, \{0, 1\}), \quad n = 1, 2, \dots$$

and the functions \mathcal{L}_n satisfy the following measurability assumption: the associated maps

$$\Omega^n \times \{0, 1\}^n \times \Omega \ni (\sigma, x) \mapsto \mathcal{L}_n(\sigma)(x) \in \{0, 1\}$$

are Borel (or just universally measurable). Here $\sigma = (x_1, \dots, x_n, y_1, \dots, y_n)$ is a labelled learning sample.

The data is modelled by a sequence of independent identically distributed random elements (X_n, Y_n) of $\Omega \times \{0, 1\}$, following the law $\tilde{\mu}$. Denote ς an infinite sample path. In this context, \mathcal{L}_n only gets to see the first n labelled coordinates of ς . A learning rule \mathcal{L} is *weakly consistent*, or simply *consistent*, if $\text{err}_{\tilde{\mu}} \mathcal{L}_n(\varsigma) \rightarrow \ell^*$ in probability as $n \rightarrow \infty$. If the convergence occurs almost surely (that is, along almost all sample paths $\varsigma \sim \tilde{\mu}^\infty$), then \mathcal{L} is said to be *strongly consistent*. Finally, \mathcal{L} is *universally (weakly / strongly) consistent* if it is weakly / strongly consistent under every Borel probability measure $\tilde{\mu}$ on the standard Borel space $\Omega \times \{0, 1\}$.

The learning rule we study is the k -NN classifier, defined by selecting the label $\mathcal{L}_n(\sigma)(x) \in \{0, 1\}$ by the majority vote among the values of y corresponding to the $k = k_n$ nearest neighbours of x in the learning sample σ .

If k is even, then a voting tie may occur. This is of lesser importance, and can be broken in any way. For instance, by always assigning the value 1 in case of a voting tie, or by choosing the value randomly. The

consistency results usually do not depend on it. Intuitively, if voting ties keep occurring asymptotically at a point x along a sample path, it means that $\eta(x) = 1/2$ and so any value of the classifier assigned to x would do.

It may also happen that the smallest closed ball containing k nearest neighbours of a point x contains more than k elements of a sample (distance ties). This situation is more difficult to manage and requires a consistent tie-breaking strategy, whose choice may affect the consistency results.

Given k and $n \geq k$, we define $r_{k\text{-NN}}^{\zeta_n}(x)$ as the smallest radius of a closed ball around x containing at least k nearest neighbours of x in the sample ζ_n :

$$r_{k\text{-NN}}^{\zeta_n}(x) = \min\{r \geq 0: \#\{i = 1, 2, \dots, n: x_i \in \bar{B}_r(x)\} \geq k\}. \tag{2.1}$$

As the corresponding open ball around x contains at most $k - 1$ elements of the sample, the ties may only occur on the sphere.

We adopt the combinatorial notation $[n] = \{1, 2, \dots, n\}$. If $\sigma \in \Omega^n$ and $\sigma' \in \Omega^k$, $k \leq n$, the symbol

$$\sigma' \sqsubset \sigma$$

means that there is an injection $f: [k] \rightarrow [n]$ such that

$$\forall i = 1, 2, \dots, k, \sigma'_i = \sigma_{f(i)}.$$

A k nearest neighbour map is a function

$$k\text{-NN}^\sigma: \Omega^n \times \Omega \rightarrow \Omega^k$$

with the properties

- (1) $k\text{-NN}^\sigma(x) \sqsubset \sigma$, and
- (2) all points x_i in σ that are at a distance strictly less than $r_{k\text{-NN}}^{\zeta_n}(x)$ to x are in $k\text{-NN}^\sigma(x)$.

The mapping $k\text{-NN}^\sigma$ can be deterministic or stochastic, in which case it will depend on an additional random variable, independent of the sample.

An example of the former kind is based on the natural order on the sample, $x_1 < x_2 < \dots < x_n$. In this case, from among the points belonging to the sphere of radius $r_{k\text{-NN}^\sigma}(x)$ around x we choose points with the smaller index: $k\text{-NN}^\sigma(x)$ contains all the points of σ in the open ball, $B_{r_{k\text{-NN}^\sigma}(x)}(x)$, plus a necessary number (at least one) of points of $\sigma \cap S_{r_{k\text{-NN}^\sigma}(x)}(x)$ having smallest indices.

An example of the second kind is to use a similar procedure, after applying a random permutation of the indices first. A random learning input will consist of a pair (W_n, P_n) , where W_n is a random n -sample and P_n is a random element of the group of permutations of rank n . An equivalent (and more common) way would be to use a sequence of i.i.d. random elements Z_n of the unit interval or the real line, distributed according to the uniform (resp. gaussian) law, and in case of a tie, give a preference to a realization x_i over x_j provided the value z_i is smaller than z_j .

Now, a formal definition of the $k\text{-NN}$ learning rule can be given as follows:

$$\begin{aligned} \mathcal{L}_n^{k\text{-NN}}(\sigma, \epsilon)(x) &= \chi_{[0, \infty)} \left[\frac{1}{k} \sum_{x_i \in k\text{-NN}^\sigma(x)} \epsilon_i - \frac{1}{2} \right] \\ &= \chi_{[0, \infty)} \left[\mathbb{E}_{\mu_{k\text{-NN}^\sigma(x)}} \epsilon - \frac{1}{2} \right]. \end{aligned}$$

Here, $\chi_{[0,\infty)}$ is the Heaviside function, the sign of the argument:

$$\chi_{[0,\infty)}(t) = \begin{cases} 1, & \text{if } t \geq 0, \\ 0, & \text{if } t < 0. \end{cases}$$

The empirical measure $\mu_{k\text{-NN}^\sigma(x)}$ is a uniform measure supported on the set of k nearest neighbours of x within the sample σ , and the label ϵ is seen as a function $\epsilon: \{x_1, x_2, \dots, x_n\} \rightarrow \{0, 1\}$.

The expression appearing under the argument,

$$\eta_{n,k} = \frac{1}{k} \sum_{x_i \in k\text{-NN}^\sigma(x)} \epsilon_i, \tag{2.2}$$

is the *empirical regression function*. In the presence of a law of labelled points, it is a random variable, and so we have the following immediate, yet important, observation.

Proposition 2.1. *Let (μ, η) be a learning problem in a separable metric space (Ω, d) . If the values of the empirical regression function, $\eta_{n,k}$, converge to η in probability (resp. almost surely) in the region*

$$\Omega_\eta = \left\{ x \in \Omega: \eta(x) \neq \frac{1}{2} \right\},$$

then the k -NN classifier is consistent (resp. strongly consistent) under (μ, η) .

We conclude this section by recalling an important technical tool.

Theorem 2.2 (Cover-Hart lemma [3]). *Let Ω be a separable metric space, and let μ be a Borel probability measure on Ω . Almost surely, the function $r_{k\text{-NN}}^{\sigma_n}$ (Eq. (2.1)) converges to zero uniformly over any precompact subset $K \subseteq \text{supp } \mu$.*

Proof. Let A be a countable dense subset of $\text{supp } \mu$. A standard argument shows that, almost surely, for all $a \in A$ and each rational $\epsilon > 0$, the open ball $B_\epsilon(a)$ contains an infinite number of elements of a sample path. Consequently, the functions $r_{k\text{-NN}}^{\sigma_n}: \Omega \rightarrow \mathbb{R}$ almost surely converge to zero pointwise on A as $n \rightarrow \infty$. Since these functions are easily seen to be 1-Lipschitz and in particular form a uniformly equicontinuous family, we conclude. □

3. EXAMPLE OF PREISS

Here we will discuss a 1979 example of Preiss [17]. Preiss’s aim was to prove that the Lebesgue–Besicovitch differentiation property (Eq. (1.1)) fails in the infinite-dimensional Hilbert space ℓ^2 . However, as already suggested in [2], his example can be easily adapted to prove that the k -NN learning rule is not universally consistent in the infinite-dimensional separable Hilbert space ℓ^2 either.

Recall the notation $[n] = \{1, 2, \dots, n\}$. Let (N_k) be a sequence of positive natural numbers ≥ 2 , to be selected later. Denote by

$$Q = \prod_{k=1}^{\infty} [N_k]$$

the Cartesian product of finite discrete spaces equipped with the product topology. It is a Cantor space (the unique, up to a homeomorphism, totally disconnected compact metrizable space without isolated points).

Let π_k denote the canonical coordinate projections of Q on the k -dimensional cubes $Q_k = \prod_{i=1}^k [N_i]$. Denote $Q^* = \cup_{k=1}^\infty Q_k$ a disjoint union of the cubes Q_k , and let $\mathcal{H} = \ell^2(Q^*)$ be a Hilbert space spanned by an orthonormal basis $(e_{\bar{n}})$ indexed by elements \bar{n} of this union.

For every $\bar{n} = (n_1, \dots, n_i, \dots) \in Q$ define

$$f(\bar{n}) = \sum_{i=1}^\infty 2^{-i} e_{(n_1, \dots, n_i)} \in \mathcal{H}.$$

The map f is continuous and injective, thus a homeomorphism onto its image. Denote ν the Haar measure on Q (the product of the uniform measures on all $[N_k]$). Let $\mu_1 = f_*(\nu)$ be the direct image of ν , a compactly-supported Borel probability measure on \mathcal{H} . If $r > 0$ satisfies $2^{-k} \leq r^2 < 2^{-k+1}$, then for each $\bar{n} = (n_1, n_2, \dots) \in Q$,

$$\begin{aligned} \mu_1(B_r(f(\bar{n}))) &= \nu(\pi_{k+1}^{-1}(\bar{n})) \\ &= (N_1 N_2 \dots N_{k+1})^{-1}. \end{aligned}$$

Now, for every k and each $\bar{n} = (n_1, \dots, n_k) \in Q_k \subseteq Q^*$ define in a similar way

$$f(\bar{n}) = \sum_{i=1}^k 2^{-i} e_{(n_1, \dots, n_i)} \in \mathcal{H}.$$

Note that the closure of $f(Q^*)$ contains $f(Q)$ (as a proper subset). Now define a purely atomic measure μ_0 supported on the image of Q^* under f , having the following special form:

$$\mu_0 = \sum_{k=1}^\infty \sum_{\bar{n} \in Q_k} a_k \delta_{\bar{n}}.$$

The weights $a_k > 0$ are chosen so that the measure is finite:

$$\sum_{k=1}^\infty a_k < \infty. \tag{3.1}$$

Since for r satisfying $2^{-k} \leq r^2 < 2^{-k+1}$ and $\bar{n} \in Q$ the ball $B_r(f(\bar{n}))$ contains in particular $f(n_1 n_2, \dots, n_k)$, we have

$$\mu_0(B_r(f(\bar{n}))) \geq a_k.$$

Assuming in addition that

$$a_k N_1 N_2 \dots N_k N_{k+1} \rightarrow \infty \text{ as } k \rightarrow \infty, \tag{3.2}$$

we conclude:

$$\frac{\mu_1(B_r(f(\bar{n})))}{\mu_0(B_r(f(\bar{n})))} \leq \frac{(N_1 N_2 \dots N_{k+1})^{-1}}{a_k} \rightarrow 0 \text{ when } r \downarrow 0.$$

Clearly, the conditions (3.1) and (3.2) can be simultaneously satisfied by a recursive choice of (N_k) and (a_k) .

Now renormalize the measures μ_0 and μ_1 so that $\mu = \mu_0 + \mu_1$ is a probability measure, and interpret μ_i as the distribution of points labelled $i = 0, 1$. Thus, the regression function is deterministic, $\eta = \chi_C$, where we are learning the concept $C = f(Q) = \text{supp } \mu_1$, $\mu_1(C) > 0$.

For a random element $X \in \mathcal{H}$, $X \sim \mu$, the distance $r_k(X)$ to the k th nearest neighbour within an i.i.d. n -sample goes to zero almost surely when $k/n \rightarrow 0$, according to a lemma of Cover and Hart, and the convergence is uniform on the precompact support of μ . It follows that the probability of one of the k nearest neighbours to a random point $X \in \mathcal{H}$ to be labelled one, conditionally on $r_{k\text{-NN}}^{S_n} = r$, converges to zero, uniformly in r . The k -NN learning rule will almost surely predict a sequence of classifiers converging to the identically zero classifier, and so is not consistent.

4. CLASSICAL THEOREM OF CHARLES J. STONE

4.1. The case of continuous regression function

Proposition 2.1 and the Cover–Hart lemma 2.2 together imply that the k -NN classifier is universally consistent in a separable metric space whenever the regression function η is continuous. In view of Proposition 2.1, it is enough to make the following observation.

Lemma 4.1. *Let (Ω, μ) be a separable metric space equipped with a Borel probability measure, and let η be a continuous regression function. Then*

$$\eta_{n,k} \rightarrow \eta$$

in probability, when $n, k \rightarrow \infty$, $k/n \rightarrow 0$.

Proof. It follows from the Cover–Hart lemma that the set $k\text{-NN}^{S_n}(x)$ of k nearest neighbours of x almost surely converges to x , for almost all $x \in \text{supp } \mu$, and since η is continuous, the set of values $\eta(k\text{-NN}(x))$ almost surely converges to $\eta(x)$ in an obvious sense: for every $\varepsilon > 0$, there exists N such that

$$\forall n \geq N, \quad \eta(k\text{-NN}(x)) \subseteq (\eta(x) - \varepsilon, \eta(x) + \varepsilon), \tag{4.1}$$

where k depends on n . Let $\varepsilon > 0$ and N be fixed, and denote $P_{\varepsilon, N}$ the set of pairs (ς, x) consisting of a sample path ς and a point $x \in \Omega$ satisfying equation (4.1). Select N with the property $\mu(P_{\varepsilon, N}) > 1 - \varepsilon$. Let $\epsilon = (\epsilon_i)_{i=1}^\infty$ denote the sequence of labels for ς , which is a random variable with the joint law $\otimes_{n=1}^\infty \{\eta(x_i), 1 - \eta(x_i)\}$. By the above, whenever $(\varsigma, x) \in P_{\varepsilon, N}$ and $n \geq N$, if x_i is one of the k nearest neighbours of x in ς_n , we have $\mathbb{E}\epsilon_i = \eta(x_i) \in (\eta(x) - \varepsilon, \eta(x) + \varepsilon)$. According to a version of the Law of Large Numbers with Chernoff’s bounds, the probability of the event

$$\frac{\sum_{x_i \in k\text{-NN}(x)} \epsilon_i}{k} \notin (\eta(x) - 2\varepsilon, \eta(x) + 2\varepsilon) \tag{4.2}$$

is exponentially small, bounded above by $2 \exp(-2\varepsilon^2 k)$. Thus, when $n \geq N$, $P[|\eta_{n,k} - \eta| \geq \varepsilon] < \varepsilon + 2 \exp(-2\varepsilon^2 k)$, and we conclude. \square

Remark 4.2. In the most general case (with the uniform tie-breaking) we can only infer the almost sure convergence if $k = k_n$ grows fast enough as a function in n , for otherwise the series $\sum_{n=1}^\infty 2 \exp(-2\varepsilon^2 k_n)$ may be divergent.

4.2. Stone’s geometric lemma for \mathbb{R}^d

In the case of a general Borel regression function η , which can be discontinuous μ -almost everywhere, where μ , as before, is the sample distribution on a separable metric space, a version of the classical Luzin theorem of real analysis says that for any $\varepsilon > 0$ there is a closed precompact set K of measure $\mu(K) > 1 - \varepsilon$ upon which η is

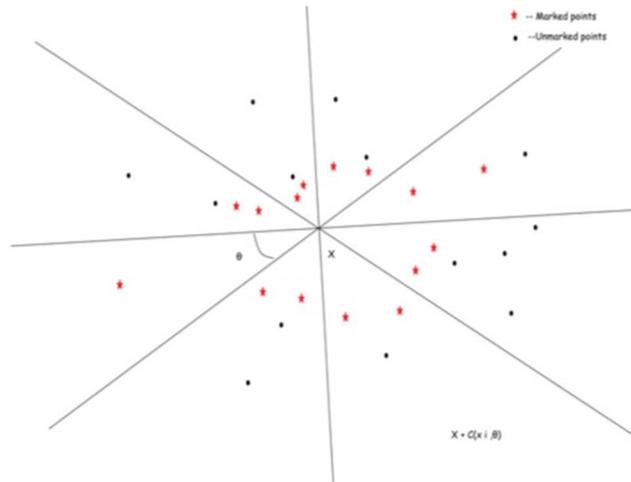


FIGURE 2. To the proof of Stone’s geometric lemma (case $k = 2$).

continuous. (See Appendix.) Now we have control over the behaviour of those k -nearest neighbours of a point x that belong to K : the mean value of the regression function η taken at those k -nearest neighbours will converge to $\eta(x)$. However, we have no control over the behaviour of the values of η at the k -nearest neighbours of x that belong to the open set $U = \Omega \setminus K$. The problem is therefore to limit the influence of the remaining $\approx \varepsilon n$ sample points belonging to U . Intuitively, as the example of Preiss shows, in infinite dimensions the influence of the few points outside of K can become overwhelming, no matter how close the measure of K is to one.

In the Euclidean case, this goal is achieved with the help of Stone’s geometric lemma, which uses the finite-dimensional Euclidean structure of the space in a beautiful way.

Lemma 4.3 (Stone’s geometric lemma for \mathbb{R}^d). *For every natural d , there is an absolute constant $C = C(d)$ with the following property. Let*

$$\sigma = (x_1, x_2, \dots, x_n), \quad x_i \in \mathbb{R}^d, \quad i = 1, 2, \dots, n,$$

be a finite sample in $\ell^2(d)$ (possibly with repetitions), and let $x \in \ell^2(d)$ be any. Given $k \in \mathbb{N}_+$, the number of i such that $x \neq x_i$ and x is among the k nearest neighbours of x_i inside the sample

$$x_1, x_2, \dots, x_{i-1}, x, x_{i+1}, \dots, x_n \tag{4.3}$$

is limited by Ck .

Proof. Cover \mathbb{R}^d with $C = C(d)$ cones of central angle $< \pi/3$ with vertices at x . Inside each cone mark the maximal possible number $\leq k$ of the nearest neighbours of x , that are set-theoretically different from x . (The strategy for possible distance tie-breaking is unimportant.) In this way, up to Ck points are marked (See Fig. 2). Let now i be any, such that $x_i \neq x$ as a point. If x_i has not been marked, this means the cone containing x_i has k points, different from x , that have been marked. Consider any of the marked points inside the same cone, say y . A simple argument of planimetry, inside an affine plane passing through x , x_i , and y , shows that

$$\|x_i - x\| > \|x_i - y\|, \tag{4.4}$$

and so the k nearest neighbours of x_i inside the sample in equation (4.3) will all be among the marked points, excluding x . □

Remark 4.4. Note that in the statement of Stone’s geometric lemma neither the order of the sample $x_1, x_2, \dots, x_{i-1}, x, x_{i+1}, \dots, x_n$, nor the tie-breaking strategy are of any importance.

Remark 4.5. If the cones have central angle $\pi/3 = 60^\circ$, then the displayed inequality in the proof of the lemma (Eq. (4.4)) is no longer strict. This is less convenient in case of distance ties.

4.3. Proof of Stone’s theorem

Theorem 4.6 (Charles J. Stone, 1977). *Let $k, n \rightarrow \infty, k/n \rightarrow 0$. Then the k -NN classification rule in the finite-dimensional Euclidean space \mathbb{R}^d is universally consistent.*

Let us begin with the case where the domain Ω is an arbitrary separable metric space, μ is a probability measure on Ω , and η is a regression function. Let $\varepsilon > 0$ be any. Using a variation of Luzin’s theorem (Thm. A.1 in Appendix), choose a closed precompact set $K \subseteq \Omega$ with $\mu(K) > 1 - \varepsilon$ and $\eta|_K$ continuous. Denote $U = \Omega \setminus K$. By the Tietze–Urysohn extension theorem (see e.g., [7], 2.1.8), the function $\eta|_K$ extends to a continuous $[0, 1]$ -valued function ψ on all of Ω .

Denote $\psi_{n,k}$ the empirical regression function (Eq. (2.2)) corresponding to ψ viewed as a regression function on its own right. We have:

$$\mathbb{E} |\eta - \eta_{n,k}| \leq \underbrace{\mathbb{E} |\eta - \psi|}_{(I)} + \underbrace{\mathbb{E} |\psi - \psi_{n,k}|}_{(II)} + \underbrace{\mathbb{E} |\psi_{n,k} - \eta_{n,k}|}_{(III)},$$

where (I) $\leq \mu(U) < \varepsilon$ and (II) $\rightarrow 0$ in probability by virtue of Lemma 4.1. It only remains to estimate the term (III).

With this purpose, let now $\Omega = \ell^2(d)$, let K be a compact subset of \mathbb{R}^d , and let $U = \Omega \setminus K$. Let X, X_1, X_2, \dots, X_n be i.i.d. random elements of \mathbb{R}^d following the law μ . We will estimate the expected number of the random elements X_i of an n -sample that (1) belong to U , and (2) are among the k nearest neighbours of a random element X belonging to K . Applying the symmetrization with a transposition of the coordinates $\tau_i: X \leftrightarrow X_i$, as well as Stone’s geometric Lemma 4.3, we obtain:

$$\begin{aligned} \mathbb{E} \frac{1}{k} \#\{i = 1, 2, \dots, n\} \mathbb{1}_{X_i \in k\text{-NN}(X)} \mathbb{1}_{X \in K} \mathbb{1}_{X_i \notin K} &= \mathbb{E} \frac{1}{k} \sum \chi_U(X_i) \mathbb{1}_{X_i \in k\text{-NN}(X)} \mathbb{1}_{X \in K} \mathbb{1}_{X_i \notin K} \\ &\leq \mathbb{E} \frac{1}{k} \sum \chi_U(X_i) \mathbb{1}_{X_i \in k\text{-NN}(X)} \mathbb{1}_{X_i \neq X} \\ &\leq \mathbb{E} \frac{1}{k} \sum \chi_U(X) \mathbb{1}_{X \in k\text{-NN}(X_1, X_2, \dots, X_{i-1}, X, X_{i+1}, \dots, X_n)} \mathbb{1}_{X_i \neq X} \\ &\leq \mathbb{E} \chi_U(X) \frac{1}{k} (kC(d)) \\ &= C(d)\mu(U). \end{aligned}$$

Getting back to the term (III), in the case $\Omega = \mathbb{R}^d$ we obtain:

$$\begin{aligned} (III) &\leq \mathbb{E} |\psi_{n,k}(X) - \eta_{n,k}(X)| \mathbb{1}_{X \in K} + \mathbb{E} |\psi_{n,k}(X) - \eta_{n,k}(X)| \mathbb{1}_{X \in U} \\ &\leq C(d)\mu(U) + \mu(U) \\ &< (C(d) + 1)\varepsilon. \end{aligned}$$

Since $\varepsilon > 0$ is as small as desired, we conclude that $\eta_{n,k}(X) \rightarrow \eta(X)$ in probability, and so the k -NN classifier in $\ell^2(d)$ is universally (weakly) consistent in $\ell^2(d)$.

5. NAGATA DIMENSION OF A METRIC SPACE

Recall that a family γ of subsets of a set Ω has multiplicity $\leq \delta$ if the intersection of more than δ different elements of γ is always empty. In other words,

$$\forall x \in \Omega, \sum_{V \in \gamma} \chi_V(x) \leq \delta.$$

Definition 5.1. Let $\delta \in \mathbb{N}$, $s \in (0, +\infty]$. We say that a metric space (Ω, d) has Nagata dimension $\leq \delta$ on the scale $s > 0$, if every finite family γ of closed balls of radii $< s$ admits a subfamily γ' of multiplicity $\leq \delta + 1$ which covers the centres of all the balls in γ . The smallest δ with this property, if it exists, and $+\infty$ otherwise, is called the Nagata dimension of Ω on the scale $s > 0$. A space Ω has Nagata dimension δ if it has Nagata dimension δ on a suitably small scale $s \in (0, +\infty]$. Notation: $\dim_{\text{Nag}}^s(\Omega) = \delta$, or simply $\dim_{\text{Nag}}(\Omega) = \delta$.

Sometimes the following reformulation is more convenient.

Proposition 5.2. A metric space (Ω, d) has Nagata dimension $\leq \delta$ on the scale $s \in (0, +\infty]$ if and only if it satisfies the following property. For every $x \in \Omega$, $r < s$, and a sequence $x_1, \dots, x_{\delta+2} \in \bar{B}_r(x)$, there are i, j , $i \neq j$, such that $d(x_i, x_j) \leq \max\{d(x, x_i), d(x, x_j)\}$.

Proof. Necessity (\Rightarrow): from the family of closed balls $B_{d(x_i, x)}(x_i)$, $i = 1, 2, \dots, \delta + 2$, all having radii $< s$, extract a family of $\delta + 1$ balls covering the centres. One of those balls, say with centre at x_i , must contain some x_j with $i \neq j$, which means $d(x_i, x_j) \leq d(x_i, x) \leq \max\{d(x, x_i), d(x, x_j)\}$.

Sufficiency (\Leftarrow): let γ be a finite family of closed balls of radii $< s$. Suppose it has multiplicity $> \delta + 1$. Then there exist a point $x \in \Omega$ and $\delta + 2$ balls in γ with centres that we denote $x_1, \dots, x_{\delta+2}$, all containing x . Denote $r = \max_i\{d(x, x_i)\}$. Then $r < s$, and by the hypothesis, there are i, j , $i \neq j$, with $d(x_i, x_j) \leq \max\{d(x, x_i), d(x, x_j)\}$. Without loss in generality, assume $d(x_i, x_j) \leq d(x_i, x)$, that is, x_j belongs to the ball with centre in x_i . Now the ball centred at x_j can be removed from the family γ , with the remaining family still covering all the centres and having the cardinality $|\gamma| - 1$. After finitely many steps, we arrive at a subfamily of multiplicity $\leq \delta + 1$ covering all the centres. \square

Example 5.3. The property of a metric space Ω having Nagata dimension zero on the scale $+\infty$ is equivalent to Ω being a non-archimedean metric space, that is, a metric space satisfying the strong triangle inequality, $d(x, z) \leq \max\{d(x, y), d(y, z)\}$.

Indeed, $\dim_{\text{Nag}}^{+\infty}(\Omega) = 0$ means exactly that for any sequence of $\delta + 2 = 2$ points, x_1, x_2 , contained in a closed ball $B_r(x)$, we have $d(x_1, x_2) \leq \max\{d(x, x_1), d(x, x_2)\}$.

Example 5.4. It follows from Proposition 5.2 that $\dim_{\text{Nag}}(\mathbb{R}) = 1$. Let x_1, x_2, x_3 be three points contained in a closed ball, that is, an interval $[x - r, x + r]$. Without loss in generality, assume $x_1 < x_2 < x_3$. If $x_2 \leq x$, then $|x_1 - x_2| \leq |x_1 - x|$, and if $x_2 \geq x$, then $|x_3 - x_2| \leq |x_3 - x|$.

The following example suggests that the Nagata dimension is relevant for the study of the k -NN classifier, as it captures in an abstract context the geometry behind Stone’s lemma.

Example 5.5. The Nagata dimension of the Euclidean space $\ell^2(d)$ is finite, and it is bounded by $C(d) - 1$, where $C(d)$ is the value of the constant in Stone’s geometric lemma (Lem. 4.3).

Indeed, let $x_1, \dots, x_{C(d)+1}$ be points belonging to a ball with centre x . Using the argument in the proof of Stone’s geometric lemma with $k = 1$, mark $\leq C(d)$ points x_i belonging to the $\leq C(d)$ cones with apex at x . At least one point, say x_j , has not been marked; it belongs to some cone, which therefore already contains a marked point, say x_i , different from x_j , and $\|x_i - x_j\| \leq \|x_j - x\|$.

Example 5.6. A similar argument shows that every finite-dimensional normed space has finite Nagata dimension.

Remark 5.7. In $\mathbb{R}^2 = \mathbb{C}$ the family of closed balls of radius one centred at the vectors $\exp(2\pi ki/5)$, $k = 1, 2, \dots, 5$, has multiplicity 5 and admits no proper subfamily containing all the centres. Therefore, the Nagata dimension of $\ell^2(2)$ is at least 5. Since the plane can be covered with 6 cones having the central angle $\pi/3$, Example 5.5 implies that $\dim_{\text{Nag}}(\ell^2(2)) = 5$.

Remark 5.8. The problem of calculating the Nagata dimension of the Euclidean space $\ell^2(d)$ is mentioned as “possibly open” by Nagata [15], page 9 (where the value $\dim_{\text{Nag}} + 2$ is called the “crowding number”). Nagata also remarks that $\dim_{\text{Nag}} \mathbb{R}^1 = 1$ and $\dim_{\text{Nag}} \ell^2(3) = 5$ (without a proof).

Remark 5.9. Notice that the property of the Euclidean space established in the proof of Stone’s geometric lemma is strictly stronger than the finiteness of the Nagata dimension. There exists a finite δ (in general, higher than the Nagata dimension) such that, given a sequence $x_1, \dots, x_{\delta+2} \in \bar{B}_r(x)$, $r < s$, there are i, j , $i \neq j$, such that $d(x_i, x_j) < \max\{d(x, x_i), d(x, x_j)\}$. The inequality here is strict, cf. Remark 4.5. This is exactly the property that removes the problem of distance ties in the Euclidean space. However, adopting this as a definition in the general case would be too restrictive, removing from consideration a large class of metric spaces in which the k -NN classifier is still universally consistent, such as all non-archimedean metric spaces.

Example 5.10. Let e_n denote the n th standard basic vector in the separable Hilbert space ℓ^2 , that is, a sequence whose n th coordinate is 1 and the rest are zeros. The convergent sequence $(1/n)e_n$, $n \geq 0$, together with the limit 0, viewed as a metric subspace of ℓ^2 , has infinite Nagata dimension on every scale $s > 0$. This is witnessed by the family of closed balls $B_{1/n}((a/n)e_n)$, having zero as the common point, and having the property that every centre belongs to exactly one ball of the family. Realizing \mathbb{R} as a continuous curve in ℓ^2 without self-intersections passing through all elements of the sequence as well as the limit leads to an equivalent metric on \mathbb{R} having infinite Nagata dimension on each scale.

Remark 5.11. The Nagata–Ostrand theorem [13, 16] states that the Lebesgue covering dimension of a metrizable topological space is the smallest Nagata dimension of a compatible metric on the space (and in fact this is true on every scale $s > 0$, [12]). This is the historical origin of the concept of the metric dimension.

There appears to be no single comprehensive reference to the concept of Nagata dimension. Various results are scattered in the journal papers [1, 12, 13, 15, 16, 18], see also the book [14], pages 151–154.

Metric spaces of finite Nagata dimension admit an almost literal version of Stone’s geometric lemma in case where the sample has no distance ties, that is, the values of the distances $d(x_i, x_j)$, $i \neq j$, are all pairwise distinct.

Lemma 5.12 (Stone’s geometric lemma, finite Nagata dimension, no ties). *Let Ω a metric space of Nagata dimension $\delta < \infty$ on a scale $s > 0$. Let*

$$\sigma = (x_1, x_2, \dots, x_n), \quad i = 1, 2, \dots, n,$$

be a finite sample in Ω , and let $x \in \Omega$ be any. Suppose there are no distance ties inside the sample

$$x, x_1, x_2, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n, \tag{5.1}$$

and k is such that, inside the above sample, $r_{k\text{-NN}}(x_i) < s$ for all i . The number of i having the property that $x \neq x_i$ and x is among the k nearest neighbours of x_i inside the sample above is limited by $(k + 1)(\delta + 1)$.

Proof. Suppose that $i = 1, 2, \dots, m$ is such that $x_i \neq x$ and x_i has x among the k nearest neighbours inside the sample as in equation (5.1). The family γ of closed balls $B_{r_{k\text{-NN}}(x_i)}(x_i)$, $i \leq m$, admits a subfamily γ' of multiplicity $\leq \delta + 1$ covering all the points x_i , $i \leq m$. Since there are no distance ties, every ball belonging to γ contains $\leq k + 1$ points. It follows that $\#\gamma' \geq m/(k + 1)$. All the balls in γ' contain x , and we conclude: $\#\gamma' \leq \delta + 1$. The result follows. \square

Now the same argument as in the original proof of Stone (Sect. 4.3) shows that the k -NN classifier is consistent under each distribution μ on $\Omega \times \{0, 1\}$ with the property that the distance ties occur with zero probability. Since we are going to give a proof of a more general result, we will not repeat the argument here, only mention that due to the Cover–Hart lemma, if n is sufficiently large, then with arbitrarily high probability, the k nearest neighbours of a random point inside a random sample will all lie at a distance $< s$.

6. DISTANCE TIES

In this section we will construct a series of examples to illustrate the difficulties arising in the presence of distance ties in general metric spaces that are absent in the Euclidean case. The fundamental difference between the two situations is the inequality in the equivalent definition of the Nagata dimension (Prop. 5.2) that is, unlike in the Euclidean space, no longer strict.

As we have already noted (Rem. 4.4), the conclusion of Stone’s geometric Lemma 4.3 remains valid even if we allow the adversary to break the distance ties and pick the k nearest neighbours. Our first example shows that it is no longer the case in a metric space of finite Nagata dimension.

Example 6.1. Consider a finite set $\sigma = \{x_1, x_2, \dots, x_n\}$ with $n \geq k$ points, and assume that in the metric space $\sigma \cup \{x\}$ all $n + 1$ points are pairwise at a distance one from each other. The Nagata dimension of the metric space $\sigma \cup \{x\}$ is equal to $\delta = 0$. Indeed, if a family γ of closed balls contains any ball of radius ≥ 1 , it already covers σ on its own. Otherwise, we choose one ball of radius < 1 (that is, a singleton) for each centre. The multiplicity of the selected subfamily is 1 in each case.

Now let us discuss the distance ties. For any element x_i of σ , the remaining n points of $\sigma \cup \{x\}$ are tied between themselves as the possible k nearest neighbours. The adversary may decide to always select x among them, thus invalidating the conclusion of Stone’s geometric lemma.

However, the problem is easily resolved if we break distance ties using a uniform distribution on the nearest neighbour candidates. In this case, the expected number of indices i such that x is chosen as one of the k nearest neighbours of x_i within the sample $\{x_1, x_2, \dots, x, \dots, x_k\}$ is obviously k .

Remark 6.2. It is worth observing that in the Euclidean case $\Omega = \mathbb{R}^d$ the size of a sample inheriting a 0–1 distance will be limited from above by the dimension, d .

The next example shows that Stone’s geometric lemma in finite dimensional metric spaces cannot be saved even with the uniform tie-breaking.

Example 6.3. There exists a countable metric space $\sigma = \{x_1, x_2, \dots\}$ of Nagata dimension 0, having the following property. Given $N \in \mathbb{N}$, for a sufficiently large n the expected number of points $x_i \neq x_1$ within the sample $\sigma_n = \{x_1, \dots, x_n\}$ having x_1 as the nearest neighbour under the uniform tie-breaking is $\geq N$.

We will construct σ recurrently. Let $\sigma_1 = \{x_1\}$. Add x_2 at a distance 1 from x_1 , and set $\sigma_2 = \{x_1, x_2\}$. If σ_n has been already defined, add x_{n+1} at a distance 2^n from all the existing points x_i , $i \leq n$, and set $\sigma_{n+1} = \sigma_n \cup \{x_{n+1}\}$. It is clear that the distance so defined is a metric.

We will verify by induction in n that $\dim_{\text{Nag}}(\sigma_n) = 0$, on the scale $s = +\infty$. For $n = 1$ this is trivially true. Assume the statement holds for σ_n , and let γ be a family of closed balls in σ_{n+1} . If one of those balls contains all the points, there is nothing to prove. Assume not, that is, all the balls elements of γ have radii smaller than 2^{n+1} . Choose a subfamily of multiplicity 1 consisting of balls centred in elements of σ_n and covering them all, and add one ball centred in x_{n+1} (which is a singleton). Now it follows that $\dim_{\text{Nag}} \sigma = 0$ as well.

Finally, let us show that if n is sufficiently large, then the expected number of indices i such that $x = x_1$ is the nearest neighbour of x_i under a uniform tie-breaking is as large as desired. With this purpose, for each $i \geq 2$ we will calculate the expectation of the event $x_1 \in NN(x_i)$, where $NN(x_i)$ denotes the set of nearest neighbours of x_i in the rest of the finite sample σ_n .

For x_2 , the unique nearest neighbour within σ_n is x_1 , therefore $\mathbb{E}[x_1 \in NN(x_2)] = 1$. For x_3 , there are two points in σ_n at a distance 2 from x_3 , which can be chosen each with probability $1/2$, namely x_1 and x_2 , therefore

$\mathbb{E}[x_1 \in NN(x_3)] = 1/2$. For arbitrary i , in a similar way, $\mathbb{E}[x_1 \in NN(x_i)] = 1/i$. We conclude:

$$\mathbb{E}[\#\{i = 1, \dots, n : x_1 \in NN(x_i)\}] = \sum_{i=1}^n \frac{1}{i},$$

and the sum of the harmonic series converges to $+\infty$ as $n \rightarrow \infty$.

Can it be that the distance ties are in some sense extremely rare? Even this expectation is unfounded.

Example 6.4. Given a value $\delta > 0$ (risk) and a sequence $n'_k \uparrow +\infty$, there exist a compact metric space of Nagata dimension zero (a Cantor space with a suitable compatible metric) equipped with a non-atomic probability measure, and a sequence $n_k \geq n'_k$, $k/n_k \rightarrow 0$, with the following property. With confidence $> 1 - \delta$, for every k , a random element X has $\geq n_k$ distance ties among its k nearest neighbours within a random n_{k+1} -sample σ .

The space Ω , just like in the Preiss example (Sect. 3), is the direct product $\prod_{k=1}^{\infty} [N_k]$ of finite discrete spaces, whose cardinalities $N_k \geq 2$ will be chosen recursively, and $[N_k] = \{1, 2, \dots, N_k\}$. The metric is given by the rule

$$d(\sigma, \tau) = \begin{cases} 0, & \text{if } \sigma = \tau, \\ 2^{-\min\{i : \sigma_i \neq \tau_i\}}, & \text{otherwise.} \end{cases}$$

This metric induces the product topology and is non-archimedean, so the Nagata dimension of Ω is zero (Ex. 5.3). The measure μ is the product of uniform measures μ_{N_k} on the spaces $[N_k]$. This measure is non-atomic, and in particular, μ -almost all distance ties occur at a strictly positive distance from a random element X .

Choose a sequence (δ_i) with $\delta_i > 0$ and $2 \sum_i \delta_i = \delta$. Choose N_1 so large that, with probability $> 1 - \delta_1$, $n_1 = n'_1$ independent random elements following a uniform distribution on the space $[N_1]$ are pairwise distinct. Now let $n_2 \geq n'_2$ be so large that with probability $> 1 - \delta_1$, if n_2 independent random elements follow a uniform distribution on $[N_1]$, then each element of $[N_1]$ appears among them at least n_1 times.

Suppose that $n_1, N_1, n_2, N_2, \dots, n_k$ have been chosen. Let N_k be so large that, with probability $> 1 - \delta_k$, n_k i.i.d. random elements uniformly distributed in $[N_k]$ are pairwise distinct. Choose $n_{k+1} \geq n'_{k+1}$ so large that, with probability $> 1 - \delta_k$, if n_{k+1} i.i.d. random elements are uniformly distributed within $\prod_{i=1}^k [N_i]$, then each element of $\prod_{i=1}^k [N_i]$ will appear among them at least n_k times.

Let k be any positive natural number. Choose $n_{k+1} + 1$ i.i.d. random elements $X, X_1, \dots, X_{n_{k+1}}$ of Ω , following the distribution μ . With probability $> 1 - 2\delta_k$, the following occurs: there are n_k elements in the sample X_1, X_2, \dots, X_{n_k} which have the same i th coordinates as X , $i = 1, 2, 3, \dots, k$, yet the $(k+1)$ -coordinates of X, X_1, \dots, X_{n_k} are all pairwise distinct. In this way, the distances between X and all those n_k elements are equal to 2^{-k-1} . We have n_k distance ties between k nearest neighbours of X (which are all at the same distance as the nearest neighbour of X), and $n_k \geq n'_k$, as desired.

Now, it would be tempting to try and reduce the general case to the case of zero probability of ties, as follows. Recall that the ℓ^1 -type direct sum of two metric spaces, (X, d_X) and (Y, d_Y) , is the direct product $X \times Y$ equipped with the coordinatewise sum of the two metrics:

$$d(x, y) = d_X(x_1, x_2) + d_Y(y_1, y_2).$$

Notation: $X \oplus_1 Y$.

Let Ω be a domain, that is, a metric space equipped with a probability measure μ and a regression function, η . Form the ℓ^1 -type direct sum $\Omega \oplus_1 [0, \varepsilon]$, and equip it with the product measure $\mu \otimes \lambda$ (where λ is the normalized Lebesgue measure on the interval) and the regression function $\eta \circ \pi_1$, where π_1 is the projection on the first coordinate. It is easy to see that the probability of distance ties in the space $\Omega \oplus_1 [0, \varepsilon]$ is zero, and every uniform distance tie breaking within a given finite sample will occur for a suitably small $\varepsilon > 0$. In this way, one could derive the consistency of the classifier by conditioning. However, we will now give an example of two

metric spaces of Nagata dimension 0 and 1 respectively, whose ℓ^1 -type sum has infinite Nagata dimension. This is again very different from what happens in the Euclidean case.

Example 6.5. Fix $\alpha > 0$. Let $\Omega = \{x_n : n \in \mathbb{N}\}$, equipped with the following distance:

$$d(x_i, x_j) = \begin{cases} 0, & \text{if } i = j, \\ \sum_{k=1}^j \alpha^k, & \text{if } i < j. \end{cases}$$

For $i < j < k$,

$$d(x_i, x_k) = \sum_{m=1}^k \alpha^m = d(x_j, x_k) > d(x_i, x_j) = \sum_{m=1}^j \alpha^m,$$

from where it follows that d is an ultrametric. Thus, Ω is a metric space of Nagata dimension 0.

The interval $\mathbb{I} = [0, 1]$ has Nagata dimension 1 (Ex. 5.4). Now let us consider the ℓ^1 -type sum $\Omega \oplus_1 \mathbb{I}$. Let $0 < \beta < \alpha < 1$, and $\beta < 1/2$. Consider the infinite sequence

$$z_i = (x_i, \beta^i) \in \Omega \oplus_1 \mathbb{I}$$

and the point

$$z = (x_0, 0).$$

Whenever $i < j$,

$$\begin{aligned} d(z_i, z_j) &= d(x_i, x_j) + \beta^i - \beta^j \\ &\geq d(x_i, x_0) + \alpha^j + \beta^i - \beta^j \\ &> d(x_i, x_0) + \beta^i \\ &= d(z_i, z), \end{aligned}$$

and also

$$\begin{aligned} d(z_i, z_j) &= d(x_i, x_j) + \beta^i - \beta^j \\ &> d(x_j, x_0) + \beta^j \\ &= d(z_j, z). \end{aligned}$$

Together, the properties imply: for all $i \neq j$,

$$d(z_i, z_j) > \max\{d(z_i, z_0), d(z_j, z_0)\}.$$

Thus, the Nagata dimension of the ℓ^1 -type sum $\Omega \oplus_1 \mathbb{I}$ is infinite.

The above examples show that beyond the Euclidean setting, we have to put up with the possibility that some points in a sample will appear disproportionately often among k nearest neighbours of other points. In data science, such points are known as “hubs” and the above (empirical) observation, as the “hubness phenomenon”, see *e.g.*, [19] and further references therein. Stone’s geometric lemma has to be generalized to allow for the possibility of a few of those “hubs”, whose number will be nevertheless limited. The lemma has to be reshaped in the spirit of the Hardy–Littlewood inequality in geometric measure theory.

To begin with, following Preiss [18], we will extend further our metric space dimension theory setting.

7. SIGMA-FINITE DIMENSIONAL METRIC SPACES

Definition 7.1. Say that a metric subspace X of a metric space Ω has *Nagata dimension* $\leq \delta \in \mathbb{N}$ on the scale $s > 0$ inside of Ω if every finite family of closed balls in Ω with centres in X admits a subfamily of multiplicity $\leq \delta + 1$ in Ω which covers all the centres of the original balls. The subspace X has a finite Nagata dimension in Ω if X has finite dimension in Ω on some scale $s > 0$. Notation: $\dim_{\text{Nag}}^s(X, \Omega)$ or sometimes simply $\dim_{\text{Nag}}(X, \Omega)$.

Following Preiss, let us call a family of balls *disconnected* if the centre of each ball does not belong to any other ball. Here is a mere reformulation of the above definition.

Proposition 7.2. *For a subspace X of a metric space Ω , one has*

$$\dim_{\text{Nag}}^s(X, \Omega) \leq \delta$$

if and only if every disconnected family of closed balls in Ω of radii $< s$ with centres in X has multiplicity $\leq \delta + 1$.

Proof. Necessity. Let γ be a disconnected finite family of closed balls in Ω with centres in X . Since by assumption $\dim_{\text{Nag}}^s(X, \Omega) \leq \delta$, the family γ admits a subfamily of multiplicity $\leq \delta + 1$ covering all the original centres. But only subfamily that contains all the centres is γ itself.

Sufficiency. Let γ be a finite family of closed balls in Ω with centres in X . Denote C the set of centres of those balls. Among all the disconnected subfamilies of γ (which exist, e.g., each family containing just one ball is such) there is one, γ' , with the maximal cardinality of the set $C \cap \cup \gamma'$. We claim that $C \subseteq \cup \gamma'$, which will finish the argument. Indeed, if it is not the case, there is a ball, $B \in \gamma$, whose centre, $c \in C$, does not belong to $\cup \gamma'$. Remove from γ' all the balls with centres in $B \cap C$ and add B instead. The new family, γ'' , is disconnected and contains $(C \cap \cup \gamma') \cup \{c\}$, which contradicts the maximality of γ' . \square

In the Definition 7.1, as well as in the Proposition 7.2, closed balls can be replaced with open ones. In fact, the statements remain valid if some balls in the families are allowed to be closed, other, open. We have the following.

Proposition 7.3. *For a subspace X of a metric space Ω , the following are equivalent.*

- (1) $\dim_{\text{Nag}}^s(X, \Omega) \leq \delta$,
- (2) every finite family of balls (some open, others closed) in Ω with centres in X and radii $< s$ admits a subfamily of multiplicity $\leq \delta + 1$ in Ω which covers all the centres of the original balls,
- (3) every finite family of open balls in Ω having radii $< s$ with centres in X admits a subfamily of multiplicity $\leq \delta + 1$ in Ω which covers all the centres of the original balls,
- (4) every disconnected family of open balls in Ω of radii $< s$ with centres in X has multiplicity $\leq \delta + 1$,
- (5) every disconnected family of balls (some open, others closed) in Ω of radii $< s$ with centres in X has multiplicity $\leq \delta + 1$.

Proof. (1) \Rightarrow (2): Let γ be a finite family of balls in Ω with centres in X , of radii $< s$, where some of the balls may be open and others, closed. For every element $B \in \gamma$ and each $k \geq 2$, form a closed ball B_k as follows: if B is closed, then $B_k = B$, and if B is open, then define B_k as having the same centre and radius $r(1 - 1/m)$, where r is the radius of B . Thus, we always have $B = \cup_{k=2}^{\infty} B_k$. Select recursively a chain of subfamilies

$$\gamma \supseteq \gamma_1 \supseteq \gamma_2 \supseteq \dots \supseteq \gamma_k \supseteq \dots$$

with the properties that for each k , the family of closed balls B_k , $B \in \gamma_k$ has multiplicity $\leq \delta + 1$ in Ω and covers all the centres of the balls in γ . Since γ is finite, starting with some k , the subfamily γ_k stabilizes, and

now it is easy to see that the subfamily γ_k itself has the desired multiplicity, and of course covers all the original centres.

(2) \Rightarrow (3): trivially true.

(3) \Rightarrow (4): same argument as in the proof of necessity in Proposition 7.2.

(4) \Rightarrow (5): let γ be a disconnected family of balls in Ω , some of which may be open and others, closed, having radii $< s$ and centred in X . For each $B \in \gamma$ and $\varepsilon > 0$, denote B_ε an open ball equal to B if B is open, and concentric with B and of the radius $r + \varepsilon$, where r is the radius of B , if B is closed. For a sufficiently small $\varepsilon > 0$, the family $\{B_\varepsilon : B \in \gamma\}$ is disconnected, and its radii are all strictly less than s , therefore this family has multiplicity $\leq \delta + 1$ by assumption. The same follows for γ .

(5) \Rightarrow (1): the condition (5) is formally even stronger than an equivalent condition for (1) established in Proposition 7.2. □

Proposition 7.4. *Let X be a subspace of a metric space Ω , satisfying $\dim_{\text{Nag}}^s(X, \Omega) \leq \delta$. Then $\dim_{\text{Nag}}^s(\bar{X}, \Omega) \leq \delta$, where \bar{X} is the closure of X in Ω .*

Proof. Let γ be a finite disconnected family of open balls in Ω of radii $< s$, centred in \bar{X} . Let $y \in \Omega$, and let γ' consist of all balls in γ containing y . Choose $\varepsilon > 0$ so small that the open ε -ball around y is contained in every element of γ' . For every open ball $B \in \gamma'$, denote y_B the centre and r_B the radius. We can also assume that $\varepsilon < r_B$ for each $B \in \gamma'$. Denote B' an open ball of radius $r_B - \varepsilon > 0$, centred at a point $x_B \in X$ satisfying $d(x_B, y_B) < \varepsilon$. Then $B' \subseteq B$, so the family $\{B' : B \in \gamma'\}$ is disconnected, and has radii $< s$. Therefore, y only belongs to $\leq \delta + 1$ balls $B', B \in \gamma'$, consequently the cardinality of γ' is bounded by $\delta + 1$. □

Proposition 7.5. *If X and Y are two subspaces of a metric space Ω , having finite Nagata dimension in Ω on the scales s_1 and s_2 respectively, then $X \cup Y$ has a finite Nagata dimension in Ω , with $\dim_{\text{Nag}}(X \cup Y, \Omega) \leq \dim_{\text{Nag}}(X, \Omega) + \dim_{\text{Nag}}(Y, \Omega)$, on the scale $\min\{s_1, s_2\}$.*

Proof. Given a finite family of balls γ in Ω of radii $< \min\{s_1, s_2\}$ centred in elements of $X \cup Y$, represent it as $\gamma = \gamma_X \cup \gamma_Y$, where the balls in γ_X are centered in X , and the balls in γ_Y are centred in Y . The rest is obvious. □

Definition 7.6. A metric space Ω is said to be *sigma-finite dimensional in the sense of Nagata* if $\Omega = \cup_{i=1}^\infty X_n$, where every subspace X_n has finite Nagata dimension in Ω on some scale $s_n > 0$ (where the scales s_n are possibly all different).

Remark 7.7. Due to Proposition 7.4, in the above definition we can assume the subspaces X_n to be closed in Ω , in particular Borel subsets.

Remark 7.8. A good reference for a great variety of metric dimensions, including the Nagata dimension, and their applications to measure differentiation theorems, is the article [1].

Now we will develop a version of Stone’s geometric lemma for general metric spaces of finite Nagata dimension.

8. FROM STONE TO HARDY–LITTLEWOOD

Lemma 8.1. *Let $\sigma = \{x_1, x_2, \dots, x_n\}$ be a finite sample in a metric space Ω , and let X be a subspace of finite Nagata dimension δ in Ω on a scale $s > 0$. Let $\alpha \in (0, 1]$ be any. Let $\sigma' \sqsubseteq \sigma$ be a sub-sample with m points. Assign to every $x_i \in \sigma$ a ball, B_i (which could be open or closed), centred at x_i , of radius $< s$. Then*

$$\#\{i = 1, 2, \dots, n : x_i \in X, \#(B_i \cap \sigma') \geq \alpha \#(B_i \cap \sigma)\} \leq \alpha^{-1}(\delta + 1)m.$$

Proof. Denote I the set of all $i = 1, 2, \dots, n$ such that

$$\#(B_i \cap \sigma') \geq \alpha \#(B_i \cap \sigma).$$

According to the assumption $\dim_{\text{Nag}}^s(\Omega) = \delta$, there exists $J \subseteq I$ such that the subfamily $\{B_i: i \in J\}$ has multiplicity $\leq \delta + 1$ and covers all the centres $x_i, i \in I$. In particular, each point of σ' belongs, at most, to $\delta + 1$ balls $B_i, i \in I$. Consequently,

$$\begin{aligned} \alpha \#I &= \alpha \#\{x_i: i \in I\} \\ &\leq \sum_{i \in I} \alpha \#(B_i \cap \sigma) \\ &\leq \sum_{i \in I} \#(B_i \cap \sigma') \\ &\leq \# \sigma' (\delta + 1) \\ &= m(\delta + 1), \end{aligned}$$

whence the conclusion follows. □

Remark 8.2. In applications of the lemma, $B_i = B_{\epsilon_{k\text{-NN}}}(x)$ (sometimes the ball will need to be open, sometimes closed, depending on the presence of distance ties).

Lemma 8.3. Let $\alpha, \alpha_1, \alpha_2 \geq 0, t_1, t_2 \in [0, 1], t_2 \leq 1 - t_1$. Assume that $\alpha_1 \leq \alpha$ and

$$t_1 \alpha_1 + (1 - t_1) \alpha_2 \leq \alpha.$$

Then

$$\frac{t_1 \alpha_1 + t_2 \alpha_2}{t_1 + t_2} \leq \alpha.$$

Proof. If $\alpha_2 \leq \alpha$, the conclusion is immediate. Otherwise, $\alpha_2 > \alpha$, and it follows that

$$t_1 \alpha_1 + t_2 \alpha_2 \leq \alpha - (1 - t_1 - t_2) \alpha_2 \leq (t_1 + t_2) \alpha.$$

□

Lemma 8.4. Let x, x_1, x_2, \dots, x_n be a finite sample (possibly with repetitions), and let $\sigma' \sqsubset \sigma$ be a subsample. Let $\alpha \geq 0$, and let B be a closed ball around x of radius $r_{k\text{-NN}}(x)$ which contains K elements of the sample,

$$\#\{i = 1, 2, \dots, n: x_i \in B\} = K.$$

Suppose that the fraction of points of σ' found in B is no more than α ,

$$\#\{i: x_i \in \sigma', x_i \in B\} \leq \alpha K,$$

and that the same holds for the corresponding open ball, B° ,

$$\#\{i: x_i \in \sigma', x_i \in B^\circ\} \leq \alpha \#\{i: x_i \in B^\circ\}.$$

Under the uniform tie-breaking of the k nearest neighbours, the expected fraction of the points of σ' found among the k nearest neighbours of x is less than or equal to α .

Proof. We apply Lemma 8.3 with α_1 and α_2 being the fractions of the points of σ' found in the closed ball B and on the sphere $S = B \setminus B^\circ$ respectively, $t_1 = \#B^\circ/\#B$, and t_2 being the fraction of the points of the sphere $S = B \setminus B^\circ$ to be chosen uniformly and randomly as the k nearest neighbours of x that are still missing in the open ball B° . Now it is enough to observe that the expected fraction of the points of σ' among the k nearest neighbours that belong to the sphere is also equal to α_2 , because they are being chosen randomly, following a uniform distribution. \square

Now we can give a promised alternative proof of the principal result along the same lines as Stone’s original proof in the finite-dimensional Euclidean case.

Theorem 8.5. *The k nearest neighbour classifier under the uniform distance tie-breaking is universally consistent in every metric space having sigma-finite Nagata dimension, when $n, k \rightarrow \infty$ and $k/n \rightarrow 0$.*

Proof. Represent $\Omega = \cup_{l=1}^\infty Y_n$, where Y_n have finite Nagata dimension in Ω . According to Proposition 7.5, we can assume that Y_n form an increasing chain, and Proposition 7.4 allows to assume that Y_n are Borel sets. Let a Borel probability measure μ and a measurable regression function η be any on Ω . Given $\epsilon > 0$, there exists l such that $\mu(Y_l) \geq 1 - \epsilon/2$. According to Luzin’s theorem (Thm. A.1), there is a closed precompact subset $K \subseteq Y_l$, such that $\eta|_K$ is continuous and $\mu(K) \geq 1 - \epsilon$. Applying the Tietze–Urysohn extension theorem ([7, 2.1.8]), we extend the function $\eta|_K$ to a continuous function g over Ω .

In the spirit of the proof of Stone’s Theorem 4.6, it is enough to limit the term

$$\begin{aligned} (B) &= \mathbb{E} \frac{1}{k} \sum_{i=1}^n |\eta(X_i) - g(X_i)| \mathbb{1}_{X_i \in k\text{-NN}(X)} \mathbb{1}_{X \in K} \mathbb{1}_{X_i \notin K} \\ &= \mathbb{E} \mathbb{E}_{j \sim \mu_{\#}} \frac{1}{k} \sum_{i \in \{0, 1, \dots, n\} \setminus \{j\}} |\eta(X_i) - g(X_i)| \mathbb{1}_{X_i \in k\text{-NN}(X_j)} \mathbb{1}_{X_j \in K} \mathbb{1}_{X_i \notin K}, \end{aligned}$$

where $\mu_{\#}$ is the uniform measure on the set $\{0, 1, 2, \dots, n\}$, and we denote $X_0 = X$.

Let $s > 0$ denote the scale on which Y_l has finite Nagata dimension, $\delta \in \mathbb{N}$. Since by the Cover–Hart lemma 2.2

$$\mathbb{E}_{j \sim \mu_{\#}} r_{k\text{-NN}}^{\{X_0, \dots, X_n\} \setminus \{X_j\}}(X_j) \rightarrow 0,$$

it suffices to estimate the term

$$(B') = \mathbb{E} \mathbb{E}_{j \sim \mu_{\#}} \frac{1}{k} \sum_{i \in \{0, 1, \dots, n\} \setminus \{j\}} |\eta(X_i) - g(X_i)| \mathbb{1}_{X_i \in k\text{-NN}(X_j)} \mathbb{1}_{X_j \in K} \mathbb{1}_{X_i \notin K} \mathbb{1}_{r_{k\text{-NN}}^{\{X_0, \dots, X_n\} \setminus \{X_j\}}(X_j) < s}.$$

We will treat the above as a sum of two expectations, (B_1) and (B_2) , according to whether the k nearest neighbours of X_j inside the sample $\{X_0, X_1, \dots, X_{j-1}, X, X_{j+1}, \dots, X_n\}$ contain more or less than $\sqrt{\epsilon}k$ elements belonging to $U = \Omega \setminus K$.

Applying Lemma 8.1 to the closed balls of radius $k\text{-NN}(X_j)$ as well as the corresponding open balls, together with Lemma 8.4, we get in the first case

$$\begin{aligned} (B_1) &= \mathbb{E} \mathbb{E}_{j \sim \mu_{\#}} \frac{1}{k} \sum |\eta(X_i) - g(X_i)| \mathbb{1}_{X_i \in k\text{-NN}(X_j)} \mathbb{1}_{X_j \in K} \mathbb{1}_{X_i \notin K}, \quad i \in \{0, 1, \dots, n\} \setminus \{j\} \mathbb{1}_{X_i \notin K} \mathbb{1}_{r_{k\text{-NN}}^{\{X_0, \dots, X_n\} \setminus \{X_j\}}(X_j) < s} \\ &\quad \times \mathbb{1}_{\#\{i: X_i \in k\text{-NN}(X_j), X_i \notin K\} \geq k\sqrt{\epsilon}} \\ &\leq \mathbb{E} \frac{1}{k} k\epsilon^{-1/2} (\delta + 1) \frac{1}{n} \#\{i = 0, 1, \dots, n: X_i \notin K\} \end{aligned}$$

$$\begin{aligned} &\leq \epsilon^{-1/2}(\delta + 1)\epsilon \\ &= \sqrt{\epsilon}(\delta + 1), \end{aligned}$$

where we have used the fact that the sum in the first line does not exceed k . In the second case,

$$\begin{aligned} (B_2) &= \mathbb{E}\mathbb{E}_{j \sim \mu_{\#}} \frac{1}{k} \sum |\eta(X_i) - g(X_i)| \mathbb{1}_{X_i \in k\text{-NN}(X_j)} \mathbb{1}_{X_j \in K} \mathbb{1}_{X_i \notin K, i \in \{0, 1, \dots, n\} \setminus \{j\}} \mathbb{1}_{r_{k\text{-NN}}^{\{X_0, \dots, X_n\} \setminus \{X_j\}}(X_j) < s} \\ &\quad \times \mathbb{1}_{\#\{i: X_i \in k\text{-NN}(X_j), X_i \notin K\} < k\sqrt{\epsilon}} \\ &\leq \frac{1}{k} k\sqrt{\epsilon} \\ &= \sqrt{\epsilon}. \end{aligned}$$

□

9. CONCLUDING REMARKS

The following question remains open. Let Ω be a separable complete metric space in which the k -NN classifier is universally consistent. Does it follow that Ω is sigma-finite dimensional in the sense of Nagata?

A positive answer would imply, modulo the results of Cérou and Guyader [2] and of Preiss [18], that a separable metric space Ω satisfies the weak Lebesgue–Besicovitch differentiation property for every Borel sigma-finite locally finite measure if and only if Ω satisfies the strong Lebesgue–Besicovitch differentiation property for every Borel sigma-finite locally finite measure, which would answer an old question asked by Preiss in [18].

Most of this investigation appears as a part of the Ph.D. thesis of one of the authors [10].

APPENDIX A. LUZIN’S THEOREM

The classical Luzin theorem admits numerous variations, of which we need the following one.

Theorem A.1 (Luzin’s theorem). *Let X be a separable metric space (not necessarily complete), μ a Borel probability measure on X , and $f: X \rightarrow \mathbb{R}$ a μ -measurable function. Then for every $\epsilon > 0$ there exists a closed precompact set $K \subseteq X$ with $\mu(K) > 1 - \epsilon$ and such that $f|_K$ is continuous.*

As we could not find an exact reference to this specific version, we are including the proof.

Theorem A.2. *Every Borel probability measure, μ , on a separable metric space Ω (not necessarily complete) satisfies the following regularity condition. Let A be a μ -measurable subset of Ω . For every $\epsilon > 0$ there exist a closed subset, F , and an open subset, U , of Ω such that*

$$F \subseteq A \subseteq U$$

and

$$\mu(U \setminus F) < \epsilon.$$

Proof. Denote \mathcal{A} the family of all Borel subsets of Ω satisfying the conclusion of the theorem: given $\epsilon > 0$, there exist a closed set, F , and an open set, U , of Ω satisfying $F \subseteq A \subseteq U$ and $\mu(U \setminus F) < \epsilon$. It is easy to see that \mathcal{A} forms a sigma-algebra which contains all closed subsets. Consequently, \mathcal{A} contains all Borel sets. Since every μ -measurable set differs from a suitable Borel set by a μ -null set, we conclude. □

Proof of Luzin's Theorem A.1. Let (ε_n) , $\varepsilon_n > 0$ be a summable sequence with $\sum_{n=1}^{\infty} \varepsilon_n = \varepsilon$. Enumerate the family of all open intervals with rational endpoints: (a_n, b_n) , $n \in \mathbb{N}$. For every n , use Theorem A.2 to select closed sets $F_n \subseteq f^{-1}(a_n, b_n)$ and F'_n , $F'_n \cap f^{-1}(a_n, b_n) = \emptyset$ so that their union $\tilde{F}_n = F_n \cup F'_n$ satisfies $\mu(\tilde{F}_n) > 1 - \varepsilon_n/2$.

The measure μ viewed as a Borel probability measure on the completion $\hat{\Omega}$ of the metric space Ω is regular, so there exists a compact set $Q \subseteq \hat{\Omega}$ with $\mu(Q) > 1 - \varepsilon/2$.

The set

$$K = \bigcap_{n=1}^{\infty} \tilde{F}_n \cap Q$$

is closed and precompact in Ω , and satisfies $\mu(K) > 1 - \varepsilon$. For each n , the set $f^{-1}(a_n, b_n) \cap K$ is relatively open in K because its complement,

$$K \setminus f^{-1}(a_n, b_n) = F'_n \cap K,$$

is closed. □

This simple proof is borrowed from [11].

Acknowledgements. The authors are thankful to the two anonymous referees of the paper, whose comments have helped to considerably improve the presentation. The remaining mistakes are of course authors' own.

REFERENCES

- [1] P. Assouad and T. Quentin de Gromard, Recouvrements, derivation des mesures et dimensions. *Rev. Mat. Iberoam.* **22** (2006) 893–953.
- [2] F. C erou and A. Guyader, Nearest neighbor classification in infinite dimension. *ESAIM: PS* **10** (2006) 340–355.
- [3] T.M. Cover and P.E. Hart, Nearest neighbour pattern classification. *IEEE Trans. Info. Theory* **13** (1967) 21–27.
- [4] L. Devroye, On the almost everywhere convergence of nonparametric regression function estimates. *Ann. Statist.* **9** (1981) 1310–1319.
- [5] Luc Devroye, L aszl  Gy orfi and G abor Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer–Verlag, New York (1996).
- [6] Hubert Haoyang Duan, *Applying Supervised Learning Algorithms and a New Feature Selection Method to Predict Coronary Artery Disease*. M.Sc. thesis, University of Ottawa. Preprint [arXiv:1402.0459](https://arxiv.org/abs/1402.0459) (2014).
- [7] R. Engelking, *General Topology*, Sigma Series in Pure Mathematics, 2nd edn. Heldermann Verlag, Berlin (1989) 6.
- [8] L. Forzani, R. Fraiman and P. Llop, Consistent nonparametric regression for functional data under the Stone–Besicovitch conditions. *IEEE Transac. Inf. Theory* **58** (2012) 6697–6708.
- [9] Stan Hatko, *k-Nearest Neighbour Classification of Datasets with a Family of Distances*. M.Sc. thesis, University of Ottawa. Preprint [arXiv:1512.00001](https://arxiv.org/abs/1512.00001) (2015).
- [10] Sushma Kumari, *Topics in Random Matrices and Statistical Machine Learning*. Ph.D. thesis, Kyoto University. Preprint [arXiv:1807.09419](https://arxiv.org/abs/1807.09419) (2018).
- [11] P.A. Loeb and E. Talvila, Luzin's theorem and Bochner integration. *Sci. Math. Japan* **60** (2004) 113–120.
- [12] J. Nagata, On a special metric characterizing a metric space of $\dim \leq n$. *Proc. Japan Acad.* **39** (1963) 278–282.
- [13] J.I. Nagata, On a special metric and dimension. *Fund. Math.* **55** (1964) 181–194.
- [14] J.-I. Nagata, *Modern dimension theory*. Bibliotheca Mathematica, Interscience Publishers, North–Holland, Amsterdam (1965) 6.
- [15] J.-I. Nagata, Open problems left in my wake of research. *Topol. Appl.* **146** (2005) 5–13.
- [16] A.P. Ostrand, A conjecture of J. Nagata on dimension and metrization. *Bull. Amer. Math. Soc.* **71** (1965) 623–625.
- [17] D. Preiss, Invalid Vitali theorems, in: *Abstracta. 7th Winter School on Abstract Analysis, Czechoslovak Academy of Sciences, Prague, Czechia* (1979) 58–60.
- [18] D. Preiss, Dimension of metrics and differentiation of measures, *General topology and its relations to modern analysis and algebra*, V (Prague, 1981) *Sigma Series of Pure Mathematics*, Heldermann, Berlin (1983) 565–568.
- [19] M. Radovanovi c, A. Nanopoulos and M. Ivanovi c, Hubs in space: popular nearest neighbors in high-dimensional data. *J. Mach. Learn. Res.* **11** (2010) 2487–2531.
- [20] C. Stone, Consistent nonparametric regression. *Ann. Stat.* **5** (1977) 595–645.
- [21] K.Q. Weinberger and L.K. Saul, Distance metric learning for large margin classification. *J. Mach. Learn. Res.* **10** (2009) 207–244.