

A Nearest Neighbor Characterization of Lebesgue Points in Metric Measure Spaces

Tommaso R. Cesari¹ and Roberto Colomboni²

¹Toulouse School of Economics (TSE), Toulouse, France
& Artificial and Natural Intelligence Toulouse Institute (ANITI)
email: tom.cesari@univ-toulouse.fr

²Istituto Italiano di Tecnologia, Genova, Italy
& Università degli Studi di Milano, Milano, Italy
email: roberto.colomboni@unimi.it

January 13, 2021

keywords: Nearest Neighbor algorithms, geometric measure theory

Abstract

The property of almost every point being a Lebesgue point has proven to be crucial for the consistency of several classification algorithms based on nearest neighbors. We characterize Lebesgue points in terms of a 1-Nearest Neighbor regression algorithm for pointwise estimation, fleshing out the role played by tie-breaking rules in the corresponding convergence problem. We then give an application of our results, proving the convergence of the risk of a large class of 1-Nearest Neighbor classification algorithms in general metric spaces where almost every point is a Lebesgue point.

1 Introduction

A point x in a metric space is a *Lebesgue point* for a function f with respect to a locally-finite measure μ if

$$\frac{1}{\mu(\bar{B}_r(x))} \int_{\bar{B}_r(x)} |f(x') - f(x)| d\mu(x') \rightarrow 0, \quad r \rightarrow 0^+$$

where $\bar{B}_r(x)$ is the closed ball of radius r centered at x . Lebesgue points are an integral generalization of continuity points. They originally found applications in Fourier analysis: Lebesgue [1905] (resp., Fatou [1906]) showed that the Fourier series of an integrable function f is Cesàro-summable (resp., non-tangentially Abel-summable) to f at all its Lebesgue points. They also find applications in

harmonic analysis [Stein and Weiss, 1971, Theorem 1.25], wavelet and spline theory [Kelly et al., 1994, Theorem 2.1 and Corollary 2.2], and are a central concept in geometric measure theory, both in \mathbb{R}^d [Evans and Gariepy, 2015, Federer, 2014, Maggi, 2012, Mattila, 1999] and in general metric spaces [Cheeger, 1999, Kinnunen and Latvala, 2002, Kinnunen et al., 2008, Björn et al., 2010]. The most famous result on Lebesgue points is probably the celebrated Lebesgue–Besicovitch differentiation theorem which states that for any Radon measure μ on a Euclidean space, μ -almost every point is a Lebesgue point for all $f \in \mathcal{L}^1_{\text{loc}}(\mu)$ (see, e.g., [Evans and Gariepy, 2015, Theorems 1.32-1.33]). Preiss [1979] showed that this result does not hold in general metric spaces and characterized those in which it does [Preiss, 1983]. These spaces include finite-dimensional Banach spaces [Loeb, 2006], locally-compact separable ultrametric spaces [Simmons, 2012, Theorem 9.1], separable Riemannian manifolds [Simmons, 2012, Theorem 9.1], and (straightforwardly) countable spaces.

The Lebesgue–Besicovitch differentiation theorem found several applications in classification, regression and density-estimation problems with Nearest Neighbor algorithms and variants thereof [Abraham et al., 2006, Biau and Devroye, 2015]. To the best of our knowledge, Devroye [1981b] was the first to show a connection between the Lebesgue–Besicovitch differentiation theorem in \mathbb{R}^d and the convergence of the risk of the Nearest Neighbor classifier in which *ties are broken lexicographically* (i.e., when ties are broken by taking the sample with smallest index). With the same tie-breaking rule, Devroye [1981a] showed that the Lebesgue–Besicovitch differentiation theorem plays a crucial role in regression problems with Nearest Neighbor and kernel algorithms. Cérou and Guyader [2006] proved that, if *ties are broken uniformly at random*, the k_m -Nearest Neighbor classifier is consistent in any Polish metric space in which the Lebesgue–Besicovitch differentiation theorem holds (see also [Forzani et al., 2012, Chaudhuri and Dasgupta, 2014]). Finally, it was shown that it is possible to combine compression techniques with 1-Nearest Neighbor classification in order to achieve consistency in essentially separable metric spaces, even if the Lebesgue–Besicovitch differentiation theorem does not hold [Hanneke et al., 2019]. Recently, Györfi and Weiss [2020] achieved the same consistency result through a simpler algorithm that does not rely on compression.

Our contributions. The main purpose of this paper is to give a characterization (Theorem 5 and Corollary 1) of Lebesgue points in terms of an \mathcal{L}^1 -convergence property of a Nearest Neighbor algorithm. More precisely, take a bounded measurable function η defined on a metric space (\mathcal{X}, d) , a random i.i.d. sample X_1, \dots, X_m on \mathcal{X} , and an arbitrary point $x \in \mathcal{X}$ in the support of the distribution of the X_k 's. The algorithm evaluates η at a point

$$X_m^x \in \operatorname{argmin}_{X' \in \{X_1, \dots, X_m\}} d(x, X')$$

with the goal of approximating $\eta(x)$.

First, we prove that, if $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , then x is a Lebesgue point for η , regardless of how ties in the definition of X_m^x are broken (Theorem 1).

Vice versa, it is known that, if ties are broken lexicographically and x is a Lebesgue point for η , then $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 [Devroye, 1981b]. By means of a novel technique (Theorem 2), we extend this result allowing more general tie-breaking rules. Under two different sufficient conditions, we show that if x is a Lebesgue point, then $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 (Theorems 3 and 4). The first one (14) is a relaxed measure-continuity condition : in this case the implication holds regardless of how ties are broken. The second one (15) bounds the bias of the tie-breaking rules in relation to the distribution of the X_k 's. In particular, we present a broad class of tie-breaking rules —which we call ISIMINs (Independent Selectors of Indices of Minimum Numbers, Definition 4)— for which this condition holds no matter how pathological the distribution of the X_k 's is (Proposition 1). At a high-level, an ISIMIN selects the smallest among finite sets of numbers, relying on an independent source to break ties. Notably, both lexicographical and random tie-breaking rules fall within this class.

Furthermore, if neither of the conditions (14) and (15) holds, we show with a counterexample (Example 2) that x being a Lebesgue point for η does *not* imply that $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , highlighting that tie-breaking rules play a role in the convergence of Nearest Neighbor algorithms.

Putting all these results together leads to our characterization of Lebesgue points x in terms of the \mathcal{L}^1 -convergence of $\eta(X_m^x)$ to $\eta(x)$, which we later extend to general measures (Theorem 6).

Moreover, the proofs of our main theorems suggest a sequential characterization of Lebesgue points, which turns out to be true for (not necessarily bounded) locally-integrable functions (Theorem 7).

Notably all of our results on Lebesgue points can also be extended to Lebesgue values (Section 4).

We then present some applications. For the broad class of tie-breaking rules defined by ISIMINs, we give a detailed proof of the convergence of the risk of the 1-Nearest Neighbor classification algorithm. This result holds in arbitrary (even non separable) metric spaces where the Lebesgue–Besicovitch differentiation theorem holds (Theorem 8). This shows in particular (Corollary 3) that the consistency of the 1-Nearest Neighbor algorithm is essentially equivalent to the realizability assumption (23). We conclude the paper with a counterexample (Example 4) showing that these convergence results do not hold (in general) without assuming the validity of the Lebesgue–Besicovitch differentiation theorem.

Outline of the paper. In Section 2 we study the relationships between x being a Lebesgue point and $\eta(X_m^x)$ converging to $\eta(x)$ in \mathcal{L}^1 in a probabilistic setting, proving our Nearest Neighbor characterization of Lebesgue points using ISIMINs. In Section 3 we extend our results to arbitrary measures, also obtaining a sequential characterization of Lebesgue points. Section 4 illustrates how our findings can be extended to Lebesgue values. In Section 5 we show an application to the convergence of the risk of binary classification with 1-Nearest Neighbor algorithms defined by ISIMINs.

2 Lebesgue vs Nearest Neighbor

In this section we study —in a probabilistic setting— the relationships between the geometric measure-theoretic concept of Lebesgue points and the \mathcal{L}^1 -convergence of a 1-Nearest Neighbor regression algorithm for pointwise estimation.

2.1 Preliminaries and Definitions

We begin by introducing our setting, notation, and definitions.

Setting 1. Fix an arbitrary metric space (\mathcal{X}, d) . Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and X, X_1, X_2, \dots a sequence of \mathcal{X} -valued \mathbb{P} -i.i.d. random variables.

For each measurable space $(\mathcal{W}, \mathcal{F}_{\mathcal{W}})$ and each random variable $W: \Omega \rightarrow \mathcal{W}$, we denote the distribution of W with respect to \mathbb{P} by $\mathbb{P}_W: \mathcal{F}_{\mathcal{W}} \rightarrow [0, 1]$, $A \mapsto \mathbb{P}(W \in A)$. In the sequel, we will use interchangeably the notations $\mathbb{P}_W(A)$ and $\mathbb{P}(W \in A)$, depending on which one is the clearest in the context. We will denote expectations with respect to \mathbb{P} by $\mathbb{E}[\cdot]$.

For any $x \in \mathcal{X}$ and each $r > 0$, we denote the open ball $\{x' \in \mathcal{X} \mid d(x, x') < r\}$ by $B_r(x)$, the closed ball $\{x' \in \mathcal{X} \mid d(x, x') \leq r\}$ by $\bar{B}_r(x)$, and the sphere $\{x' \in \mathcal{X} \mid d(x, x') = r\}$ by $S_r(x)$.

We define the *support* of \mathbb{P}_X as the set

$$\text{supp}(\mathbb{P}_X) := \left\{ x \in \mathcal{X} \mid \forall r > 0, \mathbb{P}_X(\bar{B}_r(x)) > 0 \right\}.$$

We are now ready to introduce Lebesgue points in our setting.

Definition 1 (Lebesgue point). Let $\eta: \mathcal{X} \rightarrow \mathbb{R}$ be a bounded measurable function. We say that a point $x \in \text{supp}(\mathbb{P}_X)$ is a *Lebesgue point* (for η with respect to \mathbb{P}_X) if

$$\frac{\mathbb{E}\left[\mathbb{I}_{\bar{B}_r(x)}(X) \mid \eta(X) - \eta(x) \mid\right]}{\mathbb{P}_X(\bar{B}_r(x))} \rightarrow 0, \quad \text{as } r \rightarrow 0^+.$$

Note that Lebesgue points have a natural probabilistic interpretation. Indeed, the ratio in the previous definition, which we call *Lebesgue ratio*, is simply the expectation of $|\eta(X) - \eta(x)|$ conditioned to $X \in \bar{B}_r(x)$.

The goal is to characterize Lebesgue points in terms of nearest neighbors, which we now introduce formally.

Definition 2 (Nearest neighbor). For any point $x \in \mathcal{X}$ and each $m \in \mathbb{N}$, we say that a measurable $X_m^x: \Omega \rightarrow \mathcal{X}$ is a *nearest neighbor* of x (among X_1, \dots, X_m) if, for all $\omega \in \Omega$,

$$X_m^x(\omega) \in \underset{x' \in \{X_1(\omega), \dots, X_m(\omega)\}}{\text{argmin}} d(x, x').$$

To avoid constant repetitions in our statements, we now fix some notation and the corresponding assumptions.

Assumption 1. Until the end of Section 2.2, we will assume the following:

1. $x \in \mathcal{X}$ is a point in the support of \mathbb{P}_X ;
2. $\eta : \Omega \rightarrow \mathbb{R}$ is a bounded measurable function;
3. for all $m \in \mathbb{N}$, X_m^x is a nearest neighbor among X_1, \dots, X_m of x .

For the sake of brevity, we denote $\bar{B}_r(x)$, $B_r(x)$, $S_r(x)$ simply by \bar{B}_r , B_r , S_r .

2.2 Nearest Neighbor \implies Lebesgue

In this section we show that if $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , then x is a Lebesgue point. We begin by giving a high-level overview of the proof of this result.

Note that, by definition of nearest neighbor X_m^x , for any measurable set $A \subset \mathcal{X}$,

$$\{X_m^x \in A \cap \bar{B}_r\} \supset \bigcup_{k=1}^m \left(\{X_k \in A \cap \bar{B}_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin \bar{B}_r\} \right),$$

where the key observation is that the union on the right-hand side is disjoint. Taking probabilities on both sides and integrating, one can show that

$$\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(\bar{B}_r)} \leq \frac{\mathbb{E}[|\eta(X_m^x) - \eta(x)|]}{m \mathbb{P}_X(\bar{B}_r) (1 - \mathbb{P}_X(\bar{B}_r))^{m-1}}.$$

This suggests a way to control Lebesgue ratios with the \mathcal{L}^1 -distance between $\eta(X_m^x)$ and $\eta(x)$. Tuning m in terms of r will lead to the result.

Theorem 1. *If $\mathbb{E}[|\eta(X_m^x) - \eta(x)|] \rightarrow 0$ as $m \rightarrow \infty$, then x is a Lebesgue point.*

Proof. Note that, up to a rescaling, we can (and do) assume that $|\eta - \eta(x)| \leq 1$.

If $\mathbb{P}(X = x) > 0$, then, x is a Lebesgue point by the dominated convergence theorem (with dominating function 1) and the monotonicity of the probability.

Assume then that $\mathbb{P}(X = x) = 0$. Note that, for all Borel subset A of (\mathcal{X}, d) , for all $r > 0$, and each $m \in \mathbb{N}$,

$$\{X_m^x \in A \cap \bar{B}_r\} \supset \bigcup_{k=1}^m \left(\{X_k \in A \cap \bar{B}_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin \bar{B}_r\} \right),$$

where all elements in the union are mutually disjoint, then

$$\begin{aligned} \mathbb{P}(X_m^x \in A \cap \bar{B}_r) &\geq \mathbb{P} \left(\bigcup_{k=1}^m \left(\{X_k \in A \cap \bar{B}_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin \bar{B}_r\} \right) \right) \\ &= \sum_{k=1}^m \mathbb{P} \left(\{X_k \in A \cap \bar{B}_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin \bar{B}_r\} \right) \\ &= m \mathbb{P}_X(A \cap \bar{B}_r) (1 - \mathbb{P}_X(\bar{B}_r))^{m-1} \end{aligned}$$

which in turn gives

$$\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] \geq m \mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right] (1 - \mathbb{P}_X(\bar{B}_r))^{m-1},$$

which rearranging, upper bounding $\mathbb{I}_{\bar{B}_r}(X_m^x)$ with 1, and dividing both sides by $\mathbb{P}_X(\bar{B}_r)$, yields

$$\frac{\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right]}{\mathbb{P}_X(\bar{B}_r)} \leq \frac{\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right]}{m \mathbb{P}_X(\bar{B}_r) (1 - \mathbb{P}_X(\bar{B}_r))^{m-1}}, \quad (1)$$

if $\mathbb{P}_X(\bar{B}_r) < 1$. We will show that the right hand side vanishes as m approaches ∞ . Fix any $\varepsilon > 0$. By assumption there exists $M \in \mathbb{N}$ such that, for all $m \in \mathbb{N}$, $m \geq M$,

$$\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \leq \frac{\varepsilon}{e}. \quad (2)$$

For each $m \in \mathbb{N}$, define the (smooth) auxiliary function

$$\begin{aligned} f_m &: [0, 1] \rightarrow [0, \infty), \\ t &\mapsto m t (1 - t)^{m-1}. \end{aligned}$$

By studying the sign of its first derivative, we conclude that, for each $m \in \mathbb{N}$, f_m is increasing on $[0, 1/m]$, it is decreasing on $[1/m, 1]$, and

$$\max_{t \in [0, 1]} f_m(t) = f_m\left(\frac{1}{m}\right) = \left(1 - \frac{1}{m}\right)^{m-1} > \frac{1}{e}.$$

Hence, the superlevel set $\{f_m \geq 1/e\}$ is a non-empty and closed subinterval of $[0, 1]$. For all $m \in \mathbb{N}$, we let

$$I_m := [a_m, b_m] := \{f_m \geq 1/e\}. \quad (3)$$

Note that for all $m \in \mathbb{N}$, $b_{m+1} \in I_m$, i.e., that $f_m(b_{m+1}) \geq 1/e$. Indeed, since for all $m \in \mathbb{N}$, $f_{m+1}(b_{m+1}) = 1/e$ and $b_{m+1} \geq 1/(m+1)$, we have that

$$\begin{aligned} f_m(b_{m+1}) - \frac{1}{e} &= f_m(b_{m+1}) - f_{m+1}(b_{m+1}) \\ &= m b_{m+1} (1 - b_{m+1})^{m-1} - (m+1) b_{m+1} (1 - b_{m+1})^m \\ &= b_{m+1} (1 - b_{m+1})^{m-1} ((m+1) b_{m+1} - 1) \\ &\geq b_{m+1} (1 - b_{m+1})^{m-1} \left((m+1) \frac{1}{m+1} - 1 \right) = 0. \end{aligned}$$

This implies that for all $m \in \mathbb{N}$, $a_m \leq b_{m+1} \leq b_m$, thus $I_m \cap I_{m+1} \neq \emptyset$ and $(b_m)_{m \in \mathbb{N}}$ is non-increasing. Finally, $a_m \rightarrow 0$ as $m \rightarrow \infty$, since for all $m \in \mathbb{N}$, $a_m \leq 1/m$. Hence

$$\bigcup_{m \in \mathbb{N}, m \geq M} I_m = (0, b_M].$$

Let $\delta > 0$ such that for each $r \in (0, \delta)$ we have that $\mathbb{P}_X(\bar{B}_r) \in (0, b_M]$. Thus, for each $r \in (0, \delta)$, there exists $m \in \mathbb{N}$ such that $m \geq M$ and $\mathbb{P}_X(\bar{B}_r) \in I_m$, yielding

$$\begin{aligned} \frac{\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right]}{\mathbb{P}_X(\bar{B}_r)} &\stackrel{(1)}{\leq} \frac{\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right]}{m \mathbb{P}_X(\bar{B}_r) (1 - \mathbb{P}_X(\bar{B}_r))^{m-1}} \\ &= \frac{\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right]}{f_m(\mathbb{P}_X(\bar{B}_r))} \stackrel{(3)}{\leq} e \mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \stackrel{(2)}{\leq} \varepsilon. \end{aligned}$$

Being ε arbitrary, we conclude that x is a Lebesgue point. \square

2.3 Lebesgue \implies Nearest Neighbor (sometimes)

In this section we will assume that x is a Lebesgue point and study when this implies that $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 .

We begin by addressing two trivial cases. The first one is when x is an atom for \mathbb{P}_X . In this case, the result is trivialized by the fact that X_m^x becomes eventually equal to x (almost surely). The second one is when $\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right] = 0$ for some $r > 0$. In this case, since x belongs to the support of \mathbb{P}_X , the result is trivialized by the fact that X_m^x will eventually fall inside \bar{B}_r (almost surely). These ideas are made rigorous in the following lemma.

Lemma 1. *If one of the two following conditions is satisfied:*

1. $\mathbb{P}(X = x) > 0$;
2. *there exists $r > 0$ such that $\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right] = 0$;*

then $\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \rightarrow 0$ as $m \rightarrow \infty$.

Proof. If condition 1 is satisfied, we have that

$$\begin{aligned} \mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] &= \mathbb{E}\left[\mathbb{I}_{\mathcal{X} \setminus \{x\}}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] \leq 2 \|\eta\|_\infty \mathbb{P}(X_m^x \neq x) \\ &= 2 \|\eta\|_\infty \mathbb{P}\left(\bigcap_{k=1}^m \{X_k \neq x\}\right) = 2 \|\eta\|_\infty \prod_{k=1}^m \mathbb{P}(X_k \neq x) \\ &= 2 \|\eta\|_\infty (1 - \mathbb{P}(X = x))^m \rightarrow 0, \quad \text{as } m \rightarrow \infty. \end{aligned}$$

Assume now that condition 2 is satisfied. Then, being $\mathbb{P}_{X_m^x}$ absolutely continuous with respect to \mathbb{P}_X , it follows that $\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] = 0$. Since $x \in \text{supp}(\mathbb{P}_X)$, we have that $\mathbb{P}_X(\bar{B}_r) > 0$, which in turn gives

$$\begin{aligned} \mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] &= \mathbb{E}\left[\mathbb{I}_{\bar{B}_r^c}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] \leq 2 \|\eta\|_\infty \mathbb{P}(X_m^x \notin \bar{B}_r) \\ &= 2 \|\eta\|_\infty \mathbb{P}\left(\bigcap_{k=1}^m \{X_k \notin \bar{B}_r\}\right) = 2 \|\eta\|_\infty \prod_{k=1}^m \mathbb{P}(X_k \notin \bar{B}_r) \\ &= 2 \|\eta\|_\infty (1 - \mathbb{P}_X(\bar{B}_r))^m \rightarrow 0, \quad \text{as } m \rightarrow \infty. \quad \square \end{aligned}$$

By the previous lemma, without loss of generality, we can (and do) assume that none of the two previous conditions hold.

Assumption 2. Until the end of this section, we will assume the following:

1. $\mathbb{P}(X = x) = 0$;
2. $\mathbb{E}[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|] > 0$, for all $r > 0$.

We now proceed to estimate the expectation $\mathbb{E}[|\eta(X_m^x) - \eta(x)|]$. Due to the nature of nearest neighbors, one could figure that $\mathbb{E}[|\eta(X_m^x) - \eta(x)|]$ behaves differently if X_m^x is close by, or far away from x . This idea leads to the splitting of the expectation $\mathbb{E}[|\eta(X_m^x) - \eta(x)|]$ on the region in which X_m^x belongs to a closed ball \bar{B}_r and its complement, i.e. on $\{X_m^x \in \bar{B}_r\}$ and $\{X_m^x \in \bar{B}_r^c\}$. Since ties might raise issues on spheres, we further split $\{X_m^x \in \bar{B}_r\}$ into $\{X_m^x \in B_r\}$ and $\{X_m^x \in S_r\}$.

The next lemma gives estimates of the three terms determined by this splitting.

Lemma 2. For all $r > 0$ and all $m \in \mathbb{N}$,

$$\mathbb{E}[\mathbb{I}_{B_r}(X_m^x) |\eta(X_m^x) - \eta(x)|] \leq m \mathbb{E}[\mathbb{I}_{B_r}(X) |\eta(X) - \eta(x)|], \quad (4)$$

$$\mathbb{E}[\mathbb{I}_{S_r}(X_m^x) |\eta(X_m^x) - \eta(x)|] \leq 2 \|\eta\|_\infty m \mathbb{P}_X(S_r) \exp(-(m-1) \mathbb{P}_X(B_r)), \quad (5)$$

$$\mathbb{E}[\mathbb{I}_{\bar{B}_r^c}(X_m^x) |\eta(X_m^x) - \eta(x)|] \leq 2 \|\eta\|_\infty \exp(-m \mathbb{P}_X(\bar{B}_r)). \quad (6)$$

Proof. Fix any $r > 0$ and $m \in \mathbb{N}$.

We begin by proving inequality (4). For all Borel sets A of (\mathcal{X}, d) , if $X_m^x \in A$, then at least one of the X_i 's belongs to A , i.e.,

$$\{X_i \in A\} \subset \bigcup_{k=1}^m \{X_k \in A\}.$$

This yields

$$\mathbb{P}(X_m^x \in A) \leq \mathbb{P}\left(\bigcup_{k=1}^m \{X_k \in A\}\right) \leq \sum_{k=1}^m \mathbb{P}(X_k \in A) = m \mathbb{P}(X \in A),$$

hence $\mathbb{P}_{X_m^x} \leq m \mathbb{P}_X$, which in turn gives, for any measurable function $f: \mathcal{X} \rightarrow [0, \infty]$, that $\mathbb{E}[f(X_m^x)] \leq m \mathbb{E}[f(X)]$. Then

$$\mathbb{E}[\mathbb{I}_{B_r}(X_m^x) |\eta(X_m^x) - \eta(x)|] \leq m \mathbb{E}[\mathbb{I}_{B_r}(X) |\eta(X) - \eta(x)|].$$

This proves (4).

We now prove inequality (5). Note that if $X_m^x \in S_r$ then at least one of the X_k 's belongs to S_r , while the others can't fall in B_r (and vice versa), i.e.,

$$\{X_m^x \in S_r\} = \bigcup_{k=1}^m \left(\{X_k \in S_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin B_r\} \right).$$

This yields

$$\begin{aligned}
\mathbb{E}\left[\mathbb{I}_{S_r}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] &\leq 2 \|\eta\|_\infty \mathbb{P}(X_m^x \in S_r) \\
&\leq 2 \|\eta\|_\infty \mathbb{P}\left(\bigcup_{k=1}^m \left(\{X_k \in S_r\} \cap \bigcap_{i=1, i \neq k}^m \{X_i \notin B_r\}\right)\right) \\
&\leq 2 \|\eta\|_\infty \sum_{k=1}^m \left(\mathbb{P}(X_k \in S_r) \prod_{i=1, i \neq k}^m \mathbb{P}(X_i \notin B_r)\right) \\
&= 2 \|\eta\|_\infty m \mathbb{P}_X(S_r) (1 - \mathbb{P}_X(B_r))^{m-1} \\
&\leq 2 \|\eta\|_\infty m \mathbb{P}_X(S_r) \exp(-(m-1) \mathbb{P}_X(B_r)) .
\end{aligned}$$

This proves (5).

Finally, we prove inequality (6). Note that $X_m^x \notin \bar{B}_r$ is equivalent to the fact that none of the X_k 's belong to \bar{B}_r , i.e.,

$$\{X_m^x \notin \bar{B}_r\} = \bigcap_{k=1}^m \{X_k \notin \bar{B}_r\} ,$$

then

$$\begin{aligned}
\mathbb{E}\left[\mathbb{I}_{\bar{B}_r^c}(X_m^x) |\eta(X_m^x) - \eta(x)|\right] &\leq 2 \|\eta\|_\infty \mathbb{P}(X_m^x \notin \bar{B}_r) \\
&= 2 \|\eta\|_\infty \mathbb{P}\left(\bigcap_{k=1}^m \{X_k \notin \bar{B}_r\}\right) = 2 \|\eta\|_\infty \prod_{k=1}^m \mathbb{P}(X_k \notin \bar{B}_r) \\
&= 2 \|\eta\|_\infty (1 - \mathbb{P}_X(\bar{B}_r))^m \leq 2 \|\eta\|_\infty \exp(-m \mathbb{P}_X(\bar{B}_r)) .
\end{aligned}$$

This concludes the proof. \square

We now give a high-level overview of the ideas used to prove the \mathcal{L}^1 -convergence of $\eta(X_m^x)$ to $\eta(x)$. For the sake of simplicity, assume for now that the cumulative function of $d(X, x)$ is continuous around 0. This measure continuity-condition is equivalent to:

$$\exists R > 0, \forall r \in (0, R), \mathbb{P}_X(S_r) = 0 . \quad (7)$$

In this case, the upper bound in (5) is always 0. Therefore, Lemma 2 implies, for each $m \in \mathbb{N}$ and all $r \in (0, R)$, that

$$\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \leq m \mathbb{E}\left[\mathbb{I}_{B_r}(X) |\eta(X) - \eta(x)|\right] + 2 \|\eta\|_\infty \exp(-m \mathbb{P}_X(\bar{B}_r)) . \quad (8)$$

This bound might seem pointless (under Assumption 2) since for all *fixed* $r > 0$, the first term on the right hand side diverges as m approaches infinity. The idea is then to pick a sequence of radii $r = r_m$ that vanishes in a way that this first term goes to zero. This is easily achievable, e.g. by selecting $(r_m)_{m \in \mathbb{N}}$ that decreases to zero very quickly. However, in this case it is now the second term that may

not vanish, since $m\mathbb{P}_X(\bar{B}_{r_m})$ may not diverge to ∞ . The key is then to find a trade-off between the two competing terms. This would be achieved if one could pick a sequence $(r_m)_{m \in \mathbb{N}}$ so that

$$m \sqrt{\mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}(x)}(X) |\eta(X) - \eta(x)|]} \sqrt{\mathbb{P}_X(\bar{B}_{r_m})} = 1. \quad (9)$$

Indeed, in this case, inequality (8) together with the measure-continuity condition (7) yields

$$\begin{aligned} \mathbb{E} \left[|\eta(X_m^x) - \eta(x)| \right] &\leq \left(\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(\bar{B}_{r_m})} \right)^{1/2} \\ &\quad + 2 \|\eta\|_\infty \exp \left(- \left(\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(\bar{B}_{r_m})} \right)^{-1/2} \right) \end{aligned}$$

which vanishes if x is a Lebesgue point.

Under the current assumptions, one can indeed show the existence of a sequence $(r_m)_{m \in \mathbb{N}}$ satisfying (9). However, if the measure-continuity condition (7) does not hold, this might no longer be the case. In order to address more general cases, we introduce the following definition.

Definition 3 (α -sequence). Fix any $\alpha \in (0, 1)$. For all $r > 0$, define

$$M_\alpha(r) := \left(\mathbb{E}[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|] \right)^\alpha \left(\mathbb{P}_X(\bar{B}_r) \right)^{1-\alpha}.$$

Define $m_1 := \lceil 1/M_\alpha(1) \rceil$. For all $m \in \mathbb{N}$ such that $m \geq m_1$, define

$$r_m := \sup \left\{ r > 0 \mid M_\alpha(r) < \frac{1}{m} \right\}.$$

We say that $(r_m)_{m \in \mathbb{N}, m \geq m_1}$ is the α -sequence (for η with respect to \mathbb{P}_X) at x .

The following lemma states several useful properties of α -sequences.

Lemma 3. Fix any $\alpha \in (0, 1)$. Then, the α -sequence $(r_m)_{m \in \mathbb{N}, m \geq m_1}$ is a well-defined vanishing sequence of strictly positive numbers. Moreover, for all $m \in \mathbb{N}$ with $m \geq m_1$, we have

$$m \leq \frac{1}{\mathbb{E}[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)|]} \left(\frac{\mathbb{E}[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(B_{r_m})} \right)^{1-\alpha} \quad (10)$$

$$m \geq \frac{1}{\mathbb{P}_X(\bar{B}_{r_m})} \left(\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(\bar{B}_{r_m})} \right)^{-\alpha} \quad (11)$$

where we stress that the balls B_{r_m} in (10) are open.

Proof. Note that the function $r \mapsto M_\alpha(r)$ is non-decreasing, right-continuous (by the continuity from above of finite measures), and for each $R > 0$, it satisfies

$$M_\alpha(r) \uparrow \left(\mathbb{E} \left[\mathbb{I}_{B_R}(X) |\eta(X) - \eta(x)| \right] \right)^\alpha \left(\mathbb{P}_X(B_R) \right)^{1-\alpha}, \quad \text{as } r \uparrow R \quad (12)$$

(by the continuity from below of measures), where we stress that the balls B_R in the previous formula are open. Furthermore, by Assumption 2 and the fact that x belongs to the support of \mathbb{P}_X , we have that $0 < M_\alpha(r) \downarrow 0$ as $r \downarrow 0$. This implies that the α -sequence $(r_m)_{m \in \mathbb{N}, m \geq m_1}$ is well-defined, strictly positive, and $r_m \downarrow 0$ as $m \uparrow \infty$.

Thus, for each $m \in \mathbb{N}$ such that $m \geq m_1$, if $r \in (0, r_m)$, then $M_\alpha(r) < 1/m$, which in turn yields

$$\lim_{r \uparrow r_m} M_\alpha(r) \leq 1/m.$$

Using (12) and rearranging gives (10).

Analogously, for each $m \in \mathbb{N}$ such that $m \geq m_1$, if $r > r_m$, then $M(r) \geq 1/m$, which in turn yields

$$M(r_m) = \lim_{r \downarrow r_m} M(r) \geq 1/m.$$

Rearranging gives (11) and completes the proof. \square

Before proceeding with the main results of the section, we need one simple (geometric measure theory flavored) lemma.

Lemma 4. *If x is a Lebesgue point and $(r_m)_{m \in \mathbb{N}}$ is a vanishing sequence of strictly positive real numbers, then*

$$\frac{\mathbb{E} \left[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(B_{r_m})} \rightarrow 0, \quad m \rightarrow \infty \quad (13)$$

where we stress that the balls B_{r_m} in the previous formula are open.

Proof. By the continuity from below of measures, we have that for all $m \in \mathbb{N}$, there exists $\rho_m > 0$ such that $|r_m - \rho_m| \leq 1/m$ and

$$\left| \frac{\mathbb{E} \left[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(B_{r_m})} - \frac{\mathbb{E} \left[\mathbb{I}_{\bar{B}_{\rho_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(\bar{B}_{\rho_m})} \right| \leq \frac{1}{m}.$$

Since $r_m \rightarrow 0^+$ as $m \rightarrow \infty$, we have that also $\rho_m \rightarrow 0^+$, as $m \rightarrow \infty$. Being x a Lebesgue point, it follows that

$$\frac{\mathbb{E} \left[\mathbb{I}_{\bar{B}_{\rho_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(\bar{B}_{\rho_m})} \rightarrow 0, \quad \text{as } m \rightarrow \infty,$$

which in turn implies, for $m \rightarrow \infty$,

$$\frac{\mathbb{E} \left[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(B_{r_m})} \leq \frac{1}{m} + \frac{\mathbb{E} \left[\mathbb{I}_{\bar{B}_{\rho_m}}(X) |\eta(X) - \eta(x)| \right]}{\mathbb{P}_X(\bar{B}_{\rho_m})} \rightarrow 0.$$

\square

The following result showcases the usefulness of α -sequences.

Theorem 2. *Let $(r_m)_{m \in \mathbb{N}, m \geq m_1}$ be an α -sequence for some $\alpha \in (0, 1)$. If x is a Lebesgue point, then the following are equivalent:*

1. $\mathbb{E}[|\eta(X_m^x) - \eta(x)|] \rightarrow 0$, as $m \rightarrow \infty$;
2. $\mathbb{E}[\mathbb{I}_{S_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \rightarrow 0$, as $m \rightarrow \infty$.

Proof. We prove the non-trivial implication 2 \Rightarrow 1. Recall that $r_m \rightarrow 0^+$ as $m \rightarrow \infty$ by Lemma 3. For each $m \in \mathbb{N}$, if $m \geq m_1$, we have

$$\begin{aligned} \mathbb{E}[|\eta(X_m^x) - \eta(x)|] &= \mathbb{E}[\mathbb{I}_{B_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \\ &\quad + \mathbb{E}[\mathbb{I}_{S_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \\ &\quad + \mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}^c}(X_m^x) |\eta(X_m^x) - \eta(x)|] \\ &=: \text{(I)} + \text{(II)} + \text{(III)}. \end{aligned}$$

We will show that all three terms above approach 0 as $m \rightarrow \infty$.

To show that the first vanishes, we apply inequality (4) in Lemma 2, inequality (10) in Lemma 3, and Lemma 4:

$$\begin{aligned} \text{(I)} &= \mathbb{E}[\mathbb{I}_{B_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \stackrel{(4)}{\leq} m \mathbb{E}[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)|] \\ &\stackrel{(10)}{\leq} \left(\frac{\mathbb{E}[\mathbb{I}_{B_{r_m}}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(B_{r_m})} \right)^{1-\alpha} \stackrel{(13)}{\rightarrow} 0, \quad \text{as } m \rightarrow \infty. \end{aligned}$$

The second term vanishes by assumption.

To show that the third term vanishes, we apply inequality (6) in Lemma 2 and inequality (11) in Lemma 3:

$$\begin{aligned} \text{(III)} &= \mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}^c}(X_m^x) |\eta(X_m^x) - \eta(x)|] \stackrel{(6)}{\leq} 2 \|\eta\|_\infty \exp(-m \mathbb{P}_X(\bar{B}_{r_m})) \\ &\stackrel{(11)}{\leq} 2 \|\eta\|_\infty \exp\left(-\left(\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_{r_m}^c}(X) |\eta(X) - \eta(x)|]}{\mathbb{P}_X(\bar{B}_{r_m})}\right)^{-\alpha}\right) \rightarrow 0, \end{aligned}$$

as $m \rightarrow \infty$. □

Thanks to the previous theorem, to prove the \mathcal{L}^1 -convergence of $\eta(X_m^x)$ to $\eta(x)$, we only need to control the expectation on spheres along an α -sequence. Recall the measure-continuity condition (7). By the previous theorem, under this condition, if x is a Lebesgue point, then automatically $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 . Theorem 3 will show that the same thing holds under the following weaker condition:

$$\exists K \geq 0, \exists R > 0, \forall r \in (0, R), \mathbb{P}_X(S_r) \leq K \mathbb{P}_X(B_r). \quad (14)$$

To give some intuition on condition (14), consider the cumulative F of the random variable $d(x, X)$ that evaluates the distance between x and X . For all $r > 0$, we have $F(r) = \mathbb{P}(d(x, X) \leq r) = \mathbb{P}_X(\bar{B}_r)$. Then, condition (14) can be restated as

$$\exists K \geq 0, \exists R > 0, \forall r \in (0, R), \frac{F(r)}{F(r^-)} \leq K + 1,$$

where $F(r^-) := \lim_{\rho \rightarrow r^-} F(\rho) = \mathbb{P}_X(B_r)$. This is a relaxation on the continuity of F in a neighbor of 0—which is precisely the case $K = 0$, corresponding to the measure-continuity condition (7)—and it allows infinite discontinuity points around 0.

In order to prove Theorem 3 as well as the last theorem of the section, we will need the following technical lemma.

Lemma 5. *Let $(m_j)_{j \in \mathbb{N}}$ be a strictly monotone sequence of natural numbers and $(\rho_j)_{j \in \mathbb{N}}$ a sequence of strictly positive real numbers. If both of the following conditions hold:*

1. *there exists $K \geq 0$ such that, for all $j \in \mathbb{N}$, $\mathbb{P}_X(S_{\rho_j}) \leq K \mathbb{P}_X(B_{\rho_j})$;*
2. *$m_j \mathbb{P}_X(\bar{B}_{\rho_j}) \rightarrow \infty$ as $j \rightarrow \infty$;*

then

$$\mathbb{E} \left[\mathbb{I}_{S_{\rho_j}}(X_{m_j}^x) |\eta(X_{m_j}^x) - \eta(x)| \right] \rightarrow 0, \quad \text{as } j \rightarrow \infty.$$

Proof. We show first that condition 2 still holds if closed balls \bar{B}_{ρ_j} are replaced by open balls B_{ρ_j} . Using conditions 1 and 2 yields the strengthened convergence condition:

$$m_j \mathbb{P}_X(B_{\rho_j}) \geq \frac{1}{K+1} \left(m_j \mathbb{P}_X(S_{\rho_j}) + m_j \mathbb{P}_X(B_{\rho_j}) \right) = \frac{1}{K+1} m_j \mathbb{P}_X(\bar{B}_{\rho_j}) \rightarrow \infty,$$

as $j \rightarrow \infty$. Hence, recalling inequality (5), using condition 1, and applying the strengthened convergence condition, we have

$$\begin{aligned} \mathbb{E} \left[\mathbb{I}_{S_{\rho_j}}(X_{m_j}^x) |\eta(X_{m_j}^x) - \eta(x)| \right] &\leq 2 \|\eta\|_\infty m_j \mathbb{P}_X(S_{\rho_j}) \exp\left(-m_j \mathbb{P}_X(B_{\rho_j})\right) \\ &\leq 2K \|\eta\|_\infty m_j \mathbb{P}_X(B_{\rho_j}) \exp\left(-m_j \mathbb{P}_X(B_{\rho_j})\right) \rightarrow 0, \quad \text{as } j \rightarrow \infty. \quad \square \end{aligned}$$

We can now prove one of our main results which shows that, under the weakened measure-continuity assumption (14), if x is a Lebesgue point, then $\eta(X_m^x)$ converges to x in \mathcal{L}^1 .

Theorem 3. *If condition (14) is satisfied and x is a Lebesgue point, then*

$$\mathbb{E} [|\eta(X_m^x) - \eta(x)|] \rightarrow 0, \quad \text{as } m \rightarrow \infty.$$

Proof. Take any α -sequence $(r_m)_{m \in \mathbb{N}, m \geq m_1}$, for some $\alpha \in (0, 1)$. By Theorem 2, it suffices to prove that $\mathbb{E} [\mathbb{I}_{S_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \rightarrow 0$, as $m \rightarrow \infty$. Recall that $0 < r_m \rightarrow 0$ as $m \rightarrow \infty$ (Lemma 3). By assumption (14) and Lemma 5 it is then sufficient to prove that $m \mathbb{P}_X(\bar{B}_{r_m}) \rightarrow \infty$ as $m \rightarrow \infty$. Being x a Lebesgue point, this follows immediately by (11). \square

The previous result shows that there is a large class of distributions \mathbb{P}_X for which $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , no matter how ties are broken in the definition of the nearest neighbor X_m^x .

Vice versa, we will show there exists a large class of tie-breaking rules for which $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , no matter how pathological \mathbb{P}_X is. This is a consequence of Theorem 4, together with the forthcoming Proposition 1. The theorem relies on the following condition, whose purpose is to bound the bias of the tie-breaking rule (for a concrete example, see Example 1):

$$\begin{aligned} \exists C > 0, \exists R > 0, \exists M > 0, \forall r \in (0, R), \forall m \in \{M, M+1, \dots\}, (\mathbb{P}_X(S_r) > 0) \\ \implies \mathbb{E}[|\eta(X_m^x) - \eta(x)| \mid X_m^x \in S_r] \leq C \mathbb{E}[|\eta(X) - \eta(x)| \mid X \in S_r]. \end{aligned} \quad (15)$$

Note that the previous expression is well-defined since the condition $\mathbb{P}_X(S_r) > 0$ is equivalent to $\mathbb{P}_{X_m^x}(S_r) > 0$, as shown in the following lemma.

Lemma 6. *Let $r > 0$. Then the following are equivalent:*

1. $\mathbb{P}_X(S_r) > 0$;
2. there exists $m \in \mathbb{N}$ such that $\mathbb{P}_{X_m^x}(S_r) > 0$;
3. for all $m \in \mathbb{N}$, we have that $\mathbb{P}_{X_m^x}(S_r) > 0$.

Proof. The results follow by the fact that, for all $m \in \mathbb{N}$, we have

$$\bigcup_{k=1}^m \{X_k \in S_r\} \supset \{X_m^x \in S_r\} \supset \bigcap_{k=1}^m \{X_k \in S_r\},$$

which in turn implies

$$m \mathbb{P}_X(S_r) \geq \mathbb{P}_{X_m^x}(S_r) \geq (\mathbb{P}_X(S_r))^m. \quad \square$$

We can now prove the last theorem of this section.

Theorem 4. *If condition (15) is satisfied and x is a Lebesgue point, then*

$$\mathbb{E}[|\eta(X_m^x) - \eta(x)|] \rightarrow 0, \quad \text{as } m \rightarrow \infty.$$

Proof. Up to a rescaling, we can (and do) assume that $|\eta - \eta(x)| \leq 1$. Take any α -sequence $(r_m)_{m \in \mathbb{N}, m \geq m_1}$, for some $\alpha \in (0, 1)$. By Theorem 2, it suffices to show that $\mathbb{E}[\mathbb{I}_{S_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|] \rightarrow 0$, as $m \rightarrow \infty$. We will prove that this is the case by showing that each subsequence of $(\mathbb{E}[\mathbb{I}_{S_{r_m}}(X_m^x) |\eta(X_m^x) - \eta(x)|])_{m \in \mathbb{N}, m \geq m_1}$ has another subsequence that converges to zero. Take any subsequence $(m_l)_{l \in \mathbb{N}}$ of $(m)_{m \in \mathbb{N}, m \geq m_1}$. If there exists $L \in \mathbb{N}$ such that, for all $l \in \mathbb{N}$, if $l \geq L$, the following identity holds:

$$\mathbb{E}[\mathbb{I}_{S_{r_{m_l}}}(X_{m_l}^x) |\eta(X_{m_l}^x) - \eta(x)|] = 0,$$

then the claim is trivially true. Up to taking another subsequence, then, we can assume that, for all $l \in \mathbb{N}$,

$$\mathbb{E}[\mathbb{I}_{S_{r_{m_l}}}(X_{m_l}^x) |\eta(X_{m_l}^x) - \eta(x)|] > 0,$$

which in turn gives

$$\mathbb{P}(X_{m_l}^x \in S_{r_{m_l}}) \geq \mathbb{E} \left[\mathbb{I}_{S_{r_{m_l}}}(X_{m_l}^x) |\eta(X_{m_l}^x) - \eta(x)| \right] > 0.$$

Moreover, recalling that $0 < r_m \rightarrow 0$ as $m \rightarrow \infty$ (by Lemma 3), we have that there exists $M_1 \in \mathbb{N}$ such that, for all $l \in \mathbb{N}$, if $l \geq M_1$, then $r_{m_l} < R$. Then, by Lemma 6 and condition (15), we have, for each $l \in \mathbb{N}$, if $l \geq \max\{M, M_1\}$,

$$\begin{aligned} \mathbb{E} \left[\mathbb{I}_{S_{r_{m_l}}}(X_{m_l}^x) |\eta(X_{m_l}^x) - \eta(x)| \right] &= \mathbb{E} \left[|\eta(X_{m_l}^x) - \eta(x)| \mid X_{m_l}^x \in S_{r_{m_l}} \right] \mathbb{P}_{X_{m_l}^x}(S_{r_{m_l}}) \\ &\stackrel{(15)}{\leq} C \mathbb{E} \left[|\eta(X) - \eta(x)| \mid X \in S_{r_{m_l}} \right] \mathbb{P}_{X_{m_l}^x}(S_{r_{m_l}}) \\ &\leq C \mathbb{E} \left[|\eta(X) - \eta(x)| \mid X \in S_{r_{m_l}} \right] \end{aligned}$$

and thus, if $\mathbb{E} \left[|\eta(X) - \eta(x)| \mid X \in S_{r_{m_l}} \right] \rightarrow 0$ as $l \rightarrow \infty$, the theorem is proven. Otherwise, up to taking another subsequence, we can (and do) assume that there exists $\varepsilon > 0$ such that, for all $l \in \mathbb{N}$,

$$\mathbb{E} \left[|\eta(X) - \eta(x)| \mid X \in S_{r_{m_l}} \right] \geq \varepsilon.$$

Thus, for each $l \in \mathbb{N}$, we get

$$\begin{aligned} 0 < \varepsilon \mathbb{P}_X(S_{r_{m_l}}) &\leq \mathbb{E} \left[|\eta(X) - \eta(x)| \mid X \in S_{r_{m_l}} \right] \mathbb{P}_X(S_{r_{m_l}}) \\ &= \mathbb{E} \left[\mathbb{I}_{S_{r_{m_l}}}(X) |\eta(X) - \eta(x)| \right]. \end{aligned}$$

Equivalently, for any $l \in \mathbb{N}$,

$$\frac{1}{\mathbb{P}_X(S_{r_{m_l}})} \geq \varepsilon \frac{1}{\mathbb{E} \left[\mathbb{I}_{S_{r_{m_l}}}(X) |\eta(X) - \eta(x)| \right]}$$

which in turn implies

$$\begin{aligned} \frac{\mathbb{P}_X(B_{r_{m_l}})}{\mathbb{P}_X(S_{r_{m_l}})} + 1 &= \frac{\mathbb{P}_X(B_{r_{m_l}}) + \mathbb{P}_X(S_{r_{m_l}})}{\mathbb{P}_X(S_{r_{m_l}})} = \frac{\mathbb{P}_X(\bar{B}_{r_{m_l}})}{\mathbb{P}_X(S_{r_{m_l}})} \\ &\geq \varepsilon \frac{\mathbb{P}_X(\bar{B}_{r_{m_l}})}{\mathbb{E} \left[\mathbb{I}_{S_{r_{m_l}}}(X) |\eta(X) - \eta(x)| \right]} \geq \varepsilon \frac{\mathbb{P}_X(\bar{B}_{r_{m_l}})}{\mathbb{E} \left[\mathbb{I}_{\bar{B}_{r_{m_l}}}(X) |\eta(X) - \eta(x)| \right]} \end{aligned}$$

and the last term diverges to infinity as $l \rightarrow \infty$ because x is a Lebesgue point and $0 < r_{m_l} \rightarrow 0$ as $l \rightarrow \infty$. Therefore, if $l \rightarrow \infty$

$$\frac{\mathbb{P}_X(B_{r_{m_l}})}{\mathbb{P}_X(S_{r_{m_l}})} \rightarrow \infty, \quad \text{or equivalently} \quad \frac{\mathbb{P}_X(S_{r_{m_l}})}{\mathbb{P}_X(B_{r_{m_l}})} \rightarrow 0.$$

Hence, there exists $K \geq 0$ such that, for all $l \in \mathbb{N}$, $\mathbb{P}_X(S_{r_{m_l}}) \leq K \mathbb{P}_X(B_{r_{m_l}})$. Furthermore, being $(r_{m_l})_{l \in \mathbb{N}}$ a subsequence of an α -sequence and x a Lebesgue point, we have that $m_l \mathbb{P}_X(\bar{B}_{r_{m_l}}) \rightarrow \infty$ as $l \rightarrow \infty$ by (11). Then, the assumptions of Lemma 5 are satisfied, implying that $\mathbb{E}[\mathbb{I}_{S_{r_{m_l}}}(X_{m_l}^x) |\eta(X_{m_l}^x) - \eta(x)|] \rightarrow 0$ as $l \rightarrow \infty$ and concluding the proof. \square

We now illustrate with a concrete example that Assumption (15) (hence Theorem 4) can still hold even if the nearest neighbor rule is arbitrarily (but finitely) biased.

Example 1. Let $\mathcal{X} := \{-1, -\frac{1}{2}, -\frac{1}{3}, \dots\} \cup \{0\} \cup \{\dots, \frac{1}{3}, \frac{1}{2}, 1\}$, d be the distance induced on \mathcal{X} by the Euclidean metric, \mathcal{B} be the Borel σ -algebra of (\mathcal{X}, d) , and $\mu: \mathcal{B} \rightarrow [0, 1]$ be the unique probability measure such that, for all $n \in \mathbb{N}$,

$$\mu\left(\left\{-\frac{1}{n}\right\}\right) = \frac{1}{R} \frac{n-1}{n} \frac{1}{2^{2^n}}, \quad \mu\left(\left\{\frac{1}{n}\right\}\right) = \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}}, \quad \text{where } R := \sum_{n \in \mathbb{N}} \frac{1}{2^{2^n}}.$$

Fix $C > 1$ and $(p_n)_{n \in \mathbb{N}} \subset (0, 1)$ such that, for all $n \in \mathbb{N}$, $p_n \leq (C-1)/n$. We define $\Omega = \mathcal{X} \times \mathcal{X}^{\mathbb{N}} \times \{0, 1\}^{\mathbb{N}}$, $\mathcal{F} = \mathcal{B} \otimes \bigotimes_{k \in \mathbb{N}} \mathcal{B} \otimes \bigotimes_{h \in \mathbb{N}} 2^{\{0, 1\}}$, and \mathbb{P} as the unique probability measure on \mathcal{F} such that for all $B, B_1, B_2, \dots \in \mathcal{B}$ and $\vartheta_1, \vartheta_2, \dots \in \{0, 1\}$,

$$\mathbb{P}(B \times B_1 \times B_2 \times \dots \times \{\vartheta_1\} \times \{\vartheta_2\} \times \dots) = \mu(B) \left(\prod_{i=1}^{+\infty} \mu(B_i) \right) \prod_{n=1}^{+\infty} p_n^{\vartheta_n} (1-p_n)^{1-\vartheta_n}.$$

It is well-known that such a probability measure exists [Halmos, 2013, Chapter VII, Section 38, Theorem B]. We define $X: \Omega \rightarrow \mathcal{X}$, $(x, x_1, x_2, \dots, \vartheta_1, \vartheta_2, \dots) \mapsto x$, and for all $k \in \mathbb{N}$, $X_k: \Omega \rightarrow \mathcal{X}$, $(x, x_1, x_2, \dots, \vartheta_1, \vartheta_2, \dots) \mapsto x_k$ and $\Theta_k: \Omega \rightarrow \{0, 1\}$, $(x, x_1, x_2, \dots, \vartheta_1, \vartheta_2, \dots) \mapsto \vartheta_k$. By definition, $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space. By construction, X, X_1, X_2, \dots are \mathbb{P} -independent random variables with common distribution $\mathbb{P}_X = \mu$, and they are \mathbb{P} -independent of $\Theta_1, \Theta_2, \dots$, which are Bernoulli random variables with parameters p_1, p_2, \dots respectively. Let

$$A := \mathcal{X} \cap (0, \infty), \quad \eta := \mathbb{I}_A, \quad \text{and } x := 0.$$

Finally, for all $m \in \mathbb{N}$, let X_m^x the nearest neighbor such that, whenever $\{-\frac{1}{n}, \frac{1}{n}\} = \operatorname{argmin}_{X' \in \{X_1, \dots, X_m\}} X'$ for some $n \in \mathbb{N}$, then $X_m^x = \frac{1}{n}$ if and only if $\Theta_n = 1$.

At a high-level:

- given that $\operatorname{argmin}_{X' \in \{X_1, \dots, X_m\}} X' = \{-\frac{1}{n}, \frac{1}{n}\}$ for some $n \in \mathbb{N}$, by definition of μ , the odds of hitting $\frac{1}{n}$ against hitting $-\frac{1}{n}$ are approximately $1 : n$, a fact that a fair tie-breaking rule should reflect; however, if we choose (for some $c > 0$ and sufficiently large n 's) $p_n \geq c/n$, the odds become at least $c : n$; i.e., we can artificially increase by at least c -times the odds of picking positive (over negative) values, with c large (if C is also large);
- (15) holds because $p_n \leq (C-1)/n$ poses a bound on the bias of the tie-breaking rule;

- (14) does not hold because the measure $\mathbb{P}_X(S_{1/n})$ of spheres $S_{1/n}$ decreases super-exponentially as $n \rightarrow \infty$;
- x is a Lebesgue point for η because the measure $\mathbb{P}_X(\bar{B}_{1/n})$ is more and more biased towards negative values as n approaches ∞ ;

We begin by proving rigorously that x is a Lebesgue point. For each $r > 0$, defining $n_r := \min\{n \in \mathbb{N} \mid 1/n \leq r\}$, we have that

$$\begin{aligned} \frac{\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right]}{\mathbb{P}_X(\bar{B}_r)} &= \frac{\mathbb{P}_X(A \cap \bar{B}_r)}{\mathbb{P}_X(\bar{B}_r)} = \frac{\sum_{k \geq n_r} \mathbb{P}(X = \frac{1}{k})}{\sum_{k \geq n_r} \mathbb{P}(X = \pm \frac{1}{k})} \\ &= \frac{\sum_{k \geq n_r} \frac{1}{k} \mathbb{P}(X = \pm \frac{1}{k})}{\sum_{k \geq n_r} \mathbb{P}(X = \pm \frac{1}{k})} \leq \frac{\sum_{k \geq n_r} \frac{1}{n_r} \mathbb{P}(X = \pm \frac{1}{k})}{\sum_{k \geq n_r} \mathbb{P}(X = \pm \frac{1}{k})} = \frac{1}{n_r} \rightarrow 0, \quad \text{as } r \rightarrow 0^+, \end{aligned}$$

hence x is a Lebesgue point.

While it is immediate that (14) does not hold, to show rigorously that (15) does requires some care. Without loss of generality, we can (and do) assume that $r = 1/n$, for some $n \in \mathbb{N}$. Let $m, n \in \mathbb{N}$ and $E := \operatorname{argmin}_{X' \in \{X_1, \dots, X_m\}} X' = \{-\frac{1}{n}, \frac{1}{n}\}$. Since

$$\begin{aligned} \mathbb{P}\left(E = \left\{\frac{1}{n}\right\}\right) &= \sum_{\emptyset \neq I \subset \{1, \dots, m\}} \left(\frac{1}{n} \mathbb{P}\left(|X| = \frac{1}{n}\right)\right)^{|I|} \left(\mathbb{P}\left(|X| > \frac{1}{n}\right)\right)^{m-|I|}, \\ \mathbb{P}\left(E \subset \left\{-\frac{1}{n}, \frac{1}{n}\right\}\right) &= \sum_{\emptyset \neq I \subset \{1, \dots, m\}} \left(\mathbb{P}\left(|X| = \frac{1}{n}\right)\right)^{|I|} \left(\mathbb{P}\left(|X| > \frac{1}{n}\right)\right)^{m-|I|}, \end{aligned}$$

we have that

$$\begin{aligned} \mathbb{E}\left[|\eta(X_m^x) - \eta(x)| \mid X_m^x \in S_{1/n}\right] &= \frac{\mathbb{P}\left(\left(E = \left\{\frac{1}{n}\right\}\right) \cup \left(E = \left\{-\frac{1}{n}, \frac{1}{n}\right\} \cap \Theta_n = 1\right)\right)}{\mathbb{P}\left(E \subset \left\{-\frac{1}{n}, \frac{1}{n}\right\}\right)} \\ &\leq \frac{\mathbb{P}\left(E = \left\{\frac{1}{n}\right\}\right)}{\mathbb{P}\left(E \subset \left\{-\frac{1}{n}, \frac{1}{n}\right\}\right)} + \frac{C-1}{n} \leq \frac{1}{n} + \frac{C-1}{n} = C \mathbb{E}\left[|\eta(X) - \eta(x)| \mid X \in S_{1/n}\right] \end{aligned}$$

and so, by Theorem 15, we know that $\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \rightarrow 0$ as $m \rightarrow \infty$.

In Theorems 3 and 4 we proved that if conditions (14) or (15) are satisfied and x is a Lebesgue point, then $\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \rightarrow 0$, as $m \rightarrow \infty$. The reader might be wondering if this implication holds with no assumptions other than that x is a Lebesgue point and X_m^x is a nearest neighbor (among X_1, \dots, X_m), i.e., if the result holds (in general) when neither condition (14) nor condition (15) are satisfied. A modification of the previous example shows that this is not the case.

Example 2. Consider the same setting as in Example 1, where we now define, for all $m \in \mathbb{N}$, X_m^x as the nearest neighbor such that, whenever $\{-\frac{1}{n}, \frac{1}{n}\} \subset \operatorname{argmin}_{X' \in \{X_1, \dots, X_m\}} X'$ for some $n \in \mathbb{N}$, then $X_m^x = \frac{1}{n}$. At a high-level:

- in contrast to the nearest neighbor defined in Example 1, here X_m^x is “infinitely” biased towards positive values, breaking ties always in their favor.
- as in Example 1, x is a Lebesgue point for η because the measure $\mathbb{P}_X(\bar{B}_{1/n})$ is more and more biased towards negative values as n approaches ∞ ;
- $\eta(X_m^x)$ does not converge to $\eta(x)$ in \mathcal{L}^1 because of the interplay between the tie-breaking rule being biased towards the direction of positive values and the measure $\mathbb{P}_X(S_{1/n})$ of spheres $S_{1/n}$ decreasing super-exponentially as $n \rightarrow \infty$.

The same computation as in Example 1 shows that x is a Lebesgue point. We show now that $\eta(X_m^x)$ does not converge to $\eta(x)$ in \mathcal{L}^1 . First, for all $m \in \mathbb{N}$, we have

$$\begin{aligned}
\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] &= \mathbb{E}\left[\mathbb{I}_A(X_m^x)\right] = \mathbb{P}(X_m^x \in A) = \sum_{n \in \mathbb{N}} \mathbb{P}\left(X_m^x = \frac{1}{n}\right) \\
&\geq \sum_{n \in \mathbb{N}} \mathbb{P}\left(\bigcup_{k=1}^m \left\{X_k = \frac{1}{n}\right\} \cap \bigcap_{\substack{i=1 \\ i \neq k}}^m \left(\left\{|X_i| > \frac{1}{n}\right\} \cup \left\{X_i = -\frac{1}{n}\right\}\right)\right) \\
&= m \sum_{n \in \mathbb{N}} \mathbb{P}\left(X = \frac{1}{n}\right) \left(\mathbb{P}\left(|X| > \frac{1}{n}\right) + \mathbb{P}\left(X = -\frac{1}{n}\right)\right)^{m-1} \\
&= m \sum_{n \in \mathbb{N}} \mathbb{P}\left(X = \frac{1}{n}\right) \left(\sum_{k=1}^{n-1} \mathbb{P}\left(X = \pm \frac{1}{k}\right) + \mathbb{P}\left(X = -\frac{1}{n}\right)\right)^{m-1} \\
&= m \sum_{n \in \mathbb{N}} \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}} \left(\sum_{k=1}^n \frac{1}{R} \frac{1}{2^{2^k}} - \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}}\right)^{m-1} \\
&= m \sum_{n \in \mathbb{N}} \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}} \left(\left(1 - \sum_{k=n+1}^{\infty} \frac{1}{R} \frac{1}{2^{2^k}}\right) - \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}}\right)^{m-1} \\
&\geq m \sum_{n \in \mathbb{N}} \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}} \left(1 - \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}}\right)^{m-1}.
\end{aligned}$$

This inequality will be used to prove that for countably many $m \in \mathbb{N}$, the quantity $\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right]$ is bounded away from zero. To do so, for all $j \in \mathbb{N}$, we fix an $m_j \in \mathbb{N}$ such that

$$\frac{1}{2m_j} \leq \frac{1}{R} \frac{1}{j} \frac{1}{2^{2^j}} \leq \frac{2}{m_j}.$$

Note that $m_j \rightarrow \infty$ as $j \rightarrow \infty$. Then, for all $j \in \mathbb{N}$, we have

$$\begin{aligned} \mathbb{E}\left[|\eta(X_{m_j}^x) - \eta(x)|\right] &\geq m_j \sum_{n \in \mathbb{N}} \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}} \left(1 - \frac{1}{R} \frac{1}{n} \frac{1}{2^{2^n}}\right)^{m_j-1} \\ &> m_j \frac{1}{R} \frac{1}{j} \frac{1}{2^{2^j}} \left(1 - \frac{1}{R} \frac{1}{j} \frac{1}{2^{2^j}}\right)^{m_j-1} \\ &\geq m_j \frac{1}{2m_j} \left(1 - \frac{2}{m_j}\right)^{m_j-1} \\ &= \frac{1}{2} \left(1 - \frac{2}{m_j}\right)^{m_j-1} \rightarrow \frac{1}{2e^2} \quad \text{as } j \rightarrow \infty. \end{aligned}$$

This argument shows that there exist countably many $m \in \mathbb{N}$ such that the inequality $\mathbb{E}\left[|\eta(X_m^x) - \eta(x)|\right] \geq 1/(4e^2)$ holds, which in turn implies that $\eta(X_m^x)$ does *not* converge to $\eta(x)$ in \mathcal{L}^1 .

2.4 A broad class of tie-breaking rules

In this section we introduce a broad class of tie-breaking rules for which condition (15) is always satisfied, regardless of how pathological \mathbb{P}_X is. To do so, we introduce Independent Selectors of Indices of Minimum Numbers (or ISIMINs for short).

Definition 4 (ISIMIN). We say that a sequence of pairs $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an *Independent Selector of Indices of Minimum Numbers (ISIMIN)*¹ if, for all $m \in \mathbb{N}$, there exists a measurable space $(\mathcal{Z}_m, \mathcal{G}_m)$ such that $\Theta_m: \Omega \rightarrow \mathcal{Z}_m$ is a random variable \mathbb{P} -independent from X_1, \dots, X_m and $\psi_m: [0, \infty)^m \times \mathcal{Z}_m \rightarrow \{1, \dots, m\}$ is a measurable function satisfying, for all $r_1, \dots, r_m \geq 0$ and all $z \in \mathcal{Z}_m$,

$$\psi_m(r_1, \dots, r_m, z) \in \underset{k \in \{1, \dots, m\}}{\operatorname{argmin}} r_k.$$

At a high-level, the function ψ_m selects the smallest among m numbers, relying on an independent source Θ_m to break ties.

One of the classical ways to break ties in Nearest Neighbor algorithms [Devroye, 1981a] is lexicographically. More precisely, ties are broken by selecting the smallest index among all indices minimizing the distance from x , i.e.,

$$\min \left(\underset{k \in \{1, \dots, m\}}{\operatorname{argmin}} d(x, X_k) \right). \quad (16)$$

We can easily represent such a selection through an ISIMIN by taking, for all $m \in \mathbb{N}$, $\mathcal{Z}_m := \{0\}$, $\mathcal{G}_m := \{\{0\}, \emptyset\}$, $\Theta_m \equiv 0$, and for all $r_1, \dots, r_m \geq 0$,

$$\psi_m(r_1, \dots, r_m, 0) := \min \left(\underset{k \in \{1, \dots, m\}}{\operatorname{argmin}} r_k \right).$$

¹ISIMIN is pronounced “easy-min”.

We prove that such a sequence $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an ISIMIN. With a slight abuse of notation, for all $m \in \mathbb{N}$ and all $i \in \{1, \dots, m\}$, we denote the coordinate map $(r_1, \dots, r_m, z) \mapsto r_i$ simply by r_i . Then, for each $m \in \mathbb{N}$ and all $k \in \{1, \dots, m\}$, we have that $\{\psi_m = k\}$ is measurable, being the Cartesian product of a measurable set (intersection of open and closed sets) with \mathcal{Z}_m :

$$\{\psi_m = k\} = \left(\bigcap_{j=1}^{k-1} \{r_j > r_k\} \cap \bigcap_{j=k+1}^m \{r_j \geq r_k\} \right) \times \mathcal{Z}_m .$$

Therefore ψ_m is measurable. Moreover, for each $m \in \mathbb{N}$, the constant function $\Theta_m \equiv 0$ is trivially \mathbb{P} -independent of X_1, \dots, X_m . Thus, the sequence $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an ISIMIN. With this choice, we can reproduce the deterministic tie-breaking rule (16) by taking, for each $m \in \mathbb{N}$,

$$\min \left(\operatorname{argmin}_{k \in \{1, \dots, m\}} d(x, X_k) \right) = \psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m) .$$

Another classical way to break ties is picking one of the closest points uniformly at random [C erou and Guyader, 2006]. More precisely, one draws a number ϑ_k in $[0, 1]$ uniformly at random, and independently of everything else, for each X_k . If there exist multiple $X_k \in \{X_1, \dots, X_m\}$ with $d(x, X_k) = \min_{j \in \{1, \dots, m\}} d(x, X_j)$, then ties are broken by picking the smallest X_k with the smallest value of ϑ_k . In the zero-probability event in which multiple closest X_k have the same (smallest) value of ϑ_k , one of them is chosen arbitrarily, e.g., lexicographically. This selection can also be represented by an ISIMIN. Indeed, for all $m \in \mathbb{N}$, let $\mathcal{Z}_m := [0, 1]^m$, \mathcal{G}_m be the Borel σ -algebra of \mathcal{Z}_m , and we assume there exist m uniform random variables $\vartheta_1, \dots, \vartheta_m: \Omega \rightarrow [0, 1]$ independent of X_1, \dots, X_m and also of each others (if they do not exist, one can simply enlarge the original probability space to accommodate them with a construction analogous to the one we present in Section 3.1). Then for all $m \in \mathbb{N}$ we take $\Theta_m := (\vartheta_1, \dots, \vartheta_m)$ and for all $r_1, \dots, r_m \geq 0, t_1, \dots, t_m \in [0, 1]$,

$$\psi_m(r_1, \dots, r_m, t_1, \dots, t_m) := \min \left(\begin{array}{c} \operatorname{argmin} \\ j \in \operatorname{argmin}_{k \in \{1, \dots, m\}} r_k \end{array} t_j \right) .$$

We prove that such a sequence $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an ISIMIN. With a slight abuse of notation, for all $m \in \mathbb{N}$ and all $i \in \{1, \dots, m\}$, we denote the coordinate map $(r_1, \dots, r_m, t_1, \dots, t_m) \mapsto r_i$ simply by r_i , and similarly, the coordinate map $(r_1, \dots, r_m, t_1, \dots, t_m) \mapsto t_i$ simply by t_i . Then, for each integer $m \in \mathbb{N}$ and any $k \in \{1, \dots, m\}$ we have that $\{\psi_m = k\}$ is measurable, being the union of an intersection of open and closed sets (in the following formula, think of A as the set of indices that tie with k):

$$\bigcup_{\substack{A \subset \{1, \dots, m\} \\ A \not\ni k}} \left(\bigcap_{\substack{j \in \{1, \dots, m\} \\ j \notin (\{k\} \cup A)}} \{r_k < r_j\} \cap \bigcap_{j \in A} \{r_k = r_j\} \cap \bigcap_{\substack{j \in A \\ j < k}} \{t_k < t_j\} \cap \bigcap_{\substack{j \in A \\ j > k}} \{t_k \leq t_j\} \right) ,$$

where we recall that the intersection over an empty set of indices is the universe $[0, \infty)^m \times \mathcal{Z}_m$. Therefore, for all $m \in \mathbb{N}$, the function ψ_m is measurable. Moreover, for each $m \in \mathbb{N}$, the function $\Theta_m = (\vartheta_1, \dots, \vartheta_m)$ is \mathbb{P} -independent of X_1, \dots, X_m by construction. Thus, the sequence $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is a ISIMIN. With this choice, we can reproduce the random tie-breaking rule in [C erou and Guyader, 2006] by taking, for each $m \in \mathbb{N}$,

$$\psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m) .$$

We now define nearest neighbors according to arbitrary ISIMINs.

Definition 5 (Nearest neighbors according to an ISIMIN). Let $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ be an ISIMIN. For each $m \in \mathbb{N}$, the *nearest neighbor* of x (among X_1, \dots, X_m), according to (ψ_m, Θ_m) is defined by

$$X_m^x := X_{\psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m)} .$$

Note that, being $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ a sequence of measurable pairs, nearest neighbors defined according to ISIMINs are also measurable, i.e., they are actually nearest neighbors according to Definition 2. The next proposition shows that these nearest neighbors always satisfy an even stronger condition than (15), regardless of pathological \mathbb{P}_X is. The intuition behind it is that the distribution of X_m^x is obtained by “pulling” the distribution of X towards x . Thus, if we prevent the “pull” towards x by conditioning X_m^x to have a constant distance from x , the distributions of X_m^x and X might coincide (at least if X_m^x does not break ties with directional preferences, since, as we saw in Example 2, directional preferences might distort the distribution of X_m^x badly). This is indeed the case for nearest neighbors defined according to ISIMINs. Indeed, ISIMINs hide directions, basing decisions only on distances and (independent random choices of) indices.

Proposition 1. *Let $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ be an ISIMIN. Assume that, for all $m \in \mathbb{N}$, the nearest neighbor X_m^x (among X_1, \dots, X_m) is defined according to (ψ_m, Θ_m) . If $r > 0$ is such that $\mathbb{P}_X(S_r) > 0$, then, for any Borel subset A of (\mathcal{X}, d) , it holds that*

$$\mathbb{P}(X_m^x \in A \mid X_m^x \in S_r) = \mathbb{P}(X \in A \mid X \in S_r) .$$

This, in particular, implies that the condition (15) holds.

Proof. Let $r > 0$ be such that $\mathbb{P}_X(S_r) > 0$. Fix any $m \in \mathbb{N}$. Recall that, by Lemma 6, we have that $\mathbb{P}_{X_m^x}(S_r) > 0$. Then, both conditional probabilities are well-defined.

To simplify the notation, we define the auxiliary functions

$$R_1 := d(x, X_1), \dots, R_m := d(x, X_m)$$

and, for all $i, j \in \{1, \dots, m\}$, we let $R_{i:j} := R_i, R_{i+1}, \dots, R_j$ with the understanding that $R_{i:j}$ does not appear if $j < i$. With this notation, we have

$$X_m^x = X_{\psi_m(R_{1:m}, \Theta_m)} .$$

For each $k \in \{1, \dots, m\}$, note that $R_1, \dots, R_{k-1}, X_k, R_{k+1}, \dots, R_m, \Theta_m$ are \mathbb{P} -independent random variables. Then for each Borel set A of (\mathcal{X}, d) we have that

$$\begin{aligned}
\mathbb{P}(X_m^x \in A \cap S_r) &= \mathbb{P}(X_{\psi_m(R_{1:m}, \Theta_m)} \in A \cap S_r) \\
&= \sum_{k=1}^m \mathbb{P}\left(\{X_{\psi_m(R_{1:m}, \Theta_m)} \in A \cap S_r\} \cap \{\psi_m(R_{1:m}, \Theta_m) = k\}\right) \\
&= \sum_{k=1}^m \mathbb{P}\left(\{X_k \in A \cap S_r\} \cap \{\psi_m(R_{1:k-1}, r, R_{k+1:m}, \Theta_m) = k\}\right) \\
&= \sum_{k=1}^m \mathbb{P}(X_k \in A \cap S_r) \mathbb{P}(\psi_m(R_{1:k-1}, r, R_{k+1:m}, \Theta_m) = k) \\
&= \mathbb{P}_X(A | S_r) \sum_{k=1}^m \mathbb{P}(R_k = r) \mathbb{P}(\psi_m(R_{1:k-1}, r, R_{k+1:m}, \Theta_m) = k) \\
&= \mathbb{P}_X(A | S_r) \sum_{k=1}^m \mathbb{P}\left(\{\psi_m(R_{1:k-1}, r, R_{k+1:m}, \Theta_m) = k\} \cap \{R_k = r\}\right) \\
&= \mathbb{P}_X(A | S_r) \sum_{k=1}^m \mathbb{P}\left(\{\psi_m(R_{1:m}, \Theta_m) = k\} \cap \{X_{\psi_m(R_{1:m}, \Theta_m)} \in S_r\}\right) \\
&= \mathbb{P}_X(A | S_r) \mathbb{P}(X_{\psi_m(R_{1:m}, \Theta_m)} \in S_r) = \mathbb{P}_X(A | S_r) \mathbb{P}(X_m^x \in S_r) ,
\end{aligned}$$

which in turn gives

$$\mathbb{P}(X_m^x \in A | X_m^x \in S_r) = \frac{\mathbb{P}(X_m^x \in A \cap S_r)}{\mathbb{P}(X_m^x \in S_r)} = \mathbb{P}_X(A | S_r) .$$

Being r and m arbitrary, integrating we get the second part of the result. \square

2.5 Lebesgue \iff Nearest Neighbor

In this section we collect some of the results we presented above in a theorem and a corollary. The theorem gives a characterization of Lebesgue points in term of the \mathcal{L}^1 -convergence of $\eta(X_m^x)$ to $\eta(x)$ if some mild conditions are satisfied. The corollary gives a concrete setting in which this characterization holds.

Theorem 5. *If condition (14) or condition (15) hold, the following are equivalent:*

1. x is a Lebesgue point;
2. $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , as $m \rightarrow \infty$.

Proof. The implication $2 \Rightarrow 1$ follows directly from Theorem 1. The vice versa is a restatement of Theorems 3 and 4. \square

Recall that —while the implication $2 \Rightarrow 1$ is always true— if none of the conditions (14) and (15) are satisfied, the implication $1 \Rightarrow 2$ does not hold in general (Example 2).

We now state the aforementioned concrete version of the previous result.

Corollary 1. *If $(X_m^x)_{m \in \mathbb{N}}$ is defined according to an ISIMIN (Definitions 5 and 4), the following are equivalent:*

1. x is a Lebesgue point;
2. $\eta(X_m^x)$ converges to $\eta(x)$ in \mathcal{L}^1 , as $m \rightarrow \infty$.

Proof. The result follows immediately by Theorem 5 and Proposition 1. \square

As we previously noted in Section 2.4, the two common tie-breaking rules in [Devroye, 1981b] (ties are broken lexicographically) and [C erou and Guyader, 2006] (ties are broken uniformly at random) are instances of ISIMINs. Therefore, Corollary 1 applies in particular to both cases.

Moreover, we remark that the lexicographic version of the 1-Nearest Neighbor algorithm can be formulated as a (memory/computationally-efficient) online algorithm (Algorithm 1).

Algorithm 1: Online Nearest Neighbor

Input: $x \in \mathcal{X}$
Initialization: let $R \leftarrow \infty$ and $Y \leftarrow 0$
1 **for** $m = 1, 2, \dots$ **do**
2 observe X_m
3 **if** $d(x, X_m) < R$ **then**
4 observe $\eta(X_m)$
5 let $R \leftarrow d(x, X_m)$ and $Y \leftarrow \eta(X_m)$

In particular, this gives an *online* characterization of Lebesgue points: x is a Lebesgue point if and only if $Y \rightarrow \eta(x)$ in \mathcal{L}^1 .

3 Beyond probability measures

In this section we present some consequences and extensions of the results we proved in Section 2, shifting the focus on the characterization of Lebesgue points defined in terms of general measures. We begin by fixing some notation and definitions that will be used throughout Sections 3.1 and 3.2.

Let (\mathcal{X}, d) be a metric space and $x \in \mathcal{X}$ an arbitrary point. As per previous sections, we will denote for all $r > 0$, the balls $\bar{B}_r(x)$ and $B_r(x)$ by \bar{B}_r and B_r . Let \mathcal{B} be the Borel σ -algebra of (\mathcal{X}, d) , $\eta: \mathcal{X} \rightarrow \mathbb{R}$ a measurable function, and $\mu: \mathcal{B} \rightarrow [0, \infty]$ a (Borel) measure.

To avoid constant repetitions in our results, we now explicitly state our (mild) assumptions.

Assumption 3. Until the end of Section 3, we will assume the following:

1. x is in the *support* of μ , i.e., for all $r > 0$, we have $\mu(\bar{B}_r) > 0$;
2. μ is *locally-finite* at x , i.e., there exists an $R > 0$ such that $\mu(\bar{B}_R) < \infty$.

From this point to the end of Section 3, we fix such an R and we remark that all our results hold for any other $R' \in (0, R)$.

We now state the general definition of Lebesgue points.

Definition 6 (Lebesgue point). We say that x is a Lebesgue point (for η with respect to μ) if

$$\frac{1}{\mu(\bar{B}_r)} \int_{\bar{B}_r} |\eta(x') - \eta(x)| \, d\mu(x') \rightarrow 0, \quad \text{as } r \rightarrow 0^+.$$

As before, we call the ratios in the the previous definition *Lebesgue ratios*.

3.1 Lebesgue points and nearest neighbors

In this section we show how to build an instance of a Nearest Neighbor algorithm in order to obtain a characterization of Lebesgue points in our general metric measure space (\mathcal{X}, d, μ) .

Take an arbitrary sequence of probability spaces $(\mathcal{Z}_m, \mathcal{G}_m, \nu_m)_{m \in \mathbb{N}}$. Let

$$\Omega := \mathcal{X} \times \mathcal{X}^{\mathbb{N}} \times \prod_{m \in \mathbb{N}} \mathcal{Z}_m \quad \text{and} \quad \mathcal{F} := \mathcal{B} \otimes \bigotimes_{k \in \mathbb{N}} \mathcal{B} \otimes \bigotimes_{m \in \mathbb{N}} \mathcal{Z}_m.$$

Since $\mu(\bar{B}_R(x)) < \infty$, we can define the probability measure

$$\begin{aligned} \nu: \mathcal{B} &\rightarrow [0, 1] \\ A &\mapsto \frac{\mu(A \cap \bar{B}_R(x))}{\mu(\bar{B}_R(x))} \end{aligned}$$

Let $\mathbb{P}: \mathcal{F} \rightarrow [0, 1]$ be the unique probability measure such that for all $A, A_1, A_2, \dots \in \mathcal{B}$, $G_1 \in \mathcal{G}_1, G_2 \in \mathcal{G}_2, \dots$, we have

$$\mathbb{P}(A \times A_1 \times A_2 \times \dots \times G_1 \times G_2 \times \dots) = \nu(A) \nu(A_1) \nu(A_2) \dots \nu_1(G_1) \nu_2(G_2) \dots.$$

It is well-known that such a probability measure exists [Halmos, 2013, Chapter VII, Section 38, Theorem B]. We define $X: \Omega \rightarrow \mathcal{X}$, $(x, x_1, x_2, \dots, z_1, z_2, \dots) \mapsto x$, for all $k \in \mathbb{N}$, $X_k: \Omega \rightarrow \mathcal{X}$, $(x, x_1, x_2, \dots, z_1, z_2, \dots) \mapsto x_k$, and for all $m \in \mathbb{N}$, $\Theta_m: \Omega \rightarrow \mathcal{Z}_m$, $(x, x_1, x_2, \dots, z_1, z_2, \dots) \mapsto z_m$. Finally, take an arbitrary sequence $(\psi_m)_{m \in \mathbb{N}}$ such that, for all $m \in \mathbb{N}$, $\psi_m: [0, \infty)^m \times \mathcal{Z}_m \rightarrow \{1, \dots, m\}$ is a measurable function satisfying, for all $r_1, \dots, r_m \geq 0$ and all $z \in \mathcal{Z}_m$, $\psi_m(r_1, \dots, r_m, z) \in \operatorname{argmin}_{k \in \{1, \dots, m\}} r_k$.

By construction, X, X_1, X_2, \dots are \mathbb{P} -independent random variables with common distribution $\mathbb{P}_X = \nu$ and $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an ISIMIN (Definition 4).

For the remainder of this section, for all $m \in \mathbb{N}$, X_m^x will be the nearest neighbor (among X_1, \dots, X_m) according to (ψ_m, Θ_m) (Definition 5).

Theorem 6. *If the restriction of η to \bar{B}_R is bounded, the following are equivalent:*

1. x is a Lebesgue point for η with respect to μ ;
2. $\mathbb{E}[|\eta(X_m^x) - \eta(x)|] \rightarrow 0$ as $m \rightarrow \infty$.

Proof. By Corollary 1, condition 2 is equivalent to x being a Lebesgue point for η with respect to \mathbb{P}_X . Since, for all $r \in (0, R]$, we have

$$\begin{aligned} \frac{1}{\mu(\bar{B}_r)} \int_{\bar{B}_r} |\eta(x') - \eta(x)| \, d\mu(x') &= \frac{1}{\mu(\bar{B}_r)/\mu(\bar{B}_R)} \int_{\bar{B}_r} |\eta(x') - \eta(x)| \, d\frac{\mu(x')}{\mu(\bar{B}_R)} \\ &= \frac{1}{\nu(\bar{B}_r)} \int_{\bar{B}_r} |\eta(x') - \eta(x)| \, d\nu(x') \\ &= \frac{\mathbb{E}\left[\mathbb{I}_{\bar{B}_r}(X) |\eta(X) - \eta(x)|\right]}{\mathbb{P}_X(\bar{B}_r)}, \end{aligned}$$

then x a Lebesgue point for η with respect to μ if and only if x is a Lebesgue point for η with respect to \mathbb{P}_X , which implies the result. \square

3.2 Sequential convergence

In this section we want to characterize Lebesgue points in terms of convergence to zero of Lebesgue ratios along suitable sequences of vanishing radii. Note that x is trivially a Lebesgue point if $\mu(\{x\}) > 0$ (by the dominated monotonicity of measures) or if there exists an $r > 0$ such that $\int_{\bar{B}_r} |\eta - \eta(x)| \, d\mu = 0$. Thus, for the remainder of this section we will focus on the non-trivial case in which $\mu(\{x\}) = 0$ and for all $r > 0$, $\int_{\bar{B}_r} |\eta - \eta(x)| \, d\mu > 0$.

We now generalize to arbitrary measures the definition of α -sequences introduced in Section 2.3 (Definition 3) which will play a central role in our sequential characterization of Lebesgue points.

Definition 7 (α -sequence). Fix any $\alpha \in (0, 1)$. For all $r > 0$, define

$$M(r) := \left(\int_{\bar{B}_r} |\eta - \eta(x)| \, d\mu \right)^\alpha \left(\mu(\bar{B}_r) \right)^{1-\alpha}.$$

Let $R_1 \in (0, R)$ such that $M(R_1) < M(R)$ and $m_1 := \lceil 1/M(R_1) \rceil$. For all $m \in \mathbb{N}$ such that $m \geq m_1$, define

$$r_m := \sup \left\{ r > 0 \mid M(r) < \frac{1}{m} \right\}.$$

We say that $(r_m)_{m \in \mathbb{N}, m \geq m_1}$ is the α -sequence (for η with respect to μ) at x .

Note that α -sequences are well-defined, strictly positive, and vanishing by our assumptions that $\mu(\{x\}) = 0$ and for all $r > 0$, $\int_{\bar{B}_r} |\eta - \eta(x)| \, d\mu > 0$ (as they were in Lemma 3).

We now take a closer look to the proofs of Theorems 2, 3, 4. Note that, there, it is redundant to assume that x is a Lebesgue point. Indeed, we merely used the fact that Lebesgue ratios with respect to closed (see after Definition 1) and open (the ratios appearing in (13)) balls are vanishing along the radii given by an α -sequence, for some $\alpha \in (0, 1)$. By Theorem 1, this proves that, at least when μ

is a probability measure, a construction like the one we presented in Section 3.1 gives that x is a Lebesgue point if and only if Lebesgue ratios with respect to both closed and open balls are vanishing along the radii given by an α -sequence, for some $\alpha \in (0, 1)$. This is very surprising, since, in general, if Lebesgue ratios (with respect to both closed and open balls) vanish along the radii given by a sequence, x is not necessarily a Lebesgue point, as shown by the following counterexample.

Example 3. Let $\mathcal{X} := [0, 1]$, d be the Euclidean distance, μ be the Lebesgue measure, and $x := 0$. For all $n \in \mathbb{N}$, define $\vartheta_n := \frac{1}{2^{2n}}$. Consider the function $\eta := \sum_{n=1}^{\infty} \mathbb{I}_{(\vartheta_{2n}, \vartheta_{2n-1}]}$. For any $m \in \mathbb{N}$, take $\rho_m := \vartheta_{2m}$ and $\rho'_m := \vartheta_{2m+1}$. We show now that the Lebesgue ratios with respect to both closed and open balls vanish along the sequence of radii $(\rho_m)_{m \in \mathbb{N}}$, but x is not a Lebesgue point since its Lebesgue ratios do *not* vanish along the sequence of radii $(\rho'_m)_{m \in \mathbb{N}}$. Indeed, if $m \rightarrow \infty$, we have that

$$\begin{aligned} \frac{1}{\mu(B_{\rho_m})} \int_{B_{\rho_m}} |\eta - \eta(x)| \, d\mu &= \frac{1}{\mu(\bar{B}_{\rho_m})} \int_{\bar{B}_{\rho_m}} |\eta - \eta(x)| \, d\mu = \frac{1}{\vartheta_{2m}} \int_0^{\vartheta_{2m}} \eta(t) \, dt \\ &= \frac{1}{\vartheta_{2m}} \int_0^{\vartheta_{2m+1}} \eta(t) \, dt \leq \frac{\vartheta_{2m+1}}{\vartheta_{2m}} = \frac{1}{2^{2^{2m}}} \rightarrow 0 \end{aligned}$$

but, at the same time,

$$\begin{aligned} \frac{1}{\mu(\bar{B}_{\rho'_m})} \int_{\bar{B}_{\rho'_m}} |\eta - \eta(x)| \, d\mu &= \frac{1}{\vartheta_{2m+1}} \int_0^{\vartheta_{2m+1}} \eta(t) \, dt \geq \frac{1}{\vartheta_{2m+1}} \int_{\vartheta_{2m+2}}^{\vartheta_{2m+1}} 1 \, dt \\ &= 1 - \frac{\vartheta_{2m+2}}{\vartheta_{2m+1}} = 1 - \frac{1}{2^{2^{2m+1}}} \rightarrow 1. \end{aligned}$$

This highlights that α -sequences are very special, because each one of them contains in itself enough information to characterize the convergence to zero of Lebesgue ratios in the continuum. We will prove that even more is true but before presenting the result, we give a handy definition.

Definition 8 (Lebesgue point along a sequence). Take any $k \in \mathbb{N}$ and any vanishing sequence of strictly positive numbers $(\rho_m)_{m \in \mathbb{N}, m \geq k}$. We say that x is a Lebesgue point (for η with respect to μ) along $(\rho_m)_{m \in \mathbb{N}, m \geq k}$ if

$$\frac{1}{\mu(\bar{B}_{\rho_m})} \int_{\bar{B}_{\rho_m}} |\eta - \eta(x)| \, d\mu \rightarrow 0, \quad \text{as } m \rightarrow \infty.$$

We remark that, unlike the results we proved in Section 2, the following theorem holds for more general measures and (possibly) unbounded integrands, with the minimal assumption that η is locally-integrable around x . The equivalence will be proved directly, without relying on nearest neighbor techniques.

Theorem 7. *If $\int_{\bar{B}_R} |\eta| \, d\mu < \infty$, then the following are equivalent:*

1. x is a Lebesgue point along an α -sequence, for some $\alpha \in (0, 1)$;
2. x is a Lebesgue point.

Proof. We prove the non-trivial implication. Let $\alpha \in (0, 1)$ and assume that x is a Lebesgue point along the α -sequence $(r_m)_{m \in \mathbb{N}, m \geq m_1}$. To lighten the notation, we define the auxiliary function

$$\begin{aligned} f: \mathcal{X} &\rightarrow \mathbb{R} \\ x' &\mapsto |\eta(x') - \eta(x)|. \end{aligned}$$

With a straightforward adaptation of (12), (11), and (10), we can prove that for all $m \in \mathbb{N}$, $m \geq m_1$, we have

$$\left(\int_{B_{r_m}} f \, d\mu \right)^\alpha \left(\mu(B_{r_m}) \right)^{1-\alpha} \leq \frac{1}{m} \leq \left(\int_{\bar{B}_{r_m}} f \, d\mu \right)^\alpha \left(\mu(\bar{B}_{r_m}) \right)^{1-\alpha}. \quad (17)$$

Assume by contradiction that there exists a vanishing sequence $(s_m)_{m \in \mathbb{N}}$ of strictly positive numbers and a $\delta > 0$ such that, for all $m \in \mathbb{N}$, we have that

$$\frac{1}{\mu(\bar{B}_{s_m})} \int_{\bar{B}_{s_m}} f \, d\mu \geq \delta. \quad (18)$$

Without loss of generality, we can assume that for all $m \in \mathbb{N}$,

$$\left(\int_{\bar{B}_{s_m}} f \, d\mu \right)^\alpha \left(\mu(\bar{B}_{s_m}) \right)^{1-\alpha} < 1.$$

For all $m \in \mathbb{N}$, let $n_m \in \mathbb{N}$ such that

$$\frac{1}{n_m + 1} \leq \left(\int_{\bar{B}_{s_m}} f \, d\mu \right)^\alpha \left(\mu(\bar{B}_{s_m}) \right)^{1-\alpha} < \frac{1}{n_m}. \quad (19)$$

Note that $n_m \rightarrow \infty$ as $m \rightarrow \infty$, since the middle term vanishes as $m \rightarrow \infty$. Hence, there exists $m_2 \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ with $m \geq m_2$, we have that $n_m \geq m_1$. From the first inequality in (17) and the first inequality in (19), for all $m \in \mathbb{N}$ such that $m \geq m_2$, we have that $r_{n_m+1} \leq s_m$, which in turn gives

$$\mu(\bar{B}_{r_{n_m+1}}) \leq \mu(\bar{B}_{s_m}). \quad (20)$$

Hence, for all $m \in \mathbb{N}$ such that $m \geq m_2$, we have that

$$\begin{aligned} \mu(\bar{B}_{r_{n_m+1}}) &\stackrel{(20)}{\leq} \mu(\bar{B}_{s_m}) = 1 \cdot \mu(\bar{B}_{s_m}) \stackrel{(18)}{\leq} \left(\frac{1}{\delta} \frac{1}{\mu(\bar{B}_{s_m})} \int_{\bar{B}_{s_m}} f \, d\mu \right)^\alpha \mu(\bar{B}_{s_m}) \\ &= \frac{1}{\delta^\alpha} \left(\int_{\bar{B}_{s_m}} f \, d\mu \right)^\alpha \left(\mu(\bar{B}_{s_m}) \right)^{1-\alpha} \stackrel{(19)}{<} \frac{1}{\delta^\alpha} \frac{1}{n_m}. \end{aligned} \quad (21)$$

Finally, for all $m \in \mathbb{N}$ such that $m \geq m_2$, we have that

$$\begin{aligned} \left(\frac{1}{\mu(\bar{B}_{r_{n_m+1}})} \int_{\bar{B}_{r_{n_m+1}}} f \, d\mu \right)^\alpha &= \left(\int_{\bar{B}_{r_{n_m+1}}} f \, d\mu \right)^\alpha \left(\mu(\bar{B}_{r_{n_m+1}}) \right)^{1-\alpha} \frac{1}{\mu(\bar{B}_{r_{n_m+1}})} \\ &\stackrel{(17)}{\geq} \frac{1}{n_m + 1} \frac{1}{\mu(\bar{B}_{r_{n_m+1}})} \stackrel{(21)}{\geq} \delta^\alpha \frac{n_m}{n_m + 1} \xrightarrow{m \rightarrow \infty} \delta^\alpha > 0, \end{aligned}$$

which, since $n_m \rightarrow \infty$ as $m \rightarrow \infty$, implies that x is not a Lebesgue point along a subsequence of $(r_m)_{m \in \mathbb{N}, m \geq m_1}$, contradicting the fact that x is a Lebesgue point along $(r_m)_{m \in \mathbb{N}, m \geq m_1}$. \square

4 From Lebesgue points to Lebesgue values

In this brief section, we provide a straightforward but useful² generalization of the previous results, shifting the focus from Lebesgue points to Lebesgue values.

Definition 9 (Lebesgue value). Let (\mathcal{X}, d) be a metric space, μ a locally-finite Borel measure of (\mathcal{X}, d) , $\eta : \mathcal{X} \rightarrow \mathbb{R}$ a locally-integrable function with respect to μ and $x \in \mathcal{X}$ a point in the support of μ . We say that $l \in \mathbb{R}$ is the *Lebesgue value* of η at x (with respect to μ) if

$$\frac{1}{\mu(\bar{B}_r(x))} \int_{\bar{B}_r(x)} |\eta(x') - l| d\mu(x') \rightarrow 0, \quad r \rightarrow 0^+.$$

We point out that we use the word “the” in the definition of Lebesgue values since if $l_1, l_2 \in \mathbb{R}$ are Lebesgue values for η at x with respect to μ , then the triangle inequality gives immediately $l_1 = l_2$.

We now make two key observations about Lebesgue values. The first one is that if $\mu(\{x\}) > 0$, then η admits $\eta(x)$ as its Lebesgue value at x . The second one is that if $\mu(\{x\}) = 0$ and η admits $l \in \mathbb{R}$ as its Lebesgue value at x , then we can define

$$\begin{aligned} \tilde{\eta} : \mathcal{X} &\rightarrow \mathbb{R} \\ x' &\mapsto \begin{cases} l & \text{if } x' = x \\ \eta(x') & \text{otherwise} \end{cases} \end{aligned}$$

and obtain that $\tilde{\eta}$ has a Lebesgue point at x . With these two observations in mind, we can restate appropriately *every condition and every result we obtained so far*, using Lebesgue values instead of Lebesgue points.

We illustrate with an example how this translation works for e.g., Corollary 1. Consider the same setting as in Section 2.

Corollary 2. *If $(X_m^x)_{m \in \mathbb{N}}$ is defined according to an ISIMIN (Definitions 5 and 4) and $l \in \mathbb{R}$, the following are equivalent:*

1. l is the Lebesgue value for η at x with respect to \mathbb{P}_X ;
2. $\eta(X_m^x)$ converges to l in \mathcal{L}^1 , as $m \rightarrow \infty$.

Proof. Assume that $\eta(X_m^x)$ converges to l in \mathcal{L}^1 , as $m \rightarrow \infty$. If $\mathbb{P}_X(\{x\}) > 0$ then $l = \eta(x)$ and the results follows directly from Corollary 1. If $\mathbb{P}_X(\{x\}) = 0$ then, since also $\mathbb{P}_{X_m^x}(\{x\}) = 0$, using $\tilde{\eta}$ as above, we have that

$$\mathbb{E}[|\tilde{\eta}(X_m^x) - \tilde{\eta}(x)|] = \mathbb{E}[|\eta(X_m^x) - l|] \rightarrow 0, \quad m \rightarrow \infty,$$

and Corollary 1 applied to $\tilde{\eta}$ yields

$$\frac{\mathbb{E}[\mathbb{I}_{\bar{B}_r(x)}(X) |\eta(X) - l|]}{\mathbb{P}_X(\bar{B}_r(x))} = \frac{\mathbb{E}[\mathbb{I}_{\bar{B}_r(x)}(X) |\tilde{\eta}(X) - \tilde{\eta}(x)|]}{\mathbb{P}_X(\bar{B}_r(x))} \rightarrow 0, \quad r \rightarrow 0^+.$$

²For an application, see e.g., Example 4.

Vice versa, assume that l is the Lebesgue value for η at x with respect to \mathbb{P}_X . If $\mathbb{P}_X(\{x\}) > 0$ then $l = \eta(x)$ and the results follows from Corollary 1. If $\mathbb{P}_X(\{x\}) = 0$ then, using $\tilde{\eta}$ as above,

$$\frac{\mathbb{E}\left[\mathbb{I}_{\tilde{B}_r(x)}(X) |\tilde{\eta}(X) - \tilde{\eta}(x)|\right]}{\mathbb{P}_X(\tilde{B}_r(x))} = \frac{\mathbb{E}\left[\mathbb{I}_{\tilde{B}_r(x)}(X) |\eta(X) - l|\right]}{\mathbb{P}_X(\tilde{B}_r(x))} \rightarrow 0, \quad r \rightarrow 0^+,$$

and since also $\mathbb{P}_{X_m^x}(\{x\}) = 0$, Corollary 1 gives us

$$\mathbb{E}[|\eta(X_m^x) - l|] = \mathbb{E}[|\tilde{\eta}(X_m^x) - \tilde{\eta}(x)|] \rightarrow 0, \quad m \rightarrow \infty. \quad \square$$

5 ISIMINs and nearest neighbor classification

In this section we present applications of ISIMINs to nearest neighbor classification. Our results extend what is known for lexicographical tie-breaking rules [Devroye, 1981b] to the more general tie-breaking rules given by ISIMINs. In particular, this implies that our results hold for the (uniformly) random tie-breaking rule in [C erou and Guyader, 2006].

We consider the same setting as in Section 2 (more specifically, of Subsection 2.4), with the following differences:

1. there is no fixed $x \in \mathcal{X}$;
2. there exists a sequence of random variables $Y_1, Y_2, \dots : \Omega \rightarrow \{0, 1\}$ such that $(X, Y), (X_1, Y_1), (X_2, Y_2), \dots$ are \mathbb{P} -i.i.d.;
3. $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ is an ISIMIN such that for each $m \in \mathbb{N}$ we have that Θ_m is \mathbb{P} -independent of $(X, Y), (X_1, Y_1), \dots, (X_m, Y_m)$;
4. rather than being an arbitrary bounded function, $\eta : \Omega \rightarrow \mathbb{R}$ is the *regression function* of Y with respect to X , i.e., η is any measurable function such that $\eta(X) = \mathbb{E}[Y | X]$ (whose existence is guaranteed by the Doob–Dynkin Lemma [Rao and Swift, 2006, Chapter 1.2, Proposition 3]), that we can assume $[0, 1]$ -valued.

We now define nearest neighbor classification by means of an ISIMIN.

Definition 10 (Nearest neighbor classification according to an ISIMIN). For all $m \in \mathbb{N}$, the nearest neighbor classification (with training set $(X_1, Y_1), \dots, (X_m, Y_m)$ and test data point X) according to (ψ_m, Θ_m) is defined by Y_{Ψ_m} , where Ψ_m is the random index of the random point among X_1, \dots, X_m that is closest to X according to (ψ_m, Θ_m) , i.e., $\Psi_m := \psi_m(d(X, X_1), \dots, d(X, X_m), \Theta_m)$.

We will prove the convergence of the classification risk $\mathbb{P}(Y_{\Psi_m} \neq Y)$ of the nearest neighbor Y_{Ψ_m} whenever the Lebesgue–Besicovitch differentiation theorem holds for η , i.e., if \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for η . This will follow from our results on ISIMINs and Proposition 2. This proposition is similar in spirit to other known results for plug-in decisions, but it is tailored to our 1-Nearest Neighbor classification problem. E.g., in [Devroye et al., 1996, Theorem 2.2], the comparison term on the right hand side is the Bayes risk $\mathbb{E}[\min(\eta(X), 1 - \eta(X))]$,

since the goal there is to obtain consistency. However, it has long been known that the 1-Nearest Neighbor algorithm is not consistent without additional assumptions (see, e.g., [Devroye et al., 1996, Theorem 5.4 and subsequent remark]). Nevertheless, one can study the convergence of its risk. With this goal in mind, the appropriate quantity to compare $\mathbb{P}(Y_{\Psi_m} \neq Y)$ to is not the Bayes risk, but rather a “surrogate risk”, which turns out to be $2\mathbb{E}[\eta(X)(1 - \eta(X))]$. A similar formula appears also in [Devroye, 1981b].

Proposition 2. *For all $m \in \mathbb{N}$, we have*

$$\left| \mathbb{P}(Y_{\Psi_m} \neq Y) - 2\mathbb{E}[\eta(X)(1 - \eta(X))] \right| \leq \int_{\mathcal{X}} \mathbb{E}[|\eta(X_m^x) - \eta(x)|] d\mathbb{P}_X(x),$$

where for each $x \in \mathcal{X}$, X_m^x is the nearest neighbor of x (among X_1, \dots, X_m), according to (ψ_m, Θ_m) , i.e., $X_m^x := X_{\psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m)}$.

Proof. Fix any $m \in \mathbb{N}$. Note that

$$\mathbb{P}(Y_{\Psi_m} \neq Y) = \mathbb{P}(\{Y_{\Psi_m} = 1\} \cap \{Y = 0\}) + \mathbb{P}(\{Y_{\Psi_m} = 0\} \cap \{Y = 1\}).$$

We analyze the first term. The second one can be computed similarly. Note that

$$\mathbb{P}(\{Y_{\Psi_m} = 1\} \cap \{Y = 0\}) = \mathbb{E}[Y_{\Psi_m}(1 - Y)] = \mathbb{E}[\mathbb{E}[Y_{\Psi_m}(1 - Y) | X]] = (\star).$$

Now, if we let $\mathcal{V} = \mathcal{X}$, $\mathcal{W} = (\mathcal{X} \times \{0, 1\})^m \times \mathcal{Z}_m$, $U = 1 - Y$, $V = X$, $W = ((X_1, Y_1), \dots, (X_m, Y_m), \Theta_m)$, and

$$f: \mathcal{V} \times \mathcal{W} \rightarrow [0, 1],$$

$$(x, ((x_1, y_1), \dots, (x_m, y_m), z_m)) \mapsto y_{\psi_m(d(x, x_1), \dots, d(x, x_m), z_m)},$$

applying Lemma 7 (see Appendix A), we get

$$\begin{aligned} \mathbb{E}[(1 - Y)Y_{\Psi_m} | X] &= \mathbb{E}[Uf(V, W) | V] \stackrel{(24)}{=} \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V] \\ &= \mathbb{E}[1 - Y | X] \mathbb{E}[Y_{\Psi_m} | X] = (1 - \mathbb{E}[Y | X]) \mathbb{E}[Y_{\Psi_m} | X] \\ &= (1 - \eta(X)) \mathbb{E}[Y_{\Psi_m} | X] = \mathbb{E}[Y_{\Psi_m}(1 - \eta(X)) | X]. \end{aligned}$$

Thus

$$\begin{aligned} (\star) &= \mathbb{E}[\mathbb{E}[Y_{\Psi_m}(1 - \eta(X)) | X]] = \mathbb{E}[Y_{\Psi_m}(1 - \eta(X))] \\ &= \sum_{k=1}^m \mathbb{E}[Y_k(1 - \eta(X)) \mathbb{I}_{\{\Psi_m = k\}}] \\ &= \sum_{k=1}^m \mathbb{E}[\mathbb{E}[Y_k(1 - \eta(X)) \mathbb{I}_{\{\Psi_m = k\}} | X_k]] = (\star\star) \end{aligned}$$

Now, if for each $k \in \{1, \dots, m\}$, we let $\mathcal{V} = \mathcal{X}$, $\mathcal{W} = \mathcal{X}^m \times \mathcal{Z}_m$, $U = Y_k$, $V = X_k$, $W = (X, X_1, \dots, X_{k-1}, X_{k+1}, \dots, X_m, \Theta_m)$, applying Lemma 7 (Appendix A) to the function

$$f: \mathcal{V} \times \mathcal{W} \rightarrow [0, 1],$$

$$\left(x_k, ((x, x_{1:k-1}, x_{k+1:m}), z_m) \right) \mapsto (1 - \eta(x)) \mathbb{I}_{\{\psi_m=k\}}(d(x, x_1), \dots, d(x, x_m), z_m),$$

where $x_{1:k-1}$ (resp., $x_{k+1:m}$) is a shorthand for x_1, \dots, x_{k-1} (resp., x_{k+1}, \dots, x_m)—with the obvious adjustments for $k = 1$ (resp., $k = m$)—yields

$$\begin{aligned} \mathbb{E}\left[Y_k (1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}} \mid X_k\right] &= \mathbb{E}\left[U f(V, W) \mid V\right] \stackrel{(24)}{=} \mathbb{E}[U \mid V] \mathbb{E}[f(V, W) \mid V] \\ &= \mathbb{E}[Y_k \mid X_k] \mathbb{E}\left[(1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}} \mid X_k\right] \\ &= \eta(X_k) \mathbb{E}\left[(1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}} \mid X_k\right] \\ &= \mathbb{E}\left[\eta(X_k) (1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}} \mid X_k\right]. \end{aligned}$$

Thus

$$\begin{aligned} (\star\star) &= \sum_{k=1}^m \mathbb{E}\left[\mathbb{E}\left[\eta(X_k) (1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}} \mid X_k\right]\right] \\ &= \sum_{k=1}^m \mathbb{E}\left[\eta(X_k) (1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}}\right] \\ &= \sum_{k=1}^m \mathbb{E}\left[\eta(X_{\Psi_m}) (1 - \eta(X)) \mathbb{I}_{\{\Psi_m=k\}}\right] = \mathbb{E}\left[\eta(X_{\Psi_m}) (1 - \eta(X))\right]. \end{aligned}$$

So

$$\mathbb{P}\left(\{Y_{\Psi_m} = 1\} \cap \{Y = 0\}\right) = \mathbb{E}\left[\eta(X_{\Psi_m})(1 - \eta(X))\right].$$

Analogously, we can prove

$$\mathbb{P}\left(\{Y_{\Psi_m} = 0\} \cap \{Y = 1\}\right) = \mathbb{E}\left[\eta(X)(1 - \eta(X_{\Psi_m}))\right].$$

Hence

$$\begin{aligned}
& \left| \mathbb{P}(Y_{\Psi_m} \neq Y) - 2\mathbb{E}[\eta(X)(1 - \eta(X))] \right| \\
&= \left| \mathbb{E}[\eta(X_{\Psi_m})(1 - \eta(X))] + \mathbb{E}[\eta(X)(1 - \eta(X_{\Psi_m}))] - 2\mathbb{E}[\eta(X)(1 - \eta(X))] \right| \\
&= \left| \mathbb{E}[(\eta(X_{\Psi_m}) - \eta(X))(1 - \eta(X)) + \eta(X)((1 - \eta(X_{\Psi_m})) - (1 - \eta(X)))] \right| \\
&= \left| \mathbb{E}[(1 - 2\eta(X))(\eta(X_{\Psi_m}) - \eta(X))] \right| \leq \mathbb{E} \left[\left| (1 - 2\eta(X))(\eta(X_{\Psi_m}) - \eta(X)) \right| \right] \\
&\leq \mathbb{E} \left[|\eta(X_{\Psi_m}) - \eta(X)| \right] = \mathbb{E} \left[\mathbb{E} \left[|\eta(X_{\Psi_m}) - \eta(X)| \mid X \right] \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[|\eta(X_{\psi_m(d(X, X_1), \dots, d(X, X_m), \Theta_m)}) - \eta(X)| \mid X \right] \right] \\
&\stackrel{(26)}{=} \mathbb{E} \left[\mathbb{E} \left[|\eta(X_{\psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m)}) - \eta(x)| \right]_{x=X} \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[|\eta(X_m^x) - \eta(x)| \right]_{x=X} \right] = \int_{\mathcal{X}} \mathbb{E} \left[|\eta(X_m^x) - \eta(x)| \right] d\mathbb{P}_X(x),
\end{aligned}$$

where the second to last inequality follows by the Freezing Lemma (Lemma 8 in Appendix A). \square

We can now state the most important result of this section. It guarantees the convergence of the classification risk of the 1-Nearest Neighbor algorithm, using ISIMINs, on arbitrary metric spaces, with the only assumption that \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for the regression function η with respect to \mathbb{P}_X .

Theorem 8. *If \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for η with respect to \mathbb{P}_X , then*

$$\mathbb{P}(Y_{\Psi_m} \neq Y) \rightarrow 2\mathbb{E}[\eta(X)(1 - \eta(X))], \quad m \rightarrow \infty. \quad (22)$$

Proof. As in Proposition 2, we define, for each $m \in \mathbb{N}$ and any $x \in \mathcal{X}$, X_m^x as the nearest neighbor of x (among X_1, \dots, X_m), according to (ψ_m, Θ_m) , i.e., $X_m^x := X_{\psi_m(d(x, X_1), \dots, d(x, X_m), \Theta_m)}$. Take $N \subset \mathcal{X}$ be a Borel set of (\mathcal{X}, d) such that $\mathbb{P}_X(N) = 0$ and, for all $x \in \mathcal{X} \setminus N$, we have that x is a Lebesgue point. Then, Proposition 2 implies, for all $m \in \mathbb{N}$

$$\begin{aligned}
\left| \mathbb{P}(Y_{\Psi_m} \neq Y) - 2\mathbb{E}[\eta(X)(1 - \eta(X))] \right| &\leq \int_{\mathcal{X}} \mathbb{E} \left[|\eta(X_m^x) - \eta(x)| \right] d\mathbb{P}_X(x) \\
&= \int_{\mathcal{X} \setminus N} \mathbb{E} \left[|\eta(X_m^x) - \eta(x)| \right] d\mathbb{P}_X(x) \rightarrow 0, \quad \text{as } m \rightarrow \infty,
\end{aligned}$$

where the last term vanishes by Corollary 1 and the dominated convergence theorem. \square

The previous result gives immediately the consistency of the 1-Nearest Neighbor classifier Y_{Ψ_m} in arbitrary metric spaces under the *realizability assumption*

$$\exists f: \mathcal{X} \rightarrow \{0, 1\}, f \text{ measurable, such that } \mathbb{P}(f(X) = Y) = 1. \quad (23)$$

Before stating the result, we recall that a classification algorithm is said \mathbb{P} -consistent if its classification risk converges to the *Bayes risk*

$$L^* := \mathbb{E} \left[\min \left(\eta(X), (1 - \eta(X)) \right) \right].$$

In our setting, the \mathbb{P} -consistency condition for Y_{Ψ_m} is

$$\mathbb{P}(Y_{\Psi_m} \neq Y) \rightarrow L^*, \quad m \rightarrow \infty.$$

Corollary 3. *Assume that \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for η with respect to \mathbb{P}_X . Then, the following are equivalent:*

1. Y_{Ψ_m} is \mathbb{P} -consistent and $\mathbb{P}(\eta(X) = 1/2) = 0$;
2. the realizability assumption holds.

Proof. We know from the previous theorem that the classification risk of Y_{Ψ_m} converges to $2\mathbb{E}[\eta(X)(1 - \eta(X))]$, as $m \rightarrow \infty$. Since for any $y \in (0, \frac{1}{2}) \cup (\frac{1}{2}, 1)$ it holds that $2y(1 - y) > \min(y, 1 - y)$, to have $2\mathbb{E}[\eta(X)(1 - \eta(X))] = L^*$ it is necessary and sufficient that $\mathbb{P}(\eta(X) \in \{0, \frac{1}{2}, 1\}) = 1$. Note that, since Y is $\{0, 1\}$ -valued, there exists $F \in \mathcal{F}$ such that $\mathbb{I}_F = Y$.

Suppose that the realizability assumption holds. Let $f: \mathcal{X} \rightarrow \{0, 1\}$ be a measurable function such that $\mathbb{P}(f(X) = Y) = 1$. Since

$$\eta(X) = \mathbb{E}[Y | X] = \mathbb{E}[f(X) | X] = f(X) = Y = \mathbb{I}_F, \quad \mathbb{P}\text{-almost surely,}$$

it holds that $\mathbb{P}(\eta(X) \in \{0, 1\}) = 1$. Thus, we have that Y_{Ψ_m} is \mathbb{P} -consistent and $\mathbb{P}(\eta(X) = 1/2) = 0$.

Vice versa, suppose that Y_{Ψ_m} is \mathbb{P} -consistent and $\mathbb{P}(\eta(X) = 1/2) = 0$. Then $\mathbb{P}(\eta(X) \in \{0, 1\}) = 1$. Let $E := (\eta(X))^{-1}(\{1\})$. Since E is $\sigma(\eta(X))$ -measurable, it is also $\sigma(X)$ -measurable. Then, by definition, there exists a Borel subset A of (\mathcal{X}, d) such that $E = X^{-1}(A)$, and so $\mathbb{I}_E = \mathbb{I}_{X^{-1}(A)} = \mathbb{I}_A(X)$. We show that the realizability assumption holds with $f = \mathbb{I}_A$. Now, note that

$$\mathbb{P}(E) = \mathbb{E}[\mathbb{I}_E] = \mathbb{E}[\mathbb{I}_E \mathbb{I}_E] = \mathbb{E}[\mathbb{I}_E \mathbb{E}[Y | X]] = \mathbb{E}[\mathbb{I}_E Y] = \mathbb{E}[\mathbb{I}_E \mathbb{I}_F] = \mathbb{P}(E \cap F),$$

where we used the fact that $\mathbb{P}(\mathbb{I}_E = \eta(X)) = 1$. Furthermore, note that

$$0 = \mathbb{E}[0] = \mathbb{E}[\mathbb{I}_{E^c} \mathbb{I}_E] = \mathbb{E}[\mathbb{I}_{E^c} \mathbb{E}[Y | X]] = \mathbb{E}[\mathbb{I}_{E^c} Y] = \mathbb{E}[\mathbb{I}_{E^c} \mathbb{I}_F] = \mathbb{P}(E^c \cap F).$$

Then

$$\begin{aligned} \mathbb{E} \left[\left| \mathbb{I}_A(X) - Y \right| \right] &= \mathbb{E} \left[\left| \mathbb{I}_E - \mathbb{I}_F \right| \right] = \mathbb{P}(E \cap F^c) + \mathbb{P}(E^c \cap F) \\ &= \mathbb{P}(E \cap F^c) = \mathbb{P}(E) - \mathbb{P}(E \cap F) = 0, \end{aligned}$$

that is equivalent to $\mathbb{P}(\mathbb{I}_A(X) = Y) = 1$. \square

As we pointed out in the introduction, if $(\mathcal{V}, d_{\mathcal{V}})$ is a Euclidean space then μ -almost every $v \in \mathcal{V}$ is a Lebesgue point for f with respect to μ , for every Borel probability measure μ and each bounded measurable function $f: \mathcal{V} \rightarrow \mathbb{R}$. This is a consequence of Lebesgue–Besicovitch differentiation theorem [Evans and Gariepy, 2015, Theorems 1.32-1.33]. The same holds true if $(\mathcal{V}, d_{\mathcal{V}})$ is a finite dimensional Banach space [Loeb, 2006], a locally-compact separable ultrametric space [Simmons, 2012, Theorem 9.1], a separable Riemannian manifold [Simmons, 2012, Theorem 9.1], or the straightforward case where \mathcal{V} is (at most) countable. In particular, if (\mathcal{X}, d) is one of the previous metric spaces, then \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for η with respect to \mathbb{P}_X , regardless of which specific Borel probability measure \mathbb{P}_X is. This does not hold in every metric space. Indeed, Preiss [1979] showed that if $(\mathcal{V}, d_{\mathcal{V}})$ is an infinite dimensional separable Hilbert space, then counterexamples exist, even if the underlying Borel probability measure is Gaussian. Preiss [1983] also characterized the metric spaces where Lebesgue–Besicovitch differentiation theorem holds true, in terms of a notion of σ -finite dimensionality of the space. Cérou and Guyader [2006] used Preiss’ counterexample to show that the k_m -Nearest Neighbor classification algorithm is not necessarily consistent when the Lebesgue–Besicovitch differentiation theorem does not hold. The same idea works in our context.

Example 4. Preiss [1979] showed that there exists a Polish metric space (\mathcal{X}, d) (which is actually an infinite dimensional separable Hilbert space), a Borel probability measure μ (which is actually Gaussian) and a compact set K of (\mathcal{X}, d) such that $\mu(K) > 0$ and, for all $x \in \mathcal{X}$, $x \in \text{supp}(\mu)$ and

$$\frac{1}{\mu(\bar{B}_r(x))} \int_{\bar{B}_r(x)} \mathbb{I}_K d\mu \rightarrow 0, \quad \text{as } r \rightarrow 0^+.$$

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, X, X_1, X_2, \dots be \mathbb{P} -i.i.d. random variables with common distribution $\mathbb{P}_X = \mu$, and $(\psi_m, \Theta_m)_{m \in \mathbb{N}}$ be an ISIMIN (for the existence of this setting, see Section 3.1). Define $Y := \mathbb{I}_K(X), Y_1 := \mathbb{I}_K(X_1), Y_2 := \mathbb{I}_K(X_2), \dots$. As before, we denote for any $m \in \mathbb{N}$ and all $x \in \mathcal{X}$, the nearest neighbor (among X_1, \dots, X_m) according to (ψ_m, Θ_m) by $X_m^x := X_{\psi(d(x, X_1), \dots, d(x, X_m), \Theta_m)}$. Note that, for all $x \in \mathcal{X}$, we have that

$$\frac{\mathbb{E} \left[\mathbb{I}_{\bar{B}_r(x)}(X) |\mathbb{I}_K(X) - 0| \right]}{\mathbb{P}_X(\bar{B}_r(x))} = \frac{1}{\mu(\bar{B}_r(x))} \int_{\bar{B}_r(x)} \mathbb{I}_K d\mu \rightarrow 0, \quad \text{as } r \rightarrow 0^+.$$

By Corollary 2, this implies that, for all $x \in \mathcal{X}$,

$$\mathbb{E}[\mathbb{I}_K(X_m^x)] = \mathbb{E} \left[|\mathbb{I}_K(X_m^x) - 0| \right] \rightarrow 0, \quad \text{as } m \rightarrow \infty,$$

and then also

$$\mathbb{E}[\mathbb{I}_{K^c}(X_m^x)] = 1 - \mathbb{E}[\mathbb{I}_K(X_m^x)] \rightarrow 1, \quad \text{as } m \rightarrow \infty.$$

Thus, the Freezing lemma (Lemma 8) and Lebesgue’s dominated convergence theorem yield

$$\begin{aligned}
\mathbb{P}(Y_{\Psi_m} \neq Y) &= \mathbb{P}(\{Y_{\Psi_m} = 0\} \cap \{Y = 1\}) + \mathbb{P}(\{Y_{\Psi_m} = 1\} \cap \{Y = 0\}) \\
&= \mathbb{P}(\{X_{\Psi_m} \in K^c\} \cap \{X \in K\}) + \mathbb{P}(\{X_{\Psi_m} \in K\} \cap \{X \in K^c\}) \\
&= \mathbb{E}[\mathbb{I}_{K^c}(X_{\Psi_m})\mathbb{I}_K(X)] + \mathbb{E}[\mathbb{I}_K(X_{\Psi_m})\mathbb{I}_{K^c}(X)] \\
&= \mathbb{E}\left[\mathbb{E}[\mathbb{I}_{K^c}(X_{\Psi_m})\mathbb{I}_K(X) \mid X]\right] + \mathbb{E}\left[\mathbb{E}[\mathbb{I}_K(X_{\Psi_m})\mathbb{I}_{K^c}(X) \mid X]\right] \\
&\stackrel{(26)}{=} \mathbb{E}\left[\mathbb{E}[\mathbb{I}_{K^c}(X_m^x)\mathbb{I}_K(x)]_{x=X}\right] + \mathbb{E}\left[\mathbb{E}[\mathbb{I}_K(X_m^x)\mathbb{I}_{K^c}(x)]_{x=X}\right] \\
&= \int_{\mathcal{X}} \mathbb{E}[\mathbb{I}_{K^c}(X_m^x)\mathbb{I}_K(x)] \, d\mathbb{P}_X(x) + \int_{\mathcal{X}} \mathbb{E}[\mathbb{I}_K(X_m^x)\mathbb{I}_{K^c}(x)] \, d\mathbb{P}_X(x) \\
&= \int_{\mathcal{X}} \mathbb{I}_K(x)\mathbb{E}[\mathbb{I}_{K^c}(X_m^x)] \, d\mathbb{P}_X(x) + \int_{\mathcal{X}} \mathbb{I}_{K^c}(x)\mathbb{E}[\mathbb{I}_K(X_m^x)] \, d\mathbb{P}_X(x) \\
&\quad \rightarrow \mathbb{P}_X(K) + 0 = \mathbb{P}_X(K) = \mu(K), \quad \text{as } m \rightarrow +\infty.
\end{aligned}$$

Now, note that $\mathbb{I}_K(X) = \mathbb{E}[Y \mid X]$, and so we can choose $\mathbb{I}_K(X)$ as the regression function η , leading us to $2\mathbb{E}[\eta(X)(1 - \eta(X))] = 2\mathbb{E}[\mathbb{I}_K(X)\mathbb{I}_{K^c}(X)] = 0 < \mu(K)$, which implies that

$$\mathbb{P}(Y_{\Psi_m} \neq Y) \not\rightarrow 2\mathbb{E}[\eta(X)(1 - \eta(X))], \quad \text{as } m \rightarrow \infty.$$

Therefore, Corollary 3 does not hold without requiring that \mathbb{P}_X -almost every $x \in \mathcal{X}$ is a Lebesgue point for η with respect to \mathbb{P}_X .

Acknowledgments

An earlier version of this work appeared in Roberto Colomboni’s master’s thesis, written under the supervision of Nicolò Cesa-Bianchi. Both Tom and Rob gratefully acknowledge Nicolò’s helpful advice. Rob also thanks Guglielmo Beretta for the many enlightening discussions. This work has benefitted from the AI Interdisciplinary Institute ANITI. ANITI is funded by the French “Investing for the Future – PIA3” program under the Grant agreement n. ANR-19-PI3A-0004.³

References

- Christophe Abraham, Gérard Biau, and Benoît Cadre. On the kernel rule for function classification. *Annals of the Institute of Statistical Mathematics*, 58(3): 619–633, 2006.
- Paolo Baldi. Stochastic calculus. In *Stochastic Calculus*, pages 215–254. Springer, 2017.

³<https://aniti.univ-toulouse.fr/>

- G erard Biau and Luc Devroye. *Lectures on the nearest neighbor method*. Springer, 2015.
- Anders Bj orn, Jana Bj orn, and Mikko Parviainen. Lebesgue points and the fundamental convergence theorem for superharmonic functions on metric spaces. *Revista Matem tica Iberoamericana*, 26(1):147–174, 2010.
- Fr d ric C rou and Arnaud Guyader. Nearest neighbor classification in infinite dimension. *ESAIM: Probability and Statistics*, 10:340–355, 2006.
- Kamalika Chaudhuri and Sanjoy Dasgupta. Rates of convergence for nearest neighbor classification. In *Advances in Neural Information Processing Systems*, pages 3437–3445, 2014.
- Jeff Cheeger. Differentiability of Lipschitz functions on metric measure spaces. *Geometric & Functional Analysis GFA*, 9(3):428–517, 1999.
- Luc Devroye. On the almost everywhere convergence of nonparametric regression function estimates. *The Annals of Statistics*, 9(6):1310–1319, 1981a.
- Luc Devroye. On the inequality of Cover and Hart in nearest neighbor discrimination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 75–78, 1981b.
- Luc Devroye, L szl  Gy rffi, and G bor Lugosi. *A probabilistic theory of pattern recognition*. Springer, 1996.
- Lawrence Craig Evans and Ronald F Gariepy. *Measure theory and fine properties of functions*. CRC press, 2015.
- Pierre Fatou. S ries trigonom triques et s ries de Taylor. *Acta mathematica*, 30:335–400, 1906.
- Herbert Federer. *Geometric measure theory*. Springer, 2014.
- Liliana Forzani, Ricardo Fraiman, and Pamela Llop. Consistent nonparametric regression for functional data under the Stone–Besicovitch conditions. *IEEE transactions on information theory*, 58(11):6697–6708, 2012.
- L szl  Gy rffi and Roi Weiss. Universal consistency and rates of convergence of multiclass prototype algorithms in metric spaces. *arXiv preprint arXiv:2010.00636*, 2020.
- Paul R Halmos. *Measure theory*, volume 18. Springer, 2013.
- Steve Hanneke, Aryeh Kontorovich, Sivan Sabato, and Roi Weiss. Universal bayes consistency in metric spaces. *arXiv preprint arXiv:1906.09855*, 2019.
- Susan E Kelly, Mark A Kon, and Louise A Raphael. Pointwise convergence of wavelet expansions. *Bulletin of The American Mathematical Society*, 30(1):87–94, 1994.

- Juha Kinnunen and Visa Latvala. Lebesgue points for Sobolev functions on metric spaces. *Revista matemática iberoamericana*, 18(3):685–700, 2002.
- Juha Kinnunen, Riikka Korte, Nageswari Shanmugalingam, and Heli Tuominen. Lebesgue points and capacities via the boxing inequality in metric spaces. *Indiana university mathematics journal*, pages 401–430, 2008.
- Serge Lang. *Real and functional analysis*, volume 142. Springer Science & Business Media, 2012.
- Henri Lebesgue. Recherches sur la convergence des séries de Fourier. *Mathematische Annalen*, 61(2):251–280, 1905.
- Peter A Loeb. The microscopic behavior of measurable functions. *Nonstandard methods and applications in mathematics*, 25:123, 2006.
- Francesco Maggi. *Sets of finite perimeter and geometric variational problems: an introduction to Geometric Measure Theory*. Cambridge University Press, 2012.
- Pertti Mattila. *Geometry of sets and measures in Euclidean spaces: fractals and rectifiability*. Cambridge university press, 1999.
- David Preiss. Invalid Vitali theorems. *Abstracts. 7th Winter School on Abstract Analysis*, pages 58–60, 1979.
- David Preiss. Dimension of metrics and differentiation of measures. *General topology and its relations to modern analysis and algebra, V (Prague, 1981)*, 3: 565–568, 1983.
- Malempati M Rao and Randall J Swift. *Probability theory with applications*, volume 582. Springer Science & Business Media, 2006.
- David Simmons. Conditional measures and conditional expectation; Rohlin’s disintegration theorem. *Discrete Contin. Dyn. Syst*, 32(7):2565–2582, 2012.
- Elias M Stein and Guido Weiss. *Introduction to Fourier Analysis on Euclidean Spaces*, volume 1. Princeton University Press, 1971.

A Useful Probabilistic Results

In this section we present two useful probability lemmas that we used several times throughout the paper. The first one is needed to avoid relying on conditional probabilities to obtain independence properties (allowing us to state results in non-separable metric spaces). The second one is the classic “freezing lemma”.

Lemma 7. *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $(\mathcal{V}, \mathcal{F}_{\mathcal{V}})$ and $(\mathcal{W}, \mathcal{F}_{\mathcal{W}})$ be two measurable spaces. Let $U: \Omega \rightarrow [0, \infty]$, $f: \mathcal{V} \times \mathcal{W} \rightarrow [0, \infty]$, $V: \Omega \rightarrow \mathcal{V}$, $W: \Omega \rightarrow \mathcal{W}$ be four measurable functions. If (U, V) and W are \mathbb{P} -independent, then*

$$\mathbb{E}[Uf(V, W) | V] = \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V] . \quad (24)$$

Proof. If Z is a random variable, we will denote by $\sigma(Z)$ the σ -algebra generated Z . Since $\mathbb{E}[U | V] \mathbb{E}[f(V, W) | V]$ is a $\sigma(V)$ -measurable and non negative random variable, then, by definition of conditional expectation, we only need to prove that, for all $A \in \mathcal{F}_{\mathcal{V}}$, we have

$$\mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V]\right] = \mathbb{E}\left[\mathbb{I}_A(V) U f(V, W)\right]. \quad (25)$$

Assume first that $U = \mathbb{I}_F$, for some $F \in \mathcal{F}$.

We begin by further assuming that for all $(v, w) \in \mathcal{V} \times \mathcal{W}$, $f(v, w) = \mathbb{I}_B(v) \mathbb{I}_C(w)$, for some $B \in \mathcal{F}_{\mathcal{V}}$ and $C \in \mathcal{F}_{\mathcal{W}}$. For each $A \in \mathcal{F}_{\mathcal{V}}$, we have

$$\begin{aligned} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V]\right] &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F | V] \mathbb{E}[\mathbb{I}_B(V) \mathbb{I}_C(W) | V]\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_B(V) \mathbb{E}[\mathbb{I}_F | V] \mathbb{E}[\mathbb{I}_C(W) | V]\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_B(V) \mathbb{E}[\mathbb{I}_F | V] \mathbb{E}[\mathbb{I}_C(W)]\right] \\ &= \mathbb{E}\left[\mathbb{I}_{A \cap B}(V) \mathbb{E}[\mathbb{I}_F | V]\right] \mathbb{P}(W \in C) \\ &= \mathbb{E}\left[\mathbb{I}_{A \cap B}(V) \mathbb{I}_F\right] \mathbb{P}(W \in C) \\ &= \mathbb{P}(\{V \in A\} \cap \{V \in B\} \cap F) \mathbb{P}(W \in C) \\ &= \mathbb{P}(\{V \in A\} \cap \{V \in B\} \cap F \cap \{W \in C\}) \\ &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F \mathbb{I}_B(V) \mathbb{I}_C(W)\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) U f(V, W)\right]. \end{aligned}$$

This proves (25) under these assumptions.

Then assume that, for all $(v, w) \in \mathcal{V} \times \mathcal{W}$

$$f(v, w) = \sum_{i=1}^n a_i \mathbb{I}_{B_i}(v) \mathbb{I}_{C_i}(w)$$

for some $n \in \mathbb{N}$, $a_1, \dots, a_n > 0$, $B_1, \dots, B_n \in \mathcal{F}_{\mathcal{V}}$, and $C_1, \dots, C_n \in \mathcal{F}_{\mathcal{W}}$. For each $A \in \mathcal{F}_{\mathcal{V}}$, we have

$$\begin{aligned} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V]\right] &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F | V] \mathbb{E}\left[\sum_{i=1}^n a_i \mathbb{I}_{B_i}(V) \mathbb{I}_{C_i}(W) | V\right]\right] \\ &= \sum_{i=1}^n a_i \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F | V] \mathbb{E}[\mathbb{I}_{B_i}(V) \mathbb{I}_{C_i}(W) | V]\right] \\ &= \sum_{i=1}^n a_i \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F \mathbb{I}_{B_i}(V) \mathbb{I}_{C_i}(W)\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F \sum_{i=1}^n a_i \mathbb{I}_{B_i}(V) \mathbb{I}_{C_i}(W)\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) U f(V, W)\right], \end{aligned}$$

where the third equality follows by (25), which is true in this case for what we proved above. This proves (25) under these assumptions.

Next, assume that $f = \mathbb{I}_D$, where D belongs to the product σ -algebra $\mathcal{F}_V \otimes \mathcal{F}_W$. Let \mathcal{G} be the algebra generated by $\Pi := \{B \times C \mid B \in \mathcal{F}_V, C \in \mathcal{F}_W\}$. By [Lang, 2012, Theorem 6.3, Chapter 6 (The general integral/Approximations)], there exists a sequence $(f_n)_{n \in \mathbb{N}}$ such that for all $n \in \mathbb{N}$, there exist $m_n \in \mathbb{N}$, $a_{1,n}, \dots, a_{m_n,n} > 0$, $G_{1,n}, \dots, G_{m_n,n} \in \mathcal{G}$ such that $f_n = \sum_{k=1}^{m_n} a_{k,n} \mathbb{I}_{G_{k,n}}$ and $\|f - f_n\|_{\mathcal{L}^1(\mathbb{P}(V,W))} \rightarrow 0$, as $n \rightarrow \infty$, i.e. $\mathbb{E}\left[|f(V,W) - f_n(V,W)|\right] \rightarrow 0$, as $n \rightarrow \infty$. Since Π is a π -system, we have that the elements of \mathcal{G} are finite unions of disjoint elements of Π , and so for each $n \in \mathbb{N}$ and each $k \in \{1, \dots, m_n\}$ there exist $l_{n,k} \in \mathbb{N}$ and $B_{1,n,k} \times C_{1,n,k}, \dots, B_{l_{n,k},n,k} \times C_{l_{n,k},n,k} \in \Pi$ mutually disjoint such that $G_{k,n} = \bigcup_{j=1}^{l_{n,k}} B_{j,n,k} \times C_{j,n,k}$. Therefore for each $n \in \mathbb{N}$

$$f_n = \sum_{k=1}^{m_n} a_{k,n} \mathbb{I}_{G_{k,n}} = \sum_{k=1}^{m_n} a_{k,n} \mathbb{I}_{\bigcup_{j=1}^{l_{n,k}} B_{j,n,k} \times C_{j,n,k}} = \sum_{k=1}^{m_n} \sum_{j=1}^{l_{n,k}} a_{k,n} \mathbb{I}_{B_{j,n,k} \times C_{j,n,k}} .$$

Then, $\mathbb{E}\left[\left|\mathbb{E}[f_n(V,W) \mid V] - \mathbb{E}[f(V,W) \mid V]\right|\right] \rightarrow 0$, as $n \rightarrow \infty$. Similarly, since \mathbb{I}_F is bounded, $\mathbb{E}\left[\left|\mathbb{E}[\mathbb{I}_F f_n(V,W) \mid V] - \mathbb{E}[\mathbb{I}_F f(V,W) \mid V]\right|\right] \rightarrow 0$ as $n \rightarrow \infty$. Moreover, being \mathbb{I}_F bounded, its conditional expectation $\mathbb{E}[\mathbb{I}_F \mid V]$ is also bounded, which in turn yields $\mathbb{E}\left[\left|\mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}[f_n(V,W) \mid V] - \mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}[f(V,W) \mid V]\right|\right] \rightarrow 0$ as $n \rightarrow \infty$. Thus, for each $A \in \mathcal{F}_V$,

$$\begin{aligned} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U \mid V] \mathbb{E}[f(V,W) \mid V]\right] &= \lim_{n \rightarrow \infty} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}[f_n(V,W) \mid V]\right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F f_n(V,W)\right] \\ &= \mathbb{E}\left[\mathbb{I}_A(V) U f(V,W)\right] , \end{aligned}$$

where the third equality follows by (25), which is true in this case for what we proved above. This proves (25) under these assumptions.

Let now $f = \sum_{i=1}^n a_i \mathbb{I}_{D_i}$, for some $n \in \mathbb{N}$, $a_1, \dots, a_n > 0$, and $D_1, \dots, D_n \in$

$\mathcal{F}_V \otimes \mathcal{F}_W$. Then, for each $A \in \mathcal{F}_V$,

$$\begin{aligned}
& \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U \mid V] \mathbb{E}[f(V, W) \mid V]\right] \\
&= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}\left[\sum_{i=1}^n a_i \mathbb{I}_{D_i}(V, W) \mid V\right]\right] \\
&= \sum_{i=1}^n a_i \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}[\mathbb{I}_{D_i}(V, W) \mid V]\right] \\
&= \sum_{i=1}^n a_i \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F \mathbb{I}_{D_i}(V, W)\right] \\
&= \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F \sum_{i=1}^n a_i \mathbb{I}_{D_i}(V, W)\right] \\
&= \mathbb{E}\left[\mathbb{I}_A(V) U f(V, W)\right],
\end{aligned}$$

where the third equality follows by (25), which is true in this case for what we proved above. This proves (25) under these assumptions.

Now, if f is general, we can get a sequence $(f_n)_{n \in \mathbb{N}}$ such that for each $n \in \mathbb{N}$ there exists $m_n \in \mathbb{N}$, $a_{1,n}, \dots, a_{m_n,n} > 0$, and $D_{1,n}, \dots, D_{m_n,n} \in \mathcal{F}_V \otimes \mathcal{F}_W$ such that for each $n \in \mathbb{N}$ we have that $f_n = \sum_{k=1}^{m_n} a_{k,n} \mathbb{I}_{D_{k,n}}$ and $f_n \uparrow f$ pointwise, as $n \uparrow \infty$. Hence, by the monotone convergence theorem for the conditional expectation, we have that, for each $A \in \mathcal{F}_V$,

$$\begin{aligned}
\mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[U \mid V] \mathbb{E}[f(V, W) \mid V]\right] &= \lim_{n \rightarrow \infty} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_F \mid V] \mathbb{E}[f_n(V, W) \mid V]\right] \\
&= \lim_{n \rightarrow \infty} \mathbb{E}\left[\mathbb{I}_A(V) \mathbb{I}_F f_n(V, W)\right] \\
&= \mathbb{E}\left[\mathbb{I}_A(V) U f(V, W)\right].
\end{aligned}$$

Now, suppose $U = \sum_{i=1}^n a_n \mathbb{I}_{F_i}$ for some $n \in \mathbb{N}$, for n distinct $a_1, \dots, a_n > 0$ and $F_1, \dots, F_n \in \mathcal{F}$. For each $i \in \{1, \dots, n\}$, we have that $F_i = \{U = a_i\}$ so $\sigma(\mathbb{I}_{F_i}, V) \subset \sigma(U, V)$, and since (U, V) is \mathbb{P} -independent from W we also have that

(\mathbb{I}_{F_i}, V) is \mathbb{P} -independent from W . Then, for each $A \in \mathcal{F}_{\mathcal{V}}$, we have that

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{I}_A(V) \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V] \right] \\
&= \mathbb{E} \left[\mathbb{I}_A(V) \mathbb{E} \left[\sum_{i=1}^n a_i \mathbb{I}_{F_i} | V \right] \mathbb{E}[f(V, W) | V] \right] \\
&= \sum_{i=1}^n a_i \mathbb{E} \left[\mathbb{I}_A(V) \mathbb{E}[\mathbb{I}_{F_i} | V] \mathbb{E}[f(V, W) | V] \right] \\
&= \sum_{i=1}^n a_i \mathbb{E} \left[\mathbb{I}_A(V) \mathbb{I}_{F_i} f(V, W) \right] \\
&= \mathbb{E} \left[\mathbb{I}_A(V) \sum_{i=1}^n a_i \mathbb{I}_{F_i} f(V, W) \right] \\
&= \mathbb{E} \left[\mathbb{I}_A(V) U f(V, W) \right],
\end{aligned}$$

where in the third equality we used the previous case. Finally, if U is general, for each $n \in \mathbb{N}$ and $i \in \{1, \dots, 2^{2n-1} - 1\}$ define $a_{i,n} = \frac{i2^n}{2^{2n-1}}$ and $F_{i,n} = \left\{ U \in \left[\frac{i2^n}{2^{2n-1}}, \frac{(i+1)2^n}{2^{2n-1}} \right) \right\}$. For each $n \in \mathbb{N}$, define

$$U_n = \sum_{i=1}^{2^{2n-1}} a_{i,n} \mathbb{I}_{F_{i,n}}.$$

Then, for each $n \in \mathbb{N}$, we have that $a_{1,n}, \dots, a_{2^{2n-1},n} > 0$ are distinct, that $F_{1,n}, \dots, F_{2^{2n-1},n} \in \mathcal{F}$ are mutually disjoint. Also, for each $n \in \mathbb{N}$, we have that $\sigma(U_n) \subset \sigma(\mathbb{I}_{F_{1,n}}, \dots, \mathbb{I}_{F_{2^{2n-1},n}}) \subset \sigma(U)$ and so (U_n, V) is \mathbb{P} -independent from W since (U, V) is \mathbb{P} -independent from W . Furthermore, we have that $U_n \uparrow U$ pointwise as $n \uparrow \infty$ and so, since $f(V, W) \geq 0$, also $U_n f(V, W) \uparrow U f(V, W)$ pointwise as $n \uparrow \infty$ and $\mathbb{E}[U_n | V] \mathbb{E}[f(V, W) | V] \uparrow \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V]$, \mathbb{P} -almost everywhere as $n \uparrow \infty$. Hence, by what we observed and the monotone convergence theorem, we have that, for each $A \in \mathcal{F}_{\mathcal{V}}$,

$$\begin{aligned}
\mathbb{E} \left[\mathbb{I}_A(V) \mathbb{E}[U | V] \mathbb{E}[f(V, W) | V] \right] &= \lim_{n \rightarrow \infty} \mathbb{E} \left[\mathbb{I}_A(V) \mathbb{E}[U_n | V] \mathbb{E}[f(V, W) | V] \right] \\
&= \lim_{n \rightarrow \infty} \mathbb{E} \left[\mathbb{I}_A(V) U_n f(V, W) \right] \\
&= \mathbb{E} \left[\mathbb{I}_A(V) U f(V, W) \right].
\end{aligned}$$

where in the second equality we used the previous case. This concludes the proof. \square

The next result can be proven with the same approach as the previous lemma. Alternatively, a proof is given in [Baldi, 2017, Lemma 4.1].

Lemma 8 (The “freezing lemma”). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $(\mathcal{V}, \mathcal{F}_{\mathcal{V}})$ and $(\mathcal{W}, \mathcal{F}_{\mathcal{W}})$ be two measurable spaces. Let $f: \mathcal{V} \times \mathcal{W} \rightarrow [0, \infty]$, $V: \Omega \rightarrow \mathcal{V}$, $W: \Omega \rightarrow \mathcal{W}$ be three measurable functions. If V and W are \mathbb{P} -independent, then*

$$\mathbb{E}[f(V, W) \mid V] = \left[\mathbb{E}[f(v, W)] \right]_{v=V} \quad (26)$$

\mathbb{P} -almost surely, where the right hand side is the composition

$$\left[\mathbb{E}[f(v, W)] \right]_{v=V} = \left(v \mapsto \mathbb{E}[f(v, W)] \right) \circ V .$$