Electronic Journal of Statistics
Vol. 7 (2013) 1716–1746
ISSN: 1935-7524
DOI: 10.1214/13-EJS822

Fast rates for empirical vector quantization

Clément Levrard

Université Paris Sud 11 and Paris 6
e-mail: clement.levrard@math.u-psud.fr

Abstract: We consider the rate of convergence of the expected loss of empirically optimal vector quantizers. Earlier results show that the mean-squared expected distortion for any fixed probability distribution supported on a bounded set and satisfying some regularity conditions decreases at the rate $\mathcal{O}(\log n/n)$. We prove that this rate is actually $\mathcal{O}(1/n)$. Although these conditions are hard to check, we show that well-clustered distributions with continuous densities supported on a bounded set are included in the scope of this result.

Keywords and phrases: Quantization, clustering, localization, fast rates.

Received January 2012.

1. Introduction

Empirical vector quantizer design is a way to answer the problem of identifying groupings of similar points that are relatively away from one another, or, in other words, to partition the data into dissimilar groups of similar items. For a comprehensive introduction to this topic, the reader is referred to the monograph of Graf and Luschgy [13]. To isolate meaningful groups from a cloud of data is a topic of interest in many fields, from social science to biology. In fact this issue originates in the theory of signal processing in the late 40's, known as the quantization issue, or lossy data compression (a good introduction to this field can be found in the book of Gersho and Gray [11]).

To be more precise, let P denote a probability distribution over the Euclidean space $\mathbb{R}^d$. A k-point quantizer Q, also called k-level quantizer in the case where $d = 1$, is a map from $\mathbb{R}^d$ to $\mathbb{R}^d$, whose image set is made of exactly k points, that is $\left|Q(\mathbb{R}^d)\right| = k$. By considering the preimages of these points, such a map partitions the whole space into k groups, and assigns each group a representative.

For any P-integrable function $f : \mathbb{R}^d \longrightarrow \mathbb{R}$, we will denote by Pf the integral of f with respect to P. To measure how well a quantizer Q performs in representing the source distribution P, we introduce the distortion

$$R(Q) := P\|x - Q(x)\|^2,$$

when $P\|x\|^2 < \infty$. This choice of distortion function is convenient, since it takes advantage of the underlying Euclidean structure. Note however that several authors such as Graf and Luschgy in [13] or Fischer in [10] deal with more general distortion functions.

For a k-point quantizer Q with images $Q(\mathbb{R}^d) = \{c_1, \dots, c_k\}$, we will call code points of Q the points $c_1, \dots, c_k$, and codebook of Q an arbitrary vector $\mathbf{c}$ of $(\mathbb{R}^d)^k$, the components of which are the code points, e.g. $\mathbf{c} := (c_1, \dots, c_k)$. Without loss of generality we restrict our attention to nearest neighbor quantizers, namely quantizers satisfying the condition $\|x - Q(x)\| = \min_{c \in Q(\mathbb{R}^d)} \|x - c\|$. Intuitively, a nearest neighbor quantizer sends any vector x to the nearest code point c_i to x. Note that a nearest neighbor quantizer is determined by its codebook $\mathbf{c}$ with ties arbitrarily broken. Since we only deal with continuous distributions, how ties are broken will not matter.

Throughout this paper, quantizers will be represented by their codebook. This choice will allow us to handle vectors rather than maps, in order to take advantage of the underlying Euclidean structure. From this point of view, the distortion function takes the form

$$R(\mathbf{c}) := P\|x - Q(x)\|^2 = P \min_{j=1,\dots,k} \|x - c_j\|^2,$$

when $P\|x\|^2 < \infty$.

Let $X_1, \dots, X_n$ be a independent and identically distributed sample with distribution P. The goal here is to find a codebook $\hat{\mathbf{c}}_n$, drawn from the data $X_1, \dots, X_n$, whose distortion is as close as possible to the optimal distortion $R^* := \inf_{\mathbf{c} \in (\mathbb{R}^d)^k} R(\mathbf{c})$. To solve the problem, most approaches to date attempt to implement the principle of empirical error minimization in the vector quantization context. According to this principle, good code points can be found by searching for ones that minimize the empirical distortion over the training data, defined by

$$\hat{R}_n(\mathbf{c}) := \frac{1}{n} \sum_{i=1}^{n} (X_i - Q(X_i))^2 = \frac{1}{n} \sum_{i=1}^{n} \min_{j=1,\dots,k} \|X_i - c_j\|^2.$$

The existence of such empirically optimal codebooks has been formally established by Graf and Luschgy [13, Theorem 4.12], following the approach of Pollard in the proof of [24, Lemma 8]. Denote by $\hat{\mathbf{c}}_n$ one of these empirically optimal codebooks. If the training data represents the source well, $\hat{\mathbf{c}}_n$ will hopefully also perform near optimally on the real source. Roughly, this means that we expect $R(\hat{\mathbf{c}}_n) \approx R^*$. The problem of quantifying how good empirically designed codebooks are, compared to the truly optimal ones, has been extensively studied, see for instance Linder [16].

To reach the latter goal, a standard route is to exploit the Wasserstein distance between the empirical distribution and the source distribution, to derive upper bounds on the average distortion of empirically optimal codebooks. Following this approach, Pollard [21] proved that, if $P\|x\|^2 < \infty$, then $R(\hat{\mathbf{c}}_n) - R^* \longrightarrow 0$ almost surely, as $n \to \infty$. Using techniques borrowed from statistical learning theory, Linder, Lugosi and Zeger [17], and Biau, Devroye and Lugosi [5] showed that if the support of P is bounded, then $\mathbb{E}\left(R(\hat{\mathbf{c}}_n) - R^*\right) = \mathcal{O}(1/\sqrt{n})$, where the expectation is taken over the training sample $X_1, \dots, X_n$. Bartlett, Linder and Lugosi established in [4] that this rate is minimax over distributions supported on a finite set of points. More recently, Antos [1] improved the numerical constants mentioned in this minimax result, and also proved that the minimax rate over distributions over bounded sets with continuous densities is still $1/\sqrt{n}$.

However, faster rates individual convergence rates can be achieved, under certain conditions. For example, it was shown by Chou [9], following a result of Pollard [23], that if the source distribution satisfies some regularity conditions, then $R(\hat{\mathbf{c}}_n) - R^* = \mathcal{O}_{\mathbb{P}}(1/n)$, where we recall that a sequence of random variables $Y_n = \mathcal{O}_{\mathbb{P}}(1/n)$ if, for all positive real number M, $\mathbb{P}(n|Y_n| \geq M) \to 0$ as $n \to \infty$. Nevertheless, this consistency result does not provide any information on how many training samples are needed to ensure that the average distortion of empirically optimal codebooks is close to the optimum. Antos, Györfi and György established in [2] that $\mathbb{E}(R(\hat{\mathbf{c}}_n) - R^*) = \mathcal{O}(\log n/n)$, under other conditions, paying a $\log n$ factor to derive a non-asymptotic bound. It is worth pointing out that the conditions cannot be checked in practice, and consequently remain of theoretical nature. Moreover, the rate of $\mathcal{O}(1/n)$ for the average distortion can be achieved when the source distribution is supported on a finite set of points (see Antos, Györfi and György [2]). Consequently, an open question is to know whether this optimal rate can be attained for more general distributions, and under what set of conditions.

In the present paper, we improve previous results of Antos, Györfi and György [2], by getting rid of the $\log n$ factor, adding some minor regularity conditions on P. To this aim we use statistical learning arguments, and prove that the average distortion of empirically optimal codebooks decreases at the rate $\mathcal{O}(1/n)$, under certain conditions. To get this result we use techniques such as the localization principle borrowed from Blanchard, Bousquet and Massart [6] or Koltchinskii [15]. The condition we offer can be easily interpreted as a margin-type condition, similar to the ones of Massart and Nedelec in [19], showing a clear connection between statistical learning theory and vector quantization.

Furthermore, we offer equivalences between different sets of regularity conditions which guarantee that the distortion of the empirically optimal codebook decreases at a fast rate. More precisely, we prove that conditions Pollard required in [23], conditions Antos, Györfi and György required in [2], and conditions we required, are equivalent, in the case where P has a continuous density. It is worth pointing out that all conditions mentioned above are of theoretical nature, and remain hard to understand. We also give in this paper a more reader-friendly sufficient condition.

The paper is organized as follows. In Section 2 we introduce notation and definitions of interest. In Section 3 we offer our main results. These results are discussed in Section 4, and illustrated on examples such as Gaussian mixtures or quasi-finite distributions. Finally, proofs are gathered in Section 5.

2. The quantization problem

Throughout the paper, $X_1, \ldots, X_n$ is a sequence of independent $\mathbb{R}^d$-valued random variables with distribution P. To frame the quantization problem as a statistical learning one, we first have to consider quantization as a contrast minimization issue. To this aim we introduce the following notation. Consider a nearest neighbor quantizer with codebook $\mathbf{c} = (c_1, \ldots, c_n)$. The contrast function γ is defined as

$$\gamma : \begin{cases} \left(\mathbb{R}^d\right)^k \times \mathbb{R}^d & \longrightarrow & \mathbb{R} \\ (\mathbf{c}, x) & \longmapsto & \min_{j=1,\ldots,k} \|x - c_j\|^2 \end{cases} .$$

Within this framework, the risk $R(\mathbf{c})$ takes the form $R(Q) = R(\mathbf{c}) = P\gamma(\mathbf{c}, .)$, where $Pf(.)$ means integration of the function f with respect to P. Denote by P_n the empirical distribution that is induced on $\mathbb{R}^d$ by the n-sample $X_1, \ldots, X_n$, namely, for any measurable subset $A \subseteq \mathbb{R}^d$, $P_n(A) = |\{i | X_i \in A\}|$. Once P_n introduced, the empirical risk $\hat{R}_n(\mathbf{c})$ can be expressed as $P_n\gamma(\mathbf{c}, .)$. Remark that an optimal $\mathbf{c}^*$ minimizes $P\gamma(\mathbf{c}, .)$, whereas $\hat{\mathbf{c}}_n \in \arg\min_{\mathbf{c} \in (\mathbb{R}^d)^k} P_n\gamma(\mathbf{c}, .)$. It is worth pointing out that, if $P\|x\|^2 < \infty$, then the existence of both $\hat{\mathbf{c}}_n$ and $\mathbf{c}^*$ are guaranteed by Graf and Luschgy [13, Theorem 4.12].

Let $\mathcal{M}$ denote the set of optimal codebooks, and let $\mathbf{c}^* \in \mathcal{M}$ be an optimal codebook, with code points c_i^*. The Voronoi cell V_i^*, also called quantization cell, is defined as the subset of $\mathbb{R}^d$ made of the points which are closer to c_i^* than any other c_j^*, i.e.

$$V_i^* := \left\{x \in \mathbb{R}^d | \quad \forall j \neq i \quad \|x - c_i^*\| \leq \|x - c_j^*\|\right\}.$$

It may be noted that many authors prefer to define the Voronoi cell V_i^* as the open set

$$\left\{x \in \mathbb{R}^d | \quad \forall j \neq i \quad \|x - c_i^*\| < \|x - c_j^*\|\right\}.$$

However, this choice of convention will not matter, since every boundary of an optimal Voronoi cell has zero P-measure, that is

$$P\left(\left\{x \in \mathbb{R}^d | \|x - c_i^*\| = \|x - c_j^*\|\right\}\right) = 0$$

for $i \neq j$ (see [13, Theorem 4.2]).

It is well known that any optimal codebook satisfies the centroid condition (see, e.g., [11, Section 6.2]), which states that each optimal code point is chosen to minimize the distortion over its associated cell, or, in other words

$$P\left(\|x - c_i^*\|^2 \mathbf{1}_{x \in V_i^*}\right) = \inf_{c \in \mathbb{R}^d} P\left(\|x - c\|^2 \mathbf{1}_{x \in V_i^*}\right),$$

where, for any measurable subset $A \subset \mathbb{R}^d$, $\mathbf{1}_A$ denotes the indicator function of the set A. An immediate consequence of the centroid condition is that every c_i* satisfies

$$c_i^* = \frac{P\left(x \mathbf{1}_{V_i^*}\right)}{p_i^*},$$

where $p_i^* := P(V_i^*)$ is nonzero, according to [13, Theorem 4.1]. In the case where P has a density, it is proved in [23, Lemma A] that $\mathbf{c} \longmapsto P\gamma(\mathbf{c}, .)$ is differentiable. In this case, it is easy to show that the centroid condition takes the form

$$\nabla P\gamma(\mathbf{c}^*, .) = 0,$$

where ∇f denotes the differential of f for any differentiable map $f : \left(\mathbb{R}^d\right)^k \longrightarrow \mathbb{R}$ (see, for instance, [11, Section 6.2]).

Let $\mathbf{c} \in \left(\mathbb{R}^d\right)^k$ be a $k \times d$ vector, and let $\mathbf{c}^* \in \mathcal{M}$ be an optimal codebook. We introduce the loss, or distortion redundancy, to compare the performance of $\mathbf{c}$ and $\mathbf{c}^*$, namely

$$\ell(\mathbf{c}, \mathbf{c}^*) := R(\mathbf{c}) - R(\mathbf{c}^*) = P\left(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .)\right).$$

Throughout the paper we will use the following assumptions on the source distribution P. For $c \in \mathbb{R}^d$ and any positive real number $M > 0$, let $\mathcal{B}(c, M)$ denote the closed ball of radius M and center c in $\mathbb{R}^d$. To be precise

$$\mathcal{B}(c, M) = \left\{x \in \mathbb{R}^d | \quad \|x - c\| \leq M\right\}.$$

Assumption 1 (**Peak Power Constraint**). *The distribution P is such that $P(\mathcal{B}(0, 1)) = 1$,*

For convenience we only consider distributions whose support is included in $\mathcal{B}(0, 1)$. However, it is important to note that our results hold for distributions whose support is included in $\mathcal{B}(0, M)$, for an arbitrary M. In fact, a distribution supported on $\mathcal{B}(0, M)$ can easily be turned into a distribution supported on $\mathcal{B}(0, 1)$, via an homothetic transformation. Therefore, we will only state results for distributions for which the support is included in $\mathcal{B}(0, 1)$. Note that Assumption 1 is stronger than the requirement $P\|x\|^2 < \infty$, as it imposes that P is supported on a bounded subset of $\mathbb{R}^d$. However, it is likely that our results can be extended to the case where just the weaker assumption $P\|x\|^2 < \infty$ is required, using techniques such as in [20] or [8].

Pollard introduced in [23] the following regularity requirement, which was initially used to derive an asymptotic rate of convergence for the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$.

Assumption 2 (**Pollard's regularity condition**). *The distribution P satisfies the following two conditions:*

1. *P has a continuous density f with respect to the Lebesgue measure on $\mathbb{R}^d$,*
2. *The Hessian matrix of $\mathbf{c} \longmapsto P\gamma(\mathbf{c}, .)$ is positive definite for all optimal codebooks $\mathbf{c}^*$.*

It may be noted that Condition 1 of Assumption 2 does not guarantee the existence of a second derivative for the function $\mathbf{c} \longmapsto P\gamma(\mathbf{c}, .)$. Nevertheless, if Assumption 1 and Condition 1 of Assumption 2 are satisfied, then it can be proved that $\mathbf{c} \longmapsto P\gamma(\mathbf{c}, .)$ is twice differentiable (see, e.g., [23, Lemma C]).

Let V_i be the Voronoi cell associated with c_i, for $i = 1, \ldots, k$. The Hessian matrix is composed of the following $d \times d$ blocks:

$$H(\mathbf{c})_{i,j} = \begin{cases} 2P(V_i) - 2\sum_{\ell \neq i} r_{i\ell}^{-1} \sigma\left[f(x)(x - c_i)(x - c_i)^t \mathbf{1}_{\partial(V_i \cap V_\ell)}\right] & \text{for} \quad i = j \\ 2{r_{ij}}^{-1} \sigma\left[f(x)(x - c_i)(x - c_j)^t \mathbf{1}_{\partial(V_i \cap V_j)}\right] & \text{for} \quad i \neq j \end{cases}, \tag{1}$$

where $r_{ij} = \|c_i - c_j\|$, $\partial(V_i \cap V_j)$ denotes the possibly empty common face of V_i and V_j, and σ means integration with respect to the $(d-1)$-dimensional Lebesgue measure. For a proof of that statement, we refer to Pollard [23].

Assumption 2 is hard to check in general. However, there are some cases where it can be proved that $H(\mathbf{c}^*)$ is positive definite for every optimal codebook $\mathbf{c}^*$. For example, Antos, Györfi and György proved in [2, Corollary 2] that, if $d = 1$ and P has a strictly log-concave density, then P satisfies Assumption 2. As will be shown in Corollary 3.1, this is also the case when the density is small enough at the boundaries of optimal Voronoi cells.

When Assumption 1 and Assumption 2 are satisfied, Chou [9] proved that $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) = \mathcal{O}_{\mathbb{P}}(1/n)$. This result relies on the previous result of Pollard [23], who established the asymptotic normality of $\sqrt{n}(\hat{\mathbf{c}}_n - \mathbf{c}^*)$. To get this asymptotic result, Pollard used conditions under which the distortion and the Euclidean distance are connected, and used chaining arguments to bound from above a term which looks like a Rademacher complexity, constrained on an area around an optimal codebook. Note that Koltchinskii [15] used a similar method to apply the localization principle.

The following Assumption 3 is the assumption we require to obtain our main result. It demands direct connections between the Euclidean distance, the loss $\ell(\mathbf{c}, \mathbf{c}^*)$ and the variance $\mathrm{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .))$, taken with respect to P.

Assumption 3. *The distribution P satisfies the following two technical conditions:*

(H1) $$\exists A_1 > 0 \quad \forall \mathbf{c} \in \mathcal{B}(0,1) \quad \ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c})) \geq A_1 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2,$$

where $\mathbf{c}^*(\mathbf{c}) \in \underset{\mathbf{c}^* \in \mathcal{M}}{\arg\min} \|\mathbf{c} - \mathbf{c}^*\|$, *and*

(H2)
$$\exists A_2 > 0 \quad \forall \mathbf{c} \in \mathcal{B}(0,1) \quad \forall \mathbf{c}^* \in \mathcal{M} \quad \mathrm{Var}_P\left(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .)\right) \leq A_2 \|\mathbf{c} - \mathbf{c}^*\|^2.$$

Notice that, contrary to Assumption 2, Assumption 3 does not require P to have a continuous density.

When considering several optimal codebooks $\mathbf{c}^*$ to be compared to a general $\mathbf{c}$, it is natural to choose one among the closest to $\mathbf{c}$, hence the choice of $\mathbf{c}^*(\mathbf{c})$. Furthermore, since for all $\mathbf{c}^*$ in $\mathcal{M}$, $\ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c})) = \ell(\mathbf{c}, \mathbf{c}^*)$, we will write $\ell(\mathbf{c}, \mathbf{c}^*)$ without specifying which $\mathbf{c}^* \in \mathcal{M}$ is at stake.

It is worth pointing out that Antos, Györfi and György [2, Corollary 1] proved that, if Assumption 1 is satisfied, then Assumption 2 implies Assumption 3. In the same paper, these authors only require that P satisfies the following Assumption 4.

Assumption 4 (**Condition of Antos, Györfi and György**). *There exists $A > 0$ such that*

$$\forall \mathbf{c} \in \mathcal{B}(0,1) \quad \mathrm{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .)) \leq A\ell(\mathbf{c}, \mathbf{c}^*).$$

Assumption 4 is at first sight weaker than Assumption 3, since it only requires a comparison between $\ell(\mathbf{c}, \mathbf{c}^*)$ and $\mathrm{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .))$, without comparing them to the intermediate $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$.

Antos, Györfi and György proved in [2, Theorem 2] that if P satisfies Assumption 1 and Assumption 4, then $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) = \mathcal{O}(\log(n)/n)$. As explained before the statement of Assumption 4, it has been proved by Antos, Györfi and György in [2, Corollary 1] that, provided that Assumption 1 is satisfied, Assumption 2 implies Assumption 4. Consequently, if P denotes a distribution satisfying Assumption 1 and Assumption 2, the result of Chou [9] states that $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ converges in probability to 0 at the rate $1/n$, whereas the result of Antos, Györfi and György [2] indicates that $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ converges to 0 at the rate $\log(n)/n$ in expectation. Therefore, a question of interest is to know whether these two rates are truly different, or whether the $\log(n)$ factor is artificial.

To be more precise, Antos, Györfi and György used in [2] a concentration inequality based on the fact that the variance and the expectation of the distortion are connected to get their result. Interestingly, this point of view has been developed by Blanchard, Bousquet and Massart [6] to get bounds on the classification risk of the SVM, using the localization principle. That is the approach that will be followed in the present paper.

3. Main results

The conditions we require to obtain our main result differ from the conditions Pollard required in [23], and from those Antos, Györfi and György proposed in [2, Theorem 2]. Consequently it is natural to make connections between these different sets of conditions clear. This is the aim of the following proposition.

Proposition 3.1. *Let P be a distribution on $\mathcal{B}(0,1)$. Then the following sets of conditions on P are equivalent:*

$$\left\{ \begin{array}{c} \textit{Assumption 1} \\ \textit{Assumption 2} \end{array} \right. \Leftrightarrow \left\{ \begin{array}{c} P \textit{ has a continuous density} \\ \mathcal{M} \textit{ is finite} \\ \textit{Assumption 1} \\ \textit{Assumption 3} \end{array} \right. ,$$

and

$$\left\{ \begin{array}{c} \textit{Assumption 1} \\ \textit{Assumption 2} \end{array} \right. \Leftrightarrow \left\{ \begin{array}{c} P \textit{ has a continuous density} \\ \mathcal{M} \textit{ is finite} \\ \textit{Assumption 1} \\ \textit{Assumption 4} \end{array} \right. .$$

Roughly, Proposition 3.1 states that, provided that P has a density which is continuous and whose support is bounded, all the conditions which ensure fast rates of convergence for the average distortion of the empirically optimal quantizer are equivalent. The following theorem offers a new bound on the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$, when P satisfies any of the three sets of conditions proposed in Proposition 3.1.

Theorem 3.1. *Suppose that P has a density f and $\mathcal{M}$ is a finite set. Assume that Assumption 1 and Assumption 3 are satisfied. Then, denoting by $\hat{\mathbf{c}}_n$ an empirical risk minimizer, we have*

$$\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq \frac{C_0}{n},$$

where C_0 is a positive constant depending on P, k and d.

This result shows that a convergence rate of $1/n$ can be achieved in expectation, at the price of a few more conditions on the source distribution than those Antos, Györfi and György [2, Theorem 2] required. To be more precise, their result only requires that Assumption 1 and Assumption 4 are satisfied. According to Proposition 3.1, provided that P has a continuous density and $\mathcal{M}$ is finite, the set of conditions required by Antos, Györfi and György in [2] turns out to be equivalent to the set of conditions Pollard requires in [23] or to the the set of conditions mentioned in Theorem 3.1.

As illustrated by the proof in Section 5.3, the constant C_0 mentioned in Theorem 3.1 strongly depends on constants A_1 and A_2 introduced in Assumption 3. Consequently, to understand how C_0 depends on k, d or P, the exact dependency of A_1 and A_2 on P, k and d has to be known. Unfortunately, the existence of such an A_1 often derives from compactness arguments. Thus we are not able in this paper to explain how C_0 depends on the other parameters k, d and P.

The technical result, from which derives Theorem 3.1, is presented in Section 5.3, Theorem 5.1. It is based on a version of Talagrand's inequality due to Bousquet [7] and its application to localization, following the approach of Massart and Nedelec [19]. It is important to note that drawing connections between the Euclidean distance $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|$ and the loss $\ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c}))$ is essential in the proof of Theorem 3.1, as it allows us to use chaining arguments as in [23].

The conditions required in Theorem 3.1 remain hard to check, and cannot be easily interpreted. In fact, Assumption 3 and Assumption 4 demands that the distribution P is such that a technical inequality is satisfied for every $\mathbf{c}$ in $\mathcal{B}(0, M)$, which cannot be checked in practice. Assumption 2 involves second derivatives of the distortion. Consequently, checking Assumption 2, even theoretically, remains a hard issue. Theorem 3.2 below offers a more interpretable condition regarding the L_∞-norm of the density f on the boundaries of optimal Voronoi cells, for the distribution P to satisfy Assumption 2. We recall that $\mathcal{M}$ denotes the set of all possible optimal codebooks $\mathbf{c}^*$.

Theorem 3.2. *Suppose that Assumption 1 is satisfied, $\mathcal{M}$ is finite, and P has a continuous density f. For an optimal codebook $\mathbf{c}^*$, denote by V_i^* the optimal Voronoi cell associated with the code point c_i^*. Let $N^* = \bigcup_{\mathbf{c}^* \in \mathcal{M}, i \neq j} \partial(V_i^* \cap V_j^*)$ denote the union of all possible boundaries of optimal Voronoi cells with respect to all possible optimal codebooks $\mathbf{c}^*$, and denote by Γ the Gamma function. At last, let $f_{|N^*}$ denote the restriction of the function f to the subset N^*, and define $B = \inf_{\mathbf{c}^* \in \mathcal{M}, i \neq j} \|c_i^* - c_j^*\|$.*

Suppose that

$$\|f_{|N^*}\|_\infty < \frac{\Gamma\left(\frac{d}{2}\right)B}{2^{d+3}\pi^{d/2}} \inf_{\mathbf{c}^* \in \mathcal{M}, i=1,\ldots,k} P(V_i^*).$$

Then P satisfies Assumption 2.

The proof is given in Section 5.6. Remark that, for general distributions supported on $\mathcal{B}(0,M)$, we can state a similar theorem, involving M^{d+1} in the right-hand side of the inequality in Theorem 3.2. Combining Theorem 3.2, Theorem 3.1, and the connections between different sets of conditions leads to the following corollary.

Corollary 3.1. *Suppose that Assumption 1 is satisfied and P has a continuous density. Then there exists an explicit constant $\kappa > 0$, depending only on k and d, such that, if $\|f_{|N^*}\|_\infty < \kappa$, then*

$$\mathbb{E}(\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)) = \mathcal{O}\left(\frac{1}{n}\right).$$

This corollary emphasizes the idea that, if P is well concentrated around its optimal code points, then some localization conditions can be satisfied and therefore it is a favorable case. The intuition behind this result is given by the extremal case where optimal Voronoi cells boundaries are empty with respect to P. This case is described in detail in Section 4. Moreover, the notion of a well-clustered distribution looks like margin-type conditions for the classification case, as described by Massart and Nedelec [19]. This confirms the intuition of an easy-to-quantize distribution, when the poles are well-separated.

Since the location of $\mathbf{c}^*$ are not easy to find for general distributions, the conditions required in Corollary 3.1 are not that simple to satisfy. However, the condition we offer in Theorem 3.2 is valid even if $d \geq 2$, and is not as technical as Assumption 2 or Assumption 4. Moreover, as will be described in Section 4.2, this condition is relevant in the case where P is a mixture distribution, and can be turned into a condition on parameters of the mixture which can be easily inferred from the training sample.

4. Examples and discussion

4.1. A toy example

In this subsection we intend to understand which conditions on the density f can guarantee that the Hessian matrices H are positive definite. Some light is shed on the problem by the extremal case in which the density is zero at every boundary of optimal Voronoi cells. Indeed, in this case, equation 1 guarantees that the matrices H are diagonal matrices with positive elements, thus positive definite.

The following proposition offers an intuitive example of such an extremal case.

Proposition 4.1. *Let $z_1, \ldots, z_k$ be vectors in $\mathbb{R}^d$. Let ρ be a positive number and $R = \inf_{i \neq j} \quad \|z_i - z_j\|$ be the smallest possible distance between these vectors. Define the triangular function t on $\mathbb{R}^d$ as follows: $t(x_1, \ldots, x_d) = (1-r)\mathbf{1}_{r \leq 1}$, where $r = \sqrt{x_1^2 + \ldots + x_d^2}$ is the Euclidean norm of x. Then we define the distribution P_ρ and its density f_ρ as follows*

$$f_\rho(x) = \frac{1}{kN_\rho} \sum_{i=1}^{k} t\left(\frac{x - z_i}{\rho}\right),$$

where N_ρ is such that $P(\mathcal{B}(z_i, \rho)) = 1/k$, for $k = 1, \ldots, k$.

Suppose that $\rho < R/2$. If $\left(\frac{R}{2} - 3\rho\right)^2 \geq \frac{2\rho^2 d(d+1)}{(d+2)(d+3)}$, then the optimal k-codebook is $(z_1, \ldots, z_k)$.

The proof of Proposition 4.1, which is given in Section 5, is inspired from a proof of Bartlett, Linder and Lugosi [4, Step 3]. It is interesting to note that Proposition 4.1 can be extended to the situation where we assume that the underlying distribution is supported on k small enough subsets. In this context, if each subset has a not too small P-measure, and if the subsets are far enough from each other, it can be proved in the same way that an optimal codebook has a code point in every small subset.

Let us now consider the distribution described in Proposition 4.1, with relevant values for ρ and R. We immediately see that, if $R/2 > \rho$, then every boundary of the Voronoi cells for the optimal codebook lies in a null-measured area. Thus, for this distribution,

$$H(\mathbf{c}^*) = \begin{pmatrix} \frac{1}{k} I_d & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \frac{1}{k} I_d \end{pmatrix},$$

which is clearly positive definite.

This short example illustrates the idea behind Theorem 3.2. Namely, if the density of the distribution is not too big at the boundaries of Voronoi cells associated with every optimal codebook, then the Hessian matrix H will roughly behave as a positive diagonal matrix. In this situation, Pollard's condition (Assumption 2) will hopefully be satisfied.

This most favorable case is in fact derived from the special case where the distribution is supported on k points. Here we spread the atoms into small balls to give a density to the distribution and match regularity conditions. Antos, Györfi and György [2] proved that if the distribution has only a finite number of atoms, then the expected distortion $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ is at most C/n, where C is a constant. Proposition 4.1 guarantees that the convergence rate of $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ remains $1/n$ when the distribution P we offer in this proposition is close enough to a distribution supported on k points.

Proposition 4.1 also illustrates the main difficulty in applying Theorem 3.2: to locate the optimal codebooks $\mathbf{c}^*$. Although some results about geometrical properties of optimal codebooks have been obtained by Tarpey in [25] in the special case where P is a strongly symmetric distribution, or by Junglen [14] when k grows to infinity, there are few results about the exact location of optimal codebooks in general.

However, when the probability distribution P has k natural clusters, as in Proposition 4.1, it is possible to give an approximative location of the optimal codebooks of P. The following Section offers another example of such a well-clustered distribution.

4.2. Quasi-Gaussian mixture example

The aim of this subsection is to apply our results to the Gaussian mixtures in dimension $d = 2$. The Gaussian mixture model is a typical and well-defined clustering example. However we will not deal with the clustering issue but rather with its theoretical background.

In general, a Gaussian mixture distribution $\tilde{P}$ is defined by its density

$$\tilde{f}(x) = \sum_{i=1}^{\tilde{k}} \frac{p_i}{2\pi\sqrt{|\Sigma_i|}} e^{-\frac{1}{2}(x-m_i)^t \Sigma_i^{-1}(x-m_i)},$$

where $\tilde{k}$ denotes the number of component of the mixture, and the p_i's denote the weights of the mixture, thus satisfy $\sum_{i=1}^{k} p_i = 1$. Moreover, the m_i's denote the means of the mixture, so that $m_i \in \mathbb{R}^2$, and the Σ_i's are the 2×2 variance matrices of the components.

We restrict ourselves to the case where the number of components $\tilde{k}$ is known, and match the size k of the codebooks. To ease the calculation, we make the additional assumption that every component has the same diagonal variance matrix $\Sigma_i = \sigma^2 I_2$. Note that a similar result to Proposition 4.2 can be derived for distributions with different variance matrices Σ_i, at the cost of more computing.

Since the distribution support of a Gaussian random variable is not bounded, we define the "quasi-Gaussian" mixture model as follows, truncating each Gaussian component. Let the density f of the distribution P be defined by

$$f(x) = \sum_{i=1}^{k} \frac{p_i}{2\pi\sigma^2 N_i} e^{-\frac{\|x-m_i\|^2}{2\sigma^2}} \mathbf{1}_{\mathcal{B}(0,1)},$$

where N_i denotes a normalization constant for each Gaussian variable.

To ensure this model to be close to the Gaussian mixture model, we assume that there exists a constant $\varepsilon \in [0,1]$ such that, for $i = 1, \ldots, k$, $N_i \geq 1 - \varepsilon$.

Denote by $\tilde{B} = \inf_{i \neq j} \|m_i - m_j\|$ the smallest possible distance between two different means of the mixture. To avoid boundary issues we suppose that, for all $i = 1, \ldots, k$, $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0,1)$.

For such a model, Proposition 4.2 below offers a sufficient condition for P to be well-clustered.

Proposition 4.2. *Denote by* $p_{min} = \min_{i=1,\dots,k} p_i$ *and* $p_{max} = \max_{i=1,\dots,k} p_i$. *Suppose that*

$$\frac{p_{min}}{p_{max}} \geq \max\left(\frac{288k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/288\sigma^2})}, \frac{24k}{(1-\varepsilon)\sigma^2\tilde{B}(e^{\tilde{B}/72\sigma^2}-1)}\right).$$

Then P *satisfies Assumption 2.*

The inequality we propose as a condition in Proposition 4.2 can be decomposed as follows. If

$$\frac{p_{min}}{p_{max}} \geq \frac{288k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/288\sigma^2})},$$

then the optimal codebook $\mathbf{c}^*$ is close to the vector of means of the mixture $\mathbf{m} = (m_1, \dots, m_k)$. Knowing that, we can locate the boundaries of Voronoi cells associated with $\mathbf{c}^*$, and apply Theorem 3.2. This leads to the second term of the maximum in Proposition 4.2.

This condition can be interpreted as a condition on the polarization of the mixture. A favorable case for vector quantization seems to be when the poles of the mixtures are well-separated, which is equivalent to σ is small compared to $\tilde{B}$, when considering Gaussian mixtures. Proposition 4.2 just explained how σ has to be small compared to $\tilde{B}$, in order to satisfy Assumption 2, and therefore apply Corollary 3.1, for the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ to reach an improved convergence rate of $1/n$.

Notice that Proposition 4.2 can be considered as an extension of Proposition 4.1. In these two propositions a key point is to locate $\mathbf{c}^*$, which is possible when the distribution P is well-clustered. The definition of a well-clustered distribution takes two similar forms when looking at Proposition 4.1 or Proposition 4.2. In Proposition 4.1 the good case is when every pole of the distribution is far enough from the other, separated by an empty area with respect to P, which ensures that the Hessian matrices $H(\mathbf{c}^*)$ are positive definite (in this case they are diagonal matrices). When slightly perturbing the framework of Proposition 4.1, it is quite natural to think that the Hessian matrices $H(\mathbf{c}^*)$ should remain positive definite. Proposition 4.2 is an illustration of this idea: the empty separation area between poles is replaced with an area where the density f is small compared to its value around the poles. The condition on σ and $\tilde{B}$ we offer in Proposition 4.2 gives a theoretical definition of a well-clustered distribution for quasi-Gaussian mixtures.

It is important to note that our result is valid when k is known and match exactly the number of components of the mixture. When the number of code points k is different from the number of components $\tilde{k}$ of the mixture, we have no general idea of where the optimal code points can be located.

Moreover, suppose that there is only one optimal codebook $\mathbf{c}^*$, up to re indexation, and that we are able to locate this optimal codebook $\mathbf{c}^*$. As explained in the proof of Proposition 4.2, the quantity at stake is in fact $B = \inf_{i \neq j} \|c_i^* - c_j^*\|$.

In the case where $\tilde{k} \neq k$, there is no simple relation between $\tilde{B}$ and B. Consequently, a condition like in Proposition 4.2 could not involve the natural parameter of the mixture $\tilde{B}$.

It is also worth pointing out that there exist cases where the set of optimal codebooks is not finite. For example, suppose that P is a truncated rotationally symmetric Gaussian distribution, and $k = 2$. Since every rotation of an optimal codebook leads to an other optimal codebook, there exists an infinite set of optimal codebooks. This ensures that at least one Hessian matrix $H(\mathbf{c}^*)$ cannot be positive definite, in fact none is positive definite.

The two assumptions $N_i \geq 1 - \varepsilon$ and $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0,1)$ can easily be satisfied when P is constructed via an homothetic transformation. To see this, take a generic Gaussian mixture on $\mathbb{R}^2$, denote by $\bar{m}_i, i = 1, \ldots, k$, its means and by $\bar{\sigma}^2$ its variance. For a given $\varepsilon > 0$, choose $M > 0$ such that, for all $i = 1, \ldots, k$, $\int_{\mathcal{B}(0,M)} e^{-\|x-m_i\|^2/2\sigma^2} dx \geq 2\pi\sigma^2(1-\varepsilon)$ and $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0, M)$. Denote by P_0 the "quasi-Gaussian mixture" we obtain on $\mathcal{B}(0, M)$ for such an M. Then, applying an homothetic transformation with coefficient $1/M$ to P_0 provides a quasi-Gaussian mixture on $\mathcal{B}(0,1)$, with means $m_i = \bar{m}_i/M, i = 1, \ldots, k$ and variance $\sigma^2 = \bar{\sigma}^2/M^2$. This distribution satisfies both $N_i \geq 1 - \varepsilon$ and $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0,1)$.

5. Proofs

5.1. *Proof of Proposition 3.1*

Half of the equivalences of Proposition 3.1 derives from the following interesting lemma.

Lemma 5.1. *Suppose that Assumption 1 is satisfied, $\mathcal{M}$ is finite, and P has a continuous density. Then there exist two constants $C_- > 0$ and $C_+ > 0$ such that*

$$\forall \mathbf{c} \in \mathcal{B}(0,1)^k \quad C_-\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2 \leq \mathrm{Var}_P\left(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .)\right) \leq C_+\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2.$$

Lemma 5.1 ensures that, if P is smooth enough, $\mathrm{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .)$ is equivalent to the squared Euclidean distance between $\mathbf{c}$ and $\mathbf{c}^*(\mathbf{c})$, which provides a direct connection between Assumption 3 and Assumption 4. Notice that we require the density of P to be smooth in order to apply a theorem of Baddeley [3, Theorem 1]. Improving the conditions of this theorem could be a way to soften the smoothness requirements on P.

It is important to note that Assumption 1 too is crucial to make the result of Proposition 3.1 valid, since it allows us to turn local differentiation arguments into global properties. The other half of the equivalences of Proposition 3.1 follows from the following result, due to Antos, Györfy and György in the proof of [2, Corollary1].

Lemma 5.2. *Suppose that P satisfies Assumption 1, then, there exists a constant $A_2 > 0$ such that for a fixed $\mathbf{c}^* \in \mathcal{M}$*

$$\sup_{\mathbf{c} \notin \mathcal{M}} \frac{\operatorname{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .))}{\|\mathbf{c} - \mathbf{c}^*\|^2} \leq A_2.$$

Moreover, if P satisfies Assumption 1 and Assumption 2, and if P has a continuous density, then there exists a constant $A_1 > 0$ such that

$$\inf_{\mathbf{c} \notin \mathcal{M}} \frac{\ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c}))}{\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2} \geq A_1.$$

Equipped with these two lemmas, we are now in position to prove Proposition 3.1. Antos, Györfi and György [2] pointed out that, if P satisfies Assumption 2 and Assumption 1, then $\mathcal{M}$ is finite. If not, due to a compactness argument it can be proved that $\mathcal{M}$ has an accumulation point, which ensures that the Hessian matrix H at this accumulation point cannot be positive definite.

Now suppose that P has a continuous density. Then, according to Pollard [23, Lemma C], $P\gamma(\mathbf{c}, .)$ is differentiable twice at every point $\mathbf{c}$. Furthermore, if $\mathcal{M}$ is finite and if $\ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c})) \geq A_1 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$, then the Hessian matrices $H(\mathbf{c}^*)$ have to be positive definite for every $\mathbf{c}^*$ in $\mathcal{M}$. This leads to the following equivalence:

$$\left\{ \begin{array}{c} \text{Assumption 1} \\ \text{Assumption 2} \end{array} \right. \Leftrightarrow \left\{ \begin{array}{c} f \text{ has a continuous density} \\ \mathcal{M} \text{ is finite} \\ \text{Assumption 1} \\ \text{Assumption 3} \end{array} \right. .$$

The other equivalence relies on Lemma 5.1. Since Assumption 3 obviously implies Assumption 4, the direct implication is proved. Now suppose that P has a continuous density, $\mathcal{M}$ is finite, and satisfies Assumption 1 and Assumption 4. Since Assumption 1 is satisfied and $\mathcal{M}$ is finite, the first part of Lemma 5.2 provides us with a global $A_2 > 0$ such that

$$\operatorname{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .)) \leq A_2 \|\mathbf{c} - \mathbf{c}^*\|^2.$$

Combining Lemma 5.1 with Assumption 4 ensures the existence of $A_1 > 0$ such that

$$\ell(\mathbf{c}, \mathbf{c}^*(\mathbf{c})) \geq A_1 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2.$$

Then, we can deduce that

$$\left\{ \begin{array}{c} f \text{ has a continuous density} \\ \mathcal{M} \text{ is finite} \\ \text{Assumption 1} \\ \text{Assumption 3} \end{array} \right. \Leftrightarrow \left\{ \begin{array}{c} f \text{ has a continuous density} \\ \mathcal{M} \text{ is finite} \\ \text{Assumption 1} \\ \text{Assumption 4} \end{array} \right. .$$

5.2. Proof of Lemma 5.1

Lemma 5.1 relies on the following technical lemma, which provides some differentiation arguments in order to connect $\mathrm{Var}_P(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.))$ to the squared Euclidean distance $\|\mathbf{c} - \mathbf{c}^*\|^2$.

Lemma 5.3. *Let $\mathbf{c}^* \in \mathcal{M}$ be fixed. Let f be the real-valued function defined by*

$$f : \left\{ \begin{array}{ccc} (\mathbb{R}^d)^k & \longrightarrow & \mathbb{R} \\ \mathbf{c} & \longmapsto & \mathrm{Var}_P(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)) \end{array} \right. .$$

Then f is differentiable twice at the point $\mathbf{c} = \mathbf{c}^*$, and its Hessian matrix F is made of the following $d \times d$ blocks

$$F_{i,j} = \left\{ \begin{array}{ccc} 8\int_{V_i^*} f(x)(x - c_j^*)(x - c_j^*)^t & if & i = j \\ 0 & if & i \neq j \end{array} \right. \qquad i,j = 1,\ldots,k.$$

Furthermore, the matrix F is positive definite.

Proof of Lemma 5.3. First we write

$$f(\mathbf{c}) = P\left(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)\right)^2 - \left(P\left(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)\right)^2\right).$$

Using almost the same argument as in [23, Lemma A], f is differentiable at every point $\mathbf{c}$ in $\mathcal{B}(0,1)^d$, with gradient

$$\begin{aligned} \nabla f(\mathbf{c}) &= 2P\left(\Delta(\mathbf{c},.)(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.))\right) - 2P\Delta(\mathbf{c},.)P(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)) \\ &:= 2g_1(\mathbf{c}) - 2g_2(\mathbf{c}) \end{aligned}$$

where $\Delta(\mathbf{c},x)$ is the point wise gradient function defined as in [23]

$$\Delta(\mathbf{c},x) = -2((x - c_1)\mathbf{1}_{V_1},\ldots,(x - c_k)\mathbf{1}_{V_k}).$$

First we deal with g_1. Writing

$$g_1(\mathbf{c}) = \left(-2\int_{\mathbb{R}^d}\left(\|x - c_i\|^2 - \gamma(\mathbf{c}^*,x)\right)(x - c_i)\mathbf{1}_{V_i}(x)\right)_{i=1,\ldots,k},$$

we use [3, Theorem 1] to prove that g_1 is differentiable at every $\mathbf{c}$, with derivatives matrix denoted by H_1. Through computation like in [23, Lemma C], we get the following decomposition in $d \times d$ blocks for H_1:

$$\begin{aligned} H_1(\mathbf{c})_{i,i} &= 4\int_{V_i} f(x)(x - c_i)(x - c_i)^t dx + 2\int_{V_j} f(x)\left(\gamma(\mathbf{c},x) - \gamma(\mathbf{c}^*,x)\right)dx \\ &\quad - 2\sum_{p \neq i}\|c_i - \mathbf{c}_j\|^{-1}\int_{\partial(V_i \cap V_p)} f(x)\left(\gamma(\mathbf{c},x) - \gamma(\mathbf{c}^*,x)\right)(x - c_i)(x - c_i)^t dx, \end{aligned}$$

for diagonal blocks. For other blocks we have, with $i \neq j$,

$$H_1(\mathbf{c})_{i,j} = 2\|c_i - c_j\|^{-1} \int_{\partial(V_i \cap V_j)} f(x) \left(\gamma(\mathbf{c}, x) - \gamma(\mathbf{c}^*, x)\right) (x - c_i)(x - c_j)^t dx.$$

Using the same argument (see [3, Theorem1]), g_2 is differentiable. Recalling that H denotes the Hessian matrix of $\mathbf{c} \longmapsto P\gamma(\mathbf{c}, .)$, elementary calculation shows that, if H_2 denotes the matrix of derivatives of g_2,

$$H_2(\mathbf{c}) = P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .))H(\mathbf{c}) + P\Delta(\mathbf{c}, .)(P\Delta(\mathbf{c}, .))^t.$$

Hence we deduce that f is differentiable twice at point $\mathbf{c} = \mathbf{c}^*$, with Hessian matrix $F(\mathbf{c}^*) = H_1(\mathbf{c}^*) + H_2(\mathbf{c}^*)$. Taking $\mathbf{c} = \mathbf{c}^*$ in the above calculations leads to the result for the expression of F.

It remains to prove that F is positive definite. Let $\mathbf{h} = (h_1, \dots, h_k)$ be a $k \times d$ vector, with $\mathbf{h} \neq 0$. We notice that

$$\mathbf{h}^t F \mathbf{h} = \sum_{i=1}^{k} h_i^t F_{i,i} h_i = 8 \sum_{i=1}^{k} \int_{V_i^*} f(x) \left\langle h, x - c_i^* \right\rangle^2 dx.$$

Suppose that $h^t F h = 0$, then, for $i = 1, \dots, k$, $h_i^t F_{i,i} h_i = 0$. Since $\mathbf{h} \neq 0$, we can assume without loss of generality that $h_1 \neq 0$. We denote by $h_1^\perp$ the hyperplane in $\mathbb{R}^d$ orthogonal to h_1. Since $h_1^t F_{1,1} h_1 = 0$, we deduce that $P(V_1^* \setminus \left(c_1^* + h_1^\perp\right)) = 0$. Taking into account that P has a density, we get $P(V_1^*) = 0$, which is impossible, according to [13, Theorem 4.1]. □

Now we turn to the proof of Lemma 5.1. Suppose that P has a continuous density f, $\mathcal{M}$ is finite, and P satisfies Assumption 1. Since P satisfies Assumption 1 and $\mathcal{M}$ is finite, the first part of Lemma 5.2 provides $C_+ > 0$ such that

$$\operatorname{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .)) \leq C_+ \|\mathbf{c} - \mathbf{c}^*\|^2.$$

Consequently, we only have to deal with the lower bound. To do this, suppose that P is such that

$$\inf_{\mathbf{c} \notin \mathcal{M}} \frac{\operatorname{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*(\mathbf{c}), .))}{\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2} = 0,$$

then there exists a sequence $(\mathbf{c}_n)_{n \geq 1}$, such that $\mathbf{c}_n \notin \mathcal{M}$ and

$$\frac{\operatorname{Var}_P(\gamma(\mathbf{c}_n, .) - \gamma(\mathbf{c}^*(\mathbf{c}_n), .))}{\|\mathbf{c}_n - \mathbf{c}^*(\mathbf{c}_n)\|^2} \longrightarrow 0,$$

as $n \longrightarrow \infty$. Since Assumption 1 is satisfied, we can assume without loss of generality that there exists $\mathbf{c}$ in $\mathcal{B}(0, 1)^d$, such that $\mathbf{c}_n \longrightarrow \mathbf{c}$.

We have to prove that $\mathbf{c} \in \mathcal{M}$. To do this, suppose that $\mathbf{c} \notin \mathcal{M}$, and denote by $\mathbf{c}^*$ the closest optimal codebook to $\mathbf{c}$. Since $\mathbf{c} \notin \mathcal{M}$, there exists i such that $c_i^* \neq c_j$, for $j = 1, \dots, k$. Furthermore, since $\mathcal{M}$ is a finite set, $\mathbf{c}^*(\mathbf{c}_n) = \mathbf{c}^*$ for n large enough. Since $\operatorname{Var}_P(\gamma(\mathbf{c}_n, .) - \gamma(\mathbf{c}^*(\mathbf{c}_n), .)) \longrightarrow 0$, we deduce that $\operatorname{Var}_P(\gamma(\mathbf{c}, .) - \gamma(\mathbf{c}^*, .)) = 0$, which in turn leads to $\gamma(\mathbf{c}, x) = \gamma(\mathbf{c}^*, x) + a$, , for a constant $a > 0$, P-almost surely in x. Let $x \in V_i^*$. We denote by G and G_j the following sets of points:

$$\begin{array}{rcl} G & = & \{x \in V_i^* | \quad \gamma(\mathbf{c}, x) = \gamma(\mathbf{c}^*, x) + a\} \\ G_j & = & \left\{x \in V_i^* | \quad \|x - c_j\|^2 = \|x - c_i^*\|^2 + a\right\} \end{array}.$$

Formally we have $G \subset \bigcup_{j=1}^{k} G_j$, and $P(V_i^*) = P(G)$. However, since G_j is an affine space with dimension $d-1$ and P has a density, it follows that $P(G_j) = 0$ for all $j = 1, \ldots, k$. Hence we deduce that $P(G) = 0$, so that $P(V_i^*) = 0$, which is not possible for an optimal codebook $\mathbf{c}^*$ (see [13, Theorem 4.1]). Hence we deduce that $\mathbf{c} \in \mathcal{M}$.

Then we can assume that $\mathbf{c}_n \longrightarrow \mathbf{c}^*$, for some fixed $\mathbf{c}^* \in \mathcal{M}$, and, since $\mathcal{M}$ is a finite set, without loss of generality, $\mathbf{c}^*(\mathbf{c}_n) = \mathbf{c}^*$ for $n \geq 1$. According to Lemma 5.3, there exists $C_+ > 0$ such that

$$\mathrm{Var}_P(\gamma(\mathbf{c}_n, .) - \gamma(\mathbf{c}^*, .)) \geq C_+\|\mathbf{c}_n - \mathbf{c}^*\|^2 + o(\|\mathbf{c}_n - \mathbf{c}^*\|^2),$$

and so

$$\frac{\mathrm{Var}_P(\gamma(\mathbf{c}_n, .) - \mathbf{c}^*(\mathbf{c}_n))}{\|\mathbf{c}_n - \mathbf{c}^*(\mathbf{c}_n)\|^2} \geq C_+ + o(1),$$

which leads to a contradiction.

5.3. Proof of Theorem 3.1

The proof strongly relies on the localization principle and its application by Blanchard, Bousquet and Massart [6]. We start with the following definition.

Definition 5.1. *Let Φ be a real-valued function. Φ is called a sub-α function if and only if Φ is non-decreasing and the map $x \mapsto \Phi(x)/x^\alpha$ is non-increasing.*

The next theorem is an adaptation of the result of Blanchard, Bousquet and Massart [6, Theorem 6.1]. For the sake of clarity its proof is given in Subsection 5.4.

Theorem 5.1. *Let $\mathcal{F}$ be a class of bounded measurable functions such that there exist $b > 0$ and $\omega : \mathcal{F} \longrightarrow \mathbb{R}^+$ satisfying*

(i) $\forall f \in \mathcal{F} \quad \|f\|_\infty \leq b,$
(ii) $\forall f \in \mathcal{F} \quad \mathrm{Var}_P(f) \leq \omega(f).$

Let K be a positive constant, Φ a sub-α function, $\alpha \in [1/2, 1[$. Then there exists a constant $C(\alpha)$ such that, if D is a constant satisfying $D \leq 6KC(\alpha)$, and δ^* is the unique solution of the equation $\Phi(\delta) = \delta/D$, the following holds. Assume that

$$\forall \delta \geq \delta^* \qquad \mathbb{E}\left(\sup_{\omega(f) \leq \delta} |(P - P_n)f|\right) \leq \Phi(\delta).$$

Then, for all $x > 0$, with probability larger than $1 - e^{-x}$,

$$\forall f \in \mathcal{F} \quad Pf - P_n f \leq K^{-1}\left(\omega(f) + \left(\frac{6KC(\alpha)}{D}\right)^{\frac{1}{1-\alpha}}\delta^* + \frac{(9K^2 + 16Kb)x}{4n}\right).$$

As explained in the proof, an optimal choice for $C(\alpha)$ is

$$C(\alpha) = \inf_{x>1}\left(1 + x^{\alpha}\left(\frac{1}{2} + \frac{1}{x^{1-\alpha}-1}\right)\right).$$

This theorem provides a sharp concentration inequality in the case where it is possible to control the maximal deviation between P and P_n over a set of functions whose variance with respect to P is constrained within a ball. The main point is to find a suitable control function for the variance of the process. Here the interesting set is

$$\mathcal{F} = \left\{\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.), \mathbf{c} \in \mathcal{B}(0,1)^k, \mathbf{c}^* \in \mathcal{M}\right\}.$$

Since P satisfies Assumption 3, the relevant control function for the variance of the process $\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)$ is $\omega(\mathbf{c},\mathbf{c}^*) = A_2\|\mathbf{c}-\mathbf{c}^*\|^2$, where A_2 is defined in Assumption 3.

Thus it remains to bound from above the quantity

$$\mathbb{E}\left(\sup_{\mathbf{c}^*\in\mathcal{M}, A_2\|\mathbf{c}-\mathbf{c}^*\|^2\leq\delta} |(P_n - P)(\gamma(\mathbf{c}^*,.) - \gamma(\mathbf{c},.))|\right).$$

This is done in the following proposition.

Proposition 5.1. *Suppose that P has a density and satisfies Assumption 1. Furthermore we assume that $\mathcal{M}$ is finite. Then*

$$\mathbb{E}\left(\sup_{\mathbf{c}^*\in\mathcal{M}, A_2\|\mathbf{c}-\mathbf{c}^*\|^2\leq\delta} |(P_n - P)(\gamma(\mathbf{c}^*,.) - \gamma(\mathbf{c},.))|\right) \leq \sqrt{\delta}\frac{\Xi}{\sqrt{n}},$$

where Ξ is a constant depending on k, d, and P.

Since we assume that $\mathcal{M}$ is finite, P has a density and Assumption 3 is satisfied, we can apply Theorem 5.1, with $\omega(\mathbf{c},\mathbf{c}^*) = A_2\|\mathbf{c}-\mathbf{c}^*\|^2$, $b = 8$, and $\Phi(\delta) = \sqrt{\delta}\Xi/\sqrt{n}$. Noticing that the solution of the equation $\delta = \Phi(\delta)/D$ is $\Xi^2 D^2/n$, for an arbitrary $D > 0$, we get the following result.

Lemma 5.4. *Suppose that P has a density, satisfies Assumption 1 and Assumption 3, and $\mathcal{M}$ is finite. Let $D > 0$. For all $\mathbf{c}^* \in \mathcal{M}$, $x > 0$ and $D \leq 6KC(1/2)$, we have, with probability larger than $1 - e^{-x}$,*

$$(P - P_n)(\gamma(\mathbf{c},.) - \gamma(\mathbf{c}^*,.)) \leq K^{-1}A_2\|\mathbf{c}-\mathbf{c}^*\|^2 + \frac{36KC(1/2)^2\Xi^2}{n} + \frac{2K+32}{n}x.$$

Take $\mathbf{c}^* = \mathbf{c}^*(\mathbf{c})$, a nearest optimal codebook to $\mathbf{c}$, and use (H2) to connect $\|\mathbf{c}-\mathbf{c}^*(\mathbf{c})\|^2$ to $\ell(\mathbf{c},\mathbf{c}^*(\mathbf{c}))$. Choosing $K = 2A_1A_2$, $D = 6KC(1/2)$, we get, with probability larger than $1 - e^{-x}$,

$$1/2(P - P_n)(\gamma(\hat{\mathbf{c}}_n,.) - \gamma(\mathbf{c}^*(\hat{\mathbf{c}}_n),.)) \leq \frac{C_1}{n} + \frac{C_2}{n}x,$$

for some constants $C_1 > 0$ and $C_2 > 0$. Since $P_n(\gamma(\hat{\mathbf{c}}_n, .) - \gamma(\mathbf{c}^*(\hat{\mathbf{c}}_n), .)) \leq 0$, taking expectation leads to, for all $\mathbf{c}^* \in \mathcal{M}$,

$$\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq \frac{C_0}{n},$$

for a constant $C_0 > 0$ depending only on k, d, and P.

5.4. *Proof of Theorem 5.1*

This proof is a modification of the proof of Blanchard, Bousquet and Massart [6, Theorem 6.1]. For $\delta \geq 0$, set

$$\Omega_\delta = \sup_{f \in \mathcal{F}} \left| (P - P_n) \frac{f}{\omega(f) + \delta} \right|.$$

We start with a modified version of the so-called peeling lemma:

Lemma 5.5. *Under the assumptions of Theorem 5.1, there exists a constant $C(\alpha)$ depending only on α such that, for all $\delta > 0$,*

$$\mathbb{E}\left(\Omega_\delta\right) \leq C(\alpha) \frac{\Phi(\delta)}{\delta}.$$

Furthermore, we have $C(\alpha) \underset{\alpha \to 1}{\longrightarrow} \infty$.

Proof of Lemma 5.5. Let $x > 1$ be a real number. We may write

$$\begin{aligned} \sup_{f \in \mathcal{F}} \left| (P - P_n) \frac{f}{\omega(f) + \delta} \right| \leq & \sup_{\omega(f) \leq \delta} \left| (P - P_n) \frac{f}{\omega(f) + \delta} \right| \\ & + \sum_{k \geq 0} \sup_{\delta x^k < \omega(f) \leq \delta x^{k+1}} \left| (P - P_n) \frac{f}{\omega(f) + \delta} \right|. \end{aligned}$$

Since $\sup_{\delta x^k < \omega(f) \leq \delta x^{k+1}} |(P - P_n) f| \geq 0$, and $\omega(f) + \delta > 0$, taking expectation on both sides leads to

$$\mathbb{E}(\Omega_\delta) \leq \frac{\Phi(\delta)}{\delta} + \sum_{k \geq 0} \frac{\Phi(\delta x^{k+1})}{\delta(1 + x^k)}.$$

Recalling that Φ is a sub-α function, we may write $\Phi(\delta x^{k+1}) \leq x^{\alpha(k+1)} \Phi(\delta)$. Hence we get

$$\begin{aligned} \mathbb{E}(\Omega_\delta) & \leq \frac{\Phi(\delta)}{\delta} + \frac{\Phi(\delta)}{\delta} \sum_{k \geq 0} \frac{x^{\alpha(k+1)}}{1 + x^k} \\ & \leq \frac{\Phi(\delta)}{\delta} \left(1 + x^\alpha \left(\frac{1}{2} + \frac{1}{x^{1-\alpha} - 1} \right) \right). \end{aligned}$$

Taking $C(\alpha) = \inf_{x > 1} \left(1 + x^\alpha \left(\frac{1}{2} + \frac{1}{x^{1-\alpha} - 1} \right) \right)$ proves the result. ∎

We are now in a position to prove Theorem 5.1. Using the inequality of Talagrand for a supremum of bounded variables that Bousquet [7] offered, we have, with probability larger than $1 - e^{-x}$,

$$\Omega_\delta \leq \mathbb{E}(\Omega_\delta) + \sqrt{\frac{x}{2\delta n}} + 2\sqrt{\frac{xb\mathbb{E}(\Omega_\delta)}{n\delta}} + \frac{bx}{3\delta n}.$$

Using Lemma 5.5 and the inequality $2ab \leq a^2 + b^2$,

$$\Omega_\delta \leq \frac{2C(\alpha)\Phi(\delta)}{\delta} + \sqrt{\frac{x}{2\delta n}} + \frac{4}{3}\frac{bx}{\delta n}.$$

Let δ^* be the solution of $\Phi(\delta) = \frac{\delta}{D}$. If $\delta \geq \delta^*$, then $\frac{\Phi(\delta)}{\delta} \leq \left(\frac{\delta^*}{\delta}\right)^{1-\alpha} \frac{1}{D}$. For such an δ we have

$$\Omega_\delta \leq \beta_1 \delta^{-(1-\alpha)} + \beta_2 \delta^{-1/2} + \beta_3 \delta^{-1},$$

with

$$\begin{cases} \beta_1 = \dfrac{2C(\alpha)(\delta^*)^{1-\alpha}}{D} \\ \beta_2 = \sqrt{\dfrac{x}{2n}} \\ \beta_3 = \dfrac{4bx}{3n} \end{cases}.$$

We want to find a suitable δ such that $\delta \geq \delta^*$ and $\Omega_\delta \leq 1/K$. To this aim, it suffices to see that if $\delta \geq (3K\beta_1)^{\frac{1}{1-\alpha}} + (3K\beta_2)^2 + 3K\beta_3$, and $\delta \geq \delta^*$, then $\Omega_\delta \leq 1/K$ using the previous upper bound on Ω_δ.

It remains to check that the condition $(3K\beta_1)^{\frac{1}{1-\alpha}} + (3K\beta_2)^2 + 3K\beta_3 \geq \delta^*$ holds. To see this just recall that

$$(3K\beta_1)^{\frac{1}{1-\alpha}} = \delta^* \times \left(\frac{6KC(\alpha)}{D}\right)^{\frac{1}{1-\alpha}}.$$

Thus, we deduce that, if $D \leq 6KC(\alpha)$, the choice $\delta = (3K\beta_1)^{\frac{1}{1-\alpha}} + (3K\beta_2)^2 + 3K\beta_3$ guarantees $\Omega_\delta \leq K^{-1}$ and, consequently, with probability larger than $1 - e^{-x}$,

$$\begin{aligned} Pf - P_n f &\leq |(P - P_n)f| \\ &\leq \left|(P - P_n)\frac{f}{\omega(f) + \delta^*}\right| \times (\omega(f) + \delta^*) \\ &\leq \Omega_{\delta^*}(\omega(f) + \delta^*) \\ &\leq \frac{1}{K}\left(\omega(f) + \left(\frac{6KC(\alpha)}{D}\right)^{\frac{1}{1-\alpha}} \delta^* + \frac{(9K^2 + 16Kb)x}{4n}\right). \end{aligned}$$

5.5. Proof of Proposition 5.1

Following the approach of Pollard in [23], we notice that, for any $\mathbf{c} \in (\mathbb{R}^d)^k$ and $\mathbf{c}^* \in \mathcal{M}$, P-almost surely in x,

$$\gamma(\mathbf{c}, x) = \gamma(\mathbf{c}^*, x) + \langle \mathbf{c} - \mathbf{c}^*, \Delta(\mathbf{c}^*, x)\rangle + \|\mathbf{c} - \mathbf{c}^*\| R(\mathbf{c}^*, \mathbf{c} - \mathbf{c}^*, x),$$

where, with use of Pollard's [23] notation

$$\begin{cases} \Delta(\mathbf{c}^*, x) = -2((x - c_1^*)\mathbf{1}_{V_1^*}, \dots, (x - c_k^*)\mathbf{1}_{V_k^*}) \\ R(\mathbf{c}^*, \mathbf{c} - \mathbf{c}^*, x) = \displaystyle\sum_{i,j=1,\dots,k} \mathbf{1}_{V_i^*}\mathbf{1}_{V_j}\|\mathbf{c} - \mathbf{c}^*\|^{-1}\left[2(c_i - c_j)^t x + \|c_i^*\|^2 \right. \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \left. -2(c_i^*)^t c_i + \|c_j\|^2\right] \end{cases}$$

We recall that V_i^* denotes the Voronoi cell associated with the code point $\mathbf{c}_i^*$, where c_i^* is a coordinate of $\mathbf{c}^*$, and that $\mathbf{1}_{V_i^*}(x)$ takes the value 1 if $x \in V_i^*$, 0 elsewhere.

Splitting the expectation in two parts, we obtain

$$\begin{aligned} &\mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c} - \mathbf{c}^*\|^2 \le \delta} |(P_n - P)(\gamma(\mathbf{c}^*, .) - \gamma(\mathbf{c}, .))|\right) \\ &\qquad \le \mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c} - \mathbf{c}^*\|^2 \le \delta} |(P_n - P)\langle -(\mathbf{c} - \mathbf{c}^*), \Delta(\mathbf{c}^*, .)\rangle|\right) \\ &\qquad + \sqrt{\delta}\,\mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c} - \mathbf{c}^*\|^2 \le \delta} |(P_n - P)(-R(\mathbf{c}^*, \mathbf{c} - \mathbf{c}^*, .))|\right) \\ &\qquad := A + B. \end{aligned} \tag{2}$$

5.5.1. Term A: Complexity of the model

Term A in inequality (2) is at first sight the dominant term in the expression $\Phi(\delta)$. The upper bound we obtain below is rather accurate, due to the finite-dimensional Euclidean space structure. Indeed, we have to bound a scalar product when the vectors are contained in a ball, thus it is easy to see that the largest value of the product matches in fact the largest value of the coordinates of the gradient term. We recall that $\mathcal{M}$ denotes the finite set of optimal codebooks. Let $\mathbf{x} = (x_1, \dots, x_k)$ be a vector in $(\mathbb{R}^d)^k$. We denote by x_{j_r} the r-th coordinate of x_j, and name it the (j, r)-th coordinate of $\mathbf{x}$. Moreover, denote by $e_{j,r}$ the vector whose (j, r)-th coordinate is 1, and other coordinates are 0.

Taking into account that every $\mathbf{c}^*$ in $\mathcal{M}$ satisfies the centroid condition, that is $P\Delta(\mathbf{c}^*, .) = 0$, we may write

$$
\begin{aligned}
&\sup_{\mathbf{c}^*\in\mathcal{M},\|\mathbf{c}-\mathbf{c}^*\|\leq\sqrt{\delta}} |\langle \mathbf{c}-\mathbf{c}^*, (P_n-P)(-\Delta(\mathbf{c}^*))\rangle| \\
&\qquad = \sup_{\mathbf{c}^*\in\mathcal{M},j=1,\dots,k,r=1,\dots,d} \left| \frac{1}{n}\sum_{i=1}^{n}(X_i-c_j^*)\mathbf{1}_{V_j^*}(X_i)\right|_r \times \sqrt{\delta}. \\
&\qquad = \sup_{\mathbf{c}^*\in\mathcal{M},j=1,\dots,k,r=1,\dots,d,\varepsilon=\pm 1} \left\langle \varepsilon\sqrt{\delta}e_{j,r}, P_n(\Delta(\mathbf{c}^*,.))\right\rangle
\end{aligned}
$$

Therefore we can reduce the set of $\mathbf{c}$'s and $\mathbf{c}^*$'s of interest to a finite set we denote by $\mathcal{H}_{\mathcal{M}}$, which contains $|\mathcal{M}|2^{kd}$ elements. Taking into account that, for every $\mathbf{c}^*$ in $\mathcal{M}$, $P\Delta(\mathbf{c}^*,.)=0$, and that, for every fixed $\mathbf{c}$ and $\mathbf{c}^*$, the quantity $\langle \mathbf{c}-\mathbf{c}^*, P_n(\Delta(\mathbf{c}^*,.))\rangle$ is a sub-Gaussian random variable with variance $16\delta/n$, we get, by a maximal inequality due to Massart, [18, Lemma 2.3]:

$$
\begin{aligned}
&\mathbb{E}\left(\sup_{\mathbf{c}^*\in\mathcal{M},\|\mathbf{c}-\mathbf{c}^*\|^2\leq\delta} |(P_n-P)\left\langle -(\mathbf{c}-\mathbf{c}^*),\Delta(\mathbf{c}^*,.)\right\rangle|\right) \\
&\qquad = \mathbb{E}\left(\sup_{(\mathbf{c},\mathbf{c}^*)\in\mathcal{H}_{\mathcal{M}}} |(P_n-P)\left\langle -(\mathbf{c}-\mathbf{c}^*),\Delta(\mathbf{c}^*,.)\right\rangle|\right) \\
&\qquad \leq \sqrt{2\frac{16\delta}{n}\log(|\mathcal{H}_{\mathcal{M}}|)} \\
&\qquad \leq 4\sqrt{2kd\log(2|\mathcal{M}|)}\frac{\sqrt{\delta}}{\sqrt{n}}.
\end{aligned}
$$

Therefore, the expected dominant term involves the complexity of the model in a way which is proportional to the square root of the complexity. In our case, this complexity is the dimension of the codebook space.

5.5.2. Bound on B

To bound the second term in inequality (2), we follow the approach of Pollard [23], using complexity arguments such as Dudley's entropy integral.

Let $\mathcal{F}$ be a set of functions defined on $\mathcal{X}$ with envelope F. Let $S\subset\mathcal{X}$ be a finite set and f a function. We denote $\|f\|_{l^2(S)} = \left(1/n\sum_{x\in S} f^2(x)\right)^{1/2}$, where $n=|S|$, and by $N_F(\varepsilon,S,\mathcal{F})$ the smallest integer m such that there exist $\phi_1,\dots,\phi_m$, m functions on $\mathcal{X}$ satisfying $\min_{i=1,\dots,m}\|f-\phi_i\|^2_{l^2(S)}\leq\varepsilon^2\|F\|^2_{l^2(S)}$. Also define $H(\varepsilon)=\sup_{|S|<\infty}\log N_F(\varepsilon,S,\mathcal{F})$, and $m(\varepsilon)=e^{H(\varepsilon)}$, so that for any subset $S\subset\mathcal{X}$ there exists a $\varepsilon\|F\|_{l^2(S)}$-chaining of $\mathcal{F}$ with at most $m(\varepsilon)$ elements.

Pollard proved in [23, p.921], using a result proposed in [22, Theorem 9], that, for the class of functions

$$
\mathcal{F}=\left\{R(.,\mathbf{c}^*,\mathbf{c}-\mathbf{c}^*),\mathbf{c}^*\in\mathcal{M},\mathbf{c}\in\mathcal{B}(0,1)^k\right\},
$$

there exist $C>0$ depending on k and d such that $F(x)=C(1+\|x\|)$ is an envelope for $\mathcal{F}$. Furthermore, for this envelope, we have

$$
H(\varepsilon)\leq\log(A)-W\log(\varepsilon),
$$

where A is a positive constant, and W depends only on the pseudo-dimension of $\mathcal{F}$. We will use a classical chaining argument to bound term B. Let $\tilde{\mathbf{c}}$ denote the pair $(\mathbf{c}, \mathbf{c}^*) \in (\mathcal{B}(0,1))^k \times \mathcal{M}$. For practical, let $f_{\tilde{\mathbf{c}}}$ denote the function $R(., \mathbf{c}^*, \mathbf{c} - \mathbf{c}^*)$. We set $\varepsilon_0 = 1$ and $\varepsilon_j = 2^{-j}\varepsilon_0$.

Let $X_1, \ldots, X_n$ be fixed, and denote by S_n the random set $\{X_1, \ldots, X_n\}$. For any $f_{\tilde{\mathbf{c}}}$, let $f_{\tilde{\mathbf{c}}_j}$ be a function such that $\|f_{\tilde{\mathbf{c}}} - f_{\tilde{\mathbf{c}}_j}\|^2_{l^2(S_n)} \leq \varepsilon_j^2 \|F\|^2_{l^2(S_n)}$. Making use of the result of Pollard [23, p.921] mentioned above, we may write $\left|\left\{f_{\tilde{\mathbf{c}}_j} | \mathbf{c} \in \mathcal{B}(0,1)^k, \mathbf{c}^* \in \mathcal{M}\right\}\right| \leq m(\varepsilon_j) \leq A\varepsilon_j^{-W}$.

Since Assumption 1 holds, F is bounded from above by a constant C_F. By dominated convergence Theorem we have $f_{\tilde{\mathbf{c}}_j} \underset{j\to\infty}{\overset{L^1, a.s}{\longrightarrow}} f_{\tilde{\mathbf{c}}}$, and thus

$$(P_n - P)f_{\tilde{\mathbf{c}}} = (P_n - P)f_{\tilde{\mathbf{c}}_0} + \sum_{j=1}^{\infty}(P_n - P)(f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}}).$$

Therefore

$$\begin{aligned}\mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c}-\mathbf{c}^*\| \leq \sqrt{\delta}} |(P_n - P)f_{\tilde{\mathbf{c}}}|\right) &\leq \mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c}-\mathbf{c}^*\| \leq \sqrt{\delta}} |(P_n - P)f_{\tilde{\mathbf{c}}_0}|\right) \\ &\quad + \sum_{j>0}\mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c}-\mathbf{c}^*\| \leq \sqrt{\delta}} \left|(P_n - P)(f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}})\right|\right).\end{aligned}$$

Here we use a symmetrization inequality to bound from above the last term with a Rademacher complexity. Symmetrization inequalities were introduced by Giné and Zinn in [12], however we rather use the approach developed by Koltchinskii in [15, Section 2.2]. In fact, introducing some Rademacher random variables σ ($\sigma = \pm 1$ with probability $1/2$), we get, for the first term:

$$\begin{aligned}\mathbb{E}\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c}-\mathbf{c}^*\| \leq \sqrt{\delta}} |(P_n - P)f_{\tilde{\mathbf{c}}_0}|\right) &\leq 2\mathbb{E}_X\mathbb{E}_\sigma\left(\sup_{\mathbf{c}^* \in \mathcal{M}, \|\mathbf{c}-\mathbf{c}^*\| \leq \sqrt{\delta}} \frac{1}{n}\sum_{i=1}^{n}\sigma_i f_{\tilde{\mathbf{c}}_0}(X_i)\right) \\ &\leq 2\sqrt{2}\mathbb{E}_X\left(\sqrt{\sup_{\mathbf{c}^*, \mathbf{c}} \|f_{\tilde{\mathbf{c}}_0}\|^2_{l^2(S_n)} \log(m(\varepsilon_0))}\right) \\ &\leq 2\sqrt{2}\mathbb{E}_X\left(\sqrt{\|F\|^2_{l^2(S_n)} \log(m(\varepsilon_0))}\right) \\ &\leq 2\sqrt{2}\mathbb{E}_X\left(\sqrt{C_F^2 \log(m(\varepsilon_0))}\right) \\ &\leq \frac{\kappa_A}{\sqrt{n}},\end{aligned}$$

where κ_A depends on k, d and P. In the second line of this inequality, we used the maximal inequality for random processes depending only on Rademacher variables given by Massart [18, Lemma 2.3].

It remains to bound the second term. Using the same approach (symmetrization and maximal inequality for Rademacher variables) we get, for every $j > 0$,

$$\begin{aligned}\mathbb{E}&\left(\sup_{\mathbf{c},\mathbf{c}^*}\left|(P_n - P)(f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}})\right|\right)\\ &\leq 2\mathbb{E}_X\left(\sqrt{\frac{2}{n}\log(m(\varepsilon_j)m(\varepsilon_{j-1}))\sup_{\mathbf{c},\mathbf{c}^*}\|f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}}\|^2_{l^2(S_n)}}\right).\end{aligned}$$

However $\|f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}}\|_{l^2(S_n)} \leq \varepsilon_j\|F\|_{l^2(S_n)}$, consequently

$$\begin{aligned}\|f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}}\|^2_{l^2(S_n)} &\leq 4\varepsilon^2_{j-1}\|F\|^2_{l^2(S_n)}\\ &\leq 4C_F^2\varepsilon^2_{j-1}.\end{aligned}$$

Comparing a sum with an integral, we obtain

$$\sum_{j>0}\mathbb{E}\left(\sup_{\mathbf{c}^*\in\mathcal{M},\|\mathbf{c}-\mathbf{c}^*\|\leq\sqrt{\delta}}\left|(P_n - P)(f_{\tilde{\mathbf{c}}_j} - f_{\tilde{\mathbf{c}}_{j-1}})\right|\right) \leq \frac{32}{\sqrt{n}}\int_0^{\varepsilon_1}\sqrt{\log(m(\varepsilon))}d\varepsilon,$$

which, by assumption on $m(\varepsilon)$, can be bounded from above by $\frac{\kappa_B}{\sqrt{n}}$, where κ_B depends on k, d and P.

We are now in position to prove Proposition 5.1. From the two above subsections we deduce that

$$\mathbb{E}\left(\sup_{\mathbf{c}^*\in\mathcal{M},\|\mathbf{c}-\mathbf{c}^*\|^2\leq A_2\delta}|(P_n - P)(\gamma(\mathbf{c}^*,.) - \gamma(\mathbf{c},.))|\right) \leq \sqrt{\delta}\frac{\Xi}{\sqrt{n}}.$$

This concludes the proof.

5.6. Proof of Theorem 3.2

Let $\mathbf{x} = (x_1, \ldots, x_k)$ be a $k \times d$ vector, $V_1^*, \ldots, V_k^*$ the Voronoi cells associated with an optimal codebook $\mathbf{c}^*$. We state here a sufficient condition for the Hessian matrix $H(\mathbf{c}^*)$ to be positive. Denote $r_{ij} = \|c_i^* - c_j^*\|$. It holds

$$\langle H\mathbf{x}, \mathbf{x}\rangle = \sum_{i=1}^k\left[\langle H_{i,i}x_i, x_i\rangle + \sum_{j\neq i}\langle H_{i,j}x_j, x_i\rangle\right].$$

Recalling the expression of $H_{i,i}$ and $H_{i,j}$ given in equation 1,

$$H_{i,j} = \left\{\begin{array}{lll} 2P(V_i) - 2\sum_{\ell\neq i}r_{i\ell}^{-1}\sigma\left[f(x)(x-c_i)(x-c_i)^t\mathbf{1}_{\partial(V_i\cap V_\ell)}\right] & \text{for} & i = j\\ 2{r_{ij}}^{-1}\sigma\left[f(x)(x-c_i)(x-c_j)^t\mathbf{1}_{\partial(V_i\cap V_j)}\right] & \text{for} & i \neq j\end{array}\right.,$$

we may write, for $i = 1, \dots, k$,

$$\begin{aligned}\langle H_{i,i}x_i, x_i\rangle + \sum_{j\neq i}\langle H_{i,j}x_j, x_i\rangle = \quad & 2P(V_i^*)\|x_i\|^2 \\ & - 2x_i^t\left(\sum_{j\neq i} r_{i,j}^{-1}\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_i^*)^t du\right)x_i \\ & \qquad + 2x_i^t\sum_{j\neq i} r_{i,j}^{-1}\left(\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_j^*)^t du\right)x_j.\end{aligned}$$

The support of P is included in $\mathcal{B}(0,1)$, thus we can replace $\partial(V_i^*\cap V_j^*)$ with $\partial(V_i^*\cap V_j^*)\cap\mathcal{B}(0,1)$ in the equations above. However, to lighten notation, we will omit the indication and implicitly assume that every set we consider is contained in $\mathcal{B}(0,1)$. Let $p_{i,j} = \int_{\partial(V_i^*\cap V_j^*)} f(u)du$ be the $d-1$-dimensional P-measure of the boundary between V_i^* and V_j^*. Recalling that the underlying norm is the Euclidean norm, even for matrices, we may write

$$\begin{aligned}\langle H_{i,i}x_i, x_i\rangle + \sum_{i\neq j}\langle H_{i,j}x_j, x_i\rangle \geq \quad & 2P(V_i)\|x_i\|^2 \\ & - 2\|x_i\|^2\left\|\sum_{j\neq i} r_{i,j}^{-1}\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_i^*)^t du\right\| \\ & \quad - 2\|x_i\|\left\|\sum_{j\neq i} r_{i,j}^{-1}\left(\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_j^*)^t du\right)x_j\right\|,\end{aligned}$$

with

$$\begin{aligned}&\left\|\sum_{j\neq i} r_{i,j}^{-1}\left(\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_j^*)^t du\right)x_j\right\| \\ &\qquad\qquad \leq \sum_{j\neq i} r_{i,j}^{-1}\left\|\left(\int_{\partial(V_i^*\cap V_j^*)} f(u)(u-c_i^*)(u-c_j^*)^t du\right)x_j\right\| \\ &\qquad\qquad \leq \sum_{j\neq i} r_{i,j}^{-1}\left(\int_{\partial(V_i^*\cap V_j^*)} f(u)\|u-c_i^*\|\|u-c_j^*\|du\right)\|x_j\| \\ &\qquad\qquad \leq \sum_{j\neq i} r_{i,j}^{-1}p_{i,j}4\|x_j\|.\end{aligned}$$

Next,

$$\langle H_{i,i}x_i, x_i\rangle + \sum_{j\neq i}\langle H_{i,j}x_j, x_i\rangle \geq \left(2P(V_i^*) - \frac{8}{B}\sum_{i\neq j} p_{i,j}\right)\|x_i\|^2 - \frac{8}{B}\sum_{j\neq i} p_{i,j}\|x_i\|\|x_j\|,$$

where we recall that $B = \inf_{i \neq j, \mathbf{c}^* \in \mathcal{M}} \|c_i^* - c_j^*\|$. Making use of the inequality $2\|x_i\|\|x_j\| \leq \|x_i\|^2 + \|x_j\|^2$, and summing with respect to i leads to

$$\langle H\mathbf{x}, \mathbf{x} \rangle \geq \sum_{i=1}^{k} \left(2P(V_i) - \frac{16}{B} \sum_{j \neq i} p_{i,j} \right) \|x_i\|^2.$$

The last step is to derive bounds for $p_{i,j}$ from the conditions on f. Denote $\lambda = \|f\|_\infty$, we see that

$$\sum_{j \neq i} p_{i,j} = \int_{\partial V_i^*} f(u) du.$$

V_i^* is a regular convex set included in $\mathcal{B}(c_i^*, 2)$. Therefore, by a direct application of Stokes Theorem, the surface of ∂V_i^* is smaller than the surface of $\mathcal{S}_{d-1}(c_i^*, 2)$ (the sphere of radius 2). Consequently

$$\sum_{j \neq i} p_{i,j} \leq \lambda \frac{2\pi^{d/2}}{\Gamma(d/2)} 2^{d-1}.$$

It follows that $\lambda < \frac{B\Gamma(d/2)}{2^{d+3}\pi^{d/2}} \inf_{i=1,\ldots,k} P(V_i^*)$ is enough to ensure that the Hessian matrix $H(\mathbf{c}^*)$ is positive definite.

5.7. *Proof of Proposition 4.1*

We consider a distribution on $\mathbb{R}^d$, distributed over small balls away from one another, and whose density inside each ball is a small cone, for continuity reasons. Denote by V_i the Voronoi cell associated with z_i in $(z_1, \ldots, z_k)$. Let Q be a k-quantizer, Q^* the expected optimal quantizer which maps V_i to z_i for all i. Denote finally, for all $i = 1, \ldots, k$, $R_i(Q) = \int_{V_i} \|x - Q(x)\|^2 dx$ the contribution of the i-th Voronoi cell to the risk of Q.

Let S denote the surface of the unit ball in $\mathbb{R}^d$. Taking into account that $N_\rho = \frac{kS\rho^d}{d(d+1)}$ we have

$$\begin{aligned} R_i(Q^*) &= \frac{1}{kN_\rho} \int_0^\rho S r^{d+1} \left(1 - \frac{r}{\rho}\right) dr \\ &= \frac{\rho^2 d(d+1)}{k(d+3)(d+2)}. \end{aligned}$$

Let i be an integer between 1 and k. Let $m_i^{in} = |Q(\mathcal{B}_d(z_i, \rho)) \cap V_i|$ be the number of images of V_i sent by Q inside V_i, and let $m_i^{out} = |Q(\mathcal{B}_d(z_i, \rho)) \cap V_i^c|$ be the number of images of V_i sent outside V_i. The three situations of interest are the following ones:

→ If $m_i^{in} = 1$ and $m_i^{out} = 0$, it is clear that $R_i(Q) \geq R_i(Q^*)$.

- → If $m_i^{in} \geq 2$ and $m_i^{out} = 0$, then we just can see that $R_i(Q) \geq R_i(Q^*) - \frac{\rho^2 d(d+1)}{k(d+2)(d+3)} = 0$.
- → At last, suppose that $m_i^{out} \geq 1$. Then there exists $x \in \mathcal{B}_d(z_i, \rho)$ such that $Q(x) \notin V_i$. Since Q is a nearest neighbor quantizer, for such an x we have

$$\begin{cases} \|Q(x) - x\| \leq \inf_{c \in Q(\mathcal{B}_d(z_i,\rho))} \|x - c\| \\ \|Q(x) - x\| \geq d(z_i, V_i^c) - \rho \geq \dfrac{R}{2} - \rho \end{cases}.$$

Let $c \in Q(\mathcal{B}_d(z_i, \rho))$. Then

$$\begin{aligned} \|c - z_i\| &\geq \|c - x\| - \rho \\ &\geq \|Q(x) - x\| - \rho \\ &\geq \frac{R}{2} - 2\rho. \end{aligned}$$

Then, we deduce that, for every $y \in \mathcal{B}_d(z_i, \rho)$ and codepoint $c \in Q(\mathcal{B}_d(z_i, \rho))$, $\|y - c\| \geq \frac{R}{2} - 3\rho$. Therefore

$$\begin{aligned} R_i(Q) &\geq \frac{\left(\frac{R}{2} - 3\rho\right)^2}{k} \\ &\geq R_i(Q^*) + \frac{1}{k}\left(\left(\frac{R}{2} - 3\rho\right)^2 - \frac{\rho^2 d(d+1)}{(d+2)(d+3)}\right). \end{aligned}$$

Now suppose that $m_i^{in} \geq 2$. Then at least two code points of Q lies in V_i. Therefore, there exists j such that no code point of Q lies in V_j, so that $m_j^{out} \geq 1$. We straightforward deduce that the number of cells V_i for which $m_i^{in} \geq 2$ is smaller than the number of cells for which $m_j^{out} \geq 1$.

Taking into account all contributions of Voronoi cells, we get

$$\begin{aligned} R(Q) = &\sum_{\{i; m_i^{in} \geq 2, m_i^{out} = 0\}} R_i(Q) + \sum_{\{i; m_i^{out} \geq 1\}} R_i(Q) + \sum_{\{i; m_i^{in} = 1, m_i^{out} = 0\}} R_i(Q) \\ &\geq R(Q^*) + \sum_{\{i; m_i^{in} \geq 2, m_i^{out} = 0\}} \frac{1}{k}\left(\left(\frac{R}{2} - 3\rho\right)^2 - \frac{2\rho^2 d(d+1)}{(d+2)(d+3)}\right), \end{aligned}$$

from which we deduce a sufficient condition to get $R(Q) \geq R(Q^*)$.

5.8. Proof of Proposition 4.2

We begin with a lemma which ensures that every possible optimal code point c_i^* is close to at least one mean m_j of the mixture, in the case where the ratio p_{min}/p_{max} is large enough.

Lemma 5.6. *Let* $\mathbf{c}^*$ *be an optimal codebook. Suppose that*

$$\frac{p_{min}}{p_{max}} \geq \frac{288k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/288\sigma^2})}.$$

Then, for every $j = 1, \dots, k$, *there exists* $i \in \{1, \dots k\}$ *such that* $\|m_j - c_i^*\| \leq \frac{\tilde{B}}{6}$.

Proof of Lemma 5.6. Denote by $\mathbf{m}$ the codebook $(m_1, \dots, m_k)$, and by M_i the Voronoi cell associated with m_i. We bound from above the quantity $P\gamma(\mathbf{m}, .)$:

$$\begin{aligned} P\gamma(\mathbf{m}, .) &= \sum_{i=1}^{k} \frac{p_i}{2\pi\sigma^2 N_i} \int_{M_i} \|x - m_i\|^2 e^{-\|x-m_i\|^2/2\sigma^2} dx \\ &\leq \sum_{i=1}^{k} \frac{p_i}{2\pi\sigma^2 N_i} \int_{\mathbb{R}^2} \|x - m_i\|^2 e^{-\|x-m_i\|^2/2\sigma^2} dx \\ &\leq \frac{2kp_{max}\sigma^2}{(1-\varepsilon)}. \end{aligned}$$

Let $\mathbf{c}$ be a codebook such that there exists j satisfying, for all $i = 1, \dots, k$, $\|m_j - c_i\| > \tilde{B}/6$. We will prove that $P\gamma(\mathbf{c}, .) > P\gamma(\mathbf{m}, .)$, which implies that $\mathbf{c} \notin \mathcal{M}$. In fact we have, for all $i = 1, \dots, k$ and for all $x \in \mathcal{B}(m_j, \tilde{B}/12)$, $\|x - c_i\| > \tilde{B}/12$. Hence, a lower bound for $P\gamma(\mathbf{c}, .)$ is

$$\begin{aligned} P\gamma(\mathbf{c}, .) &\geq \int_{\mathcal{B}(m_j, \tilde{B}/12)} \min_{i=1,\dots,k} \|x - c_i\|^2 f(x) dx \\ &> \frac{\tilde{B}^2}{144} \sum_{i=1}^{k} \frac{p_i}{2\pi\sigma^2 N_i} \int_{\mathcal{B}(m_j, \tilde{B}/12)} e^{-\|x-m_i\|^2/2\sigma^2} dx \\ &> \frac{\tilde{B}^2 p_j}{288\pi\sigma^2 N_j} \int_{\mathcal{B}(m_j, \tilde{B}/12)} e^{-\|x-m_j\|^2/2\sigma^2} dx \\ &> \frac{p_{min}\tilde{B}^2}{144} \left(1 - e^{-\tilde{B}^2/288\sigma^2}\right) \\ &> P\gamma(\mathbf{m}, .). \end{aligned}$$

Hence we deduce that every optimal codebook has a code point close to every mean m_j of the mixture, of at most $\tilde{B}/6$. □

Suppose that the ratio p_{min}/p_{max} satisfies the assumption of Proposition 4.2. In particular p_{min}/p_{max} satisfies the assumption of Lemma 5.6. Then we deduce that, up to a re indexation, for every $\mathbf{c}^* \in \mathcal{M}$, $\|c_i^* - m_i\| \leq \tilde{B}/6$. We conclude that $2\tilde{B}/3 \leq B \leq 4\tilde{B}/3$.

Since, for all $i = 1, \dots, k$, $\mathcal{B}(c_i^*, B/2) \subset V_i^*$, it is easy to see that $\mathcal{B}(m_i, B/4) \subset \mathcal{B}(c_i^*, B/2) \subset V_i^*$, which leads to $N^* \subset \left(\bigcup_{i=1}^{k} \mathcal{B}(m_i, B/4)\right)^c$. Consequently, in order to apply Theorem 3.2, we just have to prove that

$$\|f_{\left|\left(\bigcup_{i=1}^{k} \mathcal{B}(m_i, B/4)\right)^c\right.}\|_\infty \leq \frac{\Gamma(1) B}{2^5\pi} \inf_{i=1,\dots,k} P\left(\mathcal{B}(m_i, B/4)\right).$$

First we derive a lower bound for the right-hand side. For every $i = 1, \dots, k$,

$$\begin{aligned} P(\mathcal{B}(m_i, B/4)) &\geq \frac{p_i}{N_i} \frac{1}{2\pi\sigma^2} \int_{\mathcal{B}(0,B/4)} e^{-\frac{\|x\|^2}{2\sigma^2}} dx \\ &\geq \frac{p_i}{N_i} \frac{1}{2\pi\sigma^2} \times 2\pi \int_0^{B/4} r e^{-\frac{r^2}{2\sigma^2}} dr \\ &\geq p_{min} \left(1 - e^{-\frac{B^2}{32\sigma^2}}\right). \end{aligned}$$

Then, we deal with the left-hand side. Let x be at distance from every m_i of at least $B/4$. Then

$$\begin{aligned} f(x) &\leq \sum_{i=1}^{k} \frac{p_i}{N_i} \frac{1}{2\pi\sigma^2} e^{-\frac{B^2}{32\sigma^2}} \\ &\leq \frac{k p_{max}}{2\pi\sigma^2(1-\varepsilon)} e^{-\frac{B^2}{32\sigma^2}}. \end{aligned}$$

The rest of the proof follows from straightforward computation, using the assumption of Proposition 4.2 and the relationship between B and $\tilde{B}$: $2\tilde{B}/3 \leq B \leq 4\tilde{B}/3$.

Remark. A careful reader should have noticed that the k factor is suboptimal in the previous inequality. In fact we are able in this case to bound from above $f(x)$ with $\frac{1}{2\pi\sigma^2(1-\varepsilon)} e^{-\frac{B^2}{32\sigma^2}}$. However, this bound does not involve p_{max}, and so involve a condition not on the ratio of extremal proportions of the mixture, but rather on the minimal proportion of the mixture, which is less natural. Moreover, the p_{max}-free bound is valid only in the equal variance case, namely when the variance σ_i^2 of any element of the mixture is the same. In general it is not the case and a condition as in Proposition 4.2 for that kind of mixture would naturally involve the ratio p_{min}/p_{max}.

Acknowledgement

The author would like to thank two referees for valuable comments and suggestions.

References

[1] András Antos. Improved minimax bounds on the test and training distortion of empirically designed vector quantizers. *IEEE Trans. Inform. Theory*, 51(11):4022–4032, 2005. MR2239018

[2] András Antos, László Györfi, and András György. Individual convergence rates in empirical vector quantizer design. *IEEE Trans. Inform. Theory*, 51(11):4013–4022, 2005. MR2239017

[3] Adrian Baddeley. Integrals on a moving manifold and geometrical probability. *Advances in Appl. Probability*, 9(3):588–603, 1977. MR0471017

[4] Peter L. Bartlett, Tamás Linder, and Gábor Lugosi. The minimax distortion redundancy in empirical quantizer design. *IEEE Trans. Inform. Theory*, 44(5):1802–1813, 1998. MR1664098

[5] Gérard Biau, Luc Devroye, and Gábor Lugosi. On the performance of clustering in Hilbert spaces. *IEEE Trans. Inform. Theory*, 54(2):781–790, 2008. MR2444554

[6] Gilles Blanchard, Olivier Bousquet, and Pascal Massart. Statistical performance of support vector machines. *Ann. Statist.*, 36(2):489–531, 2008. MR2396805

[7] Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *C. R. Math. Acad. Sci. Paris*, 334(6):495–500, 2002. MR1890640

[8] Benoît Cadre and Quentin Paris. On Hölder fields clustering. *TEST*, 21(2):301–316, 2012. MR2935361

[9] Philip A. Chou. The distortion of vector quantizers trained on n vectors decreases to the optimum as $\mathcal{O}_p(1/n)$. In *Proc. IEEE Int. Symp. Inf. Theory*, page 457, Trondheim, Norway, 1994.

[10] Aurélie Fischer. Quantization and clustering with Bregman divergences. *J. Multivariate Anal.*, 101(9):2207–2221, 2010. MR2671211

[11] Allen Gersho and Robert M. Gray. *Vector quantization and signal compression*. Kluwer Academic Publishers, Norwell, MA, USA, 1991.

[12] Evarist Giné and Joel Zinn. Some limit theorems for empirical processes. *Ann. Probab.*, 12(4):929–998, 1984. With discussion. MR0757767

[13] Siegfried Graf and Harald Luschgy. *Foundations of quantization for probability distributions*, volume 1730 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2000. MR1764176

[14] Stefan Junglen. *Geometry of optimal codebooks and constructive quantization*. PhD thesis, Universität Trier, Universitätsring 15, 54296 Trier, 2012.

[15] Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Statist.*, 34(6):2593–2656, 2006. MR2329442

[16] Tamás Linder. Learning-theoretic methods in vector quantization. In *Principles of nonparametric learning (Udine, 2001)*, volume 434 of *CISM Courses and Lectures*, pages 163–210. Springer, Vienna, 2002. MR1987659

[17] Tamás Linder, Gábor Lugosi, and Kenneth Zeger. Rates of convergence in the source coding theorem, in empirical quantizer design, and in universal lossy source coding. *IEEE Trans. Inform. Theory*, 40(6):1728–1740, 1994. MR1322387

[18] Pascal Massart. *Concentration inequalities and model selection*, volume 1896 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard. MR2319879

[19] Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *Ann. Statist.*, 34(5):2326–2366, 2006. MR2291502

[20] Neri Merhav and Jacob Ziv. On the amount of statistical side information required for lossy data compression. *IEEE Trans. Inform. Theory*, 43(4):1112–1121, 1997. MR1454940

[21] David Pollard. Strong consistency of k-means clustering. *Ann. Statist.*, 9(1):135–140, 1981. MR0600539

[22] David Pollard. A central limit theorem for empirical processes. *J. Austral. Math. Soc. Ser. A*, 33(2):235–248, 1982. MR0668445

[23] David Pollard. A central limit theorem for k-means clustering. *Ann. Probab.*, 10(4):919–926, 1982. MR0672292

[24] David Pollard. Quantization and the method of k -means. *IEEE Transactions on Information Theory*, 28(2):199–204, 1982. MR0651814

[25] Thaddeus Tarpey. Principal points and self-consistent points of symmetric multivariate distributions. *J. Multivariate Anal.*, 53(1):39–51, 1995. MR1333126