



HAL
open science

Non Asymptotic Bounds for Vector Quantization in Hilbert Spaces

Clément Levrard

► **To cite this version:**

| Clément Levrard. Non Asymptotic Bounds for Vector Quantization in Hilbert Spaces. 2013. <hal-00877564v2>

HAL Id: hal-00877564

<https://hal.science/hal-00877564v2>

Preprint submitted on 26 May 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

NON ASYMPTOTIC BOUNDS FOR VECTOR QUANTIZATION IN HILBERT SPACES

BY CLÉMENT LEVRARD

Université Paris Sud, UPMC and INRIA

Recent results in quantization theory show that the convergence rate for the mean-squared expected distortion of the empirical risk minimizer strategy, for any fixed probability distribution satisfying some regularity conditions, is $\mathcal{O}(1/n)$, where n is the sample size (see, e.g., [8] or [17]). However, the dependency of the average distortion on other parameters is not known, and these results are valid for distributions over finite dimensional Euclidean spaces.

This paper deals with the general case of distributions over separable, possibly infinite dimensional, Hilbert spaces. A condition is proposed, which may be thought of as a margin condition (see, e.g., [21]), under which a non asymptotic upper bound on the expected distortion rate of the empirically optimal quantizer is derived. The dependency of the distortion on other natural parameters of the quantization issue is then discussed, in particular through a minimax lower bound.

1. Introduction. Quantization, also called lossy data compression in information theory, is the problem of replacing a probability distribution with an efficient and compact representation, that is a finite set of points. To be more precise, let \mathcal{H} denote a separable Hilbert space, and let P denote a probability distribution over \mathcal{H} and k a positive integer. A so-called k -points quantizer Q is a map from \mathcal{H} to \mathcal{H} , whose image set is made of exactly k points, that is $|Q(\mathcal{H})| = k$. For such a quantizer, every image point $c_i \in Q(\mathcal{H})$ is called a code point, and the vector composed of the code points (c_1, \dots, c_k) is called a codebook, denoted by \mathbf{c} . By considering the preimages of its code points, a quantizer Q partitions the separable Hilbert space \mathcal{H} into k groups, and assigns each group a representative. General references on the subject are to be found in [14], [13] and [19] among others.

The quantization theory was originally developed as a way to answer signal compression issues in the late 40's (see, e.g., [13]). However, unsupervised classification is also in the scope of its application. Isolating meaningful groups from a cloud of data is a topic of interest in many fields, from social science to biology. Classifying points into dissimilar groups of similar items is more interesting as the amount of accessible data is large. In many cases

Keywords and phrases: quantization, localization, fast rates, margin conditions

data need to be preprocessed through a quantization algorithm in order to be exploited.

If the distribution P has a finite second moment, the performance of a quantizer Q is measured by the risk, or distortion

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

where Pf means integration of the function f with respect to P . The choice of the squared norm is convenient, since it takes advantages of the Hilbert space structure of \mathcal{H} . Nevertheless, it is worth pointing out that several authors deal with more general distortion functions. For further information on this topic, the interested reader is referred to [14] or [12].

In order to minimize the distortion introduced above, it is clear that only quantizers of the type $x \mapsto \arg \min_{c_1, \dots, c_k} \|x - c_i\|^2$ are to be considered. Such quantizers are called nearest-neighbor quantizers. With a slight abuse of notation, $R(\mathbf{c})$ will denote the risk of the nearest-neighbor quantizer associated with a codebook \mathbf{c} .

Provided that P has a bounded support, there exist optimal codebooks minimizing the risk R (see, e.g., Corollary 3.1 in [12] or Theorem 1 in [15]). The aim is to design a codebook $\hat{\mathbf{c}}_n$, according to an n -sample drawn from P , whose distortion is as close as possible to the optimal distortion $R(\mathbf{c}^*)$, where \mathbf{c}^* denotes an optimal codebook.

To solve this problem, most approaches to date attempt to implement the principle of empirical risk minimization in the vector quantization context. Let X_1, \dots, X_n denote an independent and identically distributed sample with distribution P . 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

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

If the training data represents the source well, then $\hat{\mathbf{c}}_n$ will hopefully also perform near optimally on the real source, that is $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) = R(\hat{\mathbf{c}}_n) - R(\mathbf{c}^*) \approx 0$. The problem of quantifying how good empirically designed codebooks are, compared to the truly optimal ones, has been extensively studied, as for instance in [19] in the finite dimensional case.

In the case where $\mathcal{H} = \mathbb{R}^d$, for some $d > 0$, it has been proved in [20] that $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) = \mathcal{O}(1/\sqrt{n})$, provided that P has a bounded support. This result has been extended to the case where \mathcal{H} is a separable Hilbert space in [5]. However, this upper bound has been tightened whenever the source

distribution satisfies additional assumptions, in the finite dimensional case only.

When $\mathcal{H} = \mathbb{R}^d$, for the special case of finitely supported distributions, it is shown in [2] that $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) = \mathcal{O}(1/n)$. There are much more results in the case where P is not assumed to have a finite support.

In fact, different sets of assumptions have been introduced in [2], [24] or [17], to derive fast convergence rates for the distortion in the finite dimensional case. To be more precise, it is proved in [2] that, if P satisfies a technical inequality for every codebook \mathbf{c} , namely

$$(1) \quad \ell(\mathbf{c}, \mathbf{c}^*) \geq a \operatorname{Var} \left(\min_{j=1, \dots, k} \|X - c_j\|^2 - \min_{j=1, \dots, k} \|X - c_j^*\|^2 \right),$$

for some $a > 0$, then $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C(k, d, P) \log(n)/n$, where $C(k, d, P)$ depends on the natural parameters k and d , but also on the technical parameter a . However, in the continuous density and unique minimum case, it has been proved in [9], following the approach of [24], that, provided that the Hessian matrix of $\mathbf{c} \mapsto R(\mathbf{c})$ is positive definite at the optimal codebook, $n\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ converges in distribution to a law, depending on the Hessian matrix. As proved in [17], the technique used in [24] can be slightly modified to derive a non-asymptotic bound of the type $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C/n$ in this case, for some unknown $C > 0$.

As shown in [17], these different sets of assumptions turn out to be equivalent in the continuous density case to a technical condition, similar to that used in [23] to derive fast rates of convergence in the statistical learning framework.

Thus, a question of interest is to know whether some margin type conditions can be derived for the source distribution to satisfy the technical condition mentioned above, as has been done in the statistical learning framework in [21]. This paper provides a condition, which can clearly be thought of as a margin condition in the quantization framework, under which the condition (1) is satisfied, where the technical constant a has an explicit expression in term of natural parameters of the quantization issue, such as the smallest distance between two optimal code points. It is worth mentioning that this margin condition does not require \mathcal{H} to have a finite dimension, or P to have a continuous density. In the finite dimensional case, this condition does not require either that there exists a unique optimal codebook, as required in [24], hence seems easier to check.

Moreover, a non asymptotic bound of the type $\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C(k, P)/n$ is derived for distributions satisfying this margin condition, where $C(k, P)$ is explicitly given in terms of natural parameters of the quantization issue.

This bound is also valid in the case where \mathcal{H} has an infinite dimension. This point may be of interest for curve quantization, as done in [3].

In addition, a minimax lower bound is given which allows to discuss the influence of the different parameters mentioned in the upper bound. It is worth pointing out that this lower bound is valid over a set of probability distributions with uniformly bounded continuous densities and unique optimal codebook, such that the minimum eigenvalue of the second derivative matrices, at the optimal codebook, is uniformly lower bounded. This result refines the previous minimax bounds obtained in [1] or [4].

The paper is organized as follows. In Section 2 some notation and definition are introduced, along with some basic results for quantization in a Hilbert space. The so-called margin condition is then introduced, and the main results are exposed in Section 3: firstly an oracle inequality on the loss is stated, along with a minimax result. Then it is shown that Gaussian mixtures are in the scope of the margin conditions. Finally, proofs are gathered in Section 4.

2. Notation and Definitions. Throughout the paper, for $M > 0$ and a in \mathcal{H} , $\mathcal{B}(a, M)$ will denote the closed ball with center a and radius M . With a slight abuse of notation, P is said to be M -bounded if its support is included in $\mathcal{B}(0, M)$. Furthermore, it will also be assumed that the support of P contains more than k points.

To frame the quantization issue as an empirical risk minimization issue, the following contrast function γ is introduced as

$$\gamma : \begin{cases} (\mathcal{H})^k \times \mathcal{H} & \longrightarrow \mathbb{R} \\ (\mathbf{c}, x) & \longmapsto \min_{j=1, \dots, k} \|x - c_j\|^2, \end{cases}$$

where $\mathbf{c} = (c_1, \dots, c_k)$ denotes a codebook, that is a kd -dimensional vector if $\mathcal{H} = \mathbb{R}^d$. Throughout the paper, only the case $k \geq 2$ will be considered. The risk $R(\mathbf{c})$ then takes the form $R(\mathbf{c}) = R(Q) = P\gamma(\mathbf{c}, \cdot)$, where we recall that Pf denotes the integration of the function f with respect to P . Similarly, the empirical risk $\hat{R}_n(\mathbf{c})$ can be defined as $\hat{R}_n(\mathbf{c}) = P_n\gamma(\mathbf{c}, \cdot)$, where P_n is the empirical distribution associated with X_1, \dots, X_n , in other words $P_n(A) = 1/n |\{i | X_i \in A\}|$, for every measurable subset $A \subset \mathcal{H}$.

It is worth pointing out that, if P is M -bounded, for some $M > 0$, then there exist such minimizers $\hat{\mathbf{c}}_n$ and \mathbf{c}^* (see, e.g., Corollary 3.1 in [12]). In the sequel the set of minimizers of the risk $R(\cdot)$ will be denoted by \mathcal{M} . Since every permutation of labels of an optimal codebook provides an optimal codebook, \mathcal{M} contains more than $k!$ elements. To address the issue of a

large number of optimal codebooks, $\bar{\mathcal{M}}$ is introduced as a set of codebooks which satisfies

$$\begin{cases} \forall \mathbf{c}^* \in \mathcal{M} & \exists \bar{\mathbf{c}} \in \bar{\mathcal{M}} & \{c_1^*, \dots, c_k^*\} = \{\bar{c}_1, \dots, \bar{c}_k\}, \\ \forall \bar{\mathbf{c}}^1, \bar{\mathbf{c}}^2 \in \bar{\mathcal{M}} & & \{\bar{c}_1^1, \dots, \bar{c}_k^1\} \neq \{\bar{c}_1^2, \dots, \bar{c}_k^2\}. \end{cases}$$

In other words, $\bar{\mathcal{M}}$ is a subset of the set of optimal codebooks which contains every element of \mathcal{M} , up to a permutation of the labels, and in which two different codebooks have different sets of code points. It may be noticed that $\bar{\mathcal{M}}$ is not uniquely defined. However, when \mathcal{M} is finite, all the possible $\bar{\mathcal{M}}$ have the same cardinality.

Let c_1, \dots, c_k be a sequence of code points. A central role is played by the set of points which are closer to c_i than to any other c_j 's. To be more precise, the Voronoi cell, or quantization cell associated with c_i is the closed set defined by

$$V_i(\mathbf{c}) = \{x \in \mathcal{H} \mid \forall j \neq i \quad \|x - c_i\| \leq \|x - c_j\|\}.$$

It may be noted that $(V_1(\mathbf{c}), \dots, V_k(\mathbf{c}))$ does not form a partition of \mathcal{H} , since $V_i(\mathbf{c}) \cap V_j(\mathbf{c})$ may be non empty. To address this issue, a Voronoi partition associated with \mathbf{c} is defined as a sequence of subsets $(W_1(\mathbf{c}), \dots, W_k(\mathbf{c}))$ which forms a partition of \mathcal{H} , and such that for every $i = 1, \dots, k$,

$$\bar{W}_i(\mathbf{c}) = V_i(\mathbf{c}),$$

where $\bar{W}_i(\mathbf{c})$ denotes the closure of the subset $W_i(\mathbf{c})$. The open Voronoi cell is defined the same way by

$$\overset{o}{V}_i(\mathbf{c}) = \{x \in \mathcal{H} \mid \forall j \neq i \quad \|x - c_i\| < \|x - c_j\|\}.$$

Given a Voronoi partition $W(\mathbf{c}) = (W_1(\mathbf{c}), \dots, W_k(\mathbf{c}))$, the following inclusion holds, for i in $\{1, \dots, k\}$,

$$\overset{o}{V}_i(\mathbf{c}) \subset W_i(\mathbf{c}) \subset V_i(\mathbf{c}),$$

and the risk $R(\mathbf{c})$ takes the form

$$R(\mathbf{c}) = \sum_{i=1}^k P(\|x - c_i\|^2 1_{W_i(\mathbf{c})}(x)),$$

where 1_A denotes the indicator function associated with A . In the case where (W_1, \dots, W_k) are fixed subsets such that $P(W_i) \neq 0$, for every $i = 1, \dots, k$, it is clear that

$$P(\|x - c_i\|^2 1_{W_i(\mathbf{c})}(x)) \geq P(\|x - \eta_i\|^2 1_{W_i(\mathbf{c})}(x)),$$

with equality only if $c_i = \eta_i$, where η_i denotes the conditional expectation of P over the subset $W_i(\mathbf{c})$, that is

$$\eta_i = \frac{P(x1_{W_i(\mathbf{c})}(x)}){P(W_i(\mathbf{c}))}.$$

Moreover, it is proved in Proposition 1 of [15] that, for every Voronoi partition $W(\mathbf{c}^*)$ associated with an optimal codebook \mathbf{c}^* , and every $i = 1, \dots, k$, $P(W_i(\mathbf{c}^*)) \neq 0$. Consequently, any optimal codebook satisfies the so-called centroid condition (see, e.g., Section 6.2 of [13]), that is

$$\mathbf{c}_i^* = \frac{P(x1_{W_i(\mathbf{c}^*)}(x)}){P(W_i(\mathbf{c}^*))}.$$

As a remark, the centroid condition ensures that, for every \mathbf{c}^* in \mathcal{M} and $i \neq j$,

$$\begin{aligned} P(V_i(\mathbf{c}^*) \cap V_j(\mathbf{c}^*)) &= P(\{x \in \mathcal{H} \mid \forall i' \quad \|x - c_i^*\| = \|x - c_j^*\| \leq \|x - c_{i'}^*\|\}) \\ &= 0. \end{aligned}$$

A proof of this statement can be found in Proposition 1 of [15]. According to this remark, it is clear that, for every optimal Voronoi partition $(W_1(\mathbf{c}^*), \dots, W_k(\mathbf{c}^*))$,

$$(2) \quad \begin{cases} P(W_i(\mathbf{c}^*)) &= P(V_i(\mathbf{c}^*)), \\ P_n(W_i(\mathbf{c}^*)) &\stackrel{a.s.}{=} P_n(V_i(\mathbf{c}^*)). \end{cases}$$

The following quantities are of importance in the bounds exposed in Section 3.1:

$$\begin{cases} B &= \inf_{\mathbf{c}^* \in \mathcal{M}, i \neq j} \|c_i^* - c_j^*\|, \\ p_{min} &= \inf_{\mathbf{c}^* \in \mathcal{M}, i=1, \dots, k} P(V_i(\mathbf{c}^*)). \end{cases}$$

It is worth noting here that $B \leq 2M$ whenever P is M -bounded, and $p_{min} \leq 1/k$. If \mathcal{M} is finite, it is clear that p_{min} and B are strictly positive. The following proposition ensures that this statement remains true when \mathcal{M} is not assumed to be finite.

PROPOSITION 2.1. *Suppose that P is M -bounded. Then both B and p_{min} are strictly positive quantities.*

A proof of Proposition 2.1 is given in Section 4. The role of the boundaries between optimal Voronoi cells may be compared to the role played by

the critical value $1/2$ for the regression function in the statistical learning framework. To draw this comparison, the following set is introduced, for any $\mathbf{c}^* \in \mathcal{M}$,

$$N(\mathbf{c}^*) = \bigcup_{i \neq j} V_i(\mathbf{c}^*) \cap V_j(\mathbf{c}^*).$$

Next, the critical region N^* is defined as

$$N^* = \bigcup_{\mathbf{c}^* \in \mathcal{M}} N(\mathbf{c}^*).$$

This region seems to be of importance when considering the conditions under which the empirical risk minimization strategy for the quantization issue achieves faster rates of convergence, as exposed in [17]. However, to fully draw the comparison between the margin conditions for the statistical learning issue (see, e.g., [21]) and quantization, the neighborhood of this region has to be introduced. For this purpose the t -neighborhood of the critical region N^* is defined as

$$N_t^* = \{x \in \mathcal{H} \mid d(x, N^*) \leq t\}.$$

Intuitively, if $P(N_t^*)$ is small enough, then the source distribution P is concentrated around its optimal codebook, and may be thought of as a slight modification of the probability distribution with finite support made of an optimal codebook \mathbf{c}^* . To be more precise, let us introduce the following key assumption:

DEFINITION 2.1 (Margin condition). *Denote by $p(t) = P(N_t^*)$. Then P satisfies a margin condition with radius r_0 if and only if*

- i) P is M -bounded,*
- ii) for all $0 \leq t \leq r_0$,*

$$(3) \quad p(t) \leq \frac{B p_{\min}}{128M^2} t.$$

Note that, since $p(2M) = 1$, $p_{\min} \leq 1/k$, $k \geq 2$ and $B \leq 2M$, (3) implies that $r_0 < 2M$. It is worth pointing out that Definition 2.1 does not require P to have a density or a unique optimal codebook, up to relabeling, contrary to the conditions introduced in [24].

It may be mentioned that the margin condition introduced here only requires a local control of the weight of the neighborhood of the critical

region N^* . The parameter r_0 may be thought of as a gap size around N^* , as illustrated by the following example:

Example 1: Assume that there exists $r > 0$ such that $p(x) = 0$ if $x \leq r$ (for instance if P is supported on k points). Then P satisfies (3), with radius r .

It is also worth pointing out that the condition mentioned in [21] requires a control of the weight of the neighborhood of the critical value $1/2$ with a polynomial function with degree larger than 1. In the quantization framework, the special role played by the exponent 1 leads to only consider linear controls of the weight function. This point is explained by the following example:

Example 2: Assume that P is M -bounded, and that there exists $Q > 0$ and $q > 1$ such that $p(x) \leq Qx^q$. Then P satisfies (3), with

$$r_0 = \left(\frac{p_{\min} B}{128 M^2 Q} \right)^{1/(q-1)}.$$

In the case where P has a density and $\mathcal{H} = \mathbb{R}^d$, the condition (3) can be thought of as a generalization of the condition mentioned in Theorem 3.2 of [17], which requires the density of the distribution to be small enough over the critical region N^* . In fact, provided that P has a continuous density, a uniform bound on the density over N^* provides a local control of $p(t)$ with a polynomial function of degree 1. This idea is developed in the following example:

Example 3(Continuous densities, $\mathcal{H} = \mathbb{R}^d$): Assume that $\mathcal{H} = \mathbb{R}^d$, P has a continuous density f and is M -bounded, and that \mathcal{M} is finite. In this case, $p(t)$ is differentiable at 0, with derivative

$$p'(0) = \int_{N^*} f(u) d\lambda_{d-1}(u),$$

where λ_{d-1} denotes the $(d-1)$ dimensional Lebesgue measure, considered over the $(d-1)$ dimensional space N^* . Therefore, if P satisfies

$$(4) \quad \int_{N^*} f(u) d\lambda_{d-1}(u) < \frac{B p_{\min}}{128 M^2},$$

then there exists $r_0 > 0$ such that P satisfies (3). It can easily be deduced from (4) that a uniform bound on the density located at the critical region N^* can provide a sufficient condition for a distribution P to satisfy a margin condition. Such a result has to be compared to Theorem 3.2 of [17], where it was required that

$$\|f|_{N^*}\|_{\infty} \leq \frac{\Gamma\left(\frac{d}{2}\right) B}{2^{d+5} M^{d+1} \pi^{d/2}} p_{\min},$$

where Γ denotes the Gamma function, and $f|_{N^*}$ denotes the restriction of f to the set N^* . Note however that the uniform bound mentioned above ensures that the Hessian matrices of the risk function R , at optimal codebooks, are positive definite. This does not necessarily implies that (3) is satisfied.

Another interesting parameter of the quantization issue is the following separation factor, which quantifies the difference between optimal codebooks and local minimizers of the risk.

DEFINITION 2.2. *Denote by $\tilde{\mathcal{M}}$ the set of local minimizers of the map $\mathbf{c} \mapsto P\gamma(\mathbf{c}, \cdot)$. Let $\varepsilon > 0$, then P is said to be ε -separated if*

$$(5) \quad \inf_{\mathbf{c} \in \tilde{\mathcal{M}} \cap \mathcal{M}^c} \ell(\mathbf{c}, \mathbf{c}^*) = \varepsilon.$$

It may be noticed that local minimizers of the risk function satisfy the centroid condition. Whenever $\mathcal{H} = \mathbb{R}^d$, P has a density and $P\|x\|^2 < \infty$, it can be proved that the set of minimizers of R coincides with the set of codebooks satisfying the centroid condition, also called stationary points (see, e.g., Lemma A of [24]). However, this result cannot be extended to non continuous distributions, as proved in Example 4.11 of [14].

The main results of the present paper are based on the following proposition, which connects the margin condition stated in Definition 2.1 to the condition introduced in Theorem 2 of [2]. It is recalled here that only the case $k \geq 2$ is considered.

PROPOSITION 2.2. *Assume that P satisfies a margin condition with radius r_0 , and is ε -separated. Then, for all $\mathbf{c} \in \mathcal{B}(0, M)$,*

i) there exists $\mathbf{c}^(\mathbf{c}) \in \mathcal{M}$ such that*

$$\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\| = \arg \inf_{\mathbf{c}^* \in \mathcal{M}} \|\mathbf{c} - \mathbf{c}^*\|,$$

ii) $\|\mathbf{c} - \mathbf{c}^(\mathbf{c})\|^2 \leq \kappa_0 \ell(\mathbf{c}, \mathbf{c}^*)$,*

where $\kappa_0 = 4kM^2 \left(\frac{1}{\varepsilon} \vee \frac{64M^2}{p_{\min} B^2 r_0^2} \right)$.

Moreover, if $\mathcal{H} = \mathbb{R}^d$, then \mathcal{M} is finite.

As mentioned in [8] or [17], the connection between the loss and the squared distance can be thought of as a technical margin condition. It is worth pointing out that the dependency of κ_0 on different parameters of the quantization issue is known. This fact allows us to roughly discuss how κ_0 should scale with the parameters k , d and M , in the finite dimensional case. According to Theorem 6.2 of [14], $R(\mathbf{c}^*)$ scales like $k^{-2/d}$, at least

in the density case. Furthermore, it is likely that $r_0 \sim B$ (see, e.g., the distributions exposed in Section 3.2). Considering that $\varepsilon \sim R(\mathbf{c}^*) \sim k^{-2/d}$, $r_0 \sim B \sim Mk^{-1/d}$, and $p_{min} \sim 1/k$ leads to

$$\kappa_0 \sim k^{2+4/d}.$$

At first sight κ_0 does not scale with M , and seems to decrease with the dimension, at least in the finite dimensional case. However, there is no result on how κ_0 should scale in the infinite dimensional case.

It is worth mentioning that, if $\mathcal{H} = \mathbb{R}^d$, P has a unique optimal codebook up to relabeling, and has a continuous density, Proposition 2.2 ensures that the second derivative matrix of R at the optimal codebook is positive definite, with minimum eigenvalue larger than $p_{min}/2$. This is the condition required in [24] for $n\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ to converge in distribution. This Proposition allows us to derive explicit upper bounds on the excess risk in the following section.

3. Results.

3.1. *Risk bound.* The main result of this paper is the following:

THEOREM 3.1. *Let k be larger than 2. Assume that \mathcal{M} is finite, P satisfies a margin condition with radius r_0 , and is ε -separated. Let κ_0 be defined as*

$$\kappa_0 = 4kM^2 \left(\frac{1}{\varepsilon} \vee \frac{64M^2}{p_{min}B^2r_0^2} \right).$$

If $\hat{\mathbf{c}}_n$ is an empirical risk minimizer, then, with probability larger than $1 - e^{-x}$,

$$(6) \quad \ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C'_0 \kappa_0 \frac{|\bar{\mathcal{M}}|^2 M^2 k}{n} + \kappa_0 \frac{144M^2}{n} x + \frac{64M^2}{n} x,$$

where C'_0 is an absolute constant.

Moreover, if $\mathcal{H} = \mathbb{R}^d$, with probability larger than $1 - e^{-x}$,

$$(7) \quad \ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C_0 \kappa_0 \frac{M^2 k d \left(\log(4|\bar{\mathcal{M}}|\sqrt{kd}) + 1 \right)}{n} + \kappa_0 \frac{144M^2}{n} x + \frac{64M^2}{n} x,$$

where C_0 is an absolute constant.

This result is in line with Theorem 3.1 in [17] or Theorem 1 in [8], concerning the dependency on the sample size n of the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$. The main

advance lies in the dependency on other parameters of the loss of $\hat{\mathbf{c}}_n$, which provides a non-asymptotic bound for the excess risk.

At first sight, in the finite dimensional case, (6) seems to outperform (7) when d is large. However the dependency on the number of optimal codebooks is dramatically worse in (6) than in (7). This difference can be explained by the two different methods used to derive these bounds.

In fact, most of the proof of (7) relies on the application of Dudley's entropy bound. This technique was already the main argument in [24], [17] or [8], and makes a classical dimension factor kd appear. This result slightly improves the asymptotic bound exposed in [24], since it offers an explicit calculation of the metric entropy used to derive this result.

As suggested in [5], the use of metric entropy techniques to derive bounds on the convergence rate of the distortion may be suboptimal, as it does not take advantage of the Hilbert space structure of the squared distance based quantization. This issue can be addressed using a more general chaining technique, such as the generic chaining principle developed in [25]. The second upper bound (6) is derived that way.

Another interesting point is that Theorem 3.1 does not require that P has a density or is distributed over points, contrary to the requirements of the previous bounds in [24], [2] or [8] which achieved the optimal rate of $\mathcal{O}(1/n)$. Up to our knowledge, the more general result is to be found in Theorem 2 of [2], which derives a convergence rate of $\mathcal{O}(\log(n)/n)$ without any requirement on the regularity of the distribution P . It may also be noted that, in the finite dimensional case, contrary to the results exposed in [24], Theorem 3.1 does not require that $\bar{\mathcal{M}}$ contains a single element. According to Proposition 2.2, only (3) has to be proved for P to satisfy the assumptions of Theorem 3.1. Since proving that $|\bar{\mathcal{M}}| = 1$ may be difficult, even for simple distributions, it seems easier to check the assumptions of Theorem 3.1 than the assumptions required in [24]. An illustration of this point is given in Section 3.3.

It is also worth mentioning that the dependency in ε surprisingly turns out to be sharp when $\varepsilon \sim n^{-1/2}$, as will be shown in Proposition 3.1. In fact, tuning this separation factor is the core of the demonstration of the minimax results in [4] or [1].

3.2. Minimax lower bound. Theorem 1 in [4] ensures that the minimax convergence rate over the M -bounded distributions of any empirically designed codebook can be bounded from below by $\Omega(1/\sqrt{n})$. A question of interest is to know whether this lower bound can be refined when considering only distributions satisfying some fast convergence condition. A partial

answer is given by Corollary 2 in [1], where it is proved that the minimax rate over distributions with uniformly bounded continuous densities, unique optimal codebook (up to relabeling), and such that the second derivative matrices at the optimal codebook $H(\mathbf{c}^*)$ are positive definite, is still $\Omega(1/\sqrt{n})$. According to [1], a natural question is to know whether a uniform upper bound of the type $o(1/\sqrt{n})$ may be derived, with the additional requirement that the minimum eigenvalue of the second derivative matrices $H(\mathbf{c}^*)$ is uniformly bounded from below.

This subsection is devoted to obtaining a minimax lower bound on the excess risk over a set of distributions with continuous densities, unique optimal codebook, and satisfying the margin condition defined in Definition 2.1, in which some parameters, such as p_{min} are fixed or uniformly lower-bounded. Since, in this case, the minimum eigenvalues of $H(\mathbf{c}^*)$ are larger than $p_{min}/2$, such a minimax lower bound provides an answer to the question mentioned above.

Throughout this subsection, only the case $\mathcal{H} = \mathbb{R}^d$ is considered, and $\hat{\mathbf{c}}_n$ will denote an empirically designed codebook, that is a map from $(\mathbb{R}^d)^n$ to $(\mathbb{R}^d)^k$. Let k be an integer such that $k \geq 3$, and $M > 0$. For simplicity, k is assumed to be divisible by 3. Let us introduce the following quantities:

$$\begin{cases} m &= \frac{2k}{3}, \\ \Delta &= \frac{5M}{32m^{1/d}}. \end{cases}$$

To focus on the dependency on the separation factor ε , the quantities involved in Definition 2.1 are fixed as:

$$(8) \quad \begin{cases} B &= \Delta, \\ r_0 &= \frac{7\Delta}{16}, \\ p_{min} &\geq \frac{1}{2k}. \end{cases}$$

Denote by $\mathcal{D}(\varepsilon)$ the set of probability distributions which are ε -separated, have a continuous density and a unique optimal codebook, and which satisfy a margin condition with parameters defined in (8). The minimax result is the following:

PROPOSITION 3.1. *Assume that $k \geq 3$. Then, for any empirically designed codebook,*

$$\sup_{P \in \mathcal{D}(c_1/\sqrt{n})} \mathbb{E} \ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \geq c_0 M^2 \frac{\sqrt{k^{1-\frac{4}{d}}}}{\sqrt{n}},$$

where c_0 is an absolute constant, and

$$c_1 = \frac{(5M)^2}{4(32m^{\frac{1}{4}+\frac{1}{d}})^2}.$$

Proposition 3.1 is in line with the previous minimax lower bounds obtained in Theorem 1 of [4] or Theorem 4 in [1]. In fact, the classes of distributions used in both these results satisfy a uniform margin condition, without specification of the separation factor. Proposition 3.1, as well as these two previous results, emphasizes the fact that fixing the parameters of the margin condition uniformly over a class of distributions does not guarantee an optimal uniform convergence rate. This shows that a uniform separation assumption is needed to derive a sharp uniform convergence rate over a set of distributions.

Furthermore, as mentioned above, Proposition 3.1 also proves that the minimax distortion rate over the set of distributions with continuous densities, unique optimal codebook, and such that the minimum eigenvalues of the Hessian matrices $H(\mathbf{c}^*)$ are uniformly lower bounded by $1/4k$, is still $\Omega(1/\sqrt{n})$.

This minimax lower bound has to be compared to the upper risk bound obtained in Theorem 3.1 for the empirical risk minimizer $\hat{\mathbf{c}}_n$ over the set of distributions $\mathcal{D}(c_1/\sqrt{n})$. To be more precise, Theorem 3.1 ensures that, provided that n is large enough,

$$\sup_{P \in \mathcal{D}(c_1/\sqrt{n})} \mathbb{E} \ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq \frac{g(k, d, M)}{\sqrt{n}},$$

where $g(k, d, M)$ depends only on k , d and M . In other words, the dependency of the upper bounds stated in Theorem 3.1 on ε turns out to be sharp whenever $\varepsilon \sim n^{-\frac{1}{2}}$. Unfortunately, Proposition 3.1 can not be easily extended to the case where $\varepsilon \sim n^{-\alpha}$, with $0 < \alpha < 1/2$. Consequently an open question is whether the upper bounds stated in Theorem 3.1 remains accurate with respect to ε in this case.

3.3. Quasi-Gaussian mixture example. The aim of this subsection is to illustrate the results offered in Section 3 with 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

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

where \tilde{k} denotes the number of components of the mixture, and the θ_i 's denote the weights of the mixture, which satisfy $\sum_{i=1}^k \theta_i = 1$. 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 3.2 can be derived for distributions with different variance matrices Σ_i , at the cost of more computing.

Since the 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

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

where N_i denotes a normalization constant for each Gaussian variable.

Let ε be defined as $\varepsilon = 1 - \min_{i=1,\dots,k} N_i$. Roughly, the model proposed above will be close to the Gaussian mixture model when ε is small. 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 assume that, for all $i = 1, \dots, k$, $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0, M)$.

It is worth noticing that the assumption $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0, M)$ can easily be satisfied as soon as M is chosen large enough. For such a model, Proposition 3.2 offers a sufficient condition for P to satisfy a margin condition.

PROPOSITION 3.2. *Let $\theta_{min} = \min_{i=1,\dots,k} \theta_i$, and $\theta_{max} = \max_{i=1,\dots,k} \theta_i$. Assume that*

$$(9) \quad \frac{\theta_{min}}{\theta_{max}} \geq \max \left(\frac{2048k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/2048\sigma^2})}, \frac{2048k^2M^3}{(1-\varepsilon)7\sigma^2\tilde{B}(e^{\tilde{B}^2/32\sigma^2}-1)} \right).$$

Then P satisfies a margin condition with radius $\frac{\tilde{B}}{8}$.

It is worth mentioning that P has a continuous density, and that, according to Proposition 2.2, the second derivative matrices of the risk function, at the optimal codebooks, must be positive definite. Thus, P might be in the scope of the result in [24]. However, there is no elementary proof of the fact that $|\bar{\mathcal{M}}| = 1$, whereas \mathcal{M} is finite is guaranteed by Proposition

2.2. This shows that the margin condition given in Definition 2.1 may be easier to check than the condition offered in [24]. The condition (9) can be decomposed as follows. If

$$\frac{\theta_{min}}{\theta_{max}} \geq \frac{2048k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/2048\sigma^2})},$$

then every optimal codebook \mathbf{c}^* must be close to the vector of means of the mixture $\mathbf{m} = (m_1, \dots, m_k)$. Therefore, it is possible to approximately locate N^* , and to derive an upper bound on the weight function $p(t)$ defined in Definition 2.1. This leads to the second term of the maximum in (9).

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 3.2 gives details on how σ has to be small compared to \tilde{B} , in order to satisfy the requirements of Proposition 2.2. This ensures that the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ reaches an improved convergence rate of $1/n$.

It may be noticed that Proposition 3.2 offers almost the same condition than Proposition 4.2 in [17]. In fact, since the Gaussian mixture distributions have a continuous density, making use of (4) in Example 3 ensures that the margin condition for Gaussian mixtures is equivalent to a bound on the density over the critical region N^* .

It is important to note that this 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 exists only one optimal codebook \mathbf{c}^* , up to relabeling, and that we are able to locate this optimal codebook \mathbf{c}^* . As mentioned in Proposition 2.2, the key quantity 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 3.2 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, assume that P is a truncated rotationally symmetric Gaussian distribution, and $k = 2$. Since every rotation of an optimal codebook leads to another optimal codebook, there exists an infinite set of optimal codebooks. Since, in this case, $N^* = \mathbb{R}^2$, condition (3) can not be satisfied.

4. Proofs.

4.1. *Proof of Proposition 2.1.* The lower bound on B follows from a compactness argument for the weak topology on \mathcal{H} , stated in the following lemma. For the sake of completeness, it is recalled that a sequence c_n of elements in \mathcal{H} weakly converges to c , denoted by $c_n \rightharpoonup_{n \rightarrow \infty} c$, if, for every continuous linear real-valued function f , $f(c_n) \rightarrow_{n \rightarrow \infty} f(c)$. Moreover, a function ϕ from \mathcal{H} to \mathbb{R} is weakly lower semi-continuous if, for all $\lambda \in \mathbb{R}$, the level sets $\{c \in \mathcal{H} \mid \phi(c) \leq \lambda\}$ are closed for the weak topology.

LEMMA 4.1. *Let \mathcal{H} be a separable Hilbert space, and suppose that P is M -bounded. Then*

- i) $\mathcal{B}(0, M)^k$ is weakly compact,*
- ii) $\mathbf{c} \mapsto P\gamma(\mathbf{c}, \cdot)$ is weakly lower semi-continuous.*

PROOF OF LEMMA 4.1. A more general statement of Lemma 4.1 can be found in Section 5.2 of [12], for quantization with Bregman divergences. However, since the proof is much more simple in the special case of the squared-norm based quantization on a Hilbert space, it is briefly recalled here.

Since \mathcal{H} is reflexive, according to Banach-Alaoglu-Bourbaki's Theorem (see, e.g., Theorem 3.16 in [7]), combined with Tychonoff's Theorem (see, e.g., Theorem 2.2.8 in [10]), $\mathcal{B}(0, M)^k$ is a compact subset of \mathcal{H}^k for the weak topology. This proves *i*).

Let x be a fixed element of \mathcal{H}^k . Since $\mathbf{c} \mapsto \|x - c_i\|^2$ is weakly lower semi-continuous (see, e.g., Proposition 3.13 in [7]), $\mathbf{c} \mapsto \gamma(\mathbf{c}, x)$ is weakly lower semi-continuous over $\mathcal{B}(0, M)^k$. Let \mathbf{c}_n be a sequence of $\mathcal{B}(0, M)^k$ such that $\mathbf{c}_n \rightharpoonup_{n \rightarrow \infty} \mathbf{c}$, for the weak topology, for some $\mathbf{c} \in \mathcal{B}(0, M)^k$. Then

$$\gamma(\mathbf{c}, x) \leq \liminf_{n \rightarrow \infty} \gamma(\mathbf{c}_n, x).$$

Applying Fatou's Lemma (see, e.g., Lemma 4.3.3 in [10]) yields that

$$R(\mathbf{c}) \leq \liminf_{n \rightarrow \infty} R(\mathbf{c}_n).$$

Hence *ii*) is proved. It is worth noting that this proves the existence of optimal codebooks for bounded distributions. \square

Let \mathbf{c}'_n be a sequence of optimal codebooks such that $\|c'_{1,n} - c'_{2,n}\| \rightarrow B$, as $n \rightarrow \infty$. Then, according to Lemma 4.1, there exists a subsequence \mathbf{c}_n and an optimal codebook \mathbf{c}^* , such that $\mathbf{c}_n \rightharpoonup_{n \rightarrow \infty} \mathbf{c}^*$, for the weak topology. Then it is clear that $(c_{1,n} - c_{2,n}) \rightharpoonup_{n \rightarrow \infty} (c_1^* - c_2^*)$.

Since $u \mapsto \|u\|$ is weakly lower semi-continuous on \mathcal{H} (see, e.g., Proposition 3.13 in [7]), it follows that

$$\|c_1^* - c_2^*\| \leq \liminf_{n \rightarrow \infty} \|c_{1,n} - c_{2,n}\| = B.$$

Noting that \mathbf{c}^* is an optimal codebook, and the support of P has more than k points, Proposition 1 of [15] ensures that $\|c_1^* - c_2^*\| > 0$.

The uniform lower bound on p_{min} follows from the argument that, since the support of P contains more than k points, then $R_k^* < R_{k-1}^*$, where R_j^* denotes the minimum distortion achievable for j -points quantizers (see, e.g., Proposition 1 in [15]). Denote by α the quantity $R_{k-1}^* - R_k^*$, and suppose that $p_{min} < \frac{\alpha}{4M^2}$. Then there exists an optimal codebook of size k , $\mathbf{c}^{*,k} = (c_1^{*,k}, \dots, c_k^{*,k})$, such that $p_{min} = P(V_1(\mathbf{c}^{*,k}))$. Let $\mathbf{c}^{*,k-1}$ denote an optimal codebook of size $(k-1)$, and define the following k -points quantizer

$$\begin{cases} Q(x) = c_1^{*,k} & \text{if } x \in V_1(\mathbf{c}^{*,k}), \\ Q(x) = c_j^{*,k-1} & \text{if } x \in V_j(\mathbf{c}^{*,k-1}) \cap (V_1(\mathbf{c}^{*,k}))^c. \end{cases}$$

Since $P(\partial V_1(\mathbf{c}^{*,k})) = P(\partial V_j(\mathbf{c}^{*,k-1})) = 0$, for $j = 1, \dots, k-1$, Q is defined P almost surely. Then it is easy to see that

$$R(Q) \leq p_{min} 4M^2 + R_{k-1}^* < R_k^*.$$

Hence the contradiction. Therefore we have $p_{min} \geq \frac{\alpha}{4M^2}$.

4.2. *Proof of Proposition 2.2.* According to Lemma 4.1, \mathcal{M} is weakly compact. Since, according to Proposition 3.13 in [7], $\mathbf{c}^* \mapsto \|\mathbf{c} - \mathbf{c}^*\|$ is weakly lower semi continuous, its minimum is attained over \mathcal{M} . This proves *i*).

The proof of *ii*) is based on the following lemma.

LEMMA 4.2. *Let \mathbf{c} and \mathbf{c}^* be in $\mathcal{B}(0, M)^k$, and $x \in V_i(\mathbf{c}^*) \cap V_j(\mathbf{c}) \cap \mathcal{B}(0, M)$, for $i \neq j$. Then*

$$(10) \quad \left| \left\langle x - \frac{c_i + c_j}{2}, c_i - c_j \right\rangle \right| \leq 4\sqrt{2}M \|\mathbf{c} - \mathbf{c}^*\|,$$

$$(11) \quad d(x, \partial V_i(\mathbf{c}^*)) \leq \frac{4\sqrt{2}M}{B} \|\mathbf{c} - \mathbf{c}^*\|.$$

The two statements of Lemma 4.2 emphasize the fact that, provided that \mathbf{c} and \mathbf{c}^* are quite similar, the areas on which the label may differ with respect to \mathbf{c} and \mathbf{c}^* should be close to the boundary of Voronoi diagrams. This idea is mentioned in the proof of Corollary 1 in [2]. Nevertheless we provide here a simpler proof.

PROOF OF LEMMA 4.2. Let x be in $V_i(\mathbf{c}^*) \cap V_j(\mathbf{c}) \cap \mathcal{B}(0, M)$, then $\|x - c_j\|^2 \leq \|x - c_i\|^2$, which leads to $\left\langle c_i - c_j, x - \frac{c_i + c_j}{2} \right\rangle \leq 0$. Since $\|x - c_i^*\| \leq \|x - c_j^*\|$, we may write

$$\|x - c_i\| \leq \|x - c_j\| + \|c_i - c_i^*\| + \|c_j - c_j^*\|.$$

Taking square on both sides leads to

$$\begin{aligned} \|x - c_i\|^2 - \|x - c_j\|^2 &\leq 2\|x - c_j\|(\|c_i - c_i^*\| + \|c_j - c_j^*\|) \\ &\quad + (\|c_i - c_i^*\| + \|c_j - c_j^*\|)^2 \\ &\leq 8M(\|c_i - c_i^*\| + \|c_j - c_j^*\|) \\ &\leq 8\sqrt{2}M\|\mathbf{c} - \mathbf{c}^*\|. \end{aligned}$$

Since $\|x - c_i\|^2 - \|x - c_j\|^2 = -2\left\langle x - \frac{c_i + c_j}{2}, c_i - c_j \right\rangle$, (10) is proved.

To prove (11), remark that, since $x \in V_i(\mathbf{c}^*)$, $d(x, \partial V_i(\mathbf{c}^*)) \leq d(x, h_{i,j}^*)$, where $h_{i,j}^*$ is the hyperplane defined by $\left\{ x \in \mathcal{B}(0, M) \mid \|x - c_i^*\| = \|x - c_j^*\| \right\}$. Using quite simple geometric arguments, we deduce that

$$d(x, h_{i,j}^*) = \left| \left\langle x - \frac{c_i^* + c_j^*}{2}, \frac{c_i^* - c_j^*}{\|c_i^* - c_j^*\|} \right\rangle \right|.$$

The same arguments as in the proof of (10) guarantee that

$$\begin{aligned} \left| \left\langle x - \frac{c_i^* + c_j^*}{2}, \frac{c_i^* - c_j^*}{\|c_i^* - c_j^*\|} \right\rangle \right| &= \left\langle x - \frac{c_i^* + c_j^*}{2}, \frac{c_i^* - c_j^*}{\|c_i^* - c_j^*\|} \right\rangle \\ &\leq \frac{4\sqrt{2}M}{B}\|\mathbf{c} - \mathbf{c}^*\|. \end{aligned}$$

□

Equipped with Lemma 4.2, we are in a position to prove *ii*). Let \mathbf{c} be in $\mathcal{B}(0, M)^k$, and $(W_1(\mathbf{c}), \dots, W_k(\mathbf{c}))$ be a Voronoi partition associated with \mathbf{c} , as defined in Section 2. Let \mathbf{c}^* be in \mathcal{M} , then $\ell(\mathbf{c}, \mathbf{c}^*)$ can be decomposed as follows:

$$\begin{aligned} P\gamma(\mathbf{c}, \cdot) &= \sum_{i=1}^k P(\|x - c_i\|^2 1_{W_i(\mathbf{c})}) \\ &= \sum_{i=1}^k P(\|x - c_i\|^2 1_{V_i(\mathbf{c}^*)}) + \sum_{i=1}^k P(\|x - c_i\|^2 (1_{W_i(\mathbf{c})} - 1_{V_i(\mathbf{c}^*)})). \end{aligned}$$

Since, for all $i = 1, \dots, k$, $P(x1_{V_i(\mathbf{c}^*)}(x)) = P(V_i(\mathbf{c}^*))c_i^*$ (centroid condition), we may write

$$P(\|x - c_i\|^2 1_{V_i(\mathbf{c}^*)}) = P(V_i(\mathbf{c}^*))\|c_i - c_i^*\|^2 + P(\|x - c_i^*\|^2 1_{V_i(\mathbf{c}^*)}),$$

from which we deduce

$$P\gamma(\mathbf{c}, \cdot) = P\gamma(\mathbf{c}^*, \cdot) + \sum_{i=1}^k P(V_i(\mathbf{c}^*))\|c_i - c_i^*\|^2 + \sum_{i=1}^k P(\|x - c_i\|^2 (1_{W_i(\mathbf{c})} - 1_{V_i(\mathbf{c}^*)})),$$

which leads to

$$\ell(\mathbf{c}, \mathbf{c}^*) \geq p_{min}\|\mathbf{c} - \mathbf{c}^*\|^2 + \sum_{i=1}^k \sum_{j \neq i} P\left(\left(\|x - c_j\|^2 - \|x - c_i\|^2\right) 1_{V_i(\mathbf{c}^*) \cap W_j(\mathbf{c})}\right).$$

Since $x \in W_j(\mathbf{c}) \subset V_j(\mathbf{c})$, $\|x - c_j\|^2 - \|x - c_i\|^2 \leq 0$. Thus it remains to bound from above

$$\sum_{i=1}^k \sum_{j \neq i} P\left(\left(\|x - c_i\|^2 - \|x - c_j\|^2\right) 1_{V_i(\mathbf{c}^*) \cap W_j(\mathbf{c})}\right).$$

Noticing that

$$\|x - c_i\|^2 - \|x - c_j\|^2 = 2 \left\langle c_j - c_i, x - \frac{c_i + c_j}{2} \right\rangle,$$

and using Lemma 4.2, we get

$$\sum_{i=1}^k P(\|x - c_i\|^2 (1_{W_i(\mathbf{c})} - 1_{V_i(\mathbf{c}^*)})) \geq -8\sqrt{2}M\|\mathbf{c} - \mathbf{c}^*\|p \left(\frac{4\sqrt{2}M}{B}\|\mathbf{c} - \mathbf{c}^*\| \right).$$

Consequently, if P satisfies (3), then, if $\|\mathbf{c} - \mathbf{c}^*\| \leq \frac{Br_0}{4\sqrt{2}M}$,

$$(12) \quad \ell(\mathbf{c}, \mathbf{c}^*) \geq \frac{p_{min}}{2}\|\mathbf{c} - \mathbf{c}^*\|^2.$$

Now turn to the case where $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\| \geq \frac{Br_0}{4\sqrt{2}M}$. Let $\mathcal{B}^o(\mathbf{c}^*, r)$ denote the open ball with radius r and center \mathbf{c}^* . Since, according to Lemma 4.1, $\mathcal{B}(0, M)^k \cap \left(\bigcup_{\mathbf{c}^* \in \mathcal{M}} \mathcal{B}^o(\mathbf{c}^*, \frac{Br_0}{4\sqrt{2}M})\right)^c$ is weakly compact, and $\mathbf{c} \mapsto P\gamma(\mathbf{c}, \cdot)$ is weakly lower semi-continuous, its minimum over this set is attained. Such

a minimizer is a local minimizer, or is at the boundary $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\| = \frac{Br_0}{4\sqrt{2}M}$. Hence we deduce

$$\begin{aligned} \ell(\mathbf{c}, \mathbf{c}^*) &\geq \varepsilon \wedge \frac{p_{\min} B^2 r_0^2}{64M^2} \\ &\geq \left(\varepsilon \wedge \frac{p_{\min} B^2 r_0^2}{64M^2} \right) \frac{\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2}{4kM^2}. \end{aligned}$$

Note that, since $B \leq 2M$ and $r_0 \leq 2M$, $\left(\varepsilon \wedge \frac{p_{\min} B^2 r_0^2}{64M^2} \right) / 4kM^2 \leq p_{\min}/2$. This proves *ii*).

In the case where $\mathcal{H} = \mathbb{R}^d$, the weak topology coincides with the usual topology. Consequently $\mathcal{B}(0, M)^k$ is compact. Suppose that \mathcal{M} is not finite. Then there exists a sequence \mathbf{c}_n of optimal codebooks, and an optimal codebook \mathbf{c}^* , such that $\|\mathbf{c}_n - \mathbf{c}^*\| \rightarrow_{n \rightarrow \infty} 0$. For n large enough, we have

$$\|\mathbf{c}_n - \mathbf{c}^*\| \leq \frac{Br_0}{4\sqrt{2}M},$$

and $\ell(\mathbf{c}_n, \mathbf{c}^*) = 0$. This contradicts (12).

4.3. Proof of Theorem 3.1. Throughout this subsection P is assumed to satisfy a margin condition with radius r_0 , and to be ε -separated. A non decreasing map $\Phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is called subroot if $x \mapsto \frac{\Phi(x)}{\sqrt{x}}$ is non increasing.

The following localization theorem, derived from Theorem 6.1 in [6], is the main argument of our proof.

THEOREM 4.1. *Let \mathcal{F} be a class of bounded measurable functions such that there exist $b > 0$ and $\omega : \mathcal{F} \rightarrow \mathbb{R}^+$ satisfying*

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

Let K be a positive constant, Φ a sub-root function. Then if r^ is the unique solution of the equation $\Phi(r) = r/24K$, the following holds. Assume that*

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

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

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

A proof of Theorem 4.1 is given in Section 5.3 of [17].

4.3.1. *Proof of (7).* We begin with the finite dimensional case. The proof of (7) follows from the combination of Proposition 2.2 and a direct application of Theorem 4.1. To be more precise, let \mathcal{F} denote the set

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

Since, for all $i \in \{1, \dots, k\}$,

$$\| \|x - c_i\|^2 - \|x - c_i^*(\mathbf{c})\|^2 \| \leq 4M \|c_i - c_i^*(\mathbf{c})\|,$$

it follows that, for every $f \in \mathcal{F}$,

$$\begin{cases} \|f\|_\infty & \leq 8M^2, \\ \text{Var}_P(f) & \leq 16M^2 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2. \end{cases}$$

Define $\omega(f) = 16M^2 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$. It remains to bound from above the complexity term. This is done in the following proposition, derived from the proof of Theorem 1 in [8].

PROPOSITION 4.1. *One has*

(13)

$$\mathbb{E} \sup_{f \in \mathcal{F}, \omega(f) \leq \delta} |(P - P_n)f| \leq \frac{(2\sqrt{2} + 64)\sqrt{kd}}{\sqrt{n}} \left(\sqrt{\log(4|\mathcal{M}|\sqrt{kd})} + 1 \right) \sqrt{\delta}.$$

The proof of Proposition 4.1 derives from classical chaining arguments, and can be found in Section 5.1 of [18]. Let Φ be defined as the right-hand side of (13). Observing that $\Phi(\delta)$ takes the form $\Phi(\delta) = \Xi \sqrt{\delta/n}$, the solution δ^* of the equation $\Phi(\delta) = \delta/24K$ may be written, for any $K > 0$,

$$\delta^* = \frac{K^2 \Xi^2}{n}.$$

Applying Theorem 4.1 to \mathcal{F} leads to, with probability larger than $1 - e^{-x}$,

$$\begin{aligned} (P - P_n)(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*(\mathbf{c}), \cdot)) & \leq K^{-1} 16M^2 \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2 \\ & \quad + \frac{K \Xi^2}{n} + \frac{9K + 128M^2}{4n} x. \end{aligned}$$

Introducing the inequality $\kappa_0 \ell(\mathbf{c}, \mathbf{c}^*) \geq \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$ provided by Proposition 2.2, choosing $K = 32M^2 \kappa_0$ leads to (7).

4.3.2. *Proof of (6).* Similarly to the proof of (7), the proof of (6) is based on an application of Theorem 4.1 to the set \mathcal{F} , defined in the above subsection. However, the technique used to bound the complexity term is slightly different, and leads to the following result.

PROPOSITION 4.2. *There exists a universal constant C such that,*

$$(14) \quad \mathbb{E} \sup_{f \in \mathcal{F}, \omega(f) \leq \delta} |(P - P_n)f| \leq C \frac{|\bar{\mathcal{M}}| 2\sqrt{2k}}{\sqrt{n}} \sqrt{\delta}.$$

This proof relies on the generic chaining principle introduced by Fernique in [11] and developed by Talagrand in [25]. We postpone it to the following subsection. Let Φ' be defined as the right-hand side of (14), and let δ' denote the solution of the equation $\Phi'(\delta) = \delta/24K$, for some positive $K > 0$. Then δ' can be expressed as

$$\delta' = C^2 \frac{8 |\bar{\mathcal{M}}|^2 24^2 K^2 k}{n}.$$

As in the proof of (7), choosing $K = 32\kappa_0 M^2$, applying Theorem 4.1 and combining it with Proposition 2.2 leads to the result.

4.3.3. *Proof of Proposition 4.2.* As mentioned above, this proof relies on the generic chaining principle. As will be shown below, avoiding Dudley's entropy argument by introducing some Gaussian random vectors allows us to take advantage of the underlying Hilbert space structure. The first step is to decompose the complexity term according to optimal codebooks, in the following way

$$\begin{aligned} & \mathbb{E} \sup_{\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2 \leq \delta/16M^2} |(P - P_n)(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*(\mathbf{c}), \cdot))| \\ & \leq \sum_{\mathbf{c}^* \in \bar{\mathcal{M}}} \mathbb{E} \sup_{\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2} |(P - P_n)(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*, \cdot))|. \end{aligned}$$

Next we bound from above every term of the right-hand side. Let \mathbf{c}^* be fixed, and let $\sigma_1, \dots, \sigma_n$ denote some independent Rademacher variables. According to the symmetrization principle (see, e.g., Section 2.2 of [16]),

$$\begin{aligned} & \mathbb{E} \sup_{\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2} |(P - P_n)(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*, \cdot))| \\ & \leq 2\mathbb{E}_{X, \sigma} \sup_{\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2} \frac{1}{n} \sum_{i=1}^n \sigma_i (\gamma(\mathbf{c}, X_i) - \gamma(\mathbf{c}^*, X_i)), \end{aligned}$$

where \mathbb{E}_Y denote integration with respect to the distribution of Y . The main argument of this proof is the following Theorem 2.1.5 of [25].

THEOREM 4.2. *Let $Y_v, v \in \mathcal{V}$ denote a centered stochastic process indexed by \mathcal{V} , and X_v denote a centered Gaussian process indexed by the same set \mathcal{V} . Let d be a pseudo-distance over \mathcal{V} such that*

- i) $\forall v, v' \in \mathcal{V}$ $Y_v - Y_{v'}$ is subgaussian with variance $d^2(v, v')$,*
- ii) $\forall v, v' \in \mathcal{V}$ $\text{Var}(X_v - X_{v'}) = d^2(v, v')$.*

Then there exists a universal constant C such that

$$\mathbb{E} \sup_{v \in \mathcal{V}} (Y_v - Y_{v_0}) \leq C \mathbb{E} \sup_{v \in \mathcal{V}} (X_v - X_{v_0}),$$

where v_0 is a fixed element of \mathcal{V} .

For the sake of completeness, it is recalled that a variable Z is subgaussian with variance d^2 if it satisfies the following tail inequality

$$\mathbb{P}(Z \geq t) \leq \exp\left(-\frac{t^2}{2d^2}\right),$$

for all $t > 0$. For a complete introduction to subgaussianity and its applications in empirical processes theory, the interested reader is referred to [22]. It is worth pointing out that Dudley's entropy bound can be retrieved from Theorem 4.2, up to an absolute constant (see, e.g., Proposition 1.2.1 in [25]).

In order to apply Theorem 4.2, we define the process

$$Z_{\mathbf{c}} = \sum_{i=1}^n \sigma_i \gamma(\mathbf{c}, X_i),$$

over the set $\mathcal{V}(\delta) = \mathcal{B}(\mathbf{c}^*, \frac{\sqrt{\delta}}{4M})$, where \mathbf{c}^* is a fixed optimal codebook. For $i = 1, \dots, n$, $\mathbf{c}, \mathbf{c}' \in \mathcal{V}(\delta)$, we have

$$\begin{aligned} (\gamma(\mathbf{c}, X_i) - \gamma(\mathbf{c}', X_i))^2 &\leq \sup_{j=1, \dots, k} (\|X_i - c_j\|^2 - \|X_i - c'_j\|^2)^2 \\ &\leq \sup_{j=1, \dots, k} (-2 \langle c_j - c'_j, X_i \rangle + \|c_j\|^2 - \|c'_j\|^2)^2 \\ &\leq \sup_{j=1, \dots, k} \left(8 \langle c_j - c'_j, X_i \rangle^2 + 2(\|c_j\|^2 - \|c'_j\|^2)^2 \right) \\ &\leq 8 \sum_{j=1}^k \langle c_j - c'_j, X_i \rangle^2 + 8M^2 \sum_{j=1}^k (\|c_j\| - \|c'_j\|)^2, \end{aligned}$$

where the last inequality follows from $(\|c_j\|^2 - \|c'_j\|^2) \leq 2M(\|c_j\| - \|c'_j\|)$. Let the pseudo-distance d be defined by

$$d^2(\mathbf{c}, \mathbf{c}') = 8 \sum_{i=1}^n \sum_{j=1}^k \langle c_j - c'_j, X_i \rangle^2 + 8M^2 \sum_{i=1}^n \sum_{j=1}^k (\|c_j\| - \|c'_j\|)^2.$$

It can be easily checked that d satisfies the triangle inequality, using Cauchy-Schwarz inequality. Let X_1, \dots, X_n be fixed, then Hoeffding's inequality (see, e.g., Proposition 2.7 in [22]) ensures that, for every \mathbf{c} and $\mathbf{c}' \in \mathcal{V}(\delta)$, $Z_{\mathbf{c}} - Z_{\mathbf{c}'}$ is subgaussian with variance $d^2(\mathbf{c}, \mathbf{c}')$. Define now the Gaussian process $X_{\mathbf{c}}$ by

$$X_{\mathbf{c}} = 2\sqrt{2} \sum_{i=1}^n \sum_{j=1}^k \langle c_j, X_i \rangle \xi_{i,j} + 2\sqrt{2}M \sum_{i=1}^n \sum_{j=1}^k \|c_j\| \xi'_{i,j},$$

where the ξ 's and ξ' 's are independent standard Gaussian variables. It is straightforward that $\text{Var}(X_{\mathbf{c}} - X_{\mathbf{c}'}) = d^2(\mathbf{c}, \mathbf{c}')$. Therefore, applying Theorem 4.2 leads to

$$\begin{aligned} \mathbb{E}_{\sigma} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} (Z_{\mathbf{c}} - Z_{\mathbf{c}^*}) &\leq C \mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} (X_{\mathbf{c}} - X_{\mathbf{c}^*}) \\ (15) \quad &\leq C2\sqrt{2} \mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{i=1}^n \sum_{j=1}^k \langle c_j - c_j^*, X_i \rangle \xi_{i,j} \\ &\quad + C2\sqrt{2}M \mathbb{E}_{\xi'} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{i=1}^n \sum_{j=1}^k (\|c_j\| - \|c_j^*\|) \xi'_{i,j}. \end{aligned}$$

Using almost the same technique as in the proof of Theorem 2.1 in [5], the first term of the right-hand side of (15) can be bounded as follows:

$$\begin{aligned} \mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{i=1}^n \sum_{j=1}^k \langle c_j - c_j^*, X_i \rangle \xi_{i,j} &= \mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{j=1}^k \left\langle c_j - c_j^*, \left(\sum_{i=1}^n \xi_{i,j} X_i \right) \right\rangle \\ &\leq \mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \|\mathbf{c} - \mathbf{c}^*\| \sqrt{\sum_{j=1}^k \left\| \sum_{i=1}^n \xi_{i,j} X_i \right\|^2} \\ &\leq \frac{\sqrt{\delta}}{4M} \sqrt{\sum_{j=1}^k \mathbb{E}_{\xi} \left\| \sum_{i=1}^n \xi_{i,j} X_i \right\|^2} \\ &\leq \frac{\sqrt{k\delta}}{4M} \sqrt{\sum_{i=1}^n \|X_i\|^2}. \end{aligned}$$

Then, applying Jensen's inequality ensures that

$$\mathbb{E}_X \sqrt{\sum_{i=1}^n \|X_i\|^2} \leq \sqrt{n}M.$$

Similarly, the second term of the right-hand side of (15) is bounded from above by

$$\begin{aligned} \mathbb{E}_{\xi'} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{i=1}^n \sum_{j=1}^k (\|c_j\| - \|c_j^*\|) \xi'_{i,j} &\leq \mathbb{E}_{\xi'} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{j=1}^k (\|c_j\| - \|c_j^*\|) \left(\sum_{i=1}^n \xi'_{i,j} \right) \\ &\leq \mathbb{E}_{\xi'} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sqrt{\sum_{j=1}^k (\|c_j\| - \|c_j^*\|)^2} \sqrt{\sum_{j=1}^k \left(\sum_{i=1}^n \xi'_{i,j} \right)^2} \\ &\leq \frac{\sqrt{k\delta}}{4M} \sqrt{n}. \end{aligned}$$

Combining these two bounds yields that, for a fixed \mathbf{c}^* ,

$$\mathbb{E}_{X,\sigma} \sup_{\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2} (Z_{\mathbf{c}} - Z_{\mathbf{c}^*}) \leq C\sqrt{2k\delta n}M,$$

which leads to the desired result.

4.4. Proof of Proposition 3.1. Throughout this subsection, $\mathcal{H} = \mathbb{R}^d$, and, for a codebook \mathbf{c} , let Q denote the associated nearest-neighbor quantizer. In the general case, such an association depends on how the boundaries are allocated. However, since the distributions involved in the minimax result have densities, how boundaries are allocated will not matter.

Let $k \geq 3$ be an integer. For convenience k is assumed to be divisible by 3. Let $m = 2k/3$. Let z_1, \dots, z_m denote a 6Δ -net in $\mathcal{B}(0, M - \rho)$, where $\Delta > 0$, and w_1, \dots, w_m a sequence of vectors such that $\|w_i\| = \Delta$. Finally, denote by U_i the ball $\mathcal{B}(z_i, \rho)$ and by U'_i the ball $\mathcal{B}(z_i + w_i, \rho)$. Slightly anticipating, define $\rho = \frac{\Delta}{16}$.

To get the largest Δ such that for all $i = 1, \dots, k$, U_i and U'_i are included in $\mathcal{B}(0, M)$, it suffices to get the largest Δ such that there exists a 6Δ -net in $\mathcal{B}(0, M - \Delta/16)$. Since the cardinal of a 6Δ -net is larger than the largest number of balls of radius 6Δ which can be packed into $\mathcal{B}(0, M - \Delta/16)$, a sufficient condition on Δ to guarantee that a 6Δ -net can be found is given by

$$m \leq \left(\frac{M - \Delta/16}{6\Delta} \right)^d.$$

Since $\Delta \leq M$, Δ can be chosen as

$$\Delta = \frac{5M}{32m^{1/d}}.$$

For such a Δ , ρ takes the value $\rho = \frac{\Delta}{16} = \frac{5M}{512m^{1/d}}$. Therefore, it only depends on k , d , and M .

Let $z = (z_i)_{i=1,\dots,m}$ and $w = (w_i)_{i=1,\dots,m}$ be sequences as described above, such that, for $i = 1, \dots, k$, U_i and U'_i are included in $\mathcal{B}(0, M)$. For a fixed $\sigma \in \{-1, +1\}^m$ such that $\sum_{i=1}^m \sigma_i = 0$, let P_σ be defined as

$$\begin{cases} P_\sigma(U_i) &= \frac{1+\sigma_i\delta}{2m}, \\ P_\sigma(U'_i) &= \frac{1+\sigma_i\delta}{2m}, \\ P_\sigma &\underset{U_i}{\sim} (\rho - \|x - z_i\|)1_{\|x-z_i\|\leq\rho}d\lambda(x), \\ P_\sigma &\underset{U'_i}{\sim} (\rho - \|x - z_i - w_i\|)1_{\|x-z_i-w_i\|\leq\rho}d\lambda(x), \end{cases}$$

where λ denotes the Lebesgue measure. These cone-shaped distributions has been designed to have a continuous density, as in Theorem 4 in [1]. To be more precise, for τ in $\{-1, +1\}^{\frac{m}{2}}$, $\sigma(\tau)$ is defined as the sequence in $\{-1, +1\}^m$ such that

$$\begin{cases} \sigma_i(\tau) &= \tau_i, \\ \sigma_{i+\frac{m}{2}}(\tau) &= -\sigma_i(\tau), \end{cases}$$

for $i = 1, \dots, \frac{m}{2}$. Finally, for a quantizer Q let $R(Q, P_\sigma)$ denote the distortion of Q in the case where the source distribution is P_σ .

Similarly, for σ in $\{-1, +1\}^m$ satisfying $\sum_{i=1}^m \sigma_i = 0$, let Q_σ denote the quantizer defined by $Q_\sigma(U_i) = Q_\sigma(U'_i) = z_i + \omega_i/2$ if $\sigma_i = -1$, $Q_\sigma(U_i) = z_i$ and $Q_\sigma(U'_i) = z_i + \omega_i$ if $\sigma_i = +1$. Let \mathcal{Q} denote the set of such quantizers. It can be proved that only quantizers in \mathcal{Q} have to be considered.

PROPOSITION 4.3. *Assume that $\delta \leq 1/3$, $\Delta > 0$, and $\rho \leq \frac{\Delta}{16}$. Then, for every quantizer Q there exists a quantizer Q_σ in \mathcal{Q} such that*

$$\forall P_{\sigma'} \quad R(Q_\sigma, P_{\sigma'}) \leq R(Q, P_{\sigma'}).$$

The proof of Proposition 4.3 follows the proof of Step 3 of Theorem 1 in [4], replacing distributions supported on a finite set with distributions supported on small balls. Provided that the radius of these balls are small enough, the results are nearly the same in the two cases. The proof of Proposition 4.3 can be found in Section 5.4 of [18].

Since, for $\sigma \neq \sigma'$, $R(Q'_{\sigma}, P_{\sigma}) > R(Q_{\sigma}, P_{\sigma})$, Proposition 4.3 ensures that the P_{σ} 's have a unique optimal codebook, up to relabeling.

For any σ and σ' in $\{-1, +1\}^m$, denote by $\rho(\sigma, \sigma') = \sum_{i=1}^m |\sigma_i - \sigma'_i|$, and by $H(P_{\sigma}, P_{\sigma'})$ the Hellinger distance between P_{σ} and $P_{\sigma'}$. To apply Assouad's Lemma to the set $\{P_{\sigma(\tau)}\}_{\tau \in \{-1, +1\}^{\frac{m}{2}}}$, the following lemma is needed:

LEMMA 4.3. *Let τ and τ' denote two sequences in $\{-1, +1\}^{\frac{m}{2}}$ such that $\rho(\tau, \tau') = 2$, then*

$$H(P_{\sigma(\tau)}^{\otimes n}, P_{\sigma(\tau')}^{\otimes n}) \leq \frac{4n\delta^2}{m},$$

where $P^{\otimes n}$ denotes the product law of a n -sample drawn from P .

Furthermore, for any σ and σ' in $\{-1, +1\}^m$,

$$R(Q_{\sigma'}, P_{\sigma}) = R(Q_{\sigma}, P_{\sigma}) + \frac{\Delta^2 \delta}{8m} \rho(\sigma, \sigma').$$

A proof of Lemma 4.3 is given in Section 5.5 of [18]. Equipped with Lemma 4.3, a direct application of Assouad's Lemma as in Theorem 2.12 of [26] yields that, provided that $\delta = \frac{\sqrt{m}}{2\sqrt{n}}$,

$$\sup_{\tau \in \{-1, +1\}^{\frac{m}{2}}} \mathbb{E} \left(R(\hat{Q}_n, P_{\sigma(\tau)}) - R(Q_{\sigma(\tau)}, P_{\sigma(\tau)}) \right) \geq c_0 M^2 \sqrt{\frac{k^{1-\frac{4}{d}}}{n}},$$

for any empirically designed quantizer \hat{Q}_n , where c_0 is an explicit constant.

Finally, it may be noticed that, for every $\delta \leq \frac{1}{3}$ and σ , P_{σ} satisfies a margin condition as in (8), and is ε -separated, with

$$\varepsilon = \frac{\Delta^2 \delta}{2m}.$$

This concludes the proof of Proposition 3.1.

4.5. *Proof of Proposition 3.2.* As mentioned below Proposition 3.2, the inequality

$$\frac{\theta_{\min}}{\theta_{\max}} \geq \frac{2048k\sigma^2}{(1-\varepsilon)\tilde{B}^2(1-e^{-\tilde{B}^2/2048\sigma^2})},$$

ensures that, for every j in $\{1, \dots, k\}$, there exists i in $\{1, \dots, k\}$ such that $\|c_i^* - m_j\| \leq \tilde{B}/16$. To be more precise, let \mathbf{m} denote the vector of means

(m_1, \dots, m_k) , then

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

Assume that there exists i in $\{1, \dots, k\}$ such that, for all j , $\|c_j^* - m_i\| \geq \tilde{B}/16$. Then

$$\begin{aligned} R(\mathbf{c}) &\geq \frac{\theta_i}{2\pi\sigma^2} \int_{\mathcal{B}(m_i, \tilde{B}/32)} \frac{\tilde{B}^2}{1024} e^{-\frac{\|x - m_i\|^2}{2\sigma^2}} \\ &\geq \frac{\tilde{B}^2 \theta_{min}}{2048\pi\sigma^2} \int_{\mathcal{B}(m_i, \tilde{B}/32)} e^{-\frac{\|x - m_i\|^2}{2\sigma^2}} \\ &> \frac{\tilde{B}^2 \theta_{min}}{1024} \left(1 - e^{-\frac{\tilde{B}^2}{2048\sigma^2}}\right) \\ &> R(\mathbf{m}). \end{aligned}$$

Hence the contradiction. Up to relabeling, it is now assumed that for $i = 1, \dots, k$, $\|m_i - c_i^*\| \leq \tilde{B}/16$. Take y in $N^*(x)$, for $x \leq \frac{\tilde{B}}{8}$, then, for every i in $\{1, \dots, k\}$,

$$\|y - m_i\| \geq \frac{\tilde{B}}{4},$$

which leads to

$$\sum_{i=1}^k \frac{\theta_i}{2\pi\sigma^2 N_i} \|y - m_i\|^2 e^{-\frac{\|y - m_i\|^2}{2\sigma^2}} \leq \frac{k\theta_{max}}{(1-\varepsilon)2\pi\sigma^2} e^{-\frac{\tilde{B}^2}{32\sigma^2}}.$$

Since the Lebesgue measure of $N^*(x)$ is smaller than $4k\pi Mx$, it follows that

$$P(N^*(x)) \leq \frac{2k^2 M \theta_{max}}{(1-\varepsilon)\sigma^2} e^{-\frac{\tilde{B}^2}{32\sigma^2}} x.$$

On the other hand, $\|m_i - c_i^*\| \leq \tilde{B}/16$ yields that

$$\mathcal{B}(m_i, 3\tilde{B}/8) \subset V_i(\mathbf{c}^*).$$

Therefore,

$$\begin{aligned} P(V_i(\mathbf{c}^*)) &\geq \frac{\theta_i}{2\pi\sigma^2 N_i} \int_{\mathcal{B}(m_i, 3\tilde{B}/8)} e^{-\frac{\|x-m_i\|^2}{2\sigma^2}} dx \\ &\geq \theta_i \left(1 - e^{-\frac{9\tilde{B}^2}{128\sigma^2}}\right), \end{aligned}$$

hence $p_{min} \geq \theta_{min} \left(1 - e^{-\frac{9\tilde{B}^2}{128\sigma^2}}\right)$. Consequently, provided that

$$\frac{\theta_{min}}{\theta_{max}} \geq \frac{2048k^2 M^3}{(1-\varepsilon)7\sigma^2 \tilde{B}(e^{\tilde{B}^2/32\sigma^2} - 1)},$$

direct calculation shows that

$$P(N^*(x)) \leq \frac{Bp_{min}}{128M^2}x.$$

This ensures that P satisfies (3). According to Proposition 2.2, since $\mathcal{H} = \mathbb{R}^2$, \mathcal{M} is finite.

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

REFERENCES

- [1] ANTOS, A. (2005). Improved minimax bounds on the test and training distortion of empirically designed vector quantizers. *IEEE Trans. Inform. Theory* **51** 4022–4032. [MR2239018 \(2007b:94149\)](#)
- [2] ANTOS, A., GYÖRFI, L. and GYÖRGY, A. (2005). Individual convergence rates in empirical vector quantizer design. *IEEE Trans. Inform. Theory* **51** 4013–4022. [MR2239017 \(2007a:94125\)](#)
- [3] AUDER, B. and FISCHER, A. (2012). Projection-based curve clustering. *J. Stat. Comput. Simul.* **82** 1145–1168. [MR2958779](#)
- [4] BARTLETT, P. L., LINDER, T. and LUGOSI, G. (1998). The minimax distortion redundancy in empirical quantizer design. *IEEE Trans. Inform. Theory* **44** 1802–1813. [MR1664098 \(2001f:94006\)](#)
- [5] BIAU, G., DEVROYE, L. and LUGOSI, G. (2008). On the performance of clustering in Hilbert spaces. *IEEE Trans. Inform. Theory* **54** 781–790. [MR2444554 \(2009m:68221\)](#)
- [6] BLANCHARD, G., BOUSQUET, O. and MASSART, P. (2008). Statistical performance of support vector machines. *Ann. Statist.* **36** 489–531. [MR2396805 \(2009m:62085\)](#)
- [7] BREZIS, H. (2011). *Functional analysis, Sobolev spaces and partial differential equations. Universitext.* Springer, New York. [MR2759829 \(2012a:35002\)](#)
- [8] CHICHIGNOUD, M. and LOUSTAU, S. (2013-06). Adaptive Noisy Clustering.
- [9] CHOU, P. A. (1994). 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* 457.

- [10] DUDLEY, R. M. (2002). *Real analysis and probability*. *Cambridge Studies in Advanced Mathematics* **74**. Cambridge University Press, Cambridge Revised reprint of the 1989 original. [MR1932358 \(2003h:60001\)](#)
- [11] FERNIQUE, X. (1975). Régularité des trajectoires des fonctions aléatoires gaussiennes. In *École d'Été de Probabilités de Saint-Flour, IV-1974* 1–96. Lecture Notes in Math., Vol. 480. Springer, Berlin. [MR0413238 \(54 ##1355\)](#)
- [12] FISCHER, A. (2010). Quantization and clustering with Bregman divergences. *J. Multivariate Anal.* **101** 2207–2221. [MR2671211 \(2012c:62188\)](#)
- [13] GERSHO, A. and GRAY, R. M. (1991). *Vector quantization and signal compression*. Kluwer Academic Publishers, Norwell, MA, USA.
- [14] GRAF, S. and LUSCHGY, H. (2000). *Foundations of quantization for probability distributions*. *Lecture Notes in Mathematics* **1730**. Springer-Verlag, Berlin. [MR1764176 \(2001m:60043\)](#)
- [15] GRAF, S., LUSCHGY, H. and PAGÈS, G. (2007). Optimal quantizers for Radon random vectors in a Banach space. *J. Approx. Theory* **144** 27–53. [MR2287375 \(2007k:41084\)](#)
- [16] KOLTCHINSKII, V. (2006). Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Statist.* **34** 2593–2656. [MR2329442 \(2009h:62060\)](#)
- [17] LEVRARD, C. (2013). Fast rates for empirical vector quantization. *Electron. J. Stat.* **7** 1716–1746.
- [18] LEVRARD, C. (2013). Margin conditions for vector quantization. 43 pages.
- [19] LINDER, T. (2002). Learning-theoretic methods in vector quantization. In *Principles of nonparametric learning (Udine, 2001)*. *CISM Courses and Lectures* **434** 163–210. Springer, Vienna. [MR1987659 \(2004f:68128\)](#)
- [20] LINDER, T., LUGOSI, G. and ZEGER, K. (1994). Rates of convergence in the source coding theorem, in empirical quantizer design, and in universal lossy source coding. *IEEE Trans. Inform. Theory* **40** 1728–1740. [MR1322387 \(96b:94005\)](#)
- [21] MAMMEN, E. and TSYBAKOV, A. B. (1999). Smooth discrimination analysis. *Ann. Statist.* **27** 1808–1829. [MR1765618 \(2001i:62074\)](#)
- [22] MASSART, P. (2007). *Concentration inequalities and model selection*. *Lecture Notes in Mathematics* **1896**. Springer, Berlin. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard. [MR2319879 \(2010a:62008\)](#)
- [23] MASSART, P. and NÉDÉLEC, É. (2006). Risk bounds for statistical learning. *Ann. Statist.* **34** 2326–2366. [MR2291502 \(2009e:62282\)](#)
- [24] POLLARD, D. (1982). A central limit theorem for k -means clustering. *Ann. Probab.* **10** 919–926. [MR672292 \(84c:60047\)](#)
- [25] TALAGRAND, M. (2005). *The generic chaining*. *Springer Monographs in Mathematics*. Springer-Verlag, Berlin Upper and lower bounds of stochastic processes. [MR2133757 \(2006b:60006\)](#)
- [26] TSYBAKOV, A. B. (2009). *Introduction to nonparametric estimation*. *Springer Series in Statistics*. Springer, New York. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats. [MR2724359 \(2011g:62006\)](#)