

NONASYMPTOTIC BOUNDS FOR VECTOR QUANTIZATION IN HILBERT SPACES

BY CLÉMENT LEVRARD

Université Paris Sud, UPMC and INRIA

Recent results in quantization theory show that the mean-squared expected distortion can reach a rate of convergence of $\mathcal{O}(1/n)$, where n is the sample size [see, e.g., *IEEE Trans. Inform. Theory* **60** (2014) 7279–7292 or *Electron. J. Stat.* **7** (2013) 1716–1746]. This rate is attained for the empirical risk minimizer strategy, if the source distribution satisfies some regularity conditions. However, the dependency of the average distortion on other parameters is not known, and these results are only 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., *Ann. Statist.* **27** (1999) 1808–1829], under which a nonasymptotic upper bound on the expected distortion rate of the empirically optimal quantizer is derived. The dependency of the distortion on other parameters of distributions 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} . For a positive integer k , 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 pre-images 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 [13, 14] and [20] among others.

The quantization theory was originally developed as a way to answer signal compression issues in the late 1940s (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 data need to be preprocessed through a quantization algorithm in order to be exploited.

Received November 2013; revised November 2014.

MSC2010 subject classifications. 62H30.

Key words and phrases. Quantization, localization, fast rates, margin conditions.

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 [20] in the finite-dimensional case.

If $\mathcal{H} = \mathbb{R}^d$, for some $d > 0$, it has been proved in [21] 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 [6]. 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, 25] or [18], 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 has a support bounded by M and

satisfies a technical inequality, namely for some fixed $a > 0$, for every codebook \mathbf{c} , there is a \mathbf{c}^* optimal codebook such that

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

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 , and also on P , but only through M and the technical parameter a . However, in the continuous density and unique minimum case, it has been proved in [11], following the approach of [25], that provided 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 [18], the technique used in [25] can be slightly modified to derive a nonasymptotic 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 [18], these different sets of assumptions turn out to be equivalent in the continuous density case to a technical condition, similar to that used in [24] 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 [22]. This paper provides a condition, which can clearly be thought of as a margin condition in the quantization framework, under which condition (1) is satisfied. The technical constant a has then an explicit expression in terms of natural parameters of P from the quantization point of view. 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 demand either that there exists a unique optimal codebook, as required in [25], hence seems easier to check.

Moreover, a nonasymptotic 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 parameters of P . This bound is also valid if \mathcal{H} is infinite dimensional. This point may be of interest for curve quantization, as done in [3].

In addition, a minimax lower bound is given which allows one 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 codebooks, such that the minimum eigenvalues of the second derivative matrices of the distortion, at the optimal codebooks, are uniformly lower bounded. This result generalizes the previous minimax bound obtained in Theorem 4 of [1] for $k \geq 3$ and $d > 1$.

This paper is organized as follows. In Section 2, some notation and definitions 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: first 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 condition. Finally, the main results are proved in Section 4 and the proofs of several supporting lemmas are deferred to the supplementary material [19].

2. Notation and definitions. Throughout this paper, for $M > 0$ and a in \mathcal{H} , $\mathcal{B}(a, M)$ and $\mathcal{B}^o(a, M)$ will denote, respectively, the closed and open ball with center a and radius M . For a subset A of \mathcal{H} , $\bigcup_{a \in A} \mathcal{B}(a, M)$ will be denoted by $\mathcal{B}(A, M)$. 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 quantization 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 $\mathcal{H} = \mathbb{R}^d$. In this 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 any 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 will be denoted by \mathcal{M} . Since every permutation of the 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

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

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

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

Note 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 nonempty. 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

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

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

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

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

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

where $\mathbb{1}_A$ 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

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

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,

$$\eta_i = \frac{P(x \mathbb{1}_{W_i(\mathbf{c})}(x))}{P(W_i(\mathbf{c}))}.$$

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,

$$\mathbf{c}_i^* = \frac{P(x \mathbb{1}_{W_i(\mathbf{c}^*)}(x))}{P(W_i(\mathbf{c}^*))}.$$

As a remark, the centroid condition ensures that $\mathcal{M} \subset \mathcal{B}(0, M)^k$, and, for every \mathbf{c}^* in \mathcal{M} , $i \neq j$,

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

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 $(W_1(\mathbf{c}^*), \dots, W_k(\mathbf{c}^*))$,

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

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 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 positive.*

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 (for a comprehensive explanation of this statistical learning point of view, see, e.g., [24]). To draw this comparison, the following set is introduced, for any $\mathbf{c}^* \in \mathcal{M}$,

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

The region is of importance when considering the conditions under which the empirical risk minimization strategy for quantization achieves faster rates of convergence, as exposed in [18]. However, to completely translate the margin conditions given in [22] to the quantization framework, the neighborhood of this region has to be introduced. For this purpose, the t -neighborhood of the region $N_{\mathbf{c}^*}$ is defined by $\mathcal{B}(N_{\mathbf{c}^*}, t)$. The quantity of interest is the maximal weight of these t -neighborhoods over the set of optimal codebooks, defined by

$$p(t) = \sup_{\mathbf{c}^* \in \mathcal{M}} P(\mathcal{B}(N_{\mathbf{c}^*}, t)).$$

It is straightforward that $p(0) = 0$. Intuitively, if $p(t)$ 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). A distribution P satisfies a margin condition with radius $r_0 > 0$ if and only if:

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

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

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 [25].

Moreover, the margin condition introduced here only requires a local control of the weight function $p(t)$. The parameter r_0 may be thought of as a gap size around every $N_{\mathbf{c}^*}$, as illustrated by the following example:

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

Note also that the condition mentioned in [22] 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

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

In the case where P has a density and $\mathcal{H} = \mathbb{R}^d$, the condition (3) may be considered as a generalization of the condition stated in Theorem 3.2 of [18], which requires the density of the distribution to be small enough over every $N_{\mathbf{c}^*}$. In fact, provided that P has a continuous density, a uniform bound on the density over every $N_{\mathbf{c}^*}$ 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, for every \mathbf{c}^* , $F_{\mathbf{c}^*}(t) = P(\mathcal{B}(N_{\mathbf{c}^*}, t))$ is differentiable at 0, with derivative

$$F'_{\mathbf{c}^*}(0) = \int_{N_{\mathbf{c}^*}} f(u) d\lambda_{d-1}(u),$$

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

$$(4) \quad \int_{N_{\mathbf{c}^*}} f(u) d\lambda_{d-1}(u) < \frac{Bp_{\min}}{128M^2},$$

for every \mathbf{c}^* , 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 $\bigcup_{\mathbf{c}^*} N_{\mathbf{c}^*}$ 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 [18], where it was required that, for every \mathbf{c}^* ,

$$\|f|_{N_{\mathbf{c}^*}}\|_\infty \leq \frac{\Gamma(d/2)B}{2^{d+5}M^{d+1}\pi^{d/2}}P_{\min},$$

where Γ denotes the Gamma function, and $f|_{N_{\mathbf{c}^*}}$ denotes the restriction of f to the set $N_{\mathbf{c}^*}$. 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 imply that (4) is satisfied.

Another interesting parameter of P from the quantization viewpoint is the following separation factor. It 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 distortion $\mathbf{c} \mapsto P\gamma(\mathbf{c}, \cdot)$. 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, or have empty cells. 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 [25]). However, this result cannot be extended to noncontinuous distributions, as proved in Example 4.11 of [14].

The main results of this paper are based on the following proposition, which connects the margin condition stated in Definition 2.1 to the previous conditions in [25] or [2]. Recall that $k \geq 2$.

PROPOSITION 2.2. Assume that P satisfies a margin condition with radius r_0 , then the following properties hold.

(i) For every \mathbf{c}^* in \mathcal{M} and \mathbf{c} in $\mathcal{B}(0, M)^k$, if $\|\mathbf{c} - \mathbf{c}^*\| \leq \frac{Br_0}{4\sqrt{2}M}$, then

$$(6) \quad \ell(\mathbf{c}, \mathbf{c}^*) \geq \frac{P_{\min}}{2} \|\mathbf{c} - \mathbf{c}^*\|^2.$$

(ii) \mathcal{M} is finite.

(iii) There exists $\varepsilon > 0$ such that P is ε -separated.

(iv) For all \mathbf{c} in $\mathcal{B}(0, M)^k$,

$$(7) \quad \frac{1}{16M^2} \text{Var}(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*(\mathbf{c}), \cdot)) \leq \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2 \leq \kappa_0 \ell(\mathbf{c}, \mathbf{c}^*),$$

where $\kappa_0 = 4kM^2(\frac{1}{\varepsilon} \vee \frac{64M^2}{P_{\min}B^2r_0^2})$, and $\mathbf{c}^*(\mathbf{c}) \in \arg \min_{\mathbf{c}^* \in \mathcal{M}} \|\mathbf{c} - \mathbf{c}^*\|$.

As a consequence, (7) ensures that (1) is satisfied, with known constant, which is the condition required in Theorem 2 of [2]. Moreover, if $\mathcal{H} = \mathbb{R}^d$, P has a unique optimal codebook up to relabeling, and has a continuous density, (6) 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 [11] for $n\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$ to converge in distribution.

It is worth pointing out that the dependency of κ_0 on different parameters of P 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 $M^2k^{-2/d}$, when P has a density. Furthermore, it is likely that $r_0 \sim B$ (see, e.g., the distributions exposed in Section 3.2). Considering that $\varepsilon \sim R(\mathbf{c}^*) \sim M^2k^{-2/d}$, $r_0 \sim B \sim Mk^{-1/d}$, 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. Proposition 2.2 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. *Assume that $k \geq 2$, and that P satisfies a margin condition with radius r_0 . 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 $1 - e^{-x}$,

$$(8) \quad \ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C_0\kappa_0 \frac{(k + \log(|\bar{\mathcal{M}}|))M^2}{n} + (9\kappa_0 + 4) \frac{16M^2}{n}x,$$

where C_0 is an absolute constant.

This result is in line with Theorem 3.1 in [18] or Theorem 1 in [10], concerning the dependency on the sample size n of the loss $\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*)$. The main advance lies in the detailed dependency on other parameters of the loss of $\hat{\mathbf{c}}_n$. This provides a nonasymptotic bound for the excess risk.

To be more precise, Theorem 3.1 in [18] states that

$$\mathbb{E}\ell(\hat{\mathbf{c}}_n, \mathbf{c}^*) \leq C(k, d, P)M^2/n,$$

in the finite-dimensional case, for some unknown constant $C(k, d, P)$. In fact, this result relies on the application of Dudley's entropy bound. This technique was

already the main argument in [25] or [10], and makes use of covering numbers of the d -dimensional Euclidean unit ball. Consequently, $C(k, d, P)$ strongly depends on the dimension of the underlying Euclidean space in these previous results. As suggested in [6] or [9], 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 by using a technique based on comparison with Gaussian vectors, as done in [9]. Theorem 3.1 is derived that way, providing a dimension-free upper bound which is valid over separable Hilbert spaces.

It may be noticed that most of results providing slow convergence rates, such as Theorem 2.1 in [6] or Corollary 1 in [21], give bounds on the distortion which do not depend on the number of optimal codebooks. Theorem 3.1 confirms that $|\bar{\mathcal{M}}|$ is also likely to play a minor role on the convergence rate of the distortion in the fast rate case.

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 [2, 25] or [10] 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 the requirement that P has a density. It may also be noted that Theorem 3.1 does not require that $\bar{\mathcal{M}}$ contains a single element, contrary to the results stated in [25]. 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 [25]. An illustration of this point is given in Section 3.3.

As will be shown in Proposition 3.1, the dependency on ε turns out to be sharp when $\varepsilon \sim n^{-1/2}$. 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. 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 a margin condition, in which some parameters, such as p_{\min} are fixed or uniformly lower-bounded. It has been already proved in Theorem 4 of [1] that the minimax distortion over distributions with uniformly bounded continuous densities, unique optimal codebooks (up to relabeling), and such that the minimum eigenvalues of the second derivative matrices at the optimal codebooks are uniformly lower-bounded, is $\Omega(1/\sqrt{n})$, in the case where $k = 3$ and $d = 1$. Extending the distributions used in Theorem 4 of [1], Proposition 3.1 below generalizes this result in arbitrary dimension d , and provides a lower bound over a set of distributions satisfying a uniform margin condition.

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

$$(9) \quad \begin{cases} B = \Delta, \\ r_0 = \frac{7\Delta}{16}, \\ p_{\min} \geq \frac{3}{4k}. \end{cases}$$

Denote by $\mathcal{D}(\varepsilon)$ the set of probability distributions which are ε -separated, have continuous densities and unique optimal codebooks, and which satisfy a margin condition with parameters defined in (9). The minimax result is the following.

PROPOSITION 3.1. *Assume that $k \geq 3$ and $n \geq 3k/2$. Then, for any empirically designed codebook,*

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

where $c_0 > 0$ is an absolute constant, and

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

Proposition 3.1 is in line with the previous minimax lower bounds obtained in Theorem 1 of [4] or Theorem 4 of [1]. 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 confirms that the minimax distortion rate over the set of distributions with continuous densities, unique optimal codebooks, and such that the minimum eigenvalues of the Hessian matrices are uniformly lower bounded by $3/8k$, is still $\Omega(1/\sqrt{n})$ in the case where $d > 1$ and $k \geq 3$.

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 $\varepsilon \sim n^{-1/2}$. Unfortunately, Proposition 3.1 cannot 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 exposed in Section 3 with Gaussian mixtures in dimension $d = 2$. The Gaussian mixture model is a typical and well-defined clustering example.

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

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

where \tilde{k} denotes the number of components of the mixture, and the θ_i 's denote the weights of the mixture, which satisfy $\sum_{i=1}^{\tilde{k}} \theta_i = 1$. 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

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

where N_i denotes a normalization constant for each Gaussian variable.

Let η be defined as $\eta = 1 - \min_{i=1, \dots, k} N_i$. Roughly, the model proposed above will be close the Gaussian mixture model when η 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 $i = 1, \dots, k$, $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0, M)$.

Note that the assumption $\mathcal{B}(m_i, \tilde{B}/3) \subset \mathcal{B}(0, M)$ can easily be satisfied if 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*

$$(10) \quad \frac{\theta_{\min}}{\theta_{\max}} \geq \frac{2048k}{(1-\eta)\tilde{B}} \max \left(\frac{\sigma^2}{\tilde{B}(1 - e^{-\tilde{B}^2/2048\sigma^2})}, \frac{kM^3}{7\sigma^2(e^{\tilde{B}^2/32\sigma^2} - 1)} \right).$$

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 (i) in 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 [25]. However, there is no elementary proof of the fact that $|\tilde{\mathcal{M}}| = 1$, 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 presented in [25]. The condition (10) can be decomposed as follows. If

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

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 the $N_{\mathbf{c}^*}$'s, and to derive an upper bound on the weight function $p(t)$ defined above Definition 2.1. This leads to the second term of the maximum in (10).

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 Definition 2.1.

It may be noticed that Proposition 3.2 offers almost the same condition as Proposition 4.2 in [18]. 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 $\bigcup_{\mathbf{c}^*} N_{\mathbf{c}^*}$.

It is important to note that this result is valid when k is known and matches 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 re-labeling, and that we are able to locate this optimal codebook \mathbf{c}^* . As stated 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} .

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} , exposed 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 $\{c \in \mathcal{H} | \phi(c) \leq \lambda\}$ are closed for the weak topology.

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

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

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 simpler in the special case of the squared-norm based quantization in a Hilbert space, it is briefly recalled in Section A.1 (supplementary material [19]).

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 [8]), 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 $P(V_1(\mathbf{c}^{*,k})) < \frac{\alpha}{4M^2}$. 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

$$R(Q) \leq P(V_1(\mathbf{c}^{*,k}))4M^2 + R_{k-1}^* < R_k^*.$$

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

4.2. *Proof of Proposition 2.2.* The proof of (i) in Proposition 2.2 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*

$$(11) \quad \left| \left\langle x - \frac{c_i + c_j}{2}, c_i - c_j \right\rangle \right| \leq 4\sqrt{2}M \|\mathbf{c} - \mathbf{c}^*\|,$$

$$(12) \quad d(x, \partial V_i(\mathbf{c}^*)) \leq \frac{4\sqrt{2}M}{B} \|\mathbf{c} - \mathbf{c}^*\|.$$

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 labels 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 a simpler proof in Section A.2 (supplementary material [19]).

Equipped with Lemma 4.2, we are in a position to prove (6). 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:

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

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

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

from which we deduce that

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

which leads to

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

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

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

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

$$(\|x - c_i\|^2 - \|x - c_j\|^2) \mathbb{1}_{V_i(\mathbf{c}^*) \cap W_j(\mathbf{c})}(x) \\ \leq 8\sqrt{2}M \|\mathbf{c} - \mathbf{c}^*\| \mathbb{1}_{V_i(\mathbf{c}^*) \cap W_j(\mathbf{c}) \cap \mathcal{N}_{\mathbf{c}^*}((4\sqrt{2}M/B) \|\mathbf{c} - \mathbf{c}^*\|)}(x).$$

Hence,

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

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

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

which proves (i).

Suppose that \mathcal{M} is not finite. According to Lemma 4.1, there exists a sequence \mathbf{c}_n of optimal codebooks and an optimal codebook \mathbf{c}^* such that for all n ,

$\mathbf{c}_n \neq \mathbf{c}^*$ and $\mathbf{c}_n \xrightarrow{n \rightarrow \infty} \mathbf{c}^*$. Assume that there exists i in $\{1, \dots, k\}$ such that $\liminf_n \|c_{n,i}\|^2 > \|c_i\|^2$. Then $\liminf_n \|x - c_{n,i}\|^2 > \|x - c_i\|^2$, for every x in \mathcal{H} .

Let x be in $\overset{\circ}{V}_i(\mathbf{c})$, and $j \neq i$, then

$$\liminf_{n \rightarrow \infty} \|x - c_{n,j}\|^2 \geq \|x - c_j\|^2 > \|x - c_i\|^2,$$

which leads to $\liminf_n \gamma(\mathbf{c}_n, x) > \gamma(\mathbf{c}, x)$. Since $P(\overset{\circ}{V}_i(\mathbf{c})) > 0$, it easily follows that

$$\liminf_{n \rightarrow \infty} P\gamma(\mathbf{c}_n, \cdot) \geq P \liminf_{n \rightarrow \infty} \gamma(\mathbf{c}_n, \cdot) > P\gamma(\mathbf{c}, \cdot),$$

which is impossible. Hence, there exists a subsequence $\bar{\mathbf{c}}_n$ of \mathbf{c}_n such that, for $i = 1, \dots, k$, $\|\bar{c}_{n,i}\| \rightarrow_{n \rightarrow \infty} \|c_i^*\|$. Since Hilbert spaces are uniformly convex spaces, hence satisfy the Radon–Riesz property (see, e.g., Propositions 5.1 and 3.32 in [8]), it follows that $\bar{\mathbf{c}}_n \rightarrow_{n \rightarrow \infty} \mathbf{c}^*$. This contradicts (6) and proves (ii).

The proof of (iii) is based on the following two lemmas.

LEMMA 4.3. *Let \mathbf{c} be in $\mathcal{B}(\mathcal{M}, \frac{Br_0}{4\sqrt{2}M})$. If \mathbf{c} satisfies the centroid condition, then \mathbf{c} is in \mathcal{M} .*

Lemma 4.3 ensures that no local minimizer with nonempty cells can be found in a neighborhood of \mathcal{M} . We postpone its proof to Section A.3 (supplementary material [19]). Lemma 4.4 below shows that the infimum distortion over codebooks which are away from \mathcal{M} is achieved.

LEMMA 4.4. *For every $r > 0$, there exists \mathbf{c}_r in $\mathcal{B}(0, M+r)^k \setminus \mathcal{B}^o(\mathcal{M}, r)$ such that*

$$\inf_{\mathcal{H}^k \setminus \mathcal{B}^o(\mathcal{M}, r)} P\gamma(\mathbf{c}, \cdot) = P\gamma(\mathbf{c}_r, \cdot).$$

The proof of Lemma 4.4 is given in Section A.4 (supplementary material [19]). Let $\tilde{\mathbf{c}} \notin \mathcal{M}$ be a local minimizer of the distortion. If $\tilde{\mathbf{c}}$ has empty cells, then $P\gamma(\tilde{\mathbf{c}}, \cdot) \geq R_{k-1}^* > R_k^*$. Assume that $\tilde{\mathbf{c}}$ has no empty cells. Then $\tilde{\mathbf{c}}$ satisfies the centroid condition, thus Lemma 4.3 ensures that $\|\tilde{\mathbf{c}} - \mathbf{c}^*\| \geq r$, for every optimal codebook \mathbf{c}^* and for $r = \frac{Br_0}{4\sqrt{2}M}$. Lemma 4.4 provides \mathbf{c}_r such that $P\gamma(\tilde{\mathbf{c}}, \cdot) \geq P\gamma(\mathbf{c}_r, \cdot) > 0$. Hence, (iii) is proved.

The left part of (7) follows from the elementary inequality

$$(13) \quad \forall x \in \mathcal{B}(0, M) \quad |\gamma(\mathbf{c}, x) - \gamma(\mathbf{c}^*(\mathbf{c}), x)| \leq 4M \max_{i=1, \dots, k} \|c_i - c_i^*(\mathbf{c})\|.$$

According to (6), if $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\| \leq \frac{Br_0}{4\sqrt{2}M}$, then $\ell(\mathbf{c}, \mathbf{c}^*) \geq \frac{P_{\min}}{2} \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$. Now turn to the case where $\|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\| \geq \frac{Br_0}{4\sqrt{2}M} = r$. Then Lemma 4.4 provides \mathbf{c}_r such that $\ell(\mathbf{c}, \mathbf{c}^*) \geq \ell(\mathbf{c}_r, \mathbf{c}^*)$. Such a \mathbf{c}_r is a local minimum of $\mathbf{c} \mapsto P\gamma(\mathbf{c}, \cdot)$, or satisfies

$\|\mathbf{c}_r - \mathbf{c}^*(\mathbf{c}_r)\| = r$. Hence, we deduce

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

Note that, since $B \leq 2M$ and $r_0 \leq 2M$, $(\varepsilon \wedge \frac{p_{\min} B^2 r_0^2}{64M^2})/4kM^2 \leq p_{\min}/2$. This proves (7).

4.3. *Proof of Theorem 3.1.* Throughout this subsection, P is assumed to satisfy a margin condition with radius r_0 , and we denote by ε its separation factor. A nondecreasing map $\Phi: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is called sub-root if $x \mapsto \frac{\Phi(x)}{\sqrt{x}}$ is nonincreasing.

The following localization theorem, derived from Theorem 6.1 in [7], is the main argument of our proof.

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

$$\forall f \in \mathcal{F}, \quad \text{Var}(f) \leq \omega(f).$$

Assume that

$$\forall r > 0, \quad \mathbb{E} \left(\sup_{\omega(f) \leq r} |(P - P_n)f| \right) \leq \Phi(r),$$

for some sub-root function Φ . Let K be a positive constant, and denote by r^* the unique solution of the equation $\Phi(r) = r/24K$.

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

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

A proof of Theorem 4.1 is given in Section 5.3 of [18]. The proof of (8) 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

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

According to (13), it is clear that, for every $f \in \mathcal{F}$,

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

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.

PROPOSITION 4.1. *One has*

$$(14) \quad \mathbb{E} \sup_{f \in \mathcal{F}, \omega(f) \leq \delta} |(P - P_n)f| \leq \frac{(4\sqrt{\pi k} + \sqrt{2 \log(|\bar{\mathcal{M}}|)})}{\sqrt{n}} \sqrt{\delta}.$$

The proof of Proposition 4.1 relies on the use of Gaussian complexities combined with Slepian's lemma (see, e.g., Theorem 3.14 in [23]), as done in [9]. 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

$$\delta = \frac{576K^2}{n} (4\sqrt{\pi k} + \sqrt{2 \log(|\bar{\mathcal{M}}|)})^2 \leq C \frac{K^2(k + \log(|\bar{\mathcal{M}}|))}{n} := \frac{K^2 \Xi}{n},$$

where $C = 18,432\pi$, and $\Xi = C(k + \log(|\bar{\mathcal{M}}|))$. Applying Theorem 4.1 to \mathcal{F} leads to, with probability larger than $1 - e^{-x}$,

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

Introducing the inequality $\kappa_0 \ell(\mathbf{c}, \mathbf{c}^*) \geq \|\mathbf{c} - \mathbf{c}^*(\mathbf{c})\|^2$ provided by Proposition 2.2, and choosing $K = 32M^2 \kappa_0$ leads to (8).

4.3.1. *Proof of Proposition 4.1.* As mentioned above, this proof relies on the use of Gaussian complexities (see, e.g., [5]). 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:

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

Let $Y_{\mathbf{c}}^*$ denote the random variable defined by

$$Y_{\mathbf{c}}^* = \sup_{\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2} |(P - P_n)(\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*, \cdot))|,$$

for every \mathbf{c}^* in $\bar{\mathcal{M}}$. It easily follows that

$$(15) \quad \mathbb{E} \sup_{\mathbf{c}^* \in \bar{\mathcal{M}}} Y_{\mathbf{c}}^* \leq \mathbb{E} \sup_{\mathbf{c}^* \in \bar{\mathcal{M}}} (Y_{\mathbf{c}}^* - \mathbb{E} Y_{\mathbf{c}}^*) + \sup_{\mathbf{c}^* \in \bar{\mathcal{M}}} \mathbb{E} Y_{\mathbf{c}}^*.$$

Since, for a fixed \mathbf{c}^* , $\|\gamma(\mathbf{c}, \cdot) - \gamma(\mathbf{c}^*, \cdot)\|_\infty \leq \sqrt{\delta}$ when $\|\mathbf{c} - \mathbf{c}^*\|^2 \leq \delta/16M^2$, the bounded difference inequality (see, e.g., Theorem 5.1 in [23]) ensures that $Y_{\mathbf{c}^*}$ is a sub-Gaussian random variable, with variance bounded from above by δ/n , that is,

$$\begin{cases} \mathbb{P}\left(Y_{\mathbf{c}^*} - \mathbb{E}Y_{\mathbf{c}^*} \geq \sqrt{\frac{2\delta x}{n}}\right) \leq e^{-x}, \\ \mathbb{P}\left(\mathbb{E}Y_{\mathbf{c}^*} - Y_{\mathbf{c}^*} \geq \sqrt{\frac{2\delta x}{n}}\right) \leq e^{-x}, \end{cases}$$

for every \mathbf{c}^* in $\bar{\mathcal{M}}$ and every positive x . For a more general definition of sub-Gaussian random variables, the interested reader is referred to [23]. Applying Lemma 6.3 in [23] to the special case of sub-Gaussian random variables leads to

$$(16) \quad \mathbb{E} \sup_{\mathbf{c}^* \in \bar{\mathcal{M}}} (Y_{\mathbf{c}^*} - \mathbb{E}Y_{\mathbf{c}^*}) \leq \sqrt{\frac{2 \log(|\bar{\mathcal{M}}|) \delta}{n}}.$$

Next, we bound from above the quantities $\mathbb{E}Y_{\mathbf{c}^*}$. 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 denotes integration with respect to the distribution of Y . Let g_1, \dots, g_n denote some independent standard Gaussian variables. Applying Lemma 4.5 in [17] leads to

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

To derive bounds on the Gaussian complexity defined above, the following comparison result between Gaussian processes is needed.

THEOREM 4.2 (Slepian’s lemma). *Let X_t and Z_t , t in \mathcal{V} , be some centered real Gaussian processes. Assume that*

$$\forall s, t \in \mathcal{V}, \quad \text{Var}(Z_s - Z_t) \leq \text{Var}(X_s - X_t),$$

then

$$\mathbb{E} \sup_{t \in \mathcal{V}} Z_t \leq 2\mathbb{E} \sup_{t \in \mathcal{V}} X_t.$$

A proof of Theorem 4.2 can be found in Theorem 3.14 of [23]. For a fixed sample X_1, \dots, X_n , define the Gaussian process $Z_{\mathbf{c}}$ by

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

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

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

Define now the Gaussian process $X_{\mathbf{c}}$ by

$$X_{\mathbf{c}} = 2\sqrt{2} \sum_{i=1}^n \sum_{j=1}^k \langle c_j - c_j^*, X_i \rangle \xi_{i,j} + \sqrt{2n} \sum_{j=1}^k (\|c_j\|^2 - \|c_j^*\|^2) \xi'_j,$$

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

$$\begin{aligned} \mathbb{E}_g \sup_{\mathbf{c} \in \mathcal{V}(\delta)} Z_{\mathbf{c}} &\leq 2\mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} X_{\mathbf{c}} \\ (17) \quad &\leq 4\sqrt{2}\mathbb{E}_{\xi} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{i=1}^n \sum_{j=1}^k \langle c_j - c_j^*, X_i \rangle \xi_{i,j} \\ &\quad + 2\sqrt{2n}\mathbb{E}_{\xi'} \sup_{\mathbf{c} \in \mathcal{V}(\delta)} \sum_{j=1}^k (\|c_j\|^2 - \|c_j^*\|^2) \xi'_j. \end{aligned}$$

Using almost the same technique as in the proof of Theorem 2.1 in [6], the first term of the right-hand side of (17) can be bounded as follows:

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

$$\begin{aligned} &\leq \frac{\sqrt{\delta}}{4M} \sqrt{\sum_{j=1}^k \mathbb{E}_{\xi} \left\| \sum_{i=1}^n \xi_{i,j} X_i \right\|^2} \\ &\leq \frac{\sqrt{k\delta}}{4M} \sqrt{\sum_{i=1}^n \|X_i\|^2}. \end{aligned}$$

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 (17) can be bounded from above by

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

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

$$\mathbb{E}_{X,g} \sup_{\|\mathbf{c}-\mathbf{c}^*\|^2 \leq \delta/16M^2} Z_{\mathbf{c}} \leq 2\sqrt{2kn}\sqrt{\delta},$$

which leads to

$$(18) \quad \mathbb{E}Y_{\mathbf{c}^*} \leq \frac{4\sqrt{k\pi\delta}}{\sqrt{n}}.$$

Combining (16) and (18) into (15) gives the 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 $i = 1, \dots, m$, 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 which can be packed in $\mathcal{B}(0, M - \Delta/16)$. Since the cardinal of a maximal 6Δ -net is larger than the smallest number of balls of radius 6Δ which together cover $\mathcal{B}(0, M - \Delta/16)$, 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 $i = 1, \dots, m$, 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

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

where λ denotes the Lebesgue measure and $\delta \leq 1/3$. These distributions have been designed to have continuous cone-shaped densities, as in Theorem 4 of [1].

Similarly, 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 $\sigma_i = +1$. At last, for τ in $\{-1, +1\}^{m/2}$, $\sigma(\tau)$ is defined as the sequence in $\{-1, +1\}^m$ such that

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

for $i = 1, \dots, \frac{m}{2}$, and the set of corresponding $Q_{\sigma(\tau)}$'s is denoted by \mathcal{Q} .

Given a quantizer Q , let $R(Q, P_\sigma)$ and $\ell(Q, P_\sigma)$ denote, respectively, the distortion and loss of Q in the case where the source distribution is P_σ . Proposition 4.2 below shows that only quantizers in \mathcal{Q} may be considered in order to derive lower bounds on R .

PROPOSITION 4.2. *Let σ and σ' be in $\{-1, +1\}^m$ such that $\sum_{i=1}^m \sigma_i = \sum_{i=1}^m \sigma'_i = 0$, and let $\rho(\sigma, \sigma')$ denote the distance $\sum_{i=1}^m |\sigma_i - \sigma'_i|$. Then*

$$(19) \quad R(Q_{\sigma'}, P_\sigma) = R(Q_\sigma, P_\sigma) + \frac{\Delta^2 \delta}{8m} \rho(\sigma, \sigma').$$

Furthermore, for every nearest neighbor quantizer Q , there exists σ and τ such that

$$\forall P_{\sigma(\tau')} \quad R(Q, P_{\sigma(\tau')}) \geq R(Q_\sigma, P_{\sigma(\tau')}) \geq \frac{1}{2}R(Q_{\sigma(\tau)}, P_{\sigma(\tau')}).$$

At last, if $Q \neq Q_\sigma$, then the first inequality is strict, for every $P_{\sigma(\tau')}$.

The proof of Proposition 4.2 follows the proof of step 3 of Theorem 1 in [4], and can be found in Section B.1 (supplementary material [19]).

Since, for $\sigma \neq \sigma'$, $R(Q'_{\sigma'}, P_\sigma) > R(Q_\sigma, P_\sigma)$, Proposition 4.2 ensures that the $P_{\sigma(\tau)}$'s have unique optimal codebooks, up to relabeling. According to Proposition 4.2, the minimax lower-bound over empirically designed quantizer may be reduced to a lower-bound on empirically designed τ 's, that is,

$$\begin{aligned} (20) \quad & \inf_{\hat{Q}_n} \sup_{\tau \in \{-1, +1\}^{m/2}} \mathbb{E} \ell(\hat{Q}_n, P_{\sigma(\tau)}) \\ & \geq \frac{1}{2} \inf_{\hat{\tau}} \sup_{\tau \in \{-1, +1\}^{m/2}} \mathbb{E} \ell(Q_{\sigma(\hat{\tau})}, P_{\sigma(\tau)}) \\ & \geq \frac{\Delta^2 \delta}{8m} \inf_{\hat{\tau}} \sup_{\tau \in \{-1, +1\}^{m/2}} \mathbb{E} \rho(\hat{\tau}, \tau), \end{aligned}$$

where the inequality $\rho(\sigma(\tau), \sigma(\tau')) = 2\rho(\tau, \tau')$ has been used in the last inequality.

Let us define, for two distributions P and Q with densities f and g , the Hellinger distance

$$H^2(P, Q) = \int_{\mathbb{R}^d} (\sqrt{f} - \sqrt{g})^2(x) d\lambda(x).$$

To apply Assouad's lemma to Q , the following lemma is needed.

LEMMA 4.5. *Let τ and τ' denote two sequences in $\{-1, +1\}^{m/2}$ such that $\rho(\tau, \tau') = 2$. Then*

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

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

The proof of Lemma 4.5 is given in Section B.2 (supplementary material [19]). Equipped with Lemma 4.5, a direct application of Assouad's lemma as in Theorem 2.12 of [26] yields, provided that $\alpha \leq 2$,

$$\inf_{\hat{\tau}} \sup_{\tau \in \{-1, +1\}^{m/2}} \mathbb{E} \rho(\hat{\tau}, \tau) \geq \frac{m}{4} (1 - \sqrt{\alpha(1 - \alpha/4)}).$$

Taking $\delta = \frac{\sqrt{m}}{2\sqrt{n}}$ ensures that $\alpha \leq 2$. For this value of δ , it easily follows from (20) that

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

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

Finally, note that, for every $\delta \leq \frac{1}{3}$ and σ , P_σ satisfies a margin condition as in (9), and is ε -separated, with

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

This completes the proof of Proposition 3.1.

4.5. *Proof of Proposition 3.2.* As mentioned below in 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

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

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

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

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 $N_{\mathbf{c}^*}(x)$, for some \mathbf{c}^* in \mathcal{M} and 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

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

Since the Lebesgue measure of $N_{\mathbf{c}^*}(x)$ is smaller than $4k\pi Mx$, it follows that

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

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

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

Therefore,

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

hence $p_{\min} \geq \theta_{\min} (1 - e^{-9\tilde{B}^2 / (128\sigma^2)})$. 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

$$P(N_{\mathbf{c}^*}(x)) \leq \frac{B p_{\min}}{128M^2} x.$$

This ensures that P satisfies (3). According to (ii) in Proposition 2.2, \mathcal{M} is finite.

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

SUPPLEMENTARY MATERIAL

Appendix: Remaining proofs (DOI: [10.1214/14-AOS1293SUPP](https://doi.org/10.1214/14-AOS1293SUPP); .pdf). Due to space constraints, we relegate technical details of the remaining proofs to the supplement [19].

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](#)
- [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](#)
- [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](#)
- [5] BARTLETT, P. L. and MENDELSON, S. (2002). Rademacher and Gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.* **3** 463–482. [MR1984026](#)
- [6] BIAU, G., DEVROYE, L. and LUGOSI, G. (2008). On the performance of clustering in Hilbert spaces. *IEEE Trans. Inform. Theory* **54** 781–790. [MR2444554](#)
- [7] BLANCHARD, G., BOUSQUET, O. and MASSART, P. (2008). Statistical performance of support vector machines. *Ann. Statist.* **36** 489–531. [MR2396805](#)
- [8] BREZIS, H. (2011). *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, New York. [MR2759829](#)
- [9] CAÑAS, G. D., POGGIO, T. and ROSASCO, L. (2012). Learning manifolds with k -means and k -flats. *CoRR* abs/1209.1121.
- [10] CHICHIGNOUD, M. and LOUSTAU, S. (2014). Adaptive noisy clustering. *IEEE Trans. Inform. Theory* **60** 7279–7292.
- [11] 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*. IEEE, Trondheim.
- [12] FISCHER, A. (2010). Quantization and clustering with Bregman divergences. *J. Multivariate Anal.* **101** 2207–2221. [MR2671211](#)
- [13] GERSHO, A. and GRAY, R. M. (1991). *Vector Quantization and Signal Compression*. Kluwer Academic, Norwell, MA.
- [14] GRAF, S. and LUSCHGY, H. (2000). *Foundations of Quantization for Probability Distributions. Lecture Notes in Math.* **1730**. Springer, Berlin. [MR1764176](#)
- [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](#)
- [16] KOLTCHINSKII, V. (2006). Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Statist.* **34** 2593–2656. [MR2329442](#)
- [17] LEDOUX, M. and TALAGRAND, M. (1991). *Probability in Banach Spaces: Isoperimetry and Processes. Ergebnisse der Mathematik und Ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]* **23**. Springer, Berlin. [MR1102015](#)
- [18] LEVRARD, C. (2013). Fast rates for empirical vector quantization. *Electron. J. Stat.* **7** 1716–1746. [MR3080408](#)
- [19] LEVRARD, C. (2015). Supplement to “Nonasymptotic bounds for vector quantization in Hilbert spaces.” DOI:[10.1214/14-AOS1293SUPP](#).
- [20] LINDER, T. (2002). Learning-theoretic methods in vector quantization. In *Principles of Non-parametric Learning (Udine, 2001). CISM Courses and Lectures* **434** 163–210. Springer, Vienna. [MR1987659](#)
- [21] 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](#)
- [22] MAMMEN, E. and TSYBAKOV, A. B. (1999). Smooth discrimination analysis. *Ann. Statist.* **27** 1808–1829. [MR1765618](#)
- [23] MASSART, P. (2007). *Concentration Inequalities and Model Selection. Lecture Notes in Math.* **1896**. Springer, Berlin. [MR2319879](#)

- [24] MASSART, P. and NÉDÉLEC, É. (2006). Risk bounds for statistical learning. *Ann. Statist.* **34** 2326–2366. [MR2291502](#)
- [25] POLLARD, D. (1982). A central limit theorem for k -means clustering. *Ann. Probab.* **10** 919–926. [MR0672292](#)
- [26] TSYBAKOV, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer, New York. Revised and extended from the 2004 French original. Translated by Vladimir Zaiats. [MR2724359](#)

BÂTIMENT ALAN TURING
CAMPUS DE L'ÉCOLE POLYTECHNIQUE
INRIA
91120 PALAISEAU
FRANCE
E-MAIL: clement.levrard@inria.fr