Small Width, Low Distortions: Quantized Random Embeddings of Low-complexity Sets

Laurent Jacques*

November 15, 2016

Abstract

Under which conditions and with which distortions can we preserve the pairwise-distances of low-complexity vectors, *e.g.*, for *structured sets* such as the set of sparse vectors or the one of low-rank matrices, when these are mapped (or embedded) in a finite set of vectors?

This work addresses this general question through the specific use of a quantized and dithered random linear mapping which combines, in the following order, a sub-Gaussian random projection in $\mathbb{R}^M$ of vectors in $\mathbb{R}^N$, a random translation, or *dither*, of the projected vectors and a uniform scalar quantizer of resolution $\delta > 0$ applied componentwise.

Thanks to this quantized mapping we are first able to show that, with high probability, an embedding of a bounded set $\mathcal{K} \subset \mathbb{R}^N$ in $\delta\mathbb{Z}^M$ can be achieved when distances in the quantized and in the original domains are measured with the ℓ_1- and ℓ_2-norm, respectively, and provided the number of quantized observations M is large before the square of the "Gaussian mean width" of $\mathcal{K}$. In this case, we show that the embedding is actually *quasi-isometric* and only suffers of both multiplicative and additive distortions whose magnitudes decrease as $M^{-1/5}$ for general sets, and as $M^{-1/2}$ for structured set, when M increases. Second, when one is only interested in characterizing the maximal distance separating two elements of $\mathcal{K}$ mapped to the same quantized vector, *i.e.*, the "consistency width" of the mapping, we show that for a similar number of measurements and with high probability this width decays as $M^{-1/4}$ for general sets and as $1/M$ for structured ones when M increases. Finally, as an important aspect of our work, we also establish how the non-Gaussianity of sub-Gaussian random projections inserted in the quantized mapping (*e.g.*, for Bernoulli random matrices) impacts the class of vectors that can be embedded or whose consistency width provably decays when M increases.

1 Introduction

There exists an ever-growing trend in high (or "big") dimensional data processing to design new procedures (or to simplify existing ones) using linear dimensionality reduction (LDR) methods in order to get faster or memory-efficient algorithms. Provided this reduction does not bring too much distortion between the initial data space and the "reduced" domain, as often allowed by the intrinsic "low-dimensionality" properties of the input data, many techniques, such as nearest-neighbor search in big databases [1, 3], classification [5], regression [38], filtering [17], manifold processing [7] or compressed sensing [11, 21] can be developed in this reduced domain with controlled loss of accuracy, as well as stability with respect to data corruption (*e.g.*, noise).

*Image and Signal Processing Group (ISPGroup), ELEN Department, ICTEAM institute, Université catholique de Louvain (UCL), Belgium. The author is funded by Belgian National Science Foundation (F.R.S.-FNRS)

Most often, those LDR tools rely on defining a random projection matrix (sometimes called *sensing* matrix) with fewer rows M than columns N, whose multiplication with data represented as a set of vectors in $\mathbb{R}^N$ provides a reduced representation (or *sketch*) of the latter. This is the scheme implicitly promoted for instance by the celebrated Johnson-Lindenstrauss (JL) lemma for finite sets of vectors $\mathcal{S} \subset \mathbb{R}^N$, *i.e.*, with $|\mathcal{S}| < \infty$ [31]. This cornerstone result and its subsequent developments [1, 15] showed that, given a resolution $\epsilon > 0$, if $M \geqslant C\epsilon^{-2} \log S$ where $S = |\mathcal{S}|$ is the cardinality of $\mathcal{S}$ and $C > 0$ is a general constant, then a random matrix $\boldsymbol{\Phi} \in \mathbb{R}^{M \times N}$ whose entries are independently and identically distributed (i.i.d.) as a centered sub-Gaussian distribution with unit variance defines an isometric mapping that preserves pairwise-distances between points in $\mathcal{S}$ up to a multiplicative distortion ϵ. In other words, $\boldsymbol{\Phi}$ defines an ϵ-*isometry* between $(\mathcal{S}, \ell_2)$ and $(\boldsymbol{\Phi}\mathcal{S}, \ell_2)$, *i.e.*, with high probability, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{S}$,

$$(1 - \epsilon)\|\boldsymbol{x} - \boldsymbol{y}\| \leqslant \tfrac{1}{\sqrt{M}}\|\boldsymbol{\Phi x} - \boldsymbol{\Phi y}\| \leqslant (1 + \epsilon)\|\boldsymbol{x} - \boldsymbol{y}\|. \tag{1}$$

Equivalently, one observes that keeping the probability of success constant with respect to the random generation of $\boldsymbol{\Phi}$ and inverting the requirement linking M and ϵ, such an isometry has a distortion ϵ decaying as $1/\sqrt{M}$ when M increases, *i.e.*, this distortion vanishes when $M/\log S$ is large. Notice that variants of this embedding result exist with different "input/output" norms; see, *e.g.*, [36] for a unified treatment over a family of *interpolation* norms including ℓ_2 and ℓ_1 as special cases.

The JL lemma has been later generalized to any subsets $\mathcal{K} \subset \mathbb{R}^N$, not only finite, whose typical "dimension" can be considered as small with respect to N (see, *e.g.*, [7, 19, 39]). In other words, as soon as $\mathcal{K}$ displays some internal structure that makes it somehow parametrisable with much fewer parameters than N, as for the set of sparse or compressible signals, the set of low-rank matrices, signal manifolds, or a set given as a union of low-dimensional subspaces, an ϵ-isometry like (1) can be defined for all pairs of vectors in $\mathcal{K}$. This is for instance the essence of the restricted isometry property (RIP) and its link with the JL lemma, where (1) holds with high probability for all K-sparse vectors provided $M \geqslant CK \log N/K$ [6, 11].

However, these embeddings have one strong limitation. Except in very specific situations, such as for discrete sub-Gaussian random matrices $\boldsymbol{\Phi}$ (*e.g.*, Bernoulli) and finite sets $\mathcal{K}$, the set $\boldsymbol{\Phi}\mathcal{K} \subset \mathbb{R}^M$ is not finite. An infinite number of bits is thus required if one needs to store, process or transmit $\boldsymbol{\Phi x}$ without information loss for any possible $\boldsymbol{x} \in \mathcal{K}$. Moreover, knowing how many bits are required to represent such projections is also important theoretically for assessing and measuring the level of information contained in the reduced data space or for improving specific data retrieval and processing algorithms. Additionally, if this measure of information can be achieved, nothing prevents us to take $M \geqslant N$, as the sought "dimensionality reduction" can be aimed at minimizing the number of bits rather than the dimensionality M. For instance, [3] defines locality-sensitive hashing (LSH) as a procedure to turn data vectors into quantized *hashes* that preserve locality, so that close vectors induce, with high probability, close hashes. However, this method is specifically designed for boosting nearest-neighbor searches over a finite set of vectors and not to define an isometry similar to (1).

As a more practical solution, the embedding realized by a random projection $\boldsymbol{\Phi}$ is often followed by a scalar quantization procedure, *e.g.*, with a uniform scalar quantizer $\mathcal{Q} : \mathbb{R} \to \delta\mathbb{Z}$ with resolution $\delta > 0$, applied componentwise on the image of $\boldsymbol{\Phi}$. A direct impact of this sequence of operations is to induce a new additive distortion in (1) related to δ, as discussed in [10]. Indeed, assuming $\boldsymbol{\Phi}$ respects (1) for all $\boldsymbol{x}$ and $\boldsymbol{y}$ in a certain subset $\mathcal{K} \subset \mathbb{R}^N$, given a uniform quantizer $\mathcal{Q}(\cdot) := \delta\lfloor \frac{\cdot}{\delta} + \frac{1}{2} \rfloor$ of resolution $\delta > 0$ applied componentwise on vectors of $\mathbb{R}^M$ we would have $|\mathcal{Q}(\lambda) - \lambda| \leqslant \delta/2$ for all $\lambda \in \mathbb{R}$, which involves $\|\mathcal{Q}(\boldsymbol{u}) - \boldsymbol{u}\| \leqslant \sqrt{M}\delta/2$ for any $\boldsymbol{u} \in \mathbb{R}^M$. Therefore, a simple manipulation of (1) provides

$$(1-\epsilon)\|\boldsymbol{x}-\boldsymbol{y}\| - \delta \ \leqslant\ \tfrac{1}{\sqrt{M}}\|\mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{x}) - \mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{y})\| \leqslant\ (1+\epsilon)\|\boldsymbol{x}-\boldsymbol{y}\| + \delta. \tag{2}$$

In other words, as described in Sec. 2, the quantized mapping $\boldsymbol{A}(\cdot) := \mathcal{Q}(\boldsymbol{\Phi}\,\cdot)$ defines now a *quasi-isometric* embedding between $(\mathcal{K} \subset \mathbb{R}^N, \ell_2)$ and $(\boldsymbol{A}(\mathcal{K}) \subset \delta\mathbb{Z}^M, \ell_2)$.

However, while (2) displays a constant additive distortion, several works in this context have observed that such an additive error actually decays as M increases. First, when distances in the reduced space are measured with the ℓ_1-norm and when $\mathcal{Q}$ is combined with a *dithering*[1], a quasi-isometry similar to (2) holds with high probability for all vectors in a *finite* set $\mathcal{K} = \mathcal{S}$ [27] . The additive distortion reads then $c\delta\epsilon$ for some absolute constant $c > 0$ and this error also decays as $1/\sqrt{M}$, as does the multiplicative error ϵ. Second, when combined with universal quantization [10], *i.e.*, with a periodic scalar quantizer $\mathcal{Q}$, an exponential decay of this distortion as M grows can be reached; for the moment, this has been proved only for sparse signal sets. Finally, recent works related to 1-bit compressed sensing (CS) have shown that for a quantization $\mathcal{Q}$ reduced to a sign operator (*i.e.*, $\mathcal{Q}(\boldsymbol{\Phi}\,\cdot) = \operatorname{sign}(\boldsymbol{\Phi}\,\cdot)$) the angular distance between any pair of vectors of a low-dimensionality set $\mathcal{K}$ is close to the Hamming distance of their mappings up to an additive error decaying as $1/M^{1/q}$ for some $q \geqslant 2$. This is true for random Gaussian matrices and for the set of sparse signals [29, 44], for any sets with "low dimensionality" as measured by their Gaussian mean width [44, 46] (see below) and even for sub-Gaussian random matrices provided the projected vectors are not "too sparse" [2], *i.e.*, for vectors whose ℓ_∞-norm is much smaller than their ℓ_2-norm.

Contributions: Considering these last observations, the main results of this paper show that:

(i) quasi-isometric embeddings can be obtained with high probability from scalar (dithered) quantization after linear random projection; for such embeddings both multiplicative and additive distortions co-exist when, as in [27], distances between mapped vectors are measured with the ℓ_1-norm[2];

(ii) random sensing matrices for such embeddings are allowed to be generated from symmetric sub-Gaussian distributions provided embedded vector differences are not "too sparse" (as in the 1-bit case [2]);

(iii) the results above actually hold with high probability for *any* subset $\mathcal{K}$ of $\mathbb{R}^N$ as soon as M is large compared to its typical dimension, *i.e.*, to its squared Gaussian mean width.

(iv) with high probability, the biggest distance separating two *consistent* vectors in $\mathcal{K}$ (*i.e.*, characterized by identical quantized mappings), that is what we call the *consistency width*, decays when M increases at a faster rate than what could be predicted by using just the implications of a quasi-isometry. This extends to any set $\mathcal{K}$ the works of [28, 47], that were valid only for sparse signals;

(v) for particular *structured* sets, *e.g.*, the set of (bounded) sparse vectors or the set of (bounded) low-rank matrices, the minimal values of M necessary to specify a quantized embedding or a small consistency width can be strongly reduced compared to those required for a general set;

[1] That is, when the quantizer input is randomly shifted *inside* the quantization bin by a random translation adjusted to the quantizer resolution [24] (see Sec. 2 and Eq. (4)).

[2] Notice that for binary embeddings the Hamming distance separating the binary mapping of two vectors, as used in [25, 44], is also the half of their ℓ_1-distance.

Moreover, we aim at optimizing whenever it is possible the requirements on M (*e.g.*, with respect to ϵ and δ) that guarantee those results.

Methodology: As an important aspect of our developments, we study the conditions for obtaining quasi-isometric embeddings of any bounded subsets $\mathcal{K} \subset \mathbb{R}^N$ into $\delta\mathbb{Z}^M$. Following key procedures established in other works [44, 45], the typical dimension of these sets is measured by the *Gaussian mean width*, *i.e.*,

$$w(\mathcal{K}) := \mathbb{E} \sup_{\boldsymbol{u} \in \mathcal{K}} |\boldsymbol{g}^\top \boldsymbol{u}|,$$

with $\boldsymbol{g} \sim \mathcal{N}^N(0,1)$. This quantity, also known as Gaussian complexity, has been recognized as central for instance in characterizing random processes [51], shrinkage estimators in signal denoising and high-dimensional statistics [12], linear inverse problem solving with convex optimization [13] or classification efficiency for randomly projected signal sets [5]. More specifically, the minimal number of measurements M necessary to induce, with high probability, an ℓ_2/ℓ_2-isometric embedding of any subset $\mathcal{K} \subset \mathbb{S}^{N-1}$ into $\mathbb{R}^M$ from sub-Gaussian random projections is known to be proportional to $w(\mathcal{K})^2$ [39]. Therefore, since $w(\mathcal{K})^2 \lesssim \log|\mathcal{K}|$ for some finite set $\mathcal{K}$, we recover the condition defining the Johnson-Lindenstrauss lemma by imposing $M \gtrsim \log|\mathcal{K}|$ [31], while for the set of bounded K-sparse vectors in an orthonormal basis (ONB) $\boldsymbol{\Psi} \in \mathbb{R}^{N\times N}$, $w(\mathcal{K})^2 \lesssim K \log N/K$, which characterizes the conditions of the restricted isometry property (RIP) for sub-Gaussian random matrices [6]. The interested reader can find a summary of the main properties of the Gaussian mean width in Table 1, with explicit references to their origin. This table could be helpful also to keep trace of these properties while reading our proofs.

In our developments, we sometimes complete the characterization of sets provided by the Gaussian mean width with another important measure: the Kolmogorov ϵ-entropy of a set $\mathcal{K} \subset \mathbb{R}^N$ that we denote $\mathcal{H}(\mathcal{K}, \epsilon)$ [35]. This is defined as the logarithm of the size of the smallest ϵ-net of $\mathcal{K}$, *i.e.*, a set $\mathcal{C}_\epsilon(\mathcal{K}) \subset \mathcal{K}$ such that any vector of $\mathcal{K}$ cannot be farther than ϵ from its closest vector in $\mathcal{C}_\epsilon(\mathcal{K})$. By the Sudakov inequality, this entropy is connected to the Gaussian mean width as $\mathcal{H}(\mathcal{K}, \epsilon) \leqslant w(\mathcal{K})^2/\epsilon^2$.

However, in specific cases this last inequality is too loose with respect to ϵ. As summarized in [42], this is the case of the *structured sets* $\mathcal{K}$ defined hereafter, for which this work will provide separated and tighter results.

Definition 1 (Structured sets[3] [42])**.** *A bounded set $\mathcal{K} \subset \mathbb{R}^N$ with diameter $d = \|\mathcal{K}\| := \max\{\|\boldsymbol{u}\| : \boldsymbol{u} \in \mathcal{K}\} < \infty$ is* structured *iff there exists a quantity $\bar{w}(\mathcal{K})$, independent of d, for which we have both*

$$\mathcal{H}(\mathcal{K}, \epsilon) \leqslant \bar{w}(\mathcal{K})^2 \log(1 + \tfrac{d}{\epsilon}), \tag{3a}$$

$$w(d^{-1}\,\mathcal{K}_{\epsilon d})^2 = w\big((d^{-1}\mathcal{K} - d^{-1}\mathcal{K}) \cap \epsilon\mathbb{B}^n\big)^2 \leqslant \epsilon^2\, \bar{w}(\mathcal{K})^2, \tag{3b}$$

for any $\epsilon > 0$, where $\mathcal{K}_{\epsilon'} := (\mathcal{K} - \mathcal{K}) \cap \epsilon'\mathbb{B}^n$ is the local set *of $\mathcal{K}$ of radius $\epsilon' > 0$.*

For instance, if $\mathcal{K}'$ is a subspace of $\mathbb{R}^N$, a union of subspaces (such as the set $\Sigma_K^{\boldsymbol{\Psi}}$ of K-sparse signals in an orthonormal basis or in a redundant dictionary $\boldsymbol{\Psi}$ of $\mathbb{R}^N$), the set of rank-r matrices $\mathcal{M}_r$ in $\mathbb{R}^{N_1\times N_2}$, or even the set of group-sparse signals, then $\mathcal{K}'$ is a *cone*, *i.e.*, $\lambda\mathcal{K}' \subset \mathcal{K}'$ for any $\lambda > 0$, and the set $\mathcal{K} := \mathcal{K}' \cap d\,\mathbb{B}^N$ is structured for any diameter $d > 0$ [42].

[3] Notice that in [42] $\mathcal{K}$ is assumed to be a subset of the sphere $\mathbb{S}^{N-1}$ so that $d = 1$. However, this slight difference does not change the bound on the Kolmogorov entropy or the Gaussian mean width of the structured sets considered in [42] and in this paper.

	Names	Properties
(P1)	Definition	$w(\mathcal{A}) = \mathbb{E}\sup_{\boldsymbol{x}\in\mathcal{A}} \lvert\langle \boldsymbol{g}, \boldsymbol{x}\rangle\rvert$ for $\boldsymbol{g} \sim \mathcal{N}^N(0,1)$.
(P2)	Homogeneity [13, Sec. 3.2]	$w(\lambda A) = \lambda w(\mathcal{A})$ for $\lambda > 0$.
(P3)	Set inclusion [13, Sec. 3.2]	if $\mathcal{A} \subset \mathcal{B}$, $w(\mathcal{A}) \leqslant w(\mathcal{B})$.
(P4)	Set difference [44, Sec. 5.3]	$w(\mathcal{A} - \mathcal{A}) \leqslant 2w(\mathcal{A})$.
(P5)	Modularity [13, Sec. 3.2]	$w(\mathcal{A}\cup\mathcal{B}) + w(\mathcal{A}\cap\mathcal{B}) = w(\mathcal{A}) + w(\mathcal{B})$, if $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{A}\cup\mathcal{B}$ are convex.
(P6)	Convex hull [13, Sec. 3.2]	$w(\mathrm{conv}(\mathcal{A})) = w(\mathcal{A})$.
(P7)	Subspace [13, Sec. 3.2]	if $\mathcal{A}_K$ is a K-dimensional subspace of $\mathbb{R}^N$, then $w(\mathcal{A}_K \cap \mathbb{S}^{N-1}) = w(\mathcal{A}_K \cap \mathbb{B}^N) \leqslant \sqrt{K}$.
(P8)	Subspace addition [13, Eq. (15)]	$w((\mathcal{A}_K \oplus \mathcal{B}) \cap \mathbb{S}^{N-1})^2 \leqslant K + w(\mathcal{B} \cap \mathbb{S}^{N-1})^2$.
(P9)	Link with diameter*	for $\lVert\mathcal{A}\rVert := \sup_{\boldsymbol{u}\in\mathcal{A}} \lVert\boldsymbol{u}\rVert$, $(\frac{2}{\pi})^{1/2}\, \lVert\mathcal{A}\rVert \leqslant w(\mathcal{A}) \leqslant \sqrt{N}\, \lVert\mathcal{A}\rVert$.
(P10)	Symmetrization*	$w(\mathcal{A}) - (\frac{2}{\pi})^{1/2} \inf_{\boldsymbol{u}\in\mathcal{A}} \lVert\boldsymbol{u}\rVert \leqslant \mathbb{E}\sup_{\boldsymbol{x}\in\mathcal{A}-\mathcal{A}} \lvert\langle\boldsymbol{g},\boldsymbol{x}\rangle\rvert \leqslant 2w(\mathcal{A})$.
(P11)	Translation*	$w(\mathcal{A}) - (\frac{2}{\pi})^{1/2} \lVert\boldsymbol{t}\rVert \leqslant w(\mathcal{A} + \{\boldsymbol{t}\}) \leqslant w(\mathcal{A}) + (\frac{2}{\pi})^{1/2} \lVert\boldsymbol{t}\rVert$, for $\boldsymbol{t} \in \mathbb{R}^N$.
(P12)	Invariance under $\mathcal{O}_N$ [45, Prop. 2.1]	For all $\boldsymbol{B} \in \mathcal{O}_N := \{\boldsymbol{C} \in \mathbb{R}^{N\times N} : \boldsymbol{C}\boldsymbol{C}^\top = \boldsymbol{C}^\top\boldsymbol{C} = \mathbb{1}_N\}$, $w(\boldsymbol{B}\,\mathcal{A}) = w(\mathcal{A})$.
(P13)	Translation on origin (from (P9) & (P11))	$w(\mathcal{A}) - (\frac{2}{\pi})^{1/2}\lVert\boldsymbol{x}_0\rVert \leqslant w(\mathcal{A} - \{\boldsymbol{x}_0\}) \leqslant 2w(\mathcal{A})$ for $\boldsymbol{x}_0 \in \mathcal{A}$ with $\lVert\boldsymbol{x}_0\rVert \geqslant \inf_{\boldsymbol{u}\in\mathcal{A}} \lVert\boldsymbol{u}\rVert$.
(P14)	Sudakov inequality [44, Sec. 1.7]	For an ϵ-net $\mathcal{G}_\epsilon \subset \mathcal{A}$, $\log\lvert\mathcal{G}_\epsilon\rvert \lesssim \epsilon^{-2}\, w(\mathcal{A})^2$.

	Special sets	Widths
(P15)	Finite [44, Sec. 1.4]	$w(\mathcal{S})^2 \lesssim \log\lvert\mathcal{S}\rvert$.
(P16)	Sphere and ball [44, Sec. 1.4]	$w(\mathbb{S}^{N-1}) \leqslant \sqrt{N}$ and $w(\mathbb{B}^N) \leqslant \sqrt{N}$.
(P17)	Sparse signals [44, Sec. 1.3]	For $\Sigma_K := \{\boldsymbol{u} : \lVert\boldsymbol{u}\rVert_0 \leqslant \sqrt{K}\}$, $w(\Sigma_K \cap \mathbb{B}^N)^2 \lesssim K\log(2N/K)$.
(P18)	"Compressible signals" [44, Sec. 1.3]	For $\mathcal{K}_{N,K} := \{\boldsymbol{u} : \lVert\boldsymbol{u}\rVert_1 \leqslant \sqrt{K}, \lVert\boldsymbol{u}\rVert \leqslant 1\}$, $w(\mathcal{K}_{N,K})^2 \lesssim K\log(2N/K)$.
(P19)	Low-rank matrices [32, Lemma 21]	For $\mathcal{M}_r := \{\boldsymbol{U} \in \mathbb{R}^{N_1\times N_1} : \mathrm{rank}(\boldsymbol{U}) \leqslant r\}$, $w(\mathcal{M}_r)^2 \lesssim r\,(N_1 + N_2)$.

Table 1: Useful properties of the Gaussian mean width. If not otherwise noted, all sets are subsets of $\mathbb{R}^N$. $*$: (P9) is obtained by a simple use of the Jensen and Cauchy-Schwartz inequalities, (P11) is a simple consequence of the triangular inequality and of $\mathbb{E}\lvert\langle\boldsymbol{g},\boldsymbol{t}\rangle\rvert = (\frac{2}{\pi})^{1/2}\, \lVert\boldsymbol{t}\rVert$.

Indeed, focusing first on (3b), if $\mathcal{K}'$ is one of the sets listed above, $\mathcal{K}'' := \mathcal{K}' - \mathcal{K}'$ is also a cone and $\mathcal{K}'' \supset d^{-1}(\mathcal{K} - \mathcal{K})$. Therefore $w(d^{-1}\mathcal{K}_{\epsilon d})^2 \leqslant w(\mathcal{K}'' \cap \epsilon\mathbb{B}^N)^2 = \epsilon^2 w(\mathcal{K}'' \cap \mathbb{B}^N)^2$. This last quantity is easily bounded since $\mathcal{K}''$ often shares the same structure than $\mathcal{K}'$, *e.g.*, $\mathcal{K}'' = \Sigma^{\boldsymbol{\Psi}}_{2K}$ if $\mathcal{K}' = \Sigma^{\boldsymbol{\Psi}}_K$, and in fact $w(\mathcal{K}'' \cap \mathbb{B}^N) \simeq w(\mathcal{K}/\lVert\mathcal{K}\rVert)$ showing that $\bar{w}(\mathcal{K})$ can be set to $w(\mathcal{K}/\lVert\mathcal{K}\rVert)$ in (3b).

Second, for (3a), the Komogorov entropy of such a set $\mathcal{K}'$ can often be tightly bounded by decomposing it into a union of subspaces or subdomains restricted to $d\,\mathbb{B}^N$, so that a global ϵ-net of small cardinality could be reached by the union of the ϵ-nets of all of these subparts [6, 42, 43], *i.e.*, justifying the bound $\mathcal{H}(\mathcal{K}' \cap d\mathbb{B}^N) \leqslant \bar{w}(\mathcal{K})^2 \log(1 + \frac{d}{\epsilon})$. Actually, concerning (3a), it occurs that for all the structured sets listed above we have that either $\bar{w}(\mathcal{K})^2 \simeq w(\mathcal{K}/\lVert\mathcal{K}\rVert)^2$ or both $\bar{w}(\mathcal{K})^2$ and $w(\mathcal{K}/\lVert\mathcal{K}\rVert)^2$ have the same simplified closed-form upper bound, *e.g.*, they are both upper bounded by $K\log(N/K)$ when $\mathcal{K}' = \Sigma^{\boldsymbol{\Psi}}_K$.

Thus, due to the observations made above, we will consider that $\bar{w}(\mathcal{K})$ can be bounded similarly to the actual Gaussian mean width $w(\lVert\mathcal{K}\rVert^{-1}\,\mathcal{K})$ of the normalized set $\lVert\mathcal{K}\rVert^{-1}\,\mathcal{K}$, *i.e.*, with the same simplified upper bound. An example of this fact for the set of bounded K-sparse vectors is provided at the end of Sec. 6.

Paper organization: The rest of the paper is structured as follows. In Sec. 2, we define the construction of our quantized sub-Gaussian random mapping. Additionally, this section characterizes the sub-Gaussianity of its linear ingredient, *i.e.*, its random projection matrix, and its interplay with the "anti-sparse" nature of the mapped vectors. We also formalize and motivate the main objectives of the paper, *e.g.*, explaining the shape and the origins of the targeted quasi-isometric embedding with its two specific distortions. Sec. 3 provides the main results of this work, namely, *(i)* the possibility to create with high probability a quasi-isometric sub-Gaussian embedding from our quantized mapping (Prop. 1), and *(ii)* a study of this mapping's *consistency width* behavior (Prop. 2). Sec. 4 discusses those two propositions, analyzing them in a few specific settings in comparison with related works in the fields of dimensionality reduction and 1-bit compressed sensing. Sec. 5 questions the necessity of dithering in the mapping $\boldsymbol{A}$ and shows that, from an appropriate counterexample, our results do not hold in full generality without such a dither. Finally, Sec. 6 and Sec. 7 contain the proofs of Prop. 1 and Prop. 2, respectively, the auxiliary Lemmas being demonstrated in appendix.

Conventions: We find useful to summarize here our mathematical notations. Domain dimensions are denoted by capital roman letters, *e.g.*, $M, N, \ldots$ Vectors and matrices are associated to bold symbols, *e.g.*, $\boldsymbol{\Phi} \in \mathbb{R}^{M\times N}$ or $\boldsymbol{u} \in \mathbb{R}^M$, while lowercase light letters are associated to scalar values. The identity matrix in $\mathbb{R}^D$ reads $\mathbb{1}_D$ while $\mathbb{I}[A] \in \{0,1\}$ is the indicator function of a set $A \subset \mathbb{R}^D$. An "event" is a set whose definition depends on the realization of some random variables, *e.g.*, if $X \in \mathbb{R}$ is a random variable, the event $A = \{X \leqslant 0\}$ has probability $\mathbb{P}(X \leqslant 0) = \mathbb{E}\,\mathbb{I}[A]$. The i^{th} component of a vector (or of a vector function) $\boldsymbol{u}$ reads either u_i or $(\boldsymbol{u})_i$, and the vector $\boldsymbol{u}_i$ may refer to the i^{th} element of a set of vectors. The set of indices in $\mathbb{R}^D$ is $[D] = \{1, \cdots, D\}$. The cardinality of a finite set $\mathcal{J}$ reads $|\mathcal{J}|$. For any $p \geqslant 1$, the ℓ_p-norm of $\boldsymbol{u}$ is $\|\boldsymbol{u}\|_p^p = \sum_i |u_i|^p$ with $\|\cdot\| := \|\cdot\|_2$. The "ℓ_0-norm" of a vector $\boldsymbol{u} \in \mathbb{R}^N$ is $\|\boldsymbol{u}\|_0 = |\operatorname{supp} \boldsymbol{u}|$, with $\operatorname{supp} \boldsymbol{u} = \{i : u_i \neq 0\}$ the support of $\boldsymbol{u}$. The $(N-1)$-sphere in $\mathbb{R}^N$ is $\mathbb{S}^{N-1} = \{\boldsymbol{x} \in \mathbb{R}^N : \|\boldsymbol{x}\| = 1\}$ while the unit ball is denoted $\mathbb{B}^N = \{\boldsymbol{x} \in \mathbb{R}^N : \|\boldsymbol{x}\| \leqslant 1\}$. The diameter of a bounded set $\mathcal{A} \subset \mathbb{R}^N$ is written $\|\mathcal{A}\| = \sup\{\|\boldsymbol{u}\| : \boldsymbol{u} \in \mathcal{A}\}$. The set of K-sparse signals in $\mathbb{R}^N$ is defined as $\Sigma_K := \{\boldsymbol{u} \in \mathbb{R}^N : \|\boldsymbol{u}\|_0 \leqslant K\}$ while the set of K-sparse signals in an orthonormal basis (ONB) $\boldsymbol{\Psi} \in \mathbb{R}^{N\times N}$, *i.e.*, with $\boldsymbol{\Psi}\boldsymbol{\Psi}^\top = \boldsymbol{\Psi}^\top\boldsymbol{\Psi} = \mathbb{1}_N$, reads $\Sigma_K^{\boldsymbol{\Psi}} = \boldsymbol{\Psi}\Sigma_K$. The positive thresholding function is defined by $(\lambda)_+ := \frac{1}{2}(\lambda + |\lambda|)$ for any $\lambda \in \mathbb{R}$. For $t \in \mathbb{R}$, $\lfloor t \rfloor$ (resp. $\lceil t \rceil$) is the largest (smallest) integer smaller (greater) than t. A random matrix $\boldsymbol{\Phi} \sim \mathcal{P}^{M\times N}(\Theta)$ is a $M \times N$ matrix with entries distributed as $\Phi_{ij} \sim_{\text{i.i.d.}} \mathcal{P}(\Theta)$ given the distribution parameters Θ of $\mathcal{P}$ (*e.g.*, $\mathcal{N}^{M\times N}(0,1)$ or $\mathcal{U}^{M\times N}([0,1])$). A random vector in $\mathbb{R}^M$ following $\mathcal{P}(\Theta)$ is defined by $\boldsymbol{v} \sim \mathcal{P}^M(\Theta)$. Given two random variables X and Y, the notation $X \sim Y$ means that X and Y have the same distribution. Since our developments do not focus on sharp bounds, we denote by C, c, c' or c'' (possibly large) constants whose value can change between lines. In a few places, for simplicity, we write $f \lesssim g$ if there exists a constant $c > 0$ such that $f \leqslant c\,g$, and correspondingly for $f \gtrsim g$. Moreover, $f \simeq g$ means that $f \lesssim g$ and $g \lesssim f$. Finally, for asymptotic relations, we use the common Landau family of notations, *i.e.*, the symbols O, Ω and Θ [34].

2 Quantized Sub-Gaussian Random Mapping

In this work, given a quantization resolution $\delta > 0$, we focus on the interaction between a random projection of $\mathbb{R}^N$ into $\mathbb{R}^M$ and the following uniform (dithered) quantizer[4] $\mathcal{Q}(t) = \delta\lfloor \frac{t}{\delta} \rfloor \in \delta\mathbb{Z}$, applied componentwise on vectors in $\mathbb{R}^M$. In other words, for some random matrix $\boldsymbol{\Phi} \in \mathbb{R}^{M\times N}$ whose distribution is specified below, we study the properties of the mapping $\boldsymbol{A} : \mathbb{R}^N \to \delta\mathbb{Z}^M$ with

$$\boldsymbol{A}(\boldsymbol{x}) := \mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{x} + \boldsymbol{\xi}), \tag{4}$$

[4] Hereafter, our developments could be adapted to any quantizer defined as $\mathcal{Q}'(t) := \delta(\lfloor \frac{t+q_0}{\delta} \rfloor + r_0) \in \delta\mathbb{Z}$, for some $q_0 \in [0,\delta)$ and $r_0 \in [0,1)$, *e.g.*, for the quantizer mentioned in the Introduction with $r_0 = 0$ and $q_0 = \delta/2$.

where $\boldsymbol{\xi} \in \mathcal{U}^M([0,\delta])$ is a uniform *dithering* that stabilizes the action of $\mathcal{Q}$ [9, 24, 27].

We specialize the mapping (4) on projection (or sensing) matrices $\boldsymbol{\Phi}$ with entries independently and identically drawn from a symmetric *sub-Gaussian* distribution. We recall that a random variable (r.v.) X is sub-Gaussian if its *sub-Gaussian norm* (or ψ_2-norm) [52]

$$\|X\|_{\psi_2} \;:=\; \sup_{p\geqslant 1}\; p^{-1/2}(\mathbb{E}|X|^p)^{1/p}. \tag{5}$$

is finite[5]. Examples of sub-Gaussian r.v.'s are Gaussian, Bernoulli, uniform or bounded r.v.'s, as

$$\|X\|_{\psi_2} \;\leqslant\; \|X\|_\infty := \inf\{t \geqslant 0 : \mathbb{P}(|X| \leqslant t) = 1\}.$$

Sub-Gaussian r.v.'s are endowed with several interesting properties described, *e.g.*, in [52]. Their tail is for instance bounded as the one of a Gaussian r.v., *i.e.*, there exists a $c > 0$ such that for all $\epsilon \geqslant 0$ and for a sub-Gaussian r.v. X,

$$\mathbb{P}(|X| > \epsilon) \;\lesssim\; e^{-c\,\epsilon^2/\|X\|^2_{\psi_2}}. \tag{6}$$

Moreover, since $\|X - \mathbb{E}X\|_{\psi_2} \leqslant \|X\|_{\psi_2} + \|\mathbb{E}X\|_{\psi_2} = \|X\|_{\psi_2} + |EX| \leqslant \|X\|_{\psi_2} + E|X| \leqslant 2\|X\|_{\psi_2}$, centering X has no effect on its sub-Gaussianity.

By a slight abuse of notation, we denote collectively the distributions of *symmetric* sub-Gaussian r.v. with zero expectation, unit variance and finite sub-Gaussian norm α by $\mathcal{N}_{\mathrm{sg},\alpha}(0,1)$, with $\alpha \geqslant 1/\sqrt{2}$ from (5). This means that if $X \sim \mathcal{N}_{\mathrm{sg},\alpha}(0,1)$, we do not fully specify the pdf of X but we know that X is centered, has unit variance and sub-Gaussian norm α.

In this context, for a sub-Gaussian random matrix $\boldsymbol{\Phi} = (\boldsymbol{\varphi}_1, \cdots, \boldsymbol{\varphi}_M)^\top \sim \mathcal{N}^{M\times N}_{\mathrm{sg},\alpha}(0,1)$, each row $\boldsymbol{\varphi}_i$ is also *isotropic*, *i.e.*, for all $i \in [M]$ and all $\boldsymbol{u} \in \mathbb{R}^N$,

$$\mathbb{E}|\langle \boldsymbol{\varphi}_i, \boldsymbol{u}\rangle|^2 = \|\boldsymbol{u}\|^2.$$

However, conversely to the Gaussian case where $\mathbb{E}|\langle \boldsymbol{g}, \boldsymbol{u}\rangle| = (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{u}\|$ for $\boldsymbol{g} \sim \mathcal{N}^N(0,1)$ and $\boldsymbol{u} \in \mathbb{R}^N$ (since $\langle \boldsymbol{g}, \boldsymbol{u}\rangle \sim \mathcal{N}(0, \|\boldsymbol{u}\|^2)$), we do not necessarily have $\mathbb{E}|\langle \boldsymbol{\varphi}, \boldsymbol{u}\rangle| = c\|\boldsymbol{u}\|$ for $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$ and some absolute constant $c > 0$.

As will be clear below, we must anyway determine the deviations to this last equality. Interestingly, as noted in [2], any sub-Gaussian random vector $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$ satisfies

$$\int_0^{+\infty} \big|\mathbb{P}(|\langle \boldsymbol{\varphi}, \boldsymbol{u}\rangle| \geqslant t) \;-\; \mathbb{P}(|\langle \boldsymbol{g}, \boldsymbol{u}\rangle| \geqslant t)\big|\,\mathrm{d}t \;\leqslant\; \kappa_{\rm sg}\,\|\boldsymbol{u}\|_\infty, \quad \forall \boldsymbol{u} \in \mathbb{R}^N, \tag{7}$$

for some constant $\kappa_{\rm sg} \geqslant 0$ depending only the distribution of $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$. While we have obviously $\kappa_{\rm sg} = 0$ if $\boldsymbol{\varphi} \sim \mathcal{N}^N(0,1)$, it is possible to bound this constant in full generality. Indeed, up to a simple change of variable $t \to t\|\boldsymbol{u}\|$ in the integral, (7) is sustained by the Berry-Esseen central limit theorem (as described in a simplified form in [2, Theorem 4.2]). This result shows basically that, for $\boldsymbol{u} \in \mathbb{S}^{N-1}$, the LHS of (7) is bounded by $9\,\mathbb{E}|\varphi|^3\,\|\boldsymbol{u}\|_3^3 \leqslant 9\sqrt{27}\,\alpha^3\|\boldsymbol{u}\|_\infty$ for $\varphi_i \sim_{\text{i.i.d.}} \varphi \sim \mathcal{N}_{\mathrm{sg},\alpha}(0,1)$. This means that $\kappa_{\rm sg} \leqslant 9\sqrt{27}\,\alpha^3$ for any $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$. Notice, however, that this bound can be loose for many sub-Gaussian distributions.

Thanks to assumption (7), we can establish the behavior of the *first absolute moment* function

$$\mu_{\rm sg}(\boldsymbol{u}) := \mathbb{E}|\langle \boldsymbol{\varphi}, \boldsymbol{u}\rangle|. \tag{8}$$

[5]Notice that other equivalent definitions for sub-Gaussian r.v. exist, see *e.g.*, [39].

Since $\mathbb{E}|X| = \int_0^\infty \mathbb{P}(|X| \geqslant t)\,\mathrm{d}t$ for any r.v. X and using Jensen's inequality, we indeed observe that

$$\mu_{\rm sg}(\boldsymbol{u}) \;\leqslant\; (\mathbb{E}|\langle \boldsymbol{\varphi}, \boldsymbol{u}\rangle|^2)^{1/2} = \|\boldsymbol{u}\|, \tag{9}$$

$$\big|\,\mu_{\rm sg}(\boldsymbol{u}) \;-\; (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{u}\|\,\big| \;\leqslant\; \kappa_{\rm sg}\,\|\boldsymbol{u}\|_\infty, \tag{10}$$

for all $\boldsymbol{u} \in \mathbb{R}^N$. The last property, which is also considered in 1-bit CS with non-Gaussian projections [2], is key for characterizing quantized embeddings from sub-Gaussian projections.

Having now fully described the elements composing our random quantized mapping $\boldsymbol{A}$, we formally address the objectives defined in the Introduction by observing "when", *i.e.*, under which conditions with respect to M, there exist two small distortions $\Delta_\oplus, \Delta_\otimes \geqslant 0$ such that the pseudo-distance $\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) := \frac{1}{M}\,\|\boldsymbol{A}(\boldsymbol{x}) - \boldsymbol{A}(\boldsymbol{y})\|_1$ is involved in the quasi-isometric relation

$$\big|\,\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) \;-\; (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|\,\big| \;\leqslant\; \Delta_\otimes\,\|\boldsymbol{x}-\boldsymbol{y}\| \;+\; \Delta_\oplus, \tag{11}$$

for all pair of vectors taken in a general subset $\mathcal{K} \subset \mathbb{R}^N$.

In particular, we aim to control the distortions $\Delta_\oplus$ and $\Delta_\otimes$ with respect to M, N, the non-Gaussian nature of $\boldsymbol{\Phi}$ (*i.e.*, through α and $\kappa_{\rm sg}$), the typical dimension of $\mathcal{K}$ (*i.e.*, its Gaussian mean width) and possible additional requirements on $\boldsymbol{x}$ and $\boldsymbol{y}$.

Let us justify and comment the specific form taken by (11). First, $\mathcal{D}$ is associated to a ℓ_1-distance in the image of $\boldsymbol{A}$. As detailed in Sec. 6, this choice establishes an equivalence between the evaluation of $\mathcal{D}$ and a specific counting procedure, *i.e.*, a count of the number of quantization *thresholds* separating each components of the randomly-projected vectors. However, it is not clear if our developments can be extended to a ℓ_2-based pseudo-distance, even if this holds, with additional distortion, in the case of Gaussian random projections and for finite sets $\mathcal{K}$ [27] (see Sec. 4).

Second, as explained in the Introduction, a special case where both non-zero $\Delta_\oplus$ and $\Delta_\otimes$ appear specifies the constant $(\frac{2}{\pi})^{1/2}$ in (11). When $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, [27] has proved a quantized version of the Johnson Lindenstrauss (JL) Lemma showing that for a finite set $\mathcal{S} \subset \mathbb{R}^N$ of size S, provided $M \gtrsim \epsilon^{-2}\log S$, one has

$$|\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\,\|\boldsymbol{x}-\boldsymbol{y}\|| \;\lesssim\; \epsilon\,\|\boldsymbol{x}-\boldsymbol{y}\| + \epsilon\delta,$$

for all pairs $\boldsymbol{x},\boldsymbol{y} \in \mathcal{S}$ with a probability at least $1 - e^{-\epsilon^2 M}$. As a direct impact of the loss of information induced by the quantization, we also observe here that $\boldsymbol{A}$ realizes a *quasi-isometric* mapping between $(\mathcal{S} \subset \mathbb{R}^N, \ell_2)$ and $(\boldsymbol{A}(\mathcal{S}) \subset \delta\mathbb{Z}^M, \ell_1)$ with $\Delta_\otimes = \epsilon$ and $\Delta_\oplus = \delta\epsilon$.

Finally, as will be clearly established in Sec. 3.1, the anti-sparse nature of $\boldsymbol{x}-\boldsymbol{y}$ must be involved in the characterization of the right-hand side of (11) in the case of a general sub-Gaussian matrix $\boldsymbol{\Phi}$. Indeed, let us consider a matrix with i.i.d. Bernoulli distributed random entries, *i.e.*, $\Phi_{ij} \sim_{\rm iid} \mathcal{B}(\frac{1}{2})$ with $\mathbb{P}(\Phi_{ij} = 1) = \mathbb{P}(\Phi_{ij} = -1) = 1/2$ for all $1 \leqslant i \leqslant M$ and $1 \leqslant j \leqslant N$, the vectors $\boldsymbol{x} = (1, 0, \cdots, 0)^\top \in \mathbb{R}^N$ and $\boldsymbol{y} = \boldsymbol{0} \in \mathbb{R}^N$ and assume $\boldsymbol{x},\boldsymbol{y} \in \mathcal{K}$, *e.g.*, with $\mathcal{K} = \Sigma_K \cap \mathbb{B}^N$ and $K \geqslant 1$. Then, taking $\delta = 1$, we clearly have $\boldsymbol{A}(\boldsymbol{x}) \in \{\pm 1\}^M$ and $\boldsymbol{A}(\boldsymbol{y}) = 0$, so that $\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) = 1$ and $\|\boldsymbol{x}-\boldsymbol{y}\| = 1$. Consequently, if (11) is expected to hold on any pair of vectors in $\mathcal{K}$, inserting $\boldsymbol{x}$ and $\boldsymbol{y}$ inside it gives $\Delta_\oplus + \Delta_\otimes \geqslant 1 - (\frac{2}{\pi})^{1/2} > 0.202$. This limits our hope to have $\Delta_\oplus + \Delta_\otimes$ as small as we want by, *e.g.*, increasing M.

In fact, between the two distortions, it is actually $\Delta_\otimes$ that should depend on the configuration of $\boldsymbol{x}-\boldsymbol{y}$. As proved in App. A,

$$\mathbb{E}|\lfloor x+\xi\rfloor - \lfloor y+\xi\rfloor| = |x-y|, \quad \forall x,y \in \mathbb{R},\ \xi \sim \mathcal{U}([0,1]). \tag{12}$$

Therefore, by definition of $\mathcal{Q}$, from the independence of each component of $\boldsymbol{A}$ and using the law of total expectation over $\boldsymbol{\xi}$ and $\boldsymbol{\Phi}$ we have

$$\mathbb{E}\,\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) = \mathbb{E}_{\boldsymbol{\varphi}}\mathbb{E}_{\xi}|\mathcal{Q}(\boldsymbol{\varphi}^\top\boldsymbol{x}+\xi) - Q(\boldsymbol{\varphi}^\top\boldsymbol{y}+\xi)| = \mathbb{E}_{\boldsymbol{\varphi}}|\boldsymbol{\varphi}^\top(\boldsymbol{x}-\boldsymbol{y})| = \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y}), \tag{13}$$

with $\boldsymbol{\varphi}\sim\mathcal{N}^N_{{\rm sg},\alpha}(0,1)$ and $\xi\sim\mathcal{U}([0,\delta])$. From the assumption (10) and given $K_0\in\mathbb{R}$, we then observe that

$$|\mathbb{E}\,\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\,\|\boldsymbol{x}-\boldsymbol{y}\|| \;=\; |\,\mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\,\|\boldsymbol{x}-\boldsymbol{y}\|| \;\leqslant\; \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}\,\|\boldsymbol{x}-\boldsymbol{y}\|, \tag{14}$$

for all vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ such that $\boldsymbol{x}-\boldsymbol{y}$ belongs to the set[6]

$$\rotatebox[origin=c]{180}{Σ}_{K_0} := \{\boldsymbol{u}\in\mathbb{R}^N : K_0\|\boldsymbol{u}\|^2_\infty \leqslant \|\boldsymbol{u}\|^2\}. \tag{15}$$

This last set amounts to considering vectors that are not "too sparse", *i.e.*, if $\boldsymbol{u}\in\rotatebox[origin=c]{180}{Σ}_{K_0}$ then $\|\boldsymbol{u}\|_0\geqslant K_0$, which determines our notation $\rotatebox[origin=c]{180}{Σ}_{K_0}$ as opposed to Σ_K. However, the converse is not true and $\rotatebox[origin=c]{180}{Σ}_{K_0}\neq\Sigma^{\mathsf{c}}_{\lfloor K_0\rfloor}$. Since belonging to $\rotatebox[origin=c]{180}{Σ}_{K_0}$ prevents sparsity, we say that a vector $\boldsymbol{u}\in\rotatebox[origin=c]{180}{Σ}_{K_0}$ is an *anti-sparse* vector of level $K_0\geqslant 0$.

Actually (14) states that, for vectors $\boldsymbol{x}-\boldsymbol{y}\in\rotatebox[origin=c]{180}{Σ}_{K_0}$, the expectation of $\mathcal{D}(\boldsymbol{x},\boldsymbol{y})$ is close to the one obtained with Gaussian random projections, *i.e.*, close to the expectation $(\frac{2}{\pi})^{1/2}\,\|\boldsymbol{x}-\boldsymbol{y}\|$ associated to $\kappa_{\rm sg}=0$. Thus, if we expect to show that, for all vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\mathcal{K}$, $\mathcal{D}(\boldsymbol{x},\boldsymbol{y})$ concentrates around $(\frac{2}{\pi})^{1/2}\,\|\boldsymbol{x}-\boldsymbol{y}\|$, we must take into account the anti-sparse nature of the difference $\boldsymbol{x}-\boldsymbol{y}$, *i.e.*, we would need enforcing this vector to belong to $\rotatebox[origin=c]{180}{Σ}_{K_0}$ for a sufficiently large K_0.

Combining these three observations, and anticipating over the next section, we can now refine the meaning of (11). We are actually going to show that, if M is bigger than some M_0 growing with the typical dimension of $\mathcal{K}$ and decreasing with ϵ (see Sec. 3), then, with high probability,

$$((\tfrac{2}{\pi})^{1/2} - \epsilon - \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}})\,\|\boldsymbol{x}-\boldsymbol{y}\| - c\epsilon\delta \;\leqslant\; \mathcal{D}(\boldsymbol{x},\boldsymbol{y}) \;\leqslant\; ((\tfrac{2}{\pi})^{1/2} + \epsilon + \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}})\,\|\boldsymbol{x}-\boldsymbol{y}\| + c\epsilon\delta,$$

for all $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}$ and $\boldsymbol{x}-\boldsymbol{y}\in\rotatebox[origin=c]{180}{Σ}_{K_0}$.

Remark: As will be cleared later, our developments benefit of the tools and techniques developed in [44] where it is shown that, for a 1-bit mapping $\boldsymbol{A}':\mathbb{R}^N\to\{\pm1\}^M$ such that $\boldsymbol{A}'(\boldsymbol{x}) = {\rm sign}\,(\boldsymbol{\Phi x})$ with a random Gaussian matrix $\boldsymbol{\Phi}\sim\mathcal{N}^{M\times N}(0,1)$, and for the normalized Hamming distance $\mathcal{D}'(\boldsymbol{x},\boldsymbol{y}) = M^{-1}\sum_i\mathbb{I}[A'_i(\boldsymbol{x})\neq A'_i(\boldsymbol{y})]$, one has, provided $M\gtrsim\epsilon^{-4}w(\mathcal{K})^2$ and with probability exceeding $1-e^{-\epsilon^2 M}$, that for all $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}$,

$$|\mathcal{D}'(\boldsymbol{x},\boldsymbol{y}) - \arccos(\tfrac{\boldsymbol{x}^\top\boldsymbol{y}}{\|\boldsymbol{x}\|\|\boldsymbol{y}\|})| \;\lesssim\; \epsilon.$$

Our extension to non-Gaussian sensing matrices is also inspired by similar developments realized in [2] for binary mappings and other generalized linear models.

3 Main Results

3.1 Quasi-Isometric Quantized Embedding

In regards to the context explained in the previous section, our first main result can be stated as follows.

[6] That could be pronounced "amgis".

Proposition 1 (Quantized sub-Gaussian quasi-isometric embedding)**.** *Given $\delta > 0$, $\epsilon \in (0,1)$, $K_0 > 0$, a bounded subset $\mathcal{K} \subset \mathbb{R}^N$ and a sub-Gaussian distribution $\mathcal{N}_{\rm sg,\alpha}$ respecting* (10) *for $0 \leqslant \kappa_{\rm sg} < \infty$, there exist some values $c, c' > 0$, only depending on α, such that, if*

$$M \ \gtrsim\ \tfrac{1}{\delta^2\epsilon^5} w(\mathcal{K})^2, \tag{16}$$

for a general set $\mathcal{K}$, or

$$M \ \gtrsim\ \tfrac{1}{\epsilon^2}\, \bar{w}(\mathcal{K})^2 \, \log(1 + \tfrac{\|\mathcal{K}\|}{\delta\sqrt{\epsilon^3}}), \tag{17}$$

for structured sets $\mathcal{K}$ (see Def. 1 for the definition of $\bar{w}$), such as the set of bounded K-sparse signals or the one of bounded rank-r matrices, then, for $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}_{\rm sg,\alpha}(0,1)$, a dithering $\boldsymbol{\xi} \sim \mathcal{U}^M([0,\delta])$ and the associated quantized mapping $\boldsymbol{u} \in \mathbb{R}^N \to \boldsymbol{A}(\boldsymbol{u}) = \mathcal{Q}(\boldsymbol{\Phi u} + \boldsymbol{\xi})$, we have with probability at least $1 - e^{-c'\epsilon^2 M}$ and for all pairs $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{K}$ with $\boldsymbol{x} - \boldsymbol{y} \in \mho_{K_0}$,

$$\big((\tfrac{2}{\pi})^{1/2} - \epsilon - \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}\big)\,\|\boldsymbol{x} - \boldsymbol{y}\| - c\epsilon\delta \ \leqslant\ \mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) \ \leqslant\ \big((\tfrac{2}{\pi})^{1/2} + \epsilon + \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}\big)\,\|\boldsymbol{x} - \boldsymbol{y}\| + c\epsilon\delta. \tag{18}$$

In the Gaussian case, i.e.*, for $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, the conditions remain the same and* (18) *is simplified with $\kappa_{\rm sg} = 0$,* i.e.*, there is no additional requirement on the anti-sparse nature of $\boldsymbol{x} - \boldsymbol{y}$ in* (18) *since K_0 can be set to 1 and $\mho_{K_0} = \mathbb{R}^N$.*

In Prop. 1, as shown in Sec. 2, the constant part $\kappa_{\rm sg}/\sqrt{K_0}$ of the multiplicative distortion appearing in both sides of (18) is unavoidable in the case of non-Gaussian projections (with $\kappa_{\rm sg} \neq 0$). Actually, we can show that this distortion cannot decay faster than $\Omega(1/K_0)$ for non-Gaussian (but sub-Gaussian) random matrices when the level of anti-sparsity K_0 of $\boldsymbol{x} - \boldsymbol{y}$ increases. To see this, it is sufficient to study $\mathcal{D}(\boldsymbol{x}, \boldsymbol{y})$ for an asymptotically large M, *i.e.*, $\mathbb{E}\mathcal{D}(\boldsymbol{x}, \boldsymbol{y})$ by the law of large numbers, and to observe how the relative error between $\mathbb{E}\mathcal{D}(\boldsymbol{x}, \boldsymbol{y})$ and $(\frac{2}{\pi})^{1/2}\|\boldsymbol{x} - \boldsymbol{y}\|$ behaves when that level K_0 increases.

Taking $\delta = 1$ by simplicity, notice first that, from the observation made in (12),

$$\mathbb{E}\mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) = \mathbb{E}_{\boldsymbol{\varphi}}|\boldsymbol{\varphi}^\top(\boldsymbol{x} - \boldsymbol{y})| = \mu_{\rm sg}(\boldsymbol{x} - \boldsymbol{y}),$$

where $\boldsymbol{\varphi} \sim \mathcal{B}(\frac{1}{2})^N$ and $\mu_{\rm sg}$ was introduced in (8).

Let us then take $\boldsymbol{x}$ and $\boldsymbol{y}$ such that the vector $\boldsymbol{w} := \boldsymbol{x} - \boldsymbol{y}$ is equal to 1 on its first K_0 components and zero elsewhere, *i.e.*, $\boldsymbol{w} \in \mho_{K_0}$. In this case and if $\boldsymbol{\Phi}$ is a random Bernoulli matrix, $\mu_{\rm sg}(\boldsymbol{w})$ is actually twice the *mean absolute deviation* (MAD) of a Binomial distribution $\mathrm{Bin}(K_0, \frac{1}{2})$ with K_0 degrees of freedom and success probability $p = 1/2$ since

$$\mu_{\rm sg}(\boldsymbol{w}) = \mathbb{E}|\textstyle\sum_{j=1}^{K_0} \varphi_j| = 2\mathbb{E}|(\sum_{j=1}^{K_0} X_j) - \frac{1}{2}K_0| = 2\mathbb{E}|\beta_{K_0} - \mathbb{E}\beta_{K_0}|,$$

with, for $1 \leqslant j \leqslant K_0$ and $X_j := \frac{1}{2}(\varphi_j + 1) \sim_{\rm iid} \mathcal{B}(\{0,1\}, 1/2)$ a Bernoulli random variable such that $\mathbb{P}(X_j = 0) = 1/2$, and $\beta_{K_0} \sim \mathrm{Bin}(K_0, \frac{1}{2})$.

However, from [8, 30, 53] we can show that (see App. G for details)

$$|\mathbb{E}|\beta_{K_0} - \mathbb{E}\beta_{K_0}| - (\tfrac{2}{\pi})^{1/2}\tfrac{\sqrt{K_0}}{2}| \geqslant C\tfrac{\sqrt{K_0}}{2}K_0^{-1},$$

for $C = 1/7$. Consequently, for our choice of $\boldsymbol{w} = \boldsymbol{x} - \boldsymbol{y}$ such that $\|\boldsymbol{w}\| = \sqrt{K_0}$, this shows that

$$\big|\mathbb{E}\mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x} - \boldsymbol{y}\|\big| \ \geqslant\ 2C\,\|\boldsymbol{x} - \boldsymbol{y}\|\,K_0^{-1},$$

and proves that, even if we reached an asymptotic regime in M, a multiplicative distortion between $\mathcal{D}(\boldsymbol{x}, \boldsymbol{y})$ and $(\frac{2}{\pi})^{1/2}\|\boldsymbol{x} - \boldsymbol{y}\|$ would remain, and this one could decay faster than $1/K_0$ when K_0 increases. It is therefore unclear if our decay in $1/\sqrt{K_0}$ is optimal.

To conclude this section, let us observe that Prop. 1 improves a proof of existence of a quantized embedding given in [44, Theorem 1.10] where it was showed that, provided $M \gtrsim \epsilon^{-12} w(\mathcal{K}-\mathcal{K})^2$, there exists an arrangement of M affine hyperplanes in $\mathbb{R}^N$ and a scaling factor λ such that

$$|\lambda \mathcal{D}_{\rm c}(\boldsymbol{x},\boldsymbol{y}) - \|\boldsymbol{x}-\boldsymbol{y}\|| \leqslant \epsilon,$$

where $\mathcal{D}_{\rm c}$ denotes the fraction of affine hyperplanes that separate the two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$.

For reasons explained in Sec. 6, each element $\delta^{-1}|A_i(\boldsymbol{x})-A_i(\boldsymbol{y})|_1$ appearing in $\delta^{-1}\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) = \frac{1}{\delta M}\sum_{i=1}^M |A_i(\boldsymbol{x})-A_i(\boldsymbol{y})|_1$ actually counts the number of parallel affine hyperplanes in $\mathbb{R}^N$ normal to $\boldsymbol{\varphi}_i$ and far apart by δ, with a dithering that randomly displaces the origin. Therefore, Prop. 1 basically constructs, in a random fashion, an arrangement of M such parallel hyperplane bundle, *i.e.*, in M different directions $\{\boldsymbol{\varphi}_i/\|\boldsymbol{\varphi}_i\|, i \in [M]\}$. Considering a Gaussian matrix $\boldsymbol{\Phi}$ (with $\kappa_{\rm sg}=0$), we have therefore proved that there with a minimal M that grows like ϵ^{-5} rather than ϵ^{-12} when ϵ decays (as expressed in (16)). This is even reduced to ϵ^{-2} for pairs of vectors taken in a structured set.

3.2 Consistency Width Decay

As a second important result, we optimize the decay law (as M increases) of the distance of any pair of vectors $\boldsymbol{x},\boldsymbol{y} \in \mathcal{K}$ whose difference is "not too sparse" when those are mapped by $\boldsymbol{A}$ on the same quantization point in $\delta\mathbb{Z}^M$, *i.e.*, when they are *consistent*. We refer to this distance as the *consistency width* of $\boldsymbol{A}$.

This width could be characterized from Prop. 1 when $\mathcal{D}(\boldsymbol{x},\boldsymbol{y})=0$, which provides $\|\boldsymbol{x}-\boldsymbol{y}\| \lesssim \epsilon \simeq M^{-1/5}$ (or $M^{-1/2}$ if $\mathcal{K}$ is a structured set) for large M respecting (16) (resp. (17)), δ fixed and $\kappa_{\rm sg}/\sqrt{K_0}$ small. However, focusing on the conditions guaranteeing the consistency of $\boldsymbol{x}$ and $\boldsymbol{y}$, and considering all quantities fixed but M, our result below reaches the improved decay $\epsilon = O(M^{-1/4})$ for a general set $\mathcal{K}$ and $\epsilon = O(1/M)$ for a structured one. We prove the following proposition in Sec. 7.

Proposition 2 (Consistency width upper bound)**.** *Let us take a quantization resolution $\delta > 0$, an accuracy $\epsilon \in (0,1)$, a sub-Gaussian distribution $\mathcal{N}_{{\rm sg},\alpha}(0,1)$ respecting* (10) *for $0 \leqslant \kappa_{\rm sg} < \infty$, $K_0 > 0$ such that $\sqrt{K_0} \geqslant 16\kappa_{\rm sg}$ and a bounded subset $\mathcal{K} \subset \mathbb{B}^N$ of $\mathbb{R}^N$. For a value $c > 0$ depending only on α, provided*

$$M \gtrsim \tfrac{(2+\delta)^4}{\delta^2\epsilon^4}\, w(\mathcal{K})^2 \tag{19}$$

for a general set $\mathcal{K}$, or

$$M \gtrsim \tfrac{2+\delta}{\epsilon}\, \bar{w}(\mathcal{K})^2 \log\big(1 + \tfrac{(2+\delta)^{3/2}\|\mathcal{K}\|}{\delta\epsilon^{3/2}}\big), \tag{20}$$

for a structured set $\mathcal{K}$, the map $\boldsymbol{A}$ defined in (4) *with $\boldsymbol{\Phi} \sim \mathcal{N}_{{\rm sg},\alpha}^{M\times N}(0,1)$ and $\boldsymbol{\xi} \sim \mathcal{U}^M([0,\delta])$ is such that, with probability exceeding $1-2\exp(-c\epsilon M/(1+\delta))$,*

$$\boldsymbol{A}(\boldsymbol{x}) = \boldsymbol{A}(\boldsymbol{y}) \quad \Rightarrow \quad \|\boldsymbol{x}-\boldsymbol{y}\| \leqslant \epsilon, \tag{21}$$

for all $\boldsymbol{x},\boldsymbol{y} \in \mathcal{K}$ with $\boldsymbol{x}-\boldsymbol{y} \in \mho_{K_0}$. In the Gaussian case, i.e.*, for $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}(0,1)$, the conditions above remain the same with $\kappa_{\rm sg} = 0$,* i.e.*, with no additional requirement on the anti-sparse nature of $\boldsymbol{x}-\boldsymbol{y}$ in* (21)*.*

Unfortunately, we were unable to produce a convincing counter example of a pair of vectors both with difference not in $\mho_{K_0}$ and failing to meet (21) under the conditions of Prop. 2. Therefore, it is not clear if the condition $\boldsymbol{x}-\boldsymbol{y} \in \mho_{K_0}$ is an artifact of the proof or if removing it could worsen then dependence in ϵ in (19).

4 Discussions and Perspectives

Before delving into the proofs of Prop. 1 and Prop. 2 (see Sec. 6 and Sec. 7, respectively), let us discuss their meaning and limitations, providing also some perspectives for future works.

On the impact of the diameter of structured sets: For the structured sets considered in the Introduction, it is known that if the linear embedding (1) holds with high probability for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{K} \subset \mathbb{S}^{N-1}$ with some distortion $\epsilon > 0$, then, since (1) is homogeneous, a simple rescaling argument proves that the same relation actually holds for all points in $\mathcal{K}' = \cup_{\lambda>0}\mathcal{K}$, or equivalently for all points in the cone $\mathcal{K}'$ if $\mathcal{K} = \mathcal{K}' \cup \mathbb{S}^{N-1}$ [6, 39]. In particular, since such a linear embedding occurs with high probability for sub-Gaussian random matrices provided $M \gtrsim \epsilon^{-2}\, w(\mathcal{K})^2$ [39], this requirement remains unchanged for reaching the embedding of vectors in $\mathcal{K}'$.

Obviously, in the case of a quantized embedding such as (18), the non-linear nature of $\mathcal{Q}$ prevents this rescaling argument from holding. However, an interesting phenomenon occurs anyway in this case through the requirements (17) and (20) of Prop. 1 and Prop. 2, respectively. Indeed, we see there that the diameter of the set $\mathcal{K}$ has only a logarithmic impact on the minimal value of M needed for these propositions to hold, since $\bar{w}$ does not depend on the diameter of $\mathcal{K}$ (see Def. 1 and the subsequent explanations). This really slow increase approaches the scale-invariant requirement obtained by linear embedding of structured sets, and is anyway strikingly slower than the quadratic amplification of the minimal number of measurements provided by (16) and (19) in the case of a general set $\mathcal{K}$, as involved by (P2) when $\mathcal{K}$ is expanded like $\mathcal{K} \to \lambda\mathcal{K}$ for $\lambda > 1$.

Mitigating the anti-sparsity requirement: For both propositions, we can be concerned by the restriction that the vector difference must be "not too sparse", *i.e.*, for $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{K}$ there must be a sufficiently big K_0, either for having $\boldsymbol{x} - \boldsymbol{y} \in \daleth_{K_0}$ and minimizing the distortion $\kappa_{\rm sg}/\sqrt{K_0}$ in (18), or for satisfying $\sqrt{K_0} \geqslant 16\kappa_{\rm sg}$ in Prop. 2. However, in certain cases, it is possible to adapt the sensing matrix as to increase this K_0.

Indeed, assuming without loss of generality that the vectors $\boldsymbol{x} - \boldsymbol{y} \in \mathcal{K} - \mathcal{K}$ are expected to be "too sparse" only in $\boldsymbol{\Psi} = \mathbb{1}$ when the sensing matrix is non-Gaussian (*i.e.*, $\kappa_{\rm sg} \neq 0$), we can always "rotate"[7] $\mathcal{K}$ with an ONB $\boldsymbol{\Psi}_0$ of $\mathbb{R}^N$ so that elements of $\mathcal{K}' - \mathcal{K}'$ with $\mathcal{K}' := \boldsymbol{\Psi}_0\mathcal{K}$ have a higher anti-sparse degree than those of $\mathcal{K} - \mathcal{K}$, *i.e.*,

$$\begin{aligned} \max\{K_0 : (\mathcal{K}' - \mathcal{K}') \cap \daleth_{K_0} \neq \emptyset\} &= \min_{\boldsymbol{u} \in \mathcal{K}-\mathcal{K}} \tfrac{\|\boldsymbol{u}\|^2}{\|\boldsymbol{\Psi}_0\boldsymbol{u}\|_\infty^2} \\ \geqslant \quad \min_{\boldsymbol{u} \in \mathcal{K}-\mathcal{K}} \tfrac{\|\boldsymbol{u}\|^2}{\|\boldsymbol{u}\|_\infty^2} &= \max\{K_0 : (\mathcal{K} - \mathcal{K}) \cap \daleth_{K_0} \neq \emptyset\} \end{aligned} \tag{22}$$

possibly trying to maximize the left hand side in the selection of $\boldsymbol{\Psi}_0$.

Therefore, while the requirements imposed on M in Prop. 1 and Prop. 2 are unchanged between $\mathcal{K}$ and $\mathcal{K}'$ in Prop. 1 (by the invariance (P12) of $w(\mathcal{K})$ in Table 1) and since $\|\boldsymbol{x}' - \boldsymbol{y}'\| = \|\boldsymbol{x} - \boldsymbol{y}\|$ for $\boldsymbol{x}' = \boldsymbol{\Psi}_0\boldsymbol{x}$ and $\boldsymbol{y}' = \boldsymbol{\Psi}_0\boldsymbol{y}$, "rotating" $\mathcal{K}$ with $\boldsymbol{\Psi}_0$ helps to lighten the condition imposed on $\boldsymbol{x} - \boldsymbol{y}$. Moreover, this rotation is of course equivalent to directly build a sensing matrix $\boldsymbol{\Phi}' = \boldsymbol{\Phi}\boldsymbol{\Psi}_0$ to quasi-isometrically embed the set $\mathcal{K}$ with the mapping $\boldsymbol{A}(\cdot) := \mathcal{Q}(\boldsymbol{\Phi}'\cdot)$. Actually, in the case where $\boldsymbol{\Psi} = \mathbb{1}$ as above, a good choice for $\boldsymbol{\Psi}_0$ is the DCT basis, *i.e.*, using the incoherence of those two bases that prevents a sparse signal to be sparse in the frequency domain, also taking advantage of the fast FFT-based matrix-vector multiplication offered by the DCT. Notice, however, that the procedure above cannot work if $\mathcal{K}$ is expected to generate

[7]Strictly speaking, while $|\det\boldsymbol{\Psi}_0| = 1$, $\boldsymbol{\Psi}_0 \in \mathcal{O}_N$ is a rotation only if its determinant is 1.

differences of vectors that are sparse in different bases, *e.g.*, a union of incoherent bases such as $\mathbb{1}$ and the DCT basis. In such a case, it could be hard to maximize the right-hand side of (22) over $\boldsymbol{\Psi}_0$.

Interestingly, a similar procedure to the one described above has been developed recently in [41, Theorem 2.3] in the context of fast circulant binary embeddings of finite sets of vectors. The requirement on the anti-sparse nature of the mapped vectors is there mitigated by taking $\boldsymbol{\Psi}_0$ as the product of a Hadamard transform with a diagonal matrix with random Rademacher entries, which can provably reduce the *coherence* $\|\boldsymbol{\Psi}_0\boldsymbol{u}\|_\infty^2/\|\boldsymbol{u}\|$ of too sparse $\boldsymbol{u}$ with high probability.

Intrinsic "anti-sparse" distortion limit: We can notice that for non-Gaussian random measurements, the term $\kappa_{\rm sg}/\sqrt{K_0}$ in (18) is actually lower bounded. This is simply due to the relation $\|\boldsymbol{u}\|^2 \leqslant N\|\boldsymbol{u}\|_\infty^2$, which implies $K_0 \leqslant N$ whatever the properties of the vector $\boldsymbol{u} \in \mathcal{K} - \mathcal{K} \subset \mathbb{R}^N$. Consequently,

$$\tfrac{\kappa_{\rm sg}}{\sqrt{K_0}} \geqslant \tfrac{\kappa_{\rm sg}}{\sqrt{N}},$$

which limits our hope to tighten the multiplicative error of quantized non-Gaussian quasi-isometric embeddings, except if one considers asymptotic regimes where N can be considered as being much larger than $\kappa_{\rm sg}^2$.

Distortion regimes: As already noticed in [27], Prop. 1 allows us to distinguish different regimes of the quasi-isometric embedding. If $\delta \simeq 0$, the quantization operator tends to the identity function and (18) converges to a ℓ_2/ℓ_1 variant of the RIP generalized to any sets $\mathcal{K}$ and to sub-Gaussian random matrices, as characterized in [44, 49] for general sets and in [25] for sparse signal sets only. For $\delta \gg 2\|\mathcal{K}\|$ the embedding becomes purely quasi-isometric and, keeping the context defined in Prop. 1, (18) involves

$$(\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\| - c(\epsilon\delta + \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}) \;\leqslant\; \mathcal{D}(\boldsymbol{x},\boldsymbol{y}) \;\leqslant\; (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\| + c(\epsilon\delta + \tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}), \tag{23}$$

for some absolute constant $c > 0$. However, in this case, the quantization becomes essentially binary. In fact, it is exactly binary for random matrices whose entries are generated from a bounded symmetric sub-Gaussian distribution, *i.e.*, from $\varphi \sim \mathcal{N}_{{\rm sg},\alpha}(0,1)$ with $\|\varphi\|_\infty \leqslant F$ for some $F > 0$. In this case, since $\mathcal{K}$ is assumed bounded, for all $\boldsymbol{u} \in \mathcal{K}$, $|(\boldsymbol{\Phi}\boldsymbol{u})_i| \leqslant F\|\mathcal{K}\|$ and the components of $\boldsymbol{A}(\boldsymbol{u}) = \mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{u} + \boldsymbol{\xi})$ with $\boldsymbol{\xi} \sim \mathcal{U}^M([0,\delta])$ can only take two values, *e.g.*, $\{-1, 0\}$ if $0 \in \mathcal{K}$. Moreover, if φ is unbounded and $0 \in \mathcal{K}$, its sub-Gaussian nature is so that the fraction of quantized measurements that do not belong to $\{-1, 0\}$ can be made arbitrarily close to 0 when δ increases. In conclusion, similarly to [33], we have basically defined a one-bit quantized embedding that preserves the norm of the projected vectors, as opposed to the mapping $\boldsymbol{A}'(\cdot) = {\rm sign}\,(\boldsymbol{\Phi}\,\cdot)$ that loses this information [29, 46]. Notice there that the role of our dithering can be compared to the one of the threshold inserted in the sign quantization in [33]. Conversely to that work, however, we do not provide any algorithm to reconstruct a signal from its quantized mapping by $\boldsymbol{A}$.

Towards an ℓ_2/ℓ_2 quasi-isometric embedding? It is not clear if Prop. 1 could be turned into a quasi-isometric embedding between $(\mathcal{K} \subset \mathbb{R}^N, \ell_2)$ and $(\boldsymbol{A}(\mathcal{K}) \subset \delta\mathbb{Z}^M, \ell_2)$. As said earlier, for Gaussian random matrices and for finite sets $\mathcal{K}$, an approximate quasi-isometric embedding can be found by integrating a non-linear distortion of the ℓ_2-distance, *i.e.*, in (18) for $\kappa_{\rm sg} = 0$, $\|\boldsymbol{x}-\boldsymbol{y}\|$ is replaced by $g_\delta(\|\boldsymbol{x}-\boldsymbol{y}\|)$ for some non-decreasing function $g_\delta : \mathbb{R}_+ \to \mathbb{R}_+$. Interestingly, $|g_\delta(\lambda) - \lambda| = O(\sqrt{\delta\lambda})$ for $\lambda \gg \delta$ and $|g_\delta(\lambda) - (\sqrt{2}\lambda/\sqrt{\pi})^{1/2}| = O(\lambda)$ for $\lambda < \delta$, so that for small δ or large λ, $g_\delta(\lambda) \approx \lambda$. Therefore, as soon as $\|\boldsymbol{x}-\boldsymbol{y}\| \gg \delta$, we get approximately a ℓ_2/ℓ_2 quasi-isometric embedding. Knowing if this extends to any subset $\mathcal{K}$ and to sub-Gaussian random matrices is left for a future work.

Reconstructing low-complexity vectors from quantized compressive observation? Beyond the mere analysis of the quasi-isometric properties of our quantized mapping and closer to the context of quantized compressed sensing, this paper does not say anything on the reconstruction algorithms that could be developed for recovering a signal $\boldsymbol{x}$ from its observations $\boldsymbol{z} = \mathcal{Q}(\boldsymbol{\Phi x})$. A few algorithms exist for realizing this operation, some when δ is small compared to the expected dynamic of $\|\boldsymbol{\Phi x}\|$ [14, 25, 54], others in the 1-bit CS setting [4, 29, 45, 46]. However, for the first category, their stability (or convergence) does not rely on a quasi-isometric embedding property but rather on the restricted isometry property [11, 14, 37] or on variations involving other norms [25, 26]. In future research, it will be appealing to find a proof of the instance optimality of those algorithms, *e.g.*, for the basis pursuit dequantizer (BPDQ), using the quasi-isometry property promoted by Prop. 1, even if recent interesting results show that an optimal "non-RIP" proof can be developed for BPDQ [20].

Extension to fast and universal quantized embeddings? We conclude this section by mentioning that it would be useful to prove Prop. 1 for structured random matrices, *e.g.*, for random Fourier or random Hadamard ensembles [22], as recently obtained in [41] for the binary embedding of finite sets. This would lead to a fast computation of quantized mappings, with potential application in nearest-neighbor search for databases of high-dimensional signals. An open question is also the possibility to extend this work to universally-quantized embeddings [9, 10, 48], *i.e.*, taking a periodic quantizer $\mathcal{Q}$ in (4). This could potentially lead to quasi-isometric embeddings with (exponentially) decaying distortions on vectors sets with small Gaussian width and using sub-Gaussian random matrices.

5 On the necessity to dither the quantization

Considering the main results of this paper, namely Prop. 1 and Prop. 2, we could ask ourselves if a quantized mapping that would not include a dithering could also verify (18) and (21) under equivalent conditions on M and on the anti-sparse nature of $\boldsymbol{x} - \boldsymbol{y}$ for any vectors $\boldsymbol{x}, \boldsymbol{y}$ in $\mathcal{K}$.

The answer is, however, negative in full generality, *i.e.*, it is possible to define a quantized and undithered map $\boldsymbol{A} : \boldsymbol{x} \to \mathcal{Q}(\boldsymbol{\Phi x})$ for some appropriate quantizer resolution δ and sub-Gaussian random matrix $\boldsymbol{\Phi}$ that is *incompatible* with the definition of a quasi-isometric embedding with arbitrarily small additive distortion or with an arbitrarily small consistency width.

To see this, let us set $\delta = 1$, $\mathcal{Q}(\lambda) := \operatorname{argmin}_{\lambda' \in \mathbb{Z}} |\lambda - \lambda'| = \lfloor \lambda + \frac{1}{2} \rfloor$ (applied componentwise[8]), and take $\boldsymbol{\Phi}$ to be a Bernoulli random matrix, *i.e.*, $\boldsymbol{\Phi}_{ij} \in \{\pm 1\}$. Given the value $\kappa_{\rm sg} > 0$ associated to the distribution of $\boldsymbol{\Phi}$, we also set arbitrarily an integer K_0 such that $(\frac{2}{\pi})^{1/2} - (\kappa_{\rm sg}/\sqrt{K_0}) \geqslant 1/2$. In fact, we can compute that $\alpha = 1$ for a Bernoulli r.v., so that $\kappa_{\rm sg} \leqslant 9\sqrt{27} < 47$ from the bound given in Sec. 2. Therefore, $K_0 > (160)^2$ certainly works.

We then define two K_0-sparse vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^N$ with $\boldsymbol{u}$ equal to 1 on it first K_0 components and 0 elsewhere, and $\boldsymbol{v} := (1 + sK_0^{-1})\,\boldsymbol{u}$ for some fixed $0 < |s| < 1/2$. Clearly, when $K \geqslant K_0$ these two vectors belong to the structured set $\mathcal{K} := \Sigma_K \cap r_0 \mathbb{B}^N$ with $r_0 := \frac{3}{2}\sqrt{K_0}$. Moreover, from our definition of K_0, the difference vector $\boldsymbol{w} := \boldsymbol{u} - \boldsymbol{v} = sK_0^{-1}\,\boldsymbol{u}$ is adjustably "anti-sparse" since it lies in Ξ_{K_0} with $\|\boldsymbol{w}\|_2^2/\|\boldsymbol{w}\|_\infty^2 = K_0$. Interestingly, $\boldsymbol{u}$ and $\boldsymbol{v}$ are also consistent with respect to $\boldsymbol{A}$ since $\mathcal{Q}(\boldsymbol{\Phi u}) = \mathcal{Q}(\boldsymbol{\Phi v}) = \mathcal{Q}(\boldsymbol{\Phi u} + sK_0^{-1}\boldsymbol{\Phi u})$. This is due to the nature of quantization (*i.e.*, a rounding to the closest integer) and to the fact that both $\boldsymbol{\Phi u} \in \mathbb{Z}^M$ and $\|sK_0^{-1}\boldsymbol{\Phi u}\|_\infty \leqslant s < 1/2$.

[8] It is easy, but slightly more technical, to adapt our development here to the quantizer $\mathcal{Q}(\cdot) = \delta\lfloor \cdot/\delta \rfloor$ defined in Sec. 2. We thus prefer to select $\mathcal{Q}$ as a rounding operation for the sake of clarity.

Let us now assume, as involved by Prop. 1, that for $\epsilon := \frac{s}{4(c+s)\sqrt{K_0}}$, it is possible to find M arbitrarily large before $\epsilon^{-2}\bar{w}(\mathcal{K})^2\log(1+\frac{r_0}{\epsilon})$ so that, with high probability and for all $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}$ with $\boldsymbol{x}-\boldsymbol{y}\in\mho_{K_0}$,

$$\big((\tfrac{2}{\pi})^{1/2}-\epsilon-\tfrac{\kappa_{\rm sg}}{\sqrt{K_0}}\big)\,\|\boldsymbol{x}-\boldsymbol{y}\|-c\epsilon\ \leqslant\ \tfrac{1}{M}\|\mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{x})-\mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{y})\|_1,$$

with the constant $c>0$ defined in (18).

However, by taking the consistent vectors $\boldsymbol{x}=\boldsymbol{u}$ and $\boldsymbol{y}=\boldsymbol{v}$, this inequality leads by construction to

$$0=\tfrac{1}{M}\|\boldsymbol{A}(\boldsymbol{x})-\boldsymbol{A}(\boldsymbol{y})\|_1\ \geqslant\ \big((\tfrac{2}{\pi})^{1/2}-\epsilon-\tfrac{\kappa_{sg}}{\sqrt{K_0}}\big)\|\boldsymbol{x}-\boldsymbol{y}\|-c\epsilon\geqslant(\tfrac{1}{2}-\epsilon)\|\boldsymbol{x}-\boldsymbol{y}\|-c\epsilon.$$

In other words, since $\|\boldsymbol{x}-\boldsymbol{y}\|=s/\sqrt{K_0}\leqslant s$

$$\epsilon\geqslant\tfrac{1}{2}\tfrac{\|\boldsymbol{x}-\boldsymbol{y}\|}{c+\|\boldsymbol{x}-\boldsymbol{y}\|}\ \geqslant\ \tfrac{s}{2(c+s)\sqrt{K_0}}=2\epsilon,$$

which is a clear contradiction. We can similarly show that the same pair of consistent vectors $\boldsymbol{x}=\boldsymbol{u}$ and $\boldsymbol{y}=\boldsymbol{v}$ is incompatible with Prop. 2 as then the consistency width cannot be arbitrarily small, even for asymptotically large M.

Remark: Interestingly, the counter-example above is easily hijacked to show that it is impossible for the un-dithered quantized mapping $\boldsymbol{A}(\cdot):=\mathcal{Q}(\boldsymbol{\Phi}\,\cdot)$ to respect the following property for an arbitrarily small $\epsilon>0$ and provided M is large enough,

$$\big(C-\epsilon-g(K_0)\big)\,\|\boldsymbol{x}-\boldsymbol{y}\|-c\epsilon\ \leqslant\ h\big(\mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{x}),\mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{y})\big),\quad\forall\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}\text{ with }\boldsymbol{x}-\boldsymbol{y}\in\mho_{K_0},$$

where $C,c>0$ are some universal constants, $h:\mathbb{R}^M\times\mathbb{R}^M\to\mathbb{R}_+$ is any positive function vanishing on equal inputs (*e.g.*, a norm, a pseudo-norm or any metric) and g is any monotonically decreasing function with $\lim_{t\to+\infty}g(t)=0$. However, if $\mathcal{Q}$ is replaced by a sign operator as in [29, 44], then the known binary ϵ-stable embedding (or BϵSE) relates the *angular* distance between $\boldsymbol{x}$ and $\boldsymbol{y}$ to the Hamming distance of their mappings, *i.e.*, two distances that are equal to zero in our counter-example above, which removes the contradiction.

Remark: The question whether dithering is necessary in the special case of a quantized mapping with a Gaussian random matrix $\boldsymbol{\Phi}$ remains open.

6 Proof of Proposition 1

The architecture of this proof is inspired by the one developed in [44] for characterizing a 1-bit random mapping $\boldsymbol{A}':\mathbb{R}^N\to\{\pm1\}^M$, $\boldsymbol{u}\in\mathbb{R}^N\mapsto\boldsymbol{A}'(\boldsymbol{u})=\operatorname{sign}(\boldsymbol{\Phi}\boldsymbol{u})$. As will be clear below, some of the ingredients developed there had of course to be adapted to the specificities of $\boldsymbol{A}$ and of our scalar quantization. Compared to [44] we have also paid attention to optimize the dependency of M to the desired level of distortions induced by $\boldsymbol{A}$ in (4).

Prop. 1 is proved as a special case of a more general proposition based on a "softer" variant of $\mathcal{D}$. This new pseudo-distance is established as follows. Defining the random mapping $\boldsymbol{u}\in\mathbb{R}^N\mapsto\boldsymbol{\Phi}^{\xi}(\boldsymbol{u}):=\boldsymbol{\Phi}\boldsymbol{u}+\boldsymbol{\xi}$, with Φ^{ξ}_i its $i^{\rm th}$ component, we observe that for any $\boldsymbol{x},\boldsymbol{y}\in\mathbb{R}^N$,

$$\mathcal{D}(\boldsymbol{x},\boldsymbol{y})=\tfrac{\delta}{M}\textstyle\sum_{i=1}^M\sum_{k\in\mathbb{Z}}\mathbb{I}[\mathcal{E}(\Phi^{\xi}_i(\boldsymbol{x})-k\delta,\Phi^{\xi}_i(\boldsymbol{y})-k\delta)],\tag{24}$$

with the *distinct sign event* $\mathcal{E}(a,b):=\{\operatorname{sign}a\neq\operatorname{sign}b\}$. In words, for each $i\in[M]$, the sum over k above simply counts the number of thresholds in $\delta\mathbb{Z}$ separating $\Phi^{\xi}_i(\boldsymbol{x})=\boldsymbol{\varphi}_i^{\top}\boldsymbol{x}+\xi_i$ and

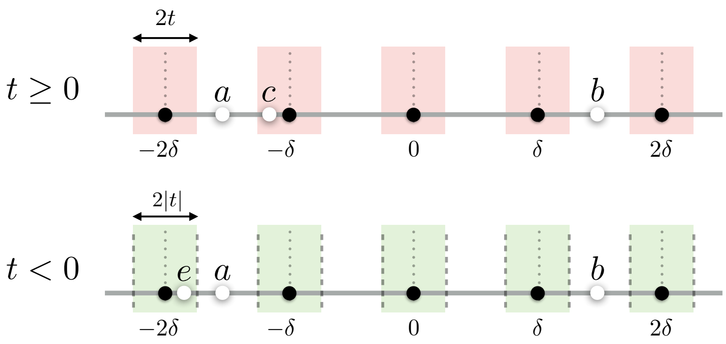

Figure 1: Behavior of the distance $d^t(a,b)$ for $a,b \in \mathbb{R}$. *On the top*, $t \geqslant 0$ and forbidden areas determined by $\mathcal{F}^t$ are created when counting the number of thresholds $k\delta$ separating a and b. For instance, for an additional point $c \in \mathbb{R}$ as on the figure, $d^0(a,b) = d^0(c,b) = 3\delta$ but $3\delta = d^t(a,b) = d^t(c,b) + \delta \geqslant d^t(c,b)$ as c lies in one forbidden area. *On the bottom* figure, $t \leqslant 0$ and threshold counting procedure operated by d^t is relaxed. Now $d^t(a,b)$ counts the number of limits (in dashed) of the green areas determined by $\mathcal{F}^t$, recording only one per thresholds $k\delta$, that separate a and b. Here, for $e \in \mathbb{R}$ as on the figure, $d^0(a,b) = d^0(e,b) = 3\delta$ but $4\delta = d^t(e,b) = d^t(a,b)+\delta \geqslant d^t(a,b)$.

$\Phi^\xi_i(\boldsymbol{y}) = \boldsymbol{\varphi}_i^\top \boldsymbol{y} + \xi_i$ on the real line, since $\mathbb{I}[\mathcal{E}(\Phi^\xi_i(\boldsymbol{x}) - k\delta, \Phi^\xi_i(\boldsymbol{y}) - k\delta)]$ is equal to 1 for those and 0 for any other thresholds.

Notice that the decomposition (24) also justifies the observation made at the end of Sec. 3.2, namely the existence of uniform random tessellations of $\mathbb{R}^N$. Indeed, from the definition of $\boldsymbol{A}$, for each $i \in [M]$, $\sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{E}(\Phi^\xi_i(\boldsymbol{x}) - k\delta, \Phi^\xi_i(\boldsymbol{y}) - k\delta)]$ also counts the number of parallel affine hyperplanes $\Pi_i := \{\boldsymbol{u} \in \mathbb{R}^N : \exists k \in \mathbb{Z},\ \boldsymbol{\varphi}_i^\top \boldsymbol{u} + \xi_i - k\delta = 0\}$, all normal to $\boldsymbol{\varphi}_i$ and $\delta/\|\boldsymbol{\varphi}_i\|$ far apart, separating $\boldsymbol{x}$ and $\boldsymbol{y} \in \mathbb{R}^N$. In other words, $\mathbb{R}^N$ is here tessellated with multiple so-called "hyperplane wave partitions" $\{\Pi_i : i \in [M]\}$ [23, 50] with random orientations, periods and dithered origin.

Based on this observation, and as a generalization of an equivalent distance given in [44, Sec. 5] for binary mappings, we introduce for some $t \in \mathbb{R}$ the new pseudo-distance

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) := \tfrac{\delta}{M} \textstyle\sum_{i=1}^M \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}^t(\Phi^\xi_i(\boldsymbol{x}) - k\delta, \Phi^\xi_i(\boldsymbol{y}) - k\delta)], \tag{25}$$

by defining the set

$$\mathcal{F}^t(a,b) = \{a > t,\ b \leqslant -t\} \cup \{a < -t,\ b \geqslant t\}. \tag{26}$$

The pseudo-distance $\mathcal{D}^t$ is a non-increasing function of t, with $\mathcal{F}^0(a,b) = \mathcal{E}(a,b)$ and

$$\mathcal{D}^{|t|}(\boldsymbol{x},\boldsymbol{y}) \leqslant \mathcal{D}(\boldsymbol{x},\boldsymbol{y}) \leqslant \mathcal{D}^{-|t|}(\boldsymbol{x},\boldsymbol{y}).$$

The behavior of $\mathcal{D}^t$ is best understood by introducing the one-dimensional distance

$$d^t(a,b) := \delta \textstyle\sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}^t(a - k\delta, b - k\delta)] \ \in\ \delta\mathbb{N}, \quad \text{for } a,b \in \mathbb{R}, \tag{27}$$

so that

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) = \tfrac{1}{M} \textstyle\sum_{i=1}^M d^t(\Phi^\xi_i(\boldsymbol{x}), \Phi^\xi_i(\boldsymbol{y})). \tag{28}$$

Fig. 1 explains how $d^t(a,b)$ evolves for positive and negative t, observing that, for each $k \in \mathbb{Z}$, $\mathcal{F}^t(a - k\delta, b - k\delta)$ determines forbidden or relaxed areas around the thresholds $k\delta$ separating a and b and counted by $d^t(a,b)$. Moreover, the next Lemma, proved in App. B, provides a first evaluation of the impact of the distance "softening", by observing that, essentially, $d^t(a,b)$ is not very far from both $|a-b|$ and $d^s(a,b)$ for s close to t.

Lemma 1. *For any $a,b \in \mathbb{R}$ and $t,s \in \mathbb{R}$,*

$$\big|d^t(a,b) - d^s(a,b)\big| \leqslant\ 4(\delta + |t-s|), \tag{29}$$

$$\big|d^t(a,b) - |a-b|\big| \leqslant\ 4(\delta + |t|). \tag{30}$$

As announced above, we aim now at proving the next proposition whose special case $t = 0$ leads to Prop. 1.

Proposition 3. *Given $\delta > 0$, $\epsilon \in (0,1)$, $t \in \mathbb{R}$, $K_0 > 0$, a bounded subset $\mathcal{K} \subset \mathbb{R}^N$ and a sub-Gaussian distribution $\mathcal{N}_{\mathrm{sg},\alpha}$ respecting* (10) *for $0 \leqslant \kappa_{\mathrm{sg}} < \infty$, there exist some values $C, c, c' > 0$, only depending on α, such that, if*

$$M \;\geqslant C \max(\epsilon^{-2}\mathcal{H}(\mathcal{K}, \sqrt{\delta^2\epsilon^3}), \tfrac{1}{\delta^2\epsilon^3}\, w(\mathcal{K}_{\sqrt{\delta^2\epsilon^3}})^2), \tag{31}$$

with $\mathcal{H}(\mathcal{K},\eta)$ the Kolmogorov η-entropy of $\mathcal{K}$ and the local set $\mathcal{K}_\eta := (\mathcal{K}-\mathcal{K}) \cap \eta\,\mathbb{B}^N$ for $\eta > 0$, then for $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}_{\mathrm{sg},\alpha}(0,1)$, a dithering $\boldsymbol{\xi} \sim \mathcal{U}^M([0,\delta])$, and the associated mapping $\boldsymbol{A}$ defined in (4)*, we have with probability exceeding $1 - e^{-c\epsilon^2 M}$ that for all pairs $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{K}$ with $\boldsymbol{x} - \boldsymbol{y} \in \daleth_{K_0}$,*

$$\big|\ \mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|\ \big| \;\leqslant\; (\epsilon + \tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\| \;+\; c'(|t| + \delta\epsilon). \tag{32}$$

Proof. The proof sketch of Prop. 3 is as follows: *(i)* given $\boldsymbol{x},\boldsymbol{y} \in \mathbb{R}^N$, we first show that the r.v. $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y})$ concentrates with high probability around $(\frac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|$ up to a systematic bias $\frac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\|$ due to the sub-Gaussian nature of $\boldsymbol{\Phi}$ and controlled by the anti-sparse level of $\boldsymbol{x}-\boldsymbol{y}$; *(ii)* we take a finite covering of $\mathcal{K}$ by a η-net $\mathcal{G}_\eta \subset \mathcal{K}$ (for $\eta > 0$) and we extend the concentration of $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y})$ to all vectors of $\mathcal{G}_\eta$ by union bound; *(iii)* we show that the softened pseudo-distance $\mathcal{D}^t$ is sufficiently continuous in a neighborood of each pair of vectors in $\mathcal{G}_\eta$, which then allows us to extend (32) to all pair of vectors in $\mathcal{K}$, as stated by Prop. 3.

(i) Concentration of $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y})$: Given a fixed pair $\boldsymbol{x},\boldsymbol{y} \in \mathbb{R}^N$, we show that $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y})$ concentrates around its mean by bounding its sub-Gaussian norm as defined in (5). From (28), $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) = M^{-1}\sum_i Z_i^t$ with the M random variables $Z_i^t := d^t(\boldsymbol{\varphi}_i^\top\boldsymbol{x}+\xi_i, \boldsymbol{\varphi}_i^\top\boldsymbol{y}+\xi_i)$ for $1 \leqslant i \leqslant M$. However, the sum of D independent sub-Gaussian random variables $\{X_1, \cdots, X_D\}$ is approximately invariant under rotation [52], which means that

$$\|\sum_i (X_i - \mathbb{E}X_i)\|^2_{\psi_2} \;\lesssim\; \sum_i \|X_i - \mathbb{E}X_i\|^2_{\psi_2}. \tag{33}$$

Therefore, from (33), we find

$$\big\|\textstyle\sum_{i=1}^M (Z_i^t - \mathbb{E}Z_i^t)\big\|^2_{\psi_2} \;\lesssim\; \textstyle\sum_{i=1}^M \|Z_i^t - \mathbb{E}Z_i^t\|^2_{\psi_2} = M\|Z_1^t - \mathbb{E}Z_1^t\|^2_{\psi_2} \;\leqslant 4M\|Z_1^t\|^2_{\psi_2}. \tag{34}$$

As shown in the following lemma (proved in App. C by using Lemma 1) $\|Z_1^t\|_{\psi_2}$ can be upper bounded (and with it, the sub-Gaussian norm of $\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y})$).

Lemma 2. *Let us take $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$ and $\xi \sim \mathcal{U}([0,\delta])$. For a fixed $t \in \mathbb{R}$, the random variable $Z^t := d^t(\boldsymbol{\varphi}^\top\boldsymbol{x} + \xi, \boldsymbol{\varphi}^\top\boldsymbol{y} + \xi)$ is sub-Gaussian with ψ_2-norm bounded by*

$$\|Z^t\|_{\psi_2} \lesssim \delta + |t| + \|\boldsymbol{x}-\boldsymbol{y}\|. \tag{35}$$

Moreover,

$$\big|\mathbb{E}\, Z^t - \mu_{\mathrm{sg}}(\boldsymbol{x}-\boldsymbol{y})\big|\ \lesssim\ |t|, \tag{36}$$

with $\mu_{\mathrm{sg}}(\boldsymbol{x}-\boldsymbol{y}) = \mathbb{E}|\langle\boldsymbol{\varphi}, \boldsymbol{x}-\boldsymbol{y}\rangle| = (\frac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|$ if $\boldsymbol{\varphi} \sim \mathcal{N}^N(0,1)$.

Consequently, from (34) and (35), $X := \frac{1}{\sqrt{M}}\sum_{i=1}^M (Z_i^t - \mathbb{E}Z_i^t)$ is itself sub-Gaussian with $\|X\|_{\psi_2} \lesssim \delta + |t| + \|\boldsymbol{x}-\boldsymbol{y}\|$. Therefore, from the tail bound (6), there exists a $c>0$ such that for any $\epsilon > 0$

$$\mathbb{P}\big[\big|\tfrac{1}{M}\sum_i (Z_i^t - \mathbb{E}Z_i^t)\big| > \epsilon\,(\delta + |t| + \|\boldsymbol{x}-\boldsymbol{y}\|)\big] \leqslant 2\exp\big(-c\epsilon^2 M\big).$$

Since $\mathbb{E}Z_i^t = \mathbb{E}Z_1^t$ and $\mathbb{E}Z_i^0 = \mathbb{E}|\langle\boldsymbol{\varphi}, \boldsymbol{x}-\boldsymbol{y}\rangle| = \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y})$ for all $i\in[M]$, (36) provides

$$\begin{aligned}\big|\tfrac{1}{M}\textstyle\sum_i (Z_i^t - \mathbb{E}Z_i^t)\big| &= |\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - \mathbb{E}Z_1^t| \\ &\geqslant |\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y})| - |\mathbb{E}Z_1^t - \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y})| \\ &\geqslant |\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y})| - c'|t|,\end{aligned}$$

for some constant $c'>0$, and

$$\mathbb{P}\big[\big|\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - \mu_{\rm sg}(\boldsymbol{x}-\boldsymbol{y})\big| > c'|t| + \epsilon\,(\delta + |t| + \|\boldsymbol{x}-\boldsymbol{y}\|)\big] \leqslant 2\exp\big(-c\epsilon^2 M\big). \tag{37}$$

(ii) Extension to a covering of $\mathcal{K}$: Given a radius $\eta > 0$ to be specified later, let $\mathcal{G}_\eta$ an η-net of $\mathcal{K}$, *i.e.*, a finite vector set such that for any $\boldsymbol{x}\in\mathcal{K}$ there exists a $\boldsymbol{x}_0 \in \mathcal{G}_\eta$ with $\|\boldsymbol{x}-\boldsymbol{x}_0\| \leqslant \eta$. In particular, any vectors $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}$ can then be written as

$$\boldsymbol{x} = \boldsymbol{x}_0 + \boldsymbol{x}', \quad \boldsymbol{y} = \boldsymbol{y}_0 + \boldsymbol{y}', \tag{38}$$

for some $\boldsymbol{x}_0,\boldsymbol{y}_0\in\mathcal{G}_\eta$ and $\boldsymbol{x}',\boldsymbol{y}'\in(\mathcal{K}-\mathcal{K})\cap\eta\mathbb{B}^N$. We also assume that the size of $\mathcal{G}_\eta$ is minimal so that, by definition, $\log|\mathcal{G}_\eta| = \mathcal{H}(\mathcal{K},\eta)$, with $\mathcal{H}$ the *Kolmogorov η-entropy* of $\mathcal{K}$.

Since there are no more than $|\mathcal{G}_\eta|^2$ distinct pairs of vectors in $\mathcal{G}_\eta$, given $t\in\mathbb{R}$, a standard union bound over (37) shows that there exist some constant $C, c', c'' > 0$ such that, if $M \geqslant C\epsilon^{-2}\mathcal{H}(\mathcal{K},\eta)$

$$\begin{aligned}&\mathbb{P}\big[\,\forall \boldsymbol{x}_0,\boldsymbol{y}_0\in\mathcal{G}_\eta,\ \big|\mathcal{D}^t(\boldsymbol{x}_0,\boldsymbol{y}_0) - \mu_{\rm sg}(\boldsymbol{x}_0-\boldsymbol{y}_0)\big| \leqslant c'|t| + \epsilon\,(\delta + |t| + \|\boldsymbol{x}_0-\boldsymbol{y}_0\|)\big] \\ &\quad\geqslant\ 1\ -\ |\mathcal{G}_\eta|^2\exp\big(-c\epsilon^2 M\big)\ \geqslant\ 1 - 2\exp\big(-c''\epsilon^2 M\big).\end{aligned} \tag{39}$$

(iii) Extension to $\mathbb{R}^N$ by continuity of $\mathcal{D}^t$: We can extend the event characterized in (39) to all pairs of vectors in $\mathcal{K}$ by analyzing the continuity property of $\mathcal{D}^t$ in a limited neighborhood around the considered vectors. We propose here to analyze this continuity with respect to ℓ_2-perturbations of those vectors, as compared to ℓ_1-perturbations in [44]. As will be clearer later, this allows us to reach a better control over M with respect to ϵ.

Lemma 3 (Continuity with respect to ℓ_2-perturbations)**.** *Let $\boldsymbol{x}_0,\boldsymbol{y}_0,\boldsymbol{x}',\boldsymbol{y}'\in\mathbb{R}^N$. We assume that $\|\boldsymbol{\Phi}\boldsymbol{x}'\| \leqslant \eta\sqrt{M}$, $\|\boldsymbol{\Phi}\boldsymbol{y}'\| \leqslant \eta\sqrt{M}$ for some $\eta>0$. Then for every $t\in\mathbb{R}$ and $P\geqslant 1$ one has*

$$\mathcal{D}^{t+\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) - 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}})\ \leqslant\ \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}')\ \leqslant\ \mathcal{D}^{t-\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) + 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}). \tag{40}$$

The proof is given in App. D. Interestingly, the following proposition proved in App. E shows that $\|\boldsymbol{\Phi}\boldsymbol{x}'\|$ and $\|\boldsymbol{\Phi}\boldsymbol{y}'\|$ can indeed be bounded uniformly for all $\boldsymbol{x}',\boldsymbol{y}'\in\mathcal{K}_\eta := (\mathcal{K}-\mathcal{K})\cap\eta\mathbb{B}^N$.

Lemma 4 (Diameter stability under random projections)**.** *Let $\mathcal{R}\subset\mathbb{R}^N$ be bounded,* i.e., $\|\mathcal{R}\| := \sup_{\boldsymbol{u}\in\mathcal{R}}\|\boldsymbol{u}\| < \infty$ *and assume $\mathcal{R}\ni\boldsymbol{0}$. Then, for some $c>0$, if*

$$M \gtrsim \tfrac{\alpha^4 w(\mathcal{R})^2}{\|\mathcal{R}\|^2},$$

for $\mathbf{\Phi} \sim \mathcal{N}_{\mathrm{sg},\alpha}^{M\times N}(0,1)$ *and with probability at least* $1-\exp(-c\,\alpha^{-4}M)$*, we have for all* $\boldsymbol{x}\in\mathcal{R}$

$$\tfrac{1}{M}\|\mathbf{\Phi}\boldsymbol{x}\|^2 \leqslant \|\mathcal{R}\|^2, \tag{41}$$

i.e., $\|\mathbf{\Phi}\mathcal{R}\| \leqslant \sqrt{M}\,\|\mathcal{R}\|$.

For the sake of simplicity, we consider below the sub-Gaussian parameter α as fixed and integrate it in explicit or hidden constants, as in the notations "$\lesssim$" or "$\gtrsim$". Noting that $\|\mathcal{K}_\eta\| \leqslant \eta$ and using a union bound over (39) and (41), we get that if

$$M \gtrsim \max(\epsilon^{-2}\mathcal{H}(\mathcal{K},\eta), \eta^{-2}w(\mathcal{K}_\eta)^2),$$

with probability higher than $1-4\exp(-c'\epsilon^2 M)$, for all $\boldsymbol{x}_0,\boldsymbol{y}_0\in\mathcal{G}_\eta$ and all $\boldsymbol{x}',\boldsymbol{y}'\in\mathcal{K}_\eta$,

$$\big|\mathcal{D}^{t-\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) - \mu_{\mathrm{sg}}(\boldsymbol{x}_0-\boldsymbol{y}_0)\big| \leqslant c|t-\eta\sqrt{P}| + \epsilon\,(\delta + |t-\eta\sqrt{P}| + \|\boldsymbol{x}_0-\boldsymbol{y}_0\|), \tag{42}$$

$$\big|\mathcal{D}^{t+\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) - \mu_{\mathrm{sg}}(\boldsymbol{x}_0-\boldsymbol{y}_0)\big| \leqslant c|t+\eta\sqrt{P}| + \epsilon\,(\delta + |t+\eta\sqrt{P}| + \|\boldsymbol{x}_0-\boldsymbol{y}_0\|), \tag{43}$$

$$\|\mathbf{\Phi}\boldsymbol{x}'\|_2 \leqslant \eta\sqrt{M}, \quad \|\mathbf{\Phi}\boldsymbol{y}'\|_2 \leqslant \eta\sqrt{M}, \tag{44}$$

for some $C,c,c'>0$ depending only on α.

Therefore, for any $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}$, using sequentially (38), (44), the upper bound given in Lemma 3 and (42) provides

$$\begin{aligned}\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) &\leqslant \mathcal{D}^{t-\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) + 4(\tfrac{\delta}{P}+\tfrac{\eta}{\sqrt{P}})\\ &\leqslant (c+\epsilon)|t-\eta\sqrt{P}| + \mu_{\mathrm{sg}}(\boldsymbol{x}_0-\boldsymbol{y}_0) + \epsilon\|\boldsymbol{x}_0-\boldsymbol{y}_0\| + \epsilon\delta + 4(\tfrac{\delta}{P}+\tfrac{\eta}{\sqrt{P}}).\end{aligned}$$

However, given $\boldsymbol{\varphi}\sim\mathcal{N}_{\mathrm{sg},\alpha}^N(0,1)$, using Jensen's inequality, the reverse triangular inequality and (9), we find

$$\begin{aligned}\big|\mu_{\mathrm{sg}}(\boldsymbol{x}_0-\boldsymbol{y}_0) - \mu_{\mathrm{sg}}(\boldsymbol{x}-\boldsymbol{y})\big| &= \big|\mathbb{E}|\langle\boldsymbol{\varphi},\boldsymbol{x}_0-\boldsymbol{y}_0\rangle| - \mathbb{E}|\langle\boldsymbol{\varphi},\boldsymbol{x}-\boldsymbol{y}\rangle|\big|\\ &\leqslant \mathbb{E}|\langle\boldsymbol{\varphi},\boldsymbol{x}'\rangle| + \mathbb{E}|\langle\boldsymbol{\varphi},\boldsymbol{y}'\rangle| \leqslant 2\eta.\end{aligned}$$

Moreover, $|\|\boldsymbol{x}_0-\boldsymbol{y}_0\| - \|\boldsymbol{x}-\boldsymbol{y}\|| \leqslant 2\eta$, so that,

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - \mu_{\mathrm{sg}}(\boldsymbol{x}-\boldsymbol{y}) \leqslant \epsilon\|\boldsymbol{x}-\boldsymbol{y}\| + (c+\epsilon)(|t|+\eta\sqrt{P}) + 2\eta + 2\epsilon\eta + \epsilon\delta + 4(\tfrac{\delta}{P}+\tfrac{\eta}{\sqrt{P}}).$$

If $\boldsymbol{x}-\boldsymbol{y}\in\mho_{K_0}$, then (14) induces $|\mu_{\mathrm{sg}}(\boldsymbol{x}-\boldsymbol{y}) - (\frac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|| \leqslant \kappa_{\mathrm{sg}}\|\boldsymbol{x}-\boldsymbol{y}\|/\sqrt{K_0}$ and assuming $\epsilon<1$, there exists a $c>0$ such that

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\| \leqslant (\epsilon+\tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\| + c(|t| + \eta\sqrt{P} + \eta + \epsilon\delta + \tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}). \tag{45}$$

Taking $P=\epsilon^{-1}\geqslant 1$ and $\eta=\delta\epsilon^{3/2}<\delta\epsilon$, which gives $\eta\sqrt{P}=\delta\epsilon$ and $\eta/\sqrt{P}=\delta\epsilon^2\leqslant\delta\epsilon$, we find for another $c>0$

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\| \leqslant (\epsilon+\tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\| + c(|t|+\delta\epsilon).$$

Similarly, using (38), (44), the lower bound given in Lemma 3 and (43), we obtain

$$\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\| \geqslant -(\epsilon+\tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\| - c(|t|+\delta\epsilon).$$

Finally, we have thus shown that there exist some $c,c'>0$ such that for

$$M \gtrsim \max(\epsilon^{-2}\mathcal{H}(\mathcal{K},\sqrt{\delta^2\epsilon^3}), \tfrac{1}{\delta^2\epsilon^3}\,w(\mathcal{K}_{\sqrt{\delta^2\epsilon^3}})^2), \tag{46}$$

with probability at least $1-4\exp(-c'\epsilon^2 M)$ the bound

$$\big|\mathcal{D}^t(\boldsymbol{x},\boldsymbol{y}) - (\tfrac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|\big| \leqslant (\epsilon+\tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{x}-\boldsymbol{y}\| + c(|t|+\delta\epsilon)$$

holds for all $\boldsymbol{x},\boldsymbol{y}\in\mathcal{K}\cap\mho_{K_0}$, which finishes the proof of Prop. 3. $\square$

As mentioned earlier, Prop. 1 is thus obtained by simplifying the requirement (31) appearing in Prop. 3. First, for a general bounded set $\mathcal{K}$, since the Sudakov inequality in (P14) provides $\mathcal{H}(\mathcal{K},\eta) \lesssim \frac{w(\mathcal{K})^2}{\eta^2}$, noticing that $\mathcal{K}_\eta \subset (\mathcal{K}-\mathcal{K})$ and that (P3) and (P4) provide $w(\mathcal{K}_\eta) \leqslant w(\mathcal{K}-\mathcal{K}) \leqslant 2w(\mathcal{K})$, we deduce that (46) holds if

$$M \gtrsim \tfrac{1}{\delta^2\epsilon^5} w(\mathcal{K})^2,$$

as imposed in (16).

Second, in the case of a quantized embedding of the structured sets defined in the Introduction (see Def. 1), we can even reach a much weaker condition on M. Indeed, for such a set $\mathcal{K}$ with $d = \|\mathcal{K}\|$, from (3b) and the definition of $\bar{w}$, we have for any $\eta > 0$

$$w(\mathcal{K}_\eta)^2 = w\big((\mathcal{K}-\mathcal{K}) \cap \eta\mathbb{B}^N\big)^2 = d^2 w\big((d^{-1}\mathcal{K} - d^{-1}\mathcal{K}) \cap (d^{-1}\eta\mathbb{B}^N)\big)^2 \leqslant \eta^2\, \bar{w}(\mathcal{K})^2,$$

so that, from (3a), the right-hand side of (31) can be bounded as

$$\begin{aligned}\max(\epsilon^{-2}\mathcal{H}(\mathcal{K},\sqrt{\delta^2\epsilon^3}), \tfrac{1}{\delta^2\epsilon^3}\, w(\mathcal{K}_{\sqrt{\delta^2\epsilon^3}})^2) &\leqslant \max(\epsilon^{-2}\bar{w}(\mathcal{K})^2 \log(1+\tfrac{\|\mathcal{K}\|}{\sqrt{\delta^2\epsilon^3}}), \bar{w}(\mathcal{K})^2)\\ &\leqslant \epsilon^{-2}\bar{w}(\mathcal{K})^2 \log(1+\tfrac{\|\mathcal{K}\|}{\sqrt{\delta^2\epsilon^3}}).\end{aligned}$$

This explains the simpler requirement (17) needed for structured sets in Prop. 1.

Example: Let us conclude this section by deducing an upper bound on $\bar{w}^2(\mathcal{K})$ for the set $\mathcal{K} := \Sigma_K^{\boldsymbol{\Psi}} \cap d\,\mathbb{B}^N$ (with $d = \|\mathcal{K}\| > 0$) of bounded K-sparse vectors in an orthonormal basis $\boldsymbol{\Psi} \in \mathbb{R}^{N\times N}$ of $\mathbb{R}^N$. We first notice that since $w(\Sigma_K^{\boldsymbol{\Psi}} \cap d\,\mathbb{B}^N) = w(\Sigma_K \cap d\,\mathbb{B}^N)$ by invariance over the orthogonal group $\mathcal{O}_N$ (see (P12) in Table 1) and from (P17),

$$w(\mathcal{K}/\|\mathcal{K}\|)^2 \lesssim K \log N/K.$$

Moreover, the Kolmogorov entropy is also invariant under $\mathcal{O}_N$, *i.e.*, $\mathcal{H}(\Sigma_K \cap d\,\mathbb{B}^N,\eta) = \mathcal{H}(\Sigma_K^{\boldsymbol{\Psi}} \cap d\,\mathbb{B}^N,\eta)$ and it is known that (see, *e.g.*, [16])

$$\mathcal{H}(\Sigma_K \cap d\,\mathbb{B}^N,\eta) \lesssim \log(\tbinom{N}{K}(1+\tfrac{2d}{\eta})^K) \leqslant K\log(\tfrac{eN}{K}(1+\tfrac{2d}{\eta})) \;\lesssim K\log(\tfrac{N}{K})\log(1+\tfrac{d}{\eta}),$$

by using Stirling's bound. This shows that $\mathcal{H}(\mathcal{K},\eta) \leqslant \bar{w}(\mathcal{K})^2\log(1+\frac{d}{\eta})$ with $\bar{w}(\mathcal{K})^2 \lesssim K\log N/K$. Additionally, since $\Sigma_K^{\boldsymbol{\Psi}}$ is invariant under dilation, $d^{-1}\mathcal{K} - d^{-1}\mathcal{K} \subset \Sigma_K^{\boldsymbol{\Psi}} - \Sigma_K^{\boldsymbol{\Psi}} \subset \Sigma_{2K}^{\boldsymbol{\Psi}}$ and

$$\begin{aligned}w\big((d^{-1}\mathcal{K} - d^{-1}\mathcal{K}) \cap \epsilon\mathbb{B}^N\big)^2 &\leqslant w\big(\Sigma_{2K}^{\boldsymbol{\Psi}} \cap \epsilon\mathbb{B}^N\big)^2 = \epsilon^2 w\big(\Sigma_{2K}^{\boldsymbol{\Psi}} \cap \mathbb{B}^N\big)^2\\ &\lesssim \epsilon^2 2K\log(N/2K) \lesssim \epsilon^2 K\log(N/K),\end{aligned}$$

showing again, by matching with (3b), that we have $\bar{w}(\mathcal{K})^2 \lesssim K\log(N/K)$.

This confirms that $\bar{w}(\mathcal{K})^2$ has the same upper bound than $w(\mathcal{K}/\|\mathcal{K}\|)^2$. Therefore, for the structured set $\mathcal{K}$ of bounded K-sparse vectors, (46) (and therefore (17)) is then satisfied if

$$M \gtrsim \tfrac{1}{\epsilon^2} K\log(\tfrac{N}{K})\log(1+\tfrac{\|\mathcal{K}\|}{\delta\sqrt{\epsilon^3}}).$$

7 Proof of Proposition 2

Using the context defined in Prop. 2 and for M satisfying (19), we are going to show the contraposition of (21), *i.e.*, that with probability at least $1-2e^{-c\epsilon M/(1+\delta)}$ for some $c > 0$ and for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{K}$ with $\boldsymbol{x} - \boldsymbol{y} \in \Sigma_{K_0}$, having $\|\boldsymbol{x} - \boldsymbol{y}\| > \epsilon$ involves $\mathcal{Q}(\boldsymbol{\Phi x} + \boldsymbol{\xi}) \neq \mathcal{Q}(\boldsymbol{\Phi y} + \boldsymbol{\xi})$, or equivalently that

$$\|\boldsymbol{x} - \boldsymbol{y}\| > \epsilon \quad \Rightarrow \quad \mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) \geqslant \tfrac{\delta}{M}, \tag{47}$$

from the definition of $\mathcal{D}$ in (24).

The proof sketch is a follows. First, for some $\eta > 0$, we create a finite η-covering of the set $\bar{\mathcal{K}} \subset \mathcal{K} \times \mathcal{K}$ of vector pairs whose difference belongs to Σ_{K_0}. Second, in order to show (47), we leverage the continuity of the pseudo-distance $\mathcal{D}^t$ under ℓ_2-perturbations (Lemma 3), as it happens that all points of $\bar{\mathcal{K}}$ are obtained by ℓ_2-perturbations of the η-covering and that, moreover, those perturbations are stable under projections by $\boldsymbol{\Phi}$ (Lemma 4). Finally, we adjust η and some additional parameters to show that, with high probability, the softened distance $\mathcal{D}^t(\boldsymbol{x}_0, \boldsymbol{y}_0)$, for some t depending on η, is large enough over all pairs $(\boldsymbol{x}_0, \boldsymbol{y}_0)$ of the covering compatible with $\|\boldsymbol{x} - \boldsymbol{y}\| \geqslant \epsilon$, hence inducing (47).

Let us define the set $\bar{\mathcal{K}} = \{(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{K} \times \mathcal{K} : \boldsymbol{x} - \boldsymbol{y} \in \Sigma_{K_0}\} \subset \mathcal{K} \times \mathcal{K}$. We introduce a minimal η-net $\bar{\mathcal{G}}_\eta \subset \bar{\mathcal{K}}$ of $\bar{\mathcal{K}}$ with $0 < \eta < \epsilon/2$ to be specified later, such that for all $(\boldsymbol{x}, \boldsymbol{y}) \in \bar{\mathcal{K}}$, there exists a $(\boldsymbol{x}_0, \boldsymbol{y}_0) \in \bar{\mathcal{G}}_\eta$ with

$$\|(\boldsymbol{x}, \boldsymbol{y}) - (\boldsymbol{x}_0, \boldsymbol{y}_0)\| \leqslant \eta,$$

which also involves $\|\boldsymbol{x} - \boldsymbol{x}_0\| \leqslant \eta$ and $\|\boldsymbol{y} - \boldsymbol{y}_0\| \leqslant \eta$.

The size of this minimal η-net is bounded as $\log |\bar{\mathcal{G}}_\eta| \leqslant 2\mathcal{H}(\mathcal{K}, \eta/\sqrt{2})$. Indeed, by the semi-additivity of the Kolmogorov entropy [35, Theorem 2], $\bar{\mathcal{K}} \subset \mathcal{K} \times \mathcal{K}$ involves that $\mathcal{H}(\bar{\mathcal{K}}, \rho) \leqslant \mathcal{H}(\mathcal{K} \times \mathcal{K}, \rho)$ for any $\rho > 0$. Since a ρ-net of $\mathcal{K} \times \mathcal{K}$ can be obtained by the product $\mathcal{G}_{\rho'} \times \mathcal{G}_{\rho'}$, with $\rho' = \rho/\sqrt{2}$ and $\mathcal{G}_{\rho'}$ a ρ'-net covering of $\mathcal{K}$, we obtain $\mathcal{H}(\bar{\mathcal{K}}, \rho) \leqslant 2\mathcal{H}(\mathcal{K}, \rho/\sqrt{2})$.

As for the proof of Prop. 1 in Sec. 6, by construction, all $(\boldsymbol{x}, \boldsymbol{y}) \in \bar{\mathcal{K}}$ can also be written as

$$(\boldsymbol{x}, \boldsymbol{y}) = (\boldsymbol{x}_0, \boldsymbol{y}_0) + (\boldsymbol{x}', \boldsymbol{y}'),$$

with $(\boldsymbol{x}_0, \boldsymbol{y}_0) \in \bar{\mathcal{G}}_\eta$, $(\boldsymbol{x}', \boldsymbol{y}') \in (\bar{\mathcal{K}} - \bar{\mathcal{K}}) \cap \eta \mathbb{B}^{2N}$. Notice that we have also $\boldsymbol{x}', \boldsymbol{y}' \in \mathcal{K}_\eta := (\mathcal{K} - \mathcal{K}) \cap \eta \mathbb{B}^N$, since $\boldsymbol{x}, \boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{y}_0 \in \mathcal{K}$ and $\max(\|\boldsymbol{x}'\|, \|\boldsymbol{y}'\|) \leqslant \|(\boldsymbol{x}', \boldsymbol{y}')\| \leqslant \eta$.

As stated by Lemma 4, the diameter of the *local set* $\mathcal{K}_\eta$ is stable with respect to random projections. Since $\|\mathcal{K}_\eta\| \leqslant \eta$, there exist indeed two values $C, c > 0$, only depending on the sub-Gaussian norm α, such that if

$$M \geqslant C \eta^{-2} w(\mathcal{K}_\eta)^2 \tag{48}$$

and $\boldsymbol{\Phi} \sim \mathcal{N}^{M \times N}_{\mathrm{sg},\alpha}(0, 1)$, we have with probability at least $1 - 2\exp(-cM)$,

$$\|\boldsymbol{\Phi} \mathcal{K}_\eta\| := \sup_{\boldsymbol{u} \in \mathcal{K}_\eta} \|\boldsymbol{\Phi u}\| \leqslant \sqrt{M} \|\mathcal{K}_\eta\| \leqslant \eta \sqrt{M}. \tag{49}$$

Therefore, $\|\boldsymbol{\Phi x}'\| \leqslant \eta\sqrt{M}$ and $\|\boldsymbol{\Phi y}'\| \leqslant \eta\sqrt{M}$ under the same conditions.

Moreover, if the previous event occurs, then, Lemma 3 for $t = 0$ shows that for any $P \geqslant 1$,

$$\mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) = \mathcal{D}^0(\boldsymbol{x}_0 + \boldsymbol{x}', \boldsymbol{y}_0 + \boldsymbol{y}') \geqslant \mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{x}_0, \boldsymbol{y}_0) - 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}). \tag{50}$$

Consequently, for reaching $\mathcal{D}(\boldsymbol{x}, \boldsymbol{y}) \geqslant \delta/M$ as expressed in (47), since $\|\boldsymbol{x} - \boldsymbol{y}\| \geqslant \epsilon$ involves $\|\boldsymbol{x}_0 - \boldsymbol{y}_0\| \geqslant \epsilon - 2\eta$, the proof can be deduced if we can guarantee that, for all $(\boldsymbol{u}, \boldsymbol{v}) \in \bar{\mathcal{G}}_\eta$ with $\|\boldsymbol{u} - \boldsymbol{v}\| \geqslant \epsilon - 2\eta$, the probability that $\mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u}, \boldsymbol{v}) \geqslant 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}) + \tfrac{\delta}{M}$ tends (exponentially) to one with M.

Let us upper bound the corresponding probability of failure. We can first observe the following result on a fixed pair of vectors. This one is proved in App. F.

Lemma 5. *Let $\boldsymbol{u}, \boldsymbol{v}$ be in $\mathbb{R}^N$ with $\boldsymbol{u} - \boldsymbol{v} \in \mathfrak{Z}_{K_0}$ for some $K_0 > 0$ and $\|\boldsymbol{u} - \boldsymbol{v}\| \leqslant \epsilon_0$ for $\epsilon_0 > 0$. For $\delta > 0$, $t \geqslant 0$, $r \in [\lfloor np \rfloor]$, $\boldsymbol{\Phi} \sim \mathcal{N}^{M\times N}_{\mathrm{sg},\alpha}(0,1)$, $\boldsymbol{\xi} \sim \mathcal{U}^M([0,\delta])$ and the pseudo-distance $\mathcal{D}^t$ defined in* (25)*, we have*

$$\mathbb{P}[\mathcal{D}^t(\boldsymbol{u}, \boldsymbol{v}) \leqslant \tfrac{\delta}{M} r] \leqslant \exp(-\tfrac{(Mp-r)^2}{2Mp}), \tag{51}$$

with $p := \mathbb{P}\big[d^t(\boldsymbol{\varphi}^\top \boldsymbol{u} + \xi, \boldsymbol{\varphi}^\top \boldsymbol{v} + \xi) \neq 0\big]$, $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$ and $\xi \sim \mathcal{U}([0,\delta])$. Moreover, if $\sqrt{K_0} \geqslant 16\kappa_{\rm sg}$,

$$p \geqslant \tfrac{1}{16(\delta+\epsilon_0)} \|\boldsymbol{u} - \boldsymbol{v}\| - \tfrac{2t}{\delta+\epsilon_0}. \tag{52}$$

From the discrete nature of $\mathcal{D}^t$, the previous lemma (with t set to $\eta\sqrt{P}$) shows that for a fixed pair of vectors $\mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u}, \boldsymbol{v}) \geqslant \frac{\delta}{M}(r+1)$ holds with probability at least $1 - \exp(-(Mp - r)^2/(2Mp))$. Moreover, if

$$\tfrac{\delta}{M} r \geqslant 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}), \tag{53}$$

we have

$$\mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u}, \boldsymbol{v}) \geqslant \tfrac{\delta}{M}(r+1) \ \Rightarrow\ \mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u}, \boldsymbol{v}) \geqslant 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}) + \tfrac{\delta}{M}.$$

Therefore, setting $r = \lceil Mp/2 \rceil \geqslant Mp/2$, (51) gives

$$\mathbb{P}\big[\mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u}, \boldsymbol{v}) \geqslant 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}) + \tfrac{\delta}{M}\big] \geqslant 1 - \exp(-\tfrac{(Mp-r)^2}{2Mp}) > 1 - 2\exp(-\tfrac{Mp}{8}),$$

if, from (53),

$$p \geqslant \tfrac{8}{\delta}(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}). \tag{54}$$

Thus, we have to adjust P and η in order to satisfy (54). Noting that $\epsilon - 2\eta \leqslant \|\boldsymbol{u} - \boldsymbol{v}\| \leqslant 2$ if $\mathcal{K} \subset \mathbb{B}^N$, *i.e.*, that we can set $\epsilon_0 = 2$ in Lemma 5, this adjustment can be done from (52) by imposing $B \geqslant C$ in

$$p \underset{\text{by (52)}}{\geqslant} B := \tfrac{1}{16(\delta+2)}(\epsilon - 2\eta) - \tfrac{2\eta\sqrt{P}}{\delta+2} \ \geqslant\ C := 8(\tfrac{1}{P} + \tfrac{\eta}{\delta\sqrt{P}}). \tag{55}$$

A solution is to set, for some $c \geqslant 1$ and $d > 0$ to be specified later, $P = c^2 \frac{2+\delta}{\epsilon} \geqslant 1$ and $\eta = d\frac{\epsilon^{3/2}}{\sqrt{2+\delta}} \leqslant d\epsilon$. Then

$$\epsilon - 2\eta \geqslant (1-2d)\epsilon, \quad \eta\sqrt{P} = cd\,\epsilon, \quad \tfrac{1}{P} = \tfrac{1}{c^2(2+\delta)}\epsilon, \quad \tfrac{\eta}{\delta\sqrt{P}} = \tfrac{d}{c}\tfrac{\epsilon^2}{\delta(2+\delta)} \leqslant \tfrac{d}{c}\tfrac{2}{\delta(2+\delta)}\,\epsilon,$$

so that

$$B \geqslant \tfrac{1-2d-32cd}{16(\delta+2)}\,\epsilon, \quad C \leqslant \tfrac{8}{c^2(2+\delta)}(1 + cd\tfrac{2}{\delta})\,\epsilon.$$

Fixing $d = \frac{1}{2}(32)^{-2}\frac{\delta}{\delta+2} < \frac{1}{2}(32)^{-2}$ and $c = 32$, a few estimations show finally that

$$\epsilon^{-1} B \geqslant \tfrac{1-(32)^{-2}-\frac{1}{2}}{16(\delta+2)} \geqslant \tfrac{1}{33(\delta+2)}, \quad \epsilon^{-1} C \leqslant \tfrac{8}{(32)^2(2+\delta)}(1 + \tfrac{2}{(64)(\delta+2)}) < \tfrac{1}{64(\delta+2)},$$

proving that for our choice of parameters, *i.e.*, for $P = (32)^2\,\frac{2+\delta}{\epsilon} \geqslant 1$ and $\eta = \frac{1}{2}(32)^{-2}\delta(\frac{\epsilon}{2+\delta})^{3/2}$, (54) can be satisfied since $B \geqslant C$. Moreover, for this choice of parameters, (54) provides

$$p \geqslant \tfrac{\epsilon}{33(2+\delta)}.$$

We are now ready to complete the proof. Using the previous developments, defining $\bar{\mathcal{G}}'_\eta := \{(\boldsymbol{u}, \boldsymbol{v}) \in \bar{\mathcal{G}}_\eta : \|\boldsymbol{u} - \boldsymbol{v}\| \geqslant \epsilon - 2\eta\} \subset \bar{\mathcal{G}}_\eta$ with $\eta \simeq \delta\epsilon^{3/2}(2+\delta)^{-3/2}$ fixed as above and $\log|\bar{\mathcal{G}}'_\eta| \leqslant$

$\log|\bar{\mathcal{G}}_\eta| \leqslant 2\mathcal{H}(\mathcal{K}, \eta/\sqrt{2})$ as explained before, by a simple union bound there exist some constants $C, c, c' > 0$ such that if

$$M \geqslant C\,\tfrac{2+\delta}{\epsilon}\,\mathcal{H}\big(\mathcal{K}, c\,\delta(\tfrac{\epsilon}{2+\delta})^{3/2}\big),$$

then the event

$$\mathcal{D}^{\eta\sqrt{P}}(\boldsymbol{u},\boldsymbol{v}) \geqslant 4(\tfrac{\delta}{P} + \tfrac{\eta}{\sqrt{P}}) + \tfrac{\delta}{M}, \quad \forall \boldsymbol{u},\boldsymbol{v} \in \bar{\mathcal{G}}'_\eta, \tag{56}$$

holds with probability at least

$$1 - 2\exp(\,2\mathcal{H}(\mathcal{K}, \tfrac{\eta}{\sqrt{2}}) - \tfrac{Mp}{8}) \;\geqslant\; 1 - 2\exp(\,2\mathcal{H}(\mathcal{K}, \tfrac{\eta}{\sqrt{2}}) - \tfrac{M\epsilon}{33(2+\delta)}) \;\geqslant\; 1 - 2\exp(-c'\,\tfrac{M\epsilon}{2+\delta}).$$

Remembering that for having (50) the diameter of $\mathcal{K}_\eta$ must remain small under random projections by $\boldsymbol{\Phi}$ (as stated in (49)), so by imposing (48), we find again by union bound that for some other constants $C, c, c' > 0$, if

$$M \geqslant C\max\bigg(\tfrac{(2+\delta)^3}{\delta^2\epsilon^3}w(\mathcal{K}_{c\,\delta(\frac{\epsilon}{2+\delta})^{3/2}})^2, \tfrac{2+\delta}{\epsilon}\,\mathcal{H}\big(\mathcal{K}, c\,\delta(\tfrac{\epsilon}{2+\delta})^{3/2}\big)\bigg), \tag{57}$$

then, with probability at least $1 - 4\exp(-c'M\epsilon/(2+\delta))$, for all $\boldsymbol{x},\boldsymbol{y} \in \mathcal{K}$ with $\boldsymbol{x}-\boldsymbol{y} \in \mho_{K_0}$ and $\|\boldsymbol{x}-\boldsymbol{y}\| \geqslant \epsilon$, (50) combined with (56) provides

$$\mathcal{D}(\boldsymbol{x},\boldsymbol{y}) \geqslant \tfrac{\delta}{M},$$

as requested at the beginning.

We conclude the proof by simplifying the general condition (57). First, for a general bounded set $\mathcal{K}$, Sudakov inequality (P. 14) and Sec. 6 provide $\mathcal{H}(\mathcal{K},\eta) \lesssim \frac{w(\mathcal{K})^2}{\eta^2}$ and $w(\mathcal{K}_\eta) \leqslant 2w(\mathcal{K})$, so that (57) holds if

$$M \geqslant C\,\tfrac{(2+\delta)^4}{\delta^2\epsilon^4}\,w(\mathcal{K})^2,$$

for another constant $C > 0$.

Second, if the set $\mathcal{K}$ is structured, then, from (3) and the same simplifications used for Prop. 3 to reach Prop. 1, the right-hand side of (57) can be bounded by

$$\begin{aligned}\max\big((\delta s)^{-2}w(\mathcal{K}_{c\,\delta s})^2, s^{-2/3}\,\mathcal{H}\big(\mathcal{K}, c\,\delta s)\big) &\leqslant \max\big(c^2\bar{w}(\mathcal{K})^2, \tfrac{2+\delta}{\epsilon}\bar{w}(\mathcal{K})^2\log(1+\tfrac{\|\mathcal{K}\|}{c\,\delta s})\big)\\ &\lesssim \tfrac{2+\delta}{\epsilon}\,\bar{w}(\mathcal{K})^2\log\big(1+\tfrac{(2+\delta)^{3/2}\|\mathcal{K}\|}{\delta\epsilon^{3/2}}\big),\end{aligned}$$

with $s := \epsilon^{3/2}/(2+\delta)^{3/2}$, which explains the requirement (20).

8 Acknowledgements

We wish to gladly thank Holger Rauhut and Sjoerd Dirksen for interesting and enlightening discussion on quantized random projections during a short stay end of January 2015 in RWTH Aachen University, and Jerry Veeh (Auburn University, AL, USA) for interesting discussions on the error bounds of the Stirling's approximation, as deduced in his lecture notes [30], and for having pointed out the work [40]. We also thank Valerio Cambareri (UCLouvain, Belgium) for interesting discussions on quantized embeddings and for his advices on the writing of this paper.

A On the absolute expectation of a difference of dithered floors

This short appendix proves the equality

$$\mathbb{E}|\lfloor x+\xi\rfloor - \lfloor y+\xi\rfloor| = |x-y|, \quad \forall x,y\in\mathbb{R},\ \xi\sim\mathcal{U}([0,1]).$$

Denoting $a=\lfloor x\rfloor\in\mathbb{Z}$, $b=\lfloor y\rfloor\in\mathbb{Z}$, $x'=x-a\in[0,1)$ and $y'=y-b\in[0,1)$, since $\lfloor\lambda-n\rfloor=\lfloor\lambda\rfloor-n$ for any $\lambda\in\mathbb{R}$ and $n\in\mathbb{Z}$, we can always write

$$\mathbb{E}|\lfloor x+\xi\rfloor - \lfloor y+\xi\rfloor| = \mathbb{E}|a-b+X|,$$

with $X=\lfloor x'+\xi\rfloor-\lfloor y'+\xi\rfloor$. Without loss of generality, we can assume that the r.v. X is positive, *i.e.*, $x'\geqslant y'$ (just flip the role of x and y if this is not the case). Moreover, since $x',y'\in[0,1)$, $X\in\{0,1\}$ and

$$\begin{aligned}\mathbb{P}(X=0) &= \mathbb{P}(x'+\xi<1, y'+\xi<1)+\mathbb{P}(x'+\xi\geqslant 1, y'+\xi\geqslant 1)\\ &= \mathbb{P}(x'+\xi<1)+\mathbb{P}(y'+\xi\geqslant 1) = 1-x'+y'.\end{aligned}$$

Therefore,

$$\begin{aligned}\mathbb{E}|a-b+X| &= (|a-b|-|a-b+1|)\,\mathbb{P}(X=0)+|a-b+1|\\ &= |a-b|-(x'-y')(|a-b|-|a-b+1|).\end{aligned} \tag{58}$$

If $x'=y'$, then $\mathbb{E}|a-b+X|=|a-b|=|x-y|$. Let us consider now the case $x'>y'$. If $x-y\geqslant 0$, then $a-b\geqslant y'-x'>-1$ since $x'<1$, *i.e.*, $a-b\geqslant 0$ since $a-b\in\mathbb{Z}$. Consequently, (58) provides $\mathbb{E}|a-b+X|=a-b+x'-y'=x-y$. When $x-y<0$, $b-a>x'-y'>0$, *i.e.*, $a-b\leqslant a-b+1\leqslant 0$, and we get $\mathbb{E}|a-b+X|=b-a-(x'-y')=x-y$. In summary, $\mathbb{E}|a-b+X|=|x-y|$ in all cases, which proves the result.

B Proof of Lemma 1

We start by observing that

$$\begin{aligned}\tfrac{1}{\delta}\left|d^t(a,b)-d^s(a,b)\right| &\leqslant \textstyle\sum_{k\in\mathbb{Z}}\left|\mathbb{I}[\mathcal{F}^t(a-k\delta,b-k\delta)]-\mathbb{I}[\mathcal{F}^s(a-k\delta,b-k\delta)]\right|\\ &\leqslant \textstyle\sum_{k\in\mathbb{Z}}\mathbb{I}[\mathcal{H}^{t,s}(a-k\delta,b-k\delta)]\end{aligned}$$

with

$$\mathcal{H}^{t,s}(a,b) \ :=\ \mathcal{F}^t(a,b)\triangle\mathcal{F}^s(a,b)\ :=\ \big(\mathcal{F}^t(a,b)\cup\mathcal{F}^s(a,b)\big)\setminus\big(\mathcal{F}^t(a,b)\cap\mathcal{F}^s(a,b)\big).$$

For $t\geqslant s$, $\mathcal{F}^t(a,b)\subset\mathcal{F}^s(a,b)$ and $\mathcal{H}^{t,s}(a,b)=\mathcal{F}^s(a,b)\setminus\mathcal{F}^t(a,b)$, while for $t<s$, $\mathcal{H}^{t,s}(a,b)=\mathcal{F}^t(a,b)\setminus\mathcal{F}^s(a,b)$. Moreover, a careful piecewise analysis made on the different sign combinations for s and t show that $\mathcal{H}^t(a,b)\subset\{|a|\in[r_-,r_+]\}\cup\{|b|\in[r_-,r_+]\}$ with $r_+:=\max(|s|,|t|)$ and r_- equals to $\min(|s|,|t|)$ if $ts\geqslant 0$ and 0 otherwise. Consequently, writing $r=r_+-r_-\leqslant|t-s|$,

$$\begin{aligned}\left|d^t(a,b)-d^s(a,b)\right| &\leqslant \delta\textstyle\sum_{k\in\mathbb{Z}}\mathbb{I}\big[\{|a-k\delta|\in[r_-,r_+]\}\cup\{|b-k\delta|\in[r_-,r_+]\}\big]\\ &\leqslant 2\delta(\tfrac{2r}{\delta}+2)=4(|t-s|+\delta).\end{aligned}$$

Moreover, if $s=0$, since then $r_-=0$ and $r_+=r=|t|$,

$$\textstyle\sum_{k\in\mathbb{Z}}\mathbb{I}\big[\{|a-k\delta|\leqslant|t|\}\cup\{|b-k\delta|\leqslant|t|\}\big]\leqslant 2\delta(\frac{2|t|}{\delta}+1)=4|t|+2\delta,$$

and we find

$$\begin{aligned}\left|d^t(a,b)-|a-b|\right| &\leqslant \left|d^t(a,b)-d(a,b)\right|+\left|d(a,b)-|a-b|\right|\\ &= \left|d^t(a,b)-d^0(a,b)\right|+\big||\mathcal{Q}(a)-\mathcal{Q}(b)|-|a-b|\big|\\ &\leqslant (4|t|+2\delta)+2\delta\ =4\,(|t|+\delta).\end{aligned}$$

C Proof of Lemma 2

Let us define $\tilde{Z} := |\boldsymbol{\varphi}^\top(\boldsymbol{x}-\boldsymbol{y})| = |a-b|$ with the two r.v.'s $a = \boldsymbol{\varphi}^\top \boldsymbol{x} + \xi$ and $b = \boldsymbol{\varphi}^\top \boldsymbol{x} + \xi$. From (13), $\mathbb{E}\tilde{Z} = \mathbb{E}Z^0$. Moreover, from the approximate rotational invariance property (33), $\tilde{Z}$ is sub-Gaussian with $\|\tilde{Z}\|_{\psi_2} = \|\boldsymbol{\varphi}^\top(\boldsymbol{x}-\boldsymbol{y})\|_{\psi_2} \lesssim \|\boldsymbol{x}-\boldsymbol{y}\|$, and using Lemma 1 and the bound $\|\cdot\|_{\psi_2} \leqslant \|\cdot\|_\infty$, we find

$$\begin{aligned}
\|Z^t\|_{\psi_2} &\leqslant \|Z^t - \tilde{Z}\|_{\psi_2} + \|\tilde{Z}\|_{\psi_2}\\
&\lesssim \|d^t(a,b) - |a-b|\|_{\psi_2} + \|\boldsymbol{x}-\boldsymbol{y}\|\\
&\lesssim \delta + |t| + \|\boldsymbol{x}-\boldsymbol{y}\|,
\end{aligned}$$

which demonstrates the sub-Gaussianity of Z^t.

For the expectation, writing $a = a' + \xi$ and $b = b' + \xi$ with $a' = \boldsymbol{\varphi}^\top \boldsymbol{x}$ and $b' = \boldsymbol{\varphi}^\top \boldsymbol{y}$, by Jensen's inequality and the law of total expectation, we find

$$\big|\mathbb{E}Z^t - EZ^0\big| \leqslant \mathbb{E}|Z^t - Z^0| = \mathbb{E}_{\boldsymbol{\varphi}}\mathbb{E}_\xi |d^t(a'+\xi, b'+\xi) - d(a'+\xi, b'+\xi)|.$$

However, reusing some elements of the proof of Lemma 1 and considering $\boldsymbol{\varphi}$ fixed,

$$\begin{aligned}
&\mathbb{E}_\xi\big|d^t(a'+\xi, b'+\xi) - d(a'+\xi, b'+\xi)\big|\\
&\leqslant \delta \textstyle\sum_{k\in\mathbb{Z}} \mathbb{E}_\xi \mathbb{I}\big[\{|a'+\xi-k\delta| \leqslant |t|\} \cup \{|b'+\xi-k\delta| \leqslant |t|\}\big]\\
&\leqslant \delta \textstyle\sum_{k\in\mathbb{Z}} \mathbb{E}_\xi \mathbb{I}\big[\{|a'+\xi-k\delta| \leqslant |t|\}\big] + \delta \textstyle\sum_{k\in\mathbb{Z}} \mathbb{E}_\xi \mathbb{I}\big[\{|b'+\xi-k\delta| \leqslant |t|\}\big].
\end{aligned}$$

Moreover, since $\xi \sim \mathcal{U}([0,\delta])$,

$$\begin{aligned}
\delta \textstyle\sum_{k\in\mathbb{Z}} \mathbb{E}_\xi \mathbb{I}\big[\{|a'+\xi-k\delta| \leqslant |t|\}\big] &= \textstyle\sum_{k\in\mathbb{Z}} \int_0^\delta \mathbb{I}\big[\{|a'+s-k\delta| \leqslant |t|\}\big]\,\mathrm{d}s\\
&= \textstyle\int_{\mathbb{R}} \mathbb{I}\big[\{|a'+s| \leqslant |t|\}\big]\mathrm{d}s \;=\; 2|t|,
\end{aligned}$$

which provides also $\delta \sum_{k\in\mathbb{Z}} \mathbb{E}_\xi \mathbb{I}\big[\{|b'+\xi-k\delta| \leqslant |t|\}\big] = 2|t|$. Consequently, since these two quantities do not depend on $\boldsymbol{\varphi}$, we find $\big|\mathbb{E}Z^t - EZ^0\big| \lesssim |t|$. Finally, if $\boldsymbol{\varphi} \sim \mathcal{N}^N(0,1)$, $Z^0 \sim \mathcal{N}(0, \|\boldsymbol{x}-\boldsymbol{y}\|^2)$, and $\mathbb{E}|Z^0| = (\frac{2}{\pi})^{1/2}\|\boldsymbol{x}-\boldsymbol{y}\|$.

D Proof of Lemma 3

We adapt the proof of Lemma 5.5 in [44] to both ℓ_2-perturbations (instead of ℓ_1 ones) of $\boldsymbol{x}_0$ and $\boldsymbol{y}_0$, and to the context of uniform dithered quantization instead of 1-bit (sign) quantization. By assumption, we have $\|\boldsymbol{\Phi}\boldsymbol{x}'\| \leqslant \eta\sqrt{M}$ and $\|\boldsymbol{\Phi}\boldsymbol{y}'\| \leqslant \eta\sqrt{M}$. Therefore, the set

$$T := \{i \in [M] : |(\boldsymbol{\Phi}\boldsymbol{x}')_i| \leqslant \eta\sqrt{P}, |(\boldsymbol{\Phi}\boldsymbol{y}')_i| \leqslant \eta\sqrt{P}\}$$

is such that $|T^{\mathrm{c}}| \leqslant 2M/P$ as $2\eta^2 M \geqslant \|\boldsymbol{\Phi}\boldsymbol{x}'\|^2 + \|\boldsymbol{\Phi}\boldsymbol{y}'\|^2 \geqslant \|(\boldsymbol{\Phi}\boldsymbol{x}')_T\|^2 + \|(\boldsymbol{\Phi}\boldsymbol{y}')_T\|^2 + |T^{\mathrm{c}}|P\eta^2 \geqslant |T^{\mathrm{c}}|P\eta^2$. Considering the definition of $\mathcal{F}^t$ in (26), we have, for all $i \in T$ and any $\lambda \in \mathbb{R}$,

$$\mathcal{F}_i^{t+\eta\sqrt{P}}(\boldsymbol{x}_0, \boldsymbol{y}_0, \lambda) \subset \mathcal{F}_i^t(\boldsymbol{x}_0 + \boldsymbol{x}', \boldsymbol{y}_0 + \boldsymbol{y}', \lambda) \subset \mathcal{F}_i^{t-\eta\sqrt{P}}(\boldsymbol{x}_0, \boldsymbol{y}_0, \lambda),$$

with $\mathcal{F}_i^t(\boldsymbol{x}_0, \boldsymbol{y}_0, \lambda) := \mathcal{F}^t(\boldsymbol{\varphi}_i^\top \boldsymbol{x}_0 + \xi_i - \lambda, \boldsymbol{\varphi}_i^\top \boldsymbol{y}_0 + \xi_i - \lambda)$.

Denoting $a_i = \max(|\boldsymbol{\varphi}_i^\top \boldsymbol{x}'|, |\boldsymbol{\varphi}_i^\top \boldsymbol{y}'|)$, we find

$$\begin{aligned}
&\mathcal{D}^{t+\eta\sqrt{P}}(\boldsymbol{x}_0, \boldsymbol{y}_0) = \tfrac{\delta}{M} \textstyle\sum_{i=1}^M \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}_i^{t+\eta\sqrt{P}}(\boldsymbol{x}_0, \boldsymbol{y}_0, k\delta)]\\
&\leqslant \tfrac{\delta}{M} \textstyle\sum_{i\in T} \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}_i^t(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)] + \tfrac{\delta}{M} \sum_{i\in T^{\mathrm{c}}} \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}_i^{t+\eta\sqrt{P}-a_i}(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)]\\
&\leqslant \tfrac{\delta}{M} \textstyle\sum_{i\in T} \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}_i^t(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)] + \tfrac{\delta}{M} \sum_{i\in T^{\mathrm{c}}} \sum_{k\in\mathbb{Z}} \mathbb{I}[\mathcal{F}_i^t(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)]\\
&\qquad + \tfrac{1}{M} \textstyle\sum_{i\in T^{\mathrm{c}}} \delta \sum_{k\in\mathbb{Z}} \big|\mathbb{I}[\mathcal{F}_i^{t+\eta\sqrt{P}-a_i}(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)] - \mathbb{I}[\mathcal{F}_i^t(\boldsymbol{x}_0+\boldsymbol{x}', \boldsymbol{y}_0+\boldsymbol{y}', k\delta)]\big|.
\end{aligned}$$

Using (29) to bound the last sum of the last expression and since, by definition of T, $a_i \geqslant \eta\sqrt{P}$ for $i \in T^{\mathsf{c}}$, we find

$$\begin{aligned}\mathcal{D}^{t+\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) &\leqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}') + \tfrac{4}{M}\textstyle\sum_{i\in T^{\mathsf{c}}}(\delta + a_i - \eta\sqrt{P})\\ &\leqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}') + \tfrac{4\delta}{P} + \tfrac{4}{M}\textstyle\sum_{i\in T^{\mathsf{c}}}(a_i - \eta\sqrt{P})\\ &\leqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}') + \tfrac{4\delta}{P} + \tfrac{4}{M}\textstyle\sum_{i\in T^{\mathsf{c}}} a_i - \tfrac{4|T^{\mathsf{c}}|}{M}\eta\sqrt{P}.\end{aligned}$$

However,

$$\tfrac{1}{M}\textstyle\sum_{i\in T^{\mathsf{c}}} a_i \leqslant \tfrac{1}{M}(\|(\boldsymbol{\Phi}\boldsymbol{x}')_{T^{\mathsf{c}}}\|_1 + \|(\boldsymbol{\Phi}\boldsymbol{y}')_{T^{\mathsf{c}}}\|_1) \leqslant \tfrac{\sqrt{|T^{\mathsf{c}}|}}{M}(\|(\boldsymbol{\Phi}\boldsymbol{x}')_{T^{\mathsf{c}}}\| + \|(\boldsymbol{\Phi}\boldsymbol{y}')_{T^{\mathsf{c}}}\|) \leqslant 2\eta\sqrt{\tfrac{|T^{\mathsf{c}}|}{M}},$$

and since $f(t) = 2t - t^2\sqrt{P} \leqslant 1/\sqrt{P}$ for all $t \in \mathbb{R}$, we find

$$\begin{aligned}\mathcal{D}^{t+\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) &\leqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}') + \tfrac{4\delta}{P} + 4\eta\,(2\sqrt{\tfrac{|T^{\mathsf{c}}|}{M}} - \tfrac{|T^{\mathsf{c}}|}{M}\sqrt{P})\\ &\leqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}') + \tfrac{4\delta}{P} + 4\tfrac{\eta}{\sqrt{P}},\end{aligned}$$

which provides the lower bound of (40).

For the upper bound,

$$\begin{aligned}&\mathcal{D}^{t-\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0) = \tfrac{\delta}{M}\textstyle\sum_{i=1}^M\sum_{k\in\mathbb{Z}}\mathbb{I}[\mathcal{F}_i^{t-\eta\sqrt{P}}(\boldsymbol{x}_0,\boldsymbol{y}_0,k\delta)]\\ &\geqslant \tfrac{\delta}{M}\textstyle\sum_{i\in T}\sum_{k\in\mathbb{Z}}\mathbb{I}[\mathcal{F}_i^{t}(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}',k\delta)] + \tfrac{\delta}{M}\textstyle\sum_{i\in T^{\mathsf{c}}}\sum_{k\in\mathbb{Z}}\mathbb{I}[\mathcal{F}_i^{t-\eta\sqrt{P}+a_i}(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}',k\delta)]\\ &\geqslant \mathcal{D}^t(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}')\\ &\qquad - \tfrac{1}{M}\textstyle\sum_{i\in T^{\mathsf{c}}}\delta\sum_{k\in\mathbb{Z}}\big|\mathbb{I}[\mathcal{F}_i^{t}(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}',k\delta)] - \mathbb{I}[\mathcal{F}_i^{t-\eta\sqrt{P}+a_i}(\boldsymbol{x}_0+\boldsymbol{x}',\boldsymbol{y}_0+\boldsymbol{y}',k\delta)]\big|,\end{aligned}$$

and, as above, the last sum can be upper-bounded by $\frac{4\delta}{P} + \frac{4\eta}{\sqrt{P}}$ using (29).

E Proof of Lemma 4

We use here a similar proposition of Mendelson[9] *et al.* in [39] for subsets of $\mathbb{S}^{N-1}$ that we lift to subsets of $\mathbb{R}^{N+1}$ thank to some tools developed in [44] for other purposes.

We fix $t = \|\mathcal{R}\|/\sqrt{6}$ and form the set $\mathcal{R}' := \{\boldsymbol{u}/\|\boldsymbol{u}\| : \boldsymbol{u} \in \mathcal{R}\oplus t\}$ with $\mathcal{R}\oplus t := \{\binom{\boldsymbol{x}}{t} : \boldsymbol{x} \in \mathcal{R}\} \subset \mathbb{R}^{N+1}$. As $\mathcal{R}' \subset \mathbb{S}^N$, we know from [39, Theorem 2.1] that for $0 < \epsilon < 1$,

$$M \ \gtrsim\ \tfrac{\alpha^4}{\epsilon^2} w(\mathcal{R}')^2$$

and $\boldsymbol{\Phi}' \sim \mathcal{N}_{\mathrm{sg},\alpha}^{M\times(N+1)}(0,1)$,

$$\mathbb{P}\big[\sup_{\boldsymbol{x}'\in\mathcal{R}'} |\tfrac{1}{M}\|\boldsymbol{\Phi}'\boldsymbol{x}'\|^2 - 1| \geqslant \epsilon\big] \leqslant \exp(-c\tfrac{\epsilon^2 M}{\alpha^4}).$$

However, for $\boldsymbol{g} \sim \mathcal{N}^N(0,1)$ and $\gamma \sim \mathcal{N}(0,1)$, as observed similarly in [44],

$$\begin{aligned}w(\mathcal{R}') = \mathbb{E}\sup_{\boldsymbol{x}\in\mathcal{R}}(\|\boldsymbol{x}\|^2 + t^2)^{-1/2}|\langle\boldsymbol{g},\boldsymbol{x}\rangle + t\gamma| \ &\leqslant \tfrac{1}{t}\,(\mathbb{E}\sup_{\boldsymbol{x}\in\mathcal{R}}|\langle\boldsymbol{g},\boldsymbol{x}\rangle| + t(\tfrac{2}{\pi})^{1/2})\\ \leqslant \tfrac{\sqrt{6}}{\|\mathcal{R}\|}w(\mathcal{R}) + (\tfrac{2}{\pi})^{1/2} \ &\leqslant\ 4\tfrac{w(\mathcal{R})}{\|\mathcal{R}\|},\end{aligned}$$

[9] Where a totally equivalent sub-Gaussian norm is used, *i.e.*, $\|X\|_{\psi_2}^{(\mathrm{Mend.})} := \inf\{s : \mathbb{E}\exp(X^2/s^2) \leqslant 2\}$ with $\|X\|_{\psi_2}^{(\mathrm{Mend.})} \simeq \|X\|_{\psi_2}$ [52].

since, for all $\boldsymbol{x} \in \mathcal{R}$, $w(\mathcal{R}) \geqslant (\frac{2}{\pi})^{1/2}\|\boldsymbol{x}\|$, *i.e.*, $w(\mathcal{R}) \geqslant (\frac{2}{\pi})^{1/2}\|\mathcal{R}\|$. Therefore, fixing $\epsilon = 1/2$, if $M \gtrsim \alpha^4 w(\mathcal{R})^2/\|\mathcal{K}\|^2$, with probability at least $1 - e^{-c\alpha^{-4}M}$, we have, for all $\boldsymbol{x} \in \mathcal{R}$,

$$\sqrt{\tfrac{3}{2}} \geqslant \tfrac{1}{\sqrt{M}}\|\boldsymbol{\Phi}'\boldsymbol{x}'\| \geqslant \tfrac{1}{t\sqrt{M}}\|\boldsymbol{\Phi}\boldsymbol{x} + t\boldsymbol{\phi}\| \ \geqslant \tfrac{1}{t\sqrt{M}}(\|\boldsymbol{\Phi}\boldsymbol{x}\| - t\|\boldsymbol{\Phi}'(\begin{smallmatrix}\mathbf{0}\\1\end{smallmatrix})\|) \ \geqslant \tfrac{1}{t\sqrt{M}}\|\boldsymbol{\Phi}\boldsymbol{x}\| - \sqrt{\tfrac{3}{2}},$$

where $\boldsymbol{x}' = \|(\begin{smallmatrix}\boldsymbol{x}\\t\end{smallmatrix})\|^{-1}(\begin{smallmatrix}\boldsymbol{x}\\t\end{smallmatrix}) \in \mathcal{R}'$, $\boldsymbol{\phi} \in \mathbb{R}^M$ is the last column of $\boldsymbol{\Phi}'$ and using the fact that $(\begin{smallmatrix}\mathbf{0}\\1\end{smallmatrix}) \in \mathcal{R}'$ since $0 \in \mathcal{R}$. Therefore, replacing t by its value, we find with the same probability,

$$\tfrac{1}{\sqrt{M}}\|\boldsymbol{\Phi}\boldsymbol{x}\| \leqslant \|\mathcal{R}\|,$$

for all $\boldsymbol{x} \in \mathcal{R}$, *i.e.*, $\|\boldsymbol{\Phi}\mathcal{R}\| \leqslant \sqrt{M}\|\mathcal{R}\|$.

F Proof of Lemma 5

From the relation $\mathcal{D}^t(\boldsymbol{u},\boldsymbol{v}) = \frac{1}{M}\sum_{i=1}^M d^t(\Phi_i^\xi(\boldsymbol{u}), \Phi_i^\xi(\boldsymbol{v}))$ established in Sec. 6 between $\mathcal{D}^t$ and $d^t \in \delta\mathbb{N}$ defined in (27), and associated to the vectorial mapping $\boldsymbol{u} \in \mathbb{R}^N \to \boldsymbol{\Phi}^\xi(\boldsymbol{u}) = \boldsymbol{\Phi}\boldsymbol{u} + \boldsymbol{\xi}$ whose components are independent, we reach the bound (51) with the cdf of a binomial distribution: since

$$\begin{aligned}\mathbb{P}\big[\tfrac{M}{\delta}\mathcal{D}^t(\boldsymbol{u},\boldsymbol{v}) \leqslant r\big] &\leqslant \mathbb{P}\big[\big|\{j \in [M] :\ d^t(\Phi_i^\xi(\boldsymbol{u}), \Phi_i^\xi(\boldsymbol{v})) \neq 0\}\big| \leqslant r\big]\\ &= \textstyle\sum_{k=0}^r \binom{M}{k} p^k (1-p)^{M-k},\end{aligned}$$

Chernoff's inequality can upper bound this binomial cdf with

$$\mathbb{P}\big[\tfrac{M}{\delta}\mathcal{D}^t(\boldsymbol{u},\boldsymbol{v}) \leqslant r\big] \leqslant \exp(-\tfrac{(Mp-r)^2}{2Mp}). \tag{59}$$

Let us now lower bound p. Defining $\boldsymbol{w} = \boldsymbol{u} - \boldsymbol{v} \in \Sigma_{K_0}$ and $\hat{\boldsymbol{w}} = \boldsymbol{w}/\|\boldsymbol{w}\|$, the action of dithering $\xi \sim \mathcal{U}([0,\delta])$ allows us to compute easily that,

$$p = \mathbb{E}_{\boldsymbol{\varphi}}\,\mathbb{P}_\xi\big[d^t(\boldsymbol{\varphi}^\top\boldsymbol{u} + \xi, \boldsymbol{\varphi}^\top\boldsymbol{v} + \xi) \neq 0\big] = \mathbb{E}\min\big(1, \delta^{-1}(|\boldsymbol{\varphi}^\top\boldsymbol{w}| - 2t)_+\big).$$

In order to avoid any further singularity when $\delta \to 0$, we can benefit from the fact that $p \geqslant 1$ and work with this slightly looser bound:

$$p \geqslant \mathbb{E}\min\big(1, (\epsilon_0 + \delta)^{-1}(|\boldsymbol{\varphi}^\top\boldsymbol{w}| - 2t)_+\big).$$

Moreover, with $\alpha = \|\boldsymbol{u} - \boldsymbol{v}\|/(\delta + \epsilon_0)$,

$$p \geqslant \mathbb{E}\min(1, \alpha|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}| - \tfrac{2t}{\delta+\epsilon_0}) \geqslant \mathbb{E}\min(1, \alpha|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}|) - \tfrac{2t}{\delta+\epsilon_0},$$

so that

$$p \ \geqslant \ \mathbb{E}\min(1, \alpha|g|) - \tfrac{2t}{\delta+\epsilon_0} - A, \tag{60}$$

where $g \sim \mathcal{N}(0,1)$ and $A := |\mathbb{E}\min(1, \alpha|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}|) - \mathbb{E}\min(1, \alpha|g|)|$.

We can upper bound A from our assumptions on the sub-Gaussian vector $\boldsymbol{\varphi} \sim \mathcal{N}^N_{\mathrm{sg},\alpha}(0,1)$:

$$\begin{aligned}A = &\ \big|\textstyle\int_0^1 \mathbb{P}(\min(1, \alpha|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}|) \geqslant u) - \mathbb{P}(\min(1, \alpha|g|) \geqslant u)\,\mathrm{d}u\,\big|\\ = &\ \big|\textstyle\int_0^1 \mathbb{P}(\alpha\,|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}| \geqslant u) - \mathbb{P}(\alpha|g| \geqslant u)\,\mathrm{d}u\,\big|\\ \leqslant &\ \alpha\textstyle\int_0^{+\infty}\big|\mathbb{P}(|\boldsymbol{\varphi}^\top\hat{\boldsymbol{w}}| \geqslant u) - \mathbb{P}(|g| \geqslant u)\big|\,\mathrm{d}u\\ \leqslant &\ \tfrac{\kappa_{\mathrm{sg}}}{\delta+\epsilon_0}\,\|\boldsymbol{w}\|_\infty \ \leqslant \ \tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}}\,\alpha,\end{aligned}$$

where the last inequalities rely on assumption (7) (setting $\boldsymbol{u} = \hat{\boldsymbol{w}}$) and on the fact that $\boldsymbol{w} \in \mathcal{Z}_{K_0}$.

Moreover, for lower-bounding $\mathbb{E}\min(1, \alpha|g|)$ in (60), we observe that $\min(1, \alpha x) = \alpha x - \alpha(x - 1/\alpha)_+$ for $x \in \mathbb{R}$. Therefore, defining $F(x) := \frac{1}{2}\alpha x^2 - \frac{1}{2}\alpha(x - 1/\alpha)_+^2 = \int_0^x \min(1, \alpha u)\mathrm{d}u$ and integrating by parts, we find

$$\mathbb{E}\min(1, \alpha|g|) = \mathbb{E}(|g| F(|g|)) \geqslant (\tfrac{2}{\pi})^{1/2} F((\tfrac{2}{\pi})^{1/2})$$

where in the last inequality we used Jensen's inequality and the convexity of $x \in \mathbb{R}_+ \mapsto xF(x)$. It is easy to see that $2F(x) \geqslant \alpha x^2/(1 + \alpha x)$ so that

$$\mathbb{E}\min(1, \alpha|g|) \geqslant \tfrac{1}{2}(\tfrac{2}{\pi})^{1/2} \frac{\frac{2}{\pi}\alpha}{1 + (\frac{2}{\pi})^{1/2}\alpha} \geqslant \tfrac{1}{4}\tfrac{\alpha}{1+\alpha}.$$

Finally,

$$p \geqslant \tfrac{1}{4}\tfrac{\alpha}{1+\alpha} - \tfrac{2t}{\delta+\epsilon_0} - \tfrac{\alpha\kappa_{\mathrm{sg}}}{\sqrt{K_0}} \geqslant \tfrac{1}{4}\tfrac{1}{\delta+2\epsilon_0}\|\boldsymbol{u} - \boldsymbol{v}\| - \tfrac{2t}{\delta+\epsilon_0} - \tfrac{\alpha\kappa_{\mathrm{sg}}}{\sqrt{K_0}} \geqslant \tfrac{1}{\delta+\epsilon_0}(\tfrac{1}{8} - \tfrac{\kappa_{\mathrm{sg}}}{\sqrt{K_0}})\|\boldsymbol{u} - \boldsymbol{v}\| - \tfrac{2t}{\delta+\epsilon_0},$$

the last expression providing (52) if $\sqrt{K_0} \geqslant 16\kappa_{\mathrm{sg}}$.

G A lower bound on the approximation error of the Mean Absolute Difference of a binomial random variable

This small section establishes a lower bound on the approximation error of the MAD $M_n := \mathbb{E}|\beta_n - \mathbb{E}\beta_n|$ of a binomial random variable $\beta_n \sim \mathrm{Bin}(n, 1/2)$ by a fraction of its standard deviation $\sigma_n := (\mathbb{E}|\beta_n - \mathbb{E}\beta_n|^2)^{1/2} = \sqrt{n}/2$. Curiously enough, we were unable to find a similar result in the literature while an upper bound on this approximation error in $O(1/n)$ when n increases is well known (see *e.g.*, [8, 18]). Specifically, we want to prove that

$$|M_n - (\tfrac{2}{\pi})^{1/2}\sigma_n| \;\geqslant\; C\sigma_n\, n^{-1},$$

for some absolute constant $C > 0$ and all $n \geqslant 1$.

We start from the Stirling's approximation of the factorial with an error bound due to R. W. Gosper [53] and redeveloped more clearly in [30] (see also [40] for a similar bound):

$$n^n e^{-n}\sqrt{2\pi(n + \tfrac{1}{6})} \;\leqslant\; n! \;\leqslant\; n^n e^{-n}\sqrt{2\pi(n + \tfrac{1}{5})}. \tag{61}$$

However, De Moivre gave the following exact formula for M_{2n} [18],

$$M_{2n} := n\, 2^{-2n}\tbinom{2n}{n} = n\, 2^{-2n}\tfrac{(2n)!}{(n!)^2}.$$

Therefore, applying (61) on this formula and using $\sqrt{1+x} \leqslant 1 + \frac{1}{2}x$ for $x \geqslant -1$, we find for $n \geqslant 1$

$$M_{2n} \leqslant n\, 2^{-2n}\,\frac{(2n)^{2n}e^{-2n}\sqrt{2\pi(2n+\frac{1}{5})}}{n^{2n}e^{-2n}2\pi(n+\frac{1}{6})} \;=\; \frac{n\sqrt{n+\frac{1}{10}}}{\sqrt{\pi}(n+\frac{1}{6})} \;<\; (\tfrac{2}{\pi})^{1/2}\sigma_{2n}\sqrt{\frac{n}{n+\frac{1}{6}}}$$

$$= (\tfrac{2}{\pi})^{1/2}\sigma_{2n}\sqrt{1 - \tfrac{1}{6n+1}} \leqslant (\tfrac{2}{\pi})^{1/2}\sigma_{2n}\,(1 - \tfrac{1}{12n+2}) \leqslant (\tfrac{2}{\pi})^{1/2}\sigma_{2n}\,(1 - \tfrac{1}{14n}),$$

or equivalently

$$(\tfrac{2}{\pi})^{1/2}\sigma_{2n} - M_{2n} \geqslant C\,\sigma_{2n}\,(2n)^{-1},$$

with $C = 1/7$, which provides the result.

References

[1] D. Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of computer and System Sciences*, 66(4):671–687, 2003.

[2] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin. One-bit compressed sensing with non-gaussian measurements. *Linear Algebra and its Applications*, 441:222–239, 2014.

[3] A. Andoni, M. Datar, N. Immorlica, P. Indyk, and V. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. *Nearest neighbor methods in learning and vision: Theory and practice (book)*, 2006.

[4] S. Bahmani, P. T Boufounos, and B. Raj. Robust 1-bit Compressive Sensing via Gradient Support Pursuit. *arXiv preprint arXiv:1304.6627*, 2013.

[5] A. Bandeira, D. G Mixon, and B. Recht. Compressive classification and the rare eclipse problem. *arXiv preprint arXiv:1404.3203*, 2014.

[6] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

[7] R. Baraniuk and M. Wakin. Random projections of smooth manifolds. *Foundations of Computational mathematics*, 9(1):51–77, 2009.

[8] C. R Blyth. Expected absolute error of the usual estimator of the binomial parameter. *The American Statistician*, 34(3):155–157, 1980.

[9] P. T. Boufounos. Universal rate-efficient scalar quantization. *IEEE Trans. Info. Theory*, 58(3):1861–1872, March 2012.

[10] P. T. Boufounos and S. Rane. Efficient coding of signal distances using universal quantized embeddings. In *Proc. Data Compression Conference (DCC)*, Snowbird, UT, March 20-22 2013.

[11] E. J Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.

[12] V. Chandrasekaran and M. I. Jordan. Computational and statistical tradeoffs via convex relaxation. *arXiv preprint arXiv:1211.1073*, 2012.

[13] V. Chandrasekaran, B. Recht, P. A Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.

[14] W. Dai, H. V. Pham, and O. Milenkovic. Distortion-Rate Functions for Quantized Compressive Sensing. Technical Report arXiv:0901.0749, 2009.

[15] S. Dasgupta and A. Gupta. An elementary proof of the Johnson-Lindenstrauss Lemma. Technical Report TR-99-006, International Computer Science Institute, Berkeley, CA, 1999.

[16] M. A. Davenport, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. A simple proof that random matrices are democratic. Technical report, Rice University ECE Department Technical Report TREE-0906, Houston, TX, November 2009.

[17] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk. Signal processing with compressive measurements. *Selected Topics in Signal Processing, IEEE Journal of*, 4(2):445–460, 2010.

[18] P. Diaconis and S. Zabell. Closed Form Summation for Classical Distributions: Variations on a Theme of De Moivre *Statistical Science*, 6(3):284–302, 1991.

[19] S. Dirksen. Dimensionality reduction with subgaussian matrices: a unified theory. *arXiv preprint arXiv:1402.3973*, 2014.

[20] S. Dirksen, G. Lecué, and H. Rauhut. On the gap between rip-properties and sparse recovery conditions. *arXiv preprint arXiv:1504.05073*, 2015.

[21] D. L. Donoho. Compressed Sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

[22] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing.* Springer, 2013.

[23] V. K Goyal, M. Vetterli, and N. T. Thao. Quantized overcomplete expansions in $\mathbb{R}^N$: Analysis, synthesis, and algorithms. *IEEE Trans. Info. Theory*, 44(1):16–31, 1998.

[24] R.M. Gray and D.L. Neuhoff. Quantization. *IEEE Trans. Info. Theory*, 44(6):2325–2383, Oct 1998.

[25] L. Jacques, D. K. Hammond, and M. J. Fadili. Dequantizing Compressed Sensing: When Oversampling and Non-Gaussian Constraints Combine. *IEEE Trans. Inf. Theory*, 57(1):559–571, January 2011.

[26] L. Jacques, D. K. Hammond, and M. J. Fadili. Stabilizing Nonuniformly Quantized Compressed Sensing With Scalar Companders. *IEEE Trans. Inf. Theory*, 59(12):7969 – 7984, January 2013.

[27] L. Jacques. A Quantized Johnson Lindenstrauss Lemma: The Finding of Buffon's Needle. *IEEE Trans. Inf. Theory*, 61(9):5012 – 5027, September 2015

[28] L. Jacques. Error decay of (almost) consistent signal estimations from quantized random gaussian projections. *IEEE Trans. Inf. Theory*, 62(8):4696 – 4709, August 2016

[29] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, 2013.

[30] Jerry A. Veeh. Stirling's Formula with Error Bounds Lecture notes, 2014, `https://javeeh.net/lecnotes/stirling.pdf`

[31] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[32] M. Kabanava, R. Kueng, H. Rauhut, U. Terstiege, Stable low-rank matrix recovery via null space properties. *arXiv preprint arXiv:1507.07184*, 2015.

[33] K. Knudson, R. Saab, and R. Ward. One-bit compressive sensing with norm estimation. *arXiv preprint arXiv:1404.6853*, 2014.

[34] D. E. Knuth. Big omicron and big omega and big theta. *ACM Sigact News*, 8(2):18–24, 1976.

[35] A. N. Kolmogorov and V. M. Tikhomirov. ε-entropy and ε-capacity of sets in function spaces. *Amer. Math. Soc. Transl. Ser.*, 17(2):277–364, 1961. (Russian version) Uspekhi Mat. Nauk , 14 : 2 (1959) pp. 386.

[36] F. Krahmer and R. Ward. A unified framework for linear dimensionality reduction in l1. *arXiv preprint arXiv:1405.1332*, 2014.

[37] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk. Democracy in action: Quantization, saturation, and compressive sensing. *App. Comp. and Harm. Anal.*, 31(3):429–443, November 2011.

[38] O. Maillard and R. Munos. Linear regression with random projections. *Journal of Machine learning Research*, 13:2735–2772, 2012.

[39] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Uniform uncertainty principle for bernoulli and subgaussian ensembles. *Constructive Approximation*, 28(3):277–289, 2008.

[40] C. Mortici On Gospers formula for the gamma function *Journal of Mathematical Inequalities*, 5(4):611–614, 2011.

[41] S. Oymak. Near-Optimal Sample Complexity Bounds for Circulant Binary Embedding. *arXiv preprint arXiv:1603.03178*, 2016.

[42] S. Oymak and B. Recht. Near-Optimal Bounds for Binary Embeddings of Arbitrary Sets. *arXiv preprint arXiv:1512.04433*, 2015.

[43] G. Pisier. The volume of convex bodies and Banach space geometry. Volume 94. Cambridge University Press, 1999.

[44] Y. Plan and R. Vershynin. Dimension reduction by random hyperplane tessellations. *arXiv preprint arXiv:1111.4452*, 2011.

[45] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory, to appear.*, 2012.

[46] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66(8):1275–1297, 2013.

[47] A. M. Powell and J. T. Whitehouse. Error bounds for consistent reconstruction: Random polytopes and coverage processes. *arXiv preprint arXiv:1405.7094*, 2013.

[48] S. Rane, P. T. Boufounos, and A. Vetro. Quantized embeddings: An efficient and universal nearest neighbor method for cloud-based image retrieval. In *Proc. SPIE Applications of Digital Image Processing XXXVI*, San Diego, CA, August 25-19 2013.

[49] G. Schechtman. Two observations regarding embedding subsets of Euclidean spaces in normed spaces. *Advances in Mathematics*, 200(1):125135, 2006.

[50] N. T. Thao and M. Vetterli. Lower bound on the mean-squared error in oversampled quantization of periodic signals using vector quantization analysis. *IEEE Trans. Info. Theory*, 42(2):469–479, March 1996.

[51] A. W. Vaart and J. A. Wellner. *Weak convergence and empirical processes.* Springer, 1996.

[52] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *In book "Compressed Sensing Theory and Applications"*, 2012. arXiv preprint arXiv:1011.3027.

[53] W. R. Gosper. Decision procedure for indefinite hypergeometric summation. *Proceedings of the National Academy of Sciences*, 75(1):40-42, 1978.

[54] A. Zymnis, S. Boyd, and E. Candes. Compressed sensing with quantized measurements. *Signal Processing Letters, IEEE*, 17(2):149–152, 2010.