The Rare Eclipse Problem in Quantised Random Embeddings: a Matter of Consistency?

Valerio Cambareri, Chunlei Xu and Laurent Jacques
ISPGroup, ICTEAM/ELEN, Université catholique de Louvain, Louvain-la-Neuve, Belgium.
E-mail: {valerio.cambareri, chunlei.xu, laurent.jacques}@uclouvain.be

***Abstract*—We study the problem of verifying when two disjoint closed convex sets remain separable after the application of a quantised random embedding, as a means to ensure exact classification from the signatures produced by this non-linear dimensionality reduction. An analysis of the interplay between the embedding, its quantiser resolution and the sets' separation is presented in the form of a convex problem; this is completed by its numerical exploration in a special case, for which the phase transition corresponding to exact classification is easily computed.**

I. Problem Statement

Non-linear dimensionality reduction techniques play an important role in simplifying statistical learning on very large-scale datasets. Among such techniques, we focus on *quantised random embeddings* obtained by a non-linear map A applied to $\boldsymbol{x} \in \mathcal{K} \subset \mathbb{R}^n$, that is

$$\boldsymbol{y} = \mathsf{A}(\boldsymbol{x}) := \mathcal{Q}_\delta(\boldsymbol{\Phi x} + \boldsymbol{\xi}) \tag{1}$$

with $\boldsymbol{\Phi} \in \mathbb{R}^{m \times n}$ a random *sensing matrix*, $\mathcal{Q}_\delta(\cdot) := \delta \lfloor \frac{\cdot}{\delta} \rfloor$ a *uniform scalar quantiser* of *resolution* $\delta > 0$ (applied component-wise), and the *signature* $\boldsymbol{y} \in \delta\mathbb{Z}^m$. In (1), the *dithering*[1] $\boldsymbol{\xi} \sim \mathcal{U}^m([0, \delta])$ is a well-known means to stabilise the action of the quantiser [1], [2].

The non-linear map (1) is a non-adaptive dimensionality reduction that yields compact signatures for storage and transmission, while retaining a notion of *quasi-isometry* that enables the approximation of $\boldsymbol{x}$ [2], [3]. Consequently, distance-based learning tasks preserve their accuracy if run on $\mathsf{A}(\mathcal{K})$ rather than $\mathcal{K}$, provided some requirements are met on m, δ, the distribution of $\boldsymbol{\Phi}$ and the "dimension" of $\mathcal{K}$ as measured, *e.g.*, by its *Gaussian mean width* $w(\mathcal{K}) := \sup_{\boldsymbol{x} \in \mathcal{K}} |\boldsymbol{g}^\top \boldsymbol{x}|$ with $\boldsymbol{g} \sim \mathcal{N}^n(0, 1)$ (see, *e.g.*, [2]). In this context we aim to show that, given two *classes* described by some sets $\mathcal{C}_1, \mathcal{C}_2 \subset \mathcal{K} : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ and $\boldsymbol{x} \in \mathcal{C}_1 \cup \mathcal{C}_2 \subset \mathcal{K}$, classifying whether $\boldsymbol{x}$ belongs to $\mathcal{C}_1$ or $\mathcal{C}_2$ is still possible from $\boldsymbol{y} = \mathsf{A}(\boldsymbol{x})$. For linear embeddings such as $\boldsymbol{y} = \boldsymbol{\Phi x}$, Bandeira *et al.* [4] approach the above classification problem as follows.

Problem 1 (Rare Eclipse Problem (from [4]))**.** *Let* $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ *be closed convex sets,* $\boldsymbol{\Phi} \sim \mathcal{N}^{m \times n}(0, 1)$*. Given* $\eta \in (0, 1)$*, find the smallest* m *so that* $p_0 := \mathbb{P}[\boldsymbol{\Phi}\mathcal{C}_1 \cap \boldsymbol{\Phi}\mathcal{C}_2 = \emptyset] \geq 1 - \eta$.

Prob. 1 amounts to ensuring for all $\boldsymbol{x}' \in \mathcal{C}_1, \boldsymbol{x}'' \in \mathcal{C}_2$ that their images $\boldsymbol{\Phi x}' \neq \boldsymbol{\Phi x}''$. Using the *difference set* $\mathcal{C}^- := \mathcal{C}_1 - \mathcal{C}_2 = \{\boldsymbol{z} := \boldsymbol{x}' - \boldsymbol{x}'' : \boldsymbol{x}' \in \mathcal{C}_1, \boldsymbol{x}'' \in \mathcal{C}_2\}$ we see the above problem equals

$$\mathbb{P}[\forall \boldsymbol{z} \in \mathcal{C}^-, \boldsymbol{\Phi z} \neq \boldsymbol{0}_m] = 1 - \mathbb{P}[\exists \boldsymbol{z} \in \mathcal{C}^- : \boldsymbol{\Phi z} = \boldsymbol{0}_m] \geq 1 - \eta.$$

This requires a bound on the probability that the kernel of $\boldsymbol{\Phi}$ "collides" with $\mathcal{C}^-$, *i.e.*, $\mathbb{P}[\mathrm{Ker}(\boldsymbol{\Phi}) \cap \mathcal{C}^- \neq \emptyset] \leq \eta$, and [4] shows that η is small if m is large compared to the "dimension" of $\mathcal{C}^-$ as measured by $w_\cap^2 := w^2\big((\mathbb{R}_+\mathcal{C}^-) \cap \mathbb{S}^{n-1}\big)$ with $\mathbb{R}_+\mathcal{C}^-$ the cone generated by $\mathcal{C}^-$.

From this standpoint, extending such existing results on Prob. 1 to non-linear maps as (1) is non-trivial. Applying A to each closed convex set $\mathcal{C}_1, \mathcal{C}_2$ would produce two countable sets $\mathsf{A}(\mathcal{C}_1), \mathsf{A}(\mathcal{C}_2) \subset \delta\mathbb{Z}^m$, and assessing if they still "collide" is our key question below.

Problem 2 (Quantised Eclipse Problem)**.** *In the setup of Prob. 1, given* $\eta \in (0, 1)$*, find the smallest* m *so that* $\mathbb{P}[\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset] \geq 1 - \eta$, i.e.*,* $p_\delta := \mathbb{P}[\forall \boldsymbol{x}' \in \mathcal{C}_1, \boldsymbol{x}'' \in \mathcal{C}_2, \mathsf{A}(\boldsymbol{x}') \neq \mathsf{A}(\boldsymbol{x}'')] \geq 1 - \eta$.

Note that the event in Prob. 2 requires $\mathbb{P}[\exists\ \boldsymbol{x}' \in \mathcal{C}_1, \boldsymbol{x}'' \in \mathcal{C}_2 : \mathsf{A}(\boldsymbol{x}') = \mathsf{A}(\boldsymbol{x}'')] \leq \eta$, *i.e.*, a bound on the probability of existence of two *consistent* vectors (through the mapping A) that do not belong to the same set.

The authors are funded by the F.R.S.-FNRS and by the project ALTERSENSE (MIS-FNRS). All authors have equally contributed to the realisation of this work.

[1] A matrix denoted by $\boldsymbol{M} \sim X^{d_1 \times d_2}$ has entries $M_{ij} \sim_{\text{i.i.d.}} X$ for a r.v. X.

We here leverage the quantised restricted isometry property (QRIP) introduced in [2] to estimate η and the conditions on m. The QRIP establishes some conditions on m that ensure $\frac{1}{m}\|\mathsf{A}(\boldsymbol{x}') - \mathsf{A}(\boldsymbol{x}'')\|_1 \geq (c' - \varepsilon)\|\boldsymbol{z}\| - c\delta\varepsilon' \geq (c' - \varepsilon)\sigma - c\delta\varepsilon' =: H$ for some controllable distortions $\varepsilon, \varepsilon' > 0$, with $\sigma \leq \|\boldsymbol{z}\|$ and some constants $c, c' > 0$. Thus $\mathsf{A}(\boldsymbol{x}') \neq \mathsf{A}(\boldsymbol{x}'')$ if $H > 0$. In particular, we deduce the following proposition whose proof is postponed to an extended version of this work.

Proposition 1. *In the setup of Prob. 2, let* $r_i := \mathrm{rad}(\mathcal{C}_i)$*,* $i \in \{1, 2\}$*,* $r := r_1 + r_2$*, and* A *defined in* (1) *with* $\delta > 0$*. Given* $\eta \in (0, 1)$*, if*

$$m \gtrsim (w_\cap^2 + n\tfrac{\delta^2}{\sigma^2})(1 + \log(1 + \tfrac{rm}{\delta n}) + w_\cap^{-2} \log \tfrac{1}{\eta}), \tag{2}$$

then $p_\delta \geq 1 - \eta$.

Numerically testable but stronger conditions ensuring $p_\delta > 1 - \eta$ in Prob. 2 can be deduced as follows. We first note that if $\boldsymbol{\Phi z} = \boldsymbol{0}_m$ for a given $\boldsymbol{\Phi}$ and any $\boldsymbol{z} \in \mathcal{C}^-$, *i.e.*, $\mathrm{Ker}(\boldsymbol{\Phi}) \cap \mathcal{C}^- \neq \emptyset$, then $p_\delta = 0$ for all $\delta > 0$ since then $\boldsymbol{\Phi x}' + \boldsymbol{\xi} = \boldsymbol{\Phi x}'' + \boldsymbol{\xi}$. Second, since $\mathsf{A}(\boldsymbol{x}') = \mathsf{A}(\boldsymbol{x}'')$ induces $\|\boldsymbol{\Phi z}\|_\infty \leq \delta$ for $\boldsymbol{z} := \boldsymbol{x}' - \boldsymbol{x}'' \in \mathcal{C}^-$, proving $\bar{p}_\delta = \mathbb{P}[\forall \boldsymbol{z} \in \mathcal{C}^-, \|\boldsymbol{\Phi z}\|_\infty > \delta] \geq 1 - \eta$ will solve Prob. 2 since $p_\delta \geq \bar{p}_\delta$.

We define accordingly a *consistency margin* $\tau := \|\boldsymbol{\Phi z}^\star\|_\infty$ with

$$\boldsymbol{z}^\star := \mathrm{argmin}_{\boldsymbol{z} \in \mathcal{K}} \|\boldsymbol{\Phi z}\|_\infty \text{ s.t. } \boldsymbol{z} \in \mathcal{C}^- := \mathcal{C}_1 - \mathcal{C}_2, \tag{3}$$

i.e., as a property of $\boldsymbol{\Phi}$ and $\mathcal{C}^-$ so that, if $\tau > \delta$, we necessarily have $\mathsf{A}(\boldsymbol{x}') \neq \mathsf{A}(\boldsymbol{x}'')$ for $\boldsymbol{x}', \boldsymbol{x}''$ in different classes, *i.e.*, $\bar{p}_\delta := \mathbb{P}[\tau > \delta]$. Intuitively, $\boldsymbol{z}^\star$ is related to the minimal separation σ between $\mathcal{C}_1$ and $\mathcal{C}_2$.

Note that (3) is clearly convex if $\mathcal{K}$ and $\mathcal{C}^-$ are convex. We anticipate that the construction of a *certificate* for this problem will provide a bound on $\tau > \delta$ when $\mathcal{C}^-$ is known, and analyse an exemplary case afterwards.

II. Numerical Test for Two Disjoint ℓ_2-Balls

We consider the simple, yet broadly applicable convex case of two balls $\mathcal{C}_1 := r_1\mathbb{B}^n_{\ell_2} + \boldsymbol{c}_1$ and $\mathcal{C}_2 := r_2\mathbb{B}^n_{\ell_2} + \boldsymbol{c}_2$. Then, $\mathcal{C}^- = r\mathbb{B}^n_{\ell_2} + \boldsymbol{c}$ with $\boldsymbol{c} := \boldsymbol{c}_1 - \boldsymbol{c}_2$ and $r := r_1 + r_2$ (see Fig. 1). In this context $\|\boldsymbol{c}\| = \sigma + r$ and $\frac{w_\cap}{\sqrt{n}} \lesssim \frac{r}{\|\boldsymbol{c}\|} \leq \frac{r}{\sigma}$ in Prop. 1 [4].

By solving random instances of this problem[2] w.r.t. $\boldsymbol{\Phi} \sim \mathcal{N}^{m \times n}(0, 1)$, with $n = 2^8$ and $m \in [2^1, 2^8]$, we are able to compute the consistency margin for each $\boldsymbol{\Phi}$ on $\mathcal{C}^-$, which is varied by fixing $r = 2$ and taking $\sigma = \|\boldsymbol{c}\| - r \in [2^0, 2^9]$. Then, we collect $\tau_{\min}$, *i.e.*, the smallest τ resulting from 2^7 trials for each configuration (Fig. 2a), and also estimate on the same trials the probability $\bar{p}_\delta = \mathbb{P}[\tau > \delta := 1]$ in Fig. 2b.

Fig. 2a reports several level curves of $\tau_{\min}$. For each curve, the event $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$ holds if $\delta := \tau_{\min}$. While this condition is necessary but not sufficient, these level curves are compatible with the points $(\frac{m}{n}, \sigma)$ that satisfy the rule $m \approx \frac{\delta^2}{\sigma^2}n$ (up to log factors) induced by (2) in Prop. 1. Fig. 2b displays a sharp *phase transition* in the contours of $\bar{p}_\delta$. Despite the fact that $p_\delta \geq \bar{p}_\delta$, the contours are also approximately aligned with the iso-probability curves that can be deduced from (2), *i.e.*, $m - c\frac{r^2 + \delta^2}{\sigma^2}n \approx C\log(\frac{1}{\eta})$, with $\bar{p}_\delta \approx 1 - \eta$ for some $C, c > 1$.

III. Conclusion and Open Questions

The fundamental limits of learning tasks with embeddings are being tackled in several studies [5]–[8]. Our contribution expands the requirements for exact classification from the signatures produced by two closed convex sets after quantised random embeddings. We shall also specify this analysis to low-complexity structured sets $\mathcal{K}$ (*e.g.*, selecting disjoint "clusters" of sparse signals).

[2] By uniformity of $\mathrm{Ker}(\boldsymbol{\Phi})$, $\boldsymbol{\Phi} \sim \mathcal{N}^{m \times n}(0, 1)$ over the Grassmannian at the origin, it is legitimate to fix a randomly drawn direction $\boldsymbol{c}/\|\boldsymbol{c}\|$ for the simulations.

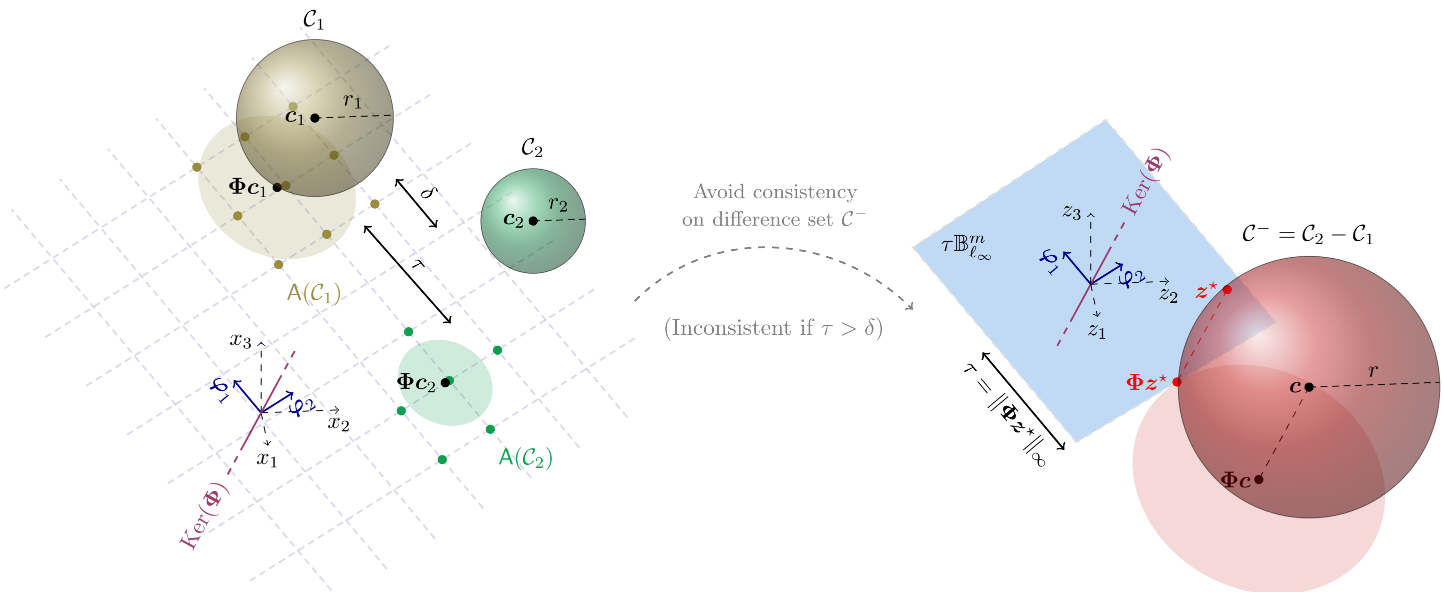

Figure 1. Geometrical intuition on the quantised eclipse problem for two disjoint ℓ_2-balls and $n = 3$, $m = 2$: (left) $\mathcal{C}_1$ and $\mathcal{C}_2$ are projected on $\boldsymbol{\Phi}$, identified by the unit vectors $\boldsymbol{\varphi}_1, \boldsymbol{\varphi}_2$; on these directions, we construct the lattice $\delta\mathbb{Z}^m$, with a shift $\boldsymbol{\xi}$ of the origin due to dithering; the finite sets $\mathsf{A}(\mathcal{C}_1)$, $\mathsf{A}(\mathcal{C}_2)$ are also reported, along with the consistency margin τ; (right) ensuring that $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$ requires that any $\boldsymbol{z} \in \mathcal{C}^-$ is so that its image under $\boldsymbol{\Phi}$ has $\|\boldsymbol{\Phi z}\|_\infty > \tau$; taking the smallest of such values on the difference set yields the consistency margin, which is $\tau = 0$ when $\mathrm{Ker}(\boldsymbol{\Phi}) \cap \mathcal{C}^- \neq \emptyset$.

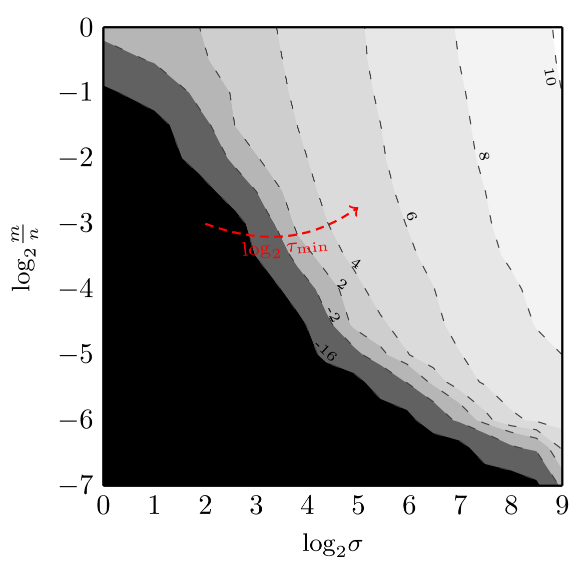

(a) Minimum consistency margin $\tau_{\min}$

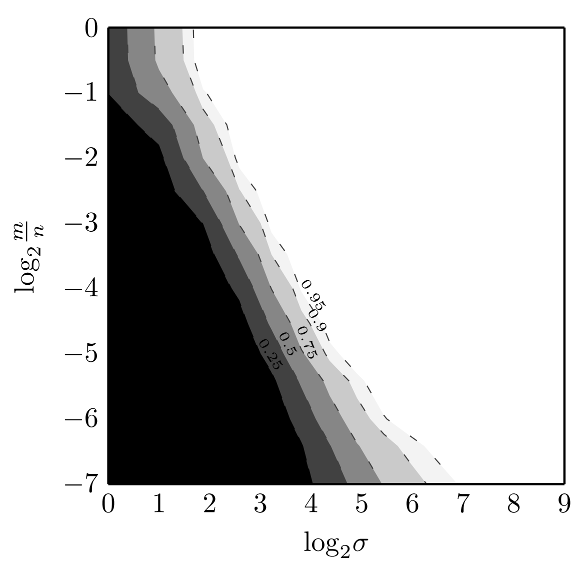

(b) Probability that $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$ at resolution $\delta := 1$

Figure 2. Empirical phase transitions of the quantised eclipse problem for the case of two disjoint ℓ_2-balls; for several random instances of $\boldsymbol{\Phi}$ and as a function of σ and the dimensionality reduction rate $\frac{m}{n}$, we report (a) the contours of $\log_2 \tau_{\min}$; (b) the contours of $\bar{p}_\delta = \mathbb{P}[\tau > \delta] \approx 1 - \eta$ for $\delta := 1$. In (a), the level curves of $\tau_{\min}$ are compatible, up to log factors, with the points $\{(\frac{m}{n}, \sigma) : m \approx \delta^2 n / \sigma^2\}$ deduced from (2) in Prop. 1. In (b), the level curves of $\bar{p}_\delta$ are also approximately aligned with the iso-probability curves $m - c\frac{r^2+\delta^2}{\sigma^2} n \approx C \log(\frac{1}{\eta})$, also deduced from (2), once we set $\bar{p}_\delta \approx 1 - \eta \in \{0.25, 0.5, 0.75, 0.9, 0.95\}$ for some $C, c > 1$.

References

[1] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, 1998.

[2] L. Jacques and V. Cambareri, "Time for dithering: fast and quantized random embeddings via the restricted isometry property," *Information and Inference: A Journal of the IMA*, Apr. 2017, accepted and in press.

[3] A. Moshtaghpour, L. Jacques, V. Cambareri, K. Degraux, and C. De Vleeschouwer, "Consistent Basis Pursuit for Signal and Matrix Estimates in Quantized Compressed Sensing," *IEEE Trans. Signal Process.*, vol. 23, no. 1, pp. 25–29, 2016.

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

[5] R. Giryes, G. Sapiro, and A. M. Bronstein, "Deep Neural Networks with Random Gaussian Weights: A Universal Classification Strategy?" *IEEE Trans. Signal Process.*, vol. 64, no. 13, pp. 3444–3457, 2015.

[6] H. Reboredo, F. Renna, R. Calderbank, and M. R. Rodrigues, "Bounds on the Number of Measurements for Reliable Compressive Classification," *IEEE Trans. Signal Process.*, vol. 64, no. 22, pp. 5778–5793, 2016.

[7] N. Keriven, A. Bourrier, R. Gribonval, and P. Pérez, "Sketching for large-scale learning of mixture models," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6190–6194.

[8] N. Keriven, N. Tremblay, Y. Traonmilin, and R. Gribonval, "Compressive K-means," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.