

# When compressive learning fails: blame the decoder or the sketch?

Vincent Schellekens<sup>1</sup> and Laurent Jacques<sup>1</sup>.

<sup>1</sup>ICTEAM/ELEN, UCLouvain.

**Abstract**— In compressive learning, a mixture model (a set of centroids or a Gaussian mixture) is learned from a sketch vector, that serves as a highly compressed representation of the dataset. This requires solving a non-convex optimization problem, hence in practice approximate heuristics (such as CLOMPR) are used. In this work we explore, by numerical simulations, properties of this non-convex optimization landscape and those heuristics.

## 1 Introduction

The general goal of machine learning is to infer a model, usually represented by parameters  $\theta$ , from a dataset  $X = \{\mathbf{x}_i\}_{i=1}^n$  of  $n$  training examples  $\mathbf{x}_i \in \mathbb{R}^d$ . When  $n$  is very large (as in many modern datasets), conventional learning algorithms—that operate in several passes over  $X$ —require enormous computational resources (both in memory and time). Compressive learning (CL) tackles this critical challenge by first (in a single parallelizable pass) compressing  $X$  into a lightweight vector  $\mathbf{z} \in \mathbb{C}^m$  called the *sketch* [1]. The parameters  $\hat{\theta} = \Delta[\mathbf{z}]$  are then learned by a *decoder* algorithm  $\Delta$  from only this compressed proxy for  $X$ , using potentially much less resources than operating on the raw dataset  $X$ .

In general, the sketch of a distribution  $\mathcal{P}$  is a linear operator, constructed as a set of generalized moments  $\mathcal{A}(\mathcal{P}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} \Phi(\mathbf{x})$  where the feature map  $\Phi$  is typically randomly generated; the sketch of a dataset is then the *empirical* average those features. In this work we focus on the random Fourier features sketch where  $\Phi(\mathbf{x}) := e^{i\Omega^T \mathbf{x}}$ , *i.e.*, the sketch of  $X$  is

$$\mathbf{z} := \mathcal{A}\left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}\right) = \frac{1}{n} \sum_{i=1}^n e^{i\Omega^T \mathbf{x}_i} \in \mathbb{C}^m, \quad (1)$$

where  $\delta_{\mathbf{x}}$  is the Dirac measure at  $\mathbf{x}$  and  $\Omega = (\omega_j \in \mathbb{R}^d)_{j=1}^m$  is a set of  $m$  “frequencies” generated i.i.d. according to some law. To avoid over-sampling large frequencies (due to the curse of dimensionality), a good heuristic [2] is to set  $\omega = R\varphi$ , where  $\varphi \sim \mathcal{U}(S^{d-1})$  is a normed random direction and  $R \sim p_R(R; \sigma)$  is the norm of  $\omega$ . In [2], this latter distribution is either (FG) a folded Gaussian  $p_R \propto e^{-(\sigma R)^2}$ , or (AR) the adapted radius distribution defined as  $p_R \propto \left((\sigma R)^2 + \frac{(\sigma R)^4}{4}\right)^{\frac{1}{2}} e^{-(\sigma R)^2}$ ; in both cases  $\sigma$  is a scale parameter, which should be adjusted to the current dataset either by prior knowledge or from a fast heuristic (see Sec. 2.2).

Learning from the sketch is formulated as an inverse problem, where the “signal” to recover is a density  $\mathcal{P}_\theta$ ; for example, in Gaussian mixture modeling  $\mathcal{P}_\theta = \sum_{k=1}^K \alpha_k \mathcal{N}(\boldsymbol{\mu}_k, \Sigma_k)$  with  $\theta = (\alpha_k, \boldsymbol{\mu}_k, \Sigma_k)_k$  the GMM parameters [2], and in k-means clustering  $\mathcal{P}_\theta = \sum_{k=1}^K \alpha_k \delta_{\mathbf{c}_k}$  with  $\theta = (\alpha_k, \mathbf{c}_k)_k$  the set of (weighted) centroids  $\mathbf{c}_k$  [3]. The model  $\hat{\theta}$  estimated from  $\mathbf{z}$  is then found using the “sketch matching” cost function  $\mathcal{L}$ ,

$$\hat{\theta} \in \arg \min_{\theta} \mathcal{L}(\theta; \mathbf{z}) \text{ where } \mathcal{L}(\theta; \mathbf{z}) := \|\mathbf{z} - \mathcal{A}(\mathcal{P}_\theta)\|_2^2. \quad (2)$$

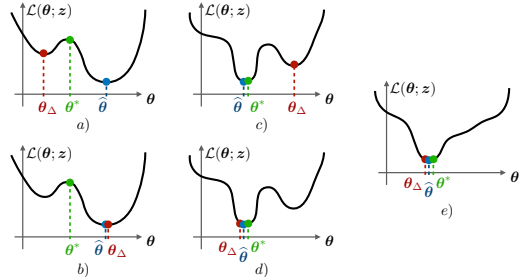


Figure 1: Five scenarii in compressive learning: the cost function  $\mathcal{L}$  (determined by the sketch  $\mathbf{z}$ ) can be bad (when  $\hat{\theta} \neq \theta^*$ , cases (a) and (b)), acceptable (when  $\hat{\theta} \simeq \theta^*$ , but with a small basin of attraction, cases (c) and (d)), or good ( $\hat{\theta} \simeq \theta^*$  with a large basin of attraction, case (e)). Given this cost function, the decoder  $\Delta$  can either fail to find the global optimum (when  $\theta_\Delta \neq \hat{\theta}$ , cases (a) and (c)) or “succeed” (when  $\theta_\Delta \simeq \hat{\theta}$ , the latter being possibly different from  $\theta^*$ , cases (b), (d) and (e)).

The cost  $\mathcal{L}(\theta; \mathbf{z})$  is non-convex, hence challenging to optimize exactly. In practice, a heuristic algorithm is used to solve it approximately, *i.e.*, a decoder  $\Delta : \mathbf{z} \mapsto \theta_\Delta \simeq \arg \min_{\theta} \mathcal{L}(\theta; \mathbf{z})$ . For k-means and GMM, the standard decoder is CLOMPR (CL by Orthogonal Matching Pursuit with Replacement). It starts from an empty solution and greedily adds new atoms  $\theta'_k$  (*i.e.*, a single Gaussian or centroid) to the current solution  $\theta$ , where new atoms  $\theta'_k$  are found by maximizing the non-convex cost  $\langle \mathcal{A}(\mathcal{P}_{\theta'_k}), \mathbf{z} - \mathcal{A}(\mathcal{P}_\theta) \rangle$  starting from a random initial point. To increase the chances of success,  $T$  trials of CLOMPR can be run independently and the solution with the lowest cost is selected, an approach we call CLOMPR<sub>XT</sub>. This decoder showed convincing empirical success for both GMM [2] and k-means [3]. For k-means specifically, the CLAMP decoder (Compressive Learning with Approximate Message Parsing) was also been proposed [4], which allowed to recover the centroids with lower sketch sizes; in this work we focus mainly on CLOMPR.

Previous works thus observed that compressive learning (*e.g.*, using CLOMPR) performed well under the right circumstances (*e.g.*, when the sketch size  $m$  is large enough, and its frequency sampling pattern  $\Lambda$  is well-chosen). *But is it possible to improve the existing CL schemes further? And if so, how? To perform constructive research on these question, we must first understand when and why existing compressive learning schemes fail.* In this work, we provide some insights on this question, based on observations from numerical simulations.

## 2 Results

**CL failure scenarii** In Fig. 1, we classify the different outcomes (a) – (e) that are possible when one runs a compressive learning decoder on a given sketch, sorted by “what could go wrong”. If the sketch was poorly designed, then  $\mathcal{L}$  does not indicate (through its global minimum  $\hat{\theta}$ ) the desired parameters  $\theta^*$ , as shown in cases (a) – (b). In [1], the authors derive theoretical guarantees on (an upper bound for) the probability




Figure 2: **Blue**: relative SSE (with respect to  $k$ -means) of the compressively learned centroids  $\theta_{\text{CLOMPR}\times 10}$ , as a function of the sketch size  $m$  (log scale, median over 50 draws of  $\Omega$ ). **Red**: number of times a failure of the type  $\mathcal{L}(\theta_{\text{CLOMPR}\times 10}; \mathbf{z}) > \mathcal{L}(\theta^*; \mathbf{z})$  was detected. The data is  $n = 5 \times 10^4$  points drawn from  $K = 10$  Gaussians in  $d = 10$ .

that this failure occurs, as a function of the sketch size  $m$  and the compatibility between the sketching function and the target learning task. However, even if the sketch matching cost function aligns with the ideal parameters, it is possible that the decoder fails to solve the non-convex problem, as shown in case (c) (the existing decoders do not have any guarantee). Whether or not the decoder succeeds depends on the general shape of the cost  $\mathcal{L}$  (besides the position of its global optimum). In CLOMPR for example, convergence to  $\hat{\theta}$  is ensured if the random initialization point falls into its *basin of attraction*. It was shown (for the centroid selection in  $k$ -means) that the size of this basin of attraction increases with the sketch size  $m$  [5]. To understand how to improve current CL schemes, we would like to classify practical failures of decoders into either scenarios (a) – (b) (where we know that performance improvement is possible—and, for (b), necessary—by changing the sketch function) or scenario (c) (where we know that improvement is possible by changing the decoder). In the follow, we analyze how this classification is modified when two parameters of the sketch are affected, namely the sketch size  $m$  (Sec. 2.1) and scale parameter  $\sigma$  in the frequency sampling (Sec. 2.2).

## 2.1 Influence of the sketch size

In this section we focus on the  $k$ -means problem, where the goal is to position  $K$  centroids  $c_k$  such that the Sum of Squared Errors,  $\text{SSE}(\{c_k\}_{k=1}^K) := \sum_i \min_k \|x_i - c_k\|_2^2$ , is minimized. In Fig. 2 we plot the quality of compressively learned centroids (by CLOMPR with 10 trials) as a function of  $m$ ; as expected from previous works, they are always perfectly recovered when  $m \geq 10Kd$ , which corresponds to case (e). We also compared the cost of those centroids  $\mathcal{L}(\theta_{\text{CLOMPR}\times 10})$  with the cost of the ground-truth centroids  $\mathcal{L}(\theta^*)$ , and reported the average number of times that  $\mathcal{L}(\theta_{\text{CLOMPR}\times 10}) > \mathcal{L}(\theta^*)$ . This allows us to identify an additional regime: when  $m \leq Kd$ , we know that we are in case (a – b), i.e., the cost is ill-defined (note that here the condition  $m \leq Kd$  coincides with the over-parameterized regime where there are less measurements than degrees of freedom). In between ( $1 \leq m/Kd \leq 10$ ), the situation is less clear: as the number of measurements increase, we transition from (a – b) to (c) with gradually more (d), then finally to (e), but it’s not clear at which point the (a) – (b) switches to (c).

To investigate this transition further, we perform another experiment, where this time we keep  $\Omega$  fixed to analyze only the performances of the decoder for a fixed cost: this is shown Fig. 3. Moreover, we would like to approach the global minimizer  $\hat{\theta}$  as much as possible. For this purpose, we propose a new genetic-algorithm-based decoder, that we name GenetiCL. Its principle is to maintain a population of “chromosome” candidate solutions  $\theta$  that explore the optimization




Figure 3: **Left**: cost function  $\mathcal{L}(\theta; \mathbf{z})$  as a function of an increasing sketch size  $m$  (for a fixed draw of  $\Omega$ ), evaluated on the centroids  $\theta_\Delta$  obtained by different decoders ( $\Delta \in \{\text{CLOMPR}, \text{CLOMPR}\times 10, \text{GenetiCL}\}$ ) as well as the on “ground-truth” centroids  $\theta^*$ . **Right**: relative SSE corresponding to those centroids  $\theta_\Delta$ . Data are  $n = 10^5$  samples from  $K = 10$  Gaussians in  $d = 5$ .




Figure 4: Empirical success rate of CLOMPR $\times 3$  when solving  $k$ -means (blue, defining success as  $\text{RSSE} \leq 1.3$ ) and GMM fitting (orange, defining success as  $\mathcal{P}_\theta(X) \geq (1.3)^{-1} \mathcal{P}_{\theta^*}(X)$  with  $\mathcal{P}(X)$  the likelihood of  $X$  given  $\mathcal{P}$ ) from a sketch with  $m = 20Kd$  frequencies drawn from the FG (plain) and AR (dashed) distributions with varying scale  $\sigma$ . The heuristic proposed in [2] for GMM (resp. in [4] for  $k$ -means) yields  $\sigma_1$  (resp.  $\sigma_2$ ), indicated by black triangles. Data are  $n = 10^5$  samples from  $K = 6$  Gaussians in  $d = 5$ .

landscape by random “mutations” (noise) and combinations by “crossover” (swapping centroids) based on a fitness score that we set to be  $\|z - \mathcal{A}(\mathcal{P}_\theta)\|_2^{-\gamma}$  for  $\gamma > 0$ . While orders of magnitude slower than CLOMPR (we discourage using it for anything other than research purposes), it finds more promising (local) minimizers to  $\mathcal{L}$ , see Fig. 3. For  $m \geq 2Kd$ , GenetiCL performs better, and we are for sure in case (c) for CLOMPR. However, uncertainty remains in the zone  $Kd \leq m \leq 2Kd$ , as the ground truth cost is lower than what the decoders can find, hence it’s not certain whether we are in case (c) or (a).

## 2.2 Influence of the scale and task

We now analyse how the cost function  $\mathcal{L}$  changes with the sketch scale  $\sigma$ , as well as with the considered task (solving  $k$ -means or GMM from a same sketch), shown Fig. 4. First, notice that there exist scales ( $\sigma \in [1, 2]$  in our case) where  $k$ -means succeeds but GMM fails. Intuitively, the culprit is the sketch (case (b)), which captures a very coarse view of the data distribution; it cannot estimate the clusters shape precisely, but can still locate them. Conversely, there are also scales ( $\sigma \in [-3, -1.5]$  in our case) where GMM succeeds but  $k$ -means fails. This is more surprising, as the successfully extracted GMM contains the centroids (in the means of the Gaussians); we thus expect that the decoder failed (case (c)), which would mean that switching the task from  $k$ -means to GMM leads to a better cost function (going from (c) to (e)). Plotting the detected failures (as in Fig. 2, not shown here for conciseness) supports these explanations: when  $\sigma \in [-3, -1.5]$  we detect failures of the  $k$ -means decoder (meaning that (c) is possible), but when  $\sigma \in [1, 2]$  not such failures are detected from the GMM decoder (meaning we are for sure in (a) – (b)).

## References

- [1] R. Gribonval et al. “Compressive Statistical Learning with Random Feature Moments”. In: *ArXiv e-prints* (June 2017). arXiv: 1706.07180 [stat.ML].
- [2] N. Keriven et al. “Sketching for large-scale learning of mixture models”. In: *Information and Inference: A Journal of the IMA* 7.3 (2018), pp. 447–508.
- [3] N. Keriven et al. “Compressive K-means”. In: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2017, pp. 6369–6373.
- [4] E. Byrne et al. “Sketched clustering via hybrid approximate message passing”. In: *IEEE Transactions on Signal Processing* 67.17 (2019), pp. 4556–4569.
- [5] Y. Traonmilin and J.-F. Aujol. “The basins of attraction of the global minimizers of the non-convex sparse spike estimation problem”. In: *Inverse Problems* (2019).