



HAL
open science

Subspace clustering in high-dimensions: Phase transitions & Statistical-to-Computational gap

Luca Pesce, Bruno Loureiro, Florent Krzakala, Lenka Zdeborová

► To cite this version:

Luca Pesce, Bruno Loureiro, Florent Krzakala, Lenka Zdeborová. Subspace clustering in high-dimensions: Phase transitions & Statistical-to-Computational gap. 2022. hal-04019738

HAL Id: hal-04019738

<https://hal.science/hal-04019738>

Preprint submitted on 8 Mar 2023

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

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

Subspace clustering in high-dimensions: Phase transitions & Statistical-to-Computational gap

Luca Pesce¹, Bruno Loureiro^{1,2}, Florent Krzakala¹, and Lenka Zdeborová³

¹Ecole Polytechnique Fédérale de Lausanne (EPFL). Information, Learning and Physics (IdePHICS) lab.
CH-1015 Lausanne, Switzerland.

²Département d'Informatique, École Normale Supérieure - PSL & CNRS, 45 rue d'Ulm,
F-75230 Paris cedex 05, France.

³Ecole Polytechnique Fédérale de Lausanne (EPFL). Statistical Physics of Computation (SPOC) lab.
CH-1015 Lausanne, Switzerland.

{luca.pesce, bruno.loureiro, florent.krzakala, lenka.zdeborova}@epfl.ch

Abstract

A simple model to study subspace clustering is the high-dimensional k -Gaussian mixture model where the cluster means are sparse vectors. Here we provide an exact asymptotic characterization of the statistically optimal reconstruction error in this model in the high-dimensional regime with extensive sparsity, i.e. when the fraction of non-zero components of the cluster means ρ , as well as the ratio α between the number of samples and the dimension are fixed, while the dimension diverges. We identify the information-theoretic threshold below which obtaining a positive correlation with the true cluster means is statistically impossible. Additionally, we investigate the performance of the approximate message passing (AMP) algorithm analyzed via its state evolution, which is conjectured to be optimal among polynomial algorithm for this task. We identify in particular the existence of a statistical-to-computational gap between the algorithm that require a signal-to-noise ratio $\lambda_{\text{alg}} \geq k/\sqrt{\alpha}$ to perform better than random, and the information theoretic threshold at $\lambda_{\text{it}} \approx \sqrt{-k\rho \log \rho}/\sqrt{\alpha}$. Finally, we discuss the case of sub-extensive sparsity ρ by comparing the performance of the AMP with other sparsity-enhancing algorithms, such as sparse-PCA and diagonal thresholding.

1 Introduction

With the growing size of modern data, clustering techniques play an important role in reducing the dimensionality of the features used in modern Machine Learning pipelines. Indeed, in many tasks of interest ranging from DNA sequence analysis to image classification, the relevant features are known to live in a lower-dimensional space (intrinsic dimension) than their raw acquisition format (extrinsic dimension) [1]. In these cases, identifying these features can help saving computational resources while significantly improving on learning performance. But given a corrupted embedding of low-dimensional features in a high-dimensional space, is it always *statistically possible* to retrieve them? And if yes - how can reconstruction be achieved *efficiently* in practice? In this manuscript we address these two fundamental questions in a simple model for subspace clustering: a k -cluster Gaussian mixture model with sparse centroids. In this model, the low-dimensional hidden features are given by the sparse centroids, which are embedded in a higher dimensional space and corrupted by additive Gaussian noise. We assume that the number of non-zero components of the centroids as well as the number of samples scales linearly with the dimension of the embedding space. Given a finite sample from the mixture, the goal of the statistician is to cluster the data, i.e. estimate the centroids (or features) as well as possible.

2 Model & setting

Let $\mathbf{x}_\nu \in \mathbb{R}^d$, $\nu \in [n] := \{1, \dots, n\}$ denote n i.i.d. data points drawn from an isotropic k -cluster Gaussian mixture:

$$\mathbf{x}_\nu \sim_{i.i.d.} \sum_{c \in \mathcal{C}} p_c \mathcal{N}\left(\sqrt{\lambda/s} \boldsymbol{\mu}_c, \mathbf{I}_d\right), \quad \nu \in [n] \quad (1)$$

where p_c are the class probabilities, \mathcal{C} is the index set ($|\mathcal{C}| = k$), $\boldsymbol{\mu}_c \in \mathbb{R}^d$ are the s -sparse vectors representing the means of the clusters and λ is a measure of the signal-to-noise ratio (SNR) for this model. In the following, we focus in the balanced case for which $p_c = 1/k$. We select the subspace of relevant features thanks to the introduction of the vector

$\mathbf{v}_i \in \mathbb{R}^k$. The projection of all the cluster means, on a given dimension $i \in [d]$, will be completely determined by a linear combination of the components of \mathbf{v}_i . We consider for this purpose a Gauss-Bernoulli distribution:

$$\mathbf{v}_i \sim_{i.i.d.} \rho \mathcal{N}(0, \mathbf{I}_k) + (1 - \rho) \delta_0 \quad (2)$$

where $\rho := s/d$ is the density of non-zero elements. For the convenience of the theoretical analysis that follows, it will be useful to work with a particular encoding of the class labels \mathcal{C} . For a given sample $\nu \in [n]$ belonging to the class $c \in \mathcal{C}$, define the following label indicator vector $\mathbf{u}'_c \in \{-\frac{1}{k}, \frac{k-1}{k}\}^k$:

$$\mathbf{u}'_c = \frac{1}{k} (-1, \dots, -1, \underbrace{k-1}_{c\text{-th element}}, -1, \dots, -1) \quad (3)$$

For a given draw $(\mathbf{x}_\nu)_{\nu \in [n]}$, define the matrices $\mathbf{X} \in \mathbb{R}^{d \times n}$ with columns given by \mathbf{x}_ν and $\mathbf{U} \in \mathbb{R}^{n \times k}$, $\mathbf{V} \in \mathbb{R}^{d \times k}$ with rows given by \mathbf{u}'_c and \mathbf{v}_i .

A crucial point that we shall exploit in this paper is that, with this notation, our model for subspace clustering can be written as a matrix factorization problem, where the data has been generated as:

$$\mathbf{X} = \sqrt{\frac{\lambda}{s}} \mathbf{V} \mathbf{U}^\top + \mathbf{W} \quad (4)$$

with $\mathbf{W} \in \mathbb{R}^{d \times n}$ a Gaussian matrix with elements $W_{i\nu} \sim \mathcal{N}(0, 1)$. This formulation is completely equivalent to the one in eq. (1), by identifying the cluster means as given component-wise by: $\mu_c^{(i)} = \mathbf{v}_i^\top \mathbf{u}_c$, $i \in [d]$. The centroids will have, in expectation, only s non-zero components; as discussed in the introduction, the s -sparse class means $\{\mu_c\}_{c \in \mathcal{C}}$ can be thought of as low-dimensional features embedded in a higher-dimensional space \mathbb{R}^d , that have been corrupted by isotropic additive Gaussian noise with variance $\sim \lambda^{-1}$. Given a finite draw $(\mathbf{x}_\nu)_{\nu \in [n]}$ generated from this model with class means \mathbf{V}_* and labels \mathbf{U}_* , the goal of the statistician is to perform *clustering*, i.e. to estimate \mathbf{U}_* from \mathbf{X} , which is equivalent to retrieving the class label for each sample in the data. However, note that there is a clear class symmetry in this problem: the labels can only be estimated up to a permutation $\pi \in \text{Sym}(\mathcal{C})$ reshuffling the columns of \mathbf{U}_* . Taking this symmetry into account we can assess the performance of an estimator $\hat{\mathbf{U}}$ through the averaged symmetrized mean-squared error:

$$\text{MSE}(\hat{\mathbf{U}}) = \min_{\pi \in \text{Sym}(\mathcal{C})} \frac{1}{n} \mathbb{E} \|\pi(\hat{\mathbf{U}}) - \mathbf{U}_*\|_F \quad (5)$$

where $\|\cdot\|_F$ denotes the matrix Frobenius norm. In particular, we will be interested in characterizing reconstruction in the *proportional high-dimensional limit* where the number of samples n , the ambient dimension d and the sparsity level s go to infinity $n, d, s \rightarrow \infty$ at fixed density $\rho = s/d$, sample complexity $\alpha = n/d$, number of clusters $k \geq 2$ and signal strength $\lambda > 0$.

Note that the clustering problem above is closely related to the problem of estimating the class means / features \mathbf{V}_* . Indeed, written as in eq. (4) the problem of estimating both the labels and centroids $(\mathbf{U}_*, \mathbf{V}_*)$ boils down to a low-rank matrix factorization problem. In this manuscript, we have chosen to frame the discussion in terms of the clustering, but all our results could be presented also in terms of the reconstruction of the class means.

In this manuscript we provide a sharp asymptotic characterization of when reconstruction, as measured by positive correlation with the ground truth, is possible in high-dimensions for this model, both *statistically* and *algorithmically*. In particular, our **main contributions** are:

- We map the subspace clustering problem to an asymmetric matrix factorization problem. We compute the closed-form asymptotic formula for the performance of the Bayesian-optimal estimator in the high-dimensional limit, building on a statistical physics inspired approach [2] that has been rigorously proven in this case [3, 4]. This allows us to provide a sharp threshold below which reconstruction of the features is *statistically* impossible as a function of the parameters of the model.
- To estimate the algorithmic limitations of reconstruction, we analyse a tailored approximate message passing (AMP) algorithm [2, 5, 6] approximating the posterior marginals in this problem, and derive the associated *state evolution equations* characterising its asymptotic performance. Such algorithms are optimal among first order methods [7, 8] and therefore their reconstruction threshold provides a bound on the algorithmic complexity clustering in our model.
- The two results above allow us to paint a full picture of the *statistical-to-algorithmic* trade-offs in high-dimensional subspace clustering for the sparse k -Gaussian mixture model, and in particular to identify an *algorithmically hard* region of the sparsity level $(1 - \rho)$ vs. signal-to-noise ratio (λ) plane for which reconstruction is possible statistically but not algorithmically, see Fig. 1. Further, we provide a detailed analysis in the high sparsity regime $(\rho \rightarrow 0^+)$ of how the algorithmically hard region grows as we increase the number of clusters and the sparsity level. In particular, the *information theoretical* transition arises at $\lambda_{\text{it}} \approx \sqrt{-k\rho \log \rho} / \sqrt{\alpha}$, and the *algorithmic* one at $\lambda_{\text{alg}} \geq k / \sqrt{\alpha}$.

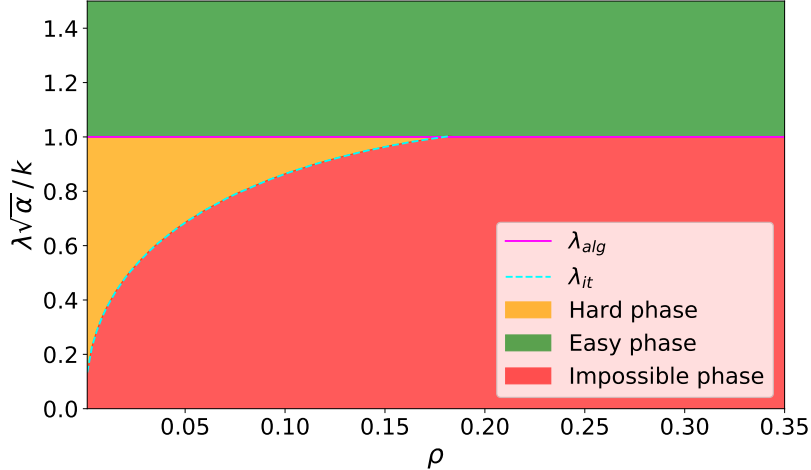


Figure 1: Phase diagram for the subspace clustering of two-clusters GMM at fixed $\alpha = 2$. We plot the SNR λ as a function of ρ and we rescale the y-axis by $\sqrt{\alpha}/k$. We colour different region of the figure according to the associated phase. The algorithmic threshold λ_{alg} is the solid line in magenta while the information-theoretic threshold λ_{it} is the dashed line in cyan. In the *impossible* region, no method can perform better than a random guess. In the *hard* region, a partial reconstruction of the signal is theoretically possible, but we do not know of any polynomial time algorithm that can do it. In the *easy* region, however, AMP, can achieve positive correlation with the ground truth (and actually achieves Bayes MMSE, except very close to the tri-critical point, see the discussion in App. A)

- The analysis for AMP optimality relies on the finite ρ assumption as $n, d \rightarrow \infty$. We thus also investigate the complementary case when the number of non-zero components is of the order $s \lesssim n$ and indeed see that sparse principal component analysis (SPCA) and Diagonal thresholding [9] can perform better than random in this region. We find, however, that this requires $s \leq \sqrt{n}$ (thus $\rho = o(1)$). We rephrase our findings in terms of existing literature on the subspace clustering for two-classes Gaussian mixtures [10, 11].

Related works: Subspace clustering is a well-studied topic in classical statistics, with a wide range of methods used in practice, see [12] for a review. Closer to this work is the theoretical line of work studying the limitations of clustering in high-dimensions. Baik, Ben Arous and P  ch   have shown that PCA for Gaussian mixture clustering fails to correlate with the mixture means below certain threshold known as the *BBP transition* [13]. For dense Gaussian means, the statistical and algorithmic limitations of clustering were analysed in different regimes of interest. Our approach to study subspace clustering relies on a mapping to a low-rank matrix factorization problem. Low-rank matrix factorization has been extensively studied in the literature, and its asymptotic Bayes-optimal performance was characterized in [2–5, 14–16]. The construction of AMP algorithms and the associated state evolutions for matrix factorization was done in [2, 5, 6, 17]. In this work, we leverage these general results on matrix factorization. The closest to our work is perhaps [2, 18], where the authors characterize the asymptotic reconstruction thresholds for the dense case ($\rho = 1$) in the proportional limit where the number of samples and input dimensions diverge at a constant rate. Non-asymptotic results were also discussed in [19] which considered a modification of Lloyd’s algorithm [20] achieving minimax optimal rate and proving a computational lower bound. To the best of our knowledge, the case in which the means are sparse has only been analysed in the regime where the number of non-zero components is sub-extensive with respect to the input dimension. Lower bounds for the statistically optimal recovery threshold in this regime were given in [21, 22], while computational lower bounds are studied in [23, 24]. Clustering of two component Gaussian mixtures has been studied in order to shed light on comparison between statistical and algorithmic tractability in a sparse scenario in [10, 11, 25]. In particular, [10] conjectured and [11] proved algorithmic bounds for this problem. They claim that even below the BBP threshold they can build an algorithm achieving exponentially small misclustering error, given that (up-to log factors) we have $s \lesssim \sqrt{n}$. We relate our findings to their work exploring the extreme sparsity regime in detail.

3 Main theoretical results

In this section, we introduce the two main technical results allowing us to characterize the limitations of clustering reconstruction (both statistically and algorithmically) for the model introduced above in the proportional high-dimensional

limit.

Statistical reconstruction: First, note that up to the permutation symmetry, the estimator minimizing the averaged mean-squared error in eq. (5) admits a closed-form solution given by the marginals of the posterior distribution:

$$\hat{U}_{\text{bo}} = \arg \min_{U \in \mathbb{R}^{n \times k}} \text{MSE}(U) = \mathbb{E}[U|X] \quad (6)$$

where the posterior distribution for the model defined in eq. (4) explicitly reads:

$$P(U|X) = \frac{1}{Z(X)} \prod_{\nu=1}^n P_u(\mathbf{u}_\nu) \int_{\mathbb{R}^k} \prod_{i=1}^d (d\mathbf{v}_i P_v(\mathbf{v}_i)) \prod_{\nu=1}^n \prod_{i=1}^d e^{-\frac{1}{2} (X_{\nu i} - \sqrt{\frac{\lambda}{s}} \mathbf{u}_\nu^\top \mathbf{v}_i)^2} \quad (7)$$

and for convenience we defined the vectors $\mathbf{u}_\nu, \mathbf{v}_i \in \mathbb{R}^k$ which are the rows of U, V , and with the prior distribution P_u being the uniform distribution over the indicator vectors defined in eq. (3) and P_v given by the Gauss-Bernoulli distribution defined in eq. (2).

Although in principle it would be possible to compute the minimum mean-squared error (MMSE) given by the Bayes-optimal estimator in eq. (6) by sampling from the posterior from eq. (7), this is impractical when d, n are large. For instance, simply computing the different integrals involved in the evidence Z scale exponentially with the dimensions. The first result consists precisely in a closed-form solution for the asymptotic performance of the Bayes-optimal estimator:

Main theoretical result 1. *In the proportional high-dimensional limit where $n, d, s \rightarrow \infty$ with fixed ratios $\rho = s/d$, $\alpha = n/d$ and fixed λ, k , the minimum mean-squared error for the reconstruction of $U \in \mathbb{R}^{n \times k}$ is given by:*

$$\lim_{n \rightarrow \infty} \text{MMSE} = \frac{k-1}{k} - \text{Tr} M_u^* \quad (8)$$

where $M_u^* \in \mathbb{R}^{k \times k}$ is the solution of the following minimization problem:

$$M_u^* = \arg \min_{M_u \in \mathbb{R}^{k \times k}} \left\{ \max_{M_v \in \mathbb{R}^{k \times k}} \Phi_{rs}(M_u, M_v) \right\}. \quad (9)$$

with:

$$\begin{aligned} \Phi_{rs}(M_u, M_v) &= \frac{\alpha \lambda \text{Tr} M_u M_v}{2\rho} - \mathbb{E}_{\mathbf{v}_*, \mathbf{w}} \left[\log Z_v \left(\frac{\alpha \lambda M_u}{\rho}, \frac{\alpha \lambda M_u \mathbf{u}_*}{\rho} + \sqrt{\frac{\alpha \lambda M_u}{\rho}} \mathbf{w} \right) \right] \\ &\quad - \alpha \mathbb{E}_{\mathbf{u}_*, \mathbf{w}} \left[\log Z_u \left(\frac{\lambda M_v}{\rho}, \frac{\lambda M_v \mathbf{v}_*}{\rho} + \sqrt{\frac{\lambda M_v}{\rho}} \mathbf{w} \right) \right] \end{aligned} \quad (10)$$

where we introduced $\mathbf{u}_* \sim P_u, \mathbf{v}_* \sim P_v, \mathbf{w} \sim \mathcal{N}(0, I_k)$, and we defined $Z_{u/v}$ as:

$$Z_u(A, \mathbf{b}) = \frac{1}{k} \sum_{c \in \mathcal{C}} \exp \left(\mathbf{b}^\top \mathbf{u}_c - \frac{1}{2} \mathbf{u}_c^\top A \mathbf{u}_c \right) \quad (11)$$

$$Z_v(A, \mathbf{b}) = 1 - \rho + \rho \exp \left(\frac{\mathbf{b}^\top (I_k + A)^{-1} \mathbf{b}}{2} \right) \sqrt{\det (I_k + A)^{-1}} \quad (12)$$

Result 1 follows from our mapping of the subspace clustering problem introduced in Sec. 2 to a low-rank matrix factorization form eq. (4). Indeed, this mapping allows us to leverage a closed-form formula characterizing the asymptotic MMSE for low-rank matrix estimation with generic priors that was derived heuristically [2, 18] using the replica method from Statistical Physics and was rigorously proven in a series of works [3–5, 14] in the context of subspace clustering. Our contribution resides in making this connection and drawing the consequences for the subspace clustering problem, a non-trivial endeavour given the complexity of the resulting formulas. Note that the minimization problem eq. (9) is fundamentally different from the one in eq. (6). Indeed, the first involves only low-dimensional variables, and can be easily solved in a computer, while the latter is a high-dimensional problem which is computationally intractable for large d, n . The other parameter $M_v^* \in \mathbb{R}^{k \times k}$ solving eq. (9) can be used to characterize the MMSE reconstruction error on V .

Algorithmic reconstruction: While result 1 allow us to sharply characterize when clustering is statistically possible, it does not provide us a practical way to estimate the true class labels $U_* \in \mathbb{R}^{n \times k}$ from the data $X \in \mathbb{R}^{d \times n}$. In order to provide a bound in the algorithmic complexity clustering, we consider an *approximate message passing* (AMP) algorithm for our problem. Message passing algorithms are a class of first order algorithms (scaling as $O(nd)$, the dimensions of the input matrix X) designed to approximate the marginals of a target posterior distribution, and which have two very important features. First, for a large class of random problems (such as the clustering problem studied here) AMP provides the best known first order method in terms of estimation performance [26–30], and has been rigorously proven to be the optimal for certain problems [7, 8]. Second, the asymptotic performance of AMP can be tracked by a set of low-dimensional *state evolution equations* [31], meaning that its reconstruction performance can be sharply computed without having to run a high-dimensional instance of the problem. For low-rank matrix factorization problems an associated AMP algorithm can be derived [5, 6, 17, 32, 33]. Therefore, yet again we leverage the mapping of subspace clustering to a low-rank matrix factorization problem to derive the associated AMP algorithm 1 with denoising functions η_v, η_u :

$$\eta_u(\mathbf{A}, \mathbf{b}) = \frac{1}{\sum_{c=1}^k \exp\left(\mathbf{b}^\top \mathbf{u}_c - \frac{\mathbf{u}_c^\top \mathbf{A} \mathbf{u}_c}{2}\right)} \sum_{c=1}^k \mathbf{u}_c \exp\left(\mathbf{b}^\top \mathbf{u}_c - \frac{\mathbf{u}_c^\top \mathbf{A} \mathbf{u}_c}{2}\right) \quad (13)$$

$$\eta_v(\mathbf{A}, \mathbf{b}) = \frac{(\mathbf{I}_k + \mathbf{A})^{-1} \mathbf{b}}{\rho + (1 - \rho) \sqrt{\det(\mathbf{I}_k + \mathbf{A})} \exp\left(\frac{-\mathbf{b}^\top (\mathbf{I}_k + \mathbf{A})^{-1} \mathbf{b}}{2}\right)} \quad (14)$$

As mentioned above, one of they key features of Algorithm 1 is that its asymptotic performance can be tracked exactly by a set of low-dimensional equations.

Main theoretical result 2. *In the proportional high-dimensional limit where $n, d, s \rightarrow \infty$ with fixed ratios $\rho = s/d$, $\alpha = n/d$ and fixed λ, k , the correlation between the ground truth (U_*, V_*) and the AMP estimators $(\hat{U}_{amp}^t, \hat{V}_{amp}^t)$ at iterate t ,*

$$M_u^t = \frac{1}{n} U_*^\top \hat{U}_{amp}^t, \quad M_v^t = \frac{1}{n} V_*^\top \hat{V}_{amp}^t \quad (15)$$

satisfy the following state evolution equations:

$$M_u^{t+1} = \mathbb{E}_{\mathbf{u}_* \sim P_u, \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}_k, \mathbf{I}_k)} \left[\eta_u \left(\frac{\alpha \lambda M_u}{\rho}, \frac{\alpha \lambda M_u \mathbf{u}_*}{\rho} + \sqrt{\frac{\alpha \lambda M_u}{\rho}} \mathbf{w} \right) \mathbf{u}_*^\top \right] \quad (16)$$

$$M_v^{t+1} = \mathbb{E}_{\mathbf{v}_* \sim P_v, \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}_k, \mathbf{I}_k)} \left[\eta_v \left(\frac{\lambda M_v}{\rho}, \frac{\lambda M_v \mathbf{v}_*}{\rho} + \sqrt{\frac{\lambda M_v}{\rho}} \mathbf{w} \right) \mathbf{v}_*^\top \right] \quad (17)$$

Moreover, the asymptotic performance of is simply given by:

$$\lim_{n \rightarrow \infty} \text{MSE}(\hat{U}_{amp}^t) = \frac{k-1}{k} - \text{Tr} M_u^t \quad (18)$$

Result 2 is a consequence from the general theory connecting AMP algorithms to their state evolution (SE) [31]. In the context of low-rank matrix factorization, a derivation of the state evolution equations above from Algorithm 1 were first provided for the rank-one case in [5, 6] and were extended to general rank and denoising functions in [2, 18]. A crucial observation is that Result 2 is intimately related to Result 1. Indeed, the state evolution equations (17) coincide exactly with running gradient descent on the potential defined in eq. (10). While the performance of the statistically optimal estimator \hat{U}_{bo} is given by the fixed point with the minimal value of the potential Φ_{rs} , the performance of the AMP estimator \hat{U}_{amp} is described by the closest minima to the initialization $M_u^{t=0}$. Therefore, studying both the statistical and algorithmic limitations of clustering in high-dimensions boils down to the study of the minima of the potential Φ_{rs} , see App. A for further details. The identification of a subspace clustering problem with a matrix factorization one allows us to leverage the rich literature from this field. However, analyzing these formulas is far from trivial and constitutes a considerable technical challenge. The major difficulty is to find a suitable parametrization of the overlap matrices that reduces the number of parameters to be tracked, we discuss this in detail in Sec. 5. Moreover, we find the scaling ansatz in the high-sparsity limit that closes the equations on amenable quantities in order to find analytically the statistical-to-computational gap for the large-rank setting, see Sec. 6. However, we stress that we did not take all the necessary precautions to claim full rigor. e.g. prove formally that the minimum is unique (although we checked all these both analytically and numerically).

Algorithm 1 low-rAMP

Input: Data $X \in \mathbb{R}^{d \times n}$ Initialize $\hat{\mathbf{v}}_i^{t=0}, \hat{\mathbf{u}}_\nu^{t=0} \sim \mathcal{N}(\mathbf{0}_k, \epsilon \mathbf{I}_k)$, $\hat{\sigma}_{u,\nu}^{t=0} = \mathbf{0}_{k \times k}$, $\hat{\sigma}_{v,i}^{t=0} = \mathbf{0}_{k \times k}$.**for** $t \leq t_{\max}$ **do**

$$\mathbf{A}_u^t = \frac{\lambda}{s} \left(\hat{\mathbf{U}}^t \right)^\top \hat{\mathbf{U}}, \quad \mathbf{A}_v^t = \frac{\lambda}{s} \left(\hat{\mathbf{V}}^t \right)^\top \hat{\mathbf{V}}$$

$$\mathbf{B}_v^t = \sqrt{\frac{\lambda}{s}} X \hat{\mathbf{U}}^t - \frac{\lambda}{s} \sum_{\nu=1}^n \sigma_{u,\nu}^t \hat{\mathbf{V}}^{t-1}, \quad \mathbf{B}_u^t = \sqrt{\frac{\lambda}{s}} X^\top V - \frac{\lambda}{s} \sum_{i=1}^d \sigma_{v,i}^t \hat{\mathbf{U}}^{t-1}$$

Take $\{\mathbf{b}_{v,i}^t \in \mathbb{R}^k\}_{i=1}^d, \{\mathbf{b}_{u,\nu}^t \in \mathbb{R}^k\}_{\nu=1}^n$ rows of $\mathbf{B}_v^t, \mathbf{B}_u^t$

$$\hat{\mathbf{v}}_i^{t+1} = \eta_v(\mathbf{A}_v^t, \mathbf{b}_{v,i}^t), \quad \hat{\mathbf{u}}_\nu^{t+1} = \eta_u(\mathbf{A}_u^t, \mathbf{b}_{u,\nu}^t)$$

$$\hat{\sigma}_{v,i}^{t+1} = \partial_{\mathbf{b}} \eta_v(\mathbf{A}_v^t, \mathbf{b}_{v,i}^t), \quad \hat{\sigma}_{u,\nu}^{t+1} = \partial_{\mathbf{b}} \eta_u(\mathbf{A}_u^t, \mathbf{b}_{u,\nu}^t)$$

Here $\hat{\mathbf{U}}^t \in \mathbb{R}^{n \times k}$, $\hat{\mathbf{V}}^t \in \mathbb{R}^{d \times k}$, $\mathbf{B}_u^t \in \mathbb{R}^{n \times k}$, $\mathbf{B}_v^t \in \mathbb{R}^{d \times k}$, $\mathbf{A}_u^t \in \mathbb{R}^{k \times k}$, $\mathbf{A}_v^t \in \mathbb{R}^{k \times k}$ **end for****Return:** Estimators $\hat{\mathbf{v}}_{\text{amp},i}, \hat{\mathbf{u}}_{\text{amp},\nu} \in \mathbb{R}^k$, $\hat{\sigma}_{u,\nu}, \hat{\sigma}_{v,i} \in \mathbb{R}^{k \times k}$

4 Reconstruction limits for sparse clustering

These results allow us to paint a full picture for when sparse subspace clustering is possible in the model defined in Sec. 2 as a function of the quantity of data α , the sparsity $1 - \rho$, the number of clusters k and the SNR λ . Figure 1 summarizes the different reconstruction regimes in the (ρ, λ) plane for a two-clusters problem at fixed sample complexity $\alpha = 2$, also known as a *phase diagram*. Moreover the general considerations on the reconstruction limits for sparse clustering, given by analyzing the two-clusters problem, are easily generalizable to the general mixture case and not restricted to that particular model, see Sec. 5. For a fixed sparsity $1 - \rho$, we identify the following regions in Fig. 1:

Impossible phase: There is not enough information in the data matrix X handed to the statistician to assign cluster membership better than chance for *any algorithm*. The Bayes-optimal MMSE is not better than a random guess. Clustering (reconstruction of \mathbf{U}_* better than chance) is impossible.

Hard phase: The MMSE is non-trivial, and clustering is statistically possible to some extent, but the best known polynomial time algorithm, AMP, fails to correlate better than chance with the true cluster assignment \mathbf{U}_* . Any polynomial-time algorithm is conjectured to fail in this region.

Easy phase: In the easy phase, not only clustering is statistically possible, but AMP is able to achieve positive correlation with \mathbf{U}_* . One can also investigate when AMP achieves the Bayes-optimal MMSE (instead of just positive correlation) leading to the same transition (except from a subtle correction very close to the tri-critical point, see the discussion in App. A).

In Fig. 2, we investigate the different phases presented above by varying the sparsity level $1 - \rho$ and SNR λ at fixed α . We compare the performance of AMP with popular spectral algorithms: Principal Component Analysis (PCA) and Sparse Principal Component Analysis (SPCA). We can initialize AMP and the SE equations in two different ways: we call *uninformative initialization* a choice for the first iterates of AMP and SE which assumes no knowledge on the ground truth values; conversely with *informative initialization* we consider that the statistician has some prior knowledge on the the ground truth signal. Note that in a real-life scenario the statistician does not have access to the ground truth. Yet, as a theoretical tool the informative initialization provides important information about the algorithmically hard phase, see App A for a discussion. The initialization strategy we considered in the uninformed case for Algorithm 1 is not the only possible choice, there are smarter ways of initialising which can lead to a considerable improvement without explicitly assuming any information about the signal, e.g. spectral initialization [34]. Looking at Fig. 2 we see that SE with uninformed initialization tracks AMP. Moreover we note that increasing the sparsity level, i.e. decreasing ρ , the problem becomes algorithmically harder. This is reflected in a discontinuous jump in the MSE at λ_{alg} which becomes larger. We observe along the same lines a neat advantage by imposing the sparsity constraint in the spectral algorithm (SPCA) with respect to vanilla one (PCA) as the sparsity grows. We discuss the details on the numerical simulations in App. D and the code is available at <https://github.com/lucpoisson/SubspaceClustering>.

5 Stability analysis and algorithmic threshold

In this section we provide a detailed analytical derivation of the threshold λ_{alg} characterizing the algorithmic reconstruction as a function of the number of clusters and the sparsity. First, note due to the permutation symmetry of the clusters

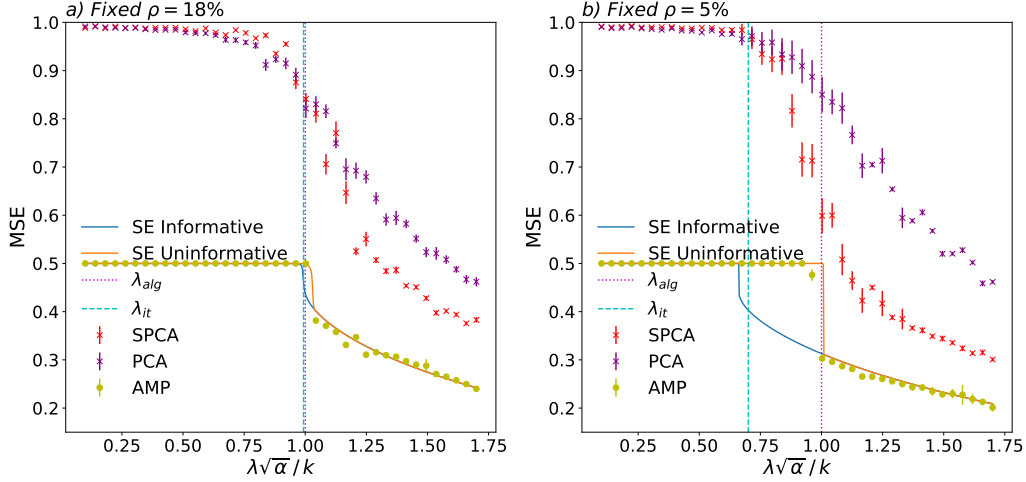


Figure 2: We compare the performance for clustering of two-classes GMM, as measured by the MSE of AMP, SPCA, PCA and SE informed and uninformed. We plot the MSE as a function of the SNR λ and we rescale the x-axis by $\sqrt{\alpha}/k$. For each algorithm considered the error bars are built using the standard deviation over fifty runs with parameters ($n = 8000, d = 4000$), i.e. $\alpha = 2$. We plot in vertical line the theoretical values for the Information-Theoretic threshold λ_{it} (dashed cyan line), and the algorithmic threshold λ_{alg} (dotted line in magenta). The theoretical values coincide with the experimental one. The SE with uninformed initialization follows AMP as expected. *Left*: The sparsity is fixed with parameter $\rho = 18\%$. Both SPCA and PCA have a worse performance with respect to AMP and in this sparsity regime we have only a marginal advantage by using SPCA with respect to PCA. *Right*: The sparsity is fixed with parameter $\rho = 5\%$. Increasing the sparsity level the width of the algorithmically hard phase becomes bigger and SPCA performs clearly better than PCA.

the overlap matrices admit the following parametrization:

$$M_u^t = \frac{m_u^t}{k} I_k - \frac{m_u^t}{k^2} J_k \quad M_v^t = m_v^t I_k - \frac{m_v^t}{k} J_k \quad (19)$$

where J_k is the $k \times k$ -matrix with all elements equal to one. This parametrization is preserved under the SE iterations. Therefore, inserting it in the SE equations yield equations for (m_u^t, m_v^t) :

$$m_u^{t+1} = f_u^{(k)}(\lambda m_v^t / \rho) \quad m_v^{t+1} = f_v^{(k)}(\alpha \lambda m_u^t / \rho) \quad (20)$$

where we introduced the following update functions:

$$f_u^{(k)}(z) := \frac{k}{k-1} \mathbb{E}_{\omega \sim \mathcal{N}(\mathbf{0}_k, I_k)} \left[\frac{e^{z+w_1\sqrt{z}}}{e^{\lambda m_v^t + w_1\sqrt{z}} + \sum_{l=2}^k e^{w_l\sqrt{z}}} \right] - \frac{1}{k-1} \quad (21)$$

$$f_v^{(k)}(z) := \frac{\rho^2 z}{k(k+z)} \int_0^{+\infty} \frac{S_{k-1}}{(2\pi)^{\frac{k}{2}} \rho + (1-\rho) \left(\frac{k+z}{k}\right)^{\frac{k}{2}} e^{-\frac{\xi^2 z}{2k}}} d\xi \quad (22)$$

where S_{k-1} is the surface of the k -dimensional unitary hypersphere. Recall that $(m_u^t, m_v^t) \in [0, 1]^2$ fully characterize the reconstruction performance, with $(m_u^\infty, m_v^\infty) = (0, 0)$ corresponding to the performance of a random guess. Conversely, $(m_u^\infty, m_v^\infty) = (1, \rho)$ corresponds to perfect reconstruction of the cluster membership and the sparse cluster means. One can check that the *trivial fixed point* $(m_u, m_v) = (0, 0)$ is always a fixed point of the SE equations. Moreover, note that its stability is crucially connected to the algorithmic threshold λ_{alg} . Indeed, if the trivial fixed point $(m_u, m_v) = (0, 0)$ is stable, AMP with an uninformed initialization will always converge to this point - and therefore achieve random guessing performance. To study its stability, we expand the update functions around $(m_u, m_v) = (0, 0)$ up to the second order. Expressing everything in terms of m_u , we obtain:

$$m_u^{t+1} = F_{se}(m_u^t) = \frac{\lambda^2 \alpha}{k^2} m_u^t + \left(\frac{(k-4)\alpha^2 \lambda^4}{2k^4} - \frac{\lambda^3 \alpha^2}{k^3} \right) (m_u^t)^2 + o((m_u^t)^2) \quad (23)$$

This immediately tells us that the trivial fixed point becomes unstable at the *algorithmic threshold* $\lambda_{alg} = k/\sqrt{\alpha}$. This transition is a well-known result in random matrix theory and goes under the name of BBP transition [13]. Despite the

fact that we expand around the trivial fixed point, the perturbative method also provides information also about the Bayes-optimal performance, thanks to the general properties that the phase diagrams in Bayes-optimal inference problems must respect [35]. Indeed, a sufficient criterion for the presence of an algorithmically hard phase requires at the algorithmic threshold: $F''_{\text{se}}(0) > 0$. The local study predicts that the phase diagram will present an algorithmically hard region for $k > k_{\text{hard}} = 4 + 2\sqrt{\alpha}$. The criterion nevertheless is not necessary in this setting, as we immediately see from the phase diagram for the two-components GMM in Fig. 1: we clearly observe an algorithmically hard region for high-sparsity although $k < k_{\text{hard}}$. In fact, the *information-theoretic threshold* λ_{it} , defined as the value of the SNR at which the problem becomes statistically possible, cannot be found simply thanks to the expansion around the trivial fixed point. Indeed the exact computation of λ_{it} requires evaluating the potential function Φ in eq. (10): one needs to find the threshold value of the SNR at which the informative fixed point has a lower value of the free energy with respect to the uninformative fixed point. The technical discussion on the calculation of λ_{it} , which we see plotted in Fig. 1 for the two-classes case, is given in App. B, while a general overview on the different thresholds is given in App. A. It is interesting to stress that the linearization of the SE equation around the trivial fixed point in eq. (23) coincide to study the gradient of the potential Φ_{rs} in eq. (10) around that point, and hence the stability of the trivial fixed point is determined by the potential landscape. This connection is at the roots of the conjectured link between the hardness of an inference task and AMP: the statistical-to-computational gap opens up as the fixed point with minimal value of the potential in eq. (10), describing the performance of the Bayes-optimal estimator, is attained far from a locally stable region around the uninformative fixed point, where the AMP algorithm gets stuck.

6 Large sparsity regime

We characterize in this section the scaling of the thresholds in the very sparse (and most interesting) regime when $\rho \rightarrow 0^+$. We highlight the main passages of the computation and focus on the principal results, see Appendix C for a detailed analysis. The starting point is the following change of variables:

$$m_u = \tilde{m}_u \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad m_v = \tilde{m}_v \rho \quad \lambda = C(k)k \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad (24)$$

Inserting these expressions into eq. (20), we obtain simplified SE equations for the rescaled overlaps $(\tilde{m}_u, \tilde{m}_v)$ without any residual dependence on (ρ, α) :

$$\tilde{m}_u = C(k)\tilde{m}_v \quad \tilde{m}_v = T_k(C(k)\tilde{m}_u) \quad (25)$$

where we introduced the auxiliary function $T_k(\cdot)$ defined as:

$$T_k(z) = \int_0^{+\infty} \frac{S_{k-1}}{k(2\pi)^{k/2}} \xi^{k+1} e^{-\xi^2/2} \Theta\left(\frac{z\xi^2}{2} - 1\right) d\xi$$

where $\Theta(\cdot)$ is the Heaviside theta function. By considering the large k expansion of $T_k(\cdot)$, we can derive the scaling of the IT threshold with (k, ρ, α) , see Appendix C for more details. Putting together with our previous result for λ_{alg} from Sec. 5, we obtain the following fundamental result:

$$\lambda_{\text{it}} \approx \sqrt{\frac{-k\rho \log \rho}{\alpha}} \quad \lambda_{\text{alg}} = \frac{k}{\sqrt{\alpha}}. \quad (26)$$

These equations completely characterize the behavior at large rank & small (but finite) sparsity. The statistical-to-computational gap, where AMP is not able to exploit the information on the sparse nature of the cluster means, grows with both the sparsity and the rank. In App. C, we further show that the large rank expression for the IT threshold λ_{it} is accurate already at moderate $k \approx 10$, see Fig. 9. It is interesting at this point to reframe our results on the high sparsity regime in the context of the existing literature. The results for the IT transition for the detection of two-classes sparse GMM were discussed in [22, 23], and we obtain a consistent scaling in the small ρ behaviour. Despite this fact, the algorithmic bounds of different relevant works for the same problem [10, 11, 23] seem, at first sight, to not agree with our findings. In particular [11] proves rigorously the existence of an algorithm that achieves minimax rate under the BBP threshold for clustering of two-classes sparse GMM. This apparent inconsistency is related to different sparsity regimes analyzed. Indeed, here we investigate the extensive sparsity regime, i.e. $\rho = O(1)$, while the guarantees for the efficient algorithms working under the BBP threshold require a very high sparsity level, i.e. $\rho = o(1)$. This difference is crucial. In Fig. 3 we illustrate it by considering the performance of two popular algorithms for this problem: *Diagonal Thresholding* (DT) [9] and SPCA. For extremely large sparsity, i.e. $s = O(1)$, these algorithms indeed provide estimators with positive correlation with the true classes below the BBP threshold! However, as soon as we increase the density of non-zero components the performance strongly deteriorates. In fact, the transition to random chance performance takes place as the number of non-zero components approach $s \lesssim \sqrt{n}$, in agreement with the literature [10, 11].

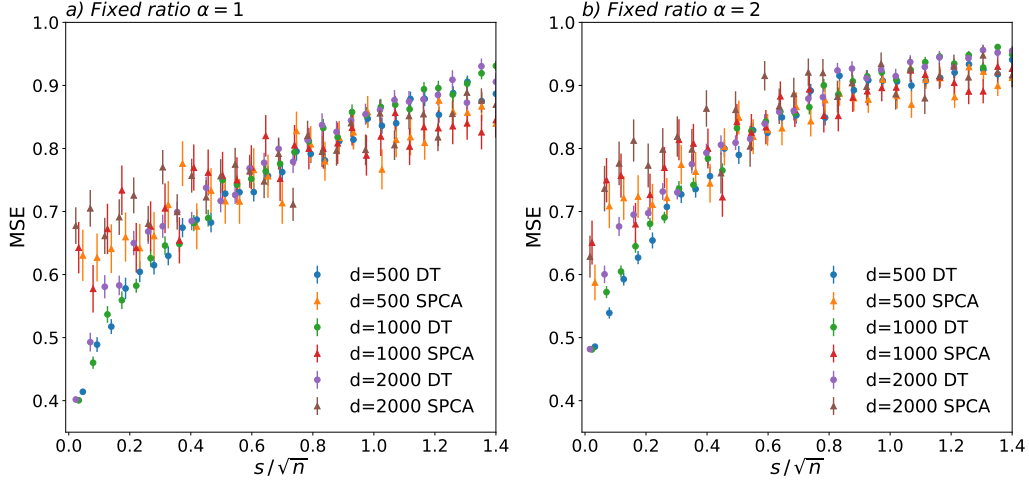


Figure 3: We compare the performance of diagonal thresholding (DT) and Sparse PCA (SPCA), as measured by the MSE, for clustering of two-classes GMM for two different parameter α . We plot the MSE vs the number of non-zero component s and we rescale the x-axis by the square-root of the number of samples. The left plot is done at $\alpha = 1$ while the right one is for $\alpha = 2$. The SNR is tuned such that we are always under the BBP threshold. With this choice we always work in the sub-extensive sparsity regime $\rho = o(1)$ and we can verify numerically what has been claimed in the literature [10, 11]: efficient algorithms in the high dimensional limit need a number of non-zero component (up-to log factors) $s \lesssim \sqrt{n}$ in order to beat random guessing under the BBP threshold.

7 Conclusion

In this work we considered the problem of clustering k homogeneous Gaussian mixtures with sparse means. Mapping this to a low-rank matrix factorization problem, we have provided an exact asymptotic characterization of the MMSE in the high dimensional limit. The Bayes-optimal performance was compared to AMP, the best known polynomial time algorithm for this problem in the studied regime. In the large sparsity regime, we uncovered a large statistical-to-computational gap as the sparsity level grows, and unveiled the existence of a computationally *hard phase*. In particular, we have shown that the SNR threshold below which recovery is statistically impossible is given by $\lambda_{\text{it}} \approx \sqrt{-k\rho \log \rho} / \sqrt{\alpha}$, while the one for which AMP positively correlates with the ground truth classes is given by $\lambda_{\text{alg}} \geq k/\sqrt{\alpha}$. Our result for the existence of an algorithmically hard region was compared with the existing literature for this problem, solving an apparent contradiction due to the scaling assumption of the sparsity level with the dimension of the features. We corroborated our findings with the help of algorithms for subspace clustering such as sparse principal component analysis and diagonal thresholding. The mapping of the subspace clustering problem to a low-rank matrix factorization one is flexible and extending this work to more general scenarios is definitely interesting. The main limitation of the current setting is the Gaussian assumption for the noise. One natural idea to go beyond this limitation is to consider other flavours of message passing algorithms, such as vector AMP (VAMP) [36] to study more general rotationally invariant noise distributions. A further prospective direction is to consider inhomogeneous mixture models, i.e. $p_c \neq 1/k$ in eq. (1), in order to study the effect of unbalancedness on the statistical-to-computational gap.

Acknowledgements

We thank Maria Refinetti for discussions. This work started as a part of the doctoral course *Statistical Physics For Optimization and Learning* taught at EPFL in spring 2021. We acknowledge funding from the ERC under the European Union's Horizon 2020 Research and Innovation Program Grant Agreement 714608-SMiLe, and by the Swiss National Science Foundation grant SNFS OperaGOST, 200021_200390.

References

- [1] Phillip Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning, 2021.
- [2] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(7):073403, Jul 2017.
- [3] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017*, volume 65 of *Proceedings of Machine Learning Research*, pages 1297–1301. PMLR, 2017.
- [4] Léo Miolane. Fundamental limits of low-rank matrix estimation: the non-symmetric case, 2017.
- [5] Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse pca. In *2014 IEEE International Symposium on Information Theory*, pages 2197–2201. IEEE, 2014.
- [6] Alyson K Fletcher and Sundeep Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7(3):531–562, 2018.
- [7] Michael Celentano, Andrea Montanari, and Yuchen Wu. The estimation error of general first order methods, 2020.
- [8] Michael Celentano, Chen Cheng, and Andrea Montanari. The high-dimensional asymptotics of first order methods with random data, 2021.
- [9] Iain M. Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009. PMID: 20617121.
- [10] Jiashun Jin, Zheng Tracy Ke, and Wanjie Wang. Phase transitions for high dimensional clustering and related problems, 2016.
- [11] Matthias Löffler, Alexander S. Wein, and Afonso S. Bandeira. Computationally efficient sparse clustering, 2021.
- [12] Lance Parsons, Ehtesham Haque, and Huan Liu. Subspace clustering for high dimensional data: a review. *SIGKDD Explor.*, 6(1):90–105, 2004.
- [13] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643 – 1697, 2005.
- [14] Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, Lenka Zdeborová, et al. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. *Advances in Neural Information Processing Systems*, 29, 2016.
- [15] Ahmed El Alaoui and Michael I Jordan. Detection limits in the high-dimensional spiked rectangular model. In *Conference On Learning Theory*, pages 410–438. PMLR, 2018.
- [16] Ahmed El Alaoui, Florent Krzakala, and Michael Jordan. Fundamental limits of detection in the spiked wigner model. *The Annals of Statistics*, 48(2):863–885, 2020.
- [17] Ryosuke Matsushita and Toshiyuki Tanaka. Low-rank matrix reconstruction and clustering via approximate message passing. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [18] Thibault Lesieur, Caterina de Bacco, Jess Banks, Florent Krzakala, Cris Moore, and Lenka Zdeborova. Phase transitions and optimal algorithms in high-dimensional gaussian mixture clustering. *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sep 2016.
- [19] Mohamed Ndaoud. Sharp optimal recovery in the two-component gaussian mixture model, 2020.
- [20] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [21] Martin Azizyan, Aarti Singh, and Larry Wasserman. Minimax theory for high-dimensional gaussian mixtures with sparse mean separation, 2013.

- [22] Nicolas Verzelen and Ery Arias-Castro. Detection and feature selection in sparse mixture models, 2016.
- [23] Jianqing Fan, Han Liu, Zhaoran Wang, and Zhuoran Yang. Curse of heterogeneity: Computational barriers in sparse mixture models and phase retrieval, 2018.
- [24] Matthew Brennan and Guy Bresler. Average-case lower bounds for learning sparse mixtures, robust estimation and semirandom adversaries, 2019.
- [25] Martin Azizyan, Aarti Singh, and Larry Wasserman. Efficient sparse clustering of high-dimensional non-spherical gaussian mixtures, 2014.
- [26] David L. Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [27] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Phys. Rev. X*, 2:021005, May 2012.
- [28] Jean Barbier, Florent Krzakala, Nicolas Macris, Leo Miolane, and Lenka Zdeborova. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- [29] Benjamin Aubin, Bruno Loureiro, Antoine Maillard, Florent Krzakala, and Lenka Zdeborová. The spiked matrix model with generative priors. *IEEE Transactions on Information Theory*, 67(2):1156–1181, 2021.
- [30] Benjamin Aubin, Bruno Loureiro, Antoine Baker, Florent Krzakala, and Lenka Zdeborová. Exact asymptotics for phase retrieval and compressed sensing with random generative priors. In Jianfeng Lu and Rachel Ward, editors, *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, volume 107 of *Proceedings of Machine Learning Research*, pages 55–73. PMLR, 20–24 Jul 2020.
- [31] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling, 2012.
- [32] Ryosuke Matsushita and Toshiyuki Tanaka. Low-rank matrix reconstruction and clustering via approximate message passing. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [33] Andrea Montanari and Ramji Venkataramanan. Estimation of low-rank matrices via approximate message passing. *The Annals of Statistics*, 49(1):321 – 345, 2021.
- [34] Marco Mondelli, Christos Thrampoulidis, and Ramji Venkataramanan. Optimal combination of linear and spectral estimators for generalized linear models, 2020.
- [35] Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Typology of phase transitions in bayesian inference problems. *Physical Review E*, 99(4), Apr 2019.
- [36] Sundeep Rangan, Philip Schniter, and Alyson K. Fletcher. Vector approximate message passing, 2016.

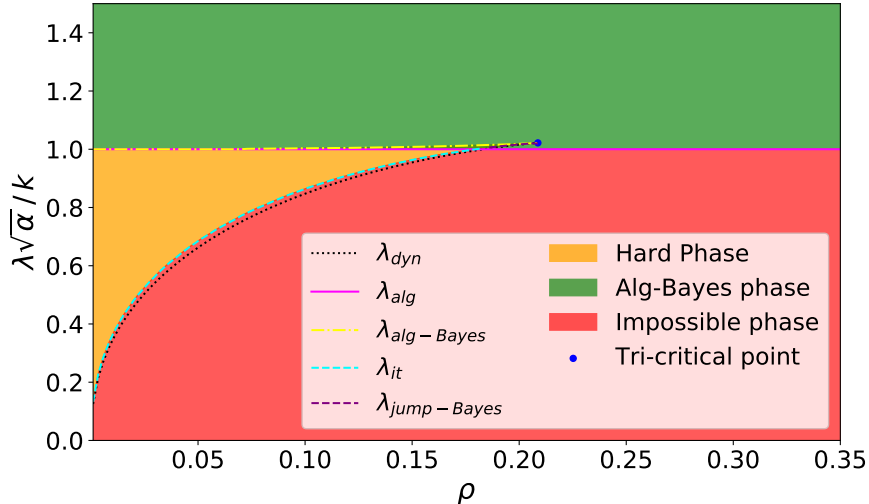


Figure 4: Enriched phase diagram for the subspace clustering of two-clusters GMM at fixed $\alpha = 2$. We plot the SNR λ as a function of ρ and we rescale the y-axis by $\sqrt{\alpha}/k$. We colour different region of the figure according to the associated phase. We introduce in the black dotted line the dynamical spinodal threshold λ_{dyn} , in the yellow dashdotted line the Bayes-algorithmical threshold $\lambda_{alg-Bayes}$ and in purple dashed line the jump-Bayes one $\lambda_{jump-Bayes}$. It is not visible, due to the choice of the axis, the *easy* region, in which AMP performs better than random but not Bayes-optimally. We analyze this in Fig. 6.

Appendix

A Analysis of the thresholds

We identified in Sec. 4 different *reconstruction phases* for the subspace clustering problem, characterizing completely the Bayes-optimal and algorithmical performances. We detail in this section the definition of the reconstruction phases and the consequences for the computational and statistical limits of subspace clustering. We discuss in particular an interesting link between the thresholds separating these phases and the potential function Φ_{rs} in eq. (10). First, we enlarge the picture on the reconstruction phases we offered in Sec. 3. Along with the impossible, hard and easy phases we can define a further region. We call the *Alg-Bayes* phase, the region of parameters in which the performance of AMP is, not only achieving positive correlation with the ground truth, but achieves the Bayes-optimal performance. We summarize now the complete description:

Impossible phase: There is not enough information in the data matrix X handed to the statistician in order to assign cluster membership better than chance for *any algorithm*, and the Bayes-optimal MMSE is not better than a random guess. Clustering (i.e. reconstruction of U_* better than chance) is impossible.

Hard phase: The MMSE is non-trivial, and clustering is statistically possible to some extent, but the best known polynomial time algorithm, AMP, fails to correlate better than chance with the true cluster assignment U_* . Any polynomial-time algorithm is conjectured to fail in this region.

Easy phase: In the easy phase, not only clustering is statistically possible, but AMP is able to achieve positive correlation with U_* .

Alg-Bayes phase: In the alg-Bayes phase AMP is able to achieve Bayes-optimal positive correlation with U_* .

Following the introduction of the alg-Bayes phase, we find an enriched version of the phase diagram for the two-classes subspace clustering at fixed α , see Fig. 4. When we cross from one region to an other we have a *phase transition*. Each phase transition is characterized by different *thresholds*: values of the parameters which signal the onset of a new phase. In Fig. 4 we see different thresholds which were not present in the previous plot in Fig. 1: $\{\lambda_{dyn}, \lambda_{alg-Bayes}, \lambda_{jump-Bayes}\}$. In order to define these quantities, it is useful to highlight the relationship between the thresholds and the minima of the "free energy" Φ_{rs} in eq. (10). Let us fix a sparsity level $1 - \rho$, say $\rho = 0.05$, and move vertically on the y-axis starting from the bottom on the y-axis, i.e. $\lambda\sqrt{\alpha}/k \ll 1$. As we increase the value of the SNR we can identify the following thresholds:

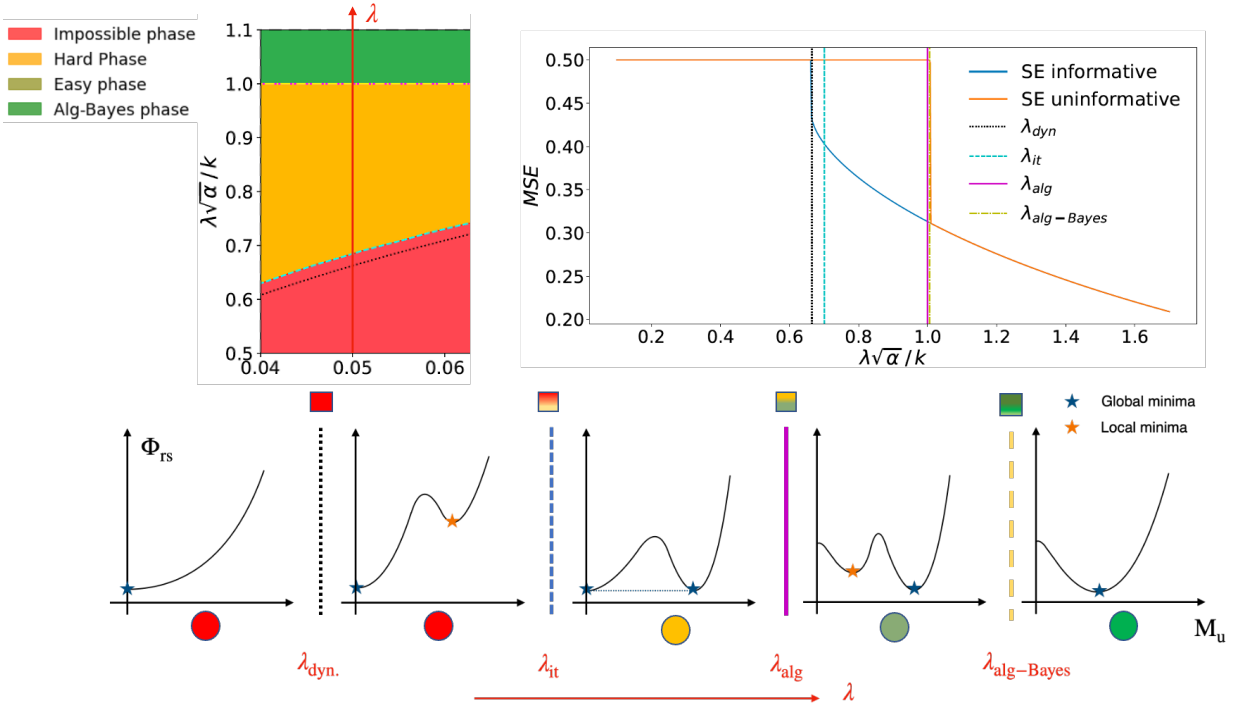


Figure 5: Evolution of the MSE and the potential Φ_{rs} for fixed $\rho \simeq 0.05$ and $\alpha = 2$. Top: We take vertical cross section of the phase diagram in Fig. 4 for $\rho \simeq 0.05$, as explained in the left panel. In the right panel we analyze the MSE via SE both informed and uninformed as a function of the SNR λ and we rescale the x-axis by $\sqrt{\alpha}/k$. In vertical line we plot the different thresholds we encounter as we increase the SNR. Bottom: Cartoon plot of the minima of the potential Φ_{rs} as we follow the cross section of the phase diagram at $\rho \simeq 0.05$. The blue star denotes the global minimum of Φ_{rs} , which corresponds to the correlation of the Bayesian-optimal estimator, while the orange dot denote local minima. We label the phase in which we are at a given stage by colouring the circle below every plot. We plot the thresholds as vertical lines separating the different subplots.

- $\lambda < \lambda_{\text{dyn}}$: The only minimum of eq. (9) is the trivial minimum corresponding to zero correlation $M_u = 0$ with U_\star . Therefore, below this threshold reconstruction is impossible.
- $\lambda \in (\lambda_{\text{dyn}}, \lambda_{\text{it}})$: A second minima with higher Φ_{rs} (i.e. a local minimum) and correlation appears, but the trivial minimum $M_u = 0$ is still the global one. Therefore, AMP with a uninformed initialization $M^{t=0} = 0$ will converge to the trivial minimum, and reconstruction is only possible with a strong informed initialization. We call the threshold value for the emergence of this local minimum the *dynamical spinodal* λ_{dyn} .
- $\lambda \in (\lambda_{\text{it}}, \lambda_{\text{alg}})$: As the SNR is increased, the local minimum goes down in energy Φ_{rs} , and at a certain λ_{it} , it crosses the trivial minimum. Therefore, in this interval the non-trivial minimum is the global one, while the trivial minimum becomes local. Although reconstruction is statistically possible in this region, AMP with a uninformed initialization $M^{t=0} = 0$ is stuck at the trivial minima. Therefore, in this region we enter the hard phase.
- $\lambda \in (\lambda_{\text{alg}}, \lambda_{\text{alg-Bayes}})$ As the SNR is further increased, AMP with a uninformed initialization starts to achieve positive correlation with U_\star , although strictly lower than the Bayes-optimal estimator. In terms of the potential Φ_{rs} , this corresponds to the trivial minimum continuously becoming a local maximum, and another local minimum corresponding to higher correlation continuously appearing. This new local minima coexists with the global one, which corresponds to the Bayes-optimal performance. We enter the easy phase.
- $\lambda > \lambda_{\text{alg-Bayes}}$: Finally, as the SNR is further increased the local minima disappears, and there is only one minimum with high-correlation with the signal left. In this region, AMP with a uninformed initialization achieves the same performance as the Bayes-optimal estimator. We enter the alg-Bayes phase.

The discussion above is summarized in Fig. 7. We plot together with the evolution of the minima of Φ_{rs} , the performance of SE with both informative and uninformed initialization to analyze the behaviour of the MSE.

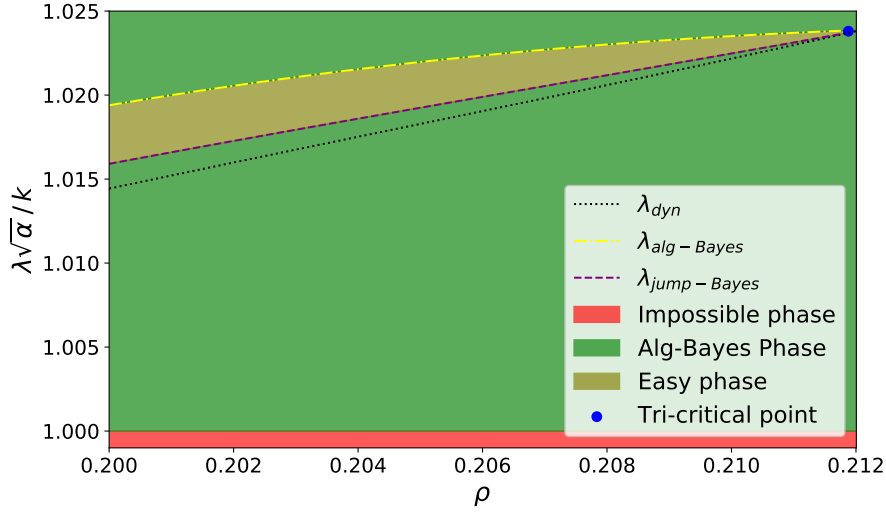


Figure 6: Zoom around the tri-critical point of Fig. 4.

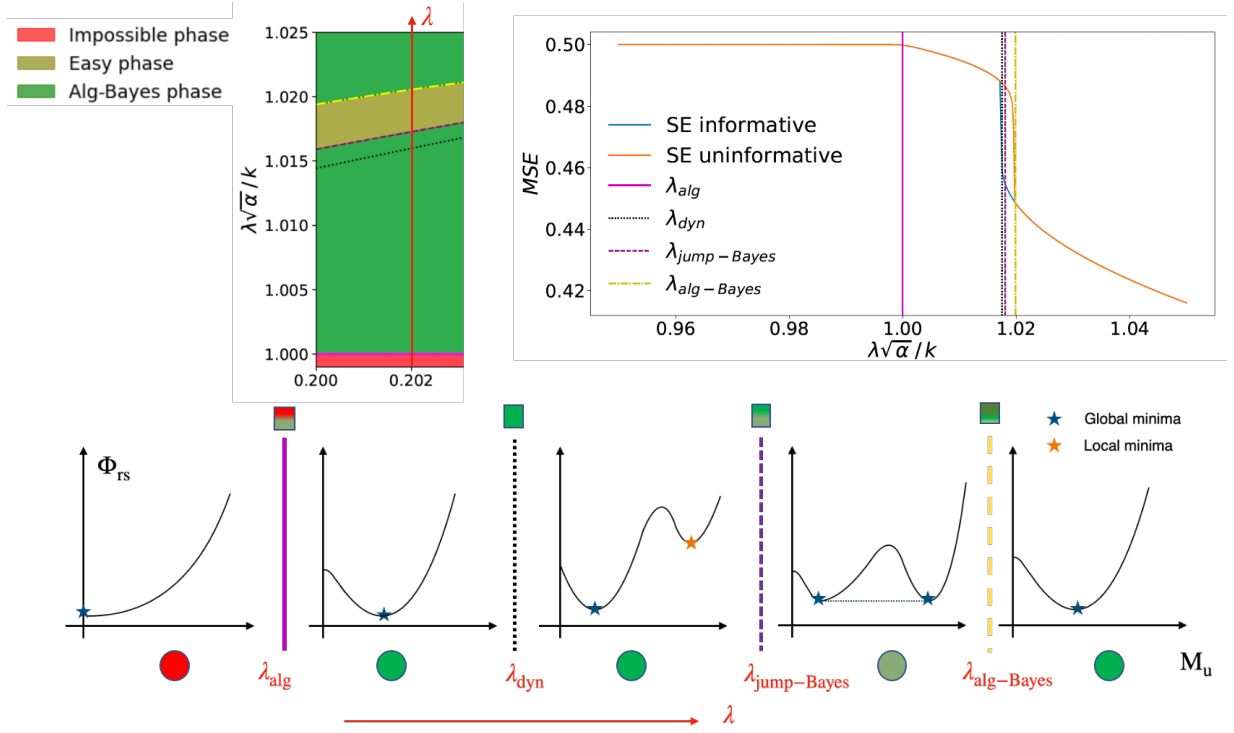


Figure 7: Evolution of the MSE and the potential Φ_{rs} for fixed $\rho \simeq 0.202$ and $\alpha = 2$. Top: We take a vertical cross section of the phase diagram in Fig. 4 for $\rho \simeq 0.202$, as explained in the left panel. In the right panel we analyze the MSE via SE both informed and uninformed as a function of the SNR λ and we rescale the x-axis by $\sqrt{\alpha}/k$. In vertical line we plot the different thresholds we encounter as we increase the SNR. Bottom: Cartoon plot of the minima of the potential Φ_{rs} as we follow the cross section of the phase diagram at $\rho \simeq 0.202$. The blue star denotes the global minimum of Φ_{rs} , which corresponds to the correlation of the Bayesian-optimal estimator, while the orange dot denote local minima. We label the phase in which we are, at a given stage, by colouring the circle below every plot. We plot the thresholds as vertical lines separating the different subplots.

The full characterization of the subspace clustering problem both from an algorithmic and statistical perspective, boils down to the analysis of the evolution of the critical points of Φ_{rs} as we vary the meaningful parameters in the problem. We note from Fig. 4 that the performance of AMP is, when it achieves positive correlation with the ground truth, almost everywhere Bayes-optimal apart from a small region around the *tri-critical* point. This point is defined - at fixed (α, k) - as the tuple of parameters $(\lambda_T(\alpha, k), \rho_T(\alpha, k))$, such that the "spinodal" thresholds meet, i.e. $\lambda_{\text{alg-Bayes}} = \lambda_{\text{dyn}}$. We discuss why these thresholds are called *spinodals*, and how to compute them practically in Sec. B. We can analyze the vicinity of the tri-critical point to analyze the non-trivial interplay between the easy and alg-Bayes phase, where AMP does not achieve always Bayes-optimal performance. Imagine to repeat the same steps as before considering the zoom around the tri-critical point of the phase diagram, see Fig. 6. Fix a sparsity level $1 - \rho$, say $\rho = 0.202$, and move vertically on the y-axis starting from the bottom on the y-axis, i.e. $\lambda\sqrt{\alpha}/k \ll 1$. As we increase the SNR we can repeat the previous analysis, obtaining now:

- $\lambda < \lambda_{\text{alg}}$: The only minimum of eq. (9) is the trivial minimum corresponding to zero correlation $M_u = 0$ with \mathbf{U}_* . Therefore, below this threshold reconstruction is impossible.
- $\lambda \in (\lambda_{\text{alg}}, \lambda_{\text{dyn}})$: The trivial minimum becomes unstable and AMP achieves positive correlation with the ground truth. The minimum is unique and also SE with a positive initialization would end up there. The phase is alg-Bayes.
- $\lambda \in (\lambda_{\text{dyn}}, \lambda_{\text{jump-bayes}})$: As the SNR is increased, a new local minimum appears. The reconstruction phase is still alg-Bayes since the non-trivial minimum has higher free energy than the global one.
- $\lambda \in (\lambda_{\text{jump-Bayes}}, \lambda_{\text{alg-Bayes}})$ As the SNR is further increased, the free energy of the informative minimum goes down and becomes equal to the uninformative one. We enter the easy phase, nevertheless AMP achieves positive correlation with the truth, the Bayes optimal performance is superior. The Bayes-optimal MSE have a first order phase transition at $\lambda_{\text{jump-Bayes}}$, hence the name *jump-Bayes*.
- $\lambda > \lambda_{\text{alg-Bayes}}$: Finally, as the SNR is further increased the local minima disappears, and there is only one minimum with high-correlation with the signal left. In this region, AMP with a uninformed initialization achieves the same performance as the Bayes-optimal estimator. We enter the alg-Bayes phase.

The analysis of the evolution of the minima of Φ_{rs} and the consequences on the MSE is done in Fig. 7.

B Building the phase diagram

We build in this section the phase diagram for the two-classes GMM shown in Fig. 1, explaining the steps which are easily generalizable to the general mixture case. First we note that we can simplify the model for two-mixtures GMM in eq. (4) even further by mapping it to an easier rank $k = 1$ version of the matrix factorization problem. It suffices to replace the matrices (U, V) in eq. (4) by the following quantities:

$$\mathbf{u} \sim \text{Rad}(n) \in \{-1, +1\}^n \quad \mathbf{v} \sim_{i.i.d.} \rho\mathcal{N}(0, \mathbf{I}_d) + (1 - \rho)\delta_0 \in \mathbb{R}^d \quad (27)$$

The two formulations of the problem are indeed formally equivalent up to a rescaling of the parameters such that at fixed α the quantity λ/k is the same in two settings. The mapping simplify significantly the computation. First, in order to compute the Bayes-optimal performance in eq. (8), we must compute the *partition functions* $Z_{u/v}$ in eqs. (11),(12) for the new simplified model. We shall exploit the following general relation, as a function of the prior distribution on (U, V) :

$$Z_{u/v}(\mathbf{A}, \mathbf{b}) = \mathbb{E}_{\mathbf{x} \sim P_{u/v}} \left[\exp \left(-\mathbf{b}^\top \mathbf{x} + \frac{\mathbf{x}^\top \mathbf{A} \mathbf{x}}{2} \right) \right] \quad (28)$$

thus exploiting the explicit expression of the prior in eq. (27) we obtain:

$$Z_u(\mathbf{A}, \mathbf{b}) = e^{-\frac{A}{2}} \cosh b \quad Z_v(\mathbf{A}, \mathbf{b}) = 1 - \rho + \frac{\rho}{\sqrt{1+A}} \exp \left(\frac{b^2}{2(1+A)} \right) \quad (29)$$

We study the computational limits for the subspace clustering problem deriving the associated AMP to this simplified low-rank matrix factorization problem. We have to compute the *denoising functions* for the simplified model, written in eqs. (13),(14) for the general rank case. We shall exploit the general formula relating them with the partition functions computed above:

$$\boldsymbol{\eta}_{u/v}(\mathbf{A}, \mathbf{b}) = \partial_{\mathbf{b}} \log Z_{u/v}(\mathbf{A}, \mathbf{b}) \quad (30)$$

thus using the prior for the simplified $k = 1$ model in eq. (27) we obtain:

$$\eta_u(A, b) = \tanh(b) \quad \eta_v(A, b) = \frac{\rho b}{1 + A} \frac{1}{\rho + (1 - \rho)\sqrt{1 + Ae^{-\frac{b^2}{2(1+A)}}}} \quad (31)$$

where now $(A, b) \in \mathbb{R}^2$. We can write at this point the SE equations for the overlaps (m_u^t, m_v^t) which are now scalar variables. Let us introduce the following notation:

$$m_u^{t+1} = \mathbb{E}_{u_* \sim P_u, \xi \sim \mathcal{N}(0,1)} \left[\eta_u \left(\frac{\lambda m_v^t}{\rho}, \frac{\lambda m_v^t u_*}{\rho} + \sqrt{\frac{\lambda m_v^t}{\rho}} \xi \right) u_* \right] := \mathcal{U}(\lambda m_v^t / \rho) \quad (32)$$

$$m_v^{t+1} = \mathbb{E}_{v_* \sim P_v, \xi \sim \mathcal{N}(0,1)} \left[\eta_v \left(\frac{\alpha \lambda m_u^t}{\rho}, \frac{\alpha \lambda m_u^t v_*}{\rho} + \sqrt{\frac{\alpha \lambda m_u^t}{\rho}} \xi \right) v_* \right] := \mathcal{V}(\alpha \lambda m_u^t / \rho) \quad (33)$$

The perturbative method we presented in Sec 5 would pinpoint easily the expected algorithmical threshold $\lambda_{\text{alg}} = 1/\sqrt{\alpha}$. Despite this, as we discussed in Sec. 5, it would not guarantee the presence of an algorithmically hard region since the number of cluster $k < k_{\text{hard}}$.

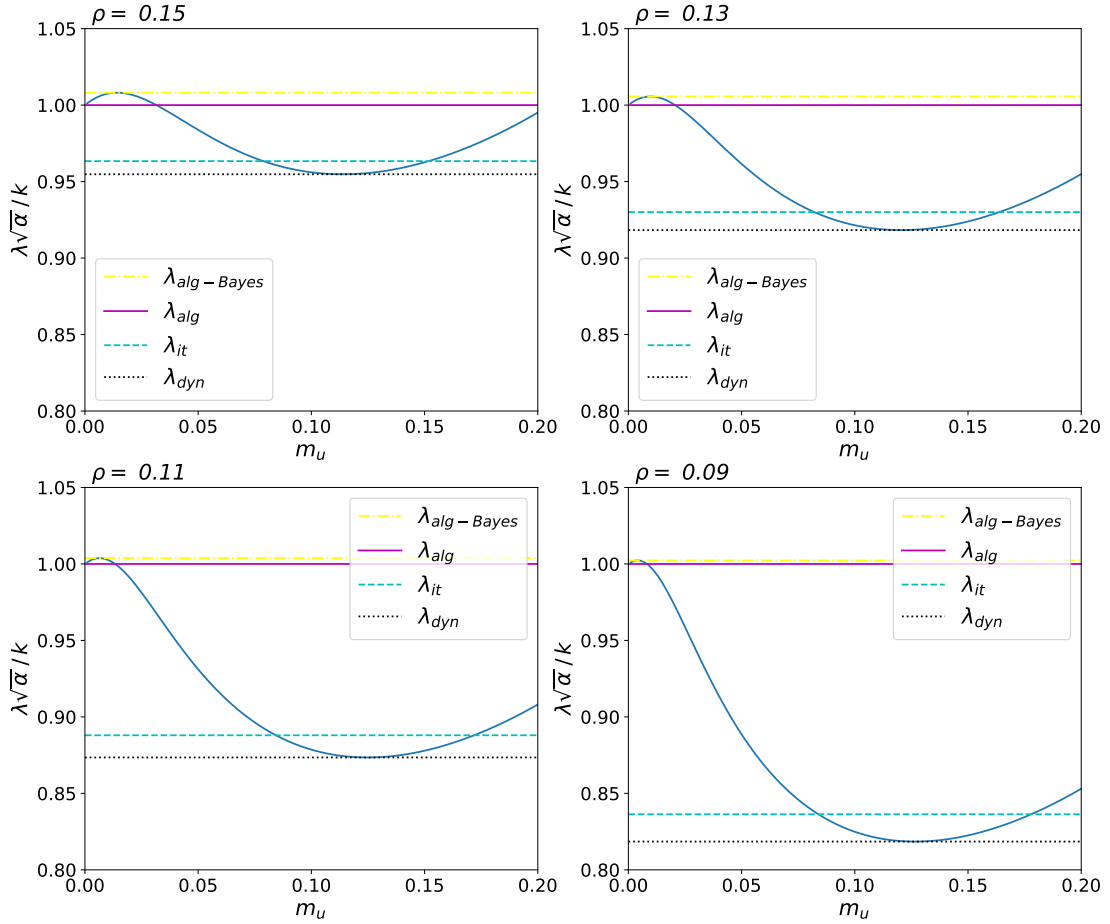


Figure 8: Evolution of the SNR λ as a function of m_u for different $\rho \in \{0.09, 0.11, 0.13, 0.15\}$. We rescale the y-axis by $\sqrt{\alpha}/k$. We plot the different transitions $\{\lambda_{\text{dyn}}, \lambda_{\text{it}}, \lambda_{\text{alg}}, \lambda_{\text{alg-Bayes}}\}$ for each different sparsity level. It is visually clear that the hard phase becomes bigger and bigger as the sparsity grows, while the gap between dynamical and IT threshold closes.

We have to resort to numerical computations for finding the exact values of the thresholds since also in this simple case we do not have a closed-form update for the iterates (m_u^t, m_v^t) , although we will treat them analytically in App. C. We keep in mind the picture in Figs. (5,7) and compute the different thresholds there defined. Let us consider first λ_{it} , the IT threshold, defined as the SNR at which the problem becomes statistically possible. We see in Fig. 5 that it coincides with the SNR level at which the free energy of the two minima (if present) are equal. We analyze for simplicity sparsity levels in which $\lambda_{\text{dyn}} < \lambda_{\text{alg}}$, i.e. we refer to Fig. 5, otherwise the criterion above would define equivalently $\lambda_{\text{Bayes-Jump}}$. We compute the difference of free energy $\Delta\Phi$ between the two minima introducing the path $\gamma : \mathbb{R} \rightarrow \mathbb{R}^2$ which follows the state evolution equations:

$$\gamma(t) = (t, \mathcal{V}(\alpha\lambda t))$$

We can use the fundamental theorem of calculus to obtain the difference of free energy between the trivial fixed point $(m_u, m_v) = (0, 0)$ and a non-trivial one $(m_u, m_v) = (x, \mathcal{V}(\alpha\lambda(x)x))$ at overlap $m_u = x$ as follows:

$$\Delta\Phi(x) = \int_0^x dq \frac{d\Phi}{dt} (m_u(t) = q, m_v(t) = \mathcal{V}(\alpha\lambda(x)q); \lambda(x), \rho, \alpha) \quad (34)$$

where we introduced $\lambda(x)$ as the SNR needed in order to have at fixed (ρ, α) an overlap $m_u = x$ defined by the self-consistent equation:

$$x = \mathcal{U}(\lambda(x)\mathcal{V}(\alpha\lambda(x)x)) \quad (35)$$

Plugging in the expression of the derivative we identify the IT threshold λ_{it} as the minimal SNR such that the following equation is satisfied:

$$\int_0^{x(\lambda_{\text{it}})} dq \mathcal{V}'(\alpha\lambda_{\text{it}}q)[q - \mathcal{U}(\lambda_{\text{it}}\mathcal{V}(\alpha\lambda_{\text{it}}q))] = 0 \quad (36)$$

Let us consider now the Bayes-algorithmical and dynamical thresholds, always referring to their pictorial representation in Figs. (5,7). From a practical standpoint they are stationary point of the the function $\lambda(m_u)$, solution of eq. (35). A reader with some statistical physics background may recognize a parallel with the theory of real gases. The curve $\lambda(m_u)$, exactly as the Pressure-Volume curve $p(v)$ for real gases, is composed of two branches called *stable* and *unstable* branch defined from the value of the derivative $\partial\lambda/\partial m$ (resp. $\partial p/\partial v$). The operative definition of these thresholds as critical points of the curve $\lambda(m_u)$ allows us to easily compute them numerically, see Fig. 8 to observe their evolution as a function of the sparsity level. We see that as we increase the sparsity, i.e. decrease ρ , the statistical-to-computation gap, measured visually by the distance between the two critical points, increase. Moreover the IT threshold collapse with the dynamical one. The same happens with the Bayes-algorithmic threshold, approaching $\lambda_{\text{alg}} \approx 1$.

C Scaling behaviour at large sparsity

We do not have in general an analytical expression for the iterates of the SE equation (m_u^t, m_v^t) appearing in eq. (20), thus we introduced in Sec. 6 a change of variables that allows us to approach analytically the problem in the large sparsity regime. We study in this section the consequences of this scaling assumption. Let us rewrite it here:

$$m_u = \tilde{m}_u \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad m_v = \tilde{m}_v \rho \quad \lambda = C(k)k \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad (37)$$

Consider the update of the parameter m_v under this parametrization:

$$m_v = \rho \tilde{m}_v = f_v^{(k)}(-\tilde{m}_u C(k) \log \rho) \quad (38)$$

We can rewrite the right hand side in the following way:

$$\rho \tilde{m}_v = \rho \frac{-\tilde{m}_u C(k) \log \rho}{(k - \tilde{m}_u C(k) \log \rho)} \int_0^{+\infty} \frac{S_{k-1}}{(2\pi)^{\frac{k}{2}}} \frac{\rho \xi^{k+1} e^{-\xi^2/2}}{\rho + (1-\rho) \left(\frac{k - \tilde{m}_u C(k) \log \rho}{k} \right)^{\frac{k}{2}} \rho^{\tilde{m}_u C(k) \xi^2/2}} d\xi \quad (39)$$

where $S_{k-1}(1)$ is the surface of the k -dimensional unitary hypersphere. Working in the small ρ limit allows us to exploit a concentration in measure over which we integrate on the right hand side. The exponent of ρ in the denominator, i.e. $\frac{\tilde{m}_u C(k) \xi^2}{2}$, will determine the large sparsity behaviour. If the exponent is greater than one, $\rho^{\tilde{m}_u C(k) \xi^2/2}$ will go to zero, otherwise it will diverge in the limit $\rho \rightarrow 0$. Thus we obtain:

$$m_v \approx \rho \int_0^{+\infty} \frac{S_{k-1}(1)}{k(2\pi)^{k/2}} \xi^{k+1} e^{-\xi^2/2} \Theta \left(\frac{\tilde{m}_u C(k)}{2} \xi^2 - 1 \right) := \rho T_k(\tilde{m}_u C(k)) \quad (40)$$

where we introduced $\Theta(x)$ as the Heavyside theta. Plugging in the above expression into the equation defining m_u^t we have:

$$m_u = f_u^{(k)}(\lambda\rho T_k(\tilde{m}_u C(k))) \quad (41)$$

thus approximating for small ρ the function $f_u^{(k)}$, expressing everything in terms of m_u and plugging in the scaling ansatz for m_u in eq. (37), we obtain the simplified SE in the large sparsity regime:

$$\tilde{m}_u = C(k)T_k(\tilde{m}_u C(k)) \quad (42)$$

We can repeat the analysis done in the previous appendix to find the thresholds in this limit. The condition defining the IT threshold λ_{it} written in eq. (34) simplifies to:

$$\int_0^{\tilde{x}} d\tilde{q} T_k'(C_k(\tilde{x})\tilde{q}) \tilde{q} = \int_0^{\tilde{x}} d\tilde{q} T_k'(C_k(\tilde{x})\tilde{q}) C_k(\tilde{x})T_k(C_k(\tilde{x})\tilde{q}) \quad (43)$$

where we defined $C_k(\tilde{x})$ as the value of the coefficient $C(k)$, solution of eq. (42) when the overlap is fixed at $m_u = \sqrt{\frac{-\rho \log \rho}{\alpha}} \tilde{x}$. This task is much easier to solve. Likewise, the computation of the dynamical spinodal simplifies greatly. We need to find the minimum SNR such that eq. (42) has a non trivial solution. By introducing the auxiliary variable $y = C_k(y)\tilde{m}$ we rewrite eq. (42) as:

$$C_k^2(y) = \frac{y}{T_k(y)} \quad (44)$$

thus the minimal SNR to obtain a non trivial solution, defined by the coefficient $C_{\text{dyn}}(k)$, to obtain a non trivial solution of the equation above is given by:

$$C_{\text{dyn}}(k) = \min_{y \in \mathbb{R}^+} \sqrt{\frac{y}{T_k(y)}} \quad (45)$$

At this stage we still need to resort to numerical inspection in order to find the coefficient $C_{\text{it}}, C_{\text{dyn}}$, but we can investigate analytically the large k behaviour. By considering the leading order of the function $T_k(\cdot)$, one can see that $T_k(z) \approx \Theta\left(z - \frac{2}{k+1}\right)$. By plugging in this expression into the definition of the coefficients ($C_{\text{it}}(k), C_{\text{dyn}}(k)$) in eqs. (43),(45) we obtain the following asymptotic result:

$$C_{\text{it}}(k) \approx \sqrt{\frac{4}{k+1}} \quad C_{\text{dyn}}(k) \approx \sqrt{\frac{2}{k+1}} \quad (46)$$

thus plugging them into the scaling assumption in eq. (37) we obtain the following scaling for the thresholds:

$$\lambda_{\text{it}} \approx \sqrt{\frac{4k^2}{k+1}} \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad \lambda_{\text{dyn}} \approx \sqrt{\frac{2k^2}{k+1}} \sqrt{\frac{-\rho \log \rho}{\alpha}} \quad (47)$$

The evolution of the coefficients $C_{\text{it}}, C_{\text{dyn}}$ as a function of the number of clusters is summarized in Fig. 9. We see in Fig. 9 that the large rank expansion is quite accurate also at moderate k , especially for the IT threshold. The easy phase and the alg-Bayes phase merge, hence we will not analyze the distinction between λ_{alg} and $\lambda_{\text{alg-Bayes}}$ in this regime.

D Details on numerical simulations

We discuss in this section the details behind the numerical simulations presented in Sec. 4. The code is available in the [Github repository](#) associated to the manuscript. First we stress an important point on the convergence of low-rAMP algorithms. Increasing the sparsity of the problem, i.e. decreasing ρ , the convergence of AMP becomes more difficult. In order to solve this problem is useful to implement *damping* to stabilize the iteration, see the modified AMP in Algorithm 2. We compared the performance of AMP with different popular algorithm in the literature. The first general-purpose algorithm we considered for subspace clustering is a modification of the sparse PCA algorithm (SPCA). Let us consider a data matrix $Y \in \mathbb{R}^{n \times d}$, where as in our notation n is the number of samples and d is the feature dimension. In the SPCA problem, the statistician wants to find directions in the space which maximize the variance of our dataset by constraining the cardinality of the new basis vectors. In vanilla PCA instead we try to find directions, called principal components $\{\hat{e}_m\}_{m=1}^d$, which maximize the variance not caring if they will be given by linear combination of all the features of

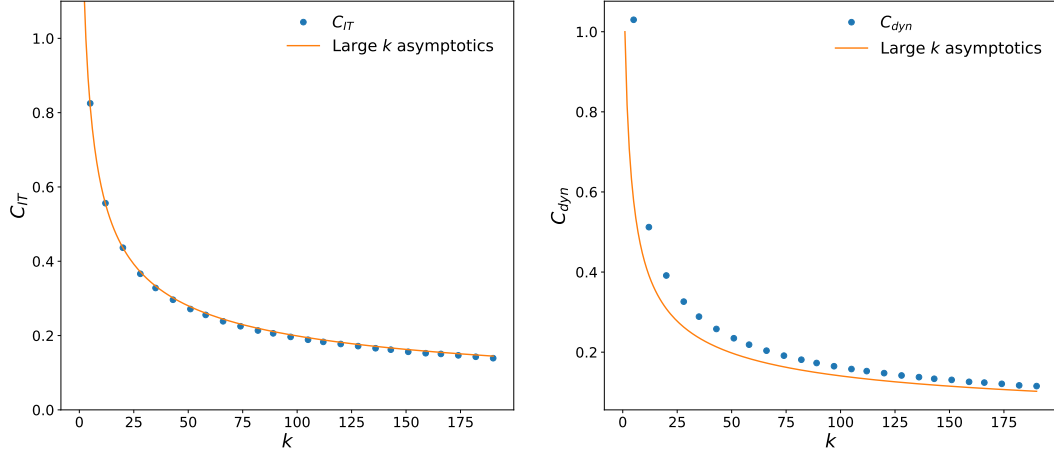


Figure 9: *Left*: Comparison of the threshold coefficients (C_{it}, C_{dyn}) (dots) with their high rank asymptotic expression (solid line).

our problem: $\hat{e}_m = \sum_{i=1}^d \alpha_i^{(m)} e_i$, where we called $\{e_i\}_{i=1}^d$ the canonical basis vectors. In SPCA we want that some of the coefficients $\alpha_i^{(m)}$ (called "loadings" in the literature) to be zero, favouring interpretability of the optimal estimator. By formulating in a variational way the problem the sparsity of the estimator is enhanced using LASSO regularization. We write the pseudocode for the program we used in the two-class subspace clustering problem in Algorithm. 3. The unregularized problem, i.e. $\Gamma = 0$, is equivalent to vanilla PCA. The comparison of the performances of the two spectral algorithms has been done in Fig. 2 and we have a clear advantage in imposing the cardinality constraint as the sparsity level increase. We considered in the sub-extensive sparsity regime in Sec. 6 the Diagonal Thresholding algorithm (DT). The main idea is to search for spatial directions with the largest variance, and threshold the sample covariance matrix accordingly, hence the name Diagonal Thresholding. The pseudocode for the algorithm we used in the two-classes subspace clustering is given in Algorithm. 4.

Algorithm 2 low-rAMP with damping

Input: Data $X \in \mathbb{R}^{d \times n}$

Initialize $\hat{\mathbf{v}}_i^{t=0}, \hat{\mathbf{u}}_\nu^{t=0} \sim \mathcal{N}(\mathbf{0}_k, \epsilon \mathbf{I}_k)$, $\hat{\sigma}_{u,\nu}^{t=0} = \mathbf{0}_{k \times k}$, $\hat{\sigma}_{v,i}^{t=0} = \mathbf{0}_{k \times k}$.

for $t \leq t_{\max}$ **do**

$$\mathbf{A}_u^{\text{tmp}} = \frac{\lambda}{s} (\hat{\mathbf{U}}^t)^\top \hat{\mathbf{U}}, \quad \mathbf{A}_v^{\text{tmp}} = \frac{\lambda}{s} (\hat{\mathbf{V}}^t)^\top \hat{\mathbf{V}}$$

$$\mathbf{B}_v^{\text{tmp}} = \sqrt{\frac{\lambda}{s}} X \hat{\mathbf{U}}^t - \frac{\lambda}{s} \sum_{\nu=1}^n \sigma_{u,\nu}^t \hat{\mathbf{V}}^{t-1}, \quad \mathbf{B}_u^{\text{tmp}} = \sqrt{\frac{\lambda}{s}} X^\top V - \frac{\lambda}{s} \sum_{i=1}^d \sigma_{v,i}^t \hat{\mathbf{U}}^{t-1}$$

Damping step with damping coefficient γ :

$$\mathbf{A}_u^t = (1 - \gamma) \mathbf{A}_u^{\text{tmp}} + \gamma \mathbf{A}_u^{t-1}, \quad \mathbf{A}_v^t = (1 - \gamma) \mathbf{A}_v^{\text{tmp}} + \gamma \mathbf{A}_v^{t-1}$$

$$\mathbf{B}_u^t = (1 - \gamma) \mathbf{B}_u^{\text{tmp}} + \gamma \mathbf{B}_u^{t-1}, \quad \mathbf{B}_v^t = (1 - \gamma) \mathbf{B}_v^{\text{tmp}} + \gamma \mathbf{B}_v^{t-1}$$

Take $\{\mathbf{b}_{v,i}^t \in \mathbb{R}^k\}_{i=1}^d, \{\mathbf{b}_{u,\nu}^t \in \mathbb{R}^k\}_{\nu=1}^n$ rows of $\mathbf{B}_v^t, \mathbf{B}_u^t$

$$\hat{\mathbf{v}}_i^{t+1} = \eta_v(\mathbf{A}_v^t, \mathbf{b}_{v,i}^t), \quad \hat{\mathbf{u}}_\nu^{t+1} = \eta_u(\mathbf{A}_u^t, \mathbf{b}_{u,\nu}^t)$$

$$\hat{\sigma}_{v,i}^{t+1} = \partial_{\mathbf{b}} \eta_v(\mathbf{A}_v^t, \mathbf{b}_{v,i}^t), \quad \hat{\sigma}_{u,\nu}^{t+1} = \partial_{\mathbf{b}} \eta_u(\mathbf{A}_u^t, \mathbf{b}_{u,\nu}^t)$$

Here $\hat{\mathbf{U}}^t \in \mathbb{R}^{n \times k}, \hat{\mathbf{V}}^t \in \mathbb{R}^{d \times k}, \mathbf{B}_u^t \in \mathbb{R}^{n \times k}, \mathbf{B}_v^t \in \mathbb{R}^{d \times k}, \mathbf{A}_u^t \in \mathbb{R}^{k \times k}, \mathbf{A}_v^t \in \mathbb{R}^{k \times k}$

end for

Return: Estimators $\hat{\mathbf{v}}_{\text{amp},i}, \hat{\mathbf{u}}_{\text{amp},\nu} \in \mathbb{R}^k, \hat{\sigma}_{u,\nu}, \hat{\sigma}_{v,i} \in \mathbb{R}^{k \times k}$

Algorithm 3 SPCA

Input: Data $\mathbf{Y} \in \mathbb{R}^{n \times d}$

Initialize $\Delta_{\text{sparsity}} = 1, \Gamma = 10^{-3}$

while $|\Delta_{\text{sparsity}}| \geq 1$ **do**

Solve variational problem: $(\hat{\mathbf{C}}, \hat{\mathbf{D}}) = \arg \min_{\mathbf{C} \in \mathbb{R}^n, \mathbf{D} \in \mathbb{R}^d} \{\|\mathbf{Y} - \mathbf{C}\mathbf{D}^\top\|_F + \Gamma \|\mathbf{D}\|_1\}$

Compute first sparse principal component \mathbf{D}

Compute the estimated sparsity $\hat{s} = \sum_{i=1}^d (1 - \delta_{\hat{v}_i,0})$

Compute sparsity mismatch $\Delta_{\text{sparsity}} = \rho d - \hat{s}$

If $\Delta_{\text{sparsity}} < 0$ decrease Γ , otherwise increase it

end while

Project the data matrix onto the first sparse principal component: $\mathbf{P} = \mathbf{Y}\mathbf{D} \in \mathbb{R}^n$

Compute cluster membership assignment: $\hat{\mathbf{U}} = \text{sign}(\mathbf{P})$

Return: $\text{MSE}(\hat{\mathbf{U}})$

Algorithm 4 Diagonal Thresholding

Input: Data $\mathbf{Y} \in \mathbb{R}^{n \times d}$

Compute the sample covariance matrix $\hat{\mathbf{K}} = \frac{1}{n} \sum_{\nu=1}^n \mathbf{y}_\nu \mathbf{y}_\nu^\top$

Find the directions with the s largest variance, with $s = \lfloor \rho d \rfloor$.

Call the subset of indices corresponding to the directions above \mathcal{S} .

Create $\tilde{\mathbf{K}}$:

$$\tilde{\mathbf{K}}_{ij} = \begin{cases} \hat{\mathbf{K}}_{ij} & \text{if } (i, j) \in \mathcal{S} \\ 0 & \text{otherwise} \end{cases}$$

Compute the largest eigenvector of the thresholded matrix $\tilde{\mathbf{K}}$ and call it $\hat{\mathbf{v}}$.

Project the data matrix onto the first sparse principal component: $\mathbf{p} = \mathbf{Y}\hat{\mathbf{v}} \in \mathbb{R}^n$

Compute cluster membership assignment: $\hat{\mathbf{u}} = \text{sign}(\mathbf{p})$

Return: $\text{MSE}(\hat{\mathbf{u}})$
