


## Cascade of phase transitions for multiscale clustering

Tony Bonnaire<sup>1,2</sup>, Aurélien Decelle<sup>2,3</sup> and Nabila Aghanim<sup>1</sup>

<sup>1</sup>Université Paris-Saclay, CNRS, Institut d'Astrophysique Spatiale, 91405 Orsay, France

<sup>2</sup>Université Paris-Saclay, TAU Team INRIA Saclay, CNRS, Laboratoire de Recherche en Informatique, 91190 Gif-sur-Yvette, France

<sup>3</sup>Departamento de Física Teórica I, Universidad Complutense, 28040 Madrid, Spain

 (Received 13 August 2020; revised 15 October 2020; accepted 2 December 2020; published 6 January 2021)

We present a framework exploiting the cascade of phase transitions occurring during a simulated annealing of the expectation-maximization algorithm to cluster datasets with multiscale structures. Using the weighted local covariance, we can extract, *a posteriori* and without any prior knowledge, information on the number of clusters at different scales together with their size. We also study the linear stability of the iterative scheme to derive the threshold at which the first transition occurs and show how to approximate the next ones. Finally, we combine simulated annealing together with recent developments of regularized Gaussian mixture models to learn a principal graph from spatially structured datasets that can also exhibit many scales.

DOI: [10.1103/PhysRevE.103.012105](https://doi.org/10.1103/PhysRevE.103.012105)

### I. INTRODUCTION

Many optimization and inference problems have been shown to have an equivalent formulation in statistical physics [1,2] that allowed a brand-new look at some long-standing problems and improved the understanding of complex systems [3,4]. In particular, the identification of the phase diagram of a model can bring interesting new insights such as knowing if a given information can be retrieved depending on the model's parameters and the dataset at hand. In the context of clustering, it has been demonstrated that the Gaussian mixture model (GMM) can be formulated as a statistical mechanics problem [5,6] where the negative log-likelihood can be interpreted as a free energy.

GMMs are extensively used in many fields of science such as physics, statistics, and machine learning, mainly for complex density estimation [7] or for unsupervised tasks such as clustering [8]. Some attempts to solve other classes of NP-hard problems such as the traveling salesman [9] or to model principal curves as an extension of principal components [10] were also conducted. Standard methods to fit model parameters rely on the expectation-maximization (EM) algorithm [11], an iterative procedure maximizing the log-likelihood with guaranteed convergence toward a local maximum [12]. However, the direct application of the EM algorithm is known to be easily trapped in local maxima of multimodal likelihoods leading to variability in the provided results depending on the initialization of the algorithm [13,14]. This problem, coupled with the main drawback of mixture modeling, i.e., the choice of the number of components to model the data, make parametric mixture models very sensitive to the initialization and the choice of hyperparameters. In that regard, the statistical physics formulation of the clustering problem helped to overcome these issues by making use of deterministic simulated annealing allowing the relaxation of the nonconvex optimization problem by solving

it iteratively while a parameter—assimilated to a temperature, in our case, the variance of all components—is slowly reduced [15].

Following these steps, we aim at showing how the latter formulation can be useful to understand and analyze the outcome of GMMs. In particular, we exploit the cascade of phase transitions occurring during annealing procedures of the EM algorithm to build a hierarchical multiscale description of a dataset. By defining an overlap between the ground truth and the inferred partitions, we show on artificial datasets how it can be interpreted as an order parameter whose value follows the sequence of phase transitions. In more general cases, where the ground truth is not known, we use a physical observable to get an *a posteriori* interpretation of what happened during the annealing. For a given mixture component, we track the hierarchy of scales that are being represented through the eigenvalues of the weighted covariance matrix and use it to extract information on the structure, i.e., the number of clusters, their scale, and their hierarchy without prior information on the “correct” number of clusters, even in high dimension when direct visualization is impossible. To represent different scales simultaneously, we propose an alternative to the classical annealing in which the temperature is the mode of a prior distribution on variances. We show that the threshold at which the first transition occurs can be computed exactly from the linear stability of the fixed-point iterative scheme in both cases. Approximate positions of successive thresholds can also be estimated when the dataset is clustered into several subsystems. The proposed methodology hence leads to a representation (i) independent from the inputted number of components, (ii) independent from the data space dimensions, and (iii) hierarchically nested from high scales to lower ones.

In Sec. II, we introduce the mathematical background of GMMs and review the main equations of the EM algorithm

to estimate parameters of the model. Section III presents the hard annealing setting and motivates the use of the procedure to learn hierarchical representations. Section IV focuses on the interest of combining a regularized version of the GMM with the annealing procedure for the study of complex datasets and provides theoretical derivations of the threshold at which the first transition occurs. Finally, in Sec. V, we illustrate how these transitions can be used together with the recently introduced graph regularization of Gaussian mixtures (GMs) to learn a principal graph structure from spatially structured datasets, even with high heteroscedasticity of the sampling.

## II. GAUSSIAN MIXTURE MODELS AND EXPECTATION-MAXIMIZATION ALGORITHM

It is not always possible to fit known probability distributions to those resulting from a set of physical measurements  $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$  with  $\mathbf{x}_i \in \mathbb{R}^D$ . GMMs naturally allow the modeling of complex probability distributions as a linear combination of  $K$  Gaussian components with unknown parameters. In what follows, we restrict ourselves to spherical and uniform GMs with means  $\{\boldsymbol{\mu}_k\}_{k=1}^K$  and variances  $\{\sigma_k^2\}_{k=1}^K$  such that

$$p(\mathbf{x}_i | \Theta) = \frac{1}{K} \sum_{k=1}^K \mathcal{N}(\mathbf{x}_i, \boldsymbol{\theta}_k), \quad (1)$$

with  $\Theta = \{\boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_k\}$  the set of parameters with  $\boldsymbol{\theta}_k = (\boldsymbol{\mu}_k, \sigma_k)$ .

EM states that the log-likelihood can be maximized through an iterative procedure involving two alternating steps by considering a set of latent variables  $\{z_i\}_{i=1}^N$  with  $z_i \in \{0, \dots, K\}$  describing the label of the cluster that generated the datapoint  $\mathbf{x}_i$ . The  $E$  step first estimates the posterior distribution of the latent variables from the current values of parameters as

$$p_{ik} = p(z_i | \mathbf{x}_i, \boldsymbol{\theta}_k) = \frac{\exp(-\|\mathbf{x}_i - \boldsymbol{\mu}_k\|_2^2 / 2\sigma_k^2)}{\sum_{j=1}^K \exp(-\|\mathbf{x}_i - \boldsymbol{\mu}_j\|_2^2 / 2\sigma_j^2)}. \quad (2)$$

The  $M$  step then refines parameter estimates based on the current values of  $p_{ik}$  as

$$\begin{aligned} \boldsymbol{\mu}_k &= \frac{\sum_{i=1}^N p_{ik} \mathbf{x}_i}{N_k}, \\ \sigma_k^2 &= \frac{\sum_{i=1}^N p_{ik} \|\mathbf{x}_i - \boldsymbol{\mu}_k\|_2^2}{DN_k}, \end{aligned} \quad (3)$$

where  $N_k = \sum_{i=1}^N p_{ik}$ .

From the statistical physics point of view, identical update equations can be obtained by deriving the free energy under a quadratic cost function for the datapoints assignment to clusters [1].

## III. PHASE TRANSITIONS IN HARD ANNEALING

Simulated annealing in the EM procedure was introduced to overcome the issue of local maxima. It consists in reducing iteratively the variance of all components such that  $\forall k \in \{1, \dots, K\}$ ,  $\sigma_k^2 = \sigma^2$ , where  $\sigma^2$  is the controlled value of the variance during the annealing. The idea behind these

approaches is to smooth the likelihood by starting with a very high variance leading to a concave function. Decreasing it slowly leads to a finer and finer description of the dataset, hence resulting in a more complex likelihood function with multiple modes appearing.

In this work, we focus on a particular aspect of the annealing based on successive phase transitions. First, we are interested in the range of  $\sigma^2$  values for which the likelihood is concave and hence all of the  $K$  components are collapsed into a single location centered at  $1/N \sum_{i=1}^N \mathbf{x}_i$ . This critical quantity, noted  $T_c^{\text{hard}}$ , is known to be the maximum eigenvalue of the data covariance matrix [5]. When  $\sigma^2 > T_c^{\text{hard}}$ , even though  $K$  components are used in the model, they are all collapsed as  $K_r = 1$  physical cluster at the center of mass (c.m.) of the dataset. When  $\sigma^2$  becomes slightly smaller than  $T_c^{\text{hard}}$ , the likelihood is deformed and centers get aligned with the first principal direction given by the data covariance matrix  $\mathbf{C}$ . When  $\sigma^2$  continues to decrease, the dataset description becomes more and more detailed and  $K_r$  takes increasing values.

We propose a way to extract information on the structure of a dataset from the annealing process by tracking the evolution of the size represented by a given component  $k$  through the maximum eigenvalue  $\Gamma_k$  of its weighted covariance matrix, namely,  $\boldsymbol{\Sigma}_k = 1/N_k \sum_{i=1}^N p_{ik} (\mathbf{x}_i - \boldsymbol{\mu}_k) (\mathbf{x}_i - \boldsymbol{\mu}_k)^\top$ .

Figure 1 illustrates the evolution of the ratio  $\Gamma_k / \sigma^2$  during the annealing for an artificial dataset with five clusters and  $K = 25$  centers colored by their end-point cluster. In Fig. 1(b), we see the cascade of transitions and successive splitting of centers when  $\sigma^2$  decreases. When two or more centers collapse, they share similar values of  $\boldsymbol{\mu}_k$  and  $p_{ik}$ , leading to similar evaluations of  $\Gamma_k$ . This is why all lines are superimposed for  $\sigma^2 > T_c^{\text{hard}}$ . Each time a curve reaches the horizontal unit line, one of the  $K_r$  subsystems made of collapsed components reached the temperature of the subdataset it represents, namely,  $\sigma^2 \simeq \Gamma_k$ . From there, we observe either a bounce or a cross of the line. When bouncing, there is a split between two populations of centers that were representing the same part of the dataset but that will take different paths. Centers thus move toward a smaller cluster and the value of  $\Gamma_k$  decreases. Crossing the line occurs when centers split inside an individual cluster due to its inner random structure. In that case, the imposed variance gets smaller than the physical one. Since the transitions are driven by the maximum eigenvalues of the empirical covariance matrices, it is this quantity that is plotted as vertical lines for the size of spherical clusters in the figures. Note that if we use instead the empirical variance, this estimate would be slightly shifted to lower values and the transition would occur earlier in the annealing than the true variance of the cluster.

Following *a posteriori* the several curves and the successive transitions provide an informative insight on the structure of the dataset. It also allows one to visualize the evolution of the local size representation of the data and the interactions between centers. In Fig. 1(b), we clearly see that {purple, red, green} sets of centers represent the same information when  $\sigma^2 > 9$  and then split into {purple} and {red, green}. This indicates the presence of a subsystem of two clusters. Later, we observe a crossing of the horizontal line

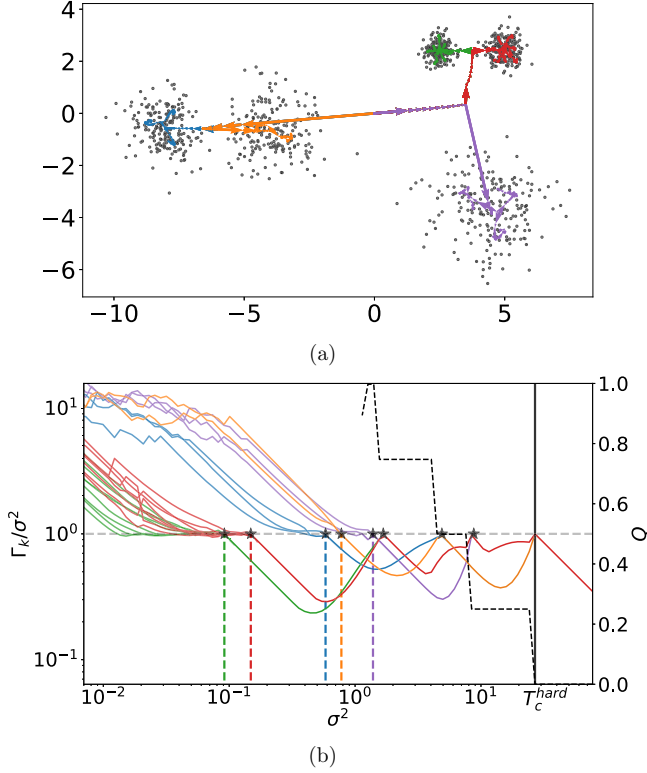


FIG. 1. (a) Displacement of  $K = 25$  centers during the annealing procedure for a dataset with five spherical Gaussian clusters. Colors indicate in which final cluster the center ends. (b) Evolution of the ratio  $\Gamma_k/\sigma^2$  as a function of  $\sigma^2$ . Black stars correspond to the scales of successive transitions, the black vertical line to  $T_c^{\text{hard}}$ , and colored ones indicate the size of the clusters as defined by the maximum eigenvalue of the empirical covariance. The black dashed curve shows the evolution of  $Q$  as defined in Eq. (4) that we identify as an order parameter. This quantity is not represented for  $\sigma \leq 1$  since the number of physical clusters  $K_r$  begins to be higher than the number of generated clusters  $q$ .

for the {purple} centers before splitting again after crossing. This indicates that the effective variance of the cluster is larger than the one fixed by the annealing and, therefore, that these centers now describe fluctuations within a “true” cluster. The {red, green} sets of centers split at lower  $\sigma^2$  followed as well by a crossing of the line at the scale of individual cluster sizes.

Successive transitions can be computed in two steps: First by identifying the  $K_r$  macrocomponents resulting from the collapse of centers based on their positions and then by assigning to each data point the label of the macrocomponent that most probably generated it. We are thus assuming, at a given iteration, a GMM with  $K_r$  components to compute responsibilities from Eq. (2). Hence, we can group data points with identical labels and compute the next transition as the maximum eigenvalue of the covariance matrix for each subsystem. These quantities, basically corresponding to successive evaluations of the critical temperature in subsystems, can be approximated during the annealing and are shown as black stars in Fig. 1(b).

We identify the overlap  $Q$ , defined as the quality of the data classification at each temperature, as an order parameter




FIG. 2. Ratio between the estimated variances  $\hat{\sigma}^2$  obtained when freezing the  $K = 25$  centers when  $\Gamma_k/\sigma^2 \simeq 1$  and the empirical ones  $\sigma_{\text{true}}^2$  from data of Fig. 1 for several values of  $\rho$ , the proportion of datapoints left for the computation.

whose value changes throughout the several phases during the annealing. Formally,

$$Q(\{\hat{z}_i\}, \{z_i\}) = \frac{\max_{\pi \in \Pi} \frac{1}{N} \sum_i \delta_{\hat{z}_i, \pi(z_i)} - 1/q}{1 - 1/q}, \quad (4)$$

where  $\delta$  is the Kronecker delta,  $\Pi$  denotes all the possible permutations of the set  $\{1, \dots, q\}$  with  $q$  the true number of clusters used to generate the data and  $\hat{z}_i = \arg\max_k p_{ik}$  the estimated latent variable for the affiliation of the datapoint  $x_i$ .  $Q$  is based on the responsibilities  $p_{ik}$ , all being  $1/K$  for a random assignation giving  $Q = 0$  and then going to either 1 or 0 when  $\forall i, \hat{z}_i = z_i$  leading to  $Q = 1$ . By doing so,  $Q$  is zero when  $\sigma^2 > T_c^{\text{hard}}$  and undergoes successive transitions as the system is cooled down, as illustrated in Fig. 1(b). During the annealing,  $Q$  remains at 1 for the range of  $\sigma^2$  between the last split of centers between two true clusters and before the first split due to the inner random structure of one of them. Note that this metric is not applicable when  $K_r > q$  since the dataset is partitioned into more clusters than actually used for the generation, and this is why the curve is not shown for  $\sigma^2 \leq 1$ .

To further assess the robustness and accuracy of the transitions, we use the dataset from Fig. 1(a) where only a fraction  $\rho$  of the datapoints is randomly kept for the computation. During the annealing, we freeze all the  $K = 25$  centers at the last split before reaching  $\Gamma_k/\sigma^2 = 1$  and then let variances evolve freely hence providing an estimate for each detected cluster that we note  $\hat{\sigma}^2$ . Figure 2 shows that the retrieved variances are, even in highly sparse sampling settings, with  $\rho \leq 30\%$ , close to the true ones of the clusters. It is worth emphasizing that all five clusters are always correctly identified and that the value of  $Q$  is always close to 1 at the end of the process, showing the ability of the method to highlight structures, even in sparse configurations.

Usual applications of GMMs for clustering are performed blindly by inputting the desired  $K$  to obtain a classification

making use of all components. Some criteria, based on information theory [16,17] or Bayesian approaches [18,19], were proposed to overcome this major drawback of unsupervised clustering. Here, we propose an approach to avoid such a  $K$ -dependent unique solution. A key aspect of the annealing is the collapse of  $\{\mu_k\}_{k=1}^K$  at the c.m. of successive subdatasets providing a hierarchical view of clustering with an increasing number of physical clusters. The proposed diagram enables one to capture this set of nested representations as shown in Fig. 3 for a five-dimensional (5D) artificial dataset made of ten Gaussian clusters, spatially appearing as three clusters at larger scale. Transitions occurring at large scales in each panel (gray vertical lines) clearly indicate that the dataset is described as three different physical clusters. Pursuing the decrease of variances leads to a finer description where each of the three macroclusters splits into smaller ones that still have physical interpretations [20]. Such information on the spatial organization of the dataset is of crucial importance when having no prior idea of its structure or the number of underlying components.

#### IV. PHASE TRANSITIONS IN SOFT ANNEALING

When the dataset is more complex with nested structures or overlapping clusters of different sizes, the previously presented analysis is not suitable. Multiple scales cannot be represented at the same time, hence biasing those of embedded structures toward higher values. To overcome this, we rely on a modified annealing acting on the mode of an *a priori* distribution on each  $\sigma_k^2$ . More particularly, we use the conjugate prior for variances, namely, an inverse-gamma distribution, with shape parameter  $1 + \lambda_\sigma$  and scale parameter  $\lambda_\sigma \sigma^2$  so that the distribution has a mode at  $\sigma^2$ . Formally, it reads

$$\ln p(\sigma_k^2) = -\lambda_\sigma \left[ \ln \sigma_k^2 + \frac{\sigma^2}{\sigma_k^2} \right] + \text{cte}, \quad (5)$$

where the constant comes from the normalization of the probability distribution. Introducing such a prior modifies the update Eq. (3) for variances as

$$\sigma_k^2 = \frac{\sum_{i=1}^N p_{ik} \|\mathbf{x}_i - \boldsymbol{\mu}_k\|_2^2 + 4\lambda_\sigma \sigma^2}{D \sum_{i=1}^N p_{ik} + 4\lambda_\sigma}. \quad (6)$$

Consequently, when  $\lambda_\sigma \rightarrow 0$ , the prior, and hence the annealing, has no effect and components update their variances as Eq. (3). Inversely, for a large enough value,  $\sigma_k^2$  will be close to  $\sigma^2$  resulting in the classical annealing procedure. Choosing intermediate values for  $\lambda_\sigma$  hence imposes a broad trend for all components but lets each of them correct the prior by the actual value of the neighboring covariance. In what follows, we refer to this procedure as “soft annealing” that we distinguish from the “hard annealing” to describe the classical procedure acting directly on the variance parameter. There is no rule to fix the hyperparameter  $\lambda_\sigma$  and, in this work, we adopt  $\lambda_\sigma = 2$  [21].

Similarly as in the hard annealing case, we can compute the threshold value  $T_c^{\text{soft}}$  such that all components are collapsed at the c.m. when  $\sigma^2 > T_c^{\text{soft}}$ . Without any loss of generality, the dataset can be considered centered, with  $\sum_{i=1}^N \mathbf{x}_i = \mathbf{0}_D$  where





FIG. 3. (a) Displacement of  $K = 25$  centers during the annealing procedure for a dataset made of ten 5D spherical Gaussian clusters visualized in the plane of the two first principal components. Colors depend on the macrocluster the component stands in at the last iteration. (b) Evolution of the ratio  $\Gamma_k/\sigma^2$  as a function of  $\sigma^2$ , the hard annealing parameter. Colored vertical lines indicate the actual size of the corresponding Gaussian cluster or macrocluster (gray dashed lines, first from the right in each panel) as defined by the maximum eigenvalue of the empirical covariance.

$\mathbf{0}_D$  is the  $D$ -dimensional zero vector. In a first step, we propose to derive the fixed-point variance  $\sigma_0^2$  of all centers when  $\sigma^2 \gg T_c^{\text{soft}}$ . Linearly Taylor expanding the expression of  $p_{ik}$  given by Eq. (2) for small perturbations  $\boldsymbol{\mu}_k \approx \mathbf{0}_D$  leads to

$$p_{ik} = \frac{1}{K} \left[ 1 + \frac{1}{\sigma_k^2} \mathbf{x}_i^T \boldsymbol{\mu}_k \right] + o(\boldsymbol{\mu}_k). \quad (7)$$

Further, assuming that  $\sigma^2 > T_c^{\text{soft}} \Rightarrow \forall k \in \{1, \dots, K\}, \sigma_k^2 = \sigma_0^2$  and injecting  $p_{ik}$  expansion in update Eq. (6) gives

$$\sigma_0^2 = \frac{4\lambda_\sigma K \sigma^2 + \sum_{i=1}^N \mathbf{x}_i^T \mathbf{x}_i}{ND + 4\lambda_\sigma K}. \quad (8)$$

This equation links  $\sigma_0^2$ , the actual variance attributed to all Gaussian components, to  $\sigma^2$ , the soft annealing parameter, and is valid in the large  $\sigma^2$  limit. By considering small perturbations  $\boldsymbol{\epsilon}_k$  and  $\delta_k$ , respectively around the fixed points  $\boldsymbol{\mu}_k = \mathbf{0}_D$  and  $\sigma_k^2 = \sigma_0^2$ , we can derive the set of equations for the vectorized perturbations  $\underline{\boldsymbol{\epsilon}} = (\boldsymbol{\epsilon}_0^T, \dots, \boldsymbol{\epsilon}_K^T)^T \in \mathbb{R}^{KD \times 1}$  and  $\boldsymbol{\delta} = (\delta_0, \dots, \delta_K) \in \mathbb{R}^{K \times 1}$ ,

$$\begin{aligned} \underline{\boldsymbol{\epsilon}}^{(t+1)} &= \frac{1}{\sigma_0^2} (\mathbf{U} \otimes \mathbf{C}) \underline{\boldsymbol{\epsilon}}^{(t)} + (\mathbf{U} \otimes \mathbf{a})^T \boldsymbol{\delta}^{(t)}, \\ \boldsymbol{\delta}^{(t+1)} &= (\mathbf{U} \otimes \mathbf{b}) \underline{\boldsymbol{\epsilon}}^{(t)} + c \mathbf{U} \boldsymbol{\delta}^{(t)}, \end{aligned} \quad (9)$$

where  $\otimes$  denotes the Kronecker product,  $\mathbf{U} = (\mathbf{I}_K - \frac{1}{K} \mathbf{J}_K)$  with  $\mathbf{I}_K$  is the  $K \times K$  identity matrix,  $\mathbf{J}_k$  the  $K \times K$  all-ones matrix, and

$$\mathbf{a} = \sum_i \frac{\|\mathbf{x}_i\|_2^2 \mathbf{x}_i^T}{2N\sigma_0^4}, \quad (10)$$

$$\mathbf{b} = \sum_i \frac{\|\mathbf{x}_i\|_2^2 \mathbf{x}_i^T}{m\sigma_0^2}, \quad (11)$$

$$c = \frac{1}{2\sigma_0^2 m} \left( \sum_i \frac{\|\mathbf{x}_i\|_2^4}{\sigma_0^2} - D \sum_i \|\mathbf{x}_i\|_2^2 \right), \quad (12)$$

with  $m = ND + 4\lambda_\sigma K$ . Putting it all together leads to the matrix representation of the system's perturbations  $\underline{\boldsymbol{\eta}} = (\underline{\boldsymbol{\epsilon}}, \boldsymbol{\delta}) \in \mathbb{R}^{K(D+1) \times 1}$ ,

$$\underline{\boldsymbol{\eta}}^{(t+1)} = (\mathbf{U} \otimes \mathbf{M}) \underline{\boldsymbol{\eta}}^{(t)}, \quad (13)$$

with  $\mathbf{M}$  the squared block matrix of order  $D + 1$ ,

$$\mathbf{M} = \left( \begin{array}{c|c} \mathbf{C}/\sigma_0^2 & \mathbf{a}^T \\ \hline \mathbf{b} & c \end{array} \right), \quad (14)$$

where  $\mathbf{C}$  is the data covariance matrix.

Since the eigenvalues of the Kronecker product are given by the product of all individual eigenvalues of the two matrices involved and that  $\mathbf{U}$  has eigenvalues 0 or 1, we can only restrict the analysis to those of  $\mathbf{M}$ . Therefore, the value of  $\sigma^2$  at which the first transition occurs, namely,  $T_c^{\text{soft}}$ , can be derived as the value of  $\sigma^2$  such that the spectral radius of  $\mathbf{M}$  is 1, leading to instabilities in the dynamic of the system.

Figure 4 illustrates the result of the soft annealing procedure on an artificial dataset made of six clusters similar to Fig. 1(a) but with more complexity such as overlapping and nested clusters. The bottom panel focuses on the evolution of the ratio between the size of the represented subsystem by a

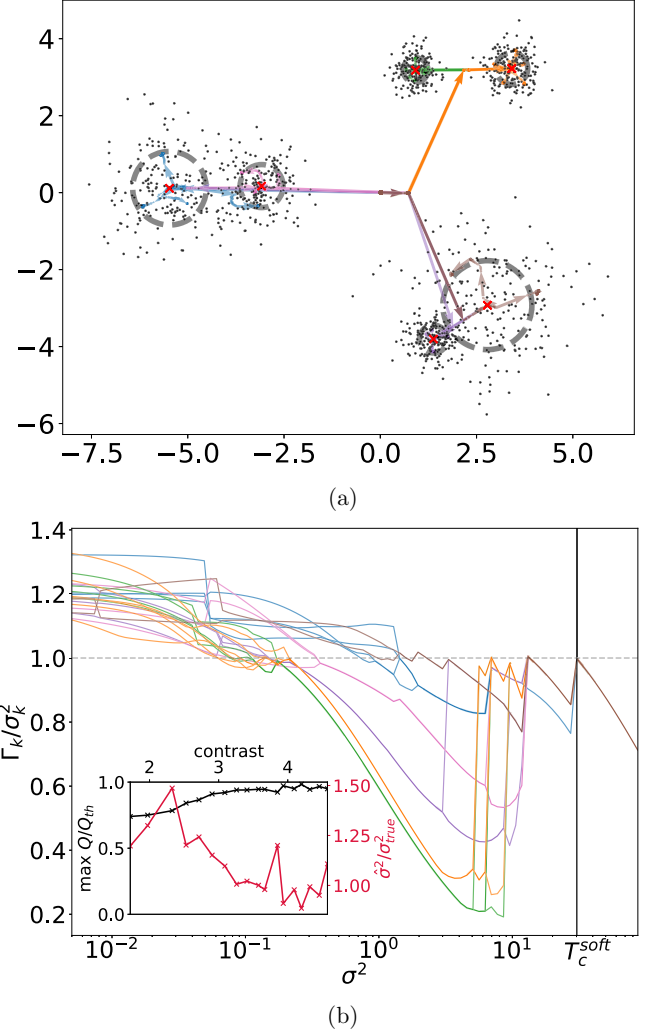


FIG. 4. (a) Arrows indicate the displacement of  $K = 25$  centers during the soft annealing procedure for a dataset made of six spherical Gaussian clusters (black points). Colors relate to the cluster in which the component ends. Red crosses and gray dashed circles respectively indicate the positions and variances fixed *a posteriori* when the center undergoes its last split before remaining above the  $\Gamma_k/\sigma_k^2 = 1$  line. (b) Evolution of the ratio  $\Gamma_k/\sigma_k^2$  as a function of  $\sigma^2$ . The vertical black line corresponds to  $T_c^{\text{soft}}$ . The inset figure shows the evolution of the ratios  $\max Q/Q_{\text{th}}$  and  $\hat{\sigma}^2/\sigma_{\text{true}}^2$ , when varying the contrast between the two nested clusters.

given component and its actual variance, namely,  $\Gamma_k/\sigma_k^2$ . This is the same physical quantity as in the hard annealing case, except that we relax the constraint on  $\sigma_k^2$  which is now varying for each component. In this soft configuration, all the  $\boldsymbol{\mu}_k$  are collapsed for  $\sigma^2 > T_c^{\text{soft}}$  followed by steep transitions when  $\sigma^2$  decreases. This relaxed annealing is especially useful for the representation of the two nested clusters. Even though we would learn that those structures are encapsulated, reaching accurate size description for the smallest component would not be possible in hard annealing because its variance would be boosted by neighboring data points. Figure 4(a) illustrates positions and variances of all  $K = 25$  components when fixing parameters *a posteriori* at the value they had during the annealing at their last transition point just before crossing the

line  $\Gamma_k/\sigma_k^2 = 1$ . Although  $K = 25$  components are used, we correctly identify  $K_r = 7$  physical clusters with their variances and means as indicated by red crosses and gray circles in Fig. 1(a).

To assess the robustness of the soft annealing procedure in clustering complex datasets, we focus on a setup restricted to the two nested clusters of Fig. 4(a) only and dilute the small one by varying its number of sampling points  $N$ . This translates into a decreasing contrast between the signal-to-noise ratios  $\sigma/\sqrt{N}$  of the two clusters. The inset of Fig. 4(b) shows the evolution of the ratio between the maximum overlap value  $Q$  obtained during the annealing and  $Q_{\text{th}}$  the theoretical overlap computed using the ground truth parameters as a function of the contrast. It can be seen that both clusters are recovered when the contrast is sufficiently high (above 1.5 in practice) while below, there is not the necessary information for the model to retrieve it. The ratio between the estimated variances is also shown and we observe an overestimated variance at lower and lower contrast which explains the decreasing  $\max Q/Q_{\text{th}}$  ratio. This effect can partly be explained by the uniform weights hypothesis being less and less true when the contrast decreases.

## V. MULTISCALE PRINCIPAL GRAPHS

Recently proposed methods explore the graph regularization of GMMs to learn a smooth graph representation from point-cloud distributions [22,23]. Based on a Gaussian prior acting on component averages, these approaches rely on a measurement of the graph smoothness through the graph Laplacian as

$$\|\Delta\boldsymbol{\mu}\|^2 = \sum_{i=1}^K \sum_{j=1}^K \mathbf{A}_{ij} \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_2^2, \quad (15)$$

where  $\mathbf{A}$  is the adjacency matrix taking value 1 when centers  $i$  and  $j$  are linked and 0 otherwise. Exploiting this prior knowledge in the GMM leads to the regularized model

$$\ln p(\Theta | \mathbf{x}_i) \propto \ln p(\mathbf{x}_i | \Theta) + \ln p(\Theta), \quad (16)$$

with  $\ln p(\Theta) = -\lambda_\mu \|\Delta\boldsymbol{\mu}\|^2/2$ . This added term on the log-likelihood acts as an attractive quadratic interaction of centers connected on the graph  $\mathbf{A}$ . This prior only impacts the  $M$ -step update of center positions of Eq. (3) as

$$\boldsymbol{\mu}_k^{(t+1)} = \frac{\sum_{i=1}^N \mathbf{x}_i p_{ik}/\sigma^2 + 2\lambda_\mu \sum_{j=1}^K \mathbf{A}_{kj} \boldsymbol{\mu}_j^{(t+1)}}{\sum_{i=1}^N p_{ik}/\sigma^2 + 2\lambda_\mu \sum_{j=1}^K \mathbf{A}_{kj}}. \quad (17)$$

As in Sec. IV, it is possible to compute the value of  $\sigma^2$  for which the high-temperature system becomes unstable, noted  $T_c^{\text{graph}}$  considering perturbations around the fixed point  $\boldsymbol{\mu}_k = \mathbf{0}_D$ . Analogous derivations as in Sec. IV show that the system is unstable when the maximum eigenvalue of  $\mathbf{M}$  is greater than 1, with

$$\mathbf{M} = \left[ \left( \mathbf{I}_K - \frac{1}{K} \mathbf{J}_K \right) \otimes \mathbf{C} \right] \times \left[ \sigma^2 \mathbf{I}_{KD} + \frac{2\lambda_\mu K \sigma^4}{N} \mathbf{L} \otimes \mathbf{I}_D \right]^{-1}, \quad (18)$$

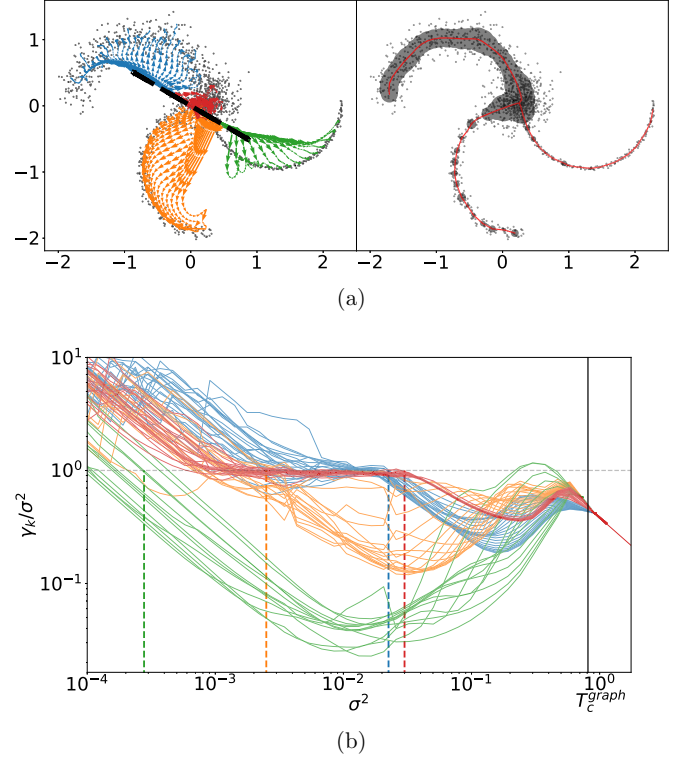


FIG. 5. (a) Left: Displacement of  $K = 100$  components during the hard annealing of a tree branches dataset with different sampling standard deviations. Black dashed line corresponds to the first principal direction. Colors refer to branches in which centers end. Right: Learned structure when stopping the annealing for components reaching the temperature  $\sigma^2 \simeq \gamma_k$ . Red lines are edges of the graph and gray shaded areas are  $1\text{-}\sigma_k$  circles. (b) Evolution of the ratio  $\gamma_k/\sigma^2$  as a function of  $\sigma^2$ . Vertical lines indicate the used variance for the generation of branches. The black vertical line corresponds to the value of  $T_c^{\text{graph}}$ .

where  $\mathbf{L}$  is the Laplacian matrix defined as  $\mathbf{L} = \mathbf{D} - \mathbf{A}$  with  $\mathbf{D}$  the diagonal  $K \times K$  degree matrix  $D_{kk} = \sum_i \mathbf{A}_{ik}$ . Since the prior is basically pulling adjacent centers, the threshold is translated toward lower values as  $\lambda_\mu$  increases.

When paving the data distribution with Gaussian clusters standing on a prior graph structure, the scale of interest is the local width of the elongated structure. This size is given locally, in our two-dimensional (2D) case, by the minimum eigenvalue  $\gamma_k$  of the weighted covariance  $\boldsymbol{\Sigma}_k$ . Figure 5 shows the result of a hard annealing procedure for  $K = 100$  components,  $\lambda_\mu = 300$ , and using a graph prior given by the minimum spanning tree construction [24], hence assuming that centers are linked together with the minimum total length. This topological prior can take different forms depending on the data under consideration and modifies the adjacency matrix  $\mathbf{A}$  and, therefore, the value of  $T_c^{\text{graph}}$ . As predicted by linear stability, centers are first aligned with the principal axis of the dataset at the beginning of the annealing and then spread over the structure to pave it more precisely, as shown in the left panel of Fig. 5(a). By tracking the evolution of the ratio  $\gamma_k/\sigma^2$  in the bottom panel, we clearly distinguish four types of behaviors, the signature of four distinct scales for structures in the dataset. It is also interesting to observe the absence of

sharp phase transitions or splits within this continuous dataset tending to smooth out the evolution of the energy when the temperature is decreasing. By imposing, for each component, the variance  $\sigma^2$  during the annealing at the moment  $\gamma_k \simeq \sigma^2$ , we obtain the graph of the right panel of Fig. 5(a), showing multiple adaptive scales, even though branches have one order of magnitude difference in sampling standard deviation.

## VI. CONCLUSIONS

In this work, we exhibited the cascade of phase transitions occurring when tracking the evolution of eigenvalues of the successive covariance matrices of mixture components during a simulated annealing of the EM algorithm in multiple cases. We saw that hard annealing has nice theoretical advantages with the ability to compute successive theoretical transitions and we showed that a relaxed version can provide an accurate description for datasets exhibiting multiple scales with nested clusters. We also illustrated how hard annealing can be used jointly with graph regularization of GMs to learn a multiscale principal graph independently on the initialization. By revisiting the widely known problem of unsupervised clustering formulated in the framework of statistical physics, the proposed way of exploring the data through simulated annealing does not lead to an automatic and blind solution using all the inputted components. Instead, it builds a hierarchical

description independent from  $K$  providing a qualitative and quantitative insight on the structure of the dataset at different scales and without requiring any prior knowledge. The 2D diagram allows the analysis of the annealing independently from the data dimensionality and highlights characteristic scales at which physical transitions occur. The information on the number of components, their scale, and hierarchy can be used *a posteriori* for data exploration before running blind clustering methods.

Interestingly, we saw that the latent variables of the GMM are directly related to the values taken by the order parameter. Since the GMM can be recast into a restricted Boltzmann machine [25] using a soft-max prior on the hidden nodes, making it a particular type of autoencoder [26], our work can be seen as the learning of a latent representation of the phase transitions, as is discussed in [27,28].

## ACKNOWLEDGMENTS

This research has been supported by the funding for the ByoPiC project from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program Grant No. ERC-2015-AdG 695561. A.D. was supported by the Comunidad de Madrid and the Complutense University of Madrid (Spain) through the Atracción de Talento program (Ref. 2019-T1/TIC-13298).

- 
- [1] K. Rose, *Proc. IEEE* **86**, 2210 (1998).
  - [2] M. Mezard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Inc., New York, 2009).
  - [3] M. Mézard and T. Mora, *J. Physiol. Paris* **103**, 107 (2009).
  - [4] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, *Phys. Rev. E* **84**, 066106 (2011).
  - [5] K. Rose, E. Gurewitz, and G. C. Fox, *Phys. Rev. Lett.* **65**, 945 (1990).
  - [6] S. Akaho and H. J. Kappen, *Neural Comput.* **12**, 1411 (2000).
  - [7] J. Q. Li and A. R. Barron, in *Advances in Neural Information Processing Systems 12*, edited by S. A. Solla, T. K. Leen, and K. Muller (MIT, Cambridge, MA, 2000), pp. 279–285.
  - [8] A. Jain, R. P. W. Duin, and J. Mao, *IEEE Trans. Pattern Anal. Mach. Intell.* **22**, 4 (2000).
  - [9] A. L. Yuille, *Neural Comput.* **2**, 1 (1990).
  - [10] R. Tibshirani, *Stat. Comput.* **2**, 183 (1992).
  - [11] A. P. Dempster, N. M. Laird, and D. B. Rubin, *J. R. Stat. Soc., Series B* **39**, 1 (1977).
  - [12] C. F. J. Wu, *Ann. Stat.* **11**, 95 (1983).
  - [13] M. Kloppenburg and P. Tavan, *Phys. Rev. E* **55**, R2089(R) (1997).
  - [14] N. Ueda and R. Nakano, *Neural Networks* **11**, 271 (1998).
  - [15] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Science* **220**, 671 (1983).
  - [16] H. Akaike, *IEEE Trans. Autom. Control* **19**, 716 (1974).
  - [17] J. J. Oliver, R. A. Baxter, and C. S. Wallace, in *Proceedings of the 13th International Conference on Machine Learning* (Morgan Kaufman, Burlington, MA, 1996), p. 364.
  - [18] G. Schwarz, *Ann. Stat.* **6**, 461 (1978).
  - [19] K. Roeder and L. Wasserman, *J. Am. Stat. Assoc.* **439**, 894 (1997).
  - [20] The separation in three panels with different colors is for visualization purposes. No prior knowledge was used in the analysis.
  - [21] In our experiments, keeping  $\lambda_\sigma \approx O(1)$  did not change the results quantitatively.
  - [22] Q. Mao, L. Wang, I. W. Tsang, and Y. Sun, *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 2227 (2017).
  - [23] T. Bonnaire, N. Aghanim, A. Decelle, and M. Douspis, *Astron. Astrophys.* **637**, A18 (2020).
  - [24] O. Borůvka, *Práce Mor. Přírodověd. Spol.* **3**, 37 (1926).
  - [25] G. E. Hinton, *Neural Comput.* **14**, 1771 (2002).
  - [26] H. Bourlard and Y. Kamp, *Biol. Cybern.* **59**, 291 (1998).
  - [27] E. P. Van Nieuwenburg, Y. H. Liu, and S. D. Huber, *Nat. Phys.* **13**, 435 (2017).
  - [28] S. J. Wetzel, *Phys. Rev. E* **96**, 022140 (2017).