

Signal Recovery Using a Spiked Mixture Model

Paul-Louis Delacour*, Sander Wahls[†], *Senior Member, IEEE*, Jeffrey M. Spraggins^{‡¶||},
Lukasz Migas*, and Raf Van de Plas^{*‡}, *Member, IEEE*

*Delft University of Technology, Delft Center for Systems and Control, Delft, Netherlands

[†]Karlsruhe Institute of Technology, Institute of Industrial Information Technology, Karlsruhe, Germany

[‡]Vanderbilt University, Dept. of Biochemistry & Mass Spectrometry Research Center, Nashville, TN, U.S.A.

[¶]Vanderbilt University, Dept. of Cell and Developmental Biology & Dept. of Chemistry, TN, U.S.A.

^{||}Vanderbilt University Medical Center, Dept. of Pathology, Microbiology, and Immunology, TN, U.S.A.

Corresponding authors: p.l.delacour@tudelft.nl & raf.vandeplas@tudelft.nl

Abstract—We introduce the spiked mixture model (SMM) to address the problem of estimating a set of signals from many randomly scaled and noisy observations. Subsequently, we design a novel expectation-maximization (EM) algorithm to recover all parameters of the SMM. Numerical experiments show that in low signal-to-noise ratio regimes, and for data types where the SMM is relevant, SMM surpasses the more traditional Gaussian mixture model (GMM) in terms of signal recovery performance. The broad relevance of the SMM and its corresponding EM recovery algorithm is demonstrated by applying the technique to different data types. The first case study is a biomedical research application, utilizing an imaging mass spectrometry dataset to explore the molecular content of a rat brain tissue section at micrometer scale. The second case study demonstrates SMM performance in a computer vision application, segmenting a hyperspectral imaging dataset into underlying patterns. While the measurement modalities differ substantially, in both case studies SMM is shown to recover signals that were missed by traditional methods such as k -means clustering and GMM.

Index Terms—machine learning, signal recovery, expectation-maximization, spiked mixture model, Gaussian mixture model, imaging mass spectrometry, hyperspectral imaging.

I. INTRODUCTION

Many advances in sensor technology and instrumentation are driven by the demand for greater specificity, sensitivity, and resolution, often implicitly leading to the acquisition of ever larger amounts of increasingly high-dimensional data. This trend can be observed across a broad range of measurement modalities and application domains, including super-resolution imaging [1], novel sensor types for computer vision and remote sensing [2], chemical assays [3], communication [4], and military applications [5]. In cases where limited signal strength or energy needs to be spread over a growing number of observations and dimensions, it can become increasingly difficult to differentiate between signals and to discern them from background variation. For example, in certain imaging techniques, sampling at smaller spatial distances can degrade the signal-to-noise ratio (SNR) of pixel-specific measurements, and adjusting the spectral resolution of a spectrometer can impact its limit-of-detection (LOD). Therefore, the ability to effectively and reliably recover signals from increasingly noisy measurements is becoming essential to unlocking the full potential of certain measurement modalities, particularly in scenarios with substantial background noise, low SNR, or high

LOD (e.g., single-cell measurements with limited molecular material to measure).

Signal recovery in noisy environments without prior knowledge of the subpopulations within the observation pool is often conducted using a Gaussian mixture model (GMM) and by maximum likelihood estimation (MLE) of the model's parameters through expectation-maximization (EM) optimization. For example, motivated by multi-reference alignment in cryo-electron microscopy (a.k.a. the orbit retrieval problem), Katsevich and Bandeira [6] studied likelihood maximization for the GMM in low SNR regimes. The use of a standard GMM implicitly assumes that the signal subpopulations or mixture components are normally distributed. Although this is a broad assumption that fits many scenarios and contributes to the popularity of this approach, certain measurement types allow for more refined assumptions on the underlying mixture components. For those measurement types, using a GMM will lead to suboptimal signal recovery, especially at low SNR.

To address the mismatch between one such signal type and the GMM, we introduce an alternative mixture model, called the spiked mixture model (SMM). In the SMM, an observation or measurement is a randomly scaled version of one of a set of underlying signals called spikes, further perturbed by additive noise. Spiked models, introduced by Johnstone [7], are a class of models characterized by the insertion of a planted vector into a random matrix. While previous work explored the statistical properties of these models under various prior distributions on the spike [8], our SMM takes a distinct approach. One can view the SMM's covariance matrix as a sum of spiked Wishart matrices with one degree of freedom and each spike sampled from a categorical prior. This differs from the multi-spiked tensor model in [9] as we focus on a mixture model rather than on linear combinations of spikes.

Although the SMM's signal spikes could potentially be recovered using a GMM, we show that GMM-based recovery is only practical in high-SNR scenarios. In contrast, an EM-based approach that directly estimates an SMM consistently outperforms the GMM, especially in noisy conditions. While one could argue that refining the Gaussian distribution assumption makes the SMM less broadly applicable than the GMM, the SMM remains widely relevant to a variety of application domains. In those domains, the SMM tends to fit

the underlying signal model better, enabling advanced signal detection and recovery. This becomes especially valuable when addressing high-noise, low-SNR measurements. Beyond its primary function of recovering signals from noisy measurements, an EM-driven approach for SMM-fitting provides additional information that classical methods, such as k -means clustering (k MC), do not. These secondary outputs from the fitting process include an implicit estimate of the observations' noise variance, mixture probabilities for each spike, and spike responsibilities, *i.e.*, the probability that a noisy observation belongs to a specific spike.

The desirable asymptotic properties of MLE make it a common choice to drive the fitting process. As the number of observations goes to infinity, MLE is an asymptotically consistent and efficient estimator, *i.e.*, it converges to the true parameter values and achieves the lowest possible variance among unbiased estimators ([10], Chap. 10). However, computing the MLE remains a challenge in many scenarios. A typical algorithmic procedure to find a MLE candidate in the presence of unobserved latent variables is EM [11] [12] [13]. EM is a sequential algorithm performing 'soft assignment' of observations to mixture components, with guarantees to converge to a local maximum. Since deriving the equations for an EM optimization is model-specific, Section II is dedicated to developing a custom EM algorithm for the SMM.

In Section III, we provide a direct comparison between standard GMM-based signal recovery and the proposed SMM. Using a synthetic dataset with known ground truth signal populations (spikes), we demonstrate that in low-noise scenarios SMM and GMM achieve equivalent recovery. However, in high-noise regimes, SMM substantially outperforms GMM.

Section IV demonstrates the SMM in real-world applications. First, we use the SMM EM-algorithm to recover underlying molecular signatures from noisy imaging mass spectrometry (IMS) measurements of a rat brain tissue section. The recovered spikes align with known biological structures, the estimated responsibilities segment the tissue according to molecular content, and SMM retrieves histological patterns missed by GMM. This case illustrates SMM's enhanced signal accuracy when recovering in the presence of sizeable noise. The second application uses SMM to segment hyperspectral imaging (HSI) measurements. This case study demonstrates SMM's ability to differentiate signal subpopulations that GMM and k MC have difficulty with. Notably, the SMM is not intrinsically related to imaging data. Its signal model also finds relevance in domains such as wireless communication, where the SMM describes a random access narrowband flat fading single-input multiple-output communication system ([14], Sec. 7.3). Imaging examples allow SMM estimation results (*e.g.*, responsibilities) to be shown as images, aiding interpretation.

II. EXPECTATION-MAXIMIZATION ALGORITHM FOR THE SPIKED MIXTURE MODEL

We study the problem of estimating a set of signals from observations that consist of randomly scaled and noisy copies of those signals. More precisely, we consider N independent




Fig. 1. Exemplary realization of a spiked mixture model as specified in (1) that underlies N observations, with $N = 6$, $d = 2$, and $K = 3$. Black lines represent the directions, or spikes, along which observations concentrate, blue dots are scaled observations along a particular spike without noise, and red circles represent one standard deviation of the Gaussian noise perturbations.

observations $\mathbf{y}_1, \dots, \mathbf{y}_N \in \mathbb{R}^d$, sampled from the model

$$\mathbf{y} = \begin{cases} \alpha \mathbf{x}_1 + \varepsilon & \text{with probability } \pi_1 \\ \vdots \\ \alpha \mathbf{x}_K + \varepsilon & \text{with probability } \pi_K \end{cases}, \quad (1)$$

$$\alpha \sim \mathcal{N}(0, 1), \quad \varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}),$$

$$\sum_{k=1}^K \pi_k = 1, \quad \mathbf{x}_1, \dots, \mathbf{x}_K \in \mathbb{R}^d,$$

where α is the random scaling factor of observation \mathbf{y} , \mathbf{x}_k is the k -th subpopulation or spike, ε is the random noise of observation \mathbf{y} , and π_k is the probability of the k -th subpopulation. Let $z \in \{1, \dots, K\}$ be a latent categorical variable indicating which of the spikes $\mathbf{x}_1, \dots, \mathbf{x}_K$ was used to generate \mathbf{y} . Given z , the random vector \mathbf{y} is normally distributed as it is a sum of independent normally distributed variables. Specifically, by computing $\mathbb{E}[\alpha \mathbf{x}_z + \varepsilon]$ and $\mathbb{E}[(\alpha \mathbf{x}_z + \varepsilon)(\alpha \mathbf{x}_z + \varepsilon)^T]$, we find that $\mathbf{y}|z \sim \mathcal{N}(\mathbf{0}, \Sigma_z := \mathbf{x}_z \mathbf{x}_z^T + \sigma^2 \mathbf{I})$. Model (1) is thus a constrained Gaussian mixture model (GMM) with density

$$p(\mathbf{y}) = \sum_{k=1}^K \pi_k p_{\mathcal{N}(\mathbf{0}, \Sigma_k)}(\mathbf{y}), \quad (2)$$

where $p_{\mathcal{N}(\mathbf{0}, \Sigma_z)}$ denotes the probability density function of a multivariate Gaussian distribution with mean $\mathbf{0}$ and covariance Σ_z . Instead of recovering the mean and covariance matrix of each Gaussian component, our goal here is to estimate the vector of parameters $\boldsymbol{\theta} = \{\mathbf{x}_1, \dots, \mathbf{x}_K, \pi_1, \dots, \pi_K, \sigma^2\}$ that defines them. We refer to this model as a *spiked mixture model*.

- The adjective *spiked* refers to observations concentrating along certain directions, called spikes, in the measurement space (Figure 1 shows an exemplary realization).
- *Mixture* refers to there being a set of K directions, spikes, or signal subpopulations within the observations.

As $\mathbf{y}|z$ is normally distributed, using Sylvester's determinant identity and the Sherman-Morrison formula to compute the

determinant and the inverse of the covariance matrix $\Sigma_z := \mathbf{x}_z \mathbf{x}_z^T + \sigma^2 \mathbf{I}$, we find the conditional density

$$p_{\theta}(\mathbf{y}|z) = \frac{1}{\sqrt{(2\pi)^d}} \exp \left[\frac{-1}{2\sigma^2} \left(\|\mathbf{y}\|^2 - \frac{(\mathbf{y}^T \mathbf{x}_z)^2}{\|\mathbf{x}_z\|^2 + \sigma^2} \right) - \frac{1}{2} \ln(\|\mathbf{x}_z\|^2 + \sigma^2) - (d-1) \ln \sigma \right], \quad (3)$$

where the subscript θ indicates that a density is determined by the parameter vector. With N observations $\mathbf{y}_1, \dots, \mathbf{y}_N$ from (1), the log-likelihood corresponds to

$$\ln p_{\theta}(\mathbf{y}_1, \dots, \mathbf{y}_N) = \sum_{i=1}^N \ln \sum_{k=1}^K \pi_k \exp \left[\frac{1}{2\sigma^2} \frac{(\mathbf{y}_i^T \mathbf{x}_k)^2}{\|\mathbf{x}_k\|^2 + \sigma^2} - \frac{1}{2} \ln(\|\mathbf{x}_k\|^2 + \sigma^2) - (d-1) \ln \sigma \right] + C,$$

where we use $p_{\theta}(\mathbf{y}_1, \dots, \mathbf{y}_N) = \prod_{i=1}^N p_{\theta}(\mathbf{y}_i)$ and exploit (2) and (3). The constant C is independent of θ . Note that one can only hope to recover the vectors $\mathbf{x}_1, \dots, \mathbf{x}_K$ in the model (1) up to a sign change since it suffers from an intrinsic symmetry. Technically, it is thus non-identifiable [15]. Indeed, for two sets of parameters, θ_1 and θ_2 , that are the same up to a sign change in front of the \mathbf{x}_k s we have:

$$p_{\theta_1}(\mathbf{y}_1, \dots, \mathbf{y}_N) = p_{\theta_2}(\mathbf{y}_1, \dots, \mathbf{y}_N).$$

The expectation-maximization (EM) algorithm [16], originally developed by [Dempster et al. in 1977] is a common approach to finding a candidate maximum likelihood estimate of θ . It is an iterative procedure that is made up of two steps. Let $\theta^{[t]} = \{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_K, \hat{\pi}_1, \dots, \hat{\pi}_K, \hat{\sigma}^2\}$ be the estimate of θ at step t . The first step, called the expectation or E -step, computes

$$\rho_z^{[i]} := p_{\theta^{[t]}}(z|\mathbf{y}_i) \quad \forall i \in \{1, \dots, N\}, z \in \{1, \dots, K\}.$$

Using Bayes rule, we have

$$p_{\theta^{[t]}}(z|\mathbf{y}_i) = \frac{p_{\theta^{[t]}}(\mathbf{y}_i|z)\pi_z}{\sum_{k=1}^K p_{\theta^{[t]}}(\mathbf{y}_i|z=k)\pi_k}.$$

Using (3) and simplifying the terms that cancel in the numerator and denominator, we get

$$\rho_z^{[i]} := p_{\theta^{[t]}}(z|\mathbf{y}_i) = \frac{\tilde{\rho}_z^{[i]}}{\sum_{k=1}^K \tilde{\rho}_k^{[i]}}, \quad \text{where}$$

$$\tilde{\rho}_z^{[i]} := \hat{\pi}_z \exp \left(\frac{1}{2\hat{\sigma}^2} \frac{(\mathbf{y}_i^T \hat{\mathbf{x}}_z)^2}{\|\hat{\mathbf{x}}_z\|^2 + \hat{\sigma}^2} - \frac{1}{2} \ln(\|\hat{\mathbf{x}}_z\|^2 + \hat{\sigma}^2) \right).$$

The second step, called the maximization or M -step, computes $\theta^{[t+1]}$ by finding the feasible θ maximizing

$$\begin{aligned} & \mathcal{Q}(\theta; \theta^{[t]}) \\ &= \sum_{i=1}^N \mathbb{E}_{z \sim p_{\theta^{[t]}}(\cdot|\mathbf{y}_i)} [\ln p_{\theta}(\mathbf{y}_i, z)] \\ &= \sum_{i=1}^N \sum_{k=1}^K \rho_k^{[i]} \left[\ln(\pi_k) - \frac{1}{2\sigma^2} \left(\|\mathbf{y}_i\|^2 - \frac{(\mathbf{y}_i^T \mathbf{x}_k)^2}{\|\mathbf{x}_k\|^2 + \sigma^2} \right) - \frac{1}{2} \ln(\|\mathbf{x}_k\|^2 + \sigma^2) - \frac{d}{2} \ln(2\pi) - (d-1) \ln(\sigma) \right]. \end{aligned} \quad (4)$$

Maximizing the expectation function \mathcal{Q} acts as a proxy for maximizing the log-likelihood and guarantees the following improvement (see (3.10) in [16]):

$$\begin{aligned} & \ln p_{\theta}(y_1, \dots, y_N) - \ln p_{\theta^{[t]}}(y_1, \dots, y_N) \\ & \geq \mathcal{Q}(\theta; \theta^{[t]}) - \mathcal{Q}(\theta^{[t]}; \theta^{[t]}). \end{aligned}$$

In other words, at every step, the log-likelihood improves by at least as much as the \mathcal{Q} function. Precisely, the M -step is

$$\theta^{[t+1]} = \arg \max_{\theta \text{ s.t. } \begin{array}{l} \mathbf{x}_k \in \mathbb{R}^d \\ \pi_k \geq 0 \\ \sum_{k=1}^K \pi_k = 1 \\ \sigma \in \mathbb{R}_+ \end{array}} \mathcal{Q}(\theta; \theta^{[t]}) \quad (5)$$

To solve this, we instead look at the less constrained problem

$$\theta^{[t+1]} = \arg \max_{\theta \text{ s.t. } \sum_{k=1}^K \pi_k = 1} \mathcal{Q}(\theta; \theta^{[t]}) \quad (6)$$

and verify that a found $\theta^{[t+1]}$ satisfies the missing constraints. The solution of this less constrained M -step is, by the Karush-Kuhn-Tucker conditions, a critical point of the Lagrangian

$$\mathcal{L} = \mathcal{Q}(\theta; \theta^{[t]}) + \lambda \left(1 - \sum_{k=1}^K \pi_k \right).$$

In what follows, we find a critical point of this Lagrangian:

$$\left(\frac{\partial \mathcal{L}}{\partial \mathbf{x}_1}, \dots, \frac{\partial \mathcal{L}}{\partial \mathbf{x}_K}, \frac{\partial \mathcal{L}}{\partial \pi_1}, \dots, \frac{\partial \mathcal{L}}{\partial \pi_K}, \frac{\partial \mathcal{L}}{\partial \sigma}, \frac{\partial \mathcal{L}}{\partial \lambda} \right) = \mathbf{0}. \quad (7)$$

In order to simplify the notation for $\mathcal{Q}(\theta; \theta^{[t]})$, and thus \mathcal{L} , we define the following for all $e \in \{1, \dots, K\}$:

$$\mathbf{A}_e := \sum_{i=1}^N \rho_e^{[i]} \mathbf{y}_i \mathbf{y}_i^T \in \mathbb{R}^{d \times d}, \quad (8)$$

$$\gamma_e := \sum_{i=1}^N \rho_e^{[i]} \in \mathbb{R}, \quad (9)$$

$$\mathbf{Y} := \begin{bmatrix} | & & | \\ \mathbf{y}_1 & \dots & \mathbf{y}_N \\ | & & | \end{bmatrix} \in \mathbb{R}^{d \times N}. \quad (10)$$

With this and the fact that $\sum_{k=1}^K \rho_k^{[i]} = 1$ for all i , we can rewrite $\mathcal{Q}(\theta; \theta^{[t]})$ as

$$\begin{aligned} \mathcal{Q}(\theta; \theta^{[t]}) &= \sum_{k=1}^K \gamma_k \ln \pi_k - \frac{\|\mathbf{Y}\|_F^2}{2\sigma^2} + \frac{1}{2\sigma^2} \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} \\ &\quad - \frac{1}{2} \sum_{k=1}^K \gamma_k \ln(\|\mathbf{x}_k\|^2 + \sigma^2) - (d-1)N \ln(\sigma) \\ &\quad - \frac{dN}{2} \ln(2\pi). \end{aligned} \quad (11)$$

a) *Derivatives with respect to π_e for $e \in \{1, \dots, K\}$*

The condition in (7) is reached when

$$\frac{\partial \mathcal{L}}{\partial \pi_e} = \frac{\gamma_e}{\pi_e} - \lambda = 0,$$

or, equivalently, when

$$\lambda = \frac{1}{\pi_e} \gamma_e \quad \forall e \in \{1, \dots, K\}. \quad (12)$$

Using (12) and the constraint $\sum_{k=1}^K \pi_k = 1$, which follows from the condition $\frac{d\mathcal{L}}{d\lambda} = 0$ in (7), we find

$$1 = \sum_{k=1}^K \pi_k = \sum_{k=1}^K \frac{\gamma_k}{\lambda} = \frac{N}{\lambda},$$

which gives $\lambda = N$, and

$$\pi_e = \frac{\gamma_e}{N} \quad \forall e \in \{1, \dots, K\}. \quad (13)$$

We note here that π_e from (13) is indeed in the range $[0, 1]$.

b) *Derivatives with respect to \mathbf{x}_e for $e \in \{1, \dots, K\}$*

We now expand the first condition in (7). The derivative of the Lagrangian w.r.t. \mathbf{x}_e is

$$\frac{\partial \mathcal{L}}{\partial \mathbf{x}_e} = \frac{1}{2\sigma^2} \left[\frac{2(\|\mathbf{x}_e\|^2 + \sigma^2) \mathbf{A}_e \mathbf{x}_e - 2(\mathbf{x}_e^T \mathbf{A}_e \mathbf{x}_e) \mathbf{x}_e}{(\|\mathbf{x}_e\|^2 + \sigma^2)^2} \right] - \frac{\gamma_e \mathbf{x}_e}{\|\mathbf{x}_e\|^2 + \sigma^2}.$$

Setting it to zero and rearranging the terms, we get

$$\mathbf{A}_e \mathbf{x}_e = \left(\frac{\mathbf{x}_e^T \mathbf{A}_e \mathbf{x}_e}{\|\mathbf{x}_e\|^2 + \sigma^2} + \sigma^2 \gamma_e \right) \mathbf{x}_e. \quad (14)$$

We distinguish two possible cases:

- either \mathbf{x}_e is the zero vector, and this is trivially satisfied,
- or \mathbf{x}_e is different from the zero vector.

In the latter case, (14) means that \mathbf{x}_e is an eigenvector of \mathbf{A}_e with eigenvalue λ_e :

$$\mathbf{A}_e \mathbf{x}_e =: \lambda_e \mathbf{x}_e. \quad (15)$$

Projecting (14) onto $\mathbf{x}_e / \|\mathbf{x}_e\|^2$, we find

$$\frac{\mathbf{x}_e^T \mathbf{A}_e \mathbf{x}_e}{\|\mathbf{x}_e\|^2} = \frac{\mathbf{x}_e^T \mathbf{A}_e \mathbf{x}_e}{\|\mathbf{x}_e\|^2 + \sigma^2} + \sigma^2 \gamma_e.$$

Substituting (15) into this equation, we get

$$\lambda_e = \lambda_e \frac{\|\mathbf{x}_e\|^2}{\|\mathbf{x}_e\|^2 + \sigma^2} + \sigma^2 \gamma_e,$$

which simplifies to

$$\lambda_e = (\|\mathbf{x}_e\|^2 + \sigma^2) \gamma_e \quad \forall e \in \{1, \dots, K\}.$$

Thus, when $\mathbf{x}_e \neq \mathbf{0}$, the critical point in (7) satisfies

$$\|\mathbf{x}_e\|^2 = \frac{\lambda_e}{\gamma_e} - \sigma^2 \quad \forall e \in \{1, \dots, K\}. \quad (16)$$

c) *Derivative with respect to σ*

Finally, we expand the third condition in (7). The corresponding derivative is

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \sigma} &= \frac{1}{\sigma^3} \|\mathbf{Y}\|_F^2 - \frac{1}{\sigma^3} \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} - \frac{1}{\sigma} \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{(\|\mathbf{x}_k\|^2 + \sigma^2)^2} \\ &\quad - \sigma \sum_{k=1}^K \frac{\gamma_k}{\|\mathbf{x}_k\|^2 + \sigma^2} - \frac{(d-1)N}{\sigma}. \end{aligned}$$

Setting the prior equation to zero and multiplying by σ^3 , yields

$$\begin{aligned} &\sigma^2(d-1)N \\ &= \|\mathbf{Y}\|_F^2 - \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} \\ &\quad - \sigma^2 \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{(\|\mathbf{x}_k\|^2 + \sigma^2)^2} - \sigma^4 \sum_{k=1}^K \frac{\gamma_k}{\|\mathbf{x}_k\|^2 + \sigma^2} \\ &= \|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} - \sigma^2 \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k \\ &\quad - \sigma^2 \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \frac{1}{\|\mathbf{x}_k\|^2 + \sigma^2} \left(\frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} + \sigma^2 \gamma_k \right). \end{aligned}$$

By rearranging the σ^2 terms, replacing γ_k in the last term with (16), and using the eigenrelation (15) twice, we get

$$\begin{aligned} &\sigma^2 \left[(d-1)N + \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k \right] \\ &= \|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \left(\frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} + \sigma^2 \frac{\lambda_k}{\|\mathbf{x}_k\|^2 + \sigma^2} \right) \\ &= \|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \left(\lambda_k \frac{\|\mathbf{x}_k\|^2 + \sigma^2}{\|\mathbf{x}_k\|^2 + \sigma^2} \right) \quad \text{by using (15)}. \end{aligned}$$

This gives a final expression for σ^2 :

$$\sigma^2 = \frac{\|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \lambda_k}{(d-1)N + \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k}.$$

Since $\sum_{k=1}^K \gamma_k = N$ implies $\sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k = N - \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k$, we can rewrite the expression for σ^2 as

$$\sigma^2 = \frac{\|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \lambda_k}{dN - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k}. \quad (17)$$

In Lemma A.1 of the Appendix, we show that estimate (17) of σ^2 is indeed positive.

In sum, the equations for critical points, $\forall k \in \{1, \dots, K\}$, are:

$$\pi_k = \frac{\gamma_k}{N} \quad \text{from (13)}$$

$$\begin{cases} \text{either } \mathbf{x}_k = \mathbf{0} \\ \text{or } \mathbf{A}_k \mathbf{x}_k = \lambda_k \mathbf{x}_k, \quad \|\mathbf{x}_k\|^2 = \frac{\lambda_k}{\gamma_k} - \sigma^2 \end{cases} \quad \text{from (15), (16)}$$

$$S = \{k \in [K] \text{ s.t. } \mathbf{x}_k \neq \mathbf{0}\}$$

$$\sigma^2 = \sigma^2(S) := \frac{\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k}{dN - \sum_{k \in S} \gamma_k} \quad \text{from (17)}$$

Now that we have established the equations that characterize critical points, we will analyze which critical point maximizes $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{[t]})$. After a series of algebraic reformulations provided in Appendix A-A we find that for a $\boldsymbol{\theta}$ satisfying the equations of a critical point, the expression of \mathcal{Q} simplifies to

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]}) &= \hat{C} - \frac{1}{2} \left[dN - \sum_{k \in S} \gamma_k \right] \ln \left(\frac{\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k}{dN - \sum_{k \in S} \gamma_k} \right) \\ &\quad - \frac{1}{2} \sum_{k \in S} \gamma_k \ln \left(\frac{\lambda_k}{\gamma_k} \right), \end{aligned}$$

with \hat{C} a constant independent of θ . Taking the derivative of the previous expression with respect to λ_k for $k \in S$ we see

$$\frac{\partial \mathcal{Q}(\theta; \theta^{[t]})}{\partial \lambda_k} \geq 0 \iff \sigma^2(S) \leq \frac{\lambda_k}{\gamma_k}. \quad (18)$$

For a $\lambda_j(\mathbf{A}_k)$ satisfying (18), we know that any $\lambda_k > \lambda_j(\mathbf{A}_k)$ also satisfies (18) since it decreases $\sigma^2(S)$ and it increases $\frac{\lambda_k}{\gamma_k}$. From the equivalency in (18), this means that $\mathcal{Q}(\theta; \theta^{[t]})$ is increasing on the set $[\lambda_j(\mathbf{A}_k), \lambda_1(\mathbf{A}_k)]$. As a result, picking $\lambda_k = \lambda_1(\mathbf{A}_k)$ for all k maximizes \mathcal{Q} .

We are now left with picking a set S that maximizes the function $\mathcal{Q}(\theta; \theta^{[t]})$ or, equivalently, that minimizes the function

$$g(S) = \left[dN - \sum_{k \in S} \gamma_k \right] \ln \sigma^2(S) + \sum_{k \in S} \gamma_k \ln \left(\frac{\lambda_k}{\gamma_k} \right).$$

We note that not all sets S satisfy the critical point equations. Indeed, for $k \in S$, the associated eigenvalue λ_k must satisfy $\frac{\lambda_k}{\gamma_k} - \sigma^2(S) \geq 0$. Therefore, we call $S \subseteq [K]$ a valid set when

$$\sigma^2(S) \leq \frac{\lambda_k}{\gamma_k} \quad \forall k \in S.$$

We will denote by \mathcal{V} the set of all valid sets:

$$\mathcal{V} = \{S \subseteq [K] \mid \forall k \in S : \sigma^2(S) \leq \lambda_k/\gamma_k\}.$$

Note that \mathcal{V} is never empty since it contains at least the empty set for which the condition is trivially satisfied. The optimal set S^* is given by

$$S^* = \underset{S \in \mathcal{V}}{\operatorname{argmin}} g(S).$$

A naive procedure trying all possible valid sets would take $\mathcal{O}(K2^K)$. In Appendix B we provide a $\mathcal{O}(K^2)$ procedure.

Algorithm I summarizes all operations performed by the EM optimization to fit a K -spiked mixture model. Standard convergence results for the EM algorithm can be applied to our algorithm. For details, see section G of the supplementary material.

III. COMPARISON WITH STANDARD GMM

We now compare our SMM method to a GMM-based approach. Since our model is a constrained Gaussian mixture model, we can use a standard GMM to estimate the covariance matrices, yielding $\hat{\Sigma}_1, \dots, \hat{\Sigma}_K$. If this recovery is successful, we expect the following relationships to hold (albeit up to a possible relabeling of the \mathbf{x}_k s):

$$\hat{\Sigma}_k = \mathbf{x}_k \mathbf{x}_k^T + \sigma^2 I \quad \forall k \in [K].$$

Once the covariances are estimated, we can extract the \mathbf{x}_k s by solving the following optimization problem:

$$\min_{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_K, \hat{\sigma}^2} \sum_{k=1}^K \|\hat{\Sigma}_k - (\hat{\mathbf{x}}_k \hat{\mathbf{x}}_k^T + \hat{\sigma}^2 I)\|_F^2. \quad (19)$$

Let $\lambda_i(\Sigma_k)$ denote the i -th largest eigenvalue of Σ_k and $\mathbf{v}_i(\Sigma_k)$ its associated eigenvector. Lemma C.2 in appendix C shows that whenever the condition

$$\lambda_1(\hat{\Sigma}_j) \geq \frac{\sum_{k=1}^K \sum_{i=2}^d \lambda_i(\hat{\Sigma}_k)}{K(d-1)} \quad \forall j \in [K] \quad (20)$$

Algorithm 1 EM fitting of a K -spiked mixture model

INITIALIZATION :

$$\mathbf{x}_1, \dots, \mathbf{x}_K, \pi_1, \dots, \pi_K, \sigma^2$$

REPEAT UNTIL CONVERGENCE:

E-Step:

$$\begin{aligned} \text{for } i = 1, \dots, N, e = 1, \dots, K \\ \tilde{\rho}_e^{[i]} &= \pi_e \cdot \exp \left[\frac{1}{2\sigma^2} \frac{(\mathbf{y}_i^T \mathbf{x}_e)^2}{\|\mathbf{x}_e\|^2 + \sigma^2} - \frac{1}{2} \ln(\|\mathbf{x}_e\|^2 + \sigma^2) \right] \\ \rho_e^{[i]} &= \frac{\tilde{\rho}_e^{[i]}}{\sum_{k=1}^K \tilde{\rho}_k^{[i]}} \end{aligned}$$

M-Step:

$$\begin{aligned} \text{for } e = 1, \dots, K : \\ \gamma_e &= \sum_{i=1}^N \rho_e^{[i]} \\ \pi_e &= \gamma_e / N \\ \mathbf{A}_e &= \sum_{i=1}^N \rho_e^{[i]} \mathbf{y}_i \mathbf{y}_i^T \\ \lambda_e, \mathbf{v}_e &\text{ largest eigenvalue/vector pair of } \mathbf{A}_e \end{aligned}$$

Find optimal S:

$$\begin{aligned} \sigma^2(S) &= [\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k] / (dN - \sum_{k \in S} \gamma_k) \\ \mathcal{V} &= \{S \subseteq [K] \text{ s.t. } \forall k \in S : \sigma^2(S) \leq \lambda_k/\gamma_k\} \\ g(S) &= [dN - \sum_{k \in S} \gamma_k] \ln \sigma^2(S) + \sum_{k \in S} \gamma_k \ln \frac{\lambda_k}{\gamma_k} \end{aligned}$$

$$\begin{aligned} S^* &= \operatorname{argmin}_{S \in \mathcal{V}} g(S) \\ \sigma^2 &= \sigma^2(S^*) \end{aligned}$$

for $e = 1, \dots, K$:

$$\begin{aligned} \text{if } e \in S^* : \\ \mathbf{x}_e &= \sqrt{\frac{\lambda_e}{\gamma_e} - \sigma^2} \cdot \mathbf{v}_e \\ \text{else:} \\ \mathbf{x}_e &= 0 \end{aligned}$$

holds, the solution to (19) is given by

$$\sigma_{\text{GMM}}^2 = \frac{1}{K} \sum_{k=1}^K \frac{\operatorname{tr}(\hat{\Sigma}_k) - \lambda_1(\hat{\Sigma}_k)}{d-1}, \quad (21)$$

$$\mathbf{x}_{\text{GMM},k} = \sqrt{(\lambda_1(\hat{\Sigma}_k) - \sigma_{\text{GMM}}^2)} \mathbf{v}_1(\hat{\Sigma}_k) \quad \forall k \in [K]. \quad (22)$$

Note that the requirement in (20) can be thought of as a large enough spectral gap for the leading eigenvalues of Σ_k , and it has been checked to hold in all the following experiments of this section.

In what follows, we compare the performance of SMM and GMM-based recovery on a synthetic dataset. Specifically, we choose $K = 3$, and sample fixed (starred) parameters:

$$\begin{cases} \theta^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_K^*, \pi_1^*, \dots, \pi_K^*, \sigma^{2*}\}, \\ \pi_k^* \in [0, 1], \sum_{k=1}^K \pi_k^* = 1 \\ \sigma^{2*} \in \mathbb{R}_+. \end{cases} \quad (23)$$

Given these parameters, we generate $N = 1500$ samples, $\mathbf{y}_1, \dots, \mathbf{y}_N$, according to the model in (I), and we do this for two different noise levels: $\sigma^2 = 0.01$ and $\sigma^2 = 0.5$. Our goal is

¹We remark in passing that this setup models, *e.g.*, a random access SIMO channel, where various sensors with one antenna sporadically transmit a symbol to a base station with many antennas at random times. The algorithm blindly estimates both the channel vectors \mathbf{x}_k and the symbols α .

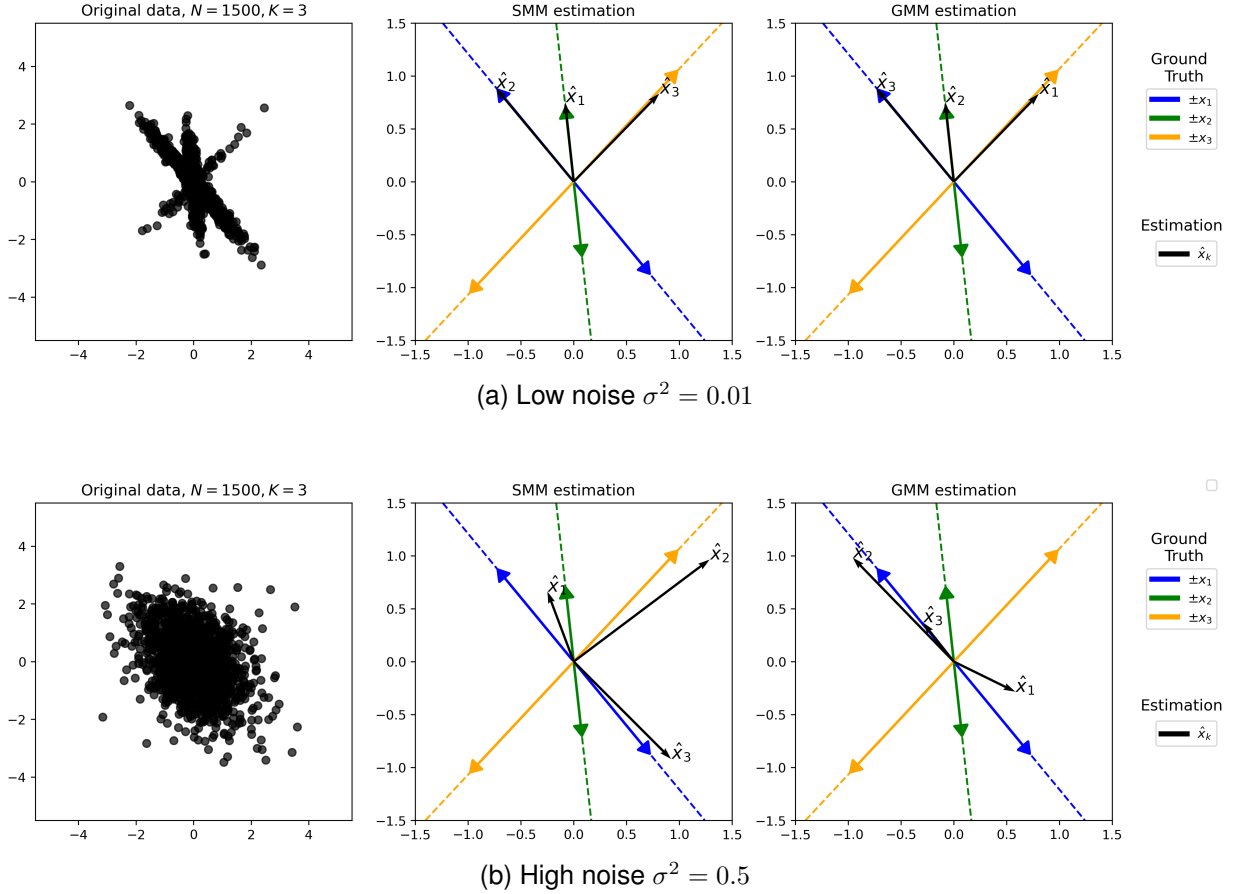


Fig. 2. Comparison between recovery by GMM versus SMM. The comparison is made at both low noise ($\sigma^2 = 0.01$) (2a) and high noise ($\sigma^2 = 0.5$) (2b). The fixed parameters are: $\mathbf{x}_1 \approx [0.75, -0.91]$, $\mathbf{x}_2 \approx [0.08, -0.75]$, $\mathbf{x}_3 \approx [-1.01, -1.08]$, $\pi_1 \approx 0.58$, $\pi_2 \approx 0.37$, $\pi_3 \approx 0.05$.

to assess how the noise level affects the accuracy of recovering the true signal vectors \mathbf{x}_1 , \mathbf{x}_2 , and \mathbf{x}_3 using SMM and GMM. We emphasize once more that, due to symmetry, the recovery of \mathbf{x}_k is just as likely as the recovery of $-\mathbf{x}_k$. The results are shown in Figure 2. While in the low noise case (Figure 2a), the performance of SMM and GMM appear to be similar, we see a clear difference for the high noise case (Figure 2b). In this regime, the accuracy of both methods decreases, but GMM only finds one vector, \mathbf{x}_1 , and produces a third estimate, which is a mixture of two ground truth signals. On the other hand, SMM delivers three clearly separate directions without the 'collapse' of estimates we see in the GMM case.

To quantify SMM's performance versus GMM's, we conducted two experiments. One compares the methods' noise estimation performance, and the other compares the distance of the estimated signals, $\hat{\mathcal{X}} = \{\hat{\mathbf{x}}_k\}_{k=1}^K$, to the true ones, $\mathcal{X} = \{\mathbf{x}_k\}_{k=1}^K$. The supplementary material, section F, furthermore contains an empirical experiment that suggests that the convergence speed of both algorithms is similar.

A. Noise estimation

We conducted signal recovery using both methods on synthetic datasets with a known ground truth noise variance σ^2 , repeating the process across 10 different noise levels in the range $[1, 30]$. Figure 3 shows the results of SMM's and GMM's noise variance estimation compared to the ground truth values.

While both methods tend to underestimate the noise variance, the bias is significantly larger for GMM than for SMM. Additionally, the norms of the estimated $\hat{\mathbf{x}}_k$ decrease as σ^2 increases (see (16), (22)), leading to overestimated norms for both methods. This overestimation is notably more pronounced in the GMM case, which further exacerbates its bias.

B. Signal estimation

To assess the fidelity of recovered signals, we performed estimation on a synthetic dataset with parameters $N = 1500$, $K = 3$, $d = 5$, and $\sigma^2 = 1.5$. The ground truth signals \mathbf{x} were compared to their estimated counterparts $\hat{\mathbf{x}}$ using the squared Euclidean and absolute cosine error distance metrics:

$$d_{\text{sqe}}(\mathbf{x}, \hat{\mathbf{x}}) := \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2,$$

$$d_{\text{abs_cos}}(\mathbf{x}, \hat{\mathbf{x}}) := 1 - \frac{|\mathbf{x}^T \hat{\mathbf{x}}|}{\|\mathbf{x}\|_2 \|\hat{\mathbf{x}}\|_2}.$$

The squared Euclidean distance, $d_{\text{sqe}}(\mathbf{x}, \hat{\mathbf{x}})$, depends on the norms of \mathbf{x} , $\hat{\mathbf{x}}$, while $d_{\text{abs_cos}}(\mathbf{x}, \hat{\mathbf{x}})$ does not. Since Section III-A highlighted a bias in the estimated norms, $d_{\text{abs_cos}}(\mathbf{x}, \hat{\mathbf{x}})$ was included for its invariance to such bias. To quantify the overall discrepancy between the set of estimated signals $\hat{\mathcal{X}}$ and the set of true signals \mathcal{X} , we used the Hausdorff distance. For a distance metric $d(\cdot)$, the Hausdorff distance

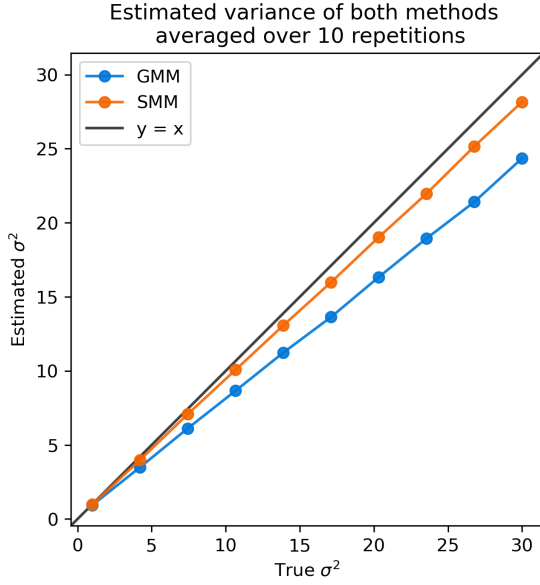


Fig. 3. Comparison of SMM's and GMM's noise variance estimation versus the ground truth value, at 10 different noise levels equally spaced in $[1, 30]$. At each noise level, 10 replicate datasets with different underlying signals are generated using $N = 1500$ samples and the parameters $d = 5$ and $K = 3$.

between sets \mathcal{X} and $\hat{\mathcal{X}}$ is defined as:

$$d_H(\mathcal{X}, \hat{\mathcal{X}}) := \max \left(\sup_{\mathbf{x} \in \mathcal{X}} d(\mathbf{x}, \hat{\mathcal{X}}), \sup_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} d(\mathcal{X}, \hat{\mathbf{x}}) \right), \text{ where}$$

$$d(\mathbf{x}, \hat{\mathcal{X}}) = \inf_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} d(\mathbf{x}, \hat{\mathbf{x}}) \quad \text{and} \quad d(\mathcal{X}, \hat{\mathbf{x}}) = \inf_{\mathbf{x} \in \mathcal{X}} d(\mathbf{x}, \hat{\mathbf{x}}).$$

The results in Figure 4 show that the Hausdorff distance between estimated signals $\hat{\mathcal{X}}$ and true signals \mathcal{X} is substantially smaller for SMM compared to GMM, regardless of whether $d_{\text{sqe}}(\mathbf{x}, \hat{\mathbf{x}})$ or $d_{\text{abs_cos}}(\mathbf{x}, \hat{\mathbf{x}})$ is used. These findings suggest that SMM achieves more accurate signal recovery than GMM.

IV. APPLICATIONS

We demonstrate EM-based fitting of a SMM in two practical applications. Given (1), the SMM is particularly well suited for measurements in which an observation is underlain by a relative signal, this signal is randomly scaled to an absolute signal by a process not under the control of the measurer, and the observation is also perturbed by external noise. Although this might seem rather particular at first, it is a remarkably common signal structure carried by quite different measurement types in different domains. For example, a biological cell has a particular relative abundance profile of the molecular species it contains. When measured by mass spectrometry, a mass spectrum can report that abundance profile, but measured absolute intensities will be scaled by the chemical matrix, *i.e.*, the overall chemical environment present at that measurement location. Similarly, a sensor aimed at an object to record its color can report an electromagnetic spectrum profile, but that relative profile can be randomly scaled, *e.g.*, by atmospheric circumstances, before it reaches the sensor.

Here, we test SMM-fitting on an imaging mass spectrometry and a hyperspectral imaging dataset. Specifically, we use SMM to recover underlying molecular signatures from noisy IMS measurements of a rat brain tissue section, and to estimate




Fig. 4. Comparison of SMM's and GMM's estimated signals versus the ground truth signals, using a synthetic dataset with $N = 1500$, $K = 3$, $d = 5$, $\sigma^2 = 1.5$, $\pi_1 \approx 0.62$, $\pi_2 \approx 0.22$, and $\pi_3 \approx 0.16$. We repeat the experiment for 100 different initializations of both methods and report the average error at every iteration. The Hausdorff distance between the estimated vectors $\{\hat{\mathbf{x}}_k\}_{k=1}^K$ and true vectors $\{\mathbf{x}_k\}_{k=1}^K$, for both the squared Euclidean distance and the absolute cosine distance, is consistently smaller for SMM compared to GMM, suggesting better signal recovery in the SMM case.

underlying color spectra from the HSI measurement of a scene picturing the Statue of Liberty. In both applications, feature-wise min-max normalization is employed to ensure all features have equal weight.

EM is typically focused on estimating subpopulations of a dataset. However, the E-Step also computes responsibility variables $\rho_k^{[i]}$ that tie an observation i to a subpopulation k , and that can be used to implicitly cluster the data. As clustering pixel measurements is equivalent to segmenting an image, we can use the implicit segmentation results to assess the quality of the estimated subpopulations (or spikes) by means of their corresponding image segment. In these examples, we choose to cluster according to maximum probability. Namely, an observation $i \in [N]$ is associated to cluster c_i , where

$$c_i := \arg \max_{c \in [K]} \rho_c^{[i]}.$$

Furthermore, we compare SMM's clustering results to ones given by traditional methods such as GMM and k -means clustering (k MC).

Note that at every step of Algorithm 1 we need to compute matrices $\mathbf{A}_k = \sum_{i=1}^N \rho_e^{[i]} \mathbf{y}_i \mathbf{y}_i^T$, $\forall k \in [K]$. Without paralleliza-

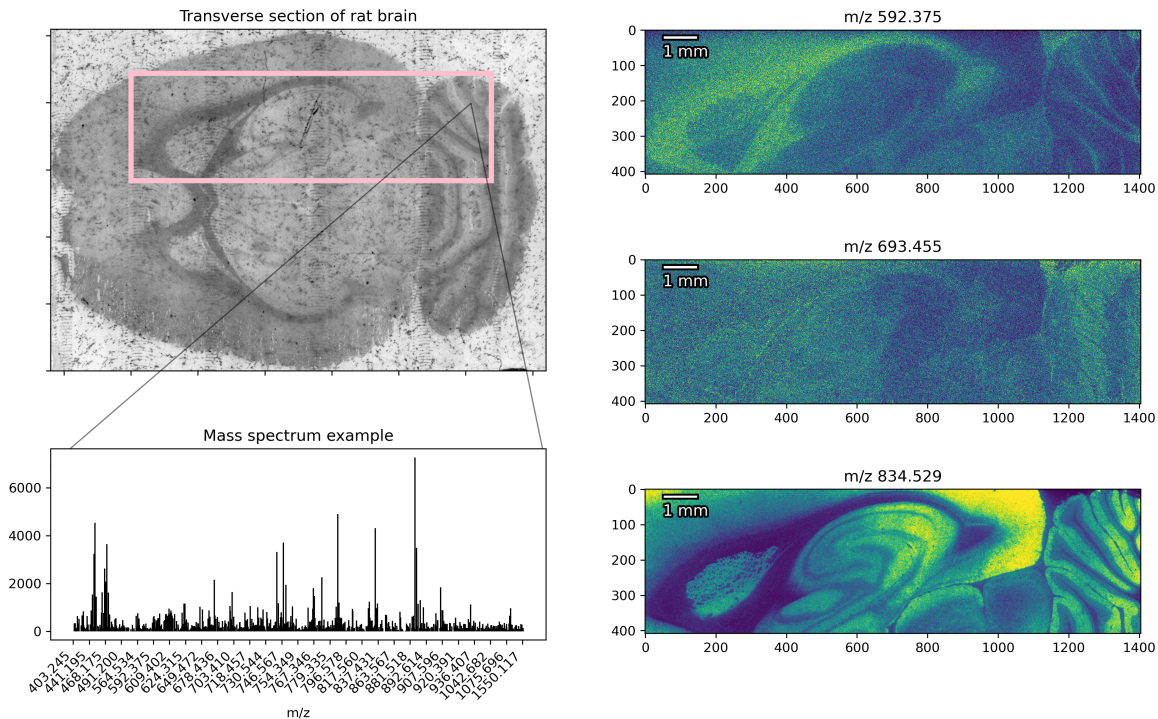


Fig. 5. Rat brain IMS dataset. (top-left) A transverse section of rat brain tissue was analyzed using IMS (analysis area outlined in pink). (right) Three of the 843 ion images acquired, showing the spatial distributions of ion species corresponding to m/z 592.375, 693.455, 834.529. (bottom-left) An example mass spectrum (peak-picked) acquired at a single pixel, reporting 843 m/z features.

tion, this takes $O(N)$ and can significantly slow down the full algorithm. The following lemma shows that the process can be sped up using an approximate matrix $\tilde{\mathbf{A}}_k$, preserving the leading eigenvalue up to a controlled precision.

Lemma IV.1. For every $\delta > 0$, let $\tau := \delta / \|\mathbf{Y}\|_F^2$. The matrix

$$\tilde{\mathbf{A}}_k = \sum_{\substack{i=1 \\ \rho_k^{[i]} \geq \tau}}^N \rho_k^{[i]} \mathbf{y}_i \mathbf{y}_i^T \quad \text{satisfies} \quad |\lambda_1(\mathbf{A}_k) - \lambda_1(\tilde{\mathbf{A}}_k)| \leq \delta.$$

The proof of this lemma uses Weyl's inequality and is provided in Section A of the Supplementary Material. Note that for $\tau \geq 0$ large enough, computing $\tilde{\mathbf{A}}_k$ is much faster since we only need to sum over the indices i satisfying $\rho_e^{[i]} \geq \tau$.

A. Results on imaging mass spectrometry

Imaging mass spectrometry (IMS) [17, 18] is a molecular imaging technique that combines spatial mapping with (mass) spectral analysis. It offers detailed chemical maps of samples such as organic tissues or biofilms, measuring the distributions of hundreds of molecular species throughout a defined measurement region. An IMS measurement can typically be considered as a 3-mode tensor (two spatial modes and one

spectral mode), where every entry reports the ion intensity at a specific spatial location and a particular m/z value. It provides a (gray scale) ion image for each recorded m/z value, visualizing where the compound corresponding to that m/z value is located in the sample. Each pixel implicitly records a full mass spectrum or m/z profile, revealing which molecular compounds are present at that particular location. Since prior labeling of compounds is not required and a single measurement simultaneously reports hundreds of molecular species, IMS has become an important imaging modality for the molecular exploration of the content of biological tissues and for elucidation of disease-related mechanisms.

In this study, a transverse section of a rat brain was measured using Quadrupole Time-of-Flight (QTOF) IMS. After preprocessing (Supplementary Material, Section E), we obtain an IMS dataset with the specifications listed in Table I. Figure 5 shows three of the 843 ion images in this dataset, which depict the distributions of three different molecular species, alongside an example mass spectrum acquired at a particular location, reporting localized abundances for 843 distinct ions there.

Figure 6 compares the clustering outputs obtained by SMM-fitting, GMM-fitting, and k -means clustering, all for $k = 12$ (results for other k in the Supplementary Material). Given that many of the ion images are relatively noisy (see m/z 592.375 and 693.455 ion images in Fig. 5), the results demonstrate SMM's ability to retrieve signals in low-SNR environments. The potential of SMM-based signal recovery is particularly illustrated in the highlighted areas of Fig. 6, where SMM is able to recover biological patterns that are missed by other methods. Specifically, in region 1 (the cerebral cortex), SMM

	Rat brain IMS	Salient HSI
Image size (pixels)	1404 × 408	1024 × 728
Nr. of spectral bands	843	81
Spectral range	[403.245, 1573.905] Da	[380, 780] nm

TABLE I
DATASET SPECIFICATIONS

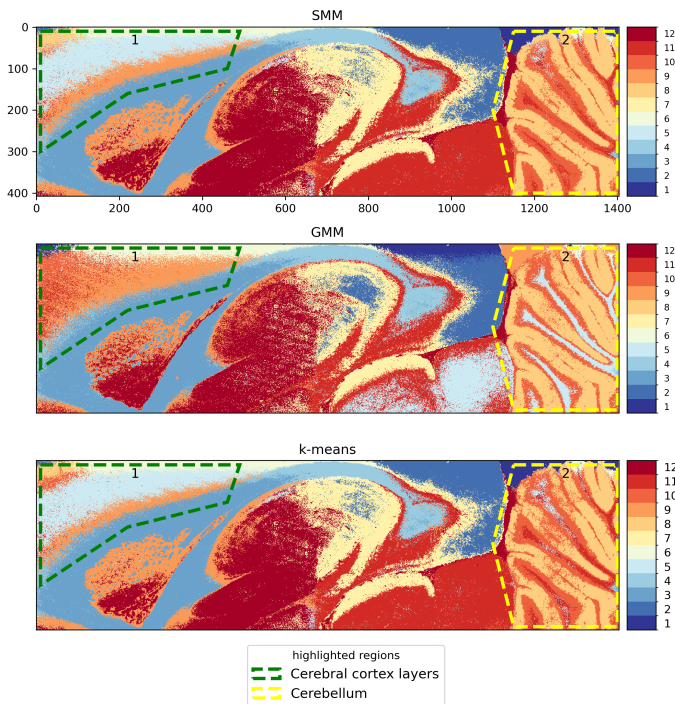


Fig. 6. Clustering results on the rat brain IMS dataset with $k = 12$, for SMM, GMM, and k MC. These results demonstrate SMM's ability to retrieve signals in low-SNR environments. (region 1 – cerebral cortex) SMM discerns a biological subdivision of the cortex layers that is known to exist, but that is missed by GMM. (region 2 – cerebellum) SMM exhibits less susceptibility to noise than k MC.

provides a subdivision of the cortex layers (e.g., molecular, granular, pyramidal, and multiiform) that are known to align parallel to the surface of the brain, but that is missed by GMM. In region 2 (the cerebellum), SMM delivers sharper delineation of key anatomical structures (e.g., white matter, molecular layer, and granule cell layer) and exhibits less susceptibility to noise than k MC. While we do not claim that SMM is superior in all cases, the difference in performance in this example is substantial enough to consider this method. Furthermore, in Fig. 12 in the Supplementary Material, we compare the estimated subpopulation spectra $\{\hat{x}_k\}_{k=1}^K$ for SMM and GMM. The signals estimated by SMM seem closer to real mass spectra than GMM-estimated signals. For example, IMS data is inherently non-negative. Without explicitly imposing non-negativity, SMM-recovered signals are largely non-negative and more similar to real mass spectra than the GMM-recovered signals, which exhibit substantial amounts of negative values.

B. Results on hyperspectral imaging

Hyperspectral imaging (HSI), like IMS, provides both spatial and spectral information, albeit of a different scale and nature with HSI reporting electromagnetic wavelengths. This study uses the salient object dataset introduced in [19], which consists of hyperspectral images of well-known objects captured under various spectral conditions. The dataset was designed to evaluate the performance and robustness of models in detecting and segmenting salient objects. However, we leverage it here because the imaged objects are familiar to most readers (e.g., the Statue of Liberty, Fig. 7), enabling easy interpretation of what constitutes an improvement in segmentation. Details




Fig. 7. RGB reference image for the HSI dataset [19].

of the HSI dataset can be found in Table II.

As in the previous section, Figure 8 compares SMM, GMM, and k MC results for $k = 10$ (results for other k in the Supplementary Material). Figure 7 provides an RGB image for reference, with three regions of interest highlighted: the torch and flame, the tablet in the hand, and the metal railing in the foreground. All three reveal noticeable differences between the methods, described in the caption of Fig. 8. In these examples, SMM exhibits a remarkable ability to discern details that GMM and k MC might not differentiate and instead view as a single signal. The estimated signal spectra are also provided in Fig. 10a of the Supplementary Material. While SMM fitting was not initially intended as a clustering method, its ability to differentiate signal subpopulations that are not discerned by approaches such as GMM and k MC is a testament to its ability for robust signal recovery in a noisy environment.

V. CONCLUSION

We introduced the SMM and a corresponding EM algorithm as a new method for signal recovery. Although it is more restrictive, for data types where the SMM applies, SMM offers substantially better recovery than GMM in low-SNR regimes. While this paper makes a comparison to the standard GMM, one potential future direction could be to compare to more robust versions of GMM [20]. Applications to real-world IMS and HSI data demonstrate more accurate signal recovery than with traditional methods and provide a powerful spatial clustering method, lifting data features from the noise that might otherwise go unrecognized. Another interesting direction for future work could be to generalize our results to settings where the noise and the random scaling distributions have heavier tails, such as the Student's t or generalized hyperbolic distributions, which may enhance robustness to outliers. Such a model could construct an even more robust finite mixture model [21] [22] [23] [24].

VI. ACKNOWLEDGMENT

Research was supported by the National Institutes of Health (NIH)'s NIDDK (U54DK134302 and U01DK133766), NEI (U54EY032442), NIAID (R01AI138581 and R01AI145992), NIA (R01AG078803), NCI (U01CA294527), and by the Chan Zuckerberg Initiative DAF (2021-240339 and 2022-309518). The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.




Fig. 8. Clustering results on the HSI dataset [19] with $k = 10$, for GMM (left), SMM (middle), and kMC (right). These results demonstrate SMM's ability to discern different underlying signals. (torch) SMM separates the flame from the torch handle and, unlike GMM and kMC , avoids the 'halo'-type signal surrounding the torch and flame. (tablet) The tablet is recovered by SMM as a signal distinct from the building behind it, while GMM and kMC both exhibit spill-over between the tablet and building signals. (railing) kMC has difficulty discerning the railing from other dark areas such as the foliage, while SMM and GMM successfully separate the foliage and railing signals. SMM furthermore surpasses GMM in delivering more railing sub-signals, separating out areas with differing amounts of sunlight and shadow (see RGB image in Fig. 7 for reference).

APPENDIX A PROOFS OF SECTION II

Lemma A.1. For fixed values of $\rho_e^{[i]}$ and $S \subseteq [K]$, the estimated variance of the noise

$$\sigma^2 = \frac{\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k}{dN - \sum_{k \in S} \gamma_k} \quad \text{is positive.}$$

Proof. For $d \geq 2$, the denominator is positive since $\sum_{k \in S} \gamma_k \leq N$. For σ^2 to be proven positive, we only need to prove that the numerator is positive. Since $\mathbf{A}_k \succeq 0$, we get:

$$\begin{aligned} \lambda_k &\leq \text{tr}(\mathbf{A}_k) & \forall k \in \{1, \dots, K\} \\ \sum_{k \in S} \lambda_k &\leq \sum_{k \in S} \text{tr}(\mathbf{A}_k) & \text{Summing over } k \in S \\ &\leq \sum_{k=1}^K \text{tr}(\mathbf{A}_k) & S \subseteq [K], \text{tr}(\mathbf{A}_k) \geq 0 \\ &= \text{tr}\left(\mathbf{Y} \sum_{k=1}^K \text{diag}(\rho_k) \mathbf{Y}^T\right) & \text{Using } \mathbf{A}_k \text{ from (8)} \\ &= \text{tr}(\mathbf{Y} \mathbf{Y}^T) & \forall i \in [N], \sum_{k=1}^K \rho_k^{[i]} = 1 \\ &= \|\mathbf{Y}\|_F^2 \end{aligned}$$

□

A. Expression of $\mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]})$ at a critical point

The following intermediate steps allow simplification of the expression of $\mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]})$ at a critical point $\boldsymbol{\theta}$. In Section II, equations for a critical point were obtained $\forall k \in \{1, \dots, K\}$:

$$\pi_k = \frac{\gamma_k}{N} \quad (24)$$

$$\begin{cases} \text{either } \mathbf{x}_k = \mathbf{0} \\ \text{or } \mathbf{A}_k = \lambda_k \mathbf{x}_k, \quad \|\mathbf{x}_k\|^2 = \frac{\lambda_k}{\gamma_k} - \sigma^2 \end{cases} \quad (25)$$

$$\sigma^2 = \frac{\|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \lambda_k}{dN - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k} \quad (26)$$

Using (25) and adding a zero, for $\mathbf{x}_k \neq \mathbf{0}$, we get

$$\begin{aligned} \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} &= \lambda_k \frac{\|\mathbf{x}_k\|^2 + \sigma^2 - \sigma^2}{\|\mathbf{x}_k\|^2 + \sigma^2} \\ &= \lambda_k - \sigma^2 \frac{\lambda_k}{\|\mathbf{x}_k\|^2 + \sigma^2} = \lambda_k - \sigma^2 \gamma_k. \end{aligned} \quad (27)$$

For $\boldsymbol{\theta}$ satisfying (24), $\mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]})$ in (II) can be written as

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]}) &= C - \frac{\|\mathbf{Y}\|_F^2}{2\sigma^2} + \frac{1}{2\sigma^2} \sum_{k=1}^K \frac{\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k}{\|\mathbf{x}_k\|^2 + \sigma^2} \\ &\quad - \frac{1}{2} \sum_{k=1}^K \gamma_k \ln(\|\mathbf{x}_k\|^2 + \sigma^2) - \frac{d-1}{2} N \ln(\sigma^2), \end{aligned}$$

where $C = \sum_{k=1}^K \gamma_k \ln \frac{\gamma_k}{N} - \frac{dN}{2} \ln(2\pi)$, a constant independent of $\boldsymbol{\theta}$.

Using (27) and (25), we can simplify further:

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]}) &= C - \frac{\|\mathbf{Y}\|_F^2}{2\sigma^2} + \frac{1}{2\sigma^2} \sum_{k: \mathbf{x}_k \neq \mathbf{0}} [\lambda_k - \sigma^2 \gamma_k] \\ &\quad - \frac{1}{2} \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k \ln \frac{\lambda_k}{\gamma_k} - \frac{1}{2} \left[(d-1)N + \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k \right] \ln \sigma^2 \\ &= C - \frac{1}{2\sigma^2} \left[\|\mathbf{Y}\|_F^2 - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \lambda_k \right] - \frac{1}{2} \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k \\ &\quad - \frac{1}{2} \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k \ln \frac{\lambda_k}{\gamma_k} - \frac{1}{2} \left[(d-1)N + \sum_{k: \mathbf{x}_k = \mathbf{0}} \gamma_k \right] \ln \sigma^2. \end{aligned}$$

Using (26) and $\sum_{k=1}^K \gamma_k = N$ in the last term, yields

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]}) &= C - \frac{1}{2} \left[dN - \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k \right] - \frac{1}{2} \sum_{k: \mathbf{x}_k \neq \mathbf{0}} \gamma_k \end{aligned}$$

$$-\frac{1}{2} \sum_{k:\mathbf{x}_k \neq \mathbf{0}} \gamma_k \ln \frac{\lambda_k}{\gamma_k} - \frac{1}{2} \left[dN - \sum_{k:\mathbf{x}_k \neq \mathbf{0}} \gamma_k \right] \ln \sigma^2.$$

Finally, denoting $C_2 = C - \frac{dN}{2}$, we get

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\theta}; \boldsymbol{\theta}^{[t]}) &= C_2 - \frac{1}{2} \left[dN - \sum_{k:\mathbf{x}_k \neq \mathbf{0}} \gamma_k \right] \ln \sigma^2 \\ &\quad - \frac{1}{2} \sum_{k:\mathbf{x}_k \neq \mathbf{0}} \gamma_k \ln \frac{\lambda_k}{\gamma_k}. \end{aligned}$$

APPENDIX B

GREEDY OPTIMIZER FOR CRITICAL POINT

We show that a greedy optimizer solves the optimization problem

$$S^* = \underset{S \in \mathcal{V}}{\operatorname{argmin}} g(S) \quad (28)$$

$$\begin{aligned} \text{with } \sigma^2(S) &= \frac{\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k}{dN - \sum_{k \in S} \gamma_k}, \\ g(S) &= \left[dN - \sum_{k \in S} \gamma_k \right] \ln \left(\frac{\|\mathbf{Y}\|_F^2 - \sum_{k \in S} \lambda_k}{dN - \sum_{k \in S} \gamma_k} \right) \\ &\quad + \sum_{k \in S} \gamma_k \ln \left(\frac{\lambda_k}{\gamma_k} \right), \\ \mathcal{V} &= \{S \subseteq [K]; \forall k \in S : \sigma^2(S) \leq \lambda_k / \gamma_k\}. \end{aligned}$$

Lemma B.1 guarantees that removing elements from a valid set only increases the objective function of (28).

Lemma B.1. For $S \subseteq [K]$ and $j \in S$ where $\sigma^2(S) \leq \lambda_j / \gamma_j$,

$$\text{we have: } g(S) \leq g(S \setminus \{j\}).$$

Proof. From the definitions of $\sigma^2(S)$ and $\sigma^2(S \setminus \{j\})$, we find

$$\sigma^2(S) = \frac{\sigma^2(S \setminus \{j\})[dN - \sum_{k \in S} \gamma_k + \gamma_j] - \lambda_j}{dN - \sum_{k \in S} \gamma_k}. \quad (29)$$

We will show that $g(S \setminus \{j\}) - g(S) \geq 0$.

$$\begin{aligned} g(S \setminus \{j\}) - g(S) &= \left[dN - \sum_{k \in S} \gamma_k + \gamma_j \right] \ln \sigma^2(S \setminus \{j\}) + \sum_{k \in S \setminus \{j\}} \gamma_k \ln \frac{\lambda_k}{\gamma_k} \\ &\quad - \left[dN - \sum_{k \in S} \gamma_k \right] \ln \sigma^2(S) - \sum_{k \in S} \gamma_k \ln \frac{\lambda_k}{\gamma_k} \\ &= \left[dN - \sum_{k \in S} \gamma_k \right] \ln \frac{\sigma^2(S \setminus \{j\})}{\sigma^2(S)} - \gamma_j \ln \frac{\lambda_j}{\gamma_j} \\ &\quad + \gamma_j \ln \sigma^2(S \setminus \{j\}) \\ &= \gamma_j \left(-\frac{dN - \sum_{k \in S} \gamma_k}{\gamma_j} \ln \frac{\sigma^2(S)}{\sigma^2(S \setminus \{j\})} - \ln \frac{\lambda_j}{\sigma^2(S \setminus \{j\})\gamma_j} \right) \end{aligned}$$

Using (29), we find the following relation:

$$\frac{\sigma^2(S)}{\sigma^2(S \setminus \{j\})} = \frac{(dN - \sum_{k \in S} \gamma_k)/\gamma_j + 1 - \lambda_j/(\sigma^2(S \setminus \{j\})\gamma_j)}{(dN - \sum_{k \in S} \gamma_k)/\gamma_j}.$$

To simplify notation, we introduce two quantities:

$$a := \frac{dN - \sum_{k \in S} \gamma_k}{\gamma_j} \geq 0, \quad b := \frac{\lambda_j}{\gamma_j \sigma^2(S \setminus \{j\})} \geq 0,$$

$$\text{leading to } \frac{\sigma^2(S)}{\sigma^2(S \setminus \{j\})} = \frac{a+1-b}{a},$$

$$\text{and } \frac{g(S \setminus \{j\}) - g(S)}{\gamma_j} = -a \ln \frac{a+1-b}{a} - \ln b.$$

We now show that the right-hand side quantity is positive. Using $1+b \leq \exp(b)$, we find that $1 + \ln(b) \leq b$. Thus,

$$\ln(b) \leq b - 1 = -a \frac{1-b}{a}. \quad (30)$$

Using again for $z \geq -1$ that $\exp(z) \geq 1+z$, and so that $z \geq \ln(1+z)$ with $z = \frac{1-b}{a}$, we get:

$$\frac{1-b}{a} \geq \ln \left(1 + \frac{1-b}{a} \right) = \ln \left(\frac{a+1-b}{a} \right).$$

$$\text{Since } a \geq 0, \text{ we get } -a \frac{1-b}{a} \leq -a \ln \left(\frac{a+1-b}{a} \right), \quad (31)$$

$$\text{and, combining (30) and (31), } \ln(b) \leq -a \ln \left(\frac{a+1-b}{a} \right).$$

This shows that $-a \ln \left(\frac{a+1-b}{a} \right) - \ln(b) \geq 0$ and, since $\gamma_j \geq 0$, that $g(S \setminus \{j\}) - g(S) \geq 0$. \square

Lemma B.2. There exists a solution \hat{S} to (28) that meets the conditions of a *saturated* set:

$$\begin{aligned} \hat{S} &\in \mathcal{V} & \text{and,} \\ \sigma^2(\hat{S}) &> \frac{\lambda_j}{\gamma_j} & \forall j \in [K] \setminus \hat{S}. \end{aligned}$$

Proof. Since \mathcal{V} is a finite set, there exists a solution to (28). Let \hat{S} be a solution to (28) that is not saturated. This means that $\hat{S} \in \mathcal{V}$, and there exist $j \in [K] \setminus \hat{S}$ such that

$$\sigma^2(\hat{S}) \leq \frac{\lambda_j}{\gamma_j}. \quad (32)$$

From the definitions of $\sigma^2(\hat{S})$ and $\sigma^2(\hat{S} \cup \{j\})$ we know that

$$\sigma^2(\hat{S} \cup \{j\}) = \frac{\sigma^2(\hat{S}) [dN - \sum_{k \in \hat{S}} \gamma_k] - \lambda_j}{dN - \sum_{k \in \hat{S}} \gamma_k - \gamma_j}.$$

Rewriting (32) as $\lambda_j \geq \sigma^2(\hat{S})\gamma_j$, and replacing it above, we find:

$$\sigma^2(\hat{S} \cup \{j\}) \leq \frac{\sigma^2(\hat{S}) [dN - \sum_{k \in \hat{S}} \gamma_k] - \sigma^2(\hat{S})\gamma_j}{dN - \sum_{k \in \hat{S}} \gamma_k - \gamma_j} = \sigma^2(\hat{S}).$$

$$\text{This shows that: } \sigma^2(\hat{S} \cup \{j\}) \leq \frac{\lambda_k}{\gamma_k} \quad \forall k \in \hat{S},$$

$$\sigma^2(\hat{S} \cup \{j\}) \leq \frac{\lambda_j}{\gamma_j} \quad \text{from (32).}$$

It thus follows from the definition of the valid sets \mathcal{V} that $\hat{S} \cup \{j\} \in \mathcal{V}$. Using Lemma B.1, we find that $g(\hat{S} \cup \{j\}) \leq g(\hat{S})$. Moreover, since \hat{S} is a solution to (28), we must have

$$g(\hat{S} \cup \{j\}) = g(\hat{S}),$$

meaning that $\hat{S} \cup \{j\}$ is also a solution. We can recursively apply this reasoning on $\hat{S} \cup \{j\}$. Since $\hat{S} \subseteq [K]$, this recursion must end, at which point we get a saturated solution. \square

Lemma B.3. *There exists a unique saturated set.*

Proof. Suppose for the sake of contradiction that there exists two different saturated sets $S_1, S_2 \in \mathcal{V}$. S_1 being saturated, it satisfies the following inequalities:

$$\sigma^2(S_1) \leq \frac{\lambda_k}{\gamma_k} \quad \forall k \in S_1, \quad (33)$$

$$\sigma^2(S_1) > \frac{\lambda_k}{\gamma_k} \quad \forall k \in [K] \setminus S_1. \quad (34)$$

Similarly, for S_2 we have:

$$\sigma^2(S_2) \leq \frac{\lambda_k}{\gamma_k} \quad \forall k \in S_2, \quad (35)$$

$$\sigma^2(S_2) > \frac{\lambda_k}{\gamma_k} \quad \forall k \in [K] \setminus S_2. \quad (36)$$

Without loss of generality, assume $\sigma^2(S_1) \leq \sigma^2(S_2)$. (37)

Combining (37) with (35), we find:

$$\sigma^2(S_1) \leq \frac{\lambda_k}{\gamma_k} \quad \forall k \in S_2.$$

Using (34), this means that: $\forall j \in S_2 : j \notin [K] \setminus S_1$, or equivalently, that $S_2 \subset S_1$ (the strict inclusion comes from the assumption $S_1 \neq S_2$). We deduce the following relation:

$$\sigma^2(S_1) = \frac{\sigma^2(S_2) [dN - \sum_{k \in S_2} \gamma_k] - \sum_{j \in S_1 \setminus S_2} \lambda_j}{dN - \sum_{k \in S_2} \gamma_k - \sum_{j \in S_1 \setminus S_2} \gamma_j} \quad (38)$$

From (36), we know: $\lambda_j < \gamma_j \sigma^2(S_2)$ for $j \in S_1 \setminus S_2$. Replacing this in (38), we find

$$\begin{aligned} \sigma^2(S_1) &> \frac{\sigma^2(S_2) [dN - \sum_{k \in S_2} \gamma_k] - \sigma^2(S_2) \sum_{j \in S_1 \setminus S_2} \gamma_j}{dN - \sum_{k \in S_2} \gamma_k - \sum_{j \in S_1 \setminus S_2} \gamma_j} \\ &= \sigma^2(S_2). \end{aligned}$$

This contradicts (37). \square

From Lemma B.2 we know that there exists a solution to (28) that is a **saturated** set. Moreover, we know from Lemma B.3 that the saturated set is unique. This means that the saturated set is a solution to (28). Algorithm 2 is a greedy algorithm finding that **saturated** solution in $\mathcal{O}(K^2)$.

Algorithm 2 Greedy set optimization for (28)

INPUT:

$$\lambda_1, \dots, \lambda_K, \gamma_1, \dots, \gamma_K, \|Y\|_F^2$$

INITIALIZATION :

$$S = \emptyset$$

$$L = [K]$$

WHILE $L \neq \emptyset$:

$$\sigma^2(S) = \left[\|Y\|_F^2 - \sum_{k \in S} \lambda_k \right] / (dN - \sum_{k \in S} \gamma_k)$$

$$L = \{k \in [K] \setminus S \text{ s.t. } \sigma^2(S) \leq \lambda_k / \gamma_k\}$$

Choose any $k^* \in L$

$$S = S \cup \{k^*\}$$

OUTPUT:

$$S$$

At every step of algorithm 2 the set S is a valid set. This can be seen by induction. For $S \in \mathcal{V}$, and $k^* \in L$, we have:

$$\sigma^2(S) \leq \lambda_{k^*} / \gamma_{k^*}, \quad \text{and thus,} \quad (39)$$

$$\begin{aligned} \sigma^2(S \cup \{k^*\}) &= \frac{\sigma^2(S) [dN - \sum_{k \in S} \gamma_k] - \lambda_{k^*}}{dN - \sum_{k \in S} \gamma_k - \gamma_{k^*}} \\ &\leq \frac{\sigma^2(S) [dN - \sum_{k \in S} \gamma_k] - \sigma^2(S) \gamma_{k^*}}{dN - \sum_{k \in S} \gamma_k - \gamma_{k^*}} = \sigma^2(S). \end{aligned}$$

Using that $S \in \mathcal{V}$, and (39) we find:

$$\sigma^2(S \cup \{k^*\}) \leq \frac{\lambda_j}{\gamma_j} \quad \forall j \in S, \text{ and}$$

$$\sigma^2(S \cup \{k^*\}) \leq \frac{\lambda_{k^*}}{\gamma_{k^*}}, \text{ showing that } S \cup \{k^*\} \text{ is also in } \mathcal{V}.$$

Algorithm 2 must terminate since set L decreases at every step. When it terminates, set S is **saturated** since L is empty.

APPENDIX C

PROOF OF SOLUTION TO GMM RETRIEVAL

Lemma C.1. *For $\mathbf{A}_1, \dots, \mathbf{A}_K$, K symmetric matrices in $\mathbb{R}^{d \times d}$, and $\mathbf{x}_1, \dots, \mathbf{x}_K \in \mathbb{R}^d$, we have the following:*

$$\sum_{k=1}^K \|\mathbf{A}_k - \mathbf{x}_k \mathbf{x}_k^T\|_F^2 \geq \sum_{k=1}^K \sum_{i=2}^d \lambda_i^2(\mathbf{A}_k), \quad (40)$$

where $\lambda_1(\mathbf{A}_k)$ is the leading eigenvalue of \mathbf{A}_k . With \mathbf{u}_k the eigenvector associated to $\lambda_1(\mathbf{A}_k)$, (40) holds with equality whenever, for all k : $\lambda_1(\mathbf{A}_k) \geq 0$ and $\mathbf{x}_k = \sqrt{\lambda_1(\mathbf{A}_k)} \mathbf{u}_k$.

Proof. Expanding the negative of the left-hand side, we find:

$$-\sum_{k=1}^K \|\mathbf{A}_k - \mathbf{x}_k \mathbf{x}_k^T\|_F^2 = \sum_{k=1}^K -\text{tr}(\mathbf{A}_k^2) + 2\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k - \|\mathbf{x}_k\|^4.$$

Combining this with $\mathbf{x}_k^T \mathbf{A}_k \mathbf{x}_k \leq \lambda_1(\mathbf{A}_k) \|\mathbf{x}_k\|^2$ for $\lambda_1(\mathbf{A}_k)$, the leading eigenvalue of \mathbf{A}_k , we find :

$$\begin{aligned} &-\sum_{k=1}^K \|\mathbf{A}_k - \mathbf{x}_k \mathbf{x}_k^T\|_F^2 \\ &\leq \sum_{k=1}^K [2\lambda_1(\mathbf{A}_k) \|\mathbf{x}_k\|^2 - \|\mathbf{x}_k\|^4 - \text{tr}(\mathbf{A}_k^2)] \end{aligned} \quad (41)$$

$$\begin{aligned} &= \sum_{k=1}^K [-(\lambda_1(\mathbf{A}_k) - \|\mathbf{x}_k\|^2)^2 + \lambda_1^2(\mathbf{A}_k) - \text{tr}(\mathbf{A}_k^2)] \\ &\leq \sum_{k=1}^K [\lambda_1^2(\mathbf{A}_k) - \text{tr}(\mathbf{A}_k^2)] = -\sum_{k=1}^K \sum_{i=2}^d \lambda_i^2(\mathbf{A}_k). \end{aligned} \quad (42)$$

Note that (41) holds with equality whenever \mathbf{x}_k is a rescaled eigenvector \mathbf{A}_k with eigenvalue $\lambda_1(\mathbf{A}_k)$. Also, the inequality in (42) can only hold with equality whenever $\lambda_1(\mathbf{A}_k) \geq 0$, and $\|\mathbf{x}_k\|^2 = \lambda_1(\mathbf{A}_k)$. \square

Lemma C.2. *Let $\hat{\Sigma}_1, \dots, \hat{\Sigma}_K \in \mathbb{R}^{d \times d}$ be symmetric positive semi-definite matrices, such that*

$$\lambda_1(\hat{\Sigma}_j) \geq \frac{\sum_{k=1}^K \sum_{i=2}^d \lambda_i(\hat{\Sigma}_k)}{K(d-1)} \quad \forall j \in [K], \quad (43)$$

where $\lambda_i(\hat{\Sigma}_k)$ is the i -th largest singular value of $\hat{\Sigma}_k$. The solution to

$$\min_{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_K, \hat{\sigma}^2} \sum_{k=1}^K \|\hat{\Sigma}_k - (\hat{\mathbf{x}}_k \hat{\mathbf{x}}_k^T + \hat{\sigma}^2 I)\|_F^2, \quad (44)$$

then is $\sigma_{GMM}^2 = \frac{1}{K} \sum_{k=1}^K \frac{\text{tr}(\hat{\Sigma}_k) - \lambda_1(\hat{\Sigma}_k)}{d-1}$,

$$\mathbf{x}_{GMM,k} = \sqrt{\left(\lambda_1(\hat{\Sigma}_k) - \sigma_{GMM}^2\right)} \mathbf{v}_1(\hat{\Sigma}_k) \quad \forall k \in [K],$$

where $\mathbf{v}_1(\hat{\Sigma}_k)$ is an eigenvector associated to $\lambda_1(\hat{\Sigma}_k)$

Proof. Using Lemma C.1 for $\mathbf{A}_k := \hat{\Sigma}_k - \sigma^2 I$, we find

$$\begin{aligned} \sum_{k=1}^K \|\hat{\Sigma}_k - (\hat{\mathbf{x}}_k \hat{\mathbf{x}}_k^T + \sigma^2 I)\|_F^2 &\geq \sum_{k=1}^K \sum_{i=2}^d \lambda_i^2(\hat{\Sigma}_k - \sigma^2 I) \quad (45) \\ &= \sum_{k=1}^K \sum_{i=2}^d (\lambda_i(\hat{\Sigma}_k) - \sigma^2)^2. \end{aligned}$$

We find the σ^2 that minimizes the previous equation by setting the derivative to zero. This leads us to

$$\sigma_{GMM}^2 = \frac{\sum_{k=1}^K \sum_{i=2}^d \lambda_i(\hat{\Sigma}_k)}{K(d-1)} = \frac{1}{K} \sum_{k=1}^K \frac{\text{tr}(\hat{\Sigma}_k) - \lambda_1(\hat{\Sigma}_k)}{d-1}.$$

Condition (43) implies that $\sigma_{GMM}^2 \leq \lambda_1(\hat{\Sigma}_k) \quad \forall k \in [K]$,

and thus that $\lambda_1(\hat{\Sigma}_k - \sigma_{GMM}^2 I) \geq 0 \quad \forall k \in [K]$. [14]

Finally, we know from Lemma C.1 that (45) is tight when

$$\hat{\mathbf{x}}_k = \mathbf{x}_{GMM,k} = \sqrt{\lambda_1(\hat{\Sigma}_k) - \sigma_{GMM}^2} \mathbf{v}_1(\hat{\Sigma}_k) \quad \forall k \in [K].$$

□

REFERENCES

- [1] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. G. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution," *Science*, vol. 313, pp. 1642 – 1645, 2006.
- [2] N. Hagen and M. W. Kudenov, "Review of snapshot spectral imaging technologies," *Optical Engineering*, vol. 52, 2013.
- [3] A. J. Baeumner, H. Cui, G. Gauglitz, M. C. Moreno-Bondi, S. Szunerits, and A. T. Woolley, "Advancements in sensor technology with innovative and significant research publications: how to write that perfect paper?" *Analytical and Bioanalytical Chemistry*, vol. 414, no. 1, pp. 21–24, 2022.
- [4] K. Senel and E. G. Larsson, "Grant-free massive mtc-enabled massive mimo: A compressive sensing approach," *IEEE Transactions on Communications*, vol. 66, no. 12, pp. 6164–6175, 2018.
- [5] G. G. Artan and G. S. Tombul, "The future trends of eo/ir systems for isr platforms," in *Image Sensing Technologies: Materials, Devices, Systems, and Applications IX*, vol. 12091. SPIE, 2022, pp. 76–88.
- [6] A. Katsevich and A. S. Bandeira, "Likelihood maximization and moment matching in low snr gaussian mixture models," *Communications on Pure and Applied Mathematics*, vol. 76, no. 4, pp. 788–842, 2023.
- [7] I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *The Annals of statistics*, vol. 29, no. 2, pp. 295–327, 2001.
- [8] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra, "Optimality and sub-optimality of pca i: Spiked random matrix models," *The Annals of Statistics*, vol. 46, no. 5, pp. 2416–2451, 2018.
- [9] G. Ben Arous, C. Gerbelot, and V. Piccolo, "Stochastic gradient descent in high dimensions for multi-spiked tensor PCA," *arXiv e-prints*, p. arXiv:2410.18162, 2024.
- [10] G. Casella and R. Berger, *Statistical inference*. CRC Press, 2024.
- [11] M. Al-Shoukairi, P. Schniter, and B. D. Rao, "A gamp-based low complexity sparse bayesian learning algorithm," *IEEE Transactions on Signal Processing*, vol. 66, no. 2, pp. 294–308, 2018.
- [12] H. Snoussi and A. Mohammad-Djafari, "Estimation of structured gaussian mixtures: The inverse em algorithm," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3185–3191, 2007.
- [13] F. Mouret, A. Hippert-Ferrer, F. Pascal, and J.-Y. Tourneret, "A robust and flexible em algorithm for mixtures of elliptical distributions with missing data," *IEEE Transactions on Signal Processing*, vol. 71, pp. 1669–1682, 2023.
- [14] D. Tse and P. Viswanath, "Fundamentals of wireless communication."
- [15] H. Holzmam, A. Munk, and T. Gneiting, "Identifiability of finite mixtures of elliptical distributions," *Scandinavian Journal of Statistics*, vol. 33, no. 4, pp. 753–763, 2006.
- [16] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the royal statistical society: series B (methodological)*, vol. 39, no. 1, pp. 1–22, 1977.
- [17] R. M. Caprioli, T. B. Farmer, and J. Gile, "Molecular imaging of biological samples: localization of peptides and proteins using maldi-tof ms," *Analytical chemistry*, vol. 69, no. 23, pp. 4751–4760, 1997.
- [18] M. E. Colley, A. B. Esselman, C. F. Scott, and J. M. Spraggins, "High-specificity imaging mass spectrometry," *Annual Review of Analytical Chemistry*, vol. 17, 2024.
- [19] N. Imamoglu, Y. Oishi, X. Zhang, G. Ding, Y. Fang, T. Kouyama, and R. Nakamura, "Hyperspectral image dataset for benchmarking on salient object detection," in *2018 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*, 2018, pp. 1–3.
- [20] S. Tadjudin and D. A. Landgrebe, "Robust parameter estimation for mixture model," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, no. 1, pp. 439–445, 2002.
- [21] P. M. Murray, R. P. Browne, and P. D. McNicholas, "Hidden truncation hyperbolic distributions, finite mixtures thereof, and their application for clustering," *Journal of Multivariate Analysis*, vol. 161, pp. 141–156, 2017.

- [22] S. X. Lee and G. J. McLachlan, "Finite mixtures of canonical fundamental skew t -distributions," *Statistics and Computing*, vol. 26, no. 3, pp. 573–589, 2016.
- [23] D. Peel and G. J. McLachlan, "Robust mixture modelling using the t distribution," *Statistics and computing*, vol. 10, pp. 339–348, 2000.
- [24] R. P. Browne and P. D. McNicholas, "A mixture of generalized hyperbolic distributions," *Canadian Journal of Statistics*, vol. 43, no. 2, pp. 176–198, 2015.
- [25] Bruker Daltonics, "Bruker sdk (v2.21)," <https://www.bruker.com/protected/de/services/software-downloads/mass-spectrometry/raw-data-access-libraries.html>.
- [26] P. Monchamp, L. Andrade-Cetto, J. Y. Zhang, and R. Henson, "Signal processing methods for mass spectrometry," *Systems Bioinformatics: An Engineering Case-Based Approach*, Artech House Publishers, 2007.
- [27] Łukasz Migas, "msalign," <https://github.com/lukasz-migas/msalign>.
- [28] G. J. McLachlan and T. Krishnan, *The EM algorithm and extensions*. John Wiley & Sons, 2008.



Paul-Louis Delacour received the B.Sc. degree in communication systems (computer science) from École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2019 and the M.Sc. degree in data science from the Federal Institute of Technology Zurich (ETH Zurich), Zurich, Switzerland, in 2022. He is currently working towards the Ph.D. degree at Delft University of Technology (TU Delft), Delft, The Netherlands. His research interest includes mathematics & machine learning with a focus on high-dimensional learning.



Sander Wahls (S'09–M'12–SM'17) received the Dipl.-Math. degree in mathematics and the Dr.-Ing. degree (summa cum laude) in electrical engineering from Technische Universität Berlin, Berlin, Germany, in 2007 and 2011, respectively.

He is currently a Full Professor of information technology at Karlsruhe Institute of Technology, Karlsruhe, Germany. From 2014 to 2023, he was with Delft University of Technology, Delft, The Netherlands, first as Assistant Professor, and later as Associate Professor. From 2012 to 2014, he was a Postdoctoral Research Fellow at Princeton University, Princeton, NJ, USA.

Dr. Wahls is a member of VDE-ITG and SIAM. He received the Johann-Philipp-Reis Preis from VDE-ITG in 2015, and a Starting Grant from the European Research Council (ERC) in 2016.



Jeffrey M. Spraggins received the B.A. degree in chemistry from the College of Wooster and the Ph.D. degree in analytical chemistry from the University of Delaware in 2010.

He is currently Associate Professor in the Department of Cell and Developmental Biology and Director of both the Mass Spectrometry Research Center (MSRC) and Biomolecular Multimodal Imaging Center (BIOMIC) at Vanderbilt University, Nashville, TN, USA. Dr. Spraggins' research focuses on two main areas: (1) the development of advanced mass spectrometry technologies to enhance imaging performance and enable molecular histology at cellular resolution, and (2) combining imaging mass spectrometry with other complementary biomedical imaging modalities to create new, integrated platforms for systems-level analysis of tissue biology. His group applies these approaches to the construction of comprehensive molecular atlases for the NIH Human Biomolecular Atlas Program (HuBMAP), the NIDDK Kidney Precision Medicine Project (KPMP), and the NCI Human Tumor Atlas Network (HTAN). In parallel, his team investigates the molecular underpinnings of Alzheimer's, kidney, and infectious diseases.



Łukasz Migas received the MChem degree in chemistry from Manchester Metropolitan University, Manchester, UK, in 2014, and the Ph.D. degree in biochemistry from the University of Manchester, Manchester, UK, in 2018. He is currently a Postdoctoral Research Fellow at the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands, and at the Mass Spectrometry Research Center, Vanderbilt University, Nashville, TN, USA. His research interests include the development of computational methods

and user-friendly graphical applications for the analysis and interpretation of imaging mass spectrometry and microscopy data.



Raf Van de Plas (Member, IEEE) received the Industriële Ingenieur degree in electronics from Groep T, Leuven, Belgium, in 2002. He received the Master in Artificial Intelligence degree and the Doctor of Engineering (Ph.D.) degree from the Katholieke Universiteit Leuven, Leuven, Belgium, in 2003 and 2010, respectively.

He was a Postdoctoral Researcher in the Department of Electrical Engineering (ESAT) at the Katholieke Universiteit Leuven in 2010, and a Research Fellow in the Department of Biochemistry through the Mass Spectrometry Research Center at Vanderbilt University, Nashville, TN, USA, from 2011 to 2012. During 2011, he was an Honorary Fulbright Scholar at Vanderbilt University. From 2012 to 2014, he joined Vanderbilt University's School of Medicine research faculty as Research Instructor. From 2014 to 2022, he was Assistant Professor in the Delft Center for Systems and Control at the Delft University of Technology (TU Delft), Delft, The Netherlands. Since 2014, he holds a position as Adjunct/Adjoint Assistant Professor of Biochemistry at the Vanderbilt University School of Medicine. He is currently Associate Professor in the Delft Center for Systems and Control at the Delft University of Technology. He works on signal processing and machine learning for spectral imaging systems (e.g., imaging mass spectrometry), molecular imaging modalities (e.g., microscopy), and data-driven multimodal fusion between imaging technologies. His research focuses on the interface between (i) mathematical engineering and machine learning; (ii) analytical chemistry and instrumentation; and (iii) life sciences and medicine.

Dr. ing. Van de Plas is a member of the American Society for Mass Spectrometry (ASMS), the Belgian Society for Mass Spectrometry (BSMS), and the Dutch Society for Mass Spectrometry (NVMS).