Statistical Inference for Fuzzy Clustering

Qiuyi Wu[1]*, Zihan Zhu[1]† and Anru R. Zhang[2]‡

January 6, 2026

Abstract

Clustering is a central tool in biomedical research for discovering heterogeneous patient subpopulations, where group boundaries are often diffuse rather than sharply separated. Traditional methods produce hard partitions, whereas soft clustering methods such as fuzzy c-means (FCM) allow mixed memberships and better capture uncertainty and gradual transitions. Despite the widespread use of FCM, principled statistical inference for fuzzy clustering remains limited.

We develop a new framework for weighted fuzzy c-means (WFCM) for settings with potential cluster-size imbalance. Cluster-specific weights re-balance the classic FCM criterion so that smaller clusters are not overwhelmed by dominant groups, and the weighted objective induces a normalized density model with scale parameter σ and fuzziness parameter m. Estimation is performed via a blockwise majorize–minimize (MM) procedure that alternates closed-form membership/centroid updates with likelihood-based updates of $(\sigma, \mathbf{w})$; the intractable normalizing constant is approximated by importance sampling with a data-adaptive Gaussian-mixture proposal. We further provide likelihood-ratio tests for comparing cluster centers and bootstrap-based confidence intervals.

*Department of Biostatistics & Bioinformatics, Duke University

†Department of Statistics & Data Science, University of Pennsylvania

‡Department of Biostatistics & Bioinformatics and Department of Computer Science, Duke University

[1]Equal contribution. [2]Corresponding author. E-mail: `anru.zhang@duke.edu`

We establish consistency and asymptotic normality of the maximum likelihood estimator, validate the method in simulations, and illustrate it on single-cell RNA-seq and Alzheimer's Disease Neuroimaging Initiative (ADNI) data. These applications show stable uncertainty quantification and biologically meaningful soft memberships, ranging from well-separated cell populations under imbalance to a graded AD–non-AD continuum consistent with disease progression.

1 Introduction

Clustering is a fundamental task in unsupervised learning, with widespread applications in bioinformatics, social science, and image analysis. In many biomedical and clinical domains, subgroup boundaries are rarely clear-cut. For instance, sepsis can exhibit overlapping phenotypes in which patients display mixed inflammatory and hemodynamic responses (Jiang et al., 2024). Similarly, neuroimaging studies often suggest disease processes that evolve along a continuum rather than a strict dichotomy, as in Alzheimer's disease and mild cognitive impairment. These examples motivate clustering methods that can explicitly represent uncertainty and gradual transitions between subgroups.

Traditional clustering algorithms such as k-means (Lloyd, 1982) and hierarchical clustering (Sokal & Sneath, 1963; Ward Jr, 1963) assign each observation to a single cluster, producing hard partitions of the data. While effective in many contexts, hard clustering can fail to reflect uncertainty and ambiguity in complex datasets. Soft clustering, in contrast, assigns each observation a vector of membership weights that quantify partial association with clusters. Among distance-based soft clustering methods, fuzzy c-means (FCM) (Bezdek, 2013) has been particularly influential, jointly estimating cluster centers and membership weights. Despite widespread empirical use, fuzzy clustering is often treated primarily as an optimization procedure, and formal statistical inference and uncertainty quantification remain underdeveloped.

1.1 Our Contributions

This paper addresses this gap by developing a framework for statistical inference in soft clustering based on a weighted extension of fuzzy c-means. Our main contributions are as follows.

We propose a *weighted* fuzzy c-means (WFCM) formulation that adjusts to settings with severe cluster-size imbalance. By introducing cluster-specific weights, the objective re-balances the classic FCM criterion so that minority (or higher-precision) clusters are not dominated by large groups. We show that this weighted objective naturally induces a normalized density model with scale parameter σ and fuzziness parameter m, under which fuzzy memberships admit closed-form, weighted FCM-style updates. The resulting log-likelihood separates into an intractable normalizing constant and the WFCM loss, linking maximum likelihood estimation to minimizing a weighted FCM criterion.

To estimate the model parameters under the induced likelihood, we develop a block-wise MM procedure that alternates between closed-form membership and centroid updates. The intractable normalizing constant is approximated via importance sampling with a data-adaptive Gaussian-mixture proposal, and we use a post-MM refinement step to further reduce the negative log-likelihood. Building on the fitted model, we provide practical inferential tools, including likelihood-ratio tests for comparing cluster centers and bootstrap-based confidence regions with label alignment. We also propose a weighted Xie-Beni index for selecting the number of clusters in a way that is consistent with the weighted membership structure.

We establish asymptotic guarantees for the proposed likelihood formulation. In particular, we prove strong consistency of the MLE up to label permutation (Theorem 1) and asymptotic normality under mild regularity conditions (Theorem 2), with covariance characterized by the (generalized) inverse Fisher information. A key technical condition is a strictly positive lower bound on component weights, which ensures identifiability of the cluster centers and enables valid inference.

Simulation studies corroborate the theoretical results and demonstrate accurate uncertainty quantification via both bootstrap and Fisher-information-based approximations. We

3

further illustrate the framework on two biomedical datasets. In scRNA-seq PBMC data (memory T cells, B cells, monocytes), the method yields tight uncertainty regions even when clusters are biologically distinct under class imbalance. In ADNI (Alzheimer's Disease Neuroimaging Initiative) neuroimaging, soft memberships reveal a graded Alzheimer's disease-to-non-Alzheimer's disease continuum and highlight intermediate subjects plausibly corresponding to prodromal or clinically ambiguous cases (e.g., mild cognitive impairment, MCI), suggesting potential utility for early-warning risk assessment.

1.2   Related Literature

Clustering is a core problem in statistics and machine learning that aims to group observations with high within-cluster similarity and low between-cluster similarity (Ball & Hall, 1965; Basu, Davidson, & Wagstaff, 2008; Duda, Hart, et al., 2006; Jain, 2010). Classical approaches include hard partitioning methods such as k-means and hierarchical clustering, as well as probabilistic model-based clustering (Jain & Dubes, 1988; Meilă, 2006).

Soft clustering assigns each observation a vector of membership weights. Two prominent paradigms are mixture-model methods (e.g., Gaussian mixture models, GMMs) and distance-based fuzzy clustering. GMMs posit a finite mixture of Gaussians and estimate parameters by maximum likelihood, typically via EM (Dempster, Laird, & Rubin, 1977; McLachlan & Peel, 2000), enabling likelihood-based uncertainty quantification but requiring covariance estimation and distributional assumptions that can be fragile in high-dimensional or small-sample regimes. In contrast, fuzzy c-means (FCM) (Bezdek, Ehrlich, & Full, 1984) is a distribution-free alternative that optimizes a distance-based objective with a fuzziness parameter controlling softness, and has been widely used in applications (Hathaway & Bezdek, 2006; Pham & Prince, 1999). Numerous extensions improve robustness and flexibility, including adaptive/metric variants, regularized formulations, spatially constrained and entropy-penalized versions, and kernel/possibilistic approaches (Fan, Zhen, & Xie, 2003; Gath & Geva, 2002; Gustafson & Kessel, 1979; Ji, Sun, & Xia, 2011; Zhang & Chen, 2003, 2004).

Despite extensive algorithmic development, formal statistical inference for optimization-

based fuzzy clustering remains relatively limited: FCM does not naturally provide a generative likelihood, and symmetry across clusters can obscure size/density differences, complicating identifiability and uncertainty quantification. Related directions include overlapping clustering and mixed-membership models (Banerjee, Krumpelman, Ghosh, Basu, & Mooney, 2005; J. Xie, Kelley, & Szymanski, 2013), as well as post-clustering or selective inference, which aims to quantify uncertainty while accounting for the data-dependent clustering step (Gao, Bien, & Witten, 2024; Taylor & Tibshirani, 2015; Yun & Foygel Barber, 2023; Zhu, Gai, & Zhang, 2024).

Finally, cluster validity indices are commonly used for model selection in fuzzy clustering. The Xie–Beni index (XBI) (X. L. Xie & Beni, 1991) measures a compactness–separation trade-off and is frequently used to select the number of clusters and the fuzziness parameter. In this paper, we derive a weighted Xie–Beni index tailored to our weighted fuzzy c-means (WFCM) framework.

1.3 Organization

The remainder of this paper is organized as follows. Section 2 introduces notation and preliminaries on classic fuzzy c-means. Section 3 introduces the new statistical weighted fuzzy c-means framework, beginning with the classic fuzzy c-means algorithm and extending it to incorporate cluster-specific weights. Section 4 details the computational procedures, including parameter estimation, importance sampling for the normalizing constant, statistical inference, and model selection via the weighted Xie–Beni index. In Section 5, we develop the theoretical properties of the proposed method, establishing its consistency and asymptotic normality. Section 6 presents simulation studies that evaluate clustering accuracy, uncertainty quantification, and robustness. In Section 7, we illustrate the method on two real-world biomedical applications: single-cell RNA sequencing data and Alzheimer's Disease Neuroimaging Initiative data. Finally, we conclude with a discussion of the implications and potential extensions of our framework.

2 Notation and Preliminaries on Classic Fuzzy C-Means

We use boldface letters, e.g., $\mathbf{x}, \mathbf{y}, \mathbf{z}$, to denote vectors, and regular letters, e.g., x, y, z, to denote scalars. $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are n samples, and $x_{i1}, \ldots, x_{id}$ are d entries of $\mathbf{x}_i$. For any vector $\mathbf{x} \in \mathbb{R}^d$, we write $\mathbf{x} = (x_1, \ldots, x_d)$ for its components and define its ℓ_2-norm as $\|\mathbf{x}\| = \sqrt{\sum_{j=1}^d x_j^2}$. For any non-negative integer n, we denote $[n] = \{1, \ldots, n\}$.

Next, we describe the classic fuzzy c-means (FCM) method as the benchmark of our work. The FCM algorithm attempts to partition a finite collection of n d-dimensional observations $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \subseteq \mathbb{R}^d$ into a collection of k fuzzy clusters with respect to a given criterion. Given a finite set of data, the algorithm returns an estimate $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$ of the cluster centroids and a partition matrix $\mathbf{U} = (u_{ij}) \in [0,1]^{n \times k}, i = 1, \ldots, n; j = 1, \ldots, k$, where the pairwise distances $d_{ij} = \|\mathbf{x}_i - \mathbf{v}_j\|, i \in [n], j \in [k]$ and the membership coefficient of point $\mathbf{x}_i$ to cluster j is

$$
\begin{aligned}
&\text{if } m > 1, \ \ u_{ij} = \left\{ \sum_{l=1}^k \left(\frac{d_{ij}}{d_{il}} \right)^{2/(m-1)} \right\}^{-1}; \\
&\text{if } m = 1, \ \ u_{ij} = \begin{cases} \frac{1}{|A|}, & d_{ij} = \min_{l' \in [k]} d_{il'}, A = \{l : d_{il} = \min_{l' \in [k]} d_{il'}\}; \\ 0 & d_{ij} \neq \min_{l' \in [k]} d_{il'}. \end{cases}
\end{aligned}
\tag{1}
$$

Here, $m \geq 1$ represents the "fuzziness" or the degree of softness: when $m = 1$, the method corresponds to the hard clustering and the k-means algorithm. The membership and cluster centers are chosen to minimize the FCM objective function, which can be rewritten as

$$
\begin{aligned}
\mathcal{J}_m^{FCM}(\mathbf{V}) &= \sum_{i=1}^n \sum_{j=1}^k u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2 \\
&= \sum_{i=1}^n \sum_{j=1}^k \frac{d_{ij}^2}{\left[\sum_{l=1}^k (d_{ij}/d_{il})^{\frac{2}{m-1}} \right]^m} = \sum_{i=1}^n \frac{1}{\left[\sum_{j=1}^k (d_{ij}^2)^{-\frac{1}{m-1}} \right]^{m-1}}.
\end{aligned}
\tag{2}
$$

We outline several significant properties of the fuzzy c-means. These properties are derived from the relationship among generalized means (Bullen, 2013).

1. When the fuzziness parameter $m \to 1$, we have $\mathcal{J}_m^{FCM} \to \sum_{i=1}^{n} \min_{j=1,\ldots,k} d_{ij}^2$. This corresponds to the objective function of k-means. Thus, FCM reduces to the classical k-means algorithm as $m \to 1$.

2. When $m \to \infty$, a rescaled FCM loss converges to the geometric mean of the squared distances d_{ij}^2:

$$k^{m-1} \mathcal{J}_m^{FCM}(\mathbf{V}) = \sum_{i=1}^{n} \left[\sum_{j=1}^{k} d_{ij}^{-\frac{2}{m-1}} / k \right]^{-m+1} \to \sum_{i=1}^{n} \prod_{j=1}^{k} d_{ij}^{2/k}.$$

3. Another intermediate case of interest in the literature is $m = 2$, in which the objective function simplifies to the harmonic mean of the squared distances

$$\mathcal{J}_2^{FCM}(\mathbf{V}) = \sum_{i=1}^{n} \frac{1}{\sum_{j=1}^{k} d_{ij}^{-2}}.$$

The procedure of classic FCM follows an iterative optimization scheme (Bezdek, 2013). It begins by randomly initializing the cluster centers $\mathbf{V}$. Given the current cluster centers, the membership coefficients u_{ij} are updated using the distance-based formula described above. With the updated memberships, each cluster center $\mathbf{v}_j$ is recalculated as the weighted average of all data points $\mathbf{v}_j = \frac{\sum_i u_{ij}^m \mathbf{x}_i}{\sum_i u_{ij}^m}$. These two steps of updating memberships and updating cluster centers are repeated alternately until the objective function or the membership matrix converges. This alternating optimization ensures that the algorithm monotonically decreases the objective at each iteration and typically converges to a local minimum.

Compared with other soft clustering approaches such as the Gaussian mixture model (GMM), FCM offers several advantages. It is a model-free method that does not assume any particular data distribution, while GMMs rely on the Gaussian assumption that may not hold in practice. FCM is also computationally simpler, as it avoids the estimation of covariance matrices and mixing proportions required in GMMs. Its fuzzy memberships have a geometric interpretation based solely on distances to cluster centers, leading to more transparent and easily interpretable results. Furthermore, FCM is readily extensible, allowing the incorporation of weights, spatial penalties, or feature regularization without altering its core structure.

However, the classic fuzzy c-means has many limitations. First, the method does not immediately yield a statistical inference or uncertainty quantification, which can hinder rigorous interpretation of clustering results. Second, it treats all clusters equally in the objective, regardless of their relative sizes or importance, whereas in real datasets some clusters may contain substantially more points than others (see Figure 1). Third, FCM does not yield a generative model like Gaussian mixture models. It does not define a likelihood function or allow the generation of synthetic samples from the learned clusters. These limitations motivate extensions of FCM, such as weighted fuzzy c-means, which incorporate cluster-specific contributions and probabilistic interpretations, offering a more flexible modeling framework and enabling statistical inference.

3 A Statistical Framework for Weighted Fuzzy C-Means

We develop statistical inference for soft clustering in settings where clusters may have markedly different sizes. Let $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^d$ be observed data. Our goal is to assign them to a small number of (potentially overlapping) latent clusters while retaining uncertainty information through fuzzy memberships.

Weighted fuzzy c-means. A well-known limitation of standard fuzzy c-means (FCM) is that all clusters are treated as equally important in the objective. When cluster sizes are imbalanced, large clusters dominate the optimization, pulling centroids toward dense regions and degrading the membership estimates of smaller clusters (Figure 1). To mitigate this, we introduce cluster-specific weights and obtain a *weighted* FCM objective.

Let $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$ denote the cluster centroids. Let

$$\mathbf{w} = (w_1, \ldots, w_k) \in \mathcal{S}_k := \left\{ (w_1, \ldots, w_k) : w_j \geq 0 \ \forall j, \ \sum_{j=1}^{k} w_j = 1 \right\}$$

be a weight vector over clusters, and write $d_{ij} = \|\mathbf{x}_i - \mathbf{v}_j\|$. w_j acts as a cluster-specific
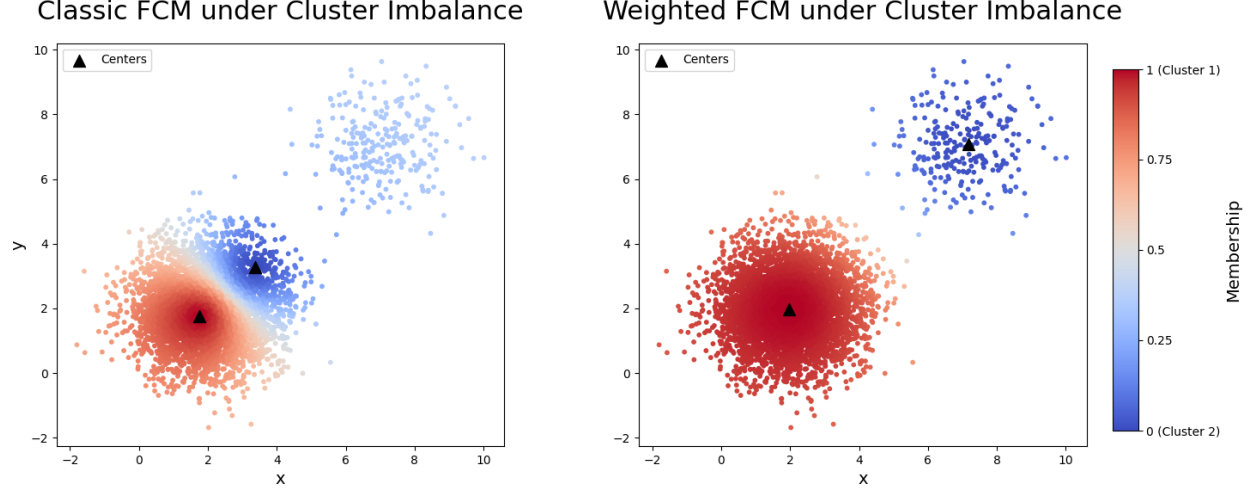
Figure 1: Illustration of cluster imbalance in classic and weighted fuzzy c-means. Due to the imbalance in cluster sizes, the larger cluster exerts stronger influence on the optimization of classic FCM, causing centroid shifts toward denser regions and reducing membership accuracy for the smaller cluster in the left figure. While our proposed WFCM incorporates density-based weights to correct for sample imbalance and thus yields more balanced membership assignments and more accurate cluster centers for the minority group.

penalty/inverse-importance factor, allowing us to upweight minority clusters. For the fuzziness parameter $m > 1$, we define the WFCM loss as

$$\mathcal{J}_m^{\mathrm{WFCM}}(\mathbf{V}, \mathbf{w}) = \sum_{i=1}^{n} \Big[\sum_{j=1}^{k} \big(w_j d_{ij}^2 \big)^{-\frac{1}{m-1}} \Big]^{-m+1}. \tag{3}$$

The weights $\mathbf{w}$ directly re-balance the contribution of each cluster to the data-fitting term.

A density induced by WFCM. To carry out statistical inference and uncertainty quantification, we embed the WFCM loss into a probabilistic model. Define a family of densities on $\mathbb{R}^d$ by

$$f(\mathbf{x}; \mathbf{V}, \mathbf{w}, m, \sigma) = C_{\mathbf{V}, \mathbf{w}, m, \sigma} \exp\left(-\frac{1}{\sigma^2 \big[\sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}} \big]^{m-1}} \right), \quad \mathbf{x} \in \mathbb{R}^d, \tag{4}$$

9

where $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$ are the centroids, $\mathbf{w} \in \mathcal{S}_k$ are the cluster weights, $m > 1$ is the fuzziness parameter, and $\sigma > 0$ is a scale parameter. The normalizing constant

$$C_{\mathbf{V},\mathbf{w},m,\sigma} = \left\{ \int_{\mathbb{R}^d} \exp\left(-\frac{1}{\sigma^2 \left[\sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}} \right]^{m-1}} \right) d\mathbf{x} \right\}^{-1} \qquad (5)$$

ensures that f integrates to one.

This model has a direct clustering interpretation: the term

$$\left[\sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}} \right]^{-(m-1)}$$

plays the same role as the pointwise contribution in (3), so that high density regions coincide with locations close to one (or more) centroids, modulated by the cluster weights.

Figure 2 illustrates the density for various values of σ^2 and m. Each panel displays the density with equal weights ($w_1 = w_2$), evaluated over a two-dimensional grid with two cluster centroids (red crosses). Smaller m values produce sharper and more localized modes, approaching hard clustering behavior, while larger m values yield smoother, overlapping regions that reflect higher fuzziness. Increasing σ^2 broadens the overall density, further flattening the cluster boundaries. The limiting behavior of the WFCM density as $m \to 1^+$, as $m \to \infty$, and at $m = 2$ is provided in Section C of the supplementary materials.

Memberships, likelihood, and estimation Under (4), the fuzzy membership of observation $\mathbf{x}_i$ in cluster j takes the closed form: $\mathbf{U} = (u_{ij})_{i \in [n], j \in [k]}$,

$$u_{ij} = \frac{w_j^{-\frac{1}{m-1}} \|\mathbf{x}_i - \mathbf{v}_j\|^{-\frac{2}{m-1}}}{\sum_{l=1}^{k} w_l^{-\frac{1}{m-1}} \|\mathbf{x}_i - \mathbf{v}_l\|^{-\frac{2}{m-1}}} = \frac{1}{\sum_{l=1}^{k} \left(\frac{w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2}{w_l \|\mathbf{x}_i - \mathbf{v}_l\|^2} \right)^{\frac{1}{m-1}}} = \frac{(w_j d_{ij}^2)^{-\frac{1}{m-1}}}{\sum_{l=1}^{k} (w_l d_{il}^2)^{-\frac{1}{m-1}}}. \qquad (6)$$

Thus, the weights $\mathbf{w}$ directly re-balance the contribution of each cluster to the data-fitting term, yielding the usual FCM membership formula (1) but now explicitly weighted by $\mathbf{w}$.

Assume $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are independent draws from $f(\cdot; \mathbf{V}, \mathbf{w}, m, \sigma)$. The joint likelihood is
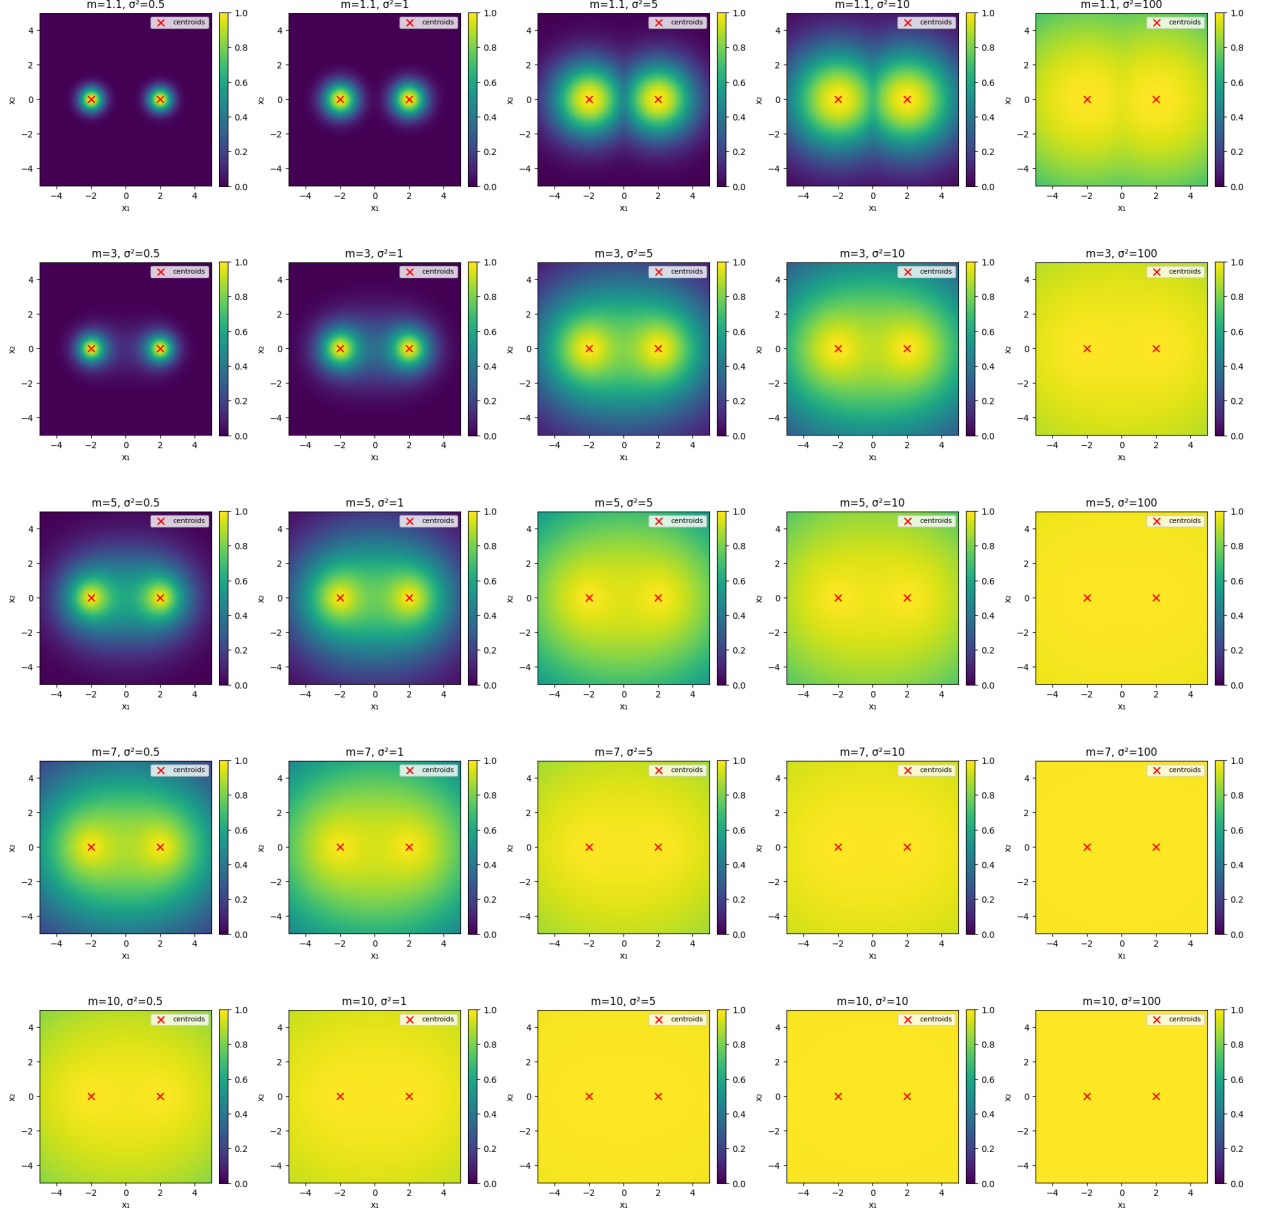
10

Figure 2: Induced density $f(\mathbf{x})$ in (4) for different values of σ^2 and m in two dimensions with two centroids (red crosses) and equal weights $w_1 = w_2$. Smaller m produces sharper, nearly hard-clustering modes; larger m yields smoother, overlapping regions. Increasing σ^2 broadens the density and flattens the boundaries.

$L(\mathbf{V}, \mathbf{w}, m, \sigma) = \prod_{i=1}^{n} f(\mathbf{x}_i; \mathbf{V}, \mathbf{w}, m, \sigma)$ and the log-likelihood can be written as

$$\ell(\mathbf{V}, \mathbf{w}, m, \sigma) = \sum_{i=1}^{n} \log f(\mathbf{x}_i; \mathbf{V}, \mathbf{w}, m, \sigma)$$

$$= \sum_{i=1}^{n} \log C_{\mathbf{V}, \mathbf{w}, m, \sigma} - \frac{1}{\sigma^2} \sum_{i=1}^{n} \frac{1}{\left[\sum_{j=1}^{k} (w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}} \right]^{m-1}}$$

$$= n \log C_{\mathbf{V}, \mathbf{w}, m, \sigma} - \frac{1}{\sigma^2} \mathcal{J}_m^{\text{WFCM}}(\mathbf{V}, \mathbf{w}). \tag{7}$$

Hence, maximizing the log-likelihood is nearly equivalent to *minimizing* the WFCM loss, differing only by the contribution of the normalizing constant.

A natural choice of estimator is the maximum likelihood estimator for the centroids $\mathbf{V}$, weights $\mathbf{w}$, fuzziness parameter m, and variance σ^2:

$$(\hat{\mathbf{V}}, \hat{\mathbf{w}}, \hat{m}, \hat{\sigma}) \in \underset{\mathbf{V}, \mathbf{w} \in \mathcal{S}_k, m>1, \sigma>0}{\operatorname{argmax}} \ell(\mathbf{V}, \mathbf{w}, m, \sigma), \tag{8}$$

with ℓ given in (7). Direct optimization, however, is nontrivial because (i) the memberships (6) depend on $(\mathbf{V}, \mathbf{w}, m)$ in a nonlinear way, and (ii) the normalizing constant $C_{\mathbf{V}, \mathbf{w}, m, \sigma}$ in (5) also depends on the parameters. To address these coupled dependencies, we adopt a blockwise MM iterative scheme. This procedure, detailed in Section 4, alternates between updating the fuzzy memberships and updating the model parameters.

Remark 1 (WFCM versus Gaussian mixture modeling) *Weighted fuzzy c-means (WFCM) and Gaussian mixture models (GMMs) both produce soft cluster assignments but differ fundamentally in their underlying principles and modeling assumptions. In GMM-based soft clustering, the data are modeled as arising from a mixture of k Gaussian components, with density*

$$p(\mathbf{x}_i) = \sum_{j=1}^{k} \pi_j \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_j, \Sigma_j),$$

where π_j are the mixing proportions, $\boldsymbol{\mu}_j$ the component means, and Σ_j the covariance matrices. The membership (or responsibility) of $\mathbf{x}_i$ for cluster j is then given by the posterior probability

$$u_{ij} = \frac{\pi_j \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_j, \Sigma_j)}{\sum_{l=1}^{k} \pi_l \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_l, \Sigma_l)}.$$

By contrast, WFCM is a distance-based, assumption-light approach that only requires a metric space and cluster centroids. It estimates the centroids $\mathbf{V} \in (\mathbb{R}^d)^k$, cluster weights $\mathbf{w} \in \mathcal{S}_k$, the fuzziness parameter m, and the scale σ, involving roughly $\mathcal{O}(kd)$ unknowns. A full-covariance GMM introduces an additional $\mathcal{O}(kd^2)$ covariance parameters, which can make estimation data-hungry and unstable in moderate to high dimensions.

GMM and WFCM also differ in the semantics of their soft assignments: WFCM memberships are proximity-based, with $u_{ij} \propto (w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2)^{-1/(m-1)}$, reflecting relative distances,

12

while GMM responsibilities represent posterior probabilities that depend additionally on component volumes $|\Sigma_j|$ and priors π_j.

Finally, the fuzziness parameter m in WFCM offers direct control over the degree of softness and yields simple, inversion-free updates, whereas the EM algorithm for GMMs is computationally heavier due to covariance updates and matrix inversions. Overall, WFCM is advantageous for metric-driven, assumption-light soft clustering, particularly for imbalanced cluster sizes or high-dimensional data, while GMM is preferable when probabilistic interpretability, likelihood-based model selection, or anisotropic cluster shapes are desired.

4   Weighted Fuzzy C-Means: Algorithm and Statistical Inference

In this section, we develop a blockwise Majorize-Minimization (MM)–type algorithm for weighted fuzzy c-means (WFCM) clustering, extending the statistical framework introduced in the previous section. The algorithm follows the general structure of fuzzy clustering but incorporates cluster-specific weights to correct for size imbalance, thereby improving estimation stability and interpretability compared to standard fuzzy c-means.

4.1   Blockwise MM-Like Fuzzy Clustering Algorithm

The blockwise MM algorithm estimates the parameters $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$, $\mathbf{w} = (w_1, \ldots, w_k)$, the scale parameter σ, and the fuzziness parameter m. Instead of jointly updating the fuzzy memberships $\mathbf{U} = \{u_{ij}\}_{i \in [n], j \in [k]}$ with other parameters, we alternate between updating $\mathbf{U}$ given the current parameters and optimizing $(\mathbf{V}, \mathbf{w}, \sigma, m)$ given $\mathbf{U}$. This separation prevents abrupt membership changes during gradient updates, leading to a smoother objective landscape and faster convergence. Since the normalizing constant $C_{\mathbf{V}, \mathbf{w}, m, \sigma}$ in (5) generally lacks a closed form, we approximate it via importance sampling.

Here we summarize the key update equations in the iteration. Given current centroids

$\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$ and fuzziness parameter $m > 1$, the membership matrix $\mathbf{U} = (u_{ij})$ is updated by minimizing, for each i, the one–point FCM objective $\sum_{j=1}^{k} w_j u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2$ subject to $\sum_{j=1}^{k} u_{ij} = 1$. Introducing a Lagrange multiplier λ_i for the simplex constraint gives

$$\mathcal{L}_i = \sum_{j=1}^{k} w_j u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2 + \lambda_i \Big(\sum_{j=1}^{k} u_{ij} - 1 \Big).$$

The KKT condition yields $m w_j u_{ij}^{m-1} \|\mathbf{x}_i - \mathbf{v}_j\|^2 + \lambda_i = 0$, so for any j, l, we have $u_{ij}/u_{il} = \big(w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2 / w_l \|\mathbf{x}_i - \mathbf{v}_l\|^2 \big)^{-\frac{1}{m-1}}$. Normalizing over j gives the closed form

$$u_{ij} = \frac{\big(w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2 \big)^{-\frac{1}{m-1}}}{\sum_{l=1}^{k} \big(w_l \|\mathbf{x}_i - \mathbf{v}_l\|^2 \big)^{-\frac{1}{m-1}}}. \tag{9}$$

Next, with memberships fixed, the centroids minimize $\sum_{i=1}^{n} \sum_{j=1}^{k} w_j u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2$. Differentiating with respect to $\mathbf{v}_j$ yields $2 w_j \sum_i u_{ij}^m (\mathbf{v}_j - \mathbf{x}_i) = 0$, which gives the closed form

$$\mathbf{v}_j = \frac{\sum_i u_{ij}^m \mathbf{x}_i}{\sum_i u_{ij}^m}. \tag{10}$$

Then parameters $(\sigma, \mathbf{w})$ are estimated by minimizing the negative log-likelihood via L-BFGS algorithm (Liu & Nocedal, 1989):

$$\text{NLL}(\mathbf{V}, \mathbf{w}, m, \sigma) = -n \log C_{\mathbf{V}, \mathbf{w}, m, \sigma} + \sum_{i=1}^{n} \frac{1}{\sigma^2} \Big[\sum_{j=1}^{k} (w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}} \Big]^{-m+1}. \tag{11}$$

We note that the normalizing constant generally lacks a closed-form expression and is therefore approximated via importance sampling. In practice, optimizing the fuzziness parameter m jointly with $(\mathbf{V}, \mathbf{w}, \sigma)$ is challenging due to the lack of a smooth, convex objective and the dependence on m. Therefore, we estimate m via a grid search over a prespecified range and select the value that minimizes the overall negative log-likelihood. The overall iterative procedure is summarized in Algorithm 1.

Initialization. To initialize the blockwise MM iteration, we set the centroids $\mathbf{V}^{(0)}$ using standard fuzzy c-means or k-means clustering results, which provide stable starting points

in practice. The cluster weights $\mathbf{w}^{(0)}$ are initialized uniformly, $w_j^{(0)} = 1/k$, and the scale parameter $\sigma^{(0)}$ is initialized based on the empirical within-cluster variance from the preliminary clustering. This initialization ensures stable convergence and mitigates local minima issues common in fuzzy optimization.

Importance Sampling for the Normalizing Constant. In the last section, we have $C_{\mathbf{V},\mathbf{w},m,\sigma}$ as the normalizing constant ensuring that density f integrates to one. However, computing this constant analytically is intractable. We therefore approximate it by importance sampling.

To estimate $C_{\mathbf{V},\mathbf{w},m,\sigma}$, we construct a proposal distribution $q(\mathbf{x})$ as a Gaussian mixture model (GMM) fitted to the observed data. The fitted GMM captures the empirical data distribution while remaining computationally convenient for both sampling and density evaluation.

Specifically, given data $\{\mathbf{x}_i\}_{i=1}^n \subset \mathbb{R}^d$, we fit a k-component GMM:

$$q(\mathbf{x}) = \sum_{j=1}^k \pi_j \mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j),$$

where π_j are mixture weights, $\boldsymbol{\mu}_j$ are component means, and $\boldsymbol{\Sigma}_j$ are full covariance matrices estimated via maximum likelihood. We then draw Monte Carlo samples $\{\mathbf{x}_r\}_{r=1}^M$ from this fitted mixture and evaluate their proposal densities $q(\mathbf{x}_r)$.

The inverse normalizing constant is approximated as

$$C_{\mathbf{V},\mathbf{w},m,\sigma}^{-1} = \int \tilde{f}(\mathbf{x};\theta)\mathrm{d}\mathbf{x} \approx \frac{1}{M} \sum_{r=1}^M \frac{\tilde{f}(\mathbf{x}_r;\theta)}{q(\mathbf{x}_r)}, \qquad \mathbf{x}_r \sim q(\mathbf{x}).$$

where $\tilde{f}(\mathbf{x}_r;\theta) = \exp\left(-\frac{1}{\sigma^2[\sum_{j=1}^k (w_j\|\mathbf{x}_r-\mathbf{v}_j\|^2)^{-\frac{1}{m-1}}]^{m-1}}\right)$. This yields a data-adaptive proposal that efficiently covers regions of high density in $f(\mathbf{x};\theta)$, improving the stability of importance-weighted estimation. This $C_{\mathbf{V},\mathbf{w},m,\sigma}$ feeds into the negative log-likelihood as $\mathrm{NLL} = -n \log C_{\mathbf{V},\mathbf{w},m,\sigma} + \sum(\text{energy terms})$.

Overall Algorithm The proposed algorithm alternates between updating soft memberships and optimizing model parameters in a blockwise MM (Majorization-Minimization)

Algorithm 1 Blockwise MM Fuzzy Clustering with Post-MM Full NLL Refinement

---

Input: Observed data $\{\mathbf{x}_i\}_{i=1}^n$, number of clusters k, initial parameters $\sigma^{(0)}, \mathbf{V}^{(0)} = (\mathbf{v}_1^{(0)}, \ldots, \mathbf{v}_k^{(0)}), \mathbf{w}^{(0)}$, tolerances ε, δ, m-grid $\mathcal{M} = \{m_1, m_2, \ldots, m_T\}$, max iters of post MM optimization K_{post}

1: **For** $m \in \mathcal{M}$:

2:     **Initialize:** Set $t \leftarrow 0$ and $\theta^{(0,m)} \leftarrow (\sigma^{(0,m)}, \mathbf{V}^{(0,m)}, \mathbf{w}^{(0,m)})$

3:     **Repeat**

4:         $t \leftarrow t + 1$

5:         **Membership update**

$$u_{ij}^{(t,m)} \leftarrow \frac{\left(w_j^{(t-1,m)} \|\mathbf{x}_i - \mathbf{v}_j^{(t-1,m)}\|^2\right)^{-\frac{1}{m-1}}}{\sum_{l=1}^k \left(w_l^{(t-1,m)} \|\mathbf{x}_i - \mathbf{v}_l^{(t-1,m)}\|^2\right)^{-\frac{1}{m-1}}} \quad \forall i \in [n], j \in [k].$$

6:         **Centroid update**

$$\mathbf{v}_j^{(t,m)} \leftarrow \frac{\sum_i (u_{ij}^{(t,m)})^m \mathbf{x}_i}{\sum_i (u_{ij}^{(t,m)})^m} \quad \forall j \in [k].$$

7:         **Scale/weight update via L-BFGS**

$$(\sigma^{(t,m)}, \mathbf{w}^{(t,m)}) \leftarrow \underset{\sigma > 0, \, \mathbf{w} \in \mathcal{S}_k}{\arg\min} \, \text{NLL}\left(\sigma, \mathbf{V}^{(t,m)}, \mathbf{w}; m\right)$$

8:     **Until** $\|\theta^{(t,m)} - \theta^{(t-1,m)}\|_2 \leq \varepsilon$ **or** $\text{NLL}^{(t-1,m)} - \text{NLL}^{(t,m)} \leq \delta$

9:     **Post-MM optimization:**

$$\theta^{(m)} \leftarrow \underset{\theta}{\arg\min} \, \text{NLL}(\theta; m) \quad \text{start from} \quad (\sigma^{(t,m)}, \mathbf{V}^{(t,m)}, \mathbf{w}^{(t,m)})$$

10:     **return** $\theta^{(m)}, \, \text{NLL}\left(\theta^{(m)}; m\right)$

Output: $m^* = \arg\min_{m \in \mathcal{M}} \text{NLL}\left(\theta^{(m)}; m\right), \quad \theta^* = (\sigma^{(m^*)}, \mathbf{V}^{(m^*)}, \mathbf{w}^{(m^*)})$

---

fashion: (A) the membership matrix $\mathbf{U}$ is updated in closed form using the current estimates of the centroids; (B) the centroids $\mathbf{V}$ are updated using a fuzzy c-means style formula; and (C) the dispersion parameter σ and weight vector $\mathbf{w}$ (and optionally the centroids again) are optimized by minimizing the negative log-likelihood, while keeping $\mathbf{U}$ fixed. This separation prevents abrupt changes in $\mathbf{U}$ during gradient-based optimization, avoiding the jaggedness typically introduced by re-evaluating soft assignments within each iteration.

This MM separation confers two key advantages: (1) stability: holding $\mathbf{U}$ fixed during the optimization results in a smoother objective surface, which improves convergence for gradient-based methods; and (2) efficiency: iterations proceed faster and more reliably than joint updates, especially in fuzzy clustering settings where membership updates can be highly sensitive. When initialized well, the algorithm consistently converges to parameter estimates close to the underlying truth.

4.2 Statistical Inference

Once the parameters have been estimated by Algorithm 1, one can carry out hypothesis tests and construct confidence intervals in the same way as in classical likelihood inference. Here we outline two common inferential tasks: testing the equality of two cluster centers and computing confidence intervals for any component of θ.

Two-sample test. Consider the hypothesis test

$$H_0 : \mathbf{v}_a = \mathbf{v}_b \quad \text{versus} \quad H_1 : \mathbf{v}_a \neq \mathbf{v}_b,$$

for some $1 \le a < b \le k$. Under H_0 the two components are constrained to share a common center $\mathbf{v}^*$. We write the restricted parameter as

$$\tilde{\theta} = (\sigma, \mathbf{v}_1, \ldots, \mathbf{v}_{a-1}, \mathbf{v}^*, \mathbf{v}_{a+1}, \ldots, \mathbf{v}_{b-1}, \mathbf{v}^*, \mathbf{v}_{b+1}, \ldots, \mathbf{v}_k, \mathbf{w}, m).$$

Let $\hat{\theta}$ be the (unrestricted) MLE and $\tilde{\theta}$ the MLE under H_0, each obtained by running our MM procedure without and with the constraint $\mathbf{v}_a = \mathbf{v}_b$, respectively. The log-likelihoods

Algorithm 2 Inference with Blockwise MM-Like Fuzzy Clustering

Input: Estimated parameters $\hat{\theta} = (\hat{\sigma}, \hat{\mathbf{V}}, \hat{\mathbf{w}}, \hat{m})$, training data $\{x_i\}_{i=1}^n$, bootstrap replicates B, level α

1: **for** $b = 1, \ldots, B$ **do**

2:    Resample indices S_b with replacement from $\{1, \ldots, n\}$; set $X^{(b)} = \{x_i : i \in S_b\}$

3:    Refit the MM model on $X^{(b)}$ to get $\hat{\theta}^{(b)} = (\hat{\sigma}^{(b)}, \hat{\mathbf{v}}_1^{(b)}, \ldots, \hat{\mathbf{v}}_k^{(b)}, \hat{w}_1^{(b)}, \ldots, \hat{w}_k^{(b)}, \hat{m}^{(b)})$

4:    Align labels of $\hat{\theta}^{(b)}$ to $\hat{\theta}$ under label permutation

5: **end for**

6: *Scalars.* For each scalar θ_j:

$$\text{CI}(\theta_j) = \left[\text{quantile}_{\alpha/2}\{\hat{\theta}_j^{(b)}\}_{b=1}^B, \text{quantile}_{1-\alpha/2}\{\hat{\theta}_j^{(b)}\}_{b=1}^B\right].$$

7: *Vectors.* For each $a \in \{1, \ldots, k\}$ and for $\mathbf{w}$, compute $\bar{\mathbf{v}}_a = \sum_{b=1}^B \hat{\mathbf{v}}_a^{(b)}/B$, $\bar{\mathbf{w}} = \sum_{b=1}^B \hat{\mathbf{w}}^{(b)}/B$ and

$$\hat{\Sigma}_a = \frac{1}{B-1}\sum_{b=1}^B (\hat{\mathbf{v}}_a^{(b)} - \bar{\mathbf{v}}_a)(\hat{\mathbf{v}}_a^{(b)} - \bar{\mathbf{v}}_a)^\top, \quad \hat{\Lambda} = \frac{1}{B-1}\sum_{b=1}^B (\hat{\mathbf{w}}^{(b)} - \bar{\mathbf{w}})(\hat{\mathbf{w}}^{(b)} - \bar{\mathbf{w}})^\top.$$

Compute $q_{1-\alpha} = \text{quantile}_{1-\alpha}\{(\hat{\mathbf{v}}_a^{(b)} - \hat{\mathbf{v}}_a)^\top \hat{\Sigma}_a^{-1}(\hat{\mathbf{v}}_a^{(b)} - \hat{\mathbf{v}}_a)\}_{b=1}^B$ and $r_{1-\alpha} = \text{quantile}_{1-\alpha}\{(\hat{\mathbf{w}}^{(b)} - \hat{\mathbf{w}})^\top \hat{\Lambda}^{-1}(\hat{\mathbf{w}}^{(b)} - \hat{\mathbf{w}})\}_{b=1}^B$. Calculate $\mathcal{C}_{1-\alpha}(\mathbf{v}_a) = \left\{\mathbf{v} : (\mathbf{v} - \hat{\mathbf{v}}_a)^\top \hat{\Sigma}_a^{-1}(\mathbf{v} - \hat{\mathbf{v}}_a) \leq q_{1-\alpha}\right\}$ and $\mathcal{C}_{1-\alpha}(\mathbf{w}) = \left\{\mathbf{w} : (\mathbf{w} - \hat{\mathbf{w}})^\top \hat{\Lambda}^\dagger(\mathbf{w} - \hat{\mathbf{w}}) \leq r_{1-\alpha}\right\}$.

Output: Percentile CIs $\{\text{CI}(\theta_j)\}$, ellipsoidal confidence regions $\{\mathcal{C}_{1-\alpha}(\mathbf{v}_j)\}$ and $\{\mathcal{C}_{1-\alpha}(\mathbf{w})\}$.

---

are

$$\ell(\hat{\theta}) = \sum_{i=1}^n \log f(x_i; \hat{\theta}), \qquad \ell(\tilde{\theta}) = \sum_{i=1}^n \log f(x_i; \tilde{\theta}),$$

and the likelihood-ratio statistic is

$$\Lambda = -2\left\{\ell(\tilde{\theta}) - \ell(\hat{\theta})\right\}.$$

Under standard regularity conditions (Theorem 2), and up to label-switching, Λ converges in distribution to χ_p^2 as $n \to \infty$, where p is the dimension of the constraint (here $p = d$, the dimension of samples). A large value of Λ leads to rejection of H_0 at level α if $\Lambda > \chi_{p,1-\alpha}^2$.

Confidence interval construction. Now we discuss the strategy to compute the confidence interval. We quantify uncertainty via a nonparametric bootstrap. Let $\hat{\theta} = (\hat{\sigma}, \hat{\mathbf{V}}, \hat{\mathbf{w}})$ be the estimate returned by our MM routine. For $b = 1, \ldots, B$, draw a bootstrap sample by resampling $\{x_i\}_{i=1}^n$ with replacement, refit the same estimation routine on the bootstrap data and obtain $\hat{\theta}^{(b)}$.

Since the mixture components are only identifiable up to permutation, we align each replicate to $\hat{\theta}$ by solving

$$\pi^{(b)} \in \underset{\pi \in \mathfrak{S}_k}{\operatorname{argmin}} \sum_{j=1}^{k} \big\|\hat{\mathbf{v}}_{\pi(j)}^{(b)} - \hat{\mathbf{v}}_j\big\|^2,$$

and then permute $(\hat{\mathbf{V}}^{(b)}, \hat{\mathbf{w}}^{(b)})$ accordingly. For any scalar component θ_j, let $\{\hat{\theta}_j^{(b)}\}_{b=1}^B$ denote the aligned bootstrap replicates. Then a two-sided $(1 - \alpha)$ percentile interval is $\big[\hat{\theta}_j^{(\alpha/2)}, \hat{\theta}_j^{(1-\alpha/2)}\big]$, where $\hat{\theta}_j^{(q)}$ is the q-quantile of $\{\hat{\theta}_j^{(b)}\}$. For vector parameters, for instance, the clustering centers, we denote the estimated centers by $\{\hat{\mathbf{v}}_a^{(b)}\}_{b=1}^B$. We compute the bootstrap mean and covariance

$$\bar{\mathbf{v}}_a = \frac{1}{B} \sum_{b=1}^{B} \hat{\mathbf{v}}_a^{(b)}, \qquad \hat{\Sigma}_a = \frac{1}{B-1} \sum_{b=1}^{B} \big(\hat{\mathbf{v}}_a^{(b)} - \bar{\mathbf{v}}_a\big)\big(\hat{\mathbf{v}}_a^{(b)} - \bar{\mathbf{v}}_a\big)^\top.$$

The $(1 - \alpha)$ bootstrap confidence region for $\mathbf{v}_a$ is the ellipsoid

$$\mathcal{C}_{1-\alpha}(\mathbf{v}_a) = \Big\{\mathbf{v} \in \mathbb{R}^d : \big(\mathbf{v} - \hat{\mathbf{v}}_a\big)^\top \hat{\Sigma}_a^{-1}\big(\mathbf{v} - \hat{\mathbf{v}}_a\big) \leq q_{1-\alpha}\Big\},$$

where $q_{1-\alpha}$ is the empirical $(1 - \alpha)$ quantile of $\big\{\big(\hat{\mathbf{v}}_a^{(b)} - \hat{\mathbf{v}}_a\big)^\top \hat{\Sigma}_a^{-1}\big(\hat{\mathbf{v}}_a^{(b)} - \hat{\mathbf{v}}_a\big)\big\}_{b=1}^B$.

4.3 Number of Clusters via Weighted Xie-Beni Index

To determine an appropriate number of clusters in our weighted fuzzy c-means framework, we propose a modified version of the Xie–Beni index (XBI), a widely used validity measure for fuzzy clustering (X. L. Xie & Beni, 1991). The XBI evaluates both the compactness and separation of fuzzy partitions, balancing within-cluster cohesion against between-cluster separation.

Given a fuzzy partition matrix $\mathbf{U} = (u_{ij})$ of size $n \times k$, where $u_{ij} \in [0, 1]$ denotes the membership degree of data point $\mathbf{x}_i$ to cluster j, and a set of cluster centers $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$,

the classical Xie-Beni index is defined as:

$$\text{XBI}(k) = \frac{\sum_{i=1}^{n} \sum_{j=1}^{k} u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2}{\min_{j \neq l} \|\mathbf{v}_j - \mathbf{v}_l\|^2},$$

where $u_{ij} = \frac{\|\mathbf{x}_i - \mathbf{v}_j\|^{-\frac{2}{m-1}}}{\sum_{l=1}^{k} \|\mathbf{x}_i - \mathbf{v}_l\|^{-\frac{2}{m-1}}}$, and $m > 1$ is the fuzziness parameter, and $\|\mathbf{x}_i - \mathbf{v}_j\|$ is the Euclidean distance between data point $\mathbf{x}_i$ and cluster center $\mathbf{v}_j$.

A smaller XBI indicates a better clustering solution. Specifically, the numerator measures the total within-cluster variance weighted by fuzzy memberships, while the denominator captures the minimum squared distance between cluster centers, promoting well-separated clusters.

In our weighted fuzzy clustering variant, where fuzzy memberships are modified by a learned weight vector or a density-driven weighting mechanism, we adapt the XBI accordingly:

$$\text{XBI}_{\text{weighted}}(k) = \frac{\sum_{i=1}^{n} \sum_{j=1}^{k} w_j \cdot u_{ij}^m \|\mathbf{x}_i - \mathbf{v}_j\|^2}{\min_{j \neq l} \|\mathbf{v}_j - \mathbf{v}_l\|^2},$$

where $u_{ij} = \frac{1}{\sum_{l=1}^{k} \left(\frac{w_j \|\mathbf{x}_i - \mathbf{v}_j\|^2}{w_l \|\mathbf{x}_i - \mathbf{v}_l\|^2} \right)^{\frac{1}{m-1}}}$, and w_j represents the weight associated with each cluster j, derived from the underlying data distribution.

We compute $\text{XBI}_{\text{weighted}}(k)$ over a range of candidate values of k, and select the number of clusters that minimizes this index:

$$k^* = \underset{k}{\arg\min} \ \text{XBI}_{\text{weighted}}(k).$$

This provides a data-driven and robust mechanism to identify the optimal cluster resolution under our soft clustering framework.

5 Theory

In this section, we present the theoretical guarantees for our weighted fuzzy clustering method. Our focus is twofold. First, we prove that the MLE is consistent for the true parameter as the sample size grows, meaning that after relabeling indistinguishable components, it recovers the correct scale, centers, and weights. Second, we show that the estimation

error, when multiplied by $\sqrt{n}$, behaves like a centered multivariate normal with covariance given by the inverse Fisher information, yielding the usual Wald approximations. Since the components in the density are exchangeable, all statements are made up to label switching. Throughout, we assume that the parameter lies in a compact set and the density is the normalized form used for estimation.

Model and parameter space. Fix an integer $k \geq 2$ and let the parameter space be

$$\Theta = \left[\sigma_{\min}, \sigma_{\max}\right] \times \left\{ w \in [\varepsilon_w, 1]^k : \sum_{j=1}^{k} w_j = 1 \right\} \times \mathcal{V}^k \times \left[m_{\min}, m_{\max}\right],$$

where $0 < \sigma_{\min} < \sigma_{\max} < \infty$, $0 < \varepsilon_w < 1/k$, $\mathcal{V} \subset \mathbb{R}^d$ is compact, and $1 < m_{\min} < m_{\max} < \infty$. For $\theta = (\sigma, \mathbf{w}, \mathbf{V}, m) \in \Theta$ write

$$f_\theta(\mathbf{x}) = C(\theta) \exp\left(-\sigma^{-2} \left[\Sigma_w(\mathbf{x})\right]^{-(m-1)}\right), \qquad \Sigma_w(\mathbf{x}) = \sum_{j=1}^{k} w_j^{-\frac{1}{m-1}} \|\mathbf{x} - \mathbf{v}_j\|^{-\frac{2}{m-1}}.$$

Theorem 1 *Let $\mathbf{x}_1, \ldots, \mathbf{x}_n$ be i.i.d. draws from the true density f_{θ_0} with $\theta_0 \in \mathrm{int}(\Theta)$. Define the MLE $\hat{\theta}_n = (\hat{\sigma}_n, \hat{\mathbf{w}}_n, \hat{\mathbf{V}}_n, \hat{m}_n) \in \mathrm{argmax}_{\theta \in \Theta} \sum_{i=1}^{n} \log f_\theta(\mathbf{x}_i)$. Then, almost surely,*

$$\hat{\theta}_n \longrightarrow \theta_0 \quad \text{as } n \to \infty$$

up to a permutation of the k component labels, namely, there exists a random permutation π_n such that $(\sigma_n, m_n) \to (\sigma_0, m_0)$ and $(\mathbf{v}_{n,j}, w_{n,j}) \to (\mathbf{v}_{0,\pi_n(j)}, w_{0,\pi_n(j)})$ for every $j = 1, \ldots, k$.

Theorem 1 says that the MLE is strongly consistent, except for the inherent ambiguity of labeling mixture components. Intuitively, the average log–likelihood converges uniformly to its population counterpart, which is maximized at the true parameter (modulo permutations). The compact parameter set and the positivity bound on the weights keep the optimizer from drifting to degenerate configurations (e.g., vanishing components or centers escaping the domain), thereby ensuring convergence of both the scale (σ, m) and the component–specific quantities $(\mathbf{v}_j, w_j)$ after relabeling.

Theorem 2 *Under the regularity conditions verified above, there exists a sequence of proper label permutations π_n such that the maximum-likelihood estimator associated with the normalized fuzzy-clustering density satisfies*

$$\sqrt{n}\big(\pi_n(\hat{\theta}_n) - \theta_0\big) \xrightarrow{d} \mathcal{N}\big(0, I(\theta_0)^{\dagger}\big), \qquad n \to \infty,$$

where $I(\theta_0)$ is the Fisher information matrix, $I(\theta_0)^{\dagger}$ denotes its Moore–Penrose pseudoinverse in the ambient parametrization, and the resulting Gaussian limit is supported on the tangent space of the simplex constraint $\sum_{j=1}^{k} w_j = 1$.

Theorem 2 provides the $\sqrt{n}$–Gaussian approximation for the MLE. After resolving label switching (e.g., by matching estimated centers to the truth via minimal total distance), the score admits a mean–zero central limit theorem, and the observed Hessian converges to the Fisher information, yielding a quadratic approximation of the log–likelihood around θ_0. Practically, this result underpins Wald–type confidence intervals and ellipsoids for smooth reparametrizations of $(\sigma, \mathbf{w}, \mathbf{V}, m)$, and it justifies our whitened QQ diagnostics in Section 6.

The lower bound $\varepsilon_w > 0$ on the component weights w_j is essential to ensure identifiability of the full parameter vector. If $w_j = 0$ for some j, then the corresponding center v_j becomes non-identifiable, since its contribution to the density $f_\theta(x)$ vanishes. This leads to a flat likelihood in the v_j direction and invalidates consistency and asymptotic normality results. By enforcing $w_j \geq \varepsilon_w > 0$, each component is guaranteed to have non-negligible influence on the model, which is necessary for consistent estimation of both weights and locations.

6 Simulation

We now conduct a suite of simulation experiments to investigate the finite-sample and asymptotic properties of the proposed MM fuzzy clustering estimator. Specifically, we aim to (i) verify finite-sample consistency, (ii) assess the empirical distribution of the estimator under increasing sample size, and (iii) compare different approaches to uncertainty quantification,

including empirical replicates, bootstrap resampling, and Fisher-information-based approximations.

Data-Generating Mechanism. We focus on a setting with three clusters in $\mathbb{R}^3$. The true parameter is specified as

$$\theta_0 = \big(\sigma_0 = 2, \mathbf{v}_{0,1} = (0,0,0), \mathbf{v}_{0,2} = (20,0,-1), \mathbf{v}_{0,3} = (-20, 2.5, 1), \mathbf{w}_0 = (0.3, 0.1, 0.6), m_0 = 2\big).$$

For each Monte Carlo replicate and each sample size n, we generate n observations from f_{θ_0} using a Metropolis–Hastings sampler with a mixture-scale proposal kernel. The sampler mixes local Gaussian moves with occasional large jumps (a scale matched to the typical distance between cluster centers), which improves exploration across multiple modes. A burn-in period of 20% is discarded to ensure stationarity.

Estimation Procedure. Given a dataset of size n, we proceed in three steps:

1. Fit a Gaussian mixture model with 2–6 components to the data and generate $M = 5000$ proposal samples. This proposal is used for importance sampling in the approximate likelihood evaluation.

2. Run the MM fuzzy clustering algorithm (Algorithm 1) to obtain parameter estimates

$$\hat{\theta}^{(b)} = (\hat{\sigma}, \hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_k, \hat{\mathbf{w}}).$$

3. Resolve label switching by aligning the estimated centers $\hat{\mathbf{v}}_j$ with the true centers $\mathbf{v}_{0,j}$. We adopt the Hungarian assignment method, which minimizes the total Euclidean distance between estimated and true centers. The estimated weights $\hat{\mathbf{w}}$ are permuted consistently with the center assignment.

This alignment step is essential because the underlying density is invariant under permutations of the cluster labels, and therefore the likelihood does not distinguish among labelings. See the discussion following Theorem 2 for further explanation.
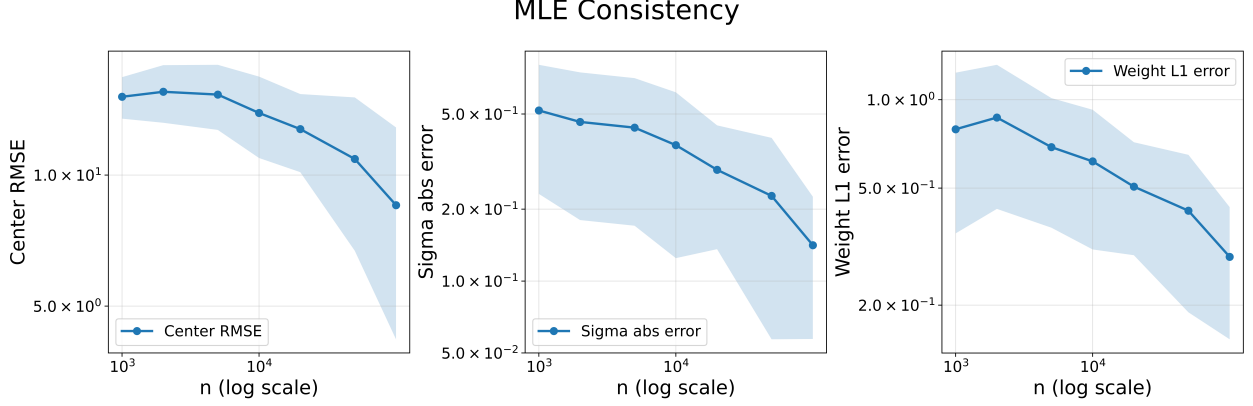
Figure 3: Finite-sample consistency of the MM fuzzy clustering estimator. Errors are shown as functions of sample size n on a log–log scale.

Consistency. To verify consistency, we compute the label-invariant errors between the estimates and the truth across sample sizes $n \in \{500, 1000, 2000, 5000, 10000, 20000\}$. Figure 3 plots the empirical root mean squared error of the centers, the absolute error of $\hat{\sigma}$, and the ℓ_1 error of the weights on a log–log scale. In each case, the empirical error decreases at approximately the $n^{-1/2}$ rate, consistent with the theoretical guarantees in Section 5. Confidence intervals are constructed across independent Monte Carlo replicates.

Asymptotic Normality. We next investigate the large-sample distributional behavior of the proposed estimator. The synthetic data are generated under the following setting: we fix the number of clusters $k = 2$, fuzziness parameter $m = 2.0$, and noise variance $\sigma_0^2 = 4$. The two true cluster centers are

$$\mathbf{v}_1 = (0, 0), \quad \mathbf{v}_2 = (3.5, 3.5),$$

and cluster weights are set to $\mathbf{w} = (0.8, 0.2)$. Given these parameters, we simulate approximately independent draws by running a Metropolis–Hastings sampler targeting the cluster-based density (with Gaussian proposal standard deviation 1.2, 20,000 iterations per chain, and a burn-in of 20%). For each sample size $n \in \{100, 200, 500, 1000, 2000\}$, we generate 100 independent datasets from the data-generating mechanism described above, each
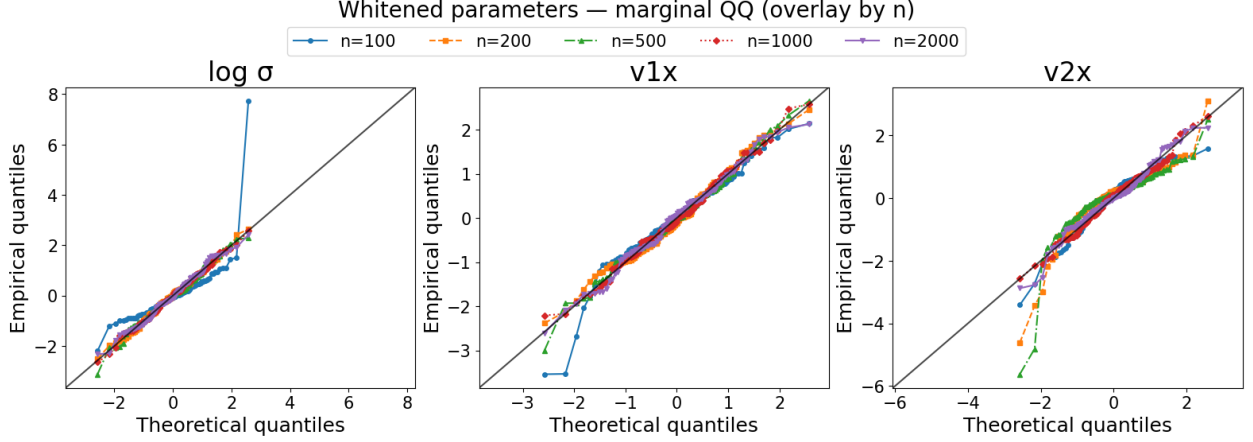
24

Whitened parameters — marginal QQ (overlay by n)

Figure 4: Marginal QQ plots of whitened parameter estimates across $n \in \{100, 200, 500, 1000, 2000\}$. Each subplot corresponds to a selected coordinate of the parameter vector. The closer the empirical quantiles lie to the $45°$ line, the stronger the agreement with the Gaussian benchmark.

using a different random seed. On each replicate we compute the aligned estimator, where label switching is resolved as in the consistency experiment.

To assess asymptotic normality, we perform the following normalization procedure. First, for each n, we center the estimated parameter vectors at the truth θ_0. Second, we rescale the centered estimates using the empirical covariance matrix across replicates, thereby whitening their joint distribution[1]. This procedure ensures that if Theorem 2 holds, the whitened estimates should approximate a multivariate standard normal distribution.

Figure 4 displays marginal quantile–quantile (QQ) plots for selected coordinates of the whitened parameter estimates, overlaid across the three sample sizes. Each subplot compares the empirical quantiles to the theoretical standard normal quantiles. For small samples ($n = 500$), visible deviations from normality remain, especially in the tails. As the sample size increases to $n = 2000$, the empirical points align increasingly closely with the $45°$ line, indicating convergence toward Gaussianity. These diagnostics confirm that the distribution of the proposed estimator approaches the Gaussian limit as predicted by Theorem 2. This

---

[1]For weight $\mathbf{w} \in \mathbb{R}^k$, we whiten in $k-1$ free coordinates (or tangent-space projection).
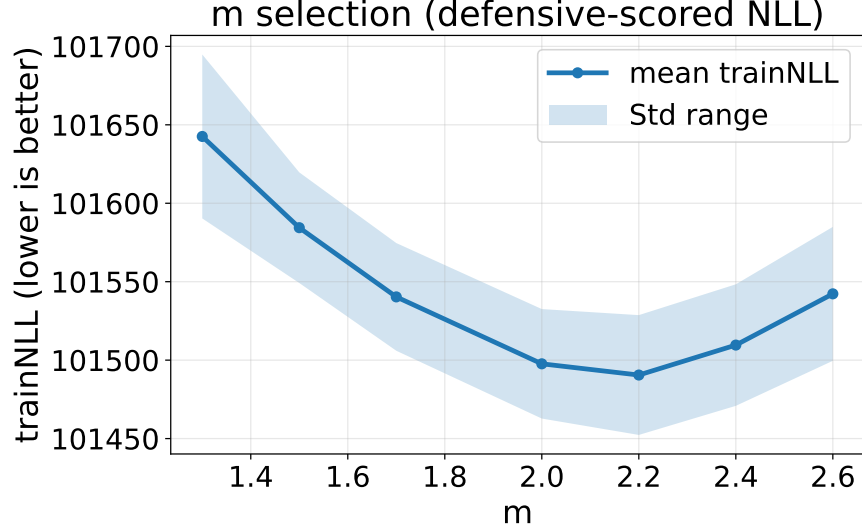
Figure 5: Selection of the fuzziness parameter m via approximate trainNLL. For each $m \in \{1.3, 1.5, 1.7, 2.0, 2.2, 2.4, 2.6\}$ we fit $(\sigma, \mathbf{V}, w)$ using the blockwise MM routine (50 iterations), and select the best m using the negative log-likelihood.

agreement provides empirical support for the validity of using asymptotic normal approximations to construct confidence intervals and hypothesis tests in practice.

Estimating the fuzziness parameter m. In addition to estimating $(\sigma, \mathbf{V}, \mathbf{w})$, we also propose a procedure that estimates the parameter m from data generated under the mechanism (three clusters in $\mathbb{R}^3$, $\sigma_0 = 2.0$, $\mathbf{v}_{0,1} = (0, 0, 0)$, $\mathbf{v}_{0,2} = (10, 0, -1)$, $\mathbf{v}_{0,3} = (-10, 2.5, 1)$, $\mathbf{w}_0 = (0.3, 0.1, 0.6)$, and $m_0 = 2$). For each dataset, we perform a grid search over

$$\mathcal{M} = \{1.3, 1.5, 1.7, 2.0, 2.2, 2.4, 2.6\}.$$

For a fixed $m \in \mathcal{M}$, parameters $(\sigma, \mathbf{V}, \mathbf{w})$ are estimated by Algorithm 1 used elsewhere in this section.

To evaluate the likelihood for model comparison across m, we use importance sampling with a data-adaptive proposal. We refer the reader to Section 4.1 for detailed explanations. Concretely, for each dataset we fit a 3-component Gaussian mixture to the observed $\mathbf{x}_{1:n}$ and draw $M = 20000$ proposal points. The same proposal is reused across all $m \in \mathcal{M}$ to avoid

favoring any single value. The training negative log-likelihood (trainNLL) for m is then computed with these shared proposal samples. To assess stability with respect to stochastic components of the routine, we repeat the entire m-grid evaluation K times with distinct random seeds and summarize the results by the mean trainNLL (with ± 1 standard-deviation bands) as a function of m. Figure 5 displays these curves, where lower values indicate better fit, and the error ribbons reflect cross-run variability of the approximate likelihood under a fixed dataset. This trainNLL-based selection is used only to choose m; once $\hat{m}$ is fixed, we refit the model at $m = \hat{m}$ using the same MM procedure for final parameter estimation and uncertainty quantification.

7 Real Data Analysis

7.1 Single-Cell RNA Sequencing Data

Single-cell RNA sequencing (scRNA-seq) data captures gene expression profiles at the resolution of individual cells, enabling researchers to explore cellular heterogeneity. A common analysis strategy involves clustering cells to infer putative cell types, followed by differential gene expression analysis between the resulting groups. However, many standard workflows rely on point-estimate cluster assignments and do not explicitly propagate clustering uncertainty to downstream analyses (Lähnemann et al., 2020; Luecken & Theis, 2019). In our analysis, we utilize a scRNA-seq dataset from peripheral blood mononuclear cells (PBMCs), originally pre-classified prior to sequencing by Zheng et al. (2017), to illustrate how our proposed soft-clustering inference framework can provide uncertainty quantification for clustering outputs while properly accounting for the fuzziness in cluster assignments.

7.1.1 Data Pre-processing

For our analysis, we focus on a subset of peripheral blood mononuclear cells consisting of memory T cells, B cells, and monocytes. Following standard single-cell RNA-seq pre-processing procedures (Duò, Robinson, and Soneson, 2018), we filter out cells with a high

proportion of mitochondrial gene expression, cells with an unusually low or high number of expressed genes, and those with low total UMI counts. To account for differences in sequencing depth/library size, we normalize the data so that each cell has the same total count equal to the average across all cells. We then apply a log-transformation of the normalized counts using $\log_2(1 + \text{count})$. To reduce dimensionality while retaining informative features, we select the top 500 genes with the highest average expression prior to normalization. This preprocessing pipeline is applied to the combined set of memory T cells, B cells, and monocytes. We construct an imbalanced dataset to evaluate our soft-clustering inference framework, consisting of 1,500 memory T cells, 1,000 B cells, and 500 monocytes, for a total of 3,000 cells. We emphasize that memory T cells, B cells, and monocytes are well-separated lineages within peripheral blood mononuclear cells and are not expected to exhibit biologically meaningful overlap (a single cell cannot simultaneously be both a T cell and a B cell, or both a lymphocyte and a monocyte). As such, this setting provides a useful baseline: soft clustering is not intended to uncover hybrid cell identities here, but rather to model uncertainty in assignments arising from technical noise, biological variability, and class imbalance. In particular, our framework allows us to quantify assignment confidence, even in cases where the underlying clusters are expected to be largely distinct.

7.1.2   Data Analysis

Figure 6 displays joint confidence ellipses for the estimated cluster centers, which capture the joint variability across bootstrap resamples. The relatively small and non-overlapping ellipses suggest that the three centers are well separated and estimated with reasonable precision. Table 1 demonstrates bootstrap summary statistics and 95% confidence intervals for the parameters, including the cluster centers (v_j^x, v_j^y), the scale parameter σ, and the weights w_j. The narrow intervals for the centers indicate stable estimation, while the wider intervals for the weights reflect the inherent imbalance in sample sizes across the three cell types.

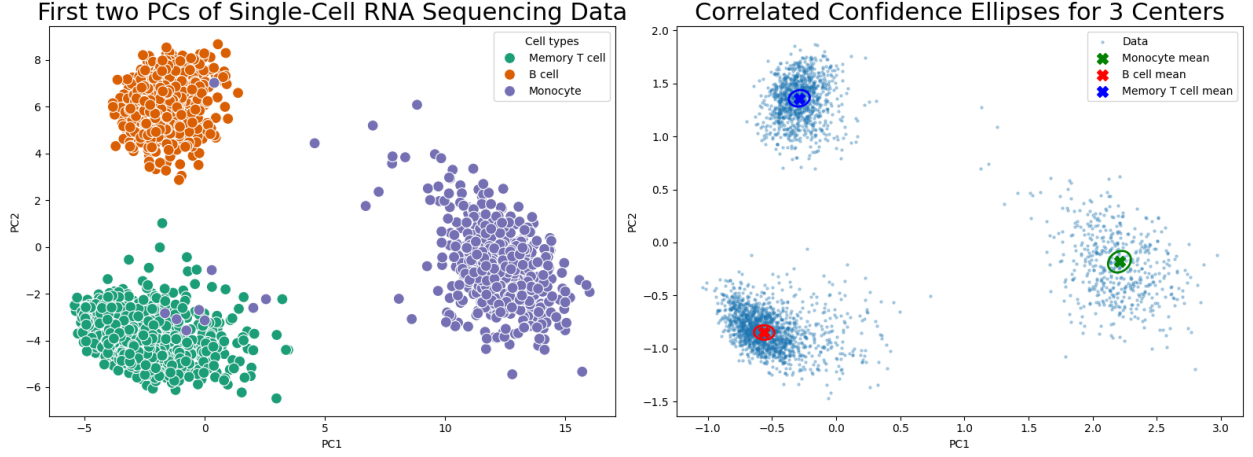We also evaluate the choice of the fuzziness parameter m using the weighted Xie–Beni

Figure 6: Left: first two PCs of Single-Cell RNA-Seq data consisting of memory T cells, B cells, and monocytes. Right: The estimated centers with 95% confidence intervals.

Table 1: scRNA-seq Bootstrap Summary Statistics and 95% Confidence Intervals

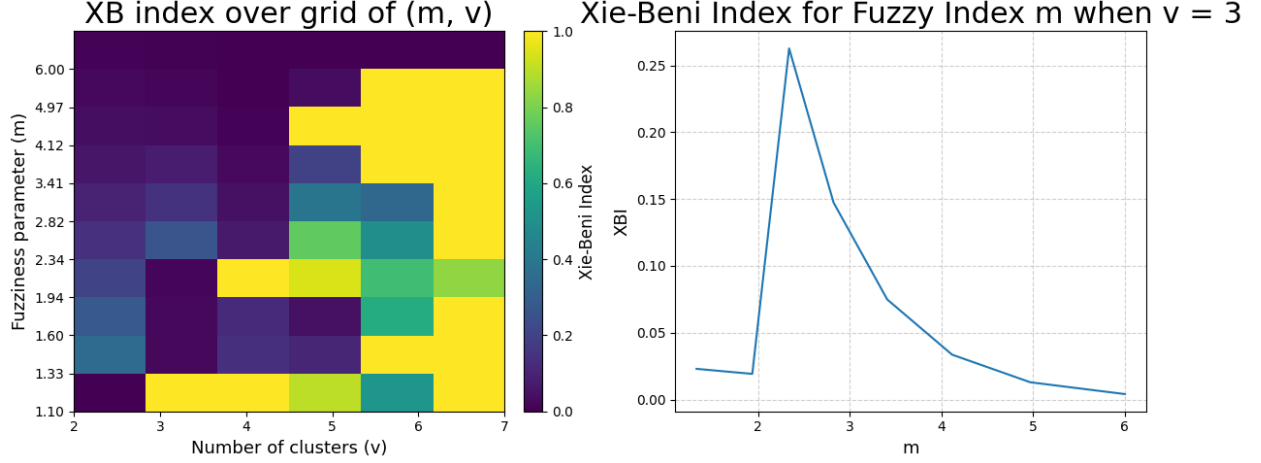| Parameter	Mean ± Std.	95% CI
σ	0.133 ± 0.009	$[0.115, \ 0.151]$
v_1^x	2.209 ± 0.036	$[2.138, \ 2.280]$
v_1^y	-0.181 ± 0.041	$[-0.262, \ -0.100]$
v_2^x	-0.286 ± 0.033	$[-0.351, \ -0.222]$
v_2^y	1.360 ± 0.031	$[1.299, \ 1.422]$
v_3^x	-0.560 ± 0.032	$[-0.624, \ -0.497]$
v_3^y	-0.849 ± 0.028	$[-0.903, \ -0.795]$
w_1	0.355 ± 0.089	$[0.181, \ 0.529]$
w_2	0.362 ± 0.074	$[0.217, \ 0.507]$
w_3	0.283 ± 0.069	$[0.147, \ 0.420]$

Figure 7: Overall grid of Weighted XBI based on number of clusters c and fuzziness parameter m. Right: Weighted XBI based on different fuzziness parameter m in three cluster scenario.

index (XBI), which balances cluster compactness and separation. By definition, XBI is always nonnegative, with smaller values indicating more compact and well-separated clusters. Its theoretical lower bound is 0, which would occur only if all points coincided exactly with their cluster centers and the centers were infinitely far apart, an impossible scenario in real data. Consequently, we should not select clusters based solely on the absolute minimum XBI to avoid degenerate or unrealistic configurations. Instead, the "elbow" region where the XBI remains relatively stable before increasing sharply provides a more practical criterion. Extremely small XBI values (e.g., $< 10^{-3}$) should be treated cautiously, particularly if the resulting clusters lack interpretability. We combine the XBI metric as shown on the left side of Figure 7 with inspection of the membership distributions and cluster centers to ensure a meaningful structure and effective cluster selection. On the right side of Figure 7, the index attains its minimum at a relatively small $m = 1.94$ when the number of clusters is 3. This suggests that the data are best represented with limited fuzziness, consistent with the fact that memory T cells, B cells, and monocytes form distinct transcriptional clusters with minimal overlap.

Figure 8 shows the estimated cluster memberships on the three-dimensional simplex. As expected given the biological distinctness of memory T cells, B cells, and monocytes, most
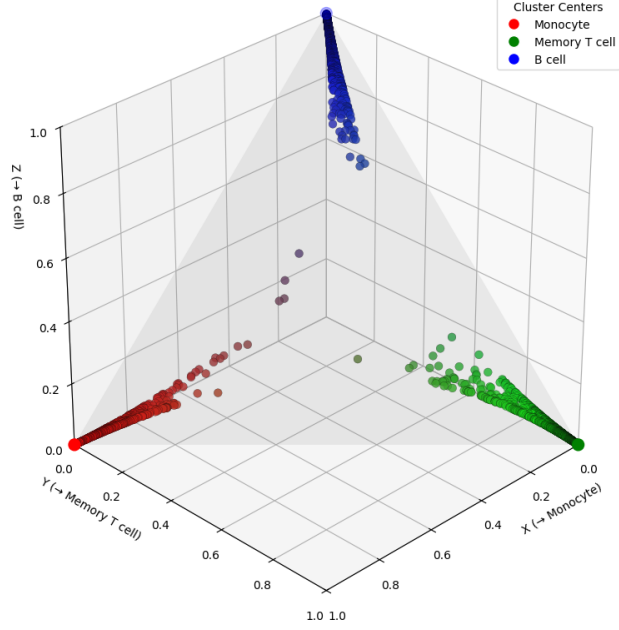
Figure 8: Estimated cluster memberships on the 3-vertex simplex

cells concentrate near the simplex vertices, reflecting near-crisp memberships; however, the soft framework still quantifies assignment probabilities and highlights the few borderline cases.

We compared the proposed WFCM with a GMM in Figure 9. Both methods achieved nearly identical cluster structures, as indicated by a high Adjusted Rand Index (ARI = 0.985), where ARI was computed by first assigning each observation to the cluster with the highest posterior membership. However, WFCM achieved this result without assuming Gaussianity or estimating covariance matrices, using only distance-based updates and the fuzziness parameter m to control softness. In contrast, GMM required full-covariance estimation and likelihood optimization (BIC = 4132.57). This highlights WFCM's light and computationally stable nature: it provides soft clustering behavior similar to GMM, but through a purely geometric criterion that scales efficiently to high-dimensional or irregularly shaped data.

From a biological perspective, these results confirm that the three immune cell populations are transcriptionally distinct, but they also demonstrate the value of our framework
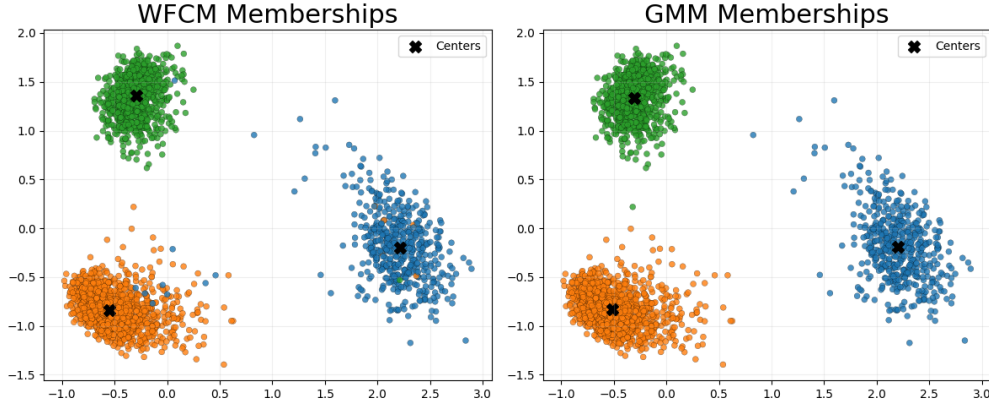
31

Figure 9: Cluster Membership Comparison between WFCM and GMM

that even in a setting where sharp boundaries exist, soft clustering provides principled measures of uncertainty in both cluster center estimation and cell-level membership assignment. This ability to quantify confidence is essential in more complex datasets where cell states transition along continua and partial memberships are biologically meaningful.

7.2 Alzheimer's Disease Neuroimaging Data

The Alzheimer's Disease Neuroimaging Initiative (ADNI) is a longitudinal multicenter study launched in 2004 to investigate biomarkers of Alzheimer's disease (AD) progression. The study enrolls cognitively normal (CN) older adults, individuals with mild cognitive impairment (MCI), and patients with AD, and follows them with repeated neuroimaging, genetic, cognitive, and clinical assessments. ADNI is widely used in methodological and clinical research because it provides a rich multimodal dataset that captures the disease continuum from normal aging to dementia. For our analysis, we focused on baseline visits to obtain a single cross-sectional feature set per subject. If multiple baseline records existed, we retained the earliest examination. This yielded a cohort of 2436 subjects.

7.2.1 Data Pre-processing

For each subject, we selected the baseline visit (VISCODE = "bl"). If multiple baseline records were available, we kept the earliest examination date. We extracted a set of de-

mographic, genetic, cognitive, and clinical variables that are commonly associated with Alzheimer's disease progression: demographics data such as age, years of education; genetic feature such as APOE4 carrier status; cognitive and functional measures such as Clinical Dementia Rating–Sum of Boxes (CDR-SB), Alzheimer's Disease Assessment Scale (ADAS11, ADAS13, ADASQ4), Mini-Mental State Exam (MMSE), Rey Auditory Verbal Learning Test (RAVLT percent forgetting, LDELTOTAL), Trail-Making Test Part B (TRABSCOR), and Functional Activities Questionnaire (FAQ); imaging identifier IMAGEUID.

Missing values were handled by median imputation for numeric variables. To ensure comparability across features, all numeric variables were standardized to mean 0 and unit variance. For evaluation against clinical diagnosis, we used ADNI diagnostic labels (DX). We relabeled dementia cases as "AD" and grouped cognitively normal (CN) and mild cognitive impairment (MCI) cases together as "non-AD."

7.2.2 Data Analysis

We projected ADNI data onto the first three principal components and compared two perspectives as shown in Figure 10. On the right, the hard labels split subjects into AD versus non-AD, creating a rigid boundary. But notice how the two groups overlap heavily—many individuals diagnosed as non-AD lie close to the AD cluster, and vice versa.

On the left, our WFCM approach produces a soft spectrum of AD membership. Here, instead of forcing each person into a binary label, we quantify the degree of belonging to the AD-like cluster. What emerges is a smooth gradient from healthy to diseased states. Importantly, this intermediate zone may represent prodromal or at-risk individuals, such as those with MCI. The FCM memberships form a smooth gradient from blue (non-AD-like) to red (AD-like). This suggests that subjects do not fall into strictly separated groups; instead, they lie on a continuum of disease burden or progression, consistent with how cognitive decline actually develops in Alzheimer's disease. The soft assignments identify a "spectrum zone" where individuals have intermediate membership values (e.g., 0.3–0.5). These may correspond to MCI subjects or clinically ambiguous cases where people who are biologically
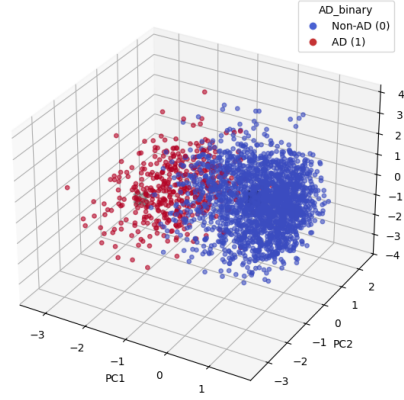
33
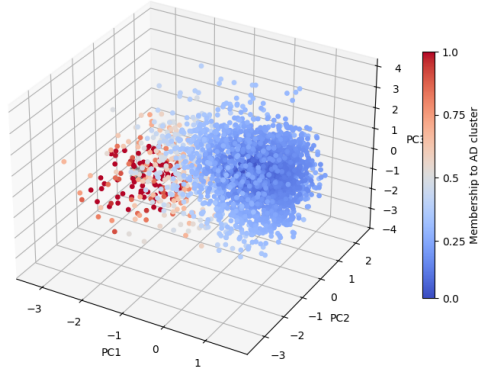
Figure 10: Left: membership weights from WFCM, representing the membership or degree of belonging to the "AD-like" cluster. Right: the binary clinical labels (AD vs. non-AD).

transitioning toward AD but not yet fully diagnosed.

Table 2 summarizes the bootstrap distribution of parameter estimates. The results show that variance estimates are relatively stable, while some cluster center coordinates and weights exhibit wider confidence intervals, reflecting greater uncertainty in those directions.

Table 2: ADNI Bootstrap Summary Statistics with 95% Confidence Intervals

Parameter	Mean $\pm$ Std.	95% CI
σ	0.438 ± 0.107	$[0.229, 0.647]$
v_1^x	0.725 ± 0.566	$[-0.384, 1.834]$
v_1^y	0.008 ± 0.280	$[-0.542, 0.557]$
v_1^z	-0.163 ± 0.508	$[-1.160, 0.833]$
v_2^x	-1.816 ± 1.046	$[-3.866, 0.234]$
v_2^y	-0.132 ± 0.469	$[-1.051, 0.787]$
v_2^z	-0.187 ± 0.662	$[-1.485, 1.111]$
w_1	0.694 ± 0.175	$[0.351, 1.000]$
w_2	0.306 ± 0.175	$[0,000, 0.649]$

Each point in the left plot of Figure 11 represents a subject, positioned along the line

segment according to their estimated membership in the two clusters: Alzheimer's disease (AD, red) and non-AD (blue). Points near the red end have high estimated membership in the AD cluster, points near the blue end have high membership in the non-AD cluster, and points in between represent subjects with intermediate or uncertain cluster assignments. This visualization illustrates the probabilistic assignments produced by soft clustering while clearly separating the two groups.

To further illustrate the soft clustering results, we project the estimated memberships for each ADNI subject onto a 2-vertex simplex, which in the case of two clusters reduces to a line segment between the Alzheimer's disease (AD) and non-AD corners (right plot of Figure 11). Each subject is represented as a point whose position reflects its estimated cluster membership. Subjects near the red corner have a high degree of membership in the AD cluster, while those near the blue corner are strongly associated with the non-AD cluster. Points located in the middle of the line indicate subjects with intermediate or uncertain cluster assignments, a pattern consistent with Mild Cognitive Impairment (MCI), reflecting the probabilistic nature of soft clustering. This representation provides a clear visualization of how soft clustering captures heterogeneity in the ADNI cohort while still distinguishing between the two major disease groups.

We compared the proposed WFCM with a Gaussian mixture model (GMM) in Figure 12. While both methods produced meaningful cluster structures, WFCM more accurately captured the biologically relevant AD/non-AD separation, as confirmed by the clinical diagnosis labels. Specifically, WFCM achieved an adjusted Rand index (ARI) of 0.535 and a classification accuracy of 88%, compared to GMM's ARI of 0.230 and accuracy of 74%, indicating substantially better alignment with the underlying disease status. Notably, WFCM attained this performance without assuming Gaussian distributions or estimating covariance matrices, relying instead on distance-based updates and the fuzziness parameter m to model partial membership. In contrast, GMM required full-covariance estimation and likelihood-based optimization (BIC = 20015.74). Overall, these results highlight the assumption-light and robust nature of WFCM, demonstrating its ability to recover biologically meaningful soft
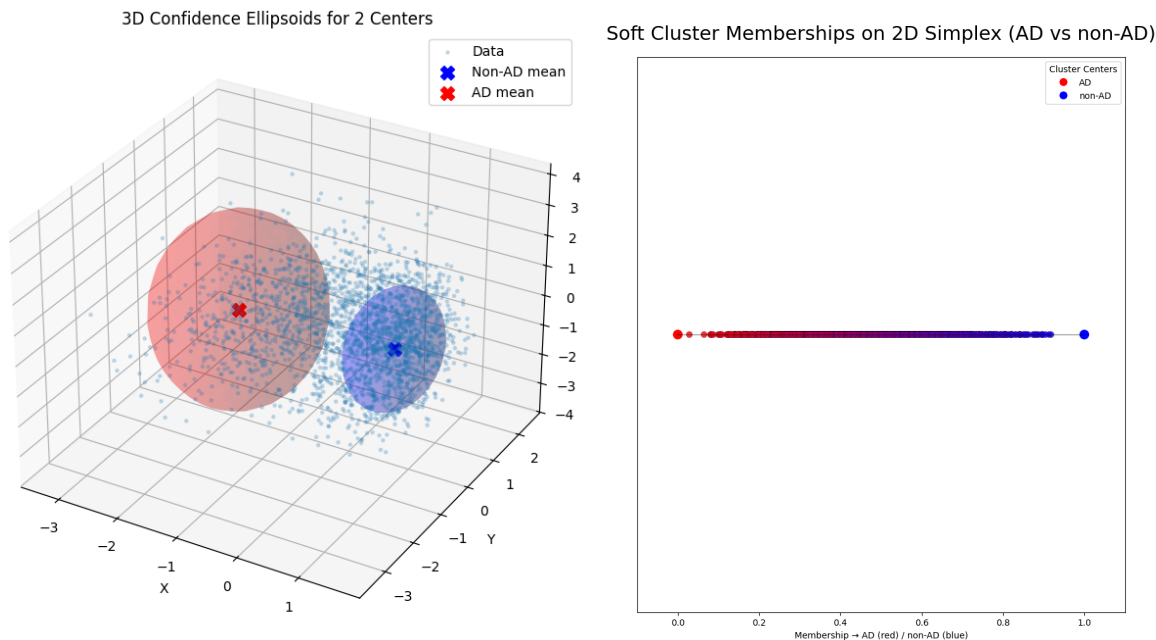
Figure 11: Left: The estimated centers with 95% confidence regions. Right: Estimated cluster memberships for ADNI subjects on the 2-vertex simplex.

cluster structures in high-dimensional, heterogeneous neuroimaging data.

By showing that AD is not well described as a dichotomy but rather as a continuum, this suggests that soft clustering can provide an early-warning signal or risk stratification score by assigning probabilistic risk even before a firm diagnosis is possible.

8 Discussion

The weighted fuzzy c-means framework proposed in this work offers a principled approach to incorporating heterogeneity into soft clustering. By formulating estimation as a block-wise MM procedure, the resulting algorithm is both interpretable and scalable. While the membership updates retain the familiar structure of classical fuzzy c-means, the introduction of weights allows the method to adapt to imbalanced or heterogeneous data, yielding a likelihood-based probabilistic interpretation.

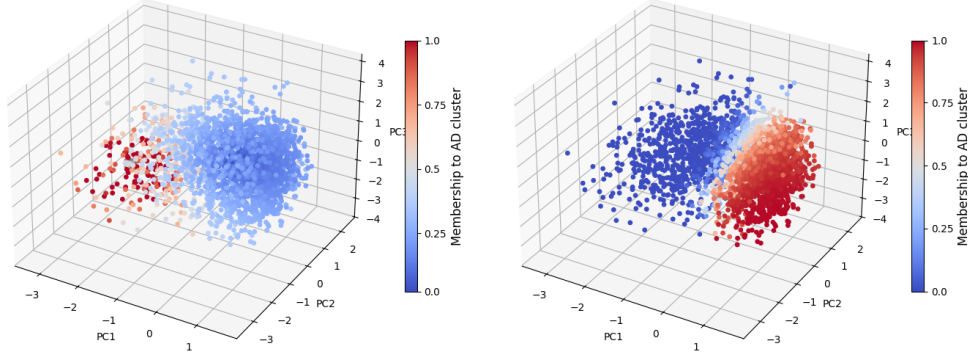Simulation studies provide empirical support for the theoretical results. Estimation er-

Figure 12: Soft AD-like cluster memberships estimated by WFCM and GMM

rors decay at approximately the $n^{-1/2}$ rate, and diagnostics of whitened estimates indicate convergence toward Gaussian limits, supporting the use of asymptotic normal approximations. Together, these findings show that the proposed estimator is consistent and amenable to uncertainty quantification through both bootstrap and Fisher-information-based methods. The real-data analyses illustrate complementary strengths of the framework. In the scRNA-seq study, the method yields well-separated clusters with tight confidence regions, while in the ADNI application it captures a continuum between Alzheimer's disease and non-AD subjects through soft membership assignments. These results highlight the value of soft clustering for quantifying uncertainty and identifying intermediate or at-risk cases in applied settings.

Several limitations warrant discussion. While the weighting mechanism increases flexibility, the interpretation of learned weights is context dependent, and more work is needed to develop general guidance for their use. In addition, although the blockwise MM-type algorithm is computationally efficient, it remains iterative and can converge to local optima; careful initialization is therefore important in practice. Future work may extend the framework by incorporating structured prior information (e.g., spatial structure, network constraints, or biological pathways) into the weighting scheme, and by developing dynamic or longitudinal extensions where memberships evolve over time.

References

Ball, G. H., & Hall, D. J. (1965). Isodata, a novel method of data analysis and pattern classification.

Banerjee, A., Krumpelman, C., Ghosh, J., Basu, S., & Mooney, R. J. (2005). Model-based overlapping clustering. In *Proceedings of the eleventh acm sigkdd international conference on knowledge discovery in data mining* (pp. 532–537).

Basu, S., Davidson, I., & Wagstaff, K. (2008). *Constrained clustering: Advances in algorithms, theory, and applications.* Chapman and Hall/CRC.

Bezdek, J. C. (2013). *Pattern recognition with fuzzy objective function algorithms.* Springer Science & Business Media.

Bezdek, J. C., Ehrlich, R., & Full, W. (1984). Fcm: The fuzzy c-means clustering algorithm. *Computers & geosciences*, *10*(2-3), 191–203.

Bullen, P. S. (2013). *Handbook of means and their inequalities* (Vol. 560). Springer Science & Business Media.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, *39*(1), 1–22.

Duda, R. O., Hart, P. E., et al. (2006). *Pattern classification.* John Wiley & Sons.

Fan, J.-L., Zhen, W.-Z., & Xie, W.-X. (2003). Suppressed fuzzy c-means clustering algorithm. *Pattern Recognition Letters*, *24*(9-10), 1607–1612.

Gao, L. L., Bien, J., & Witten, D. (2024). Selective inference for hierarchical clustering. *Journal of the American Statistical Association*, *119*(545), 332–342.

Gath, I., & Geva, A. B. (2002). Unsupervised optimal fuzzy clustering. *IEEE Transactions on pattern analysis and machine intelligence*, *11*(7), 773–780.

Gustafson, D. E., & Kessel, W. C. (1979). Fuzzy clustering with a fuzzy covariance matrix. In *1978 ieee conference on decision and control including the 17th symposium on adaptive processes* (pp. 761–766).

Hathaway, R. J., & Bezdek, J. C. (2006). Extending fuzzy and probabilistic clustering to very large data sets. *Computational Statistics & Data Analysis*, *51*(1), 215–234.

Jain, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern recognition letters*, *31*(8), 651–666.

Jain, A. K., & Dubes, R. C. (1988). *Algorithms for clustering data*. Prentice-Hall, Inc.

Ji, Z.-X., Sun, Q.-S., & Xia, D.-S. (2011). A modified possibilistic fuzzy c-means clustering algorithm for bias field estimation and segmentation of brain mr image. *Computerized Medical Imaging and Graphics*, *35*(5), 383–397.

Jiang, S., Gai, X., Treggiari, M. M., Stead, W. W., Zhao, Y., Page, C. D., & Zhang, A. R. (2024). Soft phenotyping for sepsis via ehr time-aware soft clustering. *Journal of biomedical informatics*, *152*, 104615.

Lähnemann, D., Köster, J., Szczurek, E., McCarthy, D. J., Hicks, S. C., Robinson, M. D., ... others (2020). Eleven grand challenges in single-cell data science. *Genome Biology*, *21*, 31. doi: 10.1186/s13059-020-1926-6

Liu, D. C., & Nocedal, J. (1989). On the limited memory bfgs method for large scale optimization. *Mathematical programming*, *45*(1), 503–528.

Lloyd, S. (1982). Least squares quantization in pcm. *IEEE transactions on information theory*, *28*(2), 129–137.

Luecken, M. D., & Theis, F. J. (2019). Current best practices in single-cell RNA-seq analysis: a tutorial. *Molecular Systems Biology*, *15*(6), e8746. doi: 10.15252/msb.20188746

McLachlan, G. J., & Peel, D. (2000). *Finite mixture models*. John Wiley & Sons.

Meilă, M. (2006). The uniqueness of a good optimum for k-means. In *Proceedings of the 23rd international conference on machine learning* (pp. 625–632).

Pham, D. L., & Prince, J. L. (1999). An adaptive fuzzy c-means algorithm for image segmentation in the presence of intensity inhomogeneities. *Pattern recognition letters*, *20*(1), 57–68.

Sokal, R. R., & Sneath, P. H. A. (1963). *Principles of numerical taxonomy*. San Francisco: W. H. Freeman.

Taylor, J., & Tibshirani, R. J. (2015). Statistical learning and selective inference. *Proceedings of the National Academy of Sciences*, *112*(25), 7629–7634.

Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, *58*(301), 236–244.

Xie, J., Kelley, S., & Szymanski, B. K. (2013). Overlapping community detection in networks: The state-of-the-art and comparative study. *Acm computing surveys (csur)*, *45*(4), 1–35.

Xie, X. L., & Beni, G. (1991). A new fuzzy clustering validity criterion and its application to color image segmentation. In *Proceedings of the 1991 ieee international symposium on intelligent control* (pp. 463–468).

Yun, Y.-J., & Foygel Barber, R. (2023). Selective inference for clustering with unknown variance. *Electronic Journal of Statistics*, *17*(2), 1923–1946.

Zhang, D.-Q., & Chen, S.-C. (2003). Clustering incomplete data using kernel-based fuzzy c-means algorithm. *Neural processing letters*, *18*(3), 155–162.

Zhang, D.-Q., & Chen, S.-C. (2004). A novel kernelized fuzzy c-means algorithm with application in medical image segmentation. *Artificial intelligence in medicine*, *32*(1), 37–50.

Zhu, Z., Gai, X., & Zhang, A. R. (2024). Functional post-clustering selective inference with applications to ehr data analysis. *arXiv preprint arXiv:2405.03042*.

Supplementary Materials for

"Statistical Inference for Fuzzy Clustering"

A Proofs for Section 5

Label switching convention. Because the model is invariant under permutations of the k components, we work with an identifiable representative set $\Theta^* \subset \Theta$ (e.g., lexicographic ordering of the centers $(\mathbf{v}_1, \ldots, \mathbf{v}_k)$ with a deterministic tie-break rule). On Θ^* the mapping $\theta \mapsto f_\theta$ is identifiable in the ordinary sense. Consistency on Θ^* implies consistency on Θ up to permutation.

We leverage Wald's MLE consistency theorem.

Theorem 3 *Let $\{f_\theta : \theta \in \Theta\}$ be a family of densities with respect to a dominating measure μ. Suppose*

(C1) Θ is measurable and $\sup_{\theta \in \Theta} \left| \log f_\theta(\mathbf{x}) \right| \in L^1_{\theta_0}$ (i.e., $\sup_{\theta \in \Theta} \left| \log f_\theta(\mathbf{x}) \right|$ is integrable under density f_{θ_0});

(C2) Θ is compact and $\theta \mapsto \log f_\theta(\mathbf{x})$ is continuous for every $\mathbf{x}$;

(C3) $\{\log f_\theta : \theta \in \Theta\}$ admits an integrable envelope, so that

$$\sup_{\theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n \log f_\theta(\mathbf{x}_i) - \mathbb{E}_{\theta_0} \log f_\theta(\mathbf{x}) \right| \xrightarrow{a.s.} 0;$$

(C4) $\theta \neq \theta'$ implies $f_\theta \neq f_{\theta'}$ μ-almost everywhere

and assume a global maximizer $\hat{\theta}_n \in \mathrm{argmax}_{\theta \in \Theta} \sum_{i=1}^n \log f_\theta(\mathbf{x}_i)$ exists for every n. Then $\hat{\theta}_n \to \theta_0$ almost surely as $n \to \infty$.

We verify (C1)–(C3) on Θ and verify (C4) on the identifiable subset Θ^*; the final conclusion is then lifted back to Θ as consistency up to permutation. Throughout, let $\theta_0 = (\sigma_0, \mathbf{w}_0, \mathbf{V}_0, m_0) \in \mathrm{int}(\Theta)$.

Condition (C1). Lebesgue measure dominates because every $f_\theta > 0$ and is integrable.

Fix $R := \max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{v}\|$. If $\|\mathbf{x}\| > 2R$, then $\|\mathbf{x} - \mathbf{v}_j\| \geq \|\mathbf{x}\|/2$ for all j. Since $w_j \leq 1$, we have $w_j^{-1/(m-1)} \geq 1$, hence with $q = 2/(m-1)$,

$$\Sigma_{\mathbf{w}}(\mathbf{x}) = \sum_{j=1}^{k} w_j^{-1/(m-1)} \|\mathbf{x} - \mathbf{v}_j\|^{-q} \geq \sum_{j=1}^{k} (\|\mathbf{x}\|/2)^{-q} = k \, 2^q \, \|\mathbf{x}\|^{-q}.$$

Therefore, for $\|\mathbf{x}\| > 2R$,

$$[\Sigma_{\mathbf{w}}(\mathbf{x})]^{-(m-1)} \leq \left(k \, 2^q\right)^{-(m-1)} \|\mathbf{x}\|^{q(m-1)} = \left(k \, 2^q\right)^{-(m-1)} \|\mathbf{x}\|^2.$$

Because $m \in [m_{\min}, m_{\max}]$, the factor $\left(k \, 2^q\right)^{-(m-1)}$ is uniformly bounded above by a finite constant $A < \infty$. Also $\sigma^{-2} \leq \sigma_{\min}^{-2}$, so for $\|\mathbf{x}\| > 2R$,

$$E_\theta(\mathbf{x}) = \sigma^{-2} [\Sigma_{\mathbf{w}}(\mathbf{x})]^{-(m-1)} \leq \sigma_{\min}^{-2} A \, \|\mathbf{x}\|^2.$$

Moreover, $\log C(\theta)$ is bounded on Θ (proved in (C2) below), hence for $\|\mathbf{x}\| > 2R$,

$$\sup_{\theta \in \Theta} |\log f_\theta(\mathbf{x})| \leq \sup_{\theta \in \Theta} |\log C(\theta)| + \sigma_{\min}^{-2} A \, \|\mathbf{x}\|^2.$$

Inside the ball $\|\mathbf{x}\| \leq 2R$, continuity on the compact set $\Theta \times \{\|\mathbf{x}\| \leq 2R\}$ furnishes a finite bound. Therefore $M(\mathbf{x}) := \sup_{\theta \in \Theta} |\log f_\theta(\mathbf{x})| \in L_{\theta_0}^1$, verifying (C1).

Condition (C2). Write

$$\ell_\theta(\mathbf{x}) = \log f_\theta(\mathbf{x}) = \log C(\theta) - E_\theta(\mathbf{x}), \qquad E_\theta(\mathbf{x}) = \sigma^{-2} \left[\Sigma_{\mathbf{w}}(\mathbf{x})\right]^{-(m-1)}.$$

Each constituent map is continuous on Θ: $\sigma \mapsto \sigma^{-2}$ is continuous on $(0, \infty)$; the power $u \mapsto u^{-(m-1)}$ is continuous on $(0, \infty)$ uniformly for $m \in [m_{\min}, m_{\max}]$; and $\mathbf{x} \mapsto \|\mathbf{x} - \mathbf{v}_j\|$ is continuous in $\mathbf{v}_j$.

It remains to justify continuity and boundedness of $C(\theta)$. Let $Z(\theta) :=$ $\int_{\mathbb{R}^d} \exp(-E_\theta(\mathbf{x})) \, d\mathbf{x}$, so $C(\theta) = Z(\theta)^{-1}$. We claim there exist constants $c > 0$ and $R_0 > 0$ such that for all $\theta \in \Theta$ and $\|\mathbf{x}\| > R_0$,

$$E_\theta(\mathbf{x}) \geq c\|\mathbf{x}\|^2.$$

Indeed, for $\|\mathbf{x}\| > 2R$ we have $\|\mathbf{x} - \mathbf{v}_j\| \le \|\mathbf{x}\| + R \le \frac{3}{2}\|\mathbf{x}\|$ and $w_j \ge \varepsilon_w$, so

$$\Sigma_{\mathbf{w}}(\mathbf{x}) \le \sum_{j=1}^{k} \varepsilon_w^{-1/(m_{\min}-1)} \left(\tfrac{3}{2}\|\mathbf{x}\|\right)^{-q} \le B\,\|\mathbf{x}\|^{-q},$$

with $B < \infty$ depending only on $(k, \varepsilon_w, m_{\min}, m_{\max}, R)$. Thus $[\Sigma_{\mathbf{w}}(\mathbf{x})]^{-(m-1)} \ge B^{-(m-1)}\|\mathbf{x}\|^{q(m-1)} = B^{-(m-1)}\|\mathbf{x}\|^2$, and since $\sigma^{-2} \ge \sigma_{\max}^{-2}$ and $B^{-(m-1)}$ is uniformly bounded below on $[m_{\min}, m_{\max}]$, we obtain the stated lower bound $E_\theta(\mathbf{x}) \ge c\|\mathbf{x}\|^2$.

Consequently,

$$\exp(-E_\theta(\mathbf{x})) \le \mathbf{1}_{\{\|\mathbf{x}\| \le R_0\}} + \exp(-c\|\mathbf{x}\|^2)\mathbf{1}_{\{\|\mathbf{x}\| > R_0\}},$$

an integrable dominator independent of θ. Since $E_\theta(\mathbf{x})$ is continuous in θ for each fixed $\mathbf{x}$, dominated convergence implies $\theta \mapsto Z(\theta)$ is continuous, hence $\theta \mapsto C(\theta) = 1/Z(\theta)$ is continuous. Furthermore, on compact Θ, $Z(\theta)$ is bounded away from 0 and ∞, so $\sup_{\theta \in \Theta} |\log C(\theta)| < \infty$.

Compactness of Θ is explicit in its definition. Therefore $\ell_\theta(\mathbf{x})$ is continuous in θ for every fixed $\mathbf{x}$, and (C2) holds.

Condition (C3). Define $M(\mathbf{x}) := \sup_{\theta \in \Theta} |\ell_\theta(\mathbf{x})|$. By (C2), $C_0 := \sup_{\theta \in \Theta} |\log C(\theta)| < \infty$. By the tail bound in (C1), for $\|\mathbf{x}\| > 2R$ we have $\sup_{\theta \in \Theta} E_\theta(\mathbf{x}) \le K\|\mathbf{x}\|^2$ for some $K < \infty$. Hence, for $\|\mathbf{x}\| > 2R$,

$$M(\mathbf{x}) \le C_0 + K\|\mathbf{x}\|^2.$$

On the compact set $\|\mathbf{x}\| \le 2R$, continuity on the compact $\Theta \times \{\|\mathbf{x}\| \le 2R\}$ implies $M(\mathbf{x}) \le B$ for some finite B. Thus the envelope

$$\widetilde{M}(\mathbf{x}) := B\,\mathbf{1}_{\{\|\mathbf{x}\| \le 2R\}} + (C_0 + K\|\mathbf{x}\|^2)\mathbf{1}_{\{\|\mathbf{x}\| > 2R\}}$$

satisfies $M(\mathbf{x}) \le \widetilde{M}(\mathbf{x})$ and $\mathbb{E}_{\theta_0}\widetilde{M}(\mathbf{x}) < \infty$. Therefore a uniform strong law of large numbers applies and (C3) holds.

Condition (C4) (identifiability up to permutation). Assume $f_\theta = f_{\theta'}$ Lebesgue almost everywhere. Then for all $\mathbf{x}$ away from the centers of either parameter,

$$E_\theta(\mathbf{x}) - E_{\theta'}(\mathbf{x}) = \log C(\theta) - \log C(\theta') =: \kappa.$$

Step 0: show $\kappa = 0$ and the center sets coincide. For any center $\mathbf{v}_j$ of θ, we have $\Sigma_{\mathbf{w}}(\mathbf{v}_j) = \infty$, hence $E_\theta(\mathbf{v}_j) = 0$ and $f_\theta(\mathbf{v}_j) = C(\theta)$ (interpreting the continuous extension at $\mathbf{v}_j$). Therefore,

$$C(\theta) = f_\theta(\mathbf{v}_j) = f_{\theta'}(\mathbf{v}_j) = C(\theta') \exp(-E_{\theta'}(\mathbf{v}_j)) \le C(\theta').$$

By symmetry (swap θ and θ') we obtain $C(\theta) = C(\theta')$, hence $\kappa = 0$, and also $E_{\theta'}(\mathbf{v}_j) = 0$. But $E_{\theta'}(\mathbf{x}) = 0$ iff $\mathbf{x}$ is one of the primed centers, because away from centers $\Sigma_{\mathbf{w}'}(\mathbf{x}) < \infty$ so $E_{\theta'}(\mathbf{x}) > 0$. Thus every $\mathbf{v}_j$ is a primed center. Reversing the argument shows the unordered center sets coincide:

$$\{\mathbf{v}_1, \ldots, \mathbf{v}_k\} = \{\mathbf{v}'_1, \ldots, \mathbf{v}'_k\}.$$

Fix a permutation so that $\mathbf{v}_j = \mathbf{v}'_j$ for all j.

Step 1: identify m. Fix j and write $\mathbf{v} = \mathbf{v}_j$. Let $\mathbf{x}(t) = \mathbf{v} + t\mathbf{u}$ with $\mathbf{u} \in \mathbb{S}^{d-1}$ and $t \downarrow 0$. Write $q = 2/(m-1)$ and expand

$$\Sigma_{\mathbf{w}}(\mathbf{x}(t)) = a\, t^{-q} + S + O(t), \qquad a = w_j^{-1/(m-1)}, \quad S = \sum_{i \ne j} w_i^{-1/(m-1)} \|\mathbf{v} - \mathbf{v}_i\|^{-q} > 0.$$

Then a binomial expansion yields

$$[\Sigma_{\mathbf{w}}(\mathbf{x}(t))]^{-(m-1)} = w_j\, t^2 \left(1 - (m-1)\frac{S}{a}t^q + o(t^q)\right),$$

and hence

$$E_\theta(\mathbf{x}(t)) = \sigma^{-2} w_j\, t^2 - \sigma^{-2} w_j(m-1)\frac{S}{a} t^{2+q} + o(t^{2+q}).$$

The same expansion holds for θ' with $(m', \sigma', w'_j, S', a', q')$. Since $E_\theta(\mathbf{x}(t)) \equiv E_{\theta'}(\mathbf{x}(t))$ for all small t, the second term forces $2 + q = 2 + q'$, hence $q = q'$ and $m = m'$.

Step 2: identify σ and the weights. With $m = m'$ and common centers aligned, compare the leading t^2 coefficients:

$$\sigma^{-2} w_j = \sigma'^{-2} w'_j, \qquad j = 1, \ldots, k.$$

Summing over j and using $\sum_{j=1}^{k} w_j = \sum_{j=1}^{k} w'_j = 1$ gives $\sigma^{-2} = \sigma'^{-2}$, hence $\sigma = \sigma'$. Then $w_j = w'_j$ for all j. This proves identifiability on Θ^*; on Θ it holds modulo permutations.

We have verified (C1)–(C3) on Θ and (C4) on the identifiable representative Θ^*. A global maximizer exists because Θ is compact and the log-likelihood is continuous. Applying Theorem 3 on Θ^* yields $\hat{\theta}_n \to \theta_0$ a.s. Lifting back to Θ gives consistency up to label permutation.

B Asymptotic normality of MLE

We verify (R1), (R2), and (R4) (condition (R3) was already established for Theorem 1), and then apply a standard multivariate MLE limit theorem to obtain the desired asymptotic normality.

(R1) **Interior point and smoothness.** $\theta_0 \in \text{int}\,\Theta$ and $\theta \mapsto \ell_\theta(\mathbf{x}) = \log f_\theta(\mathbf{x})$ is three-times continuously differentiable on a neighbourhood $N(\theta_0)$ for every $\mathbf{x} \notin \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$. (The score and Hessian are integrable under f_{θ_0}; see (R2).)

(R2) **Envelope for first and second derivatives.** There exist measurable $M_1, M_2 \in L^1(f_{\theta_0})$ such that, for all $\theta \in N(\theta_0)$,

$$\|\dot{\ell}_\theta(\mathbf{x})\| \leq M_1(\mathbf{x}), \qquad \|\ddot{\ell}_\theta(\mathbf{x})\| \leq M_2(\mathbf{x}).$$

(R3) **Identifiability and consistency.** The normalised model is identifiable modulo permutation and the MLE is strongly consistent (proved in the previous section), so after relabeling we may assume $\hat{\theta}_n \to \theta_0$ a.s.

(R4) **Fisher information on the tangent space (or reduced weights).** Because the weights satisfy the simplex constraint $\sum_{j=1}^{k} w_j = 1$, the original Fisher information is singular in the ambient parametrization $(w_1, \ldots, w_k)$. Let

$$\mathcal{T} := \left\{ a \in \mathbb{R}^p : \sum_{j=1}^{k} a_{w_j} = 0 \right\}$$

denote the tangent space to the simplex at $\mathbf{w}_0$ (with all other coordinates free). We assume $I(\theta_0)$ is finite and positive-definite on $\mathcal{T}$ (equivalently, invertible in a reduced parametrization $\tilde{\mathbf{w}} = (w_1, \ldots, w_{k-1})$ with $w_k = 1 - \sum_{j=1}^{k-1} w_j$).

Recall that

$$\ell_\theta(\mathbf{x}) = \log f_\theta(\mathbf{x}) = -E_\theta(\mathbf{x}) + \log C(\theta),$$

and

$$E_\theta(\mathbf{x}) = \sigma^{-2} \big[\Sigma_{\mathbf{w}}(\mathbf{x}) \big]^{-(m-1)}, \qquad \Sigma_{\mathbf{w}}(\mathbf{x}) = \sum_{j=1}^{k} w_j^{-\frac{1}{m-1}} \|\mathbf{x} - \mathbf{v}_j\|^{-q}, \qquad q = \frac{2}{m-1}.$$

Differentiating with respect to θ and using the identity

$$\partial_\theta \log C(\theta) = \frac{\partial_\theta C(\theta)}{C(\theta)} = \frac{\partial_\theta \big(Z(\theta)^{-1} \big)}{Z(\theta)^{-1}} = -\partial_\theta \log Z(\theta) = -\frac{\int (\partial_\theta E_\theta(\mathbf{x})) e^{-E_\theta(\mathbf{x})} \, d\mathbf{x}}{\int e^{-E_\theta(\mathbf{x})} \, d\mathbf{x}} = -\mathbb{E}_\theta[\partial_\theta E_\theta(\mathbf{x})]$$

gives the score decomposition

$$\dot{\ell}_\theta(\mathbf{x}) = -\partial_\theta E_\theta(\mathbf{x}) + \mathbb{E}_\theta[\partial_\theta E_\theta(\mathbf{x})].$$

A second differentiation yields

$$\ddot{\ell}_\theta(\mathbf{x}) = -\ddot{E}_\theta(\mathbf{x}) + \mathbb{E}_\theta[\ddot{E}_\theta(\mathbf{x})] + \Delta_\theta(\mathbf{x}) \Delta_\theta(\mathbf{x})^\top, \qquad \Delta_\theta(\mathbf{x}) := \partial_\theta E_\theta(\mathbf{x}) - \mathbb{E}_\theta[\partial_\theta E_\theta(\mathbf{x})].$$

Derivatives of E_θ. Write $Q(\mathbf{x}) := \Sigma_{\mathbf{w}}(\mathbf{x})$. Elementary differentiation gives

$$\partial_\sigma E_\theta(\mathbf{x}) = -2\sigma^{-3} Q(\mathbf{x})^{-(m-1)},$$

$$\partial_{w_j} E_\theta(\mathbf{x}) = \sigma^{-2} Q(\mathbf{x})^{-m} w_j^{-1/(m-1)-1} \|\mathbf{x} - \mathbf{v}_j\|^{-q},$$

$$\partial_{\mathbf{v}_j} E_\theta(\mathbf{x}) = (m-1) q \, \sigma^{-2} \, Q(\mathbf{x})^{-m} w_j^{-1/(m-1)} \|\mathbf{x} - \mathbf{v}_j\|^{-q-2} (\mathbf{x} - \mathbf{v}_j), \tag{12}$$

$$\partial_m E_\theta(\mathbf{x}) = -\sigma^{-2} Q(\mathbf{x})^{-(m-1)} \log Q(\mathbf{x}) - \sigma^{-2} (m-1) Q(\mathbf{x})^{-m} \partial_m Q(\mathbf{x}).$$

Verification of (R2). We control the derivatives in two regions: (i) the tails $\|\mathbf{x}\| \to \infty$ and (ii) neighborhoods of the centers $\mathbf{v}_j$.

(i) Tails. Let $R := \max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{v}\|$. For $\|\mathbf{x}\| > 2R$, we have $\|\mathbf{x} - \mathbf{v}_j\| \asymp \|\mathbf{x}\|$ uniformly in j. Moreover $Q(\mathbf{x}) \asymp \|\mathbf{x}\|^{-q}$ and therefore $Q(\mathbf{x})^{-(m-1)} \asymp \|\mathbf{x}\|^{q(m-1)} = \|\mathbf{x}\|^2$. A routine check

46

from (12) shows that each component of $\partial_\theta E_\theta(\mathbf{x})$ is $O(\|\mathbf{x}\|^2 \log \|\mathbf{x}\|)$ and each component of $\ddot{E}_\theta(\mathbf{x})$ is $O(\|\mathbf{x}\|^2 (\log \|\mathbf{x}\|)^2)$ uniformly over $\theta \in N(\theta_0)$.

(ii) Near centers. Fix j and write $\mathbf{x}(t) = \mathbf{v}_j + t\mathbf{u}$ with $\mathbf{u} \in \mathbb{S}^{d-1}$ and $t \downarrow 0$. Let $q_0 = 2/(m_0 - 1)$. Since $Q(\mathbf{x}(t)) \sim w_j^{-1/(m-1)} t^{-q}$ and $Q(\mathbf{x}(t))^{-m} \sim w_j^{m/(m-1)} t^{qm}$, the apparent singularities in $\|\mathbf{x} - \mathbf{v}_j\|^{-q}$ cancel. In particular, the leading orders are

$$\partial_{\mathbf{v}_j} E_\theta(\mathbf{x}(t)) = O(t), \qquad \partial_{w_j} E_\theta(\mathbf{x}(t)) = O(t^2), \qquad \partial_\sigma E_\theta(\mathbf{x}(t)) = O(t^2), \qquad \partial_m E_\theta(\mathbf{x}(t)) = O(t^2 \log(1/t)),$$

and for $\ell \neq j$, $\partial_{w_\ell} E_\theta(\mathbf{x}(t)) = O(t^{2+q})$, $\partial_{\mathbf{v}_\ell} E_\theta(\mathbf{x}(t)) = O(t^{2+q})$. Thus $\partial_\theta E_\theta(\mathbf{x})$ and $\ddot{E}_\theta(\mathbf{x})$ are locally bounded (indeed locally integrable) in punctured neighborhoods of each $\mathbf{v}_j$, uniformly over $\theta \in N(\theta_0)$.

Combining (i) and (ii), and using that $f_{\theta_0}(\mathbf{x}) \leq \exp(-c\|\mathbf{x}\|^2)$ in the tails (from the quadratic lower bound on E_{θ_0}), we obtain integrable envelopes $M_1, M_2 \in L^1(f_{\theta_0})$ such that $\|\dot{\ell}_\theta(\mathbf{x})\| \leq M_1(\mathbf{x})$ and $\|\ddot{\ell}_\theta(\mathbf{x})\| \leq M_2(\mathbf{x})$ for all $\theta \in N(\theta_0)$. This proves (R2).

Fisher information matrix. Because $\mathbb{E}_{\theta_0}[\dot{\ell}_{\theta_0}(\mathbf{x})] = 0$,

$$I(\theta_0) = \mathbb{E}_{\theta_0}\left[\dot{\ell}_{\theta_0}(\mathbf{x})\dot{\ell}_{\theta_0}(\mathbf{x})^\top\right] = \mathrm{Var}_{\theta_0}[\partial_\theta E_{\theta_0}(\mathbf{x})].$$

Verification of (R4) on the tangent space. Let $a \in \mathcal{T}$ satisfy $a^\top I(\theta_0)a = 0$. Then $a^\top \dot{\ell}_{\theta_0}(\mathbf{x}) = 0$ almost surely, hence $a^\top \partial_\theta E_{\theta_0}(\mathbf{x})$ is almost surely constant. Since f_{θ_0} is strictly positive on $\mathbb{R}^d \setminus \{\mathbf{v}_{01}, \ldots, \mathbf{v}_{0k}\}$, this implies that for Lebesgue-almost everywhere $\mathbf{x}$,

$$a^\top \partial_\theta E_{\theta_0}(\mathbf{x}) \equiv \text{constant.} \tag{13}$$

Step 1: eliminate center directions. Fix j and set $\mathbf{x}(t) = \mathbf{v}_{0j} + t\mathbf{u}$ with $\mathbf{u} \in \mathbb{S}^{d-1}$. From the local expansion above,

$$\partial_{\mathbf{v}_{0j}} E_{\theta_0}(\mathbf{x}(t)) = 2\sigma_0^{-2} w_{0j} \, t \, \mathbf{u} + O(t^{1+q_0}),$$

while all other blocks in $\partial_\theta E_{\theta_0}(\mathbf{x}(t))$ are $O(t^2)$ or smaller. Dividing (13) by t and letting $t \downarrow 0$ yields $a_{\mathbf{v}_{0j}} = 0$. Repeating for all j gives $a_{\mathbf{V}} = 0$.

47

Step 2: eliminate weight directions modulo the simplex. With $a_{\mathbf{V}} = 0$, evaluate (13) at $\mathbf{x}(t) = \mathbf{v}_{0j} + t\mathbf{u}$ and divide by t^2. Using

$$\partial_{w_j} E_{\theta_0}(\mathbf{x}(t)) = \sigma_0^{-2} t^2 + o(t^2), \qquad \partial_{w_\ell} E_{\theta_0}(\mathbf{x}(t)) = o(t^2) \ (\ell \neq j),$$

we obtain a_{w_j} is the same constant for every j, i.e. $a_{\mathbf{w}} = c(1, \ldots, 1)$. But $a \in \mathcal{T}$ implies $\sum_{j=1}^k a_{w_j} = 0$, hence $c = 0$ and $a_{\mathbf{w}} = 0$.

Step 3: eliminate (σ, m). With $a_{\mathbf{V}} = a_{\mathbf{w}} = 0$, only (a_σ, a_m) may be nonzero. Let $\|\mathbf{x}\| \to \infty$. Since $Q_0(\mathbf{x}) \asymp \|\mathbf{x}\|^{-q_0}$, we have

$$\partial_\sigma E_{\theta_0}(\mathbf{x}) \asymp \|\mathbf{x}\|^2, \qquad \partial_m E_{\theta_0}(\mathbf{x}) \asymp \|\mathbf{x}\|^2 \log \|\mathbf{x}\|,$$

up to lower-order terms. Substituting into (13), dividing by $\|\mathbf{x}\|^2$, and letting $\|\mathbf{x}\| \to \infty$ forces $a_m = 0$ (by linear independence of 1 and $\log \|\mathbf{x}\|$), and then $a_\sigma = 0$. Thus $a = 0$ on $\mathcal{T}$, proving positive definiteness of $I(\theta_0)$ on $\mathcal{T}$.

With (R1)–(R4) in place and after resolving label switching by a measurable permutation rule, a standard multivariate MLE theorem yields

$$\sqrt{n}\big(\pi_n(\hat{\theta}_n) - \theta_0\big) \xrightarrow{d} \mathcal{N}\big(0, I(\theta_0)^\dagger\big),$$

where $I(\theta_0)^\dagger$ denotes the Moore–Penrose pseudoinverse in the ambient parametrization.

C   Properties of WFCM Density

Theorem 4 *Let* $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_k)$, $\mathbf{w} = (w_1, \ldots, w_k)$ *with* $w_j > 0$ *and* $\sum_{j=1}^k w_j = 1$, *and* $\sigma > 0$. *Define the (unnormalized) WFCM density*

$$f(\mathbf{x}) \ \propto \ \exp\left(-\frac{1}{\sigma^2 \, g_m(\mathbf{x})}\right), \qquad g_m(\mathbf{x}) := \left[\sum_{j=1}^k (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\frac{1}{m-1}}\right]^{m-1},$$

for $\mathbf{x} \in \mathbb{R}^d$ *with* $\mathbf{x} \neq \mathbf{v}_j$ *for all* j. *Then the following limiting properties hold.*

(i) *As* $m \to 1^+$,

$$g_m(\mathbf{x}) \ \longrightarrow \ \max_{1 \leq j \leq k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1} = \frac{1}{\min_{1 \leq j \leq k} w_j \|\mathbf{x} - \mathbf{v}_j\|^2},$$

and hence

$$f(\mathbf{x}) \; \propto \; \exp\left(-\frac{\min_{1\leq j\leq k} w_j \|\mathbf{x} - \mathbf{v}_j\|^2}{\sigma^2}\right) = \max_{1\leq j\leq k} \exp\left(-\frac{w_j \|\mathbf{x} - \mathbf{v}_j\|^2}{\sigma^2}\right).$$

In particular, the dominant contribution comes from the centroid(s) minimizing $w_j \|\mathbf{x} - \mathbf{v}_j\|^2$ (and if w_j are equal, this reduces to the nearest centroid(s)).

(ii) *For each fixed $\mathbf{x}$ with $\mathbf{x} \neq \mathbf{v}_j$ for all j,*

$$g_m(\mathbf{x}) \;\longrightarrow\; \infty \quad (\text{indeed } g_m(\mathbf{x}) \sim k^{m-1}), \quad \text{and hence} \quad f(\mathbf{x}) \;\longrightarrow\; 1,$$

i.e., the model approaches an improper flat density on $\mathbb{R}^d$. On any bounded domain, the normalized restriction converges to the uniform distribution.

(iii) *When $m = 2$,*

$$f(\mathbf{x}) \; \propto \; \exp\left(-\frac{1}{\sigma^2 \sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1}}\right).$$

Proof Throughout, fix $\mathbf{x} \in \mathbb{R}^d$ with $\mathbf{x} \neq \mathbf{v}_j$ for all j.

(i) Limit as $m \to 1^+$. Apply Lemma 5 with $a_j = (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1}$ to get

$$\lim_{m\to 1^+} g_m(\mathbf{x}) = \max_{1\leq j\leq k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1} = \frac{1}{\min_{1\leq j\leq k} w_j \|\mathbf{x} - \mathbf{v}_j\|^2}.$$

Therefore,

$$f(\mathbf{x}) \; \propto \; \exp\left(-\frac{1}{\sigma^2 \max_j (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1}}\right) = \exp\left(-\frac{\min_j w_j \|\mathbf{x} - \mathbf{v}_j\|^2}{\sigma^2}\right) = \max_j \exp\left(-\frac{w_j \|\mathbf{x} - \mathbf{v}_j\|^2}{\sigma^2}\right).$$

(ii) Limit as $m \to \infty$. Let $\alpha_m := \frac{1}{m-1} \downarrow 0$. Then, for each j,

$$(w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\alpha_m} = \exp\left(-\alpha_m \log(w_j \|\mathbf{x} - \mathbf{v}_j\|^2)\right) \;\longrightarrow\; 1.$$

Since the sum has finitely many terms,

$$\sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\alpha_m} \;\longrightarrow\; k,$$

and thus

$$g_m(\mathbf{x}) = \left[\sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-\alpha_m}\right]^{m-1} \;\longrightarrow\; k^{m-1} \to \infty.$$

49

Consequently,

$$f(\mathbf{x}) \ \propto \ \exp\left(-\frac{1}{\sigma^2 g_m(\mathbf{x})}\right) \ \longrightarrow \ \exp(0) = 1,$$

a constant in $\mathbf{x}$, i.e., an improper flat density on $\mathbb{R}^d$. If we restrict f to a bounded measurable set and renormalize, the limit is uniform.

(iii) Case $m = 2$. Substitute $m = 2$ into the definition of $g_m(\mathbf{x})$ to obtain

$$g_2(\mathbf{x}) = \sum_{j=1}^{k} (w_j \|\mathbf{x} - \mathbf{v}_j\|^2)^{-1} = \sum_{j=1}^{k} w_j^{-1} \|\mathbf{x} - \mathbf{v}_j\|^{-2},$$

and plug into $f(\mathbf{x})$. ∎

Lemma 5 *Let* $\{a_j\}_{j=1}^{k}$ *be positive numbers and let* $\{w_j\}_{j=1}^{k}$ *satisfy* $w_j > 0$ *and* $\sum_{j=1}^{k} w_j = 1$. *Then*

$$\lim_{m \to 1^+} \left(\sum_{j=1}^{k} a_j^{\frac{1}{m-1}}\right)^{m-1} = \max_{1 \le j \le k} a_j, \qquad \lim_{m \to 1^+} \left(\sum_{j=1}^{k} w_j\, a_j^{\frac{1}{m-1}}\right)^{m-1} = \max_{1 \le j \le k} a_j.$$

Proof Let $t := m - 1 \to 0^+$ and write $r_j := a_j / a_{\max} \in (0, 1]$. Define

$$S_t := \sum_{j=1}^{k} w_j r_j^{1/t}.$$

Then

$$\left(\sum_{j=1}^{k} w_j a_j^{1/t}\right)^{t} = \left(a_{\max}^{1/t} \sum_{j=1}^{k} w_j r_j^{1/t}\right)^{t} = a_{\max} (S_t)^t.$$

It remains to show $(S_t)^t \to 1$.

Let $s_* := \sum_{j \in J^*} w_j \in (0, 1]$. For $j \in J^*$, $r_j = 1$. For $j \notin J^*$, we have $r_j < 1$; define

$$r := \max_{j \notin J^*} r_j \in [0, 1).$$

(If $J^* = [k]$, set $r := 0$; then the argument below still holds.)

Two-sided bound for S_t. Since $r_j^{1/t} = 1$ for $j \in J^*$ and $0 \le r_j^{1/t} \le r^{1/t}$ for $j \notin J^*$, we obtain

$$s_* = \sum_{j \in J^*} w_j \le \sum_{j=1}^{k} w_j r_j^{1/t} \le \sum_{j \in J^*} w_j + \sum_{j \notin J^*} w_j\, r^{1/t} = s_* + (1 - s_*) r^{1/t}.$$

Hence, for all $t > 0$,

$$s_* \leq S_t \leq s_* + (1 - s_*)r^{1/t}. \tag{14}$$

Two-sided bound for $(S_t)^t$. Raising (14) to the power t (monotonicity since all terms are positive) gives

$$(s_*)^t \leq (S_t)^t \leq \left(s_* + (1 - s_*)r^{1/t}\right)^t.$$

As $t \to 0^+$, we have $(s_*)^t \to 1$. Also $r^{1/t} \to 0$ because $0 \leq r < 1$, so $s_* + (1 - s_*)r^{1/t} \to s_*$. Therefore,

$$\lim_{t \to 0^+} \left(s_* + (1 - s_*)r^{1/t}\right)^t = \exp\left(\lim_{t \to 0^+} t \log\left(s_* + (1 - s_*)r^{1/t}\right)\right) = \exp(0) = 1,$$

since $\log\left(s_* + (1 - s_*)r^{1/t}\right) \to \log s_*$ is finite. By the squeeze theorem, $(S_t)^t \to 1$.

Finally,

$$\lim_{t \to 0^+} \left(\sum_{j=1}^{k} w_j a_j^{1/t}\right)^t = \lim_{t \to 0^+} a_{\max}(S_t)^t = a_{\max}.$$

This proves the weighted claim; the unweighted claim follows by taking $w_j \equiv 1/k$. ∎