

Nonparametric MLE for Gaussian Location Mixtures: Certified Computation and Generic Behavior

Yury Polyanskiy*

Mark Sellke†

Abstract

We study the nonparametric maximum likelihood estimator $\hat{\pi}$ for Gaussian location mixtures in one dimension. It has been known since [Lin83a] that given an n -point dataset, this estimator always returns a mixture with at most n components, and more recently [PW20] gave a sharp $O(\log n)$ bound for subgaussian data. In this work we study computational aspects of $\hat{\pi}$. We provide an algorithm which for small enough $\varepsilon > 0$ computes an ε -approximation of $\hat{\pi}$ in Wasserstein distance in time $K + Cnk^2 \log \log(1/\varepsilon)$. Here K is data-dependent but independent of ε , while C is an absolute constant and $k = |\text{supp}(\hat{\pi})| \leq n$ is the number of atoms in $\hat{\pi}$. We also certifiably compute the exact value of $|\text{supp}(\hat{\pi})|$ in finite time. These guarantees hold almost surely whenever the dataset $(x_1, \dots, x_n) \in [-cn^{1/4}, cn^{1/4}]$ consists of independent points from a probability distribution with a density (relative to Lebesgue measure). We also show the distribution of $\hat{\pi}$ conditioned to be k -atomic admits a density on the associated $2k - 1$ dimensional parameter space for all $k \leq \sqrt{n}/3$, and almost sure locally linear convergence of the EM algorithm. One key tool is a classical Fourier analytic estimate for non-degenerate curves.

1 Introduction

The *nonparametric maximum likelihood estimator* (NPMLE) has a long history in statistical problems including density estimation, regression, and mixture models (see [GW12]). This article concerns the NPMLE for the 1-dimensional Gaussian location model, which has been studied since 1950s, cf. [KW56, Rob50]. To introduce this problem, for any probability distribution π on $\mathbb{R}$ we denote $P_\pi \triangleq \pi * \mathcal{N}(0, 1)$ the probability density function of its convolution with a standard Gaussian density:

$$P_\pi(x) = \mathbb{E}^{y \sim \pi}[e^{-|x-y|^2/2}/\sqrt{2\pi}] \quad (1.1)$$

Given a finite sequence $X = (x_1, \dots, x_n) \in \mathbb{R}^n$, the NPMLE $\hat{\pi} \in \mathcal{P}(\mathbb{R})$ is a probability measure on $\mathbb{R}$ chosen so that $P_{\hat{\pi}} = \hat{\pi} * \mathcal{N}(0, 1)$ maximizes the log-likelihood

$$\ell_X(\pi) \triangleq \frac{1}{n} \sum_{i=1}^n \log P_\pi(x_i). \quad (1.2)$$

Despite the fact that $\hat{\pi}$ is defined as a solution of an infinite-dimensional optimization problem, both the problem and its solution are surprisingly well-behaved. First, the maximizer $\hat{\pi}$ exists and is unique [LR93]. Second, despite being defined as an optimum over all possible probability measures, $\hat{\pi}$ is discrete, so that the $P_{\hat{\pi}}$ becomes a finite mixture. Furthermore, even though this estimator lacks any explicit regularization, it

*Department of Electrical Engineering and Computer Science, MIT. Email: yp@mit.edu

†Department of Statistics, Harvard University. Email: msellke@fas.harvard.edu

nevertheless enjoys spectacular statistical convergence properties, for example achieving best known density estimation rates [JZ09, SG20].

The NPMLE is tractable to study in part because of its nature as an overparametrized convex relaxation. Indeed, while it is well known that the landscape of the maximization objective (1.2) if restricted to k -atomic distributions π is non-concave and has spurious local maxima [JZB⁺16], over the space of all measures $\mathcal{P}(\mathbb{R})$ the problem is concave, and thus can be characterized by local optimality conditions.

Practical solvers for maximizing (1.2) started with [Lai78] proposing a variant of the then-recently discovered EM algorithm. Due to its very slow convergence, a number of other algorithms were proposed over the years mostly differing in how new locations (atoms) are added at each iteration, see e.g. [Der86, Böh86, LK92, BSL92] and [Lin83b, Chapter 6] for a detailed survey. However, a decade ago [KM14] discovered that due to the progress in convex optimization, the (empirically) fastest and most accurate way to maximize (1.2) is to fix the support of π to be a fine equi-spaced grid $\varepsilon\mathbb{Z}$ (truncated at the range of the samples X) and maximize the concave function $\ell_X(\pi)$ of the weights of π via off-the-shelf software. This is the strategy implemented in the popular package REBayes [KG17]. More recent methods based on general-purpose convex programming have also been proposed [KCSA20, WIM25].

The ubiquity and empirical success of these approximate algorithms for finding $\hat{\pi}$ raises natural questions. If the grid-based NPMLE convex optimization algorithm hits its stopping criterion and returns a 3-atomic solution π , can we provably convince ourselves that the true NPMLE $\hat{\pi}$ has 3 and not 1 or 100 atoms? More generally, given that an algorithm (heuristic or otherwise) appears to have approximately converged to some π , can we efficiently **certify** (prove) that $\hat{\pi}$ is ε -close to π ? Do any of the iterative algorithms converge to $\hat{\pi}$ at provably efficient rates? Because the objective (1.2) is poorly conditioned, off-the-shelf convex optimization theory does not provide such guarantees.

We provide answers to all of these questions as a consequence of a new structural property of $\hat{\pi}$, perhaps of independent interest. We show that under random data the true NPMLE solution $\hat{\pi}$ has a certain *generic behavior*. Namely, suppose $x_1, \dots, x_n$ are drawn IID¹ from an absolutely continuous distribution on $\mathbb{R}$; we then say X is **generic**. As recalled below, it has been known since Lindsay [Lin83a] that $\hat{\pi}$ is characterized by always local optimality conditions. Our results show these local optimality conditions for $\hat{\pi}$ are almost surely *strict* when X is generic, as long as $\max_i |x_i| \leq cn^{1/4}$ for a small absolute constant c .

This almost sure strictness has computational consequences, enabling us to derive the first provably efficient algorithms for $\hat{\pi}$. Although it is not difficult to approximately maximize the objective ℓ_X , or to ε -approximate $\hat{\pi}$ with exponential or worse dependence on ε in run-time, neither of these yields efficient approximation of $\hat{\pi}$ (e.g. polynomial dependence on $1/\varepsilon$). Using our genericity results, we obtain almost sure locally quadratic convergence in parameter distance for generic data. This implies the same convergence rate in Wasserstein distance and exact computation of the number of atoms in $\hat{\pi}$. Importantly, our algorithms certify an explicit Wasserstein error bound and exact number of atoms in finite time, so they have a well-defined output rather than just asymptotic convergence. Our algorithms follow a simple template: first, use a convex optimization algorithm such as Frank–Wolfe to compute a sparse approximate NPMLE $\hat{\pi}_\varepsilon$ which is supported on an ε -grid such as $\varepsilon\mathbb{Z}$. Next, merge any adjacent atoms (i.e. at distance ε) within the support of $\hat{\pi}_\varepsilon$. Finally, attempt to certify that the result is close enough to $\hat{\pi}$ that Newton’s method converges quadratically, based on the local concavity of ℓ . We prove that a version of this approach succeeds almost surely for generic X , once ε is sufficiently small.

As preparation, we recall the classical stationarity conditions characterizing $\hat{\pi}$. Define the function

¹Our genericity results also hold if $x_1, \dots, x_n$ are drawn independently from different probability densities. In fact this is an immediate corollary of the IID case since the conclusions hold almost surely: given probability densities $\mu_1, \dots, \mu_n$, any event that holds almost surely for IID $x_1, \dots, x_n \sim (\mu_1 + \dots + \mu_n)/n$ also holds almost surely for independent $x_i \sim \mu_i$.

$D_{\pi,X} : \mathbb{R} \rightarrow \mathbb{R}$ by

$$D_{\pi,X}(y) = \frac{1}{n} \sum_{i=1}^n \frac{e^{-|x_i-y|^2/2}}{P_\pi(x_i)\sqrt{2\pi}} \triangleq \frac{1}{n} \sum_{i=1}^n T_{\pi,x_i}(y). \quad (1.3)$$

$D_{\pi,X}$ is the derivative of ℓ with respect to perturbations in π , and plays a central role in the following characterization of $\hat{\pi}$ (see [Lin83a]).

Proposition 1.1. *For any π and π' , denoting $\pi_t = (1-t)\pi + t\pi'$, we have*

$$\mathbb{E}^{y \sim \pi} D_{\pi,X}(y) = 1, \quad (1.4)$$

$$\left. \frac{d\ell(\pi_t)}{dt} \right|_{t=0+} = \int D_{\pi,X}(y) d(\pi' - \pi)(y). \quad (1.5)$$

The minimizer $\hat{\pi}$ of (1.2) is unique and k -atomic for some $k \leq n$, and satisfies for all $y \in \text{supp}(\hat{\pi})$:

$$\begin{aligned} D_{\hat{\pi},X}(y) &= 1, \\ D'_{\hat{\pi},X}(y) &= 0 \end{aligned} \quad (1.6)$$

Moreover $D_{\hat{\pi},X}(y) \leq 1$ for all $y \in \mathbb{R}$, so that

$$\text{supp}(\hat{\pi}) \subseteq \arg \max(D_{\hat{\pi},X}). \quad (1.7)$$

We note that uniqueness of $\hat{\pi}$ is non-obvious. Indeed, from the strong concavity of the log it is easy to see that all maximizers of $\ell(\pi)$ have the same vector $P_\pi(X) \triangleq (P_\pi(x_1), \dots, P_\pi(x_n))$. However, the linear map $\pi \mapsto P_\pi(X)$ with domain being signed measures has infinite dimensional pre-image. The surprising part of Proposition 1.1 is that intersecting this pre-image with the subset of probability distributions yields a unique $\hat{\pi}$.

1.1 Structural Results

Our first result proves essentially that (with high probability) the converse implications of Lindsay conditions hold: for any $y \notin \text{supp}(\hat{\pi})$ we have

$$D_{\hat{\pi},X}(y) < 0$$

and for every $y \in \text{supp}(\hat{\pi})$ the point y is a non-degenerate maximum:

$$D'_{\hat{\pi},X}(y) = 0, \quad D''_{\hat{\pi},X}(y) < 0.$$

To that end we provide structural results showing that whenever X is generic, both $\hat{\pi}$ and $D_{\hat{\pi},X}$ also behave “generically”. Specifically, we prove that $\hat{\pi}$ admits a density on the natural parameter space, conditional on the number of atoms. Furthermore, one would like to think of $D_{\hat{\pi},X}$ as a generic smooth function, but in light of Proposition 1.1 it may have multiple global maxima, which is not generic for a smooth function. Theorem 1.2 parts (II),(III) essentially say that $D_{\hat{\pi},X}$ behaves generically otherwise, in that it has no “spurious” global maxima, and all maxima are well-conditioned.

Below we define Π_k to consist of all $\pi = \sum_{i=1}^k p_i \delta_{y_i}$ with exactly k distinct atoms, which can be parametrized as an open subset of $\mathbb{R}^{2k-1}$; see (1.15) for a more precise definition. We also define $\Pi_{k,\varepsilon} \subseteq \Pi_k$ to consist of those π with $\min_i p_i \geq \varepsilon$ and $\min_{i < j} |y_i - y_j| \geq \varepsilon$.

Theorem 1.2. Let $X = (x_1, \dots, x_n)$ be IID from an absolutely continuous distribution μ on $\mathbb{R}$. Then for any k such that $n > (2k + 2)^2$:

- (I) The restriction of the law of $\hat{\pi}$ to Π_k is absolutely continuous. Further if μ has density supported in $[-L, L]$ with values in $[0, L']$, then the density of $\hat{\pi}$ is locally bounded, i.e. at most $C(n, k, \varepsilon, L, L')$ on $\Pi_{k, \varepsilon}$ for any $\varepsilon > 0$.
- (II) Conditional on $|\text{supp}(\hat{\pi})| = k$ (assuming k is such that $\mathbb{P}[|\text{supp}(\hat{\pi})| = k] > 0$), the function $D_{\hat{\pi}, X}$ almost surely has exactly k global maxima:

$$\text{supp}(\hat{\pi}) = \arg \max_{y \in \mathbb{R}} D_{\hat{\pi}, X}(y). \quad (1.8)$$

- (III) Conditional on $|\text{supp}(\hat{\pi})| = k$ (assuming k is such that $\mathbb{P}[|\text{supp}(\hat{\pi})| = k] > 0$), each global maximum of $D_{\hat{\pi}, X}$ is almost surely non-degenerate:

$$\max_{y \in \text{supp}(\hat{\pi})} D''_{\hat{\pi}, X}(y) < 0. \quad (1.9)$$

Using this result, we obtain further consequences on the local landscape of ℓ_X near $\hat{\pi}$.

Theorem 1.3. For $1 \leq k \leq n$, if the conclusions of Theorem 1.2 (II), (III) hold, then:

- (A) There is an open neighborhood of $\hat{\pi}$ in Π_k on which ℓ_X is locally c -strongly concave for some $c > 0$.
- (B) There is an open neighborhood of X in $\mathbb{R}^n$, such that $\hat{\pi}(\hat{X}) \in \Pi_k$ for all $\hat{X}$ in this neighborhood. Moreover, $\hat{\pi}$ is a smooth function on this neighborhood.
- (C) The expectation-maximization (EM) algorithm converges linearly in a small Π_k -neighborhood of $\hat{\pi}$. Namely for any initialization π_0 in this neighborhood, the EM iterates $\pi_0, \pi_1, \dots$ satisfy $d_{\Pi_k}(\pi_t, \hat{\pi}) \leq C \cdot (1 - \eta)^t d_{\Pi_k}(\pi_0, \hat{\pi})$ for $C, \eta > 0$. (Here d_{Π_k} , defined in (1.16), denotes parameter distance in Π_k .)

We emphasize that Theorem 1.2(II) improves the classical stationarity condition (1.7) to a strict inequality. Additionally, Theorem 1.3(A) is different from the concavity of ℓ_X on the space of probability measures, since the structure of Π_k allows averaging the locations of atoms. We also mention that Appendix C shows IID data with full support on $[0, O(k\sqrt{\log k})]$ satisfies $\mathbb{P}[|\text{supp}(\hat{\pi})| = k] > 0$ whenever $n \geq k$, so the conditional distribution of $\hat{\pi}$ on Π_k is defined.

For Theorem 1.2 to be useful, we must have $k = |\text{supp}(\hat{\pi})| \leq O(\sqrt{n})$, which is not always the case. However, following a related conjecture of [KG19], the first author showed in [PW20] with Yihong Wu that the much stronger bound $|\text{supp}(\hat{\pi})| \leq O(\log n)$ holds with high probability for sub-Gaussian data (and is sharp in natural examples). In fact the bound depends deterministically on the empirical range of the data.

Proposition 1.4 ([PW20, Theorem 1]). *There exists a universal constant C such that if $X = (x_1, \dots, x_n) \in [-L, L]^n$ for $L \geq 1$, then $|\text{supp}(\hat{\pi})| \leq CL^2$.*

Corollary 1.5 below is an immediate consequence: $\hat{\pi}$ exhibits generic behavior for generic data in $[-cn^{1/4}, cn^{1/4}]$. For example, (1.7) is almost surely an equality for such X .

Corollary 1.5. *Suppose $X = (x_1, \dots, x_n)$ is IID from an absolutely continuous distribution. Then with probability 1, either both the conclusions (1.8), (1.9) of Theorem 1.2 (II), (III) hold, or otherwise $\max_i |x_i| \geq cn^{1/4}$ holds for some absolute constant $c > 0$.*

In particular if $\sup_i \mathbb{E}[|x_i|^q] < \infty$ for some $q > 4$, then (1.8), (1.9) hold with probability $1 - O(n^{1-\frac{q}{4}})$.

Remark 1.6. The genericity conditions above can be relaxed to apply only to part of the dataset. This is relevant in settings where some datapoints have been rounded or are correlated with each other. Namely, fix arbitrary $x_{m+1}, \dots, x_n \in [-L, L]$. Then if $m \geq CL^4$ and $x_1, \dots, x_m \in [-L, L]$ are generic conditionally on $(x_{m+1}, \dots, x_n)$, all the conclusions of Theorem 1.2 still apply with essentially the same proof.

In the setting of Remark 1.6, we show in the following Theorem 1.7 that Theorem 1.2(I) is sharp in some sense. Namely with $m \leq O(k^2)$ generic datapoints, the density of $\hat{\pi}$ on Π_k may fail to be locally bounded. The proof is given in Appendix C, where we also prove Theorem 1.3. It is an interesting open problem whether Theorem 1.7 truly requires $m < n$.

Theorem 1.7. *Suppose $m \leq (k-1)(2k-1)$ and $n-m \geq (2k+2)^2$. Let $x_1, \dots, x_n$ be IID uniform on $[0, O(k\sqrt{\log k})]$. Then with positive probability, the conditional law of $\hat{\pi}$ given $x_{m+1}, \dots, x_n$ and the event $\hat{\pi} \in \Pi_k$ does not have locally bounded density on Π_k .*

1.2 Efficient Computability of $\hat{\pi}$

A related and somewhat subtle question is whether $\hat{\pi}$ is efficiently computable. Because ℓ is concave, one can find in $\text{poly}(1/\varepsilon)$ time a probability distribution π_ε such that $|\ell_X(\pi_\varepsilon) - \ell_X(\hat{\pi})| \leq \varepsilon$ (as recently observed in [FGSW23]). Uniqueness of $\hat{\pi}$ implies the convergence $\pi_\varepsilon \rightarrow \hat{\pi}$ in the space of probability measures on $\mathbb{R}$ as $\varepsilon \rightarrow 0$. However ℓ_X may be ill-conditioned, so it is not clear that one obtains a quantitative rate of convergence of π_ε to $\hat{\pi}$. Worse, this proof does not give any *stopping rule* at which one can *guarantee* that π_ε is within ε distance from $\hat{\pi}$. From the point of view of computational complexity theory, this means one does not yet have an actual algorithm to compute $\hat{\pi}$.

We will aim to compute $\hat{\pi}$ in the strong *parameter distance* in which all centers and mixing weights are well-approximated. This includes the following two questions:

1. Can $\hat{\pi}$ be certifiably approximated in Wasserstein distance to error ε in $\varepsilon^{-O(1)}$ time?
2. Is the support size $|\text{supp}(\hat{\pi})|$ a computable function of $X = (x_1, \dots, x_n)$?

We have illustrated that the first question is non-obvious, but the second merits its own discussion. Note that from a Wasserstein approximation to $\hat{\pi}$, one can *never* directly conclude the value of $|\text{supp}(\hat{\pi})|$: any atom may split onto several smaller atoms with negligible effect in Wasserstein distance. Although support size is a very brittle property, Proposition 1.4 shows a strong upper bound on $|\text{supp}(\hat{\pi})|$, so it is natural to hope that exact computation is feasible. This task requires understanding the distributional behavior of $\hat{\pi}$ at microscopic scales, a challenging task which has seen no previous work. Intuitively, it is natural to expect those $\hat{\pi}$ exhibiting potential instability to correspond to a low-dimensional submanifold in parameter space, which should not occur for generic X . Formalizing this intuition is challenging due to the multi-scale nature of the parameter space $\cup_{k \geq 1} \Pi_k$ for $\hat{\pi}$.

To formulate the computational problem in a satisfactory way, we use the notion of a Shub–Smale approximate solution to a system of equations, originally due to [SS93]. In short, π is a Shub–Smale approximate NPMLE if it is within ε parameter distance of $\hat{\pi}$, and Newton’s method exhibits quadratic convergence to $\hat{\pi}$ starting from π . To be precise, starting from any $\pi \in \Pi_k$, one may naturally define a Newton–Raphson iteration to attempt to compute $\hat{\pi} \in \Pi_k$ (viewing Π_k as an open subset of $\mathbb{R}^{2k-1}$). Recalling Proposition 1.1, the point is to view $\hat{\pi}$ as the zero of the function $\gamma_k : (p_1, \dots, p_k, y_1, \dots, y_k) \in \Pi_k \rightarrow \mathbb{R}^{2k-1}$ defined by

$$\gamma_k(\pi) = (D_\pi(y_1) - 1, \dots, D_\pi(y_k) - 1, D'_\pi(y_1), \dots, D'_\pi(y_k)).$$

(γ_k is essentially the gradient of ℓ_X in the space Π_k .) The iteration begins with $\pi^{(0)} = \pi$ and (with J the Jacobian) then recursively sets:

$$\pi^{(t+1)} = \pi^{(t)} - [J\gamma_k(\pi^{(t)})]^{-1}\gamma_k(\pi^{(t)}). \quad (1.10)$$

Since Π_k is a proper subset of $\mathbb{R}^{2k-1}$, it is possible that even if $\pi^{(t)} \in \Pi_k$ and $\hat{\pi} \in \Pi_k$, the next iterate $\pi^{(t+1)} \notin \Pi_k$ may satisfy $p_i < 0$ or $y_i > y_{i+1}$ for some i . If this occurs, we stop the iteration and declare failure.

Definition 1. We say $\pi \in \Pi_k$ is a Shub–Smale ε -approximate NPMLE if:

1. $\hat{\pi} \in \Pi_k$ and $d(\pi^{(0)}, \hat{\pi})_{\Pi_k} \leq \varepsilon$.
2. The iteration (1.10) satisfies $\pi^{(t)} \in \Pi_k$ for all t (and never declares failure).
3. We have $d(\pi^{(t)}, \hat{\pi})_{\Pi_k} \leq 2^{1-2^t} d(\pi^{(0)}, \hat{\pi})_{\Pi_k}$ for all t .

Here we use the parameter distance $d(\cdot, \cdot)_{\Pi_k}$ defined in (1.16), which is simply Euclidean distance on the parametrization $\pi = (p_1, \dots, p_k, y_1, \dots, y_k)$.

Note that the value of ε in Definition 1 is not particularly important: once the condition holds for some $\varepsilon < 1/2$, applying Newton’s method converges quadratically. Our computational results use the genericity of Theorem 1.2 to find a Shub–Smale approximate NPMLE almost surely. We note that the issue of an unknown number of parameters k is not present in [SS93]. However this arises naturally in our setting because $|\text{supp}(\hat{\pi})|$ must be known exactly for a Newton–Raphson iteration to make sense. Indeed, it underlines the point that exact computation of $\hat{\pi}$ inherently **requires** exact computation of the support size.

1.3 Naïve Brute-Force Approximation of $\hat{\pi}$

To further motivate and illustrate our algorithmic results, we discuss two naïve algorithms to approximate $\hat{\pi}$ in Wasserstein distance. The first algorithm is essentially given by the next proposition. Here and throughout the rest of the paper, for $L, \varepsilon > 0$ we let $Z_\varepsilon \subseteq \mathbb{R}$ be any set satisfying:

$$\begin{aligned} \{-L, L\} &\subseteq Z_\varepsilon \subseteq [-L, L], \\ |Z_\varepsilon| &\leq 3L/\varepsilon, \\ \min_{y \in [-L, L]} d(y, Z_\varepsilon) &\leq \varepsilon. \end{aligned} \quad (1.11)$$

The sets Z_ε will serve as supports for approximations $\hat{\pi}_\varepsilon$ to $\hat{\pi}$.²

Proposition 1.8. *For each $\varepsilon > 0$, let*

$$\hat{\pi}_\varepsilon \in \arg \max_{\text{supp}(\pi) \subseteq Z_\varepsilon} \ell_X(\pi)$$

maximize the log-likelihood for X among all distributions supported in Z_ε . Then

$$\lim_{\varepsilon \rightarrow 0} \hat{\pi}_\varepsilon = \hat{\pi},$$

where convergence is in the Wasserstein $\mathbb{W}_1$ sense.

²Throughout the paper, one can just take $Z_\varepsilon = \varepsilon\mathbb{Z} \cap [-L, L]$ for any $L, \varepsilon^{-1} \in \mathbb{N}$. However in practical algorithms, the atoms of an approximating $\hat{\pi}_\varepsilon$ are often adjusted during the optimization. Augmenting Z_ε with these additional points allows us to cover such situations.

Proof. As we show in Lemma 1.18, the map $\pi \mapsto \ell_X(\pi)$ from $\mathbb{W}_1([-L, L]) \rightarrow \mathbb{R}$ has Lipschitz constant at most $C_L \leq O(e^{4L^2})$. It follows easily that

$$\ell_X(\hat{\pi}_\varepsilon) \leq \ell_X(\hat{\pi}) + C_L \varepsilon. \quad (1.12)$$

Since the sets Z_ε are uniformly bounded, there must exist subsequential limits of $\hat{\pi}_\varepsilon$ as $\varepsilon \rightarrow 0$. Since the functional $\pi \mapsto \ell_X(\pi)$ is weakly continuous, all subsequential limits globally minimize ℓ_X . Uniqueness of $\hat{\pi}$ (see Proposition 1.1) completes the proof. $\square$

It also follows from (1.12) that $\ell_X(\pi)$ is $C_L \varepsilon^{-1}$ -Lipschitz when π is metrized by the Euclidean norm on $\mathbb{R}^{|Z_\varepsilon|}$ (via its probability mass function). Since $\pi \mapsto \ell_X(\pi)$ is concave for each X , given a choice of Z_ε , computation of $\hat{\pi}_\varepsilon$ to $\mathbb{W}_1$ to accuracy δ requires $C_L \varepsilon^{-O(1)} \log(1/\delta)$ gradient evaluations of ℓ_X using a cutting plane method (see e.g. [Bub15, Theorem 2.4]).

Unfortunately the algorithm “compute $\hat{\pi}_\varepsilon$ for small $\varepsilon > 0$ ” gives *no quantitative guarantees* for the approximation of $\hat{\pi}$ itself. That is, the above argument cannot certify upper bounds $\mathbb{W}_1(\hat{\pi}_\varepsilon, \hat{\pi}) \leq \eta$ for any nontrivial $\eta > 0$. Indeed, $\hat{\pi}_\varepsilon$ above is essentially the maximum of the concave function $\ell_X : \mathbb{W}_1([-L, L]) \rightarrow \mathbb{R}$ on the $2L/\varepsilon$ dimensional subspace of π supported in Z_ε . Although $\hat{\pi}$ is very close to this subspace, a *strong* concavity estimate would be needed to upper bound $\mathbb{W}_1(\hat{\pi}_\varepsilon, \hat{\pi})$ from this. For example the concave function

$$f(x_1, \dots, x_d) = -\left(x_1 - \frac{x_2}{10^4}\right)^2 - \frac{1}{d} \sum_{i \in [d] \setminus \{2\}} x_i^2$$

is maximized at the origin, but its minimizer on the nearby hyperplane $\{x_1 = 0.01\}$ has $x_2 = 100$.

A second try is to explicitly search the entirety of $\mathbb{W}_1([-L, L])$. Namely if $\mathcal{N}_\varepsilon \subseteq \mathbb{W}_1([-L, L])$ is an $C_L^{-1} \varepsilon$ -net, then the choice

$$\hat{\pi}_\varepsilon = \arg \min_{\pi \in \mathcal{N}_\varepsilon} \ell_X(\pi)$$

also satisfies (1.12). Moreover **if** for some $\eta > 0$ it happens to be the case that

$$\min_{\substack{\pi \in \mathcal{N}_\varepsilon \\ \mathbb{W}_1(\pi, \hat{\pi}_\varepsilon) \geq \eta - \varepsilon}} \ell_X(\pi) \geq \ell_X(\hat{\pi}_\varepsilon) + 3\varepsilon, \quad (1.13)$$

then this would immediately certify the bound $\mathbb{W}_1(\hat{\pi}_\varepsilon, \hat{\pi}) \leq \eta$. Therefore a natural approach to certifiable approximation guarantees would be to verify the estimate in (1.13). However a bit of thought reveals this approach is also impractical. Firstly, it is again non-quantitative: uniqueness of $\hat{\pi}$ implies (1.13) holds eventually (i.e. for $0 < \varepsilon \leq \varepsilon_0(\eta)$ sufficiently small depending on $\eta > 0$), but this argument does not predict any quantitative dependence between ε and η . Additionally, $|\mathcal{N}_\varepsilon|$ grows exponentially³ in $1/\varepsilon$, making this computationally inefficient even if η and ε turn out to be polynomially related.

1.4 Results on Computability and Generic Behavior

Having discussed several pitfalls, we now present our main computability results, which show that $\hat{\pi}$ can be efficiently approximated and the support size k can be exactly computed. We equip Z_ε (recall (1.11)) with the adjacent-neighbors graph structure, making it isomorphic to a path. Below we write $O_X(\cdot)$ to indicate an implicit constant factor which is random and depends on X , but not on e.g. ε .

³For example let $S_1, \dots, S_K \subseteq Z_\varepsilon = [-L, L] \cap \varepsilon\mathbb{Z}$ be IID uniformly random subsets of size $|Z_\varepsilon|/2$, and let π_i be the uniform distribution on S_i for each i . For an absolute constant $c > 0$ and $K \leq c \exp(c/\varepsilon)$, we will have $|S_i \cap S_j| \leq |Z_\varepsilon|/3$ and hence $\mathbb{W}_1(\pi_i, \pi_j) \geq \varepsilon/10$ with high probability, simultaneously for all $1 \leq i < j \leq K$.

Theorem 1.9. Assume n, L satisfy $n \geq CL^4$ for an absolute constant C . Let $X = (x_1, \dots, x_n) \in [-L, L]^n$ be generic. Let $\hat{\pi}_\varepsilon = \arg \max_{\text{supp}(\pi) \subseteq Z_\varepsilon} \ell_X(\pi)$. Then almost surely, for ε small enough depending on X :

- (a) All connected components of $\text{supp}(\hat{\pi}_\varepsilon) \subseteq Z_\varepsilon$ (w.r.t. the graph structure) have size 1 or 2.
- (b) $|\text{supp}(\hat{\pi})|$ is equal to the number of connected components of $\text{supp}(\hat{\pi}_\varepsilon)$.
- (c) $\mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/3})$,

Further, the statement in (b) and the upper bound in (c) are efficiently **certifiable** given $\hat{\pi}_\varepsilon$. Finally, the rounding of $\hat{\pi}_\varepsilon$ which merges adjacent atoms $p_j \delta_{y_j} + p_{j+1} \delta_{y_{j+1}}$ in Z_ε to $(p_j + p_{j+1}) \delta_{\frac{p_j y_j + p_{j+1} y_{j+1}}{p_j + p_{j+1}}}$ satisfies $d_{\Pi_k}(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/3})$ and is certifiably a Shub–Smale approximate NPMLE.

In proving Theorem 1.9, we consider a slight modification $\tilde{\pi}_\varepsilon$ of $\hat{\pi}_\varepsilon$ in which adjacent atoms are merged together. This in fact gives small error in the stronger parameter distance d_{Π_k} (defined in (1.16)). Since the exact maximization defining $\hat{\pi}_\varepsilon$ may not be computationally feasible, we provide another variant of Theorem 1.9 that does not require exact computation of $\hat{\pi}_\varepsilon$ (proved in the Appendix). Instead only an approximation is required, obtained using the Frank–Wolfe algorithm and a careful rounding scheme; we denote by $(\tilde{\pi}_\varepsilon, \hat{\pi}_\varepsilon)$ the corresponding analogs of $(\hat{\pi}_\varepsilon, \tilde{\pi}_\varepsilon)$.

Theorem 1.10. Under the conditions of Theorem 1.9, there exists for each $\varepsilon > 0$ a deterministic $O(Ln\varepsilon^{-11})$ time algorithm which computes $\hat{\pi}_\varepsilon \in \mathcal{P}(Z_\varepsilon)$ such that almost surely, for ε small enough depending on X :

- (a) $|\text{supp}(\hat{\pi})| = |\text{supp}(\hat{\pi}_\varepsilon)|$.
- (b) $d_{\Pi_k}(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/4})$ and $\mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/4})$.

Further, the statement in (a) and the upper bound in (b) are efficiently certifiable given $\hat{\pi}_\varepsilon$, which is a Shub–Smale approximate NPMLE.

In other words, Theorem 1.10 provides a pair of efficient algorithms. The first computes a $\mathbb{W}_1$ -approximation of $\hat{\pi}$, while the second attempts to certify the $\mathbb{W}_1$ bound and support size equality. The second either returns a checkable proof or fails. What we show is that by rerunning these algorithms with smaller and smaller ε , eventually the second algorithm will succeed; furthermore, the bound will decay as $\varepsilon^{1/4}$. Once the algorithm succeeds for some $\varepsilon = \varepsilon_0$, there is no longer a need to continue rerunning the same algorithm to decrease ε . Instead, one can simply run Newton’s method within Π_k . Hence for $\delta \ll \varepsilon_0$, the computational complexity will be $O(Ln\varepsilon_0^{-11}) + C_{\text{NR}} \log \log(\varepsilon_0/\delta)$, where $C_{\text{NR}} \leq O(nk^2)$ is the complexity of a Newton–Raphson iteration. (Computing the gradient of γ_k uses nk^2 time, inverting an $O(k) \times O(k)$ matrix takes $O(k^3) \leq O(nk^2)$ time, and other steps are faster.)

Let us explain briefly why Theorem 1.2(II) and (III) are useful towards Theorem 1.9. The chief worry in Theorem 1.9 is that although the NPMLE objective (1.2) is a concave maximization problem, it is infinite-dimensional and may be quite poorly conditioned. Because the log-likelihood can be shown to be relatively smooth, **if** it is not flat near $\hat{\pi}$, **then** one will be able to efficiently certify $\hat{\pi} \approx \tilde{\pi}$ based on local information at $\tilde{\pi}$ (namely approximate-stationarity and Hessian non-singularity). In Section 3, we employ Theorem 1.2(II) and (III) to show the necessary conditions hold once $\tilde{\pi}$ is a sufficiently accurate approximation for $\hat{\pi}$.

While Newton–Raphson iteration is appealing due to its quadratic local convergence rate, other approaches also suffice for asymptotic convergence from an approximate solution. In particular Theorem 1.3(C) shows that the EM algorithm converges linearly from suitable approximate solutions; this can be similarly made certifiable from a sufficiently good approximate solution $\hat{\pi}_\varepsilon$.

1.5 Results for the Static Support NPMLE

To illustrate the flexibility of our methods, we also consider the *static support* NPMLE. Given a fixed finite set $S \subseteq \mathbb{R}$ (independent of X), we define the static support NPMLE $\hat{\pi}_S$ as in (1.2), but restricted to π supported within S . Similarly let $\Pi_k(S)$ be the $(k-1)$ -dimensional set of k -atomic distributions on S . The following analog of Proposition 1.1 is immediate from (1.5).

Proposition 1.11. *Fix a finite $S \subseteq \mathbb{R}$ and $X = (x_1, \dots, x_n)$. Then any minimizer $\hat{\pi}_S$ of (1.2) among π supported on S satisfies for all $y \in \text{supp}(\hat{\pi}_S)$:*

$$D_{\hat{\pi}_S, X}(y) = 1.$$

Moreover $D_{\hat{\pi}_S, X}(y) \leq 1$ for all $y \in S$, so that

$$\text{supp}(\hat{\pi}_S) \subseteq \arg \max_S (D_{\hat{\pi}_S, X})$$

Uniqueness does not appear to follow from classical results; the proof in [Lin83b, Theorem 5.1] requires $\text{supp}(\hat{\pi})$ to avoid the boundary of S , which is of course impossible when S is discrete. Thus in principle, $\hat{\pi}_S$ refers to any maximizer of ℓ_X . However for bounded data, uniqueness follows by the technique of [PW20].

Proposition 1.12. *Let $L \leq cn^{1/2}$ for a small absolute constant c . Fix a L -bounded $X = (x_1, \dots, x_n)$, and finite $S \subseteq \mathbb{R}$ with $S \cap [-3L, 3L] \neq \emptyset$. Then $\hat{\pi}_S$ is unique and $|\text{supp}(\hat{\pi}_S)| \leq O(L^2)$.*

Proof. Let $\hat{\pi}_S$ be any minimizer of (1.2) among probability distributions supported in S . It is clear that $\hat{\pi}_S$ is supported in $[-10L, 10L]$. It follows from [PW20, Proof of Theorem 3] that $D_{\hat{\pi}_S, X}$ has $O(L^2)$ critical points. Since $D_{\hat{\pi}_S, X}$ takes the value 1 at each point in S , Rolle's theorem implies that $|\text{supp}(\hat{\pi}_S)| \leq O(L^2) < n/2$ (for c small). Since this support bound holds for all minimizers $\hat{\pi}_S$, [Lin83b, Lemma 6.1] implies $\hat{\pi}_S$ is unique. $\square$

Our main results on generic behavior and efficient computability both have analogs for finite S ; the proofs are similar to those of the results presented so far.

Theorem 1.13. *Fix $n \geq k^2$ and a finite set $S \subseteq \mathbb{R}$. Let X be generic and condition on the event $|\text{supp}(\hat{\pi}_S)| = k$ (assuming k is such that $\mathbb{P}[|\text{supp}(\hat{\pi}_S)| = k] > 0$). Then with $\{y_1, \dots, y_k\} = \text{supp}(\hat{\pi})$:*

- (I) *The conditional law of $\hat{\pi}_S$ on $\Pi_k(S)$ is absolutely continuous.*
- (II) *The restriction of $D_{\hat{\pi}_S, X}$ to S almost surely has exactly k global maxima $y_1, \dots, y_k$.*

Theorem 1.14. *For any $1 \leq k \leq n$, if $|\text{supp}(\hat{\pi}_S)| = k$ and Theorem 1.13(II) holds:*

- (A) *There is $c > 0$ and an open neighborhood of $\hat{\pi}_S$ in $\Pi_k(S)$ on which ℓ_X is locally c -strongly concave.*
- (B) *There is an open neighborhood of X in $\mathbb{R}^n$, such that $\hat{\pi}_S(\hat{X}) \in \Pi_k(S)$ for all $\hat{X}$ in this neighborhood. Moreover, $\hat{\pi}_S$ is a smooth function on this neighborhood.*

Theorem 1.15. *Assume n, L satisfy $n \geq CL^4$ for an absolute constant C , and let $S \subseteq \mathbb{R}$ be deterministic and finite with $S \cap [-3L, 3L] \neq \emptyset$. Let $X = (x_1, \dots, x_n) \in [-L, L]^n$ be generic. Then there exists a deterministic $C(L)\varepsilon^{-O(1)}$ time algorithm which computes $\tilde{\pi}_{S, \varepsilon}$ with the following properties. Almost surely, for sufficiently small ε , $\tilde{\pi}_{S, \varepsilon}$ is a certifiable $C(n, k, S, X)\varepsilon$ -approximation to $\hat{\pi}_S$ in both $\mathbb{W}_1$ and $d_{\Pi_k(S)}$ distance. Furthermore, $\text{supp}(\tilde{\pi}_{S, \varepsilon}) = \text{supp}(\hat{\pi}_S)$ almost surely holds certifiably for small enough ε .*

Finally, almost surely for small enough ε , $\tilde{\pi}_{S, \varepsilon}$ is a Shub–Smale approximate static support NPMLE: Newton's method within $\Pi_k(S)$ started from $\tilde{\pi}_{S, \varepsilon}$ certifiably converges to $\hat{\pi}_S$ at the rate in Definition 1.

1.6 Other Related Work

Gaussian mixture models have been studied since the pioneering work of Pearson [Pea94], which proposed that the ratio of forehead width to body length of crabs might follow such a distribution. Much work has focused on statistical recovery of such mixtures. In the 1-dimensional Gaussian location model we consider, optimal convergence rates for recovering the mixing distribution were obtained in the case of k components by [WY20], via an extension of the method of moments. See [DWYZ23] for extensions to higher dimensions. The theoretical computer science literature has also studied Gaussian mixture models since [Das99]. In the special case $k = 2$, [HP15] gave sharp bounds for parameter recovery, and showed an exponential-in- k sample complexity lower bound. This line of work led to accurate polynomial time estimators for the underlying parameters of the Gaussian mixture even in high-dimensions, under minimal assumptions that ensure statistical identifiability. These algorithms succeed even if the components may have different covariances [KMV10, MV10, HP15] and more recently if a small fraction of the data is adversarially corrupted [Kan21, LM23, BDJ⁺22, LM22] thanks to the sum-of-squares framework.

The NPMLE was introduced for general mixture models in [KW56], where its consistency was shown in quite general settings including the one we study. [GW00, GVDV01, Zha09, JZ09] upper-bounded its rate of convergence for density-estimation, for IID data generated from a mixture of unit variance Gaussians. See also [DZ16, SG20] for higher-dimensional extensions. We emphasize that we always metrize convergence based on $\hat{\pi}$ itself, rather than the convolution $\hat{\pi} * \mathcal{N}(0, 1)$ which is done in some of these works. (This yields e.g. the *smoothed* Wasserstein distance, which gives the same topology on probability measures but can be exponentially smaller.) However we emphasize that by contrast to the algorithms mentioned above, the NPMLE’s behavior can be fruitfully analyzed without assuming that the underlying data actually comes from a Gaussian mixture.

To estimate a k -component Gaussian mixture for small k , a standard approach is the expectation-maximization (EM) algorithm [DLR77]. However a key advantage of the NPMLE is that it solves a **concave** maximization problem. To take full advantage of this, one may discretize space as in [KM14, FD18, SGS24]. Namely one fixes an ε -net Z_ε and optimizes (1.2) subject to the additional constraint $\text{supp}(\hat{\pi}) \subseteq Z_\varepsilon$, which is now a finite-dimensional concave problem. It is not hard to show (see Section 3) that the resulting estimate $\hat{\pi}_\varepsilon$ converges to the true NPMLE $\hat{\pi}$ as $\varepsilon \rightarrow 0$. Our interest will be in *certifying* that a candidate $\hat{\pi}_\varepsilon$, computed in essentially arbitrary manner, approximates $\hat{\pi}$ to some explicit accuracy and additionally satisfies $|\text{supp}(\hat{\pi}_\varepsilon)| = |\text{supp}(\hat{\pi})|$.

The local convergence rate of the EM algorithm has long been of interest, and was studied in e.g. [RW84, MR94, XJ96, MXJ00, MF05, MX05]. More recently, [BWY17, ZLS20] established high-probability linear local convergence rates for well-separated Gaussian mixtures in general dimension, via perturbative analysis around the population dynamics. Theorem 1.3(C) shows almost sure local linear convergence for generic datasets, without even requiring the existence of an underlying mixture model generating the data. On the other hand, it is currently limited to dimension 1 and does not give quantitative bounds on η . Among the vast literature in this direction, we also mention a few recent works [DTZ17, WZ22, WB22] showing rapid *global* convergence of the EM algorithm for 2-component Gaussian mixtures.

We also mention the recent work [WN22] which considers rather general mixture models and uses harmonic analysis tools related to those we employ (see Section 7 therein). However their work focuses on asymptotic posterior contraction rates and uses these tools differently. Also recently, [MSS23, FGSW23] investigated properties of and algorithms for the NPMLE in high-dimensional regression, which is different from the present setting. Theorem 3.3 of the latter also provides a quantitative convergence rate for computing the NPMLE (in the same setting as the present work) via gradient flow, as measured by the log-likelihood objective (1.2). As outlined in Subsection 1.3, this does not yield any convergence rate for the NPMLE itself, while our Theorems 1.9, 1.10 give asymptotic convergence rates in Wasserstein distance.

1.7 Notations and Spaces of Measures

Denote by $\mathcal{P}(X)$ the space of probability measures on a space X , and by $\rho_\nu(x)$ the density function of an absolutely continuous distribution ν at x . Let $d_{\mathcal{H}}$ denote the Hausdorff distance between compact sets in $\mathbb{R}$:

$$d_{\mathcal{H}}(K_1, K_2) = \max \left(\max_{x_1 \in K_1} d(x_1, K_2), \max_{x_2 \in K_2} d(K_1, x_2) \right). \quad (1.14)$$

For $\pi, \pi' \in \mathcal{P}(\mathbb{R})$, denote by $\mathbb{W}_1(\pi, \pi')$ the usual Wasserstein-1 distance

$$\mathbb{W}_1(\pi, \pi') = \inf_{\Gamma \in \mathcal{C}(\pi, \pi')} \mathbb{E}^\Gamma |y - y'|.$$

Here the infimum is over all couplings $(y, y') \sim \Gamma$ with marginals $y \sim \pi$ and $y' \sim \pi'$. Define Π_k to be $2k - 1$ dimensional space of k -atomic π . We use the parametrization

$$\Pi_k \triangleq \left\{ (p_1, \dots, p_k, y_1, \dots, y_k) \in \mathbb{R}^{2k} : p_j \geq 0, \sum_{j=1}^k p_j = 1, y_1 < y_2 < \dots < y_k \right\}, \quad (1.15)$$

thus identifying $(p_1, \dots, p_k, y_1, \dots, y_k)$ with the probability distribution $\sum_{j=1}^k p_j \delta_{y_j}$. As previously mentioned, we let $\Pi_{k,\varepsilon}$ consist of all $\pi \in \Pi_k$ with $\min_i p_i \geq \varepsilon$ and $\min_{i < j} |y_i - y_j| \geq \varepsilon$. Similarly Π_k^L consists of all $\pi \in \Pi_k$ supported in $[-L, L]$, and $\Pi_{k,\varepsilon}^L = \Pi_{k,\varepsilon} \cap \Pi_k^L$. Note that (1.15) gives Π_k the structure of a smooth $2k - 1$ dimensional manifold. We metrize Π_k and its subsets via parameter distance:

$$d_{\Pi_k}(\pi, \pi')^2 = \sum_{j=1}^k (|p_j - p'_j|^2 + |y_j - y'_j|^2) \quad (1.16)$$

where π, π' correspond respectively to $(p_1, \dots, p_k, y_1, \dots, y_k), (p'_1, \dots, p'_k, y'_1, \dots, y'_k)$. Note that as a subset of $\mathbb{R}^{2k}$, (the closure of) Π_k is a codimension 1 convex polytope, hence has a natural $(2k - 1)$ -dimensional Lebesgue probability measure. The next proposition states that Π_k^L parameter distance controls $\mathbb{W}_1$ distance. There is no bound in the opposite direction because far away atoms with small probability are significant only for the former.

Proposition 1.16. *For any k -atomic $\pi, \pi' \in \Pi_k^L$,*

$$\mathbb{W}_1(\pi, \pi') \leq (L^{3/2} + 1) d_{\Pi_k}(\pi, \pi').$$

1.8 Preliminary Smoothness Estimates

For $j \geq 0$ let $H_j(t) = e^{t^2/2} \left(\frac{d}{dt} \right)^j e^{-t^2/2}$ be the j -th Hermite polynomial. The following proposition is immediate from the formula (1.3).

Proposition 1.17. *For $j \geq 0$, the j -th derivative of $D_{\pi,X}$ is given by*

$$D_{\pi,X}^{(j)}(y) = \frac{1}{n} \sum_{i=1}^n \frac{H_j(x_i - y) e^{-|x_i - y|^2/2}}{P_\pi(x_i) \sqrt{2\pi}}. \quad (1.17)$$

We next show basic estimates on these functions.

Lemma 1.18. *There exist universal constants $C > 0$ and $(C_1, C_2, \dots)$ such that the following estimates hold. Suppose $\mathbb{W}_1(\pi, \pi') \leq \delta$ where π, π' are supported in $[-L, L]$. Then for $x, z \in [-L, L]$:*

1. $C^{-1}e^{-2L^2} \leq P_\pi(x) \leq 1$.
2. $|P_\pi(x) - P_{\pi'}(x)| \leq C\delta$.
3. $|P_\pi(x)^{-1} - P_{\pi'}(x)^{-1}| \leq Ce^{4L^2}\delta$.
4. $|\ell_X(\pi) - \ell_X(\pi')| \leq Ce^{2L^2}\delta$.
5. For $j \geq 0$ the j -th derivative of T (recall (1.3)) with respect to y satisfies

$$\begin{aligned} |T_{\pi,x}^{(j)}(z)| &\leq C_j L^j e^{2L^2}, \\ |T_{\pi,x}^{(j)}(z) - T_{\pi',x}^{(j)}(z)| &\leq C_j L^j e^{4L^2} \delta. \end{aligned} \tag{1.18}$$

6. Similarly for D :

$$|D_{\pi,X}^{(j)}(z)| \leq C_j L^j e^{2L^2}, \tag{1.19}$$

$$|D_{\pi,X}^{(j)}(z) - D_{\pi',X}^{(j)}(z)| \leq C_j L^j e^{4L^2} \delta. \tag{1.20}$$

Lemma 1.19. Fix $\pi = \sum_{i=1}^k p_j \delta_{y_i}$ and $\pi' = \sum_{i=1}^k q_i \delta_{y_i}$, both in Π_k^L . Consider

$$\ell(t) \triangleq \ell_X(\pi_t) = \ell_X((1-t)\pi + t\pi')$$

If $\max_i |p_i - q_i| \leq \tau$, then

$$\sup_{0 \leq t \leq 1} \left| \frac{d}{dt} \ell(t) \right| \leq O(e^{2jL^2} \tau^j), \quad 0 \leq j \leq 3.$$

Further if $\pi_t([-10, 10]) \geq 1/10$, the improved upper bound $O(e^{0.51jL^2} \tau^j)$ holds.

The next lemma will be used to ensure that $\hat{\pi}$, or any reasonable approximation thereof, is supported in $[-L, L]$. Lemmas 1.19 and 1.20 are proved in Appendix A.

Lemma 1.20. If $X = (x_1, \dots, x_n) \in [-L, L]^n$, then for any $\pi \in \mathcal{P}(\mathbb{R})$, the function $D_{\pi,X}(\cdot)$ is strictly increasing on $(-\infty, L]$ and strictly decreasing on $[L, \infty)$.

Proof. Each term in (1.3) obeys these monotonicity conditions. □

2 Proof of Theorem 1.2

In this section we prove Theorem 1.2 on generic behavior of $\hat{\pi}$ and $D_{\hat{\pi},X}$. The different parts are established by variations of the same core argument. For concreteness, we center our discussion on absolute continuity of $\hat{\pi}$. This is natural to expect from naive dimension-counting: if $\hat{\pi}$ is k -atomic, then it varies over a $2k - 1$ dimensional parameter space Π_k and must satisfy the $2k$ equations (1.6). However one of these equations is redundant since one always has $\mathbb{E}^{y \sim \pi} D_{\pi,X}(y) = 1$. Thus intuitively, all $2k - 1$ dimensions of Π_k should be needed to solve this many equations, suggesting the law of $\hat{\pi}$ is generic. Similar remarks apply in the static support case; here the conditions on $D'_{\hat{\pi},X}$ disappear, corresponding to the reduced dimension of $\Pi_k(S)$.

Proving absolute continuity of the law of $\hat{\pi}$ amounts to upper-bounding small-ball probabilities near arbitrary $\pi_0 \in \Pi_{k,\varepsilon}$. We first reduce this to bounding the probability that the stationarity conditions (1.6) for π_0 hold approximately. Below, $B_\delta(\pi_0)$ denotes a ball in d_{Π_k} -distance around π_0 . Similarly, let $B_{\delta,S}(\pi_0)$ be the δ -neighborhood of $\pi_0 \in \Pi_k(S)$ in the space $\Pi_k(S)$. Note that $B_{\delta,S}(\pi_0) \subseteq B_\delta(\pi_0) \cap \Pi_k(S)$.

Proposition 2.1. Fix $\pi_0 = \sum_{i=1}^k p_i \delta_{y_i} \in \Pi_k([-L, L])$. Then for $\hat{\pi} = \sum_{i=1}^k \hat{p}_i \delta_{\hat{y}_i} \in B_\delta(\pi_0)$ to hold (where δ -ball is defined with respect to d_{Π_k}), π_0 must approximately solve the system (1.6) in the sense that

$$\begin{aligned} \max_{1 \leq j \leq k} |D_{\pi_0, X}(y_j) - 1| &\leq e^{O(L^2)} \delta, \\ \max_{1 \leq j \leq k} |D'_{\pi_0, X}(y_j)| &\leq e^{O(L^2)} \delta \end{aligned} \quad (2.1)$$

Moreover to have $D_{\hat{\pi}, X}(y) = 1$ for some y with $|y - y_*| \leq \delta$, or $D''_{\hat{\pi}, X}(\hat{y}_j) = 0$ for some $1 \leq j \leq k$, we must respectively have

$$\max(|D_{\pi_0, X}(y_*) - 1|, |D'_{\pi_0, X}(y_*)|) \leq e^{O(L^2)} \delta, \quad (2.2)$$

$$|D''_{\pi_0, X}(y_j)| \leq e^{O(L^2)} \delta. \quad (2.3)$$

Similarly if $\hat{\pi}_S = \sum_{i=1}^k \hat{p}_i \delta_{\hat{y}_i} \in B_{\delta, S}(\pi'_S)$ for some π'_S supported in $S \subseteq L$, then:

$$\max_{1 \leq j \leq k} |D_{\pi'_S, X}(y_j) - 1| \leq e^{O(L^2)} \delta.$$

And to have $D_{\hat{\pi}_S, X}(y) = 1$ for some y with $|y - y_*| \leq \delta$, we must have

$$|D_{\pi'_S, X}(y_*) - 1| \leq e^{O(L^2)} \delta.$$

Proof. Using (1.20) and Proposition 1.16, we obtain

$$\|D_{\hat{\pi}, X} - D_{\pi_0, X}\|_{C^2([-L, L])} \leq e^{O(L^2)} \mathbb{W}_1(\hat{\pi}, \pi_0) \leq e^{O(L^2)} d_{\Pi_k}(\hat{\pi}, \pi_0).$$

We now apply Proposition 1.1. For the first estimate, we write

$$\begin{aligned} |D_{\pi_0, X}(y_j) - 1| &\leq |D_{\pi_0, X}(y_j) - D_{\pi_0, X}(\hat{y}_j)| + |D_{\pi_0, X}(\hat{y}_j) - D_{\hat{\pi}, X}(\hat{y}_j)| + |D_{\hat{\pi}, X}(\hat{y}_j) - 1| \\ &\leq e^{O(L^2)} |y_j - \hat{y}_j| + e^{O(L^2)} d_{\Pi_k}(\hat{\pi}, \pi_0) + 0. \end{aligned}$$

Similar arguments imply the other claims. $\square$

Proposition 2.1 effectively linearizes the stationarity conditions for $\hat{\pi}$. For π_0 fixed, the function $D_{\pi_0, X}$ and its derivatives are simply IID sums over $x_i \in X$. In particular their law becomes smoother as n increases; we make this precise using classical estimates from harmonic analysis.

2.1 Harmonic Analysis and Non-Degeneracy of Exponential Curves

Definition 2. We call a function $\gamma \in C^\infty([0, 1]; \mathbb{R}^d)$ a *smooth curve*, and $\gamma \in C^\infty(\mathbb{S}^1; \mathbb{R}^d)$ a *smooth loop*. We say γ is **non-degenerate** if for each x , the vectors

$$(\gamma'(x), \gamma''(x), \dots, \gamma^{(d)}(x))$$

form a basis for $\mathbb{R}^d$. By compactness, this is equivalent to the matrix $M(x)$ with these vectors as columns having determinant bounded away from zero, uniformly over $x \in [0, 1]$ or $x \in \mathbb{S}^1$. Let μ_γ denote the pushforward of the uniform measure on $[0, 1]$ by γ . Given a continuous function f defined on the range of γ , let $\mu_{\gamma, f}$ be the signed measure with Radon–Nikodym derivative $d\mu_{\gamma, f}(x)/d\mu_\gamma(x) = f(x)$.

The next key estimate follows from [SM93, Page 334], see also [Mar88, BGG⁺07].

Proposition 2.2. *Let γ be a non-degenerate smooth loop, and f a C^1 function on its range. Then the Fourier transform $\hat{\mu}$ of μ satisfies:*

$$|\hat{\mu}(\omega)| \leq C(\gamma)(1 + \|\omega\|)^{-1/d} \cdot \|f\|_{C^1}, \quad \forall \omega \in \mathbb{R}^d.$$

We next deduce that for non-degenerate γ , at most $d^2 + 1$ self-convolutions suffice for a bounded density.

Corollary 2.3. *Let γ be a non-degenerate smooth curve or loop, and f a non-negative bounded function on its range. Then with $(\cdot)^{*j}$ denoting j -fold self-convolution, the probability measure $\mu^{*(d^2+1)}$ has bounded density on $\mathbb{R}^d$.*

Proof. First, if γ is a smooth non-degenerate curve then it easily extends to a smooth non-degenerate loop $\tilde{\gamma}$. In either case, since $\|f\|_{L^\infty} < \infty$ there is a constant (in particular C^1) function on γ or $\tilde{\gamma}$ which is point-wise larger than f . Therefore by domination, it suffices to consider the case that γ is a loop.

For γ a smooth non-degenerate loop, Proposition 2.2 implies $\mu^{*(d^2+1)}$ has integrable Fourier transform, hence bounded density by Fourier inversion, completing the proof. $\square$

Remark 2.4. When γ is a loop and f is C^1 in Corollary 2.3, $\mu^{*(d^2+1)}$ in fact has a uniformly continuous density (being the inverse Fourier transform of an integrable function). However we can only use Proposition 2.1 to *upper bound* the law of $\hat{\pi}$, so this additional information does not improve our final results.

Remark 2.5. The appearance of $\Theta(d^2)$ convolutions in Corollary 2.3 is easily seen to be sharp for any smooth γ , degenerate or not. Indeed for small $\varepsilon > 0$, the curve $\gamma([0, \varepsilon])$ is contained inside a rectangular box with $O(\varepsilon^{d(d+1)/2})$ volume, spanned by vectors $O(\varepsilon^j)\gamma^{(j)}(x)$ for $1 \leq j \leq d$. (Here implicit constant may depend on both γ and k but not ε .) Thus μ^{*k} assigns measure $\Omega(\varepsilon^k)$ to the k -dilate of this box, which still has $O(\varepsilon^{d(d+1)/2})$ volume: each independent summand lands in the box with probability $\Omega(\varepsilon)$. Thus for any smooth γ , we must have $k \geq d(d+1)/2$ for μ^{*k} to have a bounded density on $\mathbb{R}^d$.

The following lemma will be used to verify non-degeneracy for the curves relevant to $\hat{\pi}$.

Lemma 2.6. *Let $P : \mathbb{R} \rightarrow (0, \infty)$ be a smooth, strictly positive function. Let $d_1, \dots, d_k \geq 0$ be non-negative integers and let $a_{i,j}$ be real constants for $0 \leq i < d_j$ which are not all zero. Define the function*

$$F(x) = \frac{\sum_{j=1}^k \sum_{i=0}^{d_j-1} a_{i,j} x^i e^{xy_j}}{P(x)}.$$

Then with $D = \sum_{j=1}^k d_j$, F can have at most $D - 1$ on $\mathbb{R}$, counting multiplicity. (Here x is a root of multiplicity r if $F(x) = F'(x) \cdots = F^{(r-1)}(x) = 0$.)

Proof. We can set $P(x) = 1$ since it does not affect the multiplicity of a root. It remains to show $F(x) = \sum_{j=1}^k \sum_{i=0}^{d_j-1} a_{i,j} x^i e^{xy_j}$ cannot have D roots. Indeed, this function satisfies a linear ODE with degree d characteristic polynomial $z \mapsto \prod_{j=1}^k (z - y_j)^{d_j}$. We consider the corresponding multi-set of differential operators $\phi_j(f) = f' - y_j f$. Further, if F has a non-zero term, then a subset of at most $D - 1$ of these operators can be applied to reach a function $f(x) = A e^{xy_i}$ for some $A \neq 0$ and i . Indeed, ϕ_j turns $x e^{xy_j}$ into e^{xy_j} , kills the e^{xy_j} term, and scales each other term by a non-zero constant. Crucially, this resulting f has no roots. However by Rolle's theorem, each of these differential operators reduces the total number of roots of a function by at most 1. It follows that F has fewer than D roots. $\square$

A compactness argument immediately gives the following.

Corollary 2.7. *In the setting of Lemma 2.6, if $\min_{j,j'} |y_j - y_{j'}| \geq \varepsilon$, there exists $c(L, k, \varepsilon)$ such that*

$$\max_{0 \leq d \leq D-1} |F^{(d)}(x)| \geq c \cdot \max_{i,j} |a_{i,j}|, \quad \forall x \in [-L, L].$$

2.2 Distributional Regularity of NPMLE

Proposition 2.1 implies that for any $\pi_0 \in \Pi_k$, having $\hat{\pi} \in B_\delta(\pi_0)$ requires the vector

$$V_{\pi_0}(X) \triangleq (D_{\pi_0,X}(y_1), \dots, D_{\pi_0,X}(y_k), D'_{\pi_0,X}(y_1), \dots, D'_{\pi_0,X}(y_k))$$

to be within distance $C(n, L)\delta$ of the half-ones vector $(1, \dots, 1, 0, \dots, 0)$. Note that the p_j -weighted average of the first k coordinates of this vector always equals 0 by (1.4), so the image of γ_{π_0} lies in a $2k - 1$ dimensional affine subspace $U_{\pi_0} \subseteq \mathbb{R}^{2k}$. In the S -restricted case, we denote the relevant curve by $\gamma_{\pi_0}^\circ$:

$$V_{\pi_0}^\circ(X) \triangleq (D_{\pi_0,X}(y_1), \dots, D_{\pi_0,X}(y_k)) \in \mathbb{R}^k$$

obeys the same constraint, hence lies in a $k - 1$ dimensional affine subspace $U_{\pi_0}^\circ$.

Lemma 2.8. *Suppose $x_1, \dots, x_n$ are IID from a density supported in $[-L, L]$ and bounded above pointwise by L' . For $n \geq k^2$ and $\pi_0 \in \Pi_{k,\varepsilon}$, the random vector $V_{\pi_0}^\circ(X) \in \mathbb{R}^k$ has compactly supported distribution on $U_{\pi_0}^\circ$ with density uniformly at most $C(n, \varepsilon, k, L, L')$. When $n \geq 4k^2$, the same holds for $V_{\pi_0}(X) \in \mathbb{R}^{2k}$ and U_{π_0} .*

Proof. We can write $V_{\pi_0}^\circ(X)$ as the convolution $V_{\pi_0}^\circ(X) = \frac{1}{n} \sum_{i=1}^n \gamma_{\pi_0}^\circ(x_i)$ where we define

$$\gamma_{\pi_0}^\circ(x) \triangleq \left(\frac{e^{-|x-y_1|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}}, \dots, \frac{e^{-|x-y_k|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}} \right).$$

Further, $\gamma_{\pi_0}^\circ(x)$ is an (x -independent) invertible linear transformation of

$$\tilde{\gamma}_{\pi_0}^\circ(x) \triangleq \left(\frac{e^{xy_1}}{P_{\pi_0}(x)e^{x^2/2}}, \dots, \frac{e^{xy_k}}{P_{\pi_0}(x)e^{x^2/2}} \right)$$

We will apply Corollary 2.3 with $\gamma = \tilde{\gamma}_{\pi_0}^\circ$. Note that the relevant dimension is $d = k - 1$, so $n \geq k^2 \geq d^2 + 1$ as required. It suffices to check that $\tilde{\gamma}_{\pi_0}^\circ$ is non-degenerate within the subspace $U_{\pi_0}^\circ$. We show the stronger statement that within $\mathbb{R}^k$,

$$\inf_{x \in [-L, L]} \det \left(\tilde{\gamma}_{\pi_0}^\circ(x), (\tilde{\gamma}_{\pi_0}^\circ)'(x), \dots, (\tilde{\gamma}_{\pi_0}^\circ)^{(k-1)}(x) \right) > 0.$$

This follows from Lemma 2.6 with $P(x) = P_{\pi_0}(x)e^{x^2/2}$ and $d_1 = \dots = d_k = 1$. Indeed, letting $\vec{a} = (a_{0,1}, \dots, a_{0,k}) \in \mathbb{R}^k$, and with F as in Lemma 2.6, we have

$$F^{(j)}(x) = \langle (\gamma_{\pi_0}^\circ)^{(j)}(x), \vec{a} \rangle, \quad \forall 0 \leq j \leq k - 1.$$

This completes the proof. The case of V_{π_0} is similar with

$$\begin{aligned} \gamma_{\pi_0}(x) &\triangleq \left(\frac{e^{-|x-y_1|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}}, \dots, \frac{e^{-|x-y_k|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}}, \frac{(x-y_1)e^{-|x-y_1|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}}, \dots, \frac{(x-y_k)e^{-|x-y_k|^2/2}}{P_{\pi_0}(x)\sqrt{2\pi}} \right), \\ \tilde{\gamma}_{\pi_0}(x) &\triangleq \left(\frac{e^{xy_1}}{P_{\pi_0}(x)e^{x^2/2}}, \dots, \frac{e^{xy_k}}{P_{\pi_0}(x)e^{x^2/2}}, \frac{xe^{xy_1}}{P_{\pi_0}(x)e^{x^2/2}}, \dots, \frac{xe^{xy_k}}{P_{\pi_0}(x)e^{x^2/2}} \right) \end{aligned}$$

and $d_1 = \dots = d_k = 2$ (taking $\vec{a} = (a_{0,1}, \dots, a_{0,k}, a_{1,1}, \dots, a_{1,k}) \in \mathbb{R}^{2k}$). □

Proof of Theorem 1.2, Part (I). Lemma 2.8 states that uniformly in $\pi_0 \in \Pi_{k,\varepsilon}$ for any ε fixed, $V_{\pi_0}(X)$ has bounded density on the $2k - 1$ dimensional subspace U_{π_0} . Proposition 2.1 thus implies that for any $\delta > 0$,

$$\mathbb{P}[\hat{\pi} \in B_\delta(\pi_0)] \leq C(n, \varepsilon, k, L, L')\delta^{2k-1}.$$

This implies $\hat{\pi}$ has locally bounded density on each Π_k . Since $\Pi_k = \cup_{\varepsilon>0} \Pi_{k,\varepsilon}$ and using similar countable exhaustion for a general absolutely continuous data distribution, we also deduce absolute continuity in the general case. $\square$

Proof of Theorem 1.13, Part (I). As in the previous proof, by countable additivity it suffices to consider data distributions with compactly supported bounded density. Lemma 2.8 states that $V_{\pi_0}^\circ(X)$ has bounded density on the $k - 1$ dimensional subspace $U_{\pi_0}^\circ$. Proposition 2.1 thus implies that for $\pi_0 \in \Pi_{k,\varepsilon}(S)$ and any $\delta > 0$,

$$\mathbb{P}[\hat{\pi}_S \in B_{\delta,S}(\pi_0)] \leq C(n, L, \varepsilon, k)\delta^{k-1}.$$

Finally this implies absolute continuity of $\hat{\pi}_S$ within $\Pi_k(S)$, since any $\pi_0 \in \Pi_k$ is in $\Pi_{k,\varepsilon}$ for some $\varepsilon > 0$. This completes the proof. $\square$

2.3 Generic Behavior of $D_{\hat{\pi},X}$

Distributional regularity of $\hat{\pi}$ does not have direct consequences for $D_{\hat{\pi},X}$ since the latter depends on both $\hat{\pi}$ and X . We prove genericity of $D_{\hat{\pi},X}$ using the same approach as above, thus establishing the remaining statements of Theorems 1.2 and 1.13. The idea is that any non-generic behavior can be encoded into an extra constraint dimension for the curve γ or γ° . Hence the chance for such behavior to hold approximately becomes $O(\delta^{2k})$ or $O(\delta^k)$ instead of $O(\delta^{2k-1})$ or $O(\delta^{k-1})$ respectively. These probabilities are now smaller than the inverse δ -radius covering numbers for Π_k and $\Pi_k(S)$. Hence summing over such covers shows a probability upper bound $O(\delta)$; taking δ small, we deduce that such non-generic behavior occurs with probability 0.

Proof of Theorem 1.2 part (II) and (III). We first show part (II), namely that $D_{\hat{\pi},X}$ has exactly k global maxima almost surely, namely $\text{supp}(\hat{\pi})$. We fix an arbitrary $L > 0$ (for use in e.g. Proposition 2.1) and show that the probability for L -bounded generic data to violate the claims is 0, which suffices by countable additivity.

We **claim** that for $\hat{\pi} \in \Pi_{k,\varepsilon}$, and any $y_* \in [-L, L]$ with $\min_{1 \leq j \leq k} |y_* - y_j| \geq 2\varepsilon$, there is $C(n, L, \varepsilon, k)$ such that for $\delta \leq \delta_*(n, L, \varepsilon, k)$ small enough, the probability that both (2.1) and (2.2) hold is at most $C\delta^{2k+1}$. Explicitly for any fixed π_0 and y_* :

$$\begin{aligned} & \mathbb{P}[\hat{\pi} \in B_\delta(\pi_0), \exists y \in [y_* - \delta, y_* + \delta], D_{\hat{\pi}}(y) = 1] \\ & \leq \mathbb{P}[\hat{\pi} \in B_\delta(\pi_0) \text{ and } |D_{\pi_0}(y^*) - 1| \leq C'\delta \text{ and } |D'_{\pi_0}(y^*)| \leq C'\delta] \leq C\delta^{2k+1}. \end{aligned} \tag{2.4}$$

(Here $C' \leq e^{O(L^2)}$.) Note that Π_k has dimension $2k - 1$ and it suffices to take $3L/\delta$ choices of y_* to cover all possible extraneous minimizers. There are hence $(C/\delta)^{2k}$ elements of a δ -net over pairs (π_0, y_*) . So using the claim and a union bound, the probability for $\hat{\pi}$ to be (ε, k) -non-degenerate and for $D_{\hat{\pi},X}$ to have an extraneous minimizer within distance 3ε of $\text{supp}(\hat{\pi})$ is $O(\delta)$ for small enough δ , hence zero. This suffices since ε was also arbitrary.

To show the claim (2.4) above, we proceed as in Lemma 2.8 using a slightly augmented version of V_{π_0} given by

$$(V_{\pi_0}(X), D_{\pi_0,X}(y_*), D'_{\pi_0,X}(y_*)) \in \mathbb{R}^{2k+2}.$$

The claim follows by using Lemma 2.6, with $d_1 = \dots = d_{k+1} = 2$ and $y_* = y_{k+1}$, as in the proof of Lemma 2.8. Here we use the assumption $n > (2k + 2)^2$. This finishes the proof of the first assertion on global maxima.

Part (III) is handled similarly. We now need to bound the probability that (2.1) and (2.3) both hold, and we will show it is at most $C(n, L, \varepsilon, k)\delta^{2k}$. This suffices by a similar union bound since the relevant δ -nets still have cardinality $(C/\delta)^{2k-1}$. This time we consider for some fixed $1 \leq j \leq k$:

$$(V_{\pi_0}(X), D''_{\pi_0, X}(y_j)) \in \mathbb{R}^{2k+1}.$$

Applying Lemma 2.6 with $d_j = 3$ and $d_i = 2$ for $i \neq j$ completes the proof in this case. $\square$

Proof of Theorem 1.13(II). The proof is similar to the above. Again we fix arbitrary L and restrict attention to the event $S \cup X \subseteq [-L, L]$. Consider $\varepsilon > 0$ such that $\hat{\pi}_S \in \Pi_{k, \varepsilon}(S)$, and any $y_* \in S \setminus \text{supp}(\hat{\pi}_S)$. Without loss of generality, we may choose 2ε smaller than the minimum distance between distinct points of S , so $d(y_*, \text{supp}(\hat{\pi}_S)) \geq 2\varepsilon$. Then analogously to (2.5), we have

$$\begin{aligned} & \mathbb{P}[\hat{\pi}_S \in B_{\delta, S}(\pi'_S), \exists y \in [y_* - \delta, y_* + \delta], D_{\hat{\pi}_S}(y) = 1] \\ & \leq \mathbb{P}[\hat{\pi} \in B_{\delta, S}(\pi'_S) \text{ and } |D_{\pi'_S}(y_*) - 1| \leq C'\delta] \leq C\delta^k. \end{aligned} \quad (2.5)$$

In this case, the claim follows by considering

$$(V_{\pi'_S}^\circ, D_{\pi'_S, X}(y_*)) \in \mathbb{R}^{k+1}$$

and applying Lemma 2.6, with $d_1 = \dots = d_{k+1} = 1$ and $y_* = y_{k+1}$. This time $\Pi_k(S)$ has dimension $k - 1$, and now y_* ranges over a finite set, so there are $(C/\delta)^{k-1}$ elements of a δ -net over pairs (π'_S, y_*) . Hence the probability for $D_{\hat{\pi}_S}$ to have an extraneous minimizer $y_* \in S \setminus \text{supp}(\hat{\pi}_S)$ is $O(\delta)$ for small enough δ , hence zero. This completes the proof. $\square$

3 Certification

In this section we give algorithms to certifiably compute $\hat{\pi}$, for both ε -approximation in $\mathbb{W}_1$ and exact support size. Theorem 1.2 will be used to show that for generic data, they eventually succeed almost surely.

3.1 Certification of Wasserstein Approximations

In Proposition 3.1 we show how to certify a putative candidate $\mathbb{W}_1$ -approximation to $\hat{\pi}$, denoted by $\tilde{\pi}_\varepsilon$. The conditions on $\tilde{\pi}_\varepsilon$ will be satisfied by (approximately) maximizing ℓ_X over $\mathcal{P}(Z_\varepsilon)$ and merging adjacent atoms.

Proposition 3.1. *Suppose there exists $\tilde{\pi}_\varepsilon = \sum_{j=1}^k p_j \delta_{y_j} \in \mathcal{P}(Z_\varepsilon)$, which satisfies for some $c_1, c_2, \delta > 0$ the following properties.*

1. *For y such that $d(y, \text{supp}(\tilde{\pi}_\varepsilon)) \geq c_1$, we have $D_{\tilde{\pi}_\varepsilon, X}(y) \leq 1 - c_2$.*
2. $\max_{y \in \mathbb{R}} D_{\tilde{\pi}_\varepsilon, X}(y) \leq 1 + \delta$.
3. *As a function on $\mathcal{P}(\text{supp}(\tilde{\pi}_\varepsilon))$, for some $\lambda \in (0, 1)$ the empirical loss ℓ_X has Hessian*

$$\nabla^2 \ell_X(\tilde{\pi}_\varepsilon) \preceq -\lambda I_k. \quad (3.1)$$

(We treat (3.1) as vacuously true in the case $k = 1$.)

4. Let $\eta = c_1 + L\delta/c_2$. Either $\eta \leq \frac{\lambda^3}{k^3 e^{14L^2}}$ or $\tilde{\pi}_\varepsilon([-10, 10]) \geq 1/10$ and

$$\eta \leq \frac{\lambda^3}{k^3 e^{5.1L^2}}. \quad (3.2)$$

If the above properties hold, then we must have:

$$\mathbb{W}_1(\tilde{\pi}_\varepsilon, \hat{\pi}) \leq O\left(Le^{L^2}\sqrt{\eta k/\lambda}\right).$$

The intuition is that conditions 1 and 2 imply $\hat{\pi}$ can be approximated by a distribution (denoted $\tilde{\pi}_*$ below) supported on $\text{supp}(\tilde{\pi}_\varepsilon)$ and achieving a similar value of likelihood, while condition (3.1) implies any distribution on $\text{supp}(\tilde{\pi}_\varepsilon)$ achieving high likelihood must be close to $\tilde{\pi}_\varepsilon$ itself. A detailed proof is given in Appendix A.

Several conditions in Proposition 3.1 will be verified using Theorem 1.2. However one condition follows easily from uniqueness of $\hat{\pi}$.

Proposition 3.2. *For any sample $X = (x_1, \dots, x_n) \in \mathbb{R}^n$ there exists $\lambda = \lambda_X > 0$ such that condition (3.1) holds at $\hat{\pi}$. (Recall we take (3.1) to be vacuous when $k = 1$).*

Proof. Since ℓ_X is strongly convex in the vector $P_\xi(X) = (P_\pi(x_1), \dots, P_\pi(x_n))$, if the claim did not hold then the directional derivative of V would be zero along some line segment. However V depends linearly on the weights $\vec{p}$ so we conclude that $\hat{\pi}$ is not unique, contradicting Proposition 1.1. $\square$

Next we turn to the support size of $\hat{\pi}$. We note that one direction is implied by a Wasserstein approximation, as verified in the next result.

Proposition 3.3. *Suppose $\nu = \sum_{j=1}^k p_j \delta_{y_j}$. For each j let d_j be the minimum distance from y_j to a nearest (other) atom and let $\Delta = \Delta(\nu) \triangleq \min_j (p_j d_j) > 0$. Then we have*

$$\frac{\Delta}{3} \leq \inf\{\mathbb{W}_1(\nu, \mu) : \mu \text{ has } < k \text{ atoms}\} \leq \Delta.$$

In particular,

$$\mathbb{W}_1(\nu, \mu) \leq \Delta(\nu)/3 \implies |\text{supp}(\mu)| \geq |\text{supp}(\nu)|.$$

Additionally if $\mathbb{W}_1(\nu, \mu) \leq \Delta(\nu)/3$ and $|\text{supp}(\mu)| = |\text{supp}(\nu)|$, then $\mu = \sum_{j=1}^k q_j \delta_{z_j}$ where

$$\sqrt{\sum_{j=1}^k (|p_j - q_j|^2 + d(y_j, z_j)^2)} \leq \frac{12 \mathbb{W}_1(\mu, \nu)}{\Delta}. \quad (3.3)$$

Proof. For the lower bound, note that the open balls $B(y_j, d_j/3)$ of radius $d_j/3$ centered at y_j are mutually disjoint, separated by pairwise distance $\min(d_j)/3$. Hence if μ has $< k$ atoms, at least one of these balls must have μ -measure zero, and hence any coupling of μ to ν has to incur at least $\Delta/2$ cost. For the upper bound, simply remove the atom j minimizing the definition of Δ and move its mass to the nearest neighbor.

For the second assertion, it follows from the preceding paragraph that μ has exactly 1 atom in each ball $B(y_j, d_j/3)$; without loss of generality call this atom z_j . Using a similar transportation argument and then $\min(p, q) + |p - q| = \max(p, q)$, we find

$$\begin{aligned} \mathbb{W}_1(\mu, \nu) &\geq \sum_j d(y_j, z_j) \cdot \min(p_j, q_j) + \frac{\sum_j d_j |p_j - q_j|}{6} \geq \frac{\sum_j (d(y_j, z_j)p_j + d_j |p_j - q_j|)}{12} \\ &\geq \frac{\Delta}{12} \sum_j (d(y_j, z_j) + |p_j - q_j|) \geq \frac{\Delta}{12} \sqrt{\sum_j (d(y_j, z_j)^2 + |p_j - q_j|^2)}. \end{aligned} \quad \square$$

The first part of the above lemma will be used to show that the Wasserstein bound from the previous Proposition eventually certifies that

$$\mathbb{W}_1(\hat{\pi}_\varepsilon, \hat{\pi}) < \Delta(\hat{\pi}_\varepsilon)/2.$$

The second part will be used to obtain convergence in parameter distance d_{Π_k} .

Proposition 3.4. *There exists an absolute constant C such that the following holds. Suppose that $\tilde{\pi}$ is given and there exists $c > 0$ such that:*

1. $\mathbb{W}_1(\hat{\pi}, \tilde{\pi}) \leq \alpha$.
2. For C an absolute constant and all $y \in \text{supp}(\tilde{\pi})$,

$$D''_{\tilde{\pi},X}(y) \leq -CL^2e^{4L^2}(\alpha + cL).$$

3. Condition 1 of Proposition 3.1 holds with $(c_1, c_2) = (c, Ce^{4L^2}\alpha)$. I.e. for all z with $d(z, \text{supp}(\tilde{\pi})) \geq c$, we have $D_{\tilde{\pi},X}(z) \leq 1 - Ce^{4L^2}\alpha$.

Then $|\text{supp}(\hat{\pi})| \leq |\text{supp}(\tilde{\pi})|$.

Proof. We take C (polynomially) large compared to the constants C_j from Lemma 1.18. By the bound (1.19) applied with $j = 3$ and first-order Taylor expansion for D'' we conclude that $D''_{\tilde{\pi},X}(y) \leq -CL^2e^{4L^2}\alpha$ for all y in the c -neighborhood of any support element of $\tilde{\pi}$. In turn, by (1.20) with $j = 2$ we have that also $D''_{\tilde{\pi},X}(y) \leq -\frac{C}{2}L^2e^{4L^2}\alpha$ in the same neighborhood. Thus, in each such neighborhood $D_{\tilde{\pi},X}$ can have at most one local maximum, and in turn at most one y such that $D_{\tilde{\pi},X}(y) = 1$. Outside of c -neighborhoods of each support element of $\hat{\pi}_\varepsilon$ there can be no atoms of $\hat{\pi}$ because by (1.20) with $j = 0$ we should have that $D_{\tilde{\pi},X}(y) < 1$ for all $d(y, \text{supp}(\tilde{\pi})) \geq c$. $\square$

3.2 Approximability of $\hat{\pi}$

Here we synthesize the preceding results to obtain an approximation scheme for $\hat{\pi}$. Again, we assume here access to an exact maximization oracle for the concave function ℓ_X over $\mathcal{P}(Z_\varepsilon)$, deferring issues of approximate maximization to the Appendix. We first define the constants that determine the dependence on X . Let

$$A_X = - \max_{y \in \text{supp}(\hat{\pi})} D''_{\hat{\pi},X}(y) \tag{3.4}$$

and

$$\begin{aligned} B_X &= 1 - \max_{y \in S_X} D_{\hat{\pi},X}(y), \\ S_X &\triangleq \{y \in \mathbb{R} : d(y, \text{supp}(\hat{\pi})) \geq a_X/4\}, \\ a_X &\triangleq \frac{A_X}{2C_3L^2e^{4L^2}}, \end{aligned}$$

where C_j 's are from (1.20). Both A_X, B_X are almost surely strictly positive by Theorem 1.2, and we will implicitly assume so below. To approximate $\hat{\pi}$, we begin as usual with

$$\hat{\pi}_\varepsilon = \arg \max_{\text{supp}(\pi) \subseteq Z_\varepsilon} \ell_X(\pi)$$

for Z_ε as in (1.11). First, we show a version of Proposition 1.1 for $\hat{\pi}_\varepsilon$.

Lemma 3.5. *Given $\hat{\pi}_\varepsilon$ we have:*

- (a) $D_{\hat{\pi}_\varepsilon, X}(y) = 1$ for all $y \in \text{supp}(\hat{\pi}_\varepsilon)$.
- (b) $D_{\hat{\pi}_\varepsilon, X}(y) \leq 1$ for all $y \in Z_\varepsilon$.
- (c) $\max_{y \in \mathbb{R}} D_{\hat{\pi}_\varepsilon, X}(y) \leq 1 + \delta$ for $\delta = C\varepsilon^2 L^2 e^{2L^2}$.

Further, for any $y \in \text{supp}(\hat{\pi}_\varepsilon)$ there exists a local maximum $\hat{y} \in [-L, L]$ of $D_{\hat{\pi}_\varepsilon, X}$ such that $|y - \hat{y}| < \varepsilon$.

Proof. By (1.5) and optimality of $\hat{\pi}_\varepsilon$ we see that $D_{\hat{\pi}_\varepsilon, X}(y)$ is maximized at each $y \in \text{supp}(\hat{\pi}_\varepsilon)$ (possibly other y as well). Since the derivative there vanishes when $\pi' = \pi$, we see that this maximum value must be 1, yielding points (a) and (b).

Point (c) follows by smoothness estimates. Let $y \in [y_1, y_1 + \varepsilon]$ with $y_1 \in Z_\varepsilon$. We recall the following simple analytical fact: if g is a smooth function on $[0, \varepsilon]$ and $g(0), g(\varepsilon) \leq 1$ then we must have

$$\sup_{y \in [0, \varepsilon]} |g(y)| \leq 1 + \frac{\varepsilon^2}{2} \sup_{y \in [0, \varepsilon]} |g''(y)|,$$

which can be shown by a Taylor expansion. Applying this fact to $D_{\hat{\pi}_\varepsilon, X}$ we find that

$$D_{\hat{\pi}_\varepsilon, X}(y) \leq 1 + \frac{\varepsilon^2}{2} \sup_{z \in \mathbb{R}} |D_{\hat{\pi}_\varepsilon, X}''(z)|.$$

Recalling (1.19) completes the proof of (c) for $y \in [-L, L]$. Finally recall from (1.11) that we assume $\{-L, L\} \subseteq Z_\varepsilon$, and from Lemma 1.20 that $D_{\hat{\pi}_\varepsilon, X}(\cdot)$ is strictly decreasing on $[-L, \infty)$ and strictly increasing on $(-\infty, -L]$. Hence if $y > L$, then $D_{\hat{\pi}_\varepsilon, X}(y) > D_{\hat{\pi}_\varepsilon, X}(L) \leq 1 + \delta$. Proceeding similarly for $y \leq -L$ completes the proof of (c).

For the last claim, given $y \in \text{supp}(\hat{\pi}_\varepsilon)$ let $y_1 < y < y_2$ be the closest points to y on each side in $\text{supp}(\hat{\pi}_\varepsilon)$. First assuming $y \notin \{-L, L\}$ is not an extreme point of Z_ε , by definition $[y_1, y_2] \subseteq [y - \varepsilon, y + \varepsilon]$. From (a) and (b), we see that $D_{\hat{\pi}_\varepsilon}(y) \geq D_{\hat{\pi}_\varepsilon}(y_i)$ for $i \in \{1, 2\}$. Hence $\max_{z \in [y_1, y_2]} D_{\hat{\pi}_\varepsilon}(z)$ is attained within the interior of the interval, so such a $\hat{y}$ exists. Finally the boundary cases $y = \pm L$ are easily handled using Lemma 1.20 to rule out local maxima outside $[-L, L]$. $\square$

We now list several properties of $\hat{\pi}_\varepsilon$ that hold for sufficiently small ε . They do not immediately suffice for *certification* because one needs an explicit bound on how small ε must be. Instead these properties will be used to ensure the certification criteria in Propositions 3.1, 3.3 and 3.4 are applicable to $\hat{\pi}_\varepsilon$ or its modification $\tilde{\pi}_\varepsilon$, for small enough ε .

Proposition 3.6. *We have*

$$\lim_{\varepsilon \rightarrow 0} d_{\mathcal{H}}(\text{supp}(\hat{\pi}_\varepsilon), \text{supp}(\hat{\pi})) = 0, \quad (3.5)$$

where $d_{\mathcal{H}}$ denotes Hausdorff distance (recall (1.14)). Consequently, for all $\varepsilon < \varepsilon_0(X)$, the following hold:

$$\begin{aligned} d(y, \text{supp}(\hat{\pi}_\varepsilon)) &\leq a_X/2 \implies D_{\hat{\pi}_\varepsilon, X}''(y) \leq -A_X/2; \\ d(y, \text{supp}(\hat{\pi}_\varepsilon)) &\geq a_X/2 \implies D_{\hat{\pi}_\varepsilon, X}(y) \leq 1 - \frac{B_X}{2}. \end{aligned}$$

Proof. First, note that

$$\lim_{\varepsilon \rightarrow 0} \max_{y \in \text{supp}(\hat{\pi})} d(y, \text{supp}(\hat{\pi}_\varepsilon)) = 0,$$

since otherwise there would have existed an atom of $\hat{\pi}$ that is bounded away from any atom of $\hat{\pi}_\varepsilon$ contradicting $\mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) \rightarrow 0$. Second, we also have

$$\lim_{\varepsilon \rightarrow 0} \max_{y \in \text{supp}(\hat{\pi}_\varepsilon)} d(y, \text{supp}(\hat{\pi})) = 0.$$

If we take any convergent subsequence $\text{supp}(\hat{\pi}_\varepsilon) \ni y_\varepsilon \rightarrow y_*$ then we have $1 = D_{\hat{\pi}_\varepsilon, X}(y_\varepsilon) \rightarrow D_{\hat{\pi}, X}(y_*)$. But by Theorem 1.2, $D_{\hat{\pi}, X}(y_*) = 1$ (a global maximum) implies $y_* \in \text{supp}(\hat{\pi})$, thus showing the claim. Together these convergence statements yield (3.5).

Proposition 1.8 ensures that uniformly in the choice of Z_ε , we have $\lim_{\varepsilon \rightarrow 0} \mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) = 0$. By Lemma 1.19, this implies convergence of $D_{\hat{\pi}_\varepsilon, X}$ to $D_{\hat{\pi}, X}$ in the space $C^2([-L, L])$. The other two statements of the proposition now follow from the convergence of supports. $\square$

Proposition 3.7. *For $\varepsilon < \varepsilon_0(X)$ sufficiently small depending on X , any pair of points $\text{supp}(\hat{\pi}_\varepsilon)$ within distance $a_X/2$ are adjacent elements in Z_ε .*

Proof. Take a point $y \in \text{supp}(\hat{\pi}_\varepsilon)$. By Proposition 3.6 (which uses the assumption $\varepsilon < \varepsilon_0(X)$), $D_{\hat{\pi}_\varepsilon, X}'' \leq -A_X/2$ on the interval $[y - a_X/2, y + a_X/2]$. Thus, the equation $D_{\hat{\pi}_\varepsilon, X}(z) = 1$ can have at most two solutions, one of which is y . Further the concavity of $D_{\hat{\pi}_\varepsilon, X}$ on $[y - a_X/2, y + a_X/2]$ implies it is unimodal. Thus if $y' \in [y - a_X/2, y + a_X/2]$ also satisfies $D_{\hat{\pi}_\varepsilon, X}(y') = 1$, then $D_{\hat{\pi}_\varepsilon, X}(y'') = 1$ for all $y'' \in (y, y')$. Recalling Lemma 3.5 part (b), we see that no such y'' can lie in Z_ε . We conclude that y and y' must be consecutive within Z_ε . $\square$

In light of Proposition 3.7, for ε small enough atoms in $\hat{\pi}_\varepsilon$ are all separated by $a_X/2$ except those occurring in pairs of adjacent elements of Z_ε . We form $\tilde{\pi}_\varepsilon$ by taking the weighted average of each such pair, i.e. replacing $p_i \delta_{x_i}$ and $p_{i+1} \delta_{x_{i+1}}$ by $\tilde{p}_i \delta_{\tilde{x}_i}$ for

$$\tilde{p}_i = p_i + p_{i+1}, \quad \tilde{x}_i = \frac{p_i x_i + p_{i+1} x_{i+1}}{x_i + x_{i+1}}.$$

We now show crucial properties of $\tilde{\pi}_\varepsilon$ that will be used in certification. Again, the results of the following proposition by itself are insufficient since $\varepsilon_0(X)$ is not determined explicitly.

Proposition 3.8. *For all $\varepsilon < \varepsilon_0(X)$ we have that $|\text{supp}(\tilde{\pi}_\varepsilon)| = |\text{supp}(\hat{\pi})| =: k$. Furthermore, representing $\tilde{\pi}_\varepsilon = \sum_{j=1}^k \tilde{p}_j \delta_{\tilde{y}_j}$ and $\hat{\pi} = \sum_{j=1}^k p_j \delta_{y_j}$ (with both atoms sorted):*

$$\tilde{p}_j \rightarrow p_j, \quad \tilde{y}_j \rightarrow y_j$$

as $\varepsilon \rightarrow 0$. Thus, in particular, we have

$$\lim_{\varepsilon \rightarrow 0} \Delta(\tilde{\pi}_\varepsilon) = \Delta(\hat{\pi}) > 0,$$

Proof. Given $d_{\mathcal{H}}(\text{supp}(\tilde{\pi}_\varepsilon), \text{supp}(\hat{\pi}_\varepsilon)) \leq \varepsilon$ and (3.5), $\lim_{\varepsilon \rightarrow 0} d_{\mathcal{H}}(\text{supp}(\tilde{\pi}_\varepsilon), \text{supp}(\hat{\pi})) = 0$. Since the atoms of both measures are separated by at least $a_X/4$ we also conclude the two sets have the same cardinality (for small ε), and that $\tilde{y}_j \rightarrow y_j$. $\square$

Next, we show that $\tilde{\pi}_\varepsilon$ admits explicit $O_X(\varepsilon^{1/3})$ certificates for Wasserstein convergence and a certified lower bound on $|\text{supp}(\hat{\pi})|$ by virtue of Propositions 3.1 and 3.3. (The former $\mathbb{W}_1$ convergence also holds for $\hat{\pi}_\varepsilon$.)

Proposition 3.9. For $\varepsilon \leq \varepsilon_0(X)$, the conditions of Proposition 3.1 apply to $\tilde{\pi}_\varepsilon$ with:

$$\begin{aligned}\delta &= O(\varepsilon^2 L^2 e^{4L^2}), \\ c_1 &= \sqrt[3]{10L\delta/A_X}, \\ c_2 &= \sqrt[3]{L^2\delta^2 A_X}, \\ \eta &= O(\sqrt[3]{L\delta/A_X}), \\ \lambda &= \Omega_X(1)\end{aligned}\tag{3.6}$$

In particular, the estimator $\tilde{\pi}_\varepsilon$ obeys certifiable bounds of the form $\mathbb{W}_1(\hat{\pi}, \tilde{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/3})$. Furthermore, we also have a bound $|\text{supp}(\hat{\pi})| \geq |\text{supp}(\tilde{\pi}_\varepsilon)|$ certified by Proposition 3.3.

Proof. The bound on δ for $\hat{\pi}_\varepsilon$ is in Lemma 3.5. To extend it to $\tilde{\pi}_\varepsilon$, it suffices to show

$$\sup_{y \in [-L, L]} |D_{\tilde{\pi}_\varepsilon, X}(y) - D_{\hat{\pi}_\varepsilon, X}(y)| \stackrel{?}{\leq} O(\varepsilon^2 L^2 e^{4L^2}).$$

By definition of D , for this it suffices to show that

$$\sup_{y \in [-L, L]} |P_{\tilde{\pi}_\varepsilon}(y) - P_{\hat{\pi}_\varepsilon}(y)| \stackrel{?}{\leq} O(\varepsilon^2),$$

which reduces to proving that

$$\left| \left(\frac{p_i}{p_i + p_{i+1}} \right) e^{-|x_i - y|^2/2} + \left(\frac{p_{i+1}}{p_i + p_{i+1}} \right) e^{-|x_{i+1} - y|^2/2} - e^{-|\tilde{x}_i - y|^2/2} \right| \stackrel{?}{\leq} O(\varepsilon^2).$$

This holds by the following fact (applied with $f(x) = e^{-|x|^2/2}$): If $|f''(x)| \leq C$ for all $x \in [x_0, x_1]$ then for any $\lambda \in [0, 1]$ we have

$$|f(\lambda x_1 + (1 - \lambda)x_0) - \lambda f(x_1) - (1 - \lambda)f(x_0)| \leq C(x_1 - x_0)^2.$$

(This is shown by noticing that this estimate is true at $\lambda = 0$ and then expanding the right-hand side in λ to second order.)

Next, let us show that $D_{\tilde{\pi}_\varepsilon}$ satisfies the (c_1, c_2) assumption in Proposition 3.1. We already know this holds for $c_1 = a_X/2$ and $c_2 = B_X/2 - O(\varepsilon^2)$ by Proposition 3.6, but we need to improve this estimate to the case of vanishing c_1, c_2 to get a vanishing estimate on $\mathbb{W}_1$.

To that end, we claim that whenever $2\varepsilon \leq c_1 \leq \frac{a_X}{2}$, we may take

$$c_2 = \frac{A_X(c_1 - 2\varepsilon)^2}{4} - \delta\tag{3.7}$$

in Proposition 3.1 (where δ is still as in (3.6) above). Indeed, if

$$d(y, \text{supp}(\tilde{\pi}_\varepsilon)) \geq c_1,$$

then $d(y, \tilde{y}) \geq c_1 - \varepsilon$ for some $\tilde{y} \in \text{supp}(\hat{\pi}_\varepsilon)$. Note that $\tilde{y}$ is within distance ε of a local maximum $\hat{y}$ of $D_{\hat{\pi}_\varepsilon, X}$ by Lemma 3.5. This implies by Proposition 3.6 that, as claimed,

$$D_{\tilde{\pi}_\varepsilon, X}(y) \leq D_{\hat{\pi}_\varepsilon, X}(\hat{y}) - \frac{A_X(c_1 - 2\varepsilon)^2}{4} \leq 1 - \left(\frac{A_X(c_1 - 2\varepsilon)^2}{4} - \delta \right) = 1 - c_2.$$

In particular for ε small enough, we can set

$$c_1 = \sqrt[3]{10L\delta/A_X}, \quad c_2 = \sqrt[3]{L^2\delta^2A_X} \implies \eta = O(\sqrt[3]{L\delta/A_X}).$$

Further, Proposition 3.2 ensures that (3.2) holds for ε small enough, since the value λ_ε for $\hat{\pi}_\varepsilon$ converges to that of $\hat{\pi}$. (This follows from the fact that $\hat{\pi}_\varepsilon \rightarrow \hat{\pi}$ in $\mathbb{W}_1$ and in Hausdorff distance of supports, by a similar computation as in Lemma 1.19.) We conclude that Proposition 3.1 suffices to certify bounds of the form $\mathbb{W}_1(\hat{\pi}, \tilde{\pi}_\varepsilon) \leq O_X(\sqrt{\eta}) = O_X(\varepsilon^{1/3})$.

Finally, from Proposition 3.8 we know that conditions of Proposition 3.3 are eventually certifiable and hence we get a certifiable lower bound $|\text{supp}(\hat{\pi})| \geq |\text{supp}(\tilde{\pi}_\varepsilon)|$ on the support size of the NPMLE. $\square$

Proof of Theorem 1.9, except for the Shub–Smale property. Proposition 3.9 shows $\tilde{\pi}_\varepsilon$ has a certified $\mathbb{W}_1 \leq O_X(\varepsilon^{1/3})$ upper bound and that $|\text{supp}(\hat{\pi})| \geq |\text{supp}(\tilde{\pi}_\varepsilon)|$.

The upper bound on the support will be certified via Proposition 3.4, where we may certify the bound $\alpha = O_X(\varepsilon^{1/3})$ using Proposition 3.9. Lemma 1.19 then yields a (certified) estimate

$$\|D_{\hat{\pi}_\varepsilon, X} - D_{\hat{\pi}, X}\|_{C^2([-L, L])} \leq e^{O(L^2)} \alpha \leq O_X(\varepsilon^{1/3}).$$

To satisfy Condition 3 of Proposition 3.4, it suffices by (3.7) to take $c = \Theta_X(\varepsilon^{1/6})$. Then Condition 2 of Proposition 3.4 holds for ε small enough by Proposition 3.8 (regardless of the X -dependent constant factors). Recalling how $\tilde{\pi}_\varepsilon$ was constructed shows $|\text{supp}(\hat{\pi})| = |\text{supp}(\tilde{\pi}_\varepsilon)|$. Then (3.3) gives the same bound for parameter distance of the rounding $\tilde{\pi}_\varepsilon$. We refer to Appendix D for the proof that roundings are Shub–Smale approximate NPMLEs. $\square$

In the preceding proof, one can actually take $c = \Theta_X(1)$ to be an ε -independent constant. However due to the unspecified dependence on X , this does not yield an actual algorithm. By making c decay with ε , we ensure the conditions hold once $\varepsilon < \varepsilon_0(X)$ is small enough.

4 Unbounded Support Size of Higher Dimensional NPMLE

In higher-dimensions $d \geq 2$, the NPMLE $\hat{\pi}$ for a spherical Gaussian mixture need not be unique as was observed in [SGS24, Lemma 2] by taking $X = (x_1, x_2, x_3)$ to be the vertices of an equilateral triangle. We show below that $|\text{supp}(\hat{\pi})|$ may be unbounded for any $d \geq 2$ even when the points $x_i \in X$ are uniformly bounded (and $\hat{\pi}$ is chosen to minimize the support size if non-unique). In other words, Proposition 1.4 does not generalize beyond dimension 1. Both of these points indicate that different ideas are required to obtain a fine-grained understanding of the NPMLE in higher dimensions.

It will be helpful to extend (1.2) beyond discrete datasets: for any $\mu \in \mathcal{P}(\mathbb{R}^d)$, let the associated NPMLE $\hat{\pi}(\mu)$ be any maximizer of

$$\ell_\mu(\pi) \triangleq \int \log P_\pi(x) \, d\mu(x) \tag{4.1}$$

for $P_\pi = \pi * \mathcal{N}(0, I_d)$. Moreover, let $\mu_{d,r}^{\text{sp}}$ denote the uniform distribution on a centered sphere $\mathbb{S}_{d,r}$ of radius r inside $\mathbb{R}^d$.

Lemma 4.1. *For any dimension $d \geq 2$ and large enough $R \geq R_0(d)$, let $\mu = \mu_{d,R}^{\text{sp}}$. Then $\hat{\pi} = \mu_{d,r}^{\text{sp}}$ is unique, and r satisfies $0 < R - r < o_R(1)$.*

Lemma 4.1 is proved in Appendix A. It immediately implies that Proposition 1.4 does not extend to multiple dimensions.

Proposition 4.2. *For any $d \geq 2$, for large enough $R \geq R_0(d)$ there exists a sequence of data-sets $X_n = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$ satisfying $\|x_i\| \leq R$ for all $1 \leq i \leq n$ such that the following holds. If $\hat{\pi}^{(n)}$ is an NPMLE of X_n for each $n \geq 1$, then $\lim_{n \rightarrow \infty} |\text{supp}(\hat{\pi}^{(n)})| = \infty$.*

Proof. Let $(X_n)_{n \geq 1}$ be any sequence of uniformly bounded datasets such that $\frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ converges in distribution to $\mu_{d,R}^{\text{sp}}$ where R is as in Lemma 4.1 with $r > 0$. It is well-known that any NPMLE $\hat{\pi}$ for X_n supported inside the radius R ball is also supported inside the radius R ball. Note that the function $(X, \pi) \mapsto \ell_X(\pi)$ defined in (1.2) is jointly continuous for X, π supported inside the radius R ball. Therefore any subsequential limit of NPMLEs $\hat{\pi}^{(n)}$ for X_n must be an NPMLE for $\mu_{d,R}^{\text{sp}}$. By Lemma 4.1 this means $\hat{\pi}^{(n)} \xrightarrow{d} \mu_{d,r}^{\text{sp}}$. In particular their support sizes must grow to infinity, which concludes the proof. $\square$

Acknowledgment

YP was supported in part by NSF Grant CCF-2131115 and the MIT-IBM Watson AI Lab.

References

- [BDJ⁺22] A. Bakshi, I. Diakonikolas, H. Jia, D. M. Kane, P. K. Kothari, and S. S. Vempala. Robustly learning mixtures of k arbitrary gaussians. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1234–1247, 2022.
- [BGG⁺07] L. Brandolini, G. Gigante, A. Greenleaf, A. Iosevich, A. Seeger, and G. Travaglini. Average decay estimates for Fourier transforms of measures supported on curves. *J. Geom. Anal.*, 17(1):15–40, 2007.
- [Böh86] D. Böhning. A vertex-exchange-method in D -optimal design theory. *Metrika*, 33(1):337–347, 1986.
- [BSL92] D. Böhning, P. Schlattmann, and B. Lindsay. Computer-assisted analysis of mixtures (C.A.MAN): statistical algorithms. *Biometrics*, pages 283–303, 1992.
- [Bub15] S. Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [BWY17] S. Balakrishnan, M. J. Wainwright, and B. Yu. Statistical guarantees for the EM algorithm: From population to sample-based analysis. *Ann. Stat.*, 45(1):77–120, 2017.
- [Das99] S. Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.
- [Der86] R. DerSimonian. Maximum likelihood estimation of a mixing distribution. *Applied Statistics*, pages 302–309, 1986.
- [DLR77] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.
- [DTZ17] C. Daskalakis, C. Tzamos, and M. Zampetakis. Ten steps of EM suffice for mixtures of two Gaussians. In *Conference on Learning Theory*, pages 704–710. PMLR, 2017.
- [DWYZ23] N. Doss, Y. Wu, P. Yang, and H. H. Zhou. Optimal estimation of high-dimensional Gaussian location mixtures. *Ann. Stat.*, 51(1):62–95, 2023.
- [DZ16] L. H. Dicker and S. D. Zhao. High-dimensional classification via nonparametric empirical Bayes and maximum likelihood inference. *Biometrika*, 103(1):21–34, 2016.
- [FD18] L. Feng and L. H. Dicker. Approximate nonparametric maximum likelihood for mixture models: A convex optimization approach to fitting arbitrary multivariate mixing distributions. *Computational Statistics & Data Analysis*, 122:80–91, 2018.

- [FGSW23] Z. Fan, L. Guan, Y. Shen, and Y. Wu. Gradient flows for empirical Bayes in high-dimensional linear models. *arXiv preprint arXiv:2312.12708*, 2023.
- [GVDV01] S. Ghosal and A. W. Van Der Vaart. Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *Ann. Stat.*, 29(5):1233–1263, 2001.
- [GW00] C. R. Genovese and L. Wasserman. Rates of convergence for the Gaussian mixture sieve. *Ann. Stat.*, 28(4):1105–1127, 2000.
- [GW12] P. Groeneboom and J. A. Wellner. *Information bounds and nonparametric maximum likelihood estimation*, volume 19. Birkhäuser, 2012.
- [HP15] M. Hardt and E. Price. Tight bounds for learning a mixture of two Gaussians. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 753–760, 2015.
- [Jag13] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *International conference on machine learning*, pages 427–435. PMLR, 2013.
- [JZ09] W. Jiang and C.-H. Zhang. General maximum likelihood empirical Bayes estimation of normal means. *Ann. Stat.*, 37(1):1647–1684, 2009.
- [JZB⁺16] C. Jin, Y. Zhang, S. Balakrishnan, M. J. Wainwright, and M. I. Jordan. Local maxima in the likelihood of Gaussian mixture models: Structural results and algorithmic consequences. In *Advances in neural information processing systems*, pages 4116–4124, 2016.
- [Kan21] D. M. Kane. Robust learning of mixtures of gaussians. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1246–1258. SIAM, 2021.
- [KCSA20] Y. Kim, P. Carbonetto, M. Stephens, and M. Anitescu. A fast algorithm for maximum likelihood estimation of mixture proportions using sequential quadratic programming. *Journal of Computational and Graphical Statistics*, 29(2):261–273, 2020.
- [KG17] R. Koenker and J. Gu. Rebayes: an r package for empirical bayes mixture methods. *Journal of Statistical Software*, 82:1–26, 2017.
- [KG19] R. Koenker and J. Gu. Minimalist g -Modeling. *Statistical Science*, 34(2):209 – 213, 2019.
- [KM14] R. Koenker and I. Mizera. Convex optimization, shape constraints, compound decisions, and empirical Bayes rules. *Journal of the American Statistical Association*, 109(506):674–685, 2014.
- [KMV10] A. T. Kalai, A. Moitra, and G. Valiant. Efficiently learning mixtures of two Gaussians. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 553–562, 2010.
- [KW56] J. Kiefer and J. Wolfowitz. Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *Ann. Math. Stat.*, pages 887–906, 1956.
- [Lai78] N. Laird. Nonparametric maximum likelihood estimation of a mixing distribution. *Journal of the American Statistical Association*, 73(364):805–811, 1978.
- [Lin83a] B. G. Lindsay. The Geometry of Mixture Likelihoods: A General Theory. *Ann. Stat.*, pages 86–94, 1983.
- [Lin83b] B. G. Lindsay. The Geometry of Mixture Likelihoods, Part II: The Exponential Family. *Ann. Stat.*, 11(3):783–792, 1983.
- [LK92] M. L. Lesperance and J. D. Kalbfleisch. An algorithm for computing the nonparametric MLE of a mixing distribution. *Journal of the American Statistical Association*, 87(417):120–126, 1992.
- [LM22] A. Liu and A. Moitra. Learning GMMs with nearly optimal robustness guarantees. In *Conference on Learning Theory*, pages 2815–2895. PMLR, 2022.
- [LM23] A. Liu and A. Moitra. Robustly learning general mixtures of gaussians. *J. ACM*, 70(3):1–53, 2023.
- [LR93] B. G. Lindsay and K. Roeder. Uniqueness of estimation and identifiability in mixture models. *Canadian J. Stat.*, 21(2):139–147, 1993.

- [Mar88] B. Marshall. Decay rates of Fourier transforms of curves. *Trans. Amer. Math. Soc.*, 310(1):115–126, 1988.
- [MF05] J. Ma and S. Fu. On the correct convergence of the EM algorithm for Gaussian mixtures. *Pattern Recognition*, 38(12):2602–2611, 2005.
- [MR94] X.-L. Meng and D. B. Rubin. On the global and componentwise rates of convergence of the EM algorithm. *Linear Algebra and its Applications*, 199:413–425, 1994.
- [MSS23] S. Mukherjee, B. Sen, and S. Sen. A Mean Field Approach to Empirical Bayes Estimation in High-dimensional Linear Regression. *arXiv preprint arXiv:2309.16843*, 2023.
- [MV10] A. Moitra and G. Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.
- [MX05] J. Ma and L. Xu. Asymptotic convergence properties of the EM algorithm with respect to the overlap in the mixture. *Neurocomputing*, 68:105–129, 2005.
- [MXJ00] J. Ma, L. Xu, and M. I. Jordan. Asymptotic convergence rate of the EM algorithm for Gaussian mixtures. *Neural Computation*, 12(12):2881–2907, 2000.
- [Pea94] K. Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London. A*, 185:71–110, 1894.
- [PW20] Y. Polyanskiy and Y. Wu. Self-Regularizing Property of Nonparametric Maximum Likelihood Estimator in Mixture Models. *arXiv preprint arXiv:2008.08244*, 2020.
- [Rob50] H. Robbins. A generalization of the method of maximum likelihood: Estimating a mixing distribution (Abstract). In *Annals of Mathematical Statistics*, volume 21, pages 314–315, 1950.
- [RW84] R. A. Redner and H. F. Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, 1984.
- [SB13] J. Stoer and R. Bulirsch. *Introduction to Numerical Analysis*, volume 12. Springer Science & Business Media, 2013.
- [SF99] R. P. Stanley and S. Fomin. *Enumerative Combinatorics*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999.
- [SG20] S. Saha and A. Guntuboyina. On the nonparametric maximum likelihood estimator for Gaussian location mixture densities with application to Gaussian denoising. *Ann. Stat.*, 48(2):738–762, 2020.
- [SGS24] J. A. Soloff, A. Guntuboyina, and B. Sen. Multivariate, heteroscedastic empirical Bayes via nonparametric maximum likelihood. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2024.
- [SM93] E. M. Stein and T. S. Murphy. *Harmonic analysis: real-variable methods, orthogonality, and oscillatory integrals*, volume 3. Princeton University Press, 1993.
- [SS93] M. Shub and S. Smale. Complexity of Bezout’s theorem I: Geometric aspects. *J. Amer. Math. Soc.*, 6(2):459–501, 1993.
- [Tan17] Y. S. Tan. Energy optimization for distributions on the sphere and improvement to the Welch bounds. *Electron. Comm. Probab.*, 22:1 – 12, 2017.
- [WB22] N. Weinberger and G. Bresler. The EM algorithm is adaptively-optimal for unbalanced symmetric Gaussian mixtures. *Journal of Machine Learning Research*, 23(103):1–79, 2022.
- [WIM25] H. Wang, S. Ibrahim, and R. Mazumder. Nonparametric Finite Mixture Models with Possible Shape Constraints: A Cubic Newton Approach. *SIAM Journal on Mathematics of Data Science*, 7(1):163–188, 2025.

- [WN22] Y. Wei and X. Nguyen. Convergence of de Finetti’s mixing measure in latent structure models for observed exchangeable sequences. *Ann. Stat.*, 50(4):1859–1889, 2022.
- [WY20] Y. Wu and P. Yang. Optimal estimation of gaussian mixtures via denoised method of moments. *Ann. Stat.*, 48(4):1981–2007, 2020.
- [WZ22] Y. Wu and H. H. Zhou. Randomly initialized EM algorithm for two-component Gaussian mixture achieves near optimality in $O(\sqrt{n})$ iterations. *Mathematical Statistics and Learning*, 4(3):143–220, 2022.
- [XJ96] L. Xu and M. I. Jordan. On convergence properties of the EM algorithm for Gaussian mixtures. *Neural Computation*, 8(1):129–151, 1996.
- [Zha09] C.-H. Zhang. Generalized maximum likelihood estimation of normal mixture densities. *Statistica Sinica*, pages 1297–1318, 2009.
- [ZLS20] R. Zhao, Y. Li, and Y. Sun. Statistical convergence of the EM algorithm on Gaussian mixture models. *Electron. J. Stat.*, 14:632–660, 2020.

A Proofs of Lemmas

Here we provide self-contained proofs of several lemmas from the main body.

Proof of Proposition 1.16. Recalling (1.16), define a coupling (y, y') between π, π' as follows. For each j , with probability $\min(p_j, p'_j)$ we have $(y, y') = (y_j, y'_j)$. The remainder of the coupling is arbitrary. Note that by definition this remaining probability is

$$1 - \sum_{j=1}^k \min(p_j, p'_j) = \frac{1}{2} \sum_j |p_j - p'_j| \leq \frac{1}{2} \sqrt{L \sum_j |p_j - p'_j|^2}.$$

Since $\text{diam}([-L, L]) = 2L$, we obtain

$$\mathbb{W}_1(\pi, \pi') \leq L^{3/2} \sqrt{\sum_j |p_j - p'_j|^2} + \max_{1 \leq j \leq k} |y_j - y'_j| \leq (L^{3/2} + 1) d_{\Pi_k}(\pi, \pi'). \quad \square$$

Proof of Lemma 1.18. Point 1 is trivial. Letting $(y, y') \sim \Gamma$ be an optimal coupling of π, π' , point 2 follows easily from (1.1):

$$|P_\pi(x) - P_{\pi'}(x)| \leq \mathbb{E}^{(y, y') \sim \Gamma} |e^{-|x-y|^2/2} - e^{-|x-y'|^2/2}| \leq C \mathbb{E}^{(y, y') \sim \Gamma} |y - y'| \leq C\delta.$$

Point 3 follows easily by combining the previous two. For point 4, note that since 1 holds for both π and π' . Hence recalling the definition (1.2) of ℓ_X , the Lipschitz constant of the logarithm on $[C^{-1}e^{-2L^2}, \infty)$ is $O(e^{2L^2})$ which gives the claim. The bound (1.18) follows from point 1. Using (1.17) for each term $T^{(j)}$ we have

$$|T_{\pi, x}^{(j)}(z) - T_{\pi', x}^{(j)}(z)| \leq \left| \frac{H_j(x - z)e^{-|x-z|^2/2}}{P_\pi(x_i)\sqrt{2\pi}} - \frac{H_j(x - z)e^{-|x-z|^2/2}}{P_{\pi'}(x_i)\sqrt{2\pi}} \right| \leq C_j L^j \left| \frac{1}{P_\pi(x_i)} - \frac{1}{P_{\pi'}(x_i)} \right|.$$

This proves 5 in light of point 3. The analogous bounds for D follow from (1.3). $\square$

Proof of Lemma 1.19. We repeatedly differentiate. Let $v_j = q_j - p_j$ and $p_{t,j} = (1-t)p_j + tq_j$. Using (1.5) in the first line:

$$\ell'(t) = \sum_{j=1}^k v_j D_{\pi_t, X}(y_j) \stackrel{(1.3)}{=} \frac{1}{n} \sum_{m=1}^n \left(\sum_{j=1}^k \frac{v_j e^{-|x_m - y_j|^2/2}}{\sum_{\ell} p_{t,\ell} e^{-|x_m - y_\ell|^2/2}} \right).$$

Clearly it suffices to consider the case $n = 1$, with $x_1 = x$, which removes the outer average over m . Then we find:

$$\begin{aligned} \ell''(t) &= \sum_{j=1}^k v_j \frac{d}{dt} \left(\frac{e^{-|x - y_j|^2/2}}{\sum_{\ell} p_{t,\ell} e^{-|x - y_\ell|^2/2}} \right) = - \sum_{j,j'=1}^k v_j v_{j'} \left(\frac{\exp(-\frac{|x - y_j|^2 + |x - y_{j'}|^2}{2})}{(\sum_{\ell} p_{t,\ell} e^{-|x - y_\ell|^2/2})^2} \right); \\ \ell'''(t) &= 2 \sum_{j,j',j''=1}^k v_j v_{j'} v_{j''} \left(\frac{\exp(-\frac{|x - y_j|^2 + |x - y_{j'}|^2 + |x - y_{j''}|^2}{2})}{(\sum_{\ell} p_{t,\ell} e^{-|x - y_\ell|^2/2})^3} \right). \end{aligned}$$

To lower bound the denominator, note that

$$\sum_{\ell} p_{t,\ell} e^{-|x - y_\ell|^2/2} \geq \min_{x,y \in [-L,L]} e^{-|x - y|^2/2} = e^{-2L^2}.$$

This easily gives the main conclusion. When $\pi([-10, 10]) \geq 0.1$, we have $\sum_{\ell} p_{t,\ell} e^{-|x - y_\ell|^2/2} \geq e^{-(L+10)^2/2}/10 \geq \Omega(e^{-0.51L^2})$. $\square$

Proof of Proposition 3.1. We first consider the case $\tilde{\pi}_\varepsilon([-10, 10]) \geq 1/10$. Note that for $\pi_t = (1-t)\tilde{\pi}_\varepsilon + t\hat{\pi}$, by definition of $\hat{\pi}$:

$$0 \leq \frac{d}{dt} \ell(\pi_t)|_{t=0} = \int D_{\tilde{\pi}_\varepsilon, X}(y) (d\hat{\pi}(y) - d\tilde{\pi}_\varepsilon(y)) \stackrel{(1.4)}{=} \int (D_{\tilde{\pi}_\varepsilon, X}(y) - 1) d\hat{\pi}(y).$$

This implies the bulk of $\hat{\pi}$ lies near the support of $\tilde{\pi}_\varepsilon$: by Assumptions 1 and 2,

$$\mathbb{P}^{y \sim \hat{\pi}}[d(y, \text{supp}(\tilde{\pi}_\varepsilon)) \geq c_1] \leq \frac{\delta}{\delta + c_2} \leq \delta/c_2.$$

We exploit this as follows. By the previous display, there exists $\tilde{\pi}_*$ supported in $\text{supp}(\tilde{\pi}_\varepsilon)$ such that $\mathbb{W}_1(\tilde{\pi}_*, \hat{\pi}) \leq c_1 + 2L\delta/c_2 \leq 2\eta$. This implies

$$|\ell_X(\hat{\pi}) - \ell_X(\tilde{\pi}_*)| \leq O(e^{2L^2}\eta)$$

by Lemma 1.18. We write $\tilde{\pi}_* = \sum_y q_j \delta_{y_j}$.

Now that $\tilde{\pi}_*$ and $\tilde{\pi}_\varepsilon$ have the same support, it remains to show their weights are similar. We do this using strong convexity via (3.1). Set $c = \sum_j |p_j - q_j|$, which we will show is small. We write

$$\ell(t) = \ell_X((1-t)\tilde{\pi}_\varepsilon + t\tilde{\pi}_*).$$

Then we have $\ell'(0) = 0$ since $\tilde{\pi}_\varepsilon$ is an exact optimum among distributions supported in Z_ε , and $\ell''(0) \leq -c^2\lambda/k$ by (3.1). By Lemma 1.19 with $j = 3$, we find $\|\ell'''\|_\infty \leq O(c^3 e^{1.55L^2})$ and so $\ell''(t) \leq -c^2\lambda/2k$ for all

$$|t| \leq T \triangleq \Omega\left(\frac{\lambda}{ck e^{1.55L^2}}\right).$$

Thus, since $\ell(t)$ is non-increasing, we have

$$\ell_X(\tilde{\pi}_*) - \ell_X(\tilde{\pi}_\varepsilon) = \ell(1) - \ell(0) \leq \ell(1 \wedge T) - \ell(0) \leq -(1 \wedge T)^2 c^2 \lambda / 4k.$$

Combining the above, we have

$$\begin{aligned} \ell_X(\hat{\pi}) - O\left(\eta e^{2L^2}\right) &\leq \ell_X(\tilde{\pi}_*) \leq \ell_X(\tilde{\pi}_\varepsilon) - (1 \wedge T)^2 c^2 \lambda / 4k \\ &\leq \ell_X(\hat{\pi}) - (1 \wedge T)^2 c^2 \lambda / 4k. \end{aligned}$$

Therefore

$$(1 \wedge T)c \leq C e^{L^2} \sqrt{\frac{\eta k}{\lambda}}.$$

Due to the assumption (3.2), the latter estimate cannot hold with Tc on the left-hand side. Therefore we have an upper bound for c which yields:

$$\mathbb{W}_1(\hat{\pi}, \tilde{\pi}_\varepsilon) \leq \mathbb{W}_1(\hat{\pi}, \pi_t) + \mathbb{W}_1(\pi_t, \tilde{\pi}_\varepsilon) \leq Lc + O(\eta) \leq O\left(Le^{L^2} \sqrt{\frac{\eta k}{\lambda}} + \eta\right).$$

The former term dominates the latter when (3.2) holds, finishing the proof of the main bound.

Without assumption $\tilde{\pi}_\varepsilon([-10, 10]) \geq 1/10$, the proof is the same: 5.1 comes from $5.1 = 2(1 + 1.55)$, and replacing 1.55 by 6 (based on Lemma 1.19) gives 14. $\square$

We next turn to proving Lemma 4.1, which uses the following fact. It can be shown by expansion into spherical harmonics, or see [Tan17, Theorem 4.4] for an elementary approach.

Proposition A.1. *Let $F : \mathbb{R} \rightarrow \mathbb{R}$ be given by a globally absolutely convergent power series $F(t) = \sum_{k \geq 0} a_k t^k$ with strictly positive coefficients $a_k > 0$. Then for any $\nu \in \mathcal{P}(\mathbb{S}_{d,r})$,*

$$\iint F(\langle x, y \rangle) d\nu(x) d\nu(y) \geq \iint F(\langle x, y \rangle) d\mu_{d,r}^{\text{sp}}(x) d\mu_{d,r}^{\text{sp}}(y)$$

with equality if and only if $\nu = \mu_{d,r}^{\text{sp}}$.

Proof of Lemma 4.1. First, let $\hat{\pi}_{\text{sym}}$ be the spherical symmetrization of $\hat{\pi}$. It is easy to see by concavity of the logarithm that $\ell_\mu(\hat{\pi}) \leq \ell_\mu(\hat{\pi}_{\text{sym}})$. Thus $\hat{\pi}_{\text{sym}}$ is a mixture of $\mu_{d,r'}^{\text{sp}}$ for $r' \in \mathbb{R}_{\geq 0}$. We will show separately that there is a unique optimal choice r of r' , and that $\ell_\mu(\hat{\pi}) < \ell_\mu(\hat{\pi}_{\text{sym}})$ unless $\hat{\pi} = \hat{\pi}_{\text{sym}}$.

Uniqueness of Optimal Radius First, note that for any r' , the density of $\mu_{d,r'}^{\text{sp}} * \mathcal{N}(0, I_d)$ is constant on the sphere of radius R . Hence maximizing (4.1) is equivalent to maximizing this value in r' . Thus, let $f(r_1, r_2)$ be the density of $\mu_{d,r_1}^{\text{sp}} * \mathcal{N}(0, I_d)$ on the sphere of radius r_2 . Letting $\mu_{d,r,\varepsilon}^{\text{sp}}$ be the uniform distribution on the dimension d annulus $A_{d,r,\varepsilon}$ of inner and outer radii r and $r + \varepsilon$, and using ρ_ν to denote the density of an absolutely continuous distribution:

$$\begin{aligned} f(r_1, r_2) &= \frac{1}{2\pi r_2} \lim_{\varepsilon \rightarrow 0} \int_{A_{d,r_2,\varepsilon}} d\mu_{d,r_1,\varepsilon}^{\text{sp}} * \mathcal{N}(0, I_d)(x) = \lim_{\varepsilon \rightarrow 0} \int \rho_{\mu_{d,r_2,\varepsilon}^{\text{sp}}}(x) d\mu_{d,r_1,\varepsilon}^{\text{sp}} * \mathcal{N}(0, I_d)(x) \\ &= \lim_{\varepsilon \rightarrow 0} \int \rho_{\mu_{d,r_2,\varepsilon}^{\text{sp}}}(-x) d\mu_{d,r_1,\varepsilon}^{\text{sp}} * \mathcal{N}(0, I_d)(x) = \lim_{\varepsilon \rightarrow 0} \rho_{\mu_{d,r_1,\varepsilon}^{\text{sp}} * \mathcal{N}(0, I_d)} * \mu_{d,r_2,\varepsilon}^{\text{sp}}(0). \end{aligned}$$

In particular $f(r_1, r_2) = f(r_2, r_1)$, and so r is the radius on which the density $\hat{\rho} = \rho_{\mu_{d,R}^{\text{sp}} * \mathcal{N}(0, I_d)}$ is maximized. By symmetry, it suffices to consider the restriction of $\hat{\rho}$ to $\{r'v : r' \geq 0\}$ for a unit vector $v \in \mathbb{R}^d$.

The uniqueness of an optimal radius $R - o_R(1) < r < R$ is now geometrically rather clear and we outline a formal proof.

First if $|R - r'| \geq R^{1/10}$, then $\widehat{\rho}(r'v) \leq O(R^{-2d})$ which will turn out to be of lower order than the maximum value. Hence we restrict attention to $r' = R \pm O(R^{1/10})$; here the contribution to $\widehat{\rho}$ from the part of $\mu_{d,R}^{\text{sp}}$ supported at distance at least $R^{1/4}$ from v is at most $O(R^{-2d})$ in C^2 , thanks to the super-polynomial decay of the Gaussian density and its derivatives.

For $x \in \text{supp}(\mu_{d,R}^{\text{sp}})$ with $\|x - v\| \leq R^{1/4}$, we replace x by its projection x' onto the tangent hyperplane to $\mathbb{S}_{d,R}$ at Rv . It is easy to see that $\|x - x'\| \leq O(R^{-3/2})$, so this change affects $\widehat{\rho}(r'v)$ by at most $O_d(e^{-\|x' - Rv\|^2/3} R^{-1})$ in C^2 per unit mass $d\mu(x)$. Finally approximating $\widehat{\rho}(r'v)$ by integrating over x' instead of x , we get an approximation $\widetilde{\rho}(r'v)$ which is simply a Gaussian density centered at R and rescaled by a factor of $\Theta_d(R^{-(d-1)})$ (for the fraction of $\mathbb{S}_{d,R}$ within an $O(1)$ distance of Rv). The error from changing $x \rightarrow x'$ and including x at distance greater than $R^{1/4}$ is at most $O_d(R^{-d})$ in C^2 norm. Combining the above shows that a maximizing r is unique and satisfies $|R - r| \leq o_R(1)$. Finally $r < R$ simply because any NPMLE must be supported on the strict interior of $\text{supp}(\mu_{d,R}^{\text{sp}})$.

We note that at this point, $\mu_{d,r}^{\text{sp}}$ has been shown to be **an** NPMLE for $\mu_{d,R}^{\text{sp}}$. The latter part of the proof below shows it is the **only** NPMLE.

Spherical Symmetry of $\widehat{\pi}$ Given the preceding discussion, we know that any NPMLE $\widehat{\pi}$ is supported on $\mathbb{S}_{d,r}$ for some unique r . Recalling (4.1), note that without the logarithm, the quantity $\int P_{\pi}(x) d\mu(x)$ is constant over all such $\widehat{\pi}$. Therefore by concavity of the logarithm, it suffices to prove that $\mu_{d,r}^{\text{sp}}$ is the unique probability measure ν on $\mathbb{S}_{d,r}$ whose convolution with $\mathcal{N}(0, I_d)$ produces a constant density when restricted to $\mathbb{S}_{d,R}$. We will do so by proving that it uniquely minimizes the L^2 energy of the density, given by

$$\int \rho_{\nu * \mathcal{N}(0, I_d)}(w)^2 d\mu_{d,R}^{\text{sp}}(w). \quad (\text{A.1})$$

To establish the latter fact, we will expand and rearrange the integral in order to apply Proposition A.1. Crucially, note that for $\|x\| = r$ and $\|w\| = R$, we have $e^{-\|x-w\|^2/2} \propto e^{\langle x, w \rangle}$ with constant of proportionality $e^{-(r^2+R^2)/2}$ depending only on r, R . Using this observation, we expand (A.1) and interchange the order of integration:

$$\begin{aligned} \int \rho_{\nu * \mathcal{N}(0, I_d)}(w)^2 d\mu_{d,R}^{\text{sp}}(w) &\propto \iiint e^{\langle x+y, w \rangle} d\mu_{d,R}^{\text{sp}}(w) d\nu(x) d\nu(y) \\ &= \iint \left(\int e^{\langle x+y, w \rangle} d\mu_{d,R}^{\text{sp}}(w) \right) d\nu(x) d\nu(y). \end{aligned} \quad (\text{A.2})$$

For the inner integral, let

$$C_{d,k} = R^{-k} \int w_1^k d\mu_{d,R}^{\text{sp}}(w)$$

where w_1 is the first coordinate of $w = (w_1, \dots, w_d) \in \mathbb{S}_{d,R}$. Of course, $C_{d,2j+1} = 0$ while $C_{d,2j} \in (0, 1)$. Then for any $z \in \mathbb{R}^d$:

$$\int e^{\langle z, w \rangle} d\mu_{d,R}^{\text{sp}}(w) = \sum_{k \geq 0} \int \frac{\langle z, w \rangle^k}{k!} d\mu_{d,R}^{\text{sp}}(w) = \sum_{k \geq 0} \frac{C_{d,k} R^k \|z\|^k}{k!} = \sum_{j \geq 0} \frac{C_{d,2j} R^{2j} \|z\|^{2j}}{(2j)!}.$$

Recalling (A.2), we take $z = x + y$ and observe that $\|z\|^2 = 2(r^2 + \langle x, y \rangle)$. Combining,

$$\int \rho_{\nu * \mathcal{N}(0, I_d)}(w)^2 d\mu_{d,R}^{\text{sp}}(w) = \iint \left(\sum_{j \geq 0} \frac{C_{d,2j} R^{2j} (r^2 + \langle x, y \rangle)^j}{(2j)!} \right) d\nu(x) d\nu(y).$$

Recalling that $C_{d,2j} \in (0, 1)$ for all j , we find that the right-hand side has all coefficients strictly positive as a power series in $\langle x, y \rangle$. Moreover it converges absolutely on all of $\mathbb{R}$ by inspection as a power series in $w = r^2 + \langle x, y \rangle$, and global absolute convergence is invariant under an affine change of variable. Thus Proposition A.1 applies and concludes the proof. $\square$

B Approximate Stationarity Conditions

Here we explain how to compute an approximate stationary point (in the relevant sense) for the concave function $\ell_X : \mathcal{P}(Z_\varepsilon) \rightarrow \mathbb{R}$ in a provably efficient manner. Note that Proposition 1.4 implies $|\text{supp}(\hat{\pi})| \leq O(L^2) \ll |Z_\varepsilon| \asymp L/\varepsilon$, i.e. $\hat{\pi}$ is a sparse vector in $\mathcal{P}(Z_\varepsilon)$. Indeed, for us the relevant notion of approximate-stationary point π will be that

$$\left. \frac{d}{ds} (\ell_X((1-s)\pi + s\pi')) \right|_{s=0}$$

is small for all $\pi' \in \mathcal{P}(Z_\varepsilon)$. This definition is sensitive to the support of π , and in particular it is easier to satisfy when small atoms of π are rounded down to 0. This is precisely what we do below, based on the Frank–Wolfe conditional gradient method. To start, we set $\pi_0 = \delta_0$ and iteratively define:

$$\begin{aligned} \pi^{(t+1)} &= \frac{t\pi^{(t)} + 2\delta_{y_t}}{t+2}, \\ y_t &= \arg \max_{y \in Z_\varepsilon} D_{\pi^{(t)}, X}(y). \end{aligned} \tag{B.1}$$

In using this manifestation of the Frank–Wolfe algorithm, we implicitly equip $\mathcal{P}(Z_\varepsilon)$ with the Euclidean norm on its finite sequence of probability mass values (and also negated ℓ_X to make it convex). In particular the Wasserstein distance does not enter here. We also point out that [Jag13, Theorem 2], which we rely on below to understand (B.1), applies even if y_t is only an approximate maximizer of $D_{\pi^{(t)}, X}$. (This could easily be incorporated into Lemma B.1 to ensure that approximately maximal y suffice in (B.1).)

Next we modify $\pi^{(t)}$ to $\check{\pi}^{(t)}$; this will ensure $D_{\check{\pi}^{(t)}, X}(y) - 1$ is close to 1 for all $y \in \text{supp}(\check{\pi}^{(t)})$. Define for $\iota > 0$ the subset $R_{\varepsilon, \iota} = \{y \in Z_\varepsilon : p_{\pi^{(t)}}(y) \leq \iota\}$. For $\iota \leq \frac{\varepsilon}{4L}$,

$$\sum_{y \in R_{\varepsilon, \iota}} p_{\pi^{(t)}}(y) \leq 3\iota L/\varepsilon < 1/2. \tag{B.2}$$

Thus we may define another probability measure

$$\check{\pi}^{(t)}(y) = \begin{cases} 0, & y \in R_{\varepsilon, \iota} \\ \frac{\pi^{(t)}(y)}{1 - \sum_{y \in R_{\varepsilon, \iota}} \pi^{(t)}(y)}, & y \notin R_{\varepsilon, \iota}. \end{cases}$$

Lemma B.1. *Fix any $X \in [-L, L]^n$. For ε small enough depending only on L and for $t \geq \varepsilon^{-8}$, the probability measure $\check{\pi}^{(t)}$ satisfies:*

$$\begin{aligned} D_{\check{\pi}^{(t)}, X}(y) - 1 &\geq -O\left(\frac{e^{4L^2}}{\varepsilon^{1/2} \sqrt[4]{t}}\right), \quad \forall y \in \text{supp}(\check{\pi}^{(t)}), \\ D_{\check{\pi}^{(t)}, X}(y) - 1 &\leq O\left(\frac{e^{4L^2}}{\varepsilon^{1/2} \sqrt[4]{t}}\right), \quad \forall y \in Z_\varepsilon. \end{aligned}$$

Proof. We apply [Jag13, Theorem 2] to the algorithm (B.1). The conclusion is a bound

$$g(\pi^{(t)}) \leq \frac{10 \operatorname{diam}(\mathcal{P}(Z_\varepsilon))^2 \operatorname{Lip}(\nabla \ell_X)}{t+2}.$$

where (using (1.5)) we have by definition

$$g(\pi^{(t)}) = \max_{y \in Z_\varepsilon} (D_{\pi^{(t)}, X}(y) - 1) \cdot (1 - p_{\pi^{(t)}}(y)). \quad (\text{B.3})$$

It is easy to see that $\operatorname{diam}(\mathcal{P}(Z_\varepsilon))^2 \leq 2$. Meanwhile, $\nabla \ell_X(\pi) = D_{\pi, X}(\cdot)$. Combining (1.20) and the fact that $\mathbb{W}_1(\pi, \pi') \leq 2L\|\pi - \pi'\|_{TV}$, it follows that the Lipschitz constant $\operatorname{Lip}(\nabla \ell_X)$ is at most $O(Le^{4L^2})$. Altogether we find that

$$D_{\pi^{(t)}, X}(y) - 1 \leq \frac{O(Le^{4L^2})}{(t+2)(1 - p_{\pi^{(t)}}(y))}.$$

Additionally, (1.4) implies that

$$D_{\pi^{(t)}, X}(y) - 1 \leq \frac{1 - p_{\pi^{(t)}}(y)}{p_{\pi^{(t)}}(y)}.$$

Next we combine these two estimates. It is easy to see that for small β , one has $\min(\frac{p-1}{p}, \frac{\beta}{1-p}) \leq O(\sqrt{\beta})$ for all $p \in [0, 1]$ by casework on the event $p \leq 1 - \sqrt{\beta}$. Hence for (say) $t \geq e^{6L^2}$ and all $y \in Z_\varepsilon$,

$$D_{\pi^{(t)}, X}(y) - 1 \leq \max\left(\frac{O(Le^{4L^2})}{(t+2)(1 - p_{\pi^{(t)}}(y))}, \frac{1 - p_{\pi^{(t)}}(y)}{p_{\pi^{(t)}}(y)}\right) \leq O\left(\sqrt{\frac{Le^{4L^2}}{t+2}}\right). \quad (\text{B.4})$$

Next we turn to $\check{\pi}^{(t)}$, showing it approximately preserves the preceding upper bound for all $y \in Z_\varepsilon$. Note that $\|\pi^{(t)} - \check{\pi}^{(t)}\|_{TV} \leq 3\iota L/\varepsilon$, hence $\mathbb{W}_1(\pi^{(t)}, \check{\pi}^{(t)}) \leq 6\iota L^2/\varepsilon$. Using again (1.4) and the preceding upper bound, we find that for $y \notin R_{\varepsilon, \iota}$:

$$\begin{aligned} D_{\pi^{(t)}, X}(y) - 1 &\geq -p_{\pi^{(t)}}(y)^{-1} \cdot O\left(\sqrt{\frac{Le^{4L^2}}{t}}\right) \\ D_{\check{\pi}^{(t)}, X}(y) - 1 &\geq (D_{\pi^{(t)}, X}(y) - 1) + (D_{\check{\pi}^{(t)}, X}(y) - D_{\pi^{(t)}, X}(y)) \stackrel{(1.20)}{\geq} -O\left(\frac{Le^{2L^2}}{\iota\sqrt{t}} + \frac{e^{4L^2}\iota L^2}{\varepsilon}\right). \end{aligned}$$

Taking $\iota = e^{-L^2}t^{-1/4}\varepsilon^{1/2}$, we obtain

$$D_{\check{\pi}^{(t)}, X}(y) - 1 \geq -O\left(\frac{e^{4L^2}}{\varepsilon^{1/2}\sqrt[4]{t}}\right), \quad \forall y \in \operatorname{supp}(\check{\pi}^{(t)}),$$

under the condition $t \geq \varepsilon^{-8}$ for ε sufficiently small (to ensure (B.2)). Finally bounding $|D_{\check{\pi}^{(t)}, X}(y) - D_{\pi^{(t)}, X}(y)|$ in the same way, (B.4) becomes

$$D_{\check{\pi}^{(t)}, X}(y) - 1 \leq O\left(\frac{e^{4L^2}}{\varepsilon^{1/2}\sqrt[4]{t}}\right), \quad \forall y \in Z_\varepsilon. \quad \square$$

We can now show an approximate version of Lemma 3.5.

Lemma B.2. *For ε small enough depending on L there exists an algorithm with complexity $O(Ln\varepsilon^{-11})$ which returns $\check{\pi}_\varepsilon \in \mathcal{P}(Z_\varepsilon)$ obeying for $\delta = CL^2e^{4L^2}\varepsilon^2$:*

- (a) $D_{\check{\pi}_\varepsilon, X}(y) \geq 1 - \delta$ for all $y \in \text{supp}(\check{\pi}_\varepsilon)$.
- (b) $\max_{y \in \mathbb{R}} D_{\check{\pi}_\varepsilon, X}(y) \leq 1 + \delta$.
- (c) $\ell_X(\hat{\pi}) - \ell_X(\check{\pi}_\varepsilon) \leq \delta$.

Proof. Note that each Frank–Wolfe iteration (B.1) requires $O(Ln/\varepsilon)$ operations to find the maximal $y \in Z_\varepsilon$. Taking $\check{\pi}_\varepsilon = \check{\pi}^{(t)}$ as defined in Lemma B.1, the first two parts follow by setting $t = \varepsilon^{-10}$, except that the upper bound must hold for all $y \in \mathbb{R}$ instead of just $y \in Z_\varepsilon$. The extension to $[-L, L]$ follows immediately via (1.19) (and after e.g. doubling δ), and this suffices by Lemma 1.20. The last assertion follows by (B.3) since $1 - p_{\pi^{(t)}}(y) \leq 1$, since $g(\pi^{(t)})$ is a certified (dual) upper bound on the suboptimality of $\pi^{(t)}$. $\square$

Below we let $\check{\pi}_\varepsilon$ be the approximation to $\hat{\pi}$ guaranteed by Lemma B.2. We first show that Proposition 3.6 extends to this setting, with $\check{\pi}_\varepsilon$ in place of $\hat{\pi}_\varepsilon$. Recall the definitions of A_X, B_X in and just below (3.4).

Proposition B.3. *For ε small enough, we have:*

$$\begin{aligned} d(y, \text{supp}(\check{\pi}_\varepsilon)) \leq a_X/2 &\implies D''_{\check{\pi}_\varepsilon, X}(y) \leq -A_X/2; \\ d(y, \text{supp}(\check{\pi}_\varepsilon)) \geq a_X/2 &\implies D_{\check{\pi}_\varepsilon, X}(y) \leq 1 - \frac{B_X}{2}. \end{aligned}$$

Proof. Using Lemma B.2(c), the proof of Proposition 1.8 implies $\lim_{\varepsilon \rightarrow 0} \mathbb{W}_1(\hat{\pi}, \check{\pi}_\varepsilon) = 0$ uniformly in the choice of Z_ε . The rest is identical to Proposition 3.6. $\square$

Proposition B.4. *For ε sufficiently small depending on X , any pair of points in $\text{supp}(\check{\pi}_\varepsilon)$ within distance $a_X/5$ are within distance $O(\sqrt{\delta}/A_X)$.*

Proof. By Proposition B.3, $D''_{\check{\pi}_\varepsilon, X}$ is negative on the interval between any two such points (for small enough ε). Lemma B.2 and simple calculus completes the proof. $\square$

For ε small enough that $\sqrt{\delta}/A_X \ll a_X$, Proposition B.4 implies that the graph of atoms in $\check{\pi}_\varepsilon$ under the distance-at-most- $a_X/5$ graph is a union of cliques. We form $\hat{\pi}_\varepsilon$ by taking the weighted average of each such clique, i.e. replacing $p_i\delta_{x_i}, \dots, p_j\delta_{x_j}$ by $\tilde{p}_i\delta_{\tilde{x}_i}$ for

$$\tilde{p}_i = p_i + \dots + p_j, \quad \tilde{x}_i = \frac{p_ix_i + \dots + p_jx_j}{x_i + \dots + x_j}.$$

It again follows from Proposition B.4 and the Wasserstein convergence $\check{\pi}_\varepsilon \rightarrow \hat{\pi}$ that $|\text{supp}(\hat{\pi}_\varepsilon)| = |\text{supp}(\hat{\pi})|$ for ε small enough, and that

$$\lim_{\varepsilon \rightarrow 0} d_{\mathcal{H}}(\text{supp}(\hat{\pi}_\varepsilon), \text{supp}(\hat{\pi})) = 0.$$

Proposition B.5. *For $\varepsilon \leq \varepsilon_0(X)$, the conditions of Proposition 3.1 apply to $\hat{\pi}_\varepsilon$ with:*

$$\begin{aligned} \tilde{\delta} &= O(L^2e^{4L^2}\delta/A_X) = O(\varepsilon^2L^4e^{8L^2}/A_X), \\ c_1 &= \sqrt[4]{CL\delta}, \\ c_2 &= A_X\sqrt{\delta}, \\ \eta &= 2c_1, \\ \lambda_\varepsilon &\triangleq \lambda(\hat{\pi}_\varepsilon) = \lambda(\hat{\pi}) \pm o_\varepsilon(1). \end{aligned}$$

In particular, the estimator $\hat{\pi}_\varepsilon$ obeys certifiable bounds of the form $\mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/4})$ as well as $|\text{supp}(\hat{\pi})| \geq |\text{supp}(\hat{\pi}_\varepsilon)|$.

Proof. The bound on δ for $\check{\pi}_\varepsilon$ is given in Lemma B.2. To extend it to $\hat{\pi}_\varepsilon$, the analogous argument in Proposition 3.9 reduces to showing that for $|x_i - x_{i+1}| \leq O(\sqrt{\delta/A_X})$,

$$\left| \left(\frac{p_j}{p_i + \dots + p_j} \right) e^{-|x_i - y|^2/2} + \dots + \left(\frac{p_j}{p_i + \dots + p_j} \right) e^{-|x_j - y|^2/2} - e^{-|\tilde{x}_i - y|^2/2} \right| \stackrel{?}{\leq} O(\delta/A_X).$$

Taylor's theorem applied to $f(x) = e^{-|x|^2/2}$ easily again gives the bound.

Continuing, we claim that whenever $2\varepsilon \leq c_1 \leq \frac{a_X}{2}$, we may take

$$c_2 = A_X \cdot \frac{(c_1 - \varepsilon)^2 - O(\sqrt{\delta})}{4}.$$

in Proposition 3.1. Indeed, if

$$d(y, \text{supp}(\hat{\pi}_\varepsilon)) \geq c_1,$$

then $d(y, \check{y}) \geq c_1 - \varepsilon$ for some $\check{y} \in \text{supp}(\check{\pi}_\varepsilon)$, so in particular $D_{\check{\pi}_\varepsilon, X}(\check{y}) \geq 1 - \delta$ by Lemma B.2. As an intermediate step, we upper bound $|D'_{\check{\pi}_\varepsilon, X}(\check{y})|$ using Lemma B.2(b) and bound (1.19) with $j = 2$. In particular, optimizing over $y \in \mathbb{R}$ in the first line to obtain the inequality $(\dagger)$, we find

$$\begin{aligned} D_{\check{\pi}_\varepsilon, X}(y) &\geq D_{\check{\pi}_\varepsilon, X}(\check{y}) + D'_{\check{\pi}_\varepsilon, X}(\check{y})(y - \check{y}) - C_2 L^2 e^{4L^2} (y - \check{y})^2, \quad \forall y \in \mathbb{R} \\ \implies \max_{z \in \mathbb{R}} D_{\check{\pi}_\varepsilon, X}(\check{y}) + 2\delta &\stackrel{\text{Lem. B.2}}{\geq} D_{\check{\pi}_\varepsilon, X}(z) \stackrel{(\dagger)}{\geq} D_{\check{\pi}_\varepsilon, X}(\check{y}) + \Omega\left(\frac{|D'_{\check{\pi}_\varepsilon, X}(\check{y})|^2}{L^2 e^{4L^2}}\right) \\ \implies |D'_{\check{\pi}_\varepsilon, X}(\check{y})| &\leq O(Le^{2L^2} \sqrt{\delta}). \end{aligned}$$

Using Proposition B.3 we now obtain (3.7):

$$\begin{aligned} D_{\hat{\pi}_\varepsilon, X}(y) &\leq D_{\check{\pi}_\varepsilon, X}(\check{y}) + O(Le^{L^2} \sqrt{\delta} |y - \check{y}|) - \frac{A_X(c_1 - \varepsilon)^2}{4} \\ &\leq 1 - \left(\frac{A_X(c_1 - \varepsilon)^2}{4} - O(Le^{L^2} a_X \sqrt{\delta}) \right) = 1 - c_2. \end{aligned}$$

In particular for ε small enough compared to A_X and C an absolute constant, we can set

$$c_1 = \sqrt[4]{CL\delta}, \quad c_2 = A_X \sqrt{\delta} \implies \eta = c_1 + \frac{L\tilde{\delta}}{c_2} \leq 2\sqrt[4]{CL\delta}.$$

Further, Proposition 3.2 ensures that (3.2) holds for ε small enough, since the value $\lambda(\hat{\pi}_\varepsilon)$ is easily seen to converge to that of $\hat{\pi}$. We conclude that Proposition 3.1 suffices to certify bounds of the form $\mathbb{W}_1(\hat{\pi}, \hat{\pi}_\varepsilon) \leq O_X(\varepsilon^{1/4})$. Finally, this implies $|\text{supp}(\hat{\pi})| \geq |\text{supp}(\hat{\pi}_\varepsilon)|$, again for ε small enough depending on X (and in particular $\hat{\pi}$). $\square$

Proposition B.6. *For ε small enough, Proposition 3.4 certifies $|\text{supp}(\hat{\pi})| \leq |\text{supp}(\hat{\pi}_\varepsilon)|$.*

Proof. Follows exactly as in the Proof of Theorem 1.9. $\square$

Proof of Theorem 1.10, except for the Shub–Smale property. The result follows by Propositions B.5, B.6, and 3.3. See Appendix D for the Shub–Smale property. $\square$

B.1 Finite S Case

Here we prove Theorem 1.15, except for the Shub–Smale property which is addressed in Appendix D. We use exactly the same version of Frank–Wolfe as above, but restricted to S rather than Z_ε . Thus we take

$$\pi_S^{(t+1)} = \frac{t\pi_S^{(t)} + 2\delta_{y_t}}{t+2},$$

$$y_t = \arg \max_{y \in S} D_{\pi_S^{(t)}, X}(y).$$

Define for $\iota > 0$ the subset $R_{S, \varepsilon, \iota} = \{y \in S : p_{\pi_S^{(t)}}(y) \leq \iota\}$. For $\iota \leq \frac{\varepsilon}{4L}$,

$$\sum_{y \in R_{S, \varepsilon, \iota}} p_{\pi_S^{(t)}}(y) \leq 3\iota L/\varepsilon < 1/2.$$

Thus we may define another probability measure

$$\check{\pi}_S^{(t)}(y) = \begin{cases} 0, & y \in R_{S, \varepsilon, \iota} \\ \frac{\pi_S^{(t)}(y)}{1 - \sum_{y \in R_{S, \varepsilon, \iota}} p_{\pi_S^{(t)}}(y)}, & y \notin R_{S, \varepsilon, \iota}. \end{cases}$$

The proof of Lemma B.1 extends unchanged to give the following. Fix any $X \in [-L, L]^n$ and $S \subseteq \mathbb{R}$ finite. For ε small enough depending only on L and for $t \geq \varepsilon^{-8}$, the probability measure $\check{\pi}_{S, \varepsilon} = \check{\pi}_S^{(t)}$ satisfies:

$$D_{\check{\pi}_{S, \varepsilon}, X}(y) - 1 \geq -O\left(\frac{e^{4L^2}}{\varepsilon^{1/2} \sqrt[4]{t}}\right) \geq -C(L)\varepsilon^{3/2}, \quad \forall y \in \text{supp}(\check{\pi}_S^{(t)}),$$

$$D_{\check{\pi}_{S, \varepsilon}, X}(y) - 1 \leq O\left(\frac{e^{4L^2}}{\varepsilon^{1/2} \sqrt[4]{t}}\right) \leq C(L)\varepsilon^{3/2}, \quad \forall y \in S,$$

$$\ell_X(\hat{\pi}_S) - \ell_X(\check{\pi}_{S, \varepsilon}) \leq C(L)\varepsilon^2.$$

Proof of Theorem 1.15, except for the Shub–Smale property. Similarly to the main case, let

$$B_{S, X} = 1 - \max_{y \in S \setminus \text{supp}(\hat{\pi}_S)} D_{\hat{\pi}_S, X}(y) > 0.$$

Since $\hat{\pi}_S$ is unique, we have $\check{\pi}_{S, \varepsilon} \rightarrow \hat{\pi}_S$ in Wasserstein as $\varepsilon \rightarrow 0$, which implies that for ε small enough we have $D_{\check{\pi}_{S, \varepsilon}, X}(y) \leq 1 - \frac{B_{S, X}}{2}$ for all $y \in S \setminus \text{supp}(\hat{\pi}_S)$.

The remaining proof is similar to before, and we just give an outline. The initial steps follow the proof of Proposition 3.1. First one can certify that $\hat{\pi}$ has at most ε mass outside $\text{supp}(\check{\pi}_{S, \varepsilon})$, using the bound on $\|\ell''\|_\infty$ for $\ell(t) = \ell_X((1-t)\check{\pi}_{S, \varepsilon} + t\hat{\pi}_S)$. This is because

$$\ell'(0) = \mathbb{E}^{y \sim \hat{\pi}_S} [D_{\check{\pi}_{S, \varepsilon}, X}(y)].$$

Again (3.1) holds for some positive $\lambda = \lambda(S, X)$ (with S and $\hat{\pi}_S$ in place of Z_ε and $\tilde{\pi}_\varepsilon$) by uniqueness of $\hat{\pi}_S$. Then (3.1) will hold with constant $\lambda/2$ for $\check{\pi}_{S, \varepsilon}$ with small enough ε . This allows to certify that all approximate local maxima of ℓ_X supported on $\text{supp}(\check{\pi}_{S, \varepsilon})$ are within $O(\lambda)$ of $\check{\pi}_{S, \varepsilon}$ (using now the fact that ℓ has 3 bounded derivatives). Since we could also certify above that $\hat{\pi}$ has at most ε mass outside $\text{supp}(\check{\pi}_{S, \varepsilon})$,

these certificates thus combine to certify that $\mathbb{W}_1(\hat{\pi}_S, \check{\pi}_{S,\varepsilon}) \leq C(S, X)\varepsilon$ almost surely for small enough ε (again similarly to Proposition 3.1).

Since for small ε we will have $D_{\check{\pi}_{S,\varepsilon}, X}(y) \leq 1 - \frac{B_{S,X}}{2}$ for all $y \in S \setminus \text{supp}(\hat{\pi}_S)$, applying (1.20) to the Wasserstein bound certifies that

$$D_{\hat{\pi}_S, X}(y) \leq 1 - \frac{B_{S,X}}{4}, \quad \forall y \in S \setminus \text{supp}(\hat{\pi}_S).$$

In particular, this certifies that $\text{supp}(\hat{\pi}_S) \subseteq \text{supp}(\check{\pi}_{S,\varepsilon})$. The opposite inclusion is easily certified from $\mathbb{W}_1(\hat{\pi}_S, \check{\pi}_{S,\varepsilon}) \leq C(S, X)\varepsilon$ once ε is small compared to $p_* d_*$, where p_* is the smallest atom size in $\hat{\pi}_S$ and d_* is the minimum distance between points in S . Finally, a $\mathbb{W}_1$ bound immediately gives a $d_{\Pi_k(S)}$ bound as well. $\square$

C Positive Probability to be k -Atomic

We show that for any $n \geq k \geq 1$, there is a positive probability to have $\hat{\pi} \in \Pi_k$ with $L \leq O(k\sqrt{\log k})$, as mentioned in Remark 1.6. Thus shows that conditioning simultaneously on the events $\hat{\pi} \in \Pi_k$ and $\max_{i,j} (x_i - x_j) \leq cn^{1/4}$ is possible, for n polynomially large in k .

We partition $[n]$ into k parts $S_1, \dots, S_k$ of sizes $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$. We fix a large absolute constant C , and set $\tilde{x}_i = Ci\sqrt{\log(k+1)}$ for each $1 \leq i \leq k$. Let us suppose that $|x_a - \tilde{x}_i| \leq 0.1$ holds for all $a \in S_i$. This is a positive probability event, where roughly speaking the datapoints are clustered into groups of approximately equal size. We will show that this implies $|\text{supp}(\hat{\pi})| = k$. Precisely, $\hat{\pi}$ has exactly 1 atom close to each $\tilde{x}_i$, and no other atoms.

Lemma C.1. *For each $i \in [k]$, $\text{supp}(\hat{\pi}) \subseteq \bigcup_i [\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]$.*

Proof. Recall that the function $e^{-y^2/2}$ is convex outside of $[-1, 1]$. First, suppose $\hat{\pi}$ has an atom $p\delta_y$ with $\min_i |y - \tilde{x}_i| \geq 2$. Then we may replace $p\delta_y$ by $\frac{p\delta_{y+0.5} + p\delta_{y-0.5}}{2}$, and this will increase the value of $P_\pi(x_a)$ for each $a \in [n]$. This contradicts optimality of $\hat{\pi}$ in (1.2).

Next, suppose $\hat{\pi}$ has an atom $p\delta_y$ with $\min_i |y - \tilde{x}_i| \in [0.2, 2]$. Without loss of generality, suppose $y - \tilde{x}_i \in [0.2, 2]$. Then we can replace $p\delta_y$ in $\hat{\pi}$ by

$$p(1-c)\delta_{y+1} + pc\delta_{y-0.1}.$$

For c a small enough absolute constant, this will increase $P_\pi(x_a)$ for all $a \in S_i$. Then choosing C sufficiently large ensures $P_\pi(x_a)$ increases for all other $a \notin S_i$ as well. $\square$

Lemma C.2. *For each $i \in [k]$, $\hat{\pi}([\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]) \geq \frac{1}{20k}$.*

Proof. We claim that if $\hat{\pi}$ violates this condition for some i , then $D_{\hat{\pi}, X}(\tilde{x}_i) > 1$, which contradicts Proposition 1.1. Indeed in this case, Lemma C.1 implies the set $\hat{\pi}([\tilde{x}_i - 6\sqrt{\log(k+1)}, \tilde{x}_i + 6\sqrt{\log(k+1)}]) \leq \frac{1}{20k}$. Then it easily follows that $P_{\hat{\pi}}(x_a) \leq \frac{1}{10k}$ for each $a \in S_i$. We conclude that:

$$D_{\hat{\pi}, X}(\tilde{x}_i) \geq \frac{10k|S_i|}{n} \min_{x \in [\tilde{x}_i - 0.1, \tilde{x}_i + 0.1]} e^{-|x - \tilde{x}_i|^2/2} / \sqrt{2\pi} \geq 5e^{-0.005} / \sqrt{2\pi} > 1.$$

This is the desired contradiction. $\square$

Lemma C.3. *For each $i \in [k]$, $|\text{supp}(\hat{\pi}) \cap [\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]| = 1$. Thus $|\text{supp}(\hat{\pi})| = k$.*

Proof. We use the previous lemmas to show $D_{\hat{\pi},X}$ is concave on each interval $[\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]$. Fixing i , we decompose

$$D_{\hat{\pi},X}(y)n\sqrt{2\pi} = \sum_{j=1}^k \sum_{a \in S_j} e^{-|x_a - y|^2/2} / P_{\hat{\pi}}(x_a) \triangleq \sum_{j=1}^k D_j(y).$$

We first study $D_i(y)$. Trivially, $P_{\hat{\pi}}(x_a)\sqrt{2\pi} \leq 1$ for all $a \in [n]$. Furthermore, the function $f(y) = e^{-y^2/2}$ satisfies $f''(y) \leq 1/100$ on $|y| \leq 0.2$. Therefore one easily finds

$$D_i''(y) \leq -\frac{|S_i|}{100} \leq -\frac{n}{200k}, \quad \forall y \in [\tilde{x}_i - 0.2, \tilde{x}_i + 0.2].$$

Next we study $D_j''(y)$ for $j \neq i$. Here, Proposition C.2 implies that $P_{\hat{\pi}}(x_a) \geq \frac{1}{100k}$ for all $a \in [n]$. Therefore for all $y \in [\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]$:

$$D_j''(y) \geq -100k|S_j|e^{-C^2 \log(k+1)/2} \geq -\frac{S_j}{100(k+1)^{10}} \geq -\frac{n}{10(k+1)^8}.$$

Summing over $j \neq i$, it is clear that $D_i''(y)$ dominates, so $D_{\hat{\pi}}''(y) < 0$ for all $y \in [\tilde{x}_i - 0.2, \tilde{x}_i + 0.2]$. By Proposition 1.1, $\hat{\pi}$ is supported within the set of global maxima of $D_{\hat{\pi}}$, which completes the proof. $\square$

Combining the lemmas above immediately yields the claimed result.

C.1 Example of Non-Generic Behavior

Here we prove Theorem 1.7. With $n = mk$, we restrict to the positive probability event that X satisfies the conditions of Lemmas C.1, C.2, C.3, i.e. $|x_a - \tilde{x}_i| \leq 0.1$ for all $a \in S_i$. For convenience, we say such X is k -good; this condition implies $\hat{\pi} \in \Pi_k$. We claim that conditioned on $(x_{m+1}, \dots, x_n)$, the conditional law of $\hat{\pi}$ does not admit a density on Π_k . The proof is motivated by Remark 2.5. The main remaining step is the smooth dependence of $\hat{\pi}$ on $(x_1, \dots, x_k)$. To show it, we rely on the following standard version of the implicit function theorem.

Lemma C.4. *Let $U, V \subseteq \mathbb{R}^d$ be open sets and $F : U \times V \rightarrow \mathbb{R}^d$ a smooth function. Let $u_* \in U$ and $v_* \in V$ be such that $\partial_v F(u_*, v_*) \in \mathbb{R}^{d \times d}$ is invertible. Then on a neighborhood $U_0 \subseteq U$ of u_* , there exists a smooth function $v : U_0 \rightarrow V$ such that*

$$F(u, v(u)) = F(u_*, v_*)$$

holds for all $u \in U_0$.

The following lemma shows that Theorem 1.2 parts (II) and (III) imply Theorem 1.3(B). Theorem 1.3(A) also follows from the proof, as the Hessian of ℓ_X is the Jacobian of γ_k within, which is shown to be strictly negative definite.

Lemma C.5. *Suppose $\hat{\pi}$ is the NPMLE for X and satisfies the conclusions of Theorem 1.2 parts (II) and (III). Then for $\tilde{X}$ in a sufficiently small neighborhood of X , the function $\tilde{X} \mapsto \hat{\pi}(\tilde{X})$ has image in Π_k , and is smooth. Here we consider Π_k to be an open subset of a $2k - 1$ dimensional real vector space.*

Proof. It is clear that $\hat{\pi}(\tilde{X}) \in \Pi_k$ for $\tilde{X}$ close to X . We will apply Lemma C.4 with $d = 2k$. We consider the Jacobian of the map

$$\pi = \sum_{i=1}^k p_i \delta_{y_i} \mapsto \gamma_k(\pi) = (D_{\pi}(y_1) - 1, \dots, D_{\pi}(y_k) - 1, D'_{\pi}(y_1), \dots, D'_{\pi}(y_k)) \in \mathbb{R}^{2k}.$$

Here we temporarily relax the constraint $\sum_{i=1}^k p_i = 1$ so the Jacobian is a $2k \times 2k$ matrix J_{2k} . We show below that this Jacobian is invertible. Then Lemma C.4 applied to $\gamma_k(\pi, \tilde{X})$ implies existence of a smooth function $\pi(\tilde{X})$ such that $\gamma_k(\pi(\tilde{X}), \tilde{X}) = \vec{0}$ in a neighborhood of $(\hat{\pi}, X)$. Further, any such solution satisfies $\sum_{i=1}^k p_i D_\pi(y_i) = \sum_{i=1}^k p_i = 1$, so $\pi(\tilde{X})$ is always a probability measure (despite the relaxation to $2k$ dimensions in applying the inverse function theorem).

We turn to invertibility of the Jacobian. The entries of $-J_{2k}$ are, for $1 \leq i, j \leq k$:

$$\begin{aligned} -\frac{dD_\pi(y_j)}{dp_i} &= \frac{1}{n} \sum_{\ell=1}^n \frac{\exp\left(-\frac{(x_\ell - y_j)^2 + (x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)^2}; \\ -\frac{dD'_\pi(y_j)}{dp_i} &= \frac{1}{n} \sum_{\ell=1}^n \frac{(x_\ell - y_j) \exp\left(-\frac{(x_\ell - y_j)^2 + (x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)^2}; \\ -\frac{dD_\pi(y_j)}{dy_i} &= \frac{1}{n} \sum_{\ell=1}^n \frac{(x_\ell - y_i) \exp\left(-\frac{(x_\ell - y_j)^2 + (x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)^2} \\ &\quad - \underbrace{\frac{1}{n} \sum_{\ell=1}^n \frac{1_{i=j}(x_\ell - y_i) \exp\left(-\frac{(x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)}}_{=D'_\pi(y_i)=0}; \\ -\frac{dD'_\pi(y_j)}{dy_i} &= \frac{1}{n} \sum_{\ell=1}^n \frac{(x_\ell - y_i)(x_\ell - y_j) \exp\left(-\frac{(x_\ell - y_j)^2 + (x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)^2} \\ &\quad - \underbrace{\frac{1}{n} \sum_{\ell=1}^n \frac{1_{i=j}((x_\ell - y_i)^2 - 1) \exp\left(-\frac{(x_\ell - y_i)^2}{2}\right)}{P_\pi(x_\ell)}}_{=D''_\pi(y_i)<0}. \end{aligned}$$

Note the $1_{i=j}$ terms come because if $i = j$, we vary the input to D_π .

We will show $-J_{2k}$ is strictly positive definite, hence invertible. We first show the $k \times k$ submatrix $-J_k$ formed by the $\frac{dD_\pi(y_j)}{dp_i}$ terms above is strictly positive definite. Indeed, we see from the above that

$$\frac{dD_\pi(y_j)}{dp_i} = A_i A_j \sum_{\ell=1}^n C_\ell e^{x_\ell(y_i + y_j)}$$

for strictly positive constants A_i and C_ℓ . Thus $-J_k$ is positive semi-definite, represented as a positive combination of rank 1 matrices $(\vec{v}^{(\ell)})^{\otimes 2}$ for $\vec{v}_i^{(\ell)} = e^{x_\ell y_i}$. We claim $-J_k$ is *strictly* positive-definite. Indeed, Lemma 2.6 shows the vectors $\vec{v}^{(1)}, \dots, \vec{v}^{(n)}$ are linearly independent, i.e. that if the real coefficients $B_1, \dots, B_k$ satisfy

$$G(x) \triangleq \sum_{j=1}^k B_j e^{x y_j} = 0 \quad \forall x \in \{x_1, \dots, x_n\}$$

then $B_1 = \dots = B_k = 0$.

Next, using the simplifications indicated by identifying values of D'_π and D''_π above, we see that $-J_{2k}$ is positive semi-definite via another explicit representation as a sum of n rank 1 matrices, together with the

diagonal terms from D''_π appearing in the formula for $\frac{dD'_\pi(y_j)}{dy_i}$. Since all of these diagonal terms are non-zero, it easily follows that all of $-J_{2k}$ is strictly positive definite. Indeed, we can write $-J_{2k} = M + M'$ where M' includes only the diagonal D''_π terms. Then if $\langle v, (M + M')v \rangle = 0$ for some $v \in \mathbb{R}^{2k}$, we see that v cannot have any non-zero entry interacting with M' since M is already positive semi-definite. But then we are reduced to the strict positive definiteness of $-J_k$ shown above! $\square$

Lemma C.6. *Suppose $\hat{\pi}$ is the NPMLE for X and satisfies $\min_{y \in \text{supp}(\hat{\pi})} D''_{\hat{\pi}, X}(y) < 0$. Suppose further that $x_1 = \dots = x_m$. For $\varepsilon > 0$ and $Z = (z_1, \dots, z_m) \in [0, 1]^m$, let $\hat{\pi}_Z = \hat{\pi}(x_1 + \varepsilon z_1, \dots, x_m + \varepsilon z_m, x_{m+1}, \dots, x_n)$. For ε sufficiently small, $\hat{\pi}_Z$ lies in Π_k and has a Taylor expansion of the form:*

$$\hat{\pi}_Z = \hat{\pi} + \varepsilon f_1(P_1(Z)) + \varepsilon^2 f_2(P_1(Z), P_2(Z)) + \dots + \varepsilon^m f_k(P_1(Z), \dots, P_m(Z)) + O(\varepsilon^{m+1})$$

where $P_j(Z) = \sum_{i=1}^m z_i^j$ is a power sum symmetric polynomial and each function f_i is Lipschitz. Here we consider Π_k to be an open subset of a $2k - 1$ dimensional real vector space.

Proof. Lemma C.5 implies that $\hat{\pi}_Z \in \Pi_k$ depends smoothly on Z . Its k -th Taylor coefficient is a degree k polynomial in Z . Since $x_1 = \dots = x_m$, it must be symmetric in $(z_1, \dots, z_m)$, hence is a polynomial in $(P_1(Z), \dots, P_k(Z))$ (see e.g. [SF99, Chapter 7]). This easily gives the form of the Taylor expansion above. $\square$

Proof of Theorem 1.7. We begin with $(X, \hat{\pi})$ as in Lemma C.6 such that $\hat{\pi}$ has k atoms. This can be arranged by taking $x_1 = \dots = x_m$ to be fixed and the remaining data generic. Then Section C shows $|\text{supp}(\hat{\pi})| = k$ with positive probability, while since $n - m > (2k + 2)^2$, Theorem 1.2 implies the remaining hypothesis holds almost surely.

Next, take $\varepsilon > 0$ small and consider the set of achievable $\hat{\pi}_Z$ as above. Lemma C.6 implies this set of $\hat{\pi}_Z$ is contained in the $O(\varepsilon^{2k})$ neighborhood of the image of $[0, \varepsilon] \times [0, \varepsilon^2] \times \dots \times [0, \varepsilon^{2k-1}]$ under an $O(1)$ -Lipschitz function (i.e. with Lipschitz constant independent of ε). This is because the Taylor expansion above is $O(\varepsilon^j)$ -Lipschitz in $p_j(Z)$, which is bounded by a constant that does not depend on ε (thus rescaling gives the claim).

This set of possible $\hat{\pi}_Z$ can be covered by at most $O(\varepsilon^{-k(2k-1)})$ balls of radius ε^{2k-1} , hence has volume $O(\varepsilon^{(2k-1)^2 - k(2k-1)}) = O(\varepsilon^{(k-1)(2k-1)})$. On the other hand, the volume of the corresponding set of inputs $(x_1 + \varepsilon z_1, \dots, x_m + \varepsilon z_m)$ is ε^m . Hence in order for $\hat{\pi}$ to have locally bounded density, we must have $m \leq (k - 1)(2k - 1)$. This completes the proof. $\square$

D Asymptotic Convergence From Approximate Solutions

Here we explain the rapid asymptotic convergence of Newton–Raphson and Expectation–Maximization starting from a good approximate solution.

D.1 Approximate Shub–Smale Solutions

Given our existing work in proving Theorem 1.9, the Shub–Smale approximate solution property follows from standard results in numerical analysis. The notation in the following result is adapted to our setting.

Proposition D.1 ([SB13, Theorem 5.3.2]). *Let $B_r(\pi) \subseteq \Pi_k$ be the r -neighborhood of $\pi \in \Pi_k \simeq \mathbb{R}^{2k-1}$. Let $\bar{B}_r(\pi)$ be the closure of $B_r(\pi)$ in $\mathbb{R}^{2k-1}$ and assume $\bar{B}_r(\pi) \subseteq \Pi_k$ (i.e. the distance from π to the*

boundary $\partial\Pi_k$ is larger than r). Let $F : \bar{B}_r(\pi) \rightarrow \mathbb{R}^{2k-1}$ be a smooth function. Suppose the positive constants r, α, β, C satisfy:

$$h \triangleq \alpha\beta C < 1;$$

$$r = \frac{\alpha}{1-h}.$$

Suppose ℓ_X and its gradient and Hessian (in the space Π_k) satisfies:

1. $\nabla^2 \ell_X$ is C -Lipschitz on $B_r(\pi)$.
2. $\nabla^2 \ell_X \preceq -I_{2k-1}/\beta$ uniformly on $B_r(\pi)$.
3. $\|[\nabla^2 \ell_X(\pi)]^{-1} \nabla \ell_X(\pi)\| \leq \alpha$.

Then Newton–Raphson iteration starting from $\pi_0 = \pi$ converges to a limit $\pi_\infty \in \Pi_k$ such that for each $t \geq 0$:

$$d_{\Pi_k}(\pi_t, \pi_\infty) \leq \alpha \frac{h^{2^t-1}}{1-h^{2^t}}. \quad (\text{D.1})$$

Proof of Shub–Smale property in Theorems 1.9 and 1.10. We show $\hat{\pi}_\varepsilon$ is a Shub–Smale approximate NPMLE for small enough ε (the proof for the adjacent-atom rounding of $\hat{\pi}_\varepsilon$ is identical). From the proof of the previous parts of Theorem 1.9, we know that $d_{\Pi_k}(\hat{\pi}_\varepsilon, \hat{\pi}) \leq \varepsilon^{1/4}$ for small enough ε . We have shown via Lemma C.5 that Theorem 1.3(A) holds, so ℓ_X is c -strongly concave in an open Π_k -neighborhood of $\hat{\pi}$, for some $c > 0$. Additionally, this Hessian is clearly C -Lipschitz for some $C = C(L)$ thanks to e.g. Lemma 1.18. This means local $c/2$ -strong concavity of ℓ_X is certifiable in a Π_k -neighborhood of $\hat{\pi}_\varepsilon$ which certifiably contains $\hat{\pi}$, for ε small enough.

We will apply Proposition D.1 to such $\pi = \hat{\pi}_\varepsilon$ and conclude the desired result. Indeed we can take $C = C(L)$ as mentioned just above, and $\beta = 2/c$. Neither depends on ε , and so we can take $\alpha \leq \varepsilon^{1/5}$ for small enough ε . Then $h \leq 1/10 < 1$ and $r \leq 2\alpha \leq 2\varepsilon^{1/5}$ is smaller than the Π_k -distance from $\hat{\pi}_\varepsilon$ to $\partial\Pi_k$, for small ε . Then (D.1) implies quadratic convergence to a limit $\pi_\infty \in \Pi_k$ with $\nabla \ell_X(\pi_\infty) = \vec{0}$ and $d_{\Pi_k}(\pi_0, \pi_\infty) \leq 2\alpha$. From the local strong concavity, it follows that $\pi_\infty = \hat{\pi}$. The precise quantitative rate in Definition 1 holds because $h \leq 1/10$. This concludes the proof. $\square$

D.2 Static Support Case

Here we explain Theorems 1.14 and 1.15, which are the static support analogs of Theorems 1.3 and 1.9, 1.10.

Proof of Theorem 1.14. Here the Hessian of ℓ_X is just the submatrix J_k from the proof of Lemma C.5, which is strictly negative definite for all $1 \leq k \leq n$ (with no genericity conditions on D). This implies both results. $\square$

Proof of Shub–Smale property in Theorem 1.15. Given the local strong concavity around $\hat{\pi}_S$ from Theorem 1.14, the proof is identical to the corresponding parts of Theorems 1.9 and 1.10. $\square$

D.3 Locally Linear Convergence of Expectation–Maximization

The Expectation–Maximization (EM) algorithm of [DLR77] is another approach to optimize ℓ_X . Here one defines a sequence of iterates $\pi_0, \pi_1, \dots \in \Pi_k$ as follows. Given an iterate $\pi_t = \sum_{i=1}^k p_i \delta_{y_i}$, the next iterate π_{t+1} has parameters:

$$\tilde{p}_i = \frac{p_i}{n} \sum_{j=1}^n \frac{e^{-(x_j - y_i)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}, \quad (\text{D.2})$$

$$\tilde{y}_i = \frac{\sum_{j=1}^n \frac{x_j e^{-(x_j - y_i)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}}{\sum_{j=1}^n \frac{e^{-(x_j - y_i)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}} = y_i + \frac{\sum_{j=1}^n \frac{(x_j - y_i) e^{-(x_j - y_i)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}}{\sum_{j=1}^n \frac{e^{-(x_j - y_i)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}} = y_i + \frac{D'(y_i)}{D(y_i)}. \quad (\text{D.3})$$

This iteration improves the value of ℓ_X in each step; however it is not guaranteed to converge to a global optimum of ℓ_X even within the space Π_k . On the other hand, one can verify using (1.6) that the NPMLE $\hat{\pi}$ is a fixed point of the iteration.

Proof of Theorem 1.3(C). Let $F(\pi) = \tilde{\pi}$ correspond to the iteration from (D.2), (D.3). Identifying Π_k with an open subset of $\mathbb{R}^{2k}$ (i.e. not enforcing the constraint $\sum_{i=1}^k p_i = 1$), we will show that $JF(\hat{\pi})$ is strictly stable, with all eigenvalues in $(-1, 1)$. This immediately implies the same property on the $2k - 1$ dimensional space Π_k (since the output of (D.2) immediately satisfies $\sum_i p_i = 1$). It is well-known that this stability implies locally linear convergence rate for the EM algorithm; see e.g. [DLR77, MR94].

We first explicitly compute the Jacobian $JF(\hat{\pi})$ at $\hat{\pi}$:

$$\begin{aligned} \frac{\partial \tilde{p}_i}{\partial p_r} &= 1_{i=r} - \frac{p_i}{n} \sum_{j=1}^n \frac{e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}; \\ \frac{\partial \tilde{p}_i}{\partial y_r} &= \underbrace{\frac{p_i}{n} \sum_{j=1}^n \frac{(x_j - y_r) e^{-(x_j - y_r)^2/2}}{\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}}_{=D'(y_r)=0} - \frac{p_i}{n} \sum_{j=1}^n \frac{p_r (x_j - y_r) e^{-(x_j - y_r)^2/2 - (x_j - y_i)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}; \\ \frac{\partial \tilde{y}_i}{\partial p_r} &= -\frac{1}{n} \sum_{j=1}^n \frac{(x_j - y_i) e^{-(x_j - y_r)^2/2 - (x_j - y_i)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}, \\ \frac{\partial \tilde{y}_i}{\partial y_r} &= 1_{i=r} \cdot (1 + D''(y_i)) - \frac{1}{n} \sum_{j=1}^n \frac{p_r (x_j - y_r) (x_j - y_i) (e^{-(x_j - y_r)^2/2 - (x_j - y_i)^2/2})}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}. \end{aligned}$$

Here the factors of $1/n$ in the latter two lines come from $D(y_i) = 1$ for all i . We note that another term which equals 0 for $\hat{\pi}$ was omitted from the formula for $\frac{\partial \tilde{y}_i}{\partial p_r}$, again because $D'(y_r) = 0$ holds at $\hat{\pi}$.

The terms $1_{i=r}$ cancel I_{2k} . Removing the first term in $\frac{\partial \tilde{p}_i}{\partial y_r}$ as indicated above, we may write:

$$\begin{aligned} I_{2k} - JF + \text{diag}(0, \dots, 0, D''(y_1), \dots, D''(y_k)) \\ = \begin{bmatrix} p_i/n & p_i p_r/n \\ 1/n & p_r/n \end{bmatrix} \odot \sum_{j=1}^n \left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^{-2} \begin{bmatrix} e^{-(x_j - y_i)^2/2} \\ (x_j - y_i) e^{-(x_j - y_i)^2/2} \end{bmatrix}^{\otimes 2}. \end{aligned}$$

Here the former matrix is $2k \times 2k$, with (i, r) entry equal to p_i/n , $(i, k+r)$ entry equal to $\frac{p_i p_r}{n}$, etc. The latter column vector has length $2k$ (and depends on j). The first k entries are given by $\frac{p_i}{n} e^{-(x_j - y_i)^2/2}$ for

$1 \leq i \leq k$, and the next k entries are given by $p_i(x_j - y_i)e^{-(x_j - y_i)^2/2}$ for $1 \leq i \leq k$. Note that the matrix $\begin{bmatrix} p_i/n & p_i p_r/n \\ 1/n & p_r/n \end{bmatrix}$ is similar to $\begin{bmatrix} p_i/n & \frac{\sqrt{p_i p_r}}{n} \\ \frac{\sqrt{p_i p_r}}{n} & p_r/n \end{bmatrix}$ via conjugation by $M = \text{diag}(\sqrt{p_1}, \dots, \sqrt{p_k}, 1, \dots, 1)$; in particular their eigenvalues are equivalent. This conjugation extends to the preceding display, and shows

$$I_{2k} - JF + \text{diag}(0, \dots, 0, D''(y_1), \dots, D''(y_k)) \triangleq I_{2k} - JF + D$$

is similar to the positive semi-definite matrix $M^{-1}(I_{2k} - JF + D)M$ given explicitly by:

$$\sum_{j=1}^n \left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^{-2} \left[\begin{array}{c} \sqrt{\frac{p_i}{n}} e^{-(x_j - y_i)^2/2} \\ \sqrt{\frac{p_i}{n}} (x_j - y_i) e^{-(x_j - y_i)^2/2} \end{array} \right]^{\otimes 2}. \quad (\text{D.4})$$

We claim that $M^{-1}(I_{2k} - JF)M$ is strictly positive semi-definite. Since $D''(y_i) < 0$ for each i , it suffices to verify that the upper-left $k \times k$ block of (D.4) is strictly positive-definite. This amounts to proving that the n vectors $(e^{-(x_j - y_j)^2/2})_{i=1}^k$ (for $1 \leq j \leq n$) are linearly independent. As usual, this is equivalent to linear independence of the vectors $(e^{x_j y_i})_{i=1}^k$ which holds for any $k \leq n$ by Lemma 2.6.

Next we need to show all eigenvalues of JF are larger than -1 , or equivalently that all eigenvalues of $M^{-1}(I_{2k} - JF)M$ are strictly smaller than 2. For each j , we upper-bound the rank 1 term in (D.4) by replacing (p_i, p_r) entries of the form $A_{i,r,j} A_{r,i,j}$ by diagonal entries $A_{i,r,j}^2$ and $A_{r,i,j}^2$ in the (p_i, p_i) and (p_r, p_r) positions respectively, which always gives an upper bound in the positive semi-definite order. In particular, the (p_i, p_r) and (p_r, p_i) diagonal entries

$$\frac{\frac{\sqrt{p_i p_r}}{n} \cdot e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}$$

can be upper-bounded via replacement by the respective (p_i, p_i) and (p_r, p_r) entries:

$$\frac{\frac{p_r}{n} \cdot e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}, \quad \frac{\frac{p_i}{n} \cdot e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}. \quad (\text{D.5})$$

(Note that the p_i factor appears in the (p_r, p_r) position, and vice-versa.) Similarly the (p_i, y_r) and (y_r, p_i) entries

$$\frac{\frac{\sqrt{p_i p_r}}{n} \cdot (x_j - y_r) e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}$$

can be upper-bounded by the respective (p_i, p_i) and (y_r, y_r) entries

$$\frac{\frac{p_r}{n} \cdot e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}, \quad \frac{\frac{p_i}{n} \cdot (x_j - y_r)^2 e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2},$$

and the (y_i, y_r) and (y_r, y_i) entries

$$\frac{\frac{\sqrt{p_i p_r}}{n} \cdot (x_j - y_i)(x_j - y_r) e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2} \right)^2}$$

can be upper-bounded by the respective (y_i, y_i) and (y_r, y_r) entries

$$\frac{\frac{p_r}{n} \cdot e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}, \quad \frac{\frac{p_i}{n} \cdot (x_j - y_r)^2 e^{-(x_j - y_i)^2/2 - (x_j - y_r)^2/2}}{\left(\sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}\right)^2}.$$

After making all of these substitutions, the resulting diagonal matrix has (p_i, p_i) and (y_i, y_i) entries:

$$\frac{2e^{-(x_j - y_i)^2/2}}{n \sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}, \quad \frac{2(x_j - y_i)^2 e^{-(x_j - y_i)^2/2}}{n \sum_{\ell=1}^k p_\ell e^{-(x_j - y_\ell)^2/2}}.$$

Summing over j , the resulting matrix has (p_i, p_i) entry $2D(x_j) = 2$ and (y_i, y_i) entry $2 + 2D''(y_i)$. Since this matrix is an upper bound for $M^{-1}(I_{2k} - JF + D)M$ in the positive semidefinite order, we conclude that

$$M^{-1}(I_{2k} - JF)M \preceq \text{diag}(2, \dots, 2, 2 + D''(y_1), \dots, 2 + D''(y_k)). \quad (\text{D.6})$$

This shows the desired inequality, again except for strictness in the $(p_1, \dots, p_k)$ subspace. Note that for $v = (v_1, \dots, v_{2k}) \neq \vec{0}$ to satisfy

$$\langle M^{-1}(I_{2k} - JF)M, v^{\otimes 2} \rangle = \langle \text{diag}(2, \dots, 2, 2 + D''(y_1), \dots, 2 + D''(y_k)), v^{\otimes 2} \rangle,$$

we must have $v_{k+1} = \dots = v_{2k} = 0$, since $D''(y_i) < 0$ for each i . However then the contributions from (p_i, y_r) entries to the left-hand expression are zero. Hence for such v the only contribution comes from (D.5) so we actually have

$$\langle M^{-1}(I_{2k} - JF)M, v^{\otimes 2} \rangle \leq \|v\|^2.$$

Combining, we see that (D.6) is strict, i.e. the difference is strictly positive definite. This completes the proof. $\square$