

Estimating the Coefficients of a Mixture of Two Linear Regressions by Expectation Maximization

Jason M. Klusowski¹, Student Member, IEEE, Dana Yang, and W. D. Brinda²

Abstract—We give convergence guarantees for estimating the coefficients of a symmetric mixture of two linear regressions by expectation maximization (EM). In particular, we show that the empirical EM iterates converge to the target parameter vector at the parametric rate, provided the algorithm is initialized in an unbounded cone. In particular, if the initial guess has a sufficiently large cosine angle with the target parameter vector, a sample-splitting version of the EM algorithm converges to the true coefficient vector with high probability. Interestingly, our analysis borrows from tools used in the problem of estimating the centers of a symmetric mixture of two Gaussians by EM. We also show that the population of EM operator for mixtures of two regressions is anti-contractive from the target parameter vector if the cosine angle between the input vector and the target parameter vector is too small, thereby establishing the necessity of our conic condition. Finally, we give empirical evidence supporting this theoretical observation, which suggests that the sample-based EM algorithm may not converge to the target vector when initial guesses are drawn accordingly. Our simulation study also suggests that the EM algorithm performs well even under model misspecification (i.e., when the covariate and error distributions violate the model assumptions).

Index Terms—Mixture models, expectation-maximization algorithm, iterative algorithms, clustering algorithms, regression analysis.

I. INTRODUCTION

MIXTURES of linear regressions are useful for modeling different linear relationships between input and response variables across several unobserved heterogeneous groups in a population. First proposed in [1] as a generalization of “switching regressions”, this model has found broad applications in areas such as plant science [2], musical perception theory [3], [4], and educational policy [5].

In this paper, we consider estimating the model parameters in a symmetric two component mixture of linear regressions. Towards a theoretical understanding of this model, suppose

Manuscript received September 13, 2017; revised October 15, 2018; accepted December 22, 2018. Date of publication January 9, 2019; date of current version May 20, 2019. This work was supported by the Graduate Student Fellowship from Yale University.

J. M. Klusowski is with the Department of Statistics, Rutgers University—New Brunswick, Piscataway, NJ 08854-8019 USA (e-mail: jason.klusowski@rutgers.edu).

D. Yang and W. D. Brinda are with the Department of Statistics and Data Science, Yale University, New Haven, CT 06511 USA (e-mail: xiaoqian.yang@yale.edu; william.brinda@yale.edu).

Communicated by R. Balan, Associate Editor for Detection and Estimation. Digital Object Identifier 10.1109/TIT.2019.2891628

we observe data $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, where

$$Y_i = R_i \langle \theta^*, X_i \rangle + \varepsilon_i, \quad (1)$$

$X_i \stackrel{i.i.d.}{\sim} N(0, I_d)$, $\varepsilon_i \stackrel{i.i.d.}{\sim} N(0, \sigma^2)$, $R_i \stackrel{i.i.d.}{\sim} \text{Rademacher}(1/2)$, and $\{X_i\}$, $\{\varepsilon_i\}$, and $\{R_i\}$ are independent of each other. In other words, each predictor variable is Gaussian, and the response is centered at either the θ^* or $-\theta^*$ linear combination of the predictor. The two classes are equally probable, and the label of each observation is unknown. We seek to estimate θ^* (or $-\theta^*$, which produces the same model distribution).

The likelihood function of the model, $\mathcal{L}(\theta; \mathcal{D}_n)$, can be expressed as

$$\prod_{i=1}^n \left[\frac{1}{2} \psi(X_i) \psi_\sigma(Y_i - \langle \theta, X_i \rangle) + \frac{1}{2} \psi(X_i) \psi_\sigma(Y_i + \langle \theta, X_i \rangle) \right],$$

where $\psi(x) = \frac{1}{(2\pi)^{d/2}} e^{-\|x\|^2/2}$ and $\psi_\sigma(y) = \frac{1}{\sqrt{2\pi}\sigma} e^{-y^2/(2\sigma^2)}$, is a multi-dimensional, multi-modal (it has many spurious local maxima), and nonconvex objective function, and hence direct maximization (e.g., grid search) is intractable. Even the population likelihood (in the infinite data setting) has global maxima at $-\theta^*$ and θ^* , and a local minimum at the zero vector. Given these computational concerns, other less expensive methods have been used to estimate the model coefficients. For example, mixtures of linear regressions can be interpreted as a particular instance of subspace clustering, since each regressor / regressand pair $(X, Y) \in \mathbb{R}^{d+1}$ lies in the d -dimensional subspace determined by their model parameter vectors (θ^* and $-\theta^*$). When the covariates and errors are Gaussian, algebro-geometric and probabilistic interpretations of PCA [6], [7] motivate related clustering schemes, since there is an inherent geometric aspect to such mixture models.

Another competitor is the Expectation-Maximization (EM) algorithm, which has been shown to have desirable empirical performance in various simulation studies [3], [4], [8]. Introduced in a seminal paper of Dempster *et al.* [9], the EM algorithm is a widely used technique for parameter estimation, with common applications in latent variable models (e.g., mixture models) and incomplete-data problems (e.g., corrupted or missing data) [10]. It is an iterative procedure that monotonically increases the likelihood [9, Th. 1]. When the likelihood is not concave, it is well known that EM can converge to a non-global optimum [11, p. 97]. However, recent work has side-stepped the question of whether EM reaches

the likelihood maximizer, instead by directly working out statistical guarantees on its loss. For certain well-specified models, these explorations have identified regions of local contractivity of the EM operator near the true parameter so that, when initialized properly, the EM iterates approach the true parameter with high probability.

This line of research was spurred by [12], which established general conditions for which a ball centered at the true parameter would be a basin of attraction for the population version of the EM operator. For a large enough sample size, the difference (in that ball) between the sample EM operator and the population EM operator can be bounded such that the EM estimate approaches the true parameter with high probability. That bound is the sum of two terms with distinct interpretations. There is an *algorithmic convergence* term $\gamma^t \|\theta^0 - \theta^*\|$ for initial guess θ^0 , truth θ^* , and some modulus of contraction $\gamma \in (0, 1)$; this comes from the analysis of the population EM operator. The second term captures *statistical convergence* and is proportional to the supremum norm of $\sup_{\theta} \|M(\theta) - M_n(\theta)\|$, the difference between the population and sample EM operators, M and M_n , respectively. This result is also shown for a “sample-splitting” version of EM, where the sample is partitioned into batches and each batch governs a single step of the algorithm.

Our purpose here is to follow up on the analysis of [12] by proving a larger basin of attraction for the mixture of two linear models and by establishing an exact probabilistic bound on the error of the sample-splitting EM estimate when the initial guess falls in the specified region. In particular, we show that

- (a) The EM algorithm converges to the target parameter vector when it is initialized in a cone (defined in terms of the cosine similarity between the initial guess θ^0 and the target model parameter θ^*).
- (b) The EM algorithm can fail to converge to θ^* if the cosine similarity is too small.

In related works, typically some variant of the mean value theorem is employed to establish contractivity toward the true parameter and the rate of geometric decay is then determined by relying heavily on the fact that initial guess belongs to a bounded set and is not too far from the target parameter vector (i.e., a ball centered at the target parameter vector). Our technique relies on Stein’s Lemma, which allows us to reduce the problem to the two-dimensional case and exploit certain monotonicity properties of the population EM operator. Such methods allow one to be very careful and explicit in the analysis and more cleanly reveal the role of the initial conditions. These results cannot be deduced from preexisting works (such as [12]), even by sharpening their analysis. Our improvements are not solely in terms of constants. Indeed, we will show that as long as the cosine angle between the initial guess and the target parameter vector (i.e., their degree of alignment) is sufficiently large, the EM algorithm converges to the target parameter vector θ^* . In particular, the norm of the initial guess can be arbitrarily large, provided the cosine angle condition is met.

In the machine learning community, mixtures of linear regressions are known as Hierarchical Mixture of

Experts (HME) and, there, the EM algorithm has also been employed [13]. The mixtures of linear regressions problem has also drawn recent attention from other scholars (e.g., [14]–[20]), although none of them have attempted to sharpen the EM algorithm in the sense that many works still require initialization is a small ball around the target parameter vector. For example, the general case with multiple components was considered in [18], but initialization is still required to be in a ball around each of the true component coefficient vectors.

This paper is organized as follows. In Section II, we explain the model and explain how the population EM operator is contractive toward the true parameter on a cone in \mathbb{R}^d . We also show that the operator is not contractive toward the true parameter on certain regions of \mathbb{R}^d . We connect our problem to phase retrieval in Section III and borrow preexisting techniques to find a good initial guess in Section IV. Section V looks at the behavior of the sample-splitting EM operator in this cone and states our main result in the form of a high-probability bound. Section VI and Section VII are devoted to proving the contractivity of the population EM operator toward the target vector over a cone and proving our main result, respectively. A discussion of our findings, including evidence of the failure of the EM algorithm for poor initial guesses from a simulated experiment, is provided in Section VIII. A simulation study of the EM algorithm under model misspecification is also given therein. Finally, more technical proofs are relegated to Appendix.

II. THE EMPIRICAL AND POPULATION EM OPERATOR

The EM operator for estimating θ^* (see [12, p. 6] for a derivation) is

$$M_n(\theta) = \left(\frac{1}{n} \sum_{i=1}^n X_i X_i^\top \right)^{-1} \times \left[\frac{1}{n} \sum_{i=1}^n (2\phi(Y_i \langle \theta, X_i \rangle / \sigma^2) - 1) X_i Y_i \right], \quad (2)$$

where $\phi(z) = \frac{1}{1+e^{-2z}}$ is a horizontally stretched logistic sigmoid. Here $\left(\frac{1}{n} \sum_{i=1}^n X_i X_i^\top \right)^{-1}$ is the inverse of the Gram matrix $\frac{1}{n} \sum_{i=1}^n X_i X_i^\top$. In the limit with infinite data, the population EM operator replaces sample averages with expectations, and thus

$$M(\theta) = 2\mathbb{E} \left[\phi(Y \langle \theta, X \rangle / \sigma^2) XY \right]. \quad (3)$$

As we mentioned in the introduction, [12] showed that if the EM operator (2) is initialized in a ball around θ^* with radius proportional θ^* , then the EM algorithm converges to θ^* with high probability. It is natural to ask whether this good region of initialization can be expanded, possibly allowing for initial guesses with unbounded norm. The purpose of this paper is to relax the aforementioned ball condition of [12] and show that if the cosine angle between θ^* and the initial guess is not too small, the EM algorithm also converges. We also simplify the analysis considerably and use only




Fig. 1. The population EM operator $M(\theta)$ lies in the space spanned by θ and θ^* . The unit vector θ_0^\perp lies in the space spanned by θ and θ^* and is perpendicular to θ . The vector θ forms an angle α with θ^* .

elementary facts about multivariate Gaussian distributions. Our improvement is manifested in the set containment

$$\{\theta : \|\theta - \theta^*\| \leq \sqrt{1 - \rho^2} \|\theta^*\|\} \subseteq \{\theta : \langle \theta, \theta^* \rangle \geq \rho \|\theta\| \|\theta^*\|\},$$

for $\rho \in [-1, 1]$, since for all θ in the set on the left side,

$$\begin{aligned} \langle \theta, \theta^* \rangle &= \frac{1}{2} (\|\theta\|^2 + \|\theta^*\|^2 - \|\theta - \theta^*\|^2) \\ &\geq \frac{1}{2} (\|\theta\|^2 + \rho^2 \|\theta^*\|^2) \\ &= \rho \|\theta\| \|\theta^*\| + \frac{1}{2} (\rho \|\theta^*\| - \|\theta\|)^2 \\ &\geq \rho \|\theta\| \|\theta^*\|. \end{aligned} \quad (4)$$

The conditions in [12, Corollary 5] require the initial guess θ^0 to be at most $\|\theta^*\|/32$ away from θ^* , which corresponds to $\|\theta\| \leq (1 + \sqrt{1 - \rho^2}) \|\theta^*\|$ and $\rho = \sqrt{1 - (1/32)^2} \approx 0.999$, whereas our condition allows for the norm of θ to be unbounded and $\rho > 0.85$.

Let θ_0 be the unit vector in the direction of θ and let θ_0^\perp be the unit vector that belongs to the hyperplane spanned by $\{\theta^*, \theta\}$ and orthogonal to θ (i.e., $\theta_0^\perp \in \text{span}\{\theta, \theta^*\}$ and $\langle \theta, \theta_0^\perp \rangle = 0$). Let $\theta^\perp = \|\theta\| \theta_0^\perp$. We will later show in Section VI that $M(\theta)$ belongs to $\text{span}\{\theta, \theta^*\}$, as illustrated in Fig. 1. Denote the angle between θ^* and θ_0 as α , with $\|\theta^*\| \cos \alpha = \langle \theta_0, \theta^* \rangle$ and $\rho = \cos \alpha$. As we will see from the following results, as long as $\cos \alpha$ is not too small, $M(\theta)$ is a contracting operation that is always closer to the truth θ^* than θ . The next lemma allows us to derive a region of \mathbb{R}^d on which M is contractive toward θ^* . We defer its proof until Section VI.

Lemma 1: For any θ in \mathbb{R}^d with $\langle \theta, \theta^ \rangle > 0$,*

$$\|M(\theta) - \theta^*\| \leq \gamma \|\theta - \theta^*\|, \quad (5)$$

where

$$\gamma = \sqrt{\kappa} \sqrt{1 + 4 \left(\frac{|\langle \theta^\perp, \theta^* \rangle| + \sigma^2}{\langle \theta, \theta^* \rangle} \right)^2}, \quad (6)$$

and

$$\kappa^2 = \max \left\{ 1 - \frac{|\langle \theta_0, \theta^* \rangle|^2}{\sigma^2 + \|\theta^*\|^2}, 1 - \frac{\langle \theta, \theta^* \rangle}{\sigma^2 + \langle \theta, \theta^* \rangle} \right\} < 1. \quad (7)$$

If we define the input signal-to-noise ratio as $\eta' = \|\theta\|/\sigma$ and model signal-to-noise ratio (SNR) as $\eta = \|\theta^*\|/\sigma$ and use the fact that $\|\theta^*\| \cos \alpha = \langle \theta_0, \theta^* \rangle$, then the contractivity constant (6) can be rewritten as

$$\max \left\{ \left(1 - \frac{\eta^2 \cos^2 \alpha}{1 + \eta^2} \right)^{1/4}, \left(1 - \frac{\eta' \eta \cos \alpha}{1 + \eta' \eta \cos \alpha} \right)^{1/4} \right\} \times \sqrt{1 + 4 \left(\tan \alpha + \frac{1}{\eta' \eta \cos \alpha} \right)^2}. \quad (8)$$

Remark 1: If $\eta' \geq 20$, $\eta \geq 40$, and $\cos \alpha \geq 0.85$, then κ is bounded by a universal constant less than $1/2$ and γ is bounded by a universal constant less than 1 , implying the population EM operator $\theta^t \leftarrow M(\theta^{t-1})$ converges to the truth θ^ exponentially fast.*

III. RELATIONSHIP TO PHASE RETRIEVAL

The problem of estimating the true parameter vector in a mixture of two linear regressions is related to the phase retrieval problem, where one has access to magnitude-only data according to the model

$$\tilde{Y} = |\langle \theta^*, X \rangle|^2 + \varepsilon. \quad (9)$$

In the no noise case, i.e., $\varepsilon \equiv 0$, one can obtain the phase retrieval model from the symmetric two component mixture of linear regressions by squaring each response variable Y_i from (1) and visa versa by setting $Y_i = R_i \sqrt{Y'_i}$,

where $R_i \stackrel{i.i.d.}{\sim}$ Rademacher(1/2) is independent of the data $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$. Here the sample subsets giving rise to the model parameters θ^* and $-\theta^*$ are $\{i : R_i \text{sgn}(\langle \theta^*, X_i \rangle) = 1\}$ and $\{i : R_i \text{sgn}(\langle \theta^*, X_i \rangle) = -1\}$, respectively. Even in the case of noise, squaring each response variable and subtracting the variance σ^2 of the error distribution yields

$$\begin{aligned} Y'_i &= Y_i^2 - \sigma^2 = |\langle \theta^*, X_i \rangle|^2 + 2R_i \varepsilon_i \langle \theta^*, X_i \rangle + (\varepsilon_i^2 - \sigma^2) \\ &= |\langle \theta^*, X_i \rangle|^2 + \zeta(X_i, R_i, \varepsilon_i), \end{aligned} \quad (10)$$

where $\zeta(X_i, R_i, \varepsilon_i)$ is a mean zero random variable with variance $4\sigma^2 \|\theta^*\|^2 + 2\sigma^4$. This is essentially the phase retrieval model (9) with heteroskedastic errors. See also [15, Sec. 3.5] for a similar reduction to the ‘‘Noisy Phase Model’’, where the measurement error is pre-added to the inner product and then squared, viz., $|\langle \theta^*, X \rangle + \varepsilon|^2$.

Recent algorithms used to recover θ^* from (X, \tilde{Y}) include PhaseLift [21], PhaseMax [22], [23], PhaseLamp [24], [25] and Wirtinger flow [26], [27], to name a few. PhaseLift operates by solving a semi-definite relaxation of the nonconvex formulation of the phase retrieval problem. PhaseMax and PhaseLamp solve a linear program over a polytope via convex programming. Finally, Wirtinger flow is an iterative gradient-based method that requires proper initialization. Parallel to our work, [24], [25] reveal that exact recovery (when $n, d \rightarrow +\infty$) in PhaseMax is governed by a critical threshold [24, Th. 3], which is measured in terms of the cosine angle between the initial guess and the target parameter vector. Analogous to our Lemma 2 (which is asymptotic in the sense that $n \rightarrow +\infty$), they prove that recovery can fail if this cosine angle

is too small. PhaseLamp is an iterative variant of PhaseMax that allows for a smaller cosine angle criterion than the critical threshold from PhaseMax. Our setting is slightly more general than [24], [25] in that we allow for measurement error and our bounds are non-asymptotic in n and d .

IV. INITIALIZATION

Theorem 1 below requires the initial guess to have a good inner product with θ^* . But how should one initialize in practice? There is considerable literature showing the efficacy of initialization based on spectral [16]–[18] or Bayesian [4] methods. For example, inspired by the link (10) between phase retrieval and our problem, we can use the same spectral initialization method of [27] for the Wirtinger flow iterates (c.f., [16] for a similar strategy). That is, set

$$\lambda^2 = d \frac{\sum_{i=1}^n Y_i'}{\sum_{i=1}^n \|X_i\|^2}, \quad (11)$$

and take θ^0 equal to be the eigenvector corresponding to the largest eigenvalue of

$$\frac{1}{n} \sum_{i=1}^n Y_i' X_i X_i^\top, \quad (12)$$

scaled so that $\|\theta^0\| = \lambda$. According to [27, Th. 3.3], we are guaranteed that with high probability $\|\theta^0 - \theta^*\| \leq \frac{1}{8}\|\theta^*\|$, and hence by (4), $\langle \theta^0, \theta^* \rangle \geq \sqrt{1 - (1/8)^2} \|\theta^0\| \|\theta^*\| \approx 0.992 \|\theta^0\| \|\theta^*\|$ and $\|\theta^0\| \geq (7/8)\|\theta^*\|$. Provided that $\|\theta^*\| \geq (8/7)20\sigma$, we will see in Theorem 1 that this θ^0 satisfies our criteria for a good initial guess. Although the joint distributions of (X, \tilde{Y}) and (X, Y') are not exactly the same, for large n , $\frac{1}{n} \sum_{i=1}^n \zeta(X_i, R_i, \varepsilon_i) \approx 0$, and hence (11) and (12) are approximately equal to the same quantity with Y_i' replaced by \tilde{Y}_i .

The next lemma, proved in Appendix, shows that the initialization conditions in Remark 1 are essentially necessary in the sense that contractivity of M toward θ^* can fail for certain initial guesses that do not meet our cosine angle criterion. In contrast, it is known [28], [29] that the population EM operator for a symmetric mixture of two Gaussians $Y \sim \frac{1}{2}N(\theta^*, \sigma^2 I_d) + \frac{1}{2}N(-\theta^*, \sigma^2 I_d)$ is contractive toward θ^* on the entire half-plane defined by $\langle \theta, \theta^* \rangle > 0$.¹ The disparity between the EM operators for the two models is revealed in the proof of the contractivity of M toward θ^* (see Section VI). Indeed, we will see in Remark 2 that the population EM operator for mixtures of regressions is essentially a “stretched” version of the population EM operator for Gaussian mixtures.

Lemma 2: There is a subset of \mathbb{R}^d with positive Lebesgue measure, each of whose members θ satisfies $\langle \theta, \theta^ \rangle > 0$ and*

$$\|M(\theta) - \theta^*\| > \|\theta - \theta^*\|.$$

While this result does not generally imply that the empirical iterates $\theta^t \leftarrow M_n(\theta^{t-1})$ will fail to converge to θ^* for

¹Note that this is the best one can hope for: if $\langle \theta, \theta^* \rangle < 0$ (reps. $\langle \theta, \theta^* \rangle = 0$), then the population EM operator is contractive toward $-\theta^*$ (resp. the zero vector). Thus, unless $\langle \theta, \theta^* \rangle = 0$ (i.e., θ belongs to the hyperplane perpendicular to θ^*), the population EM is contractive towards either model parameter $-\theta^*$ or θ^* .

$\langle \theta^0, \theta^* \rangle > 0$, it does suggest that difficulties may arise in this regime. Indeed, the discussion in Section VIII gives empirical evidence for this theoretical observation.

V. MAIN THEOREM

As in [12], we analyze a sample-splitting version of the EM algorithm, where for an allocation of n samples and T iterations, we divide the data into T subsets of size $\lfloor n/T \rfloor$. We then perform the updates $\theta^t \leftarrow M_{n/T}(\theta^{t-1})$, using a new subset of samples to compute $M_{n/T}(\theta)$ at each iteration. The advantage of sample-splitting is purely for ease of analysis. In particular, conditional on the portion of data used to construct $M_{n/T}$ at iteration t , the distribution of θ^t depends only on the other portion of the data through θ^{t-1} . For the next theorem, let $\eta^0 = \|\theta^0\|/\sigma$ denote the initial SNR and $\eta = \|\theta^*\|/\sigma$ denote the model SNR.

Theorem 1: Let $\langle \theta^0, \theta^ \rangle > \rho \|\theta^0\| \|\theta^*\|$ for $\rho > 0.85$, $\eta^0 \geq 20$, and $\eta \geq 40$. Fix $\delta \in (0, 1)$. Suppose furthermore that $n \geq \max\{cd \log(T/\delta), c'\}$ for some positive universal constant c and positive constant $c' = c'(\rho, \sigma, \|\theta^*\|, \|\theta^0\|)$. Then there exists a universal modulus of contraction $\gamma \in (0, 1)$ and a positive universal constant C such that the sample-splitting empirical EM iterates $(\theta^t)_{t=1}^T$ based on n/T samples per step satisfy*

$$\|\theta^t - \theta^*\| \leq \gamma^t \|\theta^0 - \theta^*\| + \frac{C\sqrt{\sigma^2 + \|\theta^*\|^2}}{1 - \gamma} \sqrt{\frac{dT \log(T/\delta)}{n}},$$

with probability at least $1 - \delta$.

Note that T governs the number of iterations of the EM operator; if it is too small, the term $\gamma^t \|\theta^0 - \theta^*\|$ from Theorem 1 may fail to reach the parametric rate. Hence, T must scale like $\frac{\log(n/d)}{\log(1/\gamma)}$.

We will prove Theorem 1 in Section VII. The main aspect of the analysis lies in showing that M_n satisfies an invariance property, i.e., $M_n(\mathcal{A}) \subseteq \mathcal{A}$, where \mathcal{A} is a set on which M is contractive toward θ^* . The algorithmic error $\gamma^t \|\theta^0 - \theta^*\|$ is a result of repeated evaluation of the population EM operator $\theta^t \leftarrow M(\theta^{t-1})$ and the contractivity of M towards θ^* from Lemma 1. The stochastic error $\frac{C\sqrt{\sigma^2 + \|\theta^*\|^2}}{1 - \gamma} \sqrt{\frac{dT \log(T/\delta)}{n}}$ is obtained from a high-probability bound on $\max_{t \in [T]} \|M_{n/T}(\theta^t) - M(\theta^t)\|$, which is contained in the proof of [12, Corollary 5]).

VI. PROOF OF LEMMA 1

If $W = \langle \theta^*, X \rangle + \varepsilon$, a few applications of Stein’s Lemma [30, Lemma 1] yields

$$\begin{aligned} M(\theta) &= \mathbb{E} \left[(2\phi(W\langle \theta, X \rangle/\sigma^2) - 1) X W \right] \\ &= \theta^* \left[2\phi(W\langle \theta, X \rangle/\sigma^2) + 2(W\langle \theta, X \rangle/\sigma^2) \right. \\ &\quad \left. \times \phi'(W\langle \theta, X \rangle/\sigma^2) - 1 \right] \\ &\quad + \theta \mathbb{E} \left[2(W^2/\sigma^2) \phi'(W\langle \theta, X \rangle/\sigma^2) \right]. \end{aligned} \quad (13)$$

Letting

$$\begin{aligned} A &= \left[2\phi(W\langle \theta, X \rangle/\sigma^2) + 2(W\langle \theta, X \rangle/\sigma^2) \right. \\ &\quad \left. \times \phi'(W\langle \theta, X \rangle/\sigma^2) - 1 \right], \end{aligned} \quad (14)$$

and

$$B = \mathbb{E} \left[2(W^2/\sigma^2)\phi'(W\langle\theta, X\rangle/\sigma^2) \right], \quad (15)$$

we see that $M(\theta) = \theta^*A + \theta B$ belongs to $\text{span}\{\theta, \theta^*\} = \{\lambda_1\theta + \lambda_2\theta^*, : \lambda_1, \lambda_2 \in \mathbb{R}\}$. This is a crucial fact that we will exploit multiple times.

Observe that for any a in $\text{span}\{\theta, \theta^*\}$,

$$a = \langle\theta_0, a\rangle\theta_0 + \langle\theta_0^\perp, a\rangle\theta_0^\perp,$$

and

$$\|a\|^2 = |\langle\theta_0, a\rangle|^2 + |\langle\theta_0^\perp, a\rangle|^2.$$

Specializing this to $a = M(\theta) - \theta^*$ yields

$$\|M(\theta) - \theta^*\|^2 = |\langle\theta_0, M(\theta) - \theta^*\rangle|^2 + |\langle\theta_0^\perp, M(\theta) - \theta^*\rangle|^2.$$

The strategy for establishing contractivity of $M(\theta)$ toward θ^* will be to show that the sum of $|\langle\theta_0, M(\theta) - \theta^*\rangle|^2$ and $|\langle\theta_0^\perp, M(\theta) - \theta^*\rangle|^2$ is less than $\gamma^2\|\theta - \theta^*\|^2$. This idea was used in [28] to obtain contractivity of the population EM operator for a mixture of two Gaussians. Due to the similarity of the two problems, it turns out that many of the same ideas transfer to our (more complicated) setting.

To reduce this $(d+1)$ -dimensional problem $(X, Y) \in \mathbb{R}^{d+1}$ to a 2-dimensional problem $(Z_1, Z_2) \in \mathbb{R}^2$, we first show that

$$W\langle\theta, X\rangle/\sigma^2 \stackrel{\mathcal{D}}{=} \Lambda Z_1|Z_2| + \Gamma Z_2^2,$$

where $Z_1, Z_2 \stackrel{i.i.d.}{\sim} N(0, 1)$. The coefficients Γ and Λ are

$$\Gamma = \langle\theta, \theta^*\rangle/\sigma^2$$

and

$$\begin{aligned} \Lambda^2 &= (\|\theta\|^2/\sigma^4)(\sigma^2 + \|\theta^*\|^2) - \Gamma^2 \\ &= (\|\theta\|^2/\sigma^4)(\sigma^2 + |\langle\theta_0^\perp, \theta^*\rangle|^2). \end{aligned}$$

This is because of the distributional equality

$$(W, \langle\theta, X\rangle/\sigma^2) \stackrel{\mathcal{D}}{=} \left(\sqrt{\sigma^2 + \|\theta^*\|^2}Z_2, \frac{\Lambda}{\sqrt{\sigma^2 + \|\theta^*\|^2}}Z_1 + \frac{\Gamma}{\sqrt{\sigma^2 + \|\theta^*\|^2}}Z_2 \right). \quad (16)$$

Note further that $\Lambda Z_1 Z_2 + \Gamma Z_2^2 \stackrel{\mathcal{D}}{=} \Lambda Z_1|Z_2| + \Gamma Z_2^2$ because they have the same moment generating function. Using this, we deduce that

$$W\langle\theta, X\rangle/\sigma^2 \stackrel{\mathcal{D}}{=} \Lambda Z_1|Z_2| + \Gamma Z_2^2. \quad (17)$$

Lemma 7 implies that

$$(1 - \kappa)\langle\theta_0^\perp, \theta^*\rangle \leq \langle\theta_0^\perp, M(\theta)\rangle \leq (1 + \sqrt{\kappa})\langle\theta_0^\perp, \theta^*\rangle,$$

and consequently,

$$|\langle\theta_0^\perp, M(\theta) - \theta^*\rangle| \leq \sqrt{\kappa}|\langle\theta_0^\perp, \theta - \theta^*\rangle| \leq \sqrt{\kappa}\|\theta - \theta^*\|. \quad (18)$$

Next, we note that

$$\begin{aligned} \sigma^4|\Lambda^2 - \Gamma| &= \|\theta\|^2(\sigma^2 + |\langle\theta_0^\perp, \theta^*\rangle|^2) - \sigma^2\langle\theta, \theta^*\rangle \\ &\leq \|\theta\|^2|\langle\theta_0^\perp, \theta^*\rangle|^2 + \sigma^2|\langle\theta, \theta - \theta^*\rangle| \\ &\leq \|\theta\|(|\langle\theta_0^\perp, \theta^*\rangle| + \sigma^2)\|\theta - \theta^*\|. \end{aligned}$$

Finally, define

$$h(\alpha, \beta) = \mathbb{E}[(2\phi(\alpha|Z_2|(Z_1 + \beta|Z_2|)) - 1)(|Z_2|(Z_1 + \beta|Z_2|))].$$

Note that by definition of h , $h(\Lambda, \frac{\Gamma}{\Lambda}) = \frac{\langle\theta, M(\theta)\rangle}{\Lambda}$. In fact, h is the one-dimensional population EM operator for this model when $\theta^* = \beta$ and $\sigma^2 = 1$. By the self-consistency property of EM [31, p. 79], $h(\beta, \beta) = \beta$. Translating this to our problem, we have that $h(\frac{\Gamma}{\Lambda}, \frac{\Gamma}{\Lambda}) = \frac{\Gamma}{\Lambda} = \frac{\langle\theta, \theta^*\rangle}{\sigma^2\Lambda}$. Since $h(\Lambda, \frac{\Gamma}{\Lambda}) - h(\frac{\Gamma}{\Lambda}, \frac{\Gamma}{\Lambda}) = \int_{\frac{\Gamma}{\Lambda}}^{\Lambda} \frac{\partial h}{\partial \alpha} h(\alpha, \frac{\Gamma}{\Lambda}) d\alpha$, we have from Lemma 8,

$$\begin{aligned} |\langle\theta_0, M(\theta) - \theta^*\rangle| &\leq \frac{\sigma^2\Lambda}{\|\theta\|} \left| \int_{\frac{\Gamma}{\Lambda}}^{\Lambda} \frac{\partial}{\partial \alpha} h\left(\alpha, \frac{\Gamma}{\Lambda}\right) d\alpha \right| \\ &\leq \frac{2\sqrt{\kappa}\sigma^2\Lambda}{\|\theta\|} \left| \int_{\frac{\Gamma}{\Lambda}}^{\Lambda} \frac{d\alpha}{\alpha^2} \right| \\ &= \frac{2\sigma^2\sqrt{\kappa}|\Lambda^2 - \Gamma|}{\Gamma\|\theta\|} \\ &\leq 2\sqrt{\kappa} \left(\frac{|\langle\theta_0^\perp, \theta^*\rangle| + \sigma^2}{\langle\theta, \theta^*\rangle} \right) \|\theta - \theta^*\|. \end{aligned}$$

Combining this with inequality (18) yields (5). This completes the proof of Lemma 1.

Remark 2: The function h is related to the EM operator for the one-dimensional symmetric mixture of two Gaussians model $Y \sim \frac{1}{2}N(-\beta, 1) + \frac{1}{2}N(\beta, 1)$. One can derive that (see [32, p. 11]) the population EM operator is

$$g(\alpha, \beta) = \mathbb{E}[(2\phi(\alpha(Z_1 + \beta)) - 1)(Z_1 + \beta)].$$

Then $h(\alpha, \beta)$ is a “stretched” version of $g(\alpha, \beta)$ as seen through the identity

$$h(\alpha, \beta) = \mathbb{E}[|Z_2|g(\alpha|Z_2|, \beta|Z_2|)].$$

In light of this relationship, it is perhaps not surprising that the EM operator for the mixture of linear regressions problem also enjoys a large basin of attraction.

On the other hand, from [32, p. 11], the population EM operator \tilde{M} for the symmetric two component mixture of Gaussians $Y \sim \frac{1}{2}N(\theta^*, \sigma^2 I_d) + \frac{1}{2}N(-\theta^*, \sigma^2 I_d)$, is equal to

$$\tilde{M}(\theta) = 2\mathbb{E}[Y\phi((Y, \theta)/\sigma^2)] = \theta^*\tilde{A} + \theta\tilde{B},$$

where $\tilde{A} = \mathbb{E}[2\phi(\langle\theta, \theta^*\rangle/\sigma^2 + \|\theta\|Z_1/\sigma) - 1]$ and $\tilde{B} = 2\mathbb{E}[\phi'(\langle\theta, \theta^*\rangle/\sigma^2 + \|\theta\|Z_1/\sigma)]$.

Compare the values of \tilde{A} and \tilde{B} with A and B from (14) and (15). We see that \tilde{M} is essentially a “stretched” and “scaled” version of \tilde{M} by the random dilation factors $|Z_2|\sqrt{1 + |\langle\theta_0^\perp, \theta^*\rangle|^2/\sigma^2}$ and $|Z_2|\sqrt{1 + \|\theta^*\|^2/\sigma^2}$. As will be seen in the proof Lemma 2 in Appendix, this additional source of variability causes the repellant behavior of \tilde{M} in Lemma 2.

Remark 3: Recently in [33], the authors analyzed gradient descent for a single-hidden layer convolutional neural network structure with no overlap and Gaussian input. In this setup, we observe i.i.d. data $\{(X_i, Y_i)\}_{i=1}^n$, where $Y_i = f(X_i, w) + \varepsilon_i$ and $X_i \sim N(0, I_d)$ and $\varepsilon_i \sim N(0, \sigma^2)$ are independent of each other. The neural network has the form $f(x, w) = \frac{1}{\kappa} \sum_{j=1}^k \max\{0, \langle w_j, x \rangle\}$ and the only nonzero coordinates of

w_j are in the j^{th} successive block of d/k coordinates and are equal to a fixed d/k dimensional filter vector w . One desires to minimize the risk $\ell(w) = \mathbb{E}[(f(X, w) - f(X, w^*))^2]$. Interestingly, the gradient of $\ell(w)$ belongs to the linear span of w and w^* , akin to our $M(\theta) \in \text{span}\{\theta, \theta^*\}$ (and also in the Gaussian mixture problem [32]). This property also plays a critical role in their analysis.

VII. PROOF OF THEOREM 1

The first step of the proof is to show that the empirical EM operator satisfies $M_n(\mathcal{A}) \subset \mathcal{A}$, where \mathcal{A} is a set on which M is contractive toward θ^* . In other words, the empirical EM iterates remain in a set where $M(\theta)$ is closer to θ^* than its input θ . To this end, define the set $\mathcal{A} = \{\theta : \langle \theta, \theta^* \rangle > \rho \|\theta\| \|\theta^*\|, \|\theta\| \geq 20\sigma\}$. By Remark 1, the stated conditions on ρ , $\|\theta\|$, and $\|\theta^*\|$ ensure that M is contractive toward θ^* on \mathcal{A} and that $\kappa < 1/2$.

Next, we use Lemma 3 which shows that

$$M(\mathcal{A}) \subseteq \mathcal{B} \triangleq \{\theta : \langle \theta, \theta^* \rangle > (1 + \Delta)\rho \|\theta\| \|\theta^*\|, \|\theta^*\|(1 - \kappa) \leq \|\theta\| \leq \sqrt{\sigma^2 + 3\|\theta^*\|^2}\}.$$

The fact that $\mathcal{B} \subset \mathcal{A}$ allows us to claim that when n is large enough, $M_n(\mathcal{A}) \subset M(\mathcal{B})$, and hence $M_n(\mathcal{A}) \subseteq M(\mathcal{A}) \subseteq \mathcal{A}$. To show this, assume $\sup_{\theta \in \mathcal{A}} \|M_n(\theta) - M(\theta)\| < \epsilon$. That implies

$$\begin{aligned} \sup_{\theta \in \mathcal{A}} \left\| \frac{M_n(\theta)}{\|M_n(\theta)\|} - \frac{M(\theta)}{\|M(\theta)\|} \right\| &\leq 2 \sup_{\theta \in \mathcal{A}} \frac{\|M_n(\theta) - M(\theta)\|}{\|M(\theta)\|} \\ &< \frac{2\epsilon}{(1 - \kappa)\|\theta^*\|}. \end{aligned} \quad (19)$$

For the last inequality, we used the fact that $\|M(\theta)\| \geq \|\theta^*\|A \geq \|\theta^*\|(1 - \kappa)$ for all θ in \mathcal{A} , which follows from (13) and Lemma 7. By (19) and Lemma 3 (A.3), we have that

$$\begin{aligned} \sup_{\theta \in \mathcal{A}} \left\langle \theta^*, \frac{M_n(\theta)}{\|M_n(\theta)\|} \right\rangle &\geq \sup_{\theta \in \mathcal{A}} \left\langle \theta^*, \frac{M(\theta)}{\|M(\theta)\|} \right\rangle - \frac{2\epsilon}{(1 - \kappa)} \\ &\geq \|\theta^*\|(1 + \Delta)\rho - \frac{2\epsilon}{(1 - \kappa)} \\ &\geq \|\theta^*\|\rho, \end{aligned}$$

provided $\epsilon < (\frac{1-\kappa}{2})\Delta\rho\|\theta^*\|$ and, by (13) and Lemma 7,

$$\begin{aligned} \sup_{\theta \in \mathcal{A}} \|M_n(\theta)\| &\geq \sup_{\theta \in \mathcal{A}} \|M(\theta)\| - \epsilon \\ &\geq \|\theta^*\|(1 - \kappa) - \epsilon \\ &\geq 40\sigma(1 - \kappa) - \epsilon \\ &\geq 20\sigma, \end{aligned}$$

provided $\epsilon < 20\sigma(1 - 2\kappa)$, which is positive since $\kappa < 1/2$.

For $\delta \in (0, 1)$, let $\epsilon_M(n, \delta)$ be the smallest number such that for any fixed θ in \mathcal{A} , we have

$$\|M_n(\theta) - M(\theta)\| \leq \epsilon_M(n, \delta),$$

with probability at least $1 - \delta$. Moreover, suppose $c' = c'(\rho, \sigma, \|\theta^*\|, \|\theta^0\|)$ is a constant so that if $n \geq c'$, then

$$\epsilon_M(n, \delta) \leq \min \left\{ 20\sigma(1 - 2\kappa), \left(\frac{1 - \kappa}{2} \right) \Delta\rho\|\theta^*\| \right\}.$$

This guarantees that $M_n(\mathcal{A}) \subseteq \mathcal{A}$. For any iteration $t \in [T]$, we have

$$\|M_{n/T}(\theta^t) - M(\theta^t)\| \leq \epsilon_M(n/T, \delta/T),$$

with probability at least $1 - \delta/T$. Thus by a union bound and $M_n(\mathcal{A}) \subseteq \mathcal{A}$,

$$\max_{t \in [T]} \|M_{n/T}(\theta^t) - M(\theta^t)\| \leq \epsilon_M(n/T, \delta/T),$$

with probability at least $1 - \delta$.

Hence if θ^0 belongs to \mathcal{A} , then by Lemma 1,

$$\begin{aligned} \|\theta^t - \theta^*\| &= \|M_{n/T}(\theta^{t-1}) - \theta^*\| \\ &\leq \|M(\theta^{t-1}) - \theta^*\| + \|M_{n/T}(\theta^t) - M(\theta^t)\| \\ &\leq \gamma \|\theta^{t-1} - \theta^*\| + \max_{t \in [T]} \|M_{n/T}(\theta) - M(\theta)\| \\ &\leq \gamma \|\theta^{t-1} - \theta^*\| + \epsilon_M(n/T, \delta/T), \end{aligned}$$

with probability at least $1 - \delta$. Solving this recursive inequality yields,

$$\begin{aligned} \|\theta^t - \theta^*\| &\leq \gamma^t \|\theta^0 - \theta^*\| + \epsilon_M(n/T, \delta/T) \sum_{j=0}^{t-1} \gamma^j \\ &\leq \gamma^t \|\theta^0 - \theta^*\| + \frac{\epsilon_M(n/T, \delta/T)}{1 - \gamma}, \end{aligned}$$

with probability at least $1 - \delta$.

Finally, by a slight modification to the proof of [12, Corollary 5] that uses $M(\theta) \leq \sqrt{\sigma^2 + 3\|\theta^*\|^2}$ from (A.4), it follows that if $n \geq cd \log(T/\delta)$, then there exists a universal constant $C > 0$ such that

$$\epsilon_M(n/T, \delta/T) \leq C \sqrt{\sigma^2 + \|\theta^*\|^2} \sqrt{\frac{dT \log(T/\delta)}{n}}$$

with probability at least $1 - \delta/T$. This completes the proof of Theorem 1.

VIII. DISCUSSION

In this paper, we showed that the empirical EM iterates converge to true coefficients of a mixture of two linear regressions as long as the initial guess lies within a cone (see the condition on Theorem 1: $\langle \theta^0, \theta^* \rangle > \rho \|\theta^0\| \|\theta^*\|$).

In Fig. 2, we perform a simulation study of $\theta^t \leftarrow M_n(\theta^{t-1})$ with $\sigma = 1$, $n = 1000$, $d = 2$, and $\theta^* = (-7/25, 24/25)^\top$. All entries of the covariate vector X and the noise ε are generated *i.i.d.* from a standard Gaussian distribution. We consider the error $\|\theta^t - \theta^*\|$ plotted as a function of $\cos \alpha = \frac{\langle \theta^0, \theta^* \rangle}{\|\theta^0\| \|\theta^*\|}$ at iterations $t = 5, 10, 15, 20, 25$ (darker lines correspond to larger values of t). For each t , we choose a unit vector θ^0 so that $\cos \alpha$ ranges between -1 and $+1$. In accordance with the theory we have developed, increasing the iteration size and increasing the cosine angle decreases the overall error. According to Lemma 2, the algorithm should suffer when $\cos \alpha$ is small. Indeed, we observe a sharp transition at $\cos \alpha \approx 0.2$. The algorithm converges to the other model parameter $-\theta^* = (7/25, -24/25)^\top$ for initial guesses with cosine angle (approximately) smaller than 0.2. The plot in Fig. 3 is a zoomed-in version of Fig. 2 near this transition point.




Fig. 2. A simulation study of $\theta^t \leftarrow M_n(\theta^{t-1})$ with $\sigma^2 = 1$, $n = 1000$, $d = 2$, and $\theta^* = (-7/25, 24/25)^\top$. The values of t range from 5 to 25. The vertical axis is the error $\|\theta^t - \theta^*\|$ and the horizontal axis is the cosine angle between the initial guess θ^0 and θ^* . Darker lines correspond to larger values of t .




Fig. 3. A zoomed-in version of Fig. 2 showing the transition point at $\cos \alpha \approx 0.2$.

One of the shortcomings of the EM algorithm is that it is *model dependent*, that is, the form of the EM operator is derived from the assumption of Gaussian input X , error ε , and two component assumption. It is natural to ask how changing either distribution and using the original EM operator designed for Gaussian data performs on simulated data. As a simple illustration, the simulation results in Fig. 4 use $X \sim \text{Uniform}([- \sqrt{3}, \sqrt{3}]^d)$ and $\varepsilon \sim \text{Uniform}([- \sigma \sqrt{3}, \sigma \sqrt{3}]^d)$ (Fig. 4(a)) and $X \sim N(0, I_d)$ and $\varepsilon \sim \text{Laplace}(0, \sigma/\sqrt{2})$ (Fig. 4(b)) for $\sigma^2 = 1$. The performance is similar to Fig. 2 and Fig. 3, although note that in Fig. 4(a), a larger cosine angle is required for convergence (i.e., cosine angles at least $\cos \alpha \approx 0.4$).

More generally, future work would rigorously study the effect of EM under model misspecification. In this direction, the recent work of [34] has analyzed the EM algorithm for over-fitted mixtures.




Fig. 4. A simulation study of $\theta^t \leftarrow M_n(\theta^{t-1})$ under model misspecification with $\sigma^2 = 1$, $n = 1000$, $d = 2$, $\theta^* = (-7/25, 24/25)^\top$. The values of t range from 5 to 25. The vertical axis is the error $\|\theta^t - \theta^*\|$ and the horizontal axis is the cosine angle between the initial guess θ^0 and θ^* . Darker lines correspond to larger values of t . (a) $X \sim \text{Uniform}([- \sqrt{3}, \sqrt{3}]^d)$ and $\varepsilon \sim \text{Uniform}([- \sqrt{3}, \sqrt{3}]^d)$. (b) $X \sim N(0, I_d)$ and $\varepsilon \sim \text{Laplace}(0, \sigma/\sqrt{2})$.

APPENDIX

In this appendix, we prove Lemma 2 and all other supporting lemmas used in the body of the paper.

Proof of Lemma 2: Recall that in general, $M(\theta) = \theta^* A + \theta B$, where

$$A = [2\phi(W\langle \theta, X \rangle / \sigma^2) + 2(W\langle \theta, X \rangle / \sigma^2) \times \phi'(W\langle \theta, X \rangle / \sigma^2) - 1],$$

and

$$B = 2\mathbb{E}[(W^2 / \sigma^2) \phi'(W\langle \theta, X \rangle / \sigma^2)].$$

Suppose $\langle \theta, \theta^* \rangle = 0$. This implies that $A = 0$. To see this, note that

$$\mathbb{E}[\phi(W\langle \theta, X \rangle / \sigma^2)] = \mathbb{E}[\phi(\Lambda Z_1 | Z_2)] = \phi(0) = 1/2, \quad (\text{A.1})$$

and

$$\mathbb{E}[W\langle \theta, X \rangle \phi'(W\langle \theta, X \rangle / \sigma^2)] = \sigma^2 \mathbb{E}[\Lambda Z_1 | Z_2 \phi'(\Lambda Z_1 | Z_2)] = 0. \quad (\text{A.2})$$

The first equality (A.1) follows from the fact that if $Z \sim N(0, 1)$, then $\mathbb{E}[\phi(zZ)] = 1/2$ for all z in \mathbb{R} . This fact is easily established by noting that the derivative with respect to z is zero everywhere. The expectation in (A.2) vanishes since we are averaging an odd function with respect to a symmetric distribution. Next, observe that $B = 2(1 + \|\theta^*\|^2 / \sigma^2) \mathbb{E}[Z_2^2 \phi'(Z_1 | Z_2) (\|\theta\| / \sigma^2) \sqrt{\sigma^2 + \|\theta^*\|^2}] \rightarrow 1 + \|\theta^*\|^2 / \sigma^2 > 1$ as $\theta \rightarrow 0$. By continuity, there exists $r > 0$ such that if $\|\theta\| = r$, then $B > 1$, and hence

$$\begin{aligned} \|M(\theta) - \theta^*\|^2 &= \|\theta - \theta^*\|^2 + (B^2 - 1)\|\theta\|^2 \\ &> \|\theta - \theta^*\|^2. \end{aligned}$$

This shows that

$$\liminf_{(\theta, \theta^*) \downarrow 0, \|\theta\|=r} [\|M(\theta) - \theta^*\|^2 - \|\theta - \theta^*\|^2] > 0.$$

By continuity, it follows that there exists $r' > 0$ such that if $0 < \langle \theta, \theta^* \rangle < r'$ then $\|M(\theta) - \theta^*\|^2 > \|\theta - \theta^*\|^2$. It is easy to see that the set of all points satisfying $0 < \langle \theta, \theta^* \rangle < r'$ and $0 < \|\theta\| < r$ has positive Lebesgue measure and satisfies the stated conditions in the lemma. \square

For the following lemmas, recall the definitions

$$\begin{aligned} A &= [2\phi(W\langle \theta, X \rangle/\sigma^2) + 2(W\langle \theta, X \rangle/\sigma^2) \\ &\quad \times \phi'(W\langle \theta, X \rangle/\sigma^2) - 1], \\ B &= 2\mathbb{E}[(W^2/\sigma^2)\phi'(W\langle \theta, X \rangle/\sigma^2)], \end{aligned}$$

and

$$\begin{aligned} \kappa^2 &= \frac{1}{\frac{\Gamma}{\Lambda} \min\{\Lambda, \frac{\Gamma}{\Lambda}\} + 1} \\ &= \max\left\{1 - \frac{|\langle \theta_0, \theta^* \rangle|^2}{\sigma^2 + \|\theta^*\|^2}, 1 - \frac{\langle \theta, \theta^* \rangle}{\sigma^2 + \langle \theta, \theta^* \rangle}\right\}. \end{aligned}$$

Lemma 3: The cosine angle between θ^ and $M(\theta)$ is equal to*

$$\frac{\|\theta^*\|^2 A + \langle \theta, \theta^* \rangle B}{\sqrt{(\|\theta^*\|^2 A + \langle \theta, \theta^* \rangle B)^2 + B^2(\|\theta\|^2 \|\theta^*\|^2 - |\langle \theta, \theta^* \rangle|^2)}}. \quad (\text{A.3})$$

If $\langle \theta, \theta^* \rangle \geq \rho \|\theta\| \|\theta^*\|$, then there exists positive $\Delta = \Delta(\rho, \sigma, \|\theta^*\|, \|\theta\|)$ such that the cosine angle (A.3) is at least $(1 + \Delta)\rho$. Moreover, if $\langle \theta^*, \theta \rangle \geq 0$, then

$$\begin{aligned} \|\theta^*\|^2(1 - \kappa)^2 &\leq \|M(\theta)\|^2 = \|\theta^*\|^2 A^2 + \|\theta\|^2 B^2 + 2\langle \theta, \theta^* \rangle AB \\ &\leq \sigma^2 + 3\|\theta^*\|^2, \end{aligned} \quad (\text{A.4})$$

and

$$\langle \theta^*, M(\theta) \rangle = \|\theta^*\|^2 A + \langle \theta, \theta^* \rangle B \geq \|\theta^*\|^2(1 - \kappa). \quad (\text{A.5})$$

Proof: The stated expression (A.3) for the cosine angle between θ^* and $M(\theta)$ comes from the expression $\frac{\langle u, v \rangle}{\|u\| \|v\|} = \frac{\langle \theta^*, M(\theta) \rangle}{\|\theta^*\| \|M(\theta)\|}$ for the cosine angle between two vectors u and v , and the fact that $M(\theta) = A\theta^* + B\theta$ (see (13)).

Next, we prove the second statement about the lower bound on (A.3). Let $\tau = \frac{\|\theta^*\| A}{\|\theta\| B}$. Observe that

$$\begin{aligned} &\frac{\|\theta^*\|^2 A + \langle \theta, \theta^* \rangle B}{\sqrt{(\|\theta^*\|^2 A + \langle \theta, \theta^* \rangle B)^2 + B^2(\|\theta\|^2 \|\theta^*\|^2 - |\langle \theta, \theta^* \rangle|^2)}} \\ &= \frac{1}{\sqrt{1 + \frac{\|\theta\|^2 \|\theta^*\|^2 - |\langle \theta, \theta^* \rangle|^2}{(\|\theta^*\|^2 \frac{A}{B} + \langle \theta, \theta^* \rangle)^2}}} \\ &\geq \frac{1}{\sqrt{1 + \frac{1 - \rho^2}{(\tau + \rho)^2}}} \\ &= \frac{\rho}{\sqrt{1 - (1 - \rho^2) \frac{\tau(\tau + 2\rho)}{(\tau + \rho)^2}}} \\ &\geq \frac{\rho}{\sqrt{1 - (1 - \rho^2) \frac{\tau}{\tau + \rho}}} \\ &\geq \rho \left(1 + \frac{1}{2}(1 - \rho^2) \frac{\tau}{\tau + \rho}\right), \end{aligned} \quad (\text{A.6})$$

where the last line (A.6) follows from the inequality $1/\sqrt{1 - z} \geq 1 + z/2$ for all $z \in (0, 1)$. Next, note that from Lemma 7,

$$\frac{A}{B} \geq \frac{\sigma^2(1 - \kappa)}{2(\sigma^2 + \|\theta^*\|^2)\kappa^3}.$$

Thus, $\tau \geq \tau' \triangleq \frac{\sigma^2 \|\theta^*\| (1 - \kappa)}{2\|\theta\|(\sigma^2 + \|\theta^*\|^2)\kappa^3}$ and so we can set

$$\Delta = \frac{1}{2}(1 - \rho^2) \left(\frac{\tau'}{\tau' + \rho}\right) > 0.$$

For the statement in (A.4), the identity

$$\|M(\theta)\|^2 = \|\theta^*\|^2 A^2 + \|\theta\|^2 B^2 + 2\langle \theta, \theta^* \rangle AB$$

is an immediate consequence of $M(\theta) = A\theta^* + B\theta$. By Lemma 7, $A \geq 1 - \kappa$ and hence since $\langle \theta, \theta^* \rangle \geq 0$, we have $\|M(\theta)\|^2 \geq \|\theta^*\|^2 A^2 \geq \|\theta^*\|^2(1 - \kappa)^2$.

Next, we will show that $\|M(\theta)\|^2 \leq \sigma^2 + 3\|\theta^*\|^2$ for all θ in \mathbb{R}^d . To see this, note that by Jensen's inequality,

$$\begin{aligned} \langle \theta, M(\theta) \rangle &= \mathbb{E}[(2\phi(W\langle \theta, X \rangle/\sigma^2) - 1)W\langle \theta, X \rangle] \\ &\leq \mathbb{E}[|W\langle \theta, X \rangle|] \\ &\leq \sqrt{\mathbb{E}[|W\langle \theta, X \rangle|^2]} \\ &= \sigma^2 \sqrt{\Lambda^2 + 3\Gamma^2} \\ &= \|\theta\| \sqrt{\sigma^2 + \|\theta^*\|^2 + 2|\langle \theta_0, \theta^* \rangle|^2}. \end{aligned}$$

Next, it can be shown that $|2\phi(z) + 2z\phi'(z) - 1| \leq \sqrt{2}$ and hence $A \leq \sqrt{2}$. Using this, we have

$$\begin{aligned} \langle \theta_0^\perp, M(\theta) \rangle &= A\langle \theta_0^\perp, \theta^* \rangle \\ &\leq \sqrt{2}\langle \theta_0^\perp, \theta^* \rangle. \end{aligned}$$

Putting these two facts together, we have

$$\begin{aligned} \|M(\theta)\|^2 &= |\langle \theta_0^\perp, M(\theta) \rangle|^2 + |\langle \theta_0, M(\theta) \rangle|^2 \\ &\leq \sigma^2 + \|\theta^*\|^2 + 2|\langle \theta_0^\perp, \theta^* \rangle|^2 + 2|\langle \theta_0, \theta^* \rangle|^2 \\ &= \sigma^2 + 3\|\theta^*\|^2. \end{aligned}$$

The final statement (A.5) follows from similar arguments and so we omit them here. \square

Lemma 4: If $\langle \theta, \theta^ \rangle \geq 0$, then*

$$\mathbb{E}[W\langle \theta, X \rangle \phi'(W\langle \theta, X \rangle/\sigma^2)] \geq 0.$$

Proof: Writing $W\langle \theta, X \rangle$ according to the distributional equivalent (17), note that the statement is true if

$$\mathbb{E}[(\alpha Z + \beta)\phi'(\alpha Z + \beta)] \geq 0,$$

where $Z \sim N(0, 1)$ and $\alpha \geq 0$ and $\beta \geq 0$. This fact is proved in [32, Lemma 5] or [28, Lemma 1]. \square

Lemma 5: The following inequalities hold for all $z \in \mathbb{R}$:

$$|2\phi(z) + 2z\phi'(z) - 1| \leq 1 + \sqrt{2(1 - \phi(z))},$$

and

$$z^2 \phi'(z) \leq \sqrt{2(1 - \phi(z))}.$$

Proof: Their validity can easily be established using mathematical software. \square

Lemma 6: Let $\alpha, \beta > 0$ and $Z \sim N(0, 1)$. Then

$$\mathbb{E}[2(1 - \phi(\alpha(Z + \beta)))] \leq \exp\left\{-\frac{\beta}{2} \min\{\alpha, \beta\}\right\}.$$

Moreover,

$$\mathbb{E}[2(1 - \phi(\alpha|Z_2|(Z_1 + \beta|Z_2)))] \leq \frac{1}{\sqrt{\beta \min\{\alpha, \beta\} + 1}}.$$

Proof: The second conclusion follows immediately from the first since

$$\begin{aligned} & \mathbb{E}[2(1 - \phi(\alpha|Z_2|(Z_1 + \beta|Z_2)))] \\ &= 2\mathbb{E}_{Z_2}[\mathbb{E}_{Z_1}[1 - \phi(\alpha|Z_2|(Z_1 + \beta|Z_2))]] \\ &\leq \mathbb{E}_{Z_2}\left[\exp\left\{-\frac{Z_2^2}{2}\beta \min\{\alpha, \beta\}\right\}\right] \\ &= \frac{1}{\sqrt{\beta \min\{\alpha, \beta\} + 1}}. \end{aligned}$$

The last equality follows from the moment generating function of χ_1^2 .

For the first conclusion, we first observe that the mapping $\alpha \mapsto \mathbb{E}[\phi(\alpha(Z + \beta))]$ is increasing (see [32, Lemma 5] or [28, Lemma 1]). Next, note the inequality

$$2(1 - \phi(z)) \leq e^{-z},$$

which is equivalent to $(e^z - 1)^2 \geq 0$. If $\alpha \geq \beta$, then

$$\begin{aligned} \mathbb{E}[2(1 - \phi(\alpha(Z + \beta)))] &\leq \mathbb{E}[2(1 - \phi(\beta(Z + \beta)))] \\ &\leq \mathbb{E}\left[e^{-(\beta(Z + \beta))}\right] \\ &= e^{-\beta^2/2}. \end{aligned}$$

If $\alpha \leq \beta$, then

$$\begin{aligned} \mathbb{E}[2(1 - \phi(\alpha(Z + \beta)))] &\leq \mathbb{E}\left[e^{-\alpha(Z + \beta)}\right] \\ &= e^{\alpha^2/2 - \alpha\beta} \\ &\leq e^{-\alpha\beta/2}. \end{aligned}$$

In each case, we used the moment generating function of a Gaussian distribution to evaluate the expectations. \square

Lemma 7: The following inequalities hold:

$$1 - \kappa \leq A \leq 1 + \sqrt{\kappa},$$

and

$$B \leq 2(1 + \|\theta^*\|^2/\sigma^2)\kappa^3.$$

Proof: By Lemma 4 and Lemma 6,

$$\begin{aligned} A &= \left[2\phi(W\langle\theta, X\rangle/\sigma^2) + 2(W\langle\theta, X\rangle/\sigma^2)\right. \\ &\quad \left.\times \phi'(W\langle\theta, X\rangle/\sigma^2) - 1\right] \\ &\geq \mathbb{E}\left[2\phi(W\langle\theta, X\rangle/\sigma^2) - 1\right] \\ &\geq 1 - \kappa. \end{aligned}$$

By Lemma 5, Jensen's inequality, and Lemma 6,

$$\begin{aligned} A &= \left[2\phi(W\langle\theta, X\rangle/\sigma^2) + 2(W\langle\theta, X\rangle/\sigma^2)\right. \\ &\quad \left.\times \phi'(W\langle\theta, X\rangle/\sigma^2) - 1\right] \\ &\leq \mathbb{E}\left[1 + \sqrt{2(1 - \phi(W\langle\theta, X\rangle/\sigma^2))}\right] \\ &\leq 1 + \sqrt{\mathbb{E}[2(1 - \phi(W\langle\theta, X\rangle/\sigma^2))]} \\ &\leq 1 + \sqrt{\kappa}. \end{aligned}$$

By the inequality $\phi'(z) \leq 2(1 - \phi(z))$ for all $z \in \mathbb{R}$ and Lemma 6,

$$\begin{aligned} B &= 2\mathbb{E}\left[(W^2/\sigma^2)\phi'(W\langle\theta, X\rangle/\sigma^2)\right] \\ &\leq 2\mathbb{E}\left[2(W^2/\sigma^2)(1 - \phi(W\langle\theta, X\rangle/\sigma^2))\right] \\ &= 2(1 + \|\theta^*\|^2/\sigma^2) \\ &\quad \times \mathbb{E}_{Z_2}\left[Z_2^2\mathbb{E}_{Z_1}\left[2\left(1 - \phi\left(\Lambda|Z_2|\left(Z_1 + \frac{\Gamma}{\Lambda}|Z_2|\right)\right)\right)\right]\right] \\ &\leq 2(1 + \|\theta^*\|^2/\sigma^2)\mathbb{E}_{Z_2}\left[Z_2^2 \exp\left\{-\frac{Z_2^2}{2}\frac{\Gamma}{\Lambda} \min\left\{\frac{\Gamma}{\Lambda}, \Lambda\right\}\right\}\right] \\ &= 2(1 + \|\theta^*\|^2/\sigma^2)\left(\frac{1}{\frac{\Gamma}{\Lambda} \min\{\Lambda, \frac{\Gamma}{\Lambda}\} + 1}\right)^{3/2} \\ &= 2(1 + \|\theta^*\|^2/\sigma^2)\kappa^3. \quad \square \end{aligned}$$

Lemma 8: Define

$$h(\alpha, \beta) = \mathbb{E}[(2\phi(\alpha|Z_2|(Z_1 + \beta|Z_2)) - 1)(|Z_2|(Z_1 + \beta|Z_2))].$$

Let $\alpha, \beta > 0$. Then

$$\frac{\partial}{\partial \alpha} h(\alpha, \beta) \leq \frac{2}{\alpha^2} \left(\frac{1}{\beta \min\{\alpha, \beta\} + 1}\right)^{1/4}.$$

Proof: First, observe that

$$\frac{\partial}{\partial \alpha} h(\alpha, \beta) = \mathbb{E}\left[2\phi'(\alpha|Z_2|(Z_1 + \beta|Z_2))(|Z_2|(Z_1 + \beta|Z_2))^2\right].$$

By Lemma 5, Jensen's inequality, and Lemma 6,

$$\begin{aligned} & \mathbb{E}\left[2\phi'(\alpha|Z_2|(Z_1 + \beta|Z_2))(|Z_2|(Z_1 + \beta|Z_2))^2\right] \\ &= \frac{1}{\alpha^2} \mathbb{E}\left[2\phi'(\alpha|Z_2|(Z_1 + \beta|Z_2))(\alpha|Z_2|(Z_1 + \beta|Z_2))^2\right] \\ &\leq \frac{2}{\alpha^2} \mathbb{E}\left[\sqrt{2(1 - \phi(\alpha|Z_2|(Z_1 + \beta|Z_2)))}\right] \\ &\leq \frac{2}{\alpha^2} \sqrt{\mathbb{E}[2(1 - \phi(\alpha|Z_2|(Z_1 + \beta|Z_2)))]} \\ &\leq \frac{2}{\alpha^2} \left(\frac{1}{\beta \min\{\alpha, \beta\} + 1}\right)^{1/4}. \quad \square \end{aligned}$$

ACKNOWLEDGMENT

The authors would like to thank the Associate Editor and two anonymous reviewers whose detailed feedback led to dramatic improvements to this paper.

REFERENCES

- [1] R. E. Quandt and J. B. Ramsey, "Estimating mixtures of normal distributions and switching regressions," *J. Amer. Stat. Assoc.*, vol. 73, no. 364, pp. 730–738, 1978.
- [2] T. R. Turner, "Estimating the propagation rate of a viral infection of potato plants via mixtures of regressions," *J. Roy. Stat. Soc., C (Appl. Statist.)*, vol. 49, no. 3, pp. 371–384, 2000.
- [3] R. D. De Veaux, "Mixtures of linear regressions," *Comput. Statist. Data Anal.*, vol. 8, no. 3, pp. 227–245, Nov. 1989, doi: [10.1016/0167-9473\(89\)90043-1](https://doi.org/10.1016/0167-9473(89)90043-1).
- [4] K. Viele and B. Tong, "Modeling with mixtures of linear regressions," *Statist. Comput.*, vol. 12, no. 4, pp. 315–330, 2002, doi: [10.1023/A:1020779827503](https://doi.org/10.1023/A:1020779827503).
- [5] C. Ding, "Using regression mixture analysis in educational research," *Practical Assessment, Res. Eval.*, vol. 11, no. 11, pp. 1–11, 2006.
- [6] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (GPCA)," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 12, pp. 1945–1959, Dec. 2005.
- [7] M. E. Tipping and C. M. Bishop, "Mixtures of probabilistic principal component analyzers," *Neural Comput.*, vol. 11, no. 2, pp. 443–482, Feb. 1999.
- [8] S. Faria and G. Soromenho, "Fitting mixtures of linear regressions," *J. Stat. Comput. Simul.*, vol. 80, no. 2, pp. 201–225, 2010, doi: [10.1080/00949650802590261](https://doi.org/10.1080/00949650802590261).
- [9] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc., B (Methodol.)*, vol. 39, no. 1, pp. 1–38, 1977.
- [10] E. M. L. Beale and R. J. A. Little, "Missing values in multivariate analysis," *J. Roy. Stat. Soc., B (Methodol.)*, vol. 37, no. 1, pp. 129–145, 1975.
- [11] C. F. J. Wu, "On the convergence properties of the EM algorithm," *Ann. Statist.*, vol. 11, no. 1, pp. 95–103, Mar. 1983, doi: [10.1214/aos/1176346060](https://doi.org/10.1214/aos/1176346060).
- [12] S. Balakrishnan, M. J. Wainwright, and B. Yu, "Statistical guarantees for the EM algorithm: From population to sample-based analysis," *Ann. Statist.*, vol. 45, no. 1, pp. 77–120, 2017, doi: [10.1214/16-AOS1435](https://doi.org/10.1214/16-AOS1435).
- [13] M. I. Jordan and R. A. Jacobs, "Hierarchical mixtures of experts and the EM algorithm," *Neural Comput.*, vol. 6, no. 2, pp. 181–214, 1994.
- [14] Y. Chen, X. Yi, and C. Caramanis, "A convex formulation for mixed regression with two components: Minimax optimal rates," in *Proc. Conf. Learn. Theory*, 2014, pp. 560–604.
- [15] Y. Chen, X. Yi, and C. Caramanis. (2013). "A convex formulation for mixed regression with two components: Minimax optimal rates." [Online]. Available: <https://arxiv.org/abs/1312.7006>
- [16] X. Yi, C. Caramanis, and S. Sanghavi, "Alternating minimization for mixed linear regression," in *Proc. Int. Conf. Mach. Learn.*, Jan. 2014, pp. 613–621.
- [17] A. T. Chaganty and P. Liang, "Spectral experts for estimating mixtures of linear regressions," in *Proc. Int. Conf. Mach. Learn.*, Feb. 2013, pp. 1040–1048.
- [18] K. Zhong, P. Jain, and I. S. Dhillon, "Mixed linear regression with multiple components," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2190–2198.
- [19] H. Sedghi, M. Janzamin, and A. Anandkumar, "Provable tensor methods for learning mixtures of generalized linear models," in *Proc. Artif. Intell. Statist.*, 2016, pp. 1223–1231.
- [20] Y. Li and Y. Liang, "Learning mixtures of linear regressions with nearly optimal complexity," in *Proc. Mach. Learn. Res.*, vol. 75, S. Bubeck, V. Perchet, and P. Rigollet, Eds. Stockholm, Sweden, Jul. 2018, pp. 1125–1144. [Online]. Available: <http://proceedings.mlr.press/v75/li18b.html>
- [21] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, Aug. 2013.
- [22] T. Goldstein and C. Studer, "PhaseMax: Convex phase retrieval via basis pursuit," *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 2675–2689, Apr. 2018.
- [23] O. Dhifallah and Y. M. Lu, "Fundamental limits of phasemax for phase retrieval: A replica analysis," in *Proc. IEEE 7th Int. Workshop Comput. Adv. Multi-Sensor Adapt. Process. (CAMSAP)*, Dec. 2017, pp. 1–5.
- [24] O. Dhifallah, C. Thrampoulidis, and Y. M. Lu. (2018). "Phase retrieval via polytope optimization: Geometry, phase transitions, and new algorithms." [Online]. Available: <https://arxiv.org/abs/1805.09555>
- [25] O. Dhifallah, C. Thrampoulidis, and Y. M. Lu, "Phase retrieval via linear programming: Fundamental limits and algorithmic improvements," in *Proc. 55th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2017, pp. 1071–1077.
- [26] T. T. Cai, X. Li, and Z. Ma, "Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow," *Ann. Statist.*, vol. 44, no. 5, pp. 2221–2251, 2016.
- [27] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.
- [28] C. Daskalakis, C. Tzamos, and M. Zampetakis. (2016). "Ten steps of EM suffice for mixtures of two Gaussians." [Online]. Available: <https://arxiv.org/abs/1609.00368>
- [29] J. Xu, D. J. Hsu, and A. Maleki, "Global analysis of expectation maximization for mixtures of two Gaussians," in *Advances in Neural Information Processing Systems*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. New York, NY, USA: Curran Associates, 2016, pp. 2676–2684.
- [30] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Statist.*, vol. 9, no. 6, pp. 1135–1151, 1981. [Online]. Available: [http://links.jstor.org/sici?sici=0090-5364\(198111\)9:6<1135:EOTMOA>2.0.CO;2-5&origin=MSN](http://links.jstor.org/sici?sici=0090-5364(198111)9:6<1135:EOTMOA>2.0.CO;2-5&origin=MSN)
- [31] G. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*, vol. 382. Hoboken, NJ, USA: Wiley, 2007.
- [32] J. M. Klusowski and W. D. Brinda. (Aug. 2016). "Statistical guarantees for estimating the centers of a two-component Gaussian mixture by EM." [Online]. Available: <https://arxiv.org/abs/1608.02280>
- [33] A. Brutzkus and A. Globerson, "Globally optimal gradient descent for a ConvNet with Gaussian inputs," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 605–614.
- [34] R. Dwivedi, N. Ho, K. Khamaru, M. I. Jordan, M. J. Wainwright, and B. Yu. (2018). "Singularity, misspecification, and the convergence rate of EM." [Online]. Available: <https://arxiv.org/abs/1810.00828>

Jason M. Klusowski (S'16) received a B.S. (*hons.*) in Statistics and Mathematics from the University of Manitoba in 2013, M.A. in Statistics from Yale University in 2017, and Ph.D. in Statistics and Data Science from Yale University in 2018. He is currently Assistant Professor in the Department of Statistics at Rutgers University – New Brunswick. His research interests are high-dimensional statistics, neural networks, random forests, and network analysis.

Dana Yang received a B.S. in Mathematics from Tsinghua University in 2013 and M.A. in Statistics from Yale University in 2014. She is currently a Ph.D. student in the Department of Statistics and Data Science at Yale University. Her research interests are Bayesian analysis, fairness in machine learning, and random graph theory.

W. D. Brinda received a B.S. in Neuroscience (*cum laude*) from Tulane University in 2007, M.S. in Applied Statistics from Indiana University in 2011, and Ph.D. in Statistics and Data Science from Yale University in 2018. He is currently a lecturer and research scholar in the Department of Statistics and Data Science at Yale University. His research interests include model selection, information theory, and causality.