

---

# Provable Tensor Methods for Learning Mixtures of Generalized Linear Models

---

**Hanie Sedghi**  
Allen Institute for AI

**Majid Janzamin**  
University of California, Irvine

**Anima Anandkumar**  
University of California, Irvine

## Abstract

We consider the problem of learning mixtures of generalized linear models (GLM) which arise in classification and regression problems. Typical learning approaches such as expectation maximization (EM) or variational Bayes can get stuck in spurious local optima. In contrast, we present a tensor decomposition method which is guaranteed to correctly recover the parameters. The key insight is to employ certain feature transformations of the input, which depend on the input generative model. Specifically, we employ score function tensors of the input and compute their cross-correlation with the response variable. We establish that the decomposition of this tensor consistently recovers the parameters, under mild non-degeneracy conditions. We demonstrate that the computational and sample complexity of our method is a low order polynomial of the input and the latent dimensions.

**Keywords:** Mixture of generalized linear models, score function, spectral/tensor decomposition.

## 1 Introduction

A generalized linear model (GLM) is a flexible extension of linear regression which allows the response or the output to be a non-linear function of the input via an *activation* function. In other words, in a GLM, the linear regression of the input is passed through an activation function to generate the response. GLMs unify popular frameworks such as logistic regression and Poisson regression with linear regression. At the

same time, they can be learnt with guarantees using simple iterative methods [Kakade et al., 2011].

In many scenarios, however, GLMs may be too simplistic, and mixtures of GLMs can be much more effective since they combine the expressive power of latent variables with the predictive capabilities of the GLM. Mixtures of GLMs have widespread applicability including object recognition [Quattoni et al., 2004], human action recognition [Wang and Mori, 2009], syntactic parsing [Petrov and Klein, 2007], and machine translation [Liang et al., 2006].

Traditionally, mixture models are learnt through heuristics such as expectation maximization (EM) [Jordan and Jacobs, 1994, Xu et al., 1995] or variational Bayes [Bishop and Svensen, 2003]. However, these methods can converge to spurious local optima and have slow convergence rates for high dimensional models. In contrast, we employ a method-of-moments approach for guaranteed learning of mixtures of GLMs.

The method of moments paradigm dates back to Pearson [Pearson, 1894], and involves fitting the observed moments to parametric distributions. Recently, it has been highly successful in unsupervised learning of a wide range of latent variable models such as Gaussian mixtures, topic models, hidden Markov models [Anandkumar et al., 2014a], network community models [Anandkumar et al., 2013], mixture of ranking models [Awasthi et al., 2014, Oh and Shah, 2014], and so on. The basic idea is to find an efficient spectral decomposition of low order observed moment tensors. Under natural non-degeneracy assumptions, the tensor method is guaranteed to correctly recover the underlying model parameters with low computational and sample complexities. Moreover, in practice, these methods are embarrassingly parallel and scalable to large-scale datasets [Huang et al., 2014].

Earlier works on tensor methods [Anandkumar et al., 2014a] consider unsupervised learning and a key assumption is that the observables are linear functions of the latent variables (in expectation). However, here,

---

Appearing in Proceedings of the 19<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 41. Copyright 2016 by the authors.

we consider mixtures of GLMs, which are non-linear, and this rules out a direct application of tensor methods.

We address the above challenges with the following insight: we have additional flexibility in the regression setting since we have both the response and the input. We can therefore form different moments involving transformations of the input and the response. What are the appropriate transforms for forming the moments which are amenable to tensor decomposition methods? As detailed below, the key ingredient is using a specific feature transformation of the input, based on its probability distribution.

### 1.1 Summary of Result

The main contribution of this work is to provide a guaranteed method for learning mixtures of GLMs using score function transformations. The  $m^{\text{th}}$  order score function  $\mathcal{S}_m(x)$  is related to the normalized  $m^{\text{th}}$  order derivative of the pdf of input  $x$ , see (9). We assume knowledge of these score functions, and this can be estimated via various unsupervised learning methods using only unlabeled samples (e.g. spectral methods).

We then construct the cross-moment tensor between the response variable and the input score function. We establish that the decomposition of this tensor consistently recovers the components of the GLM mixture under some simple non-degeneracy assumptions. Let the response or the output  $y$  be generated from a mixture of GLMs:  $\mathbb{E}[y|h, x] = g(\langle Uh, x \rangle) + \langle \tilde{b}, h \rangle$ , where  $g(\cdot)$  is the activation function,  $x$  is the input and  $h$  is the hidden choice variable. Let  $r$  be the number of mixture components,  $d$  be the input dimension and  $s_{\min}(U)$  be the  $r^{\text{th}}$  largest singular value of  $U$ . Assume the weight matrix  $U = [u_1 \dots u_r]$  is full column rank. Then, we have the following result.

**Theorem 1** (Informal Result). *We recover the weight vectors  $\{u_i\}$  (up to scaling) by performing tensor decomposition on the cross-moment tensor  $\mathbb{E}[y \cdot \mathcal{S}_3(x)]$ . If we have  $n = \tilde{O}\left(\frac{d^3 r^4}{\epsilon^2 s_{\min}^2(U)}\right)$  samples, the error in recovering each weight vector  $u_i$  is bounded by  $\epsilon$ .*

The above result requires third order score function  $\mathcal{S}_3(x)$  to consistently estimate the weight vectors  $\{u_i\}$  of the GLM components in the mixture. Note that the second order score function  $\mathcal{S}_2(x)$  is only a matrix (assuming a vector input  $x$ ) and can only identify the weights  $\{u_i\}$  up to the subspace. Thus, we require at least the third order score function to consistently estimate the GLM mixture model. When the number of components exceeds the input dimension, the full column rank assumption on  $U$  is violated, and in this

case, we can resort to higher order score functions to consistently estimate the parameters.

We employ the tensor decomposition methods from [Anandkumar et al., 2014a,b] to learn the weight vectors  $u_i$  (up to scaling). The tensor method is efficient to implement and does not suffer from spurious local optima. Thus, we guarantee consistent estimation of the weight vectors of GLM mixtures through decomposition of the cross-moment tensor involving the response variable and the input score functions. Our method is shown in Algorithm 1.

### 1.2 Overview of Techniques

**Representation learning is the key:** A crucial ingredient in this work is to first learn the probabilistic model of the input, and employ transformations based on the model for learning the GLM mixture. Thus, we characterize how unsupervised learning on the input can be carried over for learning conditional models of the output via tensor methods.

The feature transformations we employ are the (higher order) score functions,<sup>1</sup> which capture local variation of the probability density function of the input. This follows a recent key result that the cross-moments between the response variable and the input score functions yield (expected) derivatives of the response, as a function of the input [Janzamin et al., 2014].

**Incorporating score functions into tensor decomposition framework:** In this paper, we exploit the above result to form the expected derivatives of the output as a function of the input. We then show that the expected derivatives have a nice relationship with the unknown parameters of the GLM mixture, and the form reduces to a tensor CP decomposition form. We require only a mild assumption on the activation function that it has non-vanishing third derivative (in expectation). For linear regression, this condition is violated, but we can easily overcome this by considering higher powers of the output in the moment estimation framework.

## 2 Problem Formulation

**Notations:** Let  $[n]$  denote the set  $\{1, 2, \dots, n\}$ . Let  $e_i \in \mathbb{R}^d$  denote the standard basis vectors in  $\mathbb{R}^d$ . Let

<sup>1</sup>In this paper, we refer to the derivative of the log of the density function with respect to the variable as the score function. In other works, typically, the derivative is taken with respect to some model parameter [Jaakkola et al., 1999]. Note that if the model parameter is a location parameter, the two quantities only differ in the sign. Higher order score functions involve higher order derivatives of the density function. For the exact form, refer to [Janzamin et al., 2014].

$I_d \in \mathbb{R}^{d \times d}$  denote the identity matrix.  $\tilde{\mathcal{O}}$  denotes the order when ignoring polylog factors. Throughout this paper,  $\nabla_x^{(m)}$  denotes the  $m$ -th order derivative w.r.t. variable  $x$  and notation  $\otimes$  represents tensor (outer) product.

A real  $p$ -th order tensor  $T \in \bigotimes_{i=1}^p \mathbb{R}^{d_i}$  is a member of the tensor product of Euclidean spaces  $\mathbb{R}^{d_i}$ ,  $i \in [p]$ . As is the case for vectors (where  $p = 1$ ) and matrices (where  $p = 2$ ), we may identify a  $p$ -th order tensor with the  $p$ -way array of real numbers  $[T_{i_1, i_2, \dots, i_p} : i_1, i_2, \dots, i_p \in [d]]$ , where  $T_{i_1, i_2, \dots, i_p}$  is the  $(i_1, i_2, \dots, i_p)$ -th coordinate of  $T$  with respect to a canonical basis.

**CP decomposition and tensor rank:** A 3rd order tensor  $T \in \mathbb{R}^{d \times d \times d}$  is said to be rank-1 if it can be written in the form  $T = a \otimes b \otimes c \Leftrightarrow T(e_i, e_j, e_l) = a(i) \cdot b(j) \cdot c(l)$ , where notation  $\otimes$  represents the *tensor product*. A tensor  $T$  is said to have a CP rank  $k \geq 1$  if it can be written as the sum of  $k$  rank-1 tensors  $T = \sum_{i \in [k]} a_i \otimes b_i \otimes c_i$ .

## 2.1 Learning Problem

Let  $y$  denote the output and  $x \in \mathbb{R}^d$  be the input. We consider both the regression setting, where  $y$  can be continuous or discrete, or the classification setting, where  $y$  is discrete. For simplicity, we assume  $y$  to be a scalar: in the classification setting, this corresponds to binary classification ( $y \in \{-1, 1\}$ ).

We consider the *realizable* setting, where we assume that the output  $y$  is drawn from an associative model  $p(y|x)$ , given input  $x$ . In addition, we assume that the input  $x$  is drawn from some continuous probability distribution with density function  $p(x)$ . We will incorporate this generative model in our algorithm for learning the associative model.

We first consider mixtures of generalized linear models (GLM) [Agarwal et al., 2014, Kakade et al., 2011] and then extend to mixture of GLMs with nonlinear transformations. The class of GLMs is given by

$$\mathbb{E}[y|x] = g(\langle u, x \rangle + b), \quad (1)$$

where  $g$  is the *activation* function,  $u$  is the weight vector, and  $b$  is the bias.  $g(\cdot)$  is usually chosen to be the logistic function, although we do not impose this limitation. In the binary classification setting, (1) corresponds to a single classifier. Note that a linear regression can be modeled using a linear activation function. Throughout this paper we assume that noise is independent of the input.

A mixture of  $r$  GLM models is then given by employing a hidden choice variable  $h \in \{e_1, e_2, \dots, e_r\}$ , where  $e_i$

---

**Algorithm 1** Learning mixture of associative models  
 $\mathbb{E}[y|x, h] = g(\langle Uh, x \rangle + \langle \tilde{b}, h \rangle)$

---

**input** Labeled samples  $(x_i, y_i), i \in [n]$ .

**input** Score function of the input  $\mathcal{S}_3(x)$  as in Equation (9).

- 1: Compute  $\widehat{M}_3 = \frac{1}{n} \sum_i y_i \cdot \mathcal{S}_3(x_i)$ , Empirical estimate of  $M_3$ .
  - 2:  $\{\hat{u}_j\}_{j \in [r]} = \text{tensor power decomposition}(\widehat{M}_3)$ . (Algorithm 4 in the Appendix)
  - 3: Recover scale and biases using EM (as in Appendix 3).
- 

is the basis vector in  $\mathbb{R}^r$  to select of the  $r$  GLM models, i.e.

$$\mathbb{E}[y|x, h] = g(\langle Uh, x \rangle + \langle \tilde{b}, h \rangle), \quad (2)$$

where  $U = [u_1 | u_2 \dots | u_r] \in \mathbb{R}^{d_x \times r}$  has the  $r$  weight vectors of component GLMs as columns and  $\tilde{b} \in \mathbb{R}^r$  is the vector of biases for the component GLMs. Let  $w := \mathbb{E}[h]$  be the probability vector for selecting the different GLMs.

We then extend our results to learning mixture of GLMs with nonlinear transformations where

$$\mathbb{E}[y|x, h] = g(\langle Uh, \phi(x) \rangle, \langle \tilde{b}, h \rangle), \quad (3)$$

for some known function  $\phi(\cdot)$ .

Given training samples  $\{x_i, y_i\}$ , our goal is to learn the parameters of the associative mixture described above. We consider a moment-based approach, which involves cross-moments of  $y$  and a function of  $x$ . We first assume that the exact moments are available, and we later carry out sample analysis, when empirical moments are used.

Throughout this paper we make the following assumptions unless otherwise stated. Derivative and expectation are interchangeable. The activation function  $g$  is differentiable up to the third order. The choice variable is independent of the input  $x$ , i.e.,  $h$  does not depend on  $x$ . The score function  $\nabla_x \log p(x)$  exists and all the entries of  $g(x) \cdot p(x)$  go to zero on the boundaries of support of  $p(x)$ .

## 3 Learning under Gaussian Input

We now present the method for learning the mixture models in (2) and (3). We first start with the simple case, where the input  $x$  is Gaussian, and we have a single GLM model, instead of a mixture, and then extend to more general cases.

### 3.1 Toy Example: single GLM

We first assume a white Gaussian input  $x \sim \mathcal{N}(0, I_d)$  to demonstrate our ideas. Assuming that  $y$  is generated from a GLM

$$\mathbb{E}[y|x] = g(\langle u, x \rangle + b),$$

we have the following result on the cross-moment  $\mathbb{E}[y \cdot x]$ .

**Lemma 2** (Moment form for Gaussian input and single GLM). *We have*

$$M_1 = \mathbb{E}[y \cdot x] = \mathbb{E}[\nabla_{x'} g(x')] \cdot u,$$

where the expectation is over  $x' := \langle u, x \rangle + b$ , and  $x \sim \mathcal{N}(0, I_d)$ .

Proof follows from Stein's identity [Stein, 1972] as discussed below. Thus, by forming the first-order cross-moment  $M_1$ , we can recover the weight vector  $u$  up to scaling. Note that the scaling and the bias  $b$  are just scalar parameters which can be estimated separately.

The main message behind Lemma 2 is that the cross-moments between the output  $y$  and the input  $x$  contain valuable information about the associative model. In the special case of Gaussian input and single GLM, the first order moment is sufficient to learn almost all the parameters of the GLM. But how general is this framework? Can we use a moment-based framework when there are mixture of GLMs? We exhibit that higher order moments can be used to learn the GLM mixture under Gaussian input in the next section. What about the case when the input is not Gaussian, but has some general distribution? We consider this setting in Section 4 and show that surprisingly we can form the appropriate cross-moments for learning under any general (continuous) input distribution.

**Stein's Identity:** The proof of Lemma 2 follows from the Stein's identity for Gaussian distribution. It states that for all functions  $G(x)$  satisfying mild regularity conditions, we have [Stein, 1972]

$$\mathbb{E}[G(x) \cdot x] = \mathbb{E}[\nabla_x G(x)]. \quad (4)$$

Thus, Lemma 2 is a direct application of the Stein's identity by substituting  $G(x)$  with  $g(\langle u, x \rangle + b)$ .

### 3.2 Learning GLM mixtures

We now consider learning mixture of GLMs

$$\mathbb{E}[y|x, h] = g(\langle Uh, x \rangle + \langle \tilde{b}, h \rangle),$$

where  $U = [u_1 | u_2 \dots | u_r]$  has the  $r$  weight vectors of component GLMs as columns and  $\tilde{b}$  is the vector of

biases for the component GLMs. Recall that  $w := \mathbb{E}[h]$  is the probability vector for selecting different GLMs.

For the mixture of GLMs, the first order moment  $M_1 := \mathbb{E}[y \cdot x]$  is now a combination of (scaled) weight vectors  $u_i$ 's, i.e.

$$M_1 := \mathbb{E}[y \cdot x] = \sum_{j \in [r]} w_j \mathbb{E}[\nabla_{x'_j} g(x'_j)] u_j,$$

where the expectation is over  $x'_j = \langle u_j, x \rangle + \tilde{b}_j$ , and  $x \sim \mathcal{N}(0, I_d)$ . Thus, the first order moment does not suffice for learning mixture of GLMs.

Now, let us look at the second order moment,

$$M_2 := \mathbb{E}[y \cdot (x \otimes x - I)] = \sum_{j \in [r]} \mathbb{E}[\nabla_{x'_j}^{(2)} g(x'_j)] w_j \cdot u_j \otimes u_j,$$

where, as before, the expectation is over  $x'_j = \langle u_j, x \rangle + \tilde{b}_j$ . If the expectations (and  $w_j$ 's) are non-zero, then we can recover the subspace spanned by the weight vectors  $u_j$ 's. However, we cannot recover the individual weight vectors  $u_j$ 's. Moreover, if the biases  $\tilde{b} = 0$  and  $g$  is a symmetric function, then the expectations are zero, and the second order moment  $M_2$  vanishes. A mirror trick is introduced in [Sun et al., 2013b] to alleviate this problem, but this still only recovers the subspace spanned by the  $u_j$ 's.

We now consider the third order moment  $M_3$  in the hope of recovering the weight vectors  $u_j$ 's for mixture of GLMs. We show that by adjusting the moment  $\mathbb{E}[y \cdot x \otimes x \otimes x]$  appropriately, we obtain a CP tensor form in terms of the weight vectors  $u_j$ 's. Specifically, consider

$$\begin{aligned} M_3 := \mathbb{E}[y \cdot x \otimes x \otimes x] &- \sum_{j \in [d]} \mathbb{E}[y \cdot e_j \otimes x \otimes e_j] \quad (5) \\ &- \sum_{j \in [d]} \mathbb{E}[y \cdot e_j \otimes e_j \otimes x] - \sum_{j \in [d]} \mathbb{E}[y \cdot x \otimes e_j \otimes e_j]. \end{aligned}$$

Note that  $M_3$  can be considered as a special case of the form  $\mathbb{E}[y \cdot \mathcal{S}_3(x)]$  for white Gaussian input  $x \sim \mathcal{N}(0, I_d)$ , where  $\mathcal{S}_3(x)$  is the third order score function of the input as defined in Section 4.1.

**Lemma 3** (Adjusted third order moments). *We have*

$$M_3 = \sum_{j \in [r]} \rho_j w_j \cdot u_j \otimes u_j \otimes u_j, \quad (6)$$

where  $\rho_j := \mathbb{E}[\nabla_{x'_j}^{(3)} g(x'_j)]$  and the expectation is over  $x'_j = \langle u_j, x \rangle + \tilde{b}_j$ .

The proof follows from Stein's Identity. See Appendix 1.1 for details. Having the CP-form allows us to recover the component weight vectors through the tensor decomposition method. We present the result below.

**Theorem 4** (Recovery of mixture of GLMs). *Assuming that the weight matrix  $U \in \mathbb{R}^{d \times r}$  is full column rank,  $\rho_j, w_j \neq 0 \ \forall j \in [r]$ , given  $M_3$ , we can recover the component weight vectors  $u_j, j \in [r]$ , up to scaling, using tensor method given in Algorithm 4 (in the Appendix).*

The proof follows from Lemma 3. The computational complexity of tensor decomposition in this factor form is  $O(nrdL)$ , where  $n$  is the number of samples and  $L$  is the number of initialization. Having recovered the normalized weight vectors, we can then estimate the scaling and the biases through expectation maximization or other methods. These are just  $2r$  additional parameters, and thus, the majority of the parameters are estimated by the tensor method.

**Theorem 5** (Sample Complexity). *Assume the conditions for Theorem 4 are met. Suppose the sample complexity*

$$n = \tilde{O} \left( \frac{d^3 r^4}{\epsilon^2 s_{\min}^2(U)} \right),$$

then for each weight vector  $u_j$ , the estimate  $\hat{u}_j$  from line 2 Algorithm 1 satisfies w.h.p

$$\|u_j - \hat{u}_j\| \leq \tilde{O}(\epsilon), \quad j \in [r].$$

**Proof outline:** From Lemma 3, we know that the exact cross-moment  $\mathbb{E}[y \cdot \mathcal{S}_3(x)]$  has rank-one components as columns of matrix  $U$ ; see Equation (6) for the tensor decomposition form. Thus given the exact moment, the theorem is proved by applying the tensor decomposition guarantees in Anandkumar et al. [2014c]. In the noisy case where the moment is empirically formed by observed samples, we use the analysis and results of tensor power iteration in Anandkumar et al. [2014d]. They show that when the perturbation tensor is small, the tensor power iteration initialized by the SVD-based Procedure 3 in the Appendix recovers the rank-1 components up to some small error. The sample complexity is also proved by applying standard matrix concentration inequalities. In particular, we matricize the error tensor between exact moment and the empirical moment, and bound its norm with matrix Bernstein’s inequality.

**Remark :** We can also handle the case when the full column rank assumption on  $U \in \mathbb{R}^{d \times r}$  is violated under some additional constraints. In the overcomplete regime, we have the latent dimensionality exceeding the input dimensionality, i.e.  $r > d$ . The tensor method can still recover the weight vectors  $u_j$ , if we assume they are incoherent. A detailed analysis of overcomplete tensor decomposition is given in [Anandkumar et al., 2014d].

**Remark :** If we assume the  $u_j$  are normalized, the above approach suffices to completely learn the parameters  $w_j$ . This is because we obtain  $w_j \rho_j$  and we have the knowledge of  $\rho_j$ , where the activation function and the input distributions are known. Otherwise, we need to perform EM to fully learn the weights. Note that initializing with our method results in performing EM in a low dimension instead of input dimension. The reason is that the only unknown parameters are the scale and biases of the components. We initialize with the output of our method (Algorithm 1) and proceed with EM algorithm as proposed by Xu et al. [1995]. For details see Appendix 3.

**Remark:** If  $\rho_j = 0$ , which is the case for mixture of linear regression, we cannot recover the weight vectors from the tensor given in (5). In this case, we form a slightly different tensor to recover the weight vectors. We elaborate on this in the next section.

**Remark:** Our results can be easily extended to multi-label and multi-class settings (one-versus-all strategy) as well as vector-valued regression problems.

### 3.3 Learning Mixtures of Linear Regression

We now consider mixtures of linear regressions:

$$\mathbb{E}[y|x, h = e_j] = w_j \langle u_j, x \rangle + b_j,$$

where  $e_j \in \mathbb{R}^r$  denotes the  $j$ -th basis vector.

In this case higher order derivatives ( $m \geq 2$ ) of the activation function vanish. Therefore, the cross-moment matrix and tensor defined in Section 3.2 can not yield the parameters. For this setting, we form

$$\begin{aligned} M_2 &:= \mathbb{E}[y^2 \cdot (x \otimes x - I)] \\ M_3 &:= \mathbb{E}[y^3 \cdot x \otimes x \otimes x] - \sum_{j \in [d]} \mathbb{E}[y^3 \cdot e_j \otimes x \otimes e_j] \quad (7) \\ &\quad - \sum_{j \in [d]} \mathbb{E}[y^3 \cdot e_j \otimes e_j \otimes x] - \sum_{j \in [d]} \mathbb{E}[y^3 \cdot x \otimes e_j \otimes e_j]. \end{aligned}$$

**Lemma 6** (Adjusted third order moments). *We have*

$$M_3 = \sum_{j \in [r]} \tilde{\rho}_j w_j \cdot u_j \otimes u_j \otimes u_j. \quad (8)$$

The proof follows from Stein’s Identity and it is provided in Appendix 1.2. Having the CP-form allows us to recover the component weight vectors through the tensor decomposition method. We present the result below.

**Theorem 7** (Recovery of linear regression mixtures). *Assuming that the weight matrix  $U \in \mathbb{R}^{d \times r}$  is full column rank,  $\rho_j, w_j \neq 0 \ \forall j \in [r]$ , given  $M_3$  as in (7), we can recover the component weight vectors  $u_j, j \in [r]$ ,*

up to scaling, using tensor method given in Algorithm 4 (in the Appendix).

The proof is similar to Lemma 6.

**Theorem 8** (Sample Complexity). *Assume the conditions for Theorem 7 are met. Suppose the sample complexity*

$$n = \tilde{O}\left(\frac{d^3 r^4}{\epsilon^2 s_{\min}^2(U)}\right),$$

then for each weight vector  $u_j$ , the estimate  $\hat{u}_j$  in line 2 Algorithm 1 satisfies w.h.p

$$\|u_j - \hat{u}_j\| \leq \tilde{O}(\epsilon), \quad j \in [r].$$

The proof follows the same approach as the one described for Theorem 5.

## 4 Learning GLM Mixtures under General Input Distribution

In the previous section, we established consistent estimation of the parameters of mixture of GLMs under Gaussian input. However, this assumption is limiting, since the input is usually far from Gaussian in any real scenario. We now extend the results in the previous section to any general (continuous) input.

### 4.1 Extensions of Stein's identity

The key ingredient that enabled learning in the previous section is the ability to compute the expected derivatives of the output as a function of the input. Stein's identity shows that these derivatives can be obtained using the cross-moments between the output and the score function of input. Is there a general unified framework where we can compute the expected derivatives under any general input distribution?

Janzamin et al. [2014] provide an affirmative answer. They show that by computing the cross-moment between the output and the (higher order) score functions of the input, we compute expected derivatives of any order. This key result allows us to extend the results in the previous section to any general input distribution.

**Definition: Score function** The score of  $x \in \mathbb{R}^d$  with pdf  $p(x)$ , denoted by  $\mathcal{S}_1(x)$ , is the random vector  $\nabla_x \log p(x)$ . Janzamin et al. [2014], define the  $m^{\text{th}}$  order score function as

$$\mathcal{S}_m(x) := (-1)^m \frac{\nabla^{(m)} p(x)}{p(x)}. \quad (9)$$

They have also shown that score function can be equivalently derived using the recursive form

$$\mathcal{S}_m(x) = -\mathcal{S}_{m-1}(x) \otimes \nabla_x \log p(x) - \nabla_x \mathcal{S}_{m-1}(x). \quad (10)$$

**Theorem 9** (Higher order derivatives [Janzamin et al., 2014]). *For random vector  $x \in \mathbb{R}^d$ , let  $p(x)$  and  $\mathcal{S}_m(x)$  respectively denote the pdf and the corresponding  $m$ -th order score function. Consider any continuously differentiable output-function  $\mathbb{E}[y|x] = g(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  satisfying some mild regularity conditions. Then we have*

$$\mathbb{E}[y \cdot \mathcal{S}_m(x)] = \mathbb{E}[g(x) \cdot \mathcal{S}_m(x)] = \mathbb{E}\left[\nabla_x^{(m)} g(x)\right].$$

For details, see [Janzamin et al., 2014]. In order to learn mixture of GLMs for general input distributions, we utilize score function  $\mathcal{S}_m(\cdot)$  of order  $m = 3$ .

### 4.2 Moment forms

We now consider the cross-moment  $M_3 := \mathbb{E}[y \cdot \mathcal{S}_3(x)]$ , which is a third order tensor. By alluding to Theorem 9, we show that the moment  $M_3$  has a CP decomposition where the components are the weight vectors  $u_j$ 's.

**Theorem 10** (Recovery of mixture of GLMs under general input). *Given score function  $\mathcal{S}_3(x)$  as in (9), we have*

$$M_3 := \mathbb{E}[y \cdot \mathcal{S}_3(x)] = \sum_{j \in [r]} \rho_j \cdot w_j \cdot u_j^{\otimes 3},$$

where  $\rho_j := \mathbb{E}[\nabla_{x'_j}^{(3)} g(x'_j)]$  and the expectation is with respect to  $x'_j := \langle u_j, x \rangle + \tilde{b}_j$ .

Assuming that matrix  $U$  is full column rank and  $\rho_j, w_j > 0, \forall j$ , we can recover the weight vectors  $u_j, j \in [r]$ , up to scaling, using tensor decomposition on  $M_3$  given in Algorithm 4 (in the Appendix).

The proof follows from Theorem 9. Thus, we have a guaranteed recovery of the weight vectors of mixture of GLMs under any general input distribution.

**Remark: Sample complexity:** For general input sample complexity can be found in a similar approach to Gaussian case. The general form is  $n \geq \tilde{O}\left(\mathbb{E}[\|H_3(x)H_3^\top(x)\|] \frac{d^{1.5} r^4}{\epsilon^2 s_{\min}^2(U)}\right)$ . Here  $H_3(x) \in \mathbb{R}^{d \times d^2}$  is the matricization of  $\mathcal{S}_3(x)$ . Theorem 5 follows from the fact that for Gaussian input  $\mathbb{E}[\|H_3(x)H_3^\top(x)\|] = O(d^{1.5})$ .

**Remark: Score function estimation:** There are various efficient methods for estimating the score

function. The framework of score matching is popular for parameter estimation in probabilistic models [Hyvärinen, 2005, Swersky et al., 2011], where the criterion is to fit parameters based on matching the data score function. For instance, Swersky et al. [2011] analyzes fitting the data to RBM (Restricted Boltzmann Machine) model. Therefore, one option is to use this method for estimating  $\mathcal{S}_1(x)$  and use the recursive form in (10) to estimate higher order score functions for the active layer.

**Remark: Computational Complexity:** If we fit the input data into an RBM model, the computational complexity of our method, when performed in parallel, is  $O(\log(\min(d, d_h)))$  with  $O(rnLdd_h/\log(\min(d, d_h)))$  processors. Here  $d_h$  is the number of neurons of the first layer of the RBM used for approximating the score function.

### 4.3 Learning Mixtures of Linear Regression

As discussed earlier, our framework can easily handle the case of mixtures of linear regression. Here, we describe it under general input distribution. Let,

$$\mathbb{E}[y|x, h = e_j] = w_j \langle u_j, x \rangle + b_j,$$

where  $x, y$  respectively denote the input and output, and  $h$  is the hidden variable that chooses the regression parameter  $u_j$  from the set  $\{u_j\}_{j \in [r]}$ ,  $w_j = p(h = e_j)$  and  $b_j$  is the bias.

**Theorem 11** (Recovery of linear regression mixtures under general input). *Given score function  $\mathcal{S}_3(x)$  as in Equation (9), we have*

$$M_3 = \mathbb{E}[y^3 \cdot \mathcal{S}_3(x)] = \sum_{j \in [r]} w_j \cdot u_j^{\otimes 3}.$$

Assuming that matrix  $U$  is full column rank,  $w_j \neq 0, \forall j$ , we can recover the weight vectors  $u_j, j \in [r]$ , up to scaling, using tensor decomposition on  $M_3$  given in Algorithm 4 (in the Appendix).

For proof, see Appendix 1.3.

**Remark: Sample complexity:** For general input sample complexity can be found in a similar approach to Gaussian case. The general form is  $n \geq \tilde{O}\left(\mathbb{E}[\|H_3(x)H_3^\top(x)\|] \frac{d^{1.5}r^4}{\epsilon^2 s_{\min}^2(U)}\right)$ . Here  $H_3(x) \in \mathbb{R}^{d \times d^2}$  is the matricization of  $\mathcal{S}_3(x)$ . Theorem 8 follows from the fact that for Gaussian input  $\mathbb{E}[\|H_3(x)H_3^\top(x)\|] = O(d^{1.5})$ .

**Remark:** Chaganty and Liang [2013] consider learning a mixture of linear regression models, using tensor decomposition approach on the higher order moments

of the output  $y$ . They model the problem as an optimization on a third order tensor and prove that the optimal tensor would have the weight vectors as its rank-1 components. Minimizing an objective function over a tensor variable is expensive (in fact, quadratic for each variable [Liu and Vandenberghe, 2009], and their computational complexity scales as  $O(nd^{12})$ . Hence their proposed method is not practical in large scale. Whereas, as discussed earlier our computational complexity is  $O(nd^2)$ . While we require the additional knowledge of the input distribution, in many scenarios, this is not a major limitation since there are large amounts of unlabeled samples which can be used for model estimation. Moreover, we can handle non-linear mixtures, while Chaganty and Liang [2013] limit to linear ones.

### 4.4 Extension to Mixture of GLMs with Nonlinear Transformations

We have so far provided guarantees for learning mixture of GLMs. We now extend the results to cover non-linear models. We consider the class of mixture of GLMs with nonlinear transformations under the realizable setting as

$$\mathbb{E}[y|x, h] = g\left(\langle Uh, \phi(x) \rangle + \tilde{b}\right), \quad (11)$$

where  $\phi(x)$  represents the nonlinear mapping of  $x$ . Assuming that  $\phi(\cdot)$  is known, we propose simple ideas to extend our previous results to the setting in (11).

The key idea is to compute the score function  $\mathcal{S}_m(\phi(x))$  corresponding to  $\phi(x)$  rather than the input  $x$ . There is a simple relationship between the scores. The connection can be made from the probability density of the transformed variable as follows. Let  $t = \phi(x)$ ,  $D_t(i, j) := \left[\frac{\partial x_i}{\partial t_j}\right]$ . We have

$$p_{\phi(x)}(t_1, \dots, t_p) = p_x(\phi_1^{-1}(t), \dots, \phi_p^{-1}(t)) |\det(D_t)|, \quad (12)$$

$$\mathcal{S}_m(t) = (-1)^m \frac{\nabla_t^{(m)} p_{\phi(x)}(t)}{p_{\phi(x)}(t)}.$$

**Theorem 12** (Recovery of mixture of GLMs with nonlinear transformations under general input). *Given score function  $\mathcal{S}_3(\phi(x))$  as in Equation (12), we have*

$$M_3 := \mathbb{E}[y \cdot \mathcal{S}_3(\phi(x))] = \sum_{j \in [r]} \rho_j \cdot w_j \cdot u_j^{\otimes 3},$$

where  $\rho_j := \mathbb{E}[\nabla_{z_j}^{(3)} g(z_j)]$  and the expectation is with respect to  $z_j := \langle u_j, \phi(x) \rangle + \tilde{b}_j$ .

Assuming that matrix  $U$  is full column rank,  $w_j, \rho_j \neq 0, \forall j$ , we can recover the weight vectors  $u_j, j \in [r]$ ,

up to scaling, using tensor decomposition on  $M_3$  given in Algorithm 4 (in the Appendix).

## 5 Related Works

**Mixture of Experts/ Regression Mixtures:** The mixture of experts model was introduced as an efficient probabilistic “divide” and “conquer” paradigm in [Jordan and Jacobs, 1994]. Since then, it has been considered in a number of works, e.g. [Xu et al., 1995, Bishop and Svensen, 2003]. Learning is carried out usually through EM [Jordan and Jacobs, 1994, Xu et al., 1995] or variational approaches [Bishop and Svensen, 2003], but the methods have no guarantees. Works with guaranteed learning of associative mixture models are fewer. Chaganty and Liang [2013] consider learning a mixture of linear regression models, using tensor decomposition approach on the higher order moments of of the label  $y$ . Yi et al. [2013] also consider mixed linear regression problem with two components and provide consistency guarantees in the noiseless setting for an alternating minimization method. Chen et al. [2014] provide an alternative convex method for the same setting under noise and established near optimal sample complexity. However, all these guaranteed methods are restricted to mixture of linear regressions and do not extend to non-linear models.

**Learning mixture of GLMs:** For the mixture of generalized linear models (GLM), Li [1992] and Sun et al. [2013a] present methods for learning the subspace of the weight vectors of the component GLMs, assuming that the input is white Gaussian distribution. Li [1992] propose the so-called principal Hessian directions (PHd), where the eigenvectors of the second-order moment matrix  $\mathbb{E}[y \cdot x \otimes x]$  are used to learn the desired subspace. However, the PHd method fails when the output  $y$  is a symmetric function of the input  $x$ , since the moment matrix vanishes in this case. Sun et al. [2013a] overcome this drawback through their clever “mirroring” trick which transforms the output  $y$  to  $r(y)$  such that the resulting second order moment  $\mathbb{E}[r(y) \cdot x \otimes x]$  matrix does not vanish.

Our work has some key differences: the works in [Li, 1992, Sun et al., 2013a] assume Gaussian input  $x$ , while we allow for any probabilistic model (with continuous density function). Another important difference between [Li, 1992, Sun et al., 2013a] and our work, is that we use tensor-based learning techniques, while [Li, 1992, Sun et al., 2013a] only operate on matrices. Operating on tensors allows us to learn the individual weight vectors (up to scaling) of the mixture components, while [Li, 1992, Sun et al., 2013a] only learn the subspace of the weight vectors.

**Spectral/Moment based methods for discrim-**

**inative learning:** Karampatziakis and Mineiro [2014] obtain discriminative features via generalized eigenvectors. They consider the tensor  $\mathbb{E}[y \otimes x \otimes x]$  and then treat  $\mathbb{E}[x \otimes x | y = i]$  as the signal for class  $i$  and  $\mathbb{E}[x \otimes x | y = j]$  as the noise due to class  $j$ . They contrast their method against classical discriminative procedures such as Fisher LDA and show good performance on many real datasets. However, their method has some drawbacks: they cannot handle continuous  $y$ , and also when  $y$  has a large number of classes  $m$  and  $x \in \mathbb{R}^d$  has high dimensionality, the method is not scalable since it requires  $m^2$  eigen-decompositions of  $d \times d$  matrices. Another line of moment based methods are the so-called *sliced inverse regression* (SIR) [Li, 1991], where input  $x$  is regressed against output  $y$ . These methods project the input to a lower dimension subspace that preserves the required information. Li [1991] consider top eigen components of the moment  $\mathbb{E}[\mathbb{E}[x|y]\mathbb{E}[x|y]^\top]$  for dimensionality reduction.

## 6 Conclusion

In this paper, we propose a tensor method for efficient learning of associative mixtures. In addition to employing the learnt weight vectors in the mixture of GLMs model for prediction, we can employ them in a number of alternative ways. For instance, in practice we can utilize the output of the tensor-based methods as initializers for likelihood based techniques such as EM. In general, these techniques can get stuck in bad local optima. Initializing with the tensor methods can lead to convergence to better local optima. Moreover, we can employ the learnt weight vectors to construct discriminative features and train a different classifier using them. Thus, our method yields discriminative information which is useful in myriad ways.

There are many future directions to consider. We assume that the choice variable for selecting the mixture components is independent of the input. This is also the assumption in a number of other works for learning regression/classifier mixtures [Sun et al., 2013a, Chaganty and Liang, 2013]. In the general mixture of experts framework, the choice variable (known as the gating variable), selects the classifier based on the input. Considering this scenario is of interest. Moreover, we have considered continuous input distributions, extending this framework to discrete input is of interest.

**Acknowledgment:** This work was done while H. Sedghi was a visiting researcher at UC Irvine and was supported by NSF Career award FG15890. M. Janzamin is supported by NSF BIGDATA award FG16455. A. Anandkumar is supported in part by Microsoft Faculty Fellowship, NSF Career award CCF-1254106, and ONR Award N00014-14-1-0665.



## References

- Alekh Agarwal, Sham M Kakade, Nikos Karampatziakis, Le Song, and Gregory Valiant. Least squares revisited: Scalable approaches for multi-class prediction. In *Proc. of ICML*, 2014.
- A. Anandkumar, R. Ge, D. Hsu, and S. M. Kakade. A Tensor Spectral Approach to Learning Mixed Membership Community Models. In *COLT*, June 2013.
- A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *J. of Machine Learning Research*, 15:2773–2832, 2014a.
- Anima Anandkumar, Rong Ge, and Majid Janzamin. Sample Complexity Analysis for Learning Overcomplete Latent Variable Models through Tensor Methods. *arXiv preprint arXiv:1408.0553*, Aug. 2014b.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014c.
- Animashree Anandkumar, Rong Ge, and Majid Janzamin. Guaranteed Non-Orthogonal Tensor Decomposition via Alternating Rank-1 Updates. *arXiv preprint arXiv:1402.5180*, Feb. 2014d.
- Pranjal Awasthi, Avrim Blum, Or Sheffet, and Aravindan Vijayaraghavan. Learning mixtures of ranking models. In *Proc. of NIPS*, 2014.
- Christopher M Bishop and Markus Svensen. Bayesian hierarchical mixtures of experts. In *Proc. of Uncertainty in Artificial Intelligence*, 2003.
- Arun Chaganty and Percy Liang. Spectral experts for estimating mixtures of linear regressions. In *Proc. of The 30th International Conference on Machine Learning*, 2013.
- Yudong Chen, Xinyang Yi, and Constantine Caramanis. A convex formulation for mixed regression with two components: Minimax optimal rates. In *Conf. on Learning Theory*, 2014.
- F. Huang, U.N. Niranjan, M. Hakeem, and A. Anandkumar. Online Tensor Methods for Learning Latent Variable Models. *Accepted to JMLR*, 2014.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. In *Journal of Machine Learning Research*, pages 695–709, 2005.
- Tommi Jaakkola, David Haussler, et al. Exploiting generative models in discriminative classifiers. In *Advances in neural information processing systems*, pages 487–493, 1999.
- Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Score Function Features for Discriminative Learning: Matrix and Tensor Frameworks. *arXiv preprint arXiv:1412.2863*, Dec. 2014.
- Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.
- Sham M Kakade, Adam Kalai, Varun Kanade, and Ohad Shamir. Efficient learning of generalized linear and single index models with isotonic regression. In *NIPS*, pages 927–935, 2011.
- Nikos Karampatziakis and Paul Mineiro. Discriminative features via generalized eigenvectors. In *Proceedings of The 31st International Conference on Machine Learning*, pages 494–502, 2014.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- Ker-Chau Li. On principal hessian directions for data visualization and dimension reduction: another application of stein’s lemma. *Journal of the American Statistical Association*, 87(420):1025–1039, 1992.
- Percy Liang, Alexandre Bouchard-Côté, Dan Klein, and Ben Taskar. An end-to-end discriminative approach to machine translation. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 761–768. Association for Computational Linguistics, 2006.
- Zhang Liu and Lieven Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2009.
- Seewong Oh and Devavrat Shah. Learning mixed multinomial logit model from ordinal data. In *Proc. of NIPS*, 2014.
- K. Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society, London, A.*, page 71, 1894.
- Slav Petrov and Dan Klein. Discriminative log-linear grammars with latent variables. In *Advances in Neural Information Processing Systems*, pages 1153–1160, 2007.
- Ariadna Quattoni, Michael Collins, and Trevor Darrell. Conditional random fields for object recognition. In *Advances in neural information processing systems*, pages 1097–1104, 2004.
- Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*,