# HIERARCHICAL MIXTURES-OF-EXPERTS FOR EXPONENTIAL FAMILY REGRESSION MODELS: APPROXIMATION AND MAXIMUM LIKELIHOOD ESTIMATION 

By Wenxin Jiang and Martin A. Tanner ${ }^{1}$<br>Northwestern University


#### Abstract

We consider hierarchical mixtures-of-experts (HME) models where exponential family regression models with generalized linear mean functions of the form $\psi\left(\alpha+\mathbf{x}^{T} \boldsymbol{\beta}\right)$ are mixed. Here $\psi(\cdot)$ is the inverse link function. Suppose the true response $y$ follows an exponential family regression model with mean function belonging to a class of smooth functions of the form $\psi(h(\mathbf{x}))$ where $h(\cdot) \in W_{2 ; K_{0}}^{\infty}$ (a Sobolev class over [0, 1] ${ }^{s}$ ). It is shown that the HME probability density functions can approximate the true density, at a rate of $O\left(m^{-2 / s}\right)$ in Hellinger distance and at a rate of $O\left(m^{-4 / s}\right)$ in Kullback-Leibler divergence, where $m$ is the number of experts, and $s$ is the dimension of the predictor $\mathbf{x}$. We also provide conditions under which the mean-square error of the estimated mean response obtained from the maximum likelihood method converges to zero, as the sample size and the number of experts both increase.


1. Introduction. Both the mixtures-of-experts (ME) model, introduced by Jacobs, Jordan, Nowlan and Hinton (1991), and the hierarchical mixtures-of-experts (HME) model, introduced by Jordan and Jacobs (1994), have received considerable attention due to flexibility in modeling, appealing interpretation and the availability of convenient computational algorithms. In contrast to the single-layer ME model, the HME model has a tree-structure and can summarize the data at multiple scales of resolution due to its use of nested predictor regions. By the way they are constructed, ME and HME models are natural tools for likelihood-based inference using the expectation maximization (EM) algorithm [Jordan and Jacobs (1994) and Jordan and Xu (1995)], as well as for Bayesian analysis based on data augmentation [Peng, Jacobs and Tanner (1996)]. An introduction and application of mixing experts for generalized linear models (GLMs) are presented in Jordan and Jacobs (1994) and Peng, Jacobs and Tanner (1996).

Both ME and HME have been empirically shown to be powerful and general frameworks for examining relationships among variables in a variety of settings [Cacciatore and Nowlan (1994), Meilă and Jordan (1995), Ghahramani and Hinton (1996), Tipping and Bishop (1997) and Jaakkola and Jordan (1998)]. Despite the fact that ME and HME have been incorporated into neural network textbooks [e.g., Bishop (1995) and Haykin (1994) which features

[^0]an HME design on the cover], there has been very little formal statistical justification [see Zeevi, Meir and Maiorov (1998)] of the methodology. In this paper we consider the denseness and consistency of these models in the generalized linear model context. Before proceeding we present some notation regarding mixtures and hierarchical mixtures of generalized linear models and one-parameter exponential family regression models.

Generalized linear models are widely used in statistical practice [ $\mathrm{McCul}-$ lagh and Nelder (1989)]. One-parameter exponential family regression models [see Bickel and Doksum (1977), page 67] with generalized linear mean functions (GLM1) are special examples of the generalized linear models, where the probability distribution can be parameterized by the mean function. In the regression context, a GLM1 model proposes that the conditional expectation $\mu(\mathbf{x})$ of a real response variable $y$ (the output) is related to a vector of predictors (or inputs) $\mathbf{x} \in \mathfrak{R}^{s}$ via a generalized linear function $\mu(\mathbf{x})=\psi\left(\alpha+\boldsymbol{\beta}^{T} \mathbf{x}\right)$, with $\alpha \in \mathfrak{R}$ and $\boldsymbol{\beta} \in \mathfrak{R}^{s}$ being the regression parameters and $\psi^{-1}(\cdot)$ being a link function. The inverse link function $\psi(\cdot)$ is often used to map the entire real axis to a restricted region which contains the mean response. For example, when $y$ follows a Poisson distribution conditional on $\mathbf{x}$, a log link is often used so that the mean is nonnegative. In general, the GLM1 probability density function (pdf) of $y$ conditional on $\mathbf{x}$ is parameterized by the conditional mean $\mu(\mathbf{x})$, having the form $p(y ; \mathbf{x})=\exp \left\{a_{*}(\mu) y+b_{*}(\mu)+\right.$ $\left.c_{*}(y)\right\}$, where $\mu=\mu(\mathbf{x})=\psi\left(\alpha+\boldsymbol{\beta}^{T} \mathbf{x}\right)$, and $a_{*}(\cdot), b_{*}(\cdot)$ and $c_{*}(\cdot)$ are some fixed functions. Such models include Poisson, binomial and exponential regression models, as well as the normal and gamma regression models with dispersion parameters regarded as known. In Remark 3 (at the end of Section 3 ), we will discuss the situation when the dispersion parameter is also estimated.

A mixtures-of-experts model assumes that the total conditional density of the response is a local mixture of the conditional densities of several GLM1 experts. It is important to note that such a model differs from standard mixture models [e.g., Titterington, Smith and Makov (1985)] in that the mixing weights depend on the input. A generic expert labeled by an index $J$, proposes that the response $y$, conditional on the input $\mathbf{x}$, follows a probability distribution with density $p_{J}(y ; \mathbf{x})=\pi\left(h_{J}(\mathbf{x}), y\right)=\exp \left\{a_{*}\left(\mu_{J}\right) y+b_{*}\left(\mu_{J}\right)+c_{*}(y)\right\}$, where $\mu_{J}=\psi\left(h_{J}(\mathbf{x})\right)$ and $h_{J}(\mathbf{x})=\alpha_{J}+\boldsymbol{\beta}_{J}^{T} \mathbf{x}$. The total probability density of $y$, after combining several experts, has the form $p(y ; \mathbf{x})=\sum_{J} g_{J}(\mathbf{x}) p_{J}(y ; \mathbf{x})$, where the local weight $g_{J}(\mathbf{x})$ depends on the input $\mathbf{x}$, and is often referred to as a gating function. The total mean response then becomes $\mu(\mathbf{x})=\sum_{J} g_{J}(\mathbf{x}) \mu_{J}(\mathbf{x})$. An example of the HME model with two layers is given in Jordan and Jacobs (1994), as illustrated in Figure 1. Note that the HME is a probabilistic decision tree, where the gating networks determine the branching probabilities as a function of the covariate $\mathbf{x}$, and the experts are identified with the leaves of the decision tree. A simple mixtures-of-experts model takes the expert label $J$ to be an integer. An HME model takes $J$ as an integer vector, with dimension equal to the number of layers in the HME decision tree.

In Figure 1, adapted from Jordan and Jacobs (1994), the expert label $J$ is a two-component vector with each component taking either value 1 or 2 ; $\mathrm{g}_{i}$ and $\mathrm{g}_{j \mid i}(i, j \in\{1,2\})$ are logistic-type local weights [see Section 2.3 and (2.6)] which are identified with the probabilities of decisions at the branches of the tree. Note that the product $\mathrm{g}_{i} \mathrm{~g}_{j \mid i}$ gives a probability $g_{J}(\mathbf{x})=\mathrm{g}_{i} \mathrm{~g}_{j \mid i}$ for the path $J=(i, j)$. At the leaves of the tree, each expert $J=(i, j)$ proposes a conditional density, $p_{J}=p_{i j}$, say, with a corresponding conditional mean $\mu_{i j}$. Summing over all path $J$ 's gives the total conditional density of the conditional mixture model: $p=\sum_{J} g_{J} p_{J}$. The corresponding summation of the conditional mean functions, $\mu=\sum_{J} g_{J} \mu_{J}$, is presented in a recursive way in the original article of Jordan and Jacobs (1994), by $\mu=\sum_{i=1}^{2} \mathrm{~g}_{i} \mu_{i}$ and $\mu_{i}=\sum_{j=1}^{2} \mathrm{~g}_{j \mid i} \mu_{i j}$, and is illustrated in Figure 1.

It is demonstrated by Zeevi, Meir and Maiorov (1998) that one-layer mixtures of linear model experts can be used to approximate a class of smooth functions as the number of experts increases, and the least-squares method can be used to estimate the mean response consistently when the sample size increases. An interesting proposition is to extend this result to HME for GLM1s with nonlinear link functions and to consider the consistency of


Fig. 1. A two-layer hierarchical mixtures-of-experts model.
maximum likelihood estimation. Jiang and Tanner (1999) show that the HME for generalized linear mean functions can be used to approximate arbitrary smooth functions in a transformed Sobolev class. The present paper, in contrast, focuses on the denseness of HME density functions and on the consistency of the maximum likelihood learning. The maximum likelihood (ML) approach has two advantages over the conventional least-squares approach. (1) The maximum likelihood approach gives the smallest asymptotic variance for the estimator of the mean response, in the case of correct model specification. (2) The convenient EM algorithm can be used naturally for maximizing the likelihood, just as in the case of ordinary mixture models. However, there are two difficulties for studying the consistency properties of a likelihood-based approach. (1) The maximum likelihood method deals with density functions rather than with mean functions. A result on the denseness of mean functions, such as the ones stated in Zeevi, Meir and Maiorov (1998) and Jiang and Tanner (1999), is not enough. We need to establish a similar result for the density functions. We show that HME for GLM1 density functions can be used to approximate density functions of the form $\pi(h(\mathbf{x}), y)$, where $h(\cdot)$ is an arbitrary smooth function in a Sobolev class. (2) The maximum likelihood method minimizes the Kullback-Leibler (KL) divergence, while the consistency properties for the estimates of mean responses are usually investigated by showing that the mean square error (MSE) of the estimated mean responses converge to zero in some fashion. We need to establish a relationship between the KL divergence of the density functions and the MSE, or the $L_{2}$ distance of the mean functions.

We also note that the parameterization of the HME, as shown in the next section, is not identifiable. Care is needed for statements about the parameter estimates, which are not unique.
2. Notation and definitions. In the following, we briefly review the oneparameter exponential family regression model with generalized linear mean function (GLM1).
2.1. GLM1. We first describe the one-parameter exponential family in a way that best fits the purpose of this paper. Let $\left(A, \mathscr{F}_{A}, \lambda\right)$ be a general measure space. A probability density function $\pi(h, \cdot)$ in the one-parameter exponential family is labeled by one real parameter $h$, and has the form

$$
\begin{equation*}
\pi(h, y)=\exp \{a(h) y+b(h)+c(y)\} \quad \text { for } y \in A \tag{2.1}
\end{equation*}
$$

such that $\int_{A} \pi(h, y) d \lambda(y)=1$ for each $h \in \Re$. The functions $a(\cdot), b(\cdot)$ and $c(\cdot)$ all have known forms; $a(\cdot)$ and $b(\cdot)$ are analytic and have nonzero derivatives on $\Re$ and $c(\cdot)$ is measurable- $\mathscr{F}_{A}$.

We list some well-known properties of the one-parameter exponential models, which will be useful later.

1. The moment generating function exists in some neighborhood of the origin, and thus moments of all orders exist. See Lehmann (1991), Theorem 1.4.2, page 31 .
2. For each positive integer $k, \mu_{(k)}(h)=\int_{A} y^{k} \pi(h, y) d \lambda$ is differentiable in $h$ up to any order, due to the analyticity of $a, b$ and Theorem 1.4.1 of Lehmann (1991), page 29. In particular, we denote $\mu_{(1)}(h)=\psi(h)=\mu$ and $\mu_{(2)}(h)=v(h)$ as the first two moments.
3. The first moment can be expressed as $\mu=\psi(h) \equiv \int_{A} y \pi(h, y) d \lambda=$ $-b^{\prime}(h) / a^{\prime}(h)$ for all real $h$ and is analytic. $\psi: \mathfrak{R} \mapsto \psi(\Re)$ forms a $b^{\infty}$ diffeomorphism.
Note that the parameterization of $a(h), b(h)$ and $c(h)$ is not unique. For our purpose we require these functions to be defined on the entire real line.

Some examples follow.
Poisson. $\mathscr{P}(\mu)$ where $\mu=e^{h}, y \in A=\{0,1,2, \ldots\}$. Then

$$
\pi(h, y)=\frac{e^{-\mu}}{y!} \mu^{y}=\exp \left\{h y-e^{h}-\log (y!)\right\}
$$

Here we can take $a(h)=h, b(h)=-e^{h}, c(y)=-\log (y!)$.
Normal ( $\sigma^{2}$ known, $>0$ ). $N\left(\mu, \sigma^{2}\right.$ ) where $\mu=h, y \in A=\mathfrak{R}$. Then

$$
\begin{aligned}
\pi(h, y) & =\frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left(-\frac{1}{2 \sigma^{2}}(y-\mu)^{2}\right) \\
& =\exp \left\{\left(\frac{h}{\sigma^{2}}\right) y-\frac{h^{2}}{2 \sigma^{2}}-\frac{y^{2}}{2 \sigma^{2}}-\frac{1}{2} \log \left(2 \pi \sigma^{2}\right)\right\}
\end{aligned}
$$

Here we can take $a(h)=h / \sigma^{2}, b(h)=-h^{2} /\left(2 \sigma^{2}\right), c(y)=-y^{2} /\left(2 \sigma^{2}\right)-(1 / 2) \log$ ( $2 \pi \sigma^{2}$ ).

The GLM1 assumes that $h=\alpha+\boldsymbol{\beta}^{T} \mathbf{x}$, which introduces the dependence of $y$ on an $s$-dimensional predictor $\mathbf{x}$ through the density function $\pi(h, y)$. In this context, the inverse of $\psi(\cdot)$ is called a link function [McCullagh and Nelder (1989)]. For a specific probability model of $y$ (say, Poisson), there could be different choices of the link function. Our paper does not restrict the choice of the link function, as long as it is "smooth" on $\Re$ and invertible. Note that the functions $a, b$ and $c$ in (2.1) correspond, respectively, to the functions $a_{*} \circ \psi$, $b_{*} \circ \psi$ and $c_{*}$ in notation of Section 1, where $\circ$ stands for composition.

Now we introduce a target family of regression models which is more flexible than the family of GLM1 by allowing $h(\cdot)$ to be an arbitrary smooth function (of $\mathbf{x}$ ) in a Sobolev class.
2.2. The family of target functions. Let $\Omega=[0,1]^{s}=\bigotimes_{q=1}^{s}[0,1]$, the space of the predictor $\mathbf{x}$, where $\otimes$ stands for the direct product. Let $A \subset \Re$ be the space of the response $y$. Let $\left(A, \mathscr{F}_{A}, \lambda\right)$ be a general measure space, $\left(\Omega, \mathscr{F}_{\Omega}, \kappa\right)$ be a probability space such that $\kappa$ has a positive continuous density with respect to the Lebesgue measure on $\Omega$ and $\left(\Omega \otimes A, \mathscr{T}_{\Omega} \otimes \mathscr{F}_{A}, \kappa \otimes \lambda\right)$ be the product measure space. Consider a random predictor-response pair $\left(\mathbf{X}_{(s \times 1)}, Y_{(1 \times 1)}\right)$. Suppose $\mathbf{X}$ has a probability measure $\kappa$, and $(\mathbf{X}, Y)$ has a probability density function (pdf) $\varphi$ with respect to $\kappa \otimes \lambda$, where $\varphi$ is a target function of the form

$$
\begin{equation*}
\varphi(\mathbf{x}, y)=\pi(h(\mathbf{x}), y) \tag{2.2}
\end{equation*}
$$

Here $\pi(\cdot, \cdot): \Re \otimes A \mapsto \Re$ has the one-parameter exponential form (2.1). In contrast to a GLM1 model, we allow a more flexible $h(\mathbf{x})$ in (2.2). Here $h: \Omega \mapsto$ $\mathfrak{R}$ is assumed to have continuous second derivatives, $\sum_{\mathbf{k}: 0 \leq|\mathbf{k}| \leq 2}\left\|D^{\mathbf{k}} h\right\|_{\infty} \leq$ $K_{0}$, where $\mathbf{k}=\left(k_{1}, \ldots, k_{s}\right)$ is an $s$-dimensional vector of nonnegative integers between 0 and $2,|\mathbf{k}|=\sum_{j=1}^{s} k_{j},\|h\|_{\infty} \equiv \sup _{\mathbf{x} \in \Omega}|h(\mathbf{x})|$ and $D^{\mathbf{k}} h \equiv$ $\left(\partial^{|\mathbf{k}|} h / \partial x_{1}^{k_{1}} \cdots \partial x_{s}^{k_{s}}\right)$. In other words, $h \in W_{2 ; K_{0}}^{\infty}$, where $W_{2 ; K_{0}}^{\infty}$ is a ball with radius $K_{0}$ in a Sobolev space with sup-norm and second-order continuous differentiability. The conditional mean function $\mu(\cdot)$, corresponding to $\varphi(\cdot, \cdot)$, is obviously

$$
\begin{equation*}
\mu(\mathbf{x})=\int_{A} y \varphi(\mathbf{x}, y) d \lambda(y)=\psi(h(\mathbf{x})) \tag{2.3}
\end{equation*}
$$

for all $\mathbf{x}$ in $\Omega$. Sobolev classes of mean functions similar to $W_{2 ; K_{0}}^{\infty}$ are also considered in Mhaskar (1996) and Zeevi, Meir and Maiorov (1998). Our family of mean functions is a transformed class $\psi\left(W_{2 ; K_{0}}^{\infty}\right)$, where $\psi^{-1}$ is the link function.

We have restricted the predictor $\mathbf{x}$ to $\Omega=[0,1]^{s}$ to simplify the exposition. The theorems of this paper actually hold for $\Omega$ being any compact subset of $\Re^{s}$. The compactness of $\Omega$ is needed in the techniques of our proof. We also note that in the situation when $\Omega$ is the direct product of $s$ closed intervals, suitable recentering and rescaling of each of the $s$ components of $\mathbf{x}$ can transform $\Omega$ into $[0,1]^{s}$.

Denote the set of all pdfs $\varphi(\cdot, \cdot)=\pi(h(\cdot), \cdot)$ defined this way as $\boldsymbol{\Phi}$. This is the set of target functions that we will consider to approximate.

Now we define the hierarchical mixtures-of-experts (HME) for GLM1s. They are the functions which we use for approximating a function in $\boldsymbol{\Phi}$.
2.3. The family of HME of GLM1s. An approximator $f$ in the HME family is assumed to have the following form:

$$
\begin{equation*}
f=f_{\Lambda}(\mathbf{x}, y ; \theta)=\sum_{J \in \Lambda} g_{J}(\mathbf{x} ; \mathbf{v}) \pi\left(h_{J}(\mathbf{x}), y\right) \tag{2.4}
\end{equation*}
$$

where $h_{J}(\mathbf{x})=\alpha_{J}+\boldsymbol{\beta}_{J}^{T} \mathbf{x}$, and $\pi(\cdot, \cdot)$ is as defined in Section 2.1. The parameters of this model include $\alpha_{J} \in \Re$ and $\boldsymbol{\beta}_{J} \in \mathfrak{R}^{s}$, as well as $\mathbf{v}$ which is some parameter for the gating function $g_{J}$ 's. We use the symbol $\theta$ to represent the grand vector of parameters containing all the components of the parameters $\mathbf{v}, \alpha_{J}$ and $\boldsymbol{\beta}_{J}$ for all $J \in \Lambda$. In (2.4), $\Lambda$ is the set of labels of all the experts in a network, referred to as a structure. Two quantities are associated with a structure: the dimension $l=\operatorname{dim}(\Lambda)$, which is the number of layers and the cardinality $m=\operatorname{card}(\Lambda)$, which is the number of experts. A HME of $l$-layers has a structure of the form $\Lambda=\bigotimes_{k=1}^{l} A_{k}$ where $A_{k}=\left\{1, \ldots, w_{k}\right\}, w_{k} \in \mathscr{N}$ and $k=1, \ldots, l$. (We use $\mathscr{N}$ to denote the set of all positive integers.) Graphically, $w_{k}=\operatorname{card}\left(A_{k}\right)$ represents the number of expert branches or the number of "splits" at layer $k, k=1, \ldots, l$. Note that in this paper we restrict attention to "rectangular-shaped" structures (corresponding to balanced trees). A generic
expert label $J$ in $\Lambda$ can then be expressed as $J=\left(j_{1}, \ldots, j_{l}\right)$ where $j_{k} \in A_{k}$ for each $k$.

To characterize a structure $\Lambda$, we often claim that it belongs to a certain set of structures. We now introduce three such sets of structures, $\mathscr{J}, \mathscr{J}_{\mathrm{m}}$ and $\mathcal{S}$, which will be used later when formulating the results. The set of all possible HME structures under consideration is $\mathcal{J}=\left\{\Lambda: \Lambda=\bigotimes_{k=1}^{l}\left\{1, \ldots, w_{k}\right\}\right.$; $\left.w_{k} \in \mathscr{N} ; k=1, \ldots, l ; l \in \mathscr{N}\right\}$. The set of all HME structures containing no more than $m$ experts is denoted as $\mathscr{J}_{m}=\{\Lambda: \Lambda \in \mathscr{L}, \operatorname{card}(\Lambda) \leq m\}$. We also introduce a symbol $\mathscr{\rho}$ to denote a generic subset of $\mathscr{J}$. This is introduced in order to formulate a major condition for some results of this paper to hold. This condition, to be formulated in the next section, will be specific to a generic subset $\mathscr{\Omega}$ of HME structures. A trivial example of $\mathscr{\Omega}$ is $\mathscr{J}$. Another example of $\mathscr{\Omega}$ is $\mathscr{I}_{L}=\{\Lambda: \Lambda \in \mathscr{L}, \operatorname{dim}(\Lambda) \leq L\}$, which includes all structures with $L$ or less layers. In particular, $\mathscr{S}_{1}$ represents the set of single-layer structures. A third example of $\mathscr{\rho}$ is $\mathscr{\mathscr { C }}_{B}=\left\{\Lambda: \Lambda=\bigotimes_{k=1}^{l}\{1,2\} ; l \in \mathscr{N}\right\}$, which represent the set of trees with binary splits.

Associated with a structure $\Lambda$ is a family of vectors of gating functions. Each member is called a gating vector and is labeled by a parameter vector $\mathbf{v} \in V_{\Lambda}, V_{\Lambda}$ being some parameter space specific to the structure $\Lambda$. Denote a generic gating vector as $G_{\mathbf{v}, \Lambda} \equiv\left(g_{J}(\cdot ; \mathbf{v})\right)_{J \in \Lambda}$. We assume the $g_{J}(\mathbf{x} ; \mathbf{v})$ 's to be nonnegative, with sum equal to unity, and continuous in $\mathbf{x}$ and $\mathbf{v}$. Note that $\int_{A} f_{\Lambda}(\mathbf{x}, y ; \theta) d \lambda(y)=1$ is ensured. Let $\mathscr{G}=\left\{G_{\mathbf{v}, \Lambda}: \mathbf{v} \in V_{\Lambda}, \Lambda \in \mathscr{\rho}\right\}$ be the family of gating vectors defined on the set of structures $\mathscr{\rho}$, which will be referred to as a gating class defined on $\mathscr{\rho}$.

In the following we define the logistic gating class $\mathscr{G}=\mathscr{L}$ on the set of all structures $\mathscr{\mathscr { L }}$. This class has been commonly used in literature [see Jordan and Jacobs (1994)]. Here, for each structure $\Lambda$ in $\mathscr{J}$ and each label $J$ in $\Lambda$, a gating function $g_{J}=g_{J}(\cdot, \mathbf{v})$ is defined recursively. Suppose $J$ is an $l$-dimensional integer $\left(j_{1}, j_{2}, \ldots, j_{l}\right)$. Then,

$$
\begin{equation*}
g_{J} \equiv g_{j_{1} j_{2} \ldots j_{l}}=\mathrm{g}_{j_{1}} \mathrm{~g}_{j_{2} \mid j_{1}} \ldots \mathrm{~g}_{j_{l} \mid j_{1} j_{2} \ldots j_{l-1}} . \tag{2.5}
\end{equation*}
$$

Here, for each $q$, the factor $\mathrm{g}_{j_{q} \mid j_{1} \ldots j_{q-1}}$ takes a multinomial logit form,

$$
\begin{equation*}
\mathrm{g}_{j_{q} \mid j_{1} \ldots j_{q-1}}=\frac{\exp \left(\xi_{j_{q} \mid j_{1} \ldots j_{q-1}}\right)}{\sum_{k=1}^{w_{q}} \exp \left(\xi_{k \mid j_{1} \ldots j_{q-1}}\right)} \tag{2.6}
\end{equation*}
$$

where $\xi_{k \mid j_{1} \ldots j_{q-1}}=\phi_{k \mid j_{1} \ldots j_{q-1}}+\gamma_{k \mid j_{1} \ldots j_{q-1}}^{T} \mathbf{x},\left(\phi_{k \mid j_{1} \ldots j_{q-1}}, \gamma_{k \mid j_{1} \ldots j_{q-1}}^{T}\right) \in \mathfrak{R}^{s+1}, k=$ $1, \ldots, w_{q}$. Usually it is assumed that

$$
\phi_{w_{q} \mid j_{1} \ldots j_{q-1}}=\gamma_{w_{q} \mid j_{1} \ldots j_{q-1}}=\xi_{w_{q} \mid j_{1} \ldots j_{q-1}}=0,
$$

since otherwise a "translation" of the parameters,

$$
\begin{aligned}
\phi_{k \mid j_{1} \ldots j_{q-1}} & \rightarrow \phi_{k \mid j_{1} \ldots j_{q-1}}+\phi_{0} \\
\gamma_{k \mid j_{1} \ldots j_{q-1}} & \rightarrow \gamma_{k \mid j_{1} \ldots j_{q-1}}+\gamma_{0} \quad \text { all } k=1, \ldots, w_{q}
\end{aligned}
$$

would leave the probability density function $f_{\Lambda}(\mathbf{x}, y ; \theta)$ unchanged. Note that the grand vector of "gating parameters" $\mathbf{v}$ includes all components of $\left(\phi_{j_{q} \mid j_{1} \ldots j_{q-1}}, \gamma_{j_{q} \mid j_{1} \ldots j_{q-1}}^{T}\right)$, where $\left(j_{1}, \ldots, j_{q-1}\right) \in \bigotimes_{r=1}^{q-1}\left\{1, \ldots, w_{r}\right\}$ and $j_{q} \in\{1$, $\left.\ldots, w_{q}-1\right\}$ for all $q=1, \ldots, l$. It is easy to see that

$$
\begin{aligned}
\operatorname{dim}(\mathbf{v}) & =(s+1)\left\{\left(w_{1}-1\right)+w_{1}\left(w_{2}-1\right)+\ldots+w_{1} \ldots w_{l-1}\left(w_{l}-1\right)\right\} \\
& =(s+1)(m-1)
\end{aligned}
$$

and the parameter space $V_{\Lambda}$ for $\mathbf{v}$ is $\mathfrak{R}^{(s+1)(m-1)}$, where $m=w_{1} \cdots w_{l}=$ $\operatorname{card}(\Lambda)$. Note that the gating functions constructed in this way are analytic for $\left(\mathbf{v}^{T}, \mathbf{x}^{T}\right) \in \mathfrak{R}^{(s+1)(m-1)} \otimes \mathfrak{R}^{s}$. The space of regression parameters (or "expert parameters") $\left(\alpha_{J}, \boldsymbol{\beta}_{J}^{T}\right)$ 's, corresponding to structure $\Lambda$, is $\mathfrak{R}^{(s+1) m}$. The space of grand parameter $\theta$ 's, corresponding to structure $\Lambda$, is $\tilde{\Theta}_{\Lambda}=\Re^{(s+1)(2 m-1)}$. Here the $(2 m-1)(s+1)$-dimensional grand parameter $\theta$ includes all components of the gating parameters from $\mathbf{v}$ and the expert parameters from $\left(\alpha_{J}, \boldsymbol{\beta}_{J}^{T}\right)_{J \in \Lambda}$.

Now we are ready to define the family of approximator functions. Let $\Pi_{\Lambda}$ be the set of all function $f_{\Lambda}$ 's of the form (2.4), specific to a structure $\Lambda$, which can be denoted as $\Pi_{\Lambda}=\left\{f_{\Lambda}(\cdot, \cdot ; \theta): \theta \in \tilde{\Theta}_{\Lambda}\right\}$. This set $\Pi_{\Lambda}$ is the set of HME functions from which an optimal function is chosen by the maximum likelihood method to approximate the truth. It is assumed that a structure $\Lambda$ is chosen a priori. In practice, people often analyze data using different choices of structures and select the best fitting model [see Fritsch, Finke and Waibel (1997) for an adaptive approach]. We consider in this paper choosing among the set of structures $\mathscr{J}_{m} \cap \mathscr{\rho}$. Denote

$$
\begin{equation*}
\boldsymbol{\Pi}_{m, \mathscr{A}}=\left\{f: f \in \Pi_{\Lambda} ; \Lambda \in \mathscr{J}_{m} \cap \mathscr{\rho}\right\} \tag{2.7}
\end{equation*}
$$

This set, $\Pi_{m, \mathscr{L}}$, is the family of HME functions for which we examine the approximation rate in $\boldsymbol{\Phi}$, as $m \rightarrow \infty$. Note that this family of HME functions is specific to $m$, the maximum number of experts, as well as to some subset $\mathscr{\mathscr { L }}$ of HME structures, which will be specified later. We do not explicitly require that $\Pi_{m, \mathscr{\rho}}$ be a subset of $\boldsymbol{\Phi}$ in this paper.

Each HME density function $f_{\Lambda}(\mathbf{x}, y ; \theta)$ generates a mean function $\mu_{\Lambda}(\mathbf{x} ; \theta)$ by

$$
\begin{equation*}
\mu_{\Lambda}(\mathbf{x} ; \theta)=\int_{A} y f_{\Lambda}(\mathbf{x}, y ; \theta) d \lambda(y)=\sum_{J \in \Lambda} g_{J}(\mathbf{x} ; \mathbf{v}) \psi\left(\alpha_{J}+\mathbf{x}^{T} \boldsymbol{\beta}_{J}\right) \tag{2.8}
\end{equation*}
$$

where $\psi(\cdot)=\int_{A} y \pi(\cdot, y) d \lambda(y)$.
The parameterization of the HME functions is not identifiable, in the sense that two different parameters $\theta$ in $\tilde{\Theta}_{\Lambda}$ can represent the same density function $f$ in $\Pi_{m, \rho}$. For example, the density functions are invariant under permutation of the expert label $J$ 's. Also, if two experts $J$ and $J^{\prime}$ propose the same output, that is, if $\alpha_{J}=\alpha_{J^{\prime}}$ and $\boldsymbol{\beta}_{J}=\boldsymbol{\beta}_{J^{\prime}}$, then the mixing proportions for these two experts can be arbitrary, as long as the sum of the two weights are unchanged. This can lead to the nonidentifiability of some components of parameter $\mathbf{v}$. Our description of the estimation procedure and the statement of
the results will take these identifiability issues into account. The identifiability issues also suggest that it makes more sense to formulate the consistency problem in terms of the estimated mean response, rather than to look at the consistency of the parameter estimates.
2.4. The method of estimation. We will use the maximum likelihood method to train the architecture. Suppose we estimate the mean response $\mu(\mathbf{x})$ based on a data set of $n$ predictor-response pairs $\left(\mathbf{X}_{i}, Y_{i}\right), \mathbf{X}_{i} \in \Omega, Y_{i} \in A, i=$ $1, \ldots, n$. Let the measure spaces $\left(\Omega, \mathscr{F}_{\Omega}, \kappa\right)$ and $\left(A, \mathscr{F}_{A}, \lambda\right)$ be as introduced in Section 2.2. Assume that $\left(\mathbf{X}_{i}, Y_{i}\right), i=1, \ldots, n$ are independent and identically distributed (i.i.d.) random vectors. The probability measure for $\mathbf{X}_{i}$ is $\kappa$. The probability measure of $Y_{i}$ conditional on $\mathbf{X}_{i}=\mathbf{x}$ has a density $\varphi(\mathbf{x}, \cdot)$ [defined in (2.2)] with respect to the measure $\lambda$, for all $\mathbf{x} \in \Omega$.

The log-likelihood function based on the HME model is

$$
\begin{equation*}
\mathrm{Ł}_{n, \Lambda}(\theta ; \omega)=n^{-1} \sum_{i=1}^{n} \log \left\{f_{\Lambda}\left(\mathbf{X}_{i}, Y_{i} ; \theta\right) / \varphi_{0}\left(\mathbf{X}_{i}, Y_{i}\right)\right\} \tag{2.9}
\end{equation*}
$$

where $f_{\Lambda}(\cdot, \cdot ; \theta) \in \Pi_{\Lambda}$ is defined in Section 2.3, $\theta \in \tilde{\Theta}_{\Lambda}$, $\omega$ is the stochastic sequence of events $\left(\mathbf{X}_{i}, Y_{i}\right), i=1, \ldots$ and $\varphi_{0}\left(\mathbf{X}_{i}, Y_{i}\right)$ can be any positive measurable function of the observed data that does not depend on the parameter $\theta$. In this paper we choose $\varphi_{0}\left(\mathbf{X}_{i}, Y_{i}\right)=\exp \left(c\left(Y_{i}\right)\right)$ where $c(\cdot)$ is from (2.1). It turns out that such a choice makes the log-likelihood function uniformly convergent to its expectation, for almost all $\omega$, in any compact subset of parameters, as $n \rightarrow \infty$. Define the maximum likelihood estimator (MLE) $\hat{\theta}_{n, \Lambda}(\omega)$ to be a maximizer (can be one out of many) of $£_{n, \Lambda}(\theta ; \omega)$ over a compact set $\tilde{B}_{\Lambda} \subset \tilde{\Theta}_{\Lambda}$, that is,

$$
\begin{equation*}
\hat{\theta}_{n, \Lambda}(\omega)=\arg \max _{\theta \in \tilde{B}_{\Lambda}}\left\{\succeq_{n, \Lambda}(\theta ; \omega)\right\} . \tag{2.10}
\end{equation*}
$$

The maximum likelihood method, in the large sample size limit, essentially searches for $\theta$ which minimizes the $\operatorname{KL}$ divergence $\operatorname{KL}\left(f_{\Lambda}, \varphi\right)$ between $f_{\Lambda}=$ $f_{\Lambda}(\cdot, \cdot ; \theta) \in \Pi_{\Lambda}$ and $\varphi=\varphi(\cdot, \cdot) \in \boldsymbol{\Phi}$, where

$$
\begin{equation*}
\mathrm{KL}(f, g) \equiv \iint_{\Omega \otimes A} g(\mathbf{x}, y) \log \left\{\frac{g(\mathbf{x}, y)}{f(\mathbf{x}, y)}\right\} d \kappa(\mathbf{x}) d \lambda(y) \tag{2.11}
\end{equation*}
$$

It turns out that the KL divergence $\operatorname{KL}\left(f_{\Lambda}, \varphi\right)$ is always well defined (see Corollary 1 in Section 2.5). Due to the nonidentifiability of the parameterization, there is $a$ set of $\theta$ 's in $\tilde{B}_{\Lambda}$ that minimize the KL divergence. Denote this set as $\Theta_{\Lambda}$, which could be expressed as

$$
\begin{equation*}
\Theta_{\Lambda}=\left\{\theta \in \tilde{B}_{\Lambda}: \theta=\arg \min _{\theta^{*} \in \tilde{B}_{\Lambda}} \operatorname{KL}\left(f_{\Lambda}\left(\cdot, \cdot ; \theta^{*}\right), \varphi\right)\right\} . \tag{2.12}
\end{equation*}
$$

Based on any MLE $\hat{\theta}_{n, \Lambda}=\hat{\theta}_{n, \Lambda}(\omega)$, an estimated mean response can be constructed as $\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}\right)$. We do not explicitly require that for two different
global MLEs the estimated mean responses be the same. The MSE of an estimated mean response is defined by

$$
\begin{equation*}
(\mathrm{MSE})_{n, \Lambda}=\mathrm{E} \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}\right)-\mu(\mathbf{x})\right\}^{2} d \kappa(\mathbf{x}) \tag{2.13}
\end{equation*}
$$

where E is the expectation taken on the MLE $\hat{\theta}_{n, \Lambda}$, and $\mu_{\Lambda}$ and $\mu$ are defined in (2.8) and (2.3), respectively.
2.5. Technical definitions. Some technical definitions are introduced below. We will use these definitions to formulate a major condition for our results to hold.

Definition 1 (Fine partition). For $\nu=1,2, \ldots$, let $\mathbf{Q}^{(\nu)}=\left\{Q_{J}^{(\nu)}\right\}_{J \in \Lambda^{(\nu)}}$, $\Lambda^{(\nu)} \in \mathcal{J}$, be a partition of $\Omega \subset \mathfrak{R}^{s}$. (This means that for fixed $\nu$, the $Q_{J}^{(\nu)}$, s are mutually disjoint subsets of $\mathfrak{R}^{s}$ whose union is $\Omega$.) Let $p_{\nu}=\operatorname{card}\left(\Lambda^{(\nu)}\right)$, ( $p_{\nu} \in \mathscr{N}$ ).

If $p_{\nu} \rightarrow \infty$, and if for all $\boldsymbol{\xi}, \boldsymbol{\eta} \in Q_{J}^{(\nu)}, \rho(\boldsymbol{\xi}, \boldsymbol{\eta}) \equiv \max _{1 \leq q \leq s}\left|(\boldsymbol{\xi}-\boldsymbol{\eta})_{q}\right| \leq c_{0} / p_{\nu}^{1 / s}$ for some constant $c_{0}$ independent of $\nu, J, \boldsymbol{\xi}, \boldsymbol{\eta}$, then $\left\{\mathbf{Q}^{(\nu)}: \nu=1,2, \ldots\right\}$ is called a sequence of fine partitions with structure sequence $\left\{\Lambda^{(\nu)}\right\}$, cardinality sequence $\left\{p_{\nu}\right\}$ and bounding constant $c_{0}$.

Definition 2 (Subgeometric). A sequence $\left\{a_{\nu}\right\}$ is subgeometric with rate bounded by $M_{2}$, if $a_{\nu} \in \mathscr{N}, a_{\nu} \rightarrow \infty$ as $\nu \rightarrow \infty$ and $1<\left|a_{\nu+1} / a_{\nu}\right|<M_{2}$ for all $\nu=1,2, \ldots$, for some finite constant $M_{2}$.

In the following we introduce some measures of the discrepancy between a $\operatorname{pdf} f$ in $\Pi_{\Lambda}$ [of the form (2.4)] and a $\operatorname{pdf} \varphi$ in $\boldsymbol{\Phi}$ [of the form (2.2)]. One of them is the $\operatorname{KL}$ distance $\operatorname{KL}(f, \varphi)$ [see (2.11)]. Another is the Hellinger distance

$$
\begin{equation*}
d_{H}(f, \varphi)=\left\{\iint(\sqrt{f}-\sqrt{\varphi})^{2} d \lambda d \kappa\right\}^{1 / 2} \tag{2.14}
\end{equation*}
$$

This is a true distance and is invariant under rescaling of the measures $\lambda$ and $\kappa$ [see Devroye and Gyoerfi (1985)]. A third description is the $L_{2}$ distance between the means,

$$
\begin{equation*}
d_{2}\left(\mu_{f}, \mu_{\varphi}\right)=\left\|\mu_{\Lambda}(\cdot, \theta)-\mu(\cdot)\right\|_{2, \kappa}=\left\{\int\left(\mu_{f}-\mu_{\varphi}\right)^{2} d \kappa\right\}^{1 / 2}, \tag{2.15}
\end{equation*}
$$

where $\mu_{f}=\int y f d \lambda$ and $\mu_{\varphi}=\int y \varphi d \lambda$, for $f$ in $\Pi_{\Lambda}$ and $\varphi$ in $\Phi$. This measure is used since it is closely related to the MSE defined in Section 2.4.

For technical convenience, we introduce a fourth measure of discrepancy between $f$ in $\Pi_{\Lambda}$ and $\varphi$ in $\boldsymbol{\Phi}$ called the upper divergence. For $f=\sum_{J \in \Lambda} g_{J}$. $\pi\left(h_{J}, y\right)$ and $\varphi=\pi(h, y)$, the upper divergence is defined as

$$
\begin{equation*}
\mathscr{Q}(f, \varphi)=\int \sum_{J \in \Lambda} g_{J}(\mathbf{x} ; \mathbf{v})\left\{h_{J}(\mathbf{x})-h(\mathbf{x})\right\}^{2} d \kappa, \tag{2.16}
\end{equation*}
$$

where $h_{J}(\mathbf{x})=\alpha_{J}+\boldsymbol{\beta}_{J}^{T} \mathbf{x}$. Note that the idea of HME approximation is to partition the input space "softly" according to the $g_{J}$ 's and use a linear function $h_{J}(\mathbf{x})$ to approximate $h(\mathbf{x})$ in each partition, so as to approximate the (conditional) pdf $\pi(h(\mathbf{x}), \cdot)$ for all $\mathbf{x}$. The upper divergence measures the quality of this softly partitioned linear approximation. The name "upper divergence" is due to the following lemma, which implies that $\mathscr{Q}$ is stronger than the other divergence measures, that is, KL, $d_{H}$ and $d_{2}\left(\mu_{f}, \mu_{\varphi}\right)$ (the proof is in Section 4).

Lemma 1 (Strength of divergence measures). Consider any structure $\Lambda$, any HME density $f=\sum_{J \in \Lambda} g_{J} \pi\left(h_{J}, \cdot\right)$ in $\Pi_{\Lambda}$ and any target density $\varphi=\pi(h, \cdot)$ in $\boldsymbol{\Phi}$. We have:
(a) $d_{2}^{2}\left(\mu_{f}, \mu_{\varphi}\right) \leq 4 M_{I} d_{H}^{2}(f, \varphi)$;
(b) $d_{H}^{2}(f, \varphi) \leq K L(f, \varphi)$. [This lemma appeared in, e.g., Haussler and Opper (1995) and Zeevi and Meir (1997).]
(c) $\mathrm{KL}(f, \varphi) \leq M_{I I} \mathscr{Q}(f, \varphi)$.

Here, $M_{I}=\sup _{|h| \leq K}\left\{\int y^{2} \pi(h, y) d \lambda\right\}$, and

$$
M_{I I}=\frac{1}{2}\left\{\sup _{|h| \leq K}\left|\int y \pi(h, y) d \lambda\right| \sup _{|h| \leq K}\left|a^{\prime \prime}(h)\right|+\sup _{|h| \leq K}\left|b^{\prime \prime}(h)\right|\right\},
$$

where $a(\cdot)$ and $b(\cdot)$ are defined as in (2.1), and $K$ is an upperbound for $\left|h_{J}(\mathbf{x})\right|$ and $|h(\mathbf{x})|$ for all $\mathbf{x}$ in $\Omega$ and all $J$ in $\Lambda$.

REMARK 1. $M_{I}$ and $M_{I I}$ are finite constants, due to the continuity of $\alpha^{\prime \prime}$, $b^{\prime \prime}$ and $\int y^{k} \pi(h, y) d \lambda(k=1,2)$, as functions of $h$.

Corollary 1. All the divergence measures $d_{2}\left(\mu_{f}, \mu_{\varphi}\right), d_{H}, \mathrm{KL}$ and $\mathscr{Q}$ are finite.

The proof is obvious from Lemma 1 . Note that $\mathscr{Q}$ is finite, since it involves an integration of a continuous function over the compact space $\Omega$ of input $\mathbf{x}$.

In Section 4, we will use Lemma 1(c) on upper divergence to prove the denseness property of the HME densities in KL divergence.
3. Results and conditions. In the following, we state some regularity conditions, as well as some results which hold under these conditions.

Condition 1. $\left(A_{\mathscr{J}, p}\right)$. For a subset $\mathscr{\Omega} \subset \mathscr{J}$, there is a fine partition sequence $\left\{\left\{Q_{J}^{(\nu)}\right\}_{J \in \Lambda_{0}^{(\nu)}} \Lambda_{0}^{(\nu)} \in \mathscr{\rho}, \nu=1,2, \ldots\right\}$ with a bounding constant $c_{0}$ and a cardinality sequence $\left\{p_{\nu}: \nu=1,2, \ldots\right\}$, such that $\left\{p_{\nu}^{1 / s}\right\}$ is subgeometric with rate bounded by a constant $M_{2}$, and for all $\nu$, for all $\varepsilon>0$, there exists $\mathbf{v}_{\varepsilon} \in V_{\Lambda_{0}^{(\nu)}}$ and a gating vector

$$
\begin{gather*}
G_{\mathbf{v}_{\varepsilon}, \Lambda_{0}^{(\nu)}}=\left\{g_{J}\left(\mathbf{x} ; \mathbf{v}_{\varepsilon}\right)\right\}_{J \in \Lambda_{0}^{(\nu)}} \in \mathscr{G}, \\
\Lambda_{0}^{(\nu)} \in \mathscr{\rho}, \text { such that } \sup _{J \in \Lambda_{0}^{(\nu)}} \| g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)-\chi_{Q_{J}^{(\nu)}(\cdot) \|_{p, \sigma} \leq \varepsilon .} \tag{3.1}
\end{gather*}
$$

Here, $\|f(\cdot)\|_{p, \sigma} \equiv\left\{\int_{\Omega}|f(\mathbf{x})|^{p} d \sigma(\mathbf{x})\right\}^{1 / p}$, where $p \in \mathscr{N} ; \sigma$ is any probability measure on $\Omega$ which has a positive continuous density with respect to the Lebesgue measure; $\chi_{B}(\cdot)$ is the characteristic function for a subset $B$ of $\Omega$, that is, $\chi_{B}(\mathbf{x})=1$ if $\mathbf{x} \in B, 0$ otherwise.

This condition is a restriction on the gating class $\mathscr{G}$ defined on a set of structures $\mathscr{\rho}$. Loosely speaking, it indicates that the vectors of local gating functions in the parametric family should arbitrarily approximate the vector of characteristic functions for a partition of the predictor space $\Omega$, as the cells of the partition become finer.

Theorem 1 (Approximation rate in Hellinger distance). Under Condition $A_{\mathscr{S}, 1}$,

$$
\sup _{\varphi \in \Phi} \inf _{f \in \boldsymbol{\Pi}_{m, \iota}} d_{H}(f, \varphi) \leq \frac{c}{m^{2 / s}}
$$

for some positive constant c independent of $m$. Here $d_{H}$ is the Hellinger distance defined in (2.14).

REMARK 2. This theorem implies the same approximation rate in the $L_{1}$ distance $d_{1}(f, \varphi)=\iint|f-\varphi| d \lambda d \kappa$, by the well-known result $d_{1}(f, \varphi) \leq$ $2 d_{H}(f, \varphi)$ [see Devroye and Gyoerfi (1985)]. The same rate in general $L_{p}$ distances ( $p \geq 2$ ) is derived in Jiang and Tanner (1998) (under Condition $A_{\mathscr{S}, p}$ ), where an extra boundedness condition is required.

Theorem 2 (Approximation rate in KL divergence). Under Condition $A_{\mathscr{\Omega}, 1}$,

$$
\sup _{\varphi \in \Phi} \inf _{f \in \boldsymbol{\Pi}_{m, \jmath}} \operatorname{KL}(f, \varphi) \leq c / m^{4 / s}
$$

for some positive constant $c$ independent of $m$. Here KL is the KL divergence defined in (2.11).

All these results depend on a major condition $A_{\mathscr{A}, 1}$. The following remark claims that it is satisfied by certain gating functions defined on certain structures.

Remarks. (a) Condition $A_{\mathscr{S}, p}$ is satisfied (for any $p \in \mathscr{N}$ ) by the logistic gating class $\mathscr{G}=\mathscr{L}$ defined on the set of structures $\mathscr{\mathscr { L }}=\mathscr{\mathscr { L }}_{B}$ for trees with binary splits (Section 2.3). This is because, roughly speaking, a logistic function from a binary split has the form $\left(1+\exp \left(-\beta\left(z-z_{0}\right)\right)\right)^{-1}$, which can approximate a step function $H\left(z-z_{0}\right)$ as $\beta$ increases, for any location of jump $z_{0}$. The gating functions in a binary tree involve products of the logistic functions (and their complements), which can approximate products of step functions which form the characteristic functions of a fine partition. In this way, Condition $A_{\mathscr{S}, p}$ can be proved. This implies that the approximation rates in the theorems stated above apply to HME of GLM1s with binary trees.
(b) Jiang and Tanner (1999), Lemma, Section 3, show that Condition $A_{\mathscr{S}, p}$ is satisfied (for any $p \in \mathscr{N}$ ) by the logistic gating class $\mathscr{L}$ defined on $\mathscr{J}_{s}$, which is the set of structures with no more than $s$ layers, where $s=\operatorname{dim}(\mathbf{x})$. Moreover, due to Jiang and Tanner (1999), Remark (i), Section 5, Condition $A_{\mathscr{S}, p}$ is also satisfied by the logistic gating class defined on the set of single-layer structures $\mathscr{I}_{1}$ (corresponding to the MEs), which implies that the approximation rates in the above theorems apply to ME of GLM1s.
(c) Another class of gating functions can be defined only on the binary trees (in $\mathscr{\mathscr { S }}_{B}$ ). There, the logistic gating functions in (2.6) are replaced by continuous cumulative distribution functions (cdfs). One example is to use the normal cdf. Then the gating factor $\mathrm{g}_{j_{q} \mid j_{1} \cdots j_{q-1}}$ of (2.6) becomes $\Phi\left(\xi_{j_{1} \cdots j_{q-1}}\right)$ if $j_{q}=1$, or $1-\Phi\left(\xi_{j_{1} \cdots j_{q-1}}\right)$ if $j_{q}=2$; where $\xi_{j_{1} \cdots j_{q-1}}=\phi_{j_{1} \ldots j_{q-1}}+\gamma_{j_{1} \ldots j_{q-1}}^{T} \mathbf{x}$. A similar argument as in part (a) of this remark shows that Condition $A_{\mathscr{S}, p}$ is satisfied for this new gating class for any $p \in \mathscr{N}$.

The next condition is useful for proving the consistency of the maximum likelihood (ML) learning method.

Condition 2 (Scope of maximum likelihood searching). The scope of the maximum likelihood (ML) searching, $\tilde{B}_{\Lambda}$, is a compact subset of $[-C, C]^{\operatorname{dim}(\theta)}$ for some large positive constant $C$, and the scope is so large that it contains a point $\theta_{\mathrm{KL}}$ which minimizes the KL divergence between $f_{\Lambda}(\cdot, \cdot ; \theta) \in \Pi_{\Lambda}$ and $\varphi(\cdot, \cdot) \in \Phi$ among all choices of $\theta$ in $\tilde{\Theta}_{\Lambda}$, where

$$
\operatorname{KL}\left(f_{\Lambda}\left(\cdot, \cdot, \theta_{\mathrm{KL}}\right), \varphi(\cdot, \cdot)\right)=\inf _{\theta \in \tilde{\Theta}_{\Lambda}} \operatorname{KL}\left(f_{\Lambda}(\cdot, \cdot, \theta), \varphi(\cdot, \cdot)\right)=\inf _{f \in \Pi_{\Lambda}} \operatorname{KL}(f, \varphi) .
$$

This condition is similar to a usual condition under correct model specification, requiring that the scope of ML search should contain the true parameter so as to make the MLE consistent. The difference here is that there is no "true parameter," since the likelihood functions are constructed based on the HME densities, which can only be used to approximate the true pdf in $\boldsymbol{\Phi}$. Condition 2 ensures that the ML searching area is big enough to contain an "optimal point" (instead of the true parameter), which minimizes the KL divergence between the true density and the HME density. This feature will be useful when proving the consistency result of the ML approach under model misspecification, when the likelihood function is constructed from the HME approximations, instead of a pdf from the true family $\boldsymbol{\Phi}$. Note that Condition 2 is hard to check in practice, although it looks plausible if a sufficiently large scope of ML search is used.

The next theorem states that the maximum likelihood method based on the HME of GLM1 models is consistent in estimating the mean functions in $\psi\left(W_{2 ; K_{0}}^{\infty}\right)$.

Theorem 3 (Consistency of the maximum likelihood method). Let $(\mathrm{MSE})_{n, \Lambda}$ be as defined in (2.13). Under regularity conditions $A_{\mathscr{\rho}, 1}$ and 2,

$$
\lim _{m \rightarrow \infty} \limsup _{n \rightarrow \infty} \inf _{\Lambda \in \mathscr{P} \cap \mathscr{F}_{m}}(\mathrm{MSE})_{n, \Lambda}=0 .
$$

Here $n$ is the sample size, $m=\sup _{\Lambda \in \mathscr{A} \cap \mathscr{I}_{m}}\{\operatorname{card}(\Lambda)\}$ and $\mathscr{J}_{m}=\{\Lambda: \Lambda \in$ $\mathscr{J}, \operatorname{card}(\Lambda) \leq m\}$ is the set of all HME structures containing no more than $m$ experts. Actually,

$$
\limsup _{n \rightarrow \infty} \inf _{\Lambda \in \mathscr{f} \cap \mathscr{Z}_{m}}(\mathrm{MSE})_{n, \Lambda} \leq \frac{c}{m^{4 / s}},
$$

where $s=\operatorname{dim}(\mathbf{x})$ and $c$ is a positive constant independent of $n, m$ and the structure $\Lambda$.

The constant $c$ 's in the above theorems can be different.

Remark 3 (Unknown shape parameter). Up to now, we have been assuming that any shape parameter $u$ of a GLM1 expert is known, or fixed at a value which is equal to the shape parameter in the true $\operatorname{pdf} \varphi$. An example of the shape parameter is $u=1 / \sigma^{2}$ for a normal expert. Now suppose the shape parameter $u$ is unknown and needs also to be estimated. We can either assume the same shape parameter for all experts, or allow it to differ. In either case, the grand parameter $\theta$ will be expanded to include additional component(s) from the $u$ ('s). We assume that the parameter space $U$ of each component $u$ is a compact subset of the positive real line. Lemma 1(c) requires a small modification, that is, a bound of the KL distance now requires an additional term proportional to the discrepancy between the true shape parameter in $\varphi$ and the "proposed" shape parameter(s) in $f$. However, using similar techniques, it is straightforward to show that all theorems on denseness and consistency are still valid.
4. Proofs and secondary lemmas. In this section we will often use the following shorthand notation: $g_{J}=g_{J}(\mathbf{x} ; \mathbf{v}), h_{J}=h_{J}(\mathbf{x}), \pi_{J}=\pi\left(h_{J}(\mathbf{x}), y\right)$, $f=\sum_{J} g_{J} \pi_{J}$ [for the HME density (2.4)], $\sum_{J}=\sum_{J \in \Lambda}, \sup _{J}=\sup _{J \in \Lambda}, h=$ $h(\mathbf{x}), \varphi=\pi(h(\mathbf{x}), y)$ [for the true density (2.2)], $\mu=\mu_{\varphi}=\int y \varphi d \lambda, \mu_{f}=$ $\int y f d \lambda$ and $\mu_{J}=\int y \pi_{J} d \lambda$.

### 4.1. Proofs of Theorems 1 and 2.

Proof of Lemma 1. Here (b) is due to Lemma 3.2 of Zeevi and Meir (1997), which is a corollary to Lemma 5 of Haussler and Opper (1995).

The following observations prove (a):

$$
\begin{aligned}
d_{2}^{2}\left(\mu_{f}, \mu_{\varphi}\right) & =\int\left\{\int y f d \lambda-\int y \varphi d \lambda\right\}^{2} d \kappa \\
& =\int\left\{\int y(\sqrt{f}-\sqrt{\varphi})(\sqrt{f}+\sqrt{\varphi}) d \lambda\right\}^{2} d \kappa \\
& \leq \int\left\{\int|y||\sqrt{f}-\sqrt{\varphi}|(\sqrt{f}+\sqrt{\varphi}) d \lambda\right\}^{2} d \kappa
\end{aligned}
$$

$$
\begin{aligned}
& \stackrel{\text { (i) }}{\leq} 2 \int\left\{\int|y||\sqrt{f}-\sqrt{\varphi}|(\sqrt{f+\varphi}) d \lambda\right\}^{2} d \kappa \\
& \stackrel{\text { (ii) }}{\leq} 2 \int\left[\left\{\int y^{2}(f+\varphi) d \lambda\right\}\left\{\int(\sqrt{f}-\sqrt{\varphi})^{2} d \lambda\right\}\right] d \kappa \\
& \stackrel{\text { (iii) }}{\leq} 4 M_{I} \iint(\sqrt{f}-\sqrt{\varphi})^{2} d \lambda d \kappa .
\end{aligned}
$$

Inequality (i) above uses $\sqrt{f}+\sqrt{\varphi} \leq \sqrt{2} \sqrt{f+\varphi}$; (ii) uses the Cauchy-Schwarz inequality. To obtain (iii), note that

$$
\begin{aligned}
\int y^{2}(f+\varphi) d \lambda & =\sum_{J} g_{J} \int y^{2} \pi_{J} d \lambda+\int y^{2} \varphi d \lambda \\
& \leq \sup _{J \in \Lambda} \sup _{\mathbf{x} \in \Omega}\left(\int y^{2} \pi_{J} d \lambda\right)\left(\sum_{J} g_{J}\right)+\sup _{\mathbf{x} \in \Omega} \int y^{2} \varphi d \lambda \\
& =\sup _{J \in \Lambda} \sup _{\mathbf{x} \in \Omega}\left(\int y^{2} \pi_{J} d \lambda\right)+\sup _{\mathbf{x} \in \Omega} \int y^{2} \varphi d \lambda \\
& =\sup _{J \in \Lambda} \sup _{\mathbf{x} \in \Omega}\left(\int y^{2} \pi\left(h_{J}, y\right) d \lambda\right)+\sup _{\mathbf{x} \in \Omega} \int y^{2} \pi(h, y) d \lambda \\
& \leq 2 \sup _{|h| \leq K}\left\{\int y^{2} \pi(h, y) d \lambda\right\}=2 M_{I} .
\end{aligned}
$$

The last inequality is true since $K$ is an upperbound for $\left|h_{J}(\mathbf{x})\right|$ and $|h(\mathbf{x})|$ for all $\mathbf{x}$ in $\Omega$ and all $J$ in $\Lambda$.

The following observations prove (c):

$$
\begin{aligned}
\mathrm{KL}(f, \varphi) & =\iint \varphi \log (\varphi / f) d \lambda d \kappa \\
& =-\iint \varphi \log \left[\sum_{J} g_{J} \exp \left(\left\{a\left(h_{J}\right)-a(h)\right\} y+\left\{b\left(h_{J}\right)-b(h)\right\}\right)\right] d \lambda d \kappa \\
& \stackrel{(\mathrm{i})}{\leq}-\iint \varphi \sum_{J} g_{J}\left[\left\{a\left(h_{J}\right)-a(h)\right\} y+\left\{b\left(h_{J}\right)-b(h)\right\}\right] d \lambda d \kappa \\
& =-\int \sum_{J} g_{J}\left[\left\{a\left(h_{J}\right)-a(h)\right\} \mu+\left\{b\left(h_{J}\right)-b(h)\right\}\right] d \kappa \\
& \stackrel{(\text { (ii) }}{=}-\frac{1}{2} \int \sum_{J} g_{J}\left\{\mu a^{\prime \prime}\left(h_{J *}\right)+b^{\prime \prime}\left(h_{J *}\right)\right\}\left(h_{J}-h\right)^{2} d \kappa \\
& \leq \frac{1}{2} \int \sum_{J} g_{J}\left|\mu a^{\prime \prime}\left(h_{J *}\right)+b^{\prime \prime}\left(h_{J *}\right)\right|\left(h_{J}-h\right)^{2} d \kappa
\end{aligned}
$$

$$
\stackrel{(\mathrm{iii})}{=} M_{I I} \int \sum_{J} g_{J}\left(h_{J}-h\right)^{2} d \kappa .
$$

Inequality (i) above is due to the convexity of $-\log (\cdot)$; (ii) is from a Taylor expansion for $h_{J}$ around $h$,

$$
-\left\{\mu a\left(h_{J}\right)+b\left(h_{J}\right)\right\}=-\{\mu a(h)+b(h)\}-\frac{1}{2}\left\{\mu a^{\prime \prime}\left(h_{J *}\right)+b^{\prime \prime}\left(h_{J *}\right)\right\}\left(h_{J}-h\right)^{2},
$$

where $h_{J *}$ is between $h_{J}$ and $h$. [The "linear term" disappears due to property (3) in Section 2.1.] To prove inequality (iii), note that $\left|h_{J *}\right| \leq \max \left\{\left|h_{J}\right|\right.$, $|h|\} \leq K$ since $K$ is an upperbound for $\left|h_{J}(\mathbf{x})\right|$ and $|h(\mathbf{x})|$ for all $\mathbf{x}$ in $\Omega$ and all $J$ in $\Lambda$. Note also that $\mu=\int y \pi(h, y) d \lambda$ where $|h| \leq K$. Then

$$
\begin{aligned}
\left|\mu a^{\prime \prime}\left(h_{J *}\right)+b^{\prime \prime}\left(h_{J *}\right)\right| & \leq|\mu|\left|a^{\prime \prime}\left(h_{J *}\right)\right|+\left|b^{\prime \prime}\left(h_{J_{*}}\right)\right| \\
& \leq \sup _{|h| \leq K}\left|\int y \pi(h, y) d \lambda\right| \sup _{|h| \leq K}\left|a^{\prime \prime}(h)\right|+\sup _{|h| \leq K}\left|b^{\prime \prime}(h)\right|,
\end{aligned}
$$

which proves (iii).
Proof of Theorem 2. Denote $\|f(\cdot)\|_{p, \kappa}=\left\{\int|f(\mathbf{x})|^{p} d \kappa\right\}^{1 / p}(p>0)$ and $\|f(\cdot)\|_{\infty}=\sup _{\mathbf{x} \in \Omega}|f(\mathbf{x})|$. Consider a fine partition sequence $\left\{\left\{Q_{J}^{(\nu)}\right\}_{J \in \Lambda_{0}^{(\nu)}: \Lambda_{0}^{(\nu)} \in, ~}^{\text {. }}\right.$ $\mathscr{S}, \nu=1,2, \ldots\}$ in Condition $A_{\mathscr{S}, 1}$, with a cardinality sequence $\left\{p_{\nu}\right\}$. We first show that for each $\nu \in\{1,2, \ldots$,$\} , for any h \in W_{2 ; K_{0}}^{\infty}$,

$$
\begin{equation*}
(\S) \equiv\left\|\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\cdot) \hat{h}_{J}(\cdot)-h(\cdot)\right\|_{2, \kappa} \leq \frac{c_{1}}{p_{\nu}^{2 / s}}, \tag{4.1}
\end{equation*}
$$

for some finite constant $c_{1}>0$, where, for each $\mathbf{x} \in \Omega$ and each $J \in \Lambda_{0}^{(\nu)}$,

$$
\begin{equation*}
\hat{h}_{J}(\mathbf{x}) \equiv \hat{\alpha}_{J}+\mathbf{x}^{T} \hat{\boldsymbol{\beta}}_{J} \equiv\left\{h\left(\boldsymbol{\xi}_{J}\right)-\boldsymbol{\xi}_{J}^{T} \nabla h\left(\boldsymbol{\xi}_{J}\right)\right\}+\mathbf{x}^{T} \nabla h\left(\boldsymbol{\xi}_{J}\right), \tag{4.2}
\end{equation*}
$$

$\xi_{J}$ being some point in the interior of $Q_{J}^{(\nu)}$. Here $\nabla h$ is the $s \times 1$ gradient column vector of a scalar function $h$. For any $h \in W_{2 ; K_{0}}^{\infty}$ and any $\hat{h}_{J}$ defined in (4.2), it is obvious that $\|h\|_{\infty} \leq K_{0}$ and $\left\|\hat{h}_{J}\right\|_{\infty} \leq K_{0}$.

To prove (4.1) note that

$$
\begin{align*}
(\S) & =\left\|\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\cdot)\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}\right\|_{2, \kappa}  \tag{4.3}\\
& \leq\|1\|_{2, \kappa} \sup _{J \in \Lambda_{0}^{(\nu)}}\left\|\hat{h}_{J}(\cdot)-h(\cdot)\right\|_{\infty}=\sup _{J \in \Lambda_{0}^{(\nu)}}\left\|\hat{h}_{J}(\cdot)-h(\cdot)\right\|_{\infty},
\end{align*}
$$

because $\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\mathbf{x})=1$ and $\kappa$ is a probability measure.

By a second-order Taylor expansion of $h(\mathbf{x})$ around $\boldsymbol{\xi}_{J}$ and the definition of $\hat{h}_{J}(\mathbf{x})$ in (4.2), we have, for all $\mathbf{x} \in Q_{J}^{(\nu)}, J \in \Lambda_{0}^{(\nu)}$,

$$
\begin{equation*}
\left|\hat{h}_{J}(\mathbf{x})-h(\mathbf{x})\right| \leq \frac{1}{2}\left(\sum_{|\mathbf{k}|=2}\left\|D^{\mathbf{k}} h\right\|_{\infty}\right)\left\{\rho\left(\mathbf{x}, \boldsymbol{\xi}_{J}\right)\right\}^{2} \leq \frac{1}{2} K_{0} \frac{c_{0}^{2}}{p_{\nu}^{2 / s}}, \tag{4.4}
\end{equation*}
$$

where $\rho\left(\mathbf{x}, \boldsymbol{\xi}_{J}\right)=\max _{1 \leq q \leq s}\left|\left(\mathbf{x}-\boldsymbol{\xi}_{J}\right)_{q}\right|, c_{0}$ is the bounding constant from Definition 1 and Condition $A_{\mathscr{\rho}, 1}$. The latter inequality holds since (1) $\mathbf{x}, \boldsymbol{\xi}_{J} \in Q_{J}^{(\nu)}$, leading to $\rho\left(\mathbf{x}, \boldsymbol{\xi}_{J}\right) \leq c_{0} / p_{\nu}^{1 / s}$ by Condition $A_{\mathcal{\rho}, 1}$ and the definition of a "fine partition"; (2) $h \in W_{2 ; K_{0}}^{\infty}$. Then, (4.4) and (4.3) lead to (4.1) (take $c_{1}=\frac{1}{2} K_{0} c_{0}^{2}$ ).

Now, by Condition $A_{\mathcal{\rho}, 1}$, for all $\varepsilon>0$, there exists $\mathbf{v}_{\varepsilon} \in V_{\Lambda_{0}^{(\nu)}}$ such that (3.1) holds for the norm $\|(\cdot)\|_{1, \sigma}$. This implies that

$$
\begin{equation*}
\sup _{J \in \Lambda_{0}^{(\nu)}}\left\|g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)-\chi_{Q_{J}^{(\nu)}}(\cdot)\right\|_{1, \kappa} \leq\left\|\frac{d \kappa}{d \sigma}\right\|_{\infty} \varepsilon \tag{4.5}
\end{equation*}
$$

where $\|d \kappa / d \sigma\|_{\infty}=\sup _{\mathbf{x} \in \Omega}\left|\left(d \kappa / d \sigma_{0}\right) /\left(d \sigma / d \sigma_{0}\right)\right|<\infty$, since both $\sigma$ and $\kappa$ have positive continuous densities with respect to the Lebesgue measure $\sigma_{0}$. Now consider

$$
\begin{aligned}
(*) \equiv & \|\underbrace{}_{\left(\dagger \in \Lambda_{0}^{(\nu)}\right.} g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2}\|_{1, \kappa} \\
\leq & \left\|\sum_{J \in \Lambda_{0}^{(\nu)}}\left\{g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)-\chi_{Q_{J}^{(\nu)}}(\cdot)\right\}\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2}\right\|_{1, \kappa} \\
& +\underbrace{\left\|\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\cdot)\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2}\right\|_{1, \kappa}}_{(\ddagger)},
\end{aligned}
$$

due to the triangular inequality.
Note that

$$
\begin{aligned}
(\ddagger) & =\int \sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\cdot)\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2} d \kappa \\
& \stackrel{(\mathrm{i})}{=} \int\left\{\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}(\cdot) \hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2} d \kappa \stackrel{(\mathrm{ii})}{\leq} \frac{c_{1}^{2}}{p_{\nu}^{4 / s}} .
\end{aligned}
$$

Here (ii) follows from (4.1), and (i) uses the properties $\sum_{J \in \Lambda_{0}^{(\nu)}} \chi_{Q_{J}^{(\nu)}}=1$ and $\chi_{Q_{J}^{(\nu)}} \chi_{Q_{I}^{(\nu)}}=\chi_{Q_{J}^{(\nu)}} \delta_{I J}$, where $\delta_{I J}=1$ if $I=J$, and 0 otherwise.

Note also that

$$
\begin{aligned}
&(\dagger) \stackrel{(\text { iii) }}{\leq} \sum_{J \in \Lambda_{0}^{(\nu)}}\left\|\left\{g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)-\chi_{Q_{J}^{(\nu)}}(\cdot)\right\}\left\{\hat{h}_{J}(\cdot)-h(\cdot)\right\}^{2}\right\|_{1, \kappa} \\
& \leq\left\{\sup _{J \in \Lambda_{0}^{(\nu)}}\left\|\hat{h}_{J}(\cdot)-h(\cdot)\right\|_{\infty}\right\}^{2} \sum_{J \in \Lambda_{0}^{(\nu)}}\left\|g_{J}\left(\cdot ; \mathbf{v}_{\varepsilon}\right)-\chi_{Q_{J}^{(\nu)}}(\cdot)\right\|_{1, \kappa} \\
& \quad \stackrel{\text { (iv) }}{\leq}\left(\frac{c_{1}}{p_{\nu}^{2 / s}}\right)^{2} p_{\nu}\left\|\frac{d \kappa}{d \sigma}\right\|_{\infty} \varepsilon .
\end{aligned}
$$

Here (iii) is due to the triangular inequality, and (iv) is due to (4.3)-(4.5).
Combining the results for $(\dagger)$ and ( $\ddagger$ ), we get

$$
\begin{equation*}
(*) \leq\left(\frac{c_{1}}{p_{\nu}^{2 / s}}\right)^{2} p_{\nu}\left\|\frac{d \kappa}{d \sigma}\right\|_{\infty} \varepsilon+\left(\frac{c_{1}}{p_{\nu}^{2 / s}}\right)^{2} \tag{4.6}
\end{equation*}
$$

Denote $\varphi=\pi(h(\cdot), y)$ and $\hat{f}=\sum_{J \in \Lambda_{0}^{(\nu)}} g_{J}\left(\cdot, \mathbf{v}_{\varepsilon}\right) \pi\left(\hat{h}_{J}(\cdot), y\right)$. Then obviously, $\hat{f} \in \boldsymbol{\Pi}_{p_{\nu}, \mathcal{S}}$ and $\varphi \in \boldsymbol{\Phi}$, and $(*)=\mathscr{Q}(\hat{f}, \varphi)$. Using Lemma 1(c), we have

$$
\begin{equation*}
\mathrm{KL}(\hat{f}, \varphi) \leq M_{I I}\left(\frac{c_{1}}{p_{\nu}^{2 / s}}\right)^{2} p_{\nu}\left\|\frac{d \kappa}{d \sigma}\right\|_{\infty} \varepsilon+M_{I I}\left(\frac{c_{1}}{p_{\nu}^{2 / s}}\right)^{2} . \tag{4.7}
\end{equation*}
$$

The upperbound $K$ in $M_{I I}$ can be taken as $K_{0}$, since both $|h|$ and $\left|\hat{h}_{J}\right|$ are bounded above by $K_{0}$. This makes $M_{I I}$ to be a constant that does not depend on $\varphi$.

By (4.7) and the arbitrariness of $\varepsilon$, we have

$$
\begin{equation*}
\inf _{f \in \boldsymbol{\Pi}_{p_{\nu}, \digamma}} \operatorname{KL}(f, \varphi) \leq \frac{M_{I I} c_{1}^{2}}{p_{\nu}^{4 / s}} \tag{4.8}
\end{equation*}
$$

Since $\left\{p_{\nu}^{1 / s}\right\}$ is subgeometric, for all $m \in \mathscr{N}$, there exists $p_{\nu}$, such that $p_{\nu} \leq m<p_{\nu+1}$, and

$$
\begin{equation*}
1 / p_{\nu}^{4 / s} \geq 1 / m^{4 / s}>1 / p_{\nu+1}^{4 / s} . \tag{4.9}
\end{equation*}
$$

By the definition in (2.7), $\Pi_{m, \mathscr{S}}$ is monotone nondecreasing in $m$, and hence

$$
\boldsymbol{\Pi}_{p_{v}, \mathscr{L}} \subset \boldsymbol{\Pi}_{m, \mathscr{L}} \subset \boldsymbol{\Pi}_{p_{v+1}, \mathscr{A}}
$$

Hence, for all $m \in \mathscr{N}$,

$$
\begin{aligned}
\inf _{f \in \boldsymbol{\Pi}_{m, \rho}} \mathrm{KL}(f, \varphi) & \leq \inf _{f \in \boldsymbol{\Pi}_{p_{\nu}, \iota}} \mathrm{KL}(f, \varphi) \leq \frac{M_{I I} c_{1}^{2}}{p_{\nu}^{4 / s}} \quad[\mathrm{by}(4.8)] \\
& \leq \frac{M_{2}^{4} M_{I I} c_{1}^{2}}{p_{\nu+1}^{4 /}} \leq \frac{M_{2}^{4} M_{I I} c_{1}^{2}}{m^{4 / s}}=c / m^{4 / s}
\end{aligned}
$$

by noting that $\left\{p_{\nu}^{1 / s}\right\}$ is subgeometric with rate bounded by $M_{2}$, and using (4.9).

By construction, $c$ does not depend on $\varphi$ in $\Phi$. Hence,

$$
\sup _{\varphi \in \boldsymbol{\Phi}} \inf _{f \in \boldsymbol{\Pi}_{m, \mathscr{}}} \mathrm{KL}(f, \varphi) \leq c / m^{4 / s} \quad \text { for all } m \in \mathscr{N}
$$

[Tracing the constant $c$ leads to $c=\left(\frac{1}{2} K_{0} c_{0}^{2} M_{2}^{2} \sqrt{M_{I I}}\right)^{2}$, which does not have explicit dependence on $s=\operatorname{dim}(\mathbf{x})$.]

The proof of Theorem 1 is obvious from Lemma 1(b) and Theorem 2.
4.2. Proof of Theorem 3. We first bound the MSE defined in (2.13) by the sum of a "stochastic" part $2 S_{n \Lambda}$ and a "deterministic" part $2 D_{n \Lambda}$. Then we will prove that each part goes to zero as the sample size and the number of experts increase:

$$
\begin{align*}
\operatorname{MSE}_{n, \Lambda} \equiv & \mathrm{E} \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}\right)-\mu(\mathbf{x})\right\}^{2} d \kappa \\
\leq & 2 \mathrm{E} \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}\right)-\mu(\mathbf{x})\right\}^{2} d \kappa  \tag{4.10}\\
& +2 \mathrm{E} \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}\right)-\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}\right)\right\}^{2} d \kappa \\
\equiv & 2 D_{n, \Lambda}+2 S_{n, \Lambda}
\end{align*}
$$

since $(a+b)^{2} \leq 2 a^{2}+2 b^{2}$ for all $a, b \in \mathfrak{R}$. Here $\theta_{n, \Lambda}$, is a point in $\Theta_{\Lambda}$ defined in (2.12).

Note that by the definition of $\theta_{n, \Lambda}$,

$$
\begin{aligned}
\theta_{n, \Lambda} & =\arg \max _{\theta \in \tilde{B}_{\Lambda}} \iint_{\Omega \otimes A} \varphi(\mathbf{x}, y) \log \left\{\frac{f_{\Lambda}(\mathbf{x}, y ; \theta)}{\varphi_{0}(\mathbf{x}, y)}\right\} d \lambda(y) d \kappa(\mathbf{x}) \\
& \equiv \arg \max _{\theta \in \tilde{B}_{\Lambda}} Ł_{\infty, \Lambda}(\theta)
\end{aligned}
$$

where we choose $\varphi_{0}(\mathbf{x}, y)=e^{c(y)}[c(\cdot)$ is from (2.1)]. There could be more than one such maximizer and the set of such maximizers is just $\Theta_{\Lambda}$.

We will treat $\theta_{n, \Lambda}$, which is used in (4.10), as a random variable depending on sample size $n$, since our choice of $\theta_{n, \Lambda}$ out of the set $\Theta_{\Lambda}$ may depend on the MLE $\hat{\theta}_{n, \Lambda}(\omega)$ that we adopt. Note that $\theta_{n, \Lambda}$ is a minimizer of the KL divergence, when the parameter varies in the scope $\tilde{B}_{\Lambda}$. However, when Condition 2 holds, $\theta_{n, \Lambda}$ also minimizes the KL divergence over the entire parameter space $\tilde{\Theta}_{\Lambda}$, that is,

$$
\begin{equation*}
\mathrm{KL}\left\{f_{\Lambda}\left(\cdot, \cdot ; \theta_{n, \Lambda}\right), \varphi(\cdot, \cdot)\right\}=\operatorname{KL}\left\{f_{\Lambda}\left(\cdot, \cdot ; \theta_{\mathrm{KL}}\right), \varphi(\cdot, \cdot)\right\}=\inf _{f \in \Pi_{\Lambda}} \mathrm{KL}(f, \varphi) \tag{4.11}
\end{equation*}
$$

In the following, we formulate the convergence of MLE in a setting with nonidentifiable parameterization. We first state and prove a proposition on the uniform convergence of the log-likelihood function.

Proposition 1 (Uniform convergence of log-likelihood). Let $\mathrm{Ł}_{n, \Lambda}$ be the log-likelihood function defined in (2.9) [with $\varphi_{0}\left(\mathbf{X}_{i}, Y_{i}\right)=\exp \left(c\left(Y_{i}\right)\right), c(\cdot)$ is from (2.1)]. We have

$$
\sup _{\theta \in \tilde{B}_{\Lambda}}\left|\coprod_{n, \Lambda}(\theta ; \omega)-Ł_{\infty, \Lambda}(\theta)\right| \rightarrow 0
$$

for almost all stochastic sequences $\omega$, and

$$
\mathrm{Ł}_{\infty, \Lambda}(\theta)=\int_{\Omega \otimes A} \varphi(\mathbf{x}, y) \log \left\{f_{\Lambda}(\mathbf{x}, y ; \theta) / \varphi_{0}(\mathbf{x}, y)\right\} d \kappa d \lambda
$$

is a continuous function of $\theta$ on $\tilde{B}_{\Lambda}$.
Proof. Choose $\varphi_{0}(\mathbf{x}, y)=e^{c(y)}$ [see (2.1)]. Denote $f_{*}=f_{\Lambda}(\mathbf{x}, y ; \theta) /$ $\varphi_{0}(\mathbf{x}, y)$. Then $f_{*}=\sum_{J} g_{J} e_{J}$, where $e_{J}=\exp \left(y a\left(h_{J}\right)+b\left(h_{J}\right)\right)$. By construction, $f_{*}$ is a measurable function of $(\mathbf{x}, y)$ for each $\theta$, and a continuous function of $\theta$ for each $(\mathbf{x}, y)$.

Note that

$$
\begin{aligned}
\sup _{J}\left|\log e_{J}\right| & \geq \sup _{J}\left(\log e_{J}\right)=\log \left(\sup _{J} e_{J}\right) \\
& \geq \log \left(\sum_{J} g_{J} e_{J}\right) \geq \sum_{J} g_{J}\left(\log e_{J}\right) \\
& \geq-\sup _{J}\left|\log e_{J}\right|
\end{aligned}
$$

where we have used the monotonicity and concavity of $\log (\cdot)$ and the fact that $g_{J}$ 's are nonnegative and have a unity sum.

Hence, for all $\theta \in \tilde{B}_{\Lambda}, \mathbf{x} \in \Omega$ and $y \in A$,

$$
\begin{aligned}
\left|\log \left\{f_{\Lambda}(\mathbf{x}, y ; \theta) / \varphi_{0}(\mathbf{x}, y)\right\}\right| & =\left|\log f_{*}\right|=\left|\log \left(\sum_{J} g_{J} e_{J}\right)\right| \\
& \leq \sup _{J}\left|\log e_{J}\right|=\sup _{J}\left|y a\left(h_{J}\right)+b\left(h_{J}\right)\right| \\
& \leq M_{a}|y|+M_{b} \equiv M(y),
\end{aligned}
$$

with $M_{a}=\sup _{J} \sup _{\mathbf{x} \in \Omega} \sup _{\theta \in \tilde{B}_{\Lambda}}\left|a\left(h_{J}\right)\right|$ and $M_{b}=\sup _{J} \sup _{\mathbf{x} \in \Omega} \sup _{\theta \in \tilde{B}_{A}}\left|b\left(h_{J}\right)\right|$ being finite, since $\alpha\left(h_{J}\right)$ and $b\left(h_{J}\right)$ are continuous functions of $\left(\mathbf{x}^{T}, \theta^{T}\right)$ on the compact set $\Omega \otimes \tilde{B}_{\Lambda}$. Next we show that $\mathrm{E}\{M(Y)\}=\iint M(y) \varphi d \lambda d \kappa<\infty$. Note that $\int_{A} y^{2} \pi(h, y) d \lambda$ is a continuous function in $h$, due to property (2) of Section 2.1. By the continuity of $h(\cdot), \int_{A} y^{2} \pi(h(\mathbf{x}), y) d \lambda$ is a continuous function of $\mathbf{x}$ in the compact set $\Omega$, leading to the finiteness of $\int_{\Omega}\left\{\int_{A} y^{2}\right.$. $\pi(h(\mathbf{x}), y) d \lambda\}^{1 / 2} d \kappa$. Therefore, $\iint|y| \varphi d \lambda d \kappa$, being bounded above by $\int_{\Omega}\left\{\int_{A} y^{2} \pi(h(\mathbf{x}), y) d \lambda\right\}^{1 / 2} d \kappa$, is also finite, and so is $\mathrm{E}\{M(Y)\}$.

Therefore, $\log \left\{f_{\Lambda}(\mathbf{x}, y ; \theta) / \varphi_{0}(\mathbf{x}, y)\right\}$ is bounded above by an integrable function. By the uniform law of large numbers [Jennrich (1969), Theorem 2, or White (1994), Theorem A.2.1], we obtain the strong uniform convergence of $Ł_{n, \Lambda}(\theta ; \omega)$, as well as the continuity of $Ł_{\infty, \Lambda}(\theta)$.

This proposition leads to the following lemma for the convergence of MLE.
Lemma 2. Let $\hat{\theta}_{n, \Lambda}(\omega), n=1,2, \ldots$, be a sequence of global maximizers of $£_{n, \Lambda}(\theta, \omega)$ in $\tilde{B}_{\Lambda}$, as defined in (2.10). Let $\Theta_{\Lambda}$ be the set of minimizers of the KL divergence between the true density and the HME density, as defined in (2.12). We have

$$
\inf _{\theta \in \Theta_{\Lambda}} \rho_{E}\left(\hat{\theta}_{n, \Lambda}(\omega), \theta\right) \rightarrow 0 \quad \text { as } n \rightarrow \infty
$$

for almost all $\omega$, where $\rho_{E}(\cdot, \cdot)$ is the Euclidean metric.
Proof. Denote

$$
d\left(\phi, \Theta_{\Lambda}\right) \equiv \inf _{\theta \in \Theta_{\Lambda}} \rho_{E}(\phi, \theta) \quad \text { for } \phi \in \tilde{B}_{\Lambda}
$$

For any $\varepsilon>0$, denote

$$
B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right) \equiv\left\{\phi \in \tilde{B}_{\Lambda}: d\left(\phi, \Theta_{\Lambda}\right) \geq \varepsilon\right\}
$$

We show
(*) For almost all $\omega$, there exists $N(\omega, \varepsilon) \in \mathscr{N}$ such that $n>N(\omega, \varepsilon)$
implies $\hat{\theta}_{n, \Lambda}(\omega) \notin B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$. Then $d\left(\hat{\theta}_{n, \Lambda}(\omega), \Theta_{\Lambda}\right) \leq \varepsilon$.
To show (*), note that by Proposition $1, \mathrm{Ł}_{n, \Lambda}(\cdot ; \omega)$ is uniformly convergent to $\mathrm{Ł}_{\infty, \Lambda}(\cdot)$ for almost all $\omega$. Hence, for any $\delta>0$, there exists $N(\omega)$ such that $n>N(\omega)$ implies that, for almost all $\omega$,

$$
\begin{aligned}
\mathrm{Ł}_{n, \Lambda}\left(\theta_{0} ; \omega\right) & >\mathrm{Ł}_{\infty, \Lambda}\left(\theta_{0}\right)-\delta / 2 \\
\mathrm{Ł}_{n, \Lambda}(\phi ; \omega) & <\mathrm{Ł}_{\infty, \Lambda}(\phi)+\delta / 2
\end{aligned}
$$

for all $\theta_{0} \in \Theta_{\Lambda}$ and all $\phi \in B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$. Hence,

$$
\begin{equation*}
\mathrm{Ł}_{n, \Lambda}\left(\theta_{0} ; \omega\right)-\mathrm{Ł}_{n, \Lambda}(\phi ; \omega)>\mathrm{Ł}_{\infty, \Lambda}\left(\theta_{0}\right)-\mathrm{Ł}_{\infty, \Lambda}(\phi)-\delta \tag{4.12}
\end{equation*}
$$

We can choose $\delta>0$ so small that

$$
\begin{equation*}
\delta<\inf _{\phi \in B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)}\left[Ł_{\infty, \Lambda}\left(\theta_{0}\right)-Ł_{\infty, \Lambda}(\phi)\right] \equiv I_{\varepsilon} \tag{4.13}
\end{equation*}
$$

Note that $I_{\varepsilon}>0$. This is because $B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$ is compact and $Ł_{\infty, \Lambda}$ is continuous (by Proposition 1), and the infimum in (4.13) is achieved for some $\phi \in B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$, where $\mathrm{Ł}_{\infty, \Lambda}\left(\theta_{0}\right)-\mathrm{Ł}_{\infty, \Lambda}(\phi)>0$. Hence, for all $\phi \in B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$, when $n>N(\omega)$, where $N(\omega)$ is as chosen above,

$$
Ł_{n, \Lambda}\left(\theta_{0} ; \omega\right)-Ł_{n, \Lambda}(\phi ; \omega)>0 \quad \text { by (4.12) and (4.13). }
$$

So $\hat{\theta}_{n, \Lambda}(\omega)=\arg \max _{\phi \in \tilde{B}_{\Lambda}} Ł_{n, \Lambda}(\phi ; \omega) \notin B_{\varepsilon}^{c}\left(\Theta_{\Lambda}\right)$, leading to the proof of $(*)$ and the lemma.

The next lemma shows that the "stochastic part" of the MLE goes to zero when the sample size increases.

LEMMA 3. There is a sequence of $\theta_{n, \Lambda} \in \Theta_{\Lambda}, n=1,2, \ldots$, possibly random, making $S_{n, \Lambda} \rightarrow 0$ as $n \rightarrow 0$, for any structure $\Lambda$.

Proof. Note that Lemma 2 implies that there exists a sequence $\theta_{n, \Lambda}(\omega) \in$ $\Theta_{\Lambda}, n=1, \ldots$, such that $\rho_{E}\left(\hat{\theta}_{n, \Lambda}(\omega), \theta_{n, \Lambda}(\omega)\right) \rightarrow 0$ for almost all $\omega$. By the definition in (2.8), $\mu_{\Lambda}(\cdot ; \cdot)$ is continuous on $\Omega \otimes \tilde{B}_{\Lambda}$, due to the continuity of the $g_{J}(\cdot ; \cdot)$ 's and $\psi(\cdot)$. Hence, we have

$$
\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}(\omega)\right)-\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)\right\}^{2} \rightarrow 0
$$

for all $\mathbf{x}$ and almost all $\omega$.
Next we show that $\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}(\omega)\right)-\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)\right\}^{2}$ is bounded above by an integrable function. This is because

$$
\begin{aligned}
& \left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}(\omega)\right)-\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)\right\}^{2} \\
& \quad \leq 2\left\{\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}(\omega)\right)\right\}^{2}+2\left\{\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)\right\}^{2} \\
& \quad \leq 4 M_{\Lambda}^{2},
\end{aligned}
$$

where $M_{\Lambda} \equiv \sup _{\mathbf{x} \in \Omega} \sup _{\theta \in \tilde{B}_{\Lambda}}\left|\mu_{\Lambda}(\mathbf{x} ; \theta)\right|<\infty$.
Therefore, by the Lebesgue's dominated convergence theorem,

$$
\mathrm{E} \int\left[\mu_{\Lambda}\left(\mathbf{x} ; \hat{\theta}_{n, \Lambda}\right)-\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}\right)\right]^{2} d \kappa(\mathbf{x}) \rightarrow 0
$$

Now we provide an upper-bound for the "deterministic part" $2 D_{n, \Lambda}$.
Lemma 4. Let $D_{n, \Lambda}=E \int_{\Omega}\left\{\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}\right)-\mu(\mathbf{x})\right\}^{2} d \kappa(\mathbf{x})$ be as defined in (4.10). If Conditions $A_{\mathscr{\rho}, 1}$ and 2 hold, then we have

$$
\inf _{\Lambda \in \mathscr{\mathscr { A }}_{m} \cap \mathscr{S}} D_{n, \Lambda} \leq c^{*} / m^{4 / s}
$$

for some finite positive constant $c^{*}$ independent of the number of experts $m$.
Proof. Consider a sequence of maximizers $\theta_{n, \Lambda}=\theta_{n, \Lambda}(\omega)$ in $\Theta_{\Lambda}$. Note that for each $\omega, n$ and $\Lambda$, we have

$$
\begin{aligned}
& \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)-\mu(\mathbf{x})\right\}^{2} d \kappa \stackrel{(\mathrm{i})}{\leq} 4 M_{I} \operatorname{KL}\left\{f_{\Lambda}\left(\cdot, \cdot ; \theta_{n, \Lambda}(\omega)\right), \varphi(\cdot, \cdot)\right\} \\
& \stackrel{(\mathrm{ii})}{=} 4 M_{I} \inf _{f \in \Pi_{\Lambda}} \operatorname{KL}(f, \varphi) .
\end{aligned}
$$

Then $D_{n, \Lambda}=\mathrm{E} \int\left\{\mu_{\Lambda}\left(\mathbf{x} ; \theta_{n, \Lambda}(\omega)\right)-\mu(\mathbf{x})\right\}^{2} d \kappa \leq 4 M_{I} \inf _{f \in \Pi_{\Lambda}} \operatorname{KL}(f, \varphi)$, and

$$
\begin{aligned}
\inf _{\Lambda \in \mathcal{A}_{m} \cap \mathscr{S}} D_{n, \Lambda} & \leq 4 M_{I} \inf _{\Lambda \in \mathscr{\mathscr { A }}_{m} \cap \mathscr{S}} \inf _{f \in \Pi_{\Lambda}} \mathrm{KL}(f, \varphi) \\
& =4 M_{I} \inf _{f \in \Pi_{m, \Omega}} \mathrm{KL}(f, \varphi) \stackrel{(\mathrm{iii})}{\leq} 4 M_{I} c / m^{4 / s}
\end{aligned}
$$

Here (i) is due to Lemma 1(a), where the constant $M_{I}$ can be made independent of $\varphi$ and $f_{\Lambda}\left(\cdot, \cdot ; \theta_{n, \Lambda}(\omega)\right)$ by using an upperbound $K=\max \left\{(s+1) C, K_{0}\right\}$, where $C$ is the constant in Condition 2 and $K_{0}$ is the radius of the Sobolev ball in Section 2.2. [Note that for all parameters in $\tilde{B}_{\Lambda},\left|h_{J}(\mathbf{x})\right|$ is less than $(s+1) C$, and for all $\varphi$ in $\boldsymbol{\Phi},|h(\mathbf{x})|$ is less than $K_{0}$.] (ii) is due to Condition 2, and equation (4.11); (iii) is due to Theorem 2 and Condition $A_{\mathscr{\rho}, 1}$.

Now we are ready to prove the consistency theorem.
Proof of Theorem 3. Note that (MSE) $n_{n, \Lambda} \leq 2 D_{n, \Lambda}+2 S_{n, \Lambda}$ from (4.10). For each $\Lambda$ in $\mathscr{J}_{m} \cap \mathscr{S}$, find a (possibly random) sequence $\theta_{n, \Lambda}$ as in Lemma 3 such that $S_{n, \Lambda} \rightarrow 0$. Then $\sup _{\Lambda \in \mathscr{f}_{m} \cap, \mathscr{J}} S_{n, \Lambda} \rightarrow 0$ as $n \rightarrow \infty$, since the cardinality of $\mathscr{J}_{m} \cap \mathscr{S}$ is finite. Then we have

$$
\begin{aligned}
& \inf _{\Lambda \in \mathscr{A}_{m} \cap \mathscr{A}}(\mathrm{MSE})_{n, \Lambda} \leq \inf _{\Lambda \in \mathscr{A}_{m} \cap \mathscr{A}}\left(2 S_{n, \Lambda}+2 D_{n, \Lambda}\right) \\
& \leq 2 \sup _{\Lambda \in \mathscr{f}_{m} \cap \mathscr{S}} S_{n, \Lambda}+2 \inf _{\Lambda \in \mathscr{A}_{m} \cap \mathscr{S}} D_{n, \Lambda} \\
& \leq o_{m}\left(n^{0}\right)+2 c^{*} / m^{4 / s},
\end{aligned}
$$

due to Lemma 4. [We denote the term $\sup _{\Lambda \in \mathscr{I}_{m} \cap \mathscr{\mathscr { ~ }}} S_{n, \Lambda}$ as $o_{m}\left(n^{0}\right)$, since it is $o(1)$ as $n \rightarrow \infty$, and is possibly dependent on $m$.]

Therefore

$$
0 \leq \limsup _{n \rightarrow \infty} \inf _{\Lambda \in \mathscr{\mathscr { H }}_{m} \cap \mathscr{S}}(\mathrm{MSE})_{n, \Lambda} \leq 2 c^{*} / m^{4 / s}
$$

and

$$
\lim _{m \rightarrow \infty} \limsup _{n \rightarrow \infty} \inf _{\Lambda \in \mathscr{\mathscr { A }}_{m} \cap \mathscr{A}}(\mathrm{MSE})_{n, \Lambda}=0
$$

5. Conclusions. We investigated the power of the HME for GLM1 experts in terms of approximating a flexible class of density functions with conditional mean functions belonging to a transformed Sobolev class. We demonstrated that the approximation rate of HME density functions is $O\left(m^{-2 / s}\right)$ in Hellinger distance and $O\left(m^{-4 / s}\right)$ in KL divergence. Here $s$ is the dimension of the predictor, and $m$ is the maximal number of experts in the network. We also showed that the maximum likelihood (ML) approach, which is associated with some optimal statistical properties and a convenient maximization algorithm, is consistent in estimating the mean response from data as the sample size and the number of experts both increase. Moreover, the approximation rates and the consistency result can be achieved within the family of HME structures with binary trees, or within the family of HME structures with one layer of experts (the MEs). We do not claim that the $O\left(m^{-2 / s}\right)$ rate is optimal. In fact, for the special case of mixing linear model experts in a single layer, Zeevi, Meir and Maiorov (1998) have shown that a better rate for approximation of mean functions can be achieved if higher than second-order continuous differentiability of the target functions is assumed. Our work is
different from Zeevi, Meir and Maiorov (1998) in the following aspects: (1) We deal with mixtures of generalized linear models instead of the mixtures of ordinary linear models. (2) We consider the setup of the HME networks instead of the single-layer mixtures of experts. (3) We consider the maximum likelihood method instead of the least-squares approach for model fitting. (4) In relation to the use of the maximum likelihood method, we obtained the approximation rate in terms of probability density functions instead of in terms of the mean response. (5) We have formulated the conditions and proofs of our results in a way that is protective of the inherent nonidentifiability problems of the parameterization.

Acknowledgments. The authors thank John Kolassa and the referees for helpful comments and Assaf J. Zeevi for suggesting a reference for Lemma 1(b).

## REFERENCES

Bickel, P. J. and Doksum, K. A. (1977). Mathematical Statistics. Prentice-Hall, Englewood Cliffs, NJ.
Bishop, C .M. (1995). Neural Networks for Pattern Recognition. Oxford Univ. Press.
Cacciatore, T. W. and Nowlan, S. J. (1994). Mixtures of controllers for jump linear and nonlinear plants. In Advances in Neural Informations Processing Systems 6 (G. Tesauro, D. S. Touretzky and T. K. Leen, eds.). Morgan Kaufmann, San Mateo, CA.

Devroye, L. and Gyoerfi, L. (1985). Nonparametric Density Estimation: The $L_{1}$ View. Wiley, New York.
Fritsch, J., Finke, M. and Waibel, A. (1997). Adaptively growing hierarchical mixtures of experts. In Advances in Neural Informations Processing Systems 9 (M. C. Mozer, M. I. Jordan and T. Petsche, eds.). MIT Press.
Ghahramani, Z. and Hinton, G. E. (1996). The EM algorithm for mixtures of factor analyzers. Technical Report CRG-TR-96-1, Dept. Computer Science, Univ. Toronto.
Haussler, D. and Opper, M. (1995). General bounds on the mutual information between a parameter and $n$ conditionally independent observations. In Proceedings of the Eighth Annual Computational Learning Theory Conference (COLT), 1995, Santa Cruz, CA. ACM Press, New York.
Haykin, S. (1994). Neural Networks. Macmillan, New York.
JaAkkola, T. S. and Jordan, M. I. (1998). Improving the mean field approximation via the use of mixture distributions. In Learning in Graphical Models (M. I. Jordan, ed.). Kluwer, Dordrecht.
Jacobs, R. A., Jordan, M. I., Nowlan, S. J. and Hinton, G. E. (1991). Adaptive mixtures of local experts. Neural Comp. 3 79-87.
Jennrich, R. I. (1969). Asymptotic properties of nonlinear least squares estimators. Ann. Math. Statist. 40 633-643.
Jiang, W. and Tanner, M. A. (1998). Hierarchical mixtures-of-experts for exponential family regression models: approximation and maximum likelihood estimation. Technical report, Dept. Statistics, Northwestern Univ., Evanston, IL.
Jiang, W. and Tanner, M. A. (1999). On the approximation rate of hierarchical mixtures-ofexperts for generalized linear models. Neural Comp. To appear.
Jordan, M. I. and Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. Neural Comp. 6 181-214.
Jordan, M. I. and Xu, L. (1995). Convergence results for the EM approach to mixtures-of-experts architectures. Neural Networks 8 1409-1431.
Lehmann, E. L. (1991). Theory of Point Estimation. Wadsworth, Monterey, CA.
McCullagh, P. and Nelder, J. A. (1989). Generalized Linear Models. Chapman and Hall, London.

Meilă, M. and Jordan, M. I. (1995). Learning fine motion by Markov mixtures of experts. A.I. Memo 1567, Artificial Intelligence Lab., Massachusetts Institute Technology.
Mhaskar, H. N. (1996). Neural networks for optimal approximation of smooth and analytic functions. Neural Comp. 8 164-177.
Peng, F., Jacobs, R. A. and Tanner, M. A. (1996). Bayesian inference in mixtures-of-experts and hierarchical mixtures-of-experts models with an application to speech recognition. J. Amer. Statist. Assoc. 91 953-960.

Tipping, M. E. and Bishop, C. M. (1997). Mixtures of probabilistic principal component analysers. Technical Report NCRG-97-003, Dept. Computer Science and Applied Mathematics, Aston Univ., Birmingham, UK.
Titterington, D. M., Smith, A. F. M. and Makov, U. E. (1985). Statistical Analysis of Finite Mixture Distributions. Wiley, New York.
White, H. (1994). Estimation, Inference and Specification Analysis. Cambridge Univ. Press.
Zeevi, A. and Meir, R. (1997). Density estimation through convex combinations: approximation and estimation bounds. Neural Networks 10 99-106.
Zeevi, A., Meir, R. and Maiorov, V. (1998). Error bounds for functional approximation and estimation using mixtures of experts. IEEE Trans. Information Theory 44 1010-1025.

Department of Statistics NORTHWESTERN UNIVERSITY
Evanston, Illinois 60208
E-MAIL: wjiang@nwu.edu tanm@neyman.stats.nwu.edu


[^0]:    Received March 1998; revised March 1999.
    ${ }^{1}$ Supported in part by NIH Grant CA35464.
    AMS 1991 subject classifications. Primary 62G07; secondary 41A25.
    Key words and phrases. Approximation rate, exponential family, generalized linear models, Hellinger distance, Hierarchical mixtures-of-experts, Kullback-Leibler divergence, maximum likelihood estimation, mean square error.

