# HISTOGRAM REGRESSION ESTIMATION USING DATA-DEPENDENT PARTITIONS ${ }^{1}$ 

By Andrew Nobel<br>University of North Carolina, Chapel Hill

We establish general sufficient conditions for the $L_{2}$-consistency of multivariate histogram regression estimates based on data-dependent partitions. These same conditions insure the consistency of partitioning regression estimates based on local polynomial fits, and, with an additional regularity assumption, the consistency of histogram estimates for conditional medians.

Our conditions require shrinking cells, subexponential growth of a combinatorial complexity measure and sublinear growth of restricted cell counts. It is not assumed that the cells of every partition be rectangles with sides parallel to the coordinate axis or that each cell contain a minimum number of points. Response variables are assumed to be bounded throughout.

Our results may be applied to a variety of partitioning schemes. We established the consistency of histograms regression estimates based on cubic partitions with data-dependent offsets, $k$-thresholding in one dimension and empirically optimal nearest-neighbor clustering schemes. In addition, it is shown that empirically optimal regression trees are consistent when the size of the trees grows with the number of samples at an appropriate rate.

1. Introduction. Let $\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right), \ldots,(X, Y) \in \mathbb{R}^{d} \times[-K, K]$ be independent and identically distributed random vectors defined on a common probability space. We may view the random vector $X \in \mathbb{R}^{d}$ as a collection of measurements that are related in a stochastic fashion to a response variable $Y \in[-K, K]$, whose value is of interest. The joint distribution of $(X, Y)$ is assumed to be unknown. We wish to estimate the regression function

$$
r(x)=E(Y \mid X=x) \in[-K, K]
$$

based on a training set of the form

$$
\begin{equation*}
T_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right) . \tag{1}
\end{equation*}
$$

Formally, a regression estimate is any function $\hat{f}_{n}\left(\cdot, T_{n}\right): \mathbb{R}^{d} \rightarrow \mathbb{R}$ that depends on the training set. The dependence of $\hat{f}_{n}$ on $T_{n}$ will be suppressed in

[^0]what follows. A sequence of estimates $\left\{\hat{f_{n}}\right\}$ is said to be strongly $L_{2}$-consistent if
$$
\int\left|\hat{f}_{n}(x)-r(x)\right|^{2} d P(x) \rightarrow 0
$$
with probability 1 . Here $P$ denotes the distribution of the random vector $X$. When the expected value of the integral tends to zero, the estimates are said to be weakly $L_{2}$-consistent.

When parametric models are not available, a natural means of estimating a multivariate regression function is to partition the observation space $\mathbb{R}^{d}$ into cells and then form estimates locally within each cell based on the response variables. In the simplest case this gives rise to histogram estimates based on data-independent partitions that consist of infinitely many congruent, rectangular cells. The size and location of the cells depends only on the cardinality of $T_{n}$, not on the geometrical or numerical properties of its constituent vectors. If the common dimensions of the rectangles shrink at an appropriate rate, results of Stone (1977) establish the weak consistency of the associated regression estimates, regardless of the underlying distribution of the data. Devroye and Györfi (1985) established the strong consistency of regression estimates based on data-independent cubic partitions, and they derived exponential bounds for the $L_{1}$-error of the estimates.

While regression estimates based on data-independent partitions are easy to implement, statistical practice suggests that estimates based on suitably chosen data-dependent partitions will provide better small-sample performance. For example, if the measurement vectors $X_{1}, \ldots, X_{n}$ fall into two distinct clusters, a data-dependent partition could allocate the majority of its cells within these clusters. By considering the response variables $Y_{i}$, a data-dependent partition can separate large and small values into separate cells or can allocate more cells to regions in which the behavior of the response variables is erratic. The flexibility of data-dependent partitions is likely to be most beneficial when the dimension $d$ of the measurement variables is large and the sample size $n$ is moderate.

In what follows we restrict our attention to measurable partitions having at most a countable number of cells. An n-sample partitioning rule is a deterministic mapping $\psi_{n}(\cdot)$ that associates each $n$-length sequence $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right) \in \mathbb{R}^{d} \times \mathbb{R}$ with a partition $\pi$ of $\mathbb{R}^{d}$. When it is applied to a training set $T_{n}$, the rule $\psi_{n}$ yields a random partition $\psi_{n}\left(T_{n}\right)$. The dependence of $\psi_{n}\left(T_{n}\right)$ on $T_{n}$ will be suppressed when no confusion will arise. With this convention in mind, $\psi_{n}[x]$ is defined to be the unique cell of $\psi_{n}\left(T_{n}\right)$ containing the vector $x$. Given a partitioning rule $\psi_{n}$ and a training set $T_{n}$, the ordinary histogram regression estimate is defined for each $x \in \mathbb{R}^{d}$ by

$$
\begin{equation*}
\hat{r}_{n}(x)=\frac{\sum_{i=1}^{n} Y_{i} I\left\{X_{i} \in \psi_{n}[x]\right\}}{\sum_{i=1}^{n} I\left\{X_{i} \in \psi_{n}[x]\right\}}, \tag{2}
\end{equation*}
$$

with $\hat{r}_{n}(x)=0$ when both the numerator and the denominator are zero. Note that $\hat{r}_{n}$ is piecewise constant on the cells of $\psi_{n}$ and that $\left|\hat{r}_{n}\right| \leq K$. Below we
present general sufficient conditions for the consistency of such estimates and assess the satisfaction of these conditions in a number of specific examples.
1.1. Discussion of related work. The simplest data-dependent partitioning methods are based on statistically equivalent blocks [Anderson (1966), Patrick and Fisher (1967)], in which each cell contains the same number of points. When $d=1$, statistically equivalent blocks reduce to $k$-spacing estimates [Parthasarathy and Bhattacharya (1961)], where the $k$ th-; $2 k$ th-; ... order statistics of $X_{1}, \ldots, X_{n}$ determine the partition of the real line. For a discussion of these and other related partitioning rules in the context of pattern recognition, we refer to the survey paper of Devroye (1988).

Theoretical evidence for the superiority of data-dependent histogram methods is suggested by Stone (1985). Stone (1977) gave necessary and sufficient conditions for the weak $L_{2}$-consistency of regression estimates based on data-dependent local averages. His conditions apply to histogram estimates based on partitioning rules $\psi_{n}\left(X_{1}, \ldots, X_{n}\right)$ that do not make use of the response variables $Y_{1}, \ldots, Y_{n}$.

In many cases of practical interest, histogram regression estimates are described by binary trees. Regression trees are produced in an iterative fashion by recursive partitioning schemes that seek at each step to minimize an empirical criterion function. The consistency of regression tress produced by means of supervised axis-parallel splitting was established by Gordon and Olshen (1984, 1980). Gordon and Olshen (1978) established similar results for classification trees.

Sufficient conditions for the weak $L_{2}$-consistency of the estimates (2) considered here can be found in Breiman, Friedman, Olshen and Stone (1984). Their result may be summarized as follows.

Theorem A [Breiman, Friedman, Olshen and Stone (1984)]. A sequence $\psi_{1}, \psi_{2}, \ldots$ of partitioning rules gives rise to weakly $L_{2}$-consistent estimates $\hat{r}_{n}$ if the following hold:
(a) each cell of $\psi_{n}$ contains at least $\alpha_{n} \log n$ points, where $\alpha_{n} \rightarrow \infty$;
(b) there is a collection of sets $\mathscr{C}$ having finite Vapnik-Chervonenkis dimension such that $\psi_{n}[x] \in \mathscr{C}$ for every $x$, every $n$ and every $T_{n}$;
(c) $P\left\{x: \operatorname{diam}\left(\psi_{n}[x]\right)>\gamma\right\} \rightarrow 0$ with probability 1 for every $\gamma>0$.

Under slightly stronger conditions, Gordon and Olshen (1984) established the strong $L_{2}$-consistency of the estimates $\hat{r}_{n}$ and showed that $\hat{r}_{n} \rightarrow r$ almost surely when the partitions $\left\{\hat{\psi}_{n}\right\}$ are nested. Chaudhuri, Huang, Loh and Yao (1994) considered regression estimates that are constructed by fitting a polynomial of fixed degree within each cell of a data-dependent partition, based on a local least squares criterion. They gave sufficient conditions under which, with probability tending to 1 , such functions give uniformly good estimates of the regression function, and a prespecified number of its derivatives, on any fixed compact set. Their conditions are similar to those of

Theorem A, with the exception of an additional regularity assumption that insures the invertibility of the matrices giving the optimal coefficients within each cell. We note here that the work cited above applied to unbounded response variables $Y_{i}$, under various moment restrictions.

Conditions (a) and (b) restrict the applicability of Theorem A and related results. Condition (a) is difficult to verify when $\psi_{n}$ is defined through minimization of an empirical criterion function, as in the tree-structured methods mentioned above. The required minimization does not generally insure that each cell contains a minimum number of points. Condition (b) requires that the cells of $\mu_{n}\left(T_{n}\right)$ be of fixed complexity, regardless of the sample size. In some cases the data may warrant increasing cell complexity as $n$ tends to infinity. In applying Theorem A to a particular method under study, satisfaction of (a), (b) and (c) is frequently accomplished by altering the method through supervisory oversight.

The results of this paper are based in part on combinatorial properties of partition families and related exponential inequalities. In this respect we follow the work of Vapnik and Chervonenkis (1971, 1981). Application of these ideas to histogram estimation originated with Breiman, Friedman, Olshen and Stone (1984) and in the later work of Zhao, Krishnaiah and Chen (1990), who considered histogram density estimates based on data-dependent partitions with rectangular cells. Adopting an approach similar to that taken here, Lugosi and Nobel (1995) established general sufficient conditions for the consistency of histogram density estimates and classification rules based on finite, data-dependent partitions.
1.2. Statement of main result. Let $\Pi$ be a family of partitions of $\mathbb{R}^{d}$. For each set $V \subseteq \mathbb{R}^{d}$, define the restricted cell count

$$
m(\Pi: V)=\max _{\pi \in \Pi}|\{A \in \pi: A \cap V \neq \varnothing\}|,
$$

where $|\{A \in \pi: A \cap V \neq \varnothing\}|$ measures the number of cells of $\pi$ that intersect $V$. The unrestricted cell count is defined by $m(\Pi)=m\left(\Pi: \mathbb{R}^{d}\right)$. The complexity of $\Pi$ is measured in terms of a combinatorial quantity proposed by Lugosi and Nobel (1995), which is similar to the growth function for classes of sets introduced by Vapnik and Chervonenkis (1971). Let $\eta=\left\{x_{1}, \ldots, x_{n}\right\}$ contain $n$ vectors in $\mathbb{R}^{d}$. Every element $\pi=\left\{A_{j}\right\} \in \Pi$ induces a partition $\left\{A_{j} \cap \eta\right\}$ of the finite set $\eta$. Let $\Delta\left(x_{1}^{h}, \Pi\right)$ be the number of distinct partitions of $\eta$ that are induced by the elements of $\Pi$. By "distinct" it is meant that the order of appearance of the individual sets is not important. The partitioning number

$$
\begin{equation*}
\Delta_{n}^{*}(\Pi)=\max \left\{\Delta\left(x_{1}^{n}, P i\right): x_{1}, \ldots, x_{n} \in \mathbb{R}^{d}\right\} \tag{3}
\end{equation*}
$$

is equal to the maximum number of different partitions of any $n$-point set that can be induced by members of $\Pi$.

Example 1. Let $\mathscr{U}(k)$ be the family of all partitions of $\mathbb{R}$ into $k$ nonempty intervals. Then $n(\mathscr{U}(k))=k$, and for any sequence of numbers $x_{1}<x_{2}<\cdots$
$<x_{n}$ an easy combinatorial argument shows that $\Delta\left(x_{1}^{n}, \mathscr{U}(k)\right)=\binom{n+k-1}{n}$. It follows that $\Delta_{n}^{*}(\mathscr{U}(k))=\binom{n+k-1}{n}$.

EXAMPLE 2. The nearest-neighbor partition of $k$ vectors $c_{1}, \ldots, c_{k} \in \mathbb{R}^{d}$ has $k$ cells $A_{1}, \ldots, A_{k}$. The cell $A_{i}$ contains those vectors $x \in \mathbb{R}^{d}$ that are closer to $c_{i}$ than any other $c_{j}$, with ties broken in favor of the vector having least index. Let $\mathscr{V}(k)$ be the family of nearest-neighbor partitions of $k$ vectors of $\mathbb{R}^{d}$. Then $m(\mathscr{V}(k))=k$. Moreover, it is readily verified that every cell of a partition $\pi \in \mathscr{V}(k)$ is the intersection of $(k-1)$ half-spaces. It is well known [c.f. Cover (1965)] that half-spaces in $\mathbb{R}^{d}$ can intersect $n$ vectors $x_{1}, \ldots, x_{n}$ in at most $n^{d}$ different ways. Thus each cell of a partition in $\mathscr{V}(k)$ can intersect $x_{1}, \ldots, x_{n}$ in at most $n^{(k-1) d}$ different ways. As every partition of $\mathscr{V}(k)$ contains $k$ cells, it follows that $\Delta_{n}^{*}(\mathscr{V}(k)) \leq n^{k^{2} d}$.

EXAMPLE 3. A tree-structured partition is described by a pair $(T, \tau)$, where $T$ is a binary tree and $\tau: T \rightarrow \mathbb{R}^{d}$ is a node function that assigns a test vector in $\mathbb{R}^{d}$ to every $t \in T$. Every vector $x \in \mathbb{R}^{d}$ is associated, through a sequence of binary comparisons, with a descending path in $T$ : beginning at the root, and at each subsequent internal node of $T, x$ moves to that child of its current node whose test vector is nearest to $x$ in Euclidean distance. In case of ties, $x$ moves to the left child of its current node. For each $t \in T$ let $U_{t}$ contain those vectors $x$ whose path includes $t$. Then $U_{t}=\mathbb{R}^{d}$ when $t$ is the root node of $T$, and if $t$ is an internval node, then $U_{t}$ is split between its children by the hyperplane that forms the perpendicular bisector of their test vectors. If $t$ is at distance $k$ from the root, then $U_{t}$ is a polytope having at most $k$ faces. The partition generated by $(T, \tau)$ is comprised of the sets $U_{t}$ associated with the leaves (terminal nodes) of $T$.

If at each internal node of $T$ the comparison between the test vectors labeling its children is based on a single coordinate of $x$, then each cell of the resulting partition is a $d$-dimensional rectangle. Tree-structured partitions of this sort, based on axis-parallel splits, are the basis for the regression trees considered by Breiman, Friedman, Olshen, and Stone (1984).

Let $\mathscr{T}(k)$ contain all the tree-structured partitions generated by binary trees $T$ having $k$ leaves. Clearly $m(\mathscr{T}(k))=k$. Consider an internal node $t$ whose children are assigned test vectors $u$ and $v$. Finding the test vector closest to $x$ is equivalent to testing the membership of $x$ in a closed half-space that is bounded by the perpendicular bisector of $u$ and $v$. A binary tree $T$ with $k$ leaves has $k-1$ internal nodes, and therefore each partition $(T, \tau)$ is based on at most $k-1$ intersecting half-spaces. As each of these half-spaces can dichotomize $n$ points in at most $n^{d}$ ways, their intersection can partition $n$ points in at most $n^{(k-1) d}$ ways. Thus $\Delta_{n}^{*}(\mathscr{T}(k)) \leq h^{(k-1) d}$.

Our principal result is stated below. Analogous conditions for the consistency of histogram classification and density estimates were established by

Lugosi and Nobel (1995). Note that the range of a partitioning rule is a family of partitions of $\mathbb{R}^{d}$.

Theorem 1. Let $\psi_{1}, \psi_{2}, \ldots$ be fixed partitioning rules, and let $\Pi_{n}$ be the range of $\psi_{n}$. Let the histogram regression estimate $\hat{r}_{n}$ be defined using $\psi_{n}$ and $T_{n}$ as in (2). Suppose that as $n$ tends to infinity the following hold:
(a) $n^{-1} m\left(\Pi_{n}: V\right) \rightarrow 0$ for every compact set $V \subseteq \mathbb{R}^{d}$;
(b) $n^{-1} \log \Delta_{n}^{*}\left(\Pi_{n}\right) \rightarrow 0$;
(c) for every $\gamma>0$ and $\delta \in(0,1)$,

$$
\begin{equation*}
\inf _{S: P(S) \geq 1-\delta} P\left\{x: \operatorname{diam}\left(\psi_{n}[x] \cap S\right)>\gamma\right\} \rightarrow 0 \quad \text { w.p.1. } \tag{4}
\end{equation*}
$$

Then $\int\left|\hat{r}_{n}-r\right|^{2} d P \rightarrow 0$ with probability 1.
Thus a sequence of partitioning rules gives rise to consistent histogram regression estimates if the cells of the selected partitions shrink, and if the rules take values in a sequence of partition families whose restricted cell counts grow sublinearly and whose partitioning numbers grow subexponentially. Note that no assumptions are made on the distribution of the pair $(X, Y) \in \mathbb{R}^{d} \times[-K, K]$. The conclusions of the theorem remain valid when conditions (a) and (b) are replaced by their natural almost-sure equivalents.

Condition (a) of Theorem 1 is significantly weaker than the corresponding condition of Theorem A, which requires that each cell contain a minimum number of points. Condition (b) allows the partitions $\psi_{n}\left(T_{n}\right)$ to become more complex as the sample size $n$ increases. In particular, the cells of each partition need not be rectangles, or polytopes with a fixed number of sides. Conditions (a) and (b) of Theorem A can be shown to imply condition (b) of Theorem 1.

In many cases of interest, the combinatorial conditions of Theorem 1 can be incorporated, without need of supervision, into the design of partitioning rules. In this case the principal task of analysis becomes verification of the shrinking cell condition (c).
1.3. Outline. A Vapnik-Chervonenkis inequality for partition families, as well as several of its consequences, is established in the next section. The proof of Theorem 1 is presented in Section 3. In Section 4 it is shown that the conditions of Theorem 1 insure the consistency of partitioning regression estimates based on local polynomial fits of fixed degree. Sufficient conditions for the consistency of histogram conditional quantile estimates are presented in Section 5.

The conditions of Theorem 1 are applicable to a variety of partitioning rules. Cubic partitions with data-dependent offsets are considered in Section 6. Section 7 is devoted to regression estimates in one dimension based on $k$-thresholding, a generalization of ordinary $k$-spacing in which the partitioning rule depends on the response variables $Y_{i}$.

Multivariate clustering schemes provide natural partitioning rules for histogram regression. It is shown in Section 8 that estimates based on empirically optimal nearest-neighbor clustering schemes are consistent when the number of cluster centers grows with the size of the training set at an appropriate rate.

It is shown in Section 9 that empirically optimal regression trees are consistent when the size $k_{n}$ of the tree grows as $o(n / \log n)$.
2. Preliminary results. A nonempty collection $\mathscr{F}$ of measurable functions $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is said to be uniformly bounded with envelope $L$ if $|f(x)| \leq L$ for every $x \in \mathbb{R}^{d}$ and every $f \in \mathscr{F}$. For each $\varepsilon>0$ and each sequence of vectors $x_{1}, \ldots, x_{n}$, the covering number $N\left(x_{1}^{n}, \varepsilon, \mathscr{F}\right)$ is the size of the smallest collection $\mathscr{G}$ such that

$$
\begin{equation*}
\min _{g \in \mathscr{G}} \frac{1}{n} \sum_{i=1}^{n}\left|f\left(x_{i}\right)-g\left(x_{i}\right)\right|<\varepsilon \tag{5}
\end{equation*}
$$

for every $f \in \mathscr{F}$. Any collection $\mathscr{G}$ satisfying (5) is said to be an $\varepsilon$-cover for $\mathscr{F}$ on $x_{1}^{n}$. If no finite $\varepsilon$-cover exists, then $N\left(x_{1}^{n}, \varepsilon, \mathscr{F}\right)=\infty$. A class $\mathscr{F}$ will be called nice if it is uniformly bounded and there is a function $\phi: \mathbb{R} \rightarrow \mathbb{R}$ such that

$$
\begin{equation*}
N\left(x_{1}^{n}, \varepsilon, \mathscr{F}\right) \leq \phi(\varepsilon) \tag{6}
\end{equation*}
$$

for every $\varepsilon>0$, every $n$ and every sequence $x_{1}, \ldots, x_{n} \in \mathbb{R}^{d}$. When (6) holds, $\phi$ will be said to majorize the covering numbers of $\mathscr{F}$.

Definition. Given a class of real-valued functions $\mathscr{G}$ on $\mathbb{R}^{d}$ and a partition family $\Pi$, define

$$
\mathscr{G} \circ \Pi=\left\{f=\sum_{A_{j} \in \pi} g_{j} I_{A_{j}}: \pi=\left\{A_{j}\right\} \in \Pi, g_{j} \in \mathscr{G}\right\}
$$

Each function $f$ in $\mathscr{G} \circ \Pi$ is obtained by applying a different function $g \in \mathscr{G}$ within each cell of a partition $\pi \in \Pi$.

Consider a sequence $X_{1}, X_{2}, \ldots \in \mathbb{R}^{d}$ of independent random vectors having a common distribution $P$. Let $\hat{P}_{n}$ be the empirical distribution of $X_{1}, \ldots, X_{n}$. For every bounded measurable function $f: \mathbb{R} \rightarrow \mathbb{R}$, define

$$
P f=\int f(x) d P(x)
$$

and

$$
\hat{P}_{n} f=\frac{1}{n} \sum_{i=1}^{n} f\left(X_{i}\right)
$$

Our results rely on exponential inequalities concerning the uniform deviation of $P$ and $\hat{P}_{n}$ over suitable classes of functions. Given a uniformly bounded class $\mathscr{F}$, define

$$
\begin{equation*}
\Lambda_{n}(\mathscr{F})=\sup _{f \in \mathscr{F}}\left|\hat{P}_{n} f-P f\right| \tag{7}
\end{equation*}
$$

If $\mathscr{H}$ is a uniformly bounded class of functions $h: \mathbb{R}^{d} \times[-K, K] \rightarrow \mathbb{R}$, then define

$$
\tilde{\Lambda}_{n}(\mathscr{H})=\sup _{h \in \mathscr{H}}\left|\frac{1}{n} \sum_{i=1}^{n} h\left(X_{i}, Y_{i}\right)-E h(X, Y)\right| .
$$

Remark. To ensure measurability of the suprema $\Lambda_{n}(\mathscr{F})$, we assume that every class $\mathscr{F}$ under consideration contains a countable subclass $\mathscr{F}_{0}$ with the property that every function in $\mathscr{F}$ is the pointwise limit of a sequence of functions in $\mathscr{F}_{0}$. We extend this condition to partition families by viewing each partition $\pi$ as a mapping from $\mathbb{R}^{d}$ into the set of natural numbers.

Proposition 1. Let $\mathscr{G}$ be class of function of $\mathbb{R}^{d}$ whose covering numbers are majorized by $\phi(\cdot)$, and let $\Pi$ be any partition family with $m(\Pi)<\infty$. For each sequence $x_{1}, \ldots, x_{n} \in \mathbb{R}^{d}$ and every $\varepsilon>0$,

$$
N\left(x_{1}^{n}, \varepsilon, \mathscr{G} \circ \Pi\right) \leq \Delta\left(x_{1}^{n}, \Pi\right) \phi(\varepsilon)^{m(\Pi)} \leq \Delta_{n}^{*}(\Pi) \phi(\varepsilon)^{m(\Pi)} .
$$

Proof. Fix $x_{1}, \ldots, x_{n} \in \mathbb{R}^{d}$ and $\varepsilon>0$. Call two elements $\pi, \pi^{\prime} \in \Pi$ equivalent if they induce the same partition of $x_{1}, \ldots, x_{n}$. If $f \in \mathscr{G} \circ \Pi$, then there is a partition $\pi=\left\{A_{j}\right\} \in \Pi$ and functions $g_{j} \in \mathscr{G}$ such that

$$
\begin{equation*}
f=\sum_{A_{j} \in \pi} g_{j} I_{A_{j}} . \tag{8}
\end{equation*}
$$

Let $\mathscr{F}_{j}$ be an $\varepsilon$-cover for $\mathscr{G}$ on $A_{j}^{\prime}=\left\{x_{1}, \ldots, x_{n}\right\} \cap A_{j}$ such that $\left|\mathscr{F}_{j}\right| \leq \phi(\varepsilon)$. To each function $g_{j}$ appearing in (8) there corresponds an approximating function $f_{j} \in \mathscr{F}_{j}$ such that

$$
\frac{1}{n_{j}} \sum_{x_{i} \in A_{j}^{\prime}}\left|g_{j}\left(x_{i}\right)-f_{j}\left(x_{i}\right)\right|<\varepsilon
$$

where $n_{j}=\left|A_{j}^{\prime}\right|$. If we define $f^{\prime}=\sum_{A_{j} \in \pi} f_{j} I_{A_{j}}$, then it is easy to see that

$$
\frac{1}{n} \sum_{i=1}^{n}\left|f\left(x_{i}\right)-f^{\prime}\left(x_{i}\right)\right|<\varepsilon .
$$

When the functions $f_{j} \in \mathscr{F}_{j}$ are suitably chosen, every function $\tilde{f} \in \mathscr{G} \circ \Pi$ defined in terms of a partition equivalent to $\pi$ can approximated by a similar estimate $f^{\prime}$. Thus the collection of all such functions $\tilde{f}$ can be covered on $x_{1}^{n}$ by no more than $\prod_{j=1}^{|\pi|}\left|\mathscr{F}_{j}\right| \leq \phi(\varepsilon)^{|\pi|}$ approximating functions. As the partitions in $\Pi$ fall into at most $\Delta\left(x_{1}^{n}, \Pi\right)$ equivalence classes, the result follows.

An application of the basic Vapnik-Chervonenkis inequality for classes of functions [cf. Vapnik and Chervonenkis (1981) or Pollard (1984)] gives the following bound.

LEMMA 1. Let $\mathscr{G}$ be a class of functions with envelope $L$ whose covering numbers are majorized by $\phi(\cdot)$. If $\Pi$ is a partition family for which $m(\Pi)<\infty$, then, for every $t>0$,

$$
\mathbb{P}\left\{\Lambda_{n}(\mathscr{G} \circ \Pi)>t\right\} \leq \Delta_{n}^{*}(\Pi) \phi(t)^{m(\Pi)} \exp \left[\frac{-n t^{2}}{32 L^{2}}\right]
$$

Proposition 2. For every uniformly bounded class $\mathscr{G}$ and every partition family П,

$$
\sup _{g \in \mathscr{G}} \sup _{\pi \in \Pi} \sum_{A \in \pi}\left|\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right| \leq \Lambda_{n}\left(\mathscr{G}^{\prime} \circ \Pi\right)
$$

where $\mathscr{G}^{\prime}=\mathscr{G} \cup(-\mathscr{G})$ contains every function $g \in \mathscr{G}$ and its additive inverse.
Proof. Fix a function $g \in \mathscr{G}$ and a parition $\pi \in \Pi$. Given $X_{1}=x_{1}, \ldots$, $X_{n}=x_{n}$, define
$\pi_{1}=\left\{A \in \pi: \hat{P}_{n}\left(g I_{A}\right) \geq P\left(g I_{A}\right)\right\} \quad$ and $\quad \pi_{2}=\left\{A \in \pi: \hat{P}_{n}\left(g I_{A}\right)<P\left(g I_{A}\right)\right\}$.
Then it is easy to see that

$$
\begin{aligned}
& \sum_{A \in \pi}\left|\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right| \\
& \quad=\left|\sum_{A \in \pi_{1}}\left(\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right)-\sum_{A \in \pi_{2}}\left(\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right)\right| \\
& \quad=\left|\hat{P}_{n} f-P f\right|
\end{aligned}
$$

where $f=\sum_{A \in \pi_{1}} g I_{A}+\sum_{A \in \pi_{2}}(-g) I_{A}$ is an element of $\mathscr{G}^{\prime} \circ \Pi$. Consequently,

$$
\sum_{A \in \pi}\left|\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right| \leq \Lambda_{n}\left(\mathscr{G}^{\prime} \circ \Pi\right)
$$

and the result follows as $g \in \mathscr{G}$ and $\pi \in \Pi$ were arbitrary.
The following inequality extends Lemma 1 of Lugosi and Nobel (1995). Although it will not be needed in what follows, it may be of independent interest.

LEMMA 2. Let $\mathscr{G}$ be a class of functions with envelope $L$ whose covering numbers are majorized by $\phi(\cdot)$. If $\Pi$ is a partition family with $m(\Pi)<\infty$, then, for every $t>0$,
$\mathbb{P}\left\{\sup _{g \in \mathscr{G}} \sup _{\pi \in \Pi} \sum_{A \in \pi}\left|\hat{P}_{n}\left(g I_{A}\right)-P\left(g I_{A}\right)\right|>t\right\} \leq 2^{m(\Pi)} \Delta_{n}^{*}(\Pi) \phi(t)^{m(\Pi)} \exp \left[\frac{-n t^{2}}{32 L^{2}}\right]$.
Proof. As $N\left(x_{1}^{n}, t, \mathscr{G}^{\prime}\right) \leq 2 N\left(x_{1}^{n}, t, \mathscr{G}^{\prime}\right) \leq 2 \phi(t)$, the stated inequality follows immediately from Lemma 1 and Proposition 2.

Lemma 3. Let $\mathscr{G}$ be a nice class of functions, and let $\Pi_{1}, \Pi_{2}, \ldots$ be a sequence of partition families. If

$$
\begin{equation*}
n^{-1} m\left(\Pi_{n}: V\right) \rightarrow 0 \tag{9}
\end{equation*}
$$

for every compact set $V \subset \mathbb{R}^{d}$, and

$$
\begin{equation*}
n^{-1} \log \Delta_{n}^{*}\left(\Pi_{n}\right) \rightarrow 0, \tag{10}
\end{equation*}
$$

then

$$
\begin{equation*}
\Lambda_{n}\left(\mathscr{G} \circ \Pi_{n}\right) \rightarrow 0 \tag{11}
\end{equation*}
$$

with probability 1.
Proof. If the unrestricted cell counts are such that $n^{-1} m\left(\Pi_{n}\right) \rightarrow 0$, then (11) is an immediate consequence of Lemma 1 and the Borel-Cantelli lemma. In the more general case, fix $\varepsilon>0$ and let $L<\infty$ be an envelope for $\mathscr{G}$. Select a compact set $V \subseteq \mathbb{R}^{d}$ such that $L \cdot P\left(V^{c}\right) \leq \varepsilon$, and define the class of functions $\mathscr{E}_{V}=\left\{g I_{V}: g \in \mathscr{G}\right\}$, which is readily seen to be nice. For each partition $\pi=\left\{A_{j}\right\} \in \Pi_{n}$ define its restriction $\pi^{\prime}=\left\{A_{j} \cap V\right\} \cup\left\{V^{c}\right\}$, and let

$$
\Pi_{n}^{\prime}=\left\{\pi^{\prime}: \pi \in \Pi_{n}\right\} .
$$

It is easily verified that $\Delta_{n}^{*}\left(\Pi_{n}^{\prime}\right) \leq \Delta_{n}^{*}\left(\Pi_{n}\right)$ and that $m\left(\Pi_{n}^{\prime}\right) \leq m\left(\Pi_{n}: V\right)+2$. By virtue of the unrestricted case above, $\Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}^{\prime}\right) \rightarrow 0$ with probability 1. Now note that

$$
\left|\Lambda_{n}\left(\mathscr{G}, \Pi_{n}\right)-\Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}\right)\right| \leq L \hat{P}_{n}\left(V^{c}\right)+L P\left(V^{c}\right)
$$

and as $\Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}\right)=\Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}^{\prime}\right)$ for each $n$,

$$
\begin{aligned}
\limsup _{n \rightarrow \infty} \Lambda_{n}\left(\mathscr{G}, \Pi_{n}\right) & \leq \limsup _{n \rightarrow \infty} \Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}\right)+2 \varepsilon \\
& =\limsup _{n \rightarrow \infty} \Lambda_{n}\left(\mathscr{G}_{V}, \Pi_{n}^{\prime}\right)+2 \varepsilon \\
& =2 \varepsilon
\end{aligned}
$$

with probability 1 . Since $\varepsilon>0$ was arbitrarily, the result follows.
Lemma 4. Let $\mathscr{G}$ be a nice class of functions with envelope $L$, and let $\Pi_{1}, \Pi_{2}, \ldots$ be partition families satisfying (9) and (10). Let $\left\{\hat{f}_{n}\right\}$ and $\left\{\hat{g}_{n}\right\}$ be regression estimates such that, for each $n$ and each training sequence $T_{n}$, the following hold:
(a) $\hat{f}_{n}\left(\cdot, T_{n}\right)$ and $\hat{g}_{n}\left(\cdot, T_{n}\right)$ lie in $\mathscr{G} \circ \Pi_{n}$;
(b) $\sum_{i=1}^{n}\left(\hat{f}_{n}\left(X_{i}\right)-Y_{i}\right)^{2} \leq \sum_{i=1}^{n}\left(\hat{g}_{n}\left(X_{i}\right)-Y_{i}\right)^{2}$.

Then $\left\{\hat{f}_{n}\right\}$ is strongly consistent if $\left\{\hat{g}_{n}\right\}$ is strongly consistent.
Proof. It is well known and easy to verify that, for every bounded function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$,

$$
\begin{equation*}
P|f-r|^{2}=E|f(X)-Y|^{2}-E|r(X)-Y|^{2} . \tag{12}
\end{equation*}
$$

For each element $\pi=\left\{A_{j}\right\} \in \Pi_{n}$, define an associated partition $\tilde{\pi}=\left\{A_{j} \times\right.$ $[-K, K]\}$ of $\mathbb{R}^{d} \times[-K, K]$, and let $\tilde{\Pi}_{n}=\left\{\tilde{\pi}: \pi \in \Pi_{n}\right\}$. It is readily verified that

$$
\begin{equation*}
\Delta_{n}^{*}\left(\tilde{\Pi}_{n}\right)=\Delta_{n}^{*}\left(\Pi_{n}\right) \quad \text { and } \quad m\left(\tilde{\Pi}_{n}: V \times[-K, K]\right)=m\left(\Pi_{n}: V\right) \tag{13}
\end{equation*}
$$

for every compact subset $V$ of $\mathbb{R}^{d}$. Let $\tilde{\mathscr{G}}$ contain all those functions $h$ : $\mathbb{R}^{d} \times[-K, K] \rightarrow \mathbb{R}$ of the form $h(x, y)=(g(x)-y)^{2}$, where $g \in \mathscr{G}$. Then $\tilde{\mathscr{G}}$ is nice, and it follows from Lemma 3 and equations (13) that

$$
\begin{equation*}
\sup _{f \in \mathscr{G} \circ \Pi_{n}}\left|\frac{1}{n} \sum_{i=1}^{n}\left(f\left(X_{i}\right)-Y_{i}\right)^{2}-E(f(X)-Y)^{2}\right|=\tilde{\Lambda}_{n}\left(\tilde{\mathscr{G}} \circ \tilde{\Pi}_{n}\right) \rightarrow 0 \tag{14}
\end{equation*}
$$

with probability 1 . Now consider the estimates $\left\{\hat{f}_{n}\right\}$ and $\left\{\hat{f}_{n}\right\}$. By virtue of (14) and the assumptions above,

$$
\begin{aligned}
P\left|\hat{f}_{n}-r\right|^{2} & =E\left|\hat{f}_{n}(X)-Y\right|^{2}-E|r(X)-Y|^{2} \\
& \leq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{f}_{n}\left(X_{i}\right)-Y_{i}\right)^{2}-E|r(X)-Y|^{2}+\tilde{\Lambda}_{n}\left(\tilde{\mathscr{G}} \circ \tilde{\Pi}_{n}\right) \\
& \leq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{g}_{n}\left(X_{i}\right)-Y_{i}\right)^{2}-E|r(X)-Y|^{2}+\tilde{\Lambda}_{n}\left(\tilde{\mathscr{G}} \circ \tilde{\Pi}_{n}\right) \\
& \leq E\left|\hat{g}_{n}(X)-Y\right|^{2}-E|r(X)-Y|^{2}+2 \tilde{\Lambda}_{n}\left(\hat{\mathscr{G}} \circ \tilde{\Pi}_{n}\right) \\
& =P\left|\hat{g}_{n}-r\right|^{2}+2 \tilde{\Lambda}_{n}\left(\hat{\mathscr{G}} \circ \tilde{\Pi}_{n}\right) .
\end{aligned}
$$

If the sequence $\left\{\hat{g}_{n}\right\}$ is strongly consistent, then the last term above tends to zero with probability 1 .

## 3. Proof of Theorem 1.

Definition. A sequence $Z_{1}, Z_{2}, \ldots$ of random variables defined on the same probability space as $(X, Y)$ is said to be of order $o^{*}(1)$, written $Z_{n}=$ $o^{*}(1)$, if as $n$ tends to infinity $Z_{n} \rightarrow 0$ with probability 1 .

Proof of Theorem 1. Define an auxiliary function

$$
\begin{equation*}
\tilde{r}_{n}(x)=\frac{\sum_{i=1}^{n} r\left(X_{i}\right) I\left\{X_{i} \in \psi_{n}[x]\right\}}{\sum_{i=1}^{n} I\left\{X_{i} \in \psi_{n}[x]\right\}} . \tag{15}
\end{equation*}
$$

Note that $\left|\tilde{r}_{n}\right| \leq K$ and that $\tilde{r}_{n}$ is piecewise constant on the cell of $\psi_{n}$. By an obvious upper bound,

$$
\begin{equation*}
P\left|r-\hat{r}_{n}\right|^{2} \leq 2 P\left|\tilde{r}_{n}-\hat{r}_{n}\right|^{2}+2 P\left|r-\tilde{r}_{n}\right|^{2} . \tag{16}
\end{equation*}
$$

The first term measures the variance of $\hat{r}_{n}$, while the second term measures its bias. We show in turn that each is of order $o^{*}(1)$.

Consider the first term in (16). An easy application of Proposition 2 and Lemma 3 shows that

$$
\begin{equation*}
\sup _{\pi \in \Pi_{n}} \sum_{A \in \pi}\left|\hat{P}_{n}(A)-P(A)\right|=o^{*}(1) . \tag{17}
\end{equation*}
$$

For each cell $A \in \psi_{n}\left(T_{n}\right)$ let $n_{A}=n \hat{P}_{n}(A)$ be the number of vectors $X_{i}$ in $A$. By virtue of (17) and the definition of $\tilde{r}_{n}$,

$$
\begin{align*}
P\left|\hat{r}_{n}-\tilde{r}_{n}\right|^{2} & \leq 2 K \cdot \sum_{A \in \psi_{n}}\left|\frac{1}{n_{A}} \sum_{X_{i} \in A}\left(Y_{i}-r\left(X_{i}\right)\right)\right| P(A) \\
& \leq 2 K \cdot \sum_{A \in \psi_{n}}\left|\frac{1}{n_{A}} \sum_{X_{i} \in A}\left(Y_{i}-r\left(X_{i}\right)\right)\right| \hat{P}_{n}(A)+o^{*}(1) \\
& =2 K \cdot \sum_{A \in \psi_{n}}\left|\frac{1}{n} \sum_{X_{i} \in A}\left(Y_{i}-r\left(X_{i}\right)\right)\right|+o^{*}(1)  \tag{18}\\
& \leq 2 K \cdot \sup _{\pi \in \Pi_{n}} \sum_{A \in \pi}\left|\frac{1}{n} \sum_{X_{i} \in A}\left(Y_{i}-r\left(X_{i}\right)\right)\right|+o^{*}(1) .
\end{align*}
$$

Let $Q$ be the joint distribution of ( $X, Y$ ), and let $\hat{Q}_{n}$ be the empirical distribution of $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. For each element $\pi=\left\{A_{j}\right\} \in \Pi_{n}$, define an associated partition $\tilde{\pi}=\left\{A_{j} \times[-K, K]\right\}$ of the set $\mathbb{R}^{d} \times[-K, K]$, and let $\tilde{\Pi}_{n}=\left\{\tilde{\pi}: \pi \in \Pi_{n}\right\}$. If $h(x, y)=y-r(x)$ is defined for $(x, y) \in \mathbb{R}^{d} \times[-K, K]$, then $h$ is bounded and $\int_{A \times \mathbb{R}} h(x, y) d Q=0$ for every measurable subset $A$ of $\mathbb{R}^{d}$. Rewriting (18) in terms of the distribution $Q$, the family $\tilde{\Pi}_{n}$ and $h$ gives

$$
\begin{aligned}
P\left|\hat{r}_{n}-\tilde{r}_{n}\right|^{2} & \leq 2 K \cdot \sup _{\tilde{\pi} \in \tilde{\Pi}_{n}} \sum_{B \in \tilde{\pi}}\left|\hat{Q}_{n}\left(h I_{B}\right)-Q\left(h I_{B}\right)\right|+o^{*}(1) \\
& \leq 2 K \cdot \tilde{\Lambda}_{n}\left(\{h,-h\} \circ \tilde{\Pi}_{n}\right)+o^{*}(1),
\end{aligned}
$$

where the second inequality follows from Proposition 2. By virtue of (13) and the conditions of the theorem, the first term above is of order $o^{*}(1)$. Thus $P\left|\hat{r}_{n}-\tilde{r}_{n}\right|=o^{*}(1)$.

Consider now the second term of (16). Fix $\varepsilon>0$ and let $g: \mathbb{R}^{d} \rightarrow[-K, K]$ be a uniformly continuous function such that $P|r-g|^{2}<\varepsilon$. Let $\gamma>0$ be chosen so that $\left|g\left(x_{1}\right)-g\left(x_{2}\right)\right|<\varepsilon^{1 / 2}$ whenever $\left\|x_{1}-x_{2}\right\|<\gamma$. Fix $\delta>0$ so small that $8 K^{2} \delta \leq \varepsilon$, and let $S$ be any set satisfying $P(S) \geq 1-\delta$. Define $\bar{g}_{n}$ by averaging $g$ over those $X_{i}$ within the cells of $\psi_{n} \cap S$,

$$
\bar{g}_{n}(x)=\frac{\sum_{i=1}^{n} g\left(X_{i}\right) I\left\{X_{i} \in \psi_{n}[x] \cap S\right\}}{\sum_{i=1}^{n} I\left\{X_{i} \in \psi_{n}[x] \cap S\right\}} .
$$

Both $\tilde{r}_{n}$ and $\bar{g}_{n}$ are constant on the cells of $\psi_{n}$. In particular, the definition of $\tilde{r}_{n}$ insures that

$$
\sum_{X_{i} \in A}\left(r\left(X_{i}\right)-\tilde{r}_{n}\left(X_{i}\right)\right)^{2} \leq \sum_{X_{i} \in A}\left(r\left(X_{i}\right)-\bar{g}_{n}\left(X_{i}\right)\right)^{2}
$$

for each cell $A \in \psi_{n}$, so that, for each training set $T_{n}$,

$$
\begin{equation*}
\hat{P}_{n}\left|r-\tilde{r}_{n}\right|^{2} \leq \hat{P}_{n}\left|r-\bar{g}_{n}\right|^{2} . \tag{19}
\end{equation*}
$$

Let $\mathscr{G}$ contain all the functions of the form $g(x)=(r(x)-a)^{2}$, where $a \in[-K, K]$. Then $\mathscr{G}$ is nice, and it is evident that $\left|r-\tilde{r}_{n}\right|^{2}$ and $\left|r-\tilde{g}_{n}\right|^{2}$ are contained in $\mathscr{G} \circ \Pi_{n}$ for each $n$. Applying Lemma 3 to the sequence $\left\{\mathscr{G} \circ \Pi_{n}\right\}$, one may deduce that

$$
\begin{align*}
& \hat{P}_{n}\left|r-\tilde{r}_{n}\right|^{2}-P\left|r-\tilde{r}_{n}\right|^{2}=o^{*}(1) \quad \text { and } \\
& \hat{P}_{n}\left|r-\tilde{g}_{n}\right|^{2}-P\left|r-\bar{g}_{n}\right|^{2}=o^{*}(1) . \tag{20}
\end{align*}
$$

It follows from (19), (20) and the choice of $S$ that

$$
\begin{aligned}
P\left|r-\tilde{r}_{n}\right|^{2} & \leq \hat{P}_{n}\left|r-\tilde{r}_{n}\right|^{2}+o^{*}(1) \\
& \leq \hat{P}_{n}\left|r-\bar{g}_{n}\right|^{2}+o^{*}(1) \\
& \leq P\left|r-\bar{g}_{n}\right|^{2}+o^{*}(1) \\
& \leq 2 P|r-g|^{2}+2 P\left|g-\bar{g}_{n}\right|^{2}+o^{*}(1) \\
& \leq 2 P\left|g-\bar{g}_{n}\right|^{2}+2 \varepsilon+o^{*}(1) \\
& =2 P\left|g-\bar{g}_{n}\right|^{2} I_{S}+2 P\left|g-\bar{g}_{n}\right|^{2} I_{S^{c}}+2 \varepsilon+o^{*}(1) \\
& \leq 2 P\left|g-\bar{g}_{n}\right|^{2} I_{S}+3 \varepsilon+o^{*}(1)
\end{aligned}
$$

If $x \in S$ and $\operatorname{diam}\left(\psi_{n}[x] \cap S\right)<\gamma$, then $\left|g(x)-\bar{g}_{n}(x)\right|^{2}<\varepsilon$. One may therefore conclude from the last inequality that

$$
P\left|r-\tilde{r}_{n}\right|^{2} \leq 8 K^{2} \cdot P\left\{x: \operatorname{diam}\left(\psi_{n}[x] \cap S\right) \geq \gamma\right\}+4 \varepsilon+o^{*}(1)
$$

As the right-hand side depends on $S$ only through its probability,

$$
P\left|r-\tilde{r}_{n}\right|^{2} \leq \inf _{S: P(S) \geq 1-\delta} 8 K^{2} \cdot P\left\{\operatorname{diam}\left(\psi_{n}[x] \cap S\right)>\gamma\right\}+4 \varepsilon+o^{*}(1)
$$

Condition (c) of the theorem guarantees that $P\left|r-\tilde{r}_{n}\right|^{2} \leq 4 \varepsilon+o^{*}(1)$, and the result follows as $\varepsilon>0$ was arbitrary.
4. Local fitting with truncated polynomials. The histogram estimate $\hat{r}_{n}$ is piecewise constant on the cells of $\psi_{n}\left(T_{n}\right)$ and has discontinuities along the boundaries of these cells. If the regression function $r(\cdot)$ is known to be smooth, or if the sample size is large, then more sophisticated local estimates may be appropriate.

The results of Theorem 1 may be extended in a natural way to regression estimates based on local polynomial fits. Fitting a suitably truncated polynomial within each cell of $\psi_{n}$ gives estimates with better approximation capabilities and potentially better performance. Discontinuities along cell boundaries remain, but may be corrected if necessary by employing weighted
averages of estimates from different cells, as in Chaudhuri, Huang, Loh and Yao (1994).

For each vector $u=\left(u_{1}, \ldots, u_{d}\right) \in \mathbb{R}^{d}$ and each sequence $\alpha=\left(\alpha_{1}, \ldots, \alpha_{d}\right)$ of nonnegative integers, let $u^{\alpha}=u_{1}^{\alpha_{1}} \cdots u_{d}^{\alpha_{d}}$ and $|\alpha|=\alpha_{1}+\cdots+\alpha_{d}$. Fix an integer $k \geq 0$ and let

$$
\mathscr{G}=\left\{g(x)=\sum_{|\alpha| \leq k} c(\alpha) x^{\alpha}: c(\alpha) \in \mathbb{R}\right\}
$$

be the class of $k$ th-order multivariate polynomials on $\mathbb{R}^{d}$. Set $L=K b$ for some $b \geq 1$, and define the class

$$
\tilde{\mathscr{G}}=\{g(\cdot) \wedge l \vee(-l): g \in \mathscr{G}, 0 \leq l \leq L\}
$$

of truncated polynomials which has envelope $L$. As $\mathscr{G}$ is a finite-dimensional vector space of real-valued functions, it follows from Lemma 28 of Pollard [(1984), Chapter 2] that $\mathscr{G}$ is a VC-graph class, and an easy argument shows that the same is true for $\tilde{\mathscr{G}}$. Lemma 25 of Pollard [(1984), Chapter 2] shows that $\mathscr{G}$ is nice.

Given a partitioning rule $\psi_{n}$ and a training set $T_{n}$, we construct a piecewise polynomial regression estimate by fitting a suitable function $g \in \tilde{\mathscr{G}}$ within each cell of $\psi_{n}\left(T_{n}\right)$. For each cell $A \in \psi_{n}\left(T_{n}\right)$, let

$$
\begin{equation*}
g_{A}=\underset{g \in \mathscr{G}}{\arg \min } \sum_{X_{i} \in A}\left(g\left(X_{i}\right)-Y_{i}\right)^{2} \tag{21}
\end{equation*}
$$

be the best $k$ th-order polynomial fit to those pairs ( $X_{i}, Y_{i}$ ) for which $X_{i} \in A$, and set

$$
l_{A}=b \cdot \max \left\{\left|Y_{i}\right|: X_{i} \in A\right\} .
$$

If the range $[-K, K]$ of the response variables $Y_{i}$ is known, then define the estimate

$$
\begin{equation*}
\hat{f}_{n}(x)=\sum_{A \in \psi_{n}} g_{A}(x) \wedge K \vee(-K) \cdot I\{x \in A\} ; \tag{22}
\end{equation*}
$$

otherwise, define

$$
\begin{equation*}
\hat{f}_{n}(x)=\sum_{A \in \psi_{n}} g_{A}(x) \wedge l_{A} \vee\left(-l_{A}\right) \cdot I\{x \in A\} \tag{23}
\end{equation*}
$$

In either case, $\hat{f_{n}}(x) \in \tilde{\mathscr{G}} \circ \Pi_{n}$. Moreover, the optimality of $g_{A}$ and the choice of truncation level insures that

$$
\sum_{i=1}^{n}\left(\hat{f}_{n}-Y_{i}\right)^{2} \leq \sum_{i=1}^{n}\left(\hat{r}_{n}-Y_{i}\right)^{2}
$$

for every training set $T_{n}$. The consistency of the estimates $\left\{\hat{f}_{n}\right\}$ follows immediately from Theorem 1 and Lemma 4. Note that truncation eliminates the need for regularity assumptions concerning the local least squares fit within each cell.

THEOREM 2. Let $\psi_{1}, \psi_{2}, \ldots$ be a sequence of partitioning rules satisfying conditions (a)-(c) of Theorem 1. If regression estimates $\hat{f}_{n}$ are defined using truncated local polynomial fits as in (22) or (23), then $P\left|\hat{f}_{n}-r\right|^{2} \rightarrow 0$ with probability 1.

When it is desirable to do so, we may center the polynomial fit within each cell, replacing $g\left(X_{i}\right)$ in (21) by $g\left(X_{i}-x_{0}\right)$, where $x_{0}$ is the average of those vectors lying in $A$, or some other centrally lcoated point of $A$.
5. Estimation of conditional medians. For each $x \in \mathbb{R}^{d}$ and $a \in$ [ $-K, K$ ], the conditional cumulative distribution function of $Y$ given $X=x$ is defined by

$$
F(a, x)=\mathbb{P}\{Y \leq a \mid X=x\}
$$

The conditional median of $Y$ given $X=x$ is defined by

$$
m(x)=\inf \left\{a: F(a, x) \geq \frac{1}{2}\right\}
$$

Here we study histogram estimates of $m(\cdot)$ that are based on local order statistics of the response variables.

Given a partitioning rule $\psi_{n}$ and a training set $T_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$, define for each $x$ a histogram estimate

$$
\hat{m}_{n}(x)=\inf \left\{a: \hat{F}_{n}(a, x) \geq \frac{1}{2}\right\}
$$

of the conditional median $m(x)$, in terms of the corresponding estimate

$$
\hat{F}_{n}(a, x)=\frac{\sum_{i=1}^{n} I\left\{Y_{i} \leq a\right\} I\left\{X_{i} \in \psi_{n}[x]\right\}}{\sum_{i=1}^{n} I\left\{X_{i} \in \psi_{n}[x]\right\}}
$$

of the conditional cumulative distribution function. Then $\hat{m}_{n}(x)$ is just the $\lfloor(k+1) / 2]$ th-order statistic among the $k$ numbers $\left\{Y_{i}: X_{i} \in \psi_{n}[x]\right\}$. If $k$ is odd, then $\hat{m}_{n}(x)$ is the ordinary median of $\left\{Y_{i}: X_{i} \in \psi_{n}[x]\right\}$.

THEOREM 3. Let $\psi_{1}, \psi_{2}, \ldots$ be partitioning rules satisfying conditions (a)-(c) of Theorem 1. If, for P-almost every $x \in \mathbb{R}^{d}$,

$$
\begin{equation*}
\mathbb{P}\{m(x)<Y \leq m(x)+\varepsilon \mid X=x\}>0 \quad \text { for every } \varepsilon>0 \tag{24}
\end{equation*}
$$

then $P\left|\hat{m}_{n}-m\right| \rightarrow 0$ with probability 1 .
Proof. Condition (24) insures that $F(a, x)$ is increasing in each neighborhood of $a=m(x)$. Fix $\varepsilon>0$ and note that $P\left|\hat{m}_{n}-m\right| \leq 2 \varepsilon+2 K P\left(A_{n}\right)+$ $2 K P\left(B_{n}\right)$, where the sets

$$
A_{n}=\left\{x: m(x) \geq \hat{m}_{n}(x)+2 \varepsilon\right\} \quad \text { and } \quad B_{n}=\left\{x: \hat{m}_{n}(x) \geq m(x)+2 \varepsilon\right\}
$$

depend on the training set $T_{n}$ through $\hat{m}_{n}$. It is enough to show that $P\left(A_{n}\right)$ and $P\left(B_{n}\right)$ are of order $o^{*}(1)$.

Fix a number $a \in[-K, K]$ and define the indicator random variable $Y_{i}^{\prime}=I\left\{Y_{i} \leq a\right\}$ for each $i \geq 1$. The regressions function of $Y_{i}^{\prime}$ is $F(a, x)$, and
its histogram estimate is just $\hat{F}_{n}(a, x)$. Although $\psi_{n}$ is a function of $T_{n}$ rather than $\left(X_{1}, Y_{1}^{\prime}\right), \ldots,\left(X_{n}, Y_{n}^{\prime}\right)$, the analysis of Theorem 1 still applies: for each $a \in[-K, K]$,

$$
\begin{equation*}
\int\left|\hat{F}_{n}(a, x)-F(a, x)\right| d P(x) \rightarrow 0 \tag{25}
\end{equation*}
$$

with probability 1.
Select numbers $-(K+\varepsilon)=a_{0}<a_{1}<\cdots<a_{l}<a_{l+1}=(K+\varepsilon)$ such that $a_{i+1}-a_{i}<\varepsilon$, for $i=0, \ldots, l$. For each $x \in \mathbb{R}^{d}$, let $j(x)$ be the unique integer in $\{1, \ldots, l-1\}$ such that $m(x) \in\left[a_{j(x)}, a_{j(x)+1}\right)$, and define $h(x)=$ $\frac{1}{2}-F\left(a_{j(x)-1}, x\right)$. The definition of $m(x)$ insures that $P\{h>0\}=1$. If $x \in A_{n}$ then $\hat{m}_{n}(x) \leq m(x)-2 \varepsilon \leq a_{j(x)-1}$, and therefore

$$
\begin{equation*}
\hat{F}_{n}\left(a_{j(x)-1}, x\right)-F\left(a_{j(x)-1}, x\right) \geq \frac{1}{2}-F\left(a_{j(x)-1}, x\right)=h(x) . \tag{26}
\end{equation*}
$$

For each $\delta>0$ define $V_{\delta}=\{x: h(x) \geq \delta\}$. By virtue of (26),

$$
\begin{aligned}
\int \max _{0 \leq j \leq l}\left|\hat{F}_{n}\left(a_{j}, x\right)-F\left(a_{j}, x\right)\right| d P(x) & \geq \int_{A_{n}} h(x) d P(x) \\
& \geq \delta P\left(A_{n} \cap V_{\delta}\right) \\
& \geq \delta\left[P\left(A_{n}\right)-P\left(V_{\delta}^{c}\right)\right],
\end{aligned}
$$

and it follows from (25) that, for every $\delta>0$,

$$
\limsup _{n \rightarrow \infty} P\left(A_{n}\right) \leq P\left(V_{\delta}^{c}\right)
$$

with probability 1 . As $\delta \rightarrow 0$, the probability $P\left(V_{\delta}^{c}\right) \rightarrow P\{x: h(x)=0\}=0$, so that $P\left(A_{n}\right)=o^{*}(1)$. A similar analysis may be carried out for the events $\left\{B_{n}\right\}$ using the regularity condition (24).

Remark. An obvious modification of the preceding arguments establishes the consistency of histogram estimates for the conditional quantiles $m_{\alpha}(x)$, with $\alpha \in(0,1)$. One need only require that (24) hold at $P$-almost every point $m_{\alpha}(x)$.
6. Cubic partitions with data-dependent offsets. A cubic partition has congruent, rectangular cells whose dimensions are specified in advance of the data. Adding an offset vector to each cell of a cubic partition shifts the cells in the direction of the vector, while maintaining their axis-parallel orientation. An application of Theorem 1 yields conditions for the consistency of regression estimates based on cubic partitions with data-dependent offsets.

Fix a sequence $\left\{h_{j n}\right\}$ of positive numbers for each $j=1, \ldots, d$. For every $n \geq 1$, define the infinite cubic partition

$$
\theta_{n}=\left\{\left[\left(r_{1}-1\right) h_{1 n}, r_{1} h_{1 n}\right) \times \cdots \times\left[\left(r_{d}-1\right) h_{d n}, r_{d} h_{d n}\right): r_{1}, \ldots, r_{d} \in \mathbb{Z}\right\}
$$

having congruent, rectangular cells with edge lengths $h_{1 n}, \ldots, h_{d n}$. Let $\Pi_{n}=$ $\left\{\theta_{n}+c: c \in \mathbb{R}^{d}\right\}$ contain all the shifts of $\theta_{n}$ by an arbitrary offset $c$.

Suppose that $d=1$, in which case each cell of $\theta_{n}$ is an interval of length $h_{n}$. Fix $x_{1}, \ldots, x_{n} \in \mathbb{R}$. If $\pi \in \Pi_{n}$, then $\pi$ and $\pi+h_{n}$ are identical, so in assessing the partitioning number $\Delta\left(x_{1}^{n}, \Pi_{n}\right)$ it is enough to consider shifts $c \in\left[0, h_{n}\right.$ ) of length at most $h_{n}$. As $c$ increases from 0 to $h_{n}$, the partition induced by $\theta_{n}+c$ changes only if the boundary of a cell crosses one of the points $x_{1}, \ldots, x_{n}$. This happens exactly $n$ times, once for each $x_{i}$, and it follows that $\Delta\left(x_{1}^{n}, \Pi_{n}\right) \leq n$.

Suppose now that $d>1$. For each $j=1, \ldots, d$, let the family $\Pi_{n}^{j}$ be generated by shifts of the partition

$$
\theta_{n}^{j}=\left\{\mathbb{R} \times \cdots \times\left[\left(r_{j}-1\right) h_{j n}, r_{j} h_{j n}\right) \times \cdots \times \mathbb{R}: r_{j} \in \mathbb{Z}\right\},
$$

each of whose cells is a slab perpendicular to the $j$ th coordinate axis. By varying the components of $c$ one at a time, it can be shown that, for each sequence $x_{1}, \ldots, x_{n} \in \mathbb{R}^{d}$,

$$
\Delta\left(x_{1}^{n}, \Pi_{n}\right) \leq \Delta_{n}^{*}\left(\Pi_{n}^{1}\right) \cdots \Delta_{n}^{*}\left(\Pi_{n}^{d}\right) .
$$

It follows from the case $d=1$ above that $\Delta_{n}^{*}\left(\Pi_{n}^{1}\right) \leq n$, and therefore the partitioning numbers $\Delta_{n}^{*}\left(\Pi_{n}\right)$ satisfy condition (b) of Theorem 1.

The other conditions of Theorem 1 lead to restrictions on the constant $h_{j n}$ that are identical to those required for fixed cubic partitions. The shrinking cell condition (c) is satisfied if and only if

$$
\begin{equation*}
\lim _{n \rightarrow \infty} h_{j n} \rightarrow 0 \text { for } j=1, \ldots, d \tag{27}
\end{equation*}
$$

As for condition (a), assume without loss of generality that $V$ is a compact set of the form $[-t, t] \times \cdots \times[-t, t]$, with $t>0$. Each cell of $\pi \in \Pi_{n}$ has volume $\prod_{j=1}^{d} h_{j n}$, and if (27) holds, then

$$
m\left(\Pi_{n}: V\right) \leq \frac{(2 t+1)^{d}}{\prod_{j=1}^{d} h_{j n}}
$$

when $n$ is sufficiently large. Thus the restricted covering numbers grow sublinearly if

$$
\begin{equation*}
n \cdot \prod_{j=1}^{d} h_{j n} \rightarrow \infty \quad \text { as } n \rightarrow \infty \tag{28}
\end{equation*}
$$

The conclusion of the following theorem holds regardless of the joint distribution of $(X, Y) \in \mathbb{R}^{d} \times[-K, K]$.

Theorem 4. For each $n$, let the regression estimate $\hat{r}_{n}$ be based on a shifted version of the partition $\theta_{n}$. If (27) and (28) hold, then $P\left|r-\hat{r}_{n}\right|^{2} \rightarrow 0$ with probability 1.
7. $K$-threshold with variable weights. In this section we consider a $Y$-dependent variant of the $k$-spacing regression estimate of Parthasarathy and Battacharya (1961). Let $d=1$, so that $X_{1}, X_{2}, \ldots$ are real-valued, and
assume that the distribution $P$ of $X_{i}$ has a density with respect to Lebesgue measure. Let $F: \mathbb{R} \rightarrow[0, \infty)$ be a nonnegative weight function that is bounded away from zero and infinity on every compact set. The partition $\psi_{n}\left(T_{n}\right)$ is found by ordering $X_{1}, \ldots, X_{n}$ and then grouping them into intervals based on the weights $F\left(Y_{i}\right)$ of the corresponding response variables. In scanning from left to right, a new interval is begun when the running sum of the weights $F\left(Y_{i}\right)$ associated with $X_{i}$ in the current interval exceeds a preassigned threshold.

Let $k_{n}>0$ be a preassigned threshold value. Let $\rho$ be the unique permutation on $\{1, \ldots, n\}$ such that $X_{\rho(1)}<X_{\rho(2)}<\cdots<X_{\rho(n)}$. (Such a permutation exists with probability 1 as $P$ has a density.) Set $m_{1}=1$ and recursively define successive threshold times

$$
m_{r+1}=1+\min \left\{m \geq m_{r}: \sum_{k=m_{r}}^{m} F\left(Y_{\rho(k)}\right) \geq k_{n}\right\}
$$

until $\sum_{k=m_{s}}^{n} F\left(Y_{\rho(k)}\right)<k_{n}$. Let $\psi_{n}\left(T_{n}\right)$ be a partition of $\mathbb{R}$ into intervals $\left\{A_{1}, \ldots, A_{s}\right\}$ such that $X_{\rho\left(m_{j}\right)}, \ldots, X_{\rho\left(m_{j+1}-1\right)} \in A_{j}$, for $j=1, \ldots, s-1$, and
 the oridinary $k$-spacing partition of $\mathbb{R}$.

For suitable sequences of threshold values, Theorem 1 shows that histogram regression estimates based on $\psi_{n}$ are consistent. The proof, which makes use of Example 1, is similar to that given in Lugosi and Nobel (1995) for $k$-spacing density estimates, and is therefore omitted.

Theorem 5. Let a regression estimate $\hat{r}_{n}$ be formed by averaging response variables within the cells of $\psi_{n}\left(T_{n}\right)$. If $k_{n} \rightarrow \infty$ and $k_{n} / n \rightarrow 0$, then $P\left|\hat{r}_{n}-r\right|^{2}$ $\rightarrow 0$ with probability 1 .

Remark. In higher dimensions, similar results can be obtained for a $Y$-independent version of the partitioning scheme proposed by Gessaman (1970). The $Y$-independent case is discussed in Lugosi and Nobel (1995).
8. Nearest-neighbor clustering. Clustering of multivariate data is a widely used method of statistical analysis. Clustering schemes typically partition the training set by minimizing an empirical error criterion. In this section we establish the consistency of regression estimates based on nearest-neighbor clustering of the (unlabeled) measurement vectors $X_{i}$. Let $\|\cdot\|$ be the usual Euclidean norm on $\mathbb{R}^{d}$.

A clustering scheme is a function $C: \mathbb{R}^{d} \rightarrow \mathscr{C}$ that associates every vector $x \in \mathbb{R}^{d}$ with one of the finite number of cluster centers $\mathscr{C}=\left\{c_{1}, \ldots, c_{m}\right\} \subseteq \mathbb{R}^{d}$. Each clustering scheme has a corresponding partition $\pi=\left\{A_{1}, \ldots, A_{m}\right\}$ with cells $A_{j}=\left\{x: Q(x)=c_{j}\right\}$. A clustering scheme is said to be nearest neighbor if, for each $x \in \mathbb{R}^{d}$,

$$
C(x)=\underset{c_{j} \in \mathscr{C}}{\arg \min }\left\|x-c_{j}\right\|,
$$

with ties broken in favor of the center $c_{j}$ having the least index. In this case the partition of $C$ is the nearest-neighbor partition of its cluster centers. See Hartigan (1975) or Gersho and Gray (1992) for more details concerning multivariate clustering and its applications.

The risk of a clustering scheme $C$ is commonly measured by $R(C)=\int \| x-$ $C(x) \|^{2} d P$, the expected squared distance between a random vector $X \sim P$ and its corresponding cluster center. The empirical risk of $C$ is given by

$$
\begin{equation*}
R_{n}(C)=\int\|x-C(x)\|^{2} d \hat{P}_{n}=\frac{1}{n} \sum_{i=1}^{n}\left\|X_{i}-C\left(X_{i}\right)\right\|^{2} . \tag{29}
\end{equation*}
$$

From a training set $T_{n}=\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ and a clustering scheme $C$ one may produce a histogram regression estimate by averaging the response variables $Y_{i}$ within the cells of $C$. We consider estimates based on empirically optimal nearest-neighbor clustering schemes.

Definition. For each $x \in \mathbb{R}^{d}$ and every $\delta>0$, let $B(x, \delta)=\{v:\|u-v\|$ $<\delta\}$. The support set of $P$ is defined by

$$
S_{P}=\{x: P(B(x, \delta))>0 \text { for every } \delta>0\} .
$$

It is easy to see that $S_{P}$ is a closed set and $P\left(S_{P}\right)=1$.
Proposition 3. For each $n \geq 1$, let $C_{n}: \mathbb{R}^{d} \rightarrow \mathscr{C}_{n}$ minimize the empirical risk $R_{n}(C)$ over all nearest-neighbor clustering schemes having $k_{n}$ cluster centers. If $E\|X\|^{2}<\infty$ and $k_{n} \rightarrow \infty$, then, for every compact set $V \subseteq \mathbb{R}^{d}$,

$$
\max _{u \in S_{P} \cap V} \min _{c \in \mathscr{C}_{n}}\|u-c\| \rightarrow 0
$$

with probability 1.
Proof. By the monotonicity of $\|\cdot\|^{2}$, the cluster centers of an empirically optimal scheme must lie within the closed convex hull of $X_{1}, \ldots, X_{n}$. As this set is compact, the continuity of $\|\cdot\|^{2}$ insures that $C_{n}$ exists.

Let $u_{1}, u_{2}, \ldots$ be dense in $\mathbb{R}^{d}$, with $u_{1}=0$. For each $n$, let $C_{n}^{\prime}$ be the nearest-neighbor clustering scheme with centers $\left\{u_{1}, \ldots, u_{n}\right\}$. For each compact set $V \subseteq \mathbb{R}^{d}$,

$$
R_{n}\left(C_{n}\right) \leq R_{n}\left(C_{n}^{\prime}\right) \leq \max _{x \in V} \min _{1 \leq j \leq k_{n}}\left\|x-a_{j}\right\|^{2}+\int_{V^{c}}\|x\|^{2} d \hat{P}_{n}
$$

and therefore

$$
\limsup _{n \rightarrow \infty} R_{n}\left(C_{n}\right) \leq \int_{V^{c}}\|x\|^{2} d P
$$

with probability 1 . As $V$ was arbitrary, $R_{n}\left(C_{n}\right) \rightarrow 0$ with probability 1 .

By applying the strong law of large numbers to a countable, dense subset of $S_{P}$, it may be established that

$$
\begin{equation*}
\mathbb{P}\left\{\liminf _{n \rightarrow \infty} P_{n}(B(u, \delta))>0 \text { for every } u \in S_{P}, \delta>0\right\}=1 \tag{30}
\end{equation*}
$$

Suppose for the moment that there is a compact set $V \subseteq \mathbb{R}^{d}$ and a number $\delta>0$ such that

$$
\begin{equation*}
\mathbb{P}\left\{\limsup _{n \rightarrow \infty}\left[\max _{u \in S_{P} \cap V} \min _{c \in \mathscr{C}_{n}}\|u-c\|\right]>\delta\right\}>0 . \tag{31}
\end{equation*}
$$

Let $\mathscr{C}_{1}, \mathscr{C}_{2}, \ldots$ be a sequence of codebooks corresponding to a sample point in this event. As $S_{P} \cap V$ is compact, there exists a sequence of vectors $\left\{u_{n_{k}}\right\}$ such that $u_{n_{k}} \rightarrow u^{*} \in S_{P}$ and $\mathscr{E}_{n_{k}} \subseteq B^{c}\left(u_{n_{k}}, \delta\right)$ for every $k$. Thus, when $k$ is sufficiently large,

$$
R_{n_{k}}\left(C_{n_{k}}\right) \geq \frac{\delta}{3} \hat{P}_{n_{k}}\left(B\left(u^{*}, \frac{\delta}{3}\right)\right)
$$

and it follows from (30) and (31) that

$$
\mathbb{P}\left\{\liminf _{\rightarrow \infty} R_{n}\left(C_{n}\right)>0\right\}>0
$$

However, this contradicts the fact that $R\left(C_{n}\right) \rightarrow 0$, and we conclude that (31) cannot hold.

Theorem 6. Let $C_{n}$ minimize the empirical risk $R_{n}(C)$ over all nearestneighbor clustering schemes with $k_{n}$ cluster centers. Define $\hat{r}_{n}$ by averaging the response variables $Y_{i}$ within the cells of $C_{n}$. If $E\|X\|^{2}$ is finite, $k_{n} \rightarrow \infty$ and $n^{-1} k_{n}^{2} \log n \rightarrow 0$, then $P\left|r-\hat{r}_{n}\right|^{2} \rightarrow 0$ with probability 1 .

Proof. The partition associated with $C$ lies in the family $\mathscr{V}\left(k_{n}\right)$ containing all nearest-neighbor partitions of $k_{n}$-vectors in $\mathbb{R}^{d}$. Example 3 (Section 1.2) shows that $m\left(\mathscr{V}\left(k_{n}\right)\right)=k_{n}$ and $\Delta_{n}^{*}\left(\nabla\left(k_{n}\right)\right) \leq n^{k_{n}^{2} d}$. Our assumptions on $k_{n}$ guarantee that

$$
n^{-1} m\left(\mathscr{V}_{k_{n}}\right) \rightarrow 0 \quad \text { and } \quad n^{-1} \log \Delta_{n}^{*}\left(\mathscr{V}_{k_{n}}\right) \rightarrow 0 .
$$

The shrinking cell condition (c) of Theorem 1 follows easily from Proposition 3 .
9. Empirically optimal regression trees. Tree-structured partitions were defined in Example 3 of Section 2. A regression tree is a function $f$ : $\mathbb{R}^{d} \rightarrow \mathbb{R}$ defined by assigning a number to each cell of a tree-structured partition ( $T, \tau$ ). Alternatively, one may define a regression tree by augmenting the pair ( $T, \tau$ ) with an additional node function $\alpha: T \rightarrow \mathbb{R}$, and setting $f(x)=\alpha(t)$ for every vector $x$ whose path through $T$ ends at the leaf node $t$. In the notation of Section 2, a $k$-node regression tree is a function $f \in \mathscr{G} \circ \mathscr{T}(k)$, where $\mathscr{G}$ is the class of constant functions taking values in $[-K, K]$.

The regression function $r(\cdot)$ minimizes the predictive risk $J(f)=E \mid f(X)$ $-\left.Y\right|^{2}$ over all functions $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$. In practice, when the training set $T_{n}$ is given but the distribution of ( $X, Y$ ) is unknown, it is common to seek an estimate $\hat{f}_{n}$ that minimizes the empirical risk

$$
J_{n}(f)=\frac{1}{n} \sum_{i=1}^{n}\left(f\left(X_{i}\right)-Y_{i}\right)^{2}
$$

over a suitable class of regression estimates. Empirically optimal regression trees yield consistent estimates of $r(\cdot)$ if the size of the trees grows with $n$ at a controller rate.

Theorem 7. Let $\hat{f}_{n}$ minimize the empirical risk $J_{n}(f)$ over all the $k_{n}$-node regression trees $f \in \mathscr{G} \circ \mathscr{T}\left(k_{n}\right)$. If $k_{n} \rightarrow \infty$ and $k_{n}=o(n / \log n)$, then $P\left|\hat{f}_{n}-r\right|^{2}$ $\rightarrow 0$ with probability 1 .

Proof. As $k_{n}$ grows without bound, there exists a fixed, tree-structured partition $\pi_{n} \in \mathscr{T}\left(k_{n}\right)$ such that, for every compact set $V \subseteq \mathbb{R}^{d}$,

$$
\begin{equation*}
\max _{A \in \pi_{n}} \operatorname{diam}(A \cap V) \rightarrow 0 . \tag{32}
\end{equation*}
$$

Let $\hat{g}_{n}$ be the histogram regression estimate produced from $T_{n}$ by averaging the response variables $Y_{i}$ within the cells of $\pi_{n}$.

Under the assumption that $k_{n}=o(n / \log n)$, the bounds derived in Example 3 (Section 1.2) show that $n^{-1} m\left(\mathscr{T}\left(k_{n}\right)\right) \rightarrow 0$ and $n^{-1} \Delta_{n}^{*}\left(\mathscr{T}\left(k_{n}\right)\right) \rightarrow 0$. It then follows from (32) and Theorem 1 that $\left\{\hat{g}_{n}\right\}$ is strongly consistent. Since, for each $n, \hat{f}_{n}$ and $\hat{g}_{n}$ are contained in $\mathscr{G} \circ \mathscr{T}\left(k_{n}\right)$, and $\hat{f}_{n}$ is empirically optimal, the consistency of the estimates $\hat{f}_{n}$ follows from Lemma 4.

We remark that the partitions $\pi_{n}$ above may be chosen to have rectangular cells. Thus the conclusion of the theorem holds for the empirically optimal regression trees employing axis-parallel splits. It should also be noted that the estimates $\hat{f}_{n}$ need not have shrinking cells in the sense of (4). When $n$ is large, the empirically optimal tree will not divide regions of the measurement space on which the regression function is constant.

Acknowledgments. The author has benefitted from numerous conversations with Richard Olshen, who helped to clarify the relationship of these results to existing work. The helpful comments of Gábor Lugosi and Amir Dembo contributed to the content and organization of the paper.

## REFERENCES

Anderson, T. W. (1966). Some nonparametric multivariate procedures based on statistically equivalent blocks. In Multivariate Analysis (P. R. Krishnaiah, ed.) 5-27. Academic Press, New York.
Brieman, L., Friedman, J. H., Olshen, R. A. and Stone, C. J. (1984). Classification and Regression Trees. Wadsworth, Belmont, CA.

Chaudhuri, P., Huang, M.-C., Loh, W.-Y. and Yao, R. (1994). Piecewise polynomial regression trees. Statist. Sinica 4 143-167.
Cover, T. M. (1965). Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. IEEE Transaction on Electronic Computers 14 326-334.
Devroye, L. (1988). Automatic pattern recognition: a study of the probability of error. IEEE Transactions on Pattern Analysis and Machine Intelligence 10 530-543.
Devroye, L. and Györfi, L. (1985). Distribution-free exponential bounds on the $l_{1}$ error of partitioning estimates of a regression function. In Proceedings of the Fourth Pannonian Symposium on Mathematical Statistics (F. Konecny, J. Mogyoródi, and W. Wertz, eds.) 67-76. Akadémiai Kiadó, Budapest.
Gersho, A. and Gray, R. M. (1992). Vector Quantization and Signal Compression. Kluwer, Dordrecht.
Gessaman, M. P. (1970). A consistent nonparametric multivariate density estimator based on statistically equivalent blocks. Ann. Math. Statist. 41 1344-1346.
Gordon, L. and Olshen, R. (1978). Asymptotically efficient solutions to the classification problem. Ann. Statist. 6 515-533.
Gordon, L. and OlShen, R. (1980). Consistent nonparametric regression from recursive partitioning schemes. J. Multivariate Anal. 10 611-627.
Gordon, L. and OlShen, R. (1984). Almost sure consistent nonparametric regression from recursive partitioning schemes. J. Multivariate Anal. 15 147-163.
Hartigan, J. A. (1975). Clustering Algorithms. Wiley, New York.
Linde, Y., Buzo, A. and Gray, R. M. (1980). An algorithm for vector quantizer design. IEEE Trans. Comm. 28 84-95.
Lugosi, G. and Nobel, A. B. (1995). Consistency of data-driven histogram methods for density estimation and classification. Ann. Statist. To appear.
Parthasarathy, K. R. and Bhattacharya, P. K. (1961). Some limit theorems in regression theory. Sankhyā Ser. A 23 91-102.
Patrick, E. A. and Fisher, F. P. (1967). Introduction to the performance of distribution-free conditional risk learning systems. Technical Report TR-EE-67-12, Purdue Univ.
Pollard, D. (1984). Convergence of Stochastic Processes. Springer, New York.
Stone, C. J. (1977). Consistent nonparametric regression (with discussion). Ann. Statist. 5 595-645.
Stone, C. J. (1985). An asymptotically optimal histogram selection rule. In Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer (L. Le Cam and R. A. Olshen, eds.) 513-520. Wadsworth, Belmont, CA.
Vapnik, V. N. and Chervonenkis, A. Ya. (1971). On the uniform convergence of relative frequencies of events to their probabilities. Theory Probab. Appl. 16 264-280.
Vapnik, V. N. and Chervonenkis, A. Ya. (1981). Necessary and sufficient conditions for the uniform convergence of means to their expectations. Theory Probab. Appl. 26 532-553.
Zhao, L. C., Krishnaiah, P. R. and Chen, X. R. (1990). Almost sure $L_{r}$-norm convergence for data-based histogram density estimates. Theory Probab. Appl. 35 396-403.

Department of Statistics
University of North Carolina
Chapel Hill, North Carolina 27599-3260
E-MAIL: nobel@stat.unc.edu


[^0]:    Received December 1993; revised May 1995.
    ${ }^{1}$ This research was completed while the author was a Beckman Institute Fellow at the Beckman Institute for Advanced Science and Technology, University of Illinois, UrbanaChampaign.

    AMS 1991 subject classifications. Primary 62G07; secondary 62H30.
    Key words and phrases. Regression estimation, histogram estimates, regression trees, clustering, Vapnik-Chervonenkis theory, data-dependent partitions.

