



2010

Small Area Estimation of the Homeless in Los Angeles: An Application of Cost-Sensitive Stochastic Gradient Boosting

Brian Kriegler

Richard A. Berk
University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/statistics_papers

 Part of the [Applied Statistics Commons](#)

Recommended Citation

Kriegler, B., & Berk, R. A. (2010). Small Area Estimation of the Homeless in Los Angeles: An Application of Cost-Sensitive Stochastic Gradient Boosting. *Annals of Applied Statistics*, 4 (3), 1234-1255. <http://dx.doi.org/10.1214/10-AOAS328>

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/statistics_papers/352
For more information, please contact repository@pobox.upenn.edu.

Small Area Estimation of the Homeless in Los Angeles: An Application of Cost-Sensitive Stochastic Gradient Boosting

Abstract

In many metropolitan areas efforts are made to count the homeless to ensure proper provision of social services. Some areas are very large, which makes spatial sampling a viable alternative to an enumeration of the entire terrain. Counts are observed in sampled regions but must be imputed in unvisited areas. Along with the imputation process, the costs of underestimating and overestimating may be different. For example, if precise estimation in areas with large homeless counts is critical, then underestimation should be penalized more than overestimation in the loss function. We analyze data from the 2004–2005 Los Angeles County homeless study using an augmentation of L_1 stochastic gradient boosting that can weight overestimates and underestimates asymmetrically. We discuss our choice to utilize stochastic gradient boosting over other function estimation procedures. In-sample fitted and out-of-sample imputed values, as well as relationships between the response and predictors, are analyzed for various cost functions. Practical usage and policy implications of these results are discussed briefly.

Keywords

homeless, boosting, statistical learning, costs, imputation, quantile estimation, small area estimation

Disciplines

Applied Statistics

SMALL AREA ESTIMATION OF THE HOMELESS IN LOS ANGELES: AN APPLICATION OF COST-SENSITIVE STOCHASTIC GRADIENT BOOSTING¹

BY BRIAN KRIEGLER AND RICHARD BERK

Econ One Research and University of Pennsylvania

In many metropolitan areas efforts are made to count the homeless to ensure proper provision of social services. Some areas are very large, which makes spatial sampling a viable alternative to an enumeration of the entire terrain. Counts are observed in sampled regions but must be imputed in unvisited areas. Along with the imputation process, the costs of underestimating and overestimating may be different. For example, if precise estimation in areas with large homeless counts is critical, then underestimation should be penalized more than overestimation in the loss function. We analyze data from the 2004–2005 Los Angeles County homeless study using an augmentation of L_1 stochastic gradient boosting that can weight overestimates and underestimates asymmetrically. We discuss our choice to utilize stochastic gradient boosting over other function estimation procedures. In-sample fitted and out-of-sample imputed values, as well as relationships between the response and predictors, are analyzed for various cost functions. Practical usage and policy implications of these results are discussed briefly.

1. Introduction. Dating as far back as the 1930s, homelessness has been a visible, public issue in the United States [Rossi (1989)]. At least over the past decade, the homeless problem has been underscored due to the rise in unemployment and foreclosures. In the 2010 census, there are no plans to perform street counts, thereby making it challenging for stakeholders (e.g., homeless service advocates and selected government agencies) to estimate the magnitude of the necessary social resources. This is especially difficult in large metropolitan areas because the homeless are often dispersed due to the changing availability of homeless services, commercial development and the government’s homeless criminalization practices [Berk, Brown and Zhao (2010)]. Areas needing these services are literally “moving targets.” Adequate spatial apportionment of homeless-related resources requires a great deal of local information that is oftentimes prohibitively expensive to obtain.

In a typical census design, people are contacted through their place of residence. With the possible exception of individuals living on private property, the homeless will not be found using this design [Rossi (1989)]. An alternative approach is to

Received March 2009; revised January 2010.

¹Supported by NSF Grant SES-0437179, “Ensemble Methods for Data Analysis in the Behavior, Social, and Economic Sciences.”

Key words and phrases. Homeless, boosting, statistical learning, costs, imputation, quantile estimation, small area estimation.

locate homeless individuals in temporary shelters or while they are receiving services (e.g., meals) from public and private agencies. It is widely known, however, that a large number of the homeless still will not be found this way because many do not use these services. Therefore, it is common for enumerators to canvas geographical areas and to count the homeless as they find them. Some metropolitan areas are very large, making spatial sampling a viable substitute to a full canvassing. One trades a reduction in the burden of data collection in exchange for the need to impute homeless counts for locales not visited by enumerators.

Estimation and imputation raise the issue of how best to represent the cost of underestimation relative to overestimation (“cost function”). The apportionment of homeless-related resources depends, at least in part, on the estimated size of the local homeless population. Some stakeholders, such as homeless service providers, are more troubled by the prospect of numbers that are too small rather than too large. This is especially true in areas where homeless counts are high, in which undercounting may carry serious consequences. Other stakeholders, such as elected city officials faced with budget constraints, may have the opposite preference. In general, one needs the flexibility to penalize overestimation and underestimation distinctly.

The homeless problem is especially serious in Los Angeles, which has a large homeless population and consists of specific areas with very densely populated homeless encampments [Berk, Kriegler and Ylvisaker (2008)]. These encampments can be a nuisance to local commerce and can compound the demand, for example, for police and hospital services [Harcourt (2005)]. One such area is “Skid Row” [Magnano and Blasi (2007)], located just outside downtown Los Angeles. Historically, this area has been marked by high crime rates in terms of drug markets, robberies, vandalism and prostitution, as well as drug and alcohol abuse [Lopez (2005)].¹ Individuals (especially the homeless) who spend significant amounts of their time in public areas of such locales have higher victimization rates than those who reside outside these areas [Koegel, Burnam and Farr (1988); Kushel et al. (2003)]. In short, the set of public and private resources dependent on the homeless population extends beyond the services dedicated to the homeless’ physical and mental health (e.g., soup kitchens, shelters, affordable housing, etc.).

In 2004–2005, the Los Angeles Homeless Services Authority (LAHSA) estimated the homeless population in Los Angeles County as the aggregate of people who were living on the streets, in shelters or who were “nearly homeless” (i.e., homeless people living on private property with the consent of its residents). At any given time, shelters cater to just a fraction of the local homeless popula-

¹In 2005, the Los Angeles Police Department tested a pilot program, called “Safer Cities Initiative” (SCI), which was designed to target specific geographical crime “hot spots” [Wilson and Kelling (1982); Bratton and Knobler (1998)]. Part of this program entailed reducing the density of homeless encampments. A full-scale version of SCI began in September 2006 [Berk and MacDonald (2010)].

tion; consequently, locating and estimating the street count was a daunting task.² It would have been prohibitively costly to canvas the entire county, which covers over 4000 square miles, includes 2054 census tracts, and is the most populous county in the United States.

A stratified spatial sampling of census tracts called for two steps. First, tracts believed to have large numbers of homeless people were visited with probability 1. There were 244 tracts of this nature, known as “hot tracts.” The second step was to visit a stratified random sample of tracts from the population of nonhot tracts. The strata were the county’s eight Service Provision Areas (SPAs), and the number of tracts drawn from each stratum was proportional to the number of tracts assigned to each SPA. In all, there were 265 tracts in the stratified random sample, leaving 1545 tracts’ counts to be imputed.³ In that analysis, the cost function was symmetric, and emphasis was placed on estimating the homeless population within each SPA, for various aggregations (e.g., cities), and for the entire county [Berk, Krieglér and Ylvisaker (2008)]. Almost certainly, symmetric costs are insufficiently responsive to the policy needs of local stakeholders because both actual and imputed counts can vary dramatically.

In this paper we re-analyze the Los Angeles data of 1810 nonhot tracts using stochastic gradient boosting [Friedman (2002)] subject to an asymmetrically weighted absolute loss function. We focus on evaluating the relationship between homeless counts and covariates in visited tracts and imputing the counts in unvisited tracts. By boosting a cost-sensitive loss function, we are able to respond to the cost functions of various stakeholders and focus on a particular region of the conditional response. Depending on which cost function is applied, widely varying fitted and imputed values can follow. We also explore how different regions of the conditional response are related to the predictors. We show that it can be practical and instructive to employ asymmetric costs when using boosting for function estimation and imputation.

The remainder of this paper consists of five sections plus an [Appendix](#). Section 2 includes a description of the Los Angeles County homeless and census data. In Section 3 we provide an overview of stochastic gradient boosting and a literature review on cost-sensitive estimation procedures. Our analysis of the homeless data, which includes comparisons between fitted and observed counts, imputed counts, and model diagnostics, is in Section 4. Section 5 includes a discussion on how our proposed methodology and analysis can have a profound effect on policy-making

²Homeless people were paid \$10 per hour to help the field researchers identify locations in which the homeless could be found. Presumably, this helped address the problem of finding “hidden homeless” [Rossi (1989)].

³This is a “small area estimation” analysis. Rao (2003) defines a domain, or area, as “small” if “the domain-specific sample is not large enough to support direct estimates of adequate precision.” In the context, homeless counts in the 265 randomly sampled tracts were used to impute the numbers of homeless people in unvisited tracts and ultimately the entire county.

decisions. In Sections 4 and 5 we stress the results based on models that place heavier penalties on underestimating, as this represents what stakeholders would likely employ to ensure proper allocation of homeless-related services. We conclude the paper in Section 6, in which we mention some aspects of cost-sensitive statistical learning to be explored. In the Appendix we derive the functional forms for the deviance, initial value, gradient and terminal node estimates when employing boosting subject to asymmetrically weighted absolute loss.

2. Data description. In the 2004–2005 Los Angeles homeless study, Berk, Kriegler and Ylvisaker (2008) considered the use of dozens of predictors in the estimation process.⁴ The 10 predictors in Table 1 were relatively important to fitting the conditional distribution of street counts, capturing information about each tract’s geographical location, land usage, socioeconomic information and ethnic demographic data. With the exception of median household income and planar coordinates, all other covariates are presented in terms of percentages. While street counts were obtained only in sampled tracts, predictor values were available for all of the county’s tracts.

Looking ahead to Section 4, none of our models are intended to necessarily suggest causal relationships. We utilized predictor information described in Table 1 primarily to estimate the conditional distribution between *StTotal* and each covariate and to construct sensible fitted and imputed street counts. Whether the predictors are causally related to homeless counts is at best a secondary concern.

TABLE 1
Names and descriptions of variables in Los Angeles County homeless data set

	Description
Response name	
<i>StTotal</i>	Homeless street count
Predictor name	
<i>Commercial</i>	% of land used for commercial purposes
<i>Industrial</i>	% of land used for industrial purposes
<i>MedianHouseholdIncome</i>	Median household income
<i>PctMinority</i>	% of population that is non-Caucasian
<i>PctOwnerOcc</i>	% of owner-occupied housing units
<i>PctVacant</i>	% of unoccupied housing units
<i>Residential</i>	% of land used for residential purposes
<i>VacantLand</i>	% of land that is vacant
<i>XCoord</i>	Planar longitude
<i>YCoord</i>	Planar latitude

⁴In that study, fitted and imputed counts were obtained using random forests [Breiman (2001)].

The distribution of *StTotal* is highly unbalanced. 75 percent of the observed counts are less than 28 people, and 22 of the 265 tracts have at least 50 homeless, of which 11 have over 100 homeless (Min = 0, Q1 = 4, Median = 12, Mean = 21.6, Q3 = 27, Max = 282). To ensure adequate local resources, stakeholders such as police departments and homeless shelter advocates may place heavy emphasis on accurately estimating the counts in areas that have large homeless populations (e.g., over 100 people). If so, one is willing to trade overall accuracy for a better fit in the right tail of the street count distribution, and underestimates are more costly than overestimates. For policy purposes, resources may still be adequate in an area with a predicted count of 30 people when in fact the count is 50. However, if the prediction is 30 and the actual count is 150, there may well be a severe shortage of local resources.

3. Estimating the conditional distribution. Let Y be a set of real response values, X be a vector of one or more real predictor variables ($1, \dots, P$), and $f(x_i)$ be a fitting function for observation i ($i = 1, \dots, N$). We seek to minimize some loss function, Ψ , to fit the conditional response distribution, $G(Y|X = x)$:

$$(3.1) \quad G(Y|X = x) = \arg \min_f E\{\Psi(Y, f(x))\}.$$

We could minimize the L_1 loss so that the estimate is

$$(3.2) \quad G_{L_1}(Y|X = x) = \arg \min_f E\{|Y - f(x)|\},$$

in which overestimating and underestimating the response are weighted symmetrically, and \hat{f} is the median of Y . But if underestimating and overestimating are not equally costly, then the loss criteria needs to be *asymmetric*. Let $L_1(\alpha)$ be the absolute loss function that weights underestimates by α and overestimates by $1 - \alpha$, where $0 \leq \alpha \leq 1$. Then $G_{L_1(\alpha)}(Y|X = x)$ is defined as

$$(3.3) \quad \begin{aligned} G_{L_1(\alpha)}(Y|X = x) \\ = \arg \min_f E\{\alpha|Y - f(x)| \cdot I(Y > f(x)) \\ + (1 - \alpha)|Y - f(x)| \cdot I(Y \leq f(x))\}, \end{aligned}$$

where $I(Y > f(x))$ and $I(Y \leq f(x))$ are mutually exclusive indicator variables. For each $i = 1, \dots, N$, if y_i is underestimated, then the former equals 1 and the latter equals 0. Conversely, if y_i is estimated perfectly or is overestimated, then these binary values are reversed. Note that $G_{L_1(\alpha)}$ reduces to G_{L_1} when $\alpha = 0.5$.

In general, $\hat{f}(x)$ from equation (3.3) is the quantile of Y , which exhibits a straightforward translation between the cost function (or “cost ratio”) and descriptions of the response distribution. For example, a 3 to 1 cost ratio implies that underestimating is three times as costly as overestimating, the ratio of underestimates to overestimates will be 3 to 1, and \hat{f} is the $3/(3+1) \times 100 = 75$ th percentile of Y . If instead the cost ratio is less than 1 to 1, then \hat{f} is less than the median of Y . Henceforth, we refer to $\alpha/(1 - \alpha)$ as the cost ratio.

3.1. *Stochastic gradient boosting: An overview.* Stochastic gradient boosting [Friedman (2002)] is a recursive, nonparametric procedure that has become one of the most popular machine learning algorithms among statisticians. It exhibits extraordinary fitting flexibility, as it can handle any differentiable and minimizable loss function. It can handle and produce highly complex functional forms, and there is growing evidence that it outperforms competing procedures (e.g., bagging [Breiman (1996)], splines, CART [Breiman et al. (1984)] and parametric regression) in terms of prediction error [Friedman (2001); Bühlmann and Yu (2003); Madigan and Ridgeway (2004)], provided that one utilizes reasonable tuning parameters.⁵ Shortly after Friedman (2001) introduced gradient boosting, Friedman (2002) augmented the algorithm by taking a random sample of observations at each iteration, thereby creating the *stochastic* gradient boosting machine. This additional feature to the algorithm resulted in marked reduction in bias and variance. Given stochastic gradient boosting’s success at estimating the center of $Y|X$, one may deduce that it also performs well at estimating other regions of the conditional response distribution.

The stochastic gradient boosting algorithm in its most general form is provided below⁶ [Friedman (2002); Ridgeway (2007); Berk (2008)]:

1. Initialize $\hat{f}(x)$ to the same constant value across all observations, $\hat{f}_0(x) = \arg \min_{\rho_0} \sum_{i=1}^N \Psi(y_i, \rho_0)$.
2. For t in $1, \dots, T$, do the following:
 - (a) For $i = 1, \dots, N$, compute the negative gradient as the working response:

$$z_{ti} = - \left[\frac{\partial \Psi(y_i, f_{t-1}(x_i))}{\partial f_{t-1}(x_i)} \right]_{f_{t-1}(x_i) = \hat{f}_{t-1}(x_i)}$$

- (b) Take a simple random sample without replacement of size N' from the data set with N observations.
- (c) Fit a regression tree with K_t terminal nodes, $g_t(x) = E(z_t|x)$ using the randomly selected observations.
- (d) Compute the optimal terminal node estimates, $\rho_{1_t}, \dots, \rho_{K_t}$, as

$$\rho_{k_t} = \arg \min_{\rho_{k_t}} \sum_{x_i \in S_{k_t}} \Psi(y_i, \hat{f}_{t-1}(x_i) + \rho_{k_t}),$$

where S_{k_t} is the set of x -values that defines terminal node k at iteration t .

⁵This is especially true when the number of predictors is large [Bühlmann and Yu (2003)].

⁶Our augmentation of stochastic gradient boosting and data analysis were conducted using gbm in R [Ridgeway (2007)]. We found four boosting libraries in R in addition to gbm: ada [Culp (2006); Culp, Michailidis and Johnson (2006)], GAMBoost [Binder (2009)], gbeV [Sexton (2009)] and mboost [Hothorn (2009)]. The respective maintainers of these packages are Mark Culp, Harald Binder, Joe Sexton and Torsten Hothorn.

(e) Again using the sampled data, update $\hat{f}_t(x)$ as

$$\hat{f}_t(x_i) \leftarrow \hat{f}_{t-1}(x_i) + \lambda \rho_{k_t(x_i)},$$

where λ is the “learning rate.”

In the [Appendix](#) we build on equation (3.3) to derive the deviance subject to $L_1(\alpha)$. Subsequently, we identify the functional form of the initial value, gradient and terminal node estimates from steps 1, 2a and 2d of the stochastic gradient boosting algorithm.

3.2. Literature review. To our knowledge, the inclusion of asymmetric costs to boosting algorithms has applied solely to classification problems. [Fan et al. \(1999\)](#) introduce an algorithm called AdaCost, a more flexible version of AdaBoost [[Freund and Schapire \(1997\)](#)].⁷ [Mease, Wyner and Buja \(2007\)](#) propose a boosting algorithm called JOUS-Boost, (**J**ittering and **O**ver/**U**nder-**S**ampling). By adding small amounts of noise to the data and weighting the probability of selection according to each class, one can obtain different misclassification rates than if using no jittering or unweighted sampling according to classes. [Berk, Krieglner and Baek \(2006\)](#) incorporate costs into a classification framework using stochastic gradient boosting by specifying a threshold between 0 and 1; observations with predicted probabilities below or above the threshold are assigned values of 0 or 1, respectively. The threshold was established so that the ratio of misclassification errors (false negatives to false positives) approximated the cost ratio.

In a regression context, we found three methods capable of handling asymmetric error costs, each building on quantile estimation. If the functional form is specifiable a priori, one can employ parametric quantile regression [[Koenker \(2005\)](#)]. However, if the functional form is not known, it is important and helpful to exploit statistical learning. Then, one could apply nonparametric quantile regression [[Takeuchi et al. \(2006\)](#)]. Yet there is evidence that ensemble procedures, such as gradient boosting, typically yield superior bias-variance tradeoffs in comparison [[Bühlmann and Hothorn \(2007\)](#)]. [Meinshausen \(2006\)](#) introduced quantile regression forests, an augmentation of random forests [[Breiman \(2001\)](#)]. The drawback to this method is that the fitted and imputed values are calculated *after* all of the trees are grown using random forests. Consequently, the conditional response function does not adapt to the cost ratio. It follows that there are no new partial dependence plots and predictor importance measurements (not even when employing L_1 , since the usual random forests algorithm estimates the conditional mean).

⁷In a follow-up study of AdaCost and other cost-sensitive variations of AdaBoost, [Ting \(2000\)](#) shows that AdaCost stumbles in certain situations, and that this could be due to the algorithm’s weighting structure.

Just as with parametric quantile regression, estimates based on $L_1(\alpha)$ stochastic gradient boosting do not necessarily increase monotonically with respect to α .⁸ Each cost function yields a different model and fitted values that minimize the $L_1(\alpha)$ loss. Therefore, a fitted (or imputed) count may be 30 when the cost ratio is 5 to 1 and 20 when the cost ratio is 10 to 1. With $L_1(\alpha)$ stochastic gradient boosting, our experience—both in this case study and with other data sets—is that (i) all (or nearly all) fitted and imputed values tend to increase with respect to α , and (ii) when decreases do occur, they tend to be small in magnitude. We found that the use of larger terminal node sizes can reduce this occurrence; however, for reasons we explain in Section 4, we purposely grew trees that potentially had small terminal node sizes. Ultimately, we were not concerned with this “side effect” because its occurrence was rare and inconsequential, and our analysis extended beyond simply calculating fitted and imputed values.

In summary, we employed $L_1(\alpha)$ stochastic gradient boosting for three main reasons. First, the functional form can be arrived at inductively. Second, we have the prospect of a good bias-variance tradeoff. Third, we can apply unequal error costs at each step of the function estimation process so that *all* of the output is properly cost-sensitive. We found $L_1(\alpha)$ stochastic gradient boosting to provide a formidable set of features for this case study, though it should not be seen as a universal preference for cost-sensitive stochastic gradient boosting in different settings.

4. Analysis. Based on our discussions with key stakeholders, including people from LAHSA and government representatives, underestimation is typically seen to be more problematic than overestimation. The prospect of having too few shelter beds, for instance, is more troubling than if a few beds are open. With this in mind, our analysis emphasizes results in which $\alpha \geq 0.5$. Output based on cost functions that penalize overestimation more heavily are also reported, primarily to demonstrate that they are employable if one desires.

All boosting models were built using the following tuning parameters: 10 splits per tree subject to at least 5 observations per terminal node k_t , a learning rate of $\lambda = 0.001$, and a maximum of $T = 6000$ trees. For stochastic gradient boosting models, we applied these same tuning parameters along with a random sample of $N' = 133$ observations (i.e., a sampling fraction of 50 percent of $N = 265$, rounded to the nearest whole number). A sensible number of iterations was determined using 10-fold cross-validation, and we found no problems in converging on a reasonable number of trees to grow in any of our cost-sensitive models.⁹

⁸Incidentally, quantile regression forests does not share this feature because the quantile estimation is performed on the distribution of each observation’s fitted values across regression trees.

⁹For example, in the stochastic models when the cost ratio $\alpha/(1 - \alpha) \in \{1 \text{ to } 10, 1 \text{ to } 1, 10 \text{ to } 1\}$, the respective “best” numbers of iterations were 436, 1843 and 1340. Small deviations from these

Using a handful of different learning rates and sampling fractions ranging from 0.001 to 0.01 and 35 to 75 percent, respectively, we saw inconsequential differences in terms of street counts estimates—both fitted and imputed—and conditional distribution diagnostics, for each α . The same held true for models subject to 1 to 10, 1 to 5, and 1 to 1 costs. By contrast, when we employed cost ratios of 5 to 1 and 10 to 1, we learned that the number of splits and the minimum terminal node size can have a substantial impact on point estimates. The `gbm` library uses the inverse of the empirical distribution to estimate quantiles, so each terminal node estimate depends on just one value. Given the unbalanced nature of *StTotal*, differences between consecutive values in the right tail within a terminal node can be very large. If employing a 10 to 1 cost function and a terminal node includes 25 points, then the estimate will be the third highest value. The use of a highly skewed cost function implies a particular interest in estimating the handful of large response values well, yet the top two values in this terminal node of this size will not factor into the estimation process. To ensure that large gradients were given ample opportunities to be terminal node estimates, we permitted large trees and small terminal node sizes. This was facilitated by tuning the number of splits and the minimum number of observations in each terminal node at each iteration.¹⁰

4.1. *Fitted and imputed street counts.* Figure 1 shows fitted versus observed street counts for the 265 visited census tracts using stochastic gradient boosting subject to 1 to 10, 1 to 1, 5 to 1, and 10 to 1 cost ratios ($\alpha \in \{1/11, 1/2, 5/6, 10/11\}$, respectively). Using 1 to 1 costs (L_1 boosting), the magnitude of the error is less than 20 people in 232 of 265 visited census tracts. In terms of resource needs, errors of this magnitude are likely tolerable. Conversely, among the 22 tracts with observed counts with at least 50 homeless, all of these tracts' counts are underestimated. The maximum fitted value is approximately 37 people, and the median error is approximately 70 people less than the true count. These large undercounts need to be reduced substantially in order to ensure adequate local resource allocation.

Figure 1 demonstrates that $L_1(\alpha)$ stochastic gradient boosting fitted values tend to increase with respect to α .¹¹ Although the overall fit worsens when the cost ratio

numbers of iterations (e.g., 1400 trees subject to a 10 to 1 ratio) yielded no substantive differences in any results. Just as one would expect when using symmetric costs, the cross-validation error exhibited a concave-up parabolic behavior that tended to decrease with respect to t , until it reached a number of iterations corresponding to the minimum cross-validation error. Beyond the minimum cross-validation error iterations, the models overfit the data [Zhang and Yu (2005)]. The key here is that these iteration estimates are well short of $T = 6000$, suggesting that we have in fact identified a sensible number of iterations.

¹⁰By default, in `gbm` each tree at each iteration has one split, subject to at least 10 observations in each terminal node.

¹¹Of the 265 visited training data observations, 10 observations' fitted values were lower for $\alpha = 10/11$ than for $\alpha = 5/6$. We did not consider this to be problematic for two reasons. The largest

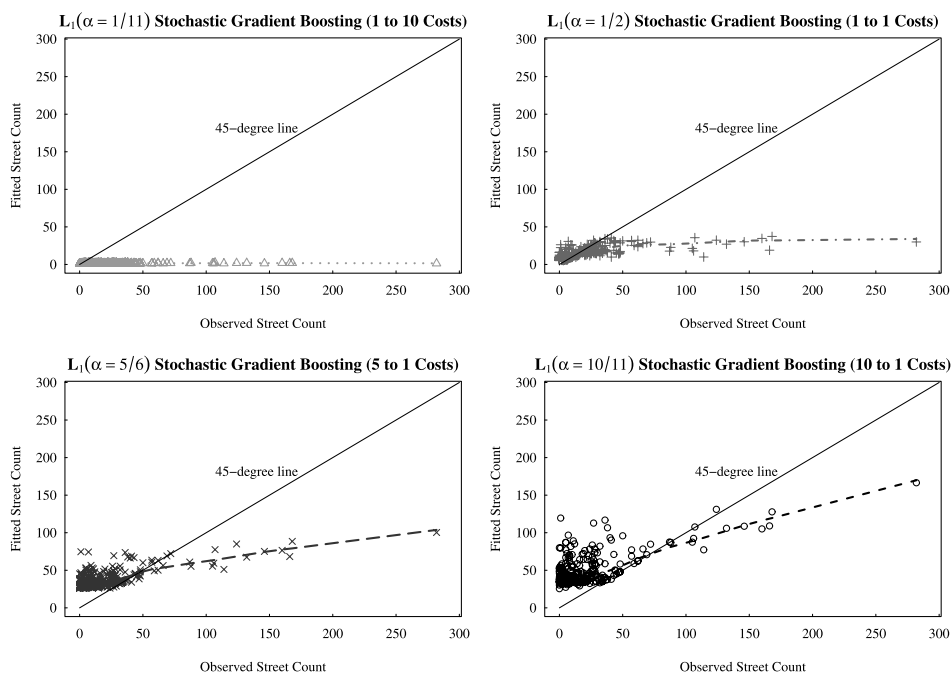


FIG. 1. Fitted versus observed census tract street counts using $L_1(\alpha)$ stochastic gradient boosting.

diverges from 1 to 1, we observe smaller errors in specific regions of the response. Using a 10 to 1 cost ratio, just 15 out of 265 tracts are underestimated. Among the 22 tracts with at least 50 people, the median difference between observed and fitted counts is 1 person, and the interquartile range is 40 people. Admittedly, most of the very large counts are still underestimated even when using a 10 to 1 cost ratio, a topic we will pick up again in Section 5.¹²

In a way, training data fitted values are irrelevant because one’s estimates of visited tracts might simply be the observed street count. Berk, Krieger and Ylvisaker (2008) employed this practice when they provided estimates to LAHSA at both the tract and aggregate levels. But provided the sampled tracts are representative of the population of all nonhot tracts and the model does not overfit the training data, fitted counts in Figure 1 reveal how close (or far) the unsampled tracts’ imputed counts are to the true counts. Figure 2 shows the distribution of imputed counts for various cost ratios. The distributions tend to shift upward with respect

of these differences was 4 people. Also, this was generally not a problem among tracts with very large counts; one tract had a street count of 62, and the next highest count was 43.

¹²Recognizing that it is in the nature of all regression models to overestimate small values and underestimate large ones, we demonstrate that the use of asymmetric costs can alleviate the problem. As the cost ratio increases, fitted values for tracts with large counts tend to move closer to the 45-degree line.




FIG. 2. *Distribution of predicted street counts in unvisited census tracts using $L_1(\alpha)$ stochastic gradient boosting.*

to α .¹³ Using 1 to 10 and 1 to 5 costs, all tracts have imputed counts of fewer than 5 people. Conversely, using 10 to 1 costs, we find that 53 of 1545 tracts have imputed counts over 100 homeless people.

Recognizing that portions of our analysis will be data set specific, one may also be interested in how $L_1(\alpha)$ boosting performs relative to other cost-sensitive methods. Figure 3 shows fitted versus observed street counts using stochastic and nonstochastic gradient boosting, and parametric quantile regression, subject to a 10 to 1 cost function.¹⁴ All three methods have a substantial number of overestimates, which is to be expected given the cost ratio of choice. Among tracts with at least 50 homeless people observed, $L_1(\alpha)$ stochastic gradient boosting performs noticeably better than the other two methods in terms of bias and variance. Nonstochastic gradient boosting exhibits a median deviation of 35 people underestimated and an IQR of 77 people. Quantile regression's median deviation and IQR are 7 and 63 people, respectively.

4.2. Conditional distribution diagnostics. With 10 predictors, a highly unbalanced response distribution and abrupt spatial variation in the data, the boosted

¹³Of the 1545 unvisited tracts, imputed values were higher using $\alpha = 5/6$ versus $\alpha = 10/11$ in 44 tracts. Over half of these deviations were less than 2 people, and the largest deviation was 6 people.

¹⁴Parametric quantile regression was performed using the `quantreg` library in R, maintained by Roger Koenker (2009).

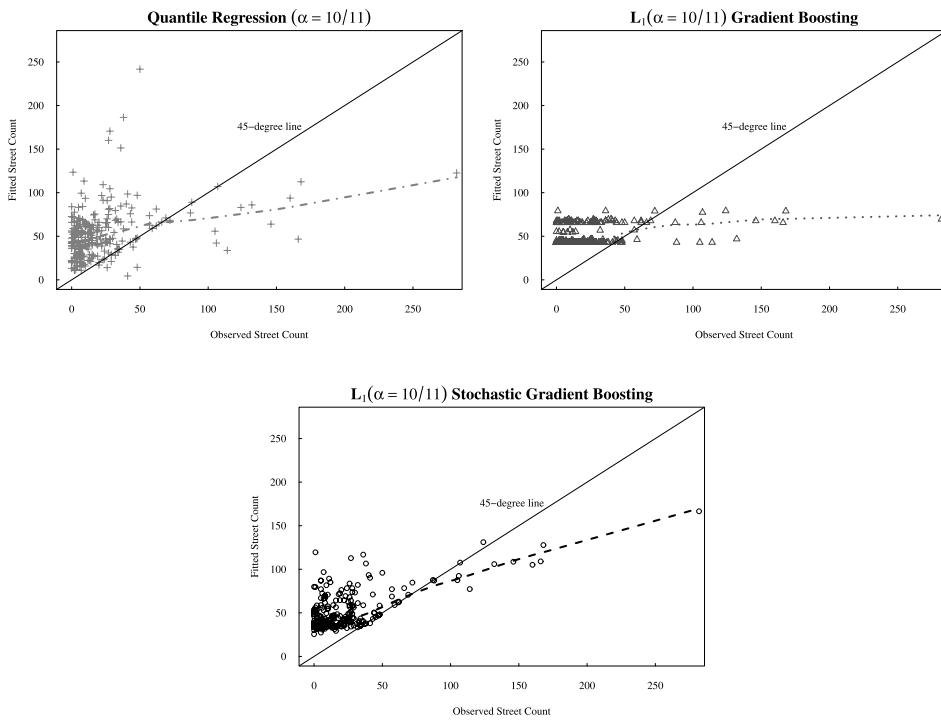


FIG. 3. Fitted versus observed street counts using quantile regression, $L_1(\alpha)$ gradient boosting and $L_1(\alpha)$ stochastic gradient boosting, subject to a 10 to 1 cost ratio ($\alpha = 10/11$).

models' conditional distribution diagnostics are practical and necessary to understanding relationships between the response and the predictors. Since the cost function is built into each step of $L_1(\alpha)$ boosting, partial plots and variable importance measures can be examined in the same manner as when employing L_1 boosting. These results are especially important if stakeholders are inclined to give causal interpretations to the associations.

One may assume that the partial relationships between the response and each predictor exhibit similar directional behavior and are nothing more than vertical shifts in the conditional response's magnitude. An analogous argument might be made regarding variable importance: if a predictor is important using symmetric costs, then perhaps the same is true using asymmetric costs. If these inferences are correct, cost-sensitive partial and predictor importance plots are less critical. Yet Figures 4 and 5 demonstrate that predictors' relationships with the response are not necessarily the same across cost ratios, underscoring the need to examine the conditional distribution diagnostics for each cost ratio of interest.

4.2.1. *Partial relationships.* To show partial relationships between the response and each predictor, Friedman (2001) describes a weighted tree traversal

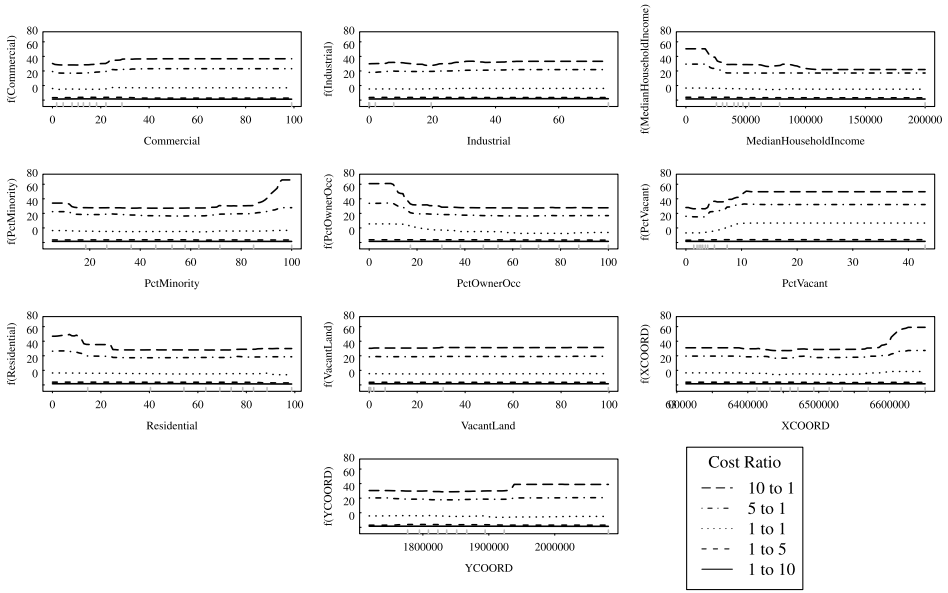


FIG. 4. Partial dependence plots from $L_1(\alpha)$ stochastic gradient boosting.




FIG. 5. Variable importance from $L_1(\alpha)$ stochastic gradient boosting.

method to “integrate out” all predictor variables, excluding the predictor(s) of interest [see also Ridgeway (2007)]. Figure 4 shows partial relationships between the response and each predictor for five different cost ratios. Since each of the predictors exhibits real values, each partial relationship is shown using a two-dimensional smoother.¹⁵ For cost ratios of 1 to 10 and 1 to 5, all of the partial relationships are nearly flat, a result consistent with the small variation in tract-level estimates reported in Figures 1 and 2. Using symmetric L_1 boosting, street counts increase with respect to *PctVacant* between 0 and 10 percent, and street counts decrease with respect to *PctOwnerOcc* between 20 and 60 percent. Pragmatically, all other partial relationships are close to null.

When underestimating *StTotal* is more costly, the conditional response can vary substantially with respect to several other predictors in addition to the housing vacancy rates and the fraction of owner-occupied units. For example, using a 10 to 1 cost function, street counts are indifferent to *PctMinority* until approximately 90 percent, but increase substantially between 90 and 100 percent. Street counts decrease in a stepwise manner with respect to *MedianHouseholdIncome*; we see plateaus for incomes between \$0 and \$15,000, \$30,000 to \$75,000, and \$100,000 and above.

4.2.2. Variable importance. One may be interested in identifying which predictors are “important” to fitting the conditional response for various cost ratios. One measure of variable importance is the reduction in loss attributed to each predictor. Friedman (2001) and Ridgeway (2007) define the “relative influence” as the empirical reduction in squared error in predicting the gradient across all node splits on predictor j , divided by the total reduction in error across all splits.

Even if the response and predictor j are completely unrelated, it is still possible for the predictor to be selected to split a regression tree node. Provided there is at least one split on predictor j , the empirical influence will not be zero. How then, does one know the extent to which a predictor’s influence is by chance? Along the same lines as in random forests [Breiman (2001)], in which importance is computed by shuffling each predictor in turn and comparing the change in error, we employed the following steps to estimate each predictor’s “baseline relative influence”:

1. For a given predictor p , randomly permute the values. Keep all other predictors’ values as is.

¹⁵The `gbm` library estimates the partial response at equally-spaced values (by default, 100) spanning the range of the predictor but independent of the predictor’s empirical density. As a result, decile rugs are shown at the bottom of each plot for each corresponding predictor to better understand the distribution of each predictor. For example, the vacancy rate is 33 percent for one tract, 43 percent for another tract and less than 20 percent for all other tracts. For *PctVacant* greater than 20 percent, it is difficult to determine the extent to which these partial smoothers are robust because they are based on so few points.

2. Construct a boosted model using the modified data in step 1 and compute the relative influence for the shuffled predictor. Apply the same tuning parameter settings and means for estimating a sensible number of iterations.
3. Repeat steps 1 and 2 many times, each time computing the relative influence of the shuffled predictor.¹⁶
4. Compute the baseline relative influence as the average relative influence from steps 1–3.
5. Repeat steps 1–4 for each predictor in turn.

Figure 5 shows each predictor’s empirical and baseline relative influence values subject to five different cost ratios. If a predictor’s baseline relative influence (denoted by a thick black line and the diagonally shaded area) is larger than its empirical influence, this suggests that the contribution to the model is happenstance. Just as in the partial plots, we learn that a predictor’s relative influence is not necessarily similar across cost functions. This can be a very important practical matter insofar as stakeholders come to accept or reject the homeless estimates depending on whether predictors “make sense.”

One should also be mindful of the difference between the overall reduction in error from $t = 0$ —at which all estimates are equal to the grand α quantile of $StTotal$ —to the “optimal” number of iterations. If the total reduction in error is very small, then the *absolute* influence will be minimal. It follows that the differences between each fitted response value and the initial constant will likely be small as well. Under these circumstances, the relative influence results are inconsequential. Such is the case for boosted models subject to 1 to 10 and 1 to 5 costs. Figures 1, 2 and 4 suggest minimal variation in fitted and predicted counts; substantively, the relationships between $StTotal$ and each predictor are null. Importance statistics subject to these two cost ratios are reported primarily for demonstrative purposes.

Using symmetric costs, *PctVacant* and *PctOwnerOcc* are relatively important, collectively accounting for nearly 35 percent of the loss reduction. *PctVacant* is also important when the cost ratio is 5 to 1 or 10 to 1, along with *PctMinority* and *XCoord*, and to a lesser extent *MedianHouseholdIncome*. These predictors’ relative influence are high compared to other predictors’ importance statistics and is well above their respective baseline influences. Conversely, *PctOwnerOcc* is much less important when underestimation is penalized more heavily, evidenced by its smaller relative influence and proximity to the baseline relative influence.

5. Discussion. $L_1(\alpha)$ stochastic gradient boosting is a potentially useful statistical tool for ensuring adequate allocation of services related to the homeless. Practitioners might find it useful to build multiple boosted models for various cost functions and examine the range of imputed counts for a specific tract in order to

¹⁶For $\alpha \in \{1/11, 1/6, 1/2, 5/6, 10/11\}$, we repeated steps 1 and 2 50 times per predictor.

make policy decisions. Suppose a homeless service provider or local police department considers it critical to identify tracts that have over 100 homeless people; the former might aspire to ensure a sufficient number of beds at the nearest shelter, and the latter may well decide to allocate additional officers to areas with high homeless counts. Assume that a particular tract's imputed count is 30 using 1 to 1 costs and 150 using 10 to 1 costs. Such stakeholders may insist on performing a full enumeration in this tract because these two imputed counts have very different resource implications. Alternatively, if the imputed counts using these respective cost ratios are 30 and 40, a full enumeration may not be worth the trouble because the difference is likely inconsequential.

Among the 11 tracts with over 100 homeless, stochastic gradient boosting subject to a 10 to 1 cost ratio yields a better prediction error than gradient boosting or parametric quantile regression. Still, 9 of the 11 tracts are underestimated, and the prediction error tends to increase with respect to the observed count. It is reasonable to assume that among unvisited tracts with over 100 homeless, imputed counts will be similarly biased. In practice, one way to further reduce this problem is by assigning larger "population weights" a priori to training data tracts with large street counts. The population weights increase the frequency of specific observations if they are selected in step 2b of the algorithm described in Section 3.1. One assumes—and perhaps rightfully so—that some tracts are inherently more important than others. If larger weights are assigned to tracts with high street counts, then fitted and imputed counts will also increase. A toy example is provided in the [Appendix](#).

In addition to evaluating imputed counts, suppose stakeholders (e.g., LAHSA) want to use response-predictor relationships to determine which unvisited tracts might require the most resources. Figure 4 suggests that areas with some combination of high non-Caucasian populations, high vacancy rates, low median household incomes and low rates of owner-occupied housing may be indicators of high homeless populations. Based on Figure 5, *PctVacant* and *PctMinority* are especially key to identifying areas potentially in need of services.

6. Conclusion. This case study features a number of characteristics that make the analysis challenging. Although there are relatively few tracts with large homeless counts, these are likely the most important tracts to fit reasonably well—without overfitting the data—so that unvisited tracts with potentially high counts are identified. In addition, Los Angeles County exhibits considerable heterogeneity and abrupt spatial changes in terms of land usage and demography. Last, the wide range of stakeholders would likely assign various costs to over/under-counting during the estimation and imputation processes. We believed that a cost-sensitive ensemble statistical learning procedure was appropriate because (i) we did not presume to understand the underlying mechanisms of the conditional street count distribution, (ii) we aspired to get favorable results in terms of prediction error for specified regions of the response, and (iii) we wanted to understand how specific

regions of the conditional response were related to the predictors. $L_1(\alpha)$ stochastic gradient boosting allowed us to address all of these issues.

There are a handful of practical statistical issues born out of this case study. First, one might argue that a “cost-sensitive Poisson” loss function is a more appropriate procedure for the homeless data because the outcome is a count. A key issue, then, is whether L_1 or L_2 loss is more responsive to the data imputation task at hand and to the quality of the data. In our case, a few very large observed counts would likely dominate the analysis under L_2 . Whether this is good or bad depends on the accuracy of the few very large counts and on the policy matter of how much those large counts should be permitted to affect the imputations. We take no strong position on either issue, but we have concerns from past research on homeless enumerations that the count data could contain significant error [Cordray and Pion (1991); Cowan (1991); Rossi (1991); Wright and Devine (1992)]. And, we find that boosting the $L_1(\alpha)$ loss function incorporates cost considerations in a straightforward and easily interpretable manner.

There is also the matter of statistical inference, a topic we glossed over in Section 4.2.2 by estimating each predictor’s baseline relative importance. To our knowledge, statistical inference remains a largely unsolved problem for stochastic gradient boosting and statistical learning in general [Leeb and Pötscher (2005, 2006); Berk, Brown and Zhao (2010)]. We have explored the properties of a procedure that wraps cost-sensitive boosting in bootstrap sampling cases. Although this seems to provide some useful information on the stability of our imputed values, we do not think it addresses the fundamental problems identified by Leeb and Pötscher (2005).

Finally, the application of $L_1(\alpha)$ boosting brings to light the issue of choosing the “right” tuning parameters, a topic explored by Mease and Wyner (2008). While the number of splits has been researched extensively [e.g., Schapire (1999); Friedman, Hastie and Tibshirani (2000); Bühlmann and Yu (2003); Ridgeway (2007)], research on the impact of different terminal node sizes is minimal thus far. Unlike estimates subject to Poisson or Gaussian loss, which are functions of *all* gradients within each terminal node, an $L_1(\alpha)$ terminal node estimate is the quantile of gradients residing in terminal node k_t . These estimates depend on just a very local region of points and can be highly dependent on the terminal node sizes and the way in which the quantile is estimated [for variants of quantile estimation, see Hyndman and Fan (1996)]. The performance of $L_1(\alpha)$ stochastic gradient boosting subject to various quantile estimation procedures remains a topic for future research.

APPENDIX: BOOSTING THE $L_1(\alpha)$ DISTRIBUTION

Ridgeway (2007) specifies the boosted L_1 (Laplace) loss function as

$$(A.1) \quad \Psi(f_t(x_i) : x_i \in S_{k_t}) = \left\{ \sum_{x_i \in S_{k_t}} |w_i(y_i - f_t(x_i))| \right\} / \sum_{x_i \in S_{k_t}} w_i,$$

where w_i is a predetermined population weight for observation i that remains constant across all iterations. Altering (A.1) to allow for unequal costs, the loss function becomes

$$(A.2) \quad \Psi(f_t(x_i) : x_i \in S_{k_t}) = \left\{ \alpha \sum_{\substack{x_i \in S_{k_t} \\ y_i > \hat{f}_t(x_i)}} |w_i(y_i - \hat{f}_t(x_i))| \right. \\ \left. + (1 - \alpha) \sum_{\substack{x_i \in S_{k_t} \\ y_i \leq \hat{f}_t(x_i)}} |w_i(y_i - \hat{f}_t(x_i))| \right\} / \sum_{x_i \in S_{k_t}} w_i,$$

which is an asymmetrically weighted absolute loss function if $\alpha \neq 0.5$.¹⁷ For shorthand, denote $\Psi(f_t(x_i) : x_i \in S_{k_t}) = \Psi$. Then, the gradient becomes¹⁸

$$(A.3) \quad z_{ti} = -\frac{\partial \Psi}{\partial f_t(x_i)} = \begin{cases} w_i \alpha : y_i > \hat{f}_{t-1}(x_i), \\ -w_i(1 - \alpha) : y_i \leq \hat{f}_{t-1}(x_i), \end{cases}$$

where the derivative is evaluated at $\hat{f}_{t-1}(x_i)$. We wish to find the value of ρ_{k_t} that minimizes Ψ subject to the loss function in (A.2):

$$(A.4) \quad \rho_{k_t} = \arg \min_{\rho_{k_t}} \left\{ \alpha \sum_{\substack{x_i \in S_{k_t} \\ y_i > \hat{f}_{t-1}(x_i) + \rho_{k_t}}} |w_i(y_i - (\hat{f}_{t-1}(x_i) + \rho_{k_t}))| \right. \\ \left. + (1 - \alpha) \sum_{\substack{x_i \in S_{k_t} \\ y_i \leq \hat{f}_{t-1}(x_i) + \rho_{k_t}}} |w_i(y_i - (\hat{f}_{t-1}(x_i) + \rho_{k_t}))| \right\},$$

where $f_t(x_i)$ is the fitted value from the previous iteration, $\hat{f}_{t-1}(x_i)$, plus the terminal node estimate from the current iteration, ρ_{k_t} . Next, we differentiate to find the value of ρ_{k_t} that minimizes Ψ :

$$(A.5) \quad \frac{\partial \Psi}{\partial \rho_{k_t}} = \left\{ -\alpha \sum_{\substack{x_i \in S_{k_t} \\ y_i > \hat{f}_{t-1}(x_i) + \rho_{k_t}}} w_i + (1 - \alpha) \sum_{\substack{x_i \in S_{k_t} \\ y_i \leq \hat{f}_{t-1}(x_i) + \rho_{k_t}}} w_i \right\} / \sum_{x_i \in S_{k_t}} w_i,$$

¹⁷With this distribution, the estimate \hat{f} is in the same units as y ; therefore, over/under-estimation are determined by comparing the two. Estimates in some distributions, such as Poisson, are in terms of logits and must be exponentiated to be on the same scale as y .

¹⁸Under the usual L_1 loss function, the gradient for observation i is the sign of the difference between the observed response (y_i) and the predicted value ($\hat{f}_t(x_i)$), multiplied by the population weight, w_i .

$$(A.6) \quad 0 = -\alpha \sum_{\substack{x_i \in S_{k_t} \\ y_i > \hat{f}_{t-1}(x_i) + \rho_{k_t}}} w_i + (1 - \alpha) \sum_{\substack{x_i \in S_{k_t} \\ y_i \leq \hat{f}_{t-1}(x_i) + \rho_{k_t}}} w_i.$$

In the right-hand side of (A.6), each summation reduces to the number of observations that are underestimated or overestimated, respectively. Let N_{k_t} denote the number of observations in terminal node k_t , and let n_{k_t} and $N_{k_t} - n_{k_t}$ be the number of underestimates and overestimates in the terminal node, respectively. For simplicity, assume that $w_i = 1$ for all i . Solving for n_{k_t} , the location parameter is

$$(A.7) \quad n_{k_t} = \alpha N_{k_t}.$$

The way in which unequal population weights affect the terminal node estimate is worthy of a toy example. Consider terminal node k_t with 5 equally-weighted observations with fitted gradients—the “working responses”—at $t - 1$ of 0, 3, 5, 6 and 15. If we are estimating the median, then the terminal node estimate is 5. Now suppose that prior to constructing the boosted model, the observation with the fitted gradient of 15 at $t - 1$ was instead assigned a population weight of 3. Then this observation’s fitted gradient from $t - 1$ will appear in node k_t three times, and the population-weighted median is 6.¹⁹

By weighting the loss function according to overestimates and underestimates, the fitted value of terminal node k_t is the α quantile of the N_{k_t} gradients. In each terminal node, there are approximately αN_{k_t} and $(1 - \alpha)N_{k_t}$ gradients above and below ρ_{k_t} , respectively. For all $i = 1, \dots, N$, $f_0(x_i)$ equals 0, and ρ_0 equals the α quantile of the response variable, y . Therefore, the fitted value for observation i after T iterations, $\hat{f}_T(x_i)$, equals²⁰

$$(A.8) \quad \hat{f}_T(x_i) = \text{quantile}_\alpha(y) + \lambda \sum_{t=1}^T \text{quantile}_\alpha(z_{ti}).$$

Because $L_1(\alpha)$ is differentiable and there exists a solution that minimizes this loss [Hastie, Tibshirani and Friedman (2001)], we are able to incorporate costs into stochastic gradient boosting where the response is quantitative, and in some sense add a distribution to those provided in Friedman (2001).

Acknowledgments. We gratefully acknowledge numerous helpful discussions with Greg Ridgeway. We would also like to thank the Editor, anonymous associate editor and anonymous reviewers for their constructive and thoughtful comments on this manuscript.

¹⁹At present, `gbm` does not allow for unequal population weights when employing the quantile distribution.

²⁰Note that $z_{ti} = 0$ if observation i is not randomly selected as one of the N' observations in step 2b of the stochastic gradient boosting algorithm described in Section 3.1.

REFERENCES

- BERK, R. A. (2008). *Statistical Learning From a Regression Perspective*. Springer, New York.
- BERK, R. A. and MACDONALD, J. (2010). Policing the homeless: An evaluation of efforts to reduce homeless-related crime. *Criminology & Public Policy*. To appear.
- BERK, R. A., BROWN, L. and ZHAO, L. (2010). Statistical inference after model selection. *Journal of Quantitative Criminology*. To appear. Available at <http://www.springerlink.com/content/fjuk84hq160181rv/?p=3ec2606a5fd946b589f2f763b384df94&pi=7>.
- BERK, R. A., KRIEGLER, B. and BAEK, J. H. (2006). Forecasting dangerous inmate misconduct: An application of ensemble statistical procedures. *Journal of Quantitative Criminology* **22** 131–145.
- BERK, R. A., KRIEGLER, B. and YLVIKAKER, D. (2008). Counting the homeless in Los Angeles County. In *Probability and Statistics: Essays in Honor of David A. Freedman. Inst. Math. Statist. Collect.* **2** 127–141. IMS, Beachwood, OH. MR2459951
- BINDER, H. (2009). GAMBBOOST: Generalized additive models by likelihood based boosting. R version 1.1. Available at <http://www.r-project.org>.
- BRATTON, W. and KNOBLER, P. (1998). *The Turnaround: How America's Top Cop Reversed the Crime Epidemic*. Random House, New York.
- BREIMAN, L. (1996). Bagging predictors. *Machine Learning* **26** 123–140.
- BREIMAN, L. (2001). Random forests. *Machine Learning* **45** 5–32.
- BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A. and STONE, C. J. (1984). *Classification and Regression Trees*. Wadsworth, Monterey, CA. MR0726392
- BÜHLMANN, P. and HOTHORN, T. (2007). Boosting algorithms: Regularization, prediction and model fitting. *Statist. Sci.* **22** 477–505. MR2420454
- BÜHLMANN, P. and YU, B. (2003). Boosting with the L_2 loss: Regression and classification. *J. Amer. Statist. Assoc.* **98** 324–340.
- CORDRAY, D. S. and PION, G. M. (1991). What's behind the numbers? Definitional issues in counting the homeless. *Housing Policy Debates* **2** 587–616.
- COWEN, D. D. (1991). Estimating census and survey undercounts through multiple service contacts. *Housing Policy Debates* **2** 869–882.
- CULP, M. (2006). ada: Performs boosting algorithms for a binary response. R package 2.0-1. Available at <http://www.r-project.org>.
- CULP, M., MICHAILEDIS, G. and JOHNSON, K. (2006). ada: An R package for stochastic boosting. *Journal of Statistical Software* **17**.
- FAN, W., STOLFO, S. J., ZHANG, J. and CHAN, P. K. (1999). Adacost: Misclassification cost-sensitive boosting. In *Machine Learning: Proceedings of the Sixteenth International Conference* 97–105. Bled, Slovenia.
- FREUND, Y. and SCHAPIRE, R. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* **55** 119–139. MR1473055
- FRIEDMAN, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Ann. Statist.* **29** 1189–1232. MR1873328
- FRIEDMAN, J. H. (2002). Stochastic gradient boosting. *Comput. Statist. Data Anal.* **38** 367–378. MR1884869
- FRIEDMAN, J. H., HASTIE, T. and TIBSHIRANI, R. (2000). Additive logistic regression: A statistical view of boosting. *Ann. Statist.* **28** 337–374. MR1790002
- HARCOURT, B. E. (2005). Policing L.A.'s skid row: Crime and real estate development in downtown Los Angeles (an experiment in real time). *Univ. Chicago Legal Forum* **2005** 325–404.
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning*. Springer, New York. MR1851606
- HOTHORN, T. (2009). mboost: Model-based boosting. R package 1.1-2. Available at <http://www.r-project.org>.

- HYNDMAN, R. J. and FAN, Y. (1996). Sample quantiles in statistical packages. *Amer. Statist.* **50** 361–365.
- KOEGEL, P., BURNAM, A. and FARR, R. K. (1988). The prevalence of specific psychiatric disorders among homeless individuals in the inner-city of Los Angeles. *Archives of General Psychiatry* **45** 1085–1092.
- KOENKER, R. (2005). *Quantile Regression*. Cambridge Univ. Press, New York. [MR2268657](#)
- KOENKER, R. (2009). `quantreg`: Quantile regression. R package 4.38. Available at <http://www.r-project.org>.
- KUSHEL, M. B., EVANS, J. L., PERRY, S., ROBERTSON, M. J. and MOSS, A. R. (2003). No door to lock: Victimization among homeless and marginally housed persons. *Arch. Intern. Med.* **163** 2492–2499.
- LEEB, H. and PÖTSCHER, B. M. (2005). Model selection and inference: Facts and fiction. *Econom. Theory* **21** 21–59. [MR2153856](#)
- LEEB, H. and PÖTSCHER, B. M. (2006). Can one estimate the conditional distribution of post-model-selection estimators? *Ann. Statist.* **34** 2554–2591. [MR2291510](#)
- LOPEZ, S. (2005). Demons are winning on Skid Row. *Los Angeles Times*, October 16.
- MADIGAN, D. and RIDGEWAY, G. (2004). Discussion of “Least angle regression” by B. Efron et al. *Ann. Statist.* **32** 465–469. [MR2060166](#)
- MAGNANO, P. F. and BLASI, G. (2007). Stuck on Skid Row. *Los Angeles Times (Opinion Section)*, October 29.
- MEASE, D. and WYNER, A. (2008). Evidence contrary to the statistical view of boosting. *J. Mach. Learn. Res.* **9** 131–156.
- MEASE, D., WYNER, A. and BUJA, A. (2007). Cost-weighted boosting with jittering and over/under-sampling: JOUS-boost. *J. Mach. Learn. Res.* **8** 409–439.
- MEINSHAUSEN, N. (2006). Quantile regression forests. *J. Mach. Learn. Res.* **7** 983–999. [MR2274394](#)
- RAO, J. N. K. (2003). *Small Area Estimation*. Wiley, Hoboken, NJ. [MR1953089](#)
- RIDGEWAY, G. (2007). `gbrm`: Generalized boosted regression models. R package 1.6-3. Available at <http://www.r-project.org>.
- ROSSI, P. H. (1989). *Down and Out in America: The Origins of Homelessness*. Univ. Chicago Press, Chicago.
- ROSSI, P. H. (1991). Strategies for homeless research in the 1990s. *Housing Policy Debates* **2** 1029–1055.
- SCHAPIRE, R. E. (1999). A brief introduction to boosting. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* **2** 1401–1406. Morgan Kaufmann, San Francisco, CA.
- SEXTON, J. (2009) `gbev`: Gradient boosted regression trees with errors-in-variables. R package 0.1.1. Available at <http://www.r-project.org>.
- TAKEUCHI, I., LE, Q. V., SEARS, T. D. and SMOLA, A. J. (2006). Nonparametric quantile regression. *J. Mach. Learn. Res.* **7** 1231–1264. [MR2274404](#)
- TING, K. M. (2000). A comparative study of cost-sensitive boosting algorithms. In *Proceedings of the Seventeenth International Conference on Machine Learning* 983–990. Morgan Kaufmann, San Francisco, CA.
- WILSON, J. Q. and KELLING, G. L. (1982). Broken windows: The police and neighborhood safety. *Atlantic Monthly* **March** 29–38.
- WRIGHT, J. D. and DEVINE, J. A. (1992). Counting the homeless: The Census Bureau’s ‘S-Night’ in five U.S. cities. *Evaluation Review* **16** 355–364.

ZHANG, T. and YU, B. (2005). Boosting with early stopping: Convergence and consistency. *Ann. Statist.* **33** 1538–1579. MR2166555

ECON ONE RESEARCH
601 W. 5TH STREET
FIFTH FLOOR
LOS ANGELES, CALIFORNIA 90071
USA
E-MAIL: bkriegler@econone.com

DEPARTMENT OF CRIMINOLOGY
SCHOOL OF ARTS AND SCIENCES
UNIVERSITY OF PENNSYLVANIA
483 MCNEIL BUILDING
3718 LOCUST WALK
PHILADELPHIA, PENNSYLVANIA 19104-6286
USA
E-MAIL: berkr@sas.upenn.edu