

# Surprise sampling: Improving and extending the local case-control sampling\*

Xinwei Shen and Kani Chen

*Department of Mathematics,  
The Hong Kong University of Science and Technology  
e-mail: [xinwei.shen@connect.ust.hk](mailto:xinwei.shen@connect.ust.hk); [shakishakister@gmail.com](mailto:shakishakister@gmail.com)*

Wen Yu

*Department of Statistics,  
School of Management,  
Fudan University  
e-mail: [wenyu@fudan.edu.cn](mailto:wenyu@fudan.edu.cn)*

**Abstract:** Fithian & Hastie [7] proposed a sampling scheme called local case-control (LCC) sampling that achieves stability and efficiency by utilizing a clever adjustment pertained to the logistic model. It is particularly useful for classification with large and imbalanced data. This paper proposes a more general sampling scheme based on a working principle that data points deserve higher sampling probability if they contain more information or appear “surprising” in the sense of, for example, a large error of pilot prediction or a large absolute score. Compared with the relevant existing sampling schemes, as reported in [7] and [1], the proposed one has several advantages. It adaptively gives out the optimal forms to a variety of objectives, including the LCC and [1] as special cases. Under same model specifications, the proposed estimator also performs no worse than those in the literature. The estimation procedure is valid even if the model is misspecified and/or the pilot estimator is inconsistent or dependent on full data. We present theoretical justifications of the claimed advantages and optimality of the estimation and the sampling design. Different from [1], our large sample theory are population-wise rather than data-wise. Moreover, the proposed approach can be applied to unsupervised learning studies, since it essentially only requires a specific loss function and no response-covariate structure of data is needed. Numerical studies are carried out and the evidence in support of the theory is shown.

**MSC2020 subject classifications:** Primary 62D05; secondary 62J12.

**Keywords and phrases:** Generalized linear models, Horvitz-Thompson estimator, local case-control sampling, model mis-specification, subsampling.

Received August 2020.

---

\*Kani Chen was supported by Hong Kong GRF grants 16309816 and 1616212117. Wen Yu was supported by the National Natural Science Foundation of China Grants (12071088).

## Contents

1	Introduction . . . . .	2455
2	Notation and problem setting . . . . .	2458
3	Surprise sampling design . . . . .	2460
	3.1 Overall prediction accuracy . . . . .	2460
	3.2 Estimation accuracy . . . . .	2461
	3.3 The algorithms . . . . .	2462
4	Large sample properties . . . . .	2463
5	Simulation studies . . . . .	2465
6	Applications . . . . .	2471
7	Conclusion and discussion . . . . .	2474
A	Appendix A . . . . .	2475
	A.1 Proof of the propositions . . . . .	2475
	A.2 Proof of the theorem and corollaries . . . . .	2477
	Acknowledgments . . . . .	2481
	References . . . . .	2481

## 1. Introduction

Nowadays, with the rapid development of data capturing and storage techniques, people meet huge amounts of data in various fields. With the growth of the data size, computational capacity becomes more and more crucial to implement efficient data analysis. Despite significant progress on computer hardware, computational ability may still become a major constraint for data analysis when the data size is sufficiently large. For instance, it might be very time and resource consuming if we want to try a variety of competing models instead of only fitting one or two predetermined models, or if we need to refit the model from time to time with new observations arriving continuously, or if we need to apply the data partition and reusing techniques such as cross-validation, bootstrapping, bagging, and so on. All these computationally intensive procedures require tremendous computational costs, especially for those large data sets.

A simple approach to reduce the computational cost is to draw a subsample from the full data set and then analyze the subsample. The easiest subsampling design is to draw the data uniformly. However, uniform subsampling might be very inefficient for some data structures. For instance, in classification problems with two classes (one class for positive examples called “cases” and the other for negative examples called “controls”), when the classes are imbalanced, that is, one of the classes (usually the class of cases) is rare and the other is dominant, uniform subsampling is unfavorable because it ignores the unequal importance of the data points. For imbalanced data, case-control sampling is a well-known subsampling design [4, 11, 12]. It draws uniform subsamples from each of the two classes with different sampling percentages. Often comparable number of cases and controls are sampled, yielding a subsample with no obvious imbalance, to increase the efficiency of estimation. [2] and [14] showed that by fitting a logistic

model on the case-control subsample and then making a simple adjustment, one can still get a consistent estimator for the regression parameters if the logistic model is correctly specified. Thus, case-control design can help reduce the computational burden and retain satisfactory estimation efficiency under the imbalanced structure. Besides classification, imbalanced data structure also appears for other data types. A survey on predictive modeling under various imbalanced distributions is provided by [3].

Subsampling designs have also received attention in epidemiological cohort studies. The cohorts to follow up usually involve a great number of subjects and the occurrence of a certain disease is interested. In most cases, the occurrence rate of the disease is low during the entire follow-up time, so the cohort is essentially imbalanced. Meanwhile, the collection of covariate information on all involved subjects can be very expensive and time consuming because of the large cohort size. Hence, subsampling designs were developed to save sampling cost and time. The widely studied ones include case-cohort design [13], case-control design, and nested case-control design [15]. More generalized case-cohort designs aiming to improve the estimation efficiency were developed later by [5, 6, 19]. The sampling probability of all these subsampling designs depends on the observed follow-up times and the censoring indicators, but not the covariates. These designs are usually used for covariates ascertainment, that is, the covariates are observed only when the subject is selected into the subsample.

Statistical models are usually imposed for data analysis, but in real applications the models are easily misspecified. For some machine learning approaches, models are just used to derive meaningful loss functions, without caring about if they are correctly specified or not. Under these circumstances, the target parameter becomes a certain population risk minimizer corresponding to the loss function used [9]. When there exists possibly model mis-specification, robust analysis procedures are often preferred. However, as [7] pointed out, in classification problems, when the logistic model is misspecified, the standard case-control sampling can not provide a consistent estimator of the target parameter. To overcome this drawback, [7] proposed a local case-control (LCC) sampling design, which depends not only on the class label, but also on the predictors and a pilot estimate (a pilot estimate is a “good” guess for the target parameter). Their design provides a clever way to remedy imbalance locally throughout the feature space and possesses potential robustness against model mis-specification. Moreover, an elegant LCC estimate mimicking the full sample maximum likelihood estimate (MLE) is proposed. They showed that when the logistic model is correctly specified and the pilot is consistent and independent of the full data, the asymptotic variance of their proposed estimate is twice the variance of the full data MLE; when the logistic model is misspecified, the proposed estimate is still consistent and asymptotically normal as long as the pilot is consistent and independent of the data.

The LCC design aims to simultaneously speed up computation and provide a simple procedure to obtain a good estimate even under model mis-specification, but the authors did not discuss the optimality of the design. More recently, [17] modified the LCC design smartly to get an optimal subsampling method that

minimizes the conditional asymptotic mean squared error (MSE) of the subsample estimator given the full data. Later on, [1] extended the method to generalized linear models with canonical links. However, they only considered the optimization criterion regarding the conditional MSE. Some other non-uniform subsampling designs for the linear model include [10] and [16].

In this paper, we propose an improved subsampling design which accommodates various types of statistical learning objectives and includes the LCC sampling and [1] as special cases. For estimating the target parameter based on the subsample, we apply the Horvitz-Thompson (HT) type estimation [8]. The new sampling design is derived by optimizing certain well-defined criteria, such as prediction accuracy, estimation accuracy, or MSE. For different criteria, the proposed design has its corresponding form adaptively. Basically, it draws a data point with a large error of pilot prediction or a large score into the subsample with higher probability. Provided a pilot estimate, such a data point is in certain sense unusual, or “surprising”, for that given pilot, so we call the proposed design surprise sampling design.

The advantages of the proposed surprise sampling are summarized as follows.

- (i) The surprise sampling design is derived based on well-defined objectives. For a specific objective, the corresponding design is optimal. Meanwhile, the proposed design flexibly adapts varying objectives. That is, for different objectives, one can derive the optimal designs adaptively. In contrast, the objective of the LCC design is not clearly defined and the optimality is not discussed in [7]. [1] and [17] only considered the objective of minimizing the conditional MSE.
- (ii) The proposed estimators are always consistent and asymptotically normal, regardless of the correctness of the model specification and the consistency of the pilot estimate. The consistency of the estimator in [7] requires the consistency of the pilot. The pilots used in [1] and [17] are also consistent.
- (iii) If the pilot estimate is consistent and the logistic model is correctly specified, the proposed estimator is no worse than that in [7] in the sense that they have the same asymptotic efficiency under the LCC sampling.
- (iv) The validity of the proposed estimation procedure does not require the pilot estimate to be independent of the full data, while [7] requires the independence. This relaxation is useful in applications especially when there is no other data source and one needs to get the pilot from the full data.
- (v) The large sample properties derived for the proposed estimators are population-wise (unconditional), while the parallel large sample properties in [1] and [17] are developed data-wise (conditional on the full data).
- (vi) The proposed approach can be generally applied to not only supervised learning such as classification and regression but also unsupervised learning tasks, because it essentially only requires a well-defined loss function with a finite-dimensional parameter. Hence the application of the approach is more than the scope of logistic model or the other generalized linear models.

We proceed as follows. Section 2 introduces the notation and describes the problem setting. In Section 3, we present the main idea of the proposed surprise sampling and the specific forms of the sampling design are respectively derived to reach various objectives, such as best prediction accuracy and estimation accuracy. Section 4 gives out the large sample properties of the HT type estimator under the surprise sampling design. In Section 5, extensive simulation studies are carried out to show the effectiveness of the proposed approach. The results of real data analysis are provided in Section 6. Section 7 concludes. All the technique details are summarized in the Appendix.

## 2. Notation and problem setting

Suppose the full data contains  $n$  subjects and let  $d_i$  be the observed data point for the  $i$ -th subject. We assume that  $d_i$ ,  $i = 1, \dots, n$ , are independent and identically distributed (i.i.d.) copies of a random element  $D$  whose distribution stands for the population. For supervised learning problems,  $D$  can be decomposed into a response, denoted by  $Y$ , and a  $q$ -dimensional predictor or covariate vector, denoted by  $X$ . Correspondingly, the observed data points  $d_i = (y_i, x_i)$ ,  $i = 1, \dots, n$ , are i.i.d. copies of  $D = (Y, X)$ . We aim to predict  $Y$  through  $X$ , or learn the regression function  $f(x) = \mathbf{E}(Y|X = x)$ , based on the observed data. Usually, a model is imposed on  $f(x)$  characterized by a  $p$ -dimensional parameter, denoted by  $\theta$  throughout the paper, and a corresponding loss function is minimized to obtain an estimate of  $\theta$ . Write the loss function as  $l(d; \theta)$ , where  $d = (y, x)$ . The loss function can take the form of squared loss, negative log-likelihood, hinge loss, Huber loss, etc. Some classical choices include logistic regression with  $l(d; \theta) = -y(\alpha + \beta^\top x) + \log(1 + \exp(\alpha + \beta^\top x))$  for a binary response, where  $\theta = (\alpha, \beta^\top)$ , Poisson log-linear model with  $l(d; \theta) = -y(\alpha + \beta^\top x) + \exp(\alpha + \beta^\top x)$  for a counting response, linear model with squared loss  $l(d; \theta) = (y - \alpha - \beta^\top x)^2 / \sigma^2$  for a continuous response, where  $\theta = (\alpha, \beta^\top, \sigma)$ , and also nonlinear models such as neural networks with squared loss or cross-entropy loss. For unsupervised learning tasks, there is no response-covariate structure of data. Typically, we derive a loss function  $l(d; \theta)$  based on a certain purpose and aim at minimizing the loss with respect to  $\theta$ . For instance, in the geometric view of principal component analysis (PCA), the goal is to find a  $k$ -dimensional affine space in  $\mathbb{R}^q$  that best approximate the  $n$  examples in terms of Euclidean distance. Parameterizing the affine space by  $\alpha + U\beta$  where  $U$  consists of  $k$ -columns of an orthogonal basis of the space, we end up with the following loss function  $l(d; \theta) = \|d - (\alpha + U\beta)\|^2$ , where  $\theta = \{\alpha, U, \beta\}$  and  $\|\cdot\|$  stands for the Euclidean norm.

The target parameter that we aim to estimate is the so-called population risk minimizer which minimizes the population risk  $R(\theta) = \mathbf{E}[l(D; \theta)]$ , that is,

$$\theta^* = \arg \min_{\theta} R(\theta).$$

In supervised learning, when the regression model is correctly specified, that is, there exists some  $\theta_0$  satisfying  $f(x) = f_{\theta_0}(x)$ , then it is easy to see that  $\theta_0 = \theta^*$ .

However, as we mentioned in Section 1, in many real applications,  $f(x)$  does not satisfy the specified model and this is the so-called model mis-specification. Meanwhile, in some unsupervised learning tasks such as PCA, there is no imposed model. It is well-known that no matter the model specification holds or not,  $\theta^*$  can be well defined under general conditions. The full sample version of the risk minimizer, denoted by  $\hat{\theta}^*$ , is given by  $\hat{\theta}^* = \arg \min_{\theta} \sum_{i=1}^n l(d_i; \theta)$ . Under suitable regularity conditions, one can show that  $\hat{\theta}^*$  is consistent for  $\theta^*$ , that is,  $\hat{\theta}^*$  converges in probability to  $\theta^*$  [9].

When  $n$  is sufficiently large, subsampling designs are preferred to save computational cost. For each  $i = 1, \dots, n$ , let  $\Delta_i$  be a 0-1 valued binary indicator, indicating whether the  $i$ th data point is sampled into the subsample ( $\Delta_i = 1$ ) or not ( $\Delta_i = 0$ ). Let  $\tilde{\theta}$  be a pilot estimate. As we mentioned in Section 1, a pilot estimate is a guess of  $\theta^*$ . It can be obtained from a pilot sample that is either a subset of the full data or any other samples such as a historical sample. Let  $\pi_i$  be the conditional sampling probability of the  $i$ th data point given the full data and the pilot estimate, that is,  $\pi_i = P(\Delta_i = 1 | d_1, \dots, d_n; \tilde{\theta})$ . Also,  $\Delta_i$ 's are generated independently with each other given the full data and the pilot.

For binary response, the LCC design proposed by [7] sets  $\pi_i = |y_i - p(\tilde{\alpha} + \tilde{\beta}^\top x_i)|$ , where  $p(t) = \exp(t) / [1 + \exp(t)]$  and  $(\tilde{\alpha}, \tilde{\beta})$  is a pilot estimate of  $(\alpha, \beta)$ . After obtaining the subsample, they fit a logistic regression to the subsample and then do a simple adjustment to get the estimator of the target parameter. The main advantage of their estimation procedure is that it basically maintains the original form of the maximum likelihood estimation with full data. However, as we mentioned in Section 1, the validity of their approach heavily relies on the consistency of the pilot estimate, yet it has not been extended to other regression models or more general learning tasks. We apply the HT type estimation. Specifically, the HT type estimator is given by

$$\hat{\theta} = \arg \min_{\theta} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} l(d_i; \theta). \quad (2.1)$$

The HT type estimation is general enough for various kinds of loss functions. The computational complexity of (2.1) is similar to that of full data. For instance, if the objective function based on the full data is convex, so is the objective function in (2.1). Meanwhile, by inverting the sampling probability, it is quite intuitive to expect the consistency of the HT type estimator regardless of the consistency of the pilot estimate. More importantly, based on the HT type estimation, we can derive an optimal form of the subsampling design for a specific objective. The details of the derivation is discussed in the following section.

Before proceeding to the next section, we introduce some more necessary notation. Let  $g(d; \theta) = \partial l(d; \theta) / \partial \theta$ . When  $l$  is the negative log-likelihood function,  $g$  becomes the score vector. Let  $G(d; \theta) = \partial^2 l(d; \theta) / \partial \theta \partial \theta^\top$  and  $A = E[G(D; \theta^*)]$ . For any column vector  $a$ ,  $a^{\otimes 2}$  stands for  $aa^\top$ .

### 3. Surprise sampling design

To give out the specific form of the proposed subsampling design, we first heuristically present the asymptotic properties of the HT type estimator  $\hat{\theta}$ . By the definition of  $\hat{\theta}$ , it is easy to see that  $\sum_{i=1}^n \Delta_i g(d_i; \hat{\theta}) / \pi_i = 0$ . Under suitable conditions, we can show that  $\hat{\theta}$  is consistent for  $\theta^*$  and that

$$\sqrt{n}(\hat{\theta} - \theta^*) = -A^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} g(d_i; \theta^*) + o_p(1). \quad (3.1)$$

Furthermore, we can show that  $\sqrt{n}(\hat{\theta} - \theta^*)$  converges in distribution to a Gaussian vector, denoted by  $Z$ , with mean zero and variance-covariance matrix  $A^{-1}V_\pi A^{-1}$ , where  $V_\pi = \mathbb{E}[g(D; \theta^*)^{\otimes 2} / \pi]$  and  $\pi$  is a probability that may depend on  $D$  and  $\theta^*$  (i.e.,  $\pi = \pi(D; \theta^*)$ ). We give out sufficient conditions that guarantee the above results in Section 4.

#### 3.1. Overall prediction accuracy

If the main purpose of the data analysis is prediction, the minimized population prediction error  $R(\theta^*) = \mathbb{E}[l(D; \theta^*)]$  is a natural criterion to measure the prediction accuracy. Based on a subsampling percentage  $\pi$  and the corresponding HT type estimator  $\hat{\theta}$ , the prediction error can be measured by  $R(\hat{\theta})$ , which is  $R(\theta)$  evaluated at  $\theta = \hat{\theta}$ . By Taylor expansion, we have that

$$\begin{aligned} nR(\hat{\theta}) - nR(\theta^*) &= n\mathbb{E}[g(D; \theta^*)](\hat{\theta} - \theta^*) \\ &\quad + \frac{n}{2}(\hat{\theta} - \theta^*)^\top \mathbb{E}[G(D; \theta^*)](\hat{\theta} - \theta^*) + o_p(1) \\ &= \frac{n}{2}(\hat{\theta} - \theta^*)^\top A(\hat{\theta} - \theta^*) + o_p(1). \end{aligned} \quad (3.2)$$

The leading term in (3.2) converges in distribution to  $Z^\top AZ/2$  according to the asymptotic property of  $\hat{\theta}$ . By some calculation, we have that

$$\begin{aligned} \mathbb{E}(Z^\top AZ) &= \text{tr}[\mathbb{E}(Z^\top AZ)] = \text{tr}[A\mathbb{E}(ZZ^\top)] = \text{tr}(V_\pi A^{-1}) \\ &= \mathbb{E} \left[ \text{tr} \left( \frac{A^{-1/2} g(D; \theta^*)^{\otimes 2} A^{-1/2}}{\pi} \right) \right] \\ &= \mathbb{E} \left( \frac{\|A^{-1/2} g(D; \theta^*)\|^2}{\pi} \right). \end{aligned} \quad (3.3)$$

Then it is natural to select  $\pi$  that minimizes (3.3), which can be treated as the average difference between the prediction error based on the subsample and the minimal population prediction error. The following proposition, proved in Appendix A.1, gives out the optimal form of  $\pi$ .

**Proposition 1.** *Let  $r \in (0, 1)$  be a constant. The optimal  $\pi$  to minimize (3.3), subject to  $E(\pi) \leq r$ , is given by*

$$\pi = \left( c \|A^{-1/2}g(D; \theta^*)\| \right) \wedge 1,$$

where  $c$  is such that  $E(\pi) = r$  and for constants  $a$  and  $b$ ,  $a \wedge b = \min\{a, b\}$ .

The constant  $r$  is used to control the subsampling rate. Note that the optimal  $\pi$  depends on some unknown quantities, such as  $A$  and  $\theta^*$ . Given a pilot estimate  $\tilde{\theta}$ , we propose the following subsampling design

$$\pi_i = \left( c \left\| \tilde{A}^{-1/2}g(d_i; \tilde{\theta}) \right\| \right) \wedge 1, \tag{3.4}$$

$i = 1, \dots, n$ , where  $\tilde{A} = n^{-1} \sum_{i=1}^n G(d_i; \tilde{\theta})$  and  $c$  is a constant selected to reach the predetermined subsampling rate. For the given subsampling rate  $r$ , an algorithm based on the bisection method to get the largest  $c$  such that  $n^{-1} \sum_{i=1}^n \pi_i \leq r$  is given in Section 3.3.

The proposed subsampling probability is proportional to the score evaluated at the pilot. When the value is relatively large for a data point, it implies that this data point has a large prediction error for the given pilot estimate, that is, it is somewhat more “surprising” than the ones with smaller score values. Under such design, the data point with a larger score value is more likely to be drawn into the subsample. Thus, we call the proposed subsampling design “surprise” sampling, or score sampling.

### 3.2. Estimation accuracy

If the main concern is the estimation accuracy of  $v^\top \hat{\theta}$ , as an estimator of  $v^\top \theta^*$ , where  $v$  is a given  $p$ -dimensional constant vector, then the objective becomes to give out an efficient subsampling scheme to increase the estimation accuracy. Without loss of generality, assume that  $\|v\| = 1$ . From (3.1), we have that  $\sqrt{n}v^\top(\hat{\theta} - \theta^*)$  converges in distribution to a Gaussian vector with mean zero and variance  $v^\top A^{-1}V_\pi A^{-1}v$ . Then we select  $\pi$  to minimize this asymptotic variance.

**Proposition 2.** *Let  $r \in (0, 1)$  be a constant. The optimal  $\pi$  to minimize  $v^\top A^{-1}V_\pi A^{-1}v$ , subject to  $E(\pi) \leq r$ , is given by*

$$\pi = \left( c |v^\top A^{-1}g(D; \theta^*)| \right) \wedge 1,$$

where  $c$  is such that  $E(\pi) = r$ .

Given a pilot  $\tilde{\theta}$ , the proposed subsampling design is given by

$$\pi_i = \left( c \left| v^\top \tilde{A}^{-1}g(d_i; \tilde{\theta}) \right| \right) \wedge 1, \tag{3.5}$$

$i = 1, \dots, n$ , where  $c$  is a constant selected to reach the predetermined subsampling rate. Again, this design is proportional to the score evaluated at  $\tilde{\theta}$ .

To elaborate the relationship between the proposed surprise sampling and the LCC sampling, we discuss the case of generalized linear models with  $D = (Y, X)$  and  $d_i = (y_i, x_i)$ . Let  $Z = (1, X^\top)^\top$  and  $\theta$  be the collection of all the regression parameters (i.e., the coefficients corresponding to  $X$  plus the intercept). For a generalized linear model, the conditional probability or density function of  $Y$  given  $X = x$ , denoted by  $p(y|x; \theta)$ , can be written as the form of  $\psi(y, \theta^\top z)$ , where  $\psi$  is a known function up to a finite-dimensional parameter and  $z = (1, x^\top)^\top$ . A natural choice of the loss function is the negative log likelihood  $l(y, x; \theta) = -\log \psi(y, \theta^\top z)$ . Define  $S(y, t) = \frac{\partial}{\partial t} \psi(y, t) / \psi(y, t)$ . Then it is easy to see that  $g(y, x; \theta) = -S(y, \theta^\top z)z$ . When the model is correctly specified, there exists  $\theta_0 = \theta^*$ , known as the true parameter value. Let  $\sigma_Z = \text{Var}[S(Y, \theta_0^\top Z)|Z]$ .

**Proposition 3.** *Set  $v = \mathbb{E}(\sigma_Z Z)$ . Then the optimal  $\pi$  to minimize the asymptotic variance of  $v^\top \hat{\theta}$  is given by  $\pi = c|S(Y, \theta_0^\top Z)| \wedge 1$ , where  $c$  is a constant controlling the subsampling rate.*

Based on Proposition 3, given a pilot estimate  $\tilde{\theta}$ , we would propose the subsampling design  $\pi_i = c|S(y_i, \tilde{\theta}^\top z_i)| \wedge 1$ , where  $z_i = (1, x_i^\top)^\top$ ,  $i = 1, \dots, n$ . For the logistic model with a binary response, we have  $\psi(y, t) = p(t)^y(1 - p(t))^{1-y}$  and  $S(y, t) = (y - p(t))$ . Then the proposed subsampling design becomes  $\pi_i = c|y_i - p(\tilde{\theta}^\top z_i)| \wedge 1$ , which is exactly the LCC sampling proposed by [7]. Thus, the LCC sampling can be viewed as a special case of the surprise sampling design, with a somewhat narrow objective to achieve optimality in estimating a certain direction of the regression parameter vector. Surprise sampling, however, is better justified, since if the data analysis objective is altered, the sampling design can be adaptively modified. We show in the next section that if logistic model is correctly specified and the pilot estimate is consistent, the proposed HT type estimator  $\hat{\theta}$  has the same asymptotic efficiency as the LCC estimator of [7].

Besides optimizing the prediction accuracy and estimation accuracy, other objectives can also be considered. The optimal selection probability  $\pi$  adaptive to various objectives can be derived similarly to those in Propositions 1 and 2. For instance, [1] and [17] mainly considered minimizing the conditional MSE of  $\hat{\theta}$  given the full data. In our proposal, for minimizing the (unconditional) MSE, the optimal design is given by  $\pi = (c\|A^{-1}g(D; \theta^*)\|) \wedge 1$ , which is similar to the optimal form derived by [17] under the logistic model and [1] under the generalized linear models.

### 3.3. The algorithms

Let  $\tilde{\pi}_i$  be the kernel part of the proposed surprise sampling design, e.g.,  $\tilde{\pi}_i = \|\tilde{A}^{-1/2}g(d_i; \tilde{\theta})\|$  for prediction and  $\tilde{\pi}_i = |v^\top \tilde{A}^{-1}g(d_i; \tilde{\theta})|$  for estimation. The implementation procedure of the surprise sampling is summarized in Algorithm 1. The constant  $c$  is obtained by Algorithm 2.

---

**Algorithm 1** Surprise Sampling

---

**Input:** data  $d_i, i = 1, \dots, n$ , a loss function  $l(d; \theta)$ , and a subsampling rate  $r$ .

**Output:** the HT type estimate  $\hat{\theta}$ .

- 1: Get a pilot estimate  $\hat{\theta}$ .
  - 2: For  $i = 1, \dots, n$ , do:
    - evaluate  $\tilde{\pi}_i$  and  $\pi_i = c\tilde{\pi}_i \wedge 1$ , where  $c$  is obtained by Algorithm 2;
    - generate independent  $\Delta_i \sim \text{Bernoulli}(\pi_i)$ .
  - 3: Obtain a subsample  $\{d_i : \Delta_i = 1\}$ .
  - 4: Compute  $\hat{\theta}$  by minimizing  $\sum_{i=1}^n \Delta_i l(d_i; \theta) / \pi_i$ .
- 

---

**Algorithm 2** Find  $c$  to reach the predetermined subsampling rate  $r$

---

**Input:**  $\{\tilde{\pi}_i, i = 1, \dots, n\}$  and the subsampling rate  $r$ .

**Output:** the constant  $c$  to control the subsample size.

- Rearrange  $\tilde{\pi}_i$ 's in ascending order  $\tilde{\pi}_{(1)} \leq \dots \leq \tilde{\pi}_{(n)}$ , and compute  $c_0 = nr / \sum_{i=1}^n \tilde{\pi}_i$ .
- if**  $c_0 \tilde{\pi}_{(n)} \leq 1$  **then**  
 Set  $c = c_0$ .
- else**  
 Find  $m$  such that  $f(1/\tilde{\pi}_{(m+1)}) \leq nr \leq f(1/\tilde{\pi}_{(m)})$  by the bisection method, where  $f(x) = \sum_{i=1}^n (x\tilde{\pi}_i \wedge 1)$ .  
 Set  $c = [nr - (n - m)] / \sum_{i=k+1}^m \tilde{\pi}_{(i)}$ .
- end if**
- 

**4. Large sample properties**

In this section we discuss the asymptotic properties of the HT type estimator under the proposed surprise sampling designs, which require some conditions listed below. The first set of conditions includes the regularity conditions on the parameter space and the underlying data distribution.

- C1.  $\Theta$  is compact and contains  $\theta^*$  as an interior point.
- C2.  $R(\theta)$  is continuous in  $\theta$  and  $\theta^*$  is the unique global minimizer of  $R(\theta)$  over  $\Theta$ .
- C3. There exists a function  $h(D)$  such that  $\sup_{\theta \in \Theta} \|l(D; \theta)\| \leq h(D)$ ,  $\sup_{\theta \in \Theta} \|g(D; \theta)\| \leq h(D)$ ,  $\sup_{\theta \in \Theta} \|G(D; \theta)\| \leq h(D)$ , and  $E[h(D)] < \infty$ .  $G(D; \theta)$  is continuous in  $\theta$  with probability one.
- C4. The matrix  $A$  is positive definite.

The second set of conditions is about the subsampling probabilities  $\pi_i$ 's.

- C5.  $n^{-1} \sum_{i=1}^n (\sup_{\theta \in \Theta} l(d_i; \theta))^2 / \pi_i = O_p(1)$ , where  $\Theta \in \mathbb{R}^p$  is the parameter space.
- C6.  $n^{-1} \sum_{i=1}^n \|G(d_i; \theta^*)\|^2 / \pi_i = O_p(1)$ .
- C7. There exists a probability  $\pi$  that may depend on  $D$  so that  $E(\|g(D; \theta^*)\|^2 / \pi) < \infty$  and  $n^{-1} \sum_{i=1}^n (1/\pi_i - 1)g(d_i; \theta^*)^{\otimes 2} \xrightarrow{p} E[(1/\pi - 1)g(D; \theta^*)^{\otimes 2}]$ .

Note that the conditions on the subsampling probabilities are stated in a general manner so that we can first present a general theorem on the large sample theory of  $\hat{\theta}$  defined in (2.1). The theorem shows the consistency and the asymptotic normality of  $\hat{\theta}$  under the above conditions.

**Theorem 1.** *If conditions C1–C7 hold, then 1)  $\hat{\theta} \xrightarrow{p} \theta^*$ , where  $\xrightarrow{p}$  means converging in probability (i.e.,  $\hat{\theta}$  is consistent) and 2)  $\sqrt{n}(\hat{\theta} - \theta^*) \xrightarrow{d} N(0, A^{-1}V_{\pi}A^{-1})$ , where  $\xrightarrow{d}$  means converging in distribution and  $N(\mu, \Sigma)$  stands for a normal vector with mean  $\mu$  and variance-covariance matrix  $\Sigma$ .*

The key step to getting the asymptotic distribution is to establish the expansion (3.1). We give out the proof in Appendix A.2. It is worthwhile to mention that the consistency and the asymptotic normality of  $\hat{\theta}$  do not require the pilot  $\tilde{\theta}$  to be consistent nor to be independent of the full data, which, however, are both required in [7]. The relaxations can be useful in practice, as we mentioned in Section 1. This is one of the main reasons that we propose to use the HT type estimator.

Based on Theorem 1, we can immediately derive the asymptotic properties of the HT type estimator under the proposed surprise sampling designs. For the surprise sampling, the specific forms of the design  $\pi_i$  are given in (3.4) or (3.5), depending on the data analysis purpose. We first consider the situation where the pilot  $\tilde{\theta}$  is consistent. Here, the probability  $\pi = (c\|A^{-1/2}g(D; \theta^*)\|) \wedge 1$  for prediction and  $\pi = (c|v^{\top}A^{-1}g(D; \theta^*)|) \wedge 1$  for estimation. We further introduce a condition which, together with C1–C4, can deduce the general conditions C5–C7.

C8.  $E(\sup_{\theta \in \Theta} \|l(D; \theta)\|^2/\pi) < \infty$  and  $E(\|G(D; \theta^*)\|^2/\pi) < \infty$ .

The following corollary presents the asymptotic properties of  $\hat{\theta}$  with the consistent pilot estimate.

**Corollary 1.** *If conditions C1–C4 and C8 hold and  $\tilde{\theta} \xrightarrow{p} \theta^*$  (i.e.,  $\tilde{\theta}$  is consistent), then 1)  $\hat{\theta} \xrightarrow{p} \theta^*$  (i.e.,  $\hat{\theta}$  is consistent) and 2)  $\sqrt{n}(\hat{\theta} - \theta^*) \xrightarrow{d} N(0, A^{-1}V_{\pi}A^{-1})$ .*

It seems that Corollary 1 only repeats the results of Theorem 1. However, one needs to notice that in Corollary 1 the probability function  $\pi$  in  $V_{\pi}$  takes the optimal form given in Proposition 1 and 2. It means that when the pilot estimate is consistent, the proposed sampling designs given in (3.4) or (3.5) are optimal in the sense that the corresponding HT type estimator  $\hat{\theta}$  asymptotically reaches the best prediction accuracy or estimation accuracy.

When the assumed model is correctly specified, the target parameter  $\theta^*$  becomes the true parameter  $\theta_0$ . For the binary response problem discussed in [7], we have shown in Proposition 3 that the LCC sampling is a special case of the proposed surprise sampling. Consider the sampling design  $\pi_i = |y_i - p(\tilde{\theta}^{\top} z_i)|$ , that is, the constant  $c = 1$ . We have the following result for the corresponding  $\hat{\theta}$ .

**Corollary 2.** *If the logistic regression model  $p(\theta^{\top} z) = \exp(\theta^{\top} z)/[1 + \exp(\theta^{\top} z)]$  holds with  $\theta_0$  being the true parameter value,  $X$  does not concentrate on a hyperplane of dimension smaller than  $q$ ,  $E(\|X\|^4/\pi) < \infty$ , and  $\tilde{\theta} \xrightarrow{p} \theta_0$ , then  $\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, 2\Sigma_{\text{full}})$ , where  $\Sigma_{\text{full}}$  stands for the asymptotic variance of the MLE for the full sample.*

Corollary 2 indicates that under the correctly specified logistic model, the proposed HT type estimator  $\hat{\theta}$  is asymptotically as efficient as the LCC estimator of [7], as long as the pilot is consistent. Again, our result does not require the independence between the pilot and the full data. Both Corollaries 1 and 2 are proved in Appendix A.2.

Next we consider the situation with an inconsistent pilot estimate. When the pilot estimate is not consistent, the LCC estimator is no longer consistent. However, based on Theorem 1, the HT type estimator can be still consistent and asymptotically normal as long as the pilot has a certain limit in probability. Specifically, suppose that  $\tilde{\theta}$  converges in probability to a limit denoted by  $\bar{\theta}$ . Let  $\bar{A} = E[G(D; \bar{\theta})]$  and  $V_{\bar{\pi}} = E[g(D; \theta^*)^{\otimes 2} / \bar{\pi}]$ , where  $\bar{\pi} = (c \|\bar{A}^{-1/2} g(D; \bar{\theta})\|) \wedge 1$  for prediction and  $\bar{\pi} = (c |v^\top \bar{A}^{-1} g(D; \bar{\theta})|) \wedge 1$  for estimation. Another condition is needed here.

C9.  $\bar{\theta}$  is an interior point of  $\Theta$ , the matrix  $\bar{A}$  is positive definite,  $E(\sup_{\theta \in \Theta} \|l(D; \theta)\|^2 / \bar{\pi}) < \infty$ ,  $E(\|g(D; \theta^*)\|^2 / \bar{\pi}) < \infty$  and  $E(\|G(D; \theta^*)\|^2 / \bar{\pi}) < \infty$ .

For the inconsistent pilot, we have the following result.

**Corollary 3.** *If conditions C1–C4 and C9 hold and  $\tilde{\theta} \xrightarrow{P} \bar{\theta}$ , then 1)  $\hat{\theta} \xrightarrow{P} \theta^*$  (i.e.,  $\hat{\theta}$  is consistent) and 2)  $\sqrt{n}(\hat{\theta} - \theta^*) \xrightarrow{d} N(0, A^{-1} V_{\bar{\pi}} A^{-1})$ .*

Again we give out the proof in Appendix A.2. When the pilot estimate is not consistent for  $\theta^*$ , the surprise sampling design is no longer optimal. However,  $\hat{\theta}$  is still consistent for  $\theta^*$  and asymptotically normally distributed. Thus, the proposed estimation procedure is robust for the model mis-specification as well as the inconsistent pilot.

Finally, a plugged-in approach can be applied to estimate the asymptotic variance-covariance matrix of  $\hat{\theta}$ . Specifically, we define

$$\hat{A} = n^{-1} \sum_{i=1}^n \Delta_i G(d_i, \hat{\theta}) / \pi_i \text{ and } \hat{V}_{\pi} = n^{-1} \sum_{i=1}^n \Delta_i g(d_i; \hat{\theta})^{\otimes 2} / \pi_i^2.$$

Then a consistent estimator for the asymptotic variance-covariance matrix is given by  $\hat{A}^{-1} \hat{V}_{\pi} \hat{A}^{-1}$ .

### 5. Simulation studies

We conduct extensive simulation studies to examine the effectiveness of the proposed approach and make some comparison with the LCC sampling. The numerical studies mainly focus on various regression type problems, but as we have already mentioned, the whole procedure is general enough to be extended to unsupervised learning tasks. We first consider binary response where the proposed HT type estimator can be compared with the LCC estimator of [7].

TABLE 1  
 Comparison of LCC and HT estimate under the correctly specified logistic model.

Pilot Estimation	Uniform MLE		WCC	
	$\widehat{\text{Bias}}^2 (\times 10^6)$	$\widehat{\text{Var}} (\times 10^3)$	$\widehat{\text{Bias}}^2 (\times 10^6)$	$\widehat{\text{Var}} (\times 10^3)$
LCC	3.888	3.682	3.347	3.613
HT	4.052	3.767	3.477	3.676

LCC: local case-control estimate; HT: Horvitz-Thompson type estimate; Uniform MLE: MLE with uniform sampling as pilot; WCC: weighted case-control estimate as pilot;  $\widehat{\text{Bias}}^2$ :  $\widehat{\text{Bias}}^2 = \|\mathbf{E}(\hat{\beta}) - \beta_0\|^2$ ;  $\widehat{\text{Var}}$ :  $\widehat{\text{Var}} = \sum_{j=1}^q \text{Var}(\hat{\beta}_j)$ .

### Simulation 1: Correctly specified logistic model

In simulation 1 we consider the scenario where the logistic model is correctly specified. We set  $q = 50$  and all the predictors  $X$  are generated independently from the standard normal distribution. Given  $X$ ,  $Y$  is generated from a Bernoulli distribution with success probability  $p(\theta^\top Z)$ , where  $\theta = (\alpha, \beta^\top)^\top$  with the first 25 components of  $\beta$  being 1, the rest 25 being 0, and  $\alpha$  being chosen to yield  $\text{P}(Y = 1) \approx 10\%$ . The entire sample size  $n$  is set to be  $10^6$ . For the pilot estimate, we use data drawn from the full sample with size  $10^4$ . We consider two types of consistent pilot estimates. The first one is to draw a random sample with uniform probabilities and get the logistic MLE to be the pilot; the second one is to draw a 50-50 split case-control sample and apply the weighted case-control approach to get the pilot. For the subsampling design, we apply the LCC sampling  $\pi_i = |y_i - p(\tilde{\theta}^\top z_i)|$ , which is a special case of the proposed surprise sampling. Both the LCC estimator and the proposed HT type estimator based on the negative log-likelihood loss are obtained for comparison. The procedure is repeated for 1000 times. We record the squared bias and variance of the estimator for  $\beta$ , denoted by  $\hat{\beta}$ , over the 1000 realizations for each of the two methods under the two pilots, respectively. The results are summarized in Table 1.

From Corollary 2, we know that when the logistic model is correctly specified, the proposed HT estimator is asymptotically as efficient as the LCC estimator under the LCC sampling. From Table 1, we see that the behaviors of the two estimates are very close to each other, which coincides with the finding of the corollary. Also, the pilot method has little impact on the behavior of the proposed estimator, as long as it is consistent.

It is worthwhile to mention that under the LCC sampling, the average size of the subsample is around 6.6% of the entire sample size. Adding the sample size used for the pilot, it means that we use about 7.6% of the full data size to reach half of the estimation efficiency. We also record the computation time. On the computer we conduct all the numerical studies (a laptop running macOS 10.14 with an Intel I7 processor and 16GB memory), it takes around 352 seconds to calculate a full sample MLE. The whole procedure of the subsample estimation, including obtaining the pilot, evaluating the sampling probabilities, drawing the subsample, and calculating the HT type estimate, however, takes around 19 seconds, only 5.4% of the time for the full sample MLE.

TABLE 2  
 Comparison of LCC and HT estimate under the incorrectly specified logistic model with consistent pilot estimate.

Pilot Estimation	Uniform MLE		WCC	
	$\widehat{\text{Bias}}^2 (\times 10^6)$	$\widehat{\text{Var}} (\times 10^3)$	$\widehat{\text{Bias}}^2 (\times 10^6)$	$\widehat{\text{Var}} (\times 10^3)$
LCC	3.170	1.643	4.222	1.226
HT	4.629	1.045	3.481	1.040

LCC: local case-control estimate; HT: Horvitz-Thompson type estimate; Uniform MLE: MLE with uniform sampling as pilot; WCC: weighted case-control estimate as pilot;  $\widehat{\text{Bias}}^2$ :  $\widehat{\text{Bias}}^2 = \|\mathbf{E}(\hat{\beta}) - \beta^*\|^2$ ;  $\widehat{\text{Var}}$ :  $\widehat{\text{Var}} = \sum_{j=1}^q \text{Var}(\hat{\beta}_j)$ .

*Simulation 2: Incorrectly specified logistic model with a consistent pilot*

In simulation 2 we turn to the scenario where the logistic model is incorrectly specified. Here  $q$  is set to be 5 and again all the predictors are generated independently from the standard normal distribution. Besides the linear combination of the 5 predictors, we add the quadratic term of the first predictor into the success probability of the Bernoulli distribution, but still fit the data by the usual logistic model with linear terms only. Thus, the fitted model is incorrectly specified. The parameters are set to yield  $P(Y = 1) \approx 1\%$ , so the data is more imbalanced than the last scenario. The entire sample size  $n = 10^6$ , and the pilot sample size is  $10^4$ . We still use the logistic MLE with uniform subsampling and the weighted case-control estimate as the pilot estimate, respectively. It is easy to see that the two pilots are consistent. Similar to simulation 1, the LCC sampling is applied and the LCC estimator and the HT estimator are obtained. The procedure is repeated for 1000 times and the results parallel to Table 1 are summarized in Table 2.

From the results we see that the LCC estimate and the HT estimate are essentially unbiased for the target parameter  $\beta^*$ . Meanwhile, in this case, the empirical variance of the proposed HT estimate is smaller than that of the LCC estimate. Under the LCC sampling, the average size of the subsample is around 1.9% of the entire sample size.

*Simulation 3: Incorrectly specified logistic model with an inconsistent pilot*

In simulation 3 we compare the two estimators under incorrectly specified logistic with an inconsistent pilot estimate. The data generation scheme is the same as that in simulation 2. For the pilot estimate, we still draw a subsample of size  $10^4$  with uniform probabilities. However, the difference here is that we use the probit MLE to be the pilot estimate. Then the resulting pilot estimate is no longer consistent for the target parameter. Using this inconsistent pilot, the LCC sampling is applied and the LCC estimator and the HT estimator are obtained. The procedure is repeated for 1000 times. Here, besides the results of the two estimates, we also report the squared bias and variance of the pilot estimate. The results are summarized in Table 3.

The consistency of the LCC estimator relies on the consistency of the pilot, but the that of the proposed HT estimator does not, as Corollary 3 illustrates. From Table 3, we see that the pilot estimate is clearly biased. Compared with

TABLE 3  
 Comparison of LCC and HT estimate under the incorrectly specified logistic model with inconsistent pilot estimate.

Estimation	$\widehat{\text{Bias}}^2 (\times 10^4)$	$\widehat{\text{Var}} (\times 10^4)$
Pilot <sup>P</sup>	935.5	102.3
LCC	6.432	7.138
HT	0.042	7.273

Pilot<sup>P</sup>: pilot estimate with the probit fitting; LCC: local case-control estimate; HT: Horvitz-Thompson type estimate;  $\widehat{\text{Bias}}^2$ :  $\widehat{\text{Bias}}^2 = \|\mathbf{E}(\hat{\beta}) - \beta^*\|^2$ ;  $\widehat{\text{Var}}$ :  $\widehat{\text{Var}} = \sum_{j=1}^q \text{Var}(\hat{\beta}_j)$ .

the HT estimator, the bias of the LCC estimator is quite obvious. Although the squared bias of the LCC estimator is not large in the absolute value, it is of the same order with its variance. The proposed HT estimator, however, has much smaller squared bias relative to its variance, showing positive evidence for Corollary 3. This is an advantage of the HT estimator over the LCC one.

*Simulation 4: Comparison of LCC sampling and surprise sampling*

Proposition 3 shows that the LCC sampling is a special case of the proposed surprise sampling design, and is optimal when the target is to estimate  $v^\top \theta^*$  with  $v = \mathbf{E}(\sigma_Z Z)$ . However, when the target direction changes, the LCC sampling is not adaptively optimal, but our surprise sampling design is. In this simulation, we again generate data by the same scheme as that in simulation 2, but suppose that the main concern is to estimate the regression coefficient of the first predictor, i.e.,  $\beta_1^*$ . Equivalently speaking, the target direction  $v$  here is a 6-dimensional column vector with the second element being 1 and all the others being 0. For this target, the LCC sampling is not optimal. Following Proposition 2, the optimal design is given by  $\pi_i = (c|v^\top \tilde{A}^{-1} z_i (y_i - p(\tilde{\theta}^\top z_i))|) \wedge 1$ , where  $\tilde{A} = n^{-1} \sum_{i=1}^n p(\tilde{\theta}^\top z_i) (1 - p(\tilde{\theta}^\top z_i)) z_i z_i^\top$ . The two subsampling designs are compared, with  $c$  decided to yield comparable average subsample sizes between the two designs. We still use the same two pilot estimates as simulation 2 and consider three estimation approaches: the LCC estimator, the HT estimator under the LCC sampling, and the HT estimator under the optimal surprise sampling. The procedure is repeated for 1000 times. For the HT estimator, besides the squared bias and variance, we also report the average of the variance estimates and the empirical coverage percentage of the 95% Wald confidence interval. The results are summarized in Table 4.

The variance of the HT estimator under the optimal surprise sampling is much smaller than that of the LCC estimator. By using comparable subsample sizes (around 1.9% of the entire sample size), the relative efficiency of the HT estimator under the optimal sampling to the LCC estimator is at least 3 in our situation. Also, the HT estimator under the optimal sampling is more efficient than that under the LCC sampling, which is expected according to Proposition 2. The plugged-in variance estimate and the Wald confidence interval give out reasonable performance.

In the remaining studies we turn to more general regression models with counting and continuous responses, where the LCC approach is no longer available.

TABLE 4  
*Estimation of  $\beta_1^*$  by using local case-control sampling and surprise sampling under the logistic model.*

Pilot	Uniform MLE				WCC			
	Bias <sup>2</sup> ( $\times 10^6$ )	Var ( $\times 10^4$ )	Var Est. ( $\times 10^4$ )	CP (%)	Bias <sup>2</sup> ( $\times 10^6$ )	Var ( $\times 10^4$ )	Var Est. ( $\times 10^4$ )	CP (%)
LCC	0.473	7.799	–	–	1.295	3.417	–	–
HT-LCC	1.669	1.684	1.444	94.1	0.139	1.540	1.394	94.4
HT-optimal	0.577	1.134	0.978	92.0	0.065	1.129	0.943	92.5

LCC: local case-control estimate; HT-LCC: Horvitz-Thompson type estimate under the local case-control sampling; HT-optimal: Horvitz-Thompson type estimate under the optimal surprise sampling; Bias<sup>2</sup>: average of the squared bias; Var: empirical variance; Var Est.: average of the variance estimate; CP: coverage probability of the 95% Wald confidence interval; Uniform MLE: MLE with uniform sampling as pilot; WCC: weighted case-control estimate as pilot.

*Simulation 5: Counting response under log-linear model*

In simulation 5 we consider the counting response. We set  $q = 2$  and the two predictors  $X = (X_1, X_2)$  are generated independently from the standard normal distribution. Two schemes for generating  $Y$  given  $X$  are considered. In the first one,  $Y$  is generated from a Poisson distribution with mean  $\exp(\alpha + \beta_1 X_1 + \beta_2 X_2)$ , while in the second one,  $Y$  is generated from a Poisson distribution with mean  $\exp(\alpha + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_1^2)$ . Thus, the Poisson log-linear model only with the linear terms is correctly specified for the first scheme, but misspecified for the second. In both schemes, the parameters are set to yield  $P(Y \leq 1) \approx 93\%$ , implying that the generated data is imbalanced in the sense that most responses are 0 or 1. The entire sample size  $n = 10^5$ . We draw a pilot sample of size 1000 from the entire sample with uniform probabilities and fit the Poisson log-linear MLE to be the pilot, denoted by  $\tilde{\theta} = (\tilde{\alpha}, \tilde{\beta}_1, \tilde{\beta}_2)^\top$ . According to Proposition 3, we use the surprise sampling design  $\pi_i = (c|y_i - \exp(\tilde{\theta}^\top z_i)|) \wedge 1$ , where  $z_i = (1, x_{1i}, x_{2i})^\top$  and  $c$  is set to make the subsample size equal to 1000. The proposed HT type estimator based on the negative log-likelihood loss is calculated. For comparison, we draw a uniform subsample of size 2000 to calculate the Poisson log-linear MLE. The size is 2000 since the proposed surprise sampling needs to pay for its pilot sample. The procedure is repeated for 1000 times. The squared bias and variance of the two estimators are recorded. For the HT estimator, we also record the average of the variance estimates and the empirical coverage percentage of the 95% Wald confidence interval. The results are summarized in Table 5.

From the results, we find out that no matter whether the model is correctly specified or not, the HT estimator based on the proposed surprise sampling is essentially unbiased for the target parameter, and is much more efficient than the uniform subsample MLE with a comparable sample size. The plugged-in variance estimate and the Wald confidence interval give out satisfactory performances.

TABLE 5  
*Estimation results under the Log-linear model.*

Model specification	Estimation Parameter	Sub-MLE		HT			
		Bias <sup>2</sup> ( $\times 10^6$ )	Var ( $\times 10^3$ )	Bias <sup>2</sup> ( $\times 10^6$ )	Var ( $\times 10^3$ )	Var Est. ( $\times 10^3$ )	CP (%)
Correct	$\alpha$	94.49	8.414	3.145	4.334	4.115	94.4
	$\beta_1$	5.979	3.200	0.307	1.034	0.963	94.6
	$\beta_2$	0.822	3.083	2.292	0.958	0.954	95.5
Incorrect	$\alpha$	45.02	11.95	12.09	5.041	4.976	94.2
	$\beta_1$	20.79	4.437	0.267	1.437	1.528	95.7
	$\beta_2$	4.642	6.372	8.191	1.714	1.774	95.4

Full MLE: full sample MLE; Sub-MLE: MLE with uniform subsample of size 2000; HT: Horvitz-Thompson type estimate under the surprise sampling; Bias<sup>2</sup>: average of the squared bias; Var: empirical variance; Var Est.: average of the variance estimate; CP: coverage probability of the 95% Wald confidence interval.

TABLE 6  
*Estimation results under the linear model.*

Model specification	Estimation Parameter	Sub-MLE		HT			
		Bias <sup>2</sup> ( $\times 10^8$ )	Var ( $\times 10^4$ )	Bias <sup>2</sup> ( $\times 10^8$ )	Var ( $\times 10^4$ )	Var Est. ( $\times 10^4$ )	CP (%)
Correct	$\alpha$	2.178	0.103	1.482	0.067	0.066	94.5
	$\beta_1$	0.774	10.77	8.392	6.651	6.665	95.5
	$\beta_2$	66.70	10.05	24.05	6.667	6.712	95.9
Incorrect	$\alpha$	2.459	0.104	2.016	0.067	0.066	94.6
	$\beta_1$	1.890	10.82	53.54	6.881	6.779	94.8
	$\beta_2$	5.766	10.14	2.796	6.745	6.818	95.9

Sub-MLE: MLE with uniform subsample of size 2000; HT: Horvitz-Thompson type estimate under the surprise sampling; Bias<sup>2</sup>: average of the squared bias; Var: empirical variance; Var Est.: average of the variance estimate; CP: coverage probability of the 95% Wald confidence interval.

*Simulation 6: Continuous response under linear model*

In the last simulation we consider the continuous response. Still  $q$  is set to be 2, and the two predictors  $X = (X_1, X_2)$  are generated independently from a normal distribution with zero mean and variance 0.01. Two schemes for generating  $Y$  given  $X$  are considered. In the first one,  $Y$  is generated from a normal distribution with mean  $\alpha + \beta_1 X_1 + \beta_2 X_2$  and variance 0.01, while in the second one,  $Y$  is generated from Poisson distribution with mean  $\alpha + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_1^2$  and variance 0.01. Thus, the linear regression model only with the linear terms is correctly specified for the first scheme, but misspecified for the second. In both schemes, the parameters are set to yield  $P(-0.5 < Y < 0.5) \approx 99.6\%$ , implying that most responses are near zero. The entire sample size  $n = 10^5$ . We draw a uniform pilot sample of size 1000 from the entire sample and fit the normal linear regression MLE to be the pilot  $\hat{\theta}$ . For the surprise sampling, we use  $\pi_i = (c|y_i - \hat{\theta}^\top z_i|) \wedge 1$ , where  $c$  is set to make the subsample size equal to 1000. The proposed HT type estimator based on the least squared loss is calculated. Similar to the last simulation, the MLE based on a uniform subsample of size 2000 is also obtained for comparison. The parallel results are summarized in Table 6.

The observations of the results are basically similar to those of Table 5, so we omit repeating the details. The last two simulations show that the proposed surprise sampling and the corresponding HT estimator are quite effective for counting and continuous data.

## 6. Applications

### *Web Spam data*

We first apply the proposed approach to the Web Spam data available on the LIBSVM website and originally from [18]. The same data was also analyzed by [7] using the LCC approach for spam filtering. The data consists of 350000 web pages with about 60% are labeled as “web spam” that designed to manipulate search engines rather than display legitimate content. Following [7], we use frequency of the 99 unigrams that appeared in at least 200 documents as features, log-transformed with an offset to reduce skew. This data set is marginally balanced, but has considerable conditional imbalance in the sense that for some feature values, the web label is easy to predict.

We keep the same subsampling procedure as that in [7]. The LCC sampling design is used and the corresponding selection percentage is about 10% for this data. To assess the sampling distribution of the estimators, we repeatedly take uniform subsample of size  $n = 100000$  from the full data for 100 times. In each replication, we use the weighted case-control approach adopted in *Simulation 1* with size 10000 to get the pilot estimate. By the LCC design, the subsample size for parameter estimation is also around 10000. We fit the full sample MLE (of size 100000), the LCC estimate, and the proposed HT estimate. Following Corollary 2, under the Logistic model, the asymptotic variance of the proposed HT estimate is twice the variance of the full sample MLE, and is the same as that of the LCC estimate. Since in this real data, it is very likely the logistic model is incorrectly specified, as [7] mentioned, the asymptotic variance of the HT estimate and the LCC estimate should be a little bit more than two times the variance of the MLE. On the other hand, the subsample size is around 20000 (including the sample for pilot estimate), i.e., 20% of the full sample size. Then a uniform subsample of size 20000 should yield variance roughly 5 times that of the full sample.

Figure 1 shows the relative variances of the subsampling estimates to the full sample MLE. Specifically, the horizontal axis indexes the 100 regression coefficients to fit, and the vertical axis stands for the variance of each estimated coefficient relative to that of the full-sample MLE of the same coefficient. The triangle is for the HT estimate and the circle is for the LCC estimate. We find that most relative variances of the two subsampling estimates are slightly larger than 2, as expected, but are substantially smaller than 5, implying the LCC sampling is more efficient than the uniform subsampling with a comparable sample size. Moreover, for most estimated coefficients, the relative variance of the HT estimate is smaller than that of the LCC estimate, implying that the proposed HT estimator is slightly more efficient than the LCC estimator. This is con-

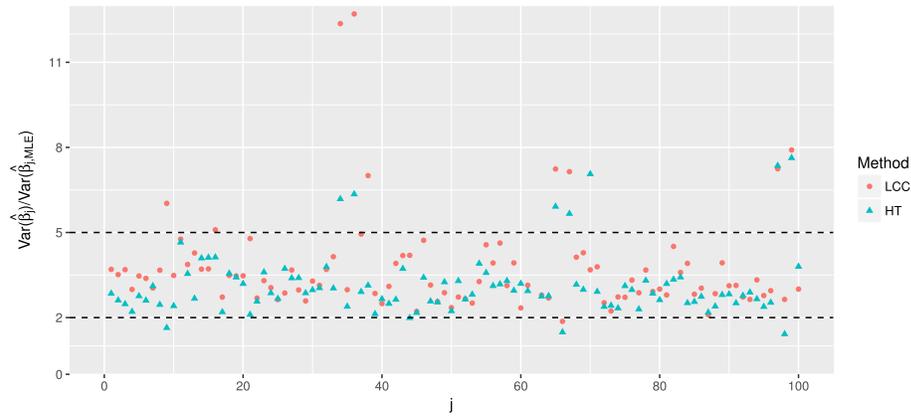


FIG 1. Relative variance of the subsampling estimate ( $\hat{\beta}_j$ ) to the full sample MLE ( $\hat{\beta}_{j,MLE}$ ). The triangle is for the HT type estimate. The round is for the LCC estimate.

sistent with the finding of *Simulation 2*. The average computation time of 100 replications for the full sample MLE is 66.64 seconds. The average computation time for the LCC estimate (including the subsampling design implementation) is 19.32 seconds and for the HT estimate is 20.04 seconds.

#### Micro-blog data

In the second example we use the proposed method to analyze a data set of micro-blog, a Chinese version of twitter. We collect 625895 tweets posted on Sina micro-blog during January to March, 2018. For each tweet, the time of posting (in hour), the number of comments, retweets, and likes are recorded. Moreover, we also record the gender, the number of followers, fans, and the number of tweets posted before of the blogger who posted the tweet. We set the number of likes to be the response and the other variables to be the predictors. Since the number of likes is a counting variable, the log-linear model is used to do analysis. The response, ranging from 0 to tens of thousands, is very right skewed, i.e., imbalanced. The mean of response is 1038 while the median is 72. Around 55% of the tweets have the number of likes less than 100, while 2.11% of the tweets have the number of likes larger than 10000. In Figure 2, we show the histogram of the response.

In this analysis we focus on predicting the response by the predictors. Usually, the full data should be randomly separated into the training part for training model and the testing part for measuring prediction accuracy. To increase the stability of the results, here we use a ten-fold cross-validation style approach, that is, the full data is randomly split into ten parts and in each time, nine parts of the data serve as the training set and the rest one part serves as the testing set. Each time we train the log-linear model based on the training set and then calculate the root mean square error (RMSE) of the predicted values

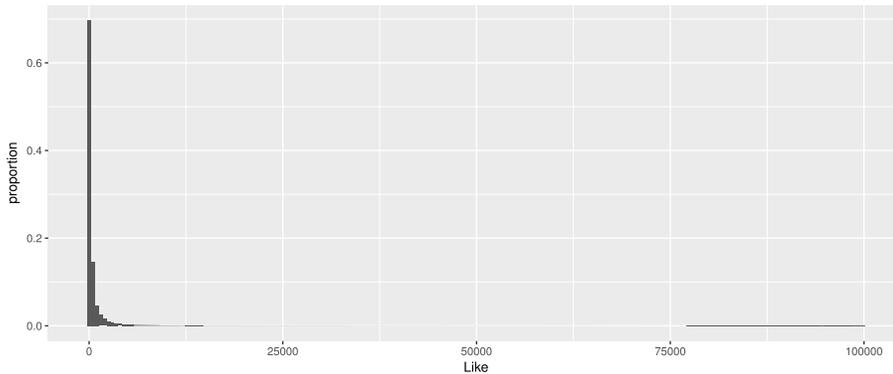


FIG 2. Histogram of the number of likes in Micro-blog data.

on the testing set, i.e.,  $\text{RMSE} = \sqrt{n_t^{-1} \sum_{i \in I_t} (y_i - \hat{y}_i)^2}$ , where  $I_t$  is the indicator set of the testing data,  $n_t$  is the size of  $I_t$ ,  $y_i$  is the response value, and  $\hat{y}_i$  is the predicted value. Finally, the average of the ten RMSEs, called ARMSE, is treated as the measure of the prediction accuracy. In the training set, we use the same sampling design as that in *Simulation 5* to get the subsample, i.e.,  $\pi_i = (c|y_i - \exp(\theta^\top z_i)|) \wedge 1$ , where  $y_i$  is the response,  $z_i = (1, x_i^\top)^\top$ ,  $x_i$  is the standardized predictor vector, and  $\theta$  is the pilot estimate which is the MLE fitted with a uniform subsample of size 10000. By taking a proper value of  $c$ , the size of the surprise subsample is set to be 10000. To make comparison, we also draw a uniform subsample of size 20000, which has a comparable sample size to the surprise subsample, to train the model by MLE. The prediction result of the full data MLE serves as the benchmark.

The ARMSE based on the full data is  $6.848 \times 10^3$ . The ARMSE based on surprise sampling data is  $7.194 \times 10^3$ , which is quite close to that of the full data. The ARMSE based on the uniform subsample data is, however,  $2.367 \times 10^{13}$ , which is much greater than that of the surprise sampling. One of the main reasons of such result is the right-skewness of the response data. With the subsampling percentage used here, the uniform subsampling has little chance to pick out the tweets with the extremely large responses, say, larger than 10000. Because of the significant influence of the large responses on the model fitting, the training models based on the uniform subsample data in some folds have quite different coefficients for some predictors from the full data models. Consequently, the prediction based on the uniform subsample models are very inaccurate for some responses in the testing set, resulting in extremely large RMSEs. By contrast, the surprise sampling is more likely to select those large responses, or in other words, the “surprising” responses, into the subsample due to its design motivation (the selection probability can be 1 for those extremely large responses compared with the pilot prediction). Consequently, the training models based on surprise sampling data take the large responses into consideration and even-

tually they have similar coefficients estimation and prediction performance with the full data models. The average computation time of 10 folds for the full data is 97.45 seconds. For the surprise sample data, the average computation time is 4.22 seconds and for the uniform sample data, is 3.54 seconds.

## 7. Conclusion and discussion

We develop an improved version of the LCC sampling, called surprise sampling design, to reduce computational burden and retain good prediction or estimation performance. The proposed sampling design is flexibly adaptive to different objectives. For a specific objective, the design has the corresponding optimality. The LCC sampling design, in this sense, can be viewed as a special case of the proposed surprise sampling design. For parameter estimation, we propose the HT type estimation approach. The resulting estimator is consistent and asymptotically normally distributed, without requiring the pilot estimator to be consistent or to be independent of the full data. For the binary response, if the logistic model is correctly specified and the pilot is consistent, the HT estimator is asymptotically as efficient as the estimator of [7] under the LCC sampling. The surprise sampling design and the HT estimation can be extended to more general responses such as counting and continuous response.

The proposed subsampling approach is more of a working principle for efficient data analysis which can be applied to various statistical learning problems. In the numerical studies, we mainly focus on regression. However, as we mentioned many times, the optimal sampling design, the algorithm described in Section 3, and the theories discussed in Section 4 can be readily extended to unsupervised learning tasks. The main reason is that our approach setup essentially only requires a specific loss function and a finite-dimensional parameter. One example of unsupervised learning is briefly given in Section 2. Some more detailed discussion on the approach and more meaningful applications of unsupervised learning are of great interest.

For the proposed approach, there are still limitations that lead to several directions worth further research. Firstly, although the sample size is much reduced for the subsample, the calculation of the subsampling probabilities may involve the full data. This can be still computationally intensive. Secondly, the pilot estimate in our approach is a guess of the target parameter, so they are of the same dimension. In some real applications, people may just want to use a relatively simpler model to get the pilot quickly and then apply a more complicated model to estimate the target parameter or do prediction. Then the pilot model used is different from the fitted model and the pilot is no longer a formal guess of the target parameter. Thus, the surprise sampling design needs adjustment and the optimality of the design requires more investigation. Thirdly, the HT type estimator defined in (2.1) is not semiparametrically efficient, but an immediate variant one is. The surprise sampling can be defined accordingly with the variant, in which the sampling probability  $\pi_i$  can be estimated by the kernel smoothed approach. It is possible to further improve the estimation

efficiency. Finally, when the dimension of the target parameter is high, one may introduce regularization penalties, such as lasso and ridge, into the objective function in (2.1). Based on the proposed framework, the regularization can be easily incorporated. The large sample properties of the resulting estimator need more exploration.

### Appendix A

In Appendix, we prove the propositions and theorems in Sections 3 and 4.

#### A.1. Proof of the propositions

To prove Propositions 1 and 2, we begin with the following lemma.

**Lemma 1.** *Given two random variables  $X$  and  $Y$  and constant  $r \in (0, 1)$ , consider the following optimization problem*

$$\min_Y \mathbb{E}(X^2/Y), \text{ s.t. } \mathbb{E}(Y) \leq r, 0 \leq Y \leq 1.$$

Suppose  $\mathbb{P}(X > 0) = 1$ . Then the solution is given by  $Y^* = cX \wedge 1$ , where  $c$  is such that  $\mathbb{E}(Y^*) = r$ .

*Proof.* For  $k > 0$ , let  $g(k) = \mathbb{E}(X \wedge k)/k = \mathbb{E}((X/k) \wedge 1)$ . Then we have that  $g(k)$  is continuous with respect to  $k$ ,  $\lim_{k \rightarrow \infty} g(k) = 0$  and  $\lim_{k \rightarrow 0} g(k) = 1$ . Hence there exists  $k$  such that  $g(k) = r \in (0, 1)$ . Let  $K$  be such that  $g(K) = r$ .

On one hand, we write  $\mathbb{E}(X^2/Y) = \mathbb{E}[(X \wedge K)^2/Y] + \mathbb{E}(Z/Y)$ , where  $Z = X^2 - (X \wedge K)^2$ . Then we have from the Cauchy-Schwarz inequality that

$$\mathbb{E} \left[ \frac{(X \wedge K)^2}{Y} \right] \geq \frac{[\mathbb{E}(X \wedge K)]^2}{\mathbb{E}(Y)} \geq \frac{[\mathbb{E}(X \wedge K)]^2}{r}$$

where the two equalities hold simultaneously if  $Y = Y^* := (rX/\mathbb{E}(X \wedge K)) \wedge 1$  with  $c = r/\mathbb{E}(X \wedge K)$ . Further we have  $\mathbb{E}(Z/Y) \geq \mathbb{E}(Z)$ . Then the whole objective is lower bounded for all  $Y$  as follows

$$\begin{aligned} \mathbb{E}(X^2/Y) &\geq \frac{[\mathbb{E}(X \wedge K)]^2}{r} + \mathbb{E}(Z) \\ &= \frac{[\mathbb{E}(X \wedge K)]^2}{r} + \frac{\mathbb{E}(X \wedge K)}{rK} [\mathbb{E}(X^2) - \mathbb{E}(X \wedge K)^2] \\ &= \frac{\mathbb{E}(X \wedge K)}{r} [\mathbb{E}(X \wedge K) + (\mathbb{E}(X^2) - \mathbb{E}(X \wedge K)^2)/K] \\ &= \frac{\mathbb{E}(X \wedge K)}{r} \left[ \mathbb{E}(X \mathbf{1}_{(X \leq K)}) + \mathbb{E} \left( \frac{X^2}{K} \mathbf{1}_{(X > K)} \right) \right]. \end{aligned} \tag{A.1}$$

On the other hand, by the definition of  $Y^*$ , we have

$$\mathbb{E}(X^2/Y^*) = \frac{\mathbb{E}(X \wedge K)}{r} \left[ \mathbb{E}(X \mathbf{1}_{(X \leq K)}) + \mathbb{E} \left( \frac{X^2}{K} \mathbf{1}_{(X > K)} \right) \right],$$

which turns out to achieve the lower bound in (A.1) and hence imply the desired result.  $\square$

*Proof of Proposition 1.* In Lemma 1, set  $X = \|A^{-1/2}g(D; \theta^*)\|$  and  $Y = \pi$ . Then it directly follows that the optimal  $\pi$  is given by  $\pi = (c\|A^{-1/2}g(D; \theta^*)\|) \wedge 1$  where  $c$  is such that  $\mathbf{E}(\pi) = r$ .  $\square$

*Proof of Proposition 2.* Note that the asymptotic variance of  $\sqrt{nv}^\top(\hat{\theta} - \theta^*)$  which we want to minimize is  $v^\top A^{-1}V_\pi A^{-1}v = \mathbf{E}(v^\top A^{-1}g(D; \theta^*)^{\otimes 2}A^{-1}v^\top/\pi) = \mathbf{E}(|v^\top A^{-1}g(D; \theta^*)|^2/\pi)$ . In Lemma 1, set  $X = |v^\top A^{-1}g(D; \theta^*)|$  and  $Y = \pi$ . Then it directly follows that the optimal  $\pi$  is given by  $\pi = (c|v^\top A^{-1}g(D; \theta^*)|) \wedge 1$  where  $c$  is such that  $\mathbf{E}(\pi) = r$ .  $\square$

*Proof of Proposition 3.* Let  $e_1$  be a  $p$ -dimensional unit vector whose first component is one and all the others are zero. We first prove that

$$v^\top A^{-1} = e_1^\top. \quad (\text{A.2})$$

When  $\theta = \theta_0$  we have that

$$\begin{aligned} \mathbf{E}(G(Y, X; \theta)|X) &= -\mathbf{E}\left(\frac{\partial^2}{\partial\theta\partial\theta^\top} \log \psi(Y, \theta^\top Z) \Big| Z\right) \\ &= \text{Cov}\left(\frac{\partial}{\partial\theta} \log \psi(Y, \theta^\top Z) \Big| Z\right) = \text{Cov}(S(Y, \theta^\top Z)Z \mid Z) \\ &= \mathbf{E}\left[(S(Y, \theta^\top Z)Z - \mathbf{E}(S(Y, \theta^\top Z)Z \mid Z))^{\otimes 2} \mid Z\right] \\ &= \mathbf{E}\left[(S(Y, \theta^\top Z) - \mathbf{E}(S(Y, \theta^\top Z) \mid Z))^2 \mid Z\right] ZZ^\top \\ &= \text{Var}[S(Y, \theta^\top Z)|Z] ZZ^\top, \end{aligned}$$

where  $\text{Cov}$  stands for the variance-covariance matrix of a random vector. Hence,  $A = \mathbf{E}[G(Y, X; \theta_0)] = \mathbf{E}[\text{Var}(S(Y, \theta_0^\top Z) \mid Z) ZZ^\top] = \mathbf{E}(\sigma_Z ZZ^\top)$ . Since the first component of  $Z$  is constant 1, we have that  $e_1^\top Z = 1$ . Also, note that  $\sigma_Z$  is a scalar. Thus,  $e_1^\top A = \mathbf{E}(\sigma_Z e_1^\top ZZ^\top) = \mathbf{E}(\sigma_Z Z^\top) = v^\top$ , which leads to (A.2).

Next, we have that

$$\begin{aligned} v^\top A^{-1}g(y, x; \theta_0) &= e_1^\top(-S(y, \theta_0^\top z)z) = -e_1^\top S(y, \theta_0^\top z)(1, x^\top)^\top \\ &= -e_1^\top(S(y, \theta_0^\top z), S(y, \theta_0^\top z)x^\top)^\top = -S(y, \theta_0^\top z). \end{aligned}$$

On one hand, by Proposition 2 we know that the asymptotic variance of  $v^\top \hat{\theta}$  based on the optimal  $\pi$  is

$$\mathbf{E}\left(\left(\frac{1}{c}|v^\top A^{-1}g(Y, X; \theta^*)|\right) \vee (v^\top A^{-1}g(Y, X; \theta^*))^2\right), \quad (\text{A.3})$$

which equals to  $\mathbf{E}((|S(Y, \theta_0^\top Z)|/c) \vee (S^2(Y, \theta_0^\top Z)))$  because of (A.2). On the other hand, the asymptotic variance of  $v^\top \hat{\theta}$  based on the  $\pi$  defined in the proposition is

$$\mathbf{E}\left(\frac{1}{\pi}|v^\top A^{-1}g(Y, X; \theta^*)|^2\right) \Big|_{\pi=c|S(Y, \theta_0^\top Z)| \wedge 1}$$

$$= \mathbb{E} \left( \left( \frac{1}{c} |S(Y, \theta_0^\top Z)| \right) \vee (S^2(Y, \theta_0^\top Z)) \right)$$

which is equal to (A.3). Therefore,  $\pi(Y, X; \theta_0) = c|S(Y, \theta_0^\top Z)| \wedge 1$  is the optimal  $\pi$  to minimize the asymptotic variance of  $v^\top \hat{\theta}$ .  $\square$

**A.2. Proof of the theorem and corollaries**

In order to make the presentation more concise, we use some abbreviation of the notation. Specifically, for  $i = 1, \dots, n$ , let  $l_i(\theta) = l(d_i; \theta)$ ,  $g_i(\theta) = g(d_i; \theta)$ , and  $G_i(\theta) = G(d_i; \theta)$ . Let  $R_n(\theta) = n^{-1} \sum_{i=1}^n \Delta_i l_i(\theta) / \pi_i$  and denote the  $\sigma$ -algebra generated by the observed data and the pilot estimate  $\tilde{\theta}$  by  $\mathcal{F}_n$ .

We now begin to prove Theorem 1. The following lemmas are needed.

**Lemma 2.** *Under conditions C1, C3, and C5  $\sup_{\theta \in \Theta} |R_n(\theta) - R(\theta)| \xrightarrow{P} 0$ .*

*Proof.* By the triangle inequality, we have that

$$\begin{aligned} & \sup_{\theta \in \Theta} |R_n(\theta) - R(\theta)| \tag{A.4} \\ & \leq \sup_{\theta \in \Theta} |R_n(\theta) - R_n^F(\theta)| + \sup_{\theta \in \Theta} |R_n^F(\theta) - R(\theta)|, \end{aligned}$$

where  $R_n^F(\theta) = n^{-1} \sum_{i=1}^n l_i(\theta)$ . By C3 and uniform law of large number, it can be shown that  $\sup_{\theta \in \Theta} \|R_n^F(\theta) - R(\theta)\| \xrightarrow{P} 0$ . Next we show the convergence of the first term in the right-hand-side of (A.4). For  $i = 1, \dots, n$ , define  $d_\delta(\Delta_i, \theta_1) = \sup_{\theta \in B(\theta_1, \delta)} \Delta_i l_i(\theta) / \pi_i - \inf_{\theta \in B(\theta_1, \delta)} \Delta_i l_i(\theta) / \pi_i$ , where  $B(\theta_1, \delta)$  is a ball in  $\mathbb{R}^p$  centered at  $\theta_1$  with radius  $\delta$ . By the continuity of  $l_i(\theta)$ , the measurability of  $l_i(\theta) / \pi_i$  to  $\mathcal{F}_n$ , and dominated convergence theorem,  $\mathbb{E}[d_\delta(\Delta_i, \theta_1) | \mathcal{F}_n] \rightarrow 0$  as  $\delta \rightarrow 0$ . Thus, for all  $\theta \in \Theta$  and  $\varepsilon > 0$ , there exists  $\delta_\varepsilon(\theta) > 0$  such that  $\mathbb{E}[d_{\delta_\varepsilon(\theta)}(\Delta_i, \theta_1) | \mathcal{F}_n] < \varepsilon$ . Since  $\Theta$  is compact, we can find a finite sequence of  $\theta_1, \dots, \theta_K$  such that  $\Theta$  is covered by  $\cup_{k=1}^K B(\theta_k, \delta_\varepsilon(\theta_k))$ . Thus, we have that

$$\begin{aligned} & \sup_{\theta \in \Theta} (R_n(\theta) - R_n^F(\theta)) \\ &= \sup_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \left[ \frac{\Delta_i}{\pi_i} l_i(\theta) - \mathbb{E} \left( \frac{\Delta_i}{\pi_i} l_i(\theta) | \mathcal{F}_n \right) \right] \\ &\leq \max_{1 \leq k \leq K} \sup_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{1}{n} \sum_{i=1}^n \left[ \frac{\Delta_i}{\pi_i} l_i(\theta) - \mathbb{E} \left( \frac{\Delta_i}{\pi_i} l_i(\theta) | \mathcal{F}_n \right) \right] \\ &\leq \max_{1 \leq k \leq K} \frac{1}{n} \sum_{i=1}^n \left[ \sup_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i}{\pi_i} l_i(\theta) - \mathbb{E} \left( \inf_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i}{\pi_i} l_i(\theta) | \mathcal{F}_n \right) \right]. \end{aligned}$$

From C5, we have that  $n^{-2} \sum_{i=1}^n (\sup_{\theta \in \Theta} l_i(\theta))^2 / \pi_i \xrightarrow{P} 0$ . Thus, by Chebyshev inequality, C3, and the weak law of large number, we can show that

$$\max_{1 \leq k \leq K} \frac{1}{n} \sum_{i=1}^n \left( \sup_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i}{\pi_i} l_i(\theta) - \mathbb{E} \left( \inf_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i}{\pi_i} l_i(\theta) | \mathcal{F}_n \right) \right)$$

$$\begin{aligned}
 &= \max_{1 \leq k \leq K} \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E} \left( \sup_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i l_i(\theta)}{\pi_i} \middle| \mathcal{F}_n \right) - \mathbb{E} \left( \inf_{\theta \in B(\theta_k, \delta_\varepsilon(\theta_k))} \frac{\Delta_i l_i(\theta)}{\pi_i} \middle| \mathcal{F}_n \right) \right] \\
 &\quad + a_n \\
 &= \max_{1 \leq k \leq K} \frac{1}{n} \sum_{i=1}^n \mathbb{E}[d_{\delta_\varepsilon}(\theta)(\Delta_i, \theta_k) | \mathcal{F}_n] + a_n \leq \varepsilon + a_n,
 \end{aligned}$$

where  $a_n$  satisfies that for all  $\varepsilon_1 > 0$ ,  $\mathbb{P}(|a_n| > \varepsilon_1 | \mathcal{F}_n) \xrightarrow{P} 0$ . Similarly, we can show that  $\inf_{\theta \in \Theta} (R_n(\theta) - R_n^F(\theta)) \geq \varepsilon + b_n$ , where for all  $\varepsilon_2 > 0$ ,  $\mathbb{P}(|b_n| > \varepsilon_2 | \mathcal{F}_n) \xrightarrow{P} 0$ . Thus, we can show that for arbitrary  $\varepsilon > 0$ ,  $\mathbb{P}(\sup_{\theta \in \Theta} |R_n(\theta) - R_n^F(\theta)| > \varepsilon | \mathcal{F}_n) \xrightarrow{P} 0$ . By Helly-Bray theorem,  $\sup_{\theta \in \Theta} |R_n(\theta) - R_n^F(\theta)| \xrightarrow{P} 0$ . The desired conclusion follows from (A.4).  $\square$

**Lemma 3.** Under condition C3 and C6, if  $\tilde{\theta} \xrightarrow{P} \theta^*$ , then

$$\frac{1}{n} \sum_{i=1}^n \frac{\Delta_i G_i(\theta^*)}{\pi_i} \xrightarrow{P} A.$$

*Proof.* Conditioning on  $\mathcal{F}_n$ ,  $\Delta_i G_i(\theta^*)/\pi_i$ ,  $i = 1, \dots, n$ , are independent with  $\mathbb{E}[\Delta_i G_i(\theta^*)/\pi_i | \mathcal{F}_n] = G_i(\theta^*)$  and  $\mathbb{E}[\|\Delta_i G_i(\theta^*)/\pi_i - G_i(\theta^*)\|^2 | \mathcal{F}_n] = (1/\pi_i - 1)\|G_i(\theta^*)\|^2$ . By C3, Chebyshev inequality, C6, and the weak law of large number, for any  $\varepsilon > 0$ ,  $\mathbb{P}(\|n^{-1} \sum_{i=1}^n \Delta_i G_i(\theta^*)/\pi_i - n^{-1} \sum_{i=1}^n G_i(\theta^*)\| > \varepsilon | \mathcal{F}_n) \xrightarrow{P} 0$ . Then, by Helly-Bray theorem, we have that for any  $\varepsilon > 0$ ,

$$\mathbb{P} \left( \left\| \frac{1}{n} \sum_{i=1}^n \frac{\Delta_i G_i(\theta^*)}{\pi_i} - \frac{1}{n} \sum_{i=1}^n G_i(\theta^*) \right\| > \varepsilon \right) \rightarrow 0,$$

that is

$$\left\| \frac{1}{n} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} G_i(\theta^*) - \frac{1}{n} \sum_{i=1}^n G_i(\theta^*) \right\| \xrightarrow{P} 0. \tag{A.5}$$

By law of large number,

$$\left\| \frac{1}{n} \sum_{i=1}^n G_i(\theta^*) - A \right\| \xrightarrow{P} 0. \tag{A.6}$$

Thus, by the triangle inequality, (A.5), and (A.6), we have

$$\frac{1}{n} \sum_{i=1}^n \frac{\Delta_i G_i(\theta^*)}{\pi_i} \xrightarrow{P} A. \tag{A.7} \quad \square$$

**Lemma 4.** Under condition C3, if  $\tilde{\theta} \xrightarrow{P} \theta^*$ , then  $\tilde{A} \xrightarrow{P} A$ .

*Proof.* Define  $A(\theta) = \mathbb{E}[G(D; \theta)]$ . By the triangle inequality, we have that

$$\|\tilde{A} - A\| \leq \|\tilde{A} - A(\tilde{\theta})\| + \|A(\tilde{\theta}) - A\|. \tag{A.7}$$

Since  $\tilde{\theta} \xrightarrow{P} \theta^*$ ,  $\|\tilde{A} - A(\tilde{\theta})\| \leq \sup_{\theta \in \mathcal{N}(\theta^*)} \|n^{-1} \sum_{i=1}^n G_i(\theta) - A(\theta)\|$  as  $n$  is sufficiently large, where  $\mathcal{N}(\theta^*)$  is a neighborhood in  $\theta$  about  $\theta^*$ . By  $C3$  and the uniform law of large number,  $\sup_{\theta \in \mathcal{N}(\theta_0)} \|n^{-1} \sum_{i=1}^n G_i(\theta) - A(\theta)\| \xrightarrow{P} 0$ , so does  $\|\tilde{A} - A(\tilde{\theta})\|$ . By the continuity of  $A(\theta)$  with respect to  $\theta$ ,  $\|A(\tilde{\theta}) - A\| \xrightarrow{P} 0$ . Thus, by (A.7),  $\tilde{A} \xrightarrow{P} A$ .  $\square$

*Proof of Theorem 1.* 1) Consistency: By  $C2$  and the uniform convergence of  $R_n(\theta)$  to  $R(\theta)$  on  $\Theta$  from Lemma 2, we can show that  $\hat{\theta} \xrightarrow{P} \theta^*$  by arguments similar to those in the proof of Theorem 5 in [7]. The details are omitted here. 2) Asymptotic normality: Since  $\sum_{i=1}^n \Delta_i g_i(\hat{\theta})/\pi_i = 0$ , by Taylor expansion, we have that

$$\begin{aligned} 0 &= \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} g_i(\theta^*) + \frac{1}{n} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} G_i(\theta^*) \cdot \sqrt{n} (\hat{\theta} - \theta^*) \\ &\quad + o_p \left( \frac{1}{n} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} G_i(\theta^*) \cdot \sqrt{n} (\hat{\theta} - \theta^*) \right). \end{aligned} \tag{A.8}$$

Write the first term in the right-hand-side of (A.8) as  $\eta_{1n} + \eta_{0n}$ , where

$$\eta_{1n} = \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\Delta_i - \pi_i}{\pi_i} g_i(\theta^*)$$

and  $\eta_{0n} = n^{-1/2} \sum_{i=1}^n g_i(\theta^*)$ . It is easy to see that conditioning on  $\mathcal{F}_n$ ,  $\eta_{1n}$  has zero mean and variance-covariance matrix equals to  $n^{-1} \sum_{i=1}^n (1/\pi_i - 1) g_i(\theta^*)^{\otimes 2}$ . Form  $C7$ , the variance-covariance matrix converges in probability to  $\mathbb{E}[(1/\pi - 1)g(D; \theta^*)^{\otimes 2}]$  which is independent of  $\mathcal{F}_n$ . Therefore,  $\eta_{1n}$  is asymptotically normal with mean zero and variance-covariance matrix  $\mathbb{E}[(1/\pi - 1)g(D; \theta^*)^{\otimes 2}]$ .

On the other hand, it is not difficult to see that  $\mathbb{E}(\eta_{1n} + \eta_{0n} | \mathcal{F}_n) = \eta_{0n}$  which is asymptotically normal with mean zero and variance-covariance matrix  $\mathbb{E}[g(D; \theta^*)^{\otimes 2}]$ . Consequently,  $\eta_{1n} + \eta_{0n}$  is asymptotically normal with mean zero and variance-covariance matrix  $\mathbb{E}[(1/\pi - 1)g(D; \theta^*)^{\otimes 2}] + \mathbb{E}[g(D; \theta^*)^{\otimes 2}] = V_\pi$ .

Combining the asymptotic normality of  $n^{-1/2} \sum_{i=1}^n \Delta_i g_i(\theta^*)/\pi_i$ , Lemma 3 and (A.8), we can derive that  $\sqrt{n}(\hat{\theta} - \theta^*) = O_p(1)$  and

$$\sqrt{n} (\hat{\theta} - \theta^*) = -A^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\Delta_i}{\pi_i} g_i(\theta^*) + o_p(1),$$

which is exactly (3.1). Using Slutsky theorem, it follows quickly that  $\sqrt{n}(\hat{\theta} - \theta^*)$  is asymptotically normal with mean zero and variance-covariance matrix  $A^{-1}V_\pi A^{-1}$ .  $\square$

Next we prove the corollaries.

*Proof of Corollary 1.* The result can be proved from Theorem 1 if one shows that C1-C4 and C8 imply C5-C7. We take  $\pi_i = (c\|\tilde{A}^{-1/2}g_i(\tilde{\theta})\|) \wedge 1$  (for prediction accuracy) as an example. The proof of the other  $\pi_i$  is of the same spirit. Note that  $\pi_i$  can be regarded as the function of  $\tilde{A}$ ,  $\tilde{\theta}$ , and the  $i$ th data point. We just write  $\pi_i$  as  $\pi(d_i; \tilde{A}, \tilde{\theta})$ , where  $\pi(d_i; \tilde{A}, \tilde{\theta}) = (c\|\tilde{A}^{-1/2}g_i(\tilde{\theta})\|) \wedge 1$ . Correspondingly, write  $\pi = (c\|A^{-1/2}g(D; \theta^*)\|) \wedge 1$  as  $\pi(D; A, \theta^*)$ . By C3, it is easy to see that  $E(\|g(D; \theta^*)\|^2/\pi) < \infty$ . Let  $V_{1n}(\mathcal{A}, \theta) = n^{-1} \sum_{i=1}^n (1/\pi(d_i; \mathcal{A}, \theta) - 1)g_i(\theta^*)^{\otimes 2}$ , and  $V_1(\mathcal{A}, \theta) = E[(1/\pi(D; \mathcal{A}, \theta) - 1)g(D; \theta^*)^{\otimes 2}]$ , where  $\mathcal{A}$  is a  $p \times p$  matrix. By the triangle inequality, we have that

$$\begin{aligned} & \|V_{1n}(\tilde{A}, \tilde{\theta}) - V_1(A, \theta^*)\| \\ \leq & \|V_{1n}(\tilde{A}, \tilde{\theta}) - V_1(\tilde{A}, \tilde{\theta})\| + \|V_1(\tilde{A}, \tilde{\theta}) - V_1(A, \theta^*)\|. \end{aligned} \tag{A.9}$$

Since  $\tilde{\theta} \xrightarrow{P} \theta^*$  and  $\tilde{A} \xrightarrow{P} A$  by Lemma 4, the first term in the right-hand-side of (A.9) is no larger than  $\sup_{(\mathcal{A}, \theta) \in \mathcal{N}(A, \theta^*)} \|V_{1n}(\mathcal{A}, \theta) - V_1(\mathcal{A}, \theta)\|$  as  $n$  is sufficiently large, where  $\mathcal{N}(A, \theta^*)$  is a neighborhood in  $(\mathcal{A}, \theta)$  about  $(A, \theta^*)$ . By C3 and the uniform law of large number,  $\sup_{(\mathcal{A}, \theta) \in \mathcal{N}(A, \theta^*)} \|V_{1n}(\mathcal{A}, \theta) - V_1(\mathcal{A}, \theta)\| \xrightarrow{P} 0$ . Moreover,  $\|V_1(\tilde{A}, \tilde{\theta}) - V_1(A, \theta^*)\| \xrightarrow{P} 0$  because of the continuity of  $V_1(\mathcal{A}, \theta)$  with respect to  $(\mathcal{A}, \theta)$ . Thus, we have  $V_{1n}(\tilde{A}, \tilde{\theta}) \xrightarrow{P} V_1(A, \theta^*)$ , which means C7 holds. By using similar arguments, we can show that  $n^{-1} \sum_{i=1}^n (\sup_{\theta \in \Theta} l_i(\theta))^2/\pi_i \xrightarrow{P} E(\sup_{\theta \in \Theta} \|l(D; \theta)\|^2/\pi)$  and  $n^{-1} \sum_{i=1}^n \|G_i(\theta^*)\|^2/\pi_i \xrightarrow{P} E(\|G(D; \theta^*)\|^2/\pi)$ . Then C8 implies that C5 and C6 hold.  $\square$

*Proof of Corollary 2.* When the covariates vector  $X$  does not concentrate on a hyperplane of dimension smaller than  $q$  and  $E(\|X\|^4/\pi) < \infty$ , conditions C3, C4, and C8 hold. When the logistic model is correctly specified, it can be showed that  $\Sigma_{\text{full}} = \{E[p(\theta_0^\top Z)(1 - p(\theta_0^\top Z))ZZ^\top]\}^{-1}$ . From Theorem 1, when  $\pi_i = |y_i - p(\tilde{\theta}^\top z_i)|$  and  $\tilde{\theta} \xrightarrow{P} \theta_0$ ,  $\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, A^{-1}V_\pi A^{-1})$ , where  $A = \Sigma_{\text{full}}$  and in  $V_\pi$ ,  $g(Y, X; \theta) = (Y - p(\theta^\top Z))Z$  and  $\pi = |Y - p(\theta_0^\top Z)|$ . Thus, to obtain the conclusion, it is efficient to show that  $V_\pi = 2\Sigma_{\text{full}}^{-1}$ . It is easy to see that

$$V_\pi = E \left[ \frac{g(Y, X; \theta_0)^{\otimes 2}}{\pi} \right] = E \left[ \frac{(Y - p(\theta_0^\top Z))^2}{|Y - p(\theta_0^\top Z)|} ZZ^\top \right] = E [|Y - p(\theta_0^\top Z)|ZZ^\top].$$

For a 0-1 binary random variable  $Y$  with  $E(Y) = p$ , it is easy to see that  $E(|Y - p|) = 2\text{Var}(Y) = 2p(1 - p)$ . Thus, we have that

$$\begin{aligned} & E [|Y - p(\theta_0^\top Z)|ZZ^\top] = E \{ E [|Y - p(\theta_0^\top Z)| | Z] ZZ^\top \} \\ = & 2E [p(\theta_0^\top Z)(1 - p(\theta_0^\top Z))ZZ^\top] = 2\Sigma_{\text{full}}^{-1}. \end{aligned}$$

The conclusion follows immediately.  $\square$

*Proof of Corollary 3.* We still take  $\pi_i = (c\|\tilde{A}^{-1/2}g(d_i; \tilde{\theta})\|) \wedge 1$  as an example.

1) Consistency: The proof of consistency of  $\hat{\theta}$  follows exactly the same step as that in Theorem 1, so we skip the details.

2) Asymptotic normality: The proof of the asymptotic normality of  $\hat{\theta}$  also follows the similar steps to those in Theorem 1. The main difference lies in that if  $\tilde{\theta} \xrightarrow{P} \bar{\theta}$ , by  $C6$ , we can show that  $\tilde{A} \xrightarrow{P} \bar{A}$  using arguments similar to those in the proof of Lemma 4. Meanwhile, it can also be shown similarly that  $V_{1n}(\tilde{A}, \tilde{\theta}) \xrightarrow{P} V_1(\bar{A}, \bar{\theta}) = V_{\bar{\pi}}$ . Thus,  $n^{-1/2} \sum_{i=1}^n \Delta_i g_i(\theta^*)/\pi_i$  is asymptotically normal with mean zero and variance-covariance matrix  $V_{\bar{\pi}}$ . Lemma 3 still holds here. It follows then that  $\sqrt{n}(\hat{\theta} - \theta^*)$  is asymptotically normal with mean zero and variance-covariance matrix  $A^{-1}V_{\bar{\pi}}A^{-1}$ .  $\square$

### Acknowledgments

The authors thank Professor Cheng Zhang and Pengfei Ma for providing the micro-blog data, and thank the referee for their constructive comments and suggestions.

### References

- [1] Ai, M., Yu, J., Zhang, H., & Wang, H. (2018). Optimal subsampling algorithms for big data generalized linear models. *arXiv preprint arXiv:1806.06761*.
- [2] Anderson, J. A. (1972). Separate sample logistic discrimination. *Biometrika*, 59(1), 19–35. [MR0345332](#)
- [3] Branco, P., Torgo, L., & Ribeiro, R. (2015). A survey of predictive modelling under imbalanced distributions. *arXiv preprint arXiv:1505.01658*.
- [4] Breslow, N. E. (1980). Statistical methods in cancer research. the analysis of case-control studies. *IARC Sci. Publ.* (pp. 32–35).
- [5] Chen, K. (2001). Generalized case-cohort sampling. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(4), 791–809. [MR1872067](#)
- [6] Chen, K. & Lo, S.-H. (1999). Case-cohort and case-control analysis with cox’s model. *Biometrika*, 86(4), 755–764. [MR1741975](#)
- [7] Fithian, W. & Hastie, T. (2014). Local case-control sampling: Efficient subsampling in imbalanced data sets. *Annals of Statistics*, 42(5), 1693. [MR3257627](#)
- [8] Horvitz, D. G. & Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260), 663–685. [MR0053460](#)
- [9] Huber, P. J. (2004). *Robust Statistics*, volume 523. John Wiley & Sons. [MR0606374](#)
- [10] Ma, P., Mahoney, M., & Yu, B. (2014). A statistical perspective on algorithmic leveraging. In *International Conference on Machine Learning* (pp. 91–99). PMLR. [MR3361306](#)
- [11] Mantel, N. & Haenszel, W. (1959). Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the National Cancer Institute*, 22(4), 719–748.

- [12] Miettinen, O. (1976). Estimability and estimation in case-referent studies. *American Journal of Epidemiology*, 103(2), 226–235.
- [13] Prentice, R. L. (1986). A case-cohort design for epidemiologic cohort studies and disease prevention trials. *Biometrika*, 73(1), 1–11.
- [14] Prentice, R. L. & Pyke, R. (1979). Logistic disease incidence models and case-control studies. *Biometrika*, 66(3), 403–411. [MR0556730](#)
- [15] Thomas, D. C. (1977). Addendum to: Methods of cohort analysis: Appraisal by application to asbestos mining. by fdk liddell, jc mcdonald and dc thomas. *Journal of the Royal Statistical Society, Series A*, 140, 469–491.
- [16] Wang, H., Yang, M., & Stufken, J. (2019). Information-based optimal sub-data selection for big data linear regression. *Journal of the American Statistical Association*, 114(525), 393–405. [MR3941263](#)
- [17] Wang, H., Zhu, R., & Ma, P. (2018). Optimal subsampling for large sample logistic regression. *Journal of the American Statistical Association*, 113(522), 829–844. [MR3832230](#)
- [18] Webb, S., Caverlee, J., & Pu, C. (2006). Introducing the webb spam corpus: Using email spam to identify web spam automatically. In *CEAS*.
- [19] Yao, Y., Yu, W., & Chen, K. (2017). End-point sampling. *Statistica Sinica* (pp. 415–435). [MR3618176](#)