

Achieving Both Valid and Secure Logistic Regression Analysis on Aggregated Data from Different Private Sources

Yuval Nardi[‡], Stephen E. Fienberg^{*}, and Robert J. Hall[†]

Abstract. Preserving the privacy of individual databases when carrying out statistical calculations has a relatively long history in statistics and had been the focus of much recent attention in machine learning. In this paper, we present a protocol for fitting a logistic regression when the data are held by separate parties—without actually combining information sources—by exploiting results from the literature on multi-party secure computation. Our protocol provides only the final result of the calculation compared with other methods that share intermediate values and thus present an opportunity for compromise of values in the individual databases. Our paper has two themes: (1) the development of a secure protocol for computing the logistic parameters, and a demonstration of its performances in practice, and (2) the presentation of an amended protocol that speeds up the computation of the logistic function. We illustrate the nature of the calculations and their accuracy using an extract of data from the Current Population Survey divided between two parties. Throughout, we build our protocol from existing cryptographic primitives, thus the novelty is in designing a concrete procedure for private computation of the logistic regression MLE rather than to propose new cryptographic constructions.

Keywords: Distributed analysis; Logistic regression; Privacy-preserving computation; Secure multiparty computation.

1 Introduction

Privacy concerns are becoming more acute, especially in the digitized world where computers with increasing processing capacities appear almost daily. These new machines together with impressive new technologies make the process of data collection, data storing, and data analysis as easy as ever. Untrusted elements may manipulate this “ease of use,” with the aim of deliberately causing harm by, for example, identifying and exposing sensitive data. The goal of privacy preserving methods is to prevent or at least lessen the chances of the occurrence of such harmful actions. In this paper we present a novel way to achieve this goal when a certain statistical analysis is required.

Preserving the privacy of individual databases when carrying out statistical calcula-

^{*}Faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology, Haifa, Israel, <mailto:ynardi@ie.technion.ac.il>

[†]Department of Statistics, Heinz College, Machine Learning Department, and Cylab, Carnegie Mellon University, Pittsburgh, PA, <mailto:fienberg@stat.cmu.edu>

[‡]Department of Statistics and Machine Learning, Carnegie Mellon University, Pittsburgh, PA, <mailto:rjhall+@cs.cmu.edu>

tions has a relatively long history in statistics and had been the focus of much attention in machine learning, e.g., see [1]. Once we merge data across sources, however, privacy becomes a far more complex matter and a number of privacy issues arise for the linked individual files that go well beyond those that associated with the data within individual sources. When the goal is the production of the results of some statistical calculation, such as a regression analysis, cf. Karr et al. [16, 17], we can often exploit results from the cryptography literature, borrowing tools such as secure multi-party computation, e.g., see [19, 24]. Secure multi-party protocols are concerned with distributed computation where each participating party, holding a private input, learns nothing but the result (see Section 3).

In this paper, we conceptualize the existence of a single combined database containing all of the information for the individuals whose data appear in the separate databases and for the union of the variables. We propose an approach that gives the results of a statistical calculation on this combined database without actually combining the data from the information sources, see [18, 19]. We focus mainly on the problem of logistic regression via maximum likelihood estimation, but our methods and tools are essentially adaptable to other statistical models and estimation methods, as we indicate in Section 9. Our approach provides only the final result of the calculation on the combined data compared with other methods that also share intermediate values and thus present an opportunity for compromise of values in the individual databases, cf. [7, 8].

We begin by developing a technique to perform the calculations required for fitting logistic regression models when the data are distributed among several parties. In our settings the parties are unwilling or are simply forbidden by laws or regulations to share their respective data. They acknowledge the fact that pooling their private data into a conceptually complete global database, and running the logistic regression on the pooled data rather than on their own data, will lead to a more accurate statistical analysis. We thus develop a secure protocol to compute the maximum likelihood estimates of the logistic regression model parameters. In principle, we could achieve the computation of this MLE (and in fact of any arithmetic function) securely by Yao's general protocol [25]. This protocol is in essence a theoretical construction, however, one which will often be impractical for large computations [24]. Therefore we focus on techniques which approximate the solution to some desired numerical accuracy, but which are more efficient and also conceptually more straightforward to implement than Yao's method. We demonstrate how most of the computation may be carried out by using standard techniques from the cryptographic literature, namely random shares and secure products. We focus on the semi-honest cryptographic model of security (see Section 3) since it is conceptually straightforward and typically permits computationally efficient protocols. However, in certain situations it may provide too weak a security guarantee, e.g., when certain agencies are willing to deviate from the protocol. Although we suspect that in the context of statistical agencies this security model is strong enough, we remark that in principle the techniques we give may be extended to handle malicious adversaries, although this is at the cost of increased computational burden.

We note that the fitting of logistic regression requires the computation of certain non-linear functions (e.g., the logistic function). Since such functions are not readily

	Protocol 1	Protocol 2
Parameters	L (number of steps in a step-function approximation of the logistic function).	k (number of steps of Eulers method used in updating the logistic values at each iteration).
Accuracy bound	$\ \beta - \hat{\beta}\ _2 \leq \frac{c_1 R}{L^\gamma \lambda_{\min}}$ (w.h.p), for some $\gamma < 1/2$. See the exact statement in Appendix 10, Theorem 1.	$\ \beta - \hat{\beta}\ _2 \leq \frac{c_2 R}{k \lambda_{\min}}$
Complexity	$O(nd^2 + d^3 \log d)$ multiplies, $O(nL)$ GT protocol	$O(nd^2 + d^3 \log d)$ multiplies.

Table 1: Summary of the two approaches. In the accuracy bound, $\hat{\beta}$ is the output parameter whereas β is what would be computed by an exact procedure, c_1, c_2 are constants, R is the radius of a ball containing the units, and λ_{\min} is the minimum eigenvalue of the Fisher information matrix.

available in the secure setting (without resorting to more expensive general purpose protocols) we first aim to approximate the logistic function by a sum of step functions, for which efficient secure protocols exist. We establish the theoretical validity of the secure protocol for computing the logistic parameters under this approximation, and show its performances in practice. In high dimensional problems with large numbers of units our protocol is more computationally burdensome. This is mainly because our approximation requires computing the predicate “greater-than,” which may take many encryptions. Indeed, evaluating this predicate by a reduction to Yao’s protocol takes roughly $O(b)$ encryptions where b is the number of bits used to represent the numbers. This calculation becomes dauntingly large in high dimensions due to the secret sharing scheme.

This leads us to investigate a second approximation to the logistic function, in which we view it as the solution to a certain ordinary differential equation, which may be integrated approximately via Euler’s method. Using this approximation we are able to perform the fitting process using only sums and products, and maintain the theoretical validity. The advantage to this is that these computations are very well studied primitives in secure multiparty computation and thus we can instantiate our method in a different secure multiparty computation scheme (e.g., [24, 10]), depending on the security demands of the data holders. We once again show that we can make the approximation arbitrarily accurate, at the expense of computational efficiency, and we present an illustrative empirical result. Accuracy and running times of the two protocols are summarized in Table 1. In both cases, computation and inversion of the Hessian matrix dominate the cost of the iterations.

Related Work

Papers focused on privacy-aware data-mining techniques have recently become popular, and the research in this area spans multiple models for privacy, some of which may be

considered orthogonal in terms of what they protect. Surveys of various branches of the resulting literature can be found in many sources, e.g., [24, 5, 19, 6].

Closest to this work are the investigations into secure computation of linear regression, first in the model of “weak security” in e.g., [16, 17], and also in the semi-honest cryptographic model [14]. We view this paper as an extension of the latter, designed to enable the handling of the non-linear logistic function, as well as the iterations needed for obtaining the solution via Newton-Raphson.

The previously published techniques, as well as those contained herein have the purpose of protecting privacy in the sense that the computation of the requisite parameter estimate leaks no information about the input data other than the estimates of parameter itself. In certain situations, however, the estimate of the parameter itself may be regarded as leaking information. For example, in a two-party setting, if one party held all of the data except for one unit (which belonged to the other party) then by comparing the output of the secure multi-party regression to the result of computing the regression on his own data in situ, he may be able to infer certain details about the unit held by the second party. Although we have not investigated the specific circumstances which would lead to such a privacy breach, it is often advisable to err on the side of caution and protect against such a leak. This is precisely what is done by Chaudhuri et al. [4], who provide a way to compute approximate logistic regression parameter estimates subject to constraints in a manner that prevents the success of such attacks, namely through the addition of random noise, which precludes identification of the units. Those authors consider a one-party setting in which the goal is for one agency to release logistic regression parameter estimates to the world (for other researchers to use). Nonetheless, their technique is amenable to the kind of secure computation we present below, and the union of the two techniques could result in a method in which both the computation itself, and the output are both privacy-preserving and the two different senses of the term.

Outline of Paper

We organize the remainder of the paper as follows: Section 2 gives a brief overview of the logistic regression and our setting, in order to set up notation. Section 3 presents the multi-party setup. In Section 4 we provide several sub-protocols which we will need. Sections 5 and 6 describe our protocol and an approach for speeding up the calculation involved, respectively. Section 7 describes implementation details. Section 8 illustrates aspects of the computation on an extract of data from the Current Population Survey divided between two parties. Section 9 discusses possible extensions. We defer all technical details to Appendices 10 and 10.

2 Logistic Regression

Logistic regression focuses on predicting binary outcomes or class membership given a set of explanatory variables or predictors. We can use the fitted model to predict class

membership for a newly obtained record consisting of only the values of the predictors. The basic framework of logistic regression treats binary responses y_1, \dots, y_n as realizations of n independent Bernoulli random variables, Y_1, \dots, Y_n , whose mean depends conditionally on a set of predictors $x_i \in \mathbb{R}^d$, as follows:

$$\mathbb{E}Y_i = \sigma(x_i^T \beta), \quad (1)$$

where $\sigma(a) = (1 + \exp(-a))^{-1}$ is the sigmoid (or the logistic) function, and β is a d -dimensional parameter vector. This makes the log odds, $\log(\mathbb{E}Y_i/(1 - \mathbb{E}Y_i))$, linear in the predictors.

A standard method for computing the maximum likelihood estimates of β is Newton-Raphson's method, since closed form expressions do not exist. The fitting process requires the user to supply the log-likelihood function associated with logistic regression, along with its first two derivatives. Suppressing dependence on the data and vector of parameters, we let ℓ be the log-likelihood, i.e., $\ell = \sum_i \{y_i x_i^T \beta - \log(1 + e^{x_i^T \beta})\}$. We also put on record the first two derivatives:

$$\nabla \ell = \sum_i \{x_i y_i - x_i \sigma(x_i^T \beta)\} \quad , \quad \nabla^2 \ell = - \sum_i \sigma(x_i^T \beta) (1 - \sigma(x_i^T \beta)) x_i x_i^T . \quad (2)$$

The gradient and the Hessian are assembled together to produce an estimate of the logistic parameters through the iterative process:

$$\beta_{t+1} = \beta_t - (\nabla^2 \ell)^{-1} \nabla \ell . \quad (3)$$

Our protocol will be structured in rounds, where each round corresponds to an iteration of Newton's method (3) followed by a convergence check. Each round involves a loop through all the cases x_i to compute the contribution to the gradient and Hessian. We keep intermediate values of β_t unshared between the parties. This is made possible by representing β_t as random shares (see Section 4).

Setting

We let X denote the $n \times d$ design matrix, and y the n -dimensional response vector. We assume the presence of $P \geq 2$ parties who are interested in computing logistic regression on the total of their data. We suppose that the union of the parties' data corresponds to the X and y of the logistic regression. In particular, we suppose that party j holds onto the pair (X_j, y_j) with $X_j \in \mathbb{R}^{n \times d}$ and $y_j \in \{0, 1\}^n$, where X_j is the j^{th} party design matrix, and y_j is her (binary) response vector.

In this work we consider a setting where each party has an "additive share" of the dataset. That is, $\sum_j X_j = X$ and $\sum_j y_j = y$ where X and y correspond to the design matrix and response vector of the combined data on which the logistic regression is performed. This subsumes all the partitioning schemes for the database (e.g., vertical and horizontal partitioning which are the cases considered in [24]) as in these cases for

each element, one party holds the value and the remaining parties hold zero. Furthermore this setup is applicable in a case where parties may have overlapping data, and the logistic regression is to be learned by using a linear function of the overlapping data (e.g., a weighted average) as a kind of measurement error model. We suppose that the union of the individual data sets gives the complete data. In cases where some data are missing, we can apply a privacy preserving imputation method such as in Jagannathan and Wright [15] as a preprocess, and then run our protocol.

We note that our method is general in the sense that it is applicable to every possible partitioning scheme, although it is clearly possible to treat specific cases such as horizontally or vertically partitioned data with more efficient specialized protocols.

3 Secure Multi-Party Computation

Ideally we would like our method to provide only the output of the calculation to the parties involved, and reveal nothing more. This is a lofty goal without the aid of trusted third parties, however it is relaxed in a useful way in the cryptographic literature. First, it is assumed that the parties are not able to quickly solve computationally hard problems (such as breaking RSA encryption). Then, a protocol is secure so long as intermediate values in the computation either contain almost no information (in the sense that the protocol would have to be re-run astronomically many times on the same input data in order to detect any information in the messages), or will only reveal information as the result of an intractable computation. We now briefly review the security model we intend to use.

We consider the “functionality” (see [10]) which maps the data of each party into the logistic regression parameter vector β :

$$\{(X_1, y_1), (X_2, y_2), \dots (X_P, y_P)\} \rightarrow \{\beta, \beta, \dots \beta\}. \quad (4)$$

The right hand side represents P copies of the parameter, so that each party receives the same output. Note that each design matrix is of the same dimensions.

A protocol for computing the functionality is just a sequence of steps consisting of parties performing local computations and sending intermediate messages to each other. In this work we build up a protocol for computing (4) which is secure in the presence of “semi-honest” parties. That is, parties who obey the protocol (and do not try to, e.g., inject malformed data) but keep a transcript of all the messages they receive. Intuitively, a protocol is secure in this setting whenever the intermediate messages give no information about the secret inputs of other parties. Formally, the “view” of the j^{th} party during the protocol is:

$$\text{view}_j((X_1, y_1), (X_2, y_2), \dots (X_P, y_P)) = \{(X_j, y_j), r, m_1, \dots m_{|m|}\}, \quad (5)$$

where r is a record of all the random draws made by party j , and m_k is the k^{th} message received by that party (we have dropped dependence of m on j for readability).

The protocol is secure so long as there exists a polynomial time algorithm which, when given only the input and output of party j , may output a random transcript of message which is *computationally indistinguishable* from view_j . See Goldreich [10] for a definition and discussion of computational indistinguishability. In essence, if the distribution of the sequence of messages depends only on the private input and output of party j , then we can simulate messages by drawing from this distribution (so long as the random number generator returns samples which are computationally indistinguishable from draws from the distribution). The existence of a simulator shows that intermediate messages do not depend on the private input of other parties, and so the protocol is secure in the sense that parties gain no more information about each other's private inputs than that revealed by the output of the protocol.

An example of a protocol which does not achieve this definition of security is one where all parties send their data to party 1, who computes the parameter locally on the combined data and then sends it back to all other parties. In this case the messages received by party 1 consist of the data of other parties; in general it is impossible to simulate these messages given only the input and output belonging to party 1.

In the next section, we present a protocol for performing Newton's method on the logistic regression objective in a way that is secure in the presence of semi-honest parties. Our protocol makes use of a specially designed approximation for the logistic function. Section 6 then describes a different approximation necessitating the operations of only sums and products, and thus speeding-up the computations.

Although we propose to use the cryptographic model for security, others exist and deserve a place in the theory of privacy preserving data analysis. The main alternatives we see are "weak" security, and *perturbation* of the data. The former comprises a body of literature summarized in Vaidya et al. [24]. The idea is that by giving weaker privacy guarantees, we can implement much more efficient protocols. Whether it is acceptable to have this weaker privacy guarantee is a question which one must consider on a case-by-case basis. Although we describe our protocol in terms of the cryptographic model, by replacing the primitive operations (in Section 4) with their weakly-secure counterparts, we convert our protocol into a weakly secure (but also computationally more efficient) one.

The second alternative is data perturbation or sanitization. The idea would be for each party to somehow perturb his data until he is happy to release it to the other parties (e.g., through the addition of random noise). Thereupon the parties would each have a noisy copy of all the data, and could locally compute whatever statistical method they wanted on the union of the data. The difficulty with this approach is that to protect privacy may require the addition of noise of such amplitude as to render the data itself useless.

4 Primitives for Secure Protocols

In this section we lay out some primitives and sub-protocols which we will combine to make a full logistic regression protocol. Details of the implementation of these primitives are in the references cited. See also [14].

4.1 Secret Sharing

In our construction we make extensive use of additive secret sharing. The idea is to divide a quantity of interest a into P random numbers a_j (one for each party) so that $\sum_j a_j = a$. If the a_j are distributed uniformly in the field then any subset of the a_j will reveal nothing about a . In fact, the sum over any subset is a random variable, the distribution of which does not depend on the secret value.

We use this construction to keep all intermediate quantities secret during the evaluation of Newton’s method (i.e., the gradient, Hessian, and intermediate parameter vectors). As long as we can construct sub-protocols which compute random shares of a quantity, from random shares of inputs, then we can compose these sub-protocols together to finally obtain random shares of the logistic regression estimate. With these in hand the parties can then exchange shares and reveal the vector itself.

Although the joint distribution of the a_j concentrates on the linear subspace corresponding to the secret value, marginally the shares are uniformly distributed and do not depend on any parameters. Hence we can easily simulate messages based on these shares since the marginal distributions are known, and we achieve security as defined in Section 3. Although this approach is intuitively appealing, when dealing with real values (as is the case in most logistic regression problems) we face two problems: the first is that we must restrict to some finite domain in order for uniform distributions to exist, and second, most cryptographic sub-protocols we may use to compute—e.g., products only operate on rings of integers modulo some large prime. Therefore, we propose to use the method of [14] in order to approximate the same computations in modular arithmetic on $\mathbb{Z}_b = \{0, 1, \dots, b-1\}$ for some large b . We use a “2s complement” approach to represent negative numbers, and then a division by a constant to represent real numbers to some fixed precision. The mapping from \mathbb{Z}_b to the fixed precision real numbers is:

$$f : \mathbb{Z}_b \rightarrow \mathbb{R}, f(a) = M^{-1} \begin{cases} a & a \leq \frac{b}{2} \\ a - b & a > \frac{b}{2}. \end{cases} \quad (6)$$

In this way, we associate each element of \mathbb{Z}_b with an element in \mathbb{R} . The constant M determines the balance between the range of values which may be represented, and the precision of the fractional quantities which may be represented. A higher value for M yields numbers with greater precision but with a smaller range.

Thus our protocol begins by appropriately rounding the data values so that they fit into this representation (this represents a negligible loss in accuracy since M may be

made large at little expense to the computation which is to follow), and then carrying out summations, products etc., using sub-protocols which operate on these shares which are integers from \mathbb{Z}_b . Finally, when the protocol ends, the shares of the final output are summed by the parties to reveal the output, which is regarded as a real number.

4.2 Computing Sums and Products with Random Shares

To implement Newton's method we must essentially perform linear algebraic operations on random shares, for example by computing shares of the Newton step from shares of the gradient and inverse Hessian. In this section we describe how to obtain random shares of sums and products of quantities that are themselves represented as shares. Using these constructions, we compute inner and outer products of vectors of random shares, and hence also matrix multiplies.

Computing shares of the sum of two secret quantities $a = \sum_j a_j$ and $b = \sum_j b_j$ is direct, as it involves only the local computation $a_j + b_j$ for each party $j = 1, \dots, P$. That is, party j simply adds his shares a_j and b_j together (in the ring \mathbb{Z}_b) to get a random share of the quantity $a + b$.

Obtaining random shares of the product of two secret quantities is more involved. Evidently it requires interaction between the parties (for computing "cross terms"). Thus we propose to use a sub-protocol which computes the functionality

$$\{(a_1, b_1), \dots, (a_P, b_P)\} \rightarrow \{c_1, \dots, c_P\},$$

in which each c_i is (marginally) distributed uniformly at random in \mathbb{Z}_b and

$$\sum_{i=1}^P c_i = \sum_{i=1}^P \sum_{j=1}^P a_i b_j,$$

in which the sums and products are taken in the ring \mathbb{Z}_b . In the setting of the semi-honest model, such a protocol may be constructed by making use of Paillier's encryption scheme (see e.g., [14]). In essence the parties each encrypt their shares using the same public key, passing the encryptions to each other whereupon the encryption of each $a_i b_j$ may be constructed via the homomorphic property. Finally these may be summed, resulting in an encryption of the product, from which the random shares may be generated.

However, note that we may not just naively apply this product protocol when the numbers constitute fixed point approximations under the scheme outlined above, since multiplication of two such numbers results in a stray factor of M (e.g., $f(ab) = Mf(a)f(b)$). Unfortunately this problem is not trivial to solve. Namely, we may not have each party divide his share by M (or multiply by the inverse M^{-1} in \mathbb{Z}_b), since this may result in a multiple of $M^{-1}b$ being added to the shared variable. A method to solve this problem is given in [14] which requires a number of encryptions which is only a small constant factor greater than required in the protocol sketched above. This protocol generates random shares of the product even if the original shares weren't themselves random, e.g., if they were due to the partitioning of the data.

We also note that dividing one secret value into another securely is much more difficult than dealing with products and requires more elaborate (and computationally demanding) protocols. Below we show how matrix inversion can be performed without any divisions.

4.3 Evaluating Interval Membership

We suppose we are able to evaluate the following predicate in a secure way:

$$1\{a \geq b\},$$

where a, b are secret values held by separate parties. This is known as Yao’s “millionaires problem,” since he described it in the context of determining which millionaire has the most money, without disclosing actual bank balances.

An example of a protocol which computes this predicate and is immediately amenable to our fixed point scheme is given by Blake and Kolesnikov [2]. We can also trivially extend it so that each party receives a random share of the output bit (i.e., each party receives a random bit, the “xor” of which yields the correct output bit). Using this technique we can also check whether a secret value (i.e., a sum of random shares) is greater or less than some constant:

$$1\{a_1 + a_2 \geq c\} = 1\{a_1 \geq c - a_2\}, \quad (7)$$

where a_1, a_2 are the random shares of a held by two parties.

4.4 Securely Inverting a Matrix

We use a matrix inversion routine built up entirely of matrix multiplications and subtractions, thus allowing us to use the constructions of the preceding sections to implement it securely. We obtain the reciprocal of a number a without necessitating any actual division by an application of Newton’s method to the function $f(x) = x^{-1} - a$. Iterations follow $x_{s+1} = x_s(2 - ax_s)$, which requires multiplication and subtraction only.

It turns out that we can apply the same scheme to matrix inversion, e.g., see [12] and references therein. A numerically stable, coupled iteration for computing A^{-1} takes the form:

$$\begin{aligned} X_{s+1} &= 2X_s - X_s M_s, & X_0 &= c^{-1}I, \\ M_{s+1} &= 2M_s - M_s^2, & M_0 &= c^{-1}A, \end{aligned} \quad (8)$$

where $M_s = X_s A$, and c is to be chosen by the user. A possible choice, leading to a quadratic convergence of $X_s \rightarrow A^{-1}$ ($M_s \rightarrow I$), is $c = \max_i \lambda_i(A)$. In our actual implementation we used instead the trace (which dominates the largest eigenvalue, as the matrix in question is positive definite), since we can compute shares of the trace from shares of the matrix locally by each party. To compute c^{-1} we use the same iteration, with scalars instead of matrices. For this iteration we initialize with an arbitrarily small

$\epsilon > 0$ (as convergence depends on the magnitude of the initial value being lower than that of the inverse we compute). We use the constructions of Section 4.2 to iterate through (8) until convergence. As $M_s \rightarrow I$, we check for convergence by considering the absolute difference between the trace of M_s and the data dimension d , and we can evaluate the function $1\{|\text{tr}(M_s) - d| > \epsilon\}$ on random shares of the trace of M_s using the same form as (7).

5 First Protocol for Logistic Regression

We recall the usual Newton-Raphson iteration expression (3). To perform the iteration we first compute random shares of the update direction: $\Delta_t = -(\nabla^2 \ell(\beta_t))^{-1} \nabla \ell(\beta_t)$, via the formulation of matrix-vector products of random shares. We can then add these random shares to the current parameters β_t to obtain random shares of β_{t+1} . To check convergence recall (from e.g., [3]) we should end if:

$$\lambda^2 = (\nabla \ell(\beta_t))^T \Delta_t \leq \epsilon. \quad (9)$$

We can compute (9) securely using the same form as (7). The result is sharable among all the parties, and the protocol ends whenever the result is 0— i.e., when λ^2 is not greater than ϵ .

By using the constructions of the previous section, we have the tools required to invert shares of the Hessian, and thus to compute the Newton step. All that we need to do is construct a secure protocol to evaluate the logistic (sigmoid) function. In principle, a specialized sub-protocol could be built up using the construction of Yao [25]. The method would be to construct circuit that evaluates the sigmoid function in the same manner that the arithmetic logic unit in a CPU would. Then we could give this circuit the secure treatment and make it into a protocol following Goldreich [10]. The disadvantage with this approach is that the circuit evaluation protocols are prohibitively expensive and thus they are not useful in practice except for trivial circuits, see e.g., Malkhi et al. [20]. Instead we use a specially crafted approximation to the logistic function in terms of indicator functions. We describe this next.

5.1 A Secure Approximation to the Logistic Function

The logistic function itself is the cumulative distribution function (CDF) of the logistic distribution. We propose to approximate this function with an “empirical CDF.” This is a function of a set of L samples z_l , taken independently from a logistic distribution:

$$\sigma(a) \approx F_L(a) = L^{-1} \sum_{l=1}^L 1\{a \geq z_l\}. \quad (10)$$

Based on the Glivenko-Cantelli theorem, and later work by Dvoretzky, Kiefer, and Wolfowitz, the rate at which the empirical CDF converges to the true CDF (i.e., the

logistic function which is of interest) is known (see e.g., [21]). Using these results, we obtain bounds on the maximum difference between the logistic function and our approximation, which hold with high probability. See the remark below in Section 5.2 about the accuracy of the approximation.

We now turn attention to obtaining random shares of the logistic function evaluated at random shares of $\beta^T x_i$. We obtain random shares of $\beta^T x_i$ by using the inner product construction for multiplying together random shares. If we denote shares of this inner product by $(\beta^T x_i)_j$ for party j , we write:

$$\sigma(\beta^T x_i) \approx L^{-1} \sum_l 1\{\beta^T x_i \geq z_l\} = L^{-1} \sum_l 1\{(\beta^T x_i)_1 + (\beta^T x_i)_2 \geq z_l\}. \quad (11)$$

Thus the problem reduces to getting random shares of the sum of indicators. Note that we can re-write each indicator function as:

$$1\{(\beta^T x_i)_1 + (\beta^T x_i)_2 \geq z_l\} = 1\{(\beta^T x_i)_1 \geq z_l - (\beta^T x_i)_2\}. \quad (12)$$

If party 2 generates the logistic random variables then we have a trivial reduction to (7). In order to restrict the view of either party to a random share, we restrict the output to random bits o_1^l , and o_2^l , such that

$$o_1^l \oplus o_2^l = \begin{cases} 1 & \text{if } 1\{a \geq z_l\} \\ 0 & \text{otherwise} \end{cases},$$

where \oplus is the exclusive or. The right-hand side of equation (11) requires random shares of the fraction of outputs with $o_1^l \oplus o_2^l = 1$. We can establish this by noticing that

$$\sum_{l=1}^L (o_1^l \oplus o_2^l) = \sum_{l=1}^L o_1^l + \sum_{l=1}^L o_2^l - 2o_1^T o_2,$$

where we denote $o_k = (o_k^1, \dots, o_k^L)$ for $k = 1, 2$. Jagannathan and Wright [15] use this method to convert xor shares into additive shares for a different privacy-preserving task.

In order for the output to behave this way, we can either use Yao's protocol directly, or take a more efficient GT protocol and modify it to give a (xor) random share. Here we use the protocol of Blake and Kolesnikov [2]. Having computed random shares of the logistic function, we can then use the constructions of Section 4.2 to compute random shares of the gradient and Hessian, and hence build a full logistic regression estimation protocol.

5.2 Quality of the Logistic Approximation

The error in the approximation (10) propagates into the error of the resulting logistic parameter estimator. This may be quantified by noticing the following inequality which relates the two errors:

$$\|\hat{\beta} - \beta\|_2 \leq \frac{R[L^{-1} + \|\sigma(\cdot) - F_L(\cdot)\|_\infty]}{\hat{\lambda}_{\min}}. \quad (13)$$

Here $\hat{\beta}$ is the optimizer of the exact log likelihood, and β is the optimizer of our approximation; $\hat{\lambda}_{\min}$ is the smallest eigenvalue of the Fisher information matrix $I(\beta) = -n^{-1}\nabla^2\ell(\beta)$ (on some interval, see Appendix 10), and R is the radius of a ball which contains all the data vectors (i.e., $\forall i, \|x_i\|_2 \leq R$). The proof of this inequality follows lemma 1 of [4] in which the two convex functions are the exact log likelihood objective, and the difference between the exact and approximate objectives.

What remains is to manage the maximum absolute error of the approximation (10). The tail behavior of this quantity is given, for every $\epsilon > 0$, by:

$$P\left(\left\|\sigma(\cdot) - L^{-1} \sum_{l=1}^L 1\{\cdot \geq z_l\}\right\|_{\infty} > \epsilon\right) \leq 2e^{-2L\epsilon^2}, \quad (14)$$

known as the Dvoretzky-Kiefer-Wolfowitz (DKW) inequality. We can use this inequality to bound the numerator of expression (13). Doing so, we may bring the parameter output by our protocol, and that output by the exact (non-private) algorithm as close as we want (except on a set of negligible probability) by increasing the parameter L . Our main result (see Theorem 1 in Appendix 10) shows that with probability tending to one exponentially fast, the following inequality holds:

$$\|\beta - \hat{\beta}\|_2 \leq \frac{c_1 R}{L^\gamma \hat{\lambda}_{\min}},$$

where $0 < \gamma < 1/2$ and c_1 is some positive constant. We refer the reader to Appendix 10 for detailed theoretical derivation.

In Section 8 we perform an experiment to show how well the method performs with reasonably small L . Note that for Newton's method to converge in this approximation, we must use the same sample of L logistic random variables each time we approximate the sigmoid. Otherwise assessing convergence would be difficult as the objective function would be constantly shifting. We propose that the parties draw L logistic variables ahead of time, and use these for all the computations.

5.3 Hessian Lower Bound Technique

Notice that the Newton Raphson method requires inverting a matrix (the Hessian of the log likelihood) at each iteration. In our setting, using our iterative inversion method, this becomes very expensive. Therefore we propose to use a well-studied approximation [22], which replaces the iteration by:

$$\beta_{t+1} = \beta_t - 4(X^T X)^{-1} \nabla \ell. \quad (15)$$

First note that under this technique the algorithm only ever needs a single matrix inversion, since $X^T X$ is constant throughout all the iterations. Second, this algorithm still eventually converges to the correct parameter value (modulo the other approximations we make in our protocol). The reason is that the inverse Hessian is always greater

than $4(X^T X)^{-1}$, in the sense that the difference is positive semi-definite, see e.g, Minka [22] for more details. What's more, this technique ensures that progress towards the optimum is monotonic, and so assessing convergence may be simpler.

5.4 Computation and Communication Complexity

First we count how many times we must run each of our primitives for each iteration of Newton's method. The approximation of Section 5.1 requires nL instances of the GT protocol per round, as L instances are required per case. Computing the gradient and the Hessian requires $n(1 + d + d^2)$ multiplications. Inverting the Hessian takes $2d^3$ multiples and one GT per iteration of (8). Since this inner iteration is quadratically convergent, it takes $O(\log d)$ iterations to converge, and thus takes $O(d^3 \log d)$ multiples and $O(\log d)$ instances of GT. In total then, each outer iteration takes $O(nd^2 + d^3 \log d)$ multiples, and $nL + O(\log d)$ invocations of the GT protocol.

Each multiplication requires a number of encryptions and decryptions; this scales quadratically with the number of parties P since they must exchange with one another. Thus the computational workload increases as the data are split into more pieces. Note that although repeated use of the cryptosystem is quite expensive, performance on normal hardware is relatively rapid. A machine dedicated to the computation and running multiple threads can do thousands of encryptions per second.

Each instance of GT using the protocol of [2] requires $O(\log B)$ encryptions and decryptions (and operations on encrypted values etc.). Therefore in total our approximation of Section 5.1 requires $O(nL \log B)$ encryptions per iteration. This may be too computationally demanding for large L . One way to reduce this cost is to run the scheme using a coarse approximation to the sigmoid (i.e., a small L) to convergence, then increase L , resample the logistic variables, and then continue Newton's method from the previous convergent parameter. Although the latter iterations will still be computationally burdensome, there will be fewer of them. Another way is to use a different approximation to the sigmoid function. This is outlined next in Section 6.

Note that the total amount of communication by all parties is also proportional to the number of multiples and GT invocations. For an invocation of either, a party must transmit $\log N$ bits to another party, and then receive a message of the same length. There are a total of $O(P^2(nd^2 + d^3 \log d) + nL)$ messages which must be sent for each iteration. If the number of parties or cases, or the granularity of the approximation is large, running the protocol over a high speed local area network would make the communication overhead manageable.

6 Second Protocol for Logistic Regression

As we mentioned above, the computation complexity of evaluating approximation (10) to the logistic function scales linearly with L , since on each of Newton's iteration we invoke Yao's protocol to compute the GT predicate, and we do it for every case i . This

may be prohibitively expensive even for a moderate L . A possible way to reduce this computational burden was briefly described in Section 5.4. Here, we provide full details of a more structured approach, which is reminiscent of Euler's method. The approach is built (again) on computing Newton's iteration (3). It would be more natural in this section to treat the logistic function in a *vectorized* fashion, i.e., $\sigma(a) = (\sigma(a_1), \dots, \sigma(a_n))$, for an n -dimensional vector $a = (a_1, \dots, a_n)$. Therefore, we use different, albeit equivalent, representations for the gradient and Hessian:

$$\nabla \ell = X^T \{y - \sigma(X\beta)\} \quad , \quad \nabla^2 \ell = -X^T \text{diag}\{\sigma(X\beta) \circ (1 - \sigma(X\beta))\} X \quad . \quad (16)$$

Here X is the design matrix whose rows are x_i^T , the units or feature vectors (see (2)). The symbol “ \circ ” denotes the element-wise product, i.e., $u \circ v = v \circ u = \text{diag}(u)v$.

We modify the iteration so that we neither explicitly compute the logistic function $\sigma(\cdot)$ which is involved in both the gradient and Hessian, nor use the approximation in expression (10). Note that throughout the procedure we may treat each unit x_i as having an associated logistic function value $\sigma(\beta_i^T x_i)$. We propose to track a vector of approximate function values $\hat{\sigma}_t \approx \sigma(X\beta_t)$ which will be updated after each iteration. Then, these approximate values will be used to compute the next iteration of β_t . Note that the derivative of the logistic function is given by:

$$\sigma'(a) = \sigma(a)(1 - \sigma(a)) \stackrel{\text{def}}{=} g(\sigma(a)) \quad . \quad (17)$$

Therefore, knowing the value $\sigma(a)$, we can determine the derivative of the logistic function around a by a single multiplication. Linearizing around some value a_0 gives:

$$\sigma(a) = \sigma(a_0) + (a - a_0)g(\sigma(a_0)) + 2^{-1}(a - a_0)^2 \sigma''(\cdot)|_{a^*} \approx \sigma(a_0) + (a - a_0)g(\sigma(a_0)) \quad , \quad (18)$$

where the second derivative is evaluated at some value a^* in the interval between a and a_0 . Denote by $\Delta_t = \beta_{t+1} - \beta_t$ as in Section 5, then make use of the approximation:

$$\hat{\sigma}_{t+1} = \hat{\sigma}_t + (X\Delta_t) \circ g(\hat{\sigma}_t) \quad , \quad (19)$$

where g is applied element-wise to $\hat{\sigma}_t$.

Over the course of the entire algorithm, the approximation $\hat{\sigma}_t$ is updated repeatedly, in a manner very similar to using Euler's method to numerically integrate the differential equation (17). It is well known that the error of this method decreases with the size of the “step” taken at each iteration. In the above, the steps are of size $X\Delta_t$, which will in general be different on each iteration, and will also be different for each unit. In order to control the error we amend this approximation by breaking down the step into k smaller steps each of size $k^{-1}X\Delta_t$, and performing k such updates. As we shall see, we may base our choice of k on some aspect of the design matrix, X , in order to reach a desired level of error in the approximation. We write this approximation as:

$$\hat{\sigma}_{t+1} = \hat{\sigma}_t + k^{-1}X\Delta_t \circ \sum_{i=1}^k g(\hat{\sigma}_i^*) \stackrel{\text{def}}{=} \hat{\sigma}_t + X\Delta_t \circ \tilde{g}_k(\hat{\sigma}_t, X\Delta_t) \quad , \quad (20)$$

where the $\hat{\sigma}_i^*$ are the intermediate values corresponding to the inner iterations, and we define \tilde{g}_k as the function which gives the average value of g evaluated on these values.

We summarize our method in the following coupled iteration:

$$\begin{aligned}
\beta_0 &= 0^{d \times 1} \\
\hat{\sigma}_0 &= 2^{-1} \cdot 1^{n \times 1} \\
\Delta_t &= 4(X^T X)^{-1} X^T (y - \hat{\sigma}_t) \\
\beta_{t+1} &= \beta_t + \Delta_t \\
\hat{\sigma}_{t+1} &= \hat{\sigma}_t + X \Delta_t \circ \tilde{g}_k(\hat{\sigma}_t, X \Delta_t),
\end{aligned} \tag{21}$$

where $0^{d \times 1}$ is the d -dimensional vector of zeros and $1^{n \times 1}$ is the n -dimensional vector of ones. The proposed iteration differs from the protocol of Section 5 (and from the usual Newton-Raphson method). The main difference is that we have replaced the logistic function approximation (10) with our Taylor approximation. Note that we are using again the bound on the Hessian (see Section 5.3), which would make computation easier. We use this technique in our method for this reason, and also since it interacts well with our Taylor approximation by ensuring that convergence towards the optimum is in a sense monotonic, as shown in Section 0.1. In keeping with our goal of using only sums and products, we recall that it is possible to invert a matrix with just these operations (see Section 4.4).

We now present a bound on the distance from our approximated regression coefficients β_t , to the true optimizer of the log-likelihood which we denote by $\hat{\beta}$, as in (13). Since our iterations are guaranteed to converge (see Section 0.1), we can choose to run the iterations until $\|X^T(y - \sigma_t)\|_2$ is smaller than some threshold b (i.e., by choosing t accordingly):

$$b \geq \|X^T(y - \hat{\sigma}_t)\|_2 \geq \|X^T(y - \sigma_t)\|_2 - \|X^T(\hat{\sigma}_t - \sigma_t)\|_2.$$

Therefore we can bound the norm of the gradient of the logistic log-likelihood taken at our final parameter estimate:

$$\|\nabla \ell(\beta_t)\|_2^2 \leq b + nRc\tau,$$

where R is the radius of a ball containing all the data vectors, exactly as in (13), c is some constant, and τ is a quantity upper bounding the maximal Euler's step size.

We can use this to construct our main result about the quality of our approximation. Suppose we choose $b \leq nRc\tau$, then from the above we have:

$$\|\beta_t - \hat{\beta}\|_2 \leq \frac{2Rc\tau}{\hat{\lambda}_{\min}}, \tag{22}$$

where $\hat{\lambda}_{\min}$ is the smallest eigenvalue of the Fisher information matrix $I(\cdot) = -n^{-1} \nabla^2 \ell(\cdot)$ in the line segment between β and $\hat{\beta}$. Note that $\hat{\lambda}_{\min} = n^{-1} \lambda_{\min}$ and the factors of n cancel.

Therefore we can make the accuracy of our approximation arbitrarily good by decreasing τ , although, as we shall see there is a tradeoff involved. A smaller τ usually means a higher k , resulting in increased computational demands. We refer the reader to Appendix 10 for complete technical details.

6.1 Choice of k

Thus far we have that the error of the approximation decreases as τ is decreased; however, this last variable is not controlled directly (as L was in protocol 1) but rather is a function of k , the number of steps taken for each outer iteration of the algorithm.

In principle, to get at a prescribed step size τ , we can choose k by noting that:

$$\begin{aligned} \tau = \max_t \|k^{-1}X\Delta_t\|_\infty &\leq k^{-1} \max_{v \in [-1/2, 1/2]^n} \|X(X^T X)^{-1}X^T v\|_\infty \\ &\leq \frac{d}{2k} \max_{v \in [-1, 1]^n, \|u\|_2=1} \|uu^T v\|_\infty = \frac{cd}{k}, \end{aligned} \quad (23)$$

where c is some universal constant. An alternative choice is to run the protocol with a small value of k , e.g., 10, and then to re-run with different values to assess the sensitivity of the computation. In Section 8 we show that this technique performs well with small k .

6.2 Computational Complexity

We can measure the overall complexity of our method in terms of the number of products that are needed, since these are the most time-consuming operations we use. First note that to construct the matrix $X^T X$ takes nd^2 products, and inversion of this matrix takes $O(d^3)$ using (8), where the constant is related to the condition of the matrix. Then on each iteration, to compute Δ_t takes $nd + d^2$ products. Our approximation to the logistic function takes nk products, for a total of $n(k + d) + d^2$ products per iteration.

We compare this with the cost of a protocol which computes the logistic function via a specially designed sub protocol based on circuit evaluation, cf. Yao [25]. If the latter may be evaluated using q encryptions, then the complexity would be $n(d + q) + d^2$ operations per iteration. As mentioned before, this number would typically be much larger than k (for example on the order of the number of bits used to represent the numbers). Therefore on each iteration we can save a multiple of n operations, which may be especially important when n is large.

7 Security Guarantees

Since our protocol runs until convergence, the number of rounds is variable and depends on the data itself. Furthermore a matrix inversion was performed by an iterative scheme

which itself took some variable number of iterations to converge. Therefore we amend the protocol so that the output for each party is a triple consisting of the convergent parameter value and the number of iterations it took to converge, and the number of iterations taken for the matrix inversion. This way the messages received from testing convergence are easily simulated (i.e., a zero on every round up until the number specified in the output, then a one on that iteration) and this clearly reveals no more information since the parties know “where they are” in the protocol at all times and could count these numbers of iterations. We remark that in principle the numbers of iterations reveal some additional information which is not present in the regression coefficients alone; namely they reveal information about the spectrum of eigenvalues of the matrix $X^T X$. Convergence may be much faster in the presence of similar eigenvalues, and slower when the ratio of the largest to smallest eigenvalue is large. Evidently the values themselves are not available, but it may be possible to conclude the approximate order of magnitude of this ratio. Whether this is extremely problematic depends on the circumstances. If it is, then a possibility is to add to the matrix some multiple of the identity matrix. This will alter the eigenspectrum so that this ratio is changed, but while preserving the theoretical properties of both procedures (since the inverse will remain a lower bound to the Hessian). Having dealt with this technicality we will consider simulating the other intermediate messages in our simulator, and consider these convergence tests already taken care of.

In both of our protocols, the messages which are transmitted are always part of some sub-protocol, namely multiplication or evaluation of the “greater than” predicate. The only exception to this is the final messages which are sent immediately before the output is reconstructed. As those messages are themselves random shares they may be simulated easily (although they must be simulated in such a way that they sum to the correct output values, but this is trivial). The messages which are passed during the sub-protocols may be simulated based on their respective input and outputs so long as the sub-protocols are cryptographically secure. Since we take care to ensure that the intermediate values are random shares, the simulators for the sub-protocols “compose” to form a simulator for the main protocol (see [10]).

8 Illustrative Experiments

We provide two illustrative experiments to demonstrate our approach. The first aims at showing the performance of our protocol from Section 5. Specifically, we examine the effect of approximation (10) on the resulting parameter values when small and large number of logistic variables L are being used. The second example takes a look at the altered protocol from Section 6, which uses the coupled iteration (21) instead of approximation (10), and reports its performances for different values of k , the number of Euler’s “steps”.

For both experiments we use an extract from the Current Population Survey (CPS) data (see <http://www.bls.gov/cps/>), which includes data on a sample of slightly more than 50,000 U.S. households. We focus on predicting whether household income is

greater than 50,000 dollars. We converted M -category features into $M - 1$ binary features, and divided age into 4 bins corresponding to 20 year intervals. Note that although we expressed our approach in terms of continuous covariates, it handles binary flags just as well, where said covariates take on e.g., 0.0 and 1.0.

Our protocol from Section 5 deviates from the exact computation in two ways, first we use an approximation to the gradient, and second we perform all the calculations in fixed-point arithmetic. Both of these approximations can be made arbitrarily tight but at the expense of computational efficiency. To demonstrate that our protocol can be implemented in an efficient manner and produce reasonably accurate results, we implemented it in a simulator and compared the results to exact logistic regression on the CPS data.

For each of $L = 100$ and $L = 500$ we ran our first protocol 100 times. The table below shows the means and standard deviations of the resulting parameter values. Evidently, as L gets bigger, the accuracy of the parameter values improve. Figures 1 and 2 show how the likelihood of the estimate maintained by the protocol increases with the number of iterations. We computed the error bars by removing the 5 samples that deviated from the mean by the greatest amount, and plotting the minimum and maximum from the remaining ones. This then corresponds to an approximate 95% confidence interval, and would become an exact interval if we were to perform more and more simulations. For the purposes of comparison, we also plotted the likelihood achieved by the exact non-private Newton Raphson algorithm, and a non-private algorithm which we referred to as “Hessian lower bound.” Both give upper bounds for what we hope to achieve, the latter is an algorithm where we just use the approximation of (15), and exact (i.e., non-private) logistic sigmoid values. We see that as L increases, the first protocol more closely approximates the Hessian lower bound technique, which converges more slowly than the exact Newton Raphson method.

For the second experiment, we ran our coupled iteration on the CPS data with $k = 5, 10$. Although each iteration of our algorithm may be cheap, all is for nought if we require many more iterations for convergence. To determine whether this happens we compared our method to the Hessian lower bound method of (15), since this represents our algorithm without the approximation. In Figure 3, we plot the likelihoods of the second protocol against the iteration number. Since there is no randomness in the second approximation, there are no error bars. Even for small values of k , much smaller than those suggested by (23), the approximation to the Hessian lower bound technique is quite good, and increasing k further (e.g., to 50) results in curves which are exactly the same as that of the Hessian lower bound method. In Table 2 we show the resulting parameter estimates for both methods.

9 Beyond Logistic Regression

We can use the construction of Section 5 to build secure protocols for similar statistical calculations, e.g., the constructions for computing shares of outer products and matrix inverses naturally yield a secure algorithm for performing linear regression, for details

	NR	P1, $L = 100$	s.d.	P1, $L = 500$	s.d.	P2, $k = 5$	P2, $k = 10$
Intercept	-10.7536	-11.6306	1.0761	-11.5732	0.5339	-10.7944	-11.2262
Child Sup	0.0002	0.0002	0	0.0002	0	0.0002	0.0002
Property Tax	0.0003	0.0003	0	0.0003	0	0.0003	0.0003
Num in Household	0.9802	0.9916	0.0863	0.9881	0.0434	0.9259	0.9601
Num Children	-1.056	-1.0721	0.0935	-1.0685	0.047	-1.0017	-1.0384
Num Married	0.0342	0.0343	0.0053	0.0343	0.0022	0.032	0.0333
Child Sup Ind.	-0.0001	-0.0001	0	-0.0001	0	-0.0001	-0.0001
Education	0.3218	0.3276	0.0295	0.3265	0.0146	0.3058	0.3172
Age	-4.6178	-3.9701	0.3451	-3.94	0.1638	-3.6202	-3.8011
	-4.2368	-3.6782	0.3148	-3.6596	0.1504	-3.4817	-3.5823
	-3.9608	-3.3355	0.2852	-3.3157	0.135	-3.1592	-3.2481
	-4.9575	-4.4069	0.3795	-4.3814	0.1829	-4.1096	-4.2624
Marital Status	0.0064	0.0051	0.0282	0.0001	0.012	-0.0035	-0.0007
	-0.5362	-0.5623	0.058	-0.5562	0.0279	-0.5205	-0.5404
	-0.4378	-0.4718	0.0819	-0.4609	0.0374	-0.3934	-0.4321
	-0.5855	-0.6096	0.0675	-0.6018	0.0309	-0.572	-0.5889
	-0.9617	-1.0283	0.1146	-1.0116	0.0549	-0.957	-0.9874
-0.7496	-0.7888	0.0852	-0.7783	0.0401	-0.736	-0.7598	
Race	0.2417	-0.2588	0.0276	-0.2565	0.0135	-0.2401	-0.2491
	-0.4981	-0.5167	0.05	-0.5128	0.0244	-0.4835	-0.4997
	-0.1658	-0.1796	0.0196	-0.1785	0.0102	-0.1631	-0.1719
Sex	-0.2763	-0.2802	0.0244	-0.2792	0.0125	-0.2613	-0.2712

Table 2: Estimates produced by the exact method (Newton Raphson), and the two protocols, for different parameter settings of the protocols.

see [14]. Furthermore we can add the “ridge regression” penalty on the weights (i.e., computing a MAP estimate under a Gaussian prior) to the protocol in a natural way for both linear and logistic regression. It is also possible to implement the coordinate ascent computation of the lasso (or sparse logistic regression) using these constructions and the GT protocol to perform soft thresholding.

Our protocol generalizes to the class of Generalized Linear Models (GLMs) with other link functions, thus going beyond linear and logistic regression. The GLM approach consists largely of a random component Y_i from an exponential family, a systematic component with a linear predictor $\eta_i = x_i^T \beta$, and a link function $\eta_i = h(\mu_i)$, where $\mu_i = \mathbb{E}Y_i$. If h makes the linear predictor $\eta_i = \theta_i$, where θ_i is the natural parameter of the exponential family, h is canonical.

For Poisson log-linear models with the canonical link, $\mu_i = \exp\{\eta_i\}$, we can approximate the exponential function similarly. For Gamma models with the canonical link, $\mu_i = 1/\eta_i$, and for inverse-Gaussian models with the canonical link, $1/\mu^2$, we can use the number inverting without division scheme. We can also extend our approach to treat binary regression with non-canonical links, such as the *probit* link function, or more generally, inverse CDF link functions. The general form of the gradient is:

$$\nabla \ell = \sum_i \frac{\{y_i x_i - x_i \mu_i\}}{\text{Var}(Y_i)} \frac{\partial \mu_i}{\partial \eta_i}. \quad (24)$$

Let F denote a given CDF ($F = \Phi$ leads to the probit link function, while, of course, $F = F_L$ leads to the logit link function). Then, $\mu_i = F(\eta_i)$, and thus $\partial \mu_i / \partial \eta_i = f(\eta_i)$, where f is the density. Therefore, we should find approximations for f as well as for F (approximation for F will follow the same idea as for F_L , i.e., using the empirical CDF).

10 Conclusion

We have demonstrated that a fully secure approach to logistic regression based on the cryptographic notion of security may be made practical for use on moderately large datasets shared between several parties. Although our protocols and approach are slower than methods with weaker security guarantees, they offer more rigorous guarantees with respect to the privacy of the input data. We emphasize that our protocol (like any cryptographic protocol) prevents leakage of information which may arise from the computation itself. It does not address any leakage which results from the output.

The problem of secure regression is far from solved however, as we have yet to deal with the problem of secure record linkage, and have implicitly assumed that the parties know how their respective datasets are aligned. Furthermore, record linkage using a statistical model may be incorrect and result in biased and more highly variable estimates of model parameters. For further details see [13].

Appendix A- Theoretical Validity of the First Protocol

Here we show how a bound on the error in the approximation (10) to the logistic function leads to a bound on the quality of the convergent parameter vector output by the protocol. Specifically, we establish the validity of (13). Let R denote a constant such that $\|x_i\|_2 \leq R$, for $i = 1, \dots, n$. Recall the expressions for the gradient $\nabla\ell$ and Hessian $\nabla^2\ell$ given in (2). Define the approximated gradient, by substituting F_L for σ :

$$\nabla\tilde{\ell}(\beta) = \sum_{i=1}^n x_i y_i - x_i F_L(x_i^T \beta). \quad (25)$$

Rewriting $\nabla\ell(\beta) = \nabla\tilde{\ell}(\beta) + \sum_{i=1}^n x_i F_L(x_i^T \beta) - x_i \sigma(x_i^T \beta)$, and applying the triangle inequality we obtain a bound on the norm of the gradient of the logistic objective:

$$\|\nabla\ell(\beta)\|_2 \leq \|\nabla\tilde{\ell}(\beta)\|_2 + nR \|F_L(\cdot) - \sigma(\cdot)\|_\infty. \quad (26)$$

Next we convert a bound in the norm of the gradient into a bound on the distance to the optimum.

Lemma 0.1. *Let $\hat{\beta}$ be the optimizer of the logistic regression objective, and let λ_{\min} denote the smallest eigenvalue of the negative Hessian in the line segment between β and $\hat{\beta}$. Then:*

$$\|\beta - \hat{\beta}\|_2 \leq \frac{\|\nabla\ell(\beta)\|_2}{\lambda_{\min}}. \quad (27)$$

Proof. We use the mean-value theorem (for vector-valued functions) to write the difference between gradient vectors at β and $\hat{\beta}$:

$$\nabla\ell(\beta) - \nabla\ell(\hat{\beta}) = \nabla\ell(\beta) - 0 = \left(\int_0^1 \nabla^2\ell(a\beta + (1-a)\hat{\beta}) da \right) (\beta - \hat{\beta}). \quad (28)$$

Now, for every (symmetric) matrix B , and a non-zero vector e , the Rayleigh quotient satisfies $e^T B e / e^T e \geq \lambda_{\min}(B)$, where $\lambda_{\min}(B)$ is the minimal eigenvalue of B . If $B = A^2$, for a positive definite (symmetric) matrix A , this reduces (after taking the square root on both sides) to $\|Ae\|_2 / \|e\|_2 \geq \lambda_{\min}(A)$. Applying this to (28), and using Weyl's inequality, we have:

$$\|\nabla\ell(\beta)\|_2 = \left\| \left(\int_0^1 \nabla^2\ell(a\beta + (1-a)\hat{\beta}) da \right) (\beta - \hat{\beta}) \right\|_2 \geq \lambda_{\min} \|\beta - \hat{\beta}\|_2. \quad (29)$$

This completes the proof. \square

Lemma 0.2. *Using the same notation we have:*

$$\min_{\beta \in \mathfrak{B}} \|\nabla\tilde{\ell}(\beta)\|_2 \leq nRL^{-1}, \quad (30)$$

where \mathfrak{B} is a (non-empty) set of logistic parameters defined in the proof.

Proof. Consider a continuous, monotonically non-decreasing function $g(\cdot)$ which satisfies $\|g(\cdot) - F_L(\cdot)\|_\infty \leq L^{-1}$. Such a function clearly exists, for example the smooth nondecreasing curve which goes through all points $(z_{(j)}, jL^{-1})$ where $1 \leq j \leq L$ (where $z_{(j)}$ is j^{th} smallest logistic variable used in F_L). Since $g(\cdot)$ is nondecreasing, it is the derivative of some convex function:

$$G(a) = \int_{-\infty}^a g(b) db. \quad (31)$$

Consider the approximation to the logistic gradient which uses g instead of F_L :

$$\nabla \bar{\ell}(\beta) = \sum_{i=1}^n x_i y_i - x_i g(x_i^T \beta). \quad (32)$$

This is the derivative of a concave function:

$$\bar{\ell}(\beta) = \sum_{i=1}^n x_i^T \beta y_i - G(x_i^T \beta), \quad (33)$$

which is indeed concave since it is a linear function minus a convex function. Hence $\bar{\ell}$ has a unique maximum somewhere. Consider the functions $g(\cdot)$ so that the maximum is in the interior of the space \mathbb{R}^d (i.e., is not at infinity). Hence for each such g we have a point $\bar{\beta} \in \mathbb{R}^d$ where the gradient is zero, i.e., $\nabla \bar{\ell}(\bar{\beta}) = 0$. Denote the set of such $\bar{\beta}$ by \mathfrak{B} , and note that \mathfrak{B} is not empty. An argument similar to the one that led to (26) shows that:

$$\|\nabla \tilde{\ell}(\beta)\|_2 = \|\nabla \bar{\ell}(\beta) + \sum_{i=1}^n x_i g(x_i^T \beta) - x_i F_L(x_i^T \beta)\|_2 \leq \|\nabla \bar{\ell}(\beta)\|_2 + nRL^{-1}. \quad (34)$$

Therefore:

$$\|\nabla \tilde{\ell}(\bar{\beta})\|_2 \leq \|\nabla \bar{\ell}(\bar{\beta})\|_2 + nRL^{-1} = nRL^{-1}, \quad (35)$$

which completes the proof. \square

We now put this all together and state the main result about our approximation F_L .

Lemma 0.3. *If our approximation $\nabla \tilde{\ell}$ is used as an approximation to the gradient of the logistic log likelihood, and numerical optimization is performed until $\|\nabla \tilde{\ell}(\beta)\|_2 \leq nRL^{-1}$, then:*

$$\|\beta - \hat{\beta}\|_2 \leq \frac{R(L^{-1} + \|F_L(\cdot) - \sigma(\cdot)\|_\infty)}{\hat{\lambda}_{\min}}, \quad (36)$$

where $\hat{\beta}$ is the optimizer of the exact logistic regression objective, β is the result of our numerical optimization, R is the radius of a ball containing all the x_i , and $\hat{\lambda}_{\min}$ is the smallest eigenvalue of the Fisher information matrix $I(\cdot) = -n^{-1}\nabla^2 \ell(\cdot)$ in the line segment between β and $\hat{\beta}$.

Proof. Notice that $\|\nabla\tilde{\ell}(\beta)\|_2 \leq nRL^{-1}$ is guaranteed in light of Lemma 0.2. The proof follows by substituting (26) into (27), and by noticing that $\hat{\lambda}_{\min} = n^{-1}\lambda_{\min}$ and the factors of n cancel. \square

Theorem 1. *Assume the conditions and notation of the previous Lemmas. Let $0 < \gamma < 1/2$. Then, with probability at least $1 - 2e^{-cL^{1-2\gamma}}$,*

$$\|\beta - \hat{\beta}\|_2 \leq \frac{c_1 R}{L^\gamma \hat{\lambda}_{\min}},$$

where c, c_1 are positive constants (with c depending on c_1).

Proof. The proof is straightforward by using (36) and the DKW inequality (see (14)):

$$\mathbb{P}\left(\|\beta - \hat{\beta}\|_2 > \frac{c_1 R}{L^\gamma \hat{\lambda}_{\min}}\right) \leq \mathbb{P}\left(\|F_L(\cdot) - \sigma(\cdot)\|_\infty > \frac{c_1}{L^\gamma} - \frac{1}{L}\right) \leq 2e^{-cL^{1-2\gamma}}.$$

\square

Appendix B- Theoretical Validity of the Coupled Iteration

Here we establish the convergence of the coupled iteration (21), and the error in our Taylor approximation of the logistic function.

0.1 Monotonicity and Convergence

We show that the update described in (21) converges monotonically towards some final value β . We relate the size of the step taken at one iteration to the size of the step in the previous iteration. We aim to show that first, these steps are always in the same directions for each unit, and secondly, the steps are monotonically decreasing and eventually the iterations converge.

Lemma 0.4. *$X\Delta_{t+1}$ element-wise has the same sign as $X\Delta_t$, in the sense that $X\Delta_{t+1} \circ X\Delta_t \geq 0$.*

Proof. If we define the idempotent matrix $M = X(X^T X)^{-1}X^T$, then we write:

$$\begin{aligned} X\Delta_{t+1} &= 4X(X^T X)^{-1}X(y - \hat{\sigma}_t) \\ &= 4M(y - \hat{\sigma}_t) \\ &= 4M[y - \hat{\sigma}_{t-1} - (X\Delta_t) \circ \tilde{g}_k(\hat{\sigma}_{t-1})] \\ &= 4MM(y - \hat{\sigma}_{t-1}) - 16M \text{diag}(\tilde{g}_k(\hat{\sigma}_{t-1}))M(y - \hat{\sigma}_{t-1}) \\ &= 4M \text{diag}(1 - 4\tilde{g}_k(\hat{\sigma}_{t-1}))M(y - \hat{\sigma}_{t-1}) \\ &= M \text{diag}(1 - 4\tilde{g}_k(\hat{\sigma}_{t-1}))X\Delta_t, \end{aligned} \tag{37}$$

where we made use of the idempotency of M . Next considering the element-wise product as the diagonal of the outer product of these two matrices,

$$X\Delta_{t+1}(X\Delta_t)^T = M \operatorname{diag}(1 - 4\tilde{g}_k(\hat{\sigma}_{t-1}))X\Delta_t\Delta_t^T X^T.$$

Since we clearly have that $1 - 4\tilde{g}_k(\hat{\sigma}_{t-1}) > 0$ no matter what value $\hat{\sigma}_{t-1}$ takes (due to the definition of g_k), we have that this matrix is the product of positive semi-definite matrices, and therefore is itself positive semi-definite. Therefore the diagonal elements are all non-negative, and we have proved the claim. \square

This result allows us to analyze our approximation to the logistic function as though we were using the forwards Euler method to integrate the differential equation (17), since all the steps for any particular unit will be in the same direction.

Lemma 0.5. *As long as each step $k^{-1}|X\Delta_t| \leq \tau < 1$ (where the inequality is element-wise), then $0 < \hat{\sigma}_t < 1$, $\forall t$ (i.e., the approximate logistic values will remain between 0 and 1).*

Proof. Suppose that the step is positive for all units and $\hat{\sigma}_t < 1$, then:

$$\hat{\sigma}_{t+1} - \hat{\sigma}_t \leq \tau \hat{\sigma}_t (1 - \hat{\sigma}_t^2) < 1 - \hat{\sigma}_t,$$

so we also have that $\hat{\sigma}_{t+1} < 1$. Likewise for units which are involved in a negative step, if they are greater than 0, then they remain so into the next iteration by an argument which is symmetric to the one above. Therefore we have that our logistic values never leave the interval $(0, 1)$. \square

With this we also have that $0 < 4\tilde{g}_k(\hat{\sigma}_t) < 1$ for all t , from the definition of g and \tilde{g}_k . Substitution into (37), yields that:

$$\|X\Delta_{t+1}\|_2 \leq \|M\|_2 \|\operatorname{diag}(1 - 4\tilde{g}_k(\hat{\sigma}_{t-1}))\|_2 \|X\Delta_t\|_2 < \|X\Delta_t\|_2, \quad (38)$$

since M has eigenvalues which are each either 0 or 1. This shows that the magnitude of the steps for the individual units is shrinking towards zero. Therefore we conclude that eventually, our approximations of the logistic values stop updating. If we assume that X has d linearly independent columns, then this also implies that Δ_t is going towards zero, and therefore our algorithm eventually converges.

0.2 Quality of the Logistic Approximation

We now analyze the error in the approximation of the logistic function values. We then use this together with the convexity of the problem to yield a bound on the error in the convergent parameters (see (22)). To aid the notation, in this section we consider the problem of estimating the logistic values for just a single case, and specifically one for which the steps are all positive. Due to the symmetry of the logistic function about 0,

we will then have the same type of bounds on the error when the approximation updates in the negative direction. We first show a loose upper bound on the supremum of the error which would be encountered if the approximation was run for an infinite number of steps of size at most τ , and then use this to bound the error after finitely many such steps.

As we have shown by the above monotonicity argument, our approximation to the logistic function is essentially analogous to using Euler's method to integrate the derivative of the logistic function. Since we consider approximating a single value, we change the names of our variables to avoid confusion with the previous vector valued approximation. If we denote by \hat{s}_t the approximated value after t steps of various sizes, $\tau_0 \dots \tau_{t-1} < \tau$. Thus $\hat{s}_t \approx s_t = \sigma(a_t)$ where $a_t = \sum_{i=0}^{t-1} \tau_i$. We compare this approximation to the exact values and consider the error:

$$\xi_t = \hat{s}_t - s_t .$$

Making use of the step (18), we evaluate the error in the next iteration:

$$\begin{aligned} \xi_{t+1} &= \hat{s}_{t+1} - s_{t+1} \\ &= \hat{s}_t + \tau_t g(\hat{s}_t) - s_t - \tau_t g(s_t) - 2^{-1} \tau_t \sigma''(\cdot) \Big|_{a_t^*} \\ &= \xi_t + \tau_t [g(\hat{s}_t) - g(s_t)] + \zeta_t \\ &= \xi_t + \tau_t (\hat{s}_t - s_t) g'(\cdot) \Big|_{s_t^*} + \zeta_t \\ &= \xi_t (1 + \tau_t g'(\cdot) \Big|_{s_t^*}) + \zeta_t \\ &= \xi_t (1 + \tau_t - 2\tau_t s_t^*) + \zeta_t \end{aligned}$$

where we have defined

$$\zeta_t = -2^{-1} \tau_t \sigma''(\cdot) \Big|_{a_t^*} \quad (39)$$

and $a_t \leq a_t^* \leq a_{t+1}$ is some value in the interval about which the second derivative is taken. Likewise s_t^* is bounded between s_t and \hat{s}_t . As we have seen from (38), as long as $\tau_t \leq \tau < 1$ then $0 < \hat{s}_t < 1$ for all t . Since we only consider positive steps $\tau_t > 0$ then we have that $2^{-1} \leq \hat{s}_t < 1$, and hence the same bound applies to s_t^* . Therefore we have that:

$$|\xi_{t+1}| \leq |\xi_t| + |\zeta_t| .$$

Therefore we see that:

$$\sup_t |\xi_t| \leq \sum_{i=1}^{\infty} |\zeta_i| . \quad (40)$$

Examining the form of $\sigma''(\cdot)$, we find it to be a function which is everywhere negative. Examining the third derivative, we find that the second derivative has exactly one stationary point in $[0, \infty)$ which is located at:

$$a^* = -\log \frac{6 - \sqrt{12}}{6 + \sqrt{12}}, \quad \sigma(a^*) = \frac{6 + \sqrt{12}}{12} .$$

Whats more, we see that $\partial^3\sigma(\cdot) < 0$ on $[0, a^*)$, and $\partial^3\sigma(\cdot) > 0$ on (a^*, ∞) . Therefore we have that a^* is the minimum of the function. Using this we bound the sum (40) by an integral:

$$-\sum_{t=0}^{\infty} \sigma''(\cdot)|_{a_t^*} \leq -\int_0^{\infty} \sigma''(a) da - 2\tau\sigma''(x^*) = 4^{-1} - 2\tau\sigma''(x^*) .$$

Substituting this into (40) and (39) we have that:

$$\max_t |\xi_t| \leq 2^{-1}\tau(4^{-1} - 2\tau\sigma''(x^*)) \stackrel{\text{def}}{=} c\tau + d\tau^2 \approx c\tau . \quad (41)$$

We can make the approximation arbitrarily tight by decreasing the step size.

Acknowledgments

This research was partially supported by Army contract DAAD19-02-1-3-0389 to Cylab, and NSF Grant BCS0941518 to the Department of Statistics, both at Carnegie Mellon University. This research was also supported by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the IDM Programme Office.

References

- [1] Aggarwal, C. and Yu, P. S., eds. (2008). *Privacy Preserving Data Mining: Models and Algorithms*. New York: Springer-Verlag.
- [2] Blake, I. and Kolesnikov, V. (2004). Strong conditional oblivious transfer and computing on intervals. In *Advances in Cryptology – ASIACRYPT 2004*, vol. 3329 of *LNCS*. Springer. 515–529.
- [3] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. New York: Cambridge University Press.
- [4] Chaudhuri, K., Monteleoni, C., and Sarwate, A. D. (2011). Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109.
- [5] Chen, B.-C., Kifer, D., LeFevre, K., and Machanavajjhala, A. (2009). Privacy-preserving data publishing. *Foundations and Trends in Databases*, 2(Nos. 1-2):1–167.
- [6] Dwork, C. (2008). Differential privacy: A survey of results. In *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation (TAMC 2008)*. Springer-Verlag. 1–19.
- [7] Fienberg, S. E., Fulp, W. J., and Slavkovic, A. B. and Wrobel, T. A. (2006). “Secure” log-linear and logistic regression analysis of distributed databases. In *Privacy in Statistical Databases: CENEX-SDC Project International Conference (PSD 2006)*, vol. 4302 of *LNCS*. Springer. 277–290.
- [8] Fienberg, S. E., Slavkovic, A. B., and Nardi, Y. (2009). Valid statistical analysis for logistic regression with multiple sources. In P. Kantor and M. Lesk, eds., *Proceedings of the Workshop on Interdisciplinary Studies in Information Privacy and Security (ISIPS 2008)*, vol. 5661 of *LNCS*. New York: Springer.
- [9] Goethals, B., Laur, S., Lipmaa, H., Mielikainen, T. (2004). On secure scalar product computation for privacy-preserving data mining. In *ISISC 2004*.
- [10] Goldreich, O. (2004). *Foundations of Cryptography: Volume 2 Basic Applications*. New York: Cambridge University Press.
- [11] Goldwasser, S. (1997). Multi-party computations: Past and present. In *Proceedings of the 16th Annual ACM Symposium on Principles of Distributed Computing (PODC '97)*. New York: ACM. 1–6.
- [12] Guo, C. and Higham N. J. (2006). A Schur-Newton method for the matrix p th root and its inverse. *SIAM Journal on Matrix Analysis and Applications*, 28(3):788–804.
- [13] Hall, R. and Fienberg, S. E. (2011). Privacy preserving record linkage. In *Privacy in Statistical Databases (PSD 2010)*, vol. 6344 of *LNCS*. Springer. 269–283.

- [14] Hall, R. and Nardi, Y. Fienberg, S. E., (2011). Secure multiple linear regression based on homomorphic encryption. *Journal of Official Statistics*, 27(4):669–691.
- [15] Jagannathan, G. and Wright, R. (2008) Privacy-preserving imputation of missing data. *Data Knowledge Engineering*, 65(1):40–56.
- [16] Karr A. F., and Lin, X., and Reiter, J. P. and Sanil, A . P. (2005) Secure regression on distributed databases. *Journal of Computational and Graphical Statistics*, 14(2):263–279.
- [17] Karr A. F., and Lin, X., and Reiter, J. P. and Sanil, A . P. (2006) Secure analysis of distributed databases. In D. Olwell, A. G. Wilson, and G. Wilson, eds., *Statistical Methods in Counterterrorism: Game Theory, Modeling, Syndromic Surveillance, and Biometric Authentication*. New York: Springer. 237–261.
- [18] Lindell, Y. and Pinkas, B. (2002). Privacy preserving data mining. *Journal of Cryptology*, 15(3):177–206.
- [19] Lindell, Y. and Pinkas, B. (2009). Secure multiparty computation for privacy-preserving data mining. *Journal of Privacy and Confidentiality*, 1(1):59–98.
- [20] Malkhi, D. and Nisan, N. and Pinkas, B. and Sella, Y. (2004). Fairplay: A secure two-party computation system. In *Proceedings of the 13th Conference on USENIX Security Symposium*. 287–302.
- [21] Massart, P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283.
- [22] Minka, T. (2003). A comparison of numerical optimizers for logistic regression. Unpublished manuscript.
- [23] Paillier, P. (1999). Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology – EUROCRYPT ’99*, vol. 1592 of *LNCS*. Springer-Verlag. 223–238.
- [24] Vaidya, J. and Zhu, Y. and Clifton, C. (2005). *Privacy Preserving Data Mining*. New York: Springer.
- [25] Yao, A. C. (1982). Protocols for secure computations. In *Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society. 160–164.
- [26] Yao, A. C. (1986). How to generate and exchange secrets. In *Proceedings of the 27th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society. 162–167.

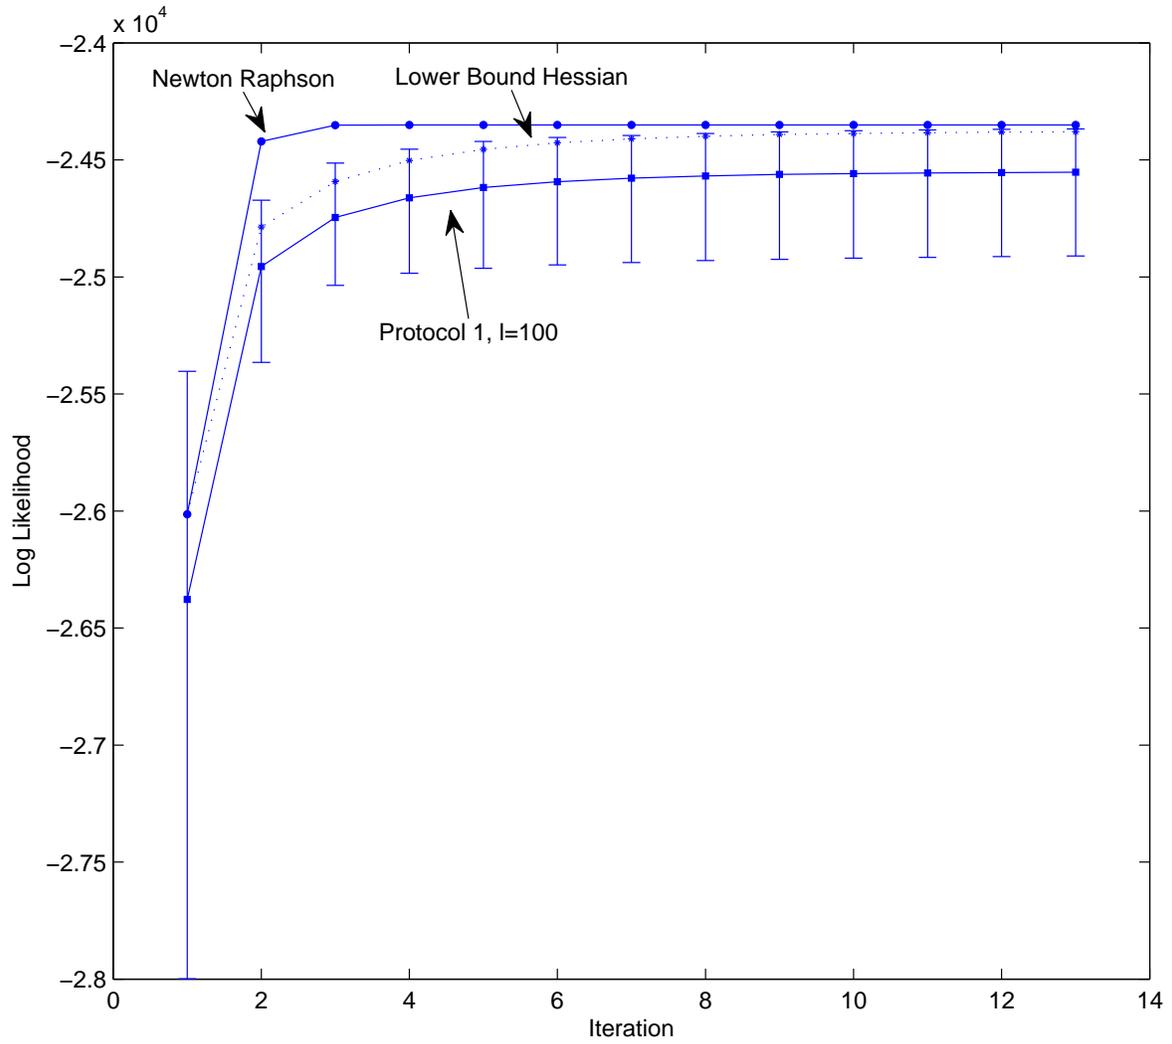


Figure 1: Log Likelihood vs iteration number for protocol 1 with $L = 100$, and that of the “Hessian Lower bound” algorithm, which is the same as protocol 1 except with exact sigmoid evaluations. We also compare to the full newton raphson method, which inverts the Hessian on each iteration.

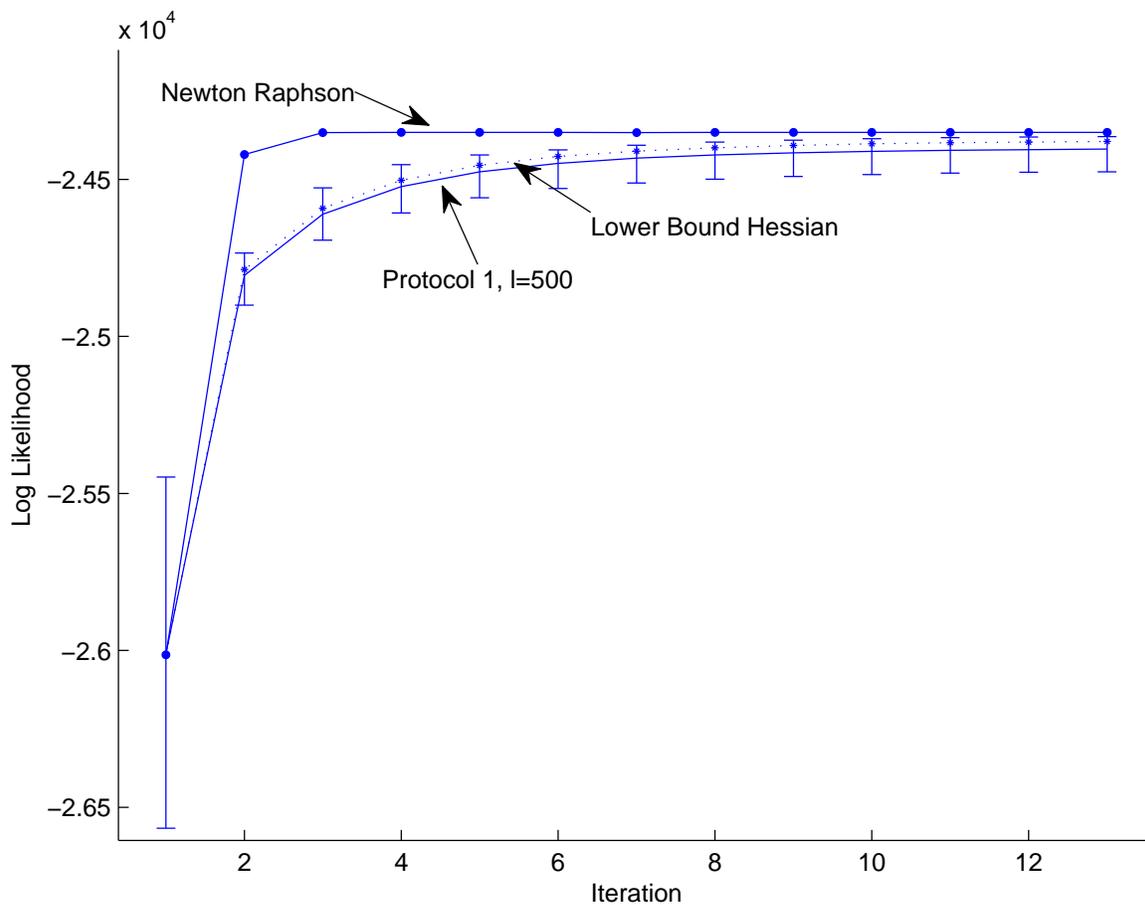


Figure 2: Log Likelihood vs iteration number for protocol 1 with $L = 500$, and that of the “Hessian Lower bound” algorithm, which is the same as protocol 1 except with exact sigmoid evaluations. We also compare to the full newton raphson method, which inverts the Hessian on each iteration.

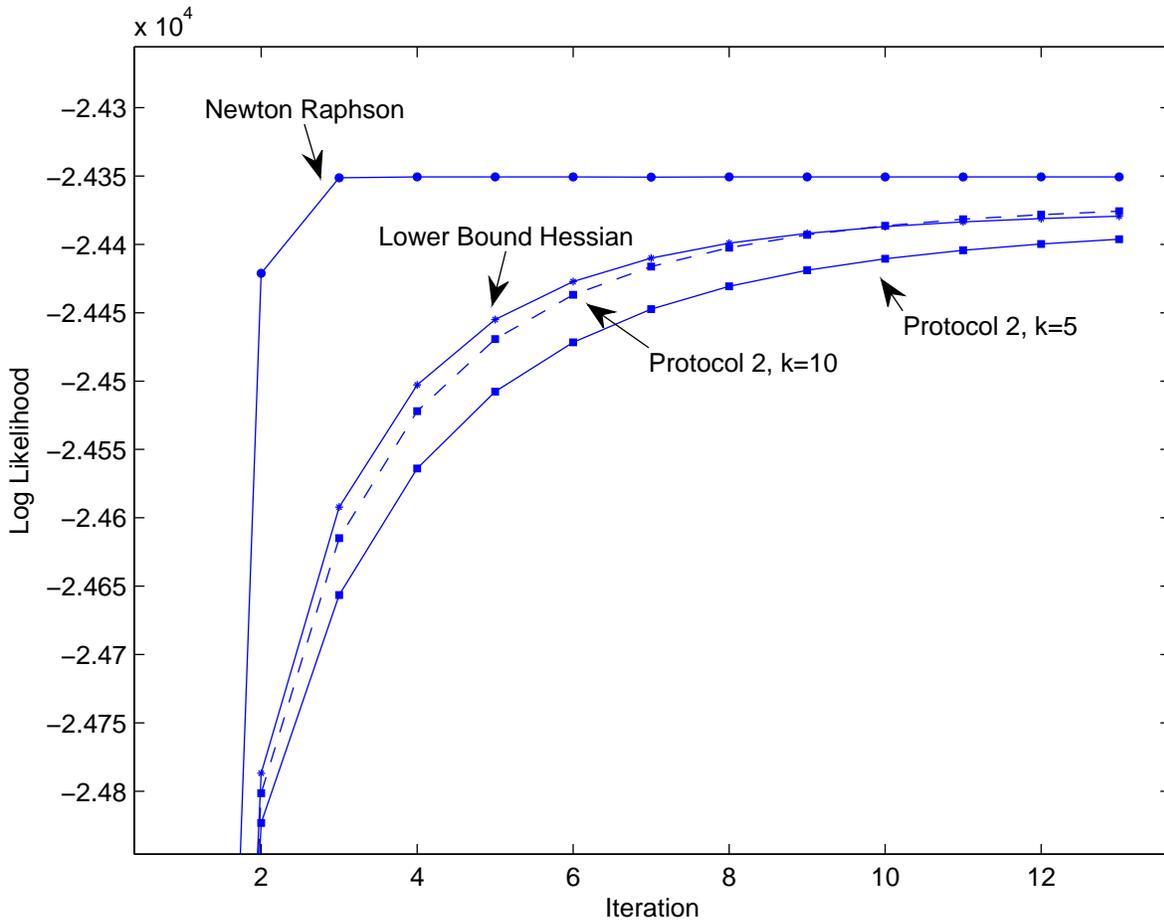


Figure 3: Log Likelihood vs iteration number for protocol 2 with $k = 5, 10$, and that of the “Hessian Lower bound” algorithm, which is the same as protocol 1 except with exact sigmoid evaluations. We also compare to the full newton raphson method, which inverts the Hessian on each iteration.