
An Estimation and Analysis Framework for the Rasch Model

Andrew S. Lan¹ Mung Chiang² Christoph Studer³

Abstract

The Rasch model is widely used for item response analysis in applications ranging from recommender systems to psychology, education, and finance. While a number of estimators have been proposed for the Rasch model over the last decades, the available analytical performance guarantees are mostly asymptotic. This paper provides a framework that relies on a novel linear minimum mean-squared error (L-MMSE) estimator which enables an exact, nonasymptotic, and closed-form analysis of the parameter estimation error under the Rasch model. The proposed framework provides guidelines on the number of items and responses required to attain low estimation errors in tests or surveys. We furthermore demonstrate its efficacy on a number of real-world collaborative filtering datasets, which reveals that the proposed L-MMSE estimator performs on par with state-of-the-art nonlinear estimators in terms of predictive performance.

1. Introduction

This paper presents a novel framework that enables an exact, nonasymptotic, and closed-form analysis of the parameter estimation error under the Rasch model. The Rasch model was proposed in 1960 for modeling the responses of students/users to test/survey items (Rasch, 1960), and has enjoyed great success in applications including (but not limited to) psychometrics (van der Linden & Hambleton, 2013), educational tests (Lan et al., 2016), crowdsourcing (Whitehill et al., 2009), public health (Cappelleri et al., 2014), and even market and financial research (Schellhorn & Sharma, 2013; Brzezińska, 2016). Mathematically, the (dichotomous) Rasch model, also known as the 1PL item

¹Department of Electrical Engineering, Princeton University

²Purdue University ³School of Electrical and Computer Engineering, Cornell University. Correspondence to: Andrew S. Lan <andrew.lan@princeton.edu>.

Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, PMLR 80, 2018. Copyright 2018 by the author(s).

response theory (IRT) model (Lord, 1980), is given by

$$p(Y_{u,i} = 1) = \Phi(a_u - d_i), \quad (1)$$

where $Y_{u,i} \in \{-1, +1\}$ denotes the response of user u to item i , where $+1$ stands for a correct response and -1 stands for an incorrect response. The parameters $a_u \in \mathbb{R}$ model the scalar *abilities* of users $u = 1, \dots, U$ and the parameters $d_i \in \mathbb{R}$ model the scalar *difficulties* of items $i = 1, \dots, Q$. The function $\Phi(x) = \int_{-\infty}^x \mathcal{N}(t; 0, 1) dt$, often referred to as the *inverse probit link function*¹, is the cumulative distribution function of a standard normal random variable, where $\mathcal{N}(t; 0, 1)$ denotes the probability density function of a standard normal random variable evaluated at t .

The literature describes a range of parameter estimation methods under the Rasch model and related IRT models; see (Baker & Kim, 2004) for an overview. However, existing analytical results for the associated *parameter estimation error* are limited; see (Tsutakawa & Johnson, 1990) for an example. The majority of existing results have been proposed in the psychometrics and educational measurement literature; see, e.g., (Carroll et al., 2006) for a survey. The proposed analysis tools rely, for example, on multiple imputation (Yang et al., 2012) or Markov chain Monte Carlo (MCMC) techniques (Patz & Junker, 1999), and are thus not analytical. Hence, their accuracy strongly depends on the available data.

Other analysis tools use the Fisher information matrix (Zhang et al., 2011; Yang et al., 2012) to obtain *lower bounds* on the estimation error. Such methods are of *asymptotic* nature, i.e., they yield accurate results only when the number of users and items tend to infinity. For real-world settings with limited data, these bounds are typically loose; As an example, in computerized adaptive testing (CAT) (Chang & Ying, 2009), a user enters the system and starts responding to items. The system maintains an estimate of their ability parameter, and adaptively selects the next-best item to assign to the user that is most informative of the ability estimate. Calculating the informativeness of each item requires an analysis of the uncertainty in the ability

¹While some publications assume the inverse logit link function, i.e., the sigmoid $\Phi(x) = \frac{1}{1+e^{-x}}$, in most real-world applications the choice of the link function has no significant performance impact. In what follows, we will focus on the inverse probit link function for reasons that will be discussed in Section 3.

estimate. Initially, after the user has only responded to a few items, these asymptotic methods lead to highly inaccurate analyses, which may lead to poor item selections.

Another family of analysis tools relies on concentration inequalities and yield *probabilistic* bounds, i.e., bounds that hold with high probability (Bunea, 2008; Filippi et al., 2010). Such results are often impractical in real-world applications. However, an exact analysis of the estimation error of the Rasch model is critical to ensure the a certain degree of reliability of assessment scores in tests (Thompson, 2002).

1.1. Contributions

We propose a novel framework for the Rasch model that enables an *exact, nonasymptotic*, and *closed-form* analysis of the parameter estimation error. To this end, we generalize a recently-proposed linear estimator for binary regression (Lan et al., 2018) to the Rasch model, which enables us to derive a *sharp upper bound* on the mean squared error (MSE) of model parameter estimates. Our analytical results are in stark contrast to existing analytical results which either provide loose lower bounds or are asymptotic in nature, rendering them impractical in real-world applications.

To demonstrate the efficacy of our framework, we provide experimental results on both synthetic and real-world data. First, using synthetic data, we show that our upper bound on the MSE is (often significantly) tighter than the Fisher information-based lower bound, especially when the problem size is small and when the data is noisy. Therefore, our framework enables a more accurate analysis of the estimation error in real-world settings. Second, using real-world student question response and user movie rating datasets, we show that our linear estimator achieves competitive predictive performance to more sophisticated, nonlinear estimators for which sharp performance guarantees are unavailable.

2. Rasch Model and Probit Regression

The Rasch model in (1) can be written in equivalent matrix-vector form as follows (Hoff, 2009):

$$\mathbf{y} = \text{sign}(\mathbf{D}\mathbf{x} + \mathbf{w}). \quad (2)$$

Here, the UQ -dimensional vector $\mathbf{y} \in \{-1, +1\}^{UQ}$ contains all user responses to all items, the Rasch model matrix $\mathbf{D} = [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U}, \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U]$ is constructed with the Kronecker product operator \otimes , identity matrices \mathbf{I} , all-ones vectors $\mathbf{1}$, and the vector $\mathbf{x}^T = [\mathbf{a}^T, -\mathbf{d}^T]$ to be estimated consists of the user abilities $\mathbf{a} \in \mathbb{R}^U$ and item difficulties $\mathbf{d} \in \mathbb{R}^Q$. The “noise” vector \mathbf{w} contains i.i.d. standard normal random variables. In this equivalent form, parameter estimation under the Rasch model can be casted as a *probit regression* problem (Bliss, 1935), for which numerous estimators have been proposed in the past.

2.1. Estimators for Probit Regression

The two most prominent estimators for probit regression are the posterior mean (PM) estimator, given by

$$\hat{\mathbf{x}}^{\text{PM}} = \mathbb{E}_{\mathbf{x}}[\mathbf{x}|\mathbf{y}] = \int_{\mathbb{R}^N} \mathbf{x} p(\mathbf{x}|\mathbf{y}) d\mathbf{x}, \quad (3)$$

and the maximum a-posteriori (MAP) estimator, given by

$$\hat{\mathbf{x}}^{\text{MAP}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} -\sum_{m=1}^M \log(\Phi(y_m \mathbf{d}_m^T \mathbf{x})) + \frac{1}{2} \mathbf{x}^T \mathbf{C}_{\mathbf{x}}^{-1} \mathbf{x}.$$

Here, $p(\mathbf{x}|\mathbf{y})$ denotes the posterior probability of the vector \mathbf{x} given the observations \mathbf{y} under the model (2), \mathbf{d}_m^T denotes the m th row of the matrix of covariates \mathbf{D} , and $\mathbf{C}_{\mathbf{x}}$ denotes the covariance matrix of the multivariate Gaussian prior on \mathbf{x} . A special case of the MAP estimator is the well-known maximum likelihood (ML) estimator, which does not impose a prior distribution on \mathbf{x} .

The PM estimator is optimal in terms of minimizing the MSE of the estimated parameters, which is defined as

$$MSE(\hat{\mathbf{x}}) = \mathbb{E}_{\mathbf{x}, \mathbf{w}} [\|\mathbf{x} - \hat{\mathbf{x}}\|^2]. \quad (4)$$

However, there are no simple methods to evaluate the expectation in (3) under the probit model. Thus, one typically resorts to Markov chain Monte Carlo (MCMC) methods (Albert & Chib, 1993) to perform PM estimation, which can be computationally intensive. In contrast to the PM estimator, MAP and ML estimation is generally less complex since it can be implemented using standard convex optimization algorithms (Nocedal & Wright, 2006; Hastie et al., 2010; Goldstein et al., 2014). On the flipside, MAP and ML estimation is not optimal in terms of minimizing the MSE in (4). In contrast to such well-established, nonlinear estimators, we build our framework on the family of *linear estimators* recently proposed in (Lan et al., 2018). There, a linear minimum MSE (L-MMSE) estimator was proposed for a certain class of probit regression problems. This L-MMSE estimator was found to perform on par with the PM estimator and outperforms the MAP estimator in terms of the MSE for certain settings, while enabling an exact and nonasymptotic analysis of the MSE.

2.2. Analytical Performance Guarantees

In the statistical estimation literature, there exists numerous analytical results characterizing the estimation errors for binary regression problems in the *asymptotic* setting. For example, (Brillinger, 1982) shows that least squares estimation is particularly effective when the design matrix \mathbf{D} has i.i.d. Gaussian entries and the number of observations approaches infinity; in this case, its performance was shown to differ from that of the PM estimator only by a constant factor. Recently, (Thrampoulidis et al., 2015) provides a related analysis in the case that the parameter vector \mathbf{x} is

sparse. Another family of *probabilistic* results relies on the asymptotic normality property of ML estimators, either in the standard (dense) setting (Gourieroux & Monfort, 1981; Fahrmeir & Kaufmann, 1985) or the sparse setting (Bunea, 2008; Bach, 2010; Ravikumar et al., 2010; Plan & Vershynin, 2013), providing bounds on the MSE with high probability. Since numerous real-world applications, such as the Rasch model, rely on deterministic, structured matrices and have small problem dimensions, existing analytical performance bounds are often loose; see Section 4 for experiments that support this claim.

3. Main Results

Our main result is as follows; the proof is given in Appendix A.

Theorem 1. *Assume that $\mathbf{x} \sim \mathcal{N}(\bar{\mathbf{x}}, \mathbf{C}_x)$ with mean vector $\bar{\mathbf{x}}$ and positive definite covariance matrix \mathbf{C}_x , and assume that the vector \mathbf{w} contains i.i.d. standard normal random variables. Consider the general probit regression model*

$$\mathbf{y} = \text{sign}(\mathbf{D}\mathbf{x} + \mathbf{m} + \mathbf{w}), \quad (5)$$

where \mathbf{D} is a given matrix of covariates and \mathbf{m} is a given bias vector. Then, the L-MMSE estimate is given by

$$\hat{\mathbf{x}}^{\text{L-MMSE}} = \mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{y} + \mathbf{b},$$

where we use the following quantities:

$$\begin{aligned} \mathbf{E} &= 2\text{diag}(\mathcal{N}(\mathbf{c}; 0, 1) \odot \text{diag}(\mathbf{C}_z)^{-\frac{1}{2}}) \mathbf{D} \mathbf{C}_x \\ \mathbf{c} &= \bar{\mathbf{z}} \odot \text{diag}(\mathbf{C}_z)^{-1/2} \\ \bar{\mathbf{z}} &= \mathbf{D}\bar{\mathbf{x}} + \mathbf{m} \\ \mathbf{C}_z &= \mathbf{D} \mathbf{C}_x \mathbf{D}^T + \mathbf{I} \\ \mathbf{C}_y &= 2(\Phi_2(\mathbf{c}\mathbf{1}^T, \mathbf{1}\mathbf{c}^T; \mathbf{R}) + \Phi_2(-\mathbf{c}\mathbf{1}^T, -\mathbf{1}\mathbf{c}^T; \mathbf{R})) \\ &\quad - \mathbf{1}_{M \times M} - \bar{\mathbf{y}}\bar{\mathbf{y}}^T \\ \mathbf{R} &= \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}) \mathbf{C}_z \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}) \\ \bar{\mathbf{y}} &= \Phi(\mathbf{c}) - \Phi(-\mathbf{c}) \\ \mathbf{b} &= \bar{\mathbf{x}} - \mathbf{E}^T \mathbf{C}_y^{-1} \bar{\mathbf{y}}. \end{aligned}$$

Here, $\Phi_2(x, y; \rho)$ denotes the cumulative density of a two-dimensional zero-mean Gaussian distribution with covariance matrix $[1 \ \rho; \rho \ 1]$ with $\rho \in [0, 1]$, defined as

$$\Phi_2(x, y; \rho) = \int_{-\infty}^x \int_{-\infty}^y \frac{1}{2\pi\sqrt{1-\rho^2}} e^{-\frac{s^2-2\rho st+t^2}{2(1-\rho^2)}} dt ds$$

and is applied element-wise on matrices. Furthermore, the associated estimation MSE is given by

$$\text{MSE}(\hat{\mathbf{x}}^{\text{L-MMSE}}) = \text{tr}(\mathbf{C}_x - \mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{E}).$$

We note that the linear estimator developed in (Lan et al., 2018, Thm. 1) is a special case of our result with $\bar{\mathbf{x}} = \mathbf{0}$ and

$\mathbf{m} = \mathbf{0}$. As we will show below, including both of these terms will be essential for our analysis.

Remark 1. *We exclusively focus on probit regression since the matrices \mathbf{E} and \mathbf{C}_y exhibit tractable expressions under this model. We are unaware of any closed-form expressions for these quantities in the logistic regression case.*

As an immediate consequence of the fact that the PM estimator minimizes the MSE, we can use Theorem 1 to obtain the following upper bound on the MSE of the PM estimator.

Corollary 2. *The MSE of the PM estimator is upper-bounded as follows:*

$$\text{MSE}(\mathbf{x}^{\text{PM}}) \leq \text{MSE}(\hat{\mathbf{x}}^{\text{L-MMSE}}). \quad (6)$$

As we will demonstrate in Section 4, this upper bound on the MSE turns out to be surprisingly sharp for a broad range of parameters and problem settings.

We now specialize Theorem 1 for the Rasch model and use Corollary 2 to analyze the associated MSE. We divide our results into two cases: (i) both the user abilities and item difficulties are unknown and (ii) one of the two sets of parameters is known and the other is unknown. Due to symmetry in the Rasch model, we will present our results with unknown/known item difficulties while the user abilities are unknown and to be estimated; a corresponding analysis on the estimation error of item parameters follows immediately.

3.1. First Case: Unknown Item Parameters

We now analyze the case in which both the user abilities and item difficulties are unknown and need to be estimated. In practice, this scenario is relevant if a new set of items are deployed with little or no prior knowledge on their difficulty parameters. We assume that there is no missing data, i.e., we observe all user responses to all items.² In the psychometrics literature (see, e.g., (Linacre, 1999)), one typically assumes that the entries of the ability \mathbf{a} and difficulty vectors \mathbf{d} are i.i.d. zero-mean Gaussian with variance σ_a^2 and σ_d^2 , respectively, i.e., $a_u \sim \mathcal{N}(0, \sigma_a^2)$ and $d_i \sim \mathcal{N}(0, \sigma_d^2)$, which can be included in our model assumptions. Thus, we can leverage the special structure of the Rasch model, since it corresponds to a special case of the generic probit regression model in (5) with $\mathbf{D} = [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U}, \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U]$ and $\mathbf{m} = \mathbf{0}$. We have the following result on the MSE of the L-MMSE estimator; the proof is given in Appendix B.

Theorem 3. *Assume that $\sigma_a^2 = \sigma_d^2 = \sigma_x^2$ and the covariance matrix of \mathbf{x} is $\mathbf{C}_x = \sigma_x^2 \mathbf{I}_{(U+Q) \times (U+Q)}$. Let*

$$s = \frac{2}{\pi} \arcsin\left(\frac{\sigma_x^2}{2\sigma_x^2 + 1}\right).$$

²Our analysis can readily be generalized to missing data; the results, however, depend on the missing data pattern.

Then, the MSE of the L-MMSE estimator of user abilities under the Rasch model is given by

$$\begin{aligned} \text{MSE}_a &= \mathbb{E}_{\mathbf{x}, \mathbf{w}}[(a_u - \hat{a}_u)^2] = \\ &= \sigma_x^2 \left(1 - \frac{2}{\pi} \frac{\sigma_x^2}{2\sigma_x^2 + 1} \frac{sQ(Q+U-3)+1}{(s(Q-2)+1)(s(Q+U-2)+1)} \right). \end{aligned} \quad (7)$$

To the best of our knowledge, Theorem 3 is the first exact, nonasymptotic, and closed-form analysis of the MSE of a parameter estimation method for the Rasch model. From (7), we see that if σ_x^2 is held constant, then the relationship between MSE_a and the numbers of users (U) and items (Q) is given by the ratio of two second-order polynomials. If the signal-to-noise ratio (SNR) is low (or, equivalently, the data is noisy), i.e., $\sigma_x^2 \ll \sigma_n^2$, then we have $\frac{\sigma_x^2}{2\sigma_x^2 + 1} \approx 0$ and hence, $s = \frac{2}{\pi} \arcsin\left(\frac{\sigma_x^2}{2\sigma_x^2 + 1}\right) \approx 0$. In this case, we have $\text{MSE}_a \approx \sigma_x^2$, i.e., increasing the number of users/items does not affect the accuracy of the ability and difficulty parameters of users and items; this behavior is as expected.

When $U, Q \rightarrow \infty$, the MSE satisfies

$$\text{MSE}_a \rightarrow \sigma_x^2 \left(1 - \frac{\sigma_x^2}{2\sigma_x^2 + 1} \arcsin^{-1}\left(\frac{\sigma_x^2}{2\sigma_x^2 + 1}\right) \right), \quad (8)$$

which is a non-negative quantity. This result implies that the L-MMSE estimator has a residual MSE even as the number of users/items grows large. More specifically, since $x \leq \arcsin(x)$ for $x \in [0, 1]$, this residual error approaches $\sigma_x^2(1 - \frac{3}{\pi})$ at high values of SNR. We note, however, this result does not imply that the L-MMSE estimator is not consistent under the Rasch model, since the number of parameters to be estimated ($U + Q$) grows with the number of the observations (UQ) instead of remaining constant.

Remark 2. The above MSE analysis is data-independent, in contrast to error estimates that rely on the responses \mathbf{y} (which is, for example, the case for method in (Carroll et al., 2006)). This fact implies that our result provides an error estimate before observing \mathbf{y} . Thus, Theorem 3 provides guidelines on how many items to include and how many users to recruit for a study, given a desired MSE level on the user ability and item difficulty parameter estimates.

3.2. Second Case: Known Item Difficulties

We now analyze the case in which the user abilities are unknown and need to be estimated; the item difficulties (\mathbf{d}) are given. In practice, this scenario is relevant if a large number of users previously responded to a set of items so that a good estimate of the item difficulties is available. Let a denote the scalar ability parameter of an user. Then, their responses to items are modeled as

$$p(\mathbf{y} = 1) = \Phi(\mathbf{1}_Q a - \mathbf{d}).$$

The following result follows from Theorem 1 by setting $\mathbf{x} = a$, $\bar{\mathbf{x}} = \bar{x}$, $\mathbf{C}_x = \sigma_x^2$, $\mathbf{D} = \mathbf{1}_Q$, and $\mathbf{m} = -\mathbf{d}$.

Corollary 4. Assume that $a \sim \mathcal{N}(\bar{x}, \sigma_x^2)$. Then, the L-MMSE estimate of user ability is given by

$$\hat{a} = \mathbf{e}^T \mathbf{C}_y^{-1} \mathbf{y} + b,$$

where

$$\begin{aligned} \mathbf{e} &= 2 \frac{\sigma_x^2}{\sqrt{\sigma_x^2 + 1}} \mathcal{N}(\mathbf{c}; 0, 1) \\ \mathbf{c} &= \bar{\mathbf{z}} \odot \text{diag}(\mathbf{C}_z)^{-1/2} \\ \bar{\mathbf{z}} &= \bar{x} \mathbf{1}_Q - \mathbf{d} \\ \mathbf{C}_z &= \sigma_x^2 \mathbf{1}_{Q \times Q} + \mathbf{I} \\ \bar{\mathbf{y}} &= \Phi(\mathbf{c}) - \Phi(-\mathbf{c}) \\ \mathbf{C}_y &= 2(\Phi_2(\mathbf{c}\mathbf{1}^T, \mathbf{1}\mathbf{c}^T, \mathbf{R}) + \Phi_2(-\mathbf{c}\mathbf{1}^T, -\mathbf{1}\mathbf{c}^T, \mathbf{R})) \\ &\quad - \mathbf{1}_{M \times M} - \bar{\mathbf{y}}\bar{\mathbf{y}}^T \\ \mathbf{R} &= \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}) \mathbf{C}_z \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}). \end{aligned}$$

The MSE of the user ability estimate is given by $\text{MSE}(\hat{a}) = \sigma_x^2 - \mathbf{e}^T \mathbf{C}_y^{-1} \mathbf{e}$.

4. Numerical Results

We now experimentally demonstrate the efficacy of the proposed framework. First, we use synthetically generated data to numerically compare our L-MMSE-based upper bound on the MSE of the PM estimator to the widely-used lower bound based on Fisher information (Zhang et al., 2011; Yang et al., 2012). We then use several real-world collaborative filtering datasets to show that the L-MMSE estimator achieves comparable predictive performance to that of the PM and MAP estimators.

4.1. Experiments with Synthetic Data

We start with synthetic data to demonstrate the exact and nonasymptotic nature of our analytical MSE expressions.

4.1.1. FIRST CASE: UNKNOWN ITEM PARAMETERS

Experimental Setup We vary the number of users $U \in \{20, 50, 100\}$ and the number of items $Q \in \{20, 50, 100, 200\}$. We generate the user ability and item difficulty parameters from zero-mean Gaussian distributions with variance $\sigma_x^2 = \sigma_a^2 = \sigma_d^2$. We vary σ_x^2 so that the signal-to-noise ratio (SNR) corresponds to $\{-10, 0, 10\}$ decibels (dB). We then randomly generate the response from each user to each item, $Y_{u,i}$, according to (1). We repeat these experiments for 1,000 random instances of user and item parameters and responses, and report the averaged results.

We compute the L-MMSE-based upper bound on the MSE of the PM estimator using Theorem 1 and the Fisher

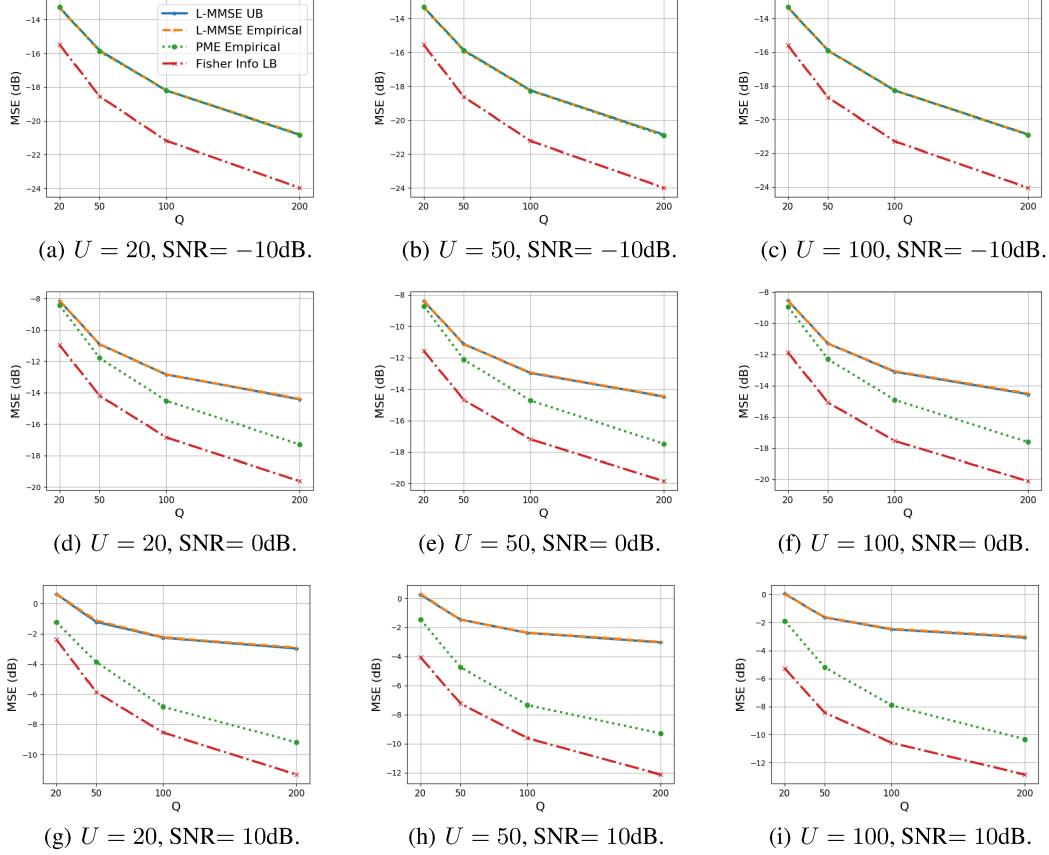


Figure 1. Empirical MSEs of the L-MMSE and PM estimators and the L-MMSE-based upper and Fisher information-based lower bounds on the MSE of the PM estimator for various SNR levels and problem sizes, when both user and item parameters are unknown. We see that the upper bound is tight at low SNR and at all SNRs when the problem size is small.

information-based lower bound using the method detailed in (Zhang et al., 2011; Yang et al., 2012). Since the calculation of the Fisher information matrix requires the true values of the user ability and item difficulty parameters (which are to be estimated in practice), we use the PM estimates of these parameters instead. We also calculate the empirical parameter estimation MSEs of the L-MMSE and PM estimators. To this end, we use a standard Gibbs sampling procedure (Albert & Chib, 1993); we use the mean of the generated samples over 20,000 iterations as the PM estimate after a burn-in phase of 10,000 iterations. We then use these estimates to calculate the empirical MSE.

Results and Discussion Fig. 1 shows the empirical MSEs of the L-MMSE and PM estimators, together with the L-MMSE-based upper bound and the Fisher information-based lower bound on the MSE of the PM estimator, for every problem size and every SNR. First, we see that the analytical and empirical MSEs of the L-MMSE estimator match perfectly, which confirms that our analytical MSE expressions are *exact*. We also see that for low SNR (i.e., the first row of Fig. 1), our L-MMSE upper bound on the MSE of the PM estimator is tight. Moreover, at all noise levels,

the L-MMSE-based upper bound is tighter at small problem sizes, while the Fisher information-based lower bound is tighter at very large problem sizes and at high SNR.

These results confirm that our L-MMSE-based upper bound on the MSE is nonasymptotic, while the Fisher information-based lower bound is asymptotic and thus only tight at very large problem sizes. Therefore, the L-MMSE-based upper bound is more practical than the Fisher information-based lower bound in real-world applications, especially for situations like the initial phase of CAT when the number of items a user has responded to is small.

4.1.2. CASE TWO: KNOWN ITEM PARAMETERS

Experimental Setup In this experiment, we randomly generate the item parameters from the standard normal distribution ($\sigma_d^2 = 1$) and treat these parameters as known; we then estimate the user ability parameters via Theorem 4. The rest of the experimental setup remains unchanged.

Results and Discussion Fig. 2 shows the empirical MSEs of the L-MMSE and PM estimators, together with the L-MMSE-based upper bound and the Fisher information-

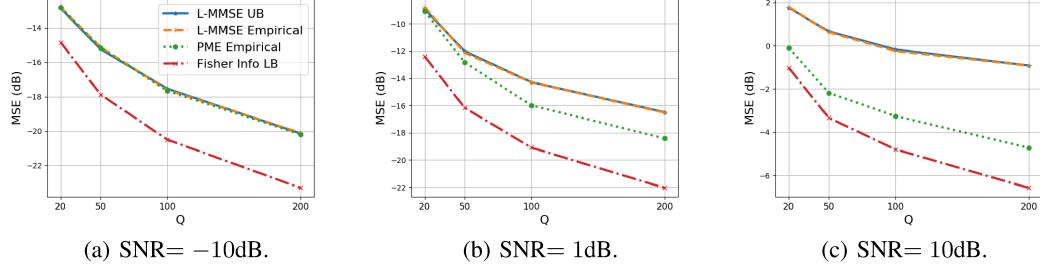


Figure 2. Empirical MSEs of the L-MMSE and PM estimators and the L-MMSE-based upper and Fisher Information-based lower bounds on the MSE of the PM estimator for various SNR levels and various problem sizes, when item parameters are known. We see that the upper bound is tight at low SNR and at higher SNRs when the problem sizes are small.

based lower bound on the MSE of the PM estimator, for every problem size and every SNR. We see that the analytical and empirical MSEs of the L-MMSE estimator match. We also see that the L-MMSE-based upper bound on the MSE is tighter than the Fisher information-based lower bound at low SNR levels (-10 dB and 1 dB), and especially when the problem size is small (less than 50 items). These results further confirm that our L-MMSE-based upper bound on the MSE is nonasymptotic, and is thus practical in the “cold-start” setting of recommender systems.

4.2. Experiments with Real-World Data

We now test the performance of the proposed L-MMSE estimator using a variety of real-world datasets. Since the noise model in real-world datasets is generally unknown, we also consider the performance of MAP estimation using the inverse logit link function (Logit-MAP).

Datasets We perform our experiments using a range of collaborative filtering datasets. These datasets are matrices that contain binary-valued ratings (or graded responses) of users (or students) to movies (or items). For these datasets, we use the probit Rasch model. The datasets include (i) “MT”, which consists of students’ binary-valued (correct/incorrect) graded responses to questions in a high-school algebra test, with $U = 99$ students’ 3,366 responses to $Q = 34$ questions, (ii) “SS”, which consists of student responses in a signals and systems course, with $U = 92$ students’ 5,692 responses to $Q = 203$ questions, (iii) “edX”, which consists of student responses in an edX course, with $U = 3241$ students’ 177,181 responses to $Q = 191$ questions, and (iv) “ML”, a processed version of the *ml-100k* dataset from the MovieLens project (Herlocker et al., 1999), with 37,175 integer-valued ratings by $U = 943$ users to $Q = 1152$ movies. We adopt the procedure used in (Davenport et al., 2014) to transform the dataset into binary values by comparing each rating to the overall average rating.

Experimental Setup We evaluate the prediction performance of the L-MMSE, MAP, PM, and Logit-MAP estimators using ten-fold cross validation. We randomly divide the

entire dataset into ten equally-partitioned folds (of user-item response pairs), leave out one fold as the held-out testing set and use the other folds as the training set. We then use the training set to estimate the learner abilities a_u and item difficulties d_i , and use these estimates to predict user responses on the test set. We tune the prior variance parameter σ_x^2 using a separate validation set (one fold in the training set). To assess the performance of these estimators, we use two common metrics in binary classification problems: prediction accuracy (ACC), which is simply the portion of correct predictions, and area under the receiver operating characteristic curve (AUC) (Jin & Ling, 2005). Both metrics have range in $[0, 1]$, with larger values indicating better predictive performance.

Results and Discussion Tables 1 and 2 show the mean and standard deviation of the performance of each estimator on both metrics across each fold. We observe that the performance of the considered estimators are comparable on the ACC metric, while the L-MMSE estimator performs slightly worse than the MAP, PM, and Logit-MAP estimators for most datasets on the AUC metric.

We find it quite surprising that a well-designed linear estimator performs on par with more sophisticated nonlinear estimators on these real-world datasets. We also note that the L-MMSE estimator is more computationally efficient than the PM estimator. As an example, on the MT and ML datasets, one run of the L-MMSE estimator takes 0.23s and 79s, respectively, while one run of the PM estimator takes 1.9s and 528s (2,000 and 10,000 iterations required for convergence) on a standard laptop computer. These observations suggest that the L-MMSE estimator is computationally efficient and thus scales favorably to large datasets.

5. Conclusions

We have generalized a recently proposed linear estimator for probit regression and applied the method to the classic Rasch model in item response analysis. We have shown that the L-MMSE estimator enables an exact, closed-form, and nonasymptotic MSE analysis, which is in stark contrast to existing analytical results which are asymptotic, probabilis-

Table 1. Mean and standard deviation of the prediction accuracy (ACC) for the L-MMSE, MAP, PM, and Logit-MAP estimators.

	L-MMSE	MAP	PM	Logit-MAP
MT	0.795 ± 0.016	0.796 ± 0.015	0.796 ± 0.016	0.794 ± 0.015
SS	0.860 ± 0.007	0.859 ± 0.007	0.859 ± 0.007	0.859 ± 0.010
edX	0.932 ± 0.001	0.934 ± 0.002	0.935 ± 0.002	0.934 ± 0.002
ML	0.715 ± 0.004	0.713 ± 0.004	0.713 ± 0.004	0.714 ± 0.004

Table 2. Area under the receiver operating characteristic curve (AUC) of the L-MMSE, MAP, PM, and Logit-MAP estimators.

	L-MMSE	MAP	PM	Logit-MAP
MT	0.840 ± 0.016	0.843 ± 0.015	0.843 ± 0.015	0.842 ± 0.015
SS	0.800 ± 0.014	0.803 ± 0.013	0.803 ± 0.013	0.802 ± 0.013
edX	0.900 ± 0.004	0.909 ± 0.004	0.909 ± 0.004	0.909 ± 0.004
ML	0.755 ± 0.005	0.756 ± 0.004	0.756 ± 0.004	0.756 ± 0.004

tic, or loose. As a result, we have shown that the nonasymptotic, L-MMSE-based upper bound on the parameter estimation error of the PM estimator under the Rasch model can be tighter than the common Fisher information-based asymptotic lower bound, especially in practical settings. An avenue of future work is to apply our analysis to models that are more sophisticated than the Rasch model, e.g., the latent factor model in (Lan et al., 2014).

A. Proof of Theorem 4

Let $\mathbf{z} = \mathbf{D}\mathbf{x} + \mathbf{m} + \mathbf{w}$. Thus, $\mathbf{z} \sim \mathcal{N}(\mathbf{D}\bar{\mathbf{x}} + \mathbf{m}, \mathbf{D}\mathbf{C}_x\mathbf{D}^T + \mathbf{I}) := \mathcal{N}(\bar{\mathbf{z}}, \mathbf{C}_z)$. The L-MMSE estimator for \mathbf{x} has the general form of $\hat{\mathbf{x}}^{\text{L-MMSE}} = \mathbf{W}\mathbf{y} + \mathbf{b}$, where $\mathbf{W} = \mathbf{E}\mathbf{C}_y^{-1}$ and $\mathbf{b} = \bar{\mathbf{x}} - \mathbf{W}\bar{\mathbf{y}}$, with

$$\mathbf{C}_y = \mathbb{E}[(\mathbf{y} - \bar{\mathbf{y}})(\mathbf{y} - \bar{\mathbf{y}})^T] = \mathbb{E}[\mathbf{y}\mathbf{y}^T] - \bar{\mathbf{y}}\bar{\mathbf{y}}^T := \tilde{\mathbf{C}}_y - \bar{\mathbf{y}}\bar{\mathbf{y}}^T$$

and

$$\mathbf{E} = \mathbb{E}[(\mathbf{y} - \bar{\mathbf{y}})(\mathbf{x} - \bar{\mathbf{x}})^T] = \mathbb{E}[\mathbf{y}\mathbf{x}^T] - \bar{\mathbf{y}}\bar{\mathbf{x}}^T := \tilde{\mathbf{E}} - \bar{\mathbf{y}}\bar{\mathbf{x}}^T.$$

We need to evaluate three quantities, $\bar{\mathbf{y}}$, $\tilde{\mathbf{C}}_y$, and $\tilde{\mathbf{E}}$.

We start with $\bar{\mathbf{y}}$. Its i th entry is given by

$$\begin{aligned} \bar{y}_i &= \int_{-\infty}^{\infty} \text{sign}(z_i) \mathcal{N}(z_i; \bar{z}_i, [\mathbf{C}_z]_{i,i}) dz_i \\ &= - \int_{-\infty}^0 \mathcal{N}(z_i; \bar{z}_i, [\mathbf{C}_z]_{i,i}) dz_i + \int_0^{\infty} \mathcal{N}(z_i; \bar{z}_i, [\mathbf{C}_z]_{i,i}) dz_i \\ &= \Phi\left(\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}\right) - \Phi\left(-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}\right). \end{aligned}$$

Next, we calculate $\tilde{\mathbf{C}}_y$. Its (i,j) th entry is given by $[\tilde{\mathbf{C}}_y]_{i,j} =$

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{sign}(z_i) \text{sign}(z_j) \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \begin{bmatrix} \bar{z}_i \\ \bar{z}_j \end{bmatrix}, \begin{bmatrix} [\mathbf{C}_z]_{i,i} & [\mathbf{C}_z]_{i,j} \\ [\mathbf{C}_z]_{j,i} & [\mathbf{C}_z]_{j,j} \end{bmatrix}\right) dz_j dz_i$$

$$\begin{aligned} &\stackrel{(a)}{=} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{sign}\left(\frac{z_i + \bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}\right) \text{sign}\left(\frac{z_j + \bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}\right) \\ &\quad \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right) dz_j dz_i \\ &= \underbrace{\int_{-\infty}^{-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}} \int_{-\infty}^{-\frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}} \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right) dz_j dz_i}_{v_1} \\ &\quad + \underbrace{\int_{-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}}^{\infty} \int_{-\frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}}^{\infty} \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right) dz_j dz_i}_{v_2} \\ &\quad - \underbrace{\int_{-\infty}^{-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}} \int_{-\infty}^{-\frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}} \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right) dz_j dz_i}_{v_3} \\ &\quad - \underbrace{\int_{-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}}^{\infty} \int_{-\infty}^{-\frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}} \mathcal{N}\left(\begin{bmatrix} z_i \\ z_j \end{bmatrix}; \mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right) dz_j dz_i}_{v_4} \\ &\stackrel{(b)}{=} 2(v_1 + v_2) - 1 \\ &= 2 \left(\Phi_2\left(\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}, \frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}, \rho\right) \right. \\ &\quad \left. + \Phi_2\left(-\frac{\bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}}, -\frac{\bar{z}_j}{\sqrt{[\mathbf{C}_z]_{j,j}}}, \rho\right) \right) - 1, \end{aligned}$$

where we have used (a) change of variable $\frac{z_i - \bar{z}_i}{\sqrt{[\mathbf{C}_z]_{i,i}}} \rightarrow z_i$ and (b) the fact that $v_1 + v_2 + v_3 + v_4 = 1$. The computation of $\tilde{\mathbf{E}}$ follows from that in (Lan et al., 2018) and is omitted.

B. Proof of Theorem 3

Recall that the expression for the MSE is $\text{tr}(\mathbf{C}_x - \mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{E})$, the critical part is to evaluate $\mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{E}$. We begin by evaluating \mathbf{C}_y^{-1} . For the Rasch model, we have

$\mathbf{D} = [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U}, \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U]$. Therefore, since $\mathbf{C}_x = \sigma_x^2 \mathbf{I}_{U+Q}$, we have

$$\begin{aligned}\mathbf{C}_z &= \mathbf{D} \mathbf{C}_x \mathbf{D}^T + \mathbf{I}_{UQ \times UQ} = \sigma_x^2 \mathbf{D} \mathbf{D}^T + \mathbf{I}_{UQ \times UQ} \\ &= \sigma_x^2 [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U} \quad \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U] \begin{bmatrix} \mathbf{1}_Q^T \otimes \mathbf{I}_{U \times U} \\ \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U^T \end{bmatrix} \\ &\quad + \mathbf{I}_{UQ \times UQ} \\ &= \sigma_x^2 (\mathbf{1}_{Q \times Q} \otimes \mathbf{I}_{U \times U} + \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_{U \times U}) + \mathbf{I}_{UQ \times UQ},\end{aligned}$$

where $\mathbf{1}_{U \times U}$ denotes an all-one matrix with size $U \times U$. Therefore, we see that the $UQ \times UQ$ matrix \mathbf{C}_z consists of three parts: (i) Q copies of the all-ones matrix $\sigma_x^2 \mathbf{1}_{U \times U}$ in its diagonal $U \times U$ blocks, (ii) copies of the matrix $\sigma_x^2 \mathbf{I}_{U \times U}$ in every other off-diagonal $U \times U$ block, plus (iii) a diagonal matrix $\mathbf{I}_{UQ \times UQ}$. Therefore, its diagonal elements are $2\sigma_x^2 + 1$ and its non-zero off-diagonal elements are σ_x^2 .

As detailed in (Lan et al., 2018, (7)), one can show that

$$\begin{aligned}\mathbf{C}_y &= \frac{2}{\pi} \arcsin(\text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}) \mathbf{C}_z \\ &\quad \times \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2})),\end{aligned}$$

we have that the term inside the arcsin function has the same structure as \mathbf{C}_z , with diagonal entries of 1 and non-zero off-diagonal entries as $\frac{\sigma_x^2}{2\sigma_x^2 + 1}$. Therefore, \mathbf{C}_y also has the same structure, with diagonal entries of 1 and non-zero off-diagonal entries as

$$s = \frac{2}{\pi} \arcsin\left(\frac{\sigma_x^2}{2\sigma_x^2 + 1}\right).$$

Since \mathbf{C}_y^{-1} satisfies $\mathbf{C}_y \mathbf{C}_y^{-1} = \mathbf{I}_{UQ \times UQ}$, it is easy to see that the entries of \mathbf{C}_y^{-1} only contain three distinct values (denoted by a , b , and c), and consists of two parts: (i) Q copies of a $U \times U$ matrix with a on its diagonal, b everywhere else, in its diagonal blocks, and (ii) copies of a $U \times U$ matrix with c on its diagonal, d everywhere else, in its other blocks. We next compute a , b , c , and d .

The first column of \mathbf{C}_y^{-1} is given by

$$[a, b\mathbf{1}_{1 \times U-1}, c, d\mathbf{1}_{1 \times U-1}, c, d\mathbf{1}_{1 \times U-1}, \dots]^T.$$

Since its inner product with the first row of \mathbf{C}_y is one (since $\mathbf{C}_y \mathbf{C}_y^{-1} = \mathbf{I}_{UQ \times UQ}$), we get

$$a + (U-1)s b + (Q-1)s c = 1.$$

Similarly, its inner products with the second, $(U+1)-th$, and $(U+2)-th$ rows are all zero; this gives

$$\begin{aligned}sa + ((U-2)s+1)b + (Q-1)s d &= 0, \\ sa + ((Q-2)s+1)c + (U-1)s d &= 0, \\ sb + sc + ((U+Q-4)s+1)d &= 0.\end{aligned}$$

Solving the linear system given by these four equations results in

$$\begin{aligned}a &= \frac{(3U^2+3Q^2-U^2Q-UQ^2+8UQ-15U-15Q+20)s^3}{r} \\ &\quad + \frac{(-U^2-Q^2-3UQ+11U+11Q-22)s^2}{r} \\ &\quad + \frac{(-2U-2Q+8)s-1}{r} \\ b &= \frac{(UQ+Q^2-3U-5Q+8)s^3+(U+2Q-6)s^2+s}{r} \\ c &= \frac{(UQ+U^2-5U-3Q+8)s^3+(2U+Q-6)s^2+s}{r} \\ d &= \frac{-(U+Q-4)s^3-2s^2}{r},\end{aligned}\tag{9}$$

where

$$r = (2s-1)((U-2)s+1)((Q-2)s+1)((Q+U-2)s+1).$$

Now, let \mathbf{A} be the $N \times N$ matrix with c on its diagonal and d everywhere else, \mathbf{B} denote the matrix with $a - d$ on its diagonal and $b - d$ everywhere else, we can write \mathbf{C}_y^{-1} as

$$\mathbf{C}_y^{-1} = \mathbf{1}_{Q \times Q} \otimes \mathbf{A} + \mathbf{I}_{Q \times Q} \otimes \mathbf{B}.\tag{10}$$

Our second task is to evaluate \mathbf{E} . Since

$$\begin{aligned}\mathbf{E} &= \sqrt{\frac{2}{\pi}} \text{diag}(\text{diag}(\mathbf{C}_z)^{-1/2}) \mathbf{D} \mathbf{C}_x = \sqrt{\frac{2}{\pi}} \frac{\sigma_x^2}{\sqrt{2\sigma_x^2 + 1}} \mathbf{D} \\ &= \frac{\sigma_x^2}{\sqrt{2\sigma_x^2 + 1}} [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U} \quad \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U],\end{aligned}$$

we have

$$\begin{aligned}\mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{E} &= \frac{2}{\pi} \frac{\sigma_x^4}{2\sigma_x^2 + 1} \begin{bmatrix} \mathbf{1}_Q^T \otimes \mathbf{I}_{U \times U} \\ \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U^T \end{bmatrix} \\ &\quad \times (\mathbf{1}_{Q \times Q} \otimes \mathbf{A} + \mathbf{I}_{Q \times Q} \otimes \mathbf{B}) [\mathbf{1}_Q \otimes \mathbf{I}_{U \times U} \quad \mathbf{I}_{Q \times Q} \otimes \mathbf{1}_U] \\ &= \frac{2}{\pi} \frac{\sigma_x^4}{2\sigma_x^2 + 1} Q(Q\mathbf{A} + \mathbf{B}),\end{aligned}$$

where we have used $(\mathbf{X} \otimes \mathbf{Y})(\mathbf{U} \otimes \mathbf{V}) = (\mathbf{X}\mathbf{U}) \otimes (\mathbf{Y}\mathbf{V})$.

Therefore, the value of entry $(1, 1)$ in $\mathbf{E}^T \mathbf{C}_y^{-1} \mathbf{E}$, i.e., the MSE of the user ability parameter estimates, is given by

$$\begin{aligned}\frac{2}{\pi} \frac{\sigma_x^4}{2\sigma_x^2 + 1} Q(a + (Q-1)c) &= \\ \sigma_x^2 \left(1 - \frac{2}{\pi} \frac{\sigma_x^2}{2\sigma_x^2 + 1} \frac{sQ(Q+U-3)+1}{(s(Q-2)+1)(s(Q+U-2)+1)}\right),\end{aligned}$$

where we have used (9), thus completing the proof.

Acknowledgments

C. Studer was supported in part by Xilinx, Inc. and by the US National Science Foundation (NSF) under grants ECCS-1408006, CCF-1535897, CCF-1652065, and CNS-1717559.

References

Albert, J. H. and Chib, S. Bayesian analysis of binary and poly-chotomous response data. *J. Am. Stat. Assoc.*, 88(422):669–679, June 1993.

Bach, F. Self-concordant analysis for logistic regression. *Electron. J. Stat.*, 4:384–414, 2010.

Baker, F. B. and Kim, S. H. *Item Response Theory: Parameter Estimation Techniques*. Marcel Dekker Inc., 2nd edition, 2004.

Bliss, C. The calculation of the dosage-mortality curve. *Ann. Appl. Biol.*, 22(1):134–167, Feb. 1935.

Brillinger, D. *A Festschrift for Erich L. Lehmann*, chapter “A generalized linear model with “Gaussian” regressor variables”, pp. 97–114. Wadsworth Statistics/Probability Series. Chapman & Hall/CRC, 1982.

Brzezińska, J. Latent variable modelling and item response theory analyses in marketing research. *J. University of Szczecin*, 16(2): 163–174, Dec. 2016.

Bunea, F. Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Electron. J. Stat.*, 2: 1153–1194, 2008.

Cappelleri, J. C., Lundy, J. J., and Hays, R. D. Overview of classical test theory and item response theory for the quantitative assessment of items in developing patient-reported outcomes measures. *Clinical Therapeutics*, 36(5):648–662, May 2014.

Carroll, R. J., Ruppert, D., Stefanski, L. A., and Crainiceanu, C. M. *Measurement Error in Nonlinear Models: A Modern Perspective*. CRC press, 2006.

Chang, H. and Ying, Z. Nonlinear sequential designs for logistic item response theory models with applications to computerized adaptive tests. *The Annals of Statistics*, 37(3):1466–1488, Jun. 2009.

Davenport, M. A., Plan, Y., van den Berg, E., and Wootters, M. 1-bit matrix completion. *Inf. Infer.*, 3(3):189–223, Sep. 2014.

Fahrmeir, L. and Kaufmann, H. Consistency and asymptotic normality of the maximum likelihood estimator in generalized linear models. *Ann. Stat.*, 13(1):342–368, Mar. 1985.

Filippi, S., Cappe, O., Garivier, A., and Szepesvári, C. Parametric bandits: The generalized linear case. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 586–594, Dec. 2010.

Goldstein, T., Studer, C., and Baraniuk, R. A field guide to forward-backward splitting with a FASTA implementation. *arXiv preprint: 1411.3406*, Nov. 2014.

Gourieroux, C. and Monfort, A. Asymptotic properties of the maximum likelihood estimator in dichotomous logit models. *J. Econometrics*, 17(1):83–97, Sep. 1981.

Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer, 2010.

Herlocker, J., Konstan, J., Borchers, A., and Riedl, J. An algorithmic framework for performing collaborative filtering. In *Proc. Ann. Int. Conf. Res. Develop. Inf. Retrieval*, pp. 230–237, Aug. 1999.

Hoff, P. D. *A First Course in Bayesian Statistical Methods*. Springer, 2009.

Jin, H. and Ling, C. Using AUC and accuracy in evaluating learning algorithms. *IEEE Trans. Knowl. Data Eng.*, 17(3):299–310, March 2005.

Lan, A. S., Waters, A. E., Studer, C., and Baraniuk, R. G. Sparse factor analysis for learning and content analytics. *J. Mach. Learn. Res.*, 15:1959–2008, June 2014.

Lan, A. S., Goldstein, T., Baraniuk, R. G., and Studer, C. Deal-breaker: A nonlinear latent variable model for educational data. In *Proc. Intl. Conf. Mach. Learn.*, pp. 266–275, June 2016.

Lan, A. S., Chiang, M., and Studer, C. Linearized binary regression. *arXiv preprint: 1802.00430*, Feb. 2018. <https://arxiv.org/abs/1802.00430>.

Linacre, J. M. Understanding Rasch measurement: Estimation methods for Rasch measures. *J. Outcome Meas.*, 3(4):382–405, 1999.

Lord, F. *Applications of Item Response Theory to Practical Testing Problems*. Erlbaum Associates, 1980.

Nocedal, J. and Wright, S. *Numerical Optimization*. Springer, 2006.

Patz, R. J. and Junker, B. W. A straightforward approach to Markov chain Monte Carlo methods for item response models. *J. Educ. Behav. Stat.*, 24(2):146–178, June 1999.

Plan, Y. and Vershynin, R. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Trans. Inf. Th.*, 59(1):482–494, Sep. 2013.

Rasch, G. *Studies in Mathematical Psychology: I. Probabilistic Models for Some Intelligence and Attainment Tests*. Nielsen & Lydiche, 1960.

Ravikumar, P., Wainwright, M., and Lafferty, J. High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *Ann. Stat.*, 38(3):1287–1319, June 2010.

Schellhorn, C. and Sharma, R. Using the rasch model to rank firms by managerial ability. *Managerial Finance*, 39(3):306–319, Mar. 2013.

Thompson, B. *Score reliability: Contemporary thinking on reliability issues*. Sage Publications, 2002.

Thrampoulidis, C., Abbasi, E., and Hassibi, B. Lasso with non-linear measurements is equivalent to one with linear measurements. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 3420–3428, Dec. 2015.

Tsutakawa, R. K. and Johnson, J. C. The effect of uncertainty of item parameter estimation on ability estimates. *Psychometrika*, 55(2):371–390, June 1990.

van der Linden, W. J. and Hambleton, R. K. *Handbook of Modern Item Response Theory*. Springer Science & Business Media, 2013.

Whitehill, J., Wu, T., Bergsma, J., Movellan, J. R., and Ruvolo, P. L. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 2035–2043, Dec. 2009.

Yang, J. S., Hansen, M., and Cai, L. Characterizing sources of uncertainty in item response theory scale scores. *Educ. Psychol. Meas.*, 72(2):264–290, Aug. 2012.

Zhang, J., Xie, M., Song, X., and Lu, T. Investigating the impact of uncertainty about item parameters on ability estimation. *Psychometrika*, 76(1):97–118, Jan. 2011.