Bayesian Penalization for an Explanatory Cognitive Diagnostic Model:

A Covariate DINA with the Lasso Prior

Yoshito Tan,^{1,2} Tetsuro Ito,³ and Kensuke Okada¹

¹Graduate School of Education, the University of Tokyo

²Japan Society for the Promotion of Science, Tokyo, Japan

³Classi Corp.

Author Note

Yoshito Tan <https://orcid.org/0000-0002-0003-445X>

Tetsuro Ito

Kensuke Okada <https://orcid.org/0000-0003-1663-5812>

This work was supported by grants from JSPS KAKENHI Grant No. 18H03612, No. 21H00936, and No. 21J21435, JST PRESTO Grant No. JPMJPR21C3, and a research grant from Classi Corp.

Correspondence concerning this article should be addressed to Yoshito Tan, Graduate School of Education, the University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo, Japan 113-0033. E-mail: yoshitotan001@gmail.com.

A preliminary study of this study was presented at the International Meeting of the Psychometric Society 2020.

Abstract

Diagnostic assessment data obtained from online learning platforms for schools are typically accompanied by student background variables and item responses. To leverage such information for cognitive diagnosis, the present study examines the applicability of the lasso prior for variable selection in a deterministic input, noisy-and-gate (DINA) model with attribute-level explanatory variables. We compared the covariate DINA model with and without the lasso prior using a real-world data analysis and a simulation study. In the real-world data analysis, which used a school-sized sample collected from an online learning platform, we found that the lasso prior selected only relatively large effects without substantially affecting the diagnostic classification and item parameter estimation. In the simulation study, we found that the lasso prior did not degrade the accuracy of the diagnostic classification or parameter estimation. Finally, we discuss the situations in which the lasso prior can be useful with the covariate DINA model.

Keywords: cognitive diagnostic models, explanatory variables, Bayesian lasso, variable selection problem

Bayesian Penalization for an Explanatory Cognitive Diagnostic Model:

A Covariate DINA with the Lasso Prior

Cognitive diagnostic assessments and cognitive diagnostic models (CDMs) are often used to provide students and teachers with feedback on the students' mastery of fine-grained skills (or *attributes*). A typical empirical interest in CDM applications is in identifying which attributes each student has mastered (Rupp et al., 2010). Meanwhile, the relationships between attribute mastery and its explanatory variables, such as students' demographics, background knowledge, and learning records, can also be a target of interest. Identifying such relationships provides teachers with useful information to plan and deliver instructions. To enable such an investigation, Ayers et al. (2013) and Park and Lee (2014) extended the deterministic input, noisy-and-gate (DINA; Junker & Sijtsma, 2001) model, which is one of the best known and interpretable CDMs, to integrate observed explanatory variables. By generalizing this model, Park et al. (2018) conceptually introduced the framework of explanatory cognitive diagnostic models (ECDMs), in which both observed and latent explanatory variables are incorporated into CDMs. Following them, we hereafter refer to the DINA model that is extended to include explanatory variables as the *covariate DINA* model.¹

The covariate DINA appears to be particularly relevant when applied to the data obtained from online learning platforms, which have become common in schools. In addition to item responses for diagnostic assessments, learning platform data typically contain various kinds of background variable data, such as students' performance for other assessments, study logs, study habits, and so forth. The covariate DINA model can be used to examine the relationship between these background variables and attribute mastery. Then, the results can be utilized to provide personalized feedback to assist students' learning. For example, a watch history of a

particular video lesson program can be an explanatory variable that positively affects the mastery of a specific attribute. If this is the case, we can recommend the program for those students who were diagnosed as not having mastered that attribute. Further, teachers can make good use of the revealed relationship to customize their instructions according to the current state of the class.

However, when many candidates of explanatory variables exist, which is often the case for the data in learning platforms, it is a significant challenge to select the explanatory variables that parsimoniously account for the dominant sources of variation in the attribute mastery probabilities in a school-sized sample. Such variable selection for model parsimony is practically desirable when schoolteachers want to know which student background variables they should focus on when designing instructions. While improving instructional plans based on the variables with large effect sizes is likely to positively affect the students' achievement, plan improvements based on variables with trivial effects may not be worth the cost. Hence, for school-sized samples, it is desirable to apply a method that selects only the variables with large effects while excluding those with negligible ones. The lasso (least absolute shrinkage and selection operator) is one of the most commonly adopted methods for such a purpose (Tibshirani, 1996), and it has been studied for variable selection in various types of regression analyses (e.g., Tutz & Schauburger, 2015). The lasso shrinks small regression coefficients toward zero by adding an L_1 -norm penalty term to the likelihood function. Still, it is challenging to find the optimal value of the parameter of the penalty term. For this problem, a Bayesian hierarchical approach allows for automatic one-step estimation of the penalty parameter simultaneously with other parameters (Park & Casella, 2008). The penalization with the fully Bayesian lasso is achieved by setting the double-exponential priors on the regression coefficients. In addition to the parsimony of the model, it is desirable that the regression coefficients be stably estimated. In the literature, it is

suggested that estimation results using the double-exponential priors provided stabler results than other penalization methods in relatively small samples (e.g. Brandt et al., 2018; Chen et al., 2021). The Bayesian approach also allows for a coherent probabilistic treatment of uncertainty by conditioning the observed data.

While the Bayesian lasso provides a promising solution to the variable selection problem, this technique has not been examined in the ECDM literature. Hence, it is essential to investigate how the lasso prior functions when it is applied to the covariate DINA. If the penalization degrades the attribute classification accuracy, which is usually vital in CDMs, the application of the lasso may not be practical. The same is true in the case that penalization deteriorates the accuracy of item parameter estimation. Bayesian inference combines the information in the data with the prior distribution based on the Bayes rule. Because the double-exponential prior is an informative prior, the information from the shrinkage priors can dominate the information from the data when the sample size is not very large. Hence, the use of such shrinkage priors may affect the marginal posteriors of other parameters of interest, such as attribute mastery and item parameters. If this is the case, the lasso prior applied to the covariate DINA model may lead to an inaccurate diagnosis. Furthermore, if the covariate DINA model with the lasso prior leads to poorer performance in estimating regression coefficients than an ordinary model, the lasso approach may not be recommended. Hence, it is necessary to comparatively investigate the parameter recovery performances between the models with and without the lasso prior in a school-sized sample. In other words, if the shrinkage leads to an accurate estimation with no major changes to the diagnostic classification and the item parameter estimation, the lasso approach can be useful for the covariate DINA. With these research objectives, we examined the utility of the lasso prior when applied to the covariate DINA with

attribute-level observed explanatory variables.

In the following sections, we first introduce the covariate DINA model with attribute-level observed explanatory variables and the priors considered in this study. Next, we present an empirical application with a school-sized sample collected through a cloud-based learning platform to investigate the differences in diagnostic classification, item parameter estimation, and variable selection between the models with and without the lasso prior. Then, we present a simulation study in which we compared the accuracy of parameter recovery between the models with and without the lasso prior under varying experimental conditions. Finally, we summarize the results and discuss when and how the lasso prior can be used for the covariate DINA model.

Models

The Covariate DINA Model with Attribute-Level Observed Explanatory Variables

Various forms of the covariate DINA have been developed, with their main difference lying in where the explanatory variables are located and whether the explanatory variables include latent variables (see Park & Lee, 2019). Because a key estimand in CDM applications is the attribute mastery probability, we focus on the attribute-level explanatory variables. In addition, for simplicity, we only consider observed explanatory variables in the present study.

The covariate DINA consists of two major model parts: item response probability and attribute mastery probability. Before we present the model, we introduce our notations. The indices used in this study are as follows: i ($= 1, \dots, N$) for students, j ($= 1, \dots, J$) for items in an assessment, k ($= 1, \dots, K$) for the attributes to be diagnosed, and m ($= 1, \dots, M$) for explanatory variables. A binary-valued element α_{ik} represents the state of attribute mastery, which equals 1 when student i has mastered attribute k . This is an element of the $N \times K$

attribute mastery profile matrix $\mathbf{A}$. A binary-valued element q_{jk} indicates whether item j measures attribute k . This is an element of the $J \times K$ Q-matrix (Tatsuoka, 1990). The Q-matrix is typically specified with the involvement of domain experts.

The item response probability of the covariate DINA is the same as that of the ordinary DINA model. Specifically, the item response probability is represented as

$$P(y_{ij} = 1 | \eta_{ij}, s_j, g_j) = g_j^{(1-\eta_{ij})} (1 - s_j)^{\eta_{ij}}, \quad (1)$$

$$\eta_{ij} = \prod_{k=1}^K \alpha_{ik}^{q_{jk}}, \quad (2)$$

where $y_{ij} \in 0,1$ denotes the observed response, and $\eta_{ij} \in 0,1$ denotes the ideal response, which represents whether student i has mastered all the required attributes for item j . Here, y_{ij} is an element of the $N \times J$ observed response matrix $\mathbf{Y}$, and η_{ij} is an element of the $N \times J$ ideal response matrix $\mathbf{H}$. $g_j = P(y_{ij} = 0 | \eta_{ij} = 1)$ is the guessing parameter and $s_j = P(y_{ij} = 0 | \eta_{ij} = 0)$ is the slip parameter for item j . Using these item parameters, the item discrimination index is represented as $d_j = 1 - s_j - g_j$ (Rupp et al., 2010). Moreover, the attribute mastery probability of the covariate DINA model is represented as

$$P(\alpha_{ik} = 1 | \zeta_i, \mathbf{x}_i, \delta_k, \boldsymbol{\gamma}_k) = \frac{1}{1 + \exp[-(\zeta_i - (\delta_k + \mathbf{x}_i^T \boldsymbol{\gamma}_k))]}, \quad (3)$$

where $\zeta_i \in \boldsymbol{\zeta} = [\zeta_1, \dots, \zeta_i, \dots, \zeta_N]'$ is a higher-order latent trait parameter of student i , $\delta_k \in \boldsymbol{\delta} = [\delta_1, \dots, \delta_k, \dots, \delta_K]'$ is an attribute difficulty parameter, and $\boldsymbol{\gamma}_k = [\gamma_{1k}, \dots, \gamma_{mk}, \dots, \gamma_{Mk}]'$ is a vector whose elements are regression coefficient parameters. The regression coefficient parameter vector is a column vector of $M \times K$ matrix $\boldsymbol{\Gamma}$. Furthermore, $\mathbf{x}_i = [x_{i1}, \dots, x_{im}, \dots, x_{iM}]'$ indicates the explanatory variable vector for student i . That is, the attribute difficulty shifts by the regression coefficient γ_{mk} for each unit of x_{im} . Therefore, the

probability of student i mastering attribute k varies depending on the explanatory variables.

Priors on the Regression Coefficient Parameters

The regression coefficient parameters of the covariate DINA can be penalized with the lasso prior. Because our purpose is to study whether the lasso prior degrades the performance of the covariate DINA, we compare the results of the models with the lasso prior to an ordinary model without any shrinkage priors. For this purpose, we prepare the latter model without any shrinkage priors. Following Ayers et al. (2013), we set the normal prior distributions with mean zero and variance 10 on the regression coefficient parameters. Specifically, the prior on each parameter is as follows:

$$\gamma_{mk} \sim \text{Normal}(0,10). \quad (4)$$

In contrast, the lasso is provided by placing a double-exponential prior for each of the coefficients of a regression model (Park & Casella, 2008). The double-exponential distribution has a penalty parameter, which can be estimated in a fully Bayesian manner by assuming a hyperprior. For this penalty parameter, following van Erp et al. (2019), we set a half-Cauchy distribution. To be more precise, in our model, the lasso prior for the regression coefficient is as follows:

$$\begin{aligned} \gamma_{mk} &\sim \text{DoubleExponential}(0, \sigma_k^2 \lambda_k^{-2}), \\ \lambda_k &\sim \text{Cauchy}_{[0,\infty)}(0,1), \end{aligned} \quad (5)$$

where the subscript $[0, \infty)$ for the prior of λ_k indicates the support of the truncated Cauchy distribution. Here, σ_k and λ_k denote the scale and penalty parameters. In the present study, we fixed the squared scale parameter $\sigma_k^2 = 10$ to compare the results between the models with priors (4) and (5). The scale of the double-exponential distribution is determined by the penalty parameter λ_k , with a larger λ_k indicating a smaller scale. When the scale becomes small, the

density of γ_{mk} is sharp at zero.

For the variable selection, we need a criterion to determine which explanatory variables are selected as relevant variables. As suggested by Park and Casella (2008), we used the credible interval (CI) of γ_{mk} to select relevant variables. For this purpose, van Erp et al. (2019) used a simulation study to investigate the optimal CI percentage in terms of the balance between true and false inclusion rates. Then, the optimal CI was used to select explanatory variables in the real-world data analyses. An explanatory variable was selected if the optimal CI did not cover zero, while the explanatory variable was not selected if the optimal CI covered zero. In the present study, we also investigated the optimal CI for each of the two prior models in the simulation study, then used that criterion to select the relevant explanatory variables in the real-world data analysis.

Priors on the Other Parameters

The priors for the other parameters were kept constant between the two models. We refer to the CDM literature (Ayers et al., 2013; Zhan et al., 2019) and set them as

$$s_j \sim \text{Beta}(1,1), \quad (6)$$

$$g_j \sim \text{Beta}_{[0,1-s_j]}(1,1),$$

$$\delta_k \sim \text{Normal}(0,10),$$

$$\zeta_i \sim \text{Normal}(0,1),$$

where the subscript $[0,1-s_j]$ for the prior of g_j indicates the support of the truncated beta distribution, which was used to satisfy the monotonicity restriction on the item parameters. Here, we denote the item parameters $\mathbf{s} = [s_1, \dots, s_j, \dots, s_J]'$ and $\mathbf{g} = [g_1, \dots, g_j, \dots, g_J]'$ as $\boldsymbol{\Omega} = [\mathbf{s}, \mathbf{g}]$. Finally, the joint posterior distribution of the covariate DINA model without the lasso

prior is given as

$$P(\mathbf{A}, \mathbf{\Gamma}, \boldsymbol{\delta}, \boldsymbol{\zeta}, \boldsymbol{\Omega} | \mathbf{Y}, \mathbf{X}) \propto P(\mathbf{Y} | \mathbf{H}, \boldsymbol{\Omega}) P(\mathbf{A} | \mathbf{X}, \mathbf{\Gamma}, \boldsymbol{\delta}, \boldsymbol{\zeta}) P(\boldsymbol{\Omega}) P(\mathbf{\Gamma}) P(\boldsymbol{\delta}) P(\boldsymbol{\zeta}), \quad (7)$$

and that of the model with the lasso prior is given as

$$P(\mathbf{A}, \mathbf{\Gamma}, \boldsymbol{\delta}, \boldsymbol{\zeta}, \boldsymbol{\Omega}, \boldsymbol{\lambda} | \mathbf{Y}, \mathbf{X}) \propto P(\mathbf{Y} | \mathbf{H}, \boldsymbol{\Omega}) P(\mathbf{A} | \mathbf{X}, \mathbf{\Gamma}, \boldsymbol{\delta}, \boldsymbol{\zeta}) P(\boldsymbol{\Omega}) P(\mathbf{\Gamma} | \boldsymbol{\lambda}) P(\boldsymbol{\delta}) P(\boldsymbol{\zeta}) P(\boldsymbol{\lambda}), \quad (8)$$

where $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_k, \dots, \lambda_K]'$ denotes the penalty parameter vector, and $\mathbf{X} =$

$[\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_N]'$ denotes the $N \times M$ explanatory variable matrix. In the following sections, we compare the above two models. In this article, we call these two Bayesian models the normal prior model and the lasso prior model.

We estimated the model parameters using the Markov chain Monte Carlo (MCMC) algorithm with JAGS (Plummer, 2017) via R2jags package (Version 0.6-1; Su & Yajima, 2020) in R (R Core Team, 2020). The R code for the simulation study is available from the Open Science Framework website (https://osf.io/6jw84/?view_only=5d47cefedfb9486192c1d631c83f06e3). [Note: during the review process, the information on the website is anonymized for the purpose of review]

Study I: Real-World Data Analysis

Method

An empirical data analysis was conducted to compare the two models. We analyzed a dataset obtained from Classi, which is a major provider of a cloud-based learning platform to middle and high schools in Japan. Data were collected from April 11, 2015, to May 15, 2018. We analyzed the mathematics assessment data obtained from one high school. This dataset includes the item responses for $J = 80$ items that measure $K = 3$ attributes obtained from $N = 395$ students. The attributes to be measured, which were predetermined by domain experts, were Basic

Calculation, Polynomial, and Positive Number and Negative Number. They were measured by 20, 40, and 20 items, respectively, which exhibit a simple structure. However, a preliminary analysis with the DINA model using the `din` function of the `CDM` package showed that, in terms of the item discrimination index, several items suffer from low-item discrimination. Because such items would lead to an insufficient model fit, we chose five items with the highest discrimination index values for each attribute. Furthermore, one of the chosen items based on the above procedure exhibited $\hat{g}_j = 0.504$, which indicates that students who did not master the required attributes could correctly answer the item more than half of the time. Because this guessing parameter value is too large to be practical, we replaced the item with the sixth item. As a result, we used $J = 15$ items with discrimination index values larger than 0.4 for our analysis.

We standardized the explanatory variables, as is typically done in penalization methods, to adjust the scales of the priors. We have 10 background variables as candidates of explanatory variables. The dataset related to all the background variables had been collected before the mathematics assessment. The 10 variables can be divided into three types: the frequencies of item solving, video watching, and learning recording. The “frequency of item solving” corresponds to the number of times each student has solved the items in the learning platform. This type contains three variables ($m = 1, 2, 3$) for different periods of a day: morning (from 4 a.m. to 12 p.m.), afternoon (from 12 p.m. to 8 p.m.), and night (from 8 p.m. to 4 a.m.). The “frequency of video watching” corresponds to the number of times each student watched the learning videos. This type contains six variables ($m = 4, 5, \dots, 9$) representing the different ways the student came to watch the videos: teachers’ assignment, the platform’s recommendation system, the platform’s learning map, the student’s search, the student’s self-study list, and others. The third type, which contained one variable ($m = 10$), was the number of times each student

recorded their learning history. The correlation coefficients between the 10 explanatory variable candidates ranged between $[-0.060, 0.275]$.

We estimated the model parameters using the MCMC algorithm. Eight MCMC chains were run on each model. We used the Gelman-Rubin diagnostic statistic ($\hat{R}$) for a convergence diagnostic and an effective number of independent draws ($\hat{n}_{\text{eff}}$) for a precision diagnostic of the MCMC simulations. We set the thresholds to be $\hat{R} < 1.1$ and $\hat{n}_{\text{eff}} > 20n_{\text{chain}}$, where n_{chain} represents the number of MCMC chains (Gelman et al., 2014). Here, because $n_{\text{chain}} = 8$, the threshold was $\hat{n}_{\text{eff}} > 160$. Furthermore, we checked the trace plots of the MCMC draws. Our preliminary investigation showed that the normal prior model required more iterations and thinning intervals than the lasso prior model to reach the thresholds for the covariate DINA model parameters. Consequently, the number of warm-up iterations per chain, the respective number of iterations per chain after the warm-up, and the number of thinning intervals were 1,000, 25,000, and 5 for the lasso prior model and 1,000, 50,000, and 10 for the normal prior model. Hence, the total number of the MCMC samples used for estimation was 40,000 for each model. Regarding the variable selection, we selected the explanatory variable if the 90% CI of its coefficient did not include zero. This size of the optimal CI was determined based on the simulation study described in the Simulation Study section.

Results

To assess whether the models' predictions adequately capture the characteristics in the empirical data, we conducted posterior predictive model checking (Rubin, 1984) for each model. For this purpose, we used the sum of the squared Pearson residual,

Table 1*Profile-Wise Attribute Classification Rate (Mixing Rate)*

Profile	Attribute			Classification rate	
	1	2	3	Normal	Lasso
1	0	0	0	0.073	0.073
2	1	0	0	0.076	0.076
3	0	1	0	0.003	0.003
4	1	1	0	0.167	0.167
5	0	0	1	0.063	0.063
6	1	0	1	0.149	0.149
7	0	1	1	0.005	0.003
8	1	1	1	0.463	0.466

$$D(\mathbf{Y}|\mathbf{A}) = \sum_{i=1}^N \sum_{j=1}^J \left(\frac{y_{ij} - P(y_{ij} = 1)}{\sqrt{P(y_{ij} = 1)(1 - P(y_{ij} = 1))}} \right)^2, \quad (9)$$

as the discrepancy measure, where $P(y_{ij} = 1)$ represents $P(y_{ij} = 1 | \eta_{ij}, s_j, g_j)$ in Equation 1 (Zhan et al., 2019). The resultant posterior predictive p -values were 0.521 for the lasso prior model and 0.521 for the normal prior model. Because these values were close to .5, we judged these models as adequate for the present dataset.

As shown in Table 1, the profile-wise attribute classification rates, also known as the mixing proportions, were overall similar across the two prior models: the largest difference, which was found in profile eight, was about 0.003. Specifically, only one student was classified differently in profiles seven and eight. Moreover, the marginal posterior distributions of the item parameters were also similar across the prior models. The resultant differences for item

parameters are shown in Supplemental Figures 1 and 2. Because of the small differences between the two models, we concluded that the lasso prior did not substantially alter the attribute classification and item parameter estimations.

Meanwhile, most of the marginal posteriors of the regression coefficient parameters were substantially different across the two models, even though some were more similar than others. Figure 1 shows the expected a posteriori (EAP) values and 90% CIs of the regression coefficient parameter $\boldsymbol{\Gamma}$ in the two models. According to the normal prior model, three, five, and two explanatory variables were respectively selected for attributes 1–3. However, according to the lasso prior model, only the frequency of item solving in the morning ($m = 1$) was selected for the first attribute (Basic Calculation), three variables were selected for the second attribute (Polynomial), and no variables were selected for the third attribute (Positive Number and Negative Number). That is, even though some variables indicated some extent of effects in the normal prior model, they were not large enough to be selected by the lasso penalization. Furthermore, the 90% CIs in the lasso prior model were narrower than those in the normal prior model. In particular, the 90% CI of the frequency of video watching assigned by teachers ($m = 4$) covered a much narrower range $[-1.129, 0.247]$ in the lasso prior model, compared to the range of the 90% CI $[-6.529, 0.117]$ in the normal prior model. The results of the lasso prior model are easier to interpret than those of the normal prior model because of the parsimony and reduction of uncertainty.

According to the results of the lasso prior model, the students who solved more items in the learning platform in the morning ($m = 1$) were more likely to master the first two attributes. Moreover, the students who watched more videos that were suggested by the recommendation system ($m = 5$) were also more likely to master the second attribute; in contrast, it was relatively

Figure 1

Expected a Posteriori (EAP) Values and the 90% Credible Intervals (CIs) of the Regression

Coefficient Γ for the Explanatory Variables in the two Compared Models

Note. The three panels correspond to the three attributes. The x-axis corresponds to the explanatory variables. For each explanatory variable, the black dot with a gray bar indicates the EAP and 90% CI for the normal prior model, and the white dot with a black bar indicates the same quantities for the lasso prior model. The box-shaped dots and the plus-shaped dots at the bottom of each panel indicate that the explanatory variable is selected based on the normal prior model and the lasso prior model, respectively. The variable selection was based on the optimal CI criterion determined from the simulation study shown in the Simulation Study section.

difficult for those who recorded relatively many learning histories ($m = 10$) to master the attribute. Although these relationships between attributes and explanatory variables are not necessarily causal and are limited to the present sample, this kind of knowledge, when accumulated, would help teachers provide the instructions and practice that each student needs.

We found that the lasso prior worked properly because the regression coefficients were shrunk toward zero without altering the results of the diagnostic classification and the item parameter estimation. However, if applying the shrinkage priors leads to considerable ill recovery of the regression coefficient parameters with relatively large effects and poor classification accuracy, it will be dangerous to use the lasso prior in the covariate DINA model with attribute-level explanatory variables. In the next section, we check the recovery of the regression coefficient parameters $\boldsymbol{\Gamma}$ and the classification accuracy in some realistic situations.

Study II: Simulation Study

Method

A simulation study was conducted to compare the parameter recovery between the two prior models. In addition to the two prior conditions, we considered two Q-matrix conditions and three multicollinearity conditions. Regarding the Q-matrix, one of the conditions employed the same structure as the Q-matrix in the previous real-world data analysis section. Hence, the size of this Q-matrix is $J = 15$ items $\times K = 3$ attributes. We call this the “simple Q” condition. The other condition employed a more complex structure with fewer items that can be possible in practice. This Q-matrix contains $J = 10$ items and $K = 3$ attributes. Seven items measured all possible attribute combinations except the pattern of (0,0,0), and the other three items measured

each different attribute, forming an identity matrix. We call this the “complex Q” condition.

Supplemental Table 1 shows the complex Q-matrix. In both Q-matrix conditions, each attribute was measured in the assessment five times. Regarding multicollinearity, we had no, middle, and strong multicollinearity conditions by setting the correlations between each pair of explanatory variables to $\rho = 0, 0.4$ and 0.8 . From the above, the total number of conditions was 12.

We used the following simulation settings across all conditions, basically following the real-world data analysis reported in the Real-World Data Analysis section. The sample size was $N = 400$. We generated both item parameters s_j and g_j from $\text{Uniform}(0,0.35)$. True values of the higher-order latent trait parameter ζ_i were generated from $\text{Normal}(0,1)$. The regression coefficient parameters were set to $\gamma_{m1} = (3, 2, 1.5, 1, 0, 0, 0, 0)'$, $\gamma_{m2} = (-3, -2, -1.5, -1, 0, 0, 0, 0)'$, and $\gamma_{m3} = (0.8, 0.5, 0.3, 0, 0, 0, 0, 0)'$. The attribute difficulty parameter δ_k was fixed at -1 . Regarding the explanatory variables, $\mathbf{X}$, we first sampled eight variables from a multivariate normal distribution with mean $\mathbf{0}$ and a covariance matrix that has all variances equal to one and all correlations equal to the value determined by the multicollinearity conditions (i.e., 0, 0.4, or 0.8). We then standardized the explanatory variables. We generated $R = 50$ datasets from the normal prior model and fit all the models to each dataset. To estimate the model parameters, four MCMC chains were run. The number of MCMC iterations per chain was 25,000, 5,000 of which were discarded as warm-up. When the convergence criterion of $\hat{R} < 1.1$ was not met, we replaced the dataset with a newly generated one.

To evaluate the classification accuracy of students' attribute mastery states, we calculated the attribute-wise classification accuracy (ACA) rate and the profile-wise classification accuracy (PCA) rate (Liu et al., 2017). They are defined as follows:

$$ACA = \frac{1}{R} \sum_{r=1}^R \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K I\{\hat{\alpha}_{ikr} = \alpha_{ikr}\}, \quad (10)$$

$$PCA = \frac{1}{R} \sum_{r=1}^R \frac{1}{N} \sum_{i=1}^N I\{\hat{\alpha}_{ir} = \alpha_{ir}\}, \quad (11)$$

where $\hat{\alpha}_{ikr} \in \hat{\alpha}_{ir} = [\hat{\alpha}_{i1r}, \dots, \hat{\alpha}_{ikr}, \dots, \hat{\alpha}_{iKr}]$ indicates an estimated attribute mastery state of attribute k by student i in replication r , and $I\{\cdot\}$ denotes an indicator variable that equals one if the estimated and the true attribute mastery states coincide, and zero otherwise. To evaluate the parameter recovery of the regression coefficients, we calculated the bias and root mean square error (RMSE):

$$\text{bias}_{mk} = \frac{1}{R} \sum_{r=1}^R (\hat{\gamma}_{mkr} - \gamma_{mkr}), \quad (12)$$

$$\text{RMSE}_{mk} = \sqrt{\frac{1}{R} \sum_{r=1}^R (\hat{\gamma}_{mkr} - \gamma_{mkr})^2}, \quad (13)$$

where $\hat{\gamma}_{mkr}$ represents the EAP estimate of the regression coefficient parameter for attribute k on the m th explanatory variable in replication r .

Additionally, for the variable selection in the real-world data analysis, we investigated the optimal CI for each of the two prior models following the procedure of van Erp et al. (2019). Specifically, from 10% to 100% CIs with an increment of 10, we calculated the distance measure:

$$\text{distance} = \sqrt{(1 - \text{correct inclusion rate})^2 + (\text{false inclusion rate})^2}, \quad (14)$$

where *correct inclusion rate* represents the rate of CIs not including zero when the true values were not zero, and *false inclusion rate* represents the rate of CIs not including zero when the true values were zero. For the optimal CI in the real-world data analysis, we selected the interval that

Table 2

Attribute-Wise Classification Accuracy Rate (ACA) and Profile-Wise Classification Accuracy Rate (PCA)

Q-matrix	Correlation	ACA		PCA	
		Normal	Lasso	Normal	Lasso
simple	0	0.979	0.980	0.940	0.940
	0.4	0.982	0.983	0.948	0.950
	0.8	0.985	0.986	0.957	0.957
complex	0	0.958	0.959	0.885	0.887
	0.4	0.964	0.965	0.900	0.902
	0.8	0.969	0.969	0.913	0.913

resulted in the lowest distance under the condition of the simple structure Q-matrix and the correlation coefficient $\rho = 0$ in each prior model because the correlation coefficients between the candidate explanatory variables ranged between $[-0.060, 0.275]$ in our dataset.

Results

Table 2 shows the values of the ACA and PCA. Both prior model conditions had almost the same classification accuracy. From this result, it is expected that the attribute classification in the real-world data analysis in the previous section was not adversely affected by the use of lasso priors. Figures 2 and 3 show the bias and RMSE of the regression coefficient parameter γ_{mk} . When the absolute values of the true values were relatively large (i.e., explanatory variables $m = 1, 2, 3, 4$ for attributes $k = 1, 2$), the lasso prior model had smaller biases than the normal prior model in all conditions. When the absolute values of the true values were relatively small

(i.e., explanatory variables $m = 1,2,3$ for attribute $k = 3$) or zero (i.e., explanatory variables $m = 5,6,7,8$ for attributes $k = 1,2$, and explanatory variables $m = 4,5,6,7,8$ for attribute $k = 3$), both prior conditions mostly resulted in close absolute values of biases. However, the lasso prior model tended to have negative biases in the regression coefficients with small true values; this is because such regression coefficients were estimated as almost zero in most of the replications. The lasso prior model had smaller or similar RMSEs to the normal prior model in all conditions and attributes. We also checked the parameter recovery of item parameters. In all item parameters, the recovery was similar for both prior conditions. Supplemental Figures 3 and 4 show the results.

As for the optimal CI in the real-world data analysis, we found that the 90% CI resulted in the lowest distance for all conditions, except that the 80% CI was selected as the optimal CI for the conditions of the lasso prior and the correlation coefficient $\rho = 0.8$. Based on the abovementioned criterion, we selected the 90% CI as the optimal CI to use for the variable selection in the real-world data analysis.

Discussion

In the present paper, we have proposed and evaluated the lasso prior for the covariate DINA model as a solution to the variable selection problem that exists with a school-sized sample. In the real-world data analysis, we found that the lasso prior shrank small effects toward zero without substantially affecting the diagnostic classification and item parameter estimation. In the simulation study, we found that the lasso prior did not degrade the accuracy of the classification and parameter estimation. Consequently, we can conclude that the lasso prior has desirable features for the variable selection problem of the DINA model with attribute-level

Figure 2

Scatter Plots Showing the Bias of the Regression Coefficient Parameters

Note. Each panel corresponds to each simulation condition. The black, gray, and white dots indicate attributes 1, 2, and 3, respectively. The triangular and square dots indicate the lasso prior model and the normal prior model, respectively.

Figure 3

Scatter Plots Showing the Root Mean Square Error (RMSE) of the Regression Coefficient

Parameters

Note. Each panel corresponds to each simulation condition. The black, gray, and white dots indicate attributes 1, 2, and 3, respectively. The triangular and square dots indicate the lasso prior model and the normal prior model, respectively.

observed explanatory variables in a school-sized sample. These results are further explained below.

Our first objective was to examine whether the shrinkage introduced by the lasso prior degrades the diagnostic classification and item parameter estimation in a school-sized sample. The results from both the real-world study and the simulation study found no major changes in the classification and estimation accuracies between the normal and lasso prior models. Hence, we conclude that the covariate DINA model with the lasso prior may be used for variable selection without concern for the deterioration of the classification and estimation.

Our second objective was to investigate whether the lasso prior in a school-sized sample leads to poorer performance of the regression coefficients parameter estimation than an ordinary model. The results of the simulation study revealed that the estimated regression coefficients in the lasso prior model were similar to or more precise than those without it. Especially when the true effects were large and the explanatory variables had low or middle correlations, the lasso prior model had less biased and more stable estimates than the normal prior model. Even when the explanatory variables were highly correlated, the lasso prior model had more stable estimates.

In the real-world data analysis, the normal prior model had an estimation challenge: the normal prior model needed more iterations and thinning intervals in the MCMC method than the lasso prior model. This could be due to the number of regression coefficient parameters: the same model did not have such an estimation challenge when fewer explanatory variables were considered in the simulation study. The proposed Bayesian lasso approach may not only provide practical variable selection but also help enable a stable MCMC estimation. Still, the relationship between the number of explanatory variables and the stableness of the MCMC estimation

requires more thorough investigation.

Finally, we examined the performance of the lasso prior in some limited situations. Research questions left for future studies include whether our results generalize to the more generalized class of ECDMs and how other shrinkage priors work in the covariate DINA model. Furthermore, more empirical studies that apply the proposed approach to other populations, assessment domains, and other Q-matrix structures will provide us with an assessment of the applicability and efficacy of our proposal in practice.

Declaration of Conflicting Interests

The authors declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: KO and YT have received a research grant from Classi Corp. TI is employed by Classi Corp.

Funding

This work was supported by grants from JSPS KAKENHI Grant No. 18H03612, No. 21H00936, and No. 21J21435, JST PRESTO Grant No. JPMJPR21C3, and a research grant from Classi Corp.

References

- Ayers, E., Rabe-Hesketh, S., & Nugent, R. (2013). Incorporating student covariates in cognitive diagnosis models. *Journal of Classification*, 30, 195-224. <https://doi.org/10.1007/s00357-013-9130-y>.
- Brandt, H., Cambria, J., & Kelava, A. (2018). An adaptive Bayesian lasso approach with spike-and-slab priors to identify multiple linear and nonlinear effects in structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, 25(6), 946-960. <https://doi.org/10.1080/10705511.2018.1474114>.
- Chen, S. M., Bauer, D. J., Belzak, W. M., & Brandt, H. (2021). Advantages of spike and slab priors for detecting differential item functioning relative to other Bayesian regularizing priors and frequentist lasso. *Structural Equation Modeling: A Multidisciplinary Journal*, 29(1), 122-139. <https://doi.org/10.1080/10705511.2021.1948335>.
- DeCarlo, L. T. (2011). On the analysis of fraction subtraction data: The DINA model, classification, latent class sizes, and the Q-matrix. *Applied Psychological Measurement*, 35, 8-26. <https://doi.org/10.1177/0146621610377081>.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). CRC Press. <https://doi.org/10.1201/b16018>.
- Junker, B. W., & Sijtsma, K. (2001). Cognitive assessment models with few assumptions, and connections with nonparametric item response theory. *Applied Psychological Measurement*, 25, 258–272. <https://doi.org/10.1177/01466210122032064>.
- Liu, R., Huggins-Manley, A. C., & Bradshaw, L. (2017). The impact of Q-matrix designs on diagnostic classification accuracy in the presence of attribute hierarchies. *Educational and Psychological Measurement*, 77(2), 220-240. <https://doi.org/10.1177/0013164416645636>.

- Park, T., & Casella, G. (2008). The Bayesian lasso. *Journal of the American Statistical Association*, 482, 681–686. <https://doi.org/10.1198/016214508000000337>.
- Park, Y. S., & Lee, Y.-S. (2014). An extension of the DINA model using covariates: Examining factors affecting response probability and latent classification. *Applied Psychological Measurement*, 38, 376-390. <https://doi.org/10.1177/0146621617738012>.
- Park, Y. & Lee, Y. (2019). Explanatory cognitive diagnostic models. In von Davier, M. & Lee, Y. (Eds.) *Handbook of diagnostic classification models: Models and model extensions, applications, software packages*. (pp. 207-222) Springer. <https://doi.org/10.1007/978-3-030-05584-4>.
- Park, Y., Xing, K., & Lee, Y. (2018). Explanatory cognitive diagnostic models: Incorporating latent and observed predictors. *Applied Psychological Measurement*, 42(5), 376-392. <https://doi.org/10.1177/0146621617738012>.
- Plummer, M. (2017). *JAGS Version 4.3.0 user manual*. Lyon, France. Retrieved from <http://sourceforge.net/projects/mcmc-jags/>
- R Core Team. (2020). *R: A language and environment for statistical computing*. Vienna, Austria: The R Foundation for Statistical Computing.
- Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*, 12(4), 1151-1172. <https://doi.org/10.1214/aos/1176346785>.
- Rupp, A. A., Templin, J., & Henson, R. A. (2010). *Diagnostic measurement. Theory, methods, and applications*. Guilford Press.
- Su, Y. S., & Yajima, M. (2020). *R2jags: Using R to run “JAGS”* (R package version 0.5-7). Retrieved from <http://cran.r-project.org/package=R2jags>.

- Tatsuoka, K. K. (1990). Toward an integration of item-response theory and cognitive error diagnosis. In N. Frederiksen, R. Glaser, A. Lesgold, & M. G. Shafto (Eds.), *Diagnostic monitoring of skill and knowledge acquisition* (p. 453–488). Lawrence Erlbaum Associates, Inc.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B*, 58, 267–288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>.
- Tutz, G. & Schauberger, G. (2015). A penalty approach to differential item functioning in Rasch models. *Psychometrika*, 80(1), 21-43. <https://doi.org/10.1007/s11336-013-9377-6>.
- van Erp, S., Oberski, D. L., & Mulder, J. (2019). Shrinkage priors for Bayesian penalized regression. *Journal of Mathematical Psychology*, 89, 31-50. <https://doi.org/10.1016/j.jmp.2018.12.004>.
- Zhan, P., Jiao, H., Man, K., & Wang, L. (2019). Using JAGS for Bayesian cognitive diagnosis modeling: A tutorial. *Journal of Educational and Behavioral Statistics*, 44, 4, 473-503. <https://doi.org/10.3102/1076998619826040>.

Footnotes

¹ Park et al. (2018) actually called the model they studied the *covariate RDINA* model because their formulation uses the reparameterized version of the DINA model, known as the RDINA model (DeCarlo, 2011). The DINA and RDINA models are mathematically equivalent. Because we use an original DINA-model formulation in our study, we refer to our model as the *covariate DINA* model.

Summary of Supplementals

Supplemental Figure 1: *Marginal Posterior Distributions of Guessing Parameters*

Supplemental Figure 2: *Marginal Posterior Distributions of Slip Parameters*

Supplemental Figure 3: *Scatterplots of Bias of Guessing and Slip Parameters for Items*

Supplemental Figure 4: *Scatterplots of Root Mean Square Error (RMSE) of Guessing and Slip Parameters for Items*

Supplemental Table 1: *Complex Q-matrix*

Figure 1*Marginal Posterior Distributions of Guessing Parameters*

Figure 2*Marginal Posterior Distributions of Slip Parameters*

Figure 3*Scatterplots of Bias of Guessing and Slip Parameters for Items*

Figure 4

Scatterplots of Root Mean Square Error (RMSE) of Guessing and Slip Parameters for Items

Table 1*Complex Q-matrix*

Item	Attribute		
	1	2	3
1	1	0	0
2	0	1	0
3	0	0	1
4	1	0	0
5	0	1	0
6	0	0	1
7	1	1	0
8	1	0	1
9	0	1	1
10	1	1	1