A MULTILEVEL LOGISTIC HIDDEN MARKOV MODEL FOR LEARNING UNDER COGNITIVE DIAGNOSIS

BY

SUSU ZHANG

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Psychology
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2018

Urbana, Illinois

Master's Committee:

Professor Hua-Hua Chang, Chair
Professor Carolyn J. Anderson

Abstract

Students who wish to learn a specific skill have increasing access to a growing number of online courses and open-source educational repositories of instructional tools, including videos, slides, and exercises. Navigating these tools is time consuming and the search itself can hinder the learning of the skill. Educators are hence interested in aiding students by selecting the optimal content sequence for individual learners, specifically which skill one should learn next and which material one should use to study. Such adaptive selection would rely on preknowledge of how the learners' and the instructional tools' characteristics jointly affect the probability of acquiring a certain skill. Building upon previous research on Latent Transition Analysis and Learning Trajectories, we propose a multilevel logistic hidden Markov model for learning based on cognitive diagnosis models, where the probability that a learner acquires the target skill depends not only on the general difficulty of the skill and the learner's mastery of other skills in the curriculum, but also on the effectiveness of the particular learning tool and the its interaction with mastery of other skills, captured by random slopes and intercepts for each learning tool. A Bayesian modeling framework and an MCMC algorithm for parameter estimation are proposed and evaluated using a simulation study.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

With the increasing popularity of online and open source education, such as massive online open courses (MOOCs), we start to see a plethora of online instructional tools targeting the learning of the same skill. Platforms such as OpenEd align instructional resources such as videos, games, and assessments to fine-grained learning objectives, such as the Common Core State Standards. Intelligent tutoring systems (Vanlehn, 2006) usually provide learners with a sequence of tasks for learning a target skill, with interactive feedback to aid problem solving and concept acquisition. A search on OpenEd for instructional tools targeting the learning of "Grade 7 Mathematics: Expressions and Equations" returned over 50 available videos, games, assessments, and so on.

With a massive amount of educational resources available online to the broad audience, people from different regions and backgrounds all receive the opportunity to take the same courses. However, due to the sheer volume of available materials, it is practically unfeasible for learners to navigate through all available materials for a course. A lot of MOOCs have reported having high attrition rate of registered students, usually with less than 10% of students who registered and completed the course. Using a survey for students who dropped out of an online course, Gütl, Rizzardini, Chang, and Morales (2014) analyzed the potential reasons for the low retention rate of MOOCs. A majority of students suggested that insufficient time, unchallenging tasks, or overly difficult contents were part of the reasons for their drop-out. To maximize the students' interest and the learning efficiency with a limited amount of time, educators are interested in evaluating the effectiveness of different materials for a student and tailoring the course content to the student's level based on his or her needs.

Such individualized learning is often achieved through an adaptive recommendation strategy (Chen, Li, Liu, & Ying, 2017), where materials are sequentially selected and recommended to the student based on currently available information. Chen, Li, et al. (2017) suggested that a good recommendation strategy should utilize the information about both the learner and the instructional tool to maximize the gain (e.g., learning efficiency), thus in the present paper, we aim at developing a learning model that simultaneously estimates the student's progress over time and evaluates the efficiency of different learning materials.

At each stage of the learning process, a student's ability at that time point cannot be directly observed. Instead, it is often measured by test items following some kind of item response model, enabling us to measure the student's latent ability. Compared to item response theory models assuming one or a few continuous latent traits, cognitive diagnosis models assume students' mastery of different skills is discrete (e.g., dichotomous), hence allowing the computationally efficient simultaneous measurement of more fine-grained skills in the curriculum. Therefore, in the current paper, we consider assessing the students' learning progress under the cognitive diagnostic modeling framework.

In the next sections of the paper, we will briefly introduce cognitive diagnosis models and hidden Markov models based on cognitive diagnosis models, provide an introduction of our current model, present a Bayesian estimation framework as well as the estimation algorithms, show the results from a simulation study on the proposed model, and, lastly, discuss the potential implications of the current model and some future directions.

1.1 Learning Models Based on Cognitive Diagnosis

Cognitive diagnosis models (CDMs), or Diagnostic Classification Models (DCMs; Rupp, Templin, & Henson, 2010), are restrictive latent class models measuring test takers' mastery of various skills based on their responses to test questions/items. These models are designed to identify students' strengths and weaknesses in learning, providing guidance for personalized, targeted remedy and support. Under a CDM, a test takers' underlying mastery status on a set of skills is usually denoted by a dichotomous vector $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_K]'$, with $\alpha_k = 1$ indicating mastery on skill k and $\alpha_k = 0$ indicating non-mastery. The probability of correct response on an item j depends on the student's mastery status on attributes required by the item, captured by a $J \times K$ Q-matrix, with $q_{jk} = 1$ if successful answering of item j depends on the student's mastery of skill k, and 0 otherwise. How the probability of a correct response is influenced by the mastery of required attributes depends on the specific CDM. Many CDMs have been previously proposed (e.g., Junker & Sijtsma, 2001; de la Torre, 2011; Henson, Templin, & Willse, 2009; von Davier, 2008). One of the simplest and most commonly used CDM is the deterministic input, noisy-"and"-gate (DINA; e.g., Macready & Dayton, 1977; Junker & Sijtsma, 2001) model, under which, the probability of a correct response is given by

$$P(X_{ij} = 1 \mid \boldsymbol{\alpha}_i, s_j, g_j, \boldsymbol{q}_j) = (1 - s_j)^{\eta_{ij}} g_j^{1-\eta_{ij}}, \tag{1.1}$$

where $\eta_{ij} = \prod_{k=1}^{K} \alpha_{ik}^{q_{jk}}$ is the ideal response, indicating whether subject i possesses all required skills to answer item j correctly, and s_j and g_j are the slipping and guessing parameters. Intuitively, the DINA model describes the case where a subject needs to have mastered all requisite skills of an item to be able to answer the item correctly with high probability $(1 - s_j)$. Missing any of the item's requisite skills would result in a probability of a correct response of g_j instead.

Whereas traditional research on CDMs focused on the assessment of the students' mastery at a single time point, there is increasing interest in assessing the students' progression of attribute mastery over time. Li, Cohen, Bottge, and Templin (2015) combined latent transition analysis (LTA; Langeheine, 1988; Collins & Wugalter, 1992) with the DINA model to estimate the students' mastery change in a longitudinal setting. At each wave, a students has a probability of transitioning from non-mastery to mastery, or from mastery to non-mastery, on each skill. The transitions on different skills are assumed to be independent. Using the estimated transition probabilities between mastery and non-mastery on each skill between waves, they compared the effectiveness of two different learning interventions.

Chen, Culpepper, Wang, and Douglas (2017) generalized Li et al.'s (2015) model to the case where attribute transitions are not necessarily independent — that is, instead of modeling the attribute-wise transitions and multiplying them to get the pattern-wise transition probabilities, they directly modelled the transition probabilities between different skill patterns. Different models for learning, including the most general unrestricted model and the monotonic model (i.e., where the probability of transitioning from mastery to non-mastery is 0), were introduced, and the cardinalities of the sets of all possible trajectories were derived.

Wang, Yang, Culpepper, and Douglas (2016) proposed the Hidden Markov Diagnostic Classification Model (HMDCM) with higher-order covariates affecting the learning outcome. Given the attribute pattern of subject i at time t, $\boldsymbol{\alpha}_{i,t} = [\alpha_{i,1,t}, \ldots, \alpha_{i,K,t}]'$, the logit of the probability of transitioning from non-master to master on attribute k is

$$\operatorname{logit}[P(\alpha_{i,k,t+1} = 1 \mid \alpha_{i,k,t} = 0, \boldsymbol{\alpha}_{i,t})] = \lambda_0 + \lambda_1 \theta_i + \lambda_2 \sum_{\forall k' \neq k} \alpha_{i,k',t} + \lambda_3 \sum_{m=1}^{t} \sum_{j=1}^{J_t} q_{j,m,k}. \quad (1.2)$$

In this model, θ_i was used to denote the overall, time-invariant learning ability of subject i. The term $\sum_{\forall k' \neq k} a_{i,k',t}$ represents how many attributes subject i has already acquired other than attribute k, and $\sum_{m=1}^{t} \sum_{j=1}^{J_t} q_{j,m,k}$ denotes the number of items involving skill k that the student has completed at previous time points, in other words, the amount of practice so far on attribute k. By using a higher order logistic model for the transition probabilities in the hidden Markov model, the effect of different factors on the probability of learning a skill can hence be examined.

1.2 Current Model

The higher-order hidden Markov modeling framework developed by Wang et al. (2016) provides a flexible tool for modeling the effect of different covariates, both of the learners and of the instructions, on the outcome of learning. We hence adopt this framework and formulate our current model as follows.

Similar to Wang et al. (2016), we assume that the learning process is monotonic, with the probability of transitioning from mastery to nonmastery equal to 0 for all time points and all skills. At a certain time point, $t \in \{0, \ldots, T\}$, a student $i \in \{1, \ldots, N\}$ is assigned to learn a skill $k \in \{1, \ldots, K\}$, by receiving an instructional tool $l \in \{1, \ldots, L_k\}$ targeting the learning of skill k. Here, the set $\{1, \ldots, L_k\}$ is used to denote the collection of all instructional tools, including exercises, slides, videos, or games, that can have an effect on the learning of skill k. We let the attribute pattern of subject i at time t be $\boldsymbol{\alpha}_{i,t}$, then the probability that the subject masters skill k at time $t+1$ is

$$P(\alpha_{i,k,t+1} = 1 \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_k, \mathbf{U}_{k,l}) = \begin{cases} 1 & \alpha_{i,k,t} = 1, \\ \frac{\exp\left(\lambda_{0,k,l} + \sum_{k' \neq k} \lambda_{k',k,l} \cdot \alpha_{i,k',t}\right)}{1+\exp\left(\lambda_{0,k,l} + \sum_{k' \neq k} \lambda_{k',k,l} \cdot \alpha_{i,k',t}\right)} & \alpha_{i,k,t} = 0. \end{cases} \tag{1.3}$$

where

$$\lambda_{0,k,l} = \gamma_{0,k,0} + U_{0,k,l}, \;\; \lambda_{k',k,l} = \gamma_{k',k,0} + U_{k',k,l}, \tag{1.4}$$

with $\mathbf{U}_{k,l} = \begin{bmatrix} U_{0,k,l} \\ U_{1,k,l} \\ \vdots \\ U_{(k-1),k,l} \\ U_{(k+1),k,l} \\ \vdots \\ U_{K,k,l} \end{bmatrix} \sim \text{M.V.N.} \left(\mathbf{0}, \boldsymbol{\Sigma}_k = \begin{bmatrix} \tau_0^2 & \ldots & \tau_{0K} \\ \vdots & \ddots & \vdots \\ \tau_{0K} & \ldots & \tau_K^2 \end{bmatrix} \right)$. The γs are used to represent the fixed effects for a skill, and the Us are used for the random effects of the individual learning materials. More specifically, $\gamma_{0,k,0}$ denotes the overall log-odds that a learner with no mastered skills switches from nonmastery to mastery on skill k after learning a material targeting k, and $\mathbf{U}_{0,k,l}$ is the incremental "approachabiliy" of material l on learning k, which either increases ($U_{0,k,l} > 0$) or decreases ($U_{0,k,l} < 0$) the log-odds of acquiring skill k. In addition, because the learning of one skill may depend on the previous mastery of another related skill, we also considered the mastery of other skills as a covariate of the learning of skill k. Thus, fixed and random slopes for each skill other than k, denoted k', were introduced. We used $\gamma_{k',k,0}$ to denote the overall effect of mastery of k' on the learning of k, and $U_{k',k,l}$ was used to represent material l's additional "reliance" on skill k'.

Intuitively, the fixed intercepts, $\gamma_{0,k,0}$s, tells us the overall probability of learning a skill k for learners who haven't mastered any skills, without consideration of the material's characteristics. Since different learning materials can differ in their approachability (e.g., one material may be easily understood by students without backgrounds, whereas another does not), the random intercepts, $U_{0,k,l}$s were introduced to capture each material's approachability. The fixed slopes, $\gamma_{k',k,0}$s, tells us overall, how much does the mastery of k' affect the learning outcome of k. While some skills may be relatively independent, others could be strongly related. For instance, it is usually perceived that the mastery of addition is a prerequisite to understanding multiplication. Therefore, the mastery of addition can have a strong influence on whether a learner could understand multiplication. Lastly, different learning materials can rely on the mastery of other skills at different degrees. As an example, we consider two video lectures on multivariate statistics, with the first one providing an introduction to linear algebra and the second one jumping directly to the matrix-based derivations of multivariate distribution means and variances. The first video would require much less previous background on linear algebra than the second. Another example is the instruction of the same skill with different approaches: We could have two instructional tools teaching the students about the "right-hand rule" in electromagnetism, with the first one teaching the students how to use their right hand to determine the direction of the magnetic force, and the second one explaining how the rule works based on vector cross products. The former would barely require any background knowledge, while the latter would require some background in linear algebra. For this reason, we introduce the random slopes for each specific learning material, $U_{k',k,l}$s, indicating to what extend would a specific material l rely on prior mastery of k' in the process of learning k.

We treat the slope and intercept of different learning materials as random effects for several reasons: First, educators are often interested in evaluating different interventions' effectiveness. Under our model, the random intercepts can directly be used to compare different materials' overall effectiveness to students in the baseline group (i.e., $\boldsymbol{\alpha}_i = \mathbf{0}$), and to get the materials' effectiveness for other groups, the effectiveness could be calculated as the random intercept plus the random slope for that material. Second, for the multilevel model, the distribution of the random effects is modelled with a multivariate normal distribution. In the learning context, we may assume that the approachability of a material could be related to the degree of its requirement of a specific prerequisite. For example, a material deemed to be "harder" for students with none of the mastered skills may have a high requirement on another skill. Therefore, we use the covariance matrix of the learning material's random effects to capture potential relationships between different characteristics of learning materials.

Under the current model, a confirmatory approach is taken — that is, at time t, if the students were assigned a material targeting a specific skill, k, then we assume that only the mastery status on skill k can change from time t to time $t+1$. Here, we focus on the case of one target skill only per material. In that case, given the attribute pattern at time t, $\boldsymbol{\alpha}_{i,t}$, the probability of transitioning to pattern $\boldsymbol{\alpha}_{i,t+1}$ after receiving a learning material l targeting skill k is

$$P(\boldsymbol{\alpha}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t}) = \begin{cases} P(\alpha_{i,k,t+1} = 1 \mid \cdot)^{\alpha_{i,k,t+1}}[1 - P(\alpha_{i,k,t+1} = 1 \mid \cdot)]^{1-\alpha_{i,k,t+1}}, & \text{if } \boldsymbol{\alpha}_{i,[k],t} = \boldsymbol{\alpha}_{i,[k],t+1}, \\ 0, & \text{otherwise.} \end{cases} \tag{1.5}$$

In the equation above, $P(\alpha_{i,k,t+1=1|\cdot}) = P(\alpha_{i,k,t+1} = 1 \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_k, \mathbf{U}_{k,l})$ denotes the conditional probability of mastering a skill after learning, as given in equation (1.3).

The transition model in equation (1.3) is used to describe the changes of the students' attribute patterns across time. However, the students' latent mastery status cannot be directly observed and needs to be measured using assessment items. In our design, we assume that for each student at each time point, a set of J_t items are administered to them to assess their mastery of skills in the curriculum. For simplicity, in the current study we used the DINA model in equation (1.1) for the responses. However, the modeling framework can be extended to other CDMs by swapping the DINA measurement model with other CDMs, such as the reduced Reparameterized Unified Model (rRUM; Hartz, 2002) or the generalized-DINA (G-DINA; de la Torre, 2011) model.

Chapter 2

Model Estimation

2.1 Bayesian Formulation

Similar to Wang et al. (2016), a Bayesian modeling framework is used to estimate the parameters of the proposed model. Let $\boldsymbol{p(\pi)}$, where $\boldsymbol{\pi} = [\pi_1, \ldots, \pi_{2^K}]$ denotes the prior distribution of the population membership probabilities of each attribute pattern at the initial stage, $t = 1$, let $\boldsymbol{p(\gamma)}$ be the prior distribution for the fixed effects, and let $\boldsymbol{p(\Sigma)}$ be the prior distribution for the covariance of the random effects, denoted as $\boldsymbol{\Sigma} = [\boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_K]$, where $\boldsymbol{\Sigma}_k$ represents the covariance matrix of the random effects of materials for skill k. We further use k_{it} and l_{it} to denote the skill and material given to student i at time t, respectively, and let $\mathbf{X} = [X_1, \ldots, X_T]$ represent the observed response data across all time points, where X_t is the $N \times J_t$ response matrix of all the N learners at time t. Then the joint likelihood of the observed responses, learning materials' random effects, latent attribute patterns, population membership probabilities, the DINA model item parameters, the learning model's fixed effects, and the covariance matrices of the random effects is

$$
\begin{aligned}
L(\mathbf{X}, \boldsymbol{\alpha}, \mathbf{U}, \boldsymbol{\pi}, \boldsymbol{\Sigma}, \boldsymbol{\gamma}, \mathbf{s}, \mathbf{g}) \quad = \quad & \prod_{i=1}^{N} \Big\{ \pi_c P(\mathbf{X}_{i1} \mid \boldsymbol{\alpha}_{i,1} = \boldsymbol{\alpha}_c, \mathbf{s}, \mathbf{g}) \\
\times \; & \prod_{t=1}^{T-1} \Big[P(\boldsymbol{\alpha}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_{k_{it}}, U_{k_{it}, l_{it}}) P(\mathbf{X}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t+1}, \mathbf{s}, \mathbf{g}) \Big] \Big\} \\
\times \; & \prod_{k=1}^{K} \prod_{l=1}^{L_k} \Big[P(\mathbf{U}_{k,l} \mid \boldsymbol{\Sigma}_k) \Big] p(\boldsymbol{\pi}) p(\boldsymbol{\gamma}) p(\boldsymbol{\Sigma}). \qquad (2.1)
\end{aligned}
$$

Based on this formulation, a learner i's attribute pattern at the initial time point, $\boldsymbol{\alpha}_{i,1}$, follows a multinomial distribution, i.e.,

$$P(\boldsymbol{\alpha}_{i,1} = \boldsymbol{\alpha}_c) = \prod_{c}^{2^K} \pi_c^{\mathcal{I}(\boldsymbol{\alpha}_{i,1}=\boldsymbol{\alpha}_c)}, \tag{2.2}$$

where the prior for the population membership probabilities at the initial time point is

$$\boldsymbol{\pi} = [\pi_1, \ldots, \pi_{2^K}] \sim \text{Dirichlet}(\boldsymbol{\delta}_0). \tag{2.3}$$

We further assign the following priors to the fixed effects,

$$\gamma_{0,k,0} \sim N(0,1) \tag{2.4}$$

$$\gamma_{k\prime,k,0} \sim \text{log-normal}(-.5, .5) \tag{2.5}$$

for $\forall k$ and $\forall k\prime \neq k$.

For the random effects $\mathbf{U}_{k,l} \sim$ M.V.N.$(\mathbf{0}, \boldsymbol{\Sigma}_k)$, we assume for all k, the prior for the covariance is given by the Inverse-Wishart distribution, where

$$\boldsymbol{\Sigma}_k^{-1} \sim \text{Wishart}(\mathbf{S}, \nu), \tag{2.6}$$

with $\mathbf{S}$ being the scale matrix and ν the degrees of freedom.

Lastly, adopting the methods in Culpepper (2015), we assign a truncated Beta prior distribution to the DINA model item parameters, i.e.

$$p(s_j, g_j) \propto s_j^{a_s-1}(1-s_j)^{b_s-1} g_j^{a_g-1}(1-g_j)^{b_g-1}\mathcal{I}(0 \leq g_j < 1-s_j \leq 1). \tag{2.7}$$

Then, the full conditional distributions of the parameters, given the observed responses at all time points, are as follows:

- For $\boldsymbol{\alpha}_{i,t}$: Similar to Wang et al. (2016) and Chen, Culpepper, et al. (2017), we use a forward-backward algorithm to sequentially update the attribute patterns of each subject at each time point. Specifically,

$$P(\boldsymbol{\alpha}_{i,t} = \boldsymbol{\alpha}_c) \propto \begin{cases} \pi_c P(X_{i1} \mid \boldsymbol{\alpha}_{i,1} = \boldsymbol{\alpha}_c) P(\boldsymbol{\alpha}_{i,2} \mid \boldsymbol{\alpha}_{i,1} = \boldsymbol{\alpha}_c, \cdot), & t = 1; \\ P(\boldsymbol{\alpha}_i = \boldsymbol{\alpha}_c \mid \boldsymbol{\alpha}_{i,t-1}, \cdot) P(X_{it} \mid \boldsymbol{\alpha}_{i,t} = \boldsymbol{\alpha}_c) P(\boldsymbol{\alpha}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t} = \boldsymbol{\alpha}_c, \cdot), & 1 < t < T; \\ P(X_{iT} \mid \boldsymbol{\alpha}_{i,T} = \boldsymbol{\alpha}_c) P(\boldsymbol{\alpha}_{i,T} = \boldsymbol{\alpha}_c \mid \boldsymbol{\alpha}_{i,T-1}, \cdot), & t = T, \end{cases} \tag{2.8}$$

where $P(\boldsymbol{\alpha}_{i,t+1} = \boldsymbol{\alpha}_c \mid \boldsymbol{\alpha}_{i,t}, \cdot)$ represents $P(\boldsymbol{\alpha}_{i,t+1} = \boldsymbol{\alpha}_c \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_{k_{i,t}}, \mathbf{U}_{k_{i,t}, l_{i,t}})$ and is given in equations (3) and (4).

- For population proportions of the attribute patterns at time 1, the full conditional distribution of $\boldsymbol{\pi}$ is still a Dirichlet distribution, with

$$\boldsymbol{\pi} \mid \boldsymbol{\alpha}_{1,1} \ldots, \boldsymbol{\alpha}_{N,1} \sim \text{Dirichlet}(\delta_0 + \tilde{N}), \tag{2.9}$$

where $\tilde{N} = [\sum_{i=1}^{N} \mathcal{I}(\boldsymbol{\alpha}_{i1} = \boldsymbol{\alpha}_1), \ldots, \sum_{i=1}^{N} \mathcal{I}(\boldsymbol{\alpha}_{i1} = \boldsymbol{\alpha}_{2^K})]$

- For fixed effect $\boldsymbol{\gamma}_k$, we have

$$P(\boldsymbol{\gamma}_k \mid \boldsymbol{\alpha}, \mathbf{U}_{k,\cdot}) \propto p(\boldsymbol{\gamma}_k) \prod_{t=1}^{K} \prod_{i:k_{i,t}=k} P(\boldsymbol{\alpha}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_k, \mathbf{U}_{k,l_{i,t}}). \tag{2.10}$$

- For random effects $\mathbf{U}_{k,l}$, we have

$$P(\mathbf{U}_{k,l} \mid \boldsymbol{\Sigma}_k, \boldsymbol{\alpha}, \boldsymbol{\gamma}_k) \propto \prod_{t=1}^{K} \prod_{i:k_{i,t}=k \ \& \ l_{i,t}=l} P(\boldsymbol{\alpha}_{i,t+1} \mid \boldsymbol{\alpha}_{i,t}, \boldsymbol{\gamma}_k, \mathbf{U}_{k,l}) P(\mathbf{U}_{k,l} \mid \boldsymbol{\Sigma}_k). \tag{2.11}$$

- For the covariance matrices of random effects, $\boldsymbol{\Sigma}_k$, conditioning on the random effects of materials targeting skill k, $\boldsymbol{\Sigma}_k$ follows an Inverse-Wishart distribution, i.e.,

$$\boldsymbol{\Sigma}_k^{-1} \mid \mathbf{U}_{1,k}, \ldots, \mathbf{U}_{L_k,k} \sim \text{Wishart}(\mathbf{S}^*, \nu^*), \tag{2.12}$$

where $\mathbf{S}^* = \mathbf{U}^T\mathbf{U} + \mathbf{S}$, and $\nu^* = L_k + \nu$, and $\mathbf{U}$ represents the matrix of the random effects corresponding to skill k, with $\mathbf{U}_{k,l}$ as the lth row.

- For the DINA model parameters, $\mathbf{s}$ and $\mathbf{g}$, the posterior distribution, given the responses and attribute patterns, is a truncated Beta distribution, with

$$p(s_j, g_j \mid \boldsymbol{\alpha}, \mathbf{X}) \propto s_j^{a_{sj}-1}(1-s_j)^{b_{sj}-1} g_j^{(a_{gj}-1)}(1-g_j)^{b_{gj}-1}\mathcal{I}(0 \leq g_j < (1-s_j) \leq 1), \quad (2.13)$$

where

$$a_{sj} = \sum_{i:X_{ij}=0} \eta_{ij} + a_s, \qquad b_{sj} = \sum_{i:X_{ij}=1} \eta_{ij} + b_s,$$

$$a_{gj} = \sum_{i:X_{ij}=1} (1-\eta_{ij}) + a_g, \qquad b_{gj} = \sum_{i:X_{ij}=0} (1-\eta_{ij}) + b_g.$$

2.2 Parameter Estimation

A Metropolis-Hastings within Gibbs algorithm is used to estimate the parameters of the proposed model. Let $\boldsymbol{\alpha}_{i,t}^0, \mathbf{U}^0, \boldsymbol{\gamma}^0, \boldsymbol{\pi}^0, \mathbf{s}^0, \mathbf{g}^0$, and $\boldsymbol{\Sigma}^0$ denote the initial values of our parameters of interest, the following procedures could be used to sample from the posterior distribution of the parameters. At each iteration of the MCMC chain:

(1) For each $i = 1, 2, \ldots, N, t = 1, 2, \ldots, T$, sample $\boldsymbol{\alpha}_{i,t}^r$, given $\mathbf{U}^{r-1}, \boldsymbol{\gamma}^{r-1}, \boldsymbol{\pi}^{r-1}$, and $\mathbf{s}^{r-1}, \mathbf{g}^{r-1}$;

(2) Sample $\boldsymbol{\pi}^r$ based on $\boldsymbol{\alpha}_{1,1}^r, \ldots, \boldsymbol{\alpha}_{N,1}^r$;

(3) For each $\gamma_{k',k,0}$, where $k, k' = 0, 1, \ldots, K$, sample $\gamma_{k',k,0}^r$ from the Uniform distribution in an interval around $\gamma_{k',k,0}^{r-1}$, and accept with probability $\frac{P(\gamma_{k',k,0}^r|\boldsymbol{\alpha}_{t+1}^r, \boldsymbol{\alpha}_t^r, \mathbf{U}_{k,\cdot}^{r-1})}{P(\gamma_{k',k,0}^{r-1}|\boldsymbol{\alpha}_{t+1}^r, \boldsymbol{\alpha}_t^r, \mathbf{U}_{k,\cdot}^{r-1})}$;

(4) For each $k = 1, 2, \ldots, K, l = 1, 2, \ldots, L_k$, sample $\mathbf{U}_{k,l}^r$ from the Uniform distribution of a small interval around $\mathbf{U}_{k,l}^{r-1}$, and accept with probability $\frac{P(\mathbf{U}_{k,l}^r|\boldsymbol{\Sigma}_k^{r-1}, \boldsymbol{\alpha}_{t+1}^r, \boldsymbol{\alpha}_t^r, \boldsymbol{\gamma}_k^r)}{P(\mathbf{U}_{k,l}^{r-1}|\boldsymbol{\Sigma}_k^{r-1}, \boldsymbol{\alpha}_{t+1}^r, \boldsymbol{\alpha}_t^r, \boldsymbol{\gamma}_k^r)}$.

(5) For each $k = 1, 2, \ldots, K$, sample $\boldsymbol{\Sigma}_k^r$ from the inverse-Wishart distribution updated based on $\mathbf{U}_{k,1}^r, \ldots, \mathbf{U}_{k,L_k}^r$.

(6) For each item j, sample s_j, g_j from the truncated Beta distribution, given $\boldsymbol{\alpha}^r$ and $\mathbf{X}$.

2.3 Model Identification

We note that for the logistic transition model given in equation (1.3), for each of the slopes and intercepts, $\lambda_{k',k,l}$ and $\lambda_{0,k,l}$ remain the same if we add a constant to the corresponding fixed effect and subtract the constant from each of the corresponding random effects. In other words,

$$\lambda_{0,k,l} = (\gamma_{0,k,0} + c) + (U_{0,k,l} - c), \qquad \lambda_{k',k,l} = (\gamma_{k',k,0} + c) + (U_{k',k,l} - c).$$

Although we are unable to exactly fix the location of the fixed and random effects, they can be softly centered by imposing a mean of $\mathbf{0}$ to the multivariate normal distribution of the random effects (e.g., Gelman et al., 2014).

Chapter 3

Simulation Study

3.1 True Parameter Generation

We conduct a series of simulation studies to evaluate the recovery of model parameters using the proposed algorithm. The design of the learning process is as folows: The curriculum is assumed to contain $K = 4$ skills. At the initial time point, $t = 1$, all students are administered $J_t = 15$ items. Following the assessment, each student is administered a random learning material targeting a random skill. Following the learning of the material, students are assessed again with 15 items, and this process iterates up to the final time point, $T = 5$. Because our main interest is the recovery of the parameters in the learning model, two factors are considered, including the the number of students who receive each specific training material ($N_{l,k} = 50$ and 100), and the number of materials available for each skill ($L_k = 10, 30,$ and 50). The L_ks under each condition are chosen to mimic the number of available online resources for each Common Core Math standard on the OpenEd repository. Based on these two factors, we can calculate the total number of subjects in each condition. For example, in the case of 50 students per material and 10 materials per skill, the total sample size would be $N = 50 \times 10 \times K = 50 \times 10 \times 4 = 2000$. Such large samples are very difficult to collect in a classroom learning environment. However, they are achievable in online learning environments, where students from different backgrounds can access the learning materials anytime and anywhere. 10 repetitions of the simulation was performed under each condition.

The initial attribute patterns of the students, $\boldsymbol{\alpha}_{\cdot,1}$, are generated following similar methods as in Chiu and Köhn (2016). Specifically, for each learner i, we first generate a K-dimensional multivariate normal variable $\boldsymbol{Z}_i \sim \text{M.V.N.}(\mathbf{0}, \boldsymbol{\Sigma}_Z)$, where the diagonal entries of $\boldsymbol{\Sigma}_Z$ are 1, and the off-diagonal entries are .2. That is, we assume that at the initial time point, the higher-order continuous traits behind each skill are correlated at $\rho = .2$. Next, $\alpha_{i,k,1} = \mathcal{I}(Z_{ik} \geq \Phi^{-1}(\frac{k}{K+1}))$ was obtained for each k, giving the true attribute pattern of subject i at the initial time point, $\boldsymbol{\alpha}_{i,t}$.

The true values of the fixed intercepts, $\gamma_{0,k,0}$s, were generated from $N(-1, 1)$, and the fixed slopes, $\gamma_{k',k,0}$s were sampled from Uniform$(0, 2)$. The covariance matrices of the random effects of each skill, $\boldsymbol{\Sigma}_k, k \in \{1, \ldots, K\}$, were assumed to have diagonal entries of .3 and off diagonal entries of .1. Based on the covariance matrices, the true random intercepts and slopes were generated from the multivariate normal distribution, with $\mathbf{U}_{k,l} \sim$ M.V.N.$(\mathbf{0}, \boldsymbol{\Sigma}_k)$. Based on the random and fixed effects of the learning model, the learners' subsequent $\boldsymbol{\alpha}$s can be generated based on transition probabilities as defined in equations (3) and (4).

Five blocks of DINA model items, with $J_t = 15$ items per block, were simulated. The true slipping and guessing parameters were generated from $U(.15, .3)$. Xu and Zhang (2016) have shown that in the static case, under the DINA model, the item parameters $\mathbf{s}, \mathbf{g}$ and the population membership probabilities $\boldsymbol{\pi}$ are identifiable if the Q-matrix contains 3 identity matrices. Thus, for each block of items, the Q-matrices were generated to include at least 3 items that exclusively measures the kth skill for all k. In addition, because we assumed the learning process to be monotonic, the proportion of students mastering a large number of skills are expected to increase as the learning process continues. In that case, the proportion of students mastering few skills may be small at later time points. To ensure sufficient sample size for the estimation of both $\mathbf{s}$ and $\mathbf{g}$ of the items from all blocks, we followed the practice in Wang et al. (2016) by using a block design for item assignment to subjects. Specifically, the N examinees were randomly assigned to 5 test design groups. For students in group 1, they were administered items in blocks 1, 2, 3, 4, and 5 at times $t = 1, 2, 3, 4, 5$, respectively. For students in group 2, they were administered item blocks $2, 3, 4, 5, 1$ at times 1 to 5. And students in group 3 were administered the item blocks in the order of $3, 4, 5, 1, 2$, and so on. Based on the students' attribute patterns at each time point, the block of items they are administered at time t, and the Q-matrix and parameters of the items in each block, the subjects' responses can be randomly generated based on the DINA model item response function.

3.2 Parameter Estimation

We assigned random initial values to the model parameters to start the MCMC algorithm. Specifically, the initial population membership probabilities, $\boldsymbol{\pi}$, were generated from the Dirichlet distribution with $\boldsymbol{\delta} = \mathbf{1}$; the covariance matrices of the random effects, $\boldsymbol{\Sigma}_k$s, were set to be the identity matrix; initial values for $\mathbf{U}, \boldsymbol{\gamma}, \mathbf{s}, \mathbf{g}$ were respectively randomly sampled from the multivariate normal and the uniform distributions, similar to how we generated the true values; lastly, the initial attribute patterns of the subjects were randomly sampled based on the initial value of $\boldsymbol{\pi}$, and the subsequent $\boldsymbol{\alpha}$s were generated based on the learning model and the initial learning model parameters. Then, we implemented the MH within Gibbs algorithm to iteratively update the parameters. For the Metropolis-Hastings sampling of the fixed and random effects, window sizes of the uniform distributions around the old values were selected so that the average acceptance rate for $\boldsymbol{\gamma}$ and for $\mathbf{U}$ were both between $20\% - 40\%$. A chain length of $20,000$ iterations was used for each condition, with burn-in of $10,000$ iterations.

3.3 Evaluation Criteria

Based on the parameter samples from the MCMC, we can calculate the expected a posteriori (EAP) estimates of the $\mathbf{U}, \boldsymbol{\gamma}, \mathbf{s}, \mathbf{g}, \boldsymbol{\Sigma}$, and $\boldsymbol{\alpha}$ using

$$\hat{\theta}_{EAP} = \frac{\sum_{r=T_{burn}+1}^{T_{tot}} \theta^r}{T_{tot} - T_{burn}}, \tag{3.1}$$

where θ^r denotes the parameter sample from the rth iteration, and T_{tot} and T_{burn} are the total length and burn-in length of the MCMC chains, respectively. For the $\boldsymbol{\alpha}$s, because for each $\alpha_{i,k,t}$, the EAP will give a value between 0 and 1, we further dichotomize each $\alpha_{i,k,t}$ and set it to 1 if the corresponding EAP is greater than .5, and 0 otherwise.

The estimation accuracy of the attribute patterns at each time point is evaluated by the attribute-wise agreement rate (AAR), where

$$AAR(\hat{\boldsymbol{\alpha}}_t) = \frac{\sum_{i=1}^{N} \sum_{k=1}^{K} \mathcal{I}(\alpha_{i,k,t} = \hat{\alpha}_{i,k,t})}{N \times K}, \tag{3.2}$$

and the estimation accuracy of $\boldsymbol{\pi}$ and $\boldsymbol{\Sigma}$ is evaluated based on their bias and RMSE, given by

$$\text{Bias}(\hat{\theta}) = \hat{\theta}_{EAP} - \theta, \qquad \text{RMSE}(\hat{\theta}) = \sqrt{\frac{\sum_{r=T_{burn}+1}^{T_{tot}} (\theta^r - \theta)^2}{T_{tot} - T_{burn}}}, \tag{3.3}$$

where θ denotes the true value of the parameter.

Because we used soft centering for the fixed and random effects, the estimated values of both might be slightly shifted from the true values. However, if the algorithm could recover the parameters accurately, the composite of the estimated fixed and random effects, λs, should be close to true values. We hence cannot use bias as a criteria for the recovery of the fixed and random effects. Instead, the Pearson correlations between the estimated and the true fixed and random effects are computed.

Chapter 4

Results

Table 1 presents the accuracy of attribute estimates at different time points in the learning process, averaged across 10 repetitions under each simulation condition. Across all conditions and all time points, the attribute-wise agreement rate between the true and estimated $\boldsymbol{\alpha}$s remained above .8, indicating relatively accurate estimates of the learners' patterns at all time points. We observed that across all conditions, the AAR in the initial time point, $t = 1$, was higher than that of the later time points. In addition, the AARs seemed to be higher for conditions with more learners per material. This suggests that part of the estimation error in the attribute patterns can be attributed to the errors in the random effect estimates for the transition model. For the initial time point, where the transition model takes a smaller weight in the estimation algorithm and for conditions with more data for the estimation of the transition model parameters, the attribute patterns were estimated slightly better.

Table 4.1: Attribute-wise agreement rates (AARs) at different time points under different simulation conditions. L_k stands for the number of available materials targeting skill k, and $N_{l,k}$ stands for the number of learners who are administered each material. Results are aggregated across 10 repetitions.

L_k	$N_{l,k}$	$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$
10	50	.884	.835	.836	.844	.839
10	100	.891	.839	.841	.844	.834
30	50	.887	.847	.844	.839	.829
30	100	.882	.838	.836	.839	.841
50	50	.892	.839	.840	.837	.830
50	100	.891	.835	.837	.838	.836

We further investigated the parameter recovery of the fixed effects, $\boldsymbol{\gamma}$, and the random effects $\mathbf{U}$ for the learning materials. Figures 1 and 2 present the scatter plots of the true and estimated fixed and random effects under different simulation conditions. We combined the parameters from different repetitions into one plot for each condition. Across all conditions, the fixed effects seemed to be accurately recovered with the estimation algorithm, with the points on the scatterplot nearly forming a straight line. With increasing $N_{l,k}$ (number of students using each learning material) and L_k (number of materials available for each skill), we see slight improvements in estimation accuracy for the fixed effects. The random effects corresponding to the learning materials demonstrated larger error. From Figure 2, we can see that although there is strong correlation between the true and estimated random parameters, there are more deviations of the points from a straight line compared to the plot of $\boldsymbol{\gamma}$. Similar to the fixed effects, for conditions with larger $N_{l,k}$ and larger L_k , the correlation between true and estimated Us appeared stronger.

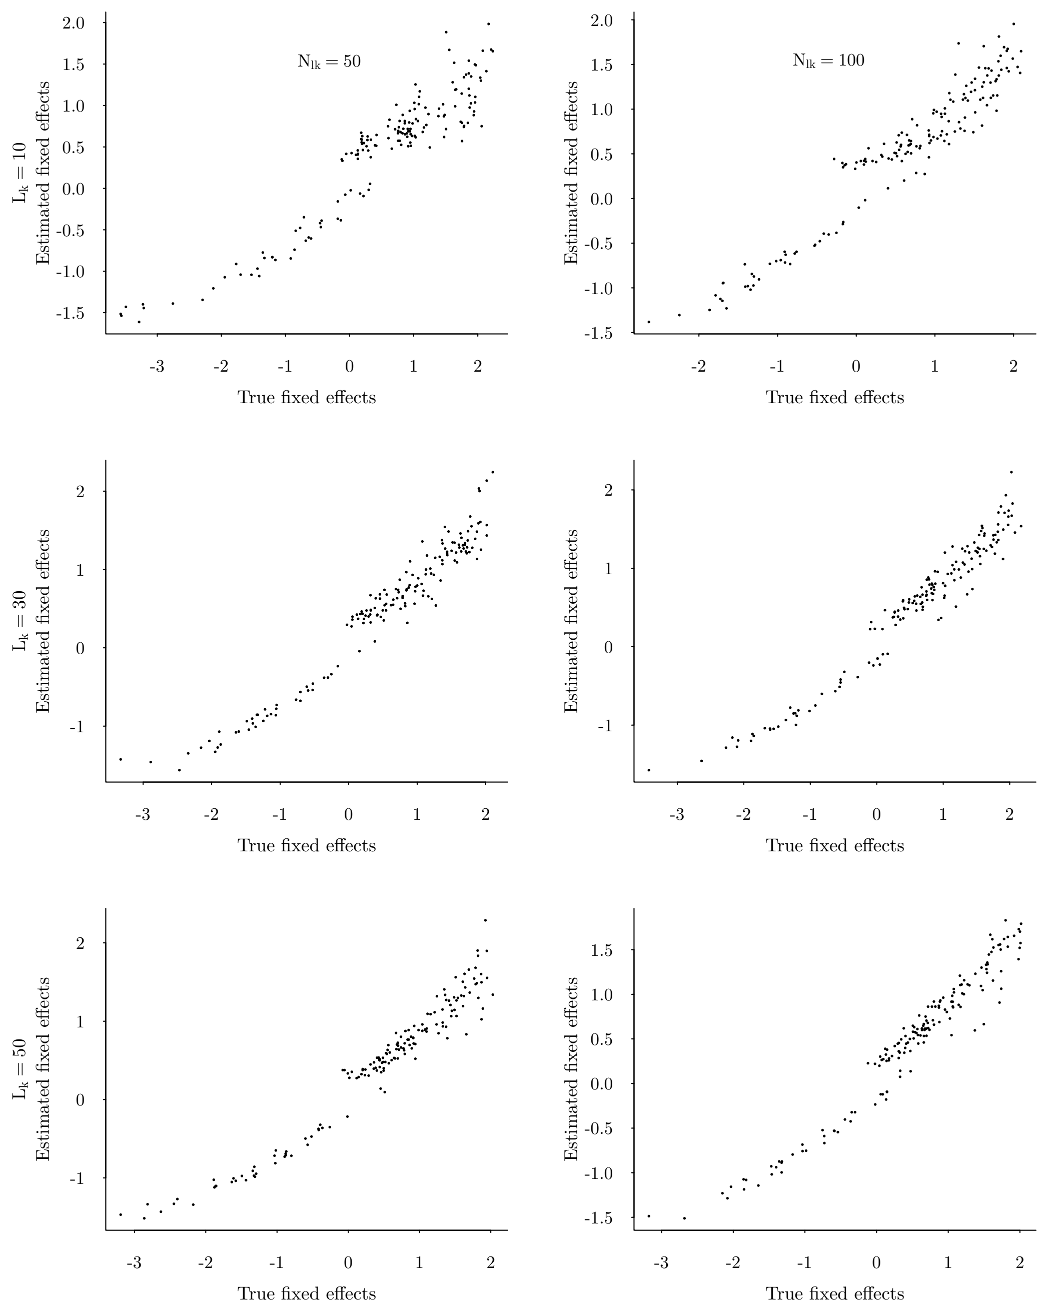

Figure 4.1: Scatterplot for the true and estimated fixed effects ($\boldsymbol{\gamma}$) under different simulation conditions.

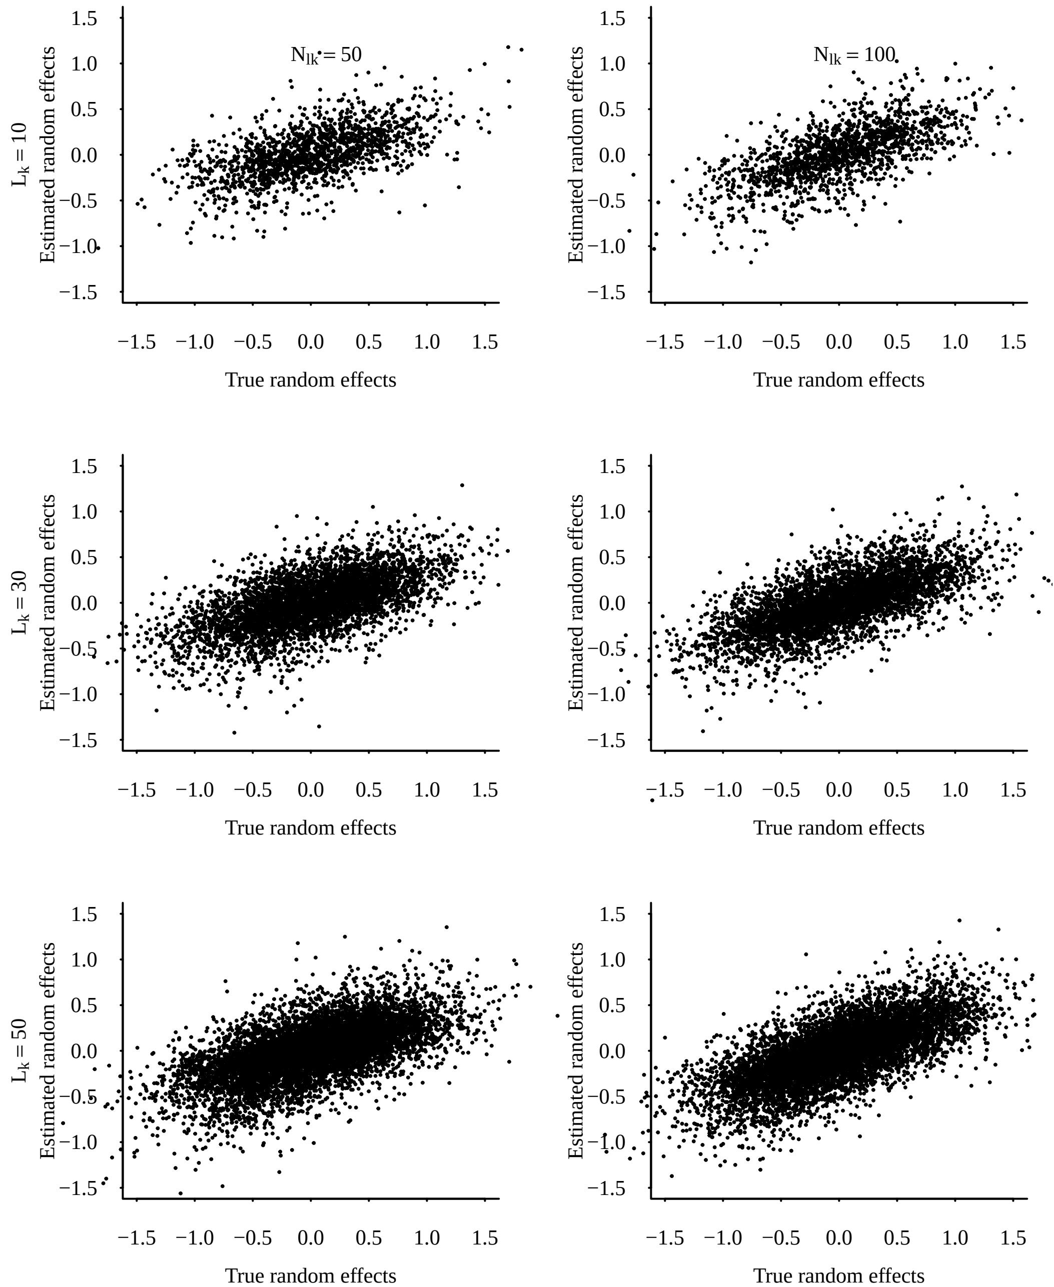

Figure 4.2: Scatterplot for the true ad estimated random effects ($\mathbf{U}$) under different simulation conditions.

Table 2 provides the Pearson correlations between true and estimated random ($\mathbf{U}$) and fixed ($\boldsymbol{\gamma}$) effects, as well as the bias and/or RMSE of the initial population membership probabilities ($\boldsymbol{\pi}$) and the covariance matrices of learning materials' random effects ($\boldsymbol{\Sigma}$) under different conditions.

We omitted the bias of $\boldsymbol{\pi}$ here, because the total probability of different classes sums to 1, so the bias would always be 0. In addition, we separately calculated the bias and RMSE of the diagonal entries (τ_k^2s) and the off-diagonal entries ($\tau_{k'k}$) of the covariance matrix of random effects, because the interpretations of the diagonal and off diagonal elements are different. Taking a coarse look at Table 2, we observe that the correlation between true and estimated fixed effects, γs, were above .94 across all conditions and increased slightly with larger $N_{l,k}$ and L_k, which is consistent with Figure 1. The correlations for the random slopes and intercepts, Us, were in the .59 to .70 range across conditions and were higher when the number of learning materials available for each skill increased, or when the number of subjects using each material increased. One explanation is that, by increasing the number of available training materials, the distribution of the random effects could be estimated better, hence improving the random effect estimates. And when the number of subjects using a particular material is large, we have more data on the transitions of the subjects on this skill using this particular material, hence improving the random effect estimates corresponding to the material. Across all conditions, the RMSE for $\boldsymbol{\pi}$ was low, but conditions with more learning materials or larger sample size per material had slightly better $\boldsymbol{\pi}$ estimates. We think this is mainly due to larger overall sample sizes in those conditions. For the covariance matrices of random effects, we observe that both the variances on the diagonal and the covariances on the off-diagonal were estimated much better for 30 or 50 materials per skill than with only 10 available materials per skill, as indicated by the smaller bias and RMSE values. This suggests that with 10 observations, the distribution of the random effects cannot be recovered very well. However, with 30 or 50 random observations, the true covariance matrix can be accurately recovered. Although the number of learners per material did not seem to have too large an effect on the bias of the covariance matrix entry estimates, the RMSE of the covariance matrix estimates were consistently lower when $N_{l,k}$ was large, indicating that the estimates were more stable when sample size for each material was larger.

Table 4.2: Parameter recovery of fixed (γs) and random (Us) effects, initial population membership probabilities ($\boldsymbol{\pi}$), and covariance matrix for random effects ($\boldsymbol{\Sigma}$).

L_k	$N_{l,k}$	ρ_γ	ρ_U	RMSE($\boldsymbol{\pi}$)	Bias(τ_k^2)	RMSE(τ_k^2)	Bias($\tau_{k'k}$)	RMSE($\tau_{k'k}$)
10	50	.947	.597	.011	.019	.221	.073	.159
10	100	.965	.671	.008	.018	.201	.069	.143
30	50	.977	.612	.008	.016	.184	.064	.122
30	100	.974	.675	.007	.018	.164	.075	.106
50	50	.978	.612	.007	.015	.178	.057	.113
50	100	.975	.700	.005	.016	.151	.064	.095

Chapter 5

Discussion

In the present study, we proposed a hidden Markov model with both learner characteristics (i.e.,previous skill pattern) and learning material characteristics (i.e., approachability and reliance on other skills) as covariates affecting the outcome of learning. By modeling the learning materials' characteristics as random effects, we could capture the underlying distribution structure of the materials' approachability and dependency on other skills, and the random slope and intercept estimates obtained for each learning material could also be used to assess their effectiveness on students with various attribute profiles. Under the proposed Bayesian estimation algorithm, model parameters of the multilevel logistic HMM were relatively accurately recovered, especially when there is a sufficient number of learning materials per skill and when there is sufficient learners using each material.

There are a few potential implications of the current model. First of all, a long term interest of educational practitioners is understanding the structural relationship between different attributes. By knowing which skills are prerequisites to others, educators can design the courses so that basic contents precede the advanced skills. In addition, if a student is assigned to learn a complex skill and fails to learn it after many trials, one possible explanation of the difficulty in learning the skill is missing prerequisites. Knowing which skills are prerequisites to an advanced, hard to learn skill may help educators identify why a student cannot learn effectively, and the student can be routed back to learn the basic prerequisite. Previous research on the hierarchical structure of attributes (e.g., Leighton, Gierl, & Hunka, 2004) usually require the attribute hierarchies to be pre-specified by content experts. Under the current model, the dependencies between skills can potentially be inferred from the fixed slope estimates in the transition model. In other words, if the fixed slope of skill k' is large in the logistic transition model for learning skill k, we can infer that whether a student can learn skill k successfully strongly depends on the previous mastery of skill k'.

A second potential application of the proposed model is the adaptive recommendation of contents to students. In adaptive learning systems, students are sequentially assigned contents to learn. To help learners achieve their learning goals (e.g., mastering all skills) as efficiently as possible, content at each time point can be adaptively selected based on previous knowledge about the student's characteristics, such as his or her previous mastery on the skills in the curriculum, level of engagement in the content, or learning style (Oxman, Wong, & Innovations, 2014). The current modeling framework could help us determine which materials could be deemed as "optimal" for a student: Under the current model, based on a learner's current attribute pattern estimate and the parameters of different learning materials, the proposed model can answer simple questions such as "What's the probability of mastering skill k if we decide to let the student learn k next and administer material l to him/her?" or "Among all materials targeting skill k, which one will be most effective to the student right now?". It can also be used to answer more complex questions, such as "Among all possible trajectories of instructions, which one can give us the highest expected attribute pattern after t stages?"

The present study has some limitations. First of all, we only considered one target skill per instructional material. We could, however, imagine materials that can simultaneously teach students multiple skills. If we assume conditional independence between the transitions on each skill, then the current model can be straightforwardly extended to the situation of multiple target skills, where the probability of a pattern-wise transition is simply the product of the probabilities of the attribute-wise transitions. However, if the assumption of conditional independence of attribute-wise transitions is relaxed, the learning model will be much more complex. In that case, the probability of a pattern-wise transition will no longer be the product of the attribute-wise transition probabilities, and a learning model will need to look at covariates contributing to the change from one pattern to another for each pair of attribute patterns. Future research can look into models for pattern-wise transitions, with learning materials' characteristics as covariates.

Another limitation of the current study is the small number of total skills we assumed in the simulation studies. A typical course usually involves a large number of skills, ranging from 10 to more than a hundred. However, fitting the current model, or even any simple cognitive diagnosis model under the Bayesian framework with more than 10 skills will be computationally intensive. By converting the estimation algorithm of the fixed and random effects from a Metropolis-Hastings sampler to a Gibbs sampler could slightly alleviate the computational burden. Another potential direction is to look into methods of partitioning the skills into small sets, as we often tend to observe that skills in one cluster (e.g., basic arithmetic operations) are relatively independent from those in another cluster (e.g., reading figures). If the parameters can be estimated within each cluster of skills, the computational intensity of fitting a model with a large number of skills might be significantly reduced.

References

Chen, Y., Culpepper, S. A., Wang, S., & Douglas, J. (2017). A hidden markov model for learning trajectories in cognitive diagnosis with application to spatial rotation skills. *Applied Psychological Measurement*, 0146621617721250.

Chen, Y., Li, X., Liu, J., & Ying, Z. (2017). Recommendation system for adaptive learning. *Applied Psychological Measurement*, 0146621617697959.

Chiu, C.-Y., & Köhn, H.-F. (2016). The reduced RUM as a logit model: Parameterization and constraints. *Psychometrika*, *81*(2), 350–370.

Collins, L. M., & Wugalter, S. E. (1992). Latent class models for stage-sequential dynamic latent variables. *Multivariate Behavioral Research*, *27*(1), 131–157.

Culpepper, S. A. (2015). Bayesian estimation of the DINA model with Gibbs sampling. *Journal of Educational and Behavioral Statistics*, *40*(5), 454–476.

de la Torre, J. (2011). The generalized DINA model framework. *Psychometrika*, *76*(2), 179–199.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (Vol. 2). CRC press Boca Raton, FL.

Gütl, C., Rizzardini, R. H., Chang, V., & Morales, M. (2014). Attrition in mooc: Lessons learned from drop-out students. In *International workshop on learning technology for education in cloud* (pp. 37–48).

Hartz, S. (2002). *A Bayesian framework for the unified model for assessing cognitive abilities: Blending theory with practicality* (Unpublished doctoral dissertation). University of Illinois at Urbana-Champaign.

Henson, R. A., Templin, J. L., & Willse, J. T. (2009). Defining a family of cognitive diagnosis models using log-linear models with latent variables. *Psychometrika*, *74*(2), 191–210.

Junker, B. W., & Sijtsma, K. (2001). Cognitive assessment models with few assumptions, and connections with nonparametric item response theory. *Applied Psychological Measurement*, *25*(3), 258–272.

Langeheine, R. (1988). Manifest and latent markov chain models for categorical panel data. *Journal of Educational Statistics*, *13*(4), 299–312.

Leighton, J. P., Gierl, M. J., & Hunka, S. M. (2004). The attribute hierarchy method for cognitive assessment: A variation on tatsuoka's rule-space approach. *Journal of Educational Measurement*, *41*(3), 205–237.

Li, F., Cohen, A., Bottge, B., & Templin, J. L. (2015). A latent transition analysis model for assessing change in cognitive skills. *Educational and Psychological Measurement*, 0013164415588946.

Macready, G. B., & Dayton, C. M. (1977). The use of probabilistic models in the assessment of mastery. *Journal of Educational Statistics*, *2*(2), 99–120.

Oxman, S., Wong, W., & Innovations, D. (2014). White paper: Adaptive learning systems. *Integrated Education Solutions*.

Rupp, A. A., Templin, J., & Henson, R. A. (2010). *Diagnostic measurement: Theory, methods, and applications*. Guilford Press.

Vanlehn, K. (2006). The behavior of tutoring systems. *International journal of artificial intelligence in education*, *16*(3), 227–265.

von Davier, M. (2008). A general diagnostic model applied to language testing data. *British Journal of Mathematical and Statistical Psychology*, *61*(2), 287–307.

Wang, S., Yang, Y., Culpepper, S. A., & Douglas, J. A. (2016). Tracking skill acquisition with cognitive diagnosis models: A higher-order, hidden markov model with covariates. *Journal of Educational and Behavioral Statistics*, 1076998617719727.

Xu, G., & Zhang, S. (2016). Identifiability of diagnostic classification models. *Psychometrika*, *81*(3), 625–649.