Flow With an Intelligent Tutor: A Latent Variable Modeling Approach to Tracking Flow
During Artificial Tutoring

Hyeon-Ah Kang¹

Adam Sales²

Tiffany A. Whittaker¹

¹University of Texas at Austin

²Worcester Polytechnic Institute

Abstract

Increasing use of intelligent tutoring system (ITS) in education calls for analytic methods that can unravel students' learning behaviors. In this study, we explore a latent variable modeling approach to tracking flow during computer-interactive intelligent tutoring. Flow is a mental state a student achieves when immersed in deep learning. A student who flows in learning shows high engagement with learning activities and tends to achieve greater academic growth. When examined in ITS, flow state can evince student's interaction with the artificial tutor and modalities of optimal and suboptimal learning. The purpose of this study is to track progression of latent flow during tutoring. We apply latent variable models that allow discrete characterization of flow state and suggest practical model-fitting strategies that accommodate assumptions, estimation constraints and characteristics of ITS implementation data. Three models are considered to demonstrate the application: the (i) latent class model, (ii) latent transition model, and (iii) hidden Markov model. Using log data from Cognitive Tutor Algebra I, we show that the models provide practical information about learning flow. The models differed in handling the stochasticity and progression of problems but were generally consistent in inferring the flow trajectory and interaction modalities. The information drawn from the models was unique enough to warrant separate attention and well complemented with each other. Based on our experiential analyses, we discuss strengths and limitations of each model and illuminate areas that need future development.

Keywords: Intelligent tutoring, Flow, Latent transition analysis, Cognitive Tutor Algebra

Flow With an Intelligent Tutor: A Latent Variable Modeling Approach to Tracking Flow During Artificial Tutoring

1. Introduction

An intelligent tutoring system (ITS; e.g., ALEKS, ASSISTments, AutoTutor, Easy with Eve, MATHia, MetaTutor, SQL-Tutor) is educational software that provides an individualized learning platform. It is embedded with an artificial tutor that guides through problem steps and helps students with customized feedback. In ITS, students directly interact with the tutor to ask for help and advance learning. Since all educational activities are achieved through direct interaction with the software, it requires minimal intervention from human teachers and enables large-scale personalized learning.

One of the important aspects to consider in using ITS is whether a student adequately flows in learning. Flow (Csíkszentmihályi, 1990, 2008) is a mental state learners experience when they are engaged in activities that appropriately challenge their cognitive levels. A student in a flow state demonstrates high levels of confidence and satisfaction, and is likely to achieve optimal learning. When it comes to ITS, the inherent design often makes the flow an achievable and measurable construct. In ITS, problems are often assigned adaptively according to student's skill levels, and the student is generally expected to flow when learning through problems. If a student deviates from flow, for example, because of ill-match in the problems or other extraneous distractors, the student is likely to exhibit patterned behaviors. Inspecting flow in this setting can help identify students at risk of poor attainment (i.e., students to whom ITS provides little service) as well as modalities of interaction patterns that lead to optimal and suboptimal learning.

The purpose of this study is to explore statistical methods that can model learning flow in artificial intelligent tutoring. Although ITS is designed to fulfill interactive learning, student's perception of tasks can differ from the projection of the system. A student can receive tasks that exceedingly challenge his skill levels or facile tasks that spawn boredom or apathy. If a student encounters extraneous factors that divert attention, the student can

also deviate from the flow and exhibit detached behaviors. The goal of this study is to identify such occurrences and investigate their commonalities such that findings can inform future ITS refinement and intervention planning.

For modeling flow, we apply latent variable models that characterize in- and out-of-flows as discrete latent states. Latent variable models have been widely used across social sciences but received relatively little attention in learning science possibly due to uniqueness of ITS design and challenges in meeting the model assumptions. While a number of studies considered some models (e.g., Dupre, Booth, Bolster, Morrison, & McKeown, 2017; Fwa & Marshall, 2018; Shih, Koedinger, & Scheines, 2010; Sinclair et al., 2018), these studies focused on modeling particular aspects of engagement (e.g., affective, behavioral) and used the canonical forms without addressing the structural variance from the problems and students. In ITS, both students and problems bring unique variance to outcome data and they need to be addressed by either explicit modeling or data cleaning. For example, students can bring about unique variance because of their background characteristics (e.g., demographics). Problems can introduce additional variance through their own measurement properties (e.g., difficulty). The adaptive ITS design can also induce extra variance because of stochastic problem assignments. The previous studies generally lacked attention to these factors and attributed all observed variance to students' latent states. In this study, we explicitly address these complications and suggest practical model-fitting strategies that accommodate the assumptions, estimation constraints, and design features of ITS.

Three models are considered for modeling flow: the (i) latent class model (LCM), (ii) latent transition model (LTM), and (iii) hidden Markov model (HMM). These models can be applied to multivariate cross-sectional time series (i.e., pooled data) and can address variance from contextual variables. The models differ in the specific ways of describing the observable and latent variables (e.g., permissible indicators, assumptions on the latent transition, covariate effects). In this study, we suggest workable solutions for each model,

addressing the related assumptions and constraints. As a working example, we use data from Cognitive Tutor Algebra (CTA; Pane, Griffin, McCaffrey, & Karam, 2014).¹ Using log data from a particular time period, we show that the suggested models can help infer learning flow. We give a didactic illustration of data preparation, model adjustment, and analysis steps, and show that the suggested strategies provide a reasonable solution for fitting the models and the outcomes give relevant information about flow. Based on the experiential analysis of CTA data, we perform a comparative model assessment and discuss relative strengths and limitations of the models in describing the ITS data.

The current study makes a number of important distinctions from the studies present in the ITS literature. Many of the existing studies examined engagement or its operationalized behaviors (e.g., Baker et al., 2012; Calvo, D’Mello, Gratch, & Kappas, 2015; Petrovica, Anohina-Naumeca, & Ekenel, 2017). They also tended to use biometric data that require separate sensors (e.g., video, eye tracker, heart rate sensor). Engagement is a complex multifaceted construct that involves multiple domains (Fredricks, Blumenfeld, & Paris, 2004; Fredricks, Filsecker, & Lawson, 2016; Martin, 2012; Redmond, Heffernan, Abawi, Brown, & Henderson, 2018) and is generally difficult to assess based only on the interaction data. Studies that focused on the particular aspects of engagement also bear only topical implications and lack consensus on addressing the identified features. Additionally, biometric devices are not typically available in usual educational settings and they can cause discomfort and interrupt learning activities. This study, by contrast, focuses on a more measurable and relevant construct, flow, and uses behavioral data that are readily accessible from ITS. Flow has been known to have a clear conception in educational research and is considered an antecedent of engagement (Boniwell, 2008). By focusing on flow, we examine the origin of (dis)engagement and inform the functioning of ITS in achieving the optimal learning—the information that can help improve future ITS design.

¹ CTA is a predecessor of MATHia. We retain the name of CTA since the data were obtained from 2007–2008 CTA administration.

Aside from the distinction in the construct and the input data, the current study gives more careful consideration to the effects of structural variables. Previous studies seldom contemplated contextual variables and applied machine learning or data mining techniques to student-level data (e.g., decision tree, naïve Bayes, support vector machine, neural networks). These procedures do not take into account systematic effects from measurements or populational variables, and attribute all observed variance to students' latent states. The inferences drawn from these analyses can be fallible if structural variables have a systematic impact on the outcome data, which is often the case in educational settings. In this study, we explore ways to incorporate contextual variables (e.g., demographics, stochastic problem assignment, measurement effects) when modeling the flow such that variance induced by these factors can be adequately addressed. When applying the models, we also address the unique characteristics of ITS design, such as stochasticity and temporality of indicators, and make practicable adjustments for each model.

The remainder of this paper is organized as follows. In Section 2, we introduce CTA and discuss ground analysis settings. We provide key information about the CTA data, variables used, and prearrangements made for the analyses. Sections 3—5 present analysis performed under each latent variable model. We discuss model formulation, data rearrangement, model-fitting process, and corresponding results. Section 6 concludes with a summary of the findings and implications for general settings. We also discuss future considerations and refinements that merit attention in developing the latent variable models for the ITS data.

2. Cognitive Tutor Algebra

2.1. Data

The study uses Cognitive Tutor Algebra (CTA) I as an experimental example of ITS. The evaluation data were collected in 2007-08 as a part of an effectiveness study at the RAND Corporation (Pane et al., 2014). The raw data contained observations from $N =$

2860 students that received tutoring between July 2007 and May 2008. The tutoring occurred during regular curricula under the supervision of teachers. The system contained a total of 637 problems across 106 sections nested within 27 units (e.g., algebra level 1, level 2; equation solver level 1, level 2). Across the study period, students received 276.826 problems ($SD = 193.642$), 43.900 sections ($SD = 32.523$), and 8.927 units ($SD = 6.746$) on average. The specific problem sets and the order of problems differed by students, teachers, and school districts. While most of the problems were prompted by an artificial tutor, teachers could reassign students to different sections and the system could also promote students to a new section if a student reaches a maximum number of problems.

2.2. Preparation

A number of challenges need to be addressed for applying the latent variable models to the CTA data. In CTA, students received different units and worked on a small portion of problem sets in various sequences over long periods of time. For examining flow, it was necessary to clean and reform the raw data. In this study, we focused on a particular unit and examined students' workings on single days. We chose one of the most frequently assigned units, *equation solver level 1* (*es1* hereinafter), and extracted interaction data that occurred on the same days. If a student worked on *es1* across multiple days, we picked the day the student attempted most problems and examined the student's flow progression during the day. Similarly, if a student worked on multiple units on the same day, only the observations from *es1* were examined to regulate variance from other units and problems.

The problem sets were deliberately fixated on *es1* because the problems within the unit showed relatively similar measurement properties and were anticipated to induce minimal variance in the indicating data (see Appendix A). Although we carefully prepared the data to have homogeneous measurement properties, we also additionally addressed the measurement noninvariance when a model allowed modification (e.g., random effect). In addition to confining the problems to the same unit, we also constrained the times of tutoring on the same days. We surmised that, if a student works over multiple days, the

student can display substantially different patterns in the number and transition of flow states. Since we view flow as a dynamic construct that can change in a short time period, we focused the analysis on the same single days.

2.3. Variables

Focusing on the *es1* unit resulted in subset data of $N = 2219$ students. The students in the final data attempted 50.236 problems on average ($SD = 16.102$) with a minimum of four and maximum of 151. For inspecting flow, we examined three indicator variables: the interaction time, the number of erroneous attempts, and the number of hints requested. Each indicator variable was transformed to meet the constraints of the calibration programs. For example, the timing variable was placed on the log metric to approximate normality. The count variables were used as observed or categorized into three ordinal categories (i.e., none, one, and more than one) and modeled via Poisson or proportional odds models. Along with the interaction indicators, we also made use of available covariates when inferring the state membership and transition behaviors. The covariates considered include: Pre- and Gain test scores on the standardized test, Race (0 = White and Asian, 1 = Black and multiethnic, 2 = Hispanic and native America), Sex (0 = Female, 1 = Male), and whether a student was enrolled in a free lunch program (0 = No, 1 = Yes).

2.4. Analysis

The cleaned data were analyzed in three stages (see Table 1 for comparison). We first applied latent class models to identify subgroups that underlie the interaction data and examined if the identified groups imply different states of learning flow. For fitting the models, we used the sample data to address the effects of the structural variables (e.g., populational characteristics, problem effects). The second stage analysis applied latent transition models and examined transition of latent states over time. We in particular focused on the development between problem sets to determine interaction patterns across different tutoring stages. The models were again fit to the sample data such that the effects of the structural variables can be accounted for. The last stage analysis applied

Table 1

Analysis Settings

Analysis	Model	Data	Variable	Measurement Invariance
Flow profiles	Random-effect LCM	Sample-level data, aggregated by knowledge components	log time, nhint and nerr on the ordinal scale (0, 1, 2)	modeled through random-effect terms in the structural and measurement models
Flow transition over tutoring stages	Random-intercept LTM	Sample-level data, aggregated for tutoring stages	log time, nhint and nerr on the ordinal scale (0, 1, 2)	modeled through random-intercept latent factors in the measurement model
Flow transition across problems	HMM	Student-level data with problem-level observations	log time, nhint, nerr	assumed weak measurement invariance and addressed possible impact of measurement noninvariance

Note. time = Task interaction time. nhint = Number of hints requested. nerr = Number of erroneous attempts. Throughout the analyses, the same set of covariates were used: Pre/Gain scores on a standardized test, Race, Sex, Enrollment on the free lunch program.

hidden Markov models. Unlike the other two analyses, HMMs were fit to the student-level data to investigate the state development over individual problems. Since problems were tailored to the student's skill levels, we assumed that the problems would exhibit weak measurement invariance and focused the analysis on each student-level interaction data. We however noted possible consequence of ignoring the measurement effects and made conscious efforts to reduce such effects.

Throughout the analyses, the models were fit assuming a few latent states. Since the study was primarily interested in modeling discrete flow states, we gave more attention to the stability of the extracted states, indicator patterns exhibited under each state, and how they evolve over different tutoring times. Where appropriate, the latent states were operationalized as two discrete states of flow and deviation-from-flow. Our supplementary analysis suggests that allowing more states tends to result in the decomposition of a flow state, indicating different problem-solving modalities. Although untangling the flow state can help distinguish different operation modes, the identified features will require

subjective interpretation and are difficult to validate beyond face validity. We therefore limit our attention to clear bimodality of flow—in- and out-of-flow. In the following sections, we discuss analysis conducted under each model, including the model formulation, analytic strategies for CTA, and corresponding results.

3. Profiling Flow

Following assumptions were made to analyze the CTA data. Across the analyses, we assumed that latent flow state, S , drives student's interaction with the CTA. The flow state is a discrete nominal state that belongs to a finite latent space, $\mathcal{S} \equiv \{1, \dots, M\}$. The specific value of S can vary by students and times. A student i ($= 1, \dots, N$) can have a series of flow states that evolve over T events, $\mathbf{S}_i = (S_{it} : t = 0, \dots, T)$, where events correspond to problems.² Note that in ITS students may receive different problem sets and T can vary by students. We define T as $\max(T_i : i = 1, \dots, N)$ and use it as a generic notation for the time length. As a student exhibits a distinct flow state at time t , S_{it} , a set of observable indicating variables, $\mathbf{Y}_{it} = (Y_{ijt} : j = 1, \dots, J)$, will exteriorize the latent state (e.g., interaction time, behavioral frequency). A collection of indicator outcomes over time, $\mathbf{Y}_i = (\mathbf{Y}_{it} : t = 1, \dots, T)$, then forms multivariate cross-sectional time series. If covariates are available for students, $\mathbf{Z}_i = (Z_{ik} : k = 1, \dots, K)$, they can be integrated when modeling the latent states or transition behaviors. The current study applies time-invariant covariates that remain constant over time (e.g., gender, ethnicity); the analyses suggested below however can be easily extended to accommodate time-variant covariates.

Several assumptions are made for modeling flow. First, we assume that elements of $\mathbf{Y}_{it}$ (i.e., the indicators at time t) are conditionally independent given the student's latent state, S_{it} . Additionally, the sequence of latent states is assumed to follow a first-order Markov process such that the later state depends only on the immediately preceding state, i.e., $P(S_t = s_t | S_0 = s_0, S_1 = s_1, \dots, S_{t-1} = s_{t-1}) = P(S_t = s_t | S_{t-1} = s_{t-1})$. Lastly, we

² The event times may be defined as problem sets depending on the analysis.

assume that the problems assigned to the students are weakly invariant across the students and times (i.e., measurement invariance). When a model can be modified to accommodate measurement noninvariance, we address the variance by making relevant alterations.

3.1. Random-Effect Latent Class Model

The first stage analysis applied the latent class models to examine the heterogeneity in the interaction data. We especially examined if the heterogeneity can be characterized by distinct states of learning flow. At each measurement point, we classify students into distinct latent states given the cross-sectional data,

$\mathbf{Y}_t = (Y_{ijt} : i = 1, \dots, N; j = 1, \dots, J)$. A sequence of latent states over time, $\mathbf{S}_i = (S_{it} : t = 0, \dots, T)$, then gives a latent profile a student went through over the course of tutoring and defines a unique latent class.

The formulation of LCM requires two sub-models: (i) a structural model that describes the probability of a latent state, $P(S_{it} = m)$ ($m = 1, \dots, M$), and (ii) a measurement model that describes the conditional probability of indicators given a latent state, $P(\mathbf{Y}_{it} | S_{it} = m)$. Unlike the other two models that will be discussed below, LCM does not model transition of latent states. It instead assumes that latent states evolved independently over time and models the entire sequence of observed data using a latent state profile. Below we discuss the formulation of each component of LCM aimed for ITS/CTA data.

The structural model describes the probability that a student exhibits a certain latent state at any moment. In regular settings, this probability is modeled under within-class homogeneity. That is, students within the same class are expected to give homogeneous performance when solving the problems. In CTA (and in many cases of ITS), this assumption is generally not tenable because students receive different problems according to their skill levels and progression. The difference in the problem sets and the assignment order can induce extra variance to the outcome data and the observed data may no longer be independent when conditioned on the latent states. A common approach

to tackling this extra variance is to allow random variation when modeling the outcome probabilities (e.g., Qu, Tan, & Kutner, 1996). A random-effect term is included in the structural model so that the heterogeneity in the outcome probabilities can be explained by the latent state as well as the random effect.

In the current setting, the random-effect LCM can be formulated as follows. Let η_i model the idiosyncratic effect of student i on the state probability. If covariates are available, the probability of a latent state m can be modeled as

$$P(S_{it} = m | \eta_i, \mathbf{Z}_i = \mathbf{z}_i) = \frac{\exp(\beta_{m0} + \boldsymbol{\beta}_m^\top \mathbf{z}_i + \eta_i)}{\sum_{l=1}^M \exp(\beta_{l0} + \boldsymbol{\beta}_l^\top \mathbf{z}_i + \eta_i)}, \quad (1)$$

where β_{m0} determines the conditional probability of the latent state m when $\mathbf{Z}_i = \mathbf{0}$ and $\eta_i = 0$, and $\boldsymbol{\beta}_m$ models the covariate effects on the logit. To identify the model parameters, we assume that $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2)$ with $\sigma_\eta^2 = \text{Var}(\eta_i; i = 1, \dots, N)$ modeling the magnitude of extra variance from the individuals (i.e., beyond covariate effects).

The measurement model describes the probability of an indicator given the latent state. Since the problem stimuli differ in the measurement properties, we again introduce problem-level random-effect terms when modeling the indicator variables. A measurement model for an indicator j is parameterized as

$$P_j(Y_{ijt} | S_{it} = m, \delta_{jt}) = f_j(\boldsymbol{\psi}_{jm}, \delta_{jt}), \quad (2)$$

where f_j takes a functional form specific to the indicator j , $\boldsymbol{\psi}_{jm}$ contains parameters of f_j , and δ_{jt} models the idiosyncratic effect of problem t on the indicator j . The functional form of f_j is defined according to the type of the variable. For binary variables, common practice is to use the Bernoulli distribution with a probit or logit link. The ordinal variables are typically modeled by cumulative probability functions such as proportional-odds models, adjacent-categories, or continuation-ratio logit models (Agresti, 2012). The count and continuous variables can be modeled by each Poisson regression and Gaussian model. The random-effect term, δ_{jt} , in (2) is parameterized for each j such that $\delta_{jt} \sim \mathcal{N}(0, \sigma_{\delta_j}^2)$ and $\sigma_{\delta_j}^2 = \text{Var}(\delta_{jt}; t = 1, \dots, T)$ models the variance across the problems in indicator j .

Integrating the two sub-models, the random-effect LCM becomes a finite mixture model:

$$P(\mathbf{Y} | \mathbf{Z}, \boldsymbol{\eta}, \boldsymbol{\delta}) = \prod_{i=1}^N \prod_{t=1}^T \sum_{S_t \in \mathcal{S}} P(S_t | \eta_i, \mathbf{Z}_i) P(\mathbf{Y}_{it} | S_t, \boldsymbol{\delta}_t), \quad (3)$$

where $\mathbf{Y}$, $\mathbf{Z}$, and $\boldsymbol{\eta}$ each denote an array of variables for all students, and $\boldsymbol{\delta} = (\boldsymbol{\delta}_t : t = 1, \dots, T)$ where $\boldsymbol{\delta}_t = (\delta_{jt} : j = 1, \dots, J)$. Note that the model (3) is specifically designed for ITS/CTA data, that is, the multivariate cross-sectional time series where students receive fractions of problems in different sequences. For estimating the parameters of (3), we apply maximum likelihood (ML) or Bayesian inference methods. The ML inference is computationally efficient, but can result in zero variance estimate (Gelman et al., 2013, p. 313).³ The Bayesian inference is computationally intensive, but gives more stable estimates once converged. In this study, we apply an ML estimator when comparing the candidate models and apply Bayesian inference when deriving the estimates of a final model. The ML estimation is performed using an L-BFGS algorithm (Nocedal & Wright, 2006) in Stan, called from R (R Core Team, 2020). The Bayesian estimation is implemented based on “No U-Turn” Markov Chain Monte Carlo (MCMC) sampler (Hoffman & Gelman, 2014) in Stan. The details of the estimation routine, including the discussion on the label switching, are presented in Appendix B.

3.2. Profiling Flow in CTA

For evaluating flow in CTA, we considered one- and two-state models. The one-state model assumes that students performed homogeneously throughout; the two-state model assumes that students behaved differently while learning with the tutor. The LCM formulation generally followed (1) and (2), but some adjustments were made in the measurement model to accommodate the CTA design. In CTA, students receive different numbers of problems in different orders. The difference in the problem assignments and the variance in the problem properties need to be addressed when modeling the effects of

³ The random-effect LCM can be seen as a hierarchical model with the random-effect terms being modeled at the higher-order level.

measurements. Our strategy in this study is to group the problems by homogeneous skill sets—termed *knowledge components* (e.g. Ritter, Anderson, Koedinger, & Corbett, 2007)—and examine latent states that underlie each knowledge mastery. Reshaping the data this way helps homogenize the impact of measurement stimuli and also alleviates computational load as it reduces the time dimensionality from the number of problems to uniformly the number of knowledge components. Applying this strategy to CTA, we identified a total of 18 knowledge components in *es1* and obtained 2219×18 ($N \times T$) sample data. We then fit the random-effect LCMs, allowing measurement noninvariance in each indicator variable (i.e., δ_{jt} ($j = 1, \dots, 3, t = 1, \dots, 18$)). The random-effect terms from the three indicators were jointly modeled by a trivariate normal distribution:

$$\boldsymbol{\delta}_t = (\delta_{1t}, \delta_{2t}, \delta_{3t})^\top \sim \mathcal{N}_3(\mathbf{0}, \boldsymbol{\Sigma}_\delta).$$

The measurement model for each indicator was specified as follows. The interaction time, Y_{i1t} , was modeled by a log-normal density:

$$\log Y_{i1t} | S_{it} = m, \delta_{1t} \sim \mathcal{N}(\delta_{1t} + \mu_{1m}, \sigma_{1m}^2) \quad (4)$$

so that both the mean and standard deviation of the distribution vary with the latent state. The count variables—the number of errors, Y_{i2t} , and the number of hints requested, Y_{i3t} —were treated as ordinal variables, with categories zero, one, and more than one, and modeled via ordinal logistic regression. For example, the probability of the number of errors was modeled as

$$P(Y_{i2t} = h | S_{it} = m, \delta_{2t}) = \begin{cases} 1 - \text{logit}^{-1}(\mu_{2m} + \delta_{2t} - c_{21}) & h = 0 \\ \text{logit}^{-1}(\mu_{2m} + \delta_{2t} - c_{21}) - \text{logit}^{-1}(\mu_{2m} + \delta_{2t} - c_{22}) & h = 1 \\ \text{logit}^{-1}(\mu_{2m} + \delta_{2t} - c_{22}) & h > 1 \end{cases} \quad (5)$$

with the location parameter μ_{2m} varying across the latent states, and the step parameters satisfying $c_{22} > c_{21} > 0$. As above, δ_{2t} is the problem-level random effect for errors ($j = 2$).

1 The inverse-logit function is $\text{logit}^{-1}(x) = 1 - \exp(-x)^{-1}$. The number of hints was modeled
 2 analogously with μ_{3m} , δ_{3t} , c_{31} , and c_{32} each replacing μ_{2m} , δ_{2t} , c_{21} and c_{22} .

3 3.3. Results

4 **Model comparison.** Table 2 reports fit statistics of four LCMs fitted under
 5 different numbers of latent states and student effects. The results were obtained from the
 6 ML estimator. The one-state model in the first line represents the hierarchical multivariate
 7 regression that regresses the three indicators on the knowledge-level random effects. Note
 8 that the one-state model assumes a single homogeneous latent state and has no structural
 9 component. It assumes that all students exhibited the same state and there is no individual
 10 difference in the performance outcome. The following lines give results for the two-state
 11 models. The models differ in the way they address the student effect. The *HomStud* model
 12 assumes that students in the same state had the same state probability, $P(S_{it} = 1) = \beta_0$.
 13 The *StudEff* model accounts for individual differences in the state probability,
 14 $P(S_{it} = 1 | \eta_i) = \beta_0 + \eta_i$. The *StudCov* models the state probability as a function of both
 15 the student random effects and student-level covariates, $P(S_{it} = 1 | \eta_i, \mathbf{z}_i) = \beta_0 + \boldsymbol{\beta}^\top \mathbf{z}_i + \eta_i$.

16 In Table 2, all information criteria preferred the two-state models over the one-state
 17 model, suggesting heterogeneity in the student behaviors. The criteria that heavily
 18 penalize complexity (i.e., CAIC, BIC, ABIC) recommended the simplest two-state model,

Table 2

Fit Statistics of LCMs

M	Condition	df	-2LL	AIC	CAIC	BIC	ABIC
1		62	-241475	483074	483672	483672	483475
2	HomStud	67	-219935	440005	440650	440650	440437
2	StudEff	2285	-216607	437783	459802	459802	452540
2	StudCov	2291	-216576	437735	459811	459811	452530

Note. M = Number of states at each time point. df = Number of free parameters. LL = -2 log likelihood. AIC = Akaike information criterion (Akaike, 1973). CAIC = Corrected AIC (Burnham & Anderson, 2002; Sugiura, 1978). BIC = Bayesian information criterion (Schwartz, 1978). ABIC = Adjusted BIC (Sclove, 1987).

HomeStud. The AIC and likelihood statistic preferred *StudCov*. A likelihood ratio test between *StudCov* and *StudEff* suggested *StudCov* ($\chi^2 = 62$, $df = 6$, $p < .001$) a better fitting model. The estimated posterior state probabilities also indicated that there exists clear heterogeneity across the student demographic profiles. Given these observations, we chose the *StudCov* model as the final latent class model. In the subsequent analysis, we examine *StudCov* more carefully to investigate the patterns across the indicator variables.

Measurement model. After identifying the final model, *StudCov*, we refit the model applying the Bayesian estimator, MCMC. The ML estimation, though computationally efficient, yielded zero variance for the student random effects (i.e., $\hat{\sigma}_\eta^2 = 0$) despite the strong indication of heterogeneity in the observed data. The Bayesian estimation, although required longer computation time, converged properly, yielding nonzero variance estimates and other point estimates essentially identical to the ML estimates.⁴ Below we present results of *StudCov* estimated from MCMC.

Table 3 reports parameter values estimated for the measurement model of *StudCov*. The first four lines give estimates of the distribution parameters for each indicator. The last line reports initial state probabilities $E[P(S_{it} = m)]$ ($m = 1, 2$). The results suggest that, about 81.02% of times, students worked under State 1, and the remaining 18.97% under State 2. The difference statistics in the last column suggest that students working in the second state requested more hints, made more errors, and spent more time than did the students working in the first state. The students in the second state also showed greater variation in the interaction time. Recall that, in CTA, problems were adaptively assigned according to the student's skill levels and students are generally expected to flow while learning through problems. Given the design of CTA, it can be concluded that State 1 represents the state of flow and State 2 represents the deviation from the flow. In State 1, students tended to progress in a timely manner, exerting adequate effort. Students in State

⁴ The ML estimation took about 15 minutes to converge and Bayesian estimation based on MCMC took over five days when run in parallel on an AMD Ryzen Threadripper 3.5 GHz 16-Core Processor.

Table 3

Parameters of the Measurement Model of the Final LCM

Indicator	Parameter	State 1	State 2	Difference
Time	μ_{1m}	-.500 (.117)	.992 (.117)	-1.493 (.008)
	σ_m	.757 (.003)	.838 (.005)	-.080 (.005)
Error	μ_{2m}	-2.092 (3.551)	2.101 (3.551)	-4.193 (.038)
	c_{21}	2.126 (3.557)		
	c_{22}	3.807 (3.557)		
Hint	μ_{3m}	-6.131 (3.807)	6.145 (3.804)	-12.275 (3.046)
	c_{31}	8.312 (3.881)		
	c_{32}	9.888 (3.882)		
	$P(S_0)$	.811 (.003)	.189 (.003)	

Note. Parameters μ , σ , and c are defined in (4) and (5); standard errors are in parentheses. All estimated differences (the final column) were significant at $p = .001$; p -values for estimates in other columns were not computed.

2 tended to exhibit prolonged learning behaviors. Note that the current results do not inform the cause of deviation; we only surmise that multiple factors could have attributed to the deviating state, for example, fatigue, excessive difficulty, disengagement, gaming, or frustration with the problems.

To understand the degree of heterogeneity between the states, we compared the difference statistics with the between-knowledge differences, as parameterized by the standard deviations (SDs) of the knowledge-level random effects, δ . The random-effects SDs were estimated as $\hat{\sigma}_{\delta_1} = .497$ (95% credible interval [.366, .676]) for time spent, $\hat{\sigma}_{\delta_2} = 1.387$ (95% CI [1.019, 1.858]) for errors, and $\hat{\sigma}_{\delta_3} = 1.529$ (95% CI [1.069, 2.128]) for hints. These estimated SDs were much smaller than the corresponding differences between the states (i.e., -1.493, -4.193, -12.275), suggesting that the difference between the states was much more important than the differences between the knowledge components. The correlations between the random effects δ_1 , δ_2 , and δ_3 implied by $\hat{\Sigma}_{\delta}$, were all fairly high: $\text{Cor}(\delta_1, \delta_2) \approx .969$, $\text{Cor}(\delta_1, \delta_3) \approx 0.711$, and $\text{Cor}(\delta_2, \delta_3) \approx .652$, suggesting that problems that involved more hint requests and errors entailed longer time.

Covariate effects. Table 4 reports estimated covariate effects. When fitting the LCM, we treated State 1 (i.e., the flow state) as a reference state and estimated loadings of covariates on State 2. The reported coefficients show the effects of covariates on the log-odds of a student working in State 2. A large value means that a student was more likely to work in the out-of-flow state. The results show that students with higher pretest scores and gain scores (i.e., whose standardized test scores increased the most from the beginning to the end of the study) were more likely to work problems in State 1, while students with lower pretest and gain scores tended to work in State 2. Black, multiethnic (RaceBN), Hispanic and Native (RaceHN) students, and to a lesser extent males, were more likely to work in State 2 than White or Asian students and females. Students who were eligible for free or reduced-price lunches tended to work in State 1 though the trend was somewhat weak. Standard deviations of the student random effects, σ_η , were estimated as .590 (95% CI [.560, .622]), which is larger in magnitude than any of the covariate coefficients in Table 4, and also larger than the standard deviation in the student-level log-odds of working a problem in State 1 predicted by covariates, $SD(\hat{\beta}^\top \mathbf{Z}) \approx .268$. This suggests that there may be important unmeasured student characteristics that predict the latent states, which we speculate to be the transitioning of latent states.

Table 4

Loading of Covariates on the Out-of-Flow State

Covariate	Est	SE	<i>p</i>
Pretest score	-.258	.020	.000
Gain score	-.132	.020	.000
RaceBM	.274	.051	.000
RaceHN	-.195	.040	.000
Sex	.081	.032	.016
Free lunch	-.020	.036	.579

Note. RaceBM = Black & Multiethnic. RaceHN = Hispanic & Native America. Sex: Male = 1, Female = 0. Free lunch: Yes = 1, No = 0.

4. Flow Across Tutoring Stages

The latent class analysis revealed that students tended to behave differently when working on the CTA. The ensuing analysis examined how the distinct interaction patterns evolved over the tutoring sessions. For the analysis, we partitioned the observed data and applied LTMs to track transition of latent states over time.

4.1. Random-Intercept Latent Transition Model

As with the LCM, the LTM includes the structural and measurement models but additionally includes a transition model to describe transition between latent states. The structural model takes a similar form with (1) but models initial state probabilities only, $P(S_0 = m)$ ($m = 1, \dots, M$), or $P(S_{i0} = m | \mathbf{Z}_i = \mathbf{z}_i)$ if covariates are applicable. The transition model then describes the ensuing latent states as a function of the preceding state(s). For example, the first-order Markov model assumes

$P(S_t | S_0, S_1, \dots, S_{t-1}) = P(S_t | S_{t-1})$ such that the current latent state depends only on the immediately preceding state. The functional form of $P(S_t | S_{t-1})$ commonly takes multinomial logistic regression. When covariates are available, the transition probability from state m to state m' can be modeled as

$$P(S_{it} = m' | S_{i(t-1)} = m, \mathbf{Z}_i = \mathbf{z}_i) = \frac{\exp(\gamma_{0m'} + \gamma_{mm'} + \beta_m^\top \mathbf{z}_i)}{\sum_{l=1}^M \exp\left(\gamma_{0l} + \sum_{l'=1}^{M-1} \gamma_{ll'} d_{l'} + \beta_m^\top \mathbf{z}_i\right)}, \quad (6)$$

where $\gamma_{0m'}$ ($m' = 1, \dots, M-1$) denotes an intercept of the logit of the transition probability; $\gamma_{mm'}$ ($m, m' = 1, \dots, M-1$) models change in the logit when the state moves from m to m' ; and β_m models the effect of covariate on the logit. The $d_{l'}$ in the denominator is a dummy variable that indicates the first $(M-1)$ states (i.e., $l' = 1, \dots, M-1$; the last state is a reference category).

The measurement model describes conditional probabilities of manifest indicators as a function of the latent state. One of the important assumptions for tracking state transition is invariance of measurements over time. That is, problems assigned across time must have constant effects on the performance probability so that variance in the observed

performance can be attributed to the corresponding underlying latent state. As alluded to above, this assumption is generally untenable in CTA because problems were adaptively administered according to the students' skill levels. To address the variance in the problem assignments and measurement properties thereof, we introduce a random intercept factor to the measurement model (B. Muthén & Asparouhov, 2020). The random intercept is parameterized for each student such that the variance in the measurements can be accounted for by the student's own characteristics. For each indicator j , the measurement model is formulated as

$$P_j(Y_{ijt} | S_{it} = m, \eta_i) = f_j(\psi_{jm} + \lambda_j \eta_i), \quad (7)$$

where f_j denotes the functional form of the probability measure for indicator j , ψ_{jm} gives the kernel function indicating the effect of state m on j , η_i denotes the student-specific random intercept, and λ_j models the extent to which students induce extra variance to indicator j . The random-intercept factor η_i is assumed to follow $\mathcal{N}(0, 1)$ so that λ_j models the size of the extra variance in each measurement. The parameterization in (7) can be seen as decomposing the observed variance into within- and between-subject variance. The slope coefficient λ_j models the average magnitude of between-subject variance across the problems.

Integrating the sub-models, the random-intercept LTM (RI-LTM) describes the marginal probability of outcome data as

$$P(\mathbf{Y} | \mathbf{Z}) = \prod_{i=1}^N \sum_{S_t \in \mathcal{S}} P(S_0 | \mathbf{Z}_i) \left(\prod_{t=1}^T P(S_{it} | S_{i(t-1)}, \mathbf{Z}_i) \right) \left(\prod_{t=1}^T \prod_{j=1}^J P(Y_{ijt} | S_{it}, \eta_i) \right). \quad (8)$$

For estimating (8), we apply regular marginal ML estimation with the expectation-maximization algorithm. The estimation is performed in the statistical software *Mplus* (L. K. Muthén & Muthén, 2017). We note that *Mplus* does not support intensive time-series data, and data cleaning is needed to meet the constraints of the estimation algorithm. Section 4.2 presents practicable strategies for fitting the RI-LTMs to ITS/CTA data.

Before we close the model formulation, we note some differences between the random-effect LCM and RI-LTM. In LCM, the heterogeneity in the problems is modeled by both the structural and measurement models. The random effect in the structural model accounts for difference in the problem assignments (and thus their effect on the state probability). The random effect in the measurement model accounts for difference in the measurement properties of the problems. The dual use of random effects in LCM is partly motivated by the absence of a transition model. The LTM, on the contrary, explicitly models the transition of latent states and attempts to model the heterogeneity in the problems through the measurement model under the assumption of conditional independence of measurements given the history of latent states (i.e., the current latent state and sequence of states). That is, the random-effect terms in LCM can be seen as modeling extra variance in both the between- and within-persons (through each structural and measurement model) whereas the random-effect term in LTM models the extra between-person variance in the measurements. The constant slope λ_j in the measurement model of RI-LTM implies the assumption of measurement invariance; if measurements induce substantial variance across time (i.e., measurement noninvariance), the extra variance can be modeled by allowing time-variant slopes, λ_{jt} (see our analysis of CTA data below). We also note that the two models assume different parameterizations for the random effects. In LCM, η_i is assumed to follow $\mathcal{N}(0, \sigma_\eta^2)$ with unit slope whereas in RI-LTM, $\eta_i \sim \mathcal{N}(0, 1)$ with distinct slopes. The latter is to allow flexibility when there exists significant measurement noninvariance.

4.2. Flow Transition in CTA

As with the preceding analysis, we used the data from one unit, *es1*, observed on the same days. The observed data consisted of intensive time series with a maximum of 151 problems. The statistical programs currently available for the latent transition analysis (e.g., Mplus, PROC LTA) assume longitudinal data with a small number of time points (e.g., 10 at most) and have limited capacity in modeling the intensive time-series data. For

modeling the CTA data, it was necessary to alter the raw data and reduce the number of measurement points. Our strategy in this study is to segment the tutoring into several sessions and aggregate the problems within each segment to create small sets of problem series. Since the problems in the evaluation data uniformly measured the same equation-solving skills, we theorized that the flow states could be reasonably evaluated under the weak measurement invariance. For the event that the problems induce significant measurement noninvariance, we also applied more general RI-LTMs that allow time-variant slopes for the random intercepts.

The data partition was achieved as follows. We first determined the number of evaluation points as three, four, and five, balancing the granularity of description and convergency of model estimation. We then split data from each student into three-, four-, and five-point time series by allocating approximately equal numbers of observations to each series. For example, when a student received 56 problems, we sequentially aggregated (18, 18, 20) problems to create four-point time series. Within each partition, observations were averaged and placed on the calibration scale. The timing variable was placed on the log scale after the average. The count variables were placed on the ordinal scale after being averaged (i.e., indicating no, one, and more than one error/hint). The analyses of the three aggregated data sets suggested that the trends in the flow states are generally consistent across the partitions. The results differed only in the model convergence and the specific parameter estimates. Below we present results from the four-point time-series data as a representative example.

4.3. Results

Model comparison. Tables 5–8 report fit statistics of the LTMs evaluated under different scenarios. In Table 5 the outcomes of the one- and two-state models are compared. The comparison shows a clear preference for the two-state model. The adjusted likelihood ratio test (Lo, Mendell, & Rubin, 2001) indicated significantly better fit for the two-state model ($\chi^2 = 23320.293$, $df = 6$, $p < .001$), implying heterogeneity in the students'

Table 5

Fit Statistics of LTMs: One- vs. Two-state Models

M	df	LL	AIC	CAIC	BIC	ABIC
1	15	-20727.74	41485.48	41485.70	41571.05	41523.4
2	21	-8995.53	18033.06	18033.48	18152.86	18086.14

Note. M = Number of states at each time point. df = Number of free parameters. LL = -2 log likelihood. AIC = Akaike information criterion. CAIC = Corrected AIC. BIC = Bayesian information criterion. ABIC = Adjusted BIC.

behaviors. We note that more complex LTMs with more latent states generally demonstrated better goodness of fit. The fitted outcomes however indicated similar patterns with respect to the out-of-flow. For example, when the LTM was fit with three states, the state identified as out-of-flow in the two-state model remained the same in the three-state model, and the flow state identified in the two-state model was separated into two different states in the three-state model. It appeared that increasing the number of latent states in LTMs disintegrates the flow state, with different behavioral strategies under the regular problem-solving mode. The inference on the deviant state remained the same across the models that assumed different numbers of latent states.

The ensuing analysis examined variants of the two-state LTM to determine the most appropriate model for CTA. We investigated the models that differ in the assumptions and constraints of the model: (i) measurement invariance of problem sets, (ii) significance of covariates, and (iii) stationarity of transition probabilities.

Measurement invariance. Table 6 reports fit statistics of the models that assume different degrees of measurement invariance. The invariance of the measurement stimuli was evaluated by comparing the three models: (i) the LTM that assumes longitudinal measurement invariance (labeled as MI), (ii) LTM that includes the random intercept (labeled as RI), and (iii) LTM that assumes time-varying effects of the random intercept (labeled as VRI). Recall that the random intercept was devised to describe different amounts of between-subject variance at each time. The loading of an indicator on the

Table 6

Fit Statistics of LTMs: Measurement Invariance

Condition	<i>df</i>	LL	AIC	CAIC	BIC	ABIC
MI	21	-8995.53	18033.06	18033.48	18152.86	18086.14
RI	24	-8700.68	17449.37	17449.91	17586.28	17510.03
VRI	33	-8674.55	17415.11	17416.14	17603.37	17498.52

Note. *df* = Number of free parameters. LL = -2 log likelihood. AIC = Akaike information criterion. CAIC = Corrected AIC. BIC = Bayesian information criterion. ABIC = Adjusted BIC. MI = LTM with measurement invariance. RI = LTM with the random intercept. VRI = LTM with the random intercept with time-varying effects.

random intercept shows the magnitude of variance caused by the measurement stimuli at the time of evaluation. The nonzero loading coefficient suggests that the problem sets induced unignorable variance in the students' behaviors. If the estimated loading coefficients are approximately equivalent across time, it means that the problems entailed similar amounts of variance across time and the constant variance can be modeled using time-invariant loading coefficients. If the estimated loading coefficients differ substantially over time, it implies that the problem sets entailed different amounts of variance across time and they must be modeled by the time-contingent loading coefficients.

In Table 6, comparison of LTMs with and without the random intercept reveals that the random-intercept model achieved significantly better fit ($\chi^2_{(MI,RI)} = 589.69$, $df = 3$, $p < .001$). This means that the problem sets administered across the tutoring induced nonzero variance in the outcome variables. When the indicators were allowed to have different loadings over time (i.e., RI vs. VRI), the LTMs yielded improved log-likelihoods and smaller AIC, CAIC, and ABIC values. Among the three indicators, the timing variable showed the largest variation across time (i.e., measurement noninvariance), followed by the number of errors, and the number of hints. All in all, the outcomes of the three modeling conditions suggested that the indicators induced different amounts of measurement variance across time and it is desirable to include the random intercept and model its effects allowing time- and indicator-specific coefficients.

Covariate effects. Table 7 reports fit statistics of the models that assume differential covariate effects. The results suggest that including covariates substantially improved the goodness-of-fit. All pairwise comparisons indicated that LTMs attained significantly higher log-likelihoods when they included the covariates (e.g., $\chi^2_{(MI,MI-C)} = 133.07$, $df = 6$, $p < .001$). The information criteria similarly indicated better fit for the covariate-integrated models, yielding smaller fit statistics.

Transition probability. The stationarity of transition probabilities was evaluated by comparing the models that posit different assumptions on the effects of the latent states and covariates. Table 8 reports fit statistics of the evaluated models—the models that assumed stationary or temporal state transition. The stationary transition models (labeled as ST) assumed that all predictors had constant effects on the transition probability across the different time points. The time-variable transition models were examined under three scenarios: (i) when both the preceding latent state and covariates had time-varying effects on the transition (labeled as VT), (ii) when only the previous latent state had a time-variant effect and the effects of covariates were fixed (labeled as VTS), and (iii) when the covariates had time-varying effects and the latent state had a constant effect (labeled

Table 7

Fit Statistics of LTMs: Covariate Effects

Condition	df	LL	AIC	CAIC	BIC	ABIC
(MI, -)	21	-8995.53	18033.06	18033.48	18152.86	18086.14
(MI, C)	27	-8928.99	17911.99	17912.68	18066.02	17980.23
(RI, -)	24	-8700.68	17449.37	17449.91	17586.28	17510.03
(RI, C)	30	-8634.15	17328.30	17329.15	17499.44	17404.13
(VRI, -)	33	-8674.55	17415.11	17416.14	17603.37	17498.52
(VRI, C)	39	-8608.02	17294.04	17295.47	17516.53	17392.62

Note. df = Number of free parameters. LL = -2 log likelihood. AIC = Akaike information criterion. CAIC = Corrected AIC. BIC = Bayesian information criterion. ABIC = Adjusted BIC. MI = LTM with measurement invariance. C = Covariates. RI = LTM with the random intercept. VRI = LTM with the random intercept with time-varying effects.

Table 8

Fit Statistics of LTMs: Stationarity of Transition Probabilities

Condition	<i>df</i>	LL	AIC	CAIC	BIC	ABIC
(VRI, C, ST)	39	-8608.02	17294.04	17295.47	17516.53	17392.62
(VRI, C, VT)	61	-8581.60	17285.19	17288.70	17633.19	17439.38
(VRI, C, VTS)	43	-8591.04	17268.08	17269.82	17513.38	17376.77
(VRI, C, VTC)	59	-8588.22	17294.44	17297.72	17631.03	17443.58

Note. *df* = Number of free parameters. LL = -2 log likelihood. AIC = Akaike information criterion. CAIC = Corrected AIC. BIC = Bayesian information criterion. ABIC = Adjusted BIC. VRI = LTM with the random intercept with time-varying effects. C = Covariates. ST = LTM with stationary transition probability. VT = LTM with time-varying transition probability. VTS = Time-varying transition due to latent states (covariate effects fixed). VTC = Time-varying transition due to covariates (state effects fixed).

as VTC). In Table 8 comparison of the stationary and temporal transition models suggests that the log-likelihood preferred the time-varying latent states and covariates whereas the information criteria preferred the time-varying latent states with fixed covariates. All in all, the results seemed to suggest that the preceding latent states had substantially different influence on the transition likelihood as the time progressed while the effects of the covariates may be modeled by constants if the state effects are modeled with time-varying coefficients.

Comparison of various LTMs suggested that the model that allows time-variant effects for the random intercept and latent states (labeled as VRI-C-VTS) demonstrates best fit for the CTA data. Based on this finding, we subsequently examined the outcomes of the final model to draw inference about the latent flow states. We note that the model identified here also yielded the best results in the other data sets—i.e., when the data were partitioned into three and five segments.

Measurement model. Table 9 reports estimated loadings of the latent variables on the measurement variables. The estimates suggest that the latent states showed distinct patterns in the indicator variables. Students in the first state spent relatively shorter time on the problems, asked for no help, and were less likely to make errors. Those in the

Table 9

Measurement Model Parameters of the Final LTM (VRI, C, VTS): Loading of Latent Variables on Indicators

Par	Latent State		Random Intercept			
	$\alpha(s = 1)$	$\alpha(s = 2)$	$\lambda(t = 1)$	$\lambda(t = 2)$	$\lambda(t = 3)$	$\lambda(t = 4)$
Time	3.788	4.141	-.387	-.261	-.246	-.224
(SE)	(.010)	(.013)	(.037)	(.031)	(.025)	(.026)
Error	.686	1.171	-.038	-.005	.006	-.035
(SE)	(.007)	(.008)	(.020)	(.020)	(.019)	(.021)
Hint	.000	1.021	-.010	-.003	-.006	-.014
(SE)	(.000)	(.004)	(.005)	(.003)	(.003)	(.005)

Note. α = Loading on the latent state. λ = Loading on the random intercept. The subscripts within the brackets indicate the latent states or time points at which the variable was conditioned.

second state spent distinctly longer time and tended to make more erroneous attempts while requesting nonzero hints. The examination of the state proportions suggested that students displayed the first attitude (68.58%) more often than the second (31.42%). The frequency of the second state tended to decrease in the first three quarters of the tutoring (36.73%, 29.02%, 27.13%) and increased in the final phase (32.81%). The patterns observed in the parameter estimates and the state frequency seemed to suggest that the first state reflects the normal problem-solving state and the second state indicates the state of out-of-flow. We surmised that the detached actions that occurred in the second state may be due to warm-up or fatigue because CTA was assigned only infrequently during regular curricula and students tended to receive a large number of problems once started (50.236 problems on average and 151 maximum).

We also add that, although the covariates were not used as immediate predictors of the latent states, a distinct pattern was observed in relation to the pretest scores. Students identified as displaying the second state tended to have substantially lower pretest scores (-.466 on average) than the overall mean (-.366) as well as the average score in the first state (-.322). Table 9 also reveals that the indicator variables showed different amounts of

measurement noninvariance. The task interaction time showed consistently large loadings on the random intercept with relatively large variation across time. The discretized error and hint variables showed somewhat smaller and constant loadings, suggesting that they induced little and approximately equal amounts of measurement noninvariance.

Transition model. In Table 10 we present transition probabilities estimated from the final model. The first entry of each transition matrix consistently showed high probability (.733 on average), suggesting that students tended to maintain the same flow state once in place. Students exhibiting the deviant state (i.e., State 2) tended to switch to the flow state in the early phase of tutoring but gradually showed a stronger tendency to stay in the same out-of-flow state as the tutoring progressed. The current results support our hypothesis on the deviating state. Earlier, we mentioned that the deviant state may reflect a warm-up or fatigue effect. The results from the transition probabilities suggest that the aberrant state identified in the early phase of tutoring is likely to reflect a warm-up or a learning period, while the state appearing in the later stage is likely related to loss of motivation (e.g., fatigue, exhaustion). The high transition rate from the deviant to the flow state in the beginning in particular suggests that students who showed slow progress at the outset began to engage in learning as they become familiar with the contents and problems. The decreased transition probability in the later phase suggests that weary students were unlikely to shift back to the flow state in the subsequent phases. The individual students' transition patterns further corroborated our hypothesis. About 63.68% of the students who

Table 10

Transition Probabilities in the Final LTM

	$P(S_2 S_1, \mathbf{Z})$		$P(S_3 S_2, \mathbf{Z})$		$P(S_4 S_3, \mathbf{Z})$		Overall	
From \ To	$S_2 = 1$	$S_2 = 2$	$S_3 = 1$	$S_3 = 2$	$S_4 = 1$	$S_4 = 2$	$S_t = 1$	$S_t = 2$
$S_{t-1} = 1$	.730	.270	.752	.248	.718	.282	.733	.267
$S_{t-1} = 2^*$	.676	.324	.674	.326	.550	.450	.633	.367

Note. S_t = Latent state at time t . $\mathbf{Z}$ = Covariates. * Suspected as the state of disengagement. $P(S_0 = 1) = .633$, $P(S_0 = 2) = .367$

showed early aberrance received *es1* as the first unit on the day. About 86.13% of the students who showed aberrancy in the final stage received *es1* as the last unit of the day.

Table 11 reports estimated effects of the latent states and covariates on the transition probabilities. The effects were evaluated for the first state (i.e., normal state) because the estimation algorithm in *Mplus* treated the second state as a baseline. The results show that the first state had constantly positive loadings on the subsequent first state. This means that students who showed the normal state in the preceding problem set tended to maintain the same state in the following problem set. The increasing loading coefficients suggest that the tendency to stay in the same flow state intensified as time progressed. The effects of covariates on the transition likelihood varied in the direction and size. The negative loadings of the two ethnicity groups suggest that Black and Multiethnic, and Hispanic and Native American students were less likely to transition to the flow state than were the White and Asian students. The strong negative loading of the Hispanic-native American students implies that they had a greater tendency to show the deviant state. The positive loadings in the other covariates suggest that students with higher scores had a stronger tendency to move to the first flow state. For example, male students showed greater likelihood to move to the flow state than female students. Students with higher pre- and gain-scores showed a greater tendency to transition to the flow state as they receive the next problem set(s). Among others, the results pointed a distinct tendency in

Table 11

Effects of Latent States and Covariates on Transition Probabilities

	$S_t[1]$ on $S_{t-1}[1]$			$S_t[1]$ on Covariates ($t = 1, 2, 3, 4$)					
	$t = 2$	$t = 3$	$t = 4$	Pre	Gain	RaceBM	RaceHN	Sex	FRL
Est	.262	.385	.745	.375	.198	-.081	-.126	.276	.021
(SE)	(.097)	(.103)	(.100)	(.042)	(.037)	(.059)	(.075)	(.049)	(.055)

Note. The base category of Race was White and Asian. RaceBM = Black and multiethnicity. RaceHN = Hispanic and native American. The baseline category of Sex was Female. FRL = Free lunch. Pre = Prescore. Gain = Gain score. All latent states and covariates except for RaceBM ($p = .207$) and free lunch ($p = .901$) were significant predictors of the transition probabilities ($p < .01$).

the Hispanic and Native American female students. The effect coefficients suggested that these students were at greater risk of showing out-of-flow states, signaling the need for remedial interventions.

State transition. We conclude our latent transition analysis with a summary of the most pronounced state patterns. Among the 16 transitional patterns, the largest group of students (27.27%) showed continued attention on the tasks with the estimated state pattern $\mathbf{S} = (1, 1, 1, 1)$. The second group of students (12.80%; $\mathbf{S} = (2, 1, 1, 1)$) showed transition from the deviant to the flow state in the early stage. The third group (9.19%; $\mathbf{S} = (1, 1, 1, 2)$) showed a shift from the flow to the out-of-flow in the final stage. The next frequent groups changed midway, showing state patterns (1, 2, 1, 1) (8.70%) and (1, 1, 2, 1) (6.17%). These results suggest that, despite the self-regulated tutoring, many students performed conscientiously in most problems and exhibited deviant behaviors only occasionally. We remark that the present analysis was performed on the aggregated problem sets. The results do not inform specific problems on which a student showed deviant behaviors nor whether the aberrant behaviors occurred continuously across the partition or on a subset of problems. In the next section, we examine the individual students' transitional patterns across the problems to draw finer-grained information.

5. Flow Across Problems

The last analysis applied hidden Markov models (HMMs; Burke & Rosenblatt, 1958) to examine the flow transition between individual problems.

5.1. Hidden Markov Model

As with the LTM, the HMM uses a sub-model to describe transition between the latent states. Consider a latent (or hidden) stochastic process, $\mathbf{S}$, that evolves over discrete time points. A series of manifest indicators is modeled jointly with the corresponding latent states, $P(\mathbf{Y}_i, \mathbf{S})$. If covariates are available, $P(\mathbf{Y}_i, \mathbf{S})$ can be modeled as

$$P(\mathbf{Y}_i, \mathbf{S} | \mathbf{Z}_i) = P(S_0 | \mathbf{Z}_i) \prod_{t=1}^T P(S_t | S_{t-1}; \mathbf{Z}_i) \prod_{t=1}^T P(\mathbf{Y}_{it} | S_t). \quad (9)$$

The formulation (9) reveals three sub-models: the (i) structural model, (ii) transition model, and (ii) measurement model. The structural model defines the probability of an initial state, $\boldsymbol{\pi} = (\pi_m : m = 1, \dots, M)$, with $\pi_m = P(S_0 = m)$. The transition model describes the transition probability between the latent states, $\mathbb{P} = (p_{mm'} : m, m' = 1, \dots, M)$ where $p_{mm'} = P(S_t = m | S_{t-1} = m')$. Observe that the transition model assumes a homogeneous first-order Markov process. The LTM similarly assumed the first-order Markov process but it allowed the transition probabilities to vary across times. The last component, the measurement model describes the emission probability of an indicator given the latent state, $P(Y_{jt} | S_t)$. Some common choices in the HMM applications include: Gaussian, Poisson, binomial, Gamma, and multinomial models. We also note that HMMs do not typically assume problem-level parameters when modeling the emission probabilities; instead, it treats all observations from the same latent state as homogeneous observations from the same emission model and estimates the state-specific emission probabilities (or corresponding distributional parameters) that apply to all within-state observations.

Let $\boldsymbol{\tau}$ contain the parameters of the emission probabilities for all indicator variables and states. HMM is then defined by a triple, $\boldsymbol{\theta} = (\boldsymbol{\pi}, \mathbb{P}, \boldsymbol{\tau})$. For estimating $\boldsymbol{\tau}$, we use an R package, depmixS4 (Visser & Speekenbrink, 2010). The model is commonly fit to person-level data under the weak measurement invariance assumption. Although the program allows sample-level fitting, our empirical analysis based on the simulated and real data suggests that the estimation tends to encounter recurrent convergence problems. To make the inference adequately reliable across the analyses, we apply HMMs only to the individual student-level data. This approach also follows the existing practice in the HMM application. Based on our empirical analysis, we discuss some possible consequence of fitting at the person level when the problems have distinct measurement effects.

5.2. Flow Progression in CTA

The analysis based on the HMM was performed as follows. Since the model allows individual-level fitting, minimal adjustment was needed for addressing the assumptions and estimation constraints. We therefore applied the HMMs to each student's problem-level data and examined the transition across the individual problems. The input data for the analysis consisted of $N = 2219$ students identified in 2.2. The analysis used the same indicator variables and covariates as before, but applied minimal transformation thanks to the flexibility in the estimation software. Specifically, we applied the log-transformation to the task interaction time and used the other variables as collected.

A number of distinctions made in the HMM analysis are worthy of note. In the preceding analyses, we categorized the count variables into three levels to accommodate the constraints in the estimation programs. The HMM analysis however used the raw data and is expected to show greater sensitivity to the manifest variables. In some cases, the large variation in the indicator values can lead to overextraction of latent states as a result of the increased heterogeneity.

Second, unlike the other models, HMMs were fit to the individual student-level data and do not account for systematic effects of the problems. In CTA, students received different problem sets in various orders, and it made it difficult to unravel the variance in the measurements with the current HMM estimation software. Our preliminary study based on the simulated data suggested that ignoring measurement variance can result in over-extraction of latent states. The current study attempted to alleviate this tendency by (i) focusing on a single unit of homogeneous problems (i.e., *es1*), (ii) conducting analysis in an exploratory manner, (iii) using a conservative information metric in the final model selection, and (iv) giving greater priority to the models with clear binary interpretation. Specifically, when fitting the HMMs, we assumed one to ten latent states at each time point and determined a final model using the most restrictive criterion, CAIC (Burnham & Anderson, 2002; Sugiura, 1978). The range of the number of latent states was determined

in the precursory analysis by weighing the interpretability and generalizability of the models. While more complex models are conceivable, we did not explore these possibilities because our primary goal in this study was to make a binary decision on the latent states that can be characterized as in- and out-of-flow. When the initial model fitting suggested multiple states, we refit the two-state HMMs to see if the estimated states can be simplified to two states.

Third, recall that all covariates used in this study were time-invariant. When HMMs are fit to the person-level data with the time-invariant covariates, the covariates become constant and play no part in inferring the transition probabilities. Our supplementary analysis suggests that, although the constant covariates function as incidental variables, they can help choose a parsimonious solution. As another way of mitigating the measurement noninvariance problem, we therefore retain the covariates when fitting the HMMs though no particular inference could be attached to the covariate effects.

The outcomes of the HMM analysis were evaluated in three aspects: (i) intricacy of the latent process, (ii) progression of latent states across tutoring (e.g., when does a student most likely lose attention?), and (iii) consistency with the preceding analyses. Below, we present results of the HMM analysis and give three representative example students that showed the typical transition patterns.

5.3. Results

After fitting the HMMs under different numbers of latent states, we determined a final model applying the information criteria, (AIC, CAIC, BIC, ABIC). Among the criteria evaluated, the ABIC and AIC suggested the most complex models with a minimum of one and a maximum of ten states. The CAIC and BIC suggested simpler models with a maximum of four states. Below we present results from CAIC that suggested the most parsimonious and clear model solutions. The outcomes of CAIC indicated that about 39.97% students ($N = 887$) displayed a homogeneous latent state as they advanced the problems, 51.83% ($N = 1150$) displayed two states, 8.07% ($N = 179$) three states, and

1 .14% students ($N = 3$) four latent states. There was weak correlation between the
 2 estimated number of states and the number of observations ($r = .241$), suggesting that
 3 students with more problem attempts tended to display more diverse states.

4 Figures 1–3 provide three example students that were suspected of showing
 5 out-of-flow during tutoring. The students in each figure were identified as displaying two,
 6 three, and four states across the evaluated time periods. The first three plots at the top of
 7 each figure present observation series from each indicator. The last plot shows the
 8 estimated state at each time point. The dotted vertical lines indicate the points at which

Figure 1. An Example Student Displaying Two States

Note. State 1 was considered a flow state. The dotted vertical lines indicate the points at which the problems were demarcated in the LTM analysis.

the problems were demarcated in the LTM analysis. Along with the figures, we present parameter estimates of the emission and transition models obtained from each student in Table 12.

The student in Figure 1 worked on a total of 57 problems over one hour and 43 minutes. As can be noted, the student maintained the same state in most cases but displayed somewhat deviant behavior toward the end of tutoring. In the early stage, the student solved most problems on his own and rarely made errors (.087 on average), resulting in average interaction time 54.49 seconds (per problem). As tutoring progressed, the student made more errors (6.909 on average) and requested more help (2.727), spending average time of 2.794 minutes. We restate that the problems assigned in CTA were adapted to the student's algebraic ability, and it is not generally expected to have a radical change in the outcome values. The continued digression from the normal state suggests that the student was either paying little attention or struggling with the problems. The transition probability estimates in Table 12 suggest that once the student adopted a particular mode, he tended to maintain the same mode in the subsequent problems. The probabilities of staying in the same state were .934 and .798 for each state.

Figure 2 presents a transition pattern of a student who displayed three states. The student worked on 79 problems in 63 minutes. As evident from the figure, the student displayed distinct behaviors in each state. When in State 1, he seldom made mistakes (.143 on average), asked for no hints, and spent little time on the problems (21.224 seconds on average). As he approaches the midpoint of tutoring, he tended to make more errors and asked for nonzero help. Specifically, when the student was in State 2 (28th–51st problems), he made .417 errors and asked for 1.792 hints on average. When he was in State 3, he showed a stronger tendency to make errors (4.167) and ask for help (3.833), and spent substantially longer time on the problems (3.584 minutes). Across the evaluated time periods, he showed State 1 most frequently (62.03%, $T = 49$ items), and States 2 (30.38%, $T = 24$) and 3 (7.59%, $T = 6$) in sequence. Combining the results, we conclude that States

Figure 2. An Example Student Displaying Three States

Note. State 1 was considered a flow state. The dotted vertical lines indicate the points at which the problems were demarcated in the LTM analysis.

- 1 2 and 3 reflect the state of out-of-flow. In both the states, the student spent more time
- 2 than usual and made distinctly more trials. These two states differed only in the
- 3 magnitude of the indicators with those from State 3 reflecting greater detachment. As we
- 4 fit the HMM confirmatively under two latent states, the States 2 and 3 indeed merged
- 5 together and formed one deviant state, corroborating our hypothesis. The transition model
- 6 in Table 12 suggests that the student showed strong a propensity to maintain the same S_1
- 7 and S_2 attitudes once adopted. The staying probabilities in these states equaled .959 and
- 8 .957, respectively, indicating high stationarity. When the student was in State 3, he

1 maintained the same state with a .598 probability and moved to State 2 or 1 with
 2 probabilities of .204 and .198, respectively.

3 Figure 3 presents a state transition pattern of a student displaying four attitudes.
 4 The student attempted 50 problems in 51 minutes. Figure 3 shows that the student rarely
 5 made errors and asked for only a few hints when working on the first three-quarters of the
 6 unit. As he worked on the last quarter, he began to show deviating behaviors, making
 7 many mistakes and asking for multiple hints. The indicator patterns in each state suggest
 8 that when the student was in State 1 or 2, he made no or very few mistakes (.250 on

Figure 3. An Example Student Displaying Four States

Note. State 1 was considered a flow state. The dotted vertical lines indicate the points at which the problems were demarcated in the LTM analysis.

Table 12

Emission and Transition Model Parameters of HMMs

	Emission			Transition				
	Task Time	Num Err	Num Hint	From \ To	$S_t = 1$	$S_t = 2$	$S_t = 3$	$S_t = 4$
Student 1 ($M = 2$; Grade 9, White-Asian, Male, Free lunch, No ESL)								
State 1	$\mathcal{N}(10.448, .907^2)$	Pois(.000)	Pois(.088)	$S_{t-1} = 1$	.934	.066		
State 2	$\mathcal{N}(11.917, .491^2)$	Pois(2.731)	Pois(6.916)	$S_{t-1} = 2$	.202	.798		
Student 2 ($M = 3$; Higher grade, Black-Multi, Male, Free lunch, ESL)								
State 1	$\mathcal{N}(9.840, .487^2)$	Pois(.000)	Pois(.143)	$S_{t-1} = 1$	.959	.000	.041	
State 2	$\mathcal{N}(10.375, .351^2)$	Pois(1.791)	Pois(.413)	$S_{t-1} = 2$	.000	.957	.043	
State 3	$\mathcal{N}(12.082, .640^2)$	Pois(3.820)	Pois(4.152)	$S_{t-1} = 3$	.198	.204	.598	
Student 3 ($M = 4$; Higher grade, White-Asian, Male, Free lunch, No ESL)								
State 1	$\mathcal{N}(9.875, .267^2)$	Pois(.000)	Pois(.000)	$S_{t-1} = 1$	.788	.068	.097	.047
State 2	$\mathcal{N}(10.706, .428^2)$	Pois(.000)	Pois(.251)	$S_{t-1} = 2$	.228	.772	.000	.000
State 3	$\mathcal{N}(11.360, .248^2)$	Pois(.824)	Pois(5.286)	$S_{t-1} = 3$	.000	.605	.197	.197
State 4	$\mathcal{N}(12.687, .387^2)$	Pois(5.000)	Pois(24.333)	$S_{t-1} = 4$	.000	.000	.667	.333

Note. M = Number of latent states.

average) and asked for zero help. When in State 3, the student made 5.333 errors and asked for more help (.833 hints on average). In State 4, the student showed a sharp increase in both the indicators, making 24.333 mistakes and asking 5 hints on average. The total task interaction time similarly indicated different degrees of lassitude. In each of the four states, the student spent on average 19.873 seconds, 49.216 seconds, 1.470 minutes, and 5.829 minutes on the problems, respectively, indicating that the student spent substantially longer time when in States 3 and 4. Across the 50 problems attempted, the student displayed States 1 and 2 most frequently (42%, $T = 21$; 40%, $T = 20$ each) and States 3 and 4 less frequently (12%, $T = 6$; 6%, $T = 3$ each). When the two-state HMM was fit to the same data, States 1 and 2 were grouped together to form one state, describing 82% of the observations ($T = 41$ problems). States 3 and 4 combined together to form the other state, explaining 18% of the observations ($T = 9$ problems). The present results suggest that States 3 and 4 are likely to indicate the deviant state. In both the

states, the student tended to make distinctly large numbers of errors and lingered on the problems longer. The transition probability estimates in Table 12 suggest that the student showed high stationarity when in State 1 or 2 (.780 probability on average). When he was in State 3 or 4, he showed great mobility shifting to each State 2 (.605) or 3 (.667).

6. Conclusion

The purpose of this study was to explore a latent variable modeling approach to tracking flow in the ITS data. We considered three models that provide discrete inference on the flow state and applied them to the log data from CTA to demonstrate practical analysis strategies. Throughout the analyses, special attention was paid to three aspects of flow: (i) the progression of flow states across tutoring, (ii) interaction modalities under the in- and out-of-flow states, and (iii) relation between the flow and student's demographic profiles. The results from the analyses revealed substantive information about students' flow states that can be of service in future ITS design and intervention planning. While the models differed in the assumptions and data constraints, they indicated similar conclusions on the trajectory of latent flow and provided information that complements each other.

The analyses were carried out in three gradual phases. In the first analysis, we applied the latent class models to identify latent profiles underlying the knowledge-level data. The evaluation data were fixed at the knowledge level to regulate dimensionality of time points (and thus that of latent space). To account for variance from the adaptive problem assignments and fluctuations in the measurements therein, we added random-effect terms to both the structural and measurement models. The results from the analysis suggested that students generally showed uniform behaviors when working on the CTA. A small group of students (18.97%) displayed deviant behaviors, spending distinctly longer time, asking for more help, and making more errors. The latent profile estimates suggested that students with lower pre- or gain-scores, students in minority groups, and male students were more prone to exhibit deviating behaviors from the flow state.

As we learned that students displayed distinct behaviors under different state profiles,

we examined the evolution of states across the segmented tutoring sessions. The state profiles identified in the latent class analysis give only the temporal portrayal of states that underlie at each point and do not provide a full picture of the flow development across tutoring. The second analysis was conducted to evaluate such transitioning of flow states across tutoring sessions. We applied the latent transition models to explicitly model the evolution of flow states across the series of problem sets. To model different degrees of measurement noninvariance, we used person-level random intercepts allowing interaction with the indicator variables. The results from the transition analysis suggested that the majority of students showed adequate persistence in general (about 68.58% of the time), spending a reasonable amount of time and making conscientious efforts. At other times, students showed deviating behaviors (31.42%), dawdling on the problems and making distinctly more attempts. The deviant behaviors appeared more prevalently in the beginning and final stages of tutoring, plausibly due to warm-up or fatigue. The transition probability estimates suggested that students with higher pre- and gain-scores and male students had a stronger tendency to stay in flow. Female students and students in the minority ethnic groups were more likely to transit to the deviant state.⁵

The last analysis examined the flow development across individual problems using the hidden Markov models. The latent transition analysis indicated heterogeneity in the students' behaviors and their flow transition over time but the analysis was conducted on the aggregated data to meet the constraints in the estimation software. HMMs allow a finer-grained description of flow transition at the individual problem levels and is more flexible in modeling different kinds of indicator variables. In this study, HMMs were applied to the student-level data under weak measurement invariance and analysis was performed in an exploratory manner to cope with possible overextraction of latent states. The outcomes of the HMM analyses pointed to the specific problems at which the students

⁵ Note that LCMs used the covariates to predict state probability at a given time point whereas LTMs used the covariates to model transition probability between the states.

1 became inattentive and deviate from the flow. Examination of individual student patterns
2 also revealed how the different flow states became exteriorized in the behavioral indicators
3 across the series of problems.

4 As described, findings from the three analyses were generally consistent and indicated
5 similar conclusions on the flow states and associated demographic profiles. The information
6 from each model was however unique enough to warrant separate attention and appeared
7 to complement each other so that they together give comprehensive overview of the
8 students' behaviors and underlying states.

9 While the models demonstrated potential for describing ITS data, there also appears
10 to be room for further improvement. As illustrated throughout the analyses, each modeling
11 approach required separate technical modifications to accommodate the model assumptions
12 and estimation constraints. The random-effect latent class model accounted for extra
13 variance inherent in the students and measurement stimuli using the random-effect terms.
14 However, the large number of random terms present in the model caused convergence
15 problems, necessitating corrective steps. In addition, the model did not explicitly describe
16 progression of latent states, limiting the inference at each time point. The latent transition
17 model improved some of the limitations of the latent class model, for example, by explicitly
18 modeling the state transition and applying a more affordable estimation routine. The
19 model however placed strict constraints on the number of time points and measurement
20 invariance. The currently available estimation programs are not generally suited for
21 intensive time-series data and require transformation or recoding of indicator variables to
22 meet the distributional assumptions of the estimation program. Among the three latent
23 variable models, the HMM had the greatest flexibility in modeling various types of indicator
24 variables and placed the least constraint. This flexibility however was achieved by ignoring
25 the systematic effects of measurement stimuli and consequently it was more susceptible to
26 random fluctuations in the indicators and was likely to overextract latent states.

27 Our empirical analysis of CTA data suggests that, while the current latent variable

1 models have some merits, the development of a more comprehensive and systematic
2 modeling framework is generally more desirable. For example, the current modeling
3 frameworks can be extended to accommodate various indicator variables and parametric
4 distributions that are commonly observed in the interactive ITS. The framework can also
5 be further advanced using a more efficient estimation program that supports the analysis of
6 intensive time-series data and allows integration of covariate information as well as prior
7 information on the latent states. Together with the efficient estimation routine, the model
8 can also be used online to signal change in behaviors in real time.

1 Open Practices Statement:

2 The program code and data that support the findings of this study are available from the
3 authors' GitHub pages:

4 <https://github.com/HyeonahKang/LatentVariableModel-IntelligentTutor>;

5 <https://github.com/adamSales/SELS>.

References

- Agresti, A. (2012). *Categorical data analysis* (3rd. ed.). New York, NY: Wiley.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Csáki (Eds.), *Proceedings of the second international symposium on information theory* (pp. 267–281). Budapest: Akadémiai Kiadó.
- Baker, R., Gowda, S. M., Wixon, M., Kalka, J., Wagner, A. Z., Salvi, A., ... Rossi, L. (2012). Towards sensor-free affect detection in cognitive tutor algebra. In K. Yacef, O. Za'iane, A. Hershkovitz, M. Yudelson, & J. Stamper (Eds.), *Proceedings of the 5th international conference on educational data mining* (pp. 126–133). Chania, Greece: www.educationaldatamining.org.
- Boniwell, I. (2008). *Living in flow: What is it and how to enter the flow state?* Retrieved from <http://positivepsychology.org.uk/living-in-flow/>.
- Burke, C. J., & Rosenblatt, M. (1958). A Markovian function of a Markov chain. *The Annals of Mathematical Statistics*, 29(4), 1112–1122.
- Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach* (2nd. ed.). New York: Springer-Verlag.
- Calvo, R., D'Mello, S. K., Gratch, J., & Kappas, A. (Eds.). (2015). *The oxford handbook of affective computing*. New York, NY: Oxford University Press.
- Csikszentmihályi, M. (1990). *Flow: The psychology of optimal experience*. New York, NY: Harper & Row.
- Csikszentmihályi, M. (2008). *Flow: The psychology of optimal experience*. New York, NY: Harper Perennial.
- Dupre, D., Booth, A., Bolster, A., Morrison, G., & McKeown, G. (2017). Dynamic analysis of automatic emotion recognition using generalized additive mixed models. In *Symposium on computational modelling of emotion: Theory and applications* (pp. 158–163). Bath: UK: <http://aisb2017.cs.bath.ac.uk/proceedings.html>.

- 1 Fredricks, J. A., Blumenfeld, P. C., & Paris, A. H. (2004). School engagement: Potential of
2 the concept, state of the evidence. *Review of Educational Research*, 74, 59–109.
- 3 Fredricks, J. A., Filsecker, M., & Lawson, M. A. (2016). Student engagement, context, and
4 adjustment: Addressing definitional, measurement, and methodological issues.
5 *Learning and Instruction*, 43, 1–4.
- 6 Fwa, H. L., & Marshall, L. (2018). Modeling engagement of programming students using
7 unsupervised machine learning technique. *GSTF Journal on Computing*, 6(1), 1–6.
- 8 Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013).
9 *Bayesian data analysis* (3rd ed.). CRC press.
- 10 Hoffman, M. D., & Gelman, A. (2014). The no-u-turn sampler: adaptively setting path
11 lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1), 1593–1623.
- 12 Lo, Y., Mendell, N. R., & Rubin, D. B. (2001). Testing the number of components in a
13 normal mixture. *Biometrika*, 88(3), 767–778.
- 14 Martin, A. J. (2012). Motivation and engagement: Conceptual, operational, and empirical
15 clarity. In S. L. Christenson, A. L. Reschly, & C. Wylie (Eds.), *Handbook of research
16 on student engagement* (pp. 303–311). Boston: Springer US.
- 17 Muthén, B., & Asparouhov, T. (2020). Latent transition analysis with random intercepts
18 (RI-LTA). *Psychological Methods*, Advance online publication, 1–18.
- 19 Muthén, L. K., & Muthén, B. O. (2017). *Mplus: Statistical analysis with latent variables:
20 User's guide (Version 8)*. Los Angeles, CA.
- 21 Nocedal, J., & Wright, S. J. (2006). *Numerical optimization* (2nd. ed.). Berlin:
22 Springer-Verlag.
- 23 Pane, J. F., Griffin, B. A., McCaffrey, D. F., & Karam, R. (2014). Effectiveness of
24 Cognitive Tutor Algebra I at scale. *Educ. Eval. Policy Anal.*, 36(2), 127–144.
- 25 Petrovica, S., Anohina-Naumeca, A., & Ekenel, H. K. (2017). Emotion recognition in
26 affective tutoring systems: Collection of ground-truth data. *Procedia Computer
27 Science*, 104, 437–444.

- 1 Qu, Y., Tan, M., & Kutner, M. H. (1996). Random effects models in latent class analysis
2 for evaluating accuracy of diagnostic tests. *Biometrics*, 52(3), 797–810.
- 3 R Core Team. (2020). R: A language and environment for statistical computing [Computer
4 software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- 5 Redmond, P., Heffernan, A., Abawi, L., Brown, A., & Henderson, R. (2018). An online
6 engagement framework for higher education. *Online Learning*, 22, 183–204.
- 7 Ritter, S., Anderson, J. R., Koedinger, K. R., & Corbett, A. (2007). Cognitive Tutor:
8 Applied research in mathematics education. *Psychonomic Bulletin & Review*, 14(2),
9 249–255.
- 10 Schwartz, G. E. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6(2),
11 461–464.
- 12 Sclove, S. L. (1987). Application of model-selection criteria to some problems in
13 multivariate analysis. *Psychometrika*, 52, 333–343.
- 14 Shih, B., Koedinger, K. R., & Scheines, R. (2010). Discovery of student strategies using
15 hidden Markov model clustering. In *Proceedings of the 6th international conference*
16 *on educational data mining*. Memphis, TN: www.educationaldatamining.org.
- 17 Sinclair, J., Jang, E. E., Azevedo, R., Lau, C., Taub, M., & Mudrick, N. V. (2018).
18 Changes in emotion and their relationship with learning gains in the context of
19 metatutor. In R. Nkambou, R. Azevedo, & J. Vassileva (Eds.), *Intelligent tutoring*
20 *systems* (pp. 202–211). Cham: Springer International Publishing.
- 21 Sugiura, N. (1978). Further analysis of the data by Akaike’s information criterion and the
22 finite corrections. *Commun. Stat. Theory Methods*, 7(1), 13–26.
- 23 Visser, I., & Speekenbrink, M. (2010). depmixS4: An R package for hidden markov
24 models. *Journal of Statistical Software*, 36(7), 1–21. Retrieved from
25 <http://www.jstatsoft.org/v36/i07/>

Appendices

Appendix A: Measurement Properties of the Problems within Units

Given the lack of supporting tools, summary statistics of the problem indicators were used to examine the measurement invariance. For count data (i.e., number of hints/errors), we examined the mean and standard deviation (SD) of the observed values, $E(X_{it} : i = 1, \dots, N)$ and $SD(X_{it} : i = 1, \dots, N)$, where i and t each indexes students and problems (administered across time). For timing data (i.e., interaction time), we examined the mean and SD of the log-transformed values, $E(\log X_{it} : i = 1, \dots, N)$ and $SD(\log X_{it} : i = 1, \dots, N)$. The (dis)similarity of the statistics across the problems within a unit was then evaluated by the variance across the problems. Table below reports the variance of the summary statistics observed from each unit. As can be seen, *es1*—the unit examined in this study—, showed the smallest variance on both the evaluation criteria. The unit was administered to a fairly large sample of students and yet showed small variance in the measurement summary statistics.

Table A1

Variance of Summary Statistics of the Problems Within Each Unit

Unit	N	Mean				SD			
		nerr	nhint	ltime	Avg	nerr	nhint	ltime	Avg
ctal 01	2841	10.258	1.285	.222	3.921	5.658	.908	.005	2.190
ctal 02	2469	55.882	4.542	.073	20.166	62.809	4.092	.002	22.301
ctal 04	2067	7.970	.800	.290	3.020	5.820	.869	.002	2.230
ctal 06	1531	53.186	3.725	.319	19.076	66.601	4.380	.000	23.661
ctal 08	1045	6.631	.705	.084	2.473	6.239	.769	.001	2.336
ctal 10	1129	3.076	.420	.071	1.189	2.126	.448	.001	.858
ctal 13	1677	3.349	.370	.112	1.277	1.901	.423	.001	.775
ctal 14	974	4.517	.749	.119	1.795	6.102	1.522	.004	2.543
ctal 12	702	12.883	1.280	.225	4.796	3.459	.944	.011	1.471
<u>es 01</u>	<u>2270</u>	<u>.354</u>	<u>.021</u>	<u>.690</u>	<u>.355</u>	<u>1.391</u>	<u>.351</u>	<u>.028</u>	<u>.590</u>
es 02	1646	1.000	.177	.442	.540	4.193	2.336	.037	2.189
es 03	1315	2.365	.144	.528	1.012	7.169	3.761	.021	3.650
es 04	1232	.491	.058	.624	.391	2.814	.799	.017	1.210
es 05	701	1.134	.070	.647	.617	3.554	.518	.016	1.363
es 07	784	14.207	4.673	.760	6.547	11.050	2.566	.015	4.544
general linear form	743	10.329	2.088	.065	4.160	14.270	.942	.051	5.088
linear inequal graphing	516	2.671	.158	.079	.969	17.273	2.077	.004	6.451
quad add area alg1	422	22.198	2.445	.053	8.232	29.401	4.255	.033	11.230
quad vertical motion	326	38.662	3.878	.076	14.205	103.537	4.516	.007	36.020
expt product simp a-es	401	.413	.034	.463	.303	4.730	.549	.028	1.769
expt quotient simp a-es	397	7.444	.708	.510	2.888	9.441	1.545	.004	3.663
polynomial arith es	279	.614	.012	.058	.228	1.371	.049	.009	.476
exponential functions	61	31.586	3.208	.095	11.630	47.612	6.037	.006	17.885
unit conversions	1								
inequality systems	1								
glf modeling	1								
ratio of irrational numbers	1								

Note. nerr = Number of errors. nhint = Number of hints. ltime = Logarithm of interaction time. Avg = Average.

Appendix B: Fitting the Random-Effect Latent Class Model in Stan

The latent class model described in Section 4 was fit in Stan, called from R via the `rstan` package. In this appendix, we will first give the full likelihood equation and prior model, and then the Stan code defining the model. Complete replication code may be found in this manuscript's companion GitHub repository.

The full likelihood is given by

$$\begin{aligned}
 P(\mathbf{Y} \mid \boldsymbol{\eta}, \mathbf{Z}, \boldsymbol{\delta}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\Sigma}_\delta, \boldsymbol{\sigma}_\eta) \\
 = \prod_{i=1}^N \prod_{t=1}^{T_i} \left\{ \sum_{m=1}^M P(S_{it} = m \mid \eta_i, \mathbf{Z}_i = \mathbf{z}_i, \boldsymbol{\beta}) \prod_{j=1}^3 P_j(Y_{ijt} = y_{ijt} \mid S_{it} = m, \delta_{jk[t]}) \right\} \times \cdots \\
 \times \phi\left(\frac{\beta_0}{5}\right) \prod_{j=1}^3 \left\{ \phi(\sigma_{\delta_j}) \prod_{m=1}^2 \phi\left(\frac{\mu_{jm}}{5}\right) \right\} \prod_{i=1}^N \phi\left(\frac{\eta_i}{\sigma_\eta}\right) \prod_{k=1}^{18} \phi_3(\boldsymbol{\Sigma}_\delta^{1/2} \boldsymbol{\delta}_k) \prod_{l=1}^6 \phi(\beta_l)
 \end{aligned}$$

where $\phi(\cdot)$ is the standard normal density, $\phi_3(\cdot)$ is the standard multivariate normal density with three components, and $\boldsymbol{\sigma}_\delta = (\sigma_{\delta_j} : j = 1, \dots, 3)$ represents the three diagonal elements of the covariance matrix $\boldsymbol{\Sigma}_\delta$.

The Stan code is written as

```

data {
  int<lower=1> nworked; //number of worked items--rows of data
  int<lower=1> nprob;    // number of items
  int<lower=1> nstud;    // number of respondents
  int<lower=1> ncov;    // number of person-level covariates
  int<lower=0> hint[nworked];
  int<lower=0> err[nworked];
  real ltime[nworked];
  int<lower=1,upper=nprob> prob[nworked];
  int<lower=1,upper=nstud> stud[nworked];
  matrix[nstud,ncov] X;
  vector[3] zeros;
  real<lower=0> sigStud;
}

```

```

1 parameters {
2   real meanTime[2];
3   real<lower=0> sigTime[2];
4   real effHint[2];
5   real effErr[2];
6   vector[3] probEff[nprob]; // hint, err, time
7   corr_matrix[3] OmegaProb;
8   vector<lower=0>[3] sigProb;
9   real alpha;
10  vector[nstud] studEff;
11  vector[ncov] beta;
12  ordered[2] cHint;
13  ordered[2] cErr;
14 }
15 transformed parameters {
16   vector[nstud] nu=inv_logit(alpha+X*beta+studEff);
17   cov_matrix[3] SigmaProb=quad_form_diag(OmegaProb, sigProb);
18 }
19 model{
20   // priors
21   meanTime~normal(0,5);
22   sigTime~normal(0,5);
23   effHint~normal(0,5);
24   effErr~normal(0,5);
25   sigProb~normal(0,1);
26   to_vector(beta)~normal(0,1);
27   studEff~normal(0,sigStud);
28   probEff~multi_normal(zeros,SigmaProb);
29

```

```

1  //measurement model
2  for(w in 1:nworked)
3    target += log_sum_exp(
4      log(nu[stud[w]])+
5      ordered_logistic_lpmf(hint[w] | probEff[prob[w]][1]+effHint[1],cHint)+
6      ordered_logistic_lpmf(err[w] | probEff[prob[w]][2]+effErr[1],cErr)+
7      normal_lpdf(ltime[w] | probEff[prob[w]][3]+meanTime[1],sigTime[1]),
8      log(1-nu[stud[w]])+
9      ordered_logistic_lpmf(hint[w] | probEff[prob[w]][1]+effHint[2],cHint)+
10     ordered_logistic_lpmf(err[w] | probEff[prob[w]][2]+effErr[2],cErr)+
11     normal_lpdf(ltime[w] | probEff[prob[w]][3]+meanTime[2],sigTime[2])
12   );
13 }

```

14
15 This model syntax was saved in a file called `lca2class.stan`, and the data was encoded in
16 R in an object named `sdat`. Then we fit the model with the code:

```

17 mod <- stan('lca2class.stan',data=sdat,iter=2000,chains=8)

```

18 That is, we fit eight separate chains, beginning from randomly-chosen initial values,
19 running for 2,000 iterations each. By default, the first 1,000 of these iterations were
20 denoted as “warm-up,” and discarded. The hope is that the Markov chains would have
21 each converged on the posterior distribution during the warm-up iterations. We checked
22 the convergence of the algorithm by examining traceplots and the Gelman-Rubin $\hat{R}$
23 statistic (Gelman et al., 2013, § 11.4). Inference was based on the remaining 1,000
24 iterations in each chain, which we consider as samples from the posterior distribution.

25 **Relabeling Classes in LCM.** Two different runs of a Markov Chain Monte Carlo
26 estimation technique can return essentially equivalent posterior estimates, but with the
27 labels of the latent classes switched. Here we illustrate our *ad hoc* solution to label
28 switching.

First, some parameters in the model are invariant to label switching, such as the optimized log posterior, denoted `lp_` in `rstan`. We began assessing model convergence by inspecting `lp_`. If this parameter has converged, we re-label the classes and assess convergence on the remaining parameters.

To re-label, we first examined the posterior means for each of the four measurement parameters, μ_1 , μ_2 , and μ_3 , location parameters for the distributions of time spent, errors made, and hints requested, respectively, and σ , parameterizing the scale of the distribution of time spent, for each of the two latent classes, as estimated by the 8 Markov chains. These are reproduced in Table A2.

Table A2

Estimated Posterior Means

Chain	μ_1		σ		μ_2		μ_3	
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
1	-0.50	0.99	0.76	0.84	-1.76	2.43	-6.29	6.00
2	0.91	-0.58	0.84	0.76	1.93	-2.26	6.18	-6.17
3	1.02	-0.48	0.84	0.76	1.85	-2.34	6.14	-6.11
4	-0.46	1.03	0.76	0.84	-2.07	2.12	-5.95	6.32
5	-0.50	0.99	0.76	0.84	-2.60	1.60	-6.14	6.14
6	1.00	-0.49	0.84	0.76	2.53	-1.66	5.70	-6.53
7	-0.51	0.98	0.76	0.84	-1.85	2.34	-5.78	6.57
8	-0.48	1.01	0.76	0.84	-2.19	2.00	-5.95	6.24

Note that for parameter μ_1 , chains 1, 4, 5, 7, and 8 have means between -.51 and -.46 for class 1 and between 0.98 and 1.03 for class 2, but that these values are reversed for chains 2, 3, and 6. Similar patterns hold for measurement parameters σ , μ_2 , and μ_3 . This discrepancy between the chains leads to the unacceptably large $\hat{R}$ values in the first row of Table A3. However, if we switch the labels of classes 1 and 2 for chains 2, 3, and 6, all chains roughly agree on all four parameters, leading to acceptable $\hat{R}$ values in the 2nd row of Table A3. Importantly, we switch the labels on the same chains—2, 3, and 6—for all parameters.

Table A3

Gelman-Rubin $\hat{R}$

	μ_{11}	μ_{12}	σ_1	σ_2	μ_{21}	μ_{22}	μ_{31}	μ_{32}
Original	6.88	7.10	11.32	10.22	1.18	1.18	1.91	1.96
Relabeled	1.05	1.05	1.00	1.00	1.00	1.00	1.00	1.00

1 Once we have determined the chains whose labels need to be switched, we relabel those
 2 chains' estimates for $Pr(S_{it} = m | \eta_i, \mathbf{Z}_i = \mathbf{z}_i, \boldsymbol{\beta})$ by subtracting the initial draws from 1, and
 3 we correct posterior draws for structural model coefficients $\boldsymbol{\beta}$ by multiplying the original
 4 draws by -1 . Once the draws have been suitably transformed, we pooled the draws across
 5 all 8 chains to calculate the estimates in Tables 3 and 11.

6 Code for this full process can be found on the Github site.