

PREDICTIVE, SCALABLE AND INTERPRETABLE KNOWLEDGE TRACING ON STRUCTURED DOMAINS

Hanqi Zhou^{1,2,4}, Robert Bamler^{1,3}, Charley M. Wu^{1,2,3*} & Álvaro Tejero-Cantero^{1,2*}

¹University of Tübingen, ²Cluster of Excellence Machine Learning, ³Tübingen AI Center, ⁴IMPRS-IS
{hanqi.zhou, robert.bamler, charley.wu, alvaro.tejero}@uni-tuebingen.de

ABSTRACT

Intelligent tutoring systems optimize the selection and timing of learning materials to enhance understanding and long-term retention. This requires estimates of both the learner’s progress (“knowledge tracing”; KT), and the prerequisite structure of the learning domain (“knowledge mapping”). While recent deep learning models achieve high KT accuracy, they do so at the expense of the interpretability of psychologically-inspired models. In this work, we present a solution to this trade-off. PSI-KT is a hierarchical generative approach that explicitly models how both individual cognitive traits and the prerequisite structure of knowledge influence learning dynamics, thus achieving interpretability by design. Moreover, by using scalable Bayesian inference, PSI-KT targets the real-world need for efficient personalization even with a growing body of learners and learning histories. Evaluated on three datasets from online learning platforms, PSI-KT achieves superior multi-step predictive accuracy and scalable inference in continual-learning settings, all while providing interpretable representations of learner-specific traits and the prerequisite structure of knowledge that causally supports learning. In sum, predictive, scalable and interpretable knowledge tracing with solid knowledge mapping lays a key foundation for effective personalized learning to make education accessible to a broad, global audience.

1 INTRODUCTION

The rise of online education platforms has created new opportunities for personalization in learning, motivating a renewed interest in how humans learn structured knowledge domains. Foundational theories in psychology (Ebbinghaus, 1885) have informed *spaced repetition* schedules (Settles & Meeder, 2016), which exploit the finding that an optimal spacing of learning sessions enhances memory retention. Yet beyond the timing of rehearsals, the sequential order of learning materials is also crucial, as evidenced by curriculum effects in learning (Dewey, 1910; Dekker et al., 2022), where exposure to simpler, prerequisite concepts can facilitate the apprehension of higher-level ideas. Cognitive science and pedagogical theories have long emphasized the relational structure of knowledge in human learning (Rumelhart, 2017; Piaget, 1970), with recent research showing that mastering prerequisites enhances concept learning (Lynn & Bassett, 2020; Karuza et al., 2016; Brändle et al., 2022). Yet, we still lack a predictive, scalable, and interpretable model of the structural-temporal dynamics of learning that could be used to develop future intelligent tutoring systems.

Here, we present PSI-KT, a novel approach for inferring interpretable learner-specific cognitive traits and a shared knowledge graph of prerequisite concepts. We demonstrate our approach on three real-world educational datasets covering structured domains, where our model outperforms existing baselines in terms of *predictive* accuracy (both within- and between-learner generalization), *scalability* in a continual learning setting, and *interpretability* of learner traits and prerequisite graphs. The paper is organized as follows: We first introduce the knowledge tracing problem and summarize related work (Sec. 2). We then provide a formal description of PSI-KT and describe the inference method (Sec. 3). Experimental evaluations are organized into demonstrations of prediction performance, scalability, and interpretability (Sec. 4). Altogether, PSI-KT bridges machine learning and cognitive science, leveraging our understanding of human learning to build the foundations for automated tutoring systems with broad educational applications.

*Equal contribution. Code at github.com/mlcolab/psi-kt

2 BACKGROUND

In this section, we begin by defining the knowledge tracing problem and then review related work.

2.1 KNOWLEDGE TRACING FOR INTELLIGENT TUTORING SYSTEMS

For almost 100 years (Pressey, 1926), researchers have developed intelligent tutoring systems (ITS) to support human learning through adaptive teaching materials and feedback. More recently, *knowledge tracing* (KT; Corbett & Anderson, 1994) emerged as a method for tracking learning progress by predicting a learner’s performance on different *knowledge components* (KCs), e.g., the ‘Pythagorean theorem’, based on past learning interactions. Here, we focus on the KT problem, with the goal of supporting the selection of teaching materials in future ITS applications.

In this setting, a learner ℓ receives exercises or flashcards for KCs $x_n^\ell \in \{0, 1, \dots, K\}$ at irregularly spaced times t_n^ℓ , whereupon the performance is recorded, often as correct/incorrect, $y_n^\ell \in \{0, 1\}$. We can formalize KT as a supervised learning problem on time-series data, where the goal of the KT model is to predict future performance (e.g., \hat{y}_{N+1}) given all or part of the interaction history $\mathcal{H}_{1:N}^\ell := \{x_n^\ell, t_n^\ell, y_n^\ell\}_{n=1}^N$ available up to time t_N^ℓ . As part of the process, a KT model may infer specific representations of learners or of the learning domain to help prediction. If these representations are interpretable, they can be valuable for downstream learning personalization.

2.2 RELATED WORK

We broadly categorize related KT approaches into psychological and deep learning methods.

Psychological methods. Focusing on interpretability, psychological methods use domain knowledge to describe the temporal decay of memory (e.g., forgetting curves; Ebbinghaus, 1885), sometimes also modeling learner-specific characteristics. *Factor-based regression* models use hand-crafted features based on learner interactions and KC properties (e.g., repetition counts and KC easiness; Pavlik Jr et al., 2009). While they model KC-dependent memory dynamics (Pavlik et al., 2021; Gervet et al., 2020; Lindsey et al., 2014; Lord, 2012; Ackerman, 2014), they ignore the relational structure between KCs. Half-life Regression (HLR; Settles & Meeder, 2016) from Duolingo uses both correct and incorrect counts, while the Predictive Performance Equation (PPE; Walsh et al., 2018) models the elapsed time of every past interaction with a power function to account for spacing effects. By using shallow regression models with predefined features, these models achieve interpretability, but sacrifice prediction accuracy. *Latent variable models* use a probabilistic two-state Hidden Markov Model (Käser et al., 2017; Sao Pedro et al., 2013; Baker et al., 2008; Yudelson et al., 2013), representing either mastery or non-mastery of a given KC. These models are limited to binary states by design, do not account for learner dynamics, and for some, their numerous parameters hinder scalability. Another probabilistic model, HKT (Wang et al., 2021) accounts for structure and dynamics by modeling knowledge evolution as a multivariate Hawkes process. Close in spirit to our PSI-KT, this approach tracks KC structure but lacks any learner-specific representations.

Deep learning methods. Deep learning methods use flexible models with many parameters to achieve high prediction accuracy. However, this flexibility also makes it difficult to interpret their learned internal representations. The first deep learning methods explicitly modeled sequential interactions with *recurrent neural networks* to overcome the dependence on fixed summary statistics in simpler regression models, with Deep Knowledge Tracing (DKT; Piech et al., 2015) pioneering the use of Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997). A similar architecture, DKTF (Nagatani et al., 2019) incorporated additional input features, whereas Shen et al. (2021) proposed an intricate modular architecture aimed at recovering interpretable learner representations, but neglecting KC relations. *Structure-aware models* leverage KC dependencies, accounting for the fact that human knowledge acquisition is structured by dependency relationships (i.e., concept maps; Hill, 2005; Koponen & Nousiainen, 2018; Lynn & Bassett, 2020). Tong et al. (2020) empirically estimate KC dependencies from the frequencies of successful transitions. AKT (Ghosh et al., 2020) relies on the attention mechanism (Vaswani et al., 2017) to implicitly capture structure (Pandey & Karypis, 2019; Choi et al., 2020; Shin et al., 2021; Liu et al., 2023), whereas GKT (Nakagawa et al., 2019) models it explicitly based on graph neural networks (Kipf & Welling, 2016). Recent work towards interpretable deep learning KT uses engineered features such as learner mastery and exercise difficulty (Minn et al., 2022), or infers them with neural networks (QIKT; Chen et al., 2023, IEKT; Long et al., 2021). While diverse approaches to interpretability exist (see Chen et al., 2023, for review), a comprehensive evaluation framework is still lacking.

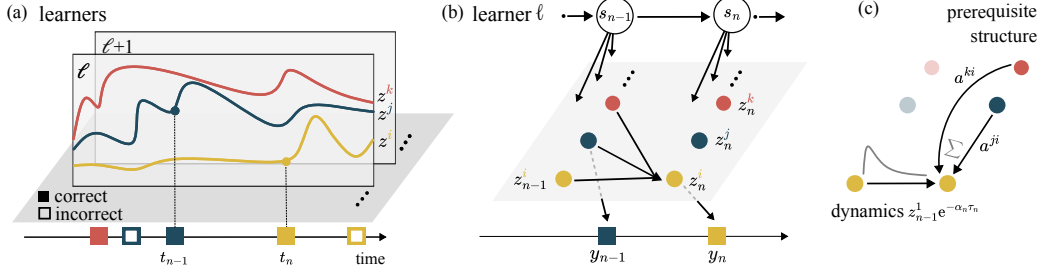


Figure 1: PSI-KT is a hierarchical probabilistic state-space model of learning. (a) Latent knowledge states for different KCs (colored curves) are inferred from observations. (b) Full hierarchical model for a single learner: cognitive traits s_n control the coupled dynamics of states z_n^k , which give rise to observations y_n . (c) The dynamics combine memory decay (Eq. 6) and structural influences (Eq. 5).

Here, we present our predictive, scalable and interpretable KT model (PSI-KT) as a psychologically-informed probabilistic deep learning approach, together with a comprehensive evaluation framework for interpretability.

3 JOINT DYNAMICAL AND STRUCTURAL MODEL OF LEARNING

In this section, we describe PSI-KT, our probabilistic hierarchical state-space model of human learning (Fig. 1). Briefly, observations of learner performance y (Fig. 1a, filled/unfilled boxes) provide indirect and noisy evidence about latent knowledge states z (colored curves, with matching dots in Fig. 1b). These latent states evolve stochastically, in line with the psychophysics of memory (temporal decay in Fig. 1c), while simultaneously being subject to structural influences from performance on prerequisite KCs (structure in Fig. 1c). We introduce a second latent level of learner-specific traits s (Fig. 1b, top), which govern the knowledge dynamics in an interpretable way.

Below, we describe the method in more detail. We start with the generative model (Sec. 3.1). Next, we discuss the joint approximate Bayesian inference of latent variables and estimation of generative parameters (Sec. 3.2). Finally, we show how to derive multi-step performance predictions (see Sec. 3.3 and Fig. 7 in Appendix A.4 for a graphical overview of inference and prediction).

3.1 PROBABILISTIC STATE-SPACE GENERATIVE MODEL

We conceptualize observations of learner performance as noisy measurements of an underlying time-dependent knowledge state, specific to each learner and KC. The evolution of knowledge states reflects the process of learning and forgetting, governed by learner-specific traits. Additionally, knowledge of different KCs informs one another according to learned prerequisite relationships. We translate these modeling assumptions into a generative model consisting of three main components: (i) the learner knowledge state across KCs, $\mathbf{z}_n^\ell = [z_n^{\ell,1} \dots z_n^{\ell,K}]^\top \in \mathbb{R}^K$ (colored curves in Fig. 1a), (ii) learner-specific cognitive traits $s_n^\ell \in \mathbb{R}^4$ (top row in Fig. 1b), and (iii) a shared static graph \mathcal{A} of KCs whose edges a^{ik} quantify the probability for a KC i to be a prerequisite for KC k (Fig. 1c).

State-space model. State-space models (SSMs) are a framework for partially observable dynamical processes. They represent the inherent noise of measurements y by an emission distribution $p(y_n | z_n)$, separate from the stochasticity of state dynamics, modeled as a first-order Markov process with transition probabilities $p(z_n | z_{n-1})$. The state dynamics are initiated by sampling from an initial prior $p(z_1)$ to iteratively feed the transition kernel, and predictions can be drawn at any time from the emission distribution. To represent the influence of individual cognitive traits over the knowledge dynamics, we additionally condition the z -transitions on the traits s (which also can be observed only indirectly). The three-level SSM hierarchy of PSI-KT consists of:

$$\text{Level 2 (latent cognitive traits):} \quad s_n^\ell \sim p_\theta(s_n^\ell | s_{n-1}^\ell) := \mathcal{N}(s_n^\ell | H s_{n-1}^\ell, R) \quad (1)$$

$$\text{Level 1 (latent knowledge states):} \quad z_n^\ell \sim p_\theta(z_n^\ell | z_{n-1}^\ell, s_n^\ell) := \prod_k \mathcal{N}(z_n^{\ell,k} | m_n^{\ell,k}, w_n^\ell) \quad (2)$$

$$\text{Level 0 (observed learner performance):} \quad \hat{y}_n^\ell \sim p(y_n^\ell | z_n^{\ell,k}) := \text{Bern}(\text{sigmoid}(z_n^{\ell,k})). \quad (3)$$

The choice of Gaussian initial priors (discussed below) and Gaussian transitions ensures tractability, while the Bernoulli emissions model the observed binary outcomes. We now unpack this model and all its parameters in detail, starting with the knowledge dynamics.

Knowledge states z . Recent KT methods (e.g., Nagatani et al., 2019) use an exponential forgetting function based on psychological theories (Ebbinghaus, 1885). Here, we augment this approach by adding stable long-term memory (Averell & Heathcote, 2011), and model the knowledge dynamics $z^{\ell,k}$ of an isolated KC k as a mean-reverting stochastic (Ornstein-Uhlenbeck; OU) process:

$$dz^{\ell,k}/dt = \alpha^\ell(\mu^\ell - z^{\ell,k}) + \sigma^\ell\eta(t). \quad (4)$$

Accordingly, the state of knowledge z^ℓ gradually reverts to a long-term mean μ^ℓ with rate α^ℓ , subject to white noise fluctuations $\eta(t)$ scaled by volatility σ^ℓ . To account for the influence of other KCs, we adjust the mean μ_n^ℓ using prerequisite weights a^{ik} (defined in Eq. 7 below), modulated by the learner’s transfer ability γ_n^ℓ :

$$\tilde{\mu}_n^{\ell,k} := \mu_n^\ell + (\gamma_n^\ell/K) \sum_{i \neq k} a^{ik} z_n^{\ell,i}. \quad (5)$$

We obtain the mean $m_n^{\ell,k}$ and variance w_n^ℓ of the transition kernel in Eq. 2 by marginalizing the OU process over one time step $\tau_n^\ell := t_n^\ell - t_{n-1}^\ell$, which can be done analytically¹,

$$m_n^{\ell,k} = r_n^\ell z_{n-1}^{\ell,k} + (1 - r_n^\ell) \tilde{\mu}_n^{\ell,k}, \quad \text{with retention ratio } r_n^\ell := e^{-\alpha_n^\ell \tau_n^\ell} \in (0, 1). \quad (6)$$

As the time since the last interaction τ_n^ℓ grows, the retention ratio r_n^ℓ decreases exponentially with rate α_n^ℓ , and the knowledge state reverts to the long-term mean $\tilde{\mu}_n^{\ell,k}$, which partly depends on the learner’s mastery of prerequisite KCs (Eq. 5). This balances short-term and long-term learning, reflecting empirical findings from memory research (Averell & Heathcote, 2011). The structural influences are accounted for in the dynamics of $z_n^{\ell,k}$, thus justifying the conditional independence assumed in Eq. 2. A Gaussian initial prior $p_\theta(z_1^{\ell,k}) = \mathcal{N}(z_1^{\ell,k} | \bar{z}, w_1)$, where $\bar{z}, w_1 \in \mathbb{R}$ are part of the generative parameters θ , completes our dynamical model of knowledge states.

Learner-specific cognitive traits s . The dynamics of knowledge states (Eqs. 4-6) are parameterized by learner-specific cognitive traits $(\alpha_n^\ell, \mu_n^\ell, \sigma_n^\ell, \gamma_n^\ell)$, which we collectively denote s_n^ℓ . Specifically, α^ℓ represents the forgetting rate (Ebbinghaus, 1885; Averell & Heathcote, 2011), μ^ℓ (via $\tilde{\mu}_k^{\ell,n}$) captures long-term memory consolidation (Meeter & Murre, 2004) for practiced KCs and expected performance for novel KCs, σ^ℓ quantifies knowledge volatility, and γ^ℓ measures transfer ability (Bassett & Mattar, 2017) from knowledge of prerequisite KCs. These traits can develop during learning according to Eq. 1, starting from a Gaussian prior $p_\theta(s_1^\ell) = \mathcal{N}(s_1^\ell | \bar{s}, R_1)$ where $\bar{s} \in \mathbb{R}^4$ and the diagonal matrices $H, R_1, R \in \mathbb{R}^{4 \times 4}$ are also part of the global parameters θ .

Shared prerequisite graph \mathcal{A} . In our model, prerequisite relations influence knowledge dynamics via the coupling introduced in Eq. 5. We now discuss an appropriate parameterization for the weight matrix of the prerequisite graph, $\mathcal{A} := \{a^{ik}\}_{i,k \in 1:K}$. We assume that prerequisites are time- and learner-independent so that, in the spirit of collaborative filtering (Breese et al., 2013), we can pool evidence from all learners to estimate them. To prevent a quadratic scaling in the number of KCs, we do not directly model edge weights but derive them from KC embedding vectors u^k in lower dimension $u^k \in \mathbb{R}^D$ with $D \ll K$, collected in embedding matrix $U_{K \times D}$. A basic integrity constraint for a connected pair is that dependence of KC i on KC k should trade off against that of k on i , i.e., no mutual prerequisites: $a^{ik} + a^{ki} = 1$. With this in mind, we exploit the factorization of a^{ik} introduced by Lippe et al. (2021) in terms of a separate probability of edge existence $p(i \circ\!\circ\! k)$ and definite directionality $p(i \rightarrow k | i \circ\!\circ\! k)$:

$$\begin{aligned} a^{ik} &:= p(i \rightarrow k | i \circ\!\circ\! k) p(i \circ\!\circ\! k) \\ &= \text{sigmoid}((u^i)^\top u^k) \text{sigmoid}((u^i)^\top (M - M^\top) u^k), \end{aligned} \quad (7)$$

where the skew-symmetric combination $M - M^\top$ of a learnable matrix M prevents mutual prerequisites. Having presented the generative model, we now turn to inference and prediction.

3.2 APPROXIMATE BAYESIAN INFERENCE AND AMORTIZATION WITH A NEURAL NETWORK

We now describe how we learn the generative model parameters θ and how we infer the latent states s, z introduced in Section 3.1 using a neural network (“inference network”). Since learner-specific latent states s and z are deducible solely from limited individual data, we expect non-negligible uncertainty. This motivates our probabilistic treatment of these states using approximate Bayesian inference. By contrast, the model parameters θ (KC parameters U, M in Eq. 7, transition parameters \bar{s}, H, R_1, R in Eq. 1, and \bar{z}, w_1 in Eq. 2) can be estimated from all learners, and we thus

¹Särkkä & Solin (2019) — the variance is $w_n^\ell = (\sigma_n^\ell)^2 (1 - e^{-2\alpha_n^\ell \tau_n^\ell}) / (2\alpha_n^\ell)$.

treat them as point-estimated parameters as described below (detailed derivation in Appendix A.1.) Here, without loss of generality, we show the inference for a single learner.

3.2.1 INFERENCE ON A FIXED LEARNING HISTORY

Here, we assume the full interaction history $\mathcal{H}_{1:N}^\ell$ is available for inferring the posterior over latents $p_\theta(\mathbf{z}_{1:N}^\ell, \mathbf{s}_{1:N}^\ell | \mathbf{y}_{1:N}^\ell)$. We approach the problem using variational inference (VI). In VI, we select a distribution family q_ϕ with free parameters ϕ to approximate the posterior p_θ by minimizing their Kullback-Leibler divergence. This can only be done indirectly, by maximizing a lower bound to the marginal probability of the data, the *evidence lower bound* (ELBO). Here, we adopt the mean-field approximation $q_\phi(\mathbf{z}_{1:N}^\ell, \mathbf{s}_{1:N}^\ell | \mathbf{y}_{1:N}^\ell) = q_\phi(\mathbf{z}_{1:N}^\ell) q_\phi(\mathbf{s}_{1:N}^\ell)$ and jointly optimize the generative θ and variational ϕ parameters using variational *expectation maximization* (EM; Dempster et al., 1977; Beal & Ghahramani, 2003; Attias, 1999). Motivated by real-world scalability, we introduce an inference network (see Appendix A.3 for the architecture) to amortize the learning of variational parameters ϕ across learners, and we employ the *reparametrization trick* (Kingma & Welling, 2014) to optimize the single-learner ELBO:

$$\begin{aligned} \text{ELBO}^\ell(\theta, \phi) = & \mathbb{E}_{q_\phi(\mathbf{s}_{1:N}^\ell)} \left[-\log q_\phi(\mathbf{s}_{1:N}^\ell) + \log p_\theta(\mathbf{s}_1^\ell) + \sum_{n=2}^N \log p_\theta(\mathbf{s}_n^\ell | \mathbf{s}_{n-1}^\ell) \right] \\ & + \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}^\ell)} \left[-\log q_\phi(\mathbf{z}_{1:N}^\ell) + \log p_\theta(\mathbf{z}_1^\ell) + \sum_{n=1}^N \log p_\theta(\mathbf{y}_n^\ell | \mathbf{z}_n^{\ell, x_n}) \right] \\ & + \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}^\ell) q_\phi(\mathbf{s}_{1:N}^\ell)} \left[\sum_{n=2}^N \log p_\theta(\mathbf{z}_n^\ell | \mathbf{z}_{n-1}^\ell, \mathbf{s}_n^\ell) \right]. \end{aligned} \quad (8)$$

The SSM emissions and transitions were introduced in Eqs. 1-3, along with the respective initial priors. To allow for a diversity of combinations of learner traits to account for the data, we model the variational posterior across learners, $q_\phi(\mathbf{s}_{1:N}^\ell)$, as a mixture of Gaussians (see Appendix A.4).

3.2.2 INFERENCE IN CONTINUAL LEARNING

In real-world educational settings, a KT model must flexibly adapt its current variational parameters ϕ_n with newly available interactions $(x_{n+1}^\ell, t_{n+1}^\ell, y_{n+1}^\ell)$. Retraining on a fixed, augmented history \mathcal{H}_{n+1}^ℓ to obtain an updated ϕ_{n+1} is possible (Eq. 8), but expensive. Instead, in PSI-KT, we use the parameters ϕ_n of the current posterior $q_{\phi_n}(\mathbf{z}_n^\ell, \mathbf{s}_n^\ell)$ to form a next-time prior,

$$\tilde{p}(\mathbf{z}_{n+1}^\ell, \mathbf{s}_{n+1}^\ell) := \mathbb{E}_{q_{\phi_n}(\mathbf{z}_n^\ell, \mathbf{s}_n^\ell | \mathbf{y}_{1:n}^\ell)} [p_\theta(\mathbf{z}_{n+1}^\ell | \mathbf{s}_n^\ell) p_\theta(\mathbf{z}_{n+1}^\ell | \mathbf{s}_{n+1}^\ell, \mathbf{z}_n^\ell)]. \quad (9)$$

Due to the Bayesian nature of our model, we can now update this prior with the new evidence \mathbf{y}_{n+1}^ℓ at time t_{n+1}^ℓ using *variational continual learning* (VCL; Nguyen et al., 2017; Loo et al., 2020), i.e., by maximizing the ELBO:

$$\begin{aligned} \text{ELBO}_{\text{VCL}}^\ell(\theta, \phi_{n+1}) = & \mathbb{E}_{q_{\phi_{n+1}}(\mathbf{s}_{n+1}^\ell)} \left[-\log q_{\phi_{n+1}}(\mathbf{s}_{n+1}^\ell) \right] \\ & + \mathbb{E}_{q_{\phi_{n+1}}(\mathbf{z}_{n+1}^\ell)} \left[-\log q_{\phi_{n+1}}(\mathbf{z}_{n+1}^\ell) + \log p_\theta(\mathbf{y}_{n+1}^\ell | \mathbf{z}_{n+1}^{\ell, x_{n+1}}) \right] \\ & + \mathbb{E}_{q_{\phi_{n+1}}(\mathbf{z}_{n+1}^\ell, \mathbf{s}_{n+1}^\ell)} \left[\log \tilde{p}(\mathbf{z}_{n+1}^\ell, \mathbf{s}_{n+1}^\ell) \right]. \end{aligned} \quad (10)$$

Maximizing this $\text{ELBO}_{\text{VCL}}^\ell$ allows us to update the parameters ϕ_{n+1} based on a new interaction $(x_{n+1}^\ell, t_{n+1}^\ell, y_{n+1}^\ell)$ directly from the previous parameters ϕ_n , i.e., without retraining.

3.3 PREDICTIONS

To predict a learner’s performance on KC x_{n+1}^ℓ at t_{n+1}^ℓ , we take the current variational distributions over \mathbf{s}_n^ℓ and \mathbf{z}_n^ℓ and transport them forward by analytically convolving them with the respective transition kernels (Eqs. 1 and 2). We then draw $\mathbf{z}_{n+1}^{\ell, x_{n+1}}$ from the resulting distribution, and predict the outcome \hat{y}_{n+1}^ℓ by Eq. 3. When predicting multiple steps ahead, we repeat this procedure without conditioning on any of the previously predicted \hat{y}_{n+m}^ℓ .

4 EVALUATIONS

We argue above that KT for personalized education must predict accurately, scale well with new data, and provide interpretable representations. We now empirically assess these desiderata, comparing PSI-KT with up to 8 baseline models across three datasets from online education platforms. Concretely, we evaluate (i) prediction accuracy,

Table 1: Dataset characteristics

Dataset →	Assist12	Assist17	Junyi15
# Learners L	46,674	1,709	247,606
# KCs K	263	102	722
# Int’s / 10^6	3.5	0.9	26




Figure 2: Within-learner prediction performance (mean \pm SEM) as a function of cohort sizes from 100 to the maximum available in each dataset (we omit HLR for legibility; see Table 2.)

quantifying both within-learner prediction and between-learner generalization (Sec. 4.1), (ii) scalability in a continual learning setting (Sec. 4.2), and (iii) interpretability of learner representations and prerequisite relations (Sec. 4.3).

Datasets. Assistments and Junyi Academy are non-profit online learning platforms for pre-college mathematics. We use Assistments’ 2012 and 2017 datasets² (Assist12 and Assist17) and Junyi’s 2015 dataset³ (Junyi15; Chang et al., 2015), which in addition to interaction data, provides human-annotated KC relations (see Table 1 and Appendix A.3.2 for details).

We select HLR from Duolingo and PPE as two influential psychologically-informed regression models. From the models that use learnable representations, we include two established deep learning benchmarks, DKT and DKTF, which capture complex dynamics via LSTM networks, as well as the interpretability-oriented QIKT.

4.1 PREDICTION AND GENERALIZATION PERFORMANCE

In our evaluations, we mainly focus on prediction and generalization when training on 10 interactions from up to 1000 learners. Good KT performance with little data is key in practical ITS to minimize the number of learners on an experimental treatment (principle of equipoise, similar to medical research; Burkholder, 2021), to mitigate the cold-start problem, and to extend the usefulness of the model to classroom-size groups. To provide ITS with a basis for adaptive guidance and long-term learner assessment, we always predict the 10 next interactions. Figure 2 shows that PSI-KT’s *within-learner prediction performance* is robustly above baselines for all but the largest cohorts ($>60k$ learners, Junyi15), where all deep learning models perform similarly. The advantage of PSI-KT comes from its combined modeling of KC prerequisite relations and individual learner traits that evolve in time (see Appendix Fig. 13 for ablations). The *between-learner generalization* accuracy of the models above, when tested on 100 out-of-sample learners, is shown in Table 2, where fine-tuning indicates that parameters were updated using (10-point) learning histories from the unseen learners. PSI-KT shows overall superior generalization except on Junyi15 (when fine-tuning).

4.2 SCALABILITY IN CONTINUAL LEARNING

In addition to training on fixed historical data, we also conduct experiments to demonstrate PSI-KT’s scalability when iteratively retraining on additional interaction data from each learner. This parallels real-world educational scenarios, where learners are continuously learning (Sec. 3.2.2). Each model is initially trained on 10 interactions from 100 learners. We then incrementally provide one data point from each learner, and evaluate the training costs and prediction accuracy. Figure 3 shows PSI-KT requires the least retraining time, retains the best prediction accuracy, and thus achieves the most favorable cost-accuracy trade-off (details in Appendix A.5.3).

4.3 INTERPRETABILITY OF REPRESENTATIONS

We now evaluate the interpretability of both learner-specific cognitive traits s^ℓ and the prerequisite graphs \mathcal{A} . We first show that our model captures learner-specific and disentangled traits that correlate with behavior patterns. Next, we show that our inferred graphs best align with ground truth graphs, and the edge weights predict causal support on downstream KCs.

4.3.1 LEARNER-SPECIFIC COGNITIVE TRAITS

For each learner, PSI-KT infers four latent traits, each with a clear dynamical role specified by the OU process (Eqs. 5-6). In contrast, high-performance baselines (AKT, DKT, and DKTF) describe learners via 16-dimensional embeddings solely constrained by network architecture and loss minimization. Another model QIKT constructs 3-dimensional embeddings with each element connected

²<https://sites.google.com/site/assistmentsdata>

³<https://psl1cdatashop.web.cmu.edu/DatasetInfo?datasetId=1198>

Table 2: Prediction accuracy. FT indicates additional fine-tuning and \uparrow indicates larger values are better. The **best model performance** is in bold and the 2nd best is underlined.

Dataset	Experiment	HLR	PPE	DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
Assist12	Within \uparrow	.54 _{.03}	.65 _{.01}	.65 _{.03}	.60 _{.01}	.55 _{.01}	<u>.67</u> _{.02}	.63 _{.03}	.63 _{.03}	.68 _{.02}
	Between \uparrow	.50 _{.03}	.50 _{.02}	.55 _{.02}	.51 _{.01}	.54 _{.00}	.58 _{.02}	.61 _{.02}	<u>.60</u> _{.02}	.61 _{.03}
	w/ FT \uparrow	.52 _{.02}	.53 _{.01}	.58 _{.00}	.55 _{.01}	.55 _{.00}	<u>.61</u> _{.00}	.62 _{.02}	.60 _{.03}	.62 _{.02}
Assist17	Within	.45 _{.01}	.53 _{.02}	.57 _{.02}	.53 _{.03}	.52 _{.03}	.56 _{.02}	.56 _{.04}	.58 _{.02}	.63 _{.02}
	Between	.33 _{.03}	<u>.51</u> _{.02}	.51 _{.00}	.48 _{.00}	<u>.51</u> _{.02}	.47 _{.01}	.53 _{.02}	.50 _{.02}	.53 _{.02}
	w/ FT	.41 _{.04}	.51 _{.00}	.51 _{.03}	.53 _{.01}	.51 _{.03}	.51 _{.02}	<u>.54</u> _{.03}	.51 _{.04}	.56 _{.02}
Junyi15	Within	.55 _{.02}	.66 _{.03}	.79 _{.03}	.78 _{.01}	.63 _{.02}	.81 _{.02}	.78 _{.02}	.81 _{.02}	.83 _{.02}
	Between	.48 _{.02}	.55 _{.02}	.76 _{.00}	.76 _{.02}	.61 _{.01}	.73 _{.01}	<u>.77</u> _{.03}	.76 _{.03}	.79 _{.03}
	w/ FT	.52 _{.00}	.65 _{.03}	.81 _{.01}	.84 _{.01}	.64 _{.03}	.83 _{.00}	.79 _{.03}	.80 _{.03}	.80 _{.02}




Figure 3: Continual learning. (Top) Cumulative training time. (Bottom) Prediction accuracy on the next 10 time steps. We omit results when time is above, or accuracy is below, the range of the axes.

to scores of knowledge acquisition, knowledge mastery, and problem-solving. We collectively refer to these learner-specific variables as *learner representations*. Here, we empirically show that PSI-KT representations provide superior interpretability. We ask that learner representations be 1) *specific* to individual learners, 2) *consistent* when trained on partial learning histories, 3) *disentangled* (i.e., component-wise meaningful, as in Bengio et al., 2013), and 4) *operationally interpretable*, so that they can be used to personalize future curricula. We evaluate desiderata 1-3 with information-theoretic metrics (Table 3; see Appendix A.6 for details), and desideratum 4 with regressions against behavioral outcomes (Table 4).

Specificity, consistency, and disentanglement.

Learner representations s should be maximally *specific* about learner identity ℓ , which can be quantified by the mutual information $\text{MI}(s; \ell) = H(s) - H(s | \ell)$ being high, where H denotes (conditional) entropy. Additionally, when we infer representations $s^{\ell_{\text{sub}}}$ from different subsets of the interactions of a fixed learner, they should be *consistent*, i.e., each $s^{\ell_{\text{sub}}}$ should be minimally informative about the chosen subset (averaged across subsets), such that $\mathbb{E}_{\ell_{\text{sub}}} \text{MI}(s^{\ell}; \ell_{\text{sub}}) = \mathbb{E}_{\ell_{\text{sub}}} [H(s | \ell) - H(s | \ell_{\text{sub}})]$ should be low. Note that sequential subsets are unsuitable for this evaluation, since representations evolve in time to track learners’ progression. Instead, we define subsets as groups of KCs whose average presentation time is approximately uniform over the duration of the experiment (see Appendix A.6.1 for details). Lastly, learner representations should be *disentangled*, such that each dimension is individually informative about learner identity. We measure disentanglement with $D_{\text{KL}}(s || \ell) := H(s) - H(s | \ell)_{\text{diag}}$, a form of specificity that ignores correlations across s^{ℓ} dimensions by estimating the conditional entropy only with diagonal covariances.

In empirical evaluations (Table 3), PSI-KT’s representations offer competitive specificity despite being lower-dimensional, and outperform all baselines in consistency and disentanglement. While disentanglement aids interpretability (Freiesleben et al., 2022), it does not itself entail domain-specific meaning for representational dimensions. We now demonstrate that PSI-KT representations correspond to clear behavioral patterns, which is crucial for future applications in educational settings.

Operational interpretability. Having shown that PSI-KT captures specific, consistent, and disentangled learner features, we now investigate whether these features relate to meaningful aspects of future

Table 3: Specificity, consistency, and disentanglement vs. best baseline.

Metric	Dataset	Baseline	PSI-KT
Specificity $\text{MI}(s; \ell) \uparrow$	Assist12	8.8	<u>8.4</u>
	Assist17	10.1	<u>10.0</u>
	Junyi15	<u>13.5</u>	14.4
Consistency ⁻¹ $\mathbb{E}_{\ell_{\text{sub}}} \text{MI}(s^{\ell}; \ell_{\text{sub}}) \downarrow$	Assist12	<u>12.2</u>	7.4
	Assist17	6.4	6.4
	Junyi15	7.7	5.0
Disentanglement $D_{\text{KL}}(s \ell) \uparrow$	Assist12	2.3	7.4
	Assist17	0.6	8.4
	Junyi15	<u>5.0</u>	11.5

behavior, which would be useful for scheduling operations for ITS. We indeed find that the learner representations of PSI-KT forecast interpretable behavioral outcomes, such as performance decay or initial performance on novel KCs. Concretely, consider the observed *one-step performance difference* $\Delta y_n^\ell := y_n^\ell - y_{n-1}^\ell$. We expect it to be lower for longer intervals $\tau_n^\ell = t_n^\ell - t_{n-1}^\ell$ due to forgetting. However, we recognize no clear trend when plotting Δy_n^ℓ over τ_n^ℓ for the Junyi15 dataset (Fig. 4, top right). We can explain this observation because different learners forget on different time scales. Plotting the same test data instead over scaled intervals $\tau_n^\ell \alpha_n^\ell$ (Fig. 4, top center) shows a clear trend against an exponential fit (solid line) with less variability, demonstrating that α_n^ℓ (derived from past data only) adjusts for individual learner characteristics and can be interpreted as a personalized rate of forgetting. Here, the choice of the factor α_n^ℓ is motivated by our inductive bias (Eq. 4). The trend is much less clear for all baselines: Fig. 4 (top left) uses the best fitting component across all learner representations from all baselines (full results in Fig. 8 in Appendix A.6.4). Analogously, when we consider *initial performance on a novel KC*, we find for PSI-KT that $\tilde{\mu}_n^{\ell,k}$ (which aggregates mastery of prerequisites for KC k at time t_n , see Eq. 5) explains it better than the best baseline Fig. 4 (bottom panels). Table 4 shows that these superior interpretability results are significant and hold across all datasets. In Appendix A.6.4, we discuss two more behavioral signatures (performance variability and prerequisite influence) and show they correspond to the remaining components γ_n^ℓ and σ_n^ℓ .

4.3.2 PREREQUISITE GRAPH

PSI-KT infers a prerequisite graph based on all learners’ data, which helps it to generalize to unseen learners. Beyond helping prediction, reliable prerequisite relations are an essential input for curriculum design, motivating our interest in their interpretability. Figure 5a shows an exemplary inferred subgraph with the prerequisites of a single KC. To quantitatively evaluate the graph, we (i) measure the alignment of the inferred vs. ground-truth graphs and (ii) correlate inferred prerequisite probability with a Bayesian measure of causal support obtained from unseen behavioral data.

Alignment with ground-truth graphs. We analyze the Junyi15 dataset, which uniquely provides human-annotated evaluations of prerequisite and similarity relations between KCs. We discuss here the alignment of prerequisites and leave similarity for Appendix A.7. The Junyi15 dataset provides both an expert-identified prerequisite for each KC, and crowd-sourced ratings (6.6 ratings on average on a 1-9 scale). To compare with expert annotations, we compute the rank of each expert-identified prerequisite relation $i \rightarrow k$ in the relevant sorted list of inferred probabilities $\{a^{jk}\}_{j=1}^K$ and take the harmonic average (mean reciprocal rank, MRR; Yang et al., 2014). Next, we compute the negative log-likelihood (nLL) of inferred edges a^{ik} using a Gaussian estimate of the (rescaled) crowd-sourced ratings for the $i \rightarrow k$ KC pair. We finally calculate the Jaccard similarity (JS) between the set of inferred edges ($a^{ik} > 0.5$) and those identified by experts as well as crowd-sourced edges with average ratings above 5. The results in Table 5 (left columns) consistently highlight PSI-KT’s superior performance across all criteria (see Appendix A.7.1 for details).

Causal support across consecutive interactions. For education applications, we are interested in how KC dependencies impact learning effectiveness. If KC i is a prerequisite of KC k , mastering KC i contributes to mastering KC k , indicating a causal connection. In this analysis, we show that inferred edge probabilities a^{ij} (Eq. 7) correspond to causal support $\text{sup}_{i \rightarrow k}$ (Eq. 11), derived from behavioral data through Bayesian causal induction (Griffiths & Tenenbaum, 2009). Specifically, we model the relationship between a candidate cause C and effect E , i.e., a pair of KCs in our case, while accounting for a constant background cause B , representing the learner’s overall ability and the influences of other KCs. We consider two hypothetical causal graphs, where Graph 0 $G_{i \rightarrow k}$ represents




Figure 4: Operational interpretability of representations, Junyi15 dataset. See text for axes labels and Appendix A.6.4 for additional results.

Behavioural signature	Dataset	Best Baseline	PSI-KT
Performance difference	Assist12	0.01, .67	0.30 , <.001
	Assist17	-0.03, .30	0.56 , <.001
	Junyi15	0.03, .06	0.72 , <.001
Initial performance	Assist12	0.04, .01	0.54 , <.001
	Assist17	0.05, .01	3.70 , <.001
	Junyi15	0.04, .02	0.92 , <.001

Table 4: Coefficients and p -values of regressions relating $\exp(-\alpha_n^\ell \tau_n^\ell)$ and $\tilde{\mu}_n^{\ell,k}$ to unseen behavioral data across datasets.

Table 5: (Left) Alignment of inferred graphs with annotated graphs for the Junyi15 dataset. (Right) Regression coefficients and p -values relating causal support to inferred edge probabilities. All baseline models either lack significance or negatively predict causal support (Appendix Fig. 12).

Metric	MRR \uparrow	JS expert \uparrow	JS crowd \uparrow	nLL \downarrow	coefficient \uparrow , p -value \downarrow		
	Junyi15				Assist12	Assist17	Junyi15
Best Baseline	.0082	.0015	.0047	3.03	1.05, .253	0.22, .792	0.42, .593
PSI-KT	.0086	.0019	.0095	4.11	1.15 , .003	0.28 , <.001	0.97 , <.001

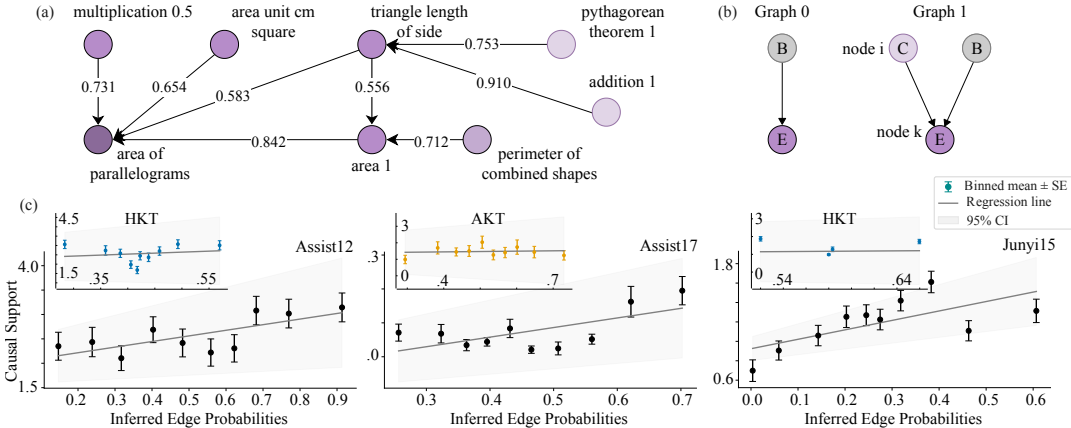


Figure 5: Graph interpretability. (a) Subgraph inferred by PSI-KT on the Junyi15 dataset, showing prerequisites of target KC ‘area of parallelograms’. (b) Hypothesized causal graphs, where Graph 1 assumes a causal relationship exists from KC i to KC k , while Graph 0 is the null hypothesis. (c) Regression of edge probabilities against causal supports. Insets show the best baseline model.

the null hypothesis of no causal relationship, and Graph 1 $G_{i \rightarrow k}$ assumes the causal relationship exists, i.e. correct performance on KC i causally supports correct performance on KC k (Fig. 5b). We estimate causal support for each pair of KCs $i \rightarrow k$ based on all consecutive interactions in the behavioral data \mathcal{H} from KC i at time t_n to KC k at time t_{n+1} , as a function of the difference in log-likelihoods of the two causal graphs (see Appendix A.7.3 for details):

$$\text{support}_{i \rightarrow k} := \log P(\mathcal{H} | G_{i \rightarrow k}) - \log P(\mathcal{H} | G_{i \nrightarrow k}). \quad (11)$$

We then use regression to predict $\text{support}_{i \rightarrow k}$ as a function of edge probabilities a^{ij} inferred from different models. The results are visualized in Figure 5c and summarized in Table 5 (right). The larger coefficients indicate that our inferred graphs possess superior operational interpretability (Sec. 4.3).

5 DISCUSSION

We propose PSI-KT as a novel approach to knowledge tracing (KT) with compelling properties for intelligent tutoring systems: superior predictive accuracy, excellent continual-learning scalability, and interpretable representations of learner traits and prerequisite relationships. We further find that PSI-KT has remarkable predictive performance when trained on small cohorts whereas baselines require training data from at least 60k learners to reach similar performance. An open question for future KT research is how to combine PSI-KT’s unique continual learning and interpretability properties with performance that grows beyond this extreme regime. We use an analytically marginalizable Ornstein-Uhlenbeck process for knowledge states in PSI-KT, resulting in an exponential forgetting law, similar to most recent KT literature. Future work should support ongoing debates in cognition by offering alternative modeling choices for memory decay (e.g., power-law; Wixted & Ebbesen, 1997), thus facilitating empirical studies at scale. And while our model already normalizes reciprocal dependencies in the prerequisite graph, we anticipate that enforcing regional or global structural constraints, such as acyclicity, may benefit inference and interpretability. Although we designed PSI-KT with general structured domains in mind, our empirical evaluations were limited to mathematics learning by dataset availability. We highlight the need for more diverse datasets for structured KT research to strengthen representativeness in ecologically valid contexts. Overall, our work combines machine learning techniques with insights from cognitive science to derive a predictive and scalable model with psychologically interpretable representations, thus laying the foundations for personalized and adaptive tutoring systems.

ACKNOWLEDGMENTS

The authors thank Nathanael Bosch and Tim Xiao for their helpful discussion, and Seth Axen for code review. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Hanqi Zhou. This research was supported as part of the LEAD Graduate School & Research Network, which is funded by the Ministry of Science, Research and the Arts of the state of BadenWürttemberg within the framework of the sustainability funding for the projects of the Excellence Initiative II. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC number 2064/1 – Project number 390727645. CMW is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A.

ETHICS STATEMENT

We evaluated our PSI-KT model on three public datasets from human learners, which all anonymize the data to protect the identities of individual learners. Although PSI-KT aims to improve personalized learning experiences, it infers cognitive traits from behavioral data instead of using learners' demographic characteristics (e.g., age, gender, and the name of schools provided in the Assistment 17 dataset), to avoid reinforcing existing disparities.

Evaluations of structured knowledge tracing in our paper are limited by dataset availability to pre-college mathematics. To ensure a broader and more ecologically valid assessment, it is essential to explore diverse datasets across various domains (e.g., biology, chemistry, linguistics) and educational stages (from primary to college level). This will allow for a more comprehensive understanding of the role of structure in learning.

REFERENCES

- Terry A Ackerman. Multidimensional item response theory models. *Wiley StatsRef: Statistics Reference Online*, 2014.
- Hagai Attias. A variational bayesian framework for graphical models. *Advances in neural information processing systems*, 12, 1999.
- Lee Averell and Andrew Heathcote. The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55:25–35, 02 2011.
- Ryan SJ d Baker, Albert T Corbett, and Vincent Aleven. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *Intelligent Tutoring Systems: 9th International Conference, ITS 2008, Montreal, Canada, June 23-27, 2008 Proceedings 9*, pp. 406–415. Springer, 2008.
- Danielle S Bassett and Marcelo G Mattar. A network neuroscience of human learning: potential to inform quantitative theories of brain and behavior. *Trends in cognitive sciences*, 21(4):250–264, 2017.
- Matthew J Beal and Zoubin Ghahramani. The variational bayesian em algorithm for incomplete data: with application to scoring graphical model structures. *Bayesian statistics*, 7:453–464, 2003.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- F Brändle, M Binz, and E Schulz. Exploration beyond bandits. In *The Drive for Knowledge: The Science of Human Information-Seeking*, pp. 147–168. Cambridge University Press, 2022.
- John S Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. *arXiv preprint arXiv:1301.7363*, 2013.
- Leslie Burkholder. Equipoise and ethics in educational research. *Theory and Research in Education*, 19(1):65–77, 2021. doi: 10.1177/14778785211009105.

- Haw-Shiuan Chang, Hwai-Jung Hsu, and Kuan-Ta Chen. Modeling exercise relationships in e-learning: A unified approach. In *EDM*, pp. 532–535, 2015.
- Jiahao Chen, Zitao Liu, Shuyan Huang, Qiongqiong Liu, and Weiqi Luo. Improving interpretability of deep sequential knowledge tracing models with question-centric cognitive representations. *arXiv preprint arXiv:2302.06885*, 2023.
- Youngduck Choi, Younghan Lee, Junghyun Cho, Jineon Baek, Byungsoo Kim, Yeongmin Cha, Dongmin Shin, Chan Bae, and Jaewe Heo. Towards an appropriate query, key, and value computation for knowledge tracing. In *Proceedings of the seventh ACM conference on learning@ scale*, pp. 341–344, 2020.
- Albert T Corbett and John R Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4:253–278, 1994.
- Ronald B Dekker, Fabian Otto, and Christopher Summerfield. Curriculum learning for human compositional generalization. *Proceedings of the National Academy of Sciences*, 119(41):e2205582119, 2022.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (methodological)*, pp. 1–38, 1977.
- John Dewey. *The child and the curriculum*. University of Chicago Press Chicago, 1910.
- Nat Dilokthanakul, Pedro AM Mediano, Marta Garnelo, Matthew CH Lee, Hugh Salimbeni, Kai Arulkumaran, and Murray Shanahan. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016.
- H. Ebbinghaus. *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Duncker & Humblot, Leipzig, 1885.
- Timo Freiesleben, Gunnar König, Christoph Molnar, and Alvaro Tejero-Cantero. Scientific inference with interpretable machine learning: Analyzing models to learn about real-world phenomena. *arXiv preprint arXiv:2206.05487*, 2022.
- Theophile Gervet, Ken Koedinger, Jeff Schneider, Tom Mitchell, et al. When is deep learning the best approach to knowledge tracing? *Journal of Educational Data Mining*, 12(3):31–54, 2020.
- Aritra Ghosh, Neil Heffernan, and Andrew S Lan. Context-aware attentive knowledge tracing. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2330–2339, 2020.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- Thomas L Griffiths and Joshua B Tenenbaum. Structure and strength in causal induction. *Cognitive psychology*, 51(4):334–384, 2005.
- Thomas L Griffiths and Joshua B Tenenbaum. Theory-based causal induction. *Psychological review*, 116(4):661, 2009.
- Lilian H Hill. Concept mapping to encourage meaningful student learning. *Adult Learning*, 16(3-4): 7–13, 2005.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Elisabeth A Karuza, Sharon L Thompson-Schill, and Danielle S Bassett. Local patterns to global architectures: influences of network topology on human learning. *Trends in cognitive sciences*, 20(8):629–640, 2016.

- Tanja Käser, Severin Klingler, Alexander G Schwing, and Markus Gross. Dynamic bayesian networks for student modeling. *IEEE Transactions on Learning Technologies*, 10(4):450–462, 2017.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International Conference on Machine Learning*, pp. 2649–2658. PMLR, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Ismo T Koponen and Maija Nousiainen. Concept networks of students’ knowledge of relationships between physics concepts: Finding key concepts and their epistemic support. *Applied network science*, 3(1):1–21, 2018.
- Robert V Lindsey, Jeffery D Shroyer, Harold Pashler, and Michael C Mozer. Improving students’ long-term knowledge retention through personalized review. *Psychological science*, 25(3):639–647, 2014.
- Phillip Lippe, Taco Cohen, and Efstratios Gavves. Efficient neural causal discovery without acyclicity constraints. *arXiv preprint arXiv:2107.10483*, 2021.
- Zitao Liu, Qionqiong Liu, Jiahao Chen, Shuyan Huang, and Weiqi Luo. simplekt: a simple but tough-to-beat baseline for knowledge tracing. *arXiv preprint arXiv:2302.06881*, 2023.
- Ting Long, Yunfei Liu, Jian Shen, Weinan Zhang, and Yong Yu. Tracing knowledge state with individual cognition and acquisition estimation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 173–182, 2021.
- Noel Loo, Siddharth Swaroop, and Richard E Turner. Generalized variational continual learning. *arXiv preprint arXiv:2011.12328*, 2020.
- Frederic M Lord. *Applications of item response theory to practical testing problems*. Routledge, 2012.
- Christopher W Lynn and Danielle S Bassett. How humans learn and represent networks. *Proceedings of the National Academy of Sciences*, 117(47):29407–29415, 2020.
- Martijn Meeter and Jaap MJ Murre. Consolidation of long-term memory: evidence and alternatives. *Psychological Bulletin*, 130(6):843, 2004.
- Sein Minn, Jill-Jënn Vie, Koh Takeuchi, Hisashi Kashima, and Feida Zhu. Interpretable knowledge tracing: Simple and efficient student modeling with causal relations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 12810–12818, 2022.
- Koki Nagatani, Qian Zhang, Masahiro Sato, Yan-Ying Chen, Francine Chen, and Tomoko Ohkuma. Augmenting knowledge tracing by considering forgetting behavior. In *The world wide web conference*, pp. 3101–3107, 2019.
- Hiromi Nakagawa, Yusuke Iwasawa, and Yutaka Matsuo. Graph-based knowledge tracing: modeling student proficiency using graph neural network. In *IEEE/WIC/ACM International Conference on Web Intelligence*, pp. 156–163, 2019.
- Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017.
- Shalini Pandey and George Karypis. A self-attentive model for knowledge tracing. *arXiv preprint arXiv:1907.06837*, 2019.

- Philip I Pavlik, Luke G Eglinton, and Leigh M Harrell-Williams. Logistic knowledge tracing: A constrained framework for learner modeling. *IEEE Transactions on Learning Technologies*, 14(5): 624–639, 2021.
- Philip I Pavlik Jr, Hao Cen, and Kenneth R Koedinger. Performance factors analysis—a new alternative to knowledge tracing. *Online Submission*, 2009.
- Jean Piaget. Science of education and the psychology of the child. trans. d. coltman. 1970.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. Deep knowledge tracing. *Advances in neural information processing systems*, 28, 2015.
- Sidney L Pressey. A simple apparatus which gives tests and scores-and teaches. *Sch. & Soc.*, 23: 373–376, 1926.
- David E Rumelhart. Schemata: The building blocks of cognition. In *Theoretical issues in reading comprehension*, pp. 33–58. Routledge, 2017.
- Michael Sao Pedro, Ryan Baker, and Janice Gobert. Incorporating scaffolding and tutor context into bayesian knowledge tracing to predict inquiry skill acquisition. In *Educational Data Mining 2013*. Citeseer, 2013.
- Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.
- Douglas Selent, Thanaporn Patikorn, and Neil Heffernan. Assisments dataset from multiple randomized controlled experiments. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*, pp. 181–184, 2016.
- Burr Settles and Brendan Meeder. A trainable spaced repetition model for language learning. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 1848–1858, 2016.
- Shuanghong Shen, Qi Liu, Enhong Chen, Zhenya Huang, Wei Huang, Yu Yin, Yu Su, and Shijin Wang. Learning process-consistent knowledge tracing. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 1452–1460, 2021.
- Yuge Shi, Brooks Paige, Philip Torr, et al. Variational mixture-of-experts autoencoders for multi-modal deep generative models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Dongmin Shin, Yugeun Shim, Hangeol Yu, Seewoo Lee, Byungsoo Kim, and Youngduck Choi. Saint+: Integrating temporal features for ednet correctness prediction. In *LAK21: 11th International Learning Analytics and Knowledge Conference*, pp. 490–496, 2021.
- Shiwei Tong, Qi Liu, Wei Huang, Zhenya Hunag, Enhong Chen, Chuanren Liu, Haiping Ma, and Shijin Wang. Structure-based knowledge tracing: An influence propagation view. In *2020 IEEE international conference on data mining (ICDM)*, pp. 541–550. IEEE, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Matthew M Walsh, Kevin A Gluck, Glenn Gunzelmann, Tiffany Jastrzembski, and Michael Krusmark. Evaluating the theoretic adequacy and applied potential of computational models of the spacing effect. *Cognitive science*, 42:644–691, 2018.
- Chenyang Wang, Weizhi Ma, Min Zhang, Chuancheng Lv, Fengyuan Wan, Huijie Lin, Taoran Tang, Yiqun Liu, and Shaoping Ma. Temporal cross-effects in knowledge tracing. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pp. 517–525, 2021.
- Zichao Wang, Angus Lamb, Evgeny Saveliev, Pashmina Cameron, Yordan Zaykov, José Miguel Hernández-Lobato, Richard E Turner, Richard G Baraniuk, Craig Barton, Simon Peyton Jones, Simon Woodhead, and Cheng Zhang. Diagnostic questions: The neurips 2020 education challenge. *arXiv preprint arXiv:2007.12061*, 2020.

John T Wixted and Ebbe B Ebbesen. Genuine power curves in forgetting: A quantitative analysis of individual subject forgetting functions. *Memory & cognition*, 25:731–739, 1997.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.

Michael V Yudelson, Kenneth R Koedinger, and Geoffrey J Gordon. Individualized bayesian knowledge tracing models. In *Artificial Intelligence in Education: 16th International Conference, AIED 2013, Memphis, TN, USA, July 9-13, 2013. Proceedings 16*, pp. 171–180. Springer, 2013.

A APPENDIX

The Appendix is organized as follows:

- Appendix A.1 and A.2 provides a detailed derivation of our objective function ELBO in two scenarios: inference involving complete learning histories (Sec. 3.2.1) and inference for in a continual learning setting (Sec. 3.2.2).
- Appendix A.3 provides in-depth descriptions of baseline models, and the details and the selection criterion of the three datasets we use for experiments.
- Appendix A.4 describes the PSI-KT architecture in full detail, including its hyperparameters.
- Appendix A.5 provides additional prediction results. We show the average accuracy, average f1-score, and their standard error of Fig. 2 in the prediction experiment given entire learning histories. We also show the average accuracy score and their standard error of Fig. 3 in the prediction experiment for continual learning setup.
- Appendix A.6 describes the details of the experiment setup and how we derive the metrics for specificity, consistency, and disentanglement. We also provide comprehensive results on the operational interpretability of baseline models.
- Appendix A.7 elaborates on the graph assessment framework, including details about the alignment metrics and a discussion of causal support, and extends the main text evaluations to all datasets.

A.1 ELBO OF THE HIERARCHICAL SSM

In this section, we derive the single-learner ELBO, Eqs. 8- 10 in the main text. For clarity, we omit the superindex ℓ in these derivations. Note that the parameters ϕ and θ are global, i.e., they are optimized based on the entire interaction data across learners.

In variational inference (VI), we approximate an intractable posterior distribution $p_\theta(z | y)$ with $q_\phi(z | y)$ from a tractable distribution family. We learn ϕ and θ together by maximizing the evidence lower bound (ELBO) of the marginal likelihood (Blei et al., 2017; Attias, 1999), given by $\log p_\theta(y) \geq \text{ELBO}(\theta, \phi) = \mathbb{E}_{q_\phi(z | y)} [-\log q_\phi(z | y) + \log p_\theta(y, z)]$. The two terms in the ELBO represent the entropy of the variational posterior distribution, $\text{H}[q_\phi(z | y)] = \mathbb{E}_{q_\phi(z | y)} [-\log q_\phi(z | y)]$, and the log-likelihood of the joint distribution of observations and latent states $\mathbb{E}_{q_\phi(z | y)} \log p_\theta(y, z)$.

We now formulate the ELBO for our hierarchical SSM (see Fig. 1) with two layers of latent states. We assume that fixed learning histories $\mathcal{H}_{1:N}$ until time t_N are available and we use capital N to represent the fixed time point. We approximate the posterior $p_\theta(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N})$ using the mean-field factorization, $q_\phi(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N}) = q_\phi(\mathbf{z}_{1:N}) q_\phi(\mathbf{s}_{1:N})$:

$$\begin{aligned} \text{ELBO}(\theta, \phi) &= \text{H}[q_\phi(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N})] + \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N})} \log p_\theta(\mathbf{y}_{1:N}, \mathbf{z}_{1:N}, \mathbf{s}_{1:N}) \\ &= \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N})} [-\log q_\phi(\mathbf{z}_{1:N}, \mathbf{s}_{1:N} | \mathbf{y}_{1:N}) + \log p_\theta(\mathbf{y}_{1:N}, \mathbf{z}_{1:N}, \mathbf{s}_{1:N})] \\ &= \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}) q_\phi(\mathbf{s}_{1:N})} [-\log q_\phi(\mathbf{z}_{1:N}) - \log q_\phi(\mathbf{s}_{1:N}) + \log p_\theta(\mathbf{y}_{1:N}, \mathbf{z}_{1:N}, \mathbf{s}_{1:N})]. \end{aligned} \quad (12)$$

In the generative model of PSI-KT, the observation y_n at time t_n depends on the particular knowledge state z_n^k associated with the interacted KC $k = x_n$. All knowledge states z_n rely on previous states z_{n-1} and cognitive traits s_n , which themselves are influenced by s_{n-1} . Thus, we can factorize the joint distribution $p_\theta(\mathbf{y}_{1:N}, \mathbf{z}_{1:N}, \mathbf{s}_{1:N})$ in Eq. 12 over all latent states and observations:

$$\begin{aligned} p_\theta(\mathbf{y}_{1:N}, \mathbf{z}_{1:N}, \mathbf{s}_{1:N}) &= p_\theta(\mathbf{s}_{1:N}) p_\theta(\mathbf{z}_{1:N} | \mathbf{s}_{1:N}) \prod_{n=1}^N p_\theta(y_n | z_n^{x_n}) \\ &= p_\theta(\mathbf{s}_1) p_\theta(\mathbf{z}_1) \prod_{n=2}^N p_\theta(s_n | s_{n-1}) p_\theta(\mathbf{z}_n | \mathbf{z}_{n-1}, s_n) \prod_{n=1}^N p_\theta(y_n | z_n^{x_n}). \end{aligned} \quad (13)$$

Here, $p_\theta(\mathbf{s}_1)$ and $p_\theta(\mathbf{z}_1)$ are the Gaussian initial priors for the latent states.

By incorporating the factorized joint distribution (Eq. 13), the ELBO for PSI-KT can be derived as follows:

$$\begin{aligned}
\text{ELBO}(\theta, \phi) &= \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}) q_\phi(s_{1:N})} \left[-\log q_\phi(\mathbf{z}_{1:N}) - \log q_\phi(s_{1:N}) + \log p_\theta(y_{1:N}, \mathbf{z}_{1:N}, s_{1:N}) \right] \\
&= \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}) q_\phi(s_{1:N})} \left[-\log q_\phi(\mathbf{z}_{1:N}) - \log q_\phi(s_{1:N}) + \log p_\theta(s_1) + \log p_\theta(\mathbf{z}_1) \right. \\
&\quad \left. + \sum_{n=2}^N \log p_\theta(s_n | s_{n-1}) + \sum_{n=2}^N \log p_\theta(\mathbf{z}_n | \mathbf{z}_{n-1}, s_n) \right. \\
&\quad \left. + \sum_{n=1}^N \log p_\theta(y_n | z_n^{x_n}) \right] \\
&= \mathbb{E}_{q_\phi(s_{1:N})} \left[-\log q_\phi(s_{1:N}) + \log p_\theta(s_1) + \sum_{n=2}^N \log p_\theta(s_n | s_{n-1}) \right] \\
&\quad + \mathbb{E}_{q_\phi(\mathbf{z}_{1:N})} \left[-\log q_\phi(\mathbf{z}_{1:N}) + \log p_\theta(\mathbf{z}_1) + \sum_{n=1}^N \log p_\theta(y_n | z_n^{x_n}) \right] \\
&\quad + \mathbb{E}_{q_\phi(\mathbf{z}_{1:N}) q_\phi(s_{1:N})} \left[\sum_{n=2}^N \log p_\theta(\mathbf{z}_n | \mathbf{z}_{n-1}, s_n) \right]. \tag{14}
\end{aligned}$$

A.2 EXTENSION TO CONTINUAL LEARNING

We now extend the ELBO to the continual learning setting, where we observe learning performances $y_{1:n}$ sequentially. Here we use the lower case n to indicate the running time index. We seek the posterior distribution $p_\theta(\mathbf{z}_n, s_n | y_{1:n})$ at each interaction time t_n given all observations so far. Usually, one would approximate the posterior with the variational posterior distribution $q_{\phi_n}(\mathbf{z}_n, s_n | y_{1:n}) = q_{\phi_n}(\mathbf{z}_n) q_{\phi_n}(s_n)$ using the mean-field factorization (Eq. 12). In that case, the inference process consists of maximizing the $\text{ELBO}(\theta, \phi_n)$ only over ϕ_n :

$$\text{ELBO}(\theta, \phi_n) = \mathbb{E}_{q_{\phi_n}(\mathbf{z}_n) q_{\phi_n}(s_n)} \left[-\log q_{\phi_n}(\mathbf{z}_n) - \log q_{\phi_n}(s_n) + \log p_\theta(y_{1:n}, \mathbf{z}_n, s_n) \right]. \tag{15}$$

However, it is challenging to calculate the joint distribution $p_\theta(y_{1:n}, \mathbf{z}_n, s_n)$ in our setup, since it requires marginalizing the full joint distribution $p_\theta(y_{1:n}, \mathbf{z}_{1:n}, s_{1:n})$ over all $\mathbf{z}_{n'}$ and $s_{n'}$ with $n' < n$.

Henceforth, we aim to reconfigure the objective function $\text{ELBO}(\theta, \phi_n)$, which involves the variational posterior distribution $q_{\phi_n}(\mathbf{z}_n, s_n | y_{1:n})$ at the time point t_n , to establish a linkage with the posterior $q_{\phi_{n-1}}(\mathbf{z}_{n-1}, s_{n-1} | y_{1:n-1})$ observed at the preceding time point t_{n-1} . By doing so, we can recursively optimize the variational parameters ϕ_n whenever new observations y_n are received (Nguyen et al., 2017), wherein the initialization draws from the parameters ϕ_{n-1} obtained at time t_{n-1} .

First, we show that the marginal joint distribution $p_\theta(y_{1:n}, \mathbf{z}_n, s_n)$ is proportional to the prior distribution $p_\theta(\mathbf{z}_n, s_n | y_{1:n-1})$ on s_n, \mathbf{z}_n at t_{n-1} :

$$p_\theta(y_{1:n}, \mathbf{z}_n, s_n) \propto p_\theta(\mathbf{z}_n, s_n | y_{1:n-1}) p_\theta(y_n | \mathbf{z}_n). \tag{16}$$

Second, we show how the prior distribution $p_\theta(\mathbf{z}_n, s_n | y_{1:n-1})$ can be formulated using the posterior $q_{\phi_{n-1}}(\mathbf{z}_{n-1}, s_{n-1} | y_{1:n-1})$ at the previous time t_{n-1} , which evolves for a single time step:

$$\begin{aligned}
p_\theta(\mathbf{z}_n, s_n | y_{1:n-1}) &= \int p_\theta(\mathbf{z}_n, \mathbf{z}_{n-1}, s_n, s_{n-1} | y_{1:n-1}) ds_{n-1} d\mathbf{z}_{n-1} \\
&= \int p_\theta(\mathbf{z}_{n-1}, s_{n-1} | y_{1:n-1}) p_\theta(s_n | s_{n-1}) p_\theta(\mathbf{z}_n | s_n, \mathbf{z}_{n-1}) ds_{n-1} d\mathbf{z}_{n-1} \\
&= \int q_{\phi_{n-1}}(\mathbf{z}_{n-1}, s_{n-1} | y_{1:n-1}) p_\theta(s_n | s_{n-1}) p_\theta(\mathbf{z}_n | s_n, \mathbf{z}_{n-1}) ds_{n-1} d\mathbf{z}_{n-1} \\
&= \int q_{\phi_{n-1}}(\mathbf{z}_{n-1}) q_{\phi_{n-1}}(s_{n-1}) p_\theta(s_n | s_{n-1}) p_\theta(\mathbf{z}_n | s_n, \mathbf{z}_{n-1}) ds_{n-1} d\mathbf{z}_{n-1} \\
&= \underbrace{\mathbb{E}_{q_{\phi_{n-1}}(s_{n-1})} [p_\theta(s_n | s_{n-1})]}_{:=\tilde{p}_{\phi_{n-1}}(s_n)} \underbrace{\mathbb{E}_{q_{\phi_{n-1}}(\mathbf{z}_{n-1})} [p_\theta(\mathbf{z}_n | s_n, \mathbf{z}_{n-1})]}_{:=\tilde{p}_{\phi_{n-1}}(\mathbf{z}_n | s_n)}. \tag{17}
\end{aligned}$$

Substituting $p_\theta(s_n, \mathbf{z}_n | y_{1:n-1}) = \tilde{p}_{\phi_{n-1}}(s_n) \tilde{p}_{\phi_{n-1}}(\mathbf{z}_n | s_n)$ using our variational approximation in Eq. 16 in the ELBO, we finally arrive at the objective function for variational continuous learning:

$$\begin{aligned}
\text{ELBO}_{\text{VCL}}(\theta, \phi_n) &= \mathbb{E}_{q_{\phi_n}(s_n) q_{\phi_n}(\mathbf{z}_n)} [-\log q_{\phi_n}(s_n) - \log q_{\phi_n}(\mathbf{z}_n) + \log \tilde{p}_{\phi_{n-1}}(s_n) \\
&\quad + \log \tilde{p}_{\phi_{n-1}}(\mathbf{z}_n | s_n) + \log p_\theta(y_n | \mathbf{z}_n^{x_n})] \\
&= \mathbb{E}_{q_{\phi_n}(s_n)} [-\log q_{\phi_n}(s_n) + \log \tilde{p}_{\phi_{n-1}}(s_n)] \\
&\quad + \mathbb{E}_{q_{\phi_n}(\mathbf{z}_n)} [-\log q_{\phi_n}(\mathbf{z}_n) + \log p_\theta(y_n | \mathbf{z}_n^{x_n})] \\
&\quad + \mathbb{E}_{q_{\phi_n}(\mathbf{z}_n) q_{\phi_n}(s_n)} [\log \tilde{p}_{\phi_{n-1}}(\mathbf{z}_n | s_n)]. \tag{18}
\end{aligned}$$

This provides a derivation of Eq. 10 as presented in the main text. Here our focus lies in the optimization of the parameters ϕ_n , while holding constant the parameters ϕ_{n-1} acquired from the preceding time step.

A.3 BASELINE MODELS AND DATASETS

A.3.1 BASELINES

KT models aim to predict the performance \hat{y}_n^ℓ of the presented KC x_n^ℓ at time t_n^ℓ for each learner ℓ , which amounts to learning the mapping $\hat{y}_n^\ell = f_\theta(\mathcal{H}_{n', < n}^\ell)$ (Sec. 2.1). Because baseline models lack learner-specific parameters, we here describe the prediction process for a single learner, and omit the superindex ℓ for clarity. Extending to multiple learners is straightforward since all parameters θ are global. We use $\tau_n := t_n - t_{n-1}$ to represent the time interval between consecutive interactions of a learner ℓ , and the KC-specific interval $\tau_n^k := t_n^k - t_{n-1}^k$ for consecutive interactions with the same KC k . The number of practice repetitions for each KC k up to time t_n is denoted as c_n^k . The dimension of embeddings D equals 16 in our experiments.

Table 6: Models. # Emb/KC is the number of learnable embeddings per KC. *Forgetting* is the functional form of memory decay, with exponential (exp) decay the most common.

Feature	HLR/PPE	DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
# Emb/KC	–	2	2	6	6	3	1	1
Forgetting	exp	–	exp	Hawkes	–	–	–	OU

We compare with eight baseline models (Sec. 4):

- a) HLR (Settles & Meeder, 2016) uses the cumulative counts of correct, incorrect, and total interactions of KC k until time t_n , collectively denoted $\mathbf{c}_n^k = [c_n^{k,1} c_n^{k,0} c_n^k]^\top \in \mathbb{R}^3$, as well as the last interval τ_n^k . When a learner interacts with KC x_n at time t_n , HLR predicts the probability of a correct performance as:

$$\hat{y}_n := 2^{-\tau_n^k / h_n^k}, \quad \text{with memory half-life } h_n^k := 2^{\theta^\top \mathbf{c}_n^k} \text{ and } k = x_n. \tag{19}$$

The learnable weights $\theta \in \mathbb{R}^3$ modulate the influences of correct, incorrect, and total interaction counts. The training process of HLR does not differentiate features from different

KCs or learners, thus HLR cannot model the relational structure of KCs or any learner-specific characteristics.

- b) PPE (Walsh et al., 2018) is similar to HLR in predicting performance as a function of interaction histories. It defines the activation m_n^k of KC k at time t_n , $m_n^k := (c_n^k)^\beta (T_n^k)^{-\alpha}$ with separate forgetting rate α and learning rate β . The forgetting term T_n^k is a function of the interaction history, which it summarizes as a weighted average of times τ_n^k elapsed between the exposures to a given KC prior to t_n :

$$T_n^k := \sum_{i=1}^{n-1} w_i^k \tau_i^k, \quad \text{with } w_i^k = (\tau_i^k)^{-\eta} \sum_{j=1}^{n-1} (\tau_j^k)^\eta. \quad (20)$$

The forgetting rate α is a function of a *stability term* κ and a cumulative average of interval durations between KC exposures modulated by the slope λ :

$$\alpha_n^k = \kappa + \lambda \left(\frac{1}{n-1} \sum_{j=1}^{n-1} \frac{1}{\ln(\tau_j^k + e)} \right). \quad (21)$$

Finally, PPE treats performance \hat{y}_n as a logistic function of m_n^k with $k = x_n$. The learnable parameters are $\theta = \{\beta, \eta, \kappa, \lambda\}$.

- c) DKT (Piech et al., 2015) infers two separate embeddings $u^k = \{u^{k,0}, u^{k,1}\}$ for each KC k , depending on performance. Here, $u^{k,0}, u^{k,1} \in \mathbb{R}^D$ represent incorrect interactions and correct interactions on KC k , respectively, and are shared across all learners, with D being the dimensionality of the embeddings. DKT trains an LSTM (Hochreiter & Schmidhuber, 1997) over $\mathcal{H}_{n' < n}$ to encode the combined information of KC indices and performance. For each learner ℓ and all time points $t_{n' < n}$, DKT takes performance embeddings $u^{x_{n'},0}$ of interacted KC $x_{n'}$ as the input when the performance $y_{n'}$ is incorrect, or $u^{x_{n'},1}$ for correct performance, i.e., DKT takes inputs $\mathbf{u}_{n'} = u^{x_{n'}, y_{n'}}$ for all $t_{n'}$ with $n' < n$. DKT then predicts the subsequent performance on all KCs $\hat{\mathbf{y}}_n = [y_n^1, \dots, y_n^K]^\top \in \mathbb{R}^K$, and chooses only the interacted one, i.e., the x_n -th dimension:

$$\begin{aligned} \mathbf{h}_n &= \text{LSTM}(\mathbf{u}_{n' < n}; \mathbf{W}_h, \mathbf{b}_h) \\ \hat{\mathbf{y}}_n &= \sigma(\mathbf{W}_{\hat{y}h} \mathbf{h}_n + \mathbf{b}_{\hat{y}}) \\ \hat{y}_n &= \hat{\mathbf{y}}_n[x_n]. \end{aligned} \quad (22)$$

Thus, θ for DKT consists of the neural network parameters $\mathbf{W}_h, \mathbf{b}_h, \mathbf{W}_{\hat{y}h}, \mathbf{b}_{\hat{y}}$.

- d) DKTF (Nagatani et al., 2019) uses the same LSTM architecture and the same combined KC-performance embeddings, $u^k = \{u^{k,0}, u^{k,1}\}$ that we described above for DKT. DKTF uses additional 3-dimensional features $\mathbf{t}_n := [\tau_n, \tau_n^k, c_n^k]$ representing the KC-unspecific and KC-specific intervals defined above and the cumulative interaction counts c_n^k for KC k until time t_n . Then, for inputs of every time point $t_{n'}$, DKTF concatenates the time information $\mathbf{t}_{n'}$ with KC performance inputs $\mathbf{u}_{n'}$ for the interacted KC $x_{n'}$. DKTF predicts future performances following the same architecture based on concatenated input $[\mathbf{t}_{n' < n}; \mathbf{u}_{n' < n}]$.
- e) HKT (Wang et al., 2021) is the most similar model to our PSI-KT. It uses a Hawkes process to model the structural influence on the state of KC k due to every other KCs state in the past interactions $i \in x_{n' < n}$ until time t_n :

$$\begin{aligned} m_n^k &= \lambda_k + \sum_{i \in x_{n' < n}} a_n^{i,k} \kappa(t_n^k - t_{n'}^i) \\ \kappa(t_n^k - t_{n'}^i) &= \exp\left(-\left(1 + \beta_n^{i,k} \log(t_n^k - t_{n'}^i)\right)\right). \end{aligned} \quad (23)$$

Here m_n^k includes a base level λ_k and all previous learned KCs' influences $a_n^{i,k}$ weighted by the a temporal exponential decay $\kappa(t_n^k - t_{n'}^i)$. The base level λ_k reflects aspects of KC k but also of the specific assignments that were interacted with at time point t_n , given that distinct assignments can provide practice for a single KC. To model cross-KC influences $a_n^{i,k}$, HKT infers embeddings $\{u_a^{k,0}, u_a^{k,1}, u_a^k\} \in \mathbb{R}^D$ for each KC k . Here $u_a^{k,0}, u_a^{k,1}$ are defined similarly to the DKT embeddings, whereas u_a^k only depends on KC identity k . When interacting with KC k at time t_n , the influence on its state due to having interacted with KC i at time $t_{n'}$ with performance $y_{n'}$ is estimated as $a_n^{i,k} = (u_a^{i, y_{n'}})^\top u_a^k$. For the coefficient $\beta_n^{i,k}$, HKT estimates three additional KC-specific embeddings $\{u_\beta^{k,0}, u_\beta^{k,1}, u_\beta^k\}$, and follows similar

calculations as for a above. HKT also predicts the performance \hat{y}_n as a logistic function of m_n^k with $k = x_n$.

- f) AKT (Ghosh et al., 2020) is a transformer-based model (Vaswani et al., 2017) that learns the structure of KCs implicitly from the self-attention weights. Unlike LSTM models, which only capture temporal information, AKT captures both temporal and structural relations. Specifically, AKT first initializes three embeddings $\{u_a^{k,0}, u_a^{k,1}, u_a^k\}$ for each KC k and a scalar μ^q for each specific assignment q , representing its difficulty level. For each KC k , these embeddings are defined as in HKT to separately reflect KC-specific correct/incorrect interactions and KC identity. However, AKT combines these representations with three additional embeddings $\{u_b^{k,0}, u_b^{k,1}, u_b^k\}$, in order to account for difficulty levels. When a learner interacts at time $t_{n'}$ with an assignment q related to KC k , the KC identity embedding becomes $u^k = u_a^k + \mu^q u_b^k$; after assessing the performance at time $t_{n'}$, the interaction embeddings are similarly updated as $u^{k,y_{n'}} = u_b^{k,y_{n'}} + \mu^q u_b^{k,y_{n'}}$. Consequently, a learner’s entire interaction history $\mathcal{H}_{n' < n}$ is represented as a sequence of these combined KC-interaction-difficulty embeddings. AKT processes these sequential embeddings as input, using KC embeddings as queries and keys, and interaction embeddings as values within its attention mechanism. To predict performance \hat{y}_n given the KC and assignment, AKT uses the KC embeddings at time t_n to compare with previous queries and keys in the learning history, and then extract the value. Details about the transformer architecture can be found in Ghosh et al. (2020).
- g) GKT (Nakagawa et al., 2019) applies a graph neural network to leverage the graph-structured nature of knowledge. Like AKT, GKT initializes three embeddings $\{u^{k,0}, u^{k,1}, u^k\}$ for each KC k , but instead of only using the embeddings to determine the KC relations, GKT learns an additional undirected KC graph, represented by its adjacency matrix \mathbf{A} . Here $a_{ij} = 1$ represents KC i and KC j are related, i.e., there is information transmission among KC i and KC j every time the model gets updated. To use the KC relations, GKT first aggregates the hidden states \mathbf{h}_n^k and embeddings for the KC reviewed at time t_n , k and its neighboring KCs i :

$$(\mathbf{h}_n^k)' = \begin{cases} [\mathbf{h}_n^k, u^{x_n, y_n}] & (i = k) \\ [\mathbf{h}_n^k, u^i] & (i \neq k \text{ with } a_{ik} = 1) \end{cases}$$

After aggregating the information from the neighboring KCs, GKT updates the hidden states based on the aggregated features and the graph structure:

$$\mathbf{h}_{n+1}^k = \begin{cases} f_\theta(\mathbf{h}_n^{k'}, \mathbf{h}_n^k) & (i = k) \\ f_\theta((\mathbf{h}_n^k)', (\mathbf{h}_n^i)', \mathbf{h}_n^k) & (i \neq k \text{ with } a_{ik} = 1). \end{cases}$$

Finally, GKT uses an MLP layer to predict the probability of a correct answer at the next time step, $\hat{y}_{n+1}^k = f_\theta(\mathbf{h}_{n+1}^k)$ where $k = x_{n+1}$.

- h) QIKT (Chen et al., 2023) focuses on assignments together with KCs, where multiple different assignments can test one KC. Inspired by item-response theory (IRT) (Lord, 2012), QIKT defines three modules, each parameterized by a neural network, to infer interpretable features, namely assignment-specific knowledge acquisition α_n , assignment-specific problem-solving ability ζ_{n+1} , and assignment-agnostic but KC-specific knowledge mastery β_n . Apart from neural network parameters, QIKT learns three sets of assignment-specific embeddings $\{v^{q,0}, v^{q,1}, v^q\}$, which have the same purpose as the KC embeddings defined in AKT, namely for correct interactions, incorrect interactions, and assignment identity. Furthermore, another set of KC-specific embeddings u^k is learned for KC-specific features. QIKT uses LSTM and sum pooling to learn the three features based on each learner’s history, $\alpha_n := f_\theta(V_{1:n}, U_{1:n}), \beta_n := f_\theta(U_{1:n})$, where $V_{1:n}$ and $U_{1:n}$ denote respectively the all the assignments and KC embeddings in the learning history. The problem-solving ability $\zeta_{n+1} := f_\theta(V_{1:n+1}, U_{1:n+1})$ is learned by including the assignment and KC information in the coming interaction. To predict performance, all three features are aggregated and input into the sigmoid function, $\hat{y}_{n+1} = \sigma(\alpha_n + \beta_n + \zeta_{n+1})$.

A.3.2 DATASETS

Here we describe the datasets that we have used for evaluation (Assist12, Assist17, and Junyi15), articulate the reasons for their selection, and discuss some of the limitations derived from this choice.

	Assist12		Assist17		Junyi15	
	All	> 50	All	> 50	All	> 50
# Interactions	6,123,270	2,431,788	942,816	942,489	25,925,992	23,907,121
# Learners	46,674	12,443	1,709	1,697	247,606	77,655
# Assignments	179,999	51,866	3,162	3,162	5,174	6,174
# KCs	265	263	102	102	722	721
KC Examples	Rounding Unit Rate Perimeter of a Polygon		substitution fraction-division prime-number		matrix_basic_distance circles_and_arcs arithmetic_means	

Table 7: Overview of Educational Datasets Assist12, Assist17, and Junyi15, including the number of interactions, learners, assignments, and KCs from overall log data (All) and the log data including learners with more than 50 interactions (> 50).

Description of the selected datasets Assist12 and Assist17 are two subsets of the ASSISTments dataset released by Worcester Polytechnic Institute (Selent et al., 2016). ASSISTments is an online educational tool widely used in U.S. mathematics classes for learners from grades 4 to 12. Predominantly, its users are middle school students (grades 6-8) from Massachusetts or its vicinity. This platform is used for both classroom and homework assignments, and can be used with or without accompanying paper materials. One of the key features of ASSISTments is the immediate feedback provided to students after they answer a question, allowing them to promptly know whether their response was correct.

Junyi Academy is a non-profit Chinese online education platform. Their Junyi15 data release reports the interactions of more than 72,000 students solving mathematics assignments over a year, totaling 16 million attempts. These interaction logs are provided along with two commissioned annotation sets (‘expert’ and ‘crowd-sourced’) concerning the structure of 837 KCs in the curriculum. Expert annotations, provided by three teachers, consist of 553 identified prerequisite relations. Crowd-sourced annotations, provided by 51 graduates from senior high school or higher, consist of both prerequisite and similarity evaluations for 1954 KC pairs, with each relation strength rated on a scale from 0 to 9 by at least 3 workers.

We report the numbers of learners, KCs, assignments, and interactions for each dataset in Table 7. In Figure 6 we complement this basic characterization with histograms of the number of per-learner interactions, KCs, and assignments, as well as histograms of elapsed time between learner interactions with arbitrary KCs as well as between interactions with the same KC.

Criteria for dataset selection In order to empirically test our model of learning in structured domains, we sought datasets from domains with a clear prerequisite structure that provide (1) identifiable KC labels, and (2) interaction times with sufficient temporal resolution. In domains where prerequisite relations between KCs are strong, the correct learning order is key for performance, so that performance data be used to uncover structural relations. Additionally, the dependencies in these domains can be identified independently by human annotators, which we use to validate model inferences about the knowledge structure.

1. Identifiable KC labels. Some datasets do not identify the specific KC reviewed at an interaction, but rather a more general assignment or task that could involve multiple unspecified KCs. While this assignment structure can be explicitly modeled (e.g. our baselines AKT, HKT, and QIKT), and we do intend to extend our model in future work to cover this setting, here we intentionally avoided modeling assignment features and concentrated directly on the underlying KCs and their dependencies, which requires KC identities.
2. Timestamped interactions with high temporal resolution. A resolution in the order of seconds or less is essential to adequately track the initial phases of the forgetting process, and to model structural influences that depend on the precise order of KC presentation (see Eq. 5).

Following these criteria, we had to exclude the Statics2011 dataset due to a lack of identified KCs. The Assistments2009 and Assistments2015 datasets lack timestamps entirely, while the 15-minute temporal resolution of the Junyi20 dataset is too coarse for our purposes. This leaves us with Assist12, Assist17, and Junyi15 as appropriate choices to evaluate KT on structured domains. Besides abundant interaction data, Junyi15 provides human-annotated KC relations that, while noisy, offer an invaluable reference to compare the inferred prerequisite graphs.

Limitations The selection of datasets is limited by design to structured domains, where we can more appropriately put to the test our structure-aware model. We acknowledge that when KCs are largely unrelated (e.g., general knowledge trivia) the inference of prerequisite structure may confer no real advantage. Mathematics, in contrast, provides an ideal testing ground, but more interaction datasets from other domains (e.g., biology, chemistry, linguistics...) and learning stages (primary school, college) are needed for a more representative assessment of the role of structure in learning. In the future, we intend to extend our model to accommodate a broader range of datasets, addressing, in particular, the common case where a single interaction, such as an assignment or a task, is associated with multiple KCs, which entails a more complex interplay of KCs than is displayed in our current dataset selection (Wang et al., 2020).

A.4 PSI-KT MODEL ARCHITECTURE

A.4.1 NETWORK DETAILS

In this section, we introduce the detailed architecture of our PSI-KT model and its hyperparameters. The inference network consists of an embedding network f_ϕ^{Emb} , the cognitive traits encoder f_ϕ^s , and the knowledge states encoder f_ϕ^z . The weights of these interconnected networks collectively constitute the inference parameters ϕ .

Interaction embedding network. The network f_ϕ^{Emb} extracts features from the learning history tuples $\mathcal{H}_{1:N}^\ell = \{x_n, y_n, t_n\}_{1:N}^\ell$, combining information about interaction time, KC identity and performance.

The *KC identity embedding* for KC x_n corresponds to the learned embedding u^{x_n} , which is part of the generative model that parameterizes the graph structure. The *performance embedding* is obtained by expanding the scalar value y_n into a vector \vec{y}_n with the same dimensionality as the time and KC embeddings so that the performance features will be represented on an equal footing. We then concatenate the KC embedding u^{x_n} with the performance embedding \vec{y}_n . The *interval embedding* is a positional encoding (Vaswani et al., 2017), $\text{PE}_n = (\sin \alpha(\tau_n); \cos \alpha(\tau_n))$. This embedding approach accommodates intervals spanning different timescales, from minutes to weeks.

Thus, the joint embedding for a learning interaction is given by $v_n = f_{\phi, \text{Emb}}([u^{x_n}; \vec{y}_n]) + \text{PE}_n$, inspired by the transformer architecture (Vaswani et al., 2017).

Latent state encoder. The network f_ϕ^z infers the parameters of the variational posterior distribution $q_\phi(z_{1:n})$. Since learning histories do not have a pre-determined length, we use an LSTM (Hochreiter & Schmidhuber, 1997) as the inference architecture. At each time point, we extract the hidden states in the LSTM, $h_{z_{1:n}} = \text{LSTM}(v_{1:n})$. Meanwhile, in the continual learning setting, information about the history is already encoded and available in the variational parameters for the last time step ϕ_{n-1} , so we use a multi-layer perceptron (MLP), $h_{z_n} = \text{MLP}(v_n)$. Finally, another MLP (similar to the encoder in Kingma & Welling, 2014) takes the hidden states h_{z_n} at every time point as inputs and produces the mean $\mu_{z_n} \in \mathbb{R}^K$ and log-variance $\log \sigma_{z_n} \in \mathbb{R}^K$ for knowledge states z_n .

Latent trait encoder. The network f_ϕ^s infers the parameters of the variational posterior distribution $q_\phi(s_{1:n})$. The resulting approximate posterior distribution enables the sampling of learner-specific traits to facilitate personalized predictions. One immediately obvious approach is to use the same architecture of f_ϕ^z . However, the unimodal Gaussian prior over the latent variables cannot account for the diversity of cognitive trait combinations that we expect to find across learners in diverse cohorts. What we need is to allow for multimodality in the distribution of s over all learners.

There is work on factorizing the joint variational posterior as a combination of isotropic posteriors, using a mixture of M experts (MoE; Shi et al., 2019), i.e., $q_\phi(s_{1:n}^\ell | \mathcal{H}_{1:n}^\ell) = 1/M \sum_m q_{\phi_m}(s_{1:n}^\ell | \mathcal{H}_{1:n}^\ell)$, assuming the different modalities are of comparable complexity. However, this may lead to overparameterization. Instead, inspired by Dilokthanakul et al. (2016), we opt for a mixture of Gaussians as a prior distribution that generalizes the unimodal Gaussian prior and provides multimodality. By assuming that the observed data arises from a mixture of Gaussians, determining the category of a data point becomes equivalent to identifying the mode of the latent distribution from which the data point originates. This approach allows us to partition our latent space into distinct categories. With these discrete variables, it is no longer possible to directly apply the reparameterization trick.

Table 8: PSI-KT architecture and hyperparameters. $FC(a, b)$ represents a fully connected layer with input dimension a and output dimension b ; K represents the number of KCs, different across datasets; C represents the number of categories in the mixture of Gaussians for s (we use $C = 10$ in our experiments); the semicolon ; separates connected layers, while the slash / separates the layer architecture for inference on entire histories from the continual learning set-up, where different.

	Inputs & Dim	Hidden Layers	Outputs
f_{ϕ}^{Emb}	KC Emb & 16 Perf Emb & 16	FC (32, 16) LeakyReLU(0.2) FC (16, 16)	v_n
f_{ϕ}^z	v_n & 16	LSTM (16, 32) / FC (16, 32) FC (32, 16); LeakyReLU(0.2) FC (16, 16); LeakyReLU(0.2) FC (16, K); FC (16, K)	$\mu_{z_n}, \log \sigma_{z_n}$
f_{ϕ}^s	v_n & 16	FC (16, 32) FC (32, 16); LeakyReLU(0.2) FC(16, 64); LeakyReLU(0.2) GumbelSoftmax(FC(64, C)) FC(32 + C , 64); LeakyReLU(0.2) FC(64, 16); LeakyReLU(0.2) FC(16); FC(64, 4)	$\mu_{s_n}, \log \sigma_{s_n}$

To solve this inference challenge, we modify the standard VAE architecture by incorporating the Gumbel-Softmax trick (Jang et al., 2016). We employ an LSTM network, taking history embeddings $v_{1:n}$ as inputs and generating one of C category labels through the Gumbel-Softmax technique, denoted as $w = \text{LSTM}(v_{1:n}) \in \text{Cat}(\pi)$. Here $\text{Cat}(\pi)$ represents the categorical distribution with probabilities $\pi \in \Delta^C$. Simultaneously, we capture hidden states at each time point as $h_{z_{1:n}}$. Subsequently, we utilize an MLP to process both the category label and hidden states as input, producing the mean $\mu_{s_n} \in \mathbb{R}^4$ and log-variance $\log \sigma_{s_n} \in \mathbb{R}^4$ of latent states s_n for each time point.

Table 8 presents an overview of the PSI-KT model architecture and hyperparameters used for all experiments.

A.5 PREDICTION AND GENERALIZATION EXPERIMENTS DETAILS

A.5.1 WITHIN-LEARNER PREDICTION RESULTS AND TRAINING HYPERPARAMETERS

In our prediction experiments, we employ a supervised training approach. For each learner, the first 10 interactions from their learning history are used for training, with the subsequent 10 interactions used as the test set. To report results, we reserve 20% of the learners as a validation set. We employ the Adam optimizer (Kingma & Ba, 2014) with an initial learning rate of 0.005 and apply gradient clipping with a threshold of 10.0. We use a linear decay schedule for the learning rate, halving it every 200 epochs. Additionally, we maintain a consistent batch size of 32 across models.

In Figure 2 in the main text, we present the average accuracy curves for comparison. For a more comprehensive overview of our training protocols, including accuracy, F1-score, and their standard deviation across 5 random seeds, please refer to the detailed results provided in Appendix Tables 9 and 10.

In our baseline models, the original approach was to predict a single time point in the future using all available historical data. However, we believe that relying solely on short-term predictions is insufficient for capturing long-term trends in learners’ performance, which is crucial for making accurate recommendations for customized learning materials. Moreover, it’s often impractical to assume that we can always access ground-truth data for immediate predictions. Therefore, we predict 10 time points into the future, using the predicted performances as inputs for each step. In other words, instead of using ground-truth data, if the model can predict \hat{y}_n based on all previous training data $y_{n' < n}$, we incorporate the predicted performance along with the historical data $[y_{n' < n}; \hat{y}_n]$ to predict \hat{y}_{n+1} .

In the evaluations, we chose to focus on prediction and generalization on a small group of learners, with numbers ranging from 100 to 1,000. This decision is based on the reality that, in educational settings, large datasets are not always available or practical. Additionally, little data is key in practical ITS to minimize the number of learners on an experimental treatment, to mitigate the cold-start

Table 9: Accuracies for within-learner prediction across numbers of learners (mean \pm SEM across random seeds).

Dataset	# Learners	HLR	PPE	DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
Assist12	100	.54 _{.03}	.65 _{.01}	.65 _{.03}	.60 _{.01}	.55 _{.01}	.67 _{.02}	.63 _{.03}	.63 _{.03}	.68 _{.02}
	200	.55 _{.02}	.63 _{.03}	.66 _{.02}	.62 _{.01}	.58 _{.01}	.67 _{.02}	.61 _{.02}	.66 _{.02}	.70 _{.02}
	300	.55 _{.01}	.66 _{.01}	.67 _{.01}	.62 _{.00}	.58 _{.01}	.69 _{.02}	.65 _{.02}	.65 _{.02}	.71 _{.01}
	400	.55 _{.01}	.65 _{.01}	.68 _{.01}	.63 _{.01}	.60 _{.02}	.67 _{.03}	.63 _{.02}	.66 _{.01}	.71 _{.01}
	500	.55 _{.01}	.64 _{.01}	.67 _{.01}	.63 _{.01}	.59 _{.03}	.67 _{.02}	.63 _{.02}	.65 _{.02}	.70 _{.01}
	1,000	.54 _{.00}	.65 _{.00}	.68 _{.01}	.63 _{.01}	.60 _{.01}	.70 _{.02}	.64 _{.01}	.64 _{.01}	.70 _{.01}
Assist17	100	.45 _{.01}	.53 _{.02}	.57 _{.02}	.53 _{.03}	.52 _{.03}	.56 _{.02}	.56 _{.04}	.58 _{.02}	.63 _{.02}
	200	.45 _{.01}	.53 _{.02}	.57 _{.02}	.54 _{.02}	.54 _{.01}	.55 _{.01}	.56 _{.02}	.60 _{.02}	.63 _{.01}
	300	.46 _{.01}	.53 _{.01}	.57 _{.02}	.55 _{.02}	.55 _{.02}	.56 _{.04}	.58 _{.02}	.61 _{.01}	.63 _{.01}
	400	.45 _{.01}	.53 _{.01}	.56 _{.01}	.57 _{.02}	.56 _{.02}	.56 _{.02}	.58 _{.02}	.61 _{.01}	.64 _{.00}
	500	.46 _{.01}	.53 _{.00}	.60 _{.01}	.58 _{.01}	.54 _{.01}	.56 _{.02}	.58 _{.01}	.61 _{.02}	.63 _{.01}
	1,000	.44 _{.01}	.55 _{.01}	.60 _{.01}	.57 _{.01}	.57 _{.01}	.61 _{.01}	.60 _{.01}	.63 _{.01}	.64 _{.00}
Junyi15	100	.55 _{.02}	.66 _{.03}	.79 _{.03}	.78 _{.01}	.63 _{.02}	.81 _{.02}	.78 _{.02}	.81 _{.02}	.83 _{.02}
	200	.57 _{.01}	.65 _{.03}	.79 _{.01}	.78 _{.02}	.68 _{.03}	.80 _{.01}	.80 _{.01}	.80 _{.01}	.84 _{.01}
	300	.56 _{.02}	.65 _{.03}	.81 _{.01}	.79 _{.01}	.70 _{.01}	.81 _{.01}	.78 _{.02}	.81 _{.01}	.85 _{.01}
	400	.61 _{.02}	.65 _{.02}	.81 _{.01}	.80 _{.02}	.69 _{.02}	.82 _{.02}	.75 _{.02}	.80 _{.01}	.85 _{.01}
	500	.61 _{.01}	.67 _{.02}	.82 _{.01}	.80 _{.02}	.70 _{.01}	.82 _{.02}	.78 _{.02}	.81 _{.01}	.85 _{.01}
	1,000	.59 _{.01}	.66 _{.02}	.81 _{.01}	.81 _{.00}	.69 _{.01}	.82 _{.01}	.79 _{.02}	.83 _{.01}	.85 _{.01}

Table 10: F1 scores for within learner prediction across learner numbers (mean \pm SEM across random seeds.)

Dataset	# Learners	HLR	PPE	DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
Assist12	100	.59 _{.02}	.77 _{.01}	.77 _{.03}	.72 _{.01}	.64 _{.01}	.79 _{.02}	.76 _{.01}	.73 _{.03}	.80 _{.01}
	200	.60 _{.02}	.74 _{.03}	.78 _{.02}	.73 _{.01}	.68 _{.01}	.76 _{.02}	.74 _{.02}	.77 _{.02}	.82 _{.01}
	300	.59 _{.02}	.77 _{.01}	.79 _{.01}	.74 _{.00}	.69 _{.01}	.73 _{.03}	.76 _{.02}	.77 _{.01}	.83 _{.01}
	400	.60 _{.02}	.77 _{.01}	.79 _{.01}	.74 _{.01}	.70 _{.03}	.73 _{.03}	.75 _{.01}	.76 _{.01}	.83 _{.01}
	500	.60 _{.01}	.76 _{.01}	.79 _{.01}	.74 _{.01}	.64 _{.10}	.74 _{.02}	.75 _{.02}	.76 _{.01}	.82 _{.01}
	1,000	.60 _{.01}	.76 _{.00}	.79 _{.00}	.74 _{.01}	.71 _{.01}	.73 _{.02}	.76 _{.01}	.76 _{.01}	.82 _{.00}
Assist17	100	.45 _{.01}	.44 _{.01}	.42 _{.02}	.40 _{.03}	.42 _{.03}	.40 _{.02}	.40 _{.02}	.41 _{.02}	.48 _{.03}
	200	.45 _{.01}	.44 _{.01}	.40 _{.03}	.42 _{.01}	.43 _{.01}	.44 _{.01}	.41 _{.02}	.43 _{.02}	.47 _{.04}
	300	.45 _{.01}	.45 _{.02}	.40 _{.02}	.41 _{.01}	.42 _{.03}	.45 _{.03}	.42 _{.01}	.44 _{.03}	.46 _{.03}
	400	.44 _{.01}	.44 _{.02}	.41 _{.01}	.42 _{.02}	.43 _{.02}	.45 _{.03}	.42 _{.02}	.45 _{.01}	.47 _{.03}
	500	.46 _{.01}	.45 _{.01}	.40 _{.01}	.42 _{.00}	.40 _{.10}	.45 _{.02}	.43 _{.02}	.45 _{.01}	.47 _{.02}
	1,000	.44 _{.01}	.44 _{.02}	.40 _{.01}	.43 _{.00}	.43 _{.03}	.46 _{.02}	.43 _{.02}	.47 _{.01}	.47 _{.04}
Junyi15	100	.53 _{.02}	.70 _{.03}	.88 _{.02}	.87 _{.01}	.75 _{.03}	.89 _{.01}	.87 _{.01}	.89 _{.01}	.92 _{.01}
	200	.54 _{.02}	.71 _{.02}	.88 _{.01}	.87 _{.01}	.80 _{.02}	.88 _{.01}	.88 _{.01}	.89 _{.01}	.91 _{.01}
	300	.53 _{.02}	.71 _{.02}	.89 _{.01}	.88 _{.01}	.80 _{.01}	.90 _{.01}	.87 _{.02}	.89 _{.01}	.92 _{.01}
	400	.52 _{.02}	.72 _{.03}	.89 _{.01}	.88 _{.01}	.80 _{.01}	.90 _{.01}	.87 _{.02}	.89 _{.01}	.92 _{.02}
	500	.53 _{.01}	.70 _{.02}	.89 _{.01}	.88 _{.01}	.74 _{.08}	.88 _{.01}	.86 _{.01}	.89 _{.01}	.92 _{.01}
	1,000	.52 _{.01}	.71 _{.02}	.90 _{.01}	.89 _{.00}	.80 _{.02}	.90 _{.01}	.85 _{.01}	.90 _{.01}	.93 _{.00}

problem, and extend the usefulness of the model to classroom-size groups. To provide ITS with a basis for adaptive guidance and long-term learner assessment, we always predict the 10 next interactions.

In order to ensure a fair evaluation of deep learning models and to avoid biasing our results, we expanded our dataset to include over 1,000 learners. This expansion was done post-filtering, where we excluded learners with fewer than 50 interactions. Additionally, 20% of these learners were designated as a validation set. The average accuracy, along with the number of learners and the number of parameters used in each model, is detailed in Table 11.

It’s crucial to recognize that deep learning models, despite benefiting from extensive datasets, face specific challenges. Firstly, PSI-KT has remarkable predictive performance when trained on small cohorts whereas baselines require training data from at least 60k learners to reach similar performance. Secondly, the deployment of these deep learning models in real-time applications is challenging due to their substantial number of parameters.

A.5.2 BETWEEN-LEARNER FINE-TUNING HYPERPARAMETERS

For between-learner generalization, we employ pre-trained models from within learner prediction, where the details can be found in Appendix A.5.1. These models are trained using data from 100 learners, and we retain the one that achieved the highest prediction accuracy on the validation set.

Table 11: Accuracy score in within-learner prediction with all learners in each dataset (mean \pm SEM across random seeds).

Dataset	# Learners	HLR	PPE	DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
Assist17	1,358	.46 _{.00}	.55 _{.00}	.58 _{.01}	.55 _{.01}	.57 _{.01}	.60 _{.01}	.60 _{.01}	.61 _{.01}	.64 _{.00}
Assist12	9,954	.44 _{.02}	.47 _{.00}	.69 _{.00}	.66 _{.00}	.66 _{.00}	.68 _{.01}	.70 _{.00}	.68 _{.00}	.70 _{.01}
Junyi15	62,124	.65 _{.01}	.71 _{.01}	.85 _{.00}	.85 _{.00}	.84 _{.01}	.86 _{.01}	.86 _{.02}	.86 _{.00}	.85 _{.01}

Then, predictions are made by randomly selecting 100 learners from the group that were not included in the training or validation sets.

In the experiment without fine-tuning, we directly apply the pre-trained models to unseen out-of-sample learners and present the results in Table 2. This entails using the pre-trained models to predict the next 10 interactions for out-of-sample learners based on their first 10 interactions as input.

In the fine-tuning experiment, we perform fine-tuning for each model using a batch size of 32. Additionally, we also set aside 20% of the learners as a validation set during this process to save the model that achieves the highest accuracy after fine-tuning. For baseline models, HLR, PPE, and HKT, which are comprised entirely of learner-independent and KC-dependent parameters, we conduct fine-tuning for all of these parameters. In this scenario, we use the pre-trained models as the initial weight values for the fine-tuning process. Conversely, for models DKT, DKTF, and AKT, we perform fine-tuning specifically on their KC embedding parameters and the last fully connected layer within the neural network, while keeping the remaining layers frozen during the fine-tuning process.

A.5.3 CONTINUAL-LEARNING RESULTS

Table 12: Continual learning accuracy. We report accuracy in predicting 10 subsequent outcomes. # Data indicate the number of interactions from each learner for training.

Dataset	# Data	10	20	30	40	50	60	70	80	90
Assist12	HLR	.54 _{.03}	.57 _{.08}	.58 _{.08}	.59 _{.09}	.57 _{.10}	.56 _{.07}	.54 _{.07}	.55 _{.06}	.57 _{.08}
	PPE	.65 _{.01}	.55 _{.07}	.53 _{.07}	.52 _{.08}	.54 _{.06}	.57 _{.06}	.59 _{.06}	.61 _{.06}	.69 _{.04}
	DKT	.65 _{.03}	.66 _{.07}	.64 _{.06}	.68 _{.05}	.69 _{.04}	.66 _{.05}	.66 _{.05}	.68 _{.03}	.65 _{.01}
	DKTF	.60 _{.01}	.67 _{.04}	.65 _{.04}	.64 _{.04}	.66 _{.03}	.62 _{.06}	.61 _{.04}	.63 _{.02}	.63 _{.02}
	HKT	.55 _{.01}	.56 _{.05}	.62 _{.04}	.62 _{.05}	.63 _{.02}	.60 _{.02}	.61 _{.03}	.61 _{.02}	.62 _{.02}
	AKT	.67 _{.02}	.66 _{.04}	.62 _{.04}	.61 _{.04}	.61 _{.05}	.65 _{.02}	.62 _{.02}	.61 _{.02}	.63 _{.02}
	GKT	.65 _{.02}	.62 _{.02}	.62 _{.01}	.64 _{.05}	.65 _{.04}	.65 _{.03}	.66 _{.06}	.65 _{.05}	.65 _{.05}
	QIKT	.70 _{.02}	.63 _{.01}	.64 _{.02}	.63 _{.01}	.62 _{.03}	.62 _{.01}	.62 _{.02}	.62 _{.02}	.63 _{.01}
	PSI-KT	.68 _{.02}	.70 _{.03}	.68 _{.03}	.72 _{.03}	.75 _{.02}	.73 _{.03}	.74 _{.02}	.74 _{.02}	.74 _{.02}
	Assist17	HLR	.45 _{.01}	.46 _{.07}	.45 _{.07}	.53 _{.06}	.55 _{.08}	.57 _{.06}	.55 _{.06}	.55 _{.04}
PPE		.53 _{.02}	.52 _{.06}	.52 _{.06}	.52 _{.07}	.52 _{.06}	.52 _{.05}	.51 _{.05}	.54 _{.04}	.56 _{.04}
DKT		.57 _{.02}	.52 _{.05}	.52 _{.05}	.52 _{.06}	.59 _{.04}	.57 _{.05}	.60 _{.04}	.63 _{.02}	.59 _{.03}
DKTF		.53 _{.03}	.58 _{.05}	.54 _{.05}	.58 _{.05}	.58 _{.04}	.55 _{.05}	.56 _{.05}	.56 _{.04}	.61 _{.02}
HKT		.52 _{.03}	.57 _{.04}	.60 _{.03}	.60 _{.03}	.62 _{.02}	.61 _{.03}	.61 _{.02}	.60 _{.02}	.61 _{.02}
AKT		.56 _{.02}	.53 _{.05}	.52 _{.06}	.54 _{.04}	.53 _{.04}	.53 _{.02}	.50 _{.03}	.51 _{.03}	.57 _{.02}
GKT		.63 _{.02}	.59 _{.05}	.54 _{.04}	.60 _{.04}	.56 _{.03}	.54 _{.02}	.57 _{.02}	.58 _{.02}	.58 _{.03}
QIKT		.65 _{.02}	.58 _{.03}	.59 _{.03}	.56 _{.05}	.58 _{.03}	.56 _{.02}	.58 _{.02}	.58 _{.01}	.56 _{.02}
PSI-KT		.63 _{.02}	.62 _{.04}	.65 _{.04}	.60 _{.05}	.60 _{.05}	.62 _{.05}	.62 _{.04}	.62 _{.04}	.64 _{.03}
Junyi15		HLR	.55 _{.02}	.43 _{.06}	.42 _{.06}	.44 _{.05}	.60 _{.04}	.63 _{.04}	.63 _{.03}	.63 _{.04}
	PPE	.66 _{.03}	.67 _{.06}	.64 _{.06}	.64 _{.05}	.62 _{.04}	.63 _{.05}	.60 _{.05}	.60 _{.03}	.61 _{.03}
	DKT	.79 _{.03}	.80 _{.04}	.78 _{.04}	.76 _{.05}	.77 _{.04}	.75 _{.04}	.84 _{.04}	.73 _{.02}	.74 _{.01}
	DKTF	.78 _{.01}	.74 _{.05}	.77 _{.05}	.74 _{.06}	.71 _{.05}	.71 _{.04}	.74 _{.03}	.71 _{.03}	.72 _{.02}
	HKT	.63 _{.02}	.63 _{.08}	.69 _{.07}	.67 _{.07}	.70 _{.04}	.73 _{.04}	.73 _{.03}	.79 _{.02}	.84 _{.02}
	AKT	.81 _{.02}	.79 _{.04}	.78 _{.05}	.79 _{.04}	.75 _{.04}	.75 _{.03}	.76 _{.03}	.74 _{.02}	.74 _{.03}
	GKT	.82 _{.01}	.80 _{.02}	.78 _{.03}	.78 _{.03}	.79 _{.04}	.79 _{.03}	.79 _{.03}	.79 _{.02}	.80 _{.02}
	QIKT	.84 _{.00}	.80 _{.02}	.80 _{.05}	.78 _{.03}	.78 _{.04}	.81 _{.03}	.80 _{.02}	.78 _{.01}	.85 _{.02}
	PSI-KT	.83 _{.02}	.81 _{.04}	.81 _{.04}	.80 _{.04}	.77 _{.06}	.81 _{.04}	.82 _{.04}	.83 _{.03}	.84 _{.03}

In this experiment, we randomly select 100 learners using five different random seeds, and collect their first 100 interactions. Initially, the models are trained using only the first 10 interactions, following the same setup as in the within-learner prediction experiment. Following the initial training, we continuously integrate new interaction data into the training process, introducing one interaction at a time for each learner. The model iteratively predicts the subsequent 10 performances. This simulates a common real-world scenario, where learners continually interact with existing or even new KCs.

In the objective function ELBO shown in Eq. 10, all historical information up to time t_n has already been fully encoded into the variational parameters ϕ_n . Additionally, to allow for the

Table 13: Specificity, consistency, and disentanglement.

Metric	Dataset	DKT	DKTF	AKT	QIKT	PSI-KT
Specificity $\text{MI}(s; \ell) \uparrow$	Assist12	8.83	6.62	6.45	2.47	<u>8.40</u>
	Assist17	8.08	7.50	10.05	2.95	<u>9.98</u>
	Junyi15	12.75	<u>13.50</u>	13.34	4.09	14.37
Consistency ⁻¹ $\mathbb{E}_{\ell_{\text{sub}}} \text{MI}(s^{\ell}; \ell_{\text{sub}}) \downarrow$	Assist12	14.13	12.24	20.15	<u>8.35</u>	7.48
	Assist17	14.95	13.11	24.47	<u>6.35</u>	6.35
	Junyi15	13.10	17.81	22.15	<u>7.66</u>	5.00
Disentanglement $D_{\text{KL}}(s \ell) \uparrow$	Assist12	-1.64	0.38	-8.17	<u>2.31</u>	7.42
	Assist17	-3.01	-0.44	-9.81	<u>0.56</u>	8.39
	Junyi15	-0.62	<u>4.96</u>	-6.65	1.57	11.49

possibility of learners encountering a new KC during their learning journey, we allow for optimization over the KC parameters in the generative model. As a result, when new interaction data $\mathcal{I} = (x_{n+1}, y_{n+1}, t_{n+1})^{1:L}$ becomes available at time t_{n+1} , we use the new data to update both the inference model parameters ϕ_{n+1} and the generative model parameters U and M which are related to KCs.

For the baseline models, which are designed to predict performances based on fixed learning histories, there is no need to update the model parameters for each new data point from each learner individually. Instead, when new interaction data, denoted as $\mathcal{I} = (x_{n+1}, y_{n+1}, t_{n+1})^{1:L}$, becomes accessible at time t_{n+1} , we update all the model parameters using all the interaction data collected up to that point, referred to as $\mathcal{H}_{1:n+1}$. This update is performed through 10 gradient descent processes. It is important to note that we do not include an additional validation set to determine when to stop training each model separately. Instead, we aim for a fair comparison among all models, ensuring that they are trained on equal footing with the same limited data and resources available to them.

A.6 LEARNER-SPECIFIC REPRESENTATIONS ANALYSIS

In this experiment, we examine temporal latent features (learner representations) derived from baseline models. When considering baseline models, it is noteworthy that only DKT, DKTF, AKT, and QIKT incorporate learner-specific temporal embedding vectors. While HKT utilizes temporal embeddings, all these embeddings originate from global parameters associated with KCs, rendering them non-learner-specific. Consequently, our comparative analysis focuses exclusively on PSI-KT compared with DKT, DKTF, AKT, and QIKT.

We initially present comprehensive results in Table 13, complementing Table 3 from Section 4.3, wherein only the results from the best-performing baseline models are displayed. Subsequent sections will detail the experimental setups and metrics employed.

A.6.1 EXPERIMENTAL SETUP FOR SPECIFICITY

In personalized learning, we assume each learner has a unique cognitive profile shaped by past experiences and educational contexts. Our first step is to connect learner representations s_n^{ℓ} with these inherent learner-specific cognitive traits, i.e., the *specificity* of learners given corresponding representations.

To quantify specificity, we employ mutual information, denoted as $\text{MI}(s; \ell) := \text{H}(s) - \text{H}(s|\ell)$ among all learners, as a measure of the information shared between learner identities and learner representations. The detailed computation of the metric $\text{MI}(s; \ell)$ is outlined as follows:

$$\begin{aligned}
 \text{MI}(s; \ell) &= \text{H}(s) - \text{H}(s|\ell) \\
 &= - \int p(s) \log p(s) - \frac{1}{L} \sum_{\ell} \int p(s|\ell) \log p(s|\ell) \\
 &= \frac{1}{2} (D(1 + \log 2\pi) + \log |\Sigma_s|) - \frac{1}{2L} \sum_{\ell} (D(1 + \log 2\pi) + \log |\Sigma_{s^{\ell}}|) \\
 &= \frac{1}{2} (\log |\Sigma_s| - \frac{1}{L} \sum_{\ell} \log |\Sigma_{s^{\ell}}|). \tag{24}
 \end{aligned}$$

Here Σ_s and $\Sigma_{s^{\ell}}$ are the covariance matrices obtained from fitting learner representations from all L learners or, respectively, a single learner with a Gaussian distribution, and D is the dimensionality

of learner representations. In experiments, we begin by randomly selecting 1,000 learners from each dataset and then extracting their first 50 interactions for training. To determine when to stop training effectively, we set aside a validation set of 20% of learners, which amounts to 200 learners in our case. This setup mirrors our approach in the prediction experiment. The metric $\text{MI}(s; \ell)$ is calculated for the learners in the training set. Since our goal here is to evaluate the model’s capacity to distill representations s^ℓ that uniquely identify learners, there is no need for a test set. Note that the baseline models have higher-dimensional learner representations (16 dimensions in our experiments), potentially allowing them to capture more information.

A.6.2 EXPERIMENTAL SETUP FOR CONSISTENCY

We proceed with a supplementary *consistency* analysis to determine, among the shared information quantified in specificity, whether the learner representations capture intricate learner attributes or merely reflect transient dynamic fluctuations. In the experiment, we split the interaction data of each learner into five separate groups, i.e., subsets. Each subset contains 30 interactions. These specific sizes were chosen to ensure we have both enough learners for robust training and enough interactions in subsets to estimate covariance matrices for our metrics. We thus exclude learners who have engaged in fewer than 150 interactions.

To form subsets, we find out the average presentation time of each KC and assign the KCs to separate subsets, so that the overall average interaction time is as similar as possible across subsets. With this, we aim to wash out, to the extent possible with the limited amount of data, systematic biases in the partition induced by the dependence of learner representations on time.

The mutual information metric $\mathbb{E}_{\ell_{\text{sub}}}\text{MI}(s^\ell; \ell_{\text{sub}}) := \mathbb{E}_{\ell_{\text{sub}}}[\text{H}(s|\ell) - \text{H}(s|\ell_{\text{sub}})]$, employed in the consistency experiments, undergoes the a similar derivation process to Eq. 24.

$$\begin{aligned} \mathbb{E}_{\ell_{\text{sub}}}\text{MI}(s^\ell; \ell_{\text{sub}}) &= \frac{1}{L} \sum_{\ell} (\text{H}(s|\ell) - \frac{1}{5} \sum_{\ell_{\text{sub}}} \text{H}(s|\ell_{\text{sub}})) \\ &= \frac{1}{L} \sum_{\ell} \left(-\mathbb{E}[\log \mathcal{N}(\mu_{s^\ell}, \Sigma_{s^\ell}^2)] + \frac{1}{5} \sum_{\ell_{\text{sub}}} \mathbb{E}[\log \mathcal{N}(\mu_{s^{\ell_{\text{sub}}}}, \Sigma_{s^{\ell_{\text{sub}}}}^2)] \right) \\ &= \frac{1}{L} \sum_{\ell} (\log |\Sigma_{s^\ell}| - \frac{1}{5} \sum_{\ell_{\text{sub}}} \log |\Sigma_{s^{\ell_{\text{sub}}}}|). \end{aligned} \tag{25}$$

We fit each sub-learner separately and quantify the divergence metric $\mathbb{E}_{\ell_{\text{sub}}}\text{MI}(s^\ell; \ell_{\text{sub}})$ between learners and their sub-learners. A lower value of $\mathbb{E}_{\ell_{\text{sub}}}\text{MI}(s^\ell; \ell_{\text{sub}})$ suggests a higher degree of consistency, reflecting the difficulty in distinguishing between sub-learners and their corresponding overarching learners given learner representations. Overall, Table 3 shows that the learner representations of PSI-KT provide comparable learner specificity and superior consistency. The lower consistency displayed by baseline models suggests that most of the representational capacity available in their higher-dimensional representations might be spent on capturing learner-unspecific characteristics of the training sample.

A.6.3 EXPERIMENTAL SETUP FOR DISENTANGLEMENT

With the insights gained from specificity, our analysis progresses to evaluating to what extent learner-specific representations, are disentangled. Disentanglement in machine learning has been characterized as the process of isolating and identifying distinct, independent, and informative generative factors of variation in the data (Bengio et al., 2013).

In our disentanglement experiments, we use the same setup for specificity, and compute the discrepancy $D_{\text{KL}}(s||\ell)$ based on 50 interactions of 1,000 learners. Our approach bears similarity to (Kim & Mnih, 2018), but we relax the unrealistic assumption of independent representations. In real-world scenarios, independence in cognitive attributes is not a priority. To assess how much information about learner identity is present in the covariance across the representation dimension, we use the divergence between full trait-vector entropy and diagonal learner-conditional trait-vector entropy.

Table 14: (regression coefficient, p -value) tuples for performance difference and initial performance across models and latents’ dimensions. If there is no significant dimension in one model and dataset ($p > 0.05$), we show the dimension with the highest regression coefficient. Bold values indicate the dimension and the baseline model with the highest statistically significant linear relationship in one dataset, with which we show the regression results in Figures 8 and 9.

Behavioural signature	Dataset	DKT	DKTF	AKT	PSI-KT
Performance difference	Assist12	(-.009, .643)	(.008, .736)	(-.009, .665)	(.300, <.001)
	Assist17	(-.008, .758)	(-.029, .304)	(-.001, .957)	(.556, <.001)
	Junyi15	(-.003, .853)	(.021, .771)	(.025, .064)	(.721, <.001)
Initial performance	Assist12	(.021, .078)	(.039, .008)	(.017, <.001)	(.544, <.001)
	Assist17	(.048, .004)	(.038, .030)	(.010, .030)	(3.705, <.001)
	Junyi15	(-.025, .034)	(.044, .021)	(.017, <.001)	(.921, <.001)

The discrepancy $D_{\text{KL}}(s|\ell)$ is estimated by the full entropy of representations $\text{H}(s)$ and the diagonal elements of the covariance matrix in the conditional entropy $\text{H}(s|\ell)$

$$D_{\text{KL}}(s|\ell) := \text{H}(s)_{\text{full}} - \text{H}(s|\ell)_{\text{diag}} = \frac{1}{2} \left(\log |\Sigma_s| - \frac{1}{L} \sum_{\ell} \sum_{i=1}^D \log (\Sigma_{s|\ell})_{ii} \right). \quad (26)$$

Small non-diagonal elements of the covariance matrix in $\text{H}(s)$ suggest low cross-correlations. This can be interpreted as a form of disentanglement. As illustrated in the third row of Table 3, the representations from PSI-KT consistently exhibit a higher degree of disentanglement across all datasets.

A.6.4 MIXED-EFFECT LINEAR REGRESSIONS IN OPERATIONAL INTERPRETABILITY

Mixed-effects regression extends linear regression to handle data with hierarchical or clustered structures, such as repeated measurements from the same subjects (learners in our case). Taking one of our experiments as an example, we conduct regressions based on $y_n^\ell \sim \tilde{\mu}_n^{\ell,k} + (1 | \text{learner})$. Here, y_n^ℓ represents the dependent variable and $\tilde{\mu}_n^{\ell,k}$ is a predictor variable at time t_n . Also, $(1 | \text{learner})$ represents the random intercept associated with each learner. This random intercept accounts for variability between learners that cannot be explained by the fixed effect $\tilde{\mu}_n^{\ell,k}$. In other words, it accounts for the fact that different learners might have different biases in their responses, allowing us to capture a more robust estimate of the group-level effect.

For regression calculations, we use the models trained in the prediction experiments, as described in Section A.5. For consistent comparisons with specificity experiments, we opt for models trained on a group of 1,000 learners. This experiment goes beyond a simple sanity check (as in Sec. A.6.1), so we use the testing data. This choice aligns with our objective of using operational interpretability to gain insights and inform future controlled experiments with unseen data. We use pre-trained models, specifically DKT, DKTF, AKT, and our PSI-KT model, selected based on their accuracy scores on the validation data.

To fairly compare with baseline models, we investigate whether any dimensions within the learner representations capture behaviors similar to our interpretable cognitive traits. Thus, we perform regression for each dimension within the learner representations of the baseline models. While Figure 4 in the main paper presents the regression results concerning the dimension featuring the most pronounced correlation among baseline models, we provide a complete list of dimensions that exhibit significant relationships with the behavioral data in Table 14.

Performance decay and forgetting rate To analyze the exponential decay of learner performances over time, we first show the relationship between performance decay Δy_n^ℓ and the raw time difference τ_n^ℓ , which is divided into 10 bins. We select bin centers to ensure an equal number of data points in each bin. This binning approach helps minimize the impact of outliers and ensures a balanced representation of data within each bin. Also, we show the relationship between decay Δy_n^ℓ and the time difference scaled by the corresponding forgetting rate α_n^ℓ at each time point, or each dimension of learner representations in the baseline models. We assume that if the forgetting rate α_n^ℓ is meaningful for each time interval and effectively controls the decay, then the standard error of behavior data Δy_n^ℓ within each bin should be smaller than the error of binning raw time differences.

This indicates that the decay is better described as a function of $\alpha_n^\ell \tau_n^\ell$ than as a function of τ_n^ℓ alone. We also compute the standard error for each dimension of learner representations in the baseline models, and we show the dimension v_n^* with the lowest standard error in Figure 8. Then, we perform mixed-effect regression over the exponential term $\exp(-\alpha_n^\ell \tau_n^\ell)$ (or $\exp(-v_n^* \tau_n^\ell)$ in baselines) to assess how well learner representations predict performance decay (as an exponential function). The results show that at least one dimension in the learner representations groups certain behavioral data and reduces the standard errors. However, none of these dimensions exhibit a statistically significant relationship with the behavioral data.

Initial performance and long-term mean We conducted a mixed-effect regression analysis between the initial performance and the long-term mean, with the results presented in Figure 9. These results indicate that, except DKT on the Assist12 dataset, at least one dimension in the baseline learner representations predicts initial performance. It is important to note that none of the dimensions exhibit a stronger effect compared to the identified trait $\tilde{\mu}_n^{\ell,k}$ in our PSI-KT. Additionally, we note that embedding dimensions in baselines are trained in a permutation-invariant manner, suggesting that these models can't route any particular generative factor of variation in the data (e.g. a behavioral signature) to a specific dimension.

Prerequisite transfer ability and learning variances In our experiments, we sought to correlate two additional cognitive traits – transfer ability γ and learning volatility σ - with behavioral data. This task proved more complex than assessing the forgetting rate and long-term mean because assessing transfer ability requires reliable annotations of prerequisite relations and learning volatility can be connected to many unconstrained factors during the learning process.

Regarding transfer ability, our hypothesis posits that given the identified prerequisite KC i for KC j , a higher transfer ability γ_n^ℓ suggests an increased likelihood of correctly transitioning from one KC i to KC j . We calculate this transition probability $p(j^+ | i^+)_n^\ell$ by observing the frequency of correct responses to KC i followed by correct responses to KC j up to a certain time t_n . This implies that learners with greater transfer abilities are more likely to answer questions related to KC i correctly after mastering KC i . However, this approach depends on accurately identifying prerequisite relationships between KCs. Therefore, we utilized the Junyi15 dataset, which includes expert-annotated and crowd-sourced prerequisite graphs, for our regression analysis. For learning volatility σ_n^ℓ , we connect the average squared mean $(\bar{\sigma}^\ell)^2$ for each learner with the variance in their performance $\text{Var}(y_{1:n}^\ell)$.

In Figure 10, we present the results of our mixed-effect regression analyses. Each regression demonstrates a significant relationship. However, due to the sparsity of the expert-annotated graph, we do not have enough data to fit the regression model effectively. Thus we choose to use the crowd-sourcing graphs and consider the edge existence if the edge weight is above 0.5.

A.6.5 VISUALIZATION OF KNOWLEDGE STATES

In this section, we display the curve of inferred knowledge states within the Junyi15 dataset. We chose sequences where the involved skills are linked by established prerequisite relations. Two such prerequisites were identified: 'alternate interior angles' as a prerequisite for 'corresponding angles', and 'number properties terminology' for 'properties of numbers'. These prerequisites were determined based on crowd-sourcing annotations, where the average score for the annotated relation exceeded half. We note that PSI-KT can estimate knowledge states at all times and not just interaction times, which allows us to use natural time in the abscissae and display knowledge states with curves instead of using the discrete color maps common in the KT literature.

A.7 GRAPH INFERENCE ANALYSIS

A.7.1 DETAILS OF THE METRICS FOR GROUND-TRUTH GRAPH COMPARISON

Here we report comprehensive evaluations of the alignment of the inferred graphs with the human-annotated graphs in the Junyi15 dataset under different metrics.

As discussed in Section 4.3.2, the Junyi15 dataset provides two types of graph annotations - crowd-sourced similarity and prerequisite ratings (with 1,954 rated edges), as well as more sparse expert-annotated prerequisite relations (837 edges). We use the following metrics to compare graph representations learned by each model against these annotations:

1. **Mean Reciprocal Rank (MRR)**. We compare the inferred graph with the expert-annotated using the MRR, defined as $|K|^{-1} \sum_{i=1}^{|K|} (\text{rank}(i))^{-1}$, where K is the total number of KCs. We compute the rank of each expert-identified prerequisite relation $i \rightarrow k$ in the relevant sorted list of inferred probabilities $\{a^{jk}\}_{j=1}^K$ and take the harmonic average.
2. **Jaccard Similarity (JS)** is a classic measure of similarity between two sets, defined as the size of the intersection of set A and set B (i.e., the number of common elements) over the size of the union of set A and set B (i.e., the number of unique elements): $JS(A, B) = |A \cap B| / |A \cup B|$. Here we define the edge sets by thresholding weights at half the scale (0.5 for probability-scaled weights, 5 for the average of the 1-9 crowd-sourced rating).
3. **Negative Log-likelihood (nLL)** of edge weights given crowd-sourced annotations. The crowd-sourced annotations provide multiple 1-9 ratings per node pair. One set of annotations rates the strength of the directed prerequisite relations, whereas the other just rates the undirected similarity of the pair of nodes. We normalize the ratings from 0 to 1 and fit them with a Gaussian distribution. Then we compute the log-likelihood of the inferred edge probability under the Gaussian. The variance of the Gaussian accounts for inter-rater disagreements when comparing a model’s inferred edge probability with the mean edge rating.
4. **Linear Regression Coefficient** between edge weights and the causal support (details are in Sec. A.7.3) from node i to node k on correctness of k if having correct interactions on i . We compute the causal support for transitions of every KC pair. However we remove the causal support of pairs of KCs that have only one transition in the dataset to avoid adding noise to our estimate.

A.7.2 QUANTITATIVE COMPARISON RESULTS ON THE JUNYI15 DATASET

Note that the graphs of baselines are based on KC embeddings (as in Sec. A.3), and thus there is no edge directionality. For the baselines that have at least two embeddings for each KC, we can use to compute the directed edges, since one embedding for KC will end up in a symmetric structure adjacency matrix (DKT, DKTF, HKT, AKT). Thus, to conduct a fair comparison with the baseline models, we leniently compute edge weights based on every combination of KC embeddings. For example, in DKT, there are two embeddings $u^{k,0}, u^{k,1} \in \mathbb{R}^D$ representing incorrect interactions and correct interactions on KC k , respectively, and embeddings are shared across all learners. We compute the edge weights a^{ik} based on two different combinations here, both $a^{ik} := u^{i,0\top} u^{k,1}$ and $a^{ik} := u^{i,1\top} u^{k,0}$, and report the graph with the best results. When extracting undirected graphs, we concatenate all KC embeddings to compute $\tilde{a}^{ik} := (u^{i,0} + u^{k,1})^\top (u^{i,0} + u^{k,1})$, in order to reflect all available information from all KC embeddings. In the case of baselines with a single embedding per KC, such as QIKT, or those using a parameterized undirected graph, like GKT, we allow their inferred graphs to be less accurate. This means that for these models, the presence of an edge between two KCs is deemed correct if there is a directed edge from either direction in annotated graphs, without the necessity for these edges to accurately indicate the directionality. We then compute the weights by min-max normalization. This normalization is necessary for computing the log-likelihood, where we also use a threshold of 0.5 to determine whether there is an edge when the comparison requires binary edges.

In Table 5, we show the comparison of ground-truth prerequisite graphs and inferred graphs from our PSI-KT, and the best baseline models on the Junyi15 dataset under the different metrics. These results demonstrate that our inferred prerequisite graph outperforms others when compared with crowd-sourced and expert-annotated graphs under different metrics.

Here we show all of the comparison results, including a comparison of the similarity (undirected) graphs and the prerequisite (directed) graphs on four metrics in Table 15. We do not report MRR ranking scores for similarity graphs because the ground-truth similarity graph does not contain an expert-annotated version.

A.7.3 CAUSAL SUPPORT

Causal induction is the problem of inferring underlying causal structures from data. Here, we use a Bayesian framework (Griffiths & Tenenbaum, 2009; 2005) to infer a singular cause-and-effect relationship between all pairs of KCs, asking how performance on one node influences performance on another, and whether the strength of the causal relationship corresponds to our inferred prerequisite

Table 15: Comparison between ground-truth graphs and inferred graphs in Junyi15 dataset. **pre** indicates evaluation against a prerequisite graph and **sim** evaluation against a similarity graph.

		DKT	DKTF	HKT	AKT	GKT	QIKT	PSI-KT
MRR \uparrow	expert pre	.0069	.0067	.0074	.0075	<u>.0082</u>	.0073	.0086
JS \uparrow	expert pre	1.46e-3	1.37e-3	<u>1.47e-3</u>	1.44e-3	1.46e-3	1.19e-3	1.86e-3
	crowd pre	4.60e-3	4.28e-3	<u>4.66e-3</u>	4.48e-3	3.44e-3	5.21e-4	9.48e-3
	crowd sim	<u>5.90e-4</u>	0.00	0.00	5.18e-4	3.43e-3	0.00	4.66e-3
nLL \downarrow	crowd pre	5.735	5.580	6.092	5.677	3.033	4.228	4.106
	crowd sim	6.598	<u>4.039</u>	4.042	9.100	9.028	10.622	2.352

graph. In this context, we model the relationship between a candidate cause C and a candidate effect E (i.e., a pair of KCs), assuming an ever-present background cause B (i.e., the learner’s general ability and the influence of other nodes). The objectives are to determine the probability of a causal relationship between C and E , known as *causal support* (Eq. 27).

In our prerequisite graph scenario, we assume that if KC i is a prerequisite of KC k , the correctness on KC i contributes to the correctness on KC k . This implies the presence of a prerequisite relationship between KC i and KC k , signified by a causal link between their correctness levels. Consequently, for every pair of nodes, a candidate cause C corresponds to performance $y_n^i = 1$ at time t_n , and an effect E corresponds to $y_{n+1}^k = 1$ at time t_{n+1} , with inputs from all remaining nodes relegated to the background cause B .

When examining elemental causal induction, we adhere to the following two-step procedure: i) We establish the nature of the relationship through causal graphical models, and ii) we quantify the strength of the relationship, provided it exists, as a problem of inferring structural parameters. In the subsequent text, C and E variables are denoted using uppercase letters, while their specific instances are represented using lowercase letters. Specifically, c^+ and e^+ indicate the presence of the cause and effect (i.e., correct performance), whereas c^- and e^- signify their absence (i.e., incorrect performance).

Causal graphical models. Causal graphical models are a formalism for learning and reasoning about causal relationships (Glymour et al., 2019). Nodes in the graph represent variables, and directed edges represent causal connections between those variables. To identify whether a causal relationship exists between a pair of variables, we consider two directed graphs denoted Graph 0 $G_{C \rightarrow E} : B \rightarrow E$ and Graph 1 $G_{C \rightarrow E} : B \rightarrow E \leftarrow C$, as shown in Figure 5b. Thus, $G_{C \rightarrow E}$ represents the null hypothesis that there is no relationship between C and E (i.e., the effect E can be accounted for by background cause B), while $G_{C \rightarrow E}$ represents the alternative hypothesis that the causal relationship exists.

In our case, the cause C and the effect E are equivalent to KC i and KC k , respectively, for every pair of KCs. The process of inferring the underlying structure between KC i and KC k , whether the learners’ behavioral learning history \mathcal{H} are generated by $G_{i \rightarrow k}$ or $G_{i \leftarrow k}$, can be cast in a Bayesian framework (Griffiths & Tenenbaum, 2009; 2005). *Causal support* quantifies the degree of evidence present in the data \mathcal{H} that favors Graph 1 $G_{i \rightarrow k}$ over Graph 0 $G_{i \leftarrow k}$:

$$\text{support} = \log \frac{P(\mathcal{H} | G_{C \rightarrow E})}{P(\mathcal{H} | G_{C \leftarrow E})} = \log \frac{P(\mathcal{H} | G_{i \rightarrow k})}{P(\mathcal{H} | G_{i \leftarrow k})}. \quad (27)$$

Intuitively, the joint presence of the cause and effect, i.e., correctness on KC i followed by correctness on KC k , offers support for a causal link from node i to node k . Conversely, the absence of the cause, i.e., incorrectness on KC i but is followed by correctness on KC k , presents evidence against the notion that KC i is a prerequisite for KC k .

Causal support. Causal graphical models depict dependencies using conditional probabilities. Defining these probabilities entails parameterizing each edge, and this parameterization determines the functional expressions that govern causal relationships.

For Graph 0 $G_{i \rightarrow k}$ and Graph 1 $G_{i \rightarrow k}$, we define $P_0(y_{n+1}^k = 1 | B) = \omega_0$ and $P_1(y_{n+1}^k = 1 | y_n^i = 1) = \omega_1$ respectively. In other words, the probability of correctness on KC k given just background causes is ω_0 , and the probability of correctness on KC k given previous correctness on KC i is ω_1 ; and

when both prerequisite KC i and background causes are present, they have independent opportunities to produce the effect.

For **Graph 0** $G_{C \rightarrow E}$, the sole parameter ω_0 denotes the likelihood of the effect being present given the background cause

$$P_0(e^+ | b^+; \omega_0) = \omega_0. \quad (28)$$

The corresponding likelihood for the data \mathcal{H} given Graph 0 $G_{i \rightarrow k}$ is accomplished by integrating over all possible parameters ω_0 with a uniform prior over ω_0 :

$$\begin{aligned} P(\mathcal{H} | G_{i \rightarrow k}) &= \int_0^1 P_0(\mathcal{H} | \omega_0, G_{i \rightarrow k}) P(\omega_0 | G_{i \rightarrow k}) d\omega_0 \\ &= \int_0^1 \omega_0^{N(e^+)} (1 - \omega_0)^{N(e^-)} d\omega_0 \\ &= \text{Beta}(N(e^+) + 1, N(e^-) + 1) \\ &= \text{Beta}(N(y_{n+1}^k = 1) + 1, N(y_{n+1}^k = 0) + 1). \end{aligned} \quad (29)$$

Here $\text{Beta}()$ is the beta function, and $N(e^+)$ and $N(e^-)$ are the marginal frequencies of the effects.

For **Graph 1** $G_{i \rightarrow k}$, the likelihood of the effect is given by:

$$P_1(e^+ | b, c; \omega_0, \omega_1) = 1 - (1 - \omega_0)^b (1 - \omega_1)^c, \quad (30)$$

where ω_0 again defines the influence of the background cause, and the additional parameter ω_1 defines the influence of the cause. Here b and c are binary, which means if cause C exists, then $c = 1$. We compute the likelihood of the data $P(\mathcal{H} | G_{i \rightarrow k})$ by integrating over parameters ω_0 and ω_1 . Each parameter value is defined by a prior probability, which when combined with the likelihood of the data, yields a joint posterior distribution over data and parameters for the structure. To determine the observed data likelihood for Graph 1 $G_{i \rightarrow k}$, we have

$$\begin{aligned} P(\mathcal{H} | G_{i \rightarrow k}) &= \int_0^1 \int_0^1 P_1(\mathcal{H} | \omega_0, \omega_1, G_{i \rightarrow k}) P(\omega_0, \omega_1 | G_{i \rightarrow k}) d\omega_0 d\omega_1 \\ &= \int_0^1 \int_0^1 \prod_{e,c} P_1(e | c, b^+; \omega_0, \omega_1)^{N(e,c)} P(\omega_0, \omega_1 | G_{i \rightarrow k}) d\omega_0 d\omega_1. \end{aligned} \quad (31)$$

Here $N(e, c)$ represents the number of occurrences. To compute $\prod_{e,c} P_1(e | c, b^+; \omega_0, \omega_1)^{N(e,c)}$, we iterate over all possible sets of (e, c) . Based on Eq. 30, we get:

$$\begin{aligned} \prod_{e,c} P_1(e | c, b^+; \omega_0, \omega_1)^{N(e,c)} &= P_1(e^+ | c^+, b^+; \omega_0, \omega_1)^{N(e^+, c^+)} P_1(e^+ | c^-, b^+; \omega_0, \omega_1)^{N(e^+, c^-)} \\ &= (\omega_0 + \omega_1 - \omega_0 \omega_1)^{N(e^+, c^+)} \omega_0^{N(e^+, c^-)}. \end{aligned} \quad (32)$$

While Eq. 31 is not analytically tractable, it can be effectively approximated using Monte Carlo simulations. With uniform priors on ω_0 and ω_1 , a reliable estimation of $P(\mathcal{H} | G_{i \rightarrow k})$ can be obtained by generating m samples of ω_0 and ω_1 from a uniform distribution spanning the interval $[0, 1]$, followed by computation of:

$$\begin{aligned} P(\mathcal{H} | G_{i \rightarrow k}) &= \frac{1}{m} \sum_{i=1}^m P_1(\mathcal{H} | \omega_{0i}, \omega_{1i}, G_{i \rightarrow k}) \\ &= \frac{1}{m} \sum_{i=1}^m \prod_{e,c} P_1(e | c, b^+; \omega_{0i}, \omega_{1i})^{N(e,c)} \\ &= \frac{1}{m} \sum_{i=1}^m (\omega_{0i} + \omega_{1i} - \omega_{0i} \omega_{1i})^{N(e^+, c^+)} \omega_{0i}^{N(e^+, c^-)} \\ &= \frac{1}{m} \sum_{i=1}^m (\omega_{0i} + \omega_{1i} - \omega_{0i} \omega_{1i})^{N(y_{n+1}^k = 1, y_n^i = 1)} \omega_{0i}^{N(y_{n+1}^k = 1, y_n^i = 0)}. \end{aligned} \quad (33)$$

	Assist12	Assist17	junyi15
PSI-KT	.68 _{.017}	.63 _{.015}	.83 _{.015}
w/o Graph	-.04 _{.005}	-.04 _{.002}	-.07 _{.002}
w/o Individual traits	-.03 _{.002}	-.06 _{.006}	-.04 _{.001}
w/o Dynamic traits	-.06 _{.001}	-.09 _{.003}	-.03 _{.002}

Table 16: The accuracy in three kinds of ablation study (mean \pm SEM across random seeds). We show the accuracy gap compared with the complete PSI-KT model.

A.8 ABLATION STUDY

To thoroughly examine the various elements of PSI-KT, including cognitive traits and the prerequisite graph, we executed three distinct ablation studies:

- Without the graph inference (w/o graph): We omit the graph inference process and the influence of prerequisite KCs on the long-term mean. Essentially, this approach treats each KC independently.
- Without individual cognitive traits (w/o individual): We alter the variational inference network in this scenario to produce a uniform distribution across all learners. This change effectively removes the consideration of individual differences in learners’ cognitive traits.
- Without dynamic cognitive traits (w/o dynamics): We remove the dynamic transition distribution over the traits in the generative model. This assumes that each learner has static traits over time.

In Table 16 and Figure 13, we present the results of our three ablation studies. We observed that the contribution of prerequisite graphs, individualized traits, and dynamic traits varied across the datasets. These findings underscore the diversity inherent in educational datasets and simultaneously reinforce the effectiveness of our unified framework.

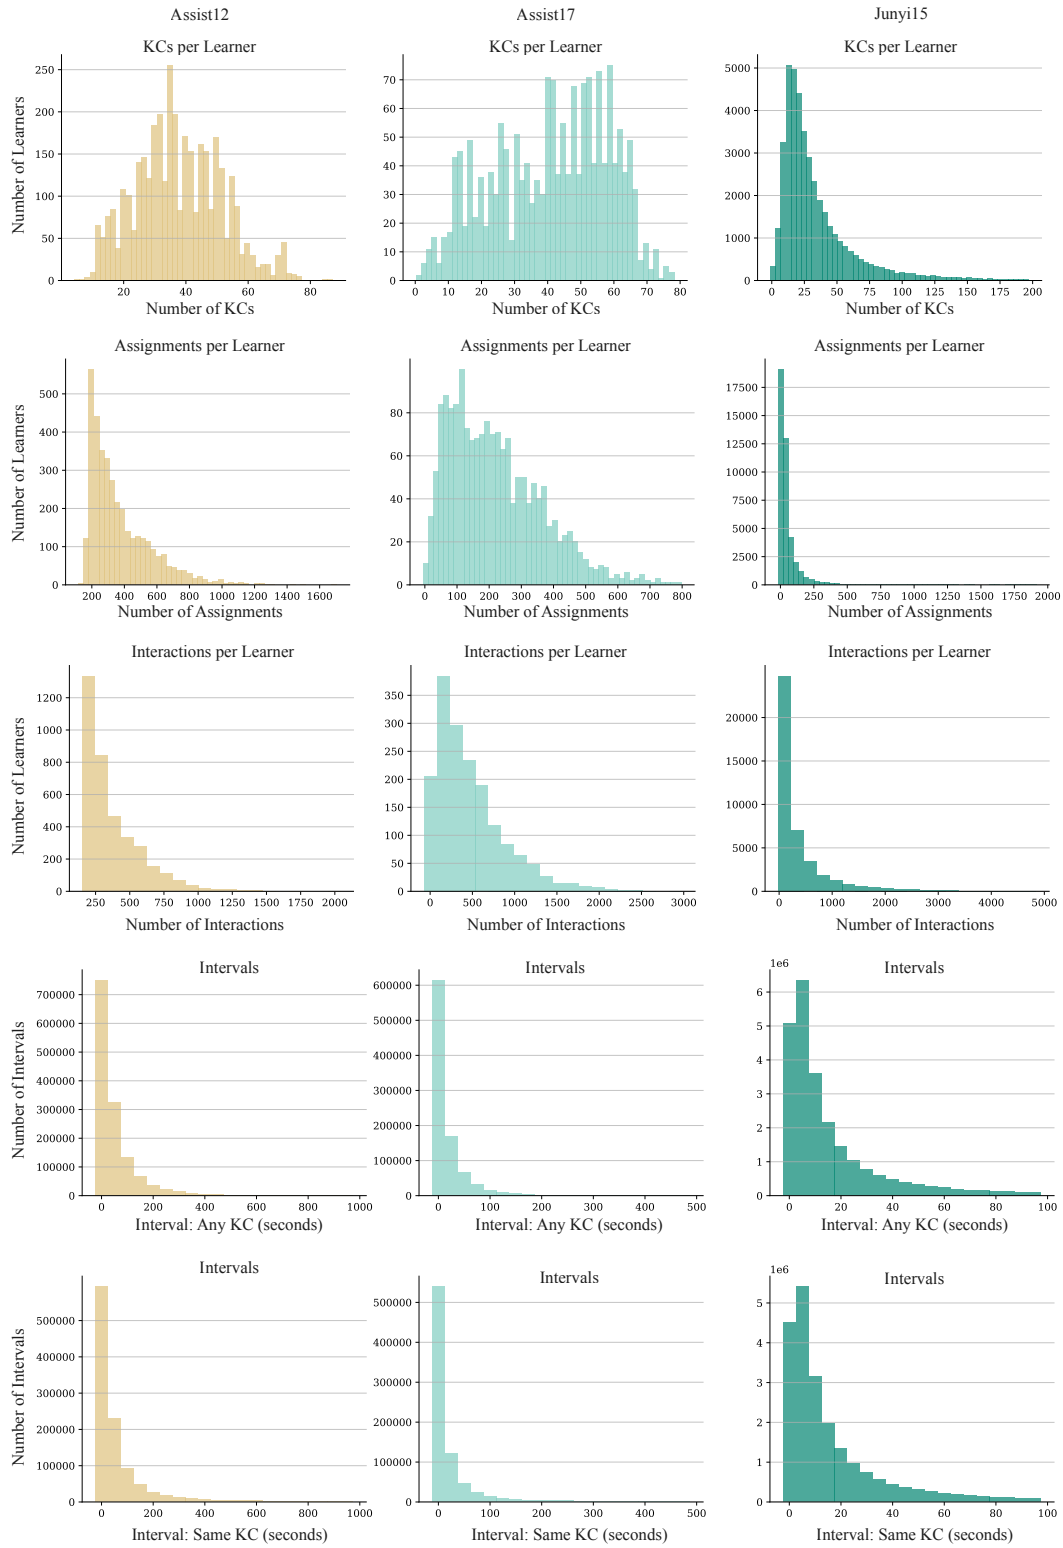


Figure 6: Histograms of key features in three datasets, including the number of interactions, KCs, and assignments per learner, and the intervals between two interactions with any KC and with the same KC.

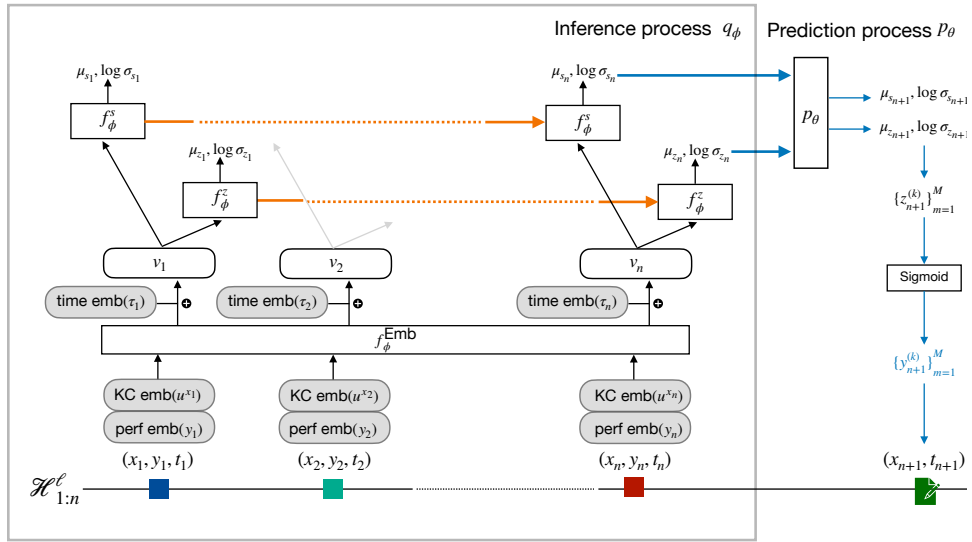


Figure 7: Inference model of PSI-KT with an example of a single learner’s history as the input. Note that all parameters ϕ, θ are shared across learners. Grey backgrounds mark inputs. Right rectangles are neural networks’ layers and rounded rectangles designate features. Layer f_ϕ^{Emb} maps the input features (time τ , KC u , and performance y , described in the text below) into embedding vectors. Layers f_ϕ^s and f_ϕ^z output the parameters of the variational posterior distribution. These are all part of the inference networks and parameterized by ϕ (surrounded by the grey box). The orange arrow is only applicable for inference on entire learning histories. Blue arrows represent the prediction stage, where during prediction M samples are drawn from the predicted distribution based on $\mu_{z_{n+1}}$ and $\log \sigma_{z_{n+1}}$.




Figure 8: The mixed-effect regressions of performance decay Δy_n^ℓ vs. scaled interval $\tau_n^\ell \alpha_n^\ell$ (scaled interval with the best dimension in baselines $\tau_n^\ell (v^*)^\ell$). The first row (a) shows the unscaled interval in the raw data. The aggregate standard error over 10 bins (SE), the regression coefficient (coef), and its p -value are reported in each panel.

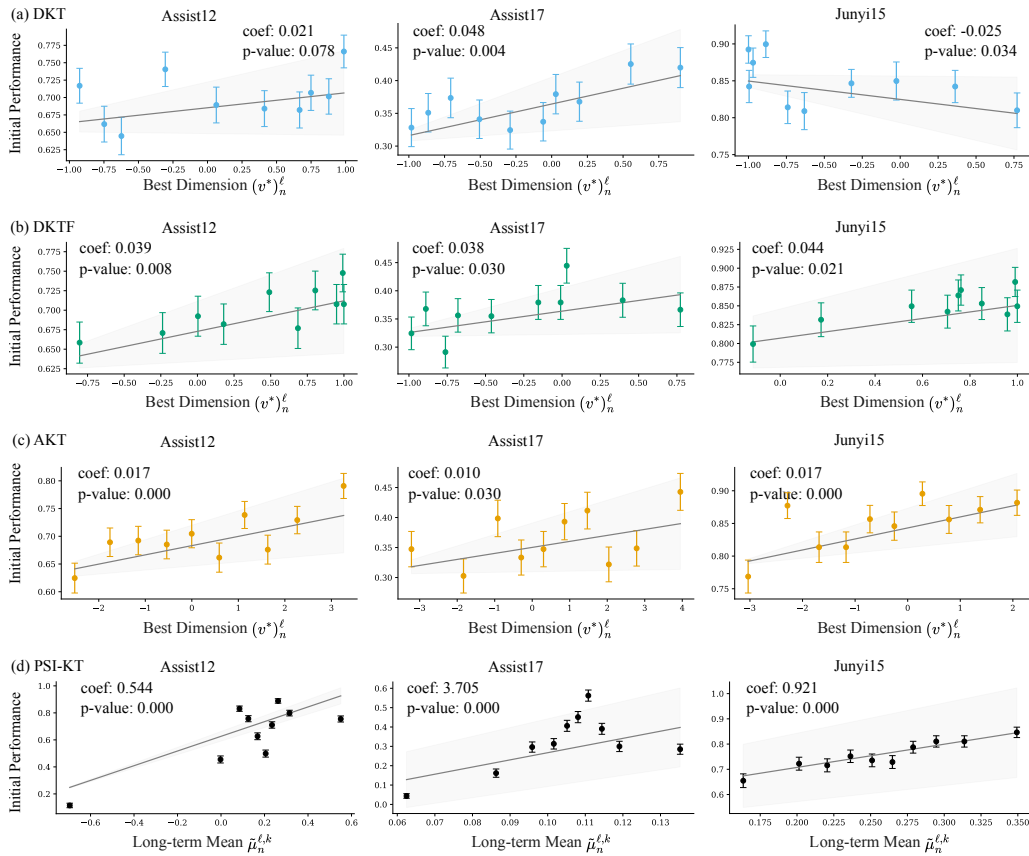


Figure 9: The mixed-effect regressions of initial performance y_n^ℓ vs. long-term mean $\tilde{\mu}_n^{\ell,k}$ (the best dimension $(v^*)^\ell_n$ in baselines). The aggregate standard error over 10 bins (SE), the regression coefficient (coef) and its p -value are reported in each panel.




Figure 10: The mixed-effect regressions of transfer ability γ with behavioral correct transition probability (a), and learning volatility σ with the variance in learning performances (b). We report the regression coefficient (coef) and its p -value in each panel, and each point illustrates the mean (\pm SEM) of the corresponding decile.




Figure 11: An example of inferred sequential knowledge states in the Junyi15 dataset.




Figure 12: Linear regressions relating causal support to the inferred edges for baseline models (a) DKT, (b) DKTF, (c) HKT, (d) AKT, (e) GKT, (f) QIKT, and (g) PSI-KT. The x-axis represents the normalized edge weights inferred by the respective baselines. The coefficient (coef) and its p-value are reported in the lower right of each panel.




Figure 13: Mean accuracy of PSI-KT vs. ablations of a) the prerequisite structure (w/o graph), b) individualized learner traits (w/o individual) and c) time-dependent learner traits (w/o dynamics). Dashed lines indicate the accuracy of the two best-performing baselines.