

PACE: A Personalized Adaptive Curriculum Engine for 9-1-1 Call-taker Training

Zirong Chen, Hongchao Zhang, Meiyi Ma

College of Connected Computing, Vanderbilt University, Nashville, Tennessee 37235, USA

{zirong.chen, hongchao.zhang, meiyi.ma}@vanderbilt.edu

Abstract

9-1-1 call-taking training requires mastery of over a thousand interdependent skills, covering diverse incident types and protocol-specific nuances. A nationwide labor shortage is already straining training capacity, but effective instruction still demands that trainers tailor objectives to each trainee’s evolving competencies. This personalization burden is one that current practice cannot scale. Partnering with Metro Nashville Department of Emergency Communications (MNDEC), we propose PACE (Personalized Adaptive Curriculum Engine), a co-pilot system that augments trainer decision-making by (1) maintaining probabilistic beliefs over trainee skill states, (2) modeling individual learning and forgetting dynamics, and (3) recommending training scenarios that balance acquisition of new competencies with retention of existing ones. PACE propagates evidence over a structured skill graph to accelerate diagnostic coverage and applies contextual bandits to select scenarios that target gaps the trainee is prepared to address. Empirical results show that PACE achieves 19.50% faster time-to-competence and 10.95% higher terminal mastery compared to state-of-the-art frameworks. Co-pilot studies with practicing training officers further demonstrate a 95.45% alignment rate between PACE’s and experts’ pedagogical judgments on real-world cases. Under estimation, PACE cuts turnaround time to merely 34 seconds from 11.58 minutes, up to 95.08% reduction.

1 Introduction

9-1-1 call-takers are first points of contact in life-threatening situations, and their guidance makes differences between life and death. Training them for this role is notoriously difficult [Chen *et al.*, 2025b]. A car crash, for instance, may begin as a routine law enforcement matter requiring only traffic blockage and vehicle description inquiries, but escalates to paramedic involvement if injuries are reported; further if unknown fluid leakage is present, fire-related safety checks become necessary as the fluid might be flammable or even explosive. Traditional training relies on human trainers who

review past performance and select subsequent training objectives accordingly [Chen *et al.*, 2025a]. As trainee cohorts grow and experienced trainers remain scarce, this model struggles to scale. In addition, trainees also differ in how they learn: some acquire skills quickly but forget without reinforcement, while others need repeated exposure but retain knowledge longer [Radvansky *et al.*, 2022]. Tailoring instruction to these differences improves outcomes substantially, but adds burden for trainers already monitoring progress across many trainees. As a result, most programs default to uniform curricula, overlooking individual learning patterns that could accelerate or hinder competency development.

Despite recent advancements in personalized learning systems, we identify following **challenges**: (1) *Engagement-Learning Discrepancy*. Educational recommendation systems often inherit objectives from information retrieval and e-commerce, such as maximizing click-through rates or session duration [Covington *et al.*, 2016; Koren *et al.*, 2009], which do not necessarily align with learning goals. While recent work incorporates learning signals [Verbert *et al.*, 2012], fundamental limitations remain. Trainees may “prefer” scenarios they can complete easily, but skill acquisition requires practice near the edge of competence. In consequences, call-takers may perform well in training simulations yet remain unprepared for high-stress real-world calls involving unfamiliar incident combinations. (2) *Inductive Knowledge Tracing*. Most knowledge tracing methods [Piech *et al.*, 2015; Ghosh *et al.*, 2020] estimate mastery from observed performance but typically treat skills as independent. In 9-1-1 call-taker training, over 1,000 distinct skills are involved, making direct assessment of each one infeasible within realistic training horizons. Inductively inferring mastery across structurally related skills is therefore critical for scalable training. (3) *Fine-Grained Tutoring*. Existing personalized tutoring systems often operate at course, module, or topic granularity [Verbert *et al.*, 2012], lumping together heterogeneous skills within broad categories. A trainee labeled “competent in medical calls” may handle routine chest pain assessments but struggle with overdose intervention or pediatric emergencies. In 9-1-1 operations, such granularity is dangerous. Because dispatching a call-taker based on coarse competence can delay or mis-specify instructions, with potentially fatal consequences. (4) *Persistent Learner Modeling*. LLM-based tutoring systems show promise for educa-

tional dialogue [Wang *et al.*, 2025; Gao *et al.*, 2025] but struggle with longitudinal training [Fatemi *et al.*, 2025]. A typical 9-1-1 program spans 6 to 8 weeks with over 200 simulated calls per trainee; encoding complete interaction histories exceeds practical context limits [Vodrahalli *et al.*, 2024]. Retrieval quality also degrades when relevant information is buried mid-context [Liu *et al.*, 2024a]. Without persistent temporal modeling, curricula fail to schedule retention practice, and previously mastered skills can erode unnoticed until a failure surfaces weeks after training.

Our **technical innovations and contributions** are: (1) We formalize 9-1-1 call-taker training as curriculum optimization over a structured graph. (2) We propose *PACE*¹, first-of-its-kind curriculum optimization system designed for 9-1-1 call-taker training that maintains probabilistic beliefs over trainee competence, estimates individual learning dynamics, and selects training scenarios based on defined tutoring context with bandits. (3) We comprehensively evaluate *PACE*’s performance against state-of-the-art frameworks in terms of time-to-competence and terminal mastery. (4) We conduct co-pilot studies with experienced training officers at a partnering 9-1-1 center, evaluating *PACE*’s alignment with expert pedagogical judgment under real-world cases.

Besides technical innovations, *PACE* delivers following **social impacts**: (1) As a co-pilot, *PACE*’s scenario recommendations achieve 95.45% agreement with experienced training officers’ selections, validating that the learned policy captures domain-relevant pedagogical principles rather than exploiting simulation artifacts. (2) Under estimation, *PACE* cuts 11.58 minutes turnaround time by 95.08% to solely 34 seconds, saving up to 11.01 minutes per call and at least 4,756.32 minutes per typical training class.

Practitioner Collaboration. *PACE* is developed in partnership with local 9-1-1 call center serving a mid-size U.S. city. Domain experts from training, quality assurance, and operations contributed throughout: validating motivating observations against national and local contexts, co-specifying the skill graph structure and annotation schema, reviewing *PACE* design choices, and participating in the co-pilot evaluation.

2 Motivating Study

To ground our motivation in operational reality, we analyzed the call-taking training manual provided by local 9-1-1 call center, along with 923 training session logs comprising conversational transcripts, evaluation rubrics, and trainer feedback annotations. We report two key observations.

Individualized feedback faces bandwidth constraints. Effective training relies on tight feedback loops in which trainees complete a simulated call, receive targeted debriefing, and apply corrections in subsequent attempts. Our analysis shows that the average turnaround time for a single simulation-debriefing cycle is 11.58 minutes, including call review, protocol cross-referencing, error diagnosis, and feedback delivery. With typical cohort structures assigning one trainer to approximately 12 trainees, each completing at least 3 sessions with at least 12 calls per session, full coverage would require over 83 hours of review time per day, which

is physically infeasible. This limitation is not specific to any single agency but reflects a structural constraint across the field. The opportunity is clear: assistance that accelerates assessment and surfaces priority issues can help trainers focus their limited time where it is most impactful.

Localized skill gaps cascade into systemic failures. We examine how deficiencies in individual skills propagate to overall call-handling performance. Using the procedural manual’s 857 discrete assessment checkpoints organized hierarchically, we generate 500 synthetic call instances with randomized applicable checkpoints and simulate the removal of individual skills. Removing a single skill orphaned, on average, 11.17% of downstream checkpoints (96 of 857) and caused 48.80% of complete call evaluations to fail (244 of 500 instances). For example, assessing patient consciousness is a foundational checkpoint shared across cardiac arrest, choking, drowning, overdose, and trauma protocols. This skill gates downstream decisions such as whether to initiate CPR, select between the Heimlich maneuver and recovery position, or assign medical dispatch priority. A trainee who inadequately masters consciousness assessment therefore underperforms across all dependent incident types rather than failing a single protocol. Because training typically progresses through incident types sequentially, such gaps may remain undetected for weeks and compound with forgetting. This interdependent structure motivates curriculum systems that track skills at finer granularity, allowing mastery in one context to inform readiness in others.

3 Problem Formulation

We formalize the curriculum design problem as decision-making over a structured skill graph.

3.1 Knowledge Graph Representation

Working with on-duty 9-1-1 call-takers and training officers, we annotate and decompose call-taking procedures into fine-grained, assessable competencies and their dependency structure. We encode this structure as a directed knowledge graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The node set $\mathcal{V} = \mathcal{V}_C \cup \mathcal{V}_Q \cup \mathcal{V}_I$ comprises *condition* nodes $\mathcal{V}_C$ (incident-state premises), *question* nodes $\mathcal{V}_Q$ (information-gathering skills), and *instruction* nodes $\mathcal{V}_I$ (directive-delivery skills). The edge set $\mathcal{E}$ captures: (1) sequential dependencies among questions/instructions (procedural order and prerequisites), (2) implication links from questions/instructions to conditions they establish, and (3) entailment links from conditions to downstream questions/instructions they activate or mandate. The resulting graph contains $|\mathcal{V}| = 1,053$ nodes and $|\mathcal{E}| = 1,283$ edges spanning 63 incident types.

3.2 Teaching under Approximation

Human-centric pedagogical frameworks involve inherently unobservable learner states [Leighton *et al.*, 2004; De La Torre, 2009]. In high-stakes domains such as emergency call-taking, two trainees may produce identical responses while possessing fundamentally different underlying competencies: one acts from procedural mastery, another guesses correctly under pressure. Trainers cannot directly inspect these latent states and must rely on observable behavior

¹Code and details: <https://github.com/AICPS-Lab/PACE-911>.

alone. Let $\Theta = \{\theta_v\}_{v \in \mathcal{V}}$ denote the trainee’s true competency state, where each $\theta_v \in [0, 1]$ represents mastery of skill v . Any curriculum system therefore maintains a belief model $\hat{\Theta} = \{\hat{\theta}_v\}_{v \in \mathcal{V}}$ that approximates Θ from observations.

Trainer’s Observations. After each simulation, structured debriefing produces observations over skills in the activated subgraph. Observations take values $o_v \in \{\top, \perp, \sim, \emptyset\}$, where $\top$ and $\perp$ indicate compliance or violation, $\sim$ indicates partial compliance, and $\emptyset$ denotes not applicable due to unmet prerequisites. These signals are noisy functions of latent mastery: higher θ_v increases the likelihood of compliance, but execution variability, stress, and context introduce observation noise. Observations accumulate in a history $\mathcal{H}$, from which beliefs $\hat{\Theta}$ are inferred.

Trainer’s Actions. Trainer selects simulations based on the current belief state $\hat{\Theta}$. Each scenario $\mathcal{S} = (\iota, \mathbf{c})$, defined by incident type $\iota \in \mathcal{I}$ (e.g., car crash) and condition configuration $\mathbf{c}$ (i.e., what nodes to activate or deactivate), activates a subgraph $V_{\mathcal{S}} \subseteq \mathcal{V}$. Training proceeds over N sessions. At each session t , the trainer selects a batch of K scenarios $\mathbf{S}_t = \{\mathcal{S}_t^{(1)}, \dots, \mathcal{S}_t^{(K)}\}$ informed by $\hat{\Theta}^{(t)}$. The curriculum design problem is to choose the sequence $\{\mathbf{S}_t\}_{t=1}^N$ that maximizes learning outcomes under approximate beliefs.

Objectives. Curriculum design balances coupled objectives. The *explicit objective* maximizes terminal competence:

$$\max_{\{\mathbf{S}_t\}} \mathbb{E} \left[\sum_{v \in \mathcal{V}} \theta_v^{(N)} \right] \quad (1)$$

The *implicit objective* minimizes the approximation gap $\delta = \frac{1}{N} \sum_i \|\mathbb{E}[\hat{\theta}_i] - \theta_i\|$, since scenario selection quality degrades as $\hat{\Theta}$ diverges from Θ [Tabibian *et al.*, 2019; Su *et al.*, 2023].

4 PACE for 9-1-1 Call-taking Training

In this section, we present *PACE* (Personalized Adaptive Curriculum Engine), first-of-its-kind system for adaptive curriculum optimization for 9-1-1 call-taking training.

4.1 System Overview

In operational settings, PACE operates as a curriculum layer within existing 9-1-1 training infrastructure that role-plays callers and conducts performance debriefing. PACE observes trainee interaction logs, including simulation instantiations, conversation transcripts, and debriefing results, and compares observed learning outcomes against target call-taking knowledge to optimize and personalize each trainee’s learning, in Figure 1. PACE consists of: (1) *Belief Tracker*, Section 4.4, that updates per-node Beta posteriors from transcripts and propagates evidence via similarity; (2) *Dynamics Estimator*, Section 4.5, that infers trainee-specific learning pace λ and forgetting rate ψ using a surrogate model; and (3) *Contextual Bandit*, Section 4.6, that selects scenario batches by balancing exploitation and exploration under the tutoring context.

4.2 Pre-runtime Preparation: Similarity Index

Skill transfer enables inference over unobserved nodes based on performance on related skills. We precompute a similar-




Figure 1: Overview of PACE’s involvement in 9-1-1 training.

ity index Φ offline. For each node v , we obtain a semantic embedding $\mathbf{s}_v \in \mathbb{R}^{384}$ by encoding its textual description using a pretrained sentence transformer [Reimers and Gurevych, 2019]. Pairwise similarity combines semantic relatedness with positional compatibility [Vaswani *et al.*, 2017]:

$$\phi(v_i, v_j) = \cos(\mathbf{s}_i, \mathbf{s}_j) \cdot \exp(-\epsilon|\mathbf{d}_i - \mathbf{d}_j|) \quad (2)$$

where $\mathbf{d}_v$ is normalized depth and ϵ controls position sensitivity. The exponential term reflects that skills at similar protocol stages (e.g., both early information-gathering) transfer more readily than skills at different stages (e.g., initial assessment vs. final instructions). Following standard practice in embedding-based retrieval [Li *et al.*, 2025], we cache only pairs exceeding a preset threshold to maintain sparsity.

4.3 PACE Runtime

The complete PACE runtime proceeds in two phases: (1) *Cold Start Phase*: During $t = 1, \dots, N_0$, PACE silently observes trainee performance on externally-assigned scenarios without making recommendations. This phase accumulates observations to initialize beliefs and estimate learning dynamics. We set $N_0 = 15$ based on local training logistics; and (2) *Adaptive Phase*, see Figure 2: From session $N_0 + 1$ onward, PACE actively selects batches with contextual bandit after observing learning outcomes and estimating learning dynamics.

4.4 Observing the Training Outcomes

The Belief Tracker maintains $\hat{\Theta}$ through observation extraction and belief updates.

Observation Extraction. An LLM extracts structured observations from call simulation transcripts [He *et al.*, 2024]. For each node $v \in V_{\mathcal{S}}$, the system outputs: outcome $o_v \in \{\top, \perp, \sim, \emptyset\}$; error type $e_v \in \{\text{SLIP}, \text{MISCONCEPTION}, \text{OMISSION}\}$ [Liu *et al.*, 2023]; and a prompted flag indicating whether the caller prompted the action versus independent performance.

Belief Updates. For each node v , we maintain a Beta posterior $\text{Beta}(\alpha_v, \beta_v)$ representing uncertainty over mastery θ_v [Shen *et al.*, 2024]. The posterior mean $\mu_v = \alpha_v / (\alpha_v + \beta_v)$ estimates mastery, while the variance σ_v^2 captures epistemic uncertainty. Given an observation $o_v \in \{\top, \perp, \sim\}$, the




Figure 2: PACE runtime under adaptive phase after cold start.

posterior is updated using weighted pseudo-counts: $\alpha_v \leftarrow \alpha_v + w^+ \mathbf{1}\{o_v \in \{\top, \sim\}\}$ and $\beta_v \leftarrow \beta_v + w^- \mathbf{1}\{o_v \in \{\perp, \sim\}\}$, where $\mathbf{1}\{\cdot\}$ denotes the indicator function and $w^+, w^- > 0$ control the strength of positive and negative evidence. Weights are modulated by error type [Pelánek, 2017]: prompted successes receive $w^+ = 0.5$; misconceptions receive $w^- = 1.5$; slips receive $w^- = 0.5$.

Inter-Node Mastery Propagation. When node v is directly observed with given mastery, we inductively update evidence to similar nodes: $\hat{\theta}_{v'}^{\text{new}} \leftarrow \text{mean}(\hat{\theta}_{v'}^{\text{curr}}, \phi(v, v') \cdot \hat{\theta}_v)$.

4.5 Estimating the Learning Dynamics

Besides widely used practice and cognitive factors [Chen *et al.*, 2023; Cheng *et al.*, 2024], the Dynamics Estimator captures individual differences with a surrogation [Tabibian *et al.*, 2019; Ma *et al.*, 2023] focusing on learning (denoted with λ) and forgetting dynamics (denoted with ψ).

Learning Pace λ . We operationalize λ as the average marginal mastery gain per practice opportunity over the trainee’s interaction history $\mathcal{H}$:

$$\lambda = \frac{1}{|\mathcal{H}|} \sum_{(v,t) \in \mathcal{H}} \Delta \theta_v^{(t)} \quad (3)$$

where $\Delta \theta_v^{(t)} := \theta_v^{(t)} - \theta_v^{(t-1)}$ denotes mastery updates for skill v at practice session t . Trainees with higher λ reliably consolidate new skills with fewer repetitions, enabling faster progression to complex or compositional scenarios.

Forgetting Model ψ . Skills decay over real time τ following a power law, one dominant model in cognitive science [Su *et al.*, 2023]. Let $\Delta \tau_v$ denote the elapsed clocked time since skill v was last practiced. We model mastery decay as:

$$\theta_v^{(\tau + \Delta \tau_v)} = \theta_v^{(\tau)} \cdot (1 + \kappa \cdot \Delta \tau_v)^{-\psi} \quad (4)$$

where $\psi > 0$ is a trainee-specific forgetting rate and $\kappa > 0$ is a constant that normalizes elapsed time into the effective scale at which forgetting accumulates. The power-law structure captures the empirically observed phenomenon that forgetting slows over time: recently acquired skills decay more

rapidly than skills that have undergone consolidation. Modeling decay over real time rather than session count is operationally important in training settings. We estimate ψ by fitting Eq. (4) to observed retention outcomes when previously mastered skills are reassessed after variable temporal gaps.

4.6 Customizing the Next Training Batch

PACE formulates scenario selection as a contextual bandit problem [Belfer *et al.*, 2022; Nguyen *et al.*, 2024], where a policy maps trainee-specific contexts to scenario choices that maximize expected learning gain under uncertainty.

Context Vector. At session t , we construct a compact context vector that summarizes the current training state:

$$\mathbf{x}_t = [\underbrace{\bar{\sigma}^2, c}_{\text{model quality}}, \underbrace{\hat{\lambda}, \hat{\psi}, \bar{w}, d}_{\text{trainee state}}, \underbrace{t/N}_{\text{progress}}] \quad (5)$$

where $\bar{\sigma}^2$ is the mean belief variance over skill mastery (direct observation of Θ is not available under deployment, we use $\bar{\sigma}^2$ as one quality indicator of the surrogate model from trainer’s side), c is current coverage, $\hat{\lambda}$ is estimated learning pace, $\hat{\psi}$ is estimated forgetting rate, $\bar{w} \in \mathbb{R}$ is the average mastery of weak skills (under preset threshold), $d \in \mathbb{R}$ is the number of skills near the forgetting threshold, and t/N denotes normalized training progress over N total sessions.

Action Space. Each action corresponds to a scenario $S = (\iota, \mathbf{c})$, defined by an incident type ι and contextual configuration $\mathbf{c}$. 297 candidate scenarios are filtered to satisfy prerequisite, operational validity, and complexity constraints. PACE selects a batch of $K = 5$ scenarios sequentially, updating the context after each selection to reflect newly covered skills.

Reward. The reward reflects observed learning gain induced by the selected scenario after applying real time decay:

$$r_t = \sum_{v \in V_{S_t}} (\hat{\theta}_v^{(t)} - \hat{\theta}_v^{(t-1)}) \quad (6)$$

where V_{S_t} is the activated subgraph by S_t . Measures of model uncertainty appear in the context rather than in the reward,

allowing the bandit to learn when exploratory actions that reduce uncertainty lead to higher downstream gains.

Algorithm. We adopt Thompson Sampling [Agrawal and Goyal, 2013], maintaining posterior distributions over expected rewards and sampling from these posteriors to select actions. This approach naturally balances exploration and exploitation under uncertainty. We model the expected reward as $\mathbb{E}[r \mid \mathbf{x}, \mathcal{S}] = \mathbf{B}_{\mathcal{S}}^{\top} \mathbf{x}$, where $\mathbf{B}_{\mathcal{S}}$ denotes scenario-specific coefficients with a Gaussian prior. Parameters are updated via Bayesian linear regression after observing rewards.

5 Evaluation

We evaluate PACE via (1) *system-level comparisons* with ablations, (2) *within-PACE analysis* across tutoring granularities and behavioral patterns, and (3) a *co-pilot study* with domain experts under operational conditions. All experiments run on a machine with 128GB RAM, an AMD Ryzen Threadripper Pro 7975WX CPU, and an NVIDIA RTX 6000 Ada GPU.

5.1 Role-played Trainee Agents

Although this study has received IRB approval, *direct evaluation with human trainees is not ideal for controlled curriculum analysis* due to: (1) training programs span multiple weeks, limiting rapid iteration; (2) trainee competency states are latent and unobservable, precluding reliable ground-truth labels; and (3) curriculum refinement requires counterfactual comparisons that are infeasible because instructional exposure irreversibly alters knowledge states. Following recent work on LLM-based educational simulation [Zhang *et al.*, 2025; Liu *et al.*, 2024b], we construct role-played trainee agents with controllable ground-truth parameters. Each agent $\mathcal{A} = (\Theta, \lambda, \psi, \mathcal{P}, \mathcal{M})$ consists of a ground-truth competency state $\Theta = \{\theta_v\}_{v \in \mathcal{V}}$, where $\theta_v \in [0, 1]$ denotes mastery of skill v , learning and forgetting rates $(\lambda, \psi) \in \mathbb{R}^2$, a behavioral persona $\mathcal{P}$ encoding response patterns, and a scratchpad $\mathcal{M}$ that maintains learning history across sessions [Park *et al.*, 2023]. Given a training scenario $\mathcal{S}$ that activates a skill subset $V_{\mathcal{S}}$, the agent generates utterances by sampling: $u \sim \text{LLM}(\mathcal{P}, \mathcal{S}, \mathcal{M}, \{(\theta_v, e_v) : v \in V_{\mathcal{S}}\})$, where e_v denotes an error type sampled with probability $1 - \theta_v$. To prevent knowledge leakage, prompts explicitly constrain the agent to act as a trainee with specified skill levels [Zhao *et al.*, 2025; Wang *et al.*, 2024]. Ecological validity is further improved by incorporating error templates extracted from anonymized training logs [Dong *et al.*, 2024]. We define four trainee archetypes parameterized by (λ, ψ) : *Fast Learner* (0.12, 0.15) with rapid acquisition and slow decay; *Moderate Learner* (0.07, 0.25) reflecting typical progression; *Struggling Learner* (0.03, 0.35) with slow acquisition and rapid decay; and *Quick Forgetter* (0.10, 0.45) with strong initial acquisition but rapid decay. The parameters (λ, ψ) are hidden from PACE during evaluation, and only conversational observations are available. Each archetype includes 10 instantiations with $\pm 15\%$ parameter noise. We use OpenAI GPT-5 for PACE and all re-implemented agentic baselines. To avoid in-family comparisons, trainee agents are instantiated using Claude Sonnet 4.5.

5.2 System-level Comparison with Ablation

Goal. This experiment configuration studies *does PACE help trainee learn effectively and efficiently?*

Baselines. We compare against: (1) *Round-Robin*, sequentially cycling through incident types; (2) *Deficit-Driven*, deterministically selecting scenarios targeting lowest-mastery skills based on debriefing; (3) *GraphRAG+LLM* [Edge *et al.*, 2024], which we selected as the strongest vanilla LLM integration after pilot trials, using graph-guided retrieval and traversal to contextualize deficits prior to selection; and (4) agentic pipelines² including GenMentor [Wang *et al.*, 2025] and Agent4Edu [Gao *et al.*, 2025].

Metrics. (1) *Coverage@t* measures the mastery cardinality fraction by session: $\text{Coverage@t} = |\{v : \theta_v^{(t)} \geq 0.85\}|/|\mathcal{V}|$; (2) *Zero-to-Hero (Z2H)* is the number of sessions required to first reach episode-level competence (scenario score ≥ 0.85), capturing learning efficiency; (3) *Random Exam (RE)* reports performance (0–100) on a held-out exam uniformly sampled across incident types, simulating certification-style slot-filling assessment used in local training practice. For each method, we train 40 simulated trainees for 50 sessions with batch size $K = 5$. Every 5 sessions, we insert a 24-hour gap to reflect realistic training schedules.

Ablation variants. (1) *PACE w/o propagation* removes similarity-based belief transfer over the skill graph; (2) *PACE w/o dynamics* replaces trainee-specific parameters (λ_i, ψ_i) with fixed population averages $(\bar{\lambda} = 0.08, \bar{\psi} = 0.30)$.

Analysis. From Table 1, on Fast Learners, PACE reaches Z2H of 22.19 compared to Agent4Edu at 27.58, corresponding to 19.50% reduction, while achieving higher coverage (C@50: 95.27% vs. 88.18%). The performance gap widens for more challenging archetypes. On Quick Forgetter ($\psi = 0.45$), PACE attains C@50 of 91.22%, whereas GenMentor degrades from 71.66% at session 10 to 67.51% at session 50, indicating negative learning progress. This regression reflects GenMentor struggles to handle high forgetting rates without explicit temporal modeling. Agent4Edu maintains positive progress but plateaus below PACE (C@50: 86.74% vs. 91.22%), as its static learner memory does not adapt to inter-session decay. Ablation results confirm complementary contributions of the two components. Removing propagation substantially increases Z2H across archetypes (e.g., Fast Learner: 46.13 vs. 22.19), since the system must directly observe each skill rather than inferring mastery via graph-based transfer. Removing dynamics estimation disproportionately harms archetypes that deviate from population norms. For Quick Forgetter, RE drops from 91.04 to 80.16 ($\Delta 10.88$) because the population-average $\bar{\psi} = 0.30$ underestimates the true $\psi = 0.45$, leading to insufficient retention scheduling. With both components enabled, propagation accelerates belief coverage while dynamics estimation calibrates beliefs to individual learning trajectories. Modeling learning as a temporal process rather than a static profile is therefore essential under realistic session gaps. In summary, *with all design choices empirically validated, PACE outperforms baseline training approaches in both effectiveness and efficiency.*

²All agentic baselines are re-implemented from published descriptions to operate on the 9-1-1 call-taking data used in this study.

Group	Method	Fast Learner ($\lambda=0.12, \psi=0.15$)					Moderate Learner ($\lambda=0.07, \psi=0.25$)				
		C@10 $\uparrow$	C@30 $\uparrow$	C@50 $\uparrow$	Z2H $\downarrow$	RE $\uparrow$	C@10 $\uparrow$	C@30 $\uparrow$	C@50 $\uparrow$	Z2H $\downarrow$	RE $\uparrow$
Heuristics	Round-Robin	42.24 $\pm$ 7.16	53.43 $\pm$ 5.57	66.11 $\pm$ 7.37	/	63.14 $\pm$ 4.11	34.19 $\pm$ 8.21	47.21 $\pm$ 7.20	55.15 $\pm$ 8.16	/	51.43 $\pm$ 6.44
	Deficit-Driven	57.40 $\pm$ 4.14	55.22 $\pm$ 6.84	56.67 $\pm$ 8.17	/	59.46 $\pm$ 6.41	49.50 $\pm$ 5.15	50.11 $\pm$ 7.13	48.00 $\pm$ 8.77	/	45.47 $\pm$ 7.17
LLM-based	GraphRAG	59.56 $\pm$ 7.72	60.50 $\pm$ 8.18	61.44 $\pm$ 9.31	/	60.22 $\pm$ 9.21	51.43 $\pm$ 7.82	48.91 $\pm$ 9.22	50.46 $\pm$ 8.91	/	49.84 $\pm$ 9.64
	GenMentor	65.67 $\pm$ 9.55	79.41 $\pm$ 8.33	84.15 $\pm$ 8.08	29.40 $\pm$ 6.60	86.87 $\pm$ 6.71	66.14 $\pm$ 10.11	73.34 $\pm$ 8.78	75.44 $\pm$ 6.41	/	78.12 $\pm$ 8.83
	Agent4Edu	68.14 $\pm$ 8.27	82.06 $\pm$ 6.58	88.18 $\pm$ 5.57	27.58 $\pm$ 5.24	90.47 $\pm$ 5.12	62.41 $\pm$ 8.10	72.59 $\pm$ 9.66	84.57 $\pm$ 6.18	48.10 $\pm$ 1.90	86.17 $\pm$ 5.21
PACE	w/o prop	66.74 $\pm$ 7.40	81.43 $\pm$ 6.18	84.73 $\pm$ 5.19	46.13 $\pm$ 3.87	86.95 $\pm$ 5.00	59.47 $\pm$ 8.20	70.25 $\pm$ 6.17	83.46 $\pm$ 5.11	49.15 $\pm$ 0.85	84.71 $\pm$ 6.83
	w/o dyn	72.93$\pm$6.11	88.71 $\pm$ 4.23	92.00 $\pm$ 4.14	34.55 $\pm$ 5.80	91.13 $\pm$ 4.06	60.15 $\pm$ 7.12	70.13 $\pm$ 8.10	84.10 $\pm$ 6.17	46.10 $\pm$ 3.90	85.03 $\pm$ 7.23
	full	72.19 $\pm$ 5.31	92.71$\pm$3.65	95.27$\pm$3.23	22.19$\pm$4.11	96.05$\pm$2.54	67.55$\pm$9.11	77.54$\pm$5.10	91.51$\pm$5.19	34.10$\pm$5.52	90.11$\pm$4.51
Group	Method	Struggling Learner ($\lambda=0.03, \psi=0.35$)					Quick Forgetter ($\lambda=0.10, \psi=0.45$)				
		C@10 $\uparrow$	C@30 $\uparrow$	C@50 $\uparrow$	Z2H $\downarrow$	RE $\uparrow$	C@10 $\uparrow$	C@30 $\uparrow$	C@50 $\uparrow$	Z2H $\downarrow$	RE $\uparrow$
Heuristics	Round-Robin	30.41 $\pm$ 5.22	44.27 $\pm$ 8.21	50.10 $\pm$ 9.81	/	53.44 $\pm$ 8.14	47.11 $\pm$ 8.93	50.12 $\pm$ 7.29	51.38 $\pm$ 7.32	/	48.19 $\pm$ 7.20
	Deficit-Driven	37.17 $\pm$ 4.75	35.20 $\pm$ 8.91	36.44 $\pm$ 7.33	/	33.73 $\pm$ 8.13	55.17 $\pm$ 5.87	50.25 $\pm$ 8.75	52.61 $\pm$ 9.51	/	51.17 $\pm$ 8.11
LLM-based	GraphRAG	35.71 $\pm$ 10.22	34.56 $\pm$ 14.80	37.07 $\pm$ 9.88	/	39.60 $\pm$ 11.36	61.08 $\pm$ 9.10	57.87 $\pm$ 12.74	58.72 $\pm$ 14.15	/	59.83 $\pm$ 10.67
	GenMentor	44.14 $\pm$ 8.81	51.68 $\pm$ 9.37	66.17 $\pm$ 10.15	/	65.73 $\pm$ 7.45	71.66$\pm$11.37	69.43 $\pm$ 9.85	67.51 $\pm$ 11.27	/	66.37 $\pm$ 9.11
	Agent4Edu	57.15$\pm$8.44	71.29 $\pm$ 8.73	80.48 $\pm$ 8.13	47.88 $\pm$ 2.12	81.15 $\pm$ 9.19	72.54 $\pm$ 8.82	77.14 $\pm$ 8.19	86.74 $\pm$ 5.51	32.55 $\pm$ 6.55	85.78 $\pm$ 6.71
PACE	w/o prop	47.17 $\pm$ 8.33	71.41 $\pm$ 7.12	80.47 $\pm$ 5.71	49.44 $\pm$ 0.56	78.83 $\pm$ 6.22	66.17 $\pm$ 8.05	76.85 $\pm$ 6.10	84.67 $\pm$ 4.11	49.50 $\pm$ 0.50	86.13 $\pm$ 4.71
	w/o dyn	45.15 $\pm$ 10.48	68.98 $\pm$ 9.73	80.19 $\pm$ 8.65	47.13 $\pm$ 2.87	82.17 $\pm$ 9.81	69.69 $\pm$ 9.14	73.43 $\pm$ 8.72	82.09 $\pm$ 6.88	48.13 $\pm$ 1.87	80.16 $\pm$ 8.85
	full	55.50 $\pm$ 9.00	77.54$\pm$8.18	86.91$\pm$5.15	44.14$\pm$5.86	88.18$\pm$5.20	71.45 $\pm$ 6.98	87.43$\pm$7.67	91.22$\pm$4.17	25.91$\pm$3.49	91.04$\pm$5.11

Table 1: System-level comparison with PACE ablations. C@ t (%) denotes coverage at session t . Z2H (zero-to-hero) denotes sessions to first competent episode. Slash means never observed, excluded for average calculation. RE (random exams) denotes held-out exam score.

5.3 Within-PACE Analysis

We observe and report PACE’s performance on different tutoring *granularity levels* and its overall *behavioral patterns*.

Granularity Analysis of PACE

Goal. This section studies *whether fine-grained belief tracking in PACE improves learning outcomes*. While prior work and motivating studies suggest benefits of finer-granularity tutoring, we seek empirical evidence under 9-1-1 training.

Granularity Levels. We compare three variants: (1) *PACE-coarse*, which tracks beliefs at the departmental level (police, fire, medical; 3 categories); (2) *PACE-medium*, which tracks beliefs at the incident-type level (e.g., car crash, structure fire; 63 categories); and (3) *PACE-fine*, which tracks beliefs at the skill level (1,053 nodes) and corresponds to the full PACE system used in prior experiments. At coarser granularities, beliefs are aggregated and scenarios are selected based on aggregate deficits. All variants are evaluated on the same four trainee archetypes over 50 sessions.

Analysis. Figure 3 shows consistent advantages for fine-grained belief tracking across all trainee archetypes. PACE-fine achieves terminal coverage (C@50) of 95.27%, 91.51%, 86.91%, and 91.22% for the four archetypes, respectively, compared to 68.49%, 56.11%, 50.93%, and 59.67% for PACE-medium, with PACE-coarse performing worst overall. PACE-medium exhibits intermediate performance but with higher variance, particularly for the Struggling and Quick Forgetter archetypes. This instability arises because incident-type granularity conflates heterogeneous skills within each category. For example, a trainee may perform well on basic choking response yet struggle with choking complications, which type-level tracking cannot distinguish. Overall, these results show that *fine-grained belief tracking in PACE is not merely a design choice but a requirement for reliable curriculum optimization in high-stakes 9-1-1 training*.

Behavioral Analysis of PACE

Goal. This section studies *does PACE correctly approximate trainee learning states and take expected actions on the fly?*




Figure 3: Coverage progression under different belief granularities.

Surfaced PACE Actions. We compare PACE’s estimated trainee mastery $\hat{\Theta}$, represented as Beta distributions over 1,053 skill nodes, against the trainee agent’s ground-truth mastery $\Theta \in [0, 1]^{|V|}$. We quantify the approximation gap using δ , defined in Section 3.2, which measures the mean absolute deviation between believed and true mastery. We also report PACE’s internal belief quality indicator $\bar{\sigma}^2 = \frac{1}{N} \sum_i \text{Var}(\hat{\theta}_i)$, which directly influences scenario selection. Each scatter point corresponds to one of 10 independent runs per archetype over 50 sessions. Point color indicates the explore to exploit ratio within the selected batch for that session.

Analysis. Figure 4 reveals two key patterns: (1) the approximation gap δ decreases across all trainee archetypes, indicating that PACE’s belief estimates converge toward ground truth. Convergence speed varies with trainee complexity: Fast Learners stabilize around session 20, Moderate Learners around session 30, and Struggling Learners around ses-




Figure 4: Convergence of approximation gap δ (scatter points, left axis) and belief variance σ^2 (gray line $\pm$ std, right axis) across sessions. Point color indicates explore:exploit ratio within each batch ($K=5$): green denotes exploration-heavy sessions, red denotes exploitation-heavy sessions.

sion 40; and (2) the explore to exploit balance adapts to belief quality. Early sessions exhibit predominantly red and orange points, corresponding to higher σ^2 values. As beliefs consolidate and σ^2 decreases, points shift toward green, indicating increased exploration. These trends show that PACE translates estimation uncertainty into appropriate action selection. When beliefs are unreliable, PACE prioritizes exploitation to calibrate suspected weaknesses. As belief confidence increases, PACE shifts toward exploration to expand coverage over uncertain regions of the skill graph. Overall, *PACE balances skill acquisition and belief calibration through uncertainty-aware curriculum control.*

5.4 Case Study: PACE as a Co-Pilot

Goal. This section studies *how well does PACE align with expert judgment and assist real-world training?*

Setup. We evaluate PACE under operational conditions at our partnering 9-1-1 call center using both quantitative alignment analysis and qualitative expert assessment. For quantitative evaluation, we collected 923 training catalog entries from recent training classes, each containing simulation logs, debriefing results, and trainer-provided comments with next-session assignments based on trainee performance. We executed PACE on this dataset and measured alignment between PACE’s curriculum recommendations and expert pedagogical decisions. For qualitative evaluation, we administered an anonymized survey to 21 domain experts, including 12 training officers and 9 active call-takers with training experience. Participants rated agreement with PACE’s design choices, such as skill-level diagnostic granularity, and overall helpfulness for call-taker training on a 5-point Likert scale from 1 (strongly disagree) to 5 (strongly agree).

Analysis. PACE achieves 95.45% alignment (881/923) with expert judgment, validating its effectiveness as a training co-pilot under practical usage. Moreover, PACE reduces average turnaround time from 11.58 minutes to 34 seconds (adaptive phase average), saving 95.08% of training officer time. Sur-

vey results further support PACE’s design and utility, with participants rating agreement with design choices at 4.62 and overall helpfulness at 4.43. Open-ended responses highlight two recurring themes. First, participants emphasized the value of fine-grained diagnosis: *“Knowing exactly which part of the protocol they missed matters more than just saying they struggled with medical calls.”* Second, participants noted reduced cognitive burden: *“It would save me from being overwhelmed by the entire training class.”* Overall, the co-pilot study confirms that *PACE’s recommendations align with expert pedagogical judgment in real-world settings while reducing training turnaround time by up to 95%.*

6 Related Work

High-Stakes and Emergency Response Training. High-stakes training requires deliberate practice near competence boundaries with tight feedback loops. In emergency response, dispatcher performance directly affects survival outcomes [Chen *et al.*, 2026], motivating simulation-based training. Recent work applies LLM-based simulation to 9-1-1 dispatcher training [Chen *et al.*, 2025b], emphasizing realistic caller role-play rather than pedagogical decision-making. **AI for Edu and Educational AI.** Knowledge tracing models learner states from interaction data [Shen *et al.*, 2024], with recent work incorporating graph structure to capture skill dependencies [Cheng *et al.*, 2024]. But these approaches primarily support diagnosis and do not implement integrated diagnose-prescribe tutoring loops. Adaptive learning systems attempt to close this gap using contextual bandits, spaced repetition, and educational recommendation [Belfer *et al.*, 2022; Su *et al.*, 2023], which correspondingly assume flat skill spaces, over-focus on retaining acquired items, and prioritize engagement over learning. Recent LLM-based tutoring agents [Wang *et al.*, 2025; Gao *et al.*, 2025] enable flexible dialogue and learner profiling, but suffers from context rot or overflow as training history enriches.

7 Summary

This paper introduces PACE, first-of-its-kind co-pilot for adaptive 9-1-1 call-taker training. By integrating inductive belief tracking over structured skill graphs, individual learning dynamics estimation, and contextual bandit-based scenario selection, PACE enables personalized curriculum optimization that accounts for both skill interdependencies and trainee-specific learning trajectories. Experimental results show improvements in time-to-competence and terminal mastery over state-of-the-art baselines. Co-pilot studies with practicing training officers further confirm alignment with expert pedagogical judgment in operational settings.

Locally verified with a 9-1-1 call-taking agency, PACE supports 9-1-1 call centers facing staffing and time constraints by surfacing diagnostic insights and curriculum recommendations that would otherwise require substantial manual effort. With over 6,000 9-1-1 call centers in the US, it potentially scales to broad deployment. As future work, beyond emergency response, PACE framework can extend to other high-stakes procedural training domains, including medical triage, air traffic control, and crisis intervention.

Acknowledgments

This work was supported in part by the U.S. National Science Foundation under Grants 2427711 and 2443803, the Google Academic Research Award, and the U.S. Department of Education under Grant R305C240010. The opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsoring agencies.

References

- [Agrawal and Goyal, 2013] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 127–135, 2013.
- [Belfer et al., 2022] Robert Belfer, Ekaterina Kochmar, and Iulian Vlad Serban. Raising student completion rates with adaptive curriculum and contextual bandits. In *Artificial Intelligence in Education (AIED)*, volume 13355 of *Lecture Notes in Computer Science*, pages 724–730. Springer, 2022.
- [Chen et al., 2023] Xiangzhi Chen, Le Wu, Fei Liu, Lei Chen, Kun Zhang, Richang Hong, and Meng Wang. Disentangling cognitive diagnosis with limited exercise labels. *Advances in Neural Information Processing Systems*, 36:18028–18045, 2023.
- [Chen et al., 2025a] Zirong Chen, Ziyang An, Jennifer Reynolds, Kristin Mullen, Stephen Maritini, and Meiyi Ma. Logidebrief: A signal-temporal logic based automated debriefing approach with large language models integration. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 9582–9590. International Joint Conferences on Artificial Intelligence Organization, 8 2025. AI and Social Good.
- [Chen et al., 2025b] Zirong Chen, Elizabeth Chason, Noah Mladenovski, Erin Wilson, Kristin Mullen, Stephen Maritini, and Meiyi Ma. Sim911: Towards effective and equitable 9-1-1 dispatcher training with an LLM-enabled simulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 27896–27904, 2025.
- [Chen et al., 2026] Zirong Chen, Yilin Liu, and Meiyi Ma. Empowering 9-1-1 calltaking training with generative ai: Experiences and lessons learned. *The 12th IEEE International Conference on Smart Computing*, 2026.
- [Cheng et al., 2024] Ke Cheng, Linzhi Peng, Pengyang Wang, Junchen Ye, Leilei Sun, and Bowen Du. DyGKT: Dynamic graph learning for knowledge tracing. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 385–396, 2024.
- [Covington et al., 2016] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys)*, pages 191–198, 2016.
- [De La Torre, 2009] Jimmy De La Torre. DINA model and parameter estimation: A didactic. *Journal of Educational and Behavioral Statistics*, 34(1):115–130, 2009.
- [Dong et al., 2024] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, et al. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, 2024.
- [Edge et al., 2024] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.
- [Fatemi et al., 2025] Bahare Fatemi, Mehran Kazemi, Anton Tsitsulin, Karishma Malkan, Jinyeong Yim, John Palowitch, Sungyong Seo, Jonathan Halcrow, and Bryan Perozzi. Test of time: A benchmark for evaluating LLMs on temporal reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Gao et al., 2025] Weibo Gao, Qi Liu, Linan Yue, Fangzhou Yao, Rui Lv, Zheng Zhang, Hao Wang, and Zhenya Huang. Agent4Edu: Generating learner response data by generative agents for intelligent education systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23923–23932, 2025.
- [Ghosh et al., 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 and Data Mining*, pages 2330–2339. ACM, 2020.
- [He et al., 2024] Xiaoxin He, Xavier Bresson, Thomas Laurent, Adam Perold, Yann LeCun, and Bryan Hooi. Harnessing explanations: LLM-to-LM interpreter for enhanced text-attributed graph representation learning. In *The Twelfth International Conference on Learning Representations*, 2024.
- [Koren et al., 2009] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [Leighton et al., 2004] Jacqueline P. Leighton, Mark J. Gierl, and Stephen M. Hunka. The attribute hierarchy method for cognitive assessment: A variation on tatsuoaka’s rule-space approach. *Journal of Educational Measurement*, 41(3):205–237, 2004.
- [Li et al., 2025] Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. From matching to generation: A survey on generative information retrieval. *ACM Transactions on Information Systems*, 43(3):1–62, 2025.
- [Liu et al., 2023] Zitao Liu, Jiahao Chen, and Weiqi Luo. Recent advances on deep learning based knowledge tracing. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 1295–1296, 2023.

- [Liu *et al.*, 2024a] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024.
- [Liu *et al.*, 2024b] Zhengyuan Liu, Stella Xin Yin, Geyu Lin, and Nancy F. Chen. Personality-aware student simulation for conversational intelligent tutoring systems. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2024.
- [Ma *et al.*, 2023] Boxuan Ma, Gayan P. Hettiarachchi, Shota Fukui, and Yutaka Ando. Each encounter counts: Modeling language learning and forgetting. In *LAK23: 13th International Learning Analytics and Knowledge Conference*, pages 79–88, 2023.
- [Nguyen *et al.*, 2024] Lam Minh Duy Nguyen, Fuhua Lin, and Maiga Chang. Generating learning sequences using contextual bandit algorithms. In *Intelligent Tutoring Systems (ITS)*, volume 14798 of *Lecture Notes in Computer Science*, pages 320–329. Springer, 2024.
- [Park *et al.*, 2023] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, pages 1–22. ACM, 2023.
- [Pelánek, 2017] Radek Pelánek. Bayesian knowledge tracing, logistic models, and beyond: An overview of learner modeling techniques. *User Modeling and User-Adapted Interaction*, 27(3-5):313–350, 2017.
- [Piech *et al.*, 2015] 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.
- [Radvansky *et al.*, 2022] Gabriel A. Radvansky, Andrea C. Doolen, Kyle A. Pettijohn, and Maureen Ritchey. A new look at memory retention and forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 48(11):1698–1723, 2022.
- [Reimers and Gurevych, 2019] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019.
- [Shen *et al.*, 2024] Shuanghong Shen, Qi Liu, Zhenya Huang, Yu Zheng, Minghao Yin, Meng Wang, and Enhong Chen. A survey of knowledge tracing: Models, variants, and applications. *IEEE Transactions on Learning Technologies*, 17:1898–1919, 2024.
- [Su *et al.*, 2023] Jingyong Su, Junyao Ye, Liqiang Nie, Yilong Cao, and Yonghui Chen. Optimizing spaced repetition schedule by capturing the dynamics of memory. *IEEE Transactions on Knowledge and Data Engineering*, 35(10):10085–10097, 2023.
- [Tabibian *et al.*, 2019] Behzad Tabibian, Utkarsh Upadhyay, Abir De, Ali Zarezade, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. Enhancing human learning via spaced repetition optimization. *Proceedings of the National Academy of Sciences*, 116(10):3988–3993, 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [Verbert *et al.*, 2012] Katrien Verbert, Nikos Manouselis, Xavier Ochoa, Martin Wolpers, Hendrik Drachsler, Ivana Bosnic, and Erik Duval. Context-aware recommender systems for learning: a survey and future challenges. *IEEE transactions on learning technologies*, 5(4):318–335, 2012.
- [Vodrahalli *et al.*, 2024] Kiran Vodrahalli, Santiago Ontanon, Nilesch Tripuraneni, Kelvin Xu, Sanil Jain, Rakesh Shivanna, Jeffrey Hui, Nishanth Dikkala, Mehran Kazemi, Bahare Fatemi, et al. Michelangelo: Long context evaluations beyond haystacks via latent structure queries. *arXiv preprint arXiv:2409.12640*, 2024.
- [Wang *et al.*, 2024] Noah Wang, Zy Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhuan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, et al. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14743–14777, 2024.
- [Wang *et al.*, 2025] Tianfu Wang, Yi Zhan, Jianxun Lian, Zhengyu Hu, Nicholas Jing Yuan, Qi Zhang, Xing Xie, and Hui Xiong. Llm-powered multi-agent framework for goal-oriented learning in intelligent tutoring system. In *Companion Proceedings of the ACM on Web Conference 2025, WWW ’25*, page 510–519, New York, NY, USA, 2025. Association for Computing Machinery.
- [Zhang *et al.*, 2025] Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhanxin Hao, Jianxiao Jiang, Jie Cao, Huiqin Liu, Zhiyuan Liu, Lei Hou, and Juanzi Li. Simulating classroom education with LLM-empowered agents. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10364–10379, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics.
- [Zhao *et al.*, 2025] Jun Zhao, Yongzhuo Yang, Xiang Hu, Jingqi Tong, Yi Lu, Wei Wu, Tao Gui, Qi Zhang, and Xuanjing Huang. Understanding parametric and contextual knowledge reconciliation within large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.