

PAPER

An Artificial Intelligence-Driven Adaptive Interaction Model for Adult Mobile Learning

Huijuan Qin¹(✉), Kun Liu²¹Hebei Open University,
Shijiazhuang, China²Hebei University of
Economics and Business,
Shijiazhuang, Chinajessie_qhj@126.com**ABSTRACT**

Adult mobile learning is characterized by fragmentation, task orientation, and context-embedded features, which pose substantial challenges to conventional adaptive systems in generating real-time interaction strategies that simultaneously accommodate learners' cognitive states and contextual constraints on mobile devices. To address these challenges, Mobile Adult Learning Adaptive Interaction Model (MAL-ADAPT), a context-cognition dual-channel adaptive interaction model driven by a lightweight on-device large language model agent, was proposed. Adaptive interaction was formulated as a constrained partially observable Markov decision process, in which learning gain, engagement, and learner autonomy were jointly optimized subject to latency, energy, memory, and privacy constraints. At the perception layer, behavioral, device, and environmental signals were integrated to construct a contextual state representation. At the cognitive layer, mastery probabilities and associated uncertainty were estimated through low-rank knowledge tracing. At the interaction layer, a unified state representation was constructed through a gated interaction mechanism. State-conditioned soft prompts were then injected into the on-device large language model, and executable adaptive actions were generated through low-rank adaptation-based parameter adaptation, action-space mapping, an adult-learning adaptation score, and feasibility masking. In addition, an edge-cloud hybrid inference gate was introduced to invoke cloud-enhanced inference under complex conditions when network connectivity permitted, while direct transmission of raw behavioral data was avoided. Six groups of experiments were conducted to evaluate component contributions, on-device performance, policy effectiveness, real-world use, robustness, and constraint sensitivity. The results demonstrate that MAL-ADAPT improved learner engagement, strategy acceptance, and self-regulated learning under mobile resource constraints while reducing response latency and dependence relative to cloud-based approaches. These findings provide a reusable technical pathway for on-device educational agents.

KEYWORDS

adult mobile learning, adaptive interaction, on-device large language model, knowledge tracing, dual-channel state representation, mobile energy efficiency

Qin, H., Liu, K. (2026). An Artificial Intelligence-Driven Adaptive Interaction Model for Adult Mobile Learning. *International Journal of Interactive Mobile Technologies (iJIM)*, 20(20), pp. 93–107. <https://doi.org/10.3991/ijim.v20i20.63556>

Article submitted 2026-06-16. Revision uploaded 2026-09-03. Final acceptance 2026-09-04.

© 2026 by the authors of this article. Published under CC-BY.

1 INTRODUCTION

Mobile learning scenarios for adult learners are characterized by temporal fragmentation, task orientation, a strong experiential foundation, and pronounced autonomy needs [1]. Learning is often initiated by adults during commuting, work breaks, or family contexts; sessions are short and frequently interrupted, and the knowledge state and the contextual state are highly misaligned [2]. Existing adaptive learning systems can dynamically adjust instructional content according to learners' cognitive, emotional, and behavioral states, but they remain primarily focused on content personalization rather than the real-time generation of interaction strategies under mobile contextual constraints [3, 4]. New possibilities for adaptive interaction have been provided by large language models [5]. However, cloud-based deployment in mobile learning scenarios remains constrained by response latency, network dependence, privacy risks, and energy consumption, whereas edge computing provides a feasible means of reducing data-transmission delays and supporting real-time local processing [6]. Local policy inference has been made possible by on-device lightweight large language models; however, existing on-device models are mostly focused on compression and acceleration and have not yet been unified with instructional state estimation, adult learning constraints, and closed-loop feedback optimization [7]. Theoretically, adaptive interaction is advanced from content recommendation to state-driven interaction strategy generation. Technically, an on-device large language model agent and a dual-channel state fusion mechanism are proposed. Practically, a low-latency, offline-capable, and privacy-friendly adaptive interaction solution is provided for adult mobile learning [8].

The limitations of existing research are mainly reflected in four aspects. First, the mobile adaptation of adaptive interaction based on cloud-based large language models is insufficient; immediate interaction in fragmented scenarios cannot be supported under conditions of high latency, strong network dependence, and privacy leakage risk. Second, context awareness and cognitive diagnosis are separated from each other. In existing work, either contextual recommendation [9] or knowledge tracing [10] is emphasized, and a unified state representation is lacking. As a result, strategy generation cannot be simultaneously constrained by the learner's knowledge state and the current context. Third, adult learning characteristics have not been transformed into computable constraints. Autonomy, experiential connectedness, task relevance, and fragmentation adaptation are mostly described qualitatively and have not been incorporated into the constraint space of strategy optimization [11]. Fourth, on-device inference and instructional strategy optimization are disconnected. Research on on-device model compression has rarely been integrated with constrained reinforcement learning, closed-loop feedback, and real-world mobile user studies [12]. Together, these limitations indicate that an adaptive interaction model is required that unifies on-device inference, dual-channel state estimation, adult learning constraints, and closed-loop optimization.

To address the aforementioned issues, Mobile Adult Learning Adaptive Interaction Model (MAL-ADAPT) is proposed, which is a contextual-cognitive dual-channel adaptive interaction closed loop driven by an on-device lightweight large language model agent. In this model, a dual-channel state representation and a gated interaction fusion mechanism are designed, through which mobile device interaction signals, environmental contextual signals, and lightweight knowledge tracing outputs are unified into a fused state. On this basis, state-conditioned soft prompt, low-rank adaptation-based parameter adaptation, and action-space constraint generation

methods are proposed for the on-device large language model agent, such that model outputs are transformed from free text into executable interaction actions. Adult learning priors are further transformed into feasibility masks and adaptation scores, a constrained partially observable Markov decision process is constructed, and an edge–cloud hybrid inference gate and a periodic low-rank adaptation update mechanism are designed. Finally, six experimental groups are designed, covering ablation, on-device performance, policy effectiveness, adult user field study, robustness, and constraint sensitivity.

2 PROBLEM FORMULATION AND MODEL OVERVIEW

2.1 Problem definition and constrained partially observable Markov decision process

Three fundamental difficulties are encountered in adaptive interaction for adult mobile learning. The true cognitive state of the learner cannot be directly observed and can only be inferred indirectly from behaviors such as question answering, navigation, and dwell. The mobile context is highly dynamic, and different constraints on interaction modality, task difficulty, and session length are imposed by commuting, work breaks, and family scenarios. For adult learners, both autonomy and external scaffolding are required; autonomy is impaired by excessive intervention, whereas discontinuation in fragmented scenarios is likely to be caused by insufficient intervention. Therefore, adaptive interaction should be modeled as a sequential decision-making problem under mobile-side resource constraints and partial observability. Specifically, the adaptive interaction process is defined as:

$$M = \langle S, A, O, T, Z, R, \gamma \rangle \quad (1)$$

where, S is the true state space, A is the interaction action space, O is the observation space, T is the state transition function, $Z(o_t | s_t, a_{t-1})$ is the observation model, R is the reward function, and $\gamma \in [0, 1)$ is the discount factor. At time t , three types of signals—learning behavior, device interaction, and environmental context—are observed by the system:

$$o_t = (x_t^b, x_t^d, x_t^e) \in O \quad (2)$$

Knowledge concepts, answer correctness, response time, and navigation count are included in learning behavior observations. Touch pressure, swipe speed, screen orientation changes, and input intervals are included in device interaction observations. Network round-trip latency, time-of-day segment, semantic location labels, and battery level are included in environmental context observations.

The cognitive state and the contextual state are estimated by the system from the observation history, and a state representation for policy decision-making is formed:

$$\hat{s}_t = (k_t, c_t) \quad (3)$$

where, $k_t \in [0, 1]^K$ denotes the mastery probability of each knowledge concept estimated from the observation history; K is the number of knowledge concepts; and $c_t \in \mathbb{R}^{d_c}$ denotes a contextual state vector encoded from multi-source observations, encoding attentional engagement, interaction intention, network conditions, degree of temporal fragmentation, and battery constraints. Together, these two components

constitute an approximate representation of the latent state, and interaction policies are generated by the system conditioned on the fused state:

$$\pi_{\phi}(a_t | o_{\leq t}) = \int \pi_{\phi}(a_t | z_t) p(z_t | o_{\leq t}) dz_t \quad (4)$$

where, $z_t = f_{\psi}(c_t, k_t)$ is the dual-channel fused state estimated from the observation history, ϕ represents the trainable parameters of the on-device large language model agent, and ψ represents the state encoding parameters.

A predefined set of interaction actions is adopted as the action space:

$$A = \{a^{(1)}, a^{(2)}, \dots, a^{(M)}\} \quad (5)$$

This includes content recommendation, difficulty adjustment, modality switching, reflective questioning, example supplementation, rest suggestion, and autonomy choice prompting. Unlike traditional adaptive systems, manner, rhythm, and modality of interaction are simultaneously determined by the actions herein. Constrained cumulative reward maximization is adopted as the optimization objective:

$$\max_{\phi} J(\phi) = E_{\tau \sim \pi_{\phi}} \left[\sum_{t=0}^T \gamma^t r_t \right] \quad (6)$$

$$\text{s.t. } E[Lat_t] \leq L_{max}, E[E_t] \leq E_{max}, M_{total} \leq M_{max} \quad (7)$$

where, Lat_t is the on-device policy generation latency, E_t is the energy consumption per interaction, and M_{total} is the total memory footprint of the on-device model. Latency, energy consumption, and memory—concerns central to mobile technology venues—are directly incorporated into the optimization framework by this constrained partially observable Markov decision process.

2.2 Model overview: Three-layer closed-loop architecture

As illustrated in Figure 1, MAL-ADAPT is composed of a perception layer, a cognition layer, and an interaction layer, forming a rapid closed loop of perception, decision-making, response, and re-perception. At the perception layer, multi-source signals from the mobile device are compressed into a contextual vector c_t , and attentional engagement is estimated. A hybrid structure of one-dimensional convolution and attention pooling is adopted at this layer to accommodate the mobile neural processing unit. Device interaction signals are used to infer attentional engagement and interaction intention, while environmental contextual signals are used to determine the type of learning scenario. The compressed contextual vector, rather than raw sensor data, is output by the perception layer, providing a basis for reducing the external transmission of raw data in edge–cloud hybrid inference.

The cognition layer is responsible for answering what the learner currently knows. A lightweight on-device knowledge tracing model is adopted by the system to estimate k_t and to compute cognitive uncertainty. When cognitive uncertainty is high, diagnostic questioning or scaffolded explanation is preferentially selected by the large language model agent; when cognitive uncertainty is low, but answer correctness fluctuates greatly, variant practice is preferentially selected. Through this design, the cognition layer is directly coupled with the interaction layer.




Fig. 1. Architecture of the MAL-ADAPT on-device contextual-cognitive dual-channel adaptive interaction model

Notes: MAL-ADAPT: Mobile Adult Learning Adaptive Interaction Model; AI: Artificial intelligence; LoRA: Low-rank adaptation.

The interaction layer constitutes the core innovation. c_t and k_t are fused into z_t through gated interaction fusion, and z_t is then injected into the on-device large language model through a state-conditioned soft prompt. Parameter-efficient adaptation is performed using low-rank adaptation, and an interaction policy constrained by the action space is generated. The model backbone parameters are frozen, and only the low-rank adaptation parameters and the policy head are trained. Model outputs are not free text but are constrained to the predefined action space. Learner responses to actions are adopted as feedback, the cognitive state and the contextual state are updated, and the immediate reward is computed.

3 METHODOLOGY

3.1 Perception layer: Multi-source contextual encoding and attentional engagement estimation

At the perception layer, high-noise, heterogeneous, and sparse signals on mobile devices are transformed into a contextual state that constrains interaction policies. Learning behavior, device interaction, and environmental contextual signals are collected within a sliding window of length w . Device interaction signals are first encoded by one-dimensional convolution, and local temporal patterns are extracted. Environmental contextual signals are transformed into low-dimensional vectors through lightweight embedding. Subsequently, weighted aggregation of the device interaction encodings within the window is performed using attention pooling, whereby the contextual vector c_t is obtained and attentional engagement att_t is estimated. When attentional engagement falls below a threshold, the agent is prohibited from delivering high-cognitive-load content, and only modality switching,

micro-quiz, or rest suggestion is permitted to be selected. The contextual channel is thereby transformed from background information into a hard constraint on interaction policies. Specifically, the contextual vector and attentional engagement are respectively given by:

$$c_t = \sum_{i=1}^w \alpha_i d_{t-w+i}, \alpha_i = \frac{\exp(q^\top W[d_{t-w+i}; e_t])}{\sum_j \exp(q^\top W[d_{t-w+j}; e_t])} \quad (8)$$

$$att_t = \sigma(w_a^\top c_t + b_a) \quad (9)$$

where, d_t denotes the device interaction encoding, and e_t denotes the environmental contextual embedding. The compressed contextual vector, rather than raw sensor data, is output by the perception layer, whereby a basis is provided for reducing the external transmission of raw data in subsequent edge-cloud hybrid inference.

3.2 Cognition layer: Lightweight knowledge tracing and cognitive uncertainty

At the cognition layer, a sliding-window deep knowledge tracing variant is adopted, and the mastery probability of the learner for each knowledge concept is dynamically estimated within each interaction segment. Knowledge concepts, answer correctness, and response time are mapped into observation embeddings, which are input into a low-rank gated recurrent unit knowledge tracer. The cognitive state is maintained across sessions by the gated recurrent unit, and the parameter count is compressed by low-rank decomposition, such that the parameter count of this layer is controlled at the million level, making it suitable for sharing mobile-side neural processing unit computation with the on-device large language model.

A cognitive state vector k_t is output by the knowledge tracer, in which each dimension represents the mastery probability of one knowledge concept. Cognitive uncertainty is further computed to determine the type of strategy required. When cognitive uncertainty is high, diagnostic questioning or scaffolded explanation is preferentially selected by the large language model agent; when cognitive uncertainty is low but large fluctuations in answer correctness are observed, variant practice is preferentially selected.

Specifically, the hidden state and the cognitive state are given by:

$$h_t = GRU(h_{t-1}, u_t), k_t = \sigma(U_k V_k h_t + b_k) \quad (10)$$

where,

$$U_k \in \mathbb{R}^{K \times r}, V_k \in \mathbb{R}^{r \times d_h}, r \ll \min(K, d_h) \quad (11)$$

The cognitive uncertainty is given by:

$$H(k_t) = -\sum_{i=1}^K [k_{t,i} \log k_{t,i} + (1 - k_{t,i}) \log (1 - k_{t,i})] \quad (12)$$

An independent low-dimensional state projection and caching mechanism is adopted by the knowledge tracer so that the additional memory footprint during parallel operation with the large language model agent is reduced.

3.3 Dual-channel fusion: Gated interaction fusion

Dual-channel fusion serves as the key connection point of the unified innovation. If the contextual vector and the cognitive vector were merely concatenated, the priority between cognitive needs and contextual constraints would be difficult for the on-device large language model agent to determine. Therefore, gated interaction fusion is adopted. The contextual state and the cognitive state are first projected into a shared dimension, and trade-off coefficients between them are computed through a gating mechanism. Subsequently, contextual features are recalibrated by a dimension-wise interaction gate so that contextual dimensions relevant to the current cognitive state are highlighted in the fused result. Both what the learner knows and the context at the present moment are encoded in the fused state z_t , and the priority in cases of conflict between the two is explicitly expressed. Specifically, the gating coefficient and the preliminary fused state are given by:

$$g_t = \sigma(W_g[W_c c_t; W_k k_t]) \quad (13)$$

$$\bar{z}_t = g_t \odot W_c c_t + (1 - g_t) \odot W_k k_t \quad (14)$$

The dimension-wise interaction gate and the final fused state are formulated as:

$$q_t = W_q k_t, v_t = W_v c_t \quad (15)$$

$$\alpha_t = \sigma(W_\alpha[q_t; v_t; q_t \odot v_t] + b_\alpha), z_t = [\bar{z}_t; \alpha_t \odot v_t] \quad (16)$$

To ensure on-device stability, the projection dimension is kept low-dimensional, and lightweight regularization is imposed on the interaction gate so that policy rigidity caused by prolonged gate saturation is avoided.

3.4 On-device large language model agent: State-conditioned soft prompt, low-rank adaptation, and policy generation

The interaction layer serves as the core carrier of innovation. The fused state z_t is first projected through a low-dimensional projection to form a state-conditioned soft prompt. This soft prompt is input into the on-device large language model together with historical interaction and system instructions. Parameter-efficient adaptation of designated linear layers in the model is performed by low-rank adaptation. The model backbone parameters are frozen, and only the soft prompt projection, low-rank adaptation parameters, and the policy head are trained. The on-device trainable parameters are compressed to a small scale, which is suitable for periodic updates. Mobile-side executability and instructional safety are ensured by mapping the hidden representations of the model to a predefined interaction action space via the policy head. Specifically, the state-conditioned soft prompt and the low-rank adaptation update are formulated as:

$$p_t = Pz_t + b_p \quad (17)$$

$$W'_l = W_l + B_l A_l, A_l \in \mathbb{R}^{r \times d_{in}}, B_l \in \mathbb{R}^{d_{out} \times r} \quad (18)$$

The hidden representation and action logits of the large language model are given by:

$$h_t^{\text{LLM}} = \text{LLM}_{\theta}([p_t; H_{t-1}; S]), l_t = W_{\text{policy}} h_t^{\text{LLM}} + b_{\text{policy}} \quad (19)$$

where, r is the low-rank adaptation rank, θ denotes the frozen backbone parameter of the large language model, H_{t-1} denotes the historical interaction, S denotes the system instruction, and $l_t \in \mathbb{R}^M$ denotes the action logits output by the policy head. The action distribution is given by:

$$\pi_{\phi}(a | z_t) = \frac{\exp(l_{t,a})}{\sum_{a' \in A} \exp(l_{t,a'})} \quad (20)$$

3.5 Adult learning prior constraints and feasibility mask

To align policies with adult learning characteristics, a feasibility mask and an adult learning adaptation score are introduced. The feasibility mask acts as a hard constraint, by which actions that are non-executable in the current context or instructionally unsafe are masked out. When attentional engagement is excessively low, high-difficulty tasks are prohibited; when network latency is excessively high, the video modality is prohibited; when cognitive uncertainty is excessively high, diagnostic questioning is prioritized; and when battery level is excessively low, high-energy-consumption interaction is prohibited.

The adult learning adaptation score acts as a soft constraint, by which actions more consistent with adult learning characteristics are prioritized among feasible actions. Task relevance, autonomy support, experiential connectedness, and a fragmentation adaptation penalty are included in this score. The final action distribution is jointly determined by the raw logits of the large language model, the feasibility mask, and the adult learning adaptation score. Specifically, the feasibility mask and the adaptation score are given by:

$$m_t(a) = \prod_{j=1}^J \mathbf{1}[g_j(z_t, a) \geq \tau_j] \quad (21)$$

$$C(z_t, a) = w_1 \text{Rel}(a, z_t) + w_2 \text{Aut}(a) + w_3 \text{Exp}(a, k_t) - w_4 \text{Frag}(a, c_t) \quad (22)$$

The final action distribution is given by:

$$\tilde{\pi}_{\phi}(a | z_t) = \text{softmax}(l_{t,a} + \log m_t(a) + \lambda C(z_t, a)) \quad (23)$$

$$a_t \sim \tilde{\pi}_{\phi}(a | z_t) \text{ or } a_t = \arg \max_a \tilde{\pi}_{\phi}(a | z_t) \quad (24)$$

Thresholds should be neither excessively strict, whereby the action space would become too small and the policy would degenerate into a rule-based system, nor excessively loose, whereby the constraint would lose its significance.

3.6 Closed-loop feedback and constrained partially observable Markov decision process optimization

The final component of the proposed framework is closed-loop feedback. Acceptance, ignorance, modification, completion, and interruption are included

in learner responses to actions. The cognitive state and the contextual state are simultaneously updated by this response, and the immediate reward is computed. Specifically, learner responses and state updates are formulated as:

$$y_t \in \{\text{acceptance, ignoring, modification, completion, interruption}\} \quad (25)$$

$$k_{t+1} = KT(k_t, a_t, y_t), c_{t+1} = Enc(o_{t+1}) \quad (26)$$

The immediate reward is given by:

$$r_t = \alpha \Delta k_t + \beta Eng_t + \eta Acc_t + \kappa Aut_t - \lambda_1 Lat_t - \lambda_e E_t - \lambda_c CL_t \quad (27)$$

where, Δk_t denotes the average knowledge gain, Eng_t denotes engagement, Acc_t denotes the strategy acceptance rate, Aut_t denotes autonomy satisfaction, Lat_t denotes the on-device policy generation latency, E_t denotes energy consumption, and CL_t denotes a cognitive load proxy. The optimization objective is given by:

$$\max_{\phi} J(\phi) = E_{\tau \sim \tilde{\pi}_{\phi}} \left[\sum_{t=0}^T \gamma^t r_t \right] \quad (28)$$

$$\text{s.t. } E[Lat_t] \leq L_{\max}, E[E_t] \leq E_{\max}, M_{\text{total}} \leq M_{\max} \quad (29)$$

In actual training, Lagrangian constrained proximal policy optimization is adopted, in which latency and energy consumption are treated as cost constraints, and policy parameters and constraint multipliers are updated jointly; the memory cap is controlled as a hard constraint during the model configuration stage. Inference is executed only on the device side, and low-rank adaptation parameters are periodically updated in the cloud after feedback from multiple devices is aggregated and are then delivered to the device side.

3.7 Edge-cloud hybrid inference and on-device adaptation

Given the difficulty of on-device training, a hybrid mechanism is adopted in which inference is performed on the device and low-rank adaptation parameters are periodically updated in the cloud. Only state encoding, knowledge tracing, and forward inference of the large language model are executed on the device side. Low-rank adaptation parameters are updated in the cloud after feedback from multiple devices is aggregated and are then delivered to the device side. For complex semantic retrieval or requests beyond on-device capability, an edge-cloud hybrid inference gate is established. Cloud-based enhancement is triggered when cognitive uncertainty or contextual estimation uncertainty exceeds a threshold, or when external knowledge retrieval is required by the current request, and the network round-trip latency and available bandwidth satisfy the conditions for cloud invocation. When triggered, only the compressed fused state and request features are uploaded to the cloud, such that direct external transmission of raw behavioral data is avoided. When not triggered, the policy is generated entirely on the device. Specifically, the hybrid inference gate is given by:

$$I_{\text{cloud}} = \mathbf{1}[(H(k_t) + U(c_t) > \tau_H \text{ or } \text{NeedRetrieval}(q_t) = 1) \wedge RTT_t \leq R_{\max} \wedge BW_t \geq B_{\min}] \quad (30)$$

where, $U(c_t)$ denotes contextual estimation uncertainty, q_t denotes the current request feature, and RTT_t and BW_t denote network round-trip latency and available

bandwidth, respectively. When $I_{cloud} = 0$, the policy is generated entirely on the device; when $I_{cloud} = 1$, only the compressed z_t and request features are uploaded to the cloud for enhancement. Model compression and runtime optimization are also involved in on-device adaptation.

4 EXPERIMENTAL DESIGN AND ANALYSIS OF RESULTS

4.1 Component ablation and dual-channel contribution analysis

The necessity of the dual-channel design and the on-device large language model agent was examined experimentally. Ablation variants were constructed as follows: full MAL-ADAPT (Model 1), MAL-ADAPT with the contextual channel removed (Model 2), MAL-ADAPT with the cognitive channel removed (Model 3), MAL-ADAPT with gated interaction fusion removed (Model 4), MAL-ADAPT with the low-rank adaptation-based large language model agent removed (Model 5), and MAL-ADAPT with adult learning constraints removed (Model 6). The experimental results are presented in Figure 2.




Fig. 2. Component ablation results

As indicated by the experimental results, the highest values in strategy acceptance rate, knowledge gain, and engagement, along with the lowest cognitive load, were achieved by full MAL-ADAPT. After the contextual channel was removed, an increase from 0.04 to 0.27 in the high-load content push rate under low attentional engagement was observed, indicating that a hard constraint effect on avoiding mismatched push is exerted by the contextual channel. After the cognitive channel was removed, a decrease in knowledge gain from 0.31 to 0.22 was observed, indicating that instructional targeting cannot be guaranteed by contextual information alone. After gated interaction fusion was removed, a decrease in strategy acceptance rate to 0.69 was observed, indicating that the conflict between cognitive needs and contextual constraints is difficult to handle effectively by simple concatenation. After the low-rank adaptation-based large language model agent was removed, the most pronounced decreases in engagement and knowledge gain were observed, demonstrating that the core component for strategy generation is the on-device large language model agent. After adult learning constraints were removed, decreases in both acceptance rate and engagement were observed, indicating that a practical contribution to strategy adaptation is made by adult learning prior constraints.

4.2 On-device deployment performance and energy-efficiency evaluation

Deployment performance of the full model, a cloud-based large language model, a distilled small model, and a rule-based policy were measured on real mobile devices. The experimental results are presented in Table 1.

Table 1. On-device deployment performance and energy efficiency

Approach	Device	P50 Latency (ms)	P95 Latency (ms)	Peak Memory (MB)	Energy Per Interaction (J)	Offline Success Rate	Cloud Invocation Rate
MAL-ADAPT	High-end A	180	320	1260	1.8	99%	8%
MAL-ADAPT	Mid-range B	260	480	1240	2.3	98%	10%
MAL-ADAPT	Low-end C	410	760	1220	3.1	96%	14%
Cloud-based large language model	Mid-range B	650	1500	Cloud-side	4.5	0%	100%
Distilled small model	Mid-range B	120	210	350	1.1	99%	5%
Rule-based policy	Mid-range B	40	70	80	0.3	100%	0%

Note: MAL-ADAPT: Mobile Adult Learning Adaptive Interaction Model.

As indicated by the experimental results, a P50 latency of 260 ms and a P95 latency of 480 ms were achieved by MAL-ADAPT on the mid-range device, both of which are lower than the 650 ms and 1500 ms observed for the cloud-based large language model. A P95 latency of 760 ms was recorded on the low-end device; although higher than that on the high-end device, quasi-real-time interaction in fragmented learning was still found to be satisfied. Offline success rates above 96% were achieved by the on-device approaches, whereas an offline success rate of 0% was observed for the cloud-based large language model, indicating that on-device deployment is irreplaceable for mobile learning scenarios. Although the lowest latency and

energy consumption were exhibited by the distilled small model, its policy quality was found to be clearly insufficient in Experiment 3. The cloud invocation rate of MAL-ADAPT was controlled at 8% to 14%, indicating that cloud dependency was effectively reduced by the hybrid inference gate. In terms of energy consumption, client-side energy consumption per completed interaction is reported in the table; 2.3 J per interaction was recorded for MAL-ADAPT on the mid-range device, which is lower than the 4.5 J for the cloud-based large language model approach.

4.3 Effectiveness of adaptive strategies and baseline comparison

A randomized controlled design was adopted, in which MAL-ADAPT was compared with a contextual bandit, collaborative filtering, traditional knowledge tracing-based recommendation, and a cloud-based large language model. The experimental results are presented in Figure 3.




Fig. 3. Policy effectiveness and baseline comparison

Note: MAL-ADAPT: Mobile Adult Learning Adaptive Interaction Model.

As indicated by the experimental results, superior performance to all baselines was achieved by MAL-ADAPT in knowledge gain, quiz accuracy, strategy acceptance rate, and self-regulated learning improvement. Compared with the cloud-based large language model, knowledge gain was increased from 0.29 to 0.31, and strategy acceptance rate was increased from 0.72 to 0.78, indicating that higher policy quality is achieved when the on-device large language model agent is combined with dual-channel constraints. In terms of cognitive load, a value of 3.2 was recorded for MAL-ADAPT, which is lower than that of most baselines, indicating that high-load content push is effectively avoided by adult learning constraints and contextual gating. In terms of session duration, 12.9 min was recorded for MAL-ADAPT, which is higher than the 12.4 min recorded for the cloud-based large language model, indicating that a longer single-session duration was exhibited by learners in this group. Although strong knowledge tracing capability was exhibited by deep knowledge tracing with rules method, a strategy acceptance rate of only 0.61 was observed, indicating that high-quality interaction strategies cannot be generated by cognitive state alone.

4.4 Adult learner field study

Employed and non-employed adult learners were recruited, and a four-week field deployment was conducted. MAL-ADAPT was deployed in the experimental group, whereas traditional adaptive recommendation was deployed in the control group. Logs, scales, and semi-structured interviews were collected. The experimental results are presented in Table 2.

Table 2. Results of the adult field study

Group	Learner Type	Engagement	Strategy Acceptance Rate	Autonomy Satisfaction	Self-Regulated Learning Improvement	Fragmentation Adaptation	Interruption Recovery Rate
Experimental group	Employed	0.79	0.77	4.3	0.27	4.4	0.72
Control group	Employed	0.61	0.54	3.6	0.12	3.5	0.51
Experimental group	Non-employed	0.74	0.73	4.1	0.23	4.2	0.68
Control group	Non-employed	0.58	0.50	3.4	0.10	3.3	0.47

As indicated by the experimental results, an engagement of 0.79 was recorded for employed learners in the experimental group, which is higher than the 0.61 recorded for the control group; a strategy acceptance rate of 0.77 was recorded, which is higher than the 0.54 recorded for the control group. In terms of autonomy satisfaction, a value of 4.3 was recorded for employed learners in the experimental group, and a value of 3.6 was recorded for the control group, indicating that adult autonomy was not sacrificed by MAL-ADAPT while scaffolding was provided. In terms of fragmentation adaptation, a value of 4.4 was recorded for employed learners in the experimental group, and a value of 3.5 was recorded for the control group, indicating that adaptation to short sessions was effectively improved by the contextual channel and adult constraints. An interruption recovery rate of 0.72 was recorded for the experimental group, and a rate of 0.51 was recorded for the control

group, indicating that learner recovery after interruption is facilitated by continuous maintenance of the cognitive state and on-device rapid inference. Similar but slightly lower trends were observed for non-employed learners, which may be related to differences in usage scenarios.

5 CONCLUSION

MAL-ADAPT, an adaptive interaction model driven by an on-device lightweight large language model agent for adult mobile learning, was proposed. Adaptive interaction was uniformly modeled as a constrained partially observable Markov decision process. Under mobile-side latency, energy consumption, memory, and privacy constraints, contextual states were generated by the perception layer, knowledge mastery and cognitive uncertainty were estimated by the cognition layer, and a unified state was formed through gated interaction fusion. Executable interaction actions were then generated by means of a state-conditioned soft prompt, low-rank adaptation parameter adaptation, action-space constraints, an adult learning adaptation score, and a feasibility mask. Cloud-based enhancement was triggered by the edge–cloud hybrid inference gate under complex contexts and favorable network conditions, while direct external transmission of raw behavioral data was avoided. Consequently, contextual awareness, knowledge tracing, policy generation, adult learning constraints, and on-device optimization were integrated into a single closed loop, and a technical pathway centered on an on-device large language model agent was established. In future research, multimodal contextual perception, federated low-rank adaptation updating, cross-device transfer, and large-scale longitudinal validation may be further explored so that the applicability of the model in broader adult learning scenarios can be examined.

6 REFERENCES

- [1] N. Moussa, M. M. Abouelnour, and T. Saali, “Digital library and adult education: Assessing acceptance, utilization, and adaptability challenges in AI-driven and mobile-enabled library services,” *International Journal of Interactive Mobile Technologies*, vol. 19, no. 24, pp. 137–148, 2025. <https://doi.org/10.3991/ijim.v19i24.59111>
- [2] P. Waterhouse, R. Samra, and M. Lucassen, “Distance education students’ satisfaction: Do work and family roles matter?” *Distance Education*, vol. 43, no. 1, pp. 56–77, 2022. <https://doi.org/10.1080/01587919.2021.2020622>
- [3] H. A. Sendeku, R. Daniel, and G. V. Gopal, “An AI-powered adaptive learning framework for personalized education,” *Information Dynamics and Applications*, vol. 4, no. 4, pp. 212–223, 2025. <https://doi.org/10.56578/ida040403>
- [4] J. Wang, “Intelligent education based on mobile learning: Transitioning from traditional classrooms to adaptive learning environments,” *International Journal of Interactive Mobile Technologies*, vol. 19, no. 11, pp. 51–65, 2025. <https://doi.org/10.3991/ijim.v19i11.56057>
- [5] E. Kasneci *et al.*, “ChatGPT for good? On opportunities and challenges of large language models for education,” *Learning and Individual Differences*, vol. 103, no. 1, Art. no. 102274, 2023. <https://doi.org/10.1016/j.lindif.2023.102274>
- [6] A. K. Pallikonda, V. K. Bandarapalli, and V. Aruna, “Enhancing performance and reducing latency in autonomous systems through edge computing for real-time data processing,” *Mechatronics and Intelligent Transportation Systems*, vol. 4, no. 3, pp. 154–165, 2025. <https://doi.org/10.56578/mits040305>

- [7] R. Wang, Z. Gao, L. Zhang, S. Yue, and Z. Gao, "Empowering large language models to edge intelligence: A survey of edge efficient LLMs and techniques," *Computer Science Review*, vol. 57, no. 1, Art. no. 100755, 2025. <https://doi.org/10.1016/j.cosrev.2025.100755>
- [8] L. Tetzlaff, F. Schmiedek, and G. Brod, "Developing personalized education: A dynamic framework," *Educational Psychology Review*, vol. 33, no. 3, pp. 863–882, 2020. <https://doi.org/10.1007/s10648-020-09570-w>
- [9] P. Vallejo-Correa, J. Monsalve-Pulido, and M. Tabares-Betancur, "A systematic mapping review of context-aware analysis and its approach to mobile learning and ubiquitous learning processes," *Computer Science Review*, vol. 39, no. 1, Art. no. 100335, 2021. <https://doi.org/10.1016/j.cosrev.2020.100335>
- [10] C. Huang, W. Jiang, K. Li, J. Wu, and J. Zhang, "Enhancing learning process modeling for session-aware knowledge tracing," *Knowledge-Based Systems*, vol. 309, no. 1, Art. no. 112740, 2024. <https://doi.org/10.1016/j.knosys.2024.112740>
- [11] Y. Lai, N. Saab, and W. Admiraal, "University students' use of mobile technology in self-directed language learning: Using the integrative model of behavior prediction," *Computers & Education*, vol. 179, no. 1, Art. no. 104413, 2022. <https://doi.org/10.1016/j.compedu.2021.104413>
- [12] O. D. Incel and S. Ö. Bursa, "On-device deep learning for mobile and wearable sensing applications: A review," *IEEE Sensors Journal*, vol. 23, no. 6, pp. 5501–5512, 2023. <https://doi.org/10.1109/JSEN.2023.3240854>

7 AUTHORS

Huijuan Qin graduated from Minzu University of China in 2009 with Ancient Literary Theory major, currently she works in College of Liberal Arts and Education, Hebei Open University as an associate professor in the College of Liberal Arts and Education. Her main research interests are digital resource construction, elderly education and online learning empowered by artificial intelligence, etc. (E-mail: jessie_qhj@126.com).

Kun Liu graduated from Minzu University of China in 2012, majoring in Ethnoecology. She is currently an Associate Professor at the College of Statistics & Mathematics, Hebei University of Economics and Business. Her main research interests are biostatistics and medical statistics (E-mail: liukun20060223@126.com).