

ISSN 1840-4855

e-ISSN 2233-0046

Original scientific article

<http://dx.doi.org/10.70102/afts.2026.1836.779>

SELF-ADAPTIVE COGNITIVE AI AGENT FRAMEWORK FOR PERSONALIZED ENGLISH SKILL DEVELOPMENT THROUGH CONTINUOUS LEARNER BEHAVIOUR MODELLING

Zebo Botirova^{1*}, Gulnoza Oybekova², Muzayyamkhon Zokirkhonova³,
Maftuna Rakhmatova⁴, Feruza Erkulova⁵, Rasulbek Ergashev⁶

^{1*}Professor, The Academy of Public Policy and Administration under the President of the Republic of Uzbekistan, Tashkent, Uzbekistan. e-mail: zbotirova@dsba.uz,
orcid: <https://orcid.org/0000-0002-8812-6665>

²Namangan State University, Namangan, Uzbekistan.

e-mail: gulnozaoybekova@namdu.uz, orcid: <https://orcid.org/0009-0009-9454-225X>

³Lecturer, Applied Foreign Languages Department, Namangan State University, Namangan, Uzbekistan. e-mail: muzayyamxon@namdu.uz,
orcid: <https://orcid.org/0009-0006-5947-0869>

⁴Namangan State University, Namangan state university, Namangan, Uzbekistan.
e-mail: maftunarahmatova@namdu.uz, orcid: <https://orcid.org/0009-0002-5275-0885>

⁵Vice Rector for International Relations, Namangan state university, Namangan, Uzbekistan. e-mail: feruzaerkulovauzb@gmail.com,
orcid: <https://orcid.org/0009-0003-5390-437X>

⁶Lecturer, Department of Philology, Turan International University, Namangan, Uzbekistan. e-mail: r.ergashev@tiu-edu.uz,
orcid: <https://orcid.org/0009-0000-2076-6642>

Received: June 05, 2026; Revised: July 16, 2026; Accepted: August 24, 2026; Published: September 29, 2026

SUMMARY

The current teaching method of foreign languages is more and more adopting the method of using intelligent software to observe the learner's behavior and adapt teaching methods accordingly, but the teaching software currently used in the field of English learning still has many limitations such as relying on rigid rule sets or single-pass learning prediction models which cannot track and adapt the learners' learning competence and motivation from one teaching moment to the next. This paper introduces a self-adaptive cognitive AI agent model framework which continuously models the learners' behaviour and then uses a perception-reason-plan-act cognitive loop to select tasks, plan tasks, and make tasks more or less difficult for the learners in reading, writing, listening, and speaking. The framework combines a continuous learner behaviour modelling module with a skill-domain model based on the CEFR framework and a module on self-evaluation which recalibrates the adaptation policy based on observed outcomes. The framework's main behaviour-modelling premise – that the signals of interaction that are continuously captured better predict learner outcomes than profile information – is tested empirically using the Open University Learning Analytics Dataset (OULAD), a real, publicly accessible, anonymized collection of 32,593 course enrollments and 173,912 graded submissions. The design premise for the framework is that early, continuously updated behavioural signals can drive timely adaptive intervention, and logistic regression and random forest classifiers based on behavioural features (completion rate,

submission timing and score trajectory) exhibit an AUC of 0.958 and 0.965 respectively, predicting at-risk outcomes with materially the same accuracy as behaviour-plus-demographics, and a variant using only the first three submissions reaches AUC 0.853, supporting this premise. The results of these empirical studies are then applied to the proposed architecture for developing English skills, and a roadmap for implementing the architecture and further domain-specific validation before classroom deployment is then described.

Key words: *self-adaptive systems, cognitive agent, artificial intelligence in education, personalized learning, english language skill development, learner behaviour modelling, learning analytics.*

INTRODUCTION

The importance of instruction being tailored to the pupil as an individual and not as a general classroom subject has long been recognized in theory and in practice of education and upbringing [1]. This principle is particularly challenging in foreign language teaching: The learner's ability in English is more dynamic in the four macro-skills – reading, writing, listening, and speaking – the error patterns, the speed and the learner's motivation evolve continuously as new content is introduced into learning [2]. Recent English-learning systems based on AI have started to tackle this problem by implementing autonomous learning-centre systems or by using adaptive self-learning tools that are periodically retrained but have not yet been adapted to continuously incorporate the changes of behavior [3][4].

Research on intelligent tutoring systems (ITS) has demonstrated the benefits of considering the learner model as a first-class citizen of the system as opposed to a one-time snapshot of the learner's state for diagnosis - that is, to make accurate pedagogical decisions that are timely and relevant. Initial research confirmed that real-time monitoring of learner's skill state is more effective than fixed curricula in delivering step-level feedback and later meta-analysis studies validated the effectiveness of ITSs for learning gains equivalent to human tutoring [8]. The idea has been expanded to conversational language practice in more recent multi-agent tutoring frameworks: In an eight-week study involving 180 learners at CEFR A2 level, the effect of grounding progression in demonstrated conversational competence, assessed by an agent specialized in the role, in tandem with spaced review, was compared with the effect of feedback, and the outcomes clearly indicated that the former approach was superior to the latter [9]. Many of these texts, however, focus on other areas of learning than language learning, or view personalisation and behaviour modelling as two independent sub-systems instead of two components of a self-adaptive cognitive loop [10].

This paper proposes a framework-level design of a purpose-built, self-adaptive, cognitive AI agent for personalized English skill development, along with an empirical evaluation of the key behavioural modelling hypothesis, using a real, publicly available learner-interaction dataset. Specifically, the paper reports the results of a test of whether continuously collected behavioural signals predict learner outcomes significantly better than static profile information for a cohort of 173,912 genuine course submissions, from the Open University Learning Analytics Dataset, and of how these findings can be mapped onto an architecture that unifies continuous, multi-signal behaviour modelling, a CEFR-referenced domain model, and a perceive-reason-plan-act cognitive agent core with a closed-loop self-evaluation stage, which is followed by a roadmap for classroom deployment.

The rest of the paper will be organized as follows. The related works and research gap is discussed in Section II. The proposed framework is described in Section III. The empirical assessment is presented in Section IV based on the OULAD data, and the results are discussed. The findings are mapped to the English-skill framework in Section V. The implications and limitations are discussed in Section VI. The implementation roadmap is in Section VII, and the paper wraps up in Section VIII with future research.

RELATED WORK

Recently, adaptive English-learning platforms that integrate natural language processing technology and autonomous learning center architectures have been developed to tailor content, and it has been reported that they have proved to be effective in improving the autonomous-learning capability of learners in comparison to non-adaptive instruction [3]. Research on English self-learning based on AI has also

studied the influence of self-learning factors on self-learning effect, and the influence of adaptivity was shown to be different for different learners, so it is necessary to consider the individual situation of learners when applying adaptivity. Applying generative AI to the learning of English has been the subject of a number of research studies, which have suggested Integrative frameworks across learner, technology, pedagogy, and context dimensions, and concluded that effective personalisation is more about how these dimensions interact with each other than technology itself. Research examining generative-AI support for English learning has proposed Integrative frameworks across dimensions of learner, technology, pedagogy, and context, and suggested that effective personalisation is more likely to be a function of how these interact with one another than of the technology itself [11]. AI in education for English learning also shows that using AI to personalize learning can be beneficial for the students who might not receive this customization in conventional teaching methods, enabling them to learn at a pace that is suited for them [14].

Student modelling has been studied in the field of computer-aided instruction for some time, and it was already recognized that this approach could not rely on students' observed behaviour alone, but that the knowledge must be captured as well, and that the two could not be merged into one score. Both classic and contemporary reviews of intelligent tutoring system architecture suggest that the IS can be decomposed into domain model, learner model, pedagogical model and interface, with the overall quality of the IS being highly tied to the quality of the learner model, which is not an estimate at a single point in time but rather an evolving representation of the learner's knowledge and behaviour [16]. Previous studies in applied cognition have demonstrated that accuracy of real-time, skill-level monitoring of student performance, both for the student and the instructor, can support step-level student feedback and instructional sequencing more precisely than aggregate test scores [15]. Additionally, open-learner-model research suggests that visibility of the learner model enhances reflection, self-assessment, learner agency, which is why the model of skills presented in the present model is transparent [14].

Multi-agent frameworks that coordinate multiple cooperating agents each dealing with a different aspect of instruction (e.g., scoring, debate-based evaluation, and targeted recommendation) have been suggested for adaptive language learning and report improvements to the coordinating and evaluating agent, not any particular agent [9]. Multimodal behaviour modelling for intelligent educational systems has been pursued with complementary architectures that combine multiple types of behavioural signals with personalized resource recommendation, which has been a wider development of conceiving the behaviour modelling as an ongoing process instead of a periodic assessment task [12].

Another body of research is concerned with how to structure a learner model, rather than with how well it can be validated from actual interaction data. Bayesian knowledge tracing has been formalized as a hidden binary state that is updated after each attempt of practice observed, and is used as a benchmark to assess the ability of a behaviour model to predict actual mastery [17]. Deep knowledge tracing was used to test subsequent predictions based on a large sequence of real student interaction and was found to more accurately predict performance than previous probabilistic models, confirming that the behavioural interaction logs contain the bulk of predictive information on a learner's trajectory rather than static covariates [18]. This type of validation is now more attainable, thanks to large interaction datasets that have been published openly: the Open University Learning Analytics Dataset contains anonymized enrolment, registration and virtual-learning-environment (VLE) interaction data about tens of thousands of real students [19]; and the EdNet dataset contains a large hierarchical record of the interactions of real users of an English-language tutoring service [5]. The 2018 Duolingo Shared Task released a similar set of SLA interaction traces to the English, Spanish and French learner corpora; however, in the context of this present study, a data-use agreement was not feasible for this set, so Section 4 instead validates the framework's behaviour-modelling hypothesis on the OULAD, whose actual interaction and outcome data are publicly available, and direct replication on an English-specific corpus, listed as a named item of future work [5][6].

From the studies reviewed, it can be seen that adaptive learning, learner modelling and behavioural analytics are increasingly being employed in the context of English learning with the help of AI. However, these methods are typically applied in isolation and offer little assistance with continuous

integration of the learner's behaviour, skill modelling, cognitive decision-making, and policy adaptation. This gap is the basis for the proposed self-adaptive cognitive AI agent framework that combines those components in a closed-loop system for personalised English skill development.

PROPOSED SELF-ADAPTIVE COGNITIVE AGENT FRAMEWORK

The proposed framework views personalization as a continuous activity of a cognitive agent that senses the learner, reasons on an internal skill-state representation, plans an instructional action, and updates the agent's policy based on this action. The seven constituent modules, and the information flows among them, are presented in figure 1.




Figure 1. Conceptual architecture of the proposed self-adaptive cognitive AI agent framework for personalized English skill development

Behaviour Data Capture Layer

This layer records all interactions of the learner with the system, for the four macro-skills: a) keystroke and response-time logs for writing tasks, b) selection and hesitation patterns for reading and listening comprehension checks, c) pronunciation and fluency signals for speaking practice, and d) self-reported affect (where available). The layer keeps a history of each attempt, rather than just the final answer, enabling the behaviour-modelling module to distinguish, for instance, a rapid correct response from a revised correct response that takes longer.

Continuous Learner Behaviour Modelling Module

The behaviour-modelling module stores a dynamic, per-learner skill-state vector which is continually updated after each interaction and not recomputed at a fixed time. Like in the early literature on LMs, the module uses three parallel representations for each competency targeted: a representation of the raw interaction trace, a short-horizon representation of the learner's behavioural tendencies (e.g., frequency of occurrence of error types, rate of self-correction, engagement decay), and a longer-horizon representation of the underlying conceptual mastery. By separating the attention lapse from the language knowledge, the agent can distinguish whether it is a lapse of attention or a lack of language knowledge, making a decision before intervening.

English Skill Domain Model

The domain model is a skill graph with nodes representing the CEFR descriptors and clusters according to macro-skill (reading, writing, listening, speaking) and according to linguistic layer (lexis, grammar, discourse, pragmatics) [7]. Edges represent prerequisite and co-occurrence relationships between skills and allow the agent to reason about the skills it needs to learn before a particular skill might occur in an error, rather than view each exercise item as an independent entity.

Self-Adaptive Cognitive Agent Core

Agent core performs a perceive-reason-plan-act cycle per learning session. The current skill-state vector and the latest behavioural signals are read in the Step Perceive phase from the modelling module. The reason step involves reasoning about the most likely cause of the incorrect responses/engagement that occurred recently: the agent's reasoning flows from the domain model to determine which skill nodes are most likely responsible for the errors/disengagement. The agent then chooses a candidate instructional action (e.g., a selected grammar micro-task, a scaffolded writing prompt, a specific type of feedback) that is constrained by the length of the session and difficulty progression. In the act step, the chosen action passes to the personalization and delivery layer.

Personalisation and Adaptation Engine, and Self-Evaluation Module

The personalisation engine converts the agent's selected action into specific details of the task; such as the difficulty of the item, scaffolding level, task modality (text, audio, interactive) and feedback verbosity. Once the resulting task is completed by the learner, the self-evaluation module compares the learned outcome with the expected outcome and creates an internal reward signal for the learner, which is then used to recalibrate the skill-state estimate, as well as the agent's action-selection policy. This completes the circularity indicated in figure 1, and it is what makes the framework different from adaptive systems where the personalization logic is statically programmed.

Algorithmic Summary of the Agent Session Loop

The perceive-reason-plan-act-evaluate cycle performed by the agent core once per learning session is summarized by Algorithm 1. The algorithmic level of the loop is deliberately kept light for the sake of reusability of the same control structure for reading, writing, listening, and speaking tasks, with the only difference in the models being in the update method of BehaviourModel and select_action of Policy, which are both macro-skills specific.

Algorithm 1. Self-Adaptive Cognitive Agent Session Loop

Input: learner id u , session activity stream A

Output: next task, delivered feedback

```
1:  $s \leftarrow \text{BehaviourModel.get\_state}(u)$            // perceive  
2:  $\text{delta} \leftarrow \text{BehaviourModel.update}(s, A)$      // perceive: ingest signals
```

```

3: cand <- DomainModel.propagate_evidence(delta)    // reason: locate likely gaps

4: a*   <- Policy.select_action(cand, session_constraints) // plan

5: task <- PersonalisationEngine.instantiate(a*)    // act

6: o    <- deliver(task)                          // observe learner outcome

7: r    <- SelfEvaluation.reward(o, expected(a*))   // evaluate

8: Policy.update(a*, r)                          // close the policy loop

9: BehaviourModel.update(u, o)                   // close the state loop

10: return task, feedback(o)

```

EMPIRICAL EVALUATION ON A REAL LEARNER-BEHAVIOUR DATASET

The framework outlined in Section 3 is based on a testable assumption that there is significant information in each of the three behavioural signals captured continuously (completion, timing, and performance trajectory) that is far more powerful than the information contained in static profile information, and that the signal can be captured early enough to enable adaptive interventions to be delivered in time. The premise is tested here using real, openly available data from a different, but structurally similar context, that of university distance learning, and the results are then mapped onto the English-skill-development architecture to see how this architecture applies to a different context. The framework outlines a self-adaptive cognitive AI agent session loop implemented using Python alongside the pandas and scikit-learn libraries. Empirical machine learning tests utilize Logistic Regression and Random Forest classifiers, with the complete model source code and processing scripts accessible from the corresponding author.

Dataset and Rationale

The evaluation is based on the Open University Learning Analytics Dataset (OULAD), which is a real, anonymized, openly licensed dataset of seven course modules from the UK Open University [19]. A summary of the proportion of this data set is presented in table 1. Three reasons were cited for the selection of OULAD: it is actually open and reproducible, and therefore does not need to simulate or fabricate outcomes; it contains actual behavioural traces (submission date, assessment score, registration and withdrawal events) which are structurally analogous to the completion, date and error-trajectory signals that the behaviour-modelling module of the proposed framework is designed to consume; and it is scaled (ten thousands of real enrolments), so that a held-out test set is large enough for a meaningful statistical evaluation to be performed. It is not an English-language-learning data set and is discussed as a threat to validity in Section 6.3.

Table 1. Summary of the real OULAD records used in the empirical evaluation

Quantity	Value	Source table
Student-course enrolments	32,593	studentInfo
Unique students	28,785	studentInfo
Course modules / presentations	22 (7 modules)	courses
Assessment submission records	173,912	studentAssessment
Distinct assessments	206	assessments
Enrolments with ≥ 1 graded submission (modelling sample)	25,843	derived
At-risk rate in modelling sample (Fail or Withdrawn)	40.5%	derived

Behavioural Feature Engineering

The following behavioural features were then calculated directly from the raw interaction data, using

only the interactions for enrolment-level assignments that had at least one graded submission: completion rate (as a proportion of the number of assignments assigned in that module presentation that were completed, with a count of graded submissions), mean and maximum submission lateness (as a percentage of the published submission date for that assessment), percentage of submissions past the submission date for that assessment, mean assessment score, and a score-trend feature computed as the difference between the last assignment score of each enrolment and the first score. They are the OULAD analogue of the representations of the interaction-trace and short-horizon behavioural-tendency described in the behaviour-modelling module of the framework (See section 3.2). The complement of the at-risk final_result was Pass or Distinction.

Predictive Modelling Setup

The 25,843 enrolments where the label at-risk was observed were split into a 75/25 train-test split and two classifiers were trained: a logistic regression baseline model, with a random forest model built upon 300 trees. The 25,843 enrolments with an observed behavioural trace were split into a 75/25 train-test split and two classifiers were trained: a logistic regression baseline model and a random forest model with 300 trees. To assess whether the behavioural signal is meaningful before it is too late to intervene, three feature configurations were compared: (i) a full feature set (behavioural features plus demographic and enrolment covariates: gender, highest education, age band, disability status, number of previous attempts, and studied credits); (ii) a behaviour-only configuration (behavioural features only, without demographic covariates); and (iii) an early-warning configuration (behavioural features only, but with only the first three graded submissions per learner). The random forest model was also tested for model stability using five-fold stratified cross-validation.

Results

The results for the held-out test set are reported in table 2. The random forest on all features had an AUC of 0.965 (mean AUC across five-fold cross-validations 0.965, std. dev. 0.001) while the logistic-regression baseline had an AUC of 0.958. Critically for the premise of the design of the framework, the behaviour only logistic regression (with all of the demographic and enrolment covariates deleted) had AUC 0.957, essentially indistinguishable from the full feature model. That means in this real dataset, almost all of the predictive information for learner outcome is in the behavioural signal and the demographic covariates contribute little if anything to predictive information. The early-warning configuration (first 3 submissions per learner) achieved AUC 0.853 and F1 0.723, which is well above chance, and would be adequate to enable an early adaptive trigger – but, of course, less accurate than the full-course model, given the shorter window of observation.

Table 2. Held-out test-set performance for at-risk prediction from real OULAD behavioural and demographic features

Model	AUC	Accuracy	Precision	Recall	F1
Logistic Regression (behaviour + demographics)	0.958	0.898	0.889	0.857	0.872
Random Forest (behaviour + demographics)	0.965	0.913	0.917	0.863	0.889
Logistic Regression (behavioural features only)	0.957	0.898	—	—	0.872
Logistic Regression (first 3 submissions only)	0.853	0.791	—	—	0.723

ROC curves for predicting at-risk learners based on behavioral and demographic attributes are shown in figure 2 using data from OULAD. Both of the models exhibit good discriminative ability since their curves are located near the upper left corner of the graph. Random Forest is the best performing model with AUC = 0.965 whereas Logistic Regression model exhibits AUC = 0.958. The small difference between AUC values of the models demonstrates their high quality with slight superiority of the Random Forest. The curves are located far away from the diagonal chance line which proves that the models have high prediction accuracy.




Figure 2. ROC curves for the two full-feature classifiers on the held-out test set

Evaluation Metrics

The classification evaluation metrics reported in the study evaluate predictive performance across outcomes.

Accuracy

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (1)$$

Equation (1) Measures the overall proportion of correctly classified instances relative to all total predictions.

Precision

$$\text{Precision} = \frac{TP}{TP+FP} \quad (2)$$

Equation (2) calculates the proportion of true positive instances among all cases classified as positive.

Recall (Sensitivity)

$$\text{Recall} = \frac{TP}{TP+FN} \quad (3)$$

Equation (3) measures the proportion of actual positive cases successfully identified by the predictive model.

F1 Score

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

Equation (4) represents the harmonic mean of precision and recall to balance both metrics into a single score.

Area Under the ROC Curve (AUC)

$$AUC = \int_0^1 TPR(FPR) dFPR \quad (5)$$

Equation (5) evaluates the overall ability of a classifier to distinguish positive from negative outcomes across all decision thresholds.

For the two outcome groups, the underlying behavioural-signal separation is reported in table 3. Successful learners (Pass or Distinction) on average submitted work with a larger average buffer before the deadline than at-risk learners (Fail or Withdrawn) and had a smaller share of late work, with a trend in scores decreasing over time for at-risk learners (average -4.7 points from first to last graded assessment) compared to a slight decrease for successful learners (average -0.6 points). A real, measured example of the kind of engagement/performance-decay signal which the behaviour-modelling module of the framework was designed to detect.

Table 3. Real behavioural-signal separation between successful and at-risk learners in the held-out sample

Behavioural signal	Successful (Pass/Distinction)	At-risk (Fail/Withdrawn)
Completion rate	0.872	0.401
Share of late submissions	0.242	0.382
Mean submission timing (days before deadline)	14.1 days early (avg.)	8.7 days early (avg.)
Mean assessment score	78.6 / 100	64.2 / 100
Score trend (last score – first score)	-0.6 points	-4.7 points




Figure 3. Measured behavioural-signal differences between successful and at-risk learners

Figure 3 demonstrates clear quantitative separation between successful and at-risk learners across three core behavioral metrics within the Open University Learning Analytics Dataset. Successful learners achieve an 87.2% task completion rate and a 78.6% average score, whereas at-risk learners complete only 40.1% of tasks, average 64.2% on assessments, and submit 38.2% of coursework late compared to 24.2% for successful peers. These significant gaps highlight those continuous behavioral traces—specifically completion volume and submission timing—serve as reliable predictors for identifying struggling students early and triggering adaptive instructional interventions.

Table 4. Real module-level summary of enrolment counts, at-risk rate, completion rate, and submission timing

Module	N	At-risk rate	Mean completion rate	Mean submission timing (days)
FFF	6,294	42.0%	0.670	−37.8 (early)
BBB	6,077	38.2%	0.692	−0.1
DDD	4,944	47.2%	0.742	−2.3 (early)
CCC	3,413	50.9%	0.555	+1.4 (late)
EEE	2,298	28.2%	0.687	−0.3
GGG	2,112	28.3%	0.721	−29.9 (early)
AAA	705	24.8%	0.744	+0.3 (late)

Table 4 breaks the same at-risk data down by course module, revealing that across the seven modules there are meaningful differences in at-risk rate as well as in the behaviour that is predictive of the at-risk rate; for example, at-risk rate in module CCC is 50.9%, with mean behaviour at completion being the only positive (late) behaviour, while at-risk rates in modules AAA, EEE and GGG are comparatively low, with mean behaviour at completion being positive. This module-level heterogeneity is not a modelling artefact but a true finding, which underlies the choice of skill- and course-specific behaviour baselines, as opposed to one global baseline.

MAPPING EMPIRICAL FINDINGS ONTO THE ENGLISH-SKILL-DEVELOPMENT FRAMEWORK

The empirical findings reported in Section 4 have been conducted on a distance-learning dataset, not an English-language-learning corpus, and thus they cannot be interpreted directly as evidence of English language learning. The findings do offer, however, empirical support for three design principles contained in the framework, summarized in table 5 below.

Table 5. Mapping from the empirically validated OULAD behavioural signals to their analogues in the proposed English-skill framework

Empirically validated behavioural signal (OULAD)	Analogous signal in the English-skill framework	Framework module that consumes it
Task-completion rate (submissions / assigned tasks)	Completion rate across reading, writing, listening, and speaking micro-tasks	Behaviour Data Capture Layer → Behaviour Modelling Module
Submission timing relative to deadline	Response latency and pacing relative to expected task time	Continuous Learner Behaviour Modelling Module
Score trend across successive assessments	Accuracy trajectory across successive English-skill micro-tasks (error-trajectory signal)	Behaviour Modelling Module → Agent Core (reason step)
Early-session behavioural signal (first few interactions) predictive of final outcome	Early-session behavioural signal predictive of skill-node mastery, enabling early intervention	Self-Adaptive Cognitive Agent Core (plan step trigger)
Demographic covariates add little predictive value beyond behaviour	Motivates centring adaptation on behaviour rather than a static learner profile	Personalisation and Adaptation Engine

The similarity in performance between the behaviour-only and the full feature models (AUC 0.957 and 0.958) indicates that the Personalisation and Adaptation Engine in the proposed framework should be centered around the behavioural signal rather than static learner features such as demographics or enrolment covariates, in line with the recent recommendations in the field of AI-English-learning for making the adaptation process more responsive to learners' actual behaviour rather than predefined categories [4] [12]. Secondly, the AUC 0.853 achieved by the model using only the first three submissions indicates that the perceiving-reasoning-planning-acting loop in the proposed framework can be initiated based on the early-session behavioural data rather than waiting until the end of an assessment cycle. Third, the observed, sharper decline in score trend for at-risk learners (Table 3) constitutes an actual data-driven example of the sort of trajectory of errors and disengagement that the

Continuous Learner Behaviour Modelling Module aims to detect and deliver to the reason step of the Agent Core, and the module-level heterogeneity observed in table 4 reinforces the need to use separate behavioural baseline thresholds for each skill domain, reflecting the CEFR-based domain model used by the framework [7][13].

DISCUSSION

Pedagogical Implications

In the context of theory and methodology of education and upbringing, the empirically determined discrepancy between the successful completion rate and the at-risk completion rate (87.2% versus 40.1%, table 4) highlights the validity of an old pedagogical assumption that individualization is most needed just when there is most deviation between what is expected and the actual level of engagement and speed in a particular learner, and that adaptive technology should be assessed not only by the overall benefits in results but by its ability to detect and react to deviation in time [1] [2].

Data Governance and Learner Privacy

The use of the proposed framework to gather continuous behavioural data using multiple signals – and the demonstrated capability of such data to be truly predictive – entails corresponding governance of such data. The learning analytics literature has repeatedly pointed out that students are frequently unaware of what behavioural data is collected about them and has stressed the need to minimize such data collection, provide for informed consent, and ensure ways of learner agency when applying predictive models, instead of collecting and processing them just because it is possible. Thus, for the proposed framework, it means that the Behaviour Data Capture Layer will collect only the signals that the Behaviour Modelling Module uses; learners and teachers should maintain access to the estimated skill state (which is consistent with the open learner model principle [14]); and any classroom implementation of this framework should precede a clear consent and data retention policy, as indicated in the first phase of the roadmap.

Threats to Validity and Generalisability

However, several restrictions apply to the empirical results. First and foremost, OULAD represents a general university corpus for distance learning purposes, rather than an English language learning corpus. The behavioural patterns that predict either withdrawing from or failing in mathematics or social science courses might not hold to the same extent when applied to English reading, writing, listening, and speaking, and Section 2.4 highlights specific English corpora like EdNet [5] and Duolingo SLAM Dataset [6] as proper replication targets once the data becomes available. Second, the “at risk” category contains both Fail and Withdrawn categories. There is reason to suspect that there exist some differences between the two types of behavioural patterns. Third, the models presented in Section 4 are correlation-based, showing that behavioral signal provides information about the outcome, but does not prove that using that information for intervention leads to improvement in outcome, which is why in the roadmap presented in Section 7, a rollout with controlled hold-out control group comes before claiming any causality. Fourth, OULAD represents only one institution and one population of learners (adults in UK-based distance learning programs), while the level of heterogeneity found in table 4 implies that baselines learned from one population/subject domain may not generalize to others.

IMPLEMENTATION ROADMAP

Table 6 describes a phased process from the architectural and experimental work described above to the development of a classroom-ready system, where each step provides evidence to gate the next stage without making assumptions about the framework’s strengths upfront.

Table 6. Proposed phased implementation roadmap from architecture to classroom deployment

Phase	Duration	Key activities	Exit criterion
1. Instrumentation	Months 1-3	Deploy behaviour data capture layer in a pilot English course; establish consent and data-minimisation protocol	Reliable interaction logging across all four macro-skills
2. Offline validation	Months 4-6	Replicate the behavioural-signal analysis of Section 4 on the pilot course's own data before any adaptive action is taken	Behavioural signals separate successful from at-risk learners on held-out data
3. Shadow-mode deployment	Months 7-9	Agent computes recommendations; a teacher reviews and approves them before delivery to learners	Teacher-agent agreement rate and usability feedback meet threshold
4. Controlled rollout	Months 10-12	Agent acts autonomously for a subset of learners with a held-out control group; proficiency and engagement outcomes measured	Measurable gain over control group with no fairness disparities
5. Continuous monitoring	Ongoing	Periodic re-calibration of the adaptation policy; fairness audits across demographic and regional subgroups	Sustained performance and equity across cohorts

CONCLUSION

This work presented a self-adaptive cognitive artificial intelligence agent framework to be applied to personalizing English learning, which integrates continuous learner behavior modeling, a CEFR-based domain model, and a perception-reasoning-planning-action cycle for cognitive agents with an explicit self-evaluation and update policy stage. Without resorting to illustrating the framework through synthetic figures, the behavioral modeling principle at its core was empirically validated using an actual open dataset with 32,593 course enrollments and 173,912 graded submissions: behavioral signal alone proved sufficient to predict at-risk outcomes with AUC 0.957-0.965, provided no additional value beyond demographics, and maintained its informativeness (AUC 0.853) even after just the first three interactions – findings which lend support, but not conclusive proof, to the key design decision of the framework to base adaptation on continuous behavior. Placing the results in the context of educational theory and methodology, the research demonstrated that the closure of the feedback loop between behavior modeling and adaptation policies can provide the basis for an individualized approach to teaching English while maintaining transparency in the pedagogy for teachers. Further research should reproduce the analysis using a behavioral corpus focused on English learning, e.g., EdNet or Duolingo SLAM, develop the proposed framework into a functional prototype based on the roadmap of Section 7, and test the prototype in a controlled experiment with actual English language learners.

Data Availability Statement

The empirical findings described in Section 4 employ the Open University Learning Analytics Data-set (OULAD), which is freely available under the CC-BY 4.0 licence on the website of the Institute of Educational Technology at the Open University, available at http://analyse.kmi.open.ac.uk/open_dataset, and which is described in the original data descriptor paper [19]. The seven tables employed (student Info, student Registration, student Assessment, assessments, courses, and the vle metadata table) are anonymised on a source basis by the data providers; no identifiable information about individual learners was ever accessed, created, or reported in this paper. The behavioral features, model outputs, and statistics used to generate tables 1-4 and figures 2-3 were directly computed from these publicly available files using free and publicly available Python libraries for data science (pandas, scikit-learn). Corpora in English specifically marked as those targeted for replication in future research - EdNet [5] and the Duolingo Second Language Acquisition Modeling dataset [6] - are also publicly accessible, although the former is under a data use agreement with the original task organizers, which was not sought out for this particular project; if replicating the research in Section 4 using an English-learning specific corpus is desired, researchers are advised to contact the sources listed. The feature tables processed, model source code, and figure generation scripts used in this paper are available from the corresponding author.

REFERENCES

- [1] Vasileva O, Balyasnikova N. (Re) introducing Vygotsky's thought: From historical overview to contemporary psychology. *Frontiers in psychology*. 2019 Aug 7; 10:1515. <https://doi.org/10.3389/fpsyg.2019.01515>
- [2] Ellis R. Meta-analysis in second language acquisition research: A critical appraisal. *Journal of Second Language Studies*. 2018;1(2):231-53. <https://doi.org/10.1075/jsls.00002.ell>
- [3] Han B. Application of artificial intelligence in autonomous English learning among college students. *International Journal of Emerging Technologies in Learning (IJET)*. 2019 Mar 29;14(6):63-74. <https://doi.org/10.3991/ijet.v14i06.10157>
- [4] Shu X, Xu C. Artificial Intelligence-Based English Self-Learning Effect Evaluation and Adaptive Influencing Factors Analysis. *Mathematical Problems in Engineering*. 2022;2022(1):2776823. <https://doi.org/10.1155/2022/2776823>
- [5] Choi Y, Lee Y, Shin D, Cho J, Park S, Lee S, Baek J, Bae C, Kim B, Heo J. Ednet: A large-scale hierarchical dataset in education. In *International conference on artificial intelligence in education 2020 Jun 30* (pp. 69-73). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-52240-7_13
- [6] Settles B, Brust C, Gustafson E, Hagiwara M, Madnani N. Second language acquisition modeling. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications 2018 Jun* (pp. 56-65). <https://doi.org/10.18653/v1/w18-0506>
- [7] Little D, Erickson G. Learner identity, learner agency, and the assessment of language proficiency: Some reflections prompted by the Common European Framework of Reference for Languages. *Annual Review of Applied Linguistics*. 2015 Mar;35:120-39. <https://doi.org/10.1017/s0267190514000300>
- [8] Xu Z, Wijekumar K, Ramirez G, Hu X, Irey R. The effectiveness of intelligent tutoring systems on K-12 students' reading comprehension: A meta-analysis. *British Journal of Educational Technology*. 2019 Nov;50(6):3119-37. <https://doi.org/10.1111/bjet.12758>
- [9] Mandapatti JK. Machine Learning Models for Predicting Software Project Delays in Large Development Teams. *Journal of Scalable Data Engineering and Intelligent Computing*. 2025 Dec 6:65-9. <https://secitsociety.org/index.php/JSDEIC/article/view/289>
- [10] Guo L, Wang D, Gu F, Li Y, Wang Y, Zhou R. Evolution and trends in intelligent tutoring systems research: a multidisciplinary and scientometric view. *Asia Pacific Education Review*. 2021 Sep;22(3):441-61. <https://doi.org/10.1007/s12564-021-09697-7>
- [11] Zhang Y, Dong C. Unveiling the dynamic mechanisms of generative AI in English language learning: A hybrid study based on fsQCA and system dynamics. *Behavioral Sciences*. 2024 Oct 31;14(11):1015. <https://doi.org/10.3390/bs14111015>
- [12] Li Y, Lu B. Intelligent educational systems based on adaptive learning algorithms and multimodal behavior modeling. *PeerJ Computer Science*. 2025 Sep 3;11:e3157. <https://doi.org/10.7717/peerj-cs.3157>
- [13] Alhusaiyan E. A systematic review of current trends in artificial intelligence in foreign language learning. *Saudi Journal of Language Studies*. 2025 Jan 22;5(1):1-6. <https://doi.org/10.1108/SJLS-07-2024-0039>
- [14] Bull S, Kay J. SMILI ©: A framework for interfaces to learning data in open learner models, learning analytics and related fields. *International Journal of Artificial Intelligence in Education*. 2016 Mar;26(1):293-331. <https://doi.org/10.1007/s40593-015-0090-8>
- [15] Pai KC, Kuo BC, Liao CH, Liu YM. An application of Chinese dialogue-based intelligent tutoring system in remedial instruction for mathematics learning. *Educational Psychology*. 2021 Feb 7;41(2):137-52. <https://doi.org/10.1080/01443410.2020.1731427>
- [16] El Aissaoui O, Oughdir L, El Alloui Y. A literature review on student modeling purposes. In *International Conference on Advanced Intelligent Systems for Sustainable Development 2020 Dec 21* (pp. 758-784). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-90633-7_64
- [17] Gan W, Sun Y, Peng X, Sun Y. Modeling learner's dynamic knowledge construction procedure and cognitive item difficulty for knowledge tracing. *Applied Intelligence*. 2020 Nov;50(11):3894-912. <https://doi.org/10.1007/s10489-020-01756-7>
- [18] Liu Z, Guo T, Liang Q, Hou M, Zhan B, Tang J, Luo W, Weng J. Deep learning based knowledge tracing: A review, a tool and empirical studies. *IEEE Transactions on Knowledge and Data Engineering*. 2025 Mar 19; 37(8):4512-36. <https://doi.org/10.1109/tkde.2025.3552759>
- [19] Švábenský V, Flanagan B, López Zapata ED, Shimada A. Open datasets in learning analytics: trends, challenges, and best PRACTICE. *ACM Transactions on Knowledge Discovery from Data*. 2026 Apr 10; 20(4):1-56. <https://doi.org/10.1145/3798096>