Ingénierie des Systèmes d'Information
Vol. 31, No. 10, October, 2026, pp. 3405-3414
Journal homepage: http://iieta.org/journals/isi

A Theoretical Context-Aware Intelligent Tutoring Framework for Corporate Microlearning: An Interpretable Decision-Tree Approach

Najlaa Maimouni*, Yassine Zaoui Seghroucheni, Soumia Ziti

IPSS, Faculty of Sciences, Mohammed V University in Rabat, Rabat 10000, Morocco

Corresponding Author Email: najlaa_maimouni@um5.ac.ma

Copyright: ©2026 The authors. This article is published by IIETA and is licensed under the CC BY 4.0 license (http://creativecommons.org/licenses/by/4.0/).

https://doi.org/10.18280/isi.311014

Received: 20 June 2026
Revised: 11 September 2026
Accepted: 18 September 2026
Available online: 31 October 2026

Keywords:
corporate microlearning, context-aware learning, intelligent tutoring systems, learner-state modelling, decision-tree recommendation, adaptive learning architecture

ABSTRACT

Corporate microlearning requires learning objects to be selected under learner-specific and situational constraints, yet many adaptive-learning designs do not specify how these constraints are represented and translated into recommendations. This paper develops a theoretical architecture for a context-aware intelligent tutoring system (ITS) tailored to corporate microlearning. The framework is derived from a structured analysis of relevant literature and explicit criteria for selecting contextual dimensions. It represents learner state through four dimensions: cognitive and behavioral, environmental and contextual, temporal, and technical capacity. These dimensions are mapped onto a modular architecture comprising learner, domain, pedagogical, and interface components, with a defined state-and-event update cycle. Decision trees are positioned as an interpretable recommendation baseline and are examined through a qualitative comparison with collaborative filtering, neural networks, contextual bandits, and reinforcement learning. The recommendation process is further formalized using entropy and information gain to distinguish model training, inference, and learner-state feedback. The proposed variables and thresholds, including acoustic interference, mobility, ambient luminosity, available time, bandwidth, and device type, are design candidates rather than validated predictors. No learner dataset, trained model, deployment, or empirical benchmark is included. Accordingly, the paper contributes a testable architectural specification, not evidence of improved learning or retention. Future work should evaluate the framework using corporate learner traces and controlled field studies, including comparisons with alternative recommendation methods and validation of contextual variables.

1. INTRODUCTION

Professional learning has been reshaped by digital transformation. Organizations increasingly require continuous, distributed and adaptive skill development rather than concentrated periodic training [1]. Microlearning responds to this constraint by delivering short, self-contained learning episodes that can be integrated into the flow of work [2-4]. Distributed practice, retrieval practice and cognitive-load management provide established pedagogical motivations for short, spaced learning activities [5-7]. Prior research [1] provides the broader educational context for this work. The system architecture and recommendation design draw on the studies discussed below.

At enterprise scale, however, microlearning creates a content-navigation problem. Organizations may hold thousands of atomic learning objects, each targeting a specific competency and designed for reuse [8, 9]. Conventional recommender systems can use role, department, or interaction history [10], but context-aware learning research shows that recommendation can also depend on the learner's current knowledge, goals, available time, and situational conditions [11]. The present paper treats context as a design resource rather than as a demonstrated causal determinant of learning performance.

The proposed framework is deliberately positioned as a theoretical architecture. Its contribution is not the invention of decision trees, learner modelling or context-aware computing, but the integration of these established ideas into an explicit corporate-microlearning design and the formal specification of the associated recommendation logic. Four design contributions are made: (i) a four-dimensional representation of learner and situational state; (ii) a modular architecture with explicit interfaces between learner, domain, pedagogical and interface components; (iii) a qualitative comparison of candidate algorithm families, extended to contextual bandits and reinforcement learning; and (iv) an information-theoretic specification of how a decision tree could be trained and subsequently traversed. No deployment, user experiment or benchmark is reported.

This paper proposes a context-aware intelligent tutoring system (ITS) that combines contextual sensing with an interpretable decision-tree recommendation engine. Context is understood in the sense of Dey [12]. The system is intended to select a learning object that is feasible and context-compatible at each microlearning checkpoint. Importantly, the environmental thresholds and expected learning effects used in the framework are design assumptions or hypotheses to be

evaluated, not established performance results.

The paper is organized as follows. Section 2 describes the methodology used to develop the framework, including the literature basis, dimension-selection criteria, design synthesis and validation pathway. Section 3 defines the four-dimensional representation of the corporate learner. Section 4 presents the Modular Intelligent Network Design (MIND) architecture and its update cycle. Section 5 compares candidate algorithm families qualitatively. Section 6 positions decision trees as an interpretable baseline and states the associated design rationale. Section 7 formalizes training, inference and learner-state feedback. Section 8 discusses the architectural contribution, validation requirements and limitations. Section 9 concludes. Figure 1 provides a conceptual overview of the proposed framework, which has yet to be implemented or experimentally validated.

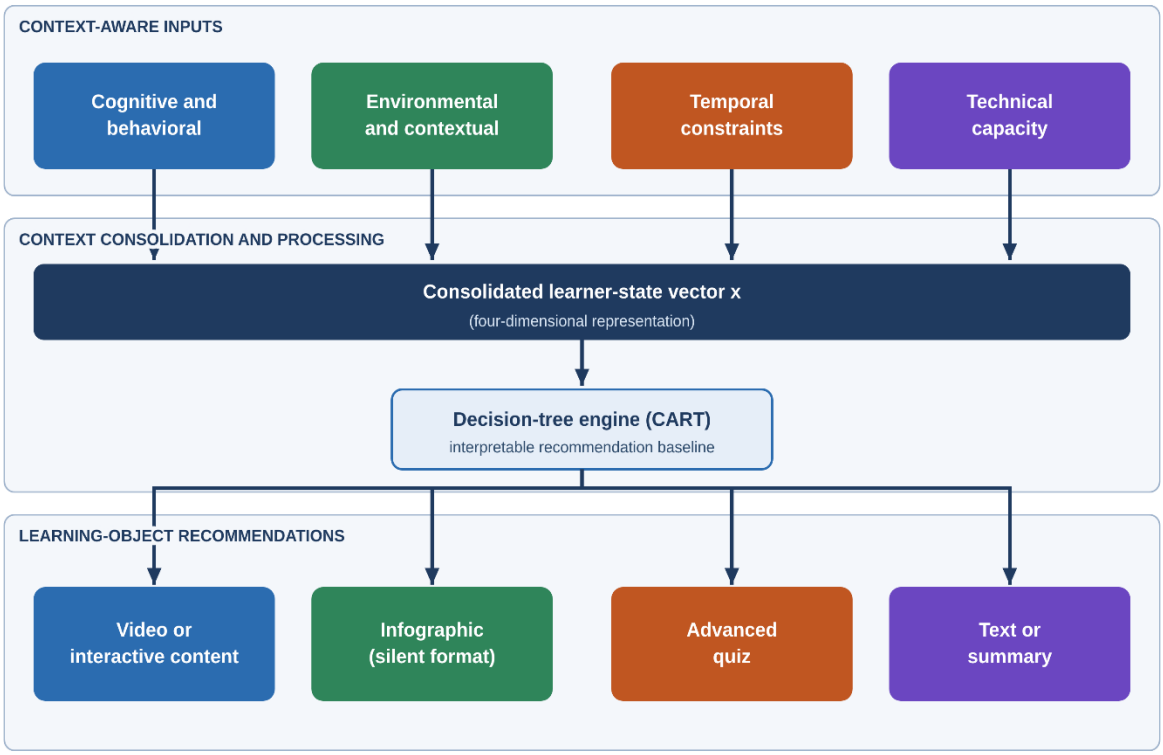


Figure 1. Conceptual research framework of the proposed context-aware intelligent tutoring system (ITS)

2. METHODOLOGY FOR FRAMEWORK DEVELOPMENT

2.1 Research objective and conceptual scope

The objective of the study is to construct a theoretical, testable architecture for context-aware corporate microlearning rather than to estimate a predictive model from a learner dataset. The framework therefore defines a candidate learner-state representation, a modular system architecture, a recommendation baseline, and an empirical validation pathway. The analysis is intentionally positioned as a conceptual design study: no corporate deployment, learner experiment or model benchmark is reported.

2.2 Literature analysis procedure

The framework was developed through a structured analysis of the literature already used in the manuscript. The analysis was organized around four evidence areas: learner modelling and knowledge-state estimation [13], context-aware and personalized learning [11, 12], ITS architecture [14], and learning-object, recommender-system and adaptive-delivery foundations [8-10]. Additional theoretical sources were used only where they address a specific architectural issue, such as contextual constraints [15, 16], interpretability [17], or sequential recommendation alternatives [18, 19]. The purpose of this analysis was to identify recurring design requirements and candidate contextual signals, not to claim a systematic literature review with quantitative screening or meta-analysis.

2.3 Criteria for contextual-dimension selection

Candidate variables were retained when they met the following design criteria: (i) direct relevance to the learner or delivery context at a recommendation checkpoint; (ii) observability or estimability in a future implementation; (iii) a plausible role in determining content feasibility, format compatibility or learner-state adaptation; and (iv) complementarity with the other variables, so that the representation captures distinct aspects of the decision context rather than multiple labels for the same construct. These criteria were applied conceptually to organize the candidate factors into cognitive and behavioral, environmental and contextual, temporal, and technical capacity dimensions. The resulting variables are therefore design candidates, not empirically validated predictors.

Table 1 summarizes the literature basis, selection rationale and evidential status of the four dimensions. It is intended as a design traceability device: each dimension is linked to a distinct architectural role, while the final column prevents the framework variables from being interpreted as validated predictors.

Table 1. Selection and design rationale for the proposed learner-state dimensions

Dimension and Candidate Variables	Literature Basis and Design Rationale	Status in Present Study
Cognitive and behavioral: knowledge mastery (Km), engagement decay (ε), and recent activity	Learner modelling and knowledge-state estimation [13]. Represents current learner state, readiness and interaction behavior for personalization decisions.	Conceptual; not empirically estimated here.
Environmental and contextual: acoustic interference (Ai), mobility index (Mx), and ambient illuminance (La)	Context-aware learning [11, 12]. International Organization for Standardization (ISO) 9921 [15] is used only for speech-communication conditions. Represents immediate situational constraints on feasible delivery formats.	Candidate signals; effects not validated here.
Temporal: available learning time (Δt) and time of day (ToD)	Timing/context rationale [20]. Represents episode-level feasibility and timing conditions at a recommendation checkpoint.	Candidate signals; effects not validated here.
Technical capacity: bandwidth (BW) and device type (DT)	Context-aware/adaptive delivery logic [11, 12]. Represents technical deliverability constraints for an otherwise suitable object.	Engineering constraints; thresholds require calibration.

2.4 Framework synthesis and design decisions

The selected dimensions were mapped to the modular MIND architecture. Cognitive and behavioral variables are exposed through the learner model; environmental, temporal and technical variables represent situational constraints; the domain model provides feasible learning objects and their metadata; and the pedagogical module applies the recommendation baseline. This mapping creates an explicit chain from contextual representation to candidate filtering and recommendation without assuming that any contextual variable has a demonstrated causal effect on learning outcomes.

2.5 Validation pathway

Validation is specified as a future empirical programme rather than a result of the present study. The planned pathway comprises offline comparison of the decision-tree baseline with alternative recommendation families, ablation of contextual dimensions, evaluation of recommendation and learning outcomes in a controlled field study, and separate testing of operational scalability and privacy or governance requirements. These elements are retained in the empirical validation plan in Section 8 and are intended to test, rather than presuppose, the theoretical advantages claimed by the architecture.

3. ADVANCED PARAMETERISATION OF THE CORPORATE LEARNER

To adapt delivery, a system must first hold a computable model of the learner's situation. We propose an ontological parameter space of four dimensions. These dimensions are design variables selected from prior work on learner modelling and context-aware learning [11-13]; their predictive relationship with corporate learning outcomes is not established by the present paper. The purpose of the representation is therefore to define a testable state space for a future implementation.

3.1 The cognitive and behavioural dimension

This dimension records what a learner knows and how ready that learner is to acquire new material. Two variables are proposed.

Knowledge mastery (Km $\in$ [0, 1]) is a continuous probability representing current competency in a knowledge area. In a future implementation, it could be updated after each micro-interaction by Bayesian Knowledge Tracing [13], which estimates a latent knowledge state from observed performance using initial-knowledge, learning, guessing, and slipping probabilities. In the present theoretical study, Km is a proposed state variable rather than an empirically estimated measure.

Engagement decay (ε) is proposed as an empirically estimable rate at which interaction with a module declines over time. It could be used to bound sequence length and calibrate subsequent difficulty, but no engagement dataset is available in this study; consequently, no value or predictive effect is claimed.

Recent activity may also be recorded, including selected format, interaction duration, and stated session goals. These observations could support inference of modality preferences and content relevance, but such preferences should be learned from data rather than assumed a priori.

3.2 The environmental and contextual dimension

Environmental conditions provide candidate just-in-time signals that distinguish the proposed architecture from purely profile-based adaptive learning [11, 12]. Three variables are included as hypotheses for future validation.

3.2.1 Acoustic interference (Ai, measured in decibels (dB))

Acoustic interference (Ai) can be captured by the device microphone. The original design used 70 dB as an illustrative engineering trigger for suppressing audio-dependent objects, informed by speech-intelligibility guidance in International Organization for Standardization (ISO) 9921 [15]. This value must not be interpreted as a validated pedagogical threshold: ISO 9921 addresses speech communication conditions, not the learning effectiveness of alternative formats. In a deployed system, the acoustic variable and any operational threshold should therefore be treated as design parameters to be calibrated against measured comprehension, task completion, and user experience.

3.2.2 Mobility index (Mx, categorical: Stationary, walking, running)

Mobility index (Mx) is conceptually derived from accelerometer and gyroscope data and classified as stationary, walking or running. The framework suggests that movement may enable voice-first or low-interaction formats, whereas stationary conditions may enable richer interaction. This relationship is to be tested.

3.2.3 Ambient illuminance (La, in lux)

The ambient illuminance (La) is measured by an ambient light sensor; we propose an adaptation to increase contrast and to reduce visually complex content under low light. These are accessibility-oriented design rules, not measures of learning performance.

3.3 The temporal dimension

In the corporate world, time is a limited resource. The available learning time (Δt, in minutes) is the time until the next commitment of the learner. It is proposed as the main feasibility filter. Only objects with estimated duration (stored in seconds and converted to minutes by dividing by 60) shorter than Δt enter the candidate pool. Time of day (ToD) is a contextual categorical feature as performance may depend on the synchrony between chronotype and timing of the task [16]. The framework does not suggest a universally optimal period. This relationship should be estimated at the learner level.

3.4 The technical capacity dimension

The dimension of technical capacity is involved in the deliverability of a pedagogically appropriate object. The original design's bandwidth categories (low <2 megabits per second (Mbps), medium 2–10 Mbps, and high >10 Mbps) and device type are retained. They set delivery constraints rather than validated thresholds for learning performance and should be calibrated to the actual content, network, and device environment.

4. THE MODULAR INTELLIGENT NETWORK DESIGN ARCHITECTURE

The MIND model structures the system as a collection of loosely coupled components. Each component has a well-defined interface so that it can be upgraded or replaced without changing the overall architecture.

This decomposition follows the classical learner-model, domain-model, pedagogical-module and interface-agent structure of ITS [14]. In the revised framework, the architecture is specified as a conceptual integration pattern rather than as a deployed system. Figure 2 depicts the resulting MIND architecture.

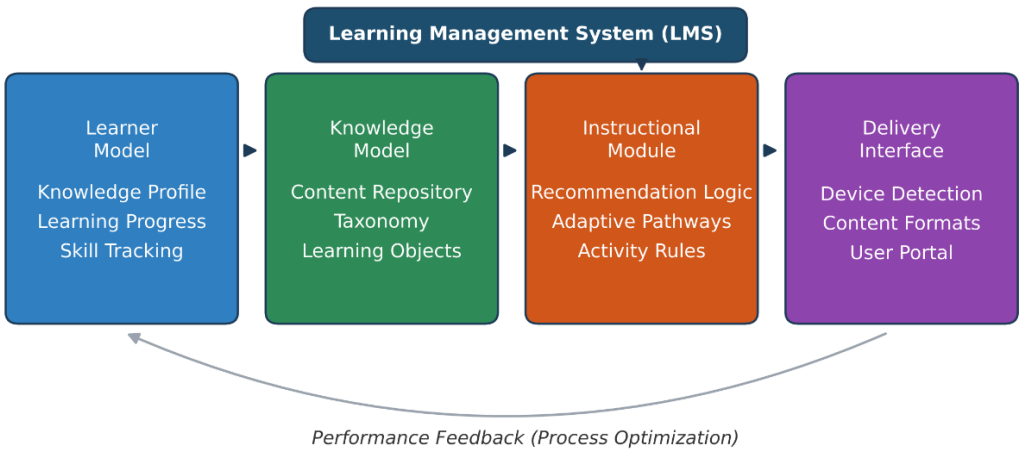


Figure 2. Conceptual Modular Intelligent Network Design (MIND) architecture of the proposed system

Figure 2 presents the conceptual architecture of the proposed MIND model; the interfaces and interactions shown have not been implemented or experimentally evaluated in this study.

4.1 The learner model

The learner model is a continually updated representation of each user's knowledge state, preferences, and interaction history. It is organized conceptually as a dynamic knowledge graph whose nodes represent competencies and whose edge weights represent inter-concept relations. After an interaction, a future implementation would update mastery using Bayesian Knowledge Tracing [13] and revise engagement estimates. The interaction event would also be appended to the learner timeline. The learner model would then expose the current state vector, x = [Km, ε, Ai, La, Mx, Δt, BW, DT, ToD], to the pedagogical module.

4.2 The domain model

The proposed ingestion pipeline and ontology-driven content graph [20, 21] are architectural specifications; their retrieval accuracy and metadata quality require empirical assessment. Table 2 lists the proposed metadata schema.

Table 2. Metadata schema of an atomic learning object

Metadata Field	Description	Example Value
difficulty_score	Semantic complexity, scaled to [0-1]	0.72
estimated_duration	Estimated time to complete the object, stored in seconds. The value is divided by 60 to convert it to minutes before comparison with Δt.	210 s (3.5 min)
prerequisite_ids	Objects required before this one	[ALO_042, ALO_055]
content_format	Delivery modality	video_interactive
topic_embedding	300-dimensional semantic vector	[0.12, −0.34, …]
last_updated	Timestamp of last revision	2025-03-01T08:00Z

4.3 The pedagogical module

The pedagogical module is the decision component. It receives the learner-state vector and a candidate list filtered by the domain model. It then traverses the decision tree specified in Section 7. The decision rule is evaluated at defined micro-checkpoints, such as a quiz submission or session start. After the learner interacts with the selected object, the resulting event is returned to the learner model. This feedback supports subsequent recommendations.

4.4 The interface agent

The interface agent converts an abstract recommendation into a deliverable learning experience. It reconciles the recommended format with current device and network constraints and applies graceful degradation and accessibility adaptations. It also returns delivery and interaction events to the learner model. Thus, inter-module communication is conceptualized as a cycle: state retrieval → candidate filtering → tree inference → interface adaptation → event logging → learner-state update. No implementation latency or interoperability test was conducted.

4.5 Inter-module communication and update cycle

At a conceptual level, the modules communicate through a state-and-event loop. The learner model exposes the current state vector. The domain model outputs learning objects with their metadata as feasible. The pedagogical module uses the trained tree, and the interface agent adapts the selected object to delivery constraints. Events are recorded and sent back to the learner model: completion, partial completion, quiz responses, abandonment. Therefore, recommendations are not reactive to each sensor change but change at certain predefined checkpoints. A production implementation would require explicit application programming interface (API) contracts, event schemas, consistency rules, and failure-handling procedures. None of these elements was evaluated in the present work.

Several algorithm families could support context-aware learning recommendations. Their suitability depends on the objective, available interaction data, need for online adaptation, and interpretability requirements. Recent work on learning context highlights the diversity of contextual indicators used in personalized learning [11], while reinforcement-learning research shows growing interest in sequential educational adaptation but also emphasizes methodological and deployment challenges [19]. Contextual bandits are another relevant paradigm because they select actions from contextual information while learning from partial feedback [18]. Table 3 therefore provides a qualitative design assessment, not an experimental ranking or performance benchmark. The ratings in Table 3 were assigned narratively, from the characteristics reported for each algorithm family in the cited literature (interpretability claims, reported data requirements, and support for online adaptation), rather than from a numerical scoring procedure or a common benchmark dataset; no evaluation source produced the labels directly, and the comparison should be read as a structured summary of qualitative literature evidence. Figure 3 visually summarizes the qualitative comparison presented in Table 3.

Table 3. Qualitative design comparison of candidate algorithm families

Algorithm Family	Interpretability	Interaction-Data Need	Online Adaptation	Inference Profile	Role in This Paper
Collaborative filtering	Low–moderate	High	Limited unless retrained	Moderate	Comparator
Neural networks (RNN/LSTM)	Low intrinsic	High	Possible	Higher	Comparator
Contextual bandits	Model-dependent	Moderate–high	Native	Potentially low	Future comparator [18]
Reinforcement learning	Model-dependent	High interaction demand	Native	Variable	Future comparator [19]
Decision trees (CART)	High intrinsic	Moderate for training	Limited unless updated	Low for fixed depth	Proposed baseline

Note: RNN = recurrent neural networks, LSTM = long short-term memory, CART = classification and regression trees.

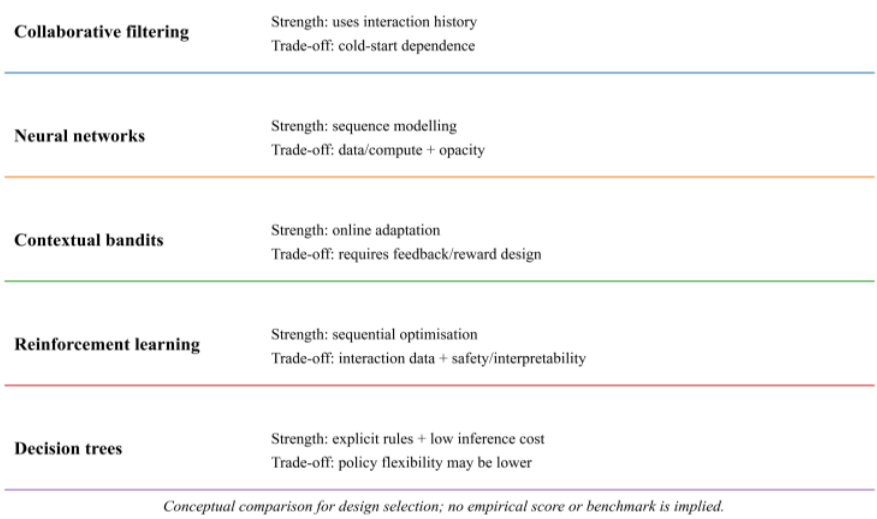


Figure 3. Conceptual comparison of candidate algorithm families

The comparison should consequently be interpreted as a design rationale rather than as evidence that decision trees outperform the other families. The proposed architecture uses decision trees as an interpretable baseline whose suitability can be tested experimentally against the alternatives in future work.

5. COMPARATIVE ANALYSIS OF ALGORITHM FAMILIES

5.1 Collaborative filtering

Collaborative filtering identifies similar users or items from interaction histories [10]. Its established strength is the ability to exploit collective behavior, but it generally depends on sufficient interaction data and can encounter cold-start problems [22]. In a learning domain, prerequisite relationships and pedagogical constraints may also require information beyond user-item similarity. These observations motivate comparison with, rather than rejection of, collaborative filtering in future experiments.

5.2 Neural networks

Recurrent neural network (RNN) and long short-term memory (LSTM) can model sequential learner behavior [23], and deep knowledge tracing demonstrated the potential of neural sequence models on large educational datasets [24]. Their possible advantages include richer temporal modelling and nonlinear interactions. The trade-offs are data requirements, computational cost, and decreased intrinsic interpretability. The present paper does not argue that these models are inferior. Their performance in the proposed corporate context remains an empirical question. Post-hoc attribution methods such as Shapley Additive Explanations (SHAP) [25] can aid explanation, but do not provide the same intrinsic rule representation as a tree.

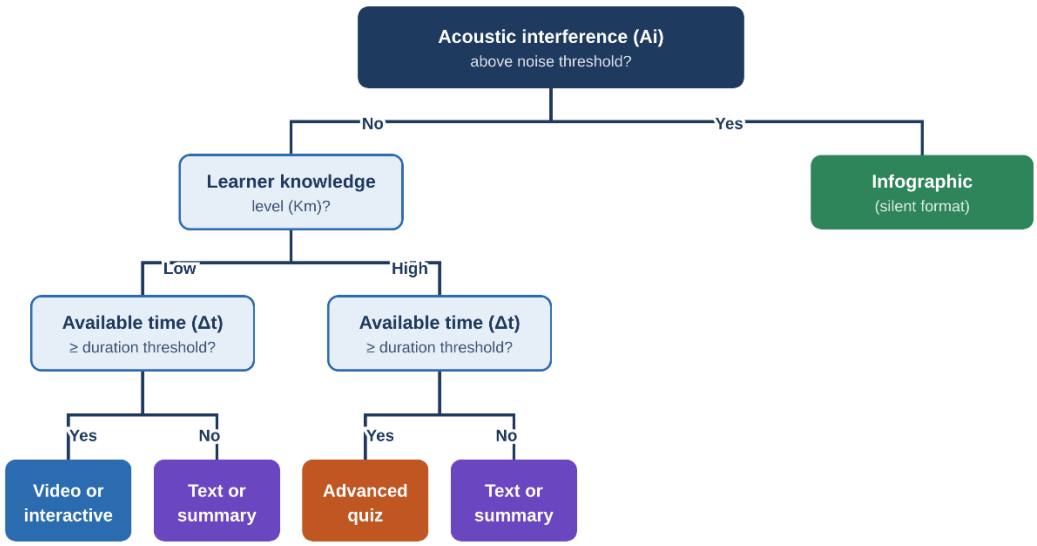


Figure 4. Illustrative decision tree for content-format recommendation

5.3 Decision trees

Decision trees partition the learner and context state space through a sequence of if–then tests, so that each terminal node corresponds to a recommended format or content category [26, 27]. In a classification and regression tree (CART), each split is chosen to reduce impurity, and depth and pruning control the balance between fit and generalization [26, 28]. Figure 4 illustrates a possible rule path for content-format recommendation based on three design signals from Section 3 (sound level, learner knowledge level and available time). The tree in Figure 4 is a hand-constructed conceptual example: it is not a trained or validated model, no information-gain values are reported, and its thresholds are placeholders to be estimated from interaction data in future work.

6. DECISION TREES AS AN INTERPRETABLE BASELINE

An example decision path for content-format recommendation is shown in Figure 4. Five theoretical considerations motivate the use of classification and regression trees (CART) as an interpretable baseline for this architecture. Each consideration addresses a design requirement rather than an empirical finding, and the suitability of the baseline remains to be tested against alternative approaches.

6.1 Explainability and regulatory compliance

Automated recommendations impacting an employee's development path should be auditable and understandable. A tree offers an intrinsic rule path that can be analyzed without a separate surrogate explanation [17]. The regulatory references [29, 30] motivate the focus on transparency, oversight by humans, and data governance, but the present paper does not claim that the proposed system is legally compliant in any specific jurisdiction.

6.2 Native handling of mixed feature types

The learner-state vector mixes continuous, categorical, and ordinal variables. Trees can partition such variables without

requiring cross-feature distance calculations or the normalization used by some distance-based methods. This is an architectural convenience, not evidence of better predictive performance.

6.3 Cold-start robustness

A trained tree can generate a recommendation based on current state variables even when no historical user-user similarity information is available [22]. This gives a plausible cold-start mechanism. However, the quality of the first recommendation depends on the validity of the rules and on the available context features; hence cold-start robustness should be evaluated empirically.

6.4 Computational efficiency

For a fixed tree depth d, traversal requires at most O(d) node evaluations. This is a computational property of tree inference, not a measurement of end-to-end system latency. Enterprise latency also includes authentication, database access, content retrieval, network communication, sensor processing, logging, and concurrent-user load. Consequently, the manuscript makes no claim that recommendation time is unchanged between 100 and 100,000 users; such a claim requires load testing.

6.5 Compatibility with prerequisite-based domains

Prerequisite relations can be incorporated into candidate filtering or tree rules so that insufficient mastery of a prerequisite routes the learner toward remedial material. This is a representational advantage, but it should be compared experimentally with graph-based sequencing and hybrid methods before any conclusion about pedagogical superiority is drawn.

7. TRAINING, INFERENCE AND FEEDBACK SPECIFICATION

We define a personalization function f(x), where x describes learner state and contextual conditions. The distinction between training, inference, and feedback is explicit. During training, entropy or impurity reduction can be used to select informative splits, and pruning or stopping criteria can be applied. During inference, a trained tree follows the learned conditional path for the current state rather than recomputing information gain for each learner. After an interaction, the learner-state variables and outcome events can be updated and retained for subsequent model retraining. The formalization below therefore separates the training criterion, inference procedure, and learner-state feedback loop.

7.1 Mathematical formalization

Let S be a training set of observations associated with learning-format outcomes. The uncertainty associated with the class distribution is quantified by Shannon entropy [31]:

$$H(S) = -\sum p_i \log_2 p_i \quad (1)$$

where, p_i is the probability of selecting format i from S. For a candidate feature a, the information gain is:

$$IG(S, a) = H(S) - \sum (|S_v| / |S|) H(S_v) \quad (2)$$

Eqs. (1) and (2) formalize the general splitting criterion used by CART. No learner dataset, labelled interaction sample or empirical split value is reported here: S denotes a generic training set, and the formalization illustrates the training principle rather than a fitted model. Reporting entropy or information-gain values computed on real interaction data is left to the future empirical validation programme described in Section 8.

7.2 Training, inference and learner-state feedback

The personalization process is specified in three conceptual stages. First, domain pruning constrains the candidate set by temporal feasibility and other hard delivery constraints. Second, a previously trained decision tree is applied to the current learner and context state, and its terminal node yields a recommended format or content category. Third, the interaction outcome is logged, and the learner-state representation is updated for future inference and model retraining. The following abstraction describes this process generically and does not represent an implemented system. Figure 5 illustrates the conceptual information-theoretic basis of a tree split underlying this process.

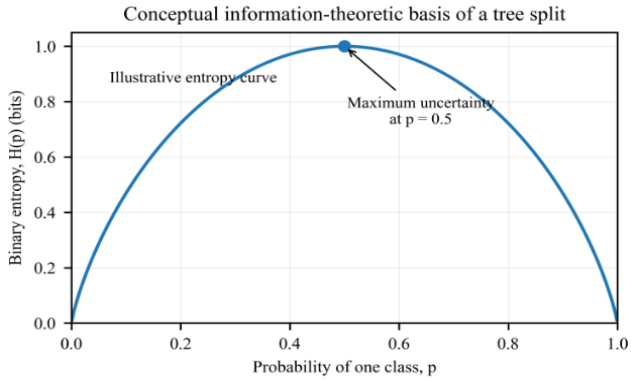


Figure 5. Conceptual information-theoretic basis of a tree split

A generic implementation would therefore train a tree from labelled interaction observations, apply the trained model at recommendation checkpoints, and record the resulting interaction outcomes for subsequent learner-state updates and model retraining. This separation makes clear that the proposed study does not define fixed hand-written recommendation rules and does not report a trained or validated tree.

8. DISCUSSION

8.1 Architectural contribution

The main contribution of the framework is architectural: it specifies how learner, environmental, temporal, and technical information could be represented and connected to a recommendation baseline within a modular intelligent tutoring architecture. Established learning principles [5-7, 32-34] motivate the broader idea of adaptive and context-compatible delivery, but they do not establish that the proposed architecture improves retention or other learning outcomes.

Such effects remain empirical questions for the validation programme described below. This qualification is applied consistently throughout the manuscript: the Abstract, Introduction, Sections 3–7 and the present discussion each state, at the point where a contextual signal, retention effect, algorithm comparison or scalability claim is introduced, that the corresponding evidence is analytical or architectural rather than experimental.

Accordingly, the expected effects of context-aware adaptation should be formulated as testable hypotheses rather than as anticipated outcomes. The present study does not estimate retention, engagement or recommendation effectiveness. The role of the architecture is to define a concrete state space, decision process and validation pathway through which these outcomes could be evaluated.

8.2 Organizational scalability

The architecture also provides explicit points for empirical validation. A fixed-depth tree has a traversal cost that depends primarily on tree depth, but this property does not imply constant end-to-end response time because database access, content retrieval, authentication, sensor processing, network conditions, logging and concurrency also contribute to system behavior. Likewise, the modular components and the proposed content metadata could facilitate incremental integration with systems supporting established packaging and tracking specifications [35, 36], but such integration and scalability remain to be tested.

8.3 Limitations and future work

The first and most important limitation is that the framework is theoretical. No corporate deployment, learner dataset, user study or controlled experiment is reported. Consequently, the paper does not establish improvements in retention, engagement, recommendation accuracy, scalability or training efficiency. The algorithm comparison is qualitative, and all contextual effects remain to be tested. Future evaluation should compare the decision-tree baseline with collaborative filtering, neural sequence models, contextual bandits and reinforcement learning using common datasets and, where feasible, a controlled corporate field study.

The second limitation is about the validity of models and thresholds. The proposed design signals (sound level, mobility, luminosity, available time, bandwidth and device type) are plausible, but their relationships with learning outcomes are not demonstrated. The 70 dB acoustic trigger is a speech-communication reference, not a validated learning threshold. Future work should estimate context-outcome relationships from learner traces, perform ablation studies, calibrate thresholds, and report uncertainty. Evaluation should include quality of recommendation, completeness, immediate and delayed knowledge tests, engagement, user satisfaction and fairness across relevant groups.

The third limitation concerns privacy and governance. Environmental and behavioral signals can constitute personal data and should be processed according to applicable law and organizational policy [29]. A practical deployment should obtain informed employee consent where required, define a clear purpose for each signal, minimize collection, process raw sensor data on device where feasible, transmit derived features rather than raw audio or location, apply pseudonymization or anonymization where appropriate, define retention and deletion periods, encrypt data in transit and at rest, enforce role-based access control, maintain audit logs and provide mechanisms for data access and correction. These controls are requirements for a future implementation, not properties already demonstrated by the prototype.

8.4 Empirical validation plan

A future empirical programme should evaluate the framework in two stages. First, offline experiments should use time-ordered learner traces to compare the decision-tree baseline with collaborative filtering, RNN/LSTM models, contextual-bandit policies and reinforcement-learning policies. The comparison should use identical candidate pools and content metadata and report recommendation accuracy or ranking quality, coverage, cold-start performance, calibration and computational cost. Context ablation experiments should quantify the incremental contribution of acoustic, mobility, luminosity, temporal and technical variables rather than assuming that every signal is useful. This design directly tests whether the proposed four-dimensional representation contributes information beyond simpler alternatives.

Second, a controlled field study with corporate learners should measure immediate and delayed knowledge outcomes, completion, time-on-task, engagement, perceived usefulness and user satisfaction. Delayed assessments at predefined intervals would permit direct estimation of retention rather than reliance on theoretical forgetting curves [37]. For scalability, load tests should separately measure tree inference, data retrieval, end-to-end latency and throughput under increasing concurrent users. Privacy and governance should be evaluated through a documented data-flow assessment covering consent, purpose limitation, retention, access control and deletion. These tests would determine whether the proposed architecture's theoretical advantages translate into measurable enterprise benefits.

9. CONCLUSION

We propose a theoretical framework for a context-aware ITS for corporate microlearning in this paper. The framework comprises a four-dimensional representation of the learner state, a modular MIND architecture, and an information-theoretic decision-tree personalization mechanism. The goal is to make design assumptions explicit and testable, not to claim a validated enterprise solution.

The decision-tree component is therefore defined as an interpretable baseline for empirical comparison, not as a demonstrated superior solution. Contextual bandits, reinforcement learning, collaborative filtering, and neural sequence models remain relevant alternatives for future benchmarking. The primary conclusion is methodological: the architecture defines a concrete, testable framework in which the contribution of individual contextual dimensions and the suitability of alternative recommendation paradigms can be evaluated empirically. Claims concerning retention, recommendation effectiveness, scalability and enterprise readiness remain hypotheses until such evaluation is conducted.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

REFERENCES

[1] Fischer, G. (2000). Lifelong learning—more than training. Journal of Interactive Learning Research, 11(3): 265-294. https://learntechlib.org/p/8380.

[2] Hug, T. (2007). Didactics of Microlearning: Concepts, Discourses and Examples. Waxmann, Münster.

[3] Leong, K., Sung, A., Au, D., Blanchard, C. (2021). A review of the trend of microlearning. Journal of Work-Applied Management, 13(1): 88-102. https://doi.org/10.1108/JWAM-10-2020-0044

[4] Shail, M.S. (2019). Using micro-learning on mobile applications to increase knowledge retention and work performance: A review of literature. Cureus, 11(8): e5307. https://doi.org/10.7759/cureus.5307

[5] Cepeda, N.J., Pashler, H., Vul, E., Wixted, J.T., Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. Psychological Bulletin, 132(3): 354-380. https://doi.org/10.1037/0033-2909.132.3.354

[6] Sweller, J., Ayres, P., Kalyuga, S. (2011). Cognitive Load Theory. Springer, New York.

[7] Roediger, H.L., Karpicke, J.D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. Psychological Science, 17(3): 249-255. https://doi.org/10.1111/j.1467-9280.2006.01693.x

[8] Polsani, P.R. (2003). Use and abuse of reusable learning objects. Journal of Digital Information, 3(4). https://jodi-ojs-tdl.tdl.org/jodi/article/view/jodi-105.

[9] Duval, E. (2001). Metadata standards: What, who and why. Journal of Universal Computer Science, 7(7): 591-601. https://www.jucs.org/jucs_7_7/metadata_standards_what_who/Duval_E.pdf.

[10] Ricci, F., Rokach, L., Shapira, B. (2015). Recommender Systems Handbook, 2nd ed. Springer, New York. https://doi.org/10.1007/978-1-4899-7637-6

[11] Abu-Rasheed, H., Weber, C., Fathi, M. (2023). Context based learning: A survey of contextual indicators for personalized and adaptive learning recommendations – A pedagogical and technical perspective. Frontiers in Education, 8: 1210968. https://doi.org/10.3389/feduc.2023.1210968

[12] Dey, A.K. (2001). Understanding and using context. Personal and Ubiquitous Computing, 5: 4-7. https://doi.org/10.1007/s007790170019

[13] Corbett, A.T., Anderson, J.R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge. User Modeling and User-Adapted Interaction, 4: 253-278. https://doi.org/10.1007/BF01099821

[14] Nkambou, R., Bourdeau, J., Mizoguchi, R. (2010). Advances in Intelligent Tutoring Systems. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-14363-2

[15] International Organization for Standardization. (2003). ISO 9921: Ergonomics — Assessment of speech communication. ISO, Geneva. https://www.iso.org/standard/33589.html.

[16] May, C.P., Hasher, L. (1998). Synchrony effects in inhibitory control over thought and action. Journal of Experimental Psychology: Human Perception and Performance, 24(2): 363-379. https://doi.org/10.1037/0096-1523.24.2.363

[17] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1: 206-215. https://doi.org/10.1038/s42256-019-0048-x

[18] Li, L., Chu, W., Langford, J., Schapire, R.E. (2010). A contextual-bandit approach to personalized news article recommendation. In Proceedings of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA, pp. 661-670. https://doi.org/10.1145/1772690.1772758

[19] Riedmann, A., Schaper, P., Lugrin, B. (2025). Reinforcement learning in education: A systematic literature review. International Journal of Artificial Intelligence in Education, 35: 2669-2723. https://doi.org/10.1007/s40593-025-00494-6

[20] Reimers, N., Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, pp. 3982-3992. https://doi.org/10.18653/v1/D19-1410

[21] Gruber, T.R. (1993). A translation approach to portable ontology specifications. Knowledge Acquisition, 5: 199-220. https://doi.org/10.1006/knac.1993.1008

[22] Schein, A.I., Popescul, A., Ungar, L.H., Pennock, D.M. (2002). Methods and metrics for cold-start recommendations. In Proceedings of the 25th ACM SIGIR Conference on Research and Development in Information Retrieval, Tampere, Finland, pp. 253-260. https://doi.org/10.1145/564376.564421

[23] Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8): 1735-1780. https://doi.org/10.1162/neco.1997.9.8.1735

[24] Piech, C., Bassen, J., Huang, J., et al. (2015). Deep knowledge tracing. In Proceedings of the 29th International Conference on Neural Information Processing Systems, Montreal, Canada, pp. 505-513. https://dl.acm.org/doi/10.5555/2969239.2969296.

[25] Lundberg, S.M., Lee, S.I. (2017). A unified approach to interpreting model predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, California, USA, pp. 4768-4777. https://dl.acm.org/doi/10.5555/3295222.3295230.

[26] Breiman, L., Friedman, J., Olshen, R.A., Stone, C.J. (1984). Classification and Regression Trees. Chapman and Hall/CRC, New York. https://doi.org/10.1201/9781315139470

[27] Quinlan, J.R. (1993). C4.5: Programs for Machine Learning. Morgan Kaufmann, San Francisco, CA, United States. https://dl.acm.org/doi/10.5555/583200.

[28] Hastie, T., Tibshirani, R., Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition. Springer, New York. https://doi.org/10.1007/978-0-387-84858-7

[29] European Parliament and Council. (2016). Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union, L119: 1-88. http://data.europa.eu/eli/reg/2016/679/oj.

[30] European Parliament and Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). Official Journal of the European Union. http://data.europa.eu/eli/reg/2024/1689/oj.

[31] Shannon, C.E. (1948). A mathematical theory of communication. The Bell System Technical Journal, 27(3): 379-423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

[32] Tulving, E., Thomson, D.M. (1973). Encoding specificity and retrieval processes in episodic memory. Psychological Review, 80(5): 352-373. https://doi.org/10.1037/h0020071

[33] Godden, D.R., Baddeley, A.D. (1975). Context-dependent memory in two natural environments: On land and underwater. British Journal of Psychology, 66(3): 325-331. https://doi.org/10.1111/j.2044-8295.1975.tb01468.x

[34] Mayer, R.E. (2019). Multimedia Learning, 3rd ed. Cambridge University Press, Cambridge.

[35] Rustici, M. (2009). SCORM 2004, 4th ed. ADL Initiative, Alexandria. https://scorm.com/blog/scorm-2004-4th-edition/.

[36] Advanced Distributed Learning. (2016). Experience API (xAPI) Specification, Version 1.0.3. ADL Initiative, Alexandria.

[37] Ebbinghaus, H. (1913). Memory: A Contribution to Experimental Psychology. Teachers College, Columbia University, New York. https://archive.org/details/memorycontributi00ebbiuoft/page/n5/mode/2up.

NOMENCLATURE

Ai	acoustic interference, dB
BW	bandwidth category, categorical
DT	device type, categorical
$f(x)$	personalisation function
Km	knowledge mastery probability, [0, 1]
La	ambient illuminance, lux
Mx	mobility index, categorical
ToD	time of day, categorical
Δt	available learning time, min
ε	engagement-decay coefficient, dimensionless