iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.81894

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

A Survey on AI-Powered Adaptive Learning Using Concept Mapping and Grounded Response Generation Technology

Prof. Nagarathna C, Arun Kumar A R, Bharath M, Girish DR, Maadhu Avati
Dept. of CSE Sapthagiri College of Engineering Bengaluru, India

Abstract: The proliferation of digital educational resources has intensified the demand for intelligent systems capable of understanding, organizing, and presenting knowledge in ways that align with individual learner needs. Conventional e-learning environments continue to rely on passive content delivery mechanisms, including static document repositories, manualrevision scheduling,andkeyword-basedretrieval,allof which failtocapturethesemanticstructureofa subject domain or the evolving retention state of a learner. This survey systematically examines the theoretical and technical foundations underlying intelligent adaptive learning, spanning Concept Graph construction, retrieval-supported generation, dense vector search, neural knowledge tracing, spaced-repetition scheduling, and AI-driven tutoring. Building on this review, the paper introduces Neuroweave, a comprehensive adaptive learning framework that ingests uploaded educational documents, extracts domain concepts, and organizes them into a living Concept Graph. Memory retention is quantified per concept through theNeuroweave Adaptive MemoryAlgorithm (NAMA), which weighs review frequency, quiz performance, graph-level reinforcement from connected concepts, and time-based exponential decay. Grounded question answering is delivered by coupling large language model generation with vector retrieval over thelearner's own studymaterial, ensuring factual accuracy and source transparency. The paper presents the architectural design, component interactions, algorithmic formulation, comparative analysis against related literature, identified limitations, and directions for future research.

Keywords: Adaptive learning, Concept Graph, retrieval-supported generation, semantic vector search, memory retention modeling, personalizededucation,spacedrepetition, knowledge tracing.

I. INTRODUCTION

Advances in computing infrastructure, large language models, and graph databases have opened transformative possibilities for educational technology. Despite these capabilities,themajorityof widelydeployedlearningplatforms remain fundamentallystatic: learnersaccess content sequenced by curriculum designers, progress through modules at a fixed pace, and receive generalized feedback that does not reflect individual knowledge states. This structural rigidity limits learningefficiency, particularlywhenstudentsmust preparefor domain-specific assessments under time pressure.

Graph-theoretic representations of knowledge provide a principled mechanism for encoding theinterdependencies that exist among academic concepts. A Concept Graph (KG) can capture that understanding of a prerequisite concept is foundational to comprehending advanced material, and that weakness in one node has cascading implications for semanticallyadjacentnodes. Thisrelationalstructureenablesa learning system to identify not merely which topic a student struggles with, but also which upstream gaps are responsible for downstream confusion.

Retrieval-supported generation (RSG) addresses a fundamental limitation of parametric language models in educational contexts: the tendency to produce plausible but unsupportedanswers.Byretrievingrelevantpassagesfrom the learner's uploaded documents before generation, RSG architectures tightly couple AI responses to verified source material. This grounding mechanism is especially valuable in high-stakes academic settings where correctness is non-negotiable.

Memorymodelingaddsathirdcriticaldimension. Cognitive science has long established that retention decays over time ina predictable pattern and can be reinforced throughstrategically timed review. Algorithms that incorporate these principlescan determinenotmerely whatastudenthasstudied, but whetherthatknowledgeisstillreliablyaccessibleandwhen review is most effective.

This survey reviews the relevant literature across these intersecting domains and presents Neuroweave, an integrated platform that synthesises Concept Graph construction, RSG-based question answering, and algorithmic memory modeling intoaunifiedadaptivelearningenvironment.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Theremainderof this paper is structured as follows: Section II discusses background and motivation; Section III surveys related literature; Section IV analyses existing systems; Section V describes the proposed system; Section VI details the methodology; Section VII discusses findings; Section VIII concludes; and Section IX lists references.

II. BACKGROUND AND MOTIVATION

A. HeterogeneityofLearnerProfiles

Learners entering a course differ substantially in prior exposure, cognitive processing speed, preferred modality, and recall durability. A one-size-fits-all curriculum necessarily underserves students at both ends of the proficiency spectrum. personalizedadaptivesystemshavethepotentialtoaddressthis heterogeneitybydynamicallyadjustingcontent sequencingand review timing to the inferred state of each individual.

B. ShortcomingsofContemporaryLearningPlatforms

Contemporary learning management systems (LMS) are primarily content warehouses. They excel at distributing materials and tracking completion metrics but offer little intelligence regarding conceptual relationships or memory trajectories. Standalone spaced-repetition tools such as Anki improverevision efficiencyformemorisation-heavytasksbut impose substantial manual overhead for card creation and require the learner to independently judge card difficulty.

General-purpose AI assistants can respond to subject-matter queriesbutfrequentlyproducefactuallyunanchoredresponses, particularlyfor specializedacademiccontentthatmaypostdate model training.

C. ConvergenceofConceptGraphs,RSG,andMemory Algorithms

The combination of Concept Graphs, retrieval-supported generation, and memory decay algorithms creates a mutually reinforcingfoundationforintelligentlearning.ConceptGraphs supply structured context for retrieval; retrieval grounds generation in verified source text; and memory algorithms prioritize which graph nodes require urgent review. This convergence is the conceptual basis of Neuroweave and motivates the literature examined in Section III.

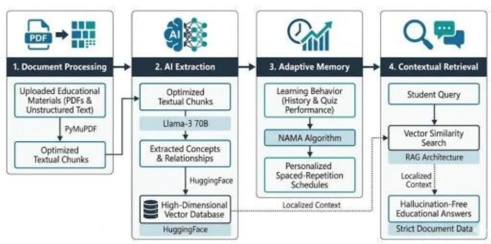

Fig.1.Thefour-phaseNeuroweavemethodology

III. LITERATURE SURVEY

A. Graph-BasedAdaptiveStudyPathDesign

AuthorsandYear:Li,Li,andGao(Electronics,2026)

Methodology: The authors conduct a structured review of Concept Graph approaches to study path design. Their taxonomy covers learner profile modeling, prerequisite extraction, and resource sequencing. They compare graph traversal strategiesandhighlighthowprerequisiterelationships can guide curriculum ordering automatically from ontological contentrepresentations.

Findings:Graph-basedpathdesignconsistentlyoutperforms purely collaborative filtering approaches when prerequisite structures are accurately captured. Systems that incorporate explicit dependency edges achieve higher learner satisfaction and lower dropout rates in controlled evaluations.

Limitation:Themajorityofreviewedsystemsrequirerichly annotated learner interaction logs that maynot be available in early deployment phases. Graph construction quality varies significantly across subjects, and transferability to disciplines with less formalized prerequisite structures remains an open challenge.

B. GroundedEducationalQuestionAnsweringwith Disciplinary Concept Maps

AuthorsandYear:MDKAGResearchGroup(Applied Sciences, 2025)

Methodology: This work constructs a multimodal disciplinary Concept Graph sourced from course syllabi, textbooks,andlecturetranscripts.Atquerytime,relevantgraph sub-graphsareretrieved andprovided as structuredcontexttoa language model fine-tuned on educational dialogue. Answer verification isperformed bycomparinggeneratedspansagainst retrieved source passages.

Findings:Groundinggenerationindomain-specificgraph context reduces hallucination rates by a measurable margin relative to graph-free RSG baselines. Question-answering accuracy improves particularly for multi-hop queries that require chaining across multiple conceptual dependencies.

Limitation: Constructing and maintaining a high-quality multimodal educational graph demands sustained curation effort. Automatic relationship extraction from diverse course materialsintroducesnoisethatcanpropagateintoretrievaland generation quality.

C. retrieval-supportedgenerationforKnowledge-Intensive NLP Tasks

AuthorsandYear:Lewisetal.(NeurIPS,2020)

Methodology:TheRSGframeworkcouplesadensepassage retriever trained with in-batch negatives to a sequence-to-sequencegenerator. Given an inputquery, theretriever returns the top-k passages from anon-parametric document index; the generator conditions on both the query and the retrieved context to produce the output. Two variants are evaluated: RSG-Sequence and RSG-Token, which differ in how retrieval is marginalized during generation.

Findings:RSGachievesnewstate-of-the-artperformanceon open-domain question answeringbenchmarksand outperforms purelyparametric baselines on knowledge-intensive tasks. The non-parametricmemorycomponentallows themodeltoaccess up-to-date information without retraining.

Limitation: Standard RSG operates over flat document chunksanddoesnotcapturehierarchicalorrelationalstructure between passages. In educational settings, this flat retrieval missesprerequisiterelationshipsthatarecriticalfor pedagogically appropriate responses.

D. Relationship-AwareRetrievalforLargeStudyCollections

AuthorsandYear:Edgeetal.(2024)

Methodology: This approach converts large document collections into entity-link structures using automated extraction. Graph communities are identified through hierarchical clustering, and community summaries are generated offline. At query time, relevant communities are rankedandtheirsummariesarepassedtothegenerationmodel, enablinganswers thatsynthesiseinformation distributedacross manydocuments.

Findings: For broad, thematic queries over large corpora, relationship-aware retrieval improves standard retrieval generation in comprehensiveness and diversity of response.Thecommunity-levelsummarizationcapturesemergentthemes that chunk-level retrieval cannot surface.

Limitation: Offline graph construction and community summarizationaddsignificantcomputationaloverheadand stoRSGe requirements. System quality is sensitive to the accuracy of LLM-assisted entity extraction, which can be unreliable for specialized academic terminology.

E. EfficientContextLinkingforRetrieval

AuthorsandYear:Guoetal.(2024)

Methodology: This method combines relationship-based search with vector matchingthrough tworetrieval paths. Local retrievalidentifiesentity-specificcontext,whileglobalretrieval identifies thematic relationships. The two paths are fused at inference time, balancing specificity with breadth.

Findings: Compared with plain retrieval generation, the method improves contextual coherence for complex queries requiring understanding of entity relationships. Retrieval latencyremainsacceptablefor interactiveapplicationsdueto lightweight graph indexing.

Limitation: Entity completeness depends on the quality of initialextraction.When sourcedocumentsarepoorlystructured or use domain-specific abbreviations, extracted graphs are sparse, reducing the advantage of relationship-aware retrieval over flat vector search.

F. DeepKnowledgeTracingwithRecurrentNeuralNetworks

AuthorsandYear:Piechetal.(NeurIPS,2015)

Methodology: Deep Knowledge Tracing (DKT) represents student knowledge as a latent hidden state maintained by a recurrentneuralnetwork.Ateachtimestep,themodelreceives the most recent exercise identifier and correctness signal and updates its internal state. The probability of correctly answering any future exercise is predicted from this hidden state.

Findings:DKToutperformsBayesianKnowledgeTracing on prediction accuracy across multiple learner datasets. The model captures complex temporal dependencies in learner behavior that parametric BKT models cannot represent.

Limitation: The latent state is not interpretable, making it difficult for educators to extract actionable insights about specific concept strengths. The model does not natively representprerequisiterelationshipsbetween skills,limitingits ability to recommend targeted interventions.

G. AdaptiveLearningThroughEducationalConceptGraph Integration

AuthorsandYear:Zhang,Wang, Ma,andWang(Frontiers in Educational Research, 2023)

Methodology: The proposed framework builds an educationalConceptGraphfromcurriculumdocumentsand links learner interaction records to graph nodes. Adaptive sequencingis achieved byselectingnext-learningresources based on mastered concepts and pending prerequisite gaps identified through graph traversal. A learner disorientation metric is introduced to detect confusion states and trigger remedial content delivery.

Findings: Students using the graph-adaptive system demonstratereduced studytimefor comparableachievement levels relative to linear curriculum controls. Disorientation eventsdecreaseover successivesessionsasthesystemlearns individual navigation preferences.

Limitation: The framework does not incorporate explicit memorydecaymodeling. As a result, the system mayover-estimate mastery of concepts reviewed long ago. Richer temporalretentionsignalsareneededtoimprovescheduling accuracy.

H. ConceptGraph-AugmentedKnowledgeTracingfor Programming Education

AuthorsandYear:AppliedSciencesResearchGroup (Applied Sciences, 2024)

Methodology: This study augments a knowledge tracing model with a programming-domain Concept Graph that encodes concept dependencies and common error patterns. personalized answer sequences are constructed by routing learnersthroughthegraphaccordingtotheirinferredmastery profile,prioritizingprerequisiteconceptsbeforeadvancingto dependent ones.

Findings: Graph-augmented knowledge tracing produces more accurate mastery predictions for programming competencies than graph-free baseline models. personalized exercisesequencesreducethenumber offailedattemptsbefore mastery is achieved.

Limitation: The approach is tuned to the structured prerequisiterelationships inherent in programming education. Its effectiveness in domains with less formalized dependency graphs,suchashumanitiesorinterdisciplinarystudies,hasnot been validated.

I. BGEDenseEmbeddingModelsforSemanticRetrieval

AuthorsandYear:BAAIFlagEmbeddingTeam(2023-2025)

Methodology:TheBGE(BAAIGeneralEmbedding)model family produces dense vector representations of text through contrastive training on large-scale pairs. BGE-small-en-v1.5 produces 384-dimensional embeddings and is optimized for resource-constrained deployment. Retrieval performance is evaluated on BEIR benchmarks covering diverse domain-specific corpora.

Findings: BGE models achieve competitive retrieval performance relative to larger embedding models while maintaininginferencelatencysuitableforreal-timeeducational applications. The small variant in particular offers a favorable accuracy-cost trade-off for edge and server deployment.

Limitation:Denseembeddingsimilaritymeasuressemantic proximity but does not ensure logical correctness or educational accuracy of retrieved passages. Incorrect or outdatedsourcematerial willberetrievedfaithfullyifit is semantically close to the query.

J. pgvector:VectorSimilaritySearchinRelationalDatabases

AuthorsandYear:pgvectorOpen-SourceProject(ongoing)

Methodology: pgvector extends PostgreSQL with a native vector data type and supports approximate nearest-neighbor search usingHNSWandIVFFlatindexes.Similarityfunctions include cosine distance, L2 distance, and inner product. The extension allows relationaland vector queries to be combined in standard SQL, eliminating the need for a separate vector store.

Findings: For moderately sized educational knowledge bases,pgvectorprovidesadequateretrievalthroughputwithout the operational overhead of a dedicated vector database.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Combiningrelational filteringwith vector search enables hybridqueriessuchasretrievingthemostsimilarconcept within a specific subject category.

Limitation:Atverylargescales,approximateindextuning becomes critical and retrieval latency can degrade. Default index parameters may not be optimal for domain-specific embedding distributions, requiring empirical calibration.

IV. EXISTING SYSTEM

Contemporarydigitallearning environmentsoccupya broad spectrumfrom simplefilerepositoriestosophisticatedlearning managementsystems,yetmostsharefundamentalarchitectural limitations that constrain their effectiveness for personalized adaptive learning.

Learningmanagementsystemssuch asMoodle,Canvas,and Blackboard organize course content, distribute assignments, and record submission events. Their internal data model is fundamentally document-centric: lectures, slides, and reading materials arestored as files with minimal semantic annotation. Progress tracking is limited to binary completion flags or assessment grades, providing no signal about conceptual understanding depth or retention durability over time.

Standalone spaced-repetition applications apply forgetting-curve scheduling to discrete flashcard items. These tools are effective for vocabulary and fact memorisation but require learners to manually create card sets, assign difficulty ratings, and maintain card decks as course content evolves. Critically, theytreat eachcardasanindependentmemoryunit,discarding the relational structure between concepts that could inform more intelligent review prioritization.

Keyword-based search engines retrieve documents containing matching terms but cannot capture semantic equivalences, identify conceptually related passages using different vocabulary, or rankresults bypedagogical relevance. Astudentsearchingfor aconceptusinginformallanguagemay fail to retrieve authoritative explanations phrased in technical terminology.

General-purpose AI assistants based on large language models can engage in subject-matter dialogue but suffer from parametric hallucination when responding to specialized queries outside their training distribution. They retain no memory of individual learner history across sessions and cannot anchor responses to specific uploaded study materials unlessexplicitlyprovidedwithcontextineachinteraction.The principal deficiencies of existing tools include absence of semantic concept modeling, lack of retention-aware scheduling,inabilitytorepresentprerequisiterelationships, high manual data-entry burden, and insufficient factual grounding of AI-generated responses.

V. PROPOSED SYSTEM

Neuroweave is an AI-powered adaptive learning platform engineeredtotransformuploadededucationaldocumentsintoa dynamic, queryable, retention-aware knowledge environment. The system addresses each identified limitation of existing toolsthroughacohesiveintegrationofmodernAIcomponents.

Upon document upload, Neuroweave extracts raw text, segments it into semantically coherent chunks, identifies key concepts and their relationships using a large language model, andstorestheresultingstructureddataasaConcept Graphina relational-vector database. This graph persists across sessions and continuously updates as the learner adds new materials, completes quizzes, and requests study plans.

The frontend is implemented in React 18 with TypeScript and Vite,providingaresponsiveinterfacestyled with Tailwind CSS and Motion for animations. Core interface modules include a personalized dashboard,an interactive 3D-style brain map for graph visualisation, an upload module, an insights panel displaying AI-generated recommendations, a study plannerwith weeklyschedules, amemoryheatmap,andanAsk Your Brain interface for grounded question answering.

ThebackendisaFastAPIapplicationdeployedon Render.It exposes REST endpoints for dashboard metrics, ConceptGraph data, document uploads, studyplan generation, memory analytics, AI insights, and RSG question answering. Supabase PostgreSQL with the pgvector extension serves as the primary data store, holding concept nodes, weighted relational edges, reviewlogs, uploadmetadata,and384-dimensionalBGE-small vector embeddings.

Language model inference for concept extraction and AI response generation is provided by the Groq API using the Llama-3.3-70B model, which delivers sub-second response latency. Embedding generation uses the BGE-small-en-v1.5 model locally, ensuring low-latency semantic search without external API dependency. The Neuroweave Adaptive Memory Algorithm(NAMA)continuouslyupdatesper-conceptmemory strength scores, driving weak-area detection, brain-map color coding, and study-plan priority ordering.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

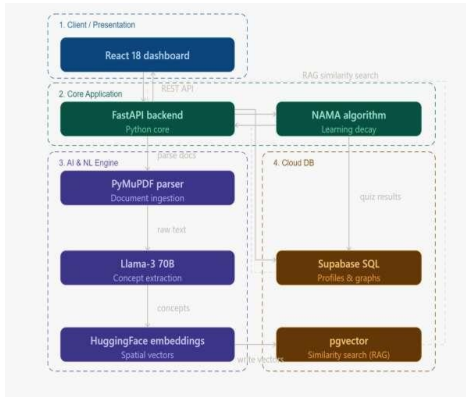

Fig.2.SystemArchitectureDiagramofNeuroweave

VI. METHODOLOGY

A. *UserInteractionandDashboardLayer*

Thelearner'sprimaryentrypointisapersonalizeddashboard that surfaces a holistic view of their learning state. Displayed metrics include an overall knowledge score, subject-level retention rates, the count of mastered versus weak concepts, scheduled upcoming reviews, and AI-generated study recommendations tailored to the current memory profile. The brain map provides an interactive visual rendering of the Concept Graph, where node color encodes NAMA-computed memorystrength:greenforhighretention,amberfor moderate, and red for critically weak concepts requiring immediate review.

B. *DocumentIngestionandTextExtraction*

The ingestion pipeline accepts PDF files, plain text documents, Markdown files, and directly pasted content. For PDF inputs, the system applies a multi-strategy extraction approach: native digital text extraction is attempted first, with an OCR fallback for scanned or image-based pages. Extracted textispre-processedtonormalizewhitespace,removeartifacts, and segment the document into overlapping chunks of configurable length. Overlapping windows ensurethat concept references spanning chunk boundaries are captured in at least one chunk.

C. *AI-AssistedConceptExtraction*

Eachtext chunkissubmittedtotheLlama-3.3-70Bmodel with a structured prompt instructing the model to identify domain concepts, assign categorical labels, summarise key ideas, and infer directional relationships between identified concepts. ThemodelreturnsastructuredJSONobject thatis validated against a schema before database insertion. In the event of API unavailabilityor response validation failure, a deterministic fallback using TF-IDF keyword extraction ensures pipeline continuity, albeit with reduced semantic richness.

D. *EmbeddingGenerationandVectorStoRSGe*

Concept labels, summaries, and source text snippets are independently embedded using BGE-small-en-v1.5, producing 384-dimensionalfloatvectors.Allembeddingsarestoredinthe pgvector-enabled Supabase PostgreSQL database alongside their associated concept records. An HNSW index is maintained on the embedding column to support approximate cosine-similarity queries at retrieval time. This vector infrastructure enables both concept relationship discovery during graph construction and document-grounded answer retrieval during RSG question answering.

E. *ConceptGraphConstructionandMaintenance*

Extracted concepts are stored as graph nodes with the following attribute schema: unique identifier, text label, domain category, source content snippet, current NAMA strength score,cumulativereviewcount,weighted quizscore, andembeddingvector.Edgesbetween conceptsarecreatedin two ways: explicitly, where the LLM identifies a directed relationship between two extracted concepts, and implicitly, where the cosine similarity between concept embeddings exceeds a configurable threshold. Each edge carries a weight reflecting relationship strength. The graph is incrementally extendedwitheachnewdocumentuploadwithoutduplicating nodes for concepts already present in the database.

F. NAMAMemoryStrengthComputation

The Neuroweave Adaptive Memory Algorithm computes a memory strength score S for each concept node by combining four weighted signals. The review component R_c increases logarithmically with cumulative review count, rewarding consistent engagement while exhibiting diminishing returns. The quiz performance component Q_c is the exponential moving aveRSGe of quiz scores on exercises linked to the concept, reflecting demonstrated recall accuracy. The graph reinforcement component G_c is computed as the weighted aveRSGestrengthofalldirectlyconnectednodes, capturingthe pedagogical principle that a concept embedded in a well-understood conceptual neighborhood is intrinsically more durable.ThetimedecaycomponentD_cappliesan exponential decay function to the interval since the most recent review, modeling the Ebbinghaus forgetting curve. These components are combinedas: S= w1 * R_c + w2 * Q_c+ w3 * G_c -w4 * D_c, where the weights w1 through w4 are globally configurable parameters subject to empirical calibration.
ConceptswithSbelowaweak-areathresholdareflaggedfor priorityscheduling.

G. RSG-BasedQuestionAnsweringandStudyPlan Generation

When a learner submits a question through the Ask Your Brain interface, the query text is embedded using BGE-small andcomparedagainstallstoredconceptembeddingsviacosine similarity. The top-k most similar concept nodes are retrieved along with their source text snippets. This retrieved context is formatted into a structured prompt that explicitly instructs the languagemodel to base its answer exclusively on theprovided passages and to cite the relevant concept when doing so. The study planner reads NAMA strength scores nightly and constructsaweeklyreviewschedulethatallocatessession slots toconcepts proportionallytotheir weakness severity, grouping related concepts within the same session to leveRSGe graph reinforcement effects during active review.

VII. DISCUSSION

Neuroweave advances the state of the art in adaptive learning by addressing the three structural deficiencies that characterise most existing platforms: absence of concept relationship modeling, inability to track retention trajectories, andlackoffactual grounding for AI-generatedresponses. The unified architecture ensures that each system component reinforces the others: Concept Graph edges provide the graph reinforcement signal consumed by NAMA, NAMA scores guidebothstudyplanningandbrain-mapvisualisation,andthe same vector embeddings that power graph construction also underpin RSG retrieval.
Theprimarystrength ofthegraph-centricdesign isitsability to surface second-order learning insights. A student who struggles with a specific concept can be shown not only the weak node but also the upstream prerequisite nodes whose incomplete mastery is likely responsible. This causal framing transforms remediation from a guessing exercise into a principled diagnostic activity.
Severalimplementationchallengesrequirecarefulattention.
Concept extraction accuracy depends substantially on the semanticcoherenceofsourcedocuments;looselystructured notes yield noisier graphs than well-organized textbooks.
Relationshipinferenceviaembeddingsimilarityintroduces false edges when domain-adjacent concepts are semantically close but pedagogically independent. The current NAMA weightconfigurationwasderivedanalyticallyratherthanfrom empirical learner data, and may not optimally reflect the memory dynamics of all subject types or learner profiles.
Privacyisa salientconcern given thatthesystem processes andstoreslearners'personalstudydocuments.Futureversions mustimplementrobust accesscontrols,end-to-endencryption for uploaded content, and clear data retention policies.
Evaluation of AI-generated answers against expert-curated groundtruthisalsoneededtoquantifyhallucinationratesin practice and inform further prompt engineering.

VIII. CONCLUSION

This survey has examined the theoretical landscape of AI-powered adaptive learning by reviewing foundational contributions in Concept Graph construction, retrieval-supportedgeneration,vectorsearch,neuralknowledgetracing, and memory modeling. Existing work demonstrates the individual value of each component but rarelyintegrates them into a cohesive, deployment-ready system accessible to individual learners without institutional infrastructure.
Neuroweave addresses this gap by providing an end-to-end platform that transforms uploaded documents into a living ConceptGraph,computesper-conceptretentionscoresthrough NAMA, and grounds AI-generated answers in verified source material via vector retrieval and large language model generation. The system reduces learner effort for material organisation, increases revision targeting precision, and provides transparent, source-anchored responses to subject-matter queries.

Avenues for future research include longitudinal validation of NAMA with diverse real-world learner cohorts, automated weightcalibrationthroughreinforcementlearningfromlearner outcome feedback, multimodal Concept Graph construction from lecture videos and diagrammatic content, privacy-preservingfederatedlearningacrossinstitutionaldeployments, andintegrationofreal-timecollaborativelearningfeaturesthat allow peer Concept Graphs to inform individual study recommendations.

REFERENCES

[1] A.Lietal.,Electronics,vol.15,no. 1,2026.
[2] P. Lewisetal.,NeurIPS, vol. 33,pp.9459-9474,2020.
[3] Recentstudyonrelationship-awareretrieval,2024.\n\n[4]Recent study on efficient context linking, 2024.
[4] C.Piech,J.Bassen,J.Huang,S.Ganguli,M.Sahami,L.J.Guibas, and J. Sohl-Dickstein, "Deep Knowledge Tracing," Advances in Neural Information Processing Systems, vol. 28, 2015.
[5] S.Zhang,X.Wang,Y.Ma,andD.Wang,"An AdaptiveLearning Method Based on Concept Graph," Frontiers in Educational Research, vol. 6, no. 6, 2023.
[6] BAAIFlagEmbeddingTeam,embeddingmodeldocumentation, 2023.
[7] PostgreSQLvector-searchextensiondocumentation.
[8] J.Shietal.,AppliedSciences,vol.15,no.16,2025.
[9] Y.Zhouetal., AppliedSciences,vol.14, no.17, 2024.
[10] A.C.Graesser,S.Lu,G.T.Jackson,H.H.Mitchell,M.Ventura,
[11] Olney, and M. M. Louwerse, "AutoTutor: A Tutor with Dialogue in Natural Language," Behavior Research Methods, Instruments,andComputers,vol.36,no.2,pp.180-193,2004.
[12] M.VanLehn,"TheRelativeEffectivenessofHumanTutoring, Intelligent Tutoring Systems, and Other Tutoring Systems," Educational Psychologist, vol. 46, no. 4, pp. 197-221, 2011.