

GNN-Enhanced Adaptive Concept Learning Module

R26-IT-006

Project Proposal Report

H. H. Jayasundera - IT22084590

B. Sc. (Hons) Degree in Information Technology Specialized in
Information Technology

Department of Information Technology
Sri Lanka Institute of Information Technology

March 2025

DECLARATION

I declare that this is my own work, and this proposal does not knowingly incorporate any material previously submitted for a degree or a diploma at any other university or higher education institution, nor does it contain any material that has been previously published or written by another person apart from where proper acknowledgment is given with the text.

Name	Student ID	Signature
H. H. Jayasundera	IT22084590	

The above candidate is conducting research for their undergraduate dissertation under my supervision.

Name	Date	Signature
Ms. Manori Gamage	20/03/2026	

ABSTARCT

Children with Autism Spectrum Disorder (ASD) often experience difficulties in developing language, communication, and comprehension skills due to differences in cognitive processing, attention patterns, and sensory sensitivity. In Sri Lanka, this challenge is further intensified because English is typically taught as a second language, requiring additional effort in vocabulary acquisition, pronunciation, reading, and concept comprehension. Existing teaching approaches rely largely on traditional classroom methods and generic digital tools that are not aligned with the Sri Lankan curriculum and do not support adaptive learning or teacher-mediated instruction. Additionally, most available autism learning applications operate within a parent–child model and lack structured instructional progression, personalized learning mechanisms, and meaningful analytics that help teachers monitor student progress.

This research proposes an adaptive digital learning support system designed to assist autistic children up to Grade 3 level in acquiring foundational English language concepts within a structured and personalized learning environment. The system introduces a three-tier concept learning architecture consisting of visual familiarization, single word recognition, and attribute introduction to ensure progressive mastery of each concept. Adaptive learning decisions are supported by a Graph Neural Network operating over a Graph Knowledge Base that models' relationships between concepts, activities, and student interaction patterns. The system dynamically adjusts activity difficulty, identifies concept confusion patterns, and recommends personalized learning paths. A rule-based scaffolding engine grounded in Applied Behaviour Analysis ensures reliable mastery progression, while a Large Language Model generates plain-language insights that help teachers interpret student performance.

The system will be implemented using React Native, Node.js, FastAPI, PostgreSQL, GNN-based learning analytics, and LLM. The expected outcome is a curriculum-aligned adaptive learning platform that supports autistic learners while enabling teachers to make data-driven instructional decisions, contributing to inclusive education in Sri Lanka.

ACKNOWLEDGMENT

I would like to express my sincere gratitude to my supervisors for their continuous guidance, encouragement, and support throughout the course of this research. I also extend my appreciation to the staff and educators who contributed valuable insights during the requirement gathering and development phases. I am thankful to my project group members for their collaboration, dedication, and teamwork throughout this study. Finally, I would like to acknowledge my family for their unwavering support, understanding, and encouragement throughout this journey

TABLE OF CONTENT

DECLARATION.....	2
ABSTARCT	3
ACKNOWLEDGMENT	4
TABLE OF CONTENT	5
LIST OF TABLES.....	7
LIST OF FIGURES	8
LIST OF ABBREVIATIONS.....	9
LIST OF APPENDICES.....	10
1. INTRODUCTION.....	11
1.1. BACKGROUND AND CONTEXT	11
1.1.1. <i>Real World Problem</i>	11
1.1.2. <i>Who Is Affected and Why It Matters</i>	11
1.1.3. <i>Statistical Evidence</i>	11
1.1.4. <i>Existing Solutions and Their Limitations</i>	12
1.1.5. <i>Proposed Solution and Its Components</i>	12
1.1.6. <i>Relevant Sustainable Development Goals</i>	13
1.1.7. <i>Research Cluster Alignment</i>	13
1.2. LITERATURE REVIEW	14
1.2.1. <i>Behavioural and Cognitive Frameworks Underlying ASD Learning Tools</i>	14
1.2.2. <i>AI and Machine Learning Approaches</i>	15
1.2.3. <i>AAC and Communication-Focused Research</i>	15
1.2.4. <i>User-Centred Design and Scheduling Applications</i>	16
1.2.5. <i>Existing Products and Their Features</i>	16
1.2.6. <i>Limitations of Existing Systems</i>	18
1.3. RESEARCH GAP	19
2. OBJECTIVES	21
2.1. MAIN OBJECTIVE	21
2.2. SPECIFIC OBJECTIVES.....	21
3. METHODOLOGY	23
3.1 RESEARCH APPROACH	23
3.2 TECHNICAL METHODS AND ALGORITHMS.....	23
3.3 TOOLS AND PLATFORMS.....	24

3.4 DATA COLLECTION AND ETHICAL CONSIDERATIONS	26
3.5 VALIDATION STRATEGY AND PERFORMANCE METRICS	27
4. HIGH-LEVEL SYSTEM ARCHITECTURE	29
5. USER REQUIREMENTS.....	30
3.1. USERS AND STAKEHOLDERS	30
3.2. FUNCTIONAL REQUIREMENTS	30
3.3. NON-FUNCTIONAL REQUIREMENTS	33
6. COMMERCIALIZATION PLAN	36
6.1. TARGET MARKET AND CUSTOMER PERSONA	36
6.2. VALUE PROPOSITION.....	36
6.3. REVENUE MODEL.....	37
6.4. COST ESTIMATION	37
6.5. PRICING STRATEGY	37
6.6. COMPETITIVE ADVANTAGE.....	38
6.7. INTELLECTUAL PROPERTY CONSIDERATIONS.....	38
7. BUDGET AND JUSTIFICATION	39
8. WORK BREAKDOWN STRUCTURE.....	41
9. GANTT CHART	42
10. REFERENCES LIST.....	43
11. APPENDICES	45

LIST OF TABLES

Table 1 - Limitations of Existing Systems	18
Table 2 - Proposed Institutional Subscription Tiers for the System.....	37
Table 3 - Estimated Research Project Budget.....	39
Table 4 - Task Breakdown	41

LIST OF FIGURES

Figure 1 - High Level Architecture of the Adaptive Concept Learning Support System for Autistic Children.....	29
Figure 2 - GANTT Chart	42
Figure 3 - Appendix A1 - Client Requirements Document.....	47

LIST OF ABBREVIATIONS

Abbreviation	Definition
AAC	Augmentative and Alternative Communication
ABA	Applied Behaviour Analysis
API	Application Programming Interface
ASD	Autism Spectrum Disorder
CAT-Q	Camouflaging Autistic Traits Questionnaire
CNN	Convolutional Neural Network
GKB	Graph Knowledge Base
GNN	Graph Neural Network
LLM	Large Language Model
PESC	Picture Exchange Communication System
PRT	Pivotal Response Treatment
R-CNN	Region-based Convolutional Neural Network
SDG	Sustainable Development Goals
SST	Software Systems and Technologies
TA	Task Analysis
UI	User Interface
UX	User Experience

LIST OF APPENDICES

1. INTRODUCTION

1.1. Background and Context

1.1.1. Real World Problem

Children with ASD often experience difficulties in developing language, communication, and comprehension skills due to differences in cognitive processing, attention, and sensory sensitivity. In Sri Lanka, this challenge is further compounded because English is typically taught as a second language, requiring additional effort in vocabulary acquisition, pronunciation, reading, and understanding basic concepts. Traditional classroom teaching methods and existing learning materials are largely designed for typically developing students and do not adequately accommodate the unique learning needs of autistic children. Furthermore, special education teachers often manage several students with varying abilities, making it difficult to provide personalized instruction tailored to each child's learning pace and preferences. As a result, many autistic children struggle to effectively learn foundational English language concepts and may find it difficult to participate in mainstream educational assessments and government examinations.

1.1.2. Who Is Affected and Why It Matters

The primary groups affected are children with ASD and special education teachers in Sri Lanka. Most autistic children currently rely on traditional classroom teaching methods, printed materials, and general educational tools primarily designed for typically developing students. These resources lack structured visual support, adaptive learning mechanisms, and individualized pacing that are critical for autistic learners. Special education teachers face the additional challenge of supporting multiple students with different learning needs without adequate tools to personalize instruction or monitor individual progress. This highlights the need for specialized educational systems capable of relational, cross-session learning, systems that can reason not just about a single student's isolated performance, but about how that student's confusion patterns relate to the structure of the concept domain itself.

1.1.3. Statistical Evidence

Globally, autism is a significant and growing concern in education systems. Studies estimate that approximately 1 in 127 children worldwide are diagnosed with ASD [1], highlighting the scale of the population that may require specialized learning support. In Sri Lanka, research suggests that about 1 in 93 children may be affected by autism [2], indicating a comparable or slightly higher prevalence within the region. Recent reports also indicate that over 9,000 children have been identified as autistic in Sri Lanka [3], demonstrating the growing number of learners who may require specialized educational assistance. These statistics highlight the increasing need for

effective and accessible learning solutions that support autistic children in developing essential academic and language skills.

1.1.4. Existing Solutions and Their Limitations

Currently, autistic children in Sri Lanka mainly rely on traditional classroom teaching methods, printed learning materials, and limited digital educational tools to learn English language fundamentals. While special education teachers attempt to adapt teaching strategies using visual aids, flashcards, and repetitive instruction, these approaches often lack structured personalization and adaptive learning mechanisms. Existing digital learning applications for autism are typically developed for global audiences and does not align with the Sri Lankan English curriculum, cultural context, or language learning needs of local students. Furthermore, many available tools do not provide teacher monitoring features, individualized progress tracking, or adaptive difficulty adjustment based on student performance. As a result, teachers face challenges in managing multiple students with different learning abilities, and autistic children may struggle to remain engaged, retain concepts, and progress effectively through foundational English language learning.

1.1.5. Proposed Solution and Its Components

The proposed solution is an adaptive digital learning support system designed to support English language concept acquisition among autistic children up to Grade 3 level in Sri Lanka, addressing the absence of contextually appropriate tools aligned to the local curriculum and special education instructional model. The system establishes a structured, personalised learning environment in which special education teachers oversee and navigate sessions while children directly interact with the learning activities. The core of the system is a student workspace through which teachers initiate sessions using curriculum aligned learning modules covering concept learning, pronunciation practice, writing activities, and conversational exercises. At the heart of the system's adaptive intelligence is a GNN operating over GKB. The GKB models the learning domain as an interconnected graph structure that represents concepts, categories, activities, and student interactions. The static curriculum layer of this graph is constructed during the design phase using the National Institute of Education Grade 1–2 English curriculum, while the dynamic interaction layer is continuously updated from student session logs so that the graph reflects the evolving learning history and behavioural patterns of each student.

The GNN performs inference over this graph to detect confusion clusters between concepts, recommend appropriate learning paths, predict concept difficulty for individual students, and identify behavioural similarities across learners. These insights are provided to the system's Content Adaptation Engine, which dynamically assembles learning activities from a modular asset library containing concept images, audio clips, distractor pools, and activity templates rather than serving static prebuilt exercises. This enables the system to personalise learning activities based on each child's observed learning patterns and interaction behaviour. Prior to a student's first

session, a onetime initial assessment captures each child's sensory and interface preferences including avatar choice, visual theme, and font style and size ensuring the interface remains consistent, predictable, and sensorially comfortable across all future sessions. The system also incorporates a rule based adaptive scaffolding mechanism grounded in Applied Behaviour Analysis principles which adjusts content difficulty, distractor levels, and visual supports in response to each child in session performance while enforcing mastery criteria before progression to new concepts. Structured evaluation activities are conducted after every five completed concepts to measure concept retention and reinforce learning. To further support teacher decision making, a large language model component interprets accumulated student interaction data together with insights generated by the GNN and produces natural language feedback for teachers after each session and evaluation. Delivered through the Gemini API, this AI assisted decision support layer provides session summaries, detected performance patterns, flagged areas of difficulty, and suggested instructional adaptations enabling teachers to make informed, data driven instructional decisions without needing to manually analyse raw performance logs.

1.1.6. Relevant Sustainable Development Goals

This research aligns with several United Nations SDGs, most prominently SDG 4 on Quality Education, which collectively call for inclusive, equitable, and disability-sensitive quality education for all children including those from vulnerable groups. By designing a curriculum-aligned, personalised digital learning system for autistic children within the Sri Lankan special education context, this research directly addresses the gap in quality English language instruction available to a population that is frequently excluded from mainstream educational support. The work further aligns with SDG 10 on Reduced Inequalities, by reducing the educational disparity between autistic learners and their neurotypical peers and addressing the broader inequity between the educational technology resources available in high-income countries and those accessible to teachers and students in Sri Lanka. SDG 3 on Good Health and Well-Being is also relevant as the system is designed with the sensory and emotional well-being of autistic children as a core design principle, incorporating calming visual environments, predictable interface behaviour, and supportive error handling mechanisms that minimise distress and promote psychological safety during learning sessions. Finally, the research aligns with SDG 17, by demonstrating the application of information and communication technology as an enabling tool for inclusive education in a developing country context, contributing to the broader goal of leveraging digital solutions to advance equitable development outcomes in underserved communities.

1.1.7. Research Cluster Alignment

This research aligns with the SST research cluster, involving the design and development of a structured digital learning system encompassing user management, teacher and student workspaces, curriculum-aligned learning modules, a GNN-powered graph intelligence layer, a rule-based adaptive scaffolding engine, and an LLM-based teacher decision support layer. The

work addresses core software systems concerns including relational data modelling, graph-based inference, personalization through persistent user preference management, performance-responsive dynamic content delivery, and learning analytics. The introduction of a GNN operating over a knowledge graph represents a substantive software systems contribution, requiring graph schema design and real-time edge weight management into a live educational application.

1.2. Literature Review

The intersection of education technology and ASD has attracted significant scholarly attention over the past decade, driven by the growing global prevalence of ASD and the persistent gap between the needs of autistic learners and the tools available to support them. The following literature review examines the recent scholarly sources alongside existing commercial products to synthesise current research approaches, evaluate the state of existing systems, and identify the limitations that the proposed system seeks to address.

1.2.1. Behavioural and Cognitive Frameworks Underlying ASD Learning Tools

A foundational contribution to understanding how autistic learners differ from neurotypical peers has been provided by Hariharan et al. [4]. Their comparative study of 196 autistic and 196 healthy adults using the CAT-Q documents significant behavioural differences across masking, compensation, and assimilation dimensions. Critically, their findings confirm that autistic individuals rely heavily on structured external supports, show low comfort in unstructured social contexts, and demonstrate a strong tendency to learn through observation and repetition rather than spontaneous social exchange. These behavioural characteristics directly inform pedagogical design decisions: structured, predictable learning sequences, repeated audio-visual exposure, and guided feedback mechanisms are not merely helpful features, but necessities grounded in the documented behavioural profile of autistic learners.

Wang and Wang [5], in their comprehensive review of ASD etiology, diagnosis, and therapy, consolidate decades of intervention research and identify ABA as one of the most extensively evidenced therapeutic frameworks for autism. They highlight that ABA-based interventions, including discrete trial training, pivotal response treatment, and early intensive behavioural intervention, share a core commitment to breaking target skills into small, sequentially mastered steps, using reinforcement to consolidate learning, and maintaining high mastery thresholds to ensure genuine retention. This aligns directly with the mastery criterion design of the proposed system, which requires 90% accuracy across consecutive trial blocks before a child progress between tiers, a threshold explicitly grounded in ABA literature by Granpeesheh et al. [6] and Fuller and Fienup [7].

1.2.2. AI and Machine Learning Approaches

Vijayan et al. [8] propose a three-layer cognitive computing architecture for an intelligent learning assistant for children with ASD. Their framework integrates an application layer with a chatbot and visual aids using automatic speech recognition and natural language processing, a data engineering layer processing facial images, brain images, phenotype data, and courseware, and an AI layer using R-CNN for facial pattern recognition, deep CNN for brain image classification, and reinforcement learning with deep Q-learning for adaptive courseware delivery. Their system represents an ambitious multi-modal approach to personalisation, demonstrating that AI-driven adaptive learning for ASD populations is technically feasible and theoretically well-motivated. However, their framework remains largely architectural, with limited evidence of empirical validation with autistic children in real classroom settings. The proposed system draws on a comparable commitment to personalised content delivery while deliberately limiting its scope to what can be practically implemented and validated in the Sri Lankan special education context.

Hayat et al. [9] apply machine learning and deep learning approaches to the problem of personalised teaching strategy suggestion for toddlers with ASD. Using a combined dataset of 2,242 instances and feature selection via ANOVA F-values, they identify executive function, language and communication, and eye contact as the strongest predictors of learning strategy suitability. Their models, including Decision Tree, Random Forest, Feedforward Neural Networks, and Deep Neural Networks, achieve accuracy rates of up to 100% on the combined dataset, and they map outputs to specific teaching strategies including PRT, TA and PECS. This work demonstrates the feasibility of data-driven strategy assignment, though the authors acknowledge that their models are trained on existing datasets rather than prospectively validated in live teaching environments, and that ground-truth labels for teaching strategy appropriateness remain dependent on clinical judgement.

Ganju et al. [10] present CogniQuest, a CNN-based iOS application for neurodivergent children aged six to twelve that uses handwriting recognition to assess and enhance cognitive skills. Their architecture employs convolutional and pooling layers trained on the Kaggle A-Z handwritten alphabet dataset, achieving a training accuracy of 98.85% and a validation accuracy of 97.70%. While their technical results are strong, their user evaluation is limited to three neurotypical children, raising significant questions about generalisation to autistic populations, whose motor profiles, cognitive processing styles, and engagement patterns differ substantially from neurotypical peers.

1.2.3. AAC and Communication-Focused Research

Kumar et al. [11] propose an AAC-based big data analytics framework for children with neural disorders, incorporating eye-tracking via a Tobii X2-60 device to enable gaze-based communication support. Their CNN model achieves 98.63% accuracy in predicting communicative intent from gaze data, and their pipeline, from eye gaze data collection through

preprocessing, feature extraction, and text-to-speech output, represents a technically sophisticated approach to augmentative communication. The work is significant in demonstrating that AAC systems can be built on robust machine learning foundations, though it presupposes access to expensive eye-tracking hardware that is unavailable in most Sri Lankan special education settings.

Hiryanto et al. [12] describe VICARA 3.0, a redesigned Bahasa Indonesia AAC mobile application for non-verbal autistic children. Their design thinking approach produces a system evaluated against nine ASD-specific UX guidelines, engaging, predictable, structured, interactive, generalizable, customizable, sense aware, attention retaining, and frustration free and confirmed compliant by a focus group including an ICT representative, an ASEAN Autism Studies representative, and a speech-language pathologist. Notable design features include Fitzgerald colour-coded picture cards, four emotion-representing mascots, five selectable visual themes, a PIN-protected parental mode, and Google Text-to-Speech integration. VICARA 3.0 is directly relevant to the proposed system's interface design principles, particularly the commitment to structured, predictable navigation and avatar-based guidance, though it focuses exclusively on AAC communication rather than curriculum-aligned language learning.

1.2.4. User-Centred Design and Scheduling Applications

Akbar et al. [13] apply User-Centred Design methodology to develop a scheduling activity application for autistic children aged seven to twelve. Their dual-interface Android application, a parent-facing content creation interface and a child-facing activity-following interface is evaluated using the QUIM usability model across ten factors in two iterative design cycles. Their final evaluation scores of 90.5% for parents and 89.9% for children confirm the effectiveness of iterative, participant-informed design. Particularly relevant is their finding that splitting the parent input process across multiple views significantly improved the effectiveness dimension score from 84% to 90%, suggesting that task decomposition is as important in the interface design for adult users as it is in the learning task design for child users. The proposed system adopts a comparable dual-perspective architecture, maintaining a strictly separated Teacher Workspace and Student Workspace, with the teacher's interface designed to support professional decision-making rather than to simulate the child's experience.

1.2.5. Existing Products and Their Features

Otsimo is a commercially available mobile learning application developed for autistic children, and its feature set was subjected to a systematic evaluation as part of the present research. The application demonstrates competence across several dimensions of instructional design. It provides a structured activity selection interface, a parent-gated settings panel secured through numeric adult authentication, and child profiles accessible exclusively within the parent-controlled environment. A premium-tiered multi-user architecture is available, and an initial onboarding

questionnaire, encompassing cognitive and social developmental indicators, generates a personalised curriculum for each child. With respect to reinforcement design, the application employs a multimodal reward system comprising confetti animations, per-step star allocation, activity-level achievement badges, and positively valenced verbal feedback including phrases such as "Good job" and "Well done." The instructional content spans a range of receptive and expressive language task types, including vocabulary matching, cloze exercises, listening comprehension, bidirectional picture-word correspondence tasks, and drag-and-drop interactions.

Notwithstanding these strengths, the feature audit surfaces a constellation of limitations with direct relevance to the proposed system. Navigation architecture lacks category-based organisation, confining the learner to sequential progression through individual activities without higher-order content grouping. No provisions exist for typographic accessibility: adjustable text size, font style selection, and high contrast display modes are uniformly absent. Auditory scaffolding options, including slow speech mode and variable speech rate are not available, nor is any form of symbol-based augmentative and alternative communication (AAC) support. The content architecture is entirely closed: custom vocabulary addition, learner-specific image upload, and editable word sets are not supported, precluding any form of curriculum modification by teaching staff or caregivers. Data governance presents a further concern, as the application's age verification mechanism does not distinguish between verified guardians and other adults, rendering its privacy protections only partially effective. The application additionally requires continuous internet connectivity, with no offline functionality available. Most significantly for the context of the present study, Otismo is not architecturally conceived as a teacher-mediated instrument; it operates as a self-contained learning environment with no alignment to any national educational framework and no mechanism by which a special education teacher can exercise professional oversight over the learning process.

Speech Blub is another mobile application oriented towards speech and language development in children. Its evaluated strengths are most apparent in the domain of interface personalization and auditory instruction. The application generates a bespoke child avatar from a submitted photograph, which subsequently appears throughout the application as a participatory model during learning activities, a design approach with theoretical grounding in social learning and self-referential processing. Pronunciation modelling is supported through slowed speech repetition, and audio instructions are delivered via the avatar at each stage of the learning sequence. Additional interface features include visual cue highlighting and generously proportioned interactive targets. A parent gate is implemented using date-of-birth verification with session timeout, and an initial assessment questionnaire spanning hearing, pragmatics, play, comprehension, and expressive language domains produces a developmental profile benchmarked against age-typical norms. A progress celebration screen, incorporating confetti and avatar imagery, provides milestone-level positive reinforcement.

However, the application presents significant functional gaps across several dimensions critical to structured educational use. Text-to-speech output, symbol-based communication support, word and sentence construction tools, and visual vocabulary card systems are absent. Category-based

content navigation is only partially realized, organized by broad competency domain rather than discrete, sequentially structured content categories. A notable safety concern in the navigation design is the absence of authenticated return logic on the back button, permitting unrestricted child-initiated navigation without adult re-verification. Typographic and visual customization, including adjustable text size, font style, high contrast mode, and background colour selection is entirely unavailable. Performance data collection is limited: neither accuracy rates nor time-on-task metrics are recorded at the activity level. Educator access to learner progress is further constrained by the absence of a dedicated reporting interface; data retrieval requires navigation through the application's settings panel. Task type coverage is incomplete, with sentence construction, picture-to-word, and word-to-picture matching activities absent from the content repertoire. In common with Otsimo, Speech Blub is not designed to support professionally mediated instruction within a structured curricular framework and provides no mechanism for educator-driven content customization or adaptive instructional decision-making.

1.2.6. Limitations of Existing Systems

Across both the scholarly literature and the commercial product landscape, several interconnected limitations emerge that the proposed system is designed to address.

Table 1 - Limitations of Existing Systems

Limitation	Otsimo	SpeechBlub
Offline accessibility	Not available; requires an internet connection for all functionality	Not available; requires an internet connection for all functionality
Category-based navigation	Not present; navigation is limited to moving between activities only, with no ability to browse or select by content category	Partially present; categories are loosely defined around competencies rather than structured content groupings
Safe return navigation	Present with verification	Not present; a child can navigate back to the home page without any adult verification
Content editability	No support for custom word addition, custom image uploads, or editable vocabulary sets; the curriculum cannot be modified	No support for custom word addition, custom image uploads, editable vocabulary sets, or custom lesson creation
Paywall restrictions on core features	Only the first activity is free; all remaining activities, multiple user accounts, detailed statistics, and exportable reports require a premium subscription	Limited activity access for free users; exportable reports are restricted to premium subscribers

Data privacy controls	Partially present; any adult can access the application after the initial login since only age verification is required, with no robust authentication mechanism	Present; access is gated through date-of-birth verification
Symbol-based communication support	Not present; no communication board or AAC-style support is provided	Not present; no communication board or AAC-style support is provided
Word and sentence building	Partially present; limited to specific activity types	Not present; no sentence construction or word building capability exists
Encouragement and celebration feedback	No encouragement messages or progress celebration screens are provided; feedback is limited to visual rewards and badges	Present but relies on confetti animations and avatar photo captures rather than contextual encouragement tied to learning progress
Highlighted word during speech output	Present but only within specific activities	Not present; only the image is displayed on screen without the corresponding written word during pronunciation
Profile customization by the child	Partially present; customization is restricted to parental control with no child-facing options beyond what the parent configures	Partially present; customization is restricted to parental control with no child-facing options beyond what the parent configures

The absence of any graph-based or relational adaptive mechanism is a cross-cutting limitation of all reviewed systems. Neither Otismo nor Speech Blub models the conceptual relationships between learning items, the structural similarity between confusion patterns, or the cross-student evidence available when multiple learners exhibit the same difficulty cluster. The proposed GNN directly addresses this absence.

1.3. Research Gap

A feature-level comparison of two widely used autism-focused learning applications, Otismo and Speech Blub, against the requirements of a structured concept learning system for autistic children reveals six critical gaps that neither application addresses.

First, neither application provides a structured, tiered concept learning architecture. Otismo offers vocabulary lessons, matching exercises, and drag-and-drop activities, while Speech Blub focuses on pronunciation and phonics. However, both deliver content as isolated activities without a

defined progression from visual familiarization through word recognition to attribute understanding, and without curriculum-aligned sequencing based on cognitive complexity. This absence of a tiered mastery framework means children are not required to demonstrate competency at one level before advancing to the next, undermining the reliability of skill acquisition.

Second, error handling in both applications remains superficial and does not constitute genuine adaptive scaffolding. Otsimo renders incorrect options unavailable upon selection, reducing the task to elimination rather than recognition. Speech Blub repeats the same pronunciation without altering the modality or stimulus. Neither application implements a graduated relearn pathway that varies the visual stimulus type across attempts, tracks which specific distractors a child repeatedly selects to identify confusion patterns or detects persistent reliance on audio support to retain scaffolding in future sessions.

Third, sensory customization options are inadequate. Otsimo does not support adjustable text size, font style, high contrast mode, or speech speed control. Speech Blub similarly lacks font, text size, and background colour customization. Neither application conducts a child-driven initial preference assessment for avatar, font, and theme selection that is applied consistently across all subsequent sessions, despite the established sensitivity of autistic children to visual and typographic presentation.

Fourth, both applications operate exclusively within a parent-child model and provide no role for special education teachers. Otsimo presents raw performance statistics behind a premium subscription without contextual interpretation. Speech Blub lacks accuracy tracking, time-per-activity metrics, and a dedicated adult-facing dashboard. Neither application supports teacher-student allocation, concept-level interpretive feedback, or mechanisms through which a teacher can exercise professional judgment over the learning pathway.

Fifth, neither application extends concept learning beyond surface-level recognition to include defining attributes such as colour, shape, or category membership. There is no pre-classification of concepts based on attribute stability, and consequently no logic to determine when attribute-based activities should be presented in full, adapted to accept multiple valid responses, or omitted where no single correct attribute exists. Sixth, neither application conducts structured post-concept evaluations that assess retention and generalization across multiple previously learned concepts, nor do they record the active difficulty level at the time of assessment to enable contextually meaningful interpretation of performance scores.

The proposed system addresses all six identified gaps within a single integrated platform through curriculum-aligned tiered concept learning with attribute classification, a graduated multi-modal relearn pathway with distractor confusion tracking, child-driven sensory UI personalization, a teacher-mediated decision layer supported by LLM-interpreted feedback, and difficulty-contextualized post-concept evaluations. This combination of capabilities does not exist in either of the reviewed applications and represents the principal contribution of the present work.

2. OBJECTIVES

2.1. Main Objective

The main objective of the proposed solution component is to design and develop a GNN-enhanced adaptive concept learning module that delivers personalized, curriculum-aligned English vocabulary instruction to children with autism in Sri Lanka (up to Grade 3 level), using a GNN driving relational adaptive decisions, supported by a rule-based pedagogical engine enforcing ABA-grounded mastery criteria, and an LLM-powered feedback system that equips special education teachers with actionable, graph-enriched, data-driven insights to guide each child's learning progression.

2.2. Specific Objectives

The specific objectives of this component are designed to address both the instructional delivery and the teacher support dimensions of the system. Together, they define the functional scope of the concept learning module and the LLM-powered feedback mechanism, ensuring that each child receives structured, personalized vocabulary instruction while teachers retain meaningful visibility and control over each learner's progress.

- To implement a three-tiered concept learning structure (Visual Familiarization, Single Word Recognition, and Attribute Introduction) that progressively builds English vocabulary understanding across ten curriculum-aligned concept categories sequenced by cognitive complexity.
- To design and populate a GKB consisting of a static curriculum layer and a dynamic interaction layer updated in real time from session interaction logs.
- To develop a GNN inference engine that performs concept confusion detection, learning path recommendation, per-concept difficulty prediction, and student learning profile clustering by traversing the GKB, and to integrate its outputs into a Content Adaptation Engine that dynamically assembles session activities from a modular asset library.
- To develop an adaptive difficulty and relearn pathway system that dynamically adjusts activity complexity based on each child's performance, applying ABA-based 90% mastery criteria to ensure reliable skill retention before progression, operating independently of GNN recommendations to preserve pedagogical integrity.
- To design concept-level attribute classification and teaching activities that accurately represent fixed, dominant, and variable concept attributes through drag-and-drop and colouring activities, with Tier 3 readiness verified by the GNN against the student's Phase 1 (Colors, Shapes) mastery edge weights before attribute activities are presented.
- To build an LLM-powered inline feedback system that generates concise, plain-language interpretive summaries for each concept directly within the teacher's view, with prompts

enriched by GNN-extracted analytics including confusion cluster summaries, attribute difficulty profiles, prerequisite gap detections, and student cluster assignments.

- To develop a per-child statistics dashboard that presents longitudinal progress data across all concept categories, with LLM-generated category-level insights enriched by GNN graph analytics, surfacing patterns not detectable from raw performance numbers alone.
- To implement an on-demand formal report generation feature allowing teachers to generate structured, period-specific LLM reports covering performance trends, GNN-identified confusion clusters, concept-level breakdowns, and actionable instructional recommendations.

3. METHODOLOGY

3.1 Research Approach

This study adopts a Mixed-Methods Research Methodology, combining quantitative system evaluation with qualitative user-centred assessment. This approach is necessary because the research problem requires evidence along two distinct dimensions that cannot be addressed by a single methodology.

The first dimension is whether the system performs as intended. The Concept Learning Module implements structured tier-based progression, adaptive difficulty adjustment, graduated relearn pathways, distractor confusion tracking, and LLM-generated teacher feedback. Each of these produces measurable outputs that can be verified through functional testing and system log analysis.

The second dimension is whether the system is genuinely usable and appropriate for its target users. The end users are young autistic children in Sri Lanka, and the proxy users are special education teachers who oversee every learning session. Whether the interface accommodates sensory sensitivities, whether the personalization features make a practical difference, whether teachers find the generated feedback useful and actionable, and whether the system supports rather than undermines professional judgment are questions that system logs alone cannot answer. These require structured usability evaluation and qualitative feedback from teachers and domain experts.

A purely quantitative approach would confirm that the system functions correctly but would leave its real-world suitability for a vulnerable population unvalidated. A purely qualitative approach would capture user experience but would not verify that the underlying learning logic and adaptive mechanisms operate to specification. The Mixed-Methods approach produces two complementary bodies of evidence: quantitative data confirming that the system works correctly, and qualitative data confirming that it works meaningfully for the people it is designed to serve.

3.2 Technical Methods and Algorithms

Every concept is taught through Tier 1 (Visual Familiarization), Tier 2 (Single Word Recognition), and Tier 3 (Attribute Introduction) in that fixed sequence. A child must demonstrate mastery at one tier before the system permits advancement to the next.

The GKB is the central data structure of the system. It has two layers. The static curriculum layer is built during the design phase and edges linking perceptually similar concepts, and attribute type flags (fixed, dominant, variable) on concept nodes. The dynamic interaction layer is populated and updated in real time from session logs. Every interaction event, correct answer, wrong distractor selected, audio replay, hint used, relearn triggered and updates the relevant edge weights in the

dynamic layer immediately. The GNN traverses the GKB to perform four inference functions as a background service before each session and after each interaction block. Concept Confusion Detection traverses the student's edges and propagates risk scores, enabling the system to preemptively flag confusion risk for concepts the student has not yet encountered. Learning Path Recommendation uses graph similarity and prerequisite of edge satisfaction to order the concept queue within a session, checking that Phase 1 mastery edge weights are sufficient before Tier 3 attribute activities are unlocked. Difficulty Prediction compares the student's learned embedding vector against each concept node's embedding to predict difficulty before the concept is introduced, pre-configuring the initial distractor count and setting early teacher-flag thresholds. Student Learning Profile Clustering groups students by the structural similarity of their interaction graphs, enabling the system to apply conservative difficulty settings for a student encountering a concept that multiple cluster peers have historically found difficult, before the new student has attempted it.

At the start of each concept session, the Content Adaptation Engine queries the GNN for predicted difficulty (to set initial distractor count), active confusion with edge weights (to exclude high-confusion distractors from early blocks), preferred image format (carried forward as an edge attribute), and audio scaffolding flag. It then dynamically assembles the session's activity from the content asset library, concept images (photograph, cartoon, animation), audio clips, colour/shape assets, distractor pools, and activity templates, rather than serving any pre-built quiz. This means every student's experience of the same concept is assembled fresh, using the most current state of their interaction graph. The rule engine operates independently of the GNN and cannot be overridden by graph outputs. It enforces the 90% mastery criterion across two consecutive blocks, tier progression gating, the three-level graduated relearn pathway and teacher handover flag issuance when persistent failure thresholds are reached. Every concept is pre-classified as fixed, dominant, or variable attribute type, stored as a node attribute on the concept node in the static GKB layer. Before Tier 3 is unlocked for a student, the GNN checks the student's mastery edge weights on Phase 1 concepts (Colors, Shapes). If foundational readiness is weak, the Content Adaptation Engine reduces the initial number of color/shape options, inserts a brief attribute name reminder, and logs the gap to the teacher dashboard.

.

3.3 Tools and Platforms

The proposed system will be implemented using a multi-layer architecture consisting of a mobile frontend, distributed backend services, machine learning components, and a structured data storage layer. The technology stack is selected to support real time interaction, adaptive intelligence, scalable data processing, and reliable system performance within an educational environment.

The frontend of the system will be developed using React Native. Both the child facing learning environment and the teacher facing management interface will be implemented from a single React Native codebase. This approach will allow the system to deliver native level performance for the

touch-based interactions that the learning activities rely on, including tap to hear playback, drag and drop concept matching, and colour-based selection tasks. The component-based architecture of React Native will also enable the system to dynamically load each student's stored sensory preferences and GNN assigned learning profile at the beginning of a session and apply these settings consistently across all screens of the learning interface.

The backend of the system will consist of two complementary services implemented using Node.js and FastAPI. The Node.js service will handle real time session management, including recording student responses, managing activity progression, updating the dynamic interaction layer of the Graph Knowledge Base stored in PostgreSQL, and serving data required by the teacher dashboard. In parallel, FastAPI will manage machine learning related operations including Graph Neural Network inference requests and Large Language Model integration. Since both graph inference operations and LLM API calls may require several seconds to complete, FastAPI's asynchronous processing capabilities will ensure that these operations do not delay session critical interactions handled by the Node.js service.

The system's adaptive intelligence will be powered by GNN. It will be trained offline using accumulated student interaction data during the later stages of system implementation. Once trained, the model will be deployed as an inference service exposed through a FastAPI endpoint. The Content Adaptation Engine will query this endpoint at the beginning of each session and after defined interaction blocks in order to obtain predictions related to concept difficulty, confusion clusters, and recommended learning paths.

The system's data layer will be implemented using PostgreSQL. The dynamic interaction layer of the GKB will be stored as a relational schema where nodes and edges are represented as tables containing typed relationships and associated weight values. This structure will allow the system to update graph relationships based on student interactions while maintaining transactional reliability. PostgreSQL's support for complex queries and concurrency control will ensure that edge weight updates generated from multiple simultaneous student sessions are processed accurately. The static curriculum layer representing the structured concept hierarchy derived from the national curriculum will be loaded from a predefined configuration file during application initialization.

To support teacher decision making, the system will integrate a LLM. It will be selected due to its ability to process structured data inputs, including graph analytics generated by the GNN, and produce contextually appropriate natural language explanations. Prompts sent to the model will include structured payloads containing GNN derived insights together with session performance statistics. Each prompt will contain explicit instructions to generate concise, plain language feedback directed toward teachers within a special education context. To ensure data privacy and ethical compliance, no personal identifiable student information will be included in any request sent to the language model.

3.4 Data Collection and Ethical Considerations

Data collection in this study occurs across two phases, corresponding to the quantitative and qualitative strands of the mixed-methods approach.

In the quantitative phase, data is collected through the system itself during controlled testing sessions. The system automatically logs all relevant interaction data generated during learning activities. This includes mastery scores per block, the number of attempts required to reach mastery at each tier, whether and at what level the relearn pathway was triggered, which distractor options the child selected incorrectly and how frequently, how often audio replay was used, which image format proved most effective, the difficulty level active at each assessment point, and post-concept evaluation scores with their associated difficulty context. No manual data recording by the researcher is required during this phase. Additional quantitative data is generated through structured functional testing, where specific system behaviours such as tier progression decisions, relearn triggers, difficulty adjustments, and LLM output quality, are tested against their expected outcomes.

In the qualitative phase, data is collected through structured usability evaluation sessions with special education teachers. Teachers interact with both the Teacher Workspace and the Student Workspace and provide feedback through task-based assessment, post-session questionnaires, and guided discussion. The focus of this feedback is on whether the interface is clear and easy to navigate, whether the LLM-generated summaries and reports are understandable and useful in practice, whether the teacher handover mechanism supports their professional workflow, and whether the overall design is appropriate for use in a special education classroom. Where ethically approved and practically feasible, guided observation sessions with autistic children may also be conducted under the direct supervision of their assigned teacher. In such sessions, data collection is limited to observational notes on engagement, task completion, and any visible signs of distress or disengagement, as identified by the supervising teacher.

Ethical considerations are given central importance throughout this study given the involvement of a vulnerable population. All research activities involving human participants will proceed only after obtaining ethical clearance from the relevant institutional review board. Informed written consent will be obtained from all participating teachers. For any child participants, written parental or guardian consent will be obtained, and assent from the child will be sought where their developmental level permits meaningful assent. Any session involving a child will be terminated immediately if the child displays signs of distress, frustration beyond what is typical for the task, or unwillingness to continue, regardless of whether the session is complete. No personally identifiable information about any child will be included in data sent to external services, including the Gemini API. All prompts sent to the language model will use anonymized session data that excludes names, identification numbers, and any other identifying details. All collected data will be stored securely, accessible only to the research team, and will not be disclosed to any third party.

Participation in all phases of the study is entirely voluntary, and any participant or guardian may withdraw at any stage without consequence.

3.5 Validation Strategy and Performance Metrics

The validation strategy is organized into two complementary strands that together provide a complete assessment of the system.

The quantitative validation strand evaluates whether the system's internal mechanisms operate as specified. This includes verifying that tier progression permits advancement only when the mastery threshold is met, that the relearn pathway triggers at the correct point and changes the learning experience appropriately at each level, that difficulty adjustments respond correctly to the child's performance pattern, that confusion detection identifies repeated distractor selections and applies the intended interface adjustments, that the attribute classification logic presents the correct type of Tier 3 activity for each concept category, that post-concept evaluations trigger at the correct interval and record difficulty context alongside scores, and that LLM-generated feedback is factually consistent with the session data, relevant to the specific concept, appropriately toned for a teacher audience, and compliant with the prompt instructions. Each of these is tested through controlled scenarios where expected outcomes are defined in advance and compared against actual system behaviour.

The qualitative validation strand evaluates whether the system is usable, understandable, and practical in a real educational setting. Special education teachers are asked to perform representative tasks within the system, selecting a student, navigating to a specific concept, reading and interpreting an LLM-generated summary, locating a particular metric on the dashboard, and generating a date-range report. Their performance on these tasks is assessed through task completion rate, time taken, and errors encountered, providing quantitative usability indicators. A standardized usability questionnaire, the System Usability Scale, provides a composite usability score. Open-ended feedback gathered through post-session interviews or questionnaire items captures the teachers' perspectives on whether the feedback is actionable, whether the interface suits their classroom context, and whether any aspects of the system would create barriers to adoption. Where child-facing evaluation is conducted, engagement is assessed observationally by the supervising teacher, focusing on whether the child interacts voluntarily, completes activity blocks without external prompting, and shows no observable signs of distress.

The specific metrics used across both strands are as follows. Mastery assessment accuracy refers to the proportion of test cases where the system correctly applies the progression or relearns decision based on the defined threshold. Relearn pathway accuracy refers to the proportion of cases where the correct relearn level is triggered at the correct point. Difficulty adjustment responsiveness refers to whether distractor count changes occur within the expected number of consecutive correct or incorrect attempts. Confusion detection accuracy refers to the proportion of

simulated confusion patterns that the system correctly identifies and acts upon. LLM feedback quality refers to the proportion of generated summaries and reports judged by evaluating teachers as accurate, relevant, and appropriately written. Task completion rate refers to the proportion of representative tasks that teachers complete successfully without assistance during usability testing. System Usability Scale score provides a standardized usability rating from the teacher participant group. Evaluation integrity refers to the proportion of post-concept evaluations that are correctly triggered, correctly composed from multi-tier activity types, and correctly annotated with difficulty level context.

4. HIGH-LEVEL SYSTEM ARCHITECTURE

The proposed system architecture is designed to support adaptive English concept learning for autistic children through a structured and intelligent digital learning environment. The architecture integrates tier based instructional design, adaptive learning mechanisms, graph driven analytics, and AI assisted teacher insights to create a personalized learning experience aligned with special education teaching practices. It combines interactive learning modules used by students with analytical and decision support components that assist teachers in monitoring progress and adjusting instruction. The system is structured to capture detailed interaction data during learning activities, analyse behavioural patterns using graph-based models, and generate meaningful insights that guide adaptive learning adjustments and instructional decision making. The overall architecture therefore ensures that concept learning is delivered in a structured, responsive, and data informed manner while remaining manageable and interpretable for special education teachers.




Figure 1 - High Level Architecture of the Adaptive Concept Learning Support System for Autistic Children

5. USER REQUIREMENTS

3.1. Users and Stakeholders

Stakeholders

Principal functions as the system administrator. They are responsible for setting up the system prior to use by creating and managing all user accounts. The principal does not interact with the learning content directly but oversees the operational structure of the system by registering teachers and students and assigning students to their respective teachers.

Users

Special Education Teachers are the primary operational users of the system. They are classified as proxy users because they navigate and manage the system on behalf of the child during learning sessions. While the child physically interacts with the on-screen activities, the teacher is responsible for initiating sessions, selecting modules, monitoring the child's responses, and reviewing performance feedback. Teachers also access their own workspace to track student progress and review LLM-generated insights following sessions and evaluations.

Autistic Children are the end users and the direct beneficiaries of the system. Although they do not independently navigate the system, they are the ones who physically interact with the learning activities tapping images, dragging and dropping elements, selecting colours, choosing avatars, and completing assessments. All learning content, interface design, and adaptive mechanisms are designed around their individual needs, preferences, and performance.

3.2. Functional Requirements

1: The system shall enable the principal to pre-register teacher and student accounts, assign each teacher up to three students, and manage all user accounts without requiring self-registration from any user.

02: The system shall authenticate teachers using their assigned Teacher ID as the username and a password created during their first login and shall present two workspace options upon successful authentication: Teacher Workspace and Student Workspace.

03: The system shall display the teacher's profile details, allocated student list, and relevant performance statistics within the Teacher Workspace dashboard.

04: The system shall allow the teacher to select an assigned student from the Student Workspace and navigate to that student's profile page displaying available learning modules.

05: The system shall conduct a one-time initial preference assessment during a student's first session, comprising avatar selection from four predefined categories (Robot, Animal, Fantasy Character, Human) and font style and size selection from autism-friendly options.

06: The system shall persistently apply the selected avatar, its associated visual theme, and the chosen font preferences across all subsequent sessions for that student without requiring reconfiguration.

07: The system shall organize concepts into ten categories sequenced across four phases of increasing cognitive complexity, derived from the Grade 1 and Grade 2 English curriculum published by the National Institute of Education of Sri Lanka.

08: The system shall pre-classify every concept according to its attribute stability as fixed, dominant, or variable, and shall use this classification to determine which Tier 3 activities are presented, adapted, or omitted for that concept.

09: The system shall deliver Tier 1 (Visual Familiarization) by displaying a concept image that pronounces the concept word upon tap for up to five repetitions, defaulting to a real-life photograph with illustrated cartoon and animation formats held in reserve as fallback stimuli.

10: The system shall assess Tier 1 mastery through an audio-visual identification activity structured in blocks of five attempts each, applying a 90% accuracy criterion across two consecutive blocks, with image positions randomized between attempts.

11: The system shall dynamically adjust the distractor count within Tier 1 activities by increasing distractors on consecutive correct answers and reducing them on consecutive incorrect answers.

12: The system shall implement graduated, non-punitive error handling within each block, progressing from a subtle visual indication on the first incorrect attempt, to distractor removal on the second, to correct answer highlighting with a hint on the third.

13: The system shall trigger a three-level relearn pathway upon failure to meet the mastery threshold, where Level 1 switches from photographs to illustrated cartoons, Level 2 introduces a short contextual animation before retrying, and Level 3 returns the child to the full familiarization screen with animation-on-tap interaction.

14: The system shall record the image format that proves most effective for each child during Tier 1 and shall use that format as the default visual stimulus in Tier 2 for the same concept.

15: The system shall deliver Tier 2 (Single Word Recognition) through two sequential activities - drag-and-drop of the image onto the corresponding word, followed by word-to-image matching - applying the same 90% mastery criterion and block structure established in Tier 1.

16: The system shall track word-visual confusion patterns during Tier 2 and, upon detecting visually similar word confusion, shall switch to more distinct fonts, increase spacing between word options, and highlight the initial letter of each word.

17: The system shall deliver Tier 3 (Attribute Introduction) by identifying the concept's category, displaying a short animation demonstrating the concept's real-world function, and introducing colour and shape attributes through interactive activities including drag-and-drop attribute matching and image colouring.

18: The system shall adapt Tier 3 activities according to the concept's attribute classification, presenting full colour and shape activities for fixed-attribute concepts, accepting multiple correct answers with positive reinforcement for dominant-attribute concepts, and skipping or replacing colour activities with category-based alternatives for variable-attribute concepts.

19: The system shall flag a concept for teacher intervention on the teacher dashboard if a child reaches persistent failure thresholds at any tier after exhausting all relearn levels and shall not loop the child indefinitely through repeated attempts.

20: The system shall provide the teacher with the ability to decide whether to return the child to a previous tier, reattempt the activity in a future session, or temporarily move forward, maintaining teacher authority over algorithmic recommendations.

21: The system shall conduct a structured post-concept evaluation after every five concepts are completed across any category, drawing on activity types from all three tiers, and shall record the active difficulty level alongside the evaluation score.

22: The system shall provide inline LLM-generated interpretive summaries of two to four sentences beneath each concept view in the teacher interface, generated on demand using current session data and written in plain, non-technical language.

23: The system shall provide a per-child statistics dashboard within the Teacher Workspace displaying category-level progress, concept-level tier completion status, cumulative relearn cycle

counts, evaluation scores with difficulty context, and LLM-generated category-level insight summaries.

24: The system shall enable the teacher to generate a formal LLM-interpreted report for a selected child by specifying a date range, producing a structured document comprising a period summary, concept-level breakdown, instructional recommendations, and a progress trend comparison against the child's historical baseline.

25: The system shall log audio replay frequency, distractor selection patterns, relearn level reached, relearn cycle count, and mastery attempt count per tier for every concept worked on by every child.

26: The system shall automatically replay the audio instruction once if a child is inactive for a predefined period during any activity attempt, and shall provide a replay button accessible at all times.

27: The system shall maintain a Graph Knowledge Base consisting of a static curriculum layer and a dynamic interaction layer, updating edge weights in the dynamic layer in real time following each interaction block, and shall expose GNN inference endpoints queried by the Content Adaptation Engine at start and after each block.

28: The GNN shall perform confusion detection, learning path recommendation, difficulty prediction, and student cluster assignment before each session, and the Content Adaptation Engine shall use these outputs to configure distractor pools, image format selection, initial difficulty level, and audio scaffolding decisions for that session.

29: The system shall verify Phase 1 (Colors, Shapes) mastery edge weight thresholds in the student's GKB node before unlocking Tier 3 attribute activities for any concept, and shall insert a brief attribute name reminder and reduce initial option count if foundational readiness is below the configured threshold.

30: The teacher dashboard shall display GNN-generated confusion cluster alerts identifying which concepts are connected by high-weight edges and whether they share a common attribute pattern and shall surface learning path visualization showing the student's current position within the concept graph.

3.3. Non-functional Requirements

3.3.1. Usability

The system shall employ calming, sensory-friendly colour palettes across all interface elements, following established colour guidelines recommended for children with autism, ensuring that no visual element causes sensory overload or distress.

The system shall maintain a minimal, uncluttered interface with no complex navigation structures, ensuring that all interactive elements are clearly visible, adequately sized for touch interaction, and free from irrelevant animated or decorative distractions.

All avatar characters, background themes, and visual accents shall be designed during the system design phase to be age-appropriate, non-threatening, and consistent with the sensory preferences of the target user population.

3.3.2. Accessibility

The system shall support multiple image formats for the same concept - real-life photographs, illustrated cartoons, and animations, to accommodate individual differences in visual processing among autistic children.

All audio output shall be clear, appropriately paced, and delivered in a consistent voice to support auditory processing, and shall be available as a voluntary replay at all times during learning activities.

3.3.3. Performance

The system shall generate inline LLM-interpreted summaries within a response time that does not disrupt the teacher's workflow, ensuring that concept-level and category-level insights are available without perceptible delay when the teacher navigates to the relevant view.

The system shall render all learning activities, including image loading, audio playback, and animation display, without latency that could interrupt the child's engagement or cause confusion during timed interaction sequences.

3.3.4. Scalability

The system shall support concurrent use by multiple teacher-student pairs without degradation in response time, data recording accuracy, or LLM feedback generation quality.

3.3.5. Data Integrity

The system shall reliably persist all session data, including mastery scores, relearn cycle records, distractor confusion logs, audio replay frequency, and UI preference selections, ensuring that no data is lost between sessions.

The system shall record the difficulty level active at the time of every evaluation score and every mastery assessment, ensuring that performance data is always interpretable in its original context.

3.3.6. Security

The system shall restrict student account creation and teacher-student allocation exclusively to the principal and shall prevent students from accessing teacher or administrative functions.

All LLM API communications shall transmit session data in a structured format via the Gemini API without exposing personally identifiable student information in prompts, and all generated feedback shall be stored within the system's secure data environment.

3.3.7. Maintainability

The concept attribute matrix, category sequencing, and curriculum content shall be configurable through a structured content configuration layer, enabling updates to concept data, attribute classifications, and category structures without requiring changes to the application codebase.

All LLM prompts shall be structured with consistent formatting and explicit instructions for plain-language, teacher-addressed output appropriate to a special education context, and shall be maintainable as a separate prompt configuration layer.

3.3.8. Availability

The system shall ensure that a child's learning session is not interrupted by background processes such as LLM feedback generation, report compilation, or data synchronization, isolating teacher-facing analytics from child-facing learning activity execution.

6. COMMERCIALIZATION PLAN

6.1. Target Market and Customer Persona

The primary market for this system will consist of government and private special education schools across Sri Lanka. Nationally, there are approximately 27 dedicated special schools and 704 special education units. The principal of each institution will serve as the primary purchasing decision-maker and will map directly to the system administrator role within the platform architecture.

A secondary procurement pathway is envisioned through the Ministry of Education Sri Lanka. Should the system demonstrate measurable outcomes at pilot scale, Ministry-level adoption could facilitate a national rollout through centralized procurement channels. The direct end users of the system will be special education teachers and autistic children aged five to twelve, who will be non-paying beneficiaries of the platform.

The target customer persona is a principal of a government-funded special education unit in a semi-urban Sri Lankan town, typically overseeing three to eight teachers and fifteen to thirty students across varying autism severity levels. This individual operates under a constrained IT budget with limited access to specialist English language teaching resources and makes procurement decisions based on demonstrable classroom impact rather than technical specifications.

6.2. Value Proposition

Existing tools for autistic children fall broadly into two categories: generic global applications that bear no alignment to the Sri Lankan curriculum, and expensive imported resources that require trained therapists to administer. This system will address both shortcomings as the only structured, adaptive English language learning platform built specifically around the Sri Lankan national syllabus and designed to be operated by a classroom teacher without any specialist technology background.

By reducing the per-child preparation burden on teachers while delivering personalised, curriculum-aligned instruction, the system will replicate outcomes that would otherwise require dedicated one-on-one therapeutic support. This will make high-quality special education more accessible and scalable across the country.

As an academic prototype, the value proposition will also extend to demonstrating a locally developed, research-backed solution that addresses a documented gap in national special education provision, with a viable path to real-world deployment following graduation.

6.3. Revenue Model

The system will operate on a Software as a Service subscription model, with an annual license fee charged per institution rather than per student. This structure will make budgeting predictable for schools operating under fixed government allocations. In addition to the paid tiers, a freemium pilot license could be made available, to allow schools to evaluate the platform before committing to a subscription, which aligns with how procurement decisions are typically made in Sri Lankan public education, where trial evidence is often required prior to budget approval.

6.4. Cost Estimation

The cost estimation covers three phases: development, deployment, and maintenance. During the development phase, costs include cloud infrastructure, LLM API integration, UI/UX design, and mobile development utilities, totalling approximately Rs. 34,500 for a six-month prototype period. For deployment, costs will scale with cloud hosting, managed database services, and content delivery requirements based on the number of institutions and students onboarded. Ongoing maintenance costs will cover server hosting, database storage, LLM API usage, and periodic curriculum-aligned content updates.

6.5. Pricing Strategy

The system will adopt an institutional subscription model, targeting special education schools and government-affiliated institutions in Sri Lanka. Three subscription tiers are proposed based on student capacity and institutional size. Tier 1 is designed for small special education units accommodating up to 30 students. Tier 2 targets mid-scale schools and larger special education units supporting between 30 and 120 students. Tier 3 is intended for large institutions with expanded administrative needs, covering 120 or more students or multi-campus institutions.

Table 2 - Proposed Institutional Subscription Tiers for the System

Tier	Student Capacity	Target Segment
Tier 1	Up to 30 students	Small special education units
Tier 2	30 – 120 students	Mid-scale schools and larger special education units
Tier 3	120+ students or multi-campus institutions	Large institutions with expanded administrative needs

6.6. Competitive Advantage

The proposed module offers several key advantages compared to existing autism learning applications. First, the platform is specifically aligned with the Sri Lankan national curriculum, whereas most existing tools are designed for global audiences without consideration for local educational requirements. Second, the system incorporates a structured three-tier learning architecture with adaptive difficulty adjustment, ensuring that each child progresses only after demonstrating mastery of previous learning stages. Third, the platform includes a teacher-centered design, providing educators with detailed performance analytics and LLM-generated interpretive feedback. Most commercial autism learning applications operate primarily as parent-child tools and do not integrate teachers into the learning process. Fourth, the system emphasizes sensory-friendly interface design, allowing personalization of avatars, visual themes, and font settings to accommodate the sensory sensitivities of autistic learners.

Finally, the system integrates learning analytics and LLM-powered insights, enabling teachers to make data-driven instructional decisions without manually interpreting raw performance data. These combined capabilities create a platform that is uniquely positioned to support structured language learning for autistic children in the Sri Lankan educational context.

6.7. Intellectual Property Considerations

The proposed system includes several intellectual property elements that can be protected to support future commercialization. The software architecture and adaptive learning algorithms can be protected through software copyright. This includes the rule-based adaptive learning engine, tier-based concept delivery mechanism, and the learning analytics framework.

The user interface design, including sensory-friendly themes, avatar-based feedback mechanisms, and activity interaction models, may also be protected under design rights. Original educational content developed for the platform, including concept images, animations, learning activities, and instructional materials, will be protected under copyright ownership.

In addition, the research methodology and system architecture described in the dissertation may serve as the basis for future patents or technology disclosures, particularly if the adaptive learning mechanisms or teacher-feedback analytics demonstrate novel technical contributions. To support long-term commercialization, intellectual property ownership may be retained by the research team or transferred to an institutional technology commercialization office depending on university policies.

7. BUDGET AND JUSTIFICATION

The following table presents the estimated budget required for the development and testing of the proposed system. As this is a university research project, the costs are limited to essential development tools, cloud services, and operational expenses needed to implement and evaluate the system prototype.

Table 3 - Estimated Research Project Budget

Description	Total Amount (Rs.)
Cloud Hosting for Backend Server (Node.js / FastAPI VPS – 6 months)	6,000
PostgreSQL Database Hosting (learning data storage – 6 months)	4,500
LLM API Usage (Gemini API for teacher feedback generation)	7,000
Mobile App Development Environment (React Native plugins / development utilities)	3,500
UI/UX Design Tools (Figma premium assets / design resources)	2,000
Internet & Utilities for Development (shared research cost)	3,000
Transport for School Visits / Usability Testing with Teachers	6,000
Miscellaneous (testing tools, adapters, unexpected research costs)	2,500
Total Budget	34,500

The proposed budget covers the essential technical and operational resources required to develop, deploy, and evaluate the adaptive learning system within a realistic environment. Cloud hosting and database services are necessary to run the backend server and store student learning data such as session logs, mastery scores, and progress records during system testing. Limited API usage costs are also included for the language model service used to generate teacher feedback summaries and reports based on student performance data. Additionally, small allocations are included for development utilities and interface design tools required for building the mobile application and designing sensory-friendly user interfaces suitable for autistic learners.

The remaining costs support the practical aspects of conducting the research and system validation. Internet and utility costs represent shared expenses associated with development and cloud deployment during the project period. Transport costs are included to facilitate visits to educational institutions or consultations with special education teachers during usability evaluation and feedback sessions. A small miscellaneous allocation is also included to account for unforeseen research-related expenses such as testing accessories or minor equipment needed during system

development. Overall, the proposed budget is modest and appropriate ensuring that sufficient resources are available to successfully implement and evaluate the proposed system.

8. WORK BREAKDOWN STRUCTURE

The overall module work breakdown is outlined below, providing a high-level overview of the key phases and tasks involved. This structure is intended to guide the development process and ensure a modular progression from design to implementation and validation.

Table 4 - Task Breakdown

Phase	Task Area	Description
1	Figma Design	Create UI/UX designs and user interface prototypes
2	User Management	Design and implement user roles and access control
3	Concept Asset Development	Create and organize learning/content assets
4	Knowledge Base & GNN Setup	Build graphical knowledge base and configure GNN
5	Backend & Database Foundation	Establish core backend architecture and database
6	Tier-Based Module Development	Develop modules based on defined system tiers
7	Teacher-Facing Features	Build features for teacher interaction and control
8	LLM Feedback Validation	Implement and validate LLM-driven feedback systems

9. GANTT CHART

The following Gantt chart outlines the planned timeline for the research project from November 2025 to December 2026, highlighting key academic milestones and component development tasks. It includes proposal submissions, progress evaluations, final report and research paper submissions, along with the implementation phases of the system such as UI design, content preparation, development of concept learning modules, adaptive learning mechanisms, and overall system integration.



Figure 2 - GANTT Chart

10. REFERENCES LIST

- [1] World Health Organization (WHO), “Autism Spectrum Disorder,” World Health Organization (WHO), 17 September 2025. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/autism-spectrum-disorders>. [Accessed 11 March 2026].
- [2] Y. M. Rohanachandra, D. M. A. Dahanayake, L. T. Rohanachandra and G. S. Wijetunge, “Knowledge About Diagnostic Features and Comorbidities of Childhood Autism Among Doctors in a Tertiary Care Hospital,” *Sri Lanka Journal of Child Health*, vol. 46, no. 1, p. 29–32, 2017.
- [3] Ada Derana, “9000 Children Identified as Autistic in Sri Lanka,” Ada Derana, 11 April 2025. [Online]. Available: <https://www.adaderana.lk/news.php?nid=107358>. [Accessed 11 March 2026].
- [4] D. Hariharan, S. K. Banerjee, A. M. A. Jinnah, S. B. Banu and G. S. P. Ghantasala, “A Comparative Study Revealing the Behavioural Difference Between Autistic and Healthy Adults,” in *International Conference on Communication and Signal Processing (ICCSP)*, Chennai, 2024.
- [5] C. Wang and H. Wang, “The growing challenge of autism spectrum disorder: a comprehensive review of etiology, diagnosis, and therapy in children,” *All Life*, vol. 17, no. 1, 2024.
- [6] M. Granpeesheh, J. Tarbox and D. R. Dixon, “Applied behavior analytic interventions for children with autism: A description and review of treatment research,” *Annals of Clinical Psychiatry*, vol. 21, p. 162–173, 2014.
- [7] J. L. Fuller and D. M. Fienup, “A Preliminary Analysis of Mastery Criterion Level: Effects on Response Maintenance,” *Behavior Analysis in Practice*, vol. 8, no. 1, p. 1–8, 2018.
- [8] A. Vijayan, . S. Janmasree, C. Keerthana and L. Baby Sylva, “A Framework for Intelligent Learning Assistant Platform Based on Cognitive Computing for Children with Autism Spectrum Disorder,” in *International CET Conference on Control, Communication, and Computing (IC4)*, 2018.
- [9] A. Hayat, T. Islam and M. Rahman, “Personalized Teaching Strategy Suggestion for Toddlers with ASD - A Machine Learning and Deep Learning Approach,” in *27th International Conference on Computer and Information Technology (ICCIT)*, 2024.

- [10] G. Ganju, A. Kumar and M. Uma, “Design And Development of Cognitive Skills Enhancement for Neurodivergent Children using CNN,” in *International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)* , 2024 .
- [11] P. V. Kumar, D. Balamurugan, K. Suriya and R. Reshma, “An Improved Learning Experience for English Language Using Big Data Analytics and AAC,” in *e International Conference on Sustainable Computing and Data Communication Systems (ICSCDS)*, 2023.
- [12] L. Hiryanto, M. Angelina, Jonathan, Hersinta, O. D. Hutagaol and Tony, “UI/UX Prototype of Visually Interactive Communication and Reading Aid for Autistic Children with Speech Disability,” in *Eighth International Conference on Informatics and Computing (ICIC)* , 2023.
- [13] G. S. Akbar, E. R. Kaburuan and V. Effendy, “User Interface (UI) Design of Scheduling Activity Apps for Autistic Children,” in *International Conference on Orange Technologies (ICOT)*, 2017 .

11. APPENDICES

This appendix presents the formal client requirements document prepared in consultation with the special education teachers and staff of Thusitha Suwasetha Special Education School, Gampaha. It outlines the fundamental system requirements identified during requirement gathering sessions, covering areas such as curriculum alignment, user management, teacher and student workspaces, interface personalization, and the concept learning module.

Development of an Intelligent Learning and Support System for Teaching English Language Fundamentals to Young Children with Autism in Sri Lanka

Institute :

Date : 12/03/2026

Sri Lanka Institute of Information Technology (SLIIT)
Malabe, Sri Lanka

Research Team :

- IT22084590 - H. H. Jayasundera
- IT22166838 - M. C. Gunawardena
- IT22216700 - R. S. Liluksha
- IT22266378 - N. Samaranayake

Client Organization :

Thusitha Suwasetha Special Education School
Gampaha, Sri Lanka

1. Purpose of the Document

This document formally records and validates the fundamental system requirements provided by the aforementioned client organization for the proposed research project.

The requirements outlined in this document represent the core expectations and needs identified by the special education Special Education Teachers and school representatives during requirement gathering discussions that were conducted by the research team of Sri Lanka Institute of Information technology (SLIIT). The research team confirms that the proposed system will be designed and developed based on these validated requirements.

The purpose of this document is to ensure a mutual understanding between the research team and the client organization regarding the scope and functionality of the proposed system.

2. Project Overview

The proposed research project aims to design and develop an interactive learning and support platform intended to assist in teaching English language fundamentals to children with Autism Spectrum Disorder (ASD).

The system should follow the Sri Lankan English language curriculum up to Grade 3 and should provide specialized learning activities tailored to the cognitive, behavioral and learning characteristics of children with autism.

The platform is expected to support special education Special Education Teachers in delivering structured learning activities, monitoring student progress and evaluating learning outcomes, while providing children with an engaging environment to develop their English language skills.

The ultimate goal of the proposed system is to support children with autism in developing the necessary English language competencies so that they can participate more confidently and inclusively in formal educational assessments and government examinations.

3. Fundamental Requirements

Based on discussions and consultations with the Special Education Teachers and staff of the aforementioned special education school, the following requirements were identified as the key expectations for the proposed system.

3.1. Curriculum Alignment

- The school requested that all learning activities included in the system align with the Sri Lankan English language syllabus currently practiced in the special education school, supporting the curriculum expectations for students up to Grade 3.
- It was also advised that the official Teacher's Guide should be carefully followed when selecting and developing the learning content, to ensure that the material taught through the system remains consistent with the curriculum used in the classroom.

3.2. System User Management

- The school requested that the system support three types of users: Principal, Special Education Teachers, and Autistic Students.
- It was indicated that the Principal will act as the System Administrator, responsible for creating and managing teacher and student accounts.
- Special Education Teachers should log into the system using accounts created by the Principal, so that they do not need to register themselves.
- Each Special Education Teacher should be able to work with a small group of assigned students in order to monitor their learning progress effectively.

3.3. Special Education Teacher Workspace

- After logging in, Special Education Teachers should have access to a workspace where they can view their assigned students and basic information related to them.
- Special Education Teachers should be able to select a student and begin a learning session through this workspace.

3.4. Student Workspace

- The system should provide a student workspace where the selected student can participate in the learning activities.
- The workspace should display the four main learning modules and activities available for the student.

3.5. User Interface Personalization

- Special Education Teachers suggested that the system should allow basic personalization options so that children feel comfortable while using the platform.
- During the first learning session, students should be able to choose an avatar that will guide them during learning activities.
- Students should also be able to select suitable font styles and sizes that are comfortable for them.
- These preferences should be saved and applied during future learning sessions.

3.6. Concept Learning Module

- Special Education Teachers requested that the system include a concept learning module to help children learn basic English vocabulary.
- The learning content should follow the Sri Lankan primary education curriculum, particularly the Grade 1 and Grade 2 Listening and Speaking Teacher's Guides.
- Special Education Teachers indicated that concepts should be taught gradually using simple and interactive learning activities.
- Learning should begin by introducing the concept of using images and audio to help children become familiar with the word.
- Students should then practice recognizing the concept and understanding simple attributes related to it.
- A new concept should be introduced after the child has completed the learning stages of the previous concept.

Figure 3 - Appendix A1 - Client Requirements Document