

**AI-POWERED MOBILE AUGMENTED REALITY
EDUCATIONAL SYSTEM FOR PRIMARY SCHOOL STUDENTS:
GEOMETRY AND ENGLISH LITERACY LEARNING**

GeoLex AR | Project ID R26-IM-005

COMMON INTEGRATED SOLUTION REPORT

Research team and component responsibilities

Research member	Student ID	Component
Himsara M.K.D.L.	IT22274298	C1: 2D-to-3D Geometry Validation
De Costa M.M.P.	IT22179876	C2: 2D Net Correction and Virtual Slicing
Abhilash K.A.T.W.	IT22081698	C3: AI Air-Writing Trajectory Tracking
Wickramasingha N.N.	IT22214102	C4: Sketch-Driven Multisensory Vocabulary Learning

Report submitted in partial fulfilment of the requirements for the
B.Sc. (Hons) Degree in Information Technology
Specialisation in Interactive Media
Department of Information Technology
Faculty of Computing
Sri Lanka Institute of Information Technology
Sri Lanka

October 2026

DECLARATION AND SUPERVISOR CERTIFICATION

We declare that this common integrated research report represents the work of the named research group and that all materials originating from other authors have been acknowledged. To the best of our knowledge, the report does not incorporate, without acknowledgement, any material previously submitted for a degree or diploma at another institution. We grant the Sri Lanka Institute of Information Technology the non-exclusive right to reproduce and distribute the dissertation in whole or in part, while retaining the right to use our original material in future academic work.

Academic integrity and drafting disclosure: AI-assisted drafting and formatting support was used in preparing this editable report. The researchers remain solely responsible for verifying the technical descriptions against the final application, checking each citation and numerical result, ensuring that any images are used with appropriate consent, and signing only after reviewing the entire document in accordance with institutional policy.

Candidate approval and signature record

Research member	Student ID	Signature	Date
Himsara M.K.D.L.	IT22274298	_____	_____
De Costa M.M.P.	IT22179876	_____	_____
Abhilash K.A.T.W.	IT22081698	_____	_____
Wickramasingha N.N.	IT22214102	_____	_____

Supervisor certification: The above candidates carried out the research project under my supervision. The final copy is subject to supervisor review and formal approval.

Supervisor and co-supervisor certification

Role and name	Signature	Date
Supervisor: Mr. Aruna Ishara Gamage	_____	_____
Co-Supervisor: Mr. Nushkan Nisme	_____	_____

ABSTRACT

Primary-school learners often find it challenging to translate two-dimensional geometric representations into three-dimensional concepts and to connect English letter movement, sketches, spelling and spoken vocabulary within engaging activities. This study designed and piloted GeoLex AR, an integrated Android-based educational suite organized around a shared learner-input–analysis–feedback–reinforcement interaction architecture. Four specialized components instantiate the design: (C1) rule-based 2D-to-3D Geometry Validation, (C2) 2D Net Correction and Virtual Slicing, (C3) camera-based AI Air-Writing Trajectory Tracking and (C4) Sketch-Driven Multisensory Vocabulary Learning. The application used Unity, C#, Android AR technologies and component-specific recognition or geometric processing. A Design Science Research approach guided development and evaluation. Fifty unique learners participated: 25 aged 9–12 evaluated both geometry components and 25 aged 6–8 evaluated both literacy components. In the recorded evaluation, C1 agreed with expected validity labels for 25/25 nets; C2 completed 20/25 Virtual Slicing sessions (80%); C3 correctly recognized 19/25 air-written letter trials (76%); and C4 completed 17/25 full spelling, pronunciation-oriented and AR-reward sequences (68%). Reported mean pre-/post-test results changed from 5/10 to 9/10 for C1, C2 and C3, and from 4/10 to 8/10 for C4. These improvements are descriptive, not causal estimates, and the four technical percentages represent different measures. The work demonstrates the feasibility of an integrated, age-specific educational interaction framework that interprets learner-generated input before providing guided feedback and AR/3D reinforcement, while identifying important needs for larger samples, independent correction evaluation, comprehensive recognition benchmarks and longitudinal study.

Keywords: Augmented Reality; Air-Writing Recognition; Geometric Net Validation; Sketch Recognition; Virtual Slicing

ACKNOWLEDGEMENTS

The research group sincerely acknowledges the guidance of supervisor Mr. Aruna Ishara Gamage and co-supervisor Mr. Nushkan Nisme throughout the conception, design, development and evaluation of the GeoLex AR research project. Their feedback contributed to the refinement of the component scope, integration narrative, experimental reporting and final academic presentation.

The group thanks the Faculty of Computing and Department of Information Technology at the Sri Lanka Institute of Information Technology for providing the academic environment for this Final Year Research Project. Appreciation is extended to the teachers, schools, families and other facilitators who supported the educational testing activities, together with the learners whose participation made evaluation of age-appropriate interaction possible. The research acknowledges the importance of parental/guardian consent, learner assent and responsible handling of information concerning minors.

The four members also recognize their mutual contribution to component implementation, testing, writing, design review and integration discussions. Any formal confirmation of individual effort, signatures and institutional approvals should be completed by the team before submission.

TABLE OF CONTENTS

1 INTRODUCTION AND PROJECT CONTEXT	1
1.1 Background and context	1
1.2 Stakeholders and user needs	2
1.3 Alignment with Sustainable Development Goal 4	2
1.4 Purpose, boundaries and evidence hierarchy	3
2 OVERALL PROBLEM DEFINITION AND FOUR SUBPROBLEMS	4
2.1 Common educational problem	4
2.2 Literature synthesis and research gap	4
2.3 Subproblem and component mapping	5
2.4 Research aim, objectives and questions	6
2.5 Scope, assumptions and constraints	7
3 OVERALL SOLUTION STRUCTURE	8
3.1 Unified educational interaction model	8
3.2 High-level architecture and deployment boundaries	9
3.3 Component 1: Geometry Validation	10
3.4 Component 2: Net Correction and Virtual Slicing	10
3.5 Component 3: Camera-Based Air-Writing	11
3.6 Component 4: Sketch-Driven Vocabulary	12
3.7 End-to-end learner and information flows	12
4 COMPONENT INTEGRATION AND COHESIVE OPERATION	14
4.1 Integration strategy and boundaries	14
4.2 Interfaces, state hand-offs and logging	15
4.3 Shared non-functional requirements	16
4.4 Error handling and recovery	17
4.5 Team ownership and dependencies	18
5 INTEGRATED IMPLEMENTATION AND PROJECT EXECUTION	20
5.1 Common technologies and rationale	20
5.2 Development evolution and design decisions	21
5.3 Deployment and research data handling	22
5.4 Verification and provenance of implementation	23
6 INTEGRATED TESTING, EVALUATION AND VALIDATION	24
6.1 Research design and participating cohorts	24
6.2 Technical and workflow evaluation	25
6.3 Age-specific learning assessment	27
6.4 Interpretation by research question	28
6.5 Objective achievement and integration evidence	29
6.6 Threats to validity and evidence boundaries	30
7 COMMERCIALISATION, SUSTAINABILITY AND RESPONSIBLE INNOVATION	31
7.1 Educational value and beneficiaries	31
7.2 Alternatives and differentiation	31
7.3 Proposed adoption and business pathways	32
7.4 Ethics, privacy, licensing and access	33
7.5 Sustainability and SDG 4 contribution	34
8 CONCLUSIONS AND FUTURE DIRECTIONS	35
8.1 Overall conclusion	35
8.2 Integrated research contribution	35
8.3 Priorities for further evaluation and deployment	36
REFERENCES	37
APPENDICES	39

LIST OF FIGURES

Figure 3.1 Unified learner-input-analysis-feedback-reinforcement architecture.....	8
Figure 3.2 Geometry and literacy workflow boundaries.....	13
Figure 4.1 Integrated C1-C2 decision and correction loop.....	18
Figure 6.1 Recorded mean pre-test and post-test scores	28
Figure 6.2 Recorded component completion measures (different metrics).....	28
Figure E.1 Representative application interface montage	42

LIST OF TABLES

Table 2.1 Four component problems, inputs and outputs.....	6
Table 2.2 Scope boundary and evidence interpretation.....	7
Table 3.1 Component processing and reinforcement responsibilities	13
Table 3.2 Supported geometry solids and slicing options	13
Table 4.1 Component interface and evidence mapping	15
Table 4.2 Shared system requirements and verification status.....	16
Table 5.1 Common implementation technology inventory	20
Table 6.1 Cohort composition and evaluation activities	24
Table 6.2 Component-level technical and workflow results	25
Table 6.3 Direction and fraction completion in C2.....	26
Table 6.4 C3 and C4 outcome detail.....	27
Table 6.5 Pre-/post-test mean scores.....	28
Table 6.6 Group objectives and associated evidence	29
Table 7.1 Adoption risks and proposed mitigations	33
Table A.1 Research evidence register	39
Table B.1 Integration acceptance checklist.....	39

LIST OF ABBREVIATIONS

Alphabetical abbreviations used in the report

Abbreviation	Expanded term
AI	Artificial Intelligence
AR	Augmented Reality
ARCore	Google mobile augmented reality platform
C1–C4	Research Components 1 to 4
CNN	Convolutional Neural Network
CSV	Comma-Separated Values
DSR	Design Science Research
HCI	Human–Computer Interaction
IEEE	Institute of Electrical and Electronics Engineers
LTS	Long-Term Support
MR	Mixed Reality
SDG	Sustainable Development Goal
SLIIT	Sri Lanka Institute of Information Technology
UAT	User Acceptance Testing
UI	User Interface
UX	User Experience
XR	Extended Reality

1 INTRODUCTION AND PROJECT CONTEXT

1.1 Background and context

Geometry and foundational English literacy are distinct domains, but each requires children to interpret representations and respond to feedback. In geometry, a learner must infer how an arrangement of two-dimensional (2D) surfaces relates to a three-dimensional (3D) object and how that object behaves when divided. In literacy, the learner coordinates letter movements, object–word associations, spelling and spoken-word practice. Static textbook diagrams, copying exercises and teacher demonstrations remain valuable, but alone they offer limited opportunities for individual input to be analyzed immediately. A supplementary learning environment can give children additional opportunities to construct, test and revise their own understanding [1], [2].

Augmented Reality (AR) can overlay 3D learning objects onto a camera view of the physical environment, making otherwise abstract forms inspectable. However, simply showing a ready-made 3D object does not ensure that the learner has understood the idea. The important educational question is what the learner does before the digital reinforcement is presented: whether the learner creates a representation, whether the application checks that representation and whether feedback helps the learner make sense of errors. A synthesis of AR education studies indicates that outcomes depend on task design, user characteristics and implementation rather than on the display technology alone [1], [2].

GeoLex AR was designed as one Android educational suite serving primary-school students aged 6–12 through two age-specific pathways. Geometry activities address learners aged 9–12, and English-literacy activities address learners aged 6–8. It provides five supported geometric solids, two linked geometry components and two different literacy activities. The research contribution is an integrated learning architecture: learner-generated input is first analyzed by a specialized process, an instructional state is produced and targeted feedback precedes 3D/AR reinforcement. The applications of geometric rules and Convolutional Neural Networks (CNNs) are different, but the educational interaction principle is shared.

1.2 Stakeholders and user needs

The primary stakeholders are children performing the activities. Geometry learners need clear feedback when a net is invalid, readable controls when working with 3D solids, and opportunities to inspect fractions and cross-sections. Younger literacy learners need accessible input activities, meaningful object associations, recognizable feedback and the ability to retry without complex text instructions. Both groups benefit from large touch targets, predictable navigation and guidance that distinguishes incomplete work from incorrect work. As noted in the K–12 AR usability literature, age suitability is a necessary design consideration rather than a secondary cosmetic feature [2].

Teachers represent a second group of stakeholders because the application is intended to supplement guided instruction. Teachers may use it to illustrate nets, discuss why a construction is incorrect, demonstrate a cut through a solid or invite children to practise symbols and vocabulary. Parents and guardians are stakeholders because they authorize children’s participation in evaluation and may support supervised learning at home. Schools and educational institutions are potential adoption partners, but this pilot did not establish school procurement willingness or large-scale deployment readiness.

Researchers and the project team require different system properties: reproducible testing states, stable navigation and clearly linked participant and session records. In particular, the geometry Virtual Slicing module recorded selection events, trace attempts and session summaries. This emphasis on both instructional output and traceable research evidence helps avoid treating a child-facing prototype as if it were already a validated mass-market learning service. The study also carries heightened ethical responsibilities because participants are minors, and identifiable information should be excluded from public reporting unless the use is covered by appropriate permission.

1.3 Alignment with Sustainable Development Goal 4

The research is aligned with Sustainable Development Goal (SDG) 4, Quality Education. Its intended contribution is to explore whether accessible Android-based interactive media can complement conventional teaching of spatial geometry and basic English literacy. The application uses phones and familiar interaction modes rather than requiring specialized headsets or wearable sensors for the evaluated activities. Such design choices may reduce some barriers to experimentation in classrooms where equipment budgets are constrained. Nevertheless, an Android AR application still depends on compatible hardware, adequate lighting and teacher support, so affordability and inclusivity cannot be assumed from platform selection alone.

The pilot contributes design knowledge and bounded observations: recorded completion, recognition and pre-/post-test means for the participating age groups. It did not measure school-wide equity, cost-effectiveness, long-term retention or educational attainment relative to a randomized control group. Claims about SDG 4 should therefore be framed as prospective educational relevance and a direction for responsible adoption rather than demonstrated population-level impact.

1.4 Purpose, boundaries and evidence hierarchy

This common report describes the integrated solution, its four differentiated pipelines, architectural relationships, common implementation, documented evaluation and plausible value for education. Detailed training hyperparameters, shape-controller internals and individual programming histories belong in the corresponding student theses and source repositories. The integrated IEEE-style research paper supplies the final quantitative result definitions, while the topic-assessment form and March 2026 proposals preserve the original research intention. Where a proposal feature was subsequently revised, the final evaluated architecture takes precedence.

Evidence is classified as verified final outcomes, implemented but unmeasured features, or historical proposals. C1 uses deterministic geometry rules; C2's numerical result concerns slicing rather than correction efficacy; C3's learner trials involve only sampled uppercase letters; and C4's completion result is not sketch recognition accuracy. The report maintains these boundaries to avoid overstating its findings.

2 OVERALL PROBLEM DEFINITION AND FOUR SUBPROBLEMS

2.1 Common educational problem

The common educational problem is limited instructional feedback on learner-generated representations across age-appropriate geometry and English-literacy activities. Geometry learners may copy a cube net without understanding why its faces can or cannot fold into a closed solid. A conventional right/wrong response identifies the outcome but not the faulty relationship. Similarly, a child who is learning an English letter can produce a recognizable image while still being unsure about its movement sequence, and a child who sketches an object may not consistently connect it with a written or spoken word. In each case, the learning task involves more than receiving information: it involves producing an input that can be interpreted.

The design problem is not that AR, CNNs or mobile apps are unavailable. The literature shows demonstrated uses of AR geometry [3]–[7], sketch-based recognition [8]–[10], vision-based air-writing [11]–[15] and AR-supported alphabet/vocabulary learning [16]–[18]. What remains of interest is an interaction framework in which a child actively creates a representation, the system evaluates or recognizes it, feedback addresses the state of the activity, and only then is an appropriate 3D/AR reinforcement or progression offered. This framework must accommodate very different data representations rather than imposing a single universal model.

2.2 Literature synthesis and research gap

Geometry research reports the usefulness of manipulable 3D representations, but an application that starts with a predefined object does not necessarily reveal whether a child can construct a mathematically correct net [3]–[5]. Intelligent AR tutoring indicates the benefit of instructional support, while work comparing different AR displays demonstrates that learning effects depend on task characteristics [6], [7]. Sketch reconstruction research supports transforming hand-drawn representations into digital forms [8], yet visual plausibility and true foldability are different properties: an attractive image is not automatically a valid net. This motivates formal validation of

topology, face counts and compatible edges before rewarding a construction with a solid.

For literacy, earlier air-writing work established camera-based hand detection and trajectory analysis as viable recognition inputs [11]–[15]. The path, stroke sequence and duration of writing carry information that is lost in a final raster of the character. Vocabulary studies suggest useful roles for AR visual association and motivation [16]–[18]. However, recognizing a sketch alone does not complete a vocabulary-learning sequence; a learning design must guide children from identified object to spelling and pronunciation-oriented participation and reinforcement. The gap addressed by GeoLex AR is therefore chiefly architectural and pedagogical: connecting component-specific interpretation with a consistent learner-generated input and feedback philosophy.

The four components are not interchangeable. C1 provides mathematical acceptance of nets; C2 interprets invalid states and supports physical-space slicing; C3 processes camera-tracked hand trajectories; and C4 interprets object sketches and manages literacy progression. An integrated environment can use common navigation and presentation principles while preserving the specialized nature and validity criteria of each procedure. This avoids the incorrect conclusion that one AI model solves all four learning problems.

2.3 Subproblem and component mapping

The first subproblem is whether a learner-created arrangement of faces meets the structural requirements of a chosen solid. C1 addresses this through rule-based analysis and valid/invalid decisions before 3D output. The second concerns both helping learners identify the incorrect relationship in an invalid 2D net and helping learners understand how a valid 3D object is partitioned; C2 addresses these with correction guidance and Virtual Slicing. The third concerns the recognition of letters written in free space using camera-observed fingertip movement; C3 processes ordered trajectories. The fourth concerns connecting a child’s own object sketch with English vocabulary through recognition, spelling, pronunciation-oriented tasks and a gated AR reward; C4 manages this sequence.

Table 2.1 Four component problems, inputs and outputs

Component / domain	Learner-generated input	Analysis mechanism	Instructional output
C1 Geometry	2D geometric net	Deterministic shape rules	Validity, 3D/AR when valid
C2 Geometry	Invalid net; direction, fraction, trace	Fault guidance; trace interpretation	Correction guide; separated 3D pieces
C3 Literacy	Camera-based fingertip trajectory	Trajectory normalization; CNN recognition	Predicted letter, retry feedback
C4 Literacy	Touchscreen object sketch and tasks	CNN object label; task gating	Vocabulary progression and AR reward

2.4 Research aim, objectives and questions

The overall aim was to design, develop and evaluate an Android AI-assisted AR suite for primary-school geometry and English literacy in which learner-generated inputs are analyzed before corrective feedback and digital reinforcement. The corresponding group objectives were: (O1) validate 2D nets and show valid solids; (O2) provide error guidance and interactive fractional slicing; (O3) classify camera-based English air-writing trajectories; (O4) link learner-created sketches to vocabulary tasks and AR rewards; (O5) connect these activities through one educational interface; and (O6) evaluate component-specific workflow, technical outcomes and descriptive learning changes with the intended age cohorts.

Six research questions operationalize the aim. RQ1 asks about agreement between C1 decisions and expected net validity. RQ2 asks about execution of Virtual Slicing and its choices, with independent net-correction efficacy recognized as unmeasured in the final quantitative log. RQ3 asks about correct air-written letter recognition and responsiveness in the evaluated sample. RQ4 asks about completion of sketch-to-spelling/pronunciation-to-AR workflow states. RQ5 examines changes in mean age-specific assessment scores. RQ6 considers observed ease of interaction through completion, retries and timings without claiming a validated usability scale was administered.

2.5 Scope, assumptions and constraints

Geometry is restricted to Cube, Cuboid, Triangular Prism, Tetrahedron and Cylinder; Sphere was removed from the final scope. Virtual Slicing offers Vertical, Horizontal and Angled directions and fractions $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{4}$ and $\frac{3}{4}$. The final study contains 25 learners aged 9–12 for the paired geometry components and a different set of 25 learners aged 6–8 for the paired literacy components. Evaluation of component implementation, observed learner workflow and educational score change must remain separate concepts. The combined 100 component-level records do not represent 100 different children.

Table 2.2 Scope boundary and evidence interpretation

Design point	Final scope / evidence constraint
Net validity	Rule-based; 25/25 recorded correct labels, not CNN accuracy
Net correction	Supported design/implementation; no independent controlled efficacy estimate
Virtual Slicing	Five solids, three directions and four fractions; 20/25 sessions completed
Air-writing	Camera-based, not touchscreen-only; 25 child-user trials, 17 uppercase targets
Vocabulary	Five object classes; 17/25 complete workflows, not sketch accuracy
Learning scores	Group means only; no paired inferential test or causal inference

3 OVERALL SOLUTION STRUCTURE

3.1 Unified educational interaction model

GeoLex AR is organized around a shared educational contract rather than a common recognition algorithm. A learner first supplies an input: a geometry net, a choice and trace on a solid, a free-space hand trajectory or a drawing of a familiar object. A component-specific processor then evaluates that input according to the relevant rules or trained model. Its output becomes an instructional state, such as valid/invalid, a correction hint, a predicted letter or a task completion condition. The system returns an age-appropriate response and, where the activity permits, reinforces the outcome through manipulable 3D content or AR visualization. The sequence can be expressed as Learner Input → Computational Analysis → Instructional Feedback → Interactive Reinforcement.




Figure 3.1 Unified learner-input-analysis-feedback-reinforcement architecture.

Figure 3.1 distinguishes the processing branches from the common presentation layer. In particular, C1 uses geometric constraints, whereas C3 and C4 contain separate AI recognition functions. This distinction avoids the impression that all learner actions

are evaluated by a single classifier. Navigation and presentation provide consistency without suppressing differences in input modality, dataset requirements, task progress or success conditions. The architecture is extensible in principle because another learning activity could adopt the same pattern while supplying its own analysis and feedback logic; the present study, however, evaluates only the four documented components.

3.2 High-level architecture and deployment boundaries

The application was developed in Unity 2022.3 LTS using C# for Android. AR Foundation and the ARCore XR Plugin provide mobile AR capabilities including plane detection, raycasting and placement for supported Android devices. The C2 development environment used Unity 2022.3.62f3 LTS, and a Xiaomi Redmi Note 9 Pro was used during development and mobile testing of that component. Mobile camera access is required for hand-trajectory acquisition and for AR world understanding where applicable. Physical-space AR and on-screen drawing should not be confused: C3 specifically collects free-space finger movement through the camera, while C4 starts with an object sketch on a drawing interface.

A layered interpretation is useful. The interaction layer presents welcome, learner entry and age-appropriate activity controls. The acquisition layer captures touch geometry, touch tracing, video-based hand positions and sketches. The domain-processing layer applies rule-based net checks, slicing and correction control, trajectory recognition, or object classification and literacy gates. The output layer renders highlights, retry prompts, solids, separated pieces, letters or AR rewards. Finally, data logging records component-relevant events and summaries. While these layers represent the integrated system, any assertion about one uniform shared database, common cloud runtime or synchronized learner-profile service would exceed the final evidence.

The original proposals described several optional technologies such as a universal cloud analytics service, fully offline inference for each AI subsystem, detailed mesh-generation internals and specific speech architectures. The final integrated paper confirms Unity, Android AR technologies and the processing elements stated in this

report, but not all proposed implementation details. Therefore, deployment boundaries are presented conservatively: local Unity/Android execution for the AR and interactive application, a documented Python-based C4 model development/integration pipeline, and component-specific research records. The final exact inference placement for each model should be checked against the release configuration before public technical documentation is finalized.

3.3 Component 1: Geometry Validation

C1 addresses whether a learner-created net can make the chosen three-dimensional solid. It uses deterministic, shape-specific geometric validation and does not depend on a CNN image classifier for the final validity decision. Checks consider face count, adjacency, connectivity, orientation, dimensional compatibility and the structural requirements of the target solid. For a Cube and Cuboid, the relationship and orientation of faces are central; a Triangular Prism requires suitable rectangular lateral surfaces with matching triangular ends; a Tetrahedron requires a compatible connected arrangement of four triangular faces; and a Cylinder requires two compatible circular ends and an appropriate rectangular lateral surface.

When the net is valid, C1 proceeds to 3D representation and interactive model inspection or AR placement. When it is invalid, the application does not treat the incorrect geometry as if it were a verified solid; instead, the invalid state is passed toward C2's correction workflow. This distinction helps preserve instructional integrity: a child should not be rewarded with a realistic-looking solid as proof that an unverified net was correct. The final recorded sample contained five trials per shape, including four expected-valid and one expected-invalid case for each solid. In all 25 recorded trials, the system result agreed with the stored expected label. That outcome supports correctness for the observed examples but does not prove universal correctness across every possible drawing or geometric-net arrangement.

3.4 Component 2: Net Correction and Virtual Slicing

C2 provides two related learning functions. For invalid nets, a 2D corrective workflow identifies a problematic geometric relationship, highlights it and offers a short explanation and dotted or other visual guide so the child can revise the construction.

This was a meaningful design revision from the original proposal: correction is applied to the 2D sketch/net before 3D generation rather than relying on comparison between incorrect and correct 3D ghost objects as the principal feedback mechanism. A ghost or reference representation may be used after correct generation as an inspection aid, but it is not the final core method for detecting a bad net.

The Virtual Slicing workspace enables learners to choose one of five supported solids, select Vertical, Horizontal or Angled slicing, choose 1/2, 1/3, 1/4 or 3/4, and trace an onscreen guide. A valid trace triggers a prepared slice outcome, displaying separated pieces that learners can inspect from alternative viewpoints. Plane detection supports AR placement. Development considered physics-based separation and falling toward a detected plane, but the final study should not be read as a controlled verification of one specific rigidbody or dynamic-mesh method. The recorded quantitative result—20 successful sessions out of 25—covers the Virtual Slicing task. It is not an 80% success rate for fault-specific net correction.

3.5 Component 3: Camera-Based Air-Writing

C3 supports the practice of English letters through actual camera-based writing in free space. Instead of drawing on a touchscreen, the child moves an index finger in front of the phone camera. Hand-landmark tracking locates the fingertip in successive frames, producing an ordered path. Preprocessing normalizes the trajectory, after which a CNN-based recognition pipeline predicts the intended letter and provides feedback so that the learner can retry. Path history, starting region, direction and stroke sequence matter as instructional information even though the final trial outcome is reported as a correct or incorrect letter prediction.

The system design supports uppercase and lowercase English letters, but only 17 sampled uppercase target letters were evaluated in the final recorded child-user study. The final logged recognition rate of 19/25 (76%) therefore has a defined denominator and cannot be generalized into complete recognition accuracy for all 52 uppercase and lowercase classes. Several details described in the March proposal, including one exact Conv1D architecture, inference format, canonical comparison thresholds and always-on ghost arrows, remain historical design details unless corroborated in final source

code. This report discusses the verified recognition pipeline and outcomes rather than presenting unverified hyperparameters as reproducible specifications.

3.6 Component 4: Sketch-Driven Vocabulary

C4 starts from a learner-generated drawing of Apple, Cat, Fish, Sun or Tree. An image-processing and CNN classification pipeline produces the predicted object label, which becomes the context for vocabulary learning. The application then presents spelling and pronunciation-oriented activities. Both completion conditions must be accepted before the AR reward becomes accessible. The corresponding 3D object can then be placed or inspected in AR, connecting the child's drawing with a written and spoken word. The focus is multisensory vocabulary learning, not a demonstrated formal phonics curriculum that teaches grapheme–phoneme decoding.

C4 model development used a Quick, Draw!-derived set of 25,000 images comprising five classes with 5,000 examples each. The recorded split was 17,500 training, 3,750 validation and 3,750 held-out test images. Python, TensorFlow/Keras, OpenCV, FastAPI, NumPy and Pillow were used in the development/integration pipeline, with Unity/C# providing the interactive experience. The final child-user study recorded 25 sketches and downstream activities. Although its mean prediction confidence was 0.832, the absence of an independent intended-object label for each evaluated child sketch prevents calculation of real child-sketch recognition accuracy from those session logs. The observed 17/25 figure is full educational workflow completion, not model accuracy.

3.7 End-to-end learner and information flows

The architecture contains two meaningful pathway boundaries rather than one serial flow through C1, C2, C3 and C4. Within geometry, C1 determines whether the learner-created net is acceptable. A rejected net can trigger C2 visual correction, revision and revalidation; a valid representation permits 3D inspection and the distinct Virtual Slicing activity. Within literacy, the child selects either air-writing or sketch-driven vocabulary tasks. C3 produces recognition feedback based on the fingertip trajectory; C4 progresses a sketch label through spelling and pronunciation-oriented acceptance into an AR reward. The shared Unity experience supports common navigation and

presentation, but there is no requirement that children complete all four activities sequentially to obtain one result.




Figure 3.2 Geometry and literacy workflow boundaries in the integrated application.

Table 3.1 Component processing and reinforcement responsibilities

Component	Primary analysis	Feedback or reinforcement
C1	Rule-based geometric net constraints	Accept valid construction or route invalid one to C2
C2	Fault guidance and slice trace state	2D hints and separated slice result
C3	Camera trajectory and CNN letter prediction	Recognition feedback and retry
C4	Sketch CNN and paired task acceptance	Vocabulary progression and AR object reward

Table 3.2 Supported geometry solids and slicing options

Parameter	Final supported configuration
Geometric solids	Cube; Cuboid; Triangular Prism; Tetrahedron; Cylinder
Slice directions	Vertical; Horizontal; Angled
Slice fractions	1/2; 1/3; 1/4; 3/4
Excluded shape	Sphere (removed from final scope)
Fraction interpretation	Task-specific selections; neither equal-volume proof nor mesh method is inferred from label alone

4 COMPONENT INTEGRATION AND COHESIVE OPERATION

4.1 Integration strategy and boundaries

A deliberate separation between domain-specific analysis and shared interaction behavior guided the integration. C1's geometric decisions have a direct downstream relationship with the C2 correction route. This means that a valid/invalid result should be accompanied by enough shape and fault context for the next screen to display appropriate guidance. A failure to convey this state can create the impression that a child is correcting the wrong solid, or can allow 3D reinforcement to appear before the mathematical check. By contrast, the literacy components do not need to transfer a single classifier output between them for their core evaluations. They share the mobile environment and educational philosophy, while preserving distinct inputs and outcomes.

Practical integration involves four types of alignment: terminology, screen progression, state continuity and data provenance. The term Triangular Prism is used consistently in shape selection, result reporting and formal documentation. A child should be able to return from a result to retry a task without resetting an unrelated component unexpectedly. The name entered at login should be accepted as any nonempty text under the revised application requirement, while a session identifier remains distinct and stable. Finally, a researcher must be able to reconstruct the sequence of choices, failed attempts and successful completion rather than receiving only a final status. These are practical design contracts for integrated educational reliability.

The project proposals contain descriptions of 3D ghost-model comparisons as a primary correction approach and of relationships between air-writing words and sketch recognition that are broader than the final tested scope. Integration documentation therefore must label such descriptions as earlier proposals where applicable. The authoritative final architecture uses 2D net fault guidance for C2, camera-driven trajectory input for C3 and completion-gated object vocabulary for C4. Conflating these stages would misrepresent component interfaces and measured results.

4.2 Interfaces, state hand-offs and logging

At a conceptual interface level, C1 produces a selected solid and a validity state, together with information needed to display an acceptable or unacceptable result. C2 consumes incorrect geometry information to construct visual guidance and supports independent slicing states for shape, direction, fraction and trace status. C3 produces the target letter, predicted letter, confidence and recorded interaction characteristics. C4 uses the recognized object label, spelling state, pronunciation-oriented acceptance state and an AR unlock flag. The final integration paper does not publish a stable cross-component JSON schema or compiled API signatures; accordingly, this report does not fabricate endpoint names or field definitions beyond the documented responsibilities.

C2's research implementation also provides a useful example of preserving the distinction between one session and multiple interaction events. A unique Session ID links choices and trace attempts. Learner login now accepts any nonempty name rather than requiring the earlier ST0001–ST0025 label format at entry, while the standardized historical evaluation datasets retain their anonymized identifiers. A repeated fraction selection is meaningful even when the final chosen fraction is unchanged; similarly, a failed trace should not disappear simply because the learner later succeeds. The C2 local CSV export concept includes event and session summaries that are copied from persistent application storage to researcher-accessible files, and an existing CSV schema should be preserved during minor UI changes.

Table 4.1 Component interface and evidence mapping

Boundary	Required state concept	Evidence recorded or verified
Login → activities	Nonempty entered name; distinct session identity	C2 session and event pipeline; exact deployed schema verify
C1 → C2 correction	Selected solid; invalid net and fault relation	Documented route; no separate efficacy trial
C1 → 3D/AR	Valid state permits 3D generation	25 recorded trials matched expected status
C2 choice → result	Shape; direction; fraction; trace status	25 Virtual Slicing logs and 20 successful outcomes
C3 recognition → feedback	Trajectory; target/prediction; attempts	25 trials, 19 correct letters
C4 tasks → AR gate	Object class; spelling and spoken task acceptance	17 complete AR-unlock and placement sessions

4.3 Shared non-functional requirements

A common child-facing AR experience requires more than algorithmic correctness. UI controls must accommodate relatively imprecise touch input, foreground text should remain legible against classroom lighting, screens must offer clear retry options and activities should progress without unnecessary data entry. In geometry, failure feedback should identify the problematic relationship rather than presenting only a red error indicator. In literacy, a nonrecognition outcome should allow repetition rather than being confused with an accepted character. Human–computer interaction considerations are especially important because the two participating cohorts have different cognitive and reading demands [2].

Mobile reliability depends on supported Android and ARCore hardware, camera permission and an adequately detected physical surface. The design should prevent meaningful slicing before appropriate placement rather than silently completing it without the AR context. Research integrity depends on avoiding accidental duplication of session IDs, retaining unsuccessful interactions where recorded and exporting copies of data without deleting a master dataset. Privacy depends on pseudonymization and restricted access to children’s records and any photographs. Data security and actual encryption behavior should be assessed from final deployment artifacts instead of assuming that statements in the proposals were fully implemented.

Table 4.2 Shared system requirements and verification status

Requirement	Reason for the integrated system	Status in the supplied evidence
Child-friendly UI	Accessible instructions, touch targets and retry	Design priority, qualitative evidence from workflow
Correctness gating	Invalid net must not generate incorrect 3D/AR	Confirmed in recorded C1 outcomes
AR plane detection	Physical placement of 3D geometry	Documented Android AR implementation
Distinct sessions	Link repeated choices and attempts	C2 session and CSV design documented
Runtime speed	Responsive classroom activity	Recorded timing statistics; full device benchmark unavailable
Data privacy	Participants are minors	Consent and anonymization documented; deployment audit needed
Accessibility	Different ages, visual contrast and text needs	Design objective; no validated accessibility trial reported

4.4 Error handling and recovery

A learning-oriented error state is valuable when it enables a correction attempt. In C1, an invalid net must not be presented as a valid solid. In C2, a learner can request a guide, adjust the net and submit it for revalidation; during slicing, the learner can repeat an incomplete or incorrect trace, reset choices, repeat a slice or change the solid. In C3, failure to predict a letter should remain a recorded unsuccessful attempt and prompt the child to try again. In C4, completion of spelling alone or pronunciation-oriented acceptance alone must not open the final AR reward. These constraints make learning progression coherent and prevent misleading successful states.




Figure 4.1 Integrated C1–C2 validation, correction and Virtual Slicing decision loop.

An implementation-oriented acceptance test should begin with an intentionally invalid geometric net and confirm the presence of a correction route, then repeat with a valid net and verify that a correct 3D output is available. Other useful tests include an incomplete C2 tracing path, two consecutive sessions with the same name, a C3 no-prediction result and a C4 session in which only one prerequisite task succeeds. The

present report can specify these as regression requirements; it does not claim that a full independent security and automated regression audit was performed when the corresponding test logs have not been supplied.

4.5 Team ownership and dependencies

The approved topic-assessment form assigns Geometry Validation to Himsara M.K.D.L., Net Correction and Virtual Slicing to De Costa M.M.P., Air-Writing to Abhilash K.A.T.W. and Sketch-Driven Vocabulary to Wickramasingha N.N. This division enables focused algorithmic development while retaining shared responsibilities for navigation, academic integration, ethical evaluation and communication of findings. C1 and C2 have a strong processing dependency because the correction pathway depends on the result of validating a net. C3 and C4 have conceptual and presentation-level relationships in literacy, although their final evaluated records and recognition tasks are distinct.

Shared project decisions are reflected in the final terminology, the five-shape geometry scope, Android/Unity integration, age-specific evaluation and report-wide treatment of metrics. Contribution statements must not be inferred from proposal authorship alone for activities such as repository administration, final user-interface design or model training that were not individually documented. Before institutional submission, all members should confirm any broader task allocation and provide signatures or an approved contribution record. Maintaining this distinction protects both fairness in group credit and reproducibility of the integrated work.

5 INTEGRATED IMPLEMENTATION AND PROJECT EXECUTION

5.1 Common technologies and rationale

Unity and C# form the common interactive foundation because the project requires multiple input interfaces, scene-based learning, 3D geometry, mobile interaction and AR placement. Unity 2022.3 LTS provides a stable development baseline, while AR Foundation and Google ARCore support the Android camera-based physical-space features. The final system is not described as four independent apps: component-specific activities are housed in the same broader educational suite with common presentation and navigation concepts. This development foundation suits an educational prototype in which repeated interaction and visual feedback have as much importance as the underlying computational models.

Different components require different processing technologies. C1's verified validation is implemented through deterministic geometric constraints rather than AI. C2 uses C# UI and slicing controllers, a trace detector, AR placement and research logging scripts; the five final solid categories were Cube, Cuboid, Triangular Prism, Tetrahedron and Cylinder. C3 combines mobile-camera hand tracking, path preprocessing and CNN-based letter recognition. C4's sketch-model development includes Python, TensorFlow/Keras, OpenCV, FastAPI, NumPy and Pillow, while Unity provides the learning UI and AR reward. The presence of FastAPI in the development pipeline does not establish that a permanent production cloud service existed, nor does it establish the final runtime locality of all inference tasks.

Table 5.1 Common implementation technology inventory

Technology / tool	Documented role	Boundary of the assertion
Unity 2022.3 LTS, C#	Shared Android interactions and 3D presentation	C2 specifically used 2022.3.62f3 LTS
AR Foundation; ARCore XR Plugin	Surface detection, placement, AR inspection	Requires compatible Android hardware
Shape rules and C# controllers	Geometric validation and slicing behaviors	Validation is deterministic, not AI classification
Camera hand-landmark tracking	C3 index-fingertip path acquisition	Exact final hand-landmark runtime not fully specified
CNN recognition	C3 letters; C4 sketch objects	Models differ in task and evidence
Python; TensorFlow/Keras; OpenCV	C4 training and image preprocessing	Supported by C4 development description

Technology / tool	Documented role	Boundary of the assertion
FastAPI; NumPy; Pillow	C4 development and inference integration utilities	Not proof of a production server
Local research CSV workflow	C2 event and session export	Exact deployed version should be verified

5.2 Development evolution and design decisions

The March 2026 topic-assessment form introduced a research program concerned with 2D-to-3D geometry, ghost-image scaffolding, fingertip air-writing and sketch-driven literacy. The four individual proposals translated these ideas into separate component objectives, architectures, software tools and commercialization concepts. The subsequent development process refined the original broad aims into a smaller, testable system with explicit age groups and outcome measures. A critical revision moved the C2 corrective comparison from a 3D ghost-model setting to the invalid 2D geometric net itself; this better aligned the error cue with the learner’s actual construction and preserved the condition that an invalid net should not generate a validated 3D solid.

The Virtual Slicing scope also became more precise. Sphere was removed, while five solids were retained. Each selected solid could be explored through the three supported slicing directions and four fraction choices. The interaction required selecting a direction, choosing a fraction and tracing a guide before a displayed result. Additional AR and physics behaviors were explored during development, including placement after plane detection and separation of pieces. The final report records the stable learner-facing workflow and prepared result geometry while not assigning unsupported validation claims to a particular mesh or rigidbody implementation. This is a necessary boundary between design evolution and measured performance.

In the literacy pathway, C3 was kept distinct from touchscreen handwriting: its defining input is movement of a fingertip in free space captured by the phone camera. C4 terminology was refined away from formal phonics instruction toward multisensory vocabulary learning. This change reflects the concrete task sequence in which a child sketch is classified, a written word activity is completed, pronunciation-oriented acceptance is recorded and a 3D AR object is unlocked only after both tasks. Correct reward gating was important enough to be enforced in data cleaning and final

reporting. Both examples show that successful integration depended on precise definitions of what the system actually measures.

5.3 Deployment and research data handling

The Android learning artifact is built from a Unity project configured for the supported ARCore platform, with runtime permission and environmental conditions required for camera and AR use. A C2 development build was tested on a Xiaomi Redmi Note 9 Pro. The application should make the detected AR surface visible or otherwise recognizable to the child, support placement and ensure that slice interactions are not offered before the necessary model state is available. After task completion, repeat and navigation controls allow the child to return to the shape-selection or interaction stage. The research paper confirms component demonstrations and logged sessions but does not provide a reproducible deployment manifest, exact hardware performance matrix or published app-store release.

Data stewardship distinguishes anonymous or pseudonymous evaluation records from convenient application login. Original test datasets used identifiers such as ST0001–ST0025 to avoid placing children’s actual names in result tables. A later user-interface revision allowed a nonempty entered name, potentially stored in an existing StudentID column to preserve the schema, while an automatically generated unique Session ID linked events from the same session. Those two representations should not be confused: an entered name is not necessarily anonymous, and any exported identifiable text requires restricted handling or transformation before publication. The session identifier is a technical joining key, not evidence that a participant is a different person from another record.

The C2 data pipeline describes locally stored master CSV records, researcher-accessible copies in an Android Downloads/C2_ResearchData location and files such as C2_Events.csv and C2_Sessions.csv. Logged events include selected shapes, directions, fractions and trace attempts as well as success and completion states. The export should copy data rather than destroy master records. The final analysis dataset for the integrated research is a completed study record; it should not be replaced with new synthetic sessions merely to make experimental coverage more even. Where event

schema, file naming or Android storage policies have changed since the documented implementation, the corresponding code and APK should be checked before the report becomes a software manual.

5.4 Verification and provenance of implementation

The supplied materials serve different evidentiary roles. The Topic Assessment Form records approved objectives and original novelty claims. Four individual proposal reports provide early architecture, methods, user needs, budget and planned evaluation. The consolidated final IEEE-style paper documents the implemented system in the study and the results presented to readers. The GeoLex AR project context supplies later scope clarifications, documented logs and dataset conventions. The common-report template defines chapter responsibilities, while the SLIIT dissertation guideline governs typography, page structure, illustration numbering and references. These sources are complementary but not equal evidence of executed code or actual tests.

When code or data files are not included, this report avoids inventing interface signatures, algorithm thresholds, automatic test outcomes, source-control activity, an exact training split for C3, student-level learning-score distributions and model deployment claims beyond the documentation. Such specificity would reduce rather than improve reproducibility. A professional handover should include the four final source-code folders, build/version records, exact data schemas, anonymized original logs, relevant consent documentation and a release-device checklist. Appendix A provides an evidence register separating items available in this source bundle from items requiring final team verification.

6 INTEGRATED TESTING, EVALUATION AND VALIDATION

6.1 Research design and participating cohorts

The integrated evaluation adopted a Design Science Research (DSR) approach in which the developed mobile educational suite was treated as an artifact designed to address the identified instructional problem. Technical correctness and interaction/workflow behavior were examined independently from changes in learning-assessment means. The reported user-testing structure included 50 unique primary-school learners in two age-specific cohorts. The geometry cohort contained 25 learners aged 9–12 who participated in both C1 and C2 activities. A separate literacy cohort contained 25 learners aged 6–8 who participated in both C3 and C4 activities. This pairing explains why the study reports 25 observations per component without claiming four independent groups totaling 100 children.

Parental or guardian informed consent and child assent were obtained, institutional ethics procedures were followed, and participant data were anonymized for evaluation. The common paper shows real user-testing evidence, but identifiable photographs should appear in public-facing versions only under permissions covering that use. The evaluation collected technical decisions, session success states, recognition outputs, attempts, interaction timings and reported 10-mark pre- and post-tests. As complete participant-level paired educational scores were not retained in the final analysis file, the educational results are described using group means only. No unsupported paired significance test or effect size is presented.

Table 6.1 Cohort composition and evaluation activities

Cohort	Unique learners	Components	Primary activity outcomes
Geometry; ages 9–12	25	C1 and C2	Net validity; Virtual Slicing completion
Literacy; ages 6–8	25	C3 and C4	Air-written letters; sketch-to-vocabulary/AR completion
Total across cohorts	50	Four components	Four sets of 25 component records (not 100 learners)

6.2 Technical and workflow evaluation

C1 tested five cases per geometric solid, with four expected-valid cases and one expected-invalid case in each category. The 25 resulting validity decisions matched expected labels: 20 true valid, five true invalid and no mismatches in the recorded sample. All 20 expected-valid cases reached 3D creation and AR placement, while all five invalid ones were prevented from producing an incorrect validated 3D state. Mean task time was 12.65 seconds with standard deviation 6.81 seconds. The observed mean task times by solid were 6.77 seconds for Cube, 6.55 for Cuboid, 11.55 for Triangular Prism, 14.23 for Cylinder and 24.14 for Tetrahedron. With only five samples per shape, these are descriptive observations and not established comparative speed benchmarks.

C2 recorded 25 Virtual Slicing sessions, of which 20 completed successfully (80%). Mean task duration was 6.40 seconds, standard deviation 1.91 seconds, with observed range 3.02–9.78 seconds. Vertical sessions were successful in all five recorded trials; Horizontal succeeded in 11 of 14; Angled succeeded in four of six. For fractions, successful sessions were four of five for $\frac{1}{2}$, three of four for $\frac{1}{3}$, six of seven for $\frac{1}{4}$ and seven of nine for $\frac{3}{4}$. Since these categories have unequal and small denominators, they do not support strong claims about one slice configuration being generally easier or more reliable. The results directly establish performance of the logged slicing workflow, not a separate accuracy score for 2D Net Correction.

Table 6.2 Component-level technical and workflow results

Component	n	Measured primary outcome	Recorded value
C1	25	Expected vs. actual net validity	25/25 agreement (100%)
C2	25	Successful Virtual Slicing sessions	20/25 (80%)
C3	25	Correct air-written letter predictions	19/25 (76%)
C4	25	Full vocabulary-to-AR workflow completion	17/25 (68%)

The four percentages in Table 6.2 do not describe four measurements of the same underlying property. C1 is a decision-agreement result for a small, balanced-by-design set of nets; C2 is interaction completion; C3 is correct character prediction including

no-prediction trials as unsuccessful; and C4 is task progression through multiple prerequisite states. Comparing those percentages as if they were ranked AI classifier accuracies would be analytically invalid. Figure 6.2 is therefore titled to emphasize distinct metrics, and its role is to communicate recorded successful outcomes rather than to rank algorithms.

Table 6.3 Direction and fraction completion in C2

Group	Recorded setting	Successful / total
Direction	Vertical	5/5
Direction	Horizontal	11/14
Direction	Angled	4/6
Fraction	1/2	4/5
Fraction	1/3	3/4
Fraction	1/4	6/7
Fraction	3/4	7/9

C3 produced 25 recorded air-writing trials from young learners. Nineteen predictions matched the target, three produced an incorrect letter and three returned no letter. The evaluated target set included A, B, C, D, F, I, J, L, N, O, P, Q, R, S, X, Y and Z. Examples of incorrect labels were A recognized as H, D as O and L as J; one trial each for N, P and R returned no prediction. Mean recognition confidence was 0.784 (standard deviation 0.314), mean drawing duration was 5.05 seconds (SD 2.07) and mean complete session duration was 8.35 seconds (SD 2.65). The results show that camera-based free-space input was possible in the study, while recognizing children's changing motion remains challenging.

C4 recorded 25 child-user sketch sessions across the five supported object classes. Spelling completion occurred in 20 sessions (80%), pronunciation-oriented acceptance in 21 (84%) and successful completion of both prerequisites in 17 (68%). Exactly the same 17 sessions unlocked and placed the AR object. Mean prediction confidence was 0.832 (SD 0.099), mean spelling attempts 1.40, mean pronunciation-oriented attempts 1.28 and mean total session duration 56.52 seconds (SD 13.22). These figures describe the recorded educational progression. Confidence is an algorithm-provided value and cannot be substituted for observed sketch accuracy when the ground-truth intended object was not recorded independently for every child's drawing.

Table 6.4 C3 and C4 outcome detail

Measure	Recorded outcome	Interpretation
C3 correctly predicted	19 / 25	76% recognition on evaluated sample
C3 incorrect prediction	3 / 25	Predicted wrong letter
C3 no prediction	3 / 25	No letter returned; counted unsuccessful
C4 spelling completed	20 / 25	Task-stage completion
C4 pronunciation accepted	21 / 25	Task-stage acceptance, not speech-model accuracy
C4 both prerequisites and AR	17 / 25	68% full workflow completion

6.3 Age-specific learning assessment

Reported group mean pre-test and post-test scores were expressed on a ten-mark scale for each component. Among geometry learners aged 9–12, C1 and C2 each increased from a mean of 5/10 to a mean of 9/10. Among literacy learners aged 6–8, C3 increased from 5/10 to 9/10 and C4 increased from 4/10 to 8/10. Each component therefore showed an absolute change of four marks, equivalent to 40 percentage points of the ten-mark scale. This does not mean a 40% relative gain in each component; for example, moving from 4/10 to 8/10 is a doubling of the score but a 40-percentage-point change on that scale.

The study reported these means as descriptive learning evidence. They are consistent with improved performance on the assessments administered after learners interacted with the relevant activities. They do not by themselves establish that the system caused the improvement because there was no randomized comparison condition, the duration was short and complete participant-level paired scores were not available for inferential analysis. The common report therefore does not claim a statistically significant treatment effect, longitudinal retention or superiority over all traditional methods. The age split should also be preserved when interpreting outcomes, since the geometry and literacy assessments measure different educational constructs.

Table 6.5 Pre-/post-test mean scores

Component	Age	Mean pre-test / 10	Mean post-test / 10	Change
C1 Geometry Validation	9–12	5	9	+4

Component	Age	Mean pre-test / 10	Mean post-test / 10	Change
C2 Net Correction & Slicing	9–12	5	9	+4
C3 Air-Writing	6–8	5	9	+4
C4 Vocabulary	6–8	4	8	+4




Figure 6.1 Recorded mean pre-test and post-test scores; descriptive cohort means only.




Figure 6.2 Recorded successful outcomes for four distinct component measures (not comparable algorithm accuracies).

6.4 Interpretation by research question

RQ1 is answered by agreement with known expected labels and correct gating of 3D/AR progression in the 25 recorded C1 trials. RQ2 is answered most directly through the 20/25 completion outcome, timing and configuration-specific summaries for C2. Although the correction workflow is described and integrated with validation, the available log does not isolate correction-specific learning or task effectiveness. This is a deliberate limitation, not an indication that the implemented guidance has no value. For RQ3, the 19/25 result quantifies letter recognition in the sampled trial set

and reports the contribution of incorrect or absent predictions, with timing providing one indication of practical responsiveness.

RQ4 is answered at the level of completing the C4 sequence: spelling, pronunciation-oriented acceptance and eventual AR placement. The 17/25 result shows that the two prerequisites were enforced in the reported evaluation, but it is neither independent sketch classification accuracy nor stand-alone pronunciation recognition accuracy. RQ5 is answered by the observed four-mark changes in each reported mean, appropriately restricted to descriptive interpretation. RQ6 is addressed provisionally by observed completion, attempts and timings; there is no verified standardized child-usability questionnaire score in the final consolidated report, so usability should not be expressed as a validated scale result.

6.5 Objective achievement and integration evidence

The component results suggest that the study successfully produced and evaluated an integrated software artifact with differentiated learning functions. The availability of 3D outputs only after valid net decisions establishes a particularly important cross-component educational constraint. The Net Correction and Virtual Slicing component provided the designed guidance stage and a successfully tested slicing interaction; its corrective benefit remains a separate evaluation opportunity. The camera-based Air-Writing component returned letter predictions from genuine free-space movement, while C4 linked real child-generated drawings to a full conditional vocabulary and AR sequence. Together they instantiate the proposed architectural contract with different recognition and validation processes.

Table 6.6 Group objectives and associated evidence

Objective	Evidence available	Judgment within the study
O1: Validate 2D nets	25/25 expected-label agreement	Supported for tested cases
O2: Correction and slicing	Correction design; 20/25 slice completions	Slicing measured; correction not independently measured
O3: Camera-based air-writing	19/25 correct sampled trials	Supported for sampled uppercase targets
O4: Sketch-to-vocabulary/AR	17/25 full sessions; task prerequisites	Supported at workflow level
O5: Integrated mobile environment	Common Unity experience and architecture	Documented prototype integration

Objective	Evidence available	Judgment within the study
O6: Technical and educational pilot evaluation	50 unique learners; descriptive pre/post means	Pilot evidence; no causal estimate

A balanced interpretation is that the artifact is feasible and that the child-user pilot captured meaningful task execution and descriptive changes. It is not demonstrated that every proposed feature worked for every possible object, that one model achieved the same accuracy on a different school population or that the observed learning changes would persist. The contribution of the common report is to preserve both sides: the educational ambition of one mobile suite and the distinct empirical limits of each technical subsystem.

6.6 Threats to validity and evidence boundaries

Internal validity is restricted because the pilot lacks a randomized non-AR comparison or other condition capable of isolating the application’s effect from practice, familiarity, teaching support or other influences. The aggregate pre/post mean records are insufficient for computing verified standard deviations of score changes, effect sizes or p-values. Construct validity requires caution because component success metrics measure different phenomena, and because completion of tasks does not directly establish conceptual mastery, recognized sketch accuracy or correct pronunciation. The period of interaction was relatively short, so durable learning or transfer to another activity was not established.

External validity is limited by the sample of 50 unique learners divided between two age bands, the 25 trial/session records for each component and unequal coverage of shapes and fractions in C2. C3 tested only 17 uppercase letters rather than all uppercase and lowercase targets. C4 child sketches lack independent ground-truth object intentions throughout the final session dataset, preventing a complete child-domain confusion matrix. Technical generalizability is limited by incomplete multi-device benchmarking, undocumented model reproduction settings and the absence of an independent Net Correction evaluation. Future studies must address these gaps without modifying genuine pilot records.

7 COMMERCIALISATION, SUSTAINABILITY AND RESPONSIBLE INNOVATION

7.1 Educational value and beneficiaries

GeoLex AR is positioned as an interactive digital learning supplement for primary-school students rather than a replacement for teachers, textbooks or a full curriculum. Its educational value is the opportunity for children to create learning input instead of only consuming predefined content. A geometry net can be rejected before incorrect 3D reinforcement, then revised with visual guidance; a child can explore sectional views of a solid; a letter can be written in free space; and an object drawing can lead to spelling, spoken practice and an AR reward. Each interaction uses the same principle of recognition or validation before feedback, giving the product a coherent instructional identity.

Potential direct users are learners and teachers, with parents/guardians and educational institutions playing supervisory or purchasing roles. In the Sri Lankan context, an Android-based design offers a practical starting point for examining lower-cost AR experiences than dedicated head-mounted equipment. However, the appropriate device requirements and support burden matter. ARCore compatibility, rear-camera performance, screen size, charging, connectivity for any external model services and classroom management requirements all affect the real cost of adoption. A future pilot would need to account for both student access and educator training to show that the system works fairly beyond the research environment.

7.2 Alternatives and differentiation

Non-digital alternatives include physical geometry nets, cut-out paper models, teacher-led demonstrations, worksheets and vocabulary cards. Their strengths include low initial technology requirements, tangible manipulation and familiarity; the proposed application should complement those strengths. Digital alternatives include general geometry visualization apps, drawing-recognition tools, tracing applications and AR flashcards. The distinctive proposition in GeoLex AR is not that every technology is unprecedented, but that four age-specific learning activities use learner-generated

input as the trigger for domain-appropriate computational interpretation and feedback. Recent AR research provides context for geometry support [3]–[7] and vocabulary interaction [16]–[18], while the paper’s contribution is the integrated design.

The strongest defensible differentiation is also the least dependent on exaggerated claims: C1 checks a child’s net mathematically before 3D display, C2 supports correction and selected fractional slicing, C3 performs actual camera-based air-writing rather than touchscreen-only tracing, and C4 gates an AR reward through two vocabulary tasks. The study does not justify marketing phrases such as clinically proven learning gain, universally superior instruction, perfect classifier accuracy or patent-protected uniqueness. A defensible commercial launch would require product usability research, real-world content validation, language and curriculum review, maintainability planning and legal advice on intellectual property rather than relying on proposal-stage promotional statements.

7.3 Proposed adoption and business pathways

Possible routes to use include supervised classroom demonstrations, teacher-led station activities, after-school learning, parent-assisted home practice and partnerships with educational technology organizations. Commercial pathways could include institutional licensing, teacher-tool bundles, school pilots or a freemium distribution model with optional content expansions. The individual proposals provide ideas for institutional pricing and market segments, but these are scenario concepts—not validated willingness-to-pay estimates or contracted revenue. Thus no financial forecast is presented as a research result. A later feasibility study should test procurement constraints, cost of Android device access, AR support requirements, teacher onboarding time and the resources needed to update recognition models and 3D assets.

A staged deployment would begin with a controlled pilot using a defined set of supported devices and supervised sessions. Once software stability and child comprehension have been assessed, a limited institutional trial could examine frequency of use, teacher confidence, content alignment and operational support. Only after those requirements are satisfied would larger distribution, ongoing customer

support and commercial pricing be defensible. Expansion should prioritize quality and accessibility within the documented five-shape/five-object scope before adding many new categories. A feature that can be demonstrated at an exhibition is not yet proof of a scalable maintained educational service.

7.4 Ethics, privacy, licensing and access

Responsible commercialization must begin with the protection of child participants and users. The completed research documented guardian consent, child assent and anonymized evaluation data. A product release would require an explicit privacy policy, age-appropriate permissions and controls that minimize personal information. If an application login accepts a child's name, this is different from storing a pseudonymous study ID; the system should provide an alias or nonidentifying approach for actual research export and prevent names from being displayed publicly. Voice and camera capture also require clear communication about whether raw recordings are stored, transmitted, processed transiently or omitted. The final source code and deployment configuration must be inspected before any claim of universal on-device processing is made.

The use of public research data, such as Quick, Draw!-derived training images, requires checking source licenses, permitted uses and any attribution obligations before commercial model distribution. Team ownership, university agreements and co-authorship rights should be clarified prior to offering the application commercially. App assets and user-interface artwork must be appropriately licensed, and identifiable photographs of participating children must never be reused in promotional material without permission covering that specific purpose. Safety also includes supervising movement around real objects during AR camera use and designing prompts that do not pressure children to share unnecessary information.

Table 7.1 Adoption risks and proposed mitigations

Potential risk	Practical consequence	Proposed response
Limited ARCore/device compatibility	Unequal access and poor experience	Publish supported device matrix; provide guided demo devices
Incomplete language and geometry coverage	Overstated curriculum alignment	Scope statements and teacher-reviewed expansion
Children's data/voice privacy	Inappropriate disclosure or trust loss	Minimize data; consent, aliases, retention and access controls

Potential risk	Practical consequence	Proposed response
Classification failures	Frustration and incorrect reinforcement	Retry flows, confidence handling and teacher assistance
Unvalidated pricing	Unsustainable commercial assumptions	Institutional interviews and costed pilots
New content licensing	Legal and distribution problems	Asset/model license audit and ownership agreement

7.5 Sustainability and SDG 4 contribution

A credible SDG 4 contribution would come from supporting inclusive, equitable and high-quality primary instruction, but the pathway from a technical pilot to such outcomes must be evaluated. The Android design could reduce dependence on specialized hardware, and interactive feedback could provide additional practice opportunities. Conversely, device ownership, digital literacy, visual accessibility and rural connectivity conditions may create new inequities. Sustainable expansion should therefore include offline-capable tasks where technically verified, assistive teacher documentation, compatibility testing and mechanisms for collecting feedback without exposing individual children.

The next stage of social evaluation should measure not only immediate completion but also classroom fit, teacher support time, access across device types, longer-term concept retention and how the digital activity works alongside conventional pedagogy. Commercial sustainability requires the same restraint. The team has documented a promising educational artifact and pilot outcomes; measurable equitable impact will require future studies designed specifically to test accessibility, continuity and learning outcomes across a broader school population.

8 CONCLUSIONS AND FUTURE DIRECTIONS

8.1 Overall conclusion

This research designed and piloted GeoLex AR, an Android-based educational suite supporting primary-school geometry and English literacy through four distinct processing pipelines under one educational interaction framework. The system turns learner-created nets, slice traces, camera-observed air-writing and object sketches into computable inputs; each component uses the relevant geometric or AI method to decide what feedback or reinforcement should follow. The most important design feature is not the presence of AR by itself, but the placement of interpretation between the child's action and the system's instructional response. This helps distinguish deliberate educational interactivity from passive 3D visualization.

The final evaluation comprised 50 unique learners, with 25 geometry learners aged 9–12 shared between C1 and C2 and 25 literacy learners aged 6–8 shared between C3 and C4. C1 matched expected validity results in 25/25 recorded cases, C2 successfully completed 20/25 Virtual Slicing sessions, C3 recognized 19/25 sampled air-written letters and C4 completed 17/25 full vocabulary-to-AR sequences. Group mean assessment scores increased by four marks on a ten-mark scale for every component. These are meaningful descriptive observations within the pilot and substantiate feasibility of the interaction design, but they do not prove causal learning effectiveness or a single common algorithmic accuracy.

8.2 Integrated research contribution

The integrated contribution is an organized, learner-generated input framework that supports two curriculum domains without forcing them into one recognition method. The geometry components demonstrate a direct dependency between deterministic correctness and guidance/3D exploration: C1 stops invalid nets before verified 3D output and C2 offers a correction path and guided slicing. The literacy components demonstrate two different routes to interpretation: temporal fingertip trajectories for letter recognition and visual object sketches for word-based tasks. The shared presentation layer allows consistent activity progression, retry and AR reinforcement

while the internal reasoning, input types and evaluation measures remain component-specific.

The study also contributes a transparent example of outcome reporting in a multi-component educational project. C1 agreement, C2 interaction success, C3 letter accuracy and C4 full workflow completion are deliberately reported as different constructs. This distinction preserves research integrity and helps future teams avoid the common error of displaying four percentages in one table as though all were model accuracies. Likewise, descriptive pre-/post-test gains are separated from technical correctness and interpreted in light of the actual sample and study design. These reporting choices are part of the project's value as a reproducible educational technology case study.

8.3 Priorities for further evaluation and deployment

The immediate technical priorities are verifying a complete release manifest, documenting current AI inference runtimes, maintaining robust session and event logging and testing AR functionality on a wider range of compatible Android devices. C1 should be challenged using a larger and more diverse library of valid and invalid nets. C2 should receive a separate user-centered evaluation of fault-specific 2D correction, including how quickly children identify errors and whether guidance helps them revise their work; its Virtual Slicing study should balance trials across shapes, directions and fractions. Any physics or gravity behavior should be documented according to the verified final application rather than assumed from a development plan.

C3 needs broader uppercase/lowercase benchmarking across learners and classroom conditions with documented training details. C4 needs independent intended-object labels for child sketches to permit class-level accuracy evaluation. Separate pronunciation-validation ground truth would be needed before making model-level pronunciation claims. Future work should retain paired scores and include delayed tests, validated child-usability instruments and suitable comparison groups.

REFERENCES

- [1] H.-Y. Chang, T. Binali, J.-C. Liang, G.-L. Chiou, K.-H. Cheng, S. W.-Y. Lee, and C.-C. Tsai, "Ten years of augmented reality in education: A meta-analysis of (quasi-) experimental studies to investigate the impact," *Computers & Education*, vol. 191, Art. no. 104641, 2022, doi: 10.1016/j.compedu.2022.104641.
- [2] E. L.-C. Law and M. Heintz, "Augmented reality applications for K-12 education: A systematic review from the usability and user experience perspective," *International Journal of Child-Computer Interaction*, vol. 30, Art. no. 100321, 2021, doi: 10.1016/j.ijcci.2021.100321.
- [3] M. B. Nadzeri, M. Musa, C. M. Chew, and I. M. Ismail, "The effects of augmented reality geometry learning applications on spatial visualization ability for primary school pupils," *International Journal of Interactive Mobile Technologies*, vol. 18, no. 16, pp. 104–118, 2024, doi: 10.3991/ijim.v18i16.47079.
- [4] W. Tarng, J.-K. Huang, and K.-L. Ou, "Improving elementary students' geometric understanding through augmented reality and its performance evaluation," *Systems*, vol. 12, no. 11, Art. no. 493, 2024, doi: 10.3390/systems12110493.
- [5] J. Wu, H. Jiang, L. Long, and X. Zhang, "Effects of AR mathematical picture books on primary school students' geometric thinking, cognitive load and flow experience," *Education and Information Technologies*, vol. 29, pp. 24627–24652, 2024, doi: 10.1007/s10639-024-12768-y.
- [6] A. Uriarte-Portillo, R. Zatarain-Cabada, M. L. Barrón-Estrada, M. B. Ibáñez, and L.-M. González-Barrón, "Intelligent augmented reality for learning geometry," *Information*, vol. 14, no. 4, Art. no. 245, 2023, doi: 10.3390/info14040245.
- [7] C. Walkington, M. J. Nathan, J. Washington, J. Hunnicutt, T. Darwin, L. Daugherty, and K. Schenck, "Comparing learning geometry using a tablet to head-mounted display augmented reality: How and when dimensionality matters," *Education and Information Technologies*, vol. 30, pp. 5397–5426, 2025, doi: 10.1007/s10639-024-13008-z.
- [8] S.-H. Zhang, Y.-C. Guo, and Q.-W. Gu, "Sketch2Model: View-aware 3D modeling from single free-hand sketches," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 6012–6021, doi: 10.1109/CVPR46437.2021.00595.
- [9] L. Zhang, "Hand-drawn sketch recognition with a double-channel convolutional neural network," *EURASIP Journal on Advances in Signal Processing*, vol. 2021, Art. no. 73, 2021, doi: 10.1186/s13634-021-00752-4.
- [10] A. K. Bhunia, P. N. Chowdhury, Y. Yang, T. M. Hospedales, T. Xiang, and Y.-Z. Song, "Vectorization and rasterization: Self-supervised learning for sketch and handwriting," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 5672–5681.
- [11] S. Mukherjee, S. A. Ahmed, D. P. Dogra, S. Kar, and P. P. Roy, "Fingertip detection and tracking for recognition of air-writing in videos," *Expert Systems with Applications*, vol. 136, pp. 217–229, 2019, doi: 10.1016/j.eswa.2019.06.034.
- [12] F. Zhang, V. Bazarevsky, A. Vakunov, A. Tkachenka, G. Sung, C.-L. Chang, and M. Grundmann, "MediaPipe Hands: On-device real-time hand tracking," in *CV4ARVR*, 2020.
- [13] M. S. Alam, K.-C. Kwon, M. A. Alam, M. Y. Abbass, S. M. Imtiaz, and N. Kim, "Trajectory-based air-writing recognition using deep neural network and depth sensor," *Sensors*, vol. 20, no. 2, Art. no. 376, 2020, doi: 10.3390/s20020376.
- [14] T. Watanabe, M. Maniruzzaman, M. A. M. Hasan, H.-S. Lee, S.-W. Jang, and J. Shin, "2D camera-based air-writing recognition using hand pose estimation and hybrid deep learning model," *Electronics*, vol. 12, no. 4, Art. no. 995, 2023, doi: 10.3390/electronics12040995.

- [15] A. Qedear, A. AlMatrafy, A. Al-Sowat, A. Saigh, and A. Alayed, "Real-time air-writing recognition for Arabic letters using deep learning," *Sensors*, vol. 24, no. 18, Art. no. 6098, 2024, doi: 10.3390/s24186098.
- [16] E. Korosidou, "The effects of augmented reality on very young learners' motivation and learning of the alphabet and vocabulary," *Digital*, vol. 4, no. 1, pp. 195–214, 2024, doi: 10.3390/digital4010010.
- [17] J. Belda-Medina and V. Marrahi-Gomez, "The impact of augmented reality (AR) on vocabulary acquisition and student motivation," *Electronics*, vol. 12, no. 3, Art. no. 749, 2023, doi: 10.3390/electronics12030749.
- [18] H. Lin and W. Wei, "A systematic review on vocabulary learning in AR and VR gamification context," *Computers & Education: X Reality*, vol. 4, Art. no. 100057, 2024, doi: 10.1016/j.cexr.2024.100057.
- [19] H. Weigand, P. Johannesson, and B. Andersson, "An artifact ontology for design science research," *Data & Knowledge Engineering*, vol. 133, Art. no. 101878, 2021, doi: 10.1016/j.datak.2021.101878.
- [20] S. G. Sutton, V. Arnold, P. Collier, and S. A. Leech, "Leveraging the synergies between design science and behavioral science research methods," *International Journal of Accounting Information Systems*, vol. 43, Art. no. 100536, 2021, doi: 10.1016/j.accinf.2021.100536.

APPENDICES

Appendix A: Research evidence register and source hierarchy

This appendix identifies the materials used for the common report and the verification actions required before institutional submission. The evidence register is not a claim that every research source file, participant consent record or application build is included in the attached document bundle. It distinguishes approved proposal intentions, final reported results, and operational artifacts that should remain available for examination.

Table A.1 Research evidence register

Research material	Function in this report	Verification note
Topic Assessment Form (R26-IM-005)	Original aim and individual assignments	Proposal-era scope; later revisions take precedence
Four March 2026 component proposals	Problems, methods, early architectures	Do not treat every proposed technology as final
Final ICAC / IEEE-style group paper	Core implemented narrative and final numeric findings	Primary source for recorded evaluation values
GeoLex AR master context	Final terminology, detailed recorded outcomes	Cross-check final APK and scripts for exact runtime claims
SLIIT documentation guidelines	Margins, type, chapter and reference format	Institutional formatting requirement
Common integrated report template	Required eight-chapter integrated structure	This report follows the supplied organization
Final APK, source code and original logs	Reproducibility and independent verification	Not supplied in this document bundle
Photographic publication permissions	Authorization for identifiable child photos	Verify scope of publication consent

Appendix B: Integrated acceptance and reproducibility checklist

The checklist below is a practical handover instrument for final verification. Items are phrased as checks, not as results from a newly conducted system audit. They can be completed against the final Unity release, anonymized test files and team records without altering the documented pilot results.

Table B.1 Integration acceptance checklist

Check	Expected behavior or evidence
C1 incorrect geometric net	No accepted 3D/AR generation; correction pathway accessible
C1 valid geometric net	Correct solid shown and AR placement available

Check	Expected behavior or evidence
C2 supported shapes	Cube, Cuboid, Triangular Prism, Tetrahedron, Cylinder present
C2 slicing options	Vertical, Horizontal, Angled; 1/2, 1/3, 1/4, 3/4 present
C2 incorrect trace	Retry and failure/attempt event available
C2 repeat selections	All meaningful direction/fraction choices remain logged
Session identity	Different runs receive unique Session IDs, even for same entered name
CSV preservation	Exports copy records without deleting underlying master data
C3 letter interaction	Letter formed by camera-based finger movement, not touch tracing
C3 unsuccessful recognition	Wrong/no-prediction cases retained as failures
C4 one prerequisite complete	AR reward remains locked
C4 both prerequisites complete	Correct AR reward unlocked and can be placed
Privacy and approvals	Publication consent and anonymization checked for all external outputs

Appendix C: Detailed result definitions and denominators

For C1, the denominator is 25 recorded geometric net validation trials, with 20 expected-valid and five expected-invalid. The numerator counts agreement with these expected labels, including both positive and negative cases. For C2, the denominator is 25 Virtual Slicing sessions; the numerator counts 20 completed slicing workflows. The shape-category counts recorded for slicing were Cube 3, Cuboid 8, Triangular Prism 9, Tetrahedron 3 and Cylinder 2, and therefore are too uneven for strong category comparisons. For C3, 19 correct recognitions out of 25 include three wrong letters and three no-result trials in the denominator. For C4, 17 complete sequences out of 25 require both spelling completion and pronunciation-oriented acceptance before AR access.

The C3 target letters sampled in the child evaluation were A, B, C, D, F, I, J, L, N, O, P, Q, R, S, X, Y and Z. The C4 supported object classes were Apple, Cat, Fish, Sun and Tree. Both lists represent evaluation and model scope statements; they should not be treated as evidence of exhaustive generalization beyond the specified test conditions. The reported means for learning assessments were computed on a ten-mark scale and should be interpreted as descriptive group aggregates. No additional

synthetic student rows, imagined raw-score distributions or unreported experimental results were introduced into this report.

Appendix D: Completion and publication review

Before the research group signs and submits this document, the members should check the official degree title and full registered names against their student records; confirm the supervisor's official title and spelling; verify the submission month; inspect all figures and Word page references after any editing; and ensure that the final technical statements still match the source repository and application release. The paper-based examples and architecture drawings are explanatory representations and must not be presented as screenshots of unverified software functionality. The present document intentionally does not include identifiable child photographs because permission for republication outside the final research paper has not been separately confirmed in the supplied bundle.

Every final reference should be checked against the source publication and citation order, and the manuscript should be proofread by the team and supervisor. The four percentages in the technical summary must remain explicitly labeled with their different denominators and meanings. Any new testing added after the finalized results should be introduced as a new, documented experiment rather than quietly changing the historical findings. This review is especially important if future work modifies the login-to-dataset mapping, the C3 inference system, the C4 pronunciation module or the C2 slicing physics.

Appendix E: Application interface evidence from the final group paper

The image below reproduces the user-interface montage supplied in the final IEEE-style group paper (Figure 5 therein). It documents representative application visual design, but it should not be read as exhaustive proof that every flow in this report appears in a single screenshot. The illustration is reproduced from the research team's own submitted paper rather than generated as a fictional software mockup. For a higher-resolution archival thesis, the authors should replace it with direct screenshots taken from the final tested APK.




Fig.5. User Interface Designs

Figure E.1 Representative application interface montage reproduced from the final research paper (original Figure 5).

The screen illustrations provide contextual design evidence. The report's measured results, in contrast, come from the final technical and session data and are described in Chapter 6. The team should independently confirm that publication permissions permit reuse of any additional user-testing photographs before including images of minors in the final dissertation.