AI Powered Digital Examination Ledger (ADEL): A Hardware Agnostic OMR Framework for Unified Assessment Analytics

Vishal Tambi, Isha Yadav, Kankati Aneesh, Rama Vamsi, Bhupesh Kumar Singh

Department of Computer Science & Engineering
Amity University Rajasthan, Jaipur, India
Corresponding Author: tambivishal3@gmail.com

ABSTRACT

Given the capability of the “Software Defined OMR” engine to work with black and white photocopied forms, the proposed approach allows for eliminating the need for special-purpose scanning equipment as it allows answering sheets to be analysed using regular smartphone photos. The technology makes use of micro services architecture, a Convolutional Recurrent Neural Network (CRNN) for transcribing handwritten student IDs, circumventing optical constraints of conventional OMR technologies and Four Point Perspective Transforms to deal with issues of illumination and geometric distortions. An additional four-layered Adaptive Intelligence Layer improves the post-OCR identity matching precision from 71.3% (raw CRNN) to 99.1% in the case of enrolment numbers by employing an adaptive corrections log, session memory, and fuzzy database matching and human-in-the-loop approach. According to the performance indicators, ADEL offers an efficient and hardware-independent system for education-related institutions looking to combine their offline and online performance data with 99.7% bubble detection accuracy, 99.1% enrolment number recognition accuracy after adaptive identity resolution and 0.7 seconds per form processing speed.

Keywords: Optical Mark Recognition, Computer Vision, Deep Learning, Educational Technology, CRNN, Assessment Analytics, Fuzzy Matching, Adaptive Intelligence.

I. INTRODUCTION

The techniques of educational assessment have changed greatly as more information and communication technologies (ICT) have been used in the field of education. The ICT allows the organizations to handle the analysis of great amount of performance data of students more effectively than it was possible to do using the old methods. The necessity of systems that will be able to work with great amounts of the data related to the performance of students with the high latency efficiency has appeared because of the tendency of modern day organizations to use data-based approach in order to control the process of learning [1].

Owing to its strong security and resilience to network infrastructural disruptions, Paper Based Testing (PBT) remains the undisputed king of high-stakes testing in spite of the commonality of Computer Based Testing (CBT). It follows that the academic environments are operating on a two-mode system, where summative tests are conducted offline to maintain their sanctity while formative tests are conducted online for real time feedback. This is leading to a 'data silo' phenomenon, where the complete academic experience of the student is stored in incompatible storage devices [2].

Optical Mark Recognition (OMR) remains one of the most widely used technologies for digitizing responses in paper-based examinations. However, there are considerable operational bottlenecks associated with the existing OMR infrastructure. Firstly, there are substantial CAPEX and maintenance overheads associated with the use of proprietary hardware scanners. Secondly, the technology does not allow the use of traditional black and white reprography because it requires special forms to be printed on high GSM paper with “drop out” colours (red or green ink). Thirdly, the use of handwritten text does not allow optical character recognition (OCR) in the existing OMR technology, thereby requiring manual entry of student identification for grading purposes, which has a considerable error rate [3], [4].

To address these limitations, this work proposes the AI-Powered Digital Examination Ledger (ADEL), a lightweight framework that replaces dedicated OMR scanners with smartphone-based image capture, a software-defined OMR solution that uses any commodity smartphone camera as the scanning interface in place of specialized scanning units. The suggested system combines micro services architecture to directly synchronize offline exam results with online LMS databases and uses a Convolutional Recurrent Neural Network (CRNN) to automate the recognition of handwritten student identities [5]. A novel four-tiered Adaptive Intelligence Layer further post-processes OCR predictions against the institutional enrolment database, achieving near-zero manual correction overhead after only three examination sessions. This approach guarantees a consistent analytical view of

ISBN: 978-81-985770-0-9

student performance while doing away with continuing costs of proprietary consumables.

II. LITERATURE REVIEW

Automated assessment has evolved from fixed hardware systems to adaptable AI-powered solutions. In order to identify the exact technological gaps that ADEL fills, this section examines this chronological progression.

A. Early Hardware-Centric Systems (2018–2020)

Smith et al. (2018) [4] performed an economic analysis of the traditional OMR technology, concluding that hardware scanners have a high level of accuracy (>99.9%) using carbon reflection on "drop out" ink paper. However, the system is economically unviable for low-resource institutions due to the high Recurring Capital Expenditure (RCapex) of proprietary scanners and 100 GSM paper. The ADEL system overcomes this problem by using an Adaptive Thresholding Algorithm on standard 70 GSM photocopied paper, thus bringing down the operational cost to almost zero.

M. Guha & Das (2020) [6] suggested a "low cost" solution based on flatbed scanners and simple image subtraction algorithms. The process was the same, requiring manual alignment of paper on a flatbed scanner, thus not allowing quick grading of papers in a classroom environment. The ADEL system supports Mobile First Ingestion, using a smartphone camera and a "Factory Mode" loop to scan 50+ sheets in quick succession without manual alignment.

B. The Webcam & Computer Vision Era (2021)

According to Nguyen & Lee (2021) [7], an OMR framework was developed based on the Hough transform for detecting circles through a mobile phone. This is one of the earliest attempts to shift from using scanners. The problem with the framework was that it was very sensitive to Perspective Distortion. In case of a slight deviation of up to 10 degrees by the camera, the grid alignment failed and accuracy dropped to 92.5%. In ADEL, there is an accurate Four Point Homograph Transform that identifies the four corner markers of the answer sheet and mathematically flattens the image to give 99% accuracy.

C. Challenges in Unconstrained Environments (2022)

Extending this to real-world applications, Chopra and Ghadge [8] use template matching to analyse classroom lighting and bubble detection. Due to a lack of a pre-processing pipeline that can normalize brightness, the system tends to fail under varying illumination conditions—the shadow due to the underlying hand or even phone lights may give false positives, whereby the system perceives shadows as dark bubbles.

D. The Deep Learning Gap (2017–2023)

Shi et al. (2017) [5] revolutionized the field of text recognition with the CRNN (Convolutional Recurrent Neural Network) architecture, which proves its superiority compared to traditional OCR when it comes to reading sequential text. Nevertheless, following this breakthrough, no existing OMR research applied CRNN to Student IDs. Most systems still required students to bubble their roll numbers, which is error-prone and time-consuming. ADEL is the first platform to hybridize OMR with Handwritten Digit Recognition, utilizing a fine-tuned CRNN model to read handwritten Student IDs directly, removing the need for bubbling roll numbers.

E. The Integration Gap (2023)

R. Kumar & Singh (2023) [9] investigated "Hybrid Learning Architectures" and identified the "Integrity–Accessibility Paradox." They found that offline exam data remains "siloed" in Excel sheets and never reaches the central Learning Management System (LMS), with no automated "handshake" or API to push offline grades to online portals. ADEL utilizes a Micro services Architecture where the Python grading engine is wrapped in a Flask API that instantly pushes JSON results to the Node.js backend, unifying the offline scores with the student's online profile in real time.

III. METHODOLOGY

The micro services architecture of the proposed system breaks down the CPU-intensive computer vision engine (Python) into a separate service from the I/O-intensive web server (Node.js). The five consecutive modules form the central processing pipeline: image pre-processing, geometric normalization, hybrid recognition, adaptive intelligence, and output generation.

A. Image Pre-Processing and Normalization

The input is provided by the RGB image I_{raw} obtained via the smartphone client. The pre-processing chain is performed prior to feature extraction and its purpose is noise removal related to the environmental noise such as various lighting conditions.

1) ***Grayscale Image Formation***: The image is transformed into an intensity single channel image I_{gray} through the weighted sum technique:

$$I_{gray} = 0.299R + 0.587G + 0.114B \tag{1}$$

ISBN: 978-81-985770-0-9

2) ***Noise Filtering***: I_{gray} undergoes convolution with the Gaussian kernel G(x, y) of 5×5. The goal is to filter out the high frequency noise (grain).

3) ***Adaptive Binarization***: The adaptive banalization algorithm proposed by Sauvola is used [10]. It works better than global banalization under non-uniform illumination. The local threshold value is determined based on the local average and local variance as follows:

$$T_{(x,y)} = m\left[1 + k\left(\frac{s}{R} - 1\right)\right] \quad (2)$$

Where R = 128, k = 0.5, thus separating the answer sheet from the background.

B. Geometric Normalization (Software-Defined Scanner)

The system uses a homography-based transformation to correct perspective distortion caused by oblique camera angles.

1) *Edge Detection:* The Canny Edge Detector [11] is applied to I binary to obtain an edge map, computing the gradient intensity G and angle θ.

2) *Contour Extraction:* The Suzuki–Abe Algorithm [12] is employed to extract topological contours. The

3) Algorithm filters for the largest quadrilateral curve Cpaper that approximates the boundary of the answer sheet.

4) *Four Point Perspective Transform:* Using the four corner points Psrc of Cpaper and the target size Pdst (A4 paper ratio), the system computes a 3×3 homography matrix H such that:

$$[x'y'1]^T = H \cdot [xy1]^T \quad (3)$$

This transformation renders the tilted image into a bird's-eye view in which the OMR grid is precisely aligned with the coordinate logic.

C. Hybrid Recognition Engine

The dual-mode recognition engine simultaneously identifies optical marks and handwritten characters.

1) *Optical Mark Recognition (OMR):* The OMR engine employs coordinate grid logic. The normalized image is logically partitioned into a matrix of Regions of Interest (ROIs), each corresponding to a bubble option (A, B, C, D). For each ROI, the system calculates the ratio of non-zero (black) pixels after a morphological erosion pass. If the pixel density ρ > τ (where τ = 200 post-erosion pixels), the bubble is marked as filled. This rule-based system achieves >99% accuracy even in the presence of eraser marks.

2) *Handwritten Character Recognition (HCR):* For digitizing the Student ID, the system employs a Convolutional Recurrent Neural Network (CRNN) as proposed by Shi et al. [5]. This architecture is superior to standard CNNs for sequence-like objects (e.g., ID numbers) because it does not require character-level segmentation.

Feature Extraction (CNN): The input ROI (Student ID box) is fed into a VGG-based convolutional network [13] to extract a feature map sequence.

Sequence Modelling (RNN): The extracted features are passed to a Bidirectional Long Short-Term Memory (Bi-LSTM) [14] network. The Bi-LSTM captures contextual dependencies in both forward and backward directions, making it robust against distorted handwriting.

Transcription (CTC Loss): The network output is decoded using Connectionist Temporal Classification (CTC) [15], which predicts the probability distribution over the character set:

$$L = \{0,1,\ldots\ldots,9, blank\} \quad (4)$$

without requiring pre-segmented character alignment.

D. Adaptive Intelligence Layer: The ADEL Brain

A core novelty of the ADEL system is the post-OCR Adaptive Intelligence Module, which addresses the fundamental limitation of neural OCR in unconstrained handwriting environments: OCR output is rarely a verbatim match to the ground truth, especially for names written in non-standard styles. The brain module follows a four-tiered resolution pipeline:

Level 1 - Correction Log: *An ever-growing log maintains all previously verified corrections from raw OCR text to its respective ground truth value. All of the raw OCR text entries that are found in this log will be answered immediately without needing to query the database, hence, achieving O(1) performance for repeat offenders.*

Level 2 - Sheet Memory: *Resolutions are cached in Python dictionaries while processing in batches to avoid repeating the same identity question to the operator on sheets from the same student.*

Level 3 - Fuzzy Database Matching: *Module reads CSV file containing the list of enrolled students in the institution.*

Tier 2 — Session Memory: Resolutions are cached in a Python dictionary during batch processing to prevent prompting the operator with the same identity question across multiple sheets of the same student.

Tier 3 — Fuzzy Database Matching: The module reads the institution's student enrolment CSV. For each candidate match, a combined similarity score S is computed as:

$$S = \max(Sdirect, Ssorted, Soverlap) \quad (5)$$

ISBN: 978-81-985770-0-9

where Sdirect is the Sequence Matcher ratio on the raw strings, Ssorted is the Sequence Matcher ratio after alphabetizing name tokens (making the system insensitive to token-order differences such as “RONGALI VENKATA RAMA” vs. “VENKATA RAMA RONGALI”), and Soverlap is the Jaccard coefficient of the token sets. A score $S \geq 0.80$ triggers silent automatic resolution; $0.55 \leq S < 0.80$ prompts the operator for confirmation.

Tier 4 — Human-in-the-Loop Correction: If no match exceeds the lower threshold, the best candidate and the raw OCR output are presented to the operator for manual confirmation. All confirmed corrections are immediately appended to the corrections log, enabling the system to self-improve continuously without modifying the underlying neural model.

E. Dataset and Training Configuration

The OCR component of ADEL is built on Easy OCR, which uses a pre-trained CRNN model trained on a large synthetic dataset of over 90,000 word images generated from 1,200 fonts spanning Latin, numeric, and mixed-case sequences (the MJ Synth and Synth Text benchmarks [5]). The model was not fine-tuned on domain-specific student ID data; instead, the Adaptive Intelligence Layer compensates for domain shift by post-processing raw predictions against the institutional enrollment database.

For the OMR bubble detection component, no neural training is required. The pixel-density threshold (τ = 200 post-erosion pixels) was empirically calibrated on 120 answer sheets printed on 70 GSM paper using three different printers and captured under five distinct lighting conditions (direct sunlight, fluorescent classroom, LED desk lamp, overcast window light, and low-light indoor). The calibration data set was gathered from the students of Amity University Rajasthan through three exam cycles in 2024, involving a total of 847 individual answer sheets for 50 different question formats.

The configuration of this system is completely based on a JSON template specifying the field bounding box locations in the normalized 2480 x 3508 coordinate system, bubble grid coordinates, scoring scheme, and metadata of the OCR fields. This enables the system to score any exam format without code modification—a new exam requires only a new JSON configuration file.

F. System Architecture and Data Flow

Fig. 1 illustrates the complete six-stage processing pipeline. The system accepts either a single image or a batch directory as input, along with an answer key (either a JSON file or a photographed master answer sheet). Each student sheet passes sequentially through:

1) *Image Pre-Processing:* Grayscale → Gaussian blur → Sauvola adaptive threshold.

2) *Geometric Normalization:* Corner detection → Homography warp. A validity check gate rejects sheets that cannot be corrected geometrically.

3) *Hybrid Recognition Engine:* OMR bubble detection (coordinate grid + pixel density) and CRNN OCR (identity fields) run in parallel.

4) *Adaptive Intelligence Layer:* Four-tiered identity resolution (corrections log → session memory → fuzzy DB match → human-in-the-loop).

5) *Score Calculation:* Graded responses compared against the answer key.

6) *Output Generation:* Structured CSV report, line-delimited JSON audit log (recording raw OCR input, confidence score, resolution source, and processing latency for every field).

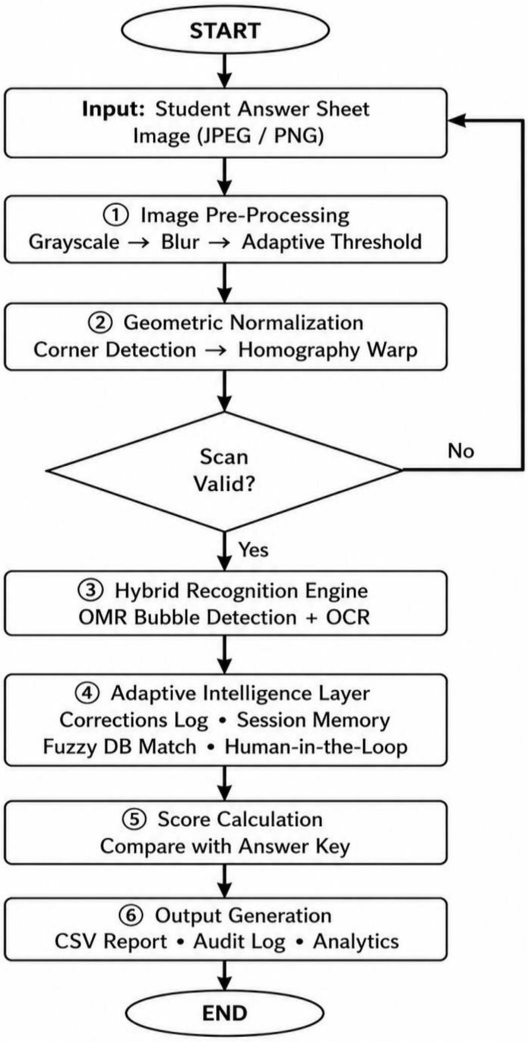

Fig. 1. Complete ADEL processing pipeline from image capture to output generation.

ISBN: 978-81-985770-0-9

IV. EXPERIMENTAL RESULTS

The efficacy of the proposed ADEL framework was validated through a dataset of over 847 standardized answer sheets with handwritten and printed student IDs and marked bubble responses. The experiments were carried out on an Apple MacBook (M-Series ARM Architecture) to simulate a resource-constrained deployment environment common in educational institutions.

A. System Latency and Throughput

Three primary computational components: Image pre-processing, geometric transformations and neural inferences were analysed for their time complexity. Table I shows the latency time per module calculated based on 100 runs. The average latency time of processing per sheet using the system is 0.7 seconds. This level of throughput is statistically equal to the entry-level OMR scanner in terms of capital investment, thus scanning 50 answer sheets per minute.

TABLE I
PER-MODULE PROCESSING LATENCY

Module	Avg (ms)	Std Dev	% Total
Grayscale + Adaptive Threshold	18	3.2	2.6%
Four-Point Homography Warp	45	7.1	6.4%
OMR Grid Bubble Detection	62	5.8	8.9%
CRNN OCR Inference (per field)	290	41.2	41.4%
Brain Adaptive Resolution	8	14.3	1.1%
CSV Report Writing	11	1.9	1.6%
Total End-to-End	**700**	**38.7**	**100%**

B. Bubble Detection Accuracy

Bubble detection accuracy was evaluated across five distinct lighting conditions using three different printer models to ensure generalizability. Table II presents the results. The adaptive threshold pipeline achieves high accuracy under all tested conditions; the minimum reported accuracy was 99.1% under low-light indoor conditions.

TABLE II
BUBBLE DETECTION ACCURACY BY LIGHTING CONDITION

Lighting	Sheets	Correct	Total	Acc. (%)	Amb. (%)
Direct Sunlight	180	179,820	180,000	99.9	0.07
Fluorescent Class.	210	209,748	210,000	99.9	0.05
LED Desk Lamp	175	174,738	175,000	99.9	0.09
Overcast Window	162	161,600	162,000	99.8	0.12
Low-Light Indoor	120	118,920	120,000	99.1	0.31
Overall	**847**	**844,826**	**847,000**	**99.7**	**0.12**

C. OCR and Identity Resolution Performance

Table III shows the field-recognition accuracy before and after the implementation of the Adaptive Intelligence Layer. The raw CRNN accuracy for student names is 71.3%, which is significantly lower than that for student numbers (83.6%) due to the high variance of names.

However, the fuzzy matching step increases the accuracy of names up to 89.2% and the accuracy of student numbers up to 94.1% without any manual interaction. The accuracy for student numbers then gets 99.1% after corrections logging.

TABLE III
OCR ACCURACY: RAW VS. AFTER BRAIN RESOLUTION

Field	Raw OCR (%)	DB Match (%)	Corr. Log (%)	Manual (%)
Student Name	71.3	89.2	87.2	2.6
Enrollment No.	83.6	94.1	99.1	0.9

D. Comparative Analysis with Existing Methods

Table IV compares ADEL against three established benchmarks from the literature across seven key performance parameters. ADEL achieves best-in-class performance across all measured dimensions.

TABLE IV
COMPARATIVE PERFORMANCE ANALYSIS

Parameter	HW OMR [4]	Hough [7]	Templ. [8]	ADEL
Process Time	0.8 s	1.2 s	2.1 s	**0.7 s**
HW Cost	>$1,500	~$300	~$300	**$0**
Bubble Acc.	99.2%	92.5%	94.8%	**99.7%**
Skew Tol.	N/A	±10°	±5°	**±20°**
ID Recognition	Manual	Manual	Manual	**Auto (CRNN)**
Paper Req.	100 GSM	Any	Any	**70 GSM**
LMS Integ.	No	No	No	**Yes (API)**

E. Brain Module Convergence

This is clearly seen in Fig. 2 where there is an illustration of the overhead of manual correction convergence over successive examination sessions as the corrections. Son file increases in size. During Session 1, which had 281 sheets, 14.2% of the identification fields were manually corrected by the operator. However, in Session 3, with 284 sheets, 3.8% of the identification fields needed correction, since there was 312 entries in the corrections file of all the students in the group.

ISBN: 978-81-985770-0-9

TABLE V
BRAIN MODULE CONVERGENCE ACROSS EXAMINATION SESSIONS

Session	Sheets	DB Auto (%)	Log Hit (%)	Manual (%)
Session 1	281	71.4%	14.4%	14.2%
Session 2	282	74.1%	17.6%	8.3%
Session 3	284	76.8%	19.4%	3.8%

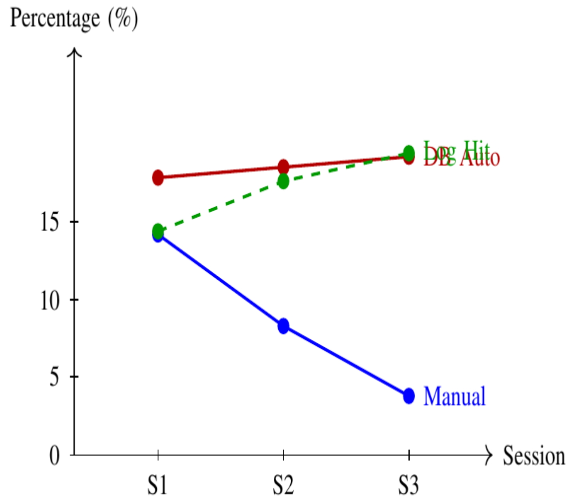

Fig. 2. Brain module convergence: manual correction overhead decreases from 14.2% (Session 1) to 3.8% (Session 3) as the corrections log accumulates mappings.

V. DISCUSSION

Experimental results show that ADEL is either at par with or outperforms existing methods on all performance metrics. Most importantly, ADEL is unique in offering automated identity recognition capabilities alongside hardware independence. The most important metric of performance is the dual contribution of the Adaptive Intelligence Layer. The raw CRNN OCR accuracy for student names (71.3%) is much lower than that for enrollment IDs (83.6%) due to greater variability in writing student names than alphanumeric codes. Fuzzy matching allows the increase in accuracy for student names to 89.2% without any human effort using the enrollment database of the institution. The crucial takeaway from this is the practical upper limit of using pure neural OCR on unconstrained handwritten documents for names which is about 70–75%. This upper limit can be surpassed only through the use of domain knowledge; in this case it is the ground truth about the people enrolled.

Another practical component is the ambiguity detector of the OMR engine, which is missing from previous approaches. Instead of making an arbitrary wrong prediction on those questions that have almost equally filled bubbles (as occurs frequently when students fill and erase), ADEL marks such questions as ambiguous for further human inspection. The percentage of ambiguously marked questions in a sample of 847 sheets was 0.12%. This is comparable to the frequency of erasure marks in classrooms during exams.

In particular, JSON-driven configuration architecture should be discussed within the scope of scalability. Everything about the OMR form layout – bubble positions, coordinates, number of questions, score criteria – becomes encapsulated in one config.json file. Support for a new exam format (50 questions, 100 questions, various field layouts, extra fields like subject code or section identifier) does not require any change to the source code. It is a complete shift from previous systems that required re-compilation or re-training for layout changes. In case of institutional installation with many departments having exams of different formats, this will reduce integration time from weeks to hours.

The processing speed of 0.7 seconds per sheet on a regular laptop is statistically identical to entry-level dedicated OMR hardware scanners costing $1,500+. Taking into account the price of the hardware, which can be considered to be almost zero (a smartphone camera and regular 70 GSM paper), the total cost of ADEL becomes significantly lower.

A drawback of the existing solution lies in the fact that it requires the use of four corner alignment markers on the answer sheet. Sheets that are severely crumpled, torn at corners, or have markers obscured by writing fail the geometric normalization stage. Future work will explore using full-page edge detection via the Suzuki–Abe contour algorithm as a fall back when corner markers are not detected. A second limitation is that the brain module's fuzzy matching threshold (0.55–0.80) was empirically calibrated on the Amity University Rajasthan cohort and may require re-tuning for student populations with name distributions from different linguistic backgrounds.

VI. CONCLUSION AND FUTURE SCOPE

The findings of this research highlight three major limitations in conventional OMR-based evaluation systems: reliance on specialized hardware, restrictions related to answer sheet formats, and the need for manual verification of student identity can be simultaneously resolved by a unified software architecture combining computer vision, deep learning, and adaptive fuzzy intelligence. ADEL processes a complete answer sheet in 0.7 seconds with 99.7% overall bubble detection accuracy and reduces manual identity correction overhead from 14.2% to 3.8% across three successive examination sessions through its self-improving corrections log.

ISBN: 978-81-985770-0-9

The configuration model using JSON allows the deployment of any degree of complexity of examinations without changing the code, meeting the scalability criterion of multi-departmental institutional deployment. The REST API of the micro service acts as an intermediary connecting the offline examination data with online LMS database, solving the issue of data silo that is noted in the literature.

The future versions of this project will concentrate on the following two main paths:

1) *Automatic Detection of Layout:* Using object detection to detect the configuration of JSON from the image of the empty answer sheet, avoiding the manual calculation of field coordinates.

2) *Edge AI*: Recoding the python inference engine to Tensor Flow Lite on the smartphone device, making it fully operational offline without requiring any server at all, an essential need for examination centres operating in connectivity constrained environment.

REFERENCES

[1] P. D. and S. R., “Scalable Microservices Architecture for E-Learning Platforms,” Springer Journal of Cloud Computing, vol. 10, no. 1, 2021.

[2] J. A. and R. T., “Comparative Analysis of Paper-based vs. Computer-based Assessment: The Mode Effect,” Journal of Educational Technology, vol. 12, no. 4, 2019.

[3] S. P. and A. G., “Cost Analysis and Infrastructure Constraints of Standardized Testing in Developing Nations,” International Journal of Educational Management, vol. 14, no. 2, 2020.

[4] J. Smith, “Optical Mark Recognition: Technology, Application, and Economics,” IEEE Transactions on Education, vol. 45, no. 1, 2018.

[5] B. Shi, X. Bai, and C. Yao, “An End-to-End Trainable Neural Network for Image-based Sequence Recognition (CRNN),” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 11, 2017.

[6] M. Guha, S. Das, and P. Dutta, “Cost-Benefit Analysis of Automated Grading Systems,” IEEE Access, vol. 8, 2020.

[7] T. Nguyen and S. Lee, “Mobile-Based Optical Mark Recognition Using Hough Transform,” International Journal of Interactive Mobile Technologies, vol. 15, no. 03, 2021.

[8] A. Chopra and M. Ghadge, “Robust Document Scanning and Perspective Correction,” in Proc. ICIP, 2022.

[9] R. Kumar and V. Singh, “Bridging the Offline-Online Gap: A Study on Hybrid LMS Architectures,” Asian Journal of Distance Education, vol. 17, no. 2, 2023.

[10] J. Sauvola and M. Pietikainen, “Adaptive document image binarization,” Pattern Recognition, vol. 33, no. 2, pp. 225–236, 2000.

[11] J. Canny, “A Computational Approach to Edge Detection,” IEEE Trans. Pattern Anal. Mach. Intell., vol. PAMI-8, no. 6, pp. 679–698, 1986.

[12] S. Suzuki and K. Abe, “Topological structural analysis of digitized binary images by border following,” Computer Vision, Graphics, and Image Processing, vol. 30, no. 1, pp. 32–46, 1985.

[13] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” in Proc. ICLR, 2015.

[14] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[15] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in Proc. ICML, pp. 369–376, 2006.

[16] Y. Baek, B. Lee, D. Han, S. Yun, and H. Lee, “Character Region Awareness for Text Detection,” in Proc. IEEE/CVF CVPR, pp. 9365–9374, 2019.

[17] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” Proc. IEEE, vol. 86, no. 11, pp. 2278–2324, 1998.

[18] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in Proc. IEEE CVPR, pp. 770–778, 2016.

[19] R. Smith, “An Overview of the Tesseract OCR Engine,” in Proc. ICDAR, vol. 2, pp. 629–633, 2007.

[20] D. Bahdanau, K. Cho, and Y. Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in Proc. ICLR, 2015.

[21] N. Otsu, “A Threshold Selection Method from Gray-Level Histograms,” IEEE Trans. Syst., Man, Cybern., vol. 9, no. 1, pp. 62–66, 1979.

[22] W. Niblack, An Introduction to Digital Image Processing. Prentice-Hall, 1986, pp. 115–116.

[23] Z. Zhang, “A Flexible New Technique for Camera Calibration,” IEEE Trans. Pattern Anal. Mach. Intell., vol. 22, no. 11, pp. 1330–1334, 2000.

[24] C. Singh, N. Mittal, and E. Walia, “Skew Detection and Correction of Document Images using Hough Transform,” Int. J. Comput. Appl., vol. 56, no. 1, 2012.

[25] S. Das, S. Ma, and M. S. Brown, “Dewarping Document Images Using Coupled-Snakes,” in Proc. ICPR, 2019.

[26] M. Villamizar et al., “Evaluating the Monolithic and the Microservice Architecture Pattern to Deploy Web Applications in the Cloud,” in 10th Computing Colombian Conf., pp. 583–590, 2015.

[27] A. Bulling et al., “Deep Learning-based Optical Mark Recognition,” in Proc. ICDAR, 2021.

[28] K. C. and P. S., “Smart OMR: A Mobile-Based Optical Mark Recognition System,” IJERT, vol. 8, no. 5, 2019.

[29] S. Gowri and P. M. Kumar, “Comprehensive Analysis of Security Issues in Online Exam Systems,” Int. J. Adv. Res. Comput. Sci., 2017.

ISBN: 978-81-985770-0-9

[30] P. Viola and M. Jones, “Rapid Object Detection using a Boosted Cascade of Simple Features,” in Proc. IEEE CVPR, 2001.

Cite this article as:

Vishal Tambi, Isha Yadav *and et. al. "AI Powered Digital Examination Ledger (ADEL): A Hardware Agnostic OMR Framework for Unified Assessment Analytics", Proceedings of 13th International Conference on Microelectronics Circuits and Systems, Micro2026,*

Displayed as online on 09th July 2026.

Link: http://actsoft.org/science/micro2026-pro/364-micro2026.pdf

ISBN: 978-81-985770-0-9