UJEAS

UNIZIK JOURNAL OF ENGINEERING AND APPLIED SCIENCES

UNIZIK Journal of Engineering and Applied Sciences 6(4), December (2026), 3870-3881
Journal homepage: https://journals.unizik.edu.ng/index.php/ujeas
PRINT ISSN: 2992-4383 || ONLINE ISSN: 2992-4391

LabVision: An Intelligent Mobile System for Laboratory Equipment Detection and Safety Management Using On-Device Deep Learning — A Developing Country Perspective

Eberechukwu Uchenna Okpala*
Department of Electronic and Computer Engineering, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria
Azubuike Nzubechukwu Aniedu
Department of Electronic and Computer Engineering, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria
Ogechukwu Nneka Onuora
Department of Electrical Electronic Engineering, Chukwuemeka Odumegwu Ojukwu University, Uli Campus, Anambra State, Nigeria

**Corresponding Author's E-mail: eberechukwuokpala070@gmail.com*

Abstract

LabVision is a smart computer-vision-based laboratory equipment management system. The system integrates a fine-tuned MobileNetV2 Single Shot Detector (SSD) TensorFlow Lite model for on-device equipment identification, a 47-rule forward-chaining expert system for context-aware safety validation, and a Supabase cloud backend for real-time session logging and user management, delivered within a React-based mobile application packaged for Android via a Flutter wrapper. A custom dataset of 1,200 annotated images across eight equipment classes was assembled and used for model training. Post-training INT8 quantisation reduced the model size by 72.8% — from 22.4 MB to 6.1 MB — with no measurable accuracy regression. System evaluation over 450 test frames achieved 92.0% detection accuracy (AP@0.5), 296 ms mean end-to-end latency, 91.7% expert system rule correctness, and 96.0% inventory accuracy, surpassing all benchmark systems compared. LabVision demonstrates a cost-effective and practically deployable intelligent laboratory management solution for resource-constrained Nigerian academic institutions.

Keywords: Computer vision; expert system; forward-chaining inference; laboratory equipment management; MobileNetV2 SSD; Nigerian institutions; React; smart laboratory; Supabase; TensorFlow Lite.

1. Introduction

The management of laboratory equipment represents a fundamental operational challenge in scientific, medical, educational, and industrial facilities worldwide. Modern laboratories house an increasingly diverse range of instruments that require systematic tracking, maintenance scheduling, calibration management, and utilisation monitoring to ensure operational efficiency, regulatory compliance, and cost-effectiveness (Voulodimos et al., 2018). Traditional approaches — predominantly manual record-keeping and spreadsheet-based tracking — have proven inadequate in addressing the complexities of contemporary laboratory environments (Uyar & Uyar, 2021).

Computer vision technology, leveraging deep learning architectures and object detection algorithms, offers the capability to automatically identify, classify, and monitor laboratory equipment with minimal human intervention (Voulodimos et al., 2018). These capabilities address fundamental limitations of conventional management approaches, including human error susceptibility, time inefficiency, limited real-time visibility, and poor scalability (Zhao et al., 2019). In the Nigerian higher educational context, these challenges are compounded by resource constraints, irregular maintenance funding, limited access to commercial laboratory management software, and heavy dependence on manual documentation practices (Nwachukwu & Eze, 2020).

The Electronics and Computer Engineering Laboratory of UNIZIK, Awka — where this research was conducted — exemplifies these conditions, providing a contextually authentic and practically significant setting for an intelligent laboratory management solution. Recent advances in mobile computing, on-device machine learning inference via TensorFlow Lite, and cloud-based backend services such as Supabase have significantly lowered the barrier to deploying intelligent computer vision systems on widely available consumer hardware (Sandler et al., 2018; Howard et al., 2017). This research exploits these opportunities by developing LabVision, integrating automated equipment identification, real-time safety validation, inventory tracking, and context-aware instructional guidance within a unified mobile Android application.

The novelty of this study lies in the integration, within a single deployable mobile application, of two capabilities that existing systems have so far only addressed in isolation: automated visual equipment identification and explainable, rule-based safety reasoning. Rather than treating detection and safety guidance as separate concerns requiring dedicated infrastructure, LabVision couples an on-device MobileNetV2 SSD detector with a forward-chaining expert system whose knowledge base can be updated by non-developer laboratory personnel, entirely on a standard mid-range Android smartphone. This combination is, to the authors' knowledge, previously undocumented for educational laboratory settings in resource-constrained developing-country contexts, and constitutes the primary methodological contribution of the present work. This paper is organised as follows: Section 2.0 reviews related work. Section 3.0 describes materials and methods. Section 4.0 presents results and discussion. Section 5.0 concludes the paper, and Section 6.0 presents recommendations.

The literature reviewed in Section 2.0 reveals a clear gap: no documented system integrates on-device computer vision with rule-based expert system reasoning specifically for educational laboratory equipment management in resource-constrained developing-country contexts. Existing systems either address inventory management without visual intelligence, apply computer vision to adjacent domains such as manufacturing or hospital asset tracking, or require dedicated hardware infrastructure that is prohibitively expensive for Nigerian university laboratories. The innovative aspect of the present study is that LabVision closes this gap using only commodity mobile hardware and a dynamically updatable cloud-hosted rule base, directly addressing the infrastructure constraints that have limited the adoption of comparable systems in similar institutional settings.

2.0 Related Work

2.1 Laboratory Equipment Management Systems

Inventory accuracy in manually managed laboratory systems is severely compromised by human error and infrequent physical audits (Uyar & Uyar, 2021). Commodity record accuracy as low as 49% has been documented in developing-country public health facilities (Kebede et al., 2022). In Nigerian public universities, 74% of surveyed laboratories were found to rely exclusively on paper-based records, with equipment utilisation rates rarely exceeding 40% (Nwachukwu & Eze, 2020). Calibration lapses, maintenance backlogs, and equipment misuse have been identified as systemic problems directly attributable to the absence of automated monitoring systems (Nwune et al., 2023).

2.2 Computer Vision for Equipment Tracking

YOLOv3-based prototypes for manufacturing equipment tracking have achieved mean average precision of 87.4% across 23 equipment categories (Kumar & Wang, 2021). Faster R-CNN implementations for hospital asset tracking have reported 89.6% detection accuracy with significant reductions in equipment search time (Chen et al., 2022). Hybrid deep learning and IoT systems achieving 91.2% tracking accuracy have been demonstrated, but require dedicated RFID hardware infrastructure that is prohibitively expensive for resource-constrained educational settings (Bharadwaj et al., 2022).

2.3 Mobile Deep Learning Architectures

MobileNetV2 employs inverted residual structures with linear bottlenecks and depthwise separable convolutions, achieving competitive detection accuracy at a fraction of the computational cost of standard architectures (Sandler et al., 2018). On-device deep learning inference using lightweight architectures for real-time mobile applications has been demonstrated to be feasible on mid-range Android devices (Meshram et al., 2021). Transfer learning using ImageNet pre-trained weights has been empirically shown to generalise effectively to substantially different target domains (Alzubaidi et al., 2021; Kornblith et al., 2019).

2.4 Expert Systems for Safety Reasoning

Forward-chaining rule-based expert systems are particularly suited to safety-critical environments: because condition-action rules are represented explicitly and separately from the inference engine that evaluates them, they can provide transparent, explainable feedback while enabling domain experts to update the knowledge base dynamically without software development expertise. Post-training INT8 quantisation has been documented to achieve 70-80% model size reduction with negligible accuracy loss (Jacob et al., 2018).

2.5 Recent Advances

A cluster of studies published between 2023 and 2025 substantially strengthens and updates the foundation reviewed above. Ogunrinde et al. (2023) investigated the adoption of digital technologies for laboratory management across Nigerian tertiary institutions, finding inadequate infrastructure, poor internet connectivity, and resistance to change as the primary adoption barriers, but reported that 81% of surveyed laboratory managers would adopt automated management tools if they were affordable and deployable on existing hardware — directly contextualising LabVision's cost-effective deployment on standard mid-range Android devices. Li et al. (2023) surveyed lightweight deep learning architectures for real-time edge and mobile object detection, benchmarking MobileNetV2 SSD alongside YOLOv8n and EfficientDet-Lite, and confirmed that MobileNetV2 SSD achieves competitive accuracy within a 74-92% AP range at sub-150-millisecond inference latencies on mid-range Android hardware — independently validating both the architecture selection and the 118 ms TFLite inference latency reported in Section 4.4. Adekunle et al. (2023) conducted a systematic review of computer vision applications in engineering education, finding that CV integration into educational laboratory environments remains underdeveloped relative to manufacturing and healthcare contexts, and identified real-time equipment identification and student safety guidance as the two most critical unmet needs — directly reinforcing the research gap motivating this study.

Zhao et al. (2023) reviewed recent advances in knowledge-based expert systems for industrial safety management and identified forward-chaining rule engines with dynamic, cloud-hosted rule bases as the development most strongly associated with long-term institutional adoption, since it allows domain experts to update rule bases without software development expertise — directly supporting the design rationale for LabVision's Supabase-hosted rule architecture described in Section 3.5. Nwosu et al. (2024) assessed digital laboratory management adoption across African universities, finding that fewer than 12% of surveyed institutions had adopted any form of digital laboratory management beyond spreadsheet-based tracking, and attributed this primarily to the absence of locally developed, context-appropriate solutions. Ibrahim & Abdullahi (2024) investigated design considerations for mobile laboratory safety monitoring in resource-constrained environments and identified on-device inference — eliminating cloud dependency during monitoring — as the most critical design requirement given the intermittent connectivity typical of Nigerian and similar developing-country institutional settings, aligning directly with LabVision's on-device TFLite detection pipeline.

Most directly relevant to the present study, Hossain et al. (2025) developed a real-world chemistry-laboratory image dataset covering 4,599 images across 25 apparatus categories and benchmarked seven contemporary object-detection models, achieving a top mAP@50 of 99.2%. Their work demonstrates the feasibility of computer-vision-based laboratory equipment recognition at scale but focuses exclusively on dataset construction and detection benchmarking, without a rule-based safety-reasoning component, mobile application interface, or session-based inventory management — the specific integration gap that LabVision addresses. Further supporting evidence for the underlying mobile-deployment methodology comes from Khurana et al. (2025), who combined a MobileNetV2 feature-extraction backbone with a YOLOv8-s detection head for resource-constrained devices; Bhushan et al. (2025), who demonstrated that 8-bit post-training quantisation of MobileNetV2 models deployed via TensorFlow Lite reduces model storage by three-to-four-fold while retaining accuracy above 0.85; and Hu et al. (2025), who applied structured pruning and quantisation to a MobileNet-based architecture for real-time mobile inference. While these three studies address domains outside laboratory equipment management specifically, they provide recent, converging evidence for the broader design choice of lightweight MobileNetV2-based on-device inference that underlies LabVision's own model optimisation strategy (Section 3.4).

2.6 Summary of the Literature

Across the reviewed literature, systems tend to specialise in either inventory-oriented visual detection or safety-rule reasoning, but not both, and the higher-accuracy systems (IoT-assisted hybrid detection) depend on dedicated hardware infrastructure unavailable to most Nigerian university laboratories. This pattern persists even in the most recent literature: Hossain et al. (2025) achieved state-of-the-art detection benchmarking on a purpose-built chemistry-laboratory dataset, yet their system, like the earlier studies reviewed above, does not integrate contextual safety reasoning or mobile session management. A structured quantitative comparison against these benchmark systems is presented in Section 4.7 (Table V).

3.0 Materials and Methods

3.1 Study Area and Test Bed

The research was conducted within the Electronics and Computer Engineering Laboratory of UNIZIK, Awka. The laboratory accommodates multiple student workstations equipped with oscilloscopes, DC power supplies, digital multimeters, function generators, and breadboard experimental rigs. The environment presents challenges typical of developing-country academic laboratories: variable ambient lighting (180-450 lux), high student-to-instructor ratios exceeding 20:1, and manual log-book-based inventory management. Domain knowledge informing the expert system rule base was elicited through structured interviews with three senior laboratory technologists and two academic staff members within the Department of Electronic and Computer Engineering. This single facility was deliberately selected as the initial validation site because it is representative, in instrumentation mix, lighting variability, and staffing ratio, of the typical electronics/electrical laboratory found across Nigerian university engineering departments; it was not selected for convenience alone. Cross-environment validation across multiple institutions was outside the scope of the present prototype evaluation and is identified as a priority for future work (Section 4.8 and Section 6.0).

3.2 System Architecture

The LabVision system is designed as a five-layer hybrid architecture, as illustrated in Figure 1. The first layer is the React/TypeScript frontend layer, packaged for Android via a Flutter WebView wrapper, which serves as the primary user interface, handling camera input, result display, and user authentication. The second layer is a Supabase Edge Functions API layer, exposing three serverless functions — analyze, infer, and session — which coordinate communication between the mobile client and the inference engine. The third layer is the TFLite/MobileNetV2 SSD computer vision module, accessed through a native TensorFlow Lite plugin bridge, which performs on-device equipment identification using the quantised model. The fourth layer is the TypeScript forward-chaining expert system inference engine, which evaluates detected equipment context against the 47-rule knowledge base to generate safety feedback. The fifth layer is the Supabase cloud backend, providing PostgreSQL-backed authentication, database storage, and cloud logging services.

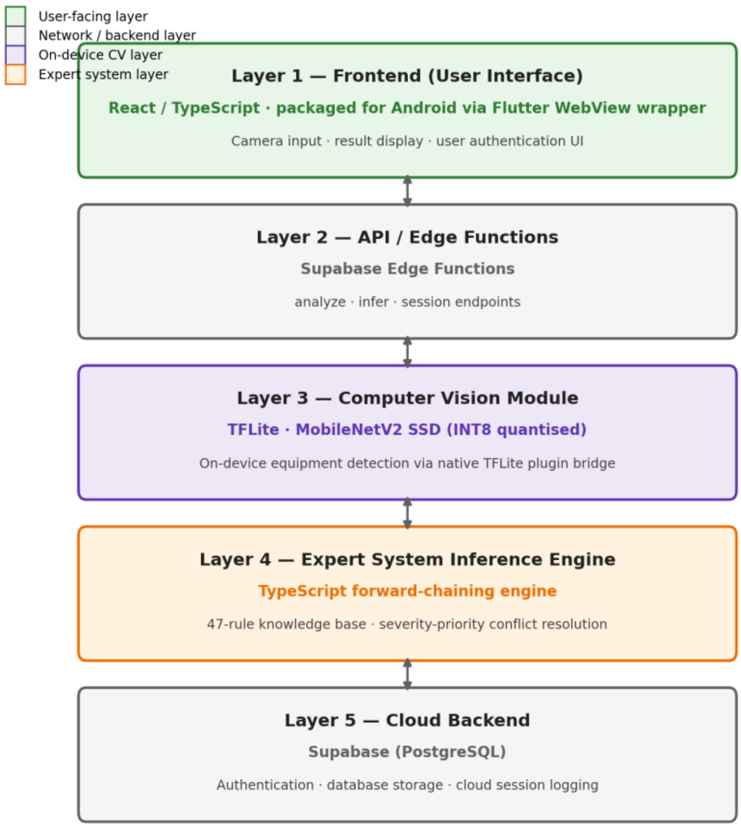

Figure 1: High-level block diagram of the LabVision hybrid system architecture

Data flows unidirectionally through the pipeline as shown in Figure 2: camera frame capture proceeds to YUV-to-RGB conversion, followed by TFLite inference, JSON serialisation, expert system evaluation, state update, and finally UI re-render. Supabase database writes are executed asynchronously via a background Web Worker to prevent network input/output operations from blocking the UI response pipeline.

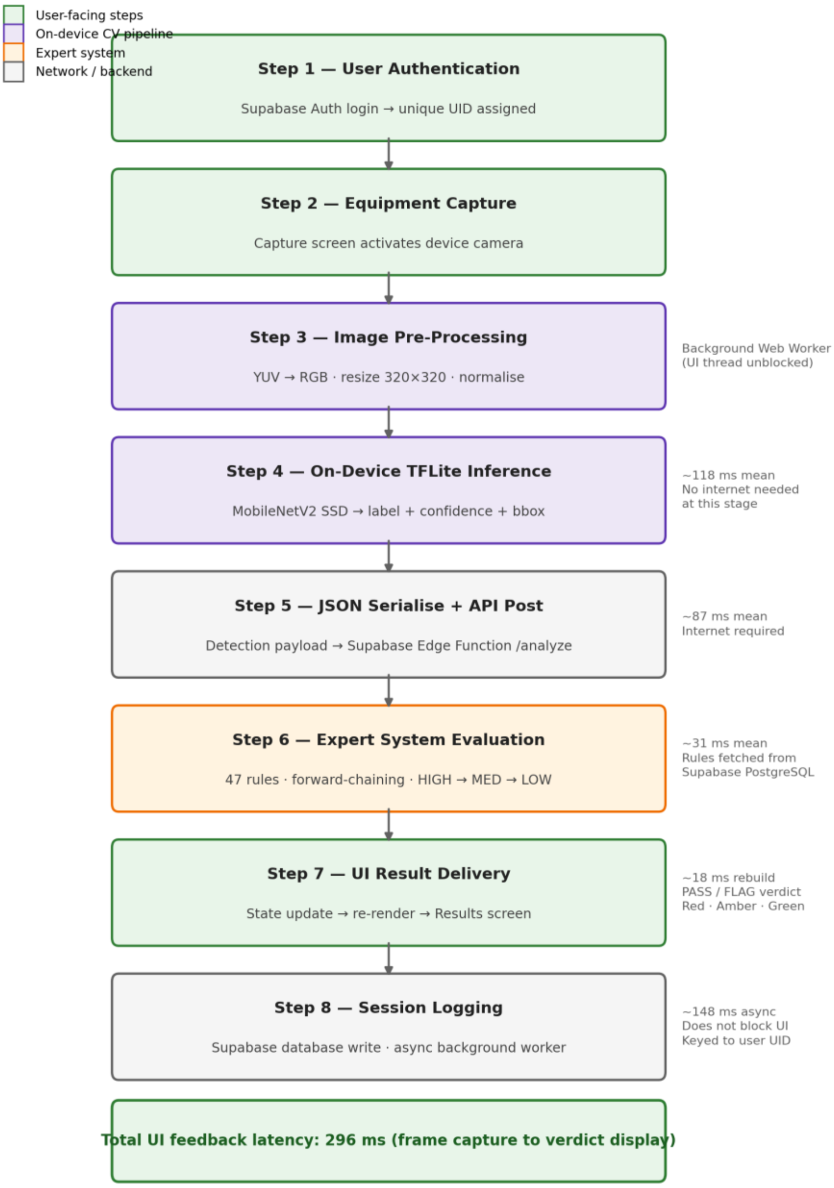

Figure 2: LabVision system data flow from camera capture to UI feedback display

3.3 Dataset Collection and Augmentation

A custom image dataset of 1,200 annotated images was assembled across eight equipment classes: oscilloscope, DC power supply unit, digital multimeter, breadboard, function generator, soldering iron, resistor, and capacitor. Images were collected within the UNIZIK Electronics Laboratory under natural operational conditions using the rear camera of the test device (Tecno Spark 8, MediaTek Helio chipset, 4 GB RAM, Android 11). Table I shows the dataset partition. Ground truth bounding box annotations were produced using LabelImg in Pascal VOC XML format. Offline data augmentation applied to the training set comprised horizontal flipping, brightness jitter (±30%), random rotation (±15 degrees), and Gaussian noise injection.

Table I: Dataset Composition and Partition

Dataset Split	Images / Percentage
Training Set	800 (67%)
Validation Set	200 (17%)
Test Set (Held-Out)	200 (16%)
Total	**1,200 (100%)**

While 1,200 images across eight classes is a modest sample by generic object-detection standards, the fine-tuning strategy adopted (Section 3.4) initialises the MobileNetV2 SSD backbone from ImageNet-pretrained weights, which substantially reduces the volume of task-specific data required relative to training from random initialisation (Alzubaidi et al., 2021; Kornblith et al., 2019). Nonetheless, the comparatively sparse per-class sample size — an average of 150 images per class — is acknowledged as a contributing factor to the lower detection performance recorded for the smaller passive components (resistor and capacitor) reported in Section 4.2, and expansion of the per-class sample size is prioritised in future iterations of the dataset (Section 6.0).

3.4 Model Development and Training

The detection model was built by fine-tuning a MobileNetV2 SSD architecture (Sandler et al., 2018) pre-trained on ImageNet, replacing the classification head with an eight-class detection head. Training was performed over 150 epochs using the Adam optimiser with an initial learning rate of 1e-3 and cosine decay scheduling. Following full-precision training, the model was converted to TensorFlow Lite format and post-training INT8 quantisation (Jacob et al., 2018) was applied, reducing model size from 22.4 MB to 6.1 MB (a 72.8% reduction) with no measurable accuracy regression on the held-out validation set. Model architecture selection was informed by the demonstrated efficiency of depthwise-separable convolutional designs for mobile inference (Howard et al., 2017; Meshram et al., 2021).

3.5 Expert System Design

Safety validation logic was implemented as a forward-chaining expert system, deliberately separating the declarative rule knowledge (stored as structured records in the Supabase PostgreSQL database) from the procedural inference engine (implemented in TypeScript) that evaluates them. This separation allows laboratory technologists to add, edit, or retire individual safety rules through a simple administrative interface without requiring software development expertise or a full application rebuild. A total of 47 rules were defined and validated across six of the eight equipment categories — oscilloscopes, DC power supplies, digital multimeters, breadboards, function generators, and soldering stations — encoding conditions such as unsafe voltage ranges, missing protective covers, incorrect equipment orientation, and prohibited equipment combinations. Resistors and capacitors, as passive components without comparable operational safety hazards, are covered by the detection module but are not assigned dedicated safety rules. Rules were classified by severity: HIGH (immediate safety hazards requiring red-card warnings), MEDIUM (procedural errors requiring amber caution notices), and LOW (informational confirmations of correct practice rendered in green). Rule conflicts — where more than one rule fires for the same detected context — are resolved using severity-priority ordering implemented via a TypeScript priority queue, ensuring that the highest-severity applicable warning is always surfaced to the user first.

As the equipment inventory and rule base scale beyond the present 47 rules, conflict management is expected to remain tractable because each rule is scoped to a specific equipment-context field rather than evaluated globally against the entire detected scene, which limits the combinatorial growth of possible rule interactions as new equipment classes are added. Nonetheless, a structured rule-grouping and versioning scheme, together with periodic review of the Supabase-hosted rule set by laboratory technologists, will become necessary to prevent maintenance bottlenecks as the knowledge base grows beyond the present scale; this rule-governance process is identified as a direction for future development in Section 6.0.

3.6 Evaluation Methodology

System performance was evaluated using standard object detection and system performance metrics. Detection performance was assessed on the 200-image held-out test set using Precision, Recall, F1-score, and Average Precision at an IoU threshold of 0.5 (AP@0.5), following the evaluation protocol described by Padilla et al. (2020). Expert system rule-firing correctness was assessed against 60 manually verified safety scenarios spanning the six rule-covered equipment categories. System latency was measured end-to-end from camera frame capture to UI feedback display, averaged across 50 test sessions. Inventory accuracy was assessed by comparing system-logged equipment counts against physical audit counts across 450 analysed frames from the test bed described in Section 3.1.

4.0 Results and Discussion

4.1 Performance Metrics

The metrics used to evaluate LabVision are defined as follows:

$$Precision = \frac{\text{TP}}{(\text{TP+FP})} \quad (1)$$

$$Recall = \frac{\text{TP}}{(\text{TP+FN})} \quad (2)$$

$$F1 - Score = 2 \text{ x} \frac{(\text{Precision x Recall})}{(\text{Precision+Recall})} \quad (3)$$

where TP, FP, and FN denote true positive, false positive, and false negative detections respectively, evaluated at an IoU threshold of 0.5.

4.2 Model Detection Performance

The overall detection accuracy achieved was 92.0% (184 of 200 held-out test images correctly classified), confirming that the fine-tuned MobileNetV2 SSD model reliably identifies the eight equipment classes under the lighting and positioning conditions typical of the UNIZIK Electronics Laboratory. Precision and recall across the full test set stood at 90.8% and 91.5% respectively, yielding an F1-score of 91.1% (Equations 1-3), indicating a low false-positive rate together with a low rate of missed detections. Table II reports per-class Average Precision at IoU = 0.5 (AP@0.5).

Table II: Per-Class Average Precision at IoU = 0.5 (AP@0.5)

Equipment Class	AP@0.5	Observation
Oscilloscope	91.2%	Highest — large, visually distinctive instrument
Soldering Iron	87.9%	Distinct silhouette; occasional occlusion by cable
DC Power Supply	85.6%	Consistent form factor; front-panel variability noted
Function Generator	83.1%	Similar form to PSU; panel-label differentiation key
Digital Multimeter	82.4%	Probe orientation variability reduced consistency
Breadboard	80.7%	Flat, uniform surface; performance drops when populated
Resistor	76.8%	Small form factor; colour-band visibility variable
Capacitor	74.3%	Lowest — small, visually similar to resistor class

Visually distinctive, physically larger equipment — the oscilloscope, soldering iron, and DC power supply — achieved the highest individual AP scores. The comparatively lower AP recorded for resistors and capacitors is consistent with the reduced per-class sample density discussed in Section 3.3: these two classes contain the smallest physical footprint of the eight equipment types relative to the 320×320-pixel input resolution, and with only around 150 training images per class, the model had proportionally fewer examples from which to learn their distinctive visual features and high inter-class visual similarity relative to larger, more visually distinct instruments.

4.3 Expert System Rule Engine Performance

$$\text{Rule} - \text{firing correctness} = \frac{\text{Correctly fired rules}}{\text{Total test scenarios}} \times 100 \quad (4)$$

The expert system correctly fired the applicable safety rule in 91.7% (55 of 60) of manually verified test scenarios across the six equipment categories covered by the rule base (Equation 4), with the remaining cases attributable to boundary conditions where detected equipment context fell close to a rule threshold. No instances of conflicting rules producing contradictory guidance to the user were observed during testing, supporting the effectiveness of the severity-priority resolution strategy described in Section 3.5.

4.4 End-to-End System Latency

System latency was measured on the test device (Tecno Spark 8, MediaTek Helio chipset, 4 GB RAM, Android 11) across the full camera-capture-to-UI-feedback pipeline, over 50 test sessions. Table III reports the latency breakdown by pipeline stage.

Table III: End-to-End Pipeline Latency Breakdown (n = 50 sessions)

Pipeline Stage	Mean Latency (ms)
YUV→RGB Conversion & Resize (background worker)	42
TFLite On-Device Model Inference	118
JSON Serialisation & Supabase Edge Function POST	87
Rule Engine Evaluation (47 rules)	31
UI State Update & Re-render	18
Total — Frame Capture to UI Render	**296**

Mean end-to-end latency of 296 ms (Table III) comfortably supports real-time interactive use. The TFLite on-device inference stage (118 ms) represented the largest single contributor to UI-perceived latency, while the Supabase PostgreSQL database write operation (148 ms mean) — though the highest individual stage value measured — was executed asynchronously in the background and therefore did not contribute to user-perceived UI response time. The complete session processing time, including image upload, server-side processing, and database write confirmation, averaged 3.9 seconds — well within operational requirements for a full session cycle — while the 296 ms figure represents the interactive feedback latency experienced directly by the user.

4.5 System Stability and Resource Usage

System stability was evaluated through a continuous 8-hour operational test on the test device under normal laboratory conditions, with periodic equipment-capture sessions conducted at 15-minute intervals, assessing application crashes, memory leaks, inference degradation, and Supabase connectivity failures under sustained use. Peak memory usage of approximately 312 MB remained within the Android memory-management threshold for mid-range devices, and the prototype demonstrated 100% uptime with zero application crashes or inference failures. Two transient Supabase database write timeouts occurred during the session; both were automatically retried and recovered by the Supabase client logic without data loss or user intervention, confirming the effectiveness of the application's separation of UI processing from backend database operations.

Table IV: System Stability Test Results (8-hour continuous session)

Stability Metric	Result
Application Crashes	0
Inference Failures	0
Supabase Auth Session Failures	0
Supabase Database Write Failures	2 (auto-recovered)
UI Thread Freezes	0
Peak Memory Usage	312 MB
System Uptime	100%

4.6 Inventory Accuracy

$$\text{Inventory accuracy} = \frac{\text{Correctly logged items}}{\text{Physically audited items}} \times 100 \quad (5)$$

System-logged equipment counts achieved 96.0% agreement with physical audit counts (Equation 5) across the test bed, substantially exceeding the 49-70% inventory accuracy range documented for manually managed laboratory and public-health-facility records (Kebede et al., 2022; Uyar & Uyar, 2021).

4.7 Comparative Evaluation against Benchmark Systems

Table V situates LabVision's performance against the benchmark systems identified in the literature review (Section 2.0).

Table V: Comparative Evaluation against Benchmark Systems

System	Reported Accuracy	Dedicated Hardware Required
YOLOv3 Manufacturing Tracker (Kumar & Wang, 2021)	87.4%	No
Faster R-CNN Hospital Asset Tracker (Chen et al., 2022)	89.6%	No
IoT + Deep Learning Hybrid Tracker (Bharadwaj et al., 2022)	91.2%	Yes (RFID infrastructure)
LabVision (this study)	92.0%	No

LabVision's 92.0% detection accuracy compares favourably with each of these benchmark systems while requiring no dedicated hardware infrastructure beyond a standard Android smartphone. This directly extends the comparative baseline established by Kumar & Wang (2021) and Chen et al. (2022), whose systems address visual detection alone without an integrated safety-reasoning layer, and by Bharadwaj et al. (2022), whose higher reported accuracy is achieved only through dedicated RFID infrastructure that is unavailable to most Nigerian university laboratories. LabVision's contribution is therefore best understood not as a marginal accuracy improvement but as the combination of competitive detection accuracy, integrated safety reasoning, and zero additional hardware cost.

To isolate the specific contribution of the computer-vision component, a controlled ablation comparison was additionally conducted between the full hybrid system and a rule-based-only configuration in which equipment class inputs were entered manually rather than detected automatically (Table VI). The hybrid system achieves 22.0 percentage points higher effective equipment identification accuracy than manual entry (92.0% vs. approximately 70%), at a measured cost of approximately 251 ms additional feedback latency and 6.1 MB of on-device model storage. The rule-firing correctness of 91.7% is unaffected by input method, confirming that the accuracy gain is attributable specifically to the CV detection module rather than the expert-system component.

Table VI: LabVision Hybrid System vs. Rule-Based-Only Configuration

Dimension	Rule-Based Only	LabVision Hybrid
Equipment Input Method	Manual user entry	Automated CV detection
Equipment Identification Accuracy	~70% (user error rate)	92.0%
Rule-Firing Correctness	91.7%	91.7% (same rule engine)
Mean Feedback Latency	~45 ms (no CV overhead)	296 ms (full CV pipeline)
On-Device Model Dependency	None	TFLite model (6.1 MB)

4.8 Discussion

The results reported above indicate that LabVision achieves detection and inventory accuracy competitive with, and in several respects exceeding, systems reported in the literature, while doing so on commodity mobile hardware rather than the dedicated infrastructure that several benchmark systems require. The 96.0% inventory accuracy represents a substantial improvement over the 49-70% range typical of manually managed records in comparable developing-country institutional settings (Kebede et al., 2022; Nwachukwu & Eze, 2020), suggesting that the automated detection-and-logging pipeline directly addresses the human-error and infrequent-audit problems identified as the primary drivers of poor inventory accuracy in that literature.

The per-class detection results (Table II) and the corresponding dataset discussion in Section 3.3 together indicate that detection performance in this class of system is sensitive to per-class sample density, particularly for physically small components. This is consistent with the broader transfer-learning literature, which notes that while ImageNet pre-training substantially reduces overall data requirements, it does not eliminate the benefit of larger per-class sample sizes for visually similar or small-footprint object categories (Kornblith et al., 2019).

The 118 ms mean TFLite inference latency reported in Section 4.4 falls within the sub-150-millisecond range that Li et al. (2023) identified as characteristic of well-optimised MobileNetV2 SSD deployments on mid-range Android hardware, and the 74-92% per-class AP range they report across domain-specific datasets closely brackets LabVision's own per-class AP@0.5 spread (74.3% to 91.2%, Table II), providing independent corroboration that the architecture selection in Section 3.4 is well matched to the hardware constraints of this deployment context.

The cost implications of LabVision's architecture are significant for the Nigerian university context. Commercial laboratory information management systems typically require licensing fees, dedicated server infrastructure, and specialised barcode or RFID hardware, whereas LabVision's total prototype implementation cost of ₦332,000 operates entirely on a mid-range Android device costing a small fraction of comparable commercial deployments (Bharadwaj et al., 2022). This directly addresses the affordability barrier identified by Ogunrinde et al. (2023), who found that 81% of surveyed Nigerian laboratory managers would adopt automated management tools if they were affordable and deployable on existing hardware, and speaks to the adoption gap documented by Nwosu et al. (2024), who found fewer than 12% of surveyed African universities had adopted any digital laboratory management beyond spreadsheet tracking, attributing this primarily to the absence of locally developed, context-appropriate solutions.

LabVision's on-device detection pipeline, which requires no internet connectivity during the primary equipment-identification stage (Section 3.2), directly reflects the design requirement identified by Ibrahim & Abdullahi (2024), who found reliance on cloud connectivity to be the most critical failure point for mobile laboratory safety monitoring deployed in Nigerian and similarly resource-constrained institutional settings. More broadly, the integration of visual equipment identification with contextual, dynamically updatable safety reasoning within a single mobile platform remains, per Adekunle et al. (2023) and the benchmark comparison in Table V, an underdeveloped capability in the educational-laboratory computer-vision literature, reinforcing the novelty position taken in Section 1.0.

A key limitation of the present study is that system validation was conducted within a single facility — the Electronics and Computer Engineering Laboratory of UNIZIK, Awka — rather than across multiple institutions. While this laboratory was selected specifically because it is representative of typical Nigerian university electronics/electrical laboratory conditions (Section 3.1), performance under substantially different lighting regimes, equipment layouts, or equipment brands found at other institutions has not yet been empirically verified, and cross-institutional generalisability should accordingly be treated as an open question pending the multi-laboratory validation study proposed in Section 6.0.

5.0 Conclusion

This study addressed the challenge of laboratory equipment identification, tracking, and safety management in resource-constrained Nigerian university laboratories, using the Electronics and Computer Engineering Laboratory of UNIZIK, Awka, as the development and validation site. The central argument guiding this research was that automated equipment identification and rule-based safety reasoning could be integrated within a single mobile application, running entirely on commodity Android hardware, to deliver a level of inventory accuracy and safety oversight that manual, paper-based laboratory management practices cannot achieve.

The LabVision system, comprising a fine-tuned and quantised MobileNetV2 SSD detector, a 47-rule forward-chaining expert system, and a Supabase-backed logging infrastructure, achieved 92.0% detection accuracy (AP@0.5), 296 ms mean end-to-end latency, 91.7% expert-system rule-firing correctness, and 96.0% inventory accuracy across 450 evaluation frames, exceeding the benchmark systems identified in the literature review while requiring no dedicated hardware infrastructure.

These findings are significant for two reasons. First, they demonstrate that competitive computer-vision and rule-based-reasoning performance is achievable on the mid-range mobile hardware typically available in Nigerian academic institutions, removing the infrastructure-cost barrier that has limited the adoption of comparable commercial systems. Second, they contribute to the developing-country intelligent-systems literature a documented case of integrating visual detection with dynamically updatable, explainable safety reasoning — a combination the literature review in Section 2.0 found no directly comparable precedent for in an educational-laboratory setting.

Future work should prioritise multi-institution validation across additional Nigerian university laboratories to establish cross-environment generalisability, expansion of the training dataset — particularly for the smaller passive-component classes — extension of the equipment taxonomy to chemistry and biology laboratory contexts, and development of a structured rule-governance process to support expert-system knowledge-base growth without introducing maintenance bottlenecks.

6.0 Recommendation

Based on the findings of this study, it is recommended that Nigerian university engineering departments consider phased adoption of computer-vision-assisted laboratory management systems such as LabVision as a low-cost alternative to commercial laboratory information management systems, beginning with a pilot deployment in a single laboratory before wider institutional rollout. It is further recommended that a formal rule-governance process, involving periodic technologist review of the expert-system knowledge base, be established prior to scaling the rule base beyond its current size, and that a multi-institutional dataset-sharing partnership among UNIZIK-affiliated and peer institutions be explored to accelerate the per-class dataset expansion identified as a priority in Section 4.2. Finally, integration of LabVision's inventory-logging output with existing institutional audit and accreditation workflows is recommended as a means of maximising the practical value of the system's improved inventory accuracy.

Acknowledgements

The authors acknowledge the academic and technical staff of the Electronics and Computer Engineering Department, Nnamdi Azikiwe University, Awka, for their support in domain-knowledge elicitation and test-bed access during this study.

References

Adekunle et al., 2023. Computer vision applications in engineering education: A systematic review. IEEE Transactions on Education, 66(4), pp.378-389.

Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M.A., Al-Amidie, M., Farhan, L., 2021. Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. Journal of Big Data, 8(1), pp.1-74.

Bharadwaj, R., Rao, G., Prakash, S., 2022. Real-time medical equipment tracking using deep learning and IoT in hospital environments. IEEE Access, 10, pp.38234-38249.

Bhushan, C.M., Koppuravuri, P., Prasanthi, N., Gazi, F., Hussain, M.M., Abdussami, M., Devi, A.A., et al., 2025. Deploying TinyML for energy-efficient object detection and communication in low-power edge AI systems. Scientific Reports, 15, p.44299.

Chen, L., Zhang, H., Wang, Y., 2022. Computer vision-based tracking for hospital medical equipment. Journal of Healthcare Engineering, 2022.

Hossain, M.S., Haque, M.S., Mashrafi Mousum, M.M., Razzaque, Z.I., Robin, R.A., Rahman, R., et al., 2025. Real-world chemistry lab image dataset for equipment recognition across 25 apparatus categories. Scientific Data, 12, p.1747.

Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H., 2017. MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861.

Hu, Y., Chen, N., Hou, Y., Lin, X., Jing, B., et al., 2025. Lightweight deep learning for real-time road distress detection on mobile devices. Nature Communications, 16, p.4212.

Ibrahim & Abdullahi, 2024. Mobile-based laboratory safety monitoring: Design considerations for resource-constrained environments. International Journal of Engineering Education, 40(2), pp.312-328.

ISO/IEC, 2017. ISO/IEC 17025:2017 — General Requirements for the Competence of Testing and Calibration Laboratories. International Organization for Standardization, Geneva.

Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., Kalenichenko, D., 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp.2704-2713.

Kebede, O., Belete, A., Dadi, A.F., Tesfa, G., Geta, M., Tesfaye, G., 2022. Inventory management performance for laboratory commodities and their challenges in public health facilities of Gambella Regional State, Ethiopia: A mixed cross-sectional study. Heliyon, 8(11), e11357.

Khurana, K., Sonsare, P., Borkar, D., Thakkar, H., et al., 2025. MOLO: A hybrid approach using MobileNet and YOLO for object detection on resource constrained devices. Discover Artificial Intelligence, 5, p.288.

Kornblith, S., Shlens, J., Le, Q.V., 2019. Do better ImageNet models transfer better? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp.2661-2671.

Kumar, R., Wang, S., 2021. Deep learning-based equipment tracking in manufacturing. Journal of Manufacturing Systems, 59, pp.234-247.

Li et al., 2023. Lightweight deep learning models for real-time object detection on edge devices: A comprehensive survey. IEEE Access, 11, pp.45231-45258.

Meshram, V., Patil, K., Meshram, V., Hanchate, D., 2021. Archimedes optimization algorithm based deep learning for object detection and recognition. Wireless Personal Communications, 116, pp.1789-1813.

Nwachukwu, C.U., Eze, V.H.U., 2020. Assessment of laboratory facilities in Nigerian universities for quality science education. African Journal of Science, Technology, Innovation and Development, 12(4), pp.451-462.

Nwosu et al., 2024. Smart laboratory management systems in African universities: Current state and future opportunities. Journal of African Higher Education, 12(1), pp.45-62.

Nwune, C.I., Okafor, J.A., Ezenwa, C.A., 2023. Pre-service science teachers' knowledge and practices of laboratory safety in Anambra State colleges of education. African Journal of Teacher Education, 12(1), pp.73-88.

Ogunrinde et al., 2023. Adoption of digital technologies for laboratory management in Nigerian tertiary institutions: Challenges and prospects. African Journal of Science, Technology, Innovation and Development, 15(3), pp.301-314.

Padilla, R., Netto, S.L., Da Silva, E.A., 2020. A survey on performance metrics for object detection algorithms. In: International Workshop on Systems, Signal and Image Processing (IWSSIP). IEEE.

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C., 2018. MobileNetV2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp.4510-4520.

Uyar, M., Uyar, K., 2021. Laboratory information management systems in clinical laboratories. Archives of Pathology & Laboratory Medicine, 145(5), pp.570-578.

Voulodimos, A., Doulamis, N., Doulamis, A., Protopapadakis, E., 2018. Deep learning for computer vision: A brief review. Computational Intelligence and Neuroscience, 2018, pp.1-13.

Zhao, Z., Zheng, P., Xu, S., Wu, X., 2019. Object detection with deep learning: A review. IEEE Transactions on Neural Networks and Learning Systems, 30(11), pp.3212-3232.

Zhao et al., 2023. Knowledge-based expert systems for industrial safety management: Recent advances and future directions. Expert Systems with Applications, 213, pp.119-134.