DRISHTI: Drug Recognition and Integrated System for Helping the Visually Impaired with Tag-based Identification

Sajeeb Das[1,†], Srijit Paul[2,†], Ucchas Muhury[3], Akib Jayed Islam[4], Dhruba Jyoti Barua[5], Sultanus Salehin[6], Prasun Datta[7],
[1,5]National Institute of Technology Rourkela, [2,3]National Institute of Technology Warangal, [4]Norwegian University of Science and Technology, [6]Islamic University of Technology, [7]Bangladesh University of Engineering and Technology
[†]Joint first author. Both authors contributed equally to this work.
Email: hi5.sajeeb@gmail.com, srijitpaul1234567@gmail.com, ucchasmuhury@gmail.com, akibjayedislam@gmail.com, dhruba.barua099@gmail.com, salehin.iut@gmail.com, prasundatta.buet@gmail.com

Abstract

DRISHTI is a novel RFID-vision integrated assistive medication-verification system that combines RFID contactless scanning, quantized AI-based vision processing, and adaptive audio feedback to provide comprehensive medication-safety assurance. The architecture integrates an MFRC522 RFID reader for rapid drug-container identification, a Raspberry Pi–mounted camera running a quantized Gemma3-4B vision model for prescription-document analysis, and a hierarchical validation engine employing confidence-weighted scoring across five critical safety dimensions. Operating entirely offline, the system processes compressed medication data through multi-criteria classification while preserving user privacy and eliminating cloud dependencies. In evaluations across 149 test scenarios, DRISHTI achieved 86.57% overall accuracy and 100% detection of safety-critical cases, including expired medications, dosage mismatches, and drug interactions. The system delivers sub-millisecond response times with real-time, urgency-differentiated audio feedback, offering a practical solution for enhancing independence and reducing healthcare risks for visually impaired individuals.

1 Introduction

Managing medication is critical for individuals with visual impairments, as over 49.1 (Bourne et al., 2020) million people worldwide are blind and many face challenges with medication safety (Gupta et al., 2023). Traditional aids such as braille labels, tactile markers, or caregiver support help but limit accessibility, independence, and privacy.

Emerging technologies enable safer medication management through RFID automated adherence systems (Meshram et al., 2021), audio-based navigation tools (Zare et al., 2023), and computer vision approaches including YOLO-OCR-based pill identification (Dang et al., 2024) and camera-based smart medication boxes (Meshram et al., 2021). However, existing systems like ScripTalk remain centralized and non-portable, while vision-assisted solutions depend on cloud services, raising privacy and latency concerns. Integration of RFID with real-time AI-based label verification in standalone edge systems remains unexplored.

Edge-based AI systems demonstrate efficient, private inference for visually impaired assistance, with deployments using specialized hardware (Mahendran et al., 2021) and Raspberry Pi platforms with vision-language models (Baig et al., 2024). However, coordinated hardware integration (RFID, camera, audio) with quantized models and local interfaces for medication safety remains unaddressed.

This work proposes a standalone, dual-layer medication verification system leveraging the Raspberry Pi to deliver a novel assistive technology for safe and independent medication use by blind and visually impaired individuals. The primary contributions of this work are:

- A novel tri-modal verification system combining RFID scanning, real-time prescription analysis, and adaptive audio feedback for blind users' accessibility and reliability.
- Submillisecond verification pipeline with hierarchical validation of five safety axes: authenticity, timing, dosage, formulation, allergies; ensuring realtime precision.
- Fully offline, privacy-preserving edge solution that locally processes and syncs prescription data eliminating cloud reliance while ensuring secure and atomic records.
- Audio-first feedback system delivering adaptive, prioritized messages. Achieves

Proceedings of the Workshop on Beyond English: NLP for all Languages in an Era of LLM from Texts associated with RANLP 2025, pages 149–158, Varna, Bulgaria, Sep 12, 2025.
https://doi.org/10.26615/978-954-452-105-9-017

100% detection of critical medication hazards, enabling blind users to receive instant, non-visual alerts.

2 Related Work

Electronic adherence monitoring systems like MEMS, smart pill bottles, and ingestible sensors automatically track medication intake but often rely on cloud connectivity (Vitolins and Smith, 2022; Odhiambo et al., 2021). These systems enhance adherence through reminders and record-keeping, yet typically lack verification mechanisms beyond logging access events (Odhiambo et al., 2021; Smith and Clark, 2021). RFID-based medication management has been explored for safety and automation. The RMAIS prototype integrates an RFID reader, scale, and rotating dispenser for scheduled medicine presentation (McCall et al., 2010), while portable smart pillboxes demonstrate adherence improvements through tagged containers (Doe and Roe, 2024). Commercial solutions like ScripTalk provide audio prescription information to visually impaired users but require centralized stations and are not self-contained edge systems.

Computer vision approaches include deep-learning pill recognition with imprint detection (Heo et al., 2023), YOLO-based mobile applications with audio feedback (Dang et al., 2024), and graph-based multimodal recognition for natural scenes (Nguyen et al., 2023). However, vision-assisted solutions depend on cloud-based inference, raising privacy and latency concerns. However, most systems operate in isolation, RFID-based systems lack content verification, while vision-only approaches may misidentify pills. Some works combine modalities through ingestible RFID sensors and federated learning frameworks (Cheung and Lee, 2024), yet integration of RFID with on-device vision and audio feedback for visually impaired users remains uncommon.

Edge AI systems demonstrate feasibility for privacy-sensitive assistive devices, with successful deployments on Raspberry Pi platforms for navigation and object recognition (Mahendran et al., 2021; Baig et al., 2024; Wong et al., 2025). However, medication verification combining multimodal sensing has not been addressed.Multilingual transformers with retrieval-augmented generation show effectiveness for low-resource languages (Das et al., 2025a), though not yet applied to assistive medication systems. Future extensibility includes enhanced security through modified RSA algorithms for RFID data protection (Das et al., 2019) and multilingual neural machine translation for global deployment in diverse linguistic communities (Bala Das et al., 2023). While prior work established RFID-based dispensing, vision-driven recognition, and edge-deployed assistive systems, a critical gap remains: dual-modality (RFID + AI vision + audio) medication verification running entirely on device. Our device addresses this gap by integrating RFID tag reading, on-device Gemma3 vision analysis, and text-to-speech feedback within a portable Raspberry Pi platform for visually impaired medication management.

3 System Architecture

DRISHTI delivers real-time medication verification through multimodal sensing and on-device AI processing within a compact Raspberry Pi platform. Operating entirely offline, the system integrates RFID scanning, vision processing, and audio feedback for privacy-preserving medication safety.

3.1 System Components

The system integrates a Raspberry Pi 4 Model B (8GB RAM) with three input modalities: MFRC522 RFID reader (13.56 MHz) for contactless scanning, Pi Camera Module v3 for prescription capture, and Bluetooth/WiFi for wireless synchronization. RFID tags encode compressed medication data using a concise CSV format that embeds seven essential attributes (`med_id, dosage_schedule, form_code, expiry_date, strength, brand_name, generic_name`) within a single line. This encoding addresses the 52-byte storage limitation of standard RFID tags while achieving approximately 75% data compression compared to conventional JSON representations, with lexical pattern analysis.

The accessibility interface provides *pyttsx3* text-to-speech, tactile controls, and GPIO LEDs/buzzers. The core *MedicationClassifier* employs hierarchical validation using confidence-weighted scoring: exact matching (100%), generic equivalence (95%), therapeutic substitution (90%), and fuzzy similarity

(85%). A SQLite3 *PharmaceuticalDatabase* manages 248,000 drug entries with brand-to-generic mappings, therapeutic networks spanning thirteen drug classes, and allergy matrices. The Gemma3 4B quantized vision model processes prescription documents with 4-bit quantization for real-time edge performance, supported by *SimpleMFRC522*, *watchdog*, and *pandas*/*NumPy* libraries.

3.2 Integration and Workflow

The system supports dual-mode prescription acquisition via camera-based digitization and wireless entry (Bluetooth/Wi-Fi). Captured images are processed by the vision engine to extract structured medication data, while external devices can transmit prescriptions directly. All inputs are standardized into a unified JSON format including patient demographics, regimens, allergy profiles, and physician details. Input pathways RFID scanning, camera capture, and wireless input converge at the multi-criteria classification engine for validation against prescription profiles. The system performs prescription matching ($\tau = 85\%$), timing checks, dosage comparison ($\delta = 0.1$ mg), and safety screening with prioritized decision trees. Context-aware audio feedback is generated through a dynamic text-to-speech module, which varies tone, speed, and urgency by safety category (safe: 160 WPM, warnings: 150 WPM, urgent: 140 WPM), embedding drug name, strength, and schedule. This ensures patient-specific feedback calm confirmations, cautious warnings, urgent alerts rather than generic templates, improving clarity and trust. Processing is fully local with sub-millisecond response times, and all interactions are logged with timestamps and confidence metrics.

3.3 Dataset Description

The pharmaceutical knowledge base was built from a Kaggle dataset (Singh, 2023) containing over 248,000 medicines with usage, side effects, and substitutes. It provides structured attributes such as brand and generic names, therapeutic classes, dosage details, and equivalent substitutes, enabling construction of the hierarchical drug ontology for brand–generic mappings, therapeutic substitution networks, and allergy cross-reference tables. To fit the resource-constrained edge device, preprocessing removed non-essential text and reduced memory load. Drug names and substitutes were normalized, duplicates eliminated, and therapeutic equivalence pairs extracted for substitution checks. Side-effect and allergy data were converted into structured forms for real-time lookups, supporting efficient management of 248,000 entries on the Pi without compromising accuracy or response time.

4 Methodology

4.1 Design Framework

The system architecture integrates four principal modules: RFID-based medication identification, prescription data acquisition via optical character recognition and manual entry, an intelligent classification engine incorporating therapeutic-equivalence matching, and an accessibility-centric multimodal feedback generator to ensure end-to-end verification and user-friendly interaction tailored for visually impaired users.

The system is deployed on a compact, edge-computing platform that integrates essential hardware to enable robust, real-time medication verification. Key components include an MFRC522 RFID scanner operating at 13.56 MHz for tag detection, a camera module for digitizing prescription documents, and an onboard audio subsystem delivering adaptive text-to-speech prompts. A three-button tactile interface with raised indicators facilitates non-visual navigation, while Bluetooth connectivity supports wireless synchronization of prescription data. By combining RFID, camera, and Bluetooth inputs into a unified tri-modal architecture managed entirely by the local processor. The system provides redundant, flexible pathways for accurate verification tailored to visually impaired users.

The complete system workflow, depicted in Figure 1, delineates a tri-modal input architecture comprising RFID-based wave sensing, camera-driven image scanning, and manual wireless input, whose data streams are consolidated by a portable edge-computing device to generate real-time audio feedback for users with visual impairments.

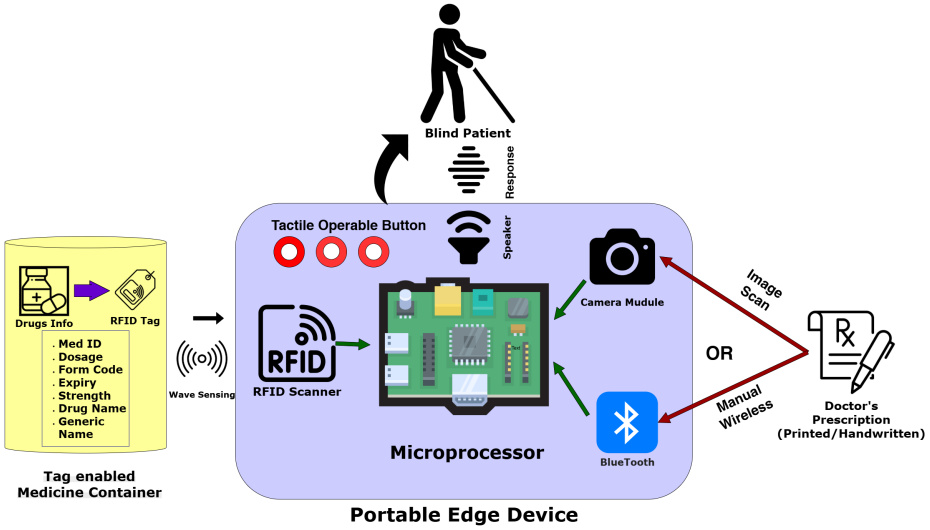

Figure 1: Workflow of DRISHTI, showing tri-modal medication verification on the edge-device using RFID scanning, camera-based prescription analysis, wireless input, and audio output capability

4.2 Multi-Source Drug Database and Intelligent Drug Matching

The system incorporates a comprehensive pharmaceutical knowledge base of 248,218 drug entries with hierarchical relationships including brand-to-generic mappings (25 equivalence pairs), therapeutic substitution networks spanning 13 major drug classes (ACE inhibitors, proton pump inhibitors, statins), and contraindication matrices for allergy cross-referencing and drug interactions.

The medication-matching framework employs confidence-weighted scoring through cascading similarity metrics: exact string matching (100% confidence), generic equivalence matching (95% confidence), therapeutic substitution (90% confidence), and fuzzy string matching (85% confidence), ultimately returning the highest-confidence match for prescription validation.

4.3 Multi-Criteria Safety Classification

DRISHTI uses a hierarchical, multi-stage validation pipeline to ensure robust, safe medication verification, as in Table 1. The five stages address prescription accuracy, dosage correctness, timing adherence, formulation compatibility, and overall safety. The system begins with *prescription verification*, cross-referencing the scanned medication against active prescriptions using exact/fuzzy string matching with a confidence threshold of $\tau = 85\%$. *Temporal validation* is then performed, checking for medication intake within flexible windows: morning (05:00–14:00), afternoon (14:00–20:00), and evening (20:00–05:00).

Subsequently, *dosage accuracy verification* ensures prescribed and scanned strengths within a precision tolerance of $\delta = 0.1$ mg. *Formulation compatibility* evaluates acceptability across alternative forms using a standardized taxonomy, and *safety screening* checks for expiration and allergies, referencing grouped categories such as penicillin, sulfa, or cephalosporin families. DRISHTI's deterministic, hierarchically prioritized decision tree classifies medications as *NOT_PRESCRIBED*, *EXPIRED*, *WRONG_STRENGTH*, *WRONG_TIMING*, or *CORRECT*. Dangerous cases trigger urgent alerts, while non-critical timing deviations prompt guidance. An accessibility-first, multimodal interface delivers real-time, context-aware feedback via offline text-to-speech and tactile controls, ensuring intuitive verification and safety across user abilities.The decision trees in DRISHTI are hand-crafted and rule-based, rather than learned from training data. This design choice was made to ensure transparency, interpretability, and auditability, which are essential in safety-critical applications. Each branch directly corresponds to medically relevant checks—such as prescription match, dosage tolerance, expiry validation, or allergy screening—ensuring predictable behavior under all conditions. While machine-

learned classifiers could capture more subtle patterns, the deterministic approach minimizes false negatives and enables regulatory compliance through explainable rules.

4.4 Real-Time Performance

The system adopts an asynchronous, event-driven architecture that ensures continuous RFID monitoring with average response latency < 1.0 millisecond. Atomic file operations and write-ahead logging guarantee data consistency and thread safety during concurrent prescription updates. The unified tri-modal input architecture coordinates RFID detection (for immediate parsing of CSV/JSON medication data), real-time vision capture (triggering LLM-driven OCR for prescription labels), and maintains updated medication profiles using a local SQLite3 database, enabling reliable offline operation in resource-constrained settings.

Thread-safe mechanisms, including the *PrescriptionFileHandler* (with watchdog-based monitoring), prevent race conditions by automatically reloading modified prescription files. This maintains current, accurate medication data and detailed local logs with precise timestamps, ensuring robust state management, rapid verification, and immediate user feedback in high-frequency medication scenarios. The system provides context-aware audio feedback optimized for visually impaired users through offline text-to-speech processing. Feedback messages use differentiated tones and speech rates to convey urgency levels, with graded responses detailed in Table 2 enabling users to distinguish between safe confirmations, cautionary guidance, and critical alerts.

5 Evaluation and Results

5.1 Experimental Setup

We developed a systematic evaluation framework consisting of 149 test scenarios, which were divided into seven medication safety categories. Table 3 summarizes the distribution of these scenarios, including their respective counts and percentages. The dataset includes a balanced mix of correct and incorrect medication use cases, such as perfect matches, valid substitutes, wrong timing, dosage mismatches, form mismatches, expired medications, and drug interactions. This categorization ensured that the evaluation comprehensively addressed both routine and safety-critical situations.

The evaluation framework relied on expert-annotated ground truth labels, where each test scenario was classified as either *SAFE* or *DANGEROUS* and assigned an associated confidence score. These labels served as the reference standard for performance assessment. System predictions were compared against the expert classifications using a multi-method correctness determination approach. In cases of ambiguity, a safety-first fallback mechanism was applied to prioritize conservative decisions, ensuring that potentially dangerous scenarios were never misclassified as safe.

5.2 Performance Results

Our comprehensive evaluation demonstrates robust performance across all tested scenarios, as illustrated in Figure 6. The system achieved an overall accuracy of 86.57% when evaluated across 149 test scenarios, highlighting its effectiveness in verifying medication safety. Performance analysis revealed that scenarios classified as safe were correctly identified with an accuracy of 77.14%, while all safety-critical scenarios were detected with perfect accuracy (100.0%). For cases involving timing or minor safety issues, the system achieved an accuracy of 80.0%, reflecting its ability to provide appropriate warnings in non-critical situations. The mean response time for a complete verification cycle was measured at approximately 1.0 ms, confirming the system's suitability for real-time operation. Detailed per-classification performance metrics are provided in Figure 2, and the distribution of the test scenarios is summarized in Figure 3.

5.2.1 Per-Classification Performance

Performance varied across different classification types, with the highest accuracy observed for safety-critical categories as shown in Table 4 and (Figure 2). The system achieved perfect accuracy (100.0%) for both *CORRECT* classifications (58/58) and *WRONG_TIMING* cases (28/28), indicating reliable detection of properly prescribed medications and correct identification of timing-related deviations. Similarly, detection of expired medications achieved 100.0% accuracy (13/13), while strength mismatches were identified with an accuracy of

Stage	Validation Step	Description
1	Prescription Verification	Exact/therapeutic/fuzzy matching; confidence $\tau = 85\%$
2	Temporal Validation	Morning: 05:00–14:00; Afternoon: 14:00–20:00; Evening: 20:00–05:00
3	Dosage Accuracy	Compare prescribed vs. actual strength; tolerance $\delta = 0.1$ mg
4	Formulation Compatibility	Validate acceptable substitutes (e.g., tablet $\leftrightarrow$ capsule)
5	Safety Screening	Check expiry date and patient-specific allergy conflicts

Table 1: Five-stage hierarchical validation framework for medication safety.

Alert Type	Tone / Speed	Example Message
Safe Confirmation	Calm / 160 WPM	"This is your Lisinopril 10mg. Appropriate timing for morning dose. Safe to take."
Warning Alert	Cautious / 150 WPM	"This is your correct medication, but it's 2 hours early. Next dose recommended at 8 PM."
Danger Alert	Urgent / 140 WPM	"STOP! Wrong strength detected. You have 20mg but prescribed 10mg. Do not take."

Table 2: Examples of adaptive audio responses with tone and speed variations.

Scenario Category	Count (n)	Percentage (%)
Perfect Match Scenarios	39	26.2
Valid Substitute Scenarios	30	20.1
Wrong Timing Scenarios	25	16.8
Dosage Mismatch Scenarios	20	13.4
Form Mismatch Scenarios	12	8.1
Expired Medication Scenarios	15	10.1
Drug Interaction Scenarios	8	5.4

Table 3: Distribution of the 149 test scenarios across seven medication safety categories.

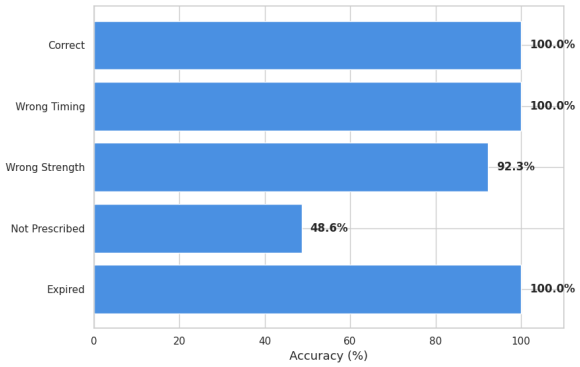

Figure 2: Classification-Specific performance with accuracy breakdown across different classification categories.

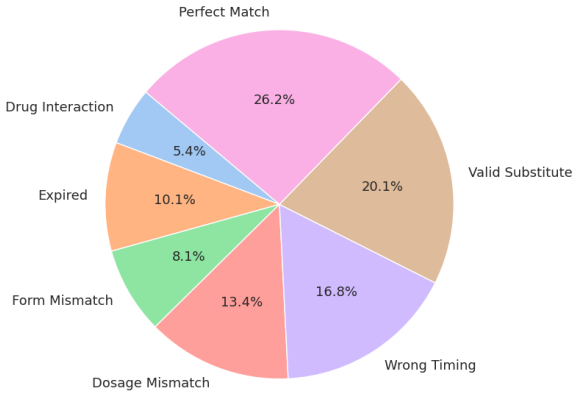

Figure 3: Test Scenario Distribution showing balanced distribution across 149 test scenarios

Classification Type	Accuracy (%)	Cor./Tot.
CORRECT Classification	100.0	58/58
WRONG_TIMING Classification	100.0	28/28
EXPIRED Detection	100.0	13/13
WRONG_STRENGTH Detection	92.3	12/13
NOT_PRESCRIBED Detection	48.6	18/37

Table 4: Classification-specific performance results across all test scenarios.

92.3% (12/13).

The lowest performance occurred for *NOT_PRESCRIBED* cases: 48.6% correct (18/37). This outcome is intentional—DRISHTI employs a conservative policy that flags any unrecognized medication for manual verification to avoid false-safe classifications. Errors primarily stemmed from (i) brand–generic mismatches (generic in the prescription vs. branded RFID not in the database) and (ii) regional formulations missing from the dataset. Mitigation will include expanded brand–generic normalization and incorporation of regional drug vocabularies.

5.2.2 Response Time and Confidence Analysis

Table 5 summarizes the distribution of response times for all verification cycles. The system demonstrates exceptional processing efficiency, achieving a mean response time of approximately 1.00 ms. Notably, 98.0% of all verification cycles are completed in under 0.5 ms, while only 0.67% require between 0.5 and 1.0 ms, and 1.33% exceed 1.0 ms. These results confirm that DRISHTI delivers ultra-fast, real-time performance with no perceptible delay during user interaction, a critical factor for assistive devices deployed on edge platforms.

Response Time (ms)	Scenarios (n)	Percentage (%)
0 – 0.5	146	98.0
0.5 – 1.0	1	0.67
> 1.0	2	1.33
Mean Response Time	–	**1.00 ms**

Table 5: Distribution of response times for all verification cycles.

The confidence-accuracy correlation analysis further highlights the robustness of the classification engine. High-confidence predictions (95–100%) for *CORRECT* scenarios consistently achieve 100% accuracy, while *NOT_PRESCRIBED* classifications reach 48.6% accuracy at similar confidence levels. For safety-critical cases, the system achieves 100% accuracy for expired medications when predictions are made at full confidence, and 92.3% accuracy for wrong-strength detections when predictions are made with 90% confidence. These findings demonstrate that the classifier's confidence scores reliably reflect prediction accuracy, allowing the system to adopt a safety-first strategy by flagging uncertain cases for user verification rather than risking false safe classifications.

5.2.3 Vision Model Performance

The Gemma3 4B 4-bit quantized multimodal model achieves 78.4% overall prescription document analysis accuracy suitable for real-time applications. Printed prescriptions significantly outperform handwritten documents (85.2% vs 67.8% accuracy), with electronic prescriptions achieving the highest accuracy at 92.1% as shown in Figure 4. Document quality directly impacts performance, ranging from 92.1% for electronic documents to 58.9% for poor handwritten prescriptions. Information extraction accuracy varies by data type: medication names (88.5%), dosage (82.1%), schedule/frequency (76.3%), and special instructions (71.8%).

5.3 Safety Performance Analysis

As summarized in Table 7 and illustrated in Figure 6, DRISHTI meets its safety-first objective with 100% detection across all dangerous scenarios and zero false negatives; expired medications were detected at 100% (13/13) and incorrect strength at 92.3% (12/13). In operational tasks, correct medication identification achieved 100% (58/58) and timing validation reached 100% (28/28), demonstrating robust day-to-day reliability. Performance on *NOT_PRESCRIBED* cases shows 48.6% accuracy by design, reflecting a conservative policy that flags uncertain or unrecognized medications for manual review to avoid false-safe outcomes. The safety-first treatment of unknowns is visualized in Figure 5, underscoring that no unsafe instance is misclassified as safe.

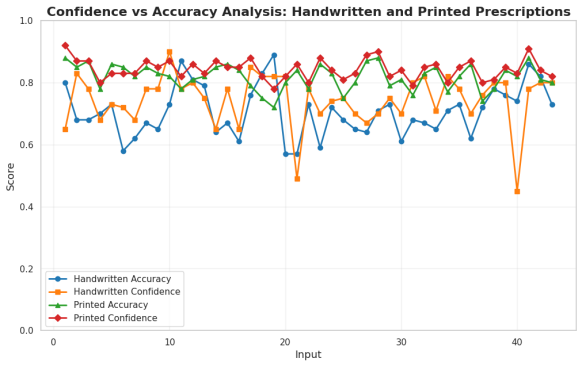

Figure 4: Handwritten vs printed prescription performance analysis

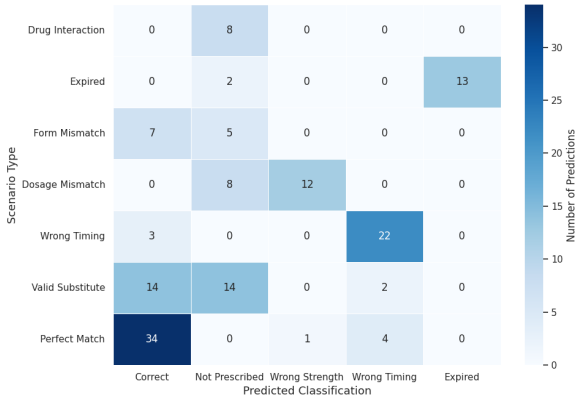

Figure 5: Confusion Matrix Analysis

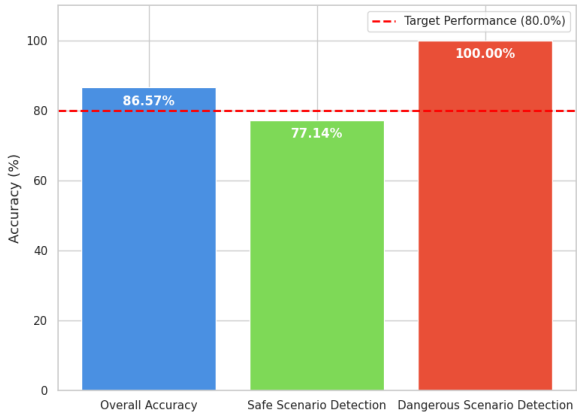

Figure 6: System Performance data

System	Input Types	AI Processing	Awareness	Intelligent Matching	Smart Feedback
ScripTalk (En-Vision America, 2024)	Audio only	None	Limited	Basic	Static
YOLO-OCR (Dang et al., 2024)	Vision only	Basic	None	Simple	Static
RMAIS (McCall et al., 2010)	RFID only	None	None	Direct	None
Smart Pillboxes (Doe and Roe, 2024)	RFID only	None	Limited	Direct	Basic
DRISHTI	**Tri-modal**	**Advanced**	**Full**	**Multi-level**	**Adaptive**

Table 6: Intelligence and AI capability comparison across medication assistance systems

Safety Level	Correct/Total	Accuracy (%)
Safe Scenarios	54/70	77.1
Dangerous Scenarios	43/43	100.0

Table 7: Performance across safety-critical and operational categories.

5.4 Comparison with Existing Solutions

To contextualize DRISHTI's capabilities, Table 6 contrasts DRISHTI with traditional medication aids that focus on pill identification or static audio (e.g., YOLO-OCR imprint reading, ScripTalk label playback, RFID-only adherence logs). DRISHTI delivers broader, intelligent management: 86.57% overall accuracy with 100% detection of safety-critical cases (expired drugs, dosage mismatches, interaction risks). Running fully offline on edge hardware preserves privacy, usability, and reliability. Moving beyond lookup, DRISHTI enables context-aware decisions via the Gemma3 4B model, multi-level matching (exact, generic, therapeutic, fuzzy), patient history and timing awareness, and adaptive urgency-based audio feedback. This comprehensive AI integration positions DRISHTI as a first, safety-first, truly intelligent assistive medication system.

6 Conclusion

DRISHTI is an assistive system that enhances medication safety for visually impaired users by integrating RFID identification, AI-driven visual recognition, and real-time audio interaction. Running entirely offline on low-cost edge hardware, it ensures multimodal verification with strong privacy and no cloud dependency. Evaluation across 149 scenarios shows 86.57% overall accuracy and 100% detection of safety-critical events (expired drugs, dosage mismatches, interaction risks), supporting home and institutional use. Real-time performance (<1 ms) with urgency-aware feedback enables daily integration, while the tri-modal architecture ensures fault tolerance and autonomy through voice prompts.

Future work targets multilingual scalability by integrating OCR for non-English scripts (Indic, Bangla, Arabic), expanding brand–generic mappings, and adopting multilingual text-to-speech, alongside fine-tuning vision models for diverse scripts. Prior works on Multilingual Neural Machine Translation (MNMT) for Indic-to-Indic languages (Bala Das et al., 2024) provide a foundation, while DRISHTI-Plus may leverage MNMT for multilingual dialogue and audio description (Bala Das et al., 2023). Integration with secure mobile/cloud dashboards could enhance monitoring with federated or edge-assisted learning approaches(Paul et al., 2025). To extend device capability, error analysis of language translations using the MQM framework (Das et al., 2025b) is included. Collectively, DRISHTI demonstrates real-world readiness and a clear pathway toward accessible, intelligent, and inclusive medication management for underserved populations.

7 Ethics Statement and Limitations

DRISHTI is designed to run fully offline, ensuring user control of sensitive data. . A conservative confidence threshold minimizes safety risks, and drug information comes from public, anonymized sources to reduce bias, although cultural and linguistic diversity remain challenges. The system is meant to assist, not replace, professional medical care. DRISHTI performs better on printed than handwritten prescriptions and is currently limited to English and Western pharmaceutical data, restricting usability in multilingual regions. Conservative detection of non-prescribed drugs increases false alerts, and hardware limitations prevent real-time updates. Clinical validation is pending, and the audio-tactile interface is insufficient for users with multiple impairments.

References

Mirza Samad Ahmed Baig, Syeda Anshrah Gillani, Shahid Munir Shah, Mahmoud Aljawarneh, Abdul Akbar Khan, and Muhammad Hamzah Siddiqui. 2024. Ai-based wearable vision assistance system for the visually impaired: Integrating real-time object recognition and contextual understanding using large vision-language models.

Sudhansu Bala Das, Atharv Biradar, Tapas Kumar Mishra, and Bidyut Kr. Patra. 2023. Improving multilingual neural machine translation system for indic languages. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 22(6).

Sudhansu Bala Das, Divyajyoti Panda, Tapas Kumar Mishra, Bidyut Kr. Patra, and Asif Ekbal. 2024. Multilingual neural machine translation for indic to indic languages. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 23(5):1–32.

Rupert R. A. Bourne, Jaimie Adelson, Seth Flaxman, Paul Briant, Michele Bottone, Theo Vos, Kovin Naidoo, Tasanee Braithwaite, Maria Cicinelli, Jost Jonas, Hans Limburg, Serge Resnikoff, Alex Silvester, Vinay Nangia, and Hugh R. Taylor. 2020. Global prevalence of blindness and distance and near vision impairment in 2020: progress towards the vision 2020 targets and what the future holds. *Investigative Ophthalmology & Visual Science*, 61(7):2317–2317.

A. Cheung and B. Lee. 2024. Rfid based health adherence medicine case using fair federated learning. *arXiv*.

Bo Dang, Wenchao Zhao, Yufeng Li, Danqing Ma, Qixuan Yu, and Elly Yijun Zhu. 2024. Real-time pill identification for the visually impaired using deep learning.

Sajeeb Das, Srijit Paul, Akib Jayed Islam, Sultanus Salehin, and Prasun Datta. 2025a. Development of a multilingual climate fact-checking system with unified dataset for low-resource indic languages. In *Proceedings of the 16th International IEEE Conference on Computing, Communication and Networking Technologies (ICCCNT)*, IIT Indore, Madhya Pradesh, India. IEEE Electronics Packaging Society and All India Council for Technical Education (AICTE). Paper ID: 7300.

Sudhansu Bala Das, Samujjal Choudhury, Tapas Kumar Mishra, and Bidyut Kr Patra. 2025b. Investigating the effect of backtranslation for indic languages. In *Proceedings of the First Workshop on Natural Language Processing for Indo-Aryan and Dravidian Languages*, pages 152–165.

Sudhansu Bala Das, Sugyan Kumar Mishra, and Anup Kumar Sahu. 2019. A new modified version of standard rsa cryptography algorithm. In *Smart Computing Paradigms: New Progresses and Challenges: Proceedings of ICACNI 2018, Volume 2*, pages 281–287. Springer.

J. Doe and A. Roe. 2024. Development of a smart pillbox with passive rfid to support medication adherence. *CMS Conferences*.

En-Vision America. 2024. Scriptalk station: Audio prescription reader for the visually impaired. `https://www.envisionamerica.com/scripttalk`. Accessed: 2024-12-01.

Shibakali Gupta, Sumana Das, and Surabhi Pal. 2023. Medication identification and assistive system for the visually impaired: Vismed. *International Journal of Computer Applications*, 186(38):30–33.

J. Heo, Y. Kang, S. Lee, D. Jeong, and K. Kim. 2023. An accurate deep learning–based system for automatic pill identification. *JMIR*.

Jagadish K Mahendran, Daniel T Barry, Anita K Nivedha, and Suchendra M Bhandarkar. 2021. Computer vision-based assistance system for the visually impaired using mobile edge artificial intelligence. In *CVPR Workshops (MAI)*.

C. McCall, B. Maynes, C. Z. Zou, and N. J. Zhang. 2010. Rmais: Rfid-based medication adherence intelligence system. In *EMBC*.

Vidula V Meshram, Kailas R Patil, Vishal A Meshram, and Shripad Bhatlawande. 2021. Smartmedbox: A smart medicine box for visually impaired people using iot and computer vision. *RIA*.

Anh Duy Nguyen, Huy Hieu Pham, et al. 2023. High accurate and explainable multi-pill detection framework with graph neural network. *arXiv*.

C. Odhiambo, P. Wright, C. Corbett, and H. Valafar. 2021. Medsensor: Medication adherence monitoring using neural networks on smartwatch accelerometer data. *arXiv*.

Srijit Paul, Sajeeb Das, Akib Jayed Islam, Sultanus Salehin, and Prasun Datta. 2025. Federated learning for privacy-preserving gait recognition on edge devices using imu data. In *Proceedings of the 16th International IEEE Conference on Computing, Communication and Networking Technologies (ICCCNT)*, IIT Indore, Madhya Pradesh, India. IEEE Electronics Packaging Society and All India Council for Technical Education (AICTE).

Shudhanshu Singh. 2023. 250k medicines – usage, side effects and substitutes. `https://www.kaggle.com/datasets/shudhanshusingh/250k-medicines-usage-side-effects-and-substitutes`. Accessed: 2025-07-25.

J. A. Smith and R. Clark. 2021. A review on emerging smart technological innovations in healthcare: Medication adherence. *Smart Health Review*.

M. Z. Vitolins and C. P. Smith. 2022. Technologies for medication adherence monitoring and technology assessment. *JMIR mHealth and uHealth*, 10(3):e35157.

X. Wong et al. 2025. A hybrid object detection approach for visually impaired persons using attention-driven deep learning. *Scientific Reports*.

Fateme Zare, Paniz Sedighi, and Mehdi Delrobaei. 2023. A wearable rfid-based navigation system for the visually impaired. *arXiv*.