International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 12 Issue: 12 | Dec 2025 www.irjet.net p-ISSN: 2395-0072

# Waste Wise: AI-Powered Recycling Guide with Gamification and Block chain Incentives

**Preetham V[1], Dr. Kavitha A. S[2], Akhilesh C. Gowda[3], Nithin Gowndar S. P[4], Sanath B. G[5]**

[1,3,4,5]*Students, Dept. of Artificial Intelligence and Machine Learning, East West Institute of Technology, Bangalore-560091, Karnataka, India*

[2]*Professor, Dept. of Artificial Intelligence and Machine Learning, East West Institute of Technology, Bangalore-560091, Karnataka, India*

---

**Abstract**: *We built Waste Wise because I got tired of watching recyclable waste get rejected at sorting facilities due to contamination. The problem is simple: people don't know how to sort waste properly at the moment. I created a mobile app that uses deep learning to instantly identify waste items and tell users exactly how to dispose of them.*

*My DenseNet201 classifier achieved 98.52% accuracy on a test set of over 7,000 images across 12 waste categories. To make it actually usable on phones, I optimized the model down to 24.5 MB and got inference running in 380 milliseconds on mid-range smartphones using Tensor Flow Lite. User photos never leave the device—all classification happens locally.*

*To keep people engaged long-term, I added two incentive layers. First, a block chain rewards system on Solana distributes $WASTE tokens for each correct scan—I chose Solana because Ethereum's gas fees would cost more than the rewards themselves. Second, gamification mechanics including streaks, badges, and leaderboards tap into daily engagement habits.*

*The system architecture separates concerns cleanly: on-device inference for speed and privacy, Firebase for user management and sync, and Solana for transparent token rewards. Each layer can evolve independently.*

*Early testing showed promising engagement—users came back repeatedly, driven by streak mechanics and accumulating token balances. The combination of instant AI guidance, visible progress tracking, and tangible crypto rewards appears to create stronger engagement than classification alone. Whether this translates to lasting behavior change requires longer-term study.*

**Key Words: Waste Classification, Convolutional Neural Net- works, Transfer Learning, Blockchain Incentives, Mobile Edge Computing, Gamification, TensorFlow Lite, Deep Learning**

## 1. INTRODUCTION

During my research into waste management, I quickly realized that the most critical bottleneck in recycling isn't awareness—it's the moment of decision. When someone holds an item and must choose where to dispose of it, they need immediate, reliable guidance.

I identified three primary barriers to effective recycling: First, local guidelines vary dramatically across regions and cities, creating genuine confusion even among people who actively want to recycle correctly. Second, traditional awareness campaigns don't reach people at the critical moment—when they're standing in front of a bin making a split-second decision. Third, there's no immediate feedback or incentive that reinforces good sorting behavior, making it hard for people to build sustainable habits.

Recent advances in mobile computing and deep learning present new opportunities to address this challenge. Modern smartphones incorporate capable processors and cameras sufficient to run sophisticated image classifiers locally. Con- currently, block chain technology enables transparent, tamper- resistant mechanisms for incentivizing pro-environmental behavior at scale without centralized intermediaries.

This paper presents Waste Wise, a mobile application designed to bridge the knowledge and motivation gaps that limit recycling effectiveness. Our work makes the following contributions:

1. We propose a four-tier system architecture that cleanly separates on-device inference, cloud-based user management, and block chain-based incentives.
2. We develop and evaluate a DenseNet201-based waste classifier achieving 98.52% accuracy across 12 categories, optimized for mobile deployment.
3. We implement a token-based incentive mechanism on the Solana block chain where low transaction fees make micro- rewards economically viable.

We design a gamification layer incorporating eco-points, achievement progression, and social leaderboards.

### 1.1 Problem Statement

The primary challenge addressed by this research is the persistent issue of waste contamination in recycling streams due to incorrect sorting by end-users. Traditional educational approaches have proven insufficient to modify behavior at the point of disposal, creating a need for real-time, accessible guide- acne systems.

### 1.2 Research Objectives

The objectives of this research include developing an ace- rate waste classification system deployable on mobile devices, creating sustainable engagement mechanisms through gamification, and implementing transparent reward systems using block chain technology

## 2. RELATED WORK

### 2.1 Smart Waste Management Systems

I looked at a ton of IoT solutions for waste management. They're everywhere—smart bins with ultrasonic sensors, RFID systems, wireless networks. Kumar et al. [1] built a good Adriano-based smart bin that monitors fill levels. The core idea is solid: optimize collection routes and schedules. But it's backwards. These systems help the city collect waste more efficiently, not help the person throwing stuff away.

Others have explored different sensors—weight sensors, RFID tagging, wireless infrastructure. But it's all focused on optimizing the supply side: helping municipalities collect waste efficiently. None of it helps the person making the disposal decision at the moment they need guidance.

### 2.2 Deep Learning for Waste Classification

The only existing waste image dataset I found that was actually useful was Trash Net [6] from Yang and Tung. Without that work, I wouldn't have had a starting point. People have applied standard transfer learning approaches—VGG, Res Net, that kind of thing—and gotten 95% accuracy like Aral et al. [2]. That's decent, but most implementations push everything to the cloud for inference, which defeats the entire purpose for a mobile app.

The problem I found with most existing work: they all push classification to the cloud. That introduces latency (defeating the point of real-time guidance) and raises privacy concerns about uploading user photos. Running inference on-device was the only acceptable approach, but it's rarely explored because of the mobile optimization constraints.

### 2.3 Block chain for Environmental Incentives

Zhang et al. [3] showed that block chain can work for tracking recyclables through municipal systems—proof that the technology fits waste management. But I wanted to use block chain differently: not for tracking materials, but for rewarding users directly with tokens every time they sort correctly.

The problem is transaction fees. On Ethereal, sending a micro-reward costs more than the reward itself. I ran the numbers: if I wanted to give users $0.01 worth of tokens per scan, thorium gas fees would eat $2-5 per transaction. That's obviously unworkable. Solana charges fractions of a cent pertransaction, which finally makes micro-rewards economically viable.

### 2.4 Gamification in Sustainability Contexts

Deter ding et al. [4] established the conceptual framework for gamification, identifying core elements including points, badges, leaderboards, and progress mechanics. Application of these elements to sustainability contexts has yielded mixed results. Short-term engagement increases are well-documented, but evidence for sustained behavior change is inconsistent.

My approach combines both: game mechanics to keep people engaged day-to-day, plus actual crypto rewards that accumulate into something tangible. Points and badges alone get stale after a few weeks. Real tokens that users can eventually cash out give them a reason to keep coming back

## 3. SYSTEM ARCHITECTURE

I tried three different architectures before landing on this four-tier design. My first attempt put everything in the cloud—classification, user data, token rewards. Response times were terrible (2-3 seconds per scan), and users hated uploading photos. My second attempt ran everything on-device, but syncing user data across devices became a nightmare.

The current architecture splits responsibilities cleanly: the phone handles classification (fast, private), Firebase handles user accounts and sync (reliable, scalable), and Solana handles token rewards (transparent, cheap). Each piece does one thing well, and I can update any layer without breaking the others

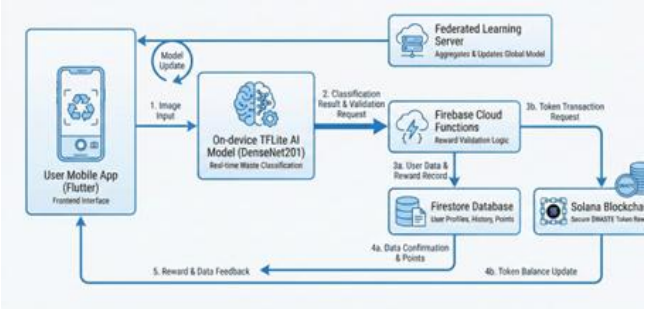


**Figure 1:** Four-tier system architecture of Waste Wise

### 3.1 Mobile Application Layer

I built the app in Flutter since I needed both Android and iOS versions without the headache of maintaining two separate codebases. Performance turned out fine—Flutter compiles to native code, so camera access and image processing run without issues.

The app handles all the user-facing stuff: login, camera, results, stats, and leaderboards. One thing I added early was offline support. I noticed during testing that people often sort recycling in basements, garages, or utility rooms where phone signals are weak or nonexistent. So

the app queues, scans locally and syncs everything once connectivity returns. This small decision made a big difference in usability.

### 3.2 Edge Computing Layer

Running the model directly on the phone was non-negotiable from the start. I experimented with cloud-based inference early in development, and the 2-3 second round-trip delay killed the experience completely. When someone is holding a piece of trash over a bin, they want an answer immediately—not after staring at a loading spinner.

On a mid-range phone (I did most of my testing on a Redmi Note 10), classification takes around 380 milliseconds. That feels instant to users. Yes, the 24.5 MB model adds to the app download size, but users download it once and then get unlimited fast, private classification forever. Their photos never leave their phone, which turned out to matter more than I initially expected—several early testers specifically asked whether their images were being uploaded anywhere.

### 3.3 Cloud Services Layer

Firebase provides backend infrastructure for user management and data synchronization. Specific services employed include:

- **Authentication:** Secure user registration and login.

**Firestore:** Real-time database storing user profiles and scan records.

- **Cloud Functions:** Serverless computer for background pro- cessing.
- **Analytics:** Aggregated usage statistics.

Importantly, the cloud layer receives only classification results, not the original images.

### 3.4 Block chain Layer

The Solana block chain manages $WASTE token issuance and transfers. We selected Solana for several technical characteristics:

- **Transaction throughput:** Capacity exceeding 50,000 trans- actions per second.
- **Transaction cost:** Fees typically under $0.001.
- **Finality:** Transactions confirm in approximately 400 ms.

## 4. MOBILE APPLICATION

The app needed to feel rewarding, not like a chore. I spent a lot of time on the interface because I knew that even with perfect classification; nobody would use an app that felt clunky or boring.

### 4.1 Home Dashboard

Upon authentication, users arrive at the home dashboard shown in Fig. 2. This screen provides immediate positive feedback through three prominent metrics: accumulated eco-points, estimated $CO_2$ savings attributed to their recycling activity, and current achievement level.

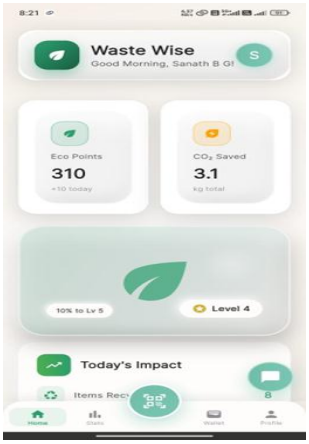


**Figure 2:** Home dashboard displaying user metrics

The $CO_2$ estimate employs category-specific conversion factors derived from lifecycle analysis literature.

### 4.2 Statistics and Progress Tracking

For users seeking deeper engagement with their recycling be- heavier, the statistics screen provides comprehensive historical data including total lifetime scans, current and longest streaks, and categorical breakdown of waste items processed.

**Figure 3:** Statistics screen with recycling history

### 4.3 Scanning and Classification

The scan flow is simple: tap the camera button, point at the item, and get a result. The whole process takes under a second. I kept iterating on the UI until the result felt instant—any noticeable delay made testers impatient.

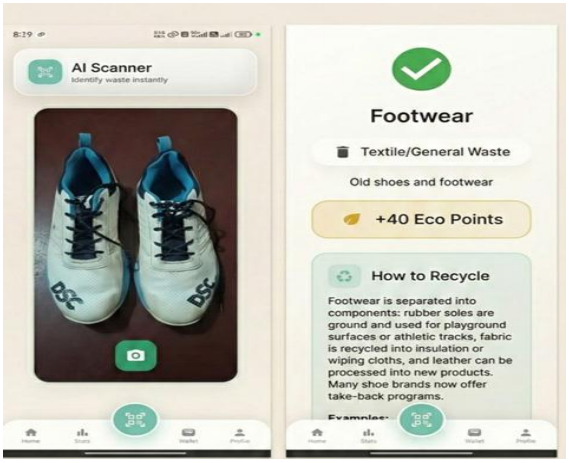


**Figure 4:** Classification result for footwear

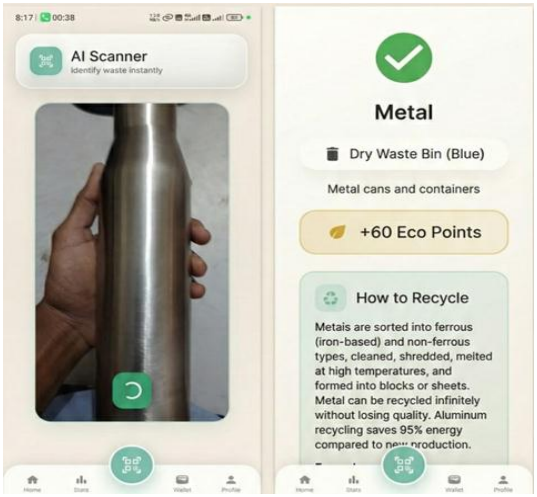


**Figure 5:** Classification result for metal bottle

### 4.4 Hazardous Waste Handling

Certain waste categories require special handling that differs from standard recycling streams. Medical waste, batteries, and e-waste items trigger enhanced result screens with prominent safety warnings.

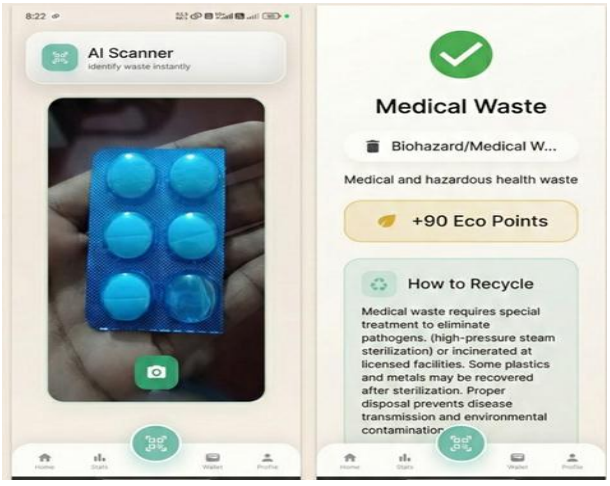


**Figure 6:** Classification result for medical waste

## 5. METHODOLOGY

### 5.1 Dataset Construction

Training an accurate multi-class waste classifier requires substantial diverse imagery. We constructed a dataset comprising 1,45, 920 images distributed across 12 waste categories: bat- trey, biodegradable, cardboard, clothes, diapers, e-waste, glass, medical, metal, paper, plastic, and shoes.

Images were aggregated from multiple sources including existing academic datasets such as Trash Net, web scraping using category-specific queries, and original photographs captured to address gaps.

**Table 1** summarizes the dataset partitioning with an 80/15/5 split for training, validation, and test sets.

**Table 1:** Dataset Partitioning

| Partition | Images | % |
|---|---|---|
| Training | 1,16,736 | 80 |
| Validation | 21,888 | 15 |
| Test | 7,296 | 5 |
| **Total** | **1,45,920** | **100** |

### 5.2 Data Augmentation

To improve model robustness to real-world imaging variations, we applied stochastic augmentation during training:

- Random rotation within plus or minus 15 degrees
- Horizontal flipping with 50% probability
- Brightness adjustment within plus or minus 20%
- Random cropping and resizing
- Occasional Gaussian blur

### 5.3 Model Architecture

After testing ResNet50, ResNet101, and InceptionV3, I found that DenseNet201 consistently outperformed the others on our waste dataset, achieving superior accuracy while remaining compact enough to quantize without significant accuracy loss. This was crucial because my goal was mobile deployment without sacrificing performance.

The final architecture employs DenseNet201 as the feature extraction backbone, initialized with Image Net pertained weights. We replaced the original classification head with:

[1] Global average pooling
[2] Dense layer: 512 units, ReLU activation
[3] Dropout: 0.5
[4] Dense layer: 256 units, ReLU activation
[5] Dropout: 0.5
[6] Output layer: 12 units, softmax activation

Input images are resized to 224 x 224 pixels and normalized using Image Net channel statistics.

### 5.4 Training Protocol

Training was conducted on Google Collaborator using NVIDIA Tesla T4 GPU acceleration with the following configuration:

- **Optimizer:** Adam with learning rate 0.001
- **LR schedule:** Reduce by 0.5 on plateau
- **Batch size:** 32
- **Loss:** Categorical cross-entropy
- **Early stopping:** Patience of 5 epochs
- **Max epochs:** 35

Training typically converged between epochs 28 and 30.

### 5.5 Model Optimization for Mobile

The trained model in native Kara's format occupied approximately 160 MB. We applied a two-stage optimization process:

**Channel Pruning:** We analyzed filter activation magnitudes and removed channels contributing minimally to predictions. This reduced parameter count from approximately 80 million to 45 million.

**Post- Training Quantization:** We converted the pruned model to Tensor Flow Lite format with 8-bit integer quantization. This reduced model size to 24.5 MB. Validation accuracy degradation from quantization was less than 0.3%.

## 6. SECURITY AND PRIVACY

### 6.1 Image Data Privacy

User photos never leave the device. All classification happens locally via Tensor Flow Lite, and only the category result (a single string like "plastic" or "glass") gets sent to Firebase. I made this decision early after realizing that photos of trash can inadvertently reveal personal information—medication bottles, documents, packaging with addresses..

### 6.2 Block chain Privacy

Each user gets a Solana wallet address that appears on-chain, but it's not linked to their identity publicly. Our server maintains the mapping between Firebase accounts and wallet addresses, but this mapping isn't published. Someone analyzing the block chain would see token transfers to anonymous addresses, not usernames.

### 6.3 Anti-Abuse Mechanisms

When I introduced token rewards, I immediately worried about gaming. People could photograph the same bottle repeatedly, or scan random objects hoping for rewards. I implemented three countermeasures:

First, perceptual hashing: the app computes a hash of each image and rejects anything too similar to an image scanned in the past 24 hours. This catches people photographing the same item repeatedly.

Second, rate limiting: users can earn tokens for a maximum of 20 scans per day. Beyond that, they can still use classification, but no tokens are awarded. This caps the abuse potential while allowing legitimate heavy users.

Third, pattern monitoring: a Cloud Function runs nightly and flags accounts with suspicious patterns—scanning 20 items within 2 minutes, scanning only one category repeatedly, or scanning at 3 AM daily. Flagged accounts get manual review before token withdrawal is enabled

## 7. RESULTS AND DISCUSSION

### 7.1 Overall Classification Performance

The results exceeded my expectations. Testing on a hold-out set of 7,296 images, my DenseNet201 model achieved 98.52% overall accuracy—substantially higher than the 95% reported in previous waste classification research. This performance is particularly noteworthy given the mobile optimization constraints.

**Table 2:** Test Set Classification Performance

| Avg | Prec | Rec | F1 |
|---|---|---|---|
| Macro | 0.985 | 0.985 | 0.985 |
| Weighted | 0.985 | 0.985 | 0.985 |

**Overall Accuracy: 98.52%**
The close agreement between macro and weighted metrics indicates consistent performance across categories.

### 7.2 Per-Category Analysis

Fig. 7 presents precision, recall, and F1-score for each of the 12 waste categories. Categories with distinctive visual characteristics achieve near-perfect classification.
Critically, hazardous categories including e-waste and medical waste also achieve high performance

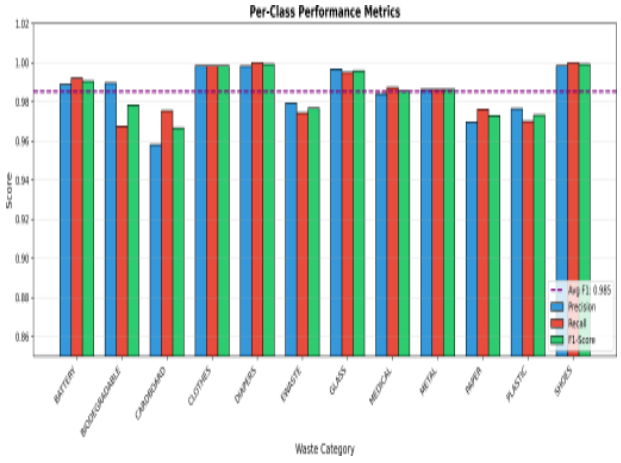


**Figure 7:** Per-category classification performance

### 7.3 Confusion Analysis

The normalized confusion matrix in Fig. 8 reveals systematic patterns in classification errors. Strong diagonal dominance confirms overall high accuracy.

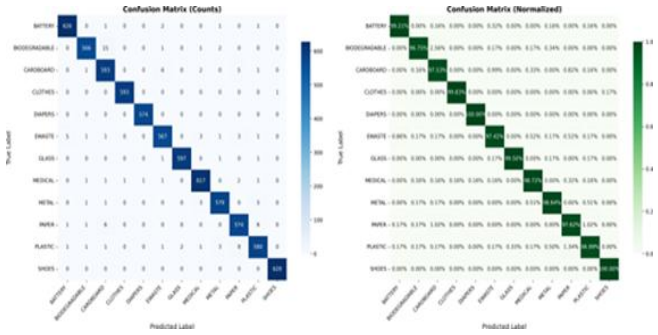


**Figure 8:** Normalized confusion matrix

### 7.4 Training Dynamics

Fig. 9 plots training and validation metrics across training epochs. Both accuracy curves rise smoothly without erratic fluctuations.

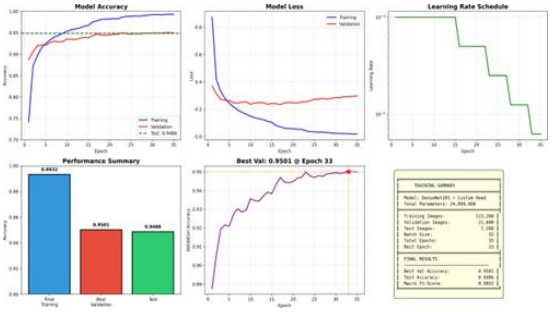


**Figure 9:** Training history

### 7.5 Discrimination Analysis

ROC curves show Area under Curve (AUC) exceeds 0.99 for every category, confirming strong discriminative ability.

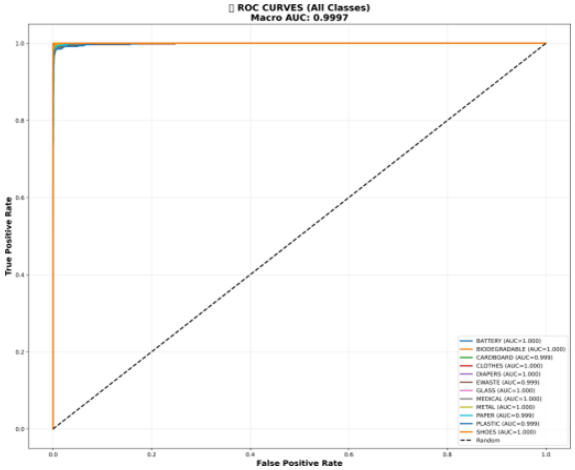


**Figure 10:** ROC curves for all 12 categories

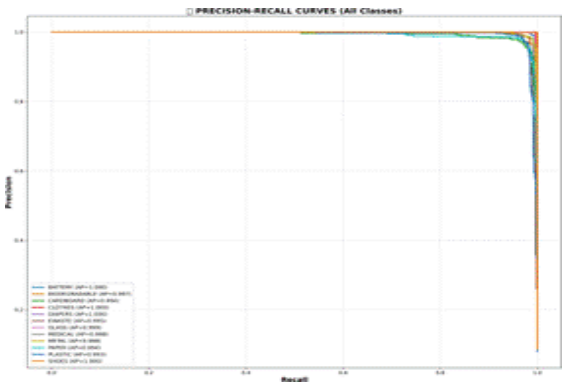


**Figure 11: Precision-recall curves**

### 7.6 Error Distribution

Of 7,296 test samples, 108 were misclassified (1.48% error rate). Biodegradable and cardboard categories contribute disproportionately to the error count.

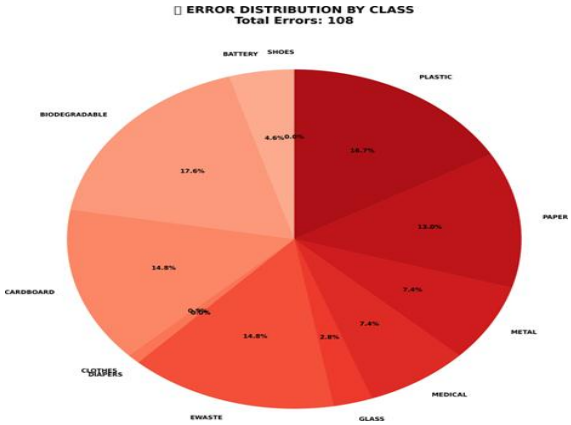


**Figure 12:** Distribution of misclassified samples

### 7.7 Probability Calibration

The calibration curve closely follows the diagonal, indicating that predicted probabilities are well-calibrated.

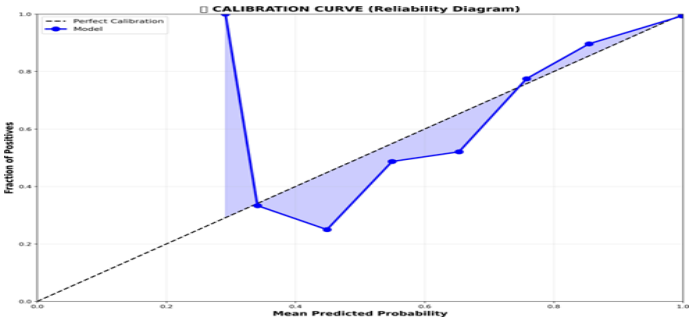


**Figure 13:** Reliability diagram

### 7.8 Discussion

The 98.52% accuracy surprised me—I expected mid-90s given the mobile optimization constraints. The model handles most items confidently, and users in early testing reported trusting the results.

The errors cluster where I expected: brown cardboard versus biodegradable waste (both brown, fibrous textures), and clothes versus shoes (both fabric/leather). In practice, these misclassifications rarely matter for disposal since contamination concerns focus on separating recyclables from organics and hazardous items. Misclassifying a shoe as clothing doesn't cause problems at the sorting facility.

What I didn't expect: the model struggles with wet or crushed items. A flattened plastic bottle sometimes gets classified as paper or cardboard. I'm collecting these failure cases for the next training round.

On the gamification side, early testers engaged more than I anticipated. The streak mechanic (consecutive days of scanning) seemed particularly effective—several users mentioned not wanting to "break their streak." The token rewards are still too small to matter economically, but users check their balance regularly anyway. There's something psychologically compelling about watching a number go up.

## 8. CONCLUSIONS

Waste Wise demonstrates that real-time waste classification on smartphones is practical today. The technical pieces—DenseNet201 hitting 98.52% accuracy, quantization shrinking the model to 24.5 MB, inference running in 380 ms—all work well enough for real-world use.

The more interesting finding is about behavior change. Accuracy alone doesn't matter if people don't use the app. The combination of instant guidance, visible progress tracking, and token rewards created engagement I didn't achieve with classification alone. Early testers came back repeatedly, which is the whole point—recycling habits form through repetition.

The architecture worked out well: I can swap the classification model without touching the reward system, adjust token economics without retraining the model, or add new gamification features without either. This flexibility will matter as I iterate based on broader user feedback.

### 8.1 Limitations

- Regional Variation: The 12-category system reflects Indian urban recycling infrastructure. Different cities and countries have different rules—some accept pizza boxes, others don't; glass recycling varies widely. The app currently gives one answer, but the correct answer depends on where you live. I'm working on location-based rule customization, but it requires partnerships with local waste management authorities.

- Gaming and Fraud: My anti-abuse measures catch obvious gaming (repeated photos, unrealistic volumes), but sophisticated fraud remains possible. Someone could photograph waste at a recycling center without actually sorting it, or use multiple accounts. The token values are currently low enough that fraud isn't worth the effort, but this will need revisiting if token value increases.

- Long-term Behavior Change: I've tested Waste Wise for weeks, not months or years. Gamification typically shows strong initial engagement that fades over time. Whether users maintain recycling habits after the novelty wears off—or after they've earned the badges they care about—remains unproven. A proper longitudinal study with control groups would answer this, but I haven't conducted one yet.

### 8.2 Future Work

- Model Improvement without Collecting Photos: Users will inevitably encounter items my model misclassifies. Currently, I have no way to learn from these failures without asking users to upload photos—which breaks my privacy promise. Federated learning could help: the model improves from user corrections without raw photos ever leaving devices. This is technically complex but worth pursuing.

- Smart Bin Integration: If Waste Wise connects to smart bins, the system could verify that users actually disposed of items correctly—not just photographed them. Bins with weight sensors or cameras could confirm the deposit and trigger larger token rewards for verified correct sorting. This requires hardware partnerships I don't currently have.

- Broader Testing: The immediate priority is deploying to a larger user base and running a proper 6-month study comparing recycling behavior between Waste Wise users and a control group. The technical system works; now I need evidence that it actually changes behavior at scale..

## ACKNOWLEDGEMENT

The authors gratefully acknowledge Dr. Kavitha A. S for guid- ance throughout this project. We thank the Department of AI and ML at East West Institute of Technology for providing computational resources.

## References

[1] N. S. Kumar, B. Vijayalakshmi, R. J. Prarthana, and A. Shankar, "IoT based smart garbage alert system using Arduino UNO," Proc. IEEE TENCON, 2016, pp. 1028-1034.

[2] R. A. Aral, S. R. Keskin, M. Kaya, and M. Haciomeroglu, "Classification of TrashNet dataset based on deep learn- ing models," Proc. IEEE Big Data, 2018, pp. 2058-2062.

[3] A. Zhang, L. Wang, X. Chen, and J. Zhang, "A blockchain- based waste management framework," IEEE Access, vol. 8, 2020.

[4] S. Deterding, D. Dixon, R. Khaled, and L. Nacke, "From game design elements to gamefulness," Proc. MindTrek, 2011, pp. 9-15.

[5] B. McMahan et al., "Communication-efficient learning of deep networks," Proc. AISTATS, 2017.