International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: *https://ijwos.com* E-ISSN : 3049-2424
Volume 3 Issue 8 , August 2026 DOI: https://doi.org/10.71366/ijwos03082629035

AgriCycle AI: An AI-Powered Smart Waste-to-Wealth Platform for Agricultural Waste Identification, Valorization and Marketplace Connectivity

Prof. CHETAN PATIL[1], Vijaylaxmi Hulipalled[2], Sanjana Kawari[3], Anupreksha Patil[4], Sushmita Mudodagi[5]
[1,2,3,4,5] *Department of Artificial Intelligence and Data Science, Angadi Institute of Technology and Management, Belagavi, India.*

Article Info

Article History:

Published: 24 August 2026

Publication Issue:
Volume 3, Issue 8
August-2026

Page Number:
555-563

Corresponding Author:
Prof. CHETAN PATIL

Abstract:
Agricultural production generates large quantities of crop residues, husks, stalks, fruit and vegetable waste, bagasse, manure and other biological by-products, which are frequently burned, dumped or left unmanaged. This paper presents AgriCycle AI, a proposed software-only, intelligent waste-to-value platform that combines computer vision, machine learning, recommendation logic, economic estimation, environmental-impact estimation and digital marketplace functionality to convert agricultural waste from a disposal problem into an economic and environmental opportunity. A farmer uploads an image of agricultural waste, which is validated, preprocessed and classified by an AI model that returns a probable waste category with a confidence score; a recommendation engine then identifies suitable utilization pathways, an indicative value is estimated, and a marketplace module connects the material with potential buyers or processors. The architecture separates the frontend, API/backend, AI inference, recommendation, marketplace, database and analytics layers so each module can be developed and validated independently. This paper reviews related work in agricultural-waste management, computer vision and digital marketplaces, and describes the proposed architecture, methodology, database design and evaluation plan for the system as an academic prototype requiring further dataset collection, model training and field validation before production deployment.

Keywords: Agricultural Waste Management, Artificial Intelligence, Computer Vision, Waste Classification, Waste Valorization, Circular Economy, Recommendation System, Marketplace, Machine Learning

1. INTRODUCTION

Agricultural residues arise at several stages of the value chain, including harvesting, threshing, processing, grading and post-harvest handling. Their physical form ranges from loose leaves and stalks to husks, shells, peels, fibrous material and wet organic residues, and the same broad waste category may require very different handling depending on moisture, cleanliness, particle size, storage duration and the availability of nearby processing facilities. The challenge is therefore not only to identify a material but also to preserve the context required for a useful decision; AgriCycle AI combines an image-based prediction with quantity, location and other contextual information so the output becomes a decision-support record rather than an isolated classification result.

A. Background and Problem Statement

Materials such as rice straw, wheat straw, sugarcane residues, corn stalks, coconut husks, banana residues and fruit and vegetable waste often contain useful organic matter, fibre and energy potential, but their value depends on quality, quantity, season, transportation and local demand. Agricultural waste is frequently managed as a disposal burden because farmers may lack information about what the material can be converted into, how much it may be worth, and who may be interested in buying it. Existing digital tools often address only one dimension of this problem —

identification, valuation or marketplace discovery — rather than connecting them; AgriCycle AI addresses this integration gap directly.

B. Motivation and Proposed Solution

The project is motivated by the idea that waste becomes a resource when information, processing knowledge and market connectivity are available together. AI can reduce the technical barrier to identifying material from images, a recommendation engine can translate a classification result into understandable possibilities, and a marketplace can provide a practical path from information to action. AgriCycle AI is accordingly designed as a web platform in which a farmer uploads an image, provides approximate quantity and location, and receives an AI-based waste category, feasible utilization pathways, an indicative value estimate, and — where a relevant buyer or processor is available — a marketplace listing.

C. Objectives and Scope

- Design an AI-assisted agricultural-waste identification workflow.
- Develop a structured waste-category and utilization knowledge base.
- Design a recommendation engine linking waste categories to feasible uses.
- Provide transparent, indicative economic and environmental estimation.
- Provide a marketplace for connecting waste generators with buyers/processors.
- Define a measurable AI-development and evaluation plan.

The scope is bounded in three ways: the AI model addresses only categories adequately represented in the project dataset; recommendations indicate potentially suitable pathways rather than replacing laboratory testing or professional advice; and the marketplace facilitates discovery and communication while commercial agreements, quality inspection, transportation and payment remain outside the initial academic prototype.

2. LITERATURE REVIEW

Effective residue management generally depends on the material's characteristics and the distance between generation and utilization; collection and storage are especially important for low-density materials because transportation can become a major part of total cost, and time and storage conditions affect the quality of wet organic residues.

A. Agricultural Waste Management and AI in Agriculture

Agricultural-waste management includes collection, storage, treatment, recycling, reuse, recovery and disposal. Because different residues have different moisture, composition, contamination risk and markets, an intelligent platform should retain category and contextual information rather than a generic ‘waste’ label. AI has been widely applied to image recognition, crop monitoring, disease detection and yield analysis in agriculture, demonstrating the value of machine learning where visual or contextual patterns are difficult to process manually; AgriCycle AI applies the same principle to post-harvest residues while extending the workflow to recommendation and market connectivity.

B. Computer Vision for Waste Identification

Computer vision provides a fast first-level identification mechanism through a pipeline of image capture, validation, preprocessing, inference and confidence-based output. Real-time object-detection frameworks such as YOLO [1] make it feasible to process field images with acceptable latency, while classical computer-vision utilities such as OpenCV [5] support the preprocessing stage. Because agricultural scenes may contain mixed material and background clutter, classification and detection formulations should be compared and the approach selected to match the collected dataset.

C. Waste Valorization and Recommendation Systems

Valorization refers to recovering useful material, energy or economic value from residues through pathways such as composting, anaerobic digestion, biomass energy, briquetting, pelletization, fibre extraction and biochar; these are best treated as candidate pathways rather than universally suitable recommendations. Recommendation systems rank options according to item properties and contextual constraints; ensemble and tree-based learning methods such as Random Forests and XGBoost, and general-purpose libraries such as scikit-learn, offer a path toward learned ranking once sufficient interaction data exist, but a transparent rule-based baseline is more appropriate for an initial system with no historical transaction data.

D. Digital Marketplace and Environmental Estimation

A digital marketplace reduces search friction by standardizing listings and buyer requirements — material category, quantity, location, availability, quality indicators and price expectations — and matching must consider transportation, since a theoretically high-value residue may become uneconomical when the buyer is far away. Environmental indicators can help users understand the potential benefit of diverting waste from open burning or unmanaged disposal, but emission factors and system boundaries must be documented, and estimated potential benefit must be clearly distinguished from measured impact.

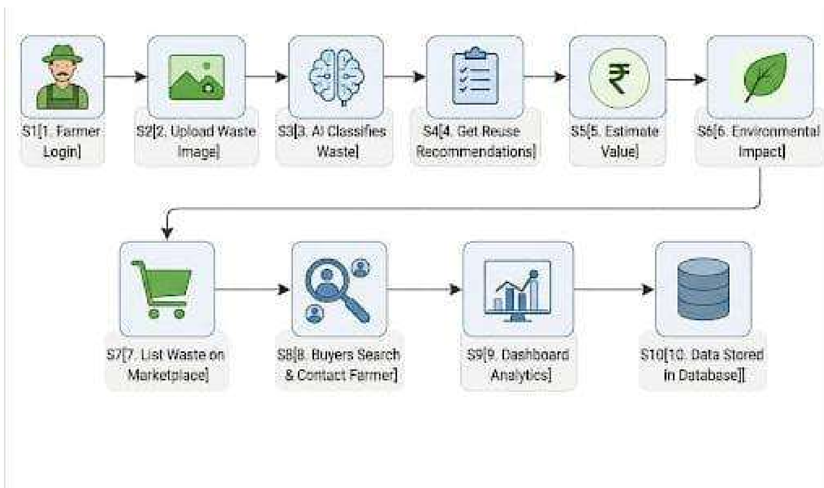

Fig. 1. System workflow diagram.

E. Research Gap

The principal gap addressed by this work is integration: identification, reuse recommendation, valuation and marketplace functions are often treated separately in existing tools and studies. AgriCycle AI proposes a single workflow in which the output of one component becomes the input to the next, beginning with a manageable AI scope and a curated recommendation knowledge base, then expanding through user feedback and validated market data.

3. SYSTEM ARCHITECTURE AND METHODOLOGY

The architecture supports a clear separation of concerns: the frontend does not contain model or valuation logic and instead requests structured results from the backend; the backend coordinates authentication, records and service calls; AI inference is isolated so a model can be replaced without redesigning the application; and the database preserves both raw inputs and derived results.

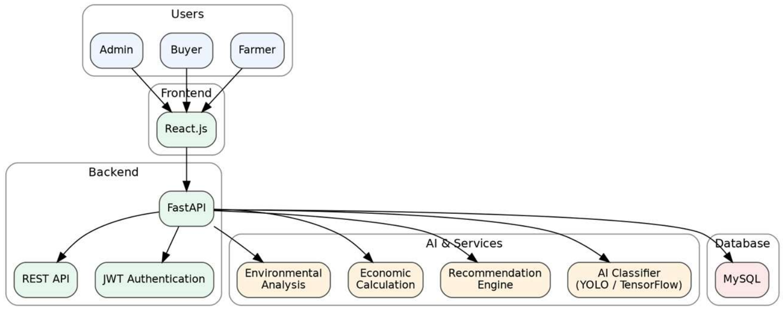

Fig. 2. System architecture diagram.

A. Architecture Overview and User Roles

The system uses a layered architecture in which the presentation layer handles user interaction, the API layer manages business logic, the AI service performs image inference, the recommendation and valuation services process the prediction, the marketplace manages listings and buyer requirements, and MySQL stores structured information. The farmer/waste-generator role creates waste records and listings, the buyer/processor role searches available material and defines requirements, and the administrator manages categories, reference data and marketplace moderation.

B. End-to-End Workflow

The planned flow is: register/login → upload image → enter quantity/location → validate and preprocess image → AI prediction → confidence check → retrieve utilization options → estimate indicative value → show environmental indicators → recommend relevant buyers/listings → optionally create a marketplace listing → buyer interest → listing status update. Each transition validates the data required by the next service — for example, valuation requires quantity and a valid reference price, and geographic matching requires a usable location.

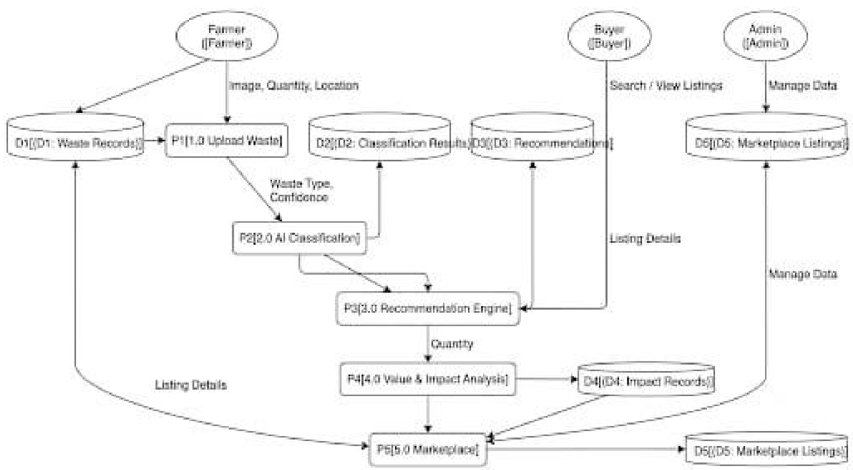

Fig. 3. Data flow diagram.

C. Error and Uncertainty Handling

A configurable confidence threshold distinguishes accepted predictions from uncertain cases, which trigger a request for a clearer image or manual category selection. Missing quantity or location data allows a draft record but blocks valuation and matching until the required fields are supplied, and stale price data is flagged with its reference date rather than presented as current.

4. TECHNOLOGY STACK AND JUSTIFICATION

The technology choices are intentionally conventional so that development effort concentrates on the AI, recommendation and marketplace logic rather than unfamiliar infrastructure. Python supports the AI and backend layers; FastAPI [7] provides a lightweight REST API; React [8] provides an interactive, component-based frontend; and MySQL [9] provides structured relational storage for the highly relational entities in the system.

Technology	Use	Reason
Python	AI, backend	Strong ML/CV ecosystem
YOLO	Detection	Real-time visual inference
OpenCV	Preprocessing	Mature CV utilities
FastAPI	REST API	Lightweight API framework
React	Frontend	Component-based UI
MySQL	Database	Structured relational data
Docker	Packaging	Reproducible deployment

Table I. Technology Stack and Justification

Uploaded images are treated as untrusted input, with server-side type and size validation rather than reliance on file extension alone. The security-oriented design includes password hashing, role-based authorization, server-side validation, upload-size limits, controlled database access and audit fields, to be implemented progressively as the prototype matures.

5. PROPOSED SYSTEM MODULES

A. Core Modules

The system comprises authentication and user management (role-specific profiles for Farmer, Buyer and Admin); an image-upload and waste-record module that stores image reference, predicted category, quantity, unit, location and status; an AI-prediction module that records category, confidence and model version; a recommendation module that ranks candidate uses with a reason and any prerequisite; valuation and environmental modules that present indicative, clearly labelled estimates; and marketplace, matching and buyer modules that support listing creation, search, filtering and staged (deterministic-then-weighted) matching. An administration module maintains reference data without requiring code changes, and a dashboard distinguishes direct counts from derived estimates.

B. Functional Requirements

ID	Requirement
FR-01–FR-03	Register/authenticate users; upload images; store quantity, unit, location

ID	Requirement
FR-04–FR-05	Generate AI prediction; display confidence/uncertainty
FR-06–FR-08	Generate recommendations; calculate indicative value; show environmental estimates
FR-09–FR-12	Manage listings; search/filter; buyer requirements; generate matches
FR-13–FR-14	Administrative controls; dashboard analytics

Table II. Functional Requirements (Condensed)

Non-functional requirements are evaluated through task completion and feedback (usability), response and inference time (performance), separable services (scalability), graceful handling of invalid inputs (reliability), modular documented code (maintainability), and authentication/authorization/validation (security).

6. PROPOSED AI MODEL AND DATA METHODOLOGY

A. Problem Definition and Dataset Design

The core AI problem is visual recognition of selected agricultural residues. Whether classification or object detection better represents real user images is decided from the collected data — detection is preferred when photographs frequently contain several distinguishable materials. The dataset should represent realistic field conditions (varied backgrounds, lighting, distances, camera devices and levels of visible damage) rather than only clean images, and is limited to a manageable initial category list: rice, wheat, sugarcane, corn, coconut, banana and fruit/vegetable residues, with additional categories introduced only once sufficient examples exist.

B. Annotation, Split and Augmentation

Annotation guidelines specify how to treat overlapping objects, mixed waste and ambiguous samples, with a sample of annotations reviewed by a second person. Data are split at the source or collection-session level to prevent near-duplicate images from appearing in both training and test sets and inflating performance estimates. Augmentation simulates realistic field conditions (brightness, orientation, scale, crops) while preserving class identity, avoiding transformations that alter the physical identity of the waste.

C. Model Development and Evaluation Metrics

A baseline model is trained first, with dataset version, hyperparameters, hardware and training duration recorded for reproducibility, following practice established for tree-based baselines such as Random Forests [3] and XGBoost [2], and evaluated with libraries such as scikit-learn [4]. Metrics are reported per class as well as overall, since an aggregate score can mask poor performance on a minority class, and the confusion matrix identifies visually similar categories needing better data or labels.

Metric	Interpretation
Precision / Recall	Correctness / completeness of predictions
F1-score	Balance of precision and recall

Metric	Interpretation
mAP	Detection performance across classes
Confusion Matrix	Class-specific error patterns
Inference Time	Practical responsiveness

Table III. Evaluation Metrics

D. Confidence Handling and Feedback Loop

A configurable confidence threshold, selected using validation data, triggers one of three states: accepted prediction, uncertain prediction requiring confirmation, or unsupported case. User corrections are reviewed before being added to a curated dataset for future model improvement, rather than automatically becoming training labels.

7. PROPOSED RECOMMENDATION, VALUATION AND MARKETPLACE

A. Knowledge Base and Recommendation Scoring

The knowledge base is implemented as structured records linking each waste category to candidate pathways, suitability conditions, buyer types and cautionary notes rather than guaranteed outcomes. An initial recommendation score combines category compatibility, quantity suitability, local processing availability, buyer demand and transportation distance, with weights finalized through expert review; the score is explainable, displaying the strongest positive factors and any disqualifying conditions.

B. Economic and Environmental Valuation

Price references include market region, date, unit and quality assumptions, since these attributes prevent an apparently precise but misleading estimate. The baseline calculation is:

Gross Value = Quantity × Reference Price

with a more realistic net estimate of:

Net Value = Gross Value − Collection − Transport − Processing − Other Costs

A generic environmental indicator is represented as Avoided Impact Estimate = Quantity Diverted × Validated Factor, where the factor depends on the comparison scenario and is drawn from a credible source; the platform does not present this as a measured reduction unless empirically verified.

C. Buyer Matching and Marketplace Workflow

Matching prioritizes feasible matches over merely the nearest ones, combining category, quantity, location, availability and price, with each match accompanied by an explanation (e.g. 'same category + within preferred distance + quantity within requested range'). The marketplace workflow proceeds from listing creation and validation, through active status, buyer interest and seller review, to coordinated collection or transfer, with real payments and logistics integration deferred to a later stage.

8. DATABASE DESIGN

MySQL [9] is proposed for the AgriCycle AI database because the core information — users, waste records, predictions, recommendations, price references, environmental factors and marketplace listings — is structured and highly relational, benefiting from referential integrity and SQL-based querying for dashboard analytics. The schema distinguishes reference data (waste categories, price references, environmental factors), which changes only through administrative action, from transactional data (waste records, predictions, listings), which is user- or system-generated.

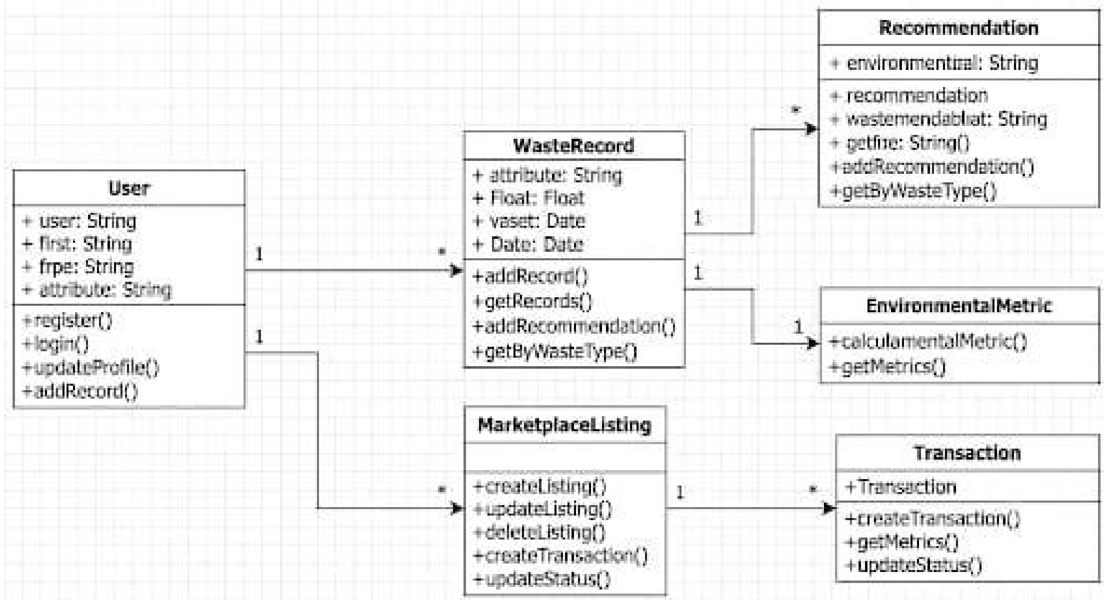

Fig. 4. Entity-relationship diagram.

A User can create multiple WasteRecords; each WasteRecord links to one Prediction and can generate multiple Recommendations. A WasteRecord can be converted into a MarketplaceListing, which can generate multiple Matches against BuyerRequirements. Economic and environmental calculations reference a time-stamped PriceReference or EnvironmentalFactor, so that later revisions do not silently alter historical estimates — preserving auditability across the identification, recommendation, valuation and marketplace pipeline.

9. EXPECTED OUTCOMES, EVALUATION AND LIMITATIONS

The expected outcome is a well-defined academic prototype demonstrating the feasibility of connecting AI recognition with practical decision support; the paper does not assume commercial success or validated environmental savings, which require implementation, external data and field evaluation.

A. Evaluation Plan

AI performance will be evaluated on a held-out test set using per-class precision/recall, F1-score, confusion matrix and mAP, together with inference time and qualitative error analysis. Recommendation quality will be rated by domain reviewers for relevance and feasibility; valuation will be compared against current local references for transparency and consistency; marketplace behaviour will be evaluated first through simulated buyer/seller scenarios; and usability will be tested with users of varied technical familiarity, focusing on whether they understand the difference between AI prediction, recommendation and indicative valuation.

B. Limitations

- No trained and validated model is claimed at this stage.
- No real marketplace transactions are claimed.
- Reference prices and environmental factors depend on regional, seasonal and source assumptions.
- Mixed waste may be difficult to classify reliably from a single image.
- Some utilization pathways require physical testing not available from an image.

10. CONCLUSION AND FUTURE SCOPE

AgriCycle AI proposes a practical extension of agricultural-waste management from simple disposal toward identification, utilization and market connectivity, integrating AI-based visual recognition, a recommendation knowledge base, indicative economic and environmental estimation, and a marketplace into a single workflow. As the project has not yet been implemented and validated, this paper deliberately describes the proposed system and methodology; subsequent work will focus on dataset collection, model training, expert validation, user testing and controlled marketplace evaluation. Future extensions — evidence-driven and prioritized after the initial prototype reveals real user needs — include a mobile application for field use, regional-language and voice-based interaction, real-time local price feeds, GPS-based buyer matching, IoT/sensor integration for moisture and quality indicators, and validated life-cycle environmental calculations.

REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 2016.

[2] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016.

[3] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[4] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," J. Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[5] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.