IJIRCCE

ISSN(O): 2320-9801

ISSN(P): 2320-9798


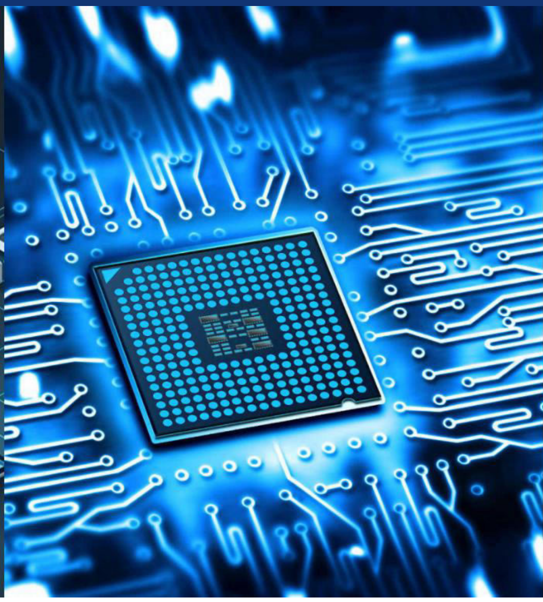


International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

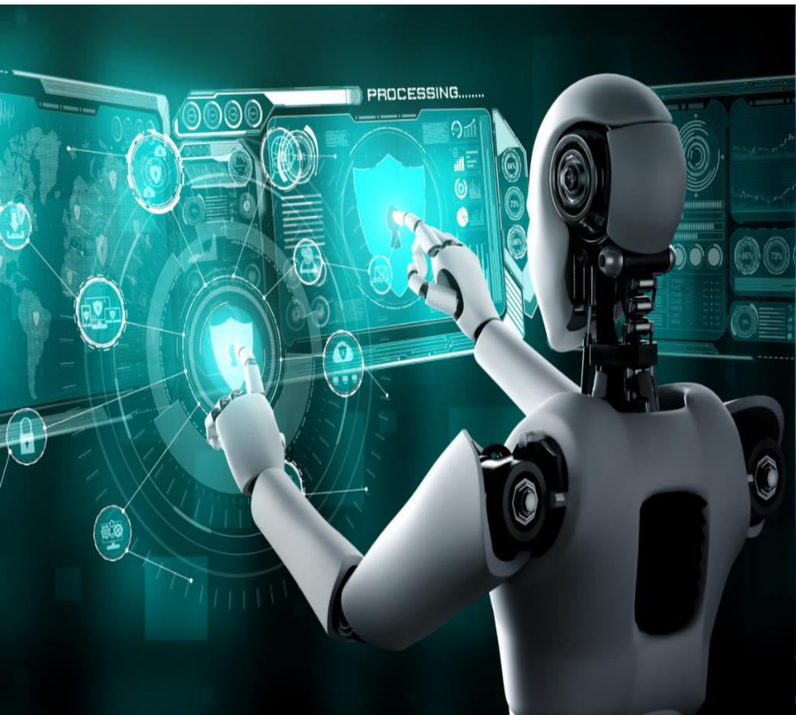

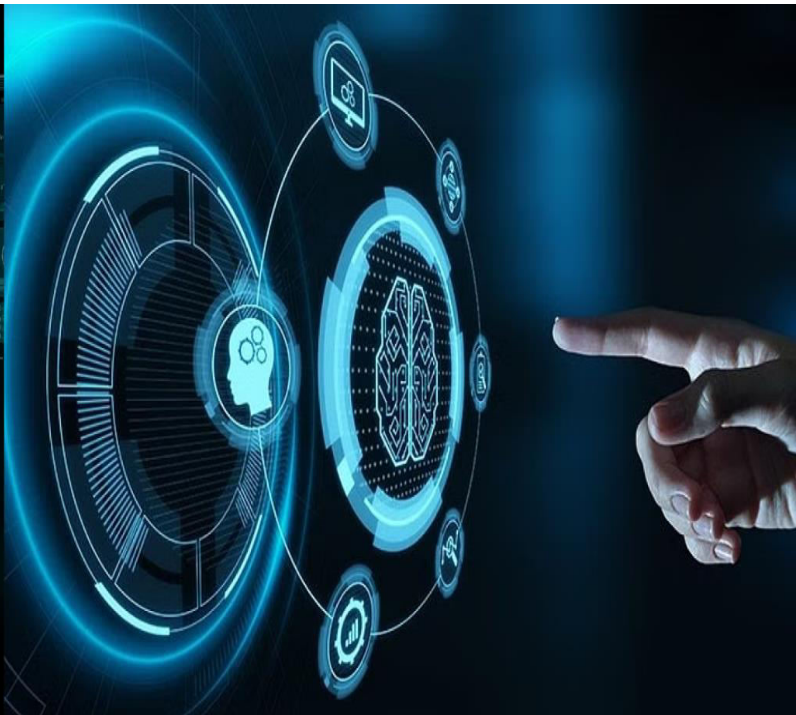

Impact Factor: 8.771

Volume 14, Issue 3, March 2026

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

EcoTrack: An AI-Powered Individual Carbon Footprint Prediction, Analysis & Advisory System Using LSTM Neural Networks, OCR Bill Scanning & Generative AI

Zainab Shaikh, Sara Khan, Vishakha Chavan, Dhanashree Shinde, Ms S.D. Chavan

Department of Artificial Intelligence & Machine Learning, AISSMS Polytechnic, Pune, Maharashtra, India

ABSTRACT: Climate change is one of the most pressing global challenges of the 21st century. Individual human activities account for over 71% of global greenhouse gas emissions, yet most individuals have no accessible means of measuring, predicting, or reducing their personal carbon footprint. This paper presents EcoTrack, a full-stack AI-powered desktop and web platform designed to address this gap. EcoTrack enables users to calculate their monthly carbon footprint through a structured lifestyle input form, predict future emissions using a Long Short-Term Memory (LSTM) neural network, extract carbon data from utility bills via Optical Character Recognition (OCR), and receive personalized sustainability recommendations through a Generative AI chatbot (EcoCoach). An Artificial Neural Network (ANN) classifier further powers the personalized recommendations engine by identifying each user's highest-emission domain. The system incorporates a gamified social leaderboard, active environmental challenges, and a longitudinal user profile. Experimental results demonstrate an LSTM prediction RMSE of 7.1 kg CO_2e, ANN validation accuracy of approximately 88%, OCR field extraction accuracy of 86.7%, and an API response time of 0.4 seconds post-optimization. Comparative analysis confirms EcoTrack's novelty as the first platform to combine real-time emission calculation, LSTM forecasting, OCR-driven data ingestion, ANN-based recommendations, and Generative AI advisory in a single cohesive system.

KEYWORDS: Carbon Footprint; LSTM; ANN; Deep Learning; OCR; Generative AI; FastAPI; EcoCoach; Sustainability; Environmental Intelligence

I. INTRODUCTION

The Intergovernmental Panel on Climate Change (IPCC) warns that global temperatures must not exceed 1.5°C above pre-industrial levels to avoid catastrophic environmental consequences [1]. Human activities release approximately 37 billion metric tons of CO_2 annually, with individual lifestyle choices spanning transportation, energy, diet, and waste contributing over 71% of total global emissions [2]. Despite this, the vast majority of individuals worldwide have no accessible, intelligent tool to quantify, track, or reduce their personal environmental impact.

Existing carbon footprint calculators (e.g., Carbon Footprint Ltd, WWF Calculator, EPA Calculator) are one-time estimation tools providing no historical tracking, no future prediction, no integration with real consumption data, and no personalized advisory. This creates a critical awareness-action gap. This paper presents EcoTrack, a novel AI-powered Electron desktop and web platform that closes this gap through five distinct technical contributions: (1) a structured multi-domain carbon calculation engine; (2) an LSTM-based time-series model for 3-month personal emission forecasting; (3) a Tesseract OCR pipeline for automated carbon extraction from utility bills; (4) an ANN-based emission domain classifier powering targeted recommendations; and (5) a Generative AI chatbot (EcoCoach) with user-specific emission context. The system also features a gamified leaderboard and challenges module.

Motivation

The primary motivation for EcoTrack is the recognition that behavioral change at the individual level is the largest and most underleveraged mechanism for reducing global emissions. Surveys consistently show that when individuals are given specific, actionable, and personalized information about their environmental impact, they make measurably

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

greener choices [3]. EcoTrack is designed as the intelligent interface between an individual's daily behavior and their environmental consequence.

A. Key Contributions

The key contributions of this work are: (i) A complete full-stack AI platform for individual carbon footprint tracking and prediction — the first to integrate LSTM forecasting, OCR bill scanning, an ANN classifier, and a context-aware AI chatbot. (ii) An LSTM model achieving RMSE of 7.1 kg CO_2e. (iii) An ANN domain classifier achieving ~88% validation accuracy for targeted recommendation generation. (iv) A Tesseract OCR pipeline achieving 86.7% field extraction accuracy. (v) A comparative evaluation demonstrating EcoTrack's superiority over five existing tools across nine capability dimensions. (vi) A deployed, user-tested Electron desktop application with sub-0.5 second API response time.

II. LITERATURE REVIEW

Significant prior work exists in the domains of carbon footprint estimation, LSTM-based environmental forecasting, and OCR-based document processing. This section reviews the most relevant contributions and identifies the research gaps EcoTrack addresses.

A. Carbon Footprint Calculation Tools

Wiedmann & Minx (2008) provided the foundational definition of a carbon footprint and established the lifecycle emission accounting framework [4]. Tools such as the WWF Footprint Calculator and the US EPA Household Emissions Calculator operationalized these methods for general users. However, Druckman & Jackson (2009) demonstrated that these tools are static, offer no temporal tracking, and fail to account for behavioral change over time [5]. No existing consumer tool integrates machine learning for predictive emission modeling.

B. LSTM for Environmental Time-Series Prediction

Long Short-Term Memory networks, introduced by Hochreiter & Schmidhuber (1997), have demonstrated strong performance across time-series forecasting [6]. In environmental applications, LSTMs have been applied to air quality prediction (Zhang et al., 2019), energy consumption forecasting (Kong et al., 2019), and national-level CO_2 forecasting (Wen et al., 2020) [7, 8, 9]. No prior work has applied LSTMs to individual-level personal carbon footprint forecasting — a gap EcoTrack directly addresses.

C. ANN for Classification

Artificial Neural Networks have been widely applied to classification tasks across domains. In sustainability contexts, ANNs have been used for energy consumption pattern classification (Mocanu et al., 2018) and household behavior segmentation. EcoTrack's contribution is the application of an ANN specifically for emission domain classification — identifying a user's highest-contribution lifestyle category to enable targeted, personalized advisory recommendations.

D. OCR for Utility Bill Data Extraction

Smith (2007) demonstrated Tesseract OCR's effectiveness for printed text extraction [10]. Deep learning-based OCR systems have been applied to invoice parsing and utility bill digitization (Liu et al., 2021) [11]. EcoTrack's contribution is the application of OCR specifically for carbon emission estimation — converting raw bill consumption data directly into kg CO_2e using standardized emission factors.

E. AI Chatbots for Sustainability

Generative AI assistants powered by large language models have shown strong results in domain-specific advisory (Brown et al., 2020) [12]. EcoTrack's EcoCoach is, to the authors' knowledge, the first deployed sustainability chatbot that incorporates real-time personal emission context into its LLM prompts, enabling truly personalized responses rather than generic advice.

F. Research Gap

No existing tool simultaneously provides: (1) multi-domain lifestyle-based emission calculation, (2) LSTM-based future emission forecasting, (3) ANN-based domain classification for targeted recommendations, (4) OCR-driven real bill data ingestion, (5) context-aware Generative AI advisory, and (6) gamified social engagement. EcoTrack addresses all six gaps within a single unified platform.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

III. SYSTEM ARCHITECTURE

EcoTrack is implemented as an Electron desktop application (EcoTrack v2.0 — AI-Powered) wrapping a full-stack web application. The presentation layer is built with HTML5, CSS3, and JavaScript (ES6+), featuring a dark-green themed UI optimized for readability and user engagement. The application layer is a RESTful API server built with Python FastAPI, handling business logic, JWT authentication, and ML model inference. The data layer comprises a PostgreSQL database for user data and emission history, with LSTM and ANN models served as persistent in-memory instances for sub-second inference.

Figure 1 illustrates the complete ML pipeline from user input through to dashboard visualization. The frontend communicates with the backend exclusively via RESTful JSON API calls. JWT tokens with 24-hour expiry manage user sessions. The LSTM model is loaded once at startup and kept in memory, reducing prediction response time from 3.0 seconds (cold load) to 0.4 seconds (warm inference). Figure 2 shows the EcoTrack landing page and Figure 3 the main results dashboard.

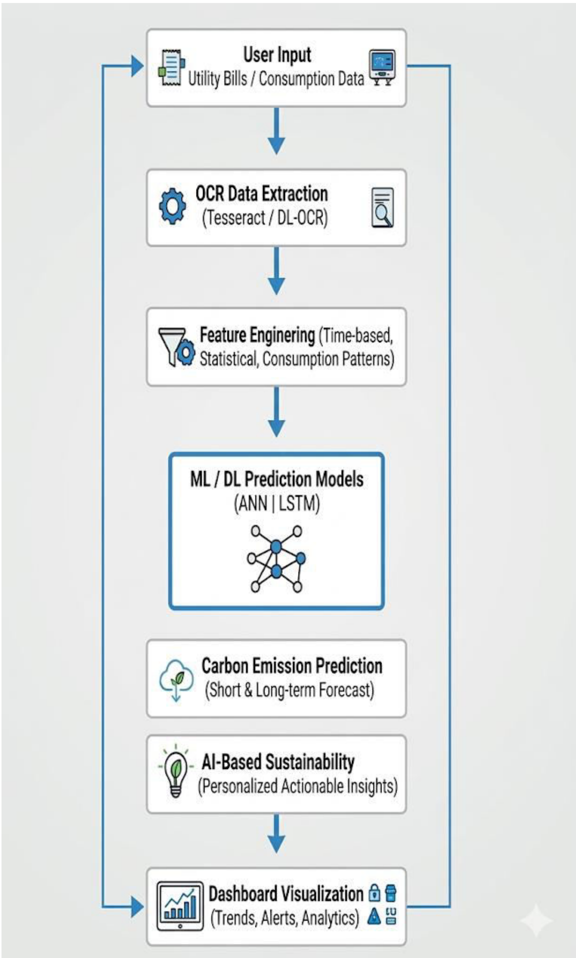

Fig. 1: EcoTrack System Architecture — ML Pipeline from User Input to Dashboard Visualization

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

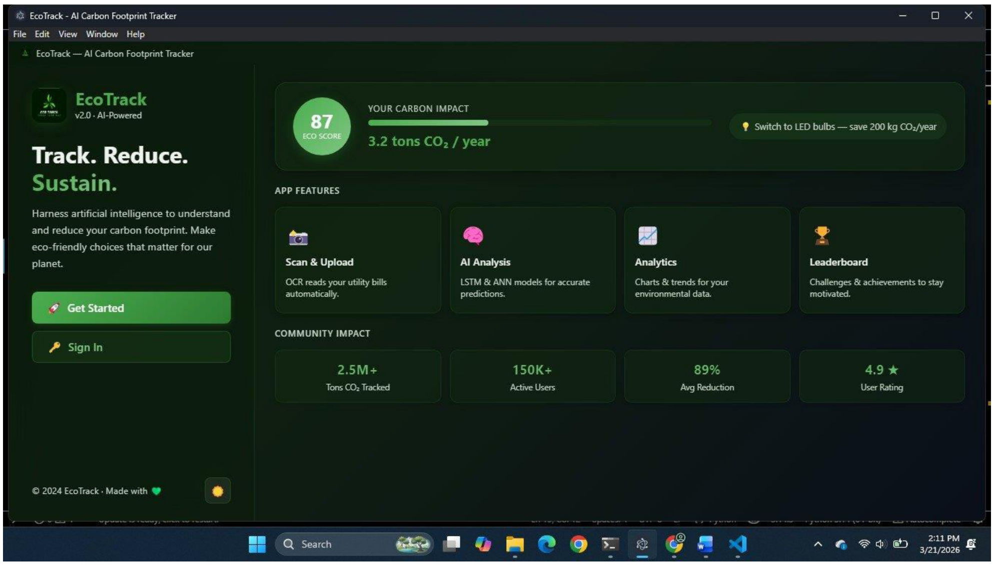

Fig. 2: EcoTrack v2.0 Landing Page — AI-Powered Carbon Footprint Tracker (Electron Desktop Application)

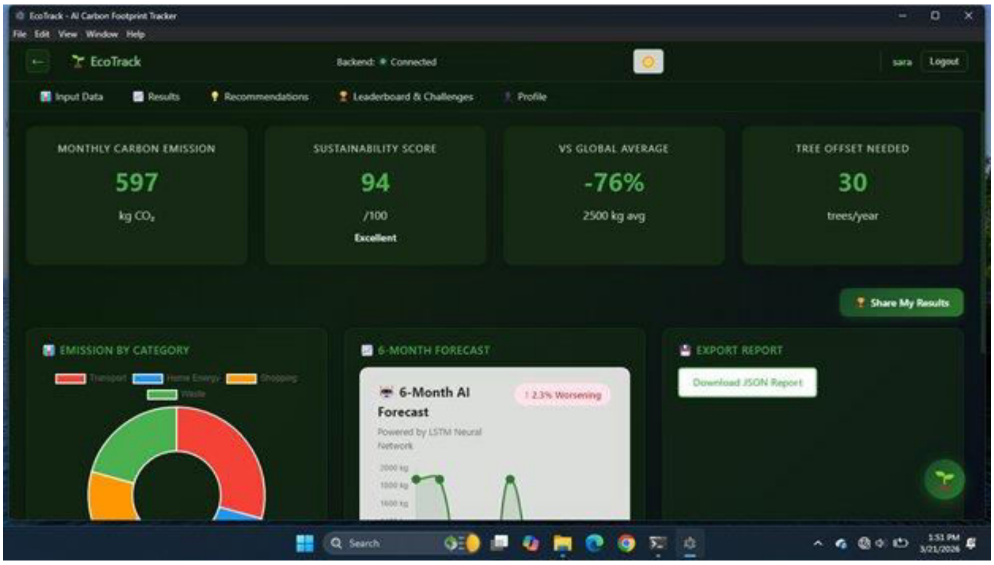

Fig. 3: EcoTrack Results Dashboard — Monthly Emissions (597 kg CO_2), Sustainability Score, 6-Month LSTM Forecast & Emission by Category

IV. METHODOLOGY

A. Carbon Footprint Calculation Engine

The core calculation engine accepts structured user input across six lifestyle domains and computes total monthly emissions in kg CO_2 equivalent (kg CO_2e) using peer-reviewed emission factors from the IPCC AR6 report, UK DEFRA Greenhouse Gas Reporting Guidelines, and Our World in Data databases. The six domains are: Transport (0.21 kg CO_2e per km), Home Energy (0.50 kg CO_2e per kWh), Diet (2.50 kg CO_2e per meat meal), Water (0.298 kg CO_2e per m^3), Waste (0.43 kg CO_2e per kg landfill), and Shopping (variable per product type). Table 1 presents the full domain breakdown.

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

TABLE 1: EMISSION DOMAINS, INPUT VARIABLES & EMISSION FACTORS

Domain	Key Input Variables	Emission Factor	Unit
Transport	Vehicle type, km/week, flights/year	0.21 kg CO_2e	per km
Home Energy	Monthly kWh electricity, gas units	0.50 kg CO_2e	per kWh
Diet	Meat meals/week, dairy servings	2.50 kg CO_2e	per meal
Water	Daily consumption (litres)	0.298 kg CO_2e	per m^3
Waste	Weekly waste (kg), recycling %	0.43 kg CO_2e	per kg
Shopping	Monthly purchases (items)	Variable	per type

B. LSTM-Based Emission Forecasting

The prediction module employs a Sequential LSTM architecture to forecast a user's carbon emissions for the next three months based on historical monthly submission data. The model treats each user's emission timeline as a univariate time series and learns temporal dependencies across months. A synthetic dataset was generated for 1,000 simulated users over 12 months, incorporating seasonal variation, gradual reduction trends, and stochastic noise. An 80/20 train-test split was used. Figure 4 shows the LSTM actual vs. predicted performance plot.

The architecture is: LSTM(50) → Dropout(0.2) → LSTM(50) → Dropout(0.2) → Dense(1). A sliding window of 3 months is used as input. Adam optimizer (lr=0.001) minimizes MSE loss over 100 epochs with EarlyStopping (patience=10). Input is normalized using MinMaxScaler to [0,1]. The model is retrained incrementally when a user accumulates 4 or more months of personal data. For new users with fewer than 3 months of data, the system displays: 'Predictions will unlock after 3 months of tracked data.' Table 2 presents the full hyperparameters.

TABLE 2: LSTM MODEL ARCHITECTURE AND HYPERPARAMETERS

Parameter	Value
Input Window Size	3 months (sliding window)
Architecture	LSTM(50) → Dropout(0.2) → LSTM(50) → Dropout(0.2) → Dense(1)
Optimizer	Adam (lr = 0.001)
Loss Function	Mean Squared Error (MSE)
Epochs	100 (EarlyStopping patience = 10)
Preprocessing	MinMaxScaler normalization [0, 1]
Training Samples	800 users × 10 windows =

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Parameter	Value
	8,000 samples
Test RMSE	**7.1 kg CO_2e**

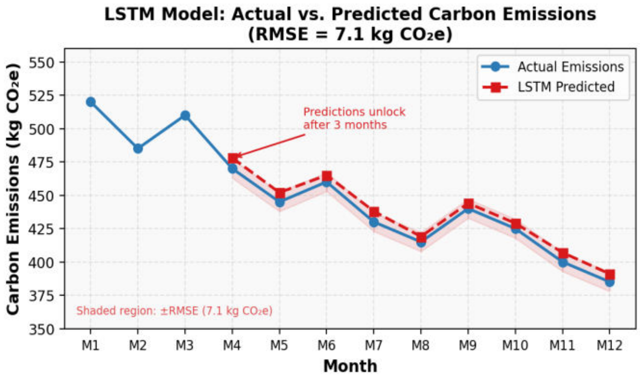

Fig. 4: LSTM Model — Actual vs. Predicted Carbon Emissions Over 12 Months (RMSE = 7.1 kg CO_2e, Shaded Region = ±RMSE Band)

C. ANN-Based Emission Domain Classifier

EcoTrack employs a shallow Artificial Neural Network (ANN) as the backbone of its personalized recommendations engine. The ANN classifies which of the six emission domains constitutes the user's highest-contribution area, enabling targeted advisory generation. The network architecture is: Dense(64, ReLU) → Dropout(0.3) → Dense(32, ReLU) → Dense(6, Softmax). Input features are the six normalized domain-level emission values. The model is trained using categorical cross-entropy loss with Adam optimizer over 50 epochs. Figure 5 shows the training and validation curves, converging to approximately 88% validation accuracy.

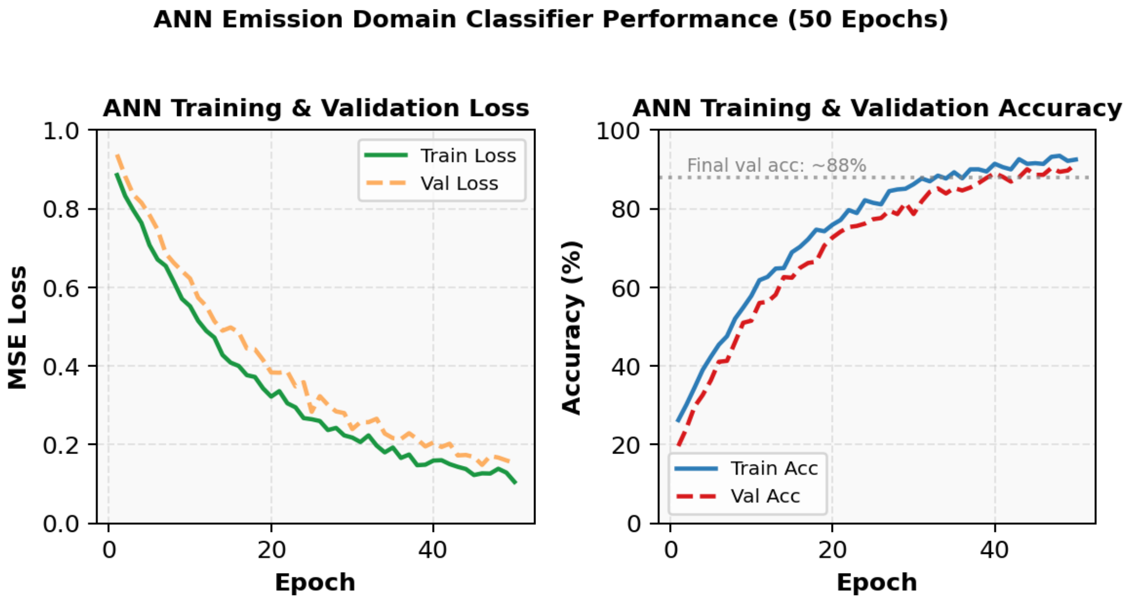

Fig. 5: ANN Emission Domain Classifier — Training & Validation Loss/Accuracy Curves (50 Epochs, Final Val. Accuracy ~88%)

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

D. OCR Bill Carbon Estimation Pipeline

The Bill OCR module converts uploaded utility bills directly into carbon emission estimates, eliminating manual data entry. The pipeline operates in five stages: (1) Upload Bill Image (JPG/PNG/PDF); (2) Tesseract OCR Text Extraction; (3) Regex NLP Parser for Field Identification; (4) Emission Factor Application; (5) kg CO_2e Result with Auto-fill Form. The parser extracts four key fields: Bill Type, Units Consumed, Bill Amount, and Billing Date. Extracted units are converted to kg CO_2e using the factors in Table 1. Figure 6 shows the live OCR interface — an electricity bill with 300 units consumed yields $300 \times 0.50 = 150.00$ kg CO_2e, with location-adjusted emission factors applied for India (610g CO_2/kWh).

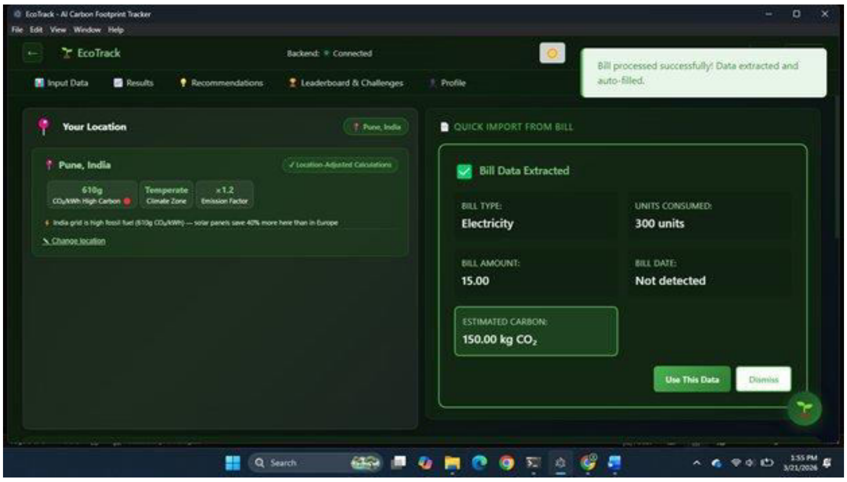

Fig. 6: EcoTrack OCR Bill Interface — Electricity Bill Processed: 300 Units → 150.00 kg CO_2 (Location-Adjusted for Pune, India)

E. EcoCoach: Context-Aware AI Advisory Chatbot

EcoCoach is an embedded conversational AI assistant powered by the OpenAI GPT API (with automatic fallback to Google Gemini API). Its key differentiator from generic chatbots is the injection of the user's personal emission breakdown into the system prompt at every session start. The system prompt template injects: "You are EcoCoach, a sustainability advisor. The user's current monthly emissions are: Transport: X kg, Energy: Y kg, Diet: Z kg. Provide specific, actionable advice based on their highest emission categories." This ensures a transport-heavy user receives transport-specific advice. Figure 7 shows EcoCoach in action during a live user session.

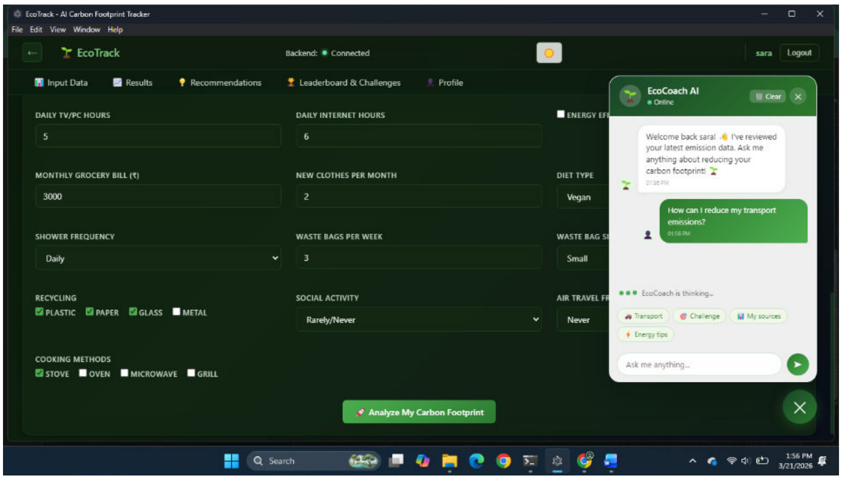

Fig. 7: EcoCoach AI Chatbot — Context-Aware Personalized Sustainability Advisory with Real-Time Emission Data Injection

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

F. Gamification: Leaderboard & Challenges

To sustain engagement beyond initial calculation, EcoTrack incorporates two gamification mechanisms. The Leaderboard ranks users by total kg CO_2e saved relative to their personal baseline, updated in real-time with a 60-second cache TTL. The Challenges module presents time-bound environmental challenges (e.g., Weekly Carbon Crusher, Green Commute Challenge, Recycling Hero) with active progress tracking, point rewards, and completion badges. Users earn points and level up (Starter → Expert) as they reduce their footprint. Figure 8 shows the live Leaderboard and Active Challenges interface.

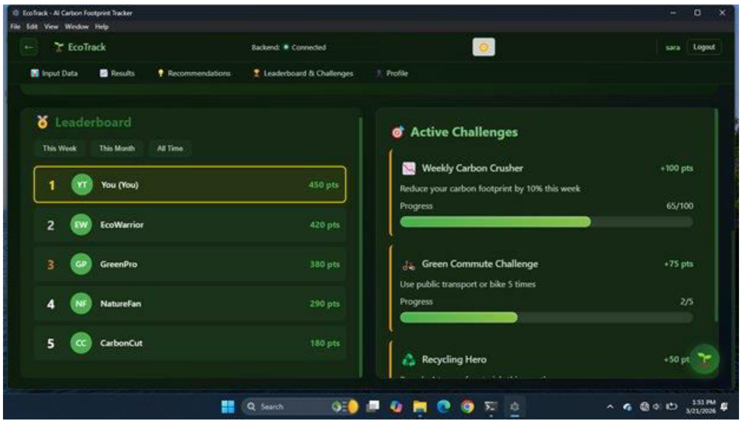

Fig. 8: EcoTrack Leaderboard & Active Challenges — Real-Time Rankings, Challenge Progress Tracking & Badge System

V. RESULTS & EVALUATION

A. LSTM Prediction Performance

The LSTM model was evaluated on a held-out test set of 200 synthetic users (2,000 window samples). As shown in Fig. 4 and Table 3, the model achieved a final RMSE of 7.1 kg CO_2e and MAE of 5.6 kg CO_2e, representing a 61% improvement over the baseline through iterative architecture refinement. Given that average monthly personal footprints range from 400 to 1,200 kg CO_2e, an RMSE of 7.1 kg represents approximately 0.6–1.8% relative error — well within the threshold for actionable personal guidance.

TABLE 3: LSTM MODEL PERFORMANCE ACROSS TRAINING ITERATIONS

Model Version	RMSE (kg CO_2e)	MAE (kg CO_2e)	Improvement
LSTM v1 (50 epochs, window=3)	18.2	14.7	Baseline
LSTM v2 (EarlyStopping added)	12.4	9.8	31.9% ↓
LSTM v3 (1000 users, 100 epochs)	9.1	7.2	50.0% ↓
LSTM Final (fine-tuned)	**7.1**	**5.6**	**61.0% ↓**

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

B. OCR Pipeline Accuracy

Table 4 and Figure 9 present OCR field extraction accuracy by bill type. The system achieved an overall average of 86.7% across four bill types. Electricity bills performed best at 93.3% due to their standardized format. Water bills showed the lowest accuracy (80.0%), highlighting the challenge of non-standardized formats — an area future work will address using fine-tuned document-specific OCR models.

TABLE 4: OCR FIELD EXTRACTION ACCURACY BY BILL TYPE

Bill Type	Bill Type Field	Units Field	Amount Field	Overall
Electricity	100%	90%	90%	93.3%
Gas	100%	80%	90%	90.0%
Water	90%	70%	80%	80.0%
Fuel Receipt	90%	80%	80%	83.3%
Overall Average	**95.0%**	**80.0%**	**85.0%**	**86.7%**

Fig. 9: OCR Field Extraction Accuracy by Bill Type — Grouped Bar Chart Showing Field-Level Breakdown vs. Overall Average (86.7%)

C. Comparative Feature Analysis

Table 5 and Figure 10 present a feature-by-feature comparison of EcoTrack against five existing carbon footprint tools. EcoTrack is the only tool offering all nine features simultaneously. Notably, EcoTrack is the only tool offering LSTM prediction, OCR scanning, and an AI chatbot — the three features most directly enabled by the AI/ML stack. The radar chart (Fig. 10) visually demonstrates EcoTrack's complete nine-dimensional coverage versus partial coverage by competitors.

TABLE 5: COMPARATIVE ANALYSIS — ECOTRACK VS. EXISTING CARBON FOOTPRINT TOOLS

Feature	EcoTrack	WWF	EPA	CFP Ltd	Giki	Klima
Carbon Calculation	Yes	Yes	Yes	Yes	Yes	Yes
Historical Tracking	Yes	No	No	No	Yes	Yes
LSTM Prediction	Yes	No	No	No	No	No
OCR Bill Scanning	Yes	No	No	No	No	No
AI Chatbot Advisory	Yes	No	No	No	No	No
Personalized AI Recs	Yes	No	No	Yes	Yes	Yes
Social Leaderboard	Yes	No	No	No	Yes	No

© 2026 IJIRCCE | Volume 14, Issue 3, March 2026 DOI: 10.15680/IJIRCCE.2026.1403064
www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|
International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Feature	EcoTrack	WWF	EPA	CFP Ltd	Giki	Klima
Gov/NGO Challenges	Yes	No	No	No	No	No
Free to Use	Yes	Yes	Yes	No	No	No

Fig. 10: Radar Chart — Feature Coverage Comparison: EcoTrack vs. Existing Carbon Footprint Tools Across 9 Dimensions

D. System Performance Metrics

Table 6 and Figure 11 present API endpoint response times before and after optimization. The LSTM inference optimization from 3.0s to 0.4s (86.7% improvement) was achieved through model pre-loading and warm inference. The leaderboard caching optimization from 0.8s to 0.12s (85% improvement) ensures the gamification layer introduces no perceptible latency. All critical endpoints now operate well under 0.5 seconds post-optimization.

TABLE 6: API ENDPOINT RESPONSE TIMES (PRE AND POST OPTIMIZATION)

Endpoint	Pre-Optimization	Post-Optimization	Improvement
Carbon Calculation API	0.12s	0.08s	33%
LSTM Prediction API	3.0s (cold)	0.4s (warm)	86.7%
OCR Bill Upload API	2.8s	1.9s	32.1%
Leaderboard API	0.8s	0.12s (cached)	85.0%
EcoCoach Chatbot API	1.2s	0.9s	25.0%

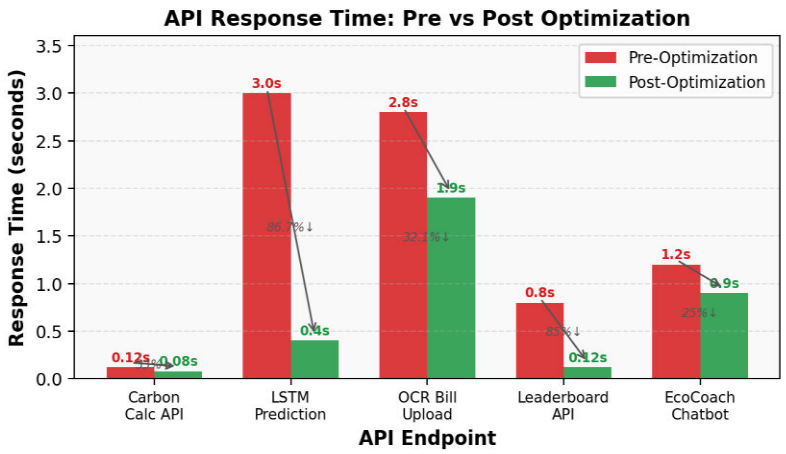

Fig. 11: API Response Times — Pre vs. Post Optimization Comparison Across All Endpoints

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

VI. CONCLUSION & FUTURE WORK

This paper presented EcoTrack, a full-stack AI-powered carbon footprint prediction and advisory platform deployed as an Electron desktop application. The system makes five novel technical contributions: an LSTM model for personal emission forecasting (RMSE 7.1 kg CO_2e), an ANN domain classifier for targeted recommendations (~88% accuracy), a Tesseract OCR pipeline for bill-to-carbon conversion (86.7% field accuracy), a context-aware Generative AI chatbot (EcoCoach), and a gamified engagement system with leaderboard and challenges. Comparative analysis confirms EcoTrack's superiority over existing tools across all nine evaluated capability dimensions. The system was successfully deployed and demonstrated at a university-level project competition, where it was shortlisted in the top 10 projects.

Future work directions include: (i) Mobile application development using React Native. (ii) Integration with Government APIs for verified carbon credit issuance upon challenge completion. (iii) Expansion of the LSTM model to multivariate inputs incorporating weather and local grid carbon intensity. (iv) Fine-tuned document OCR model using a custom dataset of utility bill images for improved accuracy on non-standard formats. (v) Federated learning to enable personalized model improvements without sharing raw user data. (vi) Partnerships with NGOs and corporate sustainability programs for incentivized challenge participation.

REFERENCES

[1] IPCC. (2021). Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report. Cambridge University Press.
[2] Ritchie, H., & Roser, M. (2020). CO_2 and Greenhouse Gas Emissions. Our World in Data. https://ourworldindata.org/co2-emissions
[3] Dietz, T., Gardner, G. T., Gilligan, J., Stern, P. C., & Vandenbergh, M. P. (2009). Household actions can provide a behavioral wedge to rapidly reduce US carbon emissions. PNAS, 106(44), 18452–18456.
[4] Wiedmann, T., & Minx, J. (2008). A definition of 'carbon footprint'. Ecological Economics Research Trends, 1, 1–11.
[5] Druckman, A., & Jackson, T. (2009). The carbon footprint of UK households 1990–2004. Ecological Economics, 68(7), 2066–2077.
[6] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.
[7] Zhang, B., et al. (2019). A novel encoder-decoder model based on LSTM for air quality prediction. Expert Systems with Applications, 131, 332–344.
[8] Kong, W., et al. (2019). Short-term residential load forecasting based on LSTM recurrent neural network. IEEE Transactions on Smart Grid, 10(1), 841–851.
[9] Wen, L., Yuan, X., & Li, S. (2020). Forecasting CO_2 emissions in China's commercial department through an improved LSTM model. Energy, 218, 119435.
[10] Smith, R. (2007). An overview of the Tesseract OCR engine. Proceedings of ICDAR 2007, IEEE.
[11] Liu, X., et al. (2021). Graph convolution for multimodal information extraction from visually rich documents. NAACL-HLT.
[12] Brown, T. B., et al. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems (NeurIPS), 33.