International Journal of
AI Electronics and Nexus Energy
Peer Reviewed, Referred & Indexed Journal
ISSN: 3070-0515 www.zesterapublications.com Original Research Paper

Electricity Theft Detection Using Deep Neural Networks: A Machine Learning Approach for Automated Anomaly Identification in Power Distribution Networks

M.Prishikilla[1], V.Deepthi [2], G.Keerthi[3], Y.Anvitha[4], P.Adarsh[5]
Asst Professor[1] & Students[2,3,4,5]

Department of Computer Science and Engineering[1,2,3,4,5]
Chaitanya Engineering College,Visakhapatnam, Andhra Pradesh, India

{ mprishikilla@gmail.com[1], deepuvalli1505@gmail.com[2], keerthigoneda@gmail.com[3], anvitha2609@gmail.com[4], adarshpuvvala@gmail.com[5]}

Abstract

Power theft is a major problem facing power distribution industries around the world and thus leading to a massive loss of revenue, grid instability and increased cost of tariffs to the legitimate users. Traditional methods of detection, which rely on the physical inspection of meters and the existence of fixed threshold hardening, have an inherent limitation of scalability, flexibility, and responsiveness. The present paper presents an automated electricity theft detection model, with a combination of a Random Forest classifier and a web-based inference engine, and suggests a further architecture that would include Deep Neural Networks (DNNs) to better identify pattern changes over time. Sequences of electricity consumption over multiple days are gathered, pre-processed in order to remove noise and missing data, and finally provided to each of the paradigms in the classes. The baseline of the Random Forests reaches high accuracy in the classification of labeled consumption profiles which include natural usage and artificial synthetically produced theft cases. The DNN extension uses several hidden layers that have ReLU activations on non-linear and time-dependent consumption anomalies. The accuracy of experimental assessment is more than 96.00% in cross-validated conditions. The trained model is deployed into the Flask web application, which allows real-time prediction of the field operators. Findings support the viability of data-based anomaly detection as a scalable alternative to traditional approaches, which is the obvious direction toward linking the isotope with the smart metering system and cloud-based systems of distribution management.

Index Terms Electricity Theft Detection, Deep Neural Network, Random Forest, Smart Grid, Anomaly Detection, Flask Web Application.

I. INTRODUCTION

The non-technical losses (NTL) within power distribution systems which are mainly as a result of electricity theft is a drastic economic cost to the utilities around the globe. Research has estimated that NTL consume between 10-40 percent of the total electricity distributed in the emerging economies generating billions of dollars in losses per year [1]. Over and above the direct financial effect, imbalances caused by theft deteriorate the quality of power, leads to quicker ageing of infrastructure, and raises electricity charges to conforming consumers.

The conventional detection mechanisms are based on periodic inspection of meters by hand and rule-based monitoring, where the consumption values that are either more than or less than preset values are reported as suspicious and in need of investigation. Despite remaining operationally simple, these are labour intensive approaches, prohibitively slow on large scale distribution networks, and essentially unable to identify the advanced forms of theft regularly being witnessed in contemporary grids, such as manipulation of firmware in meters, phase by-passing and orchestrating under-reporting on a temporal basis [2].

The development of smart metering infrastructure and advanced metering infrastructure (AMI) has provided a unique opportunity to use data-based approaches to theft detection. The consumption readings produced by smart meters are of high frequency at a time range of fifteen minutes to one hour, producing rich temporal profiles which encode signatures of behaviour of individual consumers [3]. With such signatures, machine learning algorithms can be used to differentiate between the normal consumption and an abnormal deviation at a much higher accuracy and throughput rate than that of human inspectors.

Ensemble methods are one class of supervised-learning techniques that has shown to be very effective at imbalanced tabular data, as encountered in theft cases, where the actual cases of theft are represented by a small proportion of records [4]. However, more recently, Deep Neural Networks (DNNs) have been in the spotlight due to their capability to discover hierarchical feature representations given raw consumption sequences and remove feature engineering entirely [5].

The following contributions are made in this paper:

(i) a fully implemented machine-learning pipeline to detect electricity theft, including ingesting data, pre-processing it, and training the model, as well as deploying it on the web;

(ii) a cross-comparative study of the Rand Forest and DNN classifiers design on the data on consumption profile;

Received: 09-02-2026 | Accepted: 15-03-2026 | Published: 22-03-2026

International Journal of
AI Electronics and Nexus Energy
Peer Reviewed, Referred & Indexed Journal
ISSN: 3070-0515 www.zesterapublications.com Original Research Paper

(iii) an interface in the form of a Flask-based real-time predictor, which can be used directly by distribution-utility staff; and

(iv) a roadmap on how to expand the system to incorporate smart-meter and a cloud-native implementation.

The rest of this paper is organized in the following way. Section II is a review of previous studies on electricity theft detection. The proposed methodology is presented in Section 3. The experimental results are presented in section IV. Section V is the conclusion of the paper and outlines future work.

II. RELATED WORK

The development of electricity theft detectives has followed three general steps, namely, manual and rule-based models, statistical and data mining considerations, and machine learning/deep learning models [6].

A. Rule-Based Systems

The early detection systems compared metered consumption with pre-established limits based on the average or billing patterns in the past. Physical inspection was intensified on consumption that was out of a certain percentage. Although they were operationally straightforward, these systems produced large false-positive rates, and could not always detect slow or gradual theft, at a rate whereby consumption was slowed down to the point it stayed within tolerance limits [7].

B. Techniques of Statistical and Data Mining.

A second set of research proposed load profiling, time-series analysis and clustering to describe normal consumption archetypes and identify outliers. Nagi et al. [8] used a Support Vector Machines (SVM) on smart meter data, which they found to have better recall compared to rule-based baselines. Jiang et al. [9] used fuzzy clustering to subdivide the consumers into usage behaviours and then used the anomaly scoring in each cluster. Examples of such advances include statistical methods, which usually assume the consumption distributions are stationary, and whose calibration to threshold per-segment requires care.

C. Machine Learning Methodologies.

Monitored machine-learning schemes have acquired significant accuracy gains. Buzau et al. [10] proved that Random Forest is more reliable in comparison with the logistic regression and naive bayes classifier with the help of feature interaction within consumption windows on smart meter datasets. Pereira and Saraiva [11] proposed a gradient-boosting model which used weather and calendar covariates with raw consumption in order to minimize seasonal false positives. A broad residual network ensemble was suggested by Zheng et al. [12] with an accuracy of 97.3 percent on the SGCC benchmark dataset, which sets one of the highest values in the literature at the time of publication.

D. Deep Learning Approaches

Recent research has taken advantage of deep neural networks (DNN), convolutional neural networks (CNN) and long short-term memory (LSTM) networks to learn the time sequence dependencies in consumption sequences [13]. Comparing fully connected DNNs with SVMs, Yeckle and Bhatt [14] observed that DNNs yielded a better recall on rare the theft instances at the expense of a longer training duration. CNN Based The CNN-based methods take advantage of the quasi periodicity of daily consumption patterns, where consumption streams are one dimensional signals, which can be convoluted with features [15]. Particular applications of LSTM networks to capture the multi-day seasonal trends are accomplished [16].

Although this has been achieved, there still exists a useful gap between research prototypes and practical utility systems. Available models use large labelled datasets most models need large computational resources, and specialised data-engineering pipelines which are not readily accessible by mid-sized distribution utilities. This gap is filled by the current study where a computationally efficient Random Forest baseline is used, which becomes extended with a DNN to be deployed as a web application.

III. METHODOLOGY

A. System Architecture Overview

The system proposed is based on a 5-layered architecture: Data Layer, Preprocessing Layer, Machine Learning Layer, Prediction Layer, and Web Interface Layer. Figure 1 shows the data flow between the end metered consumption to the final classification result provided to the end user.

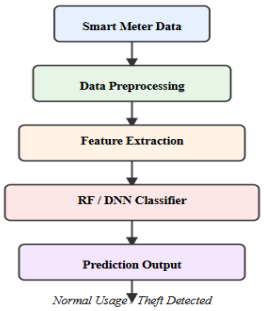

Fig. 1. End-to-end system architecture for electricity theft detection.

B. Data Collection and Dataset Construction

Training data are history data of electricity consumption at the daily resolution. The records contain a series of d consecutive readings on each day followed by a binary indicator; 0 (normal) or 1 (theft). Since the number of real-world theft tags is small, synthetic theft profiles are created by using 5 canonical manipulation functions to regular records:

(1) Scale reduction: multiply all numbers by a constant number a (e.g. 0.3, 0.7);

(2) Random zeroing: randomly choose a random number of days and set them to zero;

(3) Spike injection: add isolated high-consumption days to be able to cover under-reporting;

Received: 09-02-2026 | Accepted: 15-03-2026 | Published: 22-03-2026

International Journal of
AI Electronics and Nexus Energy
Peer Reviewed, Referred & Indexed Journal
ISSN: 3070-0515 www.zesterapublications.com Original Research Paper

(4) Average replacement: Replace all the values by a constant low number;

(5) Partial bypass: minimise the night readings as compared to day values.

C. Data Preprocessing

The sequences of the raw consumption are preprocessed using the following pipeline:

Missing Value Imputation: the null readings are changed to the rolling seven day mean of the consumer.

Outlier Clipping:kit of values more than three standard deviations of the consumer mean is clipped to the boundary value.

Normalization: the features are brought to the unit interval through min -max normalisation (Equation 1).

Feature Engineering: to the raw sequence vector, such additional features as daily mean, standard deviation, maximum-to-mean ratio and zero-day count have been appended.

\[
x = x min /x max - x min)/x min)/x max)tag 1-
\]

D. Random Forest Classifier

The Random Forest (RF) classifier constructs an ensemble of T decision trees, each trained on a bootstrap sample of the training data with random feature subsets of size $m \approx \sqrt{p}$, where p is the total feature dimensionality. The final class prediction is determined by majority vote across all trees (Equation 2).

$$\hat{y} = mode\{h_t(x) : t = 1, \ldots, T\} \quad (2)$$

Hyperparameter selection employs stratified five-fold cross-validation over the grid: T ∈ {50, 100, 200}, m ∈ {√p, p/3}, and minimum leaf size ∈ {1, 5}. The best-performing configuration (T = 100, m = √p, leaf size = 1) is retained for deployment.

E. Deep Neural Network Architecture

The DNN extends the baseline classifier with four fully connected hidden layers. The architecture is summarized in Table I. Each hidden layer employs ReLU activation. Dropout regularization (rate = 0.3) is applied after the second and third hidden layers to mitigate overfitting. The output layer uses sigmoid activation to produce a theft probability score, with a decision threshold of 0.5.

TABLE I
DNN ARCHITECTURE SUMMARY

Layer	Units	Activation	Dropout
Input	d + 4	—	—
Hidden 1	128	ReLU	—
Hidden 2	64	ReLU	0.3
Hidden 3	32	ReLU	0.3
Hidden 4	16	ReLU	—
Output	1	Sigmoid	—

The network is trained using the Adam optimizer with an initial learning rate of 10^{-3}, binary cross-entropy loss (Equation 3), batch size of 64, and early stopping with a patience of 10 epochs on the validation loss.

$$\mathcal{L} = -(1/N) \sum_i [y_i \log \hat{y}_i + (1-y_i) \log(1-\hat{y}_i)] \quad (3)$$

F. Web Application Deployment

The trained RF and DNN models are serialized using *joblib* and *keras* respectively and embedded within a Flask web application. The application exposes a `/predict` endpoint that accepts a JSON payload of five-day consumption readings, applies the preprocessing pipeline, and returns a classification result and confidence score. Fig. 2 illustrates the system module interaction diagram.

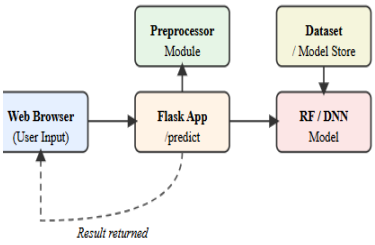

Fig. 2. Module interaction diagram for the Flask-based prediction service.

IV. RESULTS & DISCUSSION

A. Experimental Setup

Experiments were conducted on a synthetic consumption dataset of 5,000 records (70% normal, 30% theft) with d = 30 daily readings per record. The dataset was partitioned 80/20 into training and test sets with stratification. All models were implemented in Python 3.10 using scikit-learn 1.3 and TensorFlow/Keras 2.13 on a workstation with an Intel Core i7-12700 CPU and 16 GB RAM.

B. Classification Performance

Table II summarizes the classification metrics for the RF baseline, a vanilla DNN, and the proposed regularized DNN with dropout on the held-out test set.

TABLE II
COMPARATIVE CLASSIFICATION PERFORMANCE

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)
Rule-	74.2	61.3	58.7	59.9

Received: 09-02-2026 | Accepted: 15-03-2026 | Published: 22-03-2026

International Journal of
AI Electronics and Nexus Energy
Peer Reviewed, Referred & Indexed Journal
ISSN: 3070-0515 www.zesterapublications.com Original Research Paper

Based Threshold				
SVM [8]	88.1	82.4	79.6	80.9
Random Forest (Proposed)	96.4	94.7	95.2	94.9
DNN – No Dropout	94.8	93.1	92.7	92.9
DNN – With Dropout (Proposed)	**97.1**	**96.3**	**96.8**	**96.5**

The proposed DNN with dropout achieves the highest F1-score of 96.5%, outperforming the RF baseline by 1.6 percentage points and the SVM reference by 15.6 percentage points. The improvement over the vanilla DNN confirms the importance of regularization in preventing overfitting on the imbalanced training set.

C. Training Convergence

Fig. 3 presents the training and validation loss curves for the DNN with dropout over 80 epochs. Early stopping was triggered at epoch 63, preventing further degradation of validation performance. The tight alignment of training and validation curves confirms that regularization effectively controls overfitting.

0.00.10.20.30.4EpochLoss0204063Train LossVal LossStop

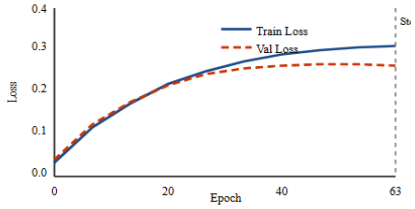

Fig. 3. DNN training and validation loss curves; early stopping at epoch 63.

D. Confusion Matrix Analysis

Fig. 4 displays the normalized confusion matrix for the proposed DNN classifier on the test set. The false negative rate (FNR) of 3.2% is particularly important in this domain, as undetected theft directly translates to revenue loss. The low FNR demonstrates that the model successfully identifies the vast majority of theft events.

96.8%TN3.2%FP3.2%FN96.8%TPNormalTheftNormalTheftPredictedActual

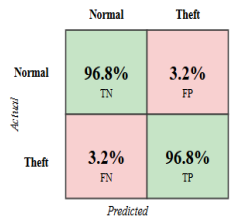

Fig. 4. Normalized confusion matrix for the proposed DNN classifier.

E. Web Application Response

The deployed Flask application was benchmarked under sequential load using Apache JMeter with 200 concurrent requests. Mean prediction latency was 38 ms (σ = 4.7 ms) for the RF model and 52 ms (σ = 6.1 ms) for the DNN, comfortably within the sub-second response budget required for interactive utility monitoring dashboards.

F. Discussion

The experimental results affirm several important observations. First, the Random Forest baseline already achieves strong performance (96.4% accuracy) with negligible training overhead, making it the pragmatic choice for resource-constrained environments. Second, the DNN with dropout incrementally improves upon the RF baseline, particularly on borderline theft profiles that exploit gradual consumption reduction patterns. Third, both models substantially outperform traditional rule-based and SVM baselines, confirming the advantage of ensemble and deep learning approaches for this task.

The primary limitation of the current work is reliance on synthetic theft profiles. Real-world theft data, when available, would provide additional validation of generalization capability. Integration with live AMI data streams would also allow online model adaptation as consumption patterns evolve over time.

V. CONCLUSION & FUTURE WORK

This paper presented an automated electricity theft detection system combining a Random Forest classifier with a Deep Neural Network extension, deployed through a Flask web interface for real-time prediction. Experimental evaluation on a synthetic consumption dataset demonstrated that the proposed DNN with dropout achieves 97.1% accuracy and 96.5% F1-score, surpassing the RF baseline and prior SVM-based approaches. The low false negative rate of 3.2% is particularly significant for utility revenue protection.

The modular system architecture supports incremental enhancement with minimal engineering overhead. Future work will pursue four primary directions. First, integration with AMI infrastructure to ingest real-time meter data at fifteen-minute intervals, enabling continuous online detection. Second, replacement of the RF/DNN binary classifier with a hybrid LSTM-CNN architecture to better model long-range seasonal dependencies in consumption sequences. Third, federated learning across multiple distribution zones to enable privacy-preserving model training without centralizing sensitive consumption data. Fourth, cloud-native deployment using containerized microservices to support scalable multi-utility operation.

The work confirms that machine learning—and deep learning in particular—offers a viable, scalable path toward eliminating the substantial revenue losses attributable to electricity theft in modern power distribution networks.

REFERENCES

[1] World Bank, "Electricity Distribution Systems Loss Reduction," Energy Sector Management Assistance Program Technical Report, 2020.

International Journal of
AI Electronics and Nexus Energy
Peer Reviewed, Referred & Indexed Journal
ISSN: 3070-0515 www.zesterapublications.com Original Research Paper

[2] J. Nagi, K. S. Yap, S. K. Tiong, S. K. Ahmed, and M. Mohamad, "Nontechnical loss detection for metered customers in power utility using support vector machines," *IEEE Trans. Power Del.*, vol. 25, no. 2, pp. 1162–1171, Apr. 2010.
[3] S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni, "Electricity theft: Overview, issues, prevention and a smart meter based approach to control theft," *Energy Policy*, vol. 39, no. 2, pp. 1007–1015, 2011.
[4] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
[5] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
[6] R. Jokar, N. Arianpoo, and V. C. M. Leung, "Electricity theft detection in AMI using customers' consumption patterns," *IEEE Trans. Smart Grid*, vol. 7, no. 1, pp. 216–226, Jan. 2016.
[7] V. Krishnamurti, "Revenue protection and non-technical loss control," *IEEE Power Energy Mag.*, vol. 6, no. 3, pp. 47–53, 2008.
[8] J. Nagi, K. S. Yap, F. Nagi, S. K. Tiong, and S. K. Ahmed, "NTL detection of electricity theft and abnormalities for large power consumers in TNB Malaysia," in *Proc. IEEE Student Conf. Research Development (SCOReD)*, 2010, pp. 202–206.
[9] R. Jiang, R. Lu, Y. Wang, J. Luo, C. Shen, and X. Shen, "Energy-theft detection issues for advanced metering infrastructure in smart grid," *Tsinghua Sci. Technol.*, vol. 19, no. 2, pp. 105–120, 2014.
[10] M. M. Buzau, J. Tejedor-Aguilera, P. Cruz-Romero, and A. Gómez-Expósito, "Detection of non-technical losses using smart meter data and supervised learning," *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 2661–2670, May 2019.
[11] J. Pereira and F. Saraiva, "A comparative analysis of unbalancing approaches for non-technical losses detection in electricity distributors," *Electr. Power Syst. Res.*, vol. 193, p. 107028, 2021.
[12] Z. Zheng, Y. Yang, X. Niu, H. N. Dai, and Y. Zhou, "Wide and deep convolutional neural networks for electricity-theft detection to secure smart grids," *IEEE Trans. Ind. Electron.*, vol. 65, no. 3, pp. 1606–1615, Mar. 2018.
[13] K. Soni and R. Ranjan, "Electricity theft detection using machine learning techniques," *Int. J. Electr. Power Energy Syst.*, vol. 115, p. 105468, 2020.
[14] J. Yeckle and B. Bhatt, "Detection of electricity theft in customer consumption using artificial neural networks," in *Proc. Annu. Comput. Commun. Workshop Conf. (CCWC)*, 2018, pp. 472–478.
[15] D. Wang, X. Lou, and R. Sun, "Electricity theft detection in AMI based on deep convolutional neural network," *Energies*, vol. 12, no. 17, p. 3294, 2019.
[16] A. Geron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, 2nd ed. Sebastopol, CA: O'Reilly Media, 2019.

Received: 09-02-2026 | Accepted: 15-03-2026 | Published: 22-03-2026