Originally published at SECRYPT link to web page: https://www.scitepress.org/HomePage.aspx

Anomaly detection in smart grid networks using power consumption data

Keywords: Smart Meter Security, Anomalies, Intrusion Detection System, False Data Injection, Deep Learning Model, Energy Data, LSTM-DAE Model.

Abstract: Smart meters, intelligent devices used for managing energy consumption of consumers, are one of the integral components of the smart grid infrastructure. The smart metering infrastructure can facilitate a two-way communications through the Internet to leverage home energy management and remote meter reading by the service providers. As consequence, the smart meters are extremely susceptible to various potential security threats, such as data tampering, distributed denial of services attack and spoofing attacks. In this paper, we put forward a scheme to detect anomalies in energy consumption data using real-world datasets. Thereby, addressing data tampering attacks. We have adapted an unsupervised machine learning method to distinguish the anomalous behaviour from the normal behaviour in the energy consumption patterns of consumers. In addition, we have proposed a robust threshold mechanism for detecting abnormalities against noise, which has not been used in smart grids before. Our proposed model shows an accuracy of 94.53% in detecting anomalous patterns in energy consumption data. This accuracy surpasses the existing benchmark in anomaly detection in energy consumption data using machine learning models (Huang and Xu, 2021).

1 INTRODUCTION

Smart meters form an integral part of the smart grid infrastructure. They play a significant role in regulating the advanced metering infrastructure(AMI) systems (Hart, 2008). The AMI enables different services, such as electronic billing, grid monitoring, grid operation and demand response for both consumers and providers. On the other hand, smart meters are deployed by electricity providers and retailers to monitor fine-grained energy consumption of households in real-time. Consequently, they are physically accessible and more prone to data tampering attacks. The demand for these smart meters is increasing with every passing day and they are being widely deployed. However, Tellbach and Li proved that cyber-attacks on these smart meters can incur huge financial losses (Tellbach and Li, 2018). These data are often communicated to the service providers over a secure channel (Erkin et al., 2013), and are required to monitor and manage the grid (Knirsch et al., 2016).

The cyber attacks in the smart grid can be destructive and cause the electronic devices like smart devices and/or generators to malfunction. Well known attacks are false data injection, spoofing, denial of services (DoS), man-in-the-middle, replay and meter by-pass attacks. We discuss the severity of these attacks below.

False Data Injection Attack: False data injection attack is launched to inject fake data or payloads into the smart meters or the advanced metering infrastructure (AMI), that modifies the power system data or state of the smart meters. A number of incidents of false data injection attacks have been performed by customers in USA, Ireland, Virginia and Hong Kong (Lo and Ansari, 2013).

Spoofing Attack: In this kind of attack, a new system element is added at one end that acts as a legitimate body (Fan et al., 2015).

DoS Attack: DoS attack is launched to flood any computer or network system with overwhelming packets through different sources or geographical locations to overflow the system buffer, thereby shattering the system and leaving it inoperable (Wang et al., 2017).

In addition, a new attack called the puppet attack on smart meters can cause DoS attacks in the metering networks (Yi et al., 2016). Security weaknesses of smart meters were discussed in the 2014 Black Hat Europe conference, where Alberto and Javier stated how an attacker can get access to the encryption keys (for e.g. a master key) by exploiting the hardware of the device (Illera and Vidal, 2014).

Attacks on smart grid seriously affect the entire ecosystems, such as smart home activities, industrial operations, hospital facilities, financial and government institutions. In 2014, an Australian utility company was seriously affected by the DDoS attack caused due to a misdirected command (Wueest, 2014). Also, the cyber-attacks on Ukrainian energy companies in 2015[1] and 2016[2] caused power blackouts in the region for several hours.

Contributions: The main contributors of this paper are summarised as follows.

- In this paper, we have proposed a novel unsupervised deep learning based Long Short time-Denoising Autoencoder(LSTM-DAE) model to detect anomalies in energy consumption data of smart meters. As a result, we have addressed the issue of real world anomaly detection. This would help in detecting the energy theft by customers, meter malfunctioning or third party attacks.
- Also, time series energy data can be appropriately handled using sequential model like LSTM(Long Short Time Memory). Since our model is built using LSTM and Auto-encoder, unlike other existing machine learning models used for anomaly detection in energy data, it is most befitting.
- Our model achieves an accuracy of 94.53% and false positive rate of 5.47%, thereby having better performance metrics than the existing models.

2 RELATED WORKS

In this section, we review some existing works (Nagi et al., 2009; Nizar et al., 2008; Yip et al., 2017; Li et al., 2020; Huang and Xu, 2021; Yip et al., 2018; Cui et al., 2021) related to the detection of anomalies in power consumption data of smart meters. They specifically focused on grid's electricity consumption data.

Yi et.al (Yip et al., 2017) designed a linear regression based detection model for energy theft and defective smart meter was used for detection of anomalies. The anomalies are considered coefficients to the power consumption values of users, sampled at different points of the day in the form of a matrix. However, the model shows numerical discrepancies whenever the rate of anomaly i.e. anomaly coefficient of a particular user vary throughout the day. They had used Irish Smart Energy Trial Data-set that was based on half-hourly samples. They acquired anomaly coefficients through t-statistics and p-values using Matlab's fitlm function. Though they introduced categorical values like off-peak and on-peak hours for coefficients of anomalies, it was not good enough to justify situations since anomalies can vary throughout different times. Also, the threshold set for anomaly coefficient to be anomaly rather than an outlier is not based on a robust mechanism since technical errors (Yip et al., 2018) or measurements errors from device can likely create the noise. Additionally, they did not provided any numerical measurements on the model's performance.

In (Yip et al., 2018) the discrepancy in numerical value of their previously mentioned LP model (Yip et al., 2017) was solved, by introducing Linear Programming where varying anomaly coefficients were considered that made the model more realistic. It further improved the threshold for anomalies from 0 to 0.05 on the same dataset. However, they still did not consider losses due to technical faults such as cables, transmission lines and distribution stations. Therefore, we still cannot rely on the improved threshold, which might not be reliable enough.

While, Li et.al (Li et al., 2020) proposed a blockchain based detection method in conjunction with unsupervised K-Nearest Neighbor(KNN) for clustering into three categories like working class, holiday class and outlier class. However, there is great uncertainty in the method of data collection from sensors and smart meters deployed by them in factories and homes. In addition, there was no justification for the selection of k value in the KNN algorithm. The concepts for relation between data using correlation coefficient and number of occurrences of data points using Poisson's distribution to address anomalies was appropriately evaluated. They neither provide a proper justification to distinguish anomalies from data-points that are simply outliers, nor deploy a robust mechanism against outliers and anomalies. The picture of their stated analysis is rather vague and thereby makes the detection rates unreliable.

Moreover, Huang et.al (Huang and Xu, 2021) used Stacked Sparse Denoising Auto-encoder for detection of data theft. The model is stated to be unsupervised with single labels of honest customers obtained from the Electricity Consumption Fujian, China dataset. However, we deem it appropriate to state that it is semi-supervised. The anomalies are obtained from the reconstruction error with a claimed optimal threshold. The threshold is set through the ROC Curve. The ROC curve in turn is dependent on the False positive rate and this is acquired from the test set which is inappropriate (Merrill and Eskandarian, 2020) because the threshold should have been determined from the training set. Consequently, we need a robust mechanism to determine thresholds and a better model for real time classification.

It is clear that almost all of them have used a

[1] https://ics-certus-cert.gov/alerts/IR-ALERT-H-16-056-01

[2] https://www.technologyreview.com/2016/12/22/5969/ukraines-power-grid-gets-hacked-again-a-worrying-sign-for-infrastructure-attacks/

supervised or semi-supervised framework for detection of data theft. However, supervised and semi-supervised machine learning algorithms cannot provide a good solution for real word scenarios.

3 Proposed Host-Based Intrusion Detection System

The Host-Based Intrusion Detection System (HIDS) is used for detecting abnormalities in smart meter energy consumption data. These abnormalities could be caused due to several reasons, such as energy theft, measurement errors, technical errors and/or faulty meters (Yip et al., 2017). Though majority of research works on anomaly detection have been carried out using supervised models, it is inappropriate since substantial amount of labeled anomalous samples is impractical. On the other hand, semi-supervised methods work a way around the requirement of labeled anomalous samples by completely relying on readily available normal samples. Thus they utilise data labeled as normal to detect anomalies and examples that do not comply with normal samples are simply flagged as anomalies. However, semi-supervised approaches are significantly susceptible to over-fitting or under-fitting,leading to poor recall or precision (Goldstein and Uchida, 2016).

This issue is daunting for all applications and specifically for grid data where we need very low false positives and false negatives (Mitchell and Chen, 2013). Since, the data may or may not contain anomalous samples. A more practical scenario is to use unlabeled data samples. As a result, unsupervised learning is the approach to achieve this (Merrill and Eskandarian, 2020). Therefore, we propose an unsupervised model for anomaly detection, relevant to any practical scenario. Our model is based on deep neural network using LSTM-AE .

3.1 Anomaly Detection Model

We present a LSTM-DAE model for anomaly detection in smart meter energy data. The model is inspired by the capability of LSTMs to predict time-series data and autoencoders in extracting features and reconstructing data as mentioned in (Huang and Xu, 2021). In our knowledge, this is the first work on anomaly detection of smart meter power data based on LSTM-AE, and significantly our approach is novel since it introduces a denoising LSTM-AE. The Denoising element is introduced to remove noise from the data in order to develop a robust Auto-encoder.

Structure of LSTM-DAE: We would first explain what is LSTM, Auto-encoder and Denoising Auto-encoder, separately.

1. LSTM is a type of neural network that was introduced to solve the vanishing gradient problem in RNNs. The vanishing gradient problem occurred when some of the weights ceased to change during the learning process. As a result, preference given to the current information would lead to neglection of past events. Therefore, the model cannot learn substantially in case of relations recurring over a long period of time. While, LSTMs were designed to control the entire information flow within neurons, through a gate that adds and deletes the information. Consequently, the model can learn long-term as well as short-term dependencies by controlling the process of forgetting unlike RNN. However, it limits the memory capacity in such a way that the output gate infers the updated cell state. It is particularly suitable for multivariate or univariate time-series data where it can be supervised or unsupervised (i.e. the data set can be with or without labels) (Lindemann et al., 2021).

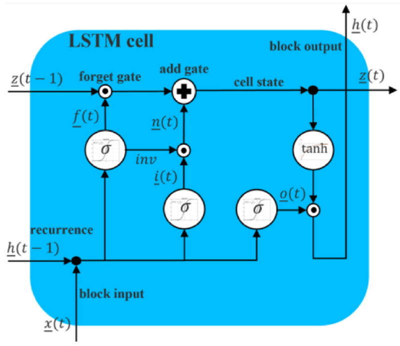

Figure 1: LSTM cell as designed by Hochreiter and Schmidhuber in 1997

2. Auto-encoders have been effective as unsupervised model for removal of outliers since they can reconstruct data efficiently with higher density. The neural network has two models called an encoder and a decoder that are trained together. The encoder compresses the initial input, thereby learning important features, while the decoder reconstructs the data from its compressed state. Therefore, the whole model can learn highly complicated data patterns (Merrill and Eskandarian, 2020).

3. When these Auto-encoders are fed with noisy inputs to reconstruct actual outputs,they are known as denoising autoencoders. These are more robust against noise and help prevent learning identity

function as in general autoencoders i.e. reconstructing X from X (Vincent et al., 2008). In this model, noise is added to the input X such that it constructs a clean output from the noisy samples i.e. $\hat{X}$. (Vincent et al., 2008). This corruption of inputs can be done in several ways such as by replacing 30% of the input values with zero, 50% of the inputs with zero, (Huang and Xu, 2021)using random noise or white Gaussian noise.

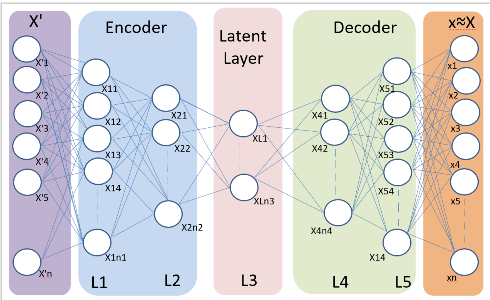

Figure 2: The proposed Denoising Auto-encoder where noise is added to the real inputs before feeding it to the model

According to Merill et.al (Merrill and Eskandarian, 2020), in a normal case of anomaly detection, the reconstruction error generated from any regular auto-encoder during prediction on test set is used for determining an anomaly score. Though this technique assumes that normal samples will be reconstructed more accurately than anomalous ones, it also assumes that anomalies cannot be reconstructed accurately. Often times, auto-encoders can actually generalize quite well to reconstruct anomalous inputs. However, this over-generalization cannot be solved through more regularization, capacity restriction or reduction in training time since all these might menace the reconstruction of normal samples.

Therefore, we proposed a denoising LSTM based Auto-encoder. This auto-encoder will detect anomalies in time series consumption data and would be robust against noise.

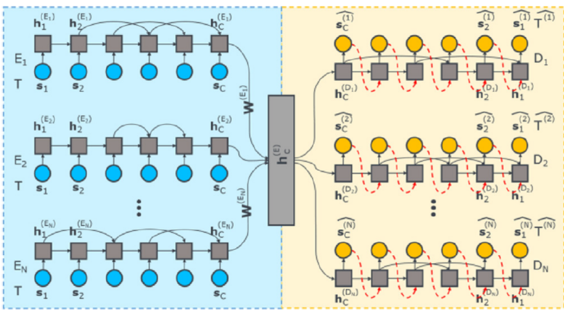

Figure 3: Pictorial depiction of an LSTM AE with skip connections according to (Kieu et al., 2019).

4 EXPERIMENTAL EVALUATION

In this section, we report the experimental findings on detecting anomalies using a smart meter use-case. We have put forward an unsupervised deep learning-enabled IDS to distinguish between the normal and anomalous behaviour in energy consumption patterns of households.

Datasets We use two different datasets for analysis as mentioned below: Our emperical evaluations are based on two different datasets present in the table.

Dataset	Period of Consumption
UCI Energy Consumption Dataset[3]	December 2006 to November 20
Irish Contracted Power Dataset[4]	1st January,2009 to June, 201

Table 1: Dataset description

The UCI Power Consumption dataset extracted from traditional meters was chosen to consider a diverse range of parameters, such as current, voltage and sub-meter data in addition to power consumption. Significantly, the dataset was unlabelled resembling any real-world dataset. However, we were unable to validate the performance of the model due to lack of a ground-truth. Therefore, we later used Irish power consumption dataset that consists of half-hourly smart meter data from honest customers only i.e. non-anomaly labels. We did not feed labels to our model but utilised the labels to calculate the various performance metrics including accuracy, false positives and false negatives.

4.1 Metrics

We define few metrics in terms of our work. It is important to understand those before we delve further into the experiment section, since it describes the way in which we have used them.

MSE- The MSE is the difference between the square of the actual and predicted data for all 'n' samples. We can represent it as

$$MSE = 1/n \sum_{i=1}^{n} (X_{actual}^2 - X_{predicted}^2) \qquad (1)$$

Model Loss- It is a scalar value that indicates how close are the predictions of the model to the actual labels. If the loss is low (ideally 0), the predictions are perfect, and close to 0 are good predictions; else if it is closer to 1, the predictions are bad.

Threshold- It is the range beyond or below which anomalies are flagged.

False Positive(fp): The number of samples that are non-anomalies while they are flagged as anomalies.

False Negative(fn): The number of samples that are anomalies, but are flagged as non-anomalies.

True Positive(tp): The numbers that state how many are samples are correctly predicted as non-anomalies.

True Negative(tn): It states how many samples are correctly predicted as anomalies.

Accuracy: It is the percentage of correct prediction of non-anomalies from the samples. It can be represented as follows.

$$Accuracy = (tp + tn)/(tp + tn + fp + fn) \quad (2)$$

4.2 Configuring Threshold

There are several methods for calculating the threshold set-up such as the use of ROC curve (Huang and Xu, 2021), $90-95\%$ on the training data (Givnan et al., 2022), mean and standard deviation method[5], Kentucky's method (Zhou et al., 2021). Though the threshold is very much dependent on the dataset, and each method might provide different results, we should choose a very robust threshold that would overcome the noise due to some outliers, such as measurement errors and technical errors. Therefore, we preferred Kentucky's method over the rest for a robust mechanism as stated in (Zhou et al., 2021). The threshold is calculated based on the training data, where we assume that the training data is not anomalous. The threshold is evaluated using Q1, Q2 and IQR. Q1 is the first quartile which means that it is the value under which 25% of data points are found when they are arranged in increasing order. Q3 is the third quartile which thereby, the value under which 75% of data points are found when arranged in increasing order. IQR is the inter-quartile range where

$$IQR = Q3 - Q1 \quad (3)$$

The formula for calculating the upper and lower thresholds respectively are as follows.

$$lowerrange = Q1 - 1.5 * IQR \quad (4)$$

$$upperrange = Q3 + 1.5 * IQR \quad (5)$$

4.2.1 Experiments on UCI Dataset

We train LSTM-AE and LSTM-DAE to compare the loss and reconstruction error for the same dataset. At first, we train the LSTM-AE for five epochs and it produces satisfactory loss value (loss value is 0.05). This is done for both training and validation set. The model is found to be a good fit since the plot of training set loss against validation seem to be converging towards the last few epochs. The loss values are substantially low indicating that the model is performing well in terms of learning.

Then, we train our model on the same dataset using noisy data. After training for 12 epochs, we observe satisfactory low loss value in the last few epochs. Thereby, indicating that the original data is recovered well from the noised input. Further, we use our LSTM-DAE model. The number of samples considered is 10,000 and that constitutes nearly 1 month of data. The model loss shows that it is considerably low i.e.,0.06, at only 12 epochs even with noisy input. Therefore, this illustrates a good learning capacity and efficient model performance.

The MSEs after noised inputs added to the training set acquired from Paris Power Consumption data, are in Figure 4.

Figure 4 shows the train MSEs on noisy inputs to the model using Paris power consumption dataset. The MSEs are low thereby, indicating that the model is predicting very well. After the reconstruction error is calculated from test set, we check if that error is beyond a selected threshold for the anomaly score. Further, the sample would be flagged as an anomaly if the error is beyond the threshold, otherwise the consumption pattern will be considered as normal.

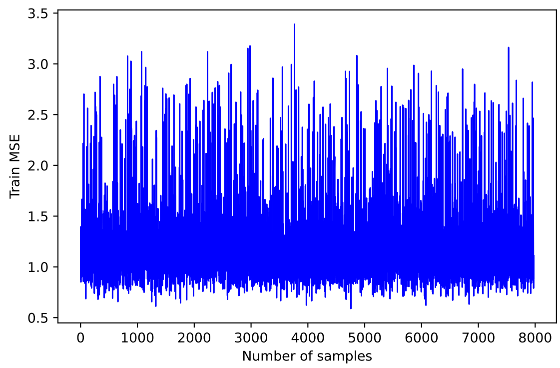

Figure 4: MSEs from the noisy input of Train set on Power Consumption Data-set

Finding Anomalies The test set for UCI dataset is predicted and the MSEs i.e., MSE per sample is calculated from the deviations of the actual test set.

Though, these MSEs in the test set are relatively higher than those in train set, they are still visibly low as shown in the y-axis of Figure 5. We tried to reconstruct the error through Keras' predict function in python. These errors are checked against the threshold. Thereby, the errors found below the lower and above the upper threshold limit are marked as anomalies. We find that out of 399 samples in the test set, 23 are flagged as anomalies. However, we are unaware, if the anomalies are correctly classified since the data-

[5]https://github.com/tensorflow/docs/

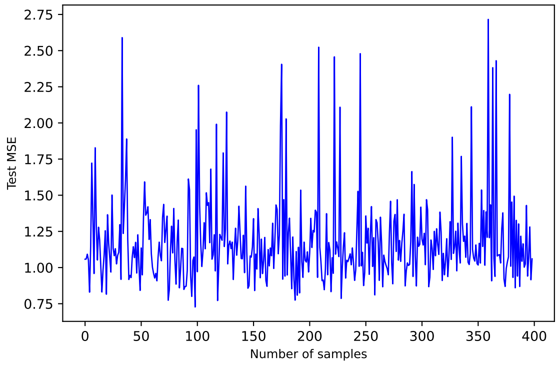

Figure 5: MSEs from the Test set on UCI Energy Dataset

set is unlabeled.

4.2.2 Experiments on Irish Power Data-set

Similar to the previous dataset, we calculate the loss for this dataset too using LSTM-DAE. The loss is found to be as low as 0.025 in this case, within just 14 epochs, thereby, yielding model consistency on low loss value and establishing the model as a good learner. Here, we have chosen 180 meters out of 6444. The data points involved with these 180 meters are 3,863,725. Therefore, we have performed our experiments on substantial amount of data rather than small to medium scale data and obtained satisfactory results on the learning capacity.

Initially, our training data is considerably less than the test data for this data-set. This is so because, we just wanted to validate the model performance in terms of reconstructing the loss and minimal error with relatively lesser data. Our model is trained using the first 30 meters ranging between 1000 and 1030 i.e. 7,00,000 samples, while our test data comprises of 150 meters i.e. 3,163,725 data points.

We plot the MSEs from training data for Irish Power Data-set(see figure 6). We find that the MSEs are relatively low as well.

Identifying Anomalies We acquire the test MSEs and anomaly scores for Irish Power Data-set samples having only healthy data. Our model is still essentially unsupervised since we train without these labels. However, we are able to use the labels for comparison after finding the anomalies. But, before acquiring the MSEs, we divide the entire test set having huge number of samples into chunks since we can achieve better visualisation with lesser data points. The MSEs for samples from meters ranging between 1031 and 1060 are low i.e. mostly within the range of 3.

In figure 7, it is seen that the samples for meters between 1061 and 1090 are mostly within the range of 3.5 and very few are beyond 8.The meters ranging between 1091 and 1120 too shows errors mostly within 3.5 and 4, while few are beyond 10. Similarly, meters between 1121 and 1150 have most of the errors in low range i.e. within 4. The last chunk for errors between 1151 and 1180 are around the range of 2 and very few are beyond 10.

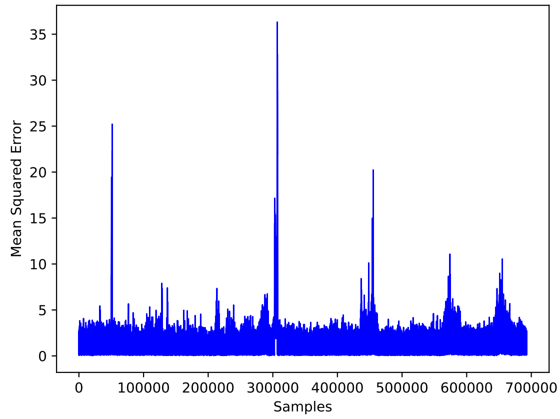

+

Figure 6: MSEs from the noisy input of Train set on Irish Power Consumption Data-set

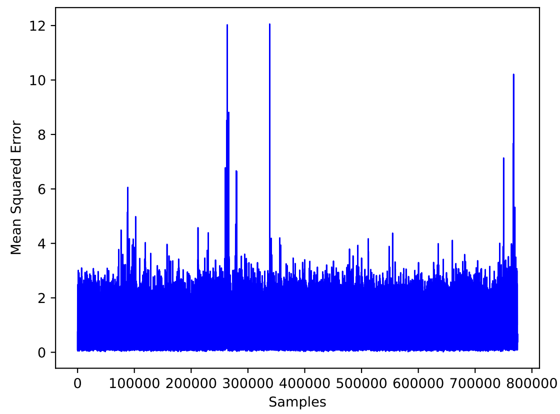

Figure 7: MSEs from the Test set obtained from Irish Power Consumption Data-set for meter samples 1061 to 1090

Therefore, with the MSEs ranging between 2 and 10, we conclude that the model performance stands out with relatively very less data for training in comparison to the testing set.

We obtained the anomaly score from training errors based on the fixed threshold for Irish Data-set. The lower and upper ranges of the threshold are -0.24896152299660124 and 1.3530767084315753 respectively. We find that 2,951,974 half-hourly data points from among 150 meters are marked to be non-anomalous.

We find that out of 3,871,203 data points, 211,751 points were marked as anomalies. Thereby, indicating the false-positives to be at 5.47%. The True Negatives i.e. non-anomalies, stand at 94.53%. As a result, the

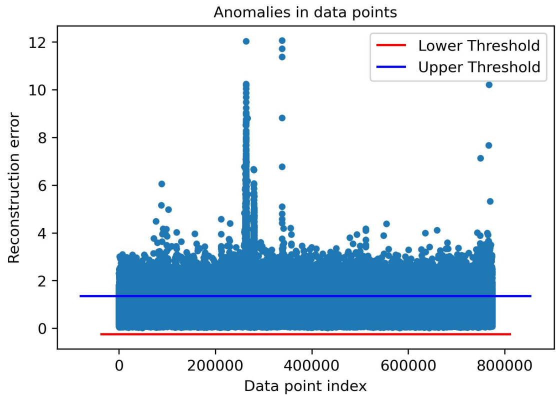

Figure 8: Anomalies from meter samples ranging between 1061 and 1090

accuracy of the model or detection rate is 94.53% with just 16.67% of data considered for training.

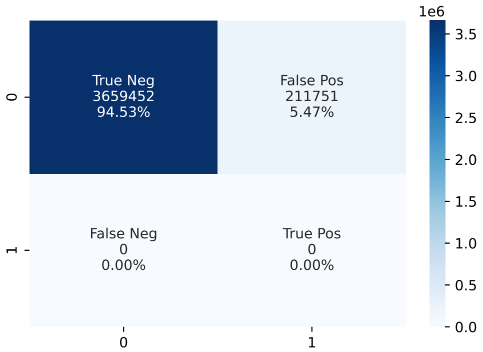

Figure 9: Confusion matrix for performance metrics based on Irish Power Consumption Data-set

4.3 Comparison and Discussion

Our model LSTM-DAE is compared to other models using two performance metrics i.e. accuracy and false positive rate. Here, we have made comparison with models like SSDAE, RDBN, PCA and SVM. We find out that our model has a higher accuracy level and lower FPR than the other models in Table 2, thereby outperforming others (Huang and Xu, 2021).

Model	Accuracy	FPR
LSTMDAE (Our Work)	0.9453	0.0547
SSDAE	0.9174	0.0719
RDBN	0.8701	0.1362
PCA	0.8582	0.1793
SVM	0.8176	0.1607

Table 2: Performance of various models on detection of anomalies in Energy consumption.

Our model LSTMDAE performs better than the ones listed in the table. This gives an indication that the model can be utilised for detecting anomalies and prove to be a good detector for smart meters.

5 CONCLUSION AND FUTURE DIRECTIONS

To conclude, we develop a robust unsupervised deep learning model to find out cohort anomalies in the power consumption data. We have considered every possible parameter to make sure that we secure our model against noise and flag the actual abnormalities. The model is reliably suitable for a real world scenario because of its unsupervised nature and has short inference time. Additionally, we consider the time series data factor through LSTM, unlike other proposed models. Therefore, it is a first of its kind for anomaly detection of smart meter data, keeping in mind their resource constrained nature. In the near future, we would focus more on the two categories of anomalies like anomalies due to faulty meter and anomalies caused by theft using LSTM-DAE.

REFERENCES

Cui, L., Guo, L., Gao, L., Cai, B., Qu, Y., Zhou, Y., and Yu, S. (2021). A covert electricity-theft cyber-attack against machine learning-based detection models. *IEEE Transactions on Industrial Informatics*.

Erkin, Z., Troncoso-Pastoriza, J. R., Lagendijk, R. L., and Pérez-González, F. (2013). Privacy-preserving data aggregation in smart metering systems: An overview. *IEEE Signal Processing Magazine*, 30(2):75–86.

Fan, Y., Zhang, Z., Trinkle, M., Dimitrovski, A. D., Song, J. B., and Li, H. (2015). A cross-layer defense mechanism against gps spoofing attacks on pmus in smart grids. *IEEE Transactions on Smart Grid*, 6(6):2659–2668.

Givnan, S., Chalmers, C., Fergus, P., Ortega-Martorell, S., and Whalley, T. (2022). Anomaly detection using autoencoder reconstruction upon industrial motors. *Sensors*, 22(9):3166.

Goldstein, M. and Uchida, S. (2016). A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PloS one*, 11(4):e0152173.

Hart, D. G. (2008). Using ami to realize the smart grid. In *Power and energy society general meeting-Conversion and delivery of electrical energy in the 21st Century*, pages 1–2.

Huang, Y. and Xu, Q. (2021). Electricity theft detection based on stacked sparse denoising autoencoder. *International Journal of Electrical Power & Energy Systems*, 125:106448.

Illera, A. G. and Vidal, J. V. (2014). Lights off! the darkness of the smart meters. *BlackHat Europe*.

Kieu, T., Yang, B., Guo, C., and Jensen, C. S. (2019). Outlier detection for time series with recurrent autoencoder ensembles. In *IJCAI*, pages 2725–2732.

Knirsch, F., Eibl, G., and Engel, D. (2016). Error-resilient masking approaches for privacy preserving data aggregation. *IEEE Transactions on Smart Grid*, 9(4):3351–3361.

Li, M., Zhang, K., Liu, J., Gong, H., and Zhang, Z. (2020). Blockchain-based anomaly detection of electricity consumption in smart grids. *Pattern Recognition Letters*, 138:476–482.

Lindemann, B., Maschler, B., Sahlab, N., and Weyrich, M. (2021). A survey on anomaly detection for technical systems using lstm networks. *Computers in Industry*, 131:103498.

Lo, C.-H. and Ansari, N. (2013). Consumer: A novel hybrid intrusion detection system for distribution networks in smart grid. *IEEE Transactions on Emerging Topics in Computing*, 1(1):33–44.

Merrill, N. and Eskandarian, A. (2020). Modified autoencoder training and scoring for robust unsupervised anomaly detection in deep learning. *IEEE Access*, 8:101824–101833.

Mitchell, R. and Chen, R. (2013). Behavior-rule based intrusion detection systems for safety critical smart grid applications. *IEEE Transactions on Smart Grid*, 4(3):1254–1263.

Nagi, J., Yap, K. S., Tiong, S. K., Ahmed, S. K., and Mohamad, M. (2009). Nontechnical loss detection for metered customers in power utility using support vector machines. *IEEE transactions on Power Delivery*, 25(2):1162–1171.

Nizar, A., Dong, Z., and Wang, Y. (2008). Power utility nontechnical loss analysis with extreme learning machine method. *IEEE Transactions on Power Systems*, 23(3):946–955.

Tellbach, D. and Li, Y.-F. (2018). Cyber-attacks on smart meters in household nanogrid: modeling, simulation and analysis. *Energies*, 11(2):316.

Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. (2008). Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103.

Wang, K., Du, M., Maharjan, S., and Sun, Y. (2017). Strategic honeypot game model for distributed denial of service attacks in the smart grid. *IEEE Transactions on Smart Grid*, 8(5):2474–2482.

Wueest, C. (2014). Targeted attacks against the energy sector. *Symantec Security Response, Mountain View, CA*.

Yi, P., Zhu, T., Zhang, Q., Wu, Y., and Pan, L. (2016). Puppet attack: A denial of service attack in advanced metering infrastructure network. *Journal of Network and Computer Applications*, 59:325–332.

Yip, S.-C., Tan, W.-N., Tan, C., Gan, M.-T., and Wong, K. (2018). An anomaly detection framework for identifying energy theft and defective meters in smart grids. *International Journal of Electrical Power & Energy Systems*, 101:189–203.

Yip, S.-C., Wong, K., Hew, W.-P., Gan, M.-T., Phan, R. C.-W., and Tan, S.-W. (2017). Detection of energy theft and defective smart meters in smart grids using linear regression. *International Journal of Electrical Power & Energy Systems*, 91:230–240.

Zhou, G., Liu, M., and Liu, X. (2021). An autoencoder-based model for forest disturbance detection using landsat time series data. *International Journal of Digital Earth*, 14(9):1087–1102.