**ORIGINAL RESEARCH** **Open access**

# AI-Enabled Energy Management in Manufacturing: A Practical Framework for Forecasting and Anomaly Detection

Andrei Popescu[1*], Mihai Ionescu[1], Elena Stan[2]

## Abstract

The urgent need to improve energy performance, as highlighted by the United Nations in its Sustainable Development Goals, underscores the imperative of reducing energy expenditure in industrial firms. Energy management structures are indispensable for meeting this objective by systematically singling out inefficiencies and surfacing prospective savings through the examination of energy usage records. With the growing deployment of Artificial Intelligence (AI), substantial scope exists to leverage cutting-edge data analytics to forecast future energy demand and identify irregularities. Nonetheless, extant scholarship gravitates toward the conceptual formulation and refinement of AI models. In contrast, the on-the-ground embedding of AI within energy management architectures in manufacturing enterprises remains limited, especially in small and medium-sized enterprises (SMEs). This case study presents a proof of concept implemented within a German manufacturing business to demonstrate how AI models can be integrated into energy management systems leveraging pre-existing data stores. Employing the Cross-Industry Standard Process for Data Mining (CRISP-DM) framework, an AI-grounded evaluative algorithm was crafted to detect energy changes in draw behavior and project future draw. The outcomes show that AI models such as Long Short-Term Memory enable the forecasting of future energy draw with striking accuracy, alongside the pinpointing of divergences that conventional systems may let slip. This investigation underlines the revolutionary capability of AI-fueled energy management systems to strengthen operational output and ease alignment with ISO 50001 stipulations. It provides a practical avenue for wider uptake of smart data analytics within energy management, with particular relevance for SMEs seeking to lay the groundwork for a more ecologically sound industrial base.

**Keywords** Anomaly detection, Machine learning, Deep learning, AI-driven energy management, Consumption forecasting, Peak-load reduction

*Correspondence:
Andrei Popescu
andrei.popescu@unibuc.ro
[1] Department of Artificial Intelligence for Smart Power Systems, Faculty of Engineering, University of Bucharest, Bucharest, Romania
[2] Department of Energy Optimization Engineering, Faculty of Engineering, Politehnica University of Bucharest, Bucharest, Romania

# Introduction

The planetary climb in energy utilization, driven by economic growth, demographic expansion and climate change, poses significant obstacles to reorienting our energy systems toward renewable sources and highlights the need to advance energy performance across many sectors [1]. In answer to these obstacles, the United Nations has instituted the Sustainable Development Goals (SDGs), designed to carve out a thoroughfare toward a durable and evenhanded future. Standing among these, the seventh goal (SDG 7) constitutes a bedrock of global energy policymaking. It works toward safeguarding inclusive access to cost-feasible, steadfast, sustainable, and up-to-date energy by the year 2030. This ambition includes substantially increasing the share of renewable sources, accelerating energy performance improvement and strengthening cross-border cooperation on clean energy engineering [2]. Given that over 75% of greenhouse gas

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

emissions are due to energy generation and utilization, significant performance gains and a reorientation toward renewable sources within this sector are vital for meeting this ambition [3].

Seen in this light, manufacturing and fabrication firms find themselves obliged to pare down their energy intake. One fundamental stride in this direction is the institution of energy management systems. These systems are designed to improve energy performance and identify latent savings by scrutinizing and weighing processes, systems, activities, products, and services in terms of their energy draw [4, 5]. Still, typical energy management systems often struggle to streamline the handling of voluminous data pools and respond nimbly to real-time swings in energy requirements [6]. Yet another stumbling block is the shortfall of decision-pertinent intelligence for leadership personnel. Whereas data points are at hand, they are recurrently left unparsed or incompletely circulated within energy management apparatuses. As a result, about 50% of economically sound energy efficiency initiatives across the industrial sector remain unexecuted [7]. These snags put the brake on the full actualization of energy management's promise within factory environments. Amidst the digital overhaul and swelling interconnectedness via the Internet of Things, the mobilization of Artificial Intelligence (AI) to chew through organizational data is taking on mounting consequence for vaulting these blockades. The integration of AI into an energy management framework enables streamlined, reliable data manipulation and analysis, with applications including forecasting future energy demand, flagging anomalies, and identifying discrete end-user load signatures [8-10]. That said, the concrete installation of AI-anchored energy management systems remains tentative. Enterprises, with special reference to small and medium-sized enterprises (SMEs), exhibit wariness about embracing such systems as a result of sundry hindrances, including an insufficiency of AI savvy, assumptions of towering expense, straitened means and patchy data underpinnings [11, 12]. Such hindrances throw into relief the necessity for tangible rollouts that lay out the way extant resources, exemplified by archival energy usage logs and present IT backbone, can be drawn upon to move past these roadblocks and smooth the path for folding AI into energy management ecosystems. Up-to-the-minute works, among them [9, 13, 14], zero in on the engineering and burnishing of AI algorithms, for instance neural nets, to elevate model competence and demonstrate technical practicability. Yet, these works habitually lack a hands-on execution playbook that offers counsel on how commercial actors can fruitfully embed these solutions into their energy management apparatuses, most pointedly within straitened-resource milieus.

Against this setting, the present case study probes the research question "How can manufacturing companies be empowered to effectively integrate AI-driven models into their energy management systems, leveraging existing resources, to reduce energy consumption?" To address this inquiry, the case study conducts a proof of concept within a German manufacturing enterprise in the mechanical engineering sector. The investigation aspires to ascertain whether AI models can be productively woven into energy management frameworks by tapping into historical electricity usage records. Two separate AI models – one geared toward anomaly identification and the other toward prediction – were put into operation to judge their capability for surfacing inefficiencies and anticipating future energy draw. The case study charts a pragmatic course for embedding AI-powered models into energy management and overcoming the hurdles connected with the lack of robust evaluation within standard energy management structures. This element is indispensable for reinforcing decision-making through targeted intelligence. This work enriches the sphere of AI-grounded energy management by delivering a workable and portable solution purpose-built for industrial milieus marked by resource scarcity. Unlike prior investigations that primarily focused on algorithm construction or performance benchmarking, this inquiry prioritizes operational applicability. Prediction and anomaly flagging were carried out via widely obtainable machine learning libraries and archival datasets. The deployment design is streamlined and cloud-dependent, eliminating any need for premium infrastructure or laborious data engineering. The realization follows the CRISP-DM schema and draws upon current IT environments without requiring deep-seated AI proficiency. First engineered within a confined pilot environment, the system was later validated in the live measurement framework of an active manufacturing installation. This modus operandi equips SMEs with a down-to-earth, resource-sparing corridor to ISO 50001-concordant energy appraisal without the requirement of weighty capital infusion into hardware or niche workforce.

# Background and Research Demand

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

The ISO 50001 norm [15] provides hands-on guidance for establishing an energy management system. It supplies an all-inclusive framework to implement measurable, parallel energy-saving steps within firms and entities, as well as attest to a conforming energy management arrangement. Implementing an energy management system aligned with ISO 50001 in manufacturing and fabrication enterprises is a powerful way to identify and measure energy flows, fine-tune energy use, and improve energy performance. The standard guides organizations through an all-embracing sequence of collecting and analyzing energy usage to pinpoint latent energy gains. The energy appraisal accordingly includes the analysis of energy use, energy consumption, and energy efficiency, together with identifying areas of significant energy use to pinpoint potential savings and follow-on interventions [16]. Since the ISO 50001 specification was introduced in 2011, the number of industrial facilities adopting energy management frameworks has swelled. The ISO tally for 2022 indicates close to 27,600 legitimate certifications worldwide [17], representing a notable augmentation when set beside the 20,260 certifications recorded in 2016 [18]. Hence, numerous corporations view energy management as an effective instrument to improve their production facilities and processes with a view to elevating energy performance. Beyond that, some enterprises are obligated to institute an energy management system. In Germany, this is stipulated by the Energy Efficiency Act, which mandates that firms whose average yearly aggregate energy intake exceeds 7.5 gigawatt hours implement an energy management system congruent with ISO 50001 [19]. That said, time-honored energy management architectures in production companies regularly wrestle with the overwhelming volume and complexity of information generated by industrial workflows [6].

What is more, energy management configurations ordinarily lean upon interval-based inspections and statistical number crunching, a state of affairs that may lead to the omission of instantaneous oscillations and an ensuing holdup in decision-making [6, 20]. This shortage of nimble modifiability and expandability curbs their punch. Uplifting energy productivity in industry thus emerges as a taxing pursuit in light of the pronounced entanglement of industrial energy webs, particularly bearing in mind the actuality that industry is presently undergoing a deep-seated reshaping to attain the target of carbon neutrality [5].

The expanding digital reshaping of the industrial domain and, most notably, the interweaving of AI signal a sharp pivot from customary routines toward more discerning, anticipatory, and self-regulating configurations [21, 22]. In consequence, the deployment of AI within energy management unlocks an entryway to streamlining energy draw, consistently projecting and examining energy uptake, and trimming expenditure [23, 24]. Disparate avenues for incorporating machine learning (ML) into energy management ecologies have been charted, commonly centering on three principal domains of deployment:

## Forecasting

Current-day literature indicates that the preponderance of AI rollouts in the energy realm gravitates toward forecasting and optimization undertakings, above all the extrapolation of upcoming power usage [25]. By projecting electricity uptake and comparing these projections against real-world readings, divergences or abrupt spikes in energy demand can be caught in the here and now [26]. Reliable load forecasting can consequently assist industrial operators in tuning production agendas around tariff peaks and troughs, raising power generation throughput and paring back energy expenditure through advance formulation of fabrication schedules and the avoidance of surplus or deficit production [27]. With respect to predictive dissection and the forecasting of energy data, deep learning (DL) constructs, exemplified by recurrent neural networks (RNNs), present themselves as the most hopeful pathways [13, 28]. Given that energy datasets are customarily arranged as chronological sequences, Long Short-Term Memory (LSTM), a distinct subtype of RNN, emerges as particularly promising and effective for predicting energy demand relative to established techniques [29].

## Anomaly detection

While forecasting aims to anticipate upcoming energy requirements and streamline tactical planning, anomaly detection focuses on identifying unexpected incidents or deviations from presumed energy consumption profiles [30]. Beyond forecasting, identifying atypical load patterns and consumption habits from past occurrences is crucial [31]. Anomaly

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

detection is a widely used method for pinpointing aberrations in archival data by isolating data points or spans where a machine's behavior drifts from customary signatures or the broader trend, potentially preventing energy dissipation and malfunctions [32, 33]. A spectrum of statistical and AI-rooted methodologies, among them machine learning, can be drawn on for this end. At this juncture, the weighing of sundry approaches against each other turns essential. To offer a case in point, Zangrando *et al.* [32] advance that classical ML methods can spot anomalies more robustly than neural-style pathways.

## Monitoring and optimization

In addition, within the realm of state surveillance, the adoption of AI-guided models has borne considerable fruit in foretelling impending equipment conditions, enabling the reduction of machinery idle periods and prospective detriments. For example, autoencoder-based structures enjoy broad application in rotating apparatus to register faint alterations in oscillation imprints, paving the way for preemptive intervention before costly malfunctions occur [34]. The engagement of AI in state surveillance showcases how live sensor data capture strengthens the flagging of functional inefficiencies and anomalies, as thrown into relief within energy-frugal manufacturing sequences [35]. These paths of attack can be ported to energy management architectures, where cognate methodologies can improve anomaly detection and set the stage for anticipatory energy refinement.

Taken as a whole, present-day investigative efforts cluster around the formulation and honing of bespoke AI algorithms and ML methodologies, such as classification algorithms [27, 36], neural networks [13, 29, 37], deep reinforcement learning [38], or combined strategies [27, 32], aimed pointedly at boosting model efficacy. Despite recent progress, material investigative gaps remain within AI-directed industrial energy management across the application landscape. A large share of models seldom leap functioning manufacturing installations [39] and possess scant substantiation under real operating conditions [40]. Adding to this, a central bottleneck lies in the reliance upon expansive, high-calibre, and multi-source data collections [41]. This undercuts the potency of high-level AI models, most acutely in prediction use cases where adapting to highly variable fabrication contexts proves essential. Similarly, fusing disparate data from varied sensors and older plant equipment throws up hurdles rooted in interoperability deficiencies, which obstruct smooth surveillance and fine-tuning [42]. These strictures land with particular weight upon SMEs, which labour under supplementary structural handicaps, including difficulties surrounding data soundness and the availability of orderly datasets [43], tight technological and human means, together with an absence of AI familiarity [44] as well as a shortfall of homegrown AI capability and hefty initial expense for AI rollouts [45]. Faced with limited data accessibility and a dearth of niche staff, SMEs routinely encounter obstacles in adopting advanced AI-based forecasting or streamlining tools, highlighting the need for more approachable and robust AI solutions within industrial energy management. While these pathways mark technical advancement, they often remain conceptual or are confined to standalone deployments, and there is sparse guidance on how manufacturing organizations can operationally embed these solutions within their production settings and existing data holdings. Thus, notwithstanding its promise, the real-world uptake of AI-anchored energy management architectures remains constrained across the industrial domain, and is felt especially sharply among SMEs.

Positioned against this canvas, the case study exhibits a practice-tested route for threading AI-steered models into the energy management apparatus of a German manufacturing enterprise by harnessing pre-existing archival electricity usage figures. With this aim, an evaluation algorithm is fashioned that marries forecasting with anomaly identification to trim energy consumption and, most notably, peak loads within factory surroundings. By anticipating energy draw and identifying divergences across both instantaneous and historical data feeds, the algorithm traces root causes of disproportionate energy utilization, enables peak-load trimming, and supports the implementation of targeted measures that simplify the decision-making sequence. This approach not only bolsters energy management capabilities but also advances the demonstration of energy-associated output in line with ISO 50001.

# Methodological Approach

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

To isolate divergences within energy usage signatures, this inquiry melds two models into an inclusive evaluation algorithm: an anomaly detection model directed at flagging past divergences and a forecasting model intended to project upcoming energy draw. The shaping and fielding of this algorithm demand a methodical and orderly procedural framework, which draws upon the processing steps of the Cross Industry Standard Process for Data Mining (CRISP-DM). CRISP-DM represents a firmly rooted, sector-neutral process paradigm devised for data mining endeavours and takes in six repeating phases: Business Understanding, Data Understanding, Data Preparation, Modelling, Evaluation, and Deployment [46, 47]. The first phase, Business Understanding, encompasses the current energy management customs within the body. Data Understanding involves statistical probing of the dataset, and Data Preparation is applied to convert raw data points into a workable structure for downstream AI models, for instance by engineering suitable features. Through the modelling phase, ML and DL structures are fielded to fulfil two overarching aims: anomaly detection geared to singling out outliers and the forecasting of forthcoming energy draw. Diverse models are subjected to inspection, undergo training, and are afterward polished and, at a subsequent juncture, assessed through appropriate evaluation benchmarks. In the final stage, the models are deployed on a cloud-based platform within the organization to ensure that the insights drawn from the analysis are available in real time. This enables rapid, evidence-based decision-making aimed at improving energy performance and ecological durability. After delineating the methodological framework, the next section will provide a fine-grained walkthrough of each phase.

# Development of an Evaluation Algorithm for Energy Consumption Data

## Business understanding

In the first phase, the emphasis is on understanding the operational imperatives, most notably the firm's current state regarding the analysis of energy consumption data. By deploying an energy management system, the organization can conduct detailed analysis and graphical rendering of energy usage information. To drive its energy output yet higher, the enterprise resolves to thread AI into its energy management infrastructure. The analytical outputs churned out by the AI-enhanced EMS are designed to undergird adherence to the ISO 50001 specification through furnishing verification of ongoing advancement in energy-linked output. The company sets out to meet this objective via its existing resource pool, for example, archival energy consumption records, without the necessity of heavy extra capital outlay.

## Data understanding

The data flows from a pilot setup, which functions as a testing ground for shaping and verifying the AI-rooted evaluation algorithm. The pilot setup corresponds to a production configuration taking in multiple amenities and pieces of machinery, mimicking the energy usage profiles of a characteristic manufacturing operation. At the start, energy consumption logs were compiled as a fixed chronological string and warehoused centrally on a server retrievable through a web-based gateway. Mechanized data capture was effected through a single principal electricity meter. The backbone includes a digital metering setup, under which electricity meters pass along readings via MBUS toward a central hub using protocol converters and Modbus TCP. This stream of information is then conveyed under secure conditions through a tiered firewall arrangement to a cloud platform, where it is served up in a structured, columnar format primed for additional dissection. The AI models were stood up and schooled employing standard Python libraries within a Jupyter Notebook workspace. This arrangement permitted straightforward probing and quick prototyping, aligning comfortably with the resource profile characteristic of SME environments. By using this configuration, the developed models and algorithms can be trialed in a hands-on setting before being rolled out at a broader scale. The focus in the following section rests on electricity draw figures, covering four years, collected from the principal electricity meter. The energy data are registered as time-sequenced streams and are initially warehoused on a locally sited server affording browser-mediated access as well as within a cloud-based repository. The data architecture consists of the energy draw stated in kWh alongside a paired time indicator. This supplies the underpinning for downstream dissection and model construction. In addition, climatic readings sourced from the internet-based databank Meteoblue were incorporated as an extraneous factor in this

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

analysis, since atmospheric variables like temperature can directly mould energy usage dynamics. Within a manufacturing building, for example, the demand for illumination may rise during brisk, heavily clouded stretches with scant natural light, translating to a higher electricity draw for electric lighting. To stay within the bounds of the organization's confidentiality and regulatory obligations, all data used in this case study has been normalized, and the broad energy data trends have been withheld from presentation.

The initial action when delving into energy consumption figures involves cultivating a deep familiarity with the dataset. This includes investigative data probing via statistical tools and graphical approaches, for instance, box plots, coupled with identifying directional drifts, periodic oscillations, and residual components to unearth latent structures within the records [46, 47]. Another operation is feature engineering. Within the present case study, extra attributes are constructed out of time markers embedded in the sequential records through recoding them into categorical variables, e.g., days of the week (Monday = 0 to Sunday = 6), the paired hours (running from 0 to 23) together with public holidays and their adjoining bridging days (true or false). To identify departures within the records or irregularities in usage signatures, it is vital to establish a benchmark condition that reflects customary energy consumption behavior. This was accomplished through statistical graphical methods along the lines of seasonal decomposition, a procedure that breaks time series data into building blocks of trend, cyclical pattern, and residual scatter. These yield an initial window into the unfolding across recent years, for example, for picking out re-emerging cyclical motifs and weekly frameworks of energy usage [48]. The seasonal decomposition lays bare a receding trajectory in energy draw across recent years. Beyond this, recognizable cyclical oscillations manifest, such as depressed draw during end-of-year public holidays. To carry the probing still further, the records were arranged by hour and inspected on a month-by-month footing and a day-by-day footing using box plots. Whilst merely faint distinctions emerged between month-to-month draw profiles, the daily granularity inspection exposed pronounced gaps between working days and weekend days. To take a case in point, weekend draw signatures diverge sharply from those captured on typical workdays. The box plot, stretching outward from the first quartile (Q1) through to the third quartile (Q3), with whiskers reaching to the outermost observations within the interquartile range (IQR = Q3 - Q1), makes it possible to isolate outliers [49]. Stated with precision, an outlier is taken to mean a data instance that lands markedly beyond the conventional span, be it underneath Q1 minus 1.5 times the IQR (lower outlier) or over Q3 plus 1.5 times the IQR (upper outlier). In **Figure 1**, the normalized energy draw over a sample day is depicted, with upper (markers ^) and lower outliers (markers v) cleanly set apart. This equips the models to latch onto particular dependencies and departures within energy draw signatures and furnishes a crisp graphical representation of how the data values fan out.

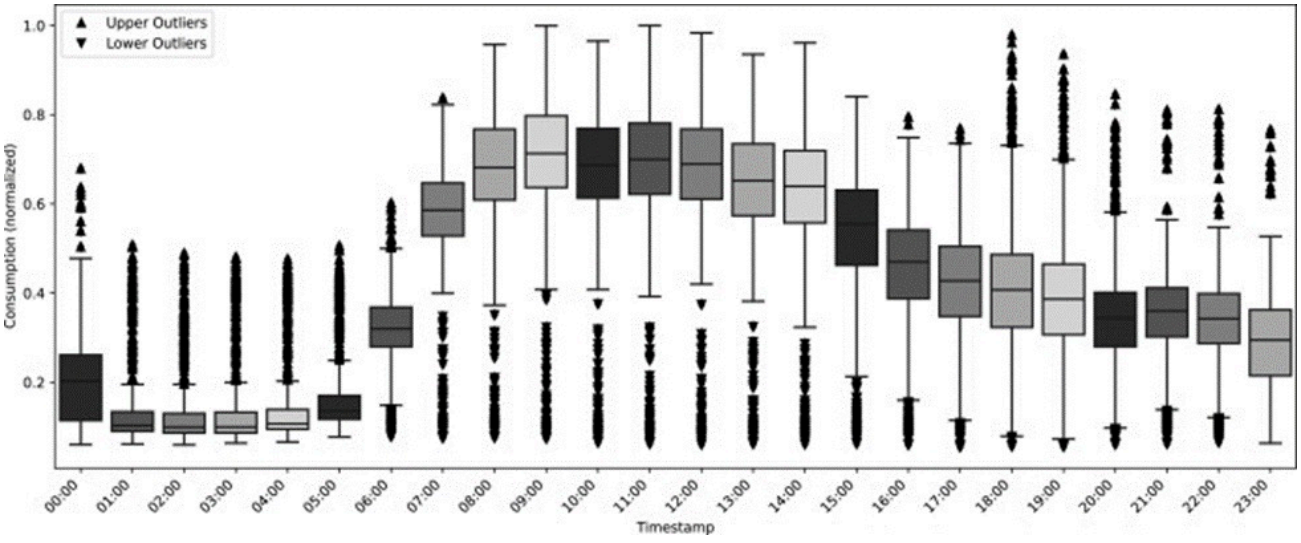


**Figure 1.** Identification of outliers for a sample working day based on boxplot analysis.

## Data preparation

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

The realizations pulled from the Data Understanding stage fed into the creation of annotated data pools geared toward outlier flagging. To train the models on a data compilation that reflects the particular draw rhythms of working days, the dataset was partitioned into weekdays (Monday to Friday) and weekend days (Saturday and Sunday). Weekend draw rhythms stand in material contrast when set side by side against those logged on working days. This arises because the firm under study operates at a lower tempo on weekends. Hence, the models set for deployment receive training on a dataset that includes working days and undergo separate training with data pools covering every Saturday and Sunday. Before placing the models into training, the dataset was cleaned and prepared to maintain uniformity. Records carrying incomplete particulars, among them absent hourly draw values or inaccessible climatic readings, were culled during the preprocessing pass. Given that the native measurements were available at minute-level resolution and subsequently rolled up to hour-wide windows for inspection, isolated absent values exerted no consequential pull on the aggregated outcomes. Because the overall portion of absent readings remained modest and no methodical voids turned up, imputation proved uncalled-for. To safeguard like-for-like comparability among attributes and to abet model convergence, every continuous numerical variable underwent scaling by way of the Min-Max scaler as realized in scikit-learn: X_std = (X - min(X))/(max(X) - min(X)), shifting each variable into a corridor bounded between 0 and 1, elevating comparability and underpinning convergence within the subsequent model training. Further, categorical variables derived from timestamps, for example weekday, month, or holiday indicators, were transformed through one-hot encoding to ensure congruence with the machine learning models while sidestepping ordinal distortion. The resulting annotated dataset **(Table 1)** groups outliers into two target variables: upper and lower outliers, rendered as Boolean values (True/False).

**Table 1.** Extract of the labelled data set based on data preparation.

| Date-Time Record | | Forecasted Consumption | Month Index | Day Index | Hour of Day | Holiday / Bridge-Day Indicator | High Outlier Flag | Low Outlier Flag |
|---|---|---|---|---|---|---|---|---|
| 2021–06–11 18:00:00 | | 0,19145 | 6 | 4 | 18 | FALSE | FALSE | FALSE |
| 2021–06–11 19:00:00 | | 0172862 | 6 | 4 | 19 | FALSE | FALSE | FALSE |
| 2021–06–11 20:00:00 | | 0116171 | 6 | 4 | 20 | FALSE | FALSE | TRUE |
| **Timestamp** | **Input Features** | **Output Variable** | | | | | | |
| Date-Time | Nominated Consumption | Month | | | | | | |
| 2021–06–11 18:00:00 | 0,19145 | 6 | | | | | | |
| 2021–06–11 19:00:00 | 0172862 | 6 | | | | | | |
| 2021–06–11 20:00:00 | 0116171 | 6 | | | | | | |

## Modelling and evaluation

The Modelling and Evaluation steps receive joint treatment inside this portion owing to their intimate coupling. This step focuses on integrating ML and DL architectures to identify irregularities in energy consumption data and sort anomalies that amount to departures from anticipated draw profiles, in addition to forecasting forthcoming energy draw. Continuous evaluation is vital throughout model construction for assessing output, fine-tuning model parameters, and implementing necessary course corrections to improve accuracy and robustness.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

## Anomaly detection for outlier classification

Based on the annotated dataset, ML models are trained to identify anomalies in future energy consumption logs reliably. Supervised learning (SL) architectures are particularly well-matched to this job, as they build on previously sorted training compilations that house both input signals and correct output values. This equips the system to churn out forecasts for datasets with unmapped output values, a capability drawn on above all for automated sorting. Representative SL pathways for classification embrace Logistic Regression, Decision Trees, Bayesian Models, and Support Vector Machines (SVMs) [50]. These four architectures were carried forward into the flagging of anomalies inside energy draw records, bearing in mind their attested particular potency in addressing sorting assignments and yielding sturdy forecasts over diverse data distributions [51]. In so doing, separate avenues get weighed against each other to zero in on the fittest algorithm for the particular deployment setting.

Before training the models, the energy draw data requires division into training and testing subsets. Because of the sequential nature of time series data, an arbitrary split falls short. A chance partition would throw cyclical swings or directional tendencies into disorder, since these temporal linkages would fail to be kept intact [52]. As a stand-in, options such as the scikit-learn TimeSeriesSplit function, which stems from k-Fold Cross-Validation, are commonly adopted. This scheme carves the data into k chunks, using k-1 chunks for training and one chunk for testing. This procedure sharpens model assessment by averaging the assessment yardsticks over k validation passes and sidesteps the call for a standalone validation dataset for hyperparameter calibration [53]. In addition to carving the dataset into training and testing pools, the class-spread weighting of the dataset calls for close attention to avoid heavy overfitting or underfitting. To accommodate the dataset's class skew, typified by markedly fewer data entries tagged as outliers (True) when compared to non-outliers (False), models are schooled with the scikit-learn parameter class_weight=balanced to underwrite evenhanded weighting across the classes.

The development of the four SL architectures relies on purpose-built modules in the scikit-learn library. In addition, a cyclical process of iterative enhancement was applied to each model. Additional attributes were trialed, including external climatic readings (temperature, degree of cloud cover), intended to amplify the models' forecasting muscle. Yet, folding in these extra weather-rooted attributes failed to yield any meaningful boost in model output. As a result, the archival energy draw data, coupled with timestamp-engineered attributes, emerged as sufficient for the onward dissection. After evaluating various attribute combinations, the most potent grouping of attributes was selected for the ongoing inquiry. After attribute selection, GridSearchCV was used to systematically tune hyperparameters for each model, aiming to improve model output, with particular emphasis on balancing model complexity and the risk of overfitting. For example, for the decision tree, the parameters Maximum Depth, Minimum Sample Split, Minimum Sample Leaf, and Minimum Impurity Decrease were run through a meticulous calibration sequence. Each model's output was appraised via an array of yardsticks, including accuracy, precision, recall, and F1-score **(Table 2)**, to support solid outlier sorting.

**Table 2.** Evaluation metrics of the implemented SL models für anomaly detection.

| Metric | Naive Bayes (Lower Outliers) | Naive Bayes (Upper Outliers) | Support Vector Machine (Lower Outliers) | Support Vector Machine (Upper Outliers) | Decision Tree (Lower Outliers) | Decision Tree (Upper Outliers) | Logistic Regression (Lower Outliers) | Logistic Regression (Upper Outliers) |
|---|---|---|---|---|---|---|---|---|
| **Classification Accuracy** | 0,51 | 0,72 | 1,0 | 1,0 | 1,0 | 1,0 | 0,89 | 0,83 |
| **Sensitivity (Recall)** | 1,0 | 1,0 | 1,0 | 1,0 | 1,0 | 1,0 | 1,0 | 0,74 |
| **Positive Predictive** | 0,02 | 0,17 | 0,98 | 0,97 | 1,0 | 1,0 | 0,01 | 0,21 |

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Value (Precision)** | | | | | | | | |
| **Harmonic Mean (F1-Score)** | 0,05 | 0,30 | 0,99 | 0,98 | 1,0 | 1,0 | 0,18 | 0,32 |

These yardsticks originate from the confusion matrix , which pairs forecast class tags with the actual ones. Under this framing, true positives (TP) represent correctly caught anomalies, while false positives (FP) denote ordinary instances mistakenly sorted as anomalies. In mirror image, false negatives (FN) constitute missed anomalies, and true negatives (TN) represent ordinary instances correctly tagged. Working outward from these tallies, accuracy conveys the overarching correctness of the sorting tool: Accuracy = (TP + TN)/(TP + TN + FP + FN), precision captures the share of forecast anomalies that hold genuine: Precision = TP/(TP + FP), and recall gauges the fraction of factual anomalies the model succeeds in catching: Recall = TP /(TP + FN). The F1-score knits precision and recall into a unified figure to capture the balance between false positives and false negatives: F1-Score = 2x Precision x Recall/(Precision + Recall). Beyond this, the separation frontiers for sorting upper and lower outliers received graphical treatment to stack the outcomes of the separate models side by side. **Figure 2** delivers, by way of a representative instance, the separation frontiers for upper outliers as isolated by the decision tree architecture. The softly shaded patches mark zones grouped as upper outliers, clearly showing how the model discriminates between normal and anomalous draw signatures at different hours of the day.

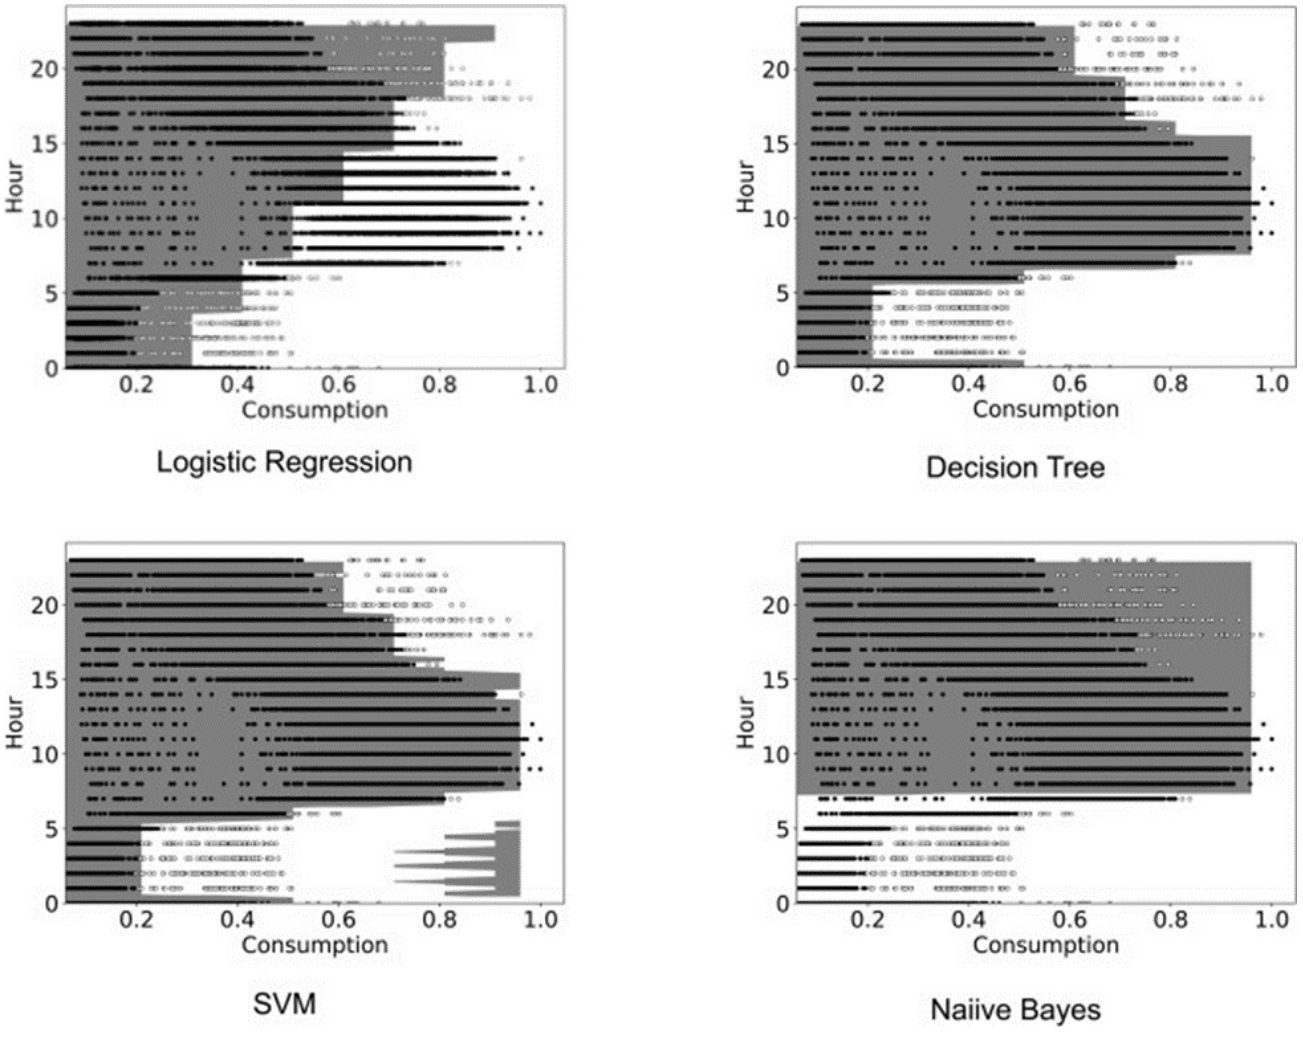


© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

**Figure 2.** Decision boundaries for upper outliers.

The Logistic Regression formulation yielded a substantial number of false positives, that is, records mistakenly labelled as irregular. This skew between recall and precision, as reflected in the disappointing F1-Score, hints at a model calibration ill-suited to the nuances of outlier discrimination within energy usage logs. The graphical outputs back this verdict, revealing a separation surface that falls short in demarcating the irregular class from the regular one. Along much the same lines, the Naive Bayes formulation's separation surface proved incapable of cleanly cleaving irregular points away from standard readings, producing an elevated false-positive frequency. Such erroneous labelling meaningfully erodes both precision and the F1 measure, pointing to a formulation that may lean toward excessive reactivity to the underrepresented outlier category.

In contrast, the SVM, deployed with a radial basis function kernel to circumvent a purely linear separation surface, was found to deliver a strong level of correctness alongside a high identification rate, coupled with adequate moulding to the separation surfaces. Even with these strong suits, the confusion matrix exposed a propensity to mis-tag regular records as irregular, prompting unease regarding the formulation's specificity and its likelihood of triggering nuisance alarms under operational conditions. This constitutes a meaningful drawback, as it undercuts the formulation's practical value in applied contexts. The decision tree formulation delivered a best-in-class showing in irregularity flagging. Its separation surfaces could latch onto irregularities with sharp precision, securing top-tier results on every evaluation yardstick. It delivered elevated precision, recall, and F1 measures, highlighting its potency in tackling the non-linear spread characteristic of energy usage logs. This strong showing positions the decision tree as well-suited for flagging irregularities within intricate data settings.

The endpoint of these undertakings produces a decision tree formulation that is meticulously dialled in and rugged enough for the oversight and probing of energy usage logs. The formulation's aptitude for detecting irregularities bears considerable promise for improving operational practices within energy management infrastructures. **Figure 3** portrays how the Decision Tree sorts irregularities within actual energy usage logs. In the depicted case, upper outliers stand marked from 01:00 to 06:00, specifically those moments where the energy draw sat in excess of the forecast bands at the respective clock positions. This graphical illustration shows how effectively the formulation separates standard readings from irregular ones while maintaining clear interpretability. The flagging and graphical rendering of such irregularities is purposed to aid the tracing of their root origins within the body and to clear the path for the rolling out of precisely aimed countermeasures to stamp them out.

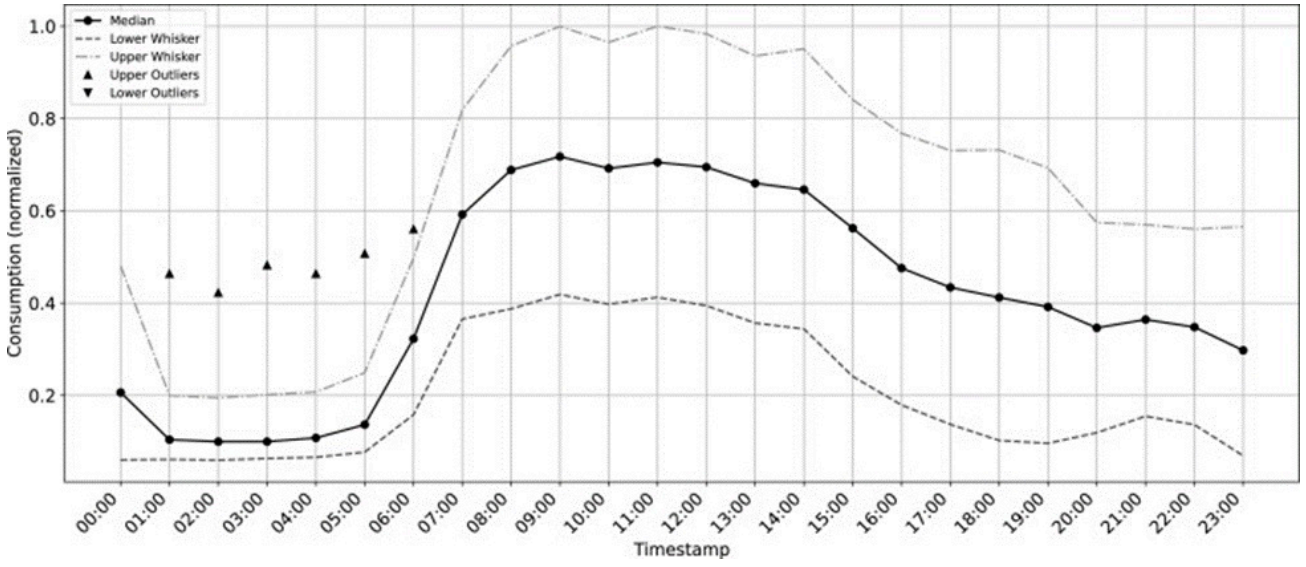


**Figure 3.** Anomaly detection based in the decision tree classifier for one exemplary workday.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

## Forecasting of future energy consumption with LSTM

Beyond the scope of irregularity flagging, which spots outliers entombed in archival records, snaring departures as they happen stands as indispensable for thoroughgoing probing. To serve this end, a predictive model was fashioned to permit ongoing juxtaposition against live feeds to catch divergences. As sketched in **Figure 4**, the goal is to set live readings (dashed trace) alongside projected figures (solid trace) to pick out divergences. For purposes of illustration, where live readings push past the projected figures over a given stretch, such an event might well betoken inefficiencies or unusual operating patterns, demanding prompt remedial intervention. This pathway clears the way for unbroken oversight and swift flagging of anomalies in energy usage. As earlier signalled in Chapter 2, LSTM architectures rank as the most capable for taming the intricacy of rhythmically recurring energy datasets and far surpass conventional avenues in predictive exactness [29]. This underscores their ruggedness in capturing stretched-out temporal motifs, which sit at the core of effective energy oversight and prediction. An LSTM is accordingly used in the section that follows to enable the projection of energy usage logs.

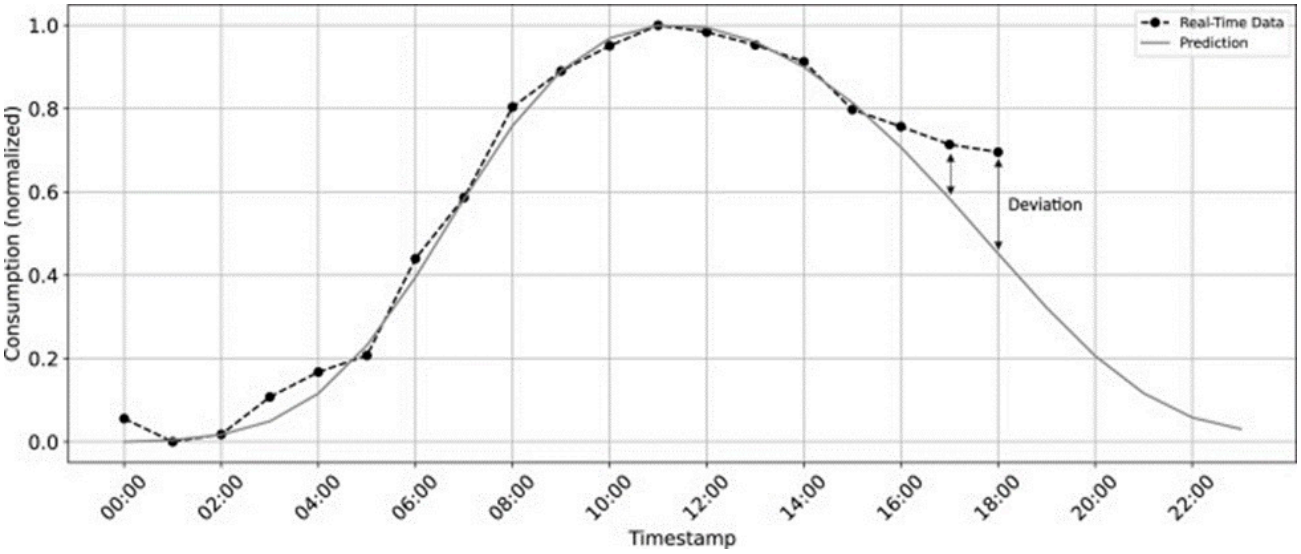


**Figure 4.** Comparison between predicted and actual values to detect deviations in real-time.

The broad procedural flow mirrors that of the ML formulation construction laid out. The formulation draws on energy usage logs and supplementary time-derived attributes, among them month, day, and hour, serving as the predictor set (X). The outcome variable (Y) likewise corresponds to energy usage, standing in for the figure slated for projection at the next temporal increment. Predominantly, the dataset underwent division into training and validation partitions, respecting the architecture of the chronological sequence. At this point, the first 80% of the dataset is allocated to training and the remainder 20% to testing. To condition the dataset for ingestion by an LSTM, the chronological logs were reshaped into tensors. Under this scheme, the tally of instances, the tally of temporal increments, and the predictor columns together form the predictor set (X), just as the tally of instances and the projection span form the outcome set (Y) [28]. When projecting chronological sequences through an LSTM, the predictor set used to estimate the value at the next temporal increment consists of readings originating from antecedent temporal increments. As a concrete illustration, 168 temporal increments (denominated in hours, corresponding to a full week) are mobilized to gauge the condition of the outcome variable at a temporal increment (a single hour) situated ahead.

The next step was to nail down the architecture of the LSTM. The LSTM subsumes an input layer that takes in the conditioned chronological sequences, buried LSTM layers that extract the data's temporal interdependencies, and an output layer entrusted with projecting forthcoming energy usage readings. Settling upon the proper tally of buried layers and neurons, alongside picking hyperparameters governing the training routine, represents a task contingent on the input dataset and must thus be resolved on a bespoke footing [54]. In keeping with this, to pin down the architecture for an

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

opening baseline formulation, established pathways for demand projection, advanced in Abbasimehr *et al.* [55], Bouktif *et al*. [56], and Mahjoub *et al*. [57], were stepwise reshaped via trial and error. This yields a baseline LSTM with 16 neurons in the input layer, a dense layer with 32 neurons using the ReLU activation function to condition the dataset, and an output layer with a single neuron whose job is to project the load readings. To assemble the formulation, the mean squared error (MSE) was appointed as the loss function and Adam as the optimiser. In addition, the mean absolute error (MAE) was used as a metric. The formulation then went through training stretched over 100 epochs and 128 batches. A first-pass graphical side-by-side already displayed a tight alignment **(Figure 5).**

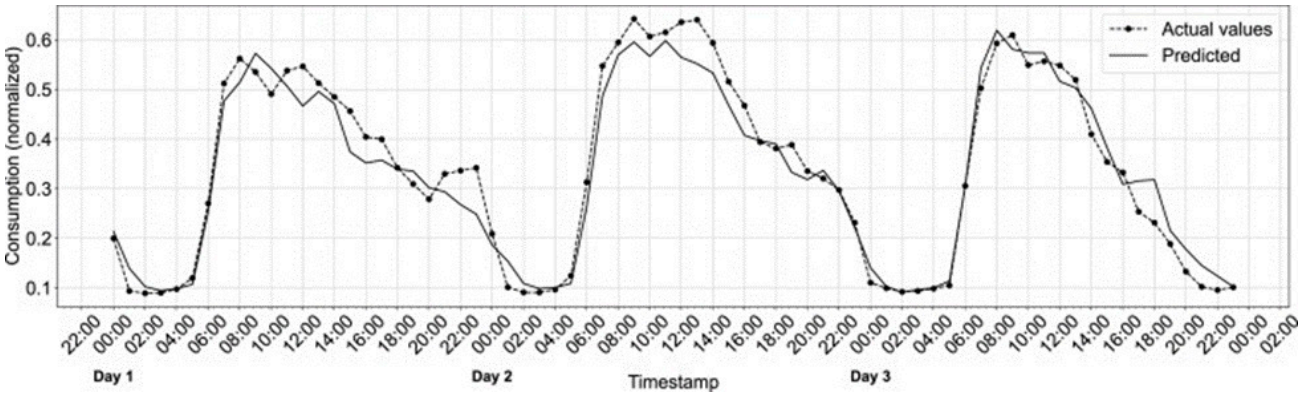


**Figure 5.** Predicted values (test data) and actual values (validation data) over three days.

To corroborate the baseline formulation, the training trajectory was carefully observed to identify possible underfitting, overfitting, or a satisfactory fit. This appraisal drew on the evolution of the loss function across epochs for both the training and the validation data partitions. As shown in **Figure 6** (left-hand panel), the loss function subsides as the epoch tally climbs for both training and validation data, signalling a pronounced moulding to the training dataset. That said, the validation loss tapers off at a sluggish tempo. It settles a shade above the optimal level, suggesting that the formulation has yet to reach full proficiency at teasing out motifs that port well to unexposed data. To gauge the formulation's showing in quantitative terms further still, several yardsticks are put to use, taking in the R-squared ($R^2 = \sum_{t=1}^{t=n}\left(y_i^{'} - \overline{y}\right)^2 / \sum_{t=1}^{t=n}\left(y_i - \overline{y}\right)^2$), Mean Absolute Error ($MAE = \frac{1}{n}\sum_{t=1}^{t=n}\left|y^{'} - y\right|$), and Mean Squared Error ($MSE = \frac{1}{n}\sum_{t=1}^{t=n}\left(y^{'} - y\right)^2$). The yardsticks for the baseline formulation are as follows: $R^2 = 0.941$,

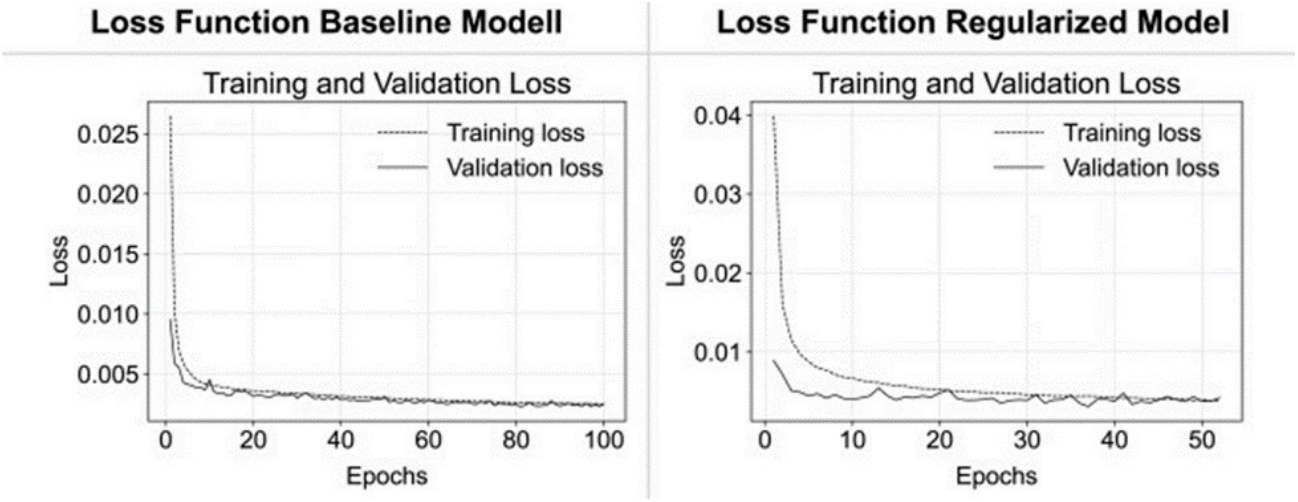


**Figure 6.** Loss function of baseline model and improved model.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

MAE = 0.033, and MSE = 0.050 (yardsticks computed on the normalized dataset). The MAE reading translates into an average absolute departure hovering near 10.4% of the median hourly draw (MAE/average consumption per hour), which serves as a marker of the formulation's respectable exactness. An $R^2$ reading approaching unity points to a sturdy moulding to the dataset, albeit the unfolding of the training and validation loss traces raises the spectre of overfitting. This faint overfitting hints that while the formulation adheres tightly to the dataset at hand, its capacity to extend to fresh, unexposed data may be somewhat weakened.

To bolster the LSTM formulation's predictive exactness and capacity for generalization, the following higher-order techniques were brought into the fold:

• Dropout: Operationalized to keep overfitting in check, dropout randomly silences a portion of neurons during training passes. This pruning of formulation intricacy prevents the formulation from fastening itself excessively to the training dataset [58, 59].

• Stacking: Now that overfitting can be warded off, it becomes feasible to expand the capacity of the CNN. Such expansion customarily proceeds by boosting the neuron count within the layers or by introducing extra layers. Stacking recurrent layers is a time-honoured avenue for building more powerful RNNs. Augmenting the formulation's capacity entails appending further neurons or layers [58].

• Bidirectional LSTM layers: Layers of this kind chew through the dataset along both orientations, forward and backward, and grant the formulation the ability to absorb cues from both antecedent and forthcoming surroundings. This double-pass handling materially lifts the formulation's aptitude for anticipating forthcoming events by leaning on a full sweep of temporal intelligence. This double-orientation handling materially lifts the formulation's aptitude for anticipating forthcoming events by leaning on a full sweep of temporal intelligence [58].

The baseline formulation is accordingly reshaped to incorporate a bidirectional LSTM layer operating as the input layer, working through the predictor set in both forward and backward orientations. The dataset next passes through further handling inside a supplementary LSTM layer. A dropout along with a recurrent dropout is affixed to each of the two LSTM layers. In addition, Keras callback mechanisms were used to steer the training routine adaptively. Of special note, early stopping wards off overfitting by bringing training to a standstill when validation loss reaches a plateau, while model checkpoint preserves ground gained for uninterrupted resumption. Dynamic learning rate adaptation hones convergence, lifting formulation pliability.

The formulation enhancements (Regularized Model) led to only modest shifts in the yardsticks. Be that as it may, the graphical portrayal of training and validation loss (**Figure 6** on the right-hand panel) reveals a reduced tendency toward overfitting and a sharpened efficiency in memory consumption. Opening trial cycles yielded encouraging outcomes ($R^2$ = 0.938, MAE = 0.035, and MSE = 0.051). The faint dip in performance yardsticks registers as a standard by-product when layering on regularization measures, as these measures dull the formulation's sensitivity to narrow training dataset motifs in the service of improved generalization. On public holidays, by contrast, when the underlying dataset would be presumed to register depressed draw, a standard trace mirroring a typical workday materialized in its stead.

In consequence, to hone the formulation's exactness when projecting energy usage on public holidays and their adjacent bridging days, these are brought in as supplementary predictor columns (Date-Responsive Model). The stitching in of the supplementary predictor column shifted the yardsticks to only a negligible degree ($R^2$ = 0.937, MAE = 0.038, and MSE = 0.054). Yet, by taking on board the supplementary predictor columns, the formulation becomes capable of dependably projecting a dip in electricity draw on a standard workday as well as on public holidays. **Figure 7** depicts the incremental honing of the LSTM. The actual unfolding of a public holiday across 24 hours is shown to convey a fuller sense of the magnitudes involved. The Date-Responsive Model lays out that public holidays at this stage are likewise registered and projected. As a culminating step, the Date-Responsive Model was trialled against an authentic dataset. This setup calls for the provision of a complete week, composed of 168 temporal increments, together with the predictor columns

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

consumption reading, month, weekday, hour, and holidays, to project a forecast reading positioned just beyond the tail end of the input data window.

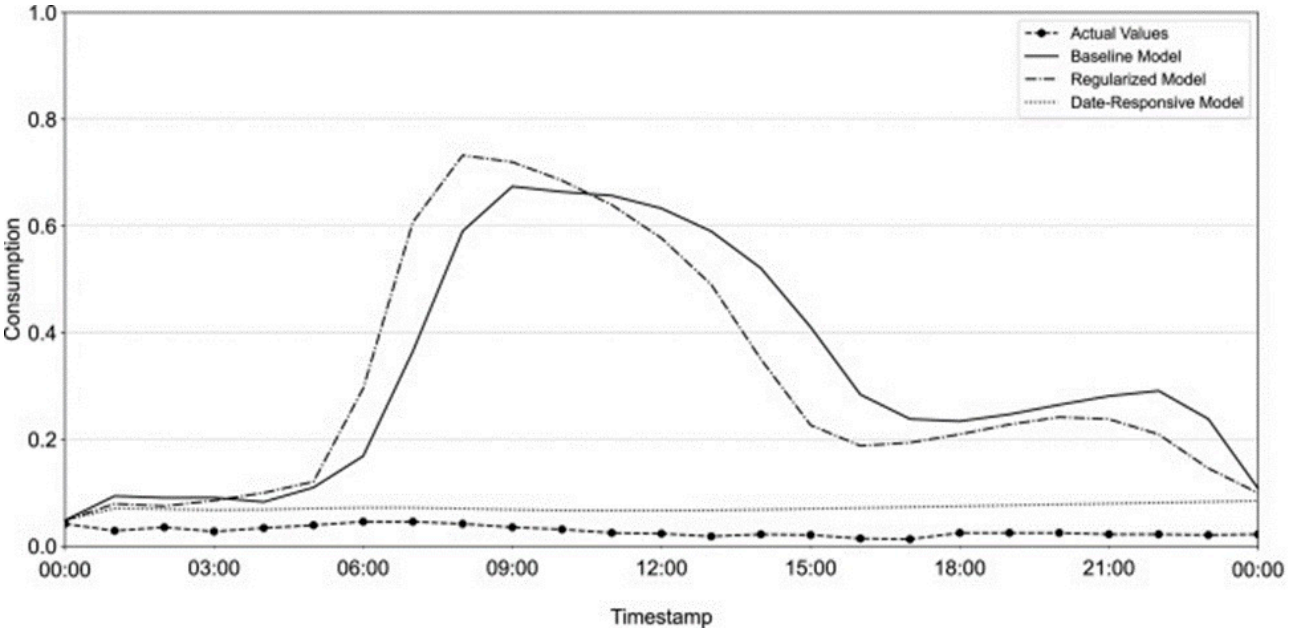


**Figure 7.** 24h-Forecasting of an exemplary public holiday.

To smooth the generation of projections that stretch out over a horizon of 24 h or longer, the input data window requires ongoing widening. One should remain mindful, nevertheless, that as the projection horizon lengthens, most pointedly over protracted spans, inaccuracies display a propensity to accumulate. Within the setting of projecting a full week (meaning, absent recourse to actual metered readings, relying exclusively on the formulation's projections inside the 168 temporal increments), a mounting error as time moves forward sits within reasonable expectation. On this account, within real-world deployments, it makes practical sense to supply the formulation with live readings at steady intervals and to confine projections to more abbreviated spans.

To situate the predictive punch of the LSTM formulation inside a broader landscape, a stripped-down persistence formulation was stood up as a reference point for head-to-head measurement. This formulation projects each temporal increment by directly reproducing the most recently sighted draw reading, a scheme that constitutes a broadly adopted benchmark in chronological sequence forecasting [60]. When assessed on the same test dataset, the persistence reference point attained passable predictive output ($R^2$ = 0.89, MAE = 0.0403, and RMSE = 0.0662), which suggests that the dataset possesses strong temporal continuity. Even so, the LSTM formulation ($R^2$ = 0.937, MAE = 0.038, and MSE = 0.054) outperformed the persistence formulation, securing better accuracy. Moreover, the LSTM architecture, especially when combined with date- and holiday-linked predictor columns, exhibited heightened robustness by reliably capturing draw anomalies linked to public holidays and off-standard load profiles—motifs that a simplistic persistence approach cannot intrinsically encapsulate. In light of this, the LSTM-anchored forecasting approach provides sharper overall predictive accuracy and greater practical utility for anticipatory energy management.

## Deployment

The engineered models are readied for real-world utilisation and stitched into a cloud-based infrastructure. This unlocks broad deployment across the enterprise and ensures that outputs remain on tap for steering decisions. The stage encompasses recording the findings, formalising operational workflows, and instituting feedback circuits to guarantee ongoing refinement. Standing up the AI-steered models atop a cloud framework demands a steady stream of current data, underscoring the imperative of fixed cadences and hands-off data harvesting. A routine is configured to draw down

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

fresh records every hour from a server-housed repository and channel them through the models stored in the cloud. To safeguard model fidelity over time, a benchmark will be instituted that focuses on cyclical retraining against newly arriving data to mirror present-day conditions with precision.

# Discussion

By way of realising a proof-of-concept within a German manufacturing enterprise, this case study lays out a practice-grounded route mapping how AI-steered models can be fruitfully embedded into energy management frameworks. Drawing upon existing assets, including archival energy draw records, and following a methodically ordered procedural path, the constructed evaluation algorithm tackles the hurdles that confront manufacturing bodies, most pointedly SMEs, such as constrained resource pools, e.g., the reachability of supplementary data channels. Over and above this, the models were built without requiring advanced AI coding fluency, showcasing that the barrier of narrow technical command can be meaningfully surmounted through the use of relatively uncomplicated AI models sourced from the ML libraries Scikit-learn and Keras. Embedding fairly straightforward AI models, including decision tree-based anomaly flagging alongside LSTM-driven forecasting, smooths peak energy draw and provides higher-level analytics that equip decision-makers to pinpoint inefficiencies and enact precisely targeted steps toward ongoing improvement. Most strikingly, the deployment of the models has shown a notable ability to project future energy draw and catch consumption anomalies before they mutate into mounting expense or systemic breakdowns.

To gauge the real-world potency of the models devised, further corroboration was undertaken against a holdout dataset covering six months. Within this interval, 125 upper outliers were flagged. These were delineated as draw incidents that overshot the hourly upper cut-off extracted from the IQR. Working alongside energy management professionals at the firm, these irregularities were rated as preventable through timely spotting and operational response. Based on the delta between the observed draw logged during these atypical incidents and their respective upper cut-offs, the total potentially preventable energy consumption reached approximately 3.0% of the aggregate draw across the half-year observation window. Even though precise draw figures remain undisclosed owing to confidentiality strictures, a 3 % curtailment equates to roughly 30,000 kWh for each gigawatt-hour drawn—amounting to, per the judgement of the firm's energy management unit, a commercially defensible case for ongoing rollout. Coupling the anomaly flagging formulation with the LSTM forecasting formulation further amplifies its operational worth. Through the sustained juxtaposition of forecast draw against live metered observations, the forecasting formulation serves as an advance-warning mechanism. Put with precision, whenever live draw observations persistently outstrip the LSTM projections, system overseers stand to be flagged at once. This forward-looking posture enables overseers to recognise and neutralise irregularities while still incipient, thereby markedly boosting the probability of averting outsized draw incidents before they mature into discernible outliers. The joint application of anomaly flagging and forecasting hence delivers a sturdy foundation for hands-on energy management, well-suited to achieving tangible energy savings within verisimilar industrial settings.

Projecting draw and its variation not only yields standards-compliant substantiation of energy-linked performance but simultaneously clears the path for instantaneous scrutiny by surfacing it on a cloud-hosted platform. This likewise eases the swift tracing of origins whenever departures or material overdraw materialise. The models deliver manifold strengths: Anomaly flagging pertains to the selective logging of departures. Outliers are sorted across the span of the day, with the timestamp of each outlier's surfacing provided in every case. The projection of draw data sets the stage for a side-by-side weighing between the forecast and the observed draw readings. In this way, state shifts can be rendered graphically over time, paving the way for the early detection of overconsumption. For example, inferences can be drawn for the corresponding load where observed archival readings outpace forecast readings over successive hours.

As a concrete example, within the pilot setup, an irregularity came to attention connected to choked pumps in an installation, entailing surplus energy requirement rooted in operational friction. The intelligence is leveraged to identify aberrant conduct, discern potential savings, and implement appropriate countermeasures. More pointedly, the scrutiny uncovered that peak energy draw habitually clusters within the hours of 09:00 to 13:00. This insight highlights the value of

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

probing load behaviour within this band and implementing targeted steps to streamline energy deployment. For example, fabrication operations might be shifted into the afternoon to realise more level capacity loading, evening out the overall load signature and trimming peak demand tariffs. These outcomes highlight the feasibility of this approach and its capacity to pare back energy draw within production operations. The integration of AI-steered models into energy management architectures has been shown to support sustained energy-linked output in alignment with ISO 50001. By isolating peak loads and capturing irregularities, these models enable enterprises to take precisely targeted steps to reduce energy consumption, thereby aligning with the Plan-Do-Check-Act cycle. In the “Check” phase, AI models play a pivotal role in tracking energy-linked output and identifying areas primed for streamlining. Within the downstream “Act” phase, the intelligence drawn from AI-steered scrutiny shapes the enactment of measures intended to streamline energy-linked output in conformity with ISO 50001.

Having noted this, the integration of AI-steered models into energy management is not without obstacles. One constraint lies in the reliance upon archival draw records as the sole data source, which, while adequate for sweeping energy performance uplifts, stands in need of supplementary data channels for higher-resolution probing, of the sort that factor in fabrication timetables or workspace staffing levels. Over and above this, uninterrupted improvement of the AI-steered models remains a necessity to ward off overfitting and attune the models to fluid conditions. To bring matters to a close, data safeguarding and privacy considerations demand due heed, above all when threading AI-steered models into the business-critical industrial backbone.

# Conclusion

The present inquiry bears out that AI-steered models can be productively threaded into the energy management frameworks of manufacturing enterprises, even those operating under slender data means, through leveraging in-place archival energy draw records and pursuing a methodically ordered procedural design. The developed AI-grounded evaluation algorithm, which fuses anomaly flagging with the projection of forthcoming energy draw, presents substantial promise for improving energy management setups by enabling precise and instantaneous dissection of energy draw logs. The establishment of anomaly flagging formulations, most notably the decision tree, is potent in latching onto intricate motifs within energy consumption. In parallel, the LSTM formulation exhibits the capacity to project forthcoming energy draw, even under atypical consumption conditions such as those encountered on public holidays. The projection of future draw and its departures functions not merely as a standards-concordant demonstration of energy-linked output but additionally unlocks instantaneous probing by surfacing it upon a cloud-hosted platform. Going further still, the probing equips decision-makers with the tooling to carry out effective energy-saving countermeasures, speaking to a shortcoming intrinsic to standard energy management setups, which routinely lack the capacity for instantaneous probing, thereby impeding the effective execution of energy-saving moves. The coupled paring back of energy draw as well as the enactment of corresponding measures to lift energy performance ought in consequence to contribute to the fostering of an ecologically sounder industrial landscape. This likewise aligns with the United Nations’ SDG 7, which aims to redouble the pace of energy performance improvement.

For onward investigation, it would stand to advantage to weigh the incorporation of supplemental shaping variables into the models that link back to the roots of the departures or over-draw. These might include markers such as workspace staffing density, fabrication capacity loading, or the order book standing. Through such means, additional states stand to be sorted with a view to corroborating departures or snaring irregularities that resist explanation via standard patterns of conduct. These setups offer pointers for paring back energy draw, which leads directly to improved energy performance. Furthermore, the threading of Large Language Models into energy management frameworks harbours weighty promise for coming-wave progress. These models can self-operate and generate and unpack energy-saving countermeasures, rooted in flagged irregularities or inefficiencies. Once knitted into a digital energy advisory tool, these capabilities could furnish enterprises with bespoke pointers to streamline energy use and implement targeted performance steps, further honing decision-making sequences. Beyond the scrutiny of state shifts within energy draw records, the yardstick sits open to being ported across to other data streams or the models being reshaped correspondingly, to take instances such as

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

pollutant release, water draw, or refuse output, to tap the possibilities for a more durable stewardship of resources throughout yet more spheres.

## Acknowledgements

None

## Conflict of interest

None

## Financial support

None

## Ethics statement

None

Received: 09 Jun 2022 Revised: 15 Aug 2022 Accepted: 31 Aug 2022
Published online: 15 June 2023

### Rights and permissions

**Open Access** The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

## References

1. van Ruijven BJ, de Cian E, Wing IS. Amplification of future energy demand growth due to climate change. Nat Commun. 2019;10(1):2762. https://doi.org/10.1038/s41467-019-10399-3.

2. Goal 7 | Department of Economic and Social Affairs. [Online]. Available: https://sdgs.un.org/goals/goal7#targets_and_indicators [accessed: 2025].

3. Energy and the Green Deal. [Online]. Available: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal/energy-and-green-deal_en [accessed: 2025].

4. Förtsch G, Meinholz H. Energiemanagement und energieaudit. In: Handbuch betrieblicher immissionsschutz. Wiesbaden: Springer Fachmedien Wiesbaden; 2020:415-84. https://doi.org/10.1007/978-3-658-29147-1_8.

5. Schulze M, Nehler H, Ottosson M, Thollander P. Energy management in industry – a systematic review of previous findings and an integrative conceptual framework. J Clean Prod. 2016;112:3692-708.

https://doi.org/10.1016/j.jclepro.2015.06.060.

6. Zhou K, Fu C, Yang S. Big data driven smart energy management: from big data to big insights. Renew Sustain Energy Rev. 2016;56:215-25. https://doi.org/10.1016/j.rser.2015.11.050.

7. Smith KM, Wilson S, Hassall ME. Barriers and drivers for industrial energy management: the frontline perspective. J Clean Prod. 2022;335:130320. https://doi.org/10.1016/j.jclepro.2021.130320.

8. Kumari A, Gupta R, Tanwar S, Kumar N. Blockchain and AI amalgamation for energy cloud management: challenges, solutions, and future directions. J Parallel Distrib Comput. 2020;143:148-66. https://doi.org/10.1016/j.jpdc.2020.05.004.

9. Shamim A, Muhamad ZBS, Syed S, Hussain R. An intelligent data-driven approach for electrical energy load management using machine learning algorithms. Energies. 2022;15(16):5789. https://doi.org/10.3390/en15165789.

10. Szczepaniuk H, Szczepaniuk EK. Applications of artificial intelligence algorithms in the energy sector. Energies. 2023;16(1):347. https://doi.org/10.3390/en16010347.

11. Friedrich R, Bauer M, Bingel M, Grimm S, Nissen D, Schulz S, et al. Potenziale der schwachen künstlichen intelligenz für die betriebliche ressourceneffizienz. Fraunhofer; 2021. https://publica.fraunhofer.de/entities/publication/e3b2fd7a-b7b0-458f-9174-bffb5b23f826/details

12. Deutsche Energie-Agentur GmbH (dena). KI in der Energiewirtschaft. 2025. [Online]. Available: https://www.dena.de/fileadmin/dena/Publikationen/PDFs/2022/dena-Umfrage_Kuenstliche_Intelligenz_in_der_Energiewirtschaft_PPT.pdf [accessed: 2025].

13. Amalou I, Mouhni N, Abdali A. Multivariate time series prediction by RNN architectures for energy consumption forecasting. Energy Rep. 2022;8:1084-91. https://doi.org/10.1016/j.egyr.2022.07.139.

14. Aguila-Leon J, Vargas-Salgado C, Chiñas-Palacios C, Díaz-Bello D. Energy management model for a standalone hybrid microgrid through a particle swarm optimization and artificial neural networks approach. Energy Convers Manag. 2022;267:115920. https://doi.org/10.1016/j.enconman.2022.115920.

15. ISO 50001:2018: Energy management systems — Requirements with guidance for use. International Organization for Standardization.

16. Brugger-Gebhardt S, Jungblut G. Nicht einfach darauf los. In: Die DIN EN ISO 50001:2018 verstehen. 2019:57-99. https://doi.org/10.1007/978-3-658-26266-2_6.

17. Umweltbundesamt. Umwelt- und Energiemanagementsysteme. 2025. [Online]. Available: https://www.umweltbundesamt.de/daten/umwelt-wirtschaft/umwelt-energiemanagementsysteme#eco-management-and-audit-scheme-emas [accessed: 2025].

18. International Organization for Standardization. The ISO Survey of Management System Standard Certifications – 2017 – Explanatory Note. 2025. [Online]. Available: https://isotc.iso.org/livelink/livelink/fetch/-8853493/8853511/8853520/18808772/00._Executive_summary_2016_Survey.pdf?nodeid=19208898&vernum=-2 [accessed: 2025].

19. Umweltbundesamt. Energiemanagementsysteme: ISO 50001. 2024. [Online]. Available: https://www.umweltbundesamt.de/energiemanagementsysteme-iso-50001#iso-50001-im-gesetzlichen-kontext [accessed:

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

2025].

20. Javied T, Rackow T, Franke J. Implementing energy management system to increase energy efficiency in manufacturing companies. Procedia CIRP. 2015;26:156-61. https://doi.org/10.1016/j.procir.2014.07.057.

21. Ahmad T, et al. Artificial intelligence in sustainable energy industry: status quo, challenges and opportunities. J Clean Prod. 2021;289:125834. https://doi.org/10.1016/j.jclepro.2021.125834.

22. Ali M, Prakash K, Hossain MA, Pota HR. Intelligent energy management: evolving developments, current challenges, and research directions for sustainable future. J Clean Prod. 2021;314:127904. https://doi.org/10.1016/j.jclepro.2021.127904.

23. Das S, Choudhury TK, Mohapatra SK. Data analytics to increase efficiency of the AI based energy consumption predictor. In: 2020 International Conference on Computational Intelligence for Smart Power System and Sustainable Energy (CISPSSE). 2020. https://doi.org/10.1109/CISPSSE49931.2020.9212249.

24. Putha S. AI-driven energy management in manufacturing: optimizing energy consumption and reducing operational costs. Distrib Learn Broad Appl Scientif Res. 2020;6:313-53. https://dlabi.org/index.php/journal/article/view/107

25. Skrobek D, Krzywanski J, Sosnowski M, Uddin GM, Ashraf WM, Nowak W. Artificial intelligence for energy processes and systems: applications and perspectives. Energies. 2023;16(8):3441. https://doi.org/10.3390/en16083441.

26. Rojek I, Mikołajewski D, Mroziński A, Macko M. Green energy management in manufacturing based on demand prediction by artificial intelligence—A review. Electronics. 2024;13(16):3338. https://doi.org/10.3390/electronics13163338.

27. Hu Y, Li J, Hong M, Ren J, Man Y. Industrial artificial intelligence based energy management system: integrated framework for electricity load forecasting and fault prediction. Energy. 2022;244:123195. https://doi.org/10.1016/j.energy.2022.123195.

28. Lazzeri F. Introduction to Neural Networks for time series forecasting. In: Mach Learn Time Series Forecast Python. 2020:137-65. https://doi.org/10.1002/9781119682394.CH5.

29. Wang JQ, Du Y, Wang J. LSTM based long-term energy consumption prediction with periodicity. Energy. 2020;197:117197. https://doi.org/10.1016/j.energy.2020.117197.

30. Nakayama K, Sharma R. Energy management systems with intelligent anomaly detection and prediction. In: 2017 Resilience Week. 2017. https://doi.org/10.1109/RWEEK.2017.8088656.

31. Wang X, Wang H, Bhandari B, Cheng L. AI-empowered methods for smart energy consumption: a review of load forecasting, anomaly detection and demand response. Int J Precis Eng Manuf-Green Tech. 2024;11(3):963-93. https://doi.org/10.1007/s40684-023-00537-0.

32. Zangrando N, Fraternali P, Petri M, Pinciroli Vago NO, Herrera González SL. Anomaly detection in quasi-periodic energy consumption data series: a comparison of algorithms. Energy Inform. 2022;5(S4):1-22. https://doi.org/10.1186/s42162-022-00230-7.

33. Kawa B, Borkowski P. Integration of machine learning solutions in the building automation system. Energies. 2023;16(11):4504. https://doi.org/10.3390/en16114504.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

34. Surucu O, Gadsden SA, Yawney J. Condition monitoring using machine learning: a review of theory, applications, and recent advances. Expert Syst Appl. 2023;221:119738. https://doi.org/10.1016/j.eswa.2023.119738.

35. Zhou Z, Yao B, Xu W, Wang L. Condition monitoring towards energy-efficient manufacturing: a review. Int J Adv Manuf Technol. 2017;91(9-12):3395-415. https://doi.org/10.1007/s00170-017-0014-x.

36. St P, Antony K, Mahmood HM, Th T, Ahmad M, Sulaiman SM. Predictive modelling for effective energy consumption in industry 4.0 using machine learning techniques. 2024. https://doi.org/10.21203/rs.3.rs-4713283/v1.

37. Takahashi K, Ooka R, Ikeda S. Anomaly detection and missing data imputation in building energy data for automated data pre-processing. J Phys Conf Ser. 2021;2069(1):012144. https://doi.org/10.1088/1742-6596/2069/1/012144.

38. Bakakeu J, Tolksdorf S, Bauer J, Klos HH, Peschke J. An artificial intelligence approach for online optimization of flexible manufacturing systems. Appl Mechan Mater. 2018;882:96-108. https://doi.org/10.4028/www.scientific.net/AMM.882.96.

39. Keramati Feyz Abadi MM, Liu C, Zhang M, Hu Y, Xu Y. Leveraging AI for energy-efficient manufacturing systems: review and future prospectives. J Manuf Syst. 2025;78:153-77. https://doi.org/10.1016/j.jmsy.2024.11.017.

40. Javed H, Eid F, El-Sappagh S, Abuhmed T. Sustainable energy management in the AI era: a comprehensive analysis of ML and DL approaches. Computing. 2025;107(6). https://doi.org/10.1007/s00607-025-01485-0.

41. Nabati EG, Alvela Nieto MT, Bode D, Schindler TF, Decker A, Thoben KD. Challenges of manufacturing for energy efficiency: towards a systematic approach through applications of machine learning. Prod. 2022;32. https://doi.org/10.1590/0103-6513.20210147.

42. Aslam S, Aung PP, Rafsanjani AS, Majeed APPA. Machine learning applications in energy systems: current trends, challenges, and research directions. Energy Inform. 2025;8(1). https://doi.org/10.1186/s42162-025-00524-6.

43. Burggräf P, Steinberg F, Sauer CR, Nettesheim P. Machine learning implementation in small and medium-sized enterprises: insights and recommendations from a quantitative study. Prod Eng Res Dev. 2024;18(5):751-64. https://doi.org/10.1007/s11740-024-01274-2.

44. Zavodna LS, Überwimmer M, Frankus E. Barriers to the implementation of artificial intelligence in small and medium sized enterprises: pilot study. JEM. 2024;46:331-52. https://doi.org/10.22367/jem.2024.46.13.

45. Yarlagadda SC. Role of artificial intelligence, automation, and machine learning in sustainable plastics packaging markets: progress, trends, and directions. IJRITCC. 2023;11(9s):818-28. https://doi.org/10.17762/ijritcc.v11i9s.9489.

46. Wirth R. CRISP-DM: towards a standard process model for data mining. In: Proceedings of the Fourth International Conference on the Practical Application of Knowledge Discovery and Data Mining. 2000:29-39.

47. Schafer F, Zeiselmair C, Becker J, Otten H. Synthesizing CRISP-DM and quality management: a data mining approach for production processes. In: 2018 IEEE Int Conf Technol Manag Oper Decis (ICTMOD). 2018:190-5. https://doi.org/10.1109/ITMC.2018.8691266.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press

48. Deb C, Zhang F, Yang J, Lee SE, Shah KW. A review on time series forecasting techniques for building energy consumption. Renew Sustain Energy Rev. 2017;74:902-24. https://doi.org/10.1016/j.rser.2017.02.085.

49. Wickham H, Stryjewski L. 40 years of boxplots. 2011. https://doi.org/10.48550/arXiv.2106.11987.

50. Kazanc H, Kauffeldt F. Künstliche Intelligenz und machine learning in der Energiewirtschaft. In: Realisierung utility 4.0 band 2. Wiesbaden: Springer Fachmedien Wiesbaden; 2020:449-63. https://doi.org/10.1007/978-3-658-25589-3_33.

51. Nassif AB, Talib MA, Nasir Q, Dakalbab FM. Machine learning for anomaly detection: a systematic review. IEEE Access. 2021;9:78658-700. https://doi.org/10.1109/ACCESS.2021.3083060.

52. Lyu Y, Li H, Sayagh M, Jiang ZM, Hassan AE. An empirical study of the impact of data splitting decisions on the performance of AIOps solutions. ACM Trans Softw Eng Methodol. 2021;30(4):1-38. https://doi.org/10.1145/3447876.

53. Miyaki K. Time Series Split with Scikit-learn. Medium. March 12, 2025. [Online]. Available: https://medium.com/keita-starts-data-science/time-series-split-with-scikit-learn-74f5be38489e [accessed: 2025].

54. Thomas L, Manoj Kumar MV, Annappa B. Discovery of optimal neurons and hidden layers in feed-forward Neural Network. In: 2016 IEEE Int Conf Emerg Technolog Innovat Business Pract Transformat Societies (EmergiTech). 2016:286-91. https://doi.org/10.1109/EMERGITECH.2016.7737353.

55. Abbasimehr H, Shabani M, Yousefi M. An optimized model using LSTM network for demand forecasting. Comput Indust Eng. 2020;143:106435. https://doi.org/10.1016/j.cie.2020.106435.

56. Bouktif S, Fiaz A, Ouni A, Serhani MA. Optimal deep learning LSTM model for electric load forecasting using feature selection and genetic algorithm: comparison with machine learning approaches. Energies. 2018;11(7). https://doi.org/10.3390/en11071636.

57. Mahjoub S, Chrifi-Alaoui L, Marhic B, Delahoche L. Predicting energy consumption using LSTM, multi-layer GRU and drop-GRU neural networks. Sensors. 2022;22(11):4062. https://doi.org/10.3390/s22114062.

58. Chollet F. Deep Learning with Python. Shelter Island, NY: Manning Publications; 2018.

59. Géron A. Hands-on Machine Learning with Scikit-Learn, Keras and TensorFlow. 2nd ed. Sebastopol, CA: O'Reilly Media; 2019.

60. Haugsdal E, Aune E, Ruocco M. Persistence Initialization: a novel adaptation of the Transformer architecture for time series forecasting. Appl Intell. 2023;53(22):26781-96. https://doi.org/10.1007/s10489-023-04927-4.

© 2023 The Author(s). | Published by Institute of Smart Energy Systems Press