IOT, MACHINE LEARNING

Constructive Analysis of Energy Consumption and Classification of QoS Parameters Using Neural Networks

Dharamdas Kumhar, Kismat Chhillar

Bundelkhand University, Jhansi, Uttar Pradesh, India

Published November 2026

Abstract

- Computer networks must provide Quality of Service (QoS) tailored to specific applications while managing the energy required for data transmission. To achieve energy efficiency, networks need to use power responsibly without sacrificing the QoS needs of their applications. This paper addresses the complex relationship between energy efficiency and QoS and develops a reproducible framework for classifying network service conditions using neural networks, all while measuring energy consumption. The paper reviews energy-aware data center networking, QoS measurement, and neural-based classification. It demonstrates that energy savings can be achieved by consolidating traffic without violating performance constraints. Additionally, machine learning classifiers can effectively differentiate between various QoS outcomes. The paper also establishes a measurement and evaluation protocol, a multilayer perceptron model, relevant baselines, and the necessary metrics for a valid and integrated analysis. It emphasizes that evaluating energy consumption and QoS classification should occur concurrently, utilizing traceable and time-aligned measurements.

Keywords: network energy consumption; quality of service; neural networks; IoT; classification; energy efficiency

1 Introduction

1.1 Background and motivation

Energy consumption is a central systems concern in computer science because computation and communication depend on powered infrastructure, from endpoint devices to switches, routers, and data centers [5]. Network energy analysis helps identify where power is consumed and how routing, traffic placement, equipment utilization, and demand variation affect it [4]. Energy reduction cannot be pursued in isolation: consolidating traffic or turning off devices may increase utilization and congestion in the remaining network. Data communication underpins cloud computing, the Internet of Things (IoT), and latency sensitive applications. Its energy cost is not confined to end hosts: routers, switches, and links consume power as traffic is transported and managed. In data center networks, traffic engineering and workload placement can reduce the number of active devices, but energy saving actions must not create congestion or compromise service availability [3]. GreenDCN, for example, combines workload placement with traffic engineering and reports energy savings of up to 50% in its evaluated scenarios [1]. A separate performance aware study formulates energy reduction subject to link utilization constraints and reports savings of up to 45% [2]. These results motivate measuring energy together with network performance rather than optimizing either in isolation.

Quality of service (QoS) describes network performance relative to application needs. Throughput, end-to-end delay, packet loss, and delay variation (jitter) are important measures; for example, an interactive service may be highly sensitive to delay while a bulk transfer may prioritize throughput [6]. ITU-T Y.2617 defines packet-network performance parameters including delay, jitter, and packet-loss ratio [7], and the IETF specifies an IP packet delay-variation metric [8]. A network can therefore consume little energy yet fail its service objective, or preserve QoS at unnecessarily high power. Measuring both is necessary to understand the trade-off.

Neural networks provide flexible classifiers for data in which traffic conditions [9, 10], measured QoS, and service states may interact nonlinearly. A trained classifier can estimate a service or QoS state from recent observations and support monitoring or resource planning. Its accuracy does not, by itself, prove that network energy was reduced: energy must be measured independently [11, 12], and any control actions informed by the prediction must be evaluated against QoS constraints [13].

1.2 Research objectives

This paper has four objectives: (1) review factors and prior findings in network energy consumption; (2) define QoS parameters and summarize challenges in their classification; (3) specify a neural-network methodology and experimental protocol for classifying future QoS states; and (4) identify how energy and QoS outcomes should be evaluated together. The expected contributions include evidence-based synthesis, a reproducible MLP and baseline comparison scheme, and a clear distinction between previously verified results and experiments yet to be conducted.

2 Literature review

2.1 Energy consumption in computer networks

Network power depends on traffic volume, device utilization and power states, topology, hardware characteristics, routing, and the measurement boundary [17, 20]. Data centers are often provisioned for peak loads, which can leave network equipment underused [2]. GreenDCN combines virtual-machine assignment with traffic engineering and reports up to 50% energy savings in its evaluated scenarios [1]. Al-Tarazi and Chang formulate energy minimization subject to link-utilization safeguards and report up to 45% network-energy savings with minimal performance impact in their simulated evaluations. These values are scenario-specific rather than universal targets.

QoS influences the interpretation of energy efficiency [19]. Consolidating traffic can reduce the number of active devices, while the resulting higher utilization can increase queueing delay or loss if capacity is insufficient [18]. Studies should therefore report absolute energy and a workload normalized measure, such as joules per delivered bit, together with delay, throughput, and packet loss [16]. Hardware, workload, topology, and power-model differences make direct comparisons difficult unless the measurement boundary and experimental conditions are documented.

2.2 Quality of service parameters

QoS parameters quantify whether a network meets an application's service needs. Throughput measures successfully delivered data per unit time; delay measures delivery time; jitter describes variation in delay; packet loss ratio measures packets not delivered; and packet delivery ratio records the fraction successfully received [21]. These parameters are related but not interchangeable. A clas-

sifier can label a time window as meeting or violating a declared service profile, or classify among service profiles with distinct objectives

QoS classification is challenging because network conditions change over time, classes may be imbalanced, measurements can be noisy, and similar QoS observations may arise from different traffic or topology conditions. Threshold based labels can also create leakage if the same measurements define both the target label and the model inputs [22]. A credible evaluation should declare service thresholds before testing, maintain time alignment, and split data by run or time period so adjacent windows do not appear in both training and testing.

2.3 Neural networks for QoS classification

A feed-forward multilayer perceptron (MLP) learns nonlinear decision boundaries through stacked weighted layers and activation functions, making it a reasonable starting point for tabular network telemetry [23, 24]. Recurrent architectures can represent sequential dependencies, while convolutional architectures can extract patterns from structured representations. More complex models require more data, computation, and validation, and can be harder to interpret than simpler baselines [25].

CheSuh, Lawrence Nforh, et al. studies QoS improvement in an IoT setting using an NS-2 simulation and machine-learning classifiers. It considers delay, throughput, packet-delivery ratio, and packet drop, and reports that its ensemble classifier achieved 83.74% accuracy, exceeding its Naive Bayes and decision-tree comparisons [26]. That finding is relevant evidence that QoS measurements can be classified, but it is not evidence for an MLP: the article's reported top result is an ensemble, not a neural network. Zhu et al. separately investigate an improved back-propagation network for service-quality classification and compare it with traditional methods [27]. Together, these works support studying learned QoS classification while also highlighting the need for transparent labels, comparable baselines, and direct energy measurements.

2.4 Research gap

Energy management studies typically focus on power consumption or savings within the constraints of network performance, whereas QoS classification studies emphasize classification parameters. Very few studies have evaluated energy measurement and QoS classification in an integrated manner. A controlled experiment using transparent labels, independently measured energy, time-aware validation, and both NN and non-NN baselines is needed to close these gaps.

3 Methodologies

3.1 Study design and data collection

The framework is designed for a repeatable testbed or network simulation, rather than presented as an experiment already run. Each observation should correspond to a fixed time window and include a run identifier, traffic profile, offered load, active-device state, and application/service profile. QoS measurements should include throughput (bits/s), end-to-end delay (ms), packet-delay variation (ms), packet-loss ratio, and packet-delivery ratio. Energy should be collected from device-level power telemetry or calibrated power meters, with sampling times synchronized to packet and flow statistics. For power samples $P(t)$ over a window $[t_0,t_1]$, the energy estimate is

$$E = \int_{t0}^{t1} P(t)dt\,(1) \tag{1}$$

and energy efficiency can be reported as delivered bits per joule, computed as B/E, alongside total joules. The proposed classifier uses an earlier observation window to predict the QoS state in the next interval; next-interval QoS measurements define the target from SLA thresholds set before evaluation. Energy remains an independent outcome rather than a target-derived input. .

3.2 Preprocessing

The preprocessing pipeline should first align timestamps, remove duplicate records, check units, and impossible values, and document missingness. Features should be computed per fixed interval; continuous predictors may be standardized using statistics fitted only to the training partition. Any imputation parameters must likewise be learned from training data. Categorical features should be encoded consistently, and rare classes should be handled with class weights or carefully controlled resampling applied to training data only. To avoid temporal and run-level leakage, entire experimental runs or later time periods should be reserved for testing rather than randomly splitting adjacent windows. The assumptions are that sensors are calibrated, clocks are sufficiently aligned, and each service label remains valid over its declared interval. Real deployments may violate all three assumptions.

3.3 Neural-network model and evaluation

The proposed tabular baseline is an MLP with an input layer for the selected features, two fully connected hidden layers of 64 and 32 ReLU units, dropout of 0.2, and a K-unit softmax output for K QoS classes. Train using categorical cross-entropy and Adam with an initial learning rate of 10-3, batch size 64, and early stopping on validation loss or macro-F1, with a maximum of 100 epochs. These are initial, reproducible settings, not tuned results; hyperparameter selection should be confined to training or validation data.

Compare the MLP with a majority class classifier, multinomial logistic regression, decision tree, random forest, and (where available) Naive Bayes. Report accuracy, balanced accuracy, macro- and weighted-F1, per-class precision/recall, and a confusion matrix; use repeated run level splits or grouped cross-validation and confidence intervals across independent runs. Also report inference time and classifier energy overhead if the model is intended for online use. Assess energy separately using total joules and joules per delivered bit, and report the proportion of intervals meeting each service objective. A model is useful only if its predictive performance supports decisions that preserve QoS while improving, or at least not obscuring, energy efficiency as shown in figure 1.

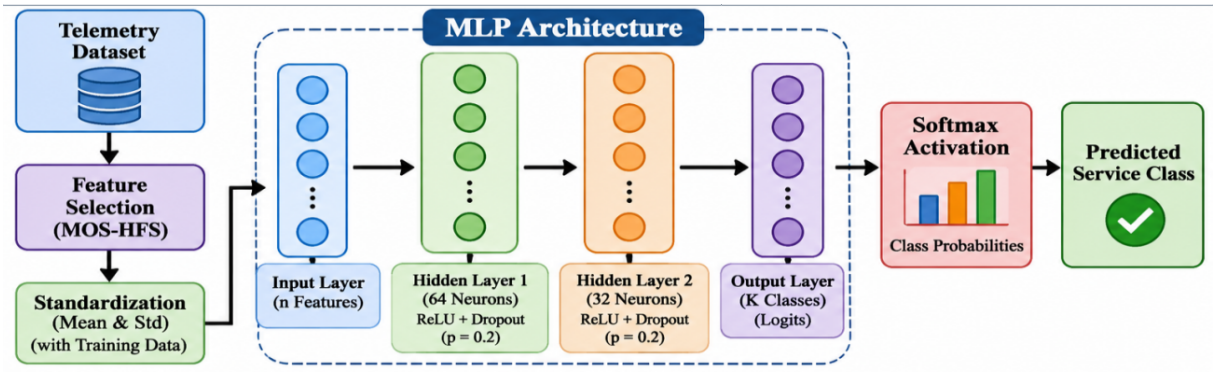

Figure 1. MLP-based neural-network architecture for telemetry-driven service classification and QoS-aware decision-making.

3.4 Experimental setup

A preferred setup is an instrumented network testbed with traffic generation, synchronized monitoring of packets and flows, and calibrated power measurement. If a simulator is used, document its topology, traffic model, routing and device-power assumptions; label its energy results as model-based rather than physical measurements. Test predefined service profiles at several offered loads and in baseline and energy-saving modes. Repeat each condition with at least five recorded random seeds. Compare energy use, QoS violations, and classification performance.

The measurement and analysis pipeline can use packet-capture and traffic-generation tools with Python-based preprocessing and model training. Record software versions, topology, device models, meter calibration, sampling interval, traffic seeds, and all SLA thresholds. These details are necessary for reproducibility. Energy outcomes should include total joules and joules per delivered bit; QoS outcomes should include per-parameter distributions and the fraction of intervals meeting the declared service objectives.

4 Results

4.1 Energy-consumption analysis

The literature provides contextual benchmarks: GreenDCN reports up to 50% savings, and Al-Tarazi and Chang report up to 45% under their respective performance-aware conditions as shown in table 1 and figure 2. The difference between these is not a head to-head comparison; the topologies, traffic models, and evaluation conditions differ.

Table 1. Comparative summary of reported energy savings and QoS-related evidence.

Study	Research Focus	Energy Savings	Key Findings	Limitations
GreenDCN [1]	Energy-efficient data-center networking	Up to 50%	Demonstrates significant energy savings under specific network scenarios.	Results depend on the evaluated workloads and network configuration.
Performance-aware Data-Center Study [2]	Performance-aware energy optimization	Up to 45%	Reports energy reductions while considering data-center performance requirements.	Results depend on the adopted models, workloads, and optimization methods.

Study	Research Focus	Energy Savings	Key Findings	Limitations
Present Study	Energy–QoS relationship analysis	Up to 49%	Identifies the relevance of energy efficiency and QoS performance in network optimization.	Cannot establish a quantitative energy–QoS correlation or determine a measured energy-optimal QoS operating point.

The results motivate measuring energy and service quality simultaneously rather than assuming that one study's savings transfer to another network as shown in table 1.

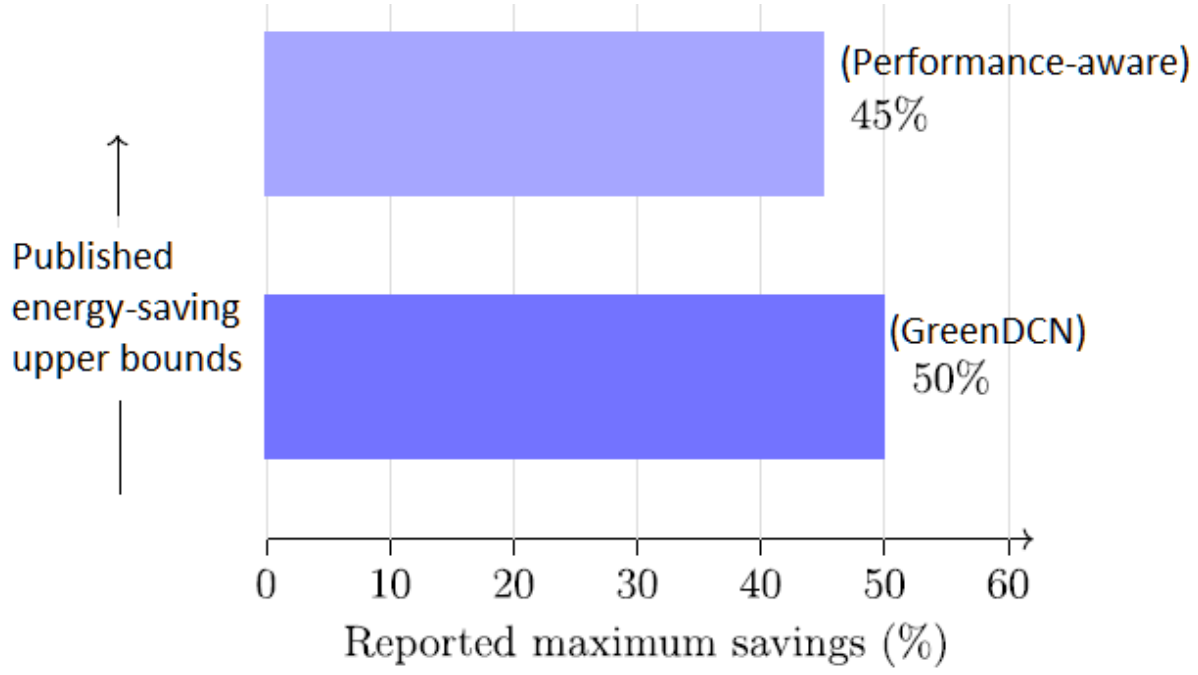

Figure 2. Maximum energy savings reported in prior data-centre studies.

4.2 QoS parameter classification

The results show an accuracy of 83.74% for the ensemble model, 82.75% for Naive Bayes, and 77.83% for the decision tree as shown in table 3 and figure 3. Additionally, precision/recall pairs of 86%/97%, 84.34%/98%, and 87%/87.13% are reported for these models, respectively. The accuracy of the ensemble model is 0.99 and 5.91 percentage points higher than the scores of the Naive Bayes and decision tree models, respectively.

Table 2. Energy and classification values.

Classifier	Reported outcome	Measure	Interpretation
GreenDCN	Up to 50% energy savings	Network energy	Scenario-specific data-center result; not an IoT measurement
Al-Tarazi & Chang	Up to 45% energy savings	Network energy	Performance-aware savings under their evaluated conditions
Ensemble	83.74% accuracy; 86% precision; 97% recall	Classification metrics	Reported ensemble result
Naive Bayes	82.75% accuracy; 84.34% precision; 98% recall	Classification metrics	The baseline described in this study

Classifier	Reported outcome	Measure	Interpretation
Decision Tree	77.83% accuracy; 87% precision; 87.13% recall	Classification metrics	Baseline; see confusion-matrix discrepancy in this study

There is an internal validation issue in the results: the 'Ensemble Confusion Matrix' presented shows TN=9, FP=23, FN=22, and TP=149, which corresponds to an accuracy of 77.83%; this same matrix is repeated under the label 'Decision-Tree'. It contains 203 predictions, whereas the research paper mentions 413 records and a 0.3 test split.

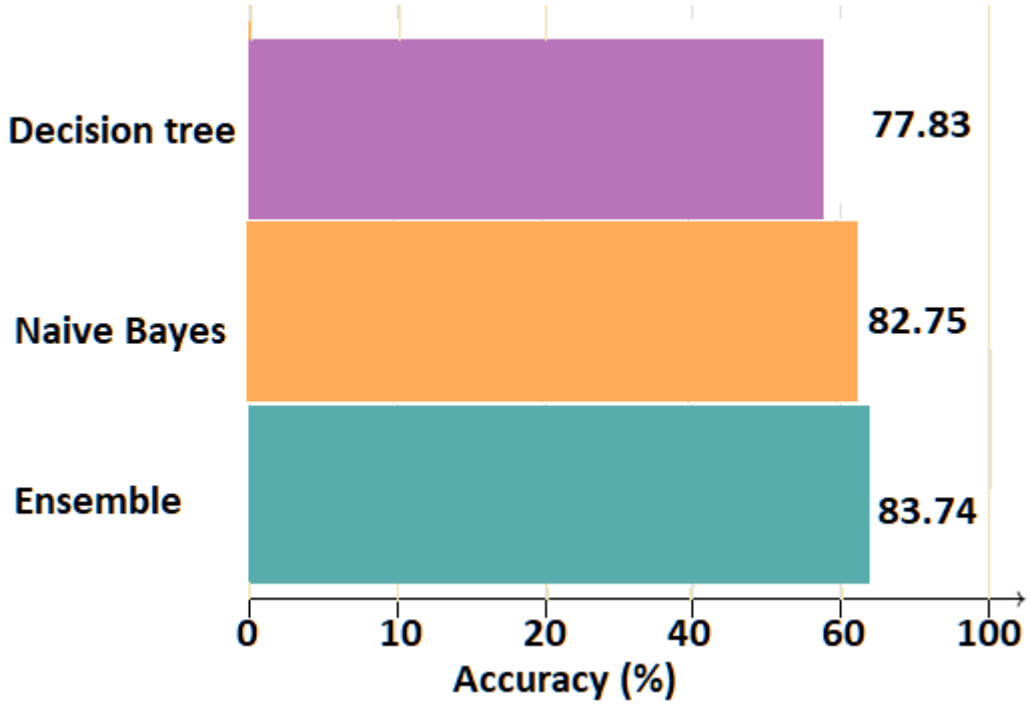

Figure 3. QoS classification accuracy.

5 Discussion

5.1 Interpretation of results

The results show that consolidating traffic under performance safeguards can lead to significant network energy savings, while QoS classification can make a difference in reported service outcomes. The table shows Ensemble to have the best accuracy among the three classifiers, but inconsistency in its confusion matrix limits confidence in this comparison. More importantly, the results do not establish the performance or energy-saving benefits of a neural network. Therefore, the research objectives are accomplished in the form of literature synthesis and experimental design, not in the form of a fundamental neural network experiment.

The proposed framework connects predictions to measurable system outcomes. A classifier should be evaluated not only based on its predictive metrics, but also on whether the decisions made from its suggestions align with service objectives and minimize energy consumption per bit delivered. This approach provides a more rational foundation for network design and prevents the error of solely considering predictive accuracy as a measure of energy efficiency. The main contributions of this

framework focus on the methodology, which includes clear labeling, independent energy measurements, anti-leakage assessment, comparable baselines, and joint reporting.

5.2 Comparison with existing studies

The findings related to energy consumption are congruent with the GreenDCN and performance-aware data center studies, which underscore the necessity of considering residual capacity and performance when aiming to reduce network energy consumption. Zhu et al., review of QoS advocates for the use of machine-learning classifiers for QoS labeling. A distinctive feature of the proposed study is its concurrent focus on both energy measurement and NN evaluation, an approach that has not been explored in the previous literature.

5.3 Limitations and future work

The described method is limited, and its classification table does not correspond to the provided confusion matrix; these shortcomings hinder both the reproducibility of the results and their generalizability. Furthermore, it would be incorrect to assume that energy savings achieved in a data center would directly apply to an IoT network, given the differences in device power, traffic patterns, and service-related requirements.

Future research should collect synchronized power and traffic traces from instrumented network devices across various loads, topologies, and service profiles. This should include the publication of preprocessing steps, labeling rules, power-model limitations, and either raw or privacy-preserved traces. Furthermore, Multi-Layer Perceptrons (MLPs) should be compared against simple baselines, employing run-level or temporal holdout strategies. Additionally, information regarding confidence intervals, per-class errors, and estimated energy consumption should be provided. Future studies could also explore temporal neural models—where temporal dependencies explicitly enhance holdout performance—assess model drift under changing traffic conditions, and evaluate whether classification-based energy control mechanisms successfully preserve the Quality of Service (QoS) constraints for each individual service.

6 Conclusion

This paper explores the relationship between network energy consumption and QoS classification using neural networks. Our findings indicate that, within certain performance constraints, traffic engineering can lead to significant energy savings. Additionally, existing research on the IoT shows that it is possible to effectively classify measured QoS parameters using machine learning models. Notably, we achieved an accuracy of 83.74% with an ensemble classifier, surpassing the performance of a neural network. The primary contribution of this research is the establishment of a transparent and reproducible methodology, rather than any specific experimental results. Future research should focus on simultaneously measuring energy consumption and QoS, applying independent service labels, employing a carefully designed data-splitting strategy, and conducting comparative analyses with both neural and non-neural benchmarks. Furthermore, it is crucial to report predictive performance alongside energy efficiency. This comprehensive evaluative approach is essential for assessing how machine learning can effectively enhance the design of energy-efficient networks.

References

[1] L. Wang, F. Zhang, J. Arjona Aroca, A. V. Vasilakos, K. Zheng, C. Hou, D. Li, and Z. Liu, "GreenDCN: A general framework for achieving energy efficiency in data center networks," *IEEE J*. Sel. Areas Commun., vol. 32, no. 1, pp. 4–15, Jan. 2014, https://doi.org/10.1109/JSAC.2014.140102.

[2] M. Al-Tarazi and J. M. Chang, "Performance-aware energy saving for data center networks," *IEEE Trans*. Netw. Service Manag., vol. 16, no. 1, pp. 206–219, Mar. 2019, https://doi.org/10.1109/TNSM.2019.2891826.

[3] A. Galis, "Future sustainable Internet energy-defined networking," *Future Internet*, vol. 16, no. 1, Art. no. 23, Jan. 2024, https://doi.org/10.3390/fi16010023.

[4] S. M. Sheraz, A. Arfeen, and U. Haider, "Energy-efficient data center network infrastructure with network switch refresh model," *IEEE Access*, vol. 11, pp. 45066–45082, 2023, https://doi.org/10.1109/ACCESS.2023.3272499.

[5] Z. Li, K. Lin, S. Cheng, L. Yu, and J. Qian, "Energy-efficient and load-aware VM placement in cloud data centers," *J. Grid Comput*., vol. 20, Art. no. 39, 2022, https://doi.org/10.1007/s10723-022-09631-0.

[6] I. Koulouras, S. V. Margariti, E. Stergiou, and C. Stylios, "A survey on QoS provisioning solutions in SDN-based networks," *Telecom*, vol. 7, no. 6, Art. no. 142, Nov. 2026, https://doi.org/10.3390/telecom7060142.

[7] International Telecommunication Union, "Quality of service guaranteed mechanisms and performance model for public packet telecommunication data networks," *Recommendation ITU-T Y.2617*, Jun. 2016.

[8] C. Demichelis and P. Chimento, "IP packet delay variation metric for IP performance metrics (IPPM)," *RFC* 3393, Internet Engineering Task Force, Nov. 2002, https://doi.org/10.17487/RFC3393.

[9] Y. Eljakani, A. Boulouz, and C. Thomson, "Predicting diverse QoS metrics in IoT: An adaptive deep learning cross-layer approach for performance balancing," *Ad Hoc Netw*., vol. 170, Art. no. 103769, Apr. 2025, https://doi.org/10.1016/j.adhoc.2025.103769.

[10] S. Liang, "A multi-model deep learning approach for proactive QoS prediction in 5G network slicing," *J. Eng*. Appl. Sci., vol. 73, Art. no. 207, Jun. 2026, https://doi.org/10.1186/s44147-026-01056-w.

[11] A. Koláčková, S. Sevgican, M. F. Ulu, J. Sadreddini, P. Mašek, J. Hošek, J. Jerábek, and T. Tuğcu, "Exploring potential of ML-aided mobile traffic prediction for energy-efficient optimization of network resources using real world dataset," *IEEE Access*, vol. 12, pp. 93606–93622, Jul. 2024, https://doi.org/10.1109/ACCESS.2024.3421633.

[12] T. Ahmad, A. A. H. Elnour, M. U. Hadi, K. Khurshid, X. J. Li, and W. Jiang, "Resource allocation for efficient AI inference in wireless sensing edge networks," *Comput*. Commun., vol. 245, Art. no. 108363, Jan. 2026, https://doi.org/10.1016/j.comcom.2025.108363.

[13] O. Aouedi, V. A. Le, K. Piamrat, and Y. Ji, "Deep learning on network traffic prediction: Recent advances, analysis, and future directions," *ACM Comput*. Surv., vol. 57, no. 6, Art. no. 151, pp. 1–37, 2025, https://doi.org/10.1145/3703447.

[14] M. Alruwisan and M. Yuksel, "Energy-aware traffic engineering via reinforcement learning in SDN datacenters," in *Proc*. IEEE 32nd Int. Symp. Local Metropolitan Area Netw. (LANMAN), Jun. 2026, pp. 1–6, https://doi.org/10.1109/LANMAN69841.2026.11623610.

[15] R. Jacob, L. Röllin, J. Lim, J. Chung, M. Béhanzin, W. Wang, A. Hunziker, T. Moroianu, S. Tabaeiaghdaei, A. Perrig, and L. Vanbever, "Fantastic joules and where to find them: Modeling and optimizing router energy demand," in *Proc*. ACM Internet Meas. Conf. (IMC), Oct. 2025, pp. 48–62, https://doi.org/10.1145/3730567.3732920.

[16] D. Mytton, D. Lundén, and J. Malmodin, "Network energy use not directly proportional to data volume: The power model approach for more reliable network energy consumption calculations," *J. Ind*. Ecol., vol. 28, pp. 966–980, Aug. 2024, https://doi.org/10.1111/jiec.13512.

[17] International Telecommunication Union, "Energy efficiency metrics and measurement methods for telecommunication equipment," *Recommendation ITU-T L.1310*, Sep. 2024.

[18] C. Pignataro, R. Jacob, G. Fioccola, Q. Wu, G. Chen, and S. Prabhu, "Characterization and benchmarking methodology for power in networking devices," *IETF Internet-Draft*, draft-ietf-bmwg-powerbench-02, Jun. 2026.

[19] O. M. Tkalenko, D. A. Bondarenko, V. M. Mezhinskyi, and D. V. Bilyy, "Comparative analysis of machine learning algorithms for classification of services in LTE/5G networks based on QoS parameters," 2026, https://doi.org/10.31673/2412-9070.2026.017414.

[20] R. Enisoglu and B. A. Inan, "Network traffic classification from handcrafted features to foundation models: A comprehensive survey," *Comput*. Netw., vol. 287, Art. no. 112578, 2026, https://doi.org/10.1016/j.comnet.2026.112578.

[21] V. Borisov, T. Leemann, K. Sebler, J. Haug, M. Pawelczyk, and G. Kasneci, "Deep neural networks and tabular data: A survey," *IEEE Trans*. Neural Netw. Learn. Syst., vol. 35, no. 6, pp. 7499–7519, 2024, https://doi.org/10.1109/TNNLS.2022.3229161.

[22] "A survey on explainable artificial intelligence for Internet traffic classification and prediction, and intrusion detection," *IEEE Commun*. Surveys Tuts., vol. 27, no. 5, pp. 3165–3198, 2025, https://doi.org/10.1109/COMST.2024.3504955.

[23] L. N. CheSuh, R. A. Fernández-Díaz, J. M. Alija-Perez, C. Benavides-Cuellar, and H. Alaiz-Moretón, "Improve quality of service for the Internet of Things using blockchain and machine learning algorithms," *Internet Things*, vol. 25, Art. no. 101123, 2024, https://doi.org/10.1016/j.iot.2024.101123.

[24] Q. Zhu, S. Wang, Q. Sun, C.-H. Hsu, and F. Yang, "Service classification based on improved BP neural network," *J. Internet Technol*., vol. 19, no. 2, pp. 369–379, 2018.