A New Explainable Power Demand Model for 4G LTE and 5G NR Base Stations

(Article begins on next page)

A New Explainable Power Demand Model for 4G LTE and 5G NR Base Stations

Greta Vallero[1], Giovanni Perin[2,3], Michela Meo[1],
Massimo Sereno Garino[4], Stefano D'Elia[4], and Davide Vaccarono[4]

[1]Politecnico di Torino, Italy, [2]University of Padova, Italy, [3]University of Brescia, Italy, [4]Vodafone, Italy

***Abstract*—To design efficient Radio Access Networks (RANs) capable of handling the increasing power demand of modern networks, it is crucial to accurately assess the power consumption of a Base Station (BS). Since existing models are outdated, we propose a new data-driven model that accurately reflects the power demand of modern BSs. We derive the model using real-world measurements from operative three-sector BSs, differing for technologies and transmission frequencies. Our findings suggest that models tailored to specific technologies and transmission frequencies outperform generalized ones. Moreover, linear regression models consistently perform up to 96% better than those based on Multilayer Perceptrons (MLPs), Recurrent Neural Networks (RNNs) and benchmarks from the literature, highlighting a predominantly linear relationship between input features and power needs. Finally, the most accurate power estimations are produced by a linear model using traffic volume, load, maximum transmission power, and cable power losses, as regressors, with errors ranging from 4 W to 38 W.**

I. Introduction

In recent years, the communication and networking community has become increasingly aware of the urgent need to reduce the power requirements of RANs [1]. This concern was driven by the fact that, in 2010, BSs were already responsible for two-thirds of the total CO_2 RAN emissions [2]. Today, with the tremendous growth of mobile services and applications, the power demand for the RAN supply keeps rising. When combined with the energy consumption of data centers, this demand is projected to reach 581 TWh in China by 2035, representing a 289% increase concerning 2020 levels [3].

In response, significant efforts were made to design RANs management strategies promoting a sustainable and power-efficient evolution. These efforts include (i) optimizing BS hardware components [4], (ii) implementing sleep mode strategies [5], (iii) enhancing radio transmission mechanisms [6], [7], (iv) optimizing network deployment and planning [8], and (v) integrating energy harvesting technologies [9].

In this context, assessing the power requirements of a BS is crucial to guide network operators toward more efficient management practices. Two simple models proposed in the past decades formalize the power needs of a BS as a linear function of load and transmitting power, respectively, with values set according to the type of BS (macro, micro, femto, and pico), as reported in [2], [10]. Due to their simplicity and accuracy, these models have been widely adopted in the literature but are now becoming outdated due to the rapid evolution of BS hardware components. Advances in power amplifiers, processing units, and cooling systems have significantly altered the power demand patterns of modern BSs. Furthermore, the introduction of recent radio technologies has made it increasingly difficult for these older models to accurately represent the power behavior of today's RAN infrastructure. Recently, a new model was presented in [11], developed using an MLP trained on data from over 7 700 antenna units. This model takes several input features, including RF configuration, number of carriers, traffic volume, load, and energy-saving strategies. However, the parameters of this model are unfortunately inaccessible. Consequently, there is a clear need for new empirical and explainable models that reflect the current state of BS hardware, similar to the model in [11], while also offering the accessibility and accuracy of earlier models like those in [2], [10].

In this paper, we propose a new data-driven and explainable model that addresses these gaps by accounting for the latest advancements in BS hardware. To realize the model, we exploited the capability of the new generation of BSs to measure the power demand for each sector and each frequency band, and to collect this information among the network counters, together with the traffic volume and the load, both in uplink (UL) and downlink (DL); we have also included in our model the power losses along the cables connecting the BS with the radiating antennas, as well as the power demand at each sector of 7 three-sector BSs. Using this data, we build different models. The results reveal that models tailored for specific technologies and transmission frequencies outperform more generalized ones. Additionally, linear regression models consistently show up to 96% better performance compared to MLPs, RNNs and literature benchmark based on [10], underlying a linear relation between regressors and power demand. Lastly, the most accurate power estimations are achieved using a linear model with traffic volume, load, maximum transmission power, and cable power losses as regressors. This model serves as a foundation for designing power-aware management strategies better aligned with the capabilities of contemporary BS technologies.

II. State of the Art

In this section, we describe existing models that quantify power demand for 4G Long Term Evolution (LTE) and 5G New Radio (NR) BS. The works presented in [10]–[12] schematize the structure of an LTE BS transceiver, identifying the key components: the antenna interface (actual antennas), the power amplifier (the most power-demanding component), the Radio Frequency (RF) transceiver section (which includes an UL receiver and a DL transmitter), the baseband unit (responsible for digital signal processing), the Direct Current to Direct Current power supply (which converts high DC voltage to the operational level), the cooling system (which maintains ambient temperature), and the main power supply (which connects the BS to the electrical grid).

In [12], the authors divided the RF section into two sub-components: the analog front-end, including analog baseband, RF system, frequency synthesizer, Digital-to-Analog Converter (DAC), and Analog-to-Digital Converter (ADC), and the digital control sub-component. Using this BS structure, in 2010 the Energy Aware Radio and neTwork tecHnologies (EARTH) project [10] proposed two models to quantify BS power demand based on measurements from LTE equipment: (i) a model defining the BS needed power at maximum load, and (ii) a model that scales linearly with the number of used radio resources. The model discussed in [13] formalizes BS power demand as the sum of load-dependent and fixed components. The fixed component is necessary to keep the BS active even if unloaded and is due to DC- and AC-DC conversions and cooling. Notably, pico BSs do not require cooling systems as their components naturally dissipate heat [13].

While these models focus on LTE BSs, power models for beamforming BSs in 5G NR are proposed in 2014 in [14]. As with LTE BSs, these models depend on the transmission power and include a static component, accounting for power used by the air conditioning system, backhaul link, rectifier, and other elements. In [11], the authors used an artificial MLP to model and estimate the power demand of a 5G BS. This model was built on a large dataset collected from 7 760 5G BSs, featuring RF configuration, number of carriers, traffic volume, load, and energy-saving strategies. Unfortunately, its parameters are not accessible and is difficult to make features intepretable and explainable. Finally, the models in [13] formalize the power requirements for Cloud-RAN and Heterogeneous Cloud-RAN, while [15] discusses power usage in Open-RAN.

As highlighted in the previous section, while these models have been largely adopted in the literature, there is a clear need for accessible and accurate new models that reflect the current state of BS hardware.

TABLE I: Technology per site.

Technology Frequency (MHz) Bandwidth (MHz)	LTE 700 10	LTE 800 10	LTE 1 800 20	LTE 2 100 15	LTE 2 600 15	NR 3 700 80
Site 1		✓				
Site 3		✓	✓	✓		
Site 4	✓	✓	✓	✓		
Site 6			✓	✓	✓	
Site 7		✓	✓	✓	✓	
Site 8	✓	✓	✓	✓	✓	
Site 9		✓	✓	✓		✓

III. Dataset Description

Since power demand plays a critical role in designing efficient mobile networks, recent BS technology advancements have enabled operators to obtain reliable power measurements. This allows them to monitor the power demand of each site and each piece of radio equipment directly from the Operating Support System (OSS), without the need for additional energy probes. Furthermore, this capability makes it possible to assess how each radio power amplifier contributes to the overall power budget of a radio site, facilitating an analysis of its relationship with traffic load and other network parameters. Leveraging this capability, power measurement readings are collected by the BS and transmitted to the OSS, typically on an hourly basis, along with other network counters used for network optimization, such as data traffic (in GByte or MByte) and traffic load (used Resource Blocks (RBs) vs. available RBs). The dataset collected from the OSS includes seven multi-frequency, multi-technology, three-sector sites located in Rome, Italy. Among these, 6 BSs employ the LTE technology, operating at $\{700, 800, 1\,800, 2\,100, 2\,600\}$ MHz. Additionally, one BS utilizes both the LTE and the NR technologies, transmitting at 3 700 MHz. The corresponding bandwidth, technology, and frequency per site are provided in Table I. For each sector of each radio module, the dataset specifies the maximum transmission power (dBm) and the losses (dB) of the cable connecting the BS to the radiating antenna. To summarize, each sector of each radio module is described by 1) technology {LTE, NR}, 2) carrier frequency F (Hz), 3) bandwidth B (Hz), 4) maximum transmission power $P_{\max}$ (dBm), and 5) cable losses A (dB).

The dataset covers one month from December 12, 2023, to January 12, 2024, with hourly granularity. For each hour, the following parameters are recorded: 1) UL traffic volume, V_{UL} (kByte), 2) DL traffic volume, V_{DL} (kByte), 3) number of available RBs in UL, 4) number of available RBs in DL, 5) percentage of used RBs in UL, RB_{UL} (%), 6) percentage of used RBs in DL, RB_{DL} (%), and 7) power demand (kW). Here, DL and UL are from the users' perspective, i.e. the former refers to the flow from the BS to the users, while UL refers to the flow from the users to the BS.

Power demand measurements aggregate data across different technologies and frequencies. Specifically, for sites hosting both 700 MHz and 800 MHz LTE BSs, as well as 1 800 MHz and 2 100 MHz LTE BSs, the power demand values are combined. Therefore, for these frequencies, we proportionally split it based on the UL and DL traffic volumes.

Fig. 1 relates the percentage of used RBs, considering both DL and UL, with the power demand, representing each sample with a marker. The x-axis corresponds to the percentage of DL and UL used RBs, while the y-axis corresponds to the power need. Specifically, Figs. 1b, 1a, 1c, 1d, 1e, and 1f depict these values for the LTE BSs that transmit at 700 MHz, 800 MHz, 1 800 Mhz, 2 100 MHz, 2 600 MHz, and the NR BS that uses 3 700 MHz, respectively, with different sites distinguished by different colors.

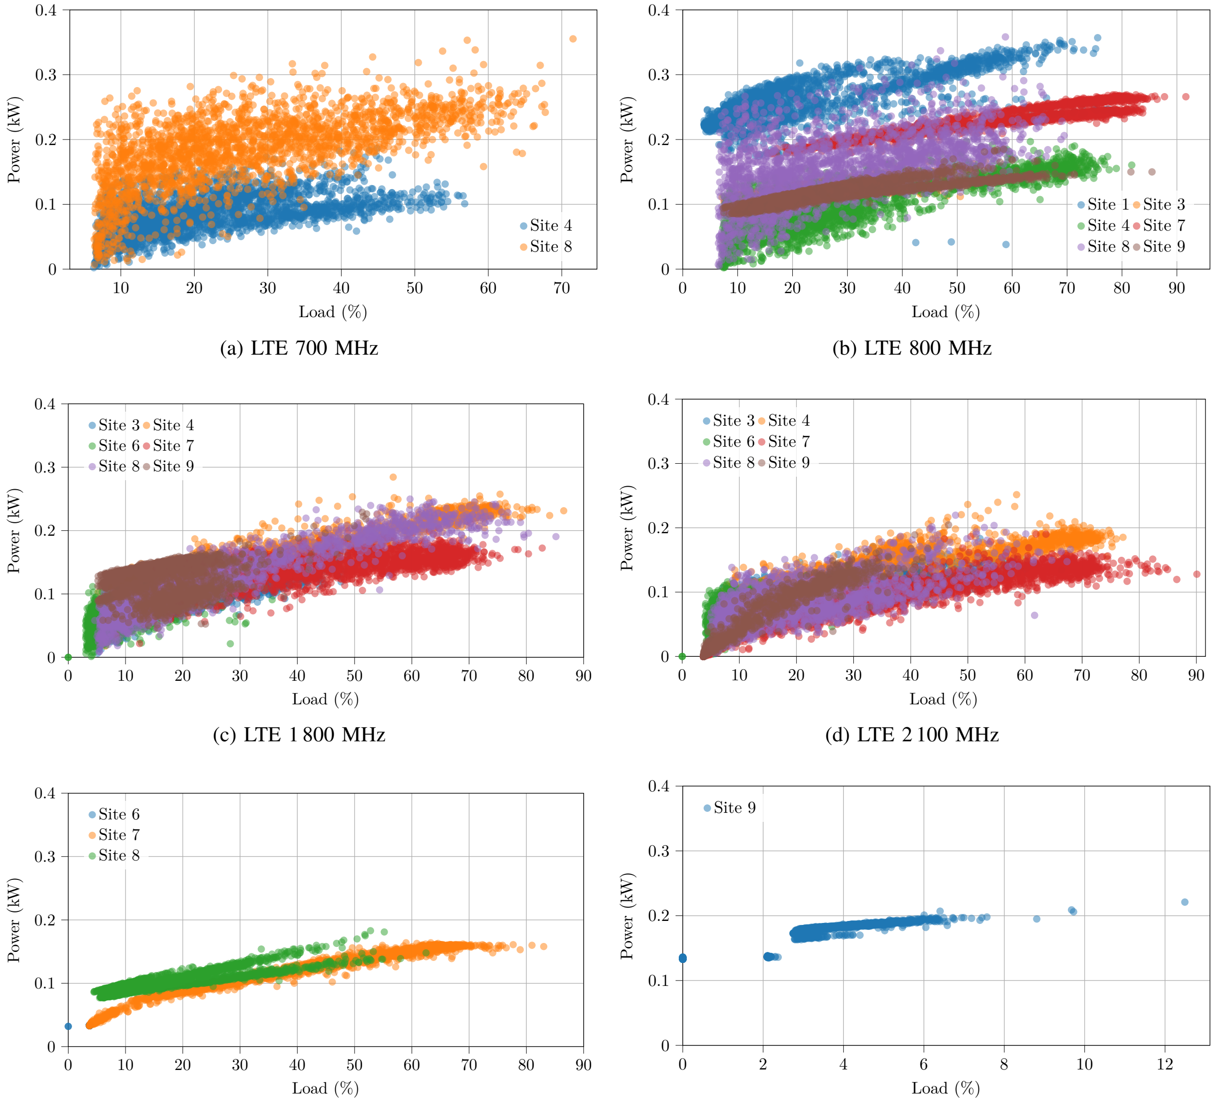

(a) LTE 700 MHz (b) LTE 800 MHz

(c) LTE 1 800 MHz (d) LTE 2 100 MHz

(e) LTE 2 600 MHz (f) NR 3 700 MHz

Fig. 1: Power demand and used RBs in UL and DL, for different technology and frequency.

The figures clearly illustrate the well-known linear relationship between used RBs and power demand, as widely recognized in the literature [11], [16]. Additionally, they highlight differences in power requirements among BSs operating with different technologies and frequencies, both in terms of absolute values and proportionality to used RBs. Moreover, BSs using the same technology and frequency can exhibit different power requirements, even when utilizing the same percentage of RBs. This variation may stem from differences in transmission power, measurement inaccuracies, variations in power loss along cables, hardware model and release date of the equipment, and small sample size (7 sites only).

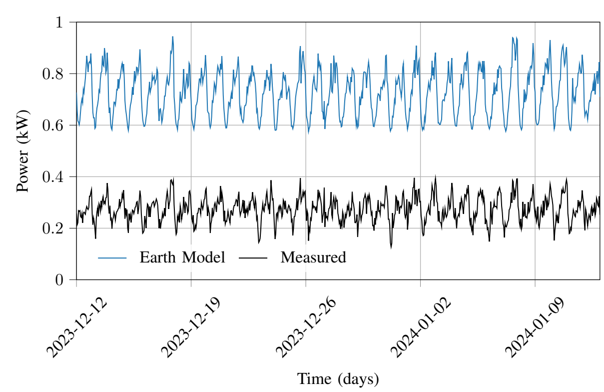

Fig. 2: Measured power demand and power demand computed by the EARTH model (in kW) for an LTE BS transmitting at 2 100 MHz.

Figure 2 shows the measured power demand (colored curve) and the power demand computed with the EARTH model (black curve), as in [10], for a LTE 2100 MHz BS. The figure reveals that the EARTH model systematically overestimates the power demand of the BS by up to 480%. This discrepancy arises because the EARTH model was developed in 2010 and fits the power demand patterns of BSs from that era and it does not account for the advancements in BS hardware in terms of power demand efficiency over the following years. A similar trend is observed for other BSs using different transmission frequencies and technologies. The corresponding plots are not shown here due to space constraints. We conclude that the EARTH model is no longer suitable for accurately quantifying the power demand of modern BSs, highlighting the urgent need for new models reflecting current power demand patterns. Therefore, this work aims to design a new model capable of quantifying the power demand of contemporary BSs.

IV. Proposed Power Demand Models

We build different models that use 60%, 15%, and 25% of the dataset as training, validation, and test set, respectively:

1) Linear Regression (LR)-based model: this model estimates the power demand of a BS as a linear combination of the inputs. We build different LR models, using different subsets of the features in the dataset, listed in section III, after standardization.
2) Multilayer Perceptrons (MLP)-based model: as in [11], we use two hidden layers, with 100 and 50 neurons respectively, using a linear activation function (relu function) and the Mean Squared Error (MSE) as loss function. As for the LR, we train different MLPs, differing for the subsets of standardized features of the dataset used as input features.
3) Recurrent Neural Network (RNN)-based model: in this case, we treat the problem as a time series prediction, exploiting the daily periodicity of the power demand, that reflects the daily periodicity of the traffic demand. Specifically, the model estimates the power demand at the current time t using an Long Short Term Memory Cell (LSTMC)-based RNN that receives the 24 previous samples as input, from $t-24$ to $t-1$. We consider an RNN composed of two hidden layers, each with 32 LSTMC units, and a dense output layer. We choose the MSE as the loss function.
4) Fixed EARTH (FE): from Fig. 2, it is evident that the EARTH model closely follows the data trend but has an inaccurate DC component. Therefore, we estimate the power demand as $aP_{EARTH} + b$, where P_{EARTH} is computed as in [10], and a and b are parameters estimated through linear regression. We use this adjusted model as a benchmark.

In this work, we use two different approaches for building a model for the power demand of a BS: (i) a model for all the BSs (M4A) and (ii) a model for each technology and frequency of the dataset (M4T).

V. Model Numerical Evaluation

In the first part of the numerical evaluation, we investigate the impact of using a single model for all the BSs (M4A), or a single model for all the BSs that implement a given technology and/or transmit at a given frequency (M4T). Figures 3a and 3b show the MAE, computed as in [17], for the M4A and M4T approaches, respectively, varying the input features of the LR-based models. The MAE is computed by grouping the BSs per technology, and different colors represent the different sets of input features used (we do not show all the possible combinations). These figures reveal that using a model per technology provides more accurate results than using a unique model for all the BSs: while with M4A the MAE ranges between 20 W and 94 W, it is never larger than 53 W if we use the M4T approach. Indeed, with M4T, each model better catches the relationship between the power demand and the input(s) characterizing each technology and/or transmitting frequency. We notice in Fig. 3b that, for LTE 800 MHz BSs and LTE 700 MHz BSs, the MAE is larger than for the others. This is because, for those technologies, high variance of the power demand measurements are observed for the same load (see Figs. 1b and 1a). When M4A is used, the model performance per technology is biased by the dataset's number of BSs of that technology. Indeed, the NR 3 700 MHz BS presents an MAE up to 93 W, 67% larger than 31 W, obtained for the LTE 1 800 MHZ BSs, which are the most numerous in the dataset (see Table I).

We now discuss the impact of the different inputs by looking at Figs. 3a and 3b. As expected, the model with the largest set of input features, namely RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A, $P_{\max}$, provides the highest accuracy, with an MAE up to 79% lower than the models with a single input feature. In the case of a single regressor, the use of the number of RBs instead of the traffic volumes yields a slightly better performance. At the same time, no significant difference is observed when uplink or downlink information is used, probably due to the high correlation of these data. For M4A, adding A and $P_{\max}$ to the regressors improves the model accuracy in most cases. Including A and $P_{\max}$ in the set of inputs adds in the M4A the information about the technology and improves its performance. Nevertheless, this does not affect the performance if M4T is used. We report the unstandardized coefficients for the M4T models that use RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A and $P_{\max}$ as input features in Table II, while for other input sets they are reported on GitHub[1]. From the table, we observe that, for LTE BSs, transmitting at $\{700, 800, 1\,800, 2\,100\}$ MHz, the intercept takes on negative values. This implies that when all input features (RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A, $P_{\max}$) are zero, the model predicts a negative power, which is unfeasible. This phenomenon may arise because the model is trained on a dataset with power demand values approaching zero, even when RB_{DL}, RB_{UL}, V_{DL}, V_{UL} A and $P_{\max}$ are non-zero (see Figs. 1a, 1b, 1c, 1d). This suggests that the model behaves non-linearly near zero for these input features, leading to a negative intercept as compensation. Therefore, when these input features are close to zero, we recommend using the model with inputs limited to RB_{DL}, RB_{UL}, V_{DL}, and V_{UL}, which yields a more meaningful result near the origin.

We now analyze the results obtained with the MLP-based models. Figures 4a and 4b show the MAE for the M4A and M4T approaches, respectively, obtained using different

[1]Standardized coefficients: https://github.com/valleGre1/Power-Consumption-Model-Standardized.git; Unstandardized coefficients: https://github.com/valleGre1/Power-Consumption-Model-Unstandardized.git

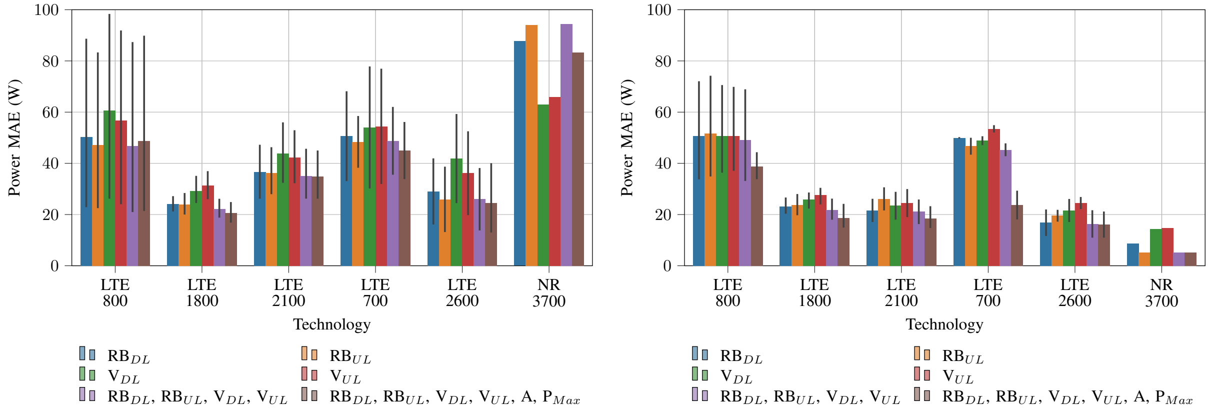

(a) M4A (b) M4T

Fig. 3: Mean Absolute Error (MAE) obtained with the LR model, using a single model for all the BSs, M4A (a), and a model for each group of BSs that transmit at the same frequency, M4T (b), with different input features.

TABLE II: Unstandardized coefficients of M4T models that use RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A and $P_{\max}$ as input features.

Technology	**RB$_{\mathrm{DL}}$**	**RB$_{\mathrm{UL}}$**	**V$_{\mathrm{DL}}$**	**V$_{\mathrm{UL}}$**	**A**	**P$_{\max}$**	**Intercept**
LTE 800	5.74×10^{-4}	1.72×10^{-3}	7.65×10^{-9}	0.00	6.12×10^{-1}	1.56×10^{-3}	-6.89×10^{-1}
LTE 1800	9.46×10^{-4}	8.70×10^{-4}	6.22×10^{-10}	0.00	5.66×10^{-2}	1.65×10^{-3}	-4.54×10^{-2}
LTE 2100	1.04×10^{-3}	0.00	1.23×10^{-9}	9.08×10^{-9}	2.11×10^{-2}	1.81×10^{-3}	-2.61×10^{-2}
LTE 700	2.05×10^{-4}	2.89×10^{-4}	2.24×10^{-8}	0.00	0.00	1.09×10^{-2}	-1.73×10^{-1}
LTE 2600	1.01×10^{-3}	9.06×10^{-4}	0.00	0.00	0.00	1.40×10^{-3}	1.58×10^{-2}
NR 3700	4.51×10^{-3}	6.17×10^{-3}	0.00	0.00	0.00	0.00	1.37×10^{-1}

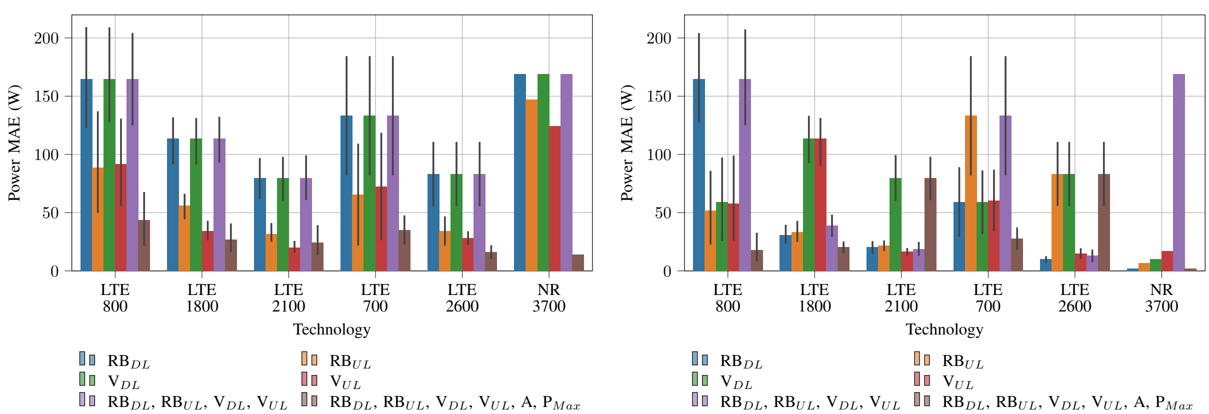

(a) M4A (b) M4T

Fig. 4: MAE obtained with MLP model, using a single model for all the BSs, M4A (a), and a model for each group of BSs that transmit at the same frequency, M4T (b), with different input features.

input features (indicated by different colors) and aggregated by BS technology and frequency. We observe that using RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A and $P_{\max}$ provides the best performance when a single model is trained for all BSs (M4A). Furthermore, employing a separate model for each BS technology and/or frequency (M4T) improves accuracy, especially when a single input feature is used. The results indicate that the lowest MAE value for each technology and frequency ranges from 2 W to 34 W, aligning with the values reported in [11]. Comparing these results with the ones in Figs. 3a and 3b, it is evident that the LR-based models provide more accurate models than MLP-based ones, that reach values of MAE up to 160 W, possibly due to the limited size of the training set.

The blue, orange, green and red bars in Figs. 5a and 5b show the MAE and MSE when the model is built using LR, MLP, RNN and FE, respectively, for each technology, using M4T. The LR and MLP models use RB_{DL}, RB_{UL}, V_{DL}, V_{UL}, A, and $P_{\max}$ as input features. The LR-based models provide the lowest MAE and MSE, never exceeding 38.8 W and 2.2 W, respectively, while with FE they increase up to 92 W and 9 W, respectively. With MLP and RNN, MAE and MSE grow up to 88 W and 15 W, respectively. As discussed above, the significant error introduced by MLP may be due to the limited size of the training set. Conversely, the RNN treats power demand estimation as a time-series forecasting problem. However, even though power demand follows a consistent

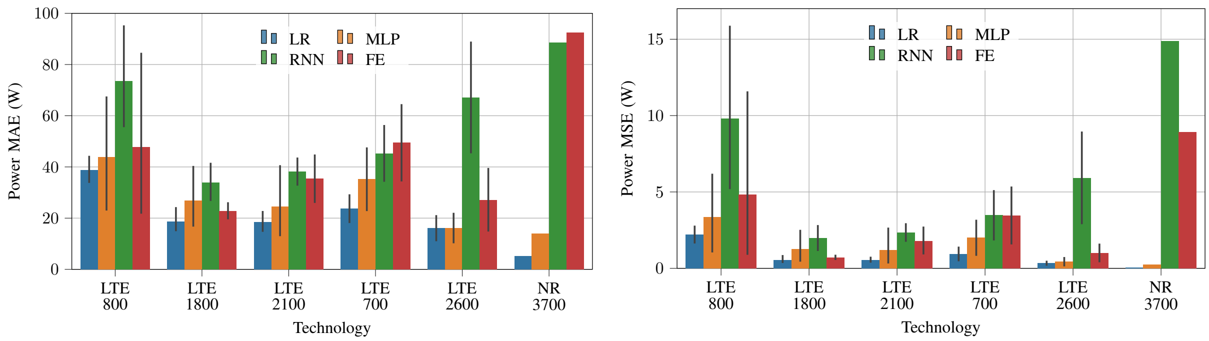

(a) MAE

(b) MSE

Fig. 5: MAE and MSE with LR, MLP, RNN and FE.

daily pattern, noise in the data (see Fig. 2) induces large errors in the forecast.

These findings highlight that the LR-based solution is well-suited to our problem and preferred over others because (i) it is simple, (ii) it is explainable, and (iii) it can be used in convex optimization.

VI. Conclusion

The objective of this paper is to address the lack of empirical models for estimating the power demand of modern BSs. To achieve this, we collect real measurements of traffic volume, load, and power requirement from various sites where three-sector LTE or NR BSs are installed, operating at different frequencies. Using this data and various methodologies, we build several models that differ in their input features and in whether non-linearity exists between these features and the estimated term. The results reveal that models specifically tailored for BSs using a given technology and transmission frequency are preferable to those without such distinctions. Additionally, models built using linear regression outperform those based on MLP, RNN, and Earth model [10], highlighting the linear relation between input features and power demand, with linear models exceeding the performance of MLPs and RNNs by up to 96%. Finally, incorporating traffic volume, load, maximum transmission power, and power losses along the cable to the BS power plug provides the most accurate power requirement estimates, with errors ranging from 4 W to 38 W. In future works, we plan to study other non-linear models while also considering additional hardware configurations and environmental conditions, such as the temperature. Additionally, we aim to apply these models to the RAN resource allocation problem to achieve power savings.

Acknowledgment

This paper was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”, Focused Project R4R).

References

[1] J. Wu *et al.*, “Energy-efficient base-stations sleep-mode techniques in green cellular networks: A survey,” *IEEE communications surveys & tutorials*, vol. 17, no. 2, pp. 803–826, 2015.

[2] M. Deruyck *et al.*, “Power consumption model for macrocell and microcell base stations,” *Transactions on Emerging Telecommunications Technologies*, vol. 25, no. 3, pp. 320–333, 2014.

[3] Keysight Technologies. (2023) Next generation wireless: A guide to the fundamentals of 6g. Accessed on: November 4, 2025. [Online]. Available: https://www.keysight.com/at/de/assets/7123-1050/ebooks/Next-Generation-Wireless-A-Guide-to-the-Fundamentals-of-6G.pdf

[4] A. Salh *et al.*, “Low computational complexity for optimizing energy efficiency in mm-wave hybrid precoding system for 5g,” *IEEE Access*, vol. 10, pp. 4714–4727, 2021.

[5] Y. Ramamoorthi *et al.*, “Resource allocation for comp in cellular networks with base station sleeping,” *IEEE Access*, vol. 6, pp. 12 620–12 633, 2017.

[6] W. Tan *et al.*, “Spectral and energy efficiency for uplink massive mimo systems with mixed-adc architecture,” *Physical Communication*, vol. 50, p. 101516, 2022.

[7] S. Alotaibi, “Key challenges of mobility management and handover process in 5g hetnets,” *International Journal of Computer Science & Network Security*, vol. 22, no. 4, pp. 139–146, 2022.

[8] G. Castellanos *et al.*, “Multi-objective optimisation of human exposure for various 5g network topologies in switzerland,” *Computer Networks*, vol. 216, p. 109255, 2022.

[9] H. Chen *et al.*, “Harvest-then-cooperate: Wireless-powered cooperative communications,” *IEEE Transactions on Signal Processing*, vol. 63, no. 7, pp. 1700–1711, 2015.

[10] G. Auer *et al.*, “D2. 3: Energy efficiency analysis of the reference systems, areas of improvements and target breakdown,” *Earth*, vol. 20, no. 10, 2010.

[11] N. Piovesan *et al.*, “Machine learning and analytical power consumption models for 5g base stations,” *IEEE Communications Magazine*, vol. 60, no. 10, pp. 56–62, 2022.

[12] C. Desset *et al.*, “Towards a flexible and future-proof power model for cellular base stations,” in *2013 24th Tyrrhenian International Workshop on Digital Communications - Green ICT (TIWDC)*, 2013, pp. 1–6.

[13] A. Israr *et al.*, “Power consumption analysis of access network in 5g mobile communication infrastructures — an analytical quantification model,” *Pervasive and Mobile Computing*, vol. 80, p. 101544, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574119222000049

[14] MAMMOET, “Massive mimo for efficient transmission: Deliverables 1.1, systems scenarios and requirements specifications,” 2014.

[15] Y. Al-Karawi *et al.*, “Power consumption evaluation of next generation open radio access network,” in *2024 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 2024, pp. 1–6.

[16] G. Auer *et al.*, “How much energy is needed to run a wireless network?” *IEEE Wireless Communications*, vol. 18, no. 5, pp. 40–49, 2011.

[17] T. O. Hodson, “Root mean square error (rmse) or mean absolute error (mae): When to use them or not,” *Geoscientific Model Development Discussions*, vol. 2022, pp. 1–10, 2022.