Predicting DC-Link Capacitor Current Ripple in AC-DC Rectifier Circuits Using Fine-Tuned Large Language Models

Mohamed Zeid, Subir Majumder, Hasan Ibrahim, Prasad Enjeti, Le Xie, Chao Tian,
Power Electronics & Power Quality Laboratory, Dept of ECE, Texas A&M University,
College Station, TX - 77845; enjeti@tamu.edu

Abstract—Foundational Large Language Models (LLMs) such as GPT 3.5 turbo allow users to refine the model based on newer information, known as "fine-tuning". This paper leverages this ability to analyze AC-DC converter behaviors, focusing on the ripple current in DC-link capacitors. Capacitors degrade faster under high ripple currents, complicating life monitoring and necessitating preemptive replacements. Using minimal invasive measurements from a full bridge rectifier and PFC-boost converter, we developed LLM-based models to predict ripple content in DC-link currents under noisy conditions. In this regard, based on simulations and experimental data from a full-bridge rectifier and a 1.5kW PFC, we have demonstrated the LLMs' ability for near accurate prediction of capacitor ripple current estimation. This study highlights the LLMs' potential in modeling nonlinear power electronic circuit behaviors and determining data requirements for precise circuit parameter predictions to predict component degradation and/or performance without any additional sensors. The final paper will have expanded results, capacitor ESR estimation based on fine tuned LLM output.

Index Terms—Power electronic Converters, Fine-tuning, Large Language Models (LLMs)

I. INTRODUCTION

The 2021 publication [1] effectively summarizes the expanding applications of artificial intelligence (AI), excluding the use of large language models (LLMs), in the field of power electronics. In the recent past, the generative nature of large language models (LLMs) and their ability to perform various natural-language processing tasks has been garnering significant attention in the scientific and industrial community [2]. Reference [3] details the application of LLMs in chip design at NVIDIA. References [4], [5] explore LLM applications in electronic design automation and power converter modulation design. References [6]–[10] discuss emerging applications of LLMs in electric power systems and smart grid. These diverse applications are enabled by transformer models, which are the foundation of LLMs, owing to their efficient and powerful pattern recognition capabilities across various tasks. This ability is facilitated by a vast number of model parameters, developed and modeled using extensive datasets. The influx of investment into LLM research and development highlights the industry's commitment to understanding and leveraging these capabilities. The premise is that if the LLMs are able to predict the next word in a sentence (a token, to be specific), they should also be capable of building a predictive model of

a nonlinear circuit with reasonable accuracy if the LLM is fine tuned using domain-specific data. Therefore, there is a strong potential for LLMs to be used to understand the nonlinear circuit behavior in the field of power electronics, a subject matter of this paper.
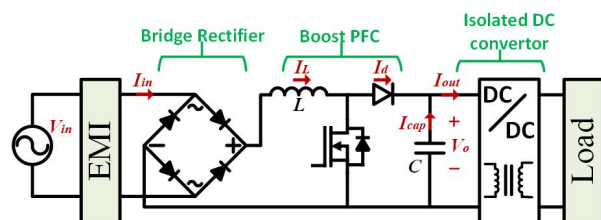

Figure 1: Single phase boost power factor correction (PFC) AC to DC rectifier schematic, V_{in}, V_o, and I_{cap} are the variables of interest

Converters based on single-phase Power Factor Correction (PFC) circuits are essential components in a wide range of applications, from laptop chargers to switching power supplies used in large data centers, because they enable unity power factor operations. These converters also have an immense growth potential from $2.3B in 2023 to $4.83B by 2031 [11]. As depicted in Fig. 1, like other AC/DC or DC/AC converters, DC-link capacitors play an important role in the PFC circuits because they help in smoothing out the voltages as the converters operate in various system conditions. In PFC circuits, DC-link capacitors behave as a decoupling element, bridging the low-frequency 50/60 Hz utility voltage and the high-frequency DC-DC conversion stage. However, these capacitors have a limited life span and can fail prematurely as they are subjected to various stressed operating conditions. One of the common causes of breakdown for power electronics circuits is the failure of capacitors [12]. The failure of a capacitor is not always immediately apparent, which can lead to increased strain on the remaining capacitors, accelerating their degradation and leading to eventual failure.

Capacitor life deteriorates faster under conditions of elevated operating temperatures, high humidity, over-voltage stress, pulsed discharge, and excessive ripple currents. In PFC circuits, capacitors experience heightened ripple current, exacerbated by real-world operational environments. This increased

ripple current elevates capacitor operating temperatures, accelerating their aging process [13]. Cyclic charging and discharging in the DC bus of a power electronic circuit induce periodic fluctuations known as ripple currents. This phenomenon results in power loss through heat dissipation, raising the internal temperature of capacitors and subsequently increasing their Equivalent Series Resistance (ESR), which further increases ESR in a positive feedback loop [14]. Prolonged exposure to elevated temperatures compromises capacitor reliability and longevity, shortening operational lifespan and accelerating degradation. Hence, understanding and accurately predicting ripple currents (a nonlinear phenomena) is critical in ensuring the reliability of power supply systems. Previous research has demonstrated the ability of using circuit harmonics in accurate ESR predictions [15]. This paper seeks to explore the practical
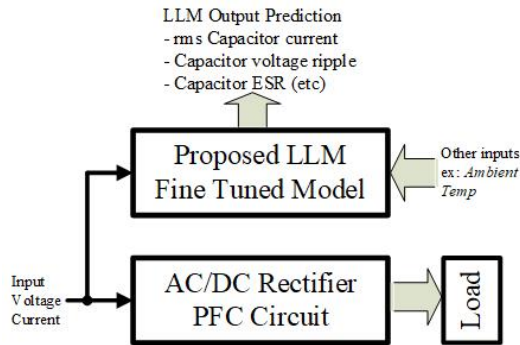

Figure 2: Block diagram of the proposed LLM fine tuned model for the ac/dc rectifier PFC circuit

applications of LLMs in predicting key behaviors of power electronic circuits, specifically ripple current content in the DC-link capacitors, under various operating conditions. Fig. 2 shows the block diagram of the proposed fine tuned LLM prediction approach. Our contributions are, therefore, twofold:

(i) Use of LLMs for the Prediction of Nonlinear Circuit Behavior: This study investigates whether LLMs, with sufficient training, can accurately model nonlinear behaviors in power electronic circuits. In this regard, we have tested with direct question answering, few-shot prompting [16], and LLM fine-tuning.

(ii) Data Requirements for Circuit Parameter Prediction: We performed both numerical simulations and hardware experiments to obtain both the Root Mean Square (RMS) and ripple content in the DC-bus capacitor voltage for various operating conditions.

II. USE OF LLM "FINE-TUNING" FOR MODELING THE BEHAVIOR OF POWER ELECTRONIC CIRCUITRY

The ability to learn from very few labeled samples is a basic attribute that separates machines from humans, and LLMs have shown a similar ability in power engineering applications [7]. However, due to the task-agnostic feature of pre-trained LLMs for natural language processing, LLMs do not do well in application-specific tasks, such as circuit analysis as can be seen in Fig. 3. Typically LLMs utilize self-consistency in their

responses, and in the current example, LLMs simply assume that DC bias as well as second, and fourth harmonic voltages do not change as the system load changes, which is incorrect.
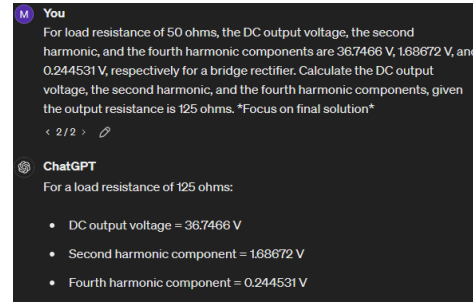

Figure 3: Example of inaccurate output in power electronic circuit analysis when using GPT 3.5, returned the same circuit values as the given example for a different circuit.
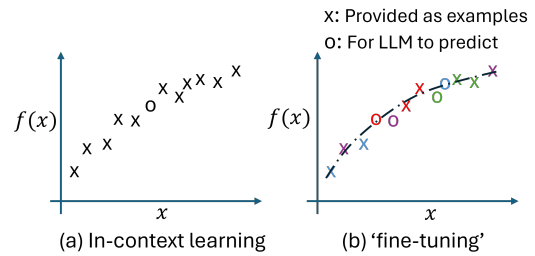

Figure 4: Motivating example for comparing in-context learning and fine-tuning for function mapping.

Next, we pose our problems as an unknown function map to alleviate the self-consistency issues regarding circuit analysis. As shown in Fig. 4(a), we provide x and $f(x)$ pairs to the LLM and query for an unknown value of x. Based on our observation, LLM (GPT 3.5 turbo in the current context) generated responses become consistent. Based on the simulated dataset containing DC, second, and fourth harmonic components corresponding to different loading conditions, it can be seen that as the LLM is provided with more and more examples, the accuracy of LLM-generated results changes. However, as shown in Fig. 5, the model accuracy may not monotonically improve as LLMs are provided with more data.

It is well-known from the literature that fine-tuning outperforms in-context learning without parameter updates [17]. Fine-tuning of LLMs involve the adaptation of pre-trained models through additional training on targeted datasets, where labeled data is provided for the model for pattern identification, thereby refining their capabilities for specific tasks or domains. This process enhances the model's performance by tailoring its general-purpose abilities to meet the nuanced demands of particular applications.

In this paper, we aim to bridge the gap where, at each training step, we expose the LLM to a few examples and compare LLM-generated responses with labeled known results to update the transformer model based on these errors. This process can be explained using Fig. 4(b). As shown in the

figure at every step (shown with different colors) we expose the LLM with a few examples, and a sample point, where the LLM can predict the function value. LLM model gets updated by comparing the LLM estimated value with actual measurements. Once the model is trained, it is evaluated on the testing set, a subsection of the data which contains unseen examples, in order to assess its performance and ability to generalize. Additionally, fine-tuning involves hyperparameter tuning to optimize the model's architecture and learning process further. Finally, after fine-tuning is complete, the fine-tuned LLM can be deployed for inference on new data, where it can generate predictions or perform other specified tasks with improved accuracy and effectiveness.
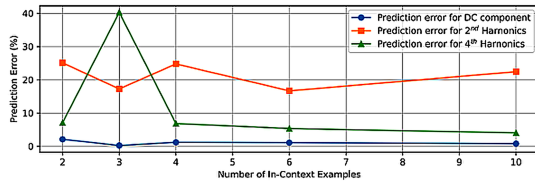

Figure 5: The effect of the number of examples on the prediction errors: As the LLM is introduced with more and more examples as contexts, the accuracy of LLM-generated results change in a non-monotonic way.


Figure 6: Sample of input prompt given to fine-tune the model, the same formatting was used for both the bridge rectifier and the Power Factor Correction circuits

A sample of an example prompt for fine-tuning is shown in Fig. 6. As in other deep-learning methods, learning by the LLM relies on several hyperparameters, such as the learning rate, batch size, and number of training epochs. Given that the LLMs are already pre-trained, and that the datasets are noisy, we need a balanced approach to the 'fine-tuning.' While a small learning rate is important in order to mitigate the effect of noisy data, it should not degrade the well-tuned weights of the pre-trained LLM. A larger batch size can reduce the variance in the computed gradient; the available hardware resources might also constrain the choice of batch size, and in our case, it gets more expensive. Thirdly, in the presence of noisy data, a larger number of training epochs may lead to overfitting the training data, which results in an increase in the training set performance and a decrease in testing performances. Therefore, one should focus on conservative updates that carefully adapt to new patterns in the data, leading to improved performance on realistic tasks without capitulating to noise-driven inaccuracies.

The metric identifying the discrepancy between the predicted outputs of the LLM and the actual labeled outcomes

during the fine-tuning process is identified as 'training loss.' Given that we utilized GPT-3.5-turbo, we utilized categorical cross-entropy loss for training [18]. Typically, the essence of the training process is to minimize this loss, which signifies a reduction in the error rate of the model's predictions against the ground truth. The optimization algorithms adjust the model's parameters (weights) to reduce the training loss, refining the LLM's ability to generalize well to unseen data. Fig. 7(b) demonstrates the training loss for the simulated full-bridge rectifier circuit. The training loss sharply decreases within the first few steps, and the loss converges. As shown in Fig. 7(a), fluctuations in the training loss during the fine-tuning phase reflect the LLM's ability to generalize unknown functions despite the dataset being noisy or complex. The overall convergence of the loss function indicates that the model learns about the underlying pattern within the dataset.
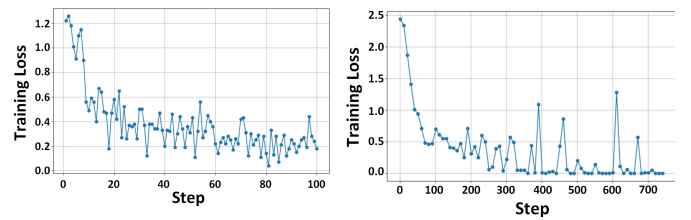
(a) Experimental bridge rectifier　　(b) Simulated bridge rectifier

Figure 7: Training loss for the Bridge rectifier fine-tuning: (a) Trained on the experimental data, constant peaks show pattern complexity or noise. (b) Trained on the simulated data, ideal relation which shows smooth convergence.

III. CAPACITOR CURRENT RIPPLE ESTIMATION

In this study, we are examining the behavior of two specific power electronic circuits: the bridge rectifier (Fig. 8) and the boost PFC circuit (Fig. 1). Specifically, we are analyzing the behavior of capacitor current for both circuits, varying the load conditions for the bridge rectifier and the input voltage for the PFC, essentially varying the power consumption of both circuits. Once the system reaches a steady state, the waveform is captured. For the experimental setups, we are using an oscilloscope along with differential probes for data measurements, a variable transformer for the power supply, and an electric load for the bridge rectifier circuit. Alternatively, the power electronic circuitry is simulated using PSIM software, and we are capturing the measurements using virtual probes. Fast Fourier transform is applied to the collected data to capture the RMS, second harmonic, and fourth harmonic current amplitudes for the bridge rectifier and PFC circuits.

To understand the accuracy of the developed model we divided the collected data into training and testing sets. While LLMs, as generative AI models, produce similar responses to the same prompt, there's still a small inherent randomness in their generation process, influenced by factors such as the model's internal state or the random sampling during generation. To address this variability, we employed a strategy

of running the model multiple times and averaging the output values. This approach helps to stabilize the predictions and mitigate the impact of randomness.
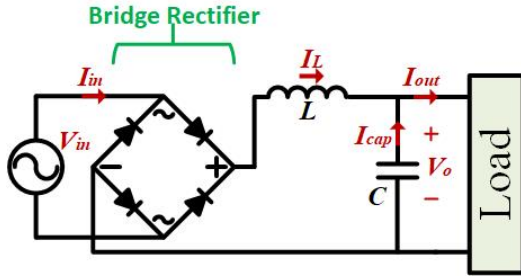
IV. ANALYSIS OF AC-DC RECTIFIER SYSTEMS

Figure 8: Single phase bridge rectifier schematic, V_{in}, V_o, and I_{cap} are the variables of interest.

A. Full-Bridge ac-dc Rectifier (Fig. 8)

As a proof of concept a full bridge rectifier circuit was first simulated using PSIM and then built in the lab following the schedmatic shown in Fig. 8 to assess the regressing capabilities of the LLM in power electronics circuits, the experimental setup of which is shown in Fig.10. The input as well as the capacitor voltage and currents, were recorded at steady state. Further computations were performed to obtain the current RMS, as well as the second and fourth harmonic components of the capacitor current. The harmonic components were obtained using the Fast Fourier Transform (FFT), a mathematical algorithm utilized for signal analysis, which decomposes a waveform into its constituent frequencies, facilitating the identification and quantification of harmonics within the signal [19].

The bridge rectifier circuit had the independent variable as the power consumption, which was varied as a result of changing the load in the experiment, while all other circuit components remained constant. A total of 50 different steady state data sets were collected for different resistance values. This data was then split into training and testing sets in order to build and assess the model performance respectively. The primary aim of the model is to predict the effect of changing the power consumption value on the current of the capacitor along with its' harmonics while maintaining an otherwise constant circuit. A total of 42 instances were used for the training purposes while 8 remained separated in the training stage as to test the resulting model capabilities.

Fig. 7 shows the training loss performance difference between the simulated and experimental models of the bridge rectifier. When comparing the training loss per step for the simulation data and the experimental data, the simulation loss function shows an initial steep decrease, followed by a general stabilization with occasional spikes, indicating moments where the model encountered anomalies in the simulated dataset. Notably, the loss in this graph rarely exceeds 0.5 after the initial drop, indicating that the simulated environment is

Power (W)	I_{cap} RMS (%)	I_{cap} 2nd Harmonic (%)	I_{cap} 4th Harmonic (%)
533	0.49	0.53	0.64
538	0.19	0.07	0.05
564	0.17	0.33	0.37
626	0.06	0.31	0.06
646	2.82	0.13	0.72
674	0.28	0.95	0.58
710	3.83	6.4	4.96
929	4.91	0.83	3.05

generally idealized. Conversely, the graph for the experimental data, although starting at a similar loss level, exhibits much greater variability and slower convergence, which is characteristic of real-world data's inherent noise and complexity. The loss values fluctuate more frequently and remain higher throughout the training process, highlighting the challenges of training models on experimental versus simulated data. These observations suggest that while the model trained on simulated data may perform well under ideal conditions, it may lack robustness when faced with real-world variability. In contrast, the model trained on experimental data, despite its fluctuations and slower convergence, may demonstrate better adaptability and robustness to diverse conditions, including non-ideal behaviors in the circuit.

$$\text{MAPE} = \left(\frac{1}{n} \sum_{i=1}^{n} \left| \frac{A_i - F_i}{A_i} \right| \right) \times 100\% \qquad (1)$$

The regression performance metric of choice is the Mean Absolute Percentage Error (MAPE), which is calculated using (1). MAPE calculates the average of absolute differences between predicted values and actual values, expressed as a percentage of the actual values, where A_i is the actual value, F_i is the predicted value, and n is the number of data points. This metric intuitively reflects the error magnitude in percentage terms, which is invaluable for comparing model performance across datasets of varying scales. Despite the relatively high training loss depicted in Fig. 7(a), the fine-tuning model performs exceptionally well as can be seen in Figs. 9 where the blue dots represent real data and the orange dots represent predictions. The fine tuning model achieves remarkable accuracy in its predictions, with none of the MAPE values exceeding 2% as can be seen in Table I.

B. Single phase boost PFC ac-dc Converter (Fig. 1)

In this section the performance of the boost PFC ac-dc converter is analyzed (Fig. 1. The capacitor current is analysed while the input power is varied similar to the bridge rectifier circuit. Similar calculations to those used for the bridge rectifier circuit to process the raw data through FFT in order to extract the capacitor current RMS as well as both the second and fourth harmonics from the measured data were applied.

Fig. 11 shows the training loss performance of the experimental PFC circuit. Initially, the fine- tuning model exhibits high variability in training loss, peaking at around 0.8, which
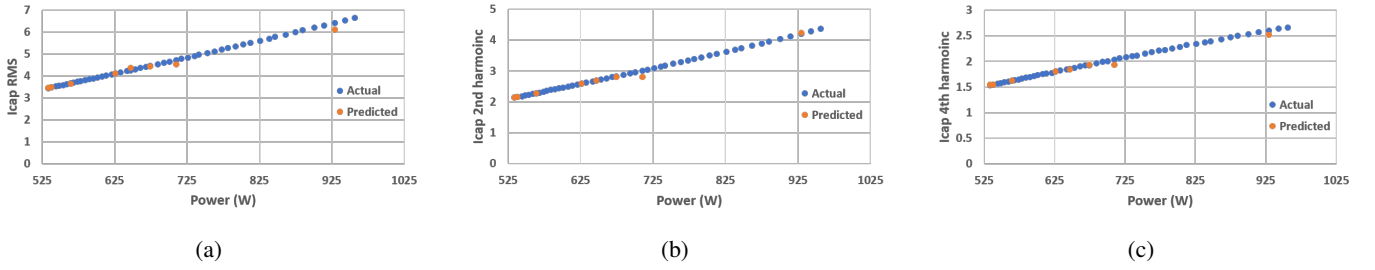
(a)　　　　　　　　　　(b)　　　　　　　　　　(c)

Figure 9: Experimental results: comparing Actual and Predicted values for the experimental FFT fine-tuned predictions for the bridge rectifier (Fig. 8). Blue dots are true values while orange are predicted using the fine-tuning model, all three graphs are smooth and have high prediction accuracy: (a) I_{cap} RMS values. (b) I_{cap} second harmonic. (c) I_{cap} fourth harmonic.
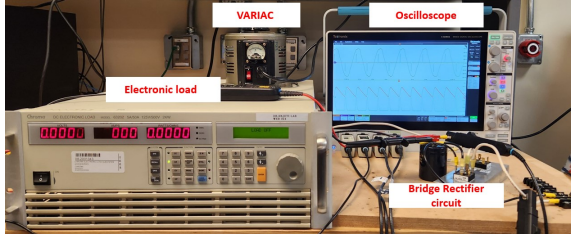

Figure 10: Bridge rectifier test setup. The figure shows the bridge rectifier circuit connected to an electronic load that precisely adjusts the load, a Variac for power supply, and an oscilloscope to monitor voltage and current.

Figure 12: Sample of collected waveform for the PFC circuit: (a) Shows the captured waveforms of I_{in} and I_{cap}, in red and green respectively. (b) Shows the computed FFT of I_{cap}, the second and fourth harmonics are the values of interest.
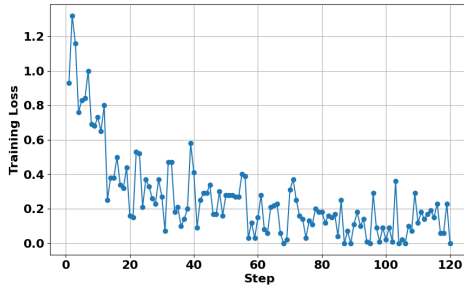

Figure 11: Fine tuning model loss of the PFC circuit as seen in Fig. 1. High fluctuations indicate noise or complexity of data, near convergence towards the later steps.

rapidly decreases within the first 20 steps. Following this, the loss demonstrates continued but less intense fluctuations, primarily ranging between 0 and 0.4. This moderation suggests an ongoing adjustment to noisy or complex data elements. Notably, the peaks of loss gradually diminish, with early training sessions reaching higher peaks, and later sessions stabilizing around 0.2. As the training progresses, an overall downward trend in loss values is observed, particularly evident beyond step 80, with a significant flattening and clustering of values below 0.3 towards the end. This behavior indicates that the model is approaching an optimal level of performance, adjusting more consistently to the intricacies of the training data and potentially nearing its learning limits given the current data and parameter settings.
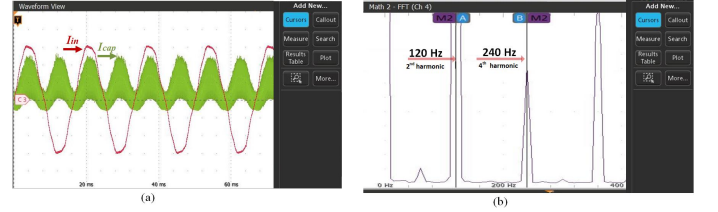
Table II contains the MAPE corresponding to the predicted

capacitor current second and fourth harmonics as well as the RMS of the experimental PFC circuit. All three variables have an average MAPE value below 3%: 1.2% for current RMS, 0.9% for current second harmonic, and 2.5% for current fourth harmonic. Fig. 12 (a) shows a sample waveform of the PFC circuit experimental data where the red waveform corresponds to I_{in} and the green waveform to I_{cap} while Fig. 12 (b) shows the computed FFT harmonics of choice. The high accuracy of the model can likely be attributed to the strong pattern recognition ability of the fine-tuning model, which can be seen from Fig. 13. The prediction performance of each of the different variables of interest are graphed in Fig. 13, where the actual true data is graphed in blue while the predicted values are graphed in orange. The final paper will have more extensive analysis by LLM on PFC and will include capacitor voltage analysis as well as ESR estimation.

TABLE II: MEAN ABSOLUTE PERCENT ERROR (MAPE) OF EXPERIMENTAL POWER FACTOR CORRECTION CIRCUIT (IN %) AS SEEN IN FIG. 1.

Power(W)	I_{cap} RMS(%)	I_{cap} 2nd Harmonic (%)	I_{cap} 4th Harmonic (%)
1151	2.75	2.97	0.31
1104	0.66	1.47	1.21
1286	2.12	0.87	3.98
1379	0.75	0.63	3.66
1258	2.26	0.14	1.52
1318	0.61	0.59	5.64
1272	0.22	0.16	2.02
1363	0.36	0.03	1.73

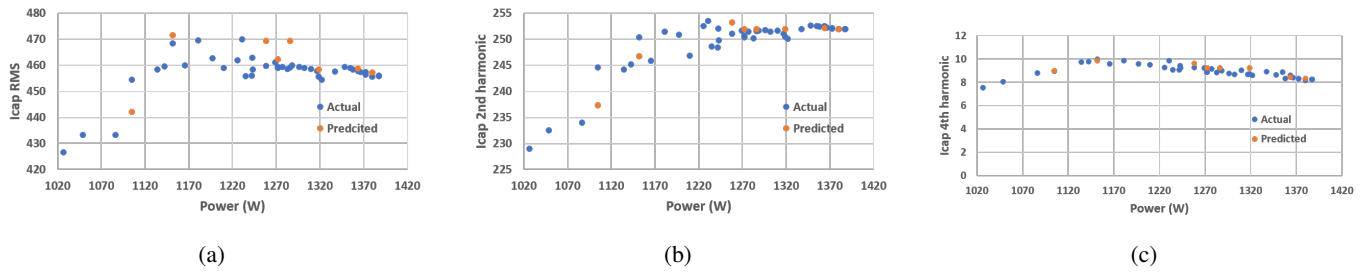|		
(a)	(b)	(c)

Figure 13: Experimental results: comparing Actual and Predicted values for the experimental FFT fine-tuned predictions for the single phase boost PFC (Fig. 1). Blue dots are true values while orange are predicted using the fine-tuning model, all three graphs are have high prediction accuracy despite the noise: (a) I_{cap} RMS values. (b) I_{cap} second harmonic. (c) I_{cap} fourth harmonic.

V. CONCLUSION

This paper explores the use of large language models (LLMs) to predict power electronic circuit behavior based on limited measurement data. Notably, the LLMs developed for the bridge rectifier and single-phase power factor correction (PFC) circuits achieved acceptable accuracy. Specifically: (a) The average mean absolute error for estimating the RMS values of the capacitor's second and fourth harmonic current components was less than 6%; (b) Experimental datasets obtained using the bridge rectifier circuit model showed an average prediction error of 1.59% for the I_{cap} RMS, 1.19% for the second harmonic, and 1.30% for the fourth harmonic component of the capacitor current; (c) Despite noisy training data (only 42 samples), the LLMs successfully captured the general trend; (d) The approach of modeling the prediction task as the mapping of an unknown function suggests broader applicability across different domains, demonstrating function generalization. In conclusion, the study highlights the potential of using LLMs with training sets to predict the rms capacitor ripple currents without the use of additional sensors. Further details on deducing capacitor ESR and heating will be presented in the final paper / conference presentation.

REFERENCES

[1] S. Zhao, F. Blaabjerg, and H. Wang, "An overview of artificial intelligence applications for power electronics," *IEEE Transactions on Power Electronics*, vol. 36, no. 4, pp. 4633–4658, 2021.

[2] K. M. Jablonka *et al.*, "14 examples of how llms can transform materials science and chemistry: A reflection on a large language model hackathon," 2023. [Online]. Available: https://arxiv.org/abs/2306.06283

[3] M. Liu *et al.*, "Chipnemo: Domain-adapted llms for chip design," *NVIDIA Research*, December 2023. [Online]. Available: https://research.nvidia.com/publication/2023-10_chipnemo-domain-adapted-llms-chip-design

[4] R. Zhong, X. Du, S. Kai, Z. Tang, S. Xu, H.-L. Zhen, J. Hao, Q. Xu, M. Yuan, and J. Yan, "Llm4eda: Emerging progress in large language models for electronic design automation," 2023. [Online]. Available: https://arxiv.org/abs/2401.12224

[5] F. Lin, J. Liu, X. Li, S. Zhao, B. Zhao, H. Ma, and X. Zhang, "Pe-gpt: A physics-informed interactive large language model for power converter modulation design," 2024. [Online]. Available: https://arxiv.org/abs/2403.14059

[6] C. Huang, S. Li, R. Liu, H. Wang, and Y. Chen, "Large foundation models for power systems," 2023. [Online]. Available: https://arxiv.org/abs/2312.07044

[7] S. Majumder, L. Dong, F. Doudi, Y. Cai, C. Tian, D. Kalathil, K. Ding, A. A. Thatte, N. Li, and L. Xie, "Exploring the capabilities and limitations of large language models in the electric energy sector," *Joule*, vol. 8, no. 6, p. 1544–1549, Jun. 2024. [Online]. Available: http://dx.doi.org/10.1016/j.joule.2024.05.009

[8] J. Ruan *et al.*, "Applying large language models to power systems: Potential security threats," 2024. [Online]. Available: https://arxiv.org/abs/2311.13361

[9] R. Xie, X. Yin, C. Li, N. Liu, B. Zhao, and Z. Dong, "Large language model-aided edge learning in distribution system state estimation," 2024. [Online]. Available: https://arxiv.org/abs/2405.06999

[10] A. Zaboli, S. L. Choi, T.-J. Song, and J. Hong, "Chatgpt and other large language models for cybersecurity of smart grid applications," 2024. [Online]. Available: https://arxiv.org/abs/2311.05462

[11] 360 Research Reports. (2022) Global power factor correction devices market report, history and forecast 2017-2028, breakdown data by companies, key regions, types and application. SKU ID: QYR-21058765; Publishing Date: 10-Jun-2022; No. of pages: 112. [Online]. Available: https://www.360researchreports.com/global-power-factor-correction-devices-market-21058765

[12] K. Zhao, P. Ciufo, and S. Perera, "Rectifier capacitor filter stress analysis when subject to regular voltage fluctuations," *IEEE Transactions on Power Electronics*, vol. 28, no. 7, pp. 3627–3635, 2013.

[13] S. J. Castillo, R. S. Balog, and P. Enjeti, "Predicting capacitor reliability in a module-integrated photovoltaic inverter using stress factors from an environmental usage model," in *North American Power Symposium 2010*, 2010, pp. 1–6.

[14] R.-J. Wai and C.-Y. Lin, "Dual active low-frequency ripple control for clean-energy power-conditioning mechanism," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 11, pp. 5172–5185, 2011.

[15] A. Deshpande, R. Sonigra, P. Enjeti, and P. Kumar, "A multi-harmonic system identification method for internal circuit element monitoring," 12 2023, pp. 114–119.

[16] Prompting Guide. Few-shot prompting. Prompting Guide. [Online]. Available: https://www.promptingguide.ai/techniques/fewshot

[17] K. Bhatia, A. Narayan, C. M. De Sa, and C. Ré, "Tart: A plug-and-play transformer module for task-agnostic reasoning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 9751–9788, 2023.

[18] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.

[19] H. J. Nussbaumer, *The Fast Fourier Transform*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1981, pp. 80–111. [Online]. Available: https://doi.org/10.1007/978-3-662-00551-4_4