MagLearn – Data-driven Machine Learning Framework for Modeling Magnetic Core Loss with Transfer and Few-shot Training

Lizhong Zhang, Tom McKeague, Binyu Cui, Navid Rasekh,
Jun Wang, Song Liu, Alfonso Martinez
University of Bristol
Jun.Wang@bristol.ac.uk

***Abstract*—In response to the MagNet Challenge 2023, this paper describes the model submitted to the competition by the University of Bristol team, which was awarded the 3rd Place Outstanding Performance among 24 competing teams worldwide. The core loss of magnetic components has been a challenge for the engineers to model due to lack of full physical models. Classic Steinmentz-Equation-based approaches show significant limitations under power electronics excitations. Data-driven approaches have emerged in the past years as a new solution to this problem, while the optimal method is still under exploration. Based on the datasets supplied by Princeton University, this work employs a machine learning framework to predict the core loss of magnetic components from a range of flux density waveforms, e.g. sinusoidal, rectangular, trapozoidal, as the input. The proposed approach builds on an LSTM network to extract features from the input waveforms and predict the power loss value. To cope with the small and imbalanced datasets supplied in the competition, special techniques are proposed in this work featuring transfer learning and few-shot training, which are realized through data augmentation and alignment. To decouple the output from the phase shift of the input waveform, a random shift/flip algorithm is applied in both pre- and post-processing blocks. The performance of the proposed approach is validated and evaluated through the experimentally measured testing sets in the competition, which demonstrates a very high accuracy.**

Keywords—MagNet Challenge, machine learning, neural networks, magnetic core loss, power electronics

I. Introduction

Nowadays, magnetic components are involved in most power electronic converters for functionality and filtering purposes. They are typically known to be the least efficient component that have a significant impact on system performance and efficiency in the size, weight and power loss factor [1][2]. However, an accurate core loss model for magnetic components that is based on the first principle remains elusive due to the non-linear feature of the magnetic material and other intercoupled factors such as dc-bias condition. Numerous research studies have been carried out to factor in the external parameters contributing to magnetic loss under nonsinusoidal excitations. The Steinmetz equation (SE), shown in (1), is widely accepted as an empirical model to calculate core loss under sinusoidal excitation. The k, α and β are the SE parameters which can be calculated by substituting the measured core loss value with the corresponding frequency and peak flux density values. This results in diminished precision of the equation due to the SE parameters demonstrating inconsistent performance across different frequency ranges.

$$P_{loss} = kf^{\alpha}B^{\beta} \quad (1)$$

To enhance the versatility and accuracy of the SE, the improved generalised Steinmetz equation (iGSE), as a modified solution, has been proposed based on SE for calculating core loss for arbitrary flux waveform under zero DC-biased condition [3]. The core concept of iGSE is to divide the complex waveform into individual B-H loops and calculate the loss respectively. However, these models generally face limitations in accuracy, particularly with certain waveform types, and tend to overlook the effects of temperature. Another approach called ‘Loss map’ is proposed in [4][5] for incorporating the pre-magnetization effects. To begin with, the operating state of one magnetic component can be described by three variables, the pre-magnetization state, the magnetic flux density swing and the flux density change rate. By measuring the B-H loop at various preset operating points, a core loss profile can be produced to cover all the operating conditions and used directly as a look-up table.

In recent years, methods using neural networks and other machine learning techniques have demonstrated excellent results in addressing nonlinear regression problems and forecasting time series data (e.g. image recognition) [6]. This technique could also be applied in core loss predicting and has proven to be more precise than the classic modelling based on the Steinmetz equation [7][8].

Three neural network models have been proposed and discussed in [7]. There is the ‘scalar-to-scalar’ model such as FNN which acts similarly to the Steinmetz equation that uses parameters such as flux density and frequency to directly predict the power loss. Overcoming the limitations of SE, the FNN-based model is proven to have a higher accuracy rate across the frequency range while covering external influencing factors such as temperature. One drawback of this model is that different models have to be trained according to the excitation waveform type which a set of scalar parameters could not fully represent. On the other hand, the ‘sequence-to-sequence’ model such as the transformer model [9] introduces the complete excitation waveform to the training process and predicts the magnetic response. While solving the issue of producing corresponding models to waveform types, the ‘sequence-to-sequence’ model normally introduces an enormous amount of parameters in both the input and output sides of the model which would lead to a longer training process and higher requirement for the training platform. To fit the sequence-to-scalars problem presented in the MagNet Challenge, the Long Short-Term Memory (LSTM) network shows strong potential. The LSTM network excels in processing regression problems with sequential input due to its key feature of capturing long-term dependencies in data and overcoming the short-memory issue prevalent in standard RNN models [10]. These characteristics are suitable for fulfilling the requirements for processing and analysing the time-series data *B(t)* and subsequently predicting the single-value output, the core loss density.

In response to the MagNet Challenge 2023, the contribution of this work is a novel machine learning framework based on LSTM network and transfer learning, which is easy to implement and outperforms almost all the competing models especially under the few-shot learning senario.

II. The MagNet Challenge

The aim of the challenge is to yield a “prediction model” for one magnetic core material, which takes in three inputs, *B, f* and *T* and outputs one volumetric loss density value. To achieve this aim, a machine learning process is expected to learn from the large database provided, which is experimentally measured loss and waveforms. This database is treated as the “ground truth” in this work. Hence the top-level idea of this work is illustrated in Fig. 1.

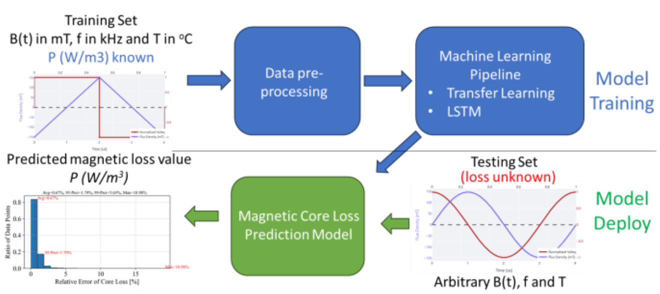

Fig. 1. The goal

The given datasets in the competition is designed to reflec real-life use senarios. There is a large-scale database of ten magnetic materials provided as the solid ground and starting point of this work. The number of samples given for the ten base materials are given in the table below.

Table I Number of samples given for 10 known materials

Material	**3C90**	**3C94**	**3E6**	**3F4**	**77**
Samples	40713	40068	6996	6564	11444
Material	**78**	**N27**	**N30**	**N49**	**N87**
Samples	11380	11396	8978	8602	40616

This large-scale database can serve as the foundation for the pre-training in a transfer learning process. Reflecting the real-world senario to adapt a trained base model to a new material, the data available for five unknown materials are given with a small and unbalanced/skewed training sets as illustrated in Fig. 2. The final testing data is given with only time-series inputs (*B*) to test out the model predictions.

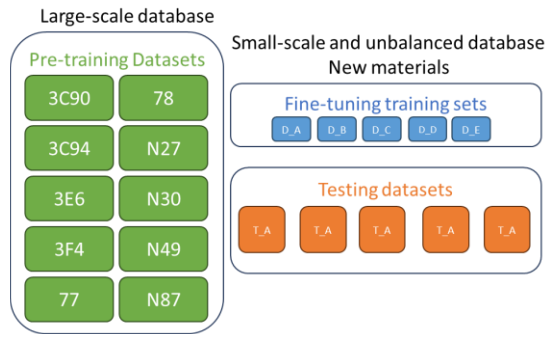

Fig. 2. The challenges in the given datasets

III. Machine Learning Framework

A. Overall structure/pipeline with Transfer Learning

The MagNet Challenge emulates practical use case by supplying rich datasets for the ten known materials and very limited and imbalanced datasets for the five unknown materials. To address this use case, this work adopts a transfer learning pipeline, which contains a pre-training stage and a fine-tuning stage. This pipeline is illustrated in Fig. 3. Starting from zero in the pre-training stage, the machine learning model is trained by a large dataset of one particular material. during the pre-training phase, all materials are first divided into training, validation, and test sets in a ratio of 8:2:1. During training, the model that performs best on the validation set is retained based on its results, not the training set results. This approach helps the model to avoid overfitting to the training set, which could result in significant errors.

As the most commonly-used and perhaps the most representative material in engineering, the training data of material 3C90 is used to generate a “fundamental model” as a generalised model to extract the key physical patterns for the general task of predicting the core loss value given the input data. The fundamental model is then tuned in the second-stage of pre-training against the ten materials’ training set, which yields 10 “base models” $\{M_i,\ i = 1\ldots10\}$. A relatively high learning rate (1e-3) and an adaptive learning algorithm (Adam) is employed in the training process.

In the fine-tuning stage, one of the base model is selected and tuned specifically for one particular unknown material (e.g. Material A) in the final test data. This stage starts with a selection logic to pick out the best tuned base models based on the minimal average error that one base model yields on the testing data, which can be considered as identifying one material out of the original ten materials that is most similar to Material A. At the end, a fine-tuned model for Material A is generated for deployment.

This pipeline is designed to make the most out of the rich training set of the original ten materials and transfer the learned knowledge into the fine-tuning step, which is a solution to the limited dataset for the testing data/materials. This idea is inspired from a machine learning concept called “meta learning”. The "meta" parameters of a neural network are pretrained using a similar material (3C90) and adapted to the target materials with further finetuning.

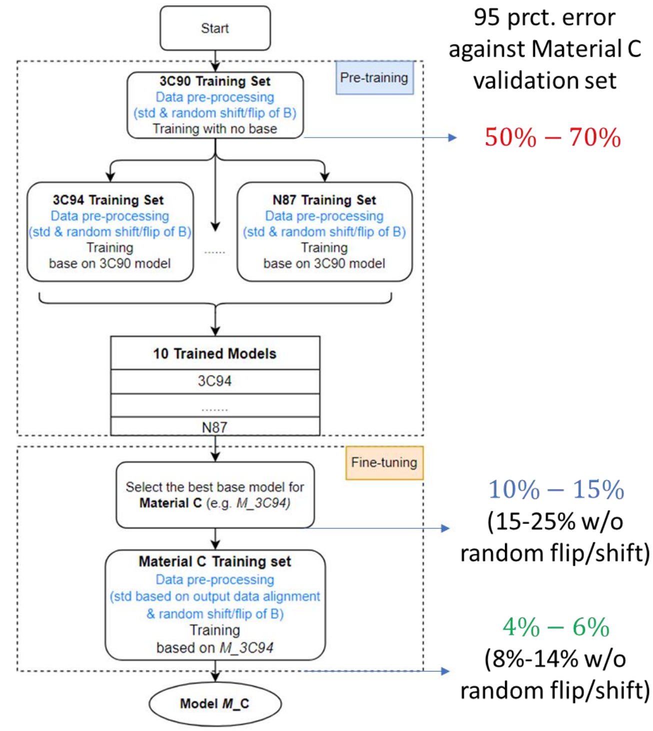

Fig. 3. Overall pipeline with transfer learning (example with 3C90 as the fundamental base and Material C as the target)

For the training process, the data are split into training set, validation set and a test set on a 70-20-10 basis. The split is performed randomly to minimize bias.

B. *Long Short-Term Memory (LSTM) network*

As the fundamental machine learning approach, a Long Short-Term Memory (LSTM) network is applied in this work, employing a 3-layer LSTM architecture with 90,653 parameters. This network is capable of processing the input time-series *B* waveform data to extract waveform features. As shown in Fig. 4, the LSTM network is dedicated to processing time-series waveform data, aiding in the preservation of waveform feature extractor parameters during transfer learning training, which enhances the efficiency of transfer learning. Subsequently, temperature and frequency data are inputted alongside waveform features into a Feedforward Neural Network (FNN). The neural network generates one scalar, the volumetric magnetic loss, as the output – the whole process is a sequence-to-scalar model.

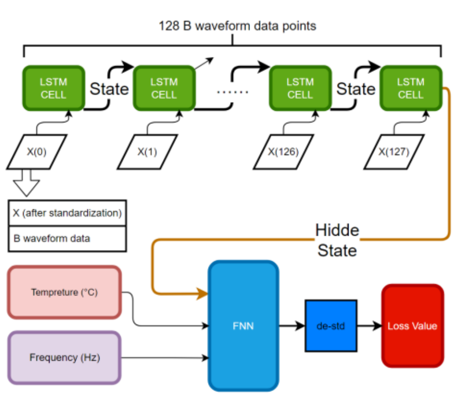

Fig. 4. Example of LSTM neural network structure

The data undergo normalization after FNN processing to compute the per-unit loss of magnetic materials. The FNN incorporates fully connected layers and employs the Exponential Linear Unit (ELU) as its activation function. The ELU function is continuous, aiding in maintaining output continuity and exhibits faster convergence during model training while avoiding saturation issues, unlike functions such as tanh.

C. *Loss function*

Ensuring effective training of the neural network and clearly defining its performance metrics are critical. Establishing a suitable performance evaluation standard is essential for accurately assessing the neural network model's capability in predicting magnetic material losses. Typically, the model's performance is evaluated by the low degree of its prediction error: the lower the percentage error of the model, the higher its performance is considered. The performance of the magnetic core loss model can be quantified by the 95th percentil error where the error is defined as

$$error = \left|\frac{pred - real}{real}\right| \quad (1)$$

In this context, determining the appropriate loss function for neural network training becomes crucial. A squared relative error loss function is selected in this work as

$$Loss = \left(\frac{pred - real}{real}\right)^2 \quad (1)$$

In this loss function, the square term is primarily used to avoid negative error values and to maintain the smoothness of the loss data. Additionally, by squaring the errors, the network is encouraged to minimize extreme values, which helps reduce the model's 95% confidence interval and the maximum error, thereby enhancing the reliability of the model's predictions. This approach not only aids in improving the overall performance of the model but also ensures the stability and credibility of the model's outputs.

IV. DATA PRE-PROCESSING

As marked in blue texts in Fig. 3, the training data has gone through a data pre-processing process before each training task takes place. To adapt input data for the neural network's requirements and ensure the data processing workflow promotes model training stability and efficiency, normalizing the input data is a critical step. It helps to avoid issues like gradient vanishing or explosion.

A. *Waveform down-sampling*

Initially, waveform data is processed through down sampling, for instance, reducing from 1024 data points to 128 using linear interpolation. This reduction lowers the model's computational load while it still maintains a level of accuracy without missing key patterns/details. This operation may also reduce the overfitting on irrelevant/insignificant details contained in the waveforms.

B. *Data standardization (linear scaling)*

The input variables temperature and frequency are linearly rescaled through $y=kx+b$ into values in the range of [0, 1]. The *B* waveform and the loss density are standardized to values are in the range of [-1, 1] while maintaining their linear relationship by setting *b* to zero. This procedure retains the polarity information and the linear correlation between the primary input

(B waveform) and the single model output (loss density) for better LSTM waveform feature extraction.

This data standardization process is observed to enhance the model performance, avoids numerical instabilities and improves model performance. For the LSTM, data standardization can improve the convergence in the gradient descent optimization process. In the case of standardising the *B* waveform, it can be considered as extracting the shape of the waveform and minimize the numerical impacts of the magnitude of data. The standardization process is illustrated in Fig. 5, where the raw data is processed into standardized data and stored together with the linear standardization coefficients, *k* (scaling factor) and *b* (bias) for each case. The scaling/standardization coefficients are determined for each material based the range of data, e.g. max value.

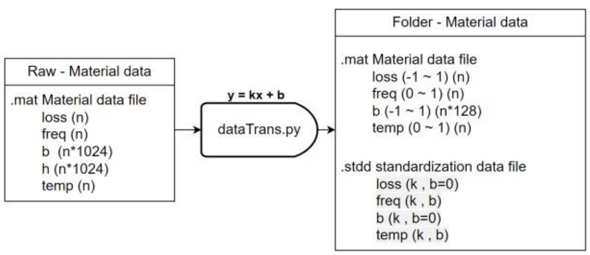

Fig. 5. Data standardization process (linear rescaling through $y=kx+b$)

C. *Random shift/flip of B waveform*

In given datasets, the input waveforms' phase all starts at zero, this differs from continuous waveforms in the real physical world, where the initial phase does not always begin at zero. To enhance the generalisation and avoid overfitting to the time-sequence waveform data, a data augmentation technique are implemented on the training data to apply random vertical and horizontal flipping as well as random phase shifting of the input *B* waveforms as illustrated in Fig. 6, before the datasets enter the training stage. The assumption is that shifting horizontally (i.e. applying a phase shift) and/or flipping (vertically and horizontally) the *B* waveform will not change the core loss density associated with this waveform – based on domain knowledge, this assumption should be valid given the core loss density is an averaged value over one whole cycle of a periodic *B* waveform repeating itself once every 1/f second. Note a rotation (pivoting) of the *B* waveform by an angle, instead of shifting or fliping, will invalid the euqal loss assumption. The phase shifting is similar to the approach in [11], while the flipping is an original contribution of this work.

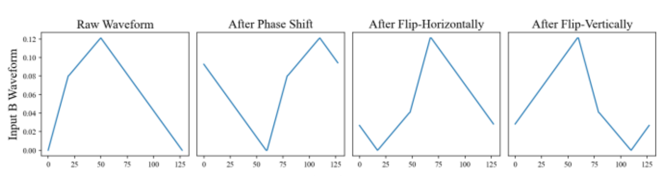

Fig. 6. Shifting and flipping operations of input *B* waveforms

These operations bring the following significant advantages in the machine learning model. 1. Decoupling the model from waveform phase dependency - through this approach, the model is not specific to inputs of a particular phase, thereby being adaptable to waveforms of any phase, or say insensitive to the phase of the waveform, which enhances the model's generalization capability. 2. Expand the training dataset - the random flip and phase shift effectively increase the size of the training set, providing the model with the opportunity to learn from a broader range of data variations. This helps the model to learn more robust waveform features, thereby improving its predictive performance.

Specifically, this data augmentation method equates to a 512-fold increase (22128) in the number of training samples, which not only significantly enhances the model's capability to process different waveforms but also strengthens its potential application in the real-world's complex environments. Through such a training strategy, the model is better equipped to understand and predict the magnetic losses caused by various waveforms of different phases and shapes, thereby enhancing the model's accuracy and reliability.

V. Few-Shot Training

In the testing stage of the challenge, the datasets for the five known materials are significantly small and skewed on waveform composition.

A. *Modified neural network*

Generally, training a neural network with a small dataset (less than 3000 samples) often leads to severe overfitting issues. To mitigate overfitting, it is crucial to select an appropriate neural network architecture, leverage pre-trained models, and employ effective data processing methods.

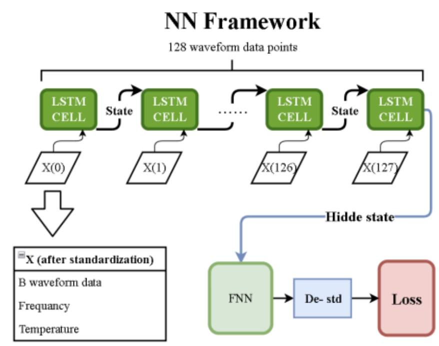

Fig. 7. Adjusted neural network structure for few-shot learning

This new neural network is specifically designed to address scenarios where the dataset is smaller than three thousand samples. It has fewer parameters, making it more suitable for training with smaller datasets. While this design necessitates some trade-offs in terms of accuracy, it significantly mitigates the problem of overfitting. The network inputs both waveform data frequency and temperature into an LSTM neural network.

B. *Base model selection*

Despite the network structure, training with a small dataset can still lead to overfitting. Thus, it is crucial to select an appropriate base model for transfer learning and employ effective data processing methods. As illustrated in Fig. 8, when training a new model *X*, one should first calculate the relative errors using ten base models and then identify the model with the smallest variance as the optimal base model for the new model. Subsequently, the model's normalization parameters are adjusted to align the data on magnetic loss.

C. *Output data alignment in the fine-tuning stage*

To cope with the small-training-set problem, a special "output data alignment" technique is applied in the fine-tuning stage as a data standardisation approach. Given the objective in the fine

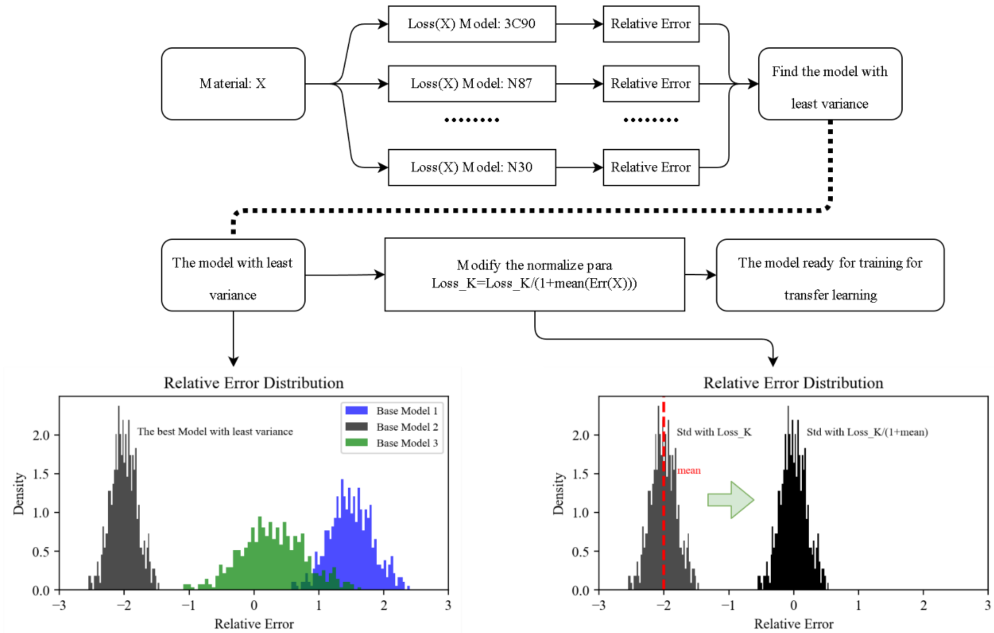

Fig. 8. Base model selection logic based on relative error distribution and output data alignment

tuning stage is to minimise the modification of the base model and unwanted overfitting towards the small training set, the output data (core loss density) to be fed into the training process is standardized in a different way – they are rescaled in a manner to align their extracted features with those of the original training data, to enhance the fine-tuned model. To align the relative error distrubition of the new datasets, a shift and rescaling is applied in the data standardization process to match the distributions in the source/training set as shown in Fig. 8.

D. *Data augmentation*

To cope with the cases with limited and unbalanced training set, a data augmentation approach is applied to artificially expand the training set in these cases. For example, the training set of Material D only contains 580 samples, with an unbalanced split between sinusoidal, trangular and trapezoidal waveforms (e.g. 145/400/35). However, the testing data of Material D is 7299 entries with a 61/2247/4991 split – the testing data has a large portion with trapezoidal waves, while the training set has very limited data for this case. To compensate this mismatch, the trapezoidal data in the training set is artificially duplicated and expanded to match the waveform split in the test set as much as possible, which leads to a 145/2000/700 split.

E. *Example on Material D*

The provided evaluation below shows an attempt to train a neural network for an unknown material, "Material *D*." As shown in Fig. 9, initializing the model with all-zero parameters or using the traditional '3C90' material as the base results in slow convergence during training and reaches a plateau without further significant error reduction. However, when the model is based on the '3F4' material, which has the smallest variance among the considered base models, it achieves lower error rates more quickly. Furthermore, when the '3F4' model is properly aligned with Material D through output data alignment, its convergence rate is significantly enhanced compared to using the unadjusted pretrained '3F4' model.

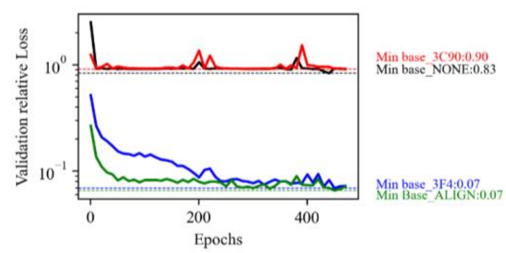

Fig. 9. Training loss of difference base models and with/without alignment

VI. Model Inference

A. *Data de-standardization*

At the end of model inference, a de-standardizaion process is applied to translate back the loss density value to the original scale based on the scaling factors (*k* and *b*) stored in the standardization stage.

B. *Averaging output with random shift/flip*

When the model is deployed, the shifting/flipping of the ***B*** waveform can still lead to variations of the predicted loss. To take this factor into account, for each entry of the test data, the *B* waveform is processed into 100 (adjustable) instances with different phase shifts, which yields 100 predicted volumetric loss values. These values are averaged in a post-processing step to obtain the final prediction for this entry.

VII. Results

A. *Base models*

The accuracy achieved for base models are listed below.

Table II Achieved accuracy for base models

Material	**3C90**	**3C94**	**3E6**	**3F4**	**77**
Avg (%)	1.41	0.97	0.72	1.00	1.45
95thErr (%)	3.68	2.77	1.64	2.84	3.09
Max (%)	11.08	12.11	18.81	8.59	8.08
Material	**78**	**N27**	**N30**	**N49**	**N87**
Avg (%)	1.35	0.91	0.56	2.87	0.89
95thErr (%)	3.07	2.53	1.41	7.66	2.63
Max (%)	11.38	9.83	24.02	29.93	8.61

It is observed that the average error for most models hovers around 1%. The 95% confidence interval is at 3%, and the maximum error is below 10%.

B. Fine-tuned models

Based on the above machine learning framework, five models are produced for the five unknown materials. The performance of these models is evaluated against the validation sets in the five cases, with the results shown in Fig. 10. Although the average of 95-percentile-error for the 10 known materials can achieve around 2.5%, the model's performance deteriorated in the cases of the five unknown materials due to the limited training sets. The average 95-percentile-error of the other four unknown materials achieves a 6.6%.

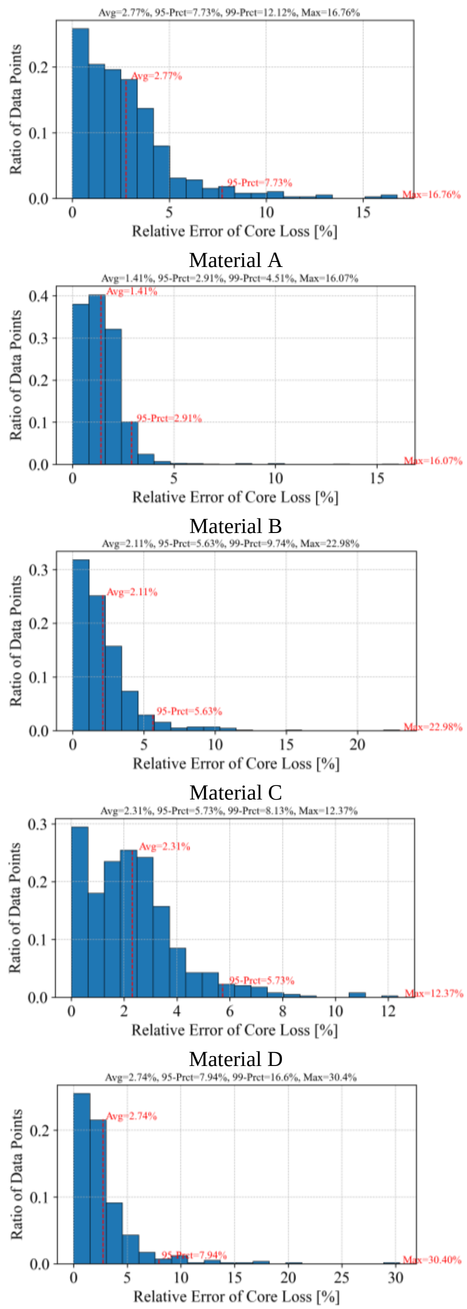

Fig. 10. Error distribution against the validation set for the target five unknown materials

The size of the five models are listed below in terms of number of parameters and file sizes.

Table III Model Size

	Material A	Material B	Material C	Material D	Material E
Number of parameters	15,653	90,653	90,653	16,449	16,449
Model Size	361 KB	361 KB	361 KB	70 KB	70 KB

VIII. CONCLUSION

This work has developed and demonstrated a machine learning pipeline in response to the MagNet Challenge 2023. An LSTM + FNN structure is applied as the core machine learning approach which features relatively low computation load. Several techniques are applied in the data pro-processing stage to enhance the model accuracy, reduce computation load and minimize bias/overfitting. Based on domain knowledge, a random shift/flip operation is applied on the *B* waveforms as pre-processing to desensitize the model from the sequence of the waveform data for feature extraction. A special data standardization and augmentation method is applied in the fine-tuning stage to realize few-shot training to cope with the small training set. Overall, excellent accuracy has been achieved by the proposed framework as evidenced by the 3rd Place performance validated against the experimentally measured testing sets. The submitted model outperforms all competing models in the few-shot learning senario (Material D), which resulted in a 15.9% 95th percentile relative error under the challenge of small and skewed dataset – in contrast, the second place's relative error is 20.6%.

REFERENCES

[1] K. Venkatachalani, C. R. Sullivan, T. Abdallah, and H. Tacca, "Accurate prediction of ferrite core loss with non-sinusoidal waveforms using only Steinmetz parameters," in Proc. IEEE Workshop Comput. Power Electron., 2002, pp. 36–41.

[2] C. R. Sullivan, "High frequency core and winding loss modeling," in Proc. IEEE Int. Electric Mach. Drives Conf., 2013, pp. 1482–1499.

[3] J. Li, T. Abdallah, and C. Sullivan, "Improved calculation of core loss with nonsinusoidal waveforms," in Proc. Conf. Rec. IEEE Ind. Appl. Conf. 36th IAS Annu. Meeting, vol. 4, 2001, pp. 2203–2210.

[4] J. Wang, K. J. Dagan, X. Yuan, W. Wang and P. H. Mellor, "A practical approach for core loss estimation of a high-current gapped in-ductor in PWM converters with a user-friendly loss map," *in IEEE Transactions on Power Electronics*, vol. 34, no. 6, pp. 5697-5710, June 2019.

[5] S. Iyasu, T. Shimizu and K. Ishii, "A novel iron loss calculation method on power converters based on dynamic minor loop," *2005 European Conference on Power Electronics and Applications*, Dresden, Germany, 2005, pp. 9 pp.-P.10.

[6] J. Schmidhuber, "Deep learning in neural networks: An overview," Neural Netw., vol. 61, pp. 85–117, Jan. 2015.

[7] H. Li et al., "How MagNet: Machine Learning Framework for Modeling Power Magnetic Material Characteristics," in *IEEE Transactions on Power Electronics*, vol. 38, no. 12, pp. 15829-15853, Dec. 2023, doi: 10.1109/TPEL.2023.3309232.

[8] N. Rasekh, J. Wang and X. Yuan, "Artificial Neural Network Aided Loss Maps for Inductors and Transformers," in *IEEE Open Journal of Power Electronics*, vol. 3, pp. 886-898, 2022, doi: 10.1109/OJP EL.2022.3223936.

[9] A. Vaswani *et al.*, "Attention Is All You Need," *arXiv.org*, Dec. 05, 2017. https://arxiv.org/abs/1706.03762

[10] X. Li and X. Wu, "Constructing long short-term memory based deep recurrent neural networks for large vocabulary speech recognition," in Proc. IEEE Int. Conf. Acoust., Speech, Signal Process., 2015, pp. 4520–4524.

[11] H. Li et al., "How MagNet: Machine Learning Framework for Modeling Power Magnetic Material Characteristics," in IEEE Transactions on Power Electronics, 2023, pp. 15829-15853.