Proceedings of the 3rd International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2025.26376

GPU Computing Resource Allocation and Prediction Based on Machine Learning Algorithms

Liao Hu

University of Illinois at Chicago, Chicago, USA
lhu231@my.trine.edu

Abstract. In this paper, a transformer-optimized BiGRU algorithm is proposed, which provides a more reliable technical path for the fine-grained allocation of GPU computing resources in deep learning models. In the study, ExtraTrees, CatBoost, XGBoost, and LightGBM were selected as comparative models. The results show that the proposed Transformer BiGRU model performs best in deep learning GPU resource regression prediction. Among them, the MSE of this model is 2.208, which is significantly lower than 4.414 of ExtraTrees, 6.71 of CatBoost, 4.464 of XGBoost, and 2.942 of LightGBM; RMSE was 1.486, which was lower than 2.101, 2.59, 2.113, and 1.715 for other models; MAE is at its lowest level of 0.969, which is better than XGBoost's 1.442 and LightGBM's 1.296; MAPE was 26.938, lower than 34.478 for ExtraTrees, 37.379 for CatBost, and 36.679 for XGBost; R^2 reaches 0.855, which is higher than LightGBM's 0.803 and XGBoost's 0.69. This indicates that the model has smaller prediction error and better GPU resource fitting effect. Its research value lies in laying an important foundation for improving the accuracy and scientificity of GPU computing resource allocation, and has positive significance for promoting the efficient utilization of resources in the field of deep learning.

Keywords: Transformer, BiGRU algorithm, GPU computing resource prediction

1. Introduction

With the rapid development of large model training and industrial intelligence, the global demand for computing power has shown an explosive growth [1]. The parameter scale of large models represented by GPT series, LLaMA, etc. has jumped from tens of billions to trillions, and the GPU time required for a single round of training can reach millions of hours. In industrial scenarios such as intelligent manufacturing and autonomous driving, the dynamic demand for computing power of real-time inference tasks presents high-frequency wave characteristics [2]. As the core carrier of deep learning computing power, the allocation efficiency of GPU resources directly restricts the process of technology landing. Under the current static allocation mode, the GPU utilization rate is generally lower than 30%. In some scenarios, there is even a contradiction between "idle computing power" and "task queuing", which not only causes a great waste of hardware resources, but also prolongs the model training cycle by more than 30% and increases the energy cost by about 40% [3]. The mismatch between supply and demand of computing power has become the core bottleneck

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.26376

restricting the large-scale application of AI technology, and it is urgent to realize the dynamic optimization of resource allocation through intelligent means.

Machine learning algorithms provide key technical support for solving the problem of GPU computing power allocation. The traditional rule-based allocation strategy is difficult to capture the nonlinear characteristics of computing power requirements, but machine learning models can construct the mapping relationship between computing power requirements and task attributes by deeply mining historical task data (such as model type, batch size, iteration number, hardware load, etc.). For example, the random forest algorithm can identify the key factors affecting the computing power demand through feature importance assessment, while the LSTM model can capture the time dependence during task execution to realize short-term computing power fluctuation prediction [4]. The application of these algorithms shifts the resource allocation from "experience-driven" to "data-driven", which can increase the GPU utilization to more than 60% and shorten the training cycle by 20%-50% in typical scenarios. By predicting the computing power demand curves of different tasks in real time, the system can schedule resources in advance, avoid redundancy or gaps in computing power, provide accurate decision-making basis for dynamic allocation, and reduce the probability of task congestion caused by resource contention [5].

Although the existing machine learning algorithms have made certain achievements in computing power prediction, there is still the problem of insufficient prediction accuracy in the face of long-term dependence and sudden demand in complex scenarios. The LSTM model exhibits significant memory decay on long sequence data, with prediction errors exceeding 25% in multitasking concurrent scenarios; Traditional temporal models are difficult to efficiently extract the global correlation between task features and computing power requirements [6]. To this end, this paper proposes a transformer-optimized BiGRU algorithm, which enhances the ability to capture long-term dependencies between task attributes and computational power requirements by introducing Transformer's self-attention mechanism. At the same time, the bidirectional recursive structure of BiGRU is used to improve the local time feature extraction efficiency. This fusion architecture not only retains the sensitivity of temporal model to dynamic changes, but also has the ability to model global features, which provides a more reliable technical path for realizing fine-grained allocation of GPU computing resources.

2. Data sources

This dataset contains 517 rows and 12 columns of data, with a total of 12 features. Some of the content of the dataset is shown in Table 1. This dataset is related to the GPU computing power requirements of deep learning models. The data covers information such as model type, model parameter size (in millions), batch size, input dimension, number of layers, iteration count, accuracy, optimizer type, whether regularization is used, whether data augmentation is used, learning rate, and GPU requirements (in gflops). The model types include RNN, CNN, ResNet, Transformer, etc., with accuracies of FP16, FP32, and optimizers including Adam, AdamW, RMSprop, etc. These data can be used to construct machine learning regression models to predict the computing power requirements of GPUs, in order to support related research and applications such as GPU computing resource allocation for deep learning tasks.

DOI: 10.54254/2755-2721/2025.26376

Table 1. Selected partial dataset

Model type	Parameters million	Batch size	Input dimension	Num layers	Iterations	Precision	Optimizer	Learning rate	Gpu demand gflops
RNN	29.89	1472	409	13	124.83	FP16	Adam	0.002335	1.89762
CNN	17.98	201	619	39	150.33	FP32	AdamW	0.008076	1.88104
ResNet	33.72	854	630	15	336.83	FP32	AdamW	0.0006	9.89094
Transformer	67.93	722	660	29	231.84	FP32	RMSprop	0.008989	10
Transformer	16.65	1423	905	29	1608.08	FP16	Adam	0.006511	10

3. Method

3.1. Transformer

Transformer is a sequence modeling architecture based on self attention mechanism, and its core breakthrough lies in breaking away from the sequence dependency of models such as RNN and achieving parallel processing. Its structure consists of an encoder and a decoder: the encoder receives an input sequence and converts it into context aware feature vectors through N stacked encoding layers; The decoder generates the final output based on the encoder output and its own generated sequence [7]. The network structure of Transformer is shown in Figure 1.

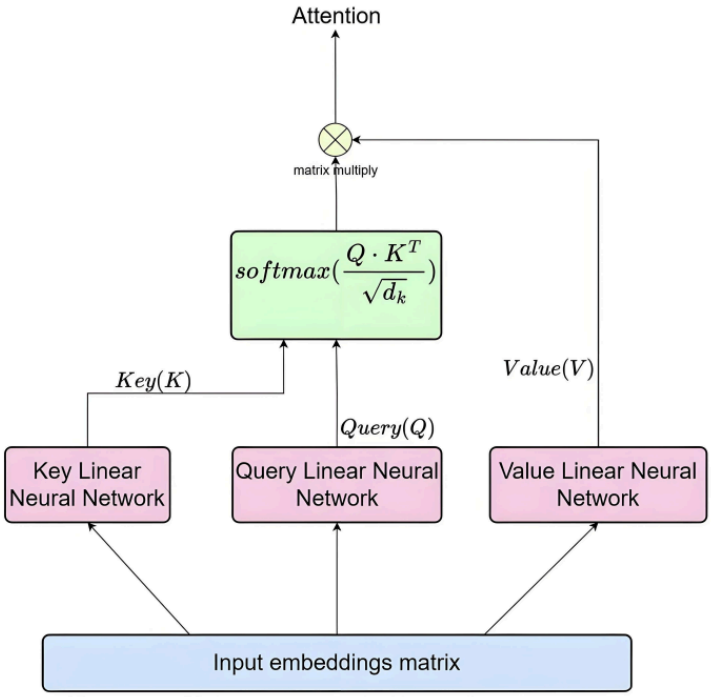

Figure 1. The network structure of Transformer

Each encoding layer includes a multi-head self-attention mechanism and a feedforward neural network: the multi-head self-attention maps inputs to multiple sets of query (Q), key (K), and value (V) matrices, calculates the relevant weights of different location tokens, and then connects multiple

DOI: 10.54254/2755-2721/2025.26376

sets of results to capture multidimensional dependencies; The feedforward network independently performs linear transformation and non-linear activation on each token. The decoding layer additionally includes encoder decoder attention, allowing the generation process to focus on key information of the input sequence. In order to preserve the order of the sequence, Transformer introduces position encoding, which injects position information into tokens at different positions through sine and cosine functions, solving the problem that self-attention is insensitive to order. Compared with the serial mode processed verbatim by RNN, Transformer can calculate the attention of the entire sequence in parallel, which greatly improves the training efficiency. This architecture, with its ability to capture remote dependencies and parallelism, has become the basis for pre-trained models such as BERT and GPT, reshaping the technical paradigm of NLP and even multimodal tasks [8].

3.2. BiGRU

The core foundation of BiGRU is GRU, which stands for Gated Recursive Unit. The network structure of BiGRU is shown in Figure 2. GRU aims to solve the problem of disappearing or exploding gradients in traditional recurrent neural networks when dealing with long sequences. The key is to flexibly control the transmission and forgetting of information through gating mechanism. Wherein the reset gate is used to determine whether to ignore the historical information. When the reset gate value is low, the model pays more attention to the current input; The update gate determines how much historical information is retained and how much new information is included, similar to the memory update mechanism [9]. This design enables GRU to efficiently capture long-term dependencies in sequences while simplifying the structure of recurrent neural networks and reducing the number of parameters. BiGRU introduces a bidirectional mechanism based on GRU, consisting of two GRUs in opposite directions. The GRU processes the data in the original order from left to right, capturing information from the past to the present; Another handles it upside down from right to left, capturing information from the future to the present. In each processing step, the two GRUs will generate corresponding outputs separately, and then combine the results of the two by cascading or fusion to form a comprehensive feature containing a bidirectional context. This bidirectional design allows the model to exploit all the information before and after a certain position in the sequence simultaneously, breaking through the limitation that unidirectional models can only rely on historical information [10].

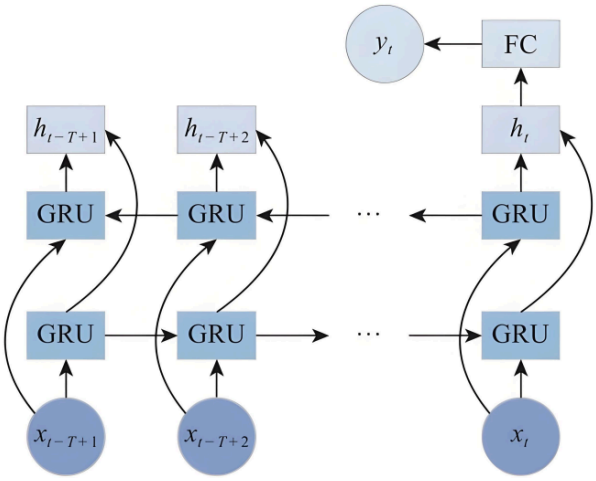

Figure 2. The network structure of BiGRU

DOI: 10.54254/2755-2721/2025.26376

The advantage of BiGRU lies in balancing the comprehensiveness of information capture with computational efficiency. Compared to unidirectional GRU, it can more accurately understand contextual information, and compared to more complex bidirectional LSTM structures, it has fewer parameters and faster training speed, making it more practical in resource limited scenarios.

3.3. Transformer-BiGRU

The principle of optimizing BiGRU with Transformer is mainly based on compensating the inherent shortcomings of BiGRU in long sequence modeling and parallel computing, while retaining its local time-aware advantages. As a bidirectional gated recursive unit, BiGRU captures sequence dependencies through forward and backward hidden layers. However, relying on a loop structure requires computations to be performed sequentially in time steps, resulting in poor parallelism and low training efficiency. Furthermore, information decay is prone to occur in long sequences, which makes it difficult to model long-range distance correlations. Transformer breaks this limitation through its self-attention mechanism: its core self-attention module can directly calculate the association weights of any two positions in the sequence, and process all time positions in parallel through matrix operation, effectively capturing global dependencies and solving the problem of long-distance information loss; The multi-head attention mechanism further maps features to multiple subspaces, learning association patterns from different scales, and enriching the dimensions of feature representation.

In the optimization practice, this paper adopts a hybrid architecture: BiGRU is used to process local short-term temporal features, and the output is used as the input of Transformer. Global dependencies are modeled through Transformer's self-attention layer, while residual joining and layer normalization are combined to enhance training stability. This not only preserves BiGRU's sensitive capture of local context, but also leverages Transformer's global modeling capabilities and parallelism efficiency to improve overall sequence modeling performance.

4. Result

In terms of experimental parameters, the dataset is divided into training and testing sets in a 7:3 ratio, and the input feature dimension is determined by the number of features in the dataset (output dimension is 1). The model includes a sequence input layer (with the number of channels equal to the number of features), a position embedding layer (with a maximum position encoding of 512), 2 self attention layers (4 heads, 32 key channels per head), a forward GRU layer (6 units), a reverse GRU layer (10 units, input after flipping the data by FlipLayer), as well as ReLU activation layer, 0.01 dropout layer, fully connected layer (output dimension 1), and regression layer. The training uses Adam optimizer with a maximum round size of 200, batch size of 256, initial learning rate of 0.01 (reduced to 0.001 after 80 rounds), L2 regularization coefficient of 0.001, and gradient clipping threshold of 10,. In terms of software and hardware configuration, the software relies on MATLAB and its deep learning toolbox. The hardware is equipped with NVIDIA GPU that supports CUDA to accelerate training, and the CPU is a multi-core processor. The specific parameter settings are shown in Table 2.

DOI: 10.54254/2755-2721/2025.26376

Table 2. The specific parameter settings

Aameter Type	Specific Parameters
Dataset split	70% training, 30% testing
Input/Output dims	Input dim determined by dataset, output dim 1
Position encoding	Max position encoding 512
Self-attention layers	2 layers, 4 heads, 32 key channels per head
GRU layers	6 forward units, 10 backward units
Other layers	ReLU activation, dropout rate 0.01
Optimizer	Adam
Training epochs	Max 200 epochs
Batch size	256
Learning rate	Initial 0.01, 0.001 after 80 epochs
Regularization	L2 coefficient 0.001
Gradient clipping	Threshold 10

In terms of comparative models, this article selects ExtraTrees, CatBoost, XGBoost, and LightGBM. In terms of model evaluation parameters, this article selects MSE, RMSE, MAE, MAPE, and R2.

Firstly, output the comparative results of ExtraTrees, CatBoost, XGBoost, LightGBM, and the proposed Transformer BiGRU on deep learning GPU computing power prediction tasks. The bar chart comparing the indicators of ExtraTrees, CatBoost, XGBoost, LightGBM, and Transformer BiGRU is shown in Figure 3, and the specific comparison results of the indicators are shown in Table 3.

Table 3. The specific comparison results of the indicators

Model	MSE	RMSE	MAE	MAPE	R^2
ExtraTrees	4.414	2.101	1.654	34.478	0.68
CatBoost	6.71	2.59	2.142	37.379	0.568
XGBoost	4.464	2.113	1.442	36.679	0.69
LightGBM	2.942	1.715	1.296	432.656	0.803
Transformer-BiGRU	2.208	1.486	0.969	26.938	0.855

DOI: 10.54254/2755-2721/2025.26376

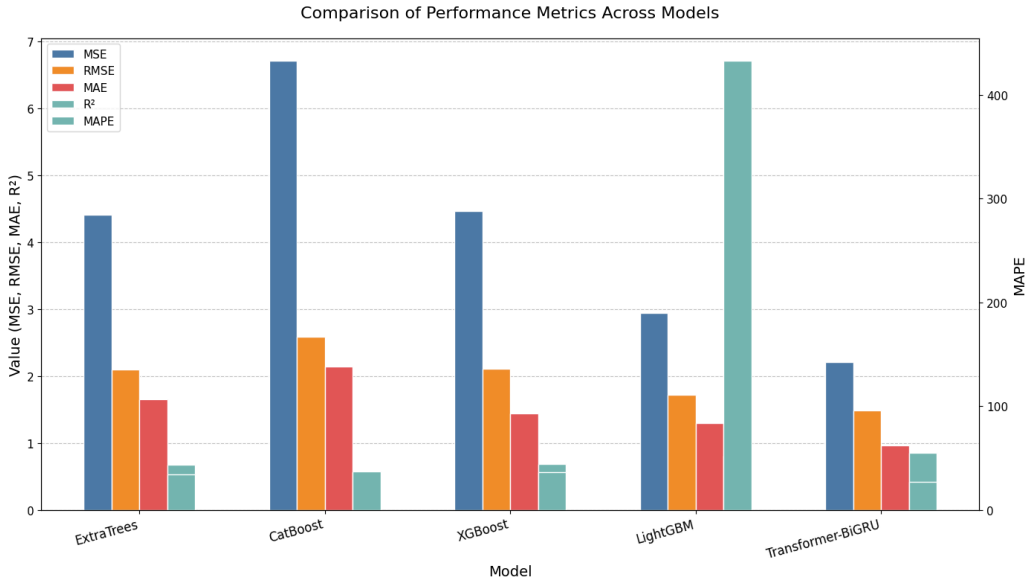

Figure 3. The bar chart comparing the indicators of ExtraTrees, CatBoost, XGBoost, LightGBM, and Transformer BiGRU

From various evaluation indicators, the Transformer BiGRU model proposed in this article performs the best in deep learning GPU resource regression prediction. Its MSE is 2.208, significantly lower than ExtraTrees' 4.414, CatBoost's 6.71, XGBoost's 4.464, and LightGBM's 2.942; The RMSE is 1.486, which is lower than the 2.101, 2.59, 2.113, and 1.715 of other models; MAE is at its lowest level at 0.969, which is better than XGBoost's 1.442 and LightGBM's 1.296; MAPE is 26.938, which is lower than ExtraTrees' 34.478, CatBoost's 37.379, and XGBoost's 36.679; The R ² reached 0.855, higher than LightGBM's 0.803 and XGBoost's 0.69, indicating that the model has smaller prediction errors for GPU resources and better fitting effects.

Output the predicted actual value curve of the Transformer BiGRU test set, as shown in Figure 4. From the predicted actual value curve, it can be seen that Transformer BiGRU's predicted values for deep learning GPU resources are very close to the actual values, which proves the predictive performance of the model.

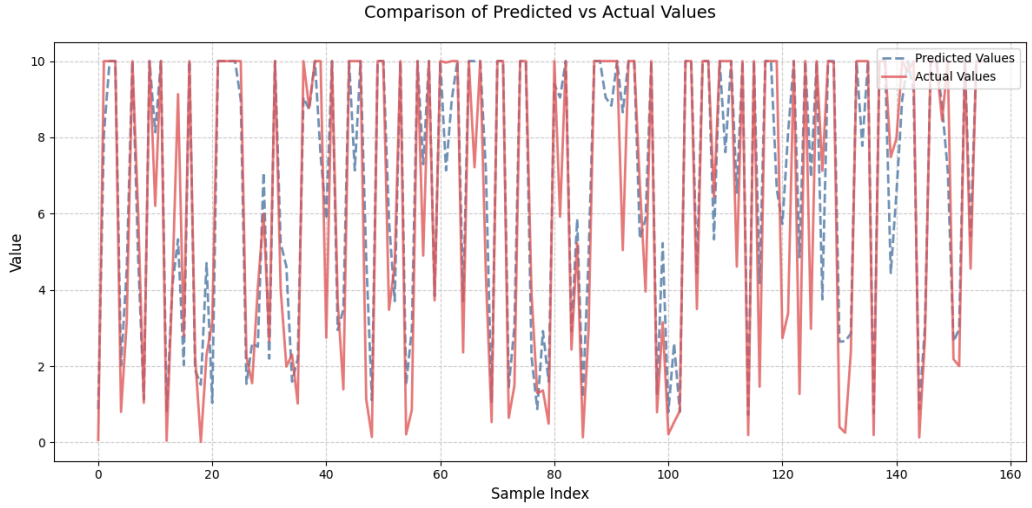

Figure 4. The predicted actual value curve of the Transformer BiGRU test set

DOI: 10.54254/2755-2721/2025.26376

5. Conclusion

The transformer-optimized BiGRU algorithm proposed in this paper lays a solid technical foundation for the accurate allocation of GPU computing resources in deep learning models. Through comparative experiments with ExtraTrees, CatBoost, XGBoost, and LightGBM models, it can be seen that the Transformer BiGRU model shows the best performance in the deep learning GPU resource regression prediction task. Specifically, it has a mean square error (MSE) of 2.208, which is significantly lower than ExtraTrees' 4.414, CatBoost6.71, XGBoost4.464, and LightGBM's 2.942; The root mean square error (RMSE) was 1.486, which was better than the 2.101, 2.59, 2.113, and 1.715 of the other models; The mean absolute error (MAE) was at its lowest level at 0.969, ahead of XGBost's 1.442 and LightGBM's 1.296; The mean absolute percentage error (MAPE) was 26.938, which was less than 34.478 for ExtraTrees, 37.379 for CatBoost, and 36.679 for XGBoost; The coefficient of determination (R^2) reaches 0.855, which is higher than 0.803 for LightGBM and 0.69 for XGBoost. These data fully prove that the prediction error of the model for GPU resources is small, and the fitting effect with the actual data is good.

The significance of this research result is that the Transformer BiGRU model not only provides an efficient solution for accurate prediction of GPU computing power resources in deep learning scenarios, but also lays a key foundation for dynamic scheduling and refined allocation of GPU resources by improving the accuracy of resource prediction. It helps to reduce the waste of computing power resources, improve the efficiency of resource utilization, and provide more solid computing power support for the large-scale application of deep learning technology in various industries.

References

[1] Li, M., et al. (2024). Deep learning and machine learning with GPGPU and CUDA: Unlocking the power of parallel computing. arXiv preprint arXiv:2410.05686.

[2] Wu, G., et al. (2015). GPGPU performance and power estimation using machine learning. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)* (pp. 1-12). IEEE.

[3] Lin, Y., et al. (2024). GPU-optimized image processing and generation based on deep learning and computer vision. *Journal of Artificial Intelligence General Science (JAIGS), 5*(1), 39-49.

[4] Dhanasekaran, S., et al. (2024). Utilizing cloud computing for distributed training of deep learning models. In *2024 Second International Conference on Data Science and Information System (ICDSIS)* (pp. 1-6). IEEE.

[5] Shank, D. B., et al. (2023). AI composer bias: Listeners like music less when they think it was composed by an AI. *Journal of Experimental Psychology: Applied, 29*(3), 676.

[6] Yang, T., & Nazir, S. (2022). A comprehensive overview of AI-enabled music classification and its influence in games. *Soft Computing, 26*(16), 7679–7693.

[7] Tchemeube, R. B., et al. (2025). Evaluating human-AI interaction via usability, user experience and acceptance measures for MMM-c: A creative AI system for music composition. *arXiv preprint arXiv:2504.14071*.

[8] Qiusi, M. (2022). Research on the improvement method of music education level under the background of AI technology. *Mobile Information Systems, 2022*(1), 7616619.

[9] El Ardeliya, V., Taylor, J., & Wolfson, J. (2024). Exploration of artificial intelligence in creative fields: Generative art, music, and design. *International Journal of Cyber and IT Service Management, 4*(1), 40–46.

[10] Zhao, H., et al. (2025). AI-driven music composition: Melody generation using Recurrent Neural Networks and Variational Autoencoders. *Alexandria Engineering Journal, 120*, 258–270.