

Predicting High-Performance Computing's Carbon Footprint: A Transformer-Based Analysis for Sustainable Development

Hao Wang^{1*†}, YiBo Dang^{2†}, ChenKai Zhang^{3†}, Richard^{4†}

¹*The Webb Schools, Claremont, USA*

²*Beijing 21st Century School, Beijing, China*

³*ZheJiang RuiAn High School, Wenzhou, China*

⁴*STaipei Wego Private Senior High School, Taipei, China*

**Corresponding Author. Email: Mike.0077hh@gmail.com*

†These authors contributed equally to this work and should be considered as co-first author.

Abstract: High Performance Computing (HPC) refers to the use of powerful computers and advanced computational techniques to solve complex problems and perform large-scale processing tasks at high speed. As HPC technology development advances, emission from HPC began to raise numerous environmental concerns, such as carbon emissions, global warming, resource consumption, etc. This study focuses on the environmental impact of the HPC industry. First, the environmental impact of HPC was evaluated. We considered the two indicators of HPC energy, Full Capacity and Average Utilization Rates, to understand and define environmental problems that HPC energy consumption can cause, such as E-Waste. Secondly, we used the time series prediction model with the Transformer architecture to predict and analyze the carbon emissions from HPC energy consumption. Among them, we first obtained the carbon emission data of different sources of HPC energy composition through data collection. Our main data comes from the Kaggle machine learning platform. Our data set mainly records the HPC energy emissions in the United States and different energy mixes. Then, after constructing the transformer model, we let the model predict HPC carbon emissions through data analysis. Finally, we calculated the global temperature increase caused by HPC carbon emissions and the subsequent environmental problems through the prediction results. We also used the Transformer model to predict the amount of carbon emissions in 2030 caused by HPC and the possible environmental impacts it might have, in order to help environmental protection organizations around the world take appropriate measures.

Keywords: HPC Carbon Emissions, Transformer Forecasting, Computational Sustainability, Energy Consumption

1. Background

The advent of High Performance Computing (HPC) marks the beginning of a new era of technological use. Its applications range from cancer treatment and genetic engineering, such as CRISPR, to cryptocurrency mining, such as Bitcoin, and to autonomous vehicle technology. Comprised of a multitude

of high-speed computer servers, it is able to divide the workload to several different servers, which allows it to complete rigorous tasks way faster. However, at the same time, its energy consumption, associated with carbon emissions, dramatically impacts the environment. The environmental impacts of HPC include, but are not limited to, excessive carbon emissions, water usage, electronic waste, damage to air quality, and resource depletion. As for carbon emissions, because of the high energy consumption of HPCS and the dependence of data centers on electricity, the main methods of producing electricity at present are burning fossil fuels or using nuclear reactions to drive heat engines. The demand for electricity by HPCs undoubtedly leads to the burning of more fossil fuels, which further leads to the release of too much greenhouse gas and accelerates the process of global warming. In addition, fossil fuels contain sulfur, which will produce toxic sulfur dioxide when burned, causing acid rain and destroying the ecological balance. In order to ensure the normal operation of HPC data centers, a large amount of water is used to cool them and keep them at a constant temperature, which puts a lot of pressure on some dry areas. For example, in Arizona, a dry region with limited water resources, Google's data center will use more than 1 billion gallons of water from the Missouri River or alluvial Missouri River in 2021. This has a negative effect on the local water supply and the environment. In addition, the used water will be discharged into the environment, and if the treatment of this water is neglected, chemicals and radioactive elements will enter the water sphere and pollute the water source. Even if some of the polluted wastewater seeps into the ground, it will cause hidden dangers for local residents to drink water. In the following research, we completed the assessment of the impact of HPC on the environment and the development trend and solutions in the future.

2. HPC environmental impacts

2.1. Impact overview

Before we can truly understand the impact of HPC carbon emissions, we need to understand two factors: Full Capacity and Utilization Rates. Full Capacity Definition: Full capacity is defined as the consumption of the HPC operating at the calculated load limit, where all components of the HPC are working at maximum efficiency.

Average Utilization Rates: The average utilization rate is the average consumption generated during HPC operation, reflecting the consumption situation more accurately since HPC does not maintain full capacity at all times.

Afterwards, we collected reliable data through official channels and analyzed environmental issues by calculating the carbon emissions caused by energy.

2.2. Data collection and variables

Before building the model, we need to collect reliable data for calculation and analysis. In the process of collecting data, we mainly collected the following variables in Table 1.

Table 1: Definitions of variables used in the calculations

Variable	Description	Unit
E	Total energy	MWh
P	Full load power	MW
U	Average utilization rate	Range (0 to 1)
T	Total number of hours per year	Hours

The full load power of the HPC system, "Frontier," located at Oak Ridge National Laboratory in the United States [1], is 21.1 MW according to Top500.org. The average utilization rate of HPC

systems typically ranges between 30% and 60%, as reported by IDC. The total number of hours in a year is 8,760.

3. Modeling

3.1. Model overview

Machine learning mainly uses a large amount of data to help us make predictions and decisions. In this study, the main task of this problem is to predict the environmental impact of HPC energy consumption and the future impacts [2]. Therefore, this study decided to mainly use machine learning and deep learning methods for prediction, which are well suited for prediction tasks.

In this study, the prediction model we mainly used is Transformer, a large deep learning model invented by Google researchers in 2017. The Transformer model was originally designed for natural language processing tasks, but because it can handle long-term dependencies and capture complex temporal patterns, it has shown significant results in time series prediction [3]. Compared with traditional models (such as ARIMA or LSTM), Transformer does not need to explicitly decompose time series or handle long sequence dependencies, and can efficiently process sequences of any length by relying on the self-attention mechanism [4].

Our team used the Transformer model to analyze a large number of data sets, and then gave the corresponding predictions for HPC carbon emissions and carbon emissions from now to 2030. The Transformer architecture has shown good performance in time series data like the ones presented, and our training is shown to be effective, with a low loss at the end. In our future predictions, we have explored more future laws of carbon emissions through Transformer, which will be discussed later. We have also added the impact of the variable "renewable energy" on carbon emissions, as well as deeper environmental impacts. Through the completion of these tasks, Transformer has demonstrated extremely strong processing capabilities.

3.2. Transformer model for HPC environmental impacts

3.2.1. Data and variables

In order to objectively determine the environmental impact of HPC energy consumption on overall carbon emissions, this study first collected data on carbon emissions in different regions of the United States. In the process of collecting data, we found reliable data on the Kaggle machine learning platform. This dataset fully meets the requirements of our model and includes the total carbon emissions generated by different energy sources in different regions from 1970 to 2018 [5]. Data on carbon dioxide emissions for several U.S. states starting in 1970 are included in this dataset. In Table 2, the information is categorized by fuel type (coal, petroleum, natural gas, and all fuels combined), state, and sector (residential, commercial, transportation, electric power, and industrial). Millions of metric tons of carbon dioxide are used to measure the emissions.

Table 2: Features of the dataset

Index	Feature	Description
1	year	The year for which the emissions data is provided (e.g., 1970)
2	state-name	The name of the U.S. state (e.g., Alabama, Alaska, Arizona)
3	sector-name	The sector for which the emissions data is provided
4	fuel-name	The type of fuel contributing to the carbon dioxide emissions
5	value	The carbon dioxide emissions value in million metric tons for the specified year, state, sector, and fuel type.

The following Table 3 is the original data of the dataset, which shows some sample data from the dataset.

Table 3: Sample data from the collected dataset

Index	year	state-name	sector-name	fuel-name	value
1	1970	Alabama	Industrial carbon dioxide emissions	Coal	26.72151
2	1988	Montana	Transportation carbon dioxide emissions	Natural Gas	0.12184
3	2004	Nevada	Electric Power carbon dioxide emissions	Coal	18.047908
4	2008	Pennsylvania	Commercial carbon dioxide emissions	Petroleum	3.218243
5	2017	Alaska	Industrial carbon dioxide emissions	Natural Gas	14.143605

In order to make the Transformer better trained, the original data set of this study was processed to make the model training more efficient. This study designed training and test data suitable for the model according to the requirements of the topic, with the following features In the Table 4, we presented features of the expected dataset. The feature "HPC Emissions" of this dataset is the target variable, and we need to build a model to successfully predict HPC emissions by analyzing emissions from different energy sources and encode "year" as a time series feature, or use positional encoding to make the model aware of the time order. Through the computer's built-in data set calculation tool, we finally succeeded in obtaining the ideal training and test data. Here are some samples in Table 5.

Table 4: Features of model training dataset and testing dataset

Index	Feature
1	year
2	Coal Emissions Global
3	Natural Gas Emissions Global
4	Petroleum Emissions Global
5	Total Emissions
6	HPC Emissions

Table 5: Sample data from the final training data and testing data

Index	year	Coal Emissions Global	Natural Gas Emissions Global	Petroleum Emissions Global	Total Emissions	HPC Emissions
1	1970	329.553989	321.6906864	561.8159956	1213.060671	0.076731953
2	1971	313.1833886	332.9252404	584.6388226	1230.747452	0.085310586
3	1972	314.6460623	325.1408023	605.6805361	1245.467401	0.094848309
4	1973	346.8467869	333.1847816	658.5826759	1338.614244	0.10545235
5	1974	328.6327001	310.9118498	612.0763559	1251.620906	0.117241923

3.2.2. Model architecture for generating outputs

The Transformer model is based on the Encoder-Decoder Architecture, but in many practical applications, such as regression tasks or sequence prediction tasks, only the encoder part is usually used. The role of the encoder is to map the input sequence to a latent space representation, and then generate the target sequence through the decoder. In order to apply it to the carbon emission prediction problem, we will only focus on the encoder part [3].




Figure 1: An overall structure of transformer model

The self-attention mechanism allows each element in the input sequence to perform a weighted sum on other elements in the sequence, allowing the model to capture long-range dependencies. The core formula of self-attention is as follows: Given an input sequence

$$X = \begin{bmatrix} \text{Coal Emissions_Global}_1 & \text{NaturalGas Emissions_Global}_1 & \dots & \text{Total Emissions}_1 \\ \text{Coal Emissions_Global}_2 & \text{NaturalGas Emissions_Global}_2 & \dots & \text{Total Emissions}_2 \\ \vdots & \vdots & \ddots & \vdots \\ \text{Coal Emissions_Global}_T & \text{NaturalGas Emissions_Global}_T & \dots & \text{Total Emissions}_T \end{bmatrix}$$

where T is the number of time steps. Here we will let x_i represent each element. $x_i \in R^d$, it is mapped into Query, Key, and Value vectors as follows:

$$Q_i = x_i W^Q, K_i = x_i W^K, V_i = x_i W^V$$

where $W^Q \in R^{d \times d_k}$, $W^K \in R^{d \times d_k}$, and $W^V \in R^{d \times d_v}$ are learnable weight matrices.

The attention score between query Q_i and key K_j is calculated using the scaled dot product:

$$\text{Attention_Score}(Q_i, K_j) = \frac{Q_i K_j^T}{\sqrt{d_k}}$$

The attention weights are obtained by applying a softmax function over the scores:

$$\alpha_{ij} = \frac{\exp\left(\frac{Q_i j^K}{\sqrt{d_k}}\right)}{\sum_{k=1}^n \exp\left(\frac{Q_i k^K}{\sqrt{d_k}}\right)}$$

where α_{ij} represents the attention weight of x_j with respect to x_i .

The output of the attention mechanism for Q_i is computed as:

$$Attention_Output(Q_i, K, V) = \sum_{j=1}^n \alpha_{ij} V_j$$

To enhance the model's capacity, multiple attention heads are computed in parallel. Each head is defined as:

$$head_h = Attention(h_Q^{QW}, h_K^{KW}, h_V^{VW})$$

where h represents the index of the head, and h_Q^W, h_K^W , and h_V^W are the weight matrices for the h -th head.

The outputs of all heads are concatenated and linearly transformed:

$$MultiHead(Q, K, V) = Concat(head_1, head_2, \dots, head_H)W^O$$

where W^O is the weight matrix for output projection.

After each layer of self-attention, we have a feed-forward neural network to further process the features of each time step. Each time step x_i is processed independently by the feed-forward network, which usually consists of two fully connected layers and an activation function.

The feed-forward network applies a two-layer transformation with a ReLU activation:

$$FFN(x) = \max(0, xW_1 + b_1)W_2 + b_2$$

where $W_1 \in R^{d \times d_{ff}}, W_2 \in R^{d_{ff} \times d}$, and b_1, b_2 are bias vectors. d_{ff} is the hidden dimension of the feed-forward network.

The Transformer model does not have a built-in mechanism for handling sequential information, so positional encoding is needed to inject the positional relationship of each time step. A common practice is to use sine and cosine functions to generate positional encodings.

To introduce positional information into the model, we use sine and cosine functions:

$$PE_{(pos, 2i)} = \sin\left(\frac{pos}{10000^{2i/d}}\right), PE_{(pos, 2i+1)} = \cos\left(\frac{pos}{10000^{2i/d}}\right)$$

where pos is the position index, i is the dimension index, and d is the input dimension.

In order to avoid gradient disappearance and speed up the training process, the Transformer model uses residual connection and layer normalization.

Each sub-layer in the Transformer is wrapped with a residual connection and layer normalization:

$$Layer_Output = LayerNorm(x + Sublayer(x))$$

Layer normalization is computed as:

$$LayerNorm(x) = \frac{x - \mu}{\sigma} \cdot \gamma + \beta$$

where μ is the mean of x , σ is the standard deviation, and γ and β are learnable parameters.

For carbon emission prediction, the mean squared error (MSE) loss is typically used:

$$L = \frac{1}{N} \sum_{i=1}^N (y_{true,i} - y_{pred,i})^2$$

where N is the number of samples, $y_{true,i}$ is the ground truth value, and $y_{pred,i}$ is the predicted value.

Once trained, the model can make predictions based on carbon emissions data. Given a series of input features (such as coal emissions, natural gas emissions, oil emissions, etc.), the model will output the corresponding HPC emissions.

4. Results and discussion

Before formally training the model, we need to preprocess the data. In the preprocessing process, we use Python's powerful built-in Pandas and Numpy libraries to help us. After importing the data, we first classified the data set, where "HPC_Emissions" is the target variable and the rest are prediction features. In addition, since each feature is a numerical type and the gap between each value is too large, in order to facilitate model training, we converted all numerical data into values in the range of $[0,1]$. The technology used is normalization.

After that, we used the Pytorch framework to obtain a dataset suitable for Transformer model training. Pytorch provides the Dataset class, which allows us to customize the way we access and manage data. Therefore, we can customize how to load data, and use DataLoader to load data in batches and automatically shuffle the order, and easily adjust the data acquisition logic, such as adding data enhancement and handling missing values [6].

We then used the Pytorch framework to write the Transformer model code, and then trained and evaluated the model, obtaining the following results [7].




Figure 2: Actual results VS. predicted results of HPC emissions

From the table, we can clearly see that the training of the model is quite effective, and the difference between the model's prediction results and the actual results is minimal, confirming the effectiveness of the Transformer architecture in this task.

Once the model successfully predicts HPC carbon emissions, we can find out the environmental impacts corresponding to different values based on the world's carbon emissions values. The main impact of carbon emissions is an increase in global temperature, so we need to know how much temperature increase is caused by each million metric tons (the unit of the dataset) of carbon emissions.

To estimate the temperature increase per million metric tons (Mt) of carbon dioxide emissions, we rely on the Transient Climate Response to Cumulative Emissions (TCRE), which quantifies the relationship between cumulative carbon dioxide emissions and global temperature rise. TCRE typically ranges from 0.8 degrees Celsius to 2.5 degrees Celsius per 1,000 Gt of carbon dioxide emissions. Here is the calculation process [8]:

$$\Delta T = \frac{TCRE}{1000} \quad (1)$$

$$\Delta T = \frac{0.8}{1000} - \frac{2.5}{1000} = 0.0008 - 0.0025 \quad (2)$$

So after calculating the carbon emissions per million metric tons, we can plug in the model's predicted data and find the temperature increase caused by HPC carbon emissions each year like in Table 6.

Table 6: Sample intervals of the temperature resulted from HPC emissions

Index	Year	HPC Emissions	Temperature Increased Interval
1	1970	0.14344907470158885	[0.00011476,0.00035862]
2	1971	8.958732	[0.007167,0.02239]
3	1972	10.605084197731285	[0.008484,0.026513]
4	1973	7.189039210726513	[0.005751,0.017973]
5	1974	0.36944066346092996	[0.00029555253,0.00092360166]

By calculating the temperature increase caused by HPC in all years of the dataset, we concluded that the temperature increase caused by HPC as a whole is in the range of [0.0001,0.5] . Therefore, according to the International Energy Agency and the Environmental Protection Organization, we can summarize the following environmental impacts caused by HPC [9]:

- Coral Bleaching: At a rise of 0.5°C , coral reefs begin to experience more frequent bleaching events, particularly in tropical regions. While this is less dramatic than a 1°C increase, it still stresses coral ecosystems, which are highly sensitive to temperature changes
- A 0.5°C temperature rise accelerates the melting of glaciers and polar ice sheets. While the contribution to sea level rise might still be small in this range, it marks the beginning of significant long-term changes .
- Even a modest increase in global temperature increases the frequency and severity of heatwaves. In the range of 0°C to 0.5°C , some regions may start seeing more frequent and longer heatwaves, especially in parts of Europe, North America, and Asia .
- A 0.5°C increase can already result in higher incidences of heat-related illnesses, particularly in urban areas. Vulnerable populations, such as the elderly and low-income groups, are at increased risk.

The 0°C to 0.5°C increase may seem small, but it marks a significant the effects of climate change begin to manifest more visibly, especially in vulnerable regions [10]. As the world approaches these levels, the cumulative effects will intensify, potentially leading to irreversible changes in ecosystems, human systems, and the climate itself.

Considering the complex trend shown in the data and the small number of training data, the model chosen for predicting carbon emission till 2030 is a transformer. The transformer is constructed with a dropout rate of 0.1 to prevent overfitting, and the Adam optimizer with a learning rate of 0.001 for training. The dataset is split with 80% of it used for training. The emission data is also scaled for

better performance. The model is trained over 50 epochs. Afterwards, the model is predicting with a sequence length of 5 and re-scaling using MinMaxScaler.

Below is a graph to showcase the predicted carbon emission from 2019 to 2030.




Figure 3: Global fuel emission estimations

As demonstrated above, we can see that future carbon emission is first met with an increase followed by reduction. According to data for the International Energy Agency (IEA), energy usage from HPC accounts for 1-1.5% of the global energy usage. We can calculate the growth of the HPC industry with Compound Annual Growth Rate (CAGR). The CAGR can be calculated with the following formula:

$$CAGR = \left(\frac{FV}{PV} \right)^{\frac{1}{n}} - 1$$

where:

- FV is the final value,
- PV is the present value,
- n is the number of periods.

With HPC having an CAGR of 11.18% [9], we can calculate the HPC usage in 2030 with the following formula:

$$U_{2030} = U_0 \times (1 + r)^t$$

where:

- U_{2030} is the projected HPC usage in 2030,
- U_0 is the current HPC usage,
- r is the CAGR (0.1118),
- t is the number of years from the current year to 2030.

We can see that in 2030 the energy usage from HPC would account for 2.6% of the global electricity, and by extension of the carbon emission. Analysing the trend from the transformer model, we can also get the source energy mix of the HPC industry, shown below:




Figure 4: Global fuel emission makeup

From the Pie Chart above, we can clearly see that Petroleum and Coal account for the largest share of energy consumption, so we believe that global factories should find ways to reduce the production of oil and coal raw materials in the future. To verify this suggestion, we did a simulation to observe the new 2030 HPC energy consumption and carbon emission forecast by reducing the production of oil and coal raw materials.




Figure 5: The new prediction of 2030 emissions

From the Figure 5, we can see that this suggestion is valid because by reducing the amount of oil and coal produced, we can effectively reduce the carbon emissions of global HPC production in the future.

5. Conclusion

5.1. Strengths of the model

- **Comprehensiveness:** This modeling analyzes the impact of HPC on the environment from multiple dimensions, including carbon emissions, resource depletion and electronic waste caused
- by HPC. The data used are sourced from official agencies, such as IEA and IDC, to ensure the authority of the data

- Accuracy of Prediction: The prediction accuracy is high, and the deviation between the predicted value and the actual value in the test stage is low. And it can accurately quantify the environmental impact of different energy structures in the future
- Actionable Insights: Specific indicators are provided. The amount of carbon emission reduction in the renewable energy scenario. This helps shape the government's policy choices, provides clear guidance, and estimates the trade-offs between mitigation policies in different circumstances.

5.2. Weaknesses of the model

- Emissions Restrictions: HPC energy consumption sometimes not only leads to carbon dioxide emissions, but also may lead to the emission of compounds such as methane, nitrous oxide, etc. However, due to data limitations, this study only considered the carbon dioxide emissions caused by HPC consumption, and chose to ignore other variables. Therefore, this is also a point that our model needs to improve.
- Ignoring Uncertainty Factors: The model simply relies on objective data sets to predict carbon emissions, without considering some uncertain factors, such as the impact of economic, military and other events on future emissions. This problem also leads to some inaccurate predictions of the model for the future.
- Experiment Environment Limits: For the training of the deep learning large model Transformer, we do not have a large dataset to help the model training, and we do not use a GPU that is more convenient for model training. Due to the limitations of our model training environment, the model did not achieve a higher performance training, which will lead to some relatively inaccurate results.
- Data Limits: Large deep learning models, especially Transformer, often require a lot of data for training. However, since most official data is inaccessible, the volume of training data is limited. This is often a disadvantage for training large models and may lead to inaccurate models, etc.

5.3. Overall conclusion

This modeling successfully quantifies the impact of HPC on energy consumption and the environment and provides valuable analysis and projections for the future. It has the advantage of comprehensive analysis, high precision and flexibility. This modeling quantifies the carbon emissions of HPC and scientifically illustrates the impact on the environment. Modeling was trained using authoritative historical data from the authorities and combined with a variety of scenarios to predict the future carbon emissions and environmental impacts of the HPC in 2030 with high accuracy, and to reveal the changes in the future environment caused by different energy mix and government choices. Through model prediction and data quantification, the ecological risks of HPC are presented and there are still areas to be improved, which provides effective help and theoretical support for the further development of the international community in this field. In summary, this modeling successfully predicted the impact of HPC on the environment and proposed a scientifically based and accurate solution.

References

- [1] NREL. (n.d.). "High-Performance Computing Data Center Power Usage Effectiveness." <https://www.nrel.gov/computational-science/measuring-efficiency-pue.html>.
- [2] ScienceDaily. (2021, September 21). "Emissions from Computing and ICT Could Be Worse Than Previously Thought." <https://www.sciencedaily.com/releases/2021/09/210910121715.htm>.
- [3] Vaswani, A. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.
- [4] Kassai, K., Dagiuklas, T., Bashir, S., & Iqbal, M. (2024, June). "GreenBytes: Intelligent Energy Estimation for Edge-Cloud." In *IET Conference Proceedings CP879*, Vol. 2024, No. 4, pp. 57-61. Stevenage, UK: The Institution of Engineering and Technology.

- [5] Kaggle. (2024, April). "CO2 Emissions (USA)." [https://www.kaggle.com/datasets/ abdelrahman16/co2-emissions-usa?resource=download](https://www.kaggle.com/datasets/abdelrahman16/co2-emissions-usa?resource=download).
- [6] (2024). "Chasing Carbon: The Elusive Environmental Footprint of Computing." <https://ar5iv.labs.arxiv.org/html/2011.02839>.
- [7] The MIT Press Reader. (2022, February 22). "The Staggering Ecological Impacts of Computation and the Cloud." <https://thereader.mitpress.mit.edu/the-staggering-ecological-impacts-of-computation-and-the-cloud/>.
- [8] International Energy Agency. (2024). "Electricity 2024." <https://www.iea.org/reports/electricity-2024>. Licence: C C BY 4.0.
- [9] Chinnici, A., Ahmadzada, E., Kor, A., De Chiara, D., Domínguez-Díaz, A., De Marcos Ortega, L., & Chinnici, M. (2024). "Towards Sustainability and Energy Efficiency Using Data Analytics for HPC Data Center." *Electronics*, 13(17), 3542.
- [10] IEEE. (2024, April). "An Efficient Energy Consumption Prediction Framework for High Performance Computing Cluster Jobs." In *IEEE Conference Publication, IEEE Xplore*. <https://ieeexplore.ieee.org/document/10603495>.