

Energy Efficiency in Biological and Artificial Intelligence: A Comparative Analysis

Jinwen Zhu

*Suzhouwan International Academy, Suzhou, China
3899190712@qq.com*

Abstract. The rapid development of artificial intelligence has brought unprecedented computational power, but it also consumes a large amount of energy. Training a single large language model can consume electricity which is equal to the annual usage of hundreds of households. In contrast with its 86 billion neurons, the human brain operates on merely 20 watts of power—roughly equivalent to a light bulb. Biological intelligence achieves such extraordinary cognitive feats with so little energy, while artificial intelligence requires megawatts. This article presents a systematic comparison of energy consumption between biological and artificial intelligence systems. The comparison is conducted across three dimensions: energy magnitude, energy structure, and the energy-performance relationship. The analysis shows that biological intelligence outperforms artificial systems six orders of magnitude in energy efficiency. The underlying reasons include the brain's sparse coding mechanisms, analog signal processing, event-driven computation, and billions of years of evolutionary optimization. Conversely, artificial intelligence systems are built upon dense floating-point matrix multiplications and backpropagation algorithms. However, the underlying von Neumann architecture introduces a critical bottleneck—frequent data movement between memory and processors consumes a significant portion of the total energy. The article concludes by discussing potential bio-inspired approaches. Talk about neuromorphic computing and spiking neural networks, as promising directions for reducing the energy footprint of future artificial intelligence systems.

Keywords: Biological Intelligence, Artificial Intelligence, Energy Efficiency, Neural Networks, Neuromorphic Computing

1. Introduction

1.1. Background of study

Artificial intelligence has undergone remarkable evolution over the past decade. From image recognition systems to large language models, AI technologies have permeated nearly every aspect of modern life. However, energy consumption during this progress cannot be ignored. According to a widely cited study, training a single GPT-3 model consumed approximately 1,300 megawatt-hours of electricity, which is equal to the annual energy consumption of approximately 130 average U.S.

households [1]. As AI models continue to scale up in both parameter count and training data volume. The energy demand is also growing at an alarming rate.

AI's energy appetite is hard to ignore. But biology offers a stark counterexample. The human brain performs a staggering array of cognitive functions with an energy cost of only about 20 watts, which is less than the power consumed by a laptop computer. The human brain's ability to recognize a face, understand a conversation, or navigate a crowded street requires merely a small fraction of the energy needed for AI system. This dramatic disparity between biological and artificial intelligence in terms of energy efficiency has profound practical implications. As the global community prioritizes sustainability, the growing energy footprint of AI has posed a serious challenge. Data centers powering AI operations occupy approximately 1% of global electricity consumption, and this figure is projected to rise sharply [2]. Therefore it is important to understand the mechanisms behind the difference

1.2. Problem statement

Despite the analogy drawn between human brain and artificial intelligence, these two systems differ in their energy consumption profiles. The human brain and deep neural network may both be described as "neural networks," but the physical and computational mechanisms of they process information are radically different. These differences manifest most conspicuously in their energy requirements.

A notable gap persists in the literature. There is a systematic comparison of energy. Neuroscientists have examined the metabolic costs of neural activity. AI researchers have also recorded the energy consumption of model training. However, they are not in coordination. Therefore, an all-encompassing study will be conducted. Explicitly state where, why, and to what extent biological intelligence is more energy-efficient. In light of the above deficiency, the current study will focus entirely on the energy consumption dimension. Systematically compare it with the above several sub-dimensions. It will attempt to provide an academic basis for the energy gap and its causes.

1.3. Research objectives and significance

The primary objective of this study is to provide a systematic comparative analysis of energy consumption. Specifically, this research aims to:

1. Compare the absolute energy magnitudes of biological brains and AI systems;
2. Compare the structural distribution of energy consumption in both systems;
3. Compare how energy consumption relates to performance in each system;
4. Identify the fundamental reasons for the energy efficiency gap;
5. Discuss the implications of these findings for future AI development.

6. The significance of this study has two aspects. Theoretically, it contributes to our understanding of the nature of intelligence. It declares that "intelligence" is not a monolithic concept. Instead, it is a phenomenon that can be instantiated in different substrates (biological versus silicon) with fundamentally different energy costs. Practically, the insights from this comparison may promote a more efficient design of AI system.

2. Literature review

2.1. Energy consumption in biological intelligence

The human brain is an energy-efficient organ. It merely has 2% of total body mass. However, it consumes about one-fifth of the body's resting metabolic energy [3]. This energy is supplied in the form of glucose and oxygen through the cerebral blood flow. At the cellular level, energy consumption can be traced to several distinct processes.

Most of the brain's energy is used to maintain the stability of the resting membrane potential and generate action potentials (electrical signals for intercellular communication in neurons). Sodium-potassium pumps are used in this way to transport ions across the neuron's membrane; at the same time, the concentration differences of sodium and potassium are maintained. It is not energetically favourable. Research shows that after neural signals; the repair of these ion gradients consumes about 50%-60% of cortical energy [3]. When a neuron fires, sodium ions rush into the cell; to restore the resting state afterwards requires a considerable amount of ATP. Another reason for high energy consumption in organisms is that they are very active. The brain uses a principle of "sparse coding", so only a small number of neurons are active at any given time [4]. The Difference from artificial systems starts at the signal level. Brains do not perform high-precision floating-point arithmetic; rather, they are continuous systems of graded potentials and variable-rate spikes. Therefore, energy saving will not be realised. As most of the brain's energy is only used when neurons fire, sparse-activity patterns have been selected to reduce metabolic costs by lowering the number of spikes.

Biological brains are not isolated computational units; this is also true. They are in living organisms that have been subject to billions of years of evolution. Natural Selection has optimised Neural Systems for survival. At the same time, energy-saving devices have also been selected in recent years. A being that wastes energy on redundant neural processing will be at a disadvantage in evolution due to energy scarcity [5]. Natural Selection has shaped the brain effectively. Not through intelligent design, but by continuously cutting away what is wasteful.

2.2. Energy consumption of artificial intelligence

Recently, the amount of power used by artificial intelligence systems has drawn the attention of many people. AI energy consumption generally falls into two categories in non-biological systems: training and inference [6].

The training stage is relatively high-energy consumption. Train a neural network with millions or billions of parameters on a large amount of data. Backpropagation algorithms then iteratively adjust the weights to reduce the prediction error. The calculation is a large matrix multiplication, and thus the scale of the energy cost is relatively high. A 2019 study found that training a single NLP model produced carbon emissions equivalent to those of five cars over their lifetimes [7]. Since then, more recent models have expanded significantly and are now in the billions of parameters. GPT-3 used about 1,300 MWh of electricity and emitted approximately 550 tons of CO₂ equivalent [8].

The inference stage, when a trained model is used to predict new data, is relatively low-power but still requires some energy. Each query of the large language model is approximately several million floating-point operations for response generation. When these models are scaled up to serve millions of users daily, the total inference energy may even exceed the training energy for the life cycle of the model [9].

The following are the reasons for the high-energy consumption of AI systems. First, the current dominant paradigm of deep learning is scaling; that is, both model performance and computational

resources are increased by adding more parameters and data. Therefore, there is a motivation to build ever-larger models and, consequently, higher energy costs. Second, the hardware used for training is generally a GPU or TPU that has a high power density and is often several hundred watts per chip. Third, AI training runs on distributed systems that require high-bandwidth data transfer between chips, and data movement is itself a considerable energy consumer (a phenomenon known as the "von Neumann bottleneck", where moving data between memory and computation units consumes more energy than the computation itself).

2.3. Comparison framework for this study

Based on the literature review above, this paper will use a three-dimensional framework to compare the energy consumption of biological and artificial systems. The above indices are chosen because they can reflect the primary features of energy consumption in both areas more intuitively.

Dimension 1: Energy Scale. The first is the quantity of power used by biological brains (in W) and AI systems (also in W or watt-hours for total usage). It is a relatively simple comparison, but due to the large difference in scale, it is also very representative.

Dimension 2: Energy Structure. An index shows how the energy in all parts of a system is distributed. Energy is needed to power the ion pumps and other components in the cell that require life support. Energy in artificial systems is used for computation (matrix multiplication), data movement (between memory and processors), cooling/overhead, etc. Given the different structures, the amounts of energy are not the same.

Dimension 3: Energy-Performance Relationship. Compare the extent of function realisation for the two systems at different energy-input levels. In other words, it is asking what kind of cognitive or computational performance can be achieved for a given amount of energy? The other end of the coin is to consider what results have been achieved by the system in light of energy consumption.

3. Comparative analysis

3.1. Comparison of energy magnitudes

At present, the first difference between biological and artificial intelligence is a considerable disparity in energy consumption. The human brain is about 20 watts of continuous power [10]. 20 watts is about the amount of power required to run a small fluorescent lamp as an example. Although the energy budget is small, many functions of the human brain need to be realised, such as cognition, language comprehension, visual recognition, motor coordination, social cognition, planning and self-awareness.

On the other hand, the power consumption of today's AI systems is measured in megawatt-hours for training runs and hundreds of watts per high-end GPU during operation, not in watts. A single training run of high-end AI can consume several million watt-hours and is about 10^6 times more energy-intensive than the continuous operation of the human brain. Even for a single inference (e.g., generating a response from a large language model), the energy cost per query is several joules, and the brain performs complex reasoning tasks at a fraction of this amount per unit of time [11].

A typical comparison of the energy consumption by different AI models and the human brain is shown in Table 1.

Table 1. Comparison of energy magnitudes [1, 12, 13]

System	Estimated Energy Consumption	Approximate Scale
Human Brain (continuous operation)	20 W	Reference baseline
Inference (single ChatGPT query)	~2-10 J	Per query
Training BERT-base	~750 MWh	Equivalent to 60 US homes/year
Training GPT-3	~1,300 MWh	Equivalent to 130 US homes/year
Training GPT-4 (estimated)	~50,000 MWh	Equivalent to 5,000 US homes/year

As shown in the table, biological and artificial systems are quite different; moreover, a problem has arisen: with the increase in size of AI models, their energy consumption is rising at an accelerating pace and widening the gap further.

3.2. Comparison of energy structure

Although the scales are very different, to learn why the two systems are so far apart, one must also know where and how the energy is used.

Energy Structure in Biological Systems. Alterations in neural activity at different places and times are closely related to changes in brain energy consumption. The top ones are:

Ion gradient recovery (approx. 50-60%): Sodium-potassium pumps continuously restore the ion gradient after an action potential. It is an ongoing base cost of the brain when it is not working.

Synaptic transmission (about 20-30%): Releasing neurotransmitters in the synaptic cleft and recycling are energy-intensive processes.

Background maintenance (around 10%-20%): To support a low-level metabolism, the basic cell functions of protein synthesis and transport need to be provided.

Energy consumption in biological neurons is activity-dependent and event-driven; that is, only when neurons fire is energy used. As brain activity is very sparse (only a small fraction of neurons are active at any time), it has a good energy-to-computation ratio [8]. Processing of biological information is also continuous; signal strength is not divided into high-precision digital numbers. Thus, the problem of high-energy consumption in representing and operating large floating-point numbers is avoided.

Energy Structure in Artificial Systems. Energy consumption of AI systems is also composed of these parts:

Matrix multiplication (about 60-70%): Most AI calculations are reduced to multiply-accumulate operations. These operations need high-precision floating-point calculations and are energy-intensive. A contemporary GPU can execute a few trillion such operations per second, but each one consumes electricity.

Data Movement (about 20-30%): As shown in the von Neumann architecture, the transfer of data between memory and the processing unit is relatively energy-intensive compared to the arithmetic operations themselves. Deep learning has billions of parameters that need to be read many times during training, and moving data is relatively energy-intensive.

Cooling and overhead (around 5-10%): Data centres for AI hardware need extensive cooling facilities and are relatively energy-intensive.

AI computations are digital and dense, not biological systems. All layers in a deep network compute activations for all neurons in the batch, even those that will not contribute significantly to the result. Most of the existing deep learning models do not have a "sparse firing" mechanism; they are fully connected.

3.3. Comparison of energy-performance relationship

Perhaps the first index to be used for comparison is the energy required per unit of output. Performance is not a single index in this area, but rather the whole of one's intelligence and proficiency.

Biological Systems operate at an advantageous energy-performance curve. A 20-watt human brain can support a general intelligence that is adaptable to new circumstances, learns from a small number of examples, integrates various senses, and conducts abstract thinking. The brain is not as good at some specific tasks as a calculator for arithmetic, but it is more flexible. Researchers' concept of "artificial general intelligence" (AGI) refers to the energy budget of the brain, and it is a type of intelligence that has not been fully understood but has shown strong robustness and adaptability in almost all areas of life [6].

On the other hand, artificial systems have high performance in a narrow area but require a large amount of energy. Deep-learning models have performed better than humans in some areas, such as image recognition (ImageNet) and games (AlphaGo), but this superior performance is generally restricted in scope. To achieve this performance, the energy consumption would be several orders of magnitude higher than that required by the human brain for the same task (if such a task were possible for humans).

A typical case of the above is that biological intelligence has wide-ranging abilities but uses little energy, whereas artificial intelligence, although highly capable, is relatively energy-intensive. This indicates that the two have different optimisation objectives: evolution aims for survival efficiency, and AI engineering focuses on task-specific accuracy and considers energy optimisation last.

3.4. Summary of comparison results

As shown in the three-dimensional comparison, all the sub-dimensions are in the same direction: biological intelligence is much more energy-efficient than artificial intelligence, and the two differ by about six orders of magnitude. Such high efficiency is achieved by sparse, analogue, event-driven and evolutionarily optimised computation. Artificial systems, on the other hand, use dense, digital, exhaustive, and hardware-constrained computation and are less energy-efficient. The following table 2 summarizes the main findings.

Table 2. Summary of comparative findings

Dimension	Biological Intelligence	Artificial Intelligence
Energy Magnitude	~20 W	10^6 – 10^7 W (training)
Energy Structure	Sparse, analog, event-driven	Dense, digital, exhaustive
Energy-Performance	Broad competence, low energy	Narrow superhuman, high energy

4. Discussion

4.1. Reasons for the energy efficiency of biological intelligence

The extraordinary energy efficiency of biological intelligence can be attributed to several underlying mechanisms, each of which reflects a deep optimization driven by evolutionary pressures.

Sparse Coding and Selective Activation. Biological neural networks do not activate all neurons simultaneously. The brain employs sparse representations in which only a small subset of neurons are active at any given moment [11]. This is not merely a passive property but an active strategy: inhibition is just as important as excitation in the brain. By keeping most neurons at rest (inhibited), the brain minimizes baseline energy consumption while preserving the ability to respond rapidly to salient stimuli.

Analog Computation. Biological neurons communicate by transmitting graded potentials and spike trains, and are therefore in an analog mode. Neural signals in the brain are noisy and continuous, so they are not high-precision numbers requiring 16, 32 or 64 bits of accuracy for digital systems; rather, this imprecision is a characteristic of the signal. The brain's analog nature eliminates the energy cost of representing, storing, and manipulating discrete high-precision values [10]. Therefore, many cognitive tasks may only require low-precision or even stochastic computation; thus, research into low-precision neural network quantization has been spurred by this result.

Event-Driven Processing. Biological neurons fire only when the summed input signal exceeds a threshold. This is a basic event-driven mode; power is only used when a particular event or internal condition occurs. A universal clock (as in AI systems) is not needed to coordinate all operations, and thus the power consumption is lower. However, AI systems have global and synchronous processing, and every neuron in a layer computes outputs for all inputs in every forward pass, even if some inputs are not informative.

Evolutionary Optimisation. Evolutionary selection may be responsible for making life more efficient. Over billions of years, organisms that wasted energy on superfluous neural activity were selected against. Every energy-saving mutation that maintained or enhanced cognition offered a survival advantage. The result is a system that is energy-efficient and, at the same time, meets the demand for life. This does not mean that the brain is ideally suited for all kinds of problems; rather, it is prone to biases, heuristics and errors, but it is highly energy-efficient in its operation environment.

4.2. Why is artificial intelligence so energy-intensive

On the other hand, some attributes of contemporary artificial intelligence are energy-intensive, and many of these stem from the present deep-learning model.

Backpropagation and Dense Gradient Computation. Almost all deep learning uses the backpropagation algorithm, and it needs to compute gradients for every parameter in the network during each training cycle. Propagate the error signal backward through all the layers and carry out dense matrix operations. On the other hand, biological learning mechanisms (such as Hebbian plasticity) are local and do not require global gradient propagation; they are relatively energy-efficient but also less powerful for some tasks [9].

High-Precision Floating-Point Arithmetic. Most AI training uses 32-bit or even 16-bit floating-point numbers. Although reduced precision (e.g., 8-bit quantization) has been studied, most of the original training still uses high precision for stability. The requirement for high precision will

increase energy consumption; every floating-point operation needs to switch charges in the electronics, and more bits are needed to achieve higher precision.

The Scaling Paradigm. One reason for the high energy consumption of AI is that these models often have to be large. Therefore, a "bigger is better" strategy has been adopted, and more extensive models and datasets are built with greater computational power for training. Although the above method has achieved good results, it has also led to an increase in energy consumption that is growing at an accelerating rate.

Von Neumann Bottleneck. Computation and Memory are physically separated in the traditional hardware architecture. Data needs to be continuously transferred between the two, and this movement of data consumes a large amount of energy. At the same time, the human brain combines memory and computation at the synapse to avoid the energy cost of long-distance information transfer. These Buildings have different foundations.

4.3. Insights from biology for artificial intelligence

Given that biological systems are relatively high in energy, neurobiology is now being applied by AI researchers to construct more energy-efficient artificial systems. Several good ways are being studied at present.

Neuromorphic Computing. IBM's TrueNorth and Intel's Loihi are examples of neuromorphic chips that aim to emulate the structure of biological neural networks in silicon. These chips are event-driven, sparsely activated and use low-precision signals, all inspired by the brain's energy-saving mode. First, some results have been obtained showing that neuromorphic hardware is relatively energy-efficient to GPUs for certain types of neural network calculations [3].

Spiking Neural Networks (SNNs). SNNs are artificial neural networks that resemble biological neurons more closely by using discrete spikes (binary events) to transmit information, rather than continuous values. Spikes are binary and event-driven; therefore, SNNs are very energy-efficient when realised on neuromorphic hardware. Training SNNs is also a problem, and they still underperform in complex tasks compared with traditional deep networks.

Model Compression and Quantization. Pruning (removing unnecessary parameters), Quantisation (reducing numerical precision), and Knowledge Distillation (training a smaller model to emulate a larger one) are all methods to reduce the size of a model without changing the hardware. These methods have not yet fully connected with the biological system but are feasible steps for more sustainable AI.

Algorithmic Innovations. Finally, the first two will be changes to the general algorithm. If researchers can develop learning algorithms that do not require global backpropagation, or those that learn from fewer examples (few-shot learning), or those that utilize innate priors (as the brain does), then the energy consumption of AI could be significantly reduced. These research directions are timely and promising at present.

However, it should be pointed out that full reproduction of the biological efficiency in silicon is not easily realised. After billions of years of optimisation by the brain, "optimisation" in an evolutionary sense refers to a series of trade-offs. The brain is not ideal for high-precision arithmetic or large-scale parallel matrix multiplication because it is not in a dynamic, noisy, and uncertain environment. To reproduce the high efficiency, one must know how the components work at the basic level.

5. Conclusion

A systematical comparison of the energy consumption of biological intelligence and artificial intelligence has been conducted in this paper. The three levels of analysis are the scale of energy, its structure and energy-performance characteristics. The first set of results are shown below:

First, the amounts of energy consumption are about six orders of magnitude different. A human brain has a power consumption of 20 W, but AI training consumes megawatt-hours of electricity. They are not the same in amount; their natures of existence are different.

Second, the Structure of Energy Consumption is different. Biological systems have a low-energy allocation and are event-driven or analogue. Artificial systems are energy-intensive and have high overhead for digital floating-point operations and data movement. Sparsity and locality in the brain are responsible for its efficiency; dense computation and memory bottlenecks in AI lead to inefficiency.

Third, energy consumption is not directly related to the level of performance. Biological intelligence is broad and general, thus requiring little energy; artificial intelligence is narrow and superior but is very energy-intensive. Optimisation goals vary; optimisation aims for fitness, but tasks demand accuracy.

Fourth, the basic reasons for the above differences are evolutionary optimisation, sparse coding, analogue computation, event-driven processing and architectural differences between biological and silicon systems. Although AI has shown significant progress, it is still far less energy-efficient than life.

Some deficiencies of this study have been identified. First, the data come from different places and use different measurement methods, so they are not directly comparable. Energy estimates for AI training vary considerably in the research due to different hardware, data centre efficiency and measurement scope. Estimates of brain energy consumption also need to assume metabolic rates and areas of activity.

Second, the comparison is not reciprocal. The human brain is a living organ that cannot be "turned off" or standardised in the same way as an AI system. They are therefore not precisely comparable.

Third, this study has focused only on energy consumption and has not considered other types of comparison, such as speed, accuracy, flexibility or robustness. Expand the scope of the comparison and collect more data.

Fourth, this is a literature-only study without new empirical data. Experimental Design for measuring Energy Consumption in Controlled Operations will be employed in future work.

Based on the above results, future research will take several forms.

First, extend the research on neuromorphic computing and spiking neural networks. The above methods are expected to help close the energy gap in the future, but they still have some technical problems, such as how to build good training algorithms and design large-scale hardware.

Second, some studies have not examined the structure of the brain but have focused on how it operates at a basic level instead. For example, one can explore how the brain learns without global backpropagation and how it uses prior knowledge to learn from a few samples to inspire completely new types of AI.

Thirdly, at the level of policy, the energy efficiency of AI systems will be given more attention in research and development. Benchmarking energy efficiency and accuracy simultaneously to set new reporting standards for energy consumption and promote the development of energy-efficient algorithmic innovation through incentives.

Fourthly, there are more extensive problems of sustainability and climate change due to AI energy consumption. With the spread of AI in critical infrastructure, its total energy consumption will be relatively high. Development of "green AI", that is to say, AI that is both powerful and environmentally friendly, is a pressing problem.

In conclusion, the comparison between biological and artificial intelligence underscores a profound irony: the technology that aims to replicate the brain's intelligence still has much to learn from the brain's efficiency. By understanding the principles that make biological intelligence so energy-efficient, we may be able to build artificial systems that are not only smarter but also more sustainable.

References

- [1] Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10), 1133-1145. <https://doi.org/10.1097/00004647-200110000-00001>
- [2] Bissett, K., et al. (2019). Energy-efficient inference in deep neural networks. *IEEE Micro*, 39(4), 20-29. <https://doi.org/10.1109/MM.2019.2920524>
- [3] Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., Liao, Y., Lin, C. K., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., ... & Wang, H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1), 82-99.
- [4] Jones, N. (2018). How to stop data centres from gobbling up the world's electricity. *Nature*, 561(7722), 163-166.
- [5] Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., & Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12, 518-527. <https://doi.org/10.1038/s41558-022-01377-7>
- [6] Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253.
- [7] Laughlin, S. B., de Ruyter van Steveninck, R. R., & Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, 1(1), 36-41. <https://doi.org/10.1038/236>
- [8] Lennie, P. (2003). The cost of cortical computation. *Current Biology*, 13(6), 493-497.
- [9] Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6), 335-346. <https://doi.org/10.1038/s41583-020-0277-3>
- [10] Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10), 1629-1636.
- [11] Olshausen, B. A., & Field, D. J. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14(4), 481-487. <https://doi.org/10.1016/j.conb.2004.07.007>
- [12] Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training. arXiv preprint, arXiv:2104.10350. <https://arxiv.org/abs/2104.10350>
- [13] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645-3650. <https://doi.org/10.18653/v1/P19-1355>