

Energy-Efficient and Sustainable Artificial Intelligence Models..

Dr. Shailesh M. Hambarde¹, Dr. Poonam Lambhate², Dr. Aparna Shailesh Hambarde³, Dr. Minakshi N. Vharkate⁴, Dr. Maroti S. Kalbande⁵, Sayali Ashok Shivarkar⁶

¹ Associate Professor, Department of Electronics and Telecommunication Engineering, JSPM's Jayawantrao Sawant College of Engineering, Savitribai Phule Pune University, Pune, Maharashtra, India.
(shaileshmhambarde@gmail.com)
ORCID: 0009-0000-0656-8607

² Associate Professor, Department of Computer Engineering, MIT Academy of Engineering, Pune, Maharashtra, India.
(poonam.lambhate@mitaoe.ac.in)
ORCID: 0009-0000-3998-7429

³ Associate Professor, Department of Computer Engineering, K. J. College of Engineering and Management Research, Pune, Maharashtra, India.
(aparnahambarde07@gmail.com)
ORCID: 0009-0003-9475-227X

⁴ Associate Professor, Department of Computer Engineering, MIT Academy of Engineering, Pune, Maharashtra, India.
(mnvharkate@mitaoe.ac.in)
ORCID: 0000-0003-4859-1762

⁵ Assistant Professor, JSPM Hadapsar, Pune, Savitribai Phule Pune University, Maharashtra, India.
(maruti.kalbande@gmail.com)
ORCID: 0009-0004-2369-5375

⁶ Assistant Professor, Department of Computer Engineering, Keystone College of Engineering, Savitribai Phule Pune University, Maharashtra, India.
(sayali.shivarkar@keystonesoe.in)
ORCID: 0009-0008-3463-5123

Abstract: Computing demand, energy consumption, and associated environmental problems have been accelerated by artificial intelligence (AI) technologies, which have seen rapid progress and widespread adoption. Despite the significant benefits that AI-driven systems provide to sectors such as healthcare, finance, transportation, and government, there are valid concerns about their long-term sustainability and viability caused by their increasing carbon footprint. This paper examines the concept of sustainable and energy-efficient AI models with a focus on methods that reduce computer complexity, optimize resource utilization, and minimize environmental costs without compromising performance. Methods like as green data centers, energy-aware algorithm design, federated and edge learning, model compression, and AI infrastructure integration of renewable energy sources are important and explored. Sustainable AI, the study continues, is an essential part of aligning technological development with global environmental goals, moral duty, and legal mandates. The article argues that artificial.

intelligence models should emphasize efficiency, scalability, and ecological balance in order to ensure responsible digital transformation and sustainable development in the era of intelligent systems

Keywords: Energy-efficient AI, Sustainable artificial intelligence, green computing, Model optimization, Carbon footprint, Ethical AI.

Introduction

The use of AI to aid in decision-making is essential to every human effort. As part of the process, optimization algorithms and prediction models are developed via trend analysis of large datasets. The use of AI has skyrocketed in the last ten years. With more than 35 billion connected devices, the Internet of Things enables the integration of devices into almost every industry, from healthcare to consumer goods to smart cities. They all provide massive amounts of data that power AI methods to boost activity efficiency. Contrary to what end users may think, AI development has been driven by massive computer capacity, regardless of advances in AI research. As a result, more advanced models and prediction systems may be developed, which has several potential uses in achieving the goals of the European Green Deal. Developing AI might potentially harm the environment, which could undermine these goals.

Concerning findings have emerged from a number of assessments of extensive AI use in recent years: Data centers that operate massive AI software account for 5-9% of the world's electricity and 2% of its carbon dioxide emissions; the use of energy for AI has increased 300,000-fold from 2012 to the present. Think tanks predict that information and communication technology will be responsible for 14% of global warming emissions by the year 2040. This increase in energy consumption is attributable, in part, to the creation of new AI-specific designs aimed at improving the performance of AI models.

But recently, many in the AI field have started looking for ways to reduce their emissions after realizing that energy efficiency is an important metric to consider while developing AI methods. Data centers, data transmission, and edge computing systems will reduce emissions using AI systems that are energy efficient. A carbon-responsive computer is a way to build computer systems that use less energy.

Environmental and other concerns need energy efficiency in the development of AI systems at the edge. The democratization of artificial intelligence is happening now, with data stored in huge computer infrastructures yet used all around the world. A lot of ancillary fields use AI. To offload tasks that demand a lot of energy, you may link to generic computer infrastructures and employ them. This application paradigm is driven by a number of factors, including the scarcity of computing resources, the energy requirements of IoT and embedded devices, the high energy consumption of AI algorithms during execution, and the ever-increasing cost of energy. Internet of Things (IoT) and embedded devices would be able to run these models and algorithms with less power consumption if AI systems used less power.

In The Quest for Energy-Efficient Ai Systems

Energy efficiency was the primary motivation for the 2008 coining and subsequent promotion of the term "green computing," which is most commonly linked with sustainable IT. However, when we narrow our attention to systems, the objectives for studying new algorithms were time and space complexity, while energy complexity was never considered. Research on carbon-responsive computing has been too focused on data center applications because to their fast proliferation, but this is not enough. A greater amount of processing power, storage, networking, control, and data management is now located outside of huge data centers thanks to the exponential growth of grid, cloud, web, pervasive, ubiquitous, autonomous, fog, and edge computing during the last twenty years.

Having said that, the low-level design of hardware components is essential to all of the aforementioned computer infrastructures. A significant area of research in integrated circuit design and manufacturing is the reduction of power and energy consumption in digital circuits. Research conducted by academic institutions, IC design firms, and IC foundries has implications for central processing units (CPUs), graphics processing units (GPUs), field-programmable gate arrays (FPGAs),

and application-specific integrated circuits (ASICs). Over the last two decades, many of publications have been uncovered by bibliographic search engines regarding IC usage.

Power consumption, or energy per unit of time, and instantaneous or cumulative energy consumption are two important measures for evaluating circuit use. The energy and environmental impact of the circuit may be better evaluated with the help of the first metric. The circuitry of devices that run on batteries determine how long the batteries last. Power consumption, on the other hand, impacts the thermal viability of circuits and the size of power supplies.

Methods to decrease the circuit's static consumption, supply voltage, and activity rate have been extensively studied in order to decrease power consumption in a central processing unit (CPU), application circuit (EC), memory, etc. The area of artificial intelligence (AI) that we are primarily concerned with is the construction of higher-level hardware and software.

Researchers use CPU-GPUs, cloud, and edge computing infrastructures to speed up AI algorithms. We explore these things. In terms of AI systems, they are all equally qualified. The picture of energy-efficient systems may be shown in Fig. 1.

CPU-GPU based systems

Performance and power consumption balance are two of the problems that arise in CPU-GPU systems that work together. Optimizing a system's cross-cutting processors lowers its power consumption and increases its usefulness.

Many multi-core systems rely on task scheduling, and there are several methods in the literature that aim to optimize energy use by focusing on this duty. To reduce power consumption during task scheduling, the authors of a paper outlining a CPU-GPU architecture offered a GPU power and performance model based on DVFS (Dynamic Voltage and Frequency Scaling). Each task's optimal GPU voltage and frequency is determined using mathematics. A system that combines CPUs and GPUs can handle a comparable problem. This approach utilizes heuristic-based job scheduling in conjunction with hybrid particle swarm optimization. Goals include evaluating and reducing energy usage during work schedule.

proposed a universal CPU-GPU work allocation technique. With the goal of enhancing performance while decreasing the energy-delay product.



Figure. 1: Brief overview of energy-efficient systems.

A approach was created by Henderson et al. to analyze the energy and carbon emissions of ML algorithms in real-time. They used four baseline approaches and two RL benchmarks in their use case. Emissions of carbon dioxide were determined for each experimental location. Energy consumption is measured using Nvidia-smi and Running Average Power Limit. You may find out how much power your CPU, DRAM, and GPU are using with Power Usage Effectiveness. The authors propose a multi-objective evaluation strategy for generative models, with a focus on their ability to generate raw waveforms, in. The approach incorporates energy consumption and model quality. For iterative ML workloads, the authors presented a live GPU energy optimization method that can automatically determine the best energy configuration.

Cloud computing

Finding the best approach to use the cloud while reducing energy use is a common goal of study. Because of the high expense of cooling systems for large cloud environments, all-encompassing methods are required for managing cooling and other factors that impact the cost and energy consumption of computing. The majority of solutions made use of meta-heuristics for these methods. Conceived and implemented an all-encompassing resource management system for cooling, servers, networks, and storage in order to reduce carbon emissions and energy consumption in cloud data centers. Based on cooling considerations and the data center's physical location, a thermal model may estimate server temperatures and associated cooling costs.

Edge computing

With edge computing, data, devices, sensors, and end users are all brought closer to the cloud. To combine grid, cloud, and Internet of Things technologies, it employs Cloudlets, Fog Computing, and Mobile Edge Computing. An important area of study is edge computing, which is necessary because there will soon be hundreds of millions of IoT and 5G devices close to the end user. These devices will include portable ones with great processing power and memory capacity, which could be utilized to deploy AI-based applications on a massive scale. While there is a lot of research on energy efficiency in cloud settings, there isn't nearly as much that has been done to identify major challenges in edge systems due to the different components. When implementing AI solutions, keep in mind that embedded systems or edge computing devices lack the capacity of large data centers. Typically, tasks are distributed hierarchically. Bringing low-energy tasks to the edge and moving those with high energy requirements to the cloud. Since the majority of data is generated on edge devices, consumption of data sent from the edge to the cloud demonstrates that hierarchical task allocation is inefficient.

Network modeling enhances energy efficiency while investigating wireless communications and mobile computing on the edge. Energy constraints in mobile and pervasive computing, as well as environmental sustainability, have been studied by certain researchers, but their potential impact on future advances in artificial intelligence has received less attention. Biomedical, image recognition, and machine translation are some of the mobile AI applications. Companies have developed Neural Processing Units and APIs for mobile devices since most of them employ Neural Network models (Android NN). The challenge is in making efficient use of them in terms of energy usage. Nevertheless, autonomous vehicles, spacecraft, and mobile devices might all make use of AI models developed for specific devices. In certain cases, mobile devices don't measure up to specialized gear. The Field-Programmable Gate Array (FPGA) technology allows AI models to run on these devices. It is possible to use FPGAs to build circuits for highly efficient NN models. Because of its built-in integer arithmetic capability, it can train and infer with smaller representations, and its customization options allow users to trim, share weights, or change the network's topology, all of which contribute to better energy efficiency. Hardware design methods like as clock gating, power domains, deep pipelines, and dynamic voltage and frequency scaling may further enhance them.

A Full-Stack Approach to Energy-Efficient Artificial Intelligence

A lot of effort has gone into making AI systems more efficient, as shown before. Research

balancing energy costs with system performance is few. Sustainable artificial intelligence systems are a must for the European Green Deal. To achieve a happy medium between energy efficiency and the accuracy of predictions and optimizations (responses), we think AI systems can and should be built. From selecting the target architecture, arithmetic precision, compilation optimizations, tuning AI algorithms to the energy consumption of the AI-specific architecture, or optimizing hyperparameters, programming language, or energy-aware algorithms during model execution, all of these steps in the AI system design flow should optimize energy. The four steps of evaluating, analyzing, and optimizing energy use are shown in Figure 2. At every stage, we suggest using the industry-standard tools, measurements, and procedures for power and energy estimation. Due to the fact that different users and roles are involved at each level, we analyze these roles and how they relate to the levels in the parts that follow.

From algorithms to user-friendly AI packages

The way that AI users use when pursuing energy efficiency is determined by their responsibilities and aims. To illustrate the difference, consider the following: solving a problem using a publicly accessible AI package and actually creating it. There may be some disagreement among AI researchers as to what constitutes an efficient model or algorithm. Efficient operation is essential when assessing power use. Consequently, we detail the available tools and the potential outcomes for each position below. Figure 2 shows that we employ a full-stack approach to tailor our recommendations and energy use viewpoint to every customer. The AI practitioner is our first point of departure; they need not be AI experts but rather tool users. Next, we'll take a look at the responsibilities of AI software developers, and lastly, the tasks of AI researchers.

Energy efficient measures of interest at different levels of the full-stack approach

Many AI software tools and packages are open-source, so keep that in mind while assessing the levels in Figure 2. So, anybody looking to solve optimization or machine learning problems may choose one, encode it in the tool they've chosen based on the approaches, run it, and see what happens. The technique, parameter settings, and available time to run the algorithm all have a role in the final solution quality, and these factors might differ across different hardware platforms. Since outside firms have developed the algorithms and software, the consumer has little control over their energy usage. Despite this, parameter values can still be changed (which, as shown earlier, might impact the power consumption required to find solutions) and the algorithm can still be executed on hardware platforms when there are many available.

First and foremost, we recommend that anybody working with AI research and utilize software that displays an algorithm's power consumption. It is recommended to check whether this data is available on different hardware platforms since computer architectures might vary. Because of this need from consumers, software engineers will include energy metrics and other standard outputs in their products. Secondly, we suggest looking at how the parameters are defined and how it impacts the final solution, training energy consumption, time to solution, and result quality. Below, we provide recommendations for the ideal proposal and talk about the connection between the studied levels and the involved users: reduction of AI power consumption

AI algorithms design / AI researchers: A problem-solving algorithm is the goal of the AI researcher. The investigator is in a hurry and needs high-quality responses. Nonetheless, reducing energy consumption must now be a part of algorithm design. It may take a long time for optimization and machine learning (ML) approaches to resolve real-world problems. For decades, everyone has been focusing on improving the quality of the results while minimizing computability issues (such as avoiding exponential costs and accepting polynomial time whenever practical) and time complexity. A new strategy is required to zero in on the amount of energy required to achieve a certain quality, supposing that there is a direct correlation between algorithm time complexity and energy usage.

Software packages and algorithms parameterization / AI package developers and users: AI package designers can encode algorithms. Few algorithms are created or improved by this guy. They aim to incorporate a prominent AI algorithm in software. Our future goals will include energy usage reduction. Data from the run is used to show final users how much power was required to identify

solutions.

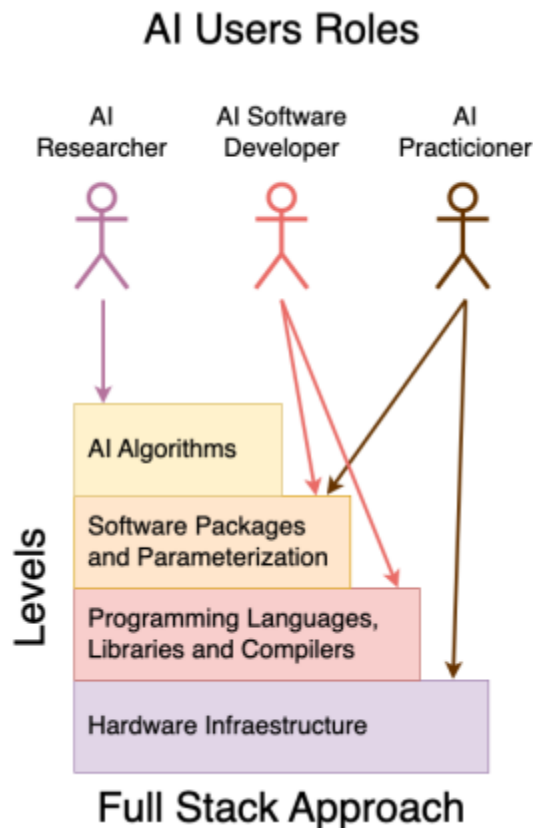


Figure 2: There are four levels in which energy consumption evaluation, analysis, and optimization should be considered and the AI user roles identified

OBJECTIVES

To create a sustainable AI corpus with an energy emphasis and carry out effective preprocessing and post-processing by combining LDA, BERT, and clustering algorithms for precise topic identification and interpretation.

To illustrate important trends in the field of energy research, seven identified clusters will be mapped into logical subjects (cluster 0 as topic 1 through cluster 6 as topic 7). This will provide a worldwide thematic picture of energy-efficient and sustainable AI research.

RESEARCH METHODOLOGY

Researchers generally believe that both quantitative and qualitative research techniques have intrinsic strengths and shortcomings. Therefore, it is recommended to combine both approaches to guarantee that the findings complement each other. In our study, we included and drew from four complementary sets of research approaches. Three of these—BERT, LDA topic modeling, and clustering—are associated with text mining techniques. In addition to these numerical data, we also incorporated a qualitative content analysis based on topics. Our mixed-methods approach stands out for three reasons. Additional approaches to evaluating qualitative articles include narrative reviews, critical reviews, and mapping reviews (Paré and Kitsiou, 2017). According to Saheb et al. (2021), our topic-based content analysis method looked at the best publications on each subject to fill in the gaps in knowledge. We used computational methods such as BERT, LDA, and clustering to collect research on sustainable AI in energy and determine their subject content. Topic modeling may make use of a number of techniques, such as LDA, LSA, and clustering using the k-means algorithm. Modern methods combine language models with LDA and LSA algorithms. Compared to other methods, language models perform better because they consider not just the meaning of individual words and

phrases but also the connections between them. Compared to other methods, language models perform better because they consider not just the meaning of individual words and phrases but also the connections between them. Consequently, the results were better and the silhouette score was higher after utilizing this method. In addition, we were able to attain the maximum degree of semantic integrity within each subject utilizing this method, which made the subjects more relevant and unique.

Additionally, a qualitative approach known as content analysis was used to conduct a comprehensive assessment of the recovered topics. Third, this study integrated the LDA and BERT topic modeling approaches to get the greatest level of accuracy in topic identification. To better understand quantitative and computational studies, our suggested mixed-method approach may be used by researchers in many domains to conduct topic-based content analysis.

Xie et al. (2020) states that LDA is predicated on the premise that words are more often found in some themes relative to others, and that texts are made up of subjects. Although LDA has been used often by academics to determine topics, it has several limitations due to the fact that it is a word co-occurrence analysis and cannot consider the whole content of the sentence. In addition, short texts are not well-represented (Qiang et al., 2017). On top of that, LDA findings could be hard for individuals to grasp and implement (Song et al., 2009). Conversely, BERT topic modeling seeks to uncover semantic similarity by merging topics with previously taught contextual representations (Peinelt et al., 2020). It greatly enhances the consistency of neural topic models by including contextual information into the topic modeling process (Bianchi et al., 2020). The Autoencoder method, an encoder that scans the input text, is part of the Transformer library that BERT employs. In this study, we combined the LDA and BERT vectors to enhance topic identification and categorization. Because dealing with huge dimensions is a particularly difficult part of word-sentence embedding, we also used the Uniform Manifold Approximation and Projection (UMAP) approach. When compared to alternative methodologies for manifold learning, UMAP is among the most successful (McInnes et al., 2018).

We encountered various computational challenges in our work. The originality of the BERT algorithm was the first major challenge. To streamline our one-of-a-kind method, we made use of the Hugging Face transformer library, which facilitates the creation of thousands of pre-trained models for various modalities such as text, vision, and audio. Simultaneously, there were challenges with pre-processing texts, such as the identification of stopwords. It was finished with the help of a specialist. Next, we needed to figure out how many themes would work best. The elbow approach yields the optimal number of clusters (k) by fitting the model with a range of k.

Corpus Building

The terms "artificial intelligence" OR "AI" AND "sustainable" OR "sustainability" AND "energy" were searched on May 29, 2021, inside the title, keywords, and abstract. The Scopus database provided the data. In order to concentrate just on AI, we eliminated data science, algorithms, machine learning, and deep learning from our study.

As a result of this search, 981 documents were found. After that, we limited the language to English and the content type to articles. There were 296 articles as a consequence of this exclusion. After that, we carefully assessed the article titles and abstracts to choose the most relevant ones that looked at AI's role in guaranteeing the sustainability of the energy industry. 182 articles from 2004 to 2022 were found during this scan. Since research article abstracts provide the most concise overview of important concepts (Delgosha et al., 2021), we included the final publications' abstracts into the study's corpus.

Topic Modeling

To carry out the subject modeling, we used the following libraries: The dataset is read using Pandas, LDA is performed using Gensim, BERT is performed by Transformers, auto-encoding is done by Keras, and the results are shown using Seaborn and Matplotlib. We imported KMeans from the Scikit-learn cluster and the TFID vectorizer from the Scikit-learn feature extraction. While BERT was used to generate the sentence embedding vector, LDA was used to construct the probabilistic topic assignment vector. First, we modeled the themes using the TF-IDF, LDA, and BERT (Figure 3). The information content of each vector was then balanced by concatenating the LDA and BERT vectors. To develop a lower-dimensional latent space representation for the concatenated vector, we processed the

auto-encoder using the Keras package. We computed the Silhouette Score, which was 0.566 and close to one for LDA+BERT+ Clustering, to make sure the clusters were of high quality. Figure 4 shows that TFIDF+clustering scored 0.048 and BERT+clustering scored 0.095. Cluster quality is measured using the Silhouette Score (Shahapure and Nicholas, 2020). SSE, Davies-Bouldin, and the Dunn index may alternatively be employed as substitutes for the silhouette score. Using silhouette was more suitable since our database had outliers. Since silhouette can detect and locate outliers, the score it provides makes more sense. As a result, silhouette offers a more accurate representation of the general quality and coherence of each subject. The range of the score is -1 to 1. The cluster is thick and well-separated from nearby clusters if the score is close to one. The clusters are of exceptional quality because LDA BERT Clustering is closer to 1 than other topic modeling algorithms.

RESULT

Figure 3 shows the integrated topic modeling approach used in this research, which underpins the findings section. The graphic shows probabilistic and contextual representations merged for topic discovery. The Latent Dirichlet Allocation (LDA) component on the left creates an interpretable topic distribution from global, corpus-level thematic structures based on word co-occurrence patterns. In parallel, the BERT-based representation in the bottom center preserves textual data's contextual semantics beyond frequency information. Concatenation and encoding combine statistical topic distributions and contextual embeddings as a unifying layer. The unified representation is clustered to create well-defined topic clusters on the right side of the picture. Contextual topic extraction renames each cluster as a cohesive subject. The hybrid LDA–BERT framework increases topic coherence, interpretability, and semantic richness compared to single topic modeling techniques, as seen in this picture.




Figure 3: LDA/BERT vector encoding for contextual theme extraction

Figure 5 shows the final topic identification that was produced using LDA+BERT+Clustering Algorithms. Dimension reductions and an eight-topic setting were implemented using the UMAP package. Another option besides UMAP is Principal Component Analysis (PCA), which is similar to TSNE. When compared to the other two, UMAP excels in handling sparse datasets, processing massive data more quickly, and managing large data sets generalized. Various topic clustering methods, such as 10, 4, and 6, were also tested.

Distant Insights Sustainable AI in Energy

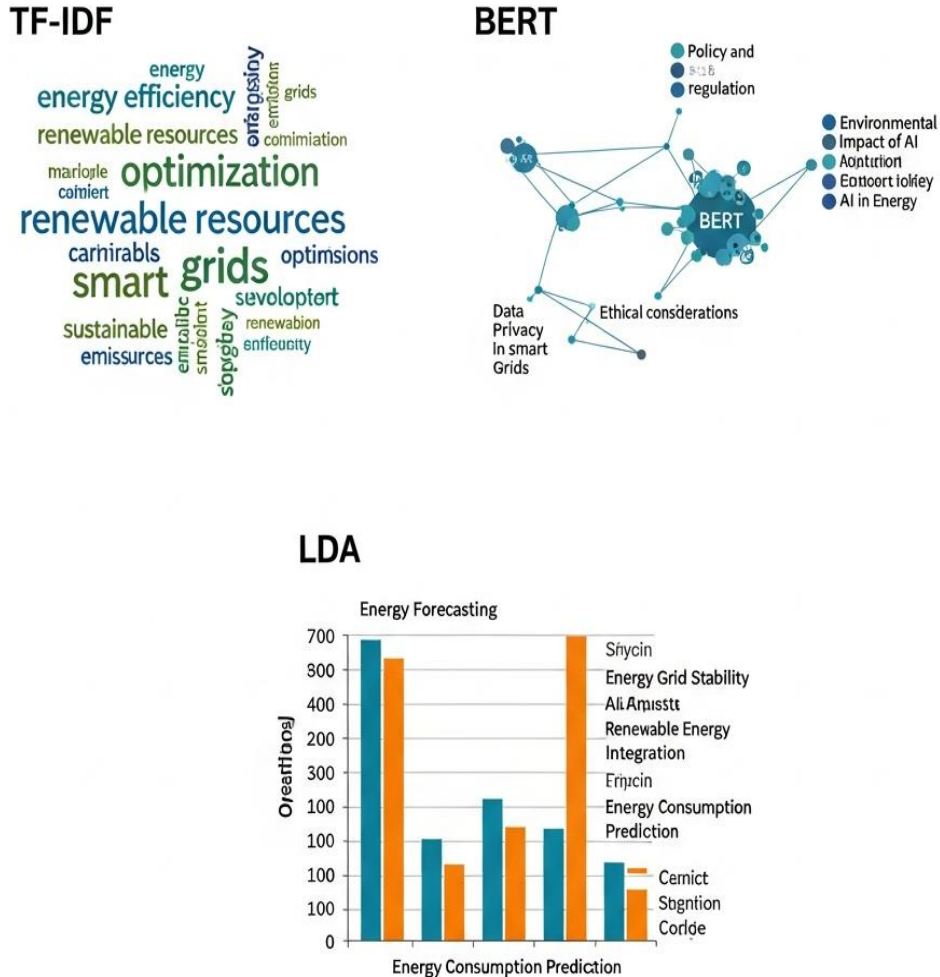


Figure 4: Topic modeling of sustainable AI in energy research using TF-IDF, BERT, and LDA algorithms gives independent results.

Our initial test included eleven clusters, although several overlapped. Four or six broad-scope groups were formed, and each cluster included so many themes that precise grouping was practically impossible. We chose eight groups since they were neither too wide nor too narrow. Eight themes were better separated and had higher density within each subject, demonstrating superior clustering. About 11% of texts are in subject 0 and 16% in topic 1. Clustering's great quality was confirmed by a balanced document distribution within each subject.

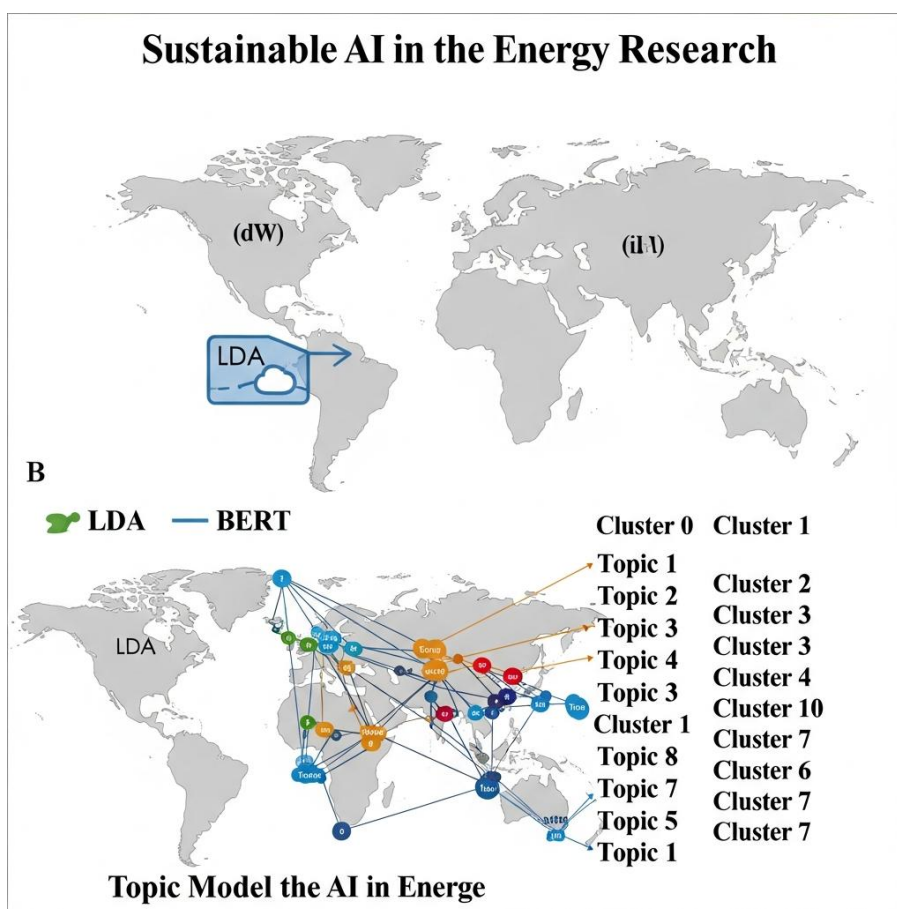


Figure 5: Sustainable AI in energy research: a global perspective. We used LDA, BERT, and clustering to identify topic modeling. We changed clusters to topics and began with topic 1, thus cluster 0 is subject 1, cluster 1 is topic 2. This covers all 7 clusters.

Topic modeling is used to examine and organize sustainable AI research in the energy domain to promote energy-efficient and sustainable AI models.

The top map shows the global distribution of sustainable AI in energy studies study by area. The graphic shows how LDA and BERT-based models are used to effectively interpret vast amounts of energy-related research data using standard and advanced AI approaches.

The bottom portion depicts a topic network where LDA and BERT AI techniques discover energy and sustainability themes and clusters. Clusters of related research show how AI models may eliminate human analysis, maximize computing resources, and extract relevant insights with reduced energy use.

The graphic shows that energy-efficient and sustainable AI models may be used to assess global energy research in a scalable, resource-efficient, and ecologically responsible way.

CONCLUSION

Systematic corpus production, robust preprocessing pipelines, and advanced topic modeling were identified as potential tools for developing and evaluating AI models with low energy consumption and long-term sustainability, according to this study. During the corpus generation stage, academic literature on sustainable AI and energy was meticulously chosen and arranged to reflect global research trends while maintaining data relevance and quality. Significant information extraction and reduced redundancy were made possible by a well-defined corpus, which improved computer speed. Prior to and after processing, the model has to be optimized for performance and sustainability. Computer cost was decreased by text normalization, tokenization, noise removal, and dimensional reduction; interpretability was enhanced through topic refinement and validation. These measures brought the analytical process in line with environmentally friendly and energy-efficient AI systems by optimizing data input and output to reduce energy consumption. International research contributions

exhibited coherent and well-defined subject structures when looking at the sustainable AI topic model in energy research from a global viewpoint. This study integrates clustering, LDA, and BERT for contextual semantic understanding with probabilistic topic modeling. By methodically renaming clusters into topics, starting with Topic 1 corresponding to Cluster 0 and continuing across all seven clusters, clearly thematic interpretation and comparison analysis were made feasible. Sophisticated, energy-conscious AI models were able to capture intricate study settings using this hybrid approach while maintaining scalability and sustainability.

References

- S. R.C. Group. (2018) The smart economy: The internet of everything. Available: <https://semico.com/content/smarteconomy-internet-everything>
- J. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016, <http://www.deeplearningbook.org>
- Nguyen, P.H., Kling, W.L., Ribeiro, P.F., Venayagamoorthy, G.K., Croes, R., 2013. Role of proactive behaviour enabled by advanced computational intelligence and ICT in Smart Energy Grids.
- P. Gailhofer, A. Herold, J. P. Schemmel, C. Scherf, C. Urrutia, A. Kohler, " and S. Braungardt, "The role of artificial intelligence in the european green deal," 2021.
- S. Engineering. (2021) 11 ways to reduce ai energy consumption. [Online]. Available: <https://semiengineering.com/11-ways-to-reduce-aienergy-consumption/>
- E. Comission, "Shaping Europe's digital future," European Union, Tech. Rep., 2020. Available: https://ec.europa.eu/info/sites/default/files/communicationshaping-europes-digital-future-feb2020_en_4.pdf
- Monteleoni, C., Schmidt, G.A., McQuade, S., 2013. Climate informatics: Accelerating discovering in climate science with machine learning. *Comput. Sci. Eng.* 15, 32–40. <https://doi.org/10.1109/MCSE.2013.50>
- L. Belkhir and A. Elmeligi, "Assessing ict global emissions footprint: Trends to 2040 & recommendations," *Journal of Cleaner Production*, vol. 177, pp. 448–463, 2018
- E. S. D. Nafus and K. Burch, "Carbon-responsive computing: Changing the nexus between energy and computing," *Energies*, vol. 14, no. 6917, 2021.
- L. Luce, *Democratization and Impacts of AI*. Berkeley, CA: Apress, 2019, pp. 185–195.
- Paré, G., Kitsiou, S., 2017. Chapter 9 Methods for Literature Reviews, in: Lau F, K.C. (Ed.), *Handbook of EHealth Evaluation: An Evidence-Based Approach*.
- Shah, A.S., Nasir, H., Fayaz, M., Lajis, A., Shah, A., 2019. A review on energy consumption optimization techniques in IoT based smart building environments. *Inf.* <https://doi.org/10.3390/info10030108>
- McInnes, L., Healy, J., Melville, J., 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction.
- van Wynsberghe, A., 2021. Sustainable AI: AI for sustainability and the sustainability of AI. *AI Ethics* 1, 3. <https://doi.org/10.1007/s43681-021-00043-6>
- Rose, D.C., Chilvers, J., 2018. Agriculture 4.0: Broadening Responsible Innovation in an Era of Smart Farming. *Front. Sustain. Food Syst.* 2, 87. <https://doi.org/10.3389/fsufs.2018.00087>
- Liess, S., Agrawal, S., 2017. A teleconnection between the West Siberian Plain and the ENSO region. *J. Clim.* 30, 301–315.
- Chu, S., Majumdar, A., 2012. Opportunities and challenges for a sustainable energy future. *Nature*. <https://doi.org/10.1038/nature11475>
- Ceylan, Z., 2020. Assessment of agricultural energy consumption of Turkey by MLR and Bayesian optimized SVR and GPR models. *J. Forecast.* 39, 944–956. <https://doi.org/10.1002/for.2673>
- Bianchi, F., Terragni, S., Hovy, D., 2020. Pre-training is a Hot Topic: Contextualized Document Embeddings Improve Topic Coherence.