Indian Journal of Computer Science and Technology
https://www.doi.org/10.59256/indjcst.20250403014
Volume 4, Issue3 (September-December 2025), PP: 72-78.
www.indjcst.com
ISSN No: 2583-5300

Deep Learning Approaches to Multimodal Sustainable Report Analysis

Sarah Chepkogei Sawe[1], Anthony Wanjoya[2], Andrew Kipkebut[3]
[1] *Department of Computer Science and Information Technology, Co-operative University of Kenya, Kenya.*
[2,3] *Department of Mathematical Sciences, Co-operative University of Kenya, Nairobi, Kenya.*

To Cite this Article: *Sarah Chepkogei Sawe[1],Anthony Wanjoya[2],Andrew Kipkebut[3], "Deep Learning Approaches to Multimodal Sustainable Report Analysis", Indian Journal of Computer Science and Technology, Volume 04, Issue 03 (September-December 2025), PP: 72-78.*

Abstract: *This review paper explores the application of deep learning techniques for analysing multimodal sustainable reports, with a particular focus on enhancing categorization accuracy and deriving deeper ESG insights. Sustainable reporting integrates environmental, social, and governance (ESG) data, often presenting information in diverse formats, including text, tables, images, and charts. Traditional analysis methods, relying heavily on manual coding and keyword-based algorithms, struggle with the complexity, heterogeneity, and dynamic nature of such data, leading to inaccuracies and inefficiencies. Deep learning, with its capacity to learn intricate patterns from high-dimensional and varied data sources, offers promising avenues for automated, comprehensive, and insightful analysis. This paper provides an overview of existing deep learning models (e.g., Transformer Models, Recurrent Neural Networks, Convolutional Neural Networks, Fully Connected Neural Networks) and architectures pertinent to multimodal data integration and analysis, identifies current challenges and limitations in their application to sustainable reports (such as data scarcity and the need for explainability), and proposes future research directions to enhance the efficiency and accuracy of ESG data extraction and interpretation, aiming to establish a standard for automatic sustainability report processing.*

Key Words: *Analyze how the integration of multimodal data from sustainability reports impacts categorization accuracy and provides deeper ESG insights.*

I.INTRODUCTION

The increasing global emphasis on corporate sustainability has led to a surge in the volume and complexity of sustainable reports. These reports, often published annually by companies, provide stakeholders with crucial information regarding their environmental impact, social responsibility initiatives, and governance structures. In contrast to traditional financial reports, sustainable reports are multimodal by default, containing a narrative text, structured tables, visuals (e.g., graphs, images, infographics), and, occasionally, audio/video data (Johnson et al., 2020). Extracting and manually analysing insights into these varying data types is time- consuming, inconsistent, and in many cases, does not fully reveal the scope of cross-connections in the report (Smith & Green, 2021). Conventional approaches like rule-based algorithms tend to be redundant, time-wasting, and ineffective since they cannot cover the entire width and breadth of multimodal data (Durkin, 1994).

Deep learning Deep learning is a sub-branch of machine learning that is loosely based on how the brain processes and interprets information, and has been remarkably effective in various aspects, such as natural language processing (NLP) and computer vision (CV) (Brown & Harris, 2020). Its ability to learn hierarchical representations and integrate multi-modal information makes it a decent candidate when it comes to dealing with the challenges of sustainable report analysis. This review is aimed to outline the current pool of deep learning applications in this young field, determine the key methodologies, describe the challenges, and outline the potential research directions to increase automation and efficiency in the process of categorizing

sustainability reports, reduce the number of human resources required, and provide valuable information to facilitate ESG decision-making and policymaking.

II.BACKGROUND

To fully appreciate the scope of deep learning applications in multimodal sustainable report analysis, it is essential to understand the foundational concepts:

2.1. Deep Learning Fundamentals

The use of Deep Learning models, which are Convolutional Neural Networks (CNNs) to process images or Recurrent Neural Networks (RNNs) and Transformers to process sequential data like text, has transformed the process of handling complex data. They are robust in that they learn features automatically without any manual feature engineering. Unlike the simpler ANNs, deep neural networks have several hidden layers that allow them to learn complicated relations and representations in high-dimensional data (Razavi, 2021). The deep learning theoretical framework informs the creation of models capable of working with various sources of data in a single architecture, providing unmatched capability to model non-linear, intricate relationships (Kufel et al., 2023).

2.2. Multimodal Learning

Multimodal learning focuses on building models that can process and relate information from multiple modalities. Key challenges include:

- **Representation:** How to represent information from different modalities in a unified way (e.g., joint representations, coordinated representations).
- **Translation:** How to translate information from one modality to another.
- **Alignment:** How to identify direct relationships between elements from different modalities.
- **Fusion:** How to combine information from multiple modalities for prediction or decision-making (e.g., early fusion, late fusion, hybrid fusion) (Baltrusaitis et al., 2019). The ability to use numerous data types makes multimodal approaches an upgrade of previous models limited to only single forms of data (Firmansyah, 2021).

2.3. Sustainable Reporting Frameworks

Sustainable reports will tend to follow the guidelines adopted which include the Global Reporting Initiative (GRI) Standards, Sustainability Accounting Standards Board (SASB), and Task Force on Climate-related Financial Disclosures (TCFD). Such frameworks play a key role in interpreting the points being referred to and the reporting standards. To illustrate, the suggested AI tool will aggregate criteria into two categories: Minimum and Advanced, which will be linked to the basic disclosures and advanced technical features respectively.

2.4. Research Objectives and Questions

The general objective of this research is to develop a multimodal deep learning model that will enhance the categorization of sustainability reports. The specific objectives are:

1. Analyze how the integration of multimodal data from sustainability reports impacts categorization accuracy and provides deeper ESG insights.
2. Design and implement a multimodal deep learning framework that effectively integrates multiple data modalities for sustainability report categorization.
3. Develop and optimize an ensemble deep learning model to enhance categorization rigor and efficiency.
4. Review the performance of the multimodal deep learning model against traditional text-only classification models across diverse ESG datasets and industries.

The research questions guiding this study are: a. How does integrating multimodal data from sustainability reports improve categorization accuracy and ESG insights? b. What multimodal deep learning framework best integrates different data types to enhance the categorization process of sustainability reports? c. How can an ensemble deep learning model be developed and optimized to automate sustainability report classification? d. How does the performance of multimodal deep learning models compare to traditional text-only categorization models across diverse ESG datasets and industries?

III.DEEP LEARNING APPROACHES TO MULTIMODEL SUSTAINABLE REPORT ANALYSIS

Existing research and potential approaches can be broadly categorized by how they handle different modalities and their fusion, as proposed in the project:

Text-based Analysis

Deep learning approaches for text analysis are crucial given the narrative-heavy nature of sustainability reports. These methods move beyond keyword matching to capture deeper semantic meaning and context.

- **Transformer Models (e.g., BERT, RoBERTa):**
  - **Review:** Devlin et al. (2018) introduced BERT, a groundbreaking model pre-trained on vast text corpora, capable of understanding context at both word and document levels. Its bidirectional nature allows it to consider context from both left and right, making it highly effective for tasks like text classification and information extraction in complex documents like sustainability reports. In the proposal, the project specifies that BERT or RoBERTa models will be employed to carry out the task of analyzing textual data to obtain contextual embeddings and investigate the correlation between various segments of the reports. Du Toit & Dunaiski (2024) also advocate the use of transformer-based models in hierarchical text classification, which may represent the possibility of complex and nuanced categorization.
  - **Pros:** Transformer models provide unprecedented contextual knowledge, and thus are ideal at understanding the vague and non-obvious meanings and nuances that you often encounter in ESG stories. The fact that they have been trained on general language data makes them able to be fine-tuned on relatively less domain-specific data, which is a massive advantage to the fact that there are too few annotated ESG datasets.
  - **Cons:** They are also computationally expensive, needing a lot of GPU resources to train and even make inference on large documents. It may also be difficult to explain the reasons behind a particular categorization that can be a problem in an industry with a high level of transparency such as in ESG reporting, since the nature of their black box can make this a challenge.

- **Recurrent Neural Networks (RNNs) and Long Short-Term Memory Networks (LSTMs):**
  - **Review:** RNNs and LSTMs handle sequential data, which is a good option in capturing long- range dependencies in text. The project proposal describes their purpose to work with sequential data and capture dependencies and long-range relationships within textual and numerical sequences, and also find important insights in text-intensive parts of reports. Although more recent Transformer versions have outperformed them in most benchmarks, LSTMs are still useful in certain sequential tasks, and in some cases are more efficient in certain tasks.
  - **Pros:** LSTMs are effective in capturing time-dependencies in text information which is helpful in gaining insights on flow of information and interrelated concepts in a report. They tend to be computationally less demanding than bigger Transformers, which makes them a feasible choice in resource-scarce settings.
  - **Cons:** They may have trouble with very long sequences because of the vanishing/exploding gradients, an issue solved by the attention mechanisms in Transformers. They could also perform worse than Transformers in complex contextual understanding.

- **Intelligent Corporate Sustainability Report Scoring (Shahi et al., 2012):**
  - **Review:** While not strictly deep learning, Shahi et al. (2012) explored an "Intelligent Corporate Sustainability report scoring solution using Machine Learning approach to Text Categorization." This study highlighted a four-stage process: text extraction, feature selection, document classification, and report rating based on the GRI framework. They concluded that hierarchical document classification was a good approach, and their work laid groundwork for automating CSR report analysis. This is relevant as a comparison point for traditional text classification methods.
  - **Pros:** These rule-based or traditional machine learning methods may be more interpretable and may need less data to complete given, specific tasks. They tend to be simpler to code and debug with simpler problems.
  - **Cons:** They are rigid and depend on fixed terminologies and predetermined rules, cannot handle the implicit and heterogeneous forms of new reports, and fail to handle complex unstructured text, resulting in inaccuracy and inefficiency.

Image and Visual Content Analysis

Deep learning has significantly advanced the ability to extract meaningful information from the visual components of sustainable reports, such as images, graphs, and infographics.

- **Exploring Visual Communication (Nakao et al., 2024):**
  - **Review:** Nakao et al. (2024) in particular explored the topic of visual communication in corporate sustainability reporting through the use of images recognition and deep learning. They studied the use of images in the reports of multinational corporations taking into account such aspects as industry type and cultural preferences of the company. This shows how deep learning (image recognition) can be used to learn about strategic visual information usage in sustainability practice.
  - **Pros:** The method enables automated analysis of visual information that was previously not available and similarities in corporate communication practices and develops the general perception of a message in a report. It goes beyond writing to take a comprehensive picture.
  - **Cons:** The research showed that it was limited to addressing visual elements, lacked integration with other modalities. It is still possible to have difficulties when retrieving particular numerical information or intricate trends, with extremely specificated models and possibly large and labelled databases of visual representations being necessary.

- **Pre-trained Convolutional Neural Networks (CNNs) (He et al., 2016):**
  - **Review:** CNNs such as ResNet (He et al., 2016) or EfficientNet became central to modern computer vision but are great at extracting features out of images. According to the project proposal, these CNNs will be applied in processing visual data such as images, graphs, and tables, to identify patterns, structures, trends in charts or objects in an image of a sustainability report. Hierarchical visual features, simple edges and complex parts of objects, are automatically learnt by these models.
  - **Pros:** Pre-trained CNNs provide a powerful backbone for visual analysis, enabling high accuracy in image classification, object detection, and even feature extraction for more complex tasks like graph understanding. Their ability to learn spatial hierarchies is crucial for recognizing visual patterns.
  - **Cons:** They require significant computational resources for training and fine-tuning. Their effectiveness can be limited by the quality and diversity of the visual data in reports (e.g., low- resolution images, inconsistent chart styles). They also struggle with understanding the *meaning* of the visual content without integration with text.

- **Image Processing with OCR (Tesseract OCR, OpenCV):**
  - **Review:** While not a single study, the project proposal mentions the use of Tesseract OCR and OpenCV for image processing "in order to extract pertinent visual data from sustainability reports" (Project Proposal, Page 40). This combination is essential for converting image-based text (e.g., in scanned documents, images of tables, or infographics) into machine-readable text, which can then be fed into NLP models or processed as structured data.
  - **Pros:** OCR is fundamental for unlocking information contained within images, making it accessible for further analysis. It bridges the gap between image and text modalities, a critical step for multimodal analysis of PDFs.
  - **Cons:** OCR accuracy can vary significantly depending on image quality, font styles, and layout complexity. Errors in OCR can propagate and negatively impact subsequent analysis, highlighting the importance of robust preprocessing.

Table and Structured Data Analysis

While traditional database methods are strong for perfectly structured data, deep learning offers advantages in extracting and analyzing structured information embedded within unstructured or semi-structured documents, particularly tables.

- **Fully Connected Neural Networks (FCNs):**
  - **Review:** The project proposal states that FCNs will "process structured numerical inputs, such as ESG performance metrics" (Project Proposal, Page 7, 39-40). While FCNs themselves are general-purpose networks, their application here implies processing already extracted or inherently structured numerical data. They are effective at learning complex, non-linear relationships within numerical datasets, allowing for sophisticated analysis of metrics.
  - **Pros:** FCNs can identify intricate patterns and correlations within numerical data that might be missed by simpler statistical methods. They are flexible and can be integrated into larger multimodal architectures.
  - **Cons:** FCNs primarily work with *already structured* numerical data. The challenge lies in extracting this structured data accurately from diverse report formats (e.g., images of tables). Their performance heavily depends on the quality and completeness of the input numerical features.

- **Deep Learning for Table Detection and Structure Recognition (General Application):**
  - **Review:** Although no specific study from the provided references directly details DL for table detection in sustainable reports, the project proposal highlights that deep learning can "enhance: Table Detection and Structure Recognition; Data Validation and Anomaly Detection." In the broader field of document analysis, deep learning models (often a combination of CNNs for visual layout analysis and LSTMs/Transformers for content parsing) are used to identify tables within documents, extract their boundaries, and understand their internal structure (rows, columns, headers).
  - **Pros:** This capability automates the conversion of tables (even those embedded as images) into usable structured data, overcoming a major bottleneck in manual analysis. It significantly improves the scalability of data extraction from diverse document types.
  - **Cons:** Table recognition models can be complex to develop and train, requiring large datasets of annotated tables with diverse layouts. Irregular table structures, merged cells, or complex visual formatting can still pose significant challenges.

- **Data Validation and Anomaly Detection (General Deep Learning):**
  - **Review:** The proposal states that deep learning could be used to improve Data Validation and Anomaly Detection of structured data. In practice, deep learning models, especially autoencoders or Generative Adversarial Networks (GANs) in anomaly detection, can be trained to understand the patterns in a large numerical dataset and mark anomalies. This is important to establish discrepancy in or misstatement of reported ESG measures.
  - **Pros:** Deep learning offers a powerful way to automatically identify outliers and potential data quality issues in sustainability metrics, enhancing the reliability of the analysis. It can uncover subtle anomalies that might be missed by rule-based checks.
  - **Cons:** The anomaly detection needs a lot of usual data to be trained with. Interpretation of a flagged data point as an anomaly may be difficult with complex deep learning models, which may need other explainability methods.

Multimodal Fusion Strategies

The true power of analysing multimodal sustainable reports lies in effectively combining information from different modalities. Deep learning offers sophisticated mechanisms for this fusion.

- **Multimodal Machine Learning: A Survey (Baltrusaitis et al., 2019):**
  - **Review:** Baltrusaitis et al. (2019) include a detailed survey and taxonomy of multimodal machine learning, covering issues such as representation, translation, alignment, and fusion as important challenges. This seminal paper classifies various fusion strategies, such as early fusion, late fusion, and hybrid fusion, and thus offers a basis of thinking about how various modalities can be fused.
  - **Pros:** Provides a robust theoretical foundation for understanding how to combine different data types. Highlights the various strategies available, allowing researchers to choose the most appropriate method based on the data and task.
  - **Cons:** Not being a survey, the research does not offer tangible solutions to individual issues of sustainable reporting. The decision to use a particular fusion strategy may be problem specific and empirical.

- **Late Fusion and Attention Mechanisms:**
  - **Review:** The project proposal explicitly states, "The proposed AI tool will integrate outputs from separate models through **late fusion or attention techniques**" (Project Proposal, Page 39). Late fusion involves processing each modality independently with its own specialized model and then combining the *predictions* or high-level features at a later stage. Attention mechanisms, particularly cross-modal attention, allow models to learn to weigh the importance of different modalities or parts of modalities when making a prediction, enabling dynamic and adaptive fusion (Vaswani et al., 2017).
  - **Pros:** Late fusion is easier to code and debug, since each of the modal models may be coded and optimized independently. The attention mechanisms are also strong, which enables the model to ignore irrelevant information across modalities, and captures complex inter-modal relations (e.g. a text statement that points to a particular section of a chart).
  - **Cons:** Late fusion might miss subtle, low-level interactions between modalities that could be captured by earlier fusion. Designing effective attention mechanisms for truly diverse modalities like text and complex visual layouts can be challenging and computationally expensive.

- **Multimodal Machine Learning (Kaur Hora & Shelke, 2024):**
  - **Review:** Kaur Hora & Shelke (2024) explain that multimodal algorithms would allow extracting data in every input, learning the relationships between them, and making predictions using the relationships. They underline that integrating several modalities will result in more precise and complete predictions because of a diversified comprehensive approach to the content. This strengthens the main idea of the multimodal approach of the project.
  - **Pros:** This work highlights the intuitive benefit of multimodal fusion: mimicking human perception to achieve a more complete understanding. It supports the idea that the sum of parts is greater than the individual parts in complex analytical tasks.
  - **Cons:** The "intricate relationships" between diverse data types can be very complex to model effectively. As the number of modalities and their heterogeneity increase, the complexity of designing and training an effective fusion model grows significantly. This also raises challenges in data alignment, where different modalities may not directly correspond in time or content.

Proposed AI Tool and Its Components

The research aims to develop an automated Generative AI system that extracts and summarizes essential information from PDF reports regarding sustainable development. The tool targets existing model weaknesses by improving operational efficiency and cost-effectiveness. Key components include:

- **Transformer models (BERT or RoBERTa):** To explore relationships between different sections of the reports for complex document categorization.
- **Multimodal models:** To integrate different modalities (text, structured data, images) into a unified model.
- **Fusion Mechanism (Late Fusion / Attention Mechanism):** To combine outputs from different models for enhanced classification accuracy.

The system will categorize reports based on compliance levels, validating those with 80% compliance as sustainability reports, identifying 50% compliant reports for manual review, recommending changes for 30% compliant reports, and deeming 0% compliant reports invalid.

Evaluation Metrics

The model's effectiveness will be evaluated using metrics such as F1-score, recall, accuracy, and precision. Ablation experiments will determine the contribution of each modality to overall performance. Cross-validation methods like k-fold cross-validation will ensure robustness and reduce overfitting.

Figure 1: Screenshots of the of Analysis

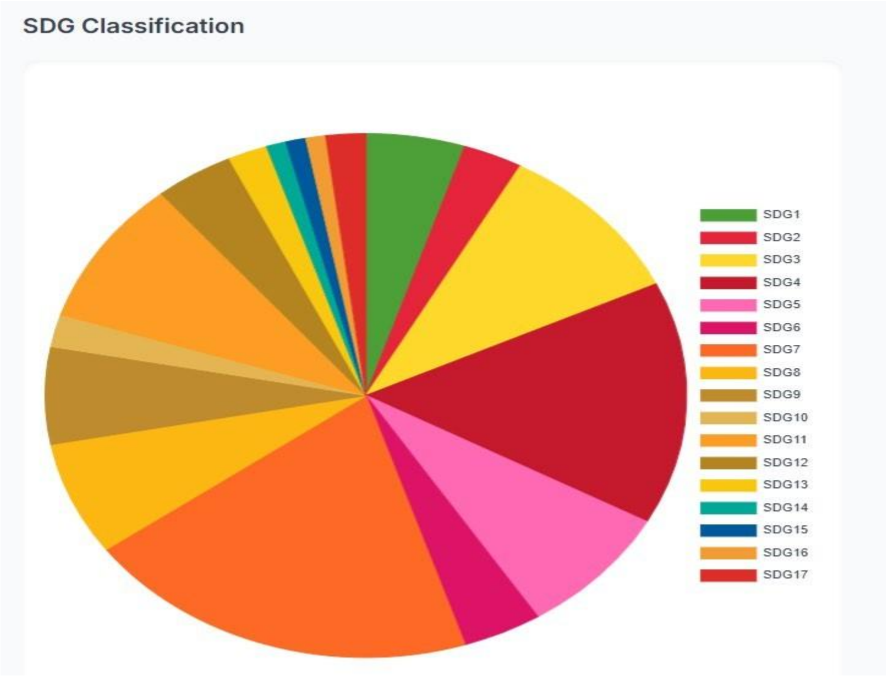

IV.CHALLENGES AND GAPS

Despite the potential, several challenges impede the widespread adoption and effectiveness of deep learning in this domain:

- **Data Scarcity and Annotation:** High-quality, multimodal, annotated datasets of sustainable reports are rare, hindering model training and evaluation. This is particularly evident in less developed nations like Kenya, where resources are limited.
- **Modality Imbalance:** The prominence of text over other modalities (e.g., few images compared to extensive text) can bias multimodal models.
- **Complex Semantic Understanding:** Sustainable concepts are often nuanced and context-dependent, requiring sophisticated reasoning beyond simple keyword matching.
- **Inter-modal Relationships:** Identifying subtle relationships between different modalities (e.g., a statement in text referring to data in a table or a trend in a graph) is difficult.
- **Interpretability and Explainability:** Deep learning models can be black boxes, making it challenging to understand how they arrive at conclusions, which is critical for trust in ESG analysis. The study aims to address this by ensuring the categorization process is transparent and explainable.
- **Dynamic Nature of Reporting Standards:** Sustainable reporting frameworks evolve, necessitating adaptable models.
- **Technological and Expertise Gaps:** In regions like Kenya, the lack of expertise and advanced computing technology (e.g., cloud computing services, GPUs) limits the adoption of deep learning models compared to simpler machine learning algorithms.
- **Data Quality Issues:** Some reports may lack proper structure or be incomplete, posing a challenge for analysis.

V.CONCLUSIONS

Deep learning holds immense promise for transforming the analysis of multimodal sustainable reports, offering a pathway to automate and enhance the extraction of critical ESG insights. The proposed multimodal deep learning model, utilizing a combination of Transformer-based models, RNNs, CNNs, and FCNs, aims to overcome the limitations of traditional text-only approaches by integrating text, image, and numerical data for more accurate and efficient categorization. While significant challenges remain, particularly concerning data availability, technological infrastructure, and model interpretability, the rapid advancements in deep learning, coupled with a growing demand for robust sustainability analysis, pave the way for innovative solutions. Future research, focusing on tailored multimodal architectures, larger annotated datasets, and improved interpretability, will be crucial in realizing the full potential of deep learning in this vital domain, especially in regions like Kenya where such tools can significantly enhance compliance and transparency in sustainability reporting.

VI.FUTURE DIRECTIONS

Based on these challenges and the evolving landscape of deep learning, future research, aligned with the project's scope, should focus on:

- **Development of Large-Scale Multimodal Datasets:** Curating and openly releasing comprehensive, annotated datasets of sustainable reports to fuel research, leveraging sources like the GRI database.
- **Novel Multimodal Architectures:** Exploring new deep learning architectures specifically designed for fusing heterogeneous sustainable report data, potentially leveraging graph neural networks for relational data or advanced attention mechanisms for cross-modal alignment.

- **Enhancing Interpretability:** Developing techniques to provide transparency into model decisions, potentially through attention visualization or rule extraction, ensuring users can understand the model's output.
- **Transfer Learning and Few-Shot Learning:** Leveraging pre-trained models from related domains (e.g., financial analysis, general NLP) and adapting them to sustainable reporting with limited annotated data.
- **Reinforcement Learning for Strategic Information Seeking:** Exploring RL agents that can learn to navigate complex reports to find specific information efficiently.
- **Integration with Domain Knowledge:** Incorporating expert knowledge and reporting standards as constraints or guiding principles for deep learning models.
- **Real-time Analysis and Monitoring:** Developing systems capable of continuously processing and updating insights from sustainable reports as they are released.
- **Addressing PDF Translation and Visual Data Extraction:** Advancing towards reading graphs and image-based tables in future developments.
- **Ensuring Algorithmic Fairness and Data Diversity:** Embracing dataset diversity to avoid categorization biasness.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... Zheng, X. (2016).TensorFlow: A system for large-scale machine learning. In 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16).
2. Baltrušaitis, T., Ahuja, C., & Morency, L. P. (2019). Multimodal machine learning: A survey and taxonomy. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(2), 423– 443.
3. Brown, T. B., Mann, B., Ryder, N., Subbiah, D., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language Models are Few-Shot Learners. arXiv preprint arXiv:2005.14165.
4. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
5. Durkin, J. (1994). Expert systems: A view of the field. IEEE Expert, 9(4), 56–63.
6. Firmansyah, B. (2021). The effectiveness of multimodal approaches in learning. EDUTEC Journal of Education and Technology, 469-479.
7. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770–778).
8. Johnson, L., Rinaldi, L., De Micco, P., Vitale, G., Cupertino, S., & Maraghini, M. P. (2020). Challenges in Sustainability Reporting. Journal of Environmental Accounting and Management, 8(3), 200–215.
9. Kaur Hora, T., & Shelke, S. (2024). Multimodal machine Learning. PICT's International Journal of Engineering and Technology (PIJET), 66–73.
10. ufel, J., Bargieł-Łączek, K., Kocot, S., Koźlik, M., Bartnikowska, W., Janik, M., ... Gruszczyńska, K. (2023). What Is Machine Learning, Artificial Neural Networks and Deep Learning? Examples of Practical Applications in Medicine. National Library of Medicine, 13(15), 2582. https://doi.org/10.3390/diagnostics13152582
11. Ozili, P. K. (2022). Sustainability and sustainable development research around the world. Managing Global Transitions.
12. Razavi, A. (2021). Introduction to Artificial Neural Networks. CRC Press.
13. Smith, J., & Green, L. (2021). The Evolution of Corporate Sustainability Reporting. Environmental Science & Policy Journal, 14(2), 150–165.
14. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research, 15(1), 1929–1958.
15. Tang, G. (2024, June 26). Application and Development of Artificial Intelligence in the ESG Investment and Financing Field. HKAIFT. Retrieved from https://hkaift.org/en/article/show.php?itemid=227.
16. Wang, Y., & Ma, Y. (2024). Application of Deep Learning Models to Predict Panel Flutter in Aerospace Structures. Aerospace, 11(5), 677.
17. Werbos, P. (1988). Backpropagation: Past and future. In IEEE Xplore.
18. Zhang, L., & Wang, Q. (2018). Artificial Intelligence in Data Analysis. International Journal of Machine Learning and Cybernetics, 9(5), 789–800.