Introduction

This research aims to identify a high-level, broad summation of the sustainability impacts of ANN. The essay ignored technical issues relating to ANN, as the technology in this field is rapidly evolving. This essay is not a literature review; a systemic approach to article extraction and filtration was not used and subjective judgment was used for article selection. Many sources of relevance have likely been ignored or overlooked. This limits the accuracy of this research.

Context

Terminology

This report focuses specifically on artificial neural networks (ANNs). The word ANN is used to distinguish this subject matter from biological neural networks (i.e., the brain), which is not within the scope of this report. ANN is a derivation of machine learning (ML) techniques, and the use of multiple layers in ANN leads to the term 'deep learning' (DL) (Hardesty, 2017.) An issue with literature relating to neural network sustainability is that the term 'AI' is used interchangeably with deep learning models, even though the term 'AI' is quite broad (Ligozat et al., 2021.) This is an issue because it implies that research relevant to neural network sustainability may have been excluded or overlooked because of the use of alternative keywords. Throughout this report, efforts were made to ensure that sources were relevant to neural networks. Incidentally, some sources may refer to the following umbrella concepts that ANN are derived from: AI, ML. Because Deep Learning is derived from ANN, the two can be considered interchangeable for this report.

Software

An Artificial Neural Network (ANN) is a computer system functionally similar to the human brain. It works by having multiple processors ('neurons') and connecting them so they can pass information to each other. ANNs are typically trained using large datasets that have non-linear relationships. The data set used for training is used as a basis for future analysis by the ANN; humans set allowable deviations from the training data. Quantitative metrics, such as maximum repetitions or processing time, determine when training stops (El-Shahat, 2018.)

Neural networks are designed to counteract hardware limitations. Neurons are connected to form layers, and synapses connect layers. Synapses can only connect two neurons. This allows information to be processed and stored parallel via different layers, reducing overall power consumption (Bouvier et al., 2020.)

There are risks associated with designing ANNs. Inputting excess values will reduce processing speeds, whereas having too few variables means information between nodes is not fully tracked. This is further complicated by other software issues beyond the scope of this research. ANN require minimal external assistance to learn. They may one day surpass human capabilities (El-Shahat, 2018.)

Another risk is overfitting - when the ANN aligns too closely with the dataset - following its general trends and anomalies. As the number of layers in an ANN increases, the computations and risk of overfitting the data increase, too (Baklacioglu, Turan & Aydin, 2018.)

Jamli & Farid (2019) compared ANN and non-ANN approaches in their study. They found that ANNs were more accurate than non-ANN models because they accounted for certain real-world complexities that the simpler models did not. Despite being increasingly inaccurate, non-ANN approaches are still widely used because the ANN approach is difficult to implement.

Machine learning algorithms are restricted by their hardware, either because they require a lot of time and power or because of memory limitations. Neural networks need to be correctly mapped to the hardware they operate on, which must be specific to allow the algorithms to evolve. Research on online learning is already underway, but this implementation would require memory allocation on hardware (Bouvier et al., 2020.)

Sustainability

Smith and Wong (2022) examined the tri-dimensional sustainability of ANNs, but only in the context of the construction industry. Of the literature reviewed, economic sustainability was the most studied, and environmental and social sustainability became a bigger focus in the late 2010s. When examining sustainability criteria, fuzzy logic has to be used. This is meant to emulate human reasoning in decision-making where uncertainty is involved. It works by assessing the accuracy of datasets which are not fully accurate or complete. One case study used weighted rankings based on expert guidance from questionnaires to form sustainability criteria against which projects were judged. The issue with this example is that they relied on the subjective opinion of experts and small sample sizes (Smith and Wong, 2022)

Social

The use of ANN was better at evaluating safety issues than subjective opinion. However, another case study, which tried to examine the conditions leading up to accidents, though effective, only used quantitative data. The absence of near-miss accidents in the datasets could be included to increase the available data. Though this would help prevent accidents, it would be difficult to quantify the improvements made specifically because of the model (Smith and Wong, 2022.)

Potential Bias and Inaccuracy

There are additional social sustainability issues to ANN: Unrepresentative training data leading to biases in models and privacy issues arising from facial and emotional recognition (van Wynsberghe, 2021.)

Smith Wong (2022) discussed how ANNs can broaden the capacity of decision-making systems, helping with data analysis and providing guidance to users. Of note is the presence of an interconnected knowledge base, which would store data related to the past and decision-making. In theory, this would make the system better over time. What the source does not mention, however, is that if historical data is biased (as other research has shown to be possible), it would warp the output of the ANN and create biased outputs.

ANNs have an issue working effectively in contexts underrepresented in their training sets. ANNs were less accurate when analysing satellite image data of poor, remote areas. This is an issue because international development programmes will favour using ANN and satellite imagery for on-the-ground data collection. In the case of humanitarian aid, there is a risk that an ANN will discriminate against the communities it is trying to help because said communities are underrepresented in the training dataset (Kim et al., 2021.)

The risk of preexisting datasets influencing AI has been documented. A study found that ML algorithms focused on language processing had a gender bias, which could be mitigated by breaking associations between words and gender in training data (Lu et al., 2020.) However, this would increase the resources needed to train such models, in addition to the substantial resource consumption already expended in their initial training. Though the author reduced gender bias, other implementations may be less effective. The risk of human bias influencing adjustments plays a factor, too. This also highlights a potential risk in datasets in general that may result in biased neural networks if they are used for training.

Baklacioglu, Turan & Aydin (2018) mention discrepancies in previous literature regarding adequate sampling data and what proportions should be used for training, cross-validation, and testing. This presents issues in the future as a lack of a universal ruleset in the implementation of ANN on data could result in variable performance and accuracy across these algorithms when applied. This has unique implications in sustainability issues, where implementing such algorithms will influence decision-making that could impact people's lives. The lack of a universal ruleset makes it difficult to have an objective framework for what constitutes 'good' ANN data.

A case study was examined to explore the limitations of statistical analysis by ANN. Jamli & Farid (2019) write that ANN could be used in engineering applications but are also critical of the suitability of ANN in FE analysis, due to deficiencies in engineering knowledge and the requirement of statistical analysis in ANN software. ANN, in this context, can only make predictions within the range of input data it was trained with. The source fails to mention that this limitation is not exclusive to ANN, as linear regression likewise is only accurate within a relevant range, indicating this is a limitation borne out of the nature of statistical analysis. Beyond the context of engineering, this implies that ANNs are still bound by the limitations of statistical analysis. It also implies that human limitations in knowledge will hamper the effectiveness of ANN. Several potential problems in wider applications can be implied by examining this case study: people without an understanding of the issues with extrapolation may apply inaccurate statistical analysis, and predictions may be warped incorrectly due to unrepresentative input data. This has negative public policy considerations if the technology is applied in that regard, as data from ANN will be used to influence people's lives.

A combination of ANN and non-ANN modelling proved successful but still required significant training data and suffered from the inaccuracy of non-ANN models (Jamli & Farid, 2019.) This implies that to eliminate the issues of non-ANN models fully, it cannot be used at all. Though theoretically viable, its myriad issues make it impractical.

Feinman & Lake (2018) looked at the development of the inductive bias in ANN, which informed interpretation beyond training data. They found that simple ANNs could develop this bias with only three images and that the conditions to develop these biases were not understood. However, it was also noted that the development of these biases aided with accelerated learning and that human brains use inductive biases in early development, too. This implies that the use of biases can aid in training efficiency.

The relative accuracy of different ANN models at gender recognition was studied. There was varying performance between models, with only two being relatively accurate. Though powerful, these models had issues. One required several powerful GPUs to train and high computer storage for its calculations. The other, though more capable than other ANNs at having multiple layers without accuracy impairment, still suffered lower accuracies with excessive layers. In this study, examining gender bias took priority. This meant that only gender recognition accuracy, and not accuracy overall, was examined in the ANN models.

Additionally, the datasets used included real and fake images. (Gwyn & Roy, 2022.) There is a possibility that the most accurate model is not the most accurate at gender recognition.

Though Gwyn & Roy (2022) tried to assess ethnic bias, they could not find a sufficiently diverse dataset and suggested that future datasets should contain demographic labels so that ANN accuracy across ethnicities can be assessed (Gwyn & Roy, 2022.)

Serna et al. (2020) write that AI can become biased through unbalanced training data and that this bias can become deeply ingrained and hidden during the training process. They also note that the trend for AI has been of increased performance and a growing influence in people's lives but low transparency. The issues raised by the author relate to a lack of transparency, discrimination and privacy issues. It is implied that the danger is that significant decisions will be made by machines whose decision-making processes we do not understand.

Though substantial literature on how bias affects different groups, not much literature has been published on how bias impacts algorithm training (Serna et al., 2020.) This implies a theoretical blind spot in our understanding of how biased and unbalanced datasets can impact AI learning.

Kim et al. (2021) mention a recent example where biases in models led to misdiagnoses of COVID-19 because the model was unable to operate effectively in areas where other diseases were common. The author attempted to eliminate bias in the ANN used in their study but noted that it was a modified version of another ANN. Since those ANNs likely had biased datasets, the bias in the modified ANN could not be fully rectified. The main finding is that in the context of identification for humanitarian aid, ANNs were least effective for the most vulnerable populations.

Though highly accurate, the performance increase of neural networks in image recognition and classification has plateaued recently (Stock and Cisse, 2018.)

This is concerning because it implies that no further significant improvements will be made suggesting that bias within image recognition will be a long-standing issue. Stock and Cisse (2018) discussed the performance of an ANN with image classification and recognition. When testing model accuracy, they found that when some groups were over or underrepresented, the model was more likely to make thematic assumptions based on human skin colour. It was unclear why the model developed the biases it did. This is concerning both from a research and a practical perspective because if we are unable to identify the logic process that forms biases, we cannot address them.

Stock and Cisse (2018) proposed a method of automated bias detection, which, though effective, may not be illustrative of its effectiveness outside of this study. The author suggests that further research into machine learning will encourage future breakthroughs. However, this is an optimistic perspective and assumes that future breakthroughs are possible, which may not be the case.

A key theme in the work of Kim et al. (2021) was the accuracy discrepancy of an ANN in rich and poor areas. The ANN was more accurate in the former and less accurate in the latter. The author mentions that the development of the ANN of concern in their study was in the developed world, but this trend will likely apply to many ANNs. This suggests that ANN will be less effective in LEDCs because the data used to train them is not representative.

Bias Variance Trade-off

A counterargument to having larger, more varied sample sizes can be found in the bias-variance trade-off, as explained by Neal et al. (2018.), who write that model complexity is inversely proportional to bias and proportional to variance and eventually leads to more error.

However, Neal et al. (2018) also write that recent research suggests this is not true for neural networks, contesting previous literature. The author cites existing literature and critically examines a previous study (Geman et al., 1992) wherein the bias-variance trade-off hypothesis was tested with modern neural networks. Their results showed that model complexity was inversely proportional to variance and bias and, thus, error. This implies that to address the issues of bias in a neural network, large, varied and ideally representative data sets need to be used for training ANNs.

Yang et al. (2020) contest this view, arguing that the bias-variance hypothesis holds merit and that while model complexity is inversely proportional to bias, variance has a weighted distribution in relation to model complexity, which increases and then decreases over time. Of note is that the model used in this study is less sophisticated than models that would be used in real life. This suggests the findings are strictly theoretical and may not be representative. Though Yang et al. (2020) were able to quantify variance behaviour concerning complexity, an explanation could not be provided, which limits the utility of this source.

Though not expressly stated in the source, their findings might suggest that careful and rigorous testing of model outputs concerning bias and variance must be performed before they become widely used. However, this is an unrealistic expectation because of the time and costs associated with it. The discrepancies in the literature regarding inaccuracy levels in ANN make it difficult to conclude how ANN training data sets should function in the future. This might suggest that consequences relating to inaccuracy due to bias or variance are unavoidable.

Environmental

Mapping environmental impacts

Baklacioglu, Turan & Aydin (2018) found that the co-relation coefficients of the model output during testing were above 0.9, indicating the models to be highly accurate in examining environmental impact indicators. The authors mention that ANN models and their starting limits must be properly optimised for their corresponding purpose. They suggest implementing similar techniques into more flight-based technologies. Contrary to this source's findings, it is worth noting the following: the success of one implementation does not warrant universal application, and though statistical techniques were used in verifying the accuracy of the model, this implementation would need to be shown to be reproducible and consistent in similar applications by other parties. Additionally, the authors state that the approach mentioned in his article had no prior reported studies, which means issues specific to their approach may not have been identified.

Water Consumption

The water consumption of AI models is not widely acknowledged and has been suppressed by some companies. Additionally, the environmental focus of AI models is usually on their carbon footprint. This happens because water is consumed to generate electricity for servers or cool them (Li et al., 2023).

The water consumption of AI models is likely to increase as models get more sophisticated, although the author does state that more evidence is needed to assert that claim. Remedies to the issue focus on supplying more water and not on reducing demand. The remedies for carbon efficiency and water efficiency are mutually exclusive. This is because the former relies on access to sunlight, whereas the latter relies on avoiding sunlight (Li et al., 2023.) If these trends continue in the future, the problem of ANN water consumption will get worse.

Carbon Emissions

ANNs have been used to track emission profiles across different power settings and torques, providing information that could be used to improve performance. This approach in environmental analysis could be used for business jets and helicopters to maximise environmental sustainability. The study used a genetic algorithm approach, which aims to mimic the principles of natural selection and adaptation to algorithms. Although this improved the algorithm's ability to generalise, it had issues: the fact that an initial sample needs to be created randomly and that the algorithm could develop in unintended ways. (Baklacioglu, Turan & Aydin, 2018.)

Van Wynsberghe (2021) proposes a proportionality framework so the tasks an AI is designated justify its carbon footprint. Policy considerations can also be made so that ANN models are not used on menial tasks. However, It could be argued that this counters one of the main advantages of ANN, which is the ability to automate menial tasks.

The increased accessibility of more powerful hardware means that more people incorporate neural networks. New models are often trained on more powerful GPUs on larger datasets for longer periods. This requires more time and energy (Lacoste et al., 2019.) If this trend continues, future advances in hardware and neural networks will increase the emission profiles of neural networks. This last point is subject to the caveat that past performance does not indicate future performance. Extrapolation is only accurate when other variables align with the past, and advances in hardware may mean that past trends with neural network emissions may not reflect what future hardware is responsible for. This remains a possibility, as Lacoste et al. (2019) highlight the power efficiency of niche GPUs over conventional ones.

Lacoste et al. (2019) found that emission profiles fluctuated based on location, with the highest emissions being in Africa and the lowest ones in North America and Europe. However, only a single data point was available for Africa, and the latter two countries have the highest ranges of the dataset. In addition, Europe had a significant variance in the quartiles of their data. The high variance and ranges suggest inconsistencies with how servers operated throughout the continent. It is also worth noting that since Africa has only one data point to compare to, it may not be reflective of how servers would operate in the rest of the continent - more data would be needed to validate this point. Although the author mentions the high variation in certain regions, they neglect to mention that this variation could only exist because certain countries had more data points to work with.

Lacoste et al. (2019) also found that fine-tuning pre-trained models for specific tasks is as effective as training models without prior training. Although not expressly stated in the research, this implies that carbon emissions can be reduced by fine-tuning pre-trained models since it would reduce the time and, thus, emissions required.

One challenge with neural networks is that both direct and incidental activities associated with their use need to be accounted for. Ligozat et al. (2021) write that 'AI services use many data and that the various stages of this data use - acquisition, storage, transfer and processing - all require energy.

Raw Material Consumption

Ligozat et al. (2021) mention that the approaches theorised for quantifying AIs focus on using data and propose examining their impact across a lifecycle. The primary impact focuses on the hardware that AIs are used on - the extraction of raw materials, hardware manufacturing, transportation and distribution, data usage (which includes acquisition, storage, transfer and processing) and hardware recycling and disposal.

Neural networks can also be trained to adapt to hardware, an example of this being software accounting for large differences in nanodevices. Research is underway that would allow hardware to mimic neuroplasticity in the brain. This would allow hardware to adapt over its lifetime, meaning it would never become redundant (Bouvier et al., 2020.) Theoretically, this could reduce the raw resource consumption of ANNs.

An example of the potential for ANNs in environmentally sustainable practices is in the automobile industry, which is trying to reduce the weight of its vehicles, improve its environmental sustainability, and increase its safety during collisions. ANNs are used in sheet metal forming to predict springback -

when a material tries to return to a previous molecular formation after deformation. This is embedded into Finite Element (FE) analysis and reduces wastage costs (Jamli & Farid, 2019.)

ANNs reduce the experiments needed to determine material behaviour, but a large amount of information is required for this ANN to be accurate, which is costly (Jamli & Farid, 2019.) This data will likely need to be collected via experiments. This means that the costs associated with requiring fewer experiments to determine material behaviour are offset by the increased costs associated with implementing an ANN. It also implies that in the short to medium term, its implementation will be exclusive to only those who can afford it. The practical restrictions on ANN implementation also reduce the potential pool for future research to explore the issue further. The author cites other studies comparing ANN and non-ANN models, but the former was not implemented even after these studies occurred. The complexity of ANN makes it difficult to conduct further research. This suggests that non-ANN approaches are favoured in the context of FE analysis, but it is unclear if this trend would extend to other applications. The difficulty of further study further hampers potential progress in expanding the implementation of ANN in the future.

Tertiary Effects

Ligozat et al. (2021) highlight that short-term gains using AI would have unintended consequences and that the literature on environmental sustainability is theoretical and neglects changes in socio-technical behaviour. Examples include:

An increase in efficiency in energy consumption can lead to increased consumption over time. The development of automated vehicles would increase vehicle use and, thus, consumption. Increasing reliance on AI would consequently increase the demand for raw materials, specifically lithium and cobalt. This would necessitate opening mines, which would increase fossil fuel emissions.

This consequential impact would imply that the increased energy efficiency using neural networks may not suffice if human behaviours are not aligned with sustainability. This is beyond the scope of the report. However, it does highlight a limitation of neural networks in sustainability - they are only as effective as the people who implement and use them.

ANN have a significant environmental impact that conflicts with social sustainability. The time and energy required to train and refine neural networks is significant - training a large neural network created an emission profile equivalent to the lifetime of five cars. This energy could be expended on other activities - such as heating or electricity. Significant energy costs are a social sustainability issue because many people in the world are still in abject poverty. In addition, in many locations, carbon offsetting is not viable (van Wynsberghe, 2021.)

GPT-3 training in Microsoft consumes between 700,000 and 21,000 litres of fresh water, depending on location. This is problematic because of global freshwater scarcity, the impact on regional water systems and the fact that servers are sometimes located in areas affected by draughts (Li et al., 2023.) This can lead to social sustainability issues, tighten resource constraints and raise ethical issues because the significant water consumption is at the expense of provision to those in dire need.

Economic

Prediction

ANN can help with cost estimation in the early stages of a project. Case-based reasoning uses the costs from previous examples to estimate future ones. Accuracy is subject to the following: the number of cases, the type of cases and the similarity of cases to future ones. Though they have been proven to predict costs accurately, it is only when projects have a high degree of similarity to each other (Smith and Wong, 2022.) This may not be the case for more complex or less similar projects.

Lobeev (2021) used an ANN to predict bankruptcy in companies with high research and development (R&D) spending, with the model correctly predicting bankruptcy 91% of the time. This was important because the author claimed there was no research on models for this specific company type predating their work and that companies with high R&D spending were more prone to risk. After all, the increased R&D spending led to reduced free cash flow and reduced financial sustainability.

Gavurova et al. (2022) were able to use an ANN to predict bankruptcy companies, noting it would be particularly effective in economies that are not fully developed. They noted that using conventional financial indicators was easy because they were standard in financial statements (the issues specific to these are discussed elsewhere.) However, the limitations of this study are the following: it only focused on select branches of certain industries in The Slovak Republic, and it was noted that different indicators would be needed for different industries. In addition, this study only examined international companies across two years in select branches of two industries. This potentially limits the applicability of their findings.

ANNs can be used to highlight financial distress in companies. This allows shareholders to be informed of potential issues. ANNs can be used for decision-making purposes for investors and creditors to mitigate potential losses. Though it is possible to use ANNs to solve complex financial problems, statistical differences between the financial services and manufacturing industries indicate that ANNs analysing economic data must be specific to their industry. This study was conducted in South Africa and may not be applicable elsewhere (Dube, Nzimande & Muzindutsi, 2023.)

Though the literature regarding stock return forecasting shows positive results, it is only because these models are tested within data sets of very short time frames. Over the long term, a study found that different models had variable prediction accuracy on stock returns and that the models were most effective when examining shorter timeframes. In the long term, the models have prediction issues because stock return predictability is episodic and unstable, partially because of how data is split and because investor behaviour influences prices. findings suggested that practical applications of ANN in stock market prediction may be difficult (Chudziak, 2023.)

Information Asymmetry

A (highly simplified) summation of the seminal work of Akerof (1970) was used for this explanation: Information asymmetry means sellers have more information than buyers. Sellers will take advantage of this to sell bad products. Customers would tell others about their bad experiences. This would deter people from paying high prices. This would deter quality goods from entering the market and only encourage bad ones. Thus, the market either is either dwindled or destroyed.

An issue with economic sustainability is information asymmetry, which can shrink or destroy markets. Companies' financial issues can have consequences for the wider economic environment. However, information on these issues can be difficult to detect by most stakeholders because financial reports can be manipulated. The use of ANN in financial distress prediction can aid in their prevention. The use of ANN and other statistical and mathematical tools resulted in a model that had superior financial distress accuracy prediction. However, because the study was focused in Taiwan, the findings may not be applicable elsewhere (Jan, 2021.)

Jan (2021) discusses the issues involved in producing financial statements. Information asymmetry is common, and it is difficult for most stakeholders to know the true financial position of a company. Fraud is becoming harder to detect as criminals become more sophisticated, and it can cause more damage than asset misappropriation or corruption. The creation of an ANN was effective in detecting financial statement fraud and suggested that failings in conventional approaches of fraud detection could be overcome. However, the author noted there was limited relevant literature on using algorithms to detect fraud. Additionally, their suggestions on increasing economic sustainability extended beyond the use of ANN, suggesting that the use of ANN alone is not sufficient to increase economic sustainability.

ANN has been used to investigate the correlation between economic and social sustainability and novel management accounting tools. Though ANN was not used to increase performance directly, its investigative role could be replicated by other organisations when new technologies and their impacts on financial and sustainability performance need to be investigated (Vărzaru et al., 2022.) The limits of this study include that it only examined Romanian companies.

Xu et al. (2019) found that ANN was effective at evaluating credit risk compared to other models. However, they also stated that AI (and, by extension, ANN) cannot explain the reasoning behind their outputs and, therefore, lack transparency. This is why they have not had more widespread adoption in situations where interpretation is necessary and why the author cautions against overreliance on AI and an overzealous pursuit of efficiency and innovation.

Inequality

Bouvier et al. (2020) write that a cloud connection to a more powerful device is usually needed, implying that most consumer-grade devices cannot process such algorithms natively. The reliance on a cloud connection presents security risks and complicates implementation. Though not mentioned in the source, one potential issue that can be inferred is that the users who can fully utilise ANN are only those with very powerful hardware, which requires significant funds. This means that existing economic inequalities will be compounded by inequalities in ANN utilisation, too.

Conclusion

This report aimed to identify the tri-dimensional impact of ANNs. ANNs have the potential to help with the collation and analysis of data, decision-making and forecasting. However, they suffer from a number of issues which hampers their sustainability. Bias and inaccuracy can be borne of training data and human/statistical limitations. They require energy and water to operate, and the mitigation strategies for both are mutually exclusive. Unintended consequences may arise, including increased resource consumption and said resources being expended as opposed to others in more dire need. Regarding finance, ANNs are not practical for long-term forecasting of share returns, though they can be effective in very specific conditions provided the ANN is tailored to the industry.

References

AKERLOF, G. A. (1978). THE MARKET FOR “LEMONS”: QUALITY UNCERTAINTY AND THE MARKET MECHANISM**The author would especially like to thank Thomas Rothenberg for invaluable comments and inspiration. In addition he is indebted to Roy Radner, Albert Fishlow, Bernard Saffran, William D. Nordhaus, Giorgio La Malfa, Charles C. Holt, John Letiche, and the referee for help and suggestions. He would also like to thank the Indian Statistical Institute and the Ford Foundation for financial support. In *Uncertainty in Economics*. https://doi.org/10.1016/b978-0-12-214850-7.50022-x

Babikir, A., Mohammed, M., Mwambi, H., Babikir, A., Mohammed, M., & Mwambi, H. (2018). Dynamic Factor Model and Artificial Neural Network Models: To Combine Forecasts or Combine Models? *Advanced Applications for Artificial Neural Networks*. https://doi.org/10.5772/INTECHOPEN.71536

Baklacioglu, T., Turan, O., & Aydin, H. (2018). Metaheuristic approach for an artificial neural network: Exergetic sustainability and environmental effect of a business aircraft. *Transportation Research Part D: Transport and Environment*, *63*, 445–465. https://doi.org/https://doi.org/10.1016/j.trd.2018.06.013

Bâra, A., & Oprea, S. V. (2018). Electricity Consumption and Generation Forecasting with Artificial Neural Networks. In *Advanced Applications for Artificial Neural Networks*. https://doi.org/10.5772/intechopen.71239

Bouvier, M., Valentian, A., Mesquida, T., Rummens, F., Reyboz, M., Vianello, E., & Beigne, E. (2020). Spiking Neural Networks Hardware Implementations and Challenges: A Survey. *ACM Journal on Emerging Technologies in Computing Systems*, *15*(2), 1–35. https://doi.org/10.1145/3304103

Chudziak, A. (2023). Predictability of stock returns using neural networks: Elusive in the long term. *Expert Systems with Applications*, *213*. https://doi.org/10.1016/j.eswa.2022.119203

Dube, F., Nzimande, N., & Muzindutsi, P. F. (2023). Application of artificial neural networks in predicting financial distress in the JSE financial services and manufacturing companies. *Journal of Sustainable Finance and Investment*, *13*(1). https://doi.org/10.1080/20430795.2021.2017257

El-Shahat, A. (2018). Introductory Chapter: Artificial Neural Networks. In A. El-Shahat (Ed.), *Advanced Applications for Artificial Neural Networks* (p. Ch. 1). IntechOpen. https://doi.org/10.5772/intechopen.73530

Feinman, R., & Lake, B. M. (2018). Learning Inductive Biases with Simple Neural Networks. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society, CogSci 2018*, 1657–1662.

Filho, A. O. B., & Viegas, I. M. A. (2018). Applications of Artificial Neural Networks in Biofuels. In *Advanced Applications for Artificial Neural Networks*. https://doi.org/10.5772/intechopen.70691

Gavurova, B., Jencova, S., Bacik, R., Miskufova, M., & Letkovsky, S. (2022). Artificial intelligence in predicting the bankruptcy of non-financial corporations. *Oeconomia Copernicana*, *13*(4). https://doi.org/10.24136/oc.2022.035

Gruber, M., Trüschel, A., & Dalenbäck, J. O. (2014). CO2 sensors for occupancy estimations: Potential in building automation applications. *Energy and Buildings*, *84*, 548–556. https://doi.org/10.1016/J.ENBUILD.2014.09.002

Gwyn, T., & Roy, K. (2022). Examining Gender Bias of Convolutional Neural Networks via Facial Recognition. *Future Internet*, *14*(12). https://doi.org/10.3390/fi14120375

Hardesty, L. (2017, April 14). *Explained: Neural networks*. MIT News Office. https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414

Jamli, M. R., & Farid, N. M. (2019). The sustainability of neural network applications within finite element analysis in sheet metal forming: A review. *Measurement*, *138*, 446–460. https://doi.org/https://doi.org/10.1016/j.measurement.2019.02.034

Jan, C. L. (2021). Financial information asymmetry: Using deep learning algorithms to predict financial distress. *Symmetry*, *13*(3). https://doi.org/10.3390/sym13030443

Jan, C. L. (2021). Detection of financial statement fraud using deep learning for sustainable development of capital markets under information asymmetry. *Sustainability (Switzerland)*, *13*(17). https://doi.org/10.3390/su13179879

Kim, D. H., López, G., Kiedanski, D., Maduako, I., Ríos, B., Descoins, A., Zurutuza, N., Arora, S., & Fabian, C. (2021). Bias in deep neural networks in land use characterization for international development. *Remote Sensing*, *13*(15). https://doi.org/10.3390/rs13152908

Lacoste, A., Luccioni, A., Schmidt, V., & Dandres, T. (2019). *Quantifying the Carbon Emissions of Machine Learning*. https://arxiv.org/abs/1910.09700v2

Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). *Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models*. https://arxiv.org/abs/2304.03271v1

Ligozat, A.-L., Lefèvre, J., Bugeau, A., & Combaz, J. (2021). Unraveling the Hidden Environmental Impacts of AI Solutions for Environment. *“Towards the Sustainability of AI; Multi-Disciplinary Approaches to Investigate the Hidden Costs of AI.”* https://arxiv.org/abs/2110.11822v2

Lobeev, I. (2021). Bankruptcy Prediction for Innovative Companies. *Journal of Corporate Finance Research / Корпоративные Финансы | ISSN: 2073-0438*, *15*(4). https://doi.org/10.17323/j.jcfr.2073-0438.15.4.2021.36-55

Lu, K., Mardziel, P., Wu, F., Amancharla, P., & Datta, A. (2020). Gender Bias in Neural Natural Language Processing. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 12300 LNCS*. https://doi.org/10.1007/978-3-030-62077-6_14

Neal, B., Mittal, S., Baratin, A., Tantia, V., Scicluna, M., Lacoste-Julien, S., & Mitliagkas, I. (2018). A Modern Take on the Bias-Variance Tradeoff in Neural Networks. *CoRR*, *abs/1810.08591*. http://arxiv.org/abs/1810.08591

Pontes-Filho, S., Olsen, K., Yazidi, A., Riegler, M. A., Halvorsen, P., & Nichele, S. (2022). Towards the Neuroevolution of Low-level artificial general intelligence. *Frontiers in Robotics and AI*, *9*, 299. https://doi.org/10.3389/FROBT.2022.1007547/BIBTEX

Serna, I., Peña, A., Morales, A., & Fierrez, J. (2020). Insidebias: Measuring bias in deep networks and application to face gender biometrics. *Proceedings - International Conference on Pattern Recognition*. https://doi.org/10.1109/ICPR48806.2021.9412443

Smith, C. J., & Wong, A. T. C. (2022). Advancements in Artificial Intelligence-Based Decision Support Systems for Improving Construction Project Sustainability: A Systematic Literature Review. In *Informatics* (Vol. 9, Issue 2). https://doi.org/10.3390/informatics9020043

Stock, P., & Cisse, M. (2018). ConvNets and imagenet beyond accuracy: Understanding mistakes and uncovering biases. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *11210 LNCS*. https://doi.org/10.1007/978-3-030-01231-1_31

SustainML. (n.d.). *Carbontracker: Predicting the Deep Learning CO2 footprint*. Retrieved August 14, 2023, from https://sustainml.eu/index.php/news/17-sustainml-carbontracker

Turchin, A. (2019). Assessing the future plausibility of catastrophically dangerous AI. *Futures*, *107*. https://doi.org/10.1016/j.futures.2018.11.007

van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics 2021 1:3*, *1*(3), 213–218. https://doi.org/10.1007/S43681-021-00043-6

Vărzaru, A. A., Bocean, C. G., Mangra, M. G., & Mangra, G. I. (2022). Assessing the Effects of Innovative Management Accounting Tools on Performance and Sustainability. *Sustainability (Switzerland)*, *14*(9). https://doi.org/10.3390/su14095585

Xu, Y. Z., Zhang, J. L., Hua, Y., & Wang, L. Y. (2019). Dynamic credit risk evaluation method for e-commerce sellers based on a hybrid artificial intelligence model. *Sustainability (Switzerland)*, *11*(19). https://doi.org/10.3390/su11195521

Yang, Z., Yu, Y., You, C., Steinhardt, J., & Ma, Y. (2020). Rethinking Bias-Variance Trade-off for Generalization of Neural Networks. In H. D. III & A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning* (Vol. 119, pp. 10767–10777). PMLR. https://proceedings.mlr.press/v119/yang20j.html