

# How to apply green AI in practice?

## *Moving from FLOPs to CO<sub>2</sub> footprint*

*Qin Zhao and Gerard Schouten*

### Abstract

Artificial intelligence (AI) is revolutionising industry and transforming society. However, the rapid growth of AI has also raised concerns about the environmental impact of this technology, as the energy consumption required to train and use AI models is increasing exponentially and can be significant. To address this issue, a new research field called *Green AI* has emerged, *which focuses on developing and deploying AI models that are both effective, efficient, and environmentally sustainable*. In this chapter, we specifically explore how to apply Green AI in practice, i.e. translate it into manageable principles that can be applied by professionals, researchers and students while building AI models. We propose an ontology of five domain-independent Green AI best practices: (i) monitor and report energy consumption as well as CO<sub>2</sub> footprint, (ii) select your model architecture with care, (iii) adapt a data-centric AI approach, (iv) apply a smart hyperparameter tuning strategy, (v) retrain your model with care. We have applied the above practices in two case studies: (i) a societal case focusing on measuring biodiversity, in particular wildflower abundance and richness ‘in the wild’ with an AI-based object detection approach, and (ii) a circular economy scenario focusing on automatic e-waste classification and outlier detection. This first exploration creates confidence that it is worthwhile having a ‘how-to’ guide for AI practitioners to take Green AI into account while building applications.

### 10.1 Introduction

The tremendous progress and fast development of AI in today’s information age enables human beings to benefit from the vast amount of data and realize breakthroughs for the economy, society, and the environment. Domains in which AI have already had a considerable impact include healthcare, finance, energy supply, and transportation. According to Joppa and Herweijer (2020) the usage of AI for environmental applications can contribute up to 4.4% of global GDP and save up to 4.0% of all greenhouse emissions worldwide by 2030.

However, as AI systems are becoming significantly more complex, their carbon footprint can no longer be ignored. Contemporary models typically contain

carbon  
footprint

millions or even billions of parameters and run on clusters of GPUs, sometimes for thousands of hours. Strubell *et al.* (2020) have studied the carbon emissions of state-of-the-art Natural Language Processing (NLP) models and has found that training large NLP models emits an amount of carbon dioxide equivalent to a trans-American flight. Moreover, as a spoiler-alert, the impact of ChatGPT and similar Large Language Models (LLMs) on our climate has been discussed in many recent blogposts.<sup>1</sup> The concept of ‘Green AI’ by Schwartz *et al.* (2020) has become a top priority on the agenda of the AI research community. The goal of Green AI research is to provide concepts and tools to monitor, report, analyse and eventually reduce the carbon footprint of the whole AI lifecycle.

AI lifecycle  
models

AI lifecycle models have many similar variants. An early example, known as named Cross Industry Standard Process for Data Mining (CRISP-DM)<sup>2</sup> is presented by Wirth and Hipp (2000). In such an AI lifecycle, depicted in Figure 10.1, we usually start with a problem definition phase together with business or associated partners. Once the objectives and requirements are defined, data acquisition will take place, including discovering available data sets, improving data quality, and deriving initial insights from the data and perspectives. Model development, training, evaluation, and refinement will be performed in the next phase. While data discovery is usually the most time-consuming process, iterative model development and refinement is the most resource exhaustive phase per actor (typically a data scientist or AI engineer). Often different model architectures need to be trained, compared, and fine-tuned to achieve the best performance. Next, the inference and deployment phase brings the optimized model into production and exposes it as an AI service to the outside world. In this phase the scale of usage is the main determinant of energy consumption. Finally, the stage of performance monitoring measures the accuracy of the trained model over time while being used by many end users or customers in real-world situations. Data drift might for instance lead to less effective models that need to be retrained to maintain a high-performant service. In today’s cloud platforms, tooling solutions, a.k.a. Machine Learning Operations (MLOps), are available to automate the AI life cycle, from streamlining data collection to carefully designing retrain intervals (Heck *et al.*, 2021).

Training and inference of AI are computationally intensive, albeit in different ways. The computational effort of AI models is usually expressed in the hardware-independent metric Floating Point Operations (FLOPs) and typically refers to one forward pass in a neural network (see Annex). Training effort is on a scale with the number of AI experiments (hyperparameter tuning strategy) that are carried out to ‘arrive at’ an accurate model. Inference effort is on a scale with the number of requests or ‘model calls’ by the community of users. Along with the cost

1 <https://www.tudelft.nl/en/stories/articles/sustainable-artificial-intelligence-from-chatgpt-to-green-ai>.

2 Data mining is a process to extract useful data from a larger set of raw data. AI takes this one step further and uses the data to solve cognitive problems associated with human intelligence.

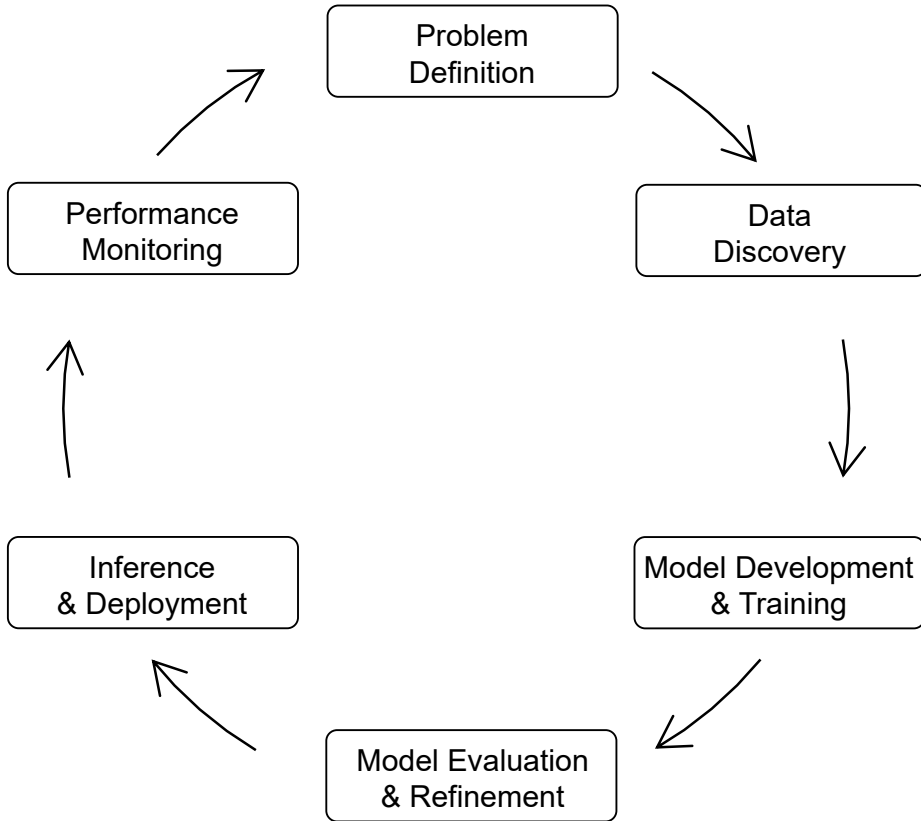


FIGURE 10.1 Typical AI lifecycle model

of hardware and consumed electricity, the carbon emission already expands from tons to several hundred tons (and even more if we consider the fast evolution of the latest LLM models). Developing and utilizing AI technologies in a sustainable way has become a major concern in a world where action is much needed to mitigate climate change. Therefore, carbon emissions of AI throughout its entire lifecycle need to be monitored, reported, and communicated transparently, and methods should be developed to benchmark and assess the AI lifecycle’s environmental impact.

In section 2 we give an overview of the background and related work and briefly sketch the existing energy-related metrics used to assess the cost of the models. Section 3 presents a recipe for how to move from FLOPs to carbon emissions of models. Section 4 discusses an ontology of Green AI best practices and provides guidelines on how to apply them. Section 5 presents two study cases and shows the achievements of applying Green AI best practices to these projects. Finally, section 6 concludes that AI practitioners should be aware of the carbon emissions associated with AI design and experimentation activities and use Green AI best practices to create sustainable solutions.

## 10.2 Background and related work

Metrics for Green AI aim at different aspects. An overview is presented in Table 10.1. FLOPs are used as a hardware-independent measure, with the aim to benchmark foundation models. Energy efficiency relates directly to the energy (in Joules) consumed by a model while running on specific hardware. CO<sub>2</sub> emission has a direct link with the impact on our environment. Note that CO<sub>2</sub> emission takes into account whether traditional fossil-based fuels or alternative green energy sources (e.g. solar, or wind energy) are exploited, and hence depends on the place on Earth where the model is running, or trained for that matter.

The most influential literature on Green AI for the above metrics is briefly discussed below.

TABLE 10.1 Metrics of Green AI

FLOPs	Floating point operations, a number that expresses the complexity (in terms of size or scale) of an algorithmic process
FLOP/s	Floating point operations per second, a unit that relates to the performance (or computational power) of a computer
Energy efficiency	A measurement of how much energy in Joules is consumed per unit (such as FLOP/s)
CO <sub>2</sub> emission	A measurement of total greenhouse gas emitted, expressed in terms of the equivalent measurement of carbon dioxide

### 10.2.1 Evolution of models in terms of FLOP/s

With the fast development of AI in recent years models are getting more complex; consequently, the energy consumption of these models increases exponentially, causing a notable environmental impact. OpenAI did extensive research on the growth trend for computing power in terms of FLOP/s needed for different AI models developed in recent years. According to OpenAI (2018) AI models have doubled the FLOP/s used every 3.4 months since 2012. This trend is depicted in Figure 10.2 on a log scale of petaflops/s-day<sup>3</sup> and has been coined ‘Moore’s Law of the AI era’.

Moore’s Law  
of the AI era

This trend won’t stop with today’s fierce competition and rapid progress in LLMs since ChatGPT (and other generative AI tools) launched so successfully in November 2022. The need to manage the risks of generative AI in a broader sense has been highlighted by researchers, scientists, and industry leaders (see e.g. Bengio, *et al.* 2023).

<sup>3</sup> FLOP/s of performing 10<sup>15</sup> neural net operations per second-day.

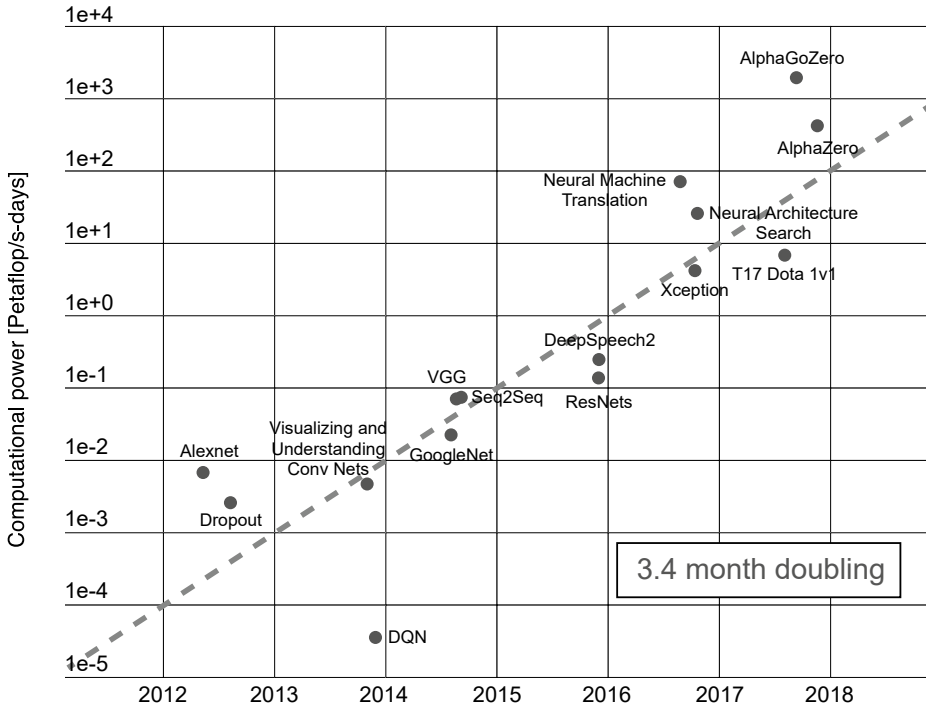


FIGURE 10.2 The computational effort of AI models increases according to Moore's Law of the AI era

10.2.2 Energy efficiency of models

A detailed analysis of the training cost and algorithmic efficiency of AI models including the latest development of vision models was carried out by Hernandez and Brown (2020). They focused on the amount of computation used to train models with the huge ImageNet dataset. Although both algorithm efficiency and hardware efficiency have been improved a lot recently, the growing demand for better performance is still driving the AI research community to build larger and more complex models. Desislavov *et al.* (2021) analysed energy efficiency with a greater number of deep neural networks with various hardware components in a longer time frame. Thompson *et al.* (2022) reported the computational demands of several deep learning applications, showing that progress in them is strongly reliant on increases in computing power.

Compared to training cost, there are a few studies reflecting on the inference cost. Canziani *et al.* (2016) compared accuracy, memory footprint, parameters, operations count, inference time and power consumption of 14 models trained on ImageNet. A similar study by Li *et al.* (2016) measured energy efficiency, Joules per image, for a single forward and backward propagation iteration. This study benchmarked four Convolutional Neural Networks (CNNs) on different hardware architectures such as different CPU and GPU configurations. Both publications analyse model efficiency, but they do this for very concrete cases.

### 10.2.3 Carbon emission of models

Some recent work has started the discussion on carbon emissions and the sustainable usage of AI models. Strubell *et al.* (2020) demonstrated the issue of carbon and energy impacts of training large NLP models by evaluating estimated power usage and carbon emissions for a set of case studies. Their results led to the conclusion that we need to reduce the carbon footprint of developing and running AI models. Schwartz *et al.* (2020) defined the term ‘Green AI’ as ‘AI research that yields novel results while taking into account the computational cost’. Henderson *et al.* (2020) proposed a framework for tracking real-time energy consumption and carbon emissions and create a leaderboard to incentivize energy-efficient research. Platforms such as Huggingface provide tools utilizing CodeCarbon made by Budennyaya *et al.* (2022) to calculate CO<sub>2</sub> emissions when performing training or pre-training.

CodeCarbon

## 10.3 Conversion sequence: from FLOPs to CO<sub>2</sub> footprint

Green AI

The term ‘Green AI’ refers to research that promotes measurement of energy efficiency for algorithms or models as a widely accepted evaluation metric alongside model accuracy. Scientific literature shows that a gamut of metrics related to energy efficiency is available. In this section we list the most used metrics and advocate that we should move from FLOPs to CO<sub>2</sub> footprint.

- i. FLOPs: The number of floating-point operations provides a direct estimation of the amount of work by the computational process (OpenAI, 2018). It is agnostic to the hardware on which the model is run. Most authors typically report the computation required to go through one forward pass of a neural network.
- ii. Number of parameters: This measure is closely correlated with the amount of memory consumed by the model. As a result, different models with a similar number of parameters often perform different amounts of work (Canziani, 2016).
- iii. Elapsed time: Authors report the time to train the whole neural network, or the time needed for inference, i.e. the time it takes to execute one forward pass. This measure is highly influenced by factors such as the underlying hardware, other jobs running on the same machine, and the number of cores used (Jeon and Kim, 2018).
- iv. Energy: The energy consumption of training or inferencing of an AI model can be obtained either by a calculation based on FLOPs, number of parameters, the elapsed real time, or by measurement with power meters on CPUs and GPUs (Desislavov *et al.*, 2021).
- v. Carbon emission: To calculate carbon emission, the carbon intensity, i.e. how many grams of carbon dioxide (CO<sub>2</sub>) are released to produce a kilowatt hour




FIGURE 10.3 Carbon intensity for an assortment of locations

(kWh) of electricity, is collected from the local grid, and used to multiply power estimation of training or inference process. It is essential to know that the local energy grid makes a huge difference when calculating carbon emissions. Figure 10.3 shows the CO<sub>2</sub> intensity (in 2014/2015) for an assortment of cloud-provider regions and energy production methods (Henderson *et al.*, 2020). It is clear that running an AI job in Quebec is much cleaner than the other regions.

### 10.4 Green AI best practices

#### 10.4.1 Monitor and report CO<sub>2</sub>

As AI practitioners we have a responsibility to develop sustainable AI. The first responsibility is to monitor and report the energy consumption as well as the carbon footprint of the AI model during training and inference phases. Several tools or packages have been developed for this purpose, such as CodeCarbon (Budennyya at al., 2022), CarbonTracker (Antony *et al.*, 2020) and ExperimentImpactTracker (Henderson *et al.*, 2020). Those packages have utilized a publicly available framework and can easily be integrated with the development code for tracking both energy consumption and carbon emission.

CodeCarbon  
CarbonTracker  
Experiment  
ImpactTracker

#### 10.4.2 Select your model wisely

It has often been a belief in AI community that to achieve greater accuracy larger and deeper models are needed with huge amounts of data to train them. However,




FIGURE 10.4 Benchmark of CNN models for image classification

larger models come with a higher financial and environmental cost. What's more, it is not always true that larger and deeper models achieve greater accuracy. Some researchers have done extensive benchmarking for the AI tasks, such as imaging classification, with various deep learning neural network models. Image classification is a fundamental task in vision recognition that aims to understand and categorize an image as a whole under a specific label. As shown in Figure 10.4 by the study of Canziani (2016), fourteen image classification models have been compared in a computation vs accuracy graph. In this figure, the x-axis represents the total number of operations in Giga FLOPs. The y-axis is the top-1 accuracy, i.e. the highest probability of correct prediction of the labelled images. Bubble size refers to the number of parameters of each model.

Figure 10.4 shows some remarkable and useful patterns, such as:

- ResNet-101 has doubled the number of FLOPs compared to ResNet-50. Note that this leads only to a minor improvement in accuracy (2%).
- For a top-1 accuracy of approximate 70% you can use several models. Note that ResNet-34 (with even a slightly higher score) has one-sixth of the FLOPs compared to VGG-16, thereby delivering more value with less energy consumption and a lower CO2 footprint.

In summary, it is always good practice to check the model benchmark and pick up the one that is most sufficient in terms of accuracy with cost efficiency for the application. In today's education on AI, the focus is still very much on accuracy and not yet on energy efficiency and carbon footprint. Our message reflects that there should be a careful per-case consideration of the trade-off between accuracy and efficiency/emission.

### 10.4.3 Adopt a data-centric AI approach

The quality of data has a huge impact on how well the whole system works, and how we fuel AI models is crucial to their success. Data-centric AI has been an emerging discipline that systematically deals with data quality to build AI systems. In Figure 10.5, we show as an educational example trend lines for noisy data and clean data. Note that, in order to reach the same level of performance, the amount of high-quality data is only a small fraction of the amount of low-quality noisy data. This practice is also strongly supported by AI pioneers, see for instance

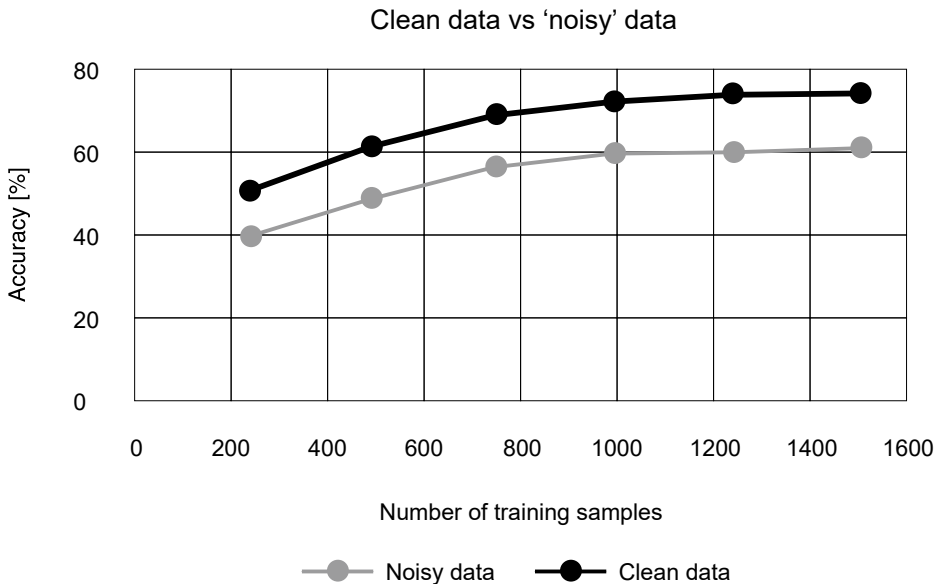


FIGURE 10.5 Clean data improves prediction accuracy

‘Andrew Ng, AI minimalist: the machine-learning pioneer says small is the new big’ (Strickland, 2022).

It is a critical and essential process to carefully check the dataset, remove data noise and balance the data as much as possible. We have also practised data-centric AI in our projects, and have obtained significant accuracy improvement and timing reduction, and consequently obtained significant energy and emission savings.

#### 10.4.4 *Tune your hyperparameters in a smart way*

Model training is a process through which a model learns its parameters (often billions). Besides this, every model also has hyperparameters (usually a few) that it cannot *learn*, but can be *tuned* for. The process of tuning hyperparameter values is called hyperparameter tuning. It is a vital aspect of increasing model performance.

There are several hyperparameter tuning strategies, such as grid search, random search and Bayesian search or optimization.

- Grid Search – A grid or set of hyperparameters are defined and every possible combination is used for training a model. This is exhaustive and computationally expensive and is used when the hyperparameter search space is restricted.
- Random Search – Instead of a grid, statistical distribution of each hyperparameter is provided and the number of iterations can be controlled. This is a suitable strategy for larger search spaces.
- Bayesian Optimization – A sequential model-based optimization that uses the results from previous iterations to decide the next hyperparameter value candidates.

Bayesian optimization methods are more efficient because they select hyperparameters in an informed manner. By prioritizing hyperparameters that appear more promising from past results, Bayesian methods can find the best hyperparameters in less time (in fewer iterations!) than both grid search and random search. Therefore, it is preferable to use a Bayesian optimization strategy for hyperparameter tuning.

#### 10.4.5 *Retrain with care*

Once the model is trained and put into production for a while, the model might need to be retrained with the following observations.

- The model’s performance metrics have deteriorated.
- The distribution of the prediction is different from those observed during training.
- The training data and the live data diverge, that is, the training data is no longer a good representation of the real world.

However, retrain is a very costly process, and has a financial, operational and environmental impact. Retrain intervals are closely associated with business cases and should be carefully analysed and determined to reduce the total cost including the emission and the environmental impact.

## 10.5 Case studies

We embrace 'AI for good' cases in various applied research projects. The first example we present is about biodiversity loss in the Netherlands. This is a major environmental issue, caused by e.g. habitat loss, intensive farming or pollution. Monitoring and classifying wild flowering plants with AI will help us better understand the changes in the biodiversity and react to it. The second example is about building a more circular economy with the use of AI. In recent decades tons of E-waste (electronic waste) have been produced worldwide. The de-assembly and recycling process is widely enforced but is still largely a manual process and could be greatly improved with AI technology.

### 10.5.1 *Automatic wildflower monitoring*

To better understand biodiversity, it is necessary to initiate and automate large-scale monitoring programs supported by AI technology. Monitoring means *identifying* and *counting* objects of interest. In this case study we applied object detection algorithms for monitoring flowering plants 'in the wild' (Heck and Schouten, 2023; Schouten *et al.*, 2024). Wildflowers are an essential component of biodiversity. They provide many eco-system services, such as medicine, building materials and food; they keep our soil healthy, purify water and mitigate climate change to a large extent by absorbing greenhouse gases and significantly lowering temperatures in cities. To this end, a unique expert-annotated reference dataset with over 2000 high resolution images, each covering approximately 1m<sup>2</sup> of soil, has been collected around the city of Eindhoven. This Eindhoven Wildflower Dataset (EWD) holds 160 flowering plant species and contains images of roadsides, rich-weed grasslands, marshland, and urban green areas (Schouten *et al.*, 2024). As with many biological datasets that are collected 'in the wild', EWD has a long-tailed distribution. Common species are overrepresented and rare or inconspicuous species are underrepresented.

We selected a state-of-the-art R-CNN object detection algorithm (Ren *et al.*, 2015) with a Resnet50 backbone algorithm from the PyTorch library and trained it with EWD images. During the AI experimentation phase, we found out that the mean average precision (mAP) varied considerably over the species and is affected to a large degree by the skewed distribution. By carefully pre-processing the image data and creating a balanced subset from the original long-tailed EWD dataset, the training time was reduced from 10 hours to 35 minutes, thereby greatly reducing the energy consumption (and CO<sub>2</sub> emission) of the training process, while the accuracy (mean average precision) even increased from 0.68 to 0.82.

This case demonstrates the importance of a data-centric AI approach. Making sure to train models with high-quality data is a major factor in achieving Green AI.




FIGURE 10.6 Clustering the E-waste dataset into device groups

### 10.5.2 *Sorting E-waste for disassembly*

This circular economy project focuses on using AI, and in particular a cluster algorithm, to identify E-waste devices with removable batteries.

Batteries, especially lithium batteries, need to be carefully recycled because they can catch fire with the slightest damage. While new-generation phones and tablets are designed with non-removable batteries to make the product tighter, slimmer and waterproof, a small proportion of old phones and gadgets still contain removable batteries. Currently, during the recycling process all E-waste devices are mixed, and then manually sorted on the conveyor belt. To distinguish devices that contain removable batteries from those that don't is a time-consuming and highly unsafe job for local workers. AI technology has the potential to improve this process. To achieve this, a dataset is collected, containing E-waste devices spanning from smart-watches to phones and tablets, with their physical dimensions as well as images.

The brute-force approach is to train an advanced image classification algorithm, such as a CNN (Convolutional Neural Network), and sort the devices automatically. A smarter approach is to use non-visual clues, in particular the dimensions (width and height) of the devices. By checking the data carefully, it was found that the 'outliers' correspond to rarely used old phones with removable batteries. An additional insight is that even if the device type is not known beforehand, it is common knowledge that smart watches, phones and tablets are intrinsically of different sizes. By

applying the DBSCAN clustering algorithm to this dataset, the data can be organized into separate groups corresponding to device types, as shown in Figure 10.6. This data pre-processing and cluster approach provides strong support for automatic decision-making for E-waste sorting.

This case study clearly illustrates that a lot can be gained by picking the right AI approach. An obvious and straightforward solution for sorting or inspecting products is to apply computationally intensive vision AI. Here we demonstrate that a lightweight cluster model also works well. Hence, rethinking the obvious solution can make an excellent contribution to Green AI.

## 10.6 Conclusions

The digital transformation of society means that more and more data are being collected and fed into AI models. Model architectures are becoming larger and more complex, and at the same time more iterations for optimization are being executed to achieve the best performance. As a consequence, the calculations executed on computers incur a heavy energy and financial cost, and the CO<sub>2</sub> emissions of AI models are no longer negligible. As AI practitioners, educators, and researchers, we need to be aware of the CO<sub>2</sub> footprint of our AI design and experimentation activities. Specifically, we should inform ourselves and others about both the positive effects and negative environmental consequences of using AI. Therefore, we should be able to measure and explicitly state the emissions of models, and finally we should take steps to reduce the carbon emissions during the development and deployment of AI solutions.

### **Annex: Explaining training and inference with an artificial neural network**

To train a model, researchers usually pre-process the data, define a model architecture, optimizing strategy and loss function, and then take an iterative process to loop over the model network with the number of epochs. A neural network is made of multiple neurons and these neurons are stacked into layers. The connections between the layers occurred through the parameters of the network. Calculations are carried out in forward propagation and backward propagation through each layer in the network until the loss function is minimized.

Unlike training, inference doesn't re-evaluate the layers and parameters of the neural network. Inference applies the trained neural network model or a forward pass to a new unknown dataset, and outputs the prediction based on the accuracy of the model. Minimizing latency issues during the inference process can pose a challenge for getting the system to make decisions in real time.

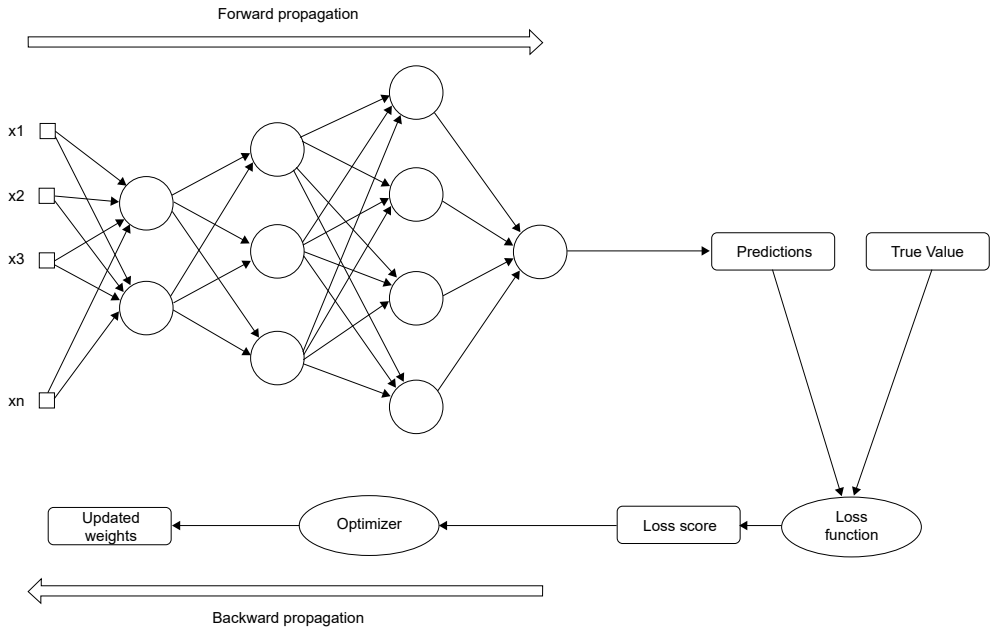


FIGURE 10.7 Training process of a neural network

## References

- Anthony, L.F.W., Kanding, B. and Selvan, R., 2020. Carbontracker: tracking and predicting the carbon footprint of training deep learning models. arXiv preprint arXiv:2007.03051. <https://github.com/lfwa/carbontracker>.
- Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., Harari, Y.N., ... and Mindermann, S., 2023. Managing AI risks in an era of rapid progress. arXiv preprint arXiv:2310.17688.
- Budenny, S.A., Lazarev, V.D., Zakharenko, N.N., Korovin, A.N., Plosskaya, O.A., Dimitrov, D.V.E., ... and Zhukov, L.E.E., 2022. Eco2AI: carbon emissions tracking of machine learning models as the first step towards sustainable AI. In: *Doklady Mathematics* (Vol. 106, No. Suppl 1, pp. S118–S128). Pleiades Publishing, Moscow, Russia. <https://github.com/mlco2/codecarbon>.
- Canziani, A., Paszke, A. and Culurciello, E., 2016. An analysis of deep neural network models for practical applications. arXiv preprint arXiv:1605.07678.
- Desislavov, R., Martínez-Plumed, F. and Hernández-Orallo, J., 2021. Compute and energy consumption trends in deep learning inference. arXiv preprint arXiv:2109.05472.
- Heck, P., Schouten, G. and Cruz, L., 2021. A software engineering perspective on building production-ready machine learning systems. In: *Handbook of research on applied data science and artificial intelligence in business and industry*. IGI Global, Hershey, PA, USA, pp. 23–54.

- Heck, P. and Schouten, G., 2023. Defining quality requirements for a trustworthy AI wildflower monitoring platform. In: IEEE/ACM 2nd International Conference on AI Engineering–Software Engineering for AI (CAIN), pp. 119–126.
- Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D. and Pineau, J., 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of machine learning research*, 21:1–43.
- Hernandez, D. and Brown, T.B., 2020. Measuring the algorithmic efficiency of neural networks. arXiv preprint arXiv:2005.04305.
- Jeon, Y. and Kim, J., 2018. Constructing fast network through deconstruction of convolution. *Advances in neural information processing systems*, 31.
- Joppa, L. and Herweijer, C., 2020. How AI can enable a sustainable future. Microsoft Corporation. <https://news.microsoft.com/wp-content/uploads/prod/sites/53/2019/04/PwC-Executive-Summary.pdf>.
- Li, D., Chen, X., Becchi, M. and Zong, Z., 2016. Evaluating the energy efficiency of deep convolutional neural networks on CPUs and GPUs. In: IEEE international conferences on big data and cloud computing (BDCloud), social computing and networking (SocialCom), sustainable computing and communications (SustainCom), pp. 477–484.
- OpenAI, 2018. AI and compute. <https://openai.com/research/ai-and-compute>.
- Ren, S., He, K., Girshick, R. and Sun, J., 2016. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39:1137–1149.
- Schouten, G., Michielsen, B.S.H.T. and Gravendeel, B., 2024. Data-centric AI approach for automated wildflower monitoring. *PloS One*, 19:e0302958.
- Schwartz, R., Dodge, J., Smith, N.A. and Etzioni, O., 2020. Green AI. *Communications of the ACM*, 63:54–63.
- Strickland, E., 2022. Andrew Ng, AI minimalist: the machine-learning pioneer says small is the new big. *IEEE spectrum*, 59:22–50.
- Strubell, E., Ganesh, A. and McCallum, A., 2020. Energy and policy considerations for modern deep learning research. In: Proceedings of the AAAI conference on artificial intelligence, 34:13693–13696.
- Thompson, N.C., Greenewald, K., Lee, K. and Manso, G.F., 2020. The computational limits of deep learning. arXiv preprint arXiv:2007.05558.
- Wirth, R. and Hipp, J., 2000. CRISP-DM: towards a standard process model for data mining. In: Proceedings of the 4th international conference on the practical applications of knowledge discovery and data mining, 1:29–39.