Assessment of carbon reduction through AI methods in inspection after reverse logistics

Marian Schlüter[1§*], Justus Caspers[2*], Clemens Briese[1*], Robert Schimanek[3*], Paul Koch [1*], Ole Kröger[1], Pinar Bilge[3], Jörg Krüger[1], Franz Dietrich[3]

[1]Department of Machine Vision, Fraunhofer Institute for Production Systems and Design Technology, Berlin, Germany
[2]Chair of Sustainable Engineering, Technische Universität Berlin, Berlin, Germany.
[3]Chair of Handling and Assembly Technology Research, Technische Universität Berlin, Berlin, Germany

*These authors contributed equally to this work
[§]Corresponding author

Email addresses:
IPK: marian.schlueter, Clemens.briese, paul.koch, ole.kroeger, joerg.krueger @ipk.fraunhofer.de
TU-SEE: j.caspers@tu-berlin.de
TU-HAMSTER: r.schimanek, p.bilge, f.dietrich@tu-berlin.de

Abstract

Although it is a goal to reduce the workforce and energy consumption required to digitize and learn Artificial Intelligence (AI) methods, it is vital to assess the balance of energy consumption versus the efficiency benefits through AI methods for reverse logistics sorting processes. Therefore, the aim of the paper is to describe how to find the return of investment in energy used for digitization and artificial intelligence. At first, energy consumption drivers introduced to the sorting process for digitization are assessed. In addition, the energy consumption of the training of AI methods and other sophisticated statistical models was monitored. Afterward, the authors discuss and select which metrics suit the comparison between energy costs of algorithms and show the influence of design decisions in artificial intelligence. An analysis shows that it is possible to achieve high recognition performance but keep energy consumption in mind.

Finally, it is shown that in an inspection use case, we can achieve a fast break-even of energy use consumption for digitization and AI used for improved processes is achievable. The demonstration compares the energy costs of more properly sorted remanufactured products compared to the opportunity costs of energy consumption for a new production of respective products.

Introduction

The circular economy is increasingly in focus in Europe and worldwide to provide sustainable answers to the problems of the consumer society as e.g. the Circular Economy Action Plan of the European Union states. [1] After repair, refurbishment, or remanufacturing, many products should be used in further life cycles. However, the processes in the circular economy are subject to many uncertainties. It is not predictable when and for what reasons a product leaves its first life phase and enters the reflux part of the cycle. [2] Static planning or classical estimations often reach their limits in reverse logistics. Therefore, in recent years, there have been increasing attempts to counter this uncertainty with the help of artificial intelligence.[3]

In particular, reverse logistics [4] and the handling of used parts are in the spotlight here, as products come back in all shapes, colors, variants, and states. Automation is only possible to a limited extent at the current state of the art and is mostly not economically viable. However, assistance systems with artificial intelligence can support employees in this process and thus optimize processes. These systems are often referred to as augmented intelligence, in which AI systems enhance the expertise of employees.

How AI can improve reverse logistics – EIBA project

In the research project EIBA an AI-based recognition algorithm for automotive components (cores) at the sorting stage for remanufacturing is being developed based on images, reverse logistics business data and user interaction. Therefore, the typical core selection workstations are enhanced with sensors, such as RGB-D cameras and a scale. The gathered information contributes to identifying the core's product code (class) independently from attached and often missing or not readable numbers or identification codes. The worker only has to place a core in a certain workbench range. Then, core images and mass are automatically acquired and fed to a learning system that predicts the core's class. The worker can confirm the prediction or disagree and pick the correct class from a prediction visualized list of cores.

In 2018 Schlüter et al. [5] have already shown that machine vision and learning approaches, particularly deep learning techniques, can recognize used cores in combination with an employee. The performance here was significantly increased by suggestions of an artificial intelligence. 2021 Schlüter et al. have shown in a succeeding work that the recognition performance could be further increased with the help of other methods and further data (delivery and customer data). [6]

AI algorithms and digitization in themselves bring their own CO_2 or energy footprint. For example, additional sensors are often used for data collection, which would not be used in a conventional process. Learning an AI is generally very energy-intensive and much discussed in the scientific community. [7] For example, the training for GPT-2, a general-purpose AI model by OpenAI and Microsoft which yielded significant progress, produced 313,000 kg of CO_2 emissions. [8] To be fair, GPT-2 is a very energy-intensive model which is, for most tasks, well oversized. But still, this raises the question of whether CO_2eq savings from AI for the circular economy are offset by the additional energy required to create the methods.

This work aims to investigate how much CO_2eq our specific AI model combinations produce and how much CO_2 eq emissions we can avoid because of prolonging life cycles of car parts. This paper considers a sub-set of cores of the EIBA project due to data availability and comparability to previous works. At the time of conducting the study, sufficient images for 511 and business data for 2,082 different starters were available. Additionally, there has been access to life cycle assessments of 4 different starter types. From that, we can get an indication of the potential impact of AI-enhanced reverse logistics. The aim of this paper is to calculate the break-even point of CO_2 eq emissions produced by the AI for a subset of automotive exchange parts. We focus on starters which we already have discussed in previous works.

Methods

This chapter analyses state-of-the-art methods in determining reductions and costs necessary for an energy consumption assessment of inspection processes after reverse logistics. First, the available methods for the assessment of gains coming with an enhanced inspection of cores are discussed. Second, different methods are discussed to estimate and measure energy consumption.

Impact on Climate of AI use in reverse logistic

From a climate perspective, the use of AI in reverse logistics has positive and negative effects. On the positive side, an increase of cores, which are brought back into the value chain, leads to decreased carbon emissions. The decrease can be seen as the remanufacturing of used cores has lower potential carbon emissions than the new production of the same or similar core. For example, a study has shown that the remanufacturing of a starter can decrease carbon dioxide equivalent (CO_2 eq) by 52 % [9].

On the negative side, related to climate effects, the energy demand and necessary hardware for calculation and application of the AI need to be considered. The AI can be implemented on-premise as well as in the cloud. The implementation in the cloud could lead to overall smaller CO_2eq emissions from the hardware production. This can be argued as the on-premise hardware is, in contrast to hardware operating the cloud, probably not used only for one specific AI task. This means that the environmental burden of the production of the cloud hardware can be shared between the different tasks it performs. As for the EIBA Project, on-premise hardware was bought, the potential savings of cloud solutions aren't considered as part of the paper. In addition, we assumed that the measurement of the energy consumption for on-premise hardware would be more valid than for off-premise. Further, in the first step, only the consumed energy is considered. This simplification is based on the results of different life cycle assessment studies of servers, which have shown that over 90 % of the potential CO_2 eq emissions were generated in the life cycle of the server or, more specifically, by the electricity consumption during the operation of the server. [10]

For the assessment of the energy consumption of the AI, it must be defined which steps of the AI compiling needs to be considered. As the AI used in the EIBA Project is pre-trained, it needs to be discussed if and to what extent the energy consumption of this training has to be included for the given example. The principles of life cycle assessment (LCA), which are defined in ISO 14040/14044, can help to decide. In theory, all contributors to the potential CO_2 eq emissions need to be assets, starting with the hardware pre-trained the AI, following the machines producing the hardware and each tool used to produce these machines, and so on. This would lead to a much too extended assessment, which is not feasible. This problem leads LCA to the principle of cut-off criteria, which are defined to intentionally exclude processes from studies based on a specification of the quantity of material or energy flow or the level of environmental significance. For example, if a process contributes less than 1 % of the total energy consumption it should be excluded from a specific study. Based on this principle, the pretraining of the AI is excluded from the assessment in this paper. This can be argued as the pre-trained AI is not only used for the EIBA project but probably also for a large number of other projects. It can be assumed that dividing the energy consumption, needed for the pretraining of the AI, over all these projects would lead to a marginal contribution to the total energy consumption in each of the projects.

This, in turn, leaves the data acquisition, AI training, and AI inference as contributions to the total energy consumption of an AI-enhanced reverse logistic process. The data acquisition is a continuous phase, starting before the rollout of any AI methods. Therefore, any energy consumption regarding the data acquisition needs to be monitored before and during the AI rollout. The AI training is an energy-expensive event-based process, which might occur after initial training with some frequency given new data. The energy consumption can be monitored within the timeframe of a training event, excluding the hardware costs as it is spread over several other projects – similar to the pretraining step. The AI inference is also an event-based process but runs at a higher frequency. Here some dedicated hardware is required to continuously host some inference service. Hence, the AI inference energy consumption is a sum of the hardware costs with respect to energy consumption, the baseline energy consumption of the hardware, and the AI inference energy consumption multiplied with some expected frequency. However, the AI inference is hosted on pre-existent hardware, which eliminates the hardware costs of the AI inference from the equation.

Based on the measured energy consumption, the CO_2eq can be calculated with an emission factor of a specific grid mix. As the project is based in Germany, the German grid mix is used. The probable change in the German grid mix will be considered with scenarios with a high share of renewables to make the result more robust.

Methods for the measurement of energy consumption of algorithms

One major goal of AI-enhanced methods is the reduction of emissions by core inspection. In order to assess the reduction, test, training, and inference costs of emissions must be determined. There are several methods for measuring and estimating the cost of emissions, for instance, by assessing energy consumption. Such methods are based on the computation effort, computation time, or power consumption. Sample methods for estimating the energy consumption and then the cost of emissions are explained in more detail below.

The drawn energy can be estimated by the **computer-independent computational effort** [11]. The respective computational effort can be expressed, for example, as multiplication and accumulation operations (MAC), or floating-point operations per second (FLOPS). Knowing the energy consumption or carbon footprint of the computer used for a FLOP or MAC (e.g., the Power-to-FLOP ratio) provides an estimate of consumption [11]. FLOPs are very well suited for comparing the computing efficiency of algorithms. The disadvantage of computer-independent methods for determining the computational effort is that they require a high degree of adaptation to the software and hardware used in each case. Therefore, other methods of energy measurement are used by the authors.

The drawn energy can be estimated by the duration of the different program parts by means of **time measurement**. With the utilization of the graphics processing unit (GPU), Dynamic Random Access Memory (DRAM), and the central processing unit (CPU) involved, the energy consumed for a computing task can then be determined. With the maximal power draw for the specific GPU denoted as P_G, DRAM denoted as P_D, CPU denoted as P_C [8]. With the power usage effectiveness (PUE) of the system, the ratio of a hardware or data center's overall energy usage to the energy provided to computing equipment, and a measurement of the full-throttle calculation time Δt, the consumed Energy E in Eq.(1) can be estimated [11]. In this paper, the authors use time measurements to estimate energy consumption.

$$\mathrm{E} = \mathrm{PUE} \cdot \Delta \mathrm{t}\ (\mathrm{P}_G + \mathrm{P}_D + \mathrm{P}_C) \tag{1}$$

The entire energy used for specific computer calculations may be assessed by an **external power** consumption **measurement (EPCM)** by calculating time. To do this, the voltage drop across the computer is measured with a voltmeter and the current consumed with an ammeter. The electrical power can be calculated from the voltage drop and the current. The electrical power can be accumulated and integrated over time or, if necessary, converted to energy consumption by measuring the duration of the computational calculations. This method has the advantage that all other (internal) consumptions of the computer are also recorded. Moreover, the method is very flexible in terms of the software and hardware used. In the context of the paper, the authors use the method of EPCM.

The various methods and their advantages and disadvantages are characterized by a variety of different computing hardware and software. Hardware, in particular, changes very quickly. Thus, a computer becomes obsolete after just a few years and is no longer state of the art. From that point of view, it makes sense to measure the energy consumption hardware independently. However, at this point, there is no objective software that would provide the same results for the calculations on other hardware for the hardware-independent measurements. In contrast, the external energy measurement is easy to imitate. Thus, it is worth considering using common hardware to compare the software on similar machines.

However, assuming a worst-case 100 % power consumption of the individual hardware components during the power estimate (see Eq. (1)) might end up far from the truth, which could lead to a false conclusion. Since the hardware used in this project is running on-premise, one can measure the power consumption directly from the external sensors. In order to avoid any wrong conclusions, the authors indicate the direct power consumption alongside a worst-case power consumption estimate.

Experiments & Results

The experiments of this paper are bound to the data provided by the EIBA project, which accompanies the work of the authors. Schlüter et al. 2021 used a subset of 1,440 digitalized parts from the EIBA project to experiment on the benefits of AI-enhanced reverse logistics. The 1,440 parts consist of two different product groups, namely automotive alternators and starter motors, which might be a unique use case for the EIBA. However, computing potential carbon footprint reduction using an AI-methods based on some product mix does not generalize to the complete reverse logistic industry. Therefore, the authors opted to limit the scope of this paper to a single parts group in order to provide a first impression of the potential C02eq saving an AI-enhanced reverse logistic might generate within a specific use case. Therefore, the part group of automotive starters has been selected as the basis for the experiments of this paper.

Experimental Setup

For the experiments, three different, locally distributed PCs were used for specific tasks. The training of the Image-AI is performed on a separate high-performance Linux machine. This machine uses GPUs to process a high amount of data and compute the complex matrix multiplications during the training of the Image-AI. The preparation, training, and inference of the business data evaluation (BDE) are made on a common Windows Laptop. A third Windows PC is employed to execute the image-AI on the workstation. Moreover, the PC is equipped with Real-Sense Cameras for Image Acquisition. All computers are equipped with a common monitor, keyboard, and mouse (with a deactivated monitor of the BDE Laptop). The most important specifications and the energy consumption of the individual components are shown in **Table 1** according to their data sheets.

The algorithms and processing logics are encapsulated in Docker containers on the application system for Image-AI. This creates some additional consumption but is closer to a real application. The direct power consumption is measured with the following shelly lugs Shelly Plus 1 PM (SP1PM) and Shelly Plug S (SPS), see Fig 1. The shelly plugs provide the current power consumption in watts via an HTTP request with a sample rate of 1s. Thereby one can approximate the statistical mean power consumption during some process, given the process is monitored long enough in order to statistically eliminate errors introduced by the sampling rate, and under the assumption that no subprocess is synchronized with the sampling rate.

Table 1 - Specifications of the PCs used and the nominal consumption of the components

Modul	Image-AI Training PC		BDE Laptop		AI Inference PC	
	Desc.	Con. (data sheet)	Desc.	Cons. (data sheet)	Desc.	Cons. (data sheet)
CPU	AMD Ryzen TR 3960X (3.8 GHz)	105 W TDP	Intel i7-6700HQ (3.1 GHz)	45 W TDP	Intel i7-5557U (3,1 GHz)	28 W TDP
GPU	2 x Geforce RTX3090 24GB	2 x 350 W TDP	Not used	-	Not used	-
DRAM	128 GB	48 W (estimated)	16GB (2.6 GHz)	2 W (estimated[1])	8GB DDR3 (1,6 GHz)	3 W (estimated[2])
Space	2 x SSD Samsung 980m Pro	2 x 7 W[3] (average)	SDD Samsung 850 EVO	2.4 W	SDD Samsung 850 EVO	2.4 W
Data acquisition	-	-	-	-	Real Sense D435	< 3.5 W (streaming mode)
O.S.	Ubuntu 20.04	176.94 W (idle, measured)	Windows 10 64bit	45.3 W (idle, measured)	Windows 10 64bit	12.4 W (idle, measured)

The following subsections are separately investigating the energy consumption of the AI training and inference for both a BDE-AI and an Image-AI. Each section describes an experiment that measures the directly consumed energy using the shelly plugs as well as indicates the estimated power consumption calculated by Eq. (1) with respect to the hardware specifications of **Table 1**.

Business Data Evaluation

Business Data Evaluation (BDE) sifts through historical data statistically. The statistical component is based on features that differentiate vehicle cores from product returns. Core features are gathered using sensors to predict the core's class, product group, and condition. The core classes depend on the supplier, the packaging, and the mass. Suppliers, for example, ship typical cores. The European Article Number (EAN) printed on the packaging can be used to determine whether a core relates to an OEM or aftermarket program. The mass of each core varies according to its dimension and, as a result, the product group and class. An example of a core is the starter motor used to start the engines of motor vehicles. In a sample station, the authors inspect 62,757 starters per year, approximately 1,206 starters per week, and

[1]https://www.micron.com/-/media/client/global/documents/products/technical-note/dram/tn4007_ddr4_power_calculation.pdf
[2]https://www.crucial.de/support/articles-faq-memory/how-much-power-does-memory-use#:~:text=Als%20Faustregel%20gilt%20jedoch%2C%20dass,die%20XMP%2DEinstellungen%20hinaus%20%C3%BCbertakten.
[3] https://semiconductor.samsung.com/consumer-storage/internal-ssd/980pro/

consequently 252 starters per day. One manual inspection process takes, on average, approximately 92 seconds, which includes customer registration, identification, damage evaluation, and intralogistics.

From the available data, the useful data reduces to 58,500 starter inspections with 2,280 different starter core classes and a median frequency of 2, as listed in Table 2. The 500 most frequent starters account for 85 % of the total number of starter inspections performed over the year 2021. In terms of the useful 58,500 records, the 500 most frequent starters account for 93 % of all usable inspection data.

Table 2 - Frequency of starter classes in 2021 at the inspection station of the research partner

Frequency of 2,082 starter classes in 2021									
Share	Ave.	Std. Dev.	Min.	Max.	25 %	50 %	75 %	Med.	#
All	24	149	1	5,480	1	2	9	2	58,500
TOP500	100	306	11	5,480	19	32	72	32	53,455

The energy impact and quality of a BDE are assessed using a cross-validation prediction accuracy study. The useful starter core data set is randomly split into 80 % training and 20 % test data 10 times by an algorithm. The training data prepare the predictions and represent a historical data set. The test data serve as samples for the generation of prediction lists. After the generation of the prediction, an algorithm checks each test sample's rank and prediction score appearing in the prediction list. 10 test runs are initiated, and the results are computed as the average of the 10 individual tests' results. The BDE class prediction delivers an average rank of 3.2 at an average score of 52 % with a TOP 1 accuracy of 69 %.

The business data evaluation predicts the class of cores and those of the starters in this study. The implemented method can create and present predictions for the identification and verification of cores during inspection. Although the simple algorithms manage to make predictions, it is not enough for an industrial application. It would be desirable if the minimum score of the test sample is 51 % on the respective prediction list, which corresponds to a 100 % occurrence on TOP 1.

The computational calculations were performed on a laptop with an Intel Core i7 6700HQ CPU (thermal design power of 45 W) and 16 GB of DRAM with an estimated power consumption of 2 W and an estimated PUE of 1. Within the calculations, the start and end times of the program parts were recorded by the CPU [i]so that the performance and energy consumption of the calculations could be measured and analyzed. A Shelly Plus 1 PM (SP1PM) was used for the external power measurement, see Figure. 1, which takes the external measurements described earlier. The SP1PM can be addressed in the network with an HTTP Get request[4]. The answer returns the current power consumption. In this way, the power consumption can be recorded and evaluated. For example, the system's baseline was measured with an average power of 45.3 W, while the system's highest measured power was 100 W.

[4] Shelly https://shelly.cloud/ last accessed 12.02.2022

Figure 1 - External power measurement means: Shelly Plus 1 PM (left), Shelly Plug S (right) (Reference: Shelly https://shelly.cloud/ last accessed 12.02.2022)

The measurements for the energy cost study are divided into preprocessing, training, and inference, as displayed in Table 3. It is noticeable that the inference consumes the largest time portion of the computation time and thus energy. It becomes evident that the measurement leads to higher energy usage in the computations. On the other hand, the measurements are partly in the order of magnitude of the estimated time of the calculation elements. The inference is assessed by an external measurement with a larger consumption. This might be owing, for example, to the hard disk's increased load and other components not acknowledged in Eq. (1).

Table 3 - Comparative energy consumption cost study for the Business Data Evaluation

		Energy in kWh			
Sensor	#	Preparation	Training	Inference	Total
SP1PM/CPU	1	4.9E-10	2.1E-8	2.0E-5	2.1E-5
SP1PM/CPU	500	2.4E-7	1.0E-5	1.0E-2	1.0E-2
Eq. (1)/CPU	1	3.0E-10	1.0E-8	9,6E-6	9.6E-6
Eq. (1)/CPU	500	1.5E-7	5.0E-5	4.8E-3	4.8E-3

Image-AI

The analyses of the business data do show that the 500 most frequent starter classes are making 93 % of the total starters being returned. Hence, one could argue that using the full span of 2,280 starter classes for the training of an Image-AI would make the problem unnecessary hard. Hereby, one would risk an overall low classification accuracy in exchange for a relatively low increase of potential starter classes. Therefore, this paper limits the Image-AI to a subset of 511 high runner starter classes. Moreover, does the high runner subset yield a subjectively sufficient amount of repetitive image data (48 per starter class) for the training, which is not the case for rarely seen starter classes. Hence, a training data set of 24,528 images is used for the Image-AI.

The dataset is split into 60 % training data, 20 % validation data in order to find the best model state during training, and 20 % test data to confirm the validation score. Following the work of Schlüter et al. 2021, a commonly used ResNet-50 architecture pre-trained on ImageNet is used as the AI model. Using a training framework developed at Fraunhofer IPK, no iterative parameter selection had to be performed. The framework contains a heuristic standard set of parameters for certain clusters of problems (such as image classification). Hereby, no hyperparameter tuning had to be conducted. Thus, the outcome is not particularly optimized but also does not rely on expert knowledge nor on additional training intervals. The trained model scored during the test a 98.99 % Top 1 accuracy. The direct power consumption of the Image-AI-Training Machine was monitored during the training via the Shelly plug S. Thereby, a total energy consumption of 10.9 kWh was measured over the duration of 22.73 hours with a mean power consumption of 486.6 W. The training peaked at a maximum of 890.2 W and had a minimum power consumption of 122.5 W. With a standard deviation of 271.5 W the training of the Image-AI is fairly inconsistent, which is due to the GPU downtimes during the data preparation processes in the pipeline. The power consumption baseline of the machine is measured to be at 177.0 W. Thus, one could argue that the training pipeline consumed an additional amount of 309.7 W. Using Eq. (1) and the numbers provided in **Table 1** to compute an estimated power consumption of the training with a PUE of 1, an energy consumption of 19.1 kWh would be expected, which is about a 75 % increase with respect to the measured energy consumption.

The energy consumption of the Image-AI during inference is measured on full replicate of a EIBA workstation. In order to simulate approximate an average inference cost the workstation is monitored for 30 min within 100 inferences were conducted with the trained Image-AI. Figure 1 shows the measurement over the 100 simulated inferences and a simple moving average (SMA) over a period of ten measurement points. The power fluctuations are clearly visible. Due to the pause intervals and the short processing time, the SMA is in the idle consumption range of about 12 W. Due to the sampling rate of 1s of the shelly plug, the energy consumption of the Image-AI inference could not be sampled consistently. Therefore, a higher sampling rate or longer time period would be needed to approximate the energy cost per inference. However, due to the relatively low intensity of the peaks with less than 30 W across all measurements, the authors opt to calculate the total energy consumption of AI-Inference with the estimated worst-case power consumption provided by Eq.(1). Here the AI-Image inference system would consume 51.1 W, as stated by Cunningham [12]. However, the system would consume an additional amount of energy for the data acquisition by the three Intel Real Sense cameras. Each camera has a maximum power consumption of 3.5 W, which would add an amount of 0.01 kW per hour. Thereby, the total worst-case power consumption of the Image-AI inference system is 0.06 kW per hour. The experiment has shown that 100 inferences can be performed in 30 min with ease. Therefore, one can expect the Image-AI inference system to be capable of handling the incoming amount of inference requests raised by a single workstation.

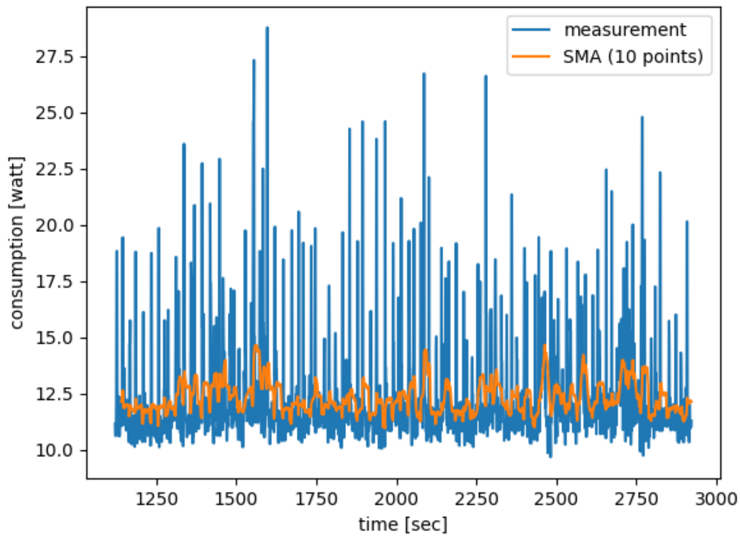

Figure 2 - Results of the Image-AI inference experiment (blue) and the simple moving average over a period of ten points (orange)

Table 4 shows the combined result of the BDE and the Image-AI Experiments for both methods of the measurement of energy consumption of algorithms. The training of the Image-AI consumed by far the most energy.

Table 4 - Results for each Experiment for analyzed 511 Cores

		Energy in kWh		
Method	Sensor	Preparation	Training	Inference
BDE	SP1PM	2.4E-7	1.0E-5	**1.0E-2**
BDE	Eq.(1)	1.5E-7	5.0E-5	**4.8E-3**
Image-AI	SP1PM	2.4E-5	**1.1E1**	1.4E-5
Image-AI	Eq.(1)	2.4E-5	**1.9E1**	1.7E-5

Conversion to CO2 equivalents and savings potential

To show what impact on the climate the introduction of the AI would have, the yearly turnover of starters was used as a basis. As the workstation has an energy consumption of 0.038 kWh and 39 cores can be sorted per hour, the sorting of 62,757 cores, which was the turnover of starters in one sorting facility in 2021, would lead to an energy consumption of 65.3 kWh. The AI is assumed to have a monthly retraining with new gathered data, which would lead to a total of 12 training cycles. As every training cycle has an energy consumption of 10.9 kWh the yearly energy demand yields 130.8 kWh. So, the total energy consumption of the implementation of the AI sums up to 196.1 kWh. To calculate the CO_2eq of that energy consumption, the CO_2eq of the used grid mix needs to be considered. Therefore, we used Germany's electricity grid mix from the LCA Database GaBi [13] , which gives 0.568 kg CO_2eq per kWh electricity. This leads to a total of 119.2 kg CO_2eq due to the use of AI for sorting the cores. The CO_2eq savings of remanufacturing a starter in comparison to a new production is used to calculate the benefit that AI brings to the sorting facility in terms of CO_2eq. For this data from literature [8] is used, which gives a potential of saving 8.8 kg CO_2eq per starter. Based on this a total of only 12.66 starters need to be sorted additionally to remanufacturing to even out the emissions of the AI. Expressed differently, the sorting rate has to be improved by 0.02 %. As the AI has the potential of reducing the rate of mistakenly dumped cores from 6 % [6] to about 1.1 %, the total potential CO_2eq saving in the case of starters is 27.446 kg CO_2eq, which is equivalent to the total CO_2eq emissions of 3.54 German citizens in one year. [14]

Discussion

In this work, we were able to show that direct energy measurement shows significant deviations compared to state-of-the-art calculation methods for estimating the energy consumption of AI. The direct measurement for the used Image-AI, for example, showed a deviation by a factor of 2. Especially for larger AI trainings that run over a longer period this difference will increase. One reason for this is that during a training process, due to the loading and unloading of data into the GPU RAM, a GPU can never run permanently under full load. It should be noted that the PUE in equation 1 is assumed to be ideally 1 by the authors because IPC and PCs were used, which is legitimate for small computing systems. For a more accurate calculation, the PUE for each PC should be modeled by a comparative measurement, for which the additional consumption of the power electronics and peripheral areas of the mainboards should be considered.

The results of the direct energy measurements show that, considering the 1 % rule, the data generation, as well as BDE, has no significant influence on the energy consideration. It proves true that training the image-processing AI emerges as the largest energy consumer in the recorded data. The bottom line is that all other energy consumptions are negligible in contrast to the training of the AI. Nevertheless, the AI manages to quickly deliver a return on invested CO_2eq despite a very high estimate of 12 training cycles per year.

Despite insufficient performance for detection accuracy, BDE has the strong advantage that it can be used ad-hoc in a digitally documented sorting process without having to collect extra data. Thus, BDE already provides predictions to a worker while the necessary data must be collected for Image-AI. It can also increase the quality of the prediction in a fusion with Image-AI. This will be investigated in another paper.

In this study, 511 starters were used. For this purpose, it was studied that 500 of the most sorted starters account for 85 % of the total volume. If an employee works with the same objects every day, it can be assumed that he can recognize them after a training period without artificial intelligence. However, an employee in a sorting facility does not only process starters but a total of approx. 30 different product groups with about 300,000 different product numbers. With high turnover rates, as is common in the logistics sector, and the usual amount of experiential knowledge, this task cannot be handled by a human. At the same time, the AI does not differentiate between high-runners and rare old parts but treats everyone with the same objectivity, regardless of location and time of day. By using the AI system, the accumulated knowledge of the company's employees becomes simultaneously available in different locations around the world.

The scalability to a changed number of learned objects is not presentable. Since the energy consumption for training an AI does not correlate linearly with the number of objects, simple scalability to more parts cannot be assumed. This behavior is the subject of future work. In addition, transferability to other product groups is limited. If for an image-processing AI, it has already been shown in other work that the methods also work for other objects, on the other hand, the CO_2eq savings from refurbished components compared to new parts is strongly divergent and must be investigated anew for each product. Nevertheless, it can be assumed that similar savings potentials of around 50 % exist for old automotive parts of similar size and comparable material.

Conclusions

We have shown in this work that from the point of view of energy / CO_2 it makes sense to invest energy in the development of AI, since this improves the processes sufficiently to generate a "return on energy". Our results show that the most significant energy expenditure is due to the use of graphics cards in the training of AI. The energy consumption costs for digitization are negligible and so are the costs for running the AI.

However, compared to the energy-intensive production processes in the automotive industry, the invested energy for AI is very low. A fast break-even point for AI-induced invested energy can be reached. With a throughput of starters improved by 4.9 % due to AI, the potential to save CO_2eq is 27,446 kg per year. In our study, the break-even point for the return on energy investment by AI methods would statistically be achieved after sorting cores for only 8 hours.

Acknowledgements

This work is part of the EIBA project, funded by the German Federal Ministry of Education and Research (033R226). We thank our partners of the project consortia for their support, especially Circular Economy Solutions GmbH for sharing data and insights. The authors acknowledge the support of Philipp Drebinger.

References

1. COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS A New Circular Economy Action Plan For a Cleaner and More Competitive Europe; 2020;
2. Kurilova-Palisaitiene, J.; Sundin, E.; Poksinska, B. Remanufacturing Challenges and Possible Lean Improvements. Journal of Cleaner Production **2018**, 172, 3225–3236, doi:10.1016/j.jclepro.2017.11.023.
3. Xing Bo; Wen-Jing Gao; Kimberly Battle; Tshildzi Marwala; Fulufhelo Nelwamondo Artificial Intelligence in Reverse Supply Chain Management: The State of the Art; Advisory Committee on Remanufacturing: Annual Report: London, 2010;
4. Flemming, F.; Balthasar, D. An Introduction to AI in Sorting Technologies Available online: https://www.recyclingtoday.com/article/an-introduction-to-ai-in-recycling-sorting-technologies/ (accessed on 11 April 2022).
5. Schlüter, M.; Niebuhr, C.; Lehr, J.; Krüger, J. Vision-Based Identification Service for Remanufacturing Sorting. Procedia Manufacturing **2018**, 21, 384–391, doi:10.1016/j.promfg.2018.02.135.
6. Schlüter, M.; Lickert, H.; Schweitzer, K.; Bilge, P.; Briese, C.; Dietrich, F.; Krüger, J. AI-Enhanced Identification, Inspection and Sorting for Reverse Logistics in Remanufacturing. Procedia CIRP **2021**, 98, 300–305, doi:10.1016/j.procir.2021.01.107.
7. Schwartz, R.; Dodge, J.; Smith, N.A.; Etzioni, O. Green AI. arXiv:1907.10597 [cs, stat] **2019**.
8. Strubell, E.; Ganesh, A.; McCallum, A. Energy and Policy Considerations for Deep Learning in NLP. arXiv:1906.02243 [cs] **2019**.
9. Köhler, D.C.F. Regenerative Supply Chains: Regenerative Wertschöpfungsketten, 2011.
10. Electronics Goes Green 2012+ (EGG 2012): Berlin, Germany, 9 - 12 September 2012 ; [Joint International Conference and Exhibition ; Proceedings]; Lang, K.-D., Fraunhofer-Institut für Zuverlässigkeit und Mikrointegration, Eds.; IEEE: Piscataway, NJ, 2012; ISBN 978-3-8396-0439-7.
11. Fu, A.; Hosseini, M.S.; Plataniotis, K.N. Reconsidering CO2 Emissions from Computer Vision. arXiv:2104.08702 [cs] **2021**.
12. Cunningham, A. Mini-Review: Intel’s Powered-up Core I7 Broadwell Mini PC Available online: https://arstechnica.com/gadgets/2015/03/mini-review-intels-powered-up-core-i7-broadwell-mini-pc/ (accessed on 11 April 2022).
13. Sphera Solutiions GaBi Ts; Sphera Solutiions GmBH: Leinfelden-Echterdingen, 2021;
14. Statista CO2-Ausstoß pro Kopf weltweit nach Ländern Available online: https://de.statista.com/statistik/daten/studie/167877/umfrage/co-emissionen-nach-laendern-je-einwohner/ (accessed on 11 April 2022).