Sustainability Budgets: A Practical Management and Governance Method for Achieving Goal 13 of the Sustainable Development Goals for AI Development

Rebecca Raper [1,*], Jona Boeddinghaus [2], Mark Coeckelbergh [1], Wolfgang Gross [2], Paolo Campigotto [2] and Craig N. Lincoln [2]

1 Department of Philosophy, University of Vienna, 1010 Vienna, Austria; mark.coeckelbergh@univie.ac.at
2 Gradient Zero, 1010 Vienna, Austria; jb@gradient0.com (J.B.); wg@gradient0.com (W.G.); pc@gradient0.com (P.C.); cl@gradient0.com (C.N.L.)
* Correspondence: rebecca.raper@univie.ac.at

Abstract: Climate change is a global priority. In 2015, the United Nations (UN) outlined its Sustainable Development Goals (SDGs), which stated that taking urgent action to tackle climate change and its impacts was a key priority. The 2021 World Climate Summit finished with calls for governments to take tougher measures towards reducing their carbon footprints. However, it is not obvious how governments can make practical implementations to achieve this goal. One challenge towards achieving a reduced carbon footprint is gaining awareness of how energy exhaustive a system or mechanism is. Artificial Intelligence (AI) is increasingly being used to solve global problems, and its use could potentially solve challenges relating to climate change, but the creation of AI systems often requires vast amounts of, up front, computing power, and, thereby, it can be a significant contributor to greenhouse gas emissions. If governments are to take the SDGs and calls to reduce carbon footprints seriously, they need to find a management and governance mechanism to (i) audit how much their AI system 'costs' in terms of energy consumption and (ii) incentivise individuals to act based upon the auditing outcomes, in order to avoid or justify politically controversial restrictions that may be seen as bypassing the creativity of developers. The idea is thus to find a practical solution that can be implemented in software design that incentivises and rewards and that respects the autonomy of developers and designers to come up with smart solutions. This paper proposes such a sustainability management mechanism by introducing the notion of 'Sustainability Budgets'—akin to Privacy Budgets used in Differential Privacy—and by using these to introduce a 'Game' where participants are rewarded for designing systems that are 'energy efficient'. Participants in this game are, among others, the Machine Learning developers themselves, which is a new focus for this problem that this text introduces. The paper later expands this notion to sustainability management in general and outlines how it might fit into a wider governance framework.

Keywords: AI; artificial intelligence; sustainability; AI governance; ethics; ethical AI; differential privacy

1. Introduction: Sustainability and AI

In 2020, carbon dioxide levels in the atmosphere were measured to be 149 percent above the pre-industrial level, Nitrogen levels at 262 percent and Nitrous Oxide at 123 percent [1]. By the end of the century, it is expected that these increased levels will contribute to a global warming of 1.5–2 degrees Celsius above the pre-industrial levels. This will not only lead to a change in climate conditions, but also contribute to melting ice caps, rising sea levels and ultimately pose risks to humans due to flooding, dangerous weather and disrupted crop growth [2]. In 2021, it might be claimed that we are already beginning to see some of the impacts of global warming, with significant heatwaves, forest fires and flooding causing global disruption in 2021 [3–5].

In 2015, the United Nations (UN) set their 'Sustainable Development Goals' (SDGs) as key global priorities for a better world by 2030 [6]. One of the key goals (Goal 13) is 'to take urgent action to combat climate change and its impacts'. More recently, at the 2021 World Climate Summit (a global conference aimed at addressing climate change), a worldwide pledge was made for countries to do more to reduce their carbon footprints.

To address these issues, governments need to take action to address the reliance on infrastructure that is fuelled by carbon. As Stein [7] outlines, Artificial Intelligence has the power to deal with some of these challenges through improved efficiency, accuracy and efficacy of systems; however, it can also be a significant contributor to some of the same issues it aims to resolve, due to the large amount of computing power that is often required to train and test modern (Machine Learning-based) AI systems. Strubell et al. [8] found that training large natural language machine learning models can consume huge amounts of energy, with the average AI development projects, training many models to find the best solution, emitting as much as 78,000 pounds of Carbon Dioxide—more than five cars in their lifetime. Moreover, Strubell et al. [8] showed this to be the case especially during the last optimization phases of AI development, i.e., increasing the performance of an already well-performing AI model a few more points, consumes disproportionally high amounts of energy for the utility of the project. Therefore, focusing on the AI development perspective is as important as considering broader management and governance.

With the 'Fourth Industrial Revolution' [9] accelerating the rate at which AI is adopted in society, and in the light of the possibility that AI is going to be used as a tool to resolve issues relating to climate change, there are serious issues regarding sustainability that need addressing, namely, the amount of energy that is consumed.

One strategy for dealing with the amount of energy expended by AI systems is to exhaustively document the amount of energy that the development and operation of an AI system uses, from idea generation, through to training, re-tuning, implementation and governance [10], and to actively work to minimise these levels. However, from a governance perspective it is not easy to see how this could be managed. In the first instance, there are minimal procedures or methods to document how much energy the production of an AI system expends (or might expend), and, even if this were known, it is not known what should be done to minimise this level. From a product manager's perspective, there is no access to the development process to understand how changes could be made to reduce the amount of energy consumption, and, from a developer perspective, they do not have the broader understanding of the system to see where allowances could be made and when energy should be reduced.

Furthermore, following the documentation, usually the idea is then to achieve minimization of energy levels by constraining developers and their companies via regulation. However, this may be politically problematic and controversial, since it raises the issue of freedom [11] and does not respect the autonomy and creativity of the developers and other stakeholders such as management and end users. The focus is also on the organisation; developers are often not directly addressed. What is required, then, is a methodology to combine and integrate these insights to allow for appropriate oversight, management, development and governance in a way that respects the autonomy and creativity of stakeholders, in particular the AI developers.

The remainder of this paper traces a method towards achieving this aim: utilising 'Sustainability Budgets', in analogy with Privacy Budgets in Differential Privacy, it develops a procedure that empowers developers, allows management to have sufficient oversight and provides a governance framework towards achieving Goal 13 of the SDGs. Section 2 introduces Differential Privacy, Privacy Budgets and then 'Sustainability Budgets'; Section 3 outlines how this can fit into a wider organisational management framework, including using Gamification to incentivise more sustainable development. In Section 4, Sustainability Budgets are outlined in relation to sustainability in society and governance towards achieving the SDGs. Finally, there is reflection on the method, along with an outline of the limitations in Section 5, before a conclusion.

2. Our Proposed Method: 'Sustainability Budgets' in Analogy with Differential Privacy

2.1. Differential Privacy and Budgets

Managing the privacy of individual data is a priority in data science. Often, within the development phase, vast quantities of personal data are used to train the AI system. Owing to regulatory frameworks, such as GDPR [12], organisations and developers have a duty to protect the personal data of individuals and to ensure it is not lost or given to the unapproved individuals. Privacy becomes particularly important in settings involving sensitive data (for example, health care settings). In the wrong hands, such data could be used to manipulate (i.e., for insurance purposes) for financial gain. Therefore, in these settings, it is even more important to ensure that data privacy is maintained. Some personal data is so sensitive that it requires specialist approval or clearance before it can be seen. It can be very time consuming and difficult from a management perspective to ensure that every developer and project team member working on development of an AI system has the right degree of clearance. One approach to resolve this is to prevent the project team having access to such sensitive data at all. However, for AI development, access to the data is required to train the models in the first place.

A resolution to this problem is to keep the data on a separate system so that training occurs 'remotely', and the data scientist does not have access to the information in the data that is being used to train but can still see the analysis outcomes. However, even if information is not immediately accessible to a data scientist, through hacking techniques (i.e., querying the information in certain ways) personal information can still be retrieved. In 2006, Cynthia Dwork et al. introduced a mathematical technique which prevents this problem, known as Differential Privacy (DP) [13]. In DP, developers can interrogate (or analyse) individual data securely, while preserving its anonymity. A differentially private algorithm guarantees that its output does not change significantly (as quantified by the DP parameter ε) if one single record in the dataset is removed or modified or if a new record is inserted in the dataset. Differential Privacy protects individuals in "that the analyst knows no more about any individual in the data set after the analysis is completed than she knew before the analysis was begun." [14].

However, if not managed effectively, the more a data set is queried, the greater the chance is of retrieving potentially private information from that data. Within DP, this is what is referred to as 'Privacy Leakages', whereby unwanted revelations about individual data points are gained by repeatedly performing queries on seemingly harmless statistics. For example, I might be able to identify who an individual person is in a medical data set, by performing repeated queries on known medical diagnoses. Within DP, this is often managed using a notion known as 'Privacy Budgets'. The Privacy Budget is a direct consequence of the privacy parameter ε as introduced by Dwork [13]. As the ε parameter in this so-called epsilon-DP defines the degree of the allowed influence of individual data points on analysis results, the privacy budget is the upper limit of accumulated epsilon values for a given dataset, effectively guaranteeing the privacy of individuals. In simple terms, the Privacy Budget defines an upper limit for information that can be disclosed for a given dataset. The important point being that every time a data analytics task is performed—be it a SQL query with aggregate statistics or a Machine Learning training job—some information about the data is revealed. The 'purported ideal' would be to have a completely anonymized data set that at the same time can be used for sensible data analysis; however, as this is not achievable, the amount of 'privacy leakage' must be balanced against the utility of the outputs. Differential Privacy, with its mathematical definition of privacy, makes this fact clear and transparent.

DP is a robust concept that tries to focus on how much information is retrieved from a data query. A particularly interesting feature of this approach is that privacy issues need to be dealt with at the time of initial development, which puts responsibility at the start of development procedure.

One aspect of sustainability and, more specifically, energy efficiency in Machine Learning development, is that it is all too often considered as an organisational issue,

where an organisation that is aligned with sustainability goals, strives to achieve a lower carbon footprint by reducing the energy consumed within its data centres [15]. While this is definitely an important endeavour, the AI development perspective is often overlooked but can contribute significantly to achieving sustainability goals.

As with 'Privacy', from a legal and governance perspective, we have a duty to manage the amount of energy that is expended when AI systems are developed. Akin to 'Privacy Budgets' within DP, we introduce the concept of a "Sustainability Budget" or "Sustainability Score" to allow sufficient oversight of this.

In DP analysis, the amount of information can—and must—be specified before each data query computation. That means, in practice, before an AI developer starts to train a Machine Learning model on a dataset, they need to detail the specified amount of 'information leakage'. The high the budget that is requested, the more accurate the model will be, but the less remaining room there will be for further computations on the dataset, and vice versa. The privacy problem is put in the hands of the developer and therefore becomes a central element of the analysis itself which leads to nothing else other than a strong requirement for data protection awareness at the development level. The developer must carefully think about how they would like to use the data; once the analysis is started (or more specifically for DP: once the analysis result is published), the decision has been made. In a similar regard, a 'Sustainability Budget' can motivate developers to think about the effects of their to-be-started analyses before they hit the run button.

Let us define the Sustainability Budget (SB) as a virtual upper limit of available compute resources for Machine Learning jobs. Likewise, a 'Sustainability Score' can be defined as the inverse: the to-be-consumed compute resource deducted from a virtual upper limit.

Such a score would not constrain developers and their organisations but instead raise awareness on the side of the developer and incentivize solutions that are less resource-intensive and therefore more sustainable.

To calculate and use this new Sustainability Budget, we need a formula for the carbon footprint of a software system and metrics that can be used to calculate it. A good template comes from the Greensoftware Foundation and their "Software Carbon Intensity (SCI) Specification". An alpha version can be found in [16].

The formula for the SCI rate is defined as:

$$SCI = ((E \times I) + M) \text{ per } R, \tag{1}$$

where E is the energy consumed by a software system, I is the location-based marginal carbon emissions, M is the embodied emissions of a software system, and R is a functional unit.

We can set R to be one ML (training) job and thereby calculate the SCI per the ML process. SCI indicates how much carbon equivalent emissions are produced per kilowatt-hour (energy, E). This must be identified with the local energy provider and data centre administration team. Similarly, the embodied emissions are calculated beforehand based on the employed hardware systems. To quantify the carbon equivalent emissions per unit, one approach is to measure the actual consumption by implementing a respective monitoring at the supplier level. However, there are other 'prediction' approaches, potentially conceptual ML techniques, that could also be conceptual.

Once the initial calculations and measurements are in place, the SCI can be used to compare different ML jobs (units R) and to publish the rates in the team's sustainability reports. On a basic level, with everything else fixed, the SCI comes down to measuring the energy consumption of a ML job. With a balanced set of compute resources this can further be refined to measure the time one ML job is running.

Having such a compute-time or SCI rate measurement in place, AI development teams can now define a Sustainability Budget (or goals for sustainability scores) for different projects. Each developer will be aware of this budget. They will plan for an efficient model training process using as little energy as possible for a model that is as high performant level (accurate, precise, private, etc.) as possible under these conditions. Since model

optimization usually follows a logarithmic curve (model performance vs. iterations), a good approach would be to estimate the expected performance gain with each step, i.e., experiment, that the machine learning developer performs while finding an optimal model. This makes the sustainability score part of the development process itself.

2.2. A Practical Case Study: Using Bayesian Optimization for Experiment Selection

If we look at the technical development process at an individual level, ML development can be described by a series of smaller experiments that are executed to develop one Machine Learning model. In this sense, each 'experiment' uses computational resources that can be described with the SCI rate introduced above. All the experiments (units R) need to fit into the given Sustainability Budget for the entire development lifecycle [10]. So, the question becomes how to structure work as an ML engineer, and what the most efficient workflow is to obtain the best model, within a given budget.

The quantity of data and size of the model are the biggest drivers for the computational cost of each experiment [17]. There is work that tries to mitigate this by working with efficient networks [18,19], compression networks [20–22] or a minimum amount of data [23,24]. Not only can this reduce computational cost during development, but it can also solve other limitations, such as computational capacity on edge devices [18,19] or the availability of data. Moreover, ideas such as Transfer Learning, etc. [25] can be used to develop models more efficiently.

These are interesting approaches, but to highlight how a sole developer might drive sustainability in their individual development, we implement Sustainability Budget approach, this section focuses on Bayesian Optimization as a case study, to guide the selection of experiments and reduce the number of experiments needed to reach the desired result.

To understand how this can be done, it is important to realise that each experiment is carried out to gain information about the model's performance for a specific set of parameters. For the ML engineer, this is the main part of their work, where they want to understand, with high certainty, how high the model's error is with different settings, in order to determine the best set of settings or to conclude how the different settings influence the results for comparison. The model's error can be described by a function, with input space over the parameters and the output space over the error of the model, where low errors represent good models. This error function is unknown and cannot be optimized directly, so optimization methods such as Gradient Descent [26] are not applicable and only black box optimization methods [27] or brute force methods (systematically calculating all possible options) can be used.

Naturally, systematic testing of different input sets could be carried out to estimate the error function for different parameters. However, this would not be favourable, because it is not very efficient and grows exponentially with the number of parameters, because each parameter is tested individually. It is very likely that some parameters have a bigger impact on the model's performance than others, but we do not know which ones beforehand. If known, the most important should be tested first. However, this becomes more complicated when we consider that the parameters are not independent at all, and that randomly selecting parameters is a better option, because it is more efficient in testing more unique combinations, which has also been proven experimentally [28]. To achieve even better efficiency, the information from previous experiments should be used to determine which parameter set it evaluates next. This can be conceptualized with Bayes Formula:

$$P(H \mid E) = P(E \mid H) \, P(H)/P(E). \tag{2}$$

Bayes Formula of conditional probabilities puts the prior probability $P(H)$ of the hypothesis H in relation with the likelihood $P(E \mid H)$ of the evidence E for a given H. With the prior and the likelihood, we can calculate the posterior probability $P(H \mid E)$, which is what is most interesting. The term $P(E)$ is the marginal likelihood of all hypotheses H and can be treated as a normalisation constant for these purposes.

Applying this to the experiment selection problem of the ML engineer, this method can be used to find the optimal set of parameters for the error function of the model. Due to the fact that the function cannot be accessed directly, it is treated as a probability function. The prior P(H), expressed as a probability function over the parameter space, gives the prior belief of the form of this error function. Prior to any experiments, it can be assumed that a uniform probability density function (pdf) would represent the belief that no set of parameters is advantageous to any other set, but the ML engineer might just as well implement any other prior beliefs that may be held over the model's error function into the equation. The evidence represents the experiments performed and the likelihood of this evidence, which indicates the probability that the evidence fits together with the hypothesis. The posterior can be calculated by combining the likelihood and the prior which can be understood as an update of the prior once new evidence (i.e., information) is present. This process can be repeated when new evidence is present. The prior of the next interaction is set equal to the posterior of the latter interaction, and the new posterior is calculated, i.e., the belief is updated. This represents an update of the belief on the ML engineer (or the system) given new evidence.

To select the best experiment to perform next, the described posterior function (which is often modelled as a gaussian process) can now be used, and the next point to evaluate can be derived (see Figure 1). During the gaussian process, a density estimation is gained for each point on the error function rather than a simple point estimate. In the plot below, the line in the middle is the mean of the estimate and the area around represents the uncertainty, i.e., the density of the estimate. The red dots are points where the function has been evaluated in some experiments, and the uncertainty is therefore zero at these locations.
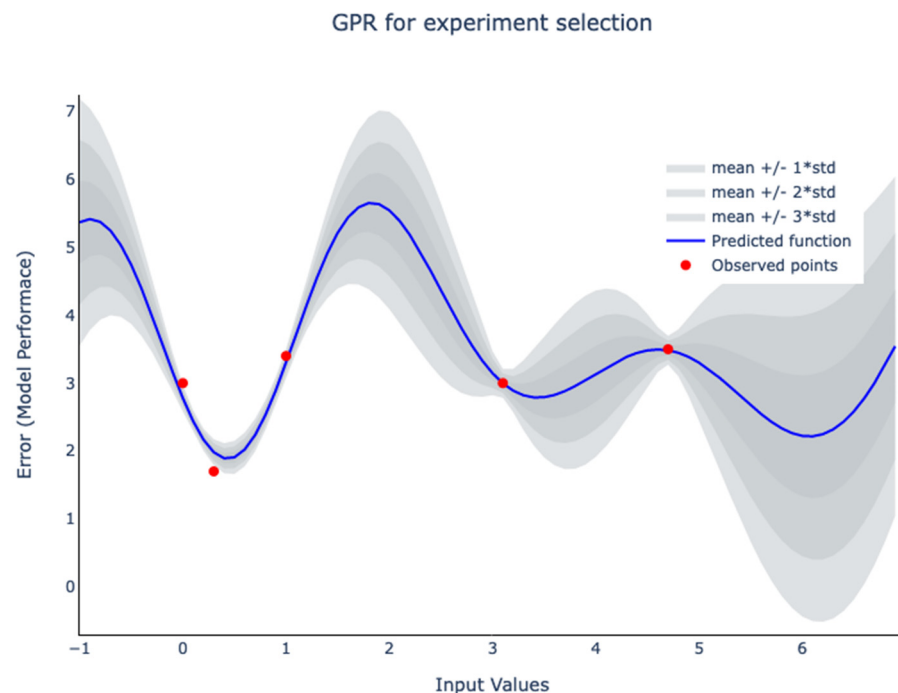

Figure 1. A Gaussian Process Regression (GPR) model that has been fitted on observed data (red dots). In the context of the case study, these observations would be the results of previous experiments. The predicted function (blue) is the result of the GPR and is simply the mean of the GPR posterior estimation. The grey area visualises the uncertainty of this function. If the parameters are desired (note that really the number of parameters is much greater than 1), in the next experiment the plot can be used to find a good trade-off between exploration and exploitation. What is seen is that the dip around 0.5 is certain, but further development would likely result in limited performance gains. Alternatively, the area around 4 or 6 could be explored, with high uncertainty, to improve the current estimation and avoid missing a better local optimum for the model performance.

Now, this provides the information to choose which point of the error function to evaluate next. Depending on the task, the current optimal point could be evaluated or a point with large uncertainty to improve the estimate there.

This strategy is an example of how to make informed decisions while optimising a machine learning algorithm. With every step, the next best experiment is decided. Therefore, the cost per unit as described above is clearly present in the form of the number of experiments to be performed. It is a good idea to keep this process a manual one (instead of trying to automate away the exploration) as it brings the desired care into the process. The ML developer is aware of the sustainability costs associated with each experiment and therefore uses a carefully chosen method to optimise the model, integrating and updating beliefs with each iteration.

3. 'Gamify!': Gamification Techniques to Manage Sustainability

We introduced Sustainability Budgets and have offered a practical way developers might use up-to-date ML techniques to achieve optimum sustainability. With these in mind, let us ask how developers might be incentivised to achieve more sustainable development. For example, while a developer might have awareness of how much energy the particular development of their system expends and the capability to minimise it, without some form of incentivisation, the whole procedure may seem pointless. Why invest time and energy in making a system more efficient if there are no benefits for doing so?

While we believe this to ultimately be a management issue, where managers decide how to maintain sustainability, Sustainability Budgets open the possibility to use 'Gamification', which has been shown to be successful in incentivising participation.

Gamification—a notion taken from the gaming industry—is the enrichment of products, services and information systems with game style elements in order to positively influence the productivity, motivation and behaviour of users [29]. An early example was the Nike+ running application [30], which encouraged individuals to engage in physical activity by making a 'game' out of running. Users were encouraged to compete against other runners for positions on leader boards, 'fastest course completions' or simply 'most runs that week'. More recently, other exercise applications such as Strava and Peloton make use of the gamification technique, and this can be extended to other areas.

Sustainability Budgets lend themselves very well to gamification techniques, with 'Sustainability Scores' or Limits usable to drive competition between developers. Therefore, a developer might be incentivised to develop their system in the most efficient way, because they are part of a 'game' with other developers, where there is a benefit for 'winning' (whether this be financial or otherwise).

As well as lending itself to competition between developers, moving one abstraction level away, we can begin to look at how this might also translate to operating at an organisational level.

As mentioned in Section 1, many organisations are motivated to reduce their carbon footprints, but this often seems to weigh against certain costs, meaning that there often has to be a compromise on sustainability in order to drive down the costs of development. However, if gamification were used as a technique to not only incentivise developers, but also to audit how systems were developed and their efficiency (through a Gamification management system), it would be easier for organisations to understand and explain the added time and costs for the delivery of projects. Ultimately, gamification systems would not only be instrumental in incentivising developers, but they could also be used by organisations to show why they are particularly sustainable

It could also be applied from an executive management perspective, and it could even be used at a national and global governance level. There could be competitions between teams of an organisation or even across different organisations or countries.

Utilising the 'Sustainability Budgets' methodology, there will be a mechanism to plan for and log the amount of energy expended per development task and also information for developers to provide explanations at a management level for different components of

the development process. For example, a developer would be able to articulate why they chose a particular training step or method over another, in terms of the balance between sustainability and effectiveness, and there would be oversight as to how much energy was expended per process. The developer would be incentivised to be more energy efficient in their development approach, and there could be incentivised competitions between different developers to find the most 'energy efficient', effective approach.

The incentivisation could be extended to include 'games' between different developers. From a management level, this would give the chance to reward more sustainable development through a points/score-based system. There would need to be a balance between sustainability and effectiveness of the algorithms, but a score-based system would allow visibility of the compromises and management accordingly. For example, as a manager, I might decide to increase my sustainability budget (i.e., the amount of energy a development project expends), because the project is particularly pertinent (from a business perspective) or because having a lower budget would make the system unsafe (i.e., in a medical context). Equally, I might decide to lower my budget because, from a resource perspective, the project allows for this. Gamification would encourage developers to create their systems in the most energy efficient ways, being rewarded for creative technical solutions.

4. Using Sustainability Budgets to Achieve the SDGs

Eventually, the proposed method might also contribute to wider political and societal goals. At the beginning of this paper, the Sustainable Development Goals (SDGs) were discussed, in particular, Goal 13 and the drive for countries to achieve carbon neutrality, or at least reduce their carbon footprints. As well as offering a mechanism for incentivisation through Gamification (or similar) approaches, at either a developer, management or governance level, what Sustainability Budgets ultimately provide is a mechanism for 'Energy Consciousness' insofar as not only do developers become more aware of the amount of energy the development of their system expends, but managers can translate this into development action, and governments can reward/recognise organisations who offer more sustainable solutions or govern organisations to be more sustainable.

Strategically, those systems that offer the greatest benefit/cost ratio would be the systems that are invested in the most. Where previously it was difficult to determine this ratio in the early phases of a project because there was no mechanism to quantify it beforehand, there can be effective management of energy expenditure at all levels of an organisation: a company, for example, but also a sector and even a nation.

For example, energy budgets could be set on an industry basis, based upon the needs of the industry. There could be regulation to ensure that industries did not go beyond their allowance, and the different levels could be managed at a political level. Organisations could be incentivised to be sustainable through energy ratings that would ensure that standards were being maintained. An organisation might receive an 'A' energy rating for particularly energy efficient processes and a lower rating ('B') for a lower energy efficiency. As with energy efficiency levels used to rate buildings and appliances (ref), the energy expended for each project would need to be measured against the size and utility of the project.

A strong contribution of the Sustainability Budget approach is to provide a way to embed the topic of environmental sustainability into ML practices and organisations from the beginning. The urgent need to address climate change is injected directly in the project development phase of machine learning and AI applications. Whereas most of the time such topics are managed purely in an abstract way with review processes that are being applied after the fact and that are followed or not, sustainability budgets put this matter at the heart of the development process itself. With the aforementioned gamification implementation, there is a direct and accessible way to address these problems.

And most of all: sustainability budgets raise awareness about this topic. Next to incentivisation, sustainability budgets and the related gamification strategies are likely to

raise awareness of energy consumption and sustainability issues, which may influence the behaviour of individuals and the organisation. Just like privacy budgets in differential privacy, sustainability budgets need to be handled before starting a training run of a machine learning task and, because of that, they force developers, project managers and policy makers to think about the impact of the project early on and continuously. It does matter if a machine learning task is solved by using the next best large network trained on all available next best data or rather designed and developed with energy consumption and sustainability in mind, carefully selecting the optimal network architecture, data subset and iterative development approach. Being aware, after all, is the first and most important step when trying to change and influence things. Sustainability budgets help raise awareness of the climate change impacts Machine Learning development can have at and before the development process itself.

5. Limits of the Methodology

It is important to note that the notion of 'budgeting' energy consumption has its limitations. As with the concept of budgeting privacy, careful calculations need to be carried out to weigh up the balance between the amount of energy that should be expended and how useful it would be to expend that energy. If the correct ratio is not established, then it becomes useless to set sustainability limits in the first place. If energy limitations are hindering innovation and project development, then the budget needs to be reassessed and project priorities need to be reassigned. This, however, is something that would need considering at the management level.

Another limitation of the presented approach is the extent to which the privacy budget metaphor extends to the case of energy budgets. It is noted that there are some dissimilarities between the two notions, namely 'privacy' as a unit of measurement is an abstract concept whereas 'energy' has real tangible values (i.e., the amount of energy that is consumed). Differential Privacy is a mathematical concept, whereas sustainability budgets stem from an organisational or management background. However, focusing on such dissimilarities distracts from the focus of the approach that has been detailed, which is an emphasis on creating an awareness of how sustainable a project might be and putting in a framework to mitigate or manage this. In terms of the methodological approach and its consequences, namely the required attention to the data protection and sustainability factors, respectively, in the development process itself, there are useful similarities that can be taken from the management of differential privacy to encouraging sustainability.

Finally, one key limitation of the approach detailed in this paper is the extent to which this is an implementable process. What has been given is a conceptual approach and a sample case study, but by no means is it ready to be placed into the software development process. Ultimately, we see this as an area for future development, and if Sustainability Budgets would be pursued, an area that would require further research. The following areas need developing further:

Calculation of the appropriate ratio for the performance vs. cost.

Investigation for how organisations ought to be governed given energy efficiency levels.

Development of a suitable gamification platform on which to record (and provide a game for) development energy consumption.

6. Conclusions

In this paper we offered a technique for managing sustainability in the development process of AI system creation, which utilises a notion we term 'Sustainability Budgets'. This both empowers the developer creating the AI system and allows the management and governments to have appropriate oversight. The notion of 'games', whereby developers can engage in competitions to achieve the most sustainable product vs. its effectiveness, is introduced as a possible incentive to encourage developers to achieve their organisational and a nation's sustainability goals. The conversation in this paper not only offers a new methodology for individuals to move closer towards achieving the SDG's, it also inspires

debate in this area, and this may lead to even more practical ways to ensure that Artificial Intelligence aids (rather than hinders) a sustainable future.

References

1.　UN News. 'No Time to Lose' Curbing Greenhouse Gases: WMO. 2021. Available online: https://news.un.org/en/story/2021/10/1103892 (accessed on 13 December 2021).
2.　UK Met Office Online. Effects of Climate Change—Met Office. 2021. Available online: https://www.metoffice.gov.uk/weather/climate-change/effects-of-climate-change (accessed on 13 December 2021).
3.　NBC News. Heat Wave 2021: Climate Scientists Warn about a New Normal. 2021. Available online: https://www.nbcnews.com/science/environment/heat-wave-2021-climate-scientists-warn-new-normal-rcna1664 (accessed on 21 March 2022).
4.　The Guardian. Fires Rage around the World: Where Are the Worst Blazes? 2021. Available online: https://www.theguardian.com/world/2021/aug/09/fires-rage-around-the-world-where-are-the-worst-blazes (accessed on 25 March 2021).
5.　BBC News. Germany Floods: Dozens Killed after Record Rain in Germany and Belgium. 2021. Available online: https://www.bbc.co.uk/news/world-europe-57846200 (accessed on 25 March 2021).
6.　Globalgoals.org. The Global Goals. 2021. Available online: https://www.globalgoals.org/ (accessed on 13 December 2021).
7.　Stein, A.L. Artificial Intelligence and Climate Change. *Yale J. Reg.* **2020**, *37*, 890.
8.　Strubell, E.; Ganesh, A.; McCallum, A. Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL), Florence, Italy, 28 July–2 August 2019.
9.　Schwab, K. The Fourth Industrial Revolution. 2017. Available online: https://www.weforum.org/about/the-fourth-industrial-revolution-by-klaus-schwab (accessed on 22 March 2022).
10.　Wynsberghe, A.V. Sustainable AI: AI for sustainability and the sustainability of AI. *AI Ethics* **2021**, *1*, 1–6. [CrossRef]
11.　Coeckelbergh, M. AI for Climate: Freedom, Justice, and other Ethical and Political Challenges. *AI Ethics* **2021**, *1*, 61–72. [CrossRef]
12.　Regulation (EU) 2016/679 of the European Parliament and of the Council, General Data Protection Regulation. Available online: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:02016R0679-20160504&from=EN (accessed on 23 February 2022).
13.　Dwork, C.; McSherry, F.; Nissim, K.; Smith, A. Calibrating noise to sensitivity in private data analysis. In Proceedings of the TCC 2006: Theory of Cryptography Conference, New York, NY, USA, 4–7 March 2006; pp. 265–284.
14.　Dwork, C.; Roth, A. The algorithmic foundations of differential privacy. Found. *Trends Theor. Comput. Sci.* **2014**, *9*, 211–407. [CrossRef]
15.　Gao, J. Machine Learning Applications for Data Center Optimization. 2014. Available online: https://www.google.com/url?q=https://research.google/pubs/pub42542/&sa=D&source=docs&ust=1645625964864718&usg=AOvVaw2sF6awAp8KZDficvpwLU5i (accessed on 23 February 2022).
16.　Hussain, A.; Gupta, A.; Time, H.-W.; Buchanan, W.; Bergman, S.; Knight, V.; Lloyd-Jones, C.; Srinivasan; Kariya, M.; Lewis-Toakley, D. Software Carbon Intensity (SCI) Specification (v.Alpha). 2021. Available online: https://github.com/Green-Software-Foundation/software_carbon_intensity/blob/main/Software_Carbon_Intensity/Software_Carbon_Intensity_Specification.md (accessed on 2 February 2022).
17.　Justus, D.; Brennan, J.; Bonner, S.; McGough, A.S. Predicting the computational cost of deep learning models. In Proceedings of the 2018 IEEE International Conference on Big Data (Big Data), Seattle, WA, USA, 10–13 December 2018; pp. 3873–3882.
18.　Tan, M.; Chen, B.; Pang, R.; Vasudevan, V.; Sandler, M.; Howard, A.; Le, Q. MnasNet: Platform-Aware Neural Architecture Search for Mobile. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16–20 June 2019; pp. 1–9.

19. Howard, A.; Menglong, Z.; Chen, B.; Kalenichenko, D.; Wang, W.; Wayand, T.; Andreeto, M.; Adam, H. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv* **2017**, arXiv:1704.04861.

20. Han, S.; Mao, H.; Dally, W. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. *arXiv* **2015**, arXiv:1510.00149.

21. LeCun, Y.; Denker, J.; Solla, S. *Advances in Neural Information Processing Systems*; Massachusetts Institute of Technology Press: Cambridge, MA, USA, 1989; pp. 598–605.

22. Tanaka, H.; Kunin, D.; Yamins, D.; Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. *arXiv* **2020**, arXiv:2006.05467.

23. Arora, S.; Du, S.; Li, Z.; Salakhutdinov, R.; Wang., R.; Yu, D. Harnessing the Power of Infinitely Wide Deep Nets on Small-data Tasks. *arXiv* **2019**, arXiv:1910.01663.

24. Tartaglione, E.; Barbano, C.A.; Berzovini, C.; Calandri, M.; Grangetto, M. Unveiling COVID-19 from chest X-ray with deep learning: A hurdles race with small data. *Int. J. Environ. Res. Public Health* **2020**, *17*, 6933. [CrossRef] [PubMed]

25. Donahue, J.; Yangqing, J.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; Darrell, T. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. In Proceedings of the 31st International Conference on International Conference on Machine Learning, Beijing, China, 21–26 June 2014; pp. 647–655.

26. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.

27. Audet, C.; Kokkolaras, M. Blackbox and derivative-free optimization: Theory, algorithms and applications. *Optim. Eng.* **2016**, *9*, 100011. [CrossRef]

28. Bergstra, J.; Bengio, Y. Random Search for Hyper-Parameter Optimization. *J. Mach. Learn. Res.* **2012**, *13*, 281–305.

29. Deterding, S.; Dixon, D.; Khaled, R.; Nacke, L. From game design elements to gamefulness: Defining "gamification". In Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments, New York, NY, USA, 28–30 September 2011; pp. 9–15.

30. Blohm, I.; Leimeister, J.M. Gamification. *Bus. Inf. Syst. Eng.* **2013**, *5*, 275–278. [CrossRef]