UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

An interdisciplinary exploration of trade-offs between energy, privacy and accuracy aspects of data

de Reus, Pepijn; Dresen, Kyra; Oprescu, A.; Irion, K.; Kolk, Ans

DOI
10.48550/arXiv.2410.00069

Publication date
2024

Document Version
Final published version

License
CC BY-NC-SA

Link to publication

Citation for published version (APA):
de Reus, P., Dresen, K., Oprescu, A., Irion, K., & Kolk, A. (2024). *An interdisciplinary exploration of trade-offs between energy, privacy and accuracy aspects of data*. (v2 ed.) ArXiv. https://doi.org/10.48550/arXiv.2410.00069

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:16 Sept 2026

arXiv:2410.00069v2 [cs.CR] 12 Oct 2024

RESEARCH ARTICLE

AN INTERDISCIPLINARY EXPLORATION OF TRADE-OFFS BETWEEN ENERGY, PRIVACY AND ACCURACY ASPECTS OF DATA

Pepijn de Reus[1] | Kyra Dresen[1] | Ana Oprescu[1] | Kristina Irion[2] | Ans Kolk[3]

[1]University of Amsterdam, Institute for Informatics, Faculty of Science

[2]University of Amsterdam, Institute for Information Law, Amsterdam Law School

[3]University of Amsterdam, Amsterdam Business School, Faculty of Economics and Business

Correspondence
Ana Oprescu, University of Amsterdam, Amsterdam, The Netherlands
Email: a.m.oprescu@uva.nl

Funding information
We acknowledge funding from our University's Interdisciplinary Research Theme Responsible Digital Transformations.

The digital era has raised many societal challenges, including ICT's rising energy consumption and protecting privacy of personal data processing. This paper considers both aspects in relation to machine learning accuracy in an interdisciplinary exploration. We first present a method to measure the effects of privacy-enhancing techniques on data utility and energy consumption. The environmental-privacy-accuracy trade-offs are discovered through an experimental set-up. We subsequently take a storytelling approach to translate these technical findings to experts in non-ICT fields. We draft two examples for a governmental and auditing setting to contextualise our results. Ultimately, users face the task of optimising their data processing operations in a trade-off between energy, privacy, and accuracy considerations where the impact of their decisions is context-sensitive.

KEYWORDS
Digital economy, energy consumption of machine learning, data privacy, privacy-enhancing techniques, synthetic data, k-anonymity, storytelling

Introduction

Besides many benefits, the digital ecosystem also brings societal challenges. Weighing the potential and the shortcomings of novel digital technologies informs policymakers and practitioners moving forward, and researchers can provide crucial insights to facilitate decision-making (Ciulli and Kolk, 2023). This paper aims to contribute by zooming in on two crucial implications of data-driven machine learning, i.e., the energy consumption and privacy. Both concerns have been raised in tandem prominently within the European Union (EU), which is the setting in which this study is placed. We undertake an interdisciplinary exploration of the trade-offs between the energy, privacy, and accuracy aspects of data processing using an experimental set-up (section 2). To facilitate decision-making we follow up this exploration with a storytelling approach (section 3), translating the technical findings to experts and decision-makers outside field of computer science. Our method, which extends earlier work in computer science that examines the change in energy consumption and accuracy of synthetic data and k-anonymity in machine learning tasks (Oprescu et al., 2022; de Reus et al., 2023), will be presented in section 2. First, however, we will indicate the broader context concerning the environmental and privacy aspects of the digital economy from the EU perspective.

A cornerstone in EU environmental policy is the European Green Deal and the "Fit for 55" package to reduce greenhouse gas (GHG) emissions by at least 55% before 2030. Interestingly, the EU has connected its overall aim to remain within 1.5 degrees of global warming to digitalisation, by stating that “the twin challenge of a green and digital transformation has to go hand-in-hand” (European Commission, 2020). Although this strategy seeks to harness digital technologies for achieving environmental goals, in this way also furthering the EU’s position in the digital realm, a concretisation is largely absent. Given that ICT’s carbon footprint is substantive, except if renewable energy is used, and will continue to rise in the future with the growth of the digital economy (Freitag et al., 2021), optimising the energy consumption of data processing is imperative. This applies especially to standard data processing operations. Nevertheless, current research to reduce the energy consumption of ICT and data processing are mostly targeted at academic readers (Verdecchia et al., 2023).

In addition to this increased focus on greening computing, data processing can trigger concerns over individuals’ privacy. In the EU, individuals’ personal data is protected by the fundamental rights to private life and to the protection of personal data. The General Data Protection Regulation (GDPR) provides for a corresponding legal framework for the processing of personal data (European Parliament and the Council, 2016). Using anonymous or anonymised data by contrast is no longer governed by the GDPR. Anonymisation techniques are considered privacy-enhancing technologies (PETs) that provide an important entry point for data operations that would otherwise be tightly regulated.

Combining these two societal concerns and EU policies raises the question of what privacy-enhancing algorithms add to the energy consumption of regular machine learning tasks. Moreover, there are perceived trade-offs between computing anonymised data and the accuracy of machine learning algorithms trained on this data. Ultimately, users face the task of optimising their data processing operations in a trade-off between privacy, accuracy, and environmental considerations. In the context of considering both the increased use of data and its climate impact, understanding which method would be preferred when using anonymisation algorithms (k-anonymity or synthetic data) in light of the energy consumed and the accuracy achieved is important. These are the aspects we will turn to in the next section.

1 | KEY CONCEPTS AND A METHOD TO ASSESS TRADE-OFFS

Several concepts from computer science form the background for this work, coherent to our approach we aim to make these comprehensible for a larger audience. This section provides an overview of anonymisation techniques that are data privacy-preserving, the energy consumption of code as a function of digital sustainability, and how the

treatment of data impacts the accuracy of machine learning models. In a subsequent step, these concepts are related to each other to explore the trade-off between the environmental, privacy, and accuracy aspects of data.

1.1 | Data privacy

Big data coins that the generation and processing of data are accelerating at scale (Mayer-Schönberger and Cukier, 2013). Training machine learning algorithms for instance requires large quantities of data. Data about individuals are called personal data and its use is in most countries regulated under personal data protection laws (Greenleaf, 2023), a development heavily influenced by the standard-setting role of the EU's GDPR (also called the 'Brussels effect', see Bradford (2020); Coche et al. (2023)). The GDPR applies to the processing of personal data belonging to an identified or identifiable individual. When the GDPR applies, processing personal data has to be lawful and thus conforms to a number of legal requirements that aim to protect individuals' rights and interests (Oostveen and Irion, 2018). Hence, the qualification as personal data or as non-personal data has important consequences for the application of the GDPR (Finck and Pallas, 2020). As a sub-set of privacy-enhancing technologies, anonymisation techniques are a means to remove personal data from datasets. After being treated with anonymisation techniques, the resulting dataset ideally becomes anonymised data and therefore non-personal data to which the GDPR no longer applies (Finck and Pallas, 2020). The advantage is evident because individuals' data privacy is no longer an issue with anonymised data that can than be processed without further restrictions. What is more of an issue is whether the anonymisation of personal data is considered effective and irreversible (AEPD and EDPS, 2020). Out of various anonymisation techniques, *k*-anonymity and synthetic data are briefly outlined below.

| *k*-anonymity

k-anonymity is an algorithm that anonymizes the data by leaving out unique data features that can be traced back to a single person (Samarati and Sweeney, 1998; Sweeney, 2002). For any given value, called *k*, the algorithm will ensure that for this value at least *k*+1 individuals are indistinguishable from each other (Sweeney, 2002). For example, a dataset with a *k*-anonymity of 5 ensures that for each data subject in the set, at least 5 other data subjects share the exact same properties. Because these 6 data subjects all have the same properties, they would thus be indistinguishable from each other. This process is exemplified in Figure 1.

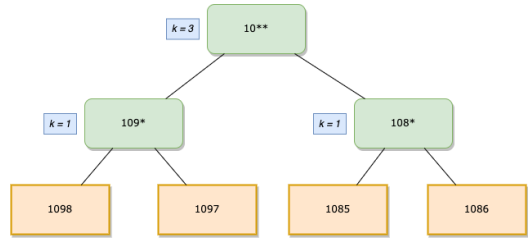

FIGURE 1 Anonymising of ZIP codes by suppression of the last digits. The more digits are suppressed, the more ZIP codes become indistinguishable from each other. The corresponding *k*-value is indicated in blue. Obtained from de Reus et al. (2023)

The *k*-anonymity algorithm utilizes generalisation and suppression which are applied to hierarchies of attributes in the data. Once the data attributes are labeled using the data types above and a value for *k* is chosen, the algorithm uses generalisation and suppression to alter the data (Samarati, 2001). To come to this *k*-anonymity, all attributes in a structured dataset should be classified as one of four data types:

- **Insensitive** attributes are deemed unimportant for privacy and will remain unaltered.

- **Sensitive** attributes are important for the subject (think of political views).
- **Identifying** attributes such as names are directly linked to a person and will be removed from the data.
- **Quasi-identifying** attributes refer to the data that could compromise privacy when linked with other attributes or datasets.

The algorithm leaves as much data as is while increasing the hierarchy of attributes in the dataset until the desired k-anonymity is obtained or until the entire attribute is suppressed. Yet, quasi-identifying attributes can still be unique trackers for this attribute. A known example is that a university has only one rector. If a university were to publish its salaries, staff functions would be quasi-identifying. However, the function of a rector is a unique tracker as just one person fulfils this job. So even though quasi-identifying features in general are not identifying, they could be in case of a unique value. This forms a risk for re-identification and, in turn, a compliance risk under the GDPR. Mindful that applying the k-anonymity algorithm is not by and in itself a guarantor for obtaining truly anonymised data, it really depends on the properties of the original dataset, the correct labeling of all attributes, setting an appropriate k-variable, and testing for re-identification.

| Synthetic data

An increasingly popular method to share data without comprising the data subjects' privacy is synthetic data (Hernandez et al., 2022). Synthetic data is a form of machine learning that learns the properties, co-variance, and distributions of a dataset and produces a representation of an original dataset. The synthetic dataset exhibits very similar statistical properties to the original dataset. Synthetic data hence mimics the properties of a dataset by creating artificial data that has a statistical resemblance to the original dataset without actually having overlap. After the synthetic data has been constructed, it should no longer contain personal data relating to individuals that are present in the original data. Because of this property, it is possible to publish and share synthetic data without compromising any individual's privacy. Most importantly, the performance of machine learning on the synthetic dataset should be similar to that of the original dataset (Abadi et al., 2016). As synthetic data bypasses privacy and security protocols, it enables organisations to access data more rapidly (Hittmeir et al., 2019).

The first use cases of synthetic data are promising. For example in healthcare settings, where doctors and researchers often work with medical data from patients, synthetic data is increasingly used to share medical data (Murtaza et al., 2023). Or the Digital Twin project from the EU, where a digital copy is made to run new models and predictions (Bauer et al., 2021).

| Digital sustainability

Global heating is the result of GHG emissions, for example during energy generation of conventional energy options (Schneider, 1989). One solution for this problem is to change to green energy generation. Another solution is to decrease energy consumption. Currently, we are operating on the basis of very rough estimations of what ICT's current share of GHG emissions is (Mytton and Ashtine, 2022). While conservative estimates attribute a 1.8%–2.8% share of global GHG emissions to ICT, other research suggests "that global emissions from ICT are as high as 2.1%–3.9%" (Freitag et al., 2021). Yet, it is machine learning which is according to first studies (de Vries, 2023) one of the biggest accelerators of ICT energy consumption and consequently ICT's claim of global GHG emissions. Where computing is powered by renewable energy this will not incur GHG emissions, however, also renewable energy is scarce and needs to be shared among all sectors of human activities.

Against this background, it is important to study ways of decreasing energy consumption of data processing operations. In the field of hardware, there is, according to Koomey's law (Koomey et al., 2010), an increase in the number of computations per Joule. However, this is not enough, because tasks are designed to need more computations to complete due to the confidence in the improvement of hardware (Verdecchia et al., 2017). For this reason, we need to look at possibilities for decreasing energy consumption from a digital perspective. This is particularly important when we consider the two-century-old Jevons' paradox, which warns that increased efficiency leads to increased consumption. The increased consumption often exceeds the initial saving of efficiency, and has at least partially been reconfirmed (Sorrell, 2009). In their research, Fouquet and Hippe (2022), analyse the speed of historical energy and communication technology transitions since 1850 which points to digitalisation outpacing the transition to renewable energy.

| Model accuracy

In the digital society progress in machine learning is appraised by looking at the accuracy of models performing specific tasks. With machine learning, we measure the performance of models in various ways. For binary predictions where an answer can be either yes (1) or no (0), prediction accuracy is often used as metric (Bratko, 1997). Accuracy can be defined as the number of correct predictions divided by the number of predictions made. Hence an accuracy of 0.9 means that in 9 out of 10 cases, the model will predict the right target value. To achieve a high accuracy, machine learning methods optimize the parameters of their model by computing a loss function. This loss function is defined over all actions, where the goal of the model is to reduce the loss to the minimal possible (Bishop and Nasrabadi, 2006). Using gradient updates, a model can reduce its loss by finding the optimal parameters. This is the learning part in machine learning: finding the optimal parameters to minimize loss and maximize the number of correct predictions.

1.2 | Technical trade-offs between environmental, privacy, and accuracy

Data privacy, energy consumption, and model accuracy are important goalposts for data processing operations. But these concepts interact and affect each other. When users explore these different goalposts, they often perceive trade-offs that may prompt them to optimize their use case for just one of these values, casting aside the remaining two. Oftentimes users make decisions based on their perceptions about how these trade-offs play out. Comparing the values for each scenario, we examine the trade-offs and their intersection. We consider several three scenarios:

- **Scenario 0, optimise all dimensions**: This is the ideal scenario, where all three goalposts are optimised simultaneously. The use of anonymisation techniques is oftentimes an escape valve from legislation such as the GDPR to obtain datasets that no longer contain personal data. In this scenario, users can choose an anonymisation technique that will produce the highest accuracy and consume the least energy. Although this field is still underexplored, recent work shows that k-anonymity can reduce the energy consumption of code as well as increase the accuracy in some cases (Oprescu et al., 2022). However, it is not clear to what extent k-anonymity can satisfy the GDPR requirements. A recent paper indicates that with k-anon the identification risk is about 0.50% (Abubakar et al., 2022), meaning that for an average of 200 individuals in the data set, 1 person could be re-identified.

- **Scenario 1, optimise for accuracy**: If datasets that have been pre-processed with anonymisation techniques are perceived to negatively affect the accuracy of a given machine-learning model, then the user may disregard using PETs to preserve data privacy. This calculus could be corroborated by the energy consumption of applying

anonymisation techniques to datasets. In this case, however, legislation on personal data protection such as the GDPR continues to apply which translates in compliance obligations and legal risks of non-compliance.

- **Scenario 2, optimise for energy**: When optimising for environmental considerations, the user would seek the least energy-consuming data processing techniques for the anonymisation of the dataset and subsequent model training. Using PETs should not be disregarded. It is also an option when the task at hand does not require the highest accuracy levels. A systemic review shows that the energy consumption of AI can often by reduced at least 50% (Verdecchia et al., 2023). In one specific dataset the energy consumption could be reduced by 77% with the accuracy decreasing from 94.3% to 93.2% (Brownlee et al., 2021) which is negligible for most use cases.

2 | EXPERIMENT

To delve into the trade-off, we set up an experiment to measure the impact of PETs on accuracy and the environment, with two leading research questions:

RQ1: Which privacy-enhancing technique yields the least loss of accuracy?

RQ2: Which privacy-enhancing technique is the least energy consuming?

In this section, we will first present the experimental setup, after which we present the results. Following these results, we link the findings to the presented trade-off.

2.1 | Experimental setup

We reuse the setup of related work (de Reus et al., 2023; Oprescu et al., 2022) to measure both the energy consumption and the impact on the accuracy for three machine learning techniques. We use two public datasets (Bache and Lichman, 2013): the Student Performance set to predict a fail or pass of a student and the Census Income set to predict whether a person's income is above or below 50 thousand dollars. By varying three different datasets - original data, anonymised data (using k-anonymity) or synthetic data - we can compare the impact of each value on the trade-off. As we keep the other variables the same, the privacy enhancement is the control variable in this experiment. The setup and code are provided online at the public code repository GitHub[1]. The results of this experiment give insights into the loss of accuracy (**RQ1**). Over this entire process we measure the energy consumption by reading the Running Average Power Limiters (RAPL) counters (see details below), which helps to answer our research question for the energy consumption (**RQ2**).

2.1.1 | RAPL counters

Computers with Intel processors from 2012 onward have energy sensors called RAPL in their Central Processing Units (CPU) (Desrochers et al., 2016; Hähnel et al., 2012). These RAPL sensors measure the energy consumption of the CPU and memory of the Intel chip. An advantage of RAPL is that the user can select which parts of the chip should be measured (García-Martín et al., 2019; Khan et al., 2018), for example solely the CPU or the memory energy consumption. Another option to measure the energy consumption would be to use a hardware-based method, also known as the power plug method. This is an external device between the socket and the used machine. Such a device is very accurate, but can measure only the entire energy input to the device and hence not divide between CPU or

[1] https://github.com/PepijndeReus/Privacy-Enhancing-ML

memory consumption. Apart from that, the devices are expensive and costly to use regarding set-up time. Hence, there are various methods to measure the energy consumption of code.

2.2 | Statistical significance

We can show the significance of the impact on energy consumption due to the addition of the code smell, proving that the samples from the original and from the smelly measurements are from different distributions for a specific problem written in a specific language. Additionally, we wish to know the direction of the impact.

The aim of a statistical test is to reject a null hypothesis in favor of an alternative hypothesis. By rejecting the null hypothesis we can claim that the alternative hypothesis is significantly true. For this purpose, we derive some statistics from our sample distribution called the p-value. This p-value is the probability of obtaining a result given that the null hypothesis is true. If this p-value is smaller than a predetermined significance level (α-value), we may reject the null hypothesis. The significance level indicates the probability of rejecting the null hypothesis when it is true. Typically, the significance level is set at 0.05.

2.2.1 | The Mann-Whitney U test

The Mann-Whitney U test (Mann and Whitney, 1947) looks at the probability that a sample from one distribution is greater than a sample from another distribution. If this probability is at 50%, the two samples belong to the same distribution. There are two Mann-Whitney U tests: the one-sided and two-sided tests. While the two-sided test has an alternative hypothesis claiming that the two samples are from different distributions, the one-sided test has an alternative hypothesis claiming that the first sample is from a greater distribution (Nachar et al., 2008).

To find the direction of the difference between the distributions, we must use the one-sided test twice. We use the greater one-sided Mann-Whitney U test to determine the probability that a sample from the first distribution is larger than a sample from the second distribution. The lesser one-sided Mann-Whitney U test determines the probability that a sample from the first distribution is smaller than a sample from the second distribution. If this p-value of the *greater* one-sided test is smaller than our significance level, we can say that the first distribution is significantly *smaller* than the second distribution, and vice versa for the lesser one-sided test.

2.3 | Technical results

In Table 1 (see next page), we present the results of the energy consumption for the Census Income dataset on the left and the Student Performance set on the right. The benchmark is presented on the first row, after which the results are shown for the models trained on either anonymised or synthetic data. We present the deviation in percentages to the benchmark for these methods. We use several exemplars of machine learning techniques: k-nearest neighbours (knn) – a non-parametric approach, logistic regression (LogReg) – a parametric approach, and Neural Network (NN) – a simple, feedforward neural network. Complementary with the energy consumption, we provide the results of the Mann-Whitney U test in Table 2. Finally, the idle energy consumption of the machine is measured. For each second of inactivity, 7.512 Joules are consumed by the machine, to be subtracted from the results. However, since we look at the trends this idle energy consumption is deemed negligible.

Table 3 shows the accuracy score of the machine learning techniques for each dataset and the privacy-enhancing technique. The benchmark reflects the accuracy scores for the unaltered or original data. To reflect the accuracies as exactly as possible we present the accuracy scores instead of the deviation to the benchmark. As k-anonymity is

obtained via generalisation and suppression, some data has been suppressed entirely. Table 4 shows the percentages of suppressed data for each dataset and respective k-value.

TABLE 1 The average run time and energy consumption for the Census Income and Student Performance datasets. On the rows, we define the input data and on the columns, we define each method, devised into run time and energy consumption. The benchmark resembles the unaltered data as obtained online, the percentages below show the deviation for each method to this benchmark.

Census Income dataset							Student Performance set						
	knn		LogReg		NN			knn		LogReg		NN	
Benchmark	8.19s	329.33J	1.41s	73.70J	9.85s	157.59J	Benchmark	0.04s	2.16J	0.06s	2.95J	3.40s	61.20J
k=3	-67%	-67%	-45%	-46%	-44%	-42%	k=3	-25%	-37%	-66%	-77%	-26%	-24%
k=10	-69%	-68%	-61%	-62%	-50%	-47%	k=10	-75%	-79%	-66%	-78%	-28%	-27%
k=27	-73%	-73%	-78%	-78%	-54%	-51%	k=27	-75%	-80%	-66%	-79%	-29%	-26%
Synthetic data	-5%	-4%	-9%	-10%	-1%	0%	Synthetic data	0%	+2%	-33%	+3%	-1%	-2%

TABLE 2 The results of the Mann-Whitney U test applied to the energy consumption presented in Table 1. The presented p-value shows the probability of the method on the row being greater than the alternative in the column.

Census Income dataset					Student Performance set				
	k=3	k=10	k=27	Synthetic data		k=3	k=10	k=27	Synthetic data
k=3	x	p <0.01	p <0.01	p >0.99	k=3	x	p <0.01	p <0.01	p >0.99
k=10	p >0.99	x	p <0.01	p >0.99	k=10	p >0.99	x	p <0.01	p >0.99
k=27	p >0.99	p >0.99	x	p >0.99	k=27	p >0.99	p >0.99	x	p >0.99
Synthetic	p <0.01	p <0.01	p <0.01	x	Synthetic	p <0.01	p <0.01	p <0.01	x

TABLE 3 The weighted accuracy over 10 measurements for each dataset, machine learning technique, and privacy-enhancing technique.

	Census Income dataset			Student Performance set		
	knn	LogReg	NN	knn	LogReg	NN
Benchmark	0.820	0.846	0.846	0.700	0.719	0.706
k=3	0.828	0.848	0.847	0.872	0.851	0.861
k=10	0.832	0.842	0.842	0.826	0.826	0.826
k=27	0.828	0.837	0.837	0.831	0.853	0.853
Synthetic data	0.802	0.828	0.828	0.728	0.756	0.755

TABLE 4 The percentage of suppressed data for the Census Income set and Student Performance with respect to the obtained k-anonymity. The percentage is defined as suppressed cells divided by the total amount of cells.

Suppressed data		
	Census Income set	Student Performance set
k=3	19%	74%
k=10	28%	80%
k=27	32%	84%

2.4 | Interpreting the results to the trade-off

The results of this experiment show that models trained on k-anonymised data consume less energy than models trained on the original data, with a similar accuracy to the unaltered data. Models trained on synthetic data have a similar energy consumption and a similar to lower accuracy compared to models trained on unaltered data.

These findings are interesting in the light of the presented trade-off. Apparently, choosing data privacy using either k-anonymity or synthetic data does not lead to a substantial reduction in accuracy. In fact, the accuracy might even improve. Similarly, machine learning models trained using PETs have a similar to lower energy consumption compared to models trained on unaltered data. As these findings form an interesting argument against the perceived

trade-off, we wish to translate these scientific results into findings suitable for non-computer science experts. And so follows our final research question: how can we translate these technical results to non-computer science experts?

3 | TRANSLATING THE FINDINGS TO NON-COMPUTER SCIENCE EXPERTS

Data can be hard to value, especially when comparing the value of data across different fields of research. For example, lawyers may prioritise compliance with regulations on data privacy over accuracy whereas data scientists might strive for the highest accuracy, but should not ignore legal compliance. To link these different valuations of the presented trade-off, we explore the economic contextualisation of our presented technical results using Slotin (2018) recommendation that when it comes to measuring the value of data often an impact-based valuation works best. The impact-based method demonstrates the relationship between data and outcomes in people's lives to determine how to value data. A typical impact-based approach is storytelling.

3.1 | Introduction to storytelling

Often a story is defined as something with a beginning, a middle, and end (Moezzi et al., 2017). In *storytelling* the emphasis is on the elicitation and construction of stories or narratives in situ. An important aspect is why the story is told in a particular context. Stories are crafted rather than pre-existing artefacts, and crafting depends on various contexts such as audience, purpose, location, etc. (Moezzi et al., 2017). Storytelling in science is important to translate scientific findings to policymakers, journalists, and the general public (Schimel, 2012). In his book on Writing Science, Schimel (2012, p 11) states: "The data are supporting actors in the story you tell. The lead actors are the questions and the larger issues you are addressing. The story grows from the data, but the data are not the story." This is why storytelling is a method that has the potential to explain the trade-offs between digital sustainability, data privacy, and model accuracy to a non-technical audience as they present themselves after our experiment.

3.2 | Stories translating the experiment

The first story offers a narrative for optimising accuracy (Section 1.2, Scenario 1). Suppose we have Alice, a data scientist working for a municipality. Alice is developing an algorithm to detect illegal holiday rentals[2]. In the current workflow, when a report is filed with the municipality that there may be a case of illegal renting at an address, the Surveillance & Enforcement employees could start an investigation. To help this team, Alice's algorithm should assist them in prioritising the reports so that their team's enforcement capacity can be used efficiently. Thus, Alice works with sensitive datasets, comparable to the datasets used in our experiment, for which she has to initially perform binary classification (whether it is indeed a case of illegal holiday rental). Basically, the dataset is anonymised before the model is trained, and then the model is used to indicate how likely it is that a report of illegal rental is correct. So, in itself, the model does not detect illegal rental, but gives advice to the human expert. Both k-anonymity (De Pascale et al., 2023) and synthetic data[3] have been used as anonymisation techniques in relation to fraud detection applications. Although the decision-making on the illegal rental case is not automated, and the municipality employee can still choose to ignore the output of the algorithm, the prediction task is high risk, as a visit from the authorities to one's household can be seen as very intrusive. In this story, both privacy and accuracy are important. Alice, who has

[2] https://algoritmeregister.amsterdam.nl/ai-system/illegal-holiday-rental-housing-risk/109

[3] https://www.kaggle.com/datasets/ealaxi/paysim1

a mathematics background, could model this decision-making problem as a Pareto-optimality problem (Kuosmanen and Kortelainen, 2005; Huppes and Ishikawa, 2005), and give some (equal) weights to the three dimensions of interest to come to a solution. She has a hard time however quantifying how many more times the privacy of a citizen is more important than the accuracy of an indication of fraud. Consequently, Alice hesitates between synthetic data and k=27-anonymity. For either of them, she would choose Logistic Regression as machine learning technique, since it delivers equally excellent accuracy as the Neural Network, but at a lower energy cost.

The second story offers a narrative about optimising energy (Section 1.2, Scenario 2). Suppose we have Bob, working as a digital enabler for an auditing company. Bob too has to perform binary classification on sensitive datasets, comparable to the datasets used in our experiment. Though the datasets are sensitive, the prediction task is low risk, in the sense that its accuracy is not crucial. Bob needs to decide what PET to use and what machine learning model to pick. Bob wants to optimise over energy and make an informed decision. Next, Bob is informed by our technical findings (de Reus et al., 2023), and he decides that he would like to use synthetic data to enhance the privacy of the subjects in the sensitive datasets. Moreover, he wishes to minimize his climate impact by picking the lowest energy consumption. Consequently, Bob picks synthetic data and Logistic Regression to predict the binary labels for his dataset.

| Limitations

We conducted the technical experiments on two structured datasets. This is a limitation in two ways: 1) using more datasets would allow to generalise findings; 2) structured data is often a utopia in data science and anonymisation methods cannot (yet) be applied to unstructured data. Moreoever, several sectors such as healthcare have special regulations and protocols with respect to personal privacy which have to be observed. Our study does not cover these regulated sectors. Moreover, k-anonymity is no guarantee that the data has been truly anonymised. Anonymisation is highly context-specific and needs to be controlled for on a case-by-case basis. Finally, our work is aimed at the GDPR and therefore a Western perspective. We acknowledge that a matured version of our work should cover universal aspects and especially the trade-offs in the Global Souths (Milan and Treré, 2019).

4 | CONCLUSION

With the EU backdrop of equal concern for protecting data privacy (GDPR) and reducing climate impact (Green Deal), it is becoming increasingly important to understand how society perceives the trade-off between these two dimensions and triangulate it with model accuracy. In this study, we make a first effort to provide insight into the trade-off between privacy and energy consumption in light of machine learning. We explore the question of what privacy-enhancing algorithms, such as synthetic data and k-anonymity, add to the energy consumption of regular machine-learning tasks. Using an experiment to answer this technical question, we then form an economic contextualisation through story-telling of our technical results. With these stories we aim to make technical results accessible to non-ICT experts and decision-makers, and as such aim to understand how society could perceive the trade-off between privacy and climate impact in the digital economy. We find that, in some contexts, such as the advertisement industry, the answers are clear-cut and simple. However, we also find that in some sensitive application domains, such as personal finance and healthcare, the legal compliance and ethical aspects require a different approach than the trade-off approach. To obtain answers in these contexts, we need to re-assess the way we typically evaluate the value of data. We provide suggestions for future work to close the gap between scientific results and policymakers.

| Future work

Policymakers from non-ICT backgrounds should be able to make informed decisions on the digital economy. To discover the scenarios relevant to them, we plan to extend our research to devise novel story-developing methods by questioning experts from various fields and stakeholders from various value chains. Using these methods, research should be aimed towards storytelling as a means to share scientific insights. Apart from that, details on the energy consumption of the digital economy are often unknown. In parallel, more research efforts should be focused on measuring the energy consumption of various AI techniques and data operations, on the impact of considering different data utility metrics, and on refining the granularity at which we measure anonymity and privacy (using tools such Anonymeter[4]. By expanding to more datasets, preferably from daily-use scenarios, we discover how the various trade-offs interact. Combining technical research with the right method to inform policymakers leads to informed decisions to reduce the impact of the digital economy on the environment.

Acknowledgements

We would like to express our gratitude towards Bart Krull and Sustainalab[5] for connecting us to the Climate Tech Cities community. We would also like to thank dr. Christine Erb from RDT for her continuous support of our ideas.

Conflict of Interest

Kristina Irion is a co-author but also one of the organisers.

Data Availability Statement

The data that support the findings of this study are openly available in UCI Machine Learning Repository at `http://www.archive.ics.uci.edu/`.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K. and Zhang, L. (2016) Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 308–318.

Abubakar, I. B., Yagnik, T. and Mohammed, K. (2022) Robustness of k-anonymization model in compliance with general data protection regulation. In *2022 5th International Conference on Computing and Big Data (ICCBD)*, 67–72. IEEE.

AEPD and EDPS (2020) Misunderstandings with regard to biometric identification and authentication. URL: `https://edps.europa.eu/sites/edp/files/publication/joint_paper_14_misunderstandings_with_regard_to_identification_and_auuthentication_en.pdf`.

Bache, K. and Lichman, M. (2013) UCI machine learning repository. URL: `http://archive.ics.uci.edu/ml`.

Bauer, P., Stevens, B. and Hazeleger, W. (2021) A digital twin of earth for the green transition. *Nature Climate Change*, **11**, 80–83.

Bishop, C. M. and Nasrabadi, N. M. (2006) *Pattern recognition and machine learning*, vol. 4. Springer.

[4] `https://github.com/statice/anonymeter`

[5] `https://sustainalab.nl`

Bradford, A. (2020) *The Brussels effect: How the European Union rules the world*. Oxford University Press, USA.

Bratko, I. (1997) Machine learning: Between accuracy and interpretability. In *Learning, networks and statistics*, 163–177. Springer.

Brownlee, A. E., Adair, J., Haraldsson, S. O. and Jabbo, J. (2021) Exploring the accuracy–energy trade-off in machine learning. In *2021 IEEE/ACM International Workshop on Genetic Improvement (GI)*, 11–18. IEEE.

Ciulli, F. and Kolk, A. (2023) International business, digital technologies and sustainable development: Connecting the dots. *Journal of World Business*, **58**, 101445.

Coche, E., Kolk, A. and Ocelík, V. (2023) Unravelling cross-country regulatory intricacies of data governance: the relevance of legal insights for digitalization and international business. *Journal of International Business Policy*, 1–16.

De Pascale, D., Cascavilla, G., Tamburri, D. A. and Van Den Heuvel, W.-J. (2023) Real-world k-anonymity applications: The kgen approach and its evaluation in fraudulent transactions. *Information Systems*, **115**, 102193.

de Reus, P., Oprescu, A. and van Elsen, K. (2023) Energy cost and machine learning accuracy impact of k-anonymisation and synthetic data techniques. In *ICT4S, June 2023, Rennes, France*, 57–65. International Conference on ICT for Sustainability.

Desrochers, S., Paradis, C. and Weaver, V. M. (2016) A validation of dram rapl power measurements. In *Proceedings of the Second International Symposium on Memory Systems*, 455–470.

European Commission (2020) Shaping Europe's digital future. *Fact Sheet*, **20**.

European Parliament and the Council (2016) Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (General Data Protection Regulation) 2016. *General Data Protection OJ L*, **119**.

Finck, M. and Pallas, F. (2020) They who must not be identified—distinguishing personal from non-personal data under the gdpr. *International Data Privacy Law*, **10**, 11–36.

Fouquet, R. and Hippe, R. (2022) Twin transitions of decarbonisation and digitalisation: A historical perspective on energy and information in european economies. *Energy Research & Social Science*, **91**, 102736.

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S. and Friday, A. (2021) The real climate and transformative impact of ict: A critique of estimates, trends, and regulations. *Patterns*, **2**.

García-Martín, E., Rodrigues, C. F., Riley, G. and Grahn, H. (2019) Estimation of energy consumption in machine learning. *Journal of Parallel and Distributed Computing*, **134**, 75–88.

Greenleaf, G. (2023) Global data privacy laws 2023: 162 national laws and 20 bills.

Hähnel, M., Döbel, B., Völp, M. and Härtig, H. (2012) Measuring energy consumption for short code paths using rapl. *ACM SIGMETRICS Performance Evaluation Review*, **40**, 13–17.

Hernandez, M., Epelde, G., Alberdi, A., Cilla, R. and Rankin, D. (2022) Synthetic data generation for tabular health records: A systematic review. *Neurocomputing*.

Hittmeir, M., Ekelhart, A. and Mayer, R. (2019) On the utility of synthetic data: An empirical evaluation on machine learning tasks. In *Proceedings of the 14th International Conference on Availability, Reliability and Security*, 1–6.

Huppes, G. and Ishikawa, M. (2005) A framework for quantified eco-efficiency analysis. *Journal of industrial ecology*, **9**, 25–41.

Khan, K. N., Hirki, M., Niemi, T., Nurminen, J. K. and Ou, Z. (2018) Rapl in action: Experiences in using rapl for power measurements. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)*, **3**, 1–26.

Koomey, J., Berard, S., Sanchez, M. and Wong, H. (2010) Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing*, **33**, 46–54.

Kuosmanen, T. and Kortelainen, M. (2005) Measuring eco-efficiency of production with data envelopment analysis. *Journal of industrial ecology*, **9**, 59–72.

Mann, H. B. and Whitney, D. R. (1947) On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, 50–60.

Mayer-Schönberger, V. and Cukier, K. (2013) *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.

Milan, S. and Treré, E. (2019) Big data from the south (s): Beyond data universalism. *Television & New Media*, **20**, 319–335.

Moezzi, M., Janda, K. B. and Rotmann, S. (2017) Using stories, narratives, and storytelling in energy and climate change research. *Energy Research & Social Science*, **31**, 1–10.

Murtaza, H., Ahmed, M., Khan, N. F., Murtaza, G., Zafar, S. and Bano, A. (2023) Synthetic data generation: State of the art in health care domain. *Computer Science Review*, **48**, 100546.

Mytton, D. and Ashtine, M. (2022) Sources of data center energy estimates: A comprehensive review. *Joule*.

Nachar, N. et al. (2008) The mann-whitney u: A test for assessing whether two independent samples come from the same distribution. *Tutorials in quantitative Methods for Psychology*, **4**, 13–20.

Oostveen, M. and Irion, K. (2018) The golden age of personal data: How to regulate an enabling fundamental right? *Personal Data in Competition, Consumer Protection and Intellectual Property Law: Towards a Holistic Approach?*, 7–26.

Oprescu, A., Misdorp, S. and van Elsen, K. (2022) Energy cost and accuracy impact of k-anonymity. In *ICT4S, June 2022, Plovdiv, Bulgaria*. International Conference on ICT for Sustainability.

Samarati, P. (2001) Protecting respondents identities in microdata release. *IEEE transactions on Knowledge and Data Engineering*, **13**, 1010–1027.

Samarati, P. and Sweeney, L. (1998) Generalizing data to provide anonymity when disclosing information. In *PODS*, vol. 98, 10–1145.

Schimel, J. (2012) *Writing science: how to write papers that get cited and proposals that get funded*. Oxford University Press USA.

Schneider, S. H. (1989) The greenhouse effect: science and policy. *Science*, **243**, 771–781.

Slotin, J. (2018) What do we know about the value of data. *May 3rd*.

Sorrell, S. (2009) Jevons' paradox revisited: The evidence for backfire from improved energy efficiency. *Energy policy*, **37**, 1456–1469.

Sweeney, L. (2002) k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, **10**, 557–570.

Verdecchia, R., Procaccianti, G., Malavolta, I., Lago, P. and Koedijk, J. (2017) Estimating energy impact of software releases and deployment strategies: the KPMG case study. In *Proceedings of the 11th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, 257–266. IEEE Press.

Verdecchia, R., Sallou, J. and Cruz, L. (2023) A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, e1507.

de Vries, A. (2023) The growing energy footprint of artificial intelligence. *Joule*.