

Scoring and Gamification to Encourage Sustainable Use of Compute Clusters

Maximilian MacDonald, Chris McCaig, Sean MacAvaney, Matthew Barr, Lauritz Thamsen
University of Glasgow, United Kingdom — {firstname.lastname}@glasgow.ac.uk

Abstract

The environmental cost of computing continues to grow, yet behaviour change remains limited. We present a composite sustainability score integrating average carbon intensity, resource utilisation, and embodied emissions into a single 0–100 metric designed for gamified feedback. Each component rewards a different dimension of sustainable behaviour: carbon-aware workload shifting, high resource utilisation, and selecting hardware that is commonly underutilised. This scoring system is built into an existing cluster management interface and underpins three dashboard conditions: raw metrics, composite score, and a gamified tree visualisation, which we are planning to evaluate in a 12-week within-subjects study with approximately 35 researchers. Furthermore, we open the discussion on the challenge of defining computational work ‘goodness’ in the context of sustainability scores.

Keywords

Low Carbon Computing, Behaviour Change, Persuasive Systems, Operational Emissions, Embodied Carbon, Cluster Computing

1 Introduction

The environmental impact of the computing sector is a growing concern, with its emissions now rivalling industries such as aviation [6]. Cluster and cloud computing energy use continues to grow at alarming rates, with data centres accounting for around 180 million tons of CO_2 per year from energy consumption alone and projected to grow to over 300 million tons by 2030¹. Machine Learning (ML) in particular has come under scrutiny for its resource consumption, leading to calls for systematic reporting of energy and carbon metrics [11].

While awareness of computing’s environmental impact is growing [23], awareness alone appears insufficient to drive meaningful behaviour change [14]. Many computing professionals are unaware of mitigation strategies [23], and even when informed, cost, time, and a lack of accepted standards limit behaviour change [25]. This awareness-action gap represents a critical barrier to achieving sustainability goals in computing, and so tools to promote behaviour change are needed.

Current approaches to computational sustainability focus primarily on measurement and reporting. Tools like CodeCarbon [18] and the Green Algorithms calculator [15] enable researchers to estimate carbon footprints of individual jobs, while cloud providers have introduced carbon dashboards showing aggregate emissions^{2 3}.

¹<https://www.iea.org/reports/energy-and-ai>

²<https://docs.cloud.google.com/carbon-footprint/docs/methodology>

³<https://learn.microsoft.com/en-us/power-bi/connect-data/service-connect-to-emissions-impact-dashboard>

Meanwhile, gamification techniques, including feedback mechanisms, goal-setting, and social comparison, have demonstrated success in related domains, with residential energy consumption studies achieving up to 20% reductions [12].

Standard energy and carbon metrics lack intuitive meaning and fail to provide actionable guidance—users see numbers but do not know what constitutes “good” or how to improve, a common problem with purely numerical metrics [16]. Single metrics like Carbon Intensity, while useful for specific optimisations such as temporal workload shifting, fail to capture the multidimensional nature of sustainable computing, which includes energy efficiency, embodied carbon, and total lifecycle use, among others. Optimising for one dimension such as embodied emissions can lead to inefficiency in another like operational emissions, highlighting the need for approaches that balance competing sustainability goals rather than optimising metrics [1, 9]. Additionally, gamification has not been applied to behaviour change in sustainability computing, a domain where users are technically proficient and receptive to optimisation challenges when provided with clear goals.

We address these challenges through a composite sustainability score that integrates multiple dimensions: average carbon intensity, resource utilisation, and hardware lifecycle considerations. We combine this into a single 0–100 scoring metric suitable for gamified feedback. We present our design for a 12-week study on a mid-sized ML research cluster with approximately 35 active users to test the effectiveness of this scoring metric. Participants rotate through three interface conditions: raw metrics, composite score, and a gamified tree visualisation. We aim to measure both attitudinal changes and behavioural outcomes to evaluate whether gamified sustainability feedback drives more sustainable computing practices.

Contributions. We present the following contributions:

- We design a gamifiable composite sustainability score that integrates average carbon intensity, resource utilisation, and a hardware bonus from underutilised resources, enabling multidimensional feedback suitable for gamification.
- We present our design for an 12-week study on a mid-sized research compute cluster that focuses on ML research.
- We open the discussion on evaluation of computational work “goodness”, and how this challenge goes beyond resource efficiency.

2 Related Work

2.1 Current Sustainability Metrics

Beyond energy consumption and raw carbon usage, metrics have been developed to capture the various dimensions of computational sustainability. Carbon Intensity is the most widely adopted of these and has shown its usefulness through temporal workload shifting, enabling users to schedule jobs during periods of low grid intensity [26] with tools like CASPER [24] reducing the carbon

footprint of workloads by up to 70%. However, Carbon Intensity ignores resource utilisation efficiency, embodied carbon, and hardware lifecycle considerations, providing little actionable guidance for users with deadline-driven work who cannot defer jobs. Embodied carbon, the scope 3 emissions associated with hardware manufacturing, represents another critical dimension. Standard accounting methods spread manufacturing emissions over projected hardware lifespans, with Zhang et al. [28] demonstrating significant variation in embodied carbon calculations depending on spatial-temporal modelling assumptions. Gupta et al. [8] quantified that hardware production often equals or exceeds lifetime operational emissions, establishing embodied carbon as a sustainability lever comparable in importance to operational efficiency. Resource utilisation likewise influences sustainability outcomes. Barroso et al. [3] highlight that computing hardware exhibits significant baseline power draw when at low utilisation. While advancements in energy proportionality have resulted in an almost linear relationship after the initial idle cost, higher utilisation still spreads both operational energy and embodied carbon across more useful work, reducing per-task environmental impact. Some tools have made use of multidimensional sustainability metrics. For example Acun et al. [1] developed Carbon Explorer, a framework integrating operational carbon, embodied carbon, and hardware lifecycle considerations for datacenter design. Sustainability dashboards, such as the Green Algorithms Dashboard [15] and Kepler [2], similarly focus on communicating energy and carbon costs, but offer little insight into actionable behaviour change, and are generally designed for cluster administrators rather than individual users. These multidimensional approaches reveal inherent tensions. Hanafy et al. [9] demonstrated that optimising for carbon efficiency can lead to energy inefficiency, while Lee et al. [17] showed that as energy grids become cleaner, the relative importance of embodied carbon increases. This remains the case even when older hardware consumes more energy per computation, and particularly when continuing to use it avoids the need to acquire new hardware. These competing priorities highlight the challenge of designing single metrics that capture sustainability's multidimensional nature. Additionally, no prior work has combined these dimensions into a single metric suitable for guiding individual job submission decisions.

2.2 Gamification for Sustainability

Gamification, defined as adding game-like elements to non-game contexts [5], has been applied to residential energy consumption, transportation, and waste reduction to positive effects. Iria et al. [12] achieved energy reductions through dashboards combining real-time feedback with social comparison, while a gamified garden representation focusing on aesthetic feedback reduced user energy consumption by 0.5 standard deviations in an office context [21].

Koivisto and Hamari [13] conducted a systematic review of 800+ gamification studies, finding that successful interventions combine extrinsic and intrinsic motivational elements, balance short-term rewards with long-term goals, and include mechanics that encourage real-world action. Meanwhile, Nicholson [20] defines Meaningful Gamification and highlights the importance of designing the gamified system with the user at its centre and allowing them to engage

with the system on their own terms. This builds on ideas of autonomy, competence, and relatedness presented in Self Determination Theory to drive behaviour change [22]. Scoring systems form a core component of gamification, providing clear targets for improvement and enabling social comparison. However, Hanus and Fox [10] demonstrated that poorly designed gamification can backfire, undermining intrinsic motivation. In sustainability contexts specifically, metrics must balance intrinsic motivations, such as environmental concern, with extrinsic rewards, like points [19]. Poorly designed systems can encourage gaming behaviours, where users appear sustainable without genuine change, or create inequitable competition disadvantaging users with legitimate high-resource needs. Related work on eco-feedback technology provides design principles for sustainability interfaces. Froehlich et al. [7] established that effective eco-feedback requires granular data, immediate feedback, and clear action-outcome relationships. Wu and Lee [27] found that interactivity and visualisation enhance learning and promote sustained behaviour change beyond passive information exposure. Despite the solid body of work on gamification for promoting sustainable behaviour no work has applied these concepts to a computing context, or applied scoring methods to sustainability computing metrics.

3 Ongoing Work

In this section we present the design of an upcoming 12-week deployment study investigating whether exposing researchers to multi-dimensional sustainability feedback can encourage more sustainable resource usage on a mid-sized ML research cluster through evaluating three dashboard interface conditions with varying levels of gamification. We then present both the composite sustainability scoring metric underpinning the intervention and the gamification techniques through which it is communicated to users. As the deployment is scheduled to begin in the coming weeks, this section focuses on the design, implementation details, and planned evaluation methodology underlying the intervention.

3.1 Study Design

The study is embedded within the daily workflows of research cluster users at the University of Glasgow School of Computing Science and employs a mixed-methods, sequential explanatory design with approximately 35 active users, in which quantitative behavioural data inform subsequent qualitative interviews. The intervention is integrated into Launcher, the custom web application cluster members use to monitor resources and deploy containerised jobs via an OKD Kubernetes interface. This integration ensures access to users' natural work environment without requiring adoption of new tools. Operational carbon emissions are calculated dynamically using real and estimated hardware power alongside real-time regional grid intensity data from the UK National Grid Carbon Intensity API⁴. GPU power consumption is measured directly via NVIDIA System Management Interface, while CPU power is estimated using a linear model to Thermal Design Power (TDP) based on utilisation. While 0-to-TDP power models are not necessarily accurate for low resource utilisation, given static power draw, they are widely available, including for our heterogeneous hardware

⁴<https://www.carbonintensity.org.uk/>

configurations. Resource utilisation of CPU, GPU, and memory is measured directly from cluster nodes at 5-minute intervals.

3.1.1 Dashboard Visualisation Configurations. Participants are exposed to three distinct dashboard configurations designed to test the effectiveness of different sustainability communication techniques. Participants rotate through three dashboard configurations in a randomised order, each presented for 3 weeks to control for temporal and learning effects, alongside 3 weeks of a control with no change to the dashboard.

- (1) **Raw Metrics Baseline (Fig.1a):** This configuration integrates fundamental, quantitative metrics directly into the user's dashboard. It displays raw energy consumption in kilowatt-hours, carbon footprint in grams of CO₂ equivalent, and the average carbon intensity over a period.
- (2) **Composite Sustainability Score (Fig.1b):** This view condenses multi-dimensional sustainability data into a single score between 0 and 100. The score balances average carbon intensity with efficient hardware utilisation across CPU, GPU, and memory allocations, whilst offering a "hardware bonus" to incentivise underutilised infrastructure. Break-downs of the score are supplied alongside the overall score.
- (3) **Gamified Representation (Fig.1c):** This configuration is still under development but aims to apply interface aesthetics and narrative themes to communicate ecological costs and provide feedback. The composite score is mapped to an interactive forest visualisation, where trees represent specific pods, with the colour and liveliness of each corresponding to hardware bonus. The weather of the background reflects the users average carbon intensity, while ground colour and gradients from each tree show utilisation contributions. Dynamic elements aim to reinforce user relatedness, and tooltips provide access to underlying pod contribution metrics for detailed inspection.

3.1.2 Surveys and Qualitative Interviews. Upon consent via in-dashboard recruitment, participants will complete a baseline survey measuring awareness of sustainable computing concepts, attitudes toward sustainability responsibility, current practices, and perceived barriers. Participants are then randomly assigned to one of six condition sequences to counterbalance order effects. Post-intervention surveys repeat key measures and add questions on score interpretation and condition preferences. Throughout the study, job submission patterns, resource allocation, utilisation, hardware selection, and dashboard engagement are continuously logged to characterise behavioural trends across conditions, providing exploratory context for the qualitative interviews. These post-study semi-structured interviews with 10–15 participants will be analysed using reflexive thematic analysis [4] to explore score interpretation, decision making, and barriers to sustainable practice.

3.2 Composite Sustainability Score Design

We designed a gamifiable composite sustainability score to provide multidimensional feedback suitable for gamification. Each component is normalised to a 0–1 scale and weighted.

3.2.1 Resource Utilisation. We measure utilisation across four dimensions: GPU usage, CPU usage, GPU memory, and CPU memory.

For each job, we calculate the average utilisation of each resource type over its reserved duration. Each utilisation dimension is scored linearly, with 0% utilisation receiving a score of 0 and 100% utilisation receiving a score of 1. In cases of over-utilisation, where actual usage exceeds requested resources, we cap the score at 1 rather than penalising users. While this creates potential for gaming the system through deliberate under-allocation or artificial over-utilisation, we assume good faith and that resources were genuinely needed. The overall score is the average of the four resource dimensions:

$$U_{\text{score}} = \frac{1}{4} (u_{\text{GPU}} + u_{\text{CPU}} + u_{\text{GPU-mem}} + u_{\text{CPU-mem}}) \quad (1)$$

where each component $u_r = \min(1.0, (\text{actual}_r / \text{allocated}_r))$ represents the utilisation efficiency for resource r .

3.2.2 Hardware Bonus. Our hardware bonus incentivises use of underutilised resources, particularly older or less popular hardware configurations. We calculate the average utilisation of each GPU type across the cluster over a rolling one-week window. The hardware bonus awards points inversely proportional to recent utilisation: the least-used GPU type receives the maximum bonus, while fully utilised GPUs receive no bonus. Specifically, for a job using GPU type g :

$$H_{\text{bonus}}(g) = \frac{1 - U_g}{1 - U_{\min}} \quad (2)$$

where U_g is the cluster-wide utilisation of GPU type g over the past week and $U_{\min}$ is the utilisation of the least-used GPU type. This ensures that the greatest incentives exists for the most underutilised resources while maintaining some reward for high but not fully utilised hardware.

3.2.3 Average Carbon Intensity (ACI). Our ACI component evaluates how well a job's timing aligns with available low-carbon periods within a realistic scheduling window. For each pod, we compare its actual ACI to the theoretically optimal intensity achievable within a constrained window. The scheduling window is limited to one week or twice the job's duration, whichever is highest. For a job of duration n timestamps, we identify the optimal scheduling window by finding the contiguous block of at least n consecutive timestamps with the lowest average carbon intensity, and the worst-case block as the contiguous sequence with the highest. The actual ACI is then mapped linearly between these bounds:

$$\text{ACI}_{\text{score}} = 1 - \frac{\text{ACI}_{\text{actual}} - \text{ACI}_{\text{optimal}}}{\text{ACI}_{\text{worst}} - \text{ACI}_{\text{optimal}}} \quad (3)$$

A score of 1 represents optimal timing and 0 represents the worst possible timing. As such users score highly by actively choosing to run their workloads at times with relatively low ACI. To aggregate across multiple jobs at the user or namespace level, we compute a time-weighted average.

3.2.4 Composite Score. The final sustainability score combines and weights these components:

$$S_{\text{total}} = w_U \cdot U_{\text{score}} + w_H \cdot H_{\text{bonus}} + w_{\text{ACI}} \cdot \text{ACI}_{\text{score}} \quad (4)$$

where w_U , w_{ACI} and w_H are weights satisfying $w_U + w_H + w_{\text{ACI}} = 1$. Initial weights are set to $w_U = 0.4$, $w_H = 0.4$, and $w_{\text{ACI}} = 0.2$, reflecting that users have the most direct control over resource utilisation




Figure 1: Screenshots of the three different dashboard configurations

and hardware selection, and least control over grid carbon intensity. We expect to refine both the scoring mechanisms and gamification methods discussed below based on the feedback and insights from the upcoming study.

3.3 Gamification Methods

Making use of this composite scoring metric, the gamified dashboard view aims to incorporate gamification through feedback, aesthetic design, and lightweight narrative elements. Rather than presenting sustainability metrics in isolation, user actions will be reflected through evolving visual states within the interface. This approach aims to create a stronger connection between user behaviour and system response. By linking visual progression directly to user action, the system seeks to support relatedness and intrinsic motivation. The visual representation takes the form of an interactive forest, whose appearance encodes multiple dimensions of the sustainability score simultaneously. Each tree represents a single pod within the selected time frame, with the colour of each tree's leaves corresponding to the hardware bonus of the pod, going from a vibrant green when healthy to brown/red when the hardware bonus is low. Additionally, each tree has flowers that first die, then fall from the tree as the hardware bonus decreases. Utilisation is represented by the ground colour, with the average utilisation over the time frame shown as the average ground colour and individual pod contributions shown through a gradient around the base of each tree. The weather of the background changes to reflect the average carbon intensity, going from sunny and clear, to overcast, to stormy with a red glow on the horizon as the score declines. The size of

each tree shows the individual contribution of each pod. The visual language is intentionally aligned with the sustainability domain, using environmentally-associated imagery and progressive visual states to reinforce the meaning of user actions through thematic congruence. Dynamic elements are intended to contribute towards a sense of immersion: leaves bob when hovered and clicked, clouds float across the scene with speed relative to the current weather, and fallen flowers can be cleared away by the user.

4 Open Challenges

Our current scoring method evaluates workloads by their carbon efficiency, resource utilisation, and hardware choice, assuming that the work being run is always valuable. However, the problem of defining 'good work' remains an open challenge. Self-evidently, the most sustainable option is to avoid running work that is not useful, for example work that leads to no significant improvement. However, determining whether work is indeed useless is difficult, particularly in a research context where trial and error and experimentation are fundamental. One approach to quantifying useful or good work could be to encourage user reflection through self-evaluation. For example, asking users to rate the usefulness of a job after it completes, which may also prompt more deliberate consideration before launching it. Conversely, care must be taken not to discourage potentially useful work. For example, running large ML tasks with minor parameter tweaks when model performance is already comparably high is arguably wasteful, but the line between exploration and inefficiency is not always clear. Perhaps the more productive framing is to encourage mindful experimentation,

rather than attempting to penalise it. More specifically there is the question of energy efficiency: how much energy is really needed to achieve a certain result? And, even if energy use is reduced is the additional time spent to increase efficiency bought from even more unsustainable sources, e.g. carbon costs associated with housing, transport and consumable goods required by the user. Even for factors that are directly measurable, like code efficiency, there are often multiple approaches to reach a result, with possibly different energy consumptions, but how optimal any specific computational load is, can be hard to capture. Ultimately, the attempt to quantify the goodness of work may be inherently hard and, hence, of limited practicality, and behaviour change in this respect may need to be culturally motivated rather than metric-driven. Hence, this work does not attempt to resolve this question, but raises it as a necessary consideration for future work on our sustainability scoring system.

Failed jobs present a particular challenge. They consume energy and hardware lifetime without producing useful output, yet determining whether a failure stems from user mistakes, such as under-provisioning memory, or infrastructure factors like hardware faults, is non-trivial. Penalising users for failures outside of their control risks discouraging legitimate resource use, while ignoring failed jobs entirely removes incentive to submit well-configured workloads. As such developing reliable fault attribution remains an open challenge.

5 Summary and Future Work

This paper presents a gamified scoring mechanism to capture usage of cluster resources in a way that can be used in future gamification practices. It aims to reflect costs related to carbon efficiency, resource utilisation, and hardware choices, and in doing so open a pathway to incentivise sustainable computing behaviour. We show how this scoring system is composed and our design for an upcoming study to test the effect of this method on user attitudes and behaviour.

Future iterations of the gamified dashboard will incorporate additional gamification features which may include social and self-comparison mechanisms, where users compete against their own historical scores, collaborate with teammates to meet shared challenges, or compete with peers for higher rankings. We also aim to introduce longitudinal progress tracking to support long-term behaviour change, and to investigate how more explicit narrative structures affect user engagement. Finally, we wish to examine what effect genuinely engaging gamification has on user motivation and sustained behaviour change beyond informational feedback.

References

- [1] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 2023. Carbon explorer: A holistic framework for designing carbon aware datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*.
- [2] Marcelo Amaral, Huamin Chen, Tatsuhiro Chiba, Rina Nakazawa, Sunyanan Choochothaew, Eun Kyung Lee, and Tamar Eilam. 2023. Kepler: A Framework to Calculate the Energy Consumption of Containerized Applications. In *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*.
- [3] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan. 2019. Energy and Power Efficiency. In *The Datacenter as a Computer: Designing Warehouse-Scale Machines*. Springer International Publishing, 99–127.
- [4] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* (2019).
- [5] Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. 2011. From game design elements to gamefulness: defining “gamification”. In *Proceedings of the 15th international academic MindTrek conference: Envisioning future media environments*.
- [6] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. 2021. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns* (2021).
- [7] Jon Froehlich, Leah Findlater, and James Landay. 2010. The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery.
- [8] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2022. Chasing Carbon: The Elusive Environmental Footprint of Computing. *IEEE Micro* (2022).
- [9] Walid A. Hanafy, Roozbeh Bostandoost, Noman Bashir, David Irwin, Mohammad Hajiesmaili, and Prashant Shenoy. 2023. The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. Association for Computing Machinery.
- [10] Michael D. Hanus and Jesse Fox. 2015. Assessing the effects of gamification in the classroom: A longitudinal study on intrinsic motivation, social comparison, satisfaction, effort, and academic performance. *Computers & Education* (2015).
- [11] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of machine learning research* (2020).
- [12] José Iria, Nuno Fonseca, Fernando Cassola, António Barbosa, Filipe Soares, António Coelho, and Aydogan Ozdemir. 2020. A gamification platform to foster energy efficiency in office buildings. *Energy and Buildings* (2020).
- [13] Jonna Koivisto and Juho Hamari. 2019. The rise of motivational information systems: A review of gamification research. *International Journal of Information Management* (2019).
- [14] Ramesh Kumar Bagla, Prachi Trivedi, and Teena Bagga. 2022. Awareness and adoption of green computing in India. *Sustainable Computing: Informatics and Systems* (2022).
- [15] Loïc Lannelongue, Jason Grealey, and Michael Inouye. 2021. Green Algorithms: Quantifying the Carbon Footprint of Computation. *Advanced Science* (2021).
- [16] Richard Larrick and Jack Soll. 2008. Economics - The MPG illusion. *Science* (New York, N.Y.) (2008).
- [17] Benjamin C Lee, David Brooks, Arthur van Benthem, Udit Gupta, Gage Hills, Vincent Liu, Benjamin Pierce, Christopher Stewart, Emma Strubell, Gu-Yeon Wei, et al. 2024. Carbon connect: An ecosystem for sustainable computing. *arXiv preprint arXiv:2405.13858* (2024).
- [18] Kadan Lottick, Silvia Susai, Sorelle A. Friedler, and Jonathan P. Wilson. 2019. Energy Usage Reports: Environmental awareness as part of algorithmic accountability. *arXiv preprint arXiv:1911.08354* (2019).
- [19] Elisa D. Mekler, Florian Brühlmann, Alexandre N. Tuch, and Klaus Opwis. 2017. Towards understanding the effects of individual gamification elements on intrinsic motivation and performance. *Computers in Human Behavior* (2017).
- [20] Scott Nicholson. 2015. *Gamification in Education and Business*. Springer International Publishing.
- [21] Divinus Oppong-Tawiah, Jane Webster, Sandy Staples, Ann-Frances Cameron, Ana Ortiz de Guinea, and Tam Y. Hung. 2020. Developing a gamified mobile application to encourage sustainable energy use in the office. *Journal of Business Research* (2020).
- [22] R. M. Ryan and E. L. Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *The American Psychologist* (2000).
- [23] Muhammad Salam and Siffat Ullah Khan. 2018. Challenges in the development of green and sustainable software for software multisourcing vendors: Findings from a systematic literature review and industrial survey. *Journal of Software: Evolution and Process* (2018).
- [24] Abel Souza, Shruti Jasoria, Basundhara Chakrabarty, Alexander Bridgwater, Axel Lundberg, Filip Skogh, Ahmed Ali-Eldin, David Irwin, and Prashant Shenoy. 2024. CASPER: Carbon-Aware Scheduling and Provisioning for Distributed Web Services. In *Proceedings of the 14th International Green and Sustainable Computing Conference*. New York, NY, USA.
- [25] Gopalakrishnan Sriraman and Shriram Raghunathan. 2023. A Systems Thinking Approach to Improve Sustainability in Software Engineering—A Grounded Capability Maturity Framework. *Sustainability* (2023).
- [26] Philipp Wiesner, Ilja Behnke, Dominik Scheinert, Kordian Gontarska, and Lauritz Thamsen. 2021. Let’s wait awhile: how temporal workload shifting can reduce carbon emissions in the cloud. In *Proceedings of the 22nd International Middleware Conference*.
- [27] Jason S. Wu and Joey J. Lee. 2015. Climate change games as tools for education and engagement. *Nature Climate Change* (2015).
- [28] Xiaoyang Zhang, Yijie Yang, and Dan Wang. 2024. Spatial-Temporal Embodied Carbon Models for the Embodied Carbon Accounting of Computer Systems. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. New York, NY, USA.