

Project Carbon Budgeting: Digital Sobriety, Gen AI, and Time for Constraint?

Since last year we have been involved in a major study investigating artificial intelligence and its generative AI (Gen AI) variant, the perceived impact on the work of information technology professionals, and resultant implications for computing education. The findings from a review of the academic literature and a set of interviews of some 47 IT professionals in New Zealand and Sweden have now been reported in an ITiCSE 2024 working group report [4,5]. So, this column reflects on a highlighted gap in the practitioners' perspectives, specifically on the sustainability of Gen AI as a technology: "Very few professionals mentioned the implications of GenAI on the environment and sustainability, despite the growing concerns associated with its use." [5] Then, first elaborating here on some of our concerns, and our responsibilities as members of professional societies, what are some actions we might feel obligated to take as computing educators?

Given our work, I was looking forward to the keynote at the Frontiers in Education Conference, on Oct. 16, 2024, titled *AI for Education: The Journey Is Just Getting Started* by Chris Daugherty, education strategy lead at Google. The keynote mooted the transformative impact of AI on education and the need for a strategy so the technology might be effectively and safely adopted. After an interesting talk in which Chris outlined some technology-led possibilities, I asked him a question about the apparently excessive power demands of Gen AI technology. While Chris acknowledged the concerns, small-scale nuclear plants were



deemed to be his novel "green computing" answer. But he did not convincingly answer a question about the passing on of incurred clean-up responsibilities, whether Google would still be in existence and take responsibility until the lengthy, and in some cases many thousand years, half-life of the radiation waste would expire. (I must acknowledge here of course my own bias in coming from nuclear-free New Zealand!

"Some isotopes decay in hours or even minutes, but others

decay very slowly. Strontium-90 and cesium-137 have half-lives of about 30 years (half the radioactivity will decay in 30 years). Plutonium-239 has a half-life of 24,000 years." [30]

So, what do we make of this rosy picture of Gen AI and the new possibilities for education painted by Google? Is this simply a revisiting of the hype surrounding the "Virtual University" movement evident at the turn of the millennium. When the death of traditional universities

was projected, an enthusiasm criticized for "the naive technological progressivism which has typified much of the commentary on the virtual campus." [10]

Is the now touted "AI revolution" simply an inevitable step in the march of technological progress we must respond to. Or with *Medium* commentator Will Lockett arguing to the contrary, are we taking a step in the wrong direction with the AI revolution? He makes several trenchant critiques, but is supported in part by others, for example cf. Chaudhary and Penn [1], Chien et al. [2], Helmore and Paul [11], and Pogrebna [24].


"This is a nonsense technology that isn't profitable, has no genuinely useful applications, is entirely based on stealing copyrighted work and individuals' labour, and uses so much energy that it threatens to derail our pathetic efforts to save the planet." [17]

But interestingly, and more positively, he does argue in favor of nuclear power. He sees it as an answer to his predictions that by 2026, AI would use power equivalent to the current energy use of Germany, and see big tech companies become some of the worst polluters on the planet:

"Nuclear power only causes, on average, 90 deaths per 1,000TWh (which includes the Chernobyl and Fukushima incidents), compared to wind power's 150, solar power's 440, natural gas's 4,000, and coal power's 100,000 deaths! This is because these other energy sources require more mining, ... as well as fossil fuel energy producing deadly emission..." [17]

So, in another response, we see reported that Microsoft similarly is not concerned with the risks and is now following the nuclear route:

"A nuclear reactor at the notorious Three Mile Island site in Pennsylvania is to be activated for the first time in five years after its owners, Constellation Energy, struck a deal to provide power to



Source:GAO, based on the Department of Energy documentation. | GAO-15-652

Microsoft's proliferating artificial intelligence operations.

"The plant was the location of the most serious nuclear meltdown and radiation leak in U.S. history, in March 1979 when the loss of water coolant through a faulty valve caused the Unit 2 reactor to overheat. More than four decades later, the reactor is still in a decommissioning phase... As part of the agreement, Three Mile Island will also be renamed the Crane Clean Energy Center." [19]

A notable element here is Microsoft's 'greenwashing' marketing strategy. But the road to GenAI nirvana will not necessarily run so smoothly, with Amazon recently thwarted in its nuclear ambitions.

"Amazon has hit a roadblock in its plans for nuclear-powered U.S. datacenters. Federal regulators rejected a deal that would let it draw more power from a Susquehanna plant to supply new bit barns next to the site, on the grounds this would set a precedent which may affect grid reliability and increase energy costs." [26]

Similarly, concerns over excessive water usage for datacenter cooling have been reported. [15]

"As Google fought to keep its water use figures secret, it found itself facing a growing public relations nightmare...When it shared its water-use figures for The Dalles... in just five years, Google's water use in the city had tripled. Its facilities used 355 million gallons of water in 2021, which was the equivalent of 29% of all the water used in the city that year. And it still wanted more to cool additional facilities. As one resident told *The Oregonian*, 'Google's become a water vampire.'" [20]

So, returning to Chris Daugherty's advocacy for Google, positively some new small modular reactor (SMR) designs are said to partly mitigate the issues I have raised around nuclear waste [12]. But Chris also failed to address some of the risks for Google. Long project durations, (for example, Rolls Royce's projected 2032-2035 SMR project completion [13]), project overruns, skill shortages, and novel technology risks for SMR's, have been highlighted in

the academic literature and news [12,13,27,28]. Strike, for instance has reported that, “The global average construction time for a 1 GW conventional reactor is about 100 months, with more extreme cases like Angra 2 in Brazil taking about 300 months or 25 years” [28], hardly a quick fix for technology companies striving urgently to meet carbon emission reduction targets!

Therefore, legitimate concerns about the risks and environmental impacts of Gen AI grow. So, societal kickback against negative impacts on citizens will begin to have increasing effect. As Marx has commented: “Those fights could become central to a wider campaign to reassert collective power and sovereignty over technology.” [20]

For us as computing educators then, and members of ACM as a professional society, what obligations do these aspects of the “AI Revolution” impose on us? In two of the codes of ethics ACM has subscribed to, we see wording relating to promoting the public good such as the following.

“1.1 Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing.

This principle, which concerns the quality of life of all people, affirms an obligation of computing professionals, both individually and collectively, to use their skills for the benefit of society, its members, and the environment surrounding them.

In addition to a safe social environment, human well-being requires a safe natural environment. Therefore, computing professionals should promote environmental sustainability both locally and globally.” [8]

“Principle 2: Public Software engineering shall,

in their professional role, act only in ways consistent with the public safety, health, and welfare. In particular, software engineers shall:

2.01 Disclose to appropriate persons or authorities any actual or potential danger to the user, a third party, or the environment, they reasonably believe to be associated with the software or related documents for which they are responsible, or merely know about.” [9]

As computing educators then, acting in accordance with these obligations, **to whom do we disclose the danger to the environment of an enthusiastic and uninformed adoption of Gen AI in our teaching**, by our students, colleagues and institutions?

In a recent *Communications of the ACM* column, former editor-in-chief Moshe Vardi has lamented, “By and large, Big Tech workers do not seem to be asking themselves hard questions, I believe, hence my conclusion that we do indeed suffer from an ethics crisis.” [32] So let us pause, instead of rushing lemming-like to adopt the newest shiny thing in the AI revolution. What hard questions do we need to ask ourselves? Given the flexibility we are privileged to exercise in our pedagogical practices, should we simply ban the use of this fashionable but profligate new technology? Or should we adopt a strategy of more conscious and measured use in what might be termed *digital sobriety*, “a lifestyle that involves not just consuming better but also, and critically, consuming less.” [23] Cooper and colleagues [6] argue likewise that New Zealand is failing to consider the ethical and environmental impacts of its data-center investments and needs a more digitally sober strategy.

Yet as Perea [23] observes, “Decreasing IT use with respect to digital sobriety seems to be a

As computing educators then, acting in accordance with these obligations—to whom do we disclose the danger to the environment of an enthusiastic and uninformed adoption of GenAI in our teaching, by our students, colleagues and institutions?

neglected topic” and criticizes the Green IT literature, which merely highlights “the potential of IT for greening the world.” That perspective fails to, significantly, recognize “the harmful consequences of IT uses.” Of these two possibilities, “IT as a solution for greening the world” and IT being held partly responsible for ecological disasters,” digital sobriety aligns with the latter. The Big Technology companies by contrast, as evangelists for the AI revolution, and through the promise of nuclear technology solutions, align with the first perspective (even if by greenwashing!).

AI researchers likewise seem to be captured by concerns about achieving higher accuracy of the technology, while research into computationally efficient algorithms has lagged. Some 3.5% of a set of research papers addressed efficiency, and a *computational sustainability* track at *AAAI 2019* comprised 0.4% of the papers as reported by Strubell and colleagues. They recommended “a concerted effort by industry and academia to promote research and development of more computationally efficient algorithms, as well as hardware that requires less energy.” [29]

Interestingly, a research paper developing a taxonomy of AI practitioners’ views of ethics in AI [22] seems to echo our findings on practitioner perspectives and the absence

of sustainability concerns [5]. While the researchers have reported that “Australia’s AI Ethics Principles address a broad spectrum of ethical concerns, spanning from human to environmental well-being” [22], the derived taxonomy from their research reporting practitioner perspectives does not include a sustainability aspect. Perhaps the question of agency at the practitioner level and the reported role of budgets is informative here:

“[A] budget can be a challenge in developing ethical AI based systems. If anybody wants us to do additional testing, which requires additional data gathering or labelling of existing data, right now we don’t have any budget set aside for that, so we need to proactively plan.” [22]

Which brings us to the challenge for this article. If we have a concern about the environmental impact of a project assigned to use Gen AI, how can we measure that impact, and can we reduce or limit it in some way? In earlier work on sustainable educational practice and green teaching with Mikey Goldweber and William Liu [16], we advocated a role for critical reflection on the environmental costs of each pedagogical decision. These include:

1. Choice of textbook or alternative, assignment submission options, and so on
2. The options and cost, for example, “A search followed by a site visit is virtually twice as expensive as utilizing a bookmark. When considering the mantra of reduce-reuse-recycle, browser bookmarks fall under the reuse category.” [16]

Researchers and practitioners are investigating metrics, risks, and costs for the use of AI and Gen AI technologies. For instance, “Lacoste

As an educator, if I were able to set a project carbon budget for each of my assignments and students, I might feel more comfortable in working within known limits rather than the open envelope of today’s GenAI terrain, with its greedy ‘evermore data for open access’ bargain ...

et al. (2019) recently published an online calculator that provides geographically-aware estimates of effective carbon emissions for users of Google, Amazon and Microsoft cloud resources. We strongly encourage ML researchers to analyze, audit and report the carbon footprint of their research using this valuable tool” [14], and the Machine Learning Emissions Calculator [7]. Also work by Luo and Sun [18], Mittal [21], and Vahdat et al. [31] has proposed relevant metrics. Weidinger and colleagues [32] investigating the risks of large language models identified as one risk area: “Environmental and socioeconomic harms”, with a set of proposed mitigations. “Technical approaches to reducing risks of environmental harm include segmenting LMs into less large LMs that search and retrieve information from a distinct data corpus ... Other work targets efficiency gains during training and inference.” [32] Techniques such as pruning, distillation, or fine-tuning might be adopted.

But perhaps the potential energy efficiencies provided by algorithmic innovations through software rather than brute-force hardware solutions will increasingly gain more sway. Interesting developments such as DeepSeek are causing consternation among U.S. technology companies and investors. Reports suggest that “DeepSeek R1 rivals the top

large language models (LLMs) such as ChatGPT-01 produced by giant AI companies, but R1’s base LLM V3 (announced a few weeks earlier) was built with a small fraction of the cost and used a much smaller number of low-grade computer chips (GPUs) than the current top LLMs.” [25]

So, while work is underway and some estimates have been provided of carbon emissions of LLM usage, in work such as reported by Mittal [21], it is acknowledged that:

“While AI like ChatGPT revolutionizes our interactions and tasks, it’s essential to recognize the hidden environmental costs associated with these tools. By being aware of our usage and supporting initiatives that prioritize sustainability, we can harness the benefits of AI without compromising the health of our planet.

“So, the next time you fire off a question to ChatGPT, remember that every little bit counts, and together, we can make a positive impact!”

So, in a call for greater digital sobriety with Gen AI, we need more research into strategies, tools, and techniques to measure the cost of our interactions, at a level of granularity that will enable us to set limits. As an educator, if I were able to set a project carbon budget for each of my assignments and students, I might feel more comfortable working within known limits rather than the open envelope of today’d GenAI terrain, with its greedy ‘evermore data for open access’ bargain [3]. In the meantime, even such proxy measures of work incurred for each assignment or project, (such as a limit on number of prompts and answers reported or in permissible response dialogues, as reported in audit logs of Gen AI project sessions), might be an option. At this stage, I cannot claim to know what will be the or even an optimal strategy. We have two master’s-level students currently investigating the applicable metrics and potential software solutions.

In many ways, this imposition of limits will reflect a back-to-the-future stage for the profligacy of modern computing (where everything has been *free*), to one where the planet's finite limits are acknowledged—where compute resources are limited and have costs which need to be estimated, measured, and tracked, so we can incorporate a project budget as a mechanism for *sober* use of computing in our student projects. ♦



TONY CLEAR

Dept. of Computer and Information Sciences
Auckland University of Technology
Private Bag 92006, Auckland
1010 New Zealand
Tony.Clear@aut.ac.nz

REFERENCES:

- Chaudhary, Y. and Penn, J. Beware the intention economy: Collection and commodification of intent via large language models. *Harvard Data Science Rev.*, Special Issue 5 (2024).
- Chien, A. A. et al. Reducing the carbon impact of generative AI inference (today and in 2035). In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM (2023), 1-7; <https://doi.org/10.1145/3604930.3605705>.
- Clear, T. Thinking issues: Large language models, the 'doctrine of discovery' and 'terra nullius' declared again? *ACM Inroads* 15, 2 (2024), 6-9; <https://doi.org/10.1145/3638564>.
- Clear, T. et al. A plan for a joint study into the impacts of AI on professional competencies of IT professionals and implications for computing students. In *Proceedings of the 2024 ACM Conf. on Innovation and Technology in Computer Science Education*. (2024) 2; <https://doi.org/https://doi.org/10.1145/3649405.3659527>.
- Clear, T. et al. AI integration in the IT professional workplace: A scoping review and interview study with implications for education and professional competencies. In *Proceedings of the 2024 ACM Conf. on Innovation and Technology in Computer Science Education*.
- Cooper, N. and Turnbull, A. NZ is encouraging the use of AI—but it's largely outsourcing the risks and societal costs. *The Conversation* (Oct. 22, 2024). <https://theconversation.com/nz-is-encouraging-the-use-of-ai-but-its-largely-outsourcing-the-risks-and-societal-costs-240666>.
- Courty, B. et al (2024). *mlco2/codecarbon: v2.7.1 (v2.7.1)*; <https://doi.org/https://doi.org/10.5281/zenodo.13744486>.
- Gotterbarn, D. et al. ACM code of ethics: A guide for positive action. *Commun. ACM* 61 (2017), 121-128.
- Gotterbarn, D., Miller, K., and Rogerson, S. Software engineering code of ethics. *Commun. ACM* 40, 11, (1997), 110-118.
- Harris, M. Virtual learning and the network society. *Information, Communication & Society*, 3, 4 (2000), 580-596.
- Helmore, E. and Paul, K. *New York Times* sues OpenAI and Microsoft for copyright infringement. *The Guardian*, (Dec. 28, 2023); <https://www.theguardian.com/media/2023/dec/27/new-york-times-openai-microsoft-lawsuit>.
- Hussein, E.M. Emerging small modular nuclear power reactors: A critical review. *Physics Open* 5 (2020), 100038.
- Jolly, J. 'A viable business': Rolls-Royce banking on success of small modular reactors. *The Guardian* (Jan. 15, 2025); <https://www.theguardian.com/environment/2025/jan/15/a-viable-business-rolls-royce-banking-on-success-of-small-modular-reactors>.
- Lacoste, A., Luccioni, A., Schmidt, V., and Dandres, T. Quantifying the carbon emissions of machine learning. (2019), *arXiv preprint arXiv:1910.09700*.
- Li, P., Yang, J., Islam, M.A., and Ren, S. Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models. (2023), *arXiv preprint arXiv:2304.03271*.
- Liu, W., Goldweber, M., and Clear, T. A definition for a sustainable computing educator. In *6th Ann. Conf. of Computing and Information Technology Research and Education*. (2015); https://citrenz.org.nz/citrenz/conferences/2015/pdf/2015CITREnz_1_Liu_SustainableEduc_v3.pdf. Accessed 2025 Jan 15.
- Lockett, W. The pathetic reality at the heart of the AI revolution. *Predict*. (Nov. 3, 2024); <https://medium.com/predict/the-pathetic-reality-at-the-heart-of-the-ai-revolution-c7fe6c603376>.
- Luo, H. and Sun, W. Addition is all you need for energy-efficient language models. *arXiv preprint arXiv:2410.00907* (2024).
- Luscombe, R. Three Mile Island nuclear reactor to restart to power Microsoft AI operations. *The Guardian*. (Sept. 9, 2024); <https://www.theguardian.com/environment/2024/sep/20/three-mile-island-nuclear-plant-reopen-microsoft>.
- Marx, P. Data vampires: Going hyperscale (Ep. 1). *Tech Won't Save Us* (Oct. 10, 2024); https://techwontsave.us/episode/241_data_vampires_going_hyperscale_episode_1.
- Mittal, A. ChatGPT: How much does each query contribute to carbon emissions? *OnGraph Technologies Limited* (Oct. 8, 2024); <https://www.linkedin.com/pulse/chatgpt-how-much-does-each-query-contribute-carbon-emissions-mittal-wjff8c/>.
- Pant, A., Hoda, R., Tantithamthavorn, C., and Turhan, B. Ethics in AI through the practitioner's view: A grounded theory literature review. *Empirical Software Engineering* 29, 3 (2024), 67.
- Pérea, C., Gérard, J., and de Benedittis, J. Digital sobriety: From awareness of the negative impacts of IT usages to degrowth technology at work. *Technological Forecasting and Social Change* 194 (2023), 122670.
- Pogrebna, G. AI is a multi-billion dollar industry. It's underpinned by an invisible and exploited workforce. *The Conversation* (Oct. 9, 2024); <https://theconversation.com/ai-is-a-multi-billion-dollar-industry-its-underpinned-by-an-invisible-and-exploited-workforce-240568>.
- Poo, M. (Ed.), Reflections on DeepSeek's breakthrough, *National Science Rev.* 12, 3 (2025).
- Robinson, D. Amazon's nuclear datacenter dreams stall as watchdog rejects power deal—Federal Energy Regulatory Commission cites grid stability concerns. *The Register* (Nov. 4, 2024); https://www.theregister.com/2024/11/04/aws_nuclear_datacenter_ferc/.
- Stewart, W.R. and Shirvan, K. Capital cost estimation for advanced nuclear power plants. *Renewable and Sustainable Energy Reviews*, 155 (2022), 111880.
- Strike, C. *Examining the Feasibility of Small Modular Nuclear Reactors as the Leading Global Source of Energy Generation*. (2021), Master of Science in Energy Policy & Climate, Johns Hopkins University, Baltimore, Maryland.
- Strubell, E., Ganesh, A., and McCallum, A. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI Conf. on Artificial Intelligence* 39, 4 (2020), 13693-13696.
- U.S. Nuclear Regulatory Commission. *Backgrounder on Radioactive Waste*. (Jan. 2024); <https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/radwaste.html>.
- Vahdat, A., Ma, X., and Patterson, D. New computer evaluation metrics for a changing world. *Commun. ACM* 17, 10 (2024), 31-33.
- Vardi, M.Y. I was wrong about the ethics crisis. *Commun. ACM* 68, 1 (2024), 5.
- Weidinger, L. et al. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conf. on Fairness, Accountability, and Transparency*, 214-229.

DOI: 10.1145/3725890

© 2025 Copyright held by the owner/author(s).