Notes for Teaching sustainably in the Age of AI

Slide 1
Artificial Intelligence began in June 1956 long before many of us were born, yet when the AI chatbot ChatGPT appeared at the end of 2022, it seemed to have come out of nowhere. But the reality was that AI's apparent overnight success was a long time in the making. AI has been part of your life for decades – you will be familiar with chatbots helping perform mundane tasks such as mobile services and banking. Some of you will have heard of the Chess Master Gary Kasparov's loss to IBM's Deep Blue in the late 1990's. Not many of us paid attention to when Musk, Altman with others launched OpenAI in 2015. But when OpenAI launched ChatGPT in late 2022, a million users joined within a week, and this forced us to take notice.

Slide 2
My lecture today is going to challenge the narratives of weightless digital technology and explain to you the environmental impacts of Generative AI. I focus on AI through a sustainability lens and will help you to think about responsible digital practice – thereby developing your digital literacy. My field of expertise is initial teacher education and technology in education; I train teachers in the field of STEM and I research use of digital technology in education.

Slide 3
So let's start with UNESCO's Education for Sustainable Development (ESD) framework which moves us beyond just 'green skills' in using technology such as AI to making ethical judgement and developing professional responsibility. I want my students to know about AI but I want them to be reflective in their use and make sustainable pedagogical choices. ESD for 2030 frames sustainability education around five core priority areas and a set of key competencies.

Slide 4
Digital tools are framed as clean and efficient with terms such as 'the Cloud' making it seem weightless. The reality is nothing like this.
The people who benefit from AI are rarely the ones who pay the price for its physical existence. We see a clean interface on our laptops, while the actual heat, noise, and massive water evaporation occur in data centres often located in rural or marginalized regions where resources are cheaper. Over the next set of slide we are going to look at **spatial displacement** referring to impacts such as water depletion, mining damage, energy use occurring far from where AI is designed, sold, or used, often affecting rural or less powerful communities.

Slide 5
So what is AI? It is a machine that can use data to learn reasoning, problem-solving, recognising patterns (in images and language), and making decisions.
As a Large Language Model, it is trained on very large datasets to recognise and generate patterns in human language. In essence, it predicts the most likely next token (word or part-word) in a context. You see examples of this at play when you use predictive text in your messaging. If you type 'app' you will get a choice of word depending upon the context. If you

type ‘I want an app’ – you will get either apple or appointment and if you type ‘I am writing an app’ – you will get appeal or application.
There are two distinct phases – training and inference which are part of the AI lifecycle:

Slide 6
When we interact with Generative AI, we are conditioned to focus solely on the **"Use" phase**—the seconds it takes for a chatbot to generate a poem or a line of code. However, a true **Lifecycle Analysis (LCA)** reveals that the environmental impact begins long before a server is turned on and lasts long after it is decommissioned.
Extraction: The story begins in the earth, with the aggressive mining of lithium, cobalt, and rare earth minerals required for high-performance Graphics Processing Units or GPUs. These are specialised processors that are designed to perform many calculations in parallel.
Production: These materials are refined in high-heat, high-energy industrial processes to create the silicon wafers and hardware that power the "Cloud."
Use: This is the visible tip of the iceberg—the massive electricity and water consumption required for model training and daily inference. Water is used both in the electricity generation happening upstream as well as for the cooling required on-site.
Disposal: Because AI hardware becomes obsolete at an accelerated rate, we are left with a mounting crisis of toxic e-waste.

Slide 7
This is an open-pit mining site in northern Chile. The copper being extracted here is used in electricity grids, data centres and AI hardware.

Slide 8
Carrying on with our aim of dismantling the weightless myth – when we think of AI, we think about writing a prompt and getting an immediate response on the screen. But through a lifecycle LCA lens, we see that the environmental debt of the tool is accumulated even before a user types their first prompt and will carry on accumulating through the operation of the system and beyond once the hardware becomes obsolete.
We will look at these in more detail over the next six slides.

Slide 9
Training GPT-3 was like turning on the lights in 120 houses and leaving them on for an entire year.
And as we move toward more 'real-time' AI low latency systems where there is very little delay between a request and the response – such as in live translation, voice assistants or classroom tools that respond instantly - this isn't a one-off burst; it's becoming a continuous drain.

Slide 10
We also have to consider **Spatial Displacement**. Data Centres aren't invisible; they are physical landmarks. By placing them in regions with low-cost power—which often means high-carbon power—we are essentially shifting the environmental burden away from the end-user and onto the communities living near the grid. More about this later.

Slide 11

This image refers to local water shortages linked to the construction of large AI data centres. In Newton County, Georgia, USA. Residents with private wells near the new facility reported that their water taps ran dry after construction began, prompting concerns about groundwater impacts and broader water stress linked to large data centres in the region
The broader issue is the **social and environmental impact of AI infrastructure**, where water-intensive data centres are often placed near rural or less politically powerful communities, shifting environmental costs away from end users and technology companies.

Slide 12
This slide represents the ultimate Spatial Displacement.
The NVIDIA H100 is described as a marvel of engineering, but it is a physical object that contains minerals like Cobalt and Lithium. The extraction of these minerals happens far away from the labs of Silicon Valley.
In the DRC, cobalt extraction relies heavily on artisanal and small-scale miners, many of whom work in hazardous and precarious conditions that pose serious risks to health and life.
n Chile, particularly in the **Atacama region**, lithium extraction requires large volumes of water to process brine and ore. Mining companies often obtain water rights that allow them to **divert or extract groundwater** that has long supported local agriculture.
Although this is sometimes framed as water being temporarily "borrowed", in practice it can **lower water tables, reduce soil moisture, and disrupt traditional farming and Indigenous livelihoods**. In an already arid environment, this redistribution of water prioritises global battery and AI supply chains over local food security and ecological balance.

Slide 13
To understand AI's true impact, we must look beyond the electricity bill of the data centre.
Emissions are commonly grouped into three **scopes**:

- **Scope 1: Direct emissions**
  Emissions from sources owned or controlled by the technology company, such as on-site diesel backup generators at data centres.
- **Scope 2: Indirect energy emissions**
  Emissions arising from the generation of purchased electricity used to power and cool AI servers.
- **Scope 3: Indirect value-chain emissions**
  The often hidden footprint. This includes emissions from across the AI lifecycle: mineral extraction for GPUs, hardware manufacturing, global transport, and end-of-life disposal. For AI systems, the **largest share of environmental impact frequently sits in Scope 3**.

Slide 14
We have traced the **lifecycle of AI** from mineral extraction to data centres. The final stage is the **hardware graveyard**. As AI models demand more GPUs and specialised hardware to run larger, more complex models, today's hardware has a short operational lifespan. In effect, we are building highly specialised, rapidly obsolete supercomputers.
This is where **spatial displacement** becomes most visible. Decommissioned chips and servers do not vanish; they are frequently exported to the Global South as electronic waste. Here, the lead, mercury, and arsenic contaminate soil and water in communities that have rarely derived any benefit from the technology itself.

Slide 15
We've looked at the **physical footprint** of AI: the mines, the data centres, the energy and water systems that make it possible. We now turn to the **human footprint**—the people whose labour, land, health, and livelihoods are embedded within these systems.
This section focuses on how the costs of AI are unevenly distributed, often falling on communities far removed from the places where AI is designed, deployed, and consumed.

Slide 16
To understand the 'weight' of the cloud, we must look at where that weight actually lands. This is the heart of **Spatial and Social Displacement**. While innovation and advanced services are concentrated in our classrooms and tech hubs in the North, the physical scars of that technology—the open-pit mines and toxic scrap heaps—are borne elsewhere.
As educators, we face an ethical risk: if we teach AI only as a tool for efficiency without acknowledging its global cost, we are reproducing extractive practices. True social justice education requires us to pull back the curtain on this supply chain, ensuring students understand that their digital tools have a human and environmental footprint that spans the entire globe

Slide 17
This slide highlights the **hidden labour that underpins AI systems**. Despite being presented as autonomous and automated, AI relies on extensive human work to label data, moderate content, and correct errors. Much of this labour is outsourced to lower-income regions, where workers undertake repetitive and emotionally demanding tasks for low pay. The slide invites us to question who benefits from AI efficiency, and whose labour and wellbeing are rendered invisible in the process.

Slide 18
Earlier, we earlier discussed energy use, water demand, hardware turnover, and hidden labour in the use of AI**.** A common response to these impacts is the claim that AI will become *more efficient* and therefore more sustainable. Jevons' Paradox challenges this assumption. As AI systems become cheaper and easier to run, they are used more frequently and, in more contexts, which can increase total resource consumption rather than reduce it. This helps explain why gains in efficiency alone are unlikely to resolve the environmental and human costs we have traced across the AI lifecycle.

Slide 19
The educators' dilemma is that we are conflicted between preparing our students for an AI-driven world and in our roles as stewards of the environment striving to protect the environmental and human systems already under strain.

Slide 20

The challenge is not whether to engage with AI, but how. Pedagogical benefits such as efficiency, personalisation, and accessibility must be weighed against environmental costs, hidden labour, and long-term ecological harm.
Convenience can no longer be treated as neutral, and innovation cannot be separated from responsibility.
This framing invites a shift from uncritical adoption towards stewardship: using AI deliberately, sparingly, and transparently, while helping our students understand the broader consequences of technological choices.
In this sense, I believe that AI literacy and planetary stewardship should be taught together, not in opposition.

Slide 21
As we near the end of our audit, we have to look at the stories we tell ourselves to justify AI adoption. We often hear that 'technology is inevitable' or that we must adopt it because 'students are already using it'.
These are **Passive Adoption Narratives**. They frame AI as a neutral, unstoppable force, much like the weather. But as we have uncovered today, AI is not a neutral fact of life—it is a physical, resource-heavy infrastructure that relies on extraction, displaced labour, and massive energy consumption.
When we say 'technology is inevitable,' we are actually making a choice to ignore these material costs.

Slide 22
Our goal as educators is to move beyond these passive stories. We must recognise that every time we integrate a new digital tool, we are making a choice—not just about pedagogy, but about environmental stewardship and global social justice.
Let's stop treating adoption as a requirement and start treating it as a responsibility.

Slide 23
We have spent the last hour looking at the massive, global costs of AI—the mining, the thirst, the carbon. You might be sitting there thinking, 'But I only use it to draft an occasional lesson plan. Surely that doesn't matter'.
This is why scale matters. In isolation, your prompt is trivial. But educational adoption is never in isolation. It happens at a population level. When a teacher integrates AI, they aren't just using a tool; they are a **multiplier**. They are modelling a lifestyle and a professional standard for every student in the room.
If we ignore the scale, we ignore the reality. Those billions of 'tiny' decisions are exactly what drive the massive e-waste piles and water shortages we discussed earlier. We are now at a point where our small daily habits have become large systemic effects. The question is: knowing this scale, how do we move toward **active, ethical stewardship?**

Slide 24
We have spent this session uncovering the heavy material reality of the 'Cloud'. Now, we must translate that knowledge into a professional framework for Auditing AI Use in Education.

This is not about being 'anti-AI'; it is about being pro-stewardship. Every time we consider a new tool, we must run it through this diagnostic filter. Is the use necessary? Does it add genuine pedagogical value that justifies its energy and water consumption?
Crucially, we must move beyond ourselves to ask: Who bears those costs?. We cannot ignore the spatial and social displacement that allows our digital innovation to exist. As architects of future practice, your role is to model a form of digital citizenship that values global equity as much as technical efficiency. Let's ensure that when we integrate AI, we aren't just 'using' it, but actively auditing its place in a sustainable future

Slide 25
Explain the custom of 'dry January'. Here I want to introduce the idea of **digital sobriety**: a deliberate and reflective approach to technology use that prioritises purpose over novelty. Digital sobriety does not mean rejecting AI but using it only where it adds clear pedagogical value, rather than as a default solution.
This involves favouring lower-carbon alternatives where they exist, and resisting the automation of tasks that are central to human learning, such as thinking, drafting, and sense-making. Crucially, the practices we model in education shape long-term professional norms. How students learn to use AI now will influence how they use it throughout their careers, with lasting implications for both society and the planet.

Slide 26 **HIDE if necessary**
While we have spent much of this lecture discussing what *you* can do as an educator, the reality of the climate crisis requires more than individual action. It requires Institutional Change.
We need to move from individual guilt to collective responsibility. This starts with how our universities and schools buy technology. If we include environmental impact in our procurement contracts, we force the multi-billion dollar AI industry to change its practices. Imagine an institutional policy where we only license tools that use 100% renewable energy or tools that provide 'Right-Sized' models for student use. By embedding these sustainability criteria into our decision-making at the highest level, we ensure that our students are learning in an environment that truly values their future. Our schools shouldn't just teach sustainability; they should be built on.

Slide 27: Summary and key takeaway

Slide 28:
Individually, review these four proposed steps to digital sobriety and select one or two actions that feel realistic within your own study or teaching practice.