

Leading Sustainability and Regeneration in Projects¹

The Weight of the Cloud: AI Data Centres and the Project Manager's Role in Balancing Digital Ambition with Environmental Reality²



Figure 1 Data Centres for improving our world

Abstract

The rapid expansion of artificial intelligence has created a tension at the heart of project management: the digital tools that promise to reduce carbon emissions across the built environment, transport, and energy systems are themselves anchored in data centres with substantial environmental footprints. This article examines the physical reality behind the metaphor of 'the Cloud,' from the extraction of over sixty chemical elements in a single

¹The author Dr. Hugo Minney is a Fellow of APM (Association for Project Management), a Member of PMI and PMI UK, Chair of APM's Benefits and Value IN and Sustainability IN, founder of APM Nuclear Industries IN and AI & Data Analytics IN, and chair of BSI's Working Group on Benefits Management. Minney Org Ltd offers consultancy on ROI and SROI. For more, see his author profile at the end of this article

² How to cite this work: Minney, H. (2026). The Weight of the Cloud, Leading Sustainability and Regeneration in Projects, series article, *PM World Journal*, Volume XV, Issue VI, June.

smartphone to the consumption of an estimated 415 TWh of electricity by data centres globally in 2024 — a figure that grew by 17% in 2025 and is projected to double by 2030.

Against these costs, the article assesses technology's enabling effects: AI-driven materials discovery for energy storage, building energy optimisation delivering measurable savings, and the paradox of AI systems now optimising the cooling of the data centres that house them. The Jevons Paradox — whereby efficiency gains drive increased total consumption — is identified as the critical risk that project governance must address.

Practical responses are proposed for the project management community: 4D Stewardship of data centre sites as long-lived material banks; colocation of Small Modular Reactors within existing Tier 4 security perimeters; a 'Two-Step Revolution' for reliable government AI deployment, illustrated through an immigration case study; and Quintuple Bottom Line stage-gate controls. The biological benchmark of the human brain — performing complex parallel processing on 20 watts — serves as a measure of the efficiency gap and a reminder that technological ambition must be governed by physical reality.

This is the 9th in a series on Leading Sustainability and Regeneration in Projects.

Introduction: The Concrete Reality of the “Cloud”

The rapid adoption of Artificial Intelligence (AI) and high-performance computing has created a tension between profit and the environment (between operational goals and planetary boundaries) ([Raworth 2022](#); [Minney 2025d](#)). To the average user and to many project teams, technology is perceived as “in the Cloud” – a term that suggests a weightless, consequence-free existence. But every “Cloud” is bolted to a building site somewhere on the planet.

The world is entering what IEA describes as the “age of electricity” ([IEA 2026a](#)). It is straightforward to distribute, increasingly generated where it is used (through rooftop solar and community energy schemes), and we are getting better at storing it. Everything (apart from some specialist requirements such as steel production and aircraft) is moving to electricity from a variety of other fuels.

Table 1 Data Centre Electricity Demand – Current and Projected

Region / Metric	2024	2025	2030 (projected)
Global DC (Data Centre) electricity (TWh)	~415	~485 (+17%)	~945 (double)
AI-focused DC share	—	Outpacing DC demand growth	Triple by 2030
UK DC electricity (TWh)	7.6		20–41
UK share of GB demand	~2.5%		~7–14%
Global share of electricity	~1.5%		~3%
Sources: (IEA 2025 ; NationalGridDSO 2025 ; NESO 2025 ; IEA 2026b , 2026a)			

This article is about data centres. The information technology sector currently consumes approximately 1.5% of global electricity, around 415 TWh in 2024, and this figure grew by 17% in 2025 alone ([IEA 2025](#), [2026b](#)). In the United Kingdom, data centres consumed an estimated 7.6 TWh in 2024 – approximately 2.5% of total GB electricity demand (roughly 300 TWh total) – yet NESO’s Future Energy Scenarios project this rising to between 20 and 41 TWh by 2035 ([NationalGridDSO 2025](#); [NESO 2025](#)). This infrastructure mirrors the ‘Iceberg’ profile observed in other sectors, where the visible interface masks a massive foundation of embodied carbon and resource extraction ([Minney 2025b](#)). In the United Kingdom, projected growth rate threatens to exceed the capacity of current grid upgrades ([NationalGridDSO 2025](#)).

The Physicality of Progress: Stewardship and Urban Mining

Manufacturing modern electronics carries real costs — to communities near mines, to waterways, soils and waste disposal, and to the atmosphere. A single smartphone contains over 60 different chemical elements, including lithium, cobalt, tantalum, and rare earth metals such as neodymium ([Bloodworth 2014](#); [Cordella, Alfieri, and Sanfelix 2021](#)). For most devices, the manufacturing phase – from mineral extraction through to component assembly – accounts for a larger share of lifetime greenhouse gas emissions than the energy consumed during several years of daily use. Ericsson’s lifecycle assessment found that production contributes approximately 50 kg CO_{2e} out of a total 57 kg CO_{2e} over the manufacture plus year 1, with the use phase adding only 7 kg CO_{2e} annually ([Ericsson 2020](#)). What this means in practice is that the decisions made before manufacturing begins — which suppliers, which materials, which recovery pathways — often matter more than anything that happens during the product’s working life ([Minney 2025d](#), [2025b](#)).

Proportionality and Lifecycle Debt

The “digital debt” becomes clearer when compared to established benchmarks in the built environment and transport.

- **Built Environment:** The construction sector is responsible for approximately 40% of annual global CO₂ emissions ([UNEP 2025](#)). In the UK, embodied carbon (the emissions from extraction and manufacturing) remains a largely unaddressed liability compared to operational heating and lighting ([UNEP 2025](#)).
- **Transport:** Lifecycle assessments for vehicles show that for internal combustion engines (ICE), the “use phase” dominates. For electric vehicles, the manufacturing phase represents a proportionally larger share, often 30–40% of total lifecycle emissions; because the use-phase emissions of an EV are so much lower per kilometre (around 1/3 that of an ICE even assuming all of the electricity is generated by fossil fuels) ([Cox Automotive 2025](#); [Recharged 2025](#)).
- **Data Centres:** the IEA estimates that servers and other equipment (disk arrays, networking, connectivity) account for approximately 60% of total facility electricity consumption, with cooling taking between 7% and 30% depending on efficiency ([IEA 2025](#)). However, when embodied carbon from construction (concrete, steel) and from

manufacturing the IT hardware itself is included, the technology inside the building can represent the majority of the total lifecycle footprint — a pattern that intensifies as hardware refresh cycles shorten ([Whitehead, Andrews, and Shah 2015](#)).

- **Operational Power Split:** In a facility with a Power Usage Effectiveness (PUE) of 1.6—a common industry average—for every 1 watt of power used for “compute” (the actual processing), 0.6 watts are consumed by the environment-maintaining infrastructure, primarily cooling and power distribution. High-efficiency facilities target a PUE of 1.1 or lower ([NationalGridDSO 2025](#); [IEA 2026a](#)).

Table 2 Lifecycle Carbon Comparison Across Sectors

Sector	Manufacturing/ Embodied Carbon	Operational/ Use- Phase Carbon	Key Insight
Built Environment	~11% of global CO2 (materials extraction + construction)	~28% global CO2 (heating, cooling, lighting)	Embodied carbon largely unregulated compared to operational
Transport (ICE)	15-20% of lifecycle	80-85% of lifecycle	Use phase dominates
Transport (EV)	30-40% of lifecycle	60-70% of lifecycle	Manufacturing share rises as use phase emissions fall
Data Centres	Significant (construction + hardware manufacture)	~60% of facility electricity to servers; cooling 7-30%	Hardware refresh cycles intensify embodied carbon

Sources: ([Whitehead, Andrews, and Shah 2015](#); [Cox Automotive 2025](#); [IEA 2025](#); [UNEP 2025](#))

Long-Term Site Stewardship – the buildings themselves

In complex industries, a cycle of build, use, and demolition destroys value. A lesson can be drawn from high-hazard infrastructure (termed the “Decommissioning Lesson” or “Sellafield Lesson” in the previous article ([Minney 2026b](#))), where project requirements lead to specialized assets with short operational lives.

Design choices for the buildings themselves: material selection, structural form, orientation; these can substantially reduce both construction-phase and whole-life embodied carbon ([Minney 2026b](#)).

Building Information Modelling (BIM) with a temporal dimension (multiple generations of use) offers the means to plan for this stewardship in practice. By modelling the land’s future use alongside its current configuration, project managers can identify the common-denominator infrastructure – ducts, roads, utility corridors, seismic-hardened foundations, and secure perimeters – that can be designed to outlive individual server refreshes and serve multiple technology generations ([Minney 2025a](#)).

The pace of technological change makes such forward planning essential. Current superconducting quantum computers, for instance, operate at temperatures close to absolute zero

(approximately 15 millikelvin). While research into higher-temperature quantum devices is advancing, these are still likely to require cryogenic infrastructure well below -100°C . Utilities and insulation capable of delivering such temperatures represent long-lived assets that will serve successive generations of quantum hardware – provided they are designed into the site from the outset rather than retrofitted.



Figure 2 Value Destruction Cycle for data centres

This requirement to plan for likely but as-yet-unspecified future change shifts the focus from individual project delivery to what we term ‘4D Stewardship’ — the management of shared assets across multiple technology generations ([Minney 2025a](#)).

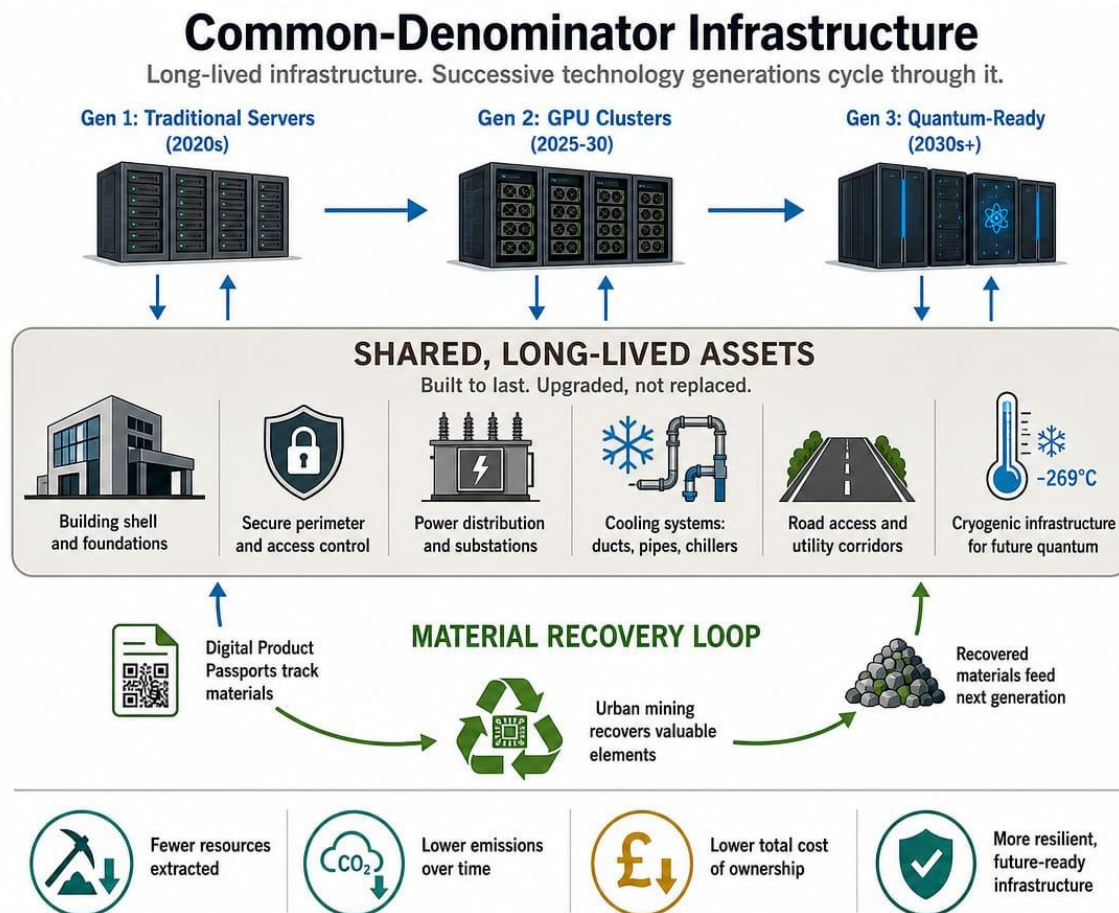


Figure 3 Identifying "Common Denominator" infrastructure

Urban Mining: From Buildings to Micro-Electronics

“Urban Mining” has traditionally focused on reusing steel or concrete from demolished buildings (Koutamanis, van Reijn, and van Bueren 2018; Arup explains 2025; SITRA 2025). However, applying this to IT is significantly more difficult because electronic components are microscopic and integrated into complex circuit boards. Currently, recovery and recycling rates for specific rare earth metals used in electronics remain below 1% of total demand for those metals, even though overall e-waste collection and formal recycling rates have reached 22.3% globally and exceed 40% in Europe (Baldé et al. 2024). The gap between collection rates and actual material recovery highlights the difficulty of extracting microscopic quantities of valuable elements from complex circuit boards.

To bring about a change, manufacturing is beginning to utilize **Digital Product Passports** (required under the EU Ecodesign for Sustainable Products Regulation) to track the provenance of materials (European data 2024). By treating a data centre as a “material bank” rather than a

temporary installation, project managers ensure that today's hardware provides the resources for tomorrow's systems, even if the components are too small for traditional recycling methods.

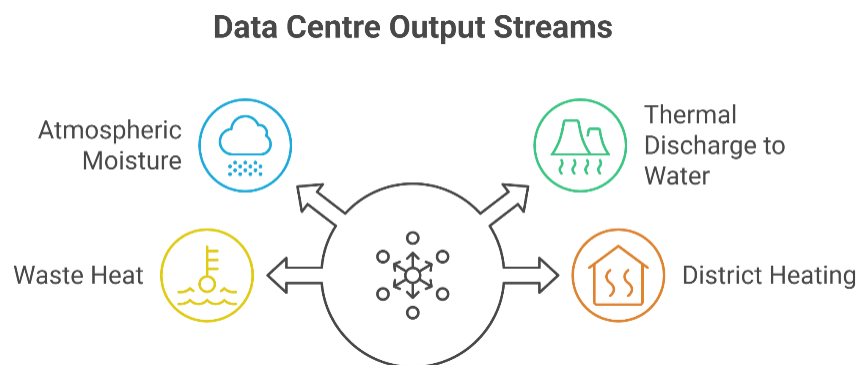


Figure 4 Data centre output streams

The Water-Energy Nexus: Cooling's Hidden Cost

The heat generated by compute operations must be dissipated to maintain system stability, and this typically involves water. A small 1 MW data centre using evaporative cooling can consume approximately 25.5 million litres of water per year ([Mytton 2021](#)). Larger hyperscale facilities can use up to 19 million litres per day – equivalent to the daily water supply of a small town ([EESI 2025](#)). The water footprint divides into three categories of concern.

- **Atmospheric effects.** Evaporative cooling towers release heat and moisture into the surrounding air. Where data centres are concentrated – as in west London or northern Virginia – this can measurably alter local humidity and micro-climate conditions. If moisture-laden exhaust is carried by prevailing wind over residential areas or temperature-sensitive facilities, it constitutes what this series has termed a ‘social pollutant’ affecting community wellbeing ([Minney 2025f](#)).
- **Aquatic effects.** Where facilities draw river or lake water for once-through cooling and return it at elevated temperatures, the ecological consequences are well documented in the power generation literature. Warmer water holds less dissolved oxygen: each 1°C rise reduces oxygen saturation capacity by approximately 2–3% ([Johnson et al. 2024](#)). This can push watercourses below the thresholds required by cold-water species such as trout and salmon, reduce invertebrate diversity, and compromise a river's capacity for natural self-purification ([Langford 1990](#); [Johnson et al. 2024](#)).
- **Siting considerations.** Project managers evaluating sites must consider prevailing wind direction, proximity to residential areas, the thermal sensitivity of receiving watercourses, and cumulative effects where multiple data centres cluster. In water-stressed regions – and even the south-east of England faces periodic water stress – the competition between data centre cooling and potable water supply is becoming a planning constraint ([Siddik, Shehabi, and Marston 2021](#)).

The Efficiency Contrast: Technology as a Catalyst

The environmental cost of IT projects might be justified when technology is used as a catalyst – reducing emissions in higher-carbon sectors, or improving the performance of data centres themselves. Regulations can drive this innovation ([Porter and Kramer 2011](#); [Minney 2025e](#)).

The Jevons Paradox: The Efficiency Trap

The Jevons Paradox is the observation that as a resource becomes more efficient to use, its total consumption often increases rather than decreases. William Stanley Jevons first identified this pattern with coal consumption in Victorian Britain ([Jevons 1865](#)). Modern scholarship has confirmed the paradox’s relevance across energy systems: Sorrell’s comprehensive review found that while direct rebound effects at the microeconomic level are typically below 100%, economy-wide ‘backfire’ – where total consumption rises despite efficiency gains – cannot be ruled out and may be the norm in rapidly growing sectors ([Alcott 2005](#); [Sorrell 2009](#)). For data centres, the pattern is already visible: as compute becomes more efficient and cheaper per operation, the volume of computation demanded grows faster still, so that total energy demand increases even as energy per task falls. The IEA reports that energy efficiency per AI task has been improving by at least an order of magnitude annually, yet total data centre electricity consumption continued to surge in 2025 ([IEA 2026b](#)). Similarly, as programming tools become more productive, the demand for software development has grown rather than shrunk ([Dans 2026](#)). Project managers must use **Total Value Contribution (TVC)** and **Net Present Social Value (NPSV)** to ensure that efficiency gains deliver actual environmental savings rather than simply enabling a higher volume of work ([Minney 2026a](#)).

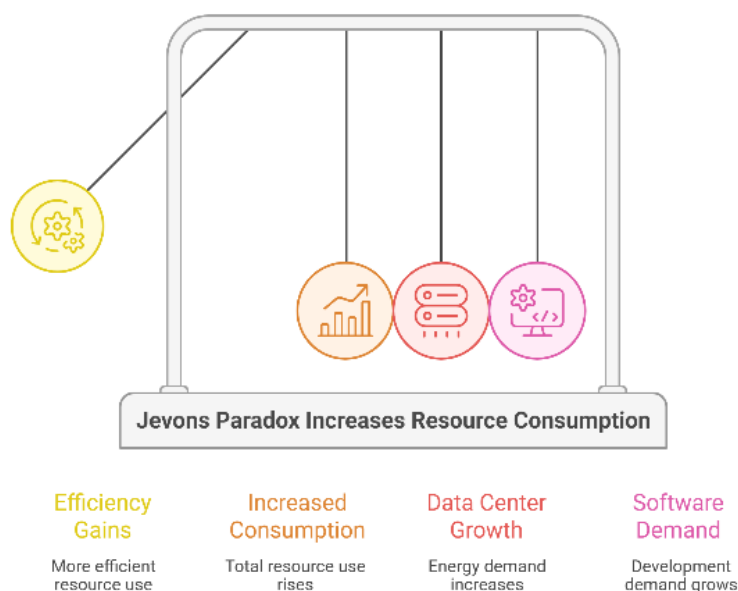


Figure 5 Jevons Paradox

Enabling Effects and Regenerative Design

The traditional approach to design in manufacturing (and possibly in the built environment) has been iterative: build a prototype, test it, learn from failure, redesign, and build again. Each cycle consumes materials and energy. AI-driven simulation changes this fundamentally, allowing project teams to run thousands or tens of thousands of design variants computationally before committing any physical resources. The carbon cost shifts from repeated physical prototyping to a smaller computational investment — and the resulting designs often outperform those achieved through conventional trial-and-error.

- **Materials Discovery for Energy Storage:** Researchers are using deep learning to accelerate the discovery of new materials for batteries and other energy applications. Merchant ([Merchant et al. 2023](#)), working at Google DeepMind, used neural networks to identify 2.2 million stable crystal structures – including 381,000 previously unknown stable materials – across the full range of applications from microchips to batteries to photovoltaics. More targeted work by Hu and his team ([Hu et al. 2025](#)) reviews how AI is now being applied specifically to rechargeable battery development, from electrode material screening to electrolyte optimisation, with machine learning reducing screening times by over 50% compared to traditional density functional theory approaches. These advances are accelerating the transition to grid-scale renewable energy storage ([Pierrel 2025](#)).
- **Building Energy Optimisation:** AI is delivering measurable improvements in building energy performance. A systematic meta-analysis of 126 studies found that reinforcement learning achieves average energy savings of 22.3% in building HVAC systems, with hybrid AI methods reaching 28.1% ([Gunasinghalge, Alazab, and Talukder 2025](#)). These approaches enable designers to meet BREEAM and LEED standards through optimised passive ventilation, thermal management, and real-time adaptive control – reducing dependence on mechanical cooling and the energy it requires.
- **Electric Vehicles:** Technology enables the design of lightweight, high-strength components and more efficient thermal management systems, increasing the “useful life” of the vehicle ([Cox Automotive 2025](#)). For example, AI-driven innovations in battery management, predictive maintenance, and charging networks, battery safety and driver augmentation (collision detect, lane departure, driver alertness) all contribute. Tesla Gigacast reduced component numbers (reducing cost to build and increasing reliability) and reduced total weight ([Munro 2022](#)). Volvo uses AI for crash simulation ([Rodrigues and Júnior 2025](#)).
- **Cooling the Compute:** In a telling example of technology optimising itself, Google’s DeepMind division demonstrated in 2016 that deep neural networks trained on sensor data from within data centres could reduce cooling energy consumption by 40% and overall Power Usage Effectiveness by 15%. By 2018, the system had moved from providing human-implemented recommendations to directly controlling cooling operations under human supervision ([Evans and Gao 2016](#); [Gamble and Gao 2018](#)). This approach has since

been adopted across Google’s global data centre fleet and represents a model for the industry.

Powering the Future: Security, SMRs, and Strategic Siting

The energy demands of high-performance computing are straining grid connections across the UK, Europe, and the United States. The contrast with China is instructive. China has invested in generation capacity such that its reserve margin has consistently remained at 80–100% nationally, meaning it maintains roughly twice the capacity it requires at any given time (Fishman, cited in [Roytburg 2025](#)). Combined with an extensive ultra-high-voltage transmission network (39 UHV projects completed as of 2025, with a further 28 planned by 2030), this allows China to site data centres in remote, energy-rich regions through its ‘East Data West Computing’ strategy without facing the grid connection bottlenecks that constrain deployment in the UK, Europe, and the United States ([Nye 2026](#); [OIES 2026](#)). The catch is carbon: China’s grid still relies heavily on coal, making its data centres among the dirtiest per unit of compute globally.

As project managers, we must look for synergies in infrastructure that solve multiple constraints simultaneously. This involves rethinking both the source of power and the geographical placement of assets to minimize the “resource debt” described earlier.

The Security Synergy: SMRs and Data Perimeters

Small Modular Reactors (SMRs) represent a significant opportunity for zero-carbon, “islanded” baseload power ([WNA 2026](#)). However, the primary barrier to their deployment is the high cost of nuclear-grade security to prevent the theft of fuel or terrorist interference ([IAEA 2022](#); [Minney 2025e](#)).

Coincidentally, Tier 4 data centres – the highest classification under the Uptime Institute’s standard – already maintain some of the most physically hardened perimeters in the commercial world. These facilities are built like fortresses: 24/7 guards, biometric entry, mantraps, and wall-to-wall surveillance come as standard ([Uptime 2019](#)). By co-locating SMRs within these existing secure perimeters, project managers can leverage existing security investments. This colocation not only secures the energy supply for the project but also reduces the transmission losses associated with drawing power from a distant national grid ([Minney 2025e](#)).

Strategic siting for cooling could be as critical as the power source, and the water-energy considerations discussed earlier in this article apply directly here. Moving projects to cold climates (Iceland, Norway, Sweden, even UK) allows for “free cooling” from the atmosphere, and placing the data centre inside a mountain is good for security (especially against kinetic weaponry).

Yet this carries its own environmental footprint, with the atmospheric and aquatic consequences discussed earlier compounding in cold, still environments where temperature inversions trap moisture. It also carries trade-offs of remoteness and latency, and social license ([Gallup 2026](#)).

Data Centres as Grid Assets: The Investment Multiplier

The IEA's April 2026 report 'Key Questions on Energy and AI' identifies a striking secondary effect of concentrated data centre demand that project managers should factor into their business cases. The sheer scale of technology companies' electricity purchasing power — capital expenditure across five major firms exceeded USD 400 billion in 2025, with a further 75% increase projected for 2026 — is now large enough to reshape energy markets rather than merely participate in them ([IEA 2026b](#)).

First, technology companies accounted for approximately 40% of all corporate renewable power purchase agreements (PPAs) signed globally in 2025. These long-term contracts provide the revenue certainty that renewable energy developers need to secure project finance, effectively underwriting wind and solar capacity that might otherwise remain uncommitted. Second, onsite battery storage is becoming a defining feature of next-generation AI data centres. Because AI workloads involve rapid, large swings in power demand that can stretch the technical capabilities of any single generation source, battery systems are being deployed both to smooth the data centre's own load profile and — with the right incentive structures — to provide demand response and frequency regulation services back to the grid. The IEA observes that this positions data centres not merely as grid loads but potentially as grid resources ([IEA 2026b](#)). Third, the pipeline of conditional offtake agreements between data centre operators and SMR nuclear projects has grown from 25 GW at the end of 2024 to 45 GW by early 2026, suggesting that data centre demand may be the commercial catalyst that brings SMR technology to market at scale — an outcome that would benefit the wider energy system far beyond the technology sector.

For the project manager assembling a business case, this means that the energy costs of a data centre project should not be assessed in isolation. The project's demand may trigger investment in renewable generation, grid-scale storage, and nuclear capacity that serves the broader community. This is the kind of cross-portfolio synergy that the Total Value Contribution approach is designed to capture ([Minney 2026a](#)).

The Biological Benchmark: Synapses vs. Silicon

There's been a certain amount of fuss about when we will reach "Singularity" — the hypothetical point where AI exceeds human intelligence. Before we get carried away, it is worth grounding expectations in biology. Current AI has a significant distance to travel before it can match the regenerative capacity of the human brain.

The Scale of Connectivity

A modern supercomputer or high-end GPU cluster operates using binary silicon gates, a system of "on" or "off" switches (1s and 0s). While we can now assemble billions of these gates, they are fundamentally linear compared to the human brain. The brain contains approximately 86 billion neurons and approximately 100 trillion (10^{14}) synapses ([Herculano-Houzel 2009](#)). A silicon gate has only one trick: on or off. A biological synapse, by contrast, responds to multiple

neurotransmitters (such as dopamine, serotonin, and glutamate) each modulating its behaviour differently. In effect, every synapse is a multi-factor logic gate.

The Power Paradox

The efficiency gap remains vast. A high-performance AI cluster can consume hundreds of megawatts of power ([Irwin 2024](#); [O'Donnell and Crownhart 2025](#); [Youell 2025](#)), while a human brain, performing more complex parallel processing, runs on approximately 20 watts. Recent experiments, such as those by Cortical Labs, have shown that a petri dish of approximately 800,000 human brain cells learned to play the game 'Pong' within five minutes of real-time gameplay – demonstrating goal-directed learning with a fraction of the energy that silicon-based AI training typically requires ([Kagan et al. 2022](#)). This has led to a new field of “Biological Computing” that seeks to bridge the gap between silicon-based binary logic and the low-power, high-density processing of organic life.

Synapses vs. Silicon: A Biological Benchmark




Figure 6 Biology vs Silicon

A Thought Experiment

It does make me wonder whether at some point in the future we will harness the power of the human brain in compute environments. I imagine data centres that look like offices, where a ‘worker’ arrives for their shift and puts on a headset containing a human-brain interface; the organisation uses the power of so many human brains for 90 minutes at a time (to allow for breaks) and the ‘worker’ goes home at the end of their shift with no new memories implanted, having done exactly what many employees do today only with more intimate technology.

The ethical questions such a scenario raises — about consent, identity, and the nature of work — would themselves require a governance framework at least as rigorous as the one proposed in this article.

Practical Application of AI: Delivering Reliable Workflows

If we cannot yet match biology's efficiency, we can at least be smarter about how we deploy the silicon we have.

The demand for “AI everything” often ignores the reality that vital government and commercial processes require consistency and reliability. Real-time AI models are designed to be creative (which can lead to “hallucinations”) and can be unpredictable when operating outside their training experience. To deliver robustness of government agent decisions, we already recognise

we must move away from the vagaries of human whim toward automated, verifiable logic; to deliver sustainability and regeneration, we need to do the same with AI.

The Two-Step Revolution

To combine efficiency with reliability, this paper proposes a two-step design process:

1. **Step 1: The AI Architect.** Use a high-power AI model to analyse vast datasets and identify the most efficient paths. The output from this process is a **human-verifiable decision tree** or a multi-factor regression formula. This stage is where “judgement” (the collation of multiple data with some missing and some contradictory) is captured and audited by human experts to ensure fairness and accuracy in the outcome of the decision tree. Racial, gender, and age bias – well-documented challenges in machine learning systems ([Buolamwini and Gebru 2018](#)) – can be systematically tested for at this stage, and the decision tree adjusted and re-validated where bias is identified. This is not a trivial exercise: it requires domain expertise, diverse test data, and ongoing monitoring. But it is substantially more tractable when applied to a static decision tree than when attempting to audit a live, opaque black-box large language model.
2. **Step 2: The Lightweight Workflow.** Once the logic is verified, we use AI to code a static, lightweight script (in C++ or Python) that follows the decision tree. This script runs with negligible power consumption compared to a real-time LLM query. Government employees use this script, still working with “dirty” data (a mass of information about a situation, some requiring confidence weights, some missing, some contradictory) that is inevitable with complex decisions such as rights to benefits, but with consistent and defensible results.

The Two-Step Revolution

Efficiency through AI. Reliability through audit.

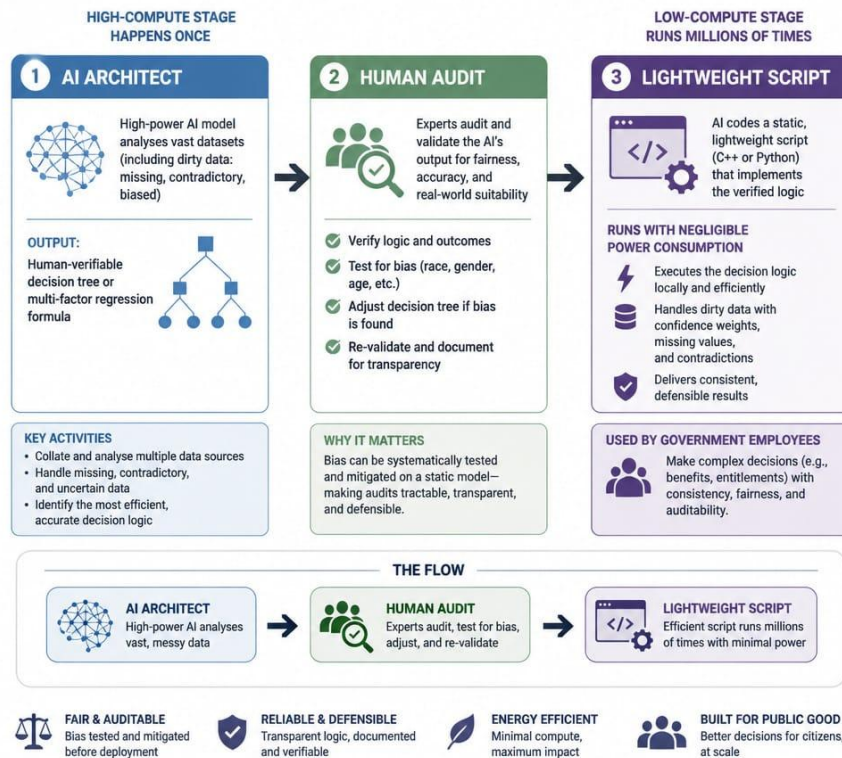


Figure 7 The Two-Step revolution - putting AI to work

To illustrate this with a practical example, here's how it could apply in the immigration system:

Step 1: Architecture

1. Past decisions, and the data collected to inform them, are analysed to create a decision-tree (using AI). The decisions can include the results of appeals and reviews, and should include the millions of cases that no individual human could possibly take into account.
2. This generates a model where information is collated, weighted and balanced, and informs a decision.
3. Immigration is especially challenging because people seeking study visas, work visas, settlement (Indefinite Leave to Remain ILR) and British Citizenship might provide information that contradicts other sources for the same person. Also it is not practical or even possible to obtain all the same information for each person, so from the model's point of view there will always be 'missing data' and the decision tree needs to be able to guide the user in spite of gaps and contradictions.

4. The model provides clarity. The model should be simple enough that it can be checked and reviewed, because reviews highlight that past decisions often reflect racial, gender and age bias, and it is these past decisions that were used to train the AI. A model can have its parameters adjusted until the reviewers are satisfied that it reflects the best current position and policy.

Step 2: Automated Workflow

1. The model is now implemented as a workflow, using AI to write the code.
2. This workflow can then be run on a test sample (in the past, researchers have generated a model using half the data and used the model on the other half to see how well it works. I think I would recommend developing the model on 90% of the data and running the model on 10% because of the numbers of cases – 2 million test cases (from 18-20 million cases reviewed since 1990) should be quite sufficient).
3. Once the workflow has been confirmed, and has a useable user interface, then this is the tool that can be used by civil servants tasked with the decisions.
4. Note that the automated workflow is still reliant on humans to interview applicants and provide a judgement score on the information obtained, and would make a recommendation from the model for a human to make the final decision. Even with this number of data samples, human-in-the-loop is a vital safeguard in sensitive situations such as those seeking access to UK.

Project Management Governance and Strategic Stewardship

The transition from traditional IT delivery to regenerative technology stewardship requires a significant shift in how project managers assemble business cases and exercise control. We have learned that sustainability is an outcome of specific decisions made during the “pre-carbon” planning stages ([Minney 2025b](#)). For the project professional, this means moving beyond the delivery of a technical output to managing the long-term consequences of that output on the wider system.

Assembling the Case for Regenerative Synergy

Project managers must develop the capability to look for synergies that exist outside the immediate project scope. A primary example is the security and energy synergy discussed earlier: while an energy project involving Small Modular Reactors (SMRs) might be stalled by the high costs of specialized security, a data centre project already provides that security as a baseline requirement, and the SMR could be scaled to provide extra power to businesses and residential nearby.

Creative thinking in the initiation phase allows a PM to advocate for colocation, effectively sharing the “resource debt” across two formerly distinct portfolios. Similarly, when considering cooling, a steward of the environment looks at next-stage consequences. Rather than simply

dumping heat into the atmosphere or water – which represents a “social pollutant” – the PM evaluates whether that heat can be captured for local district heating or industrial processes, thereby turning a waste product into a community benefit ([Minney 2025f](#)).

The Quintuple Bottom Line at the Stage Gate

The “Iron Triangle” of time, cost, and quality was never a measure of project success – only a measure of how well the project manager operated within constraints ([Minney 2025d](#)). To ensure IT projects contribute to a sustainable future, governance boards must implement a **Quintuple Bottom Line** approach: assessing People, Planet, Prosperity, Peace, and Partnership ([Minney 2025c](#)).

We propose that any project should not even pass a stage-gate without demonstrating how it addresses at least the first three pillars, to a level of detail and specificity (and identifying ‘strategic optimism’ and suggested corrections) appropriate to that stage gate:

- **People:** How does the project impact user wellbeing and digital inclusion?
- **Planet:** What is the whole-life carbon and water footprint, and other environmental impact, including decommissioning?
- **Prosperity:** Does the project strengthen the local economy or create a “value leakage” through extractive global sourcing? ([Minney 2026a](#)).

Quintuple Bottom Line Stage Gate Process

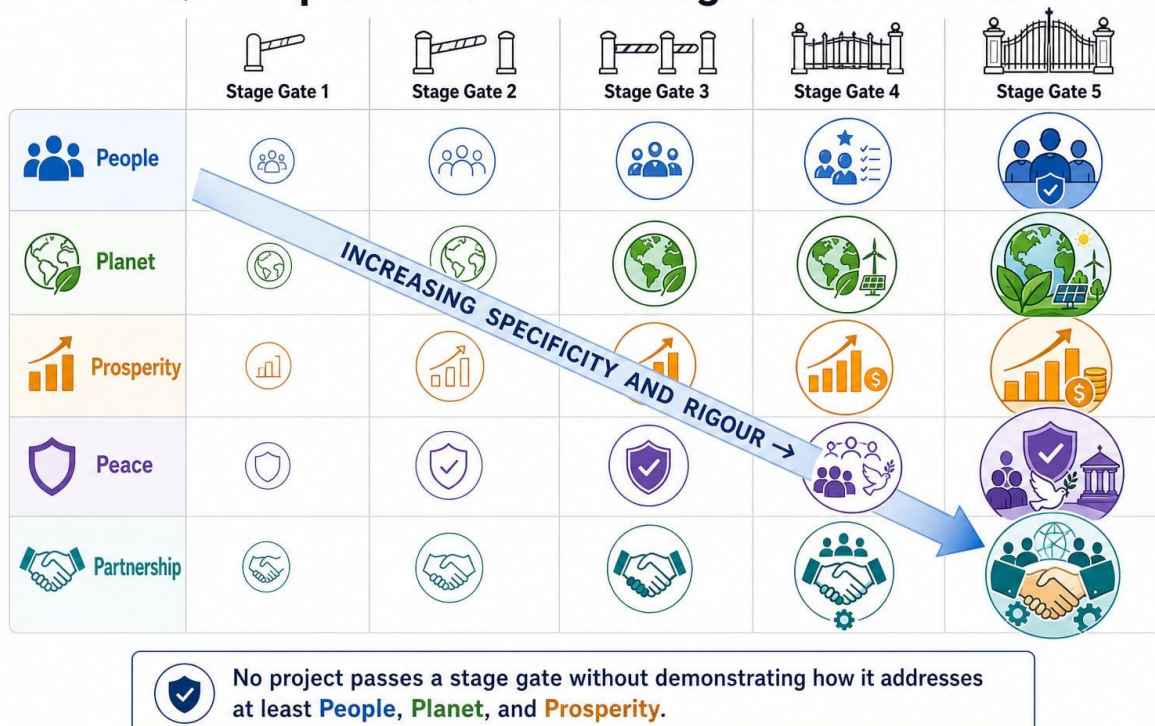


Figure 8 Quintuple Bottom Line

Table 3 Quintuple Bottom Line - specificity at each stage gate (example)

Pillar	Gate 1 (Concept)	Gate 2 (Feasibility)	Gate 3 (Design)	Gate 4 (Build)	Gate 5 (Close)
People	Identify stakeholders	Wellbeing impact assessment	Digital inclusion plan	Accessibility testing	Social outcomes measured
Planet	Carbon screening	LCA estimate; water baseline	Detailed carbon + water budget	Monitoring in place	Actual vs forecast
Prosperity	Local economy scan	Supply chain assessment	Local procurement targets	Spend tracking	Economic benefit realised
Peace	N/A at this stage	Conflict/security screening	Security design complete	Security operational	Incident review
Partnership	Identify partners	Partnership MoUs	Joint governance agreed	Collaborative delivery	Partnership review

Conclusion

The central question of this article — whether the environmental benefits of AI outweigh the costs of the data centres that house it — does not have a simple answer, and we should be suspicious of anyone who claims otherwise.

What is clear is that the metaphor of ‘the Cloud’ has done real damage by obscuring the physical reality of digital infrastructure: from the mines that supply rare earth metals to the power stations feeding hundreds of terawatt-hours into server racks, and the cooling towers evaporating billions of litres of water. These are core project costs, externalities that need to be acknowledged, and they belong in every business case.

Equally clear is that technology remains our most promising tool for addressing the sectors — built environment, transport, energy systems — that account for the vast majority of global emissions. When AI discovers a battery chemistry that transforms grid storage, or cuts a data centre's cooling bill by 40%, the environmental return can justify the initial compute investment. The operative words are ‘can’ and ‘in principle.’ I’ve spent much of this article showing why that caveat matters. Without rigorous governance — Total Value Contribution analysis, Quintuple Bottom Line stage-gate controls, and honest accounting for the Jevons Paradox — efficiency gains will simply feed demand growth, and the net environmental position will worsen.

The biological benchmark reminds us how far silicon computing has to travel: the human brain performs vastly more complex parallel processing on 20 watts than any AI cluster achieves on hundreds of megawatts. This is not an argument against AI, but it is a powerful argument for using high-compute resources sparingly and purposefully — using AI at design-time to create

lightweight, verifiable workflows rather than running energy-intensive models for every routine query.

For project managers, the practical implication is that we must move beyond delivering technical outputs and begin exercising stewardship over the long-term consequences of those outputs. This means designing sites as material banks rather than disposable containers, seeking infrastructure synergies such as SMR colocation, and treating water and carbon as first-order constraints rather than afterthoughts. The buildings, grids, and systems of the coming decades will be defined not by the hardware they contain, but by whether the project management decisions that brought them into being were wise enough to look beyond the current refresh cycle.

AI usage in researching and writing this paper – statement by the author

This article, “The Weight of the Cloud”, was prepared with the assistance of an Artificial Intelligence (AI) large language model (LLM). Under direction and control of the author, the AI LLM was used to facilitate the drafting, research, and refinement process of the article. For example, AI was guided to refine the language to ensure it aligned with British English conventions, maintained a professional yet accessible tone, and avoided common AI-generated phrasing. An AI tool was also used to assist in the generation of illustrations. The author maintained full control at all times and assumes full responsibility for the completed work.

About the Author



Dr Hugo Minney

London, UK



Dr. Hugo Minney is a Fellow of APM (Association for Project Management), a Member of PMI and PMI UK, Lead of APM's Benefits and Value IN (Interest Network) and Sustainability IN, founder of APM's Nuclear Industries IN and AI & Data Analytics IN, committee member of PMI UK's Sustainability Community of Action and Board Member of the Non-Profit PM4theWorld (none of which are paid). Minney is also chair of the British Standards Institute's working group on Benefits Management, which publishes and maintains BS 202002 (Applying benefits management on portfolios, programmes and projects) (also unpaid).

Minney is a business consultant. He analyses the benefits of change, and weighs them up against the need for effective operations to keep the lights on; he has built business cases of all types and is acutely aware of the pressures to make a single project a success at the expense of the organisation's objectives and the need to resist this; as a former executive board director in National Health Service he can take a portfolio overview and prioritise the individual benefits of projects to ensure the success of the whole organisation. Minney is now a project management consultant with a sideline chairing a charity restoring the sense of community for young people.

Minney specialises in putting a number on difficult benefits (such as sustainability and regeneration), motivating team members by reporting what they are achieving together and motivating teams to build the communities and companies we want to be part of – together. He believes in standards and is accredited as a Social Value practitioner and Chartered Project Professional.

Dr. Minney can be contacted at hugo.minney@thesocialreturnco.org

Documents Cited

- Alcott, Blake. 2005. 'Jevons' paradox', *Ecological Economics*, 54:1 9–21.
<https://www.sciencedirect.com/science/article/pii/S0921800905001084>.
- Arup explains. 2025. "What is urban mining?" In.: Arup. <https://www.arup.com/insights/what-is-urban-mining/>.
- Baldé, Cornelis P., Ruediger Kuehr, Tales Yamamoto, Rosie McDonald, Elena D'Angelo, Shahana Althaf, Garam Bel, Otmar Deubzer, Elena Fernandez-Cubillo, Vanessa Forti, Vanessa Gray, Sunil Herat, Shunichi Honda, Giulia Iattoni, Deepali S. Khatriwal, Vittoria Luda di Cortemiglia, Yuliya Lobuntsova, Innocent Nnorom, Noémie Pralat, and Michelle Wagner. 2024. "Global e-waste monitor 2024: Electronic waste rising five times faster than documented e-waste recycling." In, 148. ITU - International Telecommunications Union; UNITAR - United Nations Institute for Training and Research. <https://globalewaste.org/>; <https://unitar.org/about/news-stories/press/global-e-waste-monitor-2024-electronic-waste-rising-five-times-faster-documented-e-waste-recycling>.
- Bloodworth, Andrew. 2014. 'Resources: Track flows to manage technology-metal supply', *Nature*, 505:7481 19–20.
<https://doi.org/10.1038/505019a>.
- Buolamwini, Joy, and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, 77–91. PMLR - Proceedings of Machine Learning Research.
<https://proceedings.mlr.press/v81/buolamwini18a.html>.
- Cordella, Mauro, Felice Alfieri, and Javier Sanfelix. 2021. 'Reducing the carbon footprint of ICT products through material efficiency strategies: A life cycle analysis of smartphones', *Journal of Industrial Ecology*, 25:2 448–64. <https://doi.org/10.1111/jiec.13119>.
- Cox Automotive. 2025. "ICE vs EV: Performance, costs, and lifespan." In.: Cox Automotive.
<https://www.coxautoinc.eu/ev-hub/driver-ev-hub/ice-vs-ev-performance-costs-and-lifespan/>.
- Dans, Enrique. 2026. "AI was supposed to replace developers. So how come there are more than ever?" In.: Medium.
<https://medium.com/enrique-dans/ai-was-supposed-to-replace-developers-so-how-come-there-are-more-than-ever-00c104346085>.
- EESI. 2025. "Data Centers and Water Consumption." In, edited by Miguel Yañez-Barnuevo. EESI - Environmental and Energy Study Institute. <https://www.eesi.org/articles/view/data-centers-and-water-consumption>.
- Ericsson. 2020. "Life cycle environmental impacts of a smartphone." In *Research Paper*.
<https://www.ericsson.com/en/reports-and-papers/research-papers/life-cycle-assessment-of-a-smartphone>.
- European data. 2024. "EU's Digital Product Passport: Advancing transparency and sustainability." In.: European Union. <https://data.europa.eu/en/news-events/news/eus-digital-product-passport-advancing-transparency-and-sustainability>.
- Evans, Richard, and Jim Gao. 2016. "DeepMind AI Reduces Google Data Centre Cooling Bill by 40%." In.: Google Deepmind. <https://deepmind.google/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/>.

- Gallup. 2026. "Americans Oppose AI Data Centers in Their Area." In *Gallup Poll Social Series*, edited by Jeffrey M Jones. Gallup. <https://news.gallup.com/poll/709772/americans-oppose-data-centers-area.aspx>.
- Gamble, Chris, and Jim Gao. 2018. "Safety-first AI for autonomous data centre cooling and industrial control." In.: Google Deepmind. <https://deepmind.google/blog/safety-first-ai-for-autonomous-data-centre-cooling-and-industrial-control/>.
- Gunasinghalge, Lakpriya Udayanga Gunasena Ekanayaka, Ammar Alazab, and Md Alamin Talukder. 2025. 'Artificial intelligence for energy optimization in smart buildings: A systematic review and meta-analysis', *Energy Informatics*, 8:1 135. <https://doi.org/10.1186/s42162-025-00592-8>.
- Herculano-Houzel, S. 2009. 'The human brain in numbers: a linearly scaled-up primate brain', *Frontiers in Human Neuroscience*, 3 31. <https://www.ncbi.nlm.nih.gov/pubmed/19915731>
<https://www.frontiersin.org/journals/human-neuroscience/articles/10.3389/neuro.09.031.2009/pdf>.
- Hu, Qingyun, Junyuan Lu, Jian Hui, Ziyuan Rao, Yang Ren, and Hong Wang. 2025. 'Artificial Intelligence-Driven Development in Rechargeable Battery Materials: Progress, Challenges, and Future Perspectives', *Advanced Functional Materials*, 35:52 e08438. <https://doi.org/10.1002/adfm.202508438>.
- IAEA. 2022. "The Safety and Security of Small Modular Reactors." In.: IAEA - International Atomic Energy Agency. <https://www.iaea.org/newscenter/multimedia/videos/the-safety-and-security-of-small-modular-reactors>.
- IEA. 2025. "Energy and AI." In *World Energy Outlook Special Report*, 304. IEA - International Energy Agency. <https://www.iea.org/reports/energy-and-ai>.
- . 2026a. "Electricity 2026: Analysis and forecast to 2030." In, 225. IEA - International Energy Agency. <https://www.iea.org/reports/electricity-2026>.
- . 2026b. "Key Questions on Energy and AI." In, 138. IEA - International Energy Agency. <https://www.iea.org/reports/key-questions-on-energy-and-ai>.
- Irwin, Kate. 2024. "Energy-Hungry AI Is Eating the Planet. But Is It Also Humanity's Best Hope?" In.: PCMag UK. <https://uk.pcmag.com/ai/151956/power-hungry-ai-is-eating-the-planet-is-it-also-our-best-chance-to-save-it>.
- Jevons, W Stanley. 1865. *The Coal Question* (MacMillan and Co).
- Johnson, Matthew F., Lindsey K. Albertson, Adam C. Algar, Stephen J. Dugdale, Patrick Edwards, Judy England, Christopher Gibbins, So Kazama, Daisuke Komori, Andrew D. C. MacColl, Eric A. Scholl, Robert L. Wilby, Fabio de Oliveira Roque, and Paul J. Wood. 2024. 'Rising water temperature in rivers: Ecological impacts and future resilience', *WIREs Water*, 11:4 e1724. <https://doi.org/10.1002/wat2.1724>.
- Kagan, Brett J., Andy C. Kitchen, Nhi T. Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J. Parker, Anjali Bhat, Ben Rollo, Adeel Razi, and Karl J. Friston. 2022. 'In vitro neurons learn and exhibit sentience when embodied in a simulated game-world', *Neuron*, 110:23 3952–69.e8. <https://doi.org/10.1016/j.neuron.2022.09.001>.
- Koutamanis, Alexander, Boukje van Reijn, and Ellen van Bueren. 2018. 'Urban mining and buildings: A review of possibilities and limitations', *Resources, Conservation and Recycling*, 138 32–39. <https://www.sciencedirect.com/science/article/pii/S0921344918302441>
<https://www.sciencedirect.com/science/article/pii/S0921344918302441?via%3Dihub>.
- Langford, T.E.L. 1990. *Ecological effects of thermal discharges* (Kluwer Academic Publishers; Springer Science & Business Media; Elsevier Applied Science).

Merchant, Amil, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. 2023. 'Scaling deep learning for materials discovery', *Nature*, 624:7990 80–85.

<https://doi.org/10.1038/s41586-023-06735-9>.

Minney, Hugo. 2025a. 'Portfolio Prioritisation Through Benefits Management: Closing the Strategy-Execution Gap, Implementing BS202002: Benefits management on portfolios, programmes and projects, series article', *PM World Journal*, XIV:151:IV (Apr) 21. <https://pmworldlibrary.net/article/portfolio-prioritisation-through-benefits-management-closing-the-strategy-execution-gap/>.

———. 2025b. 'The Project Manager's Mandate: Leading Sustainability and Regeneration in Projects', *PM World Journal*, XIV:154:VII (Jul) 33. <https://pmworldlibrary.net/article/the-project-managers-mandate-leading-sustainability-and-regeneration-in-projects/>.

———. 2025c. 'Integrating sustainability pillars: a holistic approach to project management', *PM World Journal*, XIV:155:VIII (Aug). <https://pmworldlibrary.net/article/integrating-sustainability-pillars-a-holistic-approach-to-project-management/>.

———. 2025d. 'The Elephant in the Room: Questioning Continuous Growth, Donut Economics, and the Service Economy Transition, Leading Sustainability and Regeneration in Projects, series article', *PM World Journal*, XIV:156:IX (Sep). <https://pmworldlibrary.net/article/the-elephant-in-the-room-questioning-continuous-growth-donut-economics-and-the-service-economy-transition/>.

———. 2025e. 'The Impact of Environmental Regulations on Economic Growth, Leading Sustainability and Regeneration in Projects, series article', *PM World Journal*, XIV:157:X (Oct). <https://pmworldlibrary.net/article/the-impact-of-environmental-regulations-on-economic-growth-debunking-myths-and-exploring-opportunities/>.

———. 2025f. 'Building Stronger Communities Through Projects: Social Value as a Strategic Driver of Organisational Success, Leading Sustainability and Regeneration in Projects', *PM World Journal*, XIV:XII (Dec) 27. <https://pmworldlibrary.net/article/building-stronger-communities-through-projects-social-value-as-a-strategic-driver-of-organisational-success/>.

———. 2026a. 'Projects as Catalysts for Local Economies: Leveraging Social Value and Economic Benefits, Leading Sustainability and Regeneration in Projects', *PM World Journal*, XV:I (Jan) 28. <https://pmworldlibrary.net/wp-content/uploads/2026/01/pmwj160-Jan2026-Minney-Prosperity-series-article-7.pdf>.

———. 2026b. 'From Green Building to Timber Innovation, Leading Sustainability and Regeneration in Projects', *PM World Journal*, XV:V (May) 17. <https://pmworldlibrary.net/wp-content/uploads/2026/05/pmwj164-May2026-Minney-From-Green-Building-to-Timber-Innovation-8.pdf>.

Munro, Sandy. 2022. "Tesla Giga Castings: Engineering a New Paradigm in Vehicle Design." In.: Munro Associates. <https://leandesign.com/tesla-giga-castings-evolution/>.

Mytton, David. 2021. 'Data centre water consumption', *npj Clean Water*, 4:1 11. <https://doi.org/10.1038/s41545-021-00101-w>.

NationalGridDSO. 2025. "Data Centre Impact Study." In, 42. National Grid DSO. <https://dso.nationalgrid.co.uk/downloads/15083/data-centre-impact-study2.pdf>.

NESO. 2025. "Future Energy Scenarios: Pathways to Net Zero." In, 170. NESO - National Energy Systems Operator. <https://www.neso.energy/document/364541/download>.

-
- Nye, Josh. 2026. "Data centre power struggle." In. <https://www.rbcwealthmanagement.com/en-asia/insights/data-center-power-struggle>.
- O'Donnell, James, and Casey Crownhart. 2025. "We did the math on AI's energy footprint. Here's the story you haven't heard." In *Climate Change and Energy*. Technology Review. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>.
- OIES. 2026. "The China data centre advantage: hype versus reality." In, edited by Anders Hove, 22. OIES - Oxford Institute for Energy Studies. <https://www.oxfordenergy.org/wpcms/wp-content/uploads/2026/02/Comment-The-China-data-centre-advantage.pdf>.
- Pierrel, Francois. 2025. "AI propels rapid discovery of next-generation battery materials." In *News*. EnerTherm. <https://enertherm-engineering.com/blog/ai-propels-rapid-discovery-of-next-generation-battery-materials>.
- Porter, Michael E., and Mark R. Kramer. 2011. "Creating Shared Value: How to reinvent capitalism—and unleash a wave of innovation and growth." In *Corporate Social Responsibility*. <https://hbr.org/2011/01/the-big-idea-creating-shared-value>.
- Raworth, Kate. 2022. *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist* (Random House Business).
- Recharged. 2025. "EV vs ICE Efficiency: How Much More Efficient Are Electric Cars?" In *EV Education*. Recharged.com. <https://recharged.com/articles/ev-vs-ice-efficiency>.
- Rodrigues, C, and J Júnior. 2025. "Preparing Virtual Vehicle Models for Dynamic Behavior Classification Using Artificial Intelligence Techniques." In *SAE Brasil 2025 Congress*. Sao Paolo, Brazil. <https://doi.org/10.4271/2025-36-0005>.
- Roytburg, Eva. 2025. "AI experts return from China stunned: The U.S. grid is so weak, the race may already be over." In.: Fortune.com. <https://fortune.com/2025/08/14/data-centers-china-grid-us-infrastructure/>.
- Siddik, Md Abu Bakar, Arman Shehabi, and Landon Marston. 2021. 'The environmental footprint of data centers in the United States', *Environmental Research Letters*, 16:6 064017. <https://doi.org/10.1088/1748-9326/abfba1>.
- SITRA. 2025. "The circular economy." In.: SITRA - the Finnish Innovation Fund. <https://www.sitra.fi/en/a-circular-economy/>.
- Sorrell, Steve. 2009. 'Jevons' Paradox revisited: The evidence for backfire from improved energy efficiency', *Energy Policy*, 37:4 1456–69. <https://www.sciencedirect.com/science/article/pii/S0301421508007428>; <https://www.sciencedirect.com/science/article/abs/pii/S0301421508007428?via%3Dihub>.
- UNEP. 2025. "Global Status Report for Buildings and Construction 2024/2025 - not just another brick in the wall." In.: United Nations Environment Programme. <https://www.unep.org/resources/report/global-status-report-buildings-and-construction-20242025>.
- Uptime. 2019. 'Tier Classification System', The Uptime Institute. <https://uptimeinstitute.com/tiers>.
- Whitehead, Beth, Deborah Andrews, and Amip Shah. 2015. 'The life cycle assessment of a UK data centre', *The International Journal of Life Cycle Assessment*, 20:3 332–49. <https://doi.org/10.1007/s11367-014-0838-7>.
- WNA. 2026. "Small Modular Reactors." In.: World Nuclear Association. <https://world-nuclear.org/information-library/nuclear-power-reactors/small-modular-reactors/small-modular-reactors>.

Youell, Joy. 2025. "AI's Energy Crisis Is Coming for Your Electric Bill." In *AEye Spy*. Medium.
<https://medium.com/aeye-spy/ais-energy-crisis-is-coming-for-your-electric-bill-5c467d223eaa>.