ChatGPT

Sustainable AI: Environmental Impact and Governance – Comprehensive Research Outline

Executive Summary White Paper (Draft)

Overview: The sustainability of artificial intelligence has become a dual imperative, encompassing both the environmental footprint of AI systems and the policy frameworks needed to govern them responsibly. This executive summary distills key insights across two pillars – (1) **Environmental Impacts of AI** (energy use, carbon emissions, hardware lifecycle) and (2) **Policy & Governance Frameworks** (regulations, mandates, procurement standards) – with a global perspective (EU, US, and Global South comparisons). It provides a visual-forward synthesis, highlighting critical data points and scenarios (a potential "Green AI" future vs. a "Grey AI" status quo) to inform decision-makers in government and industry.

AI's Growing Environmental Footprint: Current research reveals that AI's carbon and resource footprint is exploding. By 2030, AI workloads could account for ~3–3.6% of global greenhouse gas emissions – on par with the entire aviation industry [1]. Similarly, AI's electricity consumption is projected to reach about 3.5% of global power demand by 2030 [2], doubling the energy used by countries like France. Training a single large AI model can emit **hundreds of thousands of kilograms of CO_2**; one study equated training a big language model to ~300,000 kg CO_2 – roughly **125 round-trip flights** between New York and Beijing [3]. Each use of these models incurs further costs: an average generative AI query consumes **10× more energy** than a standard Google search [4], and producing a 100-word response with GPT-4 can use ~0.5 liters of water for cooling [5]. The **lifecycle impact** extends beyond energy and carbon into **water usage and e-waste**. Data centers supporting AI may guzzle an estimated 4.2–6.6 **billion** cubic meters of water annually by 2027 – exceeding 50% of the UK's total water use [6]. Meanwhile, rapid hardware upgrades contribute to mounting electronic waste (global e-waste projected to reach 74.7 Mt by 2030, nearly double 2014 levels, with only ~17% recycled [7]). These trends paint a stark "**Grey AI**" scenario: unchecked AI growth would intensify climate change and resource strain.

Toward "Green AI" Futures: Encouragingly, a sustainable "Green AI" pathway is feasible. Advances in hardware and efficiency are already yielding gains – for instance, Google's latest AI accelerators achieved a **3× improvement in compute carbon intensity** (CO_2 per unit of compute) compared to previous generations [8]. Clean energy adoption can mitigate operational emissions: over a 6-year lifecycle, roughly **70–90%** of an AI hardware's carbon footprint comes from electricity consumption (vs. manufacturing), so powering AI with renewables can dramatically cut emissions [9]. If paired with smart policies and design choices, AI could even become an enabler of sustainability (e.g. AI-optimized energy grids and efficiency gains in other sectors have the potential to offset 5–10% of global emissions by 2030 [10]). The **Executive Summary White Paper** will feature charts and infographics contrasting these scenarios – illustrating, for example, AI's projected energy/carbon trajectory in a **"Grey AI" baseline (high emissions)** vs. a **"Green AI" scenario** where efficiency measures and renewable energy flatten the curve. It will highlight the **urgent actions** needed to bend the curve: from greener data center designs and algorithmic efficiency to robust governance that aligns AI development with climate goals.

Sustainability Scorecard / Framework (Draft)

Purpose: To provide stakeholders with a clear **scorecard for evaluating AI systems** on sustainability criteria. This draft framework compares AI models, data center setups, or AI-powered products across standardized metrics in **three dimensions: Energy, Ethics, and Scalability.** The goal is a practical tool for policymakers and builders to **rate and compare** the sustainability of AI systems at a glance.

- **Energy & Carbon Metrics:** Every AI system would be assessed on its energy efficiency (e.g. power usage per training/inference) and carbon footprint across the lifecycle. Key quantitative indicators might include **energy per 1000 inferences** (kWh/inf), **training emissions** (CO_2e per model), and **operational carbon intensity** (gCO_2 per query or per task). For example, a scorecard could list Model A vs. Model B: energy per 1000 text prompts (Model A: 50 Wh, Model B: 500 Wh), carbon per training (Model A: 1 ton, Model B: 50 tons), etc. High-level labels (*Green*, *Amber*, *Red*) can indicate performance tiers. The framework draws on initiatives like the *Green AI Index*, which combines real-time energy monitoring with lifecycle analysis [11]. It also factors in **hardware efficiency** (e.g. use of specialized chips or optimized algorithms) and data center PUE (Power Usage Effectiveness). *Example:* A cloud ML service running in a renewable-powered, state-of-the-art data center would score **excellent (Green)** on energy, whereas an equivalent model running on older GPU servers with coal-based power might score **poor (Red)** [12] [13].

- **Ethical & Social Impact Metrics:** Sustainability isn't only environmental – it encompasses ethical AI principles (fairness, transparency, inclusivity) to ensure "sustainable" AI is also socially responsible. The scorecard therefore includes qualitative metrics or checklist items for ethical design: e.g. data privacy safeguards, bias mitigation, and alignment with human-centric design (echoing EU principles that AI should enhance human well-being [14]). While harder to quantify, each AI system could carry an "Ethical AI" compliance score (e.g. meets key governance standards, yes/no) or a risk level per the EU AI Act classification. This links sustainability with long-term **social license** to operate – recognizing that an AI system that is energy-efficient but discriminatory is not truly sustainable. Scoring might reference if a system undergoes **impact assessments** (as required in the EU for high-risk AI, including environmental and social impact analysis [15]).

- **Scalability & Efficiency Metrics:** This dimension examines whether an AI system's design can **scale** to wider deployment without linear increases in resource usage, and whether it leverages resources proportionate to the value provided. Key metrics might include **compute scalability** (e.g. performance-per-watt improvements at larger scales) and **resource overhead** (does the system over-provision hardware or use idling compute?). A system using novel architectures (like sparsity, quantization, or efficient transformers) might score well by achieving the same results with fewer operations. Another factor is **lifecycle longevity**: systems designed to be modular/upgradable (extending hardware life) would get a better sustainability score than those requiring frequent hardware replacement (contributing to e-waste). For instance, using an AI hardware accelerator to its full 6-year lifespan, instead of refreshing every 2 years, improves overall sustainability. A comprehensive scorecard would thus integrate these factors into a composite rating or a dashboard-style presentation.

Draft Layout: The scorecard could be presented as a table or dashboard comparing several AI systems (e.g. *GPT-4*, *Small Open-Source Model*, *Efficient Vision AI*). Columns would list **Energy Use (kWh/query, CO_2e)**, **Hardware & Data Center (efficiency, PUE)**, **Lifecycle Emissions (manufacturing + operation)**, **Ethical Compliance (Yes/No or score)**, **Scalability (e.g. supports load without linear cost increase)**. Each system gets a color-coded rating per column, facilitating quick visual comparison. This framework will allow **institutions** to set sustainability benchmarks (e.g. minimum efficiency standards for AI procurement) and enable **engineers** to identify improvement areas (e.g. if their system is "Red" on carbon, they might prioritize optimization or cloud migration to green data centers).

Case Studies & Ecosystem Map (Draft)

Profiles of Pioneering Efforts: This section will highlight real-world **case studies** of organizations and initiatives leading the way toward sustainable AI, as well as a draft **ecosystem map** visualizing key players and their relationships. By showcasing concrete examples, it grounds the discussion in practice and illustrates the growing network of "sustainable AI" champions across tech, policy, and academia.

- **Tech Industry Leaders:** *Case Study 1:* **Google's 24/7 Carbon-Free AI Infrastructure.** Google has committed to run all its data centers on 24/7 carbon-free energy by 2030 [16]. In practice, this means efforts like matching AI compute with local renewable power on an hourly basis. Google's in-house AI hardware (TPUs) also set benchmarks in efficiency: a life-cycle assessment of Google's TPUs showed that using cleaner power and more efficient chips cuts operational emissions significantly (CCI – compute carbon intensity – improved threefold from TPU v4 to v6) [8]. This case study demonstrates how both **procurement of renewables** and **hardware innovation** can reduce AI's footprint. *Case Study 2:* **Microsoft's Sustainable AI R&D.** Microsoft, pledging to be carbon-negative by 2030, is exploring custom AI chips optimized for efficiency and investing in **liquid cooling** and other data center innovations to curb power usage [5] [17]. They also spearheaded the "AI for Earth" program, illustrating how AI can be applied to environmental solutions – reinforcing the synergy between *using* AI for sustainability and *making* AI itself sustainable. These industry cases serve as blueprints for AI builders: e.g. adopting **efficient model design** (one analysis found some chatbots like Bing's model consumed 5× more energy to train than others [18], pointing to optimization opportunities).

- **Startups and New Ventures:** A growing ecosystem of startups focuses on **"Green AI" solutions**. For example, companies offering AI model **carbon tracking tools** now help teams measure and report emissions of training runs (in line with the emerging practice of ML emissions reporting). Other startups are building **ultra-efficient AI hardware** (e.g. AI accelerators that prioritize performance-per-watt, or neuromorphic and analog chips that drastically cut energy use). *Case Study:* **Efficient Model Compression** – startups and open-source communities have pioneered techniques like knowledge distillation and model pruning to create smaller, energy-saving versions of large models without significant loss in accuracy. A notable effort is the development of **DistilBERT**, a compact version of a language model that retains ~95% of the original's performance while using far less computation, demonstrating "Green AI" in action. These case studies underscore that **innovation for efficiency** is a competitive space, aligning cost savings with sustainability.

- **Policy & Cross-Sector Initiatives:** *Case Study:* **European Union – Climate-Neutral Data Centres & AI Regulation.** The EU's policy ecosystem provides a strong push for sustainable AI. Under the European Green Deal's Digital Strategy, data center operators in Europe have pledged climate neutrality by 2030 [19]. The EU's Energy Efficiency Directive now mandates large data centers (>500 kW) to report detailed energy metrics (PUE, water usage, waste-heat recovery) starting 2024 [20]. Moreover, the forthcoming **EU AI Act** – while primarily risk-focused – includes provisions for transparency of energy use (e.g. **Article 51** requires reporting the energy consumption and carbon footprint for certain AI systems [21]) and is encouraging development of harmonized standards for AI resource efficiency [22]. This combination of *voluntary industry codes* (e.g. EU Code of Conduct for Data Centres [23]) and *regulatory mandates* (AI Act, CSRD reporting [24]) creates a policy case study in steering AI towards sustainability. *Case Study:* **African Union & Global South Innovations.** The **African Union's Continental AI Strategy** explicitly calls out environmental sustainability, urging member countries to integrate **environmental risk assessments** into AI governance [25]. In Kenya, we see a pioneering example of aligning AI with renewables: a partnership enabled a **geothermal-powered data center** (built by a global tech firm in collaboration with local regulators) to leverage Kenya's renewable energy for AI services [26]. Similarly, Malaysia's Corporate Renewable Energy Supply Scheme (CRESS) attracted major cloud providers (Google, Oracle) to invest in data centers powered by local clean energy, channeling ~$9.5B into the economy by 2030 [27]. These cases from the Global South illustrate how **policy incentives and public-private collaboration** can drive sustainable AI infrastructure, turning a potential digital divide into an opportunity.

Ecosystem Map (Visual): We envision a **visual map** that charts the landscape of sustainable AI players and efforts globally. This map would be clustered by categories such as: - **Government & Multilateral Bodies:** e.g. *EU Commission*, *UN/UNESCO (AI for Good initiatives)*, *African Union*, national AI policy task forces (highlighting those with sustainability mandates like Singapore's green data center standards [28] or Brazil's draft AI Act linking AI to sustainable development [29]). - **Industry (Tech Giants & Cloud Providers):** e.g. *Google, Microsoft, Amazon* (with renewable energy investments), *Meta* (research on efficient AI), *NVIDIA* (hardware efficiency roadmaps), etc., noting their sustainability commitments. - **Startups & NGOs:** e.g. *Green Software Foundation* (industry consortium setting "Green AI" standards), *Climate Change AI* (NGO of researchers applying AI for climate and advocating low-carbon AI), small startups in AI energy management. - **Academia & Research:** key labs and projects focusing on Green AI (for instance, MIT's Climate & AI project, the Allen Institute's work on efficiency metrics, or University initiatives tracking AI emissions). - **Cross-sector Collaboratives:** e.g. *Partnership on AI* (which has addressed environmental impacts), or the *Green AI Institute* – a collective driving the concept of a Green AI Index [11].

Lines or arrows can indicate collaborations (for example, a line connecting *Microsoft* and *ADNOC/Masdar* for their joint report on AI in energy [30], or connecting *G42 (UAE)* with *Kenyan regulators* for the geothermal data center [26]). The **ecosystem map** provides a snapshot of the multi-faceted movement toward sustainable AI, showing who is doing what and how they interrelate. It helps policymakers see potential partners and model programs, and helps industry players identify coalitions and guidelines to join.

Open Database / Dashboard (Conceptual Mock-up)

Vision: We propose an open-access **Sustainable AI Dashboard** – a conceptual tool that aggregates data on AI systems' environmental and ethical performance, allowing users to filter and compare by various criteria. This deliverable sketches the layout and features of such a dashboard, which could serve researchers, regulators, and engineers as a **one-stop transparency platform**.

- **Data & Metrics:** The dashboard would compile key metrics for a wide range of AI models and deployments. For each system (e.g. a specific ML model, an AI cloud service, or even an organization's AI program), it would display data such as: **estimated carbon footprint (CO_2e)**, **energy consumption (kWh)** for training and per inference, **water usage** (if available), **hardware type** used (e.g. GPU, TPU, edge device), and **location of data centers** (to infer grid carbon intensity). It would also include qualitative tags like **compliance certifications** (e.g. "Energy Star for servers", "ISO 14001" for data center) and link to any published **impact assessments**. The data sources could be a mix of reported figures (e.g. companies' disclosures under EU CSRD or AI Act Annex requirements [24]) and community-contributed measurements (similar to how ML researchers voluntarily share emissions estimates). A notable inspiration is the idea of a **voluntary AI environmental impact reporting system**: a recent U.S. Senate bill proposes that the EPA and NIST establish a reporting framework for AI's environmental impacts [31]. The dashboard could be the public interface of such a framework, where companies submit data and stakeholders explore it.

- **Filters and Comparison:** Users of the dashboard could filter AI systems by **Industry** (e.g. *Healthcare AI*, *Transportation AI*, *Large Language Models*, *Recommender Systems*), by **Geography** (region of deployment or data center location – acknowledging that, for instance, a data center in Norway with hydroelectric power has a very different footprint than one in a coal-dependent grid), and by **Architecture** (the type of AI or hardware – e.g. *Transformer model vs. CNN*, or *cloud vs. edge*). For example, a policymaker could use filters to find *all language models deployed in the EU over the last year* and compare their energy per query and total emissions. An engineer might filter to see *which computer vision models have the lowest carbon per inference* as a benchmark for choosing an architecture. The dashboard might also allow sorting by metrics, like listing models by descending carbon cost or by energy efficiency rating.

- **Visualization & Interaction:** The interface would present interactive charts – for instance, a bar chart comparing the carbon footprints of several AI systems, or a map highlighting where AI-related energy consumption is highest. A world map view could show concentrations of AI data center energy use and the mix of clean vs. fossil energy in those regions (building on data like the carbon intensity of electricity [32] [33]). Another feature is a **timeline** view: tracking improvements (or regressions) in AI sustainability over time. This could show, say, how the average energy per AI training run is decreasing due to efficiency gains – or conversely how total sector emissions are rising. The goal is to make the data **accessible and actionable**: a researcher could export the data for analysis, or a regulator could identify outliers (e.g. a company whose AI product has unusually high emissions) and potentially incentivize improvements.

- **Concept Mock-up Example:** Imagine logging into the dashboard and seeing a homepage: a summary statistic like "**Global AI Electricity Consumption: X TWh (year-to-date)**" and "**Data Coverage: 200 models, 50 companies reporting**". You could then select a filter like "Architecture: Large Language Models" and see a table of major LLMs (GPT-4, PaLM, LLaMA, etc) with columns for training CO_2, inference CO_2 per 1000 queries, water use per 1000 queries, etc. Clicking on a model would show a detailed profile (charts of emissions by lifecycle stage, any efficiency measures taken, and how it compares to models of similar size). The dashboard could also implement a **"Sustainability Score"** (from the prior scorecard) so each AI system has an easy-to-understand rating. In essence, this open database concept embodies *transparency* and *benchmarking*. It would support external accountability (e.g. an NGO or journalist investigating AI's footprint can draw from a credible database) and internal progress (companies can track their improvements relative to peers). Over time, such a platform could even feed into **sustainability certifications** for AI (for example, a seal of approval for AI systems that meet certain green criteria, akin to LEED certification for buildings).

Critical Insights Memo – Counterintuitive Findings

Briefing Note: In the course of this research, several **counterintuitive or overlooked insights** emerged. This memo highlights a few key findings that defy conventional wisdom or often escape notice, to inform strategic discussions among policymakers and AI developers:

1. **Operational Emissions vs. Hardware Emissions – the Dominant Factor:** It's often assumed that manufacturing AI hardware (chips, servers) is the primary source of AI's carbon footprint. In reality, **operational energy use dominates**. Over a typical AI hardware lifespan, **70–90%** of total emissions come from electricity consumed during model training and inference, far outpacing the embodied emissions of manufacturing [9]. *Insight:* Prioritizing clean energy and efficiency in operation will yield greater carbon reductions than focusing only on greener manufacturing. (This doesn't mean hardware production is trivial, but it means running AI on a coal-powered grid is far worse than the one-time emissions to make the hardware.)

2. **The Hidden Water Footprint:** Discussions on AI's sustainability focus heavily on electricity and carbon, but **water usage** is a critical and often overlooked piece of the puzzle. Data centers require vast water for cooling and also indirectly via electricity generation (for thermoelectric plants). Projections show AI's water consumption reaching **billions of cubic meters** annually in just a few years [6]. To put in perspective, AI's water use by 2027 could exceed half of the U.K.'s total annual water usage [34]. *Insight:* Water scarcity may become a limiting factor for AI growth in certain regions. Strategies like advanced cooling (air cooling, liquid cooling with recirculation), siting data centers near abundant water or using non-potable sources, and improving energy efficiency (thereby drawing less water for power plants) will be increasingly important.

3. **E-Waste and Hardware Lifecycles – A Growing Challenge:** AI's rapid progress drives **short upgrade cycles** for hardware (GPUs, TPUs, etc.), which can lead to a significant electronic waste problem. Globally, e-waste is on track to nearly **double** from 2014 to 2030, yet recycling systems are not keeping up (over 80% of e-waste is not formally recycled [7]). The pursuit of ever-more-powerful AI chips could exacerbate this, as old hardware gets decommissioned. *Insight:* Extending hardware lifetimes and improving recyclability is an often under-prioritized aspect of sustainable AI. Circular economy principles (refurbishment, component reuse, material recycling) need to be integrated into AI hardware procurement and decommissioning. Policymakers might consider incentives or requirements for tech companies to handle e-waste responsibly – for instance, **take-back programs** or minimum recycled content in new devices.

4. **Transparency vs. Market Disincentives – The EU Dilemma:** The EU's aggressive stance on AI governance includes demanding **energy transparency** from AI providers (e.g. requiring disclosure of resource use for foundation models [21]). While this is intended to spur accountability, it has had a paradoxical early effect: some companies have expressed reluctance to deploy certain AI services in Europe, citing strict and unpredictable requirements [35]. In fact, concerns arose that mandatory reporting of energy usage (Annex IX of the AI Act draft) might expose proprietary data or simply deter companies unwilling to share such details. *Insight:* Policymakers must balance the need for transparency with creating a **level playing field** globally. If only one jurisdiction requires detailed carbon reporting, companies might geo-fence advanced AI away from that market, potentially reducing local innovation. International coordination, phasing requirements in gradually, or protecting sensitive information (while still holding companies accountable) could help mitigate this risk. The bigger picture: **global standards** for AI sustainability disclosures would prevent regions with stricter rules from being at a competitive disadvantage.

5. **Not All AI is Equal – Task Variance in Energy Use:** It's easy to speak of AI's footprint in general terms, but a **granular look reveals huge variability** depending on the type of AI task. Counterintuitively, some seemingly simple AI tasks can have outsized impacts. For instance, generating images with AI (e.g. using diffusion models for art) can be *far* more carbon-intensive per output than generating text. A recent analysis found the most carbon-intensive image model emits the equivalent of driving a gasoline car for ~4 miles per just **1,000 images generated**, whereas an efficient text model emits as little as the equivalent of **a few millimeters of driving** for 1,000 sentences [36] [37]. Likewise, an LLM answering a single query might use 10× the energy of a search, but if that LLM replaces a very long process or heavy human effort, the comparison changes. *Insight:* **Efficiency opportunities lie in tailoring AI solutions to the task.** High-volume simple queries might be best served by smaller, specialized models that are much more efficient, whereas large general models should be reserved for tasks that truly require them. This also suggests that **AI architects** should consider *"right-sizing"* models to the problem – a form of sustainable design that avoids using a sledgehammer (giant model) for a nail (simple task).

6. **Global Inequities – Resource Extraction vs. Benefits:** There is a sustainability paradox in the global AI ecosystem: many **resource impacts of AI are externalized to the Global South**, even as AI benefits accrue mostly to wealthier nations. For example, minerals for AI hardware (like cobalt, lithium) are largely mined in African and Latin American countries; data centers are increasingly built where power is cheap (sometimes emerging economies), and e-waste often ships to developing countries for disposal. These countries thus face environmental degradation risks while often benefiting minimally from AI-driven economic growth [38] [39]. *Insight:* Sustainable AI must incorporate **environmental justice**. International policy should support technology transfer and capacity building so that Global South countries can participate in AI value creation, not just bear its costs. Notably, some countries are responding: Brazil's draft AI law makes sustainable development a guiding principle of AI governance [29], and the UAE is advocating for global AI sustainability standards, partly to address cross-border impacts [40]. Recognizing these asymmetries is the first step to addressing them; e.g., by including Global South voices in setting AI standards and by investing in mitigation (like funding e-waste recycling facilities in the regions affected).

These critical insights remind us that a holistic view is required – one that spans energy, water, materials, geopolitics, and beyond. They will inform the detailed findings and recommendations in the full report.

Policy & Design Recommendations

Finally, we present **tailored, actionable recommendations** for different stakeholder groups, bridging high-level policy measures with on-the-ground design practices. Building sustainable AI is a shared challenge requiring coordination between **policymakers, institutional leaders, AI engineers, designers, and startup founders**. Below, we break down guidance for these groups:

For Policymakers & Institutional Stakeholders (e.g. EU Commission, UN, Government Agencies):

- **Integrate Sustainability Criteria into AI Governance:** Update AI strategies and regulations to explicitly include environmental performance. For example, build on the EU AI Act by developing **harmonized standards for AI resource efficiency** (as anticipated in Article 40 of the Act) [22]. Define benchmarks for acceptable energy use per inference or require *eco-design* for high-risk AI systems. Governments should classify **excessive resource usage** as a risk factor in AI assessments – ensuring that an AI system's approval or funding considers its carbon footprint alongside privacy and safety concerns.

- **Strengthen Reporting and Transparency Mandates:** Require organizations to measure and disclose the environmental impact of AI systems. The EU's Corporate Sustainability Reporting Directive (CSRD) already pushes large companies to report sustainability metrics [24] ; this should encompass IT and AI operations (e.g. energy use of major algorithms). Public agencies can lead by example: the U.S. OMB's 2024 guidance on **responsible AI procurement** urges agencies to consider social risks; this can be extended to environmental criteria in procurement [41] . Governments might also implement an **AI Sustainability Label** – similar to energy efficiency labels for appliances – to certify AI products or cloud services that meet low-carbon standards.

- **Green Public Procurement & Incentives:** Leverage the power of the purse. Institutions (from the EU to city governments) should include sustainability requirements when procuring AI systems or cloud services. For instance, adopt **Green Public Procurement criteria** where bids get preference if the AI solution runs on green data centers or uses efficient algorithms [42] . Provide tax breaks, credits, or innovation grants for companies that develop *"green AI" solutions* – whether that's more efficient AI software or infrastructure powered by renewables. At the same time, consider gradually introducing a carbon price or penalty for excessive AI energy use (perhaps initially in government contracts, then more broadly), to internalize environmental costs.

- **Support R&D in Sustainable AI:** Channel funding into research for energy-efficient AI and climate-friendly tech. This includes sponsoring academic programs on *Green AI* (e.g. grants for algorithms that drastically cut computation needed) and supporting open datasets or compute resources for researchers working on efficiency (so they don't need to train giant models from scratch, saving energy). International bodies like the UN or OECD can facilitate knowledge-sharing – for example, creating an **AI Sustainability Knowledge Hub** that tracks best practices and research breakthroughs across countries [43] .

- **Global Cooperation and Standardization:** Work towards international agreements on AI sustainability. Climate change and AI are both global, so standards should be harmonized. Policymakers from the EU, US, and Global South should use forums (G20, UN "AI for Good", etc.) to draft guiding principles on AI energy use and emissions. This might involve coordinating on things like a common **AI environmental impact assessment framework** (so that an AI project in any country undergoes a similar sustainability check) and sharing data via the proposed dashboards. Capacity-building is key: help under-resourced countries in the Global South with technical support to implement these standards, ensuring they are not left behind and that AI's benefits and burdens are more equitably distributed [38] [25] .

- **Plan for Infrastructure and Grid Impacts:** Anticipate the strain AI will put on electricity grids and infrastructure. Regulatory bodies should require new data centers to be **green by design** – e.g. connected to renewable energy or waste-heat recycling systems from the outset. Urban planners and utilities must collaborate with tech companies on grid upgrades (as highlighted by estimates of needing 80 million km of new grid for AI by 2040 [44] [45]). Governments can enforce energy efficiency benchmarks (e.g. in some jurisdictions new data centers must meet a PUE <1.3 or source X% renewable power). Additionally, consider **time-of-use policies**: encouraging AI training jobs to run at times of renewable surplus (for example, via lower electricity rates at noon when solar is abundant).

For AI Industry Builders (Engineers, Designers, Startup Founders):

- **Adopt "Green AI" Design Principles:** In developing AI models and software, make **efficiency a first-class goal** alongside accuracy. This echoes the call by researchers for *Green AI*, which urges measuring and publishing the computational cost of new models [46] . Concretely, teams should track metrics like energy per training and per inference during development. Set internal targets to reduce these with each iteration (much like performance optimization). Techniques to achieve this include model compression, distillation, efficient architecture search, and algorithmic improvements that reduce complexity. For instance, prefer algorithmic optimizations that cut down on redundant computations and explore alternatives to brute-force deep learning (like neuromorphic approaches or hybrid AI that rely on smaller models plus symbolic reasoning).

- **Leverage Sustainable Infrastructure:** Wherever possible, run AI workloads on **green cloud infrastructure**. Major cloud providers offer regions powered by high percentages of renewable energy – choosing those regions (or vendors who are committed to 100% renewables) can drastically lower emissions without any code change. Also utilize cloud features like automatic **workload scheduling** that shifts flexible jobs to times/places with lower carbon intensity (some cloud platforms now integrate with tools that use live grid CO_2 data to schedule jobs). Startups building their own hardware stacks should explore **energy-efficient hardware options**: e.g. using TPUs or ASICs that have better performance-per-watt for your specific task, or the latest generation of GPUs which often improve efficiency. Consider the **embodied carbon** too: if you are a small AI company, do you need to purchase dozens of new servers, or can you use existing shared infrastructure? Using a cloud or colocation that extends hardware life (versus on-premise with rapid turnover) can be greener.

- **Optimize the Full Lifecycle:** Beyond coding, think holistically about the **AI product lifecycle**. This means prolonging the life of models and hardware. Instead of retraining models from scratch, use techniques like transfer learning or fine-tuning on existing models (saves compute). When upgrading hardware, see if older machines can be repurposed for less intensive tasks or sold to others, rather than scrapped. Implement **monitoring in production** to ensure models aren't over-provisioned – for example, scale down model instances when demand is low to save energy. Embrace DevOps/MLOps practices that incorporate efficiency (e.g. automated alerts when a job is using unusually high resources, prompting an engineer to investigate). On the hardware end, work with suppliers that have strong recycling programs, and choose modular equipment that can be partially updated (swapping a component) without replacing the whole server.

- **Sustainable UI/UX Design Choices:** For AI designers, consider **product features** that can influence sustainability. Example: give users the option to choose a "light" mode of your AI service that might be faster and use a smaller model (less compute) when ultra-high accuracy isn't needed. Educate enterprise customers on the carbon footprint of different usage patterns – e.g. "analyses run with wider scope will use X more energy." Sometimes, simply **reducing wasteful usage** is key: if an AI system generates 10 suggestions and users only need 1, can the UI be tuned to generate fewer but more relevant suggestions, cutting compute? These design considerations can subtly steer both the provider and user toward more efficient AI use.

- **Embrace Sustainability as Innovation:** View constraints not as a hindrance but as a driver for innovation. Many startups are now differentiating by saying their AI service is *"95% less carbon than the incumbent."* By building efficiency in from the ground up, you may unlock not only environmental benefits but also **cost advantages** (cloud compute bills are often a huge expense – efficiency saves money) and market appeal (environmentally conscious clients, or those in regulated industries, will prefer sustainable solutions). For instance, a startup that uses clever model architecture to achieve the same result with 10× less data and compute will have a faster, cheaper product – and a green selling point. As part of this, instill a culture of **measuring and publishing results**: include sustainability information in white papers or product docs. Some organizations now publish an "energy star" style metric for each model they release (e.g. carbon per 100 queries) to foster trust and accountability.

- **Collaborate on Standards and Ecosystems:** Industry builders should actively participate in developing the broader ecosystem of sustainable AI. This could mean contributing to open-source tools for measuring emissions (like open libraries that estimate CO_2 from cloud usage), or joining consortia like the Green Software Foundation's Green AI working group. Startups can partner with research labs to test new energy-saving techniques in real deployments. Sharing non-competitive data – for example, anonymized data on your AI energy usage – can help the community benchmark progress (feeding into the dashboard described earlier). Remember that regulators are increasingly interested in this space; engaging with them proactively (responding to consultations, piloting compliance with forthcoming rules) can give your company a voice and forewarning in shaping feasible sustainability requirements.

In Conclusion:

The pursuit of sustainable AI is not a one-time effort but a continuous journey of **innovation, policy evolution, and cross-sector cooperation**. By implementing the above recommendations, **policymakers** can create an enabling environment that aligns AI advancement with global climate goals, while **engineers and entrepreneurs** can drive technical solutions that make AI not only smarter but also cleaner. The balance of "Green AI vs. Grey AI" futures will be determined by actions taken now: whether we succeed in fostering AI that **augments human well-being** *and* respects planetary boundaries, or whether AI's unchecked growth exacerbates sustainability challenges. This comprehensive research and its modular deliverables aim to equip all stakeholders with the knowledge and tools to steer us toward the former – a future where AI is an ally in achieving environmental sustainability, not an adversary.

---

1 7 The Inconvenient Truth About AI's Environmental Footprint: What the 2024 Research Reveals
https://www.linkedin.com/pulse/inconvenient-truth-ais-environmental-footprint-what-2024-luke-hills-sh5ae

2 4 5 17 18 30 44 45 Chart of the week: AI's carbon footprint | The National
https://www.thenationalnews.com/future/2024/11/05/chart-of-the-week-ais-carbon-footprint/

3 6 25 26 27 29 34 38 39 40 The AI-Sustainability Paradox and the Global South: Regulation for Sustainable Innovation - ORF Middle East
https://orfme.org/expert-speak/the-ai-sustainability-paradox-and-the-global-south-regulation-for-sustainable-innovation/

8 9 16 Life-Cycle Emissions of AI Hardware: A Cradle-To-Grave Approach and Generational Trends
https://arxiv.org/html/2502.01671v1

10 36 37 Environmental Impact of Generative AI | Stats & Facts for 2025
https://thesustainableagency.com/blog/environmental-impact-of-generative-ai/

11 Green AI Institute - White Paper on Global Artificial Intelligence Environmental Impact
https://www.greenai.institute/whitepaper/white-paper-on-global-artificial-intelligence-environmental-impact

12 13 15 19 23 24 42 Green AI Institute - EU
https://www.greenai.institute/whitepaper/policy/eu

[14] [21] [22] [35] The EU AI Act: Insights from the Green AI Committee | GSF
https://greensoftware.foundation/articles/the-eu-ai-act-insights-from-the-green-ai-committee/

[20] [28] [32] [33] Greening AI: A Policy Agenda for the Artificial Intelligence and Energy Revolutions
https://institute.global/insights/climate-and-energy/greening-ai-a-policy-agenda-for-the-artificial-intelligence-and-energy-revolutions

[31] Text - S.3732 - 118th Congress (2023-2024): Artificial Intelligence Environmental Impacts Act of 2024 | Congress.gov | Library of Congress
https://www.congress.gov/bill/118th-congress/senate-bill/3732/text

[41] Federal sustainability focus to transform AI procurement
https://federalnewsnetwork.com/commentary/2024/08/federal-sustainability-focus-to-transform-ai-procurement/

[43] Measuring the environmental impacts of artificial intelligence ...
https://www.oecd.org/en/publications/measuring-the-environmental-impacts-of-artificial-intelligence-compute-and-applications_7babf571-en.html

[46] Green AI: The Environmental Impact and Carbon Cost of Innovation
https://www.mindfoundry.ai/blog/green-ai