Wayfinder Learning Lab
Stephen Taylor sjtlyr.net
BSc., PGCE, MA, EdD (Bath, Student)

AI Eco Impact Estimator for Students

Created Jan 2025. Last updated Jan 2026.

Stephen is Director of Innovation in Learning & Teaching at the Western Academy of Beijing and current EdD Research Student at the University of Bath.

Abstract

The AI Environmental Impact Estimator 2026 is an educational tool designed to help students and users understand and visualise the environmental impacts of artificial intelligence usage. This paper documents the development process, methodology, and research sources that informed the creation of this calculator. By providing transparent calculations based on peer-reviewed research and industry reports, the tool enables users to estimate energy consumption, carbon emissions, and water usage associated with various AI tasks including text generation, coding sessions, image generation, video generation, audio generation, data analysis, deep research, and AI search. The tool also contextualises these impacts through equivalencies such as vehicle emissions, tree offset periods, and household energy consumption, making abstract environmental metrics more tangible. This tool is connected to the (If You) USEME-AI model for school adaptation to AI and aligned with the UNESCO AI Competencies for Students (2024) and OECD AI Literacies for Students (2025).

Keywords: *Artificial Intelligence, Environmental Impact, Carbon Emissions, Energy Consumption, Water Usage, AI Sustainability, Educational Technology*

AI Transparency Statement: Both the app development and paper writing involved substantial AI use. Z.AI's GLP-4.7 Full Stack assistant supported coding, Deep Research and Perplexity Pro aided research discovery (with human verification), Perplexity guided app deployment via GitHub and Vercel, and ChatGPT-5.2 helped refine structure and concision in an otherwise human-written paper.

The original task involved 30 text queries, 2 image generations, 17 coding tasks, 22 AI searches, and around 4 hours of personal research, with a further 18 coding runs, 3 deep researches, 3 data analyses, and 10 additional hours when updating using Z.AI Full Stack mode. Using the app's own logic, this equates to about 1.52 kWh of electricity, 0.68 kg CO_2, and 2.88 L of water, roughly the same as 101 smartphone charges, 25 hours of LED lighting, a 2.7 km car journey, or 0.4 months of tree growth.

Links:

- See the app, research, methods and lesson plans here: https://sites.google.com/i-biology.net/ai-footprint-estimator/home
- Direct access to the app: https://sjtylr-ai-eco-footprints-students-2.vercel.app/
- Public repository on GitHub: https://github.com/SJTYLR/sjtylr-ai-eco-footprints-students-2026v5

Cite this document: Taylor, S. 2026. AI Environmental Impact Estimator for Students: Technical Paper. Wayfinder Learning Lab [Online]. Available from: https://sjtylr.net/2026/01/29/ai-impact-estimator-2026/ .

1. Introduction

The rapid adoption of artificial intelligence tools across educational institutions, workplaces, and personal use has created an urgent need for understanding their environmental footprint. As students increasingly rely on AI for coding assistance, research, content creation, and problem-solving, quantifying the cumulative impact of these activities becomes essential for informed decision-making and sustainable technology use.

This paper describes the development of the AI Environmental Impact Estimator 2026, a web-based application created for the Wayfinder Learning Lab. The tool serves two primary purposes:

1. to provide transparent, research-based estimates of AI environmental impacts, and
2. to educate users about the variables that influence these impacts, including energy grid composition and AI model efficiency.

The development of this tool was guided by a critical recognition: while individual AI usage may appear negligible at small scales, the cumulative impact across large user populations becomes significant. Conversely, AI tools could reduce overall environmental footprint by replacing more resource-intensive human activities. The tool aims to present this complexity honestly, encouraging critical thinking about both the costs and benefits of AI adoption.

The build of the tool was completed mostly in collaboration with Z.AI's GLP4-7 “Full Stack” mode as a coding assistant, based on deep research and the author's EdD research. This is a complement to the (If You) USEME-AI model, in which Evaluating Ethical Implications is a key consideration around AI adaptation.

2. Background and Research Context

2.1 The Complexity of AI Environmental Measurement

Measuring the environmental impact of AI tools presents significant challenges due to numerous variables including model architecture, data center efficiency, local energy grids, caching strategies, and usage patterns. As noted in multiple research sources, there is no single "correct" number for energy or water consumption per AI interaction (de Vries, 2023; Freedom, 2024).

Key complexities identified in the research literature include:

- Model Heterogeneity: Different AI models (GPT-4, Claude, Gemini, DeepSeek, local models) have vastly different energy requirements
- Context Length and Complexity: Longer prompts, system prompts, and tool calls exponentially increase energy consumption
- Caching Effects: Cached prompts reduce energy usage for repeated queries

- Data Center Efficiency: Provider-specific infrastructure varies widely in Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE)
- Geographic Variation: The same AI query powered by a clean grid (Norway, Iceland) versus a coal-heavy grid (China, Poland) produces dramatically different emissions

2.2 The Need for Educational Tools

While research on AI environmental impacts has accelerated since 2023, much of this knowledge remains inaccessible to students and general users. Academic papers, industry reports, and technical blogs often present data without context or tools for personal application. This tool bridges that gap by translating research findings into an interactive, educational interface that encourages exploration and understanding.

2.3 Connection to (If You) USEME-AI

The (If You) USEME-AI model provides a structured, learner-centred framework for making intentional decisions about when and how generative AI should be used in educational contexts, with a strong emphasis on *preserving human and learner agency*. It frames AI use as a conditional choice that requires users to clarify purpose, select tools thoughtfully, evaluate impacts, and reflect on outcomes. Within this framework, the AI Footprint Estimator supports the *Evaluate Ethical Implications* component by making the environmental costs of AI use visible, including energy, carbon, and water. By foregrounding uncertainty and scenario-based estimation, the tool encourages informed judgement rather than optimisation or compliance, prompting learners and educators to consider whether AI use is necessary, proportionate, and educationally justified in a given context.

2.4 Connection to UNESCO & OECD AI Competencies for Students

The AI Footprint Estimator aligns with both UNESCO's AI Competency Framework for Students and the OECD AI literacies by supporting informed, ethical, and responsible engagement with AI systems. It directly contributes to UNESCO's Human-Centred Mindset and Ethics of AI domains by foregrounding sustainability, environmental impact, and proportionality as integral to responsible AI use, rather than as peripheral concerns. At the same time, the tool reflects OECD emphases on AI literacy as a combination of understanding, critical evaluation, and agency, particularly in helping students interpret system-level impacts, question assumptions, and make context-sensitive decisions about AI use. By using scenario-based estimation and explicitly acknowledging uncertainty, the estimator positions AI as a sociotechnical system with real-world consequences, supporting the development of critical judgement, systems thinking, and responsible decision-making across both frameworks.

3. Development Methodology

Developing the app was a complex process of collaboration with Z.AI's Full Stack coding capabilities. Where all research and design decisions were made by the author, most of the coding was completed by GLP-4.7. This required multiple, intensive iterations.

3.1 Technology Stack

The AI Environmental Impact Estimator 2026 was built using the following modern web development technologies:

- Framework: Next.js 16 with App Router
- Language: TypeScript 5
- Styling: Tailwind CSS 4
- UI Components: shadcn/ui component library (New York style)
- Data Visualisation: Recharts library for bar charts
- State Management: React hooks (useState)
- This technology stack was chosen for its performance, accessibility features, and developer experience, ensuring the tool is responsive, accessible, and maintainable.

3.2 Development Process

The development process followed an iterative approach between human research and AI assistance:

- Research Consolidation: Comprehensive review of academic papers, industry reports, and technical blogs documenting AI energy consumption, carbon emissions, and water usage
- Task Categorisation: Identification of eight distinct AI task types with different energy profiles
- Factor Determination: Extraction of energy consumption factors from primary research sources
- Equivalency Development: Translation of abstract environmental metrics into tangible comparisons (vehicle miles, tree growth, phone charges)
- UI/UX Design: Creation of an intuitive interface with expandable sections for detailed methodology
- Validation: Cross-referencing calculations against multiple sources to ensure reasonableness
- Iterative Refinement: Incorporation of user feedback and new research (e.g., updating coding task energy to reflect session-based usage)

3.3 Key Design Decisions

Several intentional design decisions shaped the tool's approach:

Task-Based vs. Query-Based Inputs: The tool uses task-based inputs (e.g., "coding sessions" rather than "queries") to better reflect how users actually interact with AI. This decision was informed by Simon P. Couch's (2026) research showing that coding agents involve hundreds of queries per session, making per-query metrics misleading.

Grid Type Selection: Instead of requiring specific geographic location, the tool offers four representative grid types (Clean, Renewable Heavy, Mixed, Coal-Heavy) with example regions. This approach simplifies usage while demonstrating the dramatic impact of energy source on emissions.

AI Model Efficiency Tiers: Recognising that model efficiency varies by orders of magnitude, the tool includes three efficiency multipliers. This helps users understand how model choice affects environmental footprint without requiring detailed technical knowledge.

Visualisation vs Data Presentation: The app is designed for educational use. It balances some visualisation (such as the equivalencies) with qualitative data. The intention is that the data could be used and discussed with students in their own analysis and visualisation.

Expandable Learning Elements: The app is designed to expand for deeper learning, including sections on the complexities of energy measurement, calculation methodology, references and potential actions students can take to reduce their AI impacts.

Embeddable App: The app is intentionally embeddable, so it can be hosted in school learning materials and used in lessons without needing to go to an external site.

4. Calculation Methodology

Researching and refining impact calculations was by far the most time-consuming and (human) labour-intensive part of the project. As a generally opaque field, the emergence of more up-to-date research and analysis in late 2024 and 2025 was useful in iterating on the original, simpler model from Jan 2025. Despite this, there are bound to be significant errors in the calculations, and future iterations will build on further research.

4.1 Energy Consumption Calculations

For each AI task type, energy consumption is calculated using the formula:

- Base Energy (kWh) = Number of Tasks × Energy per Task (kWh/task)
- Final Energy (kWh) = Base Energy × Efficiency Multiplier
- Energy per task values are sourced from primary research and documented below in Section 5.

4.2 Carbon Emissions

Carbon emissions are calculated by converting energy consumption based on the selected grid type:

- CO_2 (kg) = Final Energy (kWh) × Grid Emissions (gCO_2/kWh) / 1000
- Grid emission factors represent typical values for each grid type:
- Clean Grid: 15 gCO_2/kWh (examples: Norway, Iceland, France)
- Renewable Heavy: 150 gCO_2/kWh (examples: Denmark, EU average)
- Mixed Grid: 450 gCO_2/kWh (examples: Global average, US)

- Coal-Heavy: 650 gCO_2/kWh (examples: Poland, China)
- These values are synthesised from EIA (2024), Ember Research (2024), Carbon Brief (2023), and Our World in Data sources.

4.3 Water Usage

Water consumption follows different methodologies based on task type:

- Text Generation (Direct Measurement): Water (L) = queries × 0.00026 mL × Efficiency Multiplier
- Other Tasks (WUE-Based):
- Water (L) = Final Energy (kWh) × 1.9 L/kWh

The 0.00026 mL per query figure for text generation comes from Hannah Ritchie's (August 2025) analysis citing Google's Gemini data. The 1.9 L/kWh Water Usage Effectiveness (WUE) factor represents an industry average synthesised from EESI (2024/2025) and University of California, Riverside (2023) research.

4.4 Equivalencies

To make abstract environmental metrics tangible, the tool calculates equivalencies:

- Vehicle Distance: km driven = CO_2 (kg) / 0.25 kg/km (miles driven = km driven × 0.6214)
- Tree Offset: Tree months = CO_2 (kg) / 1.75 kg/month (Based on 21 kg CO_2 sequestered per tree per year ÷ 12 months, sourced from ForTomorrow (2023))
- Phone Charges: Charges = Energy (kWh) / 0.015 kWh/charge
- Shower Equivalent: Showers = Water (L) / 65 L
- Lightbulb Hours: Hours = Energy (kWh) / 0.06 kWh/hour

EPA vehicle emission factors (2024) and household energy data inform these calculations.

5. Data Sources and Research Foundations

5.1 Energy Consumption per AI Task

The following research sources informed energy consumption factors:

Text Generation: 0.00027 kWh/query

- Hannah Ritchie, Our World in Data (August 2025) - 0.24-0.3 Wh per median query
- Epoch AI (2025) - Infrastructure energy analysis for GPT-4o

Coding Sessions: 0.041 kWh/session

- Simon P. Couch (January 2026) - Median code session cost of 41 Wh
- Critical insight: Coding sessions involve hundreds of longer-than-median queries through system prompts, tool descriptions, and repeated tool calls

Image Generation: 0.0014 kWh/image

- Sasha Luccioni & Jernite (Hugging Face, 2024) - 1.6 g CO_2 per image
- MIT Technology Review (2023) - Energy comparable to phone charging

Deep Research: 0.0054 kWh/query
- Estimated at 20× text generation for multi-step reasoning and context expansion

AI Search: 0.0029 kWh/query
- Kanoppi (2025) - ChatGPT uses 0.0029 kWh vs Google Search's 0.0003 kWh (10× difference)
- Also cited as 68g CO_2 per query vs 0.2g for traditional search

Video Generation: 12 kWh/minute
- Hannah Ritchie (August 2025) - ~1 kWh per 5-second Sora video
- Scaling: 1 kWh / (5/60) minutes = 12 kWh/minute

Audio Generation: 0.06 kWh/minute
- ArXiv (2025) - ~60 Wh per minute for text-to-audio models

Data Analysis: 0.0005 kWh/analysis
- Estimated between text generation and deep research for analytical tasks

5.2 Water Usage Research

Water usage is particularly challenging to evaluate because it involves complex and often opaque factors, including direct water withdrawal (such as groundwater extraction), the extent to which water is recycled within data centres, and significant variation across providers and locations. For illustrative purposes, this analysis draws on several indicative estimates: research from the University of California, Riverside (2023) highlights that AI programmes can consume large volumes of increasingly scarce water, primarily through data-centre cooling and processing; the Energy and Environment Study Institute (EESI, 2024/2025) notes substantial variability in water-use efficiency across providers, with an industry benchmark of approximately 1.8 litres per kWh; and more recently, Hannah Ritchie (August 2025) reports a rare direct estimate for Gemini, suggesting around 2 millilitres of water per text prompt.

5.3 Carbon Emissions and Grid Research

Carbon intensity estimates draw on multiple complementary sources to reflect variation in grid composition and life-cycle emissions. The app aims to recognise the complexity of student location vs the hosting location of the AI models used, giving very rough estimates.
- The U.S. Energy Information Administration (EIA, 2024) provides baseline carbon dioxide emissions per kilowatt-hour of U.S. electricity generation, disaggregated by energy source, offering a reference point for grid-average emission factors.
- At a global scale, Ember Research (2024) documents China's leadership in renewable energy deployment and broader trends in clean-energy adoption, while Carbon Brief (2023) analyses China's emissions trajectory and clean-energy transition, helping to contextualise rapidly evolving grid mixes.

- Our World in Data (2024) further situates these patterns through global renewable-energy growth data and interactive visualisations related to AI energy consumption.
- For greater regional specificity, Zhang et al. (2024) provide life-cycle carbon-emission factors for electricity generation in China, capturing upstream and downstream impacts beyond operational emissions.
- Finally, EPA (2024) equivalency research is used to translate abstract carbon figures into more tangible comparisons, supporting clearer communication and interpretation of AI-related emissions.

Greenhouse Gas Equivalencies Calculator
To support meaningful equivalencies and everyday comparisons, this analysis draws on several widely cited reference points. Average passenger vehicle emissions are estimated at approximately 0.25 kg CO_2 per kilometre (ForTomorrow, 2023), providing a familiar benchmark for translating abstract carbon figures into travel distance. Carbon sequestration is contextualised using estimates that a single mature tree absorbs roughly 21 kg of CO_2 per year (Apple, 2024), drawn from its Environmental Progress Report. Household-scale impacts are further grounded through Apple's analysis of typical residential energy-consumption patterns, while device-level energy use is informed by Jackery (2024), which provides power-consumption data for common digital devices such as laptops and smartphones, enabling clearer connections between AI use, everyday technology, and personal energy footprints.

5.5 Broader AI Environmental Research
Research on AI's environmental impact combines system-level trends, task-level measurements, and life-cycle analysis. De Vries (2023) provides a seminal overview of AI's rapidly growing energy footprint, while Freedom (2024) critically examines where environmental impacts may be overstated or misrepresented. Task-specific energy use is analysed by Luccioni and Jernite (2024), with NPR (2024) grounding these findings in real-world emissions growth reported by major cloud providers. At a modelling and policy level, the OECD's AI Emissions Scenario Generator (2024) and PlanBeEco (2024) support scenario-based analysis of AI's climate impact. Finally, Tomlinson et al. (2024) and Yu et al. (2024) extend the discussion through comparative and life-cycle perspectives, highlighting both cases where AI may reduce emissions and overlooked sources that complicate simple conclusions.

5.6 AI Model Research
Recent work also highlights rapid efficiency gains and emerging transparency in AI systems. The DeepSeek-R1 paper (2024) demonstrates how advances in reasoning optimisation can significantly improve model efficiency, illustrating how newer architectures may reduce energy per task. Industry-level reporting is beginning to improve, with Dev Sustainability (2025) analysing the environmental impact of Google Gemini as an example of growing, though still uneven, transparency in AI sustainability disclosures. Independent analysis from Epoch AI (2025) further clarifies energy consumption patterns across models such as

ChatGPT, distinguishing between the high one-time costs of training and the ongoing, cumulative impacts of inference. These developments are complemented by practitioner tools and frameworks, including those documented by Audacia (2023), which support organisations in assessing and managing the environmental impacts of AI deployment.

Methodological framework for AI impact assessment

6. How to Use the Tool

6.1 Step-by-Step Guide

1. Enter AI Usage Data: Input quantities for relevant AI tasks (text generation queries, coding sessions, images generated, etc.)
2. Users should estimate based on typical usage patterns
3. Consider both frequency and duration of AI use
4. Select Geographic Context: Choose the energy grid type most representative of user's location
5. Default: "Mixed Grid" (global average)
6. Consider local energy composition for more accurate estimates
7. Specify AI Model Efficiency: Choose the category that best matches AI tools being used
8. "Locally Hosted" for models running on personal hardware
9. "Less Efficient" for older or standard cloud models (GPT-4, Claude 3)
10. "More Efficient" for newer, optimised models (GPT-4o, Claude 3.5 Sonnet)
11. Calculate Impact: Click "Calculate Impact" to generate results
12. Review Results: Examine total energy, carbon emissions, and water usage
13. Note equivalencies for context (vehicle miles, tree offset, phone charges)
14. Expand "Task Breakdown" to see contribution by task type
15. Explore Projections: View weekly, monthly, and semester projections
16. Helps users understand cumulative impacts over time
17. Adjust Variables: Experiment with different grid types and AI model efficiencies
18. Compare impacts across scenarios
19. Understand how choices affect environmental footprint

6.2 Educational Use Cases

The simplest use-case of this tool is to visualise environmental impact estimates of student AI use, but there are further potential applications.

Introduce students to AI environmental impacts:

- Demonstrate the relationship between energy source and emissions
- Encourage critical thinking about technology sustainability

Personal Awareness: Individuals can use the tool to:
- Understand their personal AI footprint
- Make informed decisions about AI usage patterns
- Explore the impact of model choice and hosting location

Research and Projects: Students can use the tool as:
- A starting point for deeper investigation
- A reference for AI environmental methodology
- An example of translating research into interactive tools

A set of lesson plans, based on Project Zero at HGSE's Making Thinking Visible and Our World in Data are presented on the app site here:
https://sites.google.com/i-biology.net/ai-footprint-estimator/lessons

6.3 Interpretation Guidelines

Accuracy Context: Users should understand that:
- Calculations are estimates based on available research
- Actual impacts vary by specific model, provider, and usage patterns
- Results are intended for education, not precise accounting

Relative Comparisons: The tool is most useful for:
- Comparing different AI tasks and their relative impacts
- Understanding how grid type affects emissions
- Seeing the cumulative effect of sustained usage

Limitations: Users should recognise:
- Tool does not account for training energy costs
- Network transmission and device energy are excluded
- Water usage estimates have high variability

7. Limitations and Considerations

7.1 Uncertainty in Measurements

The environmental impact of AI involves significant uncertainty due to:
- Model Proprietary Information: Cloud AI providers (OpenAI, Anthropic, Google) do not publish detailed energy data
- Variability in Data Centers: Efficiency varies between and within providers
- Usage Pattern Complexity: User behavior (prompt length, frequency, complexity) dramatically affects energy
- Dynamic Grid Composition: Energy grids change over time as renewable adoption increases

Users should treat all values as order-of-magnitude estimates rather than precise measurements.

7.2 Scope Exclusions

The tool intentionally excludes:

- Training Energy: The massive energy required to train AI models (thousands of MWh per model)
- Infrastructure Energy: Data center construction, maintenance, and decommissioning
- Network Energy: Data transmission between users and data centers
- Client Device Energy: Energy consumed by user devices (laptops, phones, etc.)
- Upstream Emissions: Energy and emissions from hardware manufacturing

These exclusions mean total lifecycle impacts are higher than shown, but training and infrastructure are amortised across millions of users, making per-use estimates reasonable for user-focused analysis.

7.3 Geographic Specificity

While the tool provides four representative grid types, actual grid emissions vary continuously:

- Seasonal Variation: Renewable output varies by season (e.g., solar, hydro)
- Regional Differences: Countries have multiple grid regions with different compositions
- Time-of-Day: Grid composition changes with demand patterns
- Renewable Targets: Many countries are actively transitioning to cleaner grids

Users with specific geographic interests should consult local grid operators or environmental agencies for precise emission factors.

7.4 The Counter-Intuitive Nature of AI Efficiency

An important consideration, highlighted by Tomlinson et al. (2024) and other researchers, is that AI use can sometimes reduce overall environmental footprint:

- Efficiency Gains: AI can replace more resource-intensive human activities
- Time Savings: AI assistance reduces computing time (e.g., faster debugging, research)
- Error Reduction: Fewer mistakes mean less rework and resource waste

The creation of this tool itself exemplifies this paradox: AI coding and research assistance saved many hours of development time, likely reducing total human and computer footprint compared to traditional development. The application of AI in coding assistance allowed the author to create an app that was far beyond the level of his coding ability. Collaboration with AI resulted in significant personal learning throughout the process.

Users should consider both the direct costs of AI usage and potential efficiency gains it enables. In any use of this app, students must be encouraged to question the findings, verify their data and ask more questions.

8. Future Development and Improvements

8.1 Potential Enhancements

Several areas for future development have been identified and might be considered, though the app is currently at a level of complexity suitable for use in schools:

- Geographic Granularity: Addition of specific countries/regions with real-time grid data
- Model Selection: Direct selection of specific AI models with known energy profiles
- Water Usage Refinement: Task-specific water measurements beyond text generation
- Historical Tracking: User accounts to track usage trends over time
- Comparative Analysis: Side-by-side comparison of different scenarios
- Export Functionality: Ability to save and share calculations
- Alternative Visualisations: Ability to create different forms of comparative data visualisation.
- Multilingual Support: Ability to toggle between working languages.

8.2 Research Gaps

Several research gaps limit the accuracy of current estimates:

- Standardised Measurement: Lack of industry standards for reporting AI energy/water usage
- Real-World Usage Data: Limited studies on actual user behavior patterns
- Local Model Energy: Sparse data on energy consumption of locally-hosted AI models
- Lifecycle Analysis: Few comprehensive lifecycle assessments including training, operation, and decommissioning

Future research addressing these gaps would significantly improve estimation accuracy.

9. Conclusion

The AI Environmental Impact Estimator 2026 represents an effort to translate complex, fragmented research into an accessible educational tool. By synthesising findings from academic papers, industry reports, and technical analyses, the tool provides students and users with a foundation for understanding the environmental implications of AI adoption.

Key insights from the development process include:

- Complexity of Measurement: AI environmental impacts are highly variable and context-dependent, requiring transparent methodology and clear communication of uncertainty
- Importance of Context: Energy grid composition and AI model efficiency dramatically influence emissions, often more than task choice
- Educational Value: Interactive tools that contextualise abstract metrics (e.g., equivalencies) enhance understanding and engagement

- Balanced Perspective: While AI usage has environmental costs, it can also reduce overall footprint through efficiency gains and human time savings

The tool should be viewed as a starting point for investigation rather than a definitive answer. Users are encouraged to consult primary sources, question assumptions, and develop their own understanding through critical thinking.

As AI continues to evolve rapidly, ongoing research and tool refinement will be essential to keep estimates current and relevant. The AI Environmental Impact Estimator aims to contribute to broader efforts to make AI sustainability transparent, accessible, and actionable.

10. References

General

Taylor, S. (2024). *(If You) USEME-AI: A learner-centred framework for intentional, ethical, and agency-preserving use of generative artificial intelligence in education.* Unpublished framework, Western Academy of Beijing.

UNESCO (2023). *AI competency framework for students.* Paris: UNESCO.

OECD (2024). *AI literacies for students: Developing critical understanding, skills, and agency for responsible AI use.* Paris: OECD.

AI Environmental Impact and Energy Consumption

Audacia (2023). Sustainable AI: 3 tools to measure environmental impact of ML solutions. https://audacia.co.uk/technical-blog/tools-for-measuring-the-environmental-impact-of-ml

de Vries, A. (2023). The growing energy footprint of artificial intelligence. Joule. https://www.cell.com/joule/fulltext/S2542-4351(23)00365-3

Freedom, P. (2024). The Misrepresented Environmental Impact of AI: Separating Fact from Fiction. Medium. https://medium.com/the-simulacrum/the-misrepresented-environmental-impact-of-ai-separating-fact-from-fiction-7eb60893e4e9

Heikkilä, M. (2023). Making an image with generative AI uses as much energy as charging your phone. MIT Technology Review. https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/

Heikkilä, M. (2023). AI's carbon footprint is bigger than you think. MIT Technology Review. https://www.technologyreview.com/2023/12/05/1084417/ais-carbon-footprint-is-bigger-than-you-think/

Kanoppi (2025). Search Engines vs AI Energy Consumption Compared. https://kanoppi.co/search-engines-vs-ai-energy-consumption-compared/

Luccioni, S., & Jernite, Y. (2024). Power hungry processing: Watts driving the cost of AI deployment? Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency. https://arxiv.org/pdf/2311.16863

NPR (2024). AI brings soaring emissions for Google and Microsoft, a major contributor to climate change. https://www.npr.org/2024/07/12/g-s1-9545/ai-brings-soaring-emissions-for-google-and-microsoft-a-major-contributor-to-climate-change

OECD (2024). AI Emissions Scenario Generator. https://oecd.ai/en/catalogue/tools/ai-emissions-scenario-generator

PlanBeEco (2024). AI's carbon footprint – how does the popularity of artificial intelligence affect the climate? https://planbe.eco/en/blog/ais-carbon-footprint-how-does-the-popularity-of-artificial-intelligence-affect-the-climate/

Simon P. Couch (2026). Electricity use of AI coding agents. Blog. https://www.simonpcouch.com/blog/2026-01-20-cc-impact/

Tomlinson, J., et al. (2024). The carbon emissions of writing and illustrating are lower for AI than for humans. Nature Scientific Reports. https://www.nature.com/articles/s41598-024-54271-x

Yu, X., et al. (2024). Revisit environmental impact of artificial intelligence: overlooked carbon emission source? Environmental Science and Pollution Research. https://link.springer.com/article/10.1007/s11783-024-1918-y

AI Model Research

Couch, S. P. (2026). Electricity use of AI coding agents. https://www.simonpcouch.com/blog/2026-01-20-cc-impact/

DeepSeek (2024). DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. https://github.com/deepseek-ai/DeepSeek-R1/blob/main/DeepSeek_R1.pdf

Dev Sustainability (2025). The environmental impact of Google Gemini. https://www.devsustainability.com/p/the-environmental-impact-of-google-gemini

Epoch AI (2025). How much energy does ChatGPT use? https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use

Ritchie, H. (2025). AI footprint - August 2025. Our World in Data Substack. https://hannahritchie.substack.com/p/ai-footprint-august-2025

Our World in Data (2024). Artificial intelligence - Interactive charts and data. https://ourworldindata.org/artificial-intelligence

Energy Grid and Carbon Emissions

Carbon Brief (2023). Analysis: China's emissions set to fall in 2024 after record growth in clean energy. https://www.carbonbrief.org/analysis-chinas-emissions-set-to-fall-in-2024-after-record-growth-in-clean-energy/

EIA (2024). How much carbon dioxide is produced per kilowatthour of U.S. electricity generation? https://www.eia.gov/tools/faqs/faq.php?id=74&t=11

Ember Research (2024). China is a leader in solar, wind and electric transportation, responsible for half of global growth in 2023. https://ember-energy.org/countries-and-regions/china/

Our World in Data (2024). Renewable energy sources are growing quickly and will play a vital role in tackling climate change. https://ourworldindata.org/renewable-energy

Zhang, Y., et al. (2024). Study on life-cycle carbon emission factors of electricity in China. International Journal of Low-Carbon Technologies. https://academic.oup.com/ijlct/article/doi/10.1093/ijlct/ctae181/7762370

Water Usage

ArXiv (2025). Text-to-audio energy study. https://arxiv.org/html/2505.07615v2

EESI (2024/2025). Data centers & water consumption. https://www.eesi.se/2024/data-centers-water-consumption

University of California, Riverside (2023). AI programs consume large volumes of scarce water. https://news.ucr.edu/articles/2023/04/28/ai-programs-consume-large-volumes-scarce-water

Equivalencies and Lifestyle Comparisons

Apple (2024). Environmental progress report. https://www.apple.com/environment/

EPA (2024). Greenhouse gas equivalencies calculator.
https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

EPA (2024). Greenhouse gas emissions from a typical passenger vehicle.
https://www.epa.gov/greenhouse-gas-emissions-typical-passenger-vehicle

ForTomorrow (2023). How much CO2 does a tree absorb per year?
https://www.fortomorrow.eu/en/blog/co2-tree

Jackery (2024). How many watts does a laptop use: MacBook, Dell, Asus and more.
https://www.jackery.com/blogs/knowledge/how-many-watts-a-laptop-uses

AI vs. Human Comparative Research
PMC (2025). AI vs human programming: A correctness-controlled comparison of environmental impact. https://pmc.ncbi.nlm.nih.gov/articles/PMC12603251

Tomlinson, J., et al. (2024). The carbon emissions of writing and illustrating are lower for AI than for humans. Nature Scientific Reports.
https://www.nature.com/articles/s41598-024-54271-x

Technical Frameworks and Tools

Audacia (2023). Tools for measuring the environmental impact of ML solutions.
https://audacia.co.uk/technical-blog/tools-for-measuring-the-environmental-impact-of-ml

Acknowledgments

The AI Environmental Impact Estimator 2026 was developed for the Wayfinder Learning Lab and WAB students by Stephen Taylor. The tool was created with assistance from AI coding and research tools, demonstrating the counter-intuitive reality that AI can reduce overall development footprint through efficiency gains while introducing new energy demands.

Special thanks to the researchers, journalists, and organizations cited above for making AI environmental impact research accessible and transparent. Their work forms the foundation of this educational tool.