Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Productivity gains of artificial intelligence outweigh energy costs

Giovanni D'Elia

giovi.deliafi@gmail.com

Centre d'Economie de la Sorbonne, University Paris 1

Antoine Mandel

Centre d'Economie de la Sorbonne, Paris School of Economics, University Paris 1

Article

Keywords:

Posted Date: August 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10193525/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Productivity gains of artificial intelligience outweigh energy costs

Antoine Mandel[1,2]
Giovanni D'Elia[1]
[1]Centre d'Economie de la Sorbonne, Paris School of Economics, University Paris 1, France
[2]Climate Finance Alpha, 75009, Paris, France

June 2026

Abstract

This paper provides the first direct monetary comparison between the productivity gains and the energy costs of artificial intelligence adoption in the workplace. We estimate AI-induced time savings across 19,265 work tasks and 923 occupations. Under a quasi-full adoption scenario, aggregate wage savings reach approximately 8.87% of U.S. nominal GDP in 2024. On the energy side, we combine task-level execution frequencies with per-query energy consumption benchmarks to estimate total electricity cost savings. Energy costs remain well below 0.1% of GDP across all model and iteration assumptions, and productivity gains exceed direct operational energy costs by several orders of magnitude. Expressed in physical terms, AI assistance requires between 27 and 1,769 Wh per hour of labor saved, with a benchmark estimate of 430 Wh. Even accounting for environmental externalities, the energetic costs associated to AI adoption appear to be contained with respect to their benefits.

Introduction

Artificial intelligence is diffusing rapidly across workplaces, with adoption rates among U.S. firms rising sharply in recent years [1]. A growing body of experimental and quasi-experimental evidence documents that AI tools generate substantial productivity gains at the task level. Controlled experiments in professional writing tasks find that access to generative AI reduces task completion time by around 40%, with the largest gains concentrated among lower-performing workers [2]. Similar patterns emerge in software development, where AI-assisted programmers complete coding tasks significantly faster than unassisted controls [3], and in customer support settings [4]. These micro-level findings are broadly consistent across occupations and experimental designs, and are concentrated in cognitively intensive, white-collar tasks [5, 6].

At the same time, the energy footprint of AI systems is expanding at an unprecedented pace. Global data centre electricity consumption reached 415 TWh in 2024 and is projected to nearly double by 2030 [7]. Crucially, inference and usage now account for an estimated 60–90% of total energy consumption associated with large language models [8, 9], far exceeding the training costs that earlier studies had focused on [10, 11].

These two trends — rising AI-driven productivity and rising AI-driven energy demand — have so far been studied in isolation. On the productivity side, micro-level experimental studies quantify task-level time savings [2, 3], while macro-level approaches use task-based frameworks to estimate aggregate TFP gains, generally finding effects that are positive but modest [12]. On the energy side, existing work estimates the carbon and electricity footprints of AI infrastructure at the system or query level [13, 14], or projects future data centre demand at the sectoral and national level [7]. No study, to our knowledge, has directly compared the two on a common monetary scale for the aggregate economy.

This gap matters because the productivity benefits of AI adoption accrue primarily to firms, while the associated environmental costs are externalized onto the broader community. Even though the monetary cost of energy consumption is borne by AI-producing firms, its environmental impact falls on society at large. Without a joint assessment, neither policymakers nor firms can evaluate whether the net social return to large-scale AI adoption is positive.

Here, we provide a direct monetary comparison between the productivity gains and the energy costs of AI adoption in the U.S. workplace. On the productivity side, we adopt the task-based framework of [15] to estimate the time savings generated by AI at the occupational level. Following [5], these time savings are derived from GPT-5 evaluations of the degree to which AI can assist in performing each task in the O*NET database. Crucially, our estimates reflect a full-adoption scenario: we assume that all tasks identified as AI-exposed are effectively augmented by AI, representing an upper bound on the productivity gains that could be realised if AI were deployed at its maximum potential across U.S. workplaces. Aggregate productivity gains are then expressed in monetary terms as wage savings — the labor cost equivalent of the working time freed by AI assistance. On the energy side, we combine task-level execution frequencies with existing estimates of energy consumption per LLM query [14] to derive the total electricity cost of workplace AI adoption under the same full-adoption scenario, converted into dollars using average U.S. industrial electricity prices. In addition to this economy-wide comparison, we provide estimates of energy consumption in Wh per hour of labor saved — a normalized metric that directly captures the energetic cost of each unit of productivity gain generated by AI in the workplace.

Results

Task-level time savings

Figure 1 reports the distribution of GPT-5-based estimates of task-level time savings across all work tasks in the O*NET dataset. The distribution is markedly

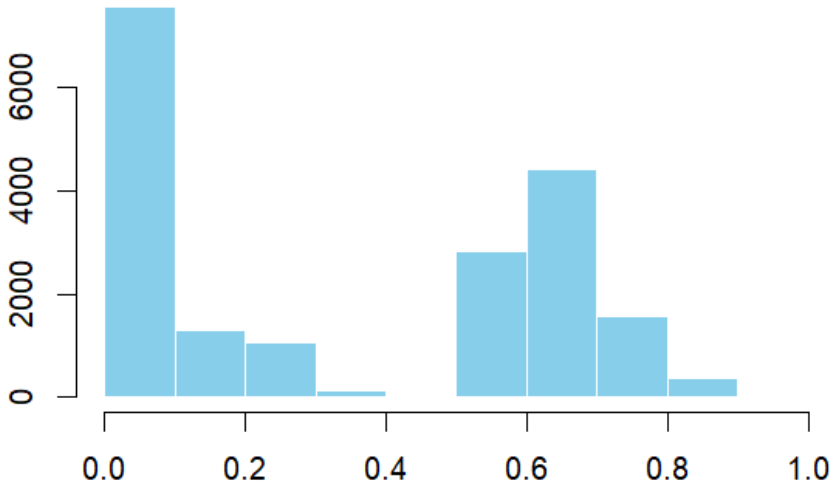


Figure 1: Task-level distribution of percentage time savings from AI usage

bimodal. A large mass of tasks clusters around values close to zero, corresponding predominantly to manual, physical, or highly context-specific activities. A second concentration emerges around time savings of 60–70%, tasks with substantial informational, cognitive, or textual content for which AI can meaningfully accelerate drafting, summarization, coding, or analytical processes. Only a very limited number of tasks display estimated time savings above 80%. Table 1 reports illustrative examples of tasks alongside their estimated time savings and AI exposure class, following the taxonomy of [5]. Tasks labelled E0 are those for which AI assistance alone yields time savings below 50%; E1 tasks exceed this threshold with access to an LLM chatbot alone; and E2 tasks fall below 50% savings with the chatbot alone, but surpass it when the model is augmented with additional software capable of processing task-specific data.

Occupation-level time savings

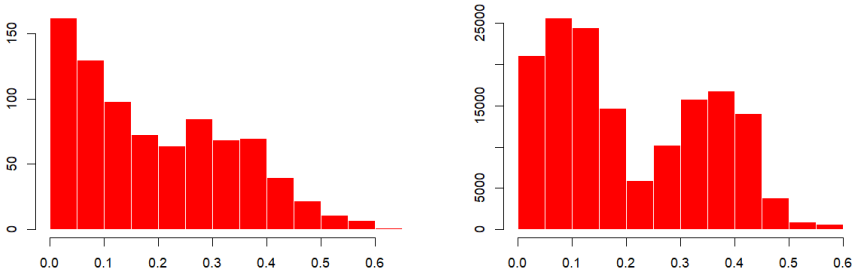
Figures 2a and 2b report the distribution of estimated time savings across the 832 occupation categories of the National Employment Matrix and across the corresponding U.S. workforce. The occupation-level distribution is right-skewed: only around 5% of occupation categories exhibit time savings exceeding 45%, while approximately half lie below 17%. The worker-level distribution displays

Table 1: Examples of task-level AI exposure classifications and estimated time savings

Occupation	Task	Class	Time savings
General Internal Medicine Physicians	Operate on patients to remove, repair, or improve the functioning of diseased or injured body parts and systems.	E0	0.00
Roofers	Install, repair, or replace single-ply roofing systems using waterproof sheet materials and asphaltic compositions.	E0	0.00
Welding, Soldering, and Brazing Machine Setters, Operators, and Tenders	Select, position, align, and bolt jigs, holding fixtures, guides, or stops onto machines using measuring instruments and hand tools.	E0	0.00
Parking Attendants	Escort customers to their vehicles to ensure their safety.	E0	0.00
Social and Community Service Managers	Participate in the definition of organizational policies regarding participant eligibility, program requirements, and benefits.	E1	0.62
Geographic Information Systems Technologists and Technicians	Develop specialized software routines, GIS databases, or business applications to customize geographic information.	E1	0.65
Fundraising Managers	Design and edit promotional publications such as brochures.	E1	0.66
Clinical Research Coordinators	Prepare study-related documentation, including procedural manuals, adverse event reports, and progress reports.	E1	0.64
Real Estate Brokers	Generate lists of properties for sale, including locations, descriptions, and financing options, using computers.	E2	0.85
Order Clerks	Calculate and compile order-related statistics and prepare reports for management.	E2	0.89
Criminal Justice and Law Enforcement Teachers, Postsecondary	Compile bibliographies of specialized materials for outside reading assignments.	E2	0.83
Order Clerks	Notify departments when supplies of specific items are low or when orders would deplete available inventories.	E2	0.90

a degree of bimodality, with mass clustering around a first group of workers in the 0–20% range and a second group centered around the 30–40% interval, suggesting heterogeneous AI exposure profiles across occupational groups.

Table 2 ranks the occupations most affected by potential AI-induced time savings, weighted by employment size. The top positions are occupied by cognitive, administrative, and managerial roles — software developers, accountants, management analysts, and clerical workers — which combine relatively high AI-exposure with large employment shares. Several occupations with lower estimated time savings also appear prominently due to workforce size alone, including home health aides and order fillers.



(a) Occupation-level **(b)** Population-weighted (thousands)

Figure 2: Distribution of AI-induced time savings at occupation level: unweighted vs population-weighted

Aggregate wage savings and energy costs

Translating occupation-level time savings into monetary terms using average occupational wages, the implied aggregate wage savings correspond to approximately 8.87% of U.S. nominal GDP in 2024 (8.78% under the alternative frequency-based task-weighting scheme). Following [5], we assume that all E1 tasks are fully AI-augmented, while E2 tasks — those requiring additional complementary software — are weighted at one half, reflecting the greater organizational and technological barriers to their adoption. This specification represents an upper bound on the productivity gains achievable under realistic near-term diffusion conditions.

Table 3 reports the estimated aggregate energy costs associated with this quasi-full adoption scenario, expressed as a share of 2024 nominal GDP. Energy consumption estimates per query are drawn from [14] and refer to long prompts of 10,000 input tokens and 1,500 output tokens, reflecting the complexity typically associated with professional tasks. Since energy consumption per query varies substantially across model architectures, and since completing a work-related task may require multiple prompt iterations, results are presented for the minimum, mean, and maximum per-query energy estimates among the models benchmarked by [14] (0.443, 7.066 and 29.078 Wh respectively), and under three alternative assumptions on the average number of query iterations per task (1, 4, and 7). Our benchmark scenario is the one with 7 query iterations per task completion and mean energy consumption across models.

Energy costs remain extremely small in macroeconomic terms across all scenarios, well below 0.1% of GDP. In absolute terms, total energy consumption under the 7-iteration scenario ranges from 1.7 to 115 TWh, with the mean being equal to 28 TWh — figures comparable to the annual electricity consumption of between 22,000 and one million and a half U.S. citizens, or equivalently between approximately 100 and 6,800 U.S. hospitals.

The contrast with the estimated wage savings is stark: productivity gains exceed direct operational energy costs by several orders of magnitude across all

Table 2: Top occupations by potential employment affected by AI-induced time savings

Rank	Occupation	Time savings	Workers affected
1	Office Clerks, General	0.32	810,984
2	Software Developers	0.42	698,348
3	Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	0.39	673,572
4	Bookkeeping, Accounting, and Auditing Clerks	0.45	649,172
5	Accountants and Auditors	0.43	617,841
6	First-Line Supervisors of Office and Administrative Support Workers	0.36	542,771
7	Business Operations Specialists, All Other	0.47	527,413
8	Sales Representatives of Services, Except Advertising, Insurance, Financial Services, and Travel	0.44	523,305
9	Project Management Specialists	0.44	443,717
10	Stockers and Order Fillers	0.15	428,300
11	Home Health and Personal Care Aides	0.11	421,056
12	Market Research Analysts and Marketing Specialists	0.47	400,816
13	Human Resources Specialists	0.43	390,124
14	Management Analysts	0.42	372,033
15	Miscellaneous Assemblers and Fabricators	0.17	252,971

model and iteration assumptions.

Finally, Table 4 reports energy consumption in Wh per hour of work saved, obtained by dividing aggregate AI-related energy use by aggregate hours saved. Under the 7-iteration scenario, these estimates range from 27 Wh/hour (LLaMA-3.1-8B) to 1769 Wh/hour (DeepSeek-R1(DS)). The average across models yields an estimate of 430 Wh, broadly comparable to 8 hours of laptop usage per hour of labor saved.

Discussion

To our knowledge, ours are the first estimates of AI-induced time savings across a comprehensive set of work tasks in the U.S. economy. The task-level results are broadly consistent with the existing literature. Tasks involving physical manipulation, direct human interaction, or highly context-specific operations — such as surgical procedures, roofing, or welding — receive zero estimated time savings, in line with evidence that AI exposure is lowest among occupations requiring manual dexterity, on-site presence, or interpersonal coordination [16, 17, 18, 19]. By contrast, tasks with substantial informational, textual, or analytical content display estimated time savings in the 60–70% range, consistent

	Min	Mean	Max
1 iteration	0.000069	0.001106	0.004551
4 iterations	0.000277	0.004423	0.018203
7 iterations	0.000485	0.007741	0.031855

Table 3: Estimated aggregate energy costs of workplace AI adoption in the U.S. under alternative per-query energy consumption scenarios and query iteration assumptions (as a percentage of 2024 U.S. GDP).

	Min	Mean	Max
1 iteration	4	61	252
4 iterations	15	246	1010
7 iterations	27	430	1769

Table 4: Estimated energy consumption in Wh per hour of work saved with AI usage in the U.S. economy under alternative per-query energy consumption scenarios and query iteration assumptions

with experimental evidence on writing [2], coding [3], and customer support [4]. The highest savings, reaching 85–90%, are concentrated among $E2$ tasks amplified by complementary software — activities that are highly routinized, information-dense, and require limited social or physical interaction.

The occupation-level rankings also reveal a pattern that pure exposure measures do not capture. The OECD AI Capability Gap index [19] identifies broadly similar occupational groups as most exposed — billing clerks, bookkeeping and auditing clerks, word processors, and data entry keyers all record a capability gap close to zero, converging with the administrative and clerical roles that feature prominently in our rankings. However, our top occupations are determined by the joint product of task-level time savings and employment size, which is why large-workforce occupations with modest individual savings — such as home health aides or stockers — appear prominently in our rankings, while they would rank low on the OECD measure. Managerial and executive roles present an instructive case: although the OECD index places them among the least exposed occupations, many of their tasks involve drafting, reporting, and information processing that can yield substantial AI-assisted time savings. They do not rank among the most affected in either framework, however, because a significant share of their work — human coordination, relationship management, strategic judgment — remains largely outside the reach of current AI systems [19]. This comparison underscores that the aggregate productivity impact of AI depends jointly on task-level exposure, workforce size, and the internal task composition of occupations, three margins that existing exposure measures typically address in isolation.

The energy estimates presented in this article refer specifically to the electricity consumed by AI-assisted task execution in U.S. workplaces under a scenario of potential — rather than realized — adoption. Following [5], we adopt the

intermediate $E1 + 0.5 \times E2$ specification: all tasks classified as $E1$ are assumed to be fully augmented by AI via LLM chatbot alone, while $E2$ tasks — those requiring additional complementary software — are weighted at one half, reflecting the greater difficulty of achieving full adoption in this category.

The closest paper to ours in scope and methodology is [20], who combine occupational AI exposure data from [5] with industry-level energy intensity data to estimate the aggregate increase in energy use and CO_2 emissions associated with AI adoption across the U.S. economy. Their top-down approach yields an aggregate estimate of 28 PJ per year ($\approx$ 7.8 TWh), equivalent to approximately 0.03% of annual U.S. energy use — a figure that falls within the range of our own estimates (1.7–115 TWh), providing indirect validation of our bottom-up approach. However, our benchmark scenario estimate is approximately four times larger than that of [20], at 28 TWh compared to their 7.8 TWh. The differences between the two estimates are consistent with methodological divergence: [20] derive energy consumption from industry-level energy intensity coefficients applied to AI-induced productivity gains, while we build our estimates from the ground up, combining task-level time savings with per-query energy consumption estimates and task execution frequencies. Additionally, their framework accounts for broader macroeconomic rebound effects — whereby productivity gains may expand economic output and thus total energy use — which are not captured in our benchmark calculation.

Compared to aggregate estimates of AI-related energy consumption in the United States, our figures are smaller. [21] report that U.S. data centers consumed 176 TWh in 2023, projected to reach 325–580 TWh by 2028, while IEA estimates place 2024 consumption at approximately 183 TWh [7]. The gap between these figures and ours is expected, for at least three reasons. First, data center consumption encompasses all digital workloads, including streaming, cloud storage, and enterprise computing, not only AI-related inference. Second, the aggregate figures include model training, which — although now accounting for a minority share of total AI energy use, estimated at 10–20% [22, 7] — still contributes substantially in absolute terms. Third, and most importantly, our estimates capture only workplace-related AI inference, excluding the rapidly growing consumer uses of AI — chatbots, AI-assisted search, image generation, and personal assistants — which represent an increasingly large share of total inference demand. Our estimates are therefore best interpreted as a lower bound on the share of U.S. data center energy consumption attributable to AI, specifically the portion linked to professional productivity gains.

At the global level, [7] estimate total data center electricity consumption at approximately 415 TWh in 2024 — about 1.5% of global electricity demand — projected to nearly double to 945 TWh by 2030. Our workplace-specific estimates for the U.S. alone amount to at most a few percent of this global total, reflecting both the scope restriction to professional tasks and the fact that the United States, while accounting for approximately 45% of global data center consumption, represents only a fraction of worldwide AI deployment.

Table 4 reports energy consumption in Wh per hour of labor saved — a metric particularly informative for evaluating the energetic efficiency of AI adop-

tion. Under the scenario of 7 query iterations per task, estimates range from 27 Wh (LLaMA-3.1-8B) to 1769 Wh (DeepSeek-R1(DS)) per hour of labor saved, with a mean of 430 Wh. To interpret these figures, consider a laptop consuming approximately 45–50 Wh per hour as a benchmark for the energy cost of computer-based work. Taking the mean model consumption scenario as reference, saving one hour of such work through AI assistance would consume 8 times the energy of that hour itself, implying that AI augmentation generates a net increase in energy consumption across virtually all occupational categories.

The case for AI adoption therefore cannot rest on direct energy savings, but must weight value of labor saved against the net energy cost. Our aggregate estimates speak to this: the monetary value of wage savings exceeds the monetary value of energy costs by several orders of magnitude at prevailing market prices. Combining our occupation-level time savings estimates with data on the energy intensity of specific job types would also allow researchers to identify where the net energetic cost of AI adoption is smallest and where productivity gains are most likely to justify the additional energy burden. To our knowledge, systematic data on energy consumption at the individual occupation level — capturing the average electricity used by a worker in performing a specific job — do not currently exist in a form suitable for this type of analysis, even though sector-level and building-level energy intensity data are available from sources such as the U.S. Energy Information Administration. Constructing such estimates would represent a valuable contribution at the intersection of labor economics and environmental accounting.

However, this comparison may be misleading. Market wages and electricity prices need not reflect the true social cost of the quantities being exchanged — in particular, the contribution of energy consumption to carbon emissions is largely not internalized in current electricity prices. In the context of an accelerating climate crisis, a purely monetary cost-benefit analysis risks systematically overstating the net benefits of large-scale AI diffusion. In order to move beyond a purely monetary assessment of AI adoption in workplaces and account for its broader societal implications, we examine how our energy consumption estimates translate into carbon emissions. Using Ember's data on grams of CO_2 lifecycle emissions per kilowatt-hour, we find that in our benchmark scenario — the one assuming 7 query iterations per task completion and mean energy consumption across models— aggregate emissions from AI adoption in workplaces amount to 10.7 Mt of CO_2. Compared to total annual U.S. emissions of approximately 4,900 Mt CO_2, these figures are remarkably small, accounting for only 0.21% of the total. Drawing on the EPA's 2023 estimate of the social cost of carbon ([23, 24]), set at \$204 per ton of CO_2, this translates into a social cost of approximately \$2.2 billion. The direct energetic costs associated with the same scenario amounted to approximately \$2.6 billion.

Including environmental externalities therefore nearly doubles our estimate of the total costs associated with AI adoption in workplaces, confirming the importance of accounting for such externalities when assessing the energy costs of AI. However, these combined costs remain far below the estimated monetary value of productivity gains, equal to approximately \$2.6 trillion. Our tentative

computations therefore suggest that the benefits of AI adoption greatly outweigh its energetic and environmental costs.

Data and Methods

Task-based framework and AI exposure classification

Following the task-based framework of [25] and [26], the analysis proceeds at the level of individual work tasks rather than occupations. This approach reflects the recognition that AI does not affect occupations as a whole, but rather the specific activities performed within them. The empirical foundation is the O*NET database, which provides detailed, survey-based descriptions of the tasks associated with each occupation in the United States. The dataset contains 19,265 tasks distributed across 923 occupations, each described by a median of approximately 20 tasks. Each task is uniquely assigned to a single occupation, allowing for precise occupational aggregation.

Each task is classified according to the three-tier exposure rubric introduced by [5]: $E0$ tasks are those for which AI assistance alone cannot reduce execution time by at least half; $E1$ tasks are those for which access to an LLM chatbot alone suffices to achieve this threshold; and $E2$ tasks are those that would not meet the threshold with the chatbot alone, but would do so with additional complementary software built on top of the LLM. This rubric is particularly well suited to the present analysis because it is intrinsically grounded in the concept of time savings: whether a task is considered AI-exposed depends directly on the degree to which AI reduces the time required for its execution. Classifications and task-level time savings are obtained via GPT-5, following the same prompt-based evaluation methodology of [5].

All occupation-level estimates reflect the $E1 + 0.5 \times E2$ scenario, in which all $E1$ tasks are assumed fully AI-augmented while $E2$ tasks are weighted at one half, reflecting the greater barriers to adoption associated with complementary software. This intermediate specification — originally proposed by [5] — is chosen because it represents the most plausible near-term adoption scenario: full $E2$ adoption requires the widespread deployment of task-specific software systems that remains more distant technologically and organizationally. Under this scenario, the estimates should be interpreted as an upper bound on the productivity gains achievable under realistic diffusion conditions.

Estimating aggregate wage savings

The aggregate monetary value of AI-induced time savings is computed as:

$$WS = \sum_{j \in J} L_j W_j \left(\sum_{i \in I_j} t_{i,j}\, h_{i,j} \right) \tag{1}$$

where $t_{i,j}$ is the estimated percentage reduction in execution time for task i in occupation j; $h_{i,j}$ is the share of annual working time devoted to that task;

L_j is total employment in occupation j; and W_j is the average hourly wage. Employment and wage data are drawn from the National Employment Matrix of the U.S. Bureau of Labor Statistics. Annual hours worked per occupation are computed from the Current Population Survey (CPS), using average weekly hours converted to annual totals.

Task weights $h_{i,j}$ are constructed using the Core–Supplemental classification provided by O*NET, assigning double weight to core tasks relative to supplemental ones, following [5] and [27]. This weighting scheme reflects the greater centrality of core tasks within each occupational bundle. As a robustness check, task weights are alternatively constructed using the O*NET Frequency variable in place of the Core–Supplemental classification. Results are largely insensitive to this choice.

It is important to stress that wage savings are employed solely as a normalization device to express productivity gains in monetary units comparable to energy costs. This does not imply that AI-induced time savings necessarily translate into reductions in employment or labor compensation. As extensively documented in the task-based literature, productivity improvements may materialize through a variety of channels — including output expansion, task reallocation, and the emergence of new activities — rather than through proportional reductions in labor demand [28, 15].

Estimating aggregate energy consumption

The aggregate energy cost of workplace AI adoption is estimated as:

$$EC = \sum_{j \in J} L_j \left(\sum_{i \in I_j} f_{i,j} \right) \cdot q \cdot r \cdot p \tag{2}$$

where $f_{i,j}$ is the annual number of executions of task i in occupation j; q is the energy consumed per AI query (in Wh); r is the average number of query iterations per task; and p is the industrial electricity price (USD per Wh). Task execution frequencies $f_{i,j}$ are derived from the O*NET Frequency variable, which reports the share of respondents indicating each of seven categories ranging from "yearly or less" to "hourly or more". These are mapped to approximate annual counts (1, 2, 20, 90, 300, 1,200, and 3,600 executions respectively), and task-level frequencies are computed as weighted averages across respondents.

Per-query energy consumption estimates are drawn from [14], who provide hardware-aware benchmarks linking model inference to electricity use across a range of LLMs. Their framework combines empirical benchmarks of model architectures with estimates of energy per token processed — accounting for both input and output sequence lengths — and maps computational requirements to data center energy intensity factors. The analysis is restricted to long-prompt scenarios (10,000 input tokens, 1,500 output tokens), reflecting the information-intensive nature of professional workplace interactions. In our calculations we employ the minimum, mean, and maximum values across the models they test.

The minimum corresponds to LLaMA-3.1-8B, the maximum to DeepSeek-R1, and the mean — 7.066 Wh per query — closely approximates the estimates for widely deployed frontier models such as GPT-4.1 and Claude 3.7 Sonnet. The substantial variation across models reflects differences in parameter count, architecture, and underlying hardware efficiency, and allows the results to span a broad range of plausible energy scenarios.

Since the number of query iterations required to complete a professional task is not directly observed, results are reported under three alternative assumptions: 1, 4, and 7 iterations. The upper bound of 7 is motivated by experimental evidence suggesting that user satisfaction in AI-assisted interactions peaks at approximately seven conversational turns [29, 30], while task-completion conversations typically involve between 3 and 7 exchanges [30]. Single-iteration estimates should be regarded as a lower bound, as professional tasks rarely require only one prompt to reach a satisfactory output. The industrial electricity price is set to 0.0813 USD/kWh, obtained from the U.S. Energy Information Administration. All data refer to 2024, the most recent year available for occupational employment and wage data.

References

[1] Kristina McElheran et al. "AI adoption in America: Who, what, and where". In: *Journal of Economics & Management Strategy* 33.2 (2024), pp. 375–415.

[2] Shakked Noy and Whitney Zhang. "Experimental evidence on the productivity effects of generative artificial intelligence". In: *Science* 381.6654 (2023), pp. 187–192.

[3] Sida Peng et al. "The impact of ai on developer productivity: Evidence from github copilot". In: *arXiv preprint arXiv:2302.06590* (2023).

[4] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. "Generative AI at work". In: *The Quarterly Journal of Economics* 140.2 (2025), pp. 889–942.

[5] Tyna Eloundou et al. "GPTs are GPTs: Labor market impact potential of LLMs". In: *Science* 384.6702 (2024), pp. 1306–1308.

[6] Emilio Colombo et al. "Towards the terminator economy: Assessing job exposure to ai through llms". In: *arXiv preprint arXiv:2407.19204* (2024).

[7] International Energy Agency. *Energy and AI*. Tech. rep. Paris: IEA, 2025. URL: `https://www.iea.org/reports/energy-and-ai`.

[8] Ana Paula Oliveira, Tânia Carraquico, and Clara Martinez-Perez. "Beyond Efficiency: A Systematic Review of Energy Consumption and Carbon Footprint Across the AI Lifecycle". In: *Sustainability* 18.3 (2026), p. 1359.

[9] Zhenya Ji and Ming Jiang. "A systematic review of electricity demand for large language models: evaluations, challenges, and solutions". In: *Renewable and Sustainable Energy Reviews* 225 (2026), p. 116159.

[10] Emma Strubell, Ananya Ganesh, and Andrew McCallum. "Energy and policy considerations for deep learning in NLP". In: *Proceedings of the 57th annual meeting of the association for computational linguistics*. 2019, pp. 3645–3650.

[11] David Patterson et al. "Carbon emissions and large neural network training". In: *arXiv preprint arXiv:2104.10350* (2021).

[12] Daron Acemoglu. "The simple macroeconomics of AI". In: *Economic Policy* 40.121 (2025), pp. 13–58.

[13] Francisco Caravaca, Ángel Cuevas, and Rubén Cuevas. "From Prompts to Power: Measuring the Energy Footprint of LLM Inference". In: *arXiv preprint arXiv:2511.05597* (2025).

[14] Nidhal Jegham et al. "How hungry is ai? benchmarking energy, water, and carbon footprint of llm inference". In: *arXiv preprint arXiv:2505.09598* (2025).

[15] Daron Acemoglu and Pascual Restrepo. "Tasks, automation, and the rise in US wage inequality". In: *Econometrica* 90.5 (2022), pp. 1973–2016.

[16] Daron Acemoglu et al. "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1 (2022), S293–S340.

[17] Edward Felten, Manav Raj, and Robert Seamans. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses". In: *Strategic management journal* 42.12 (2021), pp. 2195–2217.

[18] Michael Webb. "The impact of artificial intelligence on the labor market". In: *Available at SSRN 3482150* (2019).

[19] OECD. *The OECD AI Exposure Measure: Mapping the OECD AI Capability Indicators to Occupations*. Tech. rep. 59. Paris: OECD Publishing, 2026. DOI: 10.1787/f3da0f0a-en. URL: https://doi.org/10.1787/f3da0f0a-en.

[20] Anthony R Harding and Juan Moreno-Cruz. "Watts and bots: the energy implications of AI adoption". In: *Environmental Research Letters* 20.11 (2025), p. 114084.

[21] Arman Shehabi et al. "2024 united states data center energy usage report". In: (2024).

[22] Sasha Luccioni, Yacine Jernite, and Emma Strubell. "Power hungry processing: Watts driving the cost of AI deployment?" In: *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency*. 2024, pp. 85–99.

[23] Kevin Rennert et al. "Comprehensive evidence implies a higher social cost of CO2". In: *Nature* 610.7933 (2022), pp. 687–692.

[24] U EPA. *Epa report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances*. 2023.

[25] Daron Acemoglu and Pascual Restrepo. "The race between man and machine: Implications of technology for growth, factor shares, and employment". In: *American economic review* 108.6 (2018), pp. 1488–1542.

[26] Daron Acemoglu and Pascual Restrepo. "Robots and jobs: Evidence from US labor markets". In: *Journal of political economy* 128.6 (2020), pp. 2188–2244.

[27] Fabio Montobbio et al. "Labour-saving automation: A direct measure of occupational exposure". In: *The World Economy* 47.1 (2024), pp. 332–361.

[28] Daron Acemoglu and Pascual Restrepo. "Automation and new tasks: How technology displaces and reinstates labor". In: *Journal of economic perspectives* 33.2 (2019), pp. 3–30.

[29] Shih-Hong Huang et al. "How does conversation length impact user's satisfaction? A case study of length-controlled conversations with LLM-powered chatbots". In: *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 2024, pp. 1–13.

[30] Heung-Yeung Shum, Xiao-dong He, and Di Li. "From Eliza to XiaoIce: challenges and opportunities with social chatbots". In: *Frontiers of Information Technology & Electronic Engineering* 19.1 (2018), pp. 10–26.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- NaturearticleSupplementaryMaterials.pdf