

The Impending Cognitive Hyperabundance

Executive summary

At Thorsten Meyer AI, we use “cognitive hyperabundance” to mean a near-term regime shift in which high-quality cognitive outputs—drafts, analyses, code, designs, planning artifacts, synthesis, and decision support—become widely available at marginal costs that approach “effectively free” for many use cases. This shift is driven by two coupled curves: accelerating model capability (software) and accelerating AI infrastructure buildout (hardware, data centers, power, and cooling). Evidence from 2024–2026 strongly supports the infrastructure half of the thesis—capital formation and energy procurement are scaling at “mega-project” magnitudes—while the labor-market half is best described as rapid task reconfiguration and value compression rather than immediate “human knowledge work goes to zero.” ¹

We introduce “vesperance” as a cultural concept: wistful nostalgia for the present as an era ends. In organizational terms, vesperance is what it feels like when inherited life scripts (education → career ladder → retirement) lose predictive power faster than replacement norms can form. Multiple large surveys and institutional risk assessments are consistent with the thesis that work strain and social trust pressures are already elevated, even before any hypothetical “full automation”: (a) global employee engagement fell to 21% in 2024, with large economic costs attributed to disengagement; (b) global knowledge workers report widespread generative AI use alongside persistent workload strain; and (c) leading global risk assessments place misinformation/disinformation—amplified by AI—among the top near-term threats to social cohesion. ²

The strongest empirical evidence on knowledge work automation comes from randomized field experiments. Results are material but not uniformly transformational: one cross-industry experiment (66 firms; 7,137 knowledge workers) found that active users saved about two hours per week on email and reduced after-hours work, but did not detect large compositional task shifts over the experimental horizon; similarly, field experiments in software development show meaningful productivity gains for adopters, with heterogeneous impacts by seniority and usage intensity. ³

Macroeconomic consequences depend on whether AI behaves primarily as task substitution (automation) or task complementarity (augmentation). Task-based macro models argue the near-term aggregate productivity effect may be modest absent large complementary investments and diffusion; however, “wage collapse” scenarios are possible under aggressive assumptions where automation outpaces capital accumulation and makes labor abundant relative to AI-capital. The credible stance is scenario planning, not single-track prediction. ⁴

Actionably, the strategic response for Thorsten Meyer AI is not to mourn knowledge work, but to (1) measure cognitive hyperabundance with operational metrics (cost-per-decision, latency-to-draft, quality-per-watt, and adoption diffusion), (2) design policy and ownership proposals that broaden access to AI-capital and mitigate distributional shocks, and (3) communicate with “vesperance-aware” narratives that acknowledge grief and disorientation while offering concrete institutional pathways to agency. ⁵

Cognitive hyperabundance and vesperance

Cognitive hyperabundance is not the claim that “intelligence is solved.” It is the claim that the realized supply of useful cognition—expressed as economically actionable artifacts—will rise faster than institutions can re-price it. The proximate mechanisms are: rapid scaling of training and inference compute; falling effective cost of many cognitive tasks; and widespread tool diffusion into daily workflows for writing, planning, summarization, and software production. ⁶

The evidence for accelerating capability trends is mixed in interpretation but clear in direction. Training compute for frontier models has been estimated to grow on the order of 4–5× per year through May 2024, while broader AI Index reporting describes training compute for notable AI models doubling on roughly a five-month cadence; the same reporting notes sustained improvements in hardware performance and energy efficiency alongside declining price-performance. These trends don’t guarantee “full automation,” but they do guarantee continuously expanding feasible task coverage at a given quality threshold. ⁷

“Vesperance” functions here as a cultural symptom and a diagnostic tool: wistful nostalgia for the current social order while recognizing it is already leaving. The core claim is that such a mood becomes widespread when people confront (a) accelerating capability change, (b) ambiguous control over its direction, and (c) eroding confidence that conventional scripts yield stable outcomes. As a cultural term, it has a defined usage in online discourse aligned with this meaning, which we treat as a lens rather than as a validated clinical construct. ⁸

Infrastructure buildout and mega-project comparisons

AI’s near-term impact is becoming legible as infrastructure and energy, not merely as software. The hardest constraints are increasingly physical: grid interconnect queues, generation build times, cooling and water availability, semiconductor packaging capacity, and the capital market’s willingness to fund multi-year buildouts.

Scale and pace: investment and capex. A Reuters analysis estimated that as of 2024, investors had put nearly \$1.6 trillion into the AI boom since 2013 and that forecasts for 2025 added roughly \$375 billion—enough that a single year of AI investment could exceed the total cost of the Apollo program (as typically costed). Separately, Reuters reporting in early 2026 described planned AI-related spending splurges by major technology firms on the order of \$600 billion in 2026, while Goldman Sachs research commentary projected more than \$500 billion of AI investment as soon as 2026 and discussed power access as a gating factor. ⁹

Assumptions and comparability pitfalls. “AI investment” aggregates disparate components (venture, private equity, public markets, capex, and operating spend) and can be global in scope, while historic mega-projects were typically national, state-led, and measured in budget outlays. This paper therefore uses a transparent comparison frame:

- *Global vs national:* AI infrastructure is global and multi-actor; Apollo and the Manhattan Project were U.S.-led national programs. ¹⁰
- *Annual vs cumulative:* Many AI figures are annualized capex or yearly investment flows; Apollo and the Manhattan Project are commonly cited as cumulative costs across their lifetimes. ¹¹

- *Nominal vs real dollars:* Historic programs require inflation adjustments; published comparisons often mix methodological conventions. Where possible, we cite sources that themselves make the comparison, rather than asserting novel inflation conversions. ¹¹

Historical anchors for magnitude. For reference points: (a) the Manhattan Project cost roughly \$2 billion by 1945 (often expressed as over \$30 billion in 2023 USD), and employed around 130,000 workers at peak; (b) Apollo’s total cost is commonly reported around \$25.4–\$25.8 billion in 1960–1973 outlays; and (c) the Federal Highway Administration’s final pre-completion estimate for the Interstate System (1991) was about \$128.9 billion total cost (in then-year dollars). ¹²

Energy as the binding constraint. The International Energy Agency projects that global data center electricity consumption more than doubles to around 945 TWh by 2030 in its base case (just under 3% of global electricity), with roughly ~15% annual growth from 2024–2030—more than four times the growth of total electricity consumption from other sectors. It also projects global electricity generation to supply data centers rising from about 460 TWh in 2024 to over 1,000 TWh in 2030 and 1,300 TWh in 2035. ¹³

U.S.-specific projections illustrate the same constraint: the Electric Power Research Institute has been cited as estimating data centers could reach up to ~9% of U.S. electricity generation annually by 2030, up from recent levels—an estimate echoed in U.S. Department of Energy materials. ¹⁴

Power procurement as strategy, not compliance. Recent reporting shows hyperscalers shifting into demand-response deals and large-scale power contracting as a reliability and cost strategy. For example, Google’s demand-response agreements with multiple utilities were reported as enabling up to 1 GW of curtailable load—illustrating that compute is now negotiating directly with grid physics. ¹⁵

Cooling and density: the end of “easy scaling.” The Uptime Institute’s 2024 survey reports that average rack densities are rising but remain below 8 kW, and most facilities do not have racks above 30 kW—yet this is expected to change as AI workloads push density upward. Uptime also reports industry-average PUE holding around ~1.56, implying efficiency gains have plateaued—so incremental compute increasingly means incremental power and heat, not “free” efficiency. ¹⁶

Chip economics and vendor signals. Hardware vendor financial filings reinforce that this buildout is not hypothetical. NVIDIA reported fiscal year 2026 revenue of \$215.9 billion (up 65% year-over-year) and stated that Data Center revenue grew strongly year-over-year, driven by accelerated computing and AI platform shifts. These figures matter because they are downstream of real purchase orders, not survey intentions. ¹⁷

Evidence on AI automation of knowledge work

A rigorous evaluation requires separating (a) *capability* (what systems can do in principle) from (b) *adoption* (what organizations deploy at scale) and from (c) *labor-market incidence* (how wages, tasks, and employment adjust). The strongest evidence in 2024–2026 comes from randomized controlled trials and field experiments in realistic work settings.

Cross-industry knowledge work: time savings without immediate task collapse. A large field experiment across 66 firms and 7,137 knowledge workers randomly provided access to an integrated generative AI tool for email, meetings, and writing. In the second half of the six-month experiment, about

80% of treated workers used the tool and spent two fewer hours on email per week, with reduced after-hours work; the study did not detect shifts in the quantity or composition of tasks attributable to individual-level AI provision over the experiment window. This supports a “compression” thesis (less time per task category) more than a “near-term disappearance” thesis (roles evaporate immediately). ¹⁸

Teamwork and expertise: AI as a partial substitute for coordination. A preregistered field experiment with 776 professionals at Procter & Gamble tested individuals vs. teams, with and without generative AI, on real innovation challenges. The paper reports that individuals with AI matched the performance distribution of teams without AI and that generative AI altered the functional signature of solutions (reducing siloed R&D vs commercial patterns). Quantitatively, it reports performance improvements on the order of ~0.24 standard deviations for teams without AI relative to individuals working alone (baseline), versus ~0.37 standard deviations for individuals with AI; it also reports increased likelihood of top-decile outputs for AI-enabled teams relative to control. The implication is not “humans are useless,” but “the marginal value of certain forms of collaboration and middle-layer synthesis can compress sharply.” ¹⁹

Software engineering: measurable gains, heterogeneous by adoption and seniority. In large-scale randomized trials of a coding assistant across multiple companies, instrumental-variable estimates suggest that usage caused roughly a 26.08% increase (SE ~10.3%) in weekly completed tasks for adopters, with secondary indicators such as commits and compilations also increasing. This is substantial, but it is not infinite, and it is conditional on usage, workflow integration, and task mix. ²⁰

A Bank for International Settlements working paper reporting a field experiment in a large tech environment found the introduction of an LLM coding tool increased code output by about 55% on average, with effects concentrated among junior staff (reported ~67% increase for juniors), while measured impacts for more senior employees were not statistically significant primarily due to lower usage. This reinforces a key “hyperabundance mechanism”: AI can compress the experience gradient by enabling less experienced workers to perform closer to the frontier—changing hiring, training, and wage ladders even if total headcount doesn’t collapse immediately. ²¹

Adoption diffusion: fast, but not yet “fully utilized capability.” Survey-based evidence indicates rapid uptake of generative AI tools. Microsoft/LinkedIn’s 2024 Work Trend Index reports 75% of global knowledge workers using generative AI and high rates of “bring your own AI,” while the OECD’s worker and employer surveys report that around 80% of AI users say AI improves their performance and that AI users are more likely to report improved working conditions than worsened. These are not causal estimates, but they matter because hyperabundance requires diffusion, not just labs. ²²

What the evidence does not yet show. Across these studies, two patterns repeat: (1) productivity effects are meaningful but vary widely; (2) organizations often realize time savings before they redesign jobs. This aligns with the OECD’s synthesis stating there is “little evidence of a net negative impact of AI on the number of jobs” so far, while emphasizing substantial automation risk and the possibility of frictional unemployment due to speed of adoption. ²³

Economic scenarios and distributional outcomes

The economic question is not “can AI do tasks,” but “how do markets, organizations, and policies translate task change into wages, employment, and distribution.” The pivot variable is whether AI is deployed

primarily as labor-replacing automation or as labor-augmenting capability that expands output and creates new tasks.

Task substitution vs complementarity. OECD analysis frames AI's labor effects through substitution (AI performs tasks instead of humans, substituting labor with capital) versus augmentation (AI complements workers and shifts task composition). It emphasizes uneven adoption, concentration among large firms, and the risk that AI could widen inequality if access and complementary investments are uneven. ²⁴

Macro skepticism under realistic diffusion constraints. Task-based macroeconomic modeling emphasizes that large economy-wide productivity booms are not automatic; they depend on which tasks are automated, the cost of deployment, and complementary organizational investments. This provides an analytic "brake" on claims that cognitive hyperabundance instantly yields universal prosperity—or instant labor collapse. ²⁵

Wage-collapse scenarios under aggressive assumptions. In contrast, Korinek and Suh model scenarios for a transition to AGI-like automation in which outcomes hinge on a "race between automation and capital accumulation." Under certain trajectories (rapid automation of a very large share of tasks), wages can collapse as labor becomes effectively abundant relative to AI-capital; alternative trajectories produce sustained wage growth if there is always a thick tail of tasks that remain human-complementary or if automation is slower. This is not a forecast; it is a map of possible regimes. ²⁶

Exposure estimates and inequality risk. The IMF estimates that nearly 40% of global employment is exposed to AI, with advanced economies around ~60% exposed; it argues that roughly half of exposed jobs may benefit from AI integration while the other half could see reduced labor demand, lower wages, or reduced hiring in some cases. This is a distributional warning: even if aggregate productivity rises, incidence can be highly unequal. ²⁷

Scenario table for 2026–2035

The table below is explicitly conditional: it represents plausible alternative worlds based on key uncertainties (capability pace, energy constraints, diffusion speed, and policy response). Energy projections anchor to IEA base-case pathways; labor-market indicators are stylized ranges consistent with the modeling literature and early empirical studies.

| Scenario | Core assumption set | Timeline markers | Key indicators |
|---|--|---|---|
| Optimistic (augmentation dominates) | AI diffuses widely as a productivity tool; complementary investments (process redesign, training, governance) are made in parallel; policy expands access to AI tools and mitigates inequality | 2026–2028: rapid adoption; 2028–2030: organizational redesign wave; 2030–2035: stable AI governance and energy buildout | Data center electricity: ~945 TWh by 2030 (IEA base case), rising toward ~1,300 TWh by 2035. ²⁸ Knowledge work productivity: sustained task-level gains (e.g., email time savings ~2 hrs/week among active users; software task completion gains ~20–30% for adopters). ²⁹ Unemployment: modest frictional increases, then normalization (policy-dependent). ³⁰ |
| Baseline (mixed substitution/augmentation) | AI automates a meaningful subset of cognitive tasks; many firms bank time savings before redesigning roles; labor markets reprice skills with polarization; governance matures unevenly by sector | 2026–2027: tool saturation in office workflows; 2028–2030: skill ladder compression; 2030–2035: new norms for verification/trust emerge | Data center electricity: IEA base case remains a central planning anchor; U.S. data centers plausibly approach high-single-digit share of electricity by 2030 in some forecasts. ³¹ Labor share pressure in exposed occupations; inequality risk rises if AI-capital ownership is concentrated. ³² |
| Pessimistic (automation outruns institutions) | Rapid capability + deployment enables large-scale task substitution; capital owners capture most gains; energy and grid bottlenecks create price shocks; governance fails to contain trust erosion | 2026–2028: aggressive cost-cutting deployments; 2028–2030: labor market shock in exposed services; 2030–2035: institutional catch-up under stress | Wage-collapse becomes plausible in model space under aggressive automation trajectories, depending on “race” conditions vs capital accumulation. ²⁶ Social trust: misinformation/disinformation remains a top near-term systemic risk, amplifying polarization and governance instability. ³³ Energy volatility: grid and supply constraints force demand response and potentially higher fossil dispatch in some regions over short horizons. ³⁴ |

Projected milestones through 2035

| |
|--|
| <p>timeline</p> <p>title Projected milestones for cognitive hyperabundance (2026–2035)</p> <p>2026 : Hyperscaler capex surge; power becomes strategic constraint; agentic workflows expand in pilots</p> |
|--|

2027 : High-density AI racks and liquid cooling mainstream in new builds; governance enforcement intensifies in major jurisdictions
2028 : Widespread job redesign in early-adopter sectors; verification/trust tooling becomes mandatory in many workflows
2030 : Data center electricity demand roughly doubles versus mid-2020s baselines in leading scenarios; grid interconnect becomes pacing item
2032 : Large-scale transmission and generation additions come online; compute pricing reflects energy locality and reliability
2035 : Data center electricity supply requirement approaches ~1.3 PWh/year in base-case energy projections; mature post-labor institutions (or crisis response) become decisive

Psychological and social impacts of norm disruption

The social thesis of cognitive hyperabundance is that rapid re-pricing of cognitive labor destabilizes norms faster than new norms can form—producing vesperance as an increasingly common emotional context.

Burnout and baseline strain. The World Health Organization defines burn-out (ICD-11) as a syndrome resulting from chronic workplace stress that has not been successfully managed, characterized by exhaustion, mental distance/cynicism, and reduced professional efficacy. This is relevant because cognitive hyperabundance can increase both work intensity (more outputs possible) and evaluation pressure (more comparison points), even when tools save time. ³⁵

Engagement as a leading indicator of “script stress.” Gallup’s State of the Global Workplace reporting indicates global engagement fell to 21% in 2024 (from 23%), with large estimated productivity losses. Engagement declines at this scale are consistent with widespread disorientation at work, regardless of whether AI is the primary cause. In a vesperance frame, it is the felt mismatch between effort and stable trajectory that matters. ³⁶

AI adoption colliding with workload reality. Microsoft’s Work Trend Index reports very high generative AI use among knowledge workers (75%) and high “BYOAI” behavior, while simultaneously framing organizational readiness gaps: workers move faster than governance, process redesign, and leadership alignment. This is a recipe for norm conflict (shadow tooling, uneven expectations, trust issues in outputs) and for a cultural mood that “the old rules no longer apply.” ³⁷

Trust erosion as a macro-social risk. The World Economic Forum’s Global Risks Report 2024 identifies misinformation and disinformation as a major near-term global risk, explicitly in a context where AI lowers the cost of producing persuasive falsehoods at scale. Cognitive hyperabundance therefore includes an “epistemic externality”: abundance of content can mean scarcity of trust, which destabilizes norms in media, politics, science communication, and institutional legitimacy. ³⁸

A note of rigor: what we can and cannot attribute. It is tempting to causally attribute burnout or norm disruption to AI. The more defensible claim is that AI accelerates a pre-existing vector: information overload, attention capture, and rising performance pressure. In that sense, vesperance is not “AI sadness,” but “transition grief” under high uncertainty and low perceived control. ³⁹

Creativity and the arts under abundance

The arts sit at the emotional center of vespérance because they are where humans try to preserve meaning when instrumental scripts fail. Here the key analytic distinction is between *combinatorial possibility* (how many artifacts could be produced) and *social value* (what audiences attend to, trust, and integrate into identity).

From scarcity of production to scarcity of attention and provenance. The economic logic of the attention economy predicts that as content supply grows, attention becomes the binding constraint; value shifts toward curation, trust signals, community belonging, and authenticity. Cognitive hyperabundance makes this dynamic acute: synthetic content is not just plentiful, it is cheap and personalized. ⁴⁰

Early evidence of creative market reorganization. A working paper on U.S. artists finds little evidence of short-run earnings declines associated with LLM exposure through 2023 (point estimates near zero or modestly positive in some specifications), while employment effects are more mixed, with weaker employment growth in 2023 for more exposed artistic occupations in some models. The author frames this as early task reallocation rather than immediate earnings collapse—consistent with a “reorganization technology” view. ⁴¹

Simultaneously, field evidence from creative goods platforms has been summarized as showing that when generative AI content entered a marketplace, supply surged and human creators were squeezed—highlighting that market-level dynamics can harm human producers even if some consumers gain from variety and lower prices. ⁴²

Copyright and compensation pressure. Legal scholarship in 2024 describes digital creative content as “virtually infinite,” shifting cultural markets from scarcity to abundance and stressing that generative AI risks displacing creative labor and undermining incentives unless compensation mechanisms (e.g., royalty funds) are designed. This is an institutional design problem, not an artistic one. ⁴³

Creator income at risk: one quantified forecast. A global study commissioned by CISAC projects rapid growth in markets for generative-AI music and audiovisual content—with estimates rising from roughly €3 billion to €64 billion by 2028—and argues that creators face substantial revenue risk (e.g., 24% in music; 21% in audiovisual) under an unchanged regulatory framework. While projections should be treated cautiously, the study is directionally aligned with the “abundance → bargaining shift” thesis. ⁴⁴

What remains uniquely human (in practice). In a hyperabundant cognitive and creative environment, “uniquely human” does not mean “humans alone can generate artifacts.” It means the scarce goods are increasingly: embodied presence, trusted identity, moral accountability, local community, and culturally legitimate meaning-making. These are precisely the domains where vespérance can be metabolized into constructive creation rather than paralysis. ⁴⁵

Policy and institutional responses for the post-labor transition

Cognitive hyperabundance is not a “technology story” alone; it is a distribution and governance story. The policy question is how to shape investment, ownership, and trust frameworks so that productivity gains do not translate into destabilizing inequality and norm collapse.

Energy and infrastructure governance. Because power is a binding constraint, “AI policy” is now partly “grid policy.” Recent U.S. reporting highlights both the risk of rising fossil generation under load growth and the emergence of grid upgrades and demand-response contracting as mitigation tools. This suggests policymakers should treat data center interconnection, transmission upgrades, and efficiency standards as first-class components of AI governance. ⁴⁶

Regulation: risk-based governance is already here. The European Commission reports that the EU AI Act entered into force on August 1, 2024 and is being phased in. Whatever one’s normative stance, this establishes a template: prohibited practices, obligations for general-purpose AI, and enforceable requirements for high-risk systems on fixed timelines. Post-labor economics must assume regulatory heterogeneity across jurisdictions, not a single global rulebook. ⁴⁷

Labor market transition policies grounded in current evidence. The OECD’s 2024 brief argues that while net job loss evidence is limited so far, the risk of automation is substantial, with highly automatable occupations estimated around 27% of employment in OECD framing, and with strong worker anxiety about job loss and wage impacts. It explicitly recommends monitoring, training, worker consultation, and social protection—policies that become more urgent as cognitive hyperabundance accelerates. ²³

Distributional architecture: who owns AI-capital? The IMF warns AI may worsen inequality without policy, because exposed jobs split between complementarity and displacement. OECD analysis similarly emphasizes concentration of AI development and uneven diffusion. A post-labor agenda therefore needs concrete mechanisms for broadening the ownership of AI-capital (compute, models, data advantages) or redistributing its rents. ⁴⁸

Actionable recommendations for Thorsten Meyer AI

Research agenda (measurement before ideology).

Build a “Cognitive Hyperabundance Index” with quarterly updates, combining (a) compute supply indicators (capex, rack density proxies, power procurement), (b) effective cost metrics (cost per 1,000 high-quality tokens or equivalent task units, latency-to-first-draft), (c) adoption diffusion (survey + telemetry where available), and (d) labor outcomes (task reallocation, within-firm productivity, wage compression signals). Anchor energy and infrastructure baselines to IEA trajectories and industry operations surveys. ⁴⁹

Policy proposals (post-labor economics that can be piloted).

1) Propose an “AI-capital dividend” framework: modest levies on large-scale compute rents (or on energy-intensive inference) earmarked for (a) transition income supports, (b) retraining that includes complementary human skills, and (c) public-interest digital infrastructure (verification tooling, provenance standards). The justification is distributional: AI exposure is high and uneven; the goal is to broaden the share of gains. ⁴⁸

2) Advocate “broad-based compute ownership” pilots: municipal or national compute trusts, public-private co-investment vehicles for data center capacity tied to public service obligations (energy efficiency, transparency, auditability), and employee equity participation where feasible. This responds directly to the concentration risks emphasized in OECD analysis and the physical scarcity of power access. ⁵⁰

3) Push for “trust infrastructure” as economic infrastructure: watermarking/provenance, audit trails for high-stakes uses, and verified output pipelines in government and education—aligned with the WEF’s framing of misinformation/disinformation as a leading near-term risk. ⁵¹

Communications strategy (vesperance-aware without fatalism).

Adopt “vesperance” explicitly as a bridge term: acknowledge grief for the expiring regime of knowledge-work scarcity while offering a concrete social contract for abundance. Frame changes in three layers—(1) physics (power, chips, cooling), (2) economics (task repricing and ownership), and (3) meaning (trust, art, community). Use scenario language rather than prophecy, and publish commitments to measurement and revision as new evidence arrives. ⁵²

1 9 10 11 <https://www.reuters.com/graphics/USA-ECONOMY/AI-INVESTMENT/gkvlqbgxkpb/>
<https://www.reuters.com/graphics/USA-ECONOMY/AI-INVESTMENT/gkvlqbgxkpb/>

2 36 <https://www.gallup.com/workplace/349484/state-of-the-global-workplace.aspx>
<https://www.gallup.com/workplace/349484/state-of-the-global-workplace.aspx>

3 18 29 https://www.nber.org/system/files/working_papers/w33795/w33795.pdf
https://www.nber.org/system/files/working_papers/w33795/w33795.pdf

4 25 **The Simple Macroeconomics of AI**
https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf?utm_source=chatgpt.com

5 24 32 https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/04/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_d54e2842/8d900037-en.pdf
https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/04/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_d54e2842/8d900037-en.pdf

6 7 <https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year>
<https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year>

8 https://www.reddit.com/r/singularity/comments/16kev5c/has_anyone_had_this_peculiar_feeling_lately/
https://www.reddit.com/r/singularity/comments/16kev5c/has_anyone_had_this_peculiar_feeling_lately/

12 <https://www.nps.gov/mapr/faqs.htm>
<https://www.nps.gov/mapr/faqs.htm>

13 **Energy demand from AI**
https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai?utm_source=chatgpt.com

14 <https://www.reuters.com/business/energy/data-centers-could-use-9-us-electricity-by-2030-research-institute-says-2024-05-29/>
<https://www.reuters.com/business/energy/data-centers-could-use-9-us-electricity-by-2030-research-institute-says-2024-05-29/>

15 34 <https://www.reuters.com/sustainability/boards-policy-regulation/google-expands-utility-deals-curb-datacenter-power-use-during-peak-demand-2026-03-19/>
<https://www.reuters.com/sustainability/boards-policy-regulation/google-expands-utility-deals-curb-datacenter-power-use-during-peak-demand-2026-03-19/>

16 <https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf>
<https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf>

17 <https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm>
<https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm>

19 https://www.nber.org/system/files/working_papers/w33641/w33641.pdf
https://www.nber.org/system/files/working_papers/w33641/w33641.pdf

- 20 https://demirermert.github.io/Papers/Demirer_AI_productivity.pdf
https://demirermert.github.io/Papers/Demirer_AI_productivity.pdf
- 21 <https://www.bis.org/publ/work1208.pdf>
<https://www.bis.org/publ/work1208.pdf>
- 22 37 https://marketingassets.microsoft.com/gdc/gdcAev8aq/original?utm_source=chatgpt.com
https://marketingassets.microsoft.com/gdc/gdcAev8aq/original?utm_source=chatgpt.com
- 23 30 39 https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/using-ai-in-the-workplace_02d6890a/73d417f9-en.pdf
https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/using-ai-in-the-workplace_02d6890a/73d417f9-en.pdf
- 26 **Scenarios for the Transition to AGI**
https://www.nber.org/system/files/working_papers/w32255/w32255.pdf?utm_source=chatgpt.com
- 27 48 <https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf>
<https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf>
- 28 31 49 **Executive summary – Energy and AI – Analysis - IEA**
https://www.iea.org/reports/energy-and-ai/executive-summary?utm_source=chatgpt.com
- 33 38 51 https://www3.weforum.org/docs/WEF_The_Global_Risks_Report_2024.pdf
https://www3.weforum.org/docs/WEF_The_Global_Risks_Report_2024.pdf
- 35 <https://www.who.int/standards/classifications/frequently-asked-questions/burn-out-an-occupational-phenomenon>
<https://www.who.int/standards/classifications/frequently-asked-questions/burn-out-an-occupational-phenomenon>
- 40 45 <https://journals.sagepub.com/doi/abs/10.1177/20539517241275878>
<https://journals.sagepub.com/doi/abs/10.1177/20539517241275878>
- 41 https://www.econstor.eu/bitstream/10419/336068/1/cesifo1_wp12368.pdf
https://www.econstor.eu/bitstream/10419/336068/1/cesifo1_wp12368.pdf
- 42 <https://www.gsb.stanford.edu/insights/when-ai-generated-art-enters-market-consumers-win-artists-lose>
<https://www.gsb.stanford.edu/insights/when-ai-generated-art-enters-market-consumers-win-artists-lose>
- 43 <https://academic.oup.com/grurint/article/73/12/1137/7832810>
<https://academic.oup.com/grurint/article/73/12/1137/7832810>
- 44 <https://www.cisac.org/Newsroom/news-releases/global-economic-study-shows-human-creators-future-risk-generative-ai>
<https://www.cisac.org/Newsroom/news-releases/global-economic-study-shows-human-creators-future-risk-generative-ai>
- 46 **US power demand surge from data centers could lift fossil fuel generation, EIA says**
https://www.reuters.com/business/energy/us-power-demand-surge-data-centers-could-lift-fossil-fuel-generation-eia-says-2026-03-12/?utm_source=chatgpt.com
- 47 https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en
https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en

⁵⁰ https://www.oecd.org/en/publications/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_8d900037-en.html

https://www.oecd.org/en/publications/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_8d900037-en.html

⁵² Energy and AI – Analysis - IEA

https://www.iea.org/reports/energy-and-ai?utm_source=chatgpt.com