The Black Box That Runs on Sweat: AI's Hidden Supply Chain

Weekly Analysis — https://ainews.social

Pull up any vendor page for a contemporary AI tool and you will encounter the same rhetorical sleight of hand. The product is described as if it does the work itself. Microsoft's documentation tells you that Copilot "summarizes the report"; it does not tell you which data center spun up which racks for how many seconds, what watts were drawn, whose annotated examples taught the model what a summary is, or which subcontracted workforce in Nairobi or Manila sat through the violent, pornographic, or merely tedious training data so that the chatbot could refuse to produce it for you on demand. The tool is presented as an autonomous agent. It is in fact a node — the visible, monetizable endpoint of a long and largely invisible extraction network.

This is the genre convention of the tool documentation page, and once you notice it you cannot stop noticing it. The [6] presents the product as a discrete object with discrete features, much the way an appliance manual presents a microwave. The Spanish-language overview [30] does the same, listing capabilities — "draft," "summarize," "analyze" — as if these verbs described the artifact itself rather than a planetary apparatus of compute, labor, and data that produces them on request. The service description [20] goes further: it specifies licensing prerequisites, tenant requirements, and admin controls with admirable precision while leaving the actual content of what Copilot does — the model weights, the training corpus, the failure rates — as a black box you are licensed to query but not to inspect.

[6] Application card: Microsoft 365 Copilot

[30] ¿Qué es Microsoft 365 Copilot? | Microsoft Learn

[20] Microsoft 365 Copilot - Service Descriptions | Microsoft Learn

Kate Crawford's [27] names this rhetorical operation directly: the description of AI as "fundamentally abstract" — as software, as intelligence, as service — "distances it from the energy, labor, and capital needed to produce it." The tool's interface is engineered to absorb your attention while the supply chain stays offstage. This essay is an attempt to bring the offstage onstage: to look at the tool not as what the vendor says it is, but as what the evidence — including, importantly, the vendor's own documentation read against the grain — shows it to be.

[27] The Atlas of AI

What the Tool Page Is Actually Selling

The honest way to read a vendor documentation page is as a sales artifact dressed in the costume of technical reference. The grammar gives it away. The Power BI Copilot guide, [26], tells the user how to invoke a feature; it does not tell the user how often the resulting summary will be wrong, what kinds of analytical errors recur, or how to verify the output against the underlying data. The French Azure reference [17] presents an architectural diagram in which the LLM is a box with arrows pointing into and out of it. The interior of that box — the part that determines whether the generated document will be useful or hallucinated — is precisely what the page declines to describe.

[26] Summarize a Report With Copilot - Power BI

[17] Générer des documents à partir de vos données

This is not an oversight. It is the design. Transparency about model behavior would expose the gap between the marketing verb ("Copilot summarizes") and the empirical verb ("Copilot produces a plausible-sounding text that may or may not summarize"). The OpenAI help center collection [9] follows the same pattern: extensive guidance on how to use the product, sparse guidance on when not to trust it. The companion page [10] describes the enterprise education tier in terms of administrative privileges and data-handling commitments, but the underlying question — what the model actually does correctly, and at what rate — receives no treatment. The page is a contract. It is not a report.

[9] ChatGPT - OpenAI Help Center

[10] ChatGPT Edu at OpenAI - OpenAI Help Center

Notice what this absence is doing. A user who relies on the documentation to evaluate the tool is structurally prevented from doing so. The information the user needs to make an adoption decision is held back; the information the vendor needs to close the sale is foregrounded. This asymmetry is what Shoshana Zuboff, in [27], identifies as a defining move of the platform era: the unilateral declaration that certain facts about the product are simply none of the user's business. The documentation page is the polite face of that declaration.

[27] The Age of Surveillance Capitalism

The Free Tier as a Funnel, Not a Gift

The clearest evidence that the tool is a node in an extraction network rather than a stand-alone artifact is the structure of its pricing. GitHub's documentation on [22], and its Spanish-language twin [21], lay out a tiered architecture in which the free entry point exists specifically to seed dependency. The deeper page [23] makes the metering logic explicit: every interaction is a "request," requests are budgeted, budgets are billed, and the unit economics of your workflow eventually

[22] Plans for GitHub Copilot

[21] Planes para GitHub Copilot - Documentación de GitHub

[23] Solicitudes en GitHub Copilot - Documentación de GitHub

convert into a line item on someone's invoice.

The "free for students" and "free for teachers" framing is the most efficient version of this funnel. [2] and [1] are not philanthropy. They are user acquisition with the cost of acquisition zeroed out for the population most likely to carry the tool into a future paying job. Microsoft's training module [14] and its one-hour variant [13] operate on the same logic, dressed in the cheerful uplift of professional development. The training course is the tutorial; the tutorial is the onboarding; the onboarding is the lock-in.

[2] Access GitHub Copilot for free as a student
[1] Access Copilot Pro for free as a teacher or open source maintainer
[14] Embark on your AI journey with free AI tools from Microsoft Education
[13] Embark on your AI journey with free AI tools from Microsoft (1-hour)

When this funnel meets an institutional customer, the lock-in becomes consequential at scale. The reporting around the [8] arrangement is instructive precisely because the deal arrived as a fait accompli: an entire public university system folded into a single vendor's product without the deliberation that the scope of the commitment would seem to require. Once the workflows, the credentials, the integrations, and the muscle memory are in place, the cost of exit is no longer a contractual question; it is an organizational one. The vendor knows this. The pricing page exists because the vendor knows this.

[8] Cal State struck a surprise deal with OpenAI — but some students and faculty are pushing back

The Stack Beneath the Stack

What the tool page cannot bring itself to say is that the chatbot you are typing into is the tip of a multi-tiered industrial system. Crawford's [27] is direct about this: "Exploitative forms of work exist at all stages of the AI pipeline, from the mining sector, where resources are extracted and transported to create the core infrastructure of AI systems, to the software side, where distributed workforces" — annotators, moderators, red-teamers — produce the labeled examples and the safety filters that make the consumer product feel safe to consume. The cobalt in the GPU's circuitry, the water in the data center's cooling loop, the contracted moderators reading through the worst of the internet on piece-rate wages: none of this appears on the [15] landing page, and there is no version of that page on which it ever will.

[27] The Atlas of AI

[15] GitHub Copilot documentation - GitHub Docs

This matters for evaluation, not just for ethics. A tool whose marginal cost of operation is borne by an externalized labor force and an externalized energy market behaves differently from a tool that internalizes those costs. Its price signal is unreliable; its sustainability claims are unmeasurable from the outside; its dependency profile is opaque to the customer who is asked to bet a workflow on it. When the LinkedIn-circulated analysis [19] claims that AI assistants are now producing a third of new code at major firms, the question a careful adopter should ask is not "is the productivity real" but "what is the

[19] L'IA écrit déjà près d'un tiers du nouveau code : la productivité...

productivity priced against, and what happens to that price when the upstream subsidies — cheap labor, cheap compute, cheap capital — normalize?"

The same question applies to the open-source side of the stack, which is not as open as its rhetoric suggests. The Hugging Face documentation for [25], the [24] optimization tutorial, and [28] describe a model whose weights you can download, but whose training was conducted on hardware no individual user could afford, on data scraped at a scale no individual could replicate, by a company whose subsequent commercial pivots were not contingent on community input. "Open" in this context means inspectable, not democratic. The weights are on the table; the means of producing the weights are not.

[25] Stable Diffusion 2 · Hugging Face
[24] Stable Diffusion - Hugging Face
[28] The Stable Diffusion Guide - Hugging Face

Centralization Disguised as Convenience

Once you read the documentation pages as artifacts of a supply chain rather than descriptions of products, the consolidation move becomes visible. The integration plays — Copilot into Office, Gemini into Workspace, ChatGPT into the enterprise via [10] — are not feature additions. They are perimeter expansions. Each integration narrows the set of workflows that can be performed without traversing the vendor's infrastructure, and each narrowing increases the switching cost of the next decision.

[10] ChatGPT Edu at OpenAI - OpenAI Help Center

Google's [11] developer documentation shows the same architectural ambition from a different vendor's angle: the generative-AI application is presented as a thing you build on top of Google's stack, with Google's models, with Google's billing, with Google's identity layer. The "build" is real; the substrate is Google's. The page is not lying about what it offers. It is simply not in the business of pointing out that the offer is asymmetric.

[11] Desarrolla una aplicación de IA generativa | Generative AI | Google ...

When the Knight Institute's analysis [3] argues that AI should be understood as a social rather than purely technical phenomenon, this is the substance of the argument: a tool whose deployment reshapes the institutional relationships of the field it enters is not a tool in the way that a hammer is a tool. It is closer to what Marshall McLuhan, in [27], called a medium — an environment that restructures the perception and the practice of whoever operates within it. The hammer does not change the carpenter. Copilot, embedded in the workflow, changes the workflow, and through the workflow, the worker, and through the worker, the standards by which the work is judged.

[3] AI as Social Technology

[27] Understanding Media

This is why the [7] finding — that ChatGPT-generated referral traffic to French news sites is concentrating dramatically on Le Monde

[7] Audience générée par ChatGPT : « Le Monde » écrase la concurrence

at the expense of competitors — should worry anyone who cares about the structure of the information ecosystem. The chatbot is not neutral plumbing. It is a routing layer whose decisions about which sources to surface have measurable consequences for which sources survive. The tool centralizes, and the centralization is not a side effect; it is the business model.

Failure as a Feature, Not a Bug

A careful adopter, reading vendor documentation, has to develop a second-order reading skill: noticing what the documentation cannot afford to say. The product pages do not catalogue failure modes because cataloguing failure modes would itself be a failure mode. The failures are documented elsewhere, by users and journalists and researchers, and they are systematic enough that they should be treated as features of the technology, not as anomalies awaiting a patch.

Consider the security surface. The analysis at [5] describes a class of attacks — prompt injection, role-play exploitation, instruction smuggling — that vendor documentation handles only obliquely, because the attacks exploit the very capability the product is sold on. A model trained to follow instructions cannot be fully hardened against following the wrong instructions; the alignment is statistical, not categorical. The vendor's safety filter is a probability distribution that an attacker can reshape. The product page cannot say this without undermining the product page.

[5] AI Jailbreak Prompts: How They Work, Why They Work, and How to Stop Them

Consider performance under adversarial conditions. The French analysis [18] takes a class of generative tools through one of the more demanding standardized examinations available in the French academic system, and the results are uneven in ways that the marketing copy systematically suppresses: the model handles familiar forms well and unfamiliar ones badly, with the failure modes concentrated precisely where a naive user is least equipped to detect them. The tool is not lying. It is producing a plausible-sounding output that resembles the right answer in surface features and diverges from it in substance. This is not a bug to be fixed in the next release. It is a structural property of the technology.

[18] L'IA générative face au concours d'entrée à l'École normale supérieure

Consider the integrity surface. The Quebec public broadcaster's reporting in [29] — one in three students using AI in ways their institution would call rule-breaking — describes not a moral failure of individual users but the predictable consequence of a tool deployed without operational guardrails into an environment whose rules the tool's vendors did not write and do not enforce. Google's response,

[29] Un étudiant sur 3 transgresse les règles à l'aide de l'IA

documented at [4], is itself instructive: the platform that profits from frictionless generation is now selling, or being asked to sell, the friction. The same vendor sells the disease and the cure.

[4] AI detection in Google Classroom

Consider dependency. The Spanish reporting on [12] — measurable problematic dependency emerging among a non-trivial fraction of the user population — describes an externality the vendor has no incentive to surface and the user has limited capacity to self-diagnose. The product page does not mention it. It cannot.

[12] El 4,7% de los estudiantes ya está en riesgo de tener una dependencia problemática de la IA según un estudio de UNIR

What unites these failure modes is that each is well-documented in third-party sources and absent or attenuated in vendor sources. This is not a coincidence. The asymmetry is the design.

The Open Counter-Stack and Its Disadvantage

The honest argument for centralized commercial AI is that it works, today, in production, at scale, and that the alternatives do not. This is the argument the dominant vendors do not need to make explicitly because the deployment patterns make it for them. But it is worth being precise about what "the alternatives do not" means, because the openness gradient is wider than the marketing of either side suggests.

The Hugging Face documentation set — [25], [28], the optimization-focused [24] — represents a genuinely different relationship between user and tool. The weights are inspectable. The pipeline is modular. The user can, in principle, replace any component with a different one, including the model itself. This is the openness the dominant vendors do not offer and structurally cannot offer without undermining their pricing.

[25] Stable Diffusion 2 · Hugging Face

[28] The Stable Diffusion Guide - Hugging Face

[24] Stable Diffusion - Hugging Face

But "in principle" is doing a great deal of work in that sentence. The open stack requires technical capacity the closed stack does not require. It requires hardware budgets the closed stack rents out by the second. It requires a willingness to maintain infrastructure the closed stack maintains for you, badly and opaquely, but maintains. The result is that the openness gradient and the convenience gradient run in opposite directions, and the convenience gradient is what determines adoption at scale. The dominant vendor's product is not better. It is closer.

This is a familiar pattern. Ruha Benjamin, in [27], describes the way that technological systems encode and entrench prior distributions of power even when their explicit rhetoric is neutral or emancipatory. The open-source AI ecosystem is not exempt from this dynamic. The capacity to use it is unequally distributed; the capacity to evaluate it

[27] Race After Technology

is unequally distributed; the capacity to contribute to it is unequally distributed. The result is a counter-stack that exists, that is technically excellent, and that loses anyway, because the contest is not about technical merit. It is about which infrastructure has the capital to absorb the cost of being the default.

Noam Chomsky's [27] argued that the appearance of choice in a media system can coexist with strong structural narrowing, because the costs of operating at scale select for a small number of survivors and the small number of survivors set the terms within which "choice" is exercised. The AI tools market is reproducing this pattern with remarkable fidelity. The choices are real. The conditions under which the choices are made are not chosen.

[27] Manufacturing Consent

What a Careful Adopter Should Actually Ask

Read against the grain, the vendor documentation does provide useful information — just not the information it foregrounds. The licensing schedules in [20] tell you what data the vendor wants access to and on what terms. The metering documentation at [23] tells you what unit the vendor expects to bill against, which tells you what the vendor expects your dependency curve to look like. The training-path landing page [16] tells you what population the vendor is trying to socialize into its workflows, and how. None of this is hidden. It is simply not what the pages claim to be about.

[20] Microsoft 365 Copilot - Service Descriptions | Microsoft Learn

[23] Solicitudes en GitHub Copilot - Documentación de GitHub

[16] GitHub Copilot Fundamentals Part 1 of 2 - Training | Microsoft Learn

A careful adopter, then, is reading every documentation page twice — once for the feature claims, and once for the supply-chain commitments those feature claims encode. The questions worth asking, in order, are: What does this tool actually do, at what error rate, on what classes of input? What infrastructure does it require my organization to depend on, and what does exit from that dependency cost? Whose labor is in the loop, and under what conditions? What externalities — energy, water, attention, dependency — am I implicitly endorsing by adopting it? And what is the second-order effect of my adoption on the structure of the field I work in?

These questions are not in the documentation. They cannot be in the documentation. The documentation is a product page, and a product page exists to convert. The questions are in the seam between the product page and the world, and the world's evidence — the [5] red-team reports, the [7] traffic data, the [8] procurement reporting, the [3] framing — is the corrective the documentation cannot afford to provide.

[5] AI Jailbreak Prompts: How They Work, Why They Work, and How to Stop Them

[7] Audience générée par ChatGPT : « Le Monde » écrase la concurrence

[8] Cal State struck a surprise deal with OpenAI — but some students and faculty are pushing back

[3] AI as Social Technology

The reframing that matters, the one this week's evidence keeps

pointing toward, is that the tool is not an autonomous agent and not a neutral utility. It is a position in a network. Adopting it is taking a position in that network. The black box that runs on sweat — on cobalt, on annotation, on cooling water, on capital, on the cognitive labor of millions of users training the next model by using the current one — is not made less of a network by being painted to look like a product. It is made more obscure, which is a different thing, and which is, as Crawford's [27] suggests when it names the rhetorical move of abstraction, exactly what the painting is for. Seeing the network is the precondition for any honest evaluation of any tool in it. The documentation will not see it for you. Nothing about the structure of the documentation would let it.

[27] The Atlas of AI

References

1. Access Copilot Pro for free as a teacher or open source maintainer
2. Access GitHub Copilot for free as a student
3. AI as Social Technology
4. AI detection in Google Classroom
5. AI Jailbreak Prompts: How They Work, Why They Work, and How to Stop Them
6. Application card: Microsoft 365 Copilot
7. Audience générée par ChatGPT : « Le Monde » écrase la concurrence
8. Cal State struck a surprise deal with OpenAI — but some students and faculty are pushing back
9. ChatGPT - OpenAI Help Center
10. ChatGPT Edu at OpenAI - OpenAI Help Center
11. Desarrolla una aplicación de IA generativa | Generative AI | Google ...
12. El 4,7% de los estudiantes ya está en riesgo de tener una dependencia problemática de la IA según un estudio de UNIR
13. Embark on your AI journey with free AI tools from Microsoft (1-hour)
14. Embark on your AI journey with free AI tools from Microsoft Education

15. GitHub Copilot documentation - GitHub Docs

16. GitHub Copilot Fundamentals Part 1 of 2 - Training | Microsoft Learn

17. Générer des documents à partir de vos données

18. L'IA générative face au concours d'entrée à l'École normale supérieure

19. L'IA écrit déjà près d'un tiers du nouveau code : la productivité...

20. Microsoft 365 Copilot - Service Descriptions | Microsoft Learn

21. Planes para GitHub Copilot - Documentación de GitHub

22. Plans for GitHub Copilot

23. Solicitudes en GitHub Copilot - Documentación de GitHub

24. Stable Diffusion - Hugging Face

25. Stable Diffusion 2 · Hugging Face

26. Summarize a Report With Copilot - Power BI

27. The Atlas of AI

28. The Stable Diffusion Guide - Hugging Face

29. Un étudiant sur 3 transgresse les règles à l'aide de l'IA

30. ¿Qué es Microsoft 365 Copilot? | Microsoft Learn