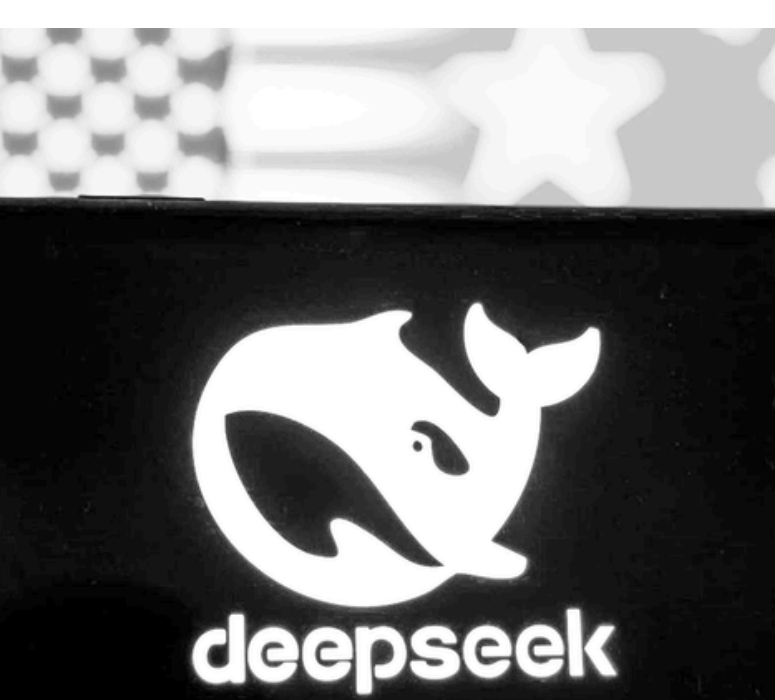


THE HUMAN COST OF DEEPSEEK

HYPE THE TECHNOLOGY, HIDE THE WORKERS



BY ANTONIO A. CASILLI, THOMAS LE BONNIEC,
JULIAN POSADA

China's AI sensation DeepSeek claims to match ChatGPT's capabilities for just 1% of the cost and a fraction of its energy consumption, marketing itself as an open-source alternative to US tech giants. Debates that focus on its technical prowess overlook a crucial factor in its success: government-subsidized data labor. Recent Chinese policies have aimed at creating sprawling data-annotation hubs in 'tier 3' cities, offering tax breaks and financial incentives to companies to sustain a vast workforce of low-wage data labelers. DeepSeek portrays these workers as expert researchers—even suggesting the CEO himself labels data—and claims a team of just 32 annotators. However, this version of events clashes with documented evidence and casts doubt on the startup's marketing narrative and technological claims. Similar to how ChatGPT's ambitious AGI prophecies were undermined by revelations of widespread human annotation networks, DeepSeek's miraculous cost and efficiency metrics may conceal less comfortable realities yet to be fully appreciated.

To cite this document:

Casilli, A. A., Le Bonniec, T., Posada, J. 2025. "The Human Cost of DeepSeek. Hype the technology, hide the workers," DiPLab Policy Memo, vol. 1, n. 1.



“Are we seeing the same old story, where Silicon Valley's perennial appetite for disruption outweighs its capacity for doubt?”

DeepSeek: The Next Theranos?

In January 2025, the Chinese AI startup DeepSeek made waves by announcing an open source chatbot that matches ChatGPT's capabilities, developed for just \$5.6 million—a mere 1% of the \$500 million price tag attached to American rival OpenAI's training costs.

DeepSeek presents its breakthrough as a triumph of optimization, achieving top performance with simpler, cheaper processors rather than relying on the latest high-end GPU. The company claims its proprietary algorithms shorten the learning curve for its AI systems, dramatically reducing the time and resources needed for training.

What's more, the company claims it uses refined data curation methods that need only minimal guidance from a small team of expert annotators.

While US big tech harbors reservations about DeepSeek's Chinese origins, it remains fascinated by the startup's technical narrative. The situation echoes previous too-good-to-be-true stories, like Theranos, whose revolutionary promises ultimately collapsed under proper scrutiny. Are we seeing the same old story, where Silicon Valley's perennial appetite for disruption outweighs its capacity for doubt?

All DeepSeek's assertions are difficult to verify. While its chatbot is technically open-source in the sense that users can create independent versions of the model, the company has not released its training data or critical portions of the code. Its bold declarations of cost and energy efficiency appear increasingly questionable, too. Recent investigations suggest costs have approached \$500 million, virtually on par with OpenAI. Most importantly, its discreet yet high-quality annotation strategy appears to contradict the evidence we will present in this memo.

Tech Bros Are Missing the Real Story

Within the AI industry, the rise to global prominence of DeepSeek has reignited old tensions and sparked new ones.

Chinese observers see the chatbot's spike in popularity as a watershed moment for their domestic AI industry. [China Daily celebrates](#) the "historical moment when Chinese LLMs surpass ChatGPT."

The reactions from American experts largely reflect anxieties about economic and national security implications. According to Trump's 'crypto czar', the US should spare no effort to defeat this threat and ["ensure American AI dominance."](#)

Commentators interpret DeepSeek's rise as a rebuke to OpenAI's resource-intensive approach, with its massive infrastructure requirements and environmental footprint. In response to this criticism, Sam Altman accuses the Chinese company of creating nothing short of "counterfeit AI" through what he describes as intellectual property theft, [appropriating OpenAI's technology via model distillation.](#)

The controversy surrounding distillation resurrects the tired xenophobic cliché that Asian innovation is merely Western ingenuity sold at discount. Importantly, while experts dissect technical specifications and national strategic implications, they overlook a fundamental truth: AI models don't train themselves. DeepSeek's foundation rests squarely on human labor. This blind spot in AI discourse reveals how technical metrics and geopolitical interests consistently overshadow the human element.

“The controversy surrounding distillation resurrects the tired xenophobic cliché that Asian innovation is merely Western ingenuity sold at discount.”

“It’s the Labor, Stupid!”

AI models like DeepSeek depend not just on algorithms but on a vast workforce that processes raw data. The spotlight often falls on a small elite of high-paid engineers, but they are merely the tip of the iceberg. DeepSeek proudly recruits recent graduates from top Chinese academic institutions like [Peking University and Tsinghua University](#), favoring the [“young, and high-potential” professionals—typically born around 1998](#) with no more than five years of experience.

This elite conceals the huge workforce made of people who transcribe text, label images, classify videos, and annotate data, playing a crucial role in shaping AI, often for meager wages.

DeepSeek's latest model is completely reliant on these data workers, despite the company's emphasis on "reinforcement learning". Unlike supervised learning, which follows an *ex ante* approach—training models with pre-tagged data before deployment—reinforcement learning works *ex post*, correcting outputs through rewards and penalties. In theory, this reduces the need for data workers.

In practice, however, reinforcement learning neither revolutionizes AI training nor significantly cuts labor costs. While it may lessen dependence on pre-labeled datasets, it still requires human oversight, as annotators must provide feedback on model outputs. Even when users casually upvote or downvote a chatbot's response, they unknowingly contribute to this process—unpaid and uncredited.

DeepSeek may promise labor-saving automation, but, like all AI models, it runs on human toil.

What We Know about AI Labor

The scale of data work is staggering and difficult to measure, as it relies on both big online platforms and small outsourcing firms. At the [Digital Platform Labor Research Group \(DiPLab\)](#), we have been at the forefront of the study of the composition of this workforce across several regions both in the North and in the global South. Recent estimates, [based partly on our methodology](#), reach impressive figures. Oxford researchers estimate around 163 million registered platform workers around the world, with approximately [16 million engaged in data work](#). A World Bank report provides an even broader range, estimating between [154 and 435 million "online gig workers"](#). If accurate, these figures would represent between 4.4% and 12.5% of the global labor force, encompassing both freelancers and data workers.

These numbers may seem overinflated, but they read as conservative when compared with platforms' own user statistics. For over a decade, [platforms have reported tens of millions of users in both the West and China, with some Chinese platforms alone claiming 10 to 15 million each](#)—not to mention the vast in-house data annotation teams of tech giants like Baidu, Tencent, and Huawei.

[DiPLab's surveys](#) have revealed a global workforce that is precarious, exploited, and wrongly dismissed as unskilled. In "peripheral" regions, workers serve the interests of technological "centers," perpetuating historical North-South inequalities.

From Madagascar workers posing as algorithmic surveillance systems to Kenyan moderators policing Meta's filters and South American data workers training robot vacuum cleaners, the cases we documented are countless. Our 2023 study on Brazilian data workers uncovered 54 platforms, with a workforce predominantly young (18-35), [female \(60%\), and systematically underpaid](#). The study we conducted in 2024 across 9 European countries similarly revealed precarious working conditions, with [women and migrants significantly represented in these digital labor markets](#).

DiPLab associate researcher Milagros Miceli has recently initiated a [Data Workers Inquiry](#) which features labelers and moderators from Syria, Kenya, Venezuela, Brazil, and Germany. Her work also illustrates how prevalent data annotation is for precarious workers, such as refugees working in foreign countries. With Julian Posada, she has conceptualized this phenomenon as a ["data production dispositif"](#) — a complex system where business process outsourcing and crowdsourcing platforms create conditions of economic dependency. Workers are often alienated, their agency limited by carefully designed interfaces and performance metrics that normalize specific modes of data interpretation. For the AI global production chain, this means that US and EU clients outsource data annotation tasks to a global workforce that must compete with other countries for cheaper and less regulated labor.

Data annotation is an AI-driven global labor market. While estimates suggest enormous numbers of workers, the reality remains fragmented, with significant regions—particularly China—remaining in a blind spot. Unlike the detailed documentation of workers in other regions, we have encountered significant barriers in accessing information about data annotators for Chinese companies. The rare cases we documented include young women in Cairo working for Chinese company Tencent on facial recognition systems, and European workers transcribing [vocal assistant conversations for a Chinese platform](#) verifying data for a US tech giant.



“Data annotation is an AI-driven global labor market (...) with significant regions —particularly China—remaining in a blind spot.”



“While Western tech giants typically outsource data work to countries in the southern hemisphere, Chinese AI companies primarily leverage domestic low-income urban regions.”

Who are China's Data Workers?

While Western tech giants typically outsource data work to countries in the southern hemisphere, Chinese AI companies primarily leverage domestic low-income urban regions. Big tech in rich cities like Shanghai, Shenzhen, and Hangzhou, where DeepSeek is located, systematically recruit workers from populous urban areas with low economic growth. Companies like DataTang, DreamDate, and DataOceanAI have been setting up massive data annotation operations in these “Tier-3” cities, where they hire workers willing to accept minimal wages for repetitive work.


As early as 2018, the New York Times reported that China's AI ambitions were fundamentally “driven by cheap labor”, with data factories emerging in remote areas where both workforce and infrastructure costs remain minimal. The sector is expanding so quickly that in 2024 China anticipates 30 million unfilled positions.

This is known as the Qiantian Houchang model—literally “front shop, back factory”, where teams of data annotators are assembled to complete assignments set out by China’s tech giants, like Alibaba, Baidu, and others.


In a landmark 2023 study, sociologists Wu Tong-yu and Xia Bing-qing, illustrate the structure of this marketplace. Specialized annotation companies serve as intermediaries, connecting AI firms with pools of annotators and managing worker training. These operations often receive local government support, with some emerging from poverty reduction programs.

However, competition among these annotation bases has triggered a wage spiral downward. Julie Chen's research, based on three years of fieldwork and dozens of interviews, describes a “Four S's” framework: Swindle, Swiping, Scold, and Study. Workers are *swindled* through scams that lure them into paying to participate in data annotation projects. Their earnings are then *swiped* through rampant wage theft, as clients can arbitrarily reject payments. Workers are frequently *scolded* for being too slow or failing to meet quality requirements. Throughout this process, they must continually *study* — applying their skills and knowledge to manage the constant learning demands across different AI training projects.

This mirrors the Western model. Companies like Meta and OpenAI face criticism for underpaying data workers in low-income countries. But AI labor exploitation isn’t just a matter of geography, with Western companies relying on foreign workers and Chinese firms extracting labor from their domestic workforce. The key difference is political: While the U.S. government considers supporting AI by potentially declaring “national energy emergencies” to power massive data centers, the Chinese government fully recognizes the centrality of data labor in driving its AI sector.



2019 estimation of China AI-related data services industry (RMB millions).
Source: iResearch



2025 outlook of China data annotation tools market. (USD millions).
Source: Grand View Research

“Data annotation now enjoys official support from the Chinese government, driving down the cost of data labor—an advantage that likely played a role in DeepSeek’s rapid ascent.”

China Big Data Annotation Push

In the last few years, China's data annotation sector has undergone a remarkable transformation, evolving from a market supported by local administrations to the object of a comprehensive national strategy for industry growth. The central government's strategic intervention, detailed in policy documents issued by the National Data Administration and the National Development and Reform Commission, has fundamentally reshaped the landscape of data labeling.

Between 2023 and 2024, a wave of policy directives and implementation guidelines has subsidized the sector while also establishing the government as both a client and a data source. By promoting partnerships between state agencies and tech companies, these measures have also facilitated the annotation and training of data for large language models tailored for government applications.

Data annotation now enjoys official support from the Chinese government, driving down the cost of data labor—an advantage that likely played a role in DeepSeek’s rapid ascent.

The strategic objective is ambitious: achieving an average annual growth rate exceeding 20% in the data annotation market by 2027. This approach positions large language model annotation, exemplified by projects like DeepSeek, as a core technological focus.

The goal extends beyond mere growth, aiming to establish a comprehensive industrial ecosystem with dedicated annotation hubs. To realize this vision, the government has deployed a multi-faceted approach: investments, corporate tax incentives, and strategic procurement. As a result, human annotation is not just expanding but appears set to become a lasting pillar of China’s economy, driven by the establishment of dedicated centers and districts. Regions are encouraged to invest in data products, with notable initiatives such as the National Data Annotation Base planned for the island of Hainan, a small yet strategically significant free-trade zone.

“DeepSeek downplays annotation as mundane while simultaneously presenting it as a noble endeavor – even the founder, Liang Wenfeng, occasionally participates in labeling data.”

The Hangzhou Thirty-Two

Despite mounting public initiatives to bring data annotation to the forefront and recognize its strategic value for AI companies, DeepSeek’s stance on data labeling wavers between secrecy and inconsistency. The AI company insists on branding its models as frugal – not just in computing power but also in their supposed minimal dependence on human data work. Yet, for all its claims of openness, the company remains conspicuously secretive about the scale and composition of its training database – and thus about the size of its pool of annotators.

Behind closed doors, however, DeepSeek’s executives candidly admit the truth: despite public claims of advanced reinforcement learning techniques, their AI relies heavily on data annotation, with supervised learning – involving numerous human labelers – underpinning their models’ performance. Every sophisticated reasoning task, at its core, still depends on human-guided learning. A technical report published by DeepSeek in January 2025, reveals how even their most recent model heavily relies on data filtered and enriched by human labor across domains like writing, factual question-answering, and self-cognition.

The startup’s narrative oscillates between trivializing and simultaneously ennobling this data work. DeepSeek downplays annotation as mundane while simultaneously presenting it as a scholarly activity. It seems that even the founder, Liang Wenfeng, occasionally participates in labeling data.

The same pattern of both belittling and elevating data work applies to those annotators who have the rare distinction of being named in the company’s scientific papers. Our analysis of available papers identifies only 31 to 32 individuals who are explicitly credited as “data annotators.” They are portrayed as an elite cadre of expert researchers, their status underscored by their co-authorship of the publications. This suggests they may be team coordinators rather than back-office data workers. The ArXiv papers indicate that their contributions are limited to specific experimental database segments rather than the full chatbot system.

Nothing is certain, but based on previous well-documented instances elsewhere, there is the possibility that these individuals coordinate expansive annotation teams in lower-tier cities, managing a workforce far larger than DeepSeek’s narrative suggests.

From this point of view, DeepSeek’s approach to work and workers mirrors that of its Western counterparts. Their obfuscation of data labor practices is standard in the tech industry, with the company not only following the conventional playbook of selective disclosure but also drawing from the same pool of global resources and methods to generate the data that annotators work on. One example is Common Crawl, the sprawling internet document repository used both for the DeepSeekMath model and... by its arch-enemy OpenAI in the development of ChatGPT.

The annotation workforce behind their technology remains shrouded in mystery, casting doubt on how such a seemingly modest team could orchestrate a model encompassing trillions of parameters. This strategy epitomizes a pervasive industry trend: grand proclamations of technological independence masking the extensive, strategically concealed human infrastructure that powers contemporary AI systems.



“Cherchez les Data Workers!”

Despite its bold claims of disrupting the AI market with a revolutionary approach, DeepSeek essentially functions as a traditional tech company. It shares core characteristics with the very competitors it seeks to challenge.

One of these competitors, OpenAI, provides an illuminating precedent. Despite positioning itself as the standard-bearer of Artificial General Intelligence, two months after the launch of ChatGPT in November 2022 a [Time magazine exposé](#) revealed that the model was reliant on hundreds of data annotators in Africa, paid less than \$2 an hour. DeepSeek is just the latest example of the same old script: hype the technology, hide the workers.

Similarly, DeepSeek’s marketing narratives appear to obscure the massive data requirements underpinning its model’s functionality.

Nevertheless, what the Chinese startup does with its data becomes the focal point of international regulatory scrutiny, including from [Belgian](#), [French](#), and [Italian data protection authorities](#). The gap between the company’s public image and its operational reliance on extensive human data labor reveals an organization whose revolutionary promises may not align with its actual technological achievements.

The intelligence in “artificial intelligence” stems largely from human effort—countless workers who screen, clean, classify, annotate, label, and “reinforce” these systems. Their concealed labor forms the foundation of what companies can present as products capable of autonomous reasoning.

Waiting for the Other Shoe to Drop

Alas, even though academics and journalists continue to expose uncomfortable truths, their revelations rarely shift public opinion or make a dent in slick marketing campaigns. Yet they reveal something crucial: The tech race between China and the US requires us to look beyond flashy tech prophecies and face hard truths. DeepSeek’s prodigious performance isn’t about engineering breakthroughs — or even intellectual property theft, as Sam Altman insinuates — but about a hidden workforce earning poverty wages.

Given this reality, should we buy DeepSeek’s advertisement? Can we seriously believe they built their miracle chatbot with just 32 annotators, even as they acknowledge their need for people to label data and supervise algorithms’ training? Even if, thanks to Chinese policies, they have been having access to that vast pool of government-subsidized labor?

Like so many AI breakthroughs before, this feels like another case of waiting for the other shoe to drop — for the glossy marketing to give way to uncomfortable realities.

“Like so many AI breakthroughs before, this feels like another case of waiting for the other shoe to drop—for the glossy marketing to give way to uncomfortable realities.”

DiPLab Policy Recommendations

Decent Work and Living Conditions for Data Annotators

China's approach includes government subsidies for data annotation, recognizing its crucial role in AI development. While this policy benefits annotation companies, its impact on workers' welfare remains uncertain.

Data workers worldwide are mobilizing to defend their rights amid challenging work conditions. Content moderators and AI trainers in Kenya, Ireland, and the US have spoken out against practices imposed by their employers under pressure from tech giants like Google, OpenAI, Meta, and Amazon, leading to lawsuits and damaged corporate reputations.

Importantly, the Fairwork initiative has established core standards for ethical data work (termed "cloudwork") through five key principles: Fair Pay, Fair Conditions, Fair Contracts, Fair Management, and Fair Representation. Using this framework, they evaluate work platforms, including those employing annotators. Their comprehensive 2023 assessment covered 270 platforms across 29 countries, with data work platforms performing poorly, achieving an average score of just 2.4 out of 10.

DiPLab advocates for fair treatment of data workers not only as an ethical imperative but also because their expertise drives technological progress. Creating a labor market free from exploitation, with appropriate working conditions and compensation, can also yield better data annotation, and ultimately improve AI quality.

Ensuring Personal Data Protection and Privacy of Both Users and Workers

European privacy regulators have been responding to the rise of AI chatbots at significantly different speeds. While OpenAI's ChatGPT took months to attract scrutiny, DeepSeek faced immediate investigation over suspicions of storing user data in China.

This swift action reflects more than potential geopolitical tensions between China and Western US-aligned countries. Arguably, it signals tech watchdog's evolution from supposedly passive observers of technological advancement to active regulators, equipped with new international and domestic frameworks to oversee AI development.

One of these frameworks, the European General Data Protection Regulation, provides guidelines limiting personal data collection. Although its enforcement remains inadequate, it is a powerful tool to minimize privacy violations across various domains, from apps monetizing intimate data to voice assistants enabling targeted advertising and intellectual property infringement. Such a regulation should extend particular attention to workers who serve dual roles as data providers and annotators.

DiPLab strongly advocates for privacy authorities to deepen their understanding of the AI challenge by embracing this latest technological shift as an opportunity to evolve their role. Rather than solely overseeing AI consumers and users, they should also position themselves as defenders of AI workers' rights.

AI Act Implementation: Data Protection and Rights Impact Assessment

The European AI Act, while groundbreaking, has been weakened by tech industry lobbying, conservative governments' interference, and law enforcement pressure, particularly in areas of oversight and fundamental rights protections for general-purpose AI systems.

Despite these limitations, the AI Act contains provisions that can protect data workers' interests. The mandatory Fundamental Rights Impact Assessments (FRIA) for high-risk AI systems may be used to evaluate effects on specific citizens' groups, including workers involved in data annotation and processing. While originally designed for AI system deployment, these assessments can foster transparency around working conditions in AI production.

DiPLab advocates extending Fundamental Rights Impact Assessments (FRIA) and Data Protection Impact Assessment (DPIA) procedures beyond high-risk systems. These assessments should scrutinize data annotators' working conditions and personal data handling throughout the AI lifecycle. For entities like DeepSeek, this would entail transparency about data sources, annotation methods, and user data collection. Moreover, when combined with the Platform Workers Directive, these provisions establish a framework for worker oversight of AI systems in the workplace.

Authors

Antonio A. Casilli, DiPLab co-director; Professor, Institut Polytechnique de Paris.

Thomas Le Bonniec, DiPLab researcher; PhD, Institut Polytechnique de Paris.

Julian Posada, DiPLab associate researcher; Assistant Professor, Yale.



For more information, please contact:

<https://diplab.eu>

contact@diplab.eu