Lesser Gods: Labor in the AI Labyrinth, Part 1

Advance Drafts #3

PHIL A. NEEL

DEC 23, 2025 · PAID

8 3 Share

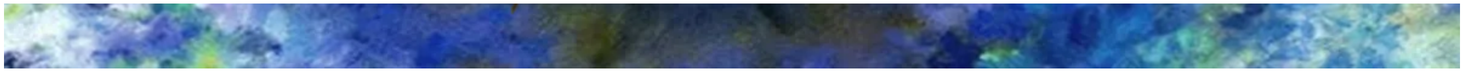

What follows is the first of three parts of an extensive draft on the AI supply chain, based on my own firsthand experience as a worker assembling models within the "AI factory." This was originally intended as a simple, one-off article detailing the boom in employment in these new AI assembly jobs and tracing out the global division of labor through which models are produced. However, given the pace of development within the industry and the sheer amount of absurd misunderstandings about the nature of the technology as well as continual controversies over data centers, energy use, military applications, etc., I soon found myself having to write out various clarifications and add expansions to the original plan. The article quickly got out of hand and now sits at roughly 30,000 words, still growing. The first part, posted below, is roughly 13,000 words.

Rather than cutting it back down to a manageable size, I've instead tentatively decided to instead write it up further into a short book. I'll then create a more polished article version, cut down from the larger content, to put forward the core argument and preview the content of the forthcoming book project. I'll be posting sections of the initial draft here, before it is formatted into its proper book form with a linked narrative. While Hellworld *carried on the project begun in* Hinterland, *which I hope to conclude with a third volume on the questions of partisanship and political organizing broached in my recent pieces for* Ill Will *and (forthcoming)* Heatwave, *this AI book is more of an aside. Think of it like a slim political pamphlet calling attention to a single industry, offering a brief bit of theory on its character, and posing some of the key questions for organizing in and around this emergent sector.*

As always, please note that this is a rough draft. It includes typos, some awkward wording, and many extraneous bits that will eventually be refined or cut out. If citing it, please cite the forthcoming piece and don't quote directly until that final version is released. Also, I'd strongly recommend taking a look at the special episode the excellent podcast Until Everyone is Free *did on the industry in Kenya:*

Introduction

Look at the computer in your hand, on your desk, hidden and humming in the depths of the server room. The sleek aluminum carapace. The glass screen like the jelly of an eye and the pixels smoldering beneath it. Look at the stream of content scrolling on the surface, pouring soundlessly into the abyss of your attention. Look, look, look. Each ripple in this stream is cast up by something moving there in the depths, unseen. Look. It is a machine, is it not? Work it hard and it will heat up like an engine. Tear it open and you will find a collection of small black boxes linked to one another by intricate circuits. Look at them. Were you to somehow unpack these boxes and peer into their microscopic interiors, you would see vast webs of even more intricate circuitry etched in eldritch sigils all the way down to the nanometer. There is doubtless something magical here. Look at what you cannot see. These cutting-edge chip transistors (2 nm) are now smaller than a strand of DNA, a mere 20 atoms in width. Gnashing gears approach the foundation of all substance. Look deeper. Perhaps hidden muscles move even further below, in the darkness beneath the waves of churning matter. Look. Look. Look. What if, behind and beyond these microscopic boxes, there are other boxes, blacker still – hidden from sight by an ideological limit, rather than an optical one? And what if these boxes are filled not with dark, trashing magics, but with people? Look at us.

Like any commodity, the computer is not simply a machine but is instead a laminate artifact composed of interwoven layers of labor and materials flowing through production chains that span society. As a result, you yourself exist in an immediate social relationship with a web of other workers stretched like pulsing veins across the planet. And this web of relationships also implies an expanding circle of ethical imperatives – moral obligations to these other people who make our lives possible, even if we have never met them. But the very nature of these relationships also makes them opaque. They appear to us as a series of monetary transactions where everyone is apparently paid what they are owed and any outstanding ethical obligation seems to have thereby been settled by the mechanics of the market before the final transaction takes place. As a result, the social character of production is effectively disavowed. We are encouraged not to see these workers at all. Only the machines that they have left behind.

As a worker in the digital supply chain, I live in one of these boxes inside your computer. And I am not alone. We are here together: the equipment operators at large-scale mines in Guinea and Australia harvesting the bauxite that will be smelted into the aluminum chassis, the debt-bonded 'artisanal' miners digging cobalt out of the hills of Kolwezi, the rural migrants at massive aluminum smelting facilities in Yunnan and West Kalimantan, the chip designers in San Diego and Silicon Valley, the clean room technicians at fabs in Hsinchu, the workers in semiconductor packaging facilities in Malaysia, the rural migrants manufacturing the glass screens in Hunan or working assembly lines in Xiamen or Bắc Giang, the optical sensor designers in Tokyo and technical staff at the affiliated factories in Kumamoto or outside Bangkok, and of course the myriad laborers on container ships, in ports, warehouses, and last-mile delivery or retail centers that bring the goods into the home.

And these relations go beyond the static materials of metal and silicon that compose the hardware of the computer. Software is equally social. Often mythologized as "immaterial," software is instead a blanket category covering a wide variety of manufactured material systems that subsist in a processual fashion at microscopic scales (i.e. as pattern that can transfer between programmable machines), constructed through the arcane medium of the coding console. Though the peculiar character of these artifacts entails a slightly different manufacturing process and makes their material character difficult to perceive, they are ultimately no less substantial than the glass and metal humming under your hands. And they require similar work to assemble. Software, accompanied by its own backend of office spaces and data centers, is just as much a product of collective labor and the channeling of physical resources as the semiconductor itself, and therefore has its own web of social relationships stretching down its own supply chain similarly divided by geography and by stratifications of technical knowledge which together manifest as different rates of pay: from the high-wage software engineer in Seattle to the middle-tier cloud database worker in Ireland to the low-paid coder in Hyderabad and the lowest-rung taskers in Nairobi. The mythos of "immateriality" simply serves to disguise this deployment of human labor and the affiliated rollout of new physical infrastructures on a massive scale.

Nonetheless, there do exist key differences in the manufacturing of hardware commodities such as the computer versus the software products implanted into them. Among the most important is the fact that software requires a one-time manufacture process so closely integrated with design that the two are often conflated, which then allows the final product to be distributed through its very sale, entailing nothing more than download and installation. Rather than requiring the mass manufacture of an individual unit by assembly workers, we find a (relatively) automated process of replication, after which the product is subject to iterative retooling via updates.[1] These updates provide material for the manufacture of new versions, much as sales and repair data would be used to inform the design of new hardware goods. However, precisely because this distribution process is relatively low-cost and multiple consumers can purchase a single piece of software, the flow of profits within the software firm tends to be structured differently than in manufacturing. Rather than deriving primarily from sales-per-unit (and interest charged on loans offered for these sales, as in the auto industry), revenues also come from a range of extraneous services and rents: the company sells both the software and the service of support, training, and maintenance (a practice established with the very first computer system sales by IBM), the consumer does not purchase the product itself but only subscription-like access to it (as with Microsoft or Adobe), the firm makes its profits through ancillary activities such as advertising or the sale of consumer data (as with Google or Meta), etc. In all of these senses, the digital economy acts as something like a hybrid of traditional manufacturing and agriculture.[2]

Ultimately, it was by seizing on and exaggerating these real differences that the mythology of the early digital age gained a veneer of legitimacy. Today, the rise of "artificial intelligence" carries a similar mystique. Then, as now, the nature of these new digital systems appears to be different, in some fundamental fashion, from what came before – whether the provision of software and services or the manufacture of tangible goods. Upon first glance, AI models even appear capable of replacing the "intellectual" workers who were once idealized in the mythology of the digital economy and whose intricate and deeply human skills were precisely what seemed to mark out the distinction between, on the one hand, a humanity which still contained some spark of the divine, however desecrated, and, on the other, the profaned world of beasts and machines. As with software before it, however, these digital systems ultimately prove to be manufactured goods just like any other, requiring vast supply chains linking together human labor. Similarly, rather than simply "replacing" workers, AI models are instead reshaping the nature of their work.

In the early industrial age, artisans found themselves thrown out of the workshop and into a maze of ever-multiplying machines which, though modelled on their own handiwork, had then taken the very tools from their hands. As these machines spit out standardized components, the intricate, human art of handicraft production was increasingly replaced by routine assembly tasks and the skilled artisan reduced to a deskilled worker churning out cheaper products of lesser quality at far grander scales. Meanwhile, new occupations arose related to the design and maintenance of these machines. Though initially artisanal in character, these new occupations were then subject to a similar process of standardization and mechanization. In much the same way, the "artistic" and "intellectual" professionals of today are being cast into a labyrinth of ever-multiplying models which, trained on their output, then spit back out a standardized slop that can be assembled by less skilled workers into generic products at an accelerated pace and for a fraction of the price. And, as in earlier periods of mechanization, new occupations are now arising alongside this industry, with professionals trained in displaced fields now specializing in the maintenance, and construction of the LLMs themselves. In this way, the workers are always forced to construct the labyrinth that imprisons them.

Part 1

Black Box

Even the use of the term "artificial intelligence" to describe the design, production, maintenance, and use of large language models invokes a certain mysticism. As historian of science Evgeny Morozov notes: "what we call 'artificial intelligence' today is neither artificial nor intelligent."[3] At root, it is nothing more than "a well-run but predictable statistical machine" in which intricate, looped, and iterative pattern-matching algorithms generate a rudimentary and superficial form of speech and logic emulation that resembles, if anything, the way that the steam engine appears to breathe like a living thing.[4] Morozov's point is echoed by philosopher of science Matteo Pasquinelli, who argues that what we call artificial intelligence today is essentially just a misleading colloquialism referring to various forms of "vectorization in multiple dimensions" which allow complex problems to be reduced down to an issue of "mathematical optimization in a vector space" and then iterated, representing "a form of *statistical intelligence*" distinct from the colloquial meaning of the word. Instead of self-awareness or even the application of logic to come to reasonable conclusions, "[t]he characteristic of 'intelligence' that is anthropomorphized in AI systems is essentially the trick of projecting data on a multidimensional space in order to perform operations of clustering, classification, and prediction."[5]

Large language models (LLMs) like those used by ChatGPT and DeepSeek appear intelligent because they perform this trick using an enormous database of human-generated text – broken up into shorter strings (called "tokens") that are more easily placed in this vector space – and provide their output in a similarly text-based form. As computer scientist Jarod Lanier argues:

> It's easy to attribute intelligence to the new systems; they have a flexibility and unpredictability that we don't usually associate with computer technology. But this flexibility arises from simple mathematics. A large language model like GPT-4 contains a cumulative record of how particular words coincide in the vast amounts of text that the program has processed. This gargantuan tabulation causes the system to intrinsically approximate many grammar patterns, along with aspects of what might be called authorial style. When you enter a query consisting of certain words in a certain order, your entry is correlated with what's in the model; the results can come out a little differently each time, because of the complexity of correlating billions of entries.[6]

This variation in tone and style and the ease with which the compilation of this vast trove of information is performed then drape the illusion of a more general "intelligence" over what is, in essence, just a layered series of mathematical encodings, transformations, and decodings, each of which is, on its own, no different than any other statistical procedure.

But the illusion extends equally well to the "artificial" part of artificial intelligence, insofar as it implies that it is the model itself automatically producing its "intelligent" outputs. Again: what at first appears to be an inert machine is, in fact, a complex nexus of social relationships. As Morozov points out, today's models ultimately "draw their strength from the work of real humans: artists, musicians, programmers and writers whose creative and professional output is now appropriated in the name of saving civilisation."[7] In contrast to the relatively mechanical, rules-based programming of 20th century AI systems, contemporary LLMs operate on flexible guidelines derived from a mass of pre-existing material produced by human hands. As Lanier describes it:

> A program like OpenAI's GPT-4, which can write sentences to order, is something like a version of Wikipedia that includes much more data, mashed together using statistics. Programs that create images to order are something like a version of online image search, but with a system for combining the pictures. In both cases, it's people who have written the text and furnished the images. The new programs mash up work done by human minds.[8]

This fact is then obscured by a series of black boxes: first, the black box of the model itself, as it is often unclear exactly how it is drawing out these particular probabilities from that mass of data; and, second, the black box of the interface, in which queries are seemingly launched out into a void from which echoes back an answer.

At a deeper level, the combination of human-made raw material, model architecture, and continual interfacing generates what Lanier calls "a giant ocean of jello – a vast mathematical mixing" that then makes it difficult to even represent the machine beneath. And, without any "definite representation of what the system 'wants,' no label for when it is doing a particular thing," it then becomes difficult to determine what it has, in fact, done, and who might be held accountable. Ultimately, however, there is no technical reason for these black boxes. As Lanier notes: "Big-model A.I. is made of people – and the way to open the black box is to reveal them." More broadly, AI systems should simply not be seen as artificial forms of intelligence at all but instead as "an innovative form of social collaboration."[9] Lanier focuses on the obvious examples of this collaboration: the "work done by human minds" in producing the raw materials fed into the model and the engineering of the model architecture by particular programmers working for particular companies. And this makes sense within his (broadly libertarian) framework, where the goal is "a more honest capitalism" in which "people could get paid for what they create, even when it is filtered and recombined through big models…"[10] Though framed as "data dignity" and often referred to as "data as labor," this approach might be better described as something like "wages for content."

Delve down inside the computer, inside the operating system, inside the browser, inside the AI model behind the chat dialogue and you will find me and many others doing the menial labor necessary to make these models work at all.

And yet even this approach retains an aura of mysticism, insofar as it focuses solely on the obviously "creative," and therefore arguably "proprietary," aspects of modern machine learning models: the people who produced the raw materials – and who could conceivably be paid some sort of residual royalty whenever their painting or blog post is skimmed by an AI system – placed alongside the people who perform the high-level design and continual reinvention of the model itself. The same logic could conceivably be extended to those suffering certain externalities, such as water scarcity from data centers or pollution from the new power plants built to run them. Additional rents charged on the platforms could be issued in the form of an ecological tax, fed back into these communities and justified precisely by inclusion in a community made to bear the costs of a profitable venture. This sounds perfectly nice. But even this approach proves to be yet another black box, focusing on the source materials and high-level design in order to disguise the grim mechanisms through which these models are produced in the first place, which is to say through hidden layers of tedious, dangerous work performed by people who are invisible precisely because they *are* being paid. And it is this black box that I myself occupy. Delve down inside the computer, inside the operating system, inside the browser, inside the AI model behind the chat dialogue and you will find me and many others doing the menial labor necessary to make these models work at all. This is not an exaggeration or playful metaphor. When you shoot that query off into the void and see the scrolling text echo back to you, there are quite literally hundreds of thousands of workers hauling those letters from the slums of Nairobi, across the fiber-optic cables that undergird the world's oceans, and through the impoverished suburbs encircling the world's wealthiest cities to finally arrive at your screen. We live inside your computer – and we hate it here.

The Hidden Abode

Lanier is ultimately correct in understanding large machine learning models as, at root, an "innovative form of social collaboration." But, in a capitalist society, collaboration is only made possible in and through domination. More specifically, the cooperative work necessary to produce any complex good occurs within supply chains controlled by profit-driven firms owned by various strata of social elites (most of whom begin from a hereditary pool of wealth) all inset within a complex, simultaneously cooperative and competitive hierarchy of market power in which downstream "lead firms" exert an inordinate influence and take an exorbitant profit relative to firms further along the chain. The classic example is that of companies like Apple or Nike, which control the design and branding (and thereby the intellectual property and retail power in high-end markets) but outsource all the actual production to large contract manufacturers, who then outsource the supply of components to others, and so on. These contract manufacturers are the companies (and behind them, the investors) that own the physical equipment, buy the materials, and pay the workers actually producing the goods.

The machines and the materials on their own are inert. It is only the workers that put them into action. At root, then, all the profit flowing through these chains is only made possible by the workers within each of the various firms that compose them – from the most menial tasks performed in the poorest reaches of production to the most creative aspects of product or machine design taking place in their upper echelons. And yet the vast majority of these workers have no proprietary right over their product, the entirety of which belongs to the company. They own and sell their physical expenditure alone. Only an elite strata of employees at the very top are given even the slightest hint of ownership, in the form of company stock – and even then, usually only an insignificant sliver. The real owners lie above: the investors who can earn a return without lifting a finger (the ultimate owners) and the executives who oversee the day-to-day functioning of the firm at its highest levels (though they are also usually invested in the firm, their function is largely to actualize this ownership through concrete infrastructures of command and control). The planetary collaboration made possible by these supply chains is therefore presided over by a class of owners whose interests are exerted through both the ambient discipline of the market (the need for the firm to sell enough at the right price to turn a profit, so as not to perish) and the specific discipline of each firm's labor regime, attuned to the requirements of each particular workplace.

Within the upper echelons of the value chain, labor regimes tend to be conciliatory and collaborative, giving special privileges and substantially higher pay to a small number of workers with more elaborate training and education – who are often viewed as a distinct "professional," "high-skill," "creative," "intellectual," etc., stratum due to their credentials – or to managerial workers who enforce the labor regime and hope to climb their way into executive positions, from whence they might eventually win a seat at the table of the ultimate owners. Below, however, labor regimes grow more blunt and more brutal. Though also complex, these lower-order regimes are all ultimately centered around concrete methods of shopfloor discipline designed to extract more output per hour of work, repressive measures meant to prevent worker revolts, and external political interventions designed to relegate any additional social costs (taxes, unemployment payments, pensions, etc.) beyond the bounds of the firm. The outsourcing of production to cheaper locales, mandatory long hours, militaristic governance of assembly lines, and even the use of outright violence are all common methods of labor discipline deployed within the lower rungs of any given supply chain. Internal divisions are equally essential, since each stratum (including "professional" and "managerial" workers) is continually threatened with losing its privileges relative to lower strata.

Despite appearances, AI is no different. Large-scale machine learning models are not simply a form of social collaboration involving user-generated resources and a carefully engineered model architecture, as Lanier's critical picture of the sector would suggest. They are also a globe-spanning labor regime with a supply chain structure that resembles that of any other industry. Put simply: every AI model is the product of a hidden factory. And, as in any other factory, the machines (here mathematical) only move because there are workers pushing the right buttons, feeding the right components in, and reaching in to clear out the gunk when the machine fails. Nonetheless, most accounts of the AI revolution stress the role played by technical innovation, tracing the boom in modern machine learning back to the late 2000s when modern image-identification algorithms first began to be trained on large datasets, driven by GPUs. But, in reality, the genius of modern machine learning arguably has less to do with the high-level design of the models – for which the basic statistical principles have been long established, with much of the foundational programming already worked out decades ago – than with the invention of new ways to exploit digital labor in a more granular fashion and at a larger scale.

Artificial Artificial Intelligence

Though precursors like the mid-20th century "perceptrons" are mentioned as ancestors, the conventional story usually starts with the advent of the ImageNet visual database in 2008. By the turn of the century, the AI industry had not yet emerged from the long winter that followed a boom in model-building in the 1980s. Most researchers were still treating machine learning just like any other form of programming, spending most of their time attempting to refine their algorithms only to find that even the best performed poorly. Toward the middle of the decade, however, Princeton computer scientist Fei-fei Li began to suspect that the field was hindered not so much by the design of the models themselves but instead by the limited data they were being trained on. Her innovation had nothing at all to do with the structure of machine learning models as such. Most of the core models had been introduced in the 1940s and 1950s, were then fully elaborated during the boom in the 1980s, and then slowly refined over subsequent decades. Li's innovation focused instead on the construction of a large and hierarchically organized relational database that would allow the best pattern-recognition models to prove themselves on a pool of validated training images an order of magnitude larger than anything available at the time. Nor was the design of the database itself particularly groundbreaking, as it drew its architecture from the established practices of the WordNet project, started in the 1980s by Princeton psychologist George Miller. Li's innovation was instead the combination of this database architecture with a particular labor regime uniquely capable of producing adequate quantities of data at scale, previously impossible.

The problem that confronted Li was very straightforward. Even with the database design decided on and a clear plan for constructing it, the limiting factor was labor cost: "Li's first idea was to hire undergraduate students for $10 an hour to manually find images and add them to the dataset. But back-of-the-napkin math quickly made Li realize that at the undergrads' rate of collecting images it would take 90 years to complete."[11] Automating the work wasn't feasible either. Having existing algorithms pull images from the internet would then bias future algorithms trained on the database, since its population of images would be skewed toward whatever images existing models were already good at identifying. Meanwhile, federal grants had failed to come through, making it impossible to hire more undergrads to do the work. In most accounts, the fix is presented as serendipitous: “A solution finally surfaced in a chance hallway conversation with a graduate student who asked Li whether she had heard of Amazon Mechanical Turk, a service where hordes of humans sitting at computers around the world would complete small online tasks for pennies.”[12] After discovering the service, Li could suddenly “corral many workers, at a fraction of the cost.”[13] The real genius of ImageNet was therefore its cheapening of high-quality, validated and labelled training data via a brutal labor regime paying workers far less than the minimum wage.

Amazon’s Mechanical Turk service had been publicly launched only a few years earlier, in 2005. Its premise was to outsource work such as data validation or annotation, content moderation, and simple research to gig workers, allowing the results to be integrated directly into the clients’ digital platforms:

> Employers, known as requesters, post batches of what are called Human Intelligence Tasks, or HITs, on Mechanical Turk’s website. A task could be transcribing an invoice, or taking part in a study, or labeling photographs to train an artificial intelligence program … Freelance workers, known informally as turkers, race to grab and do the tasks… Most tasks pay a dime or less, and there is a daily churn of tasks that pay only a penny.[14]

This allows computer programs to integrate outsourced human input in situations where humans could carry out the associated tasks much faster than existing code. Tellingly, Jeff Bezos once referred to the service as "artificial artificial intelligence," since it allowed complex queries to be solved by backroom workers who would then feed the results back into the machine system as if they were automated inputs, acting "essentially, as another part of the software, performing infinitesimal, discrete tasks, often in rapid succession."[15] More formally, the business model came to be called "crowdsourced micro-tasking" or simply "microwork" and, by the early 2010s, it had evolved into a multimillion-dollar industry.

The "crowdsourced" aspect is simply a digitized form of outsourcing, while the "micro-tasking" dimension is a conventional piece-rate. In fact, rather than being a cutting-edge form of labor deployment, "crowdsourced micro-tasking" is simply a modern iteration of one of the oldest forms of manufacturing: the "putting-out" system that dominated early industry.

But this mode of labor deployment was not really anything new. The "crowdsourced" aspect is simply a digitized form of outsourcing, while the "micro-tasking" dimension is a conventional piece-rate. In fact, rather than being a cutting-edge form of labor deployment, "crowdsourced micro-tasking" is simply a modern iteration of one of the oldest forms of manufacturing: the "putting-out" system that dominated early industry. In the classic example of the system within the textile and garment industry, merchant “clothiers” would provide households with a certain quantity of fabric that was then processed by “out-workers” operating from their own homes. These out-workers took on a variety of tasks, with certain households specializing in certain types of work and different tasks distributed between different family members:

> They worked as carders and spinners who turned wool into yarn; weavers who interlaced yarn on a loom to produce cloth; fullers who washed the cloth to remove natural oils and give it a thick baize-like finish; dyers who added colour; and shearmen who gave it a smooth finish. Clothiers brought each of them in turn the material on which they were to work and then took the product away to sell for profit.[16]

In other words, these households conducted all the artisanal work that still needed a specialized human touch (what would today be referred to as a “Human Intelligence Task”) before the end product could be fed back into the machinery of the market. One crucial aspect of this early putting-out system was therefore that the clothiers didn’t exert direct control over the production process itself. The skillsets and even the tools used in these handicraft industries were controlled by the household.

Amazon’s Mechanical Turk operates in much the same fashion, connecting merchants who need to process a raw quantity of “Human Intelligence Tasks” with workers who will use whatever methods they deem fit to convert the raw material into a completed task within the confines of their own home. As Marxist thinker Kei Ehara argues: "digitalized work organization … is a modernized form of a putting-out system, rapidly setting the norms for the 21st century. In this online work organization, labor is 'logged,' and informal and humane contact between workers is minimized."[17] Though in the history of any given product line, the "dispersed type" of work organization seen in the putting-out system sometimes precedes a move to a larger-scale, "concentrated type" of work organization – indicated by the implementation of a factory regime – the two types are always being reproduced alongside one another as commodification extends to new spheres of life. "Microtasking," gig work, artisanal production, and any number of informal services will therefore always be regenerated in new forms by capitalist development, especially within new industries where the importance of hard-to-quantify and hard-to-automate "human intelligence" skillsets ensure that work retains a handicraft character. In addition to filling in the skill gaps of the more formal factory regime, this dispersed type of work organization also offers a certain degree of flexibility within the economy as a whole, serving as an employment reservoir that sops up job losses caused by both shorter-term economic crises and long-run trends such as ongoing mechanization in the more formal manufacturing sector.

The real serendipity of the ImageNet story was therefore not a chance hallway conversation so much as the emergence of the economic crisis shortly after the advent of the new gig labor *regime.*

It is no coincidence, then, that ImageNet was only able to find a readymade pool of cheap labor in 2008, at the height of the Great Recession that followed the collapse of the US mortgage bubble. Between 2007 and 2009 the official US unemployment figures (an underestimate of the true numbers) doubled to roughly 10 percent, accompanied by an even larger increase in various forms of underemployment.[18] At the time, most Mechanical Turk workers ("turkers") were based in the US, with the next largest group based in India – the only two locales where Amazon allowed cash payment, rather than payment in Amazon gift cards.[19] These were the sort of jobs specifically marketed as not jobs at all but instead "make money in your free time" opportunities that, at best, become a form of part-time work. And yet, faced with a dire job market and stagnant wages in the wake of the crisis, "microtasking" soon became a macroscopic sector. From a niche industry employing a few thousand workers in the mid-2000s, Mechanical Turk alone had grown to employ anywhere from 100,000 to half a million workers by the late 2010s, depending on how one measures "employment" in a sector where the companies insist very strongly that their workers are not employees.[20] The real serendipity of the ImageNet story was therefore not a chance hallway conversation so much as the emergence of the economic crisis shortly after the advent of the new gig labor regime.

From the point of view of the tech firms, the new industry was a gold mine. In a talk given to an audience of fellow tech bros in 2010, Lukas Biewald, CEO and co-founder of early data labelling firm CrowdFlower, explained the market logic in blunt terms:

> Before the Internet, it would be really difficult to find someone, sit them down for ten minutes and get them to work for you, and then fire them after those ten minutes. But with technology, you can actually find them, pay them the tiny amount of money, and then get rid of them when you don't need them anymore.[21]

Bieman's firm was one of several copycat crowdworking platforms that emerged in these same years in an attempt to grab market share in the new industry. CrowdFlower (now Figure Eight) was founded in 2007 and raised tens of millions in its first few years before being sold for $300 million in 2019 to AI firm Appen. As early as 2013, the global "online outsourcing industry" was estimated to be generating $2 billion annually, with 48 million registered workers (10% of whom, or 4.8 million, were "active" at any given time).[22] For comparison, in that same year, the global steel industry was estimated to directly employ 2 million workers and 2 million contractors (4 million in total, slightly less than "active" online outsourcing workers) and to indirectly employ 50 million people via related industries such as construction, transport and industry (slightly more than the total "registered" online outsourcing workers).[23]

Throughout, firms always insist that tasking not be treated as a primary source of income or even a "job" as such. Similar to the putting-out system, the work is often used either as a subsidy to either bolster the income of a household earning their subsistence by other means or to desperately backfill a growing deficit. Actual stories from workers in the sector paint a far less rosy picture than those promoted by the platforms:

> Katie Boehm of Pittsburgh turned to turking in 2017 after her husband, who has diabetes, lost his job and insurance coverage. Her own health issues keep her from working outside the house, and turking seemed like a lifeline.
>
> She turks at least 50 hours a week, sets herself a minimum goal of $20 a day and usually makes $30 to $50.
>
> Her husband's insulin costs $1,500 a month.[24]

Another survey of participants (in 2010) found similar results: "study participants included a senior looking to supplement his or her fixed income; a laid-off accountant making $150 to $200 a week to stay afloat while looking for a new job; a schoolteacher trying to make ends meet …" And, while many used the site for "more abstract, less financially immediate reasons," the reality was that "a not-insignificant number of people, in this country and abroad, aren't just using Turk as a means to make a couple bucks here and there – they're using it to replace a job."[25]

The sector also has extremely high turnover since its maddeningly low pay and frustrating power imbalances ("requesters" often simply refuse to pay workers, and the platform often has errors that prevent payments from being made) ensure a steady outflow of novices, leaving the bulk of work to “super-turkers” who use their own automated systems to tilt the scales ever-so-slightly back in their favor. However, even these individuals are wary of the platform. According to one such “super-turker” interviewed in the late 2010s: “it’s great for what it is, a side gig … But the state of America is that people are turning to something that shouldn’t be a job and trying to make a living off it.”[26] Estimates of the average hourly income of turkers across the years have varied wildly, but the most reliable studies have shown that “only 4% earned more than $7.25/h” with a median closer to $1.50 or $2 per hour.[27] A Pew Research study focusing on US-based turkers found that they “generally report earning less than minimum wage,” with about half of those in the sample reporting “a rate of less than $5 an hour.”[28] Meanwhile, firms like Amazon charge a tax on their crowdsourcing platforms – in the case of Mechanical Turk, a “requester” paying $.01 to a worker per task will also pay $.01 to Amazon – and also use the platforms to perform their own routine work without having to hire additional staff. In fact, researchers regularly note that Amazon itself is one of the biggest “requesters” on its own platform.[29] In both senses, “crowdsourced microtasking” operates as a fairly brazen way to avoid labor laws and milk a lucrative stream of rents from the platform.

The Production Floor

As Li’s ImageNet helped ignite a race for annotation and validation work able to meet the new demands posed by the era of “Big Data,” growing social media platforms also began using the same platforms for content moderation. Soon, however, they realized that the scale and continuous nature of the work would require additional oversight. While informal, hyper-flexible forms of gig work similar in kind to the putting-out system are continually reproduced by capitalist development – rather than being historical artifacts or even signs of a renascent "feudalism" centered on the tech platforms – these forms of labor deployment are also continually being reorganized to better accord with the changing needs of global production. Broadly speaking, every new product line will gradually shed its "handicraft" character under competitive pressure, leading to new forms of mechanization, automation, and corporate reorganization that then changes the nature of the work itself. Within the online outsourcing industry, social media firms helped to drive this process, rolling out the first wave of rationalization and centralization in what had initially been a largely artisanal industry.

For work in the process of being mechanized, the first stage is simply bringing the outworkers "in" to physical work at company facilities or, digitally, to work on company-controlled platforms, where labor discipline is more strictly enforced, control over time allows for the extraction of more work in any given day, and new production techniques can be rolled out at scale. The classic example is the relocation of weaving work from the household into the early factory, which first simply gathered together workers doing the same activities that they did in the home but then, through this concentration, also enabled the technical transformation of the labor process as new machines could be rolled out and linked to new power systems (i.e. water mills, then steam engines). Of course, the actual history of this transition was more piecemeal. When one element of textile production moved into the factory, others remained in households. And yet the emerging geography of production nonetheless privileged certain households in proximity to factories or major end-markets, while falling prices induced by mechanization also exerted a subtle but persistent pressure on remaining outworkers which did, in the end, drive them to extinction in all the industrialized countries (even while new outworkers emerged in newly-industrializing countries).

Digital work is no different: from a highly-dispersed, home-based labor regime in the late 2000s, microtasking soon took on more and more features of the factory regime, even while remaining informal. By the 2010s, there had emerged a growing number of digital sweatshops scattered across the world – though primarily concentrated in poor anglophone countries like India and Kenya, given the English-language bias of the internet – specializing in data annotation and content moderation for the major social media and crowdsourcing platforms. Conservative estimates placed the total number of workers employed in content moderation alone at around 100,000 by the early 2020s, with the overall value of outsourced content moderation services estimated to be around $13 billion in those same years and rapidly growing.[30] Facebook employed roughly 30,000 digital "safety and security" workers as of 2019, while TikTok employed 40,000 content moderators and other "safety" workers as of 2025.[31] These employment figures are similar in size to other professional services such as editing and translation. Unlike most professional services, however, the work itself is brutal, both physically and mentally. Content moderators are forced to sift through flagged content, watching videos depicting violent deaths, animal abuse, and sexual assault, or reading conspiracy theories and racist hate speech. Even though psychological studies clearly demonstrate that "repeated, prolonged exposure to specific content, coupled with limited workplace support, can significantly impair the psychological well-being of human moderators," content moderators are made to work long shifts ranging from 8 to 12 hours a day.[32]

While content moderation is similar in kind to other "microtasking" jobs, the sensitive nature of the content and the need for stricter oversight over employees means that the work is often done in a centralized office where workers can be monitored and supervised via clear managerial hierarchies. The most notorious of these sites are in poorer countries like Kenya, which have seen a series of ongoing labor struggles and a number of lawsuits filed against parent firms. In 2019, for example, local content moderation and AI services subcontractor Sama (working a contract for Meta) illegally fired migrant worker Daniel Motaung after he attempted to convene a union, resulting in a lawsuit. Ironically, the press from the case then helped Motaung attract more discontented workers to successfully found the African Content Moderators Union in 2023.[33] That same year, more than a hundred workers involved in the campaign (many of them migrants from other African countries) were fired, placed on a blacklist, and refused their legally mandated backpay.[34] In 2023, the fired workers filed a lawsuit against Meta, Sama, and Majorel (the company Meta switched to after the initial controversy), alleging violations of labor law and exposure to hazardous work without adequate protections.[35] The same trends would follow in data annotation, especially as demand increased with the AI boom, leading to a similar series of controversies and the emergence of new regulations and worker's organizations such as the Data Labelers Association, formed in Kenya in 2025.[36]

Meanwhile, similar supply chain structures were also reproduced outside the anglosphere, with hundreds of thousands of migrant workers in China sourced from poorer provinces like Gansu and Guizhou via labeling centers or recruited via vocational schools (a system first piloted in manufacturing) to work as annotators, assembling the raw material that would be fed into groundbreaking models like Deepseek. In interviews performed by Xia Bingqing of East China Normal University between 2018 and 2019 among "interns" recruited through the computer science departments of vocational schools, "some of them were unpaid, while others were paid by the amount of data they processed – for example, 0.2 yuan (three cents) for each 'bounding box' they used to label an image."[37] By the 2020s, rapid increases in demand and advances into video and more complex imagery, paired with new labor laws designed to crackdown on recruitment of interns paid under the minimum wage (or not paid at all), also saw Chinese firms begin to outsource work in much the same fashion as Meta and OpenAI.[38] Given its established position in the industry and persistent unemployment youth rates as high as 67% (among those 15-34 years old), Kenya has again proven to be a major hub.

Given their revenue constraints, Chinese firms in Kenya tend to subcontract in a slightly more lean and opaque fashion than their wealthier Western counterparts, recruiting largely via word of mouth, lacking any physical office presence, and eschewing most features of formal employment. Once recruited, annotation workers at firms subcontracting for Chinese tech companies are

> onboarded via a simple Google Form, managed through WhatsApp groups of up to 30 members, and paid through the local fintech app M-Pesa. There is no formal contract, and work is executed through short-term projects of around two weeks, during which the annotators are expected to work seven days a week.[39]

Annotation itself occurs through third-party platforms and the team is managed entirely within the WhatsApp group, each of which has a supervisor. As with content moderation, these annotation groups

> work like digital factory floors: daily rankings, production charts, motivational messages, and reminders to push harder.
>
> Every day, administrators share reports ranking output and accuracy. They also hold stand-up calls several times a week to review reports, fix accuracy issues, and push the team to work faster. When one worker lags, others are ordered to pick up the slack. To get paid, workers must maintain at least 85% accuracy.[40]

The guiding logic is the same as in traditional manufacturing: greater throughput per hour, achieved by whatever means possible. As explained by one supervisor: “The bigger the project, the more people we hire and the lower the rates [we offer].”[41]

The immediate cause of any given struggle in both the content moderation and data annotation sectors sector is usually the same: some combination of low pay, authoritarian management and, in the case of content moderation, extended exposure to hazardous work.

Though Kenyan workers are currently leading the drive to formulate new forms of organization within these digital sweatshops, similar struggles are common throughout the world. In 2025, a nearly identical lawsuit was filed against Meta and Majorel by content moderators in Ghana, for example.[42] Conflicts can also be found slightly higher up the income ladder and outside the anglosphere, as in Turkey, where moderators employed by Canadian firm Telus Digital on behalf of TikTok – specializing in Kurdish, Turkish, Arabic, and Azerbaijani content and often paid under the Turkish minimum wage – organized with a local union in 2024 and, as a result, saw fifteen leading organizers fired by Telus in retaliation.[43] In response, the workers filed a lawsuit similar to those working their way through the court systems in Kenya and Ghana.[44] Soon, these emerging unions had formed an international confederation, the Global Trade Union Alliance of Content Moderators (GTUACM), in the hopes of better coordinating their activities at the global scale of the labor regime itself.[45] At its founding in the spring of 2025, the GTUACM had representatives from

Kenya, Ghana, Turkey, the Philippines, Poland, Colombia, Portugal, Morocco, and Tunisia, and was in the process of recruiting new member organizations in Ireland and Germany.[46]

The immediate cause of any given struggle in both the content moderation and data annotation sectors sector is usually the same: some combination of low pay, authoritarian management and, in the case of content moderation, extended exposure to hazardous work. In Kenya, the content moderation workers were paid a mere $1.50 an hour (for 9-hour days) throughout the 2010s, until growing international media attention placed pressure on Meta to command its subcontractors to raise the amount to $2.20 an hour in 2022 (roughly $439 a month).[47] In Turkey, workers were paid between $490 and $900 a month (the minimum monthly wage is $566), equivalent to roughly $2.7 to $5 an hour – not that much more than in Kenya, even though the cost of living is nearly twice as high. Despite being exposed to extreme images of violence, death, and a constant slew of hate speech, content moderation workers are also given little in the way of mental health support. In the voice of one of Sama's workers in Kenya: "The work that we do is a kind of mental torture … Whatever I am living on is hand-to-mouth. I can't save a cent. Sometimes I feel I want to resign. But then I ask myself: what will my baby eat."[48] Several workers were later diagnosed with PTSD and other mental illnesses due to the nature of the job and, in one prominent case, a Nigerian migrant worker in Kenya doing content moderation for TikTok – via the Kenyan subsidiary of French digital services subcontractor Teleperformance – committed suicide after her repeated requests for leave were denied.[49]

But even content moderators in wealthy countries suffer similar problems. One of the higher-end content moderation firms is Indian-American professional services vendor Cognizant, which operates the majority of its facilities in India (particularly in Chennai), but also has centers across the world, including in high-income countries like Ireland and the US. In almost every country it operates in, however, Cognizant has been plagued by a series of controversies ranging from bribery (in India) to violations of labor law (Ireland) and wage theft (in the US). Two of the firm's large content moderation offices in Phoenix (employing some 1,000 workers) and Tampa (800) both made the news in the late 2010s and early 2020s for their working conditions, with office workers paid a mere $15 an hour, offered meager mental health services, and placed under extreme stress due to repeated exposure to traumatic material, such that at least one worker (in Tampa) died on the job from a heart attack induced by the content.[50]

Similarly, in 2021, one US-based worker employed by Canadian firm Telus International, a digital services contractor, sued TikTok claiming that her content moderation work for the platform had given her PTSD, since she regularly worked 12-hour days viewing "videos of the genocide in Myanmar, mass shootings, children being raped, and animals being mutilated."[51] Moreover, workers are often not told up front about the nature of the work. In Kenya, moderators consistently report being lied to about what jobs they were being hired for. As described Angela Chukunzira, Tech and Society Fellow at the Mozilla Foundation, many of these content moderators

> have to migrate to another country for work. And… they are lured, like you're told you're going to do call center work or you're going to translate… Then when you land in the job, you realize it's a trauma-inducing job, you're exposed to graphic images all day, no psychosocial support, no therapy…[52]

Companies like Cognizant were also alleged to have tricked workers in high-income countries into taking on the work by telling these workers that they were being hired for a general marketing position and then revealing late in training the real nature of the job.[53]

Data annotation and content moderation work are where the modern division of labor in the modern digital supply chain first began to grow more intricate, filling in the space between the bottom rungs of turkers and the formal staff employed directly by the tech platforms. At the Tampa site where the worker died on the job, Cognizant reportedly even "calls the part of the building where contractors do their work 'the production floor,'" distinguishing it from the higher floors specializing in more "intellectual" services.[54] As in other supply chains, this division of labor is attended by intricate technical and geographic sub-divisions in pay: while the Kenyan content moderators were paid roughly $5,268 annually, the Turkish workers could make as much as $10,000, and the Phoenix workers earn a solid poverty wage of $28,800 complete with "nine minutes per day of 'wellness' time." Meanwhile, the average formal employee at Meta, the lead firm for whom many these people were working, had a total compensation package valued at roughly $240,000 a year.[55] And even this figure is a laughable fraction of the billion-dollar fortunes that major investors and executives were accumulating in these same years simply by owning shares in these companies.

This divide was also evident in the different character of the individual labor regimes and working environments faced by each stratum of workers. While Meta's formal employees in places like Silicon Valley and the Seattle area enjoy flexible hours in open floorplan offices with innumerable amenities and ample access to paid time off, content moderators in Phoenix and Tampa face long hours and compulsory night shifts in dirty offices infested with bedbugs where "[v]erbal and physical fights break out on a monthly basis…"[56] In the digital sweatshops in Nairobi, conditions are even worse, with companies actively seeking out the most vulnerable workers. Sama, for example, prioritized hiring from Kibera, one of the largest slums in the world, where it runs a specialized training center. Moreover:

> Sama doesn't just recruit in slums, but also from refugee camps, where people are even more desperate. Other organizations do too. M2Work is a collaboration between Nokia and the World Bank that targets jobless Palestinians – cheap, desperate labor that could be put to better use. The NGO Lifelong trains Syrian refugees to become data labelers. From Dadaab in Kenya to Shatila in Lebanon, refugees are paid the lowest wages possible to make the richest companies richer. [57]

In addition to divisions between different grades of digital labor, then, workers doing the exact same activities are assigned wildly different wage rates simply based on their location and degree of desperation. As described by Muthuri Kathure, a lawyer working with content moderators in East Africa: "…we are all Africans and we will be treated different from how the global north is being treated." In response, Ellam Brian, General Secretary of the Kenya Union of Gig Workers, compared the situation to that faced by other gig workers: "So what is the difference between an Uber driver in the UK and an Uber driver here? … Are we children of a lesser god?"[58]

Macrodata Refinement

One of the perpetual refrains in the data annotation and content moderation industries is that the work is a temporary necessity, performed by humans only until it can be effectively automated through more refined algorithms. On the surface, this initially appeared to be true. By the 2020s, for example, the early rounds of content moderation for social media were increasingly being performed by AI tools, as were simple forms of annotation once tasked out to turkers. But the catch is that these AI models required their own content moderators and annotators, especially in the early stages of training and development. As a result, the demand for content moderation work has only grown. Similarly, as AI models have grown in popularity and scope, the annotation market has exploded. By the mid-2020s, digital sweatshops across the world had reoriented away from a sole dependence on social media toward increasing integration with the AI supply chain. Today, many members of the African Content

Moderator's Union are still employed by digital service firms like Sama but, in addition to direct moderation of social media content, they now also increasingly work on contracts for firms like OpenAI.

Meanwhile, the advent of modern machine learning also brought with it an even more advanced division of digital labor: new forms of task-based annotation, ranking, editing, proofreading, translation, and other training work spanning a wide range of disciplines, all necessary to construct the very models that appear to "automate" work in these same sectors. It is at this point that the effectively handicraft model of task work and the semi-handicraft character of content moderation begin to give way to a properly industrial-scale supply chain. Meanwhile, in the same years that the news broke about the digital sweatshops, the boom in machine learning was also generating a demand for new, higher-skill inputs to provide more intricate training for the models beyond the data already available on the internet. Though some of this work entails annotation by subcontractors similar to that performed by the anonymous turkers who built the ImageNet database more than decade ago, it also includes various refinement tasks such as ranking, rewriting, and reviewing model output as well as feeding complex inputs into new models specializing in advanced reasoning, creativity, or STEM skills. Given that these tasks required individuals with higher levels of training to complete, the tech platforms saw an easy opportunity to solve two problems at once: by expanding hiring in wealthier countries, they could both obtain the skilled labor necessary for this higher-order work producing bespoke data and refining model output and, given that they'd have to pay higher wages, they could then gesture toward these elevated wage rates as evidence that they'd fixed the problem of exploiting low-end labor in digital sweatshops, even if nothing had really changed at the bottom of the supply chain.

Whether you are a coder, an editor, an artist, or even a scientist, AI is not going to "take your job." Instead, it will remake your job such that you are doing many of the same activities but now dismembered into arcane, piecemeal, and monotonous "tasks" for less pay and longer hours under a far stricter labor regime.

Ultimately, rather than replacing "intellectual" work, machine learning models are decomposing it much like early factory regime decomposed artisanal knowledge in sectors like weaving: certain portions of the work are mechanized, encoding the embedded knowledge of the craft into machine systems and, in so doing, proliferating new forms of deskilled and semi-skilled work that nonetheless require the cooperation of workers with older craft knowledge for their construction and maintenance. The bright mirage of AI as an "automated" system that will replace tens of thousands of jobs thereby serves to obscure a darker horizon. Whether you are a coder, an editor, an artist, or even a scientist, AI is not going to "take your job." Instead, it will remake your job such that you are doing many of the same activities but now dismembered into arcane, piecemeal, and monotonous "tasks" for less pay and longer hours under a

far stricter labor regime. In the end, you will be chained to a desk buried somewhere in the subterranean layers of machine systems that, far from being sentient or even "intelligent," churn forward mindlessly in an arbitrary direction set by central tendencies in macroscopic datasets, as slanted by prevailing social prejudice, crushing whatever lies ahead and burning the world in their wake.

In 2023, faced with narrowing job prospects and a shrinking number of editorial and translation contracts (due, in part, to the AI boom), I found my income falling month by month in the midst of the largest inflationary wave of the last fifty years. I effectively had an entire part-time job just writing cover letters and putting in hundreds of job applications every week for positions within in my actual fields of expertise – GIS, the social sciences, and publishing – only to receive terse rejections months later, many with a disclaimer hidden somewhere at the bottom noting that all applications were filtered by AI systems before being sent to human review. Meanwhile, the night shift warehouse job I took to patch the growing gaps only paid 70% of my rent when I started, and 60% when I had to quit after one of many nighttime theft attempts ended up totaling my truck.[59] It was in this context that a random LinkedIn recruitment message popped up, offering me up to $60 an hour for work that involved writing, editing, and proofreading specialized data to be fed into "cutting edge AI models." I navigated to the platform, signed up, verified my identity, was "hired" (as an independent contractor, not an employee, of course) and then encouraged to complete the training within a set time period for a small cash reward.

After the long training period, I was then assigned to a series of projects that had absolutely nothing to do with the information in the training. Instead, I performed a huge variety of "tasks," many of them familiar editorial services, some requiring advanced training as a social scientist. Meanwhile, the actual wage proved to be far lower and far more volatile: ranging from a low of $16 an hour to a high of $50 or so, depending on the project, and always subject to change. Sometimes I'd start a project at $35 an hour only to sign in the next day and see that the pay rate had been a "mistake," and subsequently been dropped to $20 an hour. Switching between projects every few weeks also required completing long, unpaid trainings, reading through 100+ page instruction documents, and scanning instruction updates posted to overactive Slack channels bejeweled with flashing emojis and repeated warnings to turn off your VPN, all capped by arbitrarily designed tests often riddled with errors, resulting in ejection without pay if failed and, if passed, followed by reduced pay rates for training tasks (usually less than half the advertised hourly pay) and a review period in which aggressive throttles limited the available work. One contractor working for the same company reported "that they'd spent close to 40 hour sin a single month in unpaid onboarding sessions without landing any actual work…"[60] After a week or two of these preliminary activities, the throttle would be lifted, I would grind out as many hours as possible – often working overnight, following my same warehousing schedule – only to find that the project had suddenly "finished," after which the cycle would start again.

The exact tenor of the work is difficult to capture in any matter-of-fact description. As an AI worker, I lived in a box inside the machine, my daily tasks laid out in cryptic letters dropped through a slot in the wall, which I replied to in kind. Every once in a while, I would be moved to a different box, equally descriptionless, and all connected by a maddening labyrinth of mirrored rooms and iterated hallways, most of which appear to be empty. I would never see another human being, even though I could hear them all around screaming, crying, mumbling, humming their way through their tasks. My projects were ever-shifting and inscrutable. It was never clear exactly what I was doing, what I was supposed to be doing, or whether I was doing it right. The projects had whimsical code names: White Wolf, Genesis Dolphin, Jellyfish IF, MM Biscuits. The name of the company itself seemed to regularly change, as did the platform, preferred payment methods, and the means of communication with supervisors. The supervisors rarely spoke except to berate people and seemed to cycle in and out just as quickly as the rest of us.

One day I rank AI-generated responses to real user prompts according to constantly changing rubrics. Another day I write the rubrics. I write the prompts. I rewrite the responses. Sometimes I write the prompt and then write the entire response myself according to the chatbot's stylesheet. I read a seemingly infinite number of real user queries asking an AI chatbot to produce a script for the user's favorite anime but instead of a plot the characters are boyfriend and girlfriend and they kiss and get pregnant – and I wonder quietly to myself what sort of soul the person who would ask such a question is sculpting for themself. There is another flood warning outside and the rain crashes down in what I imagine must be wrath at some cosmological unbalancing. I sit through hours of unpaid onboardings, watching automated videos narrated by a worker in the Philippines, reading through novella-length instruction documents that appear to have been designed by children. I fix the math. I write out step-by-step reasoning instructions for how to solve a problem. I fix the math again. I cannot take a bathroom break without losing pay. I shut the windows because there is another record-breaking wildfire. I upload a random image and ask a question about it. I see a random image and answer a question about it.

I record a prompt in an “angry” voice. I rank another worker’s audio recording based on whether it actually sounds “happy” or whether I think they are faking it. If I hear their children crying in the background the rules mandate that I fail the task. Life cannot intrude. Outside there is a homeless man with no shirt on and some sort of skin rot common in the Middle Ages, carrying an axe over his shoulder like a lumberjack. I listen to a voice recording of someone sobbing and it sounds real. I determine whether the model has correctly understood that the human voice it is hearing is “sad” and then I decide whether its own synthetic voice response has adopted the correct tone of “empathy.” These voice recordings are the only time I ever hear my co-workers speak. One day I am instructed to travel to a “commercial space” and record a brief prompt in an upset voice, emulating something one might say in an argument between romantic partners. It had been such a cold, dark winter, washed in the waters of endless atmospheric rivers pouring sheet after sheet of heavy rain into the flooded streets so I drive to Home Depot and find my way to the light fixture aisle where the walls glow with factory-yellowed bulbs and hundreds of lit chandeliers hang overhead glittering like a nebula and the air is warm with electricity buzzing through the filaments of all these beautiful machines. I squint into the tunnel of light and hold my phone up to my mouth and say, in a voice on the edge of tears: “I don’t want to be alone.”

If a worker were to walk by at that moment and politely ask me what I was doing, my response would be frank: “I don’t fucking know.” The first lesson in AI work is dealing with an unknowing, persistent dissonance: complete this task for no apparent reason according to these hyper-specific (but continually updated and often contradictory) instructions; you are an employee but not an employee; you have a manager, but they do not respond to your messages, tell you what to do, or decide if you will be kicked off the project; if you have a question, ask it on the forum (you have not been given access to the forum); your wage can change at any moment, surging and dipping by as much as $20 an hour; you rely on the work for a major chunk of income, but it is not your “job.” This is what is like to work in the middle rungs of the AI supply chain, completing specialized tasks designed to refine the large-scale data used to train modern machine learning models. Exactly what sort of work you are doing changes every few weeks, though there are broad similarities between projects. As in the case of Mechanical Turk, the work itself is not even described as work and the workers are not workers. Sometimes you will simply annotate for a few minutes. Sometimes you will spend hours doing research, fact-checking, and writing full essays. In each case, you are performing a “task” and are therefore described as a “tasker” or, even more dystopian, as an “attempter,” employed by a third-party company managing contracts from any number of leading tech firms.

Many tasks remain focused on issues of annotation. For this reason, it is often presented, like content moderation, as ephemeral work necessary only in the early stages of model construction. As described by tech journalist Josh Dzieza, the tech firms that need the work done tend to treat it in an extremely dismissive fashion: "You collect as much labeled data as you can get as cheaply as possible to train your model, and if it works, at least in theory, you no longer need the annotators."[61] This similarly tends to be how the handful of engineers employed at these AI firms think of the work, which is treated as more or less an afterthought – much as Li once treated the bulk of labor required to build ImageNet. However, in practice, the opposite seems to be the case: on the whole, AI training work has not proven ephemeral but has instead experienced rapid growth and, aside from the initial over-hiring common to any semi-speculative boom, the sector does not appear set to shed much of its workforce soon. Moreover, as in any industry, even if further automation makes the production of models faster (for example, by integrating lower-order models into the training process), this does not mean that human work will disappear. In other words, while much of the training work is ephemeral for each individual model, there is no reason to assume that we will stop making models over time or that their updates will not require similar labor. In fact, just the opposite is the case: just like any other sector of manufacturing (whether hardware or software), the industry is premised on the continual production of new models. From automobiles to operating systems, firms are compelled to continue iterating their core designs indefinitely. Similarly, new "edge cases" will proliferate at the margins of existing models, requiring continual re-training to update the model, and entirely new use cases for machine learning systems will also emerge, requiring new forms of human training.

For his investigative report on the sector, Dzieza obtained an annotation job, where his tasks consisted of labeling photos with clothing under a seemingly simple instruction: "DO LABEL items that are real and can be worn by humans or are intended to be worn by real people." After his long unpaid training, he then proceeds to the unpaid test that determines whether he will be allowed to continue:

> Feeling confident in my ability to distinguish between real clothes that can be worn by real people and not-real clothes that cannot, I proceeded to the test. Right away, it threw an ontological curveball: a picture of a magazine depicting photos of women in dresses. Is a photograph of clothing real clothing? *No,* I thought, *because a human cannot wear a photograph of clothing.* Wrong!

As this example makes clear, even the most basic annotation tasks often pose strange difficulties that have little to do with human reasoning. Instead, the tasker is required from the very beginning to mold their mind according to the strict requirements of the model, as communicated in cryptic instructions transmitted down in a game of telephone from frustrated engineers. Dzieza quickly learns the first rule of tasking: never trust the instructions.

After working his way through the bizarrely-designed instruction document which "basically consisted of the same direction reiterated in the idiosyncratically colored and capitalized typography of a collaged bomb threat," he reaches the actual tasks

> only to make the horrifying discovery that the instructions I'd been struggling to follow had been updated and clarified so many times that they were now a full 43 printed pages of directives: Do NOT label open suitcases full of clothes; DO label shoes but do NOT label flippers; DO label leggings but do NOT label tights; do NOT label towels even if someone is wearing it; label costumes but do NOT label armor. And so on.

This saga of (unpaid) trial and error is usually the first of many surreal experiences for new taskers, soon followed by the realization that the "help" forums are nothing but flashing warnings hovering over a flood of unanswered queries – though managers will always swoop in to warn against sharing your pay rate – or that training videos are, as a rule, recorded by either extremely fresh graduates with a deer-in-the-headlights anxiety or by middle-aged homeworkers located overseas.

But the basic annotation work Dzieza was tasked with is only a preliminary portion of the AI supply chain, today performed largely in digital sweatshops in poorer countries like Kenya. Annotation is similar in kind to the early refinement of raw materials for the manufacture of the model. After this initial feeder data is validated and gets sent to the model, the assembly work begins. This is where specialization and higher levels of generalist training are prioritized. Assembly tasks are more involved and intricate, ranging from simply ranking two model responses against one another, to systematic fact-checking and other content moderation work, to writing challenging prompts with a certain number of logical constraints, to writing your own responses to actual prompts input by users. Sometimes you will be tasked with creating entire bespoke datasets of prompt and response, pretending to be the user and then pretending to be the model. Often, you will even write the rubric. As new use-cases for machine learning proliferate, the details of this assembly work also change. Now that many AI chatbots are attempting to implement audio capabilities, tasks often involve similar ranking work applied to audio clips, with writing replaced by recording.

Though ostensibly offered a choice between projects, workers rarely have much to choose from. Similarly, you can be removed for a project at any time for no reason (almost always with a default "quality issue" message, regardless of your actual performance) and suddenly switched to lower paying or poorly managed projects that have been "prioritized." Moreover, if you refuse any project (even if it is because you do not meet the qualifications) you are more likely to be left in the limbo of having no work at all, a situation referred to as "EQ" (Empty Queue), which can last for weeks at a time. When I first started the job, after onboarding for a completely different ranking project, I was then given a first task in an unrelated project, where I was provided with a prompt – possibly a real user prompt, more likely written by another tasker – and required to write out the response myself, according to the parameters provided by the user and a complex style guide. The prompt itself was requesting a roleplay scenario in which the user tricked Hank Hill into eating meat cooked on a charcoal grill. Since the model was supposed to play the part of Hank Hill, I spent the next half hour or so figuring out a dialogue, looking up the proper formatting requirements, and then double-checking the instruction documents, only to have the entire platform crash when I hit submit, apparently realizing its error and kicking me from the project to which I had never been assigned. Through the experience, I quickly learned the next few lessons presented to every new tasker: anything that can go wrong with the platform will go wrong with the platform and, when things go wrong, you will not get paid.

Over the following months, I worked on an enormous range of projects, most of which drew on writing experience, and some of which drew on my training in the social sciences. But the platform was not particularly good at identifying specialties or even confirming that its contributors could speak adequate English. Periodically, I would get assigned geochemistry tasks simply because my resume mentioned that I had graduate degrees in "geography," which the platform automatically classes under a broader "earth sciences" label rather than alongside fields like history and sociology, or even proximate STEM fields like GIS and statistics. In each project, the work itself was essentially copy-editing, proofreading, field-specific fact checking, and other editorial services or consultant work commonly performed in para-academic settings, but at pay rates significantly lower than the standard for freelancers. Since rapid-fire deadlines, breakneck hiring, and an absolute refusal to hire sufficient administrative staff or even editors for the onboarding systems and instruction documents on the platforms all led to plummeting task quality, more experienced taskers were increasingly shifted to work as reviewers doing quality checks on other taskers' work for the same pay rates.

As time went on, the projects also grew more elaborate and their requirements more inscrutable. Soon, many required some sort of imagery: upload a picture of a specific type (charts, everyday objects, people, NO celebrities, NO personal information) and ask the model a complex question about it (spatial reasoning, data interpolation, fact-finding), then write an ideal response to that question yourself. By late 2024, an increased demand for audio data saw taskers asked to record their own voices, engage in conversations with the model or with other users, or to record fragments of dialogue in busy outdoor locations. More recently, the company launched a project to record video data, asking taskers to wear a GoPro while they work their day jobs or perform their hobbies (in particularly high demand were beekeepers and auto mechanics). However, to avoid capturing copyrighted content or personal information, the instruction document gives arcane guidelines that read like dark age superstitions: you must not listen to music, you must avoid uttering anyone's personal name, you must cover all mirrors and all photographs. NEVER speak to children. If you are a nature photographer and you hear a stranger approaching in the forest, you must look away from them immediately. DO NOT gaze upon their face. Cover any tattoos. Hide your appearance from the camera. Whatever is gestating inside the glass screen cannot be allowed any point of entry into the human soul.

[1] Though this process appears to be automatic, it in fact requires work performed either for free by the single-license user (when you install a new program, you are yourself tasked with checking for dependencies and conflicts and selecting the relevant settings) or for pay by the IT department whose responsibilities include rolling out installations and updates at scale across the enterprise. As anyone who has worked in these IT settings can attest to, even simple software updates often entail quite a bit of work when rolled out across hundreds or thousands of machines. Thus, while far more automated than the standard assembly process in manufacturing, this process really is not "automatic" in the true sense of the word.

[2] Or is perhaps most similar to modern, industrial agriculture itself, where a similar process of design-and-manufacture of new varietals precedes their sale-and-distribution, followed by a relatively automated process of growth, with profits derived mostly from control over key platforms and selling access to patented seedstock.

[3] Evgeny Morozov, "The problem with artificial intelligence? It's neither artificial nor intelligent", *The Guardian*, 30 March 2023. <https://www.theguardian.com/commentisfree/2023/mar/30/artificial-intelligence-chatgpt-human-mind>

[4] ibid

[5] Mateo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence*, New York: Verso, 2023. p.223

[6] Jarod Lanier, “There is no A.I.”, *The New Yorker*, 20 April 2023. <https://www.newyorker.com/science/annals-of-artificial-intelligence/there-is-no-ai>

[7] Morozov 2023

[8] Lanier 2023

[9] ibid

[10] ibid

[11] Dave Gershgorn, “The data that transformed AI research—and possibly the world”, *Quartz*, 26 July 2017. <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world>

[12] ibid

[13] Jessi Hempel, “Fei-Fei Li’s Quest to Make AI Better for Humanity”, *Wired*, 13 November 2018. <https://www.wired.com/story/fei-fei-li-artificial-intelligence-humanity/>

[14] Andy Newman, “I Found Work on an Amazon Website. I Made 97 Cents an Hour,” *The New York Times*, 15 November 2019. <https://www.nytimes.com/interactive/2019/11/15/nyregion/amazon-mechanical-turk.html>

[15] Ellen Cushing, “Amazon Mechanical Turk: The Digital Sweatshop”, *Utne Reader*, 7 December 2012. <https://www.utne.com/science-and-technology/amazon-mechanical-turk-zm0z13jfzlin/>

[16] Nicholass R. Amor, “The origins of the putting-out or domestic system of industrial production in England”, in Giampiero Nigro (ed.), *The knowledge economy: innovation, productivity and economic growth, 13th to 18th century*, Firenze, Italy: Firenze University Press, 2023. pp.263-285

[17] Kei Ehara, “Marx’s Ideas on Work Organization: Reinvestigating the Conceptions of Cooperation, the Division of Labor, and Machinery”, *Science & Society*, 87(4), October 2023. p.492

[18] Evan Cunningham, “Great Recession, great recovery? Trends from the Current Population Survey,” *Monthly Labor Review,* U.S. Bureau of Labor Statistics, April 2018 <https://doi.org/10.21916/mlr.2018.10>

[19] Panos Ipeirotis, “The New Demographics of Mechanical Turk”, *A Computer Scientist in a Business School*, 9 March 2010. <https://www.behind-the-enemy-lines.com/2010/03/new-demographics-of-mechanical-turk.html>

[20] Cecilia Maria Paradi-Guilford, Toks Fayomi, Saori Imaizumi, Panos Ipeirotis, Siou Chew Kuek, *The global opportunity in online outsourcing (English)*, Washington, D.C.: World Bank Group, 2015. <http://documents.worldbank.org/curated/

en/138371468000900555>

[21] Qtd. in Moshe Z.Marvit, "How Crowdworkers Became the Ghosts in the Digital Machine", *The Nation*,05 February 2014. <https://www.thenation.com/article/archive/how-crowdworkers-became-ghosts-digital-machine/>

[22] Cecilia Maria Paradi-Guilford, Toks Fayomi, Saori Imaizumi, Panos Ipeirotis, Siou Chew Kuek, *The global opportunity in online outsourcing (English)*, Washington, D.C.: World Bank Group, 2015. <http://documents.worldbank.org/curated/en/138371468000900555>

[23] According to figures from the WorldSteel Association. See, for example: WorldSteel, "World Steel in Figures 2013", *WorldSteel*, 2013, p.6. <https://worldsteel.org/wp-content/uploads/2013-World-Steel-in-Figures.pdf>

[24] Newman 2019

[25] Cushing 2012

[26] Newman 2019

[27] Kotaro Hara, Abi Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, Jeffrey P. Bigham, "A Data-Driven Analysis of Worker's Earnings on Amazon Mechanical Turk", *CHI '18: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, No. 449, 2018. p. 1 <https://arxiv.org/pdf/1712.05796>

[28] Pew, “Research in the Crowdsourcing Age, a Case Study”, *Pew Research Center*, 11 July 2016. <https://www.pewresearch.org/internet/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/>

[29] Newman 2019

[30] Note that, since these numbers focus on the Western internet market and exclude similar services in countries such as China, the total amount is likely to be higher. The employment figure comes from Miriah Steiger, Timir J. Barucha, Sukrit Venkatagiri, Martin J. Riedl, and Matthew Lease, “The Psychological Well-Being of Content Moderators: The Emotional Labor of Commercial Moderation and Avenues for Improving Support”, *CHI ’21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 07 May 2021. pp. 1-14 <https://doi.org/10.1145/3411764.3445092> ; and the valuation estimate comes from: Mike Pierides and Brooke Bruno, “Emerging Market Trend: An Overview of Content Moderation Outsourcing”, 30 June 2023. <https://www.morganlewis.com/blogs/sourcingatmorganlewis/2023/06/emerging-market-trend-an-overview-of-content-moderation-outsourcing>

[31] Casey Newton, “Bodies in Seats”, *The Verge*, 19 June 2019b. <https://www.theverge.com/2019/6/19/18681845/facebook-moderator-interviews-video-trauma-ptsd-cognizant-tampa>; Andrew Griffin, Niamh McIntyre, “TikTok workers who watched abusive and traumatic videos say they were fired after trying to unionise”, *The Independent*, 15 March 2025. <https://www.the-independent.com/tech/tiktok-videos-

moderation-content-telus-b2715766.html>

[32] Steiger et. al. 2021, p.1

[33] Billy Perrigo, “150 African Workers for ChatGPT, TikTok and Facebook Vote to Unionize at Landmark Nairobi Meeting,” *Time*, 01 May 2023. <https://time.com/6275995/chatgpt-facebook-african-workers-union/>

[34] Jody Ray, “Kenyan Courts Keep Telling Meta to Let Workers Unionize”, *Jacobin*, 14 February 2024. <https://jacobin.com/2024/02/kenya-courts-meta-content-moderation-union>

[35] Aaron Ross, Ayenat Mersie and Hereward Holland, “Content moderators sue Meta over alleged ‘union-busting’ in Kenya”, *Reuters*, 20 March 2023. <https://www.reuters.com/technology/content-moderators-sue-meta-over-alleged-union-busting-kenya-2023-03-20/>

[36] Sebastian Klovig Skelton, “Kenyan AI workers form Data Labelers Association”, *ComputerWeekly*, 14 February 2025. <https://www.computerweekly.com/news/366619321/Kenyan-AI-workers-form-Data-Labelers-Association>

[37] Viola Zhou and Caiwei Chen, “China’s AI boom depends on an army of exploited student interns,” *Rest of World*, 14 September 2023. <https://restofworld.org/2023/china-ai-student-labor/>

[38] The youth unemployment figure comes from the Federation of Kenya Employers: <https://www.fke-kenya.org/policy-issues/youth-employment>

[39] Damilare Dosunmu and Tessie Waithira, "The hidden Kenyan workers training China's AI models", *Rest of World*, 04 December 2025. <https://restofworld.org/2025/kenya-china-ai-workers/>

[40] ibid

[41] Qtd. in ibid

[42] Rachel Hall and Claire Wilmot, "Meta faces Ghana lawsuits over impact of extreme content on moderators", *The Guardian*, 27 April 2025. <https://www.theguardian.com/technology/2025/apr/27/meta-faces-ghana-lawsuits-over-impact-of-extreme-content-on-moderators>

[43] Kıvanç Eliaçık, "TikTok Moderators Fight against Trauma and for a Union", *Labornotes*, 09 April 2025. <https://labornotes.org/blogs/2025/04/tiktok-moderators-fight-against-trauma-and-union>

[44] Griffin and McIntyre 2025.

[45] Jess Weatherbed, "Content moderators are organizing against Big Tech", *The Verge*, 30 April 2025. <https://www.theverge.com/news/658566/content-moderator-union-alliance-meta-tiktok-google>

[46] Robin-Leigh Chetty, “Content moderators from across the globe have formed a union”, *hypertext*, 30 April 2025. <https://htxt.co.za/2025/04/content-moderators-from-across-the-globe-have-formed-a-union/>

[47] Billy Perrigo, “Facebook Content Moderators in Kenya to Receive Pay Rise Following TIME Investigation”, *TIME*, 2 March 2022b. <https://time.com/6153778/facebook-moderators-kenya-sama/>

[48] Billy Perrigo, “Inside Facebook’s African Sweatshop”, *Time*, 14 February 2022a. <https://time.com/6147458/facebook-africa-content-moderation-employee-treatment/>

[49] Stoneface Bombaa and Felix Omondi (Hosts), “AI Ni ya Who? The Invisible Labour Behind AI Technology”, *Until Everyone is Free*, 24 April 2025.

[50] Newton 2019b. This an other cases have been thoroughly documented and, while data annotation work remains relatively unknown, the dire conditions at content moderation plants are now widely recognized within the cultural sphere. See, for example, the 2025 thriller *American Sweatshop* by Uta Briesewitz, which follows a fictional worker in the industry.

[51] Clare Duffy, “TikTok sued by content moderator who claims she developed PTSD from reviewing disturbing content,” *CNN Business*, 29 December 2021. <https://www.cnn.com/2021/12/29/tech/tiktok-lawsuit-content-moderator/index.html>

[52] Qtd. in: Bombaa and Omondi 2025, 44:41.

[53] Newton 2019b

[54] ibid

[55] Casey Newton, “The Trauma Floor”, *The Verge*, 25 February 2019a. <https://www.theverge.com/2019/2/25/18229714/cognizant-facebook-content-moderator-interviews-trauma-working-conditions-arizona>

[56] Newton 2019b

[57] Bombaa 2025, 30:46.

[58] ibid, 36:56

[59] See: Phil A. Neel, “Quarter Pounds of Flesh: Part 1”, *The Brooklyn Rail*, Dec/Jan 2024-25. <https://brooklynrail.org/2024/12/field-notes/quarter-pounds-of-flesh/>; Phil A. Neel “Quarter Pounds of Flesh: Part 2”, *The Brooklyn Rail*, February 2025. <https://brooklynrail.org/2025/02/field-notes/quarter-pounds-of-flesh-2/>

[60] Charles Rollet and Ben Bergman, “Inside Scale AI after Meta”, *Business Insider*, 10 December 2025. <https://www.businessinsider.com/pay-cuts-poaching-pivoting-inside-scale-ai-meta-2025-12>

[61] Josh Dzieza, “AI is a Lot of Work”, *The Verge* and *New York Magazine*, 20 June 2023. <https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>

Discussion about this post