Matthew Vollo

Faking AI and AI "Magic"

"Any sufficiently advanced technology is indistinguishable from magic." This quotation from Arthur C. Clarke, though posed around 1968, still has a strong relevance today, where we are well into the era of artificial intelligence (Science). To the general public, products such as Apple's Siri, Amazon's Alexa and other artificial intelligences made for mass consumption can appear to work as if by magic. This ignorance from the general public regarding how AI is created, trained and brought to market is not a coincidence, and is in fact by design, or rather, in accordance with the wishes of tech companies and startups. It seems that tech companies such as Google, Amazon, Facebook and others wish to sell the general public the idea that machines can magically fulfill their desires, without actually telling them the amount and kind of work that goes into producing those very machines (Irani). This illusion of "magical artificial intelligence" is one that tech companies wish to maintain, as it serves as a venir that hides the true cost of the labor behind the machines.

I had first read about this concept of "faking artificial intelligence" while reading *Atlas of AI* by Kate Crawford, where she spends an entire chapter investigating the relationship between artificial intelligence and the labor that goes behind it. Rather than asking the rather oversaturated pop culture question of whether or not robots (and AI) will replace humans, Crawford was interested in the increasing ways that humans were being treated like robots (Crawford, 56). The labor behind artificial intelligence is made up of people performing monotonous, repetitive tasks to conjure up the impression that machines are doing this work. This labor is either used to get the artificial intelligence to a functioning level, possibly to the point where it can function well enough on its own or is used to completely maintain an illusion that an artificial intelligence is functioning by itself, when in reality it is being propped up and maintained by hidden laborers. I was curious as to how common this practice was within the tech sector, and in investigating its commonality, I wanted to see what ramifications this practice exactly had on the industry, and on society at large. For what does it mean for the fruits of an entire industry to possibly be illusions, and what does it mean when a society at large is ignorant of the inner machinations of an entire industry. In doing this research, I hoped to answer these questions and to figure out how often this practice was used, and what tech companies gained from hiding the labor that went into their products and in some cases, hid the hands that maintained an illusion that companies wanted the general public to see.

One thing that is unknown to most of the general public is how much work is actually behind the products that they use everyday, and specifically, behind the "magical" artificial intelligences that are presented to them and used by them on a daily basis. Specifically, there is an ignorance of how many unpaid workers are used and required to help build, maintain and test AI systems, yet they rarely receive the credit for doing so (Crawford 63). This type of automation, where humans behind the AI prop up its existence is called "human-fueled automation" (Crawford, 63) (Irani). In reality, and contrary to the narrative that tech companies try to push, hidden laborers often perform the repetitive digital tasks that underlie AI systems, with such work including labeling thousands of hours of training data and reviewing suspicious or harmful content (Crawford, 63). It was this practice that I wanted to investigate first: the concept of hidden labor training and helping to create artificial intelligence.

The usage of hidden labor to support artificial intelligence is so important that without it, AI systems could not function (Crawford 64). The AI research community relies on cheap, crowd-sourced labor in order to fine tune the artificial intelligence, as well as perform tasks that can't be done by machines (Crawford, 64). One of these tasks that human hands have to do involves the classification of cultural data, which is called "cultural data work" (Irani, 36). AI, and hence algorithms have to be trained by the generation of training data in order to teach these algorithms to recognize and match patterns. Despite the innate powers computers have for pattern matching, simply put, computers do not have the capabilities in order to interpret cultural material (Irani, 63). AI, in its current state, cannot classify cultural content, and thus is not equipped to perform content moderation in this regard. This is where the hidden human labor comes in, with this kind of work labeled by the New York Times as "janitor work" (Irani, 36). This labor is what calibrates algorithms to culture, what enables algorithms and other products to interact with the ever moving, ever complex nature of human culture (Irani, 63). This kind of cultural data work is quite demanding, and often companies have to rely on intense, monotonous labor from their own workforce, and often outside their own workforce. This resource where companies can utilize cheap labor in order to fine tune their AI is usually found in crowdwork, with the most prevalent of these being Amazon's Mechanical Turk (AMT), which we will discuss later on in this piece. For now we will focus on labor within the companies themselves that work to fine tune and maintain the illusion of an AI that can respond to things that an AI normally could not respond to, in this case, cultural data.

How would an algorithm know what constitutes hate speech? What constitutes depravity or decency? Such questions are often unanswerable by artificial intelligence, and are often deferred to human workers. One such example of human workers assisting AI in this regard is found in China. In this example, we see low wage workers applying labels to footage, in this case, surveillance and TV footage (Yuan). Due to weak privacy laws within China, the government, and hence AI startups within China, have more access to a myriad of data, giving the workers there a larger pool to train their algorithms. One such example of a company that utilized AI that was propped up by low wage laborers is the company called AInnovation. The AI company maintained an automated cashier system for a Chinese bakery chain, and the technology allowed users to put their pastries under a scanner and pay for it without the need of a human cashier. Nearly a third of the time, this system had difficulties with its computer vision aspects, for example having trouble discerning between different types of buns and different types of pastries. Taggers (low wage workers whose job it is to 'tag' and label photos and data in order to help AI with pattern matching), utilizing photos from the store's interior, managed to get the accuracy of the algorithm up to 99 percent (Yuan). This would have not been possible without the usage of taggers, and without them, the algorithm would have stayed at an inferior accuracy. This shows that AI (with the AI in question performing on its own, but at an inferior level) has to be taught by a combination of data factories and their workers (Yuan). To quote a poignant line from the article: "I used to think the machines are geniuses … Now I know we're the reason for their genius." and "all the artificial intelligence is built on human labor." (Yuan)

The above example highlights a case where a company utilizes taggers in order to better their algorithms (and hence, better their artificial intelligence, or better their product). Because of this, the second quote above is bolstered, as the algorithm showcased above was dependent on human labor in order to improve itself. But this was a case of a product having its accuracy improved in the background. What about a product that is marketed as working, while hiding its full machinations in the background? It is here where we get to the example of tech companies putting on a front, that is, when they advertise their products as self-sufficient AI while utilizing human labor in the background. With the startup company X.ai, we have a case of an artificial intelligence that isn't wholly artificial, as the AI in question was "ways from being able to operate on its own" (Bloomberg). Here we have a case where an AI company (in this case a startup) utilizes the labor of its own workers in order to push a false image of their product, in this case that their product is self-sufficient and capable of working on its own. An X.ai spokesperson said that "trainers review the 'vast majority' of information in emails so the system can improve" (Bloomberg). For employees of the company, the work that they had to do entailed a series of mote tasks (which was the case for the example of 24-year-old Willlie Calvin) by highlighting the bot's correct and incorrect responses, which was done, in Willie's case, for up to 12 hours a day (Bloomberg). The 21 employees (as described by an employee at X.ai) for some days had to work from 7am until 9:30pm and often stayed long past the end of their shift, since the service was offered 24/7 (Bloomberg). This classification and error fixing did eventually pay off, as the AI, given the name "Amy" by the former trainers at X.ai, did gradually improve and learn. The issue here arises that despite the fact that it took time before X.ai's AI could eventually learn and improve, it was not advertised as such. It was advertised to be in working and ready condition. The AI needed human calibration and support was not something that the company, X.ai, told to the public and to their investors. Another example of this was shown in Facebook's M (a personal assistant and chatbot), which was a personal assistant that existed inside Facebook Messenger that had its AI's responses reviewed, edited and sent out by a team of a few dozen contractors working from the company's campus. Facebook declined to say what hours these contractors worked and how often they issue corrections for M's guesses (Bloomberg). It was revealed from this, though it is not common knowledge that the majority of content moderation is outsourced to contract workers, often operating outside the United States. The exact manner of M's deception is programmed quite intelligently. When M's algorithms fail to figure out what the users want, it says it doesn't understand the question, all the while falling back on an invisible human to respond to the request (Simonite). The goal of Facebook's M was to test a new type of learning software called a memory network. Facebook, in the long term, hopes to use deep learning to help M process language. However, the version that they put forward for M's inception was one that relied on human labor (Simonite). Therefore, we have another example of a company putting a supposed "self sufficient" product to market, while propping it up with hidden labor in the background. The dearth of knowledge about this kind of practice illustrates tech companies' hesitancy for transparency of their products (Gray and Suri). But why is this? Why do companies want to maintain this venir that hides the hidden labor behind their product, whether the labor is from their own company, or outsourced?

And it is in this investigation of these acts, and in answering this question that we find a driving impetus for the withholding of information about the human labor that goes behind artificial intelligence. AI startups and their products in this field (such as X.ai, Clara, Operator, Mezi, etc) have raised more than $50 million in venture capital funding in the two years before 2016, which highlights the fact that this nascent industry is on the rise, and that there is a lot of revenue to be earned at stake (Bloomberg). Therefore, companies have a financial incentive to keep silent about the hidden labor that goes behind their products. These forms of labor that powers information software are obscured by the perpetual marketing claims of both the technology that surrounds this type of labor and the content that this labor allows to flow through (Downey, 146). This kind of false marketing, where companies hide the labor that goes behind their products due to financial reasons is what Crawford calls *Potemkin AI*: when tech companies demonstrate to investors what an automated system *would look like*, while relying on human labor in the background (Crawford, 65). This form of deception enables tech companies to claim a spot in the lucrative tech space (Crawford, 65). What better way to obtain funding for your product than to claim that it works as promised and to argue that it is self-sufficient? When there is a general ignorance over how the product (in this case, the artificial intelligence) gets maintained and corrected, companies will take advantage of this in order to achieve the funding they desire for their product, because in their eyes, more funding and more attention might mean that their product will eventually leave the human-supported inception stage.

This lack of transparency of tech companies is nothing new. There are notable cases of companies violating privacy in order to better their algorithms, without telling their users. This is done to better calibrate their products (and at times it is human labor that does this), but at the expense of the users as this often happens unbeknownst to them. For example, Facebook allowed outside developers access to user's data, though this stopped in 2015 (WSJ). In another case, a San Jose-based company called Edison Software had artificial intelligence engineers go through the personal emails of hundreds of users in order to improve a "smart replies" feature of their product. The company, in keeping with the tradition of lack of transparency of tech companies regarding their AI, did not tell their users that humans would view their emails in its privacy policy (Solon). As we can see, the lack of transparency that surrounds tech companies and their human-supported AI not only harms the workers (by giving them no credit for their labors), but it also harms the consumers and users as well, as they might either be fooled by the illusion of AI, or have their own personal data used to train an AI that they don't know needs improvement.

We've seen instances of companies using their own employees (for example with the case of X.ai's employees taking over 12 hour shifts in order to maintain their product [Solon]), but in my research I've discovered a practice that is far too common among tech companies looking to train their AI: the utilization of outside contract workers.These workers are often preferred over internal workers because they are paid less, are more plentiful and can be taken advantage of if the need arises. One such place where companies can acquire such cheap hidden labor in order to train their algorithms is found in Amazon's Mechanical Turk (AMT) (Gray and Suri). Crawford describes the mechanical turk illusion by giving the example that gave the company and label its name, a chess playing machine that was claimed to be able to play chess on its own when in fact inside of it there was a hidden chess master (Crawford, 67). Jeff Bezos, the CEO of Amazon described that AMT would serve as a technology that would promote "humans-as-a-service" (Irani, 225). AMT was born out of the failures of artificial intelligence to meet the needs of AI companies seeking to expand the domain of data they could store, classify and serve up online (Irani, 225). Here we have another instance of human laborers being used to AI failure to classify cultural nuances of sights, sounds and texts that filled the web (Irani, 225). On its website, AMT is advertised as a service where people can go to make more money in their spare time, therefore the technology is being advertised as a kind of side-hustle or side job (AMT Website). The website also advertises the kinds of work that an AMT worker will be doing "on their spare time" such as data processing (editing and transcribing content, translation, rating the accuracy of algorithms), data verification and cleanup and information gathering (AMT Website). So, the website advertises its services as a means for people to earn money in their spare time by doing monotonous yet useful tasks. AMT is thus advertised and operates like an online marketplace infrastructure (Irani, 227). Therefore, questions remain, such as: who's working and doing these kinds of classification, and how are these invisible workers treated?

The tasks that the "employees" do are called Human Information Tasks (HITs for short), and these tasks allow employees to input responses and train data. For example, as Andy Newman describes from his own experience, he was shown a photo of what looked like a school board meeting and was told to rate it on a scale of 1 to 5 for 23 different qualities such as "patriotic", "elitist", and so on (Newman). Such a task took three minutes. Only four percent of "turkers" made more than the federal minimum wage (the United States' minimum wage), $7.25 an hour (Newman). Employers specify the range of data for this kind of processing, and Amazon sends workers' output directly to employers' IT systems without human intermediation (Irani, 227). Therefore, the entire work of outside contractors can be done without oversight, and such a lack of accountability can subject the employers to possible abuse by the employers. Employers first filter what kind of worker they desire for the task at hand, often filtering with criteria such as an approval rating, the percentage of tasks the worker has performed and the worker's country. This filter approach allows employers, without individualized evaluation and selection, to request work from thousands of temporary workers in a matter of hours (Irani, 227). Once a worker submits work, an employer can decide whether or not to pay for it. Thus, an employer can reject work if they deem it to be insufficient, but it can also potentially lead to wage theft. This is the case since AMT's participation agreement grants full intellectual property rights over submissions regardless of rejection, and workers have no legal recourse against employers who reject work and use it anyway (Irani, 227). Therefore, through AMT, companies can abuse labor that they wouldn't ordinarily need to do within their own companies. In Newman's article, he brings up an example of turkers complaining on message boards about fighting over missing a 60-cent payment or getting a rejection of fifty cents (Newman). Because of such a policy, there is no guarantee that if a turker complains to an employer that they will get the fruit of their labor back. Amazon does not require requesters (employers) to respond to contact from a turker (Irani, 228). Thus, these hidden laborers can be abused by their recruiters, with no expected compensation if a wage theft does happen.

The question remains, how many people work for AMT? To what extent is the prevalence of this kind of work? It is predicted by economists that by 2033, due to tech innovation, 30% of today's full time occupations can be turned into augmented services to be completed "on demand" through a mix of automation and human labor (Gray and Suri). Here we have a case of AI, in this case the work that supports and maintains AI, to fundamentally redefine the kind of work humans do. But in the present moment, we must investigate who uses AMT. Though AMT is offered worldwide, in recent years, Amazon appears to have cut off international workers and focus more on US workers (Irani, 228). This is due to the fact that US workers are understood by the company to produce less "spam" work and to better understand the cultural nuances and better support the algorithms in this regard (Irani, 228). A study found that (at this specific time) 50% of AMT workers came from the US and 40% came from India, with workers from India relying on AMT as a primary means of income and workers from the US relying on AMT as a secondary means of income (Ipeirotis). For the more US-based findings, it was found that turkers are younger with 54% of them being between 21-35 years old compared to 22% of the general population, mainly female with 70% compared to 50% of the general population, and come from smaller families with 55% of the turkers not having children, compared with 40% of the general population (Ipeirotis). An interesting find for this study was that over 30% of AMT workers had a bachelors degree, and around 15% had a master's degree or above (Ipeirotis). Therefore it can be said that these workers are not mostly undereducated workers. Despite the education level of the majority of turkers, the weekly income of turkers was ranged from $1-$5 per week (over 35%) to $5-$10 per week (over 20%) (Ipeirotis). Another study found that (from the log data of 2,676 workers performing 3.8 million tasks) the median hourly wage was around two dollars an hour, and it was bolstered that around 4% of turkers make federal minimum wage (Hara). From these studies and observations of treatment of workers above, we can see that the outsource workforce of hidden labor that props up and maintains artificial intelligence in many cases is underpaid and undervalued. With the workers and their work hidden, programmers (regular programmers of tech companies who work inside the company, for example inside the company campus) can treat workers and their labor like bits of code and continue to think of themselves as builders of content, as opposed to managers of hidden labor, which is often what they are in actuality (Irani, 36). The work of these hidden laborers, especially AMT workers, is not appreciated by their employers or by their "fellow" programmers who they work in conjunction with. But why is this work hidden? Because despite the fact that this work requires highly complex skills and hard work, these skills utilized are widely available among the American people, and thus such "turkers" are replaceable because of it (Irani, 36). These workers are also kept at a distance because they pose a threat to the idealized concept of a programmer, and to the idealized concept of a tech company. These workers are kept out of sight because they threaten the idea that "computer scientists make the world a better place" (Irani, 36). This line of thinking illustrates a failure on the part of computer scientists to fulfill the "venture hyped dreams of grand technical achievement" (Irani, 36).

The usage, disrespect and lack of transparency around hidden laborers that shape AI highlights a key problem about the tech industry, the mythologization of computer scientists and tech companies. Utilizing hidden labor to create their products, tech companies can present a false view of AI, often just to get venture-capital money to fund a project that doesn't work as intended as of the time they acquire that money. This venir that is maintained by hidden labor also mythologizes tech companies as means of "magical" innovation, and this mythologizing obscures the means by which the products of these companies are actually created. In order to better educate the general public about these practices, these practices and the tech industry themselves, have to be demystified and de-masked. Tech companies currently utilize a practice which takes advantage of workers, both outsourced and internal. This practice serves as a kind of deception for the general public, which is made unaware of what's happening behind the scenes of AI, and is instead offered a narrative where AI and tech companies' products are advertised as "magical", ignoring the actual work that goes behind such "magic". The general lack of transparency and accountability can lead to the exploitation not only of the hidden laborers, but for users as well, as both parties can be exploited to better an artificial intelligence that, although it is claimed to be, is not ready for the market. In order to foster a more productive technological sector, these practices must be brought to light and made to be seen by the general public. In order to better serve the public, such practices must be demystified. The potemkin-AI must be torn down (not in a destructive way, but in an elucidating way) and the slight of human hand behind the "magic" must be revealed.

Works Cited

Crawford, Kate. *Atlas of AI*. Yale University Press, 2022.

Hara, Kotaro, et al. “A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk.” *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, https://doi.org/10.1145/3173574.3174023.

Huet, Ellen. “The Humans Hiding behind the Chatbots.” *Bloomberg.com*, Bloomberg, 18 Apr. 2016, https://www.bloomberg.com/news/articles/2016-04-18/the-humans-hiding-behind-the-chatbots.

Gray, Mary. “The Humans Working behind the AI Curtain.” *Harvard Business Review*, 23 Mar. 2017, https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain.

Ipeirotis, Panagiotis G. “Demographics of Mechanical Turk.” *SSRN*, 6 Apr. 2010, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1585030.

Irani, Lilly. “Difference and Dependence among Digital Workers: The Case of Amazon Mechanical Turk.” *South Atlantic Quarterly*, vol. 114, no. 1, 2015, pp. 225–234., https://doi.org/10.1215/00382876-2831665.

Irani, Lilly, Suri Siddarth. “The Hidden Faces of Automation.” *XRDS: Crossroads, The ACM Magazine for Students*, vol. 23, no. 2, 2016, pp. 34–37., https://doi.org/10.1145/3014390.

"Make Money in Your Spare Time Get Paid for Completing Simple Tasks." *Amazon Mechanical Turk*, https://www.mturk.com/worker.

Newman, Andy. "I Found Work on an Amazon Website. I Made 97 Cents an Hour." *The New York Times*, The New York Times, 15 Nov. 2019, https://www.nytimes.com/interactive/2019/11/15/nyregion/amazon-mechanical-turk.html.

Press, Jeff Chiu/Associated. "Tech's 'Dirty Secret': The App Developers Sifting through Your Gmail." *The Wall Street Journal*, Dow Jones & Company, 3 July 2018, https://www.wsj.com/articles/techs-dirty-secret-the-app-developers-sifting-through-your-gmail-1530544442.

Clarke's Third Law on UFO's | Science. https://www.science.org/doi/10.1126/science.159.3812.255.c.

Simonite, Tom. "Facebook's Perfect, Impossible Chatbot." *MIT Technology Review*, MIT Technology Review, 2 Apr. 2020, https://www.technologyreview.com/2017/04/14/152563/facebooks-perfect-impossible-chatbot/.

Solon , Olivia. "The Rise of 'Pseudo-AI': How Tech Firms Quietly Use Humans to Do Bots' Work." *The Guardian*, Guardian News and Media, 6 July 2018, https://www.theguardian.com/technology/2018/jul/06/artificial-intelligence-ai-humans-bots-tech-companies.

Yuan, Li. "How Cheap Labor Drives China's A.I. Ambitions." *New York Times*, 25 Nov. 2018.