

A Hegelian-Marxist Analysis of Artificial Intelligence

Cameron More

1. Introduction

In the last few years, there has been an explosion of usage and discussion of emerging Artificial Intelligence (AI) technologies. With the rise of ChatGPT, the development of an entirely new industry of AI-powered products, and the fear of Artificial General Intelligence (AGI),¹ it is imperative that the fundamental construction and mechanisms of AI be understood. Misunderstandings of what AI is, how it works, and the mysterious “black box” have led to all sorts of confusion over how AI impacts, and will continue to impact, modern society: AI should be treated like other machines, and the power of modern Large Language Models should not overpower our ability to think in concrete terms about the material, economic, and theoretical questions of how to handle such impressive tools, what role they play, and who benefits from their usage.

In order to demystify modern machine learning (ML), I show how Large Language Models (LLMs) work, from their basic elements of matrix multiplication to natural language tokenization. I show that what appears to be a potentially sentient disembodied force is really the product of labor relations and disparity between technology corporations in the Global North and data annotators in the Global South. Machine Learning is revealed not to be a paradigm shift, but only another in a long history of labor objectification (Hegel) and mechanization (Marx).

I begin by explain how machine learning works and the fact that linear algebra underlies everything an LLM can do. I then discuss the political economy behind LLMs and the data annotators behind the model. I show that LLMs are not a new, unique, paradigm shifting technology, but have the same effects as other instances of mechanization, contrary to the belief of modern Accelerationists. I finally show that the hallucinations produced by AI demonstrate a Hegelian-Marxist truth about machines: that they objectify human processes and knowledge, and when used for the purposes of capital, alienate those whose data the model was trained from.

2. Technical Brief

1 Neural Network

Neural networks are types of algorithms. An algorithm is any set of instructions which can be executed by some machine. Equations like $\mathbf{y} = \mathbf{mx} + \mathbf{b}$ are algorithms when executed by a calculator. Algorithms, most essentially, turn an input into an output.²

¹ Barry Smith and Jobst Landgrebe, *Why Machines Will Never Rule the World: Artificial Intelligence without Fear* (New York: Routledge, 2023).

² Justin Joque, *Revolutionary Mathematics: Artificial Intelligence, Statistics and the Logic of Capitalism* (London: Verso, 2022), 78.

The “neural” in neural networks is an analogy to the way neurons operate in the brain. Nodes, neurons, send signals to each other (connected via “edges”) in a decentralized manner – every node is connected to a limited amount of other nodes which have their own set of limited connections. Input enters the node, is changed (or not) in some way by that node, and exits through an edge to another node. “Training” refers to changing the relationships between the node – which edges it has.³

The analogy between computational neurons and biological ones is not accidental; the mathematical foundations of neural networks were laid out in a 1943 paper by Warren McCulloch and Walter Pitts.⁴ The neurons in their model took binary inputs and yielded binary outputs depending on how many inputs each neuron received from its edges; if the neuron received enough 1’s, then it passed along a 1 to other neurons. In 1958, Frank Rosenblatt radicalized this model by moving from binary operations to statistical ones. Rather than work through inputs and outputs of 1’s and 0’s, Rosenblatt’s neuron manipulated the inputs statistically, giving statistical outputs to other neurons. The biggest advantage of this statistical model is it can be tuned in a finer grain way.⁵

It wasn’t until the 1980s that another breakthrough occurs in the training of neural networks with the discovery that neural networks could be trained using gradient descent. This means that the modification of the values of the neurons are gradual, slow, and responsive to the output error.⁶ In the researchers’ own words, “the procedure repeatedly adjusts the weights of the connections in the network so as to measure the difference between the actual output vector of the net and the desired output vector.”⁷ When training a neural network, or any algorithm, data is split into training and testing batches. The training batch is sent through the network and its outputs (the actual output vector) are plotted against the testing data (the desired output vector). For now, suffice it to say that vectors are simply values.

Deep learning refers to neural networks that have many entire layers of neural networks feeding each other. While a simple neural network may have just one set of neurons, a deep neural network has as many hidden layers as desired. These hidden layers, and the mathematical relationships between them, are referred to as the “black box” of machine learning.

³ Bernhard Mehlig, *Machine learning with neural networks*, 27, Oct., 2021: <https://arxiv.org/abs/1901.05639>, 1.


⁴ Warren McCulloch and Walter Pitts, “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *Bulletin of Mathematical Biology* 5: 115-133.

⁵ Frank Rosenblatt, *Report No. 85-460-1: The Perceptron: A Perceiving and Recognizing Automaton* (Buffalo, NY: Cornell Aeronautical Laboratory, Inc, 1957), 14. Available at: <https://blogs.umass.edu/brain-wars/files/2016/03/rosenblatt-1957.pdf>.


⁶ Mehlig, 3.

⁷ David Rumelhart, Geoffrey Hinton, and Ronald Williams, “Learning representations by back-propogating errors,” *Nature* 323, no. 9 (October, 1986): 533-536, 533.

Simple Neural Network



Deep Learning Neural Network



● Input Layer ● Hidden Layer ● Output Layer

8

“Explainable AI” is machine learning which can be *fully* explicitly explained in mathematical terms: the weights of every neuron can be given and the relationship between every neuron in the network is calculable. In deep learning networks, with a “black box” of many hidden layers, the explainability drops because the calculations become exponentially large. While “the math” is out there – it does exist in some computer – it is not comprehensible to a human mind. We do know what is going on– nodes pass signals back and forth, forward-propagating and backward-propagating – we cannot comprehend the exact mathematical calculations in their entirety.

II Tokenization and Training

If AI is just numerical computation, how does the information turn from words, images, or sounds into numbers and vectors which are mathematical? Stepping back from how the network operates, I will explain how input data is vectorized to explain how a new technique, the Transformer, has propelled AI into its current form.

Algorithms, as mentioned, turn inputs into outputs. When dealing with numbers, the relationship between the inputs and outputs is quite obvious: some combination of operations (addition, subtraction, multiplication, linear algebra) acted upon the input to produce the output. When dealing with non-numerical inputs like words and images, some level of encoding must take place. Encoding is the changing of some non-numeric input into a numerically palpable item which can be manipulated. ML models begin by encoding the input, like the sentence “the dog wags its tail,” into some smaller more relevant chunks of information. In large language models, this happens by *tokenization*. Sentences and words are broken into tokens – “the” “dog” “wags” “its” “tail.” Some tokenizers break down sentences into words, though most break them into sub-word portions, by syllable, by punctuation (including spaces and apostrophes), or other codes. The exact choice of tokenization method depends on the desired output of the model.

Basically, to train a large language model means to find the statistical relationship between the tokens. GPT-3 uses a tokenizer roughly equivalent to 4 characters of text in English; the

⁸ <https://thedata scientist.com/what-deep-learning-is-and-isnt/>


sentence “the dog wags its tail” is tokenized as **The** **dog** **wags** **its** **tail**.⁹ The tokens are given numerical ids – in this case [1169, 3290, 266, 663, 7894] – which correspond to a list of every possible token in the English language.

The neural network is then trained on the statistical relationship between the tokens – which occur next to each other, which come before or after others. Slowly, using a gradient, the network adjusts the values of the vectors (the nodes) to more closely match with what the output should be against what the model actually produces.

Neural networks begin with arbitrary values for the nodes in the network – totally random values of vectors – which become less random as tokenized items are passed through and the output is compared to the expected output (training-testing split of the data). Networks “adjust themselves” based on how much error is calculated after it produces an output. Suppose that some token input of 1 should give an output of 5, and the model gave an output of 3. The model will adjust the nodes that that token passed through to make its output 5 instead of 3. Every model has some loss, and no model will be trained perfectly, so maybe only 4.6 or 4.8 is achieved by the corrected model.

Loss functions are built into models which impact how much and how quickly the model “learns” from its loss. If a model was off by 2 (in our case above), we don’t want it to overcorrect and yield an output of 7 or 8 (if the ceiling is 10), and given that a model is trained in many cycles (called epochs), we want to track the loss value. We want the model to correct itself just enough to move in the right direction without overcorrecting.

III Transformers and Large Language Models



10

Recursive neural networks (RNN) learn token by token. The model takes some input, adjusts it, gives an output, checks how accurate it was, and adjusts accordingly before moving on to the next token. The “recurrence” of the RNN is the step-by-step self-modification of the model to its calculated loss. In the image above, the model takes some input \mathbf{X} , gives an output \mathbf{h} , and learns from the loss to adjust the next input-output cycle.

⁹ Accessed here: <https://platform.openai.com/tokenizer> on September 30, 2023.

¹⁰ Accessed here: <https://www.datacamp.com/tutorial/tutorial-for-recurrent-neural-network> on September 30, 2023.

However, because they adjust themselves so often, RNNs have poor “long term memory.” Another more concerning problem is that of vanishing and exploding gradients. As the model adjusts itself, it balances the many layers of networks that have varying levels of loss. To say a gradient is vanishing is to say that some layers are not adjusting themselves as much as they should (the gradient of adjustment is small, and becomes smaller every iteration). To say a gradient is exploding is to say too much adjusting is occurring.

In 2017, a group of AI researchers at Google published a pre-print paper entitled “Attention Is All You Need” where they presented a novel architecture for neural networks called the Transformer.¹¹ The main two structures in this architecture were an encoder and decoder. The Transformer’s encoder not only tokenized parts of words but built matrices which also described the relationship of that token to all the other tokens in a given sequence. Not only did words become single numbers like 425, but they became matrices, such as”:

$$\begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix}$$

Each number in the matrix represents that token’s relationship to another token. By cleverly multiplying the right matrices together, a much more intelligent neural network can be created because *more information* is gathered from the data (not just the immediate relationship of one word [token] to another but the relationship of every word to every other word). By following a specific series of encodings, decodings, and normalizations, the Transformer can far more accurately and efficiently capture the truth of the training data.

The original Transformer in the 2017 paper was built for translation tasks, trained on a massive German-English dataset. Within 7 months, another breakthrough took place with the unveiling of a technique called Generative Pre-Training, GPT.¹² Developed by OpenAI, Generative Pre-Training differs from the approach taken by the Google team in that it uses vast amounts of unstructured, raw text, which is fine tuned on more clear and higher quality data. What is meant by “higher quality data” is data which has a calculable loss. This is like having a list of multiple-choice questions in one column and the answers in a second column. This file is split – often in a 70-30 fashion – to give the model examples of correct answers against which it can compare itself to. The initial unsupervised training on raw data is analogous to loading up the model with a working vocabulary and only giving a little bit of training on loss accuracy.

The generative in GPT refers to the ability to generate text. As opposed to the sequence-to-sequence task of translation one sequence of English tokens into a sequence of German tokens, encoding the tokens, running them through the neural networks and decoding them, the GPT just produces new tokens.

The decision to train a model on vast amounts of unstructured text poses problems. When models are trained on highly specified and controlled data, which is split into a training and testing batch and provides a rubric against which the model compares itself, the outputs are more accurate. Unstructured language in itself does not have a goal in the sense that it can be split into a training and testing batch, so already the model is training itself just to follow basic spelling and

¹¹ Ashish Vaswani et al., “Attention Is All You Need,” 12 June, 2017: <https://arxiv.org/pdf/1706.03762v1>.

¹² Alec Radford et al., “Improving Language Understanding by Generative Pre-Training,” 11 June, 2018, accessible here: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

grammatical conventions. Only in the fine-tuning stage does the model start to *do* things, rather than aimlessly generate new tokens (words).

There is far more unstructured text than structured data available to train models on. The human element now comes back into view when we stop to consider where all this data is actually coming from.

IV On Hallucinations

One of the most vexing problems for AI adopters is its hallucinations.

In 2013, workers at a German construction company noticed something odd about their Xerox photocopier: when they made a copy of the floor plan of a house, the copy differed from the original in a subtle but significant way. In the original floor plan, each of the house's three rooms was accompanied by a rectangle specifying its area: the rooms were 14.13, 21.11, and 17.42 square metres, respectively. However, in the photocopy, all three rooms were labelled as being 14.13 square metres in size. The company contacted the computer scientist David Kriesel to investigate this seemingly inconceivable result. They needed a computer scientist because a modern Xerox photocopier doesn't use the physical xerographic process popularized in the nineteen-sixties. Instead, it scans the document digitally, and then prints the resulting image file. Combine that with the fact that virtually every digital image file is compressed to save space, and a solution to the mystery begins to suggest itself.¹³

This is the exact same process that creates hallucinations. Through the process of statistical embedding and tokenization, the actual semantic meaning (the information) of the words used to train models like ChatGPT is lost. Hallucinations are just new text being generated with enough statistical weight that they *sound* correct.¹⁴

ChatGPT is a statistical representation of the English language, fine-tuned to complete certain tasks, nothing more. "It" knows nothing about the *meaning* of the words it produces because it knows nothing about the *meaning* of the inputs, just the tokenization of words and the relationship between the tokens. "It" is a long polynomial equation. Therefore, hallucinations are not bugs and flaws, but fundamental features of how the model "learns" in the first place. To reduce the statistical learning the model does in an effort to reduce the statistical hallucination would also reduce the amount of leaning in the first place!

These ML models are crystallizations of past artifacts – past speech, text, images, audio – which are statistically mobilized for present use.

In AI that performs analytical tasks, looking for statistical correlations, far fewer hallucinations happen because the model is not trying to generate new information. The Generative in GPT is causing the statistical missteps. Researchers have been taking steps to lessen the hallucinations, including feeding it better raw data, and especially using machine-readable

¹³ Ted Chiang, "CHATGPT IS A BLURRY JPEG OF THE WEB," *The New Yorker* 9 February, 2023, available at: <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>.

¹⁴ Lei Huang et al., "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions," *arXiv* pre-print, 9, November, 2023: <https://arxiv.org/abs/2311.05232>.

knowledge models to prevent obvious hallucinations.¹⁵ Microsoft’s Bing AI tries to solve the problem by using internet access to search the web for answers, rather than only rely on the pretraining web data.¹⁶

3. The Political Economy of AI

I Excursus on Data

What even is data? Data is information about some aspect of the world captured by an artifact – paper, computer binary, stone tablet, etc. When we speak of data for machine learning, we refer to large amounts of unstructured data, like this paper, free text, and structured data often in a spreadsheet format (Excel, csv, tsv). High quality data is data which has been turned from an unstructured to a structured form. For example, a question-answer dialogue between two people in a pdf is unstructured, but the same dialogue in an excel format, where one column corresponds to the question and the other the answer, is structured.

In the last decade, with the rise of Big Data, scientists and capitalists alike have sacrificed a notion of testable truth for statistical truth. *If the data says so, it must be true*, the refrain goes. With the rise of computer-readable data, scientific analyses are increasingly *meta*-analyses, which take the data and findings of individual experiments and surveys and perform statistical operations on them to make discoveries. On first glance, this may not seem like a problem – more data about something should give us more truth about it, right? As the saying goes, however, “there are three kinds of lies: lies, damned lies, and statistics.”

In the mid-2010s, a crisis was discovered in the life sciences called the Replication Crisis. “More than 70% of researchers have tried and failed to reproduce another scientist’s experiments, and more than half have failed to reproduce their own experiments,” an article in *Nature* claimed in 2016.¹⁷ Many researchers, it has been found, engage, knowingly or not, in “p-hacking,” a practice of selectively reporting the data they produce in their experiments. Suppose that a new drug is being tested in 10 separate clinical trials. In only one of those trials, the drug was able to boost the participant’s immunity to a particular infection. A researcher might p-hack their results by only publishing that one trial’s results. Even though the drug was ineffective in over 90% of participants in reality, it was effective in 5%, which is itself the margin of error! By putting the unwanted data in the file drawer, never to be seen, the drug’s efficacy went from a margin of error to a breakthrough.¹⁸

Modern science and machine learning relies on statistics, but even the exact relationship between statistics and probability itself is not a settled topic.¹⁹

¹⁵ Suriya Gunasekar, “Textbooks Are All You Need,” *arxiv*, 20 June, 2023: <https://arxiv.org/pdf/2306.11644.pdf>. And Konstantinos Andriopoulos and Johan Pouwelse, “Augmenting LLMs with Knowledge,” *Arxiv pre-print*, 28 September 2023: <https://arxiv.org/pdf/2309.16459.pdf>.

¹⁶ Kevin Caramacion, “News Verifiers Showdown: A Comparative Performance Evaluation of ChatGPT 3.5, ChatGPT 4.0, Bing AI, and Bard in News Fact-Checking,” *Arxiv*, 18 June, 2023: <https://arxiv.org/ftp/arxiv/papers/2306/2306.17176.pdf>.

¹⁷ Monya Baker, “Is There A Reproducibility Crisis?” *Nature* 533 (2016): 452-454, 452.

¹⁸ For more, see Megan Head et al., “The Extent and Consequences of P-Hacking in Science,” *PLoS Biology* 13, no. 3 (2015): <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002106>.

¹⁹ Joque, 102-103.

Since data is information about the world captured in some artifact, those artifacts must have been produced somewhere, at some time, by some person (via machine). Recently, OpenAI was hit with a lawsuit over violation of European privacy law. The complaint states that “it appears that OpenAI systematically ignores the provisions of [European law] regarding the processing of data for the purposes of training models within ChatGPT.”²⁰ From copywrite infringement to privacy violation, LLM companies have been tangled in legal issues, which has pushed companies to pay for custom curation of datasets.

II The Materiality of LLMs: Data Workers and the Global South

It is commonly known that lots of technological advancements and devices rely on materials extracted from the Global South. Although containing an estimated \$24 trillion in untapped resources, countries such as Nigeria and the Democratic Republic of Congo are “mired” in poverty and violence.²¹ For workers in China who assemble the finished technological product, the situation is just as concerning, where workers as young as 14 have been employed to manufacture iPhones.²² It is also the case that the Global South is a site of extraction for Big Data.

Large, carefully curated datasets do not fall from the sky. Really, “the magic of machine learning is the grind of data labeling,” Phil Jones accurately observes in his 2021, *Work Without the Worker*.²³ Even if ML models are able to annotate datasets, those models themselves must have been trained on some data to begin with; and when introducing new tasks to models, like sorting hate speech from non-hate speech, initial high-quality datasets need to be used in model training. These tasks, called “human intelligence tasks,” are primarily done by workers in the Global South.

“In the same hour it takes Amazon CEO Jeff Bezos to make \$13 million, a refugee earns mere cents teaching his algorithm to spot a car.”²⁴ Sama (formerly Samasource) trains workers in Kenya to annotate data for Meta (formerly Facebook), mostly regarding hate and explicit content. OpenAI also contracted with Sama for the same purpose – to make ChatGPT less hateful and toxic – paying Sama employed Kenyan workers less than \$2 an hour.²⁵ In 2022, a lawsuit was filed against Sama and Meta for exposing workers to such graphic content for multiple hours a day. “It is completely inaccurate,” Sama responded, “to suggest that Sama employees were hired under false pretenses or were provided inaccurate information regarding content moderation work.”²⁶

²⁰ Emma Woollacott, “OpenAI Hit With New Lawsuit Over ChatGPT Training Data,” *Forbes*, September 1, 2023, available at: <https://www.forbes.com/sites/emmawoollacott/2023/09/01/openai-hit-with-new-lawsuit-over-chatgpt-training-data/?sh=4607498a6d84>.

²¹ Katie Brigham, “How conflict minerals make it into our phones,” *CNBC* 15 February, 2023, available at: <https://www.cbc.com/2023/02/15/how-conflict-minerals-make-it-into-our-phones.html#:~:text=Tin%2C%20tantalum%2C%20tungsten%20and%20gold,the%20Democratic%20Republic%20of%20Congo>.

²² Reuters Staff, “Foxconn says underage workers used in China plant,” *Reuters* 16 October, 2012, available at: <https://www.reuters.com/article/us-foxconn-teenagers/foxconn-says-underage-workers-used-in-china-plant-idINBRE89F1U620121017>.

²³ Phil Jones, *Work Without the Worker: Labour in the Age of Platform Capitalism* (London: Verso, 2021), 2.

²⁴ Jones, 11.

²⁵ Billy Perrigo, “Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic,” *Time* 18 January, 2023, available: <https://time.com/6247678/openai-chatgpt-kenya-workers/>.

²⁶ Annie Njanja, “Meta and Sama face legal action in Kenya for alleged poor working conditions,” *TechCrunch* 31, March, 2022, available here: <https://techcrunch.com/2022/03/30/meta-and-sama-face-legal-action-in-kenya-for-alleged-poor-work-conditions/>.

In India, the situation is almost as desperate as Kenya; the Amazon Mechanical Turk and Clickworker platforms rule the data annotation market.²⁷ On the one hand, it is true that workers have access to decent paying jobs, but on the other hand, their labor is undoubtedly being exploited by the companies paying them in the first place – recall how much Jeff Bezos makes per hour. The decent rate that these workers are paid, however, can easily plummet; because lots of the work is job-based, workers often find themselves working beyond the expected time it takes to annotate some particular dataset, drastically lowering their actual rate per hour.

Tech workers themselves are unaware of the conditions of the workers that annotate the datasets they work with on an everyday basis. An Amazon review for Jones’s book reads, “I’ve worked in tech for a decade and thought I was pretty in tune with how advanced AI and ML was... Turns out, most of these advanced systems are backed by millions of underpaid and underemployed workers.”²⁸ This “uneven geographic development” is nothing new in the history of global economics.

AI researchers have begun asking about the ethics of paying for workers to curate these datasets.²⁹ But to question the ethics of this particular line of work is to ask about the ethics of capitalism generally, and its (neo)colonialist regime. It is a feature of capitalism that “centres exploit peripheries, metropolises exploit hinterlands, the first world subjugates and mercifully exploits the third,” writes Marxist geographer David Harvey.³⁰ Harvey argues that the uneven distribution of capital, land, and labor cause uneven spaces of exploitation and contradiction, such that some sites become subjugated to others. Just as capitalists within a single city exploit the proletariat of that city, entire countries exploit the labor of other countries. Capital sees the world as a planetary mine, ready to extract material from the earth and value from the workers.³¹

Artificial Intelligence seems like it saves time, but only because that time was accumulated in the data used to train the model in the first place. The “time” stored in the data is both the time spent to originally produce the content – that it took decades to put together the Wikipedia and Web data used to train ChatGPT – and the time paid for to clean up massive datasets.

III The Circuit of ML Capital

I put forward two Marxist theses: first, that the mobilization of data for the purpose of training GPT models is analogous to the commodification of natural resources, where humans come upon objects which were not built for but nevertheless have economic value; and second, that the data produced by data annotators is a paradigmatic example of commodity production.

My characterization of data accords with existing accounts of data capitalism. Jathan Sadowski, for example, describes how data can serve both as a commodity and as capital itself. On the one hand, data is a commodity when it is bought and sold in the sphere of exchange, but it

²⁷ Rina Chandran et al., “FEATURE-AI boom is dream and nightmare for workers in Global South,” *Reuters* 14 March, 2023, available at <https://www.reuters.com/article/global-tech-workers-idAFL5N2XI2X8>.

²⁸ Amazon Review, available at <https://www.amazon.com/Work-Without-Worker-Platform-Capitalism/dp/1839760435>.

²⁹ Nana Nwachukwu and Jennafer Roberts, “The Glamorisation of Unpaid Labor: AI and its Influencers,” 16 September 2023: <https://arxiv.org/pdf/2308.02399.pdf>

³⁰ David Harvey, *The Limits to Capital* (Verso: London, 2018 [1982]), 439.

³¹ Martín Arboleda, *Planetary Mine: Territories of Extraction under Late Capitalism* (London: Verso, 2020).

becomes constant capital when used in the production process – in this case, training models. “On the one hand, data is cast as a digital raw material ... on the other hand, data is cast as a commodity produced by digital labour of people posting on Facebook, clicking on Google, exercising with Fitbits,” and consciously annotating datasets.³²

Let us examine what the circuits of data looks like from these two perspectives. From the point of view of the manufacturers of data (like Sama), some initial portion of money (**M**) buys commodities (labor power **L** and means of production *mp*) which produce annotated datasets (**C'**) which are sold to companies like OpenAI and Meta for an amount of money (**M'**) greater than initially invested.

$$\mathbf{M} - \mathbf{C} \left\langle \begin{matrix} \mathbf{L} \\ \mathit{mp} \end{matrix} \right\rangle \mathbf{P} \dots \mathbf{C}' - \mathbf{M}'$$

In this circuit, those datasets act as finished commodities (**C'**), the output of the production process (**P**). In the circuit of AI companies like OpenAI, the data serves as part of the means of production (*mp*) used to produce (train) their models. OpenAI pays Sama some initial portion of money (**M**) for the datasets (*mp*) to produce the trained GPT (**C'**) which is then sold as a service to users at a rate of \$20 per month and large corporations for a high cost, like Microsoft, Slack, Coca-Cola and many others.

As a reminder, the circuit for workers (**C-M-C**) yields no surplus. Workers begin with a portion of labor they will expend throughout the day (**C**), which is sold for a wage (**M**), which allows them to buy food, shelter, and other things required for them to be able to come to work the next day, ready to expend another portion of labor (**C**).

The initial commodification of data focused on its use for *advertising*. Indeed, much of the data that is collected by the websites we browse every day is still primarily used for that purpose. Companies that pay for AI services, like Amazon, primarily use those services for increasing their ability to advertise better. Companies do not arbitrarily make money by spying on you; they do so by *selling commodities* (this point is often lost in non-Marxist discussions of internet ethics). Even when companies make money by collecting data, it is only when that data is sold to someone that it actually has value. The people who buy data use it as constant capital in the production of advertising or general business services.

Some Marxist and Leftist analyses of modern political economy have hypothesized that today’s global economy operates in such a radically different way than Marx’s time that his critique of production and commodity circulation are outdated. Maxime Ouellet, for example, has argued that “under cognitive capitalism, the labor theory of value ... gives way to knowledge or affects as the main source of value.”³³ Marx never had a labor theory of value; he had a value theory of labor – a theory of how value is regulated and “created” by social processes of both production

³² Jathan Sadowski, “When data is capital: Datafication, accumulation, and extraction,” *Big Data & Society* January-June (2019): 1-12.

³³ Maxime Ouellet, “Revisiting Marx’s Value Theory: Elements of a Critical Theory of Immaterial Labor in Information Capitalism,” *The Information Society* 31, no. 1 (2015): 20-27.

and exchange.³⁴ Further, the idea that knowledge/cognitive workers do immaterial work is nonsense, since they do, in fact, produce artifacts – spreadsheets, business presentations, perform information-services, etc. It is true that the nature of how they interact with their work is different than in Marx’s analysis – their cognitive powers are overtaken as a site of production, as opposed to the factory assembly line; however, premises of Marx’s analysis have not been disproven.³⁵

Ouellet goes on to say that “the fact remains that the increasingly obsolete nature of labor under informational capitalism has now made it objectively possible to contemplate the bases for a post-capitalist society.”³⁶ This is a drastically misinformed view because it *forgets the material reality of information*. Knowledge (or cognitive) workers are paid to perform services and produce artifacts, and anyone who has worked in one of these jobs knows that.

IV A Critique of Accelerationism

A more concerning trend in Leftist political thought has been the theory of Accelerationism. Revolutions happen when contradictions sharpen, Accelerationists argue, namely the contradiction between the forces of production and the mode of production. Marx writes in the 1859 *Contribution to the Critique of the Political Economy*,

In the social production which men carry on they enter into definite relations that are indispensable and independent of their will; these relations of production correspond to a definite stage of development of their material powers of production. The sum total of these relations of production constitutes the economic structure of society – the real foundation, on which rise legal and political superstructures and to which correspond definite forms of social consciousness.

This is a classic distinction of the middle-Marx – dividing society into an economic and political sphere, where the economic determines the political. By “relations of production,” Marx means the actual conditions for the relationships people have with each other. I have a relationship to my coworkers because we sell our labor to the same person and perform labor at the same time. I have a relation to the barista only because I am purchasing a commodity, coffee. This is what Marx means when he says that economic relations of production (and consumption) determine other spheres of social and political consciousness. He goes on:

The mode of production in material life determines the general character of the social, political and spiritual processes of life. It is not the consciousness of men that determines their existence, but, on the contrary, their social existence determines their consciousness. At a certain stage of their development, the material forces of production in society come in conflict with the existing relations of production, or – what is but a legal expression for the same thing – with the property

³⁴ For a classic investigation of the value form, see Samezo Kuruma, *Marx’s Theory of the Genesis of Money: How, Why, and Through What is a Commodity Money?* Tr. Michael Schauerte (Leiden: Brill, 2017 [1957 and 1979]), and for a recent discussion see Kiyoshi Nagatani, “The Value Form and the Fetishism of the Commodity,” *Review of Radical Economics* 0, no. 0, available at <https://doi.org/10.1177/04866134231197451>.

³⁵ Franco Berardi calls this semio-capitalism. For an excellent discussion of the cognitive shift of capital, see *Futurability: The Age of Impotence and the Horizon of Possibility* (London: Verso, 2019) and *After the Future* (Edinburgh: AK Press, 2011).

³⁶ Ouellet, 26.

relations within which they had been at work before. From forms of development of the forces of production these relations turn into fetters. Then comes the period of social revolution. With the change of the economic foundation the entire immense superstructure is more or less rapidly transformed.

The core motor of change for Marx in this passage is the development of economic relations, which inevitably change social ones. For an example, we can turn to the introduction of metal tools, particularly axes, to a tribe of Australian Aboriginals, the Yir Yoront. Stone axes were crucial for the material production of the Yir Yoront, a paleolithic culture, most importantly to cut firewood. Although women and children could use axes, two activities with them were restricted to men: gathering wild honey (an important “prized” food for the tribe), and crafting “secret paraphernalia.”³⁷ Playing an important material role, axes also generally played an important cultural role: symbolizing masculinity, serving as a totem for a clan, and being an item around which particular families identified themselves.

In the 1930s, many changes were introduced to the Yir Yoront people as they made more and more contact with white Westerners, one being the rapid increase of steel axes. A missionary, who lived with the Yir Yoront at that time, would gift axes to women if they won his favor, and this situation of a woman possessing an axe was, at times, highly disruptive. Since women and children no longer had to ask permission to borrow the axe from the men – which they would eagerly give – all kinds of kinship dynamics changed. The distribution of axes was simultaneously a distribution of ideological power.

Accelerationists claim that artificial intelligence, automated technological and mechanical systems, represent a development of the forces of production so much that the ideological social superstructure of the political capitalist regime must soon crumble. The way to get to the revolutionary society is to accelerate the contradiction between the economic base and political superstructure so that the superstructure has no choice but to bend to the will of the stewards of the base, the proletariat.

As the forces of production, machinery, Marx writes in the *Grundrisse*, is “the appropriation of living labor by objectified labor,” human shoveling by machine (object) shoveling.³⁸ The development of machinery yields a change not only in the distribution of knowledge – which moves from workers’ minds to objects, but also a change in the workers’ labor – no longer do they perform those now mechanistic tasks. The Accelerationist thesis is that modern development of machines has changed how much living labor is required, that we can automate away most of our daily work, eliminate the profit motive and surplus extraction, and foster a post-capitalist lifestyle.

The main problem with this argument is that it forgets the material production of those technologies themselves. Admittedly, with the development of machines which can make machines, we are able to reduce the amount of required labor drastically. But, the claim that artificial intelligence represents an end to cognitive work, and to labor altogether as theorists like Ouellet claim, is false.

³⁷ Lauriston Sharp, “Steel Axes for Stone-Age Australians,” *Human Organization* 11, no. 2 (Summer, 1952): 17-22, 18.

³⁸ Karl Marx, *Grundrisse*, tr. Martin Nicolaus (London: Penguin Books, 1993), 693.

4. The Ontology of AI

I From Objectification...

Recently, a group of researchers at the Harvard Business School released a working paper containing results of an experiment that measured the productivity of consultants who used AI compared to those who did not.³⁹ Out of 758 participants, workers with below average productivity had their productivity improved by 43%, and above average workers were boosted 17%.⁴⁰ The trend in LLM labor adoption is not necessarily that workers are being replaced (though that does happen), but generally workers simply become more productive. The time spent replying to emails is drastically reduced – whether a worker prompts ChatGPT to help them, or uses one of the many LLM-enabled auto-reply applications that already exist. LLM Agents, like secretaries, data analysts, and travel agents, have become a new booming business.⁴¹ One group of AI researchers even created a virtual world inhabited by 25 Generative AI agents, called Smallville.⁴²

Large Language Models are technical artifacts, tools, which, like all tools have been in the history of human material production, save time. LLMs like ChatGPT have a unique ability to generate and produce content – text, or computer code, or images. But, being artifacts, LLMs are part of the “self-propelling life of the dead,” as Hegel would say.⁴³ LLMs, because they are trained on vast amounts of input data, are congealed dead labor. What was once words for informative or entertainment purposes has been commodified (through statistics) by the Transformer architecture.

For Hegel, products of production processes are objectifications of human species-being; human abilities and knowledge are transformed into objects through the process of work. Another way of saying this is that as work is realized, made real, knowledge is brought into an objective standing.

In his famous exposition of the Master-Slave dialectic, Hegel shows that although the Slave is subject to the desires of the Master, the Slave finds freedom through their ability to objectify their knowledge. The power to turn knowledge into artifacts grants the Slave freedom, but only freedom in the mind, since, in reality, they are materially subjugated to the Master.⁴⁴ The Slave, however, becomes alienated by the fact that what they produce is owned, used, and directed by the Master, an external force.

³⁹ Fabrizio Dell’Acqua et al., “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” *Harvard Business School*, 22 September, 2023, available at: https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf.

⁴⁰ Dell’Acqua et al., 2.

⁴¹ Zhiheng Xi et al., “The Rise of Large Language Model Based Agents: A Survey,” *ArXiv*, 19 September 2023: <https://arxiv.org/pdf/2309.07864.pdf>.

⁴² Joon Sung Park et al., “Generative Agents: Interactive Simulacra of Human Behavior,” *ArXiv* 6 August, 2023: <https://arxiv.org/pdf/2304.03442.pdf>.

⁴³ Hegel, quotes in Georg Lukacs, *The Young Hegel: Studies in the Relations between Dialectics and Economics* tr. Rodney Livingstone (London: Merlin Press, 1975), 333.

⁴⁴ “Through work, the bondsman [Slave] becomes conscious of what he truly is [a consciousness with certain powers].... Through this rediscovery of himself by himself, the bondsman realizes that it is precisely in his work wherein he seemed to have only an alienated existence that he acquires a mind of his own.” Hegel, *The Phenomenology of Spirit* tr. A. V. Miller (Oxford: Oxford University Press, 1977), 118-119.

II...to Alienation

In Hegelian alienation, Spirit no longer recognizes itself in the thing which it has produced.⁴⁵ This is because, Marx argues, capitalist relations of production have separated man from the fruits of his labors in the capitalist's ownership of private property. Because the products that the laborer produces are always already owned by the capitalist, even before they exist, they only ever appear to the worker as alien things – crystallized bits of alienation in material form.

For Hegel, the objectification of labor in an artifact which is used, enjoyed, by someone else is a manifestation of alienation. For Marx, however, alienation occurs only when such objectified artifacts are separated by the laborer in the wage relation. In simpler terms, the Hegelian laborer objectifies their knowledge in some produced artifact which is given to others to freely use; the Marxist laborer is alienated from their knowledge because it is bought by a capitalist in the form of a wage. The laborer, for Marx, is always already cut off from the products of their labor because that labor has itself been bought by someone else.⁴⁶

Marx writes that in bourgeois society, the laborer is alienated in (at least) four ways, notes scholar Marcello Musto;

(1) by the product of his labor, which becomes “an alien object that has power over him”; (2) in his working activity, which he perceives as “directed against himself,” as if it “does not belong to him”; (3) by “man’s species-being,” which is transformed into “a being alien to him”; and (4) by other human beings, and in relation “to their labor and the object of the labor.”⁴⁷

Machinery, Marx writes in the *Grundrisse*, confronts the worker as an alien thing, a calcified objectified piece of knowledge which was once in the mind of another worker, now an independent material entity.⁴⁸ In a sense, machinery is a privatization and commodification of what Marx called the general intellect – the free social commons of scientific and social discourse. Although particular inventions and could be commodified, never did Marx think that the entire commons could itself be privatized. As Slavoj Žižek as pointed out, this is exactly the situation with Facebook, Microsoft, and X (formerly Twitter): our commons of social discourse is being rented to us.⁴⁹ In order to participate, we must buy smartphones or computers, and we must pay rent back to these companies for access: rent in the form of literal money payments or rent in the form of *information*, our personal data.

⁴⁵ Louis Dupre, “Hegel’s Concept of Alienation and Marx’s Reinterpretation,” *Hegel-Studien* 7 (1972): 217-236, 220.

⁴⁶ “Marx ... distinguishes sharply between objectification ... in work in general and the alienation of subject and object in the *capitalist form* of work,” Lukacs, 552.

⁴⁷ Marcello Musto, “For Karl Marx, Alienation Was Central to Understanding Capitalism,” *Jacobin* 18, December, 2021, available at: <https://jacobin.com/2021/12/alienation-reification-commodity-fetishism-philosophy-capitalism>.

⁴⁸ Marx, *Grundrisse*, 695.

⁴⁹ Slavoj Žižek, *The Relevance of the Communist Manifesto* (Cambridge: Polity Press, 2019), 5-7.

5. Conclusion

Artificial Intelligence is not a mysterious, mystical black box – far from it. AI is a complex social phenomenon, only achievable because of mass industrialization and global labor disparity. By reading this phenomenon through Marxist theory, we are able to understand the concrete social reality that the fetishized sheen of AI hides. Machines, at their best, aim at replicating human behavior – and AI is only a more effective type of this generic paradigm of machinery. Without an understanding of how a machine or machine learning model functions, how much data it needs, who is annotating that data, theorists may not understand that the world is deeper than the end-user’s computer screen.

Author

Cameron More is an independent researcher based in Buffalo, NY. He graduated with a master’s degree in philosophy from the State University of New York at Buffalo and works professionally as an applied ontologist. His work focuses on Sigmund Freud, German Idealism, Martin Heidegger, and Continental Philosophy generally.



cultural logic

marxist theory & practice



Since 1997

ISSN 1097-3087

<https://ojs.library.ubc.ca/index.php/clogic/index>

Editors

Rich Gibson
San Diego State University

Joseph G. Ramsey
University of Massachusetts Boston

Martha Gimenez
University of Colorado (retired)

E. Wayne Ross
University of British Columbia

Gregory Meyerson
North Carolina A&T State University

David Siar
Winston-Salem State University (retired)

Advisory Group

Robert Albritton, *York University, Canada*
Theodore W. Allen, *Brooklyn Public Library (1919-2005)*
Crystal Bartolovich, *Syracuse University*
Rosalyn Baxandall, *State University of New York, Old Westbury*
Vladimir Bilenkin, *North Carolina State University*
Jamie Owen Daniel, *Director of Organizing, UPI Local 4100, IFT/AFT/AFL-CIO*
Anthony Dawahare, *California State University, Northridge*
Greg Dawes, *North Carolina State University*
Teresa L. Ebert, *State University of New York, Albany*
Barbara Foley, *Rutgers University*
Marcial González, *University of California, Berkeley*
Rosemary Hennessy, *State University of New York, Albany*
Mike Hill, *State University of New York, Albany*
Peter Hitchcock, *Graduate Center-City University of New York*
Amitava Kumar, *Vassar College*
Neil Larsen, *University of California, Davis*

John McClendon, *Michigan State University*
Bill V. Mullen, *University of Texas, San Antonio*
Richard Ohmann, *Wesleyan University*
Bertell Ollman, *New York University*
Vijay Prashad, *Trinity College*
Russell Reissing, *University of Toledo*
Modhumita Roy, *Tufts University*
E. San Juan, Jr., *Philippines Cultural Study Center*
Sean Sayers, *University of Kent*
Helena Sheehan, *Dublin City University*
David Slavin, *Westfield State College & Emory University*
Paul Smith, *George Mason University*
Alan Spector, *Purdue University Calumet*
Michael Sprinker, *State University of New York, Stony Brook (1950-1999)*
Carol A. Stabile, *University of Wisconsin-Milwaukee*
Richard Stahler-Sholk, *Eastern Michigan University*
Kenneth Surin, *Duke University*
Jeffrey Williams, *Carnegie Mellon University*
Richard Wolff, *University of Massachusetts, Amherst*
Michael Yates, *University of Pittsburgh, Johnstown*