Generative AI: A Turning Point for Labor's Share?

Thanks to artificial intelligence, labor's share of national income may no longer hold steady.

Lukasz A. Drozd
Economic Advisor and Economist
FEDERAL RESERVE BANK OF PHILADELPHIA

Marina M. Tavares
Economist
THE INTERNATIONAL MONETARY FUND

The views expressed in this article are not necessarily those of the Federal Reserve and do not necessarily represent the views of the IMF, its executive board, or IMF management.

After years of slow and steady development, generative artificial intelligence (AI) technologies have exploded in popularity, and many experts believe that we are entering a new, AI-driven phase of the Industrial Revolution.[1] The advent of AI as the new engine of growth raises questions about the future of labor. Some have expressed concerns that, in the short run and the medium run, AI may lead to employment losses brought about by task automation and the skill obsolescence of the current labor force. But an additional risk of AI is in the long run: Unlike previous technologies, AI may undermine labor's share of national income, and technological innovation could, for the first time, permanently reduce the importance of labor in the economy, even if full employment is maintained.[2] The labor share—or the share of national income that labor receives as wages, salaries, and other compensation—has remained remarkably stable throughout much of the Industrial Revolution. Its permanent or even persistent decline would thus constitute a major break from past trends (Figure 1). This is troubling because, historically, automation, when combined with slower labor income growth relative to capital income, has been associated with rising income inequality and social unrest.[3]

In this article, we discuss how technology affects labor's share of income, why the labor share has been stable for so long, and why AI may threaten that stability.[4] The balanced nature of technological progress—which sometimes leads to the automation of tasks previously done by humans, but at other times improves the productivity of already-automated tasks—might have been the key to sustaining labor's stable share of income. If this is so, AI's unprecedented potential for job automation in the coming decades may disturb this long-standing equilibrium and cause the labor share to decline.[5] A declining labor share does not mean that labor income growth will stall, but to fuel labor income growth, AI will have to prove sufficiently productive in jobs it displaces us from.

AI's Pivotal Moment

Economic historians attribute much industrial growth to a handful of transformative innovations known as general purpose technologies (GPTs). A GPT has four key traits. First, it's identifiable as a generic technology or organizational system. Second, it's widely adopted across the economy. Third, it has multiple distinct applications. And fourth, it generates numerous spillovers that catalyze secondary innovations or even spawn new GPTs.[6] According to one analysis, as of the mid-2000s there had been at least 24 GPTs, ranging from language and the wheel to the steam engine and the computer.[7]

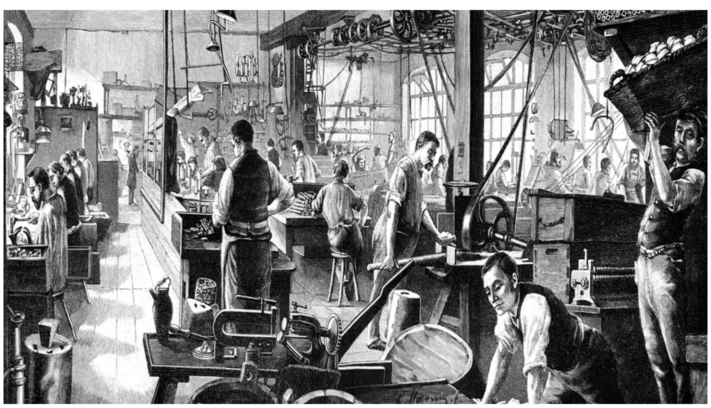

Each of these groundbreaking advances set the stage for numerous secondary innovations and drove economic growth for decades. For instance, the Industrial Revolution was kicked off by the steam engine, which revolutionized industries such as textiles, cotton agriculture, material science, and transportation. Then, in the last quarter of the 19th century, electrification and the internal combustion engine brought mass production, improved mobility, and telecommunication, leading to the Second Industrial Revolution. And beginning in the mid-20th century, the transistor gave rise to digital technologies, including computers and computer software, sparking the Third Industrial Revolution. According to many experts, AI, thanks to major advances in generative AI in the last decade, will soon join the pantheon of GPTs. Some experts even expect AI to become the leading technology of the Fourth Industrial Revolution, and for decades to come.[8]

See ***What Is Generative AI?***

The first three industrial revolutions featured very different technologies, but these revolutions' aggregate growth patterns look remarkably similar. This consistency led economist Nicholas Kaldor to formulate six stylized facts of industrial growth, known as Kaldor's facts, and it led to the development of a unifying (neoclassical) growth theory that makes no distinction between the different technological eras that drove growth.

According to one of Kaldor's facts, the incomes of capital and labor grew at similar rates even though most productivity gains could be attributed to the introduction or improvement of capital in production.[9] Economists recognized this feature as early as the 1940s, and it means that, in the long run, capital-productivity-augmenting technological progress equitably benefited both labor and the owners of capital despite ongoing automation.[10] In other words, labor somehow held its own against capital in the battle over which side would get the bigger piece of the economic pie, even though most economic growth was attributed to new or improved machinery (that is, to capital-productivity-augmenting technological progress).

If history is any guide, then, the emergence of a new GPT that drives growth need not affect growth patterns—in particular, labor's share of income.

So, what could make AI different from, say, the spinning jenny, which revolutionized the textile industry in the 19th century, or Ford's assembly line? The concerning aspect of AI, as we see it, is that it is a major GPT with the potential to broadly and persistently tilt the incoming flow of new capital-productivity-augmenting innovations toward those that automate tasks, rather than augment the productivity of capital in previously automated tasks. According to our theory, the balanced nature of technological progress in the economy as a whole—which sometimes leads to the automation of tasks previously done by humans but at other times improves the productivity of already-automated tasks—was the key to sustaining labor's stable share of income. If that's the case, AI may disturb this long-standing equilibrium and cause the labor share to decline over the coming decades.

See ***AI's Automation Potential***

Forces Affecting the Labor Share

To see how AI could affect the labor share, we must first understand how any new technology that makes machines more productive affects labor's share of income in an economy. Let's imagine the economy as a simplified system that produces a generic output—say, "a unit of gross domestic product (GDP)"—by performing a list of basic tasks.[11] The tasks are so basic, they can be executed by either machines or people. Each task needs to be completed only once, and no task can be skipped. (We discuss below how these simplifications affect the analysis.)

Because each production task can be executed by either capital (a machine) or labor (a worker), a new technology that increases capital productivity can influence a task in only one of two ways: If the task is already automated (that is, already done by a machine), it makes the machine that does that task more productive and hence cheaper to employ; but if the task is not yet automated (that is, done by a human) and the new technology makes it cost-effective to automate that task, the machine will displace a worker from that task. The impact on labor's share of income depends on which kind of technology we're talking about, and hence which kind of innovation is brought about by technological progress, which, as we shall see, turns out to be a crucial distinction for understanding how AI enters the picture.[12]

Productivity Effect of New Technologies

Consider the first case, where a technology boosts the productivity of capital in a task that is already automated. Let's assume that producing a unit of output costs a firm $100, and, due to competition, the unit of output also sells for $100. Now let's assume that a single task costs a firm $2 when performed by a machine, but, thanks to the new technology, the firm uses a next-generation machine to accomplish the same task for less, say $1.

Initially, the firm reduces its costs by $1 per unit of output, which immediately boosts the firm's profits. Since profits contribute to capital income but not wages, labor's income remains unchanged. Capital owners lose $1 in income payments for the machine, but they gain $1 in profits, and so there is no change in capital income either. Capital owners can pocket the gain because nothing else has changed, and we presume that firms pay workers just enough to keep them employed.

But if that technology becomes broadly available, other firms will adopt it as well, and all firms will seek to expand production because they can make (excess) profits. Competition between firms to expand their market share will erode profits, driving down the price of goods to a new break-even point of $99, and capital income payments will decrease by the saved dollar on the new machine. As that happens, labor can reap *all* the benefits from the new technology, for two reasons.

First, as the price of goods in the economy falls, both workers and the owners of capital pay less for the goods they purchase, leaving them more money to purchase more goods, which is to say that their real income—income adjusted for purchasing power—rises. Can labor's *share* increase because of the falling prices of goods implied by declining profit margins? Yes, and how much it increases depends on labor's (and capital's) initial share of income.

For example, if labor's initial share is two-thirds, as historically has been roughly the case, the increased purchasing power gives labor a 1 percentage point gain on two-thirds of each dollar generated via production and captured as labor income, which amounts to a total *real* income gain of two-thirds of a dollar in constant purchasing power (that is, in terms of goods one can purchase). These gains are earned on each unit of output produced. And labor's share of income is equal to the number of units of output it can purchase using the income it earned on each unit of output it helped produce. Thus, labor's increased purchasing power implies that labor's share of income has increased by two-thirds of a percentage point.[13] Simply put, to purchase goods, labor must effectively pay for its own input into production and for capital (including profits). Since capital costs less, labor can purchase more goods for the income it earns, and so its share rises at the expense of capital.

Remarkably, the benefits that accrue to labor don't end there. Since the new technology leads to a broad-based decline in prices, machines, which are also produced goods, should cost less. If the price of a machine, like the price of the good it helps produce, also falls by 1 percent, the additional cost savings associated with this decline will further drive down the price of goods, continuing the cycle of declining prices of goods and machines. This adjustment process won't stop until labor captures *all* the benefits from the new technology, with labor's share rising by as

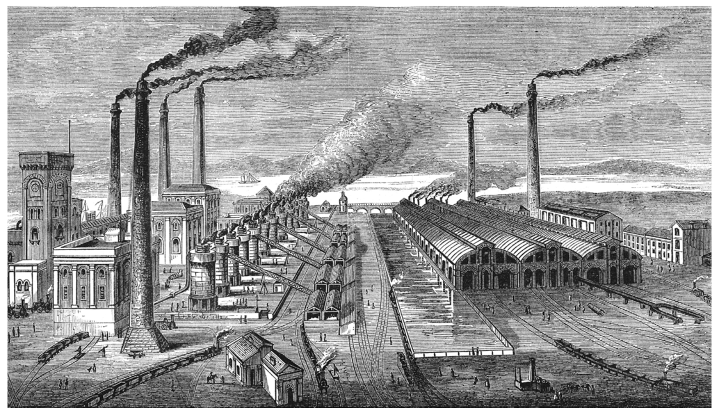

much as 1 percentage point and capital's share falling by a full percentage point.

This *productivity effect of capital deepening* increases labor's share of income; it applies when innovations make machines more productive or cheaper to employ in tasks that are already automated. Although this is a simplified example, the productivity effect arises even when we allow for limited substitutability between tasks. Of course, if one task can effectively and perfectly substitute for some other tasks, capital will displace labor from these other tasks via task substitution, and this will dampen the productivity effect. However, broad substitutability across tasks is implausible. Technological processes for producing goods and services are fairly rigid and comprise a precisely defined sequence of myriad distinct tasks. For example, just because an automotive plant can replace a human welder with a robot doesn't mean it can do away with the workers who install windshields.[14]

Displacement Effect of New Technologies

But consider a very different scenario: For a given task, a new technology makes a machine only marginally more productive than labor. Let's assume that labor earns $2 per unit of output for this task, and the new machine performs the same task at nearly the same cost—say, $1.99. That is, the arrival of a new technology (and thus a new machine) leads to the automation of the task in question, but the difference of 1 cent between the cost of labor and capital implies that it does *not* save the adopt-

FIGURE 1

Labor's Share of Income Held Steady

This was true even during an era of widespread automation.

Labor's share of national income for the UK, the U.S., and France, 1855–2017

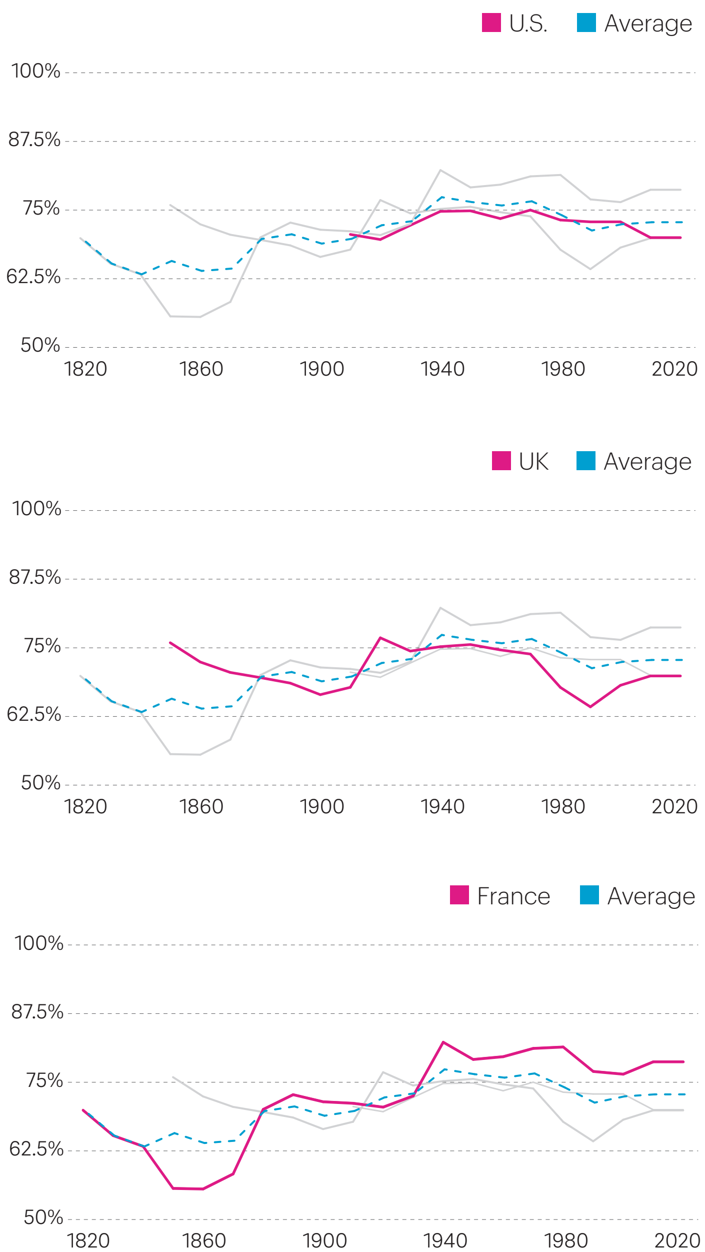

Data Sources: For all countries, 1855–2020: World Inequality Database (https://wid.world/wid-world) and Distributional National Account (Table S.A1), as detailed in Piketty, Saez, and Zucman (2018). For France, 1820–1850: Piketty (2014).

Note: Labor's share of income is defined as the ratio of compensation of employees and 70 percent of mixed income over the factor price national income. The data for this figure come from raw tax data. Labor's share of national income for all countries except the U.S., which is labor's share of personal income.

ing firm much money.

Since profits increase only marginally after automation, the price of output remains $100, and for this reason the final effect could not be more different. Payments to capital increase by almost $2 per unit of output ($1.99 to be precise); labor, by being automated out of the task, loses $2 in displaced income. Consequently, labor's share of income drops by nearly 2 percentage points, and capital's share rises by the same amount.

This *displacement effect of capital-productivity-augmenting technologies* counters the productivity effect, reducing the labor share. More-realistic situations feature both effects, but we can still decompose their net effect into the two elementary effects. Say, for example, the cost of the machine falls to $1 rather than $1.99. We can equivalently think of all such cases in two steps: The cost of using a machine first falls to $1.99, triggering the *displacement effect* associated with automation, and it then drops further to $1, triggering the *productivity effect* within an already automated task. The combined effect reduces labor's share of income by about 1 percentage point: The displacement effect lowers it by about 2 percentage points but the productivity effect then raises it by 1 percentage point.[15]

The remarkable property of such mixed cases is that labor's share of income nonetheless *always* falls—although to a varying degree depending on the strength of the offsetting productivity effect. Consequently, technologies that trigger task automation *always* displace labor's share of income, and hence for the labor share to remain constant it must be offset by innovations that raise the productivity of machines used in already-automated tasks. To see this, note that even in the most extreme case where the cost of using the new machine drops to zero instead of $1.99—that is, capital becomes infinitely more productive in our task—labor's share of income still does not increase. Although it falls by 2 percentage points due to the displacement effect, it then rises by 2 percentage points, for a roughly nil net effect.

What the Past Tells Us About the Future

Given these contrasting effects, it's clear that labor's share of income can go up or down after the arrival of innovations that lead to the emergence of new, capital-productivity-enhancing technologies in the economy. This makes it so much more surprising that the labor share lacked any distinct trend over the last century. Our research attributes this remarkable phenomenon to the random and cumulative nature of innovations across a vast number of tasks in the economy, which, if true, has consequences for how we think about the effects of AI.[16]

Innovation may spring from directed research and development (R&D) in a specific area. But at the macroeconomic level the vast number of tasks in the economy, the incentive to invest in R&D in various areas, the random outcome of research, structural transformation that reshuffles the importance of individual tasks, and the varying impact of new innovations on individual tasks implies that capital productivity within individual tasks may well be a random process that exhibits fairly stable statistical properties across both tasks and time. As a result, sometimes an innovation enhances capital productivity in an already automated task by some random percentage of its value, and at other times it enhances productivity in a task yet to be automated, with the odds of R&D producing either type of innovation, the average size of (relative) productivity increments, and their variability all being stable. If this is so, then, according to our research, labor's share of income will stabilize despite ongoing automation because, as we note above, automation, which reduces labor's share of income, will eventually be balanced out by improvements in already-automated tasks, which raise labor's share of income.

To better understand this phenomenon, imagine it's the dawn of the Industrial Revolution: All tasks are performed by labor, as machines are almost nonexistent; all innovations affect tasks that are not yet automated, simply because very few tasks are done by machines. Initially, the labor share declines, but as more tasks become automated, subsequent innovations more and more frequently improve the productivity of already-automated tasks rather than automating tasks operated by labor. As a result, the odds of either innovation stabilizes, and our research shows that the economy enters a phase of "scale invariant" growth: Even though innovations increase the productivity of the whole economy, the share of tasks that need a lot of capital relative to labor, or need little capital relative to labor, stabilizes. As this happens, the labor share of income also stabilizes.

This mechanism's relevance to the economywide effects of R&D is not that surprising—it already explains many seemingly unrelated phenomena in the natural sciences. For example, an analogous mechanism is believed to govern population growth across cities, employment growth across firms, wealth across people, the number of species across genera, the number of citations of academic papers, and even the propagation of cracks in fracturing materials. These phenomena are unrelated, but their statistical properties are strikingly similar: They all exhibit the so-called "power law property," which in all these cases is similarly attributed to scale invariant growth or propagation.[17]

Why Labor's Share of Income May Decline

But if R&D leads to a new class of innovations that affect nonautomated tasks significantly more than automated tasks relative to the historic pattern, then the natural mechanism that stabilized labor's share will have changed, and labor's share of income may decline. Our research suggests that this is one key concern about AI's long-run effect.

Until recently, one of the most significant barriers to automation has been the need for cognition. This barrier prevented machines from operating autonomously or made it costly to automate many production processes. Cognitive abilities are a significant barrier because we humans are largely incapable of supplying algorithms that would enable machines to mimic our own cognitive abilities.[18] As an example, until the advent of AI, there was no algorithm for a cognitive ability as basic as the general computer-based visual recognition of objects. It's not difficult to see how this could keep machines from replacing humans.

If AI successfully and universally overcomes this hurdle, it will unleash a torrent of long-delayed automation-focused innovations, and for some time new innovations may be heavily biased toward those that enhance the productivity of machines in tasks that have not yet been automated as opposed to those that are already automated. As that happens, labor's share of income may permanently decline despite having been constant in the past. The size of the decline in labor's share will depend on the exact nature as well as the productivity of AI-driven innovations, and it will happen irrespective of whether displaced labor finds employment.

Depending on how productive AI is, however, labor's *income* may still grow even if its *share of income* declines. As an example, consider the case where the cost of employing a machine drops to zero in tasks previously operated by labor. In that scenario, labor's share of income stays constant because the displacement effect is offset by the productivity effect, as we have discussed. But because the economy has become more productive (that is, it can produce more output from the available resources), labor's overall income must have risen. (Here it is important that all resources, including labor, remain fully employed, or else labor income will not rise.)

What bodes well for growth in labor's income is what also makes AI so good at automating jobs. AI software is a nonrival technology, meaning that its use by one entity does not preclude another entity from using it at the same time. Thanks to this nonrival nature, big economies of scale will likely make AI highly productive in many areas. For example, even though it costs billions of dollars to develop self-driving cars, the necessary software can be deployed to every car once it becomes available, and its cost will likely decline as software providers compete for customers.

But even if AI proves productive enough to fuel labor income growth, the declining labor share is a concern. A declining share means that the total purchasing power of those who supply labor grows weaker than the total purchasing power of owners of capital, contributing to inequality between these groups. In many areas where the two groups compete for a fixed supply of resources, the relative purchasing power is all that matters, and a change in relative purchasing power can trigger social unrest due to the lower welfare experienced by one group. For example, land (in good locations) is largely in fixed supply, since it cannot be manufactured. As a result, the price of land depends on the average "bidding power" of those who seek to purchase it, and even if labor income continues to grow, housing may become unaffordable for those who supply labor. It is perhaps no coincidence that labor unrest coincided with a persistent but ultimately transient decline in labor's share of income in the 19th century.[19]

Conclusion

In the early 1980s, Nobel laureate Wassily Leontief famously compared the ongoing microcomputer and software revolution to the advent of steam, electric, and later internal combustion power in production and transportation. He questioned whether humans could avoid being replaced by computers much as horses had been displaced by mechanical power. This forecast didn't come to pass during his lifetime, but the question he posed has gained renewed relevance in the face of AI's rapid development.

AI's potential seems limitless, and perhaps one day AI will take all our jobs, in which case Leontief's forecast will come to pass. But we think that's unlikely in the foreseeable future and perhaps even impossible. We humans possess a wide range of skills and can perform many different tasks. If we maintain a cost advantage over machines in a significant fraction of these tasks, we should avoid the fate experienced by horses, at the very least for quite some time.

In summary, although some jobs may be lost in the short and medium run, we agree with many commentators who argue that AI, in the long run, will transform rather than displace jobs. But this conclusion alone misses an important point. If our theory is correct, the more pressing issue is AI's likely effect on labor's share of income. We need to carefully monitor and address the growing inequality as those jobs transformed by AI receive a smaller share of the economic pie.

What Is Generative AI?

Generative AI and AI in general refer to the set of machine learning methods that use artificial neural networks (ANNs) to obtain computer algorithms for tasks normally requiring human intelligence or computer software written by a human programmer (for example, the task of facial recognition). In other words, these tools write computer software on their own.

An ANN is a mathematical structure inspired by the human brain and emulated by a computer.[20] Like the human brain, an ANN consists of an interconnected network of nodes, called artificial neurons, each of which fires a new signal based on the sum of the signals it receives from the neurons it is connected to. In an ANN, these signals are numerical; in a brain, they're electrical.[21]

Take, for example, an ANN designed to distinguish between images of dogs and cats (Figure 2). This ANN's artificial neurons take grayscale pixel values as inputs. Subsequent neurons process these signals to arrive at two classification signals returned by the last two neurons, one for "dog" and one for "cat." The more confidence this network has in its classification, the greater the signal returned by the output neuron associated with that classification (as shown in the figure). Parameters of this network, which are mostly weights that each neuron assigns to each connection, encode the underlying algorithm for classifying images, and these parameters are obtained algorithmically by computer software in the process of training.

Training an ANN involves feeding it thousands of examples (labeled images in the depicted case) and algorithmically adjusting the network's internal parameters to minimize the average prediction error on that training data set. The key principle behind this algorithm is to iteratively adjust all the parameters of the network by small increments in the direction of the steepest descent in terms of the average prediction error on the training set. Just as water flowing downhill finds the lowest point of a surface, this algorithm can obtain an error-minimizing combination of the network's parameters.

Advanced AI architectures—such as the transformer architecture powering the latest AI models—involve much more complex structures and may involve a more complex training procedure, but the basic principle of operation and training is similar. What makes ANNs so successful is still a bit of a mystery, but it almost certainly has to do with the fact that this structure is conducive to identifying the patterns in the data that do well on both the training data and the new data (for example, learning to recognize a cat by its whiskers, long flexible tail, and pricky ears, so that the ANN can classify new images as well as those in the training set).

Generative AI refers to an approach that trains ANNs to create content. For instance, generative AI could learn from images of cats how to animate an image of a cat, or how to fill in the missing details of an image. A generative AI model of this sort acquires its "knowledge" of what a cat is by being trained to fill in and predict missing pixels based on neighboring pixels, across thousands of blanked-out images that contain images of cats. By doing this, the ANN powering this model, instead of merely learning how to recognize a cat, can learn how to recreate its image when prompted to do so.

But the most revolutionary leap in generative AI methods came from natural language processing, specifically through large language models (LLMs) like OpenAI's ChatGPT and its latest product, GPT-4. Generative LLMs are trained analogously, but by using text instead of pixels, with each word represented by a unique array of numbers. The ANNs powering these models have proven useful for learning about the meaning of words, language structure, and ultimately the human world at large. They do this by being trained to predict blanked-out words based on the surrounding text. LLMs are additionally trained to use their acquired "knowledge" to generate complete responses to human queries. This is done by feeding the LLM ratings for generated responses and teaching the model, which is already pretrained on text, how to sequentially add to the text those words that will maximize the rating for the sequentially generated answer.[22]

What makes the most advanced generation of large-scale LLMs revolutionary is their apparent ability to "reason" based on concepts. By doing so, these models showcase glimmers of an early form of artificial general intelligence, or AGI—the capacity to understand, learn, and execute a broad range of intellectual tasks, much as humans do.[23] This was once considered a distant goal, and the use of AI has been thus far limited to narrowly defined tasks, such as image classification. But more people today believe that generative AI will eventually pave the way for an advanced form of AGI. It's unlikely that AGI will soon match or exceed human intelligence, but even a limited AGI-like capability can transform the economy.

FIGURE 2

An ANN Decides: Dog or Cat?

The connections between neurons show how signals travel through the network, and hence how information is processed by it.

A photo is fed into the input

Singular pixels are viewed as grayscale values

Objects within the canvas are identified and reduced to outlines, shapes, and patterns

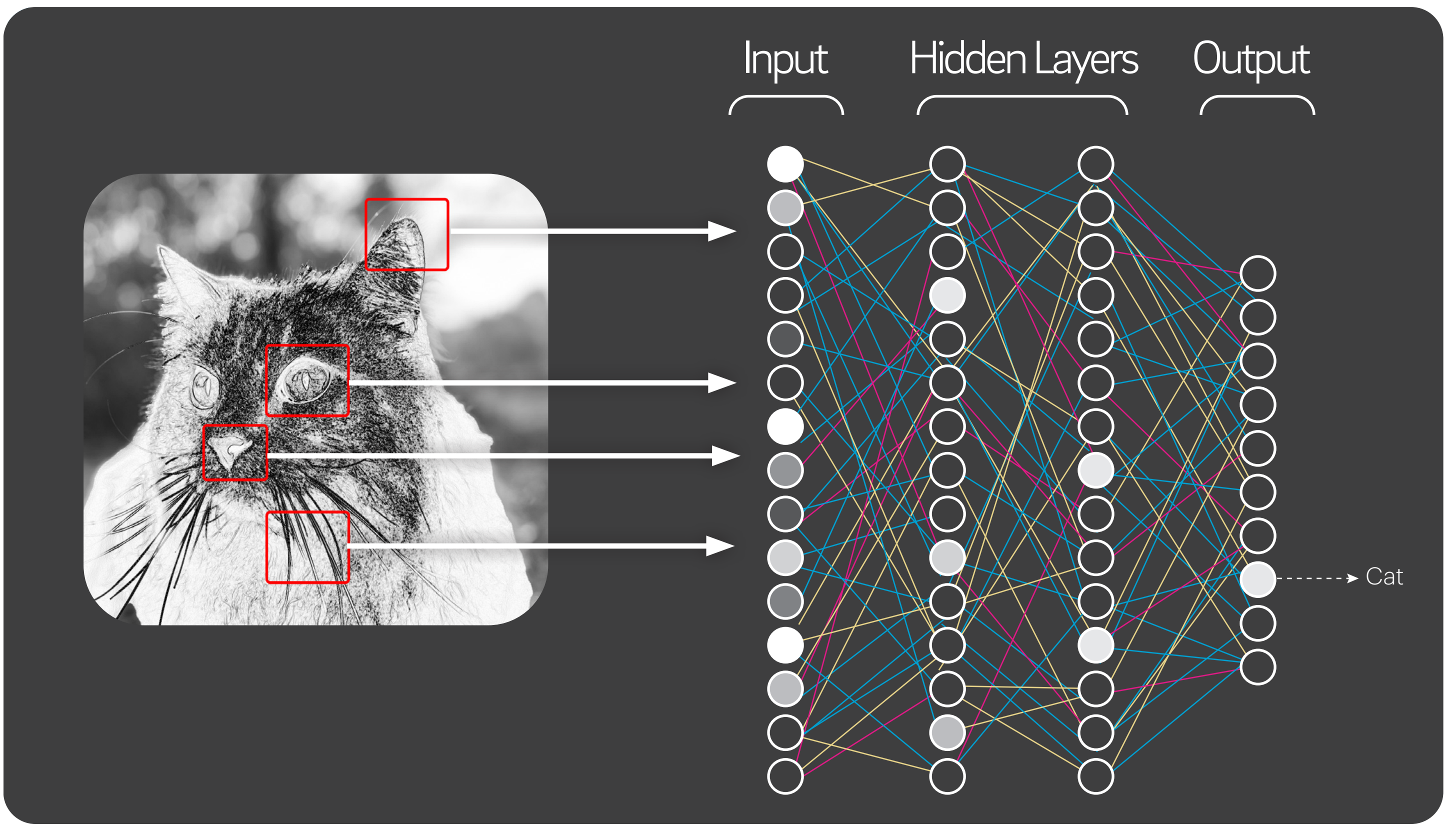

The shapes and patterns are observed and fed through the hidden layers and are compared to attributes learned in previous training, such as the eyes, ears, nose, and whiskers.

Each element is passed through the hidden layer processes and are weighted (or graded) and compiled into a summing value to provide a number that represents the likelihood of the object being either a dog or a cat.

AI’s Automation Potential

Several studies confirm AI’s immense automation potential. Given the important role of humanlike cognition in production—a role that AI research may offer a substitute for—it’s not surprising that AI may have such an effect.

According to one recent study, up to 19 percent of jobs could see at least half of their tasks automated at or near the current state of AI technology over the coming decades.[24] Overall, about a third of all tasks could be automated. According to another study, the explosive growth of generative AI tools may lead to the automation of up to 30 percent of labor hours by 2030. Generative AI increases these estimates by at least 8 percentage points, again at or near the current state of technology.[25] Higher-income jobs are more susceptible to automation—a departure from the past, when lower-income jobs were the primary target of automation.[26] Perhaps this will lessen AI’s impact on income inequality among workers, but the scale of these changes may exacerbate inequality between workers and the owners of capital.

Business people generally agree with these assessments. For example, 75 percent of companies recently surveyed by the World Economic Forum plan to use AI to automate jobs within the next five years, and they expect a reduction of approximately 20 percent of existing jobs, many of which are skilled jobs.[27] Casual evidence suggests that AI is also becoming a strategic consideration in hiring decisions across the economy. For example, IBM CEO Arvind Krishna announced that IBM may freeze hiring in all positions that are likely to be affected by generative AI. According to Krishna, if IBM follows through on the freeze, attrition will reduce IBM’s workforce by 30 percent in the next five years. This represents an estimated loss of about 7,800 jobs.[28]

These analyses do not consider whether it makes economic sense for firms to use AI, and firms may be underestimating the costs of employing AI. But the massive surge of investment in AI and the fact that AI software is a nonrival good—meaning that nothing prevents two companies from using AI software simultaneously—does not bode well for labor. Nonrival technologies can bring massive economies of scale, so even costly hurdles may eventually prove economical to overcome. And because the cost of the marginal provision of AI software to another customer is close to zero, it will be hard for labor to compete once these hurdles are overcome. Yes, there is the cost of the hardware powering AI—and for the most advanced applications, that cost can be substantial—but the cost of cloud computing has been declining for decades, and that trend is likely to continue.

Workers, in response to a major wave of AI-driven automation, will eventually acquire new skills and find new jobs. Perhaps full employment will even be maintained in the short term and the medium term. But this does not mean that AI won’t have a lasting impact on labor. Labor’s share of income may still shrink. And that’s a cause for concern separate from the employment losses or the labor force’s skill obsolescence that AI may bring about.

Notes

1 See, for example, “Pause Giant AI Experiments: An Open Letter” (2023); the press interview with AI expert Geoffrey E. Hinton, widely regarded as the “godfather of AI,” published by the *New York Times* on May 1, 2023; and hearings by the Senate Judiciary Subcommittee on Privacy, Technology, and the Law, held to discuss the safety and ramifications of AI.

2 Some income, such as income earned by proprietors, is challenging to classify. For this and other reasons, a precise measurement of the labor share is not possible and all existing measures of this share are approximate. For an overview of the challenges of measuring the labor share of income, see Gomme and Rupert (2004).

3 For a detailed account of the apparently transient labor share decline in the 19th century, see Acemoglu and Johnson (2023).

4 We draw on our recent research with Cornell University Assistant Professor of Economics Mathieu Taschereau-Dumouchel. See Drozd, Taschereau-Dumouchel, and Tavares (2022).

5 Some of the economic literature finds that labor’s share has already begun to decline due to the advent of modern computer software and modern industrial robots. For a detailed discussion of the decline in labor’s share of income over the last few decades due to software, see Aum and Shin (2020). See also Acemoglu and Restrepo (2020), who analyze the effect of industrial robots on local labor markets. Other hypotheses include the role of market power and concentration, the rise of superstar firms, the effects of trade and outsourcing, and mismeasurement.

6 This definition courtesy of Lipsey, Carlaw, and Bekar (2005).

7 Lipsey, Carlaw, and Bekar (2005).

8 As an example, see the discussion of AI in Brynjolfsson and McAfee (2016) and references in footnote 1.

9 This is supported by the decline in the relative price of capital goods in relation to consumption goods over at least a century of data. See Greenwood, Hercowitz, and Krusell (1997).

10 The stability of labor’s share in the late 1930s and 1940s has been recognized as a statistical fact in the writings of John Maynard Keynes, Paul Samuelson, Robert Solow, and Michal Kalecki. It was known as Bowley’s Law. See Krämer (2011) for a more detailed discussion.

11 This discussion is grounded in a simplified version of our more general task theory of labor share stability and growth involving automation. See footnote 4.

12 This decomposition was first proposed by Acemoglu and Restrepo (2018). Aghion, Jones, and Jones (2017) were the first to highlight the role of complementarity in the context of automation. They proposed a different naming convention because Acemoglu and Restrepo’s productivity effect is essentially the classic Baumol disease effect.

13 Say labor is initially paid $66 and capital is paid $34 on each unit of output sold for $100 in the market. Note that labor can then purchase 66/(66+34) units of output for the income it earns on each unit of output it helps produce ($66). Its share is thus 0.66, or 66 percent. Accordingly, if capital is subsequently paid $33, and labor income is unchanged, labor can purchase 66/(66+33) = 66/99 units, which is its new share. Therefore, labor's share of income increases by approximately 0.66 x 1 percent, since 66/99 = 66/100 x 100/99 ≈ 0.66 x 1.01, and so the share changes by about (0.66 x 1.01) - 0.66 = 0.66 x 1 percent.

14 Another aspect that limits substitutability across tasks is that we consumers want a variety of goods and services.

15 This is an approximation—a number that declines by x percent and subsequently rises by x percent does not return to the previous level, unless x is very small, which is implicitly assumed throughout.

16 For an alternative hypothesis for labor share's stability, see Acemoglu and Restrepo (2018). They stress the emergence of new goods and services and the fact that the new tasks associated with them may be difficult to automate at first. In the process of churning tasks in and out of existence, the labor share can be sustained despite ongoing automation.

17 For example, the resulting frequency of, say, cities with at least N inhabitants, or wealth of at least N dollars, is inversely proportional to N. In other words, cities that are double the size of Philadelphia (or larger) occur in the data half as often as cities that are as large as Philadelphia (or larger). In the context of cities, scale invariant growth means that cities will grow in population randomly and independently of their current population. See Gabaix (2009) for a more detailed discussion of these and other examples.

18 An algorithm is a set of instructions for a machine to operate autonomously. Before AI, these instructions were written exclusively by humans. In essence, AI is a set of tools that writes these instructions autonomously. See the ***What Is Generative AI?*** sidebar in this article for more details.

19 For a detailed account of this turbulent period for labor, see Acemoglu and Johnson (2023).

20 Some of the earliest and most significant research into the intersection of neuroscience and neural networks, particularly regarding the implementation of logical gates, was conducted by Warren McCulloch and Walter Pitts in 1943. They considered a simplified neural network model that used basic logic to describe neural activity. Their seminal research inspired further exploration of neural network structure and ultimately led to the development of the first general neural network model, the Perceptron, by Frank Rosenblatt in 1958. This laid the groundwork for many modern neural network concepts and algorithms. An important milestone for generalizing the Perceptron model to multilayer ANNs with neurons that can modulate signals was the backpropagation algorithm, introduced to the AI community by Rumelhart, Hinton, and Williams in 1986.

21 The known similarities likely end there, although some new research suggests that there might be a tighter connection than previously thought. See Richards (2018) and Richards et al. (2019).

22 For a more technical and detailed description of generative AI technology, see Wolfram (2023).

23 Bubeck et al. (2023) analyze AGI capabilities of generative LLMs and provide many illustrative examples. These examples showcase the ability of AI to "reason" based on the acquired understanding of concepts underlying our human world. They refer to these capabilities as "sparks of intelligence."

24 Eloundou et al. (2023).

25 Ellingrud et al. (2023).

26 See Pizzinelli et al. (2023) for a discussion of how these features may affect automation across countries depending on each country's level of development.

27 Baschuk (2023).

28 Baschuk (2023).

References

"Pause Giant AI Experiments: An Open Letter," Future of Life Institute, March 22, 2023, https://futureoflife.org/open-letter/pause-giant-ai-experiments/.

"Policymaking in the Pause: What Can Policymakers Do Now to Combat Risks from Advanced AI Systems?" Future of Life Institute, April 19, 2023, https://blog.biocomm.ai/2023/04/19/the-future-of-life-institute-policymaking-in-the-pause-what-can-policymakers-do-now-to-combat-risks-from-advanced-ai-systems-19th-april-2023/.

Acemoglu, Daron, and Pascual Restrepo. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108:6 (2018), pp. 1488–1542, https://doi.org/10.1257/aer.20160696.

Acemoglu, Daron, and Pascual Restrepo. "Robots and Jobs: Evidence from U.S. Labor Markets," *Journal of Political Economy*, 128:6 (2020), pp. 2188–2244, https://doi.org/10.1086/705716.

Acemoglu, Daron, and Simon Johnson. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs, 2023.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. "Artificial Intelligence and Economic Growth," National Bureau of Economic Research Working Paper 23928 (2017), http://www.nber.org/papers/w23928.

Aum, Sangmin, and Yongseok Shin. "Why Is the Labor Share Declining?" Federal Reserve Bank of St. Louis *Review*, 102:4 (2020), https://dx.doi.org/10.20955/r.102.413-28.

Baschuk, Bryce. "Tech, AI Driving Job Changes for Nearly a Quarter of All Workers," Bloomberg, April 30, 2023.

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Company, 2016.

Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, et al. "Sparks of Artificial General Intelligence: Early Experiments with GPT-4," Microsoft Research 2023, https://doi.org/10.48550/arXiv.2303.12712.

Drozd, Lukasz A., Mathieu Taschereau-Dumouchel, and Marina M. Tavares. "Understanding Growth Through Automation: The Neoclassical Perspective," Federal Reserve Bank of Philadelphia Working Paper 22-25 (2022), https://doi.org/10.21799/frbp.wp.2022.25.

Ellingrud, Kweilin, Saurabh Sanghvi, Gurneet Singh Dandona, et al. "Generative AI and the Future of Work in America," McKinsey Global Institute Report, July 26, 2023, https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," manuscript, March 21, 2023.

Gabaix, Xavier. "Power Laws in Economics and Finance," *Journal of Economic Perspectives*, 30:1 (2009), pp. 255–294, https://doi.org/10.1146/annurev.economics.050708.142940.

Gomme, Paul, and Peter Rupert. "Measuring Labor's Share of Income," Federal Reserve Bank of Cleveland Policy Discussion Paper 7 (2004), https://dx.doi.org/10.2139/ssrn.1024847.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87:3 (1997), pp. 342–362, https://www.jstor.org/stable/2951349.

Keynes, John M. "Economic Possibilities for Our Grandchildren (1930)," in *Essays in Persuasion*, New York: W.W. Norton & Co., 1963, pp. 358–373.

Krämer, Hagen M. "Bowley's Law: The Diffusion of an Empirical Supposition into Economic Theory," *Cahiers d'économie Politique / Papers in Political Economy*, 61:2 (2011), pp. 19–49, https://doi.org/10.3917/cep.061.0019.

Lipsey, Richard G., Kenneth I. Carlaw, and Clifford T. Bekar. *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*. Oxford, UK: Oxford University Press, 2005.

Metz, Cade. "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead," *New York Times*, May 1, 2023, https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html.

McCulloch, Warren S., and Walter Pitts. "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bulletin of Mathematical Biophysics*, 5 (1943), pp. 115–133.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. *Capital in the Twenty-First Century*. Cambridge, MA: Belknap Press of Harvard University Press, 2014.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. "Distributional National Accounts: Methods and Estimates for the United States," *Quarterly Journal of Economics*, 133:2 (2018), pp. 553–609, https://doi.org/10.1093/qje/qjx043.

Pizzinelli, Carlo, Augustus Panton, Marina M. Tavares, et al. "Labor Market Exposure to AI: Cross-country Differences and Distributional Implications," IMF Working Paper 2023/216 (2023), https://www.imf.org/en/Publications/WP/Issues/2023/10/04/Labor-Market-Exposure-to-AI-Cross-country-Differences-and-Distributional-Implications-539656.

Richards, Blake A. Speech given at the RAAIS Conference, London, June 29, 2018, https://raais.co/speakers-2018.

Richards, Blake A., Timothy P. Lillicrap, Philippe Beaudoin, et al. "A Deep Learning Framework for Neuroscience," *Nature Neuroscience*, 22 (2019), pp. 1761–1770, https://doi.org/10.1038/s41593-019-0520-2.

Rosenblatt, Frank. "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain," *Psychological Review*, 65:6 (1958), pp. 386–408, https://doi.org/10.1037/h0042519.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. "Learning Representations by Back-Propagating Errors," *Nature*, 323:6088 (1986), pp. 533–536, https://doi.org/10.1038/323533a0.

Wolfram, Stephen. "What Is ChatGPT Doing and Why Does It Work?" Wolfram Research, LLC, March 9, 2023.

World Economic Forum. *Future of Jobs Report 2023*. Geneva, Switzerland, May 2023, https://www.weforum.org/publications/the-future-of-jobs-report-2023/.