CITY JOURNAL

EYE ON THE NEWS

Yes, This Time It's Different

ChatGPT and automation come to knowledge work.

Mark P. Mills
February 15, 2023

More than 1 million users signed up within a week of the online release of the artificial intelligence (AI) tool ChatGPT on November 30, 2022, a date dystopians think will live in infamy. In just two months, 100 million users signed up, a record for any new software, app, or online tool.

ChatGPT—along with other recent, if less famous, competitors—does in fact constitute a radical advance in software, far more consequential than when IBM's Deep Blue defeated chess champion Gary Kasparov in 1997, back in the Internet's first "century." ChatGPT puts us one giant step closer to normal people having the ability to access supercomputing performance using natural language. So, the chattering classes are all a-twitter with the usual tropes about how automation will bring about the end of work, in this case for "knowledge workers."

We will, in due course, see news stories about practical "use cases" for ChatGPT. But for now, a lot of people are having a lot of fun playing with it and discovering just how much better this latest AI software is compared with, say, computer interfaces like Alexa or Siri.

Since ChatGPT was trained to recognize and produce language—a fact telegraphed by the word "chat"—I joined the rush to test the new tool, asking it to write a short blurb to promote a Caribbean cruise and to do so in the style of William Shakespeare. Within a fraction of a second, ChatGPT came back with a C+ effort:

> Hark! good sirs and fair maidens, lend an ear
> To news of grandeur and luxury rare.
> A Caribbean cruise doth now appear
> For those with hearts that do adventure dare.
> Upon the azure seas, our grand barque
> Shall bear ye to fair isles of sand and sun
> Where comfort and opulence shall embark
> And memories forever to be won.

One user gave ChatGPT a standard medical exam; it passed. Another gave it a standard test for Wharton MBA students; passed again. Impressive parlor tricks, but such stunts tell you more about the simplicity of the tests than about the brilliance of ChatGPT. Passing a written test that indicates proficiency with medical, business, or finance rules and nomenclature doesn't say anything really useful about whether the test taker will be a great surgeon or business leader. Nor would passing the written test for driving a car predict anything about whether the test-taker could be a safe driver, much less a winning Formula One racer.

Engineers love to perform stunts with new technologies. Of course, depending on the stunts, the performance *does* say something about the state of a technology. The lesson of Kasparov's loss was to see the "overnight" progress, nearly 20 years later, of Google's AlphaGo supercomputer beating the world's Go champion in that much more complex game. Now, almost a decade after that, comes the sonnet-writing, test-passing ChatGPT. The point is that even in the seemingly high-velocity world of computing, so-called exponential change takes quite a while. That's the reality of commercializing at scale all forms of engineering progress.

For a way to think about what comes next for AI, consider what followed analogous stunts in the history of aviation—a field more relevant to computing than most realize.

The feat that made it clear that an age of useful aviation was possible was Charles Lindbergh's 1927 barnstorming of all 48 contiguous U.S. states over a period of 95 days, following his better-known stunt of the first nonstop flight across the Atlantic Ocean, which was in retrospect a kind of Kasparov–Deep Blue moment. Even though, following Lindbergh's odyssey, aircraft would be used in business, industry, and warfighting, it took another three decades until the engineering was good enough to yield, in 1957, the Boeing 707, which launched the age of mass commercial aviation. From there, as entirely new industrial edifices and national infrastructures were built out, the number of

revenue-passenger-miles (to use that industry's term of art) would soar more than a hundredfold by the year 2000.

The emergence of useful, broadly available aviation brought big shifts to the structure and nature of business and employment in transportation. But it ended neither the role of, nor the expansion of uses for, ships, trains, or trucks; and it didn't end employment in those sectors. Overall, in fact, employment in U.S. transportation services doubled by the turn of the century. It's no exaggeration to frame ChatGPT as a Boeing 707 moment.

But we're being told that, well, this time is different. In part, that's because it seems somehow spookier when technology accelerates tasks performed invisibly in cyberspace —that is, cognitive rather than physical tasks. ChatGPT has reanimated the now-ancient philosophical debates about whether machines think and whether, as they get better at imitating human behaviors, they'll make a lot of humans redundant and bring on the often-predicted age of unemployable humans.

It is true that there *is* something different this time, as there is every time. The specifics of the newest machines *are* different. But what's not different is the overall effect of automation. Not to diminish the social and political challenges that all disruptions bring to markets and people, but automation has always boosted productivity and thus overall wealth *and* employment. If labor-saving technologies—namely, automation—were a net job destroyer, unemployment should have been continually rising over the course of modern history as (physical) automation inexorably expanded. It didn't. MIT economist David Autor has been particularly eloquent on the apparent paradox of seeing continued rise in employment despite advances in labor-reducing technologies, observing that "the fundamental threat [to employment growth] is not technology per se but misgovernance."

Of course, where and how most people are employed has changed over time. It's going to change again. And that is disruptive. But the central and unprecedented difference between our time and previous eras is the demographic reality of a shrinking workforce. In the near future, we will need lots of new tools to amplify the efforts of the declining labor supply. Even in our own present, despite the best efforts of the Federal Reserve to increase unemployment (that is, to reduce the pressure employers face to offer "inflationary" salaries to keep workers), job openings still outnumber people available to fill them. Demographics dictate that this gap will widen. Since most jobs in a modern economy are found in so-called knowledge work, the only way to close the labor gap will be with AI tools useful enough to amplify the efficacy of people in those areas.

AI, of course, is not a specific tool per se but a class of tools under that loosely defined term. To extend the earlier analogy, there are many radically different kinds of engines; no single engine is suitable for every class of machine, task, or vehicle—from aircraft to mining trucks. It's the same for the silicon engines at the core of all AI machines. Much of the misdirection about AI's implications comes from the sloppy term itself, "artificial intelligence." It's no more informative or accurate than calling a car an artificial horse, or an airplane an artificial bird, or an electric motor an artificial waterwheel.

While ChatGPT is a whiz with words, it wasn't trained on math and, as some users have already observed, performs poorly there. Similarly, ChatGPT couldn't drive a car, wield a hammer to drive a nail, or carry a box. One needs differently designed and trained AI tools to perform each kind of task. The category confusion about the realities of AI tools is, to put it crudely, the equivalent of seeing that a tool like a hammer makes it easier to push a nail into a board and then trying to use a hammer to drill a precise hole, weld steel, or measure voltage.

The letters GPT in ChatGPT stand for Generative Pre-trained Transformer—computer lingo for an algorithm that, when paired with a powerful computer, can be trained iteratively by looking repetitively at a very large set of samples—in this case, written texts. The same has been done for images and myriad areas where routine tasks entail patterns and rules. The "chat" in ChatGPT will doubtless find early commercial application precisely where chatbots are already used: in online commerce and with the many tasks in all businesses that involve often-confusing or arcane rules, regulations, or standards that a computer can more usefully, quickly, and accurately parse to answer questions put to it in "natural language."

The management literature is replete with analyses of the productivity-robbing burdens imposed on employees trying to comply with routine tasks in education, health care, business in general, and even in basic research. Such tasks are precisely where the narrow power of AI is most powerful. As it happens, it's also where one could free up easily re-trainable humans to be redirected to more challenging non-routine tasks.

A recent Federal Reserve analysis divided the U.S. workforce into just two high-level categories: manual and cognitive labor. No surprise that the majority are now employed in the latter. The analysis also created two sub-categories within each: routine and non-routine tasks. Thus classified, about 60 million people in the U.S. work on *routine* tasks, split almost evenly between manual and cognitive domains. Total employment in both

routine manual and routine cognitive tasks hasn't changed significantly since 1980. Meantime, the non-routine manual labor pool has risen from 15 million to 25 million people since 1980, and employment in non-routine cognitive work has grown from about 30 million to 60 million people.

Four decades of job growth has all been in non-routine tasks. If we want to find more people to take on the jobs where growth is happening—and where they can be paid more—we'll need to move people out of the routine job domains, while still ensuring that those tasks are fulfilled. That is precisely what's made possible by AI tools that can increase the efficacy of a shrinking number of people performing routine tasks. Ensuring that that can happen will require AI tools even easier to use, more accurate, and cheaper than what's available today with ChatGPT and its (jealous) competitors.

We know from history that when new technologies are found to be broadly useful, engineers drive down costs and make them easier to use. The latter is the "user interface," in the jargon of tech. Again, witness the capabilities of ChatGPT versus, say, Alexa. With Natural Language Processing (NLP), the human-machine interface makes it easier for non-experts to engage casually in computational feats previously reserved for supercomputers and the expert class. The overall effect of NLP, in addition to taking up the burden of routine tasks, will also be to reduce routine burdens for employees in non-routine types of work. It will also enable the upskilling of more people to become "knowledge workers," including even coding. It's no coincidence that AI tools are bringing greater productivity to writing computer code. One company touts that its AI-based tool can help a coder write software ten to 100 times faster.

The good news, at least from a macroeconomic perspective, is that there's been a land-rush of activity to develop mission-specific machine-learning algorithms. One measure of the scale of that activity is in the amount of private capital chasing AI deals and companies. We're in the early stages of billions of dollars directed at another tech hype cycle.

Another measure of the scale of AI activity can be found in the total quantity of the world's computer processing power used to "train" deep-learning models; it's been doubling every few *months* for the past half-dozen years. That translates into a 300,000-fold increase in computing power used for AI training over that short time. You don't need a crystal ball to predict that such prodigious efforts will soon yield a fusillade of useful AI tools to succeed ChatGPT.

Coming back to our aviation analogy, it's the inescapably physical world of energy that reveals the implications of the scale of AI and machine learning. Even AI cognoscenti are surprised to learn that the energy equivalent of the fuel used to fly a jumbo jet from Austin to Asia is gobbled up by an AI-centric computer being trained on "large language models" or other similar sets of "parameters" needed for machine learning. That's not a one-time investment; it happens every time and for each kind of similar application of learning. As "use cases" for AI expand, the proliferation of AI training will follow apace.

Ah, but for those who are anxious about energy issues, we also know that emerging and next-generation AI chips and algorithms are far more energy-efficient—some are already tenfold better. This will tamp down AI's voracious energy appetite, even as the tools improve. But it's that reality—more efficiency and higher performance—that will lead to a repeat of the trajectory of the first, pre-cloud era of the Internet.

Radical gains in efficiency have always been critical to unlocking the commercial viability of any new machine or infrastructure for society. In 1958, when Pan Am began passenger jet service with the 707, no one forecasted (much less exhibited angst about) the aggregate fuel consumption that commercial aviation would induce. Since then, aircraft have become 300 percent more energy-efficient, not to mention safer and more reliable. Those features are what enabled today's trillions of passenger-miles flown, an activity that consumes some 4 billion barrels of oil each year, compared with a trivial amount in 1958.

Similarly, decades of inexorable gains in computing energy efficiency are illustrated by the fact that, if today's smartphones operated at 1980 computing-efficiency levels, just one phone would use as much electricity as an office building. A single datacenter at 1980 efficiency levels would require the entire U.S. grid to power it. Instead, staggering reductions in energy-per-logic-operation, again along with gains in performance, are what made possible a commercial world of billions of smartphones and thousands of datacenters. And that yielded today's global cloud infrastructure—still in a pre-AI era—that already uses about as much energy as global aviation.

In a future when AI machines perform not dozens but tens of thousands of simulations entailing trillions of computing-hours, overall energy use will balloon again. And that will happen because of the economic benefits AI offers to people, businesses, and even—and especially—in the pursuit of science and new discovery.

As Regina Barzilay, an AI researcher at MIT, put it when asked about the power of AI-assisted discoveries to invent new life-saving drugs: it's "not the machine that invented the molecule. It's that the machine helped humans to scan the huge space of possibilities and zoom in on the fruitful set of hypotheses that they tested." Or, as economist Alexander Salter succinctly observed: "Data doesn't interpret itself." The AI machines are knowledge amplifiers.

Even so, we will see disruptions to the nature of jobs and businesses. Indeed, the scale of those disruptions will echo the magnitude of the opportunities that AI creates. Some educators have voiced worries about disruptions to teaching, including detecting cheaters. ChatGPT will indeed require adjustments, perhaps even a return to Socratic methods of in-class learning and testing—hardly a new idea. Nor is dealing with cheating, especially in the age of the Internet. Teachers found ways to teach math in the age of the calculator. Adaption to AI is not just possible, but arguably beneficial.

A clear-eyed recognition of benefits from any new technology doesn't constitute a Pollyanna's perspective. It's also true that AI machines won't all be useful or put to good use; such is the (sometimes sad) state of human nature. As science-fiction author and technology seer Cory Doctorow recently quipped in a long interview, "I think that the problems of A.I. are not its ability to do things well but its ability to do things badly, and our reliance on it nevertheless." His cautions—and these are a constant refrain in his dystopian fiction—center around the need to recognize the limits of any machine and the kinds of risks arising from misuses.

Coming back to where we started, looking over the long period since the emergence of the modern information era, circa 1970, Census data show a significant shift in the *structure* of employment—away from production and toward services. Economists David Autor and Anna Salomons have done pioneering work in mapping those dynamics as a kind of hollowing out of highly paid "middle-skilled" jobs that don't typically require a college degree, and a simultaneous shift toward more low- and high-skilled employment.

Autor recently posed a question as to "whether a countervailing set of economic forces will soon reverse the decline of middle-skill work?" I think the answer to that question is yes. The countervailing forces will come from the fact that computing has finally become widely useful with the advent of commercially viable AI. And that's happening just in time to rescue the economy from demographic dystopia.

Mark P. Mills is a senior fellow at the Manhattan Institute, a strategic partner in the energy-tech venture fund Montrose Lane, and author of **The Cloud Revolution: How the Convergence of New Technologies Will Unleash the Next Economic Boom and a Roaring 2020s**. *He hosts The Last Optimist podcast.*