

FOREIGN AFFAIRS

AUGUST 9, 2023

Artificial Intelligence for the Poor

How to Harness the Power of AI in the Developing
World

DANIEL BJÖRKEGREN

Artificial Intelligence for the Poor

How to Harness the Power of AI in the Developing World

DANIEL BJÖRKEGREN

Among elites in wealthy countries, a worry about artificial intelligence has taken hold: the machines will take our jobs. With the explosive popularity of ChatGPT, the remarkably lifelike chatbot, many in the West have begun to fear that it is not only truck drivers and assembly workers who are at risk of being replaced by robots but highly paid knowledge workers, too. Accountants, data analysts, coders, financial advisers, lawyers, even Hollywood screenwriters—all now worry that AI will leave them jobless.

But AI's effect on the 100-odd countries and more than four billion people in the developing world is likely to be very different. Lower-income countries employ far fewer knowledge workers, and a larger share of their populations work in sectors that are less amenable to automation, particularly agriculture. In poor countries, the big question is not how AI will affect millions of employed people but how will billions of people employ AI. The most transformative applications in the developing world

will probably not be those that replace humans; they will be those that open new possibilities for humans.

So far, nearly all discussion of how to support AI and how to mitigate its risks has centered on rich countries, which are home to the companies and universities working on the technology. But because the effects of AI—good and bad—will play out differently in poor countries, the investments and regulations these countries need are also likely to be different. Philosophers, economists, and technologists have spilled endless ink contemplating the future of AI in the developed world. It is now time to think through an AI agenda for everyone else.

MACHINE POWER

Machine learning has already touched the lives of the world's poor. Consider developments in credit. Many poor people lack financial histories and credit scores and thus have little access to formal loans. In 2010, I proposed a way of creating alternative credit scores, using machine learning to draw inferences about the likelihood of repayment from data automatically collected by cellphone networks. This method is now one of several that lenders in dozens of countries have employed to offer small loans via mobile phone to millions of people. Other researchers are applying machine learning to the same kind of data to identify which households in a given area are poorest, so that aid can be smartly targeted during a crisis. Still others are employing it on satellite images, refining population estimates based on patterns of human settlement and anticipating food shortages based on patterns of vegetation. Such programs highlight one particular value of AI in the developing world: in low-information environments, machine learning can draw signals from new sources of data.

The possibilities do not stop there. Consider schooling. Most education systems in developing countries struggle to deliver quality instruction. Personalized AI tutors—chatbots with endless patience—might someday meet the needs of curious students in remote schools. They might also help professionals transition between skills—allowing, say, repair workers to level up their skills and learn engineering. Or take health. In much of

the developing world, sound medical advice is hard to come by; AI-powered systems might offer better and more widely available diagnostics. Many communities have high rates of depression and few therapists; digital mental health tools such as chatbot therapists might fill a real need at a low cost. AI could play a similar role helping people navigate bureaucracies. An Indian entrepreneur looking to enter a new market, for example, might someday be able to rely on an AI-powered app to fill in the required permits.

The technologies that enable these potential applications will continue to improve as wealthy countries invest enormous resources in AI. The key for developing countries will be to complement that stream of investments by using the resulting technologies in products and services that meet local needs. Developing countries have much of the social infrastructure needed to start new ventures: tech hubs, universities, and entrepreneur groups. Their companies, however, have little incentive to build applications aimed at the poorest people, who are seldom profitable to serve. Some large, middle-income countries such as India can afford to overcome this problem by investing in AI technologies for the poor. But many other countries lack the resources and scale to do so. Hence, there is a role for networks of entrepreneurs, who can share learning across borders, and for international organizations such as the World Bank, which can coordinate investments between governments and philanthropies.

LEARNING CURVE

There are two main paths AI tools could take in the developing world. The first is to find a task that AI is becoming good at in wealthy countries and adapt it for poor countries. For example, many entrepreneurs are developing chatbot tutors for wealthy schools, tools that could be modified to work in places with worse Internet connectivity and higher student-to-teacher ratios. The second is to find altogether new applications for AI—new products that could meet the specific needs of the developing world. For example, an AI-powered financial planner for subsistence farmers might help them manage the risks involved in decisions about what to plant. Indeed, some innovations began in a poor

country and only reached richer ones later. Kenya's M-Pesa mobile payment system, for example, took off well before similar apps did in the United States.

While some AI tools emerging from wealthy countries may work well right out of the box in the developing world, others will require tailoring. One problem is that most AI systems have been trained on data specific to the developed world, data that is gathered from people with relatively high incomes and is usually written in English. Little of the world's corpus of written knowledge is about the poor or presented in minority languages. Moreover, AI systems are mostly trained to produce decisions and outputs that satisfy wealthy consumers in the West, so they may make faux pas when dealing with poorer ones in other places—for example, greeting customers by their first name in a culture that deems such familiarity disrespectful.

Wealthy Western societies had a head start amassing training data, so it will take time for AI models to fully represent people from the rest of the world. But the process can be hastened. Researchers can identify applications that could prove transformative, if only one could make the data behind them more representative. An AI-powered medical adviser, for example, may be good at helping a person with high blood pressure in Silicon Valley but less useful for someone in Lagos facing malaria because it lacks exposure to local medical cases. Or such a system might prove popular among English speakers but not be available in Yoruba, one of Nigeria's main native languages.

To compensate for the dearth of developing-world data, new content must be created for the models to train on. Here, crowdsourcing could help. The WikiAfrica movement, for example, coordinated the addition of African content to Wikipedia. Such initiatives are all the more valuable now that this knowledge can improve the decisions of machines. In other domains where correctness is harder to discern—such as medicine or agriculture—crowdsourcing will not be enough. Experts will have to be hired, or analog data, such as paper clinic records, will have to be digitized. Representation is only part of the puzzle because developers will have to arbitrate between groups with different values. Different religious

groups in India, for example, may disagree over what constitutes appropriate medical advice.

A second problem with importing AI to the developing world is technological. Despite vast progress, the developing world still lags behind the developed world on a number of technological benchmarks. Some AI applications will require wider access to smartphones, better Internet connectivity, or digital recordkeeping systems to track the performance of students in a school, the health of patients in a hospital, or the outcome of cases in a courthouse. For AI, as with previous waves of technological innovation, the key will be to differentiate between applications that can be valuable relatively soon and those that will remain in the realm of science fiction for the foreseeable future. That line will shift, and it will vary from one field to another. For example, medicine has a lower tolerance for the mistakes that AI systems will inevitably make, and agriculture depends on nuanced contextual factors that are intuitive to farmers but difficult to express to AI systems.

THE LIMITS OF LAWS

In the developed and developing world alike, the diffusion of AI will present risks. But developing countries face a different array of risks, and they are less able to regulate the technology. The main question is whether the technology will remain centralized—that is, controlled by a small number of tech firms. Centralized AI systems are likely to be regulated in large markets such as the United States and the EU. Smaller markets can exert only limited pressure, so they will live in the shadow of U.S. and EU regulation. Although they could shut off access to a centralized system—for example, blocking servers, just as some authoritarian governments have done with Twitter, Facebook, and YouTube—they will not be able to prevent AI-generated content from crossing borders.

It is not clear, however, whether AI will remain centralized. Open-source alternatives such as Llama (a large language model produced by Facebook's owner, Meta) and Stable Diffusion (an image generator made by the startup Stability AI) are gaining ground. These decentralized systems can be modified and run by anyone with a computer. If they

become sufficiently useful, it will be difficult for any country to regulate them directly. But such open systems can be more easily adapted to local needs because they are often free to use and because anyone can modify their code. Given the limited levers for regulation, developing countries may have to settle for adapting to new technology rather than controlling it. To mitigate harms, they may have to focus on regulating not AI itself but the industries that use it—for example, resorting to consumer protection laws that hold companies liable when a product is unsafe, regardless of whether it uses AI.

AI has kicked off a healthy debate about regulation in rich countries. But many of the proposals for addressing its risks may be insufficient in poor countries. Regulators in the West lack the ability to assess how the rules work in different contexts; a system that is certified as safe in Brussels might not work so well in Bangalore. Moreover, Western regulators' standards may be inappropriately strict in places where the existing alternatives to an AI application are much worse. Weather forecasts, for example, need not be perfect to improve on what is available to farmers in developing countries. And even in higher-stakes settings such as medicine, AI may soon be better than existing options available to the poor. One 2023 study audited clinical performance in low-income countries to find out what fraction of cases were correctly handled. The answer: less than half.

At the same time, the average person in a developing country is also more vulnerable than his or her counterpart in the developed world. Many people in the developing world have little recourse to challenge automated decisions, such as the rejection of a loan application. New AI systems often perform worse than advertised, and it is all too easy for companies to ignore problems that arise among lower-income people. That is why it will be important for regulators to ensure that consumers have adequate processes to report issues and appeal decisions.

Many people in the developing world also are new to the idea of AI and have never heard of algorithms before. So care must be taken to communicate effectively. A study I conducted with

Developing
countries may have

to settle for adapting to new technology rather than controlling it.

Joshua Blumenstock and Samsun Knight shows this is possible. We gave low-income Kenyans an app that rewarded them financially based on how they used their mobile phone, employing an algorithm similar to those that score one's creditworthiness. When subjects were given straightforward descriptions of how algorithms work, they adjusted their behavior—a concrete sign of understanding.

Political obstacles also abound. Deepfakes—realistic photos, videos, and audio clips generated by AI—can have an especially pernicious effect in developing countries, where political systems tend to be fragile, and trust between groups is often low. As people become aware that media can be generated, they may cease to believe incriminating content that is actually true. To head off these problems, civil society can play a role in building the infrastructure of trust—spreading awareness that content may be faked and establishing independent venues that develop reputations for vetting content.

AI will also enable new forms of surveillance, such as tracking people through mobile devices and facial recognition. Most developing countries in the market for high-tech surveillance tools do not develop their own but instead import them, often from China. This outsourcing means that the actual implementation of AI-powered technology may be scattershot, making it easier for the information collected to be leaked to third parties and for rights to be infringed in unpredictable ways. Once again, civil society will have a role to play, monitoring new systems and drawing attention to abuses.

BACK TO THE FUTURE

This current wave of AI has introduced challenges and opportunities with unprecedented speed. But we have seen similar technological transitions before. Although mobile phones were initially designed for wealthy consumers, they took off among the poor over the past 20 years. Developing countries benefited from the standardized hardware—antennas and handsets—made in the West. Telecom companies invented

business models that served the poor, such as pay-as-you-go cellphone plans. Entrepreneurs started new organizations that allowed people to use phones to send money, obtain credit, and check prices. These innovations allowed mobile phones to quickly reach most of the world's poor and connect them to the global economy.

It is these very links that have set the stage for the spread of AI. Yet despite the success of mobile phones, even that innovation has fallen short of its potential in the developing world. Most private-sector innovation has focused on the needs of the wealthy. Much more has been invested in apps to connect rich consumers to drivers, vacation houses, and prepared meals than in apps to connect subsistence farmers to markets and remote children to learning. Private-sector innovation in AI is likely to transform many industries, from education to health to law. But harnessing the full potential of the technology for developing countries will require formulating an expansive vision of what is possible—and paying extra attention to the people whose lives it could change.