


TECHNOLOGY FOR DEVELOPMENT

Daniel Björkegren and Joshua Blumenstock

AI must be carefully adapted to benefit the poor, shows research in Kenya, Sierra Leone, and Togo

As artificial intelligence (AI) reshapes developing economies, it raises familiar risks of disruption, misinformation, and surveillance—but also promises many potential benefits. Recent examples illustrate how AI-based technologies can target aid and credit better and improve access to tailored teaching and medical advice. But balancing these risks and opportunities means more than just plug and play of existing technology—it calls for local innovation and adaptation.

Most recent advances in artificial intelligence originated in wealthy nations—developed in those countries for local users, using local data. Over the past several



GERALD ANDERSON/ANADOLU AGENCY VIA GETTY IMAGES

Kenyan youth attend a robotics and coding class to learn future technologies.

years, we have conducted research with partners in low-income nations, working on AI applications for those countries, users, and data. In such settings, AI-based solutions will work only if they fit the local social and institutional context.

In Togo, where the government used machine learning technology to target cash aid during the COVID-19 pandemic, we found that adapting AI to local conditions was the key to successful outcomes. The government repurposed technology originally designed to target online advertising to the task of identifying the country's poorest residents. Using AI, the system processed data from satellites and mobile phone companies to identify signatures of poverty—such as villages that appeared underde-

veloped in aerial imagery and mobile subscribers with low balances on their phones. Targeting based on these signatures helped ensure that cash transfers reached people with the greatest need (Aiken and others 2022).

This application worked in Togo only because the government, in collaboration with researchers and nonprofit organizations, customized the technology to meet local needs. They built a system for distributing mobile money payments that worked for all mobile subscribers, adapted existing machine learning software to target cash transfers, and interviewed tens of thousands of beneficiaries to ensure that the system reflected the local definition of poverty. And even then, the AI-based solu-

tion was not designed to be permanent; it was to be phased out after the pandemic ended.

The AI-based program also raised another concern: algorithms that perform well in a laboratory may not be reliable when deployed for consequential decisions on the ground. For instance, in an aid-targeting system like the one in Togo, people might adapt their behavior to qualify for benefits, thereby undermining the system's ability to direct cash to the poor.

Elsewhere, machine learning is used to determine eligibility for microloans, based on mobile phone behavior (Björkegren and Grissen 2020). For example, in Kenya over a quarter of adults have taken out loans using their mobile phones. But if those with more Facebook friends are likelier to be approved for a loan, some applicants may consider adding friends quickly. Ultimately, this can make it hard for systems to target the intended people.

In a study with the Busara Center in Kenya, we found that people were able to learn and adjust their smartphone behavior in response to such algorithmic rules (Björkegren, Blumenstock, and Knight, forthcoming). We showed how a proof-of-concept adjustment to the algorithm, which anticipates these responses, performed better. However, technology alone cannot overcome problems that arise during implementation; much of the challenge of building such systems is ensuring that they are reliable in real-world conditions.

On the other hand, some systems require adaptation before they will be useful. For instance, in many lower-income countries, teachers must handle large classes with limited resources. In Sierra Leone, a local partner piloted an AI chatbot system for teachers, called TheTeacher.AI, which is similar to ChatGPT but tailored to local curriculum and instruction and accessible even when internet connections are poor. In the pilot phase, many teachers couldn't phrase questions in a way that yielded useful answers, but a small group began to use the system regularly to help with teaching concepts, planning lessons, and creating classroom materials (Choi and others 2023). It took training and experimentation for teachers to use it in practice. Uses of AI may not be immediately obvious to those who stand to benefit; discovering the many uses will depend on trial and error and sharing applications that help.

Communication barriers

Grasping the potential of AI is likely to be harder for people in lower-income countries, where literacy and numeracy are lower and residents are less familiar with digital data and the algorithms that process this information. For instance, in our field experiment in Nairobi, Kenya, we found it difficult to explain simple


“AI systems will require investment in knowledge infrastructure, especially in developing economies, where data gaps persist and the poor are digitally underrepresented.”

algorithms with negative numbers and fractions to low-income residents. But our team found simpler ways to communicate these concepts. It was clear when people responded to the algorithm that they grasped the concept. Still, complex AI systems are difficult to understand, even for AI researchers.

Some applications don't require that users know how algorithms work. For instance, Netflix movie recommendations can benefit users even if they do not understand how the algorithm selects content it thinks they will like. Likewise, in a humanitarian crisis, policymakers may deem it acceptable to use an inscrutable “black box” algorithm, as Togo's government did in response to the COVID-19 crisis.

Transparency is sometimes critical. When targeting social protections in nonemergency settings, explaining eligibility criteria to potential beneficiaries is essential. This is easier said than done: scores of interviews and focus groups showed us how norms and values around data and privacy are fundamentally different in a setting such as rural Togo than in wealthy nations, where AI-based systems are more common. For instance, few people we spoke to were worried about the government or companies accessing their data (a dominant concern in Europe and the United States), but many





Juliana Rotich works on a laptop at the i-Hub technology innovation center in Nairobi, Kenya.

wondered if and how such information would be shared with their neighbors.

As AI is more commonly deployed, populations must understand its broader societal effects. For instance, AI can generate provocative photographs that are entirely false and robocalls that mimic voices. These rapid changes will affect how much people should trust information they see online. Even remote populations must be informed about these possibilities so that they are not misled—and to ensure that their concerns are represented in the development of regulations.

Building connections

AI solutions rest on existing physical digital infrastructure: from massive databases on servers, to fiber-optic cables and cell towers, to mobile phones in people's hands. Over the past two decades, developing economies have invested heavily in connecting remote areas with cellular and internet connections, laying the groundwork for these new applications.

Even though AI applications benefit from digital infrastructure, some could make better use of existing resources. For example, many teachers in Sierra Leone struggle with poor internet access. For some tasks, it may be easier to get ideas from a chatbot

and then validate the response than to collate information from several online resources.

Some AI systems will, however, require investment in knowledge infrastructure, especially in developing economies, where data gaps persist and the poor are digitally underrepresented. AI models there have incomplete information about the needs and desires of lower-income residents, the state of their health, the appearance of the people and villages, and the structure of lesser-used languages.

Gathering these data may require integrating clinics, schools, and businesses into digital record-keeping systems; creating incentives for their use; and establishing legal rights over the resulting data.

Further, AI systems should be tailored to local values and conditions. For example, Western AI systems may suggest that teachers use expensive resources such as digital whiteboards or digital slide presentations. These systems must be adjusted to be relevant for teachers lacking these resources. Investing in the capacity and training of local AI developers and designers can help ensure that the next generation of technical innovation better reflects local values and priorities.

Artificial intelligence promises many useful applications for the poor across developing economies. The challenge is not in dreaming big—it's easy to imagine how these systems can benefit the poor—but in ensuring that these systems meet people's needs, work in local conditions, and do not cause harm. **F&D**

DANIEL BJÖRKEGREN is an assistant professor at the Columbia University School of International and Public Affairs. **JOSHUA BLUMENSTOCK** is a chancellor's associate professor at the University of California, Berkeley, School of Information and the Goldman School of Public Policy.

REFERENCES

- Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E. Blumenstock. 2022. "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid." *Nature* 603 (7903): 864–70. <https://doi.org/10.1038/s41586-022-04484-9>.
- Björkegren, Daniel, and Darrell Grissen. 2020. "Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment." *World Bank Economic Review* 34 (3): 618–34. <https://doi.org/10.1093/wber/lhz006>.
- Björkegren, Daniel, Joshua E. Blumenstock, and Samsun Knight. Forthcoming. "Manipulation-Proof Machine Learning."
- Choi, Jun Ho, Oliver Garrod, Paul Atherton, Andrew Joyce-Gibbons, Miriam Mason-Sesay, and Daniel Björkegren. 2023. "Are LLMs Useful in the Poorest Schools? theTeacherAI in Sierra Leone." Paper presented at the Neural Information Processing Systems (NeurIPS) Workshop on Generative AI for Education (GAIED).