LINKS GLOBAL
Communicating with Purpose. Powered by Insight.

THOUGHT LEADERSHIP SERIES

AI for Good

The case for proper technology adoption, and the questions "AI for good" isn't asking

A Links Global thought leadership article by M. Sanchez, founder of Links Global, a communications firm working at the intersection of trust science, narrative design, and civic engagement. Part of the Links Global article series: AI and Technology. Drawn from 25 years of field practice and evidence-based research in strategic communication, public engagement, and institutional trust.

In Brief

"AI for good" is a sincere banner and an incomplete one. This article credits what the movement gets right, then asks the questions it tends to skip: who defines good, what is being extracted, what dependency is being created, and what is being displaced. It closes with what proper technology adoption in community-serving settings actually requires.

"AI for good" has become one of the defining phrases of the current technology moment, with its own annual summit, the AI for Good Global Summit, organized by the International Telecommunication Union with more than 50 UN partners: a banner under which an expanding coalition of technology companies, foundations, research institutions, and social sector organizations has assembled to articulate a vision of artificial intelligence deployed in service of human flourishing rather than profit extraction.

The vision is compelling. The applications are often genuinely promising. And the phrase itself is doing something important: it insists on the question of purpose, demanding that AI development and deployment be evaluated not just by technical capability but by whose lives it improves, under what conditions, and at whose expense.

But "AI for good" as a field has a significant limitation that its proponents have been slow to name: the "good" in "AI for good" is doing enormous definitional work that the field has not yet done rigorously enough. Good for whom? Good by what standard, over what timeframe, assessed by what method, and with whose participation in the assessment? Good as defined by the technology developers, the institutional funders, the recipient communities, or some negotiated synthesis of all three, and if the latter, through what process, with what power balance?

These are not hostile questions. They are the questions that distinguish "AI for good" as a genuine commitment from "AI for good" as a branding exercise, and the answers to them, honestly examined, change what responsible technology adoption looks like substantially.

What "AI for Good" Gets Right

Before examining its limitations, it is worth naming what the "AI for good" frame has genuinely contributed.

It has created space for conversations about AI purpose and AI ethics that the mainstream technology development conversation was not having. In an industry whose primary optimization target for most of its history has been engagement, growth, and return on investment, the insistence that technology should be evaluated by its contribution to human wellbeing is not trivial. It is a genuine reorientation of the evaluative frame, and it has produced real consequences: the development of AI ethics as a field, the emergence of algorithmic accountability research, the creation of fairness, accountability, and transparency standards, and the increasing attention to AI bias as a technical and governance problem.

It has also produced specific applications that have demonstrated genuine value. In documented cases, AI tools applied to disease surveillance have detected outbreak signals faster than traditional epidemiological methods. In some studies, machine learning applied to climate data has identified patterns relevant to disaster risk that human analysts were missing. Natural language processing applied to electronic health records has shown promise in flagging patients at risk of deterioration in some validated settings, though results vary: a widely deployed sepsis-prediction model performed poorly on external validation (Wong et al., 2021). In documented cases, satellite imagery analysis has identified infrastructure damage in disaster zones faster than ground surveys could. These are real capabilities with real implications for real people's lives.

The "AI for good" frame, at its best, creates the institutional and discursive space in which these applications can be developed, evaluated, and improved. That is a genuine contribution to the field.

What "AI for Good" Is Not Asking

The limitation of "AI for good" as a framework is not in its aspirations. It is in the questions it tends not to ask: questions whose absence shapes what gets built, for whom, and with what consequences for the communities ostensibly being served.

Who defines the good? The most consequential unasked question in "AI for good" initiatives is whose definition of "good" is organizing the work. In most cases, the answer is: the technology developers, the institutional funders, and the research institutions designing the applications, not the communities whose lives the applications are designed to affect.

This is not malice. It is the structural consequence of where AI development capacity, institutional funding, and decision-making authority are concentrated. The organizations with the capacity to develop sophisticated AI applications are overwhelmingly based in high-income countries, staffed by

people with advanced technical education, and funded by sources that reflect the priorities of technology industries and large foundations. The communities most in need of the applications (the ones facing disease burden, environmental risk, inadequate infrastructure, and economic precarity at the highest levels) are the communities least represented in the design and governance of the tools ostensibly built for them.

The result is a pattern that the development sector recognized decades before AI arrived: tools designed by people who are not from the communities being served, based on assumptions about community needs that are not adequately grounded in community knowledge, optimized for outcomes that the designing organizations can measure rather than for outcomes that communities actually want, and governed by frameworks that community members cannot meaningfully challenge.

What are the extraction dynamics? "AI for good" initiatives frequently involve the collection of substantial data from the communities being served: health data, behavioral data, location data, economic data. This data is the raw material from which the AI tools are built and refined. It is also, in the absence of adequate data governance, a form of resource extraction: communities providing the data that makes the tools valuable without adequate participation in how that data is used, how the tools are deployed, or who benefits commercially from the capabilities the data has enabled.

The data governance frameworks that would ensure equitable benefit-sharing, meaningful community consent, and community ownership of community data are in most "AI for good" initiatives either absent or inadequate. This is not unique to AI: it is a longstanding pattern in health research, environmental data collection, and development sector data use that AI has amplified by dramatically increasing the scale and value of data collection. Researchers studying data sharing in Africa describe the same recurring features: colonial power imbalances, insufficient trust, extractive practices, and Western-centric policy. Abeba Birhane calls the broader pattern algorithmic colonization: technology built elsewhere, often unfit for the local problem, that weakens local development and creates dependence on outside infrastructure. But the amplification is significant, and the "AI for good" frame has not yet adequately grappled with its implications.

What is the dependency being created? AI tools, once deployed in community-serving institutions, create dependencies: on the tools themselves, on the data infrastructure required to operate them, and on the technical capacity to maintain and update them. In well-resourced institutions with stable funding and internal technical capacity, these dependencies are manageable. In the community health centers, small municipal agencies, and underfunded public services that are the primary target of many "AI for good" initiatives, they represent real risks.

What happens when the foundation funding that supported the AI deployment ends? When the technology company that provided the tool at below-market cost changes its pricing? When the technical capacity that the community-serving institution does not have in-house is required to maintain or update the system? The communities most in need of stable, reliable, long-term services

are the communities least equipped to manage the infrastructure dependencies that AI deployments create, and the "AI for good" framework has not yet developed adequate standards for dependency risk assessment as a condition of responsible deployment.

What is being displaced? Every AI deployment displaces something, sometimes an inefficient process that deserves displacement, sometimes a human judgment capacity or community relationship that deserves preservation. "AI for good" initiatives tend to be clear about what the AI tool adds. They tend to be less clear about what it displaces, and whether the displacement serves the communities being ostensibly helped.

Community health workers who are replaced by AI symptom-checkers take their relationships, cultural knowledge, and trust connections with them when they leave. Public health professionals who cede disease surveillance to AI systems atrophy the observational capacity that makes them effective when the AI system encounters a situation outside its training distribution. Municipal staff who are displaced by AI-powered service delivery take the contextual knowledge and interpersonal relationships that made those services work for specific community members. These losses are rarely in the "AI for good" calculation. They should be.

The Case for Proper Technology Adoption

The proper alternative to "AI for good" as currently practiced is not "no AI": it is technology adoption that is adequate to the complexity of what AI systems actually are and what their deployment in community-serving contexts actually means.

Proper technology adoption for social good requires five commitments that the field is not yet making consistently.

Community definition of the problem. Before an AI tool is designed or selected, the community that will be affected by it should have genuine participation in defining the problem the tool is meant to address. Not participation in validating a problem that has already been defined by technology developers and institutional funders. Genuine participation in asking: is this the right problem? Is this the most important problem? Are there dimensions of this problem that our community experiences that the proposed technical solution does not address?

Transparent benefit-sharing. Any AI tool deployed in community-serving contexts should come with explicit, legally enforceable commitments about how the data generated by the deployment will be used, how the community will participate in any commercial value generated by the tool or the data, and what rights the community has to modify or terminate the deployment.

Dependency risk assessment. Before any AI tool is deployed in a community-serving institution, the institution and the community should have a clear understanding of what dependencies the deployment creates and what the consequences of those dependencies are under different scenarios:

funding changes, vendor changes, technical failures. Where the dependency risks are unacceptable, the deployment should not proceed.

Genuine evaluation against community-defined outcomes. AI tools deployed for social good should be evaluated against outcomes that the community has defined as meaningful, not outcomes that are convenient for the technology developer or the institutional funder to measure. This requires the development of evaluation frameworks with genuine community participation, and the institutional willingness to report honestly when tools are not producing community-defined outcomes even when they are producing the outcomes the vendor specified.

Sunset provisions and exit ramps. Every "AI for good" deployment should include clearly specified conditions under which the deployment will be modified or terminated, and a defined process for that modification or termination that does not leave the community-serving institution worse off than before the deployment.

The Technology the World Actually Needs

The world needs AI that makes public health surveillance more effective, disaster response more equitable, municipal services more accessible, and economic opportunity more broadly distributed. These are genuine needs, and AI has genuine relevance to them.

What the world does not need is "AI for good" as a branding frame that accelerates the deployment of sophisticated tools into complex community contexts without the assessment, governance, and accountability that responsible deployment requires, and then attributes the resulting failures to the complexity of the communities being served rather than to the inadequacy of the deployment framework.

The case for proper technology adoption is not a case against technology. It is a case for technology development and deployment that takes seriously the complexity of the communities it is designed to serve, the power dynamics embedded in who designs, who deploys, and who benefits from AI systems, and the governance requirements that genuine accountability for AI's consequences demands.

"AI for good" that meets this standard would look different from most of what currently travels under that name. It would be slower. It would involve more community participation and less technical elegance as the primary measure of success. It would be more honest about limitations, more explicit about trade-offs, and more willing to conclude that specific AI applications are not appropriate for specific community contexts regardless of their technical sophistication.

It would also be more effective: in the actual communities, for the actual problems, over the actual timescales that matter for people whose wellbeing is what "AI for good" is ostensibly about.

That is the technology adoption the world needs. And the field is not yet providing it at the scale the moment requires.

Where This Fits

The questions raised here are trust questions. Communities read the behavior of the institutions deploying technology on them exactly as they read any other institution, and the Links Global Trust Audit™ and the ERA™ white paper describe what happens when that reading turns defensive.

Schedule a conversation with Links Global: www.linksglobal.org

Sources

- Abebe, R., Aruleba, K., Birhane, A., Kingsley, S., Obaido, G., Remy, S. L., & Sadagopan, S. (2021). Narratives and counternarratives on data sharing in Africa. *arXiv*. arxiv.org/abs/2103.01168
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed, 17*(2). doi.org/10.2966/scrip.170220.389
- Wong, A., et al. (2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Internal Medicine, 181*(8), 1065–1070. doi.org/10.1001/jamainternmed.2021.2626

Ownership and use. © 2026 Links Global. All rights reserved. The Links Global Trust Audit™, Narrative Stewardship Infrastructure™, NSI™, Epistemic Re-Regulation Architecture™, and ERA™ are trademarks of Links Global. You are welcome to read this article, link to it, share it, and quote brief passages with attribution to Links Global and a link to this page. Reproducing it in whole or in substantial part, adapting or translating it, or using the adoption questions and assessment approach in training, consulting, publications, or derivative works requires prior written permission from Links Global. Please cite as: Sanchez, M. (2026). *AI for Good: The Case for Proper Technology Adoption, and the Questions "AI for Good" Isn't Asking*. Links Global. www.linksglobal.org.