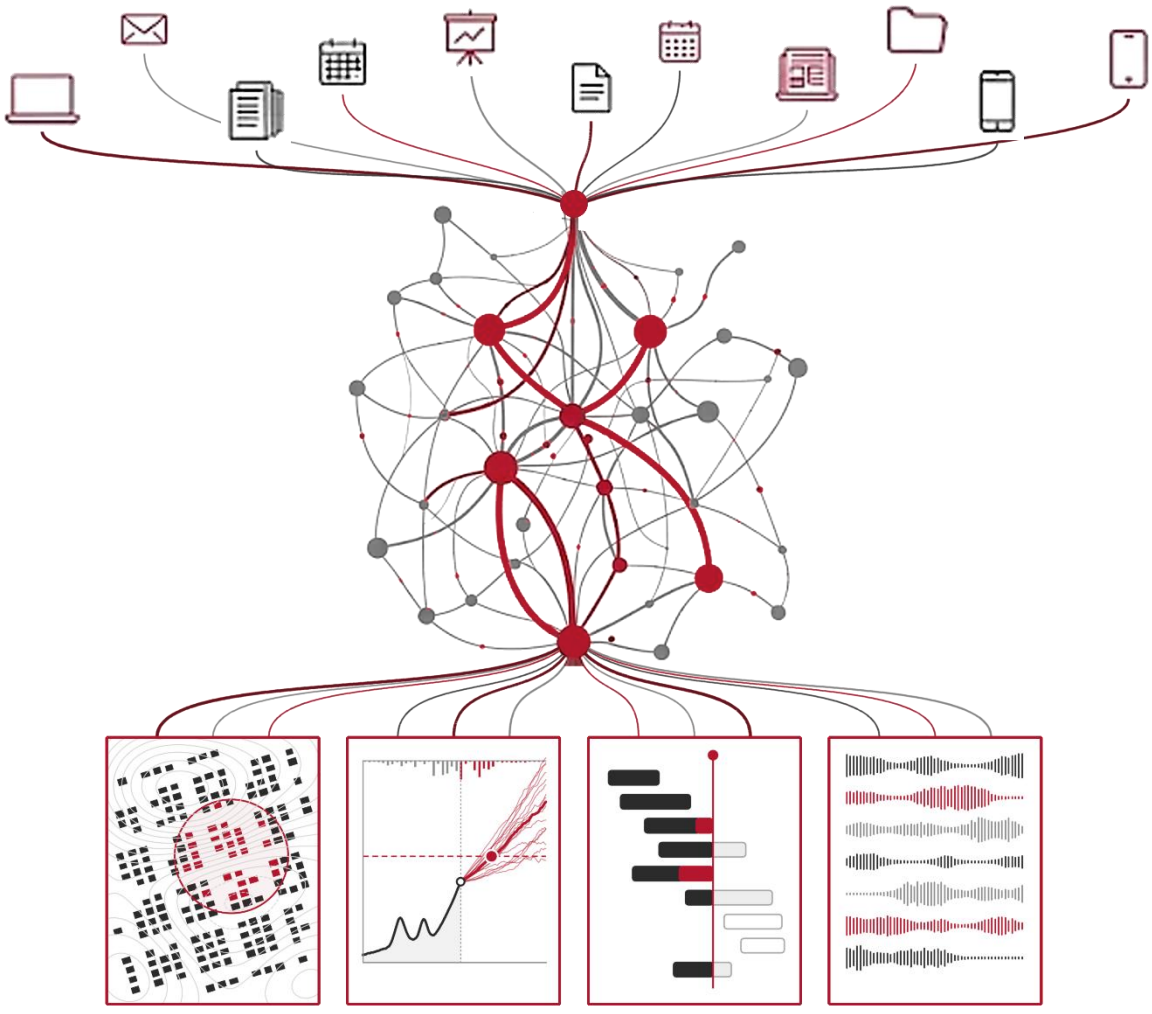
IFRC | solferino academy

From Everyday Tools to Enabling Systems

Leadership Choices for Humanitarian AI

Solferino Briefing Paper | November 2026 | Ben Ramalingam, Juriaan Lahr and Stefano Ongarello



01

Key Messages

1. Generative AI has spread faster than all previous general purpose technologies. The productivity gains are yet to be realised, in both business and in humanitarian work, where most use is in everyday internal work.

2. Published research shows humanitarian AI applications focus on needs and damage assessment, early warning, community listening and logistics. Fewer than one study in eight describes a system in real use, and co-design with affected communities is rare.

3. Technical limits and costs explain some of this limited progress. These include weak performance in many languages, scarce verified and relevant crisis data, high adoption costs and fast-changing skills requirements. The deeper explanation is organisational: as with other general purpose technologies, the value will come when organisations consciously redesign how they work with and around AI.

4. This redesign will involve changes in expertise, delegation, accountability and orchestration, all resting on trust and principled use. Leaders need to weigh the value of AI against its potential risks, robustly and continuously.

5. The paper sets out ten leadership choices across five domains, each linked to the Fundamental Principles: duty of care, crisis decision-making, accountability to communities, communications, advocacy and marketing, and finance and economics. These choices need to be made deliberately and at pace, by people working together across skills, disciplines and organisational boundaries.

02

How technologies change the world

Every general-purpose technology - electricity, the internet and other significant innovations that spread into every part of work and life - arrived first as an everyday tool that people picked up to solve immediate problems. **Each took then decades, and new ways of working, to be built into the enabling systems that changed our organisations, our economies and our world.**

People were using electric lights and trams long before electrification transformed production lines and factories, and reshaped modern capitalist production.

The internet was used mainly in military and academic circles. It then spilled out into general use in the 1990s, in the form of email and rudimentary websites. But it wasn't until the late 2000s, when broadband, the smartphone and the platform business model came together, that the internet rebuilt retail, news, music and the way people move, buy and communicate. This was several decades after the first embryonic online connections were made.

While artificial intelligence (AI) is far from new, its implementation in the form of everyday tools - specifically the various generative models - can be dated to the launch of Chat GPT in November 2022. Since then, generative AI has gone to scale at a speed without precedent, including in the humanitarian sector.

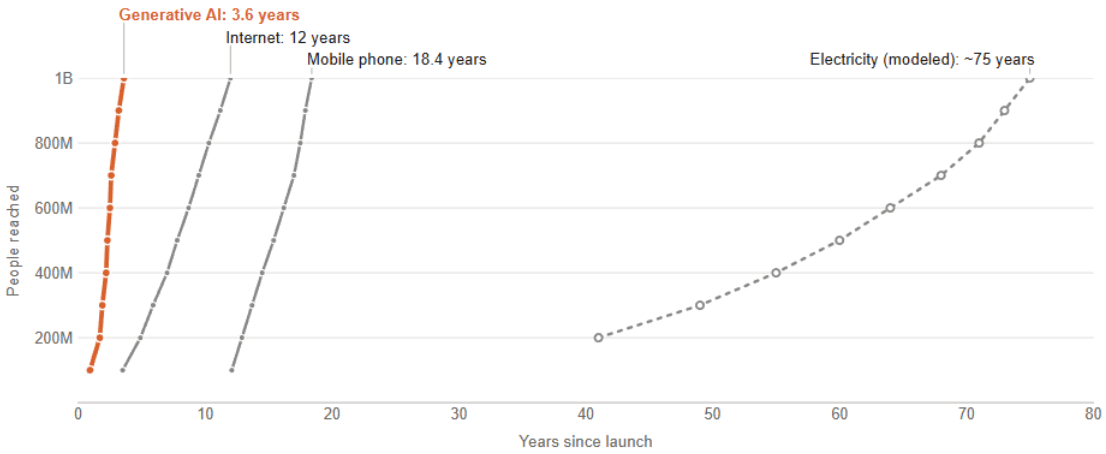
1 billion+

people using generative AI worldwide, in under four years

Generative AI is now well past 1 billion users worldwide in the four years since ChatGPT was launched. This is much faster than the growth of internet and mobile phone usage and many times faster than electricity (see Figure 1).

You would have to rewind to Pokémon Go to find comparable speed and level of dissemination, with half a billion downloads in its first two months post-launch in 2016. (At the risk of displeasing fans, that was just a game, and player numbers reportedly peaked at around 200 million within weeks and then fell sharply.)

Figure 1: Years to reach 1 billion users, for different technologies



Dates between data points are interpolated in a straight line. Generative AI: conservative estimate counting each person once (ChatGPT weekly users or Microsoft AI Diffusion, whichever is higher). Electricity has no global data before 1970; the dotted line is an S-curve fitted to 1970–2024 World Bank access data, extended back to 1920. Launch years: electricity 1882, mobile phone 1983, internet 1993, ChatGPT 2022. Sources: ITU, GSMA, World Bank, OpenAI, Microsoft. Several historical values are rounded estimates.

03

The AI productivity puzzle

But where has this accelerated uptake of everyday AI delivered? The jury is currently still out on this question.

For the most part, true positive gains are hard to find. In a survey of nearly 9,000 CEOs, CFOs and executives across seven countries published by NBER earlier this year, 83% reported no effect of AI on their productivity. Gartner global research tells a similar story: while the number of AI use cases per organisation has risen more than fourfold in two years, only one in five delivers a positive return.

In another survey, 40% of workers reported receiving low-quality AI generated work - referred to as "workslop"- in the previous month, with each instance taking nearly two hours to fix.

This is not unusual for general-purpose technologies. With computers, positive effects stayed small for over a decade after they reached office desks, prompting the economist Robert Solow's famous observation that the computer age could be seen everywhere except in the productivity statistics.

The humanitarian picture looks similar to the business one. A 2026 survey of over 1,700 humanitarians working in 120 countries found 75% using generative AI tools daily or weekly, but less than one in ten say AI is widely integrated in their organisation. While nearly two-thirds felt it had made their work more efficient, the reported uses were mostly office-based work: writing and editing, research and synthesis, summarising and translation.



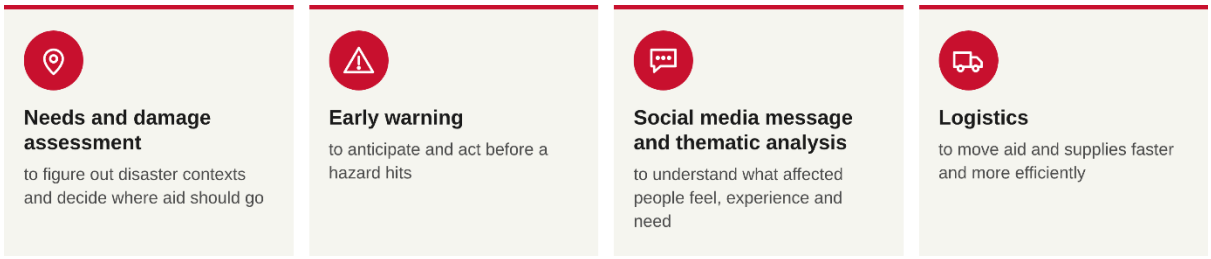
In other words, effort is focused on the everyday work, and this tells us little about the AI applications that could change the enabling systems of humanitarian work. The difference, as we see it, is whom the AI affects most.

Everyday AI helps staff with their own, internal, corporate work. AI in enabling systems shapes how humanitarian organisations engage with the people they serve or depend on - communities, donors and partners.

04

Beyond offices to frontlines

To understand these enabling applications better, the best evidence we have is from published research. Two major reviews published this year, in the Journal of Global Health and Technovation, map what AI applications have been built and tested in humanitarian delivery, and how much of this work has made it into use. These reviews find that frontline humanitarian AI is confined to relatively few domains shown below.



The applications covered focus on the preparedness and response phase of crises, with recovery almost untouched. Neither review looks at communications or fundraising, despite the growing use of AI to engage and communicate with donors and the public - an area where research has yet to catch up with practice.

Taken together, these applications reach into frontline humanitarian work and hint at what enabling systems could look like: assessing needs and damage, acting on better and faster analysis and predictions, listening to communities more systematically and at scale, redesigning work processes, and engaging donors and the public in new ways.

Very few of them, however, have made it beyond demonstrations and tests. Two-thirds of the studies were prototypes or validations against historical data, rather than real-time applications. Fewer than one study in eight described systems in real-world use. And every one of those was reported only by the organisation that had built or bought it - evidence of uptake beyond originating organisations was vanishingly rare.

And for the small number that do see use, there are further issues: co-design with affected communities appears in only around one in ten cases, and meaningful integration of human and AI capabilities is discussed in fewer than one in ten.

Research studies are only a partial proxy for AI use, since many applications are not documented and published, but they do give a reasonable sense of the state of the field. Despite the noise and light around AI in our sector, the overall picture is clear: **there are lots of AI applications as everyday tools to fix immediate problems, but very few that change the enabling systems of how we do our core work.**

05

The limits to progress

As with other technologies, technical issues can explain some of this limited progress.

AI models work badly in most of the languages in which humanitarian responders work: they are trained overwhelmingly on English and a handful of major languages and perform markedly worse in others. Connectivity and power are least reliable where need is greatest.

Verified crisis data is also scarce, and usually poorest in the places most at risk, which limits how useful AI can be. Most models are trained and tested on archives of past disasters, because data from a live crisis are fragmented and change regularly, sometimes in a matter of hours, which helps explain why so many studies stop at validation against historical data.

Closer to home, humanitarian organisations' own data often need investments to become a better enabler of AI: data management has rarely been an organisation-wide priority, partly because its value has been hard to see.

AI models also make errors, and an over-confident error in humanitarian triage is a very different matter from the same error in an email, no matter how senior the recipient list.

Then there are the tools and their economics. The cost of adopting AI is projected to be double than the cost of adopting traditional digital technologies. Many AI products are not yet ready to be deployed responsibly: vendors are building them as they go, and they are hard to integrate with existing systems and to monitor, especially for data protection and for cost. The full cost of running and evolving AI systems at scale - computing, connectivity, commercial agreements and skilled staff - is beyond what most humanitarian budgets allow.

Skilled staff are a constraint, because the roles and skills AI requires are changing faster than organisations can adapt. The total number of skills required for an average job is increasing at 4.8% annually, and 34% of the skills that were required for a job in 2025 will no longer be relevant by 2030. In the jobs most exposed to AI, the skills employers ask for are changing 66% faster than in others, and roles needing AI skills command a 56% wage premium. Humanitarian employers face this churn with human resources policies and salaries that cannot match the commercial sector. We will need to reskill staff, reshape teams and roles, and speed up the processes that grade, hire and redeploy people if we want to keep up with technological change.

The deepest part of the explanation for limited progress, however, is organisational, not technical. This is not unique to our sector. Electricity did not have a major impact on business productivity for around 40 years. The early applications simply replaced steam with electricity, keeping physical and organisational structures the same. Factories ran one giant electric motor, driving the same shafts and belts the steam engines had used.

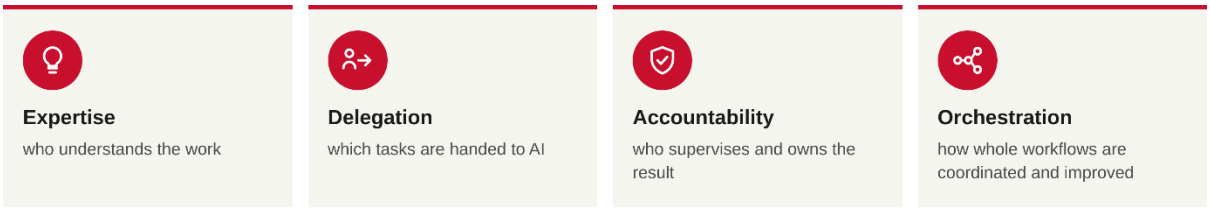
The true potential of electricity was only realised in the 1920s, and it took technical innovations - smaller motors - alongside organisational ones - the production line, factories redesigned around the flow of work and scientific management practices.

The humanitarian sector seems to be at a similar point in its engagement with AI. Most of the AI in use today is arguably the equivalent of swapping a large steam engine for a large electric motor. Writing emails, summarising meetings and drafting reports are all useful tasks, but they are not going to change the way humanitarian response works.

06

The choices facing leaders

The enabling applications described above, in needs assessment, early warning, community listening, logistics and donor engagement, are the equivalent of the small electric motors of the 1920s. They make the redesign of enabling systems possible but will only deliver on their potential if our organisations change around them. Four kinds of changes will be necessary.



Underpinning these changes are questions of trust, and of what responsible and principled use of AI looks like in humanitarian practice. None of these are technical questions but more about how we adapt to make best use of the opportunities new technologies bring.

AI makes a growing number of information tasks fast, easy and cheap. Leadership means deciding what these improvements are worth, weighing that value against the risks, and continuing to do so over time.

This last point is crucial because AI value tends to grow step by step while related risks can grow much faster. In a Gartner poll this autumn 86% of chief information officers said that AI risk in their organisation was growing around four times faster than AI value.

AI use also carries significant environmental costs. UN University researchers predict that data centres will more than double their electricity use by 2030 compared to 2025, and most of the energy AI uses going on answering everyday prompts. AI hardware depends on minerals often mined where environmental oversight is weakest. For humanitarian organisations, these costs should feature in responsible decisions about how AI is used.

Making leadership choices around AI means explicitly asking who gains what and who loses what, and being clear, transparent and principled about the trade-offs. Above all, it means making the choices consciously and deliberately, rather than on autopilot.

The table below sets out ten leadership choices across five domains, with key questions and the related opportunities and challenges for each. We have also drawn links between each domain and the most relevant Fundamental Principles of the Red Cross Red Crescent Movement. We call these leadership choices deliberately and are of the view that everyone who uses AI has a part in this leadership, wherever they sit in the organisation.

The domains begin with the people whose data and wellbeing are directly at stake, and widen step by step, through crisis decisions and accountability to communities, to communications with the public and donors, and finally to the economics that sustain the work.

All these choices reach beyond the technical teams and projects that build AI tools. Answering them requires co-creation between IT teams, AI and data specialists, and colleagues working on data governance, procurement, risk management, community engagement and accountability, public communications, law and humanitarian principles. It means treating AI tools as products, with an owner, defined users, systematic user feedback and budgets for continuous value improvement. It also means working across organisational boundaries too, with technology companies, researchers, governments and regulators, other humanitarian organisations and, above all, communities.

Of course, these are not the only leadership choices that matter for humanitarian AI, but they seem especially relevant from where we sit, working on digital transformation, strategy, foresight and knowledge in the IFRC. Colleagues elsewhere in the Network and the humanitarian system will see things differently, and we welcome feedback and debate.

We close with two points. First, on leadership. Project teams can build the tools that raise the questions, opportunities and challenges set out throughout this paper. But they cannot, on their own, change our enabling systems around the possibilities that these tools create. **This takes concerted, collaborative, cross-disciplinary and principled leadership. Treating these matters as deliberate choices, and making them transparently and accountably, is what will start to move humanitarian AI from everyday tools to enabling systems.**

The second point is about pace. Organisations had decades to adapt to electricity, computers and the internet. But AI is evolving and spreading too fast for that. **We will need to change in a matter of years if we are to shape how AI is used to enable our work. If we want to realise the potential benefits of AI for the people we serve, we cannot afford to be complacent.**

Leadership Choices

A number of challenges cut across all these domains, including overestimating what AI can do; unequal access, through the digital divide and the cost of tools; and the cybersecurity risks generated by AI systems.

Domain & Fundamental Principles	Specific Questions	Opportunities and Challenges
Duty of care Humanity and impartiality: protecting people's dignity, including what is done with their data. Aid targeted by need depends on inferences being accurate.	1. What inferences can and should be drawn from people's data about their capabilities and needs? *E.g. Existing data used to extrapolate needs and target aid*	**Opportunity.** Population conditions can be modelled from survey, phone and satellite data faster than formal assessments, so aid can be targeted and reach people sooner. **Challenge.** People are unlikely to have agreed to this use of their data, and an inference can be wrong about households and groups, with few corrective measures. Inferences drawn for one purpose can be used for others, including by actors beyond the humanitarian system. Beyond permissions, the question is whether systems are secure and the organisation keeps control of the data. Models also miss people who leave little or no data trail (see 5 below).
	2. What are the ethical, quality and safeguarding boundaries of AI assistance? *E.g. providing psychosocial care via chatbots*	**Opportunity.** Support available at any hour, in more languages, to people who would not be able or willing to access or attend formal facilities. **Challenge.** We need better evidence of who is helped, or harmed, when care can be provided without human oversight. The critical test is whether the system escalates when it should, whether human support is available, and what happens when the handover fails. Engagement and user satisfaction are no substitute for evidence of safety.
Crisis decision making Impartiality: decisions guided by need alone, with priority to the most urgent cases.	3. How far do we trust a forecast we must act on, and who answers when it is wrong? *E.g. releasing cash or supplies based on AI predictions before crises strike*	**Opportunity.** Acting before hazards hit saves lives and money, and AI-enabled forecasts make this possible at scale. **Challenge.** False alarms are visible while missed triggers are not, so accountability gets blurred. Forecasts can be confident but wrong, especially where data is weak or missing.
	4. Where do human oversight and judgement come in? *E.g. AI tools screen applications for assistance or prioritise emergency calls*	**Opportunity.** Translation, triage support and eligibility checks at a speed and volume no human team can match. Staff time freed for cases that most need human judgement. **Challenge.** Oversight drifts towards rubber-stamping, and such systems perform worst in the languages and places that already get least. Staff judgement can weaken with disuse.
Accountability to communities Impartiality and humanity: no one left out; respect for communities as participants and co-creators.	5. Whose voices count, and how do we know who is missing? *E.g. building up situational knowledge from social media or messaging data*	**Opportunity.** Social media and messaging analysis can show situational changes in close to real time. More people heard, more systematically, and at scale. **Challenge.** People with phones, data and speaking major languages are given prominence. Those left out face digital and other exclusions.
	6. How do communities get a say in the AI used on their behalf? *E.g. Locally led, user-driven AI applications*	**Opportunity.** Communities can shape what AI is used for, test it in their own languages and settings, and judge whether it works. Systems built this way are more likely to be used and trusted. **Challenge.** Involvement takes time and money that projects rarely budget for.
Communications, advocacy and marketing Neutrality and humanity: keeping the confidence of all, and portraying people with dignity.	7. What are the limits in communications, fundraising and advocacy, and how do we protect trust in real images of real crises? *E.g. AI-generated images or video in appeals*	**Opportunity.** Content can be generated in minutes, in every language, at a fraction of the cost. **Challenge.** Audiences overestimate their ability to tell real from generated, and fake examples weaken trust in real content.
	8. How far should donor modelling go before it hits ethical and reputational limits? *E.g. testing campaigns on AI simulations of real donors*	**Opportunity.** Campaigns tested before they are launched, and scarce resources routed to where they can raise most. Insight into what motivates donors at lower cost. **Challenge.** Donors agreed to give, but few consent or expect to be simulated. The commercial sector has faced numerous issues with consumer modelling, including major marketing failures and scandals.
Finance and economics Independence: keeping autonomy from commercial interests, so that gains serve people in need.	9. What should humanitarian organisations do with AI dividends? *E.g. staff time saved by AI automation of analogue processes*	**Opportunity.** Thousands of staff hours can be freed up within critical services, releasing time and money. **Challenge.** Without deliberate decisions, saved hours can disappear into existing workloads or turn into pressure to cut staff, and gains don't reach the people the organisation exists to serve. Few organisations measure the value and risk of AI well enough to know what the dividend really is.
	10. How do we build sustainable services without corporate capture? *E.g. using tools from technology companies and negotiating terms of use*	**Opportunity.** Companies offer AI capability, tools and ambition that humanitarians cannot match. **Challenge.** Vendors set the terms, the prices and the data arrangements. Humanitarians may not have leverage and negotiation skills needed to enable scale. Services can be withdrawn, repriced or changed at the vendor's discretion.

About This Paper

The IFRC Digital Transformation team, supported by the IFRC Solferino Academy, will explore the questions raised in this paper with leaders from across the IFRC network at the 25th session of the General Assembly in December 2026.

This paper was authored by:
Ben Ramalingam, Special Advisor, IFRC Solferino Academy Ben.Ramalingam@ifrc.org
Juriaan Lahr, Director, IFRC Digital Transformation Department Juriaan.Lahr@ifrc.org
Stefano Ongarello, Manager, IFRC Digital Transformation Department Stefano.Ongarello@ifrc.org

References

Aczel, M., Chamanara, S., Matin, M., Farsi, A., Marwala, T. and Madani, K. (2026). *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints.* UNU-INWEH Research Report. Richmond Hill, Ontario: United Nations University Institute for Water, Environment and Health, 3 June 2026.

Bresnahan, T.F. and Trajtenberg, M. (1995). General purpose technologies: "engines of growth"? *Journal of Econometrics*, 65(1), 83–108.

Brynjolfsson, E. and Hitt, L.M. (2000). Beyond computation: information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23–48.

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* New York: W.W. Norton.

Gartner (2025). 2025 Gartner Skills Landscape Survey.

Gartner (2026). *Gartner Says AI Risk is Rising Faster Than AI Value.* Press release, 19 October 2026.

Humanitarian Leadership Academy and Data Friendly Space (2026). *How are humanitarians using artificial intelligence? January 2026 pulse survey insights.* Briefing note, March 2026.

Jose, M., Mueen, F., Qadir, J., Abd-alrazaq, A., Kwalombota, K. and Maddah, D. (2026). From innovation to implementation: a global scoping review of artificial intelligence applications, governance, and challenges in humanitarian contexts. *Journal of Global Health*, 16, 04359.

Lythreatis, S., Acikgoz, F. and Yassine, N. (2026). Artificial intelligence in humanitarian aid: a review and future research agenda. *Technovation*, 151, 103415.

Niederhoffer, K., Kellerman, G.R., Lee, A., Liebscher, A., Rapuano, K. and Hancock, J.T. (2025). AI-generated "workslop" is destroying productivity. *Harvard Business Review*, September 2025.

PwC (2025). *The Fearless Future: 2025 Global AI Jobs Barometer.* London: PwC, 3 June 2025.

Solow, R.M. (1987). We'd better watch out. *New York Times Book Review*, 12 July 1987, p. 36.

Yotzov, I., Barrero, J.M., Bloom, N., Bunn, P., Chawla, C., Davis, S.J. et al. (2026). *Firm Data on AI.* NBER Working Paper 34836, February 2026 (revised October 2026).