ANALYSIS & COMMENTARY

Five AI Risks the Humanitarian Sector Isn't Talking About Enough

Adoption is outpacing governance — by a wide margin. Here's where the danger lies.

Aileen Voit · AI Safety & Governance for Humanitarian Organizations

Humanitarian organizations are adopting AI fast. Staff use it for translations, report writing, data analysis, and case management support. In most cases, they have no choice — workloads have expanded while teams have shrunk. AI fills the gap. But adoption is outpacing governance. By a wide margin. The Humanitarian Leadership Academy and Data Friendly Space surveyed 2,539 professionals across 144 countries and describe what they call the *humanitarian AI paradox*: individual adoption has outpaced organizational infrastructure.

95.3%	75%	30%	23.3%
have used AI for work	use commercial AI daily or weekly	now using AI agents	organizations have a formal AI policy

Source: Humanitarian Leadership Academy & Data Friendly Space — 2,539 professionals across 144 countries. Local and national organizations show the starkest gap: highest daily AI use rates, yet only 13% have an AI policy. 45% of staff report no AI policy exists at all. Only 7.8% of organizations report widespread AI integration.

The HLA/DFS research describes the landscape. But what does this adoption gap actually mean for the people humanitarian organizations serve? Having worked on protection, inclusion, and programme governance across multiple crisis responses, these are the five risks I see as most critical.

RISK 1

Invisible Exclusion

When AI is used for targeting, prioritisation, or vulnerability scoring, it can systematically exclude people without anyone noticing. An algorithm trained on incomplete data doesn't raise a flag when it drops someone from an assistance list — it just does it. In humanitarian contexts, that exclusion can mean someone doesn't receive food, shelter, or protection. The harm is silent and almost impossible to detect without deliberate oversight. WFP has already faced scrutiny over algorithmic exclusion errors in beneficiary targeting. The ICRC has put it clearly: AI systems trained on data from well-funded, digitally visible populations create blind spots that confuse historical visibility with actual humanitarian need — reinforcing cycles of over-recognition and neglect.

RISK 2

Data Exposure

Staff paste beneficiary information into ChatGPT, Claude, and other commercial tools every day — case notes, interview summaries, protection assessments. These tools were not designed for humanitarian data and were not built with the sensitivity of refugee and migrant information in mind. Once data enters a commercial system, the organization has lost control of it. 69% of humanitarian AI use relies on commercial providers, creating dependency on companies whose interests, data practices, and terms of service are not aligned with humanitarian principles. In a sector that handles some of the most sensitive personal data in the world, this should be treated as a serious protection risk — not a convenience question.

RISK 3

Loss of Human Judgement

AI decision-support tools are designed to make processes faster and more consistent. But humanitarian work — whether it's protection, MEAL, livelihoods programming, or shelter allocation — depends on contextual knowledge and professional discretion that no dataset can fully capture. A protection officer who overrides a standard output because they recognise a pattern from experience. A MEAL officer who questions an AI-generated analysis because they know the data collection context. A programme manager who adjusts a targeting recommendation because they understand

local dynamics the model doesn't. That is exactly the kind of expertise that keeps responses effective and people safe. When AI outputs become the default rather than one input among many, that expertise erodes gradually — not through a single decision, but through a slow shift in how teams relate to their own judgement. As the ICRC has written, AI systems are never fully neutral — they are designed, developed, and deployed in a specific context, by specific actors, and with specific purposes in mind. Across all areas of humanitarian work, the stakes of forgetting that are high.

RISK 4

Shadow AI

Most organizations don't know what AI tools their staff are using, what data is going into them, or what decisions are being informed by AI outputs. This isn't hypothetical — it is the structural reality of the sector right now. Access Now's March 2026 research shows how AI enters humanitarian operations not just through deliberate adoption, but through procurement workarounds, vendor updates, and embedded product features. Cloud-based AI is slowly but inexorably seeping into most internal systems through changes organizations never assessed or approved. Daily AI use is highest in Kenya (65%), Sudan (60%), and Bangladesh (59%) — often in settings where formal IT governance is minimal. Staff adopt whatever works because they're doing three jobs at once. There is also a stigma problem: in some contexts staff fear being seen as lazy; in others they worry about ethics being questioned. In both cases the result is silence. Organizations can't govern what they can't see.

RISK 5

Regulatory Exposure

The EU AI Act classifies migration and asylum as high-risk domains. There is no humanitarian exemption. AI systems used in asylum, visa, and residence permit processes face mandatory compliance requirements from August 2026. Penalties reach up to EUR 35 million or 7% of global turnover. For any organization receiving EU funding or operating in EU jurisdictions, this is not abstract. The EU Agency for Fundamental Rights has launched a dedicated 2026 research project on AI use in asylum and immigration procedures — a signal that regulatory scrutiny is tightening. The Act is strongest where states use AI in formal migration management, but much less explicit about humanitarian organizations using commercial AI in adjacent work — protection assessments, case management, translation, programming decisions. That leaves the greatest volume of real-world risk in the hands of organizations not directly captured by the most visible regulatory provisions. Most humanitarian organizations have not begun preparing for any of this.

What Needs to Happen

None of these risks require organizations to stop using AI. They require organizations to start governing it.

- Acknowledge what is actually happening — not what official policies say should be happening.
- Invest in guidance for the staff who are already using these tools daily.
- Build accountability structures that centre the people humanitarian organizations exist to serve.
- Ask hard questions: Who is using AI in your organization? What data is going in? What decisions does it influence? Who is accountable? Would the communities you serve consent to this?

There are encouraging steps at the organisational level — ICRC's AI policy, WFP's AI Sandbox, HLA and NetHope's AI Readiness Guide. But practical, field-friendly guidance that a protection officer or MEAL staff member can pick up and apply? That still largely doesn't exist. The humanitarian sector has already adopted AI at scale. What is missing is the institutionalisation of that adoption. That gap is where the risk sits.