ALNAP

BRIEFING

Explain

ESSENTIAL BRIEFINGS FOR HUMANITARIAN DECISION-MAKERS

AI IN THE HUMANITARIAN SECTOR

CLICK FOR A QUICK READ VERSION OF THIS BRIEFING

When discussing AI in the humanitarian sector, it's essential to separate potential from proven impact. While AI holds promise, decision-making must be grounded in evidence, distinguishing between applications that have been rigorously tested and those that remain speculative. The accessibility of generative AI (GenAI), amplified by Big Tech and media hype, has fuelled claims of its potential to boost efficiency and enable humanitarians to take on previously cost- and labour-prohibitive tasks like feedback collection and information sharing. However, a persistent knowledge gap on AI's challenges, implications, and proven benefits makes it difficult to assess these claims. Deployment conditions, training needs, and risk factors—such as privacy concerns—can further limit AI's value, cost savings, and applicability, challenging broad claims about its humanitarian impact. Given AI's varied applications, humanitarian decision-makers need knowledge on its actual, rather than theorised, impact, along with a clear understanding of budgeting for AI integration.

Humanitarians must stay informed about the evolving landscape of AI, recognising both its opportunities and risks. The key difference between generative AI (GenAI) and traditional AI is their function: GenAI creates new content based on trained data, while traditional AI analyses and automates. GenAI applications are broad but rely on shallow datasets, whereas traditional AI is more specialised, drawing from deep datasets. Popular GenAI tools like OpenAI's ChatGPT, Google's Gemini, Anthropic's Claude, and Hangzhou's DeepSeek use large language models (LLMs) trained on vast amounts of publicly available digital data. In the humanitarian sector, LLMs help synthesise information, generate reports, and support decision-making, particularly for language-based tasks, though some GenAI models also produce images and audio.

In contrast, traditional AI tools are trained on narrower datasets and focus on analysis and prediction, such as using machine learning to support aid targeting. While GenAI's ability to generate new outputs is promising, concerns around bias, transparency, safety, reliability, and accuracy remain. Its outputs may include incorrect information and misinformation, non-consensually sourced personal data, and proprietary information, raising ethical and practical challenges for humanitarian use.

Generative AI in the humanitarian sector is emerging, but not yet systematic. Current discussions on AI are often dominated by generative AI (GenAI), but other AI tools—such as machine learning (ML) and natural language processing (NLP) for sentiment analysis and image recognition—are also automating key processes and decisions. In practice, non-generative AI applications, particularly those using predictive analytics and machine learning, remain more prevalent.

Most humanitarian use of GenAI is limited to backend tasks like data collection, summarisation, and analysis, while other applications are still in testing, development, or pilot phases. Large-scale adoption remains distant, though interest is growing. Uncertainty around GenAI's actual benefits, concerns over risks, and a lack of technical expertise within humanitarian organisations have slowed its systematic integration. For now, organisation-wide use remains in its early stages.

When developing AI guidelines, humanitarians must consider the full spectrum of AI tools and their functions. Clearly defining where and how these tools are used in humanitarian programming helps assess their risks, benefits, and relevance in context.

1. **Efficiency gains in everyday tasks:** Most efficiency gains come from individual, ad-hoc use of commercial GenAI tools embedded in everyday platforms. Common tasks include drafting emails, note-taking, meeting summaries, report synthesis, social media posts, and copy-editing. Increasingly, GenAI is integrated into software like Power BI

and Excel to assist users. While some of these uses may become more formalised within institutions, tracking adoption remains difficult as most data is anecdotal. In this way, AI use is less humanitarian-specific and aligns more with general workplace automation trends.

2. **Augmenting institutional workflows:** This category involves a broader range of AI tools, including well-established non-generative AI, and represents more structured integration within institutions. AI is used systematically in two key areas: (1) back-end tools for internal processes and (2) automation of time-consuming humanitarian programming and service delivery tasks. These tools aim to drive efficiency at scale and are the most advanced in terms of testing, application, and humanitarian sector experience.

    - Data based functions improve internal knowledge management by cleaning data, making databases queryable, centralising information for easier processing, and streamlining task management. AI can also enhance supply chain management by integrating real-time data from multiple sources and enabling real-time tracking.

    - Big data analytics support MEAL tasks by extracting insights, synthesising data, processing surveys and feedback, generating metadata, and coding lessons learned. Catholic Relief Services (CRS) used machine learning in Malawi to proactively assess food insecurity risk.

    - Decision-making support is critical in complex humanitarian settings. AI can speed up information gathering, providing more accurate population density estimates to guide aid distribution. Semantic segmentation (a form of object recognition) helps assess building damage, directing rescue workers to overlooked sites. AI and ML also support vaccination programs and other medical applications. The European Crisis Management Laboratory has tested Large Language Models (LLMs) for extracting news, reviewing headlines, and compiling time-sensitive risk data.

3. **Enhancing community engagement:** AI tools are used for personalisation and interpretation, often integrated into existing platforms like WhatsApp and Facebook. Chatbots, widely promoted for improving communication, sharing information, and gathering feedback, vary in complexity—from simple decision-tree logic with controlled responses to more adaptive NLP- and GenAI-powered interactions.

    IRC's Signpost initiative has piloted a GenAI chatbot for information services but faced challenges in generating accurate, culturally appropriate responses in Farsi and Somali due to limited high-quality language data. These issues, along with oversight and quality control difficulties, highlight the trade-offs in AI tool sophistication. Other AI-driven engagement tools, such as sentiment analysis, help process and categorise community feedback from interviews, surveys, and focus groups more efficiently.

4. **Expanding humanitarian capabilities:** AI is enabling tasks once deemed unfeasible due to cost, personnel, or time constraints. AI-driven family reunification efforts, for instance, have shown promise in reducing search times for missing individuals. Anticipatory approaches that predict displacement could reshape humanitarian response by providing early insights into population movements. The Danish Refugee Council's Anticipatory Humanitarian Action for Displacement AHEAD model uses machine learning (ML) to anticipate displacement, while the WFP's Vulnerability Analysis and Mapping VAM tool applies ML to satellite imagery to forecast food shortages. However, these applications carry higher risks, involving predictive analytics, sensitive data (e.g., biometrics), and potential unintended consequences due to the novelty of their use.

AI tools pose direct risks, including misinformation, bias, opacity, and privacy concerns. They also introduce broader risks that impact the humanitarian mission, including by undermining the 'do no harm' principle. The lack of common standards further heightens these challenges.

- **Data gaps and misinformation:** Poor-quality data leads to inaccuracies, mistranslations, and misinterpreted nuances. Most AI tools are trained primarily in English and other data-rich languages (e.g., Spanish, Mandarin, French), leaving humanitarian contexts underrepresented. This lowers the accuracy of translations and information synthesis, potentially harming operations, creating mistrust, and endangering communities.

- **Bias and discrimination:** AI outputs often reflect biases embedded in dominant, English-language datasets. Lack of community input and local knowledge increases the risk of inaccurate or irrelevant results. Gender, racial, and socio-economic biases in large commercial AI tools can lead to discriminatory or harmful decisions.

- **Reliability challenges:** GenAI tools frequently generate false information or "hallucinations," leading to misinformed beneficiaries, service denials, or inaccurate reports. AI models also degrade over time (model drift), quickly becoming outdated and misinterpreting new data, potentially causing harm.

- **Privatisation and vendor lock-in:** Increased reliance on private AI providers risks undermining humanitarian neutrality.

- **Privacy and consent risks:** AI tools handling sensitive data may expose communities to security breaches. AI adoption pressures can lead to exploitative pilot programs in contexts where informed consent is already fragile.

- **Opacity and accountability:** AI's decision-making is often opaque, making it difficult for affected individuals to understand service denials or verification failures. Proprietary models limit transparency and oversight, while open-source tools, though more transparent, can be adapted for harmful purposes. Humanitarian actors often lack control over AI-owned data, weakening accountability.

- Human Rights **Implications:** AI's resource demands strain energy and infrastructure, while its labour-intensive data labelling process is linked to exploitation. AI-driven misinformation further complicates humanitarian efforts.

The rise of GenAI has raised the stakes for humanitarians, making shared standards more urgent than ever. The rapid evolution of AI, especially generative AI, has sparked growing interest in the humanitarian sector, yet its use remains largely unchecked, lacking clear safeguards. Previously, AI tools were deployed for specific, resource-intensive tasks, such as Save the Children's NLP tool for detecting grooming in online chats. Now, general-purpose GenAI tools offer user-friendly access to powerful capabilities, driving increased humanitarian adoption—and heightened scrutiny. Without a shared framework, AI use remains informal, experimentation is ad hoc, and learning is siloed.

To help humanitarians make informed decisions about AI-related benefits and risks, rigorous research and shared learning opportunities are needed to assess broad claims about its utility and costs in humanitarian operations. Clear boundaries for GenAI use are especially critical, as most tools are not designed or tailored specifically for humanitarian needs.

- **Stronger research on GenAI in humanitarian contexts:** Humanitarians need applied knowledge and critical literacy to assess GenAI's risks and benefits. Key concerns include compromised beneficiary data in LLM training and the lack of transparency in AI-driven aid decisions. Currently, organisations are struggling to grasp GenAI's full implications.

- **Understanding costs, skills, and risks:** Organisations must evaluate the financial and human resources needed to develop, assess, and deploy GenAI tools. Tailoring AI for humanitarian work requires ongoing testing, yet returns remain uncertain. Given funding cuts, agencies must weigh whether investing in AI experimentation is justified when costs, benefits, and risks remain unclear.

- **Supporting staff in responsible AI use:** Humanitarian staff are already using GenAI for email drafting, information synthesis, and report reviews. Organisations must equip them to critically assess AI-generated content, avoid inputting sensitive data, and recognise AI's limitations. Clear guidance and training are essential to ensure responsible use.

As humanitarians navigate responsible AI adoption, holistic assessment approaches can help capture the full scope of ethical concerns and risks. Key considerations should guide their growing scrutiny of AI tools.

1. **Keep communities at the centre:** AI adoption must not sideline the people it aims to serve. Humanitarians should prioritise participatory design, community consultation, and inclusion. Off-the-shelf AI tools may limit opportunities for meaningful community involvement in their development, raising concerns about accessibility and relevance.

2. **Establish clear standards and policies:** Defining acceptable AI use is essential. Humanitarian actors can reference technical standards (ISO/IEC 42001), the ICRC's AI policy, and NetHope's Humanitarian AI Code of Conduct for guidance. Policies should outline ethical principles, reevaluation conditions, red flags for review, testing protocols, and clear boundaries on AI use.

3. **AI is not always the answer:** Not every problem requires AI. Decision-makers should assess whether AI genuinely adds value or introduces unnecessary complexity. GenAI's broad applicability can lead to solutionism—deploying AI where simpler, safer solutions exist. Focusing on well-defined, lower-risk applications is often the better approach.

4. **Understand and assess risk:** Risk assessment is fundamental to humanitarian work and must extend to AI tools. AI systems evolve, improve, or degrade over time, affecting their risk profiles. Organisations should scrutinise AI claims, realistically assess their capacity to audit tools, and establish a culture of continuous evaluation.

5. **Account for resource demands:** AI carries significant costs—financial and otherwise. Organisations must factor in data collection, risk assessments, technical expertise, procurement, maintenance, and testing to ensure security and reliability. Overlooking these costs can lead to ineffective or unsustainable AI adoption.

6. **Share knowledge and learnings:** The humanitarian sector has a knowledge gap in AI use. Openly sharing experiences, challenges, and validated use cases will foster a more informed and responsible approach to AI adoption.

7. **Build, buy, or borrow?** Organisations must decide whether to develop bespoke AI tools, use commercial solutions, or leverage open-source models. Each option comes with trade-offs in security, customisation, and ethical considerations. More cross-sector discussion and best practice sharing are needed.

8. **Develop a humanitarian-specific AI ecosystem:** The sector should explore AI solutions tailored to humanitarian needs, such as Small Language Models (SLMs) trained on verified humanitarian data. SLMs, which can be hosted locally, require fewer resources, consume less energy and water, and allow greater control over privacy, governance, and bias mitigation.

Interest in GenAI is growing amid tight budgets, but humanitarians must ensure its adoption is driven by community needs, not just cost-cutting pressures. The push to "do more with less" risks shaping AI use in ways that conflict with humanitarian principles. Humanitarians must also learn from past innovation roll-outs when applying AI. Making responsible AI decisions requires humanitarians to critically assess claims about GenAI's potential to enhance efficiency while weighing the risks and broader implications—especially for crisis-affected communities. No matter how advanced the technology, humanitarian work must keep communities at the centre of decision-making. While AI may be new to the sector, the challenges of technological innovation are not. The humanitarian field has already built knowledge on the role of innovation, the risks of technology, and the "do no harm" principle. Longstanding concerns around data quality, protection, and literacy remain critical in AI assessments. Safe, responsible AI depends not just on the technology itself but on the broader systems that support its ethical use. Even as AI evolves, humanitarians must not lose sight of the lessons already learned.

About *Explain*

The greatest learning challenge for our sector is less about capturing lessons and experiences, but creating spaces for humanitarians to absorb and act on what is already known.

Decision makers – at all levels – are struggling to find time to engage with vital new learning and evidence.

That's why ALNAP is piloting new approaches to communicating knowledge tailored to the needs, expectations and preferences of the busiest humanitarians.

ALNAP's Explain series offers straightforward communications to help humanitarian decision-makers make sense of, and exchange on, current evidence and discourse. Key learning and links all in one place: sourced, checked and curated by ALNAP's respected global research team and subject experts from across the sector.

Explain briefings bring you up to date on the latest developments in humanitarian action and the implications for your work, providing you with a better understanding of what's out there and what you need to know.

***Explain:* simple communication, sense-making, exchange of experience.**

For more information or expert comment please contact ALNAP's Communications Team: comms@alnap.org