

NETHOPE

Harnessing AI for Humanitarian Impact: Lessons and Insights from 11 Case Studies

Anticipating Needs & Directing Aid



Forecasting Displacement

The Danish Refugee Council uses AI models to predict forced displacement from conflict and climate.



Anticipating Floods for Early Action


GiveDirectly uses AI-powered flood forecasts to deliver cash assistance before disasters strike.



Accelerating Crisis Mapping

The fAIR tool nearly doubles mapping speed for community volunteers identifying buildings and roads.

Managing Information & Supporting Communities



Analyzing Unstructured Crisis Data



Combating Health 'Infodemics'

The Africa Infodemic Response Alliance (AIRA) uses AI to counter harmful health misinformation online.

Informing Displaced People

Signpost AI pilots a chatbot to provide refugees with vital, localized answers in their language.



NetHope Center for the Digital Nonprofit

Meheret Takele Mandefro

Business Analyst, NetHope Center for the Digital Nonprofit

March 2026

Table of contents

- Executive summary2
- 1. Introduction4
- 2. Overview of the 11 AI case studies5
- 3. Thematic analysis and cross-cutting learnings.....6
- 4. Looking ahead: strategic imperatives 10
- Works cited 12
- Annex: glossary of AI terms..... 13

CDN partners

We acknowledge the confidence our esteemed partners place in our research and programming. Their investment in the pursuit of knowledge and social betterment reflects a shared commitment to make a positive difference in our world.



Special thanks to all the organizations and individuals who generously contributed their experiences and insights to these case studies.

Executive summary

The nonprofit sector has entered a defining phase in its adoption of artificial intelligence. As the scale and complexity of global crises intensify and resources tighten, organizations are moving beyond asking whether AI can help and toward determining how to make it work sustainably, equitably, and at scale.

This briefing synthesizes learnings from 11 case studies of AI initiatives implemented by organizations, documented through a partnership between [NetHope](#) and the [UK Humanitarian Innovation Hub throughout 2024-2025](#). The cases reflect a wide spectrum of AI applications, from predictive analytics for displacement forecasting to AI-assisted mapping, from chatbots serving displaced populations to tools that help staff retrieve institutional knowledge more efficiently. All findings reflect the state of each initiative at the time of documentation and may have evolved since.

Taken together, the case studies offer three major contributions: a grounded view of what AI is already achieving in humanitarian contexts, an assessment of challenges encountered, and practical recommendations for organizations at various stages of their AI journey.

Key findings from across the case studies include:

- **AI delivers measurable impact but requires realistic expectations:** AI is delivering operational impact e.g., 80% faster building identification for disaster mapping (fAIr), 83% accuracy in flood predictions used to trigger anticipatory cash transfers (GiveDirectly), 1.9M users served (Signpost AI), but demands realistic expectations, as performance varies by language/geography and augments, rather than replaces, human expertise.
- **Data readiness trumps model sophistication:** High-quality, structured, multilingual/Global South datasets with robust governance determine success far more than cutting-edge technology, positioning data infrastructure as the sector's critical investment gap.
- **Sustainability hinges on new funding paradigms:** Short-term grants mismatch AI's ongoing retraining, infrastructure, and maintenance needs, stranding proven tools like DEEP and risking widespread pilot abandonment. Without multi-year operational funding for digital tools that have already demonstrated value, the sector will continue to lose effective solutions.
- **Local adaptation requires both technical and community governance:** Off-the-shelf English-centric models falter without custom language, as evidenced by CRS's Arabic models and Signpost's 40+ language localizations. But localization is not only a technical task. Without community ownership over data, model design, and validation, technical fine-tuning risks producing tools that are technically localized but still misaligned with local realities and decision-making structures. Effective localization requires communities to shape how tools are built and governed, not just supply training data.

- **Technical capacity shortages undermine long-term viability:** Initial external partnerships enable deployment, but enduring viability requires sector-wide upskilling to reduce dependency risks.
- **Open-source collaboration requires sustained coordination:** Open source brings transparency and efficiency, but it also requires sustained coordination, documentation, and community support to be usable beyond isolated projects. While it removes licensing costs, it shifts the burden toward maintaining, adapting, and governing the tools. For many resource-constrained NGOs, this coordination effort can equal or exceed the cost of commercial options, making it a critical but often overlooked part of total cost of ownership.
- **Responsible AI governance is emerging but structurally fragmented:** Responsible AI governance shows promise (Mercy Corps Ethics Framework, Signpost Quality Assurance process) but they remain isolated. This reflects the broader pattern identified in our [Missing Middle analysis](#): governance sits in disconnected layers, with global principles above and organization-level policies below, but little sector-wide alignment in between. Without shared standards to bridge this gap, responsible AI practices will continue to vary widely across organizations.

These patterns are reinforced by evidence from a cohort of organizations applying for UKHIV AI Scale-Up Grants in early 2026, representing over 30 initiatives across humanitarian, conservation, and development sectors. The applications' distribution (70% pilot-to-scale, 15% scale-to-socialize, 15% earlier stages) underscores a sector maturing from experimentation to operationalization, while still confronting the same data, sustainability, localization, capacity, and governance barriers.

Key recommendations for the nonprofit sector:

- **For organizations starting their AI journey:** Focus on clearly defined, high-value problems, start with small pilots to learn and iterate, invest early in strong data governance, build strategic external partnerships while growing internal capacity, and co-design solutions with end users from the outset.
- **For organizations scaling AI:** Plan for financial and operational sustainability from the start, openly document and share learnings, invest in long-term skills and capacity, consider open-source and self-hosted options to reduce lock-in and costs, and embed robust ethical and quality governance frameworks to guide responsible use.
- **For donors and funders:** Provide flexible, multi-phase financing that supports iterative AI development, embeds sustained capacity building and maintenance, incentivizes structured cross organizational learning, and robustly resources localization of data, models, and community led AI solutions.
- **For the sector as a whole:** Develop shared standards and frameworks for responsible AI, data governance, and quality assurance, establish communities of practice for practitioner collaboration, pool resources to tackle common challenges like local language support and Global South data, and collectively advocate with policymakers and technology firms to embed humanitarian principles in AI governance.

1. Introduction

Nonprofit organizations have recently faced an increasing number of challenges. In 2024, the number of forcibly displaced people [exceeded 120 million worldwide](#), while climate-related disasters, conflicts, and health emergencies continue to [affect hundreds of millions more](#). In any given emergency, relevant information is spread across hundreds of documents, websites, and internal systems, all unstructured data that must be transformed into timely insights for decisionmakers.

Amid this complexity, artificial intelligence has emerged as both a promise and a source of uncertainty. The key questions remain: How can AI help nonprofits extend their reach despite limited resources? What governance needs, ethical considerations, and risks must be addressed?

Across the sector, organizations are piloting and scaling AI tools in humanitarian operations. [NetHope's Center for the Digital Nonprofit](#) has been tracking this development closely, developing tools and resources to help our members and the broader non-profit community navigate AI use responsibly and for impact. Through years of work on [AI readiness](#), creating the sector's first [governance framework](#) and capturing [practical implementation](#), such as the [AI Suitability Toolkit for Nonprofits](#) and recent guidance on the [EU's General-Purpose AI Code of Practice](#), NetHope has supported nonprofits as they navigate responsible adoption.

To deepen this understanding with real-world evidence, NetHope partnered with the [UK Humanitarian Innovation Hub](#) to document [case studies](#) from organizations already deploying AI. These initiatives span displacement forecasting, building mapping, refugee information services, misinformation response, knowledge management, market monitoring, and vulnerability assessment. Together, they offer a grounded view of what works, what doesn't, and what the sector must strengthen to use AI responsibly and effectively. Rather than a formal impact evaluation, this briefing synthesizes qualitative and self-reported operational insights across cases to identify common patterns, challenges, and practical recommendations.

This briefing is designed to inform the nonprofit sector, technology practitioners, program teams, and funders seeking practical, field-tested insights into AI adoption and to learn directly from early adopter organizations.

2. Overview of the 11 AI case studies

The case studies span a diverse range of organizations, AI technologies, and humanitarian applications. The table below provides a summary overview of the 11 documented initiatives.

Table 1: Overview of case studies reviewed

Organization	Initiative	Primary Use Case
International Rescue Committee (IRC)	Signpost AI	Information services chatbot
Danish Refugee Council (DRC)	Displacement Forecasts	Displacement prediction
Humanitarian OpenStreetMap (HOTOSM)	fAIr	AI-assisted mapping
Pacific Disaster Center (PDC)	AI for Humanity	Disaster early warning
WHO (AIRA)	Africa Misinformation Portal ¹	Health misinformation tracking
Norwegian Refugee Council (NRC)	Knowledge Retrieval Chatbot	Internal knowledge retrieval
Mercy Corps	Safe Gen AI Chatbots	Internal staff productivity tools
Catholic Relief Services (CRS)	Market Monitoring Dashboards	Food price forecasting
International Committee of the Red Cross (ICRC)	POMELO	Population mapping
GiveDirectly	AI Cash Triggers	Anticipatory cash transfers
DEEP Consortium (currently closed)	DEEP Platform ²	Secondary data review and analysis

¹ Note: Africa Misinformation Portal (AMP) platform is currently under maintenance and may not be fully operational.

² The DEEP humanitarian data analysis platform was officially discontinued in 2024 after the end of its funding cycle and the transition of its core partners to other analytical tools. [Data Friendly Space](#).

3. Thematic analysis and cross-cutting learnings

AI adoption drivers: Nonprofit organizations are adopting AI to address pressures across both internal operations and frontline service delivery: amplifying scale to meet Increasing humanitarian needs, accelerating speed for crisis response (e.g., GiveDirectly's 48-hour flood cash triggers), boosting efficiency (Mercy Corps' 40-60% document synthesis gains; Humanitarian OpenStreetMap Team's (HOT) doubled mapping throughput), enhancing precision (ICRC's POMELO tool estimates population density at 10-meter resolution, enabling aid targeting at the neighborhood level) and expanding accessibility across 200+ languages (AIRA, NRC chatbots). These drivers converge on a singular insight: AI excels at transforming unstructured data overload into actionable insights, but only when matched to acute bottlenecks rather than pursued as an end in itself. Across the cases, the most impactful deployments targeted specific bottlenecks: overwhelming volumes of unstructured crisis data (DEEP, AIRA), slow manual mapping processes (fAIr), delayed disaster response timelines (GiveDirectly), and limited multilingual information access for displaced populations (Signpost AI).

Core application categories: The AI applications in the case studies cluster into four categories where AI demonstrated the most concrete operational value. These are:

Table 2: Overview of AI application types

Category	Definition	Use cases
NLP extraction from documents/ social data	Automated extraction of structured information from unstructured text sources including field reports, surveys, and social media feeds.	DEEP: Data extraction platform AIRA: Health misinformation tracking
Predictive forecasting	Machine learning models that anticipate future conditions - displacement, market prices, or disaster risk - from historical data.	DRC displacement models: Foresight, AHEAD, SODRD, SPIN CRS market monitoring: Food price forecasting dashboards PDC DisasterAWARE: Disaster early warning GiveDirectly AI triggers: Anticipatory cash transfers
Geospatial mapping	AI-assisted analysis of satellite imagery and geographic data to map populations, infrastructure, and crisis-affected areas.	fAIr: AI-assisted mapping tool (HOT) POMELO: Population estimation (ICRC)
Generative AI chatbots	LLM-powered conversational agents that provide information, referrals,	Signpost AI: Community info service (IRC) NRC chatbot: Knowledge retrieval

and support to affected populations or staff in local languages.

Mercy Corps chatbots for Internal staff productivity

These categories are not intended to suggest that any single approach is universally dominant; rather they map where AI has shown concrete operational value across the documented cases. The relative importance of each category depends on an organization's priorities and context, e.g., actors focused on rapid decision-making may prioritize predictive forecasting and knowledge-retrieval, while logistics- and operations-focused teams may derive more value from geospatial mapping and market monitoring.

Responsible governance practices are maturing but fragmented: Organizations featured in the case studies apply varying approaches, such as Co-created ethics principles developed with 50+ stakeholders, self-hosted infrastructure to reduce vendor lock-in and protect organizational data (NRC), human-in-the-loop safeguards maintained across all initiatives, local dataset bias correction through community-generated training data (fAIr), and a defined quality assessment framework for chatbot outputs (Signpost AI). Yet without sector-wide standards to unify these efforts, ethical scaling risks stalling amid regulatory divergence and, critically, the field still lacks a shared framework for determining accountability when AI-informed decisions cause harm. Whether responsibility sits with the developer, the implementing organization, or the funder remains unresolved, underscoring why clearer, sector-level norms are urgently needed.

Cross-cutting learnings

What Works

Collaboration accelerates impact: The most successful AI initiatives across all case studies leveraged strategic partnerships that combined humanitarian expertise, technical prowess, and local knowledge. ICRC partnered with the Swiss Federal Institute of Technology Lausanne (EPFL) and ETH Zurich to develop POMELO's high-resolution population estimates; fAIr worked with Masaryk University for scientific validation while centering community-generated training data; DRC collaborated with IBM on displacement forecasting; and GiveDirectly teamed with Google Research for flood prediction. Multi-stakeholder consortia like AIRA's 50+ organization network and Signpost AI's local moderation teams further demonstrate how shared ecosystems prevent duplication and ensure contextual relevance. Partnerships accelerate impact, but several cases show these gains require deliberate sustainability planning and capacity transfer to avoid leaving operational maintenance as an unresolved risk (see Persistent Challenges).

Problem-led design delivers results: Impactful deployments began with clearly defined operational pain points rather than technology experimentation. CRS identified food procurement forecasting as a priority need before building ML models; PDC positioned AI to enhance existing disaster coordination rather than as an end in itself; and NRC's use case assessment prioritized knowledge retrieval for its immediate program benefits. This disciplined approach guarantees AI addresses genuine gaps and produces measurable impact. However, it is important to acknowledge that the problems selected for AI experimentation are rarely the most structurally complex humanitarian

challenges. The evidence base naturally reflects tractable, well-bounded use cases, which may inadvertently over-represent AI's effectiveness relative to systemic issues that resist clear problem definition.

Human oversight safeguards reliability: Every initiative maintained rigorous human-in-the-loop protocols: fAIr mappers validate AI-generated features; PDC uses AI-generated risk assessments to inform analyst workflows, but human analysts make all final alerting and escalation decisions; NRC experts verify chatbot outputs against organizational policies before adoption; and GiveDirectly refines forecasts with local data collection. Human judgment remains irreplaceable for ensuring accuracy, ethics, and contextual nuance in high-stakes humanitarian contexts. It also guards against automation bias and cognitive offloading, where users may uncritically accept AI outputs, a growing concern as these tools become more embedded in routine workflows.

Open-source transparency builds sector momentum: fAIr and POMELO published code on GitHub, while Signpost AI shares operational learnings through public blogs. These practices demystify AI systems, enable peer review, and allow resource-constrained nonprofits to adapt proven solutions rather than reinventing from scratch, creating compounding returns across the ecosystem. It is worth noting, however, that open-source does not eliminate cost, it redistributes it. Licensing fees are replaced by coordination demands: maintaining codebases, documenting tools, onboarding contributors, and governing shared infrastructure. For resource-constrained NGOs, these coordination costs can equal or exceed the price of a commercial license. The calculus shifts only when human coordination resources are abundant, as in volunteer-driven communities, but this is rarely the default in operational humanitarian contexts. Open-source is therefore best understood not as a cost-reduction strategy but as a cost-reallocation strategy, one whose net benefit depends heavily on the human infrastructure an organization can bring to bear.

Continuous evaluation strengthens performance: Several organizations embedded structured evaluation into their AI workflows. Mercy Corps employs adversarial 'red teaming' to stress-test chatbot outputs for safety and accuracy, while GiveDirectly continuously recalibrates flood forecasting triggers against observed outcomes. Signpost AI similarly tracks user satisfaction and response quality across its 17 country's deployments. This commitment to ongoing evaluation, rather than one-time testing, helps ensure AI tools remain fit for purposes as conditions change.

Persistent challenges

Technical capacity gaps threaten sustainability: ICRC struggled to maintain POMELO's specialized machine learning skills internally due to the scarcity and cost of such expertise in the humanitarian sector, resorting to contractors, Mercy Corps faces AI talent retention in competitive markets; and Signpost AI required extensive team upskilling. Without systematic capacity building, even proven tools risk becoming dependent on external expertise that may vanish when funding cycles end.

Data limitations constrain effectiveness: DRC lacked sufficient historical data for robust displacement modeling; ICRC's POMELO accuracy drops dramatically without census inputs common in crisis zones; and fAIr scalability hinges on costly, unevenly available geospatial imagery. These

gaps reveal structural deficits in humanitarian data infrastructure that no amount of model sophistication can overcome.

Localization failures undermine relevance: Global North training biases degrade fAIr performance outside familiar contexts; Signpost AI struggles with high-quality Farsi and Somali outputs; AIRA requires local triangulation for cultural nuance; and GiveDirectly must recalibrate flood models for geographic variation. Localization, whether through fine-tuning existing models with local-language data, incorporating community-generated training datasets, or establishing continuous feedback loops with local users, isn't optional. It is foundational for operational impact. Crucially, localization must not be reduced to a technical fix. Fine-tuning models for local languages is necessary but insufficient if the deeper questions of ownership and governance remain unanswered: who controls model development, who owns training data, and whose knowledge frameworks get encoded? Technical localization without community governance remains a top-down model, one that risks replicating the same extractive dynamics it claims to address. True localization requires communities to participate not just as data sources but as decision-makers in how tools are designed, validated, and governed.

Funding mismatches strand proven pilots: DEEP's closure is the clearest case in point: a widely used humanitarian analysis platform, genuinely valued by its users, discontinued not because it failed but because its funding did. This is not an isolated sustainability challenge. It is a structural indictment of how the sector finances digital infrastructure. CRS, NRC, and PDC are navigating the same pressure in real time: dashboards needing maintenance budgets, scaling requiring enterprise infrastructure, cost control forcing self-hosted workarounds. These are symptoms of a single structural flaw: the sector is using project-cycle grants to build what functions as long-term digital public infrastructure, financing roads like meals, expecting one-time investment to sustain something that requires perpetual upkeep. Until funders shift to multi-year operational endowments for proven tools, this pattern will repeat.

Scaling demands organizational transformation: CRS hit infrastructure barriers expanding across countries; ICRC faced deployment hurdles moving POMELO from academia to operations; and Signpost AI extends pilots through 2025 for maturity assurance. Technical readiness alone fails, successful scaling requires aligned leadership, change management, and enterprise architecture.

Trust gaps slow adoption: DRC users resisted "black box" predictions despite statistical validity, while fAIr invests heavily in explainability training. Technical excellence without transparency, user engagement, and demonstrated reliability breeds skepticism that stalls deployment even when solutions work. Trust resistance should also be read as a diagnostic signal rather than merely a deployment barrier. Frontline staff and affected communities bear the consequences of model errors, and their skepticism toward opaque systems reflects a legitimate institutional safeguard. This raises a deeper governance question: whether communities should have a meaningful say in whether and how predictive models are used at all.

4. Looking ahead: strategic imperatives

The documented case studies mark a promising early chapter in AI's humanitarian evolution but scaling impact demands deliberate next steps. NetHope is here to help nonprofits pivot from isolated pilots to ecosystem-wide systems, confronting four interconnected priorities with clear operational focus.

Operationalize beyond pilots: AI has proven its value, but transition bottlenecks persist. Funders must shift to multi-year operational budgets covering retraining, infrastructure, and monitoring, recognizing AI as a living infrastructure rather than one-off projects. Without this, proven tools risk abandonment, squandering sector momentum.

Center global south realities from day one: English-centric, Global North-biased models consistently underperform in target contexts, as seen across languages (Signpost's Farsi/Somali gaps) and geographies (fAIr's regional accuracy drops). Prioritize local-language datasets, regional partnerships, and community-led data collection and model adaptation by embedding localization in design, not as an afterthought, to ensure relevance where needs are greatest.

Build collective capacity through sector wide standards: No single nonprofit can manage AI's full stack amid talent shortages and resource constraints. The sector needs shared enablers, upskilling consortia, communities of practice, interoperable governance frameworks, and unified metrics for data sharing and evaluation. This aligns with our recent [AI Governance and the Nonprofit Sector: Mapping the Missing Middle](#) analysis, which shows that governance currently operates in disconnected layers: broad global principles above, organization-specific policies below, and an underdeveloped sector-wide layer in between. Strengthening this "missing middle" through common standards and shared infrastructure transforms individual vulnerabilities into collective capacity.

Institutionalize rigorous evidence generation: Current insights reflect 2024-2025 snapshots. However, sustained progress requires longitudinal tracking of outcomes, failures, and adaptations. [NetHope's AI Lighthouse](#) already serves this function by documenting real-world AI deployments, surfacing cross-cutting patterns, and sharing practical guidance with the sector. The next step is to expand this approach to make it sector-wide evidence practice that prevents siloed reinvention and builds a living knowledge base accessible to all

Invest in data as foundational infrastructure: The case studies consistently show that data quality, availability, and governance determine AI effectiveness more than model sophistication. The sector should prioritize building shared, high-quality multilingual datasets; establishing common data governance standards; and investing in the data pipelines and curation processes that make AI tools viable over the long term. Strong foundational data will not only improve individual AI deployments but will also increase the sector's collective influence in shaping how AI is developed and applied in humanitarian contexts.

Confront the accountability gap: As AI moves from pilots to operational decision-making, the sector must grapple with a question: *who is accountable when these systems cause harm?* Whether responsibility sits with the developer, the implementing organization, or the funder remains unresolved. This question sits at the intersection of responsible AI, humanitarian law, and digital rights, and it deserves dedicated research beyond the scope of this analysis. It is raised here deliberately, as a precondition for scaling rather than an afterthought, and as a commitment to address it rigorously in future work.

Conclusion

These 11 case studies move AI in the humanitarian sector beyond proof-of-concept, demonstrating tangible gains in speed, precision, and reach. Yet they equally reveal that sustained impact depends on addressing persistent structural barriers: data deficits, localization gaps, capacity shortages, and funding models misaligned with AI's ongoing needs.

These experiences call for humility in AI adoption, where technical sophistication remains secondary to clarity of purpose, rigorous human oversight, and adherence to humanitarian principles of impartiality, neutrality, and “do no harm”, ensuring AI augments contextual judgment rather than replacing it. Continuous retraining, monitoring, and adaptation remain essential in volatile crisis environments.

The organizations profiled embody this ethos through transparent learning, reinforced by [NetHope's AI Lighthouse](#), enabling nonprofits to advance collectively rather than repeat isolated missteps. Moving forward requires shifting from perpetual pilots to sustained programs through multi-year funding, prioritizing Global South-centered data and cultural adaptation, building shared upskilling consortia, aligning governance standards, and centering affected communities in design.

Thoughtfully deployed and continually refined AI emerges not as a panacea for humanitarian challenges but as a disciplined force multiplier that strengthens effective action for people enduring crisis.

Works cited

NetHope. The Center for the Digital Nonprofit.

<https://nethope.org/programs/the-center-for-the-digital-nonprofit/>

NetHope (2025). NetHope AI Case Studies. DEEP (Data Entry and Exploration Platform)

<https://nethope.org/case-studies/data-entry-and-exploration-platform-deep/>

NetHope (2025). NetHope AI Case Studies. Signpost AI

<https://nethope.org/case-studies/signpost-ai-consortium-of-irc-mercy-corps-internews-and-local-partners/>

NetHope (2025). NetHope AI Case Studies. DRC Displacement Forecast Models

<https://nethope.org/case-studies/displacement-forecast-models-danish-refugee-council/>

NetHope (2025). NetHope AI Case Studies. fAIr (Humanitarian OpenStreetMap Team)

<https://nethope.org/case-studies/fair-humanitarian-open-streetmap-team-hot/>

NetHope (2025). NetHope AI Case Studies. AI for Humanity (Pacific Disaster Center)

<https://nethope.org/case-studies/ai-for-humanity-pacific-disaster-center-pdc/>

NetHope (2025). NetHope AI Case Studies. Africa Infodemic Response Alliance (AIRA)

<https://nethope.org/case-studies/africa-infodemic-response-alliance-case-study/>

NetHope (2025). NetHope AI Case Studies. NRC Knowledge Retrieval Chatbot

<https://nethope.org/case-studies/ai-powered-knowledge-retrieval-chatbot-norwegian-refugee-council/>

NetHope (2025). NetHope AI Case Studies. Mercy Corps Safe Generative AI Chatbots

<https://nethope.org/case-studies/safe-generative-ai-chatbots-mercy-corps/>

NetHope (2025). NetHope AI Case Studies. CRS Market Monitoring Dashboards

<https://nethope.org/case-studies/market-monitoring-dashboards-catholic-relief-services-crs/>

NetHope (2025). NetHope AI Case Studies. ICRC POMELO

<https://nethope.org/case-studies/mapping-vulnerable-populations-with-ai-international-committee-of-the-red-cross/>

NetHope (2025). NetHope AI Case Studies. GiveDirectly AI Cash Triggers

<https://nethope.org/case-studies/ai-supported-triggers-for-cash-transfers-givedirectly-2/>

Annex: glossary of AI terms

Artificial Intelligence (AI)	Computer systems are designed to perform tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and language translation.
Machine Learning (ML)	A subset of AI where systems learn from data to improve performance on specific tasks without being explicitly programmed.
Natural Language Processing (NLP)	AI techniques for understanding, interpreting, and generating human language in text or speech form.
Large Language Model (LLM)	AI models trained on vast amounts of text data can generate human-like text, answer questions, and perform various language tasks. Examples include GPT, Claude, and Llama.
Generative AI	AI systems that can create new content, such as text, images, or code, based on patterns learned from training data.
Retrieval Augmented Generation (RAG)	A technique that combines language models with external knowledge retrieval, allowing AI to access and reference specific documents or databases when generating responses.
Computer Vision	AI techniques for extracting information from images and video, enabling tasks like object detection, image classification, and satellite imagery analysis.
Predictive Analytics	Using historical data, statistical algorithms, and machine learning to identify the likelihood of future outcomes.
Human-in-the-Loop	An approach where humans review and validate AI outputs before they are acted upon, maintaining direct oversight of AI decisions.
Human-on-the-Loop	An approach where AI can act autonomously, but humans monitor outputs and can intervene when needed.
Hallucination	When AI systems generate plausible-sounding but factually incorrect or fabricated information.
Open Source	Software or AI models where the source code is freely available for anyone to use, modify, and distribute.