

Responsible AI for Humanitarian Organizations

A practical guide to using AI without compromising the people you serve

Why this matters more in humanitarian work

AI tools are useful, but humanitarian work raises the stakes. You handle sensitive data about vulnerable people, often in contexts of real power imbalance, where a mistake is not an inconvenience but a risk to someone's safety. Responsible use is not a compliance checkbox here, it is part of doing no harm.

The good news: most of what responsible AI asks of you is judgment you already apply to programme decisions, extended to a new tool.

Start with the people the data is about

Before using AI on any dataset, ask who is in it and what happens to them if it leaks, is wrong, or is used for a purpose they never agreed to.

The test: would you be comfortable explaining this use to the people whose data it is?

If the answer is no, or you are not sure, stop and rethink the approach.

Data protection is not optional

Public AI tools may store, log, or train on what you paste into them. Treat anything you would not post publicly as off-limits:

- **Never paste** names, locations, contact details, or biometric and health data, or anything that could identify a beneficiary
- **Anonymize or aggregate** first: work with "Site B, 340 households", not the household list
- **Check the tool:** enterprise or self-hosted options keep data in your control; free public tools usually do not
- **Honour consent:** data collected for one purpose should not quietly become AI training input

Watch for bias

AI reflects the data it was trained on, which rarely represents the communities you work with. Bias shows up quietly:

- A model trained mostly on English or Global-North contexts can misread local realities
- Targeting or prioritization suggestions can entrench who is already visible and miss who is not
- "The model said so" can launder a biased decision into an objective-sounding one

Treat AI output as a draft to challenge, not an answer to trust, especially when it affects who receives assistance.

Keep a human in the loop

AI can suggest, summarize, and draft. It should not decide anything that affects a person's access to assistance, safety, or rights on its own. Keep a named person accountable for every consequential decision, with the authority to overrule the tool.

Be able to explain it

If you cannot explain how a tool reached a conclusion, you cannot defend it to a donor, a partner, or a community. Prefer approaches you can describe in plain language. Where a system is a black box, keep its role advisory and keep the reasoning your own.

A short guardrail checklist

Before a team adopts an AI tool, agree on:

- What data may and may not go into it
- Which decisions it may inform, and which it may never make on its own
- Who is accountable when it is used
- How you will spot and correct errors
- What your organization's AI-use and data-protection policy says, and if there is none, that is the first thing to write

Remember

Responsible AI is not about avoiding the technology. It is about keeping the people you serve at the centre of every choice about it. Used with judgment, AI frees your time for the work only humans can do. Used carelessly, it puts the most vulnerable at risk. The difference is the care you bring to it.

Need more help? Connect with Impact Node for responsible-AI guidance tailored to humanitarian contexts.

info@impactnode.org