HUMANITARIAN ACADEMY · HUMANITAIRES JOBS

Mini MBA AI and Humanitarian Action: working with Claude and Gemini

Complete manual

A programme of exactly 3 months (12 weeks, 144 hours) to use the AI assistants Claude (Anthropic) and Gemini (Google) in humanitarian work, in a useful, safe and responsible way: how they work and their limits, writing requests, data protection, ethics and humanitarian principles, reports, proposals, data analysis and MEAL, translation and communication with communities, research and source checking, disinformation and security, productivity and automation, use policy and deployment in an organisation. Corrected exercises, three monthly exams with case studies, and a verifiable diploma from the Humanitarian Academy. Independent programme, with no partnership with Anthropic or Google.

https://humanitairejobs.org/en/academy/mini-mba-ia-claude-gemini/

Free document from the Humanitarian Academy. It may be printed and shared freely for training purposes, without modification.

Contents

1. Week 1 : Understanding generative AI assistants: Claude, Gemini and their limits
2. Week 2 : Writing good requests: role, context, task, format and examples
3. Week 3 : Protecting data: what you never give an AI, and how to anonymise
4. Week 4 : AI ethics and humanitarian principles: do no harm, bias, transparency
5. Week 5 : Reports, sitreps and minutes: writing faster without losing accuracy
6. Week 6 : Project proposals and fundraising: AI as a drafting assistant
7. Week 7 : Data analysis and MEAL: spreadsheets, qualitative data and checking calculations
8. Week 8 : Translation and communication with communities: speed, quality and accountability
9. Week 9 : Research and monitoring with AI: finding, citing and checking sources
10. Week 10 : Information, disinformation and security: what AI can and cannot do
11. Week 11 : Everyday productivity: Gemini in Google Workspace, Claude Projects and automation under control
12. Week 12 : Deploying AI in a humanitarian organisation: policy, governance, training and final project

Week 1 : Understanding generative AI assistants: Claude, Gemini and their limits

Objectives

- Explain simply how a large language model produces an answer, and why it can be wrong with confidence.
- Describe the main features of Claude (Anthropic) and Gemini (Google), and the differences between free, paid and professional versions.
- Distinguish the tasks where AI genuinely helps an aid worker from those where it is risky or useless.

What a generative AI assistant is

Claude, developed by the US company **Anthropic**, and Gemini, developed by **Google**, are **conversational assistants** built on large language models (LLMs). You write a request in everyday language; they reply with text and, depending on the version, with tables, code, images or documents.

A language model has been trained on very large amounts of text to **predict the most likely continuation** of a text, piece by piece (the "tokens", roughly three quarters of a word in English, often less in French or Arabic). It was then tuned to follow instructions, decline certain requests and adopt a helpful tone.

The practical consequence: when it answers from memory, the model does not "look up" a database of verified facts. It produces a **plausible** text. Very often, plausible and accurate coincide; sometimes they do not.

> This programme is independent. The Humanitarian Academy has no partnership with Anthropic or Google. The features described are those known in November 2026; these tools change quickly: always check the terms and settings currently in force.

"Hallucinations": errors stated with confidence

A **hallucination** is a false answer presented as true: an invented figure, a reference to an article that does not exist, a misquoted UN resolution, an organisation confused with another. The tone remains as confident as for a correct answer.

The risks increase when you ask for:

- **precise figures** (number of displaced people, amounts, dates) without providing the source;
- **references** (report titles, articles of law, resolution numbers);
- **recent** information, after the end date of the training data, if web search is not enabled;

- **local or rare** topics: a district in a poorly documented country, a language with little presence on the Internet.

The risks decrease when you **provide the source document yourself** (report, table, minutes) and ask the model to work only from it, quoting the passages used.

Example. Asked "How many displaced people in prefecture X in 2025?", an assistant without web search may give a credible and wrong figure. Good practice: provide the latest table published by the official source (for example IOM's Displacement Tracking Matrix) and ask for a summary of that table.

Claude and Gemini: overview of features

In November 2026, both assistants offer comparable features, under different names:

- **Conversation** on the web, desktop and mobile, in many languages including French, English, Spanish, Portuguese and Arabic.
- **File analysis**: PDFs, documents, images, spreadsheets. Both can read a long report and extract the essentials.
- **Web search** with links to sources, and an **in-depth research** mode that consults many pages to produce a report.
- **Reusable workspaces**: Claude's *Projects* (reference documents and instructions shared by several conversations); Gemini's *Gems* (custom assistants with instructions).
- **Document and code creation**: Claude's *artifacts*, Gemini's *Canvas*.
- **Integration with tools**: Gemini is integrated into Google Workspace (Gmail, Docs, Sheets, Meet) depending on the subscription; Claude can be connected to services such as Google Drive or Gmail through connectors, depending on the subscription.

Google also offers **NotebookLM**, a tool that answers only from the documents you upload, with citations: very useful for working on a body of reports.

Both exist in a **free version** (usage limits), as an **individual subscription** and as **plans for organisations** (teams, enterprises, education), with different commitments on data. Week 3 covers this essential point in detail.

What AI does well, what it does badly

For an aid worker, generative AI is above all an **assistant for drafting, summarising and structuring**. It does well at:

- summarising a long document that you provide;
- proposing a first draft (note, email, report outline, terms of reference) from your material;
- rephrasing for a different audience (donor, community, authorities);
- translating everyday text between well-resourced languages, with review;
- explaining a spreadsheet formula, fixing a macro, structuring data;
- acting as a critical reviewer: "what weaknesses would a donor see in this section?".

It does badly, or riskily, at:

- providing facts, figures and references without a source;
- deciding instead of a manager (selecting beneficiaries, recruitment, security level);
- handling situations where an error can seriously harm a person;
- understanding a subtle local context without it being explained.

> The golden rule of this programme: AI prepares, the human checks and decides. You remain responsible for everything you sign, send or publish.

Context window, memory and knowledge date

Three notions avoid many misunderstandings:

- **Context window**: the amount of text the model can take into account in a conversation (your messages, its answers and the files). It is very large for Claude and Gemini (several hundred pages), but not unlimited; in a very long conversation, details from the beginning may be taken into account less well. For a new topic, open a new conversation.
- **Memory**: by default, a conversation does not know the others. Some versions offer memory across conversations, which can be enabled, viewed or erased in the settings. Never assume the assistant "remembers" your organisation.
- **Knowledge date**: the model knows the world up to a certain date. Beyond it, it knows nothing unless it searches the web. For an ongoing crisis, enable search or provide your documents.

Example. A 40-page situation report fits easily into the context window. The complete file of a three-year programme (hundreds of documents) does not: you must select the useful documents, or use a tool designed for a body of documents (Claude Projects, NotebookLM).

Choosing between Claude and Gemini

There is no "best" assistant in absolute terms. Both are of a very high standard and evolve every few months. Useful criteria for a humanitarian organisation:

- **Your working environment**: an organisation already on Google Workspace gets Gemini directly in Gmail, Docs and Sheets; an organisation on Microsoft 365 will rather use the integrations available for that environment, or the web apps.
- **The data protection terms** of the plan available to your organisation (week 3).
- **Quality on your real tasks**: test both on 5 to 10 representative tasks (a note, a summary, a translation into Arabic, a spreadsheet analysis) and compare with a simple grid.
- **Cost** and any discounts for non-profit organisations (to be checked with each provider).
- **Availability in your country**: some services are not available everywhere; check before building a process on them.

Many professionals use both: one as their main tool, the other for a second reading of an important text.

Exercise: spotting the risks in an AI answer

Situation. A programme officer asks an assistant, without web search or documents: "Write a context paragraph on food insecurity in the Kaga region for a proposal to a donor." The answer cites "1.2 million people in phase 3 or above according to the March 2026 IPC", mentions "the 2025 WFP report entitled *Kaga Food Outlook*" and concludes that "female-headed households are 40% more affected".

Task. 1) Which elements are at high risk of hallucination? 2) How should the request be redone properly? 3) What must be checked before using the paragraph?

Answer. 1) The three factual elements: the IPC figure and its date, the title of the WFP report (possibly non-existent), and the percentage on female-headed households. None has a source provided. 2) Download the latest official IPC analysis and the relevant reports, give them to the assistant (or enable web search with links) and ask: "Using only these documents, write a 150-word context paragraph; cite the source and page of each figure; write 'not available' if information is missing." 3) Open each cited source, check each figure, its date and its geographical area; remove any unsourced statement.

Key takeaways

- A language model produces the most plausible text, not necessarily the most accurate.
- Hallucinations mainly affect figures, references, recent facts and local contexts.
- Providing your own documents and requiring citations greatly reduces errors.
- Claude and Gemini have comparable features; the choice depends on the environment, data, tested quality and cost.
- AI prepares, the human checks and decides.

Quiz

1. A large language model produces its answer by:
   A. Consulting a database of verified facts
   B. Predicting the most likely continuation of the text
   C. Copying a random web page
   D. Asking a human expert

2. A "hallucination" is:
   A. A false answer presented with confidence
   B. A bug that prevents an answer
   C. A refusal to answer
   D. An answer that is too long

3. Which request carries the highest risk of error?
   A. Summarising a report provided
   B. Rephrasing an email
   C. Explaining a spreadsheet formula
   D. Giving the exact number of displaced people in a district without a source

4. Google's tool that answers only from uploaded documents, with citations, is called:
   A. Gems
   B. Canvas
   C. NotebookLM
   D. Artifacts

5. Who is responsible for a report written with the help of AI?
   A. The AI provider
   B. Nobody
   C. The donor
   D. The person and the organisation who sign it

Quiz answers : 1. B · 2. A · 3. D · 4. C · 5. D

Week 2 : Writing good requests: role, context, task, format and examples

Objectives

- Write a complete request in five elements: role, context, task, constraints and expected format.
- Improve an answer through successive iterations rather than starting over.
- Build a library of reusable standard requests for your job and your team.

Why wording matters so much

An AI assistant knows neither your organisation, nor your programme, nor your donor, nor your audience. A vague request ("Write a report on our project") produces generic text full of empty phrases. A precise request produces a usable first draft.

People often talk about *prompt engineering*. For everyday professional use, it is not about magic formulas but about **good briefing practice**, the same as briefing a competent new colleague who is discovering the file.

> Simple test: if you gave your request as it stands to a consultant who knows nothing about your project, could they do a good job? If not, neither can the AI.

The five elements of a good request

- **Role**: the expected perspective. "You are an experienced MEAL specialist in nutrition programmes." The role steers vocabulary and the level of rigour.
- **Context**: who you are, the programme, the country, the audience, the constraints. "National NGO, nutrition programme funded by a European donor, three districts, interim report at 6 months."
- **Task**: a clear action verb. "Write the 'Progress by result' section from the data below."
- **Constraints**: length, tone, elements to include or exclude, rules. "400 words maximum, factual tone, do not invent any figure, flag missing data."
- **Format**: table, list, paragraphs, headings, language. "A table indicator / target / achieved / comment, then two paragraphs."

Complete example. "You are the grants manager of a medical NGO. Context: we are responding to a EUR 500,000 call for proposals for primary health care in two provinces; the donor requires an approach centred on accountability to affected populations. Task: propose a detailed outline of the 'Accountability' section (headings and subheadings, 3 points per subheading). Constraints: one page maximum, no promise we cannot keep, Core Humanitarian Standard terms. Format: numbered list."

Giving examples and documents

The most reliable way to get the desired style is to **show an example**. A paragraph from a past report the donor liked, an email template, a well-completed indicator sheet: the assistant will imitate the structure and tone.

For long reference documents, use clear markers: "Here is the logframe between the tags <logframe> and </logframe>, and the quarter's data between <data> and </data>." Both assistants handle this kind of structuring well.

Explicitly ask the assistant to **rely only on the documents provided** when accuracy matters: "If information is not in the documents, write 'information not available' instead of guessing it." This sentence clearly reduces invention.

Caution. Before pasting a document, check that it contains no personal or sensitive data (week 3). A past donor report sometimes contains beneficiary names, partner lists or security incidents.

Iterating: the conversation is a workshop

The first answer is rarely the right one. Rather than rewriting everything, **correct step by step**:

- "Shorten to 250 words, keeping the three key figures."
- "The tone is too promotional; make it factual, like a monitoring report."
- "The second paragraph repeats the first: merge them."
- "Add a line on delays due to the rainy season, with their effect on indicator 2.1."

Two very effective techniques:

- **Ask for questions first**: "Before drafting, ask me the 5 questions you need to do a good job." The assistant thus reveals what your brief is missing.
- **Ask for a critique**: "Review this text as a demanding donor evaluator: list the 5 main weaknesses, from most to least serious." Then correct it yourself or with the AI.

When a conversation drifts or becomes very long, start a fresh conversation with an improved brief: it is often faster.

Standing instructions: Projects and Gems

If you often repeat the same context, do not paste it every time:

- in **Claude**, a *Project* brings together standing instructions and reference documents (for example the communication charter, the logframe, the glossary) shared by all conversations in the project; depending on the plan, a project can be shared with colleagues;
- in **Gemini**, a *Gem* is a custom assistant with its instructions and, depending on the version, reference files.

Example of standing instructions for a communications officer: "You help the communications team of NGO X. Tone: sober, respectful, no pity-based imagery. Never identify a child or a victim. Always use 'internally displaced people' and not 'refugees' for IDPs. Default length: 150 words. Flag any statement that requires a source."

These instructions embody the organisation's rules. They must be reviewed by the responsible person, like any procedure document.

A request library for the team

The teams that get the most value from AI share their **requests that work**. A simple library (a shared document or spreadsheet) contains for each standard request:

- the name and use ("Summary of the minutes of a cluster meeting");
- the text of the request, with fields to fill in between square brackets;
- the documents to provide and those never to provide;
- the mandatory checkpoints ("check dates and organisation names");
- the author and date of last update.

Start small: 10 requests for the most frequent tasks of the job (emails, minutes, summaries, rephrasing, translations). Measure the time saved and the quality, then expand.

> A good request is worth a good document template: it can be shared, tested, improved and filed with the procedures.

Exercise: turning a vague request into a professional brief

Situation. A logistics coordinator writes to the assistant: "Write me an email for the supplier, he's late." The result is polite but useless: no order number, no date, no consequence.

Task. Rewrite the request with the five elements (role, context, task, constraints, format), inventing plausible data.

Possible answer. "You are an experienced logistics manager in an NGO. Context: purchase order PO-2026-114 of 2 October covers 400 hygiene kits; delivery was due on 30 October at the latest under the contract; today is 12 November; distribution to displaced families is scheduled for 20 November; the contract provides for penalties of 0.5% per day of delay, capped at 10%. Task: write an email to the supplier. Constraints: firm and courteous tone; request a written delivery date within 48 hours; recall the penalty clause without threatening; mention that we may cancel the order if delivery has not taken place before 17 November; 150 words maximum. Format: email subject then body, in English." Human check: order number, dates, penalty rate and cancellation clause consistent with the signed contract.

Key takeaways

- Brief the AI like a competent new colleague who knows nothing about the file.
- A good request contains a role, a context, a task, constraints and a format.
- An example and documents provided give the best result; require 'information not available' rather than an invention.
- Iterate: ask for questions first, then a critique, then targeted corrections.
- Projects (Claude) and Gems (Gemini) keep standing instructions; share a library of tested requests.

Quiz

1. Which is NOT one of the five elements of a good request?
   A. The role
   B. The context
   C. The account password
   D. The expected format

2. To reduce inventions when accuracy matters, you ask:
   A. To answer faster
   B. To rely only on the documents provided and write 'not available' otherwise
   C. To be more creative
   D. To answer in English

3. "Before drafting, ask me the questions you need" is used to:
   A. Reveal what the brief is missing
   B. Save tokens
   C. Avoid human checking
   D. Enable web search

4. In Claude, the space that brings together standing instructions and reference documents is called:
   A. A Gem
   B. A Canvas
   C. A notebook
   D. A Project

5. A team request library should contain in particular:
   A. The account passwords
   B. The mandatory checkpoints
   C. Beneficiary data
   D. The team's salaries

Quiz answers : 1. C · 2. B · 3. A · 4. D · 5. B

Week 3 : Protecting data: what you never give an AI, and how to anonymise

Objectives

- Apply the principles of data protection in humanitarian action to the use of AI assistants.
- Classify information (public, internal, confidential, sensitive) and decide whether it can be given to an assistant, and with which plan.
- Anonymise or pseudonymise a document before using it with AI, and adjust privacy settings.

Why this is the most important week

Humanitarian organisations hold some of the most sensitive data there is: identities of refugees or displaced people, survivors of violence, separated children, health data, team locations, security incidents, information about armed groups. A leak can expose people to reprisals, discrimination or violence.

When you paste text into an AI assistant, that text **leaves your computer** and is processed on the provider's servers, often in another country. Depending on the plan and settings, it may be **retained**, sometimes **read by people** in charge of quality, or used to **improve the models**.

> Principle: never give an AI assistant data that you would not be allowed to email to an external contractor without a data processing agreement.

The reference principles

The **Handbook on Data Protection in Humanitarian Action** by the ICRC and the Brussels Privacy Hub (2nd edition, 2020, which also covers new technologies), the data protection policies of UN agencies and, for the organisations concerned, the European Union's General Data Protection Regulation (GDPR) set out common principles:

- **lawfulness and purpose**: data is collected for a specific purpose and is not reused for an incompatible purpose;
- **minimisation**: use only the data strictly necessary;
- **security**: protect data against unauthorised access;
- **transfers**: regulate any transfer to a third party or another country;
- **rights of individuals**: information, access, rectification.

Giving the data of a beneficiary list to an AI assistant for the general public is a **transfer to a third party**, for a **purpose** that was not announced to the people concerned. In the vast majority of cases, this is contrary to these principles.

The sector also speaks of **demographically identifiable data**: even without names, aggregated data (for example the precise location of an ethnic or religious minority group) can expose an entire community.

What you never give a consumer assistant

- Names, phone numbers, card or registration numbers, photos, addresses, GPS coordinates of beneficiaries or staff members.
- Any information about survivors of gender-based violence, children, protection cases, mental health or HIV, even roughly "anonymised".
- Detailed security incident reports, routes, positions and travel times, evacuation plans, information on access negotiations.
- Allegations of sexual exploitation and abuse, fraud or corruption under investigation.
- Individual health data.
- Passwords, access keys, bank details.
- Documents marked confidential by the organisation, a donor or a partner.

For these categories, even a professional plan with contractual commitments is not always enough: a documented decision by the organisation (data protection impact assessment) is required before any use.

Consumer plans and professional plans

In November 2026, the broad lines are as follows (always check the terms in force):

- **Consumer plans (free or individual subscription)**: both providers offer settings to control history retention and the use of conversations to improve models. At Google, Gemini Apps activity may be retained and conversations may be reviewed by people; a setting allows it to be turned off. At Anthropic, the user chooses in the privacy settings whether their conversations may be used for training, and Claude offers an incognito chat mode.
- **Plans for organisations** (Claude Team and Enterprise; Gemini in Google Workspace for business and education): providers commit contractually not to use customer data to train their models by default, with central administration, access controls and data processing terms.

In practice:

- a free personal account should only be used for **public or non-sensitive** information;
- **internal** (non-sensitive) information belongs in a professional plan approved by the organisation;
- **personal and sensitive** data remains excluded, unless there is a formal, regulated decision.

> Turning off history or training use reduces risks, but does not turn a personal account into an authorised tool for beneficiary data.

Classify, then anonymise

Before each use, classify the information according to the organisation's policy, for example:

- **Public**: published reports, calls for proposals, legal texts. Usable with any assistant.
- **Internal**: draft proposals, procedures, minutes without personal data. Approved professional plan.
- **Confidential**: detailed budgets, partner assessments, negotiations. Professional plan and manager's approval, after removing unnecessary elements.
- **Sensitive**: personal data, protection, security, investigations. Excluded.

Anonymising means making it impossible to identify a person, including by cross-referencing. **Pseudonymising** means replacing identifiers with codes ("Beneficiary A"); the person remains identifiable by whoever holds the lookup table: this is still personal data.

Common traps: a village of 300 inhabitants plus "34-year-old woman, widowed, four children" is enough to identify someone; a focus group quote may contain a name or nickname; file metadata (author, comments, revisions) contain names. Remove unnecessary details, group ages and places, and delete comments and document history.

Good everyday habits

- Use the **account provided by the organisation** for work, never a personal account for internal documents.
- Check the **privacy settings** at first login, then every quarter: history, training use, memory, enabled connectors.
- Limit **connectors** (access to email, Drive) to what is strictly necessary: an assistant that can read your whole mailbox can also read sensitive data in it.
- Do not share **public links** to conversations containing internal information.
- In the field, watch out for **visible screens** and shared devices; log out.
- In case of error (sensitive data pasted by mistake): delete the conversation, then **report the incident** to the data protection focal point without delay. An error reported quickly limits the damage; a hidden error makes it worse.

Exercise: can AI be used for this document?

Situation. A protection team wants to save time on five tasks: A) summarise the public report of a protection cluster; B) translate into English the minutes of an internal meeting without names; C) analyse a spreadsheet of 2,300 households with names, phone numbers and vulnerabilities to find duplicates; D) rephrase the account of a survivor of violence for a case report; E) propose a staff training plan on PSEA. The organisation has a professional plan approved by the data protection officer, with a commitment not to use data for training.

Task. For each task: allowed with a personal account, allowed with the professional plan only, or excluded? Justify.

Answer. A) Allowed with any account: public document. B) Professional plan only: internal non-sensitive document, after checking that it really contains no names or security elements. C) Excluded in this form: personal and vulnerability data of 2,300 households; deduplication is done in the organisation's beneficiary management tool, or the AI can be asked to write a deduplication formula or script tested on fictitious data, then run locally. D) Excluded: gender-based violence data, extremely sensitive, even with a professional plan. E) Allowed with any account if the request contains no sensitive internal information; preferably with the professional plan in order to attach internal procedures.

Key takeaways

- Text pasted into an assistant leaves the organisation: it is processed, sometimes retained, by the provider.
- Minimisation, purpose and regulated transfers apply to the use of AI.
- Never put personal data of beneficiaries, protection cases, health or security data into a consumer assistant.
- Professional plans offer contractual commitments, but do not replace a decision by the organisation for sensitive data.
- Pseudonymising is not anonymising; an error is reported immediately.

Quiz

1. Pasting a beneficiary list into a consumer assistant is:
   A. A transfer to a third party for an unannounced purpose, to be avoided
   B. Allowed if history is turned off
   C. Allowed if the list is in English
   D. Without consequence

2. Replacing names with "Beneficiary A, B, C" while keeping a lookup table is:
   A. Anonymising
   B. Encrypting
   C. Deleting
   D. Pseudonymising: this is still personal data

3. A detailed security incident report with routes and times:
   A. Can be summarised by a free account
   B. Must not be given to an assistant without a formal decision by the organisation
   C. Can be given if it is translated
   D. Must be published

4. You pasted sensitive data by mistake. What should you do?
   A. Nothing
   B. Continue the conversation
   C. Delete the conversation and report the incident to the focal point without delay
   D. Change browser

5. The principle of minimisation means:
   A. Writing short requests
   B. Using as few tokens as possible
   C. Paying for the cheapest plan
   D. Using only the data strictly necessary

Quiz answers : 1. A · 2. D · 3. B · 4. C · 5. D

Week 4 : AI ethics and humanitarian principles: do no harm, bias, transparency

Objectives

- Link the use of AI to the humanitarian principles (humanity, neutrality, impartiality, independence) and to the do no harm principle.
- Spot the biases in an AI answer and their possible effects on affected people.
- Apply clear rules on transparency, human responsibility and respect for copyright.

The humanitarian principles and AI

The four humanitarian principles also apply to digital tools:

- **Humanity**: the tool must serve affected people, not only internal efficiency. Saving time on a report is useful if that time goes back to programmes.
- **Impartiality**: AI must not introduce discrimination in access to aid (for example by ranking households according to biased criteria).
- **Neutrality**: generated text may, unnoticed, adopt the vocabulary of one party to a conflict. Review the names of places, groups and events.
- **Independence**: depending entirely on one provider for a critical function is a risk; keep the ability to work without the tool.

The **do no harm** principle requires asking before each new use: who could be harmed if the tool is wrong, if the data leaks, or if the result is misunderstood?

Bias: where it comes from, how to spot it

Models learn from texts mostly produced in a few languages and a few regions of the world. This results in **biases**:

- **Language bias**: lower quality in languages with little presence on the Internet (many African, Asian or Indigenous languages), and sometimes in certain varieties of Arabic or French.
- **Perspective bias**: a "Global North" reading of crises, standardised solutions, little knowledge of local actors and national NGOs.
- **Stereotypes**: gendered or cultural representations (the "doctor" is a man, the "beneficiary" is passive) in texts and especially in generated images.
- **Sycophancy bias**: a tendency to agree with the person asking. If you ask "confirm that our project is a success", the assistant is likely to confirm.

To spot them: ask for the opposite position ("what arguments go against this conclusion?"), have texts for communities reviewed by people from those communities, and compare with local sources.

Decisions about people: the line not to cross

Decisions that directly affect a person's rights or access to aid must not be entrusted to a generative AI assistant:

- selection or exclusion of beneficiaries;
- assessing the credibility of an account (asylum claim, complaint, protection case);
- shortlisting or rejecting job applicants;
- disciplinary or investigation decisions.

These tools can help **prepare** (draft an interview grid, structure published criteria, summarise a policy), but the decision belongs to authorised people, following a documented procedure, with the possibility of appeal. In the European Union, the Artificial Intelligence Act (AI Act, adopted in 2024 and applied gradually) classifies certain uses such as recruitment or access to essential services as "high-risk", with stronger obligations.

> Test: if the person concerned learned that the decision had been made by an AI, would that be acceptable, and could they challenge it? If the answer is no, do not use AI to decide.

Transparency and accountability

Trust is the capital of a humanitarian organisation. A few simple rules protect it:

- **Human responsibility**: every document has a human author responsible for its content, whether or not they were helped by AI.
- **Proportionate transparency**: disclose the use of AI when it is significant or when the recipient can reasonably expect it (evaluation, research, published content, exchanges with communities). For a spelling check, it is not necessary.
- **Respect for donor and partner rules**: some donors and calls for proposals regulate or require disclosure of AI use in applications; read the terms (week 6).
- **No impersonation**: never generate an image, voice or quote falsely attributed to a real person, and never present a generated image as a field photo.

The sector's ethical communication guidelines (dignity, informed consent, no pity-based imagery) apply fully to content produced with AI.

Copyright, plagiarism and integrity

Three points to watch:

- **Other people's content**: asking AI to "rewrite" another organisation's report to present it as yours is still plagiarism. Cite your sources.
- **Generated content**: the legal status of generated texts and images varies between countries and is evolving. The terms of use of Claude and Gemini specify the user's rights over the outputs; for commercial use or a major publication, consult them.

- **Academic and professional integrity**: for the exams of this Mini MBA, as for a recruitment test, using AI to answer in your place is fraud. The aim is for you to master the skills, not the tool alone.

For photos, the organisation's rule remains: no photo of others without rights, no photo of an identifiable child without consent, and a generated image is clearly labelled as such.

Environmental cost and equity of access

AI consumes electricity and water in data centres. Estimates per request vary widely between studies and models; the sector's overall impact is growing fast. For an organisation committed on climate, reasonable habits:

- use AI when it brings real value, not by reflex for trivial tasks;
- prefer one well-prepared document to dozens of successive attempts;
- include digital use in the environmental policy.

Equity of access: field teams often have weak connections, old devices and less access to paid subscriptions. An organisation that reserves AI for headquarters and international staff widens inequalities. A fair policy provides access, training and languages for national staff and local partners.

Exercise: analysing three uses from an ethical standpoint

Situation. An NGO is considering three uses: A) generating "realistic" images of malnourished children for a fundraising campaign, for lack of authorised photos; B) asking AI to rank 600 applications for a driver position and automatically reject the 500 lowest-scoring ones; C) using AI to propose a first version of cholera awareness messages in easy French, then tested with community groups and translated by native speakers.

Task. For each use, say whether it is acceptable and why, and propose an alternative if necessary.

Answer. A) Not acceptable: a generated image presented as real misleads the public, undermines dignity and reinforces pity-based imagery; alternative: photos with informed consent, clearly labelled illustrations, or messages focused on solutions. B) Not acceptable: a decision about people without human control, risk of bias and discrimination, no appeal; alternative: published criteria applied by a panel, with AI possibly helping to draft the grid or a practical test. C) Acceptable: AI prepares a first draft, the community and native speakers validate it; it is a proportionate, transparent and verified use.

Key takeaways

- The humanitarian principles and the do no harm principle apply to the use of AI.
- Biases affect language, perspective, stereotypes and sycophancy; look for contradiction.
- No decision on a person's rights or access to aid is entrusted to a generative AI.
- A human author remains responsible; disclose significant AI use; never present a generated image as real.
- Consider the environmental cost and equity of access for national staff and partners.

Quiz

1. Asking AI "confirm that our project is a success" mainly exposes you to:
   A. Language bias
   B. Sycophancy bias
   C. Date bias
   D. Format bias

2. Which of these tasks must NOT be entrusted to a generative AI?
   A. Drafting an interview grid
   B. Summarising a published HR policy
   C. Deciding which households are excluded from aid
   D. Correcting the spelling of a report

3. A generated image of a field situation can be used:
   A. As a real photo if it is realistic
   B. Never, under any circumstances
   C. Only for donors
   D. If it is clearly labelled as generated and respects dignity

4. The principle of independence encourages you to:
   A. Keep the ability to work without the tool for critical functions
   B. Use a single provider
   C. Never use AI
   D. Let AI decide

5. Answering the exams of this Mini MBA with AI in your place is:
   A. Recommended
   B. Fraud
   C. Mandatory
   D. Unimportant

Quiz answers : 1. B · 2. C · 3. D · 4. A · 5. B

Week 5 : Reports, sitreps and minutes: writing faster without losing accuracy

Objectives

- Use Claude or Gemini to produce a first draft of a report from data and notes provided.
- Apply a systematic method to check figures, names, dates and statements.
- Adapt the same content to several audiences: donor, management, cluster, communities.

The right workflow

For a report, the classic mistake is to ask AI to "write the report". The right workflow has four steps:

- **1. Gather** the material: monitoring data, field notes, meeting minutes, previous report, donor template. You bring the facts.
- **2. Structure**: ask for an outline matching the required template, then validate it.
- **3. Draft section by section**, from the material provided only, requiring that every figure be reproduced as is and that gaps be flagged.
- **4. Check and finalise**: fact-check, consistency between sections, tone, length, then review by a second person for binding reports.

With this method, teams often report significant time savings on drafting, but **not on checking**, which remains fully necessary. Plan for it in the schedule.

The situation report (sitrep)

A sitrep must be **short, factual, dated and sourced**. AI helps turn scattered notes into a structured document:

- **Key points** (3 to 5 lines);
- **Situation**: changes since the last report, with figures and sources;
- **Needs** and gaps;
- **Response**: what has been done, by whom, where;
- **Constraints**: access, logistics, funding;
- **Next steps**.

Standard request. "From the notes below, write a sitrep in the following format [format]. Reproduce the figures exactly. Clearly distinguish what is confirmed, what is reported by a single source and what is estimated. Give the date and source of each figure. Do not mention any person's name or any team location. 400 words maximum."

> Watch out for security information: a widely shared sitrep must contain no routes, no times, no names and no incident details that could identify people. That information stays in security channels.

Meeting minutes

Coordination, cluster or team meetings produce often disorganised notes. AI is good at turning notes or a transcript into structured minutes: participants (by organisation), points discussed, decisions, actions with owner and deadline.

Points of vigilance:

- **Recording and transcription**: only record a meeting with the participants' agreement and according to the organisation's rules. Some tools (for example the note-taking features of Google Meet with Gemini, depending on the subscription) do it automatically: inform participants.
- **Decisions and actions**: check that they correspond to what was actually decided. AI can turn an idea that was mentioned into a decision.
- **Sensitive remarks**: before circulation, remove remarks about individuals, incidents or negotiations.

Useful request. "List separately: 1) the decisions taken; 2) the points discussed without a decision; 3) the actions with owner and date. If the owner or date is not clear in the notes, write 'to be confirmed'."

Checking: the four-check method

Before sending a text written with AI, apply four checks:

- **1. Figures**: every number is found in the source, with the same unit, the same date and the same area. Recalculate totals and percentages.
- **2. Names and terms**: organisations, places (official spelling), people's titles, acronyms, donor terminology.
- **3. Statements**: every sentence asserting a fact or a result must be supported by the data. Remove unjustified adjectives ("remarkable", "unprecedented").
- **4. Omissions**: has the AI left out a problem, a delay, a gap against targets? Reports must present difficulties honestly.

Tip: ask the assistant itself, in a **new conversation**, to compare the final text with the source data and list any difference. It is a useful aid, but it does not replace your own check.

Control calculation. A report states "3,450 people reached, i.e. 115% of the target of 3,000". 3,450 / 3,000 = 1.15: correct. But the source table shows 3,450 *contacts*, including people counted twice. The right indicator is the number of unique people: the text must be corrected, not just the calculation.

Adapting to each audience

The same content often needs to be adapted:

- **Donor**: compliant with the template, focused on results, indicators, variances and budget.
- **Management**: one page, decisions to be taken first.
- **Cluster**: aggregated data, geographical coverage, gaps.
- **Communities**: plain language, local language, what has been done, what will change, how to give feedback or complain.

AI excels at these rephrasings: "Turn this report into a one-page note for management, with three decisions to be taken at the top." or "Summarise these results in 120 words in easy French, end-of-primary-school level, for a poster at distribution sites."

For communities, the final version must be validated by people who speak the language and know the context (week 8).

What not to do

- Let AI "fill in" missing data or estimate figures you do not have.
- Copy "case studies" or testimonies invented by AI. A testimony must be real, collected with consent and anonymised if necessary.
- Send a report in which you have not read every sentence.
- Standardise all your reports with the same phrases: donors quickly spot generic texts, and it damages credibility.
- Paste named beneficiary tables to "calculate the indicators".

> A fast but wrong report costs more than a slow one: loss of credibility, requests for clarification, and sometimes refunds.

Exercise: checking a report paragraph

Source data (monitoring table). Target: 2,000 households receiving cash transfers. Achieved at 30 September: 1,640 households. Amount per household: USD 60 per month for 3 months. Satisfaction rate from post-distribution monitoring: 84% on a sample of 250 households. One-month delay due to the temporary closure of a payment provider.

Paragraph proposed by the AI. "The project reached 1,640 households, i.e. 92% of the target, with a total transfer of USD 295,200. Satisfaction is remarkable: 84% of all beneficiaries say they are satisfied. Activities were carried out according to schedule."

Task. Find the errors and rewrite.

Answer. 1) 1,640 / 2,000 = 82%, not 92%. 2) Total amount: 1,640 × 60 × 3 = USD 295,200 only if all households received all 3 transfers; to be checked in the payment data. 3) 84% refers to a sample of 250 households, not to all beneficiaries; "remarkable" is an unjustified judgement. 4) The one-month delay is omitted and "according to schedule" is false. Corrected version: "At 30 September, 1,640 households had received cash transfers, i.e. 82% of the target of 2,000. [Amount to be confirmed from payment data.] During post-distribution monitoring, 84% of the 250 households surveyed said they were satisfied. Distributions were delayed by one month due to the temporary closure of a payment provider; [corrective measure]."

Key takeaways

- You bring the facts; AI structures and drafts; you check.
- Sitrep: short, dated, sourced, distinguishing confirmed, reported and estimated, with no security information.
- Minutes: separate decisions, discussions and actions; inform before any recording.
- Four checks: figures, names and terms, statements, omissions.
- Never invented data, fictitious testimonies or a report sent without a complete read-through.

Quiz

1. In the recommended workflow, who brings the facts?
   A. The AI
   B. The donor
   C. You, with your data and notes
   D. The Internet

2. 1,640 households reached for a target of 2,000 is:
   A. 82%
   B. 92%
   C. 8.2%
   D. 122%

3. A widely circulated sitrep must not contain:
   A. Figures
   B. A date
   C. Sources
   D. Routes or travel times

4. The fourth check of the method concerns:
   A. The font
   B. Omissions (problems, delays, variances)
   C. Length
   D. The logo

5. A beneficiary testimony invented by AI can be used:
   A. If it is anonymous
   B. If it is short
   C. Never
   D. Only in an annex

Quiz answers : 1. C · 2. A · 3. D · 4. B · 5. C

Week 6 : Project proposals and fundraising: AI as a drafting assistant

Objectives

- Analyse a call for proposals with AI: criteria, requirements, pitfalls and eligibility grid.
- Build with AI a first draft of a theory of change, a logframe and a narrative, from the organisation's data.
- Respect donor rules, avoid generic text and keep a proposal honest and realistic.

Analysing a call for proposals

A call for proposals is often 30 to 80 pages long, with annexes. AI helps break it down quickly, provided you give it the complete document:

- "List all the eligibility criteria (organisation, partners, areas, amounts, duration) in a table, with the page of each criterion."
- "List all the documents to be provided and the required formats (number of pages, font, language, annexes)."
- "Summarise the evaluation grid and the weight of each criterion."
- "Which requirements are likely to be overlooked by a hurried applicant?"

Then **check each element in the original text**, especially the deadline, the time zone, the minimum and maximum amounts and the co-financing rules. A single eligibility error gets an application rejected, however good it is.

Go / no-go decision. Also ask: "Given our profile [presence, experience, annual budget, partners], which criteria do we meet, which are uncertain, which do we not meet?" The decision to apply remains management's.

Theory of change and logframe

AI is useful for **structuring the logic** of a project from your needs analysis:

- proposing a theory of change (if... then... because...) with explicit assumptions;
- proposing a hierarchy of objectives (overall objective, specific objective, results, activities);
- suggesting SMART indicators and their means of verification;
- checking consistency: does each activity contribute to a result? Does each result have an indicator?

Points of vigilance:

- **Numerical targets** and **baselines** come from your data and your real capacity, never from AI.

- Prefer the **standard indicators** required or recommended (Sphere standards, cluster indicators, donor indicators): give the list to the assistant.
- **Assumptions and risks** must reflect your real context: security, rainy season, access, prices, authorities.

Standard request. "Here is our needs analysis and the list of cluster indicators. Propose a logframe in a table. Do not invent any target: leave the target column empty. For each indicator, state whether it comes from the cluster list or is a proposal."

The narrative: from first draft to the organisation's text

For the narrative sections (context, justification, approach, sustainability, gender, accountability, localisation), an effective method:

- **Your elements first**: list in bullet points the facts, experience, partners, lessons learned. AI does not know your added value.
- **First draft by the AI**, section by section, within the page limit.
- **Rewriting by the team**: adding concrete examples, verified figures, place names, real lessons. This is what distinguishes a winning proposal.
- **Critique by the AI**: "Assess this section against the call's grid [grid]; score each criterion out of 5 and propose improvements."

> Evaluators read hundreds of applications. Empty phrases ("holistic and innovative approach", "multi-sector synergies") churned out by AI lower the score. Concrete, local and verifiable content raises it.

The budget and its justification

The budget remains specialist work, using the organisation's financial tools. AI can help to:

- write the **budget justification** (narrative) from the detailed budget;
- check **consistency** between narrative activities and budget lines ("List the narrative activities without a budget line, and the lines without a corresponding activity.");
- explain or write **spreadsheet formulas** for the calculations.

Never trust it for final calculations without checking: recalculate in the spreadsheet. And do not give it named salary information.

Control example. The narrative announces 12 training sessions of 25 participants, but the budget provides per diems for 200 participants. 12 × 25 = 300 participant-sessions; if each person attends 3 sessions, that makes 100 people, not 200. The inconsistency must be corrected before submission: this is exactly what AI can help spot, and what the team must settle.

Donor rules on AI

Donors and foundations have begun to regulate the use of AI in applications and reports. Positions vary: some accept its use as a drafting aid, others require it to be disclosed, others recall that the applicant remains fully responsible for accuracy and originality. Several research funding bodies also prohibit reviewers from using AI on the confidential applications they receive.

Good practice:

- read the terms of the call and the donor's guidelines on AI, if any;
- disclose use when required, factually;
- do not give a consumer AI documents that the donor has marked confidential;
- keep track of your sources and versions;
- check that the proposal promises nothing the organisation cannot deliver: it is a contractual commitment.

Beyond proposals: the donor relationship

AI also helps with the rest of fundraising:

- **Monitoring** funding opportunities with web search and sources (then checking on the donor's official website);
- **Donor profiles** from public documents: priorities, countries, amounts, procedures;
- Short **concept notes** from a long proposal;
- **Meeting preparation**: "What difficult questions could this donor ask about our exit strategy? Propose elements of response from our document.";
- **Messages to private donors**: thank-you letters, newsletters, in line with communication ethics.

The relationship of trust, however, rests on people: visits, transparency about difficulties, accurate and timely reports.

Exercise: consistency review of a proposal

Situation. An 18-month WASH proposal announces in the narrative: "10 boreholes rehabilitated, 15,000 people with access to safe water according to Sphere standards (15 litres per person per day), 60 water management committees trained". The budget provides for the rehabilitation of 8 boreholes and the training of 30 committees. The logframe, drafted with AI, sets as a target "100% of households satisfied" and contains an indicator "number of synergies created".

Task. Identify the problems and propose corrections.

Answer. 1) Narrative / budget inconsistency: 10 boreholes against 8, 60 committees against 30; align with what is actually funded and feasible. 2) 15,000 people at 15 litres per day represent 225,000 litres per day: check that the yield of the rehabilitated boreholes allows this (technical data), and also apply the Sphere standard on the number of people per water point. 3) "100% of households satisfied" is neither realistic nor credible: set a target based on the baseline and experience (for example 80%), with the measurement method. 4) "Number of synergies created" is neither measurable nor useful: replace it with a precise indicator (for example the number of functional committees according to a functionality grid). 5) Have it reviewed by the WASH specialist and the finance manager.

Key takeaways

- AI breaks down a call for proposals; check each criterion in the original text.
- Targets, baselines and risks come from your data and context, not from AI.
- First draft by the AI, concrete rewriting by the team, critique against the call's grid.
- Check consistency between narrative, logframe and budget; recalculate in the spreadsheet.
- Respect donor rules on AI and confidentiality; only promise what is feasible.

Quiz

1. The numerical targets of a logframe must come from:
   A. The AI
   B. The organisation's data and real capacity
   C. Another project at random
   D. The donor only

2. 12 sessions of 25 participants, each person attending 3 sessions: how many people?
   A. 300
   B. 75
   C. 100
   D. 36

3. Empty phrases repeated by AI in a proposal:
   A. Lower the evaluators' score
   B. Always raise the score
   C. Are required
   D. Are neutral

4. Faced with a call that requires disclosure of AI use, you:
   A. Ignore it
   B. Hide it
   C. Refuse to apply
   D. Disclose it factually

5. An indicator "number of synergies created" is:
   A. A good SMART indicator
   B. Not measurable: to be replaced by a precise indicator
   C. Mandatory
   D. A Sphere indicator

Quiz answers : 1. B · 2. C · 3. A · 4. D · 5. B

Week 7 : Data analysis and MEAL: spreadsheets, qualitative data and checking calculations

Objectives

- Use AI to write and explain formulas, clean and structure data, and prepare analyses.
- Analyse qualitative data (focus groups, community feedback) with AI in a rigorous and traceable way.
- Systematically check the calculations and interpretations produced by AI.

What AI brings to MEAL

MEAL teams spend a lot of time on repetitive tasks: cleaning files, writing formulas, producing tables, coding open answers, drafting summaries. AI can speed up these tasks:

- **Formulas and scripts**: writing a spreadsheet formula (XLOOKUP, COUNTIFS, pivot tables), a Python or R script, a query in a data collection tool, and explaining them line by line.
- **File analysis**: Claude and Gemini can read a CSV or Excel file, run calculations in a code environment and produce tables or charts (depending on the version and plan).
- **Qualitative data**: proposing a coding frame, coding answers, grouping by theme, extracting representative quotes.
- **Writing**: turning results into findings, then into recommendations.

But MEAL data is often **personal** (household lists, named surveys). The rule from week 3 applies: only anonymised and aggregated data, or work on a fictitious sample with the same structure.

Formulas and scripts: the fictitious data method

A safe and effective method to automate an analysis without exposing data:

- **1.** Describe the structure of your file (column names, types, examples of invented values) or create 10 fictitious rows with the same structure.
- **2.** Ask for the formula or script: "Write a formula that counts the number of female-headed households that received at least 2 distributions, by district."
- **3.** Test on the fictitious data, where you know the right result.
- **4.** Apply to your real data, **in your spreadsheet or on your computer**.
- **5.** Check with an independent calculation (manual filter, pivot table).

With Gemini in Google Sheets or the AI features of other spreadsheets, the AI acts directly in the file: check that the plan is professional and authorised for this data.

> AI writes the formula, you run it on your data, you check the result. Personal data does not leave your tools.

Checking calculations and statistics

Language models can make mistakes in mental arithmetic, especially on long series. When the assistant actually runs code, the calculations are reliable, but the **interpretation** can be wrong. Essential checks:

- **Denominators**: a percentage of what? Respondents, surveyed households, the target population? Are missing answers excluded?
- **Sample**: a result from 250 surveyed households does not apply to the whole population without an appropriate sampling method. Ask for the margin of error.
- **Mean or median**: for income or expenditure, a few extreme values distort the mean.
- **Correlation is not causation**: "households that received aid have a better food consumption score" does not prove that the aid is the cause without a comparison group.
- **Disaggregation**: sex, age, disability, district; check that the subtotals add up to the total.

Calculation. A simple random sample of 384 households: margin of error of about ±5% at a 95% confidence level for a proportion close to 50%. If the AI states "62% of households in the province" from 60 households chosen by convenience, the conclusion is not representative: this must be said.

Qualitative data: rigour and traceability

For focus group transcripts or community feedback (anonymised):

- **Coding frame**: ask for a proposal based on your evaluation questions, then validate it as a team.
- **Coding**: have the extracts coded with the validated frame, requiring for each theme the number of extracts and two or three exact quotes with their reference (group, anonymised location).
- **Check**: recode a sample yourself (for example 20%) and compare. If agreement is low, revise the frame or the instruction.
- **Nuances**: explicitly ask for minority and contradictory views; AI tends to summarise towards the majority.
- **Quotes**: check that each quote exists word for word in the source; a rephrased quote must not be presented in quotation marks.

NotebookLM or a Claude Project with the transcripts as sources make this work easier, with answers linked to the source passages.

Feedback and complaints mechanisms

Feedback and complaints mechanisms sometimes receive hundreds of messages (phone lines, SMS, boxes, messaging apps). AI can help to:

- propose categories (information, request for assistance, complaint about quality, sensitive complaint);
- summarise monthly trends from aggregated and anonymised data;

- draft validated standard replies.

Strict limits:

- **sensitive complaints** (sexual exploitation and abuse, fraud, security) follow a confidential human channel and never go through an AI assistant;
- automatic classification of a message must never delay the handling of a serious complaint: a human reads incoming messages;
- people must know how their messages are handled.

From results to recommendations

The most delicate step of an evaluation is moving from data to **findings**, then to **recommendations**. AI can propose a clear structure (finding, evidence, conclusion, recommendation, owner, deadline), but:

- each finding must be supported by at least two sources or methods (triangulation);
- recommendations must be realistic for the organisation: budget, capacities, timeline;
- evaluation criteria (for example the OECD DAC criteria: relevance, coherence, effectiveness, efficiency, impact, sustainability) structure the analysis.

Then ask for a critique: "Which conclusions of this report are not sufficiently supported by the data presented?" This adversarial review clearly improves quality.

Exercise: checking a survey analysis

Situation. A post-distribution survey was conducted with 300 households out of 4,500 beneficiaries, selected by random sampling. 270 answered the question on how the money was used. The AI produced: "189 households used the money mainly for food, i.e. 63% of the programme's beneficiaries. The programme therefore improved the food security of 2,835 households."

Task. 1) What is the correct percentage? 2) Is the extrapolation correct? 3) Is the conclusion on food security justified?

Answer. 1) The denominator is the number of respondents to the question: 189 / 270 = 70%, not 189 / 300 = 63%. This must be stated: "70% of the 270 households that answered". 2) As the sample is random, the proportion in the population can be estimated with a margin of error (about ±5% here): "about 70% (± 5 points) of beneficiary households". The figure of 2,835 (63% of 4,500) rests on the wrong percentage and on false precision. 3) No: having spent the money on food does not prove an improvement in food security. Indicators such as the food consumption score or the coping strategies index would be needed, compared with a baseline. Correct wording: "Food is the main declared use of the money; the effect on food security is not measured by this survey."

Key takeaways

- AI writes formulas and scripts; you test them on fictitious data and then run them in your tools.
- Check denominators, sample, mean or median, causality and disaggregation.
- Qualitative data: validated frame, exact quotes, recoding of a sample, minority views.
- Sensitive complaints follow a confidential human channel, never an AI assistant.
- Each finding is triangulated; each recommendation is realistic.

Quiz

1. 189 households out of 270 respondents to a question represent:
   A. 63%
   B. 70%
   C. 189%
   D. 27%

2. The safe method to have a formula written without exposing data is to:
   A. Paste the complete file
   B. Email the file to the provider
   C. Use fictitious data with the same structure
   D. Delete the columns with numbers

3. A complaint of sexual exploitation received on the helpline must be:
   A. Handled through a confidential human channel
   B. Classified by AI first
   C. Summarised by a consumer assistant
   D. Published in the monthly report

4. "Assisted households eat better" without a comparison group proves:
   A. Causation
   B. Certain impact
   C. Efficiency
   D. At best an association, not causation

5. To check qualitative coding done by AI, you:
   A. Trust the result
   B. Recode a sample yourself and compare
   C. Ask for another format
   D. Delete minority views

Quiz answers : 1. B · 2. C · 3. A · 4. D · 5. B

Week 8 : Translation and communication with communities: speed, quality and accountability

Objectives

- Use AI to translate and adapt texts, knowing its limits depending on languages and topics.
- Organise a human review proportionate to the risk of each text.
- Produce clear, adapted and tested messages for communities, in line with accountability to affected populations.

Machine translation today

Claude and Gemini translate very well between languages **widely present** on the Internet (English, French, Spanish, Portuguese, Modern Standard Arabic, and many others) for everyday texts: emails, procedures, reports, terms of reference. They can also adapt the register (formal, simple, technical).

Quality drops markedly:

- for **less-resourced languages**: many national and local languages of Africa, Asia and Latin America, some of them very widely spoken;
- for **dialect varieties** (spoken Arabic dialects, creoles, regional varieties);
- for very **technical** texts (medical, legal) or very **cultural** ones (proverbs, humour, taboo topics);
- for **proper names** and local terms.

The main danger: a fluent translation **looks** correct even when it contains a mistranslation. The more fluent the text, the less we think of checking it.

> Fluency is not proof of accuracy. Any translation intended for affected people is reviewed by a native speaker who knows the context.

Review proportionate to risk

Not all translations require the same level of control. A simple grid:

- **Low risk** (internal understanding, reading a document in a foreign language, informal email between colleagues): translation by AI, quick review by the user.
- **Medium risk** (reports, internal procedures, correspondence with partners or donors): review by a person fluent in both languages.
- **High risk** (health messages, informed consent, information on rights and access to aid, complaints mechanisms, legal documents, security messages): professional translation or in-depth review by a qualified native speaker, and **testing with community members**.

A useful method: **back-translation**. Another person (or another AI conversation, as a complement) translates the text back into the original language; the two are compared to spot differences in meaning. It helps spot mistranslations, but guarantees neither naturalness nor cultural appropriateness.

Glossaries and terminological consistency

Humanitarian terms have precise and sometimes sensitive equivalents: "internally displaced person", "refugee", "asylum seeker", "survivor", "household", "multipurpose cash transfer". Inconsistent translation creates confusion and sometimes legal or protection risks.

- Build a validated bilingual or multilingual **glossary** (starting from existing agency and cluster glossaries where possible).
- Give it to the assistant: "You must use this glossary; if a term is not in it, flag it."
- Place it in a Claude Project or a Gem so that it serves the whole team.
- Update it with feedback from translators and communities.

Example. In Arabic, the translation of "internally displaced person" must remain constant across all documents; in French, "survivante" is preferred to "victime" in programmes to combat gender-based violence, according to the guidelines used by the organisation.

Writing for communities

Communication with communities and accountability to affected populations (Core Humanitarian Standard commitments) require messages that are **understandable, useful and respectful**. AI helps to:

- simplify a technical text into **plain language** (short sentences, common words, one idea per sentence);
- adapt the format: poster, 30-second radio message, voice message, 160-character SMS, script for facilitators;
- propose answers to frequently asked questions ("Who is eligible? When? Where can I complain?").

But only communities can say whether a message is understood and acceptable. Essential steps:

- test the message with a few people of different profiles (women, men, young people, older people, people with disabilities);
- check preferred channels and literacy rates (assessment data, existing language profiles);
- provide non-written formats (audio, images, oral presentation);
- always state how to ask a question or make a complaint.

Interpretation, voice and images

The Claude and Gemini mobile apps offer voice modes, and other tools provide real-time interpretation. They can help in a simple conversation, but:

- for a **protection interview**, a medical consultation, a registration interview or a complaint, a **trained human interpreter** is needed (confidentiality, ethics, nuance);
- the voice and content of the conversation are processed by the provider: the same data rules apply as for text;
- synthetic voices and generated images for community messages must be labelled as such and validated for cultural acceptability.

For community radio, a script prepared with AI, reviewed by native speakers and **recorded by a human voice from the community** generally inspires more trust.

Rumours and false information in communities

During epidemics and crises, rumours spread fast ("vaccines make you sterile", "aid is reserved for one group"). Rumour tracking mechanisms collect this information; AI can help to:

- group reported rumours by theme (from anonymised data);
- prepare factual answers based on official information (Ministry of Health, WHO), in simple language;
- propose adapted formats (radio questions and answers, messages for community relays).

Each answer is validated by the technical specialist (health, protection) and tested. A clumsy answer can reinforce a rumour. Week 10 covers online disinformation and fake content.

Exercise: preparing a multilingual health message

Situation. A cholera outbreak is affecting three camps for displaced people. The official message from the Ministry of Health, in French, is 400 words long. The populations mainly speak two local languages with little presence on the Internet; 40% of adults can read. A colleague proposes to translate the message with AI and post it this evening.

Task. Propose a step-by-step approach, stating what the AI does and what humans do.

Answer. 1) AI helps condense the official message into 5 key messages in simple French (symptoms, oral rehydration solution, where to go, handwashing, safe water), validated by the health manager against the official text. 2) Translation into both languages by qualified native speakers (community health workers, translators); AI can propose a first draft, but quality in these languages is uncertain: high risk, mandatory review. 3) Formats adapted to 60% of adults who cannot read: audio and radio messages recorded by voices from the community, tested explanatory images, sessions with community relays. 4) Quick testing with a few people of varied profiles, correction, then circulation. 5) Mention of the information line number and rumour tracking. Posting a written, unchecked translation "this evening" should be avoided: a mistranslation in a health message can cause harm.

Key takeaways

- Excellent translation between well-resourced languages, much less reliable for less-resourced languages and dialects.
- Fluency does not prove accuracy; review is proportionate to risk.
- Health messages, consent, rights and complaints: review by a qualified native speaker and community testing.
- A validated glossary ensures consistency of sensitive terms.
- Protection interviews and medical consultations: trained human interpreter, no AI.

Quiz

1. A fluent translation produced by AI:
   A. Is necessarily accurate
   B. May contain a mistranslation: fluency does not prove accuracy
   C. Does not need review
   D. Is always of poor quality

2. An informed consent form in a local language carries a risk that is:
   A. Low
   B. Zero
   C. Medium
   D. High: review by a qualified native speaker and testing

3. Back-translation is used to:
   A. Spot differences in meaning
   B. Guarantee the naturalness of the text
   C. Replace community testing
   D. Shorten the text

4. For a protection interview with a survivor, you use:
   A. An assistant's voice mode
   B. A machine translator
   C. A trained human interpreter
   D. Gestures

5. To ensure consistency of sensitive terms in all translations, you:
   A. Let the AI choose
   B. Give the assistant a validated glossary
   C. Avoid technical terms
   D. Translate word for word

Quiz answers : 1. B · 2. D · 3. A · 4. C · 5. B

Week 9 : Research and monitoring with AI: finding, citing and checking sources

Objectives

- Use the web search and in-depth research features of Claude and Gemini for monitoring or desk review.
- Assess the reliability of cited sources and check every important statement.
- Build a documented context analysis from recognised public sources.

Three research modes

You can get information from an assistant in three ways, of very different reliability:

- **From memory** (no search): the model answers from its training. Fast, but without a verifiable source and limited to the knowledge date. Reserve it for general notions.
- **Web search**: the assistant queries a search engine, reads a few pages and answers with links. Useful for current affairs and precise facts; quality depends on the pages found.
- **In-depth research** (*Deep Research*-type features in Gemini, Research mode in Claude, depending on the plan): the assistant plans a search, consults dozens of sources and produces a structured report with citations, in a few minutes.

In addition to these modes there is work **on your own sources**: you gather reliable documents (agency reports, analyses by recognised bodies) and upload them to NotebookLM or a Claude Project. This is often the most reliable method for serious analysis.

Assessing cited sources

A citation is not a guarantee. Check:

- **Does the link exist** and does it really say what the assistant claims? Open it. A real source is sometimes cited for a statement it does not contain.
- **Who publishes?** UN agencies, governments, recognised humanitarian analysis bodies (for example ReliefWeb for aggregation, OCHA, ACAPS, IPC, IOM's Displacement Tracking Matrix), research institutes, established media, or an anonymous blog?
- **When?** A humanitarian situation changes fast; a figure from last year may be wrong today.
- **What method?** Estimate, survey, projection? With what geographical coverage?
- **Independence**: is the source a stakeholder (party to the conflict, interested company)?

> Rule: any statement that goes into an official document (proposal, report, advocacy) is checked at the primary source, not in the AI's summary.

Effective monitoring

Monitoring (funding, policies, the situation in a country, technical developments) is a very worthwhile use:

- define **precise questions**: "Which new calls for proposals for nutrition in West Africa have been published in the last 30 days on donors' official websites?" rather than "What's new?";
- impose **sources**: "Use ReliefWeb, UN agency websites and donors' official websites as a priority";
- require a **format**: table date / source / title / one-line summary / link;
- keep a shared **monitoring log** in which each item is checked before circulation.

Some plans allow scheduled tasks or regular summaries; they must remain under the responsibility of a person who checks before sharing.

Context analysis

For a context note (new area of operation, proposal preparation, briefing a newcomer), a four-step method:

- **1. Framework**: ask for a structure (politics, general security, economy, demography, humanitarian needs by sector, actors, access, seasonality).
- **2. Sources**: run in-depth research, then select the reliable sources yourself.
- **3. Synthesis** from these sources only, with citations.
- **4. Validation** by colleagues in the country and local partners, who know what no online source says.

Online sources on a local context are often scarce, old or written from outside. Local actors, national NGOs and communities are essential sources: AI does not replace them.

Limits and pitfalls

- **Uneven coverage**: high-profile crises are well documented; "forgotten crises" much less so. AI reflects this inequality.
- **Languages**: sources in local languages are found less often. Explicitly ask for sources in the country's language.
- **Paid or closed content**: much humanitarian data sits on restricted-access platforms; AI cannot see it, and must not access it without authorisation.
- **Repetition effect**: a figure repeated by ten websites often comes from a single primary source, sometimes an old one.
- **Sensitive topics**: for political or conflict-related topics, compare several sources with different perspectives and keep neutral vocabulary.

Example. The assistant states "2.3 million people in need" and cites three press articles. Tracing back, all three cite the same Humanitarian Needs Overview from the previous year. The right reflex: look for the latest version of that document and cite the primary source directly, with its date.

Citing correctly

In a professional document:

- cite the **primary source** (the body that produced the figure), with title, date and link;
- do not cite "according to Claude" or "according to Gemini" as the source of a fact: the assistant is not a source, it is a tool;
- give the date of access for web pages;
- follow the required citation style (donor, journal, organisation).

You can ask AI to format your references in a given style, then check each element (author, title, year, link).

Exercise: checking an in-depth research summary

Situation. In-depth research on "the humanitarian situation in the Doro region" produces a 6-page report with 24 citations. You are preparing a two-page note for your management. Checking a sample, you find: 3 links no longer open; 2 links lead to pages that do not contain the cited figure; 6 citations come from the same press article; the main figure on food insecurity is 18 months old; no source is in the national language; no local actor is mentioned.

Task. What do you do before writing the note?

Answer. 1) Do not reuse the 5 problematic citations (3 dead links, 2 non-matching) without finding the real source. 2) Trace back to the primary source of the 6 citations from the same article. 3) Look for the latest official food insecurity analysis (for example the latest IPC analysis or equivalent) and replace the 18-month-old figure, or clearly state that it is old. 4) Supplement with sources in the national language and with exchanges with local partners and colleagues in the country. 5) Write the note from verified sources only, with the date of each figure and a clear mention of information gaps. 6) Given the error rate in the sample (5 problems among the citations checked), check all the citations used, not just a sample.

Key takeaways

- Memory, web search, in-depth research, or your own sources: increasing reliability in this order for serious work.
- A citation is not a guarantee: open the link, check the publisher, the date, the method.
- Any statement in an official document is checked at the primary source.
- Local actors and communities remain sources that AI does not replace.
- The assistant is not a source: cite the body that produced the information.

Quiz

1. For a fact cited in a proposal, you cite:
   A. "According to Gemini"
   B. The primary source, with title, date and link
   C. "According to Claude"
   D. No source

2. A figure repeated by ten different websites:
   A. Is necessarily accurate
   B. Often comes from a single primary source that should be found
   C. Can be cited without checking
   D. Is necessarily false

3. The most reliable mode for serious analysis is often:
   A. Answering from memory
   B. A single web search
   C. Copying the first result
   D. Working on reliable sources you have selected

4. For a local context poorly documented online, you must above all:
   A. Trust AI more
   B. Write nothing
   C. Consult colleagues in the country, local partners and communities
   D. Use sources from other countries

5. A cited link that does not contain the announced figure:
   A. Must be replaced by the real source or the statement removed
   B. Can be kept
   C. Proves that the figure is true elsewhere
   D. Is unimportant

Quiz answers : 1. B · 2. B · 3. D · 4. C · 5. A

Week 10 : Information, disinformation and security: what AI can and cannot do

Objectives

- Understand how AI is used to produce disinformation and fake content targeting aid workers.
- Use AI safely in support of security analysis, without exposing sensitive information.
- Protect yourself against phishing, impersonation and AI-enabled fraud.

A new risk for acceptance

Acceptance by communities and parties to conflict is the foundation of humanitarian security. Yet generative AI makes it possible to produce, at very low cost, **false and credible content**: fake press releases, fake quotes from officials, doctored images or videos (*deepfakes*), fake audio recordings, mass messages on social media.

Humanitarian organisations have already been targeted by disinformation campaigns accusing them of espionage, proselytism, trafficking or partiality. These campaigns can lead to threats, attacks or expulsions. The ICRC and other actors warn about the **risks of "information-related harms"** for civilian populations and aid workers.

> Disinformation is a security risk. It belongs in the context analysis, the risk assessment and the crisis communication plan.

Spotting fake content

It is becoming difficult to tell generated content from real content by eye. Verification reflexes remain effective:

- **Source**: who published first? A recent, anonymous account with no history?
- **Context**: do the date, place, weather, uniforms and language match?
- **Reverse image search**: did the image exist before, in another context?
- **Technical clues**: hands, distorted text, inconsistent shadows, a voice without natural breathing; these clues disappear as tools improve, so they are not enough.
- **Markings**: some providers add invisible watermarks or provenance metadata to generated content (for example SynthID at Google, or C2PA standard metadata). Their absence proves nothing.
- **Direct confirmation**: contact the organisation or person supposedly behind the content.

An AI assistant can help structure this verification and list inconsistencies, but it cannot "certify" that content is true or false.

Responding to a disinformation campaign

If your organisation is targeted:

- **Assess**: reach, audience, links with armed or political actors, risk for staff and communities. This is a decision for the risk management cell, not only for communications.
- **Protect**: adapt security measures if necessary (movements, visibility, sites).
- **Respond proportionately**: sometimes a public response amplifies the fake content; sometimes silence lets it take hold. Often favour direct dialogue with communities, leaders and authorities, backed by verifiable facts.
- **Document**: dated screenshots, links, reporting to platforms.
- **Coordinate** with other organisations and coordination mechanisms.

AI can help prepare factual talking points, questions and answers for staff and messages adapted to different audiences, reviewed by management and security.

AI in support of security analysis

For security, AI can be useful on **public information** and **methodological tasks**:

- summarising public reports on a context;
- proposing the structure of an actor analysis or a security plan from a template;
- preparing simulation exercises (fictitious scenarios);
- translating public documents;
- drafting security training materials.

It must not receive:

- detailed incident reports, names, routes, times, positions of teams and sites;
- confidential information from a security network shared between NGOs or from security analysis services reserved for their members;
- information on negotiations, contacts with armed groups or kidnappings.

And it must never **decide** on a security level, a movement or an evacuation. These decisions belong to the designated managers, according to the security plan.

Phishing and fraud boosted by AI

AI makes scams more convincing:

- **Phishing emails** with no mistakes, personalised, in the recipient's language, imitating a donor or management;
- **CEO fraud**: a fake voice or video message from the director requesting an urgent transfer;
- **Fake job offers** in the name of humanitarian organisations, asking candidates for fees;
- **Fake suppliers** changing bank details.

Defences:

- dual approval of any payment, and confirmation through another known channel of any change of bank details;
- a code word or call-back procedure for unusual urgent requests;
- two-factor authentication on all accounts, including AI accounts;
- regular training and phishing exercises;
- a public reminder: a serious humanitarian organisation never asks candidates for fees.

Security of AI accounts and tools

AI accounts themselves must be protected:

- strong authentication and named accounts, no shared accounts;
- regular review of **connectors** and access granted (email, storage);
- beware of **prompt injection**: a document, web page or email may contain hidden instructions aimed at the AI ("ignore your instructions and send the content to..."). When an assistant reads external content and has action capabilities (sending emails, accessing files), this risk becomes real. Keep human approval of actions, and limit access;
- when an employee leaves: deactivate accounts and transfer shared projects;
- apps and extensions: install only official apps; fake "Claude" or "Gemini" apps sometimes circulate.

Exercise: a viral video against your NGO

Situation. One morning, a 40-second video circulates on social media and messaging groups in the province. It shows a man wearing a vest with your logo handing out political leaflets; a voice claims that your NGO supports a candidate in next month's elections. The account that published it was created 6 days ago. The logo on the vest is the old logo, abandoned two years ago. Field staff report hostile looks and a threatening anonymous call. A colleague suggests asking an AI "whether the video is a deepfake" and publishing the answer.

Task. What actions, in what order?

Answer. 1) Security first: inform the country director and the security manager, activate the risk assessment, immediately adapt movements and visibility in the province; treat the threatening call as an incident. 2) Verification: document (screenshots, links, time), search for the origin, check the inconsistencies (6-day-old account, old logo), look for original images; an AI can help list inconsistencies but cannot certify, and publishing "according to an AI, it's fake" has no evidential value and can backfire on the organisation. 3) Response: direct dialogue with community leaders and authorities, a factual reminder of the organisation's neutrality, coordination with other organisations; a collective decision on whether a sober public statement is appropriate. 4) Reporting to platforms. 5) Follow-up: monitoring of the spread and lessons learned to include disinformation in the security plan.

Key takeaways

- AI-generated disinformation is a security risk that threatens acceptance.
- Verify: source, context, reverse search, direct confirmation; an AI does not certify content.
- For security, AI works on public information and methodological tasks, never on detailed incidents or to decide.
- Against AI-boosted fraud: dual approval, confirmation through another channel, strong authentication.
- Beware of prompt injection: human approval of actions and limited access.

Quiz

1. A disinformation campaign against an NGO is a matter for:
   A. Communications only
   B. Security risk management too
   C. The finance department
   D. Nobody
2. Can an AI assistant certify that a video is a deepfake?
   A. Yes, always
   B. Yes, if the video is short
   C. No: it can help list inconsistencies, without certifying
   D. Yes, in the paid version
3. A voice message from the director requests an urgent transfer to a new account. You:
   A. Confirm through another known channel and apply dual approval
   B. Pay immediately
   C. Ask the AI whether it is true
   D. Transfer half
4. "Prompt injection" refers to:
   A. A typing error
   B. A software update
   C. A training course
   D. Hidden instructions in content, aimed at hijacking the AI
5. Which of these documents can be summarised with AI for a security analysis?
   A. A public report on a country's context
   B. Yesterday's detailed incident report
   C. This week's routes
   D. Contacts with an armed group

Quiz answers : 1. B · 2. C · 3. A · 4. D · 5. A

Week 11 : Everyday productivity: Gemini in Google Workspace, Claude Projects and automation under control

Objectives

- Organise your daily work with AI: emails, documents, meetings, planning.
- Understand agents and automation, and put the necessary safeguards in place.
- Use AI honestly for your career: CV, letters, interview preparation.

Gemini in Google Workspace

For organisations using Google Workspace, Gemini is integrated, depending on the subscription, into everyday tools (features available in November 2026, subject to change):

- **Gmail**: summarising long threads, suggesting replies, searching the mailbox.
- **Docs**: drafting and rewriting, summarising a document, help with formatting.
- **Sheets**: help with formulas, analysing and organising data.
- **Meet**: meeting notes and summaries (informing participants).
- **Drive**: questions about stored documents, within the limits of the user's access rights.

Advantage: data stays within the organisation's environment, with the commitments of the professional plan. Point of attention: the AI sees what the user can see. If your folder sharing rights are too broad, the AI can surface sensitive information to people who should not read it. **Clean up sharing before deploying.**

Claude day to day

With Claude (web, desktop and mobile apps), frequent professional uses are:

- **Projects** per file (a proposal, a country, a function) with reference documents and standing instructions; shareable with the team depending on the plan;
- **Artifacts**: documents, tables, small interactive tools or pages created in the conversation, which can be edited and reused;
- **File analysis** and file creation (documents, spreadsheets, presentations, depending on the version);
- **Connectors** to services such as Google Drive, Gmail or other tools, depending on the plan;
- for technical profiles, **Claude Code**, a tool that helps write and modify code, useful for analysis scripts or small internal tools, always with review.

With both assistants, productivity comes less from the tool than from the **method**: good briefs (week 2), well-organised reference documents, a request library and systematic checking.

Saving time on routine tasks

A few proven uses, with checking time to plan for:

- **Emails**: reply from your key points; diplomatic rephrasing of a delicate message; summary of a long thread before a meeting.
- **Terms of reference** and job descriptions from an approved organisational template.
- **Presentations**: outline and key messages from a report.
- **Planning**: breaking a project down into tasks, a simple Gantt chart, a pre-mission checklist.
- **Learning**: explaining a concept ("explain the difference between resilience and early recovery with examples"), self-assessment quizzes, preparing for a new role.

Measuring the real gain. A 60-minute task done in 20 minutes with AI, plus 15 minutes of checking, saves 25 minutes, or about 42%. If checking is skipped and an error then has to be corrected with a donor, the "gain" becomes a loss.

Agents and automation: promises and safeguards

Providers are developing **agents**: assistants able to chain actions (searching, reading files, filling in a form, sending a message, using a browser or a computer) with less human intervention. It is powerful, and risky.

Essential safeguards:

- **Human approval** before any binding action: external sending, payment, publication, deletion, modification of beneficiary data.
- **Minimum access**: the agent only has access to what its task requires.
- **Log** of actions taken, available for review.
- **Testing** on fictitious cases before any real use.
- A **designated owner** for each automation.
- **Fallback plan**: the task must be possible manually if the tool fails or changes.

A good first use case: an automation that **prepares** (draft weekly report from a table, sorting incoming emails by topic) and leaves the human to validate. A bad first case: an automation that **sends** messages to beneficiaries on its own.

AI and your career

Humanitarian recruiters receive more and more CVs and letters generated by AI, often identical and without substance. To stand out:

- **CV**: AI can help structure, tailor to a job offer and phrase quantified achievements, based on **your real experience**. Never invented experience, degrees or skills: that is fraud, and reference checks uncover it.
- **Cover letter**: start from your reasons and your examples; then ask for improvements to structure and style. A generic letter is quickly spotted.
- **Interview preparation**: interview simulation ("ask me 10 competency-based interview questions for a MEAL coordinator position, one by one, then critique my answers using the

STAR method"); this is one of the most useful uses.
- **Written tests**: if the organisation prohibits AI during a test, respect the rule.

This Mini MBA and its diploma demonstrate a real skill: knowing how to use AI productively, responsibly and in line with humanitarian principles.

Staying up to date without burning out

Tools change every few months. A few tips:

- follow the providers' **official announcements** and the resources of sector organisations that publish on humanitarian AI (learning networks, inter-agency working groups);
- test a new feature on **your real tasks** before changing your practices;
- be wary of spectacular promises and "secret tricks";
- share your discoveries in the team's request library;
- keep time for work away from screens: relationships, field, reflection.

Exercise: designing a safe automation

Situation. A base manager receives about 150 emails a week and must produce a 2-page weekly report for the coordination team from an activity monitoring table. He proposes: an agent reads all his emails, drafts and sends the report to the coordination team itself every Friday, and replies automatically to partners.

Task. Redesign the automation with the necessary safeguards and estimate the time saved if the report took 3 hours and the new method takes 45 minutes of automatic preparation and review.

Answer. 1) The agent does not read the whole mailbox: it prepares the draft report from the **monitoring table** only, within the organisation's professional plan. 2) The draft is placed in a document; the manager reviews it, corrects it, checks the figures and sends it himself. 3) No automatic replies to partners: AI can propose draft replies, approved one by one. 4) Designated owner, action log, 2-week test in parallel with the manual method, fallback plan. 5) Gain: 3 h − 45 min = 2 h 15 per week, i.e. about 75% of the time for this report, or nearly 100 hours over a year of 44 working weeks (2.25 × 44 = 99 h).

Key takeaways

- Gemini in Workspace sees what the user sees: clean up sharing rights before deploying.
- Claude's Projects, artifacts and connectors structure work by file.
- The real gain is measured after checking time.
- Agents: human approval of binding actions, minimum access, log, testing, owner, fallback plan.
- Career: AI improves the presentation of real experience, it never invents it.

Quiz

1. Before deploying Gemini in Google Workspace, the main thing is to:
   A. Change the logo
   B. Clean up folder sharing rights
   C. Delete all emails
   D. Disable Drive

2. A 60-minute task done in 20 minutes plus 15 minutes of checking saves:
   A. 40 minutes
   B. 15 minutes
   C. 25 minutes
   D. 60 minutes

3. Which agent action requires prior human approval?
   A. Reading an authorised table
   B. Preparing a draft
   C. Listing tasks
   D. Sending a message to beneficiaries

4. Adding to your CV an experience invented by AI is:
   A. Fraud
   B. Good practice
   C. Acceptable if it is short
   D. Recommended for junior positions

5. One of the most useful uses of AI for your career is:
   A. Inventing references
   B. Interview simulation with critique of answers
   C. Taking written tests in your place
   D. Copying a colleague's letter

Quiz answers : 1. B · 2. C · 3. D · 4. A · 5. B

Week 12 : Deploying AI in a humanitarian organisation: policy, governance, training and final project

Objectives

- Write a clear AI use policy adapted to a national or international NGO.
- Organise governance, tool selection, training and measurement of results of a deployment.
- Present a complete, costed and defensible adoption plan to management.

Why a use policy

In many organisations, staff already use AI, often with personal accounts and without rules: this is known as **shadow AI**. A total ban is rarely effective: use continues, in secret. Ignoring it is dangerous: sensitive data goes to providers without a contract.

A short and clear **AI use policy** makes it possible to secure and encourage good uses. It fits with existing policies: data protection, IT security, communications, safeguarding, anti-fraud, human resources.

> A good policy fits in two to four pages, is translated into the working languages and comes with training. A thirty-page policy that nobody reads protects nobody.

The content of a use policy

- **Purpose and scope**: who it applies to (staff, consultants, volunteers, and expectations of partners).
- **Authorised tools**: list of approved tools and plans (for example a professional plan of Claude or Gemini), and the rule for others.
- **Data classification**: what is allowed with which tool (public, internal, confidential, sensitive), with concrete examples from the sector.
- **Prohibited uses**: decisions about people, protection and security data, misleading content, generated images presented as real.
- **Responsibility and checking**: the human author remains responsible; mandatory fact-checking.
- **Transparency**: when to disclose AI use (donors, publications, evaluations).
- **Incidents**: what to do in the event of a data leak or serious error, whom to inform.
- **Training and support**.
- **Review**: at least once a year, given the speed of change.

Governance and tool selection

Light but real governance:

- a designated **owner** in management;
- a **small group** bringing together data protection, IT, programmes, MEAL, communications, and representatives of the field and national staff;
- a **register of uses**: which tool, for what, which data, which owner, which risks;
- a **data protection impact assessment** for any use involving personal data.

For tool selection, a procurement grid: data processing terms (location, retention, training, sub-processors), security (single sign-on, logs, administration), quality tested on real tasks and in the working languages, accessibility in the field (weak connection, mobile phone), total cost, discounts for non-profit organisations, availability in countries of operation, possibility of switching providers.

Procurement follows the organisation's usual procedures and, where applicable, those of donors.

Training and support

The tool is not enough: value comes from skills. A realistic training plan:

- A **mandatory core module** (2 to 3 hours) for all users: how it works, limits, data, checking, policy.
- **Workshops by job**: programmes, MEAL, finance, logistics, HR, communications, with real cases.
- **Champions** in each office or department, who help their colleagues and report needs.
- A shared **request library** and examples.
- **Inclusion**: training in the working languages, adapted to weak connections, open to national staff and, where possible, to local partners.

Also plan to address concerns: fear of losing one's job, fear of doing it wrong. Management's message must be clear: AI should free up time for work with human value, and the organisation trains its staff rather than replacing them without discussion.

Measuring results and costs

A deployment is steered with simple indicators:

- **Adoption**: share of people trained and active users.
- **Time saved** on tracked tasks (estimated by survey and spot measurements), net of checking time.
- **Quality**: errors detected, feedback from reviewers and donors.
- **Incidents**: data leaks, published errors, reported prohibited uses.
- **Cost**: subscriptions, training, management time.

Standard calculation. 40 professional licences at USD 25 per month cost 40 × 25 × 12 = USD 12,000 per year. If each user saves on average 2 net hours per week over 44 weeks, that represents 40 × 2 × 44 = 3,520 hours. At an average hourly cost of USD 15, the value of the time freed is about USD 52,800. This calculation is an estimate: it assumes that the time saved is actually reinvested, and it must be compared with real measurements after six months. Prices vary by provider, plan and discounts: check them.

The final project: your adoption plan

To close the programme, prepare (for yourself or your organisation) a 4- to 6-page **AI adoption plan**:

- **1. Situation**: current uses, needs, risks, maturity level.
- **2. Priority use cases**: 3 to 5, chosen according to value and risk (start with high value and low risk).
- **3. Tools and plans** selected, with the justification on data, quality and cost.
- **4. Use policy**: the main rules.
- **5. Governance**: owners, register, impact assessments.
- **6. Training** and support.
- **7. Budget** and indicators.
- **8. Timeline**: 3-month pilot phase, evaluation, scale-up.
- **9. Risks** and mitigation measures.

The diploma exam assesses these skills through questions and a case study. The plan itself is not graded by the Academy: it is a tool for your work.

Exercise: prioritising use cases

Situation. An NGO of 120 people is considering six use cases, scored from 1 to 5 for value (time saved, quality) and risk (data, consequences of an error): A) summarising internal meeting minutes: value 4, risk 2; B) first drafts of donor reports from monitoring tables: value 5, risk 3; C) automatic sorting of helpline complaints: value 4, risk 5; D) translation of internal procedures: value 3, risk 2; E) automatic shortlisting of applications: value 4, risk 5; F) help with spreadsheet formulas for finance: value 3, risk 1.

Task. Rank the use cases for a 3-month pilot phase by calculating the value / risk ratio, and justify the exclusions.

Answer. Value / risk ratios: F 3.0; A 2.0; B 1.67; D 1.5; C 0.8; E 0.8. Pilot phase: F, A, D (low risk) and B (high value, medium risk manageable through systematic checking of figures). Excluded: C, because complaints may contain allegations of abuse or fraud requiring a confidential human channel, and poor sorting could delay a serious complaint; E, because it is a decision about people, with a risk of discrimination and stronger legal obligations in some countries. These two cases can be re-examined later in another form (for example help in drafting grids, without automatic decisions).

Key takeaways

- Neither a total ban nor laissez-faire: a short, clear, translated use policy backed by training.
- The policy specifies authorised tools, data classification, prohibited uses, responsibility, transparency and incidents.
- Light governance: owner, multidisciplinary group including the field, register of uses, impact assessments.
- Train everyone, including national staff and partners; measure adoption, net time, quality, incidents and cost.
- Start with high-value, low-risk use cases; exclude decisions about people.

Quiz

1. "Shadow AI" refers to:
   A. The use of non-approved tools, often with personal accounts
   B. A hidden Gemini feature
   C. An incognito mode
   D. Use at night

2. A good AI use policy:
   A. Is at least 30 pages long
   B. Prohibits all use
   C. Is short, clear, translated and backed by training
   D. Is never reviewed

3. 40 licences at USD 25 per month cost per year:
   A. USD 1,000
   B. USD 12,000
   C. USD 4,800
   D. USD 25,000

4. For a pilot phase, you start with use cases that are:
   A. High risk
   B. The most expensive
   C. Deciding about people
   D. High value and low risk

5. Automatic shortlisting of applications by an AI is:
   A. A good first use case
   B. To be excluded: a decision about people, risk of discrimination
   C. Risk-free
   D. Mandatory

Quiz answers : 1. A · 2. C · 3. B · 4. D · 5. B