HUMANITARIAN ACADEMY · HUMANITAIRES JOBS

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

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

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/semaine-1/

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

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