



Freerange Creatives (Pty) Ltd

Artificial Intelligence Use Policy

Version 2.0

Date: 28-08-2026

Owner: Chief Digital, Data & AI Officer

1. Purpose

This policy sets out how people at Freerange Creatives use generative AI tools in their work. It exists to let us use these tools well - with speed, curiosity and good judgment - while protecting our clients, our people and the trust both place in us.

AI is a tool, not a colleague and not a decision-maker. It drafts, suggests and accelerates. It does not carry accountability. A named human always does.

2. Our principles

The rules that follow flow from a handful of beliefs. When a rule doesn't cover your situation, these should.

- **People before outputs.** Real people sit behind the data we handle and the work we make. Their dignity, privacy and safety come before speed or cleverness.
- **Honest about what's real.** We don't pass off synthetic work as human, invented data as fact, or a machine's fluency as our own thinking. If something was AI-assisted and it matters, we say so.
- **Judgment doesn't get outsourced.** AI can make everything sound strangely similar when people stop thinking for themselves. We use it to sharpen our thinking, never to replace it.
- **Match the human to the stakes.** The weight of the judgment should match the weight of the decision. Small, reversible, low-stakes calls need a light touch. Uncertain, consequential or sensitive ones need real expertise and real time.
- **Fair to the people in the data.** We watch for bias and exclusion, especially in anything public-facing or anything touching vulnerable people. Fairness is a duty we act on, not a box we tick.
- **Mindful of the cost.** AI has an environmental and resource cost. We use it where it adds real value, not reflexively for its own sake.

3. Scope

This policy applies to everyone who does work for Freerange - employees, interns, freelancers and contractors - whenever they use AI tools for anything connected to the agency or its clients. It covers our commercial work and our social-impact work, including Kozani.

4. Approved tools and access

Only the following tools are approved for work use:

- **Claude** - any tool included in Freerange's paid Claude subscription (chat, Projects, and other features that form part of the plan), accessed through Freerange company accounts only.
- **Gemini** - the paid version, accessed through Freerange company accounts only.

Two rules are absolute:

- Company access only. Personal accounts and free tiers are not permitted for work. Free and consumer tiers may use your inputs to train models; our paid, company-managed accounts are configured not to.
- No other AI tools without approval. Any tool not named above - including image, video, audio, transcription or coding tools - must be approved by the CDDAIO before use on client or agency work.

5. What may and may not go into AI tools

The single most important rule in this policy:

Never enter personal identifiers into any AI tool.

That means no names, email addresses, phone numbers, ID or passport numbers, physical addresses, customer records, patient or user data, or any other information that could identify a living person - whether a client's customer, a member of the public, a colleague, or a Kozani user. Data must be de-identified before it goes anywhere near a model.

Allowed

- General, non-confidential material: public information, marketing concepts, draft copy, strategy thinking, brainstorming.
- Aggregated or anonymised data with all personal identifiers removed.
- Freerange's own internal, non-sensitive documents.

Restricted - de-identify or seek approval first

- Client data of any kind. Strip personal identifiers first. If it cannot be meaningfully de-identified, do not use AI on it without CDDAIO approval and a check on the client contract.
- Commercially sensitive client information (unreleased campaigns, pricing, strategy).

Prohibited

- Any personal or special-category data, in identifiable form.
- Kozani perinatal, clinical, mental-health, GBV-related or minor-related data of any kind. Kozani's own product use of AI runs under its own clinical and data governance - this policy is about staff tools, not the product.
- Third-party confidential information we are not authorised to share.

6. Client disclosure

We are open with clients about our use of AI. We tell clients that AI tools form part of how we work, and we disclose AI use on their work as a matter of course.

- Where a client contract, NDA or brief sets its own rules on AI, those rules override this policy for that client. Check before you start.
- Some clients or sectors require explicit consent or prohibit AI use entirely. When in doubt, ask before using AI, not after.
- Never present AI-generated work to a client as something it is not. No fabricated data, no invented sources, no synthetic content passed off as original research or human-made assets without disclosure.

- Standard disclosure wording is included in proposals and statements of work, and the CDDAIO maintains a list of clients with stricter AI terms.

7. Human accountability and review

AI output is a first draft, never a finished deliverable. A named person owns every piece of work and is responsible for checking it before it goes anywhere. But not every check is the same check - and the mistake to avoid is treating a serious judgment like a quick box-tick.

Match the review to the stakes. Three questions decide how much human judgment a piece of work needs:

- **How uncertain is it?** The less obvious the right answer, the more time someone should spend working through it.
- **How big are the consequences?** The larger the potential harm if it's wrong, the more it needs real expertise - not just any available pair of eyes.
- **How reversible is it?** Something easy to catch and undo can be checked lightly. Something that can't be taken back needs care before it goes out.

In practice: anything public-facing, anything touching vulnerable people, and anything commercially sensitive to a client sits at the high-stakes end. These get genuine expert review and genuine thinking time - not a glance. Routine, low-risk, easily-fixed work needs a lighter check. Don't spend your best judgment on a checkbox, and don't leave a serious call to one.

A human stays in the loop. Where automation or agentic workflows are involved, a person remains responsible for the outcome - and that involvement can sit in different places. Sometimes the work pauses for a person to approve it before it goes out. Sometimes a person sets the guardrails up front and reviews a sample of results afterwards. The more automated and repeatable the workflow, the more the human's real value lies in shaping it well and auditing what it produces, rather than rubber-stamping every item one by one. The test is the same: more uncertainty, consequence or irreversibility means a human closer to the decision.

- Check facts, figures, quotes and citations. AI invents things confidently. Verify before you rely.
- Check for bias and tone, especially in anything public-facing or sensitive.
- The person who puts their name to the work owns its accuracy - not the tool, and not the person who ran the prompt on their behalf.

8. Compliance

Freerange operates across South Africa and the Netherlands, so both POPIA and the GDPR apply to how we handle personal information.

- Data minimisation: only use what you need, and keep personal data out of AI tools entirely.
- Personal data is processed under a lawful basis and within our existing data-protection obligations - putting it into a model is not one of those permitted uses.
- This policy sits alongside our confidentiality, data-protection and IP obligations, not above them.

9. Intellectual property and confidentiality

- Do not paste in third-party confidential or copyrighted material as input without the right to do so.
- Treat AI-assisted outputs as drafts that require human authorship and judgment before they become Freerange work product.
- Ownership and licensing of AI-assisted deliverables follow the terms of the relevant client agreement.

10. Prohibited uses

Regardless of tool or context, the following are never acceptable:

- Entering personal identifiers or special-category data into any AI tool.
- Deepfakes, impersonation, or synthetic likeness or voice of any real person without explicit consent.
- Fabricating data, evidence, citations or results.
- Passing off undisclosed AI-generated content as human-made where disclosure is expected or required.
- Using AI in ways that discriminate, demean or cause harm.

11. Governance and breaches

- The CDDAIO owns this policy and is accountable for AI governance at Freerange.
- If you are unsure whether something is allowed, ask first. Good-faith questions are always welcome.
- Report any breach - a personal identifier entered by mistake, a client term overlooked, an output that went out unchecked - to the CDDAIO promptly. Early disclosure lets us contain harm; a mistake owned quickly is a smaller problem than one hidden.
- This policy is reviewed at least annually, and sooner as tools and regulation change.

This is a living document. As the tools change and we learn, so will this document.

Questions? Please contact: danielle.timmins@freerangecreatives.co.za