Responsible AI Use Policy

The Change Climate Project

Last Updated: 06 April 2026

The Change Climate Project (TCCP) uses AI in our operations and products. This document explains how we use it and describes the steps we take to do so responsibly.

TCCP's Director of Engineering is the policy owner and approves any policy exceptions. Questions about this policy should be sent to tech@changeclimate.org.

What This Policy Covers

This policy applies to anyone working on behalf of TCCP including internal staff, contractors, and volunteers. It covers any AI or machine learning system we use, build, or integrate into any of our tech services, such as the Business Emissions Evaluator (BEE), our workflow automations, data pipelines, and the LLMs and code assistants we use day-to-day.

"AI" in this Policy means: anything that generates predictions, recommendations, classifications, or content using statistical models or machine learning. If a vendor tool has an AI feature baked in, that counts too.

Why We Have This Policy

We are a nonprofit whose programs help boost confidence in companies' climate initiatives, and trust is a cornerstone of our work. Moreover, companies trust us with their data. To maintain this trust, we must be transparent about how we use data, create content, and make decisions. Careless use of AI will hurt our work and could interfere with our ability to adhere to our own terms of use and commitments to protect personal, confidential, sensitive information.

The field of AI is evolving rapidly, and as it evolves, is introducing many new risks, ranging from huge energy consumption to the many poorly defined ethical frontiers that users of AI must face. Because of this, it is important for us to be clear not just about how we use AI, but how we make decisions about how we use AI.

AI Code of Conduct

These are the principles we uphold across all AI use:

- **Keep humans in the loop. Higher stakes, more humans.** AI supports our work. It doesn't replace judgment, and it doesn't get the final word on decisions that affect our programs or stakeholders. AI-generated or AI-assisted content that goes out under TCCP's name is reviewed by at least one human before release. No exceptions.

- **Make AI attribution clear and obvious.** When AI contributes meaningfully to content, say so, and be explicit about how AI was used. This is especially important for external content.
- **Minimize unnecessary compute.** AI workloads consume energy. We use AI where it serves the mission and greatly accelerates our throughput – and skip it where it doesn't. We track compute usage and estimate associated emissions using provider data and grid intensity factors. Use lighter approaches where they get the job done.
- **Stay current and responsive.** AI capabilities and risks change fast. Revisit our tools, practices, and this policy regularly for needed updates.

Types of AI Use

We use AI in four ways across the organization. A single AI tool may support multiple use cases.

- Generative: AI drafts content that we use in blog posts, email drafts, program documentation, code suggestions.
- Analytic: AI processes, compares, or summarizes data and synthesizes takeaways, such as trend analysis, document comparison, or market research.
- Creative: AI produces a finished artifact: an image, a block of copy or a design that is used as-is or with minimal editing.
- Agentic: AI takes actions or makes decisions, as in the case of automated data validation, or workflow triggers.

Risk Tiers

AI use in our work may introduce different risks (High, Medium, or Low) based on three factors:

- **Whether the output could affect the integrity of the Climate Label Certification**
    - e.g.: Could the output change whether a company gets certified, affect disclosed data, or shape how The Climate Label is perceived?
- **The sensitivity of underlying data**
    - e.g.: Does the system process company-attributable emissions data, financials, or PII? Use of public or internal-only data carries less risk.
- **The potential 'blast radius' if the output is wrong**
    - e.g.: If an output is wrong, who notices and how bad is it? A wrong number in a certification report affects a company and public trust in the label. An automation error in an internal workflow might propagate before someone catches it. A bad first draft of an email wastes a few minutes but can be easily corrected before sending.

	Risk Tier		
	High	**Medium**	**Low**
Integrity	Significant integrity consequences, large negative impact in case of errors	Moderate but manageable risk to integrity / trust	Little or no effect on integrity
Data	Sensitive data at stake, including company-specific emissions data or PII	Any sensitive data is anonymized and unattributed or aggregated	No sensitive (e.g., financial, personal) data in play
Blast Radius	Large (100s+) audience for output	Mid-sized audience (5-100) with risk contained	Audience limited to small group (e.g. 1-5 people)
Who's affected	Certified or certifying companies; everyone relying on the integrity of The Climate Label	Internal staff; potential to affect broader audience if errors propagate	Internal staff; AI outputs not suited to external consumption without significant human review
Examples	Emissions calculations in the BEE; certification eligibility; automated data validation feeding certification decisions	Trend analysis/reporting; workflow automations; data pipeline transforms; outreach optimization	LLMs for drafting; code assistants; meeting summaries; internal research
Review and oversight	Human review of every output before it reaches an external stakeholder or user. Full audit trail on inputs, outputs, and reviewer decisions.	Careful human review before deployment and after significant changes. Logging of inputs and outputs.	Human review to assess meaning and relevance.
Disclosure and attribution	Disclose use of AI and attribute specific tools/models.	Disclose use of AI and attribute specific tools/models if such information is relevant to the task output.	Disclose use of AI through standardized footnote.

Example #1: AI in The Climate Label Certification

Risk Tier: High

This is the high-stakes stuff since it always incorporates company data and affects the integrity of our certification. Rules are simple:

- **Humans make certification decisions.** AI makes recommendations or assists. A person approves. **Always.**
- AI outputs in reviewer workflows are labeled as recommendations, not conclusions.
- Any company can request a fully human review of anything AI-influenced at no extra cost.

Example #2: Drafting blog content

Risk tier: Low/Medium

AI can accelerate throughput and shorten turnaround times; severity of risk depends on the nature of the communication including the subject matter and any opinions/recommendations that are being put forward.

- Always have a human review before anything goes public
- Attribute AI usage in public materials

Example #3: Code engineering assistance and review

Risk tier: Medium/High

AI can support code development process speed and quality; risk due to potential downside from errors that propagate.

- Always have a human review any code changes before being pushed to production
- Research and brainstorm with publicly available info
- Watch out for hallucinations; verify sources

Example #4: Internal meeting notes

AI can streamline notetaking with limited downside other than potential factual misrepresentation

- Footnote meeting notes with AI attribution and note whether a human reviewed them

Third-Party AI Vendors

Before onboarding any AI vendor or tool that will process TCCP or company data, the individual that manages the AI policy must evaluate the vendor's privacy and security practices against the following criteria:

1) **Must pass. Any failure is disqualifying:**

- The vendor does not use our data to train or improve their models, or provides a default opt-out that we can verify.
- The vendor can delete our data on request within 30 days, consistent with our own deletion commitments.
- Data is encrypted at rest and in transit.
- The vendor's terms are compatible with our Terms of Service and our obligations to certified companies.

2) **Must be acceptable. Evaluated case by case:**

- Data is stored in a jurisdiction consistent with our operational and legal requirements.
- The vendor holds SOC 2, ISO 27001, or equivalent security certification or can demonstrate comparable practices for their scale and maturity.
- Data retention periods are reasonable and documented. Shorter is better.
- Subprocessors, if any, are disclosed and meet the same data handling standards.
- The vendor has a documented breach notification policy with a defined timeline.

Review Schedule

This policy is reviewed annually, or sooner if:

- we deploy a significant new AI system
- provider practices change materially
- regulations change materially
- we have a notable incident
- someone on the team flags a gap

Changelog

v1.0 – April 2026 – initial policy.