

A Conscious AI Charter

Five commitments for technology that serves people.

A working framework by Tiffany Joy Basse | Mothers Shaping AI

October 2026 | Working draft for discussion

What conscious means here

Conscious AI means developing, deploying and governing artificial intelligence with deliberate attention to its human consequences. It does not mean that present-day AI systems are conscious.

The charter is an invitation to ask better questions, especially before a system affects people who may not be able to challenge it. It is not a certification, a technical audit, legal advice or a statement that outside organizations endorse the initiative.

The commitment

We ask organizations to make AI decisions that protect human agency, take children seriously, demonstrate evidence and accountability, include affected communities, and help people flourish without coercion.

The standards and guidance cited later are independent reference points. The five commitments below are Tiffany Joy Basse's proposed synthesis for public discussion.

Principles that lead to decisions

01 Human agency

People should be able to understand how a consequential decision was reached, question its basis, and obtain meaningful human review when warranted.

Practical questions: Can an affected person find the responsible human? Can an important result be appealed, corrected or reversed?

02 Care for children

Systems used by children and teens require special attention to privacy, development, learning, emotional dependency, safety and family agency.

Practical questions: Has the system been tested for the age and setting where it will be used? What extra protections apply to minors?

03 Evidence and accountability

Capability claims must be accompanied by realistic testing, disclosure of limitations, monitoring in use, and clear ownership of failures and repair.

Practical questions: Who can pause deployment? What evidence would change the decision? Who is accountable for follow-through?

04 Meaningful public participation

People affected by AI should have practical opportunities to influence its use and governance, before and after consequential decisions.

Practical questions: Whose perspective is absent? Did public or stakeholder input change a decision, or was it only collected?

05 Flourishing without coercion

AI should support learning, creativity, relationships, dignity and choice, without imposing one vision of the good life or quietly eroding autonomy.

Practical questions: Can a person disagree, opt out or choose an alternative without unreasonable penalty?

Questions for decision-makers

Builders and product teams. What shapes the outputs? Which use cases are outside acceptable limits? What risks are evaluated, reported and corrected?

Employers and institutions. Would a simpler process work? Who signs off, reviews outcomes, trains people and handles appeals?

Parents, caregivers and educators. What will children learn to expect from the system? How is personal information protected? How can adults raise a concern?

Policymakers and the public. What kinds of independent oversight and public input are missing? How do those affected gain a real voice?

Independent reference points

These are external resources for further reading. Their inclusion does not imply endorsement of this charter or initiative.

1. [NIST AI Risk Management Framework Core](#) (2023): the Govern, Map, Measure and Manage functions.
2. [NIST Generative AI Profile](#) (2024): additional considerations for generative AI.
3. [UNICEF Guidance on AI and Children](#), version 3.0 (2025).
4. [UNESCO Recommendation on the Ethics of Artificial Intelligence](#) (adopted 2021).

Begin a conversation

Use these principles to identify a real AI decision affecting people around you. Ask what is known, what remains uncertain, who can change the system, and what affected people need. Discuss the result rather than merely endorsing a slogan.

omfund.vc/books/conscious-ai.html