

The AI Ethics Self-Assessment

A free instrument from AI GOVERNANCE CHAIN™

Most organizations can name their AI ethics principles. Far fewer can prove they are practiced. This self-assessment closes that gap in twenty minutes. It asks the questions an ethics board, a regulator, or a journalist would ask, and it tells you, honestly, where you stand before someone else decides for you.

This is the free WHAT and WHY. It shows you the questions and why they matter. The paid AI Ethics Operating Kit gives you the HOW: the charters, templates, and review machinery that turn answers into practice.

Abbreviations and Acronyms

No prior knowledge is assumed. Every acronym used in this document is defined below.

AI: Artificial Intelligence

How to use this

Score each statement from 0 to 3.

0 = not true for us. 1 = partly true, informal. 2 = mostly true, documented. 3 = fully true, evidenced and reviewed. Be honest. A flattering score helps no one. Total the five sections for a score out of 60.

Section 1: Fairness and Non-Discrimination (out of 12)

1. We test our AI systems for biased or discriminatory outcomes before they go live.
2. We know which groups could be disadvantaged by each system, and we monitor for it.
3. When we find unfair outcomes, we have a defined way to correct them.
4. Fairness is treated as a moral obligation, not only a legal one.

Section 2: Transparency and Explainability (out of 12)

5. People affected by an AI decision can find out that AI was involved.
6. We can explain, in plain language, how a consequential AI decision was reached.

3. 7. We disclose the limits of our AI systems, not only their strengths.
4. 8. We do not overstate what our AI can do to customers, staff, or the public.

Section 3: Human Dignity, Autonomy, and Consent (out of 12)

1. 9. People can question or appeal an AI decision that affects them, to a human.
2. 10. We obtain meaningful consent before using personal data to train or run AI.
3. 11. Our AI supports human judgment rather than quietly replacing it where it matters.
4. 12. We have considered the human cost of automation on our own workforce.

Section 4: Accountability (out of 12)

1. 13. A named person is accountable for the ethical conduct of each AI system.
2. 14. We keep records that would let an outsider audit how a decision was made.
3. 15. There is a clear escalation path when an ethical concern is raised.
4. 16. Raising an ethical concern is safe and welcomed, not career-limiting.

Section 5: Societal and Environmental Responsibility (out of 12)

1. 17. We have asked whether each AI system should exist, not only whether it can.
2. 18. We consider dual-use and misuse risks before deployment.
3. 19. We weigh the environmental cost of training and running large models.
4. 20. We engage the communities most affected by our highest-impact systems.

Reading your score

- - 48 to 60. Principled and practiced. Ethics is operational, not aspirational. Your task is to hold the line as you scale. The paid kit helps you institutionalize and prove it.
- - 30 to 47. Committed but informal. You mean well, and it shows in places, but practice is uneven and under-documented. This is the most common and most exposed position. The Operating Kit is built for exactly this gap.
- - 12 to 29. Aspirational only. You have values on a page and little machinery behind them. A single incident would find you unprepared. Start with the kit and the Pre-Scale Audit.
- - 0 to 11. Unprotected. Ethics is not yet governed. This is urgent, and it is fixable.

The honest principle behind this tool

An organization that scores itself a perfect 60 on the first try is not finished. It is not yet looking hard enough. Ethical maturity begins the moment you can name where you fall short and act on it. That is the difference between an ethics statement and an ethics practice.

Where this points

This assessment reveals the gaps. Closing them is the work. The paid AI Ethics Operating Kit gives you the review board charter, the ethical decision templates, the ethics risk register, and the board reporting pack that move each of these twenty questions from a 1 to a 3. When the stakes are highest, the Pre-Scale Governance Audit brings decades of professional experience of judgment to the room.

Free gives you the questions. Paid gives you the practice.

Your next step: Get the [AI Ethics Operating Kit, Full Edition](https://www.aigovernancechain.com/storefront.html) in the Execution Library (www.aigovernancechain.com/storefront.html, section: The Core Governance Kits). Every instrument in the Execution Library is also included in the [Governance Vault](#), the all-access annual licence.

AI GOVERNANCE CHAIN™ | Human-First AI Governance

(c) 2026 Mathieu K. Gouanou. All rights reserved. Member, Harvard Business Review Advisory Council.

Document standard: human-first AI governance.