AI/ML Code of Ethics for Practitioners

10 RULES. ONE PAGE. NO EXCUSES.

RULE		PRINCIPLE IN THE ACM (ASSOCIATION FOR COMPUTING MACHINERY) CODE OF ETHICS, 2018
01	**Put people first.** Do not build a system if its primary effect is harm.	ACM 1.1, 1.2
02	**Tell the truth about AI.** People must know when they talk to an AI or read its output. Do not promise more than the model can do.	ACM 1.3
03	**Test for unfair results.** Test before release. Test again after each change.	ACM 1.4
04	**Respect the data.** Collect the minimum. Get consent. Do not train on private data without permission.	ACM 1.6, 1.7
05	**Respect the work of others.** Check the license on data, models, and content before you use them.	ACM 1.5
06	**Keep a human in the loop.** A person approves each decision that affects money, health, employment, or legal rights.	ACM 2.5, 3.7
07	**Evaluate before and after release.** Record the known risks and limits. Give them to the people who use the system.	ACM 2.5
08	**Stay in your area of competence.** If you do not know, say so.	ACM 2.2, 2.6
09	**Make the system secure.** Plan for prompt injection, data leaks, and misuse.	ACM 2.9
10	**Accept review and correct errors.** Give people a way to report a bad result. Then fix it.	ACM 2.4

You are responsible for what you build. The model is not.

Send comments and suggestions to info@mindfulware.com

An independent companion to the ACM (Association for Computing Machinery) Code of Ethics and Professional Conduct. ACM did not write, review, or endorse it. Principle numbers refer to the 2018 version: acm.org/code-of-ethics

LICENSE

This work is licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). You can copy, change, and share it if you give credit to Daniel T Joye. License: creativecommons.org/licenses/by/4.0

Daniel T Joye, ACM Professional Member. November 2026.