AI Human Ethics

AI Ethics

Companion Guide

Intro: Use AI Without Losing Yourself

Artificial intelligence can multiply your abilities, but it can also create blind spots. The more you rely on AI for ideas, writing, images, and decisions, the more important it becomes to think about ownership, accuracy, bias, privacy, and its impact on real people.

Ethical AI use is not just about complying with laws. It is about protecting your integrity, the trust others place in you, and the people who may be affected by what you create. This guide expands on seven key principles: ownership, accuracy, bias, safety, harm, hidden costs, and preserving your own brain. Keeping these principles in mind when using AI can help you stay human in an AI world while maximizing its benefits in your work and life.

1. Ownership: Making the Work Truly Yours

Why ownership is complicated

AI tools can generate text, images, code, and music in seconds. But the legal and ethical question, "*Who owns this?*" is still evolving. Purely AI-generated output is not clearly protected by copyright. Human authorship and meaningful creative contribution are required to transform AI output into individual intellectual property.

If you simply copy-paste AI output and present it as your own, you risk both plagiarism and an unclear legal claim over the material. Ethically, you are borrowing the work of countless human creators whose data helped train the model, without acknowledgment. There's nothing wrong with using AI to accelerate your original ideas but claiming AI output as your own thinking is wrong and deceptive.

How to use AI ethically for ownership

Treat AI output as raw material, not finished work. Edit, reorganize, and inject your own voice, decisions, and examples. Transform the original output to fully be ownable (and defendable) by you. If you can't explain every creative choice, then it's not your work.

When stakes are high (school, client work, publication), be transparent that AI assisted you, especially if a teacher, employer, or policy requires disclosure.

Build habits that keep your name attached to genuinely original thinking: add your own insights, stories, and perspectives that no model could invent for you. Ultimately, using AI *should* lead to a lot more work. You will make substantial changes but AI has helped you be more intentional and focused with those changes and may even inspire new ideas.

Key principle: Own your work by transforming AI output significantly. By using your own creativity and judgment to evolve, refine, build, and craft, the final product is 100% yours; AI simply becomes one of many tools you used to create that ownable final product.

2. Accuracy: Trust, but Verify

AI hallucinations and confident errors

Large language models are trained to be experts, but are not inherently truthful. They can generate statements that sound correct, even impressive, but are sometimes completely fabricated so that the AI still seems to be the expert… These are referred to as hallucinations. Studies and user reports show that even advanced models still produce incorrect or misleading answers in a significant share of responses. The danger is that AI sounds equally confident and reassuring with all answers. **Remember that AI is predicting, not thinking.**

A great example is when AI generates URLs based on a typical URL structure. When you ask for sources from AI, always make sure the URL works and that it corresponds to the information cited in the AI output. Similarly, AI fabricates data points that seem correct but are often unsubstantiated. Don't hesitate to ask AI to cite all of its output. You might be surprised by how often AI admits it was making up data and online resources. When pressed, AI will explain that its outputs are very close to actual data and research but not true and provable themselves.

When you copy AI-generated content into emails, slide decks, or assignments without checking it, you transfer this risk of inaccuracy and possible misrepresentation directly to yourself. One impressive-sounding false detail can damage trust with your audience and jeopardize your reputation.

Practical accuracy checks

Cross-check facts with at least one independent, reputable source (article, book, database, etc.) rather than relying solely on the model.

Ask the model to show its sources or reasoning, and then inspect those sources yourself.

Be extra careful with numbers, statistics, legal claims, medical information, and anything that could affect someone's health, money, or reputation.

If something feels slightly off, treat that as a signal to verify before you share.

Key principle: Assume AI is capable but not perfect. Double-check anything that matters before you put your name on it.

3. Bias: Human Programming Influences AI Models

Where bias comes from

AI systems learn from human data—text, images, and real-world behavior. That data reflects existing inequalities, stereotypes, and blind spots. As a result, models can unintentionally reproduce or amplify harmful patterns: rating some groups as higher risk, showing certain jobs mostly to one gender, or describing cultures in stereotyped ways.

Be careful, thoughtful, and objective about how inclusive, equitable, and modern AI answers are. Be vigilant and aware that AI training bias can cause AI answers to skew towards one group or belief system, or omit a group or belief system in ways that are not obvious.

How to read AI answers critically

Pay special attention when your prompts involve people, identity, culture, or sensitive topics.

Ask yourself: *If someone from this group read this output, would it feel fair and respectful?*

If you notice biased or harmful language, correct it explicitly in your prompt so you tell the model what to avoid next time. This improves your immediate output and helps you build toward more ethical awareness when using that prompt with AI in the future.

When in doubt, seek human perspectives, especially from people who might be affected by the decision or content.

Key principle: Never accept AI's description of people, culture, or beliefs at face value. Look for those who might be misrepresented or harmed by using empathy.

4. Safety & Privacy: Guarding What You Share

Why privacy rules matter with AI tools

When you type into an AI chat box, you are talking to a system, not a friend. Many services log inputs for quality, safety, or training unless you explicitly opt out. That means sensitive information can remain on servers indefinitely, be reviewed by staff at any time, or potentially be exposed in security breaches.

Even if a tool promises strong security, sharing your own personal or confidential details increases the risk of your or others' personally identifying information (PII) being leaked or stolen.

Practical privacy guidelines

Do not paste passwords, Social Security numbers, full addresses, medical records, financial account details, or any other sensitive information into AI tools.

Avoid sharing other people's identities or private stories without consent.

Treat anything you type as if it could one day be seen by someone else—because in some systems, it can.

AI is not therapy

For mental health crises, conflict, or trauma, seek support from trusted family, friends, and professional services, but never from an AI model.

LLMs can feel very natural, as if speaking with another person. Their speech, pacing and intonation make us feel like they are really listening, even empathizing.

But AI is simply calculating and predicting and massive scale and speed based on its training and inputs. **AI does not feel, judge nor can it relate to human struggle of any kind…**

Key principle: Never confide in AI. Anonymize what you tell it and avoid sharing sensitive information about yourself or others at all costs. Treat AI as a tool, no matter how "human" it may seem.

5. Images & Harm: Deepfakes and Respect

The impact of harmful images

AI image tools can generate photorealistic pictures and videos, including deepfakes of real people. "Deepfakes" are convincingly real images of actual people, but most often represent them in unsavory or defamatory ways. Misusing AI image tools, especially to create sexualized, degrading, or deceptive images, can cause severe harm to the reputation, safety, and mental health of the person being misrepresented.

Real cases show both public figures and private individuals facing harassment, career damage, and loss of trust after non-consensual deepfake content is circulated online. Just seeing such a realistic mischaracterization of oneself can be deeply traumatic. **Circulating these kinds of falsified images online or otherwise can lead to serious legal consequences**

Ultimately, both state and national laws have been created to mitigate this kind of abuse, and continued regulation of AI to protect people is paramount. We each have a duty to use AI and any tool at our disposal with the utmost respect for others while guided by ethics and common sense. Even when an app or website makes it simple to do this, resist the temptation and remember the damage you could cause.

Responsible image guidelines

Never create or share AI images that sexualize, humiliate, or misrepresent real people. Even if you have good intentions, never use another person’s image, name, or other personal work or ideas without their explicit consent.

Do not fabricate evidence (fake screenshots, staged "proof," edited photos) to win arguments or influence opinions.

When you share AI-generated visuals, be transparent if viewers might reasonably assume they are real. It’s best to be transparent about how you use AI when generating content.

Use a simple test: if you wouldn't want this image made of you or someone you care about, don't make it of anyone. Treat others the way you would like to be treated.

Key principle: People's dignity comes before your curiosity or comedic ideas. Deepfakes are not a harmless joke and can lead to serious legal repercussions.

6. Hidden Costs: Remembering the Physical World

Environmental and Human Costs

AI can feel weightless and instant, but each query uses real resources. Training and running large models consume substantial electricity and often large volumes of water for cooling data centers. Reports estimate that rapid AI expansion could use water and electricity equivalent to what millions of households consume in the United States alone each year.

Behind the scenes, human workers are also required to review disturbing or abusive content and label it, helping models learn what to block. This work has a real mental health cost. Being exposed to countless disturbing images on a regular basis has a lasting negative effect on the person doing that work.

Remember, these are part of the AI system's cost, even if you never see it.

Using AI thoughtfully

Use AI when it truly adds value—solving a real problem, helping you learn, or saving significant time.

For tiny tasks you could do yourself in a minute, consider skipping the AI model. This reduces environmental load and encourages you to keep basic skills sharp.

When possible, choose AI tools and settings that prioritize efficiency and sustainability.

Key principle: Every AI response has a footprint. Respect the resources and people that make AI possible.

7. Cognitive Deficit: Too much AI makes the brain lazy

The risk of outsourcing your thinking

AI can draft emails, brainstorm ideas, outline and write essays, generate images or videos, and even write code. That is powerful—but if you outsource every step from ideation to first draft, to refinement, and ultimately even the final solution, your own thinking muscles weaken. And to be frank, the quality of the final output will never be as good as if you had done the thinking yourself all along and only used AI to help you iterate and refine...

Emerging research suggests that heavy reliance on AI tools is linked to lower critical thinking skills. Over time, constantly skipping the "struggle" phase can make you more passive, less curious, and more likely to accept answers without questioning them.

Don't let AI rob you of the joy and excitement of exploring your own imagination, discovering through trial and error, and being a pioneer with your own thinking!

Healthy habits for AI use

Do the first pass yourself: brainstorm, outline, or attempt a solution before asking AI to help refine or expand it.

Use AI as a coach with tasks like: "*Here is my draft. Critique it and suggest improvements*," rather than, "*Write this better for me.*"

Regularly practice tasks without AI to keep your human edge. Write short pieces, solve problems, or study without the help of AI so you maintain confidence in your own abilities.

Notice how you feel after long AI sessions. If you feel foggy, dependent, or doubtful, that's a sign that you need to spend more time thinking on your own.

Key principle: Use AI as a tool that adds to a process you already own from start to finish. Use AI to multiply your mind, not replace it.

Your AI Ethics Quick Checklist

1. **Ownership:** Did I meaningfully transform this output so it reflects my own voice and judgment? If needed, have I been transparent enough about AI assistance?
2. **Accuracy:** Have I checked important facts, numbers, and claims with at least one independent source, especially for high-stakes topics?
3. **Bias:** Could this content be considered stereotypical? Does it exclude or harm any groups of people? Have I looked at it from their perspective and addressed any potential issues?

4. **Safety & Privacy:** Did I avoid sharing sensitive personal, financial, medical, or confidential information about myself or others?
5. **Images & Harm:** Would I be comfortable if a similar AI-generated image of me or someone I love existed? If not, then it should not be used.
6. **Hidden Costs:** Is this use of AI worth the environmental and human resources involved, or could I reasonably do this without the help of an AI model?
7. **Your Brain:** Did I think, draft, or plan on my own before asking AI for help? Am I using AI to support my growth, or am I using AI as a crutch or a shortcut?

Bottom line

Own your work. Question AI's answers. Guard your privacy. Respect the humans and resources behind the power of AI. Use AI to multiply your mind, not replace it.