NO PREP

Being a Responsible AI Citizen

A Personal Code of Use

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

Being a Responsible AI Citizen: A Personal Code of Use

Why Individual Choices Matter

When we talk about AI's impact on society, the conversation quickly scales up to the level of governments, corporations, and global policy. But the aggregate impact of AI is ultimately composed of millions of individual choices: how people use AI tools, what they share online, how they respond to AI-generated content, and what norms they enforce in their communities. Understanding the principles of responsible AI use is not just about personal virtue — it is about the contribution each person makes to the collective environment in which AI develops and operates. The concept of digital citizenship — the idea that using

Name: **Period:** **Date:**

Being a Responsible AI Citizen: A Personal Code of Use

Why Individual Choices Matter

When we talk about AI's impact on society, the conversation quickly scales up to the level of governments, corporations, and global policy. But the aggregate impact of AI is ultimately composed of millions of individual choices: how people use AI tools, what they share online, how they respond to AI-generated content, and what norms they enforce in their communities. Understanding the principles of responsible AI use is not just about personal virtue; it is about the contribution each person makes to the collective environment in which AI develops and operates. The concept of digital citizenship, the idea that using online spaces comes with responsibilities as well as rights, has been part of technology education for two decades. AI citizenship extends this concept into a new domain. It asks: given what you know about how AI works, what it can do, what it gets wrong, and how it is misused, what does responsible use look like? What do you owe to other people in how you engage with AI-generated content? What do you owe to yourself?

Honesty About AI Use

One of the most frequently contested questions in AI citizenship involves disclosure: when you use AI to generate or substantially assist with work that you present to others, do you have an obligation to say so? The debate around academic AI use (whether students should be permitted to use AI for essays, and whether they should disclose when they do) is the most visible version of this question, but it extends well beyond school. A professional who submits AI-generated work without disclosure may be misrepresenting their own capabilities or the labor involved. An author who publishes AI-generated text under their own name is making a claim about authorship that may mislead readers. A person who shares AI-generated images, audio, or video without labeling them as synthetic is potentially deceiving

their audience about the reality of what they're seeing. The principle of honest disclosure does not require labeling every AI-assisted spellcheck or grammar suggestion. It does require transparency about substantive AI contribution: when AI generated the core content, reasoning, or creative work that you are presenting as your own. The test is whether the recipient of your work would feel deceived if they knew the extent of AI's contribution. If yes, disclosure is warranted.

Verify Before You Share

AI-generated content (text, images, audio, video) can be convincing and wrong simultaneously. Large language models produce confident, fluent statements about topics they actually know nothing about. AI image generators can create photorealistic images of events that never occurred. AI voice and video can produce convincing impersonations of real people saying things they never said. The responsible AI citizen treats AI-generated information with the same skepticism they would apply to a confident stranger on the street: plausibly knowledgeable, but requiring verification before being acted on or shared. This means checking AI factual claims against primary sources when the stakes are significant, conducting reverse image searches on suspicious images, and being particularly cautious about emotionally charged political content where the incentives to deceive are highest. The 'share before you verify' reflex (the impulse to immediately forward something alarming, outrageous, or surprising) is precisely what AI misinformation is designed to exploit. Slowing down even briefly before sharing political content, surprising news, or anything that triggers a strong emotional response is one of the highest-impact individual behaviors for reducing the spread of AI misinformation.

Privacy and Data Hygiene

Every interaction with an AI system involves sharing data. What you ask about, what you share to get a response, what preferences you reveal through your queries — this information is collected, stored, and potentially used to train future AI models or shared with third parties, depending on the platform's policies. Responsible AI citizenship includes understanding what data you are sharing and with whom. Practical data hygiene with AI tools includes: not sharing sensitive personal information (medical details, financial information, passwords, confidential business information) with AI chatbots unless you have carefully reviewed the platform's data handling policies; being cautious about sharing information that could be used to

identify or harm third parties; and periodically reviewing and deleting your conversation history with AI services when available. This is particularly important for young people, whose comfort with sharing personal information online can create risks they aren't fully aware of. The voice-cloning scam risk — where audio shared on social media is used to defraud family members — is a direct example of how casual digital sharing can create security vulnerabilities with real consequences.

Supporting Human Oversight

Responsible AI citizenship also involves supporting — not undermining — human oversight of AI systems. This means supporting disclosure and labeling requirements for AI-generated content, being willing to verify and report AI-generated misinformation when you encounter it, and pushing back on the normalization of AI deception in your social environment. It also means being thoughtful about where AI-assisted decision-making is and isn't appropriate. Using AI to help draft a message is different from using AI to make a consequential decision about another person without their knowledge. Using AI to explore ideas is different from using AI to produce work you claim as your own in a context where that claim matters. The question 'Would I be comfortable if the people involved knew exactly how I used AI here?' is a useful guide to the ethics of individual AI use cases. Building a personal code of AI use is not about being anti-technology or refusing to use powerful tools that are legitimately helpful. It is about using those tools with the same ethical awareness you would bring to any powerful capability. Power and responsibility scale together — the more capable AI becomes, the more it matters that the people using it are thinking carefully about how. The environmental impact of AI use is an aspect of responsible AI citizenship that is rarely discussed but increasingly significant. Training large AI models requires substantial energy — estimates for training a single frontier model run into the equivalent of thousands of transatlantic flights. Inference (running queries against trained models) is less energy-intensive per query but occurs billions of times daily across all AI users. Responsible AI citizenship includes being thoughtful about when AI is the appropriate tool for a task versus when simpler, less energy-intensive approaches would work just as well. This is not an argument against using AI but for using it intentionally rather than reflexively. The role of AI in academic settings has generated particularly intense debate about the norms of responsible use. Schools and universities have responded with wildly

varying policies: some ban AI use entirely, others permit it freely, many have attempted to develop nuanced policies that permit some uses while prohibiting others. The practical enforcement challenges are significant — detecting AI-generated text remains unreliable — which means that whatever norms develop will depend heavily on genuine internalization rather than external enforcement. This makes the formation of personal ethical codes, rather than rule-following, particularly important for young people navigating AI in educational contexts.

Key Information: Responsible AI Citizenship

Detail	Information
AI citizenship definition	Applying digital citizenship principles to AI: using AI tools with responsibility to others
Disclosure standard	Transparency about substantive AI contribution — when AI generated the core content you're presenting
Test for disclosure	Would the recipient feel deceived if they knew the extent of AI's contribution?
Verify before share principle	Check AI factual claims against primary sources; reverse search suspicious images
Most dangerous reflex to fight	The impulse to immediately forward alarming or outrageous content before verification
Data hygiene rule	Avoid sharing sensitive personal info with AI chatbots; review platform data policies
Voice-cloning connection	Social media audio is a security vulnerability — can be used to clone voice for scams
Personal ethics test	Would I be comfortable if everyone involved knew exactly how I used AI here?

Case Study: Jaylen's History Paper

Jaylen is a junior working on a history research paper worth 30% of his grade. His teacher has said AI use is permitted with disclosure. Jaylen is genuinely interested in his topic — the civil rights movement in his city — and has done substantial original research, including an interview with a local activist and visits to two archives. He uses AI to help him organize his notes and generate an outline. He then writes the entire paper himself. He also asks an AI to proofread his final draft and suggest improvements to his sentence structure, which he reviews and accepts some of. When Jaylen submits the paper, he includes a footnote: 'AI was used to assist with organizational outlining and final proofreading. All research, analysis, and writing is my own.' His teacher comments that his disclosure was thoughtful and accurate. Another student, Marcus, submits a paper that is largely AI-generated — he asked AI to write it based on notes he provided, made a few edits, and submitted it with no disclosure. Marcus got a similar grade but his teacher later requests a conference.

1. Using the reading's 'disclosure test,' explain why Jaylen's disclosure was appropriate while Marcus's non-disclosure was an ethical violation.

2. What specific principle from the reading does Jaylen's behavior demonstrate, and how does it connect to the broader concept of AI citizenship?

Multiple Choice

1. According to the reading, what is the test for whether AI disclosure is warranted?

a) Whether the work will be graded or evaluated by another person

b) Whether the recipient would feel deceived if they knew the extent of AI's contribution

c) Whether the AI tool used requires disclosure in its terms of service

d) Whether more than 50% of the content was generated by AI rather than by the human author

2. What does the reading identify as the 'most dangerous reflex' to fight in preventing AI misinformation spread?

a) The impulse to immediately share alarming or outrageous content before verifying it

b) The automatic reflex of dismissing AI-generated content as unreliable without checking

c) The tendency to trust AI-generated content more than human-created content by default

d) The habit of using AI tools for tasks that should be completed by humans themselves

3. How does the concept of 'AI citizenship' extend the older concept of 'digital citizenship'?

a) AI citizenship replaces digital citizenship now that AI has become the dominant technology in modern educational institutions worldwide

b) AI citizenship focuses on coding skills and advanced technical programming work while digital citizenship focused primarily on online behavior

c) AI citizenship applies ethical responsibilities to AI tools specifically, including disclosure, verification, and supporting human oversight

d) AI citizenship is primarily about protecting children from AI, while digital citizenship was for all ages and all internet users

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ AI citizenship	A) The responsibility to inform recipients when AI generated the core of work you're presenting as your own
2. ___ Data hygiene	B) AI-generated core content, reasoning, or creative work — as opposed to minor assistance like spelling correction
3. ___ Substantive AI contribution	C) Checking AI factual claims against primary sources and reversing searching suspicious images before forwarding
4. ___ Disclosure obligation	D) Practices for protecting personal information when using AI tools, including reviewing platform data policies
5. ___ Verify before you share	E) Using AI tools with responsibility to others — applying digital citizenship principles in an AI context

Fill-in-the-Blank

1. Responsible AI citizenship includes supporting _____ and labeling requirements for AI-generated content in your community.

2. The concept of _____ citizenship — using online spaces responsibly — has been part of technology education for two decades.

3. The personal ethics test for AI use asks: 'Would I be _____ if everyone involved knew exactly how I used AI here?'

Short Answer

Answer each question in 2-4 complete sentences.

1. Explain why the 'verify before you share' principle has a higher impact on AI misinformation than most people might expect.

2. What does the reading mean when it says 'power and responsibility scale together' in the context of AI?

3. Write your own three-point personal code of AI use, drawing on at least two principles from the reading. Explain your reasoning for each point.