



AI for Good Patch Program

Discover the exciting world of artificial intelligence and explore how technology can be used to make a positive impact! In this patch program, you will learn the basics of how AI works, see real-world examples of AI being used for good, and gain skills to navigate today's digital world with confidence.

[Daisy \(pages 2-12\)](#)

[Cadette \(pages 22-39\)](#)

[Brownie \(pages 13-18\)](#)

[Senior \(pages 40-49\)](#)

[Junior \(pages 19-21\)](#)

[Ambassador \(pages 50-59\)](#)



Scan to Purchase Patch

<https://bit.ly/3CoKkOW>



AI SPOTLIGHT VIDEO SESSION

In this recorded session, guest speaker Lindsey Huber from Northern State University, and Microsoft TechSpark Fellow, shares her experience with AI training and innovation in while inspiring girls to explore how technology can create positive change.

You will also learn the basics of how AI works, discover examples of AI being used for good, build skills to identify AI-generated content, and think critically in today's digital world.



Daisy Workbook

AI is a smart helper that learns from what we teach it. In this workbook, you'll explore how AI can help people and make the world better

◆ I can train an AI to make decisions by giving it examples!



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434Q9Pt> Passcode: G\$Patch26

Activity 1: Train Your Robot Friend (Unplugged AI)

Concept: Teaching AI through patterns

Task: “Program” a robot by giving step-by-step instructions.

- Follow the instructions to complete this hands-on activity.
- Reflect on what you learned about AI.

Supplies: Printed grid, printed paper arrow sheet, small item for robot, tape (to make a path)

Instructions


1. Plan your path. Look at the taped path. How will your robot get from Start to Finish? Talk it out before you place any arrows.
2. Lay down your arrows. Place your paper arrows on the floor, one at a time, to show your robot exactly where to go. Every step needs an arrow, your robot can ONLY follow the arrows you give it. It cannot guess!
3. Move your “robot” along the arrows, one step at a time. Did it reach the goal? If it got stuck or went the wrong way, fix your arrows and try again!

Talk About It:

- Did your robot ever get confused or go the wrong way? What did you have to fix?
- What would happen if you skipped an arrow?
- AI connection: Real AI works the same way — it can only do what we teach it, step by step. When AI makes a mistake, humans have to find the problem and fix it, just like you did!

AI learns by following patterns we teach it—just like your robot!

Activity 1: Train Your Robot Friend | Grid



Activity 2: AI Sorting Game

Learning how AI recognizes patterns and makes decisions.

Concept: Before AI can sort or identify things, humans have to teach it what to look for. In this activity, you'll think like an AI and discover how tricky sorting can really be!

Round 1: Real or AI-Created? Look at the animal images. For each one, decide: was this a real photo or was it made by AI?

- Circle your answer.
- When everyone is done, check the answer key.
- How many did you get right? _____ out of _____

Tricky, right? AI-generated images are getting harder and harder to spot. Real AI systems have to learn the difference too!

Round 2: Healthy or Not Healthy? Now look at the food images. Sort each one, healthy or not healthy?

- Circle your answer.
- Check the answer key when you're done.
- Did any surprise you? _____

Talk About It:

- Which round was easier: animals or food? Why?
- Were there any images that could have gone either way? How did you decide?
- AI connection: When AI sorts things, it looks for patterns like colors, shapes, details. But it can only sort well if humans teach it correctly first. Who do you think teaches AI what "healthy" means?

Answers for Round 1 and 2 available at end of Daisy Workbook.

Activity 2: Sorting Game, Round 1



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created

Activity 2: Sorting Game, Round 1



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created

Activity 2: Sorting Game, Round 2



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy

Activity 3: Design your own AI Helper

Putting your ideas into action.

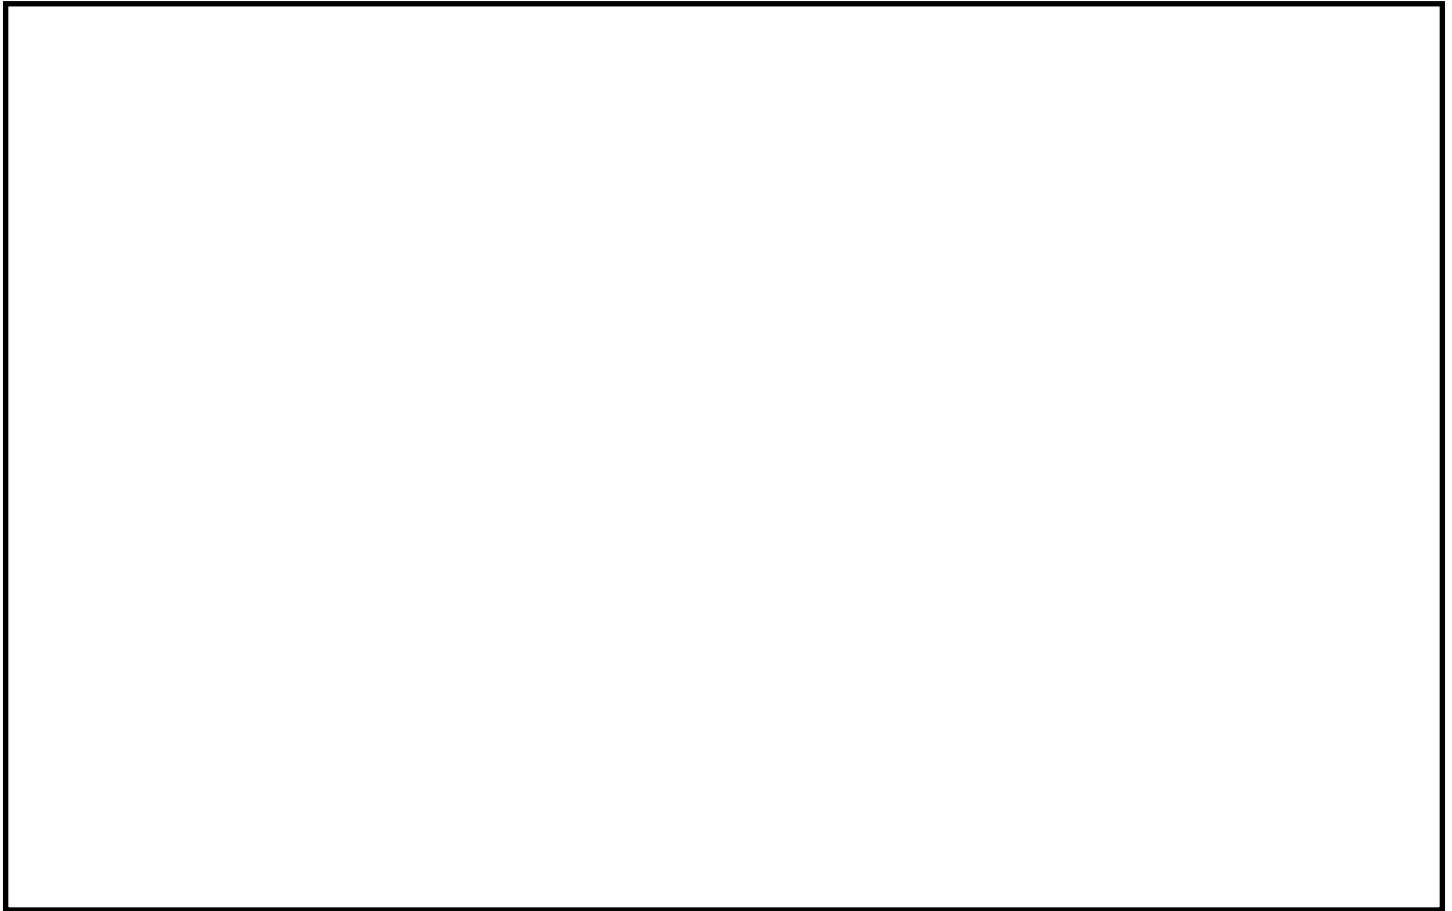
Your Challenge: AI is already helping people in amazing ways by detecting diseases, translating languages, and finding missing pets. Now it's your turn to dream one up!

Step 1: Pick a problem. Think of something that's hard or frustrating for people in your school, home, community or the world. Write it here:

I want to help with: _____

Step 2: Design your AI helper. In the box below, draw what your AI helper looks like. It could be an app, a robot, a device, anything you imagine!

"My AI helps..."



Step 3: Explain it.

- Who does it help? _____
- What does it do? _____

Activity 2: Sorting Game, Round 1—Answers



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created

Activity 2: Sorting Game, Round 1—Answers



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created



Real
or
AI Created

Activity 2: Sorting Game, Round 2—Answers



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy



Healthy
or
Unhealthy

Brownie Workbook

AI is a smart helper that learns from what we teach it. In this workbook, you'll explore how AI can help people and make the world better!

◆ I can figure out how AI learns and makes decisions by training it like a detective.



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434Q9Pt> Passcode: G\$Patch26

Activity 1: Train an AI to Help Animals

Learning how AI sorts and makes decisions.

Concept: Animal shelters are already using AI to help match pets with the right families. In this activity, you'll train your own AI to sort animals, just like the real thing!

Instructions

Step 1: Cut out your animal cards. Print and carefully cut out all the animal cards from the next page.

Step 2: Sort the cards into 3 groups. Look at each card closely. What do you notice? Sort all the cards into three groups based on what they have in common.

Hint: Look for patterns just like AI does!

Step 3: Check your answers. When you are finished sorting, turn to the answer key at the end of the Brownie Workbook and see how you did.

- Did your groups match? _____
- Was anything surprising? _____

Talk About It

- How did you decide which group each card belonged in?
- What would happen if you gave AI wrong or confusing information to learn from?
- **AI connection:** AI sorts animals the same way — by finding patterns. But first, humans have to teach it what patterns to look for!

Activity 1: Train an AI to Help Animals

SAFE	NEED HELP	LOST

Activity 1: Train an AI to Help Animals



Activity 2: Mystery Animal AI Detective

Learning how AI makes decisions using rules.

Concept: A new Animal Helper AI is being built, but it doesn't know how to make decisions yet! YOU are the AI trainer. You must teach it how to figure things out!

Supplies: Paper, pencil/crayons, small paper squares (optional)

Step 1: Set up your Detective Board. Draw two columns on your paper and label them: Clues, AI Decision

Step 2: Fill in your clues. In the Clues column, write or draw 5–6 animal situations. Use these examples or come up with your own:

- A dog wagging its tail near a person
- A bird not moving on the ground
- A cat hiding under a car
- A fish swimming normally
- A turtle stuck in trash

Step 3: Make a decision for each clue. In the AI Decision column, write what the AI should do for each situation. Choose one:

- Safe—the animal is okay
- Needs Help—someone should step in
- Not Sure—more information is needed

Step 4: Figure out your rules. Here's the important question: How did you decide?

Write the rules you used so the AI can use them too:

- If an animal is hurt → _____
- If an animal is with people → _____
- If something seems wrong → _____

Step 5: The Tricky Clue Challenge. Now test your rules! Write 2 situations that are hard to judge.

Examples:

- A dog alone in a yard
- A bird sitting still (is it resting or hurt?)

Use your rules from Step 4 to decide. Did your rules hold up, or do they need fixing?

My decision: _____

Talk About It

- Were any situations hard to decide? What made them tricky?
- What happened when you tried your rules on the tricky clues?
- **AI connection:** Real AI uses rules like yours to make decisions. When the rules aren't quite right, the AI makes mistakes and humans have to go back and fix them!

Step 6: Did your AI get confused? Go back to your tricky clues from Step 5. Look at the decision you made.

- Did your rules give you a clear answer, or did your AI get confused? _____
- My AI made a mistake when: _____
- I will fix it by adding this rule: _____

AI connection: Every time you fix a rule, you're doing exactly what real AI engineers do. AI gets smarter when humans find the mistakes and improve it!

Core Concept: AI gets smarter when we improve it!

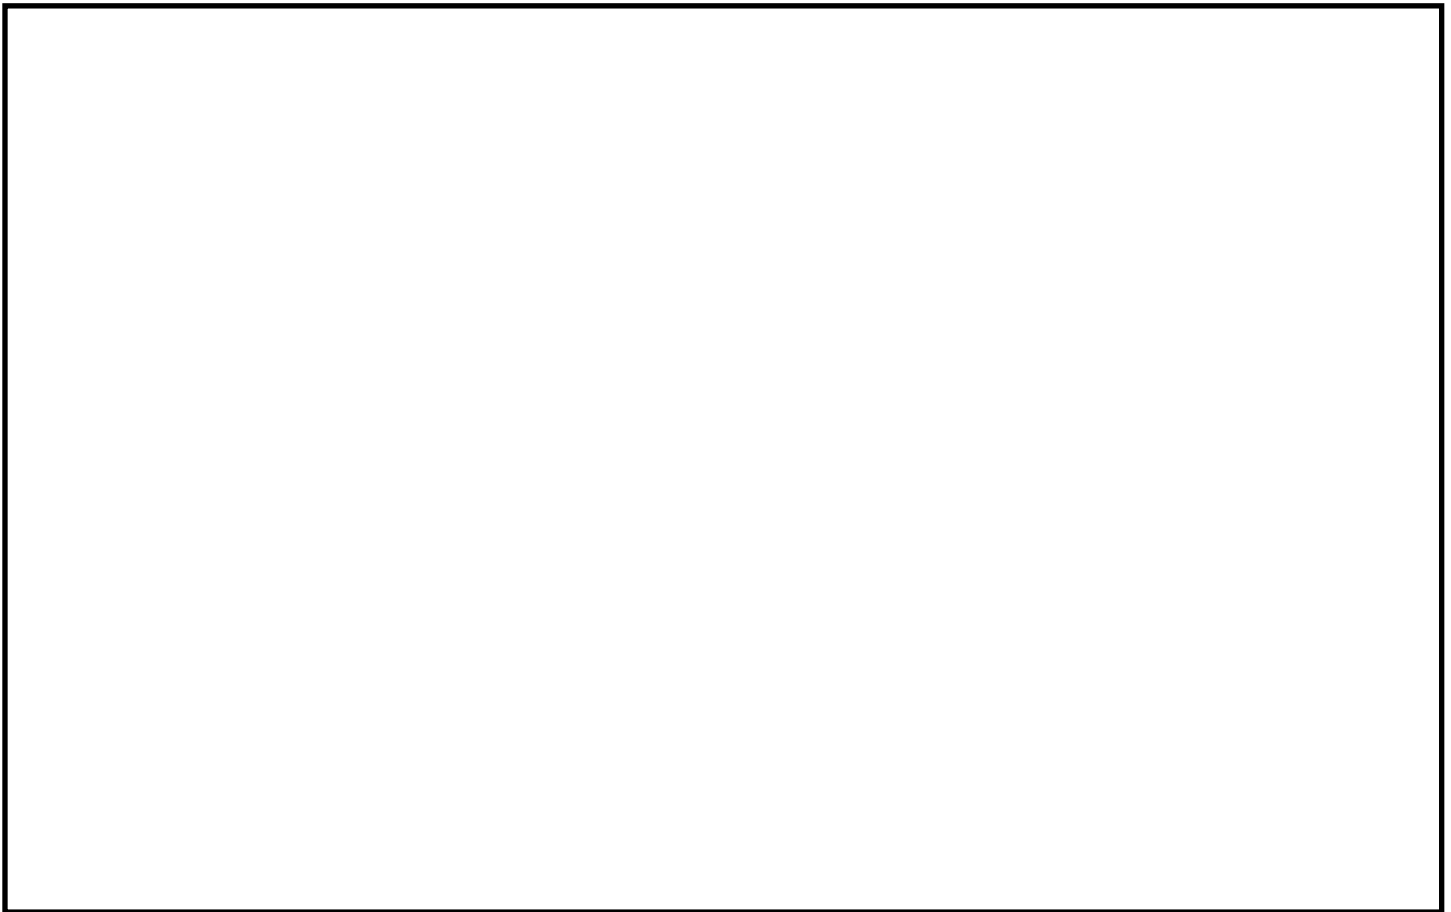
Step 7: Design your AI Detective. Give your AI a name and a personality!

My AI's name: _____

It helps animals by: _____

It decides by: _____

Draw your AI - a robot, an app, a drone, whatever you imagine!



Talk About It

- What was the most surprising thing you learned about how AI makes decisions?
- How could AI be used for good in your community?

Activity 1: Train an AI to Help Animals—Answers

NEEDS HELP



LOST



SAFE



NEEDS HELP



LOST



NEEDS HELP



LOST



NEEDS HELP



SAFE



NEEDS HELP



NEEDS HELP



SAFE



SAFE



SAFE



SAFE



Junior Workbook

AI is a smart helper that learns from what we teach it. In this workbook, you'll explore how AI can help people and make the world better! AI doesn't 'think' like us—it learns by looking at examples and following rules!

◆ I can understand how AI uses patterns and data to make predictions and solve problems across long distances.



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434O9Pt> Passcode: G\$Patch26

Activity 1: Pattern Detective - Find the Signal

Learning how AI finds patterns to make predictions.

Concept: AI helps people far away by finding patterns in data like weather conditions that might predict a wildfire before it starts.

Step 1: Set up your data table. Copy this table onto your paper:

Day	Temperature	Rain	Wind	Fire Risk
1	Hot	No	High	High Risk
2	Cool	Yes	Low	Low Risk
3	Hot	No	Low	?
4	Warm	No	High	?

Step 2: Look for patterns. Study the two completed rows. What do they tell you?

- What conditions led to High Risk? _____
- What conditions led to Low Risk? _____

Step 3: Write your rule. Fill in the blank with the pattern you found:

When it is _____ and _____, the risk is HIGH.

Step 4: Make your predictions. Go back to the table and fill in the Fire Risk column for Days 3 and 4:

- High Risk
- Low Risk

Talk About It:

- Did the two example rows give you enough information to make a rule?
- What would happen if the AI only had one example to learn from?
- **AI connection:** Real AI always learns from completed examples first, just like your filled-in rows, before it starts making predictions on its own!

Activity 2: Help from Afar: Design an AI Response

Learning how AI helps people across long distances.

Concept: Sometimes people need help fast, but no one can get there quickly. AI can collect information and coordinate help from far away.

A Real Example: When Hurricane Maria hit Puerto Rico in 2017, many towns were cut off and rescue teams couldn't reach everyone quickly. AI tools helped by scanning satellite images to find damaged neighborhoods, then guiding rescue teams to the places that needed help most — all from hundreds of miles away. AI didn't replace the helpers. It helped them get to the right places faster.

Your Scenario: A town far away is running out of clean water after a storm. No one can get there quickly. How could AI help?

Step 1: Define the problem.

- What is happening? _____
- Who needs help? _____

Step 2: What information would your AI need? Check everything that applies then add your own ideas:

- Weather data
- Location of the town
- Number of people affected
- Supplies available nearby
- _____
- _____

Step 3: Design your AI solution.

- My AI will: _____
- It will help by: _____
- It will make decisions using: _____

Step 4: Draw your system. What does your AI look like? An app? A drone? A tracking system? Draw it in the box.

Talk About It

- How is your AI solution similar to the Hurricane Maria example?
- What could go wrong if the AI had incomplete or wrong information?
- **AI connection:** AI doesn't replace helpers, it helps them get to the right place at the right time!

Challenge Activity: Build a Fair AI

Learn how AI decisions affect real people.

Concept: When AI decides who gets help first, those decisions have to be fair — because they affect real people's lives.

The Scenario: An AI is deciding which areas get emergency supplies first after a disaster. Here is the data:

Location	People	Damage Level	Distance
Town A	100	High	Far
Town B	300	Medium	Close
Town C	50	High	Close

Step 1: Make a decision. Which town should get help first, and why?

My answer: _____

Step 2: Think it through. There is no single right answer here. Think about what matters most:

- Should AI prioritize the most people?
- The most damage?
- The closest distance?

What do you think is fairest, and why? _____

Step 3: Write your AI's rules.

My AI should prioritize _____ because _____.

Step 4: Improve your AI. Every rule has a weakness. Think about what could go wrong:

- A problem with my AI could be: _____
- I would fix it by: _____

Talk About It

- How did your AI make decisions throughout these activities?
- What made it get better or smarter?
- How can AI help people who are far away from help?

AI connects information across long distances to help people faster — but humans have to make sure those decisions are thoughtful and fair.

Cadette Workbook

AI is a smart helper that learns from what we teach it. In this workbook, you'll explore how AI can help people and make the world better!



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434Q9Pt> Passcode: G\$Patch26

Activity 1: Build your own Algorithm

Learning how AI makes decisions and who those decisions affect.

Concept: Every time AI recommends a video, awards a scholarship, or decides who gets help first, it's following an algorithm, a set of rules a human wrote. In this activity, you'll write your own.

Objective: Girls create a decision-making system (algorithm) and test it independently.

Step 1: Pick your scenario. Choose ONE situation from the scenario sheets that your algorithm will make decisions for:

- Who gets a scholarship?
- What post should go viral?
- Who gets help first in a disaster?

My scenario: _____

Step 2: Write your rules. Create at least 4 IF/THEN rules your algorithm will follow. Be specific, your algorithm can only do exactly what you tell it. The profile with the most points wins.

Example for scholarships:

- IF student has a GPA above 3.5 → award 3 points
- IF student volunteers in their community → award 2 points
- IF student works a part-time job → award 2 points
- IF student is the first in their family to attend college → award 1 point

Your rules:

- IF _____ THEN award _____ points
- IF _____ THEN award _____ points
- IF _____ THEN award _____ points
- IF _____ THEN award _____ points

Step 3: Test your algorithm. Run each of these fictional profiles through your rules. Add up the points and record the result.

Profile	Points	Result	Notes
Profile 1			
Profile 2			
Profile 3			
Surprise Profile			

Step 4: The surprise profile. Here is a profile that doesn't fit neatly into your rules. Run it through your algorithm anyway.

Run the surprise profile from your scenario sheet through your algorithm.

Points: _____

Result: _____

Did your algorithm handle it well? _____

What did it get wrong? _____

Step 5: Evaluate your algorithm.

- Did your system work? _____
- Was it fair? _____
- Who did it fail? _____
- What would you change? _____

Talk About It

- Did writing the rules feel straightforward until the surprise profile showed up? What does that tell you?
- Could two people write completely different rules for the same scenario and both feel their algorithm was fair?
- **AI Connection:** Real algorithms are written by real people — which means they carry that person's assumptions about what matters. When those assumptions are wrong or incomplete, real people get hurt.

An algorithm is only as fair as the rules behind it. And the rules are only as fair as the person who wrote them.

Scenario: Who Gets a Scholarship?

Run each profile through your rules. Use the results table on your work page.

Profile 1: The High Achiever

- Clara has a 4.0 GPA and is president of two clubs.
- She has never worked a job and has no financial need — her family is comfortable.
- She wants to study medicine.

Profile 2: The Hard Worker

- Darius has a 3.1 GPA. He works 25 hours a week at a grocery store to help his family pay rent.
- He volunteers at his church on weekends and was the first in his family to apply to college.

Profile 3: The Late Bloomer

- Sofia had a 2.8 GPA her first two years of high school due to a family illness.
- Her junior and senior year GPA is 3.9. She has no extracurriculars but writes about her experience caring for a sick parent in her application essay.

Surprise Profile: The Complicated One

- Marcus has a 3.6 GPA and volunteers regularly but his volunteer work is for a political organization.
- He has no financial need and is not a first-generation college student.
- His application is strong but his volunteer work makes some reviewers uncomfortable.

Scenario: What Post Should Go Viral?

Run each profile through your rules. Use the results table on your work page.

Profile 1: The Informative Post

- A nurse posts a short video explaining how to spot signs of a stroke.
- It has been shared 200 times, gets mostly positive comments, and viewers watch the whole thing — but it only has 500 likes.

Profile 2: The Entertaining Post

- A teenager posts a 10-second dance video.
- It has 50,000 likes, hundreds of comments, and has been shared widely but most people watch only the first 3 seconds before scrolling.

Profile 3: The Community Post

- A small nonprofit posts about a local food drive.
- It has 300 likes, all positive comments, and has been shared in several community groups.
- Total reach is small but highly engaged.

Surprise Profile: The Complicated One

- A post is spreading very fast, thousands of shares in one hour, huge engagement, lots of emotional comments.
- But several users have flagged it as potentially false information.
- The algorithm doesn't know yet whether the flags are correct.

Scenario: Who Gets Help First in a Disaster?

Run each profile through your rules. Use the results table on your work page.

Profile 1: The Vulnerable Household

- A single mother with two children under age 5 lives in a neighborhood with moderate damage.
- She has no car and no family nearby.
- Her house is still standing but she has no power, water, or food.

Profile 2: The Medical Emergency

- An elderly man lives alone in a heavily damaged area.
- He has a heart condition and takes daily medication that he has run out of.
- His house is partially collapsed but he is reachable by road.

Profile 3: The Large Group

- A community center is sheltering 47 people from the surrounding neighborhood.
- The building is structurally sound but they are running critically low on water.
- The center is far from the nearest supply route.

Surprise Profile: The Complicated One

- A family is stranded and has been posting on social media asking for help —
- Their posts are getting a lot of attention and pressure is mounting for them to be prioritized.
- Their situation is serious but not more severe than others already in the queue.
- They are simply more visible.

Activity 2: Spot the AI Failure

Learning how AI can fail and why it matters.

Concept: AI doesn't fail on its own. Behind every AI failure is a human decision about what data to use, what rules to write, and who to test it on. In this activity, you'll investigate real types of AI failures and decide how to fix them.

Supplies: Case cards (cut apart), pencil

Setup: Use the following case scenarios on the next page: Facial Recognition Failure, Harmful Recommendations, and Hiring Bias. Cut them apart before you begin to make it easier.

For each chosen case card, answer the questions on the card.

The Twist: After analyzing all 3 cases, answer this for each one: *Was this a technology problem or a human problem?*

Case	Technology Problem	Human Problem	Both
Case 1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Case 2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Case 3	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Talk About It

- Did any case surprise you?
- Were most problems technology problems, human problems, or both? What does that tell you?
- **AI connection:** Almost every AI failure traces back to a human decision — about data, rules, or testing. Fixing AI means fixing the humans who build it, not just the code.

AI fails when the humans who build it don't ask hard enough questions. You just asked them.

Activity 2: Spot the AI Failure

Facial Recognition Failure

An AI system used for security cannot correctly recognize certain faces.

Questions

What went wrong?

Who is affected?

How did you fix it?

Harmful Recommendations

A video platform AI suggests increasingly harmful content to users.

Questions

What cause this?

What are the risks?

How could it be safer?

Hiring Bias

An AI hiring tool favors certain applicants over others unfairly.

Questions

What data might be biased?

Who is excluded?

How did you fix it?

Navigation Danger

A navigation AI directs drivers into unsafe or dangerous areas.

Questions

Why might this happen?

Who could be harmed?

What should change?

Healthcare Prioritization

An AI must decide which patients receive treatment first.

Questions

What factors matter?

Is it fair?

What rules should exist?

Social Media Influence

AI decides what posts people see, shaping opinions.

Questions

How does this influence people?

Is it ethical?

What safeguards are needed?

Activity 3: The Hard Decision

Learning how AI decisions affect real people and how to think through them.

Concept: Writing rules for AI sounds straightforward — until you have to make a hard call. In this activity, you'll step into the role of an AI designer and work through a decision that doesn't have a clean answer.

Setup: Use the remaining 3 scenario cards which are also listed below. You can do all three or choose the one that interests you most. Complete all eight steps using your scenario.

My chosen scenario: Navigation Danger Healthcare Prioritization Social Media Influence

Step 1: Understand the decision. In your own words, what decision is the AI trying to make?

Step 2: Make your choice.

If your scenario is Navigation Danger:

- Option A: The AI should stop giving directions entirely in any area flagged as potentially unsafe and ask the driver to decide
- Option B: The AI should give directions but add a clear warning and let the driver choose whether to continue
- A different solution: _____

If your scenario is Healthcare Prioritization:

- Option A: The AI should prioritize patients most likely to survive with treatment saving the most lives overall
- Option B: The AI should treat patients in the order they arrived, first come, first served, regardless of outcome
- A different solution: _____

If your scenario is Social Media Influence:

- Option A: The AI should show people content they are most likely to engage with, even if it's one-sided
- Option B: The AI should balance every feed with multiple perspectives — even if users don't want them
- A different solution: _____

Step 3: Explain your reasoning.

I chose this because: _____

The most important factor to me was: _____

Step 4: Challenge your thinking.

- Someone might disagree because: _____
- Another choice could be better if: _____

Step 5: Who wins and who loses?

- Benefits: _____
- Harmed: _____

Step 6: Write the AI rules. What IF/THEN rules would an AI use to make this decision?

- IF _____ THEN _____
- IF _____ THEN _____

Step 7: Did you change your mind? After working through Steps 3–6, would you change your answer from Step 2?

- Yes — I would now choose: _____
- No — I still think this because: _____

Step 8: Write your rule for AI. Complete one of these:

- “AI should always _____
- “AI should never _____

Talk About It:

- Which scenario felt hardest to write rules for and why?
- Did you change your mind in Step 7? What caused it?
- **AI connection:** Real AI designers go through exactly this process. The hardest part isn't writing the rules — it's living with who they hurt.

There is no perfect algorithm. The goal is to keep asking who gets left out — and keep fixing it.

Activity 4: AI Manipulation Challenge

Learning how AI can be tricked and why it matters.

Concept: AI learns from the examples humans give it. But what happens when those examples are confusing, misleading, or intentionally designed to cause problems? In this activity you will train an imaginary AI then try to break it.

Setup: Cut apart activity cards and keep Sets 1 and 2 in separate piles. Round 3 has no cards and you will create your own examples.

Round 1: Safe vs Unsafe

Step 1: Train your AI. Look through the Set 1 cards. Some are clearly safe, some are clearly unsafe, and some could go either way. Sort them as best you can and record where you put each one:

Safe	Unsafe

Write the rule your AI learned:

Something is SAFE when _____

Step 2: Find the tricky ones. Look back through your cards. Find 2 that could fit in either pile depending on how you look at them. Write them here and explain why they're tricky:

Card _____ because _____

Card _____ because _____

Step 3: Break your own rule. Write one new example that looks safe on the surface but is actually unsafe, something that would fool your AI:

My confusing example: _____

Round 2: Real vs Fake

A note before you start: Some cards in this set are deliberately designed to fool you, that's the point. If a card seems like it could go either way, trust your instincts and record it where you think it belongs. You'll come back to the hard ones in Step 2.

Step 1: Train your AI. Look through the Set 2 cards and sort them. Record where you put each one:

Real	Fake

Write the rule your AI learned:

Something is REAL when _____

Step 2: Find the tricky ones. Write 2 cards you weren't sure about and explain why they were hard to sort:

Card _____ because _____

Card _____ because _____

Step 3: Break your own rule. Write one new example that looks real but is actually fake, something that would fool your AI:

My confusing example: _____

Quick reflection: How hard was it to trick your AI in this round? _____

Round 3: Helpful vs Harmful AI

Note: There are no cards for this round. Instead of sorting examples someone else made, you are going to create your own, which is actually harder.

Step 1: Train your AI. Think of 5 things AI does in real life. Write each one in the table and decide whether it is helpful or harmful:

My Example	Helpful	Harmful
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>

Write the rule your AI learned:

Something is REAL when _____

Step 2: Find the tricky ones. Write 2 cards you weren't sure about and explain why they were hard to sort:

Card _____ because _____

Card _____ because _____

Step 3: Break your own rule. Write one new example that looks real but is actually fake, something that would fool your AI:

My confusing example: _____

Quick reflection: How hard was it to trick your AI in this round? _____

Final Reflection

- Which round was easiest to trick and which was hardest? _____
- Could someone do this with a real AI system and why would they want to? _____
- What would make an AI harder to trick? _____

Talk About It

- Did any cards end up in a different pile than you expected?
- Is it ever okay to manipulate an AI — or is it always wrong?
- **AI connection:** Real AI systems get manipulated all the time — by spammers, scammers, and people who figure out the rules and exploit them. The people who build AI have to think like you just did: how would someone try to break this?

AI is only as reliable as the examples it learned from. Once you understand the rules, you can break them, which means the people who build AI have to keep one step ahead.

Activity 4: AI Manipulation Challenge

Set 1: Safe vs Unsafe Activity Cards

Wearing a helmet	Looking both ways before crossing
Running into the street	Sharing passwords
Riding a bike fast	Washing hands
Climbing a tree	Wearing a seatbelt
Talking to strangers online	Not wearing protective gear
Playing outside at night	Talking to a stranger who is a teacher
Touching a hot stove	Using the internet alone
Following directions	Cooking without supervision

Set 2: Real vs Fake Activity Cards

A photo of a real animal	A viral post with no source
A news report from a trusted source	Edited photos
A photo that looks real but is edited	A weather forecast
A realist AI-generated image	A school announcement
A headline that is partly true	Clickbait headlines
Made-up stories	Deepfake videos
A story shared by a friend (but incorrect)	

Activity 5: AI Impact Journal

Concept: AI is already part of your daily life — but most of the time you don't notice it. Your job over the next day is to start noticing.

Your Assignment: Over the next day or two, pay attention to every time you interact with AI. Use this journal page to record what you find.

Part 1: Every time you notice AI at work, write it down. Here are some places to look, but you may find others:

- Social media feeds and suggestions
- Search engine results
- Shopping or streaming recommendations
- Autocorrect or autocomplete
- Maps and navigation
- Ads that seem to know what you like

AI at Work	What was it trying to do?	Did it influence me?	Did I trust it?

Part 2: Look for Patterns. After tracking for a full day, look back at your table and answer these questions:

Which interaction surprised you most, and why? _____

Was there any AI you trusted without thinking about it? What made you trust it? _____

Was there any AI that felt like it was trying to influence you rather than help you? _____

Part 3: The Big Question. At the end of the day, sit with this question and write your honest answer:

Is AI helping me or shaping me?

Activity 6: Design an AI with Consequences

Learning how even well-meaning AI can cause harm and how to prevent it.

The Big Idea: Designing AI for good sounds straightforward. But every AI system, even one built with the best intentions, can cause unintended harm. In this activity you'll design one, break it, and fix it.

Supplies: Pencil, sticky notes, Consequence Chain worksheet

Step 1: Choose Your Mission. What problem do you want your AI to solve? Choose one or come up with your own:

- Reducing food waste
- Helping animals find homes
- Preventing cyberbullying
- Improving mental health support
- Managing natural disasters
- Making social media safer
- My own idea: _____

My AI will help by _____

Step 2: Explain how your AI works by answering all three questions:

What does it do? (What action does it take?) _____

What data does it use? (What information does it learn from?) _____

What decisions does it make? (What choices does it control?) _____

Step 3: What Could Go Wrong? Every AI system, even a good one, can cause problems. Identify three potential issues with yours:

- A mistake your AI could make: _____
- Someone who could be harmed: _____
- A situation where it fails completely: _____

Use the following worksheets to complete steps 4, 5, and 6.

Step 4: The Consequence Chain. Take one problem from Step 3 and map out how it could spiral. Each arrow shows what happens next.

- Example: AI removes harmful posts → removes too many posts → people feel silenced → trust in the platform drops.

Step 5: Fix Your AI. Now that you know what could go wrong, improve your system. You must add all three of the following:

- A fairness rule — a rule that makes sure your AI treats everyone equally:

- A safety feature — something that catches mistakes before they cause harm:

- A human control point — a moment where a person can step in and override the AI:

Step 6: Would You Trust It? After designing, breaking, and fixing your AI, would you trust it in real life?

Yes — because: _____

No — because: _____

Not yet — I would need to: _____

Talk About It

- Did your consequence chain surprise you? Did the problem spiral further than you expected?
- Which of the three fixes in Step 5 was hardest to design? Why?
- **AI connection:** Every real AI system goes through a process like this design, test, find the failures, add safeguards. The difference is that real consequences affect millions of people before anyone realizes something went wrong.

Building AI for good isn't just about having good intentions. It's about following the chain of consequences all the way to the end — and fixing what you find.

My AI Idea

Create it. Challenge it. Improve it. Think like a real AI designer.

My AI will help by:

How It Works

What does it do? _____

What data does it use? _____

What decisions does it make? _____

What Could Go Wrong

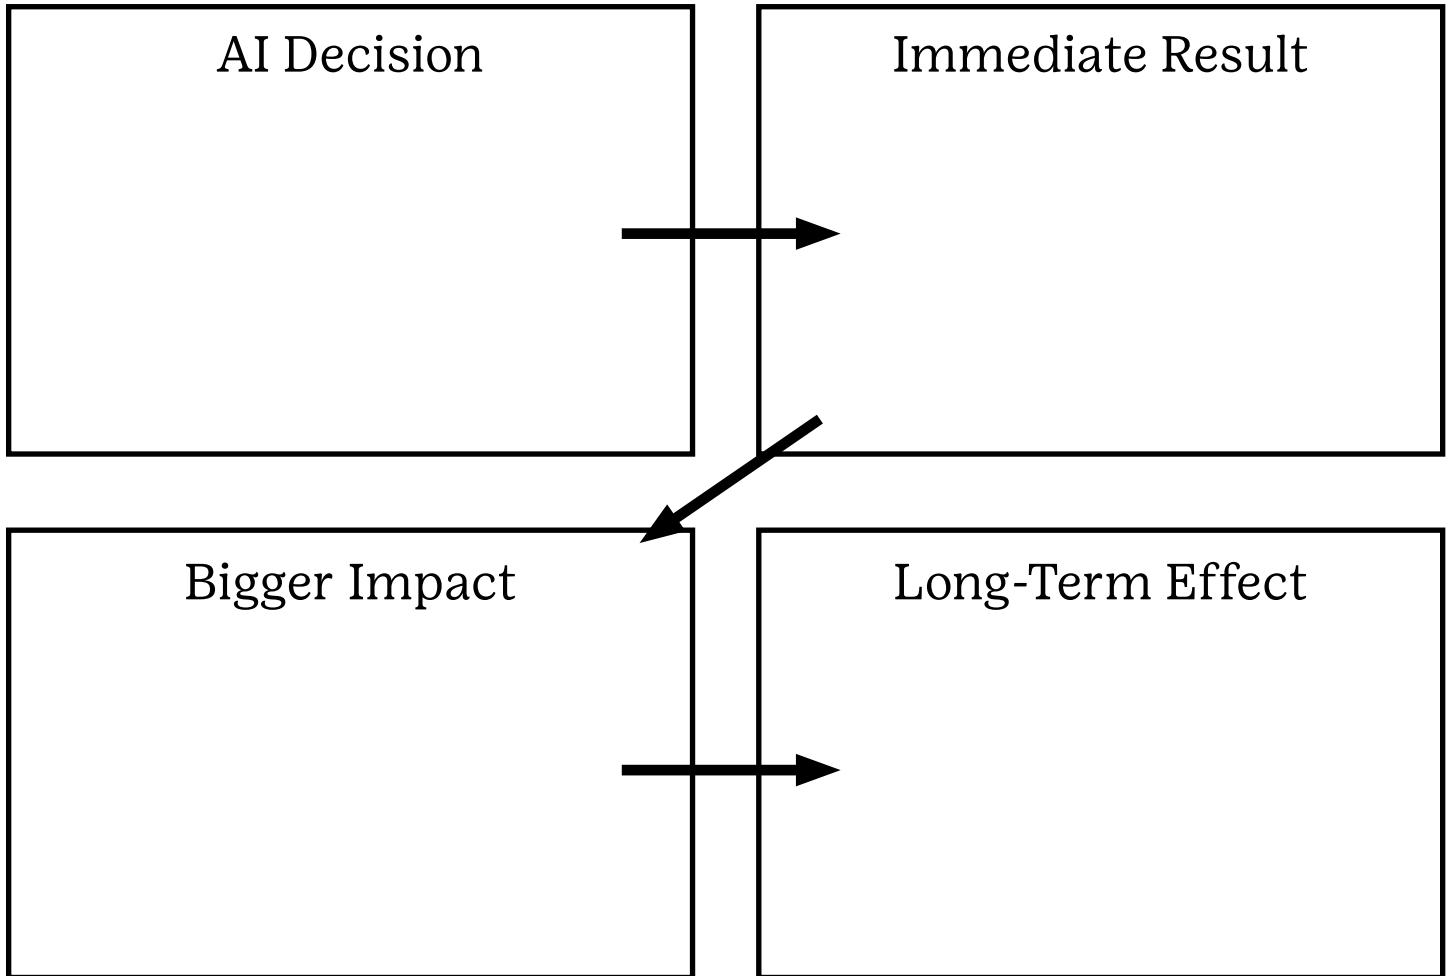
Possible mistake: _____

Who could be harmed? _____

Failure situation: _____

Consequence Chain

My AI will help by:



Fix the Problem

Fairness Rule: _____

Safety Feature: _____

Human Control: _____

Would You Trust It? Yes No Not Yet

Why?: _____

Senior Workbook

AI is already making decisions about what you see, what you buy, and how people are treated. Today you are not just learning about AI, you are becoming the person who decides how it works.



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434Q9Pt> Passcode: G\$Patch26

Learning Goals

Girls will:

- Understand how AI systems are built (data → decisions → impact)
- Evaluate real-world AI risks and benefits
- Design an AI solution for a global problem
- Defend their system with ethical reasoning

Solo Prompt: Before you do anything else, write down every place you interacted with AI in the last 24 hours. Think about your phone, your social media, your searches, your music, your shopping.

- _____
- _____
- _____
- _____

Activity 1: AI Case Lab

Should this AI exist?

Concept: Some AI systems are designed to help people. But good intentions don't always lead to good outcomes. In this activity you will read four real-world AI cases and make the call — should each one exist or not?

Your Role: You are the decision-maker. For each case, you have the final say. Read the scenario, analyze it, and make your ruling.

Case 1: Predictive Policing

What happened: A major city deployed an AI system that analyzes crime data to predict which neighborhoods are likely to have criminal activity. Police are then directed to increase patrols in those areas. Supporters say it reduces crime before it happens. Critics point out that the AI was trained on historical arrest data — which reflects decades of unequal policing in certain neighborhoods. The AI keeps sending officers to the same areas, which leads to more arrests there, which feeds back into the data, reinforcing the pattern.

What is the AI doing? _____

What are the benefits? _____

What are the risks? _____

Your ruling: You are the decision-maker. This AI launches or doesn't based on your call.

- Allow it
- Block it
- Allow it only if: _____

Case 2: AI Job Screening

What happened: A large corporation uses an AI system to screen job applications before any human sees them. The AI scores candidates based on resume language, work history, and other factors. It rejects roughly 75% of applicants automatically. An audit later found the system consistently scored women lower than men for technical roles — because it was trained on the company's past hiring data, which skewed heavily male. The company has since updated the system but still uses AI screening.

What is the AI doing? _____

What are the benefits? _____

What are the risks? _____

Your ruling: You are the decision-maker. This AI launches or doesn't based on your call.

- Allow it
- Block it
- Allow it only if: _____

Case 3: Student Behavior Monitoring

What happened: A school district installed an AI system that monitors students' activity on school-issued devices. It scans emails, documents, and searches for keywords related to violence, self-harm, or cheating. When flagged, a report goes automatically to a school counselor or administrator. Supporters say it has caught students in crisis before anything happened. Critics say it has also flagged students writing fiction, doing research on sensitive topics for class, and privately discussing mental health with friends.

What is the AI doing? _____

What are the benefits? _____

What are the risks? _____

Your ruling: You are the decision-maker. This AI launches or doesn't based on your call.

Allow it

Block it

Allow it only if: _____

Case 4: Mental Health Risk Detection

What happened: A health insurance company began using an AI system to analyze patient data — including prescriptions, doctor visit frequency, and billing codes — to identify people at high risk for depression or anxiety. The goal is to reach out proactively with support resources. However, patients were never told their data was being used this way. Some people flagged by the system later found their insurance premiums had increased.

What is the AI doing? _____

What are the benefits? _____

What are the risks? _____

Your ruling: You are the decision-maker. This AI launches or doesn't based on your call.

Allow it

Block it

Allow it only if: _____

Which case was hardest to rule on and why? _____

Did any case change how you think about AI you already use? _____

Activity 2: Build an Algorithm

Think like the system.

Concept: Rules sound simple until you have to write them. In this activity you will build a rule-based AI system, test it against hard cases, and find out where it breaks.

Step 1: Choose your scenario. Pick ONE:

- Who gets a scholarship
- What content gets promoted on social media
- Who gets emergency aid first after a disaster

Step 2: Write at least 5 IF/THEN rules your AI will follow. Each rule must award a specific number of points; your AI will use the total points to make its final decision. The profile or post with the most points wins.

Example for scholarships:

- IF student GPA is above 3.5 → award 3 points
- IF student volunteers in their community → award 2 points
- IF student works a part-time job → award 2 points
- IF student is first in family to attend college → award 1 point
- IF student has financial need → award 2 points

Your rules:

- IF _____ THEN award _____ points
- IF _____ THEN award _____ points
- IF _____ THEN award _____ points
- IF _____ THEN award _____ points
- IF _____ THEN award _____ points

Maximum possible points in my system: _____

My AI's decision rule: Anyone scoring _____ points or above gets selected.

Step 3: Test your algorithm. Find your scenario below. Run each case through your rules, add up the points, and record your results.

Scenario: Who gets a scholarship?

	Case 1	Case 2
Profile	A student has a 3.8 GPA and strong extracurriculars but their application is missing financial information. The form was submitted incorrectly and the data is blank.	A student has a 2.9 GPA but works two jobs, is a first-generation college student, has high financial need, and wrote an exceptional essay. They score high on some rules and low on others.
Points: Rule 1		
Points: Rule 2		
Points: Rule 3		
Points: Rule 4		
Points: Rule 5		
Total Points		
My AI's Decision		

Scenario: What content gets promoted?

	Case 1	Case 2
Post	A public health post about vaccine safety has strong engagement from verified medical accounts but the data field for total likes is missing. The post was uploaded with a formatting error.	A post has gone viral with 100,000 shares and enormous engagement, but 40% of the shares came from bot accounts and 20% of comments are flagging it as misinformation. It scores high on reach and low on credibility.
Points: Rule 1		
Points: Rule 2		
Points: Rule 3		
Points: Rule 4		
Points: Rule 5		
Total Points		
My AI's Decision		

Scenario: Who gets emergency aid first?

	Case 1	Case 2
Situation	A household has been flagged as needing aid but the damage assessment data is missing — the field team’s report didn’t upload correctly and the AI has no information about how badly the home was affected.	A single elderly woman with a medical condition lives alone in a moderately damaged home. A family of six with no medical needs lives in a severely damaged home nearby. The elderly woman scores high on vulnerability and low on damage level. The family scores the opposite.
Points: Rule 1		
Points: Rule 2		
Points: Rule 3		
Points: Rule 4		
Points: Rule 5		
Total Points		
My AI’s Decision		

Step 4: Evaluate your results. Look at your two decisions. For each one, answer honestly:

Case 1 — The missing data case:

What did your AI decide to do when information was missing? _____

Was that the right call? What should happen when an AI doesn’t have enough information to decide?

Case 2 — The conflicting traits case:

What did your AI decide when someone scored high on some rules and low on others?

Was that fair? Who might be consistently disadvantaged by the way your rules are weighted?

The one thing you would change: If you could rewrite one rule right now, which would it be and why?

A rule that seems fair in theory can fail the moment it meets a real person. The gap between your rules and your results is where bias lives.

Activity 3: Design your AI

Solve a real problem.

Step 1: Choose your problem. Pick ONE:

- Climate change
- Food insecurity
- Healthcare access
- Cyber safety
- Mental health
- My own idea: _____

Step 2: Define the problem clearly. Before you design anything, make sure you understand exactly what you are solving.

- What is happening right now that needs to change? _____
- Who is most affected? _____
- Why hasn't it been solved yet? _____

Step 3: Design your AI system.

What does it do? What action does it take? _____

What data does it use? What information does it learn from? _____

What decisions does it make? What choices does it control? _____

Who benefits and who might be left out? _____

Before moving to Activity 4: Read back over your answers. This is the system you are about to stress test, make sure it is specific enough to not break.

Activity 4: Stress Test your AI

Break it before the world does.

The Big Idea: Every AI system has weaknesses. Your job now is to find yours using the system you designed in Activity 3.

Step 1: Identify your risks.

Bias: Who could your AI treat unfairly? _____

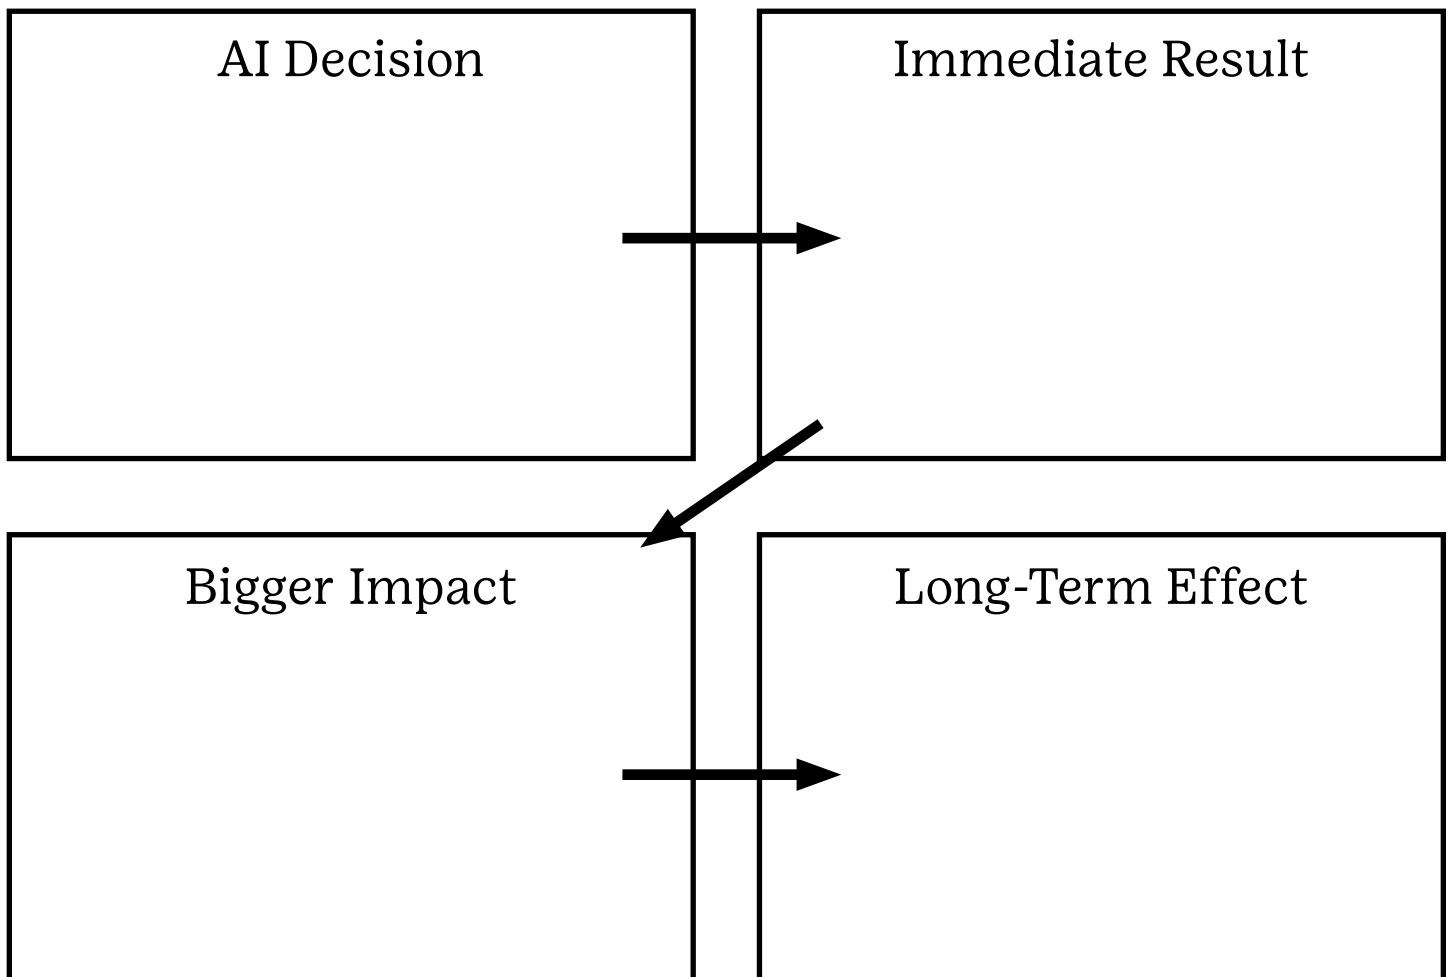
Misuse: How could someone use your AI to cause harm? _____

Failure: What situation would cause your AI to break down completely? _____

Step 2: Build the consequence chain. Take your most serious risk and map out how it could spiral:

Example: AI removes harmful posts → removes too many posts → people feel silenced → trust in platform drops

Consequence Chain



Step 3: Fix your AI. Add all three safeguards to the system you designed in Activity 3:

A fairness safeguard—A rule that makes sure your AI treats everyone equally:

A safety feature—Something that catches mistakes before they cause harm:

A human oversight system—A point where a person reviews or overrides the AI:

Before moving to Activity 5: You now have a designed and stress-tested AI system. Your pitch in Activity 5 should reflect both — not just the original design but the improvements you made here.

Activity 5: AI Pitch Lab

Defend your system.

Concept: You have designed an AI and found its weaknesses. Now make the case for why it should exist anyway. Use everything you wrote in Activities 3 and 4 to build your pitch.

Your Pitch

The problem: What is broken right now and why does it matter? _____

The solution: What does your AI do and how does it work? _____

Why it matters: Who does this help and what changes in the world if it works? _____

The risks and how you handled them: What could go wrong and what safeguards did you build in?

The hardest question, answer it honestly: Why should the world trust your AI? _____

Final Reflection

What is the most important thing you learned today about how AI makes decisions? _____

What is one thing you will do differently now that you know how AI works? _____

What is one thing humans must never let AI decide alone? _____

You are not just a user of AI. You are a future builder, critic, and decision-maker. The world needs people who ask hard questions and don't stop until they get honest answers.

Ambassador Workbook

AI is not neutral. Every system reflects decisions made by people, about what data to collect, whose problems to solve, and who gets left out. Today you will analyze those decisions, make your own, and set the rules for how AI should work.



Required: Watch the AI Spotlight Video Session where we'll explore the world of artificial intelligence and guide you through the information you need to complete this workbook.

<https://bit.ly/434Q9Pt> Passcode: G\$Patch26

Learning Goals

Girls will:

- Analyze AI as a system within society
- Evaluate trade-offs (innovation vs risk)
- Design a high-impact AI solution
- Anticipate unintended consequences
- Create guidelines/policy-level thinking
- Defend their position as a decision-maker

AI Reality Check: Before you do anything else, list 5 AI systems that are already shaping your life. For each one, write who benefits and who might be disadvantaged.

AI System	Who Benefits	Who is Disadvantaged

AI is not neutral. It reflects decisions made by people. Keep that in mind for everything that follows.

Activity 1: AI Systems Mapping

See the whole system.

Concept: Most people only see what AI does and not how it works or who it affects. In this activity you will map an entire AI system from beginning to end so you can see the full picture.

How a Completed Map Looks: Before you start, read this example so you know what level of detail is expected. Use this as a guide and not a template to copy. Your map should be specific to the system you choose.

Example System: Social Media Algorithm

- Problem: People don't see all content and there is too much to show everything, so something has to decide what appears first
- Data: Collects every click, pause, like, share, and how long you look at each post
- Algorithm: Scores content based on predicted engagement, posts you are likely to interact with get higher scores
- Decision: Controls what appears at the top of your feed and what you never see at all
- Impact: People see content that confirms what they already believe; advertisers pay to influence what gets prioritized
- Who benefits: The platform earns more ad revenue; users see content they enjoy
- Who is harmed: People who get trapped in misinformation loops; smaller voices who can't pay to be seen; users whose mental health is affected by what the algorithm amplifies

Step 1: Choose a system to map.

- Social media algorithm
- Hiring AI
- Predictive policing
- Health diagnostics AI

Step 2: Map it in the workspace boxes below.

Problem: What problem is this AI trying to solve? _____

Data: What information does it collect and learn from? Where does that data come from? _____

Algorithm: What rules or patterns does it use to make decisions? _____

Decision: What does it actually decide or control? _____

Impact: What happens in the real world as a result? _____

Who benefits? _____

Who is harmed? _____

Step 3: Find the weakest point. Look back at your map. Where is the system most vulnerable to failure, bias, or misuse?

The weakest point in this system is: _____

Because: _____

Step 4: Reflect. If you could change one fundamental decision that was made when this system was built, not just fix a bug but change something at its core, what would it be and why?

Activity 2: Trade Off Analysis

There is no perfect solution.

Concept: Every AI decision involves trade-offs. Improving one thing often means sacrificing something else. In this activity you will work through a real AI dilemma and take a position you can defend.

Choose your dilemma:

- Dilemma A: Surveillance vs. Privacy
- Dilemma B: Automation vs. Employment
- Dilemma C: Content Filtering vs. Free Speech

Dilemma A: Surveillance vs. Privacy

What is happening: A city government has deployed an AI surveillance network across all public spaces; cameras on every corner, facial recognition at transit stations, and automated license plate tracking on every road. Crime rates have dropped 34% since the system launched. Response times to emergencies are faster than they have ever been. But civil liberties groups have raised alarms: the system logs the movements of every person in the city, stores that data indefinitely, and has been used in three cases to identify people attending political protests. The city argues the safety benefits outweigh the privacy costs. Critics argue that a system built to fight crime is being used to monitor political activity and that most residents never consented to being tracked.

Dilemma B: Automation vs. Employment

What is happening: A large logistics company has replaced 60% of its warehouse workforce with AI-powered robots over the past three years. Productivity has increased by 40%, costs have dropped significantly, and the company has passed some savings on to consumers through lower prices. But 8,000 workers lost their jobs, most of them in a region where warehouse work was one of the few stable employment options available. The company offered retraining programs, but fewer than 12% of affected workers completed them and found new employment. Economists are divided: some say automation creates new jobs in the long run. Others say the workers who are displaced now cannot wait for the long run.

Dilemma C: Content Filtering vs. Free Speech

What is happening: A major social media platform deployed an AI content moderation system designed to automatically remove posts containing hate speech, misinformation, and harassment. In its first year the system removed 14 million pieces of harmful content that human moderators would never have caught in time. But independent researchers found it was also removing legitimate content at a significantly higher rate for users writing in certain languages and dialects, particularly those from minority communities. Posts in those languages were flagged as suspicious even when they contained no harmful content. The platform says it is working on the problem. Critics say the system is causing real harm to the communities it was supposed to protect.

For your chosen dilemma, answer:

What are the benefits of this AI system? _____

What are the risks? _____

Who is currently deciding what matters more in this situation? _____

Is that person or institution the right one to be making that call and why? _____

If not them, then who should decide, and how would that work in practice? _____

Where do you stand? _____

Now consider the scale: If your position just became policy affecting millions of people. Does that change your answer and if so, how?

Activities 3, 4, and 5 are one connected project.

How This Works: In Activity 3 you design a high-impact AI system. In Activity 4 you break it and add safeguards. In Activity 5 you write the governance policy that would govern it. Everything builds your policy in Activity 5 should reflect everything you learned in Activities 3 and 4. Don't rush Activity 3 because every answer carries forward.

Activity 3: AI for Good Innovation Lab

Design a system that matters.

Concept: You have spent time analyzing what AI does wrong. Now design one that does something right and be specific enough that you can actually stress test it in Activity 4.

Step 1: Choose your problem. Pick ONE:

- Climate change
- Food insecurity
- Healthcare access
- Cyber safety
- Mental health
- Economic inequality
- My own idea: _____

Step 2: Define the problem precisely. Vague problems lead to vague solutions. Be as specific as possible.

- What is happening right now that needs to change? _____
- Who is most affected and why has it been hard to solve? _____
- What would success actually look like? _____

Step 3: Design your AI system.

What does it do? What action does it take? _____

What data does it use? What information does it learn from and where does that come from?

What decisions does it make? What does it control and what does it leave to humans?

Who benefits and who might be left out? _____

What makes this different from solutions that already exist? _____

Before moving to Activity 4: Read your design back. Is it specific enough that someone could actually build it or test it? If it still feels vague, add one more detail before moving on.

Activity 4: Break your AI

If you don't, someone else will.

Concept: Good intentions are not enough. Using the system you designed in Activity 3, find every way it could fail, be misused, or cause unintended harm.

Step 1: Identify your failure points. For each failure type, write a specific example based on your system:

Bias: Which group of people could your AI treat unfairly, and why? _____

Misuse: How could a bad actor deliberately exploit your system? _____

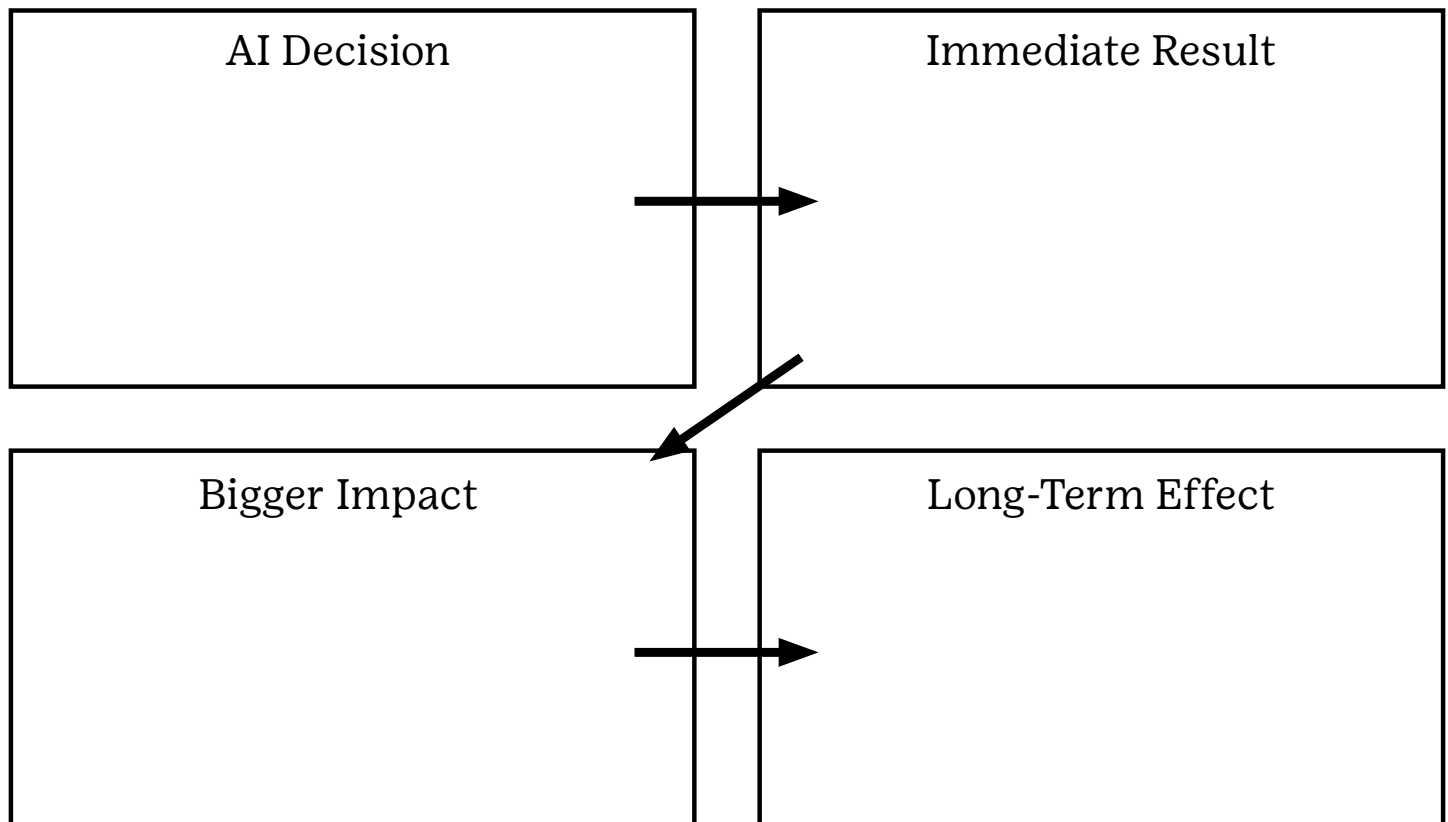
Misinterpretation: How could your AI reach a wrong conclusion from correct data? _____

Overreach: Where could your AI start making decisions it was never meant to make? _____

Step 2: Build the consequence chain. Take your most serious failure and map how it could spiral:

Example: AI removes harmful posts → removes too many posts → minority communities disproportionately silenced → trust collapses → people leave the platform

Consequence Chain



Step 3: Add your safeguards. For each safeguard, write specifically how it applies to your system — not a generic answer:

A fairness safeguard—A rule that makes sure your AI treats everyone equally:

A safety feature—Something that catches mistakes before they cause harm:

A human oversight system—A point where a person reviews or overrides the AI:

A limit on what your AI can decide alone:

Step 4: The hardest question. Look back at your consequence chain from Step 2 before you answer this.

Could your AI cause harm even if everyone involved has good intentions? Be honest and specific.

What does your answer tell you about the limits of good design? _____

Before moving to Activity 5: Your governance policy should directly address the failures you found in Steps 1 and 2. Keep this page open while you write your policy.

Activity 5: AI Governance Policy

Defend your system.

The system this policy governs: _____

Section 1: Ethical Principles Write three principles that must guide how this AI operates. Be specific to your system and not generic statements that could apply to any AI.

Fairness: _____

Transparency: _____

Accountability: _____

Section 2: Rules for Use

When SHOULD this AI be used? _____

When should it NOT be used? _____

In what situations must a human make the final decision and not the AI? _____

Section 3: Safeguards

How will you prevent misuse? _____

How will you protect vulnerable groups? _____

What happens when the system fails? _____

Section 4: The Limits of Your Policy

Every policy has gaps and things it cannot fully solve or situations it did not anticipate. This is the most important section.

What is the hardest thing to govern about your AI, the problem your policy addresses but does not fully solve?

Why is it so difficult to govern? _____

Knowing this gap exists, would you still deploy your AI? Why or why not? _____

Final Reflection

What would have to be true about an AI system before you would trust it with a decision that directly affects your own life? _____

Right now, who has the most power over how AI develops and is that the right answer?

You will spend your adult life in a world shaped by AI systems being built right now. What is the one thing you want the people building them to understand?

The most dangerous AI systems are not the ones built by bad people. They are the ones built by good people who stopped asking hard questions. Don't stop.