CHALLENGING

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

12–16 years | 50–75 min | 8–30 participants | Design challenge, teamwork and presentation

OVERVIEW

Participants design an AI-supported solution to a real social, scientific or environmental challenge. Beginning with the problem rather than the technology, teams decide whether AI is genuinely useful, what data the system would need, who might benefit or be harmed and what safeguards should be included. They present their proposal and receive feedback from other participants acting as an ethics and innovation panel.

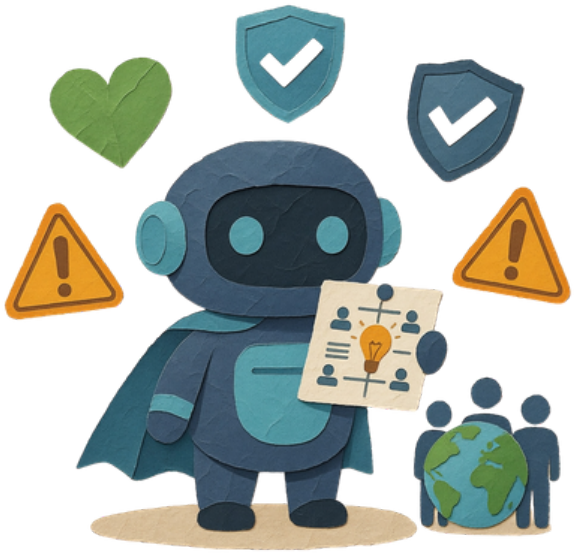

LEARNING OBJECTIVES

Participants will:

- apply their understanding of AI to a real-world challenge
- distinguish between a problem, a goal and a technological solution
- identify the data and AI capabilities a proposed system would require
- assess potential benefits, limitations and unintended consequences
- consider privacy, bias, environmental impact and human responsibility
- design practical safeguards and measures of success
- communicate and defend a responsible technology proposal

MATERIALS

- Challenge cards
- AI capability cards
- AI design worksheets or large sheets of paper
- Markers, pens and coloured pencils
- Adhesive notes
- Optional craft materials for creating a simple prototype or poster
- Benefit, risk and safeguard cards
- Panel feedback sheets
- Optional voting tokens

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

MATERIALS

Suggested challenge cards include:

- reducing food waste at school
- monitoring endangered wildlife
- improving access to learning materials
- detecting water leaks
- helping farmers identify signs of plant disease
- improving public transport routes
- reducing energy waste in buildings
- making public information more accessible
- supporting early warning for floods or extreme weather
- helping people sort waste correctly
- monitoring air or water quality
- assisting healthcare workers in areas with limited resources

Suggested AI capability cards include:

- classify images
- recognize patterns
- make predictions
- generate text or images
- translate languages
- recommend options
- detect unusual changes
- process voice commands
- analyse sensor data
- control a robot or connected device

PREPARATION

Prepare enough challenge cards for small groups of three to five participants.

Create an AI design worksheet with the following sections:

1. The problem
2. People and environments affected
3. Current solutions
4. Why AI may help
5. What the AI will do
6. Data required
7. Human role
8. Expected benefits
9. Possible errors and harms
10. Safeguards
11. Environmental cost
12. How success will be measured

Prepare feedback sheets containing four questions:

- Is the problem clearly defined?
- Is AI appropriate and necessary?
- Are the data, risks and affected people considered?
- Are the proposed safeguards sufficient?

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

INTRODUCTION

Ask the participants:

- **What does "AI for good" mean?**
- **Can a useful intention still lead to harmful results?**
- **Should every problem be solved using AI?**
- **Who should decide whether a new AI system is introduced?**
- **How can we know whether a system actually helps?**

Explain: Technology should begin with an understanding of the problem and the people affected. A responsible designer does not begin by asking, "Where can we use AI?" The first questions are, "What is the problem, what is already being done and would AI genuinely improve the situation?"

Tell participants that their goal is not to design the most futuristic system. Their goal is to develop the most useful, realistic and responsible proposal.

INSTRUCTIONS

1) SELECT AND UNDERSTAND THE CHALLENGE

Divide participants into small groups.
Give each group a challenge card or allow it to choose a locally relevant issue.

Groups discuss:

- **What exactly is happening?**
- **Who or what is affected?**
- **What may be causing the problem?**
- **What is already being done?**
- **What information is missing?**
- **What would improvement look like?**

Ask groups to write a one-sentence problem statement.

For example: Our school cafeteria throws away food because it does not know how many portions students will need each day.

The problem statement should not mention AI yet.

2) DECIDE WHETHER AI IS APPROPRIATE

Give groups the AI capability cards.

Ask them to consider:

- **Does the problem involve patterns, predictions or large amounts of information?**
- **Could a simple non-AI solution work?**
- **Would AI provide a clear advantage?**
- **Is sufficient and appropriate data available?**
- **Are the consequences of mistakes manageable?**
- **Would the people affected accept the system?**

Groups choose one of three conclusions:

- **AI could make a useful contribution.**
- **AI may help, but only as part of a larger solution.**
- **A simpler non-AI solution would be more appropriate.**

Allow groups to propose a non-AI solution if that is their reasoned conclusion. Recognizing that AI is unnecessary is a successful outcome.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

INSTRUCTIONS

3) DEFINE THE AI'S TASK

If AI is appropriate, groups define one specific task for it.

For example:

- Predict the approximate number of cafeteria portions needed
- Classify wildlife images from camera traps
- Identify possible signs of disease on plant leaves
- Detect unusual changes in water use
- Translate public information into several languages
- Recommend more efficient public-transport routes

Ask groups to avoid vague descriptions such as "The AI solves food waste."

Instead, they should complete:
"The AI helps by ..."

4) IDENTIFY THE REQUIRED DATA

Groups list the information their system needs.

For the cafeteria example, this might include:

- Previous numbers of meals served
- Day of the week
- School calendar
- Menu choices
- Attendance estimates
- Amounts of food left over

Ask:

- Where do the data come from?
- Who provides them?
- Are they accurate and representative?
- Do they contain personal or sensitive information?
- Could the system work with less data?
- Who should be allowed to access them?
- How long should they be stored?

Encourage groups to follow the principle of data minimization: collect only what the system genuinely needs.

5) MAP THE HUMAN ROLE

Groups decide what people will do before, during and after the AI system operates.

They complete:

"A person sets the goal by ..."
"The AI supports the task by ..."
"A person checks the result by ..."
"A person makes the final decision when ..."
"A person can stop or override the system if ..."

Explain that saying "A human checks it" is not enough. The reviewer must have the knowledge, authority, information and time to identify and correct problems.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

INSTRUCTIONS

6) IDENTIFY BENEFITS

Groups list the people, communities or environments that may benefit.

They should describe specific improvements, such as:

- Less food discarded
- Earlier detection of water leaks
- Faster review of wildlife images
- Improved access to information
- Reduced energy use
- Additional support for healthcare workers

Ask who might benefit most and whether any groups could be excluded.

7) CONDUCT A RISK TEST

Give groups risk cards or ask them to consider:

- Inaccurate prediction
- Biased training data
- Missing or unusual cases
- Privacy breach
- Overreliance on the system
- Unclear responsibility
- Lack of access
- Environmental cost
- System failure
- Misuse for another purpose
- People being monitored without meaningful choice

Each group selects at least three relevant risks and describes who or what could be affected.

Ask: What is the worst reasonable outcome if the system makes a mistake?

The aim is to identify realistic risks, not imagine an impossible catastrophe.

8) DESIGN SAFEGUARDS

For every identified risk, groups propose at least one safeguard.

Possible safeguards include:

- Human review
- Regular accuracy testing
- Representative training data
- Clear limitations and warnings
- Collecting less personal data
- Secure data storage
- A way to correct or challenge decisions
- A non-digital alternative
- Testing before full deployment
- Pausing the system when errors increase
- Consulting affected communities
- Measuring energy and resource use

Groups should connect each safeguard directly to a risk.

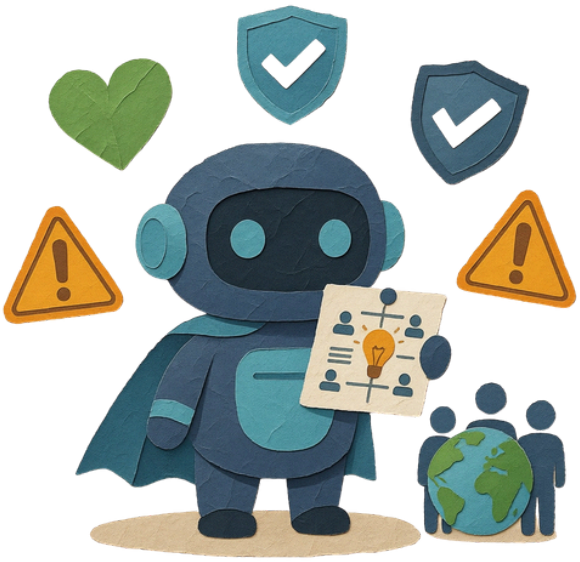

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

INSTRUCTIONS

9) MEASURE SUCCESS

Ask groups how they will determine whether the system helps.

Measures should relate to the original problem.

For example:

- Percentage reduction in food waste
- Number of water leaks detected and confirmed
- Time saved by conservation researchers
- Translation accuracy checked by speakers of the language
- Reduction in building energy consumption
- Number of incorrect alerts
- Satisfaction of the people using the service

Explain that success includes both benefits and the absence of unacceptable harm.
A system should not be judged only by speed or technical accuracy.

10) CREATE THE PROPOSAL

Groups create a poster or design sheet containing:

Project name

- Problem statement
- Description of the AI-supported solution
- Data required
- Human responsibilities
- Expected benefits
- Main risks
- Safeguards
- Environmental considerations
- Measures of success

Optional prototypes might include:
a paper interface
a flow diagram
a model sensor
a sequence of user interactions
a storyboard showing how the system works

11) PRESENT TO THE PANEL

Each group gives a three-minute presentation to the other participants, who act as an ethics and innovation panel.

The panel asks questions such as:

- Why is AI needed?
- What happens if the system is wrong?
- Whose data does it use?
- Who has been consulted?
- Who makes the final decision?
- Can someone refuse to use it?
- How will you know it works?
- What environmental resources does it require?
- Could it be misused?

After the questions, the panel recommends:

- Proceed to testing
- Revise the design
- Use a simpler solution
- Do not proceed

The recommendation must include a reason.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

INSTRUCTIONS

12) IMPROVE THE DESIGN

Give groups time to make one or two changes based on the panel feedback.

Each group finishes by stating:

"The most important improvement we made was ..."

This reinforces that responsible design is an ongoing process rather than a single decision.

DISCUSSION AND REFLECTION

Ask the participants:

- Did your group begin with the problem or the technology?
- Was AI necessary for every challenge?
- Which part of the design process was most difficult?
- What data did your system require?
- Who might be missing from those data?
- Who receives the greatest benefit?
- Who carries the greatest risk?
- What happens if people trust the system too much?
- Which safeguard was most important?
- Who should be responsible when the system causes harm?
- How should affected communities participate in the design?
- Can a project still be described as "AI for good" if it creates serious environmental or social costs?
- What evidence would convince you that the project genuinely works?

Conclude with the key idea: Responsible AI design is not only about what technology can do. It requires decisions about what it should do, who it serves, what risks are acceptable and who remains accountable for the consequences.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

CONNECTION TO ARTIFICIAL INTELLIGENCE

AI can support social, scientific and environmental work by:

- analysing large datasets
- identifying patterns that people may not easily detect
- making predictions
- classifying images or signals
- translating and generating accessible information
- monitoring changes
- supporting planning and decision-making

However, a technically impressive system is not automatically useful or responsible.
Its impact depends on:

- whether the problem has been correctly understood
- whether AI is appropriate for the task
- the quality and representativeness of the data
- how the system is used in practice
- who has access to its benefits
- who may be harmed or excluded
- whether its results can be challenged
- whether meaningful human responsibility remains
- the resources required to build and operate it

Responsible innovation involves affected people throughout the process and continues after deployment. Systems need monitoring, evaluation and revision as conditions change or new problems emerge.

EASIER VARIATION

Choose one shared challenge for the whole group, such as reducing food waste at school.

Guide participants through five questions:
1. What is the problem?
2. How might AI help?
3. What information would it need?
4. What could go wrong?
5. What rule would make it safer?

Small groups draw their solutions and present them in one minute.
The facilitator combines the strongest ideas into one shared design.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

MORE CHALLENGING VARIATION

Assign each group the same challenge but a different stakeholder perspective:

- young people
- local community
- scientists
- government
- technology company
- environmental organization
- people whose data are collected

Groups assess the same proposed AI system from their assigned perspective.

They negotiate a final project agreement covering:

- purpose
- data collection
- access and inclusion
- human oversight
- acceptable error levels
- environmental limits
- complaint and appeal procedures
- conditions under which the system must be stopped

Alternatively, groups can conduct a "red team" exercise by exchanging designs and deliberately searching for weaknesses, misuse possibilities and unintended consequences.

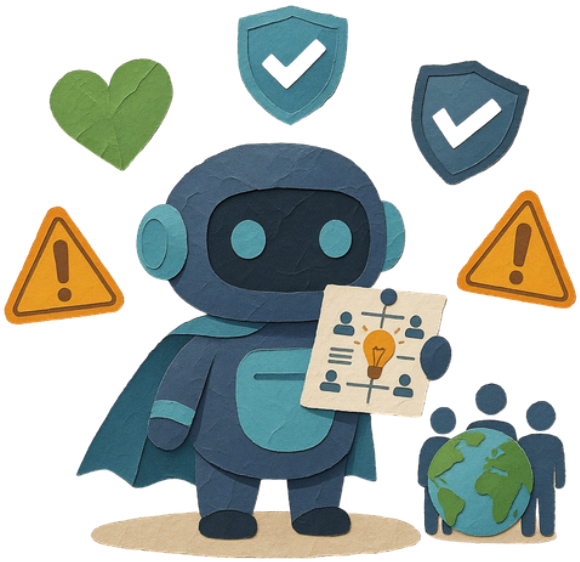

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

FACILITATOR NOTES

- Keep groups focused on a clearly defined problem.
- Do not require every group to use AI. A well-argued non-AI solution demonstrates strong critical thinking.
- Encourage participants to design one limited AI task rather than an all-purpose system.
- Avoid presenting technology as a replacement for investment, public services or human expertise.
- Ask who defines “good” and whether the people affected share the same priorities.
- Ensure that groups consider those who may lack devices, internet access or digital skills.
- Require a specific human role and clear responsibility for important decisions.
- Connect safeguards to identified risks rather than listing generic principles.
- Encourage realistic measures of success.
- Explain that high technical accuracy does not eliminate all social risks.
- Avoid collecting real personal data during the activity.
- Treat criticism as part of responsible innovation rather than opposition to technology.
- Emphasize that design continues after introduction through monitoring, feedback and revision.

DESIGN AN AI FOR GOOD

INNOVATION, ETHICS AND SOCIAL IMPACT

Science Film Festival

KEY WORDS

AI FOR GOOD	The use of AI intended to support beneficial social, scientific or environmental outcomes.
PROBLEM STATEMENT	A clear description of the issue a project aims to address.
STAKEHOLDER	A person or group affected by or able to influence a project.
DATA MINIMIZATION	Collecting only the data necessary for a specific purpose.
SAFEGUARD	A measure designed to prevent or reduce harm.
HUMAN OVERSIGHT	Meaningful human review and control of an automated system.
ACCOUNTABILITY	Responsibility for decisions, actions and consequences.
PROTOTYPE	An early model used to test and improve an idea.
EVALUATION	A structured assessment of whether a system achieves its goals and what effects it produces.
RESPONSIBLE INNOVATION	Developing technology with careful attention to its purpose, impacts, affected communities and long-term consequences.