

AI and Ethics: A Mini-Unit for High School

Activity Overview: This activity set is a week-long unit designed for 90-minute class periods utilizing curriculum and materials from the DAILY Workshop, aimed at introducing students to AI, how it works, and ethical considerations/impacts. If the class period is shorter than 90 minutes, the lessons will likely need to be broken up into two chunks.

- Day One: Students will learn about what AI is, how it works, and identify areas of their life where they use AI.
- Day Two: Students will be introduced to supervised machine learning, including classifying and generating. Students will then build their own teachable machine and examine how data sets influence results.
- Day Three: Students will learn about algorithmic bias and begin considering ethics as they related to artificial intelligence.
- Day Four: Students will learn about socio-technical systems and examine the goals of these systems by examining the recommender algorithm on YouTube.
- Day Five: Students will continue their learning about socio-technical systems, algorithms, and ethics as they redesign the recommender algorithm for a stakeholder group they select in order to make it less bias and more ethical.
- Day Six: Students will look at how they use AI in their own lives and devise a plan to use it ethically.

Extension Options: Any activities from the DAILY Workshop curriculum can be integrated into this mini unit to extend learning and/or challenge students as needed.

Target Grade Levels: 9 - 12

Target Subject: Computer science, technology, engineering (Note: Lesson is suitable for various content areas)

Notes for Educators: These lessons were originally taught in a 9th grade ELA class, but can be modified to fit any secondary classroom. Basic lesson structure and materials were taken from the DAILY Workshop and the MIT AI Ethics Education Curriculum. Links to both sites, as well as links to additional resources, can be found at the end of this unit plan.

Materials: Students will need access to a device with internet access (iPad or laptop/computer). Teacher will need access to an interactive board, such as a SmartBoard or Newline. The lessons work best if the teacher and students have access to Google Chrome and Google education suite (Slides, Drawing, Docs, etc.).

Day One – Intro to Artificial Intelligence

Lesson Resources:

- What is AI? Main Script and Slides, AI tools slides from Explore AI Journal, AI or Not activity:
 - Script: <https://docs.google.com/document/d/1FfNUyYeAhQ4JBhz4xRZC1t6oZHay6585BD0QTt4Lt6M/copy>
 - Slides: https://docs.google.com/presentation/d/1mMBoyiMDC0imjDgclXFGpvuCIL3iTp12HI2UcZf2_M/copy
- AI or Not? Activity: https://docs.google.com/drawings/d/1Z4DiMjmd4aUDFXaXCC1NApJn_Vl084CgaaeErkJZXmU/copy
- AI Journal Slide Deck: <https://docs.google.com/presentation/u/0/d/1zElK7rfmcoJrpzWAaczJCo2NKbrSlnahrguQfu-Qdl0/copy>

Guiding Questions:

- What is artificial intelligence?
- How does artificial intelligence work?
- Where do I encounter artificial intelligence in my everyday life?

Objectives:

- Understand that artificial intelligence is the science of making machines that can think like humans and have intellectual processes similar to a human – reasoning, discovering meaning, generalizing, learning from past experiences.
- Know that artificial intelligence is a specific type of algorithm and has three specific parts: dataset, learning algorithm, and prediction.
- Recognize AI systems in everyday life and be able to reason about the prediction an AI system makes and the potential datasets the AI system uses.

Catch/Hook: Students complete the “AI or Not” sorting activity.

Activity Instructions:

1. After students have completed the “AI or Not” sorting activity, have them share out as to why they put certain technologies into each category. Ask if anyone has a definition for AI that they are willing to share with the group. Using the student given definition, build upon their understanding by going over the Introduction to AI slide deck that covers algorithms, datasets, learning, and prediction.
2. Showcase the various AI tools from the Explore AI Journal slide deck. Have each student pick one they would like to check out. Allow time for students to play with their chosen AI tool, then in their groups, have them talk about what data the machine was given, how it learned, and what it predicted.
3. Have one person from each group share out about what tool they used, what the tool did, and what data/learning/prediction they think it might have used.

Review: Review the definition of AI, what an algorithm is and why it is important, and the data, learning, prediction cycle.

Assessment: Students fill out an exit ticket (see Printable Lesson Resources at the end of this document) with three things they learned and one AI tool that they use in their daily life. Bonus if they can write what they think the dataset and prediction is for that tool.

Standards:

- L1.AP.A.02 Describe how artificial intelligence algorithms drive many software and physical systems.

Day Two – Cats or Dogs? Intro to Supervised Machine Learning

Lesson Resources:

- AI Bingo (printable on page 15)
(<https://docs.google.com/document/d/1e9wx9oBg7CR0s5O7YnYHVmX7H7pnITfoDxNdrSGkp60/edit>).
- Introduction to Supervised Machine Learning
 - Script: https://docs.google.com/document/d/12qbJx6p5vHkkUqa775gIJ_4gPrQgWADEPR2DIm65U4/copy
 - Slides: <https://docs.google.com/presentation/d/1CbV4ziok2h1LvoM-PN6P1GtVwafq67MG4GOgFPU9PV8/copy>

Guiding Questions:

- What is supervised machine learning?
- What is a classification problem?
- How do data sets influence machine learning predictions/outcomes?

Objectives:

- Be able to define and explain supervised machine learning.
- Know what a classification problem is and how AI models work with it.
- Understand potential issues with classification in the supervised machine learning context.
 - Understand how the quantity of training data affects the accuracy and robustness of a supervised machine learning model.

Catch/Hook: To review the data/prediction sequence, students will participate in AI Bingo.

Activity Instructions:

1. After the students have participated in AI Bingo, review the concepts of machine learning, data set, and prediction after the activity. Tell the students that today you are going to learn more about supervised machine learning and go over the Introduction to Supervised Machine Learning slide deck through slide 23. Be sure to use version 1 of the Teachable Machine for the student activity (<https://teachablemachine.withgoogle.com/v1>). Debrief with students utilizing the questions from MIT AI Ethics Education Curriculum (linked above).
2. Let students know that in order to experience how a machine learns to classify, they are going to work with a partner to build a model that will classify dogs and cats. Work through the Dog/Cat classifier section of the MIT AI Ethics Education Curriculum, making sure to ask the questions dealing with AI bias in data.
3. Depending on the ability level of the students and the time available, do either the Musical Instrument Classification activity from the MIT curriculum or the Teachable Machine Project from the DAILY Workshop.

Review: Review the definition of a classification problem, how a machine learning model can solve it, and touch on bias in data sets.

Assessment: Students will fill out an exit ticket (see Printable Lesson Resources) where they answer the following questions about their experience with the musical instrument model or their own classification model: What did their model classify? Was their model successful in classification? What did they learn about the type of data their model needed to perform successfully? Where was their model biased?

Standards:

- L1.AP.A.01 Create a prototype that uses algorithms (e.g., searching, sorting, finding shortest distance) to provide a possible solution for a real-world problem relevant to the student.
- L1.IC.C.02 Test and refine computational artifacts to reduce bias and equity deficits.

Day Three - Algorithmic Bias

Lesson Resources:

- Investigating Bias and Exploring AI Bias activity from the DAILY Workshop

- Script: <https://docs.google.com/document/d/1LATyHyqnQ6bYHTvfv8xlgpavaNrtSgYcsd1gjAavZ5g/copy>
- Slides: https://docs.google.com/presentation/d/1Eal-m7aWBP7_JX4r0eDCy_FUuocishfpPiYWwBtgzFc/copy
- Handouts:
 - https://docs.google.com/document/d/1axNaHnXQJ54ywHB3uPVo3iqh_i2EWq_d9sAQHSm871E/copy
 - <https://docs.google.com/document/d/1MiArHt8O-zeW6ZyLYLTTCNmf3egluTu70mNNHMAKIEg/copy>
 - <https://docs.google.com/document/d/1N1DyIUBfPC4-F406cIFP-AgRGmnE6vjulRVLtm0KASo/copy>
 - https://docs.google.com/document/d/1qcH9daB9ZioyEeCc7XCr2J1GSQILsldrq5g3_32wT7c/copy
- Top four ethical principles for artificial intelligence from AIClub
 - <https://www.corp.aiclub.world/post/teaching-ai-ethics-to-high-school-students>

Guiding Questions:

- What is algorithmic bias?
- What is the difference between fair and bias?
- How can artificial intelligence be biased?
- How do we prevent bias in AI?

Objectives:

- Know the term “algorithmic bias” in the classification context.
 - Understand the effect training data has on the accuracy of a machine learning system.
 - Recognize that humans have agency in curating training datasets.
 - Understand how the composition of training data affects the outcome of a supervised machine learning system.
- Know the term “ethics” in the context of artificial intelligence/machine learning.
 - Understand how bias relates to ethics in machine learning.
 - Identify how data sets can impact the predictions of machine learning systems allowing them to make unethical predictions.
 - Recognize why ethics in relation to AI are important.
 - Explain the top four principles of ethical artificial intelligence (fairness, transparency, privacy, human centeredness) and why they are important.

Catch/Hook: Have students watch this short clip on AI, humans, and bias (<https://youtu.be/AUVcF7ehZ28>). Ask students to think of one time they have encountered AI bias in their life and write it down. If they can’t think of a specific example of AI bias, they can write about a time they have experienced bias themselves. Invite a few students to share their experiences.

Activity Instructions:

1. Remind students that last time you met, the group worked on building classification models that sorted a piece of data into one group or the other. You also talked a little bit about how AI models can be biased. Today you are going to take a deeper look at bias in machine learning. Go through the slide deck for Investigating Bias through slide 15.
2. Next, tell students that the class is going to look at different real-world examples of bias and complete the Exploring AI Bias activity utilizing slides 16 – 35 and handouts 1-4.
3. Tell students that you are going to build on their understanding of bias in machine learning by looking at the concept of ethics. Ask the student to define the word ethical, then build on their decision and give

examples of situations where an ethical choice was made versus a choice that was not ethical. How do the real-world examples looked at on Google align with ethics? Discuss the top four ethical principles of artificial intelligence (from AI Club link above). How did students see those reflected in the Google search results? Where else have they encountered these principles?

Review: Review the terms algorithmic bias and ethics. Review how AI models can be biased and how that can render models that are not ethical.

Assessment: Have students complete the exit ticket where they answer the question “What is one way an algorithm can be biased?” and then write 2-3 sentences discussing a time when they utilized technology in an ethical way (see Printable Lesson Resources).

Standards:

- L1.AP.A.02 Describe how artificial intelligence algorithms drive many software and physical systems.
- L1.IC.C.01 Evaluate the ways computing impacts personal, ethical, social, economic, and cultural practices.

Day Four – Socio-technical Systems: The YouTube Recommender

Lesson Resources:

- Socio-technical systems diagram and info from Geeks for Geeks
 - <https://www.geeksforgeeks.org/socio-technical-systems/>
- MIT AI Ethics Education Curriculum section on how algorithms can have various motives and goals (page 95-97)
 - <https://docs.google.com/document/d/1e9wx9oBg7CR0s5O7YnYHVmX7H7pnITfoDxNdrSGkp60/edit>

Guiding Questions:

- What are socio-technical systems and how do they function in our society?
- How are AI ethics and socio-technical systems connected?
- What ethical considerations need to be made when designing a socio-technical system?

Objectives:

- Understand that all technical systems are socio-technical systems. Understand that socio-technical systems are not neutral sources of information.
 - Understand the term “optimization” and recognize that humans decide the goals of the socio-technical systems they create.
 - Reason about the goals of socio-technical systems in everyday life and distinguish advertised goals from true goals (for example, the YouTube recommendation algorithm aims to make profit for the company, while it is advertised as a way to entertain users).
 - Map features in existing socio-technical systems to identified goals.

Catch/Hook: Once class has started, ask the students the following questions: How many of you have used YouTube before? How do you decide which YouTube videos to watch?

Activity Instructions:

1. Tell students that today we are going to consider socio-technical systems (put the diagram from Geeks for Geeks on the SmartBoard). On a very basic level, a socio-technical system can be referred to as the

mixture of people and technology. It is the study of how technology is used and produced and looks at people, software, hardware, data, and laws/regulations. It is important to note that socio-technical systems always have a goal in mind and are not neutral sources of information. There are often many stakeholders in a given system (think program developers, users of differing backgrounds/needs, etc.) and these types of systems can impact them in different ways.

2. Next tell students we are going to use YouTube as a socio-technical system and examine its recommender function. Utilize the MIT AI Ethics Education Curriculum section on how algorithms can have various motives and goals (starting on page 94) to lead students through an activity examining the algorithms of the YouTube recommender.
3. Lead the group in a short discussion about the possible ethical impacts of the existing algorithms of the recommender, they will continue this train of thought on their exit ticket. Are the recommender algorithms ethical or not? How can they perpetuate the spread of misinformation? How do they exclude groups of people?

Review: Review the term/concept of a socio-technical system, remind students that these systems are not bias free and usually have some sort of motive or goal. Review the main facets of AI ethics.

Assessment: Have students fill out an exit ticket where they write three to five sentences reflecting on the ethical implications of the YouTube recommender algorithms. How have they seen these implications reflected in their lives or the lives of others?

Standards:

- L1.IC.C.01 Evaluate the ways computing impacts personal, ethical, social, economic, and cultural practices.
- L2.IC.C.01 Evaluate the beneficial and harmful effects that computational artifacts and innovations have on society.

Day Five - Ethical Algorithms: Redesigning YouTube

Lesson Resources:

- YouTube Redesign from the MIT AI Ethics Education Curriculum (page 98 – 102) (<https://docs.google.com/document/d/1e9wx9oBg7CR0s5O7YnYHVmX7H7pnITfoDxNdrSGkp60/edit>)
- Redesign YouTube script and slides from DAILY Workshop
 - Script: <https://docs.google.com/document/d/1kUFD1U4r8SGVY36fbgl-JAgIGInNCVku0uU-Q6EPc3k/copy>
 - Slides: <https://docs.google.com/presentation/d/1QwViFC9hLPsfGFRDADByoPw0r1vnV4NLSN1QASMqyLs/copy>

Guiding Questions:

- Who are potential stakeholders in socio-technical systems?
- What are different ways socio-technical systems can impact stakeholders?
- What ethical considerations need to be made when designing elements of a socio-technical system?

Objectives:

- Recognize there are many stakeholders in a given socio-technical system and that the system can affect these stakeholders differentially.
 - Identify relevant stakeholders in an socio-technical system.

- Justify why an individual stakeholder is concerned about the outcome of a socio-technical system.
- Identify values an individual stakeholder has in an socio-technical system, e.g. explain what goals the system should hold in order to meet the needs of a user.
- Construct an ethical matrix around a socio-technical system.

Catch/Hook: Upon arrival, students brainstorm ways they could change the YouTube recommender to better serve their friend group. This can be done on paper or, if seated in groups, students could discuss with their tablemates.

Activity Instructions:

1. Remind the class that last time you met, the group looked at the goals of the YouTube recommender and considered what impacts it has on stakeholders. Tell them that today they are going to redesign the YouTube recommender to better serve a group of stakeholders they select.
2. Put students into groups or let them select a group to work with. Using the slide deck from DAILY Workshop in the Redesign YouTube activity and the worksheets from the MIT AI Ethics curriculum, walk students through designing their own version of the YouTube Recommender, pausing as you go along to give the groups time to make decisions and fill out their sheets.
3. Have a few groups share their new algorithms!

Review: Review the concept of a socio-technical system, and why stakeholders/users are an important part of designing these systems.

Assessment: Each group will turn in their worksheets for their new algorithm.

Standards:

- L2.AP.PD.07 Modify an existing program to add additional functionality and discuss intended and unintended implications.
- L2.IC.C.01 Evaluate the beneficial and harmful effects that computational artifacts and innovations have on society.

Day Six – AI, Ethics, and You

Guiding Questions:

- How do I use AI in my daily life?
- How do I see examples of bias and/or ethical issues in the AI around me?
- What can I do to make sure I am using AI ethically in my life?

Objectives:

- Be able to define bias and ethics and explain how they relate to artificial intelligence.
 - Provide real-world examples of each and explain why ethical issues are present
- Define various ways to use AI ethically in life.
 - Know how to identify bias and other ethical issues in AI algorithms.
 - Be able to define ways to ethically use AI.

Catch/Hook: Students will do a quick write to answer the question “How do you use AI in your daily life?”

Activity Instructions:

1. Quickly review all the concepts covered during this mini unit – what is AI, what is supervised machine learning, bias in algorithms, socio-technical systems, and what AI ethics are/why they matter.
2. Ask students to think about how they interact with AI in their daily lives. Have them write down examples of the interactions they have, then have students evaluate whether or not the AI they encounter is biased or unethical in any way.
3. In small groups, have students consider how they are using AI in their daily lives. Have each group outline three ways that students can be ethical users/creators of artificial intelligence.
4. Have each group share out the ways they suggested for students to ethically use/create AI. Consider gathering the suggestions and posting them somewhere for the group to reference as needed.

Review: Completed in the first step of the lesson.

Assessment: Have students look back on the paragraph they wrote at the beginning of class on how they use AI in their daily life. Based on what they have learned during this mini unit, have them choose one area of their life where they could improve their use of AI. Have students craft a SMART (specific, measurable, achievable, relevant, and time-bound) goal to improve their use of AI in that area in order to become a more responsible and ethical AI user.

Standards:

- L1.IC.SLE.04 Using grade level appropriate content and complexity, discuss the legal, social, and ethical impacts associated with software development and use, including both positive and malicious intent.

Resources and Acknowledgements:

The DAILY Curriculum for Middle School Students was created by the MIT Media Lab Personal Robots Group and the MIT STEP Lab. More information, including links to the full curriculum, can be found at <https://raise.mit.edu/daily/>.

The MIT AI Ethics Education Curriculum is from the Ethics of Artificial Intelligence Curriculum for Middle School Students and was created by Blakeley H. Payne with support from the MIT Media Lab Personal Robots Group, directed by Cynthia Breazeal. The full curriculum can be found at <https://docs.google.com/document/d/1e9wx9oBg7CR0s5O7YnYHVmX7H7pnITfoDxNdrSGkp60/edit>.

Additional resources can be found on the AI4K12 website at <https://ai4k12.org/> and Everyday AI at [Resources | Everyday AI \(everyday-ai.org\)](#)

The Wyoming Computer Science Content and Performance Standards can be found here <https://edu.wyoming.gov/wp-content/uploads/2021/04/2020-CS-WYCPS-with-all-PLDs-effective-04.07.21.pdf>.

Printable Lesson Resources

- Exit ticket for Day One – Intro to Artificial Intelligence
- Exit ticket for Day Two – Cats or Dogs? Intro to Supervised Machine Learning
- Exit ticket for Day Three - Algorithmic Bias
- Exit ticket for Day Four – Socio-technical Systems: The YouTube Recommender
- SMART Goal Worksheet for Day Six – AI, Ethics, and You

EXIT TICKET

WHAT I LEARNED

Here are three things I learned today:

1. _____

2. _____

3. _____

AI TOOL

An AI tool that is use in my daily life is

_____.

Bonus: I think the dataset and prediction for this tool are:

EXIT TICKET

WHAT I LEARNED

Here are three things I learned today:

1. _____

2. _____

3. _____

AI TOOL

An AI tool that is use in my daily life is

_____.

Bonus: I think the dataset and prediction for this tool are:

EXIT TICKET

CLASSIFICATION MODELS

What did your model classify?

Was your model successful with its classifications?

What did you learn about the type of data your model needed to perform successfully?

Was your model biased? If so, how?

EXIT TICKET

CLASSIFICATION MODELS

What did your model classify?

Was your model successful with its classifications?

What did you learn about the type of data your model needed to perform successfully?

Was your model biased? If so, how?

EXIT TICKET

ALGORITHMIC BIAS

What is one way that an algorithm can exhibit bias?

Write 2-3 sentences about a time when you used AI in an ethical way.

EXIT TICKET

ALGORITHMIC BIAS

What is one way that an algorithm can exhibit bias?

Write 2-3 sentences about a time when you used AI in an ethical way.

EXIT TICKET

YOUTUBE RECOMMENDER

What are some ethical implications of the YouTube recommender algorithms?

IMPLICATIONS

How have you seen these implications reflected in your life or the lives of others?

EXIT TICKET

YOUTUBE RECOMMENDER

What are some ethical implications of the YouTube recommender algorithms?

IMPLICATIONS

How have you seen these implications reflected in your life or the lives of others?

SMART GOALS

Setting realistic and achievable outcomes.

My goal is:

S
SPECIFIC

What do I want to happen?

M
MEASUREABLE

How will I know when I have achieved my goal?

A
ATTAINABLE

Is the goal realistic and how will I accomplish it?

R
RELEVANT

Why is my goal important to me?

T
TIMELY

What is my deadline for this goal?