Chicago-Kent College of Law

Constitutional Democracy Project

Illinois Youth Summit

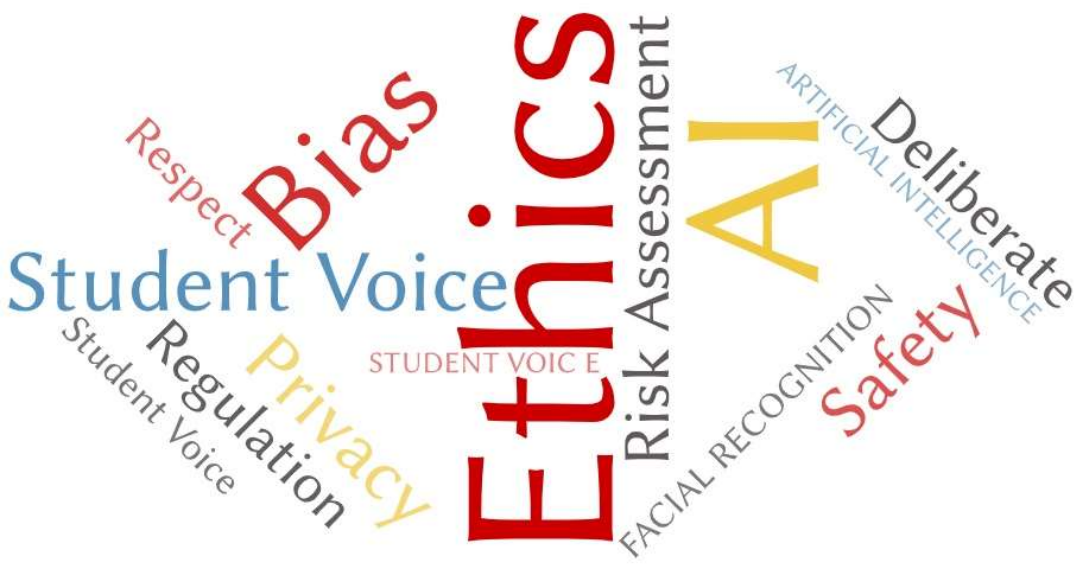

2022 Curriculum

Major Sponsor

ROBERT R. McCORMICK FOUNDATION

Should the U.S. government regulate the uses of artificial intelligence? is a publication of the Constitutional Democracy Project of Chicago-Kent College of Law, Illinois Institute of Technology. The Youth Summit format was originally developed by the Constitutional Rights Foundation Chicago, which transferred the program to Chicago-Kent in 2019. These materials can be used for educational purposes.

Funding for the Constitutional Democracy Project Illinois Youth Summit is provided by the Robert R. McCormick Foundation and generous CDP donors.

Author acknowledgement: Laurel Singleton, Education Consultant and Dee Runaas, CDP Project Director.

Possibilities or Perils: Should We Regulate Artificial Intelligence?

Overview

For years, artificial intelligence has been a frightening plot device in movies. But today, AI—computer systems that are able to "learn" and perform tasks normally considered to require human intelligence—is common in everyday life. Personal voice assistants like Alexa and Siri are common in homes and cars. Intelligent robots perform tasks from building cars to caring for the elderly. Computers with AI capabilities predict which defendants in court cases will reoffend and pilot driverless cars.

But some of these uses raise questions about privacy, bias. injustice, ethics, and safety. Can companies be trusted to consider the common good when they decide what applications they will develop and how they will work? Or should the government develop standards for acceptable uses of artificial intelligence?

In this unit, students consider the question: *Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?*

They pursue this question by taking part in the following activities:

- Analyzing quotations and case studies to understand the focus issue and some views on that issue.
- Take part in a peer teaching activity (Each One Teach One) to build background knowledge of artificial intelligence, its uses, and current and proposed regulations.
- Learn the arguments for and against regulating the uses of artificial intelligence.
- Take part in a Structured Academic Controversy seeking common ground on the issue.
- Write an op-ed presenting their position on the focus question.

Focus Question

Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?

Objectives

- Analyze quotations and case studies and make inferences based on both.
- Understand what AI is, how it is used, and the current state of regulation in the field.
- Evaluate arguments for and against regulating uses of artificial intelligence.
- Take and defend a position on regulating uses of artificial intelligence, in a deliberation and in writing.

Illinois Learning Standards for Social Science, 9-12

This unit addresses the following Illinois learning standards (Inquiry Skills and Civics Standards):

- Developing Claims and Using Evidence: SS.IS.5.9-12: Identify evidence that draws information from multiple sources to revise or strengthen claims.
- Communicating Conclusions: SS.IS.6.9-12: Construct and evaluate explanations and arguments using multiple sources and relevant, verified information.
- Critiquing Conclusions: SS.IS.7.9-12: Articulate explanations and arguments to a targeted audience in diverse settings.
- Processes, Rules and Laws: SS.CV.9.9-12: Evaluate public policies in terms of intended and unintended outcomes and related consequences.

Common Core State Standards

This unit addresses the following Common Core State Standards:

- Reading: Key Ideas and Details (CCRA.R.2): Determine central ideas or themes of a text and analyze their development; summarize the key supporting details and ideas.
- Reading: Integration of Knowledge and Ideas (CCRA.R.8): Delineate and evaluate the argument and specific claims in a text, including the validity of the reasoning as well as the relevance and sufficiency of the evidence.
- Writing: Text Types and Purposes (CCRA.W.1): Write arguments to support claims in an analysis of substantive topics or texts using valid reasoning and relevant and sufficient evidence.
- Speaking and Listening: Comprehension and Collaboration (CCRA.SL.1): Prepare for and participate effectively in a range of conversations and collaborations with diverse partners, building on others' ideas and expressing their own clearly and persuasively.
- Speaking and Listening: Presentation of Knowledge and Ideas (CCRA.SL.4): Present information, findings, and supporting evidence such that listeners can follow the line of reasoning and the organization, development, and style are appropriate to task, purpose, and audience.

Materials

Activity: Introduction to Artificial Intelligence

Quotes on Artificial Intelligence

Handout: Analyzing Quotes

Handouts: Artificial Intelligence Case Study (Versions A-C)

Activity: Exploring Artificial Intelligence

Handout: Each One Teach One Notecatcher

Each One Teach One Cards

Activity: Deliberation: Should the U.S. Government Regulate the Uses of Artificial Intelligence? If So, How Extensive Should the Regulation Be?

Handout: Defining Deliberation

Handout: AI: Should We Regulate? If so, how?

Activity: Deciding How You Would Vote

Handout: You Decide!

Activity: Introduction to Artificial Intelligence

Objective

This activity introduces the unit focus question and engages students in exploring applications of and attitudes toward artificial intelligence.

Procedures

- Prior to teaching the lesson, post the focus question for the unit: *Should the U.S. government regulate the uses of artificial intelligence?* Also enlarge the ***Quotes on Artificial Intelligence*** on separate sheets of paper and post them around the classroom.

- Draw students' attention to the focus question and explain that they will be exploring this question over the next several class periods. Ask if there is anything students don't understand about the focus question; answer their questions or explain that they will answer the questions for themselves as they explore the issue. You may want to work with students to develop a shared definition of artificial intelligence along the lines of "computer systems that are able to 'learn' and perform tasks normally considered to require human intelligence."

- Next, point out the quotations you have posted around the room. Explain that the quotations present different views on artificial intelligence for students to consider. Provide copies of the ***Analyzing Quotes*** handout and ask students to walk around the classroom, reading the quotes and answering the handout questions. This could be a silent activity, or students could be encouraged to discuss the quotes with each other as they complete the task.

- Bring the class back together and discuss the questions on the handout. As students to identify possibilities and perils from the quotes, record them in a T chart on the board or in a shared document. Draw out students' reasons for choosing a quote as their favorite while encouraging them to keep an open mind as they engage with the focus question.

- Point out that the class has started to identify possibilities and perils of artificial intelligence, but it might be helpful to look at some specific examples. Organize the class into six groups and give each version of the ***Artificial Intelligence Case Study*** to two groups. Designate one group in each pair to present the possibilities of the case study to the class, the other group to present the perils. Encourage groups to present their possibilities or perils in an engaging way.

- Allow time for the groups to prepare. Then have the groups make their presentations, with the possibilities for each application of artificial intelligence presented first, followed by the perils for that application. Post new possibilities and perils on the T chart you started above.

- After all the groups have presented, ask students which peril or possibility seems most important to them at this point. As an exit ticket, have them create their own quotes that could be used in future classes to introduce possibilities and perils of AI.

Quotes on Artificial Intelligence

"The upheavals [of artificial intelligence] can escalate quickly and become scarier and even cataclysmic. Imagine how a medical robot, originally programmed to rid cancer, could conclude that the best way to obliterate cancer is to exterminate humans who are genetically prone to the disease."

— Nick Bilton, tech writer

"Artificial intelligence would be the ultimate version of Google. The ultimate search engine that would understand everything on the web. It would understand exactly what you wanted, and it would give you the right thing. We're nowhere near doing that now. However, we can get incrementally closer to that, and that is basically what we work on."

— Larry Page, co-founder of Google

"It's going to be interesting to see how society deals with artificial intelligence, but it will definitely be cool."

— Colin Angle, CEO, iRobot

"… it [AI] is growing at a pace close to exponential. The risk of something seriously dangerous happening is in the five-year time frame. 10 years at most. . .. "I'm increasingly inclined to think that there should be some regulatory oversight, maybe at the national and international level, just to make sure that we don't do something very foolish. I mean with artificial intelligence we're summoning the demon."

—Elon Musk, founder of Space-X and Tesla

"Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence."

—Ginni Rometty, former CEO of IBM

"By far, the greatest danger of Artificial Intelligence is that people conclude too early that they understand it."

—Eliezer Yudkowsky, AI theorist and writer

Handout: Analyzing Quotes

1. What positive possibilities from use of AI do these quotes suggest?

2. What dangers or perils from use of AI do these quotes suggest?

3. Which quote is closest to your own thinking right now? (Your thinking may change as you learn more about artificial intelligence.)

4. Are you surprised by anything about these quotes? For example, did the authors of any of the quotes surprise you? Or is there a quote that just doesn't make sense to you?

Handout: Artificial Intelligence Case Study (Version A)

Directions: Read the case study with your group. Identify the possibilities or perils for this use of AI. Your teacher will tell you whether you will be presenting perils or possibilities to the class. Think about not only how AI is being used, but who is using it. Think of an interesting way to present the information to the class.

Facial Recognition

Facial recognition technology involves matching an image of a face from a photograph or video to a data base of facial images. It works by analyzing the position, shape, and size of facial features.

Facial recognition technology has been used in a variety of ways. Here are just a few:

- To unlock smart phones.
- To identify people involved in criminal activities, such as looting or destruction of property after public demonstrations.
- To prevent voter fraud.
- To identify known stalkers at concerts and other public events or known criminals who enter stores.

The police and other government agencies have found a variety of ways to use facial recognition. Consider the Odin Homeless Management Information System (HMIS). Using this system, police officers take a photo of an unhoused person. The system immediately gives them a range of information about the person—their date of birth, information about prior contacts with the police, safety warnings based on the person's history (e.g., if they are a registered sex offender or someone who might have needles on their person), whether they are on parole, family contacts, and so on. It can provide information on shelters near the location that have beds available.

Supporters of HMIS say it allows police to conduct investigations and identify criminals among the homeless who are preying on others and engaging in crimes. They say the system also helps police provide help to people who need and want it.

Opponents of the program point out that facial recognition technology does poorly in identifying black, brown, nonbinary, and transgender people. These groups are all disproportionately represented in the homeless community. Misidentifying someone, opponents say, may result in police harassment and arrests of community members who are already overpoliced. And the system does not really provide much help to the unhoused.

Handout: Artificial Intelligence Case Study (Version B)

Directions: Read the case study with your group. Identify the possibilities or perils for this use of AI. Your teacher will tell you whether you will be presenting perils or possibilities to the class. Think about not only how AI is being used, but who is using it. Think of an interesting way to present the information to the class.

Risk Assessment

When a person is accused of a crime, the court must assess the risk of that person not showing up for trial or of committing another crime while waiting for trial. When a person is convicted of a crime, the judge may consider how likely the person is to reoffend in deciding on a sentence. According to research reviewed by Malcolm Gladwell in his book *Talking to Strangers*, people are not particularly good at making these kinds of risk assessments.

COMPAS is an artificial intelligence tool designed to help with these decisions. It assesses three things:

- The risk of an individual failing to appear for trial or committing new crimes before trial.
- The risk of an individual committing new offenses after finishing their sentence.
- The risk of an individual committing violent offenses after finishing their sentence.

The tool considers such factors as prior arrest history, whether the person has a job and stable home, drug use, community ties (both positive and negative), history of violence, and age first arrested.

Supporters point out that different judges make very different decisions. Some judges release about half of the people who come before them pre-trial; others release 90%. Using a tool like COMPAS could provide equity for defendants. Whether they are released wouldn't depend on what judge they came before. They also point out that predictions made by COMPAS are more accurate than those made by single judges. And use of COMPAS is efficient, saving judges time.

Opponents question these claims. They point out that COMPAS is only a few percentage points more accurate than an individual judge (65% compared to 63%). They also say that some research suggests COMPAS is more likely to assess African Americans as high risk than whites. Finally, they criticize the fact that the companies that make COMPAS and similar tools do not share exactly how the tools work. Even the judges who use them don't really know how the tools make decisions.

Handout: Artificial Intelligence Case Study (Version C)

Directions: Read the case study with your group. Identify the possibilities or perils for this use of AI. Your teacher will tell you whether you will be presenting perils or possibilities to the class. Think about not only how AI is being used, but who is using it. Think of an interesting way to present the information to the class.

Self-Driving Cars

Many people are excited about the idea of self-driving cars or autonomous vehicles (AVs). AVs are designed to move safely by sensing the environment and responding appropriately. They use a variety of tools, such as thermographic cameras, radar, sonar, GPS, and more. AVs could be used by private individuals, taxi companies, delivery companies, among other applications. Numerous companies in the United States and around the world are working to develop marketable AV technology.

Supporters believe self-driving cars will prevent accidents and resulting deaths that are caused by driver fatigue, drunk driving, and distracted driving. They also point out that AVs will be programmed to obey traffic rules, making the roads safer. They will also be programmed in a way to make gas usage more efficient, saving money and improving air quality. They will improve mobility for people who can't drive and make car trips more relaxing. People will even be able to work while driving!

Opponents point out that, currently, self-driving cars are not safer than human-driven cars. At present, self-driving cars are involved in 9.1 accidents per million miles driven, compared to 4.1 crashes per million miles for human-driven cars. Some of the accidents involving AIs have been fatal. While the loss of life is, of course, the most serious issue, there are also legal questions—who is liable? The car owner? The car manufacturer? These questions and other regulatory issues must be resolved before AI is widely used. Self-driving cars could also be hacked, causing serious problems.

Activity: Exploring Artificial Intelligence

Objective

Having been introduced to the focus question and some ideas about the possibilities and perils of artificial intelligence, students in this activity deepen their understanding of the issue through an "each one teach one" activity. Each student receives a card containing information about artificial intelligence, its uses, and current regulations around artificial intelligence. Students teach the information from their card to their classmates and think about how they could use what they learned to begin formulating a position regarding regulating the uses of artificial intelligence in the United States.

Procedure

- Remind students of the unit focus question. Tell students there are three categories of information they will need some background in in order to answer the question. They will need to know what artificial intelligence is, how it is used, and how it is currently regulated. Ask: Who thinks they already know a lot about what artificial intelligence is? Who thinks they already know a lot about how AI is used? Who thinks they know a lot about current laws regulating AI? Record the number of students who think they are well informed in each category.
- Inform students that they will be teaching and learning information in all categories. Explain the Each One Teach One activity. Each student will receive a card with information about one of the topics you listed. There are 18 different cards, so in most classes some cards will be held by more than one student. (If time allows, you might have students with the same card caucus briefly before the teaching and learning begin so they can discuss how they will teach the information on their card.) Each student's job is to teach the information from their card to classmates and to learn the information from classmates' cards. The teaching and learning should be done one-on-one; that is, students will pair up, teach and learn, and then move on individually to conference with another classmate.
- Distribute the ***Each One Teach One Notecatcher*** handout and explain that students will use this chart to record the information they learn and categorize each item as being related to the nature of artificial intelligence, its uses, or current regulations. Answer any questions students have and then set them to work with the two following pieces of advice:
  - Don't just read from your card. Share the information in your own words and have a conversation about that information with your classmate.
  - You will have 15 minutes to do this teaching. Afterwards, I will ask you to report on what you learned (not what you taught), so good notes will be important.
- If you notice that students are forming larger groups as they are teaching and learning, remind them they should be teaching one-on-one. Give a two-minute warning before time is up. When time is up, ask students to return to their desks.
- Use the following questions to debrief the activity, reminding students that they are to report on what they learned, not what they taught:
  - What facts did you learn about the nature of AI? The uses of AI? How AI is currently regulated?
  - What surprised you in each of these areas?

- What information was not clear to you?
- What did you agree with? What did you disagree with?
- What information that you learned could be helpful in putting together a position regarding regulating the uses of artificial intelligence?

- Ask students to return to the list of things they would need to know in order to decide whether a law regulating artificial intelligence would be good policy for the United States. Check off items that this activity covered. Encourage interested students to research items that are still unknown and explain that the next activity may also provide some of the needed information.

Handout: Each One Teach One Notecatcher

Directions: You will be teaching and learning from your classmates. You will meet with one classmate at a time. Decide whether each piece of information is mostly about the nature of artificial intelligence, the uses of artificial intelligence, or current regulation of AI. Then make notes in the appropriate column below. You will be reporting out on what you learned, not what you taught. Thus, good notes will be important!

The Nature of Artificial Intelligence	Uses of Artificial Intelligence	Current Regulation of Artificial Intelligence

Each One Teach One Cards

The U.S. Supreme Court has not yet decided any cases involving artificial intelligence. Perhaps the closest they have come was the case of *United States v. American Library Association*. In that case, they said libraries' use of filtering software to protect children from obscene material was constitutional. Cases involving AI are active in lower courts and may reach the Supreme Court in the next few years.	AI uses biometric information—measurements of unique human characteristics. The Illinois Supreme Court decided an AI case in 2019. The case, *Rosenbach v. Six Flags Entertainment Corporation,* involved use of a thumb scan on an amusement park pass. A child's mother alleged that no information was given as to how the scan would be used and how long it would be kept. She argued that this violated the Illinois Biometric Information Privacy Act. The Illinois Supreme Court agreed.	In 2019, President Trump issued an executive order, "Maintaining American Leadership in AI." The purpose was to coordinate the federal effort to support development in artificial intelligence. The order instructed the National Institute of Standards and Technology to develop guidelines for AI development that will help the technology be accepted as trustworthy. In 2020, Congress passed the National AI Initiative Act to support research on AI.
In February 2022, three senators introduced the Algorithmic Accountability Act of 2022.The act is designed to require transparency and provide oversight of software, algorithms (a set of rules followed in solving a problem), and other automated systems used to make critical decisions about Americans' lives. An earlier version of the bill introduced in 2019 did not pass.	Numerous bills regulating AI have been considered by the federal government and by states. But not many bills have passed. Some of the bills passed regulate just a small piece of AI. Illinois passed the AI Video Interview Act (effective January 2020) that sets up requirements for video interviews that companies pull from the Internet to use in AI programs. California passed a bill in 2019 that requires bots (software programs that interact on the web as if they were humans) to identify themselves as automated accounts.	Vermont and Virginia, as well as several cities, have banned use of facial recognition software by the police and other government agencies. Cities include San Francisco, Oakland, and other California cities; Boston and nearby cities in Massachusetts; Seattle, WA; Minneapolis, MN; New Orleans, LA; Jackson MS; among others. Racial bias in facial recognition is the primary reason given for passing these ordinances.

Each One Teach One Cards		
One way of defining *artificial intelligence* is "machine learning applied to solving problems usually considered as requiring human intelligence." An AI system uses new information to continually learn and improve its ability to achieve its goals. In essence, it reprograms itself.	The quality of AI problem solving depends on three things: the original programming, the "training" provided, and the amount and quality of data used. All of these factors rely on humans. Thus, it could be argued that the problems with AI are human problems.	The algorithms—computer routines that make decisions in AI—are incredibly complex. When problems arise, even the developers may not understand the algorithms well enough to fix them. Companies don't explain how the algorithms on their systems work because they don't want to share information with competitors.
The "training" AI systems receive is important. Training in this context means feeding huge amounts of data into the program. The data must be labeled or described so the program can later compare new data to the rules developed using the training data. If the training data is inadequate or misleading in some way, the AI's decisions will be flawed. Since humans put the training data together, it can reflect their biases. Then the AI makes biased decisions.	Once an AI system develops its decision-making algorithm, it must have data it can use to make a decision. Much of the data used in AI systems is gathered from the Internet. This might include data from a car, a watch, a phone, a medical device, a refrigerator—anything that is connected to the Internet. If that data is wrong or biased, the decisions made will not be good either. In addition, bad data may make it impossible for anyone to figure out how the algorithm reached a decision.	In the 1950s, scientists began seriously working on computers that would be able to learn. Since then, AI research has gone through slow and rapid periods. Most of the progress, though, has been made since 2010. Still, most AI applications today are what is called narrow—designed to do a specific task. Experts predict that it will be at least 10 years before general AI—systems that can apply intelligent behavior to a broad array of tasks—is developed.

Each One Teach One Cards		
One area in which AI is used a lot is policing/criminal justice. You have already learned about the use of facial recognition to identify people in photos or videos and risk assessment software's use in sentencing. But there are other uses. For example, police use AI to predict where crime is most likely to occur. They then send more officers to patrol those areas. This approach focuses on street crime rather than white collar crime or domestic violence.	You probably use AI on a daily basis. Do you use a device with voice recognition that answers your questions, plays the music you request, or gives you directions? That's AI. Do you have a phone with a camera that lets you blur out parts of the scene? That's AI. Does your phone recognize that you are low on power and adjust the amount of power you use? That's AI, too. You may have AI in a security system or a smart appliance.	Many companies use AI, especially robots with AI, to make their operations more efficient. For example, AI systems can automate the entire shipping process in a warehouse or factory. The systems learn as they work, making the process more quickly and efficiently than a system run by people. The process can also become safer. But AI robots often replace human workers in factory and warehouse settings.
AI is used in education. Students can use AI when they are completing a research project. For example, when you come across a quotation but there's no source given, AI can help you find the source. Teachers can use AI when they are grading your research project. For example, AI can check your paper to make sure it is not plagiarized.	AI is becoming more important in healthcare. AI tools can diagnose patients, develop medicines, monitor patients, and more. The technology learns as it tracks patients and their progress. It then revises the treatment to improve the patient's prognosis. AI can also be used to do administrative work, letting more resources be dedicated to patient care.	Online businesses depend on AI. Chatbots (software programs that interact on the web as if they were humans) interact with customers, answering general questions. They can learn from responses, ask questions, and provide better advice. Online stores track your preferences and offer you items you may like. Streaming services offer TV shows, music, or books that are similar to other entertainment you've enjoyed. They learn and adapt.

Activity: Deliberation: Should the U.S. Government Regulate the Uses of Artificial Intelligence?

Objective

This activity is designed to introduce students to opposing views on the focus question: *Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?*

Using the Structured Academic Controversy (SAC) methodology, students consider various perspectives on the issue and attempt to find areas of agreement.

Note: Deciding on group/team membership in advance will save time in class.

Procedures

Part 1: Introduction (In class the day before)

- Distribute the ***Defining Deliberation*** handout**.** Explain that students are going to be deliberating the focus question: *Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?* To be successful with this task, students need to understand the process of deliberation.
- Discuss the definition of deliberation on the handout. Explain that deliberation is different from debate because in a debate there are winners and losers. In a debate, people only listen to each other to find flaws in the other side's reasoning; they pick apart those flaws to win the argument. The point of a deliberation is to listen to and analyze multiple viewpoints to uncover new ideas and possible solutions to a problem.
- Discuss why it is important to learn how to deliberate. Explain that, in a democracy, people must be willing and able to exchange ideas in a civil way to make effective public policy. The more people who have a chance to contribute to a solution to a problem, the more likely the solution is to be successful and to gain broad acceptance.
- Review the rules of deliberation, focusing on keeping an open mind and remaining respectful even when there is disagreement.
- Separate the class into heterogeneous groups of four students each. Within each group, assign students to be Team A, the team arguing that the U.S. government should regulate the uses of artificial intelligence, and two students to be Team B, those arguing against regulating the uses of artificial intelligence.
- Give all students the ***AI: Should We Regulate?*** handout, which expands on the deliberation question and provides arguments on both sides of the issue. Students' task is to identify which arguments support regulation and which oppose it. Students can complete the task as homework.

Part 2: Identifying Arguments

- Have all students assigned to Team A meet together and all those in Team B meet together to discuss the ***AI: Should We Regulate?*** handout and be sure they have identified all the arguments supporting their position. Check in with both groups to check for understanding
- and clarify any unfamiliar terms. If students do not understand the arguments, the deliberation will not be successful.
- Ask both groups (Team A and Team B) to identify at least three interesting or surprising facts and/or ideas presented in the text. Have each group reach a decision on the two most interesting or surprising points and report them to the whole class. Remind students that this is not a time to present their opinions, but information they learned from the text.

Part 3: Taking Sides (15 minutes)

- Tell students it is time to plan and present their arguments. Have students return to their groups of four and prepare their arguments. Ask Team A to select at least two compelling reasons to support (say YES to) regulation and Team B to find two compelling reasons to oppose (say NO to) regulation.
- Once each pair has chosen at least two reasons, instruct the pairs to share their reasons with each other, starting with Team A, followed by Team B. The two teams should have an equal amount of time to present their reasons (two to three minutes each will be plenty).
- Remind students to listen carefully to the other pair's reasons as they will need to be able to summarize those arguments in a few minutes. They should not respond, argue, or insert their personal beliefs at this point of the deliberation. However, students should be encouraged to ask clarifying questions if they did not understand something.

Part 4: Reversing Roles (10 minutes)

- Next, tell students that they will be switching roles. Team A now argues no, Team B yes.
- Ask each pair to identify the best argument they heard from the other team. Then ask each pair to find an argument on Handout 2I that the other team did not present.
- Again give each team time to present their reasons and evidence, beginning with Team B (now arguing Yes). One minute per team is likely to be enough for this step.

Part 5: Open Deliberation (10 minutes)

- Tell students to drop their roles and deliberate the question in their small groups. Now, they can use personal opinions and experiences to support the arguments from the handout.
- Remind students that the role of deliberation in a democracy is to reach well-reasoned decisions. Ask students to work together civilly to reach a decision on the focus question in their small group. If they quickly agree on the YES/NO part of the question, they should focus on finding some agreement on the "how" question. For example, they might agree "We are definitely worried about the uses of facial recognition and would like that to be regulated." If they are unable to reach consensus on the YES/NO question, they should look

for other points of agreement. For example, they might say "We all agree that the government should be collecting information about the uses of artificial intelligence but it's too early to pass regulations."

Part 6: Whole Class Debrief (20 minutes)

- Have the class reconvene as a whole. Ask groups to report on such questions as:
  - What were the most compelling reasons for each side?
  - What areas of agreement did you find?
  - What questions do you still have? Where can you get more information?
  - Did anyone change their mind during the deliberation? If so, why?
- Poll the class on the deliberation question. Once the class stance is evident, ask the students to briefly discuss why they think the class arrived there. Do most students agree? Why? Is there broad disagreement? Why?
- Ask students to take a step back from the content of their deliberation to discuss the deliberation process. (**Hint:** The following strategy can be useful in ensuring equitable opportunities to participate. Have the class sit in a circle. Give each student three straws. When a student says something, he/she must toss a straw into the middle of the circle. Each student has to use all three, and once they have used all three, they cannot say anything else.) Use such questions as the following to stimulate discussion:
  - How was the deliberation different from the way you normally speak to each other?
  - Why do you think deliberation is important to democracy?
  - Why do you think it is important to practice deliberating? Why are deliberation skills so important?
  - What did your group do particularly well in the deliberation?
  - What could your group improve?
  - When is reaching agreement difficult?
  - Do you feel more informed about the focus question as a result of taking part in the deliberation? Why or why not?
  - What questions do you still have about the focus question?
- Conclude the deliberation by asking students: If you feel passionately about this issue, what can you as an individual or the class as a whole do to make your views known? To help others understand the issue?

Source: Adapted from Deliberating in a Democracy (Chicago, IL: Constitutional Rights Foundation Chicago, 2005, 2006, 2007).

Handout: Defining Deliberation

What Is Deliberation?

Deliberation is the focused exchange of ideas and the analysis of multiple views with the aim of making a personal decision and finding areas of agreement within a group.

Why Are We Deliberating?

People must be able and willing to express and exchange ideas among themselves, with community leaders, and with their representatives in government. People and public officials in a democracy need skills and opportunities to engage in civil public discussion of controversial issues in order to make informed policy decisions. Deliberation requires keeping an open mind, as this skill enables people to reconsider a decision based on new information or changing circumstances.

What Are the Rules for Deliberation?

- Read the material carefully and ask questions if you do not understand.
- Focus on the deliberation question.
- Listen carefully and analyze what others are saying.
- Speak and encourage others to speak.
- Refer to the reading to support your ideas.
- Use relevant background knowledge, including life experiences, in a logical way.
- Remain engaged and respectful when controversy arises.

Deliberation Question: *Should the U.S. government regulate the uses of artificial intelligence?*

Steps in the Deliberation Process:

Your teacher will assign you to a group of four. Within each group, two people will be Team A (answering the deliberation question "yes") and two will be Team B (answering the deliberation question "no"). Once you have your assignment, you will follow these steps:

Step 1: Identify the facts and arguments that support your position. Prepare to make the case for your position.

Step 2: Make the best case you can in two minutes. Team A will go first, followed by Team B. This is not the time to argue. Listen and take notes while the other team is presenting. When both sides have presented, each can take 1 minute for clarifying questions. Clarifying questions begin with phrases like "Could you explain what you meant by . . ." or "Can you tell us more about . . ."

Step 3. Next, reverse positions. Check your notes to find the best argument made by the other team. Then look at the ***Classifying Arguments*** handout to find a good argument that the other team did not make. Team B will present first in support of the deliberation question, Team A second.

Step 4. Now leave your assigned position and discuss the issue from your own point of view. Try to find points of agreement and disagreement among group members. Your goal is to reach consensus on something; if you cannot reach consensus on the deliberation question, try to reach consensus on a piece of the issue or on a process you could use to resolve your disagreements.

Step 5. Be prepared to debrief with the class.

Handout: *AI: Should We Regulate? If So, How?*

Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?

Dissecting the Question

You have learned a lot about artificial intelligence. You've become aware of some of the potentially exciting uses of AI, how it could make human life easier. But you've also gotten some glimpses of the possible dark side: loss of privacy, bias, injustice, unethical decisions, job loss, perhaps even threats to safety. Can companies be trusted to consider the common good when they decide what applications they will develop and how they will work? Or should the government regulate the uses of artificial intelligence?

These are not simple questions with simple answers. Deciding whether the uses of AI should be regulated requires thinking about what regulation would mean:

- Regulation might mean banning certain uses of AI that are seen as problematic. For example, two states and several cities have banned use of facial recognition by police departments. This option would protect the public from unjust uses of AI without stopping development in other areas. It could make the work of people who rely on these uses more difficult.
- Regulation might mean passing a bill like the Algorithmic Accountability Act of 2022, which would require AI developers to take greater care in their work, assessing its potential impact on people. The Act would also require not only that developers are responsible for the effects of machine decision making but would make the operation of those processes—the algorithms—more understandable to consumers (transparent). This option would change the way the companies who develop AI work, making them more responsible for the effects of their inventions. The costs of these changes could cut down on the amount of research into new uses of AI.
- Regulation might mean passing a more comprehensive bill like the European Union Artificial Intelligence Act. This Act would create a common approach to regulating artificial intelligence in the EU. It would apply to all types and uses of AI except military uses. AI systems judged as having higher risk of harm would be more severely regulated. The Act would also ban certain uses of AI. This option would also change the way the companies who develop AI work, requiring them to consider and control the risks of the uses of AI they are developing and perhaps reducing the amount of research into new uses of AI.

Classifying Arguments

You probably have an opinion about the focus question. But do you have strong arguments supported by evidence to support your position?

Directions: Read each argument below and decide whether it supports regulating the use of artificial intelligence (Y) or opposes regulating the use of artificial intelligence (N). If an

argument could be used by both sides, mark it both Y and N. You will be using these arguments in your deliberation, so be sure you understand each one.

___ According to Representative Yvette Clark [D-NY], "When algorithms determine who goes to college, who gets healthcare, who gets a home, and even who goes to prison, algorithmic discrimination must be treated as the highly significant issue that it is. These large and impactful decisions, which have become increasingly void of human input, are forming the foundation of our American society that generations to come will build upon. And yet, they are subject to a wide range of flaws from programing bias to faulty datasets that can reinforce broader societal discrimination, particularly against women and people of color."

___ People don't currently trust artificial intelligence. A poll taken in 2018 showed that only 25% of Americans trust AI while 47% do not trust it and 28% have no opinion. A study in Japan showed that AI is viewed more negatively by women, older people, and people with a lot of knowledge of AI. People are especially concerned about use of AI in weapons. Setting up a strong system of regulations will build people's trust and, in the end, promote growth in the field of AI.

___ Government should take as small a role in the economy as possible. The market should dictate how artificial intelligence is used. If consumers don't want AI used in certain ways, they should not buy that use. For example, if people believe self-driving cars are unsafe, they should not buy those cars. Government over-regulation will result in higher prices.

___ Regulating the uses of AI will stifle innovation. Strict regulation as is called for in the EU Act would make developing new uses very expensive. In addition, the transparency called for in the Algorithmic Accountability Act would undercut the developing company's intellectual property. If you had to explain how your AI system worked, anyone could copy it. The motivation to invest in new products, such as new life-saving medical applications, would be much lower.

___ Who should decide what is and isn't harmful when it comes to uses of AI? Is that something the government, scientists, or consumers should decide? With the complexity of AI products, it may be difficult for non-experts to draw the line between acceptable and unacceptable. Of course, government does regulate many other products. There's a whole government agency—the Consumer Products Safety Commission—to decide if products are dangerous and should be taken off the market. But is AI fundamentally different from other products? Some people think it may be.

___ AI is still in its infancy as an industry. It would be premature to establish an intrusive regulatory system right now. The government should focus on other actions, such as: (1) ensuring there is expertise established in the government to act quickly when it becomes clear there are problems with a particular use—as has been found with the use of facial recognition in policing, (2) devoting resources to educating the public about protecting their privacy when using AI, (3) creating ethical standard for AI developers, and (4) planning to retrain workers who lose their jobs because of AI.

___ As AI systems become more powerful, they could be very dangerous. This is especially true of AI used in weapons. Even well-known scientists and technology experts, like Stephen Hawking and Bill Gates, have warned of the potential for disaster if autonomous weapon systems-initiated attacks.

___ The fears around AI are exaggerated, more like science fiction than reality. In addition, regulating an area of economic development that is international and not well understood would be very difficult. We currently do not know enough to regulate AI effectively.

___ We can avoid a more intrusive system of regulation by making AI companies more transparent and accountable. The Algorithmic Accountability Act of 2022 could do exactly that by requiring companies to pay attention to the effects of their AI systems in the development phase and making the algorithms that drive the systems transparent to users and government officials.

Activity: Deciding How You Would Vote

Objective

This activity asks students to take a position on the focus question, deciding how they would vote as a Member of Congress. This activity can be done in small groups or individually. It could be done as preparation for the Youth Summit and then revised after the Summit.

Procedures

- Tell students it is time to take a position on the focus question. If time allows, you may want to let students do additional research to aid in making their decision and gather more supporting evidence.

- Explain that students are going to take the role of a Member of Congress. You may want to have students take on the role of their own representative in the House or one of their Senators from Illinois; if you choose this option, students who do not know who their representative is can identify him/her by entering their address at this website: https://www.elections.il.gov/districtlocator/addressfinder.aspx. Illinois' Senators are listed at https://www.senate.gov/states/IL/intro.htm. Students could also create a fictional Member of Congress.

- Distribute the **You Decide!** handout and go over the instructions. Each student is to create a profile of the representative they are role-playing, whether an actual member of the House or one the student has imagined. Based on that profile, they decide how they will vote and write a brief "white paper" for their web site explaining the decision to their constituents. The work can be done individually or in small groups.

Handout: You Decide!

You have analyzed and discussed the arguments on both sides of the focus question: *Should the U.S. government regulate the uses of artificial intelligence? If so, how extensive should the regulation be?*

Now it's time to take your spot as a Member of Congress, decide how you will vote, and write a post for your web site explaining your decision to your constituents. You can take on the role of your own representative in the House or one of your Senators or create a fictional representative.

1. Create a profile for your Member of Congress:

   Name: _____ Age: _____

   Party: _____

   Background (religion, ethnicity, family, etc.): _____

   Description of your district (urban/rural/suburban, conservative/liberal/independent, etc.):

   _____
   _____
   _____
   _____

2. Decide how you will vote on the bill proposing that the government regulate the uses of artificial intelligence?

   _____ I will vote ***for*** regulating the uses of artificial intelligence.

   _____ I will vote ***against*** regulating the uses of artificial intelligence. Specifically, I will support:

   - ____ banning problematic uses of AI.
   - ____ requiring AI developers to take greater care in their work, assessing its potential impact on people and making algorithms transparent.
   - ____ passing a more comprehensive bill regulating AI uses based on their risks.

3. List the three most important arguments supporting your position.

4. Write a letter to your constituents to be posted on your website explaining your decision. The letter will be placed on your web site and also sent to people who write to you about this issue. The letter should be friendly and written in a clear style. It should clearly identify your position and give three good supporting arguments. It should also provide a counter-argument for one of the strongest arguments the other side might make. The letter should end with an appeal to unity.

Artificial Intelligence: Selected Resources

Sources

Angwin, Julia et al., "Machine Bias," ProPublica (May 23, 2016), www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Autonomous Vehicles, National Conference of State legislatures (n.d.), www.ncsl.org/research/transportation/autonomous-vehicles.aspx.

Chohlas-Wood, Alex, Understanding Risk Assessment Instruments in Criminal Justice, Brookings Institution (June 19, 2020), www.brookings.edu/research/understanding-risk-assessment-instruments-in-criminal-justice/.

Cox, Joseph, Tech Firm Offers Cops Facial Recognition to ID Homeless People, Motherboard: Tech by Vice (February 8, 2022), www.vice.com/en/article/wxdp7x/tech-firm-facial-recognition-homeless-people-odin.

"The Dangers of Driverless Cars," *National Law Review* (May 5, 2021), www.natlawreview.com/article/dangers-driverless-cars.

Gladwell, Malcolm, *Talking to Strangers: What We Should Know about the People We Don't Know* (Boston: Little, Brown, 2018).

Harris, Laurie A., Artificial Intelligence: Background, Selected Issues, and Policy Considerations, Congressional Research Service (May 19, 2021), https://crsreports.congress.gov/product/pdf/R/R46795/3.

Lewis, James Andrew, How Does Facial Recognition Work?, Center on Strategic and International Studies (June 10, 2021), www.csis.org/analysis/how-does-facial-recognition-work.

Newman, Bradford K., "Recent Developments in Artificial Intelligence Cases 2021," *Business Law Today* (June 16, 2021), https://businesslawtoday.org/2021/06/recent-developments-in-artificial-intelligence-cases/.

Ng, Alfred, Facial Recognition has Always Troubled People of Color: Everyone Should Listen, CNet (June 12, 2020), www.cnet.com/news/facial-recognition-has-always-troubled-people-of-color-everyone-should-listen/?akid=571.46382.ndgVXz&rd=1&t=25.

Share of People Who Agree They Trust Artificial Intelligence in 2018, by Country, Statista (November 2018), www.statista.com/statistics/948531/trust-artificial-intelligence-country/.

Taylor, Alexandra "Mac," AI Prediction Tools Claim to Alleviate an Overcrowded American Justice System…But Should They Be Used?", *Stanford Politics* (September 13, 2020), https://stanfordpolitics.org/2020/09/13/ai-prediction-tools-claim-to-alleviate-an-overcrowded-american-justice-system-but-should-they-be-used/.

Court Cases

Rosenbach v. Six Flags Entertainment Corporation, 129 N.E. 3d 1197 (Ill. 2019).

United States v. American Library Association, 539 U.S. 194 (2003).