TELUS Wise®

Understanding artificial intelligence (AI) tools.

Created in partnership with Media Smarts

Grades: 6-12

Duration: 60 to 90 minutes, plus extension activities

Learning goals:

Key concepts/big ideas:

- **Media are constructions:** Algorithms do not "think" like people do; they make connections by finding patterns in data.
- **Media have social and political implications:** Careful steps need to be taken to make sure that algorithms are used fairly and ethically.
- **Digital media have unanticipated audiences:** Algorithms use our personal data to make decisions about us.
- **Digital media experiences are shaped by the tools we use:** What algorithms know (or think they know) about us can affect how we interact with them.
- **Interactions through digital media have real impact:** We need to think critically and act ethically when we're using or interacting with algorithms.

Key questions:

- How do algorithms work?
- How can we tell if something was made with AI?
- How can we make sure that algorithms are used fairly?
- How can we control how algorithms use our personal information?
- What critical questions should we ask when using algorithms?
- How can we use algorithms fairly?

Misconceptions to correct:

- AI works like human intelligence, with agency and intention.
- Algorithms are always fair because they make their decisions based only on data.

Preparation:

Prepare to distribute the following:

- Handout: Understanding AI tools (page 10)
- Handout: Understanding AI tools presentation notes (page 8)
- Handout: How can I use chat AIs safely and ethically? (page 14, print this hand out double-sided)

Assignment: Understanding AI tools - rules for safe and ethical Chat AIs (page 16)

Procedure

Start by asking students: How many of you have ever used a chat AI tool or "chatbot"? Which ones have you used? (Some examples you can use as prompts: ChatGPT, Character AI, My AI, Alexa, Siri.)

Distribute the handout **Understanding AI tools** and have students complete this poll:

Do you think Chat AI tools...

- Have feelings?
- Can feel hurt or left out?
- Can choose to do or not do things?
- Are smart?
- Always give you correct information?
- Know if they are giving you correct information or not?
- Can think about whether something is right or wrong?

Go through the items in the poll and ask students if they answered "yes", "no" or "maybe". For each item, ask why they think that.

Point out that many kids, even ones their age, are unsure about these things. In one study, a quarter of eleven-year-olds thought that "Alexa" (a smart speaker that uses AI) has feelings and can think for itself, and another third thought these were "maybe" true.

Tell students that AI tools definitely seem intelligent (thus the term "artificial intelligence") and may even seem to have feelings. They even make decisions – e.g. they may refuse to do something you tell them to, and if you ask them a question they will "choose" what to include or leave out in the answer.

What is an AI tool anyway?

AI can be defined as any information technology capable of solving complex problems that would normally be attributed to humans or animals.

AI is not a single technology and can include many technologies. **This lesson plan is specifically about machine learning algorithms.**

Ask students what they think the term "AI" means. Give them a few minutes to write their answer on the handout, and ask if any would like to share their thoughts.

Explain that most of the things we call "AI" are machine learning algorithms, and ask students if they can think of any other examples of algorithms. (Prompts: search engines, recommended content on YouTube or Netflix, ads served on social media, image/content generation algorithms like MidJourney or Dall-E, autocomplete.)

Tell students that an algorithm is a series of instructions that guide the decision-making process. Using algorithms makes decisions more consistent, and it can make decisions more fair, because you're applying the same standards to everyone – but if you're not careful, it can also make it harder to realize when an algorithm **isn't** fair

Consideration 1: Is AI always fair?

Read through **Consideration 1** in the handout. Explain that because there are always more people who need kidney transplants than there are kidneys available, it can be very hard to decide who gets a kidney.

Refer to the equation KP = YL in the handout and explain that this simple algorithm was created to decide who should get priority: KP means "kidney priority" and YL means "years of life." In other words, a person is more likely to get a transplant over someone else if they are more likely to live more years than the other person. Give students a few minutes to write their answer to these questions:

- Does this seem reasonable to you?
- Does it seem fair? Why or why not?

If students don't raise the following questions, ask:

- How accurate a measurement of "YL" can you have? How can you know, or guess, how many more years someone will live if they get the kidney?
- The older someone is, the fewer extra years of life a kidney is likely to give them. Is it fair to give younger people a better chance at getting a kidney?
- Different groups of people have different life expectancies: groups that are discriminated against or those that have less access to medical care, have shorter life expectancies. Is it fair for people who already have a higher life expectancy to have first priority for kidney transplants?

At the end of the discussion, write on the board the principle that **algorithms can reinforce real-world biases.**

Explain that the people designing the kidney transplant algorithm realized the problem and decided to change it. (The real algorithm considers many more variables. For example, it also considers how healthy the donor kidney is – so that younger patients are not necessarily more likely to get a kidney, but they will get a healthier kidney.)

Consideration 2: How is AI different from simple algorithms?

Now read through **Consideration 2** and explain that almost all of the algorithms in their lives – search engines, recommendation algorithms that YouTube and TikTok use, sorting algorithms like the Instagram feed (which decides which posts to show you and in which order) – are not simple algorithms like the one you were just discussing, but machine learning algorithms or AI. Give them a few minutes to write their answer to the question: "How might this make it harder to be sure that AI makes fair decisions?" Then, take it up in class.

Make sure the following points come up:

- AI is trained on real-world data, so just like the kidney algorithm they will reflect real-world biases if we aren't careful.
- Because we don't know why they made a decision, we might not realize they're making decisions for biased reasons.

At the end of the discussion, write on the board the principle that **we don't always know how AI makes decisions.**

Consideration 3: Does AI always make good choices?

Now read through **Consideration 3.** Give students a few minutes to write down their answer and then take it up in class.

Make sure the following points come up:

- As time goes on, you will get more and more ice cream and less and less broccoli. Eventually, you will get only ice cream and no broccoli. If there are more than two elements of your dinner, eventually you will only get your favourite one and nothing else.
- You might like getting less broccoli and more ice cream, but it's not necessarily good for you!

Explain to students that most of the algorithms in our lives work the same way: they are optimized to give you things that will keep you using the app or platform that uses them.

Now ask students: What are some examples of "ice cream" (i.e., things that are not good for you, either in large amounts or at all) that different kinds of AI might give you? What are some examples of "broccoli" that AI should give you even if you don't always want them? Tell students to think about this question in the context of their social media feeds.

"Ice cream" examples (things you see often in your feed but regret reading or watching):

- "Clickbait" or "fake news" (things that seem really exciting but are untrue or misleading)
- Things that make you angry or upset
- Things that make you feel bad about yourself
- Things that make you keep watching
- Things that are dangerous (e.g. content promoting unsafe viral challenges)

"Broccoli" examples (things you have to look for in your feed, but make you feel better after reading or watching them):

- News
- Emergency alerts
- Reliable information
- Things that make you feel good about yourself
- Things that encourage you to do something healthy or meaningful

At the end of the discussion, write on the board the principle that **AI gives you what it thinks you want, not what's good for you.**

Consideration 4: Where does AI get the data it needs to make choices?

Now read through **Consideration 4.** Give students a few minutes to write down their answers and then discuss as a class.

Make sure the following points come up:

- The AI might match people's search, viewing or shopping history with more or less risky behaviours, such as buying a skateboard, or watching skateboarding videos.
- This might be good for people who were seen as being lower-risk, because they would probably pay less, but people who were seen as being higher risk would pay more.
- The data might let the AI make more accurate decisions than ones based on very broad categories like gender. But the connections it makes might be inaccurate in two ways:
  - data might be misinterpreted (e.g. if you bought the skateboard as a gift for someone else)
  - the connection between one kind of risk and another might be inaccurate (e.g. you might like risky sports but still drive very carefully)

Now explain to students that nearly all AI does work that way. They either use:

- Data that they collected about you (e.g., the prompts you enter into a chat AI, or information your computer automatically sends like your location and Internet Protocol address)
- Information the company bought (either from another company or from data brokers that buy data from lots of places and put it together into profiles)
- Or information shared between different parts of the same company (e.g., Google can use your YouTube views to decide what search results to show you, and vice versa, because they're owned by the same company)

Now write on the board the principle that **AI makes decisions based on what it thinks it knows about you.**

Consideration 5: How are Chat AIs similar to other machine learning algorithms?

Read through **Consideration 5** with the class.

It's alright if you don't feel confident answering more detailed questions about the content: as one expert put it, "no one on earth fully understands the inner workings of large language models (LLMs)." If you do want a more detailed explanation, read the article Large language models, explained with a minimum of math and jargon

What is an AI tool anyway?

Have students return to the "What is an AI tool anyway?" section in the Understanding AI tools handout and revisit the definition they wrote. Ask them if their definition has changed based on the class discussion; if it has, have them write the new definition.

Now remind students of the examples of chat AIs that were listed at the beginning of the lesson, and then the four principles that have come up through the discussion:

1. Algorithms can reinforce real-world biases;
2. We don't always know how AI makes decisions;
3. AI gives you what it thinks you want, not what's good for you; and
4. AI makes decisions based on what it thinks it knows about you.

Now divide students into groups of 3-4 and assign each group one of those principles. (Each principle can be assigned to more than one group depending on class size.)

Have each group discuss how their principle applies to chat AIs. As the groups talk, circulate and prompt them with the following questions:

Group 1: What kinds of biases might be found in the texts chat AIs were trained on? (Think about social media posts, or in old books.) How might that affect its responses?

Group 2: When designers don't fully know how AI makes decisions, it's harder to make safeguards or "guardrails" to prevent some of the other problems. How do you deal with that?

Group 3: Are there topics or information you might want chat AI to tell you, but that wouldn't be good for you? What could go wrong if you asked a Chat AI for advice, or told them about a personal problem?

Group 4: How do you think chat AIs get information about you? Do you think that's fair? Do you think it's accurate?

Now have the groups that were given the same principle come together and share their discussions. Have them identify which points were raised in both groups and discuss those that were only raised in one group. Groups don't have to agree on all points, but they do need to identify how they disagree.

Next, have the combined groups make a short presentation to the class where they explain the ways in which their principle applies to chat AIs. They should also make one suggestion for how people who make or use chat AIs could address or improve the issues they identified.

- If possible, have the groups write their points on chart paper that can be hung on the wall or blackboard after their presentation.
- Before the groups present, distribute the handout Understanding AI tools: Presentation notes (page 10). Make sure students know they will need to draw on each group's presentations in the next part of the lesson and encourage them to use the handout to take notes.

Understanding AI tools: Rules for safe and ethical chat AIs

Distribute the assignment sheet Understanding AI tools: Rules for safe and ethical chat AIs (page 16) and go through it with students.

Explain that they will think of at least two rules, codes or guidelines to address each of the four principles, for a total of at least eight.

- These will be aimed at people who use or design chat AIs or use them in their services (such as a social network or search engine that builds a chat AI into their platform).
- For each rule, they should also write a few sentences explaining why they feel it's important.

The rules and explanations should show the students' knowledge and understanding about algorithms and chat AIs gained from the Understanding AI handout and the class discussion, as well as the groups' presentations about how the four principles of AI apply to chat AIs.

Reflection: How can I use chat AIs safely and ethically?

Distribute the reflection exercise: How can I use chat AIs safely and ethically? (page 14) Have students read the prompt and spend a few minutes writing a reflection in response.

When they're done, have them turn the paper over and read the tips there. How well do those address the things they raised in their reflective writing?

Extension: Using chat AIs in the classroom

The previous activities were designed to be delivered without access to chat AI services, due to the privacy concerns associated with them. If you or your students have access to a chat AI and you would like to integrate it into the lesson, you may use some or all of the following extension activities.

Because chat AI models are always in development, it's recommended that you test out all of the exercises before using them in class.

Make sure you're familiar with your board or district's policies on using commercial online platforms in the classroom.

1. **Making AI work for you:** Point out to students that for a long time, students were not allowed to use calculators in math class at any grade level. Now, older students are allowed to use them, so they can spend more time on complex math instead of doing arithmetic.

    - Ask: What are the things in our class that are like arithmetic – the things you have to do that you have already learned? How could we use chat AIs like a calculator, to give more time to focus on critical thinking and other higher-level tasks?
    - Have students experiment with using a chat AI for the tasks you identified.
    - Have students develop a process and a code of ethics for using chat AIs in school work.

2. **Fact-checking AIs:** Ask a chat AI a question about a topic you have recently covered in class, then have students read and critique the result. (Alternatively, you can have students ask a question about a topic they have personal knowledge or expertise in.)

    - How accurate was the result? (Large language models sometimes "hallucinate" details.)
    - How well did the result show an understanding of the topic? (Large language models often

over- or under-emphasize different details compared to what a human expert would consider important. Because of how they were trained, they also frequently express widely-held misconceptions as facts.)

- Do different ways of phrasing the prompt lead to more or less accurate results? (For example, using using the phrase "write an essay on" will sometimes lead to less accurate results than asking a question.)

3. **Summary judgment:** Have students read a well-known short story, then have the class collaborate on a summary of the key points in the plot and an interpretation of its meaning. Next, have the chat AI do the same thing, and have students compare the results. How similar or different were they?

4. **Papers, please:** Ask a chat AI a question about an academic topic and ask that it give references to scholarly sources. Then ask students to look up the sources. How many are real? How many had the correct author(s) and publication date?

5. **Leveling up:** Pick a concept or idea that the class has studied, then tell the chat AI to explain it at three levels: for grade-school children, for high-school students, and for teachers who are preparing to teach the subject. Then have them compare the results. How were they different? How effective do you think each explanation would be for that audience?

Extension: Exploring algorithms

These activities expand the focus of the main lesson to look at other ways that AI affects our lives. Make sure you're familiar with your board or district's policies on using commercial online platforms in the classroom.

1. **How to train the TikTok algorithm:** TikTok, which MediaSmarts' research has found is the second-most popular online platform among Canadian youth, relies heavily on its algorithmically-sorted "For You" feed.

Ask students:

- If you use TikTok, how much do you like what the For You feed shows you?
- Do you know how you "trained" the For You feed?
- Here's how to train it intentionally:
  - If you like a video and want to see more like it, watch it all the way to the end.
  - If you want to see more videos about a specific topic, search for that topic and then watch some relevant videos.
  - If you don't like a video and would rather not see more like it, hold your finger on the video for a few seconds and then tap "Not Interested."
  - If you want the algorithm to forget everything it knows about you and start over, Profile -> Menu -> Settings and Privacy -> Content Preferences and then tap Refresh Your For You Feed.
- How close is that to how you thought it worked?
- Are there videos you're going to intentionally encourage or discourage the algorithm to show you?
- If you decided to refresh your For You feed, what would you do differently now that you know how to train the algorithm?

2. **Picturing bias:** The way that algorithms replicate (and sometimes intensify) bias may be easier to understand if it's presented in visual form. The tool Diffusion Bias Explorer allows you to compare images created by different AI image generators and explore the ways that different prompts, such as adjectives and professions, reflect bias in their training set.

Understanding AI tools.

Presentation notes

Use this sheet to record notes from the group presentations (including your own group's presentation.)
You will need these for the next part of the lesson, so ask questions if you miss or don't understand anything.

Algorithms can reinforce real-world biases:

Notes

We don't always know how AI makes decisions:

Notes

AI gives you what it thinks you want, not what's good for you:

Notes

AI makes decisions based on what it thinks it know about you:

Notes

Understanding AI tools.

Poll:

Answer the questions below with "Yes," "No" or "Maybe."

Do you think chat AI tools...

- Have feelings? _____
- Can feel hurt or left out? _____
- Can choose to do or not do things? _____
- Are smart? _____
- Always give you correct info? _____
- Know if they're giving you correct info or not? _____
- Can think about whether something is right or wrong? _____

What is an "AI" tool anyway? Write what you think the answer is here:

Now that we've talked about it in class, has your definition changed? If yes, write your new answer here:

Digging deeper

Consideration #1: Are algorithms always fair?

Kidney transplants can give people many extra years of life. But because there are always more people who need transplants than there are kidneys available, somebody has to decide who gets a kidney and who doesn't. It's important that this decision be fair, so in 2007 doctors came up with this algorithm to decide it:

KP = YL, where:

- KP means kidney priority: where someone is in the line to get a kidney. The higher your KP, the more likely you are to get one of the kidneys available.
- YL means years of life: how much longer you're likely to live if you get the kidney.

So KP=YL means that you are more likely to get a kidney if it will give you 20 years of life vs. 10 years of life.

Does this seem reasonable to you?

Does it seem fair? Why or why not?

Consideration #2: How is AI different from simple algorithms?

That means they're given a goal (like deciding who should get a kidney), then given a data set (like a list of people who got kidney transplants in the past, and how long they lived after), and then look for patterns in that data that they then use to make a decision. We don't always know why an AI tool made a particular decision. Even the people who designed it might not know!

How might this make it harder to be sure that AI makes fair decisions?

Consideration #3: Does AI always make good choices?

Another thing that makes AI different from simple algorithms is that they can change themselves. Imagine that an AI tool decided what you were going to eat, and was optimized to give you the dinner you liked the most. It might do this by watching how much you ate of the different things it fed you, and adjusting what you got based on that: if you didn't eat all of your broccoli one day, it would give you less broccoli the next, and if you ate all of your ice cream, it might give you more ice cream the next day.

Would that be good for you? Why or why not?

Consideration #4: Where does AI get the data it needs to make choices?

Imagine that you were designing an algorithm for an insurance company that would help determine who was more likely to get in an accident. (The more likely you are to get in an accident, the more you have to pay for insurance.)

Like the kidney algorithm, insurance companies already set prices based on demographic data: men pay more than women for car insurance, on average, because they have more car accidents.

But what if your AI knew what videos people watched online, what online purchases they had made, or what they had searched for on Google, and could match that with whether or not people had car accidents?

What kinds of patterns do you think it would find? How might that be good or bad for different people?

How accurate do you think that data would be? What might make it more or less accurate?

Do you think it would be fair for AI to use this data to make decisions about likelihood to get into an accident? Why or why not?

Consideration #5: How are chat AIs similar to other machine-learning algorithms?

Chat AIs are a kind of AI called a large language model.

What does that mean? Let's go through the three words in reverse order:

- **Model –** Like other machine learning algorithms, most of chat AIs' capabilities aren't programmed, but instead come from being trained on large amounts of writing. They find patterns in these to create a model of how language works.
- **Language –** Unlike earlier algorithms, such as autocomplete tools, they can read and write sentences, paragraphs and even full articles fluently. They do this mostly by looking at how similar or different words are in different ways or "dimensions."
  - For example, if we were to consider just two dimensions, roundness and redness, an apple and a fire truck might be very far apart on roundness but close together on redness, while a baseball would be close to the apple in terms of roundness but far away in redness.
  - Chat AIs can make guesses by "looking" along different dimensions: if it started at "king" and looking down the gender dimension, it would see "queen," while if it looked down the age dimension, it might see "prince," and looking in both directions might lead to "princess."
  - This lets the AI make better guesses about what words should follow each other, based on other parts of the sentence or paragraph. For example, if you wrote "Frida had a drink of chocolate," a simpler algorithm, like autocomplete, might always suggest that the next word after "chocolate" should be "chip" because that's what most often follows it.

    On the other hand, if you asked a chat AI: "What kind of chocolate did Frida drink?", it might spot the word "drink" and calculate that the word closest to "chocolate" along the liquid dimension is "milk."
- **Large –** ChatGPT was trained on a data set of around 500 billion words. Each word is given a value in up to 96 dimensions, and it does more than 9,000 operations every time it guesses a new word. Different chat AIs were trained on different data, but most were trained on what's called the "open web", such as Wikipedia and Reddit, public posts on social media, publicly accessible websites like news outlets and people's blogs, and digital texts like the public domain books at Project Gutenberg. Most of these texts were "scraped" – read without getting specific permission from their creator to be used in the training set.

How can I use chat AIs safely and ethically?

Now that you've thought about how the people who make chat AIs can make them safer and more ethical, take a few minutes to reflect on how you use them.

What are some things you should not do when using chat AIs? What are some things you **should** do?

What can you do to protect yourself and your privacy when using them?

Write your thoughts here.

Tips for using chat AIs safely and ethically

Read the list below of things to remember when using chat AIs.

How many did you think of in your reflection?

Did you think of things that aren't listed here?

Ethics

- Never present AI-written work as your own.
- Talk to your teacher before using a chat AI in your work.
- If you use AI in your work, show clearly where and how you used it.
- Don't rely on information provided by Chat AI without validating and fact checking; without double-checking for yourself, you can't know how accurate it is.

Safety and privacy

- Think carefully about what you say or write when using chat AIs. Remember that the company can save and read all of your prompts and the information you provide may affect how the AI acts later. Additionally, the things you say now could affect anything from your search results to whether you get offered a job 10 years from now.
- Find out how much data the AI collects about you and what you can do about it. Look it up at https://tosdr.org/ to see a summary of its privacy policy. If it isn't listed there, go to the Privacy Policy and look for a title like "How We Use Your Personal Information." Then look for a title like "Your Choices" or "Your Rights" to see what options you have (if any).
- Limit how much data AI has about you. On an iPhone or an iPad, use the "Ask App Not to Track" setting. You can also use apps like Blokada or DuckDuckGo and browser plug-ins like Privacy Badger or Ghostery to stop apps and websites from collecting data.
- Remember that while talking to a chat AI can feel like talking to a real person, chat AIs are just simulations that guess what you want to hear. Treat them like a toy, not a friend.
- Chat AIs can be a lot of fun, but just like everything else you do online, you have to make sure they aren't taking time away from important things like friends, family, school, sleep and exercise.

Understanding AI tools - Rules for safe and ethical chat AIs.

For this assignment, you will make rules and guidelines that could be used by people who make or use chat AIs. These rules should address the issues that were raised in the group presentations.

Review your presentation notes and then think of **two** rules to address each of the principles we've discussed:

1. Chat AIs can reinforce real-world biases;
2. We don't always know how chat AIs make decisions;
3. Chat AIs give you what they think you want, not what's good for you; and
4. Chat AIs make decisions based on what they think they know about you.

This will make a total of **eight** rules and guidelines.

For each of the rules you come up with, explain why you think it's important.

Your rules and explanations should show the knowledge and understanding you gained from the Understanding AI tools handout and the class discussion, as well as the group presentations.

Principle	Rule	Why is this rule important?
1. Chat AIs can reinforce real-world biases	1.	
	2.	
2. We don't always know how chat AIs make decisions	3.	
	4.	
3. Chat AIs give you what they think you want, not what's good for you	5.	
	6.	
4. Chat AIs make decisions based on what they think they know about you	7.	
	8.	