let's talk science est 1993 CHALLENGE

Engineering and Technology

2024

Thank you to our chapter sponsor:

FONDATION TOYOTA FOUNDATION

What is Artificial Intelligence?

To define **artificial intelligence**, or **AI**, you need to define what intelligence is. You might think it means being smart. But intelligence is far more than doing well on tests.

Intelligence is complicated. Scientists who study intelligence still have many questions to answer. Are animals intelligent? How are imagination and memory linked to intelligence? Does the size of a brain have anything to do with intelligence?

There are many different definitions of AI out there. Most include the idea that AI is the theory and development of computer systems to do tasks that normally need human intelligence. These include things like recognizing images, interpreting speech, and making decisions.

This part of the handbook will help you understand more about AI.

General and Narrow Artificial Intelligence

There are two main kinds of artificial intelligence. These are **Artificial General Intelligence** and **Artificial Narrow Intelligence**. It is important to understand the difference between the two.

Artificial General Intelligence

The goal of Artificial General Intelligence (AGI) is to be equal to human intelligence. Many robotic engineers have managed to create human-like robots. But those robots don't have intelligence equal to humans yet.

AGI is what you usually see in books and movies that take place in a **dystopian** future. Some examples of this are the Terminator and the Matrix movies. This is not quite AGI, but **superintelligence**. Superintelligence is the idea that machines could someday be smarter than humans. But many experts think this is not a threat we should worry about anytime soon.

To develop AGI, scientists look at what we know about human intelligence. They try to understand things like how we learn, or how creativity works. To do this they use things like digital models to copy human behaviours. Brain research can help AI developers. And AGI can help brain researchers understand the brain better!

Artificial Narrow Intelligence

Artificial Narrow intelligence (ANI) is AI we can use right now. Sometimes this is also called weak AI. This lets machines perform actions and make decisions in limited situations. Self-driving cars and virtual assistants are examples of ANI.

A Brief History of Artificial Intelligence

The term "artificial intelligence" was first used in 1955. Then, they thought general AI would be possible in a few years. But the technology needed to research AI was not ready.

People using AI for personalized shopping recommendations and to pay for products (Source: elenabs via iStockphoto).

Here are a few of the milestones in artificial intelligence:

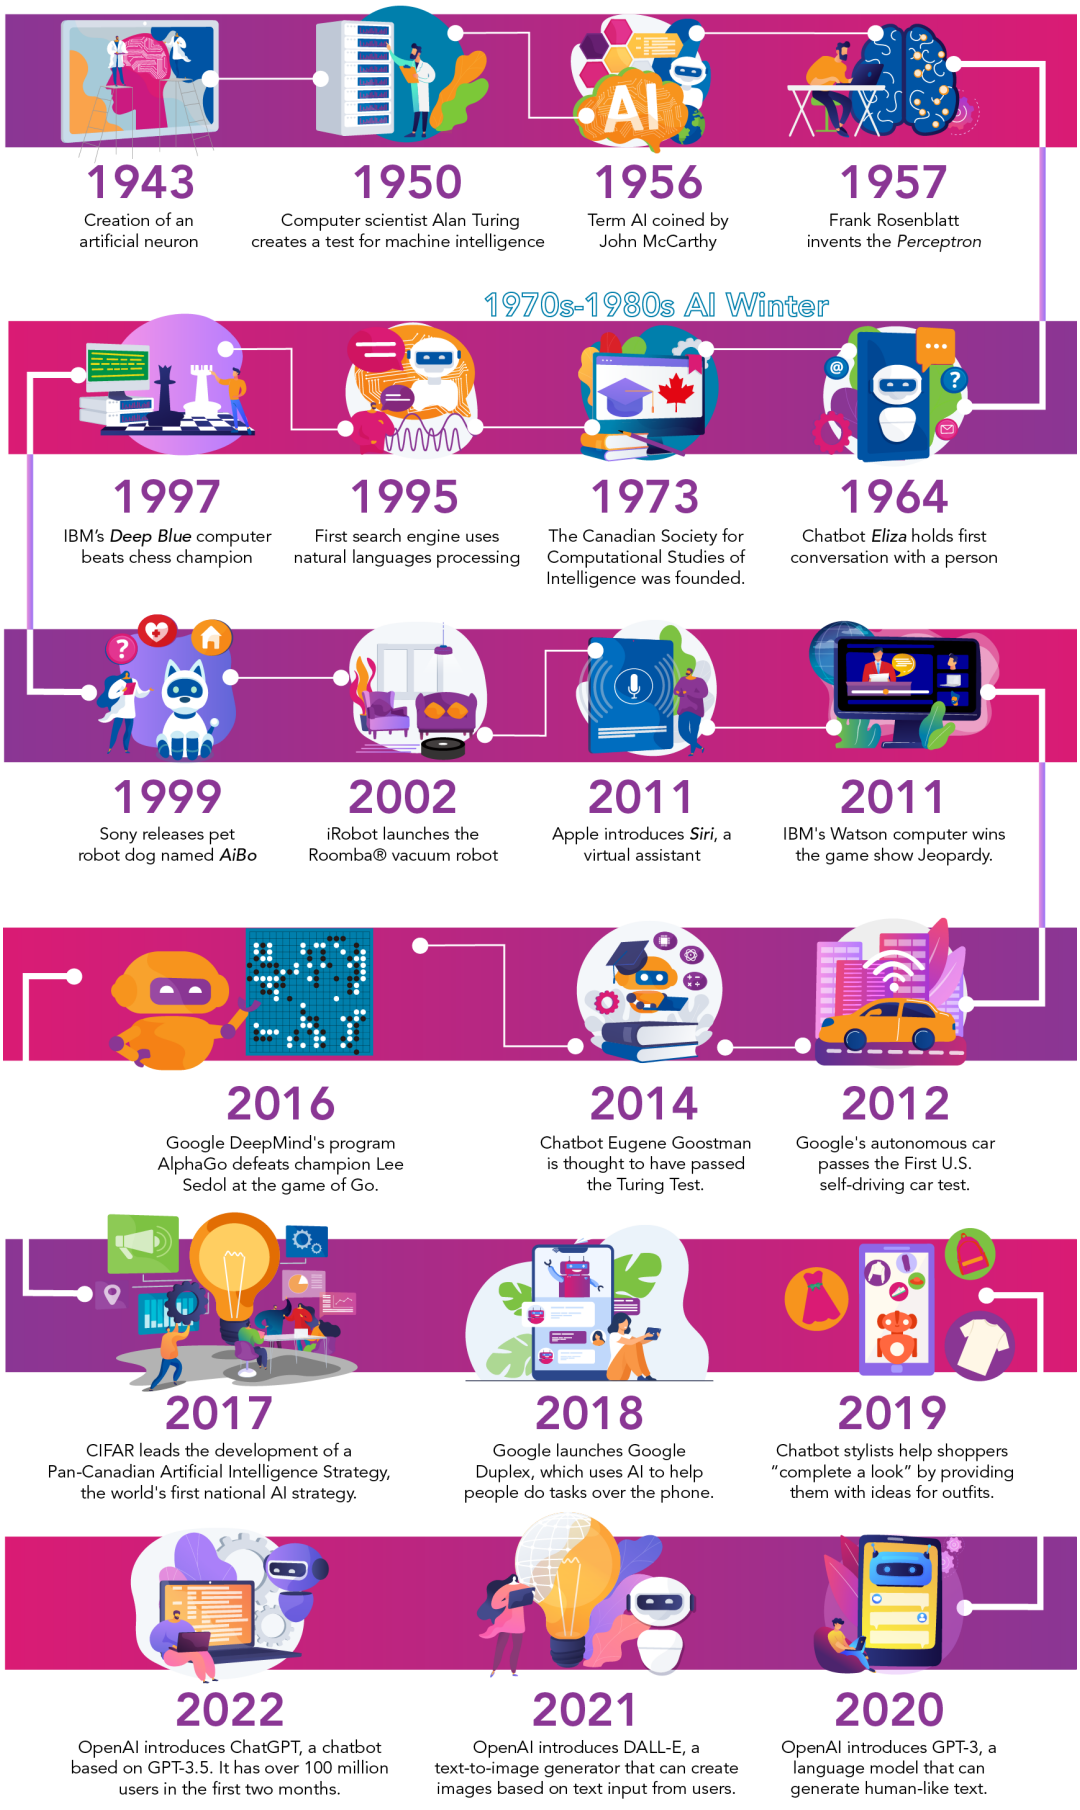

AI timeline (Let's Talk Science using an image by Visual Generation via iStockphoto).

Humans and Computers working together

“Human in the loop” AI is a way to combine the strengths of computers and humans. A computer's strength is its computing speed. Humans will never be able to do calculations as fast as computers. Human strengths include creativity and critical thinking. Humans have the ability to look at a problem from different points of view.

This means humans have many important roles in AI. For example, humans monitor self-driving cars. People responsible for hiring can make sure AI hiring tools are fair. And medical experts can help make sure an AI tool is accurate for all patients.

Artificial Intelligence Opportunities

AI is a great tool. It can help us find solutions to many problems. This is true in any area where we use data, which is almost everywhere!
In medicine, AI is being used in research. It can help researchers fight cancer or search for vaccines. Some doctors use AI diagnostic tools to help identify illnesses faster. In agriculture, AI can help farmers detect pests and diseases in plants. It can also help to reduce food waste.

AI can also provide tools to help us protect the environment. AI can use satellite images and data to help scientists make predictions. For example, they can use AI to identify where forest fires could happen, before they start. We can also use AI to predict other natural disasters. Scientists can even use AI to understand and protect endangered species. Or protect ecosystems from invasive species.

Did you know?
In 2018, Canadian researchers won the AM Turing Award. This is sometimes called the "Nobel Prize of Computing." This led to the creation of many AI research centres across Canada. And many people now consider Canada an important centre for AI development.

Artificial Intelligence Concerns

Should we be afraid that artificial intelligence might develop beyond human intelligence? Or that it could become a threat to humans? Experts have been debating these issues for a long time. Some people only see AI’s potential for innovation and problem-solving. Others, like Elon Musk and Bill Gates, think AI could be dangerous. They think humans should be very careful about how they develop and use it. That’s why governments are now working to create regulations for AI. These rules will make sure people develop AI applications with human well-being as a priority.

Jobs

AI is already changing the job market. Robots controlled by AI are doing many things that humans used to do. And people are beginning to worry that they will be replaced by machines. During the first industrial revolution, people worried about the same thing.

Some experts predict that AI will make some jobs obsolete. But they also predict it will create more jobs. Some of them call this the fourth industrial revolution. This change will mainly affect jobs that involve manual labour. But it will also affect office jobs where people make routine decisions. These tasks can be automated with AI systems.

AI has potential for many opportunities including those in healthcare, disaster relief and refugee assistance (Source: Visual Generation via iStockphoto).

Fairness

AI can solve many problems, but it can also create problems. AI is developing so quickly that some problems are only discovered when a technology is already being used. Like all new technologies, developers can't always predict how it could go wrong. This is why it's important for more people to better understand how AI works.

We also need people to follow situations closely when new AI tools are used. This is important because AI is used in more and more areas where it affects people's lives. These include justice, medicine and banking systems.

For example, the medical data used to train AI needs to represent the characteristics of a whole population. These include things like age, gender and ethnicity. If it does not, then the AI cannot make accurate predictions for that population.

It is important that AI tools don't create or increase inequalities. Experts are already working on this problem. It is important that people of all backgrounds have access to education in AI. This will help to make sure equity and diversity are better addressed in the future.

In conclusion...

The field of AI is especially important for young people. Its impact on your lives will only increase. But becoming more familiar with AI does not have to mean taking computer science classes. You can learn about it in many different ways.

We may not know exactly what the future holds, but we can be pretty sure that AI will play an important role. And many scientists in Canada are working hard to figure out how humans and machines can build a better world together. Will you join them?

Let's Talk Science appreciates the contributions of Melissa Valdez Technology Consultant from AI & Quantum for revisions to this part of the handbook.

What is Machine Learning?

From Big Data to Machine Learning

Data has a story to tell - if you know how to look for it! In the past, people analyzed data to find meaningful **patterns** and **trends**. This is fairly easy with a small amount of data. But when there is a lot of data, it can be very difficult. Using computers can make **data analysis** much easier and faster.

So, how much data is out there? Short answer - a lot! People generate lots of data, especially on the internet. This is often called "**Big Data**." Data scientists and computer engineers have had to develop bigger and better systems to deal with this. Cloud computing is one such system. Developers created it because individual computers could not process large amounts of data on their own.

All this data has changed the relationship between humans and computers. In the past, humans used computers to help organize and represent data. But humans still had to make sense of it. Now, machines are figuring out how to understand and explain data that is difficult for humans to make sense of. We call this Machine Learning (ML).

Misconception Alert
The terms machine learning and artificial intelligence are often used together. But they do not mean the same thing. Machine learning is one type of artificial intelligence.

Where is machine learning used?

Machine learning can be used in any situation that involves large amounts of data. This is almost anywhere! Businesses were among the first to use machine learning. This is mainly because they had the money to invest in the technology, which was very expensive. Now, machine learning is less expensive and easier to access. In fact, people share many machine learning programs online as open source files.

Did you know?
Open source is the term for when people freely share their code online. People can change open source code to suit their needs, and to help improve it.

Machine learning is used in autonomous vehicles, medical research, and marketing campaigns. But did you know that machine learning is also used in sports, restaurants, and even in writing rap lyrics? The possible uses of machine learning are endless. And we are just beginning to use this powerful tool.

	Supervised	Unsupervised	Reinforcement (self-supervised)
Key aspect	Computers are provided with data that has been labelled by people	Computers look for patterns in data that are difficult for people to label	Computer creates its own data and supervises its learning
Mostly used when	We have a well-defined and time-consuming task	We want to find new patterns in data	We have a goal but we may not know the best way to achieve it

How does machine learning work?

There are many types of machine learning. Each has its own strengths and weaknesses. Below and to the left is a table that summarizes the different types.

Supervised Machine Learning

The first type of machine learning is supervised machine learning. As you might guess, here people supervise the computer as it is trained. Let's look at an example.

Say you were designing a self-driving car. You would want the car to know the difference between different types of road signs. In this case, you would need to design a machine learning **model**. This is a program that can find patterns and make decisions based on a set of data.

To help the model figure out a pattern to correctly identify stop signs, people need to teach it what a stop sign looks like. First the model is given **labelled** pictures of stop signs and other road signs. Each label is either “stop” or “non-stop”. In computer language, we say that these pictures and their labels are the input data. The label name, “stop” or “non-stop”, is what we want the model to be able to identify later. This is the output data. To do this, the model may learn to look for the shape of an octagon.

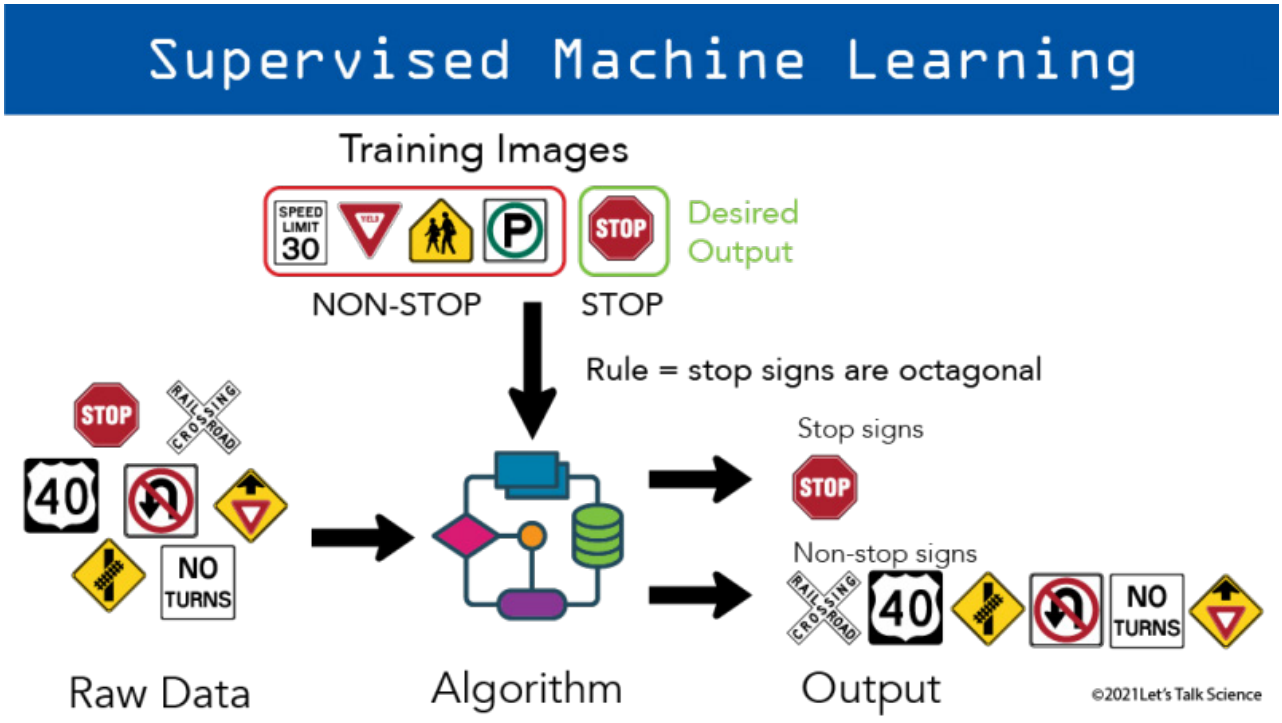

Unsupervised machine learning (Let’s Talk Science using images from Pop_jop via iStockphoto and Infografx via iStockphoto).

Once there is a model, engineers test it using new data. The model should be able to identify stop sign images it has never seen before. If it can't, then it needs more training. Sound familiar? It's a lot like learning new things yourself!

You may wonder how good machine learning is at this. We call the percentage of correct solutions accuracy. For example, if the machine learning identified stop signs 98 times out of 100, the accuracy would be 98%. But accuracy is an imperfect measure of a model's performance on its own. Imagine a system that is given 98 images of cats and 2 images of dogs. It can recognize cats with a 98% accuracy if it just guesses "cat" every time! To get meaningful results, the system needs to be given the same amount of input data for each object.

When people use images to train a computer to see, it is called computer vision. Computer vision helps self-driving cars. It also helps computers read handwriting.

Try this!
Identifying a stop sign may seem easy to you, but it's hard for a computer. Look at an object around you and try to describe it using only shapes and colours.

The previous examples used machine learning to classify things. But supervised machine learning can also be used to make predictions. For example, a company could use machine learning to predict how long people will stay with the company. Machine learning could analyze input data like an employee's education and years of experience. The output data would be how many years they stayed at the company. Once machine learning creates a model from existing employee data, it can be used to predict how long new employees will stay.

The biggest disadvantage of supervised machine learning is that it needs good labelled data to train on. A study about data labelling found that up to 80% of machine learning engineers' time was spent making sure the labels were correct.

Unsupervised Machine Learning

Unsupervised machine learning is used to find patterns in data sets that are difficult to label. One example of this type of data is human speech. All people's speech sounds different. So it is hard to tell a computer exactly how a word should sound. Unsupervised machine learning can be used to analyze spoken words.

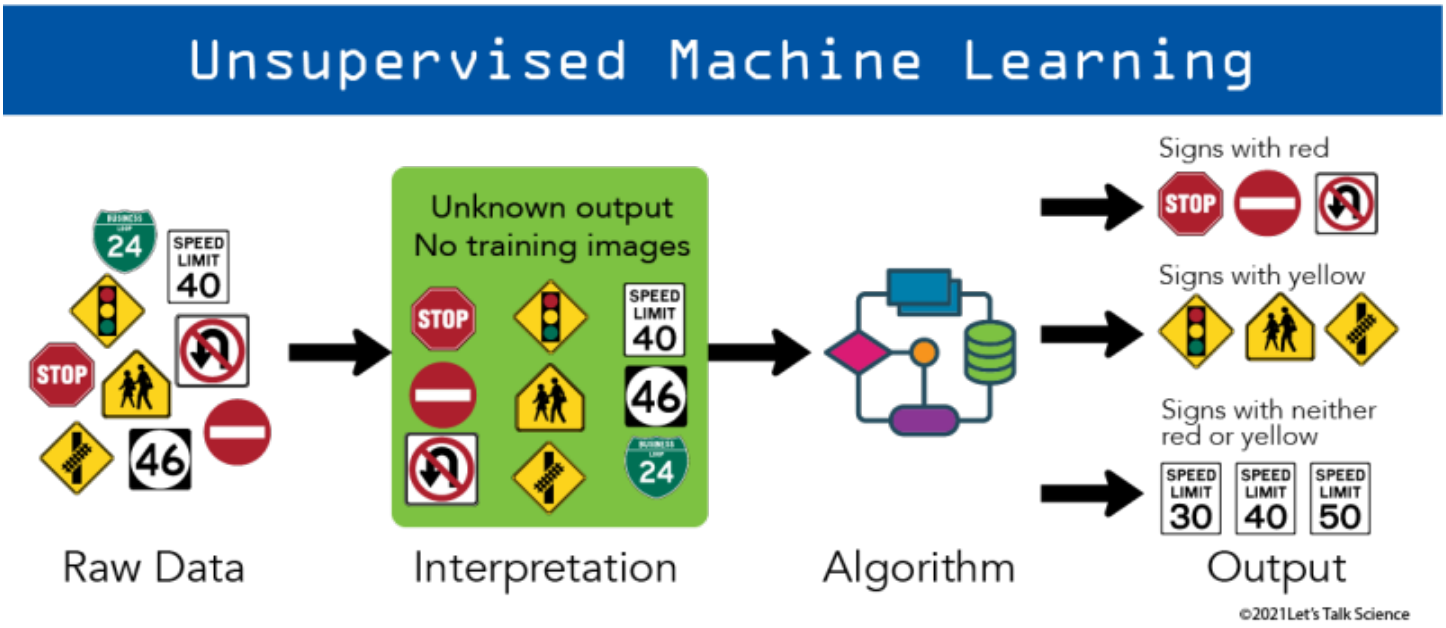

Unsupervised machine learning (Let's Talk Science using images from Pop_jop via iStockphoto and Infografx via iStockphoto).

Another example is in medicine. To help treat or cure a disease, scientists may want to figure out if it involves specific genes. Genes carry the information that makes you who you are. Each of your cells contains 25 000 to 35 000 genes. Researchers could use unsupervised machine learning to look for similarities in the genes of people who have the disease.

To see how unsupervised machine learning works, let's go back to the self-driving car example. Here, the system is not given labelled training images and the output is not pre-defined. Instead, the system takes the raw data and looks for patterns itself. Once it has found a pattern, a human can develop a model that includes these patterns. This can lead to models for specific tasks, like avoiding other vehicles on the road or dealing with traffic slowdowns.

Reinforcement or Self-supervised Machine Learning

The third type of machine learning is reinforcement or self-supervised machine learning. In this type of machine learning, the machine learns by trial and error. It creates large sets of data by running various scenarios. Then it evaluates which strategies produced the best results. This is similar to how you learn to play a video game. You try different things. Eventually you learn which strategies work best. Like a human, the machine adapts and improves its strategies based on its experiences. But it also gets some fine-tuning from actual humans.

Imagine a robotic arm that uses computer vision to toss different objects. The goal is for the robot to correctly throw the object in the bin as quickly as possible.

Again, this might seem simple to us. But a robot needs to be able to consider many things to complete this task. First, it needs to locate and pick up an object. It needs to consider its gripping force, the force of the throw, and the weight and shape of the object. This requires understanding many physics principles. It would be pretty tricky to create a program that considers all these things. But this is a good task for self-supervised machine learning. Robots using self-supervised machine learning would be useful in places like recycling plants, where robots sort materials.

Another example of machine learning is when computers beat humans at games. Computers can use self-supervised machine learning to find the fastest way to win. Two computers can even play against each other using self-supervised machine learning. In one case, two machines discovered a flaw in a game.

There are many areas where self-supervised machine learning is used to improve systems. One area is computer security or **cybersecurity**. Keeping data safe is very important when the data is confidential. This includes data used by banks and the government. To test cybersecurity systems, self-supervised machine learning can pretend to be hackers. This lets people find the flaws in a system before an actual hacker does!

So how could self-supervised machine learning be used in our self-driving car example? Such a system could use virtual driving simulations to learn the safest action to take when the camera sees a red octagon.

Is there a best type of machine learning?

Choosing the right type of machine learning will depend on the type of problem that needs to be solved. For some problems, combining types can help get even better results. Remember the example above, about genes? We could use unsupervised machine learning to identify a gene that may play a role in a disease. We could then use this information to train a supervised machine learning model to detect that gene, based on the patterns identified in the unsupervised model. We could test its accuracy using input data from people with and without the disease.

Further Thoughts About Machine Learning

We know a lot about how our brains work, but some things are still a mystery. This is a lot like machine learning. It is great that the machines do what we want. But it is not enough. We also want to understand how they work.

Without knowing how machines make decisions, how can we know if their decisions are fair and ethical? This is especially true if people use machine learning with data about people.

Being able to explain how machine learning works is called transparency or explainable artificial intelligence.

You may also be thinking, if machines can learn, will humans still be needed? The answer is yes! A machine learning model is only as good as its data. This is why we need many human experts to make sure it is accurate and appropriate. We also need qualified people to make sure these technologies are used wisely and fairly. Many people are working on this now, but even more will be needed in the future.

Let's Talk Science appreciates the contributions of Melissa Valdez, Technology Consultant from AI & Quantum, for revisions to this part of the handbook.

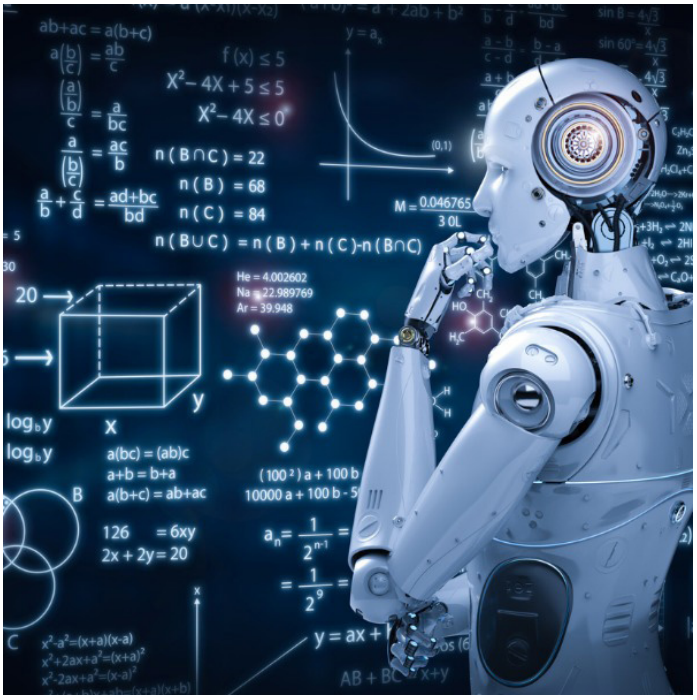

Robot thinking about a problem (PhonlamaiPhoto, iStockphoto)

AI and Human-Computer Communication

Does Artificial Intelligence (AI) make you think of a robot takeover movie? If so, you might think of AI as a future technology. But it's actually used in many areas of your life right now. In this backgrounder, we will explore some of the tools that AI has made possible. But, before you go further, you need to know what machine learning (ML) is. If you have not read the machine learning backgrounder, you will want to do it now. Most of the **applications (apps)** we will discuss involve machine learning.

Natural Language Processing

Have you ever used a search engine? Or used autocorrect to catch your spelling mistakes? If so, you have used an app that involved Natural **Language Processing (NLP)**. Natural language is the speech and text humans use to communicate with each other.

But humans and computers do not speak the same language. This is one of the biggest problems in human-computer communication. NLP is a way of trying to solve this problem.

As people, we use text, voice and gestures to communicate. Machines do not immediately recognize a lot of this.To help bridge this gap, NLP follows a number of steps.

Let's look at an example of text analysis.

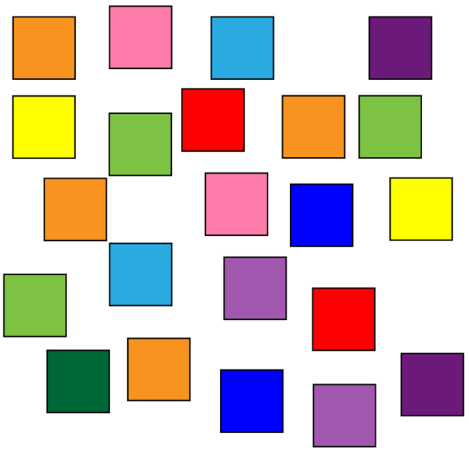

Assortment of text represented by coloured squares (©2023 Let's Talk Science).

Preprocessing

This is the first, and often most important step of NPL. Many of the choices people make in this step affect the accuracy of **models** that are built.

In preprocessing, people ask computers to prepare a text for machine reading. This involves:

1. **Tokenization**, which is splitting the text into individual words, called **tokens**, or small phrases, called **ngrams**
2. Removing punctuation
3. Making all text lowercase
4. Removing **stopwords**, which are common words like 'the', 'is', 'in', etc.

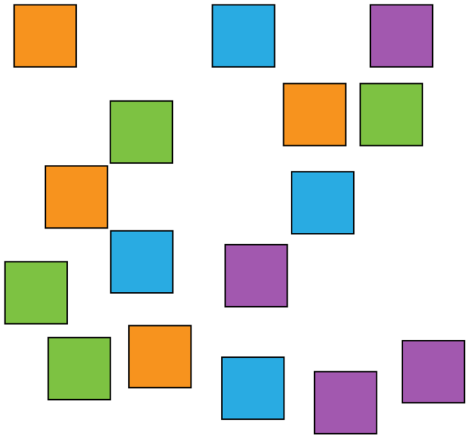

Preprocessing involves removing (squares removed) and changing text (squared recoloured) (©2023 Let's Talk Science)

Here is an example of text before preprocessing:

The sunny day filled us with joy, and we played games in the park. Even the dark clouds did not dampen the fun. It was a day we'll always remember with big smiles on our faces!

Example text after preprocessing:
"sunny day filled us joy played games park even dark clouds dampen fun day always remember big smiles faces"

Text Analysis

In this step, a person gets the machine to change the tokens and ngrams into meaningful data for a computer to use. There are a number of ways this can be done. Two common ways are semantics analysis and sentiment analysis.

Semantics analysis looks at the meaning of words in context, like the way they're used in sentences. **Sentiment analysis** uses a special kind of **dictionary**. The dictionary is a separate dataset where words are assigned scores based on their emotional impact.

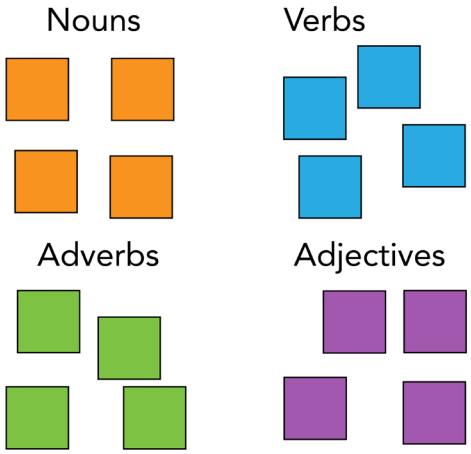

Words categorized by how they are used in context (©2023 Let's Talk Science)

Most sentiment dictionaries score words based on how positive or negative they are. For example, if we apply a sentiment analysis dictionary to our previous set of tokens:

"sunny day filled us joy played games park even dark clouds dampen fun day always remember big smiles faces"

It will provide a set of words and their scores. Below are the scores from our example paragraph. Any word not listed in the table has a sentiment score of 0.

Word	Negative	Positive
Sunny		1
Joy		1
Dark	1	
Clouds	1	
Dampen	1	
Fun		1
Smiles		1

This paragraph would have a positive score of 4 and a negative score of 3. NLP typically uses large amounts of text in its analysis, and this would be one chunk of data in a large set. The machine would use this paragraph as an example of a fairly neutral paragraph.

Did you know?
Other dictionaries break sentiment analysis down further. One example is the Canadian NRC Word-Emotion Association Lexicon, which uses eight emotion groupings.

Feature Extraction

The goal of text analysis is to turn our human-written text into structured data. A machine usually expresses this data as a table or spreadsheet. This is often called a **Document-Feature Matrix**, or DFM.

Documents form the rows of the table. A document is a collection of tokens. These can be the individual tokens themselves. But often they are paragraphs, chapters, or whole reports.

Features are the columns in the table. These are the things we want the machine to consider. In this example, the positive and negative scores would be the features.

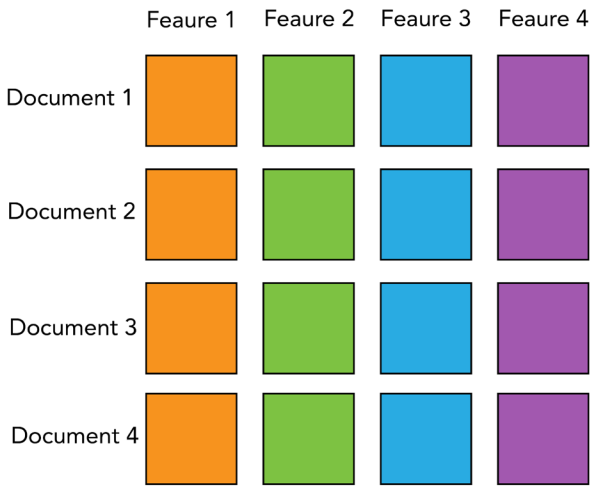

Example Document-Feature Matrix (©2023 Let's Talk Science).

If we imagine our example paragraph as just one of many paragraphs in a story, our DFM might look like this:

Document	Feature (Positive)	Feature (Negative)
Paragraph 1 (our example above)	4	3
Paragraph 2	2	8
Paragraph 3	6	1

In the structured set of data above, the positive column is one feature, and the negative column is a second feature. The individual scores, like the 4 in the positive column for paragraph 1 are called **vectors**.

An ML model for this set of data can then group the paragraphs by their features. A categorization model might create three groups. One is for strongly positive paragraphs, like paragraph 3. One is for strongly negative paragraphs, like paragraph 2. The third is for more neutral paragraphs, like paragraph 1.

A generative ML model, like ChatGPT, would then use these categories as examples. If it was asked to produce a positive text, it would use the strongly positive paragraphs as a guide on how to do it. Of course, an engine like GPT-3.5, which powers ChatGPT, uses many features to create text. But this gives you a simplified idea of how it works.

NLP is used for:

- **Text classification**, like organizing survey responses into categories and detecting spam;
- **Text generation**, like having chatbots build sentences by predicting what word will follow each generated word;
- **Translation**, like the **DeepL app**, which use predictive ML models to translate text from one language to another; and
- **Summarization**, like taking long pieces of text and creating short summaries, based on examining words in context.

Chatbots

Chatbots are one of the many apps that rely on NLP. The word "**chatbot**" comes from the words "chat" and "robot". But there are no mechanical robots here. Chatbots are software programs.

Chatbots are an app you can find on the web. Many companies use them for customer service. You may see them pop up when you go to a website. Chatbots are often the first point of contact with a company on their website or social media channels. They try to answer people's questions without putting them on hold to talk to a human agent. This means they must be able to understand a customer's question and respond with the best possible answer.

Many basic chatbots have a predefined set of questions and answers. You often have to choose from a list of options when interacting with these chatbots. The ways they can help are limited. Some chatbots are more sophisticated though. They allow you to type in your question the way you would normally ask it. This is where natural language processing comes in.

Person typing to a chatbot (Source: Blue Planet Studio via iStockphoto).

Natural language processing allows chatbots to understand different ways of asking the same thing. That may sound easy to you, but it is tricky for a computer. Think of all the ways that someone might ask, “what is the weather going to be like today?” They might say: “Do I need an umbrella?” or “What’s the temperature outside?” or “Is it going to be hot today?” Natural language models can group questions by the customer’s intent. That way the company can teach a chatbot how to answer different intents, rather than specifically worded questions.

NLP can even help companies figure out how customers are feeling based on their chat transcripts. For example, if a person used words associated with anger, a chatbot could connect them directly to a human!

NLP can only understand a person's words. More complex systems also use other information from speech, like tone. A system that analyzes a person's tone of voice can identify if they sound angry, sad or happy. Combining NLP and speech recognition gets us much closer to having computers understand people.

ChatGPT

Probably the most famous NLP chatbot right now is **ChatGPT**. This was developed by OpenAI and based on their GPT 3.5 engine. ChatGPT is designed to generate text that looks, at first glance, like it was written by a person. Like all chatbots, it can engage in conversation with a human operator. The difference is that it can understand and respond to a much greater range of input than your typical chatbot. ChatGPT can generate idea prompts, help to modify language to fit a certain audience, or work with people to fix computer code. It was trained on an enormous amount of internet text. It can generate answers with the correct context and provide further detail if prompted.

Search Tools

NLP is an essential tool for web searches and recommendation engines. **Web search tools** try to find the most relevant information based on what you are looking for. To do this, they need to understand what is contained in the many different web pages they can suggest. This involves analyzing text in news articles, blogs, and even the language in videos. NLP looks for specific keywords and phrases to help search engines find what you are looking for.

Spam Filters

NLP is also useful for email spam filters. NLP helps detect patterns commonly found in spam emails. This can make them better at catching and redirecting these emails.

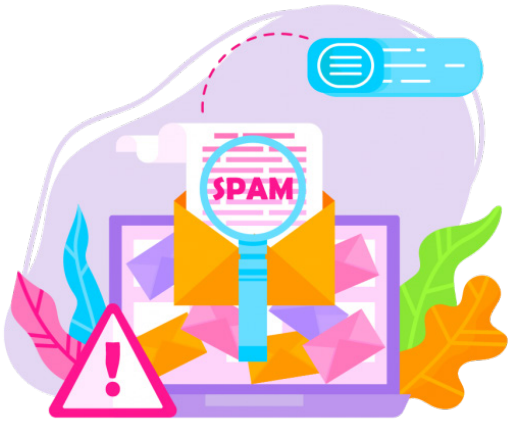

Spam filter (Source: axel2001 via iStockphoto).

Summarizing Text

NLP can also be useful for summarizing long or complicated documents. For example, scientific research is usually shared in scientific papers that are sometimes hard to read.

NLP tools can create short summaries of these. This way, scientists can find and read the most relevant papers for their work. Law firms have also been using NLP tools to summarize legal documents. NLP tools can quickly find information like dates, names, and the outcome of a trial.

Speech Recognition

Speech recognition (SR) also uses NLP. This involves turning audio files into text. Some examples of SR include talking to a chatbot on the phone, dictating a text message to your friend, or talking to a virtual assistant like Alexa or Siri. Web search tools also use speech recognition to transcribe videos and audio files like podcasts.

The most popular speech recognition apps are virtual assistants. These include Apple's Siri, Microsoft's Cortana, Amazon's Alexa and Google's Google Assistant. Virtual assistants are a sophisticated kind of chatbot. They have much better NLP and speech recognition skills. Virtual assistants can also do different types of tasks. They can play music, find a document or look up a location on a map.

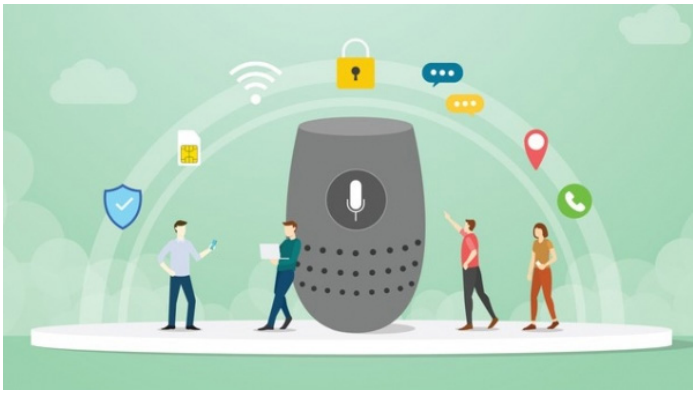

Smart virtual assistants use speech recognition (Source: ribkhan via iStockphoto).

Did you know?
Early virtual assistants had female voices. Based on feedback from users, some virtual assistants now offer voices of different genders. But the best solution might be a gender-neutral voice, such as that of Q, the virtual assistant.

Having your own personal assistant sounds pretty useful, but it comes with a price. For them to be ready to answer you, they listen to you at all times. This is a privacy risk. But it can be managed by adjusting a few settings. So be smart with your smart devices, and what you say and do around them!

Translation

NLP has allowed breakthroughs in the world of multilingual communication. Automatic translators on the web are available to everyone with an internet connection. Together, speech recognition and NLP let people who speak different languages communicate in real time.

Accessibility

NLP also helps remove communication barriers for people with disabilities. People with hearing loss can have easier access to closed captioning of recorded and live video content. AI apps exist that allow Deaf people to "see" sounds.

Combining this technology with computer vision, there is a system that can turn sign language into speech!

What the Future Might Bring

NLP is a technology that is moving very fast. In fact, we had to revise this article once already because of all the advances in AI. More and more companies are finding ways to use it.

Health care is one area where NLP can help. NLP can save doctors time by asking patients basic questions before their appointments. Tools like this could be very useful in remote areas where people do not always have access to doctors. Robots are also starting to be used to care for people, like senior citizens.. NLP could help these robots to better understand them. NLP is also a useful tool that could help assess someone's mental health.

So, the next time you search for something online, or talk to a chatbot, remember that it could not happen without AI and ML.

Let's Talk Science appreciates the contributions of Melissa Valdez, Technology Consultant from AI & Quantum, for revisions to this part of the handbook.

AI and Computer Vision

Visual and Facial Recognition Technology

Throughout history, humans have developed machines to do work for us. Recently, this has included machines that imitate our senses, like our vision. **Vision recognition technologies** are technologies that can see and label things. These let machines, robots, and apps see and understand the world as we see it.

Computer vision (CV) is a type of computer engineering. It involves teaching computers to "see" digital images like photos and videos. Engineers who work in this field have a variety of tasks. One thing they do is to find ways to use digital cameras with devices and computers. They also find ways to teach computers to recognize images and videos. This is done through coding or machine learning.

There are different types of computer vision. They depend on what the computer is trying to identify. The computer may look for text, images or faces. We will look at these three in detail.

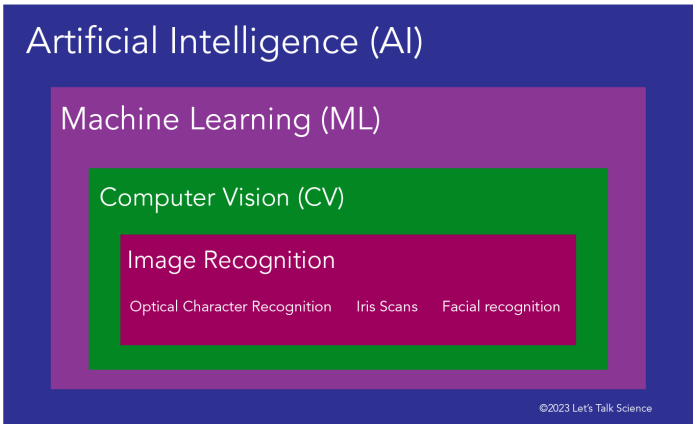

Image showing the relationship between different Artificial Intelligence and different aspects of Computer Vision (©2023 Let's Talk Science. Based on an image by deepomatic).

Optical Character Recognition

Optical character recognition (OCR) is a technology used to look for text. The text may be handwritten or in typed documents.

Let's see how it works for handwriting.

The first step in OCR is taking pictures of people's handwriting. Then, people scan them into a computer. Next, people match the handwritten text with the **characters** on a computer. A character is any letter, number, space, punctuation mark, or symbol. This teaches the computer which handwriting goes with which character. Now the computer will be able to identify and match handwriting with text.

This is an example of **supervised machine learning**. Supervised machine learning involves giving data, like images, labels, and file names. In OCR, machines learn to identify characters using many labelled images of handwritten letters. The machine can then look for patterns in all the images of the same character.

Let's take the example of the character one (1). Rules can be set to look for the following patterns in how humans write the character 1.
Pattern rules:

- Often found close to other numerals.
- A long straight vertical line, e.g. I
- An optional short line that hangs from the top backward at 45 degrees, e.g. 1
- An optional short horizontal line centred on the bottom, e.g. 1

Illustration of handwritten variations of numeral one (Source: Daranz via Wikimedia Commons).

> Try this!
> How would you describe a pattern for the numeral 3? or the numeral 9?

These types of pattern rules can be written as computer code. The code includes a step by step set of instructions, or **algorithm**. Once a computer has a code, then an OCR program can translate handwriting into computer text.

Some computer vision models can learn and record the pattern rules themselves. Then, when they see a new character, they analyse it the same way and find which group it matches with.

OCR technologies are now found in some smartphone apps. These apps take photos of your handwritten notes. Then they convert them into digital text. Being able to handwrite our notes and turn them into text is much easier than typing on a small device. Turning visual information, like your handwritten notes, into text data has many advantages. Text data can be searchable, it can be put into categories, and it takes up a lot less memory on your phone or computer!

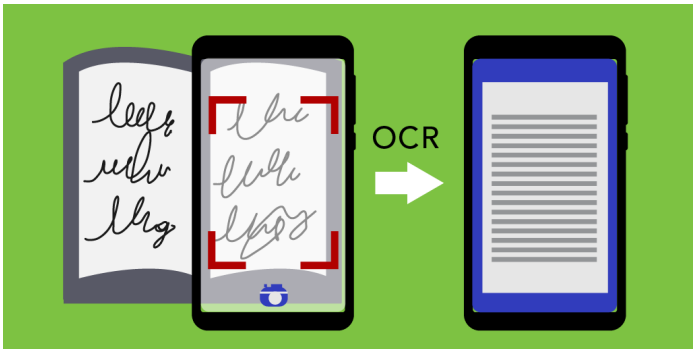

Image showing how OCR converts handwriting into typed text (Source: Piscine via iStockphoto).

Object and Visual Recognition

Many manufacturing processes involve machines and robotic systems that detect and recognize objects. Object detection can be as simple as a sensor that uses light to see if an item has passed by. Think of a labelling machine. It detects if a box moving along a conveyor belt is in the correct position. When the system 'sees' that the package is in the right location, it prints a label on it.

Today, people are developing even more complex visual recognition systems for robots. These let the robots better identify and handle objects. It is important that these systems come close to matching human abilities. For example, a robot needs to recognize and adjust its grip one way for a paper cup and a different way for a glass cup.

Simple visual object detection systems detect ***where*** something is. This is like the back-up camera in a car. It uses object detection sensors and cameras to detect objects. But it doesn't tell the driver ***what*** the objects are.

Image recognition systems figure out what objects are. This is one of the most important systems in autonomous cars. Like cars with sensors, autonomous cars need to be able to detect objects. But they also need to decide what to do, depending on the object and situation. For example, if the car recognizes a stop sign, it needs to stop. But if a car detects a person, it needs to analyze where that person is and what they're doing. Is the person safely on the sidewalk? Is the person crossing the street? You can imagine that this system needs to be really good at what it does!

3D image of self-driving cars. The rectangles show where the car detects other vehicles (Source: 3alexd via iStockphoto).

Autonomous cars are not the only systems that use image recognition. The smartphone app PlantNet is another example. It lets people find information about plants. Using your phone, you take a picture of the plant. The image recognition system compares your image to many other images of plants it already knows. It then suggests what your plant is. Leafsnap and Florist are similar apps. They help people to identify trees and flowers from images or their camera.

Facial Recognition Technologies

Facial Recognition Technology (FRT) is a technology that identifies human faces. The process it uses is like the way humans recognize each other. A computer's facial recognition system is like your facial recognition system. You see someone's face with your eyes. A smartphone takes an image of someone's face with its camera. Your brain takes the features of the face and stores it in your memory. This is what lets you remember people later. A computer does the same using algorithms.

Faces are unique. Like a fingerprint, we can measure and compare them. The term for measuring biological features is **biometrics**.

Facial biometric software measures and maps parts of a face. This includes things like the shape and colour of eyes, noses, mouths and chins. We call these measurements nodal points. A geometric map of a person's face needs about 80 nodal points.

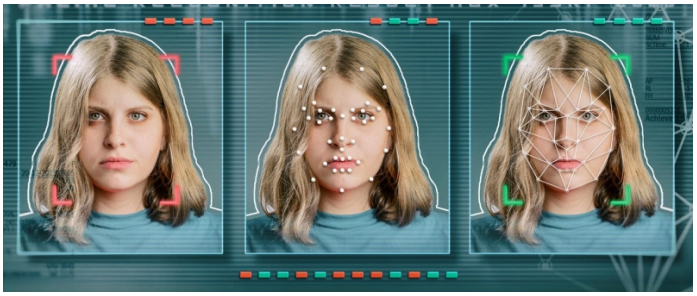

Facial recognition showing nodal points and measurements (Source: Grafissimo via iStockphoto).

The image and nodal points are then written as code. We call this code the faceprint or facial signature. Once a faceprint exists, a computer can compare it to other faceprint codes in a database of pictures. Faceprints are pretty unique, but they are not as unique as an iris scan or iris print. An iris scan is an image of a person's iris. The iris is the coloured part of your eye. Your iris is unique to you, like your fingerprint. This makes it a good means of identification.

Did you know?
Iris scanners use around 240 nodal points.

Many areas now use FRTs. The main area is security. Some smartphones and locks use faceprints or iris prints instead of passwords. The advantage of using your face is that you don't need to remember your password!

Law enforcement can use FRT to identify criminals from surveillance video footage. Governments could use FRT to confirm a person's identification. They could also use it when issuing passports, or at borders and airport security.

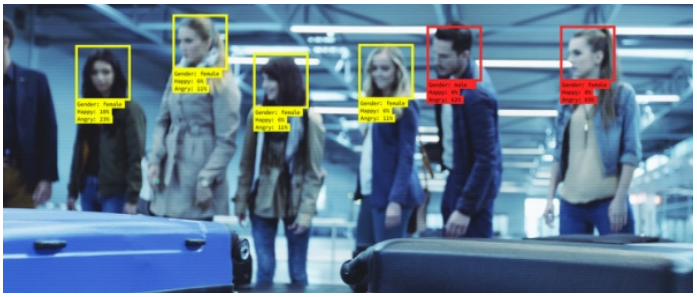

Facial recognition being used at the airport (Source: izusek via iStockphoto).

Unlike your face, your iris doesn't change over time. So it can be used to identify you throughout your life. But iris prints are not as easy to take as faceprints.

Concerns About FRT

FRT is pretty good, but it is not always accurate. One problem is that the pictures and videos we take may not be clear. Photos taken in poor lighting can affect the ability of FRT to make a positive match. Changes in glasses, jewellery, and facial hair can also affect FRT. In those situations, the matching results can be wrong. New software for both 2D and 3D images captured from video are improving FRT. Some systems even allow for changes in hair or things people use to disguise themselves. These improvements will help make FRT more accurate.

Another issue with FRT is the quality of the data given to computers. Algorithms used to analyze biometrics are given thousands of pictures of people. But sometimes computers are not fed enough data on certain groups of people.These include people who are visible minorities in North America and Europe. This problem can lead to false identification. If used in law enforcement, false identification can have serious impacts on people’s lives. This is why we need to be careful when using technologies like FRT for identifying people.

Privacy is a big concern when it comes to FRT. What we look like is a big part of our identity. In some cases, we are okay with others having images of us. This includes groups like the government who provide us with photo identification. What we do not want is people using images of us without our knowledge or permission. For example, some cities in China use FRT for shaming people. The names and pictures of people who break the law are shown on big screens. But in North America, some cities are already banning facial recognition.

One place you need to be careful of FRT is on social media. Did you know that when you post a picture on social media, you are giving the social media company permission to use it for their own purposes? Probably not. FRT allows these companies to collect and match faces with names. What they do with this information is not always clear.

More and more object and vision recognition systems are coming into our lives. These technologies can provide us with security and let us do things we could not do before. But we need to be aware that these technologies could also affect our freedom and our privacy. It is up to you to control how much information you share about yourself. This includes your face.

There are some things you can do. You can be thoughtful about who takes pictures of you and where they are posted. And you should always read the privacy policy for any social media platform you use. You should also pay attention to the news about your country's regulations on privacy. Being an informed citizen is always a smart choice!

Let's Talk Science appreciates the contributions of Melissa Valdez Technology Consultant from AI & Quantum for revisions to this backgrounder.

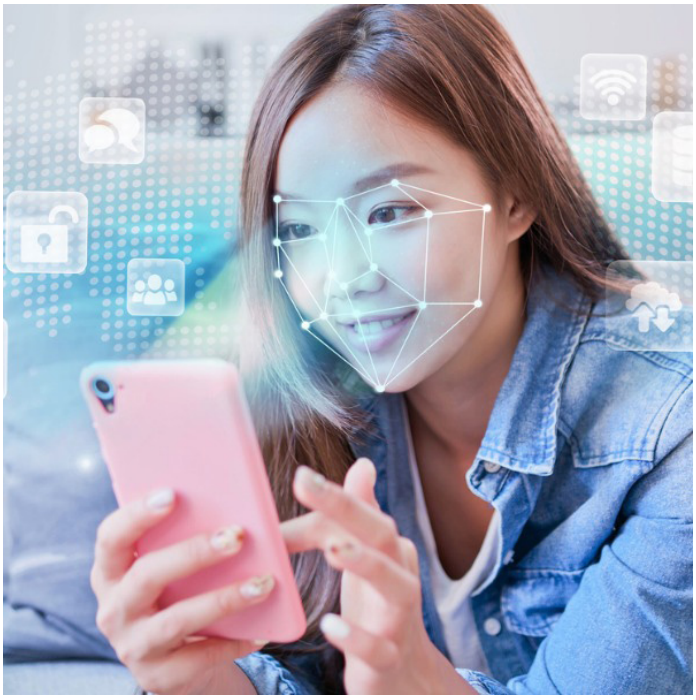

Young woman and smart phone with facial recognition (RyanKing999, iStockphoto)

AI and Personal Vehicles

Artificial Intelligence (AI) may seem new. But AI apps have been used in transportation for a while. Many modern vehicles use a **Global Positioning System (GPS)**. This uses data from satellites to figure out where the vehicle is on Earth. **Mapping algorithms** use AI to determine the best way to get from point A to point B.

To do this, AI systems have learned to predict the best routes from huge amounts of data. Then they combine this data with real-time information from users. This includes things like how fast they drive along the route. These two types of data can give people accurate and precise information about their trip. The AI can even help drivers get around traffic and avoid construction.

GPS system attached to a car's windshield (Source: Leo Malsam via iStockphoto).

Did you know?
Machine Learning (ML) is a type of AI. It is used to develop most transportation systems that use AI.

Many of the safety features in modern vehicles use AI. One example is driver assistance. Driver assistance systems help alert drivers to possible dangers. This could be something like beeping when a car is drifting out of its lane. It does this using a variety of sensors, including cameras and **infrared** sensors.

Some systems also help with the task of driving. This can include specific functions, like control systems that adjust the vehicle's speed or steering. It can also include more general functions, like using machine learning models to provide decisions based on different traffic conditions. All of these functions send data to a central data hub. Then that information is used as more training data for future models.

AI and Traffic Monitoring

No one likes being stuck in traffic! So city planners are always looking for ways to improve the flow of vehicles on roads. Installing sensors on traffic lights can help. The sensors send data to a large remote **database**. The data is then used to build different light-timing scenarios, which are analyzed to determine the best settings. City planners can also use machine learning to design better road systems. This could include things like building roundabouts instead of traffic lights.

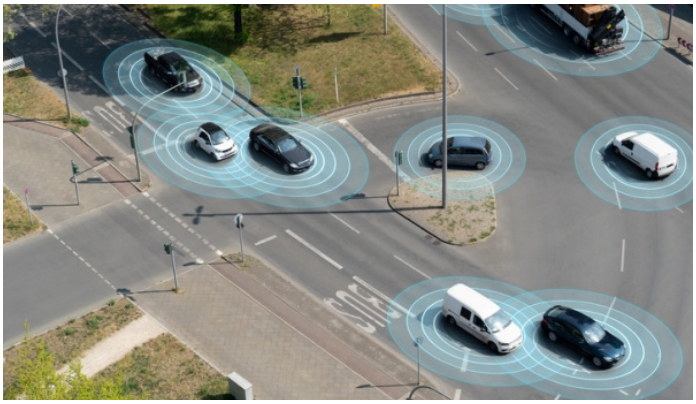

Sensors used to track vehicles at an intersection (Source: IGphotography via iStockphoto).

AI and Road Safety

Did you know that more than 1 700 people die from car accidents each year in Canada? When looking at the whole world, this number jumps to approximately 1.3 million people. And that is not counting the more than 20 million people who suffer non-fatal injuries each year.

Three common causes of car accidents are speeding, impaired driving and distracted driving. To improve road safety, we can use AI systems to identify people who are doing these things. An AI system can look for patterns in people's driving, both good and bad. We can then teach the systems to look for certain things that are dangerous, like speeding.

Did you know?
Robocar, the fastest autonomous car, reached a speed of 282.42 km/h!

AI and Self-Driving Vehicles

Unlike humans, machines do not do reckless and dangerous things. This led people to wonder if self-driving, or autonomous, vehicles could make our roads safer.

Safety is the most important reason people are developing self-driving cars. But it is not the only factor. Time is another factor. Imagine if people could use the time spent driving for other more fun or more productive tasks.

Did you know?
Driverless vehicles might seem like a solution to reduce traffic. But a study has shown that people who use driverless vehicles could actually spend more time on the road! This is not good news for the environment.

When it comes to cars, there are different levels of autonomy. Most modern vehicles have some features of Stage 2 automation. Some new cars even have Stage 3 or 4 (see the following page).

These cars can drive themselves under certain conditions, such as on the highway. It is important to remember that this technology is still new, and not perfect. People still need to keep their eyes on the road while driving in autonomous vehicles.

How do self-driving vehicles work?

To drive themselves, cars need both hardware and software. The hardware is a set of sensors and mechanical parts. It allows the car to sense its environment and provide data for the car’s computer. It is like the driver’s eyes, hands and legs. Software is computer programming. It allows the car’s computer to make decisions. It is like the driver's brain.

Autonomous cars use many kinds of technologies to sense their environment. This includes high-definition cameras, ultrasonic sensors, radar and LIDAR. Radar uses radio waves to detect objects. LIDAR is like radar except that it uses pulses of light to detect objects.

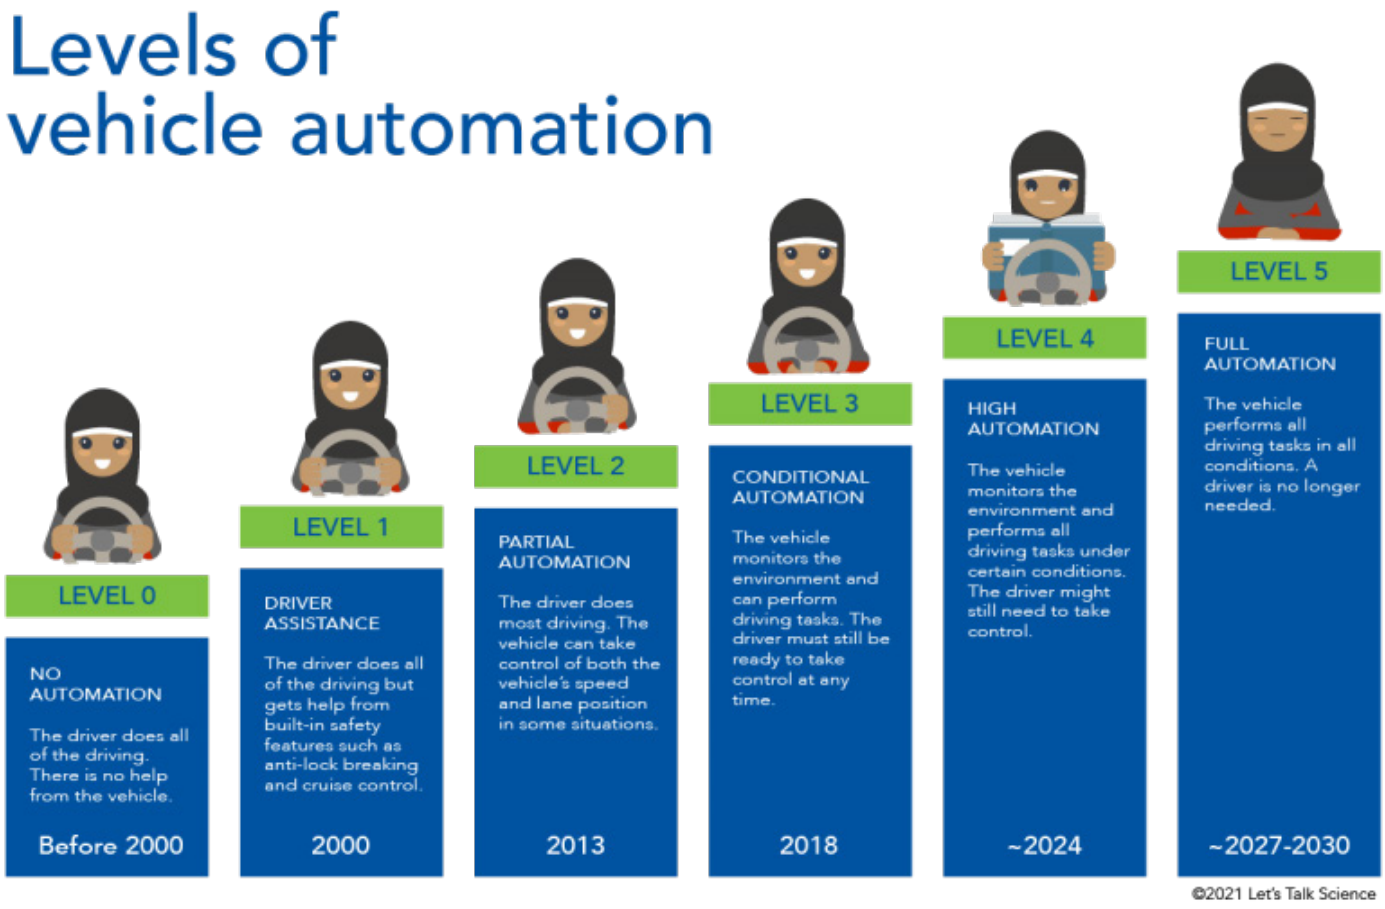

Concept image of the BWRX-300 SMR plant at the Darlington nuclear facility (Source: GE Hitachi Nuclear Energy. Used with permission).

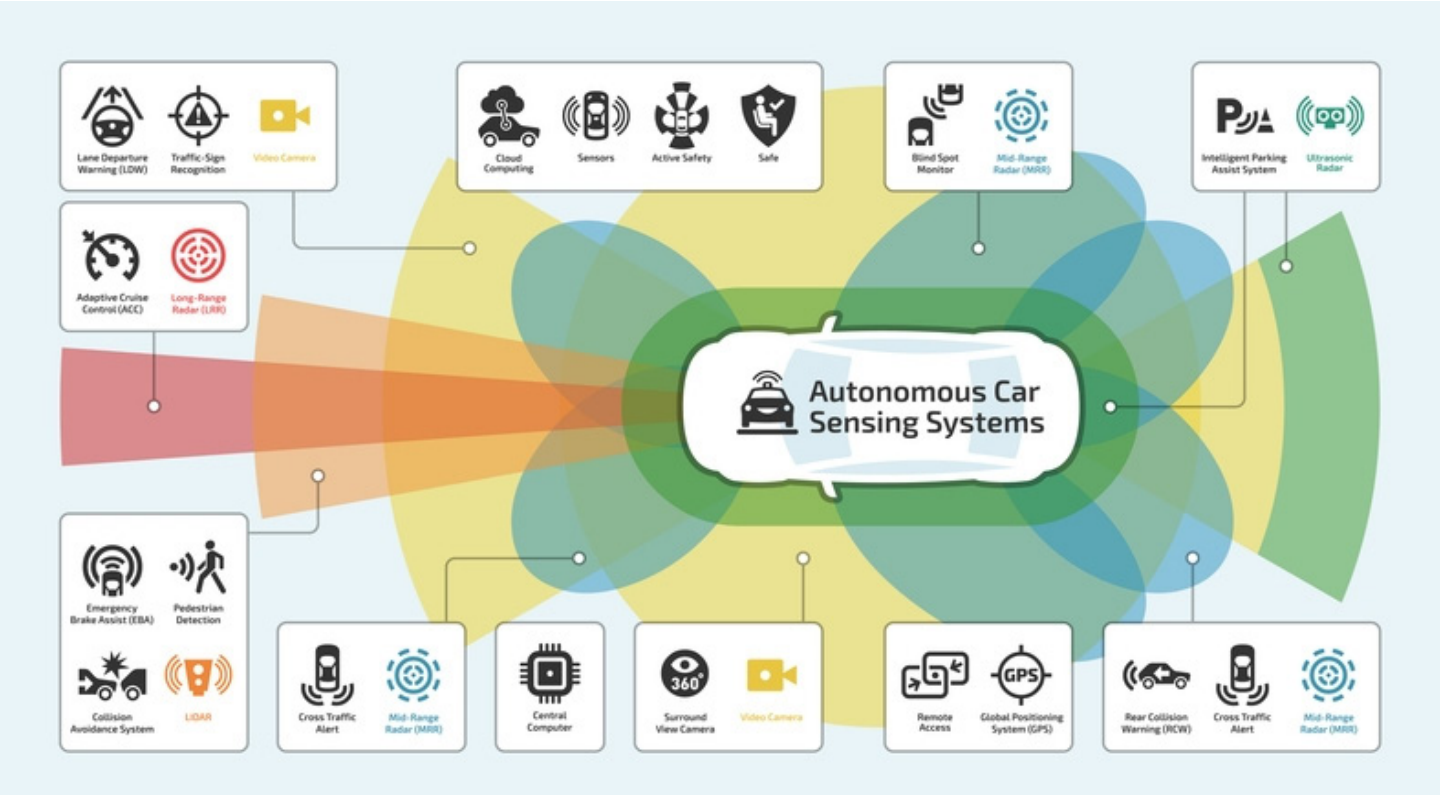

Concept image of the BWRX-300 SMR plant at the Darlington nuclear facility (Source: GE Hitachi Nuclear Energy. Used with permission).

These let the car detect traffic lights, people riding bikes, or even a squirrel crossing the street! They are especially useful when weather conditions reduce visibility.

The vehicle's software also uses information from GPS. This includes where the car is, and information like speed limits. This is a lot of information. Which is why an autonomous car needs a powerful computer. This computer must also process all this information very quickly. Delays in deciding how to move the car could be very dangerous!

Programming Driverless Vehicles

We once thought that by now, everyone would be using driverless cars. So, why is this not happening? It is pretty simple. Creating machines that can make decisions for themselves in a world of humans is tricky.

Sometimes when driving, a driver finds themselves in a difficult situation. For example, a driver suddenly sees a coyote standing in the middle of the road.

Sensors used to track vehicles at an intersection (Source: IGphotography via iStockphoto).

On the side of the road is a deep ditch. The driver hopes the animal will run off, but it is not moving. The driver will not be able to stop in time. Should the driver swerve to avoid hitting the coyote? If they do, they might end up harming themselves and their car by going into the ditch. Or should they hit the coyote? If they do, the coyote may die, but the person and their car would be okay. What would you do?

If you think that making a decision like that is hard, imagine trying to create a computer program to do it! This is exactly what AI engineers are working on. Going back to the example of the coyote, do you think that everyone made the same decision as you? How you made your decision depends on your values. In other words, what you think is important. Not everyone thinks the same things are important. A study of data gathered through MIT's moral machine proved this.

The moral machine is a set of car accident scenarios in which people decide what they would do, given a choice. The study found that people around the world made similar decisions. People preferred to save people rather than animals. They preferred to save more lives over fewer lives. And they preferred to save children instead of adults. They also noticed some differences between countries. These likely have to do with what things people in a certain country value. For example, some countries value their elders more than others.

A related issue to this is **bias** in the training data used to develop ML models. ML models are only as good as the data that goes into them. One example of this is the data used to train the systems self-driving cars use to detect and avoid pedestrians. The datasets used to train these systems need to be extremely large and diverse to cover every possible size, shape, and skin tone of human beings. If they aren't diverse enough, we could end up with systems that are better at identifying pedestrians with some skin tones than others. This could mean certain groups of people are in danger, based on their skin tone or other factors.

In summary

The role of AI in personal vehicles is increasing. Someday these vehicles will probably be the norm. The next generation of people may not even learn how to drive!

Let's Talk Science appreciates the contributions of Melissa Valdez Technology Consultant from AI & Quantum for revisions to this part of the handbook.

Canada: A Global Leader in AI Technology

Artificial Intelligence (AI) may seem new. But it has taken decades of research and development to get to where it is today. Canada started investing in this technology over 40 years ago. Since then, Canada has become a major global centre for AI development.

Early AI Research in Canada

Two of the leading AI researchers in Canada are Geoffrey Hinton and Yoshua Bengio. For decades, they have been advancing scientific research on AI and have contributed groundbreaking ideas.

Geoffrey Hinton

Geoffrey Hinton was born and raised in England. He came from a family full of famous doctors, mathematicians, and surveyors. Geoffrey was particularly interested in the human brain and neuroscience.

Neuroscience is the study of the brain and how it works. **Neurons** are nerve cells in the human brain. Each neuron can send electrical signals to other neurons. Neurons are connected or "networked" together in the brain. This network helps a person process information. There are about a hundred billion neurons in an adult human brain.

Hinton wondered if people could train a machine to learn like a human brain. This is something scientists had been thinking about since the 1940s. This type of approach was later called **machine learning (ML)**. His interest led him to complete a Ph.D. in AI in 1978. His focus was on neural networks. In 2012, it became known as **deep learning**. Deep learning is a type of machine learning based on artificial neural networks. The name "deep" means that the network has many layers. These allow the network to learn more and make better predictions. In 1987, Hinton came to Toronto, Ontario. He felt that Canada was a good place to live and do research. In 2017, he helped launch the Vector Institute in Toronto.

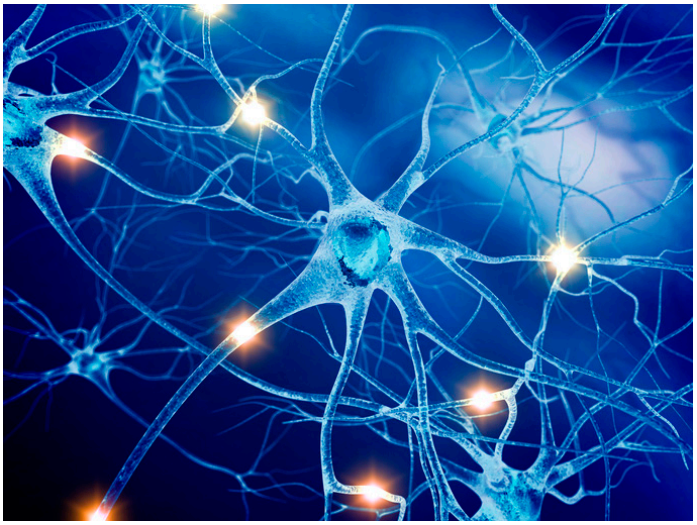

Active neurons (Source: SCIEPRO via Getty Images).

Yoshua Bengio

Yoshua Bengio was born in France, and moved to Montreal, Quebec as a child. As a teen, he loved science fiction and computers. He went on to study computer science at McGill University in Montreal. There, he discovered the work of Geoffrey Hinton. It inspired him to think about the question "What is intelligence?" Like Hinton, he believed in the potential of artificial neural networks. He helped launch Mila, the Quebec artificial intelligence institute.

> Did you know?
> In 2018, Geoffrey Hinton, Yoshua Bengio and Yann LeCun won the Turing Award. It was for their work with neural networks. The Turing Award is like the Nobel Prize for computer science.

What kind of AI work is being done in Canada?

In 2017, Canada became the first country in the world to have a national AI strategy. Many countries followed Canada's lead and developed their own. Canada's AI strategy is called the Pan-Canadian AI Strategy. It is led by the **Canadian Institute for Advanced Research (CIFAR)**.

The strategy helps provide a direction and goals for AI in Canada. This includes how AI is used in healthcare, the environment and other application areas.

The strategy also guides activities at Canada's three National AI Institutes. Each institute is led by a respected AI expert.

Alberta Machine Intelligence Institute (AMII)
(Edmonton, AB)
- Lead researcher: Richard Sutton

Vector Institute for Artificial Intelligence (Toronto, ON)
- Lead researcher: Geoffrey Hinton

Mila - Quebec AI Institute (Montreal, QC)
- Lead researcher: Yoshua Bengio

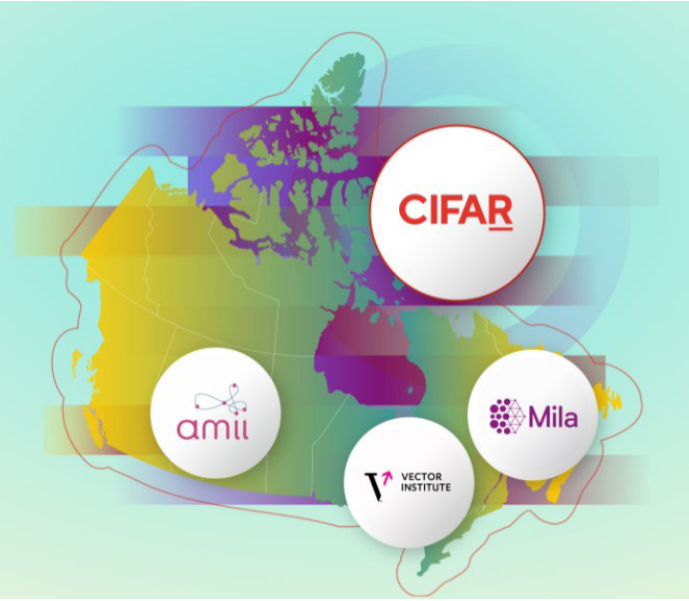

Locations of the three National AI Institutes (Source: Screen capture from CIFAR. Used with permission).

At each Institute, researchers, companies and innovators work together on responsible AI research and its applications. These centres attract people from around the world to work on AI-based solutions for many different problems.

What AI research is happening in Canada?

AI is an "enabling technology". This means that people can use it for just about any kind of challenge where data is generated. That is almost everywhere! But before it can help solve problems, AI needs to "learn" from data.

AI's ability to "learn" gives it huge potential. When used ethically, AI has the potential to change our lives. It could change the world for the better. This may happen in ways that we are just beginning to explore.

> Did you know?
> Canada has adopted a **Generative AI Code of Conduct**. It helps guide companies and organizations in the safe and fair use of AI for the good of Canada and the world.

Canadian AI researchers are doing important work for people and the environment. Let's look at a few **Canada CIFAR AI Chairs**. They have received a special award that helps support their groundbreaking research.

Parvin Mousavi is using AI to create new ways to detect and treat **cancer**.

Chair Parvin Mousavi (Source: Screen capture from CIFAR. ©CIFAR. Used with permission).

Jian Tang is using machine learning and computer models to design **proteins** and **molecules** for new drugs.

Chair Jian Tang (Source: Screen capture from CIFAR. ©CIFAR. Used with permission).

Martha White is using AI to help **water filtration systems** monitor themselves.

Chair Martha White (Source: Screen capture from CIFAR. ©CIFAR. Used with permission).

David Rolnick is using AI to combat **climate change**. His work helps communities better predict and prepare for extreme weather. He is also exploring ways to use AI to improve the health of **ecosystems** and improve **biodiversity.**

Chair David Rolnick (Source: Screen capture from CIFAR. ©CIFAR. Used with permission).

As of 2023, Canada has the third largest number of AI researchers and investments into new AI companies. This makes Canada one of the best places in the world for AI research, applications, training, and jobs.

Where can you learn about AI in Canada?

There is a wide variety of AI education options across Canada. Anyone can take a free on-line course called **Destination AI**.

You can learn:

- What AI really is (and what it's not!).
- The impact of artificial intelligence on the workplace and on society as a whole.
- And go behind the scenes of an artificial intelligence project to get familiar with deep learning and machine learning.

At the end of the course, students receive a certificate to show their school or employer.

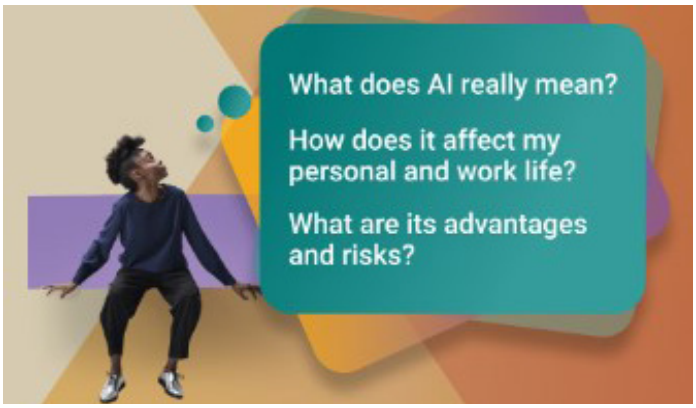

Youth pondering questions about AI (Source: Screen capture from CIFAR. ©CIFAR. Used with permission).

Certificates, bachelors, masters and doctoral studies in AI are available at many Canadian colleges and universities. And the list is growing. Amii, Mila, and the Vector Institute also partner with researchers at different post-secondary institutions. Some of these include the University of Alberta, the University of British Columbia, McGill University, Université de Montreal, Université Laval, University of Toronto, University of Waterloo, University of Guelph, Western University, University of Ottawa and Dalhousie University.

AI and Your Future

Does AI inspire you to make a difference in the world? Canada is one of the best places in the world to plan for a future career in AI. Everything you need, from education to jobs, is right on your doorstep. Canada has some of the best schools. It also has access to leading researchers. People are coming from around the world to harness AI for the good of everyone. Using AI will take many trained people. Would you like to become an AI expert to help people and make the world a better place?

Let’s Talk Science appreciated the contribution of Tibor Turi, PhD P.Eng., Project Lead, Pan-Canadian AI Compute Environment, CIFAR in the development of this part of the handbook.

How can you train a machine learning system?

Consider using the Design & Build Process with this challenge.

This activity will help build skills related to the Plan, Create, Test & Evaluate, and Reflect & Share phases of this process. It also highlights the iteration required during this process.

Materials:

- Icon Library (see below)
- Something to write with
- A partner

What to do!

Your job is to develop rules to train a machine learning system to recognize icons. In this activity, your partner will pretend to be the machine learning program. We have provided you with a library of icons to use.

1. **Plan** – Choose one of the icons from the library. Do not let your partner know which icon you chose. Look at the icons and think of different ways you could describe and sort them. This will help you when creating the rules that your partner will use to find the icon you chose. For this activity,
   - A rule cannot use the name of the thing (e.g., “It has four legs” and not “It is a cat”).
   - Each rule can only involve one characteristic (e.g., “It is alive” and not “It is alive and has legs”).
2. **Create** - Write down the rules that your partner (the computer) will use to find your icon.
3. **Test** - When you think you have a good set of rules, read them to your partner. See if they can figure out which icon you picked from the icon library.
4. **Evaluate** – Was your partner able to identify your icon? If not, revise your rules and try again.
5. **Reflect & Share** – What types of characteristics did you use for your rules? What rules helped your partner the most? Were any rules not useful? What are the fewest number of rules that could identify your icon? Would your rules have been different if the icons had been in colour?

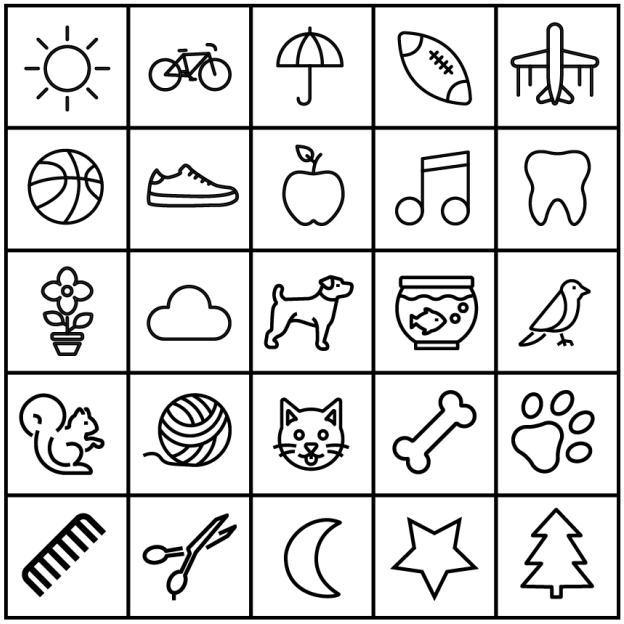

Icon library (©2023 Let’s Talk Science).

What’s happening?

A **machine learning** model is a program that can find patterns and make decisions based on a set of data. Machine learning is a form of **artificial intelligence**.

In a **supervised machine learning** model, people need to help the model figure out a pattern to correctly identify something. First the model is given labelled pictures of something, like an object. These pictures and their labels are called input data. The model uses the input data to come up with rules. Then it uses them to get the right output data. The output data is the identity of the object. Machine learning models need lots of input data to come up with accurate rules.

In an **unsupervised machine learning** model, the model is given unlabelled data. Then it is asked to find patterns. A human can then use those patterns to build a model that does the task needed. Sometimes machine learning models find unexpected patterns in data. This is why humans need to be in the loop. They help build models that work the way we want them to.

Why does it matter?

It's important to know how and why a machine learning model comes up with its output data. If the input data is flawed or incomplete, the output data could be incorrect, biased, or unethical.

Computers do not "think" the same way as people do. A task that humans think is easy can be very difficult for a computer, and vice versa!

Investigate further!

- Try your rules again with other icons. Do they still work? If not, what needs to be changed or added?
- Can you come up with rules that only use shapes to identify the icon?
- Read your rules to someone, without showing them the Icon Library. Can they tell what object you're trying to define? Can either of you think of other objects that would also fit your rules?

Melissa Valdez

Technical Consultant, Watson AI
IBM Canada

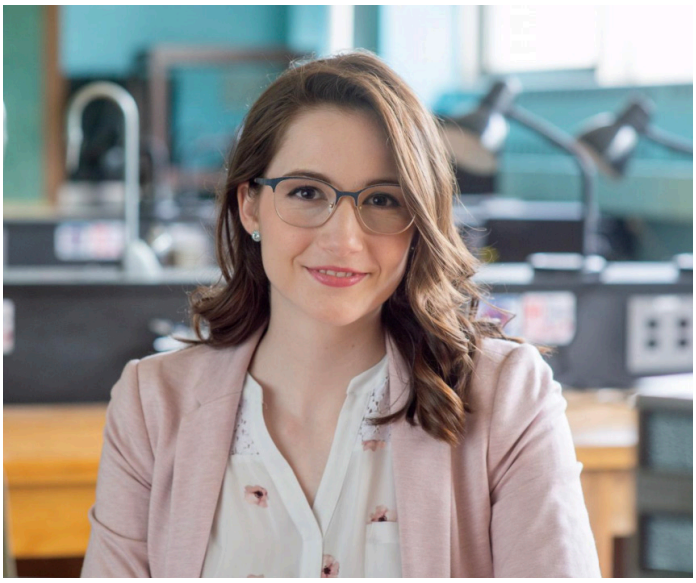

I was born in Windsor, Ontario and I now live in Toronto, Ontario. I completed my training/education at the University of Windsor (BSc 2015) and York University (MSc 2017).

What I do at work

Right now, most of my work focuses on building "Virtual Assistants" using a technology called Watson Assistant. Virtual Assistants are sophisticated chatbots. Chatbots are software programs that "talk" to humans by copying human conversation. By asking, and answering questions, they can perform tasks and assist customers. Watson Assistant is a Natural Language Processing (NLP) algorithm, a type of Artificial Intelligence (AI). It is my job to customize the AI for a specific uses. An example of this is the program that provides advice to university students on various personal and academic questions they may have.

My days are never the same. I have a home base at a Toronto IBM office, but about half the time I'm on site at one of our clients' offices for consulting work. Consultants can travel up to 80% of the time - 4 of 5 days a week! When I'm at IBM I work on "training" Watson so it can understand a user's input and respond appropriately. To do this I give examples of the types of questions people might ask and examples of the type(s) of answers that would make sense to a human.

My career path

When I was in high school, I didn't know what I wanted to do! I studied physics because it was my favourite subject. Volunteering when I was a student at university helped me expand my network. Someone I met as a volunteer was my connection to IBM and helped me get the job I have today. When I was doing research, I craved more in-person interaction. When I was working in outreach and education, I found myself wanting to do work that was more technical. Kind of like Goldilocks, I found my "just right" career in consulting. Now I get to use my technical skills and my people skills equally.

I am motivated by

I love that I get to learn new things every single day. IBM is really a big company and there are so many intelligent people working here from all corners of the globe. Consulting is a business focused on people - so there is a big emphasis on learning new skills and professional development. There are many courses, online resources, and working groups, you can join to learn new skills. That definitely keeps things interesting! Plus, no two projects are the same! We are always pushing to develop new and unique customer experiences. I personally enjoy it because it's a perfect mix of technical work and interaction with clients.

How I affect peoples' Lives

The virtual assistants we are building aim to help make peoples' lives easier. Our clients are all large companies, and the virtual assistants we build simplify interactions for potentially millions of customers, and sometimes thousands of their employees, too.

My advice to others

Jump into coding as early as you can - it's just like learning any language so the longer you practise the better you'll be! Other than that, volunteer for causes you care about to keep you engaged and meeting new people.

References

What is Artificial Intelligence?

Canadian Nuclear Safety Commission. (2021, May 04). Chui, M., Chung, R. and van Heteren A. (2019, January 21) Using AI to help achieve Sustainable Development Goals. United Nations Development Program.

Granville, V. (2017, January 2) Difference between Machine Learning, Data Science, AI, Deep Learning, and Statistics. Data Science Central.

Hello Future. How AI can help reduce inequalities (2019)

Leviathan, Y. and Y. Matias (May 8, 2018). Google Duplex: An AI System for Accomplishing Real-World Tasks Over The Phone. Google Research.

Maini, V. (2017, August 19) Machine Learning for Humans. Medium.com

Metz, C. (Nov 24, 2020). Meet GPT-3. It Has Learned to Code (and Blog and Argue). The New York Times.

O'Brien, M. (Aug. 1, 2023) Chatbots sometimes make things up. Is AI's hallucination problem fixable? Associated Press.

OpenAI (n.d.). DALL-E: Creating images from text.

Sample, I. (Oct 30, 2019). AI becomes grandmaster in 'fiendishly complex' StarCraft II. The Guardian.

Stark, L. and Zenon W. Pylyshyn. (2020, March 10). Artificial Intelligence (AI) in Canada. The Canadian Encyclopedia.

Timelines Wiki. Timeline of machine learning. (2020).

Wingate, J. (2018, April 30) Why Artificial Intelligence Will Make You Question Everything.

What is machine learning?

Heller, M. (2020, March 9). Data Labeling: AI's Human Bottleneck. Medium.

New Tech Dojo. (2018, March 6). List of Machine Learning Algorithms.

Shankar, S. (2020, May 30). Types of Machine Learning Algorithms. The Startup.

Van Loon, R. (2018, February 5). Machine learning explained: Understanding supervised, unsupervised, and reinforcement learning. Big Data Made Simple.

Vas3k Blog. (n.d.). Machine Learning for Everyone.

Wilson, A. (2019, September 29). A Brief Introduction to Supervised Learning. Towards Data Science.

AI and Human-Computer Communication

Brownlee, J. (2019, Aug. 7). What Is Natural Language Processing? Machine Learning Mastery.

Joshi, N. (2018, Dec. 23). Yes, Chatbots And Virtual Assistants Are Different! Forbes.

Mahler, L. (2015, Feb. 13). What Is NLP and Why Should Lawyers Care? Lawpracticetoday.org

Mortada, D. (2019, Mar. 11). Meet Q, The Gender-Neutral Voice Assistant. Npr.

Rebhan A. (2019). Natural language processing: How this emerging tool can improve mental health treatment. Advisory.com.

Rouse, M. (2017, Oct.). Virtual assistant (AI assistant). TechTarget.

Strong, S. (2019, May 22). Geospatial Machine Learning: Structuring Unstructured, Structured Data. Medium.com

AI and Computer Vision

AI Multiple. (2021) Image Recognition in 2021: In-depth Guide.

Bonsor, K. & Johnson, R. How Facial Recognition Systems Work. How Stuff Works.

Electronic identification. (2020) Face Recognition: how it works and its safety.

Marutitech. What is the Working of Image Recognition and How it is Used?

The Next Web. A beginner's guide to AI: Computer vision and image recognition

Panda Security. (2019, Oct) The Complete Guide to Facial Recognition Technology

Symanovich, S. (2019, Feb 8). How does facial recognition work? NortonLifeLock.

AI and Personal Vehicles

Awad, E., Dsouza, S., Kim, R. et al. The Moral Machine experiment. Nature 563, 59–64 (2018). https://doi.org/10.1038/s41586-018-0637-6

Hendrickson, J. (2020, Jan. 3). What Are the Different Self-Driving Car "Levels" of Autonomy?. How-To Geek.

IEEE Spectrum. Accelerating Autonomous Vehicle Technology.

Jensen, C. & Jensen, C. (2017, Feb. 9). States Must Prepare For Human Drivers Mixing It Up With Autonomous Vehicles. Forbes.

Rivelli, E. (2020, Jun. 30). How Do Self-Driving Cars Work and What Problems Remain?. The Simple Dollar.

Schroer, A. (2019, Dec.19). Artificial Intelligence in Cars Powers an AI Revolution in the Auto Industry. Built In.

TED-Ed. [TED-Ed]. (2019, May 13). How do self-driving cars "see"? - Sajan Saini. [Video]. Youtube.

Canada: A Global Leader in AI Technology

Association for Computing Machinery. (2018). A.M. Turing Award Laureate Geoffrey E. Hinton.

CIFAR. (2023). The impact of the Pan-Canadian AI Strategy.

Deloitte. (September 2023). Impact and opportunities: Canada's AI ecosystem - 2023.

EduCanada. Study artificial intelligence in Canada.

Government of Canada. (September 2023). Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems.