


A CPA's Introduction to AI: From Algorithms to Deep Learning, What You Need to Know



DISCLAIMER

This paper was prepared by the Chartered Professional Accountants of Canada (CPA Canada) and the American Institute of CPAs (AICPA), as non-authoritative guidance.

CPA Canada and AICPA do not accept any responsibility or liability that might occur directly or indirectly as a consequence of the use, application or reliance on this material.

© 2019 Chartered Professional Accountants of Canada

All rights reserved. This publication is protected by copyright and written permission is required to reproduce, store in a retrieval system or transmit in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise).

For information regarding permission, please contact permissions@cpacanada.ca

Executive Summary

Although not always noticeable to the general public, Artificial Intelligence (AI) has been around since the 1950s. However, advancements in AI have accelerated in the past decade due to the rise in computing power, the availability of data, and improvements in machine learning algorithms. As a result, AI capabilities have begun to infiltrate our daily lives—some more so than others. With AI adoption increasing for both consumers and businesses, it is important for CPAs to understand the fundamentals of AI and the potential impact to their organizations, clients, and themselves.

CPAs have an opportunity to use AI as an enabler of innovation and increased productivity. Leveraging AI, CPAs can perform better quality work with higher efficiency. But to do that, CPAs need a solid understanding of how AI works before they can identify these opportunities. For example, robotic process automation can increase productivity and appear to work like AI, but is it? Where do machine learning and deep learning fit into the world of AI?

This publication was developed by Chartered Professional Accountants of Canada (CPA Canada) and the American Institute of CPAs (AICPA) to be used as a foundational resource. It will explain the buzzwords and terms you have likely been hearing, discuss the evolution of data, AI and computing power and help you to begin to think about AI and how it might impact your work. Future publications dedicated to AI will refer to this resource as it is the first of a planned series of publications to explore AI and its impact on the CPA profession.

CPA Canada and the AICPA encourage all CPAs to continue reading, collaborating with colleagues and learning about AI and other technologies which may impact you as a CPA. We are all on this digital journey together!

Introduction

As you read this, a smartphone app called Climate Basic is telling farmers where to plant their crops. The app divides the land into plots, scans through local temperature and erosion records, expected rainfall, soil quality and other agricultural data to figure out how to maximize yields.

Somewhere, someone is sitting on their couch idly scrolling through Netflix's recommendations, blissfully unaware they are unique to them and were generated by an algorithm that compares their individual TV-viewing patterns with those of millions of users across 190 countries.

In the accounting profession we see automation in accounts payable where manual tasks have been replaced by systems that do everything from matching purchase orders to flagging invoices for payments. Moreover, through Optical Character Recognition and machine learning, algorithms can read those invoices and classify expense types or account types automatically with a high degree of accuracy.

Welcome to Artificial Intelligence (AI). From virtual assistant Siri to navigation app Waze to autonomous cars to humanoid robot bank ambassadors and hospital receptionists, to Google Assistant's ability to call a hair salon and set up an appointment for you, AI is changing how we live and work. AI is everywhere and we are still just scratching the surface of how it will change how we work and live.

But What Is AI, Really?

A short answer and the broadest definition: **AI is the science of teaching programs and machines to complete tasks that normally require human intelligence.**

The new space race: Global initiatives to win at AI

- Of the \$15.2 billion invested globally in AI start-ups in 2017, 48% went to China and 38% went to the U.S., as per CBInsights.¹ This is indicative of a global race to lead in AI. The **Chinese government** in its State Council paper has laid out its strategy to emerge as a global AI hub by 2030.
- On the other hand, in the **United States**, the Defense Advanced Research Projects Agency (DARPA), announced in September 2018 a \$2 billion campaign to develop the next wave of AI technologies focused on contextual reasoning capabilities.²
- **The European Union (EU)** in April 2018 launched a \$1.8 billion investment in AI research under the EU's Horizon 2020 fund.³
- **France** has announced a \$1.8 billion investment policy through 2022 focused on AI and data sharing.⁴
- **The United Kingdom** government announced a £300 million investment allocation in April 2018 and 50 leading UK businesses have come forward to set a £1 billion deal in AI.⁵
- The Government of **Canada** is all in as well. In 2017, it tasked the Canadian Institute for Advanced Research to develop a \$125 million AI strategy.⁶ The Canadian government is also investing \$950 million in the Innovation Super Cluster initiative which includes SCALE.AI (Supply Chains and Logistics Excellence.AI) an industry-led consortium that will shape a new AI-powered global supply chain platform.⁷

1 www.technologyreview.com/the-download/610271/chinas-ai-startups-scored-more-funding-than-americas-last-year/

2 www.darpa.mil/news-events/2018-09-07

3 www.forbes.com/sites/amyguttman/2018/09/28/three-innovative-european-startups-to-watch/#507861316f37

4 www.reuters.com/article/us-france-tech/france-to-spend-1-8-billion-on-ai-to-compete-with-u-s-china-idUSKBN1H51XP

5 <https://tech.newstatesman.com/business/uk-government-ai-research>

6 www.newswire.ca/news-releases/canada-funds-125-million-pan-canadian-artificial-intelligence-strategy-616876434.html

7 www.ic.gc.ca/eic/site/093.nsf/eng/home

Why Is It Important for CPAs to Stay Current with What Is Happening in AI?

Simply put, AI has the potential to significantly impact accounting and assurance jobs and businesses in the near future—if it has not already. Transactional tasks are increasingly being automated. Algorithms have the ability to spot subtle clues in data sets to pinpoint fraud or identify new opportunities such as savings, improved efficiency and profitability, and new market opportunities. All this while leveraging and making sense of a large volume of information coming from multiple sources, commonly referred to as Big Data.

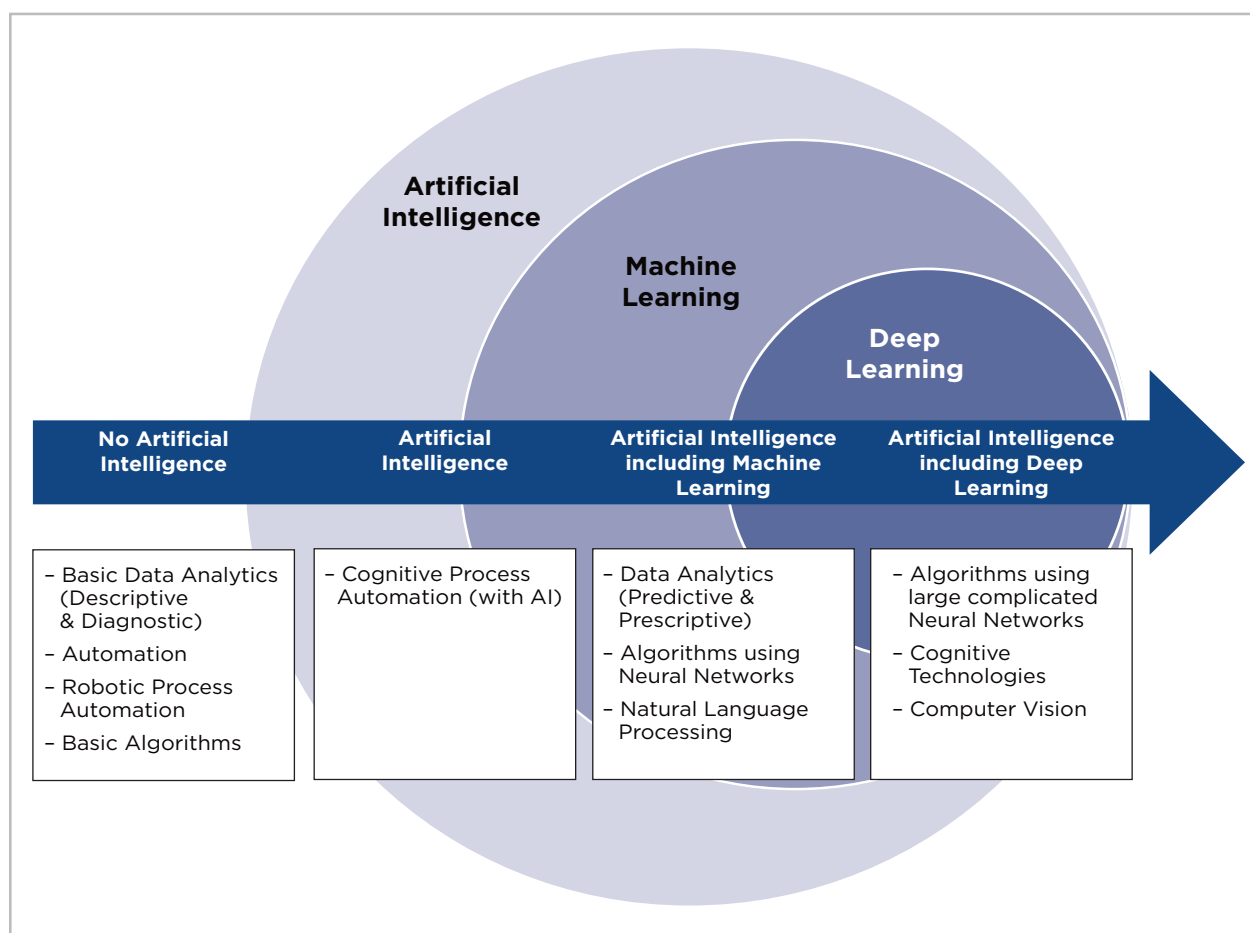
Of course, everyone has an opinion about AI. Some focus on the positive and how it will improve everything from transportation to medical diagnostics. Others see problems and ethical concerns: the potential elimination of jobs, and the violation of privacy. How far can we trust the autonomy of AI systems? For example, how does AI resolve a self-driving car's dilemma when it must decide whether to crash into an oncoming car or pedestrian when an accident is unavoidable? There is also the question of decision transparency. Would every AI system be able to explain in terms we understand, the reasons for its decisions?

CPAs have an opportunity to focus on AI as an enabler of innovation and increased productivity. Consider the Climate Basic app helping farmers. According to the U.S. Department of Agriculture, this use of AI across the industry has produced the largest crops in the country's history. And Netflix says its sophisticated recommendation engine saves the company \$1 billion in lost subscribers.⁸ CPAs have the opportunity to improve efficiencies and quality of day-to-day work through AI that reads contracts, performs automated tasks such as accurately identifying and booking expenses, or identifying risky transactions in real time.

Fast facts: The potential impact of AI is also fueling investments and research. The International Data Corporation (IDC) forecasts AI-related investments will grow from \$12 billion in 2017 to \$57.6 billion in 2021. Since 2013, machine-learning-based AI has been the third-fastest-growing category for patent filing. On the job front, employment search engine Indeed reports the share of jobs in Canada and the U.S. requiring AI skills has significantly grown since 2013.

8 www.businessinsider.com/netflix-recommendation-engine-worth-1-billion-per-year-2016-6

This publication was developed by CPA Canada and the AICPA to be used as a foundational resource. It will explain the buzzwords and terms you have likely been hearing, discuss the evolution of data, AI and computing power and help you to begin to learn about AI and think how it might impact your work. CPA Canada will also refer to this publication in future publications dedicated to AI as the first of a planned series of publications to explore the impact of AI in the CPA profession. The diagram below provides an overview of the terminology to be covered throughout this publication and helps provide a visual representation of where these technologies fit in relation to each other.



How We Got Here

It all starts with **data**. Companies have historically used data to improve their business performance (e.g., customer feedback surveys, planning resource allocation using the Critical Path Method, etc.). However, data gathering was an explicit exercise and those performing it might not have taken steps to verify that the data collected represented what was actually happening. For example, it may have been difficult to determine whether survey respondents truly represented the customer base or whether their buying habits were far different from those of more typical customers.

You can't have
Machine learning
without data.

With the Internet and the evolution to a digital world, data acquisition has gone from being an explicit and often manual exercise to an ongoing, automated passive activity. Websites, apps and social media can track everything from your shopping behaviour to your heart rate. Many business processes can be digitized leading to the collection of information to help reduce costs. Real data in real time. No surveys required to be completed by a user, no response bias: just facts.

From Distributed Computing to Cloud Computing

With all this passive data collection, data sets have become so massive and complex that traditional data processing software simply cannot keep up. Companies such as Google, whose business is based on data, set about finding a way to increase computational power in order to process Big Data. The answer: **Distributed Computing**. Rather than having one sophisticated computer handling all data processing, the data is divided and sent to hundreds or thousands of commodity computers working in concert but independently thus cutting down overall processing time.

This solved the problem for large technology companies, but smaller organizations still had no way to access and process large volumes of data effectively. Enter **Cloud Computing**, a rentable network of remote servers hosted on the Internet to store, manage and process data. The scalability of Cloud Computing has eliminated the need to buy and maintain costly computer systems; it provides access to more on-demand resources and processing power for everyone, including small organizations. As a result, all organizations can now benefit from Big Data.

A word on Data Warehouses, Data Lakes and Data Swamps

How is all this data stored? Typically, via a **Data Warehouse** or a **Data Lake**. A Data Warehouse is designed to store data and provides businesses a central repository to integrate, manage and analyze data at many levels. The data can come from internal and/or external sources and is largely structured. A Data Lake is both a way to organize huge volumes of wildly diverse data, both structured and most often unstructured, and store it. When a Data Lake becomes unmanageable, it turns into a **Data Swamp**.

When Big Data Meets AI

The basic idea behind AI is to let a machine statistically analyze all the data being collected to derive insights much faster and more accurately than otherwise possible. For example, rather than programming a computer with a simple rule such as “more than 12 hours of daylight = summer; less than 12 hours of daylight = winter,” AI analyzes the parameters of numerous summer days. Using algorithmic models, it will then determine which parameters and at what levels those parameters constitute a summer day (e.g., if temperature is > 18 degrees, humidity of x%, etc.). It will then leverage that understanding and based on probabilities determine whether a particular day is a summer day or not by assessing all those relevant parameters against what it has “learned” to be a typical summer day.

“50 times the growth in data expected from 2010 to 2020!”⁹

Experts now agree that AI has evolved to a point where it has real, practical value. The reason? The **convergence of four trends**. The explosive growth in data is made possible by:

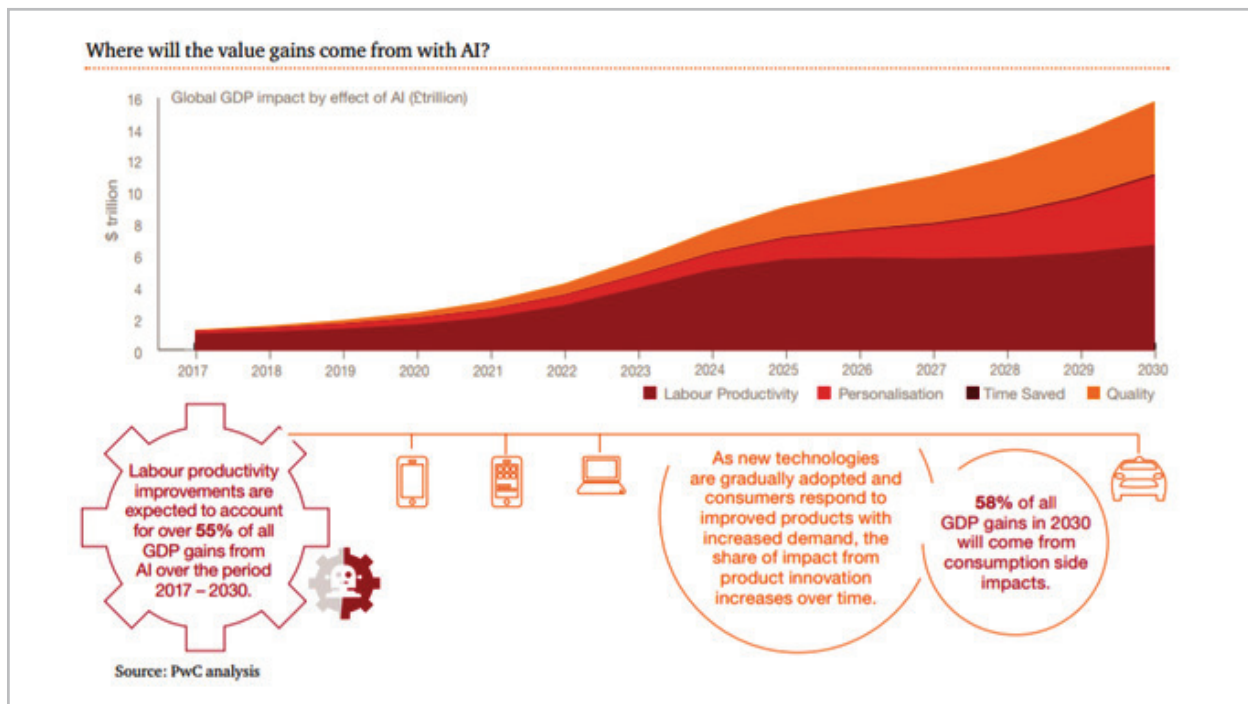
1. digitization of business processes, smartphones, Internet of Things (IoT) and social media/web 2.0
2. the advancement (and availability) in processing power to store and compute data
3. the maturity of algorithms and AI models
4. the huge upsurge in investment in AI.

9 www.slideshare.net/sfamiliar/working-with-big-data-jan-2016-part-1/7-CONTEXT_WHATS_BIG_DATA_BIG_IN

So what does all of this mean for business today?

Two concepts: improved productivity and innovation. These include simplification and automation of tasks that free up employees to focus on more stimulating and higher value-adding work. New innovative business models have been established as a result of AI, such as virtual assistants like Alexa, Google assistant or x.ai. Although not popular before AI, they are now a multi-billion dollar industry at the centre of the home automation market.

Capital-intensive sectors such as manufacturing and transportation are likely to see the largest productivity gains from AI, given that many of their repetitive operational processes are ripe for AI-enabled automation.¹⁰ AI robots are already taking over warehouses across the globe. AI is also poised to unlock efficiencies in professional services such as legal, human resources, accounting and audit where data plays a central role.



¹⁰ <https://preview.thenewsmarket.com/Previews/PWC/DocumentAssets/476830.pdf>

Within global accounting and auditing firms, work is underway to get out in front of and take advantage of the new paradigm emerging due to AI.

- BDO has successfully piloted its *BDOLexi* translation app, using neural network technology, to manage information in multiple languages during global audits.
- Deloitte has launched Omnia AI, a multi-disciplinary practice focused on delivery of AI services and solutions.
- EY is embedding AI technology across its end-to-end audit process, including the use of machine learning in EY Helix and EY Optix for predictive analytics, using drones for inventory observations and the development of business document-reading and interpretation tools.
- KPMG has launched its AI toolkit called Ignite and is working with IBM/Watson.
- PwC is investing in data platforms to securely capture, organize and facilitate analysis of data and is working with H2O.ai, a leading Silicon Valley company, to develop Audit.ai, which will provide enhanced capabilities to provide assurance.

Types of AI

There are essentially two kinds of AI: **Narrow AI** and **General AI**. As the name suggests, Narrow AI, which can also be described as weak AI, is made up of narrowly intelligent systems that can exceed humans in specific tasks, such as playing chess or making medical diagnoses. These narrow capabilities are not transferrable (i.e., an AI chess player cannot be used to perform another task such as a medical diagnosis).

General AI or strong AI refers to human-level intelligence that is able to transfer knowledge between domains. While Narrow AI is all around us in language and vision recognition systems and recommendation engines, General AI may be the stuff of science fiction and movies—for now.

How AI Works

In order for AI programs to navigate through situational complexities, different approaches to creating software, with the ability to determine different outcomes, are necessary.

The **Logic and Rules**-based approach uses conditional instructions and defined rules to carry out a task or solve a problem such as “if this, then that.” It has been in practice for a long time and has been the underlying premise for AI until recent advances in machine learning and deep learning, which are techniques within AI.

For example, to create systems to calculate tax through a digitized decision-making process, the Logic and Rules approach employs a detailed series of computer instructions, including if/then/else or probabilities/weights. These detailed series of instructions are at the heart of any AI system and are known as **algorithms**.

Machine Learning is the more prevalent AI approach using algorithms to guide its predictions. Its name is derived from the ability of algorithms to “learn” from experience (i.e., use data sets) rather than relying on a rules-based system. The algorithms create computational models that process large data sets to predict outputs and make inferences. More data leads to more examples; this helps the algorithm to finely tune its output over time. In this way, the algorithm adjusts—or “learns”—by trial and error. Each novel example is an opportunity for the algorithm to guess correctly or incorrectly and learn from its mistakes thus usually providing better insights. The insights are then fed back to refine the algorithmic models further and make them more accurate over time. Machine learning is in use in email spam filters, text, image and voice recognition systems, search rankings, spell check, and many more applications.

A Closer Look at Machine Learning

Machine learning models can broadly be grouped into two types of problems: regression and classification.

Regression problems involve the prediction of a quantity or value (e.g., the taxi fare to go from Santa Monica to downtown LA during rush hour on Friday morning).

Classification problems involve classifying input (e.g., determining whether a transaction is typical or anomalous).

Three different techniques are most commonly used for a machine to “learn” the problem and become smart at providing the answer:

- **Supervised Learning** is a method to teach AI systems by example. The systems are provided with many points of data. Each point is tied to expected outcomes so that an understanding can start to develop of how data relates to the expected outcome. Once trained, the systems can then ingest points of data and give an output that fits the learned model.

Take a classification problem as an example. In the case of rating the credit worthiness of loan applicants, the AI system is provided with data on thousands of past loan applications, each one having detailed information about the prospective borrower (i.e., credit score, income, marital status, etc.). Because the data is historic, the AI system is also supplied with information on which loans defaulted and which were repaid successfully. The AI system then uses statistical techniques to figure out a set of rules to determine when a loan seems to have gone bad and when it did not.

These rules can now be applied to gauge the probability of a new loan application being successfully repaid versus not. Both the input and output training data are labeled so that once trained, the AI system can then apply what it has learned to new data. This process shows that decisions are not based on 100% accuracy but on probabilities instead. IBM Watson's debut on the TV show *Jeopardy* is an example of this. When responding to questions, Watson calculated probabilities of different answers and replied with the highest probability answer. This approach is called supervised learning because an accurately classified training data set is used to supervise the learning process.

- **Unsupervised Learning:** In supervised learning the system is told which loans were repaid successfully and which ones were not and asked to identify what trends most affected the outcome. If only raw insights into the data are wanted, the computer is asked to find insights effectively through unsupervised learning.

For example, Netflix may ask its system to find other customers whose viewing habits are similar to yours. It can then use those results to recommend movies those similar customers liked. Here there are no trained data sets such as a Good Loan and Bad Loan label; it is just about looking at data to find clusters of customers who are "similar," even though the human customers did not specify what they meant by "similar." Detecting fraud and anomaly in transactions often falls in this realm because you are essentially looking for transactions that are "not typical."

- **Reinforcement Learning** is a technique by which an AI system learns under its own supervision without the need of trained data sets. It makes predictions, validates those predictions against reality and continually adjusts itself to have a better output next time. It is inspired by behaviourist psychology and refers to goal or reward-oriented algorithms.

Algorithms are trained to perform tasks by maximizing the rewards or points for the expected favourable outcome and continue to explore new possibilities, essentially through trial and error, until that maximized reward is found. For example, rewards could be linked to increased values in stock picking, evaluating trading strategies, points in a simulation game or beating an opponent in an adversarial game.

Moreover, in DeepMind's AlphaGo Zero the system learned to play simply by playing games against itself, starting from completely random play. It quickly surpassed the human level of play and defeated the previous champion-defeating version of AlphaGo by 100 games to 0.¹¹

11 <https://deepmind.com/blog/alphago-zero-learning-scratch/>

Regression problems are solved mostly with supervised and reinforcement learning while classification problems, depending on specifics, may be answered by using any one of the three learning techniques.

AI Algorithms

There are different strategies used to implement the three different techniques described above. These strategies are at the heart of what allows the AI system to “learn” insights from the data. As a CPA you may never have to deal with technology at this level, but a cursory look may provide some appreciation of what AI algorithms do.


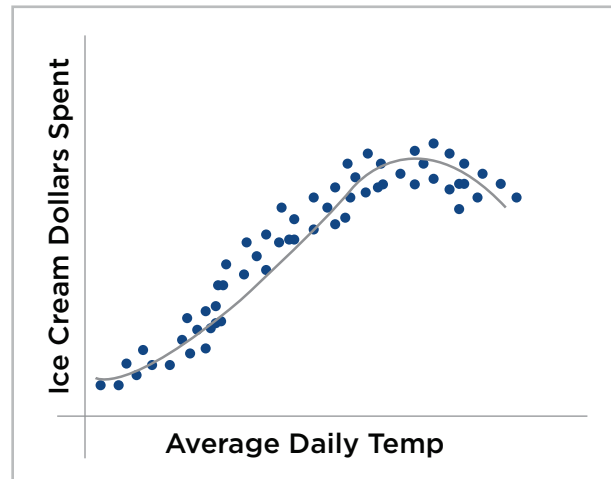
Consider a very simplistic example. Say you have data on sales per day at an ice cream shop as a function of the daily average of temperature. An AI system could identify a regression equation that provides the best fit for the provided data set and use that to predict sales for the next week using projected temperatures.

Mathematically, the problem is now about finding a curve that fits the geometry of the correlation shown above. If the AI system finds the parameters that defined the curve, it can predict any Y (\$) values for a given X (temperature).

In classification problems, the curve would typically divide two classes (e.g., good loans and bad loans).

AI algorithms are essentially all designed to find the parameters that define the curve.

Of course these simple curves are only based on one or two dimensions. The more complex and multidimensional the task, the more complex the AI algorithm.



Source: PwC Analysis

What is the black box of AI?¹²

- As algorithms become more complex, it is not always clear how or why a decision has been made by the AI system.
- The lack of interpretability and transparency of results to human users of the AI system could lead to a lack of trust and loss of confidence in the AI system.

AI and the Human Brain: A Technical Dive into Deep Learning

Deep Learning (also known as hierarchical learning) is a subset of machine learning. It is an emerging and exciting form of AI that can identify relationships and linkages in vast volumes of data that would be impossible for humans to process and apply them to similar situations. While still in the early stages of development, Deep Learning has the potential to take automation to new levels of effectiveness and improve decision making. But how? By using algorithms that roughly approximate the structures and functions of the human brain.

The idea is to create algorithms that can simulate an array of neurons in an artificial **neural network** that learns from vast sources of data. The interconnected layers of nodes or neurons provide data to each successive layer as a nested hierarchy of concepts many layers deep. The connections between these virtual neurons are represented by a number or weight. These weights reflect the importance attributed to the input data; the greater the weight the more important the input is to carrying out the desired task.

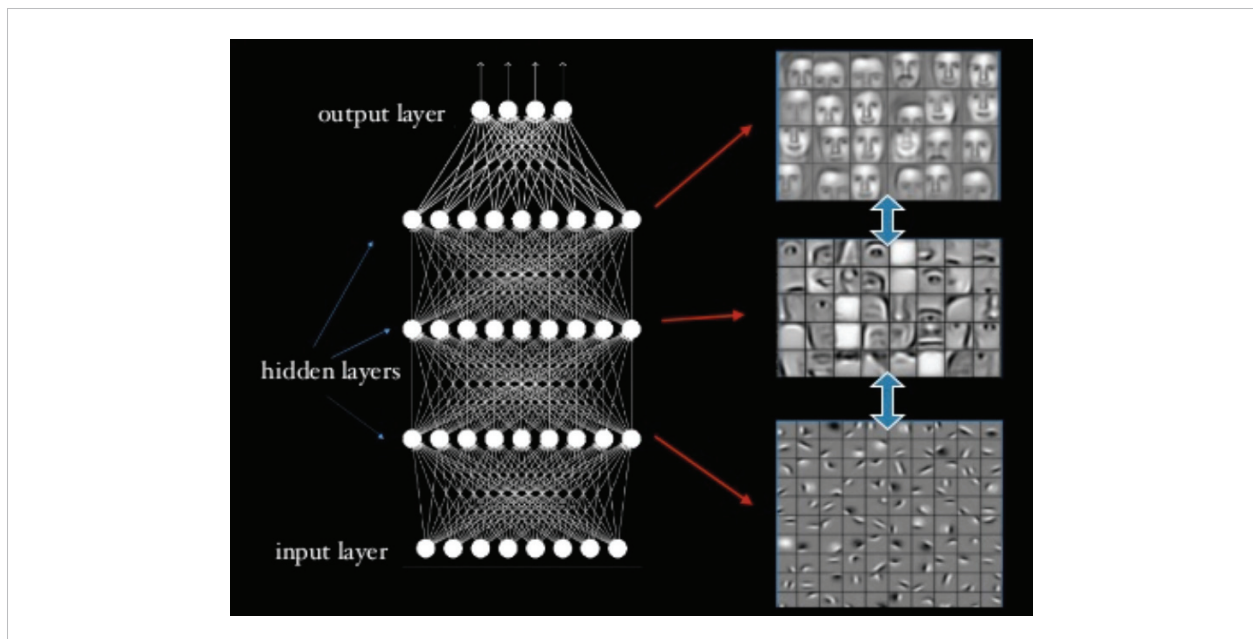
In simple terms, Deep Learning is the stacking of algorithms (both supervised and unsupervised models) on top of each other into multiple layers. This allows for creation of extremely sophisticated and finely tuned input/output models.

¹² <https://usblogs.pwc.com/emerging-technology/to-open-ai-black-box>

During the training of neural networks, weights are varied and adjusted until the desired outcome is arrived at. For example, as shown in the illustration below, when processing thousands of images of human faces, the neural network recognizes objects based on a hierarchy of simpler building blocks:

1. straight lines and curved lines at the basic level
2. then eyes, mouths, and noses
3. then entire faces and
4. finally, specific facial features.

Once the first level learns primitive features, this data is fed to the next layer, which trains itself to recognize more complex features. The process is repeated in successive layers until the system can reliably recognize a nose or a scar. Training these deep learning systems requires vast amounts of data; however, the learnings from these systems can sometimes be replicated to other situations. For example, a neural network trained to recognize facial features can be replicated in multiple situations such as within different security systems.



A neural network is a computer system loosely modelled on the human brain. Source: www.slideshare.net/roelofp/python-for-image-understanding-deep-learning-with-convolutional-neural-nets.

In addition to recognizing images and patterns, deep learning appears to be a promising way to approach complex challenges such as speech comprehension, human-machine conversation, language translation and vehicle navigation.

How AI Is Being Used

The most promising and most advanced applications of AI are in the areas of computer vision, natural language processing and generation, and audio processing.

Computer Vision

Computer vision is a big reason self-driving cars and drones are possible. AI-based pattern detection pulls relevant information from an image or sets of images in order to further group and analyze them. In addition to autonomous cars and drones, computer vision is being used in retail (e.g., Amazon Go store), agriculture (e.g., crop and soil monitoring), medical diagnostics, and for surveillance.

Natural Language Processing (NLP)

Natural Language Processing is everywhere. We interact with NLP in the form of auto-correct on our phones, spell checks, grammatical error detection and when we Google something. In essence it is the ability of a computer system or device to understand spoken or written natural language in order to provide a range of insights or answers.

Natural Language Generation (NLG)

Natural Language Generation transforms visualized data such as charts or graphs into a narrative in an understandable language. Quill by Narrative Science is an example of how NLG can transform insight at the data level into concise, easy-to-read, intelligent reports.¹³ Advancements in NLP and NLG have also given rise to an era of **chatbots**, AI-based programs able to carry on meaningful conversations. Chatbots allow us to have conversations with virtual assistants such as Siri and Alexa. They are increasingly found on company websites and are used for conversational marketing, customer support and post-sales support.

Analyzing Big Data


Thanks to the IoT, social media and the power and reach of cloud and mobile computing, there has been an explosion in the amount of data produced and the amount is only growing. According to the IDC, the world is on track to create 44 zettabytes (or ~44 billion terabytes) of data by 2020. Organizations are tapping into increasingly sophisticated analytics techniques to get closer to customers, set strategy, innovate and grow.

There are four main categories of analytics:

1. **Descriptive Analytics** provide insights into events of the past. Organizations can use this technique to evaluate their annual performance and gather insights.

¹³ <https://narrativescience.com/Resources>

2. **Diagnostic Analytics** examine data to answer why an outcome happened. It is characterized by techniques such as drill-down, data discovery and correlations.
3. **Predictive Analytics** look into the future to anticipate outcomes such as demand forecasting for a supply chain operation. In this case, existing data is used to train machine learning models to forecast what will probably happen.
4. **Prescriptive Analytics** go even further by providing possible outcome solutions that guide predictions into actions such as generating ways to optimize production or inventory. For example, when manufacturing a product, a system may detect a lack of material to complete the order. Prescriptive Analytics will analyze the inputs needed and order or launch a production run for any missing material to complete the manufacturing of the ordered goods.



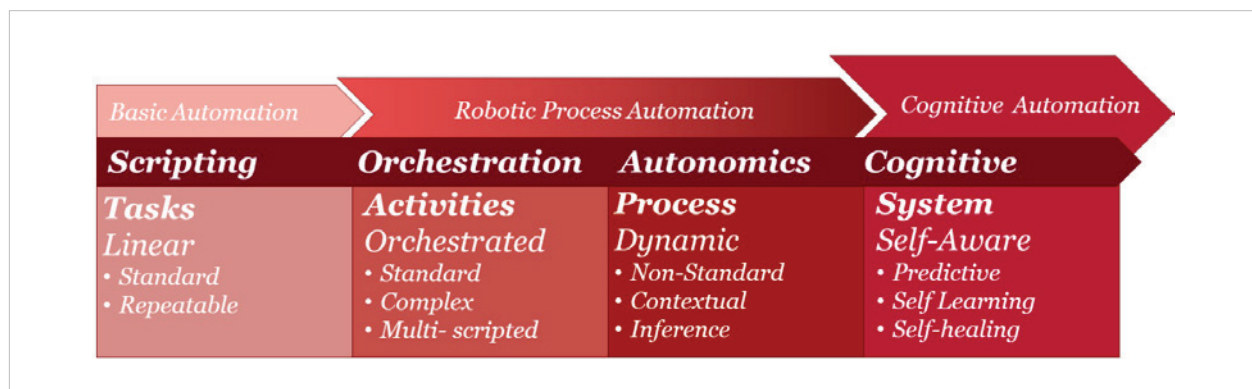
Source: PwC Analysis: Increasing Sophistication of Data & Analytics, Rao 2017. Modified from: Think big: learning contexts, algorithms and data science. Baldassarre, 2016 University of Bari, Italy

Of course, AI needs Big Data just as much as Big Data needs AI because, as has been established, AI needs lots of data to learn.

Why RPA Is Not AI

RPA or **Robotic Process Automation** is a set of capabilities of software automation that can handle high-volume, repeatable tasks such as answering questions, making calculations, maintaining records and recording transactions. All are tasks that previously required humans to perform them. RPA can mimic tasks performed by humans and automate those tasks digitally. RPA systems are easily programmed and modified by non-technical workers. However, implementing dozens of process robots requires significant involvement from human technical specialists. RPA is a waypoint in the journey of organizational automation which started off as basic scripting. For example, basic automation could involve a factory-line robot performing one simple repeatable action. An orchestration activity within RPA could involve downloading a file, taking key figures from the file, transferring those figures to another file and saving the document. All these tasks are actions performed using pre-programmed instructions. No additional decisions need be made and no learning occurs as a result of performing the tasks. Therefore, it is important to note that RPA on its own is not AI. (See the Resource list below which includes additional RPA resources).

As technology progresses, RPA systems are being coupled with algorithms to work with unstructured data related to vision, images and NLP in order to work on more judgment-related activities. For example, a bot in an HR recruitment team may help with background checks from multiple systems such as motor registration, credit agencies, police verification etc. Cognitive automation, which refers to AI techniques applied to automate specific processes, creates more capabilities. For instance, companies are using cognitive automation to perform signature verification (through digital image processing) for contract assurance in telecom operators.



Source: Institute for Robotic Process Automation and Artificial Intelligence

What Next?

If you think AI is impressive now, wait until it is powered by **Quantum computing**. Unlike classic computers which rely on circuits or bits that are either on or off to perform computations and have a computational power equal to 2x the number of bits, Quantum computers are based on quantum mechanics and use qbits (or qubits). These qbits have the ability to be both on and off at the same time. The result: qbit-based computing is inherently parallel (i.e., all possible outcomes from a calculation may be calculated at the same time). This inherent parallelism means quantum computers of sufficient size could very quickly factor large numbers and manipulate vast arrays of numbers to detect subtle patterns, a core requirement in making a powerful AI system. Imagine the possibilities of machine learning when powered by quantum computing! Researchers certainly are.

While the science and potential of AI is exciting, the more businesses explore and embrace AI, the more they become aware of the challenges it poses. A 2017 PwC survey found that 76% of CEOs were concerned with the potential of bias and lack of transparency in AI systems.¹⁴ In order to overcome bias going forward, organizations will need to develop an approach to responsible AI. They will need to establish teams and processes to look for bias in data and models and closely monitor ways malicious actors might “trick” algorithms. On the other side, how will the black box of AI impact the assurance professional’s work effort? It is not just data bias and malicious actors that have to be worried about; think privacy. What are the privacy implications when an AI algorithm suggests a product to a consumer based on the individual’s past preferences and purchasing habits? And what might the fallout be when individuals are categorized for mortgage acceptance or flagged at airport security? These questions and concerns are increasing just as our use of AI in our everyday lives is increasing.

CPA Canada and the AICPA are committed to exploring the implications of AI. Future companion publications will address the impact of AI on CPAs; from CPAs in business to assurance professionals. This publication is just a starting point. It is important that you continue to read, collaborate with colleagues and learn about AI and other technologies which may impact you as a CPA. Everyone is making this digital journey together.

The world of AI is large and growing exponentially. If there are other emerging technologies you would like to learn about or if you would like a deeper dive into the concepts you have read here, contact CPA Canada (yhakimpour@cpacanada.ca) or the AICPA (erin.mackler@aicpa-cima.com).

¹⁴ <http://usblogs.pwc.com/emerging-technology/artificial-intelligence-is-your-business-ready/>

For additional resources on understanding AI, check out:

- [McKinsey Analytics—An executive's guide to AI](#)
- [McKinsey Analytics—Applying AI: Where and how to put Artificial Intelligence to work](#)
- [PwC—Global AI Study—Sizing the prize](#)
- [2017 Deloitte State of Cognitive Survey](#)
- [Deloitte—Internal Controls Over Financial Reporting Considerations for Developing and Implementing Bots](#)
- [EY—Assurance in the age of AI](#)
- [The Economist—The workplace of the future](#)
- [ICAS article— How is the accounting and finance world using Artificial Intelligence?](#)
- [The CPA Journal article — The coming disruption of drones, robots and bots](#)
- [Rise of the Robots \(book\)](#)
- [The Second Machine Age](#)
- [The Singularity is Near: When Humans Transcend Biology](#)
- [Superintelligence: Paths, Dangers, Strategies](#)

About the Authors


CPA Canada and the AICPA would like to express their gratitude to the following professionals for their contributions to this publication: Asif Qayyum, Hrishikesh Bhagwat, and Michael Paterson, CPA, CA of PwC LLP, and Yasmine Hakimpour, CPA, CA of CPA Canada.

CPA Canada and the AICPA would also like to express their gratitude to [CPA Canada's Audit & Assurance Technology Committee](#) and the [AICPA's Assurance Services Executive Committee](#) who assisted in the authoring and review of this publication.

Glossary of Terms

Term	Definition
Artificial Intelligence (AI)	the science of teaching programs and machines to complete tasks that normally require human intelligence
Algorithms	a detailed series of computer instructions such as if/then/else or probabilities/weights for carrying out an operation or solving a problem
Analytics techniques	four main categories: descriptive, diagnostic, predictive, and prescriptive
Big Data	large and complex data sets that cannot be managed with traditional data processing software
Computer vision	ability of a computer system or device to see (e.g., identify and process images)
Deep learning/hierarchical learning	a type of machine learning using algorithms that roughly approximate the structures and functions of the human brain (e.g., algorithms that can simulate an array of neurons in an artificial neural network that learns from vast sources of data)
General AI/ strong AI	human-level intelligence allowing knowledge to be transferred between domains
Internet of Things	interconnection via the Internet of computing devices embedded in everyday objects enabling them to send and receive data
Logic and rules-based approach	uses algorithms to carry out a task or solve a problem
Machine learning	ability of algorithms to learn from experience rather than being provided with instructions. Algorithms create computational models that process large data sets to predict outputs and make inferences; more data leads to more examples, which helps the algorithm to finely tune its output/insight over time; the insights are fed back to further refine the algorithmic models making them more accurate over time
Narrow AI/weak AI	narrowly intelligent systems that can exceed humans in specific tasks such as playing chess but cannot transfer capabilities
Natural language generation (NLG)	ability of a computer system or device to transform visualized data such as charts or graphs into a narrative in an understandable language

Term	Definition
Natural language processing (NLP)	ability of a computer system or device to understand spoken or written natural language
Quantum computing	computer processing using qubits to store an enormous amount of information while using less energy than a classical computer
Reinforcement learning	an AI system that learns under its own supervision by making predictions, validating them against reality and continually adjusting itself for a better output next time
Robotic process automation (RPA)	software automation that can handle high-volume, repeatable tasks such as answering questions, making calculations, maintaining records and recording transactions
Supervised learning	a method to teach AI systems by providing them with the desired outcomes for the training data (where input and output training data are labelled) so they can connect those outputs to the data
Unsupervised learning	algorithms draw inferences from data sets by identifying patterns and looking for similarities by which that data can be grouped



CPA

CHARTERED
PROFESSIONAL
ACCOUNTANTS
CANADA

277 WELLINGTON STREET WEST
TORONTO, ON CANADA M5V 3H2 T.
416 977.3222 F. 416 977.8585
WWW.CPACANADA.CA