Brighterion
mastercard

Artificial intelligence and machine learning: **The next generation**

Contents

Understanding AI

Artificial Intelligence (AI) will soon be at the heart of every major technological system in the world, including payments, compliance, financial markets, security and defense, healthcare, Internet of Things (IoT), and marketing.

While it seems that it's captured most people's attention only recently, AI has actually been around for over 60 years. In the late 1950s, Arthur Samuel wrote a checkers-playing program that could learn from its mistakes and over time became better at playing the game. In the 1970s, MYCIN, the first rule-based expert system, was developed to diagnose blood infections based on the results of various medical tests. The MYCIN system outperformed non-specialist doctors.

Machine learning (ML) is applied in various fields such as computer vision, speech recognition, natural language processing, web search, biotech, risk management, cyber security, and many others. It is the science of getting computers to act without being explicitly programmed, but rather is "programmed by example."

Two types of learning are commonly used: supervised and unsupervised. In supervised learning, a collection of labeled patterns is provided, and the learning process is measured by the quality of labeling a newly encountered pattern. Labeled patterns are used to learn the descriptions of classes, which in turn are used to label a new pattern. In the case of unsupervised learning, the problem is to group a given collection of unlabeled patterns into meaningful categories.

It's important to understand the benefits and shortcomings of AI and ML technologies, and how Brighterion, powered by Smart Agents, is tackling those challenges in real time with its supervised and unsupervised learning.

There are two different types of supervised learning: classification and regression. In classification learning, the goal is to categorize objects into fixed specific categories. Regression learning, on the other hand, tries to predict a real value. For instance, to predict changes in the price of a stock, we may use both methods to derive insights. The classification method determines if the stock price will rise or fall, while the regression method predicts how much the price will increase or decrease.

Now that they are becoming major staples of technology, it's important to understand the benefits and shortcomings of AI and ML technologies, and how Brighterion, powered by Smart Agent technology, is tackling those challenges in real time with its supervised and unsupervised learning. Smart Agents create a virtual representation of every entity of interest, learning and building a profile from each entity's actions and activities. As the engine that drives Brighterion technology, Smart Agents overcome the limits of the legacy machine learning by adapting and updating in real time with every new piece of data. But before we look at how Smart Agents will help your organization manage and deliver intelligence when you need it, we need to understand the basic elements of machine learning.

The tools of Smart Agent technology

The tools of Smart Agent technology

Business rule management system

A *business rule management system* (BRMS) enables companies to easily define, deploy, monitor, and maintain new regulations, procedures, policies, market opportunities, and workflows. One of the main advantages of business rules is that they can be written by business analysts without the need of IT resources. Rules can be stored in a central repository and can be accessed across the enterprise. Rules can be specific to a context, a geographic region, a customer, or a process. Advanced BRMS offers role-based management authority, testing, simulation, and reporting to ensure rules are updated and deployed accurately.

One of the main advantages of business rules is that they can be written by business analysts without the need of IT resources.

Limits in business rule management systems

Business rules represent policies, procedures, and constraints regarding how an enterprise conducts business. Business rules can, for example, focus on the policies of the organization for considering a transaction as suspicious. A fraud expert writes rules to detect suspicious transactions. However, the same rules will also be used to monitor customers whose unique spending behaviors are not accounted for properly in the rule set, resulting in poor detection rates and high false positives. Additionally, risk systems based only on rules detect anomalous behavior associated with just the existing rules; they cannot identify new anomalies which may occur daily. As a result, systems based on rules are outdated almost as soon as they are implemented.

Neural network

A *neural network* (NN) is a technology loosely inspired by the structure of the brain. A neural network consists of many simple elements called artificial neurons, each producing a sequence of activations. The elements used in a neural network are far simpler than biological neurons. The number of elements and their interconnections are orders of magnitude fewer than the number of neurons and synapses in the human brain.

Backpropagation, first described by David Rumelhart in 1986, is the most popular supervised neural network learning algorithm. Backpropagation is organized into layers, and connections between the layers. The leftmost layer is called the input layer. The rightmost, or output, layer contains the output neurons.

Finally, in the middle are the hidden layers. The goal of backpropagation is to compute the gradient (a vector of partial derivatives) of an objective function with respect to the neural network parameters. Input neurons activate through sensors perceiving the environment and other neurons activate through weighted connections from previously active neurons.

Each element receives numeric inputs and transforms this input data by calculating a weighted sum over the inputs. A non-linear function is then applied to this transformation to calculate an intermediate state. While the design of the input and output layers of a neural network is straightforward, there is an art to the design of the hidden layers. Designing and training a neural network requires choosing the number and types of nodes, layers, learning rates, training data, and test sets.

Deep learning

Recently *deep learning*, a new term that describes a set of algorithms that use a neural network as an underlying architecture, has generated many headlines. The earliest deep learning-like algorithms possessed multiple layers of non-linear features and can be traced back to Ivakhnenko and Lapa in 1965. They used thin but deep models with polynomial activation functions which they analyzed using statistical methods.

Deep learning became more usable in recent years due to the availability of inexpensive parallel hardware (GPUs, computer clusters) and massive amounts of data. Deep neural networks learn hierarchical layers of representation from the input to perform pattern recognition. When the problem exhibits non-linear properties, deep networks are computationally more attractive than classical neural networks. A deep network can be viewed as a program in which the functions computed by the lower-layered neurons are subroutines. These subroutines are reused many times in the computation of the final program.

Limits of deep learning

Deep learning is currently one of the main focuses of machine learning. It has led to many speculative comments about AI and its possible impact on the future. Although deep learning garners much attention, people fail to realize that deep learning has inherent restrictions that limit its application and effectiveness in many industries and fields.

Patrick H Winston MIT Deep Neural Nets Lecture

Further examples of the limitations of deep learning are presented by Patrick Henry Winston, the former director of the MIT Artificial Intelligence Laboratory and an Artificial Intelligence professor at the MIT. These examples can be seen at the 44-minute mark of the following video.

Deep learning requires human expertise and significant time to design and train

Deep learning algorithms lack interpretability as they are not able to explain their decision-making. In mission critical applications, such as medical diagnosis, airlines, and security, people must feel confident in the reasoning behind the program. It is difficult to trust systems that do not explain or justify their conclusions.

Another limitation is minimal changes can produce big errors. For example, in vision classification, slightly changing an image that was once correctly classified in a way that is imperceptible to the human eye can cause a deep neural network to label the image as something else entirely.

Data mining

Data mining, or knowledge discovery in databases, is the nontrivial extraction of implicit, previously unknown and potentially useful information from data. Statistical methods are used that enable trends and other relationships to be identified in large databases.

The major reason that data mining has attracted attention is due to the wide availability of vast amounts of data, and the need for turning such data into useful information and knowledge. The knowledge gained can be used for applications ranging from risk monitoring, business management, production control, market analysis, engineering, and science exploration.

In general, three types of data mining techniques are used: association, regression, and classification.

Association analysis

Association analysis is a machine learning method used to discover attribute-value conditions that occur frequently together in a given set of data. Association analysis is widely used to identify the correlation of individual products within shopping carts.

Regression analysis

Regression analysis creates models that explain dependent variables through the analysis of independent variables. As an example, the prediction for a product's sales performance can be created by correlating the product price and the average customer income level.

A deeper dive into deep learning

Additional examples of the limitations of deep learning are explained in a research paper from Cornell and Wyoming Universities titled "Deep Neural Networks are Easily Fooled." Another interesting article is "Deep Learning Isn't a Dangerous Magic Genie. It's Just Math" from Oren Etzioni, a professor of Computer Science and head of the Allen Institute for Artificial Intelligence.

The data mining process consists of an iterative sequence of the following steps:

1. **Data coherence and cleaning to remove noise and inconsistent data.**
2. **Data integration such that multiple data sources may be combined.**
3. **Data selection where data relevant to the analysis are retrieved.**
4. **Data transformation where data are consolidated into forms appropriate for mining.**
5. **Pattern recognition and statistical techniques are applied to extract patterns.**
6. **Pattern evaluation to identify interesting patterns representing knowledge.**
7. **Visualization techniques are used to present mined knowledge to users.**

Classification and prediction

Classification is the process of designing a set of models to predict the class of objects whose class label is unknown. The derived model may be represented in various forms, such as if-then rules, decision trees, or mathematical formulas.

A decision tree is a flowchart-like tree structure where each node denotes a test on an attribute value, each branch represents an outcome of the test, and each tree leaf represents a class or class distribution. Decision trees can be converted to classification rules.

Classification can be used for predicting the class label of data objects. Prediction encompasses the identification of distribution trends based on the available data.

Limits of data mining

GIGO (garbage in, garbage out) is almost always referenced with respect to data mining, as the quality of the knowledge gained through data mining is dependent on the quality of the historical data. We know data inconsistencies and dealing with multiple data sources represent large problems in data management. Data cleaning techniques deal with detecting and removing errors and inconsistencies from data to improve data quality. However, detecting these inconsistencies is extremely difficult. How can we identify a transaction that is incorrectly labeled as suspicious? Learning from incorrect data leads to inaccurate models.

Another limitation is that data mining only extracts knowledge limited to the specific set of historical data, and answers can only be obtained and interpreted with regards to previous trends learned from the data. Because the decision tree is trained specifically on the historical data set, it does not account for personalization within the tree, and limits the ability to learn from new trends. Additionally, data mining (decision trees, rules, clusters) are non-incremental and do not adapt while in production.

Case-based reasoning

Case-based reasoning (CBR) is a problem-solving paradigm that is different from other major AI approaches. CBR learns from past experiences to solve new problems. Rather than relying on a domain expert to write the rules or make associations along generalized relationships between problem descriptors and conclusions, a CBR system learns from previous experience in the same way a

physician learns from his patients. A CBR system will create generic cases based on the diagnosis and treatment of previous patients to determine the disease and treatment for a new patient. The implementation of a CBR system consists of identifying relevant case features. A CBR system continually learns from each new situation. Generalized cases can provide explanations that are richer than explanations generated by chains of rules.

Limits of CBR

The most important limitations relate to how cases are efficiently represented, how indexes are created and how individual cases are generalized.

Fuzzy logic brings a middle ground where statements can be partially true and partially false.

Fuzzy logic

Traditional logic typically categorizes information into binary patterns such as black/white, yes/no, or true/false. *Fuzzy logic* brings a middle ground where statements can be partially true and partially false to account for much of day-to-day human reasoning. For example, stating that a tall person is over 6′ 2″ traditionally means that people under 6′ 2″ are not tall. If a person is nearly 6′ 2″, then common sense says the person is also somewhat tall. Boolean logic states a person is either tall or short and allows no middle ground; fuzzy logic allows different interpretations for varying degrees of height.

Neural networks, data mining, CBR, and business rules can benefit from fuzzy logic. For example, fuzzy logic can be used in CBR to automatically cluster information into categories that improve performance by decreasing sensitivity to noise and outliers. Fuzzy logic also allows business rule experts to write more powerful rules. Here is an example of a rule that has been rewritten to leverage fuzzy logic: When the number of cross border transactions is high and the transaction occurs in the evening, then the transaction may be suspicious.

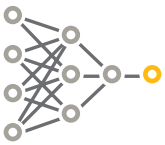

Genetic algorithms

Genetic algorithms work by simulating the logic of Darwinian selection where only the best performers are selected for reproduction. Over many generations, natural populations evolve according to the principles of natural selection.

A genetic algorithm can be thought of as a population of individuals represented by chromosomes. In computing terms, a genetic algorithm implements the model of computation by having arrays of bits or characters (binary string) to represent the chromosomes. Each string represents a potential solution. The genetic algorithm then manipulates the most promising chromosomes searching for improved solutions. A genetic algorithm operates through a cycle of three stages:

1

Build and maintain a population of solutions to a problem.

2

Choose the better solutions for recombination with each other.

3

Use their offspring to replace poorer solutions.

Genetic algorithms provide various benefits to existing machine learning technologies, such as being able to be used by data mining for the field/attribute selection, and can be combined with neural networks to determine optimal weights and architecture.

Real-time, long-term profiling

Brighterion's profiling solution creates hundreds of new features on the fly that are used for scoring. The features generated are derived fields such as grouping, mappings, sets, expressions, and more, creating real-time, long-term profiles that track entity behavior.

Brighterion's profiling uses proprietary algorithms that are database independent and can profile any entity in real time: merchant, agreement, outlet, terminal, merchant segmentations and others. It enables an organization to monitor a wide variety of merchant data such as production purchasing patterns, suspicious changes in activities, number of transactions over a window of time, entity transaction frequency, comparison of transaction versus authorization to detect anomalies, and trends in the number of chargebacks over time. There is no limit to the number of profiling criteria that can be defined.

- **Real-time profiling** – The time window for aggregation can be anywhere from several seconds to several weeks. The counters are updated in real time as transactions are processed.
- **Long-term profiling** – Profile the same entities over a long time period, from a few months to several years. Long-term profiling is used to establish the baseline behaviors for entities, such as users, IP addresses, and devices. At any time, a window can be defined for aggregation and the user can specify the refresh rate to automatically update the profiling values.
- **Recursive profiling** – Gain a full view of user behavior by being able to track and monitor the normal behavior of an entity. An example would be to compute the maximum number of times a user logs on to online banking in a week.
- **Geo-location profiling** – Compute the distance in real time between two zip codes, IP addresses, or other geo-location data to detect abnormal behaviors.
- **Multidimensional profiling** – Profile multiple entity interactions to link suspicious behaviors together and identify unknown entity links.
- **Peer comparison profiling** – Compare entities, such as merchants, to their peers in real time to detect any suspicious activity.

The next generation: Brighterion Smart Agents

Researchers have explored many different architectures for intelligent systems: neural networks, genetic algorithms, business rules, Bayesian network and data mining to name a few. Keeping in mind the most important limits of legacy machine learning techniques, the next generation of artificial intelligence is based on Smart Agents to overcome these limitations.

As we've seen, current AI and ML technologies suffer from various limits. Most importantly, they lack the capacity for:

Personalization: To successfully protect and serve customers, employees and audiences, we must know them by their unique and individual behaviors over time and not by static, generic categorization.

Adaptability: New trends and behaviors arise daily, making it inefficient to rely on models based only on historical data or expert rules.

Self-learning: Over time, an intelligent system should learn from every activity associated to each specific entity.

Controlling fraud and intrusion within a network environment requires a dynamic and continuously evolving process.

The need for autonomous tools

Consider the challenges of two important business fields – fraud prevention and network security – to see how these factors create limitations. Fraud and intrusion are perpetually changing and never remain static. Fraudsters and hackers are criminals who continuously adjust and adapt their techniques. Controlling fraud and intrusion within a network environment requires a dynamic and continuously evolving process. A static set of rules or a machine learning model developed by learning from historical data has only short-term value.

In network security, we know dozens of new malware programs with ever more sophisticated methods of embedding and disguising themselves appear on the internet daily. In most cases, after vulnerabilities are discovered a patch is released to address the vulnerability. The problem is that it is often easy for hackers to reverse engineer the patch, only to find and exploit another defect within hours.

Smart Agents across industries

In the payment industry, for example, a Smart Agent is associated with each individual cardholder, merchant, or terminal. The Smart Agents associated to an entity (such as a card or merchant) learns in real time from every transaction made and builds a profile of their specific and unique behaviors over time. There are as many Smart Agents as there are active entities in the system. So, if there are 200 million cards transacting, there will be 200 million Smart Agents analyzing and learning the behavior of each. Decision-making is specific to each cardholder and no longer relies on logic that is universally applied to all cardholders, regardless of their individual characteristics. Unlike legacy machine learning, the Smart Agents are self-learning and adaptive since they continuously update their individual profiles from each activity and action performed by the entity.

Many well-known malwares (Conficker is an example) exploit vulnerabilities for which there are known patches. They use the fact that the patch frequently is not deployed on vulnerable systems, or is not deployed in a timely manner, leaving open targets. The Aurora attack, originating in China in the fall of 2009 against Google and several other companies, was an example of exploitable dangling pointers in a Microsoft browser that had not been previously discovered.

Tools that autonomously detect new attacks against specific targets, networks or individual computers are vital. They must be able to change their parameters to thrive in new environments, learn from each individual activity, respond to various situations in different ways, and track and adapt to the specific situation/behavior of every entity of interest over time. This continuous, one-to-one behavioral analysis provides real time actionable insights. In addition to the self-learning capability, another key concept for the next generation of AI and ML systems is being reflective. Imagine a plumbing system that autonomously notifies the plumber when it finds water dripping out of a hole in a pipe and detects incipient leaks.

Creating adaptive self-learning with Brighterion

Brighterion's Smart Agents is a personalization technology that creates a virtual representation of every entity and learns/builds a profile from the entity's actions

and activities. Smart Agents is the only technology with the ability to overcome the limits of the legacy machine learning, allowing personalization, adaptability and self-learning.

In an email filtering system, Smart Agents learn to prioritize, delete, forward, and email messages on behalf of a user. They work by analyzing and learning from actions taken by the user. Smart Agents constantly make internal predictions about the actions a user will take on an email. If these predictions prove incorrect, the Smart Agents update their behavior accordingly.

In a financial portfolio management system, a multi-agent system consists of Smart Agents that cooperatively monitor and track stock quotes, financial news, and company earnings reports to continuously monitor and make suggestions to the portfolio manager.

Here is where Brighterion's platform is truly unique. Smart Agents enrich the data to make it possible to use all the technologies mentioned above to your best advantage. Unlike our competitors' products, our patented Smart Agents make unsupervised learning truly functional and optimal for your business needs.

Since they focus on updating the profile based on the actions and activities of the entity, Smart Agents store only relevant information and intelligence rather than the raw incoming data they are analyzing, which achieves enormous compression in storage.

Intelligent, self-learning

Smart Agents do not rely on pre-programmed rules and do not try to anticipate every possible scenario. Instead, Smart Agents create profiles specific to each entity and behave according to their goals, observations, and the knowledge they continuously acquire through their interactions with other Smart Agents. Each Smart Agent pulls all relevant data across multiple channels, irrespective of the type or format and source of the data, to produce robust virtual profiles. Each profile is automatically updated in real-time and the resulting intelligence is shared across the Smart Agents. This one-to-one behavioral profiling provides unprecedented, omni-channel visibility into the behavior of an entity.

Smart Agents can represent any entity and enable best-in-class performance with minimal operational and capital resource requirements. Smart Agents automatically validate the coherence of the data, and perform the features learning, data enrichment and one-to-one profiles creation. Since they focus on updating the profile based on the actions and activities of the entity, Smart Agents store only relevant information and intelligence rather than the raw incoming data they are analyzing, which achieves enormous compression in storage.

Unlimited scalability, resistant to disruption

Legacy technologies in machine learning generally rely on databases, which use tables to store structured data; however, tables cannot store knowledge or behaviors. Smart Agents bring a powerful, distributed file system specifically

designed to store knowledge and behaviors. This distributed architecture allows lightning speed response times (below 1 millisecond) on entry-level servers as well as end-to-end encryption and traceability. The distributed architecture allows for unlimited scalability and resilience to disruption as it has no single point of failure.

Next generation Artificial Intelligence and Machine Learning must be intelligent, self-learning and adaptive. For more information about Smart Agents technology, contact Brighterion to discuss how to make your organization's solution scalable, resistant to disruption, and continually adaptive.