

An Introduction to Artificial Intelligence

This book is a beginner's guide to important topics in AI, machine learning, and deep learning. The goal is to give readers an intuition for how powerful new algorithms work and how they are used, along with code examples where possible.

Advances in the field of machine learning (algorithms that adjust themselves when exposed to data) are driving progress more widely in AI. But many terms and concepts will seem incomprehensible to the intelligent outsider, the beginner, and even the former student of AI returning to a transformed discipline after years away. We hope this book helps you better understand AI, the software used to build it, and what is at stake in its development.

The line between mathematics and philosophy is blurry when we talk about artificial intelligence, because with AI, we ask the mineral called silicon to perceive and to think – actions once thought exclusive to meat, and now possible with computation. We hope that by reading this book, you will find new ways of thinking about life and intelligence, just as we have by writing it.

INTRODUCTION	1
DEEPLARNING, MACHINE LEARNING & AI USE CASES	2
Feature Introspection	2
Text	2
Images	3
ARTIFICIAL INTELLIGENCE (AI) VS. MACHINE LEARNING VS. DEEP LEARNING	3
Artificial Intelligence, Symbolic AI and GOF AI	3
Machine Learning: Programs That Alter Themselves	4
Deep Learning: More Accuracy, More Math & More Compute	5
What's Next for AI?	5
CONCLUSION	6

DEEP LEARNING, MACHINE LEARNING & AI USE CASES

Deep learning excels at identifying patterns in unstructured data, which most people know as media such as images, sound, video and text. Below is a list of sample use cases we've run across, paired with the sectors to which they pertain.

General use case	Industry
Sound	
Voice recognition	UX/UI, Automotive, Security, IoT
Voice search	Handset maker, Telecoms
Sentiment analysis	CRM
Flaw detection (engine noise)	Automotive, Aviation
Fraud detection (latent audio artifacts)	Finance, Credit Cards
Time Series	
Log analysis/Risk detection	Data centers, Security, Finance
Enterprise resource planning	Manufacturing, Auto., Supply chain
Predictive analysis using sensor data	IoT, Smart home, Hardware manufact.
Business and Economic analytics	Finance, Accounting, Government
Recommendation engine	E-commerce, Media, Social Networks
Text	
Sentiment Analysis	CRM, Social media, Reputation mgt.
Augmented search, Theme detection	Finance
Threat detection	Social media, Govt.
Fraud detection	Insurance, Finance
Image	
Facial recognition	
Image search	Social media
Machine vision	Automotive, aviation
Photo clustering	Telecom, Handset makers
Video	
Motion detection	Gaming, UX, UI
Real-time threat detection	Security, Airports

Feature Introspection

Traditional machine learning has the advantage of feature introspection – that is, it knows why it is classifying an input in one way or another, which is important for analytics. But that very advantage is what excludes it from working with unlabeled and unstructured data, as well as attaining the record-breaking accuracy of the latest deep learning models. Feature engineering is one of the major chokepoints of traditional machine learning, since so few people are able to do it well and quickly enough to adapt to changing data.

For cases where feature introspection is necessary (e.g. the law requires that you justify a decision to, say, deny a loan due to predicted credit risk), we recommend using a deep net in an ensemble with machine-learning algorithms, allowing each one to vote and relying on each for its strength. Alternatively, one can perform various analyses on the results of deep nets to form hypotheses about their decision-making.

Text

NAMED-ENTITY RESOLUTION

One use of deep-learning networks is [named-entity recognition](#), which is a way to extract from unstructured, unlabeled data certain types of information like people, places, companies or things. That information can then be stored in a structured schema to build, say, a list of addresses or serve as a benchmark for an identity validation engine.

SPEECH-TO-TEXT

With the proper data transforms, a deep network is capable of understanding audio signals. This can be used to identify snippets of sound in larger audio files, and transcribe the spoken word as text.

Image

OBJECT RECOGNITION

Object recognition is an algorithm's ability to identify arbitrary objects – such as balls, animals, or even faces – within larger images. This is typically used in engineering applications to identify shapes for modeling purposes. It's also used by social networks for photo tagging. Facebook's Deep Face is a good example of a deep-learning application in this realm.

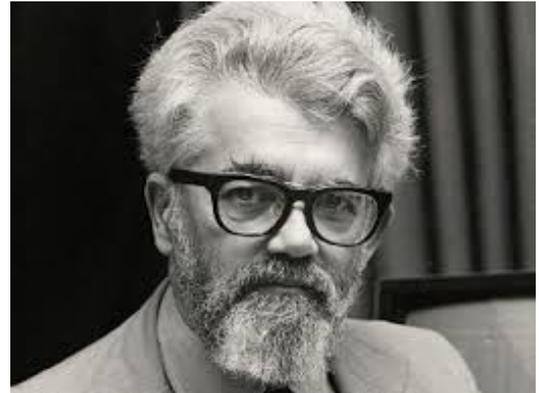
MACHINE VISION + NATURAL LANGUAGE PROCESSING

Advances in reality capture and reality computing are making virtual and real worlds converge. One application of deep learning to this newly available data is to recognize and label objects in 3D environments, and in real life. From there, it's a short step to simulated semantics, in which machines learn the nature and constraints of objects in the world, through their virtual representations, and then bring that understanding to the language they generate and ingest. We believe that is one of many futures in store for neural nets.

ARTIFICIAL INTELLIGENCE (AI) VS. MACHINE LEARNING VS. DEEP LEARNING

You can think of deep learning, machine learning and artificial intelligence as a set of Russian dolls nested within each other, beginning with the smallest and working out. Deep learning is a subset of machine learning, and machine learning is a subset of AI, which is an umbrella term for any computer program that does something smart. In other words, all machine learning is AI, but not all AI is machine learning, and so forth.

Artificial Intelligence, Symbolic AI and GOFAI



[John McCarthy](#), widely recognized as one of the godfathers of AI, defined it as "[the science and engineering of making intelligent machines.](#)"

Here are a few other definitions of artificial intelligence:

- A branch of computer science dealing with the simulation of intelligent behavior in computers.
- The capability of a machine to imitate intelligent human behavior.
- A computer system able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

There are a lot of ways to simulate human intelligence, and some methods are more intelligent than others.

AI can be a pile of if-then statements, or a complex statistical model mapping raw sensory data to symbolic categories. The if-then statements are simply rules explicitly programmed by a human hand. Taken together, these if-then statements are sometimes called [rules engines, expert systems, knowledge graphs or symbolic AI](#). Collectively, these are known as [Good, Old-Fashioned AI \(GOFAI\)](#).

The intelligence that rules engines mimic could be that of an accountant with knowledge of the tax code, who takes information you feed it, runs the information through a set of static rules, and gives you the amount of taxes you owe as a result. In the US, we call that TurboTax.

Usually, when a computer program designed by AI researchers actually succeeds at something – like winning at chess – many people say it’s “not really intelligent”, because the algorithm’s internals are well understood. The critics think intelligence must be something intangible, and exclusively human. A wag would say that true AI is whatever computers can’t do yet.

Machine Learning: Programs That Alter Themselves

Machine learning is a subset of AI. That is, all machine learning counts as AI, but not all AI counts as machine learning. For example, symbolic logic – rules engines, expert systems and knowledge graphs – could all be described as AI, and none of them are machine learning.

One aspect that separates machine learning from the knowledge graphs and expert systems is its ability to modify itself when exposed to more data; i.e. machine learning is dynamic and does not require human intervention to make certain changes. That makes it less brittle, and less reliant on human experts.

“ **A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E.**

–Tom Mitchell

In 1959, Arthur Samuel, one of the pioneers of machine learning, defined machine learning as a “field of study that gives computers the ability to learn without being explicitly programmed.” That is, machine-learning programs have not been explicitly entered into a computer, like the if-then statements above. Machine-learning programs, in a sense, adjust themselves in response to the data they’re exposed to (like a child that is born knowing nothing adjusts its understanding of the world in response to experience).



Samuel taught a computer program to play checkers. His goal was to teach it to play checkers better than himself, which is obviously not something he could program explicitly. He succeeded, and in 1962 his program beat the checkers champion of the state of Connecticut.

The “learning” part of machine learning means that ML algorithms attempt to optimize along a certain dimension; i.e. they usually try to minimize error or maximize the likelihood of their predictions being true. This has three names: an error function, a loss function, or an objective function, because the algorithm has an objective. When someone says they are working with a machine-learning algorithm, you can get to the gist of its value by asking: What’s the objective function?

How does one minimize error? Well, one way is to build a framework that multiplies inputs in order to make guesses as to the inputs’ nature. Different outputs/guesses are the product of the inputs and the algorithm. Usually, the initial guesses are quite wrong, and if you are lucky enough to have ground-truth labels pertaining to the input, you can measure how wrong your guesses are by contrasting them with the truth, and then use that error to modify your algorithm. That’s what [neural networks](#) do. They keep on measuring the error and modifying their parameters until they can’t achieve any less error.

They are, in short, an optimization algorithm. If you tune them right, they minimize their error by guessing and guessing and guessing again.

Deep Learning: More Accuracy, More Math & More Compute

Deep learning is a subset of machine learning. Usually, when people use the term deep learning, they are referring to [deep artificial neural networks](#), and somewhat less frequently to [deep reinforcement learning](#).

Deep artificial neural networks are a set of algorithms that have set new records in accuracy for many important problems, such as [image recognition](#), sound recognition, recommender systems, [natural language processing](#) etc. For example, deep learning is part of DeepMind's well-known AlphaGo algorithm, which beat the former world champion Lee Sedol at Go in early 2016, and the current world champion Ke Jie in early 2017. A more complete explanation of neural works is [here](#).

Deep is a technical term. It refers to the number of layers in a neural network. A shallow network has one so-called *hidden layer*, and a deep network has more than one. Multiple hidden layers allow deep neural networks to learn features of the data in a so-called feature hierarchy, because simple features (e.g. two pixels) recombine from one layer to the next, to form more complex features (e.g. a line). Nets with many layers pass input data (features) through more mathematical operations than nets with few layers, and are therefore more computationally intensive to train. Computational intensity is one of the hallmarks of deep learning, and it is one reason why a new kind of chip call GPUs are in demand to train deep-learning models.

So you could apply the same definition to deep learning that Arthur Samuel did to machine learning – a “field of study that gives computers the ability to learn without being explicitly programmed” – while adding that it tends to result in higher accuracy, require more hardware or training time, and perform exceptionally well on machine perception tasks that involved unstructured data such as blobs of pixels or text.

What's Next for AI?

The advances made by researchers at DeepMind, Google Brain, OpenAI and various universities are accelerating. AI is capable of solving harder and harder problems better than humans can.

This means that AI is changing faster than its history can be written, so predictions about its future quickly become obsolete as well. Are we chasing a breakthrough like nuclear fission (possible), or are attempts to wring intelligence from silicon more like trying to turn lead into gold?

There are four main schools of thought, or churches of belief if you will, that group together how people talk about AI.

Those who believe that AI progress will continue apace tend to think a lot about [strong AI](#), and whether or not it is good for humanity. Among those who forecast continued progress, one camp emphasizes the benefits of more intelligent software, which may save humanity from its current stupidities; the other camp worries about the existential risk of a superintelligence.

Given that the power of AI progresses hand in hand with the power of computational hardware, advances in computational capacity, such as better chips or quantum computing, will set the stage for advances in AI. On a purely algorithmic level, most of the astonishing results produced by labs such as DeepMind come from combining different approaches to AI, much as AlphaGo combines deep learning and reinforcement learning. Combining deep learning with [symbolic reasoning](#), [analogical reasoning](#), Bayesian and [evolutionary methods](#) all show promise.

Those who do not believe that AI is making that much progress relative to human intelligence are forecasting another [AI winter](#), during which funding will dry up due to generally disappointing results, as has happened in the past. Many of those people have a pet algorithm or approach that competes with deep learning.

Finally, there are the pragmatists, plugging along at the math, struggling with messy data, scarce AI talent and user acceptance. They are the least religious of the groups making prophecies about AI – they just know that it's hard.

Footnotes

It's interesting to note that even when certain technologies are physically impossible, they can still be regulated. In 1404, during the reign of Henry IV, the English parliament passed a law called the "Act against multiplication", outlawing the creation of gold and silver from other materials, since that was seen as a threat to the throne's control over the currency. The law was later modified to allow only certain people to create gold and silver through alchemical processes, until it was finally repealed in the 17th century. Regulations outlawing strong AI, a technology that may or may not be possible, and for which there exists no strong theoretical foundation, would be similarly absurd.

SkyMind offers AI infrastructure that enables corporate data scientists and IT teams to rapidly prototype, deploy, maintain, and retrain machine learning workflows that accelerate time to value. SkyMind bridges the gap between data science, DevOps and the big data stack.

