

Deep Learning Limitations and Flaws

Bahman Zohuri^{1*} and Masoud Moghaddam²

¹Electrical and Computer Engineering Department, University of New Mexico, Albuquerque, New Mexico, USA

²Galaxy Advance Engineering Director and Consultant, Albuquerque, USA

***Corresponding author:** Bahman Zohuri, University of New Mexico, Electrical Engineering and Computer Science Department, Albuquerque, and Galaxy Advanced Engineering (CEO), New Mexico, USA

Received:  January 21, 2020

Published:  January 29, 2020

Abstract

With today's growing interest toward Artificial Intelligence (AI) and its augmentation as part of integrated business from banking to eCommerce, medical applications and others, we are getting more and more dependency on AI in our day to day operations. However, the most sophisticated AI or Super AI (SAI) still needs to rely on its two other integrated sub-sets of components namely, Machine Learning (ML) and Deep Learning (DL). However, there certain limitation and flaws that exists within DL component of AI or SAI that will cause and error to grow way beyond control and will impact its main master component namely AI and SAI for its final processing of data and information in a trusted way of process of the precise decision making and prediction of event including any forecasting as part of its Use Case (UC) and Service Level Agreement (SLA) assigned as task to the AI or SAI. This article points few of these limitations and flaws that presently are concerns of the scientists and engineers behind the artificial intelligence technologies momentum. These sorts of limitation and flaws also would impact the Business Resilience System (BRS) from perspective of resiliency built into your daily business operations that is also pointed it out in this article.

Keywords: Resilience System; Business Intelligence; Artificial Intelligence; Super Artificial Intelligence; Image Processing; Cyber Security; Decision Making in Real Time; Machine Learning; Deep Learning

Introduction

Today's campaign and momentum that is built by engineers and scientists involved with Artificial Intelligence (AI) is at its pick of innovative approach., yet the entire AI system is not hundred percent bolt proof, yet there may exist certain flaws and limitations within sub-component of such intelligence system that we are depending on and associating our businesses around it. Application of AI is going anywhere from Medical and Health to eCommerce and Banking to give us a better chance as stakeholder and decision making at top of the organization management to perform our tasks and duties much more efficient, thus we need a flawless system that has no chance of making any error or mistake, so we can rely on it for our final call for the final process going forward in order to be more efficient. As we have learned in today's world of Artificial Intelligence (AI) a fully integrated AI system has dependency on two other sub-system and we know them as Machine Learning (ML) and Deep Learning (DL) as depicted in (Figure 1) below [1]. One of the main driving sub-layer that provides the right information from incoming data either structured or unstructured at shear volume is Deep Learning (DL), where this sub-layer will process and compares the present incoming data with historical data that

is deposited in our database in order to disorientate the right date to pass them on to the layer of Machine Learning (ML) and then the ML to Artificial Intelligence (AI). In return AI is processing this information as a knowledge learned from it to do right forecasting and prediction analytics for our next efficient move toward an efficient yet decisive decision making for our daily operation without any doubt in our mind as managers. Any simple error or flaw in any of this information will be escalated at larger scale and then the chain of decision-making managers who are in charge of their operations. Such proper and correct and not doubtful decision will make our system to be very resilience and cost effective [2,3]. This very important step and an important factor if we are going to rely in particular on an unsupervised AI as illustrated and depicted in (Figure 2).

Deep Learning Limitation

As we have learned so far in introduction of this paper and the (Figure 1, 2), the most sophisticated Artificial Intelligence (AI) of present time or futuristic Super Artificial Intelligence (SAI) with its integrated and sophisticated Neural Network (NN) type

human brain of future time, they both relay and depend on their two important sub-systems of Machine Learning (ML) and Deep Learning (DL) as their main driving elements. See (Figure 3). As per Figure 3 depiction, Neural Network (NN) models of AI process signals by sending them through a network of nodes analogous to neurons. Signals pass from node to node along links, analogs of the synaptic junctions between neurons. "Learning" improves the outcome by adjusting the weights that amplify or damp the signals each link carries. Nodes are typically arranged in a series of layers that are roughly analogous to different processing centers in the cortex. Today's computers can handle Deep Learning (DL)

networks with dozens of layers. However, given the dependency of AI performance on ML and consequently DL has raised a widely shared concerns and sentiment among AI developing engineers and practitioners, any of whom can easily rattle off a long list of deep learning's drawbacks and flaws. In addition to its vulnerability to spoofing, for example, there is its gross inefficiency. "For a child to learn to recognize a cow," says Hinton, "it's not like their mother needs to say 'cow' 10,000 times" -a number that's often required for deep-learning systems. Humans generally learn new concepts from just one or two examples [4].

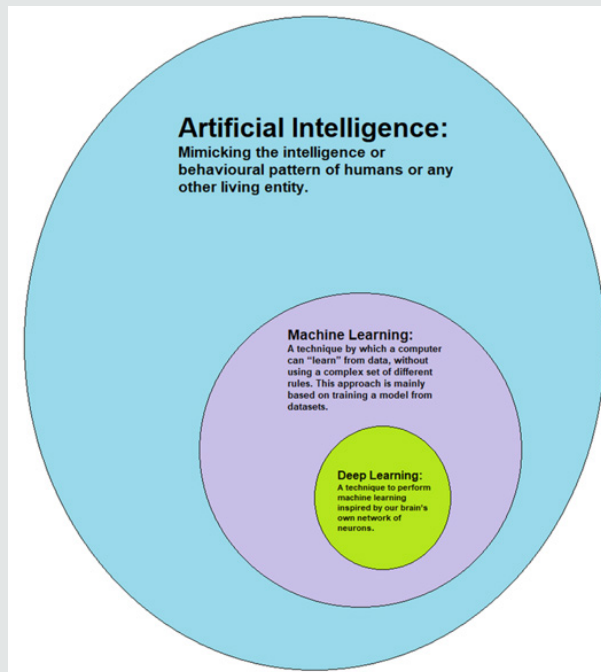


Figure 1: Depiction of Data, Information, and Knowledge is Power in Four Dimensions.

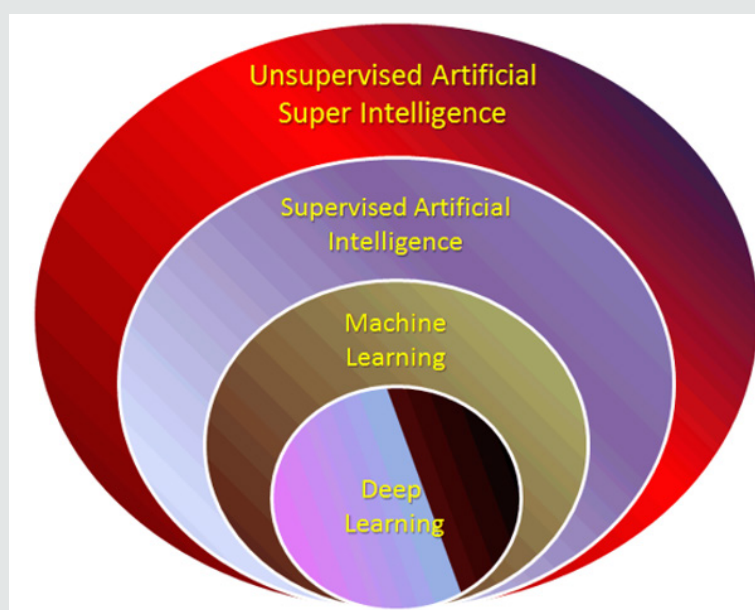


Figure 2: Schematic of A Simple Super Artificial Intelligence.

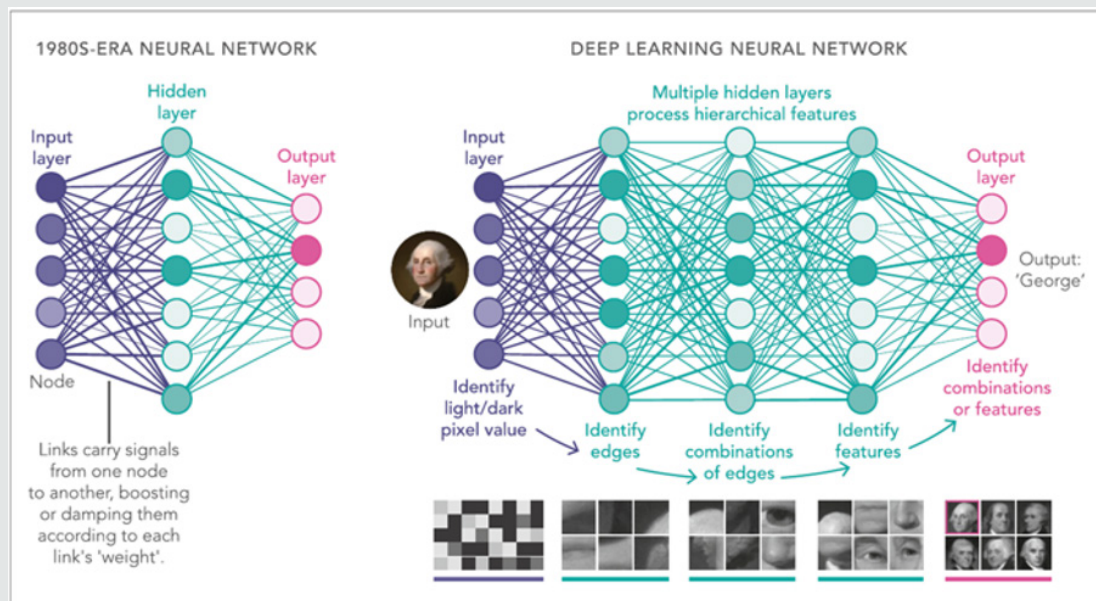


Figure 3: Typical Sophisticated Neural Network Models of AI (Image Credit: Artist Lucy Reading-Ikkanda).

Then there's the opacity problem. Once a deep learning system has been trained, it's not always clear how it's making its decisions. "In many contexts that's just not acceptable, even if it gets the right answer," says David Cox, a computational neuroscientist who heads the MIT-IBM Watson AI Lab in Cambridge, MA. Suppose a bank uses AI to evaluate your creditworthiness and then denies you a loan: "In many states there are laws that say you have to explain why," he says. And perhaps most importantly, there's the lack of common sense. Deep-learning systems may be wizards at recognizing patterns in the pixels, but they can't understand what the patterns mean, much less reason about them. "It's not clear to me that current systems would be able to see that sofas and chairs are for sitting," says Greg Wayne, an AI researcher at DeepMind, a London-based subsidiary of Google's parent company, Alphabet [4]. Increasingly, such frailties are raising concerns about AI among the wider public, as well -especially as driverless cars, which use similar deep-learning techniques to navigate, get involved in well-publicized mishaps and fatalities. "People have started to say, 'Maybe there is a problem,'" says Gary Marcus, a cognitive scientist at New York University and one of deep learning's most vocal skeptics. Until the past year or so, he says, "there had been a feeling that deep learning was magic. Now people are realizing that it's not magic". Still, there's no denying that deep learning is an incredibly powerful tool -one that's made it routine to deploy applications such as face and voice recognition that were all but impossible just a decade ago. "So, I have a hard time imagining that deep learning will go away at this point," Cox says. "It is much more likely that we will modify it or augment it" [4].

As PNAS states or asks – What are the limits of Deep Learning? [4], indicates, there are certain deep learning flaws and they are briefly listed here:

- The systems need 10,000+ examples to learn a concept like cows. Humans only need a handful of examples.

- Deep Learning cannot explain how the systems got an answer.

- Deep Learning lacks common sense. This makes the systems fragile and when errors are made, the errors can be very large.

These are part of concerns and thus, there is a growing feeling in the field that deep learning's shortcomings require some fundamentally new ideas. One solution is simply to expand the scope of the training data. In an article published in May 2018, Botvinick's DeepMind group studied 24 what happens when a network is trained on more than one task. They found that as long as the network has enough "recurrent" connections running backward from later layers to earlier ones -a feature that allows the network to remember what it's doing from one instant to the next -it will automatically draw on the lessons it learned from earlier tasks to learn new ones faster. This is at least an embryonic form of human style "meta-learning," or learning to learn, which is a big part of our ability to master things quickly. A more radical possibility is to give up trying to tackle the problem at hand by training just one big network and instead have multiple networks work in tandem. In June 2018, the DeepMind team 25 published an example they call the Generative Query Network architecture, which harnesses two different networks to learn its way around complex virtual environments with no human input.

One, dubbed the representation network, essentially uses standard image-recognition learning to identify what's visible to the AI at any given instant. The generation network, meanwhile, learns to take the first network's output and produce a kind of 3D model of the entire environment -in effect, making predictions about the objects and features the AI does not see. For example, if a table only has three legs visible, the model will include a fourth leg with the same size, shape, and color. A more radical approach

is to quit asking the networks to learn everything from scratch for every problem. A potentially powerful new approach is known as the graph network. These are deep-learning systems that have an innate bias toward representing things as objects and relations. In conclusion, generally speaking we can say that, the most surprising thing about deep learning is how simple it is. Ten years ago, no one expected that we would achieve such amazing results on machine perception problems by using simple parametric models trained with gradient descent. Now, it turns out that all you need is sufficiently large parametric models trained with gradient descent on sufficiently many examples. As Feynman once said about the universe, "It's not complicated, it's just a lot of it". In deep learning, everything is a vector, i.e. everything is a point in a geometric space. Model inputs (it could be text, images, etc.) and targets are first "vectorized", i.e. turned into some initial input vector space and target vector space. Each layer in a deep learning model operates one simple geometric transformation on the data that goes through it. Together, the chain of layers of the model forms one very complex geometric transformation, broken down into a series of simple ones. This complex transformation attempts to map the input space to the target space, one point at a time. This transformation is parametrized by the weights of the layers, which are iteratively updated based on how well the model is currently performing. A key characteristic of this geometric transformation is that it must be differentiable, which is required in order for us to be able to learn its parameters via gradient descent. Intuitively, this means that the geometric morphing from inputs to outputs must be smooth and continuous - a significant constraint.

The whole process of applying this complex geometric transformation to the input data can be visualized in 3D by imagining a person trying to uncrumple a paper ball: the crumpled paper ball is the manifold of the input data that the model starts with. Each movement operated by the person on the paper ball is similar to a simple geometric transformation operated by one layer. The full uncrumpling gesture sequence is the complex transformation of the entire model. Deep learning models are mathematical machines for uncrumpling complicated manifolds of high-dimensional data. That is the magic of deep learning: turning meaning into vectors, into geometric spaces, then incrementally learning complex geometric transformations that map one space to another. All you need are spaces of sufficiently high dimensionality in order to capture the full scope of the relationships found in the original data. Thus, the limitation of Deep Learning (DL) in the space of its applications that can be implemented with this simple strategy is nearly infinite, and yet many more applications are completely out of reach for current deep learning techniques - even given vast amounts of human-annotated data. Say, for instance, that you could assemble a dataset of hundreds of thousands - even millions - of English language descriptions of the features of a software product, as written by a product manager, as well as the corresponding source code developed by a team of engineers to meet these requirements.

Even with this data, you could not train a deep learning model to simply read a product description and generate the appropriate codebase. That's just one example among many. In general, anything that requires reasoning - like programming, or applying the scientific method - long-term planning, and algorithmic-like data manipulation, is out of reach for deep learning models, no matter how much data you throw at them. Even learning a sorting algorithm with a deep neural network is tremendously difficult.

This is because a deep learning model is "just" a chain of simple, continuous geometric transformations mapping one vector space into another. All it can do is map one data manifold X into another manifold Y , assuming the existence of a learnable continuous transform from X to Y , and the availability of a dense sampling of $X:Y$ to use as training data. So even though a deep learning model can be interpreted as a kind of program, inversely most programs cannot be expressed as deep learning models - for most tasks, either there exists no corresponding practically-sized deep neural network that solves the task, or even if there exists one, it may not be learnable, i.e. the corresponding geometric transform may be far too complex, or there may not be appropriate data available to learn it. Scaling up current deep learning techniques by stacking more layers and using more training data can only superficially palliate some of these issues. It will not solve the more fundamental problem that deep learning models are very limited in what they can represent, and that most of the programs that one may wish to learn cannot be expressed as a continuous geometric morphing of a data manifold. Moreover, the risk of anthropomorphizing Machine Learning (ML) models does exist as well.

One very real risk with contemporary Artificial Intelligence (AI) is that of misinterpreting what deep learning models do and overestimating their abilities. A fundamental feature of the human mind is our "theory of mind", our tendency to project intentions, beliefs and knowledge on the things around us. Drawing a smiley face on a rock suddenly makes it "happy" - in our minds. See Figure 4 and as it can be seen it is not obvious for DL to know if the kid is holding a baseball bat or toothbrush on first image process glance. Applied to deep learning, this means that when we are able to somewhat successfully train a model to generate captions to describe pictures, for instance, we are led to believe that the model "understands" the contents of the pictures, as well as the captions it generates. We then proceed to be very surprised when any slight departure from the sort of images present in the training data causes the model to start generating completely absurd captions. In particular, this is highlighted by "adversarial examples", which are input samples to a deep learning network that are designed to trick the model into misclassifying them. You are already aware that it is possible to do gradient ascent in input space to generate inputs that maximize the activation of some convnet filter, for instance - this was the basis of the filter visualization technique we introduced in Chapter 5 of the book on Deep Learning with Python [6], as well as the Deep Dream algorithm from Chapter 8 of the same book.



Figure 4: The Boy is Holding a Baseball Bat.

Similarly, through gradient ascent, one can slightly modify an image in order to maximize the class prediction for a given class. By taking a picture of a panda and adding to it a “gibbon” gradient, we can get a neural network to classify this panda as a gibbon. This evidence both the brittleness of these models, and the deep difference between the input-to-output mapping that they operate and our own human perception. See the illustration in (Figure 5). In short, deep learning models do not have any understanding of their input, at least not in any human sense. Our own understanding of images, sounds, and language is grounded in our sensorimotor experience as humans -as embodied earthly creatures. Machine learning models have no access to such experiences and thus cannot “understand” their inputs in any human-relatable way. By

annotating large numbers of training examples to feed into our models, we get them to learn a geometric transform that maps data to human concepts on this specific set of examples, but this mapping is just a simplistic sketch of the original model in our minds, the one developed from our experience as embodied agents -it is like a dim image in a mirror. (Figure 6) As a machine learning practitioner, always be mindful of this, and never fall into the trap of believing that neural networks understand the task they perform -they don't, at least not in a way that would make sense to us. They were trained on a different, far narrower task than the one we wanted to teach them: that of merely mapping training inputs to training targets, point by point. Show them anything that deviates from their training data, and they will break in the most absurd ways.

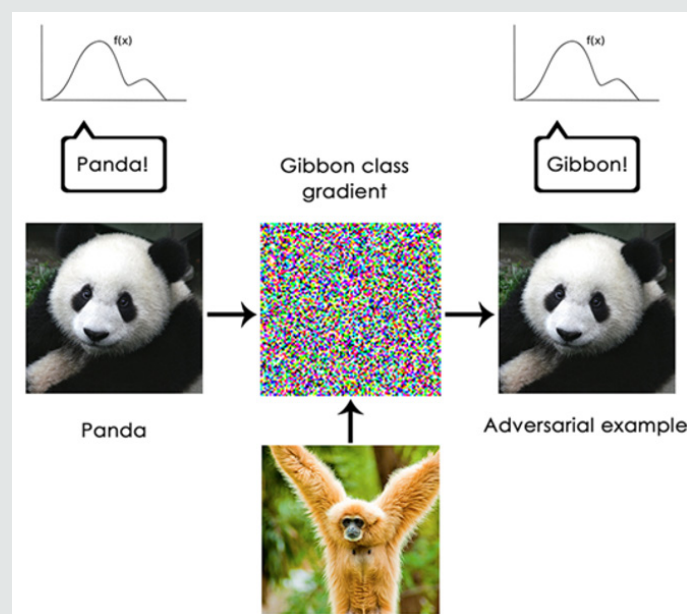


Figure 5: Gibbson Class Gradient with Panda and Monkey Involved (Image Credited: Francois Chollet).

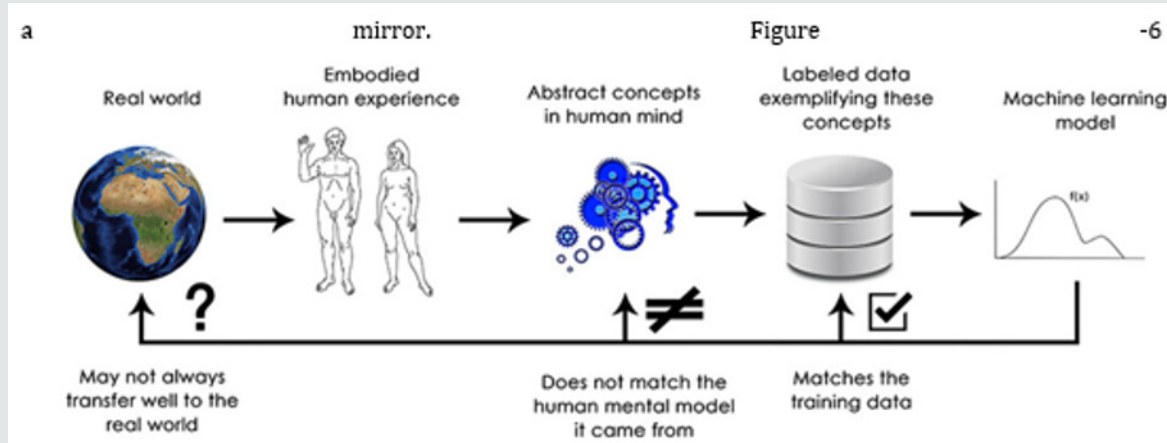


Figure 6: Machine Learning Model Illustration (Image Credited: Francois Chollet).

Local Generalization Versus Generalization

There just seems to be fundamental differences between the straightforward geometric morphing from input to output that deep learning models do, and the way that humans think and learn. It isn't just the fact that humans learn by themselves from embodied experience instead of being presented with explicit training examples. Aside from the different learning processes, there is a fundamental difference in the nature of the underlying representations. Humans are capable of far more than mapping immediate stimuli to immediate responses, like a deep net, or maybe an insect, would do. They maintain complex, abstract models of

their current situation, of themselves, of other people, and can use these models to anticipate different possible futures and perform long-term planning. They are capable of merging together known concepts to represent something they have never experienced before -like picturing a horse wearing jeans, for instance, or imagining what they would do if they won the lottery. This ability to handle hypotheticals, to expand our mental model space far beyond what we can experience directly, in a word, to perform abstraction and reasoning, is arguably the defining characteristic of human cognition. We call it "extreme generalization": an ability to adapt to novel, never experienced before situations, using very little data or even no new data at all.

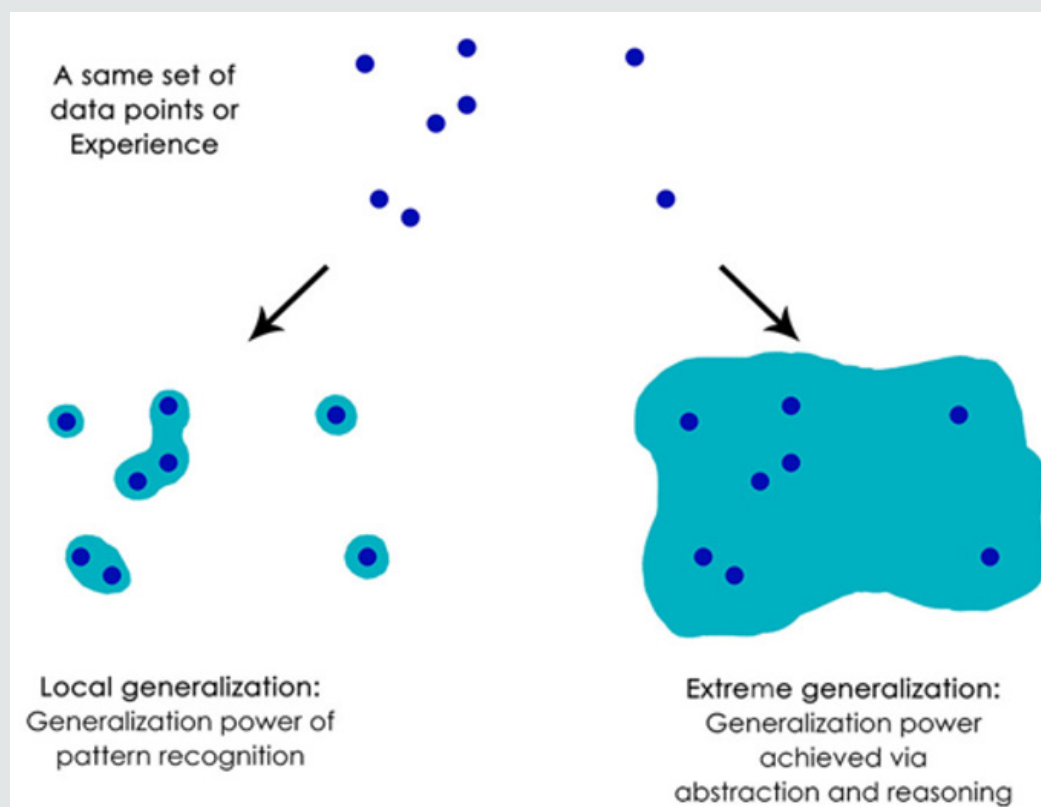


Figure 7: Local Generalization Vs. Extreme Generalization Illustration (Image Credited: Francois Chollet).

This stands in sharp contrast with what deep nets do, which I would call “local generalization”: the mapping from inputs to outputs performed by deep nets quickly stops making sense if new inputs differ even slightly from what they saw at training time. Consider, for instance, the problem of learning the appropriate launch parameters to get a rocket to land on the moon. If you were to use a deep net for this task, whether training using supervised learning or reinforcement learning, you would need to feed it with thousands or even millions of launch trials, i.e. you would need to expose it to a dense sampling of the input space, in order to learn a reliable mapping from input space to output space. By contrast, humans can use their power of abstraction to come up with physical models -rocket science -and derive an exact solution that will get the rocket on the moon in just one or few trials. Similarly, if you developed a deep net controlling a human body and wanted it to learn to safely navigate a city without getting hit by cars, the net would have to die many thousands of times in various situations until it could infer that cars are dangerous and develop appropriate avoidance behaviors. Dropped into a new city, the net would have to relearn most of what it knows. On the other hand, humans are able to learn safe behaviors without having to die even once -again, thanks to their power of abstract modeling of hypothetical situations. See Figure 7.

Summary

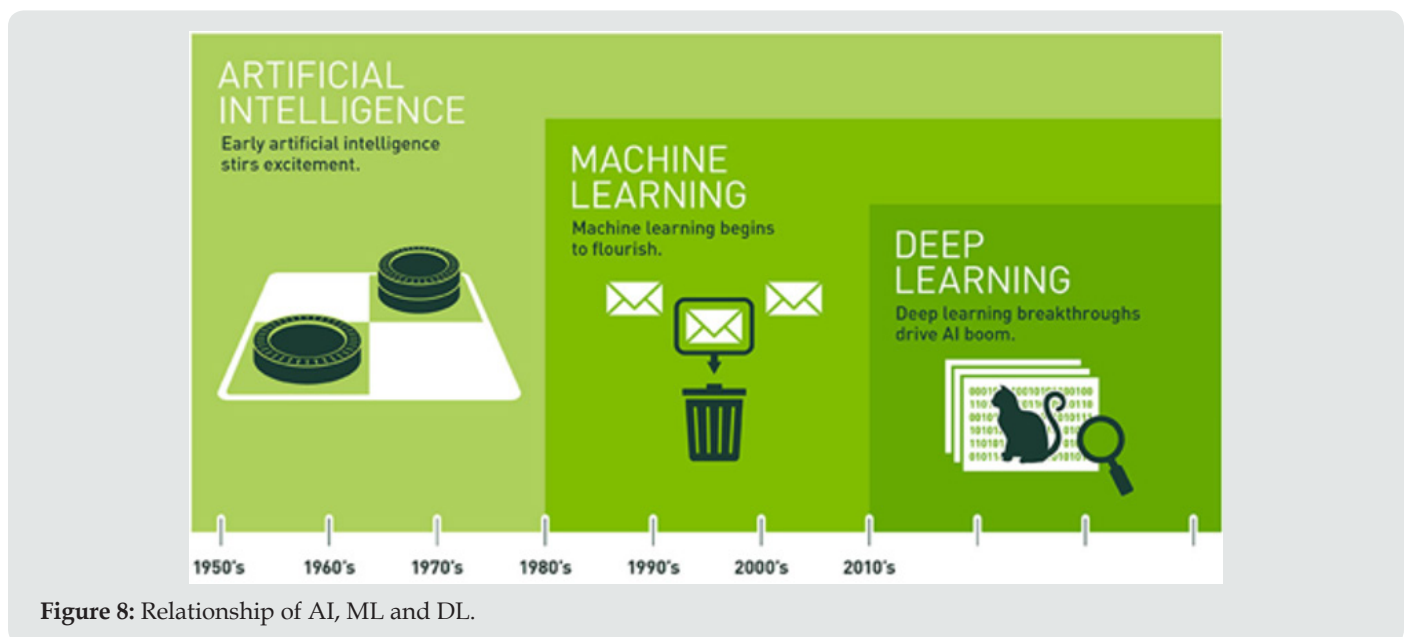


Figure 8: Relationship of AI, ML and DL.

Deep learning is a type of machine learning that trains a computer to perform human-like tasks, such as recognizing speech, identifying images or making predictions. Instead of organizing data to run through predefined equations, deep learning sets up basic parameters about the data and trains the computer to learn on its own by recognizing patterns using many layers of processing. In summary, Deep learning is a subset of AI and machine learning that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in tasks such as object detection, speech recognition, language translation and others. As Figure 3,2,6

In short, despite our progress on machine perception, we are still very far from human-level AI: our models can only perform local generalization, adapting to new situations that must stay very close from past data, while human cognition is capable of extreme generalization, quickly adapting to radically novel situations, or planning very for long-term future situations. In summary, the take-away from this discussion is as follow. Here is what you should remember: the only real success of deep learning so far has been the ability to map space X to space Y using a continuous geometric transform, given large amounts of human-annotated data. Doing this well is a game-changer for essentially every industry, but it is still a very long way from human-level AI. To lift some of these limitations and start competing with human brains, we need to move away from straightforward input-to-output mappings, and on to reasoning and abstraction. A likely appropriate substrate for abstract modeling of various situations and concepts is that of computer programs. We have said before (Note: in Deep Learning with Python) [5] that machine learning models could be defined as “learnable programs”; currently we can only learn programs that belong to a very narrow and specific subset of all possible programs. But what if we could learn any program, in a modular and reusable way. You can read for more details in this link [6].

shows, since an early flush of optimism in the 1950s, smaller subsets of a Deep learning differs from traditional machine learning techniques in that they can automatically learn representations from data such as images, video or text, without introducing hand-coded rules or human domain knowledge. Their highly flexible architectures can learn directly from raw data and can increase their predictive accuracy when provided with more data. Deep learning is responsible for many of the recent breakthroughs in AI such as Google DeepMind’s AlphaGo, self-driving cars, intelligent voice assistants and many more. With NVIDIA GPU-accelerated

deep learning frameworks, researchers and data scientists can significantly speed up deep learning training, that could otherwise take days and weeks to just hours and days. When models are ready for deployment, developers can rely on GPU-accelerated inference platforms for the cloud, embedded device or self-driving cars, to deliver high-performance, low-latency inference for the most computationally intensive deep neural networks. See Figure 8.

Deep learning is one of the foundations of Artificial Intelligence (AI), and the current interest in deep learning is due in part to the buzz surrounding AI. Deep learning techniques have improved the ability to classify, recognize, detect and describe – in one word, understand. For example, deep learning is used to classify images, recognize speech, detect objects and describe content. Systems such as Siri and Cortana are powered, in part, by deep learning.

Several developments are now advancing deep learning:

- a) Algorithmic improvements have boosted the performance of deep learning methods.
- b) New machine learning approaches have improved accuracy of models.
- c) New classes of neural networks have been developed that fit well for applications like text translation and image classification.
- d) We have a lot more data available to build neural networks with many deep layers, including streaming data from the Internet of Things (IoT), textual data from social media, physicians note and investigative transcripts.
- e) Computational advances of distributed cloud computing and graphics processing units have put incredible computing power at our disposal. This level of computing power is necessary to train deep algorithms.

At the same time, human-to-machine interfaces have evolved greatly as well. The mouse and the keyboard are being replaced with gesture, swipe, touch and natural language, ushering in a renewed interest in AI and deep learning. A lot of computational power is needed to solve deep learning problems because of the iterative nature of deep learning algorithms, their complexity as the number of layers increase, and the large volumes of data needed to train the networks. The dynamic nature of deep learning methods-their ability to continuously improve and adapt to changes in the underlying information pattern-presents a great opportunity to introduce more dynamic behavior into analytics. Greater personalization of customer analytics is one possibility. Another great opportunity is to improve accuracy and performance in applications where neural networks have been used for a long time. Through better algorithms and more computing power, we can add greater depth. While the current market focus of deep learning techniques is in applications of cognitive computing, there is also great potential in more traditional analytics applications, for example, time series analysis.

Another opportunity is to simply be more efficient and streamlined in existing analytical operations. Recently, companies

like SAS experimented with deep neural networks in speech-to-text transcription problems. Compared to the standard techniques, the word-error-rate decreased by more than 10 percent when deep neural networks were applied. They also eliminated about 10 steps of data preprocessing, feature engineering and modeling. The impressive performance gains and the time savings when compared to feature engineering signify a paradigm shift. As a conclusion of this summary, Deep Learning (DL) is a machine learning technique that teaches computers to do what comes naturally to humans: learn by example. Deep learning is a key technology behind driverless cars, enabling them to recognize a stop sign, or to distinguish a pedestrian from a lamppost. It is the key to voice control in consumer devices like phones, tablets, TVs, and hands-free speakers. Deep learning is getting lots of attention lately and for good reason. It's achieving results that were not possible before. In deep learning, a computer model learns to perform classification tasks directly from images, text, or sound. Deep learning models can achieve state-of-the-art accuracy, sometimes exceeding human-level performance. Models are trained by using a large set of labeled data and neural network architectures that contain many layers [7].

In a word, accuracy. Deep learning achieves recognition accuracy at higher levels than ever before. This helps consumer electronics meet user expectations, and it is crucial for safety-critical applications like driverless cars. Recent advances in deep learning have improved to the point where deep learning outperforms humans in some tasks like classifying objects in images. While deep learning was first theorized in the 1980s, there are two main reasons it has only recently become useful:

- A. Deep learning requires large amounts of labeled data. For example, driverless car development requires millions of images and thousands of hours of video.
- B. Deep learning requires substantial computing power. High-performance GPUs have a parallel architecture that is efficient for deep learning. When combined with clusters or cloud computing, this enables development teams to reduce training time for a deep learning network from weeks to hours or less.

Examples of Deep Learning (DL) at work listed below as:

- a. Deep learning applications are used in industries from automated driving to medical devices.
- b. Automated Driving: Automotive researchers are using deep learning to automatically detect objects such as stop signs and traffic lights. In addition, deep learning is used to detect pedestrians, which helps decrease accidents.
- c. Aerospace and Defense: Deep learning is used to identify objects from satellites that locate areas of interest and identify safe or unsafe zones for troops.
- d. Medical Research: Cancer researchers are using deep learning to automatically detect cancer cells. Teams at UCLA built an advanced microscope that yields a high- dimensional data set used to train a deep learning application to accurately identify cancer cells.

e. Industrial Automation: Deep learning is helping to improve worker safety around heavy machinery by automatically detecting when people or objects are within an unsafe distance of machines.

f. Electronics: Deep learning is being used in automated hearing and speech translation. For example, home assistance devices that respond to your voice and know your preferences are powered by deep learning applications.




Figure 9: Networks with Hidden Layers.

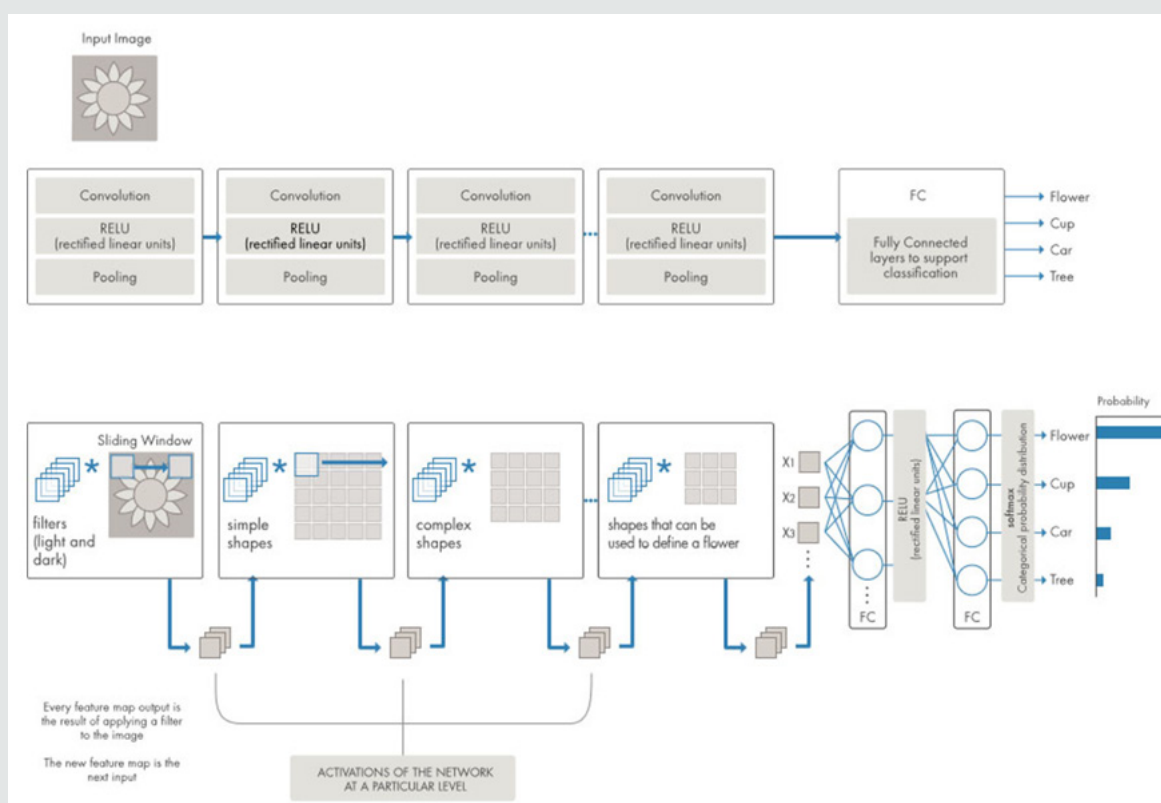


Figure 10: Network with Many Convolutional Layers Illustration.

Most deep learning methods use Neural Network (NN) architectures, which is why deep learning models are often referred to as Deep Neural Networks (DNNs). The term “deep” usually refers to the number of hidden layers in the neural network. Traditional neural networks only contain 2-3 hidden layers, while deep networks can have as many as 150. As we know, deep learning models are trained by using large sets of labeled data and neural network

architectures that learn features directly from the data without the need for manual feature extraction. See Figure 9, where Neural Networks, which are organized in layers consisting of a set of interconnected nodes. Networks can have tens or hundreds of hidden layers. One of the most popular types of deep neural networks is known as convolutional neural networks (CNN or ConvNet). A CNN convolves learned features with input data, and uses 2D convolu-

tional layers, making this architecture well suited to processing 2D data, such as images. CNNs eliminate the need for manual feature extraction [8], so you do not need to identify features used to classify images. The CNN works by extracting features directly from images. The relevant features are not pretrained; they are learned while the network trains on a collection of images. This automated feature extraction makes deep learning models highly accurate for computer vision tasks such as object classification. Figure 10 is an example of a network with many convolutional layers. Filters are applied to each training image at different resolutions, and the output of each convolved image serves as the input to the next layer. CNNs learn to detect different features of an image using tens or hundreds of hidden layers. Every hidden layer increases the complexity of the learned image features. For example, the first hidden layer could learn how to detect edges, and the last learns how to detect more complex shapes specifically catered to the shape of the object we are trying to recognize.

References

1. Zohuri B, Mossavar RF (2019) Artificial Intelligence Driven Resiliency with Machine Learning and Deep Learning Components. Journal of Communication and Computer 1-13.
2. A Model to Forecast Future Paradigms: Volume 1: Introduction to Knowledge Is Power in Four Dimensions, Apple Academic Press, a CRC Press, Taylor & Francis Group.
3. Zohuri B, Moghaddam M (2017) Business Resilience System (BRS): Driven Through Boolean, Fuzzy Logics and Cloud Computation: Real and Near Real Time Analysis and Decision-Making System 1st (edn.). Springer Publishing Company.
4. Zohuri B, Moghaddam M (2019) Artificial Intelligence Driven Resiliency with Machine Learning and Deep Learning Components. Japan Journal of Research 1(1): 2.
5. François Cholet (2017) Deep Learning with Python. pages printed in black & white. 384pp.
6. Francois Chollet (2017) the future of deep learning. The Keras Blog 3(9).
7. Mitchell MW (2019) What are the limits of deep learning? Pnas license 116: 4.




This work is licensed under Creative Commons Attribution 4.0 License

To Submit Your Article Click Here:

[Submit Article](#)

DOI: [10.32474/MAMS.2020.02.000138](https://doi.org/10.32474/MAMS.2020.02.000138)



Modern Approaches on Material Science

Assets of Publishing with us

- Global archiving of articles
- Immediate, unrestricted online access
- Rigorous Peer Review Process
- Authors Retain Copyrights
- Unique DOI for all articles