

ARTIFICIAL INTELLIGENCE AND INDUSTRY

# HOW MACHINE LEARNING IS IMPACTING INDUSTRIES

DISRUPTOR REPORT 2019  
CHAPTER 3

PRESCOUTER


# ABOUT THE AUTHORS



---

**Fatima Al-Raisi** is a Project Architect at PreScouter and a PhD researcher in the field of artificial intelligence at the School of Computer Science, Carnegie Mellon University. She works on content modeling and planning using optimization, data mining, machine learning, and text analytics. She is a Computer Science lecturer at Sultan Qaboos University (Muscat, Oman). Fatima taught computer science and discrete mathematics courses in different higher education institutions and worked in academic planning and quality assurance.

She also worked as an application engineer in Petroleum Development Oman. Her current research focuses on developing algorithms and applying machine learning, and data mining techniques for finding and generating coherent content and connecting components in unstructured data. She holds a MS degree in (Theoretical) Computer Science from the Rochester Institute of Technology and a BS degree in Computer Science with a minor in mathematics from Sultan Qaboos University, Oman. She is the author of an award-winning book in Arabic titled "Letter to a Scholar (abroad)." Fatima's research background includes machine learning, data mining, natural language processing/computational linguistics, complexity theory, and theoretical computer science.



---

**Mohamed Akrouf** earned a B.S. in Computer Science and Statistics from University of Montreal and an M.S. in Artificial Intelligence from University of Toronto. His master's topic is about applying reinforcement learning algorithms to perform skin disease symptoms checking that outperform the dermatologist-level accuracy. His research interests lie between artificial intelligence and healthcare applications. He is currently a research scientist at Triage, a startup working between AI and dermatology to make

the medical diagnosis accessible to everyone. His research focuses on translating several learning mechanisms from the brain to mathematical techniques in order to build the next generation of AI algorithms.

# Table of Contents

1	AI: HISTORICALLY AND TODAY	1
2	DEEP LEARNING	9
3	TRANSFER LEARNING	14
4	ANOMALY DETECTION	18
5	MULTIMODAL MACHINE LEARNING	20
6	ONE-SHOT AND ZERO-SHOT LEARNING	22
7	CHALLENGES AND OPPORTUNITIES	24
8	INTERVIEW WITH ZACH LIPTON	26
*	END NOTES	31
PS	ABOUT PRESCOUTER	33

# 1

# AI: Historically and Today

The term “artificial intelligence” (AI) was first coined by John McCarthy in 1956 during a summer workshop called the Dartmouth Summer Research Project on Artificial Intelligence. The workshop was held to discuss what would become the transformative field of AI. McCarthy defined AI as “the science and engineering of making intelligent machines<sup>1</sup>.” One of the key features of modern AI is algorithms based on empirical data. These algorithms can appear to emulate human performance typically by learning, coming to their own conclusions, appearing to understand complex content, engaging in natural dialogs with people, enhancing human cognitive performance (also known as cognitive computing), or replacing people in the execution of nonroutine tasks<sup>2</sup>.

Intelligence, and hence artificial intelligence, is characterized by five main properties:

1. Learning: Improving with experience and data (**evolve**)
2. Reasoning: Drawing inferences (**justify**)
3. Problem-Solving: Given input (data), find the target output(s) or patterns/relations (**act**)
4. Perception: Identifying and computing features, spaces, and transformations (**discover**)
5. Linguistic Capabilities: Identifying the “grammar” of a solution (**infer**).

Artificial intelligence is historically rooted in logic, including symbolic reasoning and first-order logic. Machine learning (ML) is rooted in statistical pattern recognition. ML is basically a way of achieving AI. Artificial intelligence was historically achieved through rule-based systems. Machine learning techniques are empirical techniques that attempt to learn from data.

The scientific history of artificial intelligence conceals a very old rivalry between its two main currents, that of Symbolists and Connectionists. The Symbolists included a host of researchers working on solving theorems such as Herbert Gelernter (creator of a geometric theorem demonstration system in 1959), Allen Newell and Herbert Simon (inventors of the General Problem Solver in 1957), John McCarthy (inventor of the name of artificial intelligence in 1955 as well as the LISP language in 1958), James Slagle (designer of the first expert system, SAINT, of treatment of mathematical formulas, in 1961), Thomas Evans (and his program ANALOGY of 1963 which could solve problems of tests of IQ) and also the Frenchmen Alain Colmerauer and Philippe Roussel (inventors of the language Prolog in 1972). These were the beginnings of attempts to emulate human reasoning in machines. In the camp of the connectionists, we had Frank Rosenblatt especially with his perceptrons, the premises of neural networks of today. Their technical limitations raised in 1969 by Marvin Minsky coupled with the weak material capabilities available prevented them from breaking through before the 21st century. Today, most state of the art ML techniques rely on neural networks to model perception and cognition by explicitly modelling the data processing that is believed to occur in the brain.

More recently, expert knowledge is being leveraged by incorporating domain knowledge into ML approaches to combine both data-driven and expert-designed solutions. The English Oxford Living Dictionary provides a more application oriented definition of AI as “the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”


Today, examples of the state of the art in AI include: ResNet<sup>3</sup>, a deep neural model for image recognition that outperforms humans<sup>4</sup> in object recognition on the ImageNet dataset; ML-based speech recognition software with an error rate that is lower than human typing<sup>5</sup>; Baidu speech synthesis models that can generate human-like sounding speech in real time<sup>6</sup>; digital pathology systems that are exceedingly successful at identifying features that accurately predict the stage of cancer<sup>7</sup> and can automatically identify skin cancer as well as dermatologists can<sup>8</sup>; and the first AI-based news anchor<sup>9</sup> which was recently introduced in China. The list of machine learning success stories will expand geometrically as development in this space continues to advance.

AI and ML technologies have already been applied in a wide range of industries and domains including education, healthcare, finance and governance, entertainment, construction, transportation, logistics, and many others. AI and ML have been notably successful in automation of digital and physical tasks using automated workflows and robotics technology, respectively. They have also been instrumental in deriving insights from the huge amounts of text and numeric data being produced by various systems using natural language processing and text mining tools. Further, AI and ML have begun to play an enormous role in decision-making support through the development of tools using summarization, information retrieval, and data mining and techniques to predict (and prevent or drive) various outcomes and events using classification, regression, and clustering techniques. In the consumer experience and marketing space, AI and ML can support personalization of different experiences using collaborative filtering techniques and recommendation systems, and optimizing processes like marketing and sales using longitudinal data, signal processing, and time series analysis techniques.

The availability of large amounts of data, both user and system-generated, has enabled the use of machine learning techniques to learn from large swaths of training data in order to solve problems in different domains. The successful application of AI and ML in different industries depends on a number of factors. These include variation in availability of data; the scale and quality of available data; increasing emphasis on issues related to privacy and sensitivity of data, especially data about human subjects; bias in data; the sensitivity of the application itself to the importance of correct interpretation of the used models; and accurate verification of results.

**THE PERCENTAGE OF JOBS  
REQUIRING AI SKILLS HAS  
GROWN 450% SINCE 2013  
ACCORDING TO STANFORD  
2017 AI INDEX REPORT.**

FIG 1.




Left: Annual venture capital investment in AI startups. Right: Number of startups developing AI systems

The global artificial intelligence market, which was valued at 16.06 billion USD in 2017, is expected to exceed 190 billion USD by 2025<sup>10</sup>. This market is projected to grow at a compounded annual growth rate (CAGR) of 37%; one of the largest CAGR numbers reported in recent years. PricewaterhouseCoopers research reports that 72% of business leaders believe artificial intelligence provides a business advantage<sup>11</sup>. One of the challenges facing these leaders will be to find the best-fitted solution and experts for their particular need.

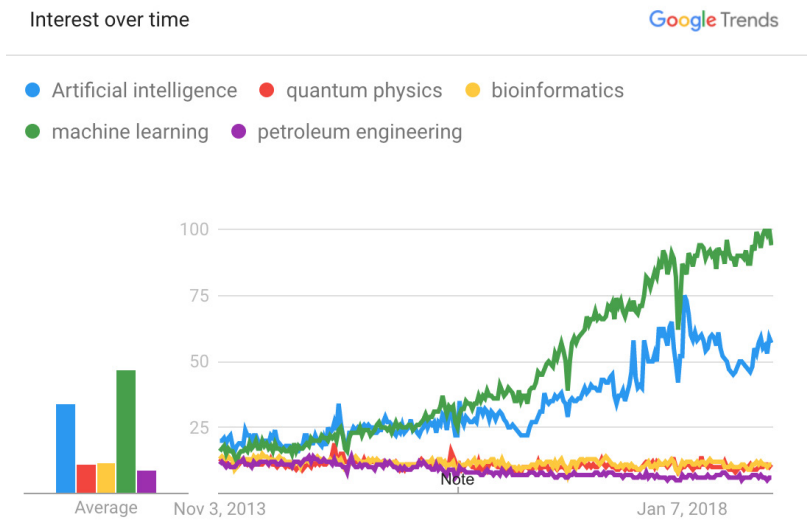
This rapid growth pattern is also seen in the number of **AI startups around the world**, as well as in the expansion of artificial intelligence and machine learning research focused on developing new tools and algorithms for tackling new and challenging problems. The percentage of jobs requiring AI skills has grown 450% since 2013.

FIG 2.



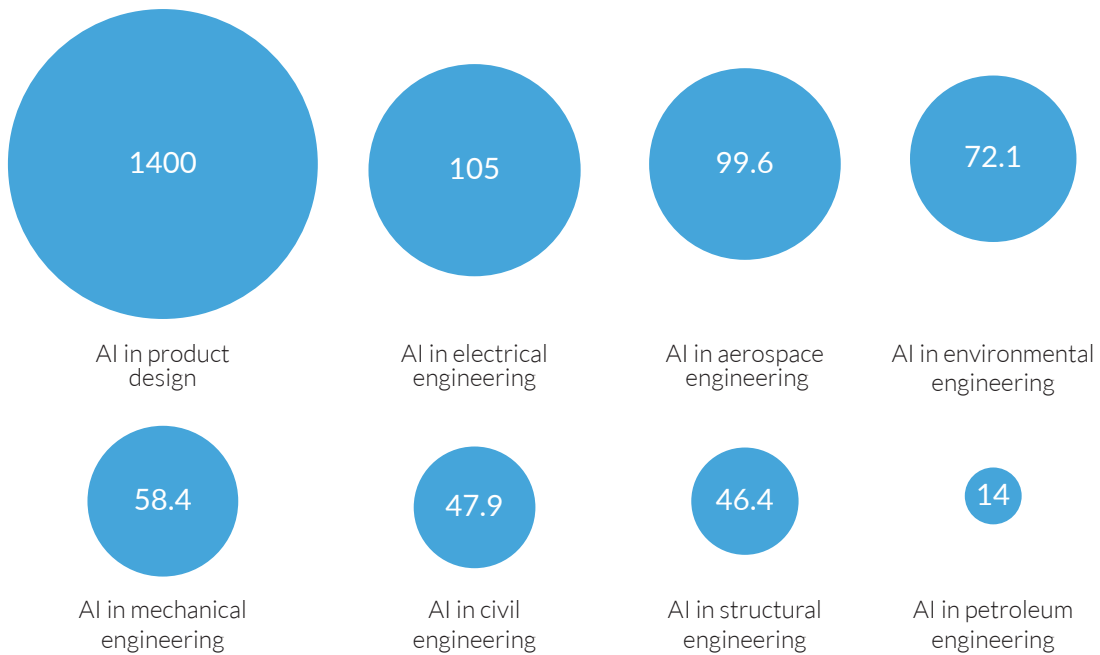
Number of machine learning papers published on e-print repository arXiv<sup>13</sup>

FIG 3.



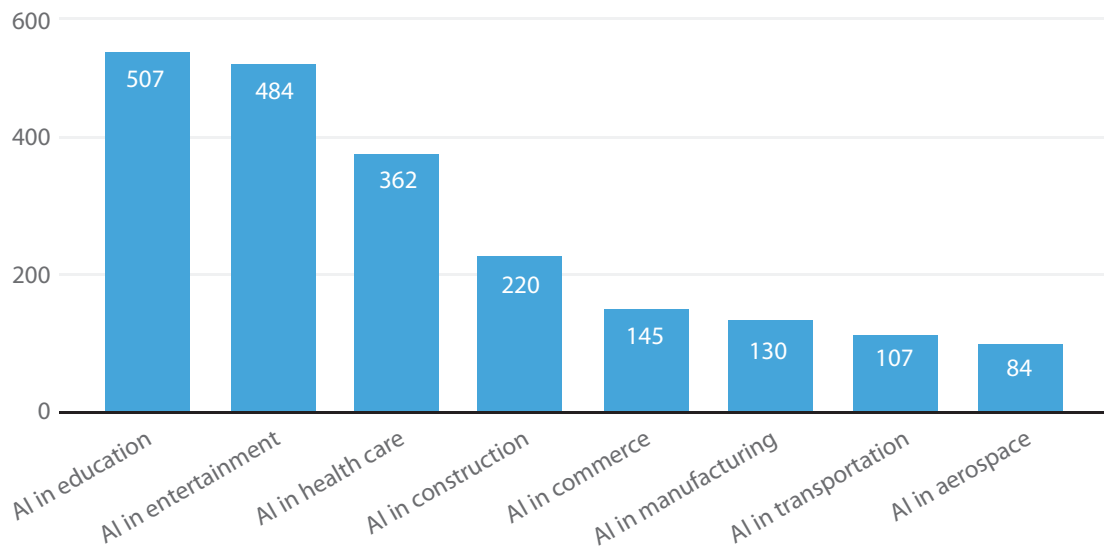
Interest in artificial intelligence compared to other fields measured in terms of search volume: Worldwide trend over the past 5 years

FIG 4.



Interest in AI in different engineering fields measured by the number of search engine hits (results returned within half a second)


FIG 5.



Number of search engine hits per AI industry term (per million)

The successful application of AI technology has led to consistent reductions in costs and increased revenues. In manufacturing, for example, predictive maintenance of equipment using machine learning leads to a 10% reduction in annual maintenance costs, up to a 20% downtime reduction and a 25% reduction in inspection costs<sup>14</sup>. In the financial sector, intelligent automation has led to 35% of financial services players reporting up to 5% increase in top line growth<sup>15</sup>.

FIG 6.



Forecast of revenues from AI for the enterprise applications market worldwide



In machine learning, available data, annotated with some label or category, is used to train machine learning algorithms to capture subtle patterns in the data that correlate with a given label or category. The model can be used to label new or unseen data once it has been trained on enough data. This learning can lead to new discoveries in the form of i) predictions, ii) detection of deviation from underlying base patterns, iii) detection of outliers, and iv) categorization of data into groups or clusters. This data-driven approach has led to less dependence on rule-based systems, where domain experts manually define rules that govern the mapping from input to output. In addition, explicit modeling of domain knowledge, which can be expensive or difficult to obtain by other methods, can be implicitly achieved with data-driven approaches.

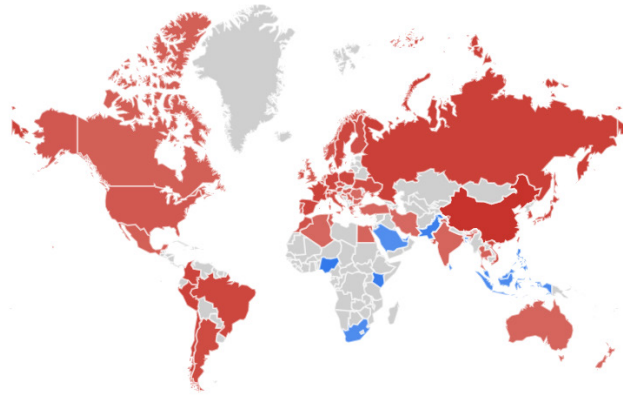
Today, data compilation, data analysis, and data management are essential parts of any successful enterprise. AI and ML systems are now being applied as methods to obtain, organize, and interpret the useful information contained in that data. Artificial intelligence is mainly achieved through machine learning and pattern recognition. In domains and problems where there is a large (or infinite) number of possible configurations and scenarios, developing a technology based on handwritten rules is not practical (or even possible). In addition to their incompleteness, rule-based systems can involve conflicting rules that need to be resolved using additional rules or heuristics. Machine learning strategies, rooted in probabilistic modeling and statistical pattern recognition, provide an elegant solution to these problems by learning automatically from the available data while allowing for the incorporation of domain knowledge.

For example, consider the optimization of software. A database management system (DBMS) is a software system that allows computers to efficiently store, retrieve, add, delete, and modify data for optimal load balancing and throughput. The challenge with DBMSs is that they have a large number of parameters and hundreds of configuration “knobs,” which makes DBMS configuration a challenging, time-consuming, and expensive task that requires vast knowledge of and experience with these complex systems. This optimization is typically carried out by highly skilled (and highly paid) experienced database administration experts. Machine learning methods have been used to automatically solve the challenging problem of finding the right configuration for a DBMS given its workload and data transaction patterns. These ML techniques can be used to automatically tune DBMS configurations by reusing training data gathered from previous tuning sessions carried out by human experts<sup>16</sup>. This approach can produce configurations that achieve up to 94% lower latency compared to their default settings or configurations generated by other tuning advisors<sup>17</sup>. The time required to entirely and automatically produce these configurations is also comparable to the time taken by human experts.

In some tasks, the ML/AI system exceeds human performance. For example, in object recognition within images the error rate in recognizing signs in the German Traffic Signs dataset is 1.16%, while an ensemble of Convolution Neural Networks achieved an error rate of 0.54%<sup>18</sup>. Similar human vs. machine learning results were found in playing video games like Go<sup>19</sup>, and in question answering and reading comprehension tasks on several datasets<sup>20</sup>.

FIG 7.


● artificial intelligence ● machine learning



Worldwide interest trend over the past 5 years in artificial intelligence and machine learning: red represents areas where interest is mostly in machine learning while blue represents areas where interest is mostly in artificial intelligence. Source: Google Trends

Interestingly, since the machine learning methods are highly empirical, they often lead to discoveries and novel findings that are not necessarily planned. For example, the Google AI team was studying the possibility of applying machine learning and computer vision technology to predict diabetic retinopathy from images of the retina.

FIG 8.



Retinal imaging can be used to collect images of retina to train machine learning algorithms to recognize and identify the stage of retinopathy

The machine learning algorithm not only performed competitively with human ophthalmologists but was also able to discover some cardiovascular risk factors that were previously unknown<sup>21</sup>.


In this information age, marked with an abundance of data and information, the potential for discovery, through the application of machine learning techniques, is at its peak. In addition to insight and discoveries from input data, these algorithms make it possible to tag or label new (unseen) data. Input data can be categorized as strongly or weakly supervised. Strong supervision is when the training data have annotations/labeling that exactly match the target output/label to be predicted on unseen test data; i.e., you have ground truth or gold labels (e.g., provided by human experts).

Weak supervision is an umbrella term for annotation/labeling methods that are noisier (e.g., provided by nonexperts, or by using heuristics, or general domain rules that label the data automatically and somewhat correctly) and only approximate the exact labels or true labels. The application of a machine learning process generates larger amounts of weakly supervised data that can be used to self-train and improve the very algorithms that produced the data, or to train other models by means of techniques such as transfer learning or zero-shot learning, which will be discussed later in this report.

# 2 Deep Learning

Artificial neural networks, widely known as neural networks, are biologically inspired computational models that were originally proposed to explain how thinking happens in the brain. The very first neural network model was proposed by McCulloch and Pitts in 1943<sup>22</sup>. Since then, the field of neural networks has advanced to the point where neural networks are now the de facto state-of-the-art models across a range of applications. A neural network consists of a collection of connected units or nodes, the “artificial neurons,” which loosely model the neurons in the brain. These units are linked with connections that resemble synapses in the brain that can transmit signals between nodes.

FIG 9.




The McCulloch-Pitts Model of neural networks

The figure above shows how these computational units work and transmit signals to connected nodes. When the sum of incoming signals exceeds a certain threshold, the computational unit, like a biological neuron, gets excited and transmits the output signal, which gets carried to connected nodes. Note that different inputs may have different weights reflecting differences in the importance of input signals or features. In recent neural models, the sum is replaced with functions that can capture more complex relations between the features and the output. Another factor is depth, expressed in terms of the number of hidden layers, which allows more patterns in the data to be recognized and modeled. This enhanced modelling capability improves the generalization power of these models when they are applied to different datasets or problems.

This ability to approximate complex functions makes neural networks more powerful and advantageous compared to other recognized machine learning methods. Traditionally, before the machine learning algorithm is applied, a lot of work is done to engineer input features, including their manual or automatic definition and selection, finding redundancies and interactions in the features, and reducing them or mapping them. Thus, the algorithm can be better applied to solve the problem when the number

of input features is very large. Neural networks are automatic feature extractors and can, in theory, approximate any function of arbitrary complexity if presented with enough training examples.


FIG 10.



*Deep neural networks are automatic feature extractors<sup>23</sup>*

Since the introduction of artificial neural network models, there have been several challenging problems and historical setbacks, including the AI winters in the 70s and 80s, that limited their application. With recent advancements in the mathematical formulations for these models and marked improvements in computational power, neural models have been resurrected and equipped with “depth” and increasingly complex architectures that have facilitated major breakthroughs in the field of artificial intelligence and machine learning. “Deep” neural models, consisting of many layers of connected units, mimic the activity in layers of neurons in the neocortex. They can learn from input examples, without being explicitly programmed with any task-based rules, to identify patterns in the input, which is typically a digital representation of text, images, speech, or other data.

FIG 11.



*A neural network: Layers of connected computational units mapping input to output*

One of the most prominent successes of deep learning is in computer vision and image recognition. Real-time image processing, segmentation, and recognition of objects in images is driving other major developments, including self-driving vehicles, digital pathology, and intraoperative tissue segmentation. The following subsections highlight some recent interesting and promising applications of deep learning.

NEURAL NETWORKS ARE AUTOMATIC FEATURE EXTRACTORS AND CAN, IN THEORY, APPROXIMATE ANY FUNCTION OF ARBITRARY COMPLEXITY IF PRESENTED WITH ENOUGH TRAINING EXAMPLES.

## STYLE TRANSFER

In the emerging technique of style transfer, the content of one image is combined with the style of another image. It is an optimization-based technique that takes a content image and a style reference image as input. A convolutional neural network (CNN) blends them together such that the content image is transformed to look like the original content image but painted in the style of the style reference image.


FIG 12.



*An example of style transfer from one image to another*

A video style transfer has extended this AI technique to a frame-to-frame transfer style providing a real-time stable style-processing method. Figure 13 shows an example of this process that was employed on a video of New York City at night.

FIG 13.

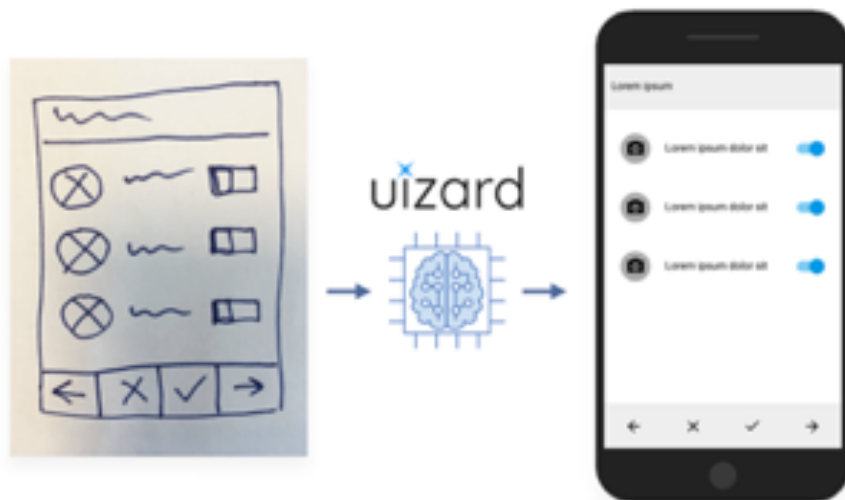


Video style transfer

## FROM MANUAL DESIGN TO GRAPHICAL INTERFACES

Deep learning has been recently applied to provide a creative solution to wireframe prototyping. The Copenhagen-based startup Uizard<sup>23</sup> is transforming the current user-friendly applications from a tedious programming and mockup design into an image upload that is processed by a neural network that automatically generates the code of a working prototype. This approach, which maps manual design to graphical interfaces with fully functional backends, improves the current workflow and may shape the future of the web/app development industry.

FIG 14.



Example of the application of the Uizard system

Such an AI technology will drastically reduce the development time of graphical interfaces and will facilitate the communication between designers and programmers during the product development process.

## AI STORYTELLER

Telling a story from a video is particularly difficult. Due to a user disability or limited network bandwidth, participation by, and communication with, certain audiences can be problematic. AI Storyteller is a deep learning-based technology that processes the sequences of frames and their context, transforming the big data of videos into a textual format that will allow people with disabilities such as blindness to enjoy the video content by reading it in braille. For entertainment companies, this technology would facilitate automated summary generation, a task that currently is exclusively carried out by humans. It would lower costs and the time required for introduction of various types of videos including films, educational videos, and documentaries.

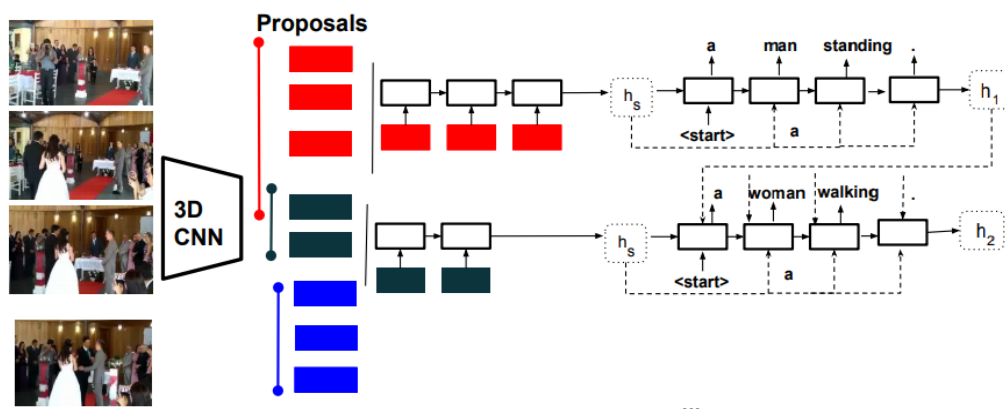
FIG 15.



### The AI generated description

A dog is standing in the middle of a house.  
 The dog runs around the room and the dog jumps up and down.  
 The dog is walking on the floor and the dog walks away.  
 The girl runs around the house and the other dog runs away.

(a) This technology has been designed by Facebook using its own VideoStory dataset, which compiles 20,000 videos with high user engagement into 123,000 descriptive sentences.



(b) Combining frames and context to generate descriptions.

Simplified architecture of the deep neural model



# 3

## Transfer Learning

Machine learning is a data-driven field. Data is the key factor in the applicability and performance of machine learning methods. In the typical setting of a machine learning problem, the goal is to learn some function from the input data, or training data, from which the machine learning algorithm learns. That learned function can then be used to make predictions or decisions about new data. Unfortunately, in many real-life problems, training data is not readily available, and it could be difficult to obtain, expensive, sensitive, or simply nonexistent. One possible answer to this problem is that the current application may have similarities to another task where data is available and more readily accessible.

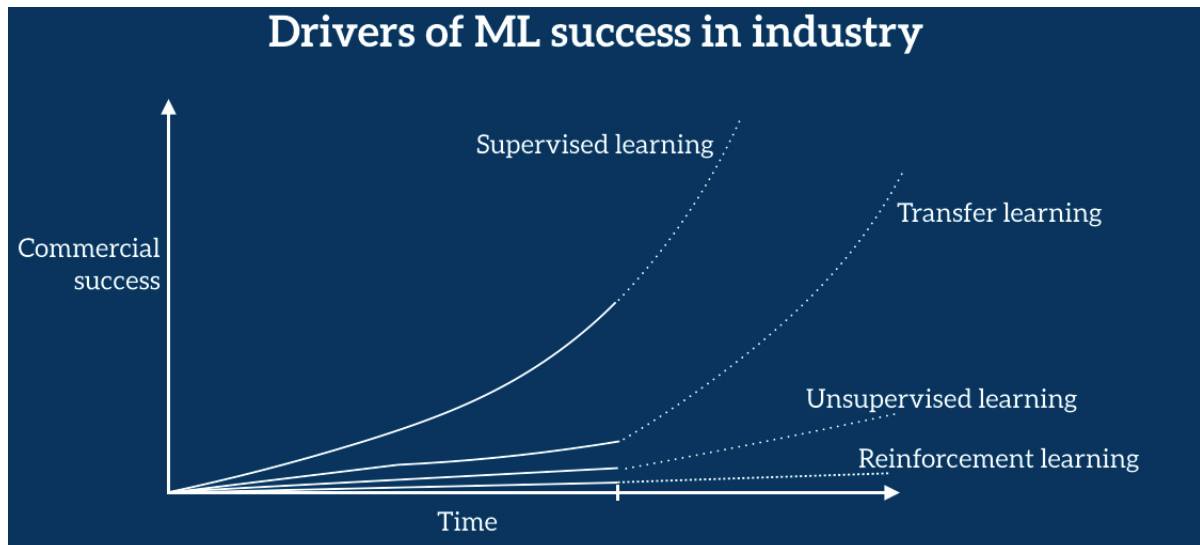
**Transfer learning** methods exploit the similarity between tasks or domains in order to solve a new task B, based on knowledge gained from previously solving a different but somewhat similar task A. The knowledge is transferred from task A to task B. The similarity can be in the problem domain, the data, the output (e.g., labels or classes in a classification task), or the tasks themselves, thus facilitating transfer and enabling machine learning algorithms to solve new problems by leveraging solutions to existing problems. The domain from which data comes for task A may be similar in nature to the domain of data in another task B. We refer to the first domain as the *source domain* and the latter as the *target domain*.

An example of domain similarity is the correspondence that exists between languages. A machine translation (MT) system may be trained with large amounts of data to translate from language  $L_1$  to language  $L_2$ . In the situation where the MT system needs to translate from language  $L_1$  to  $L_3$ , where  $L_3$  is a low-resource language and little or no training data is available to train the MT system, but  $L_3$  happens to be similar in its vocabulary, syntax, or some other aspect to  $L_2$ , then this similarity can be exploited to train the MT system to translate from  $L_1$  to  $L_3$ . This is accomplished by using the same strategy followed to translate from  $L_1$  to  $L_2$  with some adaptation or fine-tuning to  $L_3$ <sup>24</sup>.

In his popular tutorial titled “Nuts and Bolts of Building Applications using Deep Learning” at the Neural Information Processing Systems (NIPS’16) conference, arguably the most prestigious AI-related academic conference, Andrew Ng, chief scientist at Baidu, said that “transfer learning will be the next driver of ML commercial success.” Ng specifically sketched the relationship between different machine learning approaches and commercial success in industry.

“TRANSFER LEARNING WILL BE THE NEXT DRIVER OF ML COMMERCIAL SUCCESS.”  
ANDREW NG, FOUNDER AND CEO OF LANDING AI.

FIG 16.



Drivers of ML success in industry. “Nuts and Bolts of Building Applications using Deep Learning” tutorial at NIPS 2016

Transfer learning can also be applied as an initialization strategy for the deep network. A popular technique in current machine learning methods is to initialize the entire network with weights that were learned on a similar dataset, followed by fine tuning the entire deep network on the new dataset.

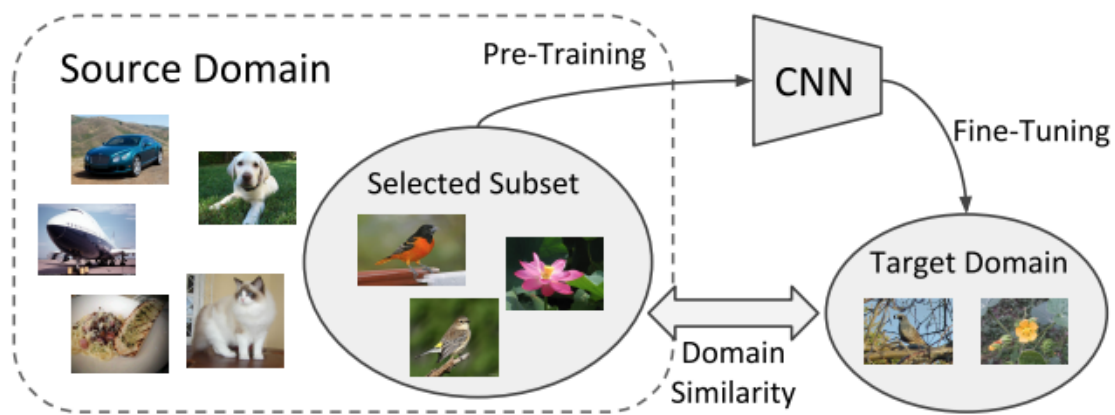
Two of the most prominent application domains of transfer learning are image recognition and natural language processing.

## TRANSFER LEARNING FOR IMAGE RECOGNITION

Transfer learning is a very popular technique in computer vision, especially for training recognition models. Since collecting huge amounts of data for a particular task is an expensive process, data scientists try to minimize the cost by applying transfer learning from a similar task. Deep models for recognition are trained on generic large-scale datasets like ImageNet, which allows the network to learn feature extraction capabilities. The final layer of the deep neural network is then replaced with a new dataset, which might be a subset of ImageNet or some other dataset whose domain is similar to that of ImageNet, and the model is trained only to update weights for the final (softmax) layer. This results in the model being trained in a very short period of time and not requiring as much data as would be needed to train it without this methodology.

Recent research on “Large Scale Fine-Grained Categorization and Domain-Specific Transfer Learning” has shown that the learning ability of state-of-the-art neural models for image recognition can be enhanced by increasing the similarity between the domain of the data used for pre-training the model and the target domain of the data used for fine tuning<sup>25</sup>.

FIG 17.



Given the target domain of interest, a convolutional neural network is trained on the selected subset from the source domain based on the proposed domain similarity measure and then fine-tuned on the target domain<sup>26</sup>.

In the figure shown above, the convolutional neural network is expected to work on a dataset of images of birds to perform recognition. Hence, this dataset of bird images is the “target domain” that the network needs to work on. In slight contrast to the conventional transfer learning technique of pretraining the CNN on a large-scale dataset like ImageNet and then fine tuning on the bird dataset, the authors of this paper show that if the CNN is pretrained only on selected classes from ImageNet that are closer to the target domain, then the transfer learning performance can be improved.

## TRANSFER LEARNING FOR NATURAL LANGUAGE PROCESSING

In natural language processing (NLP), the task of language modeling is defined as the generative task of predicting the next word when the system is provided with the previous word or words. This next-word prediction forces a model to learn long-term dependencies that capture hierarchical relations, syntactic information, and even sentiment, which are all useful for other NLP tasks such as syntactic parsing, question answering, and sentiment analysis.

Transfer learning also occurs when a model, such as a deep learning model, is trained on large amounts of data for some text analysis task and this pretrained model is found to be useful for a starting point for learning another task. Recent research in natural language processing has shown that pretrained models can be transferred to other tasks and can outperform state-of-the-art models on a number of tasks, including text classification, question answering, and textual entailment<sup>27,28</sup>.

Facebook research shows that knowledge can be transferred from the natural language inference (NLI) task, where entailment and contradiction relations are predicted between given pairs of sentences, to the universal sentence representation task, where the goal is to create a universal semantic representation of a

sentence that captures its meaning regardless of the language and its specific syntax peculiarities<sup>29</sup>. These two tasks clearly share, at some core level, the need to “understand” or model the meaning of the sentence. There has been a lot of work in academia on the NLI task where the input is a pair of sentences and the output is a prediction of whether the first sentence implies the second, contradicts it, or is neutral. Additionally, Facebook AI research shows that the knowledge gained from solving this problem can improve current machine learning models that attempt to learn universal sentence representations. Platforms and applications that support multi-linguality can benefit from a wide range of solved or nearly solved NLP problems to solve new problems by exploiting similarities in the data or the task in hand.

# 4

## Anomaly Detection

Anomaly detection aims to identify unusual patterns, anomalies, or data points that do not conform to the expected distribution. Various machine learning methods are used in anomaly detection, including threshold-based approaches, density estimation approaches, and classification approaches. **Applications of anomaly detection include fraud detection in financial transactions, fault detection in manufacturing, intrusion detection in a computer network, monitoring sensor readings in an aircraft, spotting potential risk or medical problems in health data, and predictive maintenance.**

Anomaly detection has been successfully applied to optimize operations in a number of industries. The anomaly detection market is expected to reach \$4.45 billion by 2023<sup>30</sup>.

**THE ANOMALY DETECTION MARKET IS EXPECTED TO REACH \$4.45 BILLION BY 2023.**

**Doxel** is a Silicon Valley startup providing AI-enhanced software focused on improving construction productivity<sup>31</sup>. Doxel uses rugged robots and drones equipped with cameras and LiDAR sensors to monitor and scan worksites. A Doxel robot scans construction sites every day to monitor how things are progressing, tracking what gets installed and whether it's the right equipment, at the right time, in the right place. Once a construction site shuts down for the night, the small robot deployed by Doxel can get to work. The robot scans the site and uploads data to the cloud. There, deep learning algorithms flag anything that deviates from the building plans so that a manager can fix it the day after. The robot has no problem following prescheduled paths that can include stairs, and just one of these robots can scan about 30,000 square meters over the course of a week.

The visual data collected by their robot is also processed to measure the currently installed units and the rate of production by matching acquired data against the desired planning and design parameters for the client. The company states that their AI platform can also detect errors in construction by comparing visual data from daily scans of the job site to small-scale design models. Doxel has had successful collaborations in the past, such as their project for the Kaiser Permanente Viewridge Medical Office project. Their real-time progress-tracking system was able to prompt the Viewridge Medical Office project team to take action when predefined schedule deviations were detected, eventually yielding a 38% increase in labor productivity across all teams involved in the construction, thus enabling the project to be completed 11% under budget<sup>32</sup>.

**General Electric**<sup>33</sup> launched its Brilliant Manufacturing Suite for customers, which the company had been field testing in its own factories. The system takes a holistic approach of tracking and processing everything in the manufacturing process to find possible issues before they emerge and detect inefficiencies. Their first “Brilliant Factory,” a \$200-million investment, was built in 2015 in Pune, India. GE claims it improved equipment effectiveness at this facility by 18%<sup>34,35</sup>.

**NextEra Energy** software, developed by **SpaceTime Insight**, helps with performance optimization, real-time diagnostics and troubleshooting, and maintenance crew scheduling<sup>36,37</sup>. NextEra’s ControlComm is built on a communications platform built by **AutoGrid**<sup>38</sup>, which helps business customers reduce energy bills during times of peak energy demand or high wholesale electricity prices by adjusting their energy consumption manually or with an automated solution.

**Duke Energy** developed the **SmartGen program**, which saved them more than \$4.1 million by triggering an early warning when one of Duke Energy’s steam turbines had begun to malfunction. Duke Energy also uses **SparkCognition’s** AI algorithms to safeguard its multibillion-dollar turbines by predicting potential disasters and shut downs<sup>39</sup>.

**SLAC National Accelerator** Laboratory uses AI to optimize solar and other distributed energy resources<sup>40</sup>. Recent news reports describe new research efforts from the US Energy Department’s SLAC National Accelerator Laboratory at Stanford University that use AI applications to help utilities better integrate their solar resources and make more informed planning decisions for enhancing grid reliability, resiliency, and security.

# 5

## Multimodal Machine Learning

To equip AI systems with human-like intelligence, they must be modeled and trained to perceive and understand the surrounding environment, which involves interpreting multimodal signals. Multimodal machine learning aims to build models that can process and relate information from multiple modalities, including numbers, images, video, text, audio and speech, robotic sensors, and time-series data (data collected over time and annotated with time-stamp information).

Multimodal techniques include combining and integrating features from different streams and types of data for the purpose of predicting an outcome. This integration of features can be an initial step of processing (early fusion), or it can occur later (late fusion) in the ML process. Other multimodal techniques include: learning joint representation of input data, which aims to concurrently learn from two or more modalities and map them into one “joint” model; translation between modalities; alignment of modalities; and co-learning, where the machine learning algorithm is trained with data from two or more parallel modalities, for instance, images and their captions or speech and its transcription.

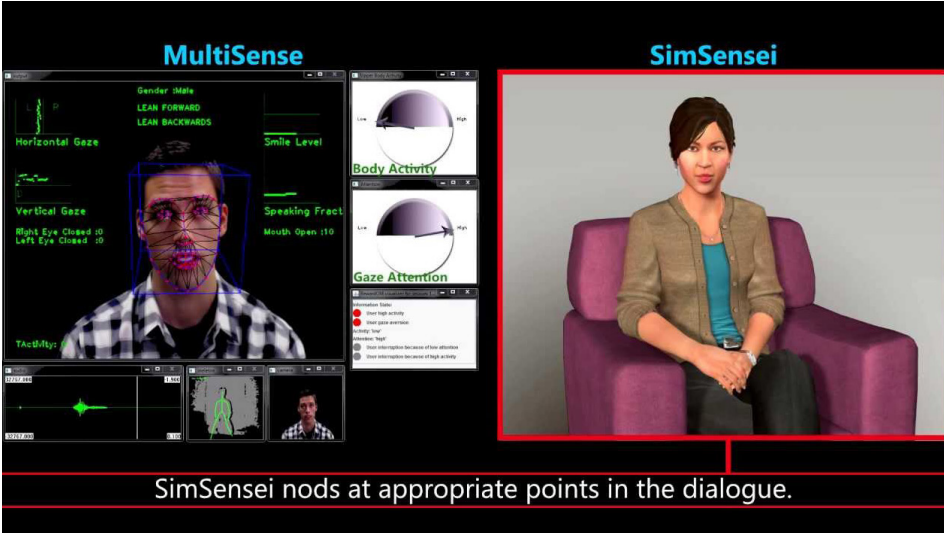
Multimodal machine learning works well for several reasons. When data is missing or noisy in one modality, data from another modality can fill the gap and help the algorithm make a decision despite incompleteness or noise in the data. For example, if visibility is low for the image modality in a vehicle recognition task, information from a thermal sensor can help detect the vehicle. Also, signals from different modalities can reinforce each other, which can boost confidence in the decision made by the algorithm. More interestingly, some patterns in the data can only be seen or understood when multiple (complementary) modalities are combined. The following applications of multimodal machine learning show the capabilities of this approach and its usefulness across different domains.

WHEN DATA IS MISSING OR NOISY IN ONE MODALITY, DATA FROM ANOTHER MODALITY CAN FILL THE GAP AND HELP THE ALGORITHM MAKE A DECISION DESPITE INCOMPLETENESS OR NOISE IN THE DATA.

**Oak Ridge National Laboratory**<sup>41</sup>, a federally funded science and technology research center managed and operated by nonprofit UT-Battelle<sup>42</sup>, applied multimodal machine learning for vehicle identification and classification<sup>43</sup>. Multimodal Sensing and signal learning were used to achieve this goal. Their algorithm was designed to parse data gathered by roadside sensors, which made it easier to identify vehicles. The lab researchers built a sensor platform to collect detailed images of cars, as well as electrical pulses and audio signals from engines, to uniquely identify vehicles.

Carnegie Mellon multimodal research group, led by assistant professor Louis-Philippe Morency<sup>44</sup>, developed the multimodal system MultiSense that tracks a person’s gaze, facial expressions, body orientation, and head position to make inferences about the psychological state of the person. This type of research provides new insights into psychological, social, and biological aspects behind facial expressions and helps develop new diagnostic tools for mental and psychological illness<sup>45</sup>. The US Department of Defense is using MultiSense as a tool to detect post-traumatic stress disorder (PTSD), and as a tool to predict behavior.


FIG 18.



MultiSense, the sensing and analysis technology created in the SimSensei project at CMU

At MIT, a multimodal approach is used to classify neuromuscular signals and infer user intention. MIT has created a headset called AlterEgo<sup>46</sup>, a wearable device that attaches to the jaw, where electrodes pick up neuromuscular signals triggered when the subject says words in their head. Recognition accuracy is reported to be 92% even when the mouth is not open. Thus, the device measures “silent speech.”

FIG 19.



AlterEgo seeks to make computing a natural extension of the user. Credit: Arnav Kapur, Neo Mohsenvand.




# 6

# One-Shot and Zero-Shot Learning

The existing AI models require large amounts of task-specific training data such as the ImageNet<sup>47</sup> and CIFAR-10 image databases<sup>48</sup>, composed of 1.2 million and 60 thousand data points, or labeled images, respectively. Learning from only a few examples remains a key challenge in artificial intelligence. The next objective is to make the standard supervised deep learning paradigm offer a satisfactory solution for learning new concepts from a small amount of data.

The answer is **one-shot learning**. The key motivation for the one-shot learning technique is that systems, like humans, can use prior knowledge about object categories to classify new objects. For example, humans can learn to differentiate various types of apples, and differentiate apples from other similar fruit, after seeing one or a very small number of images for each type. Humans can also learn to identify new objects that are similar to familiar objects if they are taught the analogy and the characterizing properties of the new object; for instance, pen vs. pencil. Children are estimated to learn 10-30 thousand object categories by the age of six<sup>49</sup>. One-shot learning and zero-shot learning attempt to mimic this ability. In one-shot learning, the machine learning algorithm is presented with one, or very few, training examples from which it must learn. In **zero-shot learning**, the machine learning algorithm may not be presented with any training example in some of the input categories but is able to correctly classify those instances using attribute-based information incorporated in the model. This is known as **attribute-based classification**<sup>50</sup>.

FIG 20.



Architecture of iterative refinement LSTM for one-shot learning drug discovery

MIT and Stanford University developed a neural network architecture to do one-shot learning for drug discovery (Figure 20). They showed that this method can lower the amount of data required to make meaningful predictions. Their architecture is based on the now widely recognized long short-term memory (LSTM) neural networks that capture long temporal patterns and long-distance dependencies in the data. Their iterative refinement long short-term memory facilitates the learning of meaningful distance metrics on small-molecule space<sup>51</sup>.

DeepMind has developed a matching nets architecture that uses neural networks augmented with memory. In these models, there is some external memory and an attention mechanism that is used to access the memory<sup>52</sup>. It's a network that learns how to learn a classifier from only a very small number of examples by estimating the probability that an input image  $x$  has an output label  $y$ .

FIG 21.

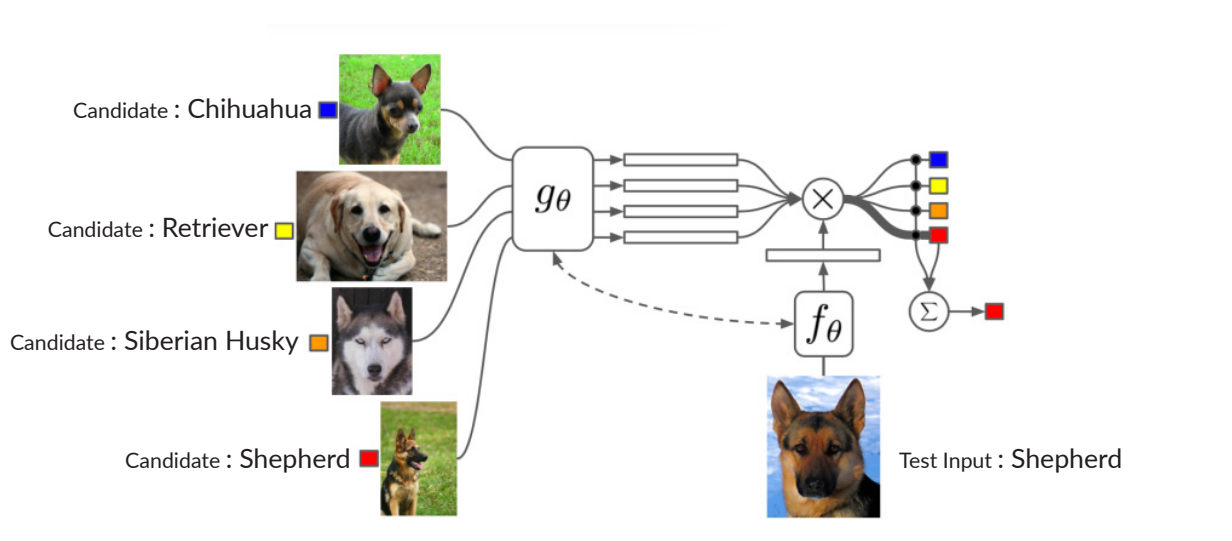


Figure 21. Matching network architecture

This algorithm improves one-shot accuracy on ImageNet from 87.6% to 93.2% compared to competing approaches.

One-shot learning may be the answer to one of AI's biggest challenges and application bottlenecks: limited data in some domains.

# 7

## Challenges and Opportunities

AI is arguably the leading driver of change and acceleration in today's world. This powerful collection of tools and techniques creates endless opportunities and yet poses some serious challenges and questions for AI researchers and professionals, as well as for users and society at large.

The rate at which computational power and memory is becoming cheaper and more available creates opportunities for wide adoption and application of AI and ML technologies across different industries. Nonetheless, many key questions and challenges need to be addressed before AI is exploited to its full potential.

As AI technology is performing competitively with humans or even outperforming humans in a number of tasks, the **need for human labor will decrease**. This does not only imply a change in demand in the job market and loss of jobs in some sectors, but also poses a greater challenge as increased automation replaces people. When automation replaces people in labor-intensive tasks, humans are left to confront the subtler or more intellectually challenging aspects of jobs. Computers will continue to be more and more effective and efficient in information processing and data mining, but insight-based decision-making remains a human-led endeavor. Machines will be handling high-level tasks such as medical diagnosis, stock market prediction, and negotiation. This new world of enhanced machine capabilities poses many important questions. Will humans have to **redefine their roles and even their purpose** in this era of increased automation? Will this have an influence on how humans develop mentally and even biologically as their roles and patterns of thinking change? Will some skills become more important than others in the transformative AI era? Will educational and professional training institutions have to **rethink topics and skills of their curricular focus** to reflect the change brought by machine learning's profound impact?

Increased automation and increased dependence on intelligent systems also leads to the question of whether **behavior of machine learning algorithms** can be guaranteed or predicted, especially in safety-critical applications. Does increased accuracy or improved performance on some metric translate to predictability and consistency of behavior? Take this simple example: a self-driving vehicle, trained only using data with no domain knowledge or rules, learns to slow down when the distance from the object in front of it decreases and learns that it's safe to speed up when no object is seen in the near horizon. That training may lead the machine to learn a positive correlation between distance from objects/obstacles and speed. Without explicitly injecting the speed limit constraint, such an algorithm may decide to continue to speed up when no object is seen in front. These **"action limits" or constraints** may not always be explicitly available to train models in order to implement them. Common-sense knowledge is specifically hard to model and incorporate because it is implicitly learned by humans and not typically expressed in naturally

occurring training data. Humans consider it to be too obvious to explicitly state; and since machine learning depends highly on data in order to learn, how does **bias in the data get absorbed or amplified** by these algorithms<sup>53</sup>? How is **fairness** achieved when these algorithms are made to make critical decisions or recommendations, like what candidate gets hired, or what patient population is more suitable for a new trial treatment? How well do machine learning models **generalize** and make reliable predictions when tested on new data or under slightly different circumstances?

**INCREASED AUTOMATION AND INCREASED DEPENDENCE ON INTELLIGENT SYSTEMS ALSO LEADS TO THE QUESTION OF WHETHER BEHAVIOR OF MACHINE LEARNING ALGORITHMS CAN BE GUARANTEED OR PREDICTED, ESPECIALLY IN SAFETY-CRITICAL APPLICATIONS.**

In addition to the issue of **verifiability** and ensuring machine learning algorithms are “well-behaved<sup>54</sup>”, the difficulty in **interpreting and explaining, in terms of performance, the predictions or decisions** made by some of the most powerful algorithms hinders their adoption in some domains. This is especially relevant since current research shows that, despite their competitive performance, some of the most powerful techniques like deep learning can be easily fooled and will label, with high confidence, unclassified data or random noise as other objects<sup>55</sup> or can fit random labeling of the data<sup>56</sup>. Research in **explainable AI** attempts to address the problems of “explainability” and “interpretability,” especially for deep learning models that are difficult to interpret or explain.

As AI and machine learning success is driven by the availability of training data, **opportunities and challenges are created around data**, and entirely new fields of study and practice emerge such as data science. How do we solve problems where there is little or no data? As researchers and industry collect data to enable machine learning, how are the rapidly increasing amounts of data managed, processed, and retrieved? How will the value and accessibility of data change as data becomes key for developing AI technologies? Will data become the new currency in AI-enabled industries? Will there be demand and pressure for developing standards governing data sharing and protection that balance rights to competitive advantage with social good on a larger scale, especially in domains like medicine and education?

***We touch upon some of these concerns and issues in the following interview.***



## Interview with Zach Lipton,

Assistant Professor in the Tepper School of Business, and affiliate faculty in the Machine Learning Department (MLD) and Heinz School of Public Policy at Carnegie Mellon University

***What is/are some misconceptions about AI and machine learning that may hinder their successful application in different domains?***

It's not a specific misconception, but the whole framing of issues in the vaguest terms; we often hear the leap from very specific technology to the all-encompassing notion of "AI" (Is the AI going to be my next boss? Is AI going to take my job?).

A more concrete example is the conversation about progress; it's captivating, and it's very easy for people to see that there is progress and follow it enthusiastically without doing critical thinking about what we're progressing at vs. what we're not progressing at exactly. For example, we're making progress in making predictions, such as building models that can make accurate predictions based on high-dimensional data and structured objects and not just producing binary labels; i.e., more complex tasks such as generating sentence and structured predictions. This may fool people into thinking that we progressed in AI broadly; however, an area where we didn't make progress is reasoning (pushback from causal inference community, societal implications, automated decision systems, death of expert systems), discussion of will your next boss or doctor be AI is obvious, making interventions; counterfactual reasoning, it's not obvious if there's a bridge from fitting data to counterfactual reasoning.



***This leads to the topic of "explainable AI" and work in interpretability and in making ML models "interpretable" as a way of adding a reasoning layer to predictions. What are your thoughts about that?***

In machine learning people focus on predictive performance; interpretability ends up being like "chronic fatigue syndrome" (which captures an array of symptoms and underlying causes). In order to solve a problem, you need to identify the problem; i.e., identify underlying concepts and go after them. Contexts in which demand for interpretability arise are legitimate and are often compelling problems, but the word is applied too broadly to have an answer to the interpretability problem. What is

its logic? What's the model reasoning? Causality has a well-defined notion of spuriousness but not empirical risk minimization; for instance, an ERM-based model may predict that people who wear sneakers are less likely to pay a loan, but we know that this is based on correlation, not causation.

*Author's comment:*

*Empirical risk minimization is a statistical principle used in the supervised learning framework where we attempt to learn and predict an underlying process or distribution of data by observing a sample of that data; and in order to make predictions that are as accurate as possible on the unknown true distribution, we must minimize the error of our predictions on the data we observe from that distribution. This error is hence termed "empirical risk." The mathematical constructs used to define the empirical risk, typically some function of the difference between the true and predicted labels/outputs, can capture correlations between data features and output but says nothing about causal relations between the two.*

Other issues include reliability and robustness; e.g., out-of-sample generalization ability, which is often confused with interpretability. For example, consider the topic of domain adaptation. The question is, what is it that might change and can I identify it? Can we build models that, given invariance between train and test data, can perform well on unseen test data? I am loosely skeptical about approaches and models that fundamentally don't address these core problems by willing to go deeper and find the questions that need to be asked. Currently we're a community that has one hammer — prediction — but prediction isn't always the answer.

## **What are your thoughts on deep learning as a disruptor technology?**

There's uncritical thinking of the ability to solve conceptually hard problems by adding skip connections.


At the same time, there's anti-hype that's equally trivial and uncritical, such as arguments about deep learning that are tautological (e.g., deep learning is not going to solve all problems).

*Author's comment:*

*Skip connections are architectural variants in deep neural network models that have been found very useful for training and improving the performance of these models. They are extra connections between nodes in different layers that skip one or more layers of nonlinear computation. The comment about the ability of solving conceptually hard problem is suggesting that although architectural "hacks" may lead to improved performance on some tasks, they are not enough to tackle conceptually hard problems or identify the source of difficulty.*

There's no doubt that pattern recognition is a really important tool if you want to build adaptive systems, and deep learning is our best pattern recognition tool right now — especially for large datasets and high-dimensional data. There are a lot of learning frameworks and approaches that are so much better, with the existence of universal function approximators. They opened plenty of opportunities. Deep reinforcement learning is a good example where most of the improvement is from the function approximators (i.e., deep learning models) combined with reinforcement learning, compared to previous work in reinforcement learning that did not employ deep learning.

FIG 22.



Skip connections<sup>57, 58</sup>

***This reminds me of a statement often made about advancements in deep learning and that work in this area is “mostly engineering rather than science.” Do you agree with that?***

I am a big fan of science and there's little science in machine learning. There's a differentiating factor between science and leaderboard chasing: science is more concerned with simple explanation. The modern culture of research is more enthralled by unnecessarily complex models. For example, a paper with 10 ideas (but only one that really works) is more likely to get published than a paper with one thing that works and nine negative results. Falsifiability (running experiments to rule out ideas) is essential in science. There are interesting ideas in papers on the theory of machine learning, but the modern practice of learning theory could be improved if we think about it more the way they do it in the natural sciences, where the mark of good theory is how well it describes the things you've observed and that you might see, or it allows you to observe things in the world that you weren't able to see before the existence of the theory (think relativity for example) – this idea of making predictions that someone wouldn't have necessarily assumed to be true. A problem I've noticed in the ML field is too much separation between theory

and practice. For example, there's a lot of work in theory that shows that neural net works may not work with no or minimal expert intervention to achieve the optimal solution due to the existence of multiple local optima that the optimization algorithm may get stuck at), but no one in practice seems to be stuck at those! A big problem in practice is generalization; that is, whether a model that was trained on data will work on unseen data. There's now more work on theory that is closer to impractical reality. There's more dialogue based on questions and observations that come from practice. Closing the loop between theory and observation in the natural sciences was the norm long before this kind of work in machine learning.

***What else, in addition to the gap between theory and practice, may inhibit wide adoption or practical application of machine learning in industry?***

When people have to make decisions, decisions don't inform the kind of data used to make the decisions, that there are no feedback loops. Consider Internet startups (almost all doing the same thing). The technology is old (a back end that interacts with a database, updates via some view in the browser or phone). What can we do with this kind of setup using machine learning? Take for example Amazon, which was launched in the mid 90s ... and now we have ride-sharing services. Machine

learning will offer modest improvement (we've limited imaginations as researchers), but practitioners will continue to identify interesting problems inspired by needs in their domain, for instance, agriculture. An area where we might encounter a lot more obstacles is automated decisions that don't account for the environment and in situations where there's potential feedback loops (even if the machine learning model is not responsible for the change, but the environment changes over time). We don't always recognize the hard part of doing AI. It's hard to define good objectives and what a good objective is. For example, in recommender systems, an objective may be maximizing clicks; but we do not really know the full effect of that. What are the consequences of showing people (or recommending) what news articles to read or items to buy?

Here's a high-level synthesis: There are simple applications where ML will continue to propagate and be useful and benign, and hard problems that are very specifically dependent on causal reasoning and where stakes are high and no organizing body will allow applying ML right away. The most danger is in the middle spot, where the ML tools are easy to scale and apply but with potential serious consequences that are not necessarily known. As a society, we have a responsibility for regulating how these tools are applied and making sure that ML is not applied where it doesn't fit; for instance, consider news and information propagation in social media and advertisements. We're now dealing with the consequences of showing people what to click, which can be catastrophic.

***Speaking of disruptor technologies and influence, what's your opinion about the next big thing? Is transfer learning the next game changer?***

Transfer learning is already the status quo. In natural language processing (NLP), it's common to

fine-tune models on target domain data, which is, transfer learning. I don't necessarily see a next breakthrough around the corner by just grabbing some pretrained weights and fine-tuning, but models that work on some variant of the task and cope naturally without additional training data are important (domain adaptation is a subset of transfer learning). Ultimately, making real progress on that is not saying, let's fine-tune pre-trained weights; there remains a question of how to be robust. For example, when a human plays a game and all colors in the game change, the human can deal with that quickly. It's important to understand the mechanism, assumed invariance between training and target tasks, and assumptions of causal structure in the problems.

***Along the lines of breakthroughs and the next game changer, machine learning has been successful at training models for solving specific tasks (or similar tasks); but we're still not at a point where we can create agents or models that can learn from the previous experience of solving different tasks and be able to use that collective experience to solve a completely new task.***

Right. What is the knowledge that we bring to bear on each successful thing that we do? It might be worth thinking about how a lot of ML is model-free, where by "model" I mean a model of the world, or, what you know and what you don't know about the world/context. Deep learning does things end to end, vs. old symbolic logic AI (that explicitly defines "knowledge" in terms of a set of predefined, static rules that are often brittle). Is there some better way to build and learn a model from experience, and which part of it to bring to a new task (are key questions). There's a question of what it is that's in the future that's the same in the past.



***Can we automate machine learning itself? Can machine learning do machine learning? Any concluding thoughts on “learning for learning”?***

Right now, “learning for learning” is a double-edged sword. On one hand, we have so many graduates who get cushy jobs in ML where a huge part of the work is to get a model we already know to work on a new dataset. Steven Boyd, author of *Convex Optimization*, stated that linear programming is a technology, but convex optimization is not yet a technology! Stochastic optimization is an art. From our performance in ML, convex optimization is not a technology. What does that mean? It’s not at the point where a nonexpert can solve it; it still requires too much work. Your car is a technology because you drive to work every day and you don’t know about automotives. Neural learning is low to medium skills for ML engineers; for instance, I’ve got a dataset of 40k images I want to train a CNN (there are now some platforms that try to offer training as a service).

So on one hand, there are great employment opportunities with machine learning (knowledge and skills). On the other hand, the path of technology is to get the thing to work (e.g., linear program) (with no or minimal expert intervention). A lot of what we call research involves mechanical work “elbow grease.” This is precisely the sort of thing that we automate (given enough instances, you should be able to press play): how to choose hyperparameters for a model or the right architecture for a neural network. Just like in traditional ML, where the goal is to fit a function and generalize to samples of unseen data, with “learning for learning,” you’ll come up with protocols that will learn and generalize to a new problem instance. There are many definitions for “meta learning” or “learning for learning”; it’s tied to discussions of transfer learning.

One is generalization across problem instances, and another is hacking into the solution architectures. The main job of humans is to be the random seed (What if I try this trick or list of tricks? Let’s try every one at every single layer). Making random choices and executing things by brute force are the kind of things machines are good at and are probably more efficient than humans at.



# ENDNOTES

1. [https://www.sciencedaily.com/terms/artificial\\_intelligence.htm](https://www.sciencedaily.com/terms/artificial_intelligence.htm)
2. <https://www.gartner.com/it-glossary/artificial-intelligence/>
3. <https://dl.acm.org/citation.cfm?id=3298188>
4. <http://karpathy.github.io/2014/09/02/what-i-learned-from-competing-against-a-convnet-on-imagenet/>
5. <https://dl.acm.org/citation.cfm?id=3161187>
6. <https://www.theverge.com/2017/10/24/16526370/baidu-deepvoice-3-ai-text-to-speech-voice>
7. <https://med.stanford.edu/news/all-news/2011/11/stanford-team-trains-computer-to-evaluate-breast-cancer.html>
8. <https://www.sciencedaily.com/releases/2017/01/170125145903.htm>
9. <https://www.cnbc.com/video/2018/11/09/worlds-first-ai-news-anchor-debuts-in-china.html>
10. <https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-market-74851580.html>
11. [https://www.pwc.com/CISAI?WT.mc\\_id=CT1-PL52-DM2-TR1-LS4-ND6-BPA1-CN\\_CIS-AI-AIsocial](https://www.pwc.com/CISAI?WT.mc_id=CT1-PL52-DM2-TR1-LS4-ND6-BPA1-CN_CIS-AI-AIsocial)
12. <http://aiindex.org/2017-report.pdf>
13. <https://medium.com/civic-analytics/signal-1-mls-growth-trend-has-surpassed-moore-s-law-fdcaf34be6e0>
14. <https://www.mckinsey.com/industries/semiconductors/our-insights/smartening-up-with-artificial-intelligence>
15. <https://www.capgemini.com/resources/the-growth-in-the-machine/?thefinancialbrand>
16. <https://www.cs.cmu.edu/~ggordon/van-aken-et-al-parameters.pdf>
17. <http://www.vldb.org/pvldb/vol11/p1910-zhang.pdf>
18. <http://benchmark.ini.rub.de/?section=gtsrb&subsection=news>
19. <https://techcrunch.com/2017/05/23/googles-alphago-ai-beats-the-worlds-best-human-go-player/>
20. <https://medium.com/the-new-nlp/ai-outperforms-humans-in-question-answering-70554f51136b>
21. <https://www.nature.com/articles/s41551-018-0195-0>
22. <https://link.springer.com/article/10.1007%2FBF02478259>
23. <https://towardsdatascience.com/why-deep-learning-is-needed-over-traditional-machine-learning-1b6a99177063>
24. <https://uizard.io/>
25. <https://aclweb.org/anthology/D16-1163.pdf>
26. <https://arxiv.org/abs/1806.06193>
27. <https://arxiv.org/pdf/1806.06193.pdf>
28. [https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf)

29. <https://arxiv.org/pdf/1801.06146.pdf>
30. <https://research.fb.com/wp-content/uploads/2017/09/emnlp2017.pdf>
31. <https://www.marketsandmarkets.com/Market-Reports/anomaly-detection-market-138133262.html>
32. <https://www.doxel.ai/>
33. <https://medium.com/@doxel/introducing-artificial-intelligence-for-construction-productivity-38a74bb-d6d07>
34. <https://www.ge.com/digital/manufacturing-execution-systems>
35. <https://www.ge.com/digital/press-releases/ge-launches-brilliant-manufacturing-suite>
36. <https://www.ge.com/reports/post/110927997125/ges-brilliant-advanced-manufacturing-plant-in/>
37. <https://spacetimeinsight.com/portfolio-posts/nextera-energy-resources/>
38. <https://youtu.be/pl6-P2lQB5Q>
39. <http://www.auto-grid.com/news/nextera-energy-services-teams-up-with-autogrid-to-offer-new-demand-response-programs-in-pjm/>
40. [https://ww2.frost.com/files/7414/6894/1903/Spark\\_Cognition\\_Award\\_Write\\_Up.pdf](https://ww2.frost.com/files/7414/6894/1903/Spark_Cognition_Award_Write_Up.pdf)
41. <https://www.energy.gov/eere/solar/articles/national-lab-uses-artificial-intelligence-optimize-solar-and-other-distributed>
42. <https://www.ornl.gov/>
43. <https://www.ut-battelle.org/>
44. <https://ieeexplore.ieee.org/document/8108568>
45. <https://www.lti.cs.cmu.edu/news/morencys-multimodal-research-helps-doctors-diagnose-mental-illness>
46. <https://dl.acm.org/citation.cfm?id=3133948>
47. [https://dam-prod.media.mit.edu/x/2018/03/23/p43-kapur\\_BRjFwE6.pdf](https://dam-prod.media.mit.edu/x/2018/03/23/p43-kapur_BRjFwE6.pdf)
48. <http://image-net.org/about-stats>
49. <https://www.cs.utoronto.ca/~kriz/cifar.html>
50. [http://geon.usc.edu/~biederman/publications/Biederman\\_RBC\\_1987.pdf](http://geon.usc.edu/~biederman/publications/Biederman_RBC_1987.pdf)
51. <https://ieeexplore.ieee.org/document/6571196>
52. <https://pubs.acs.org/doi/full/10.1021/acscentsci.6b00367>
53. <https://arxiv.org/pdf/1606.04080.pdf>
54. <https://arxiv.org/abs/1707.09457>
55. <https://arxiv.org/pdf/1708.05448.pdf>
56. <https://arxiv.org/abs/1412.1897>
57. <https://arxiv.org/abs/1611.03530>



# About PreScouter

## PRESCOUTER PROVIDES CUSTOMIZED RESEARCH AND ANALYSIS

PreScouter helps clients gain competitive advantage by providing customized global research. We act as an extension to your in-house research and business data teams in order to provide you with a holistic view of trends, technologies, and markets.

Our model leverages a network of 3,000+ advanced degree researchers at tier 1 institutions across the globe to tap into information from small businesses, national labs, markets, universities, patents, start-ups, and entrepreneurs.

### CLIENTS RELY ON US FOR:



**Innovation Discovery:** PreScouter provides clients with a constant flow of high-value opportunities and ideas by keeping you up to date on new and emerging technologies and businesses.



**Privileged Information:** PreScouter interviews innovators to uncover emerging trends and non-public information.



**Customized Insights:** PreScouter finds and makes sense of technology and market information for clients. **COMPLETED RESEARCH HOURS OF 150,000+**

