Getting Started with Machine Learning

CSC131 The Beauty & Joy of Computing

Cornell College
600 First Street SW
Mount Vernon, Iowa 52314

September 2018

Contents

Chapter 1

Applications: where machine learning is helping

1.1 Sheldon Branch

Sheldon Branch CSC131 9/3/13 Machine Learning in the Medical Field The medical field is facing a huge problem, one that affects over 14 million people every year. This problem is cancer. Because of our lack of knowledge, doctors are often unable to reliably identify cancer. Researchers at The Johns Hopkins Hospital in Baltimore reviewed tissue samples from 6,000 cancer patients across the country. They found one out of every 71 cases was misdiagnosed; for example, a biopsy was labeled cancerous when it was not. And up to one out of five cancer cases was misclassified. Human error in the medical field is responsible for hundreds of thousands of deaths each year. So where does machine learning fall into all of this? Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it learn for themselves. By using algorithms to build models that uncover connections, doctors can make better decisions without human intervention. This innovation allows calculations and diagnoses to be made without human error. Healthcare involves a lot of prediction tasks. Such tasks are perfect for an artificial intelligence whose sole purpose is to predict. Regina Barzilay is an AI pioneer who leads a team at MIT. In 2014, she was diagnosed with breast cancer. After her initial diagnosis, she had a mammogram and the doctors said that her cancer was very small. Next, the doctors did an MRI which reported that Reginas cancer was very widespread. The doctors then needed to do a biopsy on Regina, which elucidated that the MRI showed a false positive. After her treatment, her past mammograms were reviewed and small spots of cancer

were found on the scans as early as two years prior to her diagnosis. Using a machine to predict cancer instead of doing painful and expensive procedures is much more efficient. Her diagnosis and treatment fueled her desire to use her knowledge of artificial intelligence and machine learning to prevent such mistakes and improve the accuracy of cancer diagnoses. Regina aims to utilize her computer science knowledge to create an AI that can recognize patterns in medical images and doctors electronic notes in an attempt to catch cancer earlier and avoid overusing invasive treatments. A few years ago, Researchers in Brazil studied the brains of expert radiologist in order to understand how they reached their diagnoses. The researchers wanted to know if the seasoned diagnosticians were applying a mental rule book to the images, or if they applied pattern recognition. Twenty-five radiologists were asked to evaluate X-rays of the lung while inside MRI machines that could track their brain activity. Some of the X-rays contained common encounters, such as a palm shaped shadow of pneumonia or the dull, opaque wall of fluid that accumulated behind the lining of the lung. Others showed lines of drawings of animals. The final group had the outlines of letters of the alphabet. The radiologists were shown these images in a random order and then asked to call out the name of the lesion, animal, or letter as quickly as possible. In all three cases, the same areas of the brain lit up. These results support the hypothesis that a process similar to naming things in everyday life occurs when a physician promptly recognizes a characteristic and previously known lesion. Identifying a lesion was a similar process to naming an animal. When we recognize an animal, were not considering and eliminating alternative candidates nor mentally fusing animals whose combined physical attributes result in a similar figure. Instead we recognize animals as a pattern. The same is true for the radiologists when they recognized the lesions. Machine learning is enabling AI to recognize and discriminate between various images based on the characteristics of the presented objects. If such a process is so similar to identifying lesions in X-rays, AI in the medical field may become a necessity. This is an excellent and noble endeavor that can potentially save millions of lives each year. Unfortunately, artificial intelligence has a ways to go before its ready to diagnose anyone. There are several obstacles in the path of machine learning, one of the most prominent being the issue of bad data. Bad data will corrupt the prediction of the AI and make the machines diagnosis less accurate. If the hospital's pathologists miss cancer diagnoses or misdiagnose something else as cancer, the system will be trained on bad data. Its important to make sure that whatever human-generated information is used is as clean as possible. (Cats and dogs, numbers). Another type of bias is one where a certain population is underrepresented in the general patient population. That leads to higher error rates. In the hospitals Reginas team generally works with, there has been a smaller proportion of African-American women. They can address this algorithmically but they are also collecting data from hospitals that have a reasonable representation of diverse groups to compensate for the bad data. Another example of misrepresentation is the facts that all breast cancer oncology decisions made in the US today are based on the 3trials. This is not nearly enough data to make judgements on something as varied and unique as

cancer. It is also quite tedious to obtain the health data from several hospitals. To access any of this data, one would need approval from an institutional review board at the hospital to make sure its handled according to protocol. Even after this, the data needs to be anonymized. Despite these setbacks, machine learning is definitely making progress in the medical field. There is a condition in breasts called high-risk lesions. In the US, all patients with these lesions get surgery but 87benign and did not need surgery. With the help of machine learning AI, we can now identify 30We are just beginning to tap into the potential of machine learning. The road ahead is a rocky one but there is definitely a bright future on the other side.

1.2 Bram Dedrick

The articles I chose to talk about are Cade Metz' "Google Researchers Say They're Learning How Machines Learn" from the New York Times in March of 2018 and Christopher Mims' "Should Artificial Intelligence Copy the Human Brain?" form the Wall Street Journal in August of 2018. Both articles focus on neural networks, the computer algorithms that allow these artificially intelligent computer work. Back in the middle of the twentieth century the idea was first proposed as a computer that would behave like a human, thus the name neural network, derived from the neurons in our brains. Although these neural networks work very differently from the neuron in our head, we do know that we do not completely understand how each network works.

In Cade Metz' article he talks about the team at Google that is trying to go back and discover how neural networks work and how these artificial intelligences derive their answers. The neural networks often make mistakes, and until this Google team, or another team looking at the same topic, figures out the exact path the computer takes, we may never understand how to make the machines better. One way the Google team is going about this is by using visuals, something a general audience could understand. In some pictures the machine recognizes colors, lines, curves and shapes and begins to parse together what that picture is of. On occasion however, the computer can be tripped up and produce a wrong answer, and it will never realize it until a human points it out.

In Mims' article he talks to a few computer engineers who were asked if neural networks should copy the human brain. Essentially, they said no, neural networks work differently from humans in that they are a series of algorithms that when put together reflect something similar to a human thought. However, these algorithms all have a human element in that they are designed by a human at some point along the way. They all have some reflection of whoever built them, whether it be one person or a team, the process is going to be somewhat unique, as not all programmers think the same way.

Neural networks were a rather new subject to me before the start of this class, by no means do I consider myself to be an expert yet, but I think that I have gotten a pretty good grasp on how they work and what they are. In both Metz' and Mims' articles they make the information extremely approachable but also very informative, by talking to experts who are extremely well versed in the field but not filling their articles with too much technological information. I also found it very interesting how the earliest origins of machine learning date back to when computers were still just on the brink of their invention. Some of the earliest people to believe in machine learning thought that it would progress much faster than it has but that by no means deems it a failure. Machine learning has come a long way since the mid twentieth century to where it is now, but it still has a long way to go.

1.3 Tony Ferenzi

1.3.1 Benefits of Machine Learning

Machine learning is revolutionizing the world as we know it. It is simply a reliable system the reduces error over time and can track multiple sets of data. Not only is it simple, but it can also be applied to many different areas. From artificial intelligence to collecting data on a variety of different topics, machine learning has improved and is continuing to improve our lives today.

For example, many people depend on electricity for plenty of different reasons. Whether it is used to power our phones and computers, our heating and air conditioning units, or simply the lights in our homes, we all depend on electricity in some way. Due to this dependency, scientists and engineers have invented things like solar power and wind power to help generate more. Focusing on wind power, it's a simple way of generating power until you consider all the variables that come into play involving wind. Whether it is the weather or a light breeze, wind can be very unpredictable. This unpredictability actually increases the cost of running these massive turbines, and therefore increases the cost of electricity for the public. To fix or lessen this problem, researchers invented a type of machine learning system called accurate short-term wind speed forecasting (STWSF). The system is actually a hybrid machine learning system, made up of a method centered on wavelet transforming and a machine learning technique focusing on a FA network. This system essentially predicts wind speeds in order for engineers to plan ahead for the changes with an incredibly low chance of error when compared to other models. As the error lowers, so does the cost.

Machine learning has even been effective in keeping people healthy and free of diseases, like in Singapore. In 2016, there was a dengue virus outbreak in Singapore due to very high temperatures. This disease is carried by mosquitos and can spread extremely fast and far because of it. To combat this, the Singapore government has put a forecasting system into use created by Singapore's National Environment Agency. The system uses a machine learning algorithm called least absolute shrinkage and selection operator. Basically, the algorithm collects data from multiple variables like recent cases of the virus and weather to accurately pinpoint where an outbreak may soon occur. Then the government can warn the areas population to prepare for it and prevent further infections. Researchers are continuing to adjust and make improvements to the system in order to make it even more accurate and effective than it already is.

Machine learning also keeps us healthy by protecting us from certain pollutants. There is actually not much research on what the effects of mixed pollutants are on people's health. So, researchers created an ERS, or environmental risk score, which is a system that uses machine learning to predict what certain pollution interactions may do to someone's health if they are exposed to it. The machine learning algorithm again solves problems in the system by reducing and

bypassing certain errors while offering an array of tools to consider during data collection. Through these predictions, researchers can safely predict and prevent dramatic pollution problems while helping keep people and the environment safe.

Today's world is already a safer and healthier place due to machine learning fixing problems that we could never solve before. Nobody is perfect and errors can be over looked without caution. Machine learning helps us reduce these errors consistently over a wide variety of topics. Whether it be electricity or health, machine learning has changed our world for the better and hopefully will continue to do so.

1.4 William Golden

1.4.1 Humans: The Teachers of Technology

Without mankind, technology would not exist. Computers are essentially the birth of man, and as such, we must nurture and teach them as if they were our own flesh and blood. There are obviously obstacles we must overcome before we are able to literally "raise up" technology, but scientists and engineers are making great strides towards how computers learn and how we teach them.

A big obstacle is that computers don't necessarily learn the same way humans do. Even us humans don't have the perfect idea of how to create a human brain or why we think exactly the way we do. This becomes problematic when trying to build a computers knowledge. However, we are coming closer to finding out how to build computers like our own brains.

Computers are able to hold much more specific data than the human brain, so engineers suppose that one way technology can learn is by first gathering an extremely large database of information then narrowing it down to solve problems or create ideas. This process is known as a "top-down" way of learning (information becomes funneled down into a solution). On the other hand, computers may also learn by building on bits and pieces of simple data until eventually creating a solution. This method of computer learning is known as bottom-up or deep learning (information starts at a minimal level, then progressively grows larger and upward). Its important that we analyze these kinds of learning process if we are to begin building and teaching machines.

Scientists are looking at these concepts of "deep-learning" and and "top-down" ways of learning to solve the best way computers should learn. In recent studies, psychologists and engineers are working together to try to pinpoint how children learn in hopes of creating an algorithm for machines to learn effectively (Gopnik, 2017). In this specific research study, Alison Gopnik, a professor of psychology at the University of California, Berkeley, monitors the thinking process of children and how they learn about the world, and then tries to replicate the same process of developing information digitally.The experiment is just a milestone in the long road ahead to help progress machine learning, but ultimately, this idea of using human-based learning techniques for computer learning could possibly be the key to furthuring machine learning.

Other amazing discoveries humans are making on how to teach technology is through the use of body language and action. Just like parents teach a child how to throw a baseball, we could potentially teach machines the same way. In 2016, David Vogt and his team at the Freiburg University of Mining and Technology in Germany were able to successfully train a robotic arm how to build with lego

blocks by having the robot watch someone use the legos (Rutkin, 2016). His team used a camera and motion tracking devices to set up a visual learning mechanism on the robot, allowing it to follow and learn from humans.

Here is another similar example of machines learning through visual cues. A group of researchers from multiple universities in the United States and Australia taught a machine how to cook by showing it cooking videos on Youtube (Mamiit, 2015). The robot was built to be able to recognize basic cooking objects in the videos and learn how they move and interact with the hands of the humans. Next, the robot was able to mimic the movements and cook by itself. The article also talks about the dangers of AIs becoming too self-aware, but the majority of the article implies that this innovation is positive and is opening the possibilities of artificial intelligence and learning.

The ways in which we teach technology and how technology learn are expanding. Things such as developing machines to learn the way humans do is one example of innovative growth. Developing machine learning through visual cues by humans are also another way in which we are shaping the future of how man and technology can work together. Truly, the projects and articles discussed in this essay provide only a miniscule amount of information about machine learning today. It is a large and growing field and it is vital that we continue to find new ways to develop machine learning with humans so that we can better improve the quality of life for mankind.

1.5 Yuan Hong

Machine learning article Yuan

Machine learning is a computing system based on a simple set of fixed rules that can learn from input data. Learning machines are systems that implement mechanical learning. Importantly, we emphasize that mechanical learning is based on a simple set of fixed rules that distinguish it from machine learning, which is a software based on complex mathematical operations that often requires manual optimization or adjustment of the software. Here, we will discuss the basic issues and principles of the mechanical learning system and try to set a framework for further research. As in the case of Turing, in order to achieve mechanical learning, we propose two directions: one is to try to implement a learning machine, and the other is to try to describe mechanical learning. Whether the ability of the machine can surpass people's ability, one of the main arguments of many negative opinions is that the machine is man-made, its performance and movement are completely written by its designer, so its ability will not exceed the designer himself anyway. This kind of opinion is true for machines that don't have the ability to learn, but it is worth considering for a machine with learning ability, because the ability of this machine is constantly improving in the application. After a while, the designer himself I dont know the level of its ability. However, learning machines are something interesting and useful to our society in whole. For example, technology are doing successful performances in different tasks. Now it utilizes mass spectra data which performed by human language. The sensory includes smell and taste. However, a large scale of sensory is not able to complete or perform. But the experiment shows us how amazing this technology could analyze. And this the first one that uses this method to perform something like that. In the second example, human Locomotion needs graphic profiles and discrete variables. People could not present the full size of the data. So the participants developed a self-learning machine algorithm for getting the data. It successfully predicts by ninety-three percent correct. And to the author, it's the first study to optimise the algorithm. Machine learning are almost at any field of research, as the time passed on and it is more likely to become a medical tool. People could use this to preserve a better medical care, no matter men or children. Research shows that this method is getting more reliable and becoming a potential value in the future. The range of applications of various learning methods has been expanding, and some have formed commodities. Knowledge acquisition tools for inductive learning have been widely used in diagnostic subtype expert systems. Connected learning is dominant in acoustic image recognition. Analytical learning has been used to design integrated expert systems. Genetic algorithm and reinforcement learning have a good application prospect in engineering control. Neural network connection learning coupled with the symbol system will play a role in enterprise intelligent management and intelligent robot motion planning. In these articles, they are share the same quality, that they are serves to make humans life better, and to make workers in different fields work easier go through the process

of machine. All they need to do is to input the variables in the machine, and through the program, machine will get the answer you want, in a short amount of time. This not only saves people time, but also sometime will solve problems that are complicated in many ways.

1.6 Easton Jensen

Author Madhavi Ramani for Wilson Quarterly, wrote "Art(ificial Intelligence)". This article talks about the advancement of artificial intelligence in the art department. She takes the example of, "the next Rembrandt" which is a big project that takes art works from all types of sources and makes new images. In this machine there is a generator and a discriminator that makes the image and then the other makes critics to help make better images. This project is funded by Microsoft, ING Bank, and a Dutch advertising firm. Another project by Googole named Deep Dream can replicate famous artists styles so well that they have been sold for thousands of dollars. She then talks about a program called AARON. The man who made AARON is Harold Cohen. he refined AARON for over forty years. She takes an experience of Cohens that talks about the color in an image. Cohen talks about how a person can't imagine a picture with color in your mind but AARON can put color in the image it is going to make before it is even made. Ramani realizes by the end of the article that AI can't replace art. She justifies that AI can enhance people's creativity and it will never take it away.

Another article that shows the uses of artificial intelligence is "Artificial intelligence: past and future" by Moshe Y. Vardi, written in Communications of the ACM. Talks about the victories of certain AI against human competitors. First it talks abouts the victories of Deep Blue, a supercomputer by IBM, against chess champions Booby Fischer and Boris Spassky. Vardi states that most dramatic chess match was the rematch against Garry Kasparov. The first game Garyy Kasparov won and was very confident in his skills. However in the second match he got smacked by the supercomputer. Another IBM creation played Jeopardy! against two players that were again champions at what the AI program was doing. In 2000 an article was written about AI and how it was going to take away our jobs. This article wasn't minded much until in 2011 another article was written about the same thing. This caused people to start paying attention to the advances in artificial intelligence. Vardi ends the article with the concern that these people may be right and even though AI is harder than the first pioneers thought the progress made suggests that maybe there is no need for the human in the future.

Gaming has also come into the world of AI to help with driver-less cars. In the article "Gaming Machine Learning: Game simulations are driving improvements in machine learning for autonomous vehicles and other devices" talks about the research being done by playing games like TORCS and Grand Theft Auto V. The challenge is that there are so many things a driverless car has to detect, rangeing from a stop sign covered by something to snow on the road, covering everything on the road. THis work is possible because of the rise of GPUs and CNNs. Graphics processing units and convolutional neural networks allow the computer to see all the images and compare the actions and possible scenarios to make a better algorithm. A profesor from Princeton University says that

in the virtual world you can create that specific situation taht may be hard to recreate in the real world. He calls them “corner cases” and they represent the situations that lead to the greatest number of crashes. These simulations are gaining headway to help, “ eliminate the cost, time, and human resources” that it takes to set up these situations and collecting data. But driverless cars are not the only things that video games and simulations have been used for. Things like robots, drones, and agents have been looked at to improve by the virtual world.

All three of these so some ways that AI has been used today. From art that looks like an orginal, Artificial Intelligence playing against poeple in games of the mind, and video games bieng used to make better driverless computers. The world is progressing in a way that AI can be the future in the next 30 years if it kepps advancing at the rate it has been for the past 50 years.

1.7 Rodrigo Martinez

1.7.1 Machine Learning in Medicine

Machine learning in the medical field has already done a lot to help out society and yet there is still much to do with it. Having algorithms to organize files of thousands of people with their medical information or using technology to treat cancer are just some ways that machine learning helps. These articles discuss about the possibilities and the accomplishments that machine learning has made. For example, an article made by Techemergence called "7 Applications of Machine Learning in Pharma and Medicine" (July, 2018 Daniel Faggela) talks about that machine learning is able to make disease identifications and diagnosis. They mentioned how IBMs Watson and Googles deepmind health are trying to also contribute in this. Treatments also play an important role. Radiology and radiotherapy are used to help detect cancerous cells around your body with the help of machine learning. Epidemic outbreak predictions is also a thing that machine learning technologies are being used. This helps to predict where an outbreak will occur and how will it spread. By using variables such as temperature or the number of cases that occur, it is able to create a simulation about the disease outbreak.

Machine learning has also brought us to create a accurate prediction of someones plausible heart disease or diabetes. In "How Machine Learning Is Helping Us Predict Heart Disease and Diabetes" (May, 2017 Yannis Paschalidis) Boston University was trying to tackle on this problem and see if they can contribute to helping out by using machine learning. Yannis claims that they have found a way to find these diseases in "a year in advance with an accuracy rate of as much as 82%." . This is only using medical forms and records. The percentage would increase even higher if personal information is provided to them. This will be able to distinguish who will need to be hospitalized and not. Which may save the U.S 30.8 billions of dollars since it will prevent unnecessary hospitalization and benefit patients overall. 9 billion dollars goes for the patients with heart disease and 5.8 billion dollars goes for the patients with diabetes. Overall we are wasting too much money on unnecessary needs for hospitalization. This article also talks about how the team is also working alongside the Department of Surgery at The Boston Medical Center. In this case they are using machine learning to predict readmissions within the 30 days of general surgery (Yannis). Machine learning in the medical field has also helped people with disabilities in several ways. In Machine learning opens up new ways to help people with disabilities (Tom Simonite March 2017) Tom mentions examples such as how closed captions exist in television and youtube to help the deaf understand. Companies like Google and Facebook are striving to reach technologies that will benefit the disabled. IBMs Watson is also participating in this by using a tool called Content Clarifier to help disabilities such as dementia or autism. Basically this tool helps people with these disabilities understand figure of speech by replacing the terms used

into simpler direct words. Theres a similar project in which is using the same idea but is being implemented in major social medias such as Gmail or Facebook. Although it may not seem as much it is definitely having a positive impact on those who need the assistance and by time technology will only improve. This will make more of our goals achievable and possible.

Sources: https://www.techemergence.com/machine-learning-in-pharma-medicine/ https://hbr.org/2017/05/how-machine-learning-is-helping-us-predict-heart-disease-and-diabetes https://www.technologyreview.com/s/603899/machine-learning-opens-up-new-ways-to-help-disabled-people/

1.8 Matt Morrical

1.9 Ella Nelson

1.10 Koichi Okazaki

Self-driving cars

The dream of the self-driving car existed prior to the invention of the car. Leonardo De Vinci left the sketch of a rough blueprint of self-driving automobiles. In the 20th century, people developed the technology of self-driving cars such as the radio-controlled vehicle and smart highway technology. However, most of the inventions werent actually embodying the norm of self-driving cars by applying technology of machine learning. Realization of the self-driving car had been no more than dream for human beings until 1980s. A team at Japans Tukuba Mechanical Engineering laboratory equipped cameras on the car in order to analyze the surrounding environments with computer system and that gave the self-driving car sight. The prototype was able to run about 20 miles per hour and it was programmed to follow the white street markers captured by built-in cameras. After that invention, a German aerospace engineer named Ernst Dickmanns applied the technology of machine learning to the self-driving vehicle and enabled it to drive at high speed. This was achieved by the Mercedes van with cameras and sensors that offers data into a computer program to complete tasks such as adjusting the steering wheel, brake, and throttle. The innovation by Dickmann induced large investments in western countries and the market scale grew dramatically.

Because of the amount of investment and development of machine learning, self-driving cars are becoming capable for practical use. Nowadays, we can see cars which are adopting the technology of full self-driving vehicles in particular tasks such as steering and accelerating and it is effective to reduce driving errors and sudden accidents. However, the driver must be engaged with the driving tasks and observe the surrounding environments at all times. Therefore, the functions are going no further than supporting the driver. Nevertheless, it is said that fully self-driving cars will be practical in few years. For instance, Waymo the subsidiary of Alphabet Inc. completed the worlds first fully self-driving trip on public roads in 2015, and most of the major vehicle companies declared that they will develop the self-driving cars by about 2020. If fully self-driving cars became used in our lives, it will reduce accidents dramatically, people with disabilities and the elderly will be able to use vehicles safely, and the options of public transportation will increase, etc.

However, there are several obstacles to put fully self-driving cars into reality. For example, if an accident occurred and if it was unavoidable to hit someone, the program might have to choose the victims, in order to save others. Thus, this is considered as a crucial ethical issue since program is assuming to kill somebody in that situation and also where the responsibility goes isnt clear; is it the car, producer, or consumer? Additionality, since cars, trucks, and buses begin to drive themselves, people who working with these vehicles will be unemployed. According to the U.S. Bureau of labor Statistics, potential loss of jobs will be more than 2.6 million jobs. Then, most of those workers are considered as low-

skilled labors, and it would not be easy for them to get another job.

Citation

History of Self-Driving Cars, Tuan C. Nguyan, ThoutCo., April, 2018

https://www.thoughtco.com/history-of-self-driving-cars-4117191

Waymos fully self-driving vehicles are here, Waymo Team, November, 2017

https://medium.com/waymo/with-waymo-in-the-drivers-seat-fully-self-driving-vehicles-can-transform-the-way-we-get-around-75e9622e829a

The Unintended Consequences of Self-Driving Cars, Adam Hayes,

https://www.investopedia.com/articles/investing/090215/unintended-consequences-selfdriving-cars.asp

'Self-Driving Cars: An Ethical Perspective', fall2015, Penn Bioethics Journal, vol. 11, no. 2, p. 8.

1.11 Jakob Orel

"Delving into Deep Learning" written by Brian Hayes for *American Scientist* in 2014 provides an inside look on deep machine learning and the processes of complex algorithms. The article is intended for the average American who is interested in learning how these systems work. He informs the audience of the origin of the deep learning environments and breaks down the topic to the most basic input and outputs. Hayes then explains how much more complex the topic of deep learning is with the connections of the neural networks. He gives examples of how these systems are used in the real world such as Apple's Siri. Hayes also explains that the future of these deep thinking networks depends on the investment of companies and the further research needed. This article was intended to give the audience a brief idea of how deep learning machines work and how they can be applied to the real world.

"#Flu" written by Rachel Berkowitz for *Scientific American* in April of 2018 explains how machine learning can be used to predict influenza outbreaks. Berkowitz informs the average audience about how social media posts can be used to identify trends. These trends can help identify where and when new outbreaks will occur. Multiple other systems have been used to comb through flu related words, but this system is different because it sorts through non-flu related words instead. Berkowitz explains how many people believe this could be very useful in the near future when it is further developed. I believe this article was written to give an example of how machine learning can be applied to help in healthcare and around the world. It made me realize how interesting machine learning can be and its amazing uses.

Another article written for *New Scientist* titled "Voice calls combed for signs of disease" by Matt Reynolds in 2017 describes how a new program is using a collection of voices to detect diseases. A new software called Canary Speech combs through millions of calls supplied by a health insurance company with the medical history of each call attached. The software then can distinguish vocal cues that may result in diseases such as Alzheimers or even depression or anxiety. Some experts are skeptical of using this process because it uses a large collection of data instead of actual medical responses. Although skepticism remains, this could prove to be very useful to the healthcare industry to help research or diagnosis of patients. Reynolds intended the article for the average citizen without knowledge of machine learning.

The key characteristic that I found in most articles on the use of machine learning in healthcare stated that it required more research and development before it could be efficiently applied to society. Canary Speech and the flu detecting software may prove to be extremely useful to make predictions about many health related issues in our world. These deep learning networks do require more testing and development to ensure they can be accurate and precise. Machine learning and deep neural networks can greatly advance our system of healthcare by using artificial intelligence and computers in the future.

1.12 Marcellus Parks

Marcellus Parks

The applications of machine learning are quite large in number. One of the more prominent uses of it is in the healthcare field. Machine learning can be exercised all over from prevention to predicting the best treatment. Cardiovascular disease has been the world leading cause of death for the last fifteen years. This is mainly due to the fact that most people do not have access to preventative measures and do not get diagnosed early enough for an effective treatment. In the article "Mobile Personal Health Monitoring for Automated Classification of Electrocardiogram Signals in Elderly", preventative measures were taken by developing a wearable electrocardiogram monitor that was linked with a smartphone app. The test consisted of one hundred older adults This monitor was able to accurately separate normal and abnormal ECG signals. It was discovered that because of its accuracy, this monitor could help prevent, diagnose early, and effectively treat heart diseases. Also because of its mobility and accessibility, it is much more cost effective than older models of the ECG monitor. This would increase the access to life saving services. Sometimes, people are not so lucky. There are many cases of people who have diseases that could not really be prevented. Machine learning can be there to help in these scenarios as well. Precision medicine is a swiftly growing field and doctors all over the world are looking to standardize and have the ability to quickly choose the best treatments for patients. In the study "Open source machine-learning algorithms for the prediction of optimal cancer drug therapies", machine learning algorithms were combined with the recursive feature elimination approach to create personalized drug responses for sixty human cancer cell lines. Applications with their models to open ovarian cancer dataset created predictions consistent with previously successful responses reported. Also they made the algorithms open source to quickly gain improvements and be able rapidly integrate their methods into other types of cancers with the changes from those willing to help. Both of these studies use machine learning to expedite medical procedures that overall, increase the lifespan of patients. In the ECG monitor, the app is able to detect early cardiac abnormalities, which will increase the chance of whatever cardiovascular problem someone might have being treatable. While the drug therapy algorithm can quickly decipher the best treatment options available to the patient, instead of trying a large amount of different drugs, which would consume time and increase the chance of fatal condition. It would also save the patients money. Instead of having someone spend hours trying to come up with the perfect treatment, the patients doctor could just enter the patients specific cell line and within a few moments the algorithm would come up with the best treatment saving the clinic money and in turn, the patient. The improvements to medicine through machine learning will be very significant, having a much higher success rate than methods of the past. The amount of success is one thing, but the main premise of machine learning in the medical field is to create more efficiency. These two studies show that, with

the proper utilization, machine learning algorithms can create faster acting and more cost efficient health care for doctors and patients alike.

1.13 Lydia Sanchez

In our society today AI has become more and more prevalent with the creations of personal assistants like Amazons Alexa and Teslas self-driving cars. AI and machine learning comes in many different forms shown in these articles. One a machine making its own pictures and the other making mask of people based on DNA .In both of these two articles it shows where AI is used to create something visual.

The GANfather: the man whos given machines the gift of imagination by Martin Giles is about a machine that can create its own picture by training from other pictures. The GAN is made by Ian Goodfellow and some friends. Their motivation was generative models did not create very good models like there would be blurry or missing a crucial part of the face. The GAN work by putting two different networks one is the generator, which produce a artificial output, such as photos, while the other network, the discriminator, compares the photo to real photos and tries to pick out the real one and they repeat this until the discriminator cannot tell which photo is real and which is fake. The GAN is really innovative because it can be proficient in making fake celebrities by just training with a couple hundred photos. In the future this could be used for optimizing the fields of science and engineering, like in the medical field sometimes they cannot get enough patient data to tell why a drug did not work, by creating fake data that would be almost be viable as the real thing. With these good attribute there are some bad ones like people could use the GAN to generate fake news and could create problems with cyber security. This article pretty easy to read a bit technical but the average person can still read it easily.

Mask Crusader by Alec Wilkinson is about Heather Dewey-Hagborg portraits of in form of realistic masks made from the DNA of people who dropped cigarettes and chewing tobacco on the ground around New York City. Dewey-Hagborg say its not a study on how to use DNA to make faces, it is supposed to make you think about privacy. As an art major she got interested in machine learning and then continuing to graduate school to study AI, she also got interested in surveillance. After a therapy session it sparked something in her and she started noticing the cigarettes and chewing gum all over and wondered how much can I of someone by what they leave behind?. After she learned how to dissect and analyze DNA she found out you can tell a lot about a person by their DNA, like eye color and distance of eye from each other. It cannot tell what the color of your skin but can tell what your ancestry is. Then she uses a 3D Printer to make the masks. This article was very easy to read, for the average person to inform them of something new happening.

The two article Mask Crusaders and the GANfather are both very different but they do have many similarities. In both of the articles AI has to generate something new or unique, in the Mask Crusaders the AI has to make a face from and in the GANfather has to make new pictures. GANfather and the Mask Crusaders both Generate faces with AI. They both are new technology

using machine learning. In both articles they use machine learning and AI to create something visual.

1.14 Tiff Serra-Pichardo

Tiff Serra Professor Leon Tabok CSC131 3 September 2018 Applications of Machine Learning in Medicine Machine learning is a form of computer science that gives computers the ability to learn using statistical data. Machine learning is currently being used in many fields, one of the most prominent and emerging fields in which it is used is in the medical field. According to artificial intelligence (AI) researcher Regina Barzilay in an interview published by New Scientist, Virtually every aspect of life today is regulated by machine learning, whether you know it or not. The only area that isn't is healthcare, which involves a lot of prediction tasks. The medical field has notoriously been difficult to incorporate machine learning into. Health care is a highly personalized field. Barzaley best defines it as, When your doctor tries to find you a treatment, they look at different clues together and make a prediction. With personalization, which we're all trying to achieve in medicine, the goal is matching you and your unique characteristics to the correct drug. The advances in medical field have been slow and, in most cases, not worth the expensive upkeep of having an AI. A field that has made the most advances in regards to the implementation of technology, within the medical field is oncology. One of the most well-known examples of artificial intelligence having some degree of success of is Watson. Watson was developed by IBMs DeepQA project. One of the main advances of Watson is its ability to utilize and understand non-standardized text like doctors notes. Watson also does not suggest what a patients treatment should be, rather, it presents treatments in a ranking. Watsons primary goals in the field of oncology is to deal with the big data problem. By one estimate, health information – electronic health records, insurance claims, images such as CT scans, vital signs of people being remotely monitored by hospitals or smartphone, gene sequencing results – will grow to the equivalent of about 500 billion four-drawer file cabinets by 2020, from a mere 10 billion in 2011. (Hobson 13). The overwhelming data that doctors and scientists are collecting in regards to their patients must be sorted and analyzed. These large amounts of data make it impossible for anyone to look through all of it; that is what makes Watson and other AIs a necessity. While there is some controversy and concern that AIs will not be able to make diagnosis due to the highly personalized data involved in individual peoples health care, AI, as of today, has not reached the point of making proper diagnosis's and is still within its infancy. In the future there is hope that AI can permanently alter the health care industry so that the positives out way the negatives.

Citations Benios, Thania. At the Edge of Life's Code. EBSCOHost, Scientific American, Apr. 2008. Hobson, Katherine. Meet Dr. Watson. EBSCOHost, News Digital Weekly, 23 Aug. 2013. Simonite, Tom. IBM Aims to Make Medical Expertise a Commodity. EBSCOHost, MIT Technology Review, 2014. Temming, Maria. Computers Can Diagnose Eye Problems. EBSCOHost, Science News, 31 Mar. 2018. Whyte, Chelsea. I Teach Machines to Hunt down Cancer. EBSCOHost, New Scientist, 21 July 2018.

1.15 Austin Stala

There are many applications for machine learning. One of the most prominent uses for machine learning is in the medical field. One way that machine learning can be used in helping to predict the effectiveness of disorder treatments. Another way that machine learning can be used is to classify diseases and find the best treatments for them. The final use of machine learning I will cover in this essay is using machine learning in the medical field is to help with diagnosis and treatment for veterinarians. The first use of machine learning that I will talk about in this essay is using machine learning to help predict the effectiveness of treatments for substance abuse disorders. There are several ways that machine learning can predict the effectiveness of treatments. One of these ways is called super learning or SL. Super learning is a method that takes all predicted models and combines the to attempt to find the most effective treatment. A second way that machine learning can be used to predict the effectiveness of treatments of disorders is called Artificial Neural Networks or ANN. Artificial Neural Networks fundamentally work similarly to how Super learning work. The biggest difference is that Super Learning uses predictions and Artificial Neural Networks use a set of predetermined traits. Artificial Neural Networks are used to more broad outcomes to give more general treatments doctors. The second way that machine learning can be used in the medical field is using it to help classify diseases. The International Journal of Recent Research Aspects believe that the advances in machine learning can be used to help doctors sort through medical information to help classify diseases. The article written by the International Journal of Recent Research Aspects states that a way to do this is to create a desktop app or a web plugin. This can be used by selecting the symptoms/disease then it will gather and display the newest and most reliable documentation to help a doctor come to a conclusion for a treatment. The third way that machine learning can be used in the medical field is by veterinarians. There have been many ways that machine learning has been used for veterinarians. These range from treating diseases and diagnosis to the growth and production of animals. One of these is tracking herds of cows based on how likely they are to be exposed to disease. A second way is to find the factors that most affect the success of hatching eggs. There are many problems with using machine learning. One of these problems is how limited machines can be. Machines can still get things wrong by misidentifying things and coming to miss conclusions. Second they can have issues finding reliable information. Despite these limitations machine learning can be extremely useful in the medical field. In conclusion there are many uses that machine learning can have in the medical field. One of these uses is helping to predict the effectiveness of disorder treatments. The second use of machine learning talked about in this essay is that machine learning can be used is to classify diseases and find the best treatments for them. Third use for machine learning is using machine learning in the medical field is to help with diagnosis and treatment for veterinarians.

1.16 Nicole Trenholm

Nicole Trenholm September 3, 2018 Machine Learning in Medicine Imagine walking into the doctors office, reporting your symptoms to a machine and within minutes, getting a diagnosis and suggestion for treatment. That is the future of medicine with machine learning. The main focus of researchers is how they can use these machines in order to accurately diagnosis a patient based on their records, symptoms, and traits. In the articles I have found, they narrow their studies as one focuses on Post-traumatic Stress Disorder and the other one on Schizophrenia disease. By applying machine learning to these diseases, the authors have found ways to identify and treat them on top of other things. In the article Bridging a Translational Gap: Using Machine Learning to Improve the Prediction of PTSD by Karen- Inge Karstoft and colleagues, they discuss how using supervised machine learning can help to uncover interchangeable, maximally predictive combinations of early indicators of PTSD. In the medical field they have found that there are a set of risk indicators that can be observed best shortly after a traumatic experience occurs, however research has failed to propose clinically useful, personalized predictors. To overcome this set back, Karstoft and others make a goal to address methods of PTSD that accommodate multiple combinations of risk indicators or in other words use sets of data in order to accurately predict post-traumatic morbidity. The authors report about reaching this goal by applying machine learning modeling to a large dataset as they evaluated the methods ability to identify multiple, equally predictive sets of variables. With machine learning being able to increase prediction versatility, Karstoft and colleagues concluded that medicine can now take a step forward towards developing algorithmic, knowledge-based, personalized prediction of post-traumatic psychopathology. The article Machine Learning Classification of First-Episode Schizophrenia Spectrum Disorders and Controls Using Whole Brain White Matter Fractional Anisotropy by Pavol Mikolas and colleagues is geared more toward the diagnosis of schizophrenia. The authors combine apply machine learning to MRI scanning as it would allow for subtle patterns of diseases on a single subject to become easier to identify. In turn this would open up a new door into the potential of MRI in psychiatry. It will also be a breakthrough in the medical field in the sense that when combining machine learning and brain imaging data, there is a chance of diagnosing schizophrenia early which may improve prognosis and treatment outcomes. The data explored in this article is of 77 first-episode schizophrenic patients and 77 controlled, healthy patients. By looking at the subjects brain diffusion tensor imaging (DTI) data the machine learning programs were able to differentiate the first-episode participants from the controlled participants. The essay concluded that from the abilities of the simple machine learning program used, there are many improvements it can make in both the medical and psychiatric field and many more possibilities that can arise when combining this knowledge with more complex machine learning programs. After reading these articles I found myself persuaded on the huge advantages of machine learning and the

wide range of capabilities it possesses. I enjoyed reading about how much data is being collected to make it easier and more efficient to diagnosis those with these specific diseases because that work being done has so much potential to help people. Both of the articles I studied were higher level articles with a main audience of professionals in the fields of medicine and psychiatry. The content was dense but still easy to grasp for those that do not have much background knowledge. The goals of the studies done were to use machine learning as a tool to explore ways to diagnosis patients in more proficient ways. In both studies they were successful in finding ways to apply machine learning in medicine and opened the door to other potential uses of it. With the help of machine leaning, researchers are able to go to new lengths and educate others on how they can make advances in new fields. I found articles that explored new sides of medicine that I have personally never considered. The authors go past the traditional uses of machine learning and focus on finding ways to improve the lives of patients with diseases such as post-traumatic stress disorder and schizophrenia. Studies done such as the ones mentioned have potential to open new doors within not just the medical field but other fields as well. Technology has advanced dramatically in the past few years and has given opportunities to many fields and in the future, disease diagnosis will be only a small sample of the capabilities of machine learning.

1.17 Maddy Weaver

Madison Weaver How is machine learning being used today?

In todays world technology is constantly changing through the use of machine learning and artificial intelligence. Our smartphones are becoming smarter and we can reach information with the click of a button, they can even recognize our faces. Many people use technology as an outlet from their daily lives and even participate in interactions with artificial bots. If you take a closer look into the evolution of society, it is moving much faster than we even realize with the help of machine learning. Perhaps one of the most obvious advancements in machine learning lies right within our own hands; our iphones/smartphones. Why is this? Smartphones are getting more resources every year, specifically with their cameras. Cameras use to be seperated from phones, nonetheless way less smart. Farhad Manjoo describes that cameras used to be like eyes disconnected from intelligence. But now, cameras are actually recognizing things and putting one and two together. Some people use Apples newest iphone to recognize their face as a way to unlock their phones. A startup called Lighthouse Al even wants to incorporate face recognition into home security programs, so that your house will recognize you when entering your home. These resources are changing the way that we live. Not only are people using machine learning for usefulness, they are even starting to use it for intimacy. There are many programs that allow people to speak with an artificial bot. The more that these bots speak to people, the more that they learn. Why do people do this? According to John Markoff and Paul Mozur, people have admitted to not having anyone to talk to besides the bot, Xiaoice, which was created in China. This is worrisome to researchers who worry that humans are beginning to lack intimacy and society is heading in a negative directions. These changes reflect how our society is changing, and how people who feel lonely are taking advantage of it. All of these changes are a result of new machine learning and the advances that computer scientists and researchers have come across. Its amazing what kind of work that they have done. A group of researchers in Googles secretive x laboratory created a large neural network for machine learning by connecting 16000 computers to let loose on the internet. They were able to recognize cats on the internet. This type of technology did not come out of nowhere and has been processing for years. Its amazing how far we have come.

1.18 Peter Weber

Applications

Machine learning is used for situations where there are large amounts of data that need to be processed and it does not make sense for the programmers to tell the machine what to do in every situation. Instead, the machine makes its own rules to follow depending on the situation.

One specific application of machine learning is in the Large Hadron Collider where it processes information from experiments. The amount of data that is collected from particles colliding is far too much to store, so instead a computer looks through and picks out the information that should be kept. There is a lot of background data that the system filters out as well as results from more common events that they don't require data on. The computers they use have learned the subtle signs that indicate what they are looking for in the large mess of data that comes in. It records what it needs when it recognizes something important, and discards the rest.

Another application of machine learning is sorting through biological data. Scientists use machine learning to go through thousands of past experiments to learn how genes work. The machine learning program goes through the information on how genes work, comparing it to genes it already knows have been linked to a higher risk of autism, to figure out which genes are most likely to signal or cause autism. The data could eventually be used to predict autism at an earlier age or to even find the cause and potential cure for it.

Chapter 2

Recommendations: How to learn more about machine learning

2.1 Sheldon Branch

2.1.1 Course 1: Machine Learning

Link

Institution: Stanford University Instructor: Andrew Ng: Co-founder, Coursera; Adjunct Professor, Stanford University; formerly head of Baidu AI Group/Google Brain Cost: The class is free but you can choose to pay $79 if you want a Course Certificate. If purchased, the certificate will be added to your accomplishments page, from where it can be added to your LinkedIn profile. Some assignments may not be available without first enrolling for a certificate.

Course Description: About this course: Machine learning is the science of getting computers to act without being explicitly programmed. In the past decade, machine learning has given us self-driving cars, practical speech recognition, effective web search, and a vastly improved understanding of the human genome. Machine learning is so pervasive today that you probably use it dozens of times a day without knowing it. Many researchers also think it is the best way to make progress towards human-level AI. In this class, you will learn about the most effective machine learning techniques, and gain practice implementing them and getting them to work for you. More importantly, you'll learn about not only the theoretical

underpinnings of learning, but also gain the practical know-how needed to quickly and powerfully apply these techniques to new problems. Finally, you'll learn about some of Silicon Valley's best practices in innovation as it pertains to machine learning and AI. This course should be taken first because it introduces a lot of concepts as well as the software Matlab, which is used in the subsequent classes.

Difficulty: This course serves as an introductory course or a refresher course about machine learning. It will be relatively easy for someone with a computer science background but not beneath their level.

Effort: This class is not very rigorous. Every week covers 1-3 lessons, each of which starts with an introductory video. Homework assignments require you to watch short videos and read short passages about the different modules in each lesson. There is one graded assignment due at the end of each module. Duration: The classs suggested duration is 11 weeks but the course is self-paced. The class is available for 180 days after enrollment. Software: This course includes programming assignments designed to help you understand how to implement the learning algorithms in practice. There is a module that introduces you to the software, Matlab and Octave, and shows you how to submit an assignment.

Subjects: This course provides a broad introduction to machine learning, data mining, and statistical pattern recognition. Topics include: linear regression, and logistic regression, linear algebra, anomaly detection, supervised learning (parametric/non-parametric algorithms, support vector machines, kernels, and neural networks), unsupervised learning (clustering, dimensionality reduction, recommender systems, deep learning), and best practices in machine learning (bias/variance theory; innovation process in machine learning and AI. The course will also draw from numerous case studies and applications, so that you'll also learn how to apply learning algorithms to building smart robots (perception, control), text understanding (web search, anti-spam), computer vision, medical informatics, audio, database mining, and other areas.

Reviews: See reviews here.

- "I have a background in AI but have not done any work with it in many years. I found this course to be a great refresher and would recommend it both as a refresher and as an introduction."
- "A motivational introduction to the machine learning field as a whole. Clear and concise lectures with practice problems equally enjoyable and challenging."

2.1.2 Course 2: Robotics: Vision Intelligence and Machine Learning

Link

Institution: PennX

Instructors:
1. Jianbo Shi: Professor, Computer and Information Science, School of Engineering and Applied Science
2. Kostas Daniilidis: Professor, Computer and Information Science, School of Engineering and Applied Science
3. Dan Lee: Professor, Computer and Information Science and Electrical and Systems Engineering, School of Engineering and Applied Science

Course Description: How do robots "see," respond to and learn from their interactions with the world around them? This is the fascinating field of visual intelligence and machine learning. Visual intelligence allows a robot to "sense" and "recognize" the surrounding environment. It also enables a robot to "learn" from the memory of past experiences by extracting patterns in visual signals. You will understand how Machine Learning extracts statistically meaningful patterns in data that support classification, regression and clustering. Then by studying Computer Vision and Machine Learning together you will be able to build recognition algorithms that can learn from data and adapt to new environments. By the end of this course, part of the Robotics MicroMasters program, you will be able to program vision capabilities for a robot such as robot localization as well as object recognition using machine learning. Projects in this course will utilize MATLAB and OpenCV and will include real examples of video stabilization, recognition of 3D objects, coding a classifier for objects, building a perceptron, and designing a convolutional neural network (CNN) using one of the standard CNN frameworks.

Prerequisites: College-level introductory linear algebra (vector spaces, linear systems, matrix decomposition), college-level introductory calculus (partial derivatives, function gradients), and basic knowledge of computer programming (variables, functions, control flow) is preferred, but students may also choose to learn it on their own. The class projects will be carried out MATLAB/Python, with C++ as an option. This class should be taken second because the first class will introduce you to Matlab and this one expects you to use Matlab a lot. This class also expects more prior knowledge of computer science. It is also more hands-on and requires you to do more intense projects than the previous course.

Cost: The class is free but you can choose to pay $349 if you want a Verified Certificate. The instructor-signed certificate has the PennX logo on it and

it can be used to verify your completion of the class, increasing your job prospects.

Difficulty: Advanced. This course is much more focused on machine learning applications in robotics and will introduce several new concepts.

Effort: This class does not require much effort. The course will take about 8-10 hours a week.

Duration: 12 weeks

Software: Matlab and OpenCV

Subjects: This course will cover the fundamentals of image filtering and tracking, and how to apply those principles to face detection, mosaicking and stabilization, how to use geometric transformations to determine 3D poses from 2D images for augmented reality tasks and visual odometry for robot localization, and how to recognize objects and the basics of visual learning and neural networks for the purpose of classification.

2.1.3 Course 3: Machine Learning Fundamentals

Link

Course Description: Do you want to build systems that learn from experience? Or exploit data to create simple predictive models of the world? In this course, part of the Data Science MicroMasters program, you will learn a variety of supervised and unsupervised learning algorithms, and the theory behind those algorithms. Using real-world case studies, you will learn how to classify images, identify salient topics in a corpus of documents, partition people according to personality profiles, and automatically capture the semantic structure of words and use it to categorize documents. Armed with the knowledge from this course, you will be able to analyze many different types of data and to build descriptive and predictive models. All programming examples and assignments will be in Python, using Jupyter notebooks. This course should be taken last because you will be doing very complex assignments with machine learning.

Institution: UCSanDiegoX

Instructors: Sanjoy Dasgupta: Professor of Computer Science and Engineering

Cost: The class is free but you can choose to pay $350 if you want a Verified Certificate. The instructor-signed certificate has the PennX logo on it and it can be used to verify your completion of the class, increasing your job prospects. Difficulty: Advanced. This course is about applying the

knowledge from the previous classes to make machine learning programs and write algorithms.

Effort: This class does not require much effort. The course will take about 8-10 hours a week.

Duration: 10 weeks

Software: Jupityr Notebook

Subjects: Classification, regression, conditional probability estimation, generative and discriminative models, linear models and extensions to non-linearity using kernel methods, ensemble methods: boosting, bagging, random forests, and representation learning: clustering, dimensionality reduction, auto encoders, deep nets

2.2 Bram Dedrick

2.2.1 Online Courses

In this section we will look at possible online courses focused around machine learning. We will compare their costs, both in time commitment and monetary commitment, the estimated complexity, the instructor of the course and other possible benefits one course may have over another.

1. This course, offered through edX, focuses on learning the principles and algorithms necessary for machine learning. It is offered by the Massachusetts Institute of Technology through their "MITx MicroMasters Program in Statistical and Data Science" and is taught by two professors form MIT. The first is Regina Barzilay a Delta Electronics Professor in the Department of Electrical Engineering and Computer Science. The second instructor is Tommi Jaakkola a Thomas Siebel Professor of Electrical Engineering and Computer Science and the Institute for Data, Systems, and Society.

   The course is taught over fifteen weeks at a commitment of ten to fourteen hours a week and has an "Advanced" level ranking on edX. The course is free to enroll and has the option of adding a verified certificate for an additional three hundred dollars.

   Learn more here.

2. This course, offered through edX, in an introductory level course that teaches some popular algorithms used in machine learning as well as other basic principles. It is being offered by Harvard University and is taught by Rafael Irizarry a Professor of Biostatistics at Harvard. The course length is four weeks, while only allotting two to four hours a week to this study.

   It is free to register, and a verified certificate is available for forty-nine dollars.

   Learn more here.

3. This course, offered through Coursera, gives an introduction to machine learning through a number of its uses. In this course you will learn to implement and work with introductory machine learning algorithms. You will learn these processes through information through basic datamining and statistical pattern recognition. Later you will learn how machine learning is being implemented by some of Silicon Valley's premier companies. The instructor, Andrew Ng, is the cofounder of Coursera, an adjunct professor at Stanford University and was previously the head Baidu AI Group/Google Brain.

   It is taught in four sections each with videos, quizzes, problems and examples. The course will take roughly fifty-three hours to complete and it

is recommended that you do roughly seven hours a week, but deadlines are flexible, so you can complete this at whatever pace you want.

Learn more here.

2.2.2 Beginning the Class

At the beginning of the class, it may be wise to show a short video explaining what machine learning is, where it is implemented, etc. For this I found an interesting video from the Crash Course YouTube channel in association with PSB Digital Studios. The video uses a number of graphics, examples, short clips and graphs to explain how machine learning works in its most basic sense and expands on that explanation for more complicated processes.

See the video here.

2.2.3 Machine Learning Courses Offered in Other Schools

1. Stanford University is currently offering a course on machine learning, CS229. In the class, students will learn topics similar to those offered in the online courses but in a classroom environment. The class will cover a large range of machine learning basics such as supervised learning, unsupervised learning, learning theory, reinforced learning and adaptive control. The class will also cover some of the implementation of the machine learning knowledge including robotic control, datamining, autonomous navigation, bioinformatics, speech recognition and text and web data processing.

   The course it taught by two professors Dan Boneh and Andrew Ng. Dan Boneh is a professor of Computer Science and Electrical Engineering as well as the Co-director of Stanford Computer Security Lab. Andrew Ng is Co-founder of Coursera and an Adjunct Professor at Stanford University. Andrew Ng is also the teacher of one of the Coursera courses I recommended earlier.

   Learn more here.

2. This machine learning course is being offered by Cornell University as CS4780/CS5780. Like many of the other courses this class focuses on introductory machine learning algorisms and techniques. This course description also lists plenty of helpful resources that could be useful when designing a course of your own. The course is taught using the programming languages of Julia and Python. It has two textbooks, "Machine Learning A Probabilistic Perspective" by Kevin Murphy and "The Elements of Statistical Learning" by Trevor Hastie, Robert Tibshirani and Jerome Friedman.

   The class also has all of last year's homework listed with solutions, a transcript of every past lecture as well as a listed placement exam for

the class. The professor is Killian Weinberger, an Associate Professor at Cornell, who earned his Ph.D. in machine learning from the University of Pennsylvania. He is also now part of the newly established "Cornell Center Data Science for Improved Decision Making."

Learn more here.

3. This course, EECS 545, offered at the University of Michigan and is designed to be a graduate level machine learning course. The topics, like many other machine learning courses, include supervised learning, unsupervised learning, learning theory, graphical models and reinforced learning.

   This course will also include sparsity and feature selection, Bayesian techniques and deep learning as research topics. Like many of the other courses this will have an emphasis on machine learning applications such as computer vision, datamining, speech recognition, text processing, bioinformatics, and robot perception and control. The textbook for this course is Christopher M. Bishop's "Pattern Recognition and Machine Learning."

   There are two professors for this course Honglak Lee, who received his Ph.D. from Stanford University in Computer Science, and Clayton Scott, Professor of Electrical Engineering and Computer Science and Statistics.

   Learn more here.

2.3 Tony Ferenzi

Machine learning is one of the most useful things that we have ever created. It can be used in almost any career from producing power to just collecting data on trees. The fact that it can bypass difficulties and minimize errors as it learns and adapts is amazing by itself. That is why so many schools and universities are opening programs to educate those who want to learn about machine learning so that they may find a career that suits their desires and dreams. Machine learning is the future as it is always improving, even running fully functional and adaptive robots. Here are a few courses that you could teach to students to improve their understanding of machine learning, or possibly learn something yourself.

I highly recommend taking ideas from the course offered by the University of Pennsylvania to start off any advanced class in computer science. The course is twelve weeks long, eight to ten hours per week on average and is actually free unless you want the verified certificate for 349 dollars. Some prerequisites for the class are:

- College level introductory linear algebra
- College level introductory calculus
- At least knowledge of computer programming

The simple idea of learning how to program robots is extraordinary by itself. It would definitely draw attention to the class even though it would be at a more advanced level. The course is essentially a mix of visual intelligence, which allows robots to visualize their surroundings as well as remember them, and machine learning, which allows robots to adapt to new surroundings based on previous data. By the end, the course promises that you will know how to program a robot to have visual capabilities, as well as object recognition. Students will surely enjoy learning about visual intelligence and in turn will appreciate machine learning and the benefits it can offer.

Next on the list is a course offered by Columbia university. This course lasts twelve weeks in total with about eight to ten hours per week. The price for admission is free unless you want the verified certificate for 199 dollars in total. The prerequisites for this class are:

- Calculus
- Linear algebra
- Probability and statistical concepts

- Coding and knowledge on data manipulation

This is a simpler class the involves basic and advanced topics during learning. Due to this, students will be challenged but also be inspired to overcome obstacles and learn new topics. The course will teach students about specific methods and models, like those that search engines use to recommend certain websites that may interest users over others. It will also go into different perspectives, like unsupervised and supervised learning that will further branch out into more research and topics for students to investigate. The course is actually split between supervised and unsupervised learning, the first half being supervised and the second being unsupervised. Some of the methods mentioned are linear and and logistic regression, tree classifiers, hidden Markov models, Gaussian mixture models, and so many more.

UC San Diego has its own course on the fascinating topic of machine learning. It is currently ten weeks long and is estimated to be about eight to ten hours per week. It costs nothing to enter, but again specifies that if you want the verification certificate, it will cost 350 dollars in total. The prerequisites include:

- Multivariate calculus
- Linear algebra
- Previous courses in the MicroMasters program

This class is a mix of basic and advanced topics. Again, this course goes in depth into the important topics of supervised and unsupervised learning, but also the theories behind them. What is exciting and different about this course is that it uses actual case studies used in society today, which can easily capture the attention of new studies as well as broaden views on a topic. Another major topic that the course promises to teach students is ensemble methods like random forests, boosting, and bagging. All these methods are unique and therefore might be a little more difficult for newer students, but introducing them early on will only improve their understandings of the subject.

Harvard University offers has a very interesting course on machine learning. In fact, Harvard is the only course on this list that was not listed as advanced. It is an introductory course that lasts four weeks and is about two to four hours per week. It is also one of the shortest courses, but that is because it is just a basic course that focuses on simpler topics. This specific course is free unless the student would like the 49 dollar verified certificate. This would be a great class for anyone who is new to the field and wants to further their studies. Another difference compared to other courses is that this one is part of the Professional Certificate Program in Data Science, which is a series of classes involved with machine learning. The courses only prerequisites are simply taking the previous courses in the program. The class states that students who will take the course will learn how to initiate cross-validation in order to avoid

overtraining and reduce errors during the coding process. Taking lessons from this course would highly benefit teachers and students alike as the fundamentals are highlighted specifically so that new students will not make basic errors in the future. An interesting fact about this course is that students will learn about recommendation systems and how to apply them to movies.

These are only a few courses that can help those who want to learn about machine learning and other topics linked to it. There are more courses being offered from all over the globe and even some not offered by schools but by companies like Microsoft which actually offers a multitude of courses on specific machine learning topics. Those who want to learn are simply a click away and each course can offer a new way of looking at an idea. Even professors and professionals can learn something new, or at least see it from a different point of view. Given time, machine learning will be incorporated into every single career and machine in the world.

2.3.1 Robotics: Vision Intelligence and Machine Learning

This course, offered by the University of Pennsylvania, actually introduces how robots and machine learning work together. It teaches students about how visual intelligence and machine learning work together to make fantastic robots. Visual intelligence gives robots the ability to have sensations like sight so it can learn and even remember. Machine learning allows robots to simply have object recognition and adapt to new environments. There arent any difficult prerequisites either, just linear algebra, introductory calculus, and a basic knowledge of computer programming.

2.3.2 Machine Learning

Georgia Tech offers a course exclusively on machine learning. Students will learn many different methods and algorithms as the class progresses, like Bayesian learning and reinforcement learning. It will also emphasize theoretical and practical problems along the way. The course promises by the end that students will have a deeper understanding of several major topics in machine learning. This class is a much more advanced class that is recommended to those proficient in machine learning and want to make a career out of it.

2.3.3 Machine Learning Fundamentals

UC San Diego is offering a more basic course to students who want to get into machine learning. The prerequisites are education in linear algebra and multivariate calculus, along with starter courses that are also offered. This course uses actual case studies in order to help students comprehend work. Students will also be taught a number of supervised and unsupervised machine learning

algorithms. The class is interesting in that it also promises to teach students different estimations and methods like conditional probability estimation.

2.3.4 Machine Learning

Columbia University emphasizes the importance of machine learning in data analysis and has made a course to teach those who want to improve their knowledge on machine learning. In order to take the class, you must have taken calculus and linear algebra while also having knowledge on probability and statistical concepts. It also requires coding and comfort with data manipulation which may seem like a lot, but the course is encouraged for those wanting to pursue a career in machine learning. This class focuses on a wide variety of methods in machine learning, from logistic regression to Gaussian mixture models. Probabilistic versus non- probabilistic viewpoints will also be apart of the lessons along the way.

2.3.5 Data Science: Machine learning

Harvard University offers a more basic class than others with a different kind of prerequisites. This course is part of their Professional Certificate Program in Data Science and asks students to take the previous courses before this one. An interesting subject they plan to teach is how to perform cross-validation as to not over train. The course expresses that machine learning is different because it can actually predict using data. Something interesting about the course is that it will explain how to make a recommendation system.

2.4 William Golden

2.4.1 course overview

The course offered through Udacity titled Machine Learning Nanodegree is a quality program for one who is interested in knowing more about machine learning. Especially if one is to teach a course about machine learning, This class will help provide basic knowledge as well as current perspectives and concepts about machine learning wich will help prepare students seeking employment in the computer science fields. The program has several pros as well as some cons such as financing, however, the overall quality of the course would be a great asset.

Prerequisites

Lets begin with the prerequisites of the class. It is recommended on the syllabus of the Machine Learning Nanodegree course that the students have a basic knowledge of the programming software Python. It is required also that students have some prior knowledge of mathematics (statistics, algebra, calculus). If needed, the free course titled Intro to Computer Science is available to anyone who wishes to brush up on these concepts. That course is offered through Udacity as well.

Course Objectives

Next, let's go over a basic overview of the course. We will look at the education objectives as well as the material covered in the course. The purpose of Machine Learning Nanodegree is to teach learners how to become a proficient Machine Learning Engineer. It is set up to give those interested in fields such as finance, healthcare and education, a firm foundation of knowledge and experience in machine learning. Over the duration of the course, students will create machine learning models used in several different fields and learn how to apply them in real-life situations. One example of this is found on the syllabus for Term 2; Students will create an app to detect the breeds of dogs using user-supplied images, and if given a photo of a human, the app will be able to match that person with a resembling breed of dog. At the end of the course, students will leave with the experience of working with numerous types of machine programs as well as having cutting-edge knowledge in the world of machine learning.

Length

We will now take a look at the length of the program. The current class opens on September 18, 2018. The course will also become available again after a few

months of the initial class start. Length time of the class in total is six months. However, it is divided into two terms, making each term last three months. The estimated average of hours spent on the class per week is ten hours. Overall, the commitment time to the course in entirety is roughly two-hundred hours. In order for students to graduate, they must complete seven main projects during the period of these two terms.

Material

The materials required for the course will all be provided online through Udacity. The instructional tools used will consist of a variety of video lectures, personalized project reviews, and mentorship from a professor of the class. Students will be required to do all projects online and will need to be able to access programming applications, so this means having a computer setup that is currently up to date and having a decent connection to the internet.

Instructors

The class is directed by several instructors. All of whom have achieved high degrees such as PhD in there respected fields. The curriculum leader of the project is Luis Serrano who was formerly a Machine Learning Engineer at Google. He and his team of instructors offer personal experiences and mentorship to the students.

Costs

Financially, the course could be looked at as fairly pricey. The cost of the first term (Term 1, Machine Learning Engineer Nanodegree) of the Machine Learning Nanodegree program is a $999.00. The Second term (Term 2, Advanced Machine Learning) is the same price of $999.00. In total, the payment for the course would come down to $1,998.00. The program can be payed in a one-time payment of $999.00 per term, or can be purchased in monthly installments as low as $84.00. This price could be viewed as expensive, however, when considering the value of this education, the costs are not so horrendous.

User Ratings

So that is the basic summary of the course itself. It is important not only to look into what the course entails, but to also listen to the opinions of previous students who had participated in this course. The average rating of the course (found on Udacity's website) is 4.7 out of 5. The number of students who have rated the class are two thousand three hundred forty one. Debasish H. writes a review of the course saying "Very much informative, nice teaching

style, wonderful training materials. Absolutely useful for a new comer where the journey begins with basic statistical concepts along with nice visual and graphical representation and then slowly moves into machine learning with fundamental concepts and then getting into the complex level in each of the section through a consistently upward learning graph." Another user by the name of Rodrigo D. reviewed the class by saying this, "I just complete my machine Learning program, that was exactly what I was looking for. theory and practice in the right portion. This knowledge empowered me with great skills!" Alongside with the positive comments, some students added constructive criticism of the course. One student, Arjun S., said "This is an interesting program. There are certain topics (tough to keep track of which ones exactly) that would be better understood if they are explained in more basic terms." This could possibly mean there are some advanced materials in the course that are not well explained to students.

2.4.2 Conclusion

Overall, with such a high recommendation from past students and simple, detailed descriptions about the course online, I personally found the Machine Learning Nanodegree offered by Udacity to be absolutely impressive. I searched from a vast database of MOOCs for classes on machine learning, and this course by far was the easiest to understand and the most applicable. It may be a bit costly, but as stated previously before, education is invaluable. After taking this course, a teacher would be able to provide there students with fresh, current ideas about the world of Machine Learning. The course would also provide great ideas for projects and techniques in teaching computer learning to beginner students. This is a great program I hope it can provide a lot of assistance.

2.5 Yuan Hong

Machine learning has become a new edge discipline and forms a course in colleges and universities. It combines applied psychology, biology and neurophysiology, as well as mathematics, automation and computer science to form the basis of machine learning theory. Academic activities related to machine learning are unprecedentedly active. In addition to the annual machine learning seminars, there are computer learning theory conferences and genetic algorithm conferences. So it is really wise to just take this machine learning class to follow up the steps of this new era.

For machine learning, I firstly recommend the Machine Learning Course in Columbia University. This course is now very popular because nowadays the technology allows to widen the careers in data analysis. This course includes the skills and solutions, such as classification and regression, clustering methods, sequential model, matrix factorization, topic modeling and model selection.By using these methods, you can use these skills to apply them to the real problems, such as different topics in trends such as building recommendation engine, ranking sports teams and also plotting path for some of the models. While studying the first half this course, we will be learning techniques, for regression and classification. Students will predict a set of inputs with the possession of a single output. By predicting it, we will use the algorithms to perform this fundamental method. Once the approach is motivated, we will abstractly learn about different algorithms by using mathematical learning methods. Then, in the second half the course, the learning goal will shift to unsupervised learning techniques, which is a learning algorithm of artificial intelligence that is neither classified or labeled. With the unsupervised learning we will cover three fundamental learning problems of it, these will be data clustering, matrix factorization, and sequential models for order-dependent data. By doing object recommendation and topic modeling, we can apply this learning method to the real world. You could learning this material by watching videos online, and the language is all English.

The reason why I recommend this course is because this course is completely free to take it, and an extra $199 for a verified certificate to it. And it is eight to ten hours per week average for twelve weeks. And from the information above, we know that we are learning something that could apply to different studying fields and careers.

The next course I would recommend is Data Science of Machine Learning from Harvard University. Data science is tightly connected to the machine learning. Some of the most popular products using machine learning include handwritten readers implemented by postal services, speech recognition, movie recommendation systems and spam detectors. During the course, students are part of the Certificate in Data Science program and you will learn popular machine learning algorithms, principal component analysis and normalization by building a film recommendation system. As time goes on, students will learn how to train data,

a set of data to discover potential predictive relationships, and how the data appears in the form of the results we want to predict, and the features we'll use to predict this outcome. When building a movie recommendation system, student will learn how to use training data training algorithms so that you can predict the outcome of future data sets. Students will also learn about techniques for overtraining and avoiding cross-validation. All of these skills are the foundation of machine learning.

And the class is free to take, and certificate is 49 dollars. The class is extremely short and it wont bother your other activities with two to four hours per week in four weeks. I would recommend this because it is a fun class in a short period.

As we all know, machine learning run predictive models on computers and learn from existing data to predict future behaviors, outcomes, and trends. To learn more about machine learning theories, I would recommend this Principles of Machine Learning: Python Edition, which the institution is Microsoft. In this data science course, students will gain a clear explanation of machine learning theory, build, validate and deploy machine learning models based on actual scenarios and practical experience. Students will learn how to build these models and gain insights from Python and Azure notebooks. By the way, in this course, the sponsor edX provides financial assistance to students who want a certification certificate but may not be able to pay the fee. The reason that I would recommend this class is that this class focus on a different perspective of machine learning by researching the theories and studying python language. There are many useful skills, after studying the whole course, students will learn data exploration, preparation and cleaning, supervised machine learning techniques, unsupervised machine learning techniques, model performance improvement and etc. The course is six weeks long and is six to eight hours per week, so this is not a long course. The certificate for this class is 99 dollars, with the new topics and shortened time of studying, and the famous institution Microsoft, this class is really worth to take.

The next course I will recommend is Robotics: Vision Intelligence and Machine Learning. This course has a fascinating field of visual intelligence and machine learning. Robots see and respond to and learn from their interactions with the world around them by using the visual intelligence allows robots to "perceive" and "recognize" the surrounding environment. It also enables the robot to "learn" from the memories of past experiences by extracting patterns in the visual signal. Students will learn how machine learning extracts statistically significant patterns from data that supports classification, regression, and clustering. Then by working together on computer vision and machine learning, then students will be able to build recognition algorithms that can learn from the data and adapt to the new environment. At the end of the course, as part of the Robotics MicroMasters program, students will be able to program the robot's visual functions, such as robot localization and object recognition using machine learning. The projects in this course will use MATLAB and OpenCV and will include video stabilization, 3D object recognition, coding object classi-

fiers, building perceptrons, and real-world examples of designing convolutional neural networks CNNs using one of the standard CNN frameworks. By learning this course interacting with robots and different types of working environment is really helpful to students, when trying out different things, students will think of the problem in a different perspective. I recommend this course only because that's the machine learning course tight to robotics. Although the course is long, with eight to ten hours per week in twelve weeks, and the certificate is expensive, I would still encourage you to go to try this course.

With the maturity of the theory of machine learning, more and more related technologies have been applied to all aspects of life practice, and the work of machine learning has become a key consideration for many computer industry practitioners. So go to study machine learning courses online, you can study it at home, and you can pause if you dont understand some parts. Although some people are pros at programming, but lacks a lot of theoretical knowledge about machine learning is very difficult, and even the simple formulas involved in the logical regression of machine learning are a bit difficult. So taking these courses before running into the problem, it is definitely a wise choice.

2.6 Easton Jensen

2.6.1 Overview

"Machine Learning Specialization", a class from University of Washington on coursera, is for a specilization inmachine learning. This requires some experience in the field. There are 4 courses in this specilization class. The courses have a suggested order, but you can do them in whatever your order you want. After that they give you a project that is designed for you to apply and pactice the things you have just learned. At the end of the class you will gain applied experience in Machine Learning and its major fields. Those include Prediction, Clustering, Classification, and Information Retrieval. Learning to analyze complex datasets, creating systems that adapt, and build intelligent applications to make predictions from data. You also learn Python programming experience as you you use Python throughout all the courses.

2.6.2 Course 1

Machine Learning Foundations: A Case Study Approach, is the first suggested course. This course is six weeks of study with five to eight hours a week spent on the course. In the course you get hands on experience with several case-studies. The description used for the first course is a black box. By delving into the box you can form together a machine learning pipeline, which can be used to develop intelligent and practical applications. At the end of the course you should be able to Identify possible machine learning applications, describe the differences of analyses by using regression, classification, and clustering. Able to task a machine to a certian application with machine learning. use regression, classification, clustering, retrieval, recommender sytems, and deep learning in your field. Take data at turn it into a form you are able to input it into a machine learning model. Take a model and analyze and assess its quality based on relevent error metrics for each task. Procure a dataset that fits a model so a computer may analyze the data. Building a final application that has and uses machine learning at its core. Throughout the entire course you will be using Python therefore you will take all these techniques and being able to use them in Python.

2.6.3 Course 2

Course 2 is all about Regression. This course is also six weeks of study, and five to eight hours a week. The case study used in this course is predicting housing prices. You will be able to use this case study in a way to predict the price of a house based on its features such as, square footage, number of bedrooms and bathrooms and so on. There are other ways to use regression however not just

for house prices. they are used in predicting many things from health outcomes to stock prices in finance. This course takes linear regression models and large sets of data to help you learn how to analyze and sift through these datasets with optimization algorithms. Taking this course will give you the knowledge to take input and output of regression models and identify or describe it. Compare the good and bad of bias and varience when taking data. evaluate the parameters of a model using a optimization algorithm. Cross validate parameters. Assess performance of a given model. Learn parts of sparsity and how LASSO leads to sparse solutions. Learn methods to differ to different models. Take a model and make predictions. Take a housing dataset and create a regrssion model to predict price. Python is still used, so learning how to do all these things in Python will be another skill aquired by the student.

2.6.4 Course 3

The specilization in course three is about classification. This course is seven weeks long and is also five to eight hours a week. The case studies used for the third part is Analyzing Sentiment & Loan Default Prediction. the case about Analyzing Sentiment will help students create models that will predict a negative/positive sentiment from input features. These imputs will include text of the review to user profile information. In the study about Loan Default Prediction, students will take on financial data that will help predict if the loan would be risky or safe for the bank in question. This takes into acoount a big and widely used area of machine learning called classification that includes ad targeting, spam detection, medical diagnosis, and image classification. Being able to create classifiers that perform in the way intended is the goal of this course. Delving deeper than that, you will be able to create the algoriths that can learn from these models at scale. You will be able to use these techniques on real world machine learning tasks. At the end of this you should be able to display an input and output of a classification model and describe them both. You will be able to take down classification problems of both the binary and multiclass variety. Being able to use logistic regression models on large-scale classification problems. You will also be able to make non-linear models while looking at decision trees. Using boosting to imporve any models productivity. change your models scale using stochastic gradient ascent. Make predictions of sentiment in a product review dataset with a classification model you build yourself. Take data to predict loan defaults. Being able to work around and handle missing data. Even evaluate models created by using precision-recall metrics. Then by the end all these skills will be easily used in python.

2.6.5 Course 4

Course 4 is a unit about clustering and retrieval. The commitment a student has to put into the course is six weeks of study. During those weeks students

are asked to spend five to eight hours a week on the class. The case study used in the fourth course is Finding Similar Documents. This course takes students to find material that is similar to something you are looking up, like a news article. They find these documents and pile them together. Then as people go to look up things that are similar it shows them these other articles as well as what they are looking for. In this case study students will learn similarity-based algorithms for retrieval. They will learn about clustering an mixed memborship models, like LDA, better known as latent Dirichlet allocation. At the end of the course students will be able to, retrieve documents using the system k-nearest neighbors. Find metrics of various similarity for text data. rduce computation times in k-nearest neighbor by using KD-trees. Use locality sensitive hashing to generate an approximation of nearest neighbor. take learning tasks of supervised and unsupervised works to compare and contrst them. Clump together documents and files by topic using k-means. learn how to parrallelize k-means by the means of a program called MapReduce. Examineing all probabilistic clustering methods by using mixtures models. Using EM, expectation maximization, to fit a mixture of Gaussian model. Use latent Dirichlet allocation(LDA) to perform mixed membership modeling. Learn the steps to Gibbs sampler and how to draw inferences and make decisions with the output of the sampler. take diffirent initialization techniques for non-convex optimization objectives and see the pros and cons of the different techniques. During this project you will be using Pyhton. Therefore these techniques will become familiar with you in the Python language.

2.6.6 Why Take The Course

I believe you should teach some of the things in this class because it is a specialization course which teaches you different techniques for the same problem. Since there are so many different programs the same problem learning as many as you can would be good for a student going out into the world. These courses would be for a older student as you need a little bit of prior knowledge to complete these courses. Plus the courses are 6 weeks long and only 5 to 8 hours long a week. Fitting it into 18 days would take some thinking based on what exactly you want to teach in the course but I have no doubts that taking parts of this course will make a great course for aspiring computer scientists.

2.7 Rodrigo Martinez

Machine learning has been an important aspect of the advancement of computer science. Learning how machine learning works is also a very important factor to do which is why courses exist. Several companies like Microsoft and Google create opportunities for anyone to learn about machine learning. A course I recommend is called "Introduction to programming." This course this a 3 month course that you can take in order to learn the basics of programming. This course will include languages like, HTML, CSS, Javascript, and Python. These programming languages will be taught at a basic level so there isnt a need of having background knowledge about programming. The only downside to this is that the cost of the course is $599.

Another course that I recommend is created by Google. The course is called, "How Google does Machine Learning." This course does not specify on the amount of money needed however it will require 8-10 hours a week! This course is separated into 5 individual parts in which each has its own unique focus. However for this one we are focusing part 1 out of 5. This course will offer you the explanation and concept of machine learning. Theres a kick to it, you be taught Googles interpretation on machine learning. To them it is, "about logic more than just data." You will learn something called the "5 phases" of machine learning. Requirements for this course will be to have a desktop web browser that is able to run the websites interactive labs via Qwiklabs. You will be provided a syllabus to get the idea of what the course will expect and the atmosphere of the course.Google recommends people who are interested in programming, data scientist, data engineers, etc should take this course.

The next course I recommend is called "Become a machine learning Engineer." This course will take a three month term and will cost $999. This course will also offer you knowledge about machine learning however, it will be taught in a self dependent manner. There will be 2 terms for this course. Term 1 will focus on the introduction of machine learning like the basic fundamentals. As you proceed on to term 2, you will be left self dependent and unsupervised in some cases. You will also learn topics such as deep learning and reinforcement learning in term 2. A project will be expected by the end of the course demonstrating that you understood how machine learning works.

Lastly the final course I recommend is called "Deep learning." This course is free to take and will take require around 3 months to complete. It will offer you the explanation of Deep learning. You will also be taught about Deep neural networks, convolutional networks and recurrent networks. You will need at least two years of programming experience to take this course and also have background knowledge about Git and Github. This course will provide you with 4 different unique topics each having its own purpose to machine learning.

- Udacity's Introduction to Programming Nanodegree

- Coursera's Google Machine Learning course
- Udacity's Machine Learning Engineer Nanodegree
- Udacity's Deep Learning course

2.8 Matt Morrical

The field of Mechanical Learning as a subset of Computer Science is still in its infancy and thus has a few things that can help people to learn about it. I would suggest taking a course or two online and watching a few lectures.

There a few courses that I would suggest taking from MOOCs, massive open online courses. The order that I would suggest is, EDX's Data Science: Machine Learning. The price of this course is $49 if you want a Verified Certificate otherwise, it's just the two to four hours per week for the four weeks the class runs for. The credentials of this class are quite strong, it is taught by a Harvard professor and that EDx backs it, which is no small thing.

The next course I would suggest is Udacity's Intro to Machine Learning, located at https://www.udacity.com/course/intro-to-machine- learning–ud120. As for the price, it doesn't cost a cent outside of the about ten weeks that it lasts. The credentials of this class are even better than that of the previous, but it is as a more indepth class, being labeled an intermediate class but Udacity. The instructors are industry leaders, Katie Malone and Sebastian Thrun, and the course is self paced.

Overall for these course, you can expect 14 weeks of study, but if you are able to put in more work, you can make them go even faster, not to mention that they cost nothing, unless you want a certificate only availible from one course and costing only $49!

2.9 Ella Nelson

Here are my recommendations.

2.10 Koichi Okazaki

Beside searching for the course model, I considered these four points are crucial for the ideal course model.

1. Enough resources are offered in the program.
    - Having enough resources or not determines the quality and efficiency of the learning.
2. Course which enhances ones from basic to advanced knowledge of machine learning.
    - Starting from basic skill, it is ideal to expand ones knowledge to advanced level gradually.
3. Enough opportunities to reflect the learning are offered.
    - Taking quizzes, tests, or conducting project reinforces ones learning and also make realize what is the difficult points in that field of study.
4. Interesting.
    - To build an interesting class the course that I going to introduce has to be interesting too.

Based on these four factors, I concluded on that the Nanodegree program offered from Udacity fulfills these criteria the most.

2.10.1 Deep Learning Nanodegree Program from Udacity

In this Deep Learning Nanodegree program, you will develop the basic understanding of machine learning. At the same time , you are also able to study the most advanced fields such as, Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, Generative Adversarial Networks, and Network Deployment. Additionally, you have enough opportunities to work on projects, in this program five projects are offered to you, and you could verify your knowledge which you learned in the lessons. Knowledge about these fields and the learning method might be useful for building your class.

This program is four months long and you are expected to spend about 12-15 hours per week to complete both lecture and program work, and lecture will be conducted by Ian Goodfellow and Jun-YanZhu, originators of types of generative adversarial networks, and also AI professional, Sebastian Thrun and Andrew Trask. Needless to say, all of them are professionals in their fields. The course costs $83.25 per month and this is not a cheap price (writing studio).

Therefore, I can say the courses are ideal program for the person who wants to learn about deep learning.

After all is said and done, after graduating this program you are able to participate in a Guaranteed Admission Program. There are two programs: Self-Driving Car Engineer and Flying Cars, you can choose one of the programs that you would like to learn more.

Course Description

During the program you are going to be engaged in five projects, and for each of the projects there are several supporting lessons which you will take in order to bring the project to success.

Prerequisite Knowledge: In order to succeed in this program, students are recommend to have mid-level knowledge of python and basic knowledge ofprobability. Therefore, you might teach about Deep learning after you make your students are familiar with those two fields.

Through this course AI with Python Nanodegree program you could learn about those skills. (https://www.udacity.com/course/ai-programming-python-nanodegree–nd089)

Projects: Testing your skills and knowledge is very important, because you can prove and reinforce your knowledge by making mistakes, and building the project is one of the most effective way to do so. Student will be engaged in five projects though this program:

- Predicting Bike-Sharing Patterns
- Dog Breed Classifier
- Generate TV Scripts
- Generate Faces
- Deploy a Sentiment Analysis Model

Project 1: Predicting Bike-Sharing Patterns This project introduces you the fundamental skill of neural networks by using Python and NumPy. You will build a multi-layer neural network from the very first and analyze the number of bike-share users on a given day.

Supporting Lesson

- Introduction to Neural Networks
- Implementing Gradient Descent

- Training Neural Networks
- Sentiment Analysis
- Deep Learning With Pytorch

Project 2: Dog Breed Classifier This project provides you an opportunity to learn how to build Convolutional Neural Network (one kind of neural network which is mostly used for analyzing visual images.) and you will build your own algorithm which identify dogs breed.

Supporting lessons

- Cloud Computing
- Convolutional Neural Network
- CNNs IN Pytorch
- Weight Initialization
- Autoencoders
- Transfer learning in Pytorch
- Deep Learning for cancer detection

Project 3: Generate TV Script By using PyTorch you will produce your own Recurrent Networks and Long Short-Term Memory Networks. Additionally, you will do sentiment analysis and create new text, and in order to create a new text which is like a training set of TV script, you will use recurrent networks.

Support Lessons

- Recurrent Neural Networks
- Long Short-Term Memory Network
- Implementation of RNN & LSTM
- Hyperparameters
- Embeddings & Word2vec
- Sentiment Prediction RNN

Project 4: Generate Faces Through this project you will learn GAN (Generative Adversarial Networks) with Lan Goodfellow (model's inventor). And eventually you will invoke a Deep Convolutional GAN, this DCGAN is made of a pair of multi-layer neural networks which are called generator and discriminator that compete against each other until generator learns to generate realisticimages of faces.

Supporting curses

- Generate Adversarial Network
- Deep Convolutional GAN
- PIXPIX CycleGAN

Project 5: Deploy a Sentiment Analysis Model This project provides you an opportunity to train and deploy your PyTorch sentiment analysis model for analyzing movie reviews, by using Amazon SageMaker fully managed machine learning service. With Amazon SageMaker, data scientists and developers can quickly and easily build and train machine learning models, and then directly deploy them into a production-ready hosted environment on AWS. (Amazon Web Sevices)

Supporting Courses

- Introduction to deployment
- Deploy a model
- Custom models & Web hosting
- Model monitoring
- Updating a model

2.10.2 At the end

For sure, it's not cheap price which this program costs to you. However, I think this price is based on the high quality and produces self-confidence and I also think taking this project won't be a waste of money. Thus, in order to build a fascinating class, I think you should better learn from the experts of that fields, since they could offer clear and sophisticated knowledge and way of teaching about this field. This program provides you the opportunities to do so. Therefore, I will recommend this program to you.

Cited from "What is Amazon Sagemaker?." Amazon Web Services, 2018.

2.11 Jakob Orel

There are many options to consider when studying machine learning (ML). There are courses from universities, online courses, and foundations of education that may allow you to learn more about this topic. I would suggest taking an online course from Coursera titled Machine Learning Foundations: A Case Study Approach or taking an EdX program recommended by the Association of Computing Machinery titled Learning from Data from the California Institute of Technology. Extra knowledge on the material can be gained from attending a workshop by the Argonne National Laboratory.

2.11.1 Coursera

Machine Learning Foundations: A Case Study Approach is a course from the University of Washington that offers introductory yet in-depth material to students. The online course offers hands-on experience with machine learning by evaluating practical scenarios. These hands-on experience will allow you to apply machine learning to real-world situations such as advertising or predicting user suggestions. Learning in this course is interactive with videos, quizzes, and projects. You can also interact with thousands of other learners about the course material. This first course will allow you to focus on tasks of interest and use machine learning to assess a quality output. By the end of the course you should learn how to:

- Identify potential applications of machine learning in practice.
- Describe the core differences in analyses enabled by regression, classification, and clustering.
- Select the appropriate machine learning task for a potential application.
- Apply regression, classification, clustering, retrieval, recommender systems, and deep learning.
- Represent your data as features to serve as input to machine learning models.
- Assess the model quality in terms of relevant error metrics for each task.
- Utilize a dataset to fit a model to analyze new data.
- Build an end-to-end application that uses machine learning at its core.
- Implement these techniques in Python.

This Coursera course is the first installment of four of the machine learning specialization. It will take six weeks with five to eight hours of work per week.

The price of the course is unclear on the Coursera website, but a membership is required. This University of Washington course is taught by Carlos Guestrin, Amazon Professor of Machine Learning in the computer science and engineering department, and Emily Fox, assistant professor and Amazon Professor of Machine Learning in the statistics department. Carlos was a co-founder and CEO of Turi (originally GraphLab Inc.), focusing large-scale machine learning and graph analytics, which was acquired by Apple. I would recommend following an outline of this class to teach other students because it provides a clear view of how machine learning can be applied to the world.

2.11.2 EdX

Learning From Data is an online course from the California Institute of Technology that is also available on EdX.com. It was recommended by the Association of Computing Machinery as a way to learn more machine learning. This course covers the basic theory, algorithms, and applications of ML. This course offers 18 lectures covering the mathematical as well as the heuristic aspects. There is also a discussion forum and video library to access. I believe one planning to instruct a course on machine learning should master the fundamentals found in this course.

The course is taught by Caltech Professor Yaser Abu-Mostafa. Dr. Abu-Mostafa is a highly qualified Professor of Electrical Engineering and Computer Science. He is also Chairman of Machine Learning Consultants LLC and has written and published the textbook for the course on Amazon which is a bestseller. Dr. Abu-Mostafa received the Clauser Prize for his doctoral thesis and is a founding member of the IEEE Neural Network Council. He has also published several articles in Scientific American. This is a great instructor to learn the foundations of machine learning.

This online course requires some complex math knowledge. Eight homework sets and a final exam are spread throughout ten weeks. EdX suggests that 10 to 20 hours per week will be required. The textbook that correlates to 14 out of the 18 lectures costs twenty eight dollars on Amazon currently. Familiarity with quadratic programming and basic matrix operation skills are needed. I believe this course lays the groundwork for anyone interested in learning about machine learning. An instructor should have full mastery of these concepts to teach a class over ML.

2.11.3 Argonne National Laboratory

The Argonne National Laboratory is a government funded lab outside of Chicago that has done extensive research on machine learning. Their projects range from algorithms and software development to applications in science and the environment. Examples of projects include novel algorithms for Bayesian and

blackbox optimization, scalable frameworks for neural network hyperparameter optimization and tuning, predictive modeling of wide area data transfer, creation of a lightweight thermal prediction system for runtime management, and many more. Accessing published papers from Argonne National Lab that have been in peer-reviewed journals and presented at conferences would further one's experience in machine learning.

I would recommend learning from this resource because they have conducted large projects on this topic. In January of 2018, the Argonne National Lab held a Deep Learning Workshop to dive into the use of deep learning with the goal of raising competency. This workshop included speeches from senior scientists of the program Prasanna Balaprakash, Nicola Ferrier, and Justin Wozniak. This type of workshop would further ones understanding of certain topics and would be a great experience to meet scientists conducting research in the machine learning field. The uses of machine learning in the lab may be able to be incorporated into a course.

2.11.4 Takeaway

These online courses and options for learning are far from the only options to explore machine learning. There are hundreds of courses and thousands of sources of material to gain more knowledge on this broad topic. There are many courses, but I would suggest taking the course Machine Learning Foundations: A Case Study Approach from the University of Washington and modeling a similar outline of work for the course. I also recommend taking the very popular Learning From Data course from the Caltech Professor Yaser Abu-Mostafa. Learning from this course will further ones knowledge on the complex topic of machine learning. I would also advise attending a workshop at the Argonne National Laboratory. This workshop may give someone more information to add to the course that is currently being researched at a government lab. I believe these three ways of learning would provide enough material to create a new course on machine learning.

2.12 Marcellus Parks

Marcellus Parks CSC131 Recommendations

Before teaching this class students need to have taken some prerequisite courses including but not limited to: Calculus, Linear Algebra, Statistics, and one or two programming courses. They should know how use programing languages such as java, C Language, or python. This will allow students to know what is going on in your class. Without a strong foundation in math, their algorithms would not work very well. Without some programming and data manipulation knowledge, they would not be able to get much coding work done. With all of these prerequisites taken, students are able to hit the ground running.

When teaching a class on machine learning, it would be best not to dive right into it. Similar to the first few days of the class in the University of Washingtons Machine Learning Foundations course, one of the best ways to start this course it to teach students the things they could do with machine learning. Be the cause of your students imagining what the world could do in the future with machine learning. Invite students to share what they would want to do with the knowledge they gain from the class. Discuss how machine learning is in almost everything we use in life in this day in age.

Once students have a better idea of what machine learning is for and what they can do with it, they need to be taught the fundamentals. Understanding of the fundamental concepts of machine learning will allow students to be able to create machine learning programs that can be applied to many different fields of study, not just the ones they are interested in. In UC San Diegos Machine Learning Fundamentals, students are taught supervised and unsupervised learning algorithms and the theories behind them. They are shown real world case studies. A good course of action would be to emulate this course in some amount so students can see the logic behind the coding. It would also help them know how they would have to think when writing their lines of code. At this point it would be better to focus on the logic behind it than the coding itself. Without sound logic and arithmetic, the coding would be useless. This point in the class would be a good time to differentiate between the different types of machine learning algorithms. For instance, supervised learning is where a model is provided a training data set. The more the in the training set, the better the performance on unknown examples. Students should also be taught about not commonly used algorithms such as those with active learning. Active learning is when the model improves by receiving feedback from a human. Students should know the difference between these algorithms because each application of machine learning has a most efficient way to learn.

After learning what fundamentals to use in machine learning there is still a mathematical aspect. The recommended textbook for this class is The Elements of Statistical Learning by Trevor Hastie, Robert Tibshirani and Jerome Friedman. The book attempts to bring together the important new ideas in

statistical learning. Even though the book uses many mathematical models, the textbook puts emphasis the core concepts of learning more than the theoretical properties. The authors themselves want the book to be used by not just statisticians but researchers and practitioners in a multitude of other fields. The book would be a great tool to teach students the mathematical and theoretical properties of statistical learning while focusing its main concepts.

Learning the concepts are more than likely the most difficult part of machine learning. However, once they are learned they can be put to use. It is recommended that you imitate Deeplearning.ais Structuring Machine Learning Projects. The course teaches many of the abilities needed to have a successful machine learning project. Understanding how to diagnose errors is an absolute must when creating a successful model. Being able to prioritize directions in reducing error is also a significant skill. Understanding complex machine learning training sets such as comparing and surpassing human-level performance is another must have ability. The course also teaches how to apply end-to-end learning, transfer learning and multi task learning, which are very useful skills. Be sure to include other skills than the ones listed in just one class, such as feature scaling which is preprocessing data to find where someone could see where they can improve their algorithms creating a better and more efficient product.

Another good textbook to use is Introduction to Machine Learning by Alex Smola and S.V.N. Vishwanathan. The textbook talks about how Machine learning has become an integral part of the forward trend of software. It is similar in a few ways to the textbook specified earlier, but it has more of a focus on the applications of machine learning even when giving mathematical instruction. One of the main reasons the book is recommended is because the book is very well structured. It methodically connects nearly all of the skills and methods that should be taught in a machine learning class. It covers the mathematical skills better than the first book as well.

It is also recommended that you imitate Udacitys Intro to Machine Learning. Not necessarily in content, but definitely in style of teaching. The course encourages students to try and fail. Learning by doing is one of the best ways for a student to retain the knowledge you plan on giving.

A class on machine learning should cover all of these topics and probably more. Students that enter the class as a novice should come out proficient and efficient at creating machine learning models. The class should not start fast but should soon accelerate into the the topics discussed. Students will at least come out of class satisfied because they have knowledge on machine learning concepts and algorithms.

2.13 Lydia Sanchez

Program: Machine Learning Fundamentals

Sponsor: UC San Diego(edx)

Length: 10 week course

Pace: 8 to 10 hours a week

Price: Free

Certificate: yes is $350 extra

Prerequisites: previous courses in micromasters program; DSE200x and DSE210x Experience in the undergraduate level with multivariate calculus and linear algebra

Overview: the class will go other these topics,Classification, regression, and conditional probability estimation,Generative and discriminative models, Linear models and extensions to nonlinearity using kernel methods,and Representation learning: clustering, dimensionality reduction, autoencoders, deep nets. Recommendation: This looks like a good class on machine learning, probably not one to start on because it does have some prerequisites courses. It is part of the micromasters program in data science but the micromaster is $1260 but it seems the classes individually are free.

Teacher: Sanjoy Dasgupta

Learn more here.

Program: Machine Learning

Sponsor: MIT opencourseware

Length: At your own pace

Pace: At your own pace

Price: Free

Certificate: No

Overview: Introductory class on machine learning with topics on classification and linear regression, boosting support vectors machine, hidden Markov models, bayesian networks. Basic idea of of machine methods as well as a little bit of how, why and when they work.

Recommendation: The course is seems good especially since it is from MIT but one downfall is that the class is old, it is taught as it was in 2006. Another downfall is that there does not seem to have any feedback on test but you do get the class as it was taught to MIT students. All round it looks like a good class.

Teacher: Prof. Tommi Jaakkola

Learn more here.

Program: Creative Application of deep learning with TensorFlow

Sponsor: Kadenze

Length: 5 week session

Pace: 12 hours of work per session

Price: Free

Certificate: yes (included with premium membership $20/month)

Prerequisites: Be familiar with Python, take the first CADL program,

Overview: introduction into deep learning(what it means and how it works) and building AI algorithms such as deep convolutional networks.will learn to how to build codes but how to use them for creative applications.

Teacher: Geoffrey Hinton

Review: The instructor seems very active on the forums and even set up a slack for the course. It's been great, and the homework and notebooks are really easy to follow. So far it has really made me think and seems a lot more engaging than the Udacity or Coursera course. Can't wait to see where it goes!-Anonymous via Class central(This was a five star review)

Recommendation: This has class has very good reviews as seen above, seems like a good intro to machine learning according to the reviews. You get one on one time with instructor in forums, which is good if you have questions. People in the review say it is better than coursera and udacity. The reviews said it was a hard course but they said it was good class.

Learn more here.

Program: Neural Networks for Machine Learning

Sponsor: University of Toronto(Coursera)

Length: 16 weeks

Pace: 7 to 9 hours a week

Price: $ (free if you audit lectures but you can't do it for the whole course)

Certificate: yes

Prerequisites: the course is for an intermediate level

Overview: Learn about how artificial neural network is used for machine learning. Machine learning used with speech and object recognition, image segmentation,modeling language and human motion.

Teacher: Geoffrey Hinton

Recommendation: The class seems like it has some good content but the seems too long at 16 weeks. You can talk to your peers, which is a benefit because its like your taking a class in person. The reviews in the class are good.

Review: Geoff Hinton is one of the founding fathers of neural network when everyone jumped ships in the 90s.This course takes a more theoretical and math-heavy approach than Andrew Ng's Coursera course.If you are interested in the mechanisms of neural network and computer science theories in general,you should take this! An intellectually invigorating experience.-Dolly Y.,via Class central (this was a five star review)

Learn more here.

Program: Deep Learning by Google

Sponsor: Google (udacity)

Length: 12 weeks long

Pace: self paced

Price: Free

Certificate: Not for this class(this is class is part of nanodegree you can get)

Prerequisites: No

Overview: It is an intro and overview to deep learning and how it works with machine, you will learn deep neural networks and advanced architecture like convolution networks. You can go more in depth into deep learning with the nanodegree.

Recommendation: The class has a lot of good attributes, such as it as it is self paced and it is free but on the website class central there are a lot of bad reviews such as they did not like that it was intermediate class(like needing an intermediate knowledge of python programming). This class would probably should not be the first class you take in machine learning.

Teacher: N/A

Learn more here.

2.14 Tiff Serra-Pichardo

Carnegie Mellon University
10-601 Introduction to Machine Learning (master's)

Students in this class have the most diverse academic background. The course is intended for students from math backgrounds to be able to easily catch up. This course is mathematically rigorous and contains both programming and derivations in its homework. The goal of the course is to introduce the theory and practice of machine learning. Students will be able to select and apply appropiate machine learning algorithms for a given learning problem, modify existing learning algorithms to novel situations, and implement the modified algorithms, read and understand research papers about machine learning algorithms.

Coursera (Stanford University)
Machine Learning

The program covers a large number of techniques and algorthms throughout an estimated 11 weeks. Two of the weeks are spent covering neural networks and deep learning. The course is taught by Andrew Ng, he was the former head of Baidu AI/Google Brain and is a current adjunct professor. He has been noted as an excellent instructor. The assignments can be completed using MATLAB or Octave.

edX (Columbia University)
Machine Learning

This course is for the more advanced, it covers more algorithms than the Stanford course. As long as the prerequisites are completed there should not be a problem with the course difficulty. Students can use Python, Octave, or MATLAB to complete assignments. The course is free with a verified certificate available for purchase.

Udemy
Machine Learning A–Z

The incredibly detailed course includes instruction in both Python and R. The course covers all of machine learning workflow and a much larger array of algorithms compared to the previously courses. The course is more applied in its approach and much lighter on the math. This course has less prerequisites and the instructors are known for their ability to make the complex seem so simple.

2.15 Austin Stala

2.15.1 Machine Learning Background

There are many complex and in many places confusing concepts behind machine learning. Luckley there are many online courses, programs and lectures to help provide you background on machine learning for you to teach a course in it.

I recommend the following courses and programs to help you learn enough about machine learning to teach a course in it.

Machine Learning

Cost Free

Prereq None

Provider Stanford University via Coursera

Length 11 weeks

Short Description This course is a broad intro to machine learning, data mining and statistical pattern recognition.

URL https://bit.ly/1IXp8Lg

Machine Learning Fundamentals

Cost Free

Prereq MicroMasters program: DSE200x (https://bit.ly/2Qnx9oS 10 weeks, free, 8 to 10 hours per week) and DSE210x (https://bit.ly/2N3fAMy 10 weeks, free, 10 to 12 hours per week)

Provider UC San Diego via edx

Length 10 weeks (8 to 10 hours per week)

Short Description Learn the basics of machine learning and how to create several types of models for machine learning.

URL https://bit.ly/2wQVl9V

Practical Machine Learning

Cost $49

Prereq None Listed

Provider Johns Hopkins University via coursera

Length 4 weeks

Short Description This course will teach the basics of creating and applying prediction functions with a focus on practical applications for machine learning.

URL https://bit.ly/2NoOe2U

Neural Networks for Machine Learning

Cost Free

Prereq None Listed

Provider University of Toronto via coursera

Length 16 weeks

Short Description This course teaches about artificial neural networks and how they're being used for machine learning.

URL https://bit.ly/2QbE4RR

Machine Learning for Musicians and Artists

Cost Free

Prereq None

Provider Goldsmiths, University of London via Kadenze

Length 7 sessions (8 hours per session)

Short Description This course will teach the basics of machine learning and how to create programs that can sense human gestures, musical audio and other real time data.

URL https://bit.ly/2KaWeTz

Artificial Intelligence (AI)

Cost Free

Prereq Python Language experience

Provider Columbia University via edx

Length 12 weeks (8 to 10 hours per week)

Short Description Learn the history and applications of AI as well as build some basic search AI.

URL https://bit.ly/2Ct4Mmh

Deep Learning

Cost Free

Prereq None

Provider Google via Udacity

Length 3 months

Short Description Learn the basics of machine learning and its applications.

URL https://bit.ly/1WB1zAk

Machine Learning

Cost Free

Prereq None

Provider Georgia Tech via Udacity

Length 4 months

Short Description Advanced level courses on machine learning.

URL https://bit.ly/1LMGjRJ

Based on my research of these difference causes I believe that the courses listed above should provide a good understanding of machine learning. Overall these courses will provide information on things starting at the basics to some applications of machine learning. Some of these courses also provide practical learning and hands on experience creating and working with machine learning.

2.15.2 Lesson Resources

I believe that it would be best to start out a lesson on machine learning with a short and simple video lecture explaining what machine learning is. It is likely that most of the class would have some familiarity with machine learning or at least it is likely that they have at least heard the term or a related term before. The video below provided by The University of Oxford is a good and short (just over two minutes) general introduction to machine learning.

https://youtu.be/f_uwKZIAeM0

There are also some shorter online course that can be used in the lesson to help your students learn more practical applications for machine learning. Some of these lessons can also help the student by providing hands on experience with machine learning. I believe that hands on learning can greatly assist students in understanding how machine learning works and what it can do.

Machine Learning

Cost Free

Prereq Some Python Language experience

Provider Google via Google

Length 15 hours (25 sessions 40 exercises)

Short Description This is a crash course on machine learning with video lectures, real world cases and hands on experience.

URL https://developers.google.com/machine-learning/crash-course/

I believe that the course listed above can be a great resource for the first few weeks of class. There are several reasons that I believe this. The first reason its short duration only being 15 hours of work. Second it provides basic and advanced knowledge. Third this course provides hands on experience for students.

2.15.3 Conclusion

In conclusion there are many courses that can help in learning machine learning and deep learning. These courses also teach in many ways to help further and retain knowledge of machine learning. Some of the ways that these courses can be very useful. One of these is ways is hands on learning creating machine learning models and programs. There are also real work applications and studies

used in the courses which can help further understanding of machine learning and its applications. The courses also use videos to help guide learning.

The providers and levels of the courses can also help in learning about machine learning and deep learning. There are many providers of the courses listed above. Some of these include google, Georgia Tech, University of London, Columbia university, Stanford University, The University of Toronto and several others. These courses are also at different levels and cover many different topics. These courses levels range from beginner and basics to advanced and graduate level. Some of topics that these courses range from the real world applications of machine learning, the fundamentals of machine learning and how to create machine learning or deep learning models and programs.

Finally there are many ways to learn about machine learning. These range from the several courses listed above and many many more not listed. There are also many video lectures that can help with learning about machine learning. These are also many resources that can be used to create the class on machine learning. One of these resources is the short google course listed above. There are also several short videos that can be used to help teaching. One of these short crash course videos is also listed above.

2.16 Nicole Trenholm

- Machine Learning by Stanford University professor and co-founder of courser Andrew Ng. Offered through the website coursera.
- This course requires a payment of $79 and 11 weeks in order to complete. It has a combined amount of 113 videos, 229 readings, and 1 practice quiz.
  - This class will teach you how to use machine learning effectively by practicing the correct techniques and getting them to work on your own. You will know the theoretical aspects of the learning as well as the practical knowledge needed to be able to quickly and powerfully apply the correct techniques to all sorts of problems you may encounter. The main focuses of this course are machine learning, data mining, and statistical pattern recognition. There are a number of case studies and applications that are incorporated in the teaching in order to learn about how one can apply learning algorithms to building smart robots, text understanding, computer vision, medical informatics, audio, database mining as well as other areas. The lessons included in this course explore topics including (i) supervised learning (parametric/non-parametric algorithms, support vector machines, kernels, neural networks), (ii) unsupervised learning (clustering, dimensionality reduction, recommender systems, deep learning), and (iii) best practices in machine learning (bias/variance theory; innovation process in machine learning and AI).
  - This is the highest rated course on Machine Learning that is offered on the internet with an earned rating of 4.9 stars out of 5. Taken by more than 1,678,000 students, it is a reliable source to gain information about the topic of machine learning and base a course of your own off of. This is a course taken around the globe by students and professionals so the information can be applied to a beginners course or one that is more geared towards those with some background knowledge in computer programming and machine learning.
- Machine Learning A-Z: Hands- On Python and R in Data Science by data scientist and forex systems expert Kirill Eremenko and data scientist Hadelin de Ponteves offered through the website udemy. This course is currently on sale for $11.99 (originally priced at $199.99) and requires 40 hours a week to complete made up partly of 19 articles. Prerequisites include just basic high school math skills.
  - From this course you will learn how to master machine learning on Python and R, gain an understanding of machine learning models, how to predict accurately, how to use tools like reinforcement learning, NLP and Deep learning. Most importantly you will learn how

when encountering a problem to use the correct model in order to solve it. The course is broken up in 10 parts which take you through topics such as data preprocessing, simple linear regression, SVM, and hierarchical clustering. Udemy, the website this course is offered through, is a platform where instructors can build courses on any topic they desire. They can teach their courses using videos they uploaded, power point presentations, pdfs, audio, and zip files. If in need of assistance or if there are any concerns the students can interact with the teachers through live classes and online discussion boards.

  - The audience for this course is a wide range of people from anyone with high school math skills that are interested in machine learning to people looking into becoming a data scientist. Making it a good source to base a class off of due to the fact that the lessons will be easy to follow for students that may not have that much background knowledge of computer programming while still being interesting and beneficial for those that do have that knowledge. It is an even more appealing course to take due to the fact that it is significantly cheaper than many other courses offered online while still offering the same amount of lessons and topics covered.

- Principles of Machine Learning by Graeme Malcolm (senior content developer of the Microsoft Learning Experiences team), Steve Elston (data scientist with over two decades of experience using machine learning), and Cynthia Rudin (head of the Prediction Analysis lab at MIT). This course is offered through the website edx in association with Microsoft. This course is free, however there is a $99 fee if you wish to get a certificate. It is a 6 week long course, making it 3 to 4 hours per week.

  - The main topics explored in this course are classification, regression in machine learning, how to improve supervised models, details on non-linear modeling, and recommender systems. Within these topics you will learn more specific skills such as how to understand the operation of classifiers, how to use regularization on over-parameterized models, how to correctly apply and evaluate hieratical clustering model, and understanding how to evaluate recommenders. Each of these main topics have the more specific skills taught and those are followed by a lab where you apply the information that has been taught. Besides the labs the course includes a hand-on experience building, validating, and developing machine learning models.
  - Due to this being an intermediate course offered by Microsoft, there is a lot of information that can be taken and used to outline a course for students with little computer programming experience that want to learn about machine learning in an easy way. The information taught can also be applied in other aspects of life because the course examines machine learning theory combined with practical scenarios.

Not only can the information be applied in other situations in life, you can use this course to continue your computer programming career as it is part of both the Microsoft Professional Program Certificate in Data Science and the Microsoft Professional Program in Artificial Intelligence. Taking this course will not only improve your understanding of machine learning but it will give you base to continue with that knowledge and build off of it.

- Machine Learning & Artificial intelligence: Crash Course Computer Science #34 narrated by Carrie Anne Philbin, an award winning secondary Computing and ICT teacher, author, and YouTuber. This video can be viewed through YouTube from the channel crash course.

- Since YouTube is free, you can watch this video for no cost. It will take about 12 minutes to watch. Additionally, this video is one part of the computer science playlist offered by crash course with another 40 videos you can learn from.

  - This video begins by setting up a situation in which machine learning is applied. It looks at identifying unknown types of moths and how graphs can be set up and used in the identifying process. They go further and talk about other applications in the real world such as how self-driving cars and machines used for disease diagnosis. Philbin also mentions other topics such as how algorithms work in association with machine learning, decision trees and other machine learning techniques, statistics, deep learning, and narrow verses strong artificial intelligence.
  - By incorporating this video into your teaching it will make it easier for your audience to understand the applications of machine learning and how can impact the future. With the visuals such as graphs and animations it makes the video intriguing while still being informative. This is a fun and easy source that can benefit you in showing to future students learning about machine learning.

2.17 Maddy Weaver

There are various resources available for learning about machine learning. These resources include online courses, books, videos, and more. Whether its a crash course on deep learning or a six month course on data science, many of these resources deem valuable and can be very helpful to anyone who is curious about learning more aspects of machine learning.

Many people know of the famous online shopping website, Amazon. Machine Learning algorithms is very important to Amazons internal systems. From their recommendations engine, to Echo powered by Alexa, Prime Air (drone initiative), and more, machine learning is at the core. Because Amazon relies heavily on artificial intelligence, they have decided to share their learnings of machine learning with their new online program AWS. AWS allows developers, data scientists, beginners, or anyone work on their platforms for machine learning as needed. They support all of the major machine learning frameworks so that you have various options for developing models.

AWS Frameworks:

- TensorFlow
- Caffe2
- Apache MXNet

AWS has other features such as AWS DeepLens, which is the world's first deep-learning enabled camera for developers. This allows you to "get up and running with deep learning quickly and easily." AWS is a great resource for machine learning because it is tailored to fit anyone who wishes to use it and offers many options for building your own models. It is also is very secure. It offers strong encryption to keep your data secure. It also allows you to choose whether you or AWS will manage encryption keys.

Udacity is another great resource for learning online. Better yet, it offers a free course for an intro to Machine Learning. Udacity gives a clear synopsis of what the course offers and explains how machine learning is the key to careers in data analysis. This inro to machine learning course teaches you how to find useful features that represent your data best, important algorithms, and more. It is also about 10 weeks long and is split into 10 lessons. Some of their courses that are helpful to take before this class include:

- Intro to Computer Science
- inferential Statistis
- Descriptive Statistics

- Intro to Data Science

This course is great because you learn by doing and take on real-world problems. By the end of the course you will be able to analyze data with the use of machine learning. They teach you how to use helpful tools and also include instructor videos, taught by industry professionals.

If you're going about the fundamentals of machine learning, "Machine Learning in Action" would be a good book to go to. This book focuses on algorithms is very "hands on text." Along with Algorithms this book also has topics about worked examples in Python and lots of exposition rather than math. This book is heavy on algorithms so it would be a good way to introduce beginners to them. It also suggests that you take a look at Python before diving into this book. You will also learn how to implement classic algorithms such as Apirori and Adaboos. Everyday data analysis will be included as well as common ML tasks. The author of this book, Peter Harrington is a professional developer and data scientist. His work has been published in numerous academic journals, and he holds five US patents.

Standford offers a machine learning course on coursera which would be very beneficial as well. They give an overview of what Machine Learning is, how it is being used today, and more. They state that with this class, you will learn the most effective machine learning techniques, and gain practice using them and getting tem to wor. They also teach you how to address problems in the process. There are 11 sections to this course as well as various subsections, each ranging from about 1–5 hours each to complete. Here are the sections:

1. Introduction
2. Linear Regression with One Variable: Linear regression predicts a real-valued output based on an input value. This section discusses the application of it.
3. Linear Algebra Review: This optional module provides a refresher on linear algebra concepts.
4. Linear Regression with Multiple Variables
5. Octave/Matlab Tutorial
6. Logistic Regression
7. Regularization
8. Neural Networks: Representation
9. Neural Networks: Learning
10. Advice for Applying Machine Learning

11. Machine Learning System Design
12. Support Vector Machines
13. Unsupervised Learning
14. Dimensionally Reduction
15. Anomaly Detection
16. Reccomender Systems
17. Large Scale Machine Learning
18. Application Example: Photo OCR

Another valuable resource is "Deep Learning for Dummies" which is a book offered by Hewlett Packard Enterprise. This book guides beginners into understanding what AI, deep learning, and machine learning can mean for you or your oranization. The authors of this book, John Paul Mueller and Luca Massaron, understand that machine learning can be confusing but they express how important it is. Without it, fraud detection, web search results, real-time ads on web pages, credit scoring, automation, and email spam filtering would not exist. The two authors are data science experts making this book a very credible source.

Another course that would be helpful is one that Udemy has called *Deep Learning A–Z: Hands-On Artificial Neural Networks*. This course allows you to learn to create Deep Learning Algorithms in Python. You are taught by two Machine Learning & Data Science experts which make it even better. Things that you will learn in this course include:

- Understand the intuition behind Artifical Neural Networks
- Understand the intuition behind Convolutional Neural Networks
- Understand the intuition behind Recurrent Neural Networks
- Understand the intuition behind Self-Organizing Maps
- Understand the intuition behind Boltzmann Machines
- Understand the intuition behind AutoEncoders
- Apply Artifical Neural Networks in practice
- Apply Convolutional Neural Networks in practice
- Apply Recurrent Neural Networks in practice
- Apply Self-Organizing Maps in practice

- Apply Boltzmann Machines in practice
- Apply AutoEncoders in practice

In this course you will learn how to use Tensorflow and Pytorch. TensorFlow was develped by Google and is used in their speeach reconition system and more. Pytorch is just as powerful and being developed by leading universites such as Standford and Oxford. They also offer many exciting projects including Stacked Autoencoders to take on the challenge for the Netflix $1 Million prize.

As you can see, there are many resources avaliable for teaching machine learning on the web. Many will introduce you to different algorithms, tools, programs, and more concepts. Here are some to get you started!

2.18 Peter Weber

Machine Learning Recommendations.

The courses that I looked at are university courses that have a public syllabus and have online class notes and lectures available. Here are the best ones that I found:

University of Tennesee Chattanooga CPSC4430 Introduction to Machine Learning

University of Texas at Austin CS 391L: Machine Learning

Stanford Untiversity CS229

MIT OpenCourseWare Machine Learning

The Machine learning courses that I researched start with learning what machine learning can do and what it cannot do. Where is machine learning useful and what are its limitations? This can be covered in the first day of class. There are two main types of problems that can be solved with machine learning: regression problems and classification problems. Regression problems are used to predict a chararacteristic of something based on patterns learned from how other known chararacteristics interact. A basic example is trying to guess how much a fruit is worth based on its weight, this is linear regression. Classification problems are exactly what they sound like, looking at data to try to determine what something is. The example we saw it class was a classification problem where the computer had to figure out if a fruit was an apple or an orange based on its weight and texture. As the course goes on, it can cover more complex topics that I am not going to research now but here are some concepts that I saw mentioned:

- Linear Regression
- Gradient Descent
- SSE
- Decision boundaries
- Nearest neighbor methods
- Linear classifiers
- Logistic Regression
- Ensemble methods
- Decision tree
- Boosting

- Unsupervised learning
- Artificial Neural Networks

Making a course based on these courses can work well because they will give you a fairly thorough knowledge in machine learning, and between them they cover many concepts that are learned in different ways. Taking these courses will not only teach the concepts, they will also demonstrate different ways to teach the concepts and give insight on how to design a new introductory course about machine learning.

Materials used in classes:

- UTC intro to machine learning class required book:
  Ethem Alpaydin, "Introduction to Machine Learning", Second Edition
  According to the syllabus, this book covers all of the material in the course.
- University of Texas course article:
  Thomas G. Dietterich, "Ensemble Learning"
- University of Texas course textbook:
  Tom Mitchell, "Machine Learning"
  Written in 1997, the book may be somewhat outdated but seems to be more about concepts and since it is still used by this course (I am not sure how old this course is), it probably remains at least somewhat relevant.
  Description: "This book provides a single source introduction to the field. It is written for advanced undergraduate and graduate students, and for developers and researchers in the field. No prior background in artificial intelligence or statistics is assumed."

Sources:

University of Tennesee Chattanooga CPSC4430 Introduction to Machine Learning syllabus.

University of Texas at Austin CS 391L: Machine Learning syllabus.

Stanford Untiversity CS229 syllabus cs229.stanford.edu/syllabus.html/.

Rohit Singh, Tommi Jaakkola, and Ali Mohammad. *6.867 Machine Learning*. Fall 2006. Massachusetts Institute of Technology: OpenCourseWare, https://ocw.mit.edu/.

Chapter 3

Algorithms: how machine learning works

3.1 Sheldon Branch

Sheldon Branch CSC131 9/12/18 Machine Learning and Self-Driving Cars

We live in a society that is constantly bombarded with innovation. This innovation has propelled us to the age of artificial intelligence (AI). AI is the future. It can diagnose cancer patients, synthesize images, and even drive our cars. Over the last decade, car accidents have resulted in 30,000 to 35,000 deaths annually in the U.S. Automated cars could reduce that number drastically. It can remove the risk of drunk, distracted, and reckless driving, all of which are leading causes of accidents.

The on-board computer of a self-driving car has several different sources of information. Through GPS, it determines its location in the broader context. By using a roof-mounted camera, it gathers information in 360 degrees; it reads traffic signals and signs and detects road lines to stay in its lane, as well as speed and trajectory of the surrounding cards, pedestrians, animals, and various animals. Through all of these sources, a self-driving car develops a database of the environment, and obstacles totaling hundreds of thousands of miles over several years. The information from each vehicle is then combined and shared with every other self-driving car, creating an incredibly large amount of data and collective learning available to all cars.

Getting self-driving cars to recognize other cars is surprisingly difficult. Recently, however, a study has been shown that self-driving cars can be taught the rules of the road by studying virtual traffic on video ages such as Grand Theft Auto V (GTAV).

Companies like google and uber are teaching their software by physically driving millions of miles in the real world. Racking up that kind of distance with a prototype is not easy and feeding images to the software takes people a lot of time and effort. Computers need hundreds of thousands of laboriously labelled images, showing where vehicles begin and end, to make them expert vehicle measurers. This is very inefficient and time-consuming.

A very good alternative to this onerous task is the use of video games. Picking out cars in a video game is a similar task to doing it in reality, with the advantage that everything comes pre-labelled because it has been generated by the games software. A team at the University of Michigan in Ann Arbor trained an algorithm solely using GTA V and tested it against the same algorithm trained on real world images. The GTA-trained algorithm performed just as well at spotting cars in a pre-labelled data set. The video-game version needed around 100 times more training images to reach the same standard but given that 500,000 images can be generated from the game overnight, that is not a problem.

Finding the right training data is difficult. AI is usually trained on images from similar locations at similar times of day, under similar weather conditions, and then tested under similar conditions. This means that its hard to tell whether the computers can genuinely recognize cars, or whether they have just memorized that particular data set. Using video games can help because they often show a verity of vehicles and conditions but the problem is still there. For example, there is nothing in GAT V that likes like a city in Japan. This lack of data could pose a problem for self-driving cars in parts of the world that arent represented in video games.

The goal of self-driving cars is to show that driverless cars are safer than human-driven ones. In the real world, there is a car fatality for every 100 million miles driven. The first step could be to make sure every system has been tested using video- game-style simulations before it hits the road.

3.2 Bram Dedrick

Neural networks are a fundamental principal of machine learning. Neural networks were first though of over fifty years ago, but until recently they were not practical or possible or build. They were originally theorized to be used in a computer that would think like a human, using a chain of thoughts somewhat resembling the neurons in our brains, hence the name neural networks. The basic principal of neural networks is complicated, but understandable without any sort of previous knowledge in machine learning or computer science. Neural networks take in number and spit out different numbers unlike the neurons in our brains which take in chemical signals and subsequently shoot out similar chemical signals.

Neural networks can be used for classification. Classifications is term in machine learning where data is input and then classified into categories. This can be used for distinguishing species or other things based on specific traits like size or location. As humans, we can visualize these data types by using a scatterplot with 2 parameters or a 3D scatterplot for three. If we put weight on an animal on the X axis and height of an animal on Y axis we can begin to do our own classification. If we then take input sample data of the relative height and weight of cats and dogs we then have the beginnings of a classification problem that is easy to visualize. We can then draw decision boundaries on the graph which break up the data into groups based on categories. If we draw one boundary that says if the weight is greater than twelve pounds it is likely a dog, and a second boundary that states if the height is greater than twelve inches it is likely a dog. These two boundaries will have some amount correct and some wrong as it will not always account for small dogs or big cats. The amount correct and the amount incorrect can be displayed in what is called a confusion matrix. A confusion matrix is a Punnett square that on one side displays the species guessed and on the other has the actual species. For example, lets say that we had 100 cats and 100 dogs in our training data and we classified 95 cats correctly and 75 dogs correctly with our two boundaries. This would put our classifier at an 85

When doing this with more than one species, or more than three parameters the process becomes significantly more confusing as there is not a good way of visualizing it. This is where neural networks can come in very handy. A neural network can consist of many layers, but for this example I will only use three and I will stick with the analogy of cats and dogs. The first of the three layers here will be the input layer, consisting of two neurons, height and weight. The second layer is the hidden layer, this is where the classifier works to determine what to send to the final layer. The third and final layer is the output layer, this layer will consist of two neurons, the probability that the animal is a cat and the probability that the animal is a dog. The neurons between each layer are connected by links forming what looks like a network of lines.

The hidden layer is where the classification takes place so that is where most

of the computing takes place. In any specific neuron in the hidden layer it will take the numbers from the input layer, multiply them by a weight and the sum those weighted numbers. The next thing the neuron does is apply a bias, a constant number, to the result of the weighted sum plus this bias is a new number that is combined with all other outputs from the other neurons in that layer and then sent to the next layer which in this example is the output layer. The weight value and bias are initially set to random numbers, then tweaked by an algorithm, gradually over many integration, that determines who good the answer is from data it already knows the answer of. For more complex neural networks, the hidden layer can be more than just one layer deep, hence the term deep learning.

For a visual representation, check out my source video from the Crash Course YouTube channel. https://www.youtube.com/watch?v=z-EtmaFJieY

3.3 Tony Ferenzi

Machine learning is one of the more difficult topics to learn. It is even harder to figure out where to start as machine learning is such a complex and wide subject that is being added to every single day. The more machine learning expands, the more it is incorporated and used in a variety of careers. That is why this subject is becoming so popular and important to understand around the world. Even though it may be difficult for some to first get started, machine learning should not be feared or labeled to difficult to do. A list of algorithms will be discussed thoroughly in order for new students to hopefully learn and understand them as they begin their studies.

Starting with classification, these algorithms are used for things such as speech recognition, telling whether a person is male or female, and detecting spam mail. Classification is a type of supervised learning where the program is simply supplied data. It then takes that data and uses it to make predictions. Some examples of classifiers are Nave Bayes, random forest, boosted trees, decision trees, support vector machines, and logistic regression. The reason decision trees are called this is because they break down data into smaller and smaller subsets until they start to look like trees. This tree has decision nodes, which has two or more branches, and a leaf node, which represents a classification. Decision trees are what essentially make up the ensemble learning method known as random forest. Random forest works by creating multiple decision trees with the final product being the mode of the classes of each and every tree.

Another type of machine learning technique is dimensionality reduction. Algorithms in this category basically remove unimportant or repetitive variables in order for a simpler and more accurate prediction, such as principle component analysis. This also helps to avoid over training and over fitting your algorithm. Also, in principle component analysis, the predicted variables are all independent of each other which means they can be used in a linear model. Anyone can use this algorithm if they have multiple different variables to draw form, such as finding the gross domestic product product of the United States. There are three main reasons for using principle component analysis. The first reason is to reduce the number of variables but do not know which ones to keep or take out. The second is if you want to make sure all your variables are independent of each other. The third and most important reason is if you are fine with making your standalone variables less interpretable. Linear Discrimination Analysis is a type of dimensionality reduction and is actually closely related to principle component analysis, but it is different as it can be used for different applications. Another difference is that Linear Discrimination Analysis attempts to discover the differences between data sets or classes, while principle component analysis does not focus on the differences.

Regression is an important topic in machine learning. Algorithms in this class can predict a continuous-valued output, such as the popular linear regression algorithm. This algorithm is apart of the field of predictive modeling and trades

interpretation for reduced error and more accuracy. Linear regression is an extremely popular and useful algorithm that can be used for predicting house prices, forecasts, and trends for products. There are actually two types of linear regression. The first is single linear regression, where one independent variable is used to calculate and predict the dependent variable. The other is multiple linear regression, which is the opposite of single linear regression and involves multiple independent variables.

Clustering is another popular topic that many use in machine learning. This topic focuses on the automatic grouping of unique and similar objects into sets, also known as clusters. Clustering is used in search engines, such as grouping similar products and presenting them to the consumer and excluding different ones that the consumer may not be interested in. It could also group the consumers themselves into market segments in order to focus advertisements. Algorithms in this class include K-means and Latent Dirichlet Allocation. K-means clustering is a part of unsupervised learning and is mostly used when you have unspecified data. The goal of the algorithm is to obviously classify unspecified data by grouping them into sets based on the features of the data, represented by the variable K. K-means is used in inventory categorization, such as grouping sales activity and monitoring the manufacturing of products. The algorithm is also very fast, which makes it useful. There are also downsides to the algorithm though. The first downside is that when using the algorithm, you have to specify how many groups you want the algorithm to make. The second is that the algorithm chooses a random center point each time you run the algorithm and therefore it may yield different results with each use. Latent Dirichlet Allocation is actually a type of generative probabilistic model. This algorithm is used to group and classify certain text in a document to a specific topic. This is useful if you want to draw from multiple sources of text and group specific information into a certain section, such as news headlines or research papers.

The last major topic of machine learning is collaborative filtering. Collaborative filtering is basically an algorithm that automatically predicts items based on the interests of a person. It may also take into account the interests of another person who has similar tastes in products as someone else. With these similar interests, the algorithm will then advertise the products of both people to each other as they will have a higher chance of liking the product. This topic is mostly used for things such as movies and product recommendation. A specific algorithm in this topic is alternating least squares.

These are just the basics of machine learning. There are thousands of topics, methods, algorithms, and techniques used during machine learning that differ from each other in many ways. The only way to understand all of these subjects is to simply use them in practice. Quiz yourself by making programs that incorporate and connect algorithms together to explore what they do and refresh yourself on old ones. It is becoming increasingly important to understand machine learning and anyone who is new to it may feel overwhelmed. But with

determination and effort, anyone can learn and even make a career out of it.

3.4 William Golden

3.4.1 A brief introduction

Here is a collection of terms, programs, and algorithms. These pieces have been broken down into definitions and explanations that are easy to understand. Hopefully, this simplified gathering of information on computer learning will provide help to those looking to learn a little more about the world of computers and how they learn.

3.4.2 Programs for Computer Learning

Note: these programs are all free and great for applications in computer learning

- TensorFlow - made from the Google Brain team. It offers a large library of pre-built programs and the ability to create numerous systems for machine learning.
- Scikit-Learn - available through Python (computer developing software) which offers tools for computer learning. Specific functions include Clustering, Regression, Classification, Dimensionality reduction and more.
- Weka - offers tools for data mining. It offers the basic functions of classification, linear regression, clustering and more. There appears to be a library of numerous tools and algorithms for computer learning.

3.4.3 Good Terms to Know

Here are a several terms that every computer engineer should know. Some of them are more common than others, but understanding all of the terms will help students understand the complex language of programmers describing computer learning.

1. Supervised Learning - a way of how machines learn. It includes showing the AI (artificial intelligence) a set of algorithms. These algorithms contain problems with their desired solutions. After the AI views enough of the examples, it can then apply the algorithms to new incoming data.
2. Unsupervised Learning - How machines classify "stuff." The AI is given input data and then without any labels or given information, sorts the information based off of its own observations. This is different from supervised learning because there are no examples or labels given prior to inputting data.

3. Linear Regression - A way to interpret information between two variables. The linear regression shows how data relates. Through a linear regression, data can be predicted as well.

4. Method of Least Squares - A way to calculate the line of best fit (a line drawn equal and straight through the data points) on a linear regression

5. Feature Set - set of data which computers can learn from.

6. Test set/Training set - training set is a specific series of inputs for a machine. The inputs help "train" the machine to perform specific functions. The test set is slightly different input data compared to the training set. It is used to measure a programs performance and ability to solve foreign problems.

7. Classification - sorting inputs into groups of similarities.

8. Clustering - grouping data points with similar properties; a form of unsupervised learning

9. Data mining - finding patterns and relationships in large sets of data

   (a) see section 'Data Mining in Use' for more information

10. Deep Learning - a way computers learn; built to mimic the human brain's process of thinking/learning by grouping raw data into complex clusters and patterns.

11. Neural Network - a computer system built to mimic the human brain. It interprets data in many different ways and creates complex analysis of raw data.

12. K-nn (k Nearest Neighbors) - a simple way of classifying data. The computer looks at how closely data is grouped together and creates classifiers to each group of data points

13. Naive Bayes Method - a process of using a set of complicated algorithms to help computers classify data

14. Dimensionality Reduction - a series of algorithms used to remove random data (or outliers) in a set of data.

3.4.4 Data mining in use

This is a great example of how an application of machine learning known as data mining can be applied in real-life problems.

Healthcare Data Mining: Predicting Hospital Length of Stay of Dengue Patients
Wiratmadja, Iwan Inrawan, 2018

Currently, computer scientists are working with hostpitals in southeast Asia to find the estimated hospital stay of patients with Dengue, a mostiquito-borne illness. Hospitals are able to predict the time each patient needs in the hostpial by using data mining methods. To data mine the information, scientists collected a wide aray of statistics (age, sex, height, etc) about current Dengue patients and how long they stayed in hospitals. The data then was put into a data-mining algorithm to make predictions of how long future admitted patients will stay in the hospital according to that patients health information. This is helping hospitals gather more needed supply and prepare for incoming Dengue victims.

3.4.5 Algorithms

Here are two examples of helpful algorthms used in machine learning programs.

- Log-Loss Function

  The log-loss function is used in many business that wish to know the productivity of their machine learning programs. This algorithm is used to check the accuracy of a machine that makes predictions with a single probability (either it is yes or no, 1 or 0). When using the log-loss function, you will get a number which tells the reader how likely the machine is to making the correct guess on classification.

  For example: a machine predicts which houses will be sold or not. It says there is a 80% chance house 1 will be sold, 30% chance house 2 will be sold, and 10% house 3 will be sold. Later on, houses 1 and 2 were sold and 3 was not. The log-loss function can look at the machines initial prediction and compare it to the actual outcome, then produce a number saying the accuracy of the machine in future predictions.

  Note: this algorithm can ONLY be applied to machines that make predictions resulting in a yes or no result (on/off, 1/0). However, many other types of algorithms exist which can analyze jdifferent programs in this way.

- Bayes Theorem

  Bayes Theorem is used to find the relationship of conditional probabilities. What this means is that by using this algorithm, you can find the chances of something happening (A) if something else (B) occurs.

  Here is a good example:

  There is a 40% chance of it raining on Sunday. If it rains on Sunday, there is a 10% chance it will rain on Monday. If it didn't rain on Sunday, there's an 80% chance it will rain on Monday.

  It rained on Monday. What is the probability it rained on Sunday? That is where Bayes' theorem comes in. It allows us to calculate the probability of an earlier event, given the result of a later event.

3.4.6 Creating Speech Recognition

The ability for computers to understand speech has been an incredible discovery in the world of computers. Many improvements of this technology continuously emerge today. For the final part of this section, we will look at a simplified explanation on the basics of how this type of machine learning works.

The first step: a sound wave is created on a computer.

The computer program can analyze this data and interpret what is being said. However, there are many complex things that go into the programming of the computer in order to make it understand what this data means.

To begin, the program needs *features*. Features are essentially the raw data of the audio wave. Every ten milliseconds of speech can be broken down into many features (approximately thirty-nine). The computer looks at the group of features (also known as feature vector) and begins to translate.

The computer program also contains a *model*. A model is the computer's database (or how it interprets what each feature means). When the computer receives the human voice, it replicates the data using it's model and then predicts what the original speech was based off of the data made from the model.

side note on models: The most common model used in voice recognition systems is called a *Hidden Markov Model*. The Markov Model maps out random patterns and probabilities of speech particles. For example, after sound (a) is made, it knows the likelihood that sound (b) or sound (e) will follow next. This allows the machine to remain coherent and function better while translating a long line of speech.

So overall, the machine receives the audio code, breaks it down into little pieces, re-creates the pieces using its own set of tools, then tells the user what it created. Simple, right?

Unfortunately, there are still a million other tiny things programmers have to do in order to create a fast and accurate voice recognition system.

Specifications to Consider

- Speed - (if the program looked at every single feature of a audio recording, it would take too long. The program must quickly process the data by using tricks to decipher things in short amounts of time.)

- Accuracy - The program must produce accurate results, otherwise users would become infuriated. Programmers use Receiver Operating Characteristic curves (ROC curves) to view the accuracy of their product. The ROC curve is a graph with a line showing the correct/incorrect translation ratio at certain points.

3.4.7 Conclusion

We have looked over some of the current language and functions of popular mechanics in machine learning. The world of computer learning is growing deeper and deeper. While studying the basic algorithms, programs, and tech-language of this field, it is always good to remember that there are still many possibilities unexplored by computer scientists. Also, many of the procedures and functions of the field are susceptible to change.

3.5 Yuan Hong

Machine Learning Yuan

Big data and machine learning are two hot areas of rapid growth in the information and science industry. From the occlusion of information in the past to the current data explosion, the volume of data and the growth rate of data in all fields have grown at an alarming rate. According to the National Security Agency, the Internet processes 1826 petabytes per day. As of 2011, digital information has grown nine-fold in the past five years, and by 2020 this number will reach 35 trillion gigabytes. The scale of this digital data brings tremendous opportunities and potential for change, and can take advantage of the integrity of these data to help us make better decisions in all walks of life, providing a very transformative approach to data-driven research in scientific research. A good example. It enables us to use the search engine's big data forecasting ability in the field of medicine, astronomy and other fields in the scientific research of data. Compared with traditional machine learning, machine learning under big data greatly expands the number of samples, so that many problem categories are supported by rich samples. This is the advantage of big data, but it also causes many problems. Now with the continuous optimization of hardware technology and programming algorithms, the collection and magnitude of data is no longer a major problem hindering big data research. The relationship between data, that is, which data is useful, redundant, or even interferes with other data, how the data sometimes interacts with is the main challenge facing big data. Big data has enormous potential value in all aspects of our society. Getting valuable information from big data is not a simple task. It is a core goal of big data technology to dig out the laws hidden in the data and the information we need from the huge amount of data and the huge amount of data to maximize the value of the data. Traditional machine learning is generally considered to be a shallow learning structure. In contrast, deep learning refers to machine learning technology that automatically learns deep-level architectures under supervision or unsupervised and is used for classification or data mining. . Inspired by the human brain's signal processing model in nature, the concept of deep learning has been proposed and is receiving more and more attention due to its superior processing performance in many fields. Many companies are now making full use of the advantages of big data to apply them to commercial products, and have achieved great success. These companies and organizations collect massive amounts of information and analyze them on a daily basis based on a large amount of data generated, and then use the analysis results for related projects in deep learning. For example, Siri, a virtual personal assistant for iPhones, offers a variety of services such as weather forecasts, sports news, answering user questions and reminders. Google will apply massive pieces of chaotic data to deep learning algorithms. These pieces of data come from Google Translate, Android's speech recognition, Google Street View and search engines. Other industry giants are not far behind. There is currently no uniform rule for big data. Generally speaking, big data is understood as data that cannot be loaded

into the memory of the computer. This is an informal definition because each computer has a large size that cannot be loaded into the memory. data set. The industry's characteristics of big data have expanded from the original 3V model to the current 4V model, including: volume of data which is volume, variety of data which are types, low value of data which are values, and many real-time data require fast processing, which is the velocity. In response to these characteristics, knowledge analysis in the big data era, machine intelligence and human intelligence coordination work and intelligent analysis systems will play an important role. People need an intelligent analysis interface to connect humans with the computer world, otherwise they will be lost in the torrent of data. With pattern recognition, in machine learning, and deep learning represent three different schools of thought. Pattern recognition is the oldest method, as a term, it can be said to be very outdated. Machine learning is the most basic, it is one of the hot spots of current startups and research laboratories, and at last deep learning is a very new and influential frontier, and we won't even think about the post-deep learning era. The three are related to each other and different. Pattern recognition In the big data environment, due to the large volume and various structures, using machine learning to achieve its classification and recognition has many advantages, while machine learning is deep learning, the continuation and future direction of development. In the whole process of machine learning development, there have been two research directions: one is to study the learning mechanism, focus on exploring and simulating the learning mechanism of human beings; instead, it is to study how to effectively use information and focus on obtaining hidden and effective data from huge amounts of data. Knowledge that is understandable. The study of learning mechanism is the source of machine learning. However, with the continuous increase of data analysis needs in various industries in the era of big data, the efficient acquisition of valuable information through machine learning has gradually become the main driving force for the development of machine learning technology today.

3.6 Easton Jensen

Randxom Forests Random forests or random decision trees is a method to classify or find the regression of data. Random forests are a multitude of decision trees put into one system. The main focus point of a random forest is that having multiple trees to classify an object will create more accurate findings. Finding the mode of all the decision trees or the mean of each individual tree. Decision trees themselves take data and can overfit the data. In doing so the tree makes the data go wonky, but having a forest where multiple trees are working on the same problem and each tree classifies the problem. This allows the random forest to average out or take majority vote of the trees to get a more accurate answer. This technique is called bagging. Bagging is used to average out unbiased and noisy data and create a model with a low variance. Let's take the example of four trees that classify the same data and the first tree identifies it as class one. The next tree classifies it as class two. Trees three and four then classify it as class 3. The random forest then deems the data as class three. The random tree can use unsupervised learning and supervised learning as a basis to get the answer.

Unsupervised Learning Unsupervised learning takes data that has not been put into categories or been classified. This makes the results put out by the machine hard to determine because finding what is right and wrong is not straightforward. Putting in data that isn't labeled gives the machine a lot of work to do. This makes the machines predictions all its own thinking". Unsupervised learning can be used in clustering, anomaly detection and neural networks to name a few.

Supervised Learning Supervised learning takes an input and creates an output that tries to match the test data previously given to the machine by you. You give the machine inputs and corresponding outputs that give the data a basis to predict data further down the road. Giving it this training data gives the machine what it needs to run the way you want it to. Also throughout the process you will be supervising the machine just as the name implies. The data you give the machine is labeled and that is a big deal. If the data was not labeled then the machine would not be able to tell which output goes with what input.

3.7 Rodrigo Martinez

How Machine Learning Works Machine learning has several parts that you must understand in order to know how it works. Each part usually connects to another sometimes they may not but they are very crucial to know how it works when it comes to learning and understanding machine learning. Firstly, how machine learning works is basically controlled by functions called algorithms. In other terms machine learning runs by given recipes or instructions. Algorithms can be made up of one of the following expressions which include: mathematical, english, and pseudocodes. Mathematical notation basically uses multiplication, subtraction, addition, and division. English notations is the basic english typing you are seeing right now. Pseudocoding is more of the typical coding you would see in java, python, c++, etc. These Algorithms then commands the machine to do some certain things. Machine Learning algorithms can include several things such as supervised and unsupervised learning which then can connect to linear regression. Supervised and unsupervised learning works in a way in which machine learning can teach itself to draw out something. For example a graph with several random points can be drawn out by the machine by using its given resources within the algorithm. This can vary of how a graph can be drawn out. If it were using supervised learning then it would follow specific instructions on whether it can or cannot do during the process of graphing out points. If it were to have used unsupervised learning it would have to have self teach itself by trial and error and through an analytical process. Machine learning algorithms can be used create complex software, machines, etc. For example, machines can also be instructed to do other things such as recommendations to customers. This works by using certain methods which are training sets, classification, neural network, and deep learning. In order to get specific information you must collect data about your customers. This will then go to the classification process in which will be later be added to the machine. Classification is when you separate certain pieces of information into its closest groups, this helps organize the data. Once the information has been classified it will ongo to a training set. A training set helps the model/machine train itself from doing a certain task, for this example we are designing the machine to recommend customers certain things. This will require the classification data such as sex, age, race, etc. The machine will then work with organizing itself by using something called a neural network. Neural networks are computer systems in which are structured in a complex manner such as an actual biological neural network. For example, lets say theres a computer you would want to recommend your customers to buy. The machine will have to go through a set of data in order to find the right person to recommend it to. This may include data in which previous people who bought that exact computer. Lets say the data states that most males at the age around sixteen through twenty five tend to buy computers a lot more. If the customer youre trying to recommend fits the specific criteria then itll be a successful recommendation. In order for a machine to learn this type of neural network which will require a lot of classification it will have to go through

something called deep learning. Deep learning is when a machine goes through a process of learning data representations. In this case the machine will have to learn the information given about their customers. Once the machine has undergone through the deep learning from the neural network, it will be able to provide accurate recommendations to the specific customers.

3.8 Matt Morrical

Here is something about how machine learning works.

3.9 Ella Nelson

Here is something about how machine learning works.

3.10 Koichi Okazaki

3.10.1 Facial Recognition

Now a days, development of facial recognition is eye opening, and the technology is becoming familiar with are daily life. For instance, now China has a world most advanced surveillance system "SkyNet" which covers whole country. The surveillance system is mainly using 20 million AI equipped CCTV (China Central Television) facial recognition cameras. By using these cameras police are able to identify the pedestrian and makes them easy to capture and trace criminals. The system knows who is who by using the citizens national ID information stored in the Chinese government's database. However, this system has been criticized because it might cause an invasion of privacy and oppression of government against citizens. Thus, it is said that they could identify the person a second despite there are 1.3 billion citizens in China, and the accuracy is over 98%. In this point AI's ability has already overwhelmed human beings.

Then how does face recognition works? First of all, to identify a person system has to recognize, collect, and measure a face data, but how? Several ways to detect and measure a face were invented, but researchers had concluded that by using deep learning and let the computer find out the measurements to collect itself. For instance, in order to create an algorithm, we have to train Deep Convolutional Neural Network by giving data and create a measurement itself. The process will be like this.

1. Load a disciplinary face image of a known person.
2. Load another picture of the same known person.
3. Load a picture of totally different person. (Medium 2016)

By creating the measurements for those each images computer compares the results, and it then tweaks the neural network slightly so that it makes sure the measurements it generates for #1 and #2 are slightly closer while making sure the measurements for #2 and #3 are slightly further apart. (Medium 2018) After repeating the process thousands and millions of times the system establishes the way to measure and detect the face, and all we have to do is give the face images which will be used as database to the network and get the measurements for each face. Thus, if the camera found a person then classification algorithm searches for the known people's images find out the closest image and finally the system could identify the person.

Using face recognition library called face recognition in GitHub will be easy to install. In order to build the algorithm, you must have Python 3.3 or 2.7, macOS or Linux, and Dlib with python bindings — Dlib is a modern C++ toolkit containing machine learning algorithms and tools for creating complex software in C++ — Dlib web site) in your device.

Citation

China installs 20 million AI-equipped street cameras

Face recognition

Machine learning is fun—Part 4: Modern face recognition with deep learning

3.11 Jakob Orel

There are many different machine learning algorithms that can be used for a variety of processes. An algorithms is a recipe of logical and arithmetic operations the computer can comprehend when translated into machine language. Some algorithms are used to classify data. Some are used to make predictions based on trends in data. Some are used to cluster certain data together. Different machine learning algorithms are more efficient for different types of problems. Some are faster or slower at solving different types of machine learning problems.

/subsectionClassification:

Naive Bayes Method: The Naive Bayes algorithm is a classification technique based on Bayes Theorem. Bayes Theorem is a complex rule that involves independent probabilities to calculate the probability of a new event. This method uses probabilities of independent features to classify objects. Each feature is separate from the others. The algorithm uses probabilities of these independent features to classify unknown data sets. For example an apple may be classified by its red color, roundness, or smooth texture. Each one of these features has a probability that affects the probability of the fruit being an apple separately. The "naive" in the title suggests that the probabilities to do not affect each other. While this is good in a theoretical stance, it may not be practical to many real world scenarios. Naive Bayes is often the first choice to classify data sets because it is very efficient and fast to classify data sets and is used in multi-class prediction.

Simple article

KNN (K nearest neighbor): Another algorithm for classification is the KNN or K nearest neighbor algorithm. This is a commonly used method to separate data without previous assumptions. The training set does not have any general assumptions or very little training phase. A point is set in a group of data. A distance from this point (K) is set. This distance determines which points in the surrounding area are included in this class. The value of K or distance from the center point changes the class points. For example, a bank may use KNN to determine if a person is more likely to default on a loan. The algorithm may determine if a person has characteristics similar to a person who will default on a loan or to a person who will not. KNN is used because of its simplicity, ease of use, calculation time, and predictive power. It may also be used for some regression problems.

Simple article

/subsectionClustering:

K-means clustering: The K-means clustering algorithm is very useful to cluster numerical data. It uses an iterative process of averaging to find the best fit for groups of data points. First, center points(K) are placed randomly in a field of data points. The algorithm then assigns the closest data points to each K

point to separate clusters. Second, the center points are moved to the middle of this new group by averaging the distance to each data point. Then, the process iterates over again. The center points will continue to move. Once the algorithm reaches a point where no data points will change cluster groups, the program stops.

This is a type of unsupervised learning. This is the grouping of data sets that are unlabeled or uncategorized. The algorithm draws inferences based on data without labelled input responses. Supervised learning uses training sets of data that have an equal response to an input. Every input has an output response in the data. This type of learning infers a function by through labeled training data. K means is a form of unsupervised learning that focuses on cluster analysis. This algorithm is important because not all problems have labelled inputs and outputs.

Video Explanation

/subsectionPredicting:

Linear Regression: Linear regression is a statistical method to predict trends in a set of data points. To put it simply, a line of best fit is placed on a graph of data points with a certain trend. This line can be used to make relatively close predictions about future outputs of data. This process uses several concepts including Ordinary Least Squares and Gradient Descent. The sum of all squared residuals is the value of error from each data point to the line added together. The values are squared in order to keep all numbers positive in order to calculate other values. Ordinary Least Squares works to minimize the sum of squared residuals by finding the optimal values for the coefficients. A very common approach is using Gradient Descent to iteratively minimize the error of the model. After the sum of the squared errors are calculated for each pair of input and output values, a learning rate is used as a scale factor and the coefficients are updated in the direction towards minimizing the error. The process is repeated until a minimum sum squared error is achieved or no further improvement is possible. This algorithm is very common to find the best fitting line. Once the best fitting line is found it is very easy to make predictions by solving the equation for a specific set of inputs.

Article

Overall, many different algorithms are used in machine learning to teach the system how to do its job. Data is taken in many forms and has to be used accordingly. Most of these processes work well with the Python language and machine learning libraries. These different types of algorithms cluster and classify data as well as use it to make predictions on trends. The articles that I found are targeting an audience that knows very little about the machine learning concepts or mathematical concepts which makes it easier to understand.

3.12 Marcellus Parks

Marcellus Parks * Machine Learning Software Scikit-learn [https://github.com/scikit-learn/scikit-learn] - Made with Python, this software builds upon Python packages like NumPy, SciPy and Matplotlib to create many math and science programs. It is under the BSD license so it is free and reusable. Scikit-learn uses the C interface to the Basic Linear Algebra Subprograms library or CBLAS for short. The project was started in 2007 by David Cournapeau as a summer of google projected gain numerous contributors upon release. It currently maintains its current status with a group of volunteers updating it with the times.

Apache Mahout- Although applications from here may or may not be tied to Hadoop, the newer versions have supported high-performance Spark framework and they support the ViennaCL Library for GP-accelerated linear algebra. It is, by itself, free to use and reuse. Apache Mahout was made mainly to produce free implementations of machine learning algorithms that are focused mainly on fields like collaborative filtering, clustering and classification. This was built on top of Hadoop using the mapreduce program. However, the creation of 0.10.0 shifts to building backend-independent programming environment, code named "Samsara". The environment has an algebraic backend-independent optimizer and an algebraic Scala DSL unifying in-memory and distributed algebraic operators. At the time of this writing supported algebraic platforms are Apache Spark and H2O, and Apache Flink. The once monopolizing support for MapReduce algorithms is slowly being put out of commision

H20 Prediction Engine- [https://github.com/0xdata/h2o] Originally created with Hadoopin the Java programming language, these algorithms are built for business processes and fraud or trend predictions. The math platform has parallelized and distributed algorithms to make the most use out of multithreaded systems. In supervised learning, it produces a numerous amount of decision trees, or what some call a forest, with increasing refined approximations. In Unsupervised learning, it linearly transforms correlated variables to independent components. Also within H20, is H20 Flow, a web based interactive environment for big data learning. Even though it is compatible with multiple programming languages including Java and Python. It is also free to use and reuse.

GoLearn [https://github.com/sjwhitworth/golearn]- This is the machine learning library created for Googles Go language. Mainly used for big data analysis, it was originally created with two goals, one being simplicity, and the other being customizability.The creators of Googles programming products and libraries made sure that they did not have the large criticisms of other libraries and languages and such. It was is one of of the few libraries with its own language. The language and library have numerous similarities in purpose. Data structures can be easily extended in an application once all of the the different features of the Go language are learned. Pretty much all of Googles software

products are always free.

Weka- Named after the flightless bird this set of Java created machine learning algorithms have a soul purpose of data mining. It is a GNU GPLv3-licensed collection. This means it is free to use, also, that there are official and unofficial packages that can be used. This software also comes with an e-book on how the software works and techniques used to navigate it. Similar to all the other software previously stated, Weka contains tools and algorithms for data analysis and predictive modeling. Since it has a main purpose of data mining it has a graphical interface that is considered better than other software because it designers could spend more time making it easy to use. Not all of the software previously stated even has a graphical interface. Weka has access to SQL databases using Java Database Connectivity. Weka 3 has more applications including educational and research purposes. There are a few downsides to using Weka though, as it is not capable of multi-relational data mining. Also, sequence modeling is not covered by the algorithms on the older versions. Other than for those two rather minor things Weka is one of the most used data mining tools.

This is just a list of the more popular machine learning software and their uses. Because machine leanring is quickly growing field, there will almost always be more and better software to choose from. There are a few links to the free software above. Not all machine learning software is free. Some software has a specific purpose and works better doing certain functions than others. These are all something to build upon, none of these are the finished product. You must tailor these programs to your specific needs.

3.13 Lydia Sanchez

GAN or generative adversarial network, is a machine learning algorithm that uses unsupervised learning. How it is that there are two neural networks pitted against each other, the discriminative and the generative. The generative network creates new data and the discriminative network has to tell which one is real data and which is fake data. Based on the results from the discriminator, the generator adjusts its parameters for new photos and this is continued until the discriminative program can no longer tell which data set is real. One downfall of GAN machine is that the GAN can be temperamental, if the discriminator is duped too easily the generator will not make a realistic looking photo, you need a good balance of the two neural nets. Another downfall is that the GAN will not work correctly if there are not good quality photos for training data. One example is a GAN was creating picture of cats with letter in them because the GAN was training material contained cat memes with text on them from the internet. In the article the GANfather Ian Goodfellow created a GAN machine that created its own realistic photos by looking at real photos, notably realistic fake celebrities' photos. In the GANfather, the generator's goal is to create its own celebrity photos that pass the discriminator test (to lie without getting caught). While the discriminators goal is to detect which the generator's fake images. The GAN can learn to mimic any distribution of data. These are the steps a GAN takes:

- The generator takes in random numbers and returns an image.
- This generated image is given to the discriminator with a stream of images taken from the actual dataset.
- The discriminator takes in both real and fake images and returns probabilities, a number between 0 and 1, with 1 representing a prediction of authenticity and 0 representing fake.

Here is a diagram of how the GAN works.

The GAN is a Neural network, but what is a neural network. A neural network is computer system based on the humans' brain, neural network, learning from examples rather than being programmed with task specific rules. There are many different kinds of neural networks, one being the GAN, however you can group generative algorithms into 3 groups. The groups being:

- Given a label, they predict the associated features (Naive Bayes)
- Given a hidden representation, they predict the associated features (VAE, GAN)
- Given some of the features, they predict the rest (inpainting, imputation).

The GAN also uses unsupervised which is drawing references from data sets that are not labeled rather than being trained to do a specific task, such as learning from labeled example. An example of unsupervised learning is going to a party and meeting total strangers. Now you have to classify them with no prior knowledge and this classification can be on the basis of gender, age group, dressing, educational qualification or whatever way you would like. This learning is different from supervised learning, because since you didn't use any past or prior knowledge about people and classified them "on-the-go". Unsupervised learning is not only used for make pictures, it can also be used in the future to help self-driving cars learn about different road conditions without leaving the garage.

The GAN can do many other things than create realistic fake celebrities' photos, such as help in medical studies. When there is not enough real patient data, the GAN can produce fake data that is almost as good as the actual data, for example help them understand why a drug didn't work. The GAN could greatly help the medical field as well as other field such as the engineering and science field. At Yale they use a GAN machine to help predict how subatomic particles will interact, after having the GAN train on existing data.

In recent year GAN started to be able to make realistic picture based on captions given to them. They give the GAN a single sentence description such as "small yellow bird with a brown crown and a yellow tail." They use a set of the pictures and their associated captions as a training set. The data had to be encoded before giving it to the GAN with a hybrid neural network called character-level convolutional recurrent network. This network takes the image and the caption and creates and joint representation of them for the GAN. The generator makes an image that is close to original image but not the same. While the discriminator receives 3 inputs: the real image with the right text, the real image with the wrong text and, a fake image with right text. The discriminator is looking for similarities between the generator's outputs and the training set, trying to decide with image is real and which is fake. After this discriminator decision the generator and the discriminator both are updated and this process is repeated and they both continue to improve. And by the end of training the generator should be pretty good at fooling the discriminator, and making images from descriptions. Here is the GAN's images of the "small yellow bird with a brown crown and a yellow tail."

The GAN can do many things, not just fake celebrity photos, it can many other fields such as the medical field or for imaging video games, could tell you what a character looked like running in the rain. Many more things such as drug discovery, the GAN can recommend new medicine for incurable diseases by training from available medicine for curable diseases. There are still many drawbacks such as the GAN has a hard time with 3D objects, and has sometimes failed to tell how many of a particular attribute should appear in an object, such as a GAN giving a dog two heads. There is a delicate balance between the two neural network that is essential to the GAN working correctly, and keeping it

balanced is ongoing task. The GAN has a very bright future and can do many great things that would make integral in society.

3.14 Tiff Serra-Pichardo

3.14.1 Artificial Neural Networks and How They Work

An artificial neural network, also known as deep learning, can be defined as a method of machine-learning that attempts to mimic or simulate biological neural networks (the human brain). An artificial neural network is a method of machine-learning in which a computer learns by analyzing training examples. The goal of an artificial neural network is to have a computer recognize patterns and make decisions like a human would. When used, a speech recognition system can be fed thousands of audio clips and their labels, and the computer can find a consistent audio pattern that correlates with specific labels. Rather than just programming a computer, which takes quite a while, teaching a computer to learn may be more efficient.

Since a neural network is modeled after the human brain, which contains billions of neurons, an artificial neural network consists of thousands, and sometimes millions, of simple processing nodes that are tightly interconnected. Most neural networks are organized into layers of nodes that move data in only one direction, this can also be called a "feedforward network." One node can be connected to several nodes above and below it from which it receives and sends data. The middle nodes, between the input and output layers, are called hidden nodes.

When a node receives data, it will assign a number known as a "weight." When the network is active, the node will receive a different data item (a different number) through the connection and will multiply by its associated weight, creating a weighted value. All the nodes in a single layer will add up the sums of the products from before yielding a single number, if the sum is below a threshold value, no data will be passed to the next layer. If the sum is above the threshold value, the data will be "fired" or sent to the next layer. The end data being sent is the sum of the weighted inputs and all its outgoing connections.

When training a neural network, all the weights and thresholds are initially set to random values. The training data is fed into the bottom layer (the input layer) and passes through the proceeding layers above, then gets multiplied and added together in complex ways. The data arrives to the final layer (the output layer) thoroughly changed. During training, the weights and thresholds are adjusted until the training data labels match. This method of feedback is much like how a child learns when they are told what they are doing right or wrong. Once a network has been trained enough, it can reach a point where it's exposed to entirely new sets of data and is able to make accurate labels.

Artificial neural networks are mathematically complex and still have a long way to go before they become more commonly-used, but they are still a promising way and method of machine learning.

3.15 Austin Stala

Part I: Introduction

There are many ways to create and use machine learning algorithms. This essay will focus on two different methods of teaching machine learning programs. The first part of the first method that will be discussed in this essay is supervised learning. This method involves giving the algorithm data and allow it to sort through it on its own. The second part of the first method that will be discussed is unsupervised learning. This method involves giving the algorithm two sets of data first a training set to determine how to sort the data and then a validation set to confirm the algorithm can do its job properly. The second method that will be discussed in this essay is reinforced learning. With this model, you give the algorithm a set of data and a point system if the algorithm gets things correct it is rewarded with points and if it is incorrect it is punished by losing points.

Part II: Un/Supervised Learning

The first type of learning that will be discussed in this paper is supervised learning. The two most common types of supervised learning used are Classification and Regression. Each of these types has many different versions. The first of the two types that I will cover is Classification. One of the versions of Classification is Decision Tree this type of classification takes a set of data and runs it through a set of requirements to help it sort the data into different requirements. The second version of Classification is Discriminant Analysis this type of classification takes a set of data and sorts it based on the differences between them. The third type of Classification is k-Nearest Neighbors (kNN) this type of classification takes a set of data and groups it together based on how similar their value is.

The second most common type of supervised learning is Regression. One type of Regression is Linear Regression which uses a dependent variable and one or more independent variables to order data and use it on a line representing their underlying relationship. The second type of Regression is Nonlinear Regression. This type of Regression follows the same principles as Linear Regression but the underlying relationship in Non-Linear Regression is not a straight line.

The second type of learning that will be discussed in this paper is Unsupervised learning. This type of learning uses many of the same versions as supervised learning but the data is presorted for a human to confirm if the algorithm works correctly. One type of unsupervised learning is clustering which sorts data based on inherent traits. One type of Clustering is hard clustering where all of the data fits into a group completely or it does not. The second type of clustering is soft clustering. With this type of clustering, the data points not all of the data is put into clusters and the data that is not in those clusters is given a probability for which cluster it would be in.

The second type of unsupervised learning is Association. This type of unsupervised learning describes why portions of data are similar. One way of doing this is through the Apriori algorithm. This algorithm sorts through a set of data and creates categories one item at a time. A Second way of doing this is through a Frequent Pattern algorithm which searches through the data for trends and similar traits.

Part III: Reinforcement Learning

The third type of learning that will be discussed in this paper is called Reinforcement Learning. Reinforcement Learning is a type of machine learning that can be either supervised or unsupervised. Reinforcement Learning algorithms can be supervised with a human manually rewarding the Artificial Intelligence (AI) with a point for being correct. The way that a reinforcement algorithm can be unsupervised is by giving the algorithm an automated point system. One type of Reinforcement learning is Brute Force learning this type of learning takes different algorithms and tests them against what the objective is and if it meets the objective. The Second type of Reinforcement learning is Direct Policy Search. This method searches the different algorithms for what can fulfill the objective given to the AI.

Part IV: Conclusion

In conclusion there are many ways to create and use machine learning algorithms. This essay focused on two different methods of teaching machine learning programs. The first part of the first method that was discussed in this essay was supervised learning. This method involves giving the algorithm data and allow it to sort through it on its own. The second part of the first method that was discussed was unsupervised learning. This method involves giving the algorithm two sets of data first a training set to determine how to sort the data and then a validation set to confirm the algorithm can do its job properly. The second method that was discussed in this essay was reinforced learning. With this model, you give the algorithm a set of data and a point system if the algorithm gets things correct it is rewarded with points and if it is incorrect it is punished by losing points.

3.16 Nicole Trenholm

Self-driving cars, Netflix and amazon recommendations, this is what machine learning is. It is feeding information into a computer and the computer learning from those few examples. The goal is to try and get technology to behave like humans in the sense that they should be able to improve their learning over time from the information that they have been given. In order to achieve this goal we must feed the machines enough information that they can use to do tasks and learn from their successes and failures. There are many different concepts that go into the topic of machine learning as a whole. It is about understanding the algorithms and probably apply them. Some algorithms that I focus on are logistic regression, learning vector quantization, and boosting. Each of these topics serves a different purpose and preforms tasks that make machine learning as essential as it is now.

Logistic regression is a technique in machine learning that is used for binary classification problems, which is where there are elements of a given set being classified into two groups based on specific rules. It is similar to linear regression in that the goal of it is to find the values for the coefficients that weight each input variable however, the prediction for the output is transformed using a non-linear function called the logistic function. The logistic function, when graphed, is shaped like an S, with horizontal asymptotes at 0 and 1 and an x-intercept at 0.5. This function, also known as the sigmoid function, is useful since it will transform any value into the range 0 to 1 which can be applied to the output of the function and help predict a class value. So if the output is above 0.5, it will then be 1 and if it below 0.5 the output will then be classified as 0. This can be applied to probability, for example, if you are looking at your favorite soccer teams stats and you get an output of 0.72, which means there is a 72 percent change that your team will win. It can also be useful when you have a problem that you need to give a more rational prediction because the predictions made from this function will not be a random, decimal but instead it will be either 1 or 0. With this being a concept borrowed from statistics, the math behind it can get complicated, however it is used in predication of simple this or that cases. Examples of these predictions are deciding whether an email is spam or not or if a tumor is malignant or benign.

Learning Vector Quantization, LVQ for short, is an artificial neural network algorithm which allows the user to choose the number of training instances to hang onto and that is what the machine learns from. This algorithm supports both two-class or binary classification problems and multi-class classification problems. LVQ is represented by a collection of codebook vectors which is a list of numbers that have the same input and output characteristics as the data entered to train the machine. The codebook can be explained as each vector being a neuron and each characteristic of the codebook vector is a weight and the collection of these vectors are a network. Using an LVQ model is similar to K-nearest neighbor in that predictions are made for a new instance by searching for

the most similar instance throughout all the codebook vectors and summarizing the output for those instances. The way the LVQ algorithm learns the codebook vectors can be thought of as similar to the concept of gradient descent. This is because it starts with a group of random codebook vectors, they are then processed one at a time, if the output is similar to the output of the training instance, it moves closer to the training instance, if it does not match, it moves farther away. This algorithm has many positive features because with machine learning it makes many things in everyday life possible, such as medical image storage and satellite image storage and transmission.

Boosting algorithms are what grant machine learning models the power to be more precise in their predictions. The term boosting refers to a group of algorithms that convert weak learner to strong learner. In cases such as with identifying emails as spam there are a set of rules that are followed to group the emails however, the rules individually are not strong enough to classify them. These rules are then labeled as weak learner and to convert them to strong learner programmers use methods such as using average, weighted average and considering prediction has higher vote. With the help of these methods the weak can become stronger and collectively classify the emails in a more accurate way. This is just one use of the boosting algorithm and with there being three types of the algorithm, there are many more. The types of algorithm that branch off of boosting are AdaBoost, short for Adaptive Boosting, Gradient Tree Boosting, and XGBoost. All these algorithms are able to be used on Python with their own codes and correlate with other aspects of machine learning. AdaBoost for example, are decisions trees with a single split that are called decision stumps. This can be used for both classification and regression problems.

The concept of machine learning is one that should be understood for all the amazing tasks it can perform. With information on specific algorithms from papers like this one or articles that go in more depth, it can be an easy thing to comprehend. There are many advances that are constantly being made and many improvements that follow in fields such as medical, automotive, and military. These are all made possible by having machines that are able to store the data they are given and use it to learn from and improve their performance. The idea of machine learning may seem too complex to comprehend without a background in computer science however, breaking it apart and getting a grasp of these singular algorithms, you can go very far in the field.

3.17 Maddy Weaver

3.17.1 Random Forest

Before we learn about Random Forest, we must first know what a decision tree is. So what is a decision tree? A decision tree is a decision support tool that uses a tree-like graph or model of decisions and their possible consequences, including chance event outcomes, resource costs, and utility. It is one way to display an algorithm. Supervised learning is defined as a machine learning task of learning a function that maps an input to an output based on example input-output pairs, used for both classification and regression problems. The decision tree represents a flow chart-like figure and maps out the outcomes of certain events based on different circumstances.

We should also know what a classifier/classification is. Classification is the problem of identifying to which a set of categories a new observation belongs, on the basis of a training set of data containing observations whose category membership is known. Basically, it is what the decision tree is trying to figure out based on what it already knows.

What is regression? Regression is a statistical approach to find the relationship between variables. In machine learning, this is used to predict the outcome of an event based on the relationship between variables.

Random forest is a very popular algorithm for machine learning. It is a type of ensemble machine learning algorithm called Bootstrap Aggregation or bagging.

What is an ensemble machine learning algorithm? Ensemble methods combine several models which helps improve machine learning results. Decision trees help to explain this best. Ensembles are a divide and conquer approach, used to improve performance. A group of weak learners can come together to create strong learners! Each classifier is a weak learner, while many classifiers together are a strong learner. Each of them look at different attributes and come together.

The bootstrap is a statistical method that is used to estimate a quantity from a data sample, such as a mean/average. You take lots of samples, calculate the mean, then average all of your mean values to give a better estimation of the true mean value. Boosting reduces bias and also variance.

The same approach is used in bagging but rather for estimating entire statistical models instead of samples, most commonly decision trees. Various samples of your training data (training data is used to train an algorithm) are taken and then models are made for each data sample. When you need to make a prediction for new data, each model makes a prediction and the predictions are averaged to give a better estimate of the true output value. Bagging is designed to improve

the stability, and reduce variance and accuracy.

Random Forest is a variation on this approach where decision trees are created so that rather than selecting optimal split points, suboptimal splits are made by introducing randomness. Combining their prediction results in a better estimate of the true underlying output value.

Think of a random forest literally, like a bunch of trees. These trees are decision trees and the more trees that there are, the more accurate the outcome will be. There are multiple trees rather than one. The trees each come to one classification as their outcome. It is like the trees are casting their votes on an outcome to choose a winner. This would be the average of all of the trees.

What are advantages of random forests?

- Both classification and the regression task
- Handle the missing values and maintains accuracy for missing data.
- Won't overfit the model
- Handle large data set with higher dimensionality

Disadvantages?

- Good for classification, not as good for regression
- You have very little control on what the model does

Applications of random forests? You can use them in banking centers to find loyal customers based on different attributes. You can also use it in medicine to see what is best for certain patients by looking at their medical records. In the stock market, random forests are use to find out the stock behaviors.

Random Forest Pseudocode:

- Assume number of cases in the training set is N. Then, sample of these N cases is taken at random but with replacement.
- If there are M input variables or features, a number m¡M is specified such that at each node, m variables are selected at random out of the M. The best split on these m is used to split the node. The value of m is held constant while we grow the forest. Each tree is grown to the largest extent possible and there is no pruning. Pruning is defined as: a technique in machine learning that reduces the size of decision trees by removing sections of the tree that provide little power to classify instances. Pruning reduces the complexity of the final classifier, and hence improves predictive accuracy by the reduction of overfitting.

- Predict new data by aggregating the predictions of the n tree trees. (Example: majority votes for classification, average for regression).

Random tree is a very important algorithm for machine learning and very useful! Although it is a complex method, the more you know the more helpful it will be.

3.18 Peter Weber

3.18.1 Neural Networks

An artificial neural network is a method of computing that is based on the way that brains work. Artificial neural networks are, as the name implies, a network of interconnected neurons. A neuron in these artificial networks can be described as something that holds a number. The complexity of neural networks can vary, but this article describes multilayer perception neural networks.

A multilayer perception neural network is made up of layers of neurons. Each neuron holds a number, or activation. The activation of the neuron is determined by a function of the activation of neurons in the previous layer. Each neuron takes in information from other neurons and uses weights and biases to determine what its activation will be. The weights tell the neuron how much of a positive or negative effect each neuron that it is checking should have on its own activation. The bias tells the neuron what minimum value it should have to have any activation at all. The weights and biases for every connection between every neuron must be programmed before the neural network can function. Even in simple neural networks there can be thousands of connections, and machine learning is a solution to this problem.

The diagram below is of a neural network. The neurons are setup in layers, the first layer being the input and the final layer being the output. The middle layers of the network, called the hidden layers, are where most of the processing is done.

The largest layer on the left is where data comes in. The layer on the right is where data goes out. Suppose a neural network is used to identify different types of fruit. Each neuron on the input layer would be given an activation that corresponds to a certain attribute of the fruit. The neurons have an analog activation rather than a digital one (activation can be anything, not just 0 or 1), so inputs 1-3 could be the red blue and green values of the fruit, for example, and the fourth could be weight. Or the inputs could include the radius of the fruit and how far it is from being a sphere, there are many possibilities. The next layer's job is to begin to recognise patterns in the first layer. In the color example the next layer might have a neuron that would activate if the fruit was orange it would do this by having more weight on the red and green neurons from the previous layer and having low or negative weight on the blue neuron. This way if the red and green values are high and the blue is not the neuron will activate. There can be multiple hidden layers in a neural network but this diagram only has one. The output layer would then take information from the previous layer, perhaps one neuron tells that the fruit is green and another one tells that the fruit is big, and make a prediction on what the fruit is. The output layer can also have multiple neurons in it, for example a watermelon neuron and an apple neuron and an orange neuron, each one having an activation that is

the probability that it is the fruit.

The number of connections in a neural network is so massive that it would be impossible for humans to decide the weights and biases of each neuron. This is where machine learning comes in. Backpropagation is a way that artificial neural networks learn. Back propagation compares the actual result with the expected result and works backwards, noticing what connections it should have made and comparing them to the ones it did. This gives the direction that weights and biases should be adjusted. Starting with random weights an biases and repeatedly doing this technique and making adjustments is a method of gradient descent to make the network more accurate. Every time that the network does this, it improves its results. Using this technique can work well and is much faster than programming and tweaking every weight and bias by hand, but it makes it unclear what the program is actually doing, and when things do not work it can be hard to figure out why.

Sources:

- 3Blue1Brown *Deep Learning* series
- Brandon Rohrer *How Deep Neural Networks Work*
- Diagram: Jake VanderPlas

Chapter 4

Authors: Who we are and how we learn

4.1 Sheldon Branch

Sheldon Branch CSC 131 8/27/18 An experience that gave me a sense of accomplishment was taking piano lessons from my second piano instructor, David Felton. Davids teaching style was very different from that of my previous piano instructor, Tami Moates. Tami was very gentle and taught at a slow pace, whereas David taught at a rigorous pace and expected his students to put in enough effort to keep up. Because I was used to a much slower pace, taking lessons from David was a great challenge. I often considered looking for a new piano instructor but there was always a part of me that enjoyed the challenge. Eventually, I adjusted to the newer, faster pace, which allowed me to greatly improve my piano skills. Adapting to Davids pace was very fulfilling to me and left me with a feeling of pride. I was happy that my hard work paid off and I could see the results first hand. This experience also taught me about the learning environment for which I am best suited. I learn best when there is little leeway given because I am a chronic procrastinator. My work ethic is not as polished as I would like and I feel that a disciplined schedule would greatly assist in improving it.

4.2 Bram Dedrick

From my first year of elementary school, until my sophomore year of high school, I was constantly stuck in a state of indecision when it came to my academic future. I played a bunch of instruments, including the guitar, piano, flute, French horn and saxophone. I wanted to be good at each one I played but never actually enjoyed playing them let alone the process it would take to become good at any of them. The same can be said about the sports I played. I played soccer, football, basketball, track and field and golf. I wanted to be good at each of them, I practiced both with the teams, and on my own but still, just with like with the instruments I played, something seemed off and I never loved any of them. In my traditional course load, I was good at pretty much every topic, but was not great at any and did not love any.

In my sophomore year of high school, I took a Computer Science elective and I finally found the passion that I was looking for with all those classes and extracurriculars from the past. My teacher made the class approachable for someone with absolutely no experience, by making our transition to coding easy. We played logic games and used all sorts of programs that helped us with our critical thinking before we ever touched the keyboard. I went to a camp the summer after my sophomore year that taught us some basic knowledge about coding for video games. In every project, I found myself going above and beyond doing extra and spending every free block in my computer science classroom, which eventually lead to me doing a rather large programming project for my senior project and my love of programming.

4.3 Tony Ferenzi

During all four of my high school years I was a part of the Marist marching band, specifically a percussionist. It wasn't entirely a marching band as we would switch to an orchestra during the winter and spring after the football season was over. I played the bass drum during the marching season and that is the biggest drum on the field that is attached to a harness, if you didnt know what a bass drum is. Our band director was probably one of the most unique individuals I ever knew. He could be your worst nightmare and your best friend all in one day. He is really a nice guy, but teenagers will be teenagers and sometimes he has to yell or make us run laps on the track to keep us in order. My band director actually had multiple ways of teaching and practicing. Not onlydid we practice everyday, but he would make the band do the same formation over and over again until we could do it in our sleep. The most iconic phrase that anyone in the band will tell you after running a formation is the 'go back and do it again' that he would tell us for the twelfth time that day. He would also give us weekly homework assignments to complete, scheduled sectionals every week, and the monthly three hour practice videos we had to complete for a grade so our band director knew we were practicing at home. It may not have been a popular or fun way of learning or practicing, but it definitely got the job done and made us one of the most popular bands in the area.

4.4 William Golden

I have had a wide variety of learning experiences outside of high school classrooms. My biggest learning experience took place between the summer of 2016 and summer 2018. Shortly after graduating high school in Des Moines, Iowa, I signed up to become a missionary for a religious group in South Korea. The requirements for becoming a missionary were that I must speak proficient Korean and must know and teach many church doctrines fluently. My training began in Provo, Utah, in an institute called the Missionary Training Center (MTC). I spent a vigorous 9 weeks there studying both the Korean language and the sprictures. My schedule was similar throughout the 9 weeks; 6:30am wake up, 8:00am study personally, 9:00am study with teachers, 12:00 pm lunch, 1:00 pm personal study, 2:00 pm study with teacher, 5:00 pm dinner, 6:00 pm practice teach, 8:00 pm study, 9:00pm go to dorms and prepare for bed, 10:30 pm lights out. There was little to no recreational/free-time, but I really felt I learned a lot there because they repeatedly forced me to study in several different ways each day, then practice what I have learned immediately afterwards. After 9 weeks of being in the MTC, I was shipped off to Korea to begin proselytizing in the Korean language to the Korean people. I was still following a very strict schedule of scripture/language study in the morning, then from noon to 9pm I was outside or in a church applying what I learned that morning though forms of teaching and communicating. Quickly learning material then applying it immediately afterwards was an exteremely difficult yet rewarding process. I was able to master the language of Korean and become a scriptural scholar in less than 2 years. However, it was an everyday struggle of speaking and teaching. I made countless mistakes and had to push myself past what I thought I was capable of In the end I was able to grow and learn a lot.

4.5 Yuan Hong

To first think of studying outside the school, I studied calligraphy. Chinese calligraph art began in the stage of the beginning of the Chinese characters. And it is a tradition still toda, that people still write calligraphy for festivals and families. Calligraphy is a interesting field of studying, to just study calligraphy is easy to do, but in fact, the learning of doing good at calligraphy needs a long time of practice. It is very difficult to make the level of calligraphy reach a certain height. To be good at this type of art, there are six types of difficulties I need to solve: difficulty of using the pen, difficulty in the structure, difficulty in the ink method, difficulty in the chapter, difficulty in the mood, and at last, difficulty in the unity. After accomplishing the basics of this art, I was told by my teacher that I need to imitate how these great calligrapher wrote before. At first, even imitating was hard for me, I could barely imitate these calligrapher's styles. There are many trails and errors and I have to stand at a table for hours everyday, and I feel pain from my legs and shoulder, it was a tough time. However, from that time on, I learned how to do something quietly for a long time, even though it was boring, I could endure this boredness and complete the work. Once I got better after years of practice, I can finally move on, to do things by my own styles, but not imitating other people's work. It was a totaly different experience, I have to think of different styles and words to writing, and to be better from my friends who studied calligraphy with me. After that, every Chinese Spring Festival, I write calligraph arts for my family and my friends, and they all liked them a lot. It was a good piece of memory.

4.6 Easton Jensen

A little bit about me.

I have come to experience and learn from family friends. for this story they are mormon while I happen to be catholic. However I went to church with them for the last half of my junior year. After church I would go to their house for lunch. I learned a lot from them by just being near them in church and also going to their house to eat and mess around. They taught me not to look through a window and see how they live. It showed me to not judge people on their religion or how they live day by day. Another person I look up to is my football coach. He has taught me the importance of relationships and brotherhoods. My best friends are all on my football team, even though with a small school all my friends played on every sports team, and they have shown me that people need to be in your life so u can pick people up when they are down because they will do the same for you. One story happened in my senior when a fellow senior tore his achilles and spent the whole season not being able to play. He sat out all camp and had to watch on the sidelines. That year we all came together as a family and performed for him. We took it into playoffs but ended up losing in the semifinals. That year however we all played for one reason and I want to continue that with this school and people I see around campus.

4.7 Rodrigo Martinez

Rodrigo Martinez Beauty and Joy of Computing 8/27/18

Throughout my academic years I had various types of learning experiences. One example is when I took orchestra in high school. The way of learning was very different compared to your traditional classes in which can include some reading, writing, etc. I played the violin at the time and was taught how to play and read notes. You only begin to actually learn how orchestra works when you begin to play the instrument. Eventually as time passed I became invested into the world of computer science. I decided to take an after school class on coding to understand how it works. From experience, in order to learn and understand how coding works you will need to do it Hands-on. Reading a book on coding will not help as much. It is more of a learn from your mistakes method type of learning. I wanted to understand more of coding so I began to learn by myself by looking for sources that help how to code such as w3schools.com and videos. I implemented the things that I learned into a small video game project that I am creating. This helped me even more to understand how coding can be used and how it works. Going through the task of figuring out what causes the script to fail is what brings the learning process together. It helps validate what can work and what you should remember in coding. Even if it means memorizing to put a simple bracket in one of the lines of the script. In my experience, the best way to actually learn from something is by fixing the mistakes made along the way. I dealt with this in my orchestra class as well as my computer science classes I took. Yes, there may be alternatives to learn these things but learning it this way will improve yourself academically/socially/etc.

4.8 Matt Morrical

In high school in Mount Vernon, Iowa there was and still are rather limited opportunities to be able to do interesting and innovative things outside or inside the class due to its smaller size. That is where I went to hisghschool, but there is the Robotics Club. That was the single greatest thing that I did in all of my time in high school. In the club we got to do a great many thing s to help us learn how to build the robots we needed and eventually the ones we wanted, with programming to follow in the same way. The thing that I got from this is that machines are the future and one would be well advised to go into an appropriate field to prosper.

4.9 Ella Nelson

0;256;0cIn sixth grade, I joined competitive math, easily the most nerdy thing I could say about myself. Then again, it was even more nerdy when I started winning all the time. My Junior year, I made a small computer game mostly for my own enjoyment, but my librarian convinced me I should submit it to a contest. I won somehow and ended up with a thousand dollar grant and a trip to DC. In high school I also participated in science fair three out of four years and continued my math team victories. Then one day I decided I wanted to go to college and now this is the one I'm going to.

4.10 Koichi Okazaki

Koichi Okazaki

CSC 131 Writing Assignment

Speaking of my athletic experience, I used to play soft tennis -which is one kind of tennis originated in Japan and popular among Asian countries-for three years in my junior high school. Usually in japan we have a sports camp which is held by the school in each summer, and every athletic clubs participates the camp for about 1 week. These camps are aimed to reinforce skills such as teamwork, technique, physical strength and mental strength. Thus, the practice during the camp tend to be harsh, since we have to practice for long time (about 9 hours) in the very high temperature environment. The content of the practice starts form the fundamental training such as running, and it gets to be more advanced training and eventually we do some matches. According to the result of the daily match coach decides the regular members of the team. Therefore, the mental stress for students will be large. However, since we really concentrated on our spots for a week and did hard work with basic practice and by having the chance to apply it. I felt my technique, physical, and mental strength improved steeply than practicing for one week in usual way, and overcoming the difficulty with my teammates brought me the improvement of coordination.

4.11 Jakob Orel

My past learning has been shaped by my teachers and my interests. I have participated in multiple online courses, athletic teams, band, and academic team. In my online courses, I learned independently about psychology and the foundations of information technology. The classes required discussion on an online forum with other classmates and daily reflections or quizzes. These applied concepts and quizzes would relate back to the ideas in the textbook. Being a member of the cross country team in high school allowed me to learn more about myself and my limits. I worked closely with the coach and was motivated to train and lift weights outside of mandatory practice. Participating in band allowed me to develop as a leader and hone my skills. Playing trumpet in the highly ranked jazz band required me to practice and learn the instrument on my own. I had to listen to other musicians to further develop my skills. I was able to change the style I play trumpet in order to achieve a better sound while still using the scales I had previously known. As a member of the academic team in high school, I had to attend practices to learn the material. Success in the quiz bowls required all members to have a basic knowledge of most topics. My band teacher, science teacher, and my coach have influenced me to learn in a way I can build on my past knowledge or skills in order to become a better person.

4.12 Marcellus Parks

Marcellus Parks 8/28/2018 CSC131 My most successful learning experiences have been where I have been immersed in an environment. I have been to many football camps and coaches clinics. These taught me the successful concepts and strategies of the game. Most of theses things are useless until I know completely how to use them. Its trial and error way of learning helps one pick up things up quickly. It constantly test your knowledge and reactions. I am also an amaeteur musician. I like to teaching myself to play the guitar, and piano. It constantly causes me to use my prior knowledge of chords and melodies to create some new songs. I used to be a part of a band, playing the clarinet. Learning music was a little more difficult then, because I did not have the luxury of being able to fail and try again all the time. There were concerts and performances only a short time after receiving new material to play. I have taken a few college courses I do not have credit for. This was actually easier than other classes sometimes. They allowed me to learn in the way that best suited me and gave me pleny of time to do it. The class was over an entire school year and not just a semester like in college. Even after the seemingly easy learning the class still sent me into a workshop where I had to use the knowledge that I had just gained to make something on a CNC machine, 3D printer, etc.

4.13 Lydia Sanchez

I did not know this at the time, but, Girl Scouts would really change my life. When I joined in second grade, I thought this would be just a fun activity to do. What I did not know was that I would still be a Girl Scout until the end of my senior year in high school. Our leaders recently revealed that they had a master plan to get us to go on a big trip at the end of our 5th grade year, when we were in second grade. They prepared us by taking us camping and teaching us to be independent, but they always did it in the most fun way. Until one meeting our leaders brought up the idea of a big trip during the summer after fifth grade to Glacier National Park. We set off the first of August to Glacier. We drove up to Minneapolis and got on a one and a half day trip by train to Glacier. When we got there we dropped our stuff off in our room and explored a little before dinner. We went down to Lake McDonald,it was right next to where we were staying,and it was a unusual lake. What made it so unusual was that was the lake was two different temperatures split in the middle because of the different sources of water coming into the lake, one warm from a river and one cold from the glacier. The next days were filled with hikes, huckelberry products, jammer buses and a little community service, but the best thing was the snowball fight we had on the logan pass. We went up the mountain to logan pass and it was sunny and about 60 degrees, however there was snow on the ground. We hiked a little and then on the way back we had a snowball fight. This was so memorable because how many people can say that they have a snowball fight in August. This trip was the first of many, as later on we went to the Boundary Waters, the Black Hills, and many other small camping trips. Girl Scouts has taught me many things such as leadership, independence and a love of travel. Girl Scouts has taught and changed me so much. I would not be here without it.

4.14 Tiff Serra-Pichardo

During ninth grade I attended a trade school the second half of the school day. During my time at this trade school I studied electronis engineering and nanofabrication. This class gave me my first experience working with circuits and general engineering. I was able to practice my soldering skills and work at creating simple small electronics. The class allowed me to explore my interest in electronics engineering. It gave me the opportunity to create a life long passion for creating electronics and soldering.

4.15 Austin Stala

Over my career as a student, I have had many teachers and instructors. In most of my experiences with their teaching, I felt that I had no remarkable experiences with them. Although I had good experiences with most of them none of these experiences were particularly remarkable. Almost every experience that I have had with teaching has been the normal techniques and normal lessons. Despite this, I would say that I have had memorable learning experiences outside of and classes or camps. I would say that the most memorable and interesting learning experiences have been in museums and with YouTube channels. These are the most memorable learning experiences for me because of the combination of visual and audio that help me stay engaged and that these experiences have no rigid structure. The passion that behind these experiences also help me stay engaged in these learning experiences because the presenter's excitement and interest about what they are talking about helps keep me engaged. One specific experience that comes to mind is my learning about WW1 through the YouTube channel The Great War. In my career as a student, I have not learned or remembered much about the first world war. My experience with The Great War channel has had the opposite effect with the distanced but regular uploads and the enthusiasm of the host Indy Neidell has helped me learn and remember much about the First World War another events in the interwar period.

4.16 Nicole Trenholm

Growing up I was never the biggest follower of music. I enjoyed listening to it but I never made a habit of listening to it. That was until I started high school. I started listening to more and more music and now I really enjoy finding new songs and artists. That being said I am in no sense musically inclined. The only instrument I have ever played was the recorder which I did not exactly excel in. Sports was always more of a focus for me. All my life I have been playing soccer and I am now fortunate enough to be playing it here at Cornell. Throughout my life I have gone to many camps and tournaments through soccer, either with my school team or with my travel team. It always acted as an escape for me and allowed me to express myslef while also learning new life skills. The sport has taught me about teamwork, leadership, accountability, and persistence. Through soccer I also participated in strengthening activites that included lifting weights and agility drills. These helped me better prepare for the soccer season as well as learn more about the body and how I can keep myslef in the best shape possible.

With soccer being such a big part of my life, it is where I met some of the most influential people. One being my middle school soccer coah who was also my teacher in eighth grade. He helped me improve myself on and off the field. I learned lessons from him that gave me a new perspective in life. From him I began to understand that I needed to have condifedence in myslef because that only opinion of you that matters is your own. Participating in a team with him being the coach I grew so much as a person and I gained many skills that I use both on the field and in the classroom.

Soccer has also given me many opportunities to help others and meet new people. In one way being that I was able to become a referee for youth soccer teams and act as a mentor for the kids that wanted to learn more about the sport and how to play it. I also had chances to participate in volunteer work with my team. We took part in activites such as helping out at soup kitchens, collecting money and toys for children in need, and showing appreciation to the leaders of our community such as the firefighters and police officers. Through these activites I learned how to better myself as well as how to better the area around me with the help of others.

One of the most memorable activites we did is when as a team we participated in random acts of kindness wewk. This was a time when we, as an organiziation, got together and thought of small tasks that we could do to better at least one persons week. We did small acts such as complimenting strangers or helping someone carry their groceries. There was also larger activites we did such as going to a local soup kitchen and interacting with the people there and learning about there stories and experience. This gave me a new appreciation of my own life and all the people I am fortunate enough to have supporting me. The team as a whole found what we did as a life changing event and so we countinued to find new ways to help others while learning about how we can improve ourselves.

4.17 Maddy Weaver

During my sophomore year of highschool, I created a club with my friends called Cinema Club. One day after school, my friends and I were taqlking to our English teacher, Mr. Mclain, who was also our class advisor. We were talking to him about how there were so many things in the world that we didn't know and weren't taught in school. Somewhere in the midst of this conversation, we had the idea: a club were we watched movies that were important, had important messages, and generally impacted the way we thought about different things. At some point, we decided on the name Cinema Club and actually got this club started. This was a huge learning experience for me not only in terms of what it takes to start a club, but also because the content in the movies we watched was very valuable. Some of my favorite movies we watched include Short Term 12 and The Fundamentals of Caring. Along with watching weekly movies, we also attended the Beloit Internation Film Festival yearly in Wisconsin. That experience has been so meaningful for me because we watched both local and internation movies and got talk with many directors and actors/actresses, some of which have starred in/directed famous Hollywood movies.

4.18 Peter Weber

4.18.1 Best learning experiences

One of my best experiences in learning is when I went on a hiking trip during the summer of 2015. I went with my boy scout troop to New Mexico for a ten day backpacking trip. The trip was fun but very hard and we had very little time to do anything other than hiking or sleeping. Two years later, in 2017 we went again and this time we were better prepared for what we would face. We got up earlier in the morning and had practiced packing and unpacking so we did not waste time in the morning or afternoon setting up or taking down our camp. Because we were so much more efficient with our time, we had plenty of it left over to get extra rest or have fun. Because we spent less time working during the day, we also had more energy when we were done hiking. The extra energy would go into whatever we wanted or would be saved for the next day of hiking. The second trip was much more enjoyable and was much easier to do than the first one. The thing that made the difference is that we were much more prepared for the second trip. This taught me that if you are ready for it, you can face anything. Another thing that taught me this lesson is when I went camping in the winter. We slept outside on a night where the temperature was -20^o F and the windchill was -40^o. We were able to survive, and actually be comfortable, because we had what we needed to stay warm. This also taught me that whatever you are going to face, weather it is camping outside in the cold or going to a new school, it can be easy if you are prepared for what is going to come.

Bibliography

[1] Laura Acion, Diana Kelmansky, Mark van der Laan, Ethan Sahker, DeShauna Jones, and Stephan Arndt. Use of a machine learning framework to predict substance use disorder treatment success. *PLoS ONE*, pages 1–14, APR 2017.
(Austin Stala) This article talks about using AI and machine learning to help predict if treatment for substance disorder will be successful. The article states that the are m,any ways that machine learning can be used to predict the chances of success. One way that the article proposes using machine learning to predict the chance of success is by using data from other patients to compare to the current patient.

[2] Musatafa Abbas Abbood Albadr. Spoken language identification based on the enhanced self-adjusting extreme learning machine approach. *PLoS ONE*, pages 1–27, April 2018.
(Yuan Hong) LID, based on the literature, process features on i-vector based framework. The learning based on it is not efficient enough. So they build a new learning model, which is a learning machine. Although the new learning machine is not efficient enough, it is still one point two- five percent more accurate than the old one did.

[3] Gopnik alison. Making ai more human. *Scientific American*, pages 60–65, June 2017.
(William Golden) Teaching computers how to learn by anylyzing children's thinking processes and creating algorithyms.

[4] Mark Anderson. Machine learning and quantom computing become bffs. *IEEE Spectrum*, page 14, August 2017.
(Bram Dedrick).

[5] Jamais Ascio. In machine we trust. *New Scientist*, pages 26–27, 2016.
(Koichi Okazaki).

[6] Gregory P. Asner. Rapid forest carbon assessments of oceanic islands: a case study of the hawaiian archipelago. *Carbon Balance & Management*, pages 1–13, January 2016.
(Tony Ferenzi) Aboveground carbon stock assessments have been growing

in need the past years as the statistics help make policies that can reduce carbon emissions from deforestation. Yet a lot of oceanic islands have been left out during the data collection. So researchers used a combination of airborne Light Detection and Ranging and ground mapping to create a very detailed map of carbon stocks, like Hawaii and other islands.

[7] Rutkin aviva. Robots: just look at me and copy what i do. *New Scientist*, page 23, October 2016.
(William Golden) Using motion tracking from humans to teach computer robots to perform certain functions.

[8] Yoshua Bengio. Machines who lean. *Scientific American*, pages 46–51, June 2016.
(Bram Dedrick).

[9] Yoshua Bengio. Machines who learn. *Scientific American*, pages 46–51, June 2016.
(Sheldon Branch) This article discusses the history of deep learning and the progress that has been made in the field since then.

[10] Yoshua Bengio. Making who learn. *Scientific American*, pages 46–46, June 2016.
(Tiff Serra-Pichardo).

[11] Thania Benios. At the edge of life's code. *Scientific American*, pages 106–109, April 2008.
(Jakob Orel).

[12] Rachel Berkowitz. #flu. *Scientific American*, page 23, April 2018.
(Jakob Orel).

[13] Reddy Bhaskar and E. Madhusudhana. Able machine learning method for classifing disease treatment sematics relations from bio-mediacal sentences. *International Journal*, pages 223–226, march 2018.
(Lydia Sanchez).

[14] Dipankar Bhattacharya. Relevance vector-machine-based solar cell model. *International Journal of Sustainable Energy*, pages 685–692, November 2015.
(Tony Ferenzi) Machine learning has had some major developments over the past few years. There are multiple techniques for machine learning to use and adapt to, like artificial neural network or support vector machine. Each come with there own benefits and drawbacks. Yet one doesn't suffer from the downsides of others, called RVM, which is a Bayesian treatment. It is so effective that it is used for solar cells because it reduces an incredible amount of error.

[15] Bartosz Binias, Darius Myszor, and Krzysztof A Cyran. A machine learning approach to the detection of pilot's reaction to unexpected events

based on EEG signals. *Computational Intelligence Neuroscience*, pages 1–8, April 2018.
(Nicole Trenholm).

[16] Maria Bisele. Optimisation of a machine learning algorithm in human locomotion using principal conponent and discriminant function analyses. *PLoS ONE*, pages 1–19, September 2017.
(Yuan Hong) Human Locomotion needs graphic profiles and discrete variables. People could not present the full size of the data. So the participants developed a self-learning machine algorithm for getting the data. It successfully predicts by ninty-three precent correct. And to the author, it's the first study to optimise the algorithm.

[17] Bianca Bosker. The mayo mogul. *Atlantic*, pages 76–92, November 2017.
(Jakob Orel).

[18] Anne-Laure Boulesteix. A plea for neutral comparison studies in computational sciences. *PLoS ONE*, pages 1–11, April 2013.
(Yuan Hong) There are some new methods, when studies are written poor by the journalists. Survey shows the seven high ranked science journal, the purpose is the compare the new the methods focusing on the study itself. There are three criteria, that study have to be neutral.

[19] Rodney Brooks. The seven deadly sins of AI predictions. *Technology Review*, pages 79–86, Nov/Dec 2017.
(Austin Stala) There are many problems with peoples predictions about AI and machine learning. Many of these problems the unpredictable speed of advancement in the field of computer science. There are also many misconceptions about what AI and machine learning can do and how fast AI and machine learning will be implemented in the world.

[20] Rodney Brooks. The seven deadly sins of ai predictions. *MIT Technology Review*, pages 79–86, Nov/Dec 2017.
(Sheldon Branch) This article discuess the misunderstanding of the social impacts of AI.

[21] Dayra Chyzhyk, Alexandre Savio, and Manuel Grana. Computer aided diagnosis of schizophrenia on resting state fMRI data by ensembles of ELM. *Neural Networks*, pages 23–33, August 2015.
(Nicole Trenholm).

[22] Pinar Cihan, Erhan Gorkce, and Oya Kaliipsiz. A review of machine learning applications in veterinary field. *Kafkas Universitesi Veteriner Fakultesi Dergisi*, pages 673–680, Jul 2017.
(Austin Stala) This article discusses that many applications of AI machine learning in the veterinary field. One of the mentioned applications of machine learning is using software to help find treatments and identify diseases. The article states that there are many different ways for machine learning to help in the veterinary field.

[23] P DELL'AVERSANA. Application of machine learning and digital music technology to distinguish high and low gas saturated reservoirs. *Bollettino di Goefsica Tecoria ed Applicata*, pages 71–94, March 2018.
(Marcellus Parks).

[24] Bob Dickinson. Algo pop. *Art Monthly*, pages 6–10, Novemeber 2017.
(William Golden) Machine learning/algorythms influencing art and music that are based on computer assited paining and drawing machines.

[25] Benedict Diederich. Using machine-learning to optimize phase contrast in a low-cost cellphone microscope. *PLoS ONE*, pages 1–20, March 2018.
(Marcellus Parks).

[26] Chris Dixon. Monster waves. *Scientific American*, pages 64–67, August 2018.
(William Golden) Teaching computers to predict when large, surf-able waves will apear.

[27] editors. More smarts needed: We still need humans to ensure AIs are up to the job. *New Scientist*, page 5, October 2017.
(Jakob Orel).

[28] editors. Mini footballers on your dining table. *New Scientist*, page 19, June 2018.
(William Golden) Computers learning to create 3D holograms of football games through youtube videos.

[29] Li Fei-Fei. How to make a.i human-friendly. *New York Times*, pages pA27–27, March 2018.
(Maddy Weaver).

[30] Seth Fletcher. Machine learning. *Scientific American*, pages 62–68, August 2013.
(Tiff Serra-Pichardo).

[31] Joseph Frankel. This is your brain on death. *Newsweek Global*, pages 42–43, November 2017.
(William Golden) Computers learning to read depression signals in human brain.

[32] Edward Frenkel. Speeded by computation. *New York Times*, pages pD1–D2, October 2013.
(Maddy Weaver).

[33] Martin Giles. The ganfather: The man whos given machines the gift of imagination. *MIT Technology Review*, pages 48–53, Mar/Apr 2018.
(Sheldon Branch) Ian Goodfellow creates Generative Adversarial Networks (GANs). These GANs has catalyzed the progress of deep learning, allowing AI to recognize faces, drive cars, and power virtual assistants such as Siri and Alexa.

[34] Xiajing Gong. Big data toolsets to pharmacometrics:application of machine learning for time- to-event analysis. *CTS:Clinical & Translational science*, pages 305–311, May 2018.
(Yuan Hong) Add big data tools to pharmacometric problems are new. Models are evaluated based on preset scenarios, That uses different variables, and result shown that machine learning methods are doing much more better than the old method they've used. Also machines aren't sensitive to data size and cencoring rates.

[35] Alison Gopnik. Making AI more human. *Scientific American*, pages 60–65, June 2017.
(Sheldon Branch) This article discusses the mindset of a child. It proposes that AI can learn to think better by analyzing how children think.

[36] Alison Gopnik. Making ai more human. *Scientific American*, pages 60–65, June 2017.
(Tiff Serra-Pichardo).

[37] Samuel Greengard. Gaming machine learning: game simulations are driving improvements in machine learning for autonomous vehicles and other devices. *ACM*, pages 00–00, 2017.
(Easton Jensen).

[38] Ashraf UI Haque. Wind speed forecast model for wind farm based on a hybrid machine learning algorithm. *International Journal of Sustainable Energy*, pages 38–51, January 2015.
(Tony Ferenzi) Wind power is an incredibly useful source of energy that is cleaner and cheaper alternative to fossil fuels. Yet wind is a very challenging force to calculate the speed of and if calculated wrong, it could raise the price of using the system. So a hybrid machine learning system called STWSF was created to efficiently counter this problem. Using a variety of options and tools, STWF reduces the uncertainty in difficult wind patterns. It is still currently being developed to be able to calculate even more possibilities, like humidity and other forms of weather.

[39] Reed Hastings. Trust the algorithm. *New York Times*, pages A1–A4, June 2017.
(Maddy Weaver).

[40] Brian Hayes. Delving into deep learning. *American Scientist*, pages 186–189, May/June 2014.
(Jakob Orel).

[41] Douglas Heaven. The line between hate and banter. *New Scientist*, page 12, October 2017.
(Jakob Orel).

[42] Jeremy Hsu. For sale: deep learning. *IEEE spectrum*, pages 12–13, 2016.
(Easton Jensen).

[43] Cai Huang. Detailed record open source machine-learning algorithms for th prediction of optimal cancer drug therapies. *PLoS ONE*, pages 1–14, October 2017.
(Marcellus Parks).

[44] Nicolas Papernot Ian Goodfellow, Patrick McDaniel. Making machine learning robust against adversarial inputs. *Communications of the ACM*, pages 56–66, July 2018.
(Sheldon Branch) This article suggests that AI will be much more effective if they can learn from their mistakes.

[45] Jose M. Inesta, Darrell Conklin, and Rafael Ramirez. Machine learning and music generation. *Scientific American*, pages 87–91, July 2016.
(Sheldon Branch) This article discusses the possiblity of AI generting music. It presents several subtopics, such as the music representation and the structural coherence of the music through machine learning.

[46] Nicolas Jaccard. Record detection of concealed cars in complex cargo x-ray imagery using deep learning. *Journal of X-Ray & Science Technology*, pages 323–339, March 2017.
(Marcellus Parks).

[47] Dysart Joe. Artificial intelligence: Not your father's toolbox: Some new artificial intelligence business tools to help park and rec agencies. *Parks & Recreation*, pages 72–73, Aug 2017.
(Austin Stala) This article talks about how AI and Machine learning can be used to help with apps and website design. The use of AI can also help monitor website traffic to help cut costs.

[48] Karen-Inge Karstoft, Issac Galatzer-Levy, Alexander Statinikov, Li Zhiguo, and Arien Y Shalev. Bridging a translational gap: Using machine learning to improve the predictions of PTSD. *BMC Psychiatry*, pages 1–7, 2015.
(Nicole Trenholm).

[49] Micheal Kearns and Umesh Virkmar Vazirani. An introdutction to computational learning theory. *The MIT Press*, 1994.
(Nicole Trenholm).

[50] Omar F. Khan. Artificial intelligence in medicine: What oncologists need to know about its potential—and its limitations. *Oncology Exchange*, pages 8–13, November 2017.
(Tiff Serra-Pichardo).

[51] Raffi Khatchadorian. The doomsdaaay invention. *New Yorker*, pages 64–79, November 2015.
(Lydia Sanchez).

[52] Raffi Khatchdourian. The doimsday invention. *New Yorker*, pages 49–64, November 2015.
(Bram Dedrick).

[53] Will Knight. Your driverless ride is arriving. *MIT Technology Review*, pages 34–39, Nov Dec 2016.
(Peter Weber).

[54] Will Knight. Artificial intelligence that doubts itself. *Technology Review*, pages 11–12, Mar/Apr 2018.
(Austin Stala) This article talks about adding the ability for an AI with machine learning to doubt its conclusions. This is believed to be able to help an AI to come to better conclusions. This would happen because the AI would be able to consider other decisions rather than the one it deems most likely.

[55] Kashyap Kompella. Can machine learning help fight fake news? *Econtent*, pages 40–40, 2017.
(Koichi Okazaki).

[56] Kashyap Kompella. Can machine learning help fight fake news? *Econtent*, page 40, Sep/Oct 2017.
(Rodrigo Martinez) Todays social media can be filled with a bunch of fake news. In order to fight back against this we at least should implement machine learning. Although with machine learning we have many possibilities, we also have many limits. Not all bots/AIs can easily detect whether theres a fake article or not. Sometimes it is up to us to figure if something is false or not.

[57] Utku Kose. Are we safe enough in the future of artificial intellegence? *BRAIN:Broad Research in Artificial Intellegence & Neuroscience*, pages 184–197, May 2018.
(Yuan Hong) Nowadays people have anxiety that artificial intellegence may take over humanity. So people ask if they are really safe living in this environment. The author improve the awareness that poeple should notice the sciences around them, that are taking role of mankind.

[58] Koan Koseler and Matthew Stephan. Machine learning applications in baseball: A systematic literature review. *Applied Artificial Intelligence*, pages 1–8, february 2018.
(Nicole Trenholm).

[59] Marina Krakovsky. Accelerating search. *Communications of the ACM*, pages 15–16, July 2016.
(Peter Weber).

[60] Diana Kwon. Self-taught robots. *Scientific American*, pages 26–31, March 2018.
(Tiff Serra-Pichardo).

[61] Steve Lohr. A profitable idea from a victory on 'jeopardy'! *New York Times*, page P1, June 2012.
(Maddy Weaver).

[62] Susana Lopez-Aparicio. Webcrawling and machine learning as a new approach for the spatial distribution of atmospheric emissions. *PLoS ONE*, pages 1–15, July 2018.
(Yuan Hong) There are two ways to collect data in atmospheric science. It is used collect information that used as emission inventory. Which for residential Wood Conbustion. One way is to extract online data systematically. The second way is a model trained that could learn on its own by recongnizing images. Increased amount of data, add significant value to the application of emission sectors.

[63] Robert Lucky. Are engineers designing robotic replacements? *IEEE Spectrum*, page 27, May 2016.
(Rodrigo Martinez) Robots are designed to make our lives easier or at least simplify some work for us. However, with technological advancement Robots are getting smarter. Is it possible for robots to replace humans in jobs. Frey and Osbornes information shows that it is a possible thing to occur.

[64] Robert Lucky. Are engineers designing their robotic replacements? *IEEE spectrum*, page 27, 2016.
(Easton Jensen).

[65] Illias G. Maglogiannis. Emerging artificial intelligence applications in computer engineering: Real word AI system with applications in ehealth, HCI, infomation retrieval and pervasive technology. *The IOS Press*, 2007.
(Nicole Trenholm).

[66] Farhad Manjoo. Making camers eerily smart. *New York Times*, pages PB1–B4, February 2018.
(Maddy Weaver).

[67] Johm Markoff and Paul Mozur. Program knows just how you feel. *New York Times*, pages pD1–D2, August 2015.
(Maddy Weaver).

[68] John Markoff. How many computers to identify a cat? 16,000. *New York Times*, page P1, June 2012.
(Maddy Weaver).

[69] John Markoff. A learning advance in artificial intelligence rivals human abilities. *New York Times*, pages b3–b3, december 2015.
(Nicole Trenholm).

[70] Jane Mayer. Rules of play. *New Yorker*, pages 14–15, July 2018.
(Lydia Sanchez).

[71] Luis J Mena. Mobile personal health monitoring for automated classification of electrocardiogram signals in elderly. *Computational & Mathematical Methods in Medicine*, pages 1–9, May 2018.
(Marcellus Parks).

[72] Cade Metz. This a.i. can build a.i. itself. *New York Times*, pages 00–00, November 2017.
(Bram Dedrick).

[73] Cade Metz. Google researchers say they're learing how machines learn. *New York Times*, pages 00–00, March 2018.
(Bram Dedrick).

[74] Christopher Mims. Why ads are the best thing to happen to the internet. *Wall Street Journal*, pages B1–B4, July 2015.
(Lydia Sanchez).

[75] Christopher Mims. Should artificial intelligence copy the human brain? *Wall Street Journal*, pages 00–00, August 2018.
(Bram Dedrick).

[76] Kakuko Miyata, Hitoshi Yamamoto, and Yuki Ogawa. What affects the spiral of silence and the hard core on twitter? an analysis of the nuclear power issue in japan. *American Behavioral Scientist*, pages 1129–1141, JAug 2015.
(Sheldon Branch) This article discusses the social impact of Twitter on public opinion.

[77] Don Monroe. Chips for artificial intelligence: companies are racing to developing hardware that more directly empowers deep learning. *Communications of the ACM*, pages 15–17, April 2018.
(Rodrigo Martinez) Companies are trying to develop new technology hardware for AI that is efficient enough. These companies are taking different approaches in order to achieve their goal. Some companies are working together to reach their goals.

[78] Michael Morisy. How paypal boasts sercurity with artifical intelligent. *Technology review*, pages 73–74, march/april 2016.
(Lydia Sanchez).

[79] Siddhartha Mukherjee. The algorithm will see you now. *New Yorker*, pages 46–53, Apr 2017.
(Sheldon Branch) The article explores the potential for algorithms and learning machines in medical diagnosis and diagnostic techniques.

[80] Sidhartha Mukherjee. The algorithm will see you now. *New Yorker*, pages 46–53, April 2017.
(Bram Dedrick).

[81] Jurriath-Azmathi Mumith. Design and optimization of a thermoacoustic heat engine using reinforcement learning. *International Journal of Low Carbon Technologies*, pages 431–439, September 2016.
(Tony Ferenzi) A thermoacoustic heat engine is a very unique and desirable engine as it is simple and efficient. It is useful in that it can power things like refrigerators and in general can produce electtricity for any form of machine. The engine is actually being created using a type of machine learning technique called reinforcement learning. This type of learning rules out flaws and improves upon working parts to maximize energy output. Reinforcement learning is simple in that it specifies whether something is good or bad amd then rewards it for the good.

[82] N/A. For the military,AI and 3-D might make for easier maintenance: Army uses machine learning to monitor bradley fighting vehicles; marines quickly fix aircraft. *ISE: Industrial & Systems Engineering at Work*, page 17, Aug 2018.
(Austin Stala) This article discusses applications AI and machine learning by the military. One of the main applications by the military is to monitor vehicles and help with maintenance.

[83] Brian Nadel. AI accelerates: Machine learning will offer teachers deeper insight, but not replace them. *District Administration*, pages 55–58, 2018.
(Koichi Okazaki).

[84] Gautam Narula. Machine learning in gaming- building ais to conquer virtual worlds. *Techemergence*, page N/A, September 2017.
(Rodrigo Martinez) People are implementing machine learning to various things such as the medical field. People are now using machine learning to create self teaching AIs to play games. An example is a former Microsoft engineer named, Seth Hendrickson created a self teaching AI and made it play through a level in a Mario game.

[85] Neville Neil. The chess player who couldn't pass the salt. *Communications of the ACM*, pages 24–25, 2017.
(Koichi Okazaki).

[86] Floyd Norris. Inside the workings of a money machine. *New York Times*, pages pC1–C3, October 2003.
(Maddy Weaver).

[87] Yuji Nozaki. Predictive modeling for odor character of a chemical using machine learning combined with natrual language processing. *PLoS ONE*, pages 1–13, June 2018.
(Yuan Hong) Technology are doing successful performances in different tasks. Now it utilizes mass spectra data which performed by human language. The sensory includes smell and taste. However, a large scale of sensory is not able to complete or perform. But the experienment shows

us how amazing this technology could analyze. This the first one to perform something like that.

[88] ole marius hoel rindal, trinem seeberg, johannes tjonnas, pal haugnes, and oyvind sandbakk. Automatic classifcation of sub-tecnniques in classical cross country skiing using a machine learning algorithm on micro-sensor data. *sensors*, pages 1–15, january 2018.
(Lydia Sanchez).

[89] Sung Kyun Park. Construction of environmental risk score beyond standard linear models using machine learning methods: application to metal mixtures, oxidative stress and cardiovascular disease in nhanes. *Environmental Health: A Global Access Science Source*, pages 1–15, September 2017.
(Tony Ferenzi) Pollution has always been a huge problem over the past couple of years. Though steps have been taken to research and counter pollution health effects, not much is known about what health problems can be caused by multiple pollutants mixed together in an area. Researchers came up with a system known as ERS to predict and monitor these specific pollutants, along with the health problems that come with them. Machine learning was critical in bypassing certain problems and quickly producing predictions. Different methods of machine learning were also used to make specific predictions about certain elements.

[90] Roberto Pieraccini. The voice in the machine: Building computers that understand speech. *The MIT Press*, 2012.
(Nicole Trenholm).

[91] Chris Preimesberger. why AI machine learning will become mainstream in 2017. *eweek*, page 1, December 2016.
(Lydia Sanchez).

[92] Madhavi Ramani. Art(ificial intelligence). *Wilson Quarterly*, pages 00–00, 2018.
(Easton Jensen).

[93] E. Madhusudhana Reddy. Able machine learning method for classifying disease. *International Journal of Recent research aspects*, pages 223–226, March 2018.
(Yuan Hong) Machine learning are almost at any field of research, the time passed on and it is more likely to become a medical tool. People could use this to preserve a better medical care, no matter men or children. Research shows that this method is getting more reliable and becoming a potential value in the future.

[94] E. Madhusudhana Reddy and P. Bhaskar. Able machine learning method for classifying disease-treatment semantic relations from bio-medical sentences. *International Journal of Recent Research Aspects*, pages 223–226,

Mar 2018.
(Austin Stala) This article discusses how AI and machine learning can be used in medical field. One of the ways mentioned in the article is how it Can help in selecting treatment faster and more accurately. The article also states that AI with machine learning can help human doctors researching treatments and diseases.

[95] Matt Retnolds. Neural net learns words like a child. *New Scientist*, pages 12–12, November 2017.
(Tiff Serra-Pichardo).

[96] Matt Reynolds. Face ID tech can see through your disguise. *New Scientist*, page 7, September 2017.
(Jakob Orel).

[97] Adee sally. Beware tweets that lure clicks to steal info. *New Scientist*, page 21, August 2016.
(William Golden) Twitter holds a machine learning system that aids in phishing schemes.

[98] David Scheider. Deeper and cheaper machine learing. *IEEE Spectrum*, pages 43–44, January 2017.
(Bram Dedrick).

[99] David Schneider. Deeper and cheaper machine learning. *IEEE spectrum*, pages 42–43, 2017.
(Easton Jensen).

[100] David Schnieder. Deeper and cheaper machine learning. *IEEE spectrum*, pages 42–43, January 2017.
(Rodrigo Martinez) As technology advances we want to take efficient steps. Companies such as Nvidia, Intel, Microsoft and Google are finding ways to approach this. Nvidia has done this by shifting from their Maxwell Architecture chips to a Pascal Architecture.

[101] Nate Seltenrich. Singapore success. *Environmental Health Perspectives*, page 167, September 2016.
(Tony Ferenzi).

[102] Nathan Sinnott. How machine learning is changing the world and your everyday life. *Entrepreneur*, page N/A, April 2018.
(Rodrigo Martinez) Machine learning has evolved along society and is helping society grow and improve. This affects the education, labor work, law, health, and transport environments.

[103] James Somer'. Is ai riding a one-trick pony? *MIT Technology Review*, pages 28–36, November/December 2017.
(Tiff Serra-Pichardo).

[104] James Somers. Is AI riding a one-trick pony. *Technology Review*, pages 28–36, Nov/Dec 2017.
(Austin Stala) This article talks about the limitations of machine learning. The central point of the article is how the main concepts behind machine learning are over 30 years old and are reaching their limitations.

[105] Kaan Koseler; Matthew Stephan. Machine learning applications in baseball: a systematic literature review. *Taylor & Francis*, pages 745–763, 2018.
(Easton Jensen).

[106] Hakon Hapnes Strand. How plentiful are machine learning jobs in 2018? *Forbes*, page N/A, April 2018.
(Rodrigo Martinez) Machine Learning has driven the tech industry to a whole new level. This brings a whole new world of job opportunities. We take a look on how countries are doing with machine learning jobs. One example is the United States in which is the biggest machine learning countries having places such as Silicon Valley.

[107] Revell timothy. Play it for real. *New Scientist*, page 24, October 2016.
(William Golden) Scientists use the realistic driving simulation of GTA V to teach self-driving cars.

[108] unknown. Slef-driving cars—an ethical perspective. *Penn Bioethics Journal*, pages 8–8, 2015.
(Koichi Okazaki) Recently, many major companies such as Google, Benz, General Motors, and Toyota are developing the technology of self-driving cars and it is becoming capable for the practical use. One of the most beneficial point of self-driving cars is that reduces the car accident dramatically. However, one ethical issue came up, if the accident occurred and if it was unavoidable to hit someone the program might have to choose the victims, in order to save others. Thus, this is considered as a crucial issue since where the responsibility goes isnt clear, is it the car, producer, or consumer?

[109] Mikhail Urbazarv. Estimation of forest aboveground biomass and uncertainties by integration of field measurements, airborne lidar, and sar and optical satellite data in mexico. *Carbon Balance & Management*, pages 1–20, February 2018.
(Tony Ferenzi).

[110] Moshe Y. Vardi. Artificial intelligence past and future. *ACM*, page 5, 2012.
(Easton Jensen).

[111] Jaikumar Vijayan. Google for jobs employs machine learning to help people find work. *eWeek*, page 1, May 2017.
(Rodrigo Martinez) Google decides to use machine learning to help people

find Jobs in a much more efficient manner. People will be able to search for jobs in their area in which matches their criteria. With this people may have a better chance getting a job opportunity.

[112] Athanasios Voulodimos, Nikolaos Doulamis, George Bebis, and Tania Stathaki. Recent developments in deep learning for engineering applications. *Computational Intelligence & Neuroscience*, pages 1–2, 2018.
(Koichi Okazaki).

[113] John Wallace. All-optical neurak network can be 100 times faster than electronic versions. *Laser Focus World*, pages 8–13, November 2017.
(Tiff Serra-Pichardo).

[114] Chelsea Whyte. I teach machines to hunt down cancer. *New scientist*, pages 42–43, 2018.
(Easton Jensen).

[115] Chelsea Whyte. I teach machines to hunt down cancer. *New Scientist*, pages 42–43, 2018.
(Koichi Okazaki).

[116] Jeffrey A. Wilke. Battle of the brains. *Economist*, pages 61–63, December 2017.
(Jakob Orel).

[117] Alec Wilkinson. Mask crusader. *New Yorker*, pages 22–23, July 2013.
(Lydia Sanchez).

[118] Wun-ae Kim. Knowledge-based didagnosis and prediction using big data and deep learning in precision medicine. *Investigative & Clinical Urology*, pages 69–71, 2018.
(Koichi Okazaki).

[119] M Yoon. New reinforcement learning algorithm for robot soccer. *Orion*, pages 1–20, February 2017.
(Tony Ferenzi) Reinforcement learning can be used in many different ways, including in robot soccer. A specific type of reinforcement learning is being used called Temporal-Difference value iteration algorithm with state-value functions, This algorithm helped with quick decision-making for the robots in play during the game. An Artificial Neural Network is once again used to improve choices called the Multi-Layer Perceptron. In the end, the reinforcement learning proved effective in various tests involving shooting a soccer ball into a goal.

[120] Naihui Zhou. Crowdsourcing image analysis for plant phenomics to generate ground truth data for machine learning. *PLoS Comutational Biology*, pages 1–16, July 2018.
(Marcellus Parks).