How Machine Learning is Helping us Get Smarter at Healthcare

ClearDATA's Chief Technology Officer Matt Ferrari shares insights on:

- Machine learning in the public cloud
- Industry sector use cases
- Examples of cloud-managed services for machine learning applications

This whitepaper is designed for healthcare execs looking for a foundational understanding of what machine learning means to healthcare in the public cloud. These are the observations and opinions of ClearDATA's CTO Matt Ferrari and does not imply endorsement of his comments from any of the three public clouds mentioned: Amazon Web Services, Google Cloud Platform and Microsoft Azure.

ClearDATA®
SECURE • HEALTHCARE • CLOUD

Fascinating use cases for machine learning are springing up across all healthcare sectors. Machine learning promises to improve the quality and speed of care, build efficiencies, and offer better patient-provider relationships. In this paper, we explore how machine learning is being applied to different healthcare sectors across public cloud environments such as Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP). We conclude with a description of managed services in the cloud that can help your own organization replace manual or cumbersome processes with smarter, faster machine learning.

Defining Machine Learning

In 1950, Alan Turing proposed a method to determine if a machine could "think" in a manner indistinguishable from a human being. The original Turing Test required one human and one computer to communicate only via text on a computer screen, with another human acting as judge to determine which of those communicating was human, and which was machine. Since then, we've come a long way with artificial intelligence. In May 2018, attendees at Google's developer confab I/O 2018 watched as Google Duplex[1], Google's new digital assistant mimicked the human voice in conversation, and made phone calls to book appointments on behalf of the human requesting them.

So, what exactly do we mean when we refer to machine learning today? There has been some debate regarding a universally-approved definition for machine learning or deep learning, but in simplest terms, it involves "training" computer systems to find, remember and act on patterns in data. And similar to statistics, machine learning requires large data sets to reach meaningful outcomes. Healthcare is an area that has gotten very proficient at producing massive amounts of data, much more than any clinician or physician could begin to process alone. As such, machine learning outcomes stand to benefit patients more than the solo efforts of one clinician or healthcare practitioner.

Machine learning today can be supervised or unsupervised. With supervision, the variables and outcomes are known. We teach the machine (that is, create algorithms) by monitoring and ensuring the training model is reaching the desired results. We can explain how the results were reached, as the machines sort through massive data sets. With unsupervised learning, the machine's outputs are unknown at the onset. Algorithms are essentially set loose with the data to glean insights that can be very meaningful and important, but harder to explain. Supervised is more common than unsupervised, but both methods are growing, and both have merit.

In the past, activities controlled by artificial intelligence and machine learning were very linear and the end goals were quite basic: program a machine to accomplish a simple task and then simply press the Enter button. Input A and expect Output A.

The new approach is to provide the machine a start "state" with a specific end goal, then allow the system to develop its own intermediate states and let it progress from one to the next until that goal is accomplished.[2]

Which specific machine learning applications help us collect mountains of once disparate data and produce meaningful outcomes that can improve healthcare? Examples follow from each of the three major public clouds today. We look at the technology running the solution, what the solution does, and why it matters to healthcare.

Amazon Web Services

Like the other primary cloud platforms, Amazon is investing heavily in machine learning. In the last year, Amazon launched Lex, which is designed to enable developers to build conversational interfaces for applications. It is powered by the same deep learning technologies as Alexa. It uses advanced deep learning or machine learning for automatic speech recognition to convert speech to text and natural language understanding to recognize the intent of the text--all of which help create engaging user experiences.

Amazon machine learning, or AML, provides companies with a method for interpreting data. It offers visual aids and easily-digestible analytics reports and is based in part on Amazon's existing internal system algorithms. It's very accessible even to those practitioners who may not have extensive experience with data science and it attempts to help businesses build machine learning models of their own without have to invest in countless hours spent on creating the necessary code to do so otherwise.[3] By comparison, in the past, a practitioner may have used Dragon to create speech–to-text, then add natural language processing on top of that. Today, public clouds are embedding this natural language processing straight into their platforms, making the same deep learning that powers Alexa available to any developer who wants to quickly build sophisticated conversational bots or chat bots.

Lex scales automatically, taking away the need to be concerned about infrastructure capacity, and users only pay for what they use. It's an impressive achievement. Expect healthcare organizations to increasingly request this kind of functionality behind business associated-covered services and to make it HIPAA compliant. There are third-party services available that provide a secure, HIPAA compliant environment for machine learning in healthcare which we'll look at further in this paper's conclusion.

Amazon also launched SageMaker, which enables developers or data scientists to deploy machine learning models at any scale and move quickly through production, so they can gain insights from data much sooner. SageMaker bypasses the need for data scientists to go through acquisition processes, buy Cap-ex assets, and get involved with infrastructure demands that are outside of their area of expertise. Instead, these data scientists gain easy access to their data sources for analysis, deploying a machine learning model with a single click from the SageMaker console.

While SageMaker is not currently HIPAA eligible, Amazon is moving that way, and in the meantime, healthcare is benefitting by using depersonalized, anonymized data with machine learning to learn volumes about predicting and preventing disease.

Machine learning has moved from just speech and text applications to imaging. The possibilities for machine learning off images in healthcare are profound, with many use cases happening right now. Amazon Rekognition is a machine learning image recognition service that recognizes objects, people and activities in images and medical videos. Currently used predominantly for facial recognition and security to detect unsafe or inappropriate images, the possibilities for healthcare are tremendous, such as decision support in oncology and radiology.

Moreover, Amazon Machine Image [4] technology provides practitioners and researchers pre-installed tools that can scale and integrate with deep learning frameworks, whether developed by Amazon or others. For example, Cognitive Toolkit by Microsoft Tensor Flow is developed by Microsoft, but many TensorFlow projects are also running on AWS and GCP. AWS Glue is another machine learning friendly solution. It's an extract, transform and load service that enables users to easily prepare and submit data for analytics. It is a good example of one of the new "serverless" technologies being adopted by healthcare. Serverless means no infrastructure to buy, manage and maintain – the environment is automatically provisioned, and customers only pay for what they use; in this case, to do their data preparation.

Amazon's auto scaling cluster is yet another tool that speeds use of machine learning, by enabling data scientists to run a machine learning algorithm against massive data sets in minutes. An emerging use is to run queries against "data lakes,"[4] without the prior or concurrent need to move or synchronize the data.

Google Cloud Platform

Google has significant history with artificial intelligence and machine learning, including its Google Cloud Machine Learning Engine that ingests structured or unstructured data into the Google Cloud Platform (GCP). The engine's training and predictive services can be used independently or together. The interface allows data scientists to automatically design and evaluate their model architectures to reach optimal results. Google also supports the import of machine learning models that have been "trained" elsewhere.

GCP is one of the premier cloud computing platforms on the market today. When it began in 2008 as the Google App Engine, it was used mostly to allow developers to create applications and host them in Google's ecosystem. Soon thereafter, Google launched their Computer Engine application to help support the use of virtual machines. GCP now provides a very broad spectrum of complementary services including computing, storage, API systems, productivity tools, IoT services and of course the Cloud.[5]

Google Cloud also offers big data solutions, including an integrated, serverless Big Data platform that captures, processes, stores and analyzes data all within one very scalable platform free from traditional constraints. Like AWS, Google has its own speech detection and natural language processing tools. It also offers special tools for anonymizing data. Cloud Dataprep cleans structured data (such as from a SQL database) but can also use unstructured data in a serverless environment. This is a huge step forward, as the healthcare organization sidesteps investing in and maintaining servers and security. With Cloud Dataprep, data scientists can run their data prep with just clicks instead of code, simply dragging and dropping. There's a trend in this direction, which speeds the time it takes for the data scientists to start making decisions based on the anonymized data. We will talk more about use cases later in this paper, but for now it is worth noting that this has powerful implications for accelerating the time it takes to complete drug studies and clinical trials.

A well-known machine learning technology is DeepMind, the artificial intelligence venture by Google which has created neural networks that learned to play video games in a fashion like humans. Of greater importance to those of us in healthcare is DeepMind's impressive ability to process images.

This includes its work examining mammograms with the Cancer Research UK Imperial Centre[6], as well as work with macular degeneration in aging eyes, and much more. By finding early warning signs we can offer proactive healthcare options rather than attending to more disastrous consequences for patients later. And, as is always the case with machine learning, the more images DeepMind processes, the smarter it is going to become.

Micosoft Azure

Microsoft Azure was launched in 2010 and leveraged existing Microsoft technologies such as RemoteApp, Active Directory and SQL Server. Essentially it is a collection of various cloud computing services. Azure is a tempting option for companies as resource capacity is available on demand and as a product, it does not require any initial costs to implement. Microsoft only charges for total resource consumption as opposed to billing for server installation or leasing fees. The Azure umbrella has now grown to include multiple services such as an IoT suite, a managed search service and a media service for content protection and video playback. [7]

Microsoft was one of the earliest pioneers in analytics tools, based upon what has become the Cortana virtual assistant that uses voice command technology across multiple devices. Microsoft was one of the first to offer serverless computing with no infrastructure to manage. One such tool is the Azure Machine Learning Studio, which Microsoft describes as "a collaborative, drag-and-drop tool you can use to build, test, and deploy predictive analytics solutions on your data. Machine Learning Studio publishes models as web services that can easily be consumed by custom apps or BI tools such as Excel." Designed for applied machine learning with built-in templates, data scientists need not write a single line of code to utilize this machine learning platform.

In fact, one of the interesting things around Microsoft's machine learning capabilities is the plethora of pre-configured environments it offers. Most data scientists aren't interested in vetting, trialing and then deploying the various tools needed to drive decisions from machine learning. To that end, Microsoft has pre-configured environments to get started quickly, working with familiar processing systems like Anaconda Python to Apache Spark.

Microsoft also has a deep learning virtual machine straight out of the box, running on Azure. The Deep Learning Virtual Machine makes it more straightforward to use GPU-based VM instances for training deep learning models. Lastly, Azure's machine learning capabilities include enabling customers to build AI apps that sense process and act on information. This can increase speed and efficiency, and can lead to better healthcare outcomes.

As machine learning and healthcare move forward, expect Amazon, Google and Microsoft to make more of their services eligible for protected health information (PHI). Amazon, for example, adds more services all the time behind HIPAA compliance. In the meantime, healthcare organizations are anonymizing data to use public cloud machine learning capabilities. The movement has begun and is gaining momentum daily offering benefits to those we serve industrywide.

Use Cases by Sector

While we are only at the beginning of understanding all of the use cases for machine learning, we are far enough along in adoption to have many notable examples. Let's explore these by sector.

Providers

Securing Data

For doctors and hospitals, a good 'best use' case for machine learning is in cybersecurity. A primary advantage machine learning is bringing to providers today is to help defend against the growing threat of ransomware. With healthcare-based ransomware attacks happening every day[8], cybersecurity focused on machine learning can identify potential vulnerabilities and prevent attacks before malicious actors find them. Intelligent algorithms flag out-of-the-ordinary activity and send alerts for timely interventions. By calculating scores for where the data should be moving, and by whom, and in what ways, machine learning can help target problem areas long before an IT or security team member may notice the anomaly.

Providers seeking such machine learning expertise in cybersecurity should look to managed services vendors that deal exclusively in healthcare—and can prove their experience and willingness to assume risk of a breach via a comprehensive BAA.

In-depth Research

Imaging offers another best use case. Machine learning is providing consistently correct determinations from image recognition, and in many cases, to a higher degree than a human. As such, we're starting to see major investments in this area. Microsoft, for example, is succeeding in this area with biomarkers and phenotyping, and the results are driving cancer research forward as never before.[9] Provider organizations are also gathering and cleaning their data in anticipation of being able to use it for image analysis behind a business associate agreement. An ideal managed services partner for such a use case would have healthcare expertise, HITRUST Certification, and offer scalable, secure data management in any of the dominant public clouds.

Improving Patient Care

A third broad adoption example of machine learning in hospitals and clinics surrounds the use of Natural Language Processing (NLP); particularly to process information in the EHR. Clinicians can research what prescriptions a patient is on and use NLP to search the unstructured data at scale to determine the answer. They can ask how many patients a facility has with similar conditions. In addition to driving down costs, NLP can help drive up consistency of care in healthcare practices.

And for providers, the ability to use machine learning to extract significant meaning from massive amounts of unstructured text is helping with critical decision making. The doctor or nurse who enters a patient's room typically begins with a series of important questions. Whether it's being recorded on audio, notepad or iPad™, there is great potential in that data, if it gets used. Accordingly, providers are investing a lot of time and money into machine learning, so they can augment and better position diagnosis. In just one such example, machine learning is being used to distinguish between heart conditions that have very similar results. This reduces the risk of providers opting for the diagnosis they are most familiar with when a different diagnosis may be the best decision.

To summarize, machine learning is making it possible to read all the structured and unstructured data sitting in EHR applications and analyze what is really happening. Practitioners can search for commonalities or determine what kind of trauma a patient may have been through to better determine the critical care or long-term care the patient needs to manage their disease or illness. In addition to existing illness, the practitioner can more accurately predict who is at risk in the near term of certain diseases. By being more accurate in the diagnosis, there is less repeat treatment and the facility can optimize its code, reduce costs, and most importantly, improve patient outcomes.

Payers

When we pivot to insurance providers, the best use cases for machine learning change dramatically. They are highly focused on payment and revenue cycle management, participation with HIEs, and patient engagement.

Revenue Cycle Management

In one scenario, payers are using machine learning to quickly determine how soon, based on a particular diagnosis, a hospital gets paid for services rendered. How quickly does the patient get the letter in the mail that says what their financial responsibility is, and how quickly do they pay it? In another patient-oriented example, machine learning can be used by payers to help the provider better understand some things about the patients, such as which ones are likely to be "no shows," which are likely to receive services, but then not pay, and more. Payers are also using machine learning in participatory healthcare scenarios. Patients are now doing their own research before they select a doctor, and payers can help them understand which ones will provide them the best coverage on their plan.

Improving Clinical Efficiency

Merging the clinical and financial, healthcare payers are also focusing on improving clinical efficiency while they improve clinical outcomes. Here they use machine learning to look at which treatments deliver the most clinically effective outcome at the lowest cost.

Fraud detection is another best use case for payers that need to know if the person submitting a claim is in fact that person, and did in fact have that treatment or surgery, especially with ever-escalating medical identity theft. With real-time authorizations that work in a way similar to the banking industry model, payers can ensure the authenticity of a claim.

Pharma and Life Sciences

We see deep adoption of machine learning in the pharmaceutical space, primarily to enhance clinical trials by enabling pharma and life sciences companies to use big data and analytics engines to normalize and de-identify patient data. They can use this data to run machine learning-driven trials to understand outcomes of a drug, finding ways to improve it before bringing it to market, all without having to test thousands of patients who may have negative side effects from drug trials. With the ability to more quickly identify potential outcomes and iterate and retest, researchers can help speed to market life-changing treatment.

Independent Software Vendors

Healthcare technology companies are making use of machine learning to create more intelligent apps. Many ClearDATA® customers, for example, use machine learning to understand, with specificity, how the patient or doctor uses the app; in turn, the company's developers can further improve features based on usability.

Or they may feed their data into a data lake to project and predict scenarios and outcomes; for instance, to examine the path of a patient who has a specific disease and has been to a neurologist and an ophthalmologist. Machine learning can determine the next likely step to best serve this patient. And yes, all of this data can be fed back into the patient's electronic record.

Applications of Machine Learning in the Public Cloud

Now that we've looked at what machine learning is, how the clouds are expanding machine learning opportunities, and the ways in which various healthcare sectors are focusing their machine learning efforts, let's dive a little deeper into a few examples based on specific objectives.

Data preparation

Over the past several years, biomedical researchers have figured out how to use machine learning to perform studies based on de-identified patient data derived from electronic health records and other sources. The major public cloud providers – Amazon, Google and Microsoft – now enable data scientists for pharmaceutical firms to do such research in a HIPAA-compliant environment without being machine learning experts.

One of the major use cases for this approach is research that seeks to understand how genetic differences determine the responses of individuals to particular drugs. To accomplish this goal, data on the physical characteristics and medical histories of individual patients, known as phenotypes, are compared to their genetic features, known as genotypes, which are derived from genomic sequencing.

Before machine learning was used for this task, phenotypes had to be derived manually--a very laborious and time consuming[10] job. Machine learning, in contrast, allowed phenotypes to be generated using an automatic, unsupervised and high-throughput process. But researchers who wanted to use this approach in onsite data centers had to buy additional equipment, hire developers with the requisite skill set, and build their own algorithms.

Today, a pharmaceutical firm can buy the same infrastructure as a service (IaaS) in a public cloud, which can be scaled as needed. Machine learning algorithms and related services in the cloud are prebuilt for specific purposes. The Google Cloud, for example, has products called Cloud Data Lab, which explores both structured and unstructured data sets, and Cloud Data Prep, which cleans and prepares data for analysis and can also de-identify clinical data.

Large scale data analysis

The life sciences community is very focused on genomic research, which promises to revolutionize our understanding of health and illness. The challenge is the massive scale of genomic information. It takes a petabyte of data, for instance, to represent a single human genome. To store and analyze this amount of data in a data center would be very difficult for any organization. So public clouds are increasingly being utilized for genomic research and are providing machine learning tools to analyze the data.

The National Institutes of Health's All of Us project[11], in conjunction with academic medical centers, aims to analyze genetic information from a million people is the best-known of the initiatives launched to discover cures through genomic research. Many healthcare organizations are doing similar work on precision medicine, using machine learning in the cloud.

For example, the Colorado Center for Personalized Medicine (CCPM), a branch of the University of Colorado Denver, is helping doctors evaluate patients[12] at the molecular level to predict their risk for disease and to develop personalized therapies based on their DNA. This requires large-scale analysis of their genetics and the health histories of thousands of patients to find patterns.

Because on-premises technology is too expensive for this kind of research, CCPM migrated its data infrastructure to a HIPAA-compliant public cloud.

CCPM is also using an algorithmic database application to accelerate its research. This application has reduced the time required for responses to data queries by 97%.

Algorithm-driven diagnosing

Cardiology, radiology, and pathology, which use large image databases, are natural candidates for machine learning. Algorithms have been developed, for example, to identify abnormalities in images that may indicate the presence of disease. These algorithms are as accurate as and sometimes more accurate than physicians who are trained to "read" these images.

Google, for instance, attracted media attention in 2017 when it used machine learning, predictive analytics and pattern recognition to detect breast cancer by looking for cell patterns in tissue slides. A study showed[13] that machine learning could identify breast cancer with 89 percent accuracy, compared to the 73 percent score achieved by a human pathologist.

Similarly, Stanford University researchers used a deep learning algorithm to identify skin cancer[14] as well as dermatologists did. The computer scientists made a database of nearly 130,000 skin disease images and trained their algorithm to visually diagnose potential cancer. Instead of writing code to tell the computer what to look for, they trained the algorithm to decide which lesions were cancerous.

Sepsis prevention modeling

A few years ago, a large medical research university set out to predict which patients were at risk of developing sepsis, a highly dangerous condition that is one of the leading causes of mortality in hospitals. Using machine learning tools in the cloud, researchers figured out how to identify these patients 24 hours earlier than could have been done with traditional methods.

The team that developed this clinical protocol started with an expert model that used specific thresholds of temperature, heart rate, respiratory rate, and white blood count as key indicators of sepsis risk. After loading in the available data on a patient, including the data in their EHR, the computer used an algorithm to determine how closely a patient's characteristics matched those of patients who had previously developed sepsis. When a patient matched the profile, their physician received an alert, acted on it or didn't, and fed their reaction back to the algorithm to improve it. Over time, this early warning system became highly sensitive and is now used routinely at that university.

Managed Cloud Services for Machine Learning Projects

As sophisticated as the above use cases may be, almost any healthcare organization can embark on a machine learning initiative without the need to add IT infrastructure. As a full-service cloud provider, ClearDATA is at the forefront of helping organizations of every size from every sector shed the constraints of outdated, expensive infrastructure and participate in the cloud economy.

ClearDATA's multi-cloud abilities are some of its most relevant and useful features. However, that flexibility requires extra scrutiny around security. Purpose-built safeguards, automation and healthcare expertise are just a few of the tools leveraged by the platform to ensure full protection of sensitive healthcare data. As the platform was designed exclusively for use within healthcare organizations, it comes standard with comprehensive, real-time support from the ClearDATA team and is backed by an industry-leading BAA.

Additionally, the platform is also HITRUST Certified ensuring that ClearDATA's cloud computing and backup services meet the highest industry standards for managing protected health information.

Moreover, in addition to possessing multiple certifications, ClearDATA is also fully compliant with HIPAA, GxP and GDPR regulations and frameworks. Mitigating risk is easily managed via the compliance dashboard which allows users to monitors thousands of components across multiple cloud environments. The dashboard is an easy way to get a glance into the state of organizational compliance whenever it's needed. Lastly, these great tools and functions are further strengthened with the support of ClearDATA's array of professional services. Whether the challenge is developing a process design, custom development or a risk assessment, ClearDATA delivers extensive healthcare and IT expertise to provide the necessary solutions to resolve any issue. Even with potentially overlooked needs such as breach simulation, penetration testing, security audit support and application code reviews, ClearDATA applies a hands-on consultative approach to every challenge. ClearDATA also offers service consultations for various types of project specifications.

Conclusion

Much like how Amazon Prime has completely redefined consumer access to on-demand entertainment, machine learning is also redefining the state of healthcare in 2018. Machine learning allows healthcare systems to break their normal way of thinking in terms of existing models and allows clinical teams to test for potential outcomes using more complex algorithms than they ever could before. Machine learning is ultimately a game-changing technology that healthcare systems need in order to help overcome the challenges they encounter every day. [15] While there was a time when talk of machine learning was futuristic and visionary, the healthcare industry has moved past the idea to actual broad adoption. All signs point to increasingly deep use, with the opportunities only limited by our collective imaginations and expertise. The outcome stands to benefit patients as well as providers, payers and those working to serve them, as we all work to make healthcare smarter.

Resources

1 Rawes, E. (2018, July 2). What is Google Duplex? Digital Trends. Retrieved from https://www.digitaltrends.com/home/what-is-google-duplex/

2 Drepper, U. (2017, Sept 20). An introduction to machine learning today. Retrieved from https://opensource.com/article/17/9/introduction-machine-learning

3 Reese, H. (2016, Dec. 1). Should Amazon be your AI and machine learning platform? Retrieved on from https://www.zdnet.com/article/should-amazon-be-your-ai-and-machine-learning-platform/

4 AWS. (n.d.). What is a data lake? [Blog post]. Retreived from https://aws.amazon.com/big-data/datalakes-and-analytics/what-is-a-data-lake/

5 Forrest, C. (2018, May 14). Google Cloud Platform: A cheat sheet. Retrieved from https://www.techrepublic.com/article/google-cloud-platform-the-smart-persons-guide/

6 DeepMind Health. (2017 November 24). DeepMind, the Cancer Research UK Imperial Centre and OPTIMAM [Blog Post]. Retrieved from https://deepmind.com/applied/deepmind-health/working-partners/health-research-tomorrow/cancer-research-imperial-optimam/

7 Sanders, J. (2018, June 18). Microsoft Azure: A cheat sheet. Retrieved from https://www.techrepublic.com/article/microsoft-azure-the-smart-persons-guide/

8 Donovan, F. (2018, May 3). Healthcare Industry Takes Brunt of Ransomware Attacks. HealthITSecurity. Retrieved from https://healthitsecurity.com/news/healthcare-industry-takes-brunt-of-ransomware-attacks

9 Thatcher, L. (n.d.). How Machine Learning in Healthcare Saves Lives. Retrieved from https://www.healthcatalyst.com/how-machine-learning-in-healthcare-saves-lives

10 Jette Henderson, MS,1 Ryan Bridges, BS,2 Joyce C. Ho, MA, PhD,3 Byron C. Wallace, PhD,4 and Joydeep Ghosh, PhD1.(2017, Jul 26). PheKnow–Cloud: A Tool for Evaluating High-Throughput Phenotype Candidates using Online Medical Literature. AMIA Jt Summits Transl Sci Proc. Retrieved from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5543339/

11 NIH. (n.d.). All of Us Research. [Blog post]. Retrieved from https://allofus.nih.gov/

12 Google. (n.d.) Colorado Center for Personalized Medicine: Integrating genetic data and patient records in the cloud to personalize care. [Blog post]. Retreived from https://cloud.google.com/customers/colorado-center-for-personalized-medicine/

13 Stumpe, M. (2017, March 3). Assisting Pathologists in Detecting Cancer with Deep Learning
Posted by Martin, Technical Lead, and Lily Peng, Product Manager. Retreived from https://ai.googleblog.com/2017/03/assisting-pathologists-in-detecting.html

14 Kubota, T. (2017, January 25). Deep learning algorithm does as well as dermatologists in identifying skin cancer. Retreived from https://news.stanford.edu/2017/01/25/artificial-intelligence-used-identify-skin-cancer/.

15 Lee, P. (2018, Feb 28) Microsoft's focus on transforming healthcare: Intelligent health through AI and the cloud Microsoft [Blog Post]. Retrieved from https://blogs.microsoft.com/blog/2018/02/28/microsofts-focus-transforming-healthcare-intelligent-health-ai-cloud/