RAIS RESEARCH ASSOCIATION for INTERDISCIPLINARY STUDIES
June 2022

DOI: 10.5281/zenodo.6945948

AWS Corporate AI Use Cases

Brian Kan[1], Douglas Klein[2]

[1]*Crean Lutheran High School, Irvine, United States, briankan1115@gmail.com*
[2]*New Jersey City University/Department of Business, Jersey City, United States, dougkleinteacher@gmail.com*

ABSTRACT: Amazon, with $469 Billion in sales in 2021, has established itself as a world-class user of AI, utilizing Machine Learning (ML) in its search engine to deliver desired results quickly - so millions of shoppers find the products they want to buy. Amazon's affiliate, Amazon Web Services (AWS), had annual sales of $62 Billion in 2021, making it the 53rd largest company on the Fortune 500 as measured by revenues. AWS provides enterprises with a fully managed AI service with tools needed to execute every step of the ML development lifecycle in one integrated environment. By 2021, more than one hundred thousand companies utilized AWS Machine Learning - more than any other cloud platform. Outside of the traditional search engine applications (finding a product to buy, booking travel arrangements, or looking for a new job) what are some compelling and important business use cases where ML and AI have the greatest impact? Some use cases in this paper: AWS AI and Machine Learning are used by commercial landlords and industrial real estate owners to save energy and reduce carbon emissions. The World Wildlife Federation uses AWS AI tools in Indonesia to better understand the size and health of orangutan populations in their native habitat. And The Walt Disney Company uses ML and AI to organize metadata into one archival system, storing information about the stories, scenes, and characters in every second of Disney's huge catalog of shows and movies.

KEYWORDS: AWS, Amazon Web Services AI, AWS Machine Learning, AWS Business Use Cases

Introduction

Artificial Intelligence (AI) for business use cases has progressed from early adoption to the growth stage and is used by tens of thousands of corporate enterprises - large and small - globally. By 2022, one enterprise is poised to be the undisputed leader in artificial intelligence research, commercial licensing, and applications: Amazon. Specifically, the subsidiary Amazon Web Services, most commonly referred to as AWS, launched by Amazon 15 years ago. With $62.2 billion in sales for the calendar year 2021, AWS would be ranked 53 on the Fortune 500 list of largest US companies if separated from Amazon, following #50 The Disney Co. (founded 99 years ago), Archer Daniels Midland (founded 120 years ago) and Albertsons/Vons/Safeway (founded 115 years ago) (Fortune 2021).

According to the Amazon Annual Report (2021), AWS serves "developers and enterprises of all sizes, including start-ups, government agencies, and academic institutions, through AWS, which offers a broad set of on-demand technology services, including compute, storage, database, analytics, and machine learning, and other services."

Operating income in Amazon's AWS division was $18.5 billion in the calendar year 2021, an increase of 37% compared to the same period in 2020. This equates to an operating profit margin of 30% in 2021- the Management Discussion indicated that Amazon will continue to invest in AWS technology and capacity, as it clearly has liquidity. Of the total $55.4 billion investment Amazon made in technology infrastructure, the Annual Report states "the majority" went to support AWS (Amazon 2021).

What about artificial intelligence specifically? Amazon believes that "the practical applications of artificial intelligence and machine learning, will continue to improve users' experience on the Internet and increase its ubiquity in people's lives. To best take advantage of these continued advances in technology, we are investing in AWS" (Amazon 2022).

How does Amazon define Artificial Intelligence?

AWS describes AI as Machine Learning and Deep Learning, both computer science fields derived from the discipline of Artificial Intelligence. Professor Pedro Domingos, a prominent researcher in the field of AI, says there are 5 "tribes" of machine learning:

- Symbolists - origins in logic and philosophy
- Connectionists - stemming from neuroscience
- Evolutionaries - relating to evolutionary biology
- Bayesians - engaged with statistics and probability
- Analogizers - with origins in psychology (Amazon 2022).

Recently, advances in the efficiency of statistical computation have led to Bayesians being successful at furthering the field in a number of areas, under the name "machine learning". Similarly, advances in network computation have led to connectionists furthering a subfield under the name "deep learning". Broadly, these techniques are separated into "supervised" and "unsupervised" learning techniques, where "supervised" uses training data that includes the desired output, and "unsupervised" uses training data without the desired output (Amazon 2022).

Amazon's primary AI consumer products are delivered via the Amazon Echo - from 2017-2019, over 100 million devices were sold and 56 million are currently installed, representing 25% of all households in the US (Bohn 2019). . The Echo is the smart speaker housing the intelligent voice server Alexa. Amazon's Head of Devices, Dave Limp voiced the number, and Amazon's Vice President of Devices, Steve Rabuchin confirmed it. Alexa Smart Assistant and Echo products occupy 10,000 of Amazon's employees (Bohn 2019).

AWS' commercial AI products are Lex, the business version of Alexa, Polly, which turns text to speech and Rekognition, an image recognition service (Maguire 2022).

AWS' primary competitors in the AI/ML space are Google, IBM, Microsoft and Alibaba. There are at least 100 known vendors in the US with Go-to-Market strategies in the US, generally focusing on primary target markets:

- Healthcare
- Quantum Machine Learning using photonics
- Business Intelligence and Predictive Analytics
- Sales Enablement/Customer Information Systems
- Consumer Intelligence and Internet traffic tracking (next generation of cookies)
- Automated Personal Assistants
- Data Storytelling
- Scientific/Mathematical applications of algorithms for academic research ● Robotics and Smartphones
- Chatbots/Automated FAQs that allow for multi-round verbal conversation (Maguire 2022).

Machine Learning

Machine Learning (ML) describes a process in data science in which software programs repeatedly use previously gathered information in order to produce desired results based on a series of algorithms. Generally speaking, the purpose of ML is to quickly predict outcomes based on historical data. The use cases of ML are 1) to optimize a purpose-driven program by eliminating human biases 2) produce efficacious results given a certain amount of risk and unknowns 3) extract patterns and significant insights that are not readily obvious and 4) classify the inputs into succinct representations.

Machine Learning can be thought of as a next generation information science that builds on traditional software programs that yield concrete outcomes to generate open ended advanced products that the programmer didn't specifically envision. ML uses given historical data to

calculate intelligent results built from patterns the software engineers and data scientists seek to find from previously-gathered samples of information. It's important to note upfront that the accuracy of ML programs relies on the quality of the information in the historical examples and the accuracy and skills of the data scientists writing the algorithms deployed.

ML quickly leapt from computer science labs in universities and Think Tanks to the business world because in the US economy with a Gross Domestic Product of over $23 trillion annually, industries capture millions and even billions of customer data points. ML modeling analyzes large-scale business challenges to identify software language that can model predictions. The test or Proofpoint for ML models are whether or not the results provide statistical significance that is accurate - since money is involved, end-users must be confident that the predictions are accurate.

There are many reasons why data science is a growing field creating a lot of new jobs for knowledge workers and AI experts, including the need to:

- make sense of the billions of data inputs,
- write algorithms to yield perceptive outcomes
- evaluate and study correlations,
- work with business partners for cause and effect relationships
- test and analyze the results for accuracy and effectiveness.

Corporate business use cases for Machine Learning

Business organizations see Machine Learning as a tool to increase product sales by improving the results from internet search inquiries and from analyzing positive and negative consumer behavioral patterns based on the past data. Companies are interested in demographics, how to attract new customers, how to retain customers and how to predict customer profiles that become dissatisfied.

ML's objectives are to predict future outcomes in myriad business cases and complement business Intelligence (BI), which focuses on reporting past business data. The field of BI is still growing in the US in the current time frame, as the number of commercial BI vendors increases and there is greater adoption by businesses across all sectors and sizes - not just the largest corporations. BI identifies data points - building on historical reports, ML can identify irregularities to normal patterns, for financial institutions, it's important to detect fraud and possible criminal acts before they occur.

Amazon takes these steps to implement ML for their customers

First, identifying the right problem -- identifying the prediction that would benefit the business if ascertained (Amazon 2022).

Next, the data must be collected, based on historical business metrics (transactions, sales, attrition, etc.) (Amazon 2022).

Once the data is aggregated, a ML model can be built based on that data. The ML model is run and the prediction output of the model is applied back to the business system to make more informed decisions (Amazon 2022).

A few of the companies that incorporated machine learning with successful results are included in this paper: Disney, T-Mobile, and the NFL.

Disney Study

The Walt Disney Company partners with Amazon to use Machine Learning for an interesting use case. Disney has a global corporate goal to catalog and preserve every cartoon and live action character that has ever appeared in a Disney content production. This idea is an extension

of founder Walt Disney's policy going back to the 1950's to preserve all content and make it accessible to employees, artists, and researchers (Wired.com 2019). The idea is to tag every character that appears in every scene. This can be accomplished manually of course by having personnel watch the content then tag the characters second by second. By labeling the data on each pixel contained in the scene, Disney can search for the scenes and access the characters/animations at any time.

In 2018, Disney created the Direct-to-Consumer & International (DTCI) Technology team, where computer scientists unite with content experts from the Disney catalog department. (Wired.com 2019) Hardware had to be fitted with software to fulfill the dedicated purpose of tagging each character on devices. The ultimate root foundation of the DTCI effort is metadata: all the information about the plot lines, scripts, scenes and characters in every Disney show, clip and movie (Wired.com 2019). Miquel Farré, the team's technical lead, engaged AWS to automatically tag digital content with the required metadata to optimize the cataloging (Wired.com 2019).

In order for Disney's deep-learning tools to accurately produce metadata that stores the creative content, the team depends on writers and animators to delineate what characterizes each character to make them special. Remember, tagging everything with the right metadata presents a labor problem: because the DTCI Technology team doesn't have time to catalog every frame by hand. The Disney team even describes nature scenes and special effects environments that have no characters or dialogue, in order for metatags to be applied (Wired.com 2019).

AWS worked with DTCI to implement machine learning for the job of generating metadata. There are limitations to the ML capabilities, which means that Disney employees must visually approve the tagged scenes, but the goal is to reduce the workload from workers who are confirming character identification (Wired.com 2019). ML will help to reduce work required organizing the Disney library, with the goal of improving the accuracy of searches. Prior to this year, the team worked with more traditional machine-learning algorithms, which require less data than the newly developed deep-learning approach. However, this response also resulted in less flexible results (Wired.com 2019). Though fewer data inputs allowed for traditional algorithms to perform efficiently, an exponential increase in data allows the deep-learning approach to take leaps above the traditional one . As time goes on, the deep-learning ML model will improve the effectiveness of trained programs and will be refined for different purposes.

AWS has been a key partner in Disney's transition from traditional machine learning to deep learning, especially when it comes to experimentation. Elastic cloud computing EC2 instances allow the team to quickly test new versions of the model (Wired.com 2019). The metadata project was evangelized internally at Disney and recognized as a success - other operating divisions of the Disney Corporation want to replicate the concept with projects of their own. For instance, the ESPN sports networks - which have years of video and on-air productions in sports realms professional, college, youth, and amateur - want to begin a similar campaign to properly catalog and designate metadata on the articles and videos that most often appear on trending digital applications/websites (Wired.com 2019).

Beyond that, the machine-learning algorithms, and the metadata they deliver, can power more advanced AI to drive further implicit personalization (based on data relationships and behavior) overtime. Disney is using the MAP3 Media Annotation Platform to annotate and catolog data. MAP3 is a flexible and efficient approach to tagging and filtering scenes and sub-scenes in Disney's large content library (https://tagging.disney.com/#!/about) Farré sees no end, a limitless application of metadata, especially given Disney's infinitely growing library of unique content, characters, and products. He personally states, "I think we won't get bored." (Wired.com 2019).

T-Mobile Study

Through the use of predictive capabilities of ML, T-Mobile has improved customer service by exploiting the qualities of AI that complement human abilities (Aws.amazon.com 2022a).

T-Mobile sees AI as the perfect opportunity to offer customers better and faster service. This fact helps all constituencies -the company, support agents, and subscribers by creating stronger relationships (Aws.amazon.com 2022a). “Most industries have looked to use AI and machine learning to build more sophisticated Interactive Voice Response (IVR) systems and chatbots as a means to deflect for as long as possible the interaction between a human customer service agent and the customer,” says Cody Sanford, executive vice president and chief information officer at T-Mobile. T-Mobile customers immediately connect with a customer service agent that knows them, rather than talking to an IVR or chatbot. With the help of AI, these customer service agents can quickly access the information most salient to customer needs (Aws.amazon.com 2022a).

Collected company data includes hundreds of thousands of incoming customer requests a day. To properly infuse its data labeling with AI, T-Mobile turned to Amazon SageMaker Ground Truth. Ground Truth speeds up the labeling of training data nd scales it further - which is essential for machine learning models to produce predictions with high accuracy (Aws.amazon.com 2022a).

NFL Study

In the modern day, NFL’s Next Gen Stats (NGS) program utilizes sophisticated tracking technology that is collected via RFID devices placed in the shoulder pads of all players and embedded at each stadium. The little devices can capture complete data sets about the player's movement at any given moment, their specific location by inches, and the speed/direction in which they make their movements (Hardesty 2021). By partnering with Amazon Web Services, the NFL is leveraging the power of its data through sophisticated analytics and machine learning. “We had a lot of stats and wanted to find the best way to leverage them. We’re taking in so much data now with the tracking system that we’re able to use machine learning to understand what elements are relevant and what are not” (Hardesty 2021). The NGS platform deploys the machine learning tool Amazon SageMaker, which enables the NFL to rapidly create machine learning models capable of interpreting the flow of action (Hardesty 2021).

Of course, data is only useful when it can be quickly and easily accessed. Using the business intelligence tool Amazon QuickSight, the NFL is able to gain greater insight internally while also opening a window for fans to engage with data (Hardesty 2021). Patterns of plays based on data and analysis, via machine learning, could be the keys to better understanding where players are more likely to get injured and to help design rules to mitigate risk (Hardesty 2021). The end result is a better experience for fans, players and teams—analogous to the successful T Mobile customer service project. It’s a win-win scenario for everyone involved powered by AI and Machine Learning.

References

Amazon. 2021. “Amazon Annual Report 2021.” https://ir.aboutamazon.com/overview/default.aspx.

Amazon. 2022. “What is Artificial Intelligence? Machine Learning and Deep Learning.” https://aws.amazon.com/machine-learning/what-is-ai/.

Aws.amazon.com. 2022a. “At T-Mobile, AI Humanizes Customer Service.” https://aws.amazon.com/machine-learning/customers/innovators/t_mobile/.

Aws.amazon.com. 2022b. “Customer Success Stories. Increase agility, lower costs, and accelerate innovation using AWS.” https://aws.amazon.com/solutions/case-studies/.

Bohn, Dieter. 2019. “Amazon says 100 million alexa devices have been sold — what’s next?” *The Verge.* https://www.theverge.com/2019/1/4/18168565/amazon-alexa-devices-how-many-sold-number-100-million-dave-limp.

Fortune. 2021. *Fortune 500*, 2021. Search. Available at: https://fortune.com/fortune500/2021/search/.

Hardesty, Larry. 2021. “How AWS scientists help create the NFL’s Next Gen Stats.” *Amazon Science News*, February 5, 2021. https://www.amazon.science/latest-news/how-aws-scientists-help-create-the-nfls-next-gen-stats.

Maguire, James. 2022. "Top Performing Artificial Intelligence Companies of 2022." *Datamation.* https://www.datamation.com/artificial-intelligence/ai-companies/.

Wired.com, December, 2019. “With Deep Learning, Disney Sorts Through a Universe of Content.” Available at: https://www.wired.com/wiredinsider/2019/12/deep-learning-disney-sorts-universe-content/.