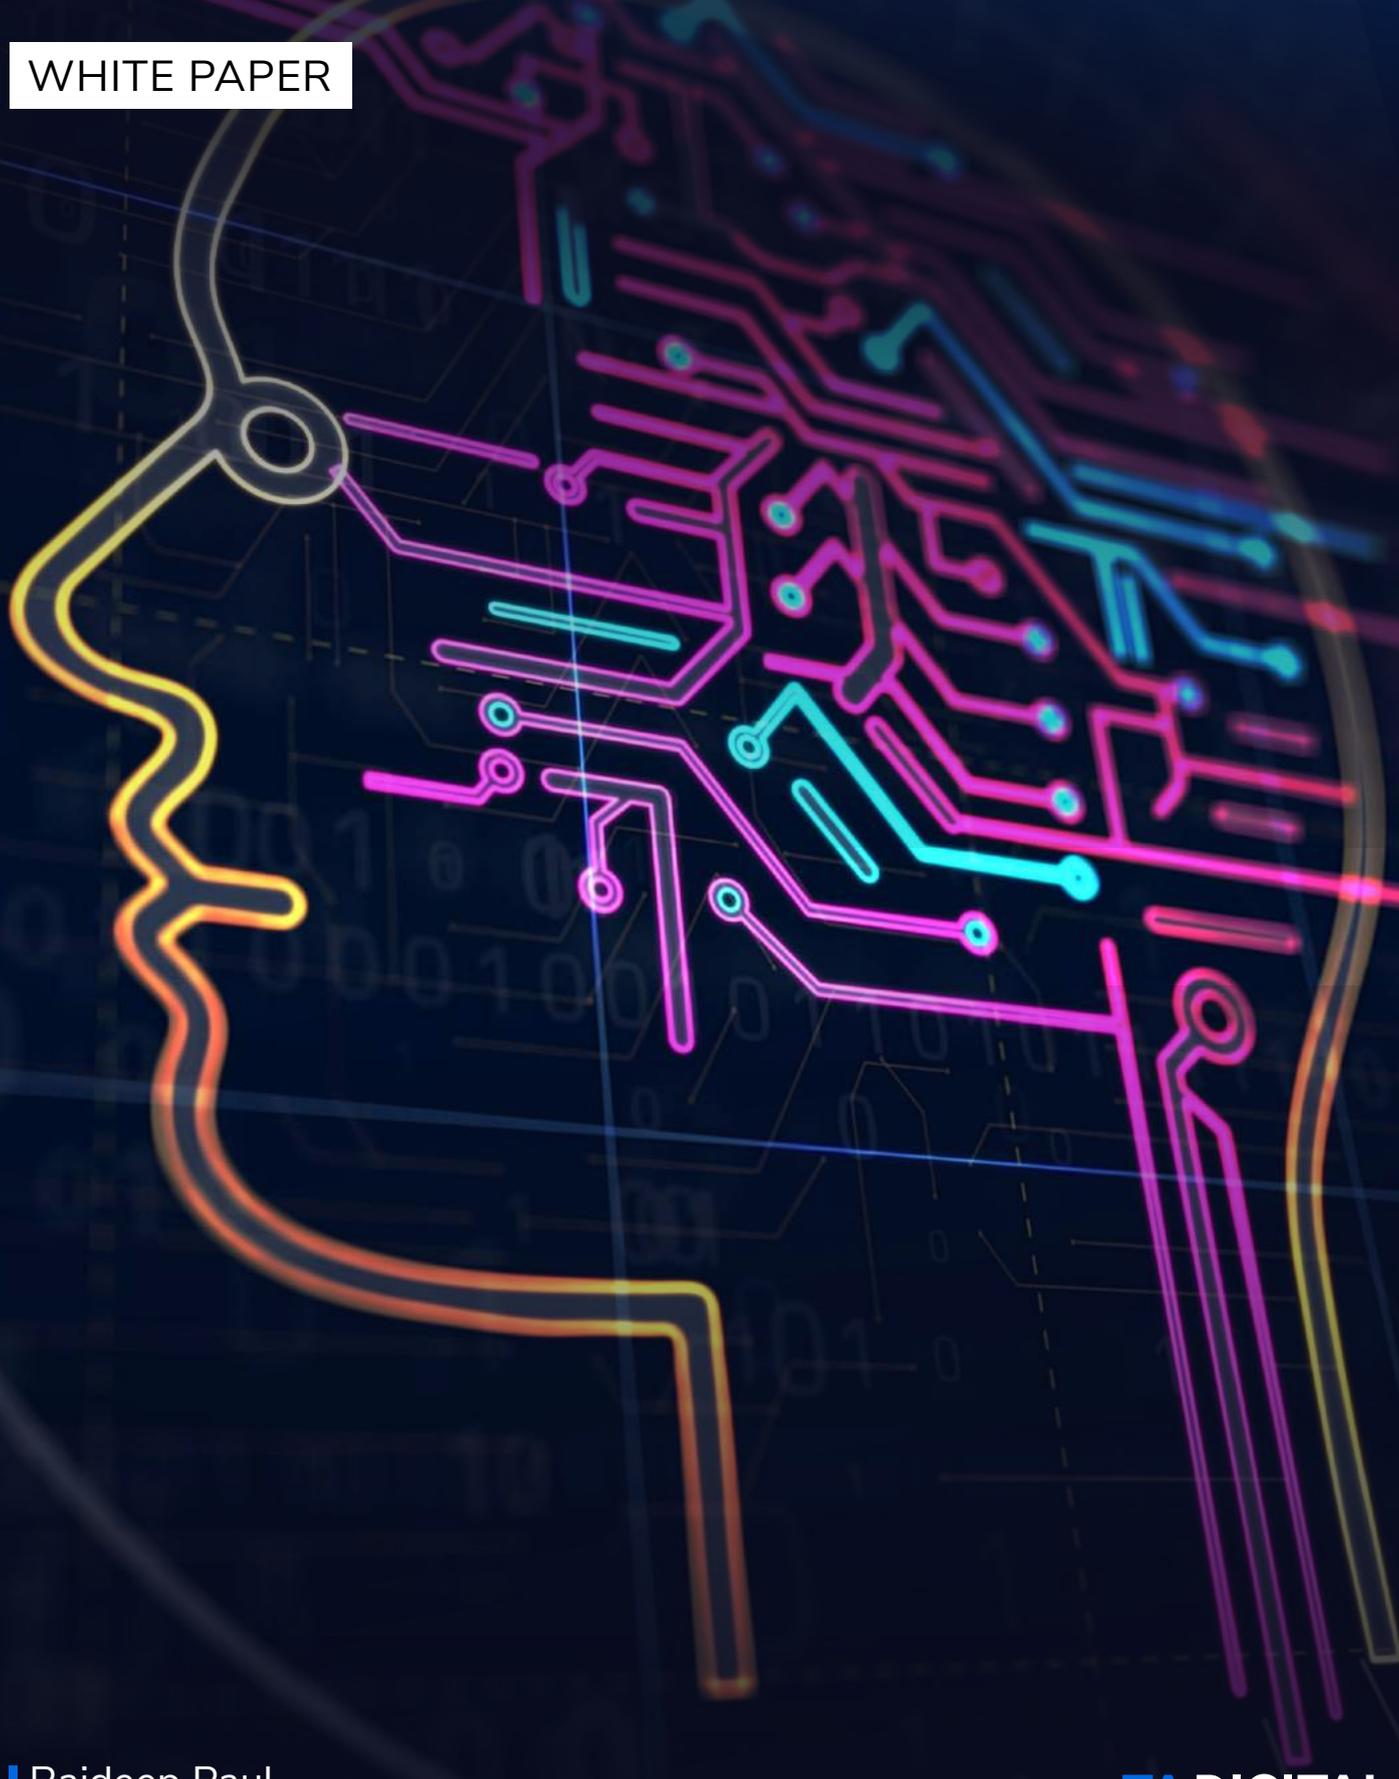


How To Leverage Machine Learning For Content Management

WHITE PAPER



Rajdeep Paul
Associate Technical Lead, Drupal

TA DIGITAL
Intelligent Transformation

TABLE OF CONTENT

| | |
|---|----|
| Introduction | 2 |
| What is machine learning? | 2 |
| Example uses of machine learning | 3 |
| What type of content can be used for machine learning? | 3 |
| Why do we need to leverage content for ML? | 5 |
| Use cases | 5 |
| 1. Analyzing and predicting customer churn | 5 |
| 2. Incorporating bots for improved customer experiences | 5 |
| 3. Real-time personalized advertising | 5 |
| 4. Improved audience insights | 5 |
| 5. Demonstrating marketing ROI | 6 |
| 6. Sentiment analysis | 6 |
| 7. Computer vision for product recognition | 6 |
| | |
| Approach to consume content | 7 |
| 1. Word embeddings | 7 |
| 2. Learn from the best, transfer to the rest | 7 |
| 3. Learn from examples | 7 |
| 4. HitL | 7 |
| | |
| How UX is important? | 8 |
| Techniques to develop ML | 9 |
| 1. Deep learning | 9 |
| 2. Classification | 10 |
| 3. Regression | 10 |
| 4. Supervised learning | 10 |
| | |
| How these techniques converge to ML? | 11 |
| Conclusion | 11 |

INTRODUCTION

At some point, we have all been amazed by our app recommendations. For example, when we book a cab using an app like Uber, it auto-populates the destinations. These recommendations change depending on parameters such as the booking time and the booking location. This is a simple case of machine learning, which analyzes our previous behavior and, based on our historical choices, provides suggestions.

Marketing teams see immense potential in machine learning and its applications. From content management to personalization and analytics, machine learning can significantly help businesses in a number of ways. It can help businesses streamline day-to-day activities by automating repetitive tasks such as audience targeting, content classification, reporting, etc.

While many IT firms are investing in machine learning, they often don't define what these techniques are. The risks are ever-present, and embedding the wrong ML technique often returns different outputs, which is undesired by the organization. In this whitepaper, we will explore machine learning techniques for content management. But before exploring these techniques, let's first understand what machine learning is and how it can leverage content for automation.

WHAT IS MACHINE LEARNING?

Machine learning can be defined as a set of techniques that enables technology to learn from data and make intelligent decisions.

CMS solutions and ML generally complement each other. Every great customer experience is built on two pillars - content and data. The role of machine learning is to use that data and create multiple automated processes.

ML is a specific set of techniques that helps machines to learn from data and make predictions. A lot of people think that at some point in time, the machine will reach a higher level of intelligence (basically, more than human) and rule the world. Fortunately, that's not an immediate risk.

We can find some systems which predict our interests based on name, and some conclude that 'man' is to 'IT Analyst' as 'woman' is to 'housemaker.' Some systems are predicting terrorists based on facial features. This biased decision-making is concerning and comes from flaws with the methodology which enables these decisions.

Our past and present data is the fuel to this prediction for these machines. Machines are only using these data to give its output. It is hard for ML to carry out these independently without flaws. People must be aware of this fact if we want ML to affect our lives in productive and ethical ways.

EXAMPLE USES OF MACHINE LEARNING

1. Netflix users are already experiencing machine learning algorithms, which are analyzing users' viewing habits and comparing them to those of millions of other viewers to suggest what the user might want to binge-watch next.

2. Today's smart cars use ML to adjust drivers' preferences, such as seat position and temperature, provide advice about road conditions, detect and fix vehicle issues.

3. AI and machine learning also power chatbots, which converse with users via a fast, friendly interface to get things done. For example:

- Pizza Hut's chatbots let customers place orders and enquire about dietary information, delivery areas, and more via Facebook Messenger and Twitter.
- FedEx is using Amazon.com APIs to build an app that lets customers ship packages by saying, "Alexa, I want to ship a package."
- Wynn Hotels plans to outfit nearly 5,000 hotel rooms with Amazon's Echo device, which will enable guests to ask Alexa for the room, hotel, and other information.

WHAT TYPE OF CONTENT CAN BE USED FOR MACHINE LEARNING?

Big corporations like Google, Facebook, and Amazon are focusing on 'perceptual AI', while some others are on 'Enterprise AI'. This 'perceptual AI' deals with a huge amount of data uploaded by the public and uses a deep learning approach for ML. But these data are not useful for 'Enterprise AI'.

These organizations that are working in Enterprise AI must be smart about working with smaller amounts of data. They have a lot of entries of customers who yield these datasets, and there are specific techniques (transfer learning, few-shot learning, and human-in-the-loop machine learning) that are useful.

Organizations that are willing to work on ML must use their existing data, which comes from their own business and treat that as a starting point. The data comprises of customers' content which yielded as the customers use their products. And these contents have a lot of varieties and huge as it comes from multiple industries. ML teams treat these contents as data and move on.

Words that are compiled to create blog posts, news articles, or product info need to be understood by machines. But machines only understand numbers, not words. This made ML techniques to come up numeric version of the words that can create content categories, may find similar content, or give suggestions of the same content. This can solve multiple content challenges for customers to make

their processes more effective. Tasks like content classification or content recommendation, automated by ML, will help the marketing team a lot. Using pre-trained word embeddings for natural language processing is one of the core techniques of machine language. ML teams need to look for new techniques and incorporate them into products to make solutions easier and faster.

Due to multiple types of machine learning techniques, the marketing team must understand these are incorporated into solutions, and eventually, how these will solve their problems. In a broader view, these techniques must treat content as data and solve customer problems.

Generic machine learning algorithms require labeled data to learn. But, most of the customers do not have labeled data. To overcome this challenge, there are some machine learning techniques like transfer learning, few-shot learning, and human-in-the-loop learning.

Transfer learning is to apply learnings from a large dataset and apply them to a small dataset. For example, the algorithm or learnings coming from pre-trained word embeddings will be applied to smaller datasets like customers' content.

Few-shot learning is like transfer learning, where the former one learns for a very few examples, while the latter depends on the quantity of data available.

Humans in the Loop (HitL) learning is to use data from review or ratings, basically using crowdsourced labeled data. The product interface plays a crucial role to build up strategies that involve a human element in the data training process, which makes increment inaccuracy of learnings.



WHY DO WE NEED TO LEVERAGE CONTENT FOR ML?

Organizations that are likely to leverage machine learning techniques must understand the available content and data. They must also plan the goals they want to achieve with that information. Goals can be like easily searchable content, user-based personalization, and lots of insights.

Machine learning can help in reducing tedious marketing tasks like content tagging, but it needs an upfront human element of labeling data to help inform machine learning algorithms. The organization needs to find ways for data to be leveraged upfront. For example, features like ratings or feedback option as a part of the product review, or tracking engagement data on content assets, will give crucial insights to marketers about purchasing behaviors and content consumption with less manual efforts.

Another example is to improve user personalized content suggestions. The marketing team would typically like to recommend content to customers based on their previous consumed content. A machine learning system should be able to read and identify both structured and unstructured content data and use a trained ML model. This will make customers avoid heavy classification or tagging so they can get started with suggestions easily.

USE CASES

1. Analyzing and predicting customer churn

With insight into the relevant data, marketers can identify how a certain segment of the market will behave in the future. But analyzing this data for millions of customers is humanly impossible. ML teams can develop an algorithm that analyzes mobile customer behavior to detect the most loyal users and predict those that are likely to move. With this kind of information, marketers can take steps to deepen customer engagement or invest in retaining specific customer segments.

2. Incorporating bots for improved customer experiences

Marketers frequently turn to bots, integrated with popular messaging apps such as Facebook Messenger or Kik, to automatically answer questions about post-purchase requests, reducing the time spent trying to track down answers to FAQs.

3. Real-time personalized advertising

With a social content marketing platform, marketers can tap into the nearly two billion posts shared on social media every day to better personalize their marketing campaigns.

4. Improved audience insights

Service providers make the mistake of clubbing their customers into one category. Segmenting

customers into groups can help improve customer engagement and generate more accurate data to use when building a customer persona to help improve personalization and target people accordingly. Software that uses machine learning, like Affinio, can help with this.

5. Demonstrating marketing ROI

Marketers must also understand and communicate the ROI of each new tool used. This is where AI can help. Brands and sports teams are turning to GumGum for its computer vision technology to determine the value of their investments in sports. Each logo displayed on TV and social media is captured and analyzed, producing a more comprehensive and accurate media valuation of their sponsorships.

6. Sentiment analysis

When a customer sends an email or direct message, an organization needs to know how they're feeling to respond properly. Machine learning can analyze text to determine whether the sentiment is positive or negative. ML system can identify social media content and alert the organization about negative content, which helps them to take preventive action. AI can also identify people happy with the products/services to help companies find social influencers and brand ambassadors.

7. Computer vision for product recognition

Software such as GumGum can help brands recognize their products in images and videos online. Such technology could help companies find millions of social posts associated with the brand. It would be nearly impossible for a human to complete this task.



APPROACH TO CONSUME CONTENT

1. Word embeddings

In 2003, Bengio proposed an idea to learn representations of words that capture semantic meaning. In 2013, Google established this with the Word2Vec algorithm. Word2Vec and other recent approaches (like GloVe from Stanford) learn from huge datasets like Google news articles or Wiki. The representations they are learning after analyzing through all of that text is numeric vectors of multiple dimensions. So, a single word representation is a long list of numbers. The fun of it is that the mathematical relationship between those vectors manages to capture the semantic relationship between the words. The popular example is king - man + woman = queen. Once these representations have been learned from a massive dataset, they can then be used in other tasks like classifying content.

2. Learn from the best, transfer to the rest

This ML technique is about knowledge gained on one task being reused in solving another task. For example, the use of pre-trained word embeddings. We can take the word embeddings trained by Google or Stanford and transfer them to be used in our tasks.

One such task is similarity-based content recommendations. If we have numeric representations of our content that capture semantics, then we automatically have a measure of similarity between pieces of content. Even if two pieces of content are talking about the same topic but using different words, they're going to be identified as similar, thanks to the character of those representations.

3. Learn from examples

We have already heard the phrase 'data is the new oil'. However, someone took it a step further at the 2017 O'Reilly's AI conference by proposing that 'labeled data is that the new 'new oil'. For classification tasks, few-shot learning is an approach that stands in contrast to plain deep learning approaches because deep learning requires enormous quantities of labeled data.

The key to having the ability to find out from only a few examples has great representations of our data. That is why transfer learning and few-shot learning often go together. We transfer the knowledge from some previous tasks and use it to make representations of the next data. Just labeling one or two examples then allows all the others to be labeled automatically. This is often our approach to automated content tagging.

4. HitL

A solution to the matter of lack of labeled training data is to urge humans to label our data. This is often called human-in-the-loop (HitL) ML, a term which was coined by the founding father of a corporation called CrowdFlower, which focuses on a crowdsourced approach to the present technique. Another company, Mighty AI, is focused specifically on training data for autonomous

vehicles. Anyone with an iPhone can earn a couple of cents a pass labeling pedestrian, lamp posts, parked cars, etc. in images.

Humans are often made a part of the loop in other, less straight-forward ways than labeling entire training sets to feed into ML algorithms. Any application or service that explicitly asks users for feedback within the sort of ratings - Netflix movie ratings, for instance - are often thought of as employing HitL. The organization StitchFix, which provides a clothing service where they send customers a daily 'fix' of clothing items selected by a stylist, gets tons of upfront data from users by requesting them to rate styles through a series of photos. The more data they will get from their users upfront, the less they need to infer through purchasing behavior. This is often important to the success of their service because without HitL initial, 'fixes' would stand a poor chance of being purchased. Companies that use HitL understand that the UI they present to the human in their loop is of important importance.

HOW UX IS IMPORTANT?

In the current wave of pleasure over ML, tons of recommendations are being offered to companies on the way to incorporate these techniques to enhance their business. Counting on who's offering it, the recommendation differs significantly. Those within the business of coaching and recruiting data scientists will say that a company needs many data scientists. In contrast, those within the business of selling machine learning as a service (MLaaS) solutions will speak the opposite. The truth lies somewhere in between.

It is important to possess people with skills to border your business' problems as data science or machine learning problems and confirm the info needed to unravel them is out there. Merely getting your engineers to feed masses of knowledge into Amazon or Google's MLaaS isn't getting to achieve considerably. On the opposite hand, data scientists alone probably can't do everything. If you're building a product, only one or two data scientists working with engineers and UX professionals are going to be much more useful than 10 data scientists. The proper mix depends on what you're trying to accomplish.

We are using ML to reinforce our SaaS offerings and have built a team focused specifically on this area. It includes data scientists, data engineers, front-end engineers, and back-end engineers. The team also works very closely with the UX team. Where we're using HitL, UX is significant to making sure we get the info we'd like to support our learning algorithms to form them as accurately as possible. Other efforts don't entail a HitL aspect but require skilled engineers to make sure that services delivering ML predictions are performant and scalable.

We have people that are conversant in the kinds of solutions that machine learning research has developed (many of which are available in open source libraries) and, therefore, the sorts of problems

to which they're best applied. This expertise, including strong engineering and UX skills, is what we'd like to execute our ML strategy. If we didn't have a well-thought-out strategy on the way to play to our strengths, make use of publicly available datasets and open source libraries, and incorporate the opposite necessary technical functions in our efforts, an AI Ph.D. would struggle to feature value.

Without even realizing it, people interact with machine learning systems daily, surfacing an entirely new set of issues for consumer-facing brands. Google's 'did you mean?' feature was built to assist users within the search process, but in practice, there's no denying it can get a touch creepy and feel invasive. On the opposite hand, if you're given a Netflix or Spotify recommendation that's off base, users get frustrated and expect a better understanding of their preferences. Today's companies are forced to toe the road between helpful and creepy.

Companies also are running into problems around AI's inability to elucidate their suggestions and proposals to users. For instance, Alexa can't explain her reasoning for recommending a replacement product or service. Although we would like to trust the devices that give us answers, at an equivalent time, we would like to know where they are available. Solving these user experience challenges is what is going to separate the leaders from the laggards within the machine learning realm.

TECHNIQUES TO DEVELOP ML

1. Deep learning

This idea was the seed for more complicated 'Artificial Neural Networks', where hidden layers of neurons between the input layer and, therefore, the output enabled great flexibility in tackling different types of problems. Under the category of 'Deep Learning', we got success in resolving problems like image classification and speech recognition, by using direct descendants of the lowly perceptron.

Developments along this fork within the path of AI research moved away not just from the methodologies traditionally employed (using rule-based learning), but also from the broader goals of these earlier systems. It had been not about producing something that would think sort of a human but was instead merely about solving practical problems.

The process of DL suffered its specific setbacks that led to years during which the very mention of the term 'neural network' was almost banned. But breakthroughs involving the smart application of some tricks from calculus eventually put it back on target as a theoretically sound approach to answering questions with data. They required tons of knowledge and tons of computing resources to coach to A level of accuracy competitive thereupon of other sorts of algorithms. Over the last few decades, computing has not only been about higher processing power (cloud computing, GPUs) but also about big data. And with cloud computing and various technologies enabling massive parallelization of algorithms, training on these massive datasets is vastly sped up.

Deep learning has seen tremendous success in the last decade, particularly within the areas of image classification and speech recognition. It's what Facebook uses to spot faces in photos, what Siri uses to know what we are saying, what Google image search uses to point out images associated with search terms.

2. Classification

Specifically, it had been focused on classification problems, learning to classify examples as belonging to at least one class or another. One very early and successful application was training a neural network to acknowledge handwritten digits. Here, each possible digit from 0-9 may be a class, and therefore the network got to take a picture of a digit just like the one pictured below as input (raw pixel data) and output the right class (digit).

The way it works is that the network is shown thousands and thousands of handwritten digits and told what all of them are, and it must learn the features that distinguish all. Once it's been trained to try to this, when it sees a fresh image of a handwritten digit, it makes a prediction about which class it belongs to (which digit it's.)

3. Regression

Meanwhile, over within the world of statistics, statisticians had, for many years, been using techniques like rectilinear regression for creating predictions about real-valued quantities (i.e., estimating numerical values as against categories) of a variable from one or more independent variables. For instance, trying to estimate the sales of a product as a function of advertising spend.

Of course, this approach is often wont to make predictions supported more variables. For instance, predicting sales supported advertising spend per channel, per market segment, and at different times of the day. This report can help the marketing team make decisions that increase sales while reducing spending.

4. Supervised learning

Regression and classification are both about taking labeled data — examples where the answer is known and using it to develop a model that can then be used to make predictions about new data. This general method is mentioned as supervised learning. The variable is typically mentioned because of the result, and thus the independent variables are mentioned as predictors or features.



Deep Learning



Classification



Regression



Supervised Learning

HOW THESE TECHNIQUES CONVERGE TO ML?

The sector of statistics had also developed methods for classifying while only deep learning (DL) researchers performing on classification problems and the statisticians performing on regression problems. The DL people found themselves reinventing techniques from statistics and data processing as they worked to enhance their neural network algorithms.

Gradually, the overlap in these fields became a field, that of machine learning. It is about defining some loss function, which outputs a measure of how wrong a given statistical model is about data and using optimization algorithms to regulate the model so on minimize that loss (cost or error). More algorithms are invented for solving both classification and regression problems. Samples include 'tree-based' methods, whereby the choice to classify an example (or estimate its real-valued output as lying within a specific range) depends on the answers to questions asked of the predictors at each branching point within the tree.

For much of the earlier varied approaches were on a comparatively level playing field - some were better for problem types but worse for others. Then some developments indirectly associated with this field of research led to the DL approach, i.e., using sophisticated neural networks, really beginning.

CONCLUSION

Keeping in mind ML's target of developing human intelligence, we use the term Strong AI to differentiate it from the (weak) AI that the majority of researchers are performing lately. An honest rule of thumb is, once we hear the term Artificial Intelligence, and it's not been qualified by the term 'strong' or 'weak', it's presumably weak AI. Nobody has solved strong AI, and then it's only ever talked about theoretically. Even in cases where we mention "an AI," meaning "an artificially intelligent entity," like Siri or Alexa, we are still talking about weak AI; it's just that it's been packaged into something that we interact within human-like ways.

Speech recognition, where Siri and Alexa find out the words we are saying supported sounds, is one ML task that then feeds into the subsequent task of tongue understanding (NLU), i.e., deciding the intent of the utterance to understand the way to respond. The tasks are combined to offer the illusion of an artificially intelligent entity. But make no mistake; Siri and Alexa are an extended way from 'waking up' as sentient beings.

Just because the present state of AI isn't producing robot overlords doesn't mean it isn't doing amazing things. Several of the ways AI is being put to use to realize spectacular successes during a wide selection of various fields, solving real-world problems.

TA Digital is the only global boutique agency that delivers the “best of both worlds” to clients seeking to achieve organizational success through digital transformation. Unlike smaller, regional agencies that lack the ability to scale or large organizations that succumb to a quantity-over-quality approach, we offer resource diversity while also providing meticulous attention to the details that enable strategic success.

Over the past 20 years, TA Digital has positioned clients to achieve digital maturity by focusing on data, customer-centricity and exponential return on investment; by melding exceptional user experience and data-driven methodologies with artificial intelligence and machine learning, we enable digital transformations that intelligently build upon the strategies we set into motion. We are known as a global leader that assists marketing and technology executives in understanding the digital ecosystem while identifying cultural and operational gaps within their business-ultimately ushering organizations toward a more mature model and profitable digital landscape.

Recognized in 2013, 2014, 2015, and 2019 Inc. 5000 list as one of the most successful technology companies in the United States, TA Digital is pleased also to share high-level strategic partnerships with world class digital experience platform companies like Adobe, SAP and Salesforce and possess global partnerships with industry leaders such as Sitecore, Episerver, Elastic Path, BigCommerce, AWS, Azure and Coveo.



www.tadigital.com | sales@tadigital.com

TA Digital™ is a registered trademark of TechAspect Solutions, Inc.

© 2020. No part of this document may be distributed, reproduced or posted without our written permission.