


# The Marketer's Field Guide to Machine Learning

The Past, the Present, and the Future of Martech



---

## AI: The Future of the Marketing Tech Stack

Marketers need to ensure they truly understand the underlying data behind the automation and AI functions, which means data literacy is key for today's marketer. (...) The best marketers will be able to leverage AI data to provide the kind of creative insight that turns data-crunching into campaign gold by delivering the authenticity that consumers demand today.

— Marc Keating, B2B Marketing

January 24, 2017

## Table of Contents

<b>Foreword</b>	<b>4</b>	<b>Machine Learning &amp; Marketing</b>	<b>13</b>
<b>Machine Learning at a Glance</b>	<b>5</b>	Benefits & Challenges	
The History of Machine Learning		Meet the Machine Learning & Marketing Innovators	
Types of Machine Learning		Acquisio is Machine Learning	
Artificial Neural Networks		BBM Becomes Acquisio Turing™	
Deep Learning (AKA Deep Neural Networks)		Results from our Latest Performance Report	
Artificial Intelligence		Acquisio Uses Machine Learning to Solve Small Business Marketing	
Inside Machine Learning		<b>The Future of Machine Learning</b>	<b>24</b>
Supervised Learning		Learning to Trust Machines	
Unsupervised Learning		Explainability & Fairness	
Reinforcement Learning		Ethics & Security	
		Making AI More Human	
		About Acquisio	

## Foreword

The promise and potential of Artificial Intelligence (AI) over the last 50 years has been met with so little clarity, whether in marketing or everyday use.

Part of the confusion has been fed by popular science-fiction movies and TV shows delivering preconceived notions of lifelike, ill-intentioned robots. The idea of human mimicry is unclear and because people don't understand it, they're afraid of it. However in the business world real companies are using intelligent technology at scale today and AI is no longer an imaginary idea of the past. Being afraid is no longer an option and embracing AI has become a necessary means for business success.

The current AI capabilities we have as marketers and business owners has complemented our processes rather than disrupted our workflow. We are able to solve more complex problems and scenarios with greater precision while using more data, faster than ever before. The results of machine learning martech have been unprecedented.

The biggest barrier to AI right now is adoption. A lack of clarity in what AI and machine learning can do breeds uncertainty; however, a computer is just a computer until you use it. AI is just a tool. Just like you can't email your boss using your air conditioner, all types of technology have a specific purpose and limitations that go along with it. Machine learning is not without limitation nor controversy.

This eBook will educate business leaders and extraordinary marketers on the past, the present, and the future of machine learning. It will look at the artificially intelligent world of marketing tools available. It will uncover a deeper dialogue around our ability to trust machines. Most importantly, it will empower the business leaders of tomorrow to solve specific problems using the magic of artificial intelligence.

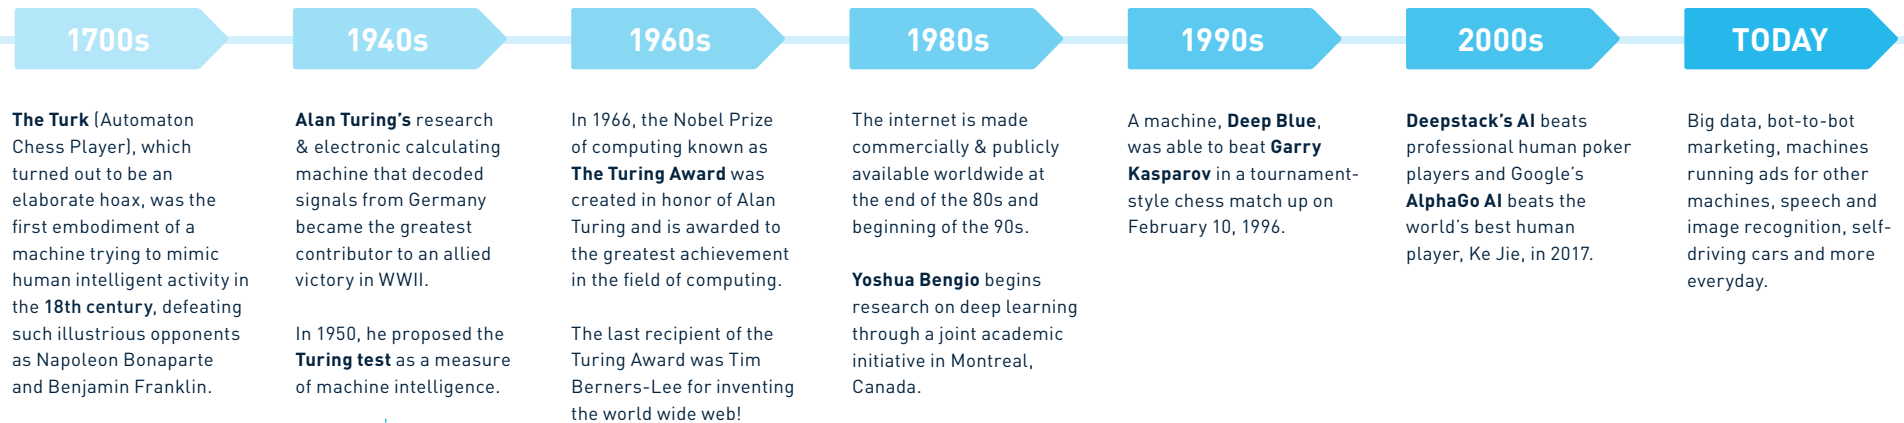


The Turk






## The History of Machine Learning



### Alan Turing 1912 — 1954

Alan Turing was a mathematician, a logician, a cryptanalyst and an early computer scientist who became a war hero. In 1950, Alan Turing proposed the Turing test as a measure of machine intelligence four years before he died. The test consists of a human observer communicating with another human and a machine, and is tasked with telling them apart. If the machine can “trick” the observer into thinking it is a human in this natural language conversation, it passes the Turing test. Since its introduction the test has been the subject of extensive debate about its applicability in defining Artificial Intelligence.



---

## How AI Allows Intelligent Marketers to Market More Intelligently

The key to customer research and data analysis, whether you're doing it manually or relying on technology, lies in asking the right questions. Using AI and CC [contextual computing] can speed up the process of learning about your customers and help optimize your efforts.

— Jennifer Steckel Elliot, Marketing Tech News

February 9, 2017

## Types of Machine Learning



### Artificial Neural Networks

The last few years have seen an explosive growth in Machine Learning technologies. The availability of large datasets accumulated in the years of Big Data finally found its match in increased computing power and advances in **Artificial Neural Network** design that brought successful Machine Learning solutions to problems previously deemed too complex for machine intelligence.



### Deep Learning (AKA Deep Neural Networks)

Techniques such as **Deep Learning** (or Deep Neural Networks) make use of huge datasets to achieve near human (or better) performance in facial recognition, image classification, voice recognition, or mastery of such games as chess, **poker**, and **go**.



### Artificial Intelligence

Artificial Neural Networks and Deep Learning have been heralded as great examples of **Artificial Intelligence (AI)**, intelligence exhibited by machines. But viewed from the angle of the Turing test these “narrow AI” applications would not fool any human; they can only exhibit the behavior of their chosen specialty. An ML system that excels at chess cannot recognize faces any better than a voice recognition system can play go. Furthermore, these systems suffer from “catastrophic forgetting”. Once they master a specific task, learning a new one would mean they completely forget the previous skill.

To date, no system even comes close to the capabilities of AI in Hollywood movies. Researchers may have to develop a different approach to achieve this type of general AI. All the success stories so far share a common theme - they are directed by humans. The machines are setup to solve problems specified by humans and the data is provided by humans (including, in many cases, the correct label of the outcome). Humans learn a lot by observing the world around them, studying cause and effect relationships (eg. objects fall towards the ground if unsupported), and coming up with generalizations from those learnings. Recent research from the University of California, Berkeley explored a different learning paradigm by **injecting curiosity into the AI learning process**.

Machines are successful at many of these problems because they can consume data at a rate humans simply can't, and then process that data to uncover patterns that an average human couldn't even detect. This makes them ideal for augmenting our human activities, becoming the tool that takes over the boring and mundane tasks of number crunching and pattern discovery while leaving the creative work to us.

ML is becoming part of everyday life whether we know it or not. Most of us have been exposed to machine learning in our day to day activities in some form or another (eg. product recommendations based upon items in the shopping basket, automated email response messages, internet advertising optimization, etc). Hardly a day goes by without some company announcing the fruits of their ML efforts. Whether they are labeled ML or AI, they are here to stay. **The real question is what our relationship with machine intelligence looks like now and how we are going to cooperate with the thinking machine in the future.**

**One of the most commonly used definitions of Machine Learning describes it as “the technology that gives the ability to computers to learn without being explicitly programmed.”**



## **Data Virtuality**

The combination of agile data integration like logical data warehouse and sophisticated machine learning algorithms will be the key for companies to be smarter and quicker than their competitors. The optimized, insightful, real-time models for data visualization and content analysis will revolutionize sales and marketing across industries.

— Nick Golovin, PhD, CEO & Founder of Data Virtuality

June 2017

## Inside Machine Learning

Machine Learning (ML) was born out of the research into Artificial Intelligence and the desire for machines to learn from data. Much of its original influence came from statistics and probability theory. From gathering and curating data to the selection of appropriate modeling techniques and the tuning of their parameters, humans are involved.

Instead of being explicitly programmed like the computers we know, machine learning technologies are tuned and directed by humans for a specific purpose. Software developers are also needed to integrate machine learning APIs into an organization's production systems and to maintain the continuous-learning framework.

The machine learning technique used to solve a problem depends on the question being asked.

Machine Learning problems can be classified into **three** broad categories:


- **Supervised Learning:** the computer is given curated, labeled data where both the input and the output are available and the task is to learn the mapping between the known inputs and outputs.
- **Unsupervised Learning:** unlike in supervised learning, here no labels are given for the input data. The computer's goal is to find some structure in the data.
- **Reinforcement Learning:** while reinforcement learning does not involve output labels either, it uses a rewards and punishments mechanism to teach the computer how to perform a certain goal in a dynamic environment.

At Acquisio we use both supervised and unsupervised machine learning.

## Supervised Learning

Supervised learning relies on curated, labeled data where the machine is learning from the available input-output pairs and then applies that learning to new, unlabeled inputs to predict the output. If the output is chosen from a finite number of outcomes (discrete value) the problem is a classification problem. If the output can take on any potential number (continuous value) that makes it a regression problem.


In classification, the inputs are divided into two or more classes and the computer must learn how to assign inputs to the right output class. For example, one can measure different aspects of their PPC account performance and plot each observation as shown in the figure below.



Here poorly-performing account observations are shown in orange. Well-performing accounts are shown in blue. The task is to learn how to tell them apart, which in this case is just a linear separation shown by the line. New observations above this line are declared poor performers, those that fall below the line are pronounced good performers.

Other examples of classification include separating real news from fake news, or classifying images based upon their content. A well-known early example of image classification includes **Google's AI learning** to detect videos that have cats in them.

In regression, the goal is to derive the best estimate of a continuous value, such as stock prices, daily high and low temperatures, or potential daily ad spend in search advertising on a given day. The data consists of observed inputs and the corresponding output value at those input observations. The model learned by the computer is then applied to previously unseen input observations to forecast the desired output.



The figure above shows the learned model (orange curve) of a house price - square footage relationship based upon the available observations (blue dots).

## Unsupervised Learning

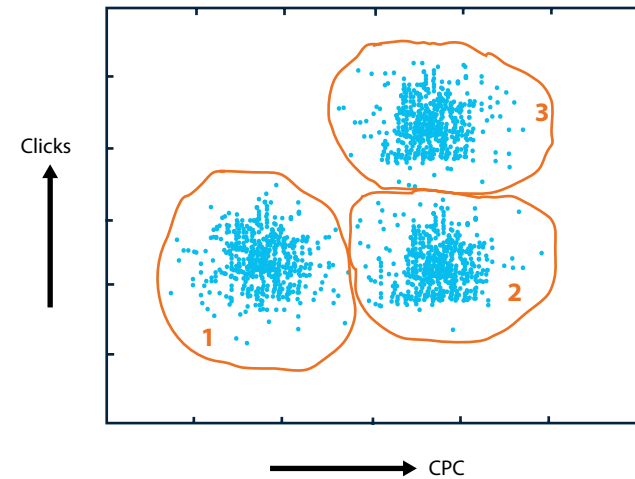
In an unsupervised learning problem no ground truth is available; the task is to try to discover some structure in the data without being directed on how to find it.

In the image below, the goal is to find similarly performing advertising accounts based upon the number of clicks they receive in a month and the average cost-per-click (CPC) during that period. The dots in the diagram represent each account at their observed average CPC and number of clicks.

There appears to be 3 clusters of these accounts in the data:

1. Accounts with low CPC and low click volume
2. Accounts with high CPC and low click volume
3. Accounts with high CPC and high click volume

The clustering method only finds the structure, the descriptive labels are given to them after the fact.



Other examples of unsupervised learning include Google News where many published articles are grouped together based upon their content and then presented to the users.

## Reinforcement Learning

In reinforcement learning there are no correct input/output pairs. Instead, the machine learns by performing certain steps (making decisions at each point) until a final stage is reached. At this point it is presented with either a reward or a punishment. This reward/punishment is then split among the steps that it took along the way so as to learn from experience what the rewarding strategies are.


Lucy, the cognitive companion featured in the next section, uses a star based rating system to collect feedback on its accuracy in responses - this is an example of a rewards based form of reinforcement learning in marketing.

Reinforcement learning is not unlike the lab rat situation where the rat is taught the path through the maze that leads it to a block of cheese. These techniques are used in AI systems that can play Atari video games, chess, go, etc. There are even attempts at training a self-driving car this way by letting the system observe human drivers to learn the right strategy from them.

Within each of the Machine Learning categories there are a number of different techniques, each with their own strengths and weaknesses. The selection of the right method is still mainly done by data analysts in a manual process of trial and error. The same process extends to tuning the model parameters once the desired model is selected. There are a number of initiatives to try and automate some of this process, including **using AI to train AI models**.



# Machine Learning & Marketing



---

## 10 Things Marketers Need to Know About AI

In the next few years, we're going to see AI and machine learning become part of the core 'fabric' of software platforms in the marketing industry.

— Dharmesh Shah, co-founder and CTO of HubSpot

February 13, 2017



## Benefits & Challenges

AI and machine learning in marketing will become what technology is already to modern life - an acceptable component powering our daily lives in the background, with built-in consequences perceivably outweighed by the accompanying benefits. With all talk of job security aside, there are in fact a laundry-list of benefits to enhancing marketing systems with continuously-improving algorithms.

Generally, marketers profit from beneficial trends fundamental to machine learning advancements like increased data collection and analysis capability leading to a greater number of insights – be it about brand-consumer interaction or post-campaign analytics. Marketers stand to gain both knowledge and efficiency broadly. Additional potential benefits that ML martech tools promise marketers range from customer experience to channel tech.

Today marketers are using machine learning and AI to unlock deeper customer insights, allowing them to deliver personalized one-to-one messaging and in return achieve greater campaign ROI. Injecting AI into backbone platforms like CRMs allows marketers to personalize the best buyer discount to offer automatically, for example. In the ongoing quest to better understand and cater to consumers, marketers at the very least reduce the time spent analyzing customer behavior.

Marketers who work with channel martech providers already will have an easier time transitioning to additional technological advancements. From email campaigns to cross channel distribution, marketers are already embracing the intelligent tools around them. Consider increasing open rates on email campaigns by automatically determining the best time of day to send an email, based on an individual's behavior, explains **Marketing Technology News**.

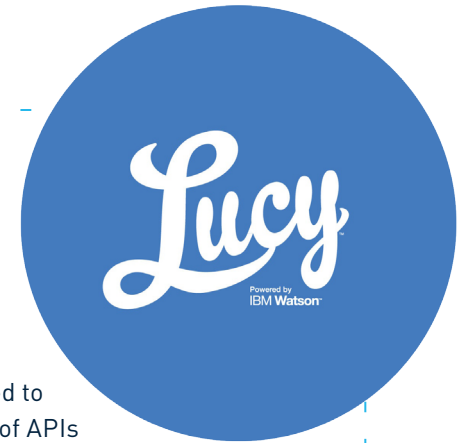
With infinite marketing possibility, **highlights** in the AI martech landscape include **autonomous audience targeting, personalization, media buying, cross channel execution, testing, optimization and retargeting**. Not to mention these massive marketing silos playing out effortlessly at greater speeds with more efficiency than ever before.

Automating data collection and enhancing analysis around marketing activity pre and post campaign creates a more intelligent marketing tool kit and customer ecosystem, ripe for results. And with a greater volume of data and a greater depth in analysis lies the opportunity for increased velocity in everything marketers do. Marketers are already making leaps and bounds alongside intelligent innovation in martech. From email marketing to search engine marketing, those who can leverage the emerging pieces of the ML and AI marketing puzzle will have an advantage over their competition.

However with great opportunity often lies uncertainty and boundary. One of the biggest challenges in adopting machine learning technology in marketing is that each version will solve a different problem and therefore will look and behave differently - each will have its own set of challenges.

Practical challenges arise for marketers when they think literally about how they would purchase and adopt such a system. Marketers have never done this before. Despite a general enthusiasm for AI systems to solve business challenges, marketers have never completed RFPs for a “cognitive agent,” and there's no mark in pre-existing budgets **says** Scott Litman, the Managing Director of **Equals 3**. Early adopters of AI will trail-blaze, having no experience purchasing machine learning technology, setting new precedent and shaping the vision of what such a technology will look like. And yet AI and machine learning algorithms are in use with systems we interact with everyday like Google, Alexa and Facebook.

To better understand how machine learning is benefiting the lives of marketers it's useful to analyze prominent examples. Equals 3's Lucy is helping marketers gather information using machine learning while an inventory of different martech players, including Acquisio, are developing and offering tools using similar advancements for better business outcomes.



## Case Study: Meet Lucy, IBM's Step-Child

When IBM's Watson showed up on Jeopardy, Scott Litman, the Managing Director of Equals 3, was excited to learn that it would be available to developers. Watson is a well-funded startup within IBM, whose series of APIs act like lego blocks for companies like E3 to solve some kind of challenge. Scott thought about all the marketing platforms he had seen over the years and wondered what he could do if he took all the data from predominant marketing platforms and brought it together. That is how "Lucy" was born.

**Lucy** is a "cognitive companion," built off Watson, for the marketing professional, says Scott. She, if you will, is set up to solve the content problem that marketers have so often in so many places and can also help with persona modeling and media planning. You can ask her questions like "What is the latest info on self-driving cars?"

Built off of an IBM intelligence engine, Lucy still has to learn about her companion before she can be of any help. Lucy is deployed quickly but it's up to the user to train their companion to be effective, which can take time. Therefore there must be a business agent of some kind to train Lucy, out of the box you won't get anywhere. Since Lucy is trained on a per-company or per-agency basis, the data given to her is unique to who she is serving and therefore each version deployed would be slightly different. Users of all levels can also interact with Lucy to give her performance feedback based on the accuracy of her responses using a star-rating system. This simplified grading system is another way the customized marketing engine can self-improve.

Lucy uses structured (ComScore, Kantar etc.) and unstructured data to make conclusions. Lucy reads a million new pieces of content per day to give brand insights and can break down silos to answer omni-channel business questions. Through one login marketers can access data from so many systems. At the moment, Lucy is unique in what she does. The ability to do research as well as audience persona and media modeling is new.



---

## 10 Things Marketers Need to Know About AI

Eighty percent of marketing executives believe AI will revolutionize marketing by 2020, according to a Demandbase and Wakefield Research survey of 500 business-to-business marketers published in December. But just 26 percent are highly confident they understand how AI is actually used in marketing, and only 10 percent say that they are using AI in their marketing programs.

— **James A. Martin, IBM**

February 13, 2017

## Meet the Machine Learning & Marketing Innovators

Companies like Microsoft, Salesforce, Hubspot, Equals 3 and more are pivoting to focus on their AI products. Below is an alphabetical shortlist of new and leading martech providers using AI:



### Adobe Sensei

Using AI machine learning technology and advanced algorithms, Adobe Sensei powers intelligent features across the Adobe universe to improve design and workflow in popular programs like Photoshop. Working with Adobe Cloud, Sensei can leverage content and data across Adobe products and can be embedded into any kind of application. As it learns from millions of photo and video editing sessions, Sensei will be able to help users make the most of the platform.

### Albert AI

Albert is an artificially intelligent platform for enterprise marketing departments or agencies intended to replace certain marketing efforts with greater efficiency and speed than ever before. The platform creates actionable data-driven cross-channel insights to help marketers make smarter decisions faster and reduce the cost of operations.



### Appier, CrossX

Focusing on enterprise business, Appier is a cross-device optimization system for eCommerce, gaming and brand building that allows them to deploy "cross-screen campaigns." By analyzing device ownership and audience attributes like browsing habits, interests and behaviors, the system simplifies communications across devices for better results with predictive capability.

### Automat

With over 17 AI patents, ranging from speech recognition to virtual assistants, Automat is a recognized brand solution for personalized conversational marketing at scale. The technology allows marketers to have personalized customer service conversations without intrusively monitoring their online behavior, addressing concerns about consumer data and privacy.



### Aviso

Using advanced machine learning algorithms and portfolio management frameworks, Aviso's Total Revenue Intelligence cloud applications help businesses make the best decisions surrounding revenue assets, risk mitigation, and more.



### Boomtrain, Boomtrain Messenger

Engaging with customers in real-time based on behavioral patterns, Boomtrain Messenger is an online chat tool that offers an alternative to sending yet another email to leads and customers, while making multi-channel recommendations.



### BrightFunnel

BrightFunnel's predictive marketing intelligence platform generates actionable revenue analytics, forecasting, attribution, and benchmarking for B2B marketers. Companies get a leg up on future revenue decisions and align their marketing and overall business goals accordingly.



### Collective[i]

The combination of advanced analytics, decision support tools, and optimization helps teams improve their performance in sales, marketing, customer service and customer support. Packaged in a user friendly platform, Collective[i] lets any business user easily make data-driven decisions using machine learning.



### Conversica

The cloud-based conversational AI Sales Assistant makes life easier on marketers and salespeople by simplifying the mundane parts of the sales funnel. The automated sales assistant qualifies, engages, and follows up with leads in a natural way, then hands them off to a salesperson - meaning salespeople have more time to focus on closing the deal.

## enigma **Enigma AI**

By using machine learning Enigma makes data accessible and actionable for smart operational data management in the financial, pharmaceutical, and insurance industries. They support the open data community by making their massive data library available to the public.



## **Google, Cloud ML Platform**

Google Cloud includes an artificially intelligent unit to process photos, voice search, translations and Gmail Smart Reply. The unit makes use of image recognition and other advanced techniques and is being led by Professor Fei-Fei Li from Stanford University. Customers can now make use of Google's Cloud Machine Learning.



## **InsideSales.com**

As an artificially intelligent growth platform for business, InsideSales.com uses machine learning to predict and recommend sales activity from the moment a lead enters the funnel until it closes.

## **LAYER 6**

A deep learning platform for customer personalization and marketing predictions suited for enterprise. The platform leverages data to make real time customer recommendations in finance, e-commerce and media. Part of the Vector Institute for Artificial Intelligence, Layer 6 has an open call for PhD data scientists. They pride themselves as one of the first companies offering deep learning to address enterprise business challenges.



## **MailChimp, Omnivore**

Omnivore is an initiative for email abuse prevention. When the system detects suspicious user activity, it can issue a warning and/or suspend accounts. Using machine learning, the system analyzes billions of emails to determine bad behavior campaigns and can identify them before they're even launched.



## **MindMeld by Cisco**

An advanced conversation platform, MindMeld AI intends to deliver cross-device and cross-channel conversational experiences. Businesses use MindMeld's conversation assistant to help their customers find information they're looking for.



## **Pinterest AI**

In order to show content results that accurately reflect a person's search, Pinterest uses machine learning that determines the best image to show out of 150 million image possibilities. The system uses image recognition to match queries to past user pins.



## **Salesforce, Einstein AI**

With a team of 175 data scientists using deep learning, natural language processing, and computer vision, the Einstein AI platform is available for sales, marketing, analytics and more. Each application has its own set of features. For example, the Einstein AI Sales Cloud helps teams decipher which leads are most likely to convert, while also alerting the team when that lead is considering competitor products or services.



## **Sas**

The largest independent vendor in the business intelligence market, Sas uses advanced analytics to turn data into information that makes sense and lets businesses make fact-based decisions. Being able to anticipate opportunities empowers businesses to take the right action and make a lasting impact.

---

## **Artificial Intelligence Marketing: Where Does Artificial Intelligence Fit in Your Marketing Stack?**

Addressing the myriad needs of a large marketing organization with a single suite of tools can be very frustrating, especially when working with legacy hardware and software. Before building your stack, it's key that you take stock of which tasks in each marketing structure can be handled by artificial intelligence and which ones must be delegated to other tech platforms.

— Albert.AI Whitepaper

## Acquisio is Machine Learning

Acquisio has coded predictive algorithms that learn how to optimize marketing results across channels and ad platforms.

The first intelligent algorithms we created were responsible for making very high-frequency bid and budget adjustments, which is why these became known as BBM – bid and budget management. These algorithms are still in use today by many of our clients. BBM earned a name for itself as the fastest SEM optimization technology on the market and is an example of algorithms not being explicitly programmed, but having the ability to quickly learn and adjust to new situations.


### BBM Becomes Acquisio Turing™

Over the past five years Acquisio's ML algorithms evolved past bid and budget optimization, forming a new foundation for innovation within marketing tech that our team is harvesting. BBM became part of our growing collection of machine learning algorithms. Today our machine learning layer powers every Acquisio product including our legacy platform, Dynamo, Marketplace and our platform for local businesses called Promote.

With new algorithms being developed for different products at Acquisio, our machine learning layer needed a name. We were looking for a name that would not only reflect our development of intelligent algorithms that are the first of their kind in our industry, but one that we believed in - a name that embodied our rigorous testing and development process as well as our quest for marketing excellence. Today Acquisio's machine learning layer is appropriately and fondly called Acquisio Turing™.



Acquisio Turing was built in Montreal amongst a **globally-recognized hub** for artificial intelligence research, innovation and development. Using both supervised and unsupervised learning, the layer self-improves by analyzing enormous sets of live campaign data, surpassing its own level of intelligence daily. Our team houses every new machine learning algorithm we develop under the Acquisio Turing umbrella, which continuously injects intelligent marketing insights into our products and platforms.



Our paid media platform is the first application of advanced machine learning technology to optimize and scale local marketing. With many other machine learning players entering the marketing space, Acquisio provides an essential piece of the AI-marketing puzzle with an intelligent SEM solution for any size budget. Acquisio Turing joins the self-improving cluster of machine-intelligence applied to marketing, aiming to deliver significantly better results than even the smartest humans.

The outcomes of our machine learning technology in the marketplace have been significant. Campaign managers see better campaign performance, even for low budget or volume spends. Our technology automates workflows, offering dramatic improvements to operational efficiency. We've also proven that it can decrease OPEX, increase LTV, lower customer churn rates and improve customer satisfaction. Most importantly our machine learning technology gets our customer's accounts more clicks and conversions for less money.

## Results from our Latest Performance Report

We recently released our 2017 **Acquisio Turing Performance Report** that analyzed over 30 000 SEM accounts either using machine learning or not. Across all business types and budget sizes globally, the accounts using Acquisio's machine learning technology are performing better than those that aren't, giving advertisers an advantage.

While we saw improvements across the board for KPIs like CPC, CPA etc., one metric matters most in any campaign - conversions. Advertisers can get a million clicks on one campaign in one day, at the lowest cost, but if they don't have even a single conversion to show for it, what does it matter? We observed that of the accounts who were optimized for conversions, **those using our machine learning experienced an average conversion increase of 71% over those who weren't!** This is huge! With these latest findings we're urging our clients to take advantage of our machine learning technology built into our paid search platform and products.

Other findings of the report focused on budget attainment. We found that clients using our machine learning were a lot better at budgeting. To be precise, accounts using our machine learning on AdWords were on average 3 times more likely to pace and spend their budget as intended than accounts that were not using machine learning. On Bing these accounts were 11 times more likely to attain their budget. While there are **many arguments supporting budget attainment**, PPC marketers too often run out of money with a few days left in the month. This is the worst time to run out of money since CPCs drop near the end of the month when there is less competition.

Aside from the usual PPC KPIs and budget attainment, our Report measured the impact that machine learning could have on the lifetime of an account. When an account stays longer on the platform, it suggests that clients are happy with the results they're getting from it. As we state in the report:

**The lifetime of each account contributes to any business's bottom line. The longer an account stays the more revenue they provide. For large SEM resellers with thousands of SMB accounts, an extra month of business across all those customers represents significant annual revenue.**

The accounts using machine learning technology lived three months longer on AdWords and two months longer on Bing than those that weren't, resulting in significantly more revenue for the advertiser. Budgets with less than \$500 in ad spend per month performed the best.

## Acquisio Uses Machine Learning to Solve Small Business Marketing

The problem with applying machine learning to local SMB advertising is the lack of resources. In fact, most think that small businesses don't have enough data or money to make use of machine learning. Getting a few clicks a day or less is very challenging to optimize - data sets are difficult to make viable and the technology is comprehensive and costly. That's why Acquisio developed Promote, a new machine learning product to support small businesses so they can be competitive - SMBs can be significant players in their local market with the help of AI.

We're solving a problem that was never previously able to be solved because no one can effectively optimize campaigns on small budgets. With Promote, we're helping SMBs thrive in the digital economy.

---

## Acquisio Turing Performance Report

We're not waiting for the future of martech, we're building it.

— Marc Poirier, CEO, Acquisio

September 2017



au fond de l'âme à la bouche





Fig. 24 (1/4) Fig. 25 (1/4)




Fig. 36

Fig. 28 (1/4)






Fig. 35




Fig. 37




Fig. 52 (1/4)





Fig. 51 (1/4)

Fig. 55





Fig. 27 (1/4)




Fig. 29 (1/4)





Fig. 38

Axe de l'âme



---

## How Artificial Intelligence Will Revolutionize Marketers' Lives

Don't send yourself to retirement too quickly. Thankfully these abilities mean nothing without a person's creative touch, extensive experience and strategic thinking. In fact, the more machines are used, the more humans are needed to interpret their output. (...) Larger, more emotional purchases require a human marketer to push the sale over the line. The AI will place the product in front of the consumer via a PDA [personal digital assistants], but the campaign's creativity will be needed to connect emotionally with the customer to close the deal. The perfect blend of art and science.

— Steven Ledgerwood, Marketing Tech News

January 27, 2017



## Learning to Trust Machines

As AI is becoming more and more pervasive, human-machine interaction and cooperation become commonplace. For such a relationship to be fruitful there must be some level of trust, otherwise the human actors keep second guessing everything the machine does. How does an AI system gain and maintain trust? How do we humans build up that trust and how easily can it be lost?

Both as individuals and as a society we have to work out our trust mechanisms. We know how much we trust people around us, we know when the trust deepens or when it is being lost. When we board a bus we implicitly trust that the company operating it keeps them well maintained and that the driver they hired is a safe driver. When we hear about accidents involving transit vehicles we might be shaken but typically don't lose trust in the entire transit system. But what if the vehicle involved in the accident was a self-driving bus? What if the driver was an AI system, not certified by some driving license agency? When it comes to humans making an error we are more likely to forgive them because they are humans and can learn from their mistakes. When it comes to machines we expect them to be perfect and not to make any errors.

Certainly there are times when a machine makes a mistake, but so do humans. While many human mistakes can be attributed to fatigue, distractedness, or lack of attention to details, machines don't err for those reasons. The most likely culprit in a machine's case is not having seen a particular situation before and not having learned how to handle those cases correctly. But let's not forget that AI is built on machine learning, where "learning" is just as important as it being a machine. On the one hand, being (and acting) like a human includes admitting not knowing things or even making mistakes, while on the other hand our expectation towards intelligent machines is that they be all knowing and always perfect.

Researchers are still hard at work making AI systems better and more accurate for the task they are trying to solve. These activities focus on improving the system's predictions; making the self-driving car require fewer and fewer driver interruptions, making the voice assistant understand more accents and more complex sentences, etc. These activities go a long way toward establishing trust, as near flawless execution will win over skeptics.

But at the same time these AI agents must also pass more general tests, such as that of explainability, fairness, ethics, and security to fully gain human trust.

### Explainability

When a person arrives at a decision they typically follow some logic that they can verbalize for the most part. Sure, there is room for instinct and spur-of-the-moment improvisation, but by and large their actions can be explained to someone else. When it comes to AI, certain models work like a black box. Even their creators could not formulate an answer to a question about the "logic" used to arrive at a prediction. The answer is hidden deep inside the neural network among its many connections. AlphaGo is proving **again** that its abilities now surpass that of the best human player, but if somebody asked it to teach them how to play Go the machine would be entirely helpless. The best answer to even explain why it picked a certain unexpected move would be that in the large number of prior games it observed, that step led to the most profitable outcome most of the time. A more important situation could arise when an applicant is denied a loan and would like to find out what caused that denial. If the system cannot provide an answer, people may lose faith in its decisions as if they were just arbitrary.

### Fairness


Any machine learning algorithm cannot be better than the data used to train it. As they say, garbage in, garbage out. The poor quality of the data is often due to errors in measurement, not recording vital information, or just incorrectly processing the raw information. A more subtle error stems from the data that includes human bias. If a job recommendation software observes that the majority of successful job applicants for C-level jobs are male, it is more likely to recommend those type of jobs to male applicants because that might be the factor most correlated with the outcome. If the training dataset includes racial bias for loan applications, the system is likely to learn that as a good indicator of loan worthiness because it just learns what it can observe in the data. It is the task of the engineers creating the system to ensure that the data is free of such biases and is a representative sample.

## Ethics

To date, ML and AI have been mainly deployed to perform lower level, mundane tasks to help improve productivity. But soon they could be in a position to literally decide life or death. A self-driving car will be in charge of not only getting its occupants safely to their destination, but keeping everyone else around them safe. It is only a matter of time before a self-driving car finds itself in an impossible situation. For example, if an accident is unavoidable and the self-driving car's only choice is between steering towards pedestrian A on the left or pedestrian B on the right, how would the AI system under the hood decide which action to take? Based upon size? Age? Social status? And when the accident investigators try to determine what influenced the outcome, what would the logic be?

## Security

If the recent waves of computer virus attacks have taught us anything, it's that when it comes to computers 100% security does not exist. In our internet-connected-everything-world it isn't very difficult to take control of connected cars or even smart toasters. As AI systems run on some sort of computer infrastructure, their creators are acutely aware of the need for amped up computer security. But is the AI algorithm itself hacker proof at least? Not necessarily. As **reported by OpenAI**, one can easily fool an image recognition system by altering an image slightly, making it the same to the human observer but for the algorithm it forces a very different determination.



Source

Here the original panda image is altered by a custom-made noise pattern that fools the algorithm into detecting a gibbon. This same technique could be used to fool self-driving cars trying to read traffic signs. One could design an error pattern on clear film that if placed over a Do-Not-Enter sign would be read as

something else by the self-driving car. A human observer wouldn't notice. Today's voice assistants are also lacking authentication. Anyone could potentially place an order via someone else's Alexa assistant (unless the ordering feature has been disabled) or obtain a person's calendar events or other personal information from Google Home or Siri.

## Making AI More Human

With discussion of trust, error, explainability, bias, ethics and security on the table, the idea that machines should be more human is fundamental to the concept of artificial intelligence. As machines get smarter and better at what we do, many are concerned about their jobs, but more likely they will work alongside us, augmenting human activity by taking over some of the less creative tasks initially.

In MIT's game-playing **experiment**, humans and machines weren't able to work together until they exchanged simple communication. At first, cooperation in the experiment's game was non-existent but after a few sentences were shared, the human and the machine began working together. As part of the human experience, we need to get to know something before we trust it. It is possible that if we just view these AI systems as pure machines, the interaction and trust level will reflect that.

From Hollywood's portrayal to our own set of standards in our workflow, we want all forms of AI and machine learning to impress us and never make a mistake. We want them to be more like us and even be named like us, but to be 100% objective. We're not at the point where we can let go of the steering wheel and just let the machine do all the work, but we can't forget that we are at the tip of the AI/ML iceberg. While as a society we are still navigating through this dynamic discussion filled with constant innovation, we do know one thing: As AI and ML automate decisions in increasing parts of our human lives through technological innovation, when it comes to marketing, the need for human intervention couldn't be more vital. Machines are not currently perfect nor will they ever be, but they offer a new possibility for better results and indeed the fourth industrial revolution - the digital economy, the age of information, the era of artificial intelligence.

---

## About Acquisio

At Acquisio, we come to work everyday inspired to help businesses thrive in today's ever changing digital economy. Our team uses Acquisio Turing, an award-winning suite of machine learning technology to optimize even small advertising budgets.

The Acquisio platform is known for facilitating customer acquisition using ad platforms such as Google AdWords, Facebook Ads, and Bing Ads. As one of the first SaaS companies to apply advanced machine learning technology to advertising, Acquisio's high-frequency optimization algorithms perform nearly 2.5 million campaign adjustments per day for 400+ clients running over 300,000 campaigns.

Acquisio is recognized as being one of the fastest-growing companies in North America, winning the Deloitte Tech Fast 500 and Fast 50 awards for four consecutive years. Acquisio has its headquarters in Montreal, and has offices in New York, Seattle, Austin, London, and Istanbul.

For more information, visit [www.acquisio.com](http://www.acquisio.com)

[START YOUR FREE TRIAL](#)



---

## **Artificial Intelligence (AI) and the Future of Marketing: 6 Observations From Inbound 2016**

The “marketing conversation” will become a human-machine conversation.

— **Gil Press, Forbes**

November 2016