

Machine Learning Overview: Understanding The 'Gold Rush'

Originally published at <https://www.iflexion.com/blog/machine-learning-overview>

Published: April 23, 2020 | By Martin Anderson



Over the last twenty years, the advent of global network connectivity has combined with exponential increases in computing capabilities to transform data science from the dry study of actuarial tables and census records into a hyper-detailed locus of rapidly updating and high-volume information about the human race.


What we think, what we buy, what we believe, who we like, love and hate...these, and hundreds of thousands of previously unknown decisions, preferences and minor events in the continuum of a human life are now directly readable via our relationships with networked devices, or else inferable from the same data, to those corporations, governments or researchers that consider such insights useful, valuable, or even essential.

From weblogs through to fitness tracking data, historical data digitization via OCR, and the JSON feeds of IoT devices, the information piled up far more quickly than even the advancing technologies of data science could cope with — at least, initially. This machine learning overview will look at how data-centricity has come to power our society at large, and where machine learning, among other data processing technologies, is positioned in this new reality.

Machine Learning Returns

Due to the staggering new volumes of data available, and better technology to process it, the early 2000s saw a resurgence of academic and commercial interest in machine learning and [AI software development](#) — a sector that had retreated back to academia after the second AI winter of 1987-1995.

The high research budgets of major players such as Google led to the advent of open source machine learning architectures such as [Hadoop](#), rapidly capturing the headlines and the imagination of businesses and governments alike, and quickly evolving into an ongoing symbiosis between academic research and the needs of the civic and private sectors.



AI Hype v. 3.0

This revived interest in machine learning and artificial intelligence was, perhaps inevitably, accompanied by the return of the hyperbolic 'AI hype cycle'. Like its predecessors, this new strain of 'AI fever' was fed by three main contenders:

- Academic projects that promised much, but were in fact highly speculative and unlikely to yield commercial results for many years (if ever).
- Existing vendors keen to associate their established technologies with the promise of machine learning, even where there was only a cursory association, or none.
- The more pragmatic exploiters of genuine, deployable machine learning solutions. Though they tended to overstate their case, at least they had a working product to put on the table, based on the current machine learning state-of-the-art, rather than technologies that might

only bear fruit later on in the AI/ML development timeline.

Momentarily we'll take a closer look at these three perennial proponents of the AI hype cycle, and then examine ways to identify real-world applicable machine learning technologies through the high signal-to-noise ratio of [AI hype](#).

First, let's see how the machine learning gold rush played out in previous eras.

A Brief History of Machine Learning and AI Hype

In each case, the state of the art was advanced only for the massive forward impetus to ultimately be hamstrung by a combination of technological bottlenecks (such as processor speed limitations or scarcity of datasets) and disillusionment around the potential of current approaches. In each case, many over-invested corporations and companies suffered when the bubble burst; and in each case, the most obvious persistent cultural debris of that gold rush was a stream of speculative sci-fi movies elaborating the potential dangers of high-level artificial intelligence.

AI Gold Rush #1: 1956-1974

The Dartmouth Artificial Intelligence conference of 1956 was to define the 18 years of AI's first 'golden age', as well as first to coin the phrase 'artificial intelligence'. Seminal ideas emerged which were to be fully exploited during the 1960s, including 'natural language' interfaces with computers and the full possibility of neural networks capable of processing information in similar ways to humans.

In this period, the prospect of job automation and computer-driven authority fed both the Hollywood studios and the headlines and TV reports of the day:



The reality proved less exhilarating. By 1974, it had become clear that even endowing an artificial intelligence with a basic understanding of context, or an ability to learn from its own mistakes, was proving an uphill struggle.

Decline triggers: Lack of adequately performant machines and systems; Marvin Minsky's criticism on the applicability of Frank Rosenblatt's perceptrons; the UK government's 1974 Lighthill Report, which was also highly critical of the real potential of machine learning after a decade of excitement over the topic; DARPA's shift to applied research in the early 1970s, after many years of blank checks to open-ended AI research projects, including a failed speech-recognition system for pilots.

AI Gold Rush #2: 1980-1987

Along with the advent of more affordable business and personal computers, the second machine learning gold rush revived the topic of connectionism, which had been an early casualty of the first AI winter. However, it drew most of its impetus from the promise of rule-based 'expert systems'.

Once again, the hype promised 'computers that think like us'.

The second wave did not last half as long as the first, and the decline would ultimately affect a broad range of mid- and high-level industry players. By 1988, recurrent quarterly losses across the applied and research sector in AI were making for devastating business reading:

THE SECOND AI WINTER

FRIDAY, MARCH 4, 1988

Copyright © 1988 The New York Times

Business Day

The New York Times

Setbacks for Artificial Intelligence

Companies Are Hurt By Poor Decisions

By ANDREW POLLACK

Some who have been "over-represented" in the artificial intelligence industry, claiming that they would prosper by putting the technology to use in "commercial" fields.

Some of the setbacks stem from the culture of artificial intelligence to quickly live up to its promise of making machines that can understand English, recognize objects or reason like a human expert — to be used for such purposes as diagnosing machinery breakdowns or deciding whether to authorize a loan. Despite this, the technology is making slow but steady progress, and now is being widely incorporated into more conventional computer programs.

Weak Business Moves

For more of the artificial intelligence industry's problems result from poor business decisions by companies that were heavily weighted with technology rather than business sense. Their main mistake was trying to use special computers for artificial intelligence; the machines were too expensive and did not mesh well with those used by general computers.

"People believed their own hype," said David Knight, chief of one leading artificial intelligence company, Expertwise Inc., and now president of the G2 Corporation, a strategy software company in San Francisco. "Everyone after appears to be intelligent in their everyday applications, they generally perfect respective tasks following rigid rules set down by programs. They do not think or make creative decisions, as humans do."

Artificial intelligence aims to make computers do tasks that are said to require intelligence when people do them. One goal has been to make computers understand English and other languages, and to give them the ability to make a program governing involves conducting lengthy inter-

As the industry realigns, the companies that relied on special purpose machines are languishing.

Company/Headquarters	Description
EXPERT SYSTEM DEVELOPMENT TOOLS	
Technology Inc., Calif.	Four quarters of losses; 60 workers of 200 laid off; 1987 sales: \$20 million.
Intelligence Resources Corp., Calif.	Six quarters of losses; 30 workers of 200 laid off; 1987 sales: \$20 million.
Carnegie Group*, Pittsburgh	Losses; 20-40 workers of 200 laid off; 1987 sales: \$12 million.
Inference*, Los Angeles	Losses; 20 workers of 130 laid off; 1987 sales: \$12 million.
MACHINE MANUFACTURERS	
Symyx*, Cupertino, Calif.	Continuing losses; Third round of layoffs last fall. Chief chairman and founder, 1987 sales: \$104 million.
Lisp Machine, Andover, Mass.	Fired for bankruptcy last year; 1986 sales: \$12 million.
Xerox, Stamford, Conn.	Sluggish sales; recent reorganization of its artificial intelligence business.
Texas Instruments	Big push in artificial intelligence. Announced a chip containing Lisp to go into Macintosh.
EXPERT SYSTEM APPLICATIONS	
Symlogics*, Sunnyvale, Calif.	No layoffs; 1987 sales: \$9 million.
Applied Expert Systems*, Cambridge, Mass.	Layoffs; 1987 sales: \$4 million.
Parasoft*, Cambridge, Mass.	Layoffs; Chief chairman and founder, 1987 sales: \$6 million.

*Private company, estimates.

Source: AI Trends

Companies

Macy is expected to contest the legality of Campeau's latest tender offer for Federated Department Stores. The Macy strategy apparently will be to contend that the offer represents a new bid, not an amendment, and that it violates Federal laws and a California state takeover statute. (Page D1.)

A \$1.1 billion bid for Koppers was made by Beazer, a British construction company. Koppers is a Pittsburgh-based steel mill and manufacturer of construction materials. The offer was for \$40 a share. Koppers' shares rose to \$1.75 and \$4.50. (D1.)

Chrysler will sell Avenger Inc., its parts-making unit, but will not clear four plants, affecting about 1,000 workers. A sale had been strongly opposed by Chrysler's unions. (D1.)

A new cable entertainment service is planned by Time Warner, who has gained tentative approval of the plan from a number of cable operators. The proposal would be to provide original productions for cable. (D1.)

ES&K ordered 16 Boeing 737 jets and took options on 30 more in a deal that could cost \$2.4 billion. (D4.)

It is expected to announce its entry into wholesale cable through a joint venture with MTV, a Dutch company. (D4.)

Bain presented its tender offer for Farmers Group to \$5 a share, or \$4.5 billion. (D4.)

Leading airlines raised fares for the second time in two weeks, reflecting unexpected strength in winter traffic. (D4.)

The Economy

Readers reported unexpectedly slow sales for February, with a number of large chains posting declines for stores open more than a year. Analysts, while noting that the consumer was still a relatively strong February last year, still expressed concern. (D1.)

Vehicle sales rose 16 percent in late February. For the month, sales of domestic vehicles were 15.9 percent higher, while sales of imported vehicles fell 4 percent. Detroit's heavy discounting received most of the credit for the strong results. (D4.)

A Senate panel's compromise on a landmark banking bill is expected to help secure its passage in the full Senate this spring. But strong opposition in the House. (D1.)

The Fed's chairman advanced two changes in the tax system. Alan Greenspan told Congress that he favors the imposition of consumption taxes and a reduction in the capital gains tax. (D4.)

Readers' productivity improved last year by 4.8 percent, just half as much as in 1986, according to revised figures. (D5.)

Markets

The stock market had a mixed day, dominated by takeover activity. The Dow Jones Industrial Average fell 7.89 points, to 2,663.45, but other indexes had gains. (D1.)

Read prices were little changed in light activity as many traders waited for fresh data on the economy. (D1.)

The dollar held steady in trading described as the lightest in more than a year. Gold fell to \$450.20 an ounce, down \$2.50, in New York. (D1.)

Macy Plans Legal Step In Bidding

Status of Campeau's Last Federated Offer May Face Challenge

By MADONNE BARNHART

In an interesting struggle to win Federated Department Stores Inc., R.H. Macy & Company is expected to contest the legality of the Campeau Corporation's latest tender offer for the giant retailer.

Macy is expected to contend that Campeau's latest offer is not an amended proposal but a new one that violates Federal securities laws and state laws that are anti-takeover statutes. With that challenge, Macy would be seeking to delay the current deadline of the Campeau offer in order to improve the chances for its own offer.

According to Securities and Exchange Commission rules, reducing the price of a tender offer or changing the price in a tender offer constitutes an amended offer, allowing shares to be bought in 10 business days. But in the case of a new offer, 30 business days are required before shares can be purchased.

Eyes on Expiration Date

If the reported plan by Macy's latest decision, the March 15 expiration date of Campeau's latest proposal could be delayed, giving Macy's tender offer an opportunity to run its course without being challenged.

Each offer has an estimated value of about \$1.1 billion, but even analysts, who are used to large takeovers, are not sure of the exact difference between the two offers.

Reflecting the confusion, Federated's stock rose modestly for a day, dropped 2 1/2 cents, to \$48.25, in heavy trading.

Campeau has said it would challenge features of the Macy bid.

Trouble for Many of the Artificial Intelligence Companies

As the industry realigns, the companies that relied on special purpose machines are languishing.

Company/Headquarters	Description
EXPERT SYSTEM DEVELOPMENT TOOLS	
Technology Inc., Calif.	Four quarters of losses; 60 workers of 200 laid off; 1987 sales: \$20 million.
Intelligence Resources Corp., Calif.	Six quarters of losses; 30 workers of 200 laid off; 1987 sales: \$20 million.
Carnegie Group*, Pittsburgh	Losses; 20-40 workers of 200 laid off; 1987 sales: \$12 million.
Inference*, Los Angeles	Losses; 20 workers of 130 laid off; 1987 sales: \$12 million.
MACHINE MANUFACTURERS	
Symyx*, Cupertino, Calif.	Continuing losses; Third round of layoffs last fall. Chief chairman and founder, 1987 sales: \$104 million.
Lisp Machine, Andover, Mass.	Fired for bankruptcy last year; 1986 sales: \$12 million.
Xerox, Stamford, Conn.	Sluggish sales; recent reorganization of its artificial intelligence business.
Texas Instruments	Big push in artificial intelligence. Announced a chip containing Lisp to go into Macintosh.
EXPERT SYSTEM APPLICATIONS	
Symlogics*, Sunnyvale, Calif.	No layoffs; 1987 sales: \$9 million.
Applied Expert Systems*, Cambridge, Mass.	Layoffs; 1987 sales: \$4 million.
Parasoft*, Cambridge, Mass.	Layoffs; Chief chairman and founder, 1987 sales: \$6 million.

*Private company, estimates.

Source: AI Trends

Decline triggers: As early as 1984, fears that funding would dry up as it had done at the beginning of the first AI winter; dependence on specialized hardware, which was expensive (and later unavailable, as returns on investment failed to manifest); investors' inability to wait out the long lead times for what was then perceived as the most enticing areas of machine learning, such as natural language systems and speech recognition frameworks.

Although academic research would continue to evolve during the second AI winter, it would take an extraordinary confluence of market forces, utterly unrelated to AI, to kick-start the machine learning market again, as we'll see.

Two waves of AI gold rush dissected in this major machine learning overview spanning 1956-2020s

The State of Machine Learning Now: The 'Opportunist' Age

So here we are again, wondering if the third time is the charm. Investment and interest in machine learning and AI have ascended from the troughs of the mid-1990s to become an unavoidable business topic, most especially over the last ten years of GPU-assisted machine learning.

GLOBAL REVENUE IN AI SYSTEMS (IN BILLIONS OF DOLLARS)

Year	Revenue (Billions of Dollars)
2017	100
2018	150
2019	200
2020	250
2021	350
2022	450
2023	650
2024	800
2025	950

Data source: IHS Markit, 2019

The defining traits of the two previous booms were that the theories and frameworks proposed would require either the most expensive systems of the age, or else dedicated, custom-built hardware — which was even more expensive, and would have required sustained market confidence in order to achieve economy-of-scale.

Instead, the current boom in machine learning has been fueled by a unique brand of opportunism, coincidence, and exploitation of technologies that were intended, at least initially, for other industry sectors.

Distributed and High-Volume Computing (including Parallel Computing)

Distributed and parallel computing frameworks such as [Hadoop and Spark](#) were developed initially to analyze the petabytes of data that were being generated from internet searches from the turn of the century onwards. This approach involved the marshalling of many low-cost computer units into high-volume storage clusters (distributed storage), and letting each unit calculate and contribute a portion of a workload back to a central organizing entity that would then collate the final sum of all the units' activity (parallel computing). Improvements in Graphic Processing Units (GPUs), driven by the gaming and leisure industries, would eventually offer machine learning an exponential boost in performance and capacity.

3

Open-source and Collaborative Development Models

The open source revolution freed machine learning and AI R&D from the confines of the patent silo. In the process, it transformed the economic model of machine learning systems research into a service-driven approach, favoring those who can best build on, commodify and exploit the common locus of knowledge while still contributing back to it.

Unprecedented, Actionable, High-Volume Dataset Availability

Above all, the current boom in machine learning is fueled by the spontaneous and relatively unrelated creation of huge digital data repositories that can be treated as datasets:

The digitization of information, from newspapers, archives, and historic texts to on-demand digitization of new documents — all cases where language can be indexed for the purposes of NLP.

- The social media phenomenon, wherein some of the biggest players, including Facebook and Twitter, are actively engaged in machine learning research, and provide APIs that allow machine learning projects to exploit public data; or else where screen-scraping and other furtive approaches are nonetheless able to derive data from more 'closed' social media systems.
- The explosion in online video, where publicly-available audio can be analyzed as waveforms for research in speech recognition, NLP and other pursuits around categorization and comparison, and where frames can be utilized for image recognition research.
- Device-driven information flows, including coordinate and temporal information from mobile and IoT devices, which can be exploited for research around [augmented and virtual reality development](#), as well as used in medical and statistical research.
- Network-based user analysis, including weblogs and Global Positioning System (GPS) data, for sentiment inference, recommender systems, fraud analysis systems, and targeted marketing.

An 'AI Autumn'?

On the general principle that history tends to repeat itself, and considering that memories of the first AI winter undermined the second boom during the 1980s, it's reasonable to wonder if a third AI winter is coming.

After all, we have all the same cultural signs as evidenced in the previous booms: media headlines that oscillate between paranoia and extreme skepticism, and a slew of AI-related movies, including *The Matrix* (1999), *A.I. Artificial Intelligence* (2001), *Her* (2013), *Eagle Eye* (2008), *Ex Machina* (2014), and the ongoing *Terminator* franchise entries, among many others.

However, it might be more pertinent to ask if the bursting of the AI bubble would actually mean the same thing this time, and what the persistent achievements are of this latest and longest 'golden age' of machine learning.

Some leading researchers propose the possibility of an 'AI autumn', wherein — for the third time — we accept that Artificial General Intelligence (AGI) is either conceptually flawed or otherwise unobtainable on any known current research pipeline; but that the more 'narrow' range of machine learning technologies from the last twenty years will continue to prosper on their own terms.


The first AI winter undermined the second boom during the 1980s, so it's only reasonable to wonder if a third AI winter is coming

Identifying Applicable Machine Learning Technologies Amid the Hype

Neither of the two previous machine learning booms delivered the sheer range of practical and applicable technologies that have conquered business markets over the last ten years. The machine learning state now is characterized with genuine commercial products and services rather than long-term ambitions that yield little fruit in the interim, as was the case in the two previous AI bubbles.

However, since the AI hype machine will likely continue to conflate old, current and putative technologies under attention-grabbing headlines, the business imperative is to stay on top of the terms, and seek out the signs of a genuine value proposition.

The Gartner Hype Cycle takes an annual look at the promises and deliveries for technologies related to machine learning, artificial intelligence and related terms.



While this is a useful crib to distinguish between legacy, high-value and speculative technologies, we should also learn to make these distinctions for ourselves as new products and frameworks emerge.

Therefore, let's now extend this machine learning overview to the three most common levels of capability in the AI hype cycle.

Three Players in The Machine Learning Hype Cycle

Discounting the claimed 40% of European AI startups that apparently use no AI at all according to the 2019 State of AI Divergence report by MMC, let's class the various capabilities of technologies that are currently marketed under the aggregated headings of 'machine learning' or 'artificial intelligence'.

The Drudge

At one end there's Robotic Process Automation (RPA), wherein machines observe human actions and repeat them in a rote manner, either on request or to a schedule. Companies such as UiPath, BluePrism, and WorkFusion currently provide stable and robust solutions in this sector. The fundamental concepts behind RPA have changed little since the industrial automation revolution of the 1970s.

Machine learning has only a casual foothold in RPA, mostly regarding facial recognition modules and cognitive automation bots capable of developing rudimentary 'principles' through active engagement with humans — albeit with a high level of oversight, supervision, and manual intervention.

The God

At the far opposite end, there's the popular conception of 'artificial intelligence', wherein a synthetic construct becomes capable of useful abstract thought and ethical decision-making, among other traits previously thought exclusive to the human domain. A 'singularity' of this nature, some contend, could be argued to have achieved our best definition of 'consciousness'.

The nearest real-world analogue to this ideal is the case of experimental self-driving vehicle (SDV) systems. Autonomous vehicles may hypothetically have to perform a kind of moral triage, for instance deciding instantly between the safety of its passengers and that of someone who has unexpectedly stepped out into the road in front of it.

However, if only for political reasons, such parameters will inevitably be pre-configured by humans, rendering this 'ethical capacity' the least 'self-determining' part of SDV architecture. Therefore, there is no 'AI God' in any current roadmap.

The Productive Researcher

Here, in the middle ground, lies practically every commercially viable application of machine learning, and most of the newer emerging technologies and approaches.

Applications, frameworks and architectures in this class toil and grind through petabyte-level reservoirs of big data in order to identify and classify, sift and refine, report and advise. Though they need meaningful human instructions, these solutions can nonetheless iterate through huge quantities of otherwise ungovernable raw information.


They don't understand why you wanted the information and they don't understand the significance of the results, but they'll hand them over at a speed and volume no human team could match.

The results they offer can be used for aggregate statistical analysis and reporting (such as dashboards for business intelligence), or to train algorithms, from Netflix-style recommender systems and HR candidate evaluation applications through to visual recognition algorithms in [computer vision software](#) and fraud identification systems).

The Drudge, the God and the Productive Researcher – what other types of machine learning capabilities can be singled out right now?

An Elementary Comparison

We can also think of these three tiers of analytical capability in terms of the famous creations of Arthur Conan Doyle:



Doctor Watson (left) is the tireless factotum and arch-researcher who lacks Sherlock Holmes' insight, but makes up for it with diligent trips to libraries and record offices, unearthing the core facts and figures that constitute the case.

Inspector Lestrade (center), is the plodding face of RPA-style automation, macros, conditional scripting and branching logic routines. He's been around a long time, and he gets the job done; but he needs a lot of help, and rarely innovates.

Sherlock Holmes (right) is the visionary genius susceptible to insights and new paradigms of thought that ordinary people can't aspire to. He's brilliant, but — like true AI/AGI — he's unavailable for the foreseeable future.

For our purposes, Doctor Watson represents the fruitful and achievable current state of machine learning. The genius he serves, for the moment, must remain our own.

Applied machine learning, RPA, and AI could well be thought of as Doctor Watson, Inspector Lestrade and Sherlock Holmes. Doctor Watson is the one representing the achievable current state of machine learning

Is Machine Learning Just a Macro On Steroids?

If one perceives Lestrade, Watson and Holmes as three points on a scale of efficacy and performance, is there any real difference between them, except their individual speed and performance thresholds? Can you overclock Lestrade until he reaches Watson's level of competence, or turbo-boost Watson until he attains the greater genius of his old friend?

If so, is the entire development of machine learning and AI substantially different from the evolution in computing throughput, clock speeds or capacity? After all, lack of adequate computing power played a notable role in foreshortening both of the previous AI booms.

These questions assume, as the AI hype machine does, that high level AGI/AI might eventually be capable of the same kind of serendipity that we ourselves experience in moments of unexplainable inspiration; that such insights can be considered to transcend the sum of the host (whether the host is a person or an AI architecture); that the continued, linear progression of computing capabilities will lead logically to such 'meta-human' insights; and that such insights would be in some way profitable to us.

By inference, they further assume that we ourselves may be more than the sum of our own 'programming'.

For the most part, these are unanswerable metaphysical questions. From Lestrade's 30-line CRON script to Holmes' zettabyte-level neural network, all three levels of automation represent machines executing our instructions — some more quickly, profitably or autonomously than others.

If such machines were to genuinely challenge our beliefs — as a 'self-aware' AI might — they'd likely meet the same resistance and skepticism as occurs when real people do this to us. In most cases, this would confound the usual business model for adopting technologies.

Therefore, our immediate task is to identify and enlist machine systems that can help us in our purposes, and leave these more abstract concerns to a later date.

Machine Intelligence = Business Intelligence

In order to attract general attention, the AI/ML hype cycle often conflates the three classes of capability into one overly generic super-sector, freely using (and abusing) the distinctive terms for each type of technology.

In this way, we are being marketed to buy into established 'legacy' or older technologies to which the new buzzwords are being deceptively applied; or else invited to invest in nascent machine learning pursuits that are currently more speculative and academic than applicable.

Since the true meaning of such popular terms remains clouded, it's useful to think of the phrase 'machine intelligence' not as a simile for computer consciousness, but rather in the sense of ['business intelligence'](#) — useful and actionable information that has been derived systematically through machine learning techniques from otherwise opaque and vast volumes of data.

Machine learning sectors with applications that yield 'machine intelligence' in this sense include the following.

Natural Language Processing

One of the oldest lodestones of machine learning research, Natural Language Processing (NLP) seeks to derive the correct context and meaning for the written and spoken word across a range of languages.

The open-source NLP libraries that dominate this field are helping researchers and business applications to create and feed a number of machine learning sectors, including:

- **Machine translation** systems, which have latency requirements similar to self-driving vehicles (see below), and require massive datasets and computation to advance the state of the art.
- **Speech recognition** systems, a corollary of machine translation systems, essential for voice-based AI assistants, [chatbot development](#) and voice-based authentication frameworks, among other uses.
- **Sentiment analysis**, which seeks to gauge user state and intent from natural language, and must identify and overcome enigmatic linguistic eccentricities such as sarcasm.
- **Spam filtering**, wherein common language patterns are used to identify and block unwanted commercial communications, even where spam actors are using elliptical language to bypass the filters.
- **Market intelligence systems** such as [QlikView and Tableau](#), capable of inferring and anticipating customer intent not only from customer behavior but also from user-generated content.
- **Machine summaries** of long and often complex documents that may lack metadata or structurally obvious natural language summaries.

Computer Vision

Teaching computers to understand the world around them in visual terms, and in some cases to interact with it in a creative, useful and safe manner, is one of the compelling economic use cases for ML/AI research. The market for [computer vision applications](#), hardware and services is estimated to reach a value of more than \$26 billion by 2025, across sectors that include:

- **Vision-based robotics systems:** pre-programmed 'positional' robotics workflows were the driver of the factory automation revolution of the 1970s, as well as new and revolutionary motion-control systems in the movie industry at that time. Machine learning, however, promises not only to make robots capable of acting based on understanding the changing environment around them (object segmentation, see below), but also of learning autonomously from their own mistakes and successes.
- **Facial recognition** is among the best-funded sector in machine learning due to its high applicability in the context of national and private sector security solutions. It drives secondary pursuits such as image manipulation (see below) in much the same way that the high-value videogame industry has inadvertently driven machine learning via advances in GPU technologies.
- **Self-driving vehicle (SDV)** systems, which must collate and interpret multiple sources of input including external video (see 'Object segmentation' below), data from LIDAR and analogous systems, temperature and geolocation data — and manage these inputs in the same fraction of a second as afforded to a human driver.
- **Object segmentation**, which transcends the single mission of facial and object recognition by attempting to identify and classify multiple components of an image across domains (i.e. people, cars, roads, etc.), and is a critical component of SDV and vision-based robotics systems.
- **Image manipulation**, including style transfer, reconstruction, upscaling, colorization and synthesis, applicable both to individual images and to video footage.

Predictive Maintenance

Predicting structural, organizational or technical failure is among the most prized goals in machine learning: in 2019 the US Army commissioned an AI-driven system to anticipate technical failure in its aircraft, vehicle and weapons systems; industrial Supervisory Control and Data Acquisition (SCADA) systems are beginning to adopt machine learning; and ML is also being used to decide the Remaining Useful Life (RUL) of technologies, allowing them to be decommissioned before critical failures occur.

Static Analysis and Security Scanning

Various companies now offer [security testing](#) for application source code, backed up with machine learning techniques. Anti-malware and anti-virus products have long used heuristics and behavioral detection techniques to attempt to identify rogue code that's behaving in unusual ways, but machine learning can catalogue and define such behavior at a far quicker and more efficient rate.

Machine Learning Overview: Conclusion

The above are only select examples of some of the real-world machine learning solutions that have emerged from the current AI/ML 'gold rush'. All are driven by the advent of big data, and supported by serendipitous developments in other industries and sectors — two crucial factors absent in the previous AI booms of the 1960s and 1980s.

While we may continue to overestimate the prospect of 'sentient' machines, we also perhaps underestimate the extent to which the *economically valuable* parts of our thought processes can be imitated, replaced or even improved with machine learning technologies. Most of our work is as task-based, iterative, goal-oriented and quantifiable as any in applied machine learning.

On the plus side, this may cause us to re-think the limited scope to which we choose (or are required) to apply our cognitive abilities.

If business should ultimately lose faith in the loftier promises of the current machine learning gold rush (as it has done twice over the course of the last fifty years), it will likely prove to have been the most fruitful and fortunate 'wild goose chase' in the history of technological advancement.

In order to survive the growing wave of machine automation, our own task is to identify pertinent and deployable machine learning technologies, and recognize that for the foreseeable future they will remain tools in a destiny that we continue to write for ourselves.

Our task now is to identify pertinent and deployable machine learning technologies, and recognize that they will remain tools in a destiny that we continue to write for ourselves.