We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

AI & Digital Platforms: The Technology [Part 2]

Örjan Larsson

Abstract

This essay aims to describe the dynamics at play in the field of industrial AI, where the significant efficiency potential is driving demand. There are rapid technological development and increasing use of AI technology within the industry. Meanwhile, practical applications rather than technical development itself are creating value. The primary purpose of the article is to spread knowledge to industry. It is also intended to form the basis of the Swedish innovation program PiiA's ongoing work around open calls and targeted strategic innovation projects. The basic approach taken is to investigate both industry demand for AI and how the supply of technology is developing. AI takes in a broad and dynamic range of concepts, but it should also be considered in the even broader context of industrial digitalisation. It is not just a question of technology development, but equally about application knowledge. Realising the full potential of AI requires the ability for change within individual companies, but also to handle exchanges and interactions in changing ecosystems. The article has been divided into two sections: *The Market*, in which we assess the development and the consequences on the factory floor; and *The Technology*, which provides a more in-depth understanding of the structures of industrial IT and machine-learning technology. The article concludes with four practical examples from the industry.

Keywords: AI, artificial, intelligence, PiiA, blue, institute, automation, algorithmization, platform, data, process, industry, IndTech, digitalization, digital, twin, ecosystem

1. Introduction

This article has two primary purposes: the first to provide the industry with an evaluation of the importance of AI development as a force for change. Second to create an internal basis for the Swedish Innovation program PiiA's future development efforts, within which AI can be described as the next phase of industry's digitalisation. Both these objectives are naturally compatible with the overall ambition of the report: to reach our target group of industry leaders and to serve as a source of knowledge for ongoing activities within relevant companies.

Technological, industrial development is awash with grand ambitions that have turned into mere passing fads and costly dead ends. With this in mind, throughout our work in assessing the development of AI, we have endeavoured to take into account the magnitude and direction of different vectors of change. On the one

hand, we have attempted to understand the power of demand for AI by assessing the economic impacts at a macro level. We have focused on productivity and qualitative values at various stages of industry value systems. On the other hand, we have attempted to assess the range of available technologies by analysing initiatives taken on a global scale and through focused academic research. We have also put considerable effort into understanding the major commercial – or applied – forces that are crucial to development, both in the short and medium-term.

We have also strived to place AI development in the context of current systemic developments, as characterised by the “platformisation” of company IT resources. By this we mean the transfer of automation and IT support to the cloud – a trend that is creating new competitive dynamics. Finally, we have attempted to translate this big picture into real impacts on the factory floor and to revisit well-known concepts such as organisational development which – with the help of the raw power of AI technology – have the potential to make the previously impossible, possible.

The project was a collaboration between PiiA and Blue Institute, with valuable input from Blue Institute’s network of CEOs and industry leaders on all levels. A big thank you is extended to everyone who contributed to this study.

2. From market to technology

2.1 One platform to rule them all...

“One Ring to rule them all” was the theme of Tolkien’s Lord of the Rings. The parallel with the power of algorithms – and thus the importance of platforms – is not so far remote from that idea. In the platform war, there is currently a battle for dominance of the market, with resource concentration an ongoing feature. At the same time, most companies have unique needs for which the general services that the big tech companies offer in their public clouds are not enough. Most companies, therefore, now have some form of a privately-owned platform (private cloud). One trend in the market is for hybridisation between one or more public clouds and a company’s private environment.

The issue of cloud complexity is becoming increasingly topical as more and more specialised domain-specific clouds/platforms are launched to the market (**Figure 1**). The concept of vertical clouds has taken hold and complements the original “general” clouds, which are now also called horizontal clouds. Vertical clouds represent industrial verticals and subprocesses within each vertical. They typically take the form of a PLM cloud, an MES cloud, an automation cloud, and so on. Here we can see, as described above, how automation vendors are now adding vertical domain-specific clouds or platforms to their offerings.

Another contemporary trend is locally distributed clouds known as edge computing or sometimes fog, which we touched upon earlier in the report. Edge computing is expected to have an increasing significance as AI technology increasingly requires a local capacity to complement the core resources of server halls.

One practical, short-term solution to get all these clouds working together through hybridisation solutions. The big hope for a long-term solution for industry in operational applications lies in standardisation which would make it possible for different environments to be combined in the same physical facilities. Standardisation work is going on within ISO, among other areas. The idea that we, like Windows, would become a de-facto standard is not likely, even though Microsoft is the provider that is currently the most successful in production-related applications.

Figure 1.
 Cloud structures are becoming increasingly complex. Public, vertical, local and distributed clouds are four varieties of the cloud with related services and infrastructures. These platforms are suitable for coordinating through hybridisation and, in the longer term, common standards.

2.2 Automation suppliers

There is currently a developmental trend that has the highly consolidated automation industry and leading companies such as ABB, Siemens, Emerson and Rockwell all moving in the same direction. Their common goal is to achieve market platforms that are specifically tailored to the industry. While these structures can be likened to operating systems for IoT, they meet the criteria for platforms as they match different types of users against each other. (See also the section on vertical clouds on the previous page).

AI is being used as a management and administration tool within these platforms and can also be used to produce products in the form of smart apps. These can be tailored to different applications and are available for purchase with one click via corporate app stores. Specialist centres are also being created for various product and industry applications and being linked to production facilities to allow for online optimisation and troubleshooting by experts.

- Schneider Electric's IoT platform is called EcoStruxure Platform and uses Microsoft Azure.
- In 2018, Emerson acquired GE's famous Predix Platform, which uses services from Microsoft Azure and Oracle, among others.
- Rockwell has Connected Enterprise and Factory Talk, which also use Microsoft Azure.

- Siemens' investment in this area is called MindSphere and rests on resources from Microsoft Azure, IBM Watson, SAP and even Amazon Web Services.
- Microsoft has taken a firm grip on the close-to-production IoT market through its partners. Microsoft Azure is its platform and includes services, tools, and infrastructures that can, among other things, simplify AI development.
- Service offerings include Microsoft Cognitive Services, a set of pre-built AI features, including vision, speech, language, and search functions. Everything is in the cloud and can be integrated into applications. Some features are customisable and can be optimised to transform and enhance organisational or industry-specific processes.
- ABB's venture is called Ability and is built on solutions from Microsoft Azure and IBM Watson.

The world market for IndTech – products and systems for industrial digitalisation and automation – is worth USD 340 billion and has a growth rate of 6–7%. The area can be divided into IT and OT (operational technology). The IT share is USD 100–110 billion, while operational technology for production and logistics accounts for USD 230–240 billion. Within OT, the distribution is 45 per cent for automation for the manufacturing industry and 55 per cent for process automation.

OT includes different types of industrial control systems and field equipment such as instrumentation, drive systems and robots. A particular growth area is industry's Internet of Things which complements traditional system environments. Several platform suppliers are now also launching dedicated and distributed systems for machine learning at a local level. Edge capacity on the factory floor can thus effectively be integrated into the cloud. For example, in 2018, Google released the third generation of the Tensor Processing Unit (TPU) chip.

This parallel development has led to several automation companies developing stand-alone AI modules with neural networks that can be plugged into the racks of the control systems. With this comes pre-custom type solutions for different processes or process objects.

2.3 Is AI expensive?

The computerisation of industry in the 1980s and 90s cost a large amount of money. Machines were replaced, and investment mainly focused on large process control systems and specialised computers, while thousands of kilometres of cable were laid. Air-conditioned computer halls were built, along with control rooms and cross-connection spaces. In short: it was expensive, but a good investment, nonetheless, as productivity and key quality figures skyrocketed. This type of primary investment will always be needed when rebuilding or when new investment is required, but the nature of digitalisation alters the equation. Consider the so-called “logic of small streams”, where a large number of spread-out, smaller contributions can come together to produce great results.

Today, you do not need to build air-conditioned computer halls and to buy servers to bring good ideas to fruition. It's easy to order as much computing power and functionality as you require – including AI tools – from the cloud at comparatively low prices. The majority of the data monitoring and collection infrastructures one might require are already in place. If sensors and hardware are needed, there are – or

will soon be – cost-effective IoT modules that will meet even the strictest precision and environmental requirements. Communication solutions are also expected to be wireless, reliable and inexpensive in the future.

These developments are leading to an evolving approach to change. It's becoming possible to be creative, to test the boundaries, and to think outside the box. Massive investments will not be required to unlock parts of the hidden value within plant and production processes and to outperform both competitors and customers' expectations.

The challenge with succeeding with AI is less about expensive investments in computer technology and more about obtaining the right skills in the right constellations. Experience shows that good AI projects are characterised by successful teams where domain and process knowledge, knowledge of analysis, and the right tools, all play a crucial role.

2.4 Concluding thoughts about the market as a framework for technology development

Our first aim was to examine the expected value-creating effects of AI and, based on the evidence we have studied, we have found that AI's impact on productivity development is likely to outpace by a factor of two previous, generic technology shifts, such as the introduction of steam, robotisation and IT. According to PwC, the global GDP in 2030 could be as much as 14 per cent higher due to the AI effect. Global impacts on industrial sectors leading up to 2030 could amount to as much as USD 2.3 trillion.

Our second ambition was to examine whether the development of AI and the associated supplier system could meet the demand for the technology that is being generated by the potential value creation it brings. Our view is that the centre of the development is now leaving the initial innovation phase and moving into the best-practice phase. This assessment is based partly on the fact that the large R&D investments being made need to yield profits, and partly on the fact that standardisation work is well on the way to show results. Besides, industry leaders around the world have come to realise the vast sums at stake in the coming transformation of the industry. Furthermore, there is also a dynamic that will arise when the three developmental hubs begin to work together, once development results reach the market. This is likely to produce an increase in torque for the entire system.

Thirdly, we wanted to understand whether this development was sustainable or whether we are currently seeing a hype effect that will wear off, with the actual market breakthrough coming much further down the track. Looking at Gartner's (often challenged) model Gartner's Hype Cycle for Emerging Technologies, there are different technical aspects of AI spread out across the different phases of the model. For example, deep learning using neural networks – which can be viewed as representative of the industrial use of AI – is currently at the "peak of inflated expectations" stage (Aug 2018) but could move to the "plateau of productivity" within 2–5 years.

Our assessment and our S-curve model suggest that digitalisation, in a broad sense, has now reached the beginning of the "plateau of productivity": best practice. It is challenging to make timing predictions around the introduction of AI to industry. On the one hand, there is a spread over different verticals with different conditions, and there are also various activities underway, ranging from management and administration to forecasting and foresight, to operational functions in production and logistics. In some areas, AI has already established technology, while in others, it is still in the developmental phase.

But we also found via our empirical evidence that machine learning is being quietly tested in many more places than you might imagine. Safety reasons dictate that the incubation period for new technology within heavier industries is much longer than in other commercial areas and longer still than the consumer area. With this in mind, 2–5 years seems not a long time, but rather a reasonable action period for translating ideas into operational benefits. A two-to five-year timeframe for the more permanent establishment of AI technology is also in line with our analyses, with the focus point shifting from a “best practice” situation to a commercial production breakthrough on the S-curve.

We conclude that the massive underlying forces driving both the demand for and supply of technology guarantee a stable development outlook for AI for industrial applications.

Within the period mentioned above, the timing seems right for a match between technology reaching the industry and the spreading of insight into the possibilities of this technology. This has the potential to create a significant industrial movement and thus deliver increased commercial demand at the company level and the achievement of previously unachievable results in production systems.

We also expect structural changes in the supplier system as entirely new concepts reach the market and previous industry and supplier limitations cease to be valid. There will be incentives for both extensive consolidation and repositioning in many areas.

Industry case study: Kone and ThyssenKrupp - Preventive Predictive Maintenance

Two hundred and sixteen centuries is a long time to wait for a lift. That figure is an estimate of the cumulative annual stopping time of the twelve million lifts in the world, moving about one billion people every day. To improve the maintenance of lifts, escalators and conveyor belts, the two suppliers Kone and ThyssenKrupp, have begun to use machine learning. According to the companies, it is just a taste of what is to come. The idea of the work is to anticipate errors before they happen, and the experience gained from these large-scale, global applications is expected to provide a valuable knowledge base for the broader development of AI in preventive maintenance – something expected to be required in most industries.

The challenge

Predictive maintenance is not a new concept. Industries that require high availability such as pulp and paper, chemistry, oil, gas and steelworks, to name a few, have long used statistical analysis tools to forecast interruptions and improve maintenance work. But machine learning provides a new level of accuracy and efficiency and makes predictable maintenance possible on a large scale on a large installed base. It is possible to identify common error patterns over hundreds of thousands of lifts, and at the same time, using algorithms, detect anomalies and specific behaviours for each lift plant. While two lift plants might be of the same model, their practical use will differ from day to day, as will the infrastructure around them.

It is simply not possible to apply simple sets of rules across such large and heterogeneous environments – which is why machine learning represents a real breakthrough in this context. Until now, predictive maintenance has involved identifying fault thresholds using a range of sensor data, which in turn can statistically indicate faults in lift plants. Machine learning involves using historical data in which fault events have been identified to allow the system to learn to find new faults – all without operators having to tell the AI what a fault pattern looks like.

The experience

In 2015, ThyssenKrupp launched a service called “Max” based on data from IoT sensors, control system data and data from the company’s ERP environment and CRM

systems from SAP and Oracle. In collaboration with Microsoft, a cloud-based data storage facility was created based on the Azure Cloud Platform. ThyssenKrupp currently provides the service to approximately 120,000 lifts and other systems, or 10 per cent of the installed base.

Open source code is used to build classification and regression models. A combination of models across different data streams and types of objects is compiled to achieve highly relevant and reliable results. The various predictive models are also gradually becoming outdated, due to lifts and escalators wearing out, being rebuilt and maintained, and so continuous re-learning is carried out.

The goal is to send field technicians to a facility before it fails. Although the maintenance system currently is not reaching that goal, the technician is often on the road when the call from the customer arrives. Once in place, the system has already done much of the troubleshooting work that would otherwise be started only once service staff are in place.

The introduction of AI technology at ThyssenKrupp has led to a review of some organisational boundaries between its service departments, IT and other functions. According to ThyssenKrupp's data, the system (used by 20,000 service technicians worldwide) has so far reduced the stoppage times of over 40,000 customers. It's not just the development of machine learning that has made this possible. Lower mobile data costs and the development of cloud technologies have also been enabling factors.

From an organisational development perspective, the project is not primarily technology-driven. From the outset, a broad group of different professions has been involved. Field technicians have been at the centre of the action, and they have been complemented with an IT team with skills in cloud and machine learning, as well as the skills to bring together and prepare the data. HR, legal, construction, production and other divisions have also become involved.

Kone – one of ThyssenKrupp's competitors – has developed its service offering along with IBM and Watson IoT systems. The partnership was launched in February 2017, and Kone has since equipped the facilities that use the service with IoT sensors to measure around 200 different parameters, such as movement, temperature, air pressure and forces within the machine.

Data is transferred to the IoT cloud platform as well as data and error status from the control systems. IBM Watson's natural language learning and machine learning processes have also been able to analyse information in maintenance logs and manuals accumulated over several decades. Watson can also be applied to images, sounds, and vibration patterns. Some generic components such as rotary machines have been modelled with general data from the installed base, and models have since been refined with more data specific to each device.

The result is that customers see significantly fewer stops and errors, and they experience a higher level of service. Plans include adding more data and expanded infrastructures that will further develop the customer offering. There is a plan to let people interact with the lifts so that the lifts, for example, sense when someone leaves a hotel room and then ensure that a lift is available on the right floor. Kone sees great quantifiable benefits from applications like these, which will drive AI to integrate into other applications as the technology improves.

Kone and ThyssenKrupp show that technical industries with large dispersed installation bases can use AI and machine learning technologies in the cloud to build predictive models that plan and make maintenance more efficient. But just as we have noted before, there are challenges when it comes to developing and applying the computer models, especially concerning data management. This means that these techniques are likely to be mainly used, where the return is most significant.

Sources: ThyssenKrupp, Kone, Computer Weekly.

2.5 The birth of modern AI

In the summer of 1956, a select group of researchers met at a seminar at Dartmouth College in the United States. The topic was Artificial Intelligence, and the optimistic goal of the summer meeting was to “achieve significant advances in the AI field”.

Convening the Dartmouth Summer Research Project on Artificial Intelligence was a young assistant mathematics professor who would later become legendary. His name was John McCarthy, in his invitation, he wrote: “We propose that a two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed based on the conjecture that every aspect of learning or any other feature of intelligence can, in principle, be so precisely described that a machine can be made to simulate it. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

The seminar marked the beginning of a favourable period for American AI research. Support, in the form of considerable funding from the defence authorities during the Cold War, led to great academic freedom and a creative research climate. But the 1970s brought challenges. AI researchers had underestimated the challenges, and a series of setbacks followed. Both American and British-oriented AI-foundational research was deprived of its funding. The research was criticised for lack of realism and lack of results.

The 1970s came to be called the first AI winter. Several setbacks would come later, as the hype around the field repeatedly rose and fell. The so-called LISP machines and expert systems of the 1980s were market failures that once again reduced development grants and led to new freeze periods. However, in the late 1990s and early twenty-first century, AI development quietly began to make progress. Expert systems using the technology could be used commercially for logistics and medical diagnosis. These successes came from better methods and more, and cheaper, computational power. The twenty-first century, and especially the period after 2010, has shown that AI is now an established commercial field that is growing rapidly.

Consumer applications from Amazon, Google, Microsoft and Apple are being rolled out on a broad scale, while AI support is now built into finance, media, trade and industrial applications. AI for language management is expected to grab the largest share of the market in the coming years, while health care applications are forecast to have the highest growth rate. Meanwhile, industrial applications are also expected to grow rapidly, a trend that is supported by our analysis in the first part of this report.

The current commercial breakthrough of AI technology is the result of the simultaneous coming to maturity of several underlying fields. The rapid expansion of the Internet from the 1990s onwards means that large amounts of data are readily available today. Data is the raw material of AI technology and is transformed into money and growth using algorithms.

In parallel, other developments mean that the cost of computer capacity is now rarely a limiting factor. Algorithm technology has undergone a similar process. All this comes on top of unprecedented growth within tech companies and their large appetites for investing thousands of billions of dollars in AI development, as well as general conditions that have allowed for the spread of AI. The overall picture is clear: AI in the form of machine learning is an established field of commercial technology that is achieving significant breakthroughs within all verticals in all markets.

The purpose of this concluding section of the report is to provide a somewhat deeper technical perspective to complement the report’s market focus in the first part.

We will start with an overall picture of IndTech and the scope of industrial IT, automation and digitalisation in general. This will be followed with a description of machine learning technology, and finally, a discussion of the data challenge, the concept of collaborative intelligence and the future of AI.

2.6 IndTech: an overview of the industry as an application area

The concept of IndTech brings together IT with both operational technology on the factory floor and digital development. It has a special significance in that it is where technologies from a range of different fields and periods come together (**Figure 2**). In addition to helping to transform the industry, the IndTech movement is creating a world market for industrial technology worth SEK 3.5 trillion per year [1]. IndTech is a hidden and yet giant industry and a field of excellence for Sweden, with numerous renowned companies operating in the area across the world.

The installed base of automation and industrial IT in the world is estimated to be SEK 50 trillion. This is where technology with roots in the 1980s meets with digital innovations generally not even developed for industry; something that’s hardly surprising given that a range of other sectors encountered digitalisation far earlier. The picture of the field that is emerging is thus one of the great opportunities but also significant challenges.

The traditional view of system support for the industry has been a pyramid-shaped hierarchy, with operational technology closest to production, and IT

Figure 2.
The model for IndTech: traditional and new technologies come together and make “smart industry” possible. Classic automation and industrial IT meet digitalisation and create new digital platforms and business ecosystems. Source: Blue Institute, 2019.

for administrative processes located above it. The idea that this hierarchy, the Automation Pyramid, might be dissolved in favour of more flexible structures has long been the subject of discussion. How this might happen has been less clear.

Incremental change scenarios seem the most likely given industry’s installed base of 1990s technology, much of which has a significant remaining life span, and the need for extensive standardisation work. In the short term, the focus may be on removing silos through better, more practical integration between computers and organisations, both within companies and in supply chains. In the longer term, the focus is likely to be on interoperability in the shape of the full interchangeability of information, without manual intervention, based on accepted industry standards.

To understand the general impact of digitalisation on the industry, it’s essential to consider which existing structures could simply be replaced by new technologies (a less common scenario), and which are likely to go through incremental changes over a long period (the more common scenario). The challenge going forward will be to use digital platforms and information transparency to address market fluctuations with new organisational approaches and ways of doing business (**Figure 3**).

Industry’s experience with previous technology shifts has demonstrated the importance of creating an overall conceptual picture, as well as having clear objectives from the outset and working towards them one step at a time. These objectives should include at a bare minimum: having digital infrastructure delivered through one, or several, specialised cloud services from different providers; using AI analysis for automation, augmentation and a collaborative approach between people and machines; using of the Internet of Things as a general application platform to lowers prices and simplify hardware and software.

Together, these three verticals form a digital platform with the potential to resolve information hierarchies over time. One of these verticals pertains to advanced analytics, an area in which machine learning if applied correctly, can be a potent tool. We will now examine this field in more detail.

Figure 3. Development can be summarised as integration in vertical and horizontal directions, and through new technology fields that both complement, improve and challenge the traditional environments and hierarchies. Source: Blue Institute, 2019.

2.7 AI analytics with machine learning

Artificial intelligence is often seen as something almost supernatural, and the media is often prone to highlighting its more sensational aspects. But as we will see, machine learning might just as well be called data analysis or applied mathematical statistics. The principles are very logical, even if the calculation processes are wide-ranging and complex.

Advances in AI development typically based around machine learning being applied to larger and larger sets of data and the development and efficiency of learning algorithms. Machine learning is, therefore, the technology behind most types of AI we see today. While traditional computer programs adhere to pre-determined explicit program instructions, machine learning algorithms scan data to detect patterns and then learn to make predictions. The algorithms adapt gradually, and the experience they gain is utilised and improves efficiency over time.

The mechanism behind machine learning centres on how tasks are presented as an input to a matrix-like structure; a neural computer network inspired by the functioning of the brain.

A machine learning algorithm expresses a function between the data it is fed, and the data produced by the model: $y = f(x)$. This function is always unknown, as it cannot be precisely determined mathematically, and this is where the finesse lies in machine learning: estimating the target function as accurately as possible. Correspondingly, if it is possible to determine the function in some other way, machine learning is not needed.

The output of the network, the prediction, depends on how the junction points in the network where the data meets during the process are given different values, called weights. These weights are the secret to the system's learning. (The junction points can be likened to the neurons of the brain (**Figure 4**)).

The problem lies in how to calculate the weights. The most common way is to start by giving them random values and seeing how significant the errors emerging from the model are. Each error is measured and then used to gradually change the weights and eventually approach a solution where the error is as little as possible—in other words, minimising the function's cost. A central part of the learning process is a mechanism called “backpropagation” that tells the network which mistakes it makes.

Figure 4.
A machine learning algorithm expresses a function between the data it is fed, and the data produced by the model: $y = f(x)$.

A tremendous amount of data is required to train and validate a model. Some models can automatically separate data into different clusters and see the context and patterns themselves, but many forms of neural networks require data with guidance. This includes examples of what should be entered and what the expected results should be. For this purpose, collections of open training and test data are created of various kinds, such as those for traffic images, with a label that classifies them as representing a traffic light, a pedestrian, etc.

As we have often returned to in this study, the amount, structure and quality of data are the most challenging parts of machine learning, which are both time consuming and costly.

For industry, the technology is useful in optimising the sourcing and the supply of materials; optimising internal and external logistics; planning production and forecasting demand and capacity utilisation; for process management and energy optimisation; for creating maintenance plans and working with preventive maintenance; for understanding customer behaviours; and for simulating cash flows. In summary, for progress in operational development.

The key to success using analysis as a method for operational development lies in good domain knowledge, that is knowledge of the company's operations and processes, and in the ability to create an analysis culture with a solid understanding of both mathematics and statistics. The tools needed are rapidly being commercialised and are becoming both cheaper and easier to use.

One of the simplest methods of classifying items through supervised machine learning, and also one of the most accurate, is called the "nearest neighbour" method. The technique is to measure the difference between two objects, or the distance between the objects. A large number of items are collected with each object labelled with a class affiliation. This is called the reference quantity. When a new unknown object is found, it is compared to the reference quantity until the object that differs the least from the new one is found. The unknown object is then considered to be of the same class as its nearest neighbour from the reference quantity.

Regression analysis, or regression, is a branch of statistics where the goal is to create a function that best fits the observed data. Linear regression is a method commonly used in machine learning contexts that has its limitations but compensates for these with simplicity, interpretability and efficiency. Simple linear regression assumes that a straight line can be adapted to the data and the regression equation can be described as $y = a + b x$. The intercept with the y-axis A and the slope B is calculated so that the error compared to the observed data is minimal. The error can be calculated using, for example, the least square method or maximum likelihood.

Logistic regression is an appropriate method of analysis when the dependent variable is binary. Like all regression analyses, logistic regression is a predictive analysis. Logistic regression is used to explain the relationship between a dependent binary variable and one or more independent variables.

With clustering, the aim is to divide the inputs into several groups. A difference between clustering and classification is that with clustering, it is not clear what the groups are in advance. This is typical of unsupervised learning.

2.8 How data becomes money: the machine learning process

A fundamental difference between neural networks and conventional computer programs is that the former develops in two stages. In the first stage, which can partly consist of regular programming, the width and depth of the network are determined, along with how it is to be provided with data and how it will be connected with the rest of the application and the process to be automated or optimised. The next stage is that the network begins to be trained.

The machine learning process (**Figure 5**) – the pipeline – begins with data collection in a procedure called ingesting and includes the cleaning and normalisation of the data so that, for example, numerical scales of values are aligned with each other. This is a time-consuming part of the process and can take as much as 80% of the project time.

The data sets need to be representative, and it is essential to analyse how bias can affect the model. The critical issue is how data is selected and how it is normalised. Distortions and prejudices built-in by algorithms is one of the most significant risks of machine learning because these undermine the entire purpose of the technology. The old truth about “you put garbage in, you get garbage out” applies in the highest, amplified, degree to machine learning.

In many cases, the process involves working with streaming data. In that way, it is possible to choose to first save the data in a database or to collect the data continuously to fine-tune existing models. The alternative is to build new models and train them with new data occasionally. The decision affects the choice of algorithms, as some algorithms are suitable for fine-tuning and others are not.

The next phase is comprised of training the model, or to put another way: determining the weights in the function relationship so that the model delivers the best possible results. The procedure for setting the weights is called hyperparameterisation. A hyperparameter is a setting that controls how a model is to be created based on an algorithm.

In reality, the process of teaching a model by seeking the “correct” weights can include millions, perhaps billions of iterations. To increase performance during modelling, there need to be multiple, parallel work processes running. That is copies of a program that run simultaneously at different locations. The parallelisation calculations utilise special hardware. CPUs originally used for graphic drivers (GPUs) have proven to be excellent in these cases.

In the summer of 2019, there was an emerging discussion over the impact of machine learning technology on the environment and the climate, given the energy-intensive GPUs that run the learning processes. A recent article from the University of Massachusetts [2] has found it is the marginal fine-tuning of models, in particular, that consume energy, thus leaving an imprint on the climate if the computers are driven by, for instance, coal power. This is also one of the reasons why Sweden is a country of interest to the localisation of data centres.

Figure 5.
The machine learning process or pipeline that begins with the collection of data and continues with data preparation and training of the model. The model is then implemented and run to provide accurate predictions.
Source: InfoWorld.

The final phase of the process is to use the pre-trained model. The model is now run with new, live data to make predictions that can then be translated into intermediate values such as quality, time and efficiency, which in turn can be assigned a price. Data has thus been transformed into money.

2.9 Typical problem types and methods of analysis

1. Classification, which means that based on a set of training data, new input data is categorised into one of several different categories. An example of classification is identifying whether an image contains a specific type of object or product of acceptable quality from a manufacturing line.
2. Continuous estimates calculate the next numeric value in a sequence based on a set of training data. These types of problems are sometimes described as “predictions”, mainly when applied to time-series data. An example of continuous estimates might be to forecast the sales demand for a product based on inputs such as previous sales, consumer preferences and the weather situation.
3. Cluster comparisons require systems that create sets of categories where the data instances have common or similar characteristics. An example of cluster formation is different consumer segments based on data from individual consumers, including demographics, general preferences and consumer behaviour.
4. Anomaly detection, which, with a set of training data, determines whether specific input data falls outside of a norm. For example, a system that has been trained with historical vibration data from a machine can determine whether a new data batch suggests there is a fault in the machine. Anomaly detection can be considered a subcategory of the classification problem.
5. Ranking involves algorithms being used for information retrieval problems where the results of a request need to be set against a criterion. Recommendation systems that, for example, suggest prioritised purchases of products use these types of algorithms to sort the suggestions by relevance before they are presented.
6. Recommendations are systems that provide recommendations based on a set of training data. A typical example is a system that suggests a “next purchase” for a specific customer based on the buying patterns of similar people and the observed behaviour of the particular person.
7. Data generation requires a system that can generate appropriate new data based on the training data. For example, a music composition system can be used to create music pieces in a particular style after being trained on pieces of music in that style.
 - Bad - means that the quality of available data is substandard, even though it has a clear physical significance. This makes it difficult to compensate for flaws in quality by adding more data of more or less the same type. The latter is a method that can work for applications using deep learning, such as image recognition.

- Broken - means that data that has been collected to train a machine learning model lacks the essential qualities of validity/relevance and contain error conditions. This then leads to false positives or negatives in the online implementation of the model. This is a serious problem because even a few or occasional erroneous statements can endanger the reliability of the system, and industrial AI applications typically have significant potential to impacts on assets and personal safety.
- Background – means the data patterns in industrial contexts can be transient. The process involved is volatile, fluctuating and fast. Interpreting such data often requires in-depth domain knowledge, and it's not enough to simply dig for more numerical data. In addition to precision around predictions and quality of performance, an ability to find the roots of possible anomalies is also required.

2.10 The data challenge

One of the biggest challenges with AI concerns the quality of the data needed to make predictions, create forecasts, and recognise patterns. It is a widespread issue, and a great deal of monotonous, routine work takes place behind the sometimes simplified depictions of AI that we see.

In autumn of 2018, BBC News [3] brought attention to a new concept: labelling farms. This is a rapidly growing global sector involving data centres that have been located in low-cost countries for economic reasons. Labelling farms today employ thousands of people whose only task is to help AI algorithms interpret data.

Pixel by pixel, the content of millions of images is classified; a car is identified as a car, a dog as a dog, a road sign as a road sign, and so on so that self-driving cars can recognise real-world objects. Similar data challenges are being encountered everywhere that AI is to be applied. The high cost of data preparation means that there are financial incentives to solve the data problem, and many projects are being carried out with the help of even more AI in a bid to find new solutions and better methods.

Industrial AI involves transforming raw data into “intelligent” predictions to make decisions. In industrial processes – in a steel mill or a paper mill – quick decisions are made in real-time at the millisecond level in models representing physical reality. Several challenges arise in such processes. Real-time requirements mean that the cost-effective and almost endless resources of the cloud need to be supplemented with locally distributed computational and storage capacity, also known as “edge”. But the most fundamental challenges also concern the availability and the quality of the data.

Since the 1980s, industrial control systems have been producing enormous amounts of information. Industrial Big Data is available in every factory, and yet while industrial data is generally well structured, it often lacks quality. You sometimes hear talk of the “three B’s” of Industrial Big Data: Bad, Broken & Background.

Teams of people who possess both process knowledge and computer science are required for the development of good adaptive models. There is also a need for method development, with experience teaching us that data preparation demands a disproportionate amount of work. This is a serious issue that needs to be continuously addressed and prioritised, lest it becomes an obstacle in releasing industrial value.

2.11 Solutions to data deficiency and manual intervention

Much of the success of modern AI applications are based upon bottom-up strategies within which models are trained using large, well-structured data sets typically collected via the Internet. For example, the GPT-2 text bot was trained using a data set of eight million web pages. Intelligent assistants like Apple's Siri or Amazon's Alexa use thousands of terabytes of data to perform their tasks, and self-driving cars consume about forty terabytes per eight hours of driving, according to INTEL.

For operational industrial applications, large amounts of information are being collected. However, critical processes, in particular, lack the volumes needed to train good models. There is a lack of data in marginal or edge cases, and it is not always easy to deliberately address such deficiencies (by inducing errors in physical processes). The errors they represent correspond to high costs due to significant production disruptions. This is a problem that also applies to other, normally data-rich applications. One of the considerable challenges in the development of autonomous vehicles is managing the most unusual of operating cases. Another characteristic of today's AI technology is that it tends to easily become "confused" if circumstances deviate significantly from what is expected.

Methods are in development to overcome these weaknesses. Similar to human intelligence, they involve working in a more flexible, top-down manner, which allows for reduced data requirements and enhanced speed. There are a number of trends related to the development of more natural systems worth keeping an eye out for soon.

2.12 Eight trends

1. The first trend involves giving robots conceptual properties (both physical and artificial) that in turn give them a greater ability to perceive themselves – and their environment. See the text box on page 45, describing how researchers at the University of Columbia have succeeded in giving a robot such properties.
2. Another developmental avenue involves something of a renaissance of the concept of "expert systems" within which computers become better at doing what human process operators do by making adjustments in real-time to optimise processes.

Siemens has developed data-efficient methods such as these based on "reinforcement learning" to control the company's gas turbines. In this area, traditional neural networks would take up to a hundred years to learn the complex combustion processes. The method has subsequently been developed to increase the efficiency of the company's wind turbines. Google is also using technology to reduce the energy consumption of its data centres successfully.

3. A third way to address the weaknesses of today's AI algorithms is to give computers more common sense. According to an article in Harvard Business Review [4], the Allen Institute for Artificial Intelligence is working on developing test data that can be used to verify what common sense means to a machine. Meanwhile, DARPA is investing USD 2 billion in AI research through, among other things, creating models that mimic human cognition. And Microsoft and McGill University have jointly developed a system for distinguishing ambiguities in natural language; a challenge needs to be solved if, among other things, computers are to be able to communicate with human beings in a human way.

4. A fourth track is the possibility of letting computers make similar balances of probability assumptions to those that humans intuitively make. This is being made possible through stochastic Gaussian processes that can function and recognise patterns within limited data sets and learn from experiences. Another feature of this method is that processes are traceable if something goes wrong, unlike with the black boxes of neural networks.
5. Yet another method of advancement is Probabilistic Programming for the applications described above. This method brings together the best practices for mimicking human intelligence such as probability theory for modelling, statistical methods for drawing conclusions, and neural pattern recognition networks, along with symbolic program languages that hold the system together.
6. “Explainable AI” is an adjacent developmental track. The black box phenomenon of machine learning can be problematic. Therefore, systems must be able to justify how they have arrived at their conclusions. It’s also essential to ensure that human beings can have trust in the way that such systems arrive at their results and decisions when, for example, traffic situations, legal support or medical diagnosis become automated.
7. Federated machine learning is another method showing promise. The idea was launched in 2017 by Google as a concept within which the ability to train a model is decoupled from the up-until-now necessary central storage of data in the cloud. The method can train a single machine learning algorithm over several decentralised servers that store data, without actually exchanging data with other servers. It allows multiple actors to build a common, robust machine learning model without sharing data, thereby addressing critical issues such as data privacy, data security, data access rights, and access to heterogeneous data. It also enables capacity in distributed applications.

This way of working is based on the idea that a distributed device, such as a phone, downloads an existing shared model, improves it by learning the data that is locally available on the phone/device and then summarises the changes as a small concentrated update. Only the update is sent to the cloud, via encrypted communications, where it is immediately computed and integrated with other user updates to improve the shared model. All training data remains on the local unit, and no individual updates are stored in the cloud.

The technology can contribute to breakthroughs for industrial operational applications based on conventional automation or IoT, where distributed capacity is both a prerequisite and natural in the current concept of control and monitoring.

8. Finally, AutoML or Automatic Machine Learning looks like the holy grail for solving the many and long routine steps found in today’s state of the art technology. Automated machine learning involves automating the process from end-to-end. As we have noted, typical machine learning projects involve extensive pre-processing before the dataset can be made available for actual machine learning.

Pre-processing is followed by a selection of an algorithm, hyperparameterisation and fine-tuning to maximise predictive performance in the final model. In addition, many of these steps require both experience and specialist knowledge. What could

be more logical, then, than to suggest AutoML as an artificial intelligence-based solution to these growing challenges? Automating the process would be an effective productivity-enhancing method, which besides would be likely to provide solutions and models that exceed manually designed ones.

AutoML solutions with drag-and-drop-based user interfaces, and that do not require any coding in the ordinary sense are now on the market and are offered by all major platform providers such as Google, MS Azure and IBM, along with many specialised smaller companies. The technology is evolving rapidly and will further lower the **threshold for users**.

Industry case study: BillerudKorsnäs in Gävle, Sweden - Deep Process Learning.

DEEP is a project that will show how deep process learning (deep learning) can be used for the next step in process automation. The project takes advantage of data that already exists in process control systems and uses it to suggest the measures required to improve selected key performance figures. The project, which is a collaboration with PiiA, consists of a consortium between BillerudKorsnäs, Peltarion, PulpEye and FindIT.

“The forestry sector has an advanced supply chain with multiple levels of complexity and difficult, resource-intensive processes...”

The process industry accounts for almost half of all industrial production in Sweden. So, the achievement of efficiency and productivity improvements within it is certain to have a significant impact on the Swedish economy. The forestry sector has an advanced supply chain with multiple levels of complexity and difficult, resource-intensive processes: from felling to barking and chipping the wood, to boiling, washing and bleaching the pulp before it reaches the paper machine to be refined to produce paper and cardboard of various grades.

Process industries produce vast amounts of data and have a high degree of automation, but they also face a variety of challenges. These challenges cannot always be addressed through traditional analysis methods. As such, the data produced can be a valuable asset, capable of being refined through AI to generate insights, predictions and automation algorithms – thus creating the next stage of productivity, quality and automation.

BillerudKorsnäs is a forestry company that supplies packaging materials and packaging solutions. The company has three divisions: Division Board, which manufactures and sells liquid and non-liquid packaging board, as well as fluting and liner; Division Paper, which produces and sells high-quality kraft and sack paper; and the Solutions Division, which meets the needs of brand owners for efficient packaging solutions and systems.

During a feasibility study for DEEP, BillerudKorsnäs and Peltarion jointly led a machine learning project to predict the kappa number of pulp after boiling. The kappa number is a measure of residual lignin in the pulp and determines the boiling process required for different pulp qualities. The project was successful and resulted in a useful technique for predicting the kappa number. This success encouraged further development of the approach in other process steps.

“The successful use of machine learning as a tool is based upon a deep understanding of the processes that are to be optimised”.

The challenge

Paper machine 4 in Gävle is a cardboard machine that manufactures liquid packaging board for juice and milk packaging, among other things. The purpose of the DEEP project has been to realise the efficiency potential identified in the manufacturing process by proposing optimal machine operational parameters. An essential feature of the finished liquid packaging board is the carton's bending stiffness. This property is determined by complex relationships between the different stages of the manufacturing process, not least by the pulp's fibre properties. The goal is to produce strong packaging using less raw material.

To meet the quality objectives at optimum production speed, process settings must be continuously evaluated and adjusted. In the DEEP project, data is being collected to support the online optimisation of such decisions. The data used in the project consists of high-resolution microscopy images from PulpEye's analyser which provides information about the pulp's fibre properties and camera images from the drying cylinder which provides information about the dewatering of the pulp, in combination with measurement values from different sensors in the system. In the next step, data will be used to develop a suitable model to predict quality properties.

During the DEEP project, many different attempts were made using various methods, including deep learning with Peltarion's self-developed platform.

The experience

BillerudKorsnäs has formed a digitisation team with different competences from different parts of the organisation, and that initiates and runs transformation projects. The company's various AI initiatives are part of that transformation process.

BillerudKorsnäs' experience shows that deep learning technology is ripe for use in various types of classification problems and for further increasing the degree of process automation. The process industry is characterised by a combination of large amounts of data and a high degree of automation, which partly produces conditions that differ from other fields that apply deep learning and machine learning. Over time, the technology will find its place in process analysis and control and will solve many more problems that affect efficiency, quality and logistics.

One of the essential takeaways from BillerudKorsnäs' AI projects is the need for domain knowledge and the ability to formulate the right problems. The successful use of machine learning as a tool is based on a deep understanding of the processes that are to be optimised.

BillerudKorsnäs is continuing its work on developing processes with the help of AI, and another project will be launched in collaboration with PiiA in the spring of 2019. This time, it will be led together with Finnish Quva OY as a data analytics provider.

Sources: PiiA, BillerudKorsnäs, Peltarion.

2.13 Man and machine: collaborative intelligence

It was once said that we should, "Let the machine take care of the details and let the man think and dream". And, as Anders Ynnerman, a professor at Linköping University, states, "for every AI system that we have where we add on the human aspect, we get a much better system."

At the same time, there is a fear that AI will eventually push people out of the labour market. The latter is hardly inevitable or even the most plausible outcome. Never before have digital tools been better suited for collaboration with people. And while AI will surely change the way work tasks are performed, and who performs them, the role of machines in future will be to reinforce and supplement human abilities rather than to replace them.

The concepts of collective and collaborative intelligence are also worth bearing in mind. Models where people's intellectual capacity can be increased through smart, collaborative methods, either working with other people or with machines, will have a significant impact on industrial development. Man's abilities in leadership, teamwork, creativity and social interactions will complement AI's speed, scalability and quantitative ability to keep track of large complex data sets. Industrial activities require both.

But the above line of reasoning also demonstrates the need for changed processes and in many cases, radical transformations of both business activities and the way people and machines interact on a practical level. An article in Harvard Business Review [5] notes that the business effects of artificial intelligence depend

on the ability to “rethink” activities so that they both incorporate AI and cultivate related abilities in human employees, in addition to allowing creative experimentation and having clear AI strategies. Last but not least, it is crucial to managing data in both a relevant and responsible manner.

"For every AI system that we have where we add on the human aspect, we get a much better system."

Anders Ynnerman, Professor at Linköping University

AI will lead to more automation and more advanced automation. One of the significant advantages of automation is avoiding errors caused by people not being able to repeat tasks efficiently. A robot that is asked to do the same motion a thousand times makes the same motion a thousand times – as long as the sensors and the mechanics work. A person might be able to perform it three times but is at the same time a master at interpreting their senses and dealing with new, unexpected situations.

The process of how this might happen is not yet clear and making machines that act in a human-like manner is a complex matter. The recent accidents involving one of a highly advanced Boeing aircraft model have, in a frightening way, also shown that for every human mistake that a machine eliminates, there is a risk that a new one will be introduced. There are endless possibilities for misunderstandings to occur between human intelligence and machines. In the industrial context, the challenge boils down to establishing collaborative intelligence, and how well the interface between human and machine works. This developmental field is known as UX – user experience – or in the AI context, it's perhaps more appropriately called MMC, man-machine communication.

Issues with misunderstandings and mistakes have the potential to intensify as the degree of automation increases further. Humans will no longer have full control over machines. Overall, this will lead to a decrease in those parts of industry domain knowledge that include artisanal process knowledge. At an operational level, the challenge for the machine operator will be to monitor a process over a significant amount of time and to be prepared to take over the moment something goes wrong. Problems in this area have the potential to be costly in the process industry and utterly catastrophic within aviation.

One conclusion that can be drawn is that machines that do not allow people to keep up with the processes they are managing aren't optimal in events where people are forced to take over. Another conclusion is that the best kind of automation is not necessarily where the computer automatically does most of the work, but rather where there is an optimal distribution and a realisation that people and machines will probably never understand each other perfectly. We have two pilots in the cockpit and two operators in the control rooms, and unfortunately, both can sometimes be expected to do unexpected things.

2.14 Looking ahead

In this section, we have skimmed over some of the concepts and constructs that may come in handy from an applied industrial perspective. Of course, there are countless other aspects of AI that could potentially be taken into account when assessing a technology which proponents claim to be “intelligent.” Many of these issues relate to morals and ethics. As society and industry move ahead, we will likely encounter machines with questionable intentions and distorted development, whose intent is to benefit individual stakeholders. AI will influence people's attitudes; false correlations and self-reinforcing feedback will eventuate – and

algorithms may influence reality to gain even more influence, even though their base assumptions are false. The origin and quality of data will continue to be an issue and, last but not least, we will face uncertainty around what is real and true: will we, in future, be able to trust what we see and hear? Will we be able to trust pictures, movies and sounds?

From an industry development point of view, our hopes for AI and machine learning might be for them to provide greater flexibility than that currently found in our simple neural networks which are only capable of performing one task at a time and are expensive and arduous to retrain. We might also hope to see significant productivity gains in system development, while there is also room for improvement in the deployment of models.

But we can rest assured that these are areas that are currently being addressed by research. Likewise, the substantial data requirements, the need for manual intervention, and the problems with edge data all need to be addressed. The actual learning process, with its hyperparameterisation, needs to be further automated. Another potentially growing concern is the lack of transparency in neural networks, which, for the most part, resemble black boxes.

It's impossible to know how and when these issues will be addressed. It could take years, or there may be sudden breakthroughs, such as when the AlphaGo defeated one of the world's best Go players with the help of reinforcement learning. But it does not change the fact that AI and machine learning are already powerful enough tools to change the industry, and that those who acquire knowledge, experience and an upper hand when it comes to applying the technology have everything to gain.

"Will we in the future be able to trust what we see and hear? Will we be able to trust pictures, movies and sounds?"

2.15 Glossary

The article contains some terms that may need clarification. Key terms include:

Artificial Intelligence (AI)

The term "Artificial Intelligence" (AI) does not have any clear definitions or delineations. AI research itself is both specialised and spread across many subfields. For this analysis, we have chosen the definition also used by Vinnova in the study of Artificial intelligence in Swedish business and society, 2018".

This is: "The ability of a machine to mimic intelligent human behaviour. Artificial intelligence is also the designation of the science and technology field that aims to study, understand and develop computers and software with intelligent behaviour."

When we talk about AI in an industrial context, we are primarily referring to machine learning technology with neural networks.

Algorithmisation

Algorithmisation is a mega-trend within which more and more value-adding activities are managed and controlled by algorithms instead of human beings.

IndTech

IndTech is used to describe the development, companies and markets that arise when traditional automation and industrial IT meet digitisation. IndTech companies include:

- Suppliers of industrial automation, such as ABB or Siemens.
- Suppliers of industrial IT software, such as SAP or IBM.
- Providers of digital platforms, such as Microsoft or Amazon Web Services.
- IoT providers, such as Ericsson or Nokia, and operators, such as Telia or Telenor.
- System integrators and machine suppliers who base their process or mechanical engineering offerings on digital technology. These include companies such as Sandvik, Epiroc, Valmet and many more.

Platformisation

Platformisation can be used to describe the general movement of various companies' automation and IT support to the cloud, and also to describe the movement of platforms created by open standards to platforms owned and controlled by a particular actor. Because the value of a platform tends to increase for all involved as more people use it, there is a tendency for already-large platforms to grow even bigger.

Operational Development - OD

We have chosen to use the term operational development to encompass the operational changes in processes or in organisations that lead to increased efficiency or increased customer values. Within this area, AI can be a potent tool.

Author details

Örjan Larsson

Blue Institute and Strategic Innovation Program PiiA, Västerås, Sweden

*Address all correspondence to: orjan.larsson@blueinst.com

IntechOpen

© 2020 The Author(s). Licensee IntechOpen. Distributed under the terms of the Creative Commons Attribution - NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited.

References

[1] Blue Institute. PiiA Insight. Västerås, Sweden: PiiA; 2018

[2] Strubell E et al. University of Massachusetts. In: Energy and Policy Considerations for Deep Learning in NLP. Ithaca, NY, USA: Cornell University; 2019

[3] BBC News. Why Big Tech Pays Poor Kenyans to Teach Self-Driving Cars. United Kingdom: BBC; 2018

[4] James Wilson H et al. The future of AI will be about less data. In: Not More. Brighton, Massachusetts: Harvard Business Review; 2019

[5] Collaborative Intelligence: Humans and AI Are Joining Forces. Brighton, Massachusetts: Harvard Business Review; 2018