

AI in Logistics and Supply Chain Management

Robert N. Boute, Maxi Udenio

June 2021

To appear in edited book *Global Logistics and Supply Chain Strategies for the 2020s: Vital Skills for the Next Generation*, Springer Series in Supply Chain Management.

Abstract Artificial Intelligence (AI) applications are pervading logistics and supply chain management. The widespread availability of data, combined with sustained improvements in computing power, provide new opportunities to improve supply chain decision-making. Data may originate from digital logistics applications, or the connectivity of assets through Internet of Things technologies. AI can also facilitate the automation of well-defined workflows. Yet, while AI may appropriate certain tasks, we believe it will not make the job of the logistics planner obsolete. AI empowers and augments human capabilities. This chapter explores and demystifies the opportunities of AI for logistics and supply chain management.

1 Introduction: Digital Logistics

Artificial intelligence (AI) refers to the ability of a computer or computer-controlled robot to perform tasks commonly associated with human beings. The *intelligence* in the term implies that the task being performed by a machine, script, or algorithm would require the use of intelligence, were a human to do it.

Although AI has been around since the late fifties, it has only become a mainstream concept since the last decade. AI now powers our smartphones, our TV recommendations, and even our photo libraries. Also in business, AI has transformed from an obscure term to a buzzword. In a 2021 survey by Accenture, 77% of

Robert N. Boute
KU Leuven and Vlerick Business School, Belgium, e-mail: robert.boute@kuleuven.be

Maxi Udenio
KU Leuven, Belgium, e-mail: maxi.udenio@kuleuven.be

executives state that their IT architecture is becoming critical to the overall success of the organization.¹ In the healthcare sector—transformed in 2020 by the COVID pandemic—the confidence in AI is even stronger. Reportedly, 98% of healthcare executives have developed an AI-strategy plan, among which 44% have already implemented an AI strategy.² Other business surveys are similarly strong. McKinsey, for example, notes that nearly 58% of executives surveyed have already embedded at least one AI capability in their company.³ The message from the industry is clear: AI is here to stay, and the companies that learn how to adopt and scale it are poised to enjoy a competitive advantage.

In logistics and supply chain management, analytics and computer support have been around for decades. Supply chain planners, for instance, use software tools that process historical data to forecast demand; many enterprise resource planning (ERP) systems automate the decision of when and how much to order; and warehouse and transportation management systems optimize storage and transportation operations. Each of these supply chain support tools can be run as a siloed application, or integrated with other business operations, such as financial accounting or supplier relationship management. Integration facilitates data sharing to a common data platform. When the platform is also accessible over the Internet, web-based tools provide remote access and connectivity with third-party applications through application programming interface (API) software. Cloud-based service offers the additional flexibility to scale up IT infrastructure to accommodate temporary computing needs.

The recent breakthrough in digitizing logistics operations comes from real-time connectivity of assets to the data platform: machines, vehicles, and devices can now be monitored via sensor technologies that capture all sorts of data in real-time. In addition, when sensors become impractical, operators can provide feedback information through mobile and wearable devices. This extensive connectivity is known as the fourth industrial revolution, also referred to by the term *Industry 4.0*.

Such connectivity provides (quasi) real-time visibility over all workflows. A “digital control tower,” in analogy to the airport control tower, can provide visual alerts that warn of inventory shortfalls, or process bottlenecks, before they happen. Using simple control algorithms, teams on the front line can course-correct even before potential problems become actual ones. Furthermore, the availability of historical data can give rise to increasingly sophisticated algorithms which add additional intelligence to the control rule: predictive analytics learn from historical data to obtain patterns and correlations not evidently detected by humans. By means of a *digital twin* of its physical operation, real-time analysis and optimization can even prescribe decision-making where users make decisions based on what intelligent agents recommend.

The digital control tower providing real-time information, potentially augmented by predictive diagnostics and analytics, may support logistics and supply chain man-

¹ Accenture: Technology Trends 2021. <https://acntu.re/3nN0ABW>

² 3rd Annual Optum Survey on AI in Healthcare. <https://bit.ly/3azuVOU>

³ McKinsey global AI survey. <https://mck.co/3auUtN5>

agers in their decision-making. As we discuss in this chapter, these ‘smart’ decisions have the potential to bring about more efficient, more resilient, and even more sustainable supply chains. Observe, however, that these AI tools do not necessarily perform an entire workflow. Instead, each delivers a predictive component to assist someone in making a decision. AI can take over some, but not all, tasks. In fact, in the majority of supply chain AI implementations to date, humans still have the last word. AI does not imply—by itself—autonomous decision-making (Boute and Van Mieghem, 2021).

2 Smart Logistics

Back in 2017, The Economist published an article titled, “The world’s most valuable resource is no longer oil, but data” (The Economist, 2017). The use of digital applications as well as the connectivity of assets through, e.g. sensors and digital control towers, generates large amounts of data (possibly in real-time). The question that now arises is how such data can be leveraged to improve the level of *intelligence* of logistics and supply chain decision-making. Notice that the use of data in logistics is not new—we have been transporting goods around the world based on data-driven forecasts for decades. What is new is the sheer volume of data that we now generate, store, and share. These data have the potential to make logistics and supply chain control more adaptive and *smarter*.

In traditional data-driven applications, one typically uses one or—at most—a few sources of data, such as historical demands or current inventory levels. As long as the input variables remain ‘countable,’ one can implement (or even program) if-then instructions to support (or even automate) decision-making. The integration of various digital applications, in contrast, generates a data pool of different sources, collected automatically through sensors (*Internet of Things*) as well as manually through mobile and wearable communication devices (known as the *Internet of People*). When the number of data sources grows rapidly, the ensuing mountain of data makes the explicit enumeration of if-then instructions infeasible.

This is where machine learning comes into the picture. Whereas AI is the umbrella term for all computer rules that mimic human intelligence (including if-then instructions), machine learning is the subset of AI where an algorithm *learns* to mimic human behavior and makes its own decisions. Machine learning algorithms are in essence *prediction machines* that perform a task without using explicit instructions (Agrawal et al., 2018).

A milestone for the mainstream use of machine learning is the victory of the algorithm AlphaGo over the world champion of the Chinese board game Go, Lee Sedol, in March 2016.⁴ The ancient game of Go is played on a board with $19 \times 19 = 361$ positions, each of which can contain a black, white or no stone. It therefore has $3^{361} \approx 10^{172}$ possible states; several orders of magnitude more than chess⁵, and

⁴ A breathtaking documentary of this victory is available at <https://www.alphagomovie.com/>

even more than the number of atoms in the universe. Due to the sheer number of possible states, devising explicit instructions (prescriptions) for how to play each state is impossible. Hence, Go was considered the holy grail of AI.

Instead of using brute force to calculate all possibilities, Alpha Go used machine learning and neural networks to mimic the learning process of a human brain. The system was not pre-programmed. Instead, it was fed data of historic games and allowed to play itself so as to improve its win rate through trial and error.

The same machine learning power can also be applied to logistics and supply chain management. To understand possible applications of machine learning, it is instrumental to differentiate between the different forms it can take. Broadly speaking, one can divide machine learning into supervised, unsupervised and reinforcement learning.

Supervised learning

Supervised learning is perhaps the best known (and most frequently used) type of machine learning. An extensive set of labeled training data with the ‘correct’ answers, which serve as “supervisors,” is used to estimate a ‘mapping function’ f that *predicts* output Y based on a set of input data variables X . The mapping function is determined so as to minimize the prediction error ε ,

$$Y = f(X) + \varepsilon.$$

The algorithm makes predictions on training data (i.e., data for which both X and Y are known) and continuously improves its mapping function by comparing its output to the ‘correct’ answers in the training data. Learning stops when the algorithm achieves an acceptable level of performance.

The goal is to approximate the mapping function f so well that you can predict the outcome Y for *new* input data X . It is possible that the mapping function performs well on the training data and poorly on new data that was not encountered during its training. In such a case, there is a good chance that the model is overfitted. The most obvious remedy to avoid overfitting is to enlarge the training data set.

Supervised learning is, for instance, used to classify images. In such a case, an algorithm learns the mapping function based on a set of labeled images. Each time the algorithm is fed with new training data, the mapping function can be improved. This explains the importance of (lots of) data to train a good learning algorithm. The next time you are prompted with a “Captcha”⁶ when filling out a form on the Internet to prove that you are fully human, do know that your humanoid clicks are feeding supervised learning algorithms behind the scenes and, in doing so, improving object recognition of traffic lights, street signs, etc. The more labeled data you feed the

⁵ Approximations place the number of possible states in chess at around 10^{46} .

⁶ CAPTCHA stands for the Completely Automated Public Turing test to tell Computers and Humans Apart. Captchas are used to differentiate between real users and automated users, such as bots.

algorithm, the more accurate it becomes. Thus, with each Captcha we complete, we are implicitly contributing towards autonomous driving.

Supervised learning can be used in logistics to predict a multitude of observations:

- *Demand forecasting*: given enough computation time and data instances, supervised learning can learn how sales are influenced by a wide variety of features, such as the marketing mix (price, promotions, discounts, advertising), seasonality, calendar events, weather forecasts, lagged sales data (sales from previous periods), and even social media reviews, using tools such as text mining and natural language processing (Cui et al., 2018).
- *ETA prediction*: based on data from a transport control tower that tracks real-time information of the trucks and keeps track of the realization of the planning, one can build a predictive model to classify whether a truck will be on time or not. These data can additionally be complemented with external data variables, such as weather or traffic information (Kolner, 2019).
- *Throughput times at customs*: based on a historical data set of packages that passed customs administration, one can predict waiting times at customs based on the characteristics of the package, such as source of origin, weight, size, etc. (Flows, 2019)
- *Downtime prediction*: the time-to-failure, or remaining useful life, of (parts in) truck cabins, rail wagons or machines can be predicted by integrating measurements from specific variables such as condition monitoring (e.g., vibration) and operational (e.g., usage) data in the estimation process (Si et al., 2011).

Supervised learning can also be used to predict cases in which a prescription made by a computer system is likely to be overruled by a human, and adjust the original prediction accordingly. For instance, packing workers at e-commerce warehouses sometimes deviate from the order packing instructions (which items to pack in which sequence and in which box) prescribed by the system. These human adjustments are typically necessary to pack the box, but they increase packing time and reduce operational efficiency. By tracking when packing workers deviate from the system prescriptions, a machine learning algorithm can predict when workers are more likely to switch to larger (or different sized) boxes. By pro-actively adjusting the algorithmic prescriptions of those ‘targeted packages,’ the rate of switching to larger boxes—and thus the average packing time—can be reduced (Sun et al., 2021).

Another example where human decision-makers approve or override algorithmic recommendations is in the review of sales forecasts. Even in cases where the algorithm is well-tuned (and thus overrides are rare), the human decision-maker is still burdened with reviewing a potentially large number of recommendations. By analyzing a history of sales forecast reviews, a supervised learning algorithm can predict whether or not the decision-maker will modify the recommendation and whether such a modification will improve or impair the performance of the system. Using these predictions, a significant portion of the order recommendations can then be automated with little, or even a positive, impact on performance, thereby freeing up the decision-makers’ time for other value-added activities (Imdahl et al., 2021).

Unsupervised learning

Unsupervised learning algorithms *describe* patterns or groupings of data given a set of unlabeled observations, i.e., without knowing the correct answers. These algorithms are basically left to their own devices to discover patterns and information that was previously undetected. The goal of such analyses is to group a set of data points, such that data points in the same group or cluster are more similar to each other than to those in other clusters. Often, such clusters represent data groups with distinctive characteristics for which specific operational policies can be designed.

Customer data can be screened to discover groupings of customers according to, for instance, similar purchasing behavior or common profiles through similar combinations of customer characteristics. Likewise, product data can be used to cluster products into groups according to, for instance, their product life cycle stage. Such clustering can be useful for customer or product segmentation where a distinct logistics distribution approach or different target inventory levels are required. One can also customize the prediction model for each group to achieve higher accuracy (possibly combined with supervised learning). Policies or predictions based on clustering typically provide superior performance compared to unclustered prescriptions.

Unsupervised learning is also used to analyze customer orders to identify relationships or groups of items that are frequently purchased together. Such ‘market basket analysis’ can, for example, be helpful to optimize product placement and layout in an e-commerce warehouse to, amongst other objectives, improve order picking productivity. A higher order picking productivity is one of the key drivers to reduce the cost to serve an e-customer.

Finally, unsupervised learning can be used to identify observations which do not conform to an expected pattern, known as ‘anomaly detection’ or ‘outlier analysis’, such as abnormal delays in transportation times. This detection can spur further analysis to predict—or even better: prevent—future delays.

Reinforcement learning

Reinforcement learning is different from (un)supervised learning: rather than predicting or describing an outcome, it *prescribes* which decision or action to take, based on the current state of the system, while taking the future impact of these decisions into account.⁷ One could say it tries to predict the optimal action given the current situation. Reinforcement learning also requires training for the algorithm to *learn* how to convert inputs into outputs. However, instead of comparing the output directly to the ‘correct’ answers (as in supervised learning), training a reinforcement learning algorithm relies on trial and error by simulating sequences of states, actions, and rewards. These simulations can be fed by either historical (real) data or simulated data, provided an accurate data generation engine. It is up to the model to figure out how to perform the task to maximize the reward, starting from totally random trials

⁷ In mathematical terms, it formulates a problem as a Markov decision process, where an action taken in a given state transitions the system to a new state and generates a reward (or cost).

and finishing with sophisticated tactics and superhuman skills. By leveraging the power of search, performing numerous trials and reinforcing specific actions that generate high rewards (or low costs), the algorithm learns which actions provide the best results in any given state. In contrast to human beings, a reinforcement learning algorithm can gather experience from thousands of parallel trial runs if it is run on a sufficiently powerful computer infrastructure.

The aforementioned machine learning algorithm, AlphaGo, uses reinforcement learning. The same algorithms have the potential to also be applied in logistics, where the optimal decisions are unknown due to the sheer problem complexity and thus cannot be captured—or programmed—by simple if-then instructions, for instance:

- *Multi-source or multi-mode replenishment*: when you have access to multiple sources to replenish your inventory, reinforcement learning can support the decision of how much to replenish from a cheap offshore supply and how much to source locally at higher cost. Similarly, it can be used to combine multiple transport modes in parallel, where part of the shipment is shipped using a slow, but more carbon-friendly transport mode such as rail or waterways, and part of the shipment is shipped using a more responsive mode such as road or air freight (Gijbrecchts et al., 2020).
- *Joint replenishment and collaborative shipping*: to synchronize the replenishment cycles of individual products or companies and facilitate collaborative shipping, machine learning algorithms can be used in a control tower setting that tracks the supply chain flows in real-time (Vanvuchelen et al., 2020).
- *Perishable inventory management*: managing inventory of products with an expiration date is notoriously complex, as one should not only take into account the inventory levels (and those in transit), but also the age distribution of the goods in inventory. As the optimal inventory policy for this problem is intractable in many cases, reinforcement learning algorithms can develop good performing heuristics through learning (De Moor et al., 2021).
- *Omni-channel supply chains*: when managing inventory for multiple channels, reinforcement learning can prescribe how much should be stocked centrally to leverage inventory pooling benefits, which products to stock locally to ensure fast delivery, and from which warehouse different customer orders should be filled.

These AI applications not only support logistics planners to improve their logistics costs, they also build more resilience into the supply chain through higher responsiveness and agility to real-time events or disruptions. Moreover, as we discuss in the next section, they may reduce the carbon footprint of logistic distribution systems without compromising on service levels.

Implementing Machine Learning Algorithms

Thanks to the open nature of the machine learning community, code for most (un)supervised and reinforcement algorithms is freely available online. Most algorithms can also be re-used in different problem settings with minimal changes. As

a result, machine learning may be seen as a general-purpose technology where one is only required to outline the prediction objective along with the data sources (input variables) needed to make the prediction. There is no (or little) need to program the algorithm itself.

Moreover, as software development kits (a collection of software development tools in one package) are being designed to facilitate the interaction between programmers on the one hand and end users on the other, the focus of machine learning applications is reverting to the collection of (clean) data. Machine learning models thrive on the availability of (a large amount of) data. The more labelled data that can be fed to an algorithm, the better it becomes at generating accurate predictions.

What all machine learning algorithms have in common is that training an algorithm to learn good predictions or prescriptions can be computationally very expensive. Therefore, most applications resort to cloud computing. The flexibility and scalability of the cloud make it possible to temporarily utilize peak capacity in computing power in only those periods where the algorithms need to be trained.

Scale provides a significant advantage in machine learning. More data improves the models which, in turn, improves the data. Thus, larger firms with the necessary resources, as opposed to smaller ones, are typically at an advantage to develop superior models and improve their data collection, simultaneously creating a self-reinforcing loop. The best known examples of this scale advantage are Amazon and Google. We are aware that data-driven markets can lead to a “winner takes all” situation. The fact that industry giants, such as Amazon, have orders of magnitude more computing power at their disposal than the large majority of firms on the planet combined, forces us to ask the question at what point the scale advantage is just too much.

3 Sustainable Logistics

The value of AI stems from its ability to (semi-)autonomously process data to produce predictions or prescriptions. In logistics and supply chain management, these are typically used with the intent to optimize operational parameters, such as customer service levels or inventory holdings. The same characteristics also make AI a powerful tool to improve other objectives—most notably sustainability development goals (SDGs).

The argument for AI within sustainable supply chain management is mostly based around efficiency improvements. Some relevant smart logistics examples from the previous section include increased demand forecasting accuracy, reduced inventory/production management redundancy, and improved supply chain network designs. All potentially achievable through AI, they directly decrease pressure on the environment *in addition to* its operational advantages.

Improved forecasts reduce waste in two important ways: by reducing the amount of safety inventory required to account for uncertainty, and by reducing secondary flows and waste stemming from returns and obsolete stocks. Moreover, given that

demand of perishable items is notoriously difficult to forecast, and given that AI has shown promise in the area, there is an argument to be made for the potential environmental impact of AI models for demand forecasting—and by extension, inventory management—of perishable goods.

The use of control towers to combine geographically compatible shipments, either by bundling thereof or through effective backhauling, relieves the carbon footprint of logistics distribution systems. The main reason being that improving the load factor, through bundling or effective backhauling, reduces the number of vehicles on the roads. Fewer vehicles, in turn, means lower emissions of harmful greenhouse gases, less congestion and fewer chances of accidents. Collaborative shipping, where shipments from different companies are bundled, also facilitates the shift to greener transport modes, such as rail or inland waterways, thanks to the economies of scale that may be achieved as a result.

The potential of AI to solve complex optimization problems also plays a role in sustainable supply chain network design. Thanks to the availability of (almost) real-time data and powerful computing capabilities—in other words, thanks to smart logistics—networks with more efficient distribution routes are potentially achievable through AI. Current estimates attribute roughly one-quarter of global greenhouse gas emissions to transportation—thus, any improvement in transportation efficiency has a direct positive environmental impact. Once more, as enabler of smart distribution networks, AI shows environmental, in addition to operational, potential.

AI can also be used for sustainable supplier auditing and selection. Large, global firms typically struggle to achieve visibility across their entire supply chain. In addition, modern supply chains simply involve too many partners to allow for fruitful “manual” auditing efforts.⁸ Data analytics now enable access to massive amounts of data regarding multiple sustainability dimensions of a company’s suppliers, their suppliers’ suppliers, and so on. Massive amounts of data, however, are challenging to process efficiently. Thus, leading suppliers in the environmental quality management space, such as EcoVadis and Intalex, have recently started to offer AI solutions as a way to generate actionable insights, such as supplier selection and development based on sustainable development standards, from these massive data sets.

All this said, the application of AI for sustainable supply chain goals is still in its infancy. One of the main challenges for the adoption of sustainable AI models within the industry is their dependence on being complementary to operational gains. In other words, it is difficult to imagine (barring government intervention) firms adopting sustainable solutions which do not simultaneously improve their operational efficiency. Thus, sustainability runs the risk of being reduced to a positive side-effect of AI implementations, instead of a primary objective in itself.

This fact, coupled with the speed of adoption of AI on the consumer-facing side (e.g., in marketing and revenue management) has prompted some warnings (Dauvergne, 2020). One fear is that, even if the adoption of AI within supply chain management results in reduced environmental impact at the unit level, the adoption of AI in other areas (e.g., marketing) may, at the same time, stimulate consumption

⁸ Ford Motor Company, for example, has over 10 tiers of suppliers; the number of suppliers in their first tier alone exceeds 1 000 (Simchi-Levi et al., 2015).

to such an extent that the combined effect is negligible or even negative. This is not unlike Jevon's paradox⁹: sustainability through efficiency gains runs the risk of being undone by the sheer increase in consumption.

In addition, widespread adoption of AI and the subsequent requirements for computing equipment places significant stress on the demand for rare earth metals and other raw materials required to power the (significant) additional computing resources. In effect, the increased efficiency at the production and distribution sides of supply chains may be offset by an increase in environmental pressure at the resource extraction side. This relationship is particularly important given the reliance on developing and/or poorly regulated countries for extraction of the required resources.

In summary, AI as a tool shows promise to bring about efficiency in logistics and supply chain management, to the extent that these improved efficiencies translate into a decrease in environmental pressure for the firms adopting AI. However, environmental systems need to be analyzed from a global perspective. Thus, looking forward, it is important for firms to be aware of the environmental impact of their entire supply chain, and not only of their local logistics processes.

4 Towards autonomous supply chains?

The digitization of workflows, including their work instructions, may also enable the automation of certain tasks, allowing the work to be performed by a machine instead of a human. With new digital tools providing visibility into real-time supply chain data and sophisticated algorithms capable of processing these data to prescribe decision-making, some even argue that the supply chain function is rapidly growing obsolete (Lyall et al., 2018).

We believe that human planners will *not* become obsolete, although their job content may likely change in light of the technological evolution. Indeed, one should draw a distinction between automation and autonomy. Automation implies that the task is performed “without thinking” i.e., by a machine (robot) or software application (bot). Autonomy, however, means that the task is capable of operating “without external intervention” i.e., using its own control rules. In office environments, the automation of tasks is referred to as *Robotic Process Automation (RPA)*. RPA is a software application (or bot) that performs automated tasks. By interacting with applications similar to how a human would, software bots can open email attachments, complete electronic forms, record and re-key data, and perform various other tasks that mimic human action. RPA is particularly efficient in automating very specific, highly repetitive tasks that follow predefined rules. When RPA bots run autonomously (unattended), their workflows should be pre-programmed such that human involvement is not required in the processes that they perform. In order to

⁹ The observation that improvements in efficiency are often associated with increased resource utilization due to an increase in total demand. William Stanley Jevon originally observed this effect with regards to the consumption of fuel in his 1865 book, “The Coal Question”.

work independently, these bots follow a rules-based process to completion. Such a process is only possible for relatively simple tasks.

Complex logistics tasks arising from judgmental decisions which require context awareness (and perhaps even qualitative factors), however, are much less amenable to explicit if-then instructions used in automating repetitive, predefined tasks. For such more complex tasks, RPA bots currently run in supervised (attended) mode. In such a case, human employees work side by side with the bots, which may be likened to virtual assistants providing support to an individual employee with their tasks so as to boost productivity. Employees trigger a bot and interact with it as the bot provides assistance. Attended RPA bots therefore increase the productivity of the planner by performing time-consuming repetitive tasks, while leaving the contextual and judgemental decision-making to the human.

Smart control rules may gradually increase autonomy in the supply chain as data are increasingly captured through digitally connected smart systems (the aforementioned Industry 4.0). In 2019, an Uncrewed Surface Vessel (USV) journeyed from the UK to deliver oysters in Ostend, Belgium. The USV then returned to the UK with Belgian beer. In order to safely navigate what is one of the busiest shipping lanes in the world, the USV relied on input from multiple installed sensors, including Sonar, Radar, Lidar, Camera, IR-camera and GPS. According to the USV's owner, the trip was the first commercial crossing of the North Sea by an autonomous vessel. Throughout the trip, an operator was able to remotely access video footage, thermal imaging and radar transmitted by the USV, along with audio transmissions of the USV's surroundings. The operator was even able to communicate with other vessels in the USV's vicinity (Amos, 2019). Although the presence of multiple data sources is instrumental to increase the level of autonomy afforded to an AI system, the end result is still 'supervised' automation. In order to realize full autonomous automation, algorithms will increase in complexity and sophistication (in addition to standard if-then-else rules), similar to Tesla's pursuits of incorporating state-of-the-art machine learning technology into their vehicle operating systems.

5 The Human Aspect (Going Forward)

The previous sections have demonstrated how AI is shaping a future where well-defined tasks are automated and how algorithms can solve problems that may be too complex for humans to analyze. These uses of AI raise the question *what is the role that we—as humans—are going to play in all this?* Looking back, technological revolutions have changed our ways of life in profound ways. We settled down during the agricultural revolution, we moved in large numbers to bigger cities during the industrial revolution and we unlocked the key for real-time worldwide communication during the digital revolution. Looking forward, will Industry 4.0 bring about a change of the same order of magnitude? If AI and real-time connectivity are indeed driving us towards a fourth industrial revolution, the consequences could be far-

fetching and difficult to predict.¹⁰ One could also argue, however, that AI is integral to an ongoing evolution in the way we conduct business, such that resulting changes will be gradual and organizations will—to a certain extent—be able to anticipate them. There is a fundamental difference in the human-machine interaction within the new AI paradigm, compared to the “old” decision support systems of the past few decades. Thus far, computers have been solving models explicitly developed by researchers for a particular problem at hand. Most of these models are theory-based, relying on assumptions about causality and on abstractions about how the physical world behaves. For example, a decision support system that calculates when and how to order from suppliers requires the modeling of a demand structure and calculating the optimal trade-off in terms of the cost of having too much inventory on hand as opposed to too little. In addition to their practical implementations, the underlying theoretical models also allow us to extract insights, and thus learn about the intuition behind the system being modeled. Extracting such intuition is the reason why these models are taught in supply chain education programs; they allow students to build-up their own intuition, learn how to interpret results and, eventually, how to derive their own solution methods for specific problems at hand. Such a skill set is valuable even when the problems faced are much more complex than the stylized versions seen in class.

AI in general, and machine learning in particular, do not require such models. Machine learning models are atheoretical and thus typically provide a prediction without requiring prior knowledge (or explicit model) regarding the problem nor additional insights or information from the user. Machine learning models are seen as a black box, because it is usually not possible to understand how they arrive at a solution/prediction. Thus, while humans are typically in full control at all stages in “legacy” decision support systems, our role within future AI decision support systems will be limited to ensuring availability of any data required by the computer to run its model and, eventually, to fine tune the necessary reward functions and parameters.

AI is likely to bring about a reconfiguration of the core competencies of firms and, consequently, also a reconfiguration of jobs (Agrawal et al., 2018). In this view, a rethinking of the core-competencies of jobs and firms is in order. Should machines take over prediction jobs, such as demand forecasting, the boundaries of a supply chain analyst’s job will change: such a professional will go from creating and maintaining forecasts to possibly just approving and archiving them. Taking this thinking to a company level, once processes are automated and do not require in-house expertise, they can potentially be outsourced. In effect, Amazon is already offering ready-to-go solutions for processes that are still commonly considered core competences of supply chain management, such as forecasting and predictive maintenance. It is not difficult to imagine a future where sales forecasts are computed by an outsourced specialized firm making use of the latest machine learning models run on state-of-the-art server farms and are received every morning.

¹⁰ Who could have predicted the loss of importance of good handwriting in an educated person 100 years ago? Or the loss of the importance of a computer user understanding the concept of manually saving a local copy of a file even 10 years ago?

When day-to-day requirements change, the composition of the workforce and their skill requirements will also change. With routine prediction and data processing tasks being taken over by machines, organizations will most likely demand workers specialized in analyzing computer output as well as workers with ‘soft’ interpersonal skills; necessary for sales, consulting, and coaching functions (Ernst et al., 2019).

For workers, new tasks will emerge that may require a different combination of skills. As AI serves up improved and cheaper predictions, there is a need to clearly think about and determine how best to develop and use those predictions. Agrawal et al. (2018) refer to the job of determining how algorithmic predictions are evaluated as “reward function engineering.” A reward function engineer will determine the way in which the predictions are evaluated and consequently adapted or adopted. Successfully performing reward function engineering requires an understanding of the needs of the organization and the capabilities of the algorithm. This understanding requirement also extends to management. Managers will require basic AI-literacy—knowing what an AI system can and cannot do, and knowing what their strong and weak points are. No coding necessary. Indeed, in the same way in which no coding knowledge is required to interact with a spreadsheet, no coding or knowledge of the underlying technology will be needed to interact with and supervise AI models. With widespread adoption, and thanks to the above mentioned availability of open source resources, coding/development of machine learning frameworks may be completely dissociated from its application.

Even if a firm outsources certain supply chain processes to machines, an analyst will still be required to interpret AI output and to sign-off on decisions. For all the advantages AI has over human decision making in terms of data processing, humans still trounce AI in the nuances of understanding context and qualitative information. A human decision-maker can navigate new situations through context clues; AI is typically as good as its training data. Things that are still easy for a human to spot, are not possible for an AI system if not explicitly trained. For example, in 2018 Reuters reported that Amazon had to abandon an AI recruiting tool—capable of grading job candidates based on their resumes—because it discriminated against women. The AI model had been trained with resumes from previous applicants, which skewed heavily male. The machine learned the implicit biases in the recruitment practices and female traits were seen as negative by the system.

It follows that, in outsourcing/automating predictions and recommendations to AI algorithms, the role of humans will necessarily move towards that of being gatekeepers. Even though it is still early days in the adoption of AI in the corporate world, Acemoglu et al. (2020) recently performed an empirical study using vacancy data from 2010–2018. They found that (1) AI-related jobs increased steeply after 2014, (2) there was evidence of job replacement with AI-jobs replacing other, now-automated, jobs, and (3) there is no effect (yet) of AI on the aggregate level of employment.

In light of all this evolution, one thing has become clear: data and data management will become increasingly important. Whereas, historically, data has been siloed within a company, where each decision-maker compiles and stores their data independently, we foresee—especially in applications where data is a competitive

advantage—the adoption of a “data manager,” or chief data officer (CDO) with an overview over all the data retrieval and storage processes in a company.

Finally, we note that, despite its many productivity improvements, the widespread adoption of AI in the workplace does not come without its risks. Ernst et al. (2019) performed a detailed analysis of the implications of AI on the future of work. Even though their conclusions are optimistic, they note that certain challenges will require smart policy-making and conscientious adoption of the technology. From a supply chain perspective, they warn about an increase in inequality due to the automation of operational tasks. Of particular concern is the impact of AI in developing countries. Increased use of capital-intensive technologies can increase inequality in countries with large percentages of under-utilized, under-educated labor. In contrast to the agricultural revolution, in which automation in agriculture led workers to low-skilled jobs in manufacturing and eventually to increased income and education, during the current industrial (digital) evolution manual jobs are being increasingly replaced by higher-skilled jobs in service, design, and supporting roles. The move to AI adoption and automation contributes to these trends. The main fear is inequality in job creation, where mid-level jobs are replaced by low-end and high-end ones. To the extent that highly educated people are better at learning new skills, and the skills to succeed with AI will change over time, then the educated managers will benefit disproportionately.

Data, data, data¹¹

*Hans Thibau, Vice President Logistics
Compressor Technique, Atlas Copco Airpower*

Atlas Copco is a Swedish multinational industrial company that manufactures industrial tools and equipment. In 2019, global revenues totaled SEK 104 billion (approx. €10 billion), and by the end of that year the company employed nearly 40,000 people. Atlas Copco companies develop, manufacture, service, and rent industrial tools, air compressors (of which it is the world's leading producer), construction and assembly systems. Atlas Copco is global market leader in most of its segments. It manufactures critical equipment components in-house and co-operates with business partners for non-critical components. Atlas Copco co-engineers and buys a lot of parts and only produces core components in house. A main contribution of the revenues is generated from service (spare and consumables, maintenance and repairs, uptime contracts, air as a service, accessories, rental). Service is the responsibility of divisions in each business area. Atlas Copco is worldwide represented and has own customer centers in about 71 countries.

As an engineering company, Atlas Copco has always embraced innovation and new technologies to continuously improve its products and processes. The company is also leading when it comes to sustainability goals and limited impact on environment. To date, more than 170K of our installed compressors worldwide are connected through Internet of Things technologies, sending information about their operating condition. This can vary between every minute and weekly information. This massive amount of data is processed through cloud computing to guarantee up-time, machine efficiency and lowest total cost of ownership for the customer. Well-timed preventive maintenance interventions avoid costly machine downtime, and prevent that the machine's components are replaced too early. At the same time the lowest energy consumption for the customer can be guaranteed.

Accurate prediction of maintenance interventions is also instrumental to have the right spare parts in stock at the right place at the moment when they are needed. The use of *big* data in logistics goes a step further than the IoT application for preventive maintenance. As a machine state is directly linked to sensors, an accuracy of 98% on AI algorithms for machine failures brings substantial benefits on boosting the up-time further beyond the normal conditions. In logistics, however, 98% of the demanded stock keeping units are 'relatively' easy to forecast; the complexity is in predicting the remaining 2% of the (very) slow moving products. Unfortunately, those products are only rarely used, and thus also only few data points are generated. It thus turns out that for those stock keeping units where accurate data would add most value, only *small* data are available.

Digital solutions also support Atlas Copco's transport operations. Containers shipped by boat or truck are equipped with sensors to track and trace incoming

¹¹ "Data! Data! Data! I can't make bricks without clay." Sherlock Holmes, *The Adventure of the Copper Beeches*.

products throughout their journey using digital control towers. The transparency in knowing when the container has left its origin, when it encounters a delay, and when it is planned to arrive on site, provides multiple insights. It can for instance lead to emergency air shipments if the component is highly urgent or it can suggest re-planning the production line based on the expected arrival time of the missing components. When the part is intended for customers or for own service engineers, they can be timely informed about the status of arrival. Finally, by analyzing historical trajectory data, one can learn that—or when—some trajectories take longer than planned, such that the corresponding replenishment parameters may be updated. Although the transport control tower is today still its infancy, it is a growing opportunity area.

There is also opportunity to optimize picking locations in the warehouse by analyzing historical orders. By grouping products that are frequently purchased together in each other's vicinity, picking productivity can be enhanced. The same analysis can be instrumental to optimize service levels per order line, depending on the combination of products that are often included in the same order.

All these opportunities, however, stand or fall by the accuracy of the data. Bad data results in poor quality decision-making, which makes any attempt to adopt AI moot (*garbage in, garbage out*). The quality of the master data is important in order for it be useful. This makes a strong case of integrated data systems.

Data and analytics have always been important at Atlas Copco. To leverage the true benefits of AI, data scientists are connected to the business. Logistics and supply chain executives do not necessarily need to be computer scientists, but they do require an analytical mindset and a data science acumen to understand how AI can bring value to their profession. And most importantly, they need to be critical on the data, and—together with the data science team—co-own responsibility on data quality.

References

- Acemoglu D, Autor D, Hazell J, Restrepo P (2020) AI and jobs: Evidence from online vacancies. Tech. rep., National Bureau of Economic Research
- Agrawal A, Gans J, Goldfarb A (2018) Prediction machines: the simple economics of artificial intelligence. Harvard Business Press
- Amos J (2019) Autonomous boat makes oyster run. <https://www.bbc.com/news/science-environment-48216966>
- Boute R, Van Mieghem J (2021) Digital operations: Autonomous automation and the smart execution of work. *Management and Business Review* 1(1)
- Cui R, Gallino S, Moreno A, Zhang DJ (2018) The operational value of social media information. *Production and Operations Management* 27(10):1749–1769
- Dauvergne P (2020) Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy* pp 1–23
- De Moor BJ, Gijsbrechts J, Boute RN (2021) Reward shaping to improve the performance of deep reinforcement learning in inventory management, Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3804655>
- Ernst E, Merola R, Samaan D (2019) Economics of artificial intelligence: Implications for the future of work. *IZA Journal of Labor Policy* 9(1):1–35
- Flows (2019) Gentse ML6 voorspelt wachttijden douane met artificiële intelligentie. <https://www.flows.be/nl/logistics/gentse-ml6-voorspelt-wachttijden-douane-met-artificiele-intelligentie>
- Gijsbrechts J, Boute RN, Van Mieghem JA, Zhang DJ (2020) Can Deep Reinforcement Learning Improve Inventory Management? Performance and Implementation of Dual Sourcing-Mode Problems, Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3302881>
- Imdahl C, Hoberg K, Schmidt W (2021) Targeted automation of order decisions using machine learning, Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3822131>
- Kolner T (2019) Applying machine learning on the data of a control tower in a retail distribution landscape. Master's thesis, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science
- Lyll A, Mercier P, Gstettner S (2018) The death of supply chain management. <https://hbr.org/2018/06/the-death-of-supply-chain-management>
- Si XS, Wang W, Hu CH, Zhou DH (2011) Remaining useful life estimation – a review on the statistical data driven approaches. *European Journal of Operational Research* 213(1):1–14
- Simchi-Levi D, Schmidt W, Wei Y, Zhang PY, Combs K, Ge Y, Gusikhin O, Sanders M, Zhang D (2015) Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces* 45(5):375–390
- Sun J, Zhang D, Hu H, Van Mieghem J (2021) Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations. *Management Science* To appear
- The Economist (2017) The world's most valuable resource is no longer oil, but data. May 6

Vanvuchelen N, Gijsbrechts J, Boute R (2020) Use of proximal policy optimization for the joint replenishment problem. *Computers in Industry* 119:103239