How AI will Revolutionise Construction Site Management

Novade

Table of Contents

Executive Summary

The construction industry stands out as a heavyweight within any nation's economic and societal fabric. However, it trails most other sectors in the maturity of its structure and practices, the quality of its workforce development and the magnitude of its investment in innovation and technology. This has led to lacklustre performance at best, reflected in the widespread inability to deliver projects on time and within budget. Testament to the sector's inefficiency on a global scale, construction productivity has grown by a mere 20% over the past twenty years while manufacturing productivity has nearly doubled over the same period.

Construction, albeit a very late adopter, is finally embracing digitalisation. Digital collaboration platforms are slowly finding their way to the market. Workflow digitalisation, real-time data capture, continuous site status monitoring and construction data analytics are gradually disseminating through the industry. Against this backdrop, another technology is already drawing considerable attention, destined to impact this sector as it has started to radically transform others: Artificial Intelligence.

Beyond the hype, Artificial Intelligence (AI) is indeed one of the most powerful and impactful general-purpose technologies of all time. Four of its techniques are particularly relevant to the construction sector.

- AI identifies patterns and correlations in data that humans are oblivious to. Given enough input variables, it turns them into predictions and prescriptions.
- AI can also explore millions or even billions of scenarios and come up with optimised solutions devoid of human preconceived ideas.
- AI far surpasses humans in sensory perception and particularly vision. It instantly detects visual characteristics sometimes imperceptible to humans.
- AI is improving by the day at Natural Language Processing. For instance, it is capable of extracting intent or broad meaning from human speech or writing.

AI will be pervasive in construction in the medium and long-term. Three applications are worth highlighting as they will improve the most relevant performance indicators in construction.

- AI will increase productivity by a series of scheduling-related actions forming a virtuous cycle: better prediction of task lead time, real-time feedback about variance from the master schedule, enhanced site coordination, master schedule stabilisation, master schedule optimisation, and reduced labour idle time.
- AI will lower the cost of total quality: the progressive automation of fault detection, particularly with computer vision, will increase the effectiveness and the efficiency of the process. Then, as the drivers of non-quality are identified, quality issues will become predictable and preventable.
- AI will provide a much-awaited boost to construction site safety: the collection of large amounts of data in case of incidents or lack thereof will enable prediction of the occurrence and the severity of safety accidents. This in turn will lead to comprehensive sets of preventative measures and eventually to the reduction of incidents. AI-powered guardian angels could appear in the long-term future. Lives will be saved.

The impact of AI will not be felt overnight. The journey will give rise to a new set of challenges, epitomised by the highly sensitive ethical challenges which need to be addressed with utmost care and expertise. Most importantly, the bulk of AI applications will require a large quantity of data, which needs to be collected as early as possible. For any contractor wanting to reap the immense benefits of Artificial Intelligence, this entails starting the digitisation of their construction activities immediately.

A Juggernaut in Search of Solutions

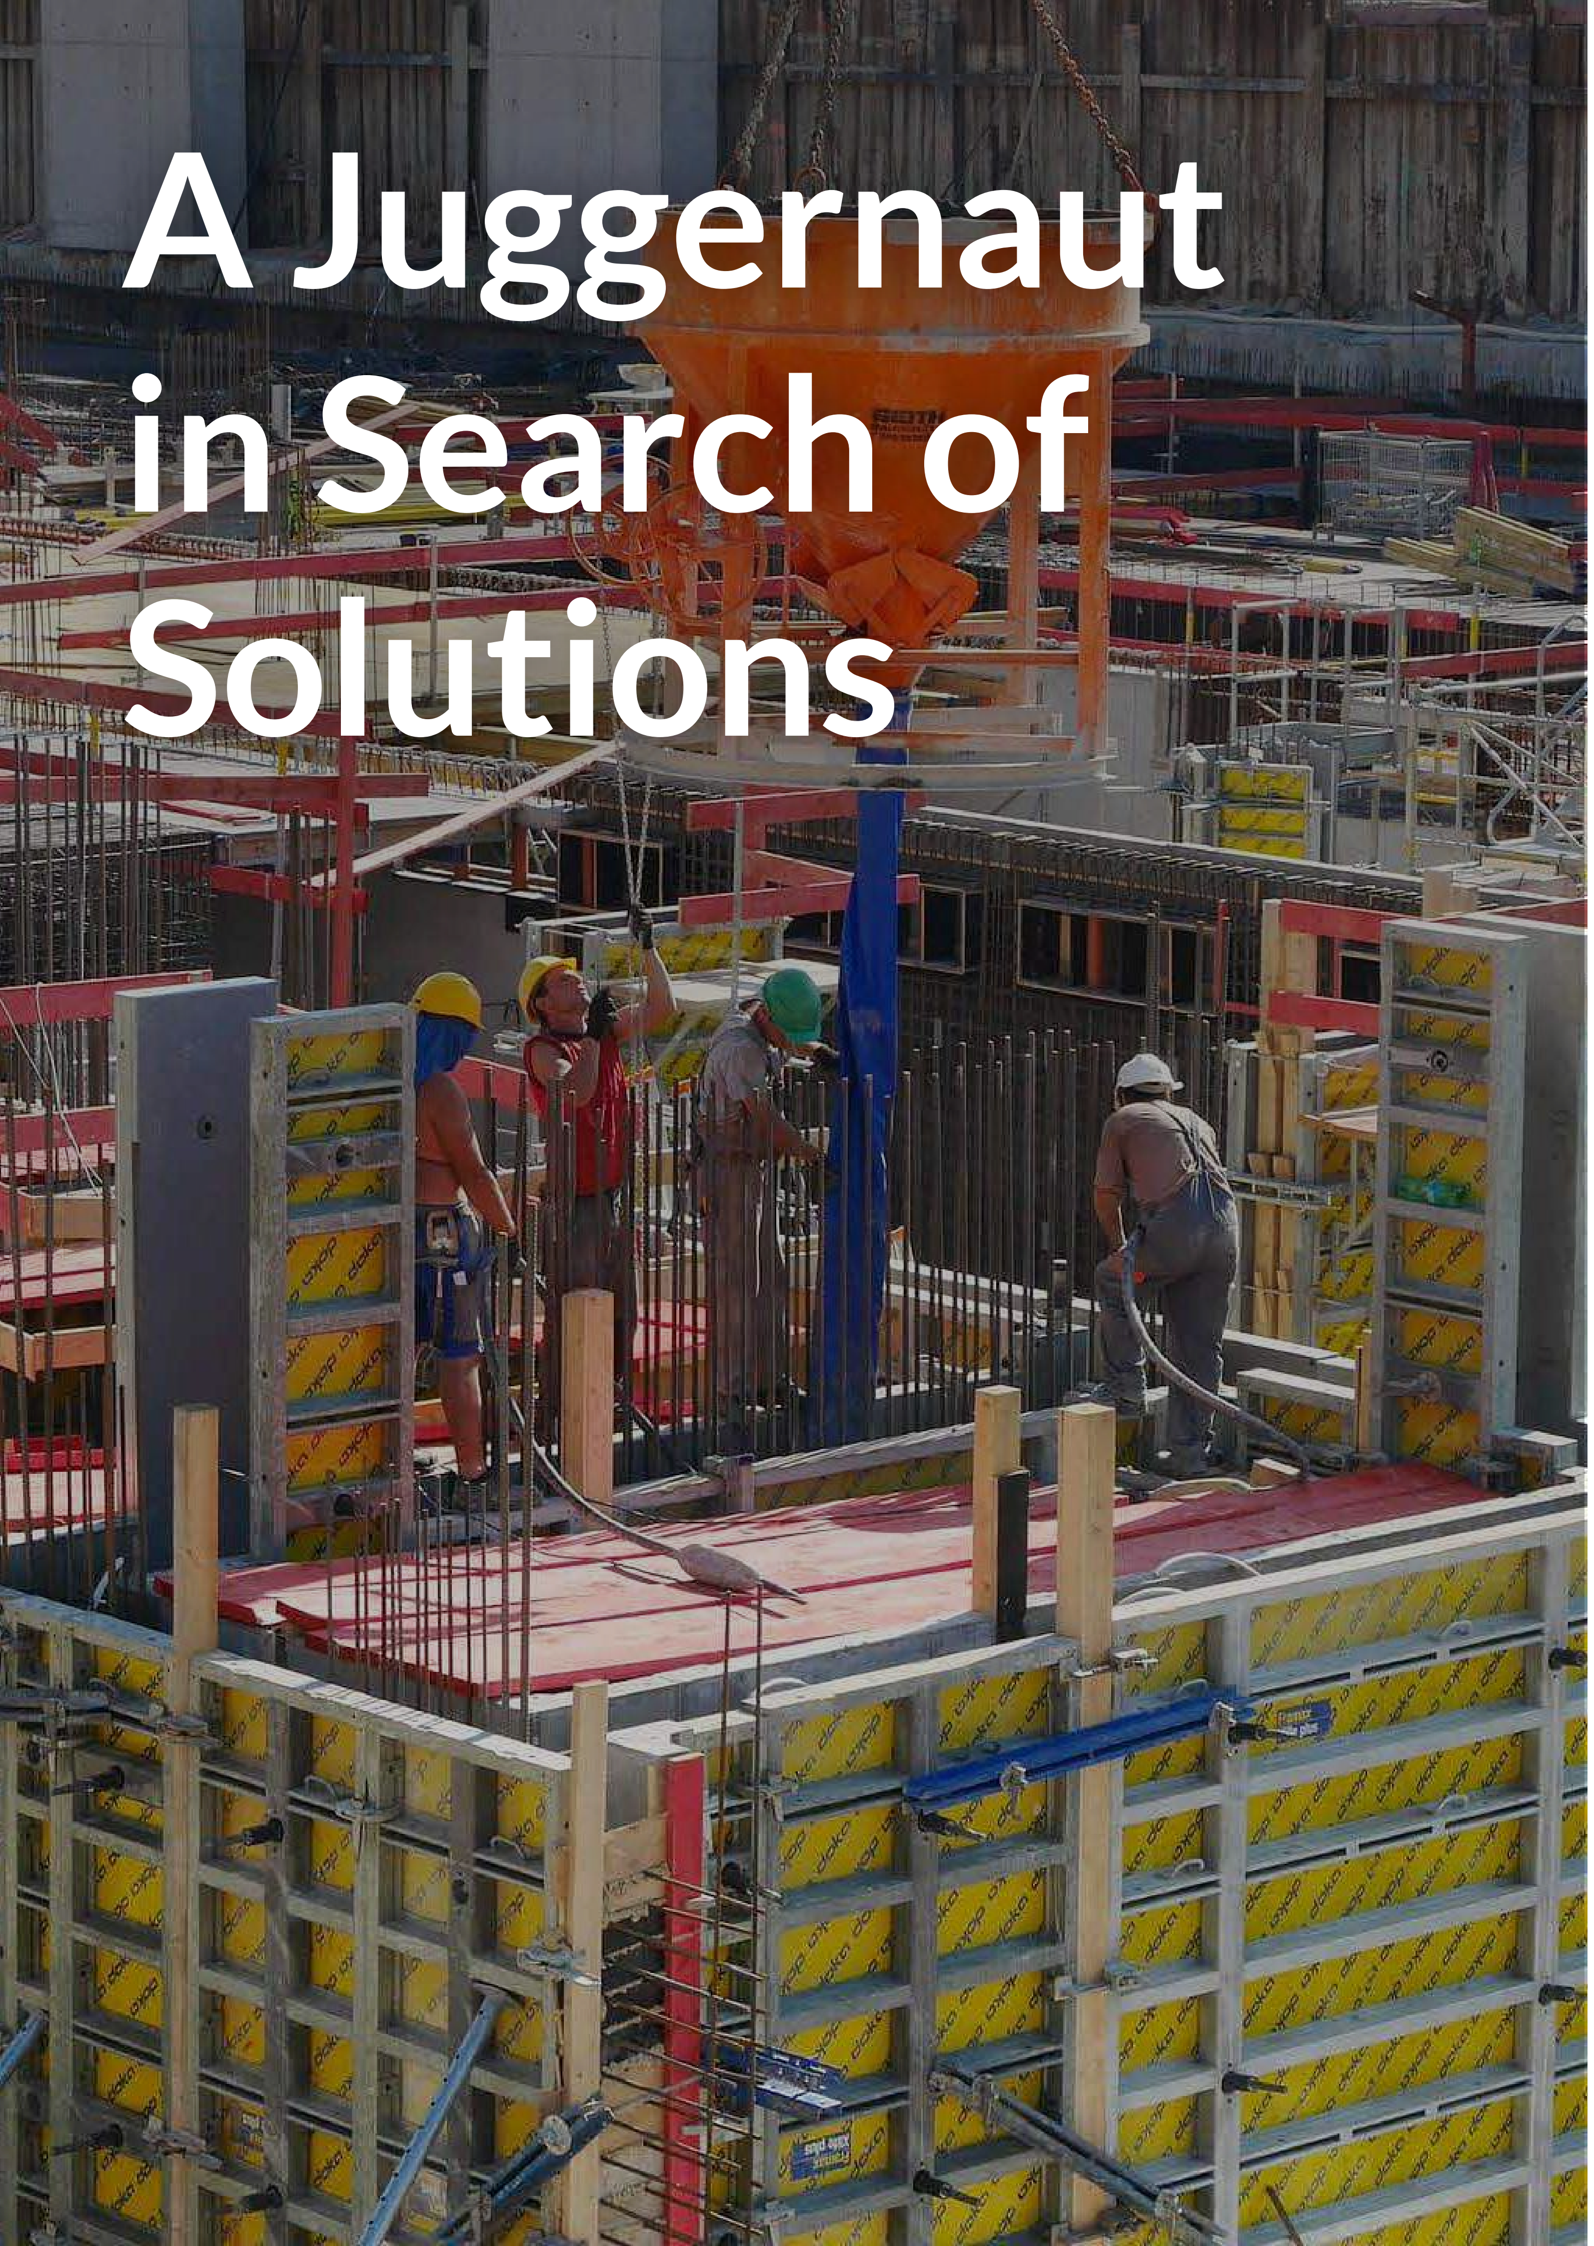

A Sector that Matters

Few economic sectors can boast as important a role in society and nation building as construction. Invariably, construction shapes and produces the environment for all human activities to take place: from residential dwellings for living, to commercial, industrial and tertiary buildings for work, to recreational centres for play. Day after day, construction erects the required infrastructure for children to grown and learn; for businesses to invent, develop and trade; for people to work, eat, heal and travel; for cities to run and become ever smarter – in a word, for all to thrive.

The sector carries a staggering economic weight. Construction-related spend accounts for approximately 13% of global GDP and is forecast to grow at more than 3% annually for the next few years[1]. The sector employs 7% of the world's working-age population, making it one of the biggest job providers.

Finally, the rapid aging of the world population, particularly in developed economies, and the addition of 2 billion people by 2050, mostly in newer economies, entail an essential role for construction regarding two Sustainable Development Goals (SDGs). SDG number 9 calls for resilient infrastructure and SDG number 11 for sustainable cities; construction will be instrumental in achieving both.

Pervasive Inefficiencies

Yet, the construction sector is mired with inefficiencies. Alongside promising pockets of improvement, the landscape remains by and large one of fragmentation and opacity, archaic project management practices and people development, and underinvestment in innovation and technology.

As a result, all stakeholder categories stand to lose. Construction as a sector is in the lowest quadrant of profitability, penalising Engineering and Construction companies. Wages remain low[2] and conditions harsh, making the trade less appealing to the workforce. Costs of construction have generally outpaced the growth in consumer price index and demand for construction may not be met, causing significant dissatisfaction amongst owners. Overruns are commonplace, at a global average of 70% over original budget and 61% over original schedule[3]. Variability in these indicators is high while predictability is dismal.

Such characteristics, which other sectors or individuals might deem utterly unacceptable, have gradually become accepted as a fact of life, as if they were inevitable. Compounding the difficulties, resistance to change and sometimes misaligned interests have strengthened the status quo. But it doesn't have to be this way.

Digitisation, at Last

Although one of the least digitised sectors in the world, the sector has finally seen the emergence of numerous and innovative construction technology firms. These new entrants raised US$ 10 billion of funding between 2011 and early 2017[4]. They collectively operate across the whole construction lifecycle, from design and preconstruction, to construction itself, then to operations and management. This paper will keep its focus on construction itself and more specifically on onsite execution.

> "Digitalisation is turning from a business enabler into an invaluable source of actionable insights."

Take, for instance, the digitisation of workflows. It started as a simple business enabler. Digital workflows are required to be documented, whereas they were not necessarily formalised before. Every activity performed by any worker is recorded. This provides real-time visibility to project management and facilitates contract compliance monitoring. For work authorisation workflows, no more paper needs to travel from person-in-charge to next-level person-in-charge amidst the buzzing construction environment; lead time is drastically reduced from days or hours to minutes, worker idle time (or work without authorisation) plummets, the information is safely recorded and archived, and literally tonnes of paper are spared. For quality workflows, all defects lodged are now certain to be monitored and solved, and probably much quicker than they would have been before.

But digitalisation is turning from a business enabler into an invaluable source of actionable insights, answering the most practical questions:

- Which workflows need to be transformed because they encompass the most idle time, and how must scheduling be altered? Real-time crew deployment tracking and subsequent analytics thus improve field productivity.
- Which quality issues are the most detrimental on a consolidated basis, because of their occurrence rate or their impact? What should be done to contain the issue, mitigate it and eradicate it once and for all? Rigorous monitoring and multi-dimensional analyses uncover and solve unsuspected teething issues.
- Which subcontractors cause the highest number of non-conformities or safety incidents, and how can we improve their awareness and train them – or should more radical action be taken? Supplier analysis across quality and safety data enables appropriate corrective action to be taken.

For all the benefits digitalisation has already yielded, this is only the beginning. A very special technology is poised to take the construction sector by storm: Artificial Intelligence.

Did You Say "Artificial Intelligence?"

Beyond the Hype

Are you confused by the conflicting information you are bombarded with regarding Artificial Intelligence, from unsubstantiated hype to equally excessive downplaying? You are not the only one.

In essence, Artificial Intelligence is the ability of machines to mimic, and increasingly often to surpass, the outcomes of human intelligence. A more sophisticated definition would refer to machines that can sense their environment, think, learn, and act, in response to what they sense and their programmed objectives[5]. AI's most widespread subset today is Machine Learning, a remarkable algorithmic class, in which machines learn to effectively perform a specific task without using explicit instructions. Their performance increases with experience; in other words, the more they are used, the better they work. Finally, within Machine Learning, Deep Learning is an algorithm learning mode using Artificial Neural Networks, loosely inspired by the functioning of the biological brain. Deep Learning accounts for many of the extraordinary achievements of AI in recent years. More specifically, much of the recent exponential performance increase stems from the combination of algorithmic progress, increase in computing power and massive availability of data.

Although still in its infancy and facing many challenges, AI has started exhibiting formidable results. It is highly unlikely that it will hit another of the "winters" that has mired its development in the past half-century. A detailed comparative analysis suggests that AI is nearing or surpassing human performance in an increasing number of human mental processes – the building blocks of all decisions we make and actions we take[6].

The following broad elements summarise the big picture of Artificial Intelligence:

- **AI is arguably one of the most powerful and impactful general-purpose technologies of all time.** AI is destined to permeate all layers of society through all kinds of physical objects or intangible services. Like electricity, Artificial Intelligence is everywhere.

Directionally, AI is constantly progressing. According to the most competent experts[7], we are more likely than not to have invented, by the end of the century, Artificial General Intelligence – an intelligence as versatile and efficient as that of humans, which will profoundly change the fate of our species.

- **AI performs but does not understand.** AI is purely a set of cold mathematical algorithms. Despite its breath-taking performance often surpassing that of humans, AI has no understanding of the decisions it recommends, the signals or images it recognises, the sentences it translates or the answers it gives back.
- **AI is highly beneficial, but also presents dangers.** At work, AI frees us from an increasing number of mundane, painful, or menial tasks while giving us access to insights and predictions that are beyond human computing capabilities. It also makes our lives more convenient and contributes to solving humankind's grand challenges.

But like all technologies, Artificial Intelligence is double-edged. Leaving aside the potential malevolent use by humans, AI at this nascent stage comes with its fair share of ethical problems such as design dilemmas or induced biases. In the medium run, the social impact of AI on employment will inevitably be felt.

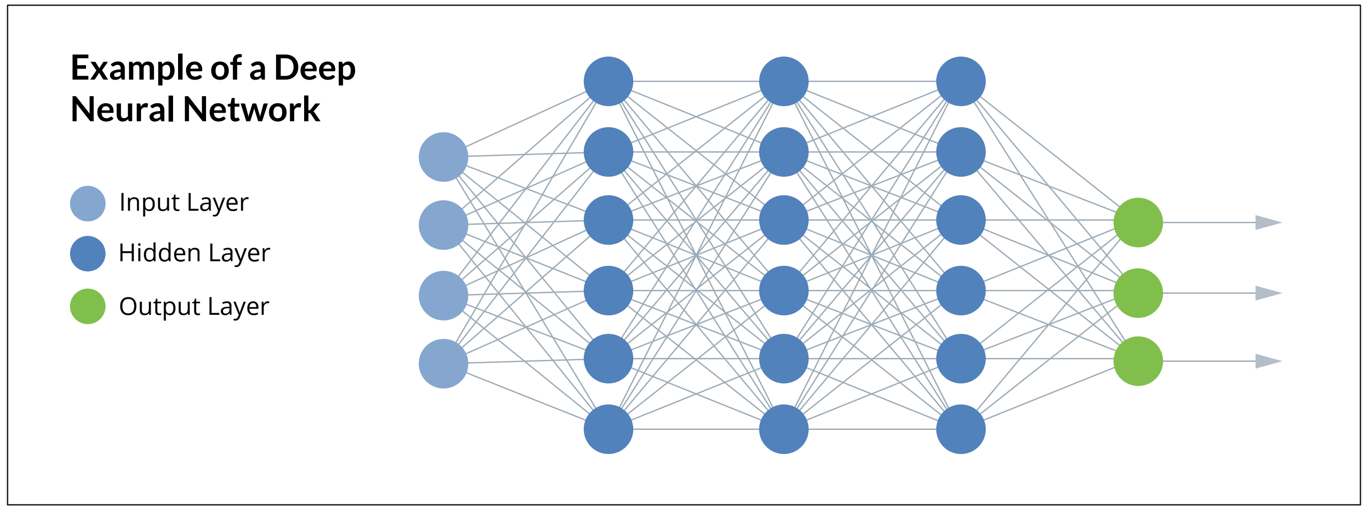

Four Relevant Techniques

Amongst a long list of Artificial Intelligence techniques, four at least will be found of particular relevance to the construction sector in the short and medium term.

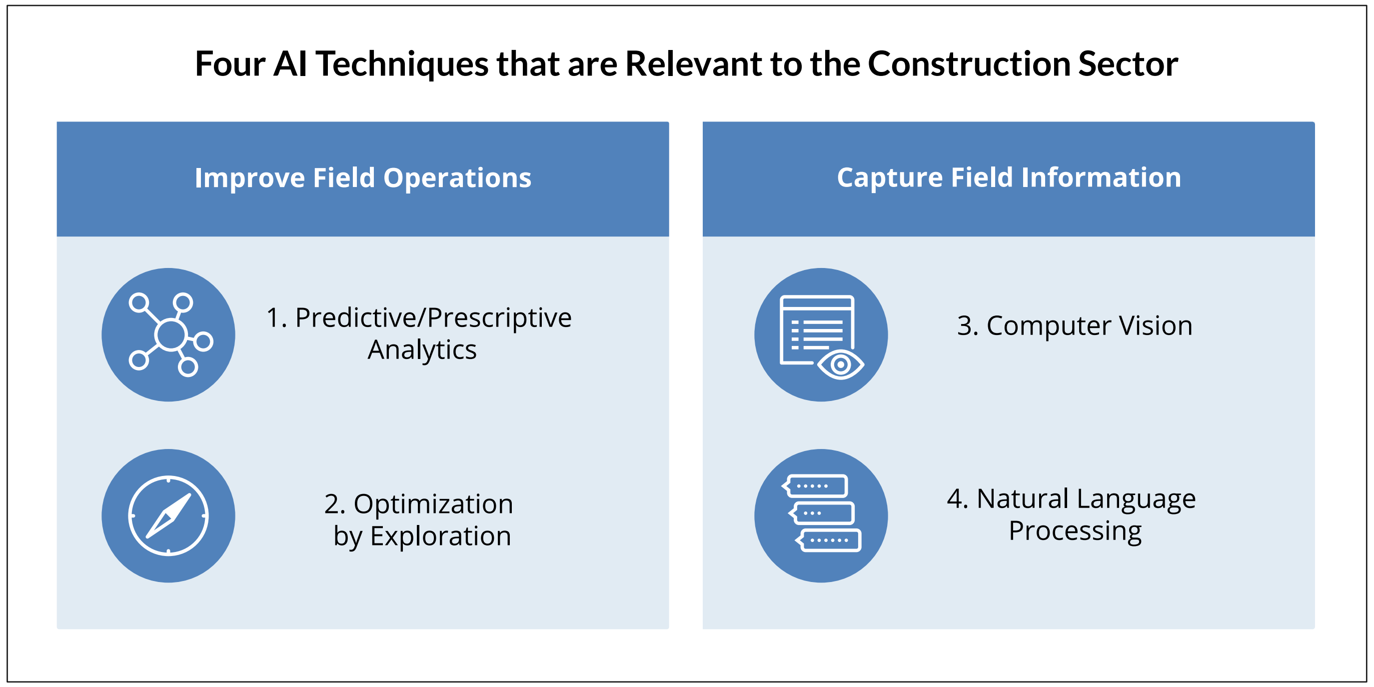

Predictive / Prescriptive Analytics

AI is a prediction machine and uses its predictions for estimation and optimisation. It does so through an unmatched ability to spot correlations in data, to find patterns amongst millions of parameters that humans are simply unable to process. The idea of precision medicine, for instance, is to gather all the information on a patient: the patient's genetic heritage, family history, physiological data, real-time immune system status, medication intake, and medical test results, but also the patient's environment, behaviour and social relationships, conversations and state of mind, and finally all the medical literature that is constantly being updated. AI draws on the insights it has gained from a huge database to make predictions on the evolution of the subject's health.

By the same token, AI can forecast the price of real estate assets or guess whether an employee is likely to resign or not. It can spot the subtle signals of upcoming defects in a complex industrial system long before humans. AI can determine and recommend an optimal route for a car, knowing the starting point, destination and traffic conditions. Note the important shift in nature here: from being predictive (describing what the situation will be), AI becomes prescriptive (advising on the steps to take for a positive outcome). It goes on to suggest the best time to irrigate your crops and predict the date your flowers will bloom. It can improve predictions of earthquakes or volcanic eruptions. It can identify new craters on the moon and new exoplanets in the galaxy.

This is relevant to construction site management for predicting numerous events such as quality issues, safety incidents, issue resolution lead time, task completion date, final building cost, and many others.

Optimisation by Exploration

AI excels at optimising any variable by exploring the realm of possibilities, using a technique called Reinforcement Learning. This technique of exploration does not require prior knowledge or human intervention: the machine is essentially rewarded for exploring by itself the strategies that look most promising and dropping the others. It is thus devoid of human biases and prone to exploring radically new paths.

This technique gained immense popularity in 2017 when DeepMind's AlphaGo Zero defeated the world's best player at Go, an extremely complex Chinese board game[8]. So great was the feat that when historians start looking for an event marking the beginning of the Age of Artificial Intelligence, this victory will be a very serious contender. In the aftermath, the technique was used again to solve such complex problems as protein folding – determining the minimum energy and hence the spatial configuration of proteins from amongst billions of possibilities[9], thereby predicting their crucially important biological properties.

This is relevant to construction site management for optimising, say, scheduling based on earliest completion, lowest cost or any other criteria.

> "This technique of exploration does not require prior knowledge or human intervention: the machine is essentially rewarded for exploring by itself the strategies that look most promising and dropping the others."

Computer Vision

AI surpasses humans at recognising objects. Using a technique called Supervised Learning, it is fed with large quantities of labelled images, which it learns to recognise by progressively modifying the parameters of its algorithm although not told explicitly how to do so.

AI's quest to "see" started in the late 1980s when Optical Character Recognition enabled machines to recognise postal codes. Subsequent to algorithmic improvement in 2012, AI went from almost blind to superhuman in the field of computer vision. In the blink of an eye, AI can now filter out tomatoes that are too green or non-compliant, assess the composition of a forest into different tree species, read lips[10] and accurately diagnose skin cancer[11].

This is relevant to construction site management for any process currently carried out by human vision – there are plenty, such as quality control or risk prevention.

Natural Language Processing

AI is making spectacular progress at Natural Language Processing – an umbrella term referring to all tasks involving language as it is naturally spoken or written by humans. Such are the rate and magnitude of improvements that the very benchmarks used to assess them need to be changed on a nearly yearly basis.

AI beats humans at recognising speech ("speech to text") as well as voicing speech in any language or accent ("text to speech"). It nears human performance at translating – but of course surpasses any individual human at all languages. It is becoming increasingly good at extracting meaning from a text ("semantic understanding") or at generating texts that make sense ("speech generation"). AI enables a live conversation between two people not speaking each other's language. AI can process large amounts of documents, extract the information sought almost instantly, and generate a synthetic report. Remember: AI performs but does not understand as humans do.

Natural Language Processing is relevant to construction site management as voice is a natural, yet currently underused interface between human and machine. NLP is also widely used to make sense of all the previously handwritten reports that now need to be processed on an industrial scale.

AI in Construction

AI is bound to play a highly significant role in construction, just as it will in every other sector. Its role will be felt in capturing data and rendering analyses, optimising, predicting, and prescribing. This is set to help CEOs and owners of building firms solve many of the issues at hand. Technical capabilities are maturing very fast. From China to the US via Japan, Korea, Taiwan, Singapore, France and Greece, to name a few, academia has started to research and publish papers developing and evaluating AI algorithms designed to improve performance in construction.

While all construction technology firms have AI aspirations, few have already designed actual roadmaps and even less have developed any meaningful solution yet – but the most advanced are gearing towards it. At the end of the day, the main unknown for the development of the field remains the adoption rate of AI by the sector as a whole.

Within construction, the focus of this paper remains construction site management. We deliberately choose to ignore upstream processes such as design or downstream processes such as facility management, although they will be substantially impacted by AI. We will not dwell either on peripheral processes such as logistics and warehousing, or back office operations such as finance and accounting or HR management, as their disruption by AI is ongoing and already well documented. We will thus concentrate on Productivity, Quality and Safety.

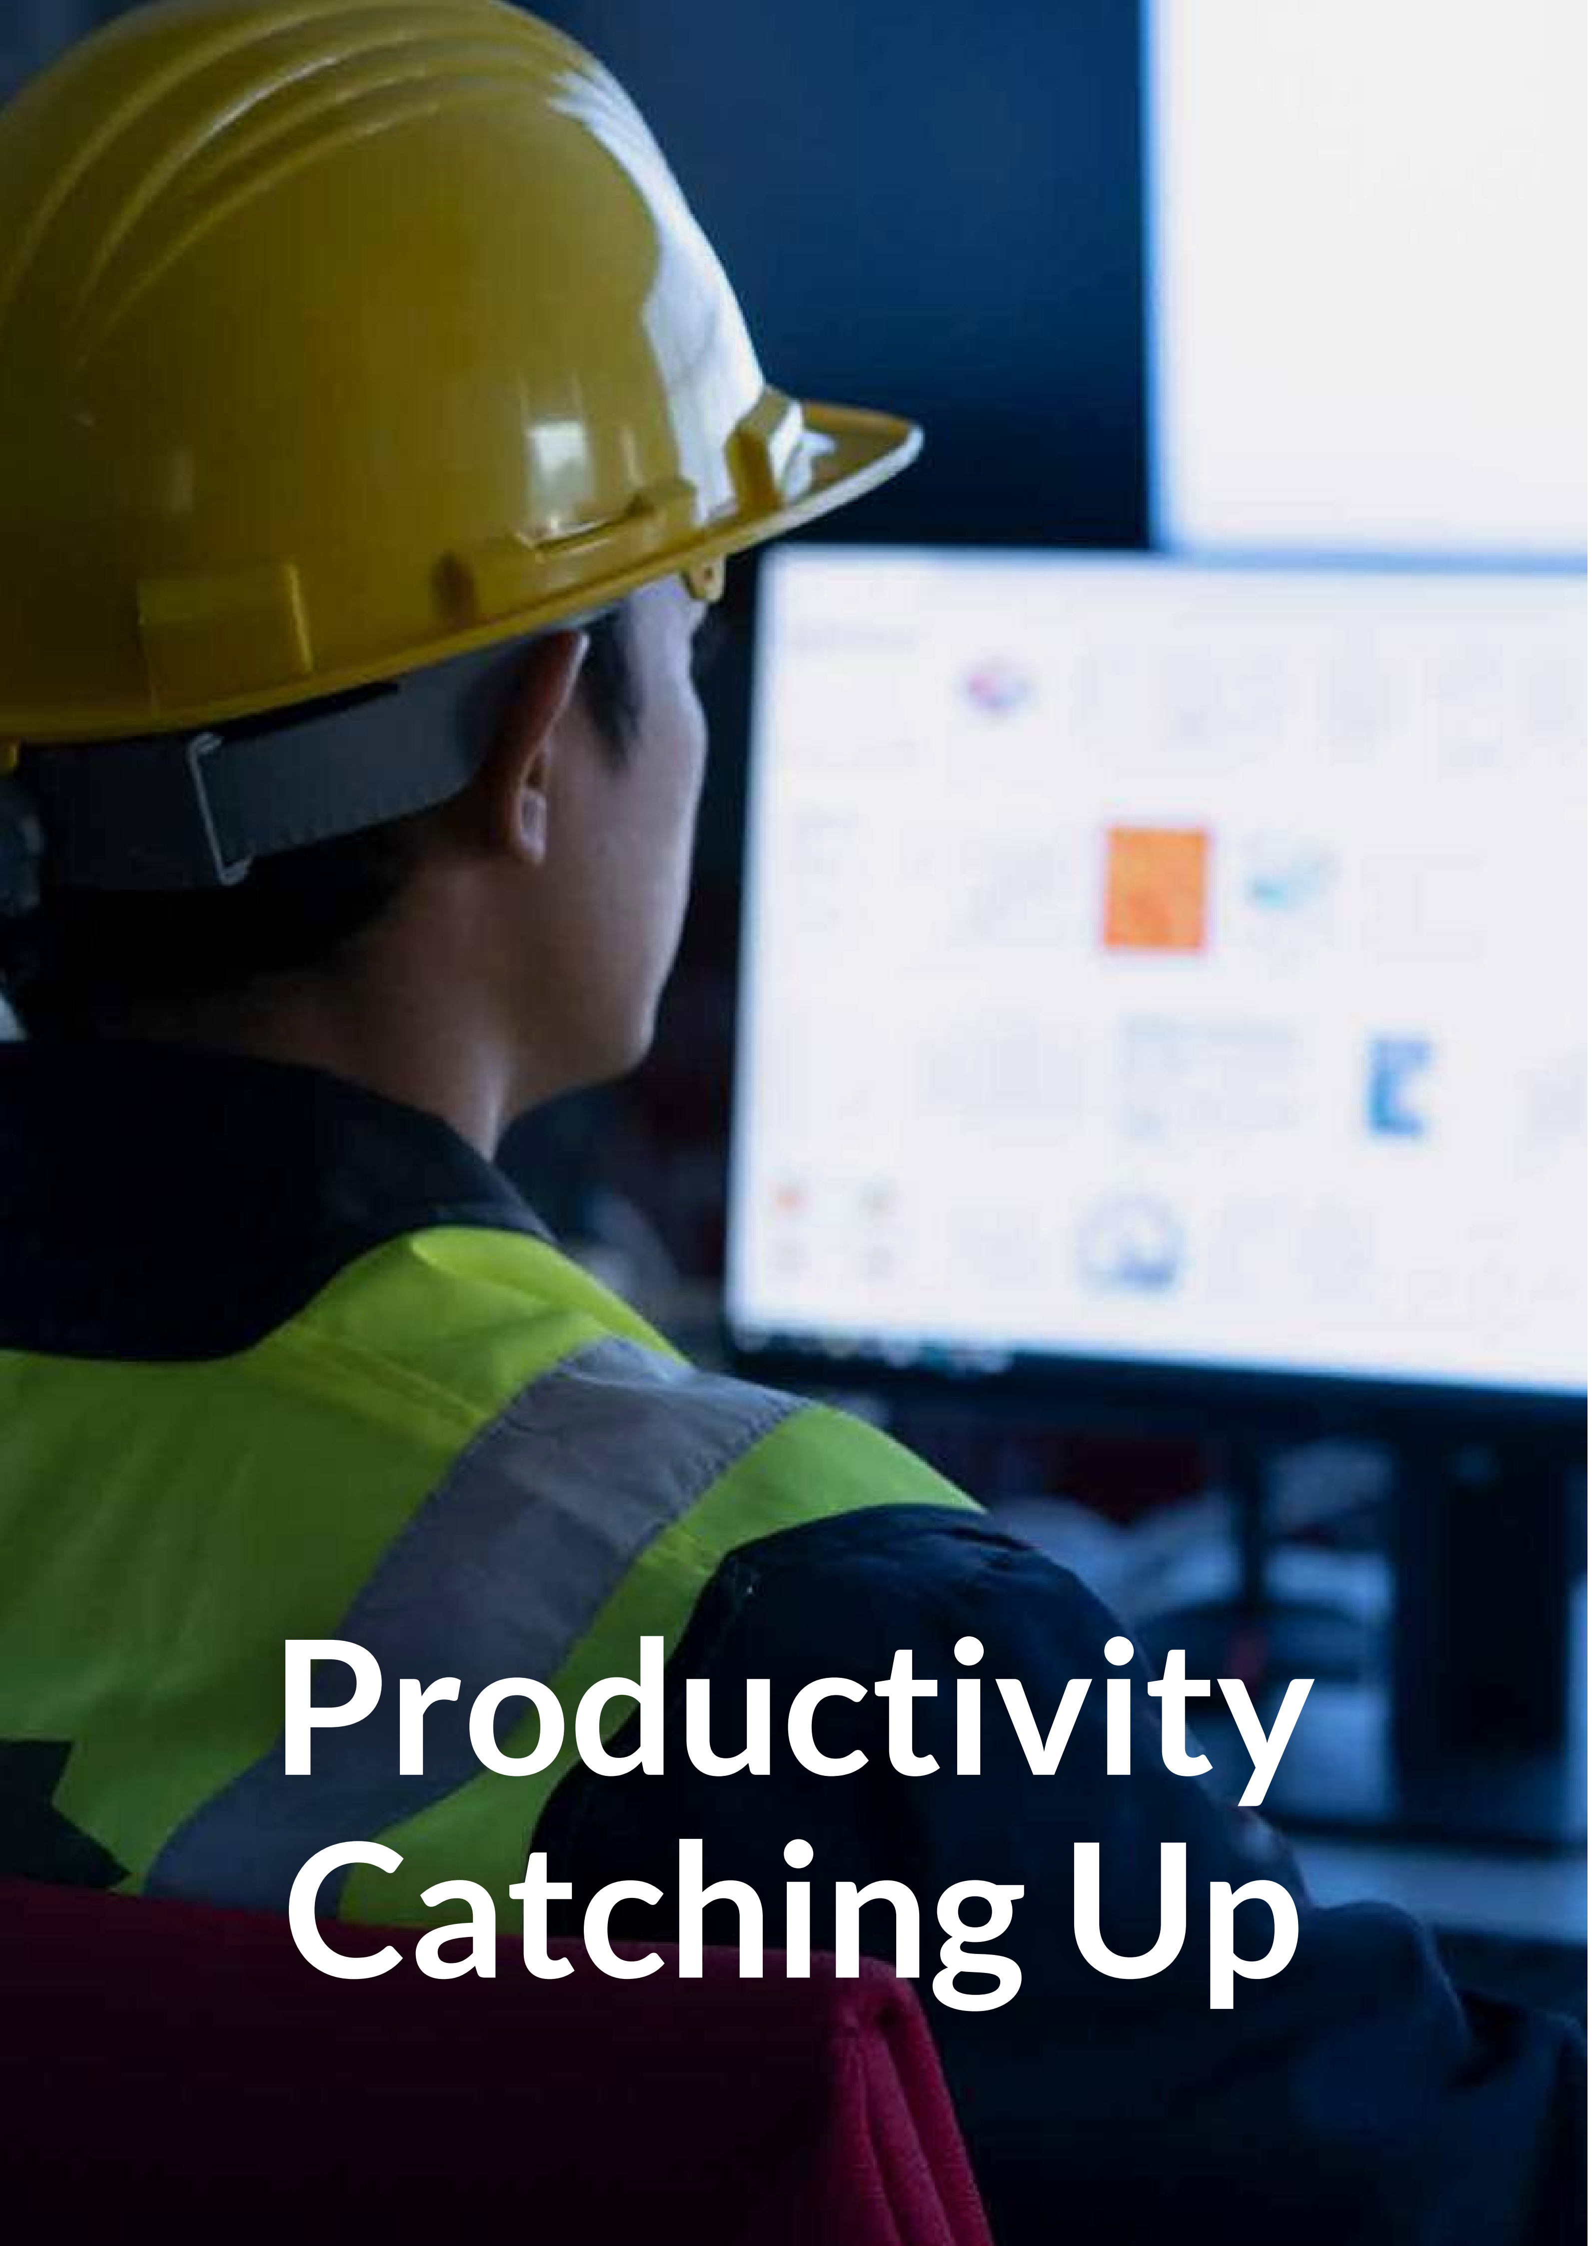

Productivity Catching Up

The Flatlining Indicator

Globally, construction labour productivity has grown by a paltry 1% annually over the past twenty years, significantly lagging behind the productivity of other sectors. In the extreme case of the US, construction productivity has declined since 1968! In the most up-to-date case of Singapore, construction productivity has grown at less than 0.3% annually between 1992 and 2018[12]. Compounded over the 1995 to 2014 period, global productivity has nearly doubled in manufacturing while increasing by a mere 20% in construction.

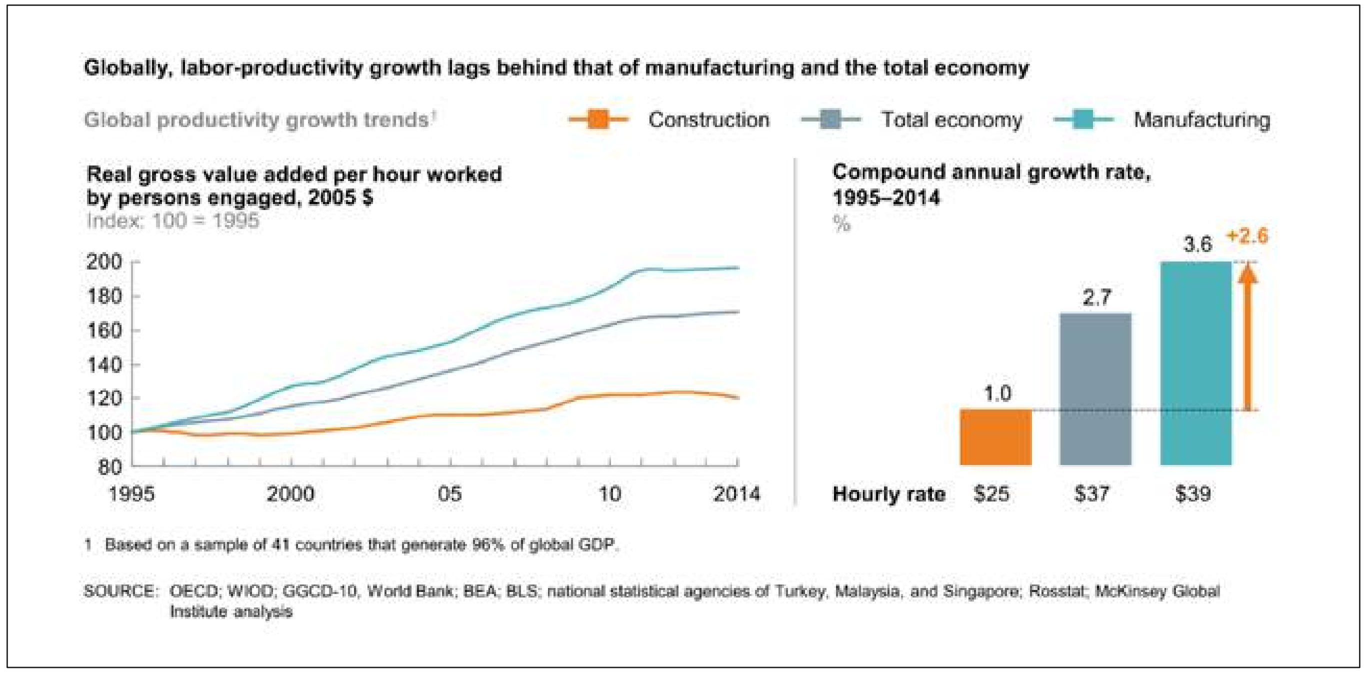

Above: Construction's global productivity growth between from 1995 to 2014. (source: McKinsey Global Institute Analysis).

Construction productivity is a critical parameter because labour costs account for 30% to 50% of the total cost of a construction project. If the productivity in construction were to catch up with the productivity of the global economy, construction value added would increase by almost 50% while using the same resources[13].

According to McKinsey, it is indeed possible for construction productivity to catch up by gaining 50–60%. Roughly a third could be gained upstream of construction firms, particularly in design and engineering, which determines constructability. But the balance would be in the hands of construction firms, in supply-chain management, on-site execution, technology and capability building. This could lead to a 30–40% improvement in productivity and amount to a 15–25% cost saving. Moreover, selected pockets of the construction sector could even increase their productivity by a factor of 5 to 10 by relying on production system techniques and standardisation (for instance using off-site modular pre-fabrication).

Technology & AI to the Rescue

Productivity improvement is best achieved through a systems approach comprising technical, social and managerial components[14]; its implementation involves policies, processes, practices and tools in a people-rich environment. Indeed, it is not only about technology; however, all other economic sectors have exhibited a correlation between the level of digitisation and the growth in productivity.

In the case of construction, digital collaboration platforms will eventually become the norm, enhancing ideas, reducing miscommunication, speeding up information sharing, increasing buy-in, and decreasing later rework and change orders. Drones and other unmanned aerial or ground vehicles, equipped with sensing equipment such as optical cameras or LIDARs, will carry out surveying and monitoring activities, even in previously inaccessible areas. Autonomous pieces of construction equipment will gradually make their way to the job sites: hoisting and lifting machines, brick-laying robots. 3D wall printing is expected to spread, particularly for off-site prefabricated modules.

The workforce will be managed by digital and AI solutions, for recruitment, attendance monitoring, task allocation, productivity and quality performance analysis, detection of training needs, and training effectiveness evaluation.

Sensors, often supplemented by AI, will proliferate as construction adopts the Internet of Things (IOT), allowing for asset and worker geolocation as well as continuous site monitoring. Data will be ubiquitous, and advanced analytics will be used to understand patterns in it. This will increase overall efficiency and effectiveness by deploying real-time solutions, preventing breakdowns and downtime, and cutting waste and idle time. One case is worth detailing.

“In the case of construction, digital collaboration platforms will eventually become the norm.”

Scheduling Right

Ineffective planning and scheduling are often cited as a primary cause of poor productivity, rather than an unmotivated or unskilled workforce[15]. Specifically, top-down master planning fails for two mutually reinforcing reasons: individual tasks are subject to significant variance relative to their planned characteristics, and real-time information on this variability is not fed back from the construction site into the master planning. Consequently, the master plan starts drifting because of individual tasks' lead time being out of control, increasing labour idle time and reducing productivity. Because of the lack of effective updates, issues ripple down the master schedule. Coordination issues are inevitably amplified along the way. Instead of tackling the root causes, a common reaction is to build more labour capacity or material inventory, further reducing productivity in the process, until the master schedule is eventually dropped. Digitisation and AI can help with both root causes and deliver additional productivity benefits.

Better Prediction of the Actual Duration of Each Task

A scheduling AI first learns to predict the actual duration of each single task based on hundreds of input variables. These inputs might come from digital site reports updated in real time (supply delivery, material consumption, equipment usage) or mobile applications tracking tasks assigned to each worker. Predictions could consider weather conditions (extreme temperatures lower productivity); individual worker characteristics (including skill level, performance track record, and recent overtime); and supervisor characteristics

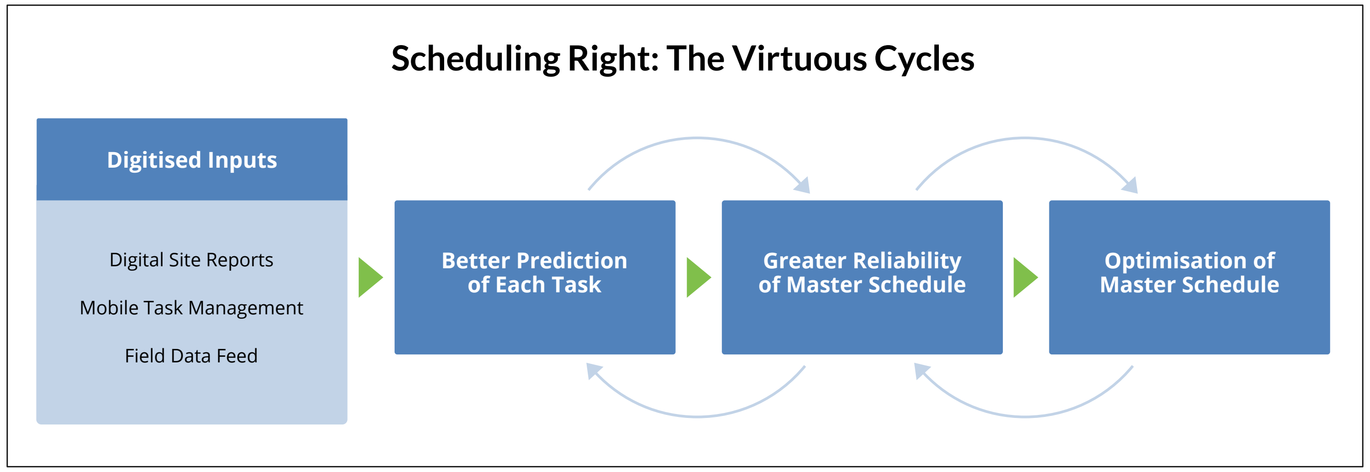

(leadership style, skill and motivational ability). The scheduling AI thus identifies meaningful lead time drivers and reduces the variance of the actual task lead time relative to the plan.

Higher Reliability of Master Schedule

In addition, real-time collection of site data enables the AI to immediately feed information back into the master schedule. Potential drifts are captured from the start and contained. The mechanism is self-reinforcing: higher reliability at task level allows for higher reliability at project level, enabling more precise resource allocation and reduction of idle time or interruption time, thereby strengthening dependability.

Optimisation of Master Schedule

This is only the beginning of the virtuous cycle. As the process becomes better controlled at an individual task level, AI insights on lead-time drivers in conjunction with lean methodology practices enable progressive reduction of task lead time. Once individual tasks are under control and optimised, the scheduling AI can start working on optimising the master schedule itself. Determining the shortest schedule is an optimisation problem under constraints of capacity, resources, and sequencing. The scheduling AI will explore millions of combinations to optimise for a performance indicator – typically lowest total cycle time or cost. Like a digital twin, simulations will be conducted to evaluate different scenarios (for instance different workforce sizes) and select the most suitable option.

It is important to point out that powerful schedule optimisation algorithms already exist today. The paradigm shift here will be driven by the first step: higher reliability at task level. The prerequisite is for site data to be rigorously captured, something achievable now with the deployment of mobile applications, localisation technology, IOT among others.

Let us be perfectly clear on the scheduling AI's likely characteristics and modus operandi. The scheduling AI will of course not compress schedules by excessively shortening incompressible individual tasks; it will instead find the best sequence under constraints for these tasks, which human experience and expertise might have overlooked. The schedule will have to be robust; in other words, its performance should remain relatively unaffected by small variations in input. The scheduling AI, in any case, will run on a frequent basis, possibly daily or even at shorter intervals, so as to take into account changes in the site environment (for instance unexpected weather), project-related events (for example supply chain issues) or construction-related developments (such as change orders).

Experts might argue that scheduling is an art so complex it cannot be given to machines. Humans will indeed be required at the beginning to evaluate the result and spot scheduling constraints or incompatibilities previously unaccounted for. In other words, experts will contribute to the continuous training of the algorithm. But odds are high that the AI would quickly surpass humans[16].

At the end of the day, the well-known correlation between higher productivity, shorter and more reliable schedules will once again be observed.

Fixing Quality

The Indicator that Made a Difference

Quality can make a significant difference to a construction project's performance indicators, commonly bringing about cost and schedule overruns. Although measuring the cost of total quality is a tricky exercise, its level can be estimated between 5% and 20% of total construction cost – depending on the type of construction, the contractors involved, and of course the measurement methodology[17].

Non-quality is often broken down into internal quality failures, detected before the project is handed over, and external failures, detected after the project is handed over. Internal quality failures remain within the builder's control. It is likely that the cost of rework is underestimated, as it seldom reflects the induced variability on task lead time and the indirect impact on overall project schedule. External quality failures harm the contractors' reputation and reduce the owners' goodwill. They represent a missed opportunity for contractors to command premium prices. As for the cost of good quality, it encompasses all the systems implemented to generally avoid non-quality, to contain quality failures when they happen, then to mitigate and eliminate non-quality altogether. Overall, the labour component of quality management is significant, be it for detection and remediation of non-quality or total quality management including preventative training.

The sheer volume of tasks related to non-quality can quickly become staggering. Take, for instance, a residential compound. The builder maintains total quality management during the building phase and conducts a pre-inspection before handover to the developer. A so-called final inspection takes place jointly with the builder and the developer. Lastly, at handover to the unit owner, quality is checked again. In spite of the successive checks, it is not uncommon to have 20 issues lodged per unit at this final stage. This quickly adds up to 10,000 issues for 500 units. These 10,000 issues need to be identified, monitored, remedied, and closed.

The Sharpest Eye

Spotting quality defects, confirming their nature and recording them in proper IT systems are currently tedious and labour-intensive tasks. A defect can easily be missed if environmental conditions are not ideal – for instance working in insufficient lighting – or if the quality inspector cannot give their full attention because they are tired or distracted. Quality issues may also end up inaccurately reported and poorly tracked.

Machines do not suffer from these human weaknesses. As previously described, AI has no match when it comes to sensory perception and more specifically visual perception. The only limitation in today's state-of-the-art technology is that models must have been previously trained – AI cannot currently recognise what it hasn't seen before. Train your model with enough examples of a particular object, shape, colour, defect or imperfection, and it will surpass humans at identifying them.

In the relative short-term, instead of sending skilled quality inspectors throughout the construction site, optical sensors will be mounted on handheld devices and entrusted to less-skilled personnel. In each area to be inspected, data will be captured almost instantaneously on-site and analysed by AI. A simple visible-light camera sensor will spot scratches on doors, shade variations on windows, or missing tile fragments on the floor. A LIDAR will take measurements and check the correctness of dimensions, identifying walls that are not perfectly vertical, wrongly positioned pipes or improperly fitted appliances. An infrared camera will detect abnormal heat patterns and spot insulation defects. An X-ray camera will visualise faulty steel reinforcement in concrete beams.

The Short-Term Impact: 'Augmented' Inspectors

Above: Example of an augmented reality app for construction.

As a next step, instead of mounting sensors on handheld devices, they will be fitted on robots, or even better, on autonomous vehicles. The inspection will be more productive and more comprehensive while removing all risks that would otherwise be taken by humans. Drones might fly up and down external walls of high-rise buildings, looking for cracks in façades, defects in paintjobs, or issues in roofs. Animal-like robots may roam around your site, tirelessly going up and down the stairs like actual dogs, venturing near open edges, walking over or under obstacles, exploring hazardous areas. All the while, cameras would scan the building in a systematic way.

The Brain Behind the Eye

Once digitised, quality issues are automatically fed into rectification workflows. The quicker a quality issue is solved and closed, the lower its impact on overall project cost and schedule overrun. In contrast, the amount of rework, the spread of quality issues, and the disturbance of material flow (with mounting inventory or waste) and human flow (with idle time or zero value-added work) can grow to a considerable extent if remedial action is delayed. In the simple residential case, think of the benefit of detecting water leaks in new units before plastering is done, let alone before damp starts building up.

The Long-Term Impact: Fully Automated Inspections with Field Robots & Drones

Above, left: Robot for construction (source: Boston Dynamics). Above, right: Drone being used on a construction site.

Traditional data science extracts valuable insights from non-quality data using descriptive analytics. Which quality issues are the most common? How long does remedial action take on average, and in extreme cases? Which contractors or workers generate the most issues? Do we fundamentally have a task problem or a labour one? This information is traditionally used to derive action plans.

AI and advanced analytics yield more sophisticated information. For instance, with clustering techniques, AI might be able give a completely new perspective by grouping issues in novel categories, each requiring a specific action plan.

By analysing numerous input variables and finding patterns in the data, AI might warn of other faults, not yet detectable or possibly overlooked, but likely to have already occurred. Alternatively, it might predict upcoming faults, even if the subtle signals are not yet perceptible by humans. But AI can observe that a specific combination of circumstances

typically leads to the occurrence of a given defect. The prediction can then be used to take preventative measures, fine-tuning or resetting the equipment, changing the material used, adjusting procedures and work packages, or retraining or replacing labour.

When a fault occurs, a trained AI algorithm can attempt to forecast the remedial time and cost. Combined with the scheduling AI, it is also best positioned to measure the impact on the overall schedule and consequently decide on the magnitude of the resources to allocate to fix the issue.

Through anticipation of upcoming defects and a swifter and better-informed reaction to quality issues, the true cost of quality is bound to decrease thanks to Artificial Intelligence.

Taming Safety Demons

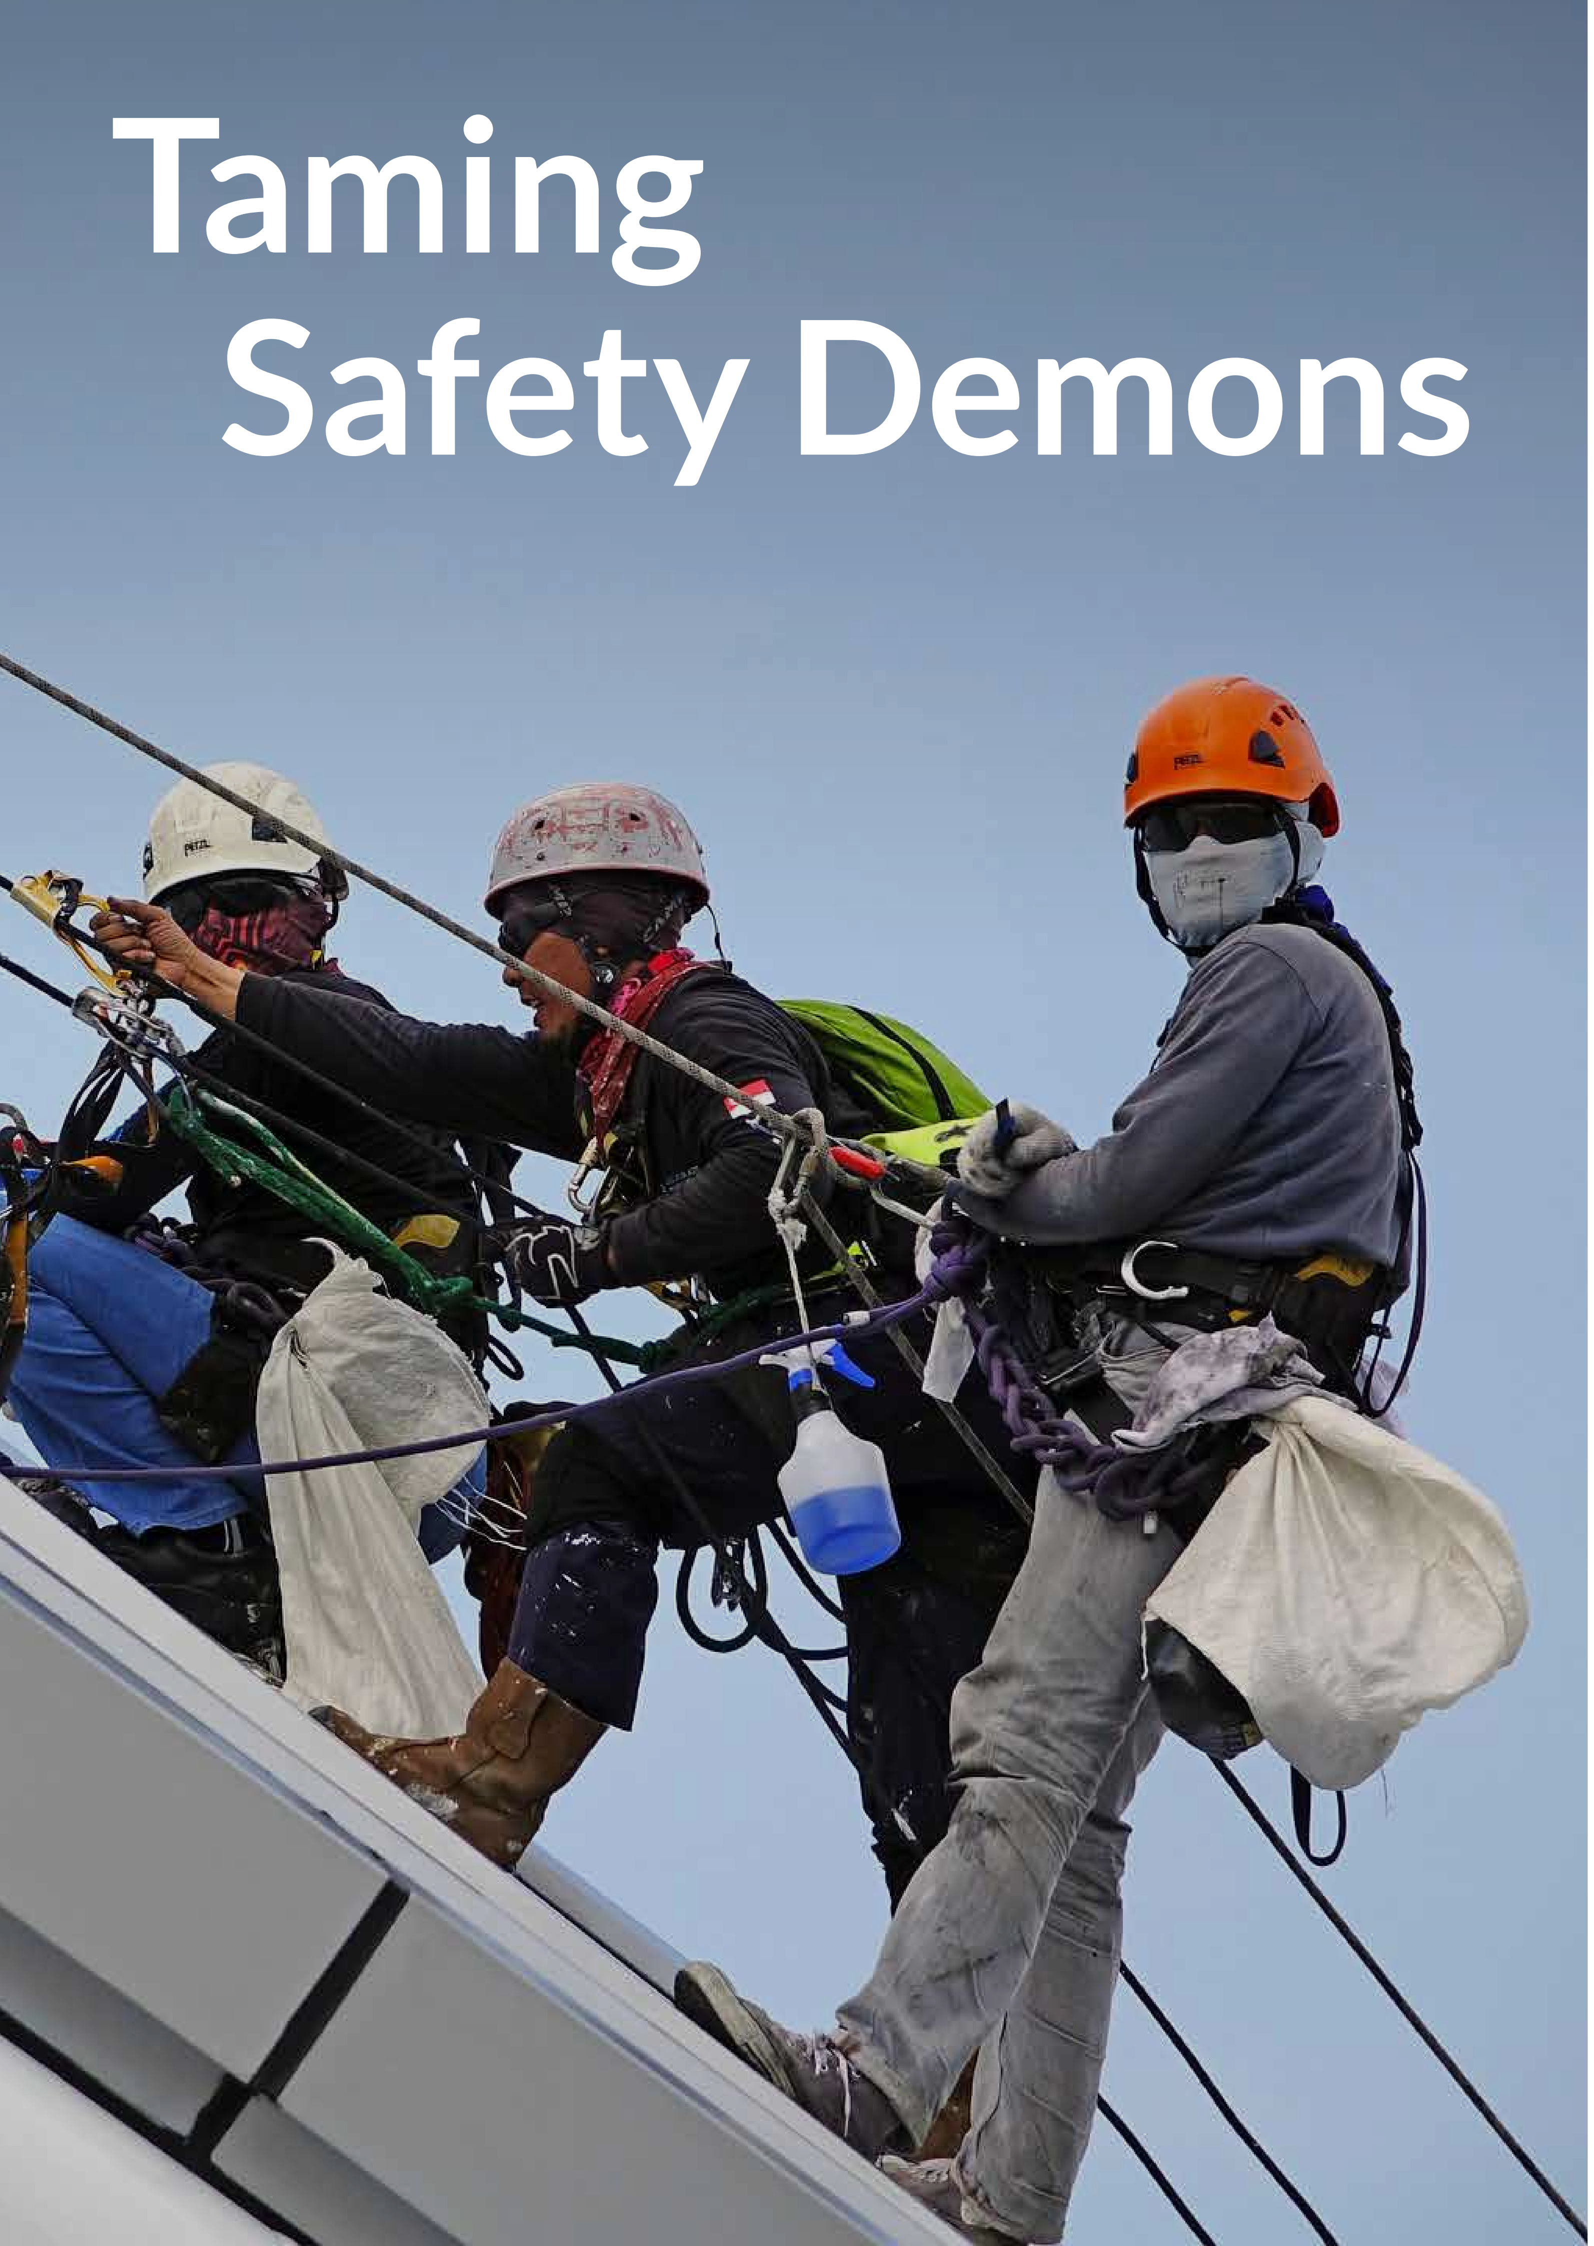

The Industry’s Dark Side

Despite numerous industry and regulatory efforts, construction remains one of the most dangerous sectors to work in. In the United States alone, 971 construction workers lost their lives in 2017. This is a fifth of all worker deaths in the US that year[18]. In the European Union, there were 815 fatal accidents in 2015 (7 per 100,000 persons employed), and almost 400,000 non-fatal accidents resulting in at least four days of absence from work[19]. In Singapore, with 14 fatalities in 2018, the construction industry remains the highest contributor of workplace fatalities despite a significant decrease over the years[20]. Worldwide, of all fatal construction accidents, falls from height take the highest toll by far (between a third and a half of the total, depending on the country), followed by being struck by a falling or moving object, being crushed or trapped, and electrocution[21].

Why have safety statistics improved so little on a global basis in spite of all the safety initiatives? Human expertise has until now been the main weapon to fight accidents. But no matter how experienced any single safety manager or construction worker may be, they have only come across tens or hundreds of accidents at the very most during their entire career. More examples or more time might not have helped anyway: humans are limited in their ability to process large amounts of data. They can only handle so many incidents or consider so many variables for each single incident. Furthermore, all individuals are subject to many cognitive biases. All in all, it has become nearly impossible for humans to better recognise hazards, correctly estimate risks, and avoid the occurrence of accidents. Enter Artificial Intelligence.

Precision Safety

Precision medicine, as we have seen above[22], can theoretically predict a subject’s health evolution and hopefully influence it for the better by precise, fully personalised treatment. Similarly, precision agriculture is already practically capable of feeding plants precisely with the nutrients they need, where and when they need them, resulting in increased effectiveness for plant development and up to an 80% reduction in nutrients used[23]. Why not aim to build precision safety, where the availability of multiple variables would enable AI to predict construction site incidents and hopefully prevent them by the right intervention?

Research labs around the world have accepted the challenge. For example, at the University of Colorado[24], researchers used a list of nearly 5,300 injury reports gathered from 470 contractors and representing millions of work hours. They were able to skilfully predict, based on 80 input variables, the type of injury and the body part affected in construction incidents. In another study conducted in Singapore[25], a large contractor provided data from 27 construction projects including 785 monthly safety inspection records. Based on 13 input variables, the algorithm was able to predict with high reliability the occurrence and severity of accidents with high reliability.

The research is nascent and not fully robust yet. The sets of attributes used to characterise incidents and predict future occurrence vary from model to model, and each is clearly incomplete. But one can safely assume that in the foreseeable future, precision safety will avail of data sets as rich as those aspired to by precision medicine.

Data gathering will start with overall project characteristics (e.g. type of construction, ownership, size in cost and manpower) and consider its progress at the time of the incident (e.g. completion stage, actual delay). The general safety record of the company and of the project will be included (e.g. safety inspection results and past incidents). The intrinsic characteristics of the task causing the incident will be fully recorded (e.g. precise task description, equipment and energy used, and location type). The specific circumstances at the time of the incident will be captured (e.g. site condition, headcount and equipment in the vicinity, time and date, air temperature) as well as the actual execution of the task (e.g. body position and movement, availability of work permit, compliance with instructions). The worker will be entirely described (e.g. age, experience, qualification, general training, task training, and safety track record) as well as their particular state just before the incident (e.g. body temperature, blood pressure, heart rate, breathing rate, last statements, sleeping and eating records). Finally, the incident outcome itself will be depicted with great care (e.g. type and severity of injury, and body part affected).

With all this information available, the precision safety algorithm – let us call it SafetyAI – would be trained and would figure out which combination of factors is most likely to lead to an accident. The result would not necessarily be the one expected – this is precisely the point of resorting to a system capable of spotting correlations where humans cannot, or better still, capable of deducting causality.

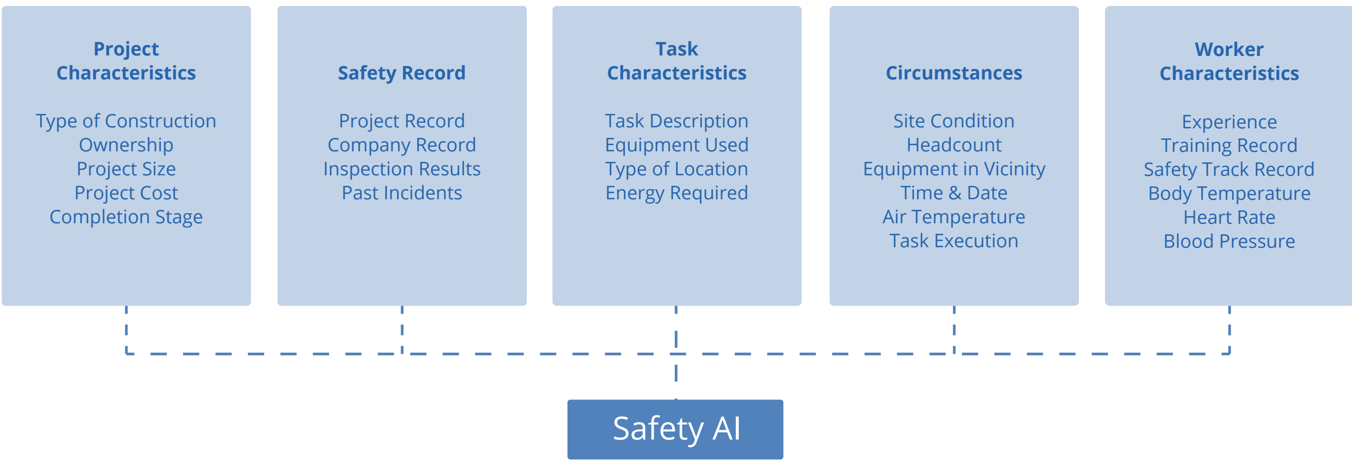

Saving Lives

With the precision safety algorithm switched to prediction mode, SafetyAI would be able to determine the risk of occurrence of accidents and provide further information on the immediate cause, the affected body part, and the severity. This ability to predict outcomes means that, given the right level of preparation and reaction, accidents could be anticipated, mitigated, and at best, avoided altogether. SafetyAI would simulate the effect of a precision treatment.

For instance, regular preparation routines would involve running SafetyAI ahead of the workweek. Managers could be informed, with high accuracy, of accidents most likely to happen given general site conditions and expected work packages. They would also receive suggestions of remedial action. Such actionable feedback could lead to highlighting specific risks during pre-job safety meetings, providing extra training and supervision, reinforcing targeted safety measures, reassigning work to well-suited workers, or revising the work package altogether, thereby decreasing risk in a preventative manner.

SafetyAI would also monitor the construction site in real time, using live data.The changing conditions – whether of site or workers – would thus be continuously fed into the algorithm for analysis and warning whenever needed. Consider, say, workers operating at height, on a hot August day at 11am. Vitals would be monitored, in case the body temperature should increase so much, or the blood pressure or the heart rate fall so low that in combination with the other factors, the risk of dizziness and fall would become unacceptable. In addition, workers would constantly be located on the site. They would set off alarms when venturing into dangerous areas or if heavy equipment moved too close for comfort, increasing the calculated risk level in conjunction with other variables tracked.

The variables feeding SafetyAI in real time would be either manually input or automatically captured by digital sensors, for instance to track weather conditions – heat, rain – or site luminosity. AI would undoubtedly come to the rescue for data capture as well. For instance, computer vision would identify workers wearing – or not wearing – their helmets or being

Examples of Incident Prevention with SafetyAI

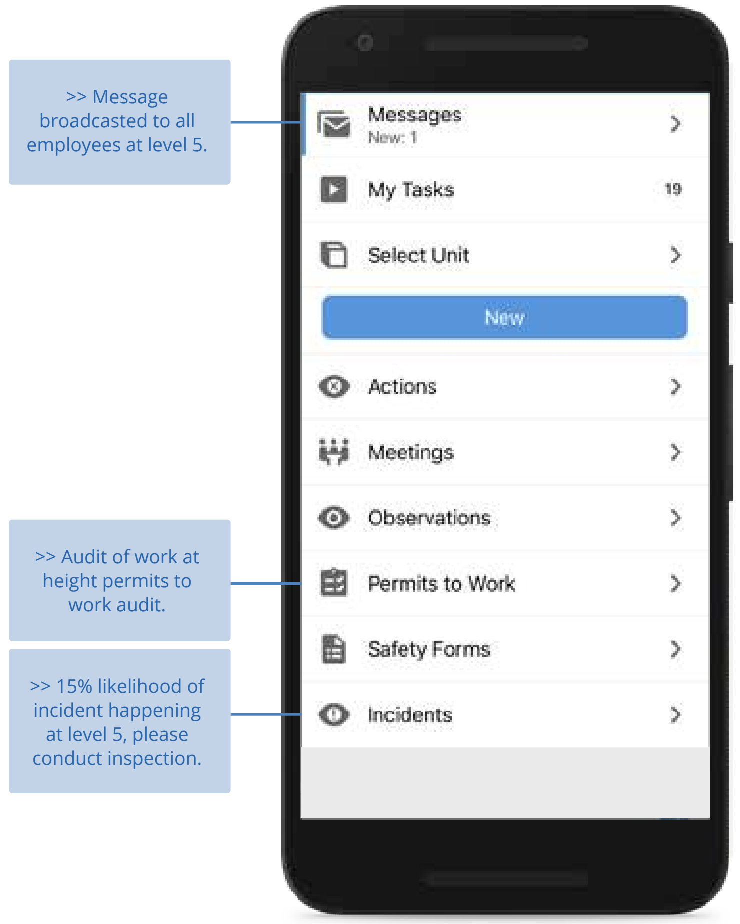

properly attached – or not – to safety harnesses. Computer vision could also identify the presence or the absence of safety rail guards, and more generally the existence of fall-protection systems or conversely the presence of risky open edges.

It may even happen in the future that workers will have a guardian angel in the form of a drone hovering not far from them. Before any potentially difficult work the guardian angel would check certain parameters, ask safety questions and give instructions via an earplug. During the work phase, it would constantly monitor the worker's status as well as the site conditions, anticipating possible difficulties given their work package and their personal track record. As risk increased, the guardian angel might advise or warn the worker, before triggering an alarm should risk levels exceed acceptable thresholds. Lives would be saved.

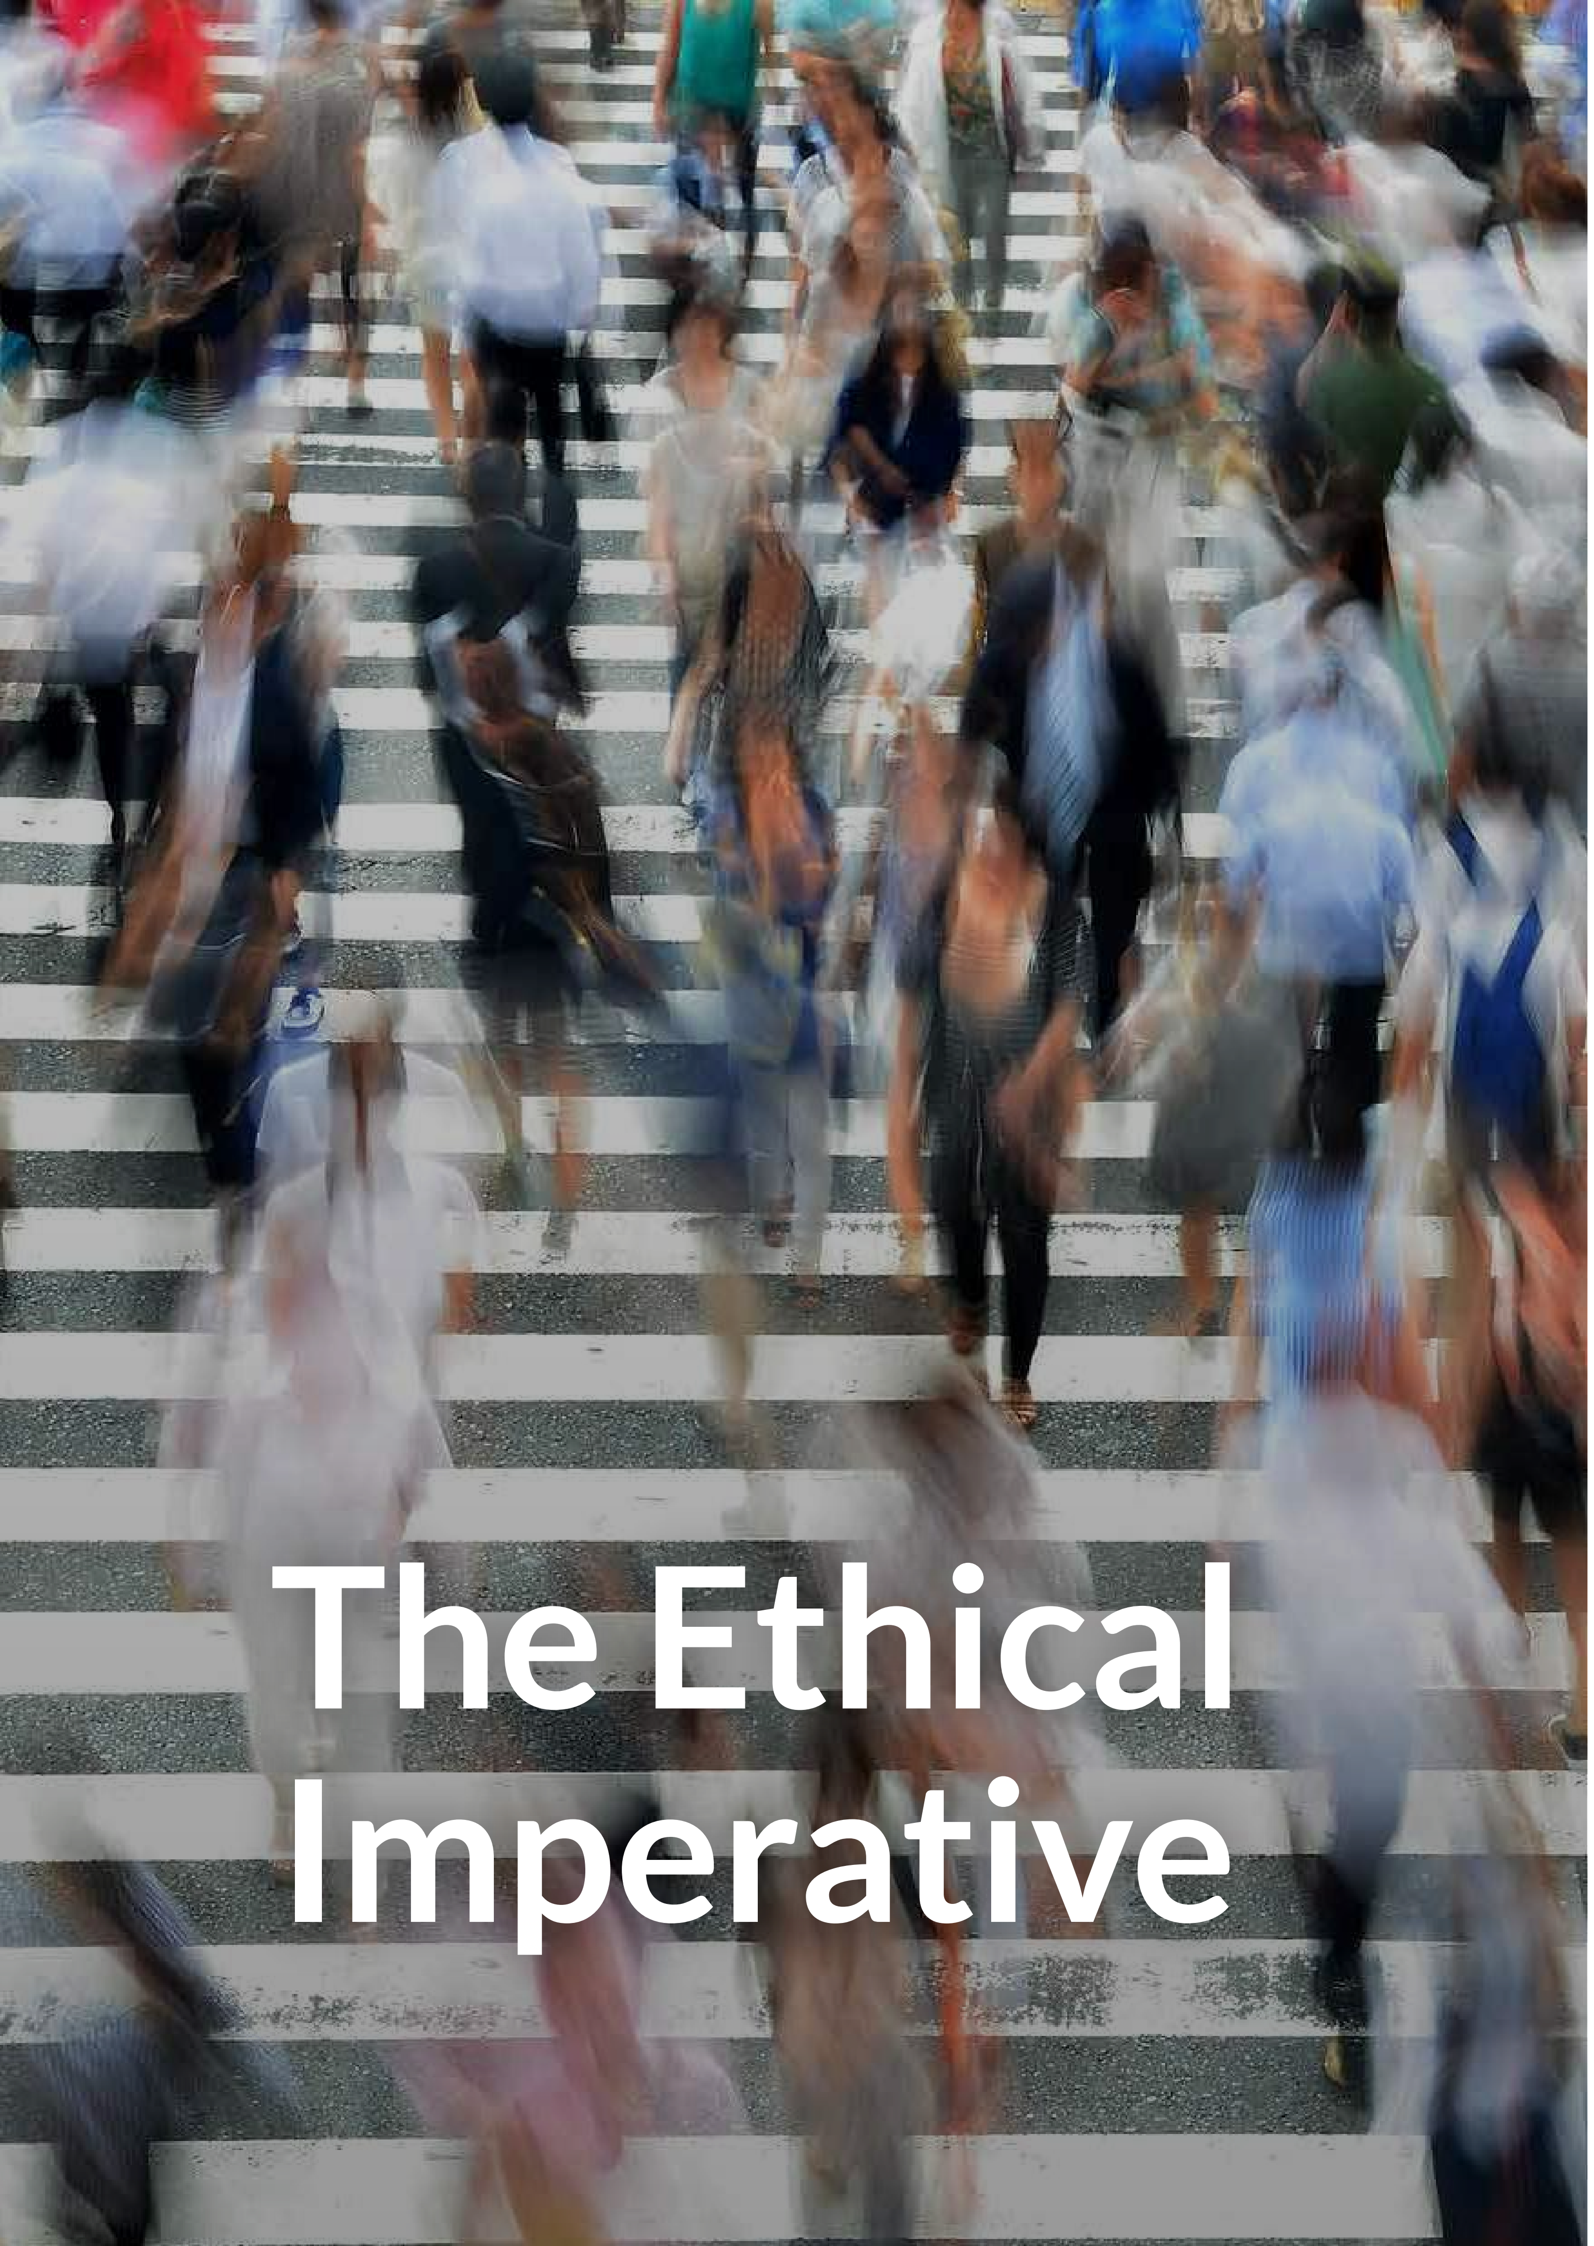

The Ethical Imperative

Trustworthy AI

There is much more to AI implementation than a simple, ordinary digitisation project. AI is becoming more powerful than nuclear reactions, which can be used both to produce energy for society and to manufacture lethal weapons. However, instead of necessitating the world's smartest scientists and rare fissile material to operate, AI is available to virtually anyone with access to computing power and an internet connection. Harm can be caused in multiple ways, especially unintentionally. In order to prevent the dangers of AI from materialising, many across the world have called for ethical safeguards. The question is thorny since ethics is a highly cultural and personal concept.

Europe is at the forefront of ethical thinking about AI. In 2019 the European Commission released its Ethics Guidelines for Trustworthy Artificial Intelligence[26]. According to the Guidelines, trustworthy AI should be (1) lawful, respecting all applicable laws and regulations, (2) ethical, respecting ethical principles and values, and (3) robust, from a technical and social perspective. The Guidelines put forward a set of seven key requirements that AI systems should meet in order to be deemed trustworthy.

Human Agency & Oversight: AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights. At the same time, proper oversight mechanisms need to be ensured, which can be achieved through human-in-the-loop, human-on-the-loop, and human-in-command approaches.

Technical Robustness & Safety: AI systems need to be resilient and secure. They need to be safe, ensuring a fall-back plan in case something goes wrong, as well as being accurate, reliable and reproducible. That is the only way to ensure that harm, including unintentional harm, can be minimised and prevented.

Privacy & Data Governance: Besides ensuring full respect for privacy and data protection, adequate data governance mechanisms must also be ensured, considering the quality and integrity of the data, and ensuring legitimised access to data.

Transparency: The data, system and AI business models should be transparent. Traceability mechanisms can help achieve this. Moreover, AI systems and their decisions should be explained in a manner adapted to the stakeholder concerned. Humans need to be aware that they are interacting with an AI system, and must be informed of the system's capabilities and limitations.

Diversity, Non-Discrimination & Fairness: Bias must be avoided, as it could have multiple negative implications, from the marginalisation of vulnerable groups, to the exacerbation of prejudice and discrimination. Fostering diversity, AI systems should be accessible to all, regardless of any disability, and involve relevant stakeholders throughout their entire life cycle.

Societal & Environmental Well-being: AI systems should benefit all human beings, including future generations. It must hence be ensured that they are sustainable and environmentally friendly. Moreover, they should take into account the environment, including other living beings, and their social and societal impact should be carefully considered.

Accountability: Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes. Auditability, which enables the assessment of algorithms, data and design processes, plays a key role, especially in critical applications. Moreover, adequate and accessible redress should be ensured.

Some of these key requirements are likely to raise multiple questions on the construction site if not thought through well enough.

Privacy, Transparency & Non-Discrimination

Imagine SafetyAI up and running on the construction site. By scrutinising and processing thousands of parameters on all aspects of the running site, SafetyAI is indeed saving lives. While doing so, it collects large amounts of data on Joe, the worker: Joe's body temperature, his facial expressions, his pulse, his body movements, the timing and frequency of his breaks, and of course all his medical records available to the company. Three scenarios depict what could happen next.

In our first scenario, the AI almost inadvertently discovers that Joe suffers from a rare disease. This disease does not prevent him from carrying out the work currently allocated to him, but it is unclear how his health situation will evolve over the next few years. It appears that Joe himself is not even aware, whereas his manager got hold of the information through the system. What if the manager makes use of this data to make him redundant? What if Joe stays on the job but the insurance company finds out about his health situation? Regardless of employment considerations, would Joe even want to be informed about his condition? Data privacy and governance should be worked out in advance to avoid such questions.

In our second scenario, SafetyAI suddenly warns of an imminent incident very likely to involve Joe. The operational manager immediately puts the construction site on hold, requests Joe to step aside and asks SafetyAI for more information. Unfortunately, while SafetyAI was able to forecast an impending incident with high likelihood, it is unable to explain why. Its algorithm is currently based on deep neural networks, and the explicability of this algorithmic class is very low. SafetyAI just doesn't know, and Joe is left hanging because of the black box that is SafetyAI. Different technical choices should have been made to ensure transparency.

Our third scenario, like the second one, involves the forecast of an imminent incident, the construction site is put on hold and Joe is requested to step aside by his manager. This time, it doesn't take long before the main driver for SafetyAI's prediction is found. In the past couple of months, there have been multiple safety incidents involving workers originating from the same country as Joe. It appears the country of origin was taken as a parameter. Joe was therefore penalised by the algorithm not because of who he is as an individual, but because of the country he comes from – and maybe because of his race. The algorithm clearly contains bias and is discriminatory.

The Construction Moral Machine

The emblematic case of ethical design for AI systems resides in the so-called trolley problem. Imagine, for a moment, that the brakes of an autonomous car suddenly give way, when it is driving at full speed. Pedestrians are crossing the road at that very moment. Whom should the car save amongst passengers and pedestrians? How will this autonomous system decide who should live and who should die? Addressing this problem of the "Moral Machine", a survey was conducted involving millions of participants around the world[27]. The survey demonstrates that some preferences are universal: sparing humans rather than animals, saving as many people as possible, and favouring children. But there are also individual preferences and even quite strong cultural preferences about saving women versus men, the elderly versus the young, the rich versus the poor, the law-abiding versus the jaywalkers... or just letting the car follow its initial path.

As we have seen, autonomous systems will inevitably appear on construction sites too: very large hoisting and lifting equipment, transportation vehicles, automated wall-erecting systems, load-carrying robots, welding and cutting machines, quality control drones and animal-shaped robots. Once systems on site are mostly autonomous, the rate of accidents is bound to decrease significantly. But in the interim period when machines and humans coexist on a jobsite and each bring in their own logic and behaviour, the occurrence of accidents cannot be ruled out.

Above: Autonomous crane technology (source: Intsite).

Construction autonomous systems too will raise ethical issues, ushering in the construction Moral Machine. Artificial Intelligence will be on watch, computing real-time risk thresholds. In some cases, it will manage to communicate danger levels to humans who will have enough time to make the final decision – for instance, continue operations or bring everything to a halt. In other circumstances, the machine might have to decide by itself, typically sparing someone or something at the expense of someone or something else. Difficult decisions could result in causing widespread idle time, sacrificing loads of cement or allowing an expensive piece of equipment to crash beyond repair. People rarely consciously think about these decisions, but in the case of autonomous systems, decision guides will have to be thought through and provided to machines beforehand. AI designers or implementers may have to decide what costs they are ready to bear in order to save a life – in other words, they might have to put a price tag on human life.

A Journey of a Thousand Miles

As we have seen, Artificial Intelligence will bring about a giant leap in the performance of the construction. Schedules will become shorter and more predictable. They will form a virtuous cycle together with enhanced overall coordination, more efficient supply chain management, leaner inventory, reduced labour idle time and increased productivity. This will result in higher rates of on-time and on-budget project delivery.

Quality issues will be quicker to detect as computer perception ushers in more effective and efficient processes. Faults will be remedied faster and eventually become easier to predict and prevent, in turn mitigating the impact of non-quality on project cost and schedule. For contractors, this will lead to higher profit margins, improved reputations and opportunities for premium pricing.

Safety issues will be better understood as their most impactful drivers are identified and weighted. This will enable construction firms to implement a wide range of preventative measures. After years of mixed results, safety-related efforts will finally be rewarded by the reduction of incident frequency and severity. Ultimately, lives will be saved.

But the scheduling AI, the quality AI and the safety AI are not off-the-shelf products, unlike autonomous vehicles or invoice processing software. As well defined as their engines – the algorithms – may be, they require oil to function. More specifically, they need to be trained with your data. Their performance will improve with the quantity and quality of data you can provide.

Gathering such data is a journey of a thousand miles. Like other long journeys, it starts with a single step – namely, the decision to digitise. Long before AI is eventually implemented, workflows need to be defined and relevant data must be captured. The time to start is now. Along the way, AI implementation will require expert support, not only to overcome well-documented, standard digitisation hurdles, but also to avoid AI-specific pitfalls. Ethical concerns epitomise these novel emerging challenges, which will separate average players from the best, most visionary and most respected.

Only when the multiple facets of Artificial Intelligence are thoroughly addressed will the construction sector reap the immense benefits of this revolutionary technology.

Sources Cited

[1] Orbis Research, Global construction Outlook to 2022, https://www.orbisresearch.com/reports/index/global-construction-outlook-to-2022-q4-2018-update

[2] Despite an increase in labour costs across the industry.

[3] McKinsey Global Institute, "Reinventing construction: a route to higher productivity", February 2017.

[4] McKinsey Global Institute, "The new age of engineering and construction technology", July 2017.

[5] World Economic Forum (WEF) and PWC, "Harnessing Artificial Intelligence for the Earth", Fourth Industrial Revolution for the Earth series, 2018.

[6] Bernard Golstein, "Duality: Prepare Yourselves and Your Children for the Age of Artificial Intelligence", KDP, 2019.

[7] Martin Ford, "The Architects of Intelligence", Packt, 2018.

[8] In this 2500-year-old board game, there are more combinations for the stones on the board than there are atoms in the universe. Over the centuries, strategies were crafted by human players only to be defeated by AlphaGo, decades before observers thought it might happen.

[9] According to Levinthal's paradox, these possibilities would take longer to enumerate than the age of the universe.

[10] The result is not perfect but the performance is better than that of humans: https://arxiv.org/pdf/ z1611 .05358v1.pdf

[11] One AI has matched the performance of a panel of dermatologists in 2017 in classifying 2,032 skin diseases.

[12] Computed by the authors based on the following dataset : Changes in value added per worker in 2015 dollars per detailed industry, Data.gov.sg, https://data.gov.sg/dataset/ changes-in-value-added-per-worker-in-chained-2015-dollars- by-industry-ssic-2015-annual?view_id=d47f1acb-0a84-4fe5-8058-6ba17b2c93ea&resource_id=594b0405-8395-47e4-be8c-ae2c50e69808

[13] Unless otherwise specified, most of the data from this paragraph is derived from McKinsey Global Institute, "Reinventing Construction: A Route to Higher Productivity", February 2017.

[14] Shamil Naoum et al., "A new framework for determining productivity factors on construction sites", CIB Joint International Symposium, Construction Facing Worldwide Challenges.

[15] Shamil Naoum et al, "A new framework for determining productivity factors on construction sites", CIB Joint International Symposium, Construction Facing Worldwide Challenges.

[16] The AlphaGo example mentioned on page 7 is a very good example: much to all the experts' dismay, in a highly complex game, AlphaGo invented strategies that were thought to be losers yet turned out to be the most effective.

[17] B. Abbasnedjad, "Poor quality cost in construction: Literature review", RMIT University Master's Thesis, 2013.

[18] United States Department of Labor, OSHA Commonly Used Statistics, https://www.osha.gov/data/commonstats, retrieved 11 Aug 2019.

[19] Global Data, "Safety in Construction, a review of latest statistics". CN0054WP, July 2018.

[20] Government Technology Agency of Singapore, "Annual workplace injuries", data.gov.sg

[21] Refer to sources 17 and 18.

[22] See "predictive / prescriptive analytics" page 6.

[23] The expert data was provided on condition of anonymity by a global agrochemical company.

[24] Antoine J.-P. Tixier et al, "Application of machine learning to construction injury prediction", Automation in Construction Volume 69, June 2016.

[25] Clive Q. X. Poh et al, "Safety leading indicators for construction sites: A machine learning approach", Automation in Construction, Volume 93, September 2018

[26] https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

[27] See its dedicated website http://moralmachine.mit.edu/ or the article in Nature https://www.nature.com/articles/s41586-018-0637-6 [50].

About the Authors

BERNARD GOLSTEIN dedicates himself to raising awareness among business leaders and the general public about Artificial Intelligence, the future of work and the future of education. As CEO of sharperAI, he helps organisations transform themselves. He also teaches in a major business school. Prior to that he founded and led an innovative edtech company and was VP in a multinational. Bernard holds an MBA from INSEAD and an MSc from Mines Paristech.

DENIS BRANTHONNE is the CEO and Founder of Novade, and believes in a hands-on approach to building projects. A business leader with a track record in building global teams, Denis grew up in France, obtained his Master of Science degree at MIT, and previously worked at Autodesk as VP of Emerging Markets & APAC. He remains a firm believer in the 'get it done' philosophy.

About Novade

Founded in 2014, Novade is headquartered in Singapore and operates globally. Novade offers smart field management software for the building and construction industry. Site processes including quality, safety, progress monitoring, workforce management and maintenance are digitised and automated using mobile devices. The data captured provides insights to streamline operations and drive results. Leading contractors, real estate developers, owners and operators around the world trust Novade to drive their digital transformation.

For more information about Novade Solutions, visit www.novade.net

Novade