

ARTIFICIAL INTELLIGENCE

Enabling machines to learn

.....
EXPERTS PERSPECTIVES



ARUP

CONTENT

- 1-2 *Foreword*
Chris Luebke and Deborah Smith
- 3-6 *Artificial Intelligence – enabling machines to learn*
Alvise Simondetti
- 7-10 *The natural logic of artificial intelligence*
Professor Mario Carpo
- 11-13 *What images of cats and dogs can tell us about building design*
Professor Anton van den Hengel
- 14-16 *Artificial intelligence for design – 3 generations ahead*
Mike Haley
- 17-20 *Design 2017AD*
Giulio Antonutto
- 21-24 *Photographic memory*
Yung Loo, David Gonzalez, Michael Devriendt
- 25-28 *Unlocking construction productivity through computer vision*
Felix Neufeld
- 29-32 *Protecting asset networks with neural networks –
the automated way*
Josh Symonds

FOREWORD

In 1793 an English poet, painter and printmaker William Blake published a book that documented the different stages of man's life. In particular, one image, *I Want! I Want!*, depicts a tiny figure standing before a celestial ladder that leads up to the crescent moon. The etching acts as a metaphor for humankind's ability to aspire and dream, and suggests an early fantasy of lunar travel. Who would have thought that a scant 200 years later we would be walking on the moon? The result of a dream empowered by three industrial revolutions and technological advances.

The exhibition *Artificial Intelligence - enabling machines to learn* is timely. Unlike the long wait for Blake's realisation of his dream, Artificial Intelligence (AI) development has been relatively short. Blake did not see his fantasy realised, however we stand in the present surrounded by a revolution and evolution of AI. At times this is a conscious interaction (i.e. when we consult the 'oracle' encapsulated within our devices with some random question) and at other times it is unconscious (i.e. when the photographs on our phones are automatically and autonomously meta-tagged and re-arranged). We have an inherent lack of understanding of the strengths, weaknesses, potential and pitfalls of this new capability. We hope to change this.

The exhibition is informed by investigations into the impact of AI with Machine Learning and presents various insights and provocations through four loosely clustered zones: the new approach to AI - strengths and weaknesses, fully learnt machines, machines learning at the edge, and today's existing consumer solutions. The associated publication extends the conversation to include fascinating perspectives on AI from seven industry experts.

Our deep gratitude is extended to Alvis Simondetti whose dedication to being two paces ahead of the leading edge of design tools is an inspiration. He has drawn the exhibition and publication together with colleagues from research over the last two years. The design has been realised by the energy and thoroughness of Alex Bourganou, Melissa Mak and Francesco Anselmo who have provided the spatial fabric for the presentation, with a fitting graphic design by Katharina Efremov, Katharine Hogan and Ruth Huntington. With Francesco Anselmo at the helm, assisted by Angeliki Bakogianni and Lucas Broux, an integrated lighting design lights the way to an exploration of new approaches.

A wholehearted thanks and appreciation goes to the gallery exhibition team: Toria Richardson and TPA Modelmakers who are

truly exceptional. Martin Ansley-Young has been a great liaison as well as parting waters to ensure the appropriate timing and funding was secure. We would like to thank the contributors to the publication for their stimulating essays. They include the curator Alvis Simondetti, Professor Mario Carpo, Michael Devriendt, David Gonzalez, Mike Haley, Giulio Antonutto, Professor Anton van den Hengel, Yung Loo, Felix Neufeld and Josh Symonds.

Our project was enriched and realised by many, with the generosity of the following lenders: Arup Inspire, Ambi, Comfy, Autodesk, Google Creative Lab, Mamou-Mani Architects, NVIDIA, TED and IBM Watson and Yarn.

Finally, the aspiration of learning why, what and how (may it be to reach the moon) is what compels us to gain a deeper understanding and explore further. The exhibition *Artificial Intelligence – enabling machines to learn* places us on that ladder.

Deborah Smith, Arup Curator.
Chris Luebke, Arup Fellow, Director for Global Foresight, Research + Innovation.

ARTIFICIAL INTELLIGENCE – ENABLING MACHINES TO LEARN

The press has bombarded us with rhetoric about the ‘rise of robots’ and their potential to threaten humanity. Fear of the new and unknown certainly sells papers and magazines. But Arup is always compelled to know more through sustained exploration and deeper understanding. The exhibition *Artificial Intelligence – enabling machines to learn* has provided the opportunity to share our key findings from a period of two years focussed on the impact of Artificial Intelligence (AI) with Machine Learning. It is hoped to provide colleagues, clients and society-at-large with enough information to form their own opinion about the revolution caused by machines that are enabled to learn and adapt.

A vital understanding is that while the traditional approach to AI with expert systems threatened humanity by eliminating the human from the system, the new approach of AI with Machine Learning aims to ‘make people smarter’ and ‘produce something that neither can produce alone’.

This was a wise comment by MIT’s Professor Patrick Winston, considered one of the fathers of AI, and is further proved by NVIDIA in *The Deep Learning Revolution*, (2016) where AI with Machine Learning is demonstrated to enable the blind to read or recognise a familiar face for the first time.

In a nutshell, the key challenge faced by the implementation of AI with Machine Learning is that it requires a new approach to data, a new mind-set, and a new *forma mentis* right across the industry – from data scientists to integrators, leaders and educators.

THE NEW APPROACH TO AI

Whereas the approach to AI with Machine Learning existed long before, the major breakthroughs occurred in 2012, fuelled by the exponential growth in computing power. Success in creating such artificial intelligence grew from enabling machines to learn from real-world data and the creation of their own adaptable models. This stood in contrast to when humans instructed machines to execute preconceived models, the approach often referred to as AI with Expert Systems.

Super power now comes in tiny packages. The first supercomputers to break the PetaFlops (floating-point operations per second) barrier were in 2008 and filled a warehouse, NVIDIA’S ‘supercomputer in a box’ in 2017 achieved the same performance. Processing power reserved for governments and very few companies may soon be available to consumers.

FULLY LEARNT MACHINES

In addition to establishing supercomputer power under the desk, the new approach is far more data hungry than static programming of the past. AI with Machine Learning needs quantity and diversity of data provided quickly enough to the machine so that it can adapt its model and produce more appropriate (cultural and/or environmental), more timely, and more scalable or repeatable solutions.

An example is Autodesk’s *Design Graph*, a large and growing database of three-

dimensional representations of mechanical parts that enables designers to 'shape search' in the same way as i.e. a Google text and image search. Such quick access to alternative shapes aids the designer to improve the provenance, cost and performance of any part.

MACHINES LEARNING AT THE EDGE

The transition to software that is fully learned, as opposed to the traditional scripted method of creating software, has not happened overnight. The hybrid software solution has been at the forefront – where the deterministic portion of the practice is scripted and the non-deterministic portion is learned.

The DNA of Making, the construction cable robot builds towers of timber components, demonstrates this hybrid process. Where most of the code uses a traditional set of instructions for the 'fork' to pick-up and position the components, one portion of the code is learned. Every time the 'fork' picks up a component, the robot shows it to a Kinect camera that confirms how the piece is oriented on the 'fork' and then corrects the position of the piece accordingly. This real-time feedback loop ensures that all the pieces deposited to form the tower are correctly orientated.

TODAY'S CONSUMER SOLUTIONS

The accuracy of prediction is not yet acceptable for certain applications, such as the autonomous vehicle. But the path to a

near-perfect prediction is clear, focussed on amassing tremendous quantities of data, in essence 'all' data. Google search is very accurate simply because it has amassed billions of text searches. The next frontier is doing the same with video, shapes, sketches, even spoken language (so called Natural Language Processing). Companies that operate in this field seem to have favoured a crowdsourcing model in addition to having created large collections of these new formats.

The collection of data has already happened in homes. The Amazon *Echo's* incorporated microphone and speaker are used to interface with Amazon's Alexa AI Assistant. Now shopping is as easy as hollering out your desires, and a direct purchase is authorised through the Amazon shopping database. Google *Home* is similar with voice access to the Google search database. With the need to directly interact with a computer removed, voice access broadens the pool of users of the service. Technology adverse grandparents are now part of the technology revolution. All these solutions in turn enrich the companies' data to enable better prediction.

ROBUST YET FRAGILE

AI with Machine Learning's breath-taking advancements are not without their own vulnerabilities, and early applications have exposed areas of concern. *Forgery in the 21st Century* is a reminder that computers' visual perception can be confused. Or as the press describe, 'an aural perception confusion'. A

television news station broadcasted a story of an unsupervised 6-year-old girl who ordered an expensive dollhouse through Amazon *Echo*. This triggered hundreds of *Echo* devices to attempt to order dollhouses when the news anchor simply reported what the girl had said, 'Hey Alexa, (wake word) buy me a dollhouse.'

We know with certainty that the pace of change has accelerated. It is hoped that this exhibition stimulates visitors to expand their exploration of this inevitable transformation, and to consider the impact on their lives and within their professions.

.....
Alvise Simondetti, curator of *Artificial Intelligence – enabling machines to learn* and an Associate for Arup's Foresight, Research + Innovation team.

THE NATURAL LOGIC OF ARTIFICIAL INTELLIGENCE

Most of the things we learn, we learn by doing. We are not unique in that: my parents' cat, when I was a child, was an urban dweller--born and bred in an apartment in town, he was familiar with underfloor heating, but he had never in his life experienced a strong and direct source of flameless heat. The first time he did, in a country house, he badly burnt his whiskers by getting too close to a monumental, German-made cast iron stove. He did not like that, and since that unhappy experience, he learned to enjoy the radiant heat of the stove (which he cherished) from a safe distance. Not unlike cats, we learn all kind of things by trial and error. When something hurts, or does not work, we take notice, and we don't do it again.

Unlike most cats, however, we can also save plenty of time by being told in advance how some things may play out. This is why we listen to our elders, go to school, and read books. Pre-industrial artisans at the end of the Middle Ages and in early modern Europe had to go through a laborious, rigidly regulated educational and training process, marked by exams at all steps, before they could set up shop on their own: apprentices first had to qualify to become journeymen, then take the most impervious tests to become masters in a trade, art, or craft. At every stage, all trainees were made privy to some core technical know-how, so that all members of the same guild would produce reliable objects of standard, comparable quality, at very similar costs. Yet a few craftsmen always stood out, by making better or cheaper stuff. This they did by tweaking and twisting, within limits, the

methods they had been taught: by taking risks and trying something new. Artisans can try and copy from rivals, but the best--the innovators--always learn from their own, solitary trials. You make a chair, and if it breaks you make another, and then another, until you make one that won't break. If you are smart, you can also intuit and learn some "transferable skills" in the process, so the next time you won't have to restart it all from scratch.

Trial and error is a laborious and expensive procedure, because making and breaking stuff takes time and money. Which is why, over time, scientists came up with some shortcuts. Indeed, that was the main achievement--if not the main purpose--of the modern scientific method: from the comparison and selection, generalization and abstraction of the results of many trials, scientists formulate causal laws (generally in the format of algebraic expressions) that distil in a few lines of clean mathematical script the results of many experiments, and allow scientists to predict the result of similar phenomena that may occur in the future in similar conditions.

Built on such premises, statics and the mechanics of materials are among the most successful of modern sciences. Structural designers use the laws and formulas of elasticity to calculate and predict the mechanical resistance of very complex structures, like the wings of an airplane, or the Eiffel Tower, before they are built, and in most standard cases they do that using numbers, not by making and breaking

physical models in a workshop every time anew. That's because those numbers--those laws--condense in simple mathematical notations the lore acquired through countless experiments performed over time. And as crunching numbers is cheaper than making and breaking stuff, and provides more reliable results, over time we came to trust engineers more than artisans, and number-based engineering replaced empirical craft as the driving and dominant technical logic of the modern industrial world.

Then computers came. At the beginning we thought that computers were just, as the name still suggests, calculating devices--very fast abacuses that would speed up all our traditionally slow and error-prone number-based calculations. To the contrary, it is more and more evident today that computers can produce better and more useful results if we let computers solve problems following their own logic and methods, instead of replicating ours. For computers, oddly, do not think like the engineers that designed them; they think more like the good, artisans of old--those very artisans that modern engineers have replaced, and relegated to the dustbin of technical history. Today, using computers, we can make and break on the screen--in simulations--in a few minutes more chairs than a traditional artisan would have made and broken in a lifetime; and if we are smart we can more easily intuit or learn something in this process. In fact, we may not even need to do so, because computers are proving increasingly smarter than us in learning by their own trials and errors.

Computer-based simulations are already, in most cases, perfectly reliable--and yes, they are obtained, mostly, using traditional, calculus-based or discrete mathematical tools. But computers can churn out so many of them, changing as many parameters as needed, so fast, and at so little cost, that one can easily imagine that, to the limit, computers can offer an almost infinite number of solutions for each stated problem. Among so many options, inevitably at some point one or two will show up that will be good enough to solve the matter at hand. Thus computer-based, simulated trial and error becomes a perfectly viable and effective problem-solving strategy, and computational heuristic should be seen today as a fully fledged post-scientific method: the core method of artificial intelligence. That means, for example, that no engineer needs to calculate the mechanical resistance of a chair any more, because computers can just make and break as many chairs in simulation as needed, until they find one that won't break. That's not unlike what a traditional artisan would have done: but computers can now do that so much faster that they can beat, by massive trial and error--by brute force--both the artisan's intuitions, and the engineer's demonstrations.

However, picking and choosing among so many chairs--where many may have the same mechanical resistance, and where mechanical resistance may not even be the only design requirement--would still take time. That's why we can now teach computers to compare results themselves,

and come up with smaller and smaller rosters of winners. We call that optimization, and the principles of the process have been known since the 1970s, when John Holland famously compared the optimization of mathematical algorithms to Darwin's theory of evolution by random variations and natural selection. John Holland thought that we should breed algorithms the way we breed horses: by force-mating the strongest. Since then, the science of genetic, evolutionary algorithms (but which should be more pertinently called eugenic algorithms) has largely proven its efficacy; most software of structural optimization used in engineering today proceeds by trying a huge number of solutions at random, then picking the parameters that appear to yield better results and dropping (killing) those that don't, then again and again, *ad libitum atque ad infinitum*--or, in fact, until someone is pleased with the results, or runs out of time, and pulls the plug. There is no guarantee that any lead found this way may go anywhere, but as computers can keep trying forever, that's irrelevant. Indeed, having smart intuitions or trying to orient this massively random process in any "intelligent" way is unwarranted, unnecessary, and may even be counterproductive.

And that is finally the main difference between the way computers solve problems and the way we do: our science predicts events through laws of causation; the same formula that allows us to calculate the resistance of a cantilever also offers us an explanation--or at least a rational

interpretation--of how the statics of a cantilever works. Computers don't do that, because computers are not in the business of making sense of the world. Why will this extraordinarily complex, indescribably messy structure we see on the screen stand up, and the 20,000 very similar ones just tried and discarded in simulation, won't? Who knows: nobody knows that--least of all, its designers. But we know it will stand up, which is why we can build it. The first, still inchoate applications of artificial intelligence in technology and in the arts have already produced visual and formal results than many find weird, alien, or hostile. And rightly so, as these unusual shapes and forms are the outward and visible sign of a technical logic we may master and unleash, but we can neither replicate, emulate, nor even comprehend with our mind; and the feelings of "alienation," which originally, in Marx's critique of the Industrial Revolution, meant the industrial separation of the hands of the makers from the tools of production, may just as well be applied today to the discomfort we already feel due to the ongoing postindustrial separation of the minds of the thinkers from the tools of computation.

.....
Mario Carpo, the Reyner Banham Professor of Architectural History and Theory, the Bartlett, University College London. He is the author of *The Second Digital Turn: Design Beyond Intelligence*, published by the MIT Press, and other books.

WHAT IMAGES OF CATS AND DOGS CAN TELL US ABOUT BUILDING DESIGN

Everything you can do has been done before, or at least something very similar has. Even the most creative and ground-breaking idea builds on thousands of previous ideas, and could not have been generated without them. It may be incredibly valuable, it may change our lives, but it would not be possible without the 100,000 years of good ideas that went before it.

The fundamental idea behind Machine Learning is to exploit the information contained in these past experiments, and build upon it. It aims to devise methods capable of extracting value from previous experience, of learning from data. Its methods are able to generalise from previous experience in order to make sensible decisions about new, previously unseen, situations.

In traditional approaches to automating a process, humans try to develop a set of rules that embody their understanding of the way a decision should be made. The computer then applies these rules, which represent a person's synthesis of their experience. This model has served well, but fails when humans are unable to synthesise their experience, or unable to communicate the result.

Distinguishing images of cats from images of dogs has been a key problem in Machine Learning. It is not of particular practical importance, but success would indicate the ability to synthesise the distinction between

different classes of complex input (images) into a decision rule that could be applied to previously unseen data (new images of cats or dogs).

The reason that distinguishing images of cats from images of dogs is interesting is that a human can't tell a computer how to do it. We can all do it, but we can't explain how it is done. Having a human construct a set of rules therefore doesn't work.

The Machine Learning approach to the problem is instead to show the machine a suitable set of examples and to ask it to identify the distinction itself. From a suitable set of images of cats and dogs, for example, a Machine Learning method might learn to make the distinction just by observation. The rule it comes up with for making the distinction is unlikely to be interpreted by a human, but is likely to be more accurate than one.

The threshold point at which Machine Learning methods regularly outperformed humans, was sometime in about 2015. This applies only to relatively straightforward recognition, detection, and regression tasks, but that covers an incredibly wide field of applications. Humans still outperform Machine Learning at tasks that require high-level comprehension or interpretation, and will for a long time to come.

One of the primary competitions in Machine Learning is the ImageNet Recognition Challenge. Each year the organisers pose

a test devised to evaluate and extend the current state of art in extracting semantic information from images. The test is based on the ImageNet dataset that comprised one million images, each of which is labelled as depicting one of a thousand classes of objects.

To succeed at the ImageNet Recognition Challenge requires learning the appearance of each class from the set of (approximately one thousand) exemplars. Success is measured as the percentage of test images correctly labelled. The Machine Learning approach first outperformed humans at this task when originated in 2015, and the gap has widened since.

Machine Learning doesn't outperform human tasks requiring insight, reasoning, or expertise. It is, however, far better at dealing with large volumes of data, and it's much less likely to arrive at an erroneous solution due to prejudice, laziness, or preconceived ideas.

The opportunity Machine Learning offers is not to automate the jobs we already have, but to find the new things we can do that would previously have been impossible. It can enable advances in the level of flexibility and customisation of building control, it might improve environmental efficiency significantly, and allow buildings to be designed without the constraints implied by current control systems.

Machine Learning is changing the way society shops, communicates, travels, and works. It will certainly change the way that buildings are designed, built and controlled.

.....
Professor Anton van den Hengel, Director of the Australian Centre for Visual Technologies, School of Computer Science, University of Adelaide.

ARTIFICIAL INTELLIGENCE FOR DESIGN – 3 GENERATIONS AHEAD

Artificial Intelligence (AI) technology is emerging everywhere whether you know it or not, be it auto-correcting your text messages or helping you find directions to a restaurant. But what about design software? How might we interact with the software tools to deliver solutions to the problems of the future. It is the belief that AI will infiltrate and eventually completely transform software tools in three distinct generations of sophistication.

INTELLIGENT TOOLS

In the first generation, there is no marked changes to the traditional design software user interface, but under the covers is the emergence of Intelligent Tools. Tools that solve traditional problems such as catalogue management, standards checking, symbol selection, auto-dimensioning.

However, these tools are not yesterday's implementations using inflexible hard-coded algorithms, rather software based on Machine Learning that is constantly learning better representations of the problem and coming up with better solutions. The tools look the same but they are constantly adapting and getting smarter.

INTELLIGENT COLLABORATORS

In the second generation the user experience changes as the software begins to take on slightly more human characteristics. Perhaps the interaction with the software itself becomes more

natural, incorporating voice, gestures, sketching and other very human modalities of expression. This is the generation of software known as the Intelligent Collaborator.

More importantly the software is now aware of the 'context' in which someone is working. It is perceptive to the traditional working environment and it can begin to assist with tasks beyond those conventionally associated with deterministic software. The main tasks the Intelligent Collaborator will take on will be those that distract from the creative process that the human should be best focused on. In other words, all those tasks that are traditionally a waste of time such as finding the right piece of content, producing derivative works, repeating a design aspect over and over, configuring a tool, etc.

This is exciting. Now the human designer is more empowered than ever, able to stay in the flow of design without the distraction of mundane tasks.

TRUSTED COLLABORATORS

The final generation is marked by the near anthropomorphization of the software. In other words, design software that almost feels like working with another human. Just, another 'human' that can consider billions of possibilities and work at the speed of light. At this point the method of interaction with the software is more akin to a designer working with a team of experts to solve a problem.

Just as in an architecture firm where a leading partner might head up a project with a team of designers and engineers, everybody will now have such a team. The process of design becomes one of instructing the team of what to explore and then interpreting and curating the result. The process is just as creative as before except now the intelligent agents will be exploring solutions that might never have been expected, just as an ideal human expert would. This will push the designer into areas of creative consideration that would never have happened in traditional software.

The future of design is brighter than ever and the combinations of humans and AI bring together the best of both. Empathy, creativity, problem expression, synthesis by the humans and the exploration, adaptivity, productivity and breadth of AI systems. So how will you embrace this change?

.....
Mike Haley, Senior Director of Machine Intelligence at Autodesk, Inc.

DESIGN 2017AD

THE MODEL IS THE DATA

Evidence-based design has been the pièce de résistance for many colleagues over the past decade. This contemporary approach to multidisciplinary design relies more and more on measured data rather than on simplified models (the good old “rule of thumb” some may say). The complex interdependence of variables which is experienced in today’s professional activities requires a change in approach to design.

Using “the worst case scenario” hypothesis served our predecessors well as it allowed the definition of solid and resilient solutions to problems. However, these solutions were over designed, over engineered, and inflated by hefty safety factors.

The growing need to optimise the use of available resources has given rise to a new requirement, which is to reduce such safety factors and to look at the cost of designs from a financial, environmental and social perspective. It is no longer acceptable to design something that is twice as strong as needed, as this is seen as a waste of resources.

This reduced margin for error has increased the need to rely on actual data, on benchmarks, in order to minimise risks. As such, designing is no longer only a gesture of pure imagination, but has evolved with a strong component of risk evaluation and mitigation. New techniques have been developed and a shift towards statistics –

Machine Learning as it is generally hyped – has allowed us to manage very large amounts of data to determine the optimal balance between risk mitigation and creativity.

THE VALUE OF DATA

While design has become more dependent on precedents, the value of data has increased exponentially. It is a new “gold rush” where every company is acquiring data in the hope of monetising it within some commercial scheme. Examples include email systems, self-driving cars, taxi and hotel booking services, etc. We don’t need to stretch our imagination very far to know that this will also extend to the building industry: long-term maintenance, asset management and user experience are some of the core elements that would benefit from data-driven design. But data-driven design has an intrinsic weakness, which is in the way that data is mined and processed by the algorithms.

ERROR AND MIMESIS

AlexNet is a mature convolutional neural network with 25 layers, trained on a million images, and capable of recognising one thousand objects. It is proven that *AlexNet* can fail and mis-recognise images. It can happen statistically, but can also be made to happen.

Images, which do not look like anything a human would expect can be fabricated

to misguide *AlexNet*: a pattern that at first sight looks like random noise could in reality be a fabricated image with its only purpose to confuse *AlexNet*. It is, for example, possible to recognise a traffic lamp with a confidence of no less than 100% with one of these patterns. An optimisation algorithm can be used to adjust pixel intensities in order to increase the confidence of *AlexNet* in recognising a traffic light, using the confidence score as an objective function and pixel intensities as variables.

Can this failure of *AlexNet* be used to some commercial advantage? We need to remain vigilant and still retain our designer skills and intuition as they will be our last defence before system failure happens. And as it seems even a state-of-the-art image recognition convolutional neural network can be hacked, what about a simpler network, defining orientation of a building, type of HVAC system, type of structural system or even the cost of land for new developments?

MACHINE LEARNING AND DESIGN OWNERSHIP

Using Machine Learning techniques to design does not require us to set a sequence of actions (basically writing code that accounts for all possible scenarios to operate) but rather only to set goals. The Machine Learning program will adapt and adjust parameters to create a satisfactory outcome independently from the original code. This will happen by learning from precedent data.

The fact of a program capable of adjusting and humans setting just goals will require us to give away creative control to an increased degree. It happened earlier with computational design optimisation, when an algorithm was used to define designs by iterative evolutionary processes and the human component was limited to setting acceptance criteria. Note that these criteria were rather rigid and based on the designer's experience. It is now possible that a Machine Learning program will not even need these criteria as these will be determined statistically based on precedents.

Design ownership will therefore be under discussion: does it belong to the software house, to the program maker, to the user, the designer, or to the actual machine? Are we truly ready for this?

A ROLE FOR THE CONTEMPORARY DESIGNER

If we decide to endorse the transformation of the industry towards what some call "the industrial revolution of the mind" we may need to redefine our position. What will be the role of the designer in such a future scenario? Will we need designers?

If some can successfully argue that creativity can be artificial, that it is no more our exclusive quality, it is clear that a machine can produce fallacious results if fed with the wrong data. Is it our role to just ensure data coherence, ready to be fed to the powerful algorithm?

A new direction, which seems more exciting, but also more fruitful, is to enhance our capabilities and use Machine Learning programs to speed-up iterations of design, test intuitions and look for inspiration, a new-born collaboration as suggested by Garry Kasparov. There seem to be a lot that humans can do in collaboration with machines, far more than what either would be capable of on their own. Designers do not all need to be statisticians and there is still room for creativity and discovery, just silicon enhanced.

.....
Giulio Antonutto, Arup Associate Director of Lighting.

PHOTOGRAPHIC MEMORY

Civil infrastructure forms the backbone to our functional modern society. It is taken for granted that roads will not buckle and tunnels will not collapse. However, to maintain this expectation, there is an assumption that the infrastructure itself will be maintained. Around eighty percent of the lifetime cost of an infrastructure asset accumulates post-construction¹, which accounts for over £8bn of annual expenditure in the UK infrastructure market alone². However, practical operational and maintenance tasks are often overlooked or left to chance, and on occasion this has led to catastrophic consequences, such as the devastating Sasago tunnel collapse.

The fundamental component of infrastructure maintenance is the engineer's visual observation. As a case in point, tunnels house crucial linear infrastructure such as roads, railways and utilities. In the UK, tunnels account for over 3,000 km of repeatable structures, some of which have been around since the 19th Century. During inspection, the engineering judgement of the structures condition is paramount. Much of this entails decades of repetitive and subjective observation of extensive, similar structures. This not only costs substantial manpower and time, but more importantly these conditions can give rise to inconsistency and significant risk of human error³.

The digital engineering transformation of today has made the knowledge of computational neural networks and deep learning capabilities more accessible,

having created benefits from game-changing that yet are simple innovations. Specifically, the field of computer vision as applied to civil engineering has enabled a move that has gone beyond explicit programming of routines, into the realm of Machine Learning. The coupling of Artificial Intelligence with conventional digital photographic hardware has seen the development of a 'photographic memory'. A memory which not only archives, but which is capable of training itself over time.

Characteristics of the infrastructure's condition from photographic survey data into machine memory are vast, complex, and require precise and intricate algorithm design. Initial training by manual supervision: hand-holding the neural networks that recognised true cracks and true crack growth, which are specific to the material under examination; have progressed to partial supervision; to correct the neural network when it has made a mistake; and finally, self-supervision. Because this eventually has evolved into a fully automated process, inspection campaigns are replicated at a scale of repeatability and objective efficiency that is outside the civil engineer's conventional capacity.

To train and analyse the photographic mind, image capture must be robust. Therefore, work focused on heavily streamlining this data collection phase and to ensure a baseline of high integrity imagery. This enabled more effective use of open source programming, libraries and computational

engines. These technology stacks include Hadoop and OpenCV libraries, Python and Amazon Web Services Machine Learning engines and Deep Learning NVIDIA processing. To facilitate this, work involved collaborations with technologists and academia, such as Toshiba, Smart Vid, University College London and Cambridge University.

Machine Learning and Artificial Intelligence is not a new concept. However, within the everyday practice of civil engineering, it has not yet become a mainstay ingredient. So why now? Increasing emphasis is placed on sustainability, together with whole life cost efficiency and risk reduction. Computationally, there are libraries that are more open and accessible than ever, with cloud computing and data processing systems to match. Therefore, digital capabilities are finally aligned with imagination and for opportunities to become real viable options.

Industry is gradually waking up to this idea. Increasingly, the concept has been tested by asset owners, which has provided innovative challenges and demonstrated projects that have filtered through to the supply chain. Arup, with its visionary approach, is an early adopter of these approaches – in the demonstration of and confirmation of the power of techniques on real projects for clients, namely CERN and Hertfordshire County Council.

Foreseeably, successful automation of Artificial Intelligence to routine engineering tasks has given rise to significant gains, not only in service and safety improvements, but also in the commercial bottom line of cost and productivity. The latter, an additional result of freeing up the engineer for focused time on more complex tasks and efficient application of skills. To harness these capabilities is a move away from an isolated engineering landscape to one that is increasingly translational and interdisciplinary. Furthermore, these advances can transform infrastructure maintenance into a forefront consideration – as part of an infrastructure asset’s design life cycle. However, to take this innovation further into standard practice, a step change in governance and whole life sustainable design thinking is required. So hopefully the industry embraces the opportunity, and it doesn’t become just a photographic snapshot lost in time.

.....
Yung Loo, David Gonzalez and Michael Devriendt, Arup Infrastructure.

¹ National Construction Category Strategy, 2017, Local Government Association

² Construction Industry Forecasts, Construction Products Association

³ Tunnels: inspection, assessment and maintenance, 2009, CIRIA C671



UNLOCKING CONSTRUCTION PRODUCTIVITY THROUGH COMPUTER VISION

The digital revolution has largely skirted the construction industry and the sector continues to be hampered by its long-standing challenges, such as its low-productivity growth. However, today's technological wave powered by deep learning is likely to prove a turning point for the industry. The latest advances in visual recognition technologies have made them effective even for dynamic environments like construction sites, unlocking new cutting edge ways to augment people and processes in the industry. With its industry-leading research team, Disperse strives to lead in evolving these technologies, aiming to help the construction industry maximise its potential in the digital era.

Our computer vision technology understands and analyses construction activity on many levels. With the help of on-site fixed and mobile cameras, our deep learning systems drive actionable insights and report on activities around the site, as well as seamlessly connecting to digital plans, models and other systems. This leads to significantly better control of ongoing and completed construction activities.

Disperse are making construction sites smarter by providing automated insights in real time, facilitating greater control of the project and its various aspects, boosting decision making accuracy and speed, and creating a better experience for all



stakeholders on and off the construction site. Through this convergence of on-site activities and project documentation, manufacturing-like productivity leaps for construction projects are enabled.

Our initial pilot is with a housebuilder that is currently managing the construction of 115 apartments across 3 projects and planning to double the number of apartments under construction in the next year. As our customer's business grew it adopted a more structured and digital approach to its planning but it struggled to do so on-site. The growing complexity and amount of

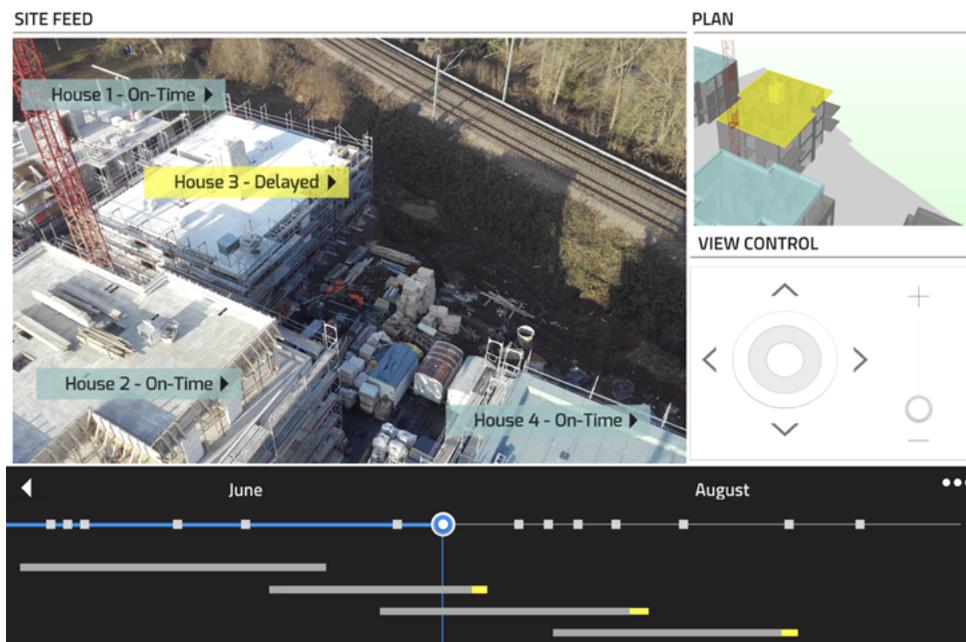
coordination and verification necessary to execute the projects grew exponentially. While the company considered applications to systematically log and validate construction activity, the dynamics of day-to-day operations led to each of those systems being quickly abandoned.

When we first suggested our fully automated approach to construction progress analysis and verification, the construction management was sceptical. However, they quickly embraced the idea of being able to focus on resolving issues while having a system in place

that enables faster and easier detection of those issues. We worked together to find the best solution to the problem and set up a series of fixed on-site cameras which automatically track the construction project against the project plan. The customer was provided with reports, alerts and a custom dashboard. By being able to detect discrepancies between project tasks and all activity and progress on site, our system gives the construction and project management teams a sense of control over their project that they felt they had lost once the projects started to become bigger and more complex. The system has helped the construction and project managers to identify several issues and bottlenecks earlier, impacting on both the bottom-line and ability to deliver on schedule.

construction process significantly closer to the way that modern manufacturing lines operate today, including the associated growth in productivity. As is the case in manufacturing, this automation is not going to replace project and construction managers, but instead augments their ability to oversee construction activity – enabling faster, seamless and resource efficient construction delivery at any scale.

.....
 Felix Neufeld, Co-Founder and CEO of Disperse. Previously he worked in construction, technology integration and digitalisation in London, Germany and Zurich.



Disperse believes computer vision and deep learning technologies have the potential to completely upend the construction industry. Already in the short term, we offer significant advantages to construction companies. Our technology allows construction and project management to automate construction progress monitoring and analysis. Doing so enables construction companies to enhance project coordination, identify issues faster and reduce delivery risks. In the long term, the true value of these systems will be unlocked once they are fully integrated with building information modelling and used to automate the coordination of construction projects. This linkage takes the

PROTECTING ASSET NETWORKS WITH NEURAL NETWORKS – THE AUTOMATED WAY

When Auckland Transport investigated the introduction of light rail to address congestion and growth, Arup and Jacobs Joint Venture (AJJV) were commissioned to create the reference design for the proposed 29km Auckland Light Rail (ALR) route. This included 24 stations, overhead wire (OHW) pole installations, depot and related infrastructure, and road realignment. This major construction project was located in a heavily congested corridor that contained multiple major utilities including gas, water and electricity, which provided essential services to the operation of Auckland City.

As with any urban infrastructure project, underground utility interfaces posed a huge risk to the cost, programme and on-site safety of Auckland Light Rail. As a result, it was critical to understand how the proposed route alignment impacted these utilities. Traditionally, the only acceptable process to identify utility clashes has been for teams of engineers to compare the alignment with horizontal and vertical vertices of individual utilities (clipping). Use of this method of alignment requires large amounts of re-work and in some cases complete reassessment of the utilities. This manual assessment was extremely tedious and costly.

Prior to Auckland Light Rail, was an attempt to automate the clash detection for the Sydney Light Rail project. This involved taking the manual "clippings" section of the process and automating it using a variety of methods. These methods included hierarchical processing of utilities and generating solid

geometry from 2D features which proved to be very successful to reduce the pressure on utility engineers.

Based on this success Auckland Transport approached Arup to provide similar support for the Auckland Light Rail project. One of the first questions the client had was 'What was learnt from other clash detection projects?' What was learnt was an improvement in the 'clipping' section of the process through the shift of the largest elements of the workload to the utility engineers who assessed the individual clashes.

This assessment process provided two classification fields. The first identified the traffic light risk values as Red, Amber or Green (RAG). While the second highlighted the need to divert, replace or protect utilities in the High Level Works Plan (HLWP). The use of the traditional method meant that this could take teams of engineers weeks and if at any time, during the analysis period, the alignment was altered the whole process would be restarted. This was a major bottle neck for the design and optioning process. It raised some hard questions for the team: How can the process be automated? How is waste reduced? How is the engineers process duplicated that would i.e. determine if a pipe should be relocated or a pit protected or too complex to automate and be flagged for manual assessment?

The initial start of the automated process asked engineers how they would go about the process, what were the big-ticket items

they would be able to pick easily? Simple things like material, asset type, age of the asset, etc., was considered and applied to some functional logic (IF, AND, OR, ELSE) in order to start to build an 'engineering logic'. This evolved from 'How can this be improved?' to 'What other information could be sourced to inform this process?'

As the automation progressed, the engineers got more and more excited, and asked for more detailed and in-depth analysis, and manual checks improved the process further and faster. Within 4 weeks an engineering logic to process utility clashes had developed. This functional logic analysis made it possible to automatically assess approximately 80% of the asset features for the Auckland Light Rail project.

During those 4 weeks, having completed my Arup University course in Data Analytics in the Built Environment that taught Machine Learning and Neural Networks, a dataset was recognised that was conducive to an applied Machine Learning algorithm. The analysed data was bigger than the assets that required a value, which made it possible to utilise supervised algorithms.

After extensive research and tests it was decided to utilise a multi-classification Neural Network to determine the location of utility clashes. This algorithm provided the best accuracy of the tested algorithms available.

For clash detection, the algorithm can be retrained to understand the risk and treatment requirements specific to individual Network Owners. Combined functional analysis allowed for further reduction of clashes required for manual project assessment.

The use of algorithms for both RAG and HLWP classifications on the Auckland Light Rail enabled a further amount of reduced assets that required manual assessment. An entirely new automatic system was created to detect clashes and the existing utilities' asset was consolidated into a federated asset information model. A Machine Learning algorithm was then applied to further reduce any manual assessments. Overall 5183 clashes were trimmed to 443, which saved approximately 790 engineering hours.

Arup brought a level of clarity of thought and execution to a hugely complex infrastructure project. The new automated process brought numerous improvements to the project that resulted in a huge reduction in programme length and cost. This has a flow-on effect and led to increased design efficiency with minimised disruption to communities and the environment.

The process has the potential to significantly impact broader infrastructure projects in Australia, New Zealand and

globally. The platform is shared within the Arup Australia and global network. Moving forward, it is anticipated that significant impacts on infrastructure projects, of varying sizes, will be extended to broader infrastructure projects with subsurface interactions, such as tunnelling.

.....
Josh Symonds, Arup Regional Leader of Spatial & Data Engineering, Transport & Resources, Australasia.

The publication is published on the occasion
of the exhibition *Artificial Intelligence – enabling
machines to learn*, curated by Alvis Simondetti.

02 October 2017 – 12 January 2018

Edited by Deborah Smith, Arup Curator.

All rights reserved.

© Arup, the authors, the photographers.

Arup
8 Fitzroy Street
London W1T 4BJ



