GROUP TECHNOLOGY & RESEARCH
POSITION PAPER

MAKING RENEWABLES SMARTER

The benefits, risks, and future of artificial intelligence in solar and wind energy

SAFER, SMARTER, GREENER

MAKING RENEWABLES SMARTER

THE BENEFITS, RISKS, AND FUTURE OF ARTIFICIAL INTELLIGENCE IN SOLAR AND WIND ENERGY

Authors:
Lars Landberg, Group Leader & Senior Principal Researcher, lars.landberg@dnvgl.com
Elizabeth Traiger, Senior Researcher, elizabeth.traiger@dnvgl.com

Acknowledgements:
The Noun Project, https://thenounproject.com
© H. Alberto Gongora, Knut M. Synstad, Meaghan Hendricks, Gregor Cresnar
Simon Fraser University School of Engineering Science/MENRVA
Teddy Seguin, Paris/France

CONTENTS

EXECUTIVE SUMMARY

The use of artificial intelligence in industries continues at an impressive rate – in manufacturing, engineering, healthcare, transportation, finance, telecommunications, services, and energy. Artificial intelligence's ability to use machine learning to analyse historical and new data, make predictions, control physical operations, and make decisions at increasingly higher levels is having an immense impact. Helping artificial intelligence on its way are advances in data processing capacity, the expansion of sensor markets, and a new world of data sources. Technologies using artificial intelligence see falling rates of investment return; and, in general, costs for artificial intelligence are lowering as the ease of its use is increasing.

In the solar and wind industries, there is an enormous amount of data; and renewables has benefited from the fact that it is relatively new and has had sensor technology installed from the beginning. As a result, most of the advances supported by artificial intelligence have been in meteorology, control, and predictive maintenance (and arguably those have been the most useful).

This paper looks at the benefits and risks to the solar and wind industries in future artificial intelligence advances. Solar and wind developers and operators, as well as the financial institutions and other companies that invest in those industries, will see artificial intelligence having immense future benefits in the uses of that data in terms of decision making and planning, condition monitoring, robotics, inspections, certifications, supply chain optimization, and generally increasing efficiency.

Renewables stakeholders also will see that embracing artificial intelligence has risks, in terms of system safety in an automated world, adverse effects on the workforce, human bias in learning systems, and the long-term effects on human expertise and judgement. These risks can be mitigated by taking advantage of the large amount of expertise and data in the solar and wind areas.

> Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.

Nils J. Nilsson, The Quest for Artificial Intelligence: A History of Ideas and Achievements[1]

[1] Nilsson, Nils J. *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge, UK: Cambridge University Press, 2010)

INTRODUCTION

Industries whose processes combine technology, data, and complex decision-making have started to employ artificial intelligence to a high degree: in healthcare, whether it is to predict patient outcomes or analyse data in pharmaceutical development; in finance, whether it is to automate investment decisions or analyse customer behaviour; in oil and gas, whether it is to maximize extraction or manage pipeline logistics; in advanced manufacturing, with smart robots; and more[2].

[2] 'One Hundred Year Study on Artificial Intelligence (AI100)', Stanford University, accessed August 1, 2016, https://ai100.stanford.edu

Artificial intelligence benefits from new applications, new data flows, and new heights of data processing capabilities. In DNV GL's 2017 "Global Opportunity Report," published in partnership with United Nations Global Compact and Sustainia, the global market for artificial intelligence was worth $126.24 billion and was projected to reach a value of $3.06 trillion by the end of 2024—an impressive compound annual growth rate (CAGR) of 36.1% between 2016 and 2024. In terms of revenue generated from the direct and indirect application of artificial intelligence software, research firm Tractica predicts that the market will grow from $1.38 billion in 2016 to $59.75 billion by 2025—a CAGR of 52%[3]. The global market for sensors (which generate data from physical objects, data used as artificial intelligence's feedstock) is expected to be $241 billion by 2022, growing at a CAGR of 11.3% from 2014, says Research and Markets[4]. Even with minor diminishment, Moore's Law—the 1965 observation by Gordon Moore that the number of transistors in a circuit doubles approximately every two years, with the result that processing power for artificial intelligence applications becomes steadily stronger—seems durable.

For the developers, operators, and financial supporters of solar and wind energy, artificial intelligence has benefited their industries immensely, particularly when it comes to weather prediction and control.

But artificial intelligence is in a constant state of advancement, and no one can read the news today without learning about a new development or success in artificial intelligence. This is also true for solar and wind technologies, whose relative newness has led to the placement of more sensors (and therefore more data sources) on its moving parts than older generation technologies. The solar and wind industry will see new benefits from artificial intelligence, among them:

- robots—flying, crawling, swimming, and sailing for remote inspection, with new benefits in maintenance and troubleshooting;
- accelerated due diligence, so that planning and analysis that today might require many human hours and thousands of documents can be reduced by an enormous factor in the future, and even enhanced;
- new efficiencies in supply chain optimization, such as the delivery of solar and wind components by self-driving trucks and even the automation of renewables construction.

Other uses are on the horizon too, some visible, some invisible. It is also certain that some uses will emerge with disruptive force as they cut costs and improve efficiencies for early adopters. Indeed, artificial intelligence is here to stay, and it will have a deep and fundamental effect on the renewables industry and the way we work.

[3] 'Artificial Intelligence Market Forecasts', Tractica, https://www.tractica.com/research/artificial-intelligence-market-forecasts/2 May 2017

[4] 'Sensor Market by Type, Industry Vertical - Global Opportunity Analysis and Industry Forecast, 2014-2022', Research and Markets, http://www.researchandmarkets.com/research/pxtkhn/sensor_market_by, October 2016

There are clear difficulties in embracing new artificial intelligence applications, including cyber security, safety issues, and the status of the workforce. Trying to gain a comprehensive understanding of the use of artificial intelligence in the renewables industry can seem overwhelming. But the fact is that artificial intelligence is accessible to all. Indeed, in the current OpenSource revolution, much artificial intelligence development software is free. Developers are creating data analysis programs specifically for processing artificial intelligence and machine learning algorithms; and those programs are available under open source licenses to encourage a fast pace of development from a wider community.

That said, it is a demanding effort to apply artificial intelligence appropriately. It is not a one-off or an off-the-shelf endeavour. There are common pitfalls regarding human bias, lack of usable data, and auditable results. And the solar and wind industries can take advantage of significant domain knowledge.

In fact, not using domain knowledge, and relying on one's own judgement, could lead to suboptimal and even incorrect results.

For the solar and wind industries, the benefits of exploring and incorporating artificial intelligence into their various processes can outweigh the risks. The longer businesses wait, however, the faster their future disruption. Solar and wind developers, operators, and investors need to consider how their industries can use artificial intelligence, what the impacts are on the industries in a larger sense, and what decisions those industries need to confront.

LEARNING TECHNIQUES

Artificial intelligence can benefit from dozens of models, algorithms, and learning techniques. For the renewables industry, there are some particularly important tools in the toolbox—and the ways of "teaching" artificial intelligence systems are relatively straightforward.

Fuelling artificial intelligence development is a symbiotic relationship between systems and data: Just like the brain of a child, an artificial intelligence system needs huge amounts of information in order to learn; and companies that sit on huge amounts of information (renewable systems generate a lot of it) usually want to minimize the human effort of analysing it.

A child learns and develops from the environment and examples set. Artificial intelligence systems learn and 'develop' from the data and algorithms used. Since computer systems work according to a prescribed set of rules, the domain expert has the responsibility to provide the right set of parameters in which to operate. Those parameters require a mix of rules, standards, and judgement priority. In renewable operations for artificial intelligence applications, this framework of rules and standards already exists: The operating environment consists of known engineering structures with detailed design specifications. The closed systems are well understood.

To analyse and use data, artificial intelligence employs machine learning—the ability to learn without being programmed—managed in two fundamental ways. With **supervised learning**, when we know the outcome related to an historical data set, we use the data with algorithms to detect patterns or trends for deeper insight into processes and systems.

A simple example is working with historical data from a marketing campaign, looking at responses, money spent, and other campaign drivers, to tailor future campaigns. As applied to the wind industry, historical data can help in reading wear and tear in the gearbox of a wind turbine's nacelle and thereby improve predictive maintenance and life extension. Supervised techniques provide powerful tools for prediction and classification.

With **unsupervised learning**, we do not know an event's outcome but are interested in finding associations or patterns in the data. For example, in cases of financial fraud, we may not know that a transaction is fraudulent until after the event. Rather than attempting to predict which transactions are fraudulent, we might want to use machine learning to identify transactions that are unusual and require further investigation. In the case of a solar panel, using data from image processing can help identify possible anomalies.

Supervised learning techniques

There are five common supervised learning techniques.

GENERALIZED LINEAR MODELS (GLM)
An advanced form of linear regression that supports different probability distributions and link functions, enabling the analyst to model the data more effectively. Enhanced with a grid search, GLM is a hybrid of classical statistics and the most advanced machine learning.

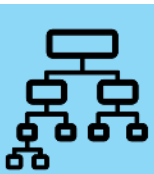

DECISION TREES
A method using a set of rules to split a population of data into progressively smaller segments that are homogeneous with respect to the target variable.

RANDOM FORESTS
A popular ensemble learning method that trains many decision trees and then averages across the trees to develop a prediction. This averaging process produces a more generalizable solution and filters out random noise in the data.

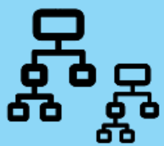

GRADIENT BOOSTING MACHINE (GBM)
A method that produces a prediction model by training a sequence of decision trees, where successive trees adjust for prediction errors in previous trees.

DEEP LEARNING
An approach that models high-level patterns in data as complex multi-layered artificial neural networks. Deep learning is the most general way to model a problem; and it has the potential to solve the most challenging problems in machine learning (see box, "Deep Learning").

Unsupervised learning techniques

Three unsupervised learning techniques are common.

PRINCIPAL COMPONENTS ANALYSIS (OR DIMENSION REDUCTION)
Principal components analysis evaluates a set of raw features and reduces them to informative indices that are independent of one another. As renewable systems capture more data with the addition of more sensors, PCA can help identify the data that provide real informational value for a particular problem.

CLUSTERING
This technique groups objects into clusters that are like one another in various metrics. There are many different clustering algorithms; the most widely used is k-means.

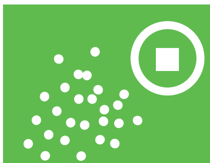

ANOMALY DETECTION
In fields like security and fraud, it is not possible to investigate exhaustively every transaction – we usually want to flag the most unusual transactions in the most systematic way. The supervised technique of deep learning also can be used for anomaly detection.

DEEP LEARNING

Deep learning models are evaluated by measuring how well they predict, treating the architecture itself as a "black box".

Deep learning works well with data that have a large number of discrete values. For example:

- Speech recognition, where a sound may be one of many possible words. In wind energy, analysis of sound anomalies could help identify problems.
- Image recognition, where a particular image belongs to a large class of images. Think of images of solar panels, with or without different kinds of cracks.
- Recommendation engines, where the optimal item to offer can be one of many. Making a move in chess is a popular example. If you were to purchase a solar or wind component, a recommendation engine might remind you to consider purchasing other components for installation (for example), in the same way that Amazon lists products "frequently bought together" when you order something.

Another strength of deep learning is its ability to learn from unlabelled data, which lack a definite "meaning" pertinent to the problem at hand. Common examples of this include untagged images, videos, news articles, tweets, and computer logs. Most data generated in the Internet of Things today are unlabelled. Deep learning can detect fundamental patterns in such data, grouping similar items together or identifying outliers for investigation.

Deep learning has disadvantages, as well. Compared to other machine learning methods, it can produce models that can be difficult to interpret. Such models may have many layers and thousands of nodes. To infer meaning from each layer and node individually is impossible. Also, deep learning is extremely data hungry. In many domains, there will never be enough data for accurate and informative predictions.

REINFORCEMENT LEARNING

Learning to act through trial and error, solely from rewards of punishments.

In reinforcement learning, algorithms, known as software agents, automatically determine the ideal behaviour within a specific context to maximize total performance or cumulative reward. Simple reward feedback, the reinforcement signal, is required for the agent to learn its behaviour. The agent's actions affect the subsequent data it receives. It only received feedback from its own actions. The agent has memory, as past results are used to determine future play. In this way, reinforcement learning differs from other machine learning techniques which assume a Markov property[5] where past performance is forgotten.

[5] In probability theory and statistics, the term Markov property refers to the memoryless property of a stochastic process (Wikipedia)

TRAINING ARTIFICIAL INTELLIGENCE

There are two general ways to approach the training problem. While we can train on a select set of examples, artificial intelligence needs to reason and deduce cases outside of those previously observed. One can feed an artificial intelligence system by either providing a data set (in a whole data approach) or streaming cases as they occur.

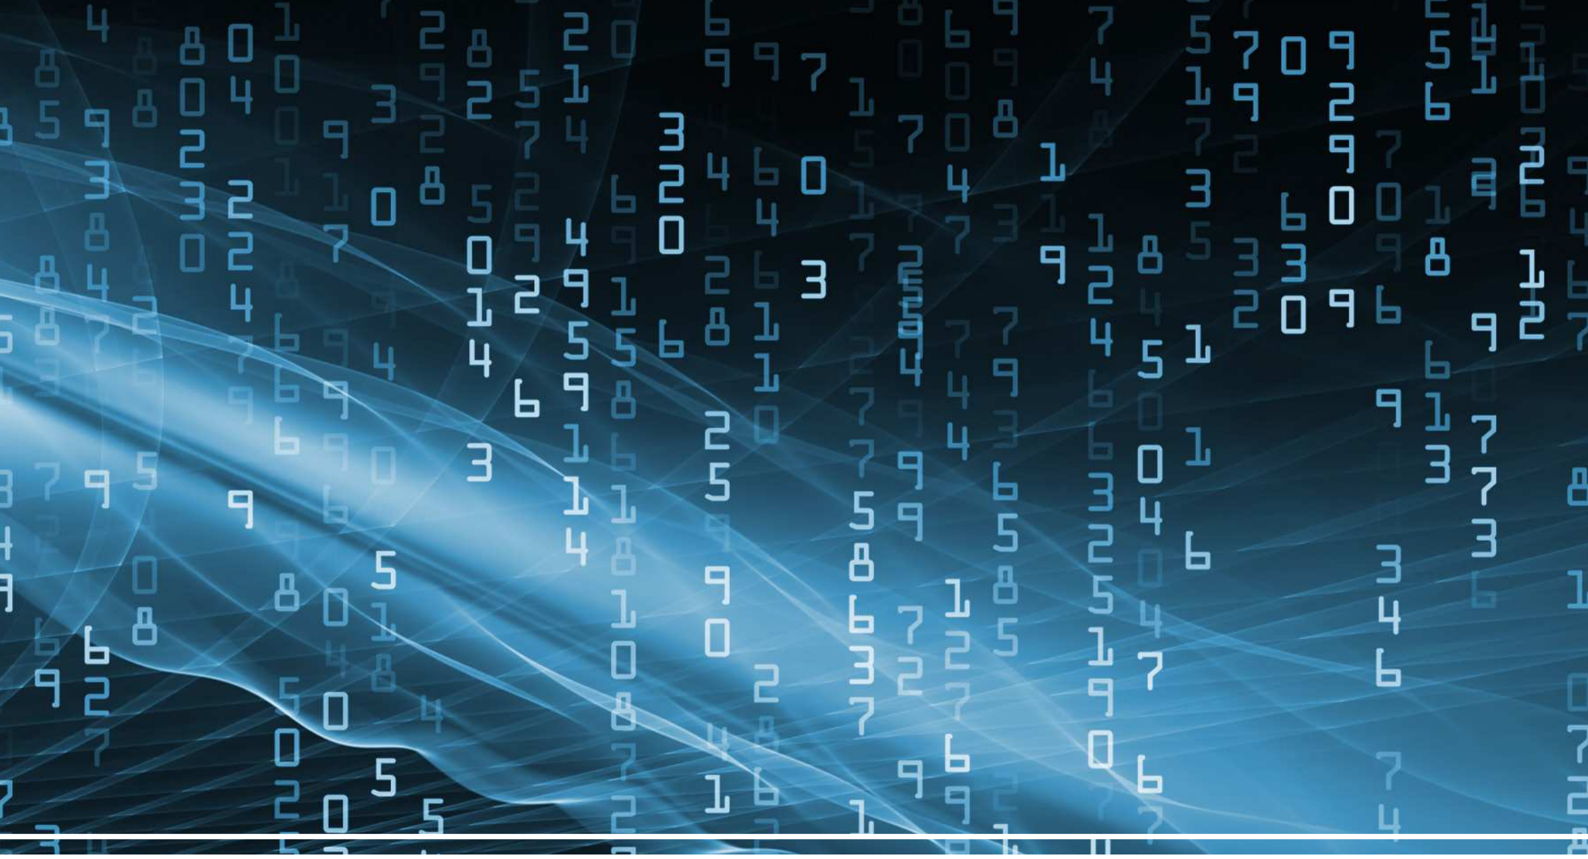

Using whole data

In this approach, the user builds a complex model, collects historical data from many sources (wind turbines, solar panels, weather, end use, etc.), and stores those data on a central server. New data are added to the central repository to refine the model.

Streaming

With streaming, the user builds the base model and exports it to local systems, which update the model with their data and send it back to the central base. The central model changes without direct collection or visibility to the user data (see box, “The Distributed Cognitive System").

THE DISTRIBUTED COGNITIVE SYSTEM

A distributed cognitive system (IBM’s Watson is a popular example) “lives” in the cloud, collecting information every time people or machines use them, everywhere. This means that the more people use them, the better the system becomes at its job.

The whole data approach simplifies feature traceability (where you can track changes) but requires a gatekeeper to store and curate the data. This raises concerns about access, security, and intellectual property. Also, because new data must be reanalysed, updates to the model can be relatively slow. In comparison, the streaming case results in quick model updates. Smaller files with relatively more pertinent information are being shared (though this “snapshot” updating tends to lose feature traceability). Not all situations will have happened, and streaming cases also can deal with rapid evolution of baseline conditions. If the status quo changes, patterns found in the historical data used in the whole data approach can lose meaning for a current application. Streaming can pick up on these behavioural shifts and maintain relevancy; by the same token, streaming can fall prey to temporary trends.

Traditionally, in the solar and wind worlds, information about individual project performance has not been shared except to demonstrate the achievement of regulatory minimum standards. For the adoption of artificial intelligence, those with access to large portfolios—with more data and the ability to stream information—are at an advantage.

ARTIFICIAL INTELLIGENCE APPLICATIONS IN RENEWABLES

A variety of models and learning techniques can provide the support for artificial intelligence in several areas of the solar and wind industries. The combination of artificial intelligence learning with new data and advanced processing capabilities will influence robotics, with impacts on inspection, supply chain, and even construction. Other areas include planning and certification.

Robotics

Any discussion about artificial intelligence in renewables quickly turns to robotics. A prime example is the use of autonomous drones for inspecting wind and PV plant. To automate the process and control the drones will require artificial intelligence—from the system being able to recognize anomalies and even to steadying a drone in mid-air.

The field of robotics is developing rapidly, so this overview should be considered as a snapshot of what is currently available. Suffice it to say that within the renewables industry, robots (those that fly, crawl, drive, sail, and dive) are used for tasks that encompass the three Ds of robotics: dangerous, dirty and dull.

Autonomous robots—which actually would make decisions—will add a new dimension of efficiency and effectiveness. With proper implementation, robots could improve quality and consistency. If the "mantra" in renewable energy is to bring down the cost of energy, artificial intelligence-controlled autonomous robots could be important contributors to this in the future.

Flying drones (or unmanned aerial vehicles) are probably the most developed of the robotics technologies used in wind and solar. Most applications have been for inspection purposes. These drones use multi-rotor technology, with most simple drones having four rotors. Larger, more industrial drones have six or more (see figure 1).

Most drones carry a high-resolution video camera. A drone flying near the blade of a wind turbine, for example, can provide images with resolution high enough to detect cracks, dirt and other anomalies.

Figure 1 - Industrial drone

Moreover, many high-end drones also have the capability to carry infrared/thermal imaging cameras—looking at wind turbines and solar panels in the infrared part of the spectrum can reveal a wealth of new information.

Whatever the camera, many gigabytes of data result from the inspection of even a few panels or turbines. Analysing this data quickly, even while the drone carries out the inspection, requires powerful data analysis tools. Recent developments in the deep learning field has demonstrated that this particular artificial intelligence technique is well suited for analysing images and identifying features[6]. Current processing power makes it possible for deep learning networks, once trained, to carry out real-time analysis of the images. Such a capability will make it possible, for example, to send out a number of drones in a large wind farm and identify the turbines that require repair or need more detailed inspection—all in less than an hour, if required. Indeed, autonomous drones with real-time artificial intelligence-supported analysis will become the primary tool for carrying out effective and efficient inspections of wind turbines and solar panels.

[6] LeCun, Yann, et al., 'Deep Learning', *Nature* magazine (no. 521, 28 May 2015)

Crawling robots
The use of autonomous crawling robots also is under rapid development but has not reached the same commercial stage as flying drones (see figure 2 for an example of an experimental crawling robot). Many non-autonomous robots already are working in the solar and wind industries in a wide range of inspection tasks.

One challenge for crawling robots is to "stick" to a vertical surface—like a wind turbine tower or the surface of a solar panel. Developers are experimenting with gecko-style nanotechnological adhesive materials, suction (using vacuum), and magnetic materials, among other technologies.

An advantage of crawling robots over drones is that they can get close to a structure's surface—and in fact touch it. This opens up possibilities for a new set of technologies, the most important being microwave and ultrasonic transmitters and receivers, which can be used to penetrate into the structure to reveal faults in materials.

With crawling robots, artificial intelligence is used mainly to control them. But, as with flying drones, crawling robots collect large amounts of data, requiring artificial intelligence in the analysis of observations. Two obvious use-cases in renewables will be the autonomous inspection of wind towers and blades and around solar panels.

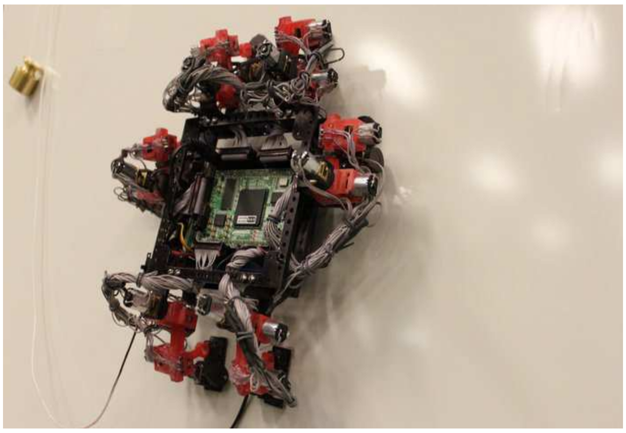

Figure 2 - 'Abigaille' wall-crawler robot (source: Simon Fraser University School of Engineering Science/MENRVA)

Driving robots
Self-driving cars take the headlines, but for industry the arrival of self-driving trucks (as well as cranes and the like) may happen more quickly, due to labour cost savings.

It is possible to imagine an onshore wind/solar farm being built entirely by autonomous robots: the parts of a wind tower and turbine or a solar array are transported from the factory by self-driving lorries, unloaded by another set of robots, attached to the foundations that yet other robots have dug and filled, and pieced together by a final set of robots and drones. Most of these robots are driving robots—or automatic guided vehicles—controlled by artificial intelligence.

Artificial intelligence's main use here pertains to obstacle avoidance and navigation, but higher level systems also are needed for the "orchestration" of the transport as well as actual construction. Still, though autonomous artificial intelligence-controlled robots in the transport and construction industry will continue to develop, the reality of an autonomously built renewable project is probably many years away.

Sailing robots
For the renewables industry, the sailing robot will transport and deliver parts, like the driving robots. And, like the driving robots, sailing robots need to use artificial intelligence to tackle the challenges of avoiding obstacles and navigating. Autonomous ships are still in the experimental stages (see figure 3 for an example). Several research projects (like MUNIN, which is looking at open-sea autonomous shipping[7], and DNV GL's ReVolt system, which is studying a coastal transportation system[8]) have investigated many aspects of autonomy. As with the other autonomous robots, the sensors used for giving the vehicle a picture of the surroundings are an essential part of these investigations.

[7] http://www.unmanned-ship.org/munin/
[8] https://www.dnvgl.com/technology-innovation/revolt/

Figure 3 - ReVolt: a vision for the unmanned, zero emission, short sea ship (source: DNV GL)

Wind farms, again, seem to be the most obvious use-case for sailing robots; and aside from transporting parts of wind turbines across the globe, the biggest role for sailing robots would be in the construction, inspection and maintenance of offshore wind farms.

Again, as above, artificial intelligence will ensure the autonomous behaviour of the individual sailing robots, but also take an essential role in coordinating the efforts of the "army" of robots carrying out the construction.

Diving robots

In the maritime industry, human controlled robots–so-called remotely operated vehicles (ROVs)–have been used extensively for inspections (much like the drones) and for carrying out various simple tasks (like using a robot arm to hammer on a structure as a way to detect corrosion).

The development of autonomous ROVs–that is, autonomous underwater vehicles (AUVs)–is in its infancy, but the direction is clear (see figure 4).

Simply automating the tasks the robots are doing now will provide lots of ideas and challenges for applying artificial intelligence.

AUVs could help in the construction and maintenance of offshore wind farms. One could imagine a situation where, after a severe storm, the onsite drones and AUVs are sent out to inspect the structures above and below the waterline and report back damage or critical issues that need attention.

Figure 4 - OceanOne bimanual human robot (source: Teddy Seguin)

Services

In tech circles, the description of artificial intelligence's impact on services productivity improvements is "10x." It is a somewhat loose measure posited by Marc Andreesen and Ben Horowitz and applies also to network effects, economies of scale, and other techno-economic topics[9]. Still, it is a good rule of thumb for artificial intelligence. In this case, its introduction provides not only incremental benefits, but also dramatic improvements, often with a gain of a factor of 10 in measures such as time spent, pages read, reports written, and so on. A simple way to think about that is to consider that if you spent ten hours on a task yesterday, now it could be done in one.

But 10x-thinking is not really to see whether actually we can accomplish things 10 times faster—the issue is that artificial intelligence's real potential is not about incremental improvements.

In solar and wind construction, enormous amounts of data and documentation inform due diligence, certification, and regulatory reporting. If the construction concept is a novel one, the criteria and standards requiring synthesis and analysis form an even larger pool than that. Several tasks can benefit immensely from assistance from artificial intelligence.

The medical sciences provide a good example of synthesizing and analysing documentation. More than 1.3 million medical journal articles were published in 2016 alone[10]. No doctor can extract the relevant information from that pile, but there are systems that can read (and structure) millions of documents in seconds—IBM's Watson reads more than 800 million pages per second, for example[11], providing the doctor with synthesized, up-to-date information from many sources.

In certification, an artificial intelligence system could provide the certifier with information referring not only to rules and standards directly, but also a synthesis of unstructured information—for example, reoccurring problems with this particular manufacturer, court cases concerning previous designs, news about the company and its technology, and so forth. Even access to the most up-to-date rules and guidelines will increase the accuracy and efficiency of the certifier's work.

The analysis of specific structures like gearboxes and transition pieces also would benefit greatly from assistance from an artificial intelligence system. Pattern recognition to detect potential problems through deep learning is a good example (see box, "Unhindered by Expertise").

UNHINDERED BY EXPERTISE

Stanford researchers have developed a cancer detection machine-learning algorithm that outperforms human epidemiologists[12]. The algorithm identifies tissue slides exhibiting specific lung cancer types with a higher accuracy than review by clinicians. Clinicians are trained to specialize in a subset of exact cancer types. By focusing in on the search for a specific cancer variation, the human clinicians may miss other cancer types appearing in same tissue sample. The machine was not constrained by the knowledge domain of one subfield of medicine; and since it has been trained in the detection of cancer from multiple specialties, it found a significantly higher number of cancerous cells with greater speed and accuracy than by clinicians in the study. The machine also ranked severity of found cancer cells with higher accuracy. By providing processed cancer identification the algorithm gives the clinicians more time for interpretation and treatment strategies. It removes some of the subjectivity from pathology.

---

9 Horowitz, Ben, *The Hard Thing About Hard Things: Building a Business When There Are No Easy Answers* (New York: HarperBusiness, 2014)

10 PubMed.gov

11 "Watson's the name, data's the game," *PC World* online (http://www.pcworld.com/article/3128917/watsons-the-name-datas-the-game.html, 7 October 2016)

12 Kun-Hsing Yu, et al., "Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features," *Nature Communications* (vol. 7, no. 12474, 16 August 2016, https://www.nature.com/articles/ncomms12474).

Artificial intelligence applications in the renewables industry

THE PROBLEMS OF ARTIFICIAL INTELLIGENCE ACCEPTANCE AND IMPACT

Apart from its many opportunities, artificial intelligence also raises a number of problems. Acceptance by potential users is fraught with "human" assumptions. Artificial intelligence in operations means trusting the machine, which has safety concerns. The increased use of artificial intelligence to make processes more effective and efficient—on the factory floor, in planning and decision-making, and throughout business and industry—has a socio-economic impact.

> Most people don't mind changing, but they hate being changed.

Attitudes

In our anecdotal research with colleagues in the renewable industry (though this goes for other industries as well), we have found some basic responses:

- Nearly everyone thinks the advance of artificial intelligence is unstoppable
- Few have actually tried or worked with artificial intelligence-based systems
- Most people can't really imagine what will come and in what form
- Many fear that the "coming of artificial intelligence" will introduce unemployment, either because competitors are able to pick it up faster than us or because the internal efficiencies will result in "redundancies"
- We see elements of a kind of not-in-my-backyard syndrome, in that most people think that their expertise is so specific that an artificial intelligence system is not able to help

Socio-economic impact

Whatever the artificial intelligence application, it is clear that it will have an influence on employment and productivity. The influence on employment most likely will be negative. Forrester Research sees automation—that combination of artificial intelligence, robotics, and processing power—replacing 16% of US jobs by 2025[13], with the equivalent of 9% jobs being created and a resultant net loss of 7% of US jobs by 2025. In a 2016 Citibank report, researchers from the Oxford Martin school concluded that automation puts 35% of jobs in the UK at risk, 47% of US jobs, an average of 57% across the OECD, and perhaps as much as 77% in China[14].

Such labour tumult occurs during every big leap in automation—the Industrial Revolution of the 18th and 19th centuries, where machines replaced large sections of labour, is an oft-cited example. There are always people who say that the introduction of new technology will result in unemployment.

But, considering the Industrial Revolution and more recent technological upheavals, it is clear that new jobs (hitherto unimagined) have sprung up almost as fast as old ones have disappeared. Concerning artificial intelligence, there are also two sides:
One side sees this revolution and thinks of it as fundamentally different and disruptive than others; and the other side considers that the introduction of new technology will create new opportunities and currently unimagined job prospects, as it always has. In that discussion, there are relative views of socio-economic demographics, education and skill levels, and so on.

In the solar and wind industries, artificial intelligence will have an impact on jobs related to planning, operations, and inspections—this has not occurred yet, but we will feel the impact soon. Obviously, it will take some time before the consequences of artificial intelligence introduction on renewables job functions will be clear—but it is time to include those considerations in all artificial intelligence-related development.

Safety[15]

A fundamental issue with artificial intelligence and self-learning systems (like artificial neural networks and deep learning) is that they function as black boxes. On a basic level, we don't know how they work—that is, we can see that once trained they give the right answers in general, but we can't open up the system and see how it came to a particular result. This is different from the way most technical people think, and it requires a certain amount of trust in the system and its results. If the system is properly trained and evaluated this might not be an issue; however, not much imagination is required to see that this can also produce unwanted and even dangerous outcomes and actions.

An example in renewables could be a system trained to control a wind farm. During the training and under normal operation the system performs perfectly.

---

[13] Forrester Research, "Robots, AI Will Replace 7% of US jobs by 2025," press release for report by Le Clair, Craig et al. "The Future Of White-Collar Work: Sharing Your Cubicle With Robots" (released June 22, 2016).

[14] Frey, Carl Benedikt, et al., "Technology at Work v2.0," a Citi GPS: Global Perspectives & Solutions report (Citi and Oxford Martin School/University of Oxford, January 2016)

[15] Amodei, Dario et al. "Concrete Problems in AI Safety" (white paper). Cornell University Library (26 July 2016, arXiv:1606.06565v2).

But, after (for example) an icing event followed by a storm, followed by a grid outage, it behaves unexpectedly and maybe even harmfully—through unanticipated energization, perhaps, or bad transformer communication leading to machine failure. Trying to understand what went wrong can prove difficult due to the system's black-box nature, leaving the operator with general doubt about the system's abilities and also requiring a retraining of the entire system.

Aviation expert Earl Wiener, coined what is known as Wiener's Laws of aviation and human error—they can well be applied to renewables and are valid for all other fields where the levels of automation are high[16]. One of those laws is this: "Digital devices tune out small errors while creating opportunities for large errors." Tim Harford rephrases it as: "Automation will routinely tidy up ordinary messes, but occasionally create an extraordinary mess."[17]

Auditing

Creating an auditable artificial intelligence system is paramount—and problematic. Natural human bias is unavoidable and works itself into any artificial intelligence system. This stems from unintentional censoring of training sets, restrictions on algorithms used, and unrecognized favouritism inherent to the users. An independent, unbiased review by a knowledgeable, impartial auditor will ensure neutrality and fairness within the system. While seen as black boxes, all artificial intelligence algorithms have a degree of transparency and traceability in their application and deployment. The auditing of assumptions is particularly important, for example for pricing algorithms for renewable energy or, solar-panel degradation rates in investment decisions.

A famous example of the need for auditing is the Black-Scholes algorithm, the basic model for financial markets that contain derivative instruments—and the model often blamed for the U.S. stock crash in 1987. The Black-Scholes formula generally provides an option price close to the observed price, but the model is not infallible (especially in markets with high volatility), and users often tweak it. In any event, such tweaks introduce bias—and the need for an impartial auditor.

Deskilling

As machine intelligence improves, there is also a concern about the effect on our own intellectual skills. Indeed, the value of human prediction skills may decrease because machine prediction will provide a cheaper and better substitute. But in an artificial intelligence world, the value of human judgement and interpretation skills ought to increase. Detailed analysis will be more frequent, and artificial intelligence will make it more convenient. Its systems can indicate a pattern, correlation, or behaviour, but not the mechanisms behind the behaviour—human intellect still must investigate and engineer solutions for that behaviour. Such interpretation will require critical thinking. Further, a greater demand for the application of ethics and formulation of best plan of action will emerge, requiring expert judgement skills.

That said, a potential consequence of artificial intelligence's advance may be that we stop learning and improving and stop maintaining our already required skills (also called deskilling), which will erode our judgement, judgement that could be necessary if/when the autonomous world goes wrong. Gary Klein, a psychologist who specializes in the study of expert and intuitive decision-making, summarizes the problem:

"When the algorithms are making the decisions, people often stop working to get better. The algorithms can make it hard to diagnose reasons for failures. As people become more dependent on algorithms, their judgement may erode, making them depend even more on the algorithms. That process sets up a vicious cycle. People get passive and less vigilant when algorithms make the decisions."[18]

Many initiatives already address these issues. On the regulation side, standards organizations have issued guidelines (see, for example, "BS 8611:2016 Robots and robotic devices, guide to the ethical design and application of robots and robotic systems," developed by the British Standards Institute), and organizations have been set up to investigate issues around artificial intelligence (for example, www.openai.com).

---

[16] Croft, John, "Wiener's Laws," *Aviation Week* magazine (28 July 2013, http://aviationweek.com/blog/wiener-s-laws)
[17] Harford, Tim, Messy: *The Power of Disorder to Transform Our Lives* (New York: Little Brown, 2016)
[18] Klein, Gary A., *Streetlights and Shadows: Searching for the Keys to Adaptive Decision-Making* (MIT Press: Cambridge, 2016)

For solar and wind companies, the impact is also on labour force training and preparation and the extent of redundancies in the system. Who runs the plant if the artificial intelligence system goes down?

An impartial subject-matter expert can help here, as well. As artificial intelligence and plant automation expands in the solar and wind industries, and developers and operators "outsource" decision-making and physical operations (and even judgement and active participation) to artificial intelligence systems, the auditor could play a greater role in risk management and generally keep management abreast of artificial intelligence developments.

THE FUTURE

"As soon as it works, no one calls it artificial intelligence anymore."[19]

"Prediction is very difficult, especially about the future," said Niels Bohr. In the case of artificial intelligence, in whatever industry, it is extremely difficult. In renewable energy, often in unpredicted ways and speeds, technology costs have fallen while efficiency has risen. It seems that it is only months and even sometimes weeks between when something we previously thought only humans could do (the multi-level strategic thinking of chess and Go) is now done better by artificial intelligence. (In an April 14, 2017 *Wall Street Journal* article, chess grandmaster Garry Kasparov, defeated by IBM's Deep Blue artificial intelligence system 20 years ago, said that it was time to embrace artificial intelligence's "liberating" potential.)

There are three levels of artificial intelligence, from the "mundane" to the near-philosophical:[20]

- **Artificial Narrow Intelligence (ANI)**
  Also known as weak artificial intelligence, ANI describes a system that is specialized in one area. This is the kind of artificial intelligence we see today every day in various digital assistants, Go- and chess playing systems, etc.
- **Artificial General Intelligence (AGI)**
  Also known as strong artificial intelligence, full artificial intelligence, or human-level artificial intelligence, AGI describes a system that can perform any intellectual task that a human can. These systems should be able to pass the Turing Test (to determine whether a machine can fool a human into thinking it is it is not a machine) and others.
- **Artificial Super Intelligence**
  An artificial intelligence system that is smarter than the smartest human being in any field.

"Super intelligence" can mean so many different things, from something we can understand, like meeting another human being who is cleverer than we are, to highly intelligent machines that think exponentially faster than we do. The behaviours of such systems are almost by definition impossible for us to predict or comprehend. Indeed, how could you measure it? Moreover, as philosopher Nick Bostrom worries, it might be a basic Darwinian error to create something that is significantly smarter that we are!

In a 2013 survey of artificial intelligence experts, asked to predict the advent of ASI, the median optimistic year (at a 10% likelihood) reported was 2022, the median realistic year (at a 50% likelihood) was 2040, and the median "pessimistic" year (at a 90% likelihood) was 2075.

---

[19] Attributed to John McCarthy (1927-2011), professor of computer science at Dartmouth and Stanford, who coined the phrase "artificial intelligence" in 1955

[20] Adapted from the website http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html

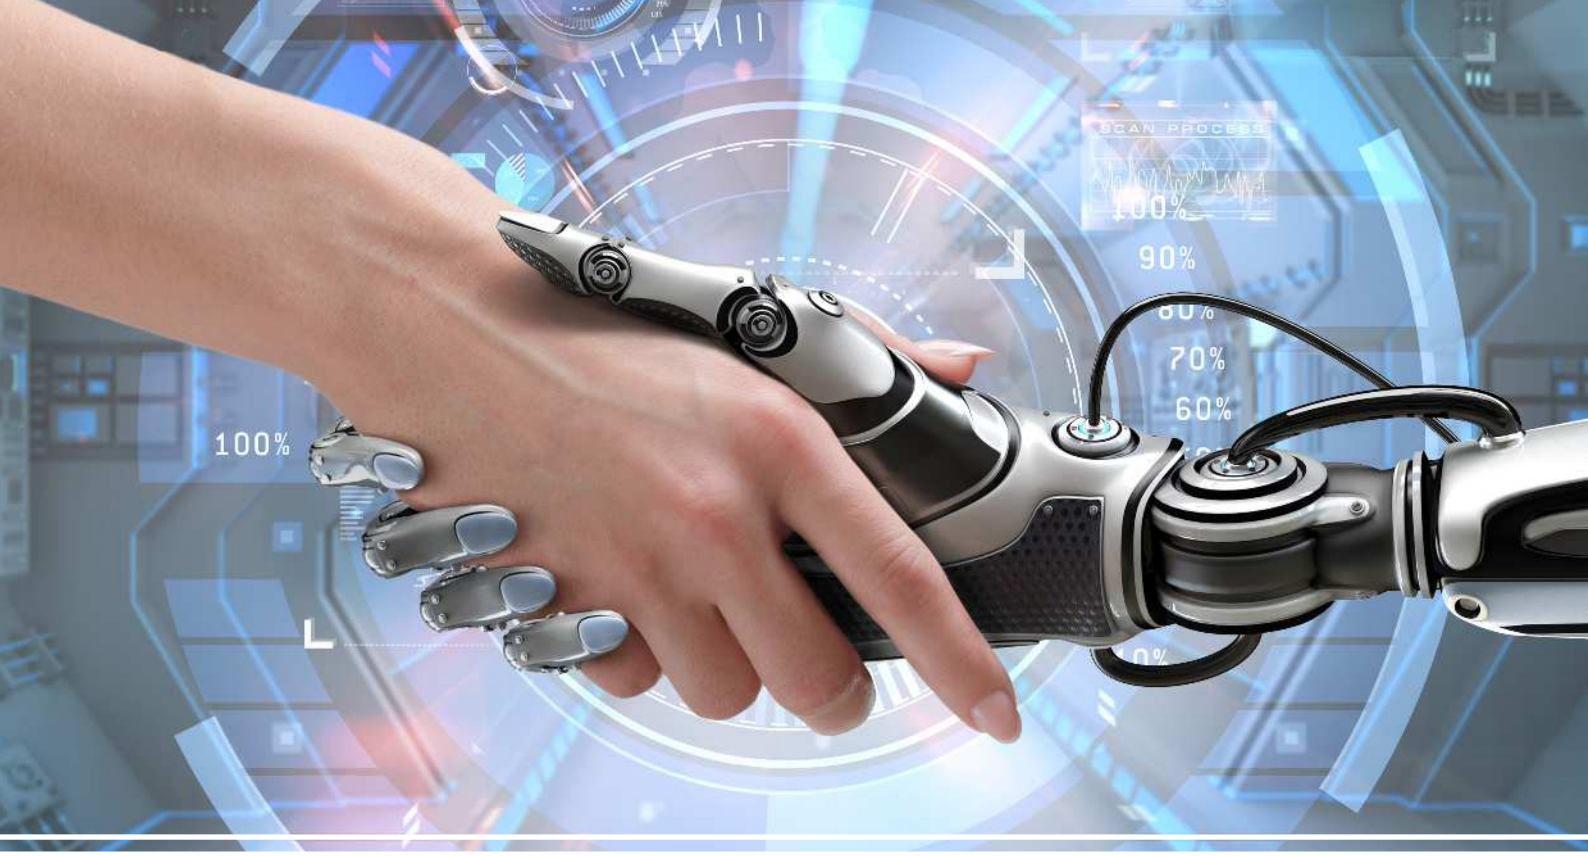

CONCLUSION

The rate of development within artificial intelligence and its application is extremely high. For the solar and wind industries, which already have benefited from artificial intelligence applications in weather forecasting and control, the positive impact (in the relative near-term) will be on inspections and certification, with further-term impacts on supply chain and even transportation and construction. How quickly these impacts occur is an issue of debate—but they surely will occur.

Further, as artificial intelligence enhances the Internet of Things, it will influence wind and solar business models and support rapid commoditization. Such disruptive models may turn out to have the greatest impact on the business. So, embracing artificial intelligence where one can is a prudent move.

That can be difficult, however, due to artificial intelligence's daunting nature. In machine learning, what is good (and safe) training? Many avenues exist, and the tools for the renewables artificial intelligence toolbox are available (see box). As all OEMs and other players use machine learning, there may be more need to review assumptions, which requires the involvement of technical experts ensuring that model bias does not "take over".

For most players in the renewables industry, staying abreast of artificial intelligent advancements, selecting, building and integrating artificial intelligence systems that are stable, progressive, and reliable requires sets of knowledge and data from across projects. From that perspective, you can create objective functions and avoid assumptions and shortcuts (what some call "reward hacking"). To devise the proper training architecture and to get response times and adoption rates right, the key is deep knowledge of the entire domain.

Larger firms offer domain knowledge—that is, larger sets of databases and topic matter expertise—which provides the broad background to use the data sources appropriately. This helps to mitigate the chance of erroneous conclusions based on data alone. Balancing the need for technical thoroughness with safe practices to avoid problems occurring during implementation or operation is crucial. Third-party auditing is also important. In closing, the most important aspect of this "new" technology is to implement adoption in all parts of the renewables' value chain.

OPEN SOURCE SOFTWARE

- Apache Singa, a general distributed deep learning platform
- Amazon Machine Learning, cloud based service with GUI tools and wizards for standard analysis
- Azure ML Studio, cloud based tools as well as ability to run more complex models
- Caffe, a deep learning framework optimized for image processing
- Deeplearning4j, an open-source, distributed deep learning framework written for the JVM
- H2O, tools to apply ML via GUI or language of choice
- Keras, deep learning neural networks which can run on top of Deeplearning4j, Tensorflow or Theano
- MLlib, Apache Spark scalable machine learning library
- OpenNN, a comprehensive C++ library implementing neural networks
- Scikit-Learn, a collection of machine learning and data manipulation libraries for Pytho
- R, many packages available for machine and deep learning from the statistical community
- TensorFlow, Google's open-source software library for machine learning
- Theano, a Python based machine learning library
- Torch, a CPU efficient ML library for use in computer vision and image processing
- Weka, a collection of machine learning algorithms for data mining tasks

SAFER, SMARTER, GREENER

DNV GL - Energy
Utrechtseweg 310-B50
6812 AR Arnhem
The Netherlands
Tel: +31 26 356 9111
Email: contact.energy@dnvgl.com
www.dnvgl.com

DNV GL
DNV GL is a global quality assurance and risk management company. Driven by our purpose of safeguarding life, property and the environment, we enable our customers to advance the safety and sustainability of their business. We provide classification, technical assurance, software and independent expert advisory services to the maritime, oil & gas, power and renewables industries. We also provide certification and supply chain services to customers across a wide range of industries. Operating in more than 100 countries, our experts are dedicated to helping customers make the world safer, smarter and greener.

In the power and renewables industry
DNV GL delivers world-renowned testing and advisory services to the energy value chain including renewables and energy efficiency. Our expertise spans onshore and offshore wind power, solar, conventional generation, transmission and distribution, smart grids, and sustainable energy use, as well as energy markets and regulations. Our experts support customers around the globe in delivering a safe, reliable, efficient, and sustainable energy supply.

Group Technology & Research
The objective of strategic research is through new knowledge and services to enable long term innovation and business growth in support of the overall strategy of DNV GL. Such research is carried out in selected areas that are believed to be of particular significance for DNV GL in the future. A position paper from Group Technology & Research is intended to highlight findings from our research programmes.