WHITE PAPER

PRODUCT∞ION INTELLIGENCE

FLANDERS MAKE

DRIVING INNOVATION IN MANUFACTURING

FLANDERS MAKE'S VISION ON AI FOR MANUFACTURING

Flanders Make believes in the value of artificial intelligence – strengthened by domain knowledge - to enhance products and production systems with additional decision supporting and/or developing features to make them smarter. However, safety, robustness, reliability, explainability and predictability are very important requirements of any product or production system, and meeting these requirements is a challenging endeavour if we want to implement AI-based solutions. In this white paper, we will provide specific solutions and explain how our product∞ion intelligence approach overcomes these challenges to bring real added value to the manufacturing industry and provide a viable steppingstone to a circular economy.

WHY PRODUCT∞ION INTELLIGENCE?

We use the infinity symbol '∞' in this brand new term because we at Flanders Make strongly believe in an approach where data and domain knowledge are used to allow product development, production and after-sales service to interact with and learn from each other "infinitely".
In this way, a company can continuously improve its product or service and at the same time extend its lifespan. Even an individual product, with its unique digital twin, can be further optimised throughout its lifecycle thanks to this approach.
Product∞ion intelligence not only gives the company a competitive advantage. It is also a major step towards a circular economy, where economic development is more balanced with caring for our planet.

SMART PRODUCTS AND PRODUCTION SYSTEMS

Consumers are increasingly expecting personalised products at the same price as mass-produced products. This requires some profound changes for our industry, that has been geared towards mass-produced generic products at lower costs. Turning products and production systems "smart" is a necessity to manufacture these personalised products at the same efficiency as mass production. This means our machines need more and better sensors, which gather more data and are able to optimise all kinds of processes. In order to do so, the product/production system "brain" needs to be able to take complex decisions, which have to be safe, robust, reliable, explainable, predictable and then act accordingly by controlling the involved actuators like an electric machine to adapt its behaviour.

This process can be used for various improvements and optimisations:

- An electric vehicle can maximise the driving range
- An agricultural machine such as a baler can adapt the plunger movement depending on the quality of the crops
- Historical data of changeover times between production runs can be used as a solid indicator how long a certain changeover can take, making it possible to sort production runs in a way that minimises changeover times between them.
- An assembly system can effectively inform and advise its operators to reduce rework
- A manufacturing system can reconfigure its operation to customer needs and changing value chains.
- ...

In the past, many of these "smart" systems were the result of programmed rules that engineers added to products and production systems. These rules applied to certain situations and could only be used if those specific situations arose again. These days, artificial intelligence (see inserted text box) is increasingly finding its way to industrial applications. With AI, we can enable machines to gather data, learn from those data and reason if a certain context or situation that occurred in the past is also applicable to a situation it can perceive now. This allows a machine to handle problem solving in a very broad range of situations, whereas it would be near impossible to foresee all these situations and manually teach them to the machine.

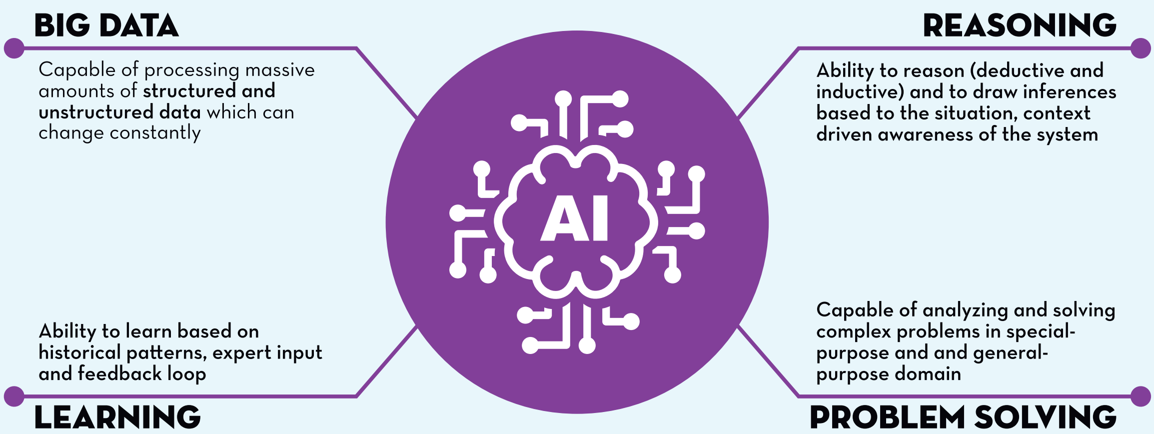

Adapted from https://www2.deloitte.com/nl/nl/pages/data-analytics/articles/part-1-artificial-intelligence-defined.html

Within the Flanders Make industrial network, the technology definitely has been gaining traction over recent years. In our recent biennial Industry 4.0 report, which gathers insights from various companies, AI and big data have seen a rise to the top spot of technologies that has a lot of potential for the near future. As much as 8 out of 10 companies are already gathering data and believe in the potential of AI for their business.

DOWNLOAD
our Industry 4.0 report

ARTIFICIAL INTELLIGENCE

Flanders Make adopts the definition of AI as stated by the European Commission. An AI-system includes four specific means to (1) perceive its environment through data acquisition, interpret the collected structured or unstructured data; (2) reason on the knowledge, or process the information, derived from the data; (3) decide on the best action to achieve a given complex goal; (4) act (in the physical or digital dimension) based upon the decision through actuators, possibly modifying the environment (see Figure).

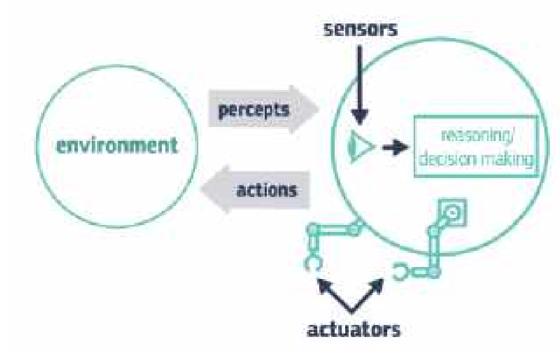

Figure: The EU has defined an AI system clearly, as a system that perceives its environment via sensors, reasons on the obtained information to take decisions, in order to take actions via actuators in its environment.

There are two traditional branches typically associated with AI that can be summarised as follows:

- Machine learning (ML): Refers to collection of methods that can learn patterns from structured and/or unstructured data. Classical machine learning methods include logistic regression, random forests and support vector machines. Recently, deep learning methods (machine learning methods that utilise artificial neural networks with many sequential layers, usually trained on large scale datasets) have been the dominant approach in addressing real-life problems that involve unstructured data. Unsupervised learning (learning patterns in unlabelled datasets) and reinforcement learning (learning and optimising the rewards received from taking decisions) are special forms of machine learning.
- Machine reasoning: Reasoning and decision making are at the very core of any AI system. Knowledge representation (i.e. transforming the data into knowledge) is a central component of machine reasoning, where such representations can be either built by symbolic rules or rules learned from the data. Machine reasoning methods are usually used when potential activities to reach the goal need to be planned and scheduled. As there are many ways to come to the best way to achieve the goal, the Machine reasoning methods search for these activities and optimise amongst possible solutions.

Sustainability and a circular economy

Besides enabling us to manufacture personalised products at the cost of mass production, the need for a more sustainable and circular economy can also greatly benefit from smarter products and production systems. This requires a significant shift in how we cope with products over their lifetime, from design up to the recycling of the product. Recycling of products should lead to a regenerative economy "giving back" the natural resources used in the past. This might become key in Europe, as we have almost no natural resources of our own. In order to cover the full value chain, we would need to develop ecosystems of companies working together to achieve circularity. We will discuss this more in-depth in our text box on end-to-end engineering. But from a single company's perspective, a lot can be done already with the help of smart and connected products and the data that can be captured during the entire lifetime of those products. Not to mention that smarter products and production systems can optimise usage of energy and resources, which not only helps cutting costs, but also offers obvious advantages to the environment.

END-TO-END ENGINEERING: DESIGN-OPERATION CONTINUUM & CROSSING THE PRODUCT-PRODUCTION DOMAIN

The products we discuss in this white paper are often mechatronic systems, used by customers in the manufacturing industries. These products are designed based on specific requirements, produced in dedicated factories and then used by customers who are maintaining them (in-house or by a service provider). When the product, the mechatronic system, is connected to the product-generating company throughout its life cycle and this company has a common and unified policy for data acquisition throughout all contributing departments of the company, a rich data set becomes available. Sharing data between product and production can lead to a better alignment of the product and its production phase. Adding data on maintenance and usage at the end customer creates context about how and in which circumstances the product is used. The latter data generate insights into how the product can be better designed with new features and adapted functionalities based on updated requirements. For example, when a product is designed for a very high operating speed, but the customer is only using it at a moderate speed, the product can be optimised and production costs can be lowered. The advantages for both customer and manufacturer are clear, we can decrease the cost of the product, but there's also an advantage for the environment in decreasing needed resources and energy during production.

In its updated roadmap Flanders Make will further leverage AI and digital twins to realise what we refer to as end-to-end engineering and realise the design-operation continuum.

It means we are working on methods that:

- Allow to bridge the product (design) and production (manufacturing) world.
  - » Product quality and behaviour is determined during manufacturing
  - » Performance of the manufacturing systems is being determined by the product design
- Exchange (data) information across the life cycle of products/production systems:
  - » Information from the design phase, typically models, will be used during the operational phase. To this end the design models will be translated into a digital twin and used to optimise efficiency and predict degradation or even machine failure.
  - » Data obtained from the validation phase of a prototype, by building a digital twin of this prototype, can be used to validate both the design models as well as the prototype. Further, during the validation phase, we will use a digital twin-based approach by applying load conditions that emulate real vehicle component and system behaviour extracted from historic data sets and system models.
  - » Data obtained from the validation (and/or operational phase) will allow to update our design models to better mimic the variety of real hardware behaviour over a wider range of operating conditions such that these digital twin-inspired models can lead to improved designs for future products.

The above use of the digital twin corresponds to several smaller and bigger loops in the infinity concept as shown in the figure[1] below.

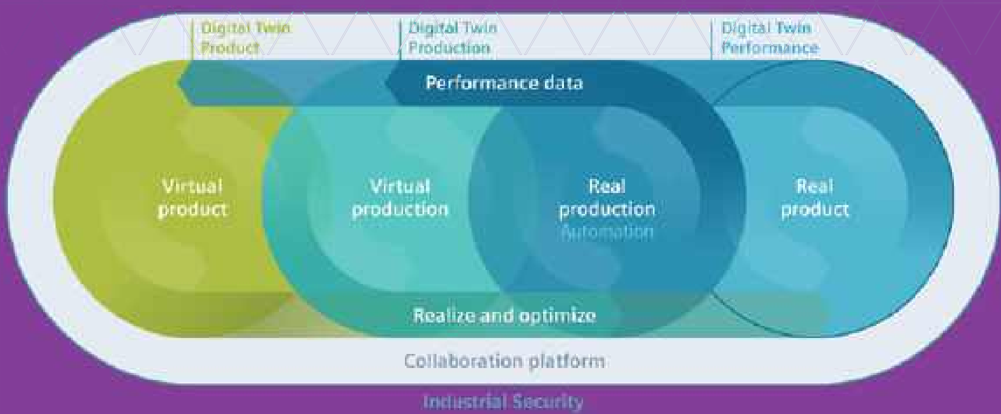

[1] Decision-making loop for the continuous optimization of the production and product. (n.d.). Siemens. https://new.siemens.com/global/en/ markets/automotive-manufacturing/ digital-twin-performance.html

CHALLENGES FOR AI IN THE MANUFACTURING INDUSTRY

Using AI in a consumer-related context, such as an online shop, giving advice based on shown interest, is quite different from using it in an industrial context. Dedicated approaches are needed to create adequate value for the industry. The next table describes the main differences for several aspects. A dedicated approach needs to be developed, as we'll gradually explain further on in this white paper.

Type	Consumer cases	Industrial cases
Available data points, training sets	Many	Few
Probability of an event to occur	High – Medium common	Low
Financial impact of e.g. wrong decisions	Low	High
Operational impact	Low to no casualties, no to low production losses	High, severe casualties, high production losses
Trained case available	Common	Frequently untrained case - corner cases
Data ownership	Available	Company-dependent/critical relationships between customers and suppliers e.g. OEM & TIER-I
Extrapolation of insights	Easy-safe	Difficult – out of trained data set – potentially unsafe

Challenges of working in the physical world

In general, when comparing AI in manufacturing to AI in consumer-oriented applications, probably one of the most important differences is that the potential economic and safety impact of the actions taken in the **physical world** are very different.

- Acting in the real world often means that **physical sensors** need to be added to products and/or production systems. These sensors come at a large cost and have to be very reliable. However, the cost of e.g. camera sensors is decreasing and these sensors are becoming increasingly advanced and more reliable. In consumer-oriented AI, the need for accurate sensors is much lower, the error margins are much wider.
- In a manufacturing context the **reasoning skills** of an AI system often need to go beyond just extracting correlations between data points, the AI system needs to recognise causalities as well. **Correlation** enables machines to recognise objects on pictures, to spot parts of low quality, to cluster load profiles, to identify anomalies, to evaluate skill levels of operators, etc. In a manufacturing context, we need an AI system to understand **causalities** as well. Perceiving a certain effect and knowing what caused it is vital to build a system that can take action to achieve a more desirable effect. What further complicates this, is that the effect of a certain action can manifest itself within milliseconds (e.g. when controlling an electric machine), but it might as well take days or months to notice an effect (for example the remaining useful lifetime of a product).

Challenges related to trustworthiness

Only by capturing, understanding and enabling reasoning on the causalities, mentioned above, trustworthy actions defined and/or supported by AI can be taken in a product or in a production environment.

The trustworthiness of an AI system should be assessed using multiple parameters (see inserted text box on trustworthy AI). **Trust** requires **actions to be safe** for a nearby operator, the actuator and the manufacturing system.
Otherwise an operator might get hurt, a drivetrain might overload or a production line might halt unplanned. Also cybersecurity is to be considered when discussing safety, having unauthorised and unpredictable actors in the systems is always a risk for the system and the humans around it.

Trusted performance is key in a manufacturing context, where a wrong decision might lead to huge economic losses ranging from tens of thousands up to millions of euros due to high scrap levels, lost production, damaged machinery because of overloading or even physical harm to humans.

Eventually, a trustworthy production system augmented by AI should lead to **acceptance** by the humans involved. Humans are still and should remain a key element in the manufacturing sector throughout the whole product-production life cycle engineering.

TRUSTWORTHY AI

The European Commission defined trustworthy AI in its High-Level Expert Group on AI (AI HLEG). Trustworthy AI has three components, which should be met throughout the system's entire life cycle and in all sectors (not just the manufacturing context):

1. It should be lawful, complying with all applicable laws and regulations in view of e.g. GDPR, data security & ownership.
2. It should be ethical, ensuring adherence to ethical principles and values.
3. It should be robust, both from a technical and social perspective, since, even with good intentions, AI systems can cause unintentional harm.

Evidently, the third bullet about robustness is the one we focus on most in the industrial context. Since customer-facing AI often processes private information, they need to do this according to the law (e.g. GDPR laws) and within ethical boundaries. In the industrial context, an AI needs to be robust, above all, in order to be considered trustworthy. Robust means it has to perform well, safe for the machine and human operator, reliable and explainable in terms of actions taken. In an industrial context, solutions should also be set up to be cybersecure as an intruder with bad intentions could cause harm to all those factors.

Humans act as:

- **Manufacturing experts,** when designing products and production systems (design engineers), verifying the quality level, tuning manufacturing settings at different manufacturing control levels and deciding on repair & maintenance activities
- **Operators** in the production systems, operating a machine, executing end-of-line tests, moving objects, etc. The experience and skill level of these operators often makes them experts, even though they will often not be able to fully document this knowledge formally.
- **AI experts,** as more and more companies hire data experts, big data specialists and AI experts, the shortage of these profiles is only growing on the job market.

Challenges in terms of data availability & computational limitations

Finally, there are challenges that relate to the **availability of data as well as the computational resources.**

- In manufacturing, data is both **big & scarce data.** With more and more connected and smart systems, high-frequency data flows across the product-production life cycle become available. A single production system can deliver information about the design in CAD files, information about the creation in its task instructions and information on how it operates in the field through measurements of sensors.
  Even in a recycling phase, more data can be gathered. Therefore, large quantities of data are available for exploitation. However, as quality levels have been increasing and preventive maintenance has been introduced, most of these data is collected from a system during nominal operation. Obviously, an enterprise wants to avoid faulty operation of any system, so in general, there is a lot less data that can offer insights when things go wrong. This makes it more complicated to predict errors in operations.
- **Computational resources are not always present.** AI-based services heavily rely on the computational resources offered by the manufacturing system. The decreasing cost of computational resources is what makes AI much more readily available in more and more sectors. But industrial applications often have very different requirements than consumer oriented AI systems. Constraints like the time to take a decision or the required accuracy can make it impossible to implement a cloud-based solution, while constraints on the needed computational power can make it impossible to implement solutions at the edge. This means it's often a challenge of provisioning a sufficient amount of computational resources to satisfy both the requirements of the AI service and operational constraints of the manufacturing system.

The above-described challenges are not minor, at Flanders Make we are tackling them every day with a well proven strategy that relies on domain knowledge.

ENTER PRODUCT∞ION INTELLIGENCE

Being experts in manufacturing from both a product and production perspective, Flanders Make was thrilled to add artificial intelligence to its set of engineering tools. However, we needed to tackle all of the above-described challenges to make full use of it. At Flanders Make, we strongly believe that the key to solving these issues is actually hidden in the fact that the behaviour of products and production systems are inherently rooted in physics. We refer to this as domain knowledge.

Domain knowledge for trustworthy AI

Domain knowledge of the physics behind these systems is typically in place in a manufacturing environment and has been used/developed over years.

In general, we identify three types of highly valuable domain knowledge:

1. Knowledge included in **static and dynamic representations,** such as physics-based models, CAD files and traditional control loops.
2. Knowledge incorporated in **simulation models and tools** giving access to synthesised/simulated information mainly related to understanding dynamic behaviour, these models incorporate state space models, differential equations, finite element simulations, etc.
3. Knowledge/expertise obtained via **humans,** manufacturing experts or operators that is not easily formalised into a physics based model. Senior experts with years of experience can label images of good/bad product quality, or have developed a gut feeling of how to operate a complex system that is very hard to model. They also offer creative solutions to sometimes complex problems.

This knowledge is increasingly represented in the form of a **digital twin,** which combines it with data coming from simulated models as well as from the physical system. The unique digital representation of the physical asset can be created for each individual physical asset (a product or production system). Digital twins are not the focus of this white paper, however in the text box we do indicate some interesting interactions between AI and digital twins.

DIGITAL TWINS

Reaching an overall definition of digital twins is a complicated matter, since the concept has been extended, misused and derived a lot over the years. However, there is a common understanding that there are three important elements to any digital twin:

- A digital twin relates to a real "asset", often a product or system in our case. We refer to it as the physical twin/asset, e.g. an axial flux machine, a drive train or a vehicle.
- The digital twin and the physical twin exchange information in at least one direction, e.g. data are captured on the physical twin and made available to the digital twin.
- For the intended use (i.e. with that use in mind) of the digital twin, the digital version sufficiently mimics some type of behaviour of this real physical asset.

The last element is key for us as it means there is no such thing as THE digital twin, it means that depending on the intended use an alternative digital twin can be envisioned: there might be a digital twin for condition monitoring, another one for validation and yet another for disassembly, ...

Flanders Make merges domain knowledge in several ways and in a diverse intimacy with AI either by using it during the design of the AI system (e.g. while training the model) and/or the operation of the AI system (e.g. classical controllers running concurrently with AI). More particularly, we:

- include **hybrid models** into the machine learning (e.g. (deep) neural networks, long short term memory, etc.) and reasoning process either by predefined features up to the intrusive introduction of models into neural networks, these hybrid models are a combination of physical and data-driven models.
- combine **classical controllers** with AI techniques such as reinforcement learning, allowing us to either retune the classical controller or even further integrate AI into the controller itself to verify its output.
- use **simulators** in order to simulate both the behaviour of the intelligent system with its actuator as well as the complexity of the environment the system will be exposed to. As such, similators are not only used to generate synthesised data, but also allow AI systems to learn fast and safely from simulations before being deployed and before further learning in a real environment.
- use dedicated **industry relevant setups** where AI routines are combined with domain knowledge to extract the real behaviour of systems.
- use **human experts**' knowledge not only to label static data sets, but also to allow AI systems to learn from complex decision-making of senior operators/designers to assist in the learning process.
- develop **virtual sensors**, since more information can be extracted from existing sensors using more advanced signal processing methods that are now available due to progress in available processing power.

While **upgrading the performance** using these techniques, we have also shown how to increase the trustworthiness of products and production systems in terms of

- **explainability/interpretability** allowing for a more effective interaction and complementary operation with AI experts, as well as operators and manufacturing-experts.
- **safety** of both **the installed manufacturing system and the humans** active in the AI-controlled environment. Having explainable, and thus expected, results is an important factor in the safety of the system.

This is a big leap from self-learning AI systems that make decisions in a black box, offering very little transparency as to why it makes those decisions.
For example, feeding a self-learning system a number of pictures of pets might teach it to make a distinction based on whether it contains a cat or a dog, but it might as well make a distinction based on the background being grass or an indoor environment.
This approach is obviously far from adequate in an industrial context, where an unexplainable result offers little confidence and might have a big impact.

All of the above remains to be complemented by our work in "hard core" manufacturing including e.g. robust and optimal control, novel hardware design of both products and production systems, etc.
We, at Flanders Make, strongly believe that the upgraded performance comes from the co-development of key industrial assets like hardware and operators, as well as intelligence systems.

PRODUCT∞ION INTELLIGENCE FOR ORGANISATIONAL LEARNING

Product∞ion intelligence is how we envision the manufacturing industry overcoming some of the challenges for AI in manufacturing. Further, in order to maximise impact across a company we further propose to embed this into a framework of organisational learning.

Organisational learning

In every organisation, the process of 'organisational learning' (cf. MIT Sloan[2]) is central to growing and thriving. Transferring knowledge within the organisation is an important step in developing knowledge among employees. Creating knowledge, transferring it between the different parts of the organisation and making sure it will be retained are what makes each company more and more specialised in what they do.

The same process happens when dealing with AI-augmented systems. In order to create impact and scalability using AI in your organisation, it is not sufficient to focus only on data, infrastructure and talent. It is additionally and even more important:

1. for critical operations to include and start from domain knowledge to set up product∞ion intelligence and complement these operations with AI.
2. to facilitate organisational learning with AI and allow the stakeholders in the organisation to learn from AI. It is therefore important to not only train the machine, but also to ensure the organisation learns from the AI.

ORGANISATIONAL LEARNING WITH AI

Organisational learning with AI involves a bi-directional learning of humans and machines, not only working together but influencing each other to adapt to contextual changes depending on the severity of the potential impact on the business case. This makes both smarter, more relevant and more effective, for creating a robust, explainable and predictable impact on the applications and business returns. This should be integrated in the strategy for adoption of the AI component of production intelligence as it might otherwise hinder the scalability of results and hence jeopardise the ROI of the use of AI. According to the MIT Sloan study , organisations that embed organisational learning have 3 essential characteristics:

1. They facilitate systematic and continuous learning between humans and machines.
2. They develop multiple ways for humans and machines to interact (e.g. generate recommendations and humans decide to use them, human solutions with AI evaluation of the solution, ...).
3. They change to learn and learn to change.

[2] RANSBOTHAM, S., KHODABANDEH, S., KIRON, D., CANDELON, F., CHU, M., & LAFOUNTAIN, B. (2020, October 20). Expanding AI's Impact With Organizational Learning. MIT Sloan Management Review. https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/

Organisational learning in manufacturing

The translation of the organisational learning concept to the manufacturing industry is depicted in the following figure and a non-exhaustive interpretation of the interaction areas is formulated.

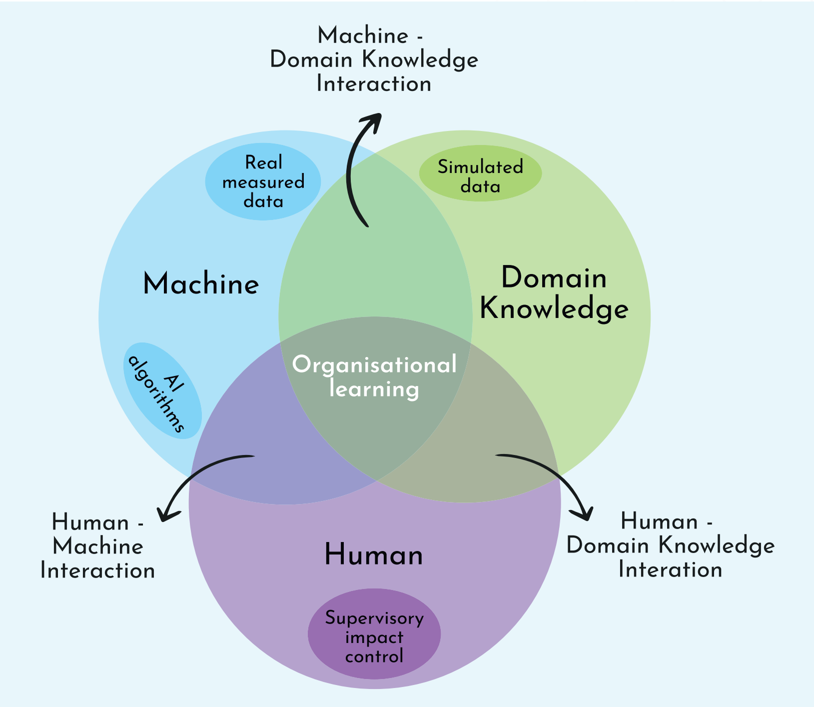

Human-Machine interaction

- Human prepares the training data for the machine and checks automated decision-making for robustness reasons and to avoid major adverse impact
- Machine proposes the actionable insights into the control mechanisms that the human can execute and analyses the human solutions to perform quality checks and further improve the automated decision making

Machine-Domain knowledge interaction

- Domain knowledge is integrated (e.g. hardcoded) in the algorithm for learning and leads to simulated (biased) data for more efficient algorithm training (e.g. only presenting real cases)
- Machine's autonomous/automated decision-making leads to new insights based on corner cases (nearly unrobust action).
  Results of using ML algorithms lead to new domain knowledge (e.g. friction model parameters)

Human-Domain knowledge interaction

- Humans add new information to extend the domain knowledge
- Domain knowledge gives insight to human on behaviour of Cyber-Physical Systems (explainability)

Organisational learning

- Between all these interactions: collaborative multi-directional organisational learning where humans, machines and domain knowledge are all enriched from what they learn on their own and from each other.

CREATING IMPACT WITH PRODUCT∞ION INTELLIGENCE

Our vision is that product∞ion intelligence embedded in a framework of organisational learning will be a key driver to increase **business impact**. This belief stems from the fact that we find product∞ion intelligence on the crossroads of two technologies that have been introduced in the context of Industry 4.0. On the one hand, we have the introduction of big data and AI, motivating enterprises to gather large amounts of data on every process. But often, these efforts fall short of also understanding the correlation between the data and the real-world results. On the other hand, there's the introduction of physics/model based intelligence, simulating how systems perform or should perform in the real world. Product∞ion intelligence is where these roads meet, combining the data that are gathered throughout every process with the knowledge that stems from simulating a process through models. By combining both, we can deliver a **smarter and more trustworthy industrial** AI system.

Often, when AI is integrated in industrial processes, it's used to streamline a specific process. This is when an organisation discovers and builds AI, without necessarily going beyond that step. While the organisation gains from their AI system, it is missing out on much larger gains by not taking the next steps as well. Scaling up AI and eventually performing organisational learning with AI are the steps we consider a part of product∞ion intelligence, the resulting smartness and trustworthiness is what the industry needs.

All of the above enables companies to increase their probability of becoming an industry leader from 2% to 73%. First of all, the basics need to be in place and then they need to be complemented with sound piloting for re-usability and scalability, as well as the cross-domain interaction of machines, humans and domain knowledge extended with AI.

Three steps that lead to a high probability of being a leader in product∞ion intelligence:

1. Defining a sound strategy, creating the necessary infrastructure, ensuring qualitative data access, recruiting and creating talents, and defining and executing first trials for testing and and familiarisation.
2. Building product∞ion intelligence based solutions as a pilot for later re-usability throughout the factory, for multiple products or product functionalities.
3. Ensuring interactions between machine, humans and domain knowledge as part of the realisation of organisational, mutual learning

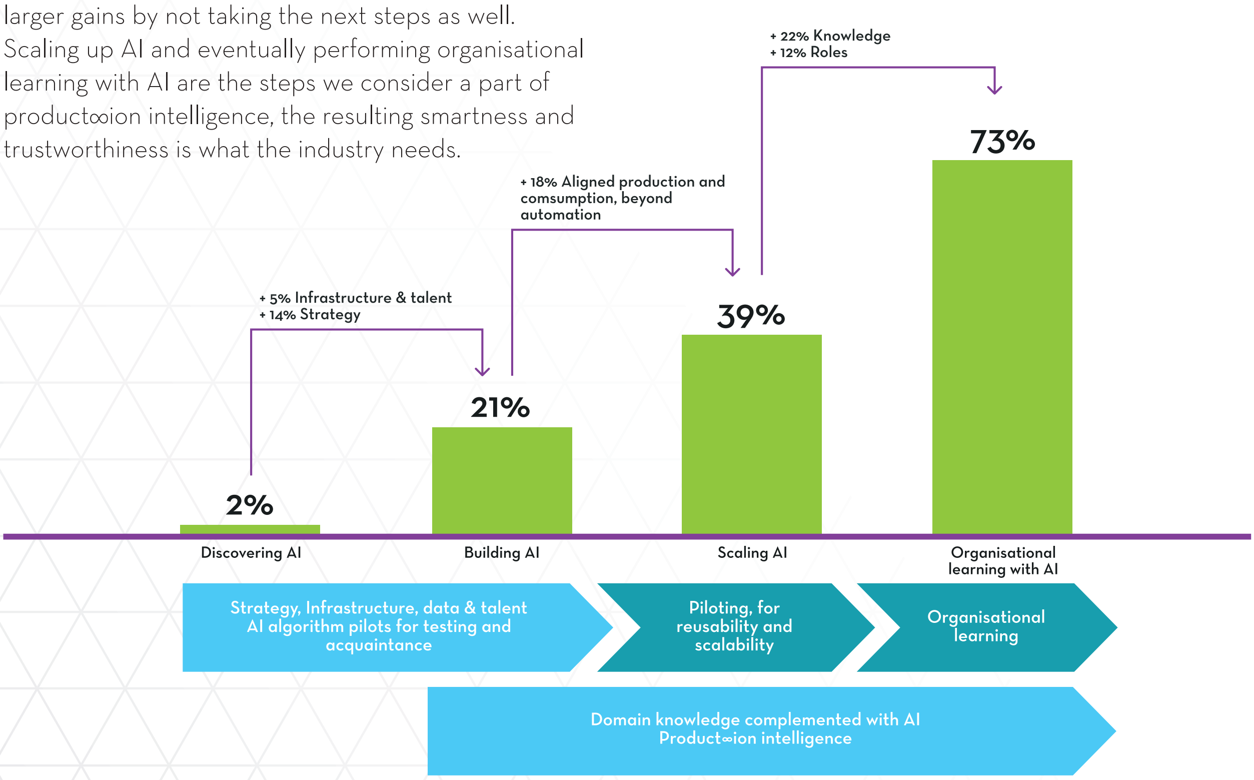

Adapted from MITSLOAN management review: Expanding AI's Impact With Organizational Learning; Oct 2020; by Sam Ransbotham, Shervin Khodabandeh, David Kiron, François Candelon, Michael Chu, and Burt LaFountain

CONCLUSION

This whitepaper positions product∞ion intelligence and the use of AI as a complement to domain knowledge for making more productive mechatronic systems, while at least remaining as safe, robust, reliable, explainable and trustworthy as the traditional alternative. Certainly the manufacturing industry but also any industry with manufacturing challenges such as food, pharma, etc. need the proposed solution to learn from data in a way that the requirements of the industrial environments are met and the economic impact is under control.

Domain knowledge is the starting point of all actions and this is complemented with one high-quality AI algorithm or a combination of sequentially used high-quality algorithms to achieve a better solution.

Let's make a comparison with the restaurant environment. An important difference between a high class restaurant and a normal restaurant is certainly the chef's domain knowledge of the food. Both chefs might use high quality ingredients (such as AI algorithms) but the master chef combines these ingredients in such a way that taste, colours and flavours are perfectly matching, well-balanced, and complementary.

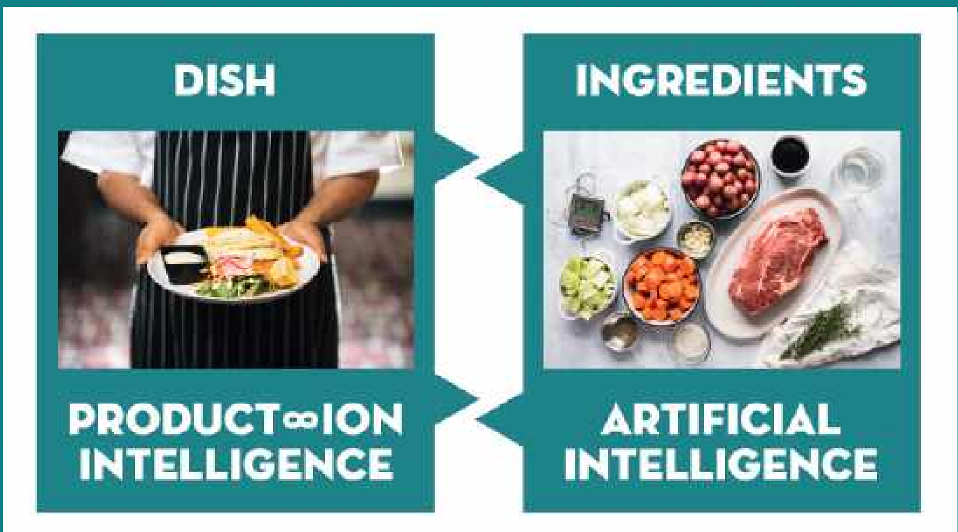

When looking back at the table of challenges we referred to early on in this white paper, product∞ion intelligence is how we solve those kinds of challenges in the industry. The challenges arise when using data alone, the domain knowledge and the use of physics and models allow us to build an AI system that is more reliable, predictable and accurate. It's first of all enriched by the experience of operators, the data from the product development stages, but most of all keeps on evolving by the constant stream of new data throughout its life cycle. This makes sure the actions taken by an AI system have no unexpected negative impact, mitigates the need for a lot of data early on, but still manages to learn from the scarce data on faulty operation during usage.

We believe product∞ion intelligence is the way forward for AI in the manufacturing industry. We will gladly show how in a few deep dives of practical examples.

DEEP DIVE INTO FLANDERS MAKE RESULTS

To further illustrate the power of product∞ion intelligence, we have gathered a few use cases wherein we use, develop and integrate several types of AI in different ways.

Product∞ion intelligence with vision

A major focus of AI has always been on the vision side of the technology. Image recognition, object recognition and navigating with camera's are fairly new technologies that are powered by an AI system reasoning about what it sees and how to interact with what it sees. In a manufacturing context, combined with product∞ion intelligence, vision can truly enrich the toolbox of modern machines.

Bin picking with AI

A common task to tackle in industrial automation is bin picking, where a robot and a vision system work together to pick objects from all sorts of containers for sorting, assembly or other purposes.
To complete this task, correctly recognising the object is of course vital. But also determining their orientation is often important, so the system knows how to pick up and place the object at the right place.

Typically, 3D-cameras are used to scan the bin and recognise the objects in it, then the robot can pick individual objects based on its geometrical features. AI algorithms are not commonly used to achieve this. However, this approach has a few limitations:

- The technology is quite expensive, a 3D-camera costs between € 10.000 and € 15.000 on average.
- The scanning is typically based on pulsed light, which is a relatively slow process. It takes a few seconds to recognise the objects because of the light pattern.
- In the computing algorithms, a fixed template with the desired geometrical features is used for it to recognise the final product. This makes it hard for it to recognise large variabilities of objects.
- The material of the objects can make matters more complicated, like reflective metals that reflect the pulsed light back to the camera. This can lead to unreliable results.

WATCH VIDEO

Flanders Make developed a different solution consisting of a standard 2D-camera and AI algorithms that are trained with datasets created from the CAD drawing of the part. By adding a light-material interaction model and surface textures, we can augment the regular CAD drawing to a photo-realistically rendered image. With that, it becomes an easy task to generate many 2D images, with detailed annotations of its contents. Manually annotating the images is no longer needed, as we know exactly which item we are rendering in which orientation, adding this information to the image has become an automatic process.

Figure 1 - Synthetic data of different industry parts created by simulation environment developed at flanders make

By using the CAD models of the industrial parts, we already incorporate a valuable source of domain knowledge that provides detailed information for the AI model training. Extending these CAD models with different variances, 3D poses, geometrical features, surface patterns and reflective patterns further improves the accuracy and robustness of the AI models. Finally, also integrating different ambient conditions, like various backgrounds and lighting, makes the AI models a lot more accurate in less strictly controlled environments.

The results that Flanders Make has achieved (named CAD2POSE) are quite spectacular and can work robustly in extreme changing conditions.

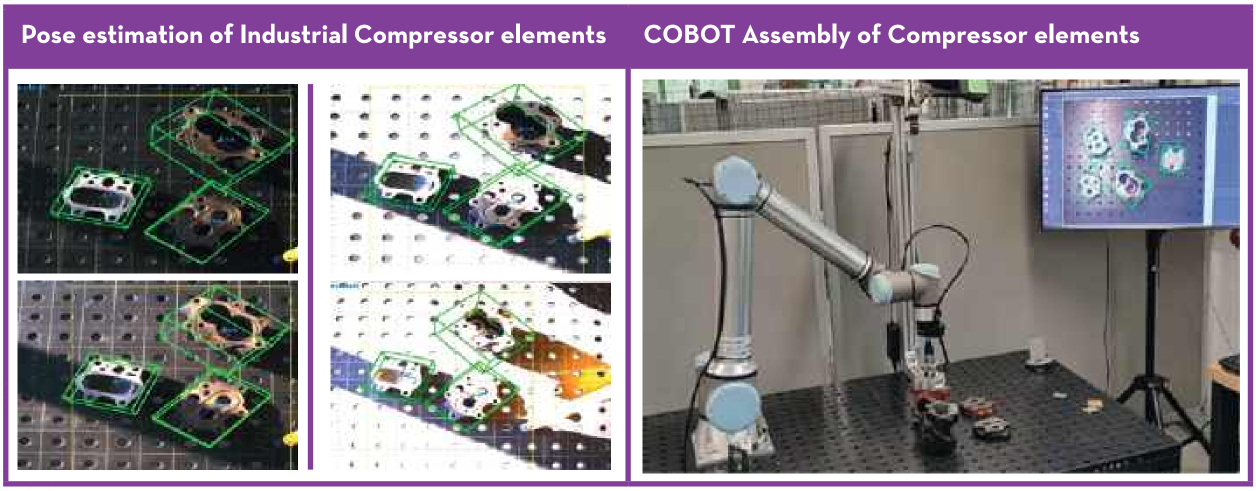

Figure 2 - 3D poses of industrial parts, using standard 2d cameras and ai models trained using augmented cad rendered data, under extreme light / contrast changes

The proposed solution has certain advantages over the more conventional approach:

- Much lower hardware costs (< € 1.000 for a suitable 2D camera)
- One-time investment to generate a large amount of training datasets with detailed annotations
- The method can be optimised to work in sub-seconds
- The accuracy can quickly increase (by > 95%) and the robustness can be highly improved by injecting more variability in the training data

Navigation

Computer vision is also used in a different application in manufacturing automation, namely as a sensor for perception on automated ground vehicles (AGV's) which are active in an industrial environment. These AGV's can perform various tasks, such as moving goods between assembly cells, warehouse inventory, inspection tasks, collaborating with humans for lifting heavy parts, and so on. For each of these tasks it is important for the AGV to know where it is and which obstacles are on its path. A powerful technique for localisation is using Simultaneous Localisation And Mapping (SLAM), in which images from a video stream are analysed for features that can be tracked over time. The goal is to recognise scenes when returning to the same spot, by comparing the locations of these features in the images taken at an earlier occasion. This allows it to navigate more confidently, since it has learned from the last time it visited this location.

The traditional SLAM techniques however are not sufficient. Driving AGV's, in an environment where humans work as well, requires safe behaviour, therefore corner cases are very important. In some cases the AGV might think it is in another location due to drift in the SLAM algorithm or bad relocations, if the features were placed on objects that are dynamic and can move from frame to frame. In current state-of-the-art versions of these algorithms, there is no discrimination between dynamic and static objects.

Flanders Make proposed to improve these algorithms with domain knowledge and segment the images of the scene in order to detect objects and classify them as static or dynamic. Simulators that mimic industrial environments and generate realistic images are used to train the algorithm, with the time and cost saving benefits of correctly labelling the contents of the image. Furthermore, an additional benefit is that these datasets with ground truth can be reused to train improved algorithms with no extra costs for generating new datasets.

AI-assisted operators

Context

The cooperation between humans and cobots is on the rise in the industrial landscape. The improved sensing capabilities of these cobots makes them a lot safer in operation together with humans. Their prices are also continuously decreasing, making them a lot more available for various tasks. At the same time, a lot of repetitive, easy tasks have already been fully automated by robots, the next obvious step is making the tasks that are typically handled by humans more efficient with the help of cobots. A cobot can, for example, lift and hold a heavy object, while the human attaches it with screws. Or a cobot can assist in certain assembly tasks that are harder for humans, like snapping together pieces that would require a lot of force for a human.

But, not every operator requires the same level of support and tasks can quickly shift in the modern assembly environment. This means that we need to add intelligent technologies to the assembly stations to deal with variability of situations and scenarios.

WATCH VIDEO

Situation-aware cobots

To further develop this, Flanders Make developed several AI techniques that can make the cobot more aware of the environment and the tasks to be executed, as well as create intuitive interfaces between human and cobot through speech and controls. Domain knowledge is utilised to increase the safety and robustness of these AI techniques. By adding relevant constraints, we can prevent unusual and unexpected output from the AI models, comparable to how guardrails prevent cars from straying too far from the expected path.

The following example is used to illustrate this AI / domain knowledge paradigm within an assembly context.

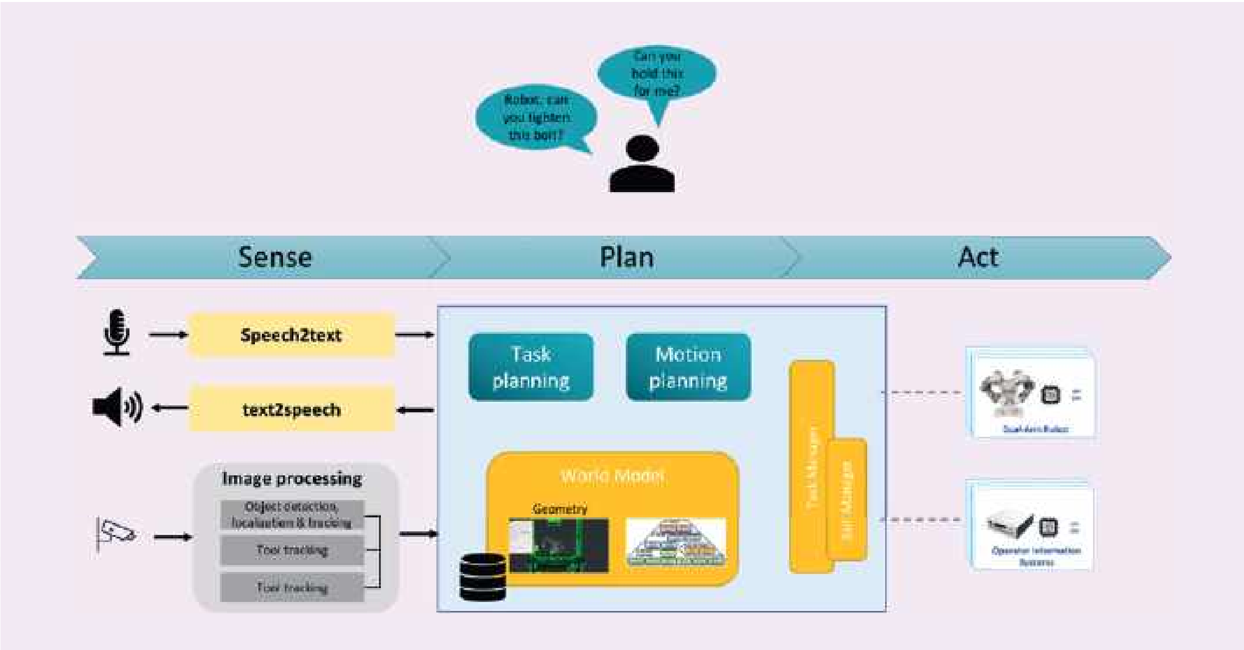

Figure 3. Sketch of AI models constrained by a formalised domain knowledge (world model)

In this example, a cobot is combined with a camera system, a voice recognition system and a human operator. The camera system is equipped with AI models to handle image processing, in order to track the progress of the assembly task and to see what the human operator is currently doing. Historical and current data is used to train this model, to improve the handling of variabilities. The speech recognition is also handled by AI models. The outputs of these systems are then used to handle the next steps after the operator finished their task, like updating the digital work instructions (DWI, Figure 4) for the next task or triggering the cobot to perform its next step. By having the system recognise what the operator is doing and responding to simple voice commands, the cobot and operator can cooperate much more fluently.

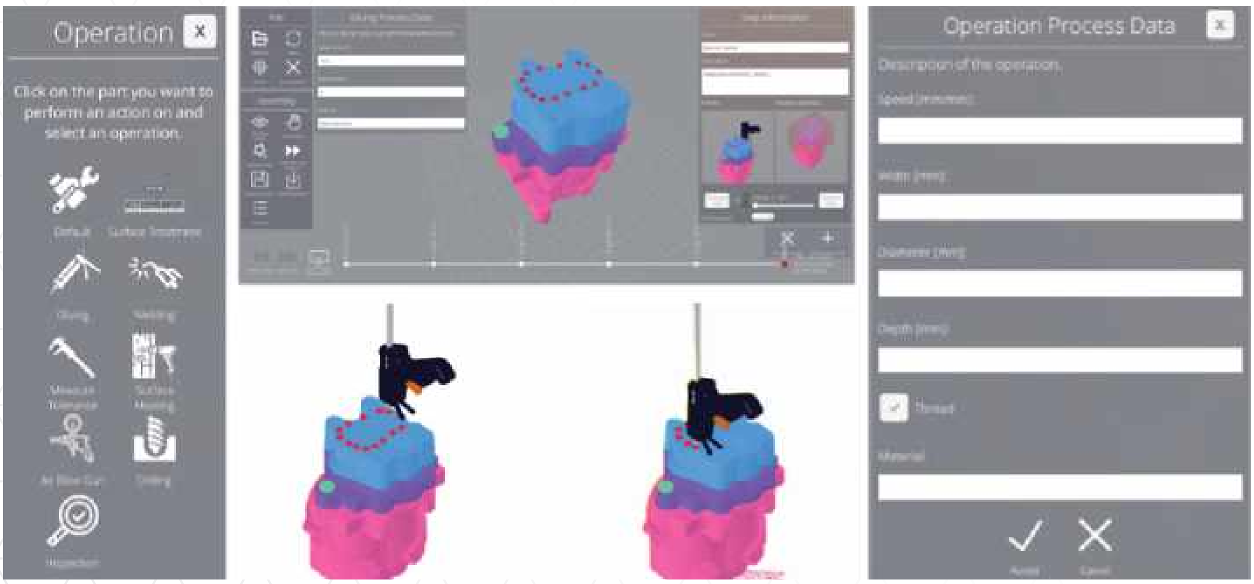

Figure 4. Example of digital work instruction (dwi) that guides operators step by step to assemble a part

When having a human and a cobot operate in the same space, safety is obviously a priority. After all, a lot of robots still operate in safety cages to make sure humans stay at a safe distance during movements. That's why we made a world model (Figure 5), based on prior domain knowledge, to formalise constraints and avoid unintended actions. For example, if the AI-based image processing model detects an object outside the safe area defined in the world model, then the system will prevent the cobot arm from moving to that area or trigger some kind of warning to the operator.

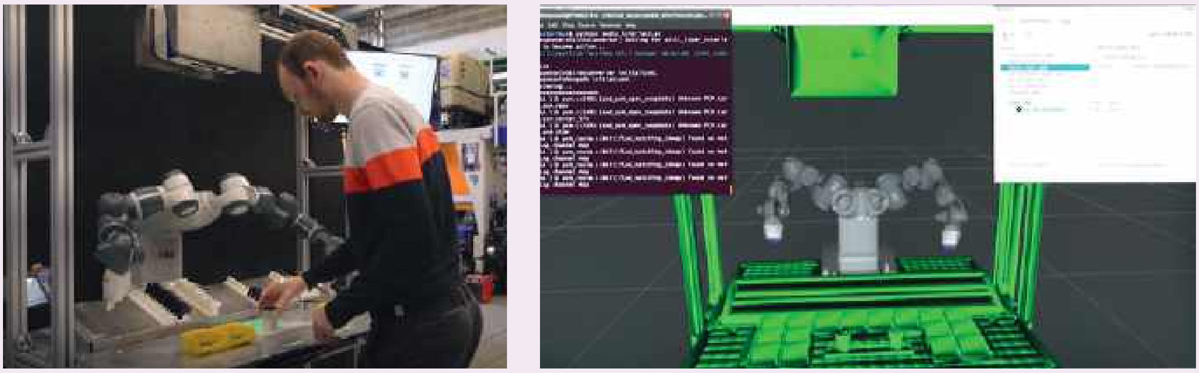

Figure 5. (Left) An assembly station with an operator, (right) a real-time digital twin derived from and updated with world model data.

Scheduling

Having these kinds of situation-aware cobots also unlocks a next step in the process. Manual assembly environments are often rather complex and difficult to manage because the time it takes to assemble a product is highly dependent on the operator. Differences in skill level, experience and level of expertise lead to differences in performance. However, by offering digital work instructions to operators that are tailored to their individual needs, their performance can be significantly improved.

One of Flanders Make's industrial partners has developed an adaptive operator monitoring and work instruction system that adapts the work instruction content to the experience level of the operator. Up until now, this system only worked when operators identified themselves. This identification step often conflicts with privacy regulations and forms an added barrier in the acceptance of the system by the operators.

To overcome this issue, Flanders Make developed an algorithm to estimate the operator's skill level for a specific task, solely based on the available detection data of the monitoring system. Because of the high-mix low-volume nature of many assembly environments, we often lack sufficient execution data for a specific tasks. Therefore, assembly tasks are decomposed into small micro tasks for which the skill and experience level is estimated. Skill level estimations for new tasks are then performed based on the available information on micro task level. This approach was validated using the data of a work station in a sheltered workplace. The neural network approach reaches an accuracy of over 85% compared to the operator identification system as currently used.

Hybrid AI – building an intimate relation

Any moving object in a smart product or production system needs a controller, designed by the **experts in the field of (classical) control engineering.** The control engineering community heavily relies upon structured, mostly state space, mathematical models of the system dynamics to design controllers. These models are built based on the engineer's prior knowledge. Due to changes in the physical system extensively interacting with its environment, a controller based on such sometimes very detailed and costly to build models still relies on guesswork during its operation. A lot of time and resources are being invested into making sure this guesswork is sufficient to make the system safe and stable.

Similarly **reinforcement learning** is also a way of defining a control policy. It is a machine learning training method where a machine is rewarded for progress towards a desired goal, without telling it how to achieve this goal. It allows to learn in an unstructured manner, through trial and error, which offers a high degree of freedom to learn and adapt. This can lead to surprising solutions a human wouldn't even think of.

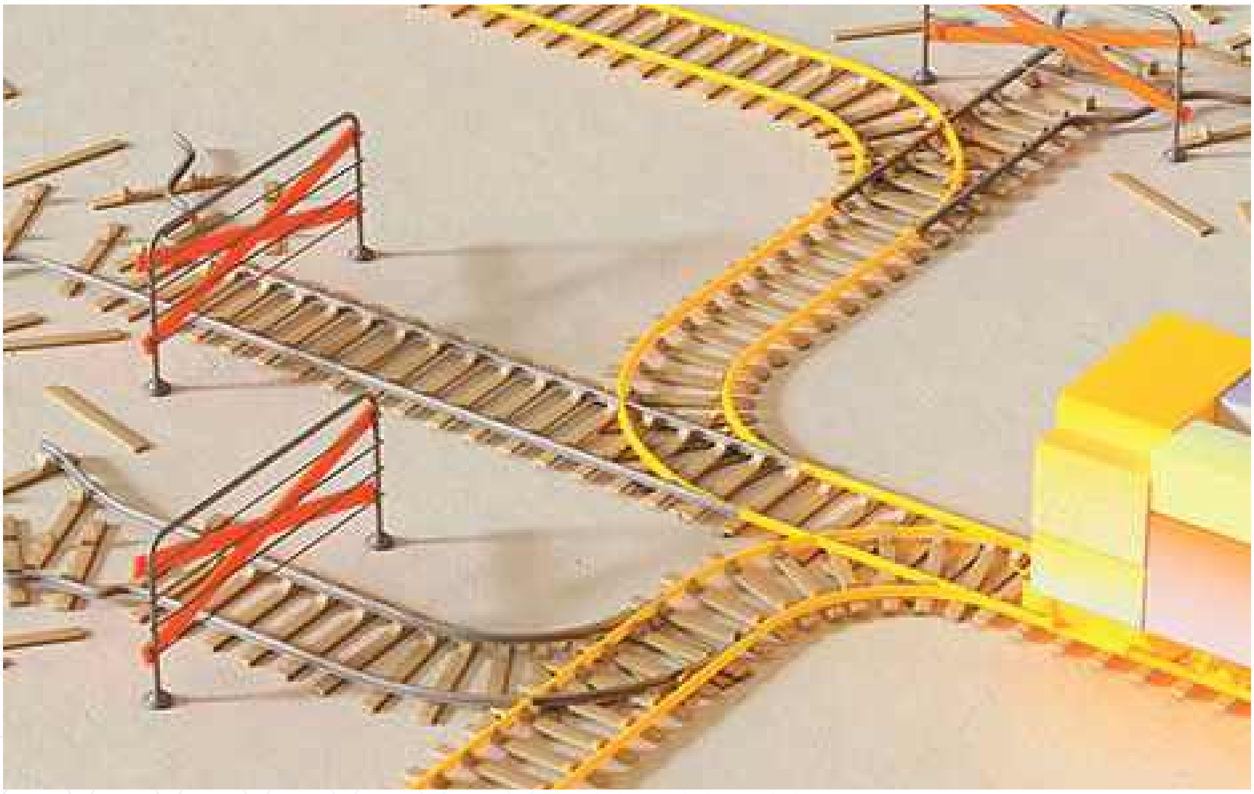

In short, while classical control focusses on maintaining safety, trust, and complex planning, in the face of model uncertainty; reinforcement learning focusses on its distinct counterpart: the iterative aspect of learning and gradual improvement by repeated interaction.

Could we get the best out of both worlds?

Recently, **residual reinforcement learning** was introduced for robot control: it trains a reinforcement learning controller residually on top of an imperfect, traditional controller. The reinforcement learning algorithm leverages the traditional controller as an initialisation to enable data-efficient reinforcement learning for tasks where traditional reinforcement learning is unmanageable, such as robotic insertion tasks where rewards are sparse. Starting from a suboptimal but adequate and robust controller, as often present in the motion control of industrial applications, we at Flanders Make have introduced the **constrained residual reinforcement learning architecture**[3] as shown in the figure below. It shows how the output of a conventional controller and actor are combined and scaled with the factor 'βr' to obtain actions 'u(s)' to drive a system. In view of human interaction, this scheme closely reflects a traditional control loop and hence increases the explainability and transparency towards control experts (Human-Machine Interaction, Machine-Domain Knowledge Interaction).

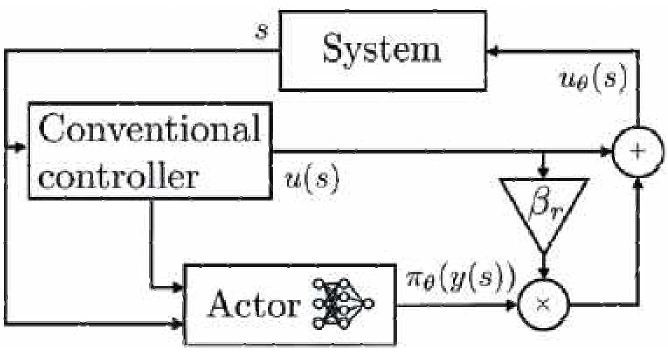

Industrial Application

Present in many industrial applications, a slider crank provides reciprocating linear motion through a rotary motor in combination with a bar linkage system. This system exhibits highly nonlinear behaviour and is often plagued by unidentified load disturbances and unknown interactions with the environment. This application is of direct relevance in various industrial systems, like compressors, hydraulic pumps, weaving looms, and presses. PID controllers are typically used in these applications e.g. to keep a constant angular reference of 60 r/min. However such PID controllers suffer from suboptimality for systems with varying loads or ambient conditions, which requires retuning the controller.

[3] T. Staessens, T. Lefebvre and G. Crevecoeur, "Adaptive Control of a Mechatronic System Using Constrained Residual Reinforcement Learning," in IEEE Transactions on Industrial Electronics, vol. 69, no. 10, pp. 10447-10456, Oct. 2022, doi: 10.1109/TIE.2022.3144565.

The figure below shows the performance of a relative CRRL controller with βr = 0.2 compared to a benchmark PI controller. ωd denotes the desired and ω the actual crank angular velocity. The performance of a PI controller on the slider crank setup varies slightly from run to run despite lab conditions with limited external disturbances. Therefore, each experiment starts with an initial run-in phase in which only the PI controller acts on the system to benchmark the results (the blue shaded region). This variability of the PI controller's performance over different runs is illustrative for the difficulties in optimally tuning a controller for all conditions. The error offset of the PI controller illustrates the inability of the PI controller to compensate for the nonlinearities throughout one revolution which amounts to the mean error shown (blue/red line).

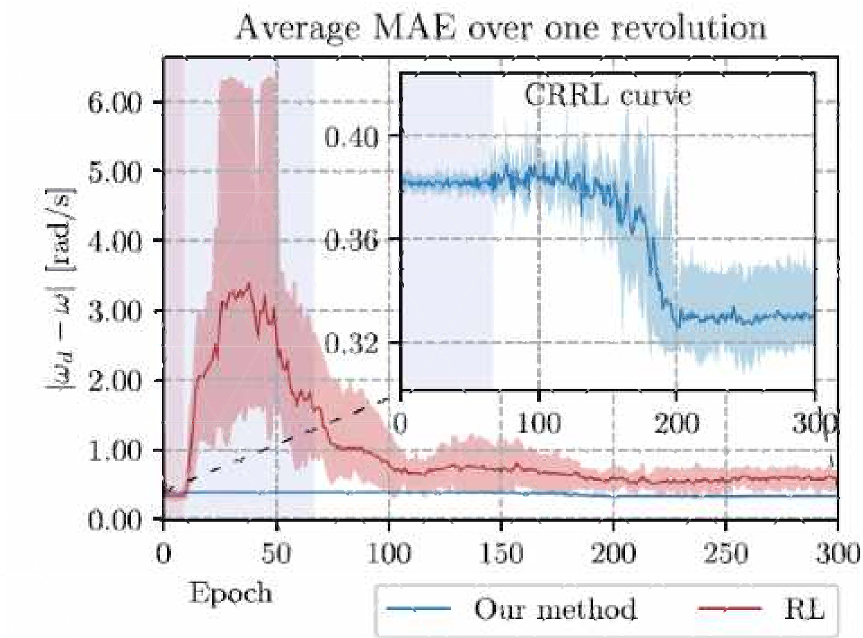

Learning performance - After the initial phase, the residual controller is activated to learn. Early on in the training process, the RL method starts to explore causing a tremendous loss in performance and even possibly unsafe controller outputs. During the exploration phase in CRRL this is strongly limited.

Operating performance - After the learning phase, the figure below on the left shows the average performance improvement after convergence in terms of MAE (mean absolute error) and MSE (mean squared error). The right side shows a boxplot of the relative decrease in performance of all epochs after activating the residual controller where the reward was lower than the average PID reward. From these, we learned that the CRRL is is beneficial to achieve a higher optimality while maintaining safe operation.

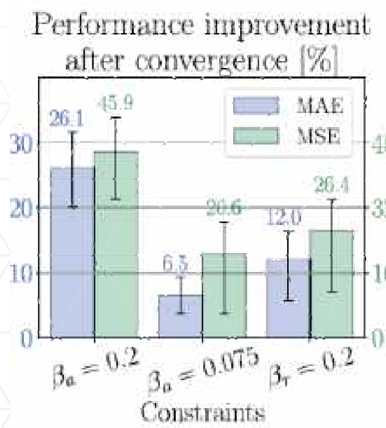

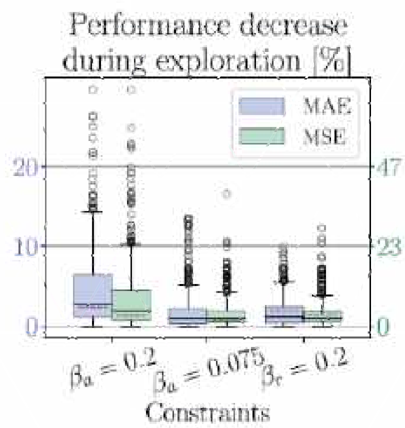

AI to improve the bonding process

Context

From an ecological and economic perspective, companies are increasingly looking for ways **to make products lighter and stronger**. It is therefore important to find the right **combination of materials** for a particular application. This means that a lot of research is being done into the way in which these materials can be **joined/bonded**. Lately, robots and cobots are used more often to improve the quality of bonding joints and support operators during difficult or repetitive assembly tasks. In order to achieve an optimal bond and bonding process, a lot of time and attention is invested in choosing the **correct settings of these cobots and robots**. For example, regular fine-tuning of the settings of the robot and the discharge unit is of great importance in order to achieve good glue bead quality. However, this is often time-consuming and based on trial & error, halting the production process. In order to speed up this setup process, Flanders Make has developed a method to optimise a number of important parameters faster without disrupting the production process, using Hybrid AI.

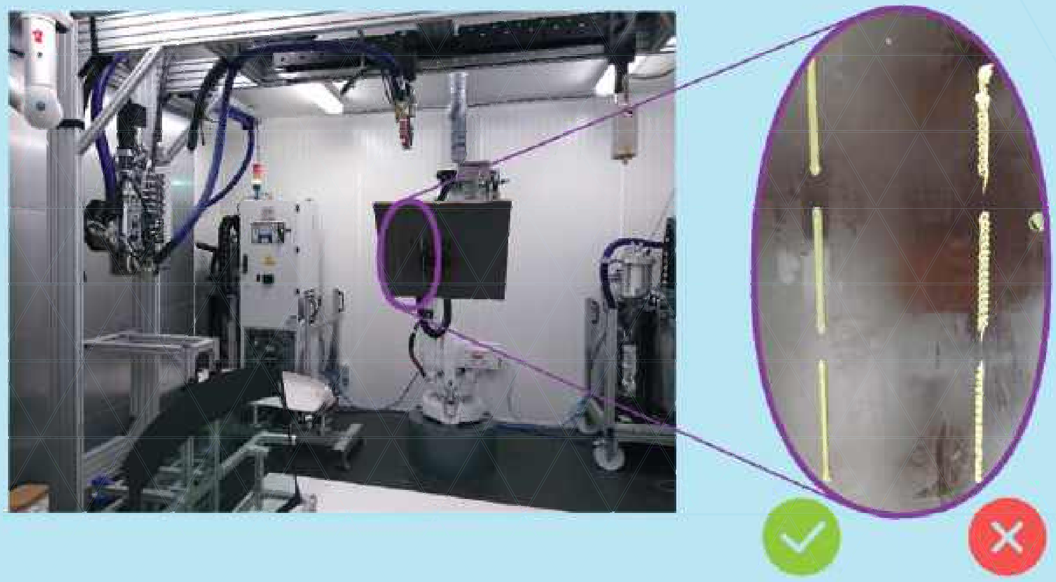

When looking for the optimal bonding technology for specific industrial applications, we can optimise various variables. These are not only the variables of the robot itself, but also the (climate) conditions of the bonding, the materials used and the dosage of the adhesive. In research that Flanders Make has recently conducted, the robot speed, pre-pressure of the dosing unit, the amount of adhesive liquid per second and the distance between the nozzle and the work object were measured in its Joining & Materials Lab. The above shows how the quality of the bonded product is directly related to the production parameters realised by another product (the robot/cobot), more and more we observe this interrelationship and crossing of the product-production domain.

Approach & impact

More specifically, we first applied **data-efficient data-driven models**, namely the Gaussian process regression in the context of Bayesian optimisation. This way, we were able to link the process control settings with the product quality. Furthermore, to optimise the process settings, a **hybrid AI model** was built based on the measurements and making use of **expert knowledge** with the routines of Flanders Make@UGent.

In addition, a calculation was made of the savings by minimising variation in the adhesive amount on the surface, the width of the adhesive bead and maximising the speed of the adhesive application.

By using a hybrid AI model on the two focus points in this study, namely optimization of altitude and speed, the setting time of these parameters was reduced **from two days to just three hours**.

In addition, **a reduction in production costs** was achieved, this also led to a more **stable and robust production process**.

Outlook

Starting from this work, we are further improving the methods, to make the tuning data-efficient for process control so that that it becomes even more viable in industry. We leverage not only data-efficient Bayesian optimisation methods that include experimental data but also **digital twin information** to construct **quality predictive models**. Next, re-tuning of process parameters is enforced by changing circumstances, like different operating conditions or welding processes, that can lead to inferior predictive capabilities if the quality predictive models are left unchanged.

We will realise this by means of transfer learning, which corresponds with hot-starting the learning/optimisation of one similar process condition from another, and meta-learning.

This is an example of how we acknowledge the importance of controlling a production process in order to optimise the performance of the resulting product and wherein we bridge the product-production domain.

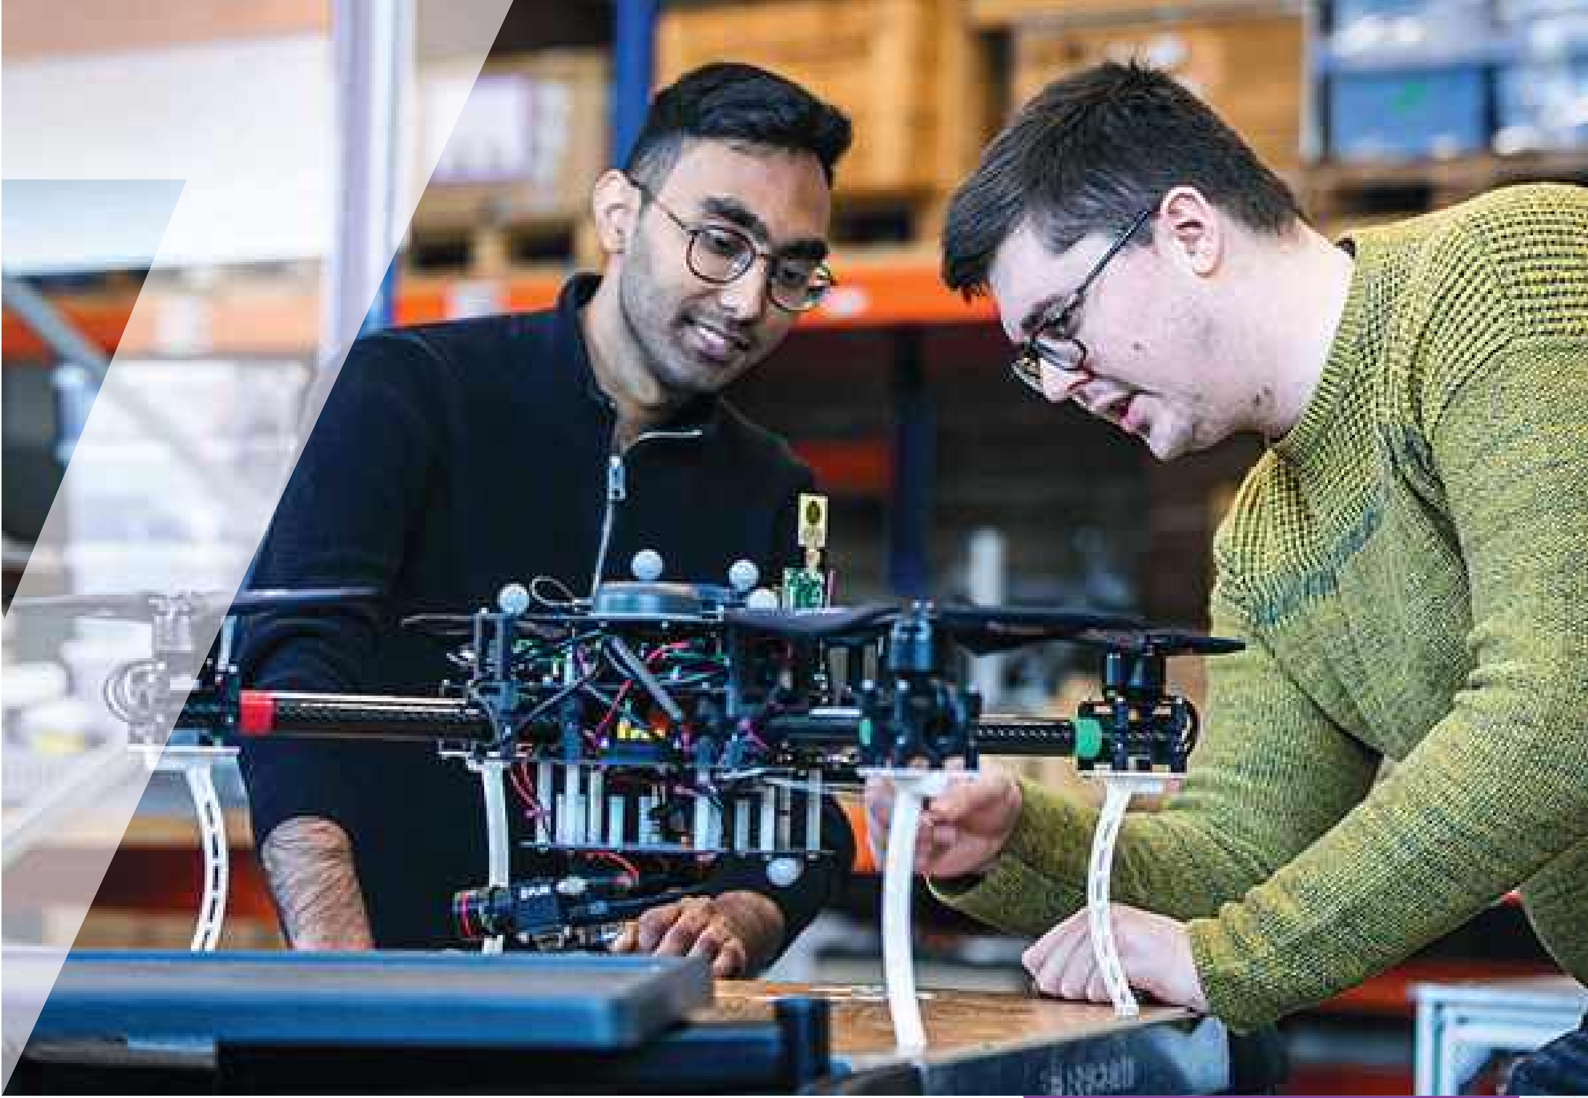

AI for the design-operation continuum

Context

For the design of moving products and production systems, historically, companies relied on physics-based models that apprehend the behaviour of mechatronic systems and rely upon domain-specific electromechanical engineering knowledge. These models have a few physical parameters that relate directly to the actual physical system and can be expressed in static input-output expressions and dynamic partial and/or ordinary differential equations. These are mostly parametric (with inertia, damping, etc. parameters). A typical way to align the models with real-world mechatronics is by identifying the physical parameter values such as backlash, damping, inertia and friction.

During the operation of these products, data are being generated. By using these data, without relying on physical priors, it is possible to find structured data-driven models such as NARMAX, Wiener-Hammerstein and other data-driven models. Currently there is a rise in the use of AI (or data-driven) models that increase the degrees of freedom to apprehend complex relationships in data. Supervised learning is typically used to train the many parameters of a data-driven model by relating input to output data. (Deep) neural networks have shown their ability to approximate various physical dynamical systems exhibiting complex and nonlinear behaviours. In manufacturing, data-driven methods have demonstrated their value at the supervisory level since they adhere well to the abstract level. However, closer to the physical processes, significant predictive errors may however appear outside their training data.

Approach

In order to combine the best of both worlds **hybrid modelling** approaches that combine physics-based with data-driven models have been suggested. In such models, physical equations are combined with black-box neural networks. The neural network mappings assist in compensating for prediction discrepancies in the physics-based model. Flanders Make has been working with several types of hybrid models ranging from Physics Inspired Neural Networks (PINNs) up to 'multiplicative' hybrid models developed by Flanders Make (see figure below[3]). Herein, we closely combine ordinary differential equations and differential-algebraic system of equations (ODEs/DAEs) with neural networks. We refer to it as multiplicative because instead of adding a residual data-driven model output to the physics-based model to correct the latter, the physical states are connected to neural network nodes, the physics and data parameters are trained.
The neural network is then not used to merely **adapt some parameters** of the physics-based model, instead it is adding **unmodelled dynamic behaviour** to the physics-based model. Such unmodelled behaviour could be either directly linked to the machine such as e.g. friction within a mechatronic transmission or to a process behaviour that cannot be directly measured or modelled such as e.g. slip or stretch of a foil in a foil winding machine. Hence, as these data add new aspects to an existing model and as such improves the model, these hybrid models function as one of the tools to build digital twins. Furthermore, at Flanders Make, we use these hybrid models as one of the tools to realise a design-operation continuum: (physical) design models are used during operation and data from operation are used to build better models that are in turn used in the design.

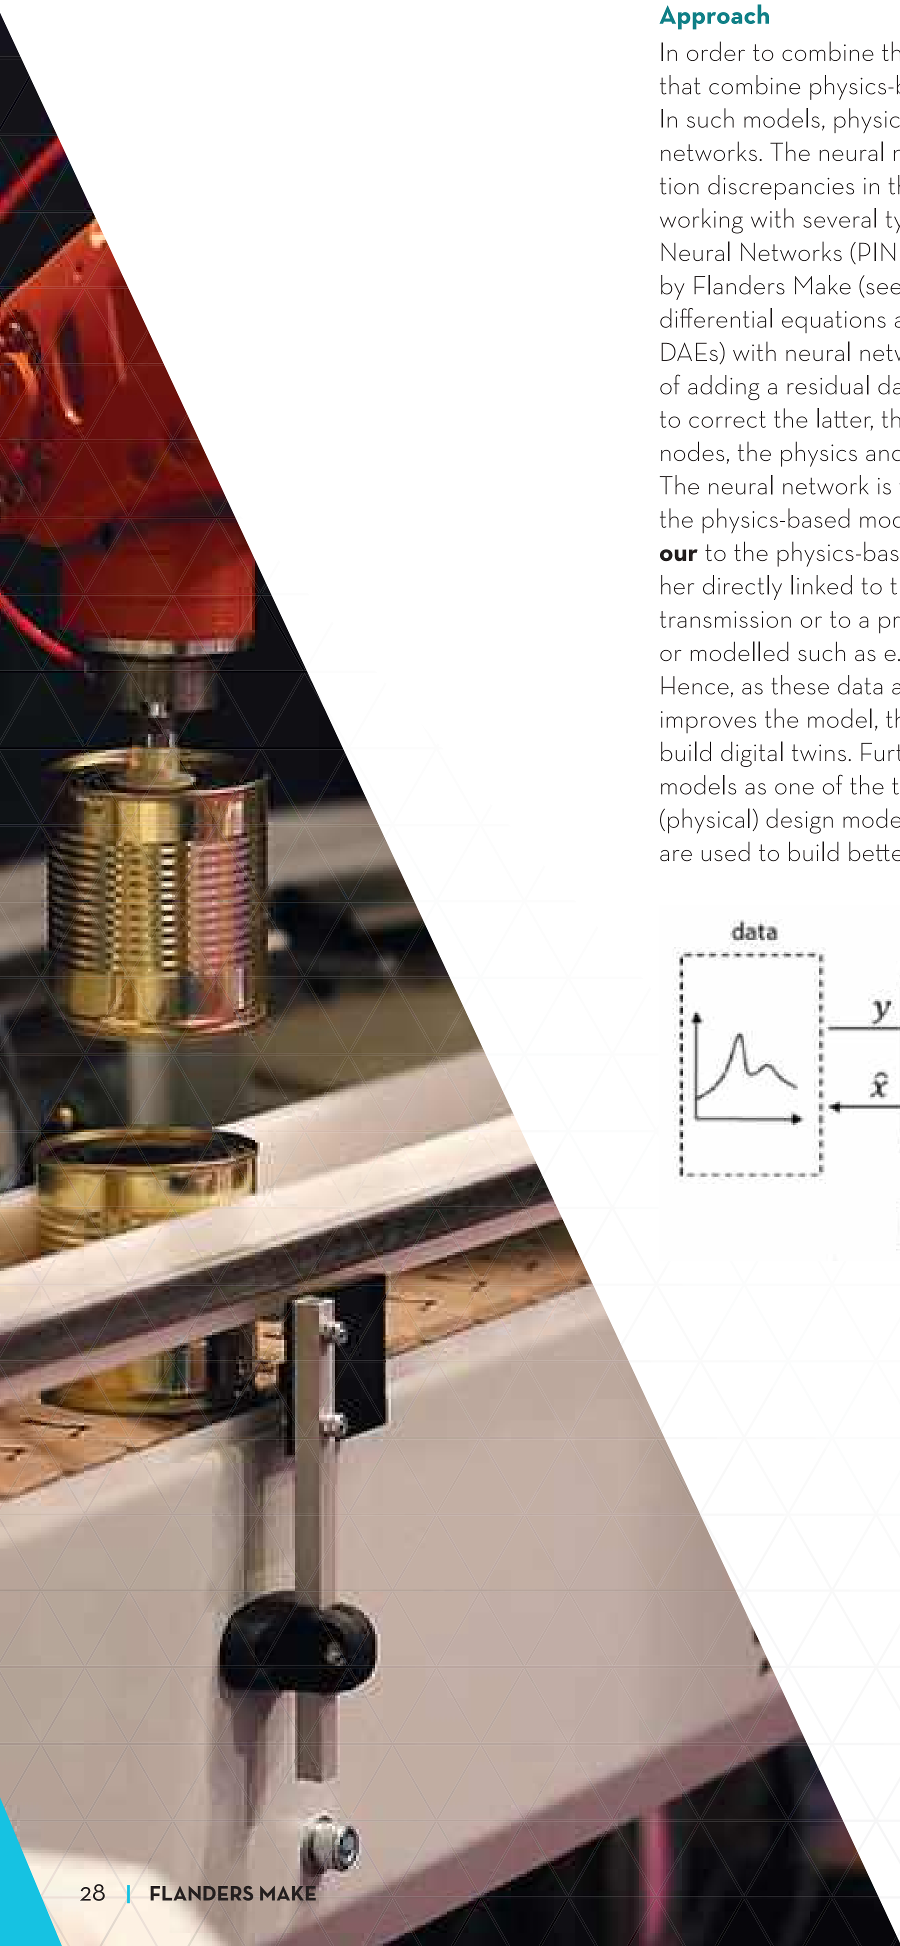

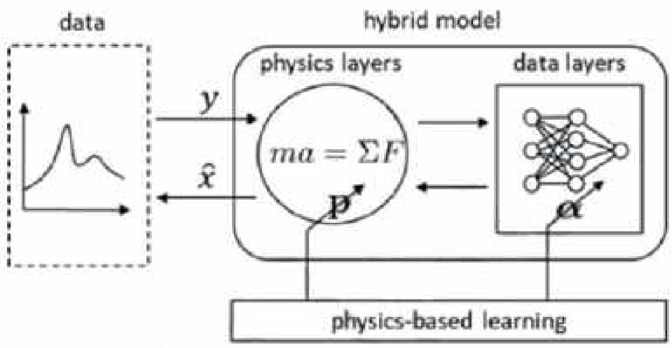

[3] W. De Groote, E. Kikken, E. Hostens, S. Van Hoecke and G. Crevecoeur, "Neural Network Augmented Physics Models for Systems With Partially Unknown Dynamics: Application to Slider-Crank Mechanism," in IEEE/ASME Transactions on Mechatronics, vol. 27, no. 1, pp. 103-114, Feb. 2022, doi: 10.1109/TMECH.2021.3058536.

Industrial application

To illustrate this, we consider a **cam-follower** mechanism which is used in many applications to translate a rotary motion into a linear displacement. Designing such a cam-follower, e.g. to derive the shape, mass, … for a large range of operating conditions such as speed, loads, … ensuring the follower trajectory is such that it does not jump is challenging. The latter is due to e.g. unknown friction, etc. and the fact that the e.g. the height of a follower jump does not depend on the momentaneous rotational measurements but relies on the energy build-up created when contact was still assured.

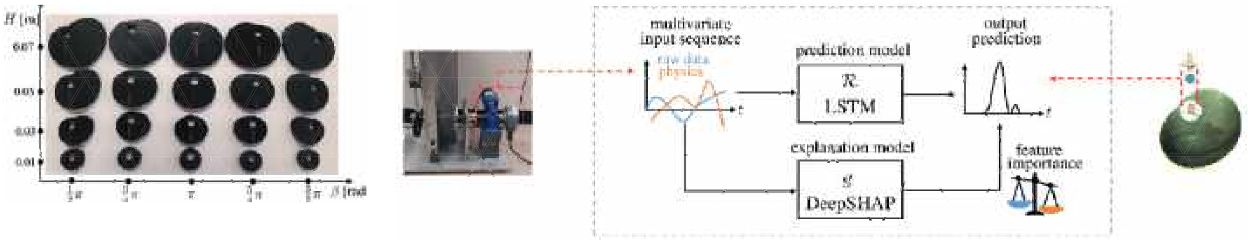

In this case, Flanders Make defined a **hybrid prediction** model, whereby the trajectory of the detachment variable was learned by an LSTM neural network model that receives a combination of direct measurements, **physics inspired features**, and system properties. We then further studied the relative importance of the inclusion of physics-inspired expert knowledge using a DeepSHAP (Deep Shapley Additive Explanations method). We have shown that, only thanks to including physics-inspired features, the generalisation capabilities of the model drastically increased; both over the operating conditions of the cam follower as well as its shape; making the resulting model ideal for further designs realising the **design-operation continuum vision**.

Major productꝏion intelligence-related research conducted by Flanders Make:

1. Product-ion Data related research
2. Research into AI algorithms selection/classification for guaranteed performance
3. Research into Domain knowledge capturing & integration
4. Research into Explainability loop // Hardcoding of domain knowledge in AI algorithm structure // physics based IO constraining //
5. Research into Robustness loop and constraints definition for higher efficiency in training networks and exploitation
6. Research into Mutual learning / (partial yet smart) learning (expert knowledge, family features) / architectural design of combinatory AI algorithms
7. Research into PI for impact creation and enabling business value // deployability (embedded, edge, real time, ..)

Acknowledgements

- Dirk Torfs – Flanders Make
- Jeroen De Maeyer – FlandersMake@UGent
- Andrei Bartic – Flanders Make
- Wouter De Nijs – Flanders Make
- Abdellatif Bey-Temsamani – Flanders Make

References

1. Decision-making loop for the continuous optimization of the production and product. (n.d.). Siemens. https://new.siemens.com/global/en/ markets/automotive-manufacturing/ digital-twin-performance.html
2. RANSBOTHAM, S., KHODABANDEH, S., KIRON, D., CANDELON, F., CHU, M., & LAFOUNTAIN, B. (2020, October 20). Expanding AI's Impact With Organizational Learning. MIT Sloan Management Review. https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/
3. W. De Groote, E. Kikken, E. Hostens, S. Van Hoecke and G. Crevecoeur, "Neural Network Augmented Physics Models for Systems With Partially Unknown Dynamics: Application to Slider–Crank Mechanism," in IEEE/ASME Transactions on Mechatronics, vol. 27, no. 1, pp. 103-114, Feb. 2022, doi: 10.1109/TMECH.2021.3058536.

More on www.flandersmake.be

Flanders Make consists of three co-creation centres,
the Flemish drone federation EUKA
and core labs at the five Flemish universities:

Co-creation centre for machine development
Gaston Geenslaan 8
3001 Heverlee

Co-creation centre for the vehicle industry
Oude Diestersebaan 133
3920 Lommel

Co-creation centre for Industry 4.0 production
Graaf Karel De Goedelaan 5
8500 Kortrijk

The Flanders Make – EUKA dronecluster
Lichtenberglaan 1090/102
3800 Sint-Truiden

Contact

For more information or if you would like to discuss future collaborations,
please send an e-mail to info@flandersmake.be
or call us at +32 11 790 590.

FLANDERS MAKE
DRIVING INNOVATION IN MANUFACTURING

Vlaamse overheid

EFRO
EUROPEES FONDS VOOR REGIONALE ONTWIKKELING

Europese Unie
Europees Fonds voor Regionale Ontwikkeling

WWW.FLANDERSMAKE.BE