MIKRONIEK

PROFESSIONAL JOURNAL ON PRECISION ENGINEERING

2018 (VOL. 58) ISSUE 6

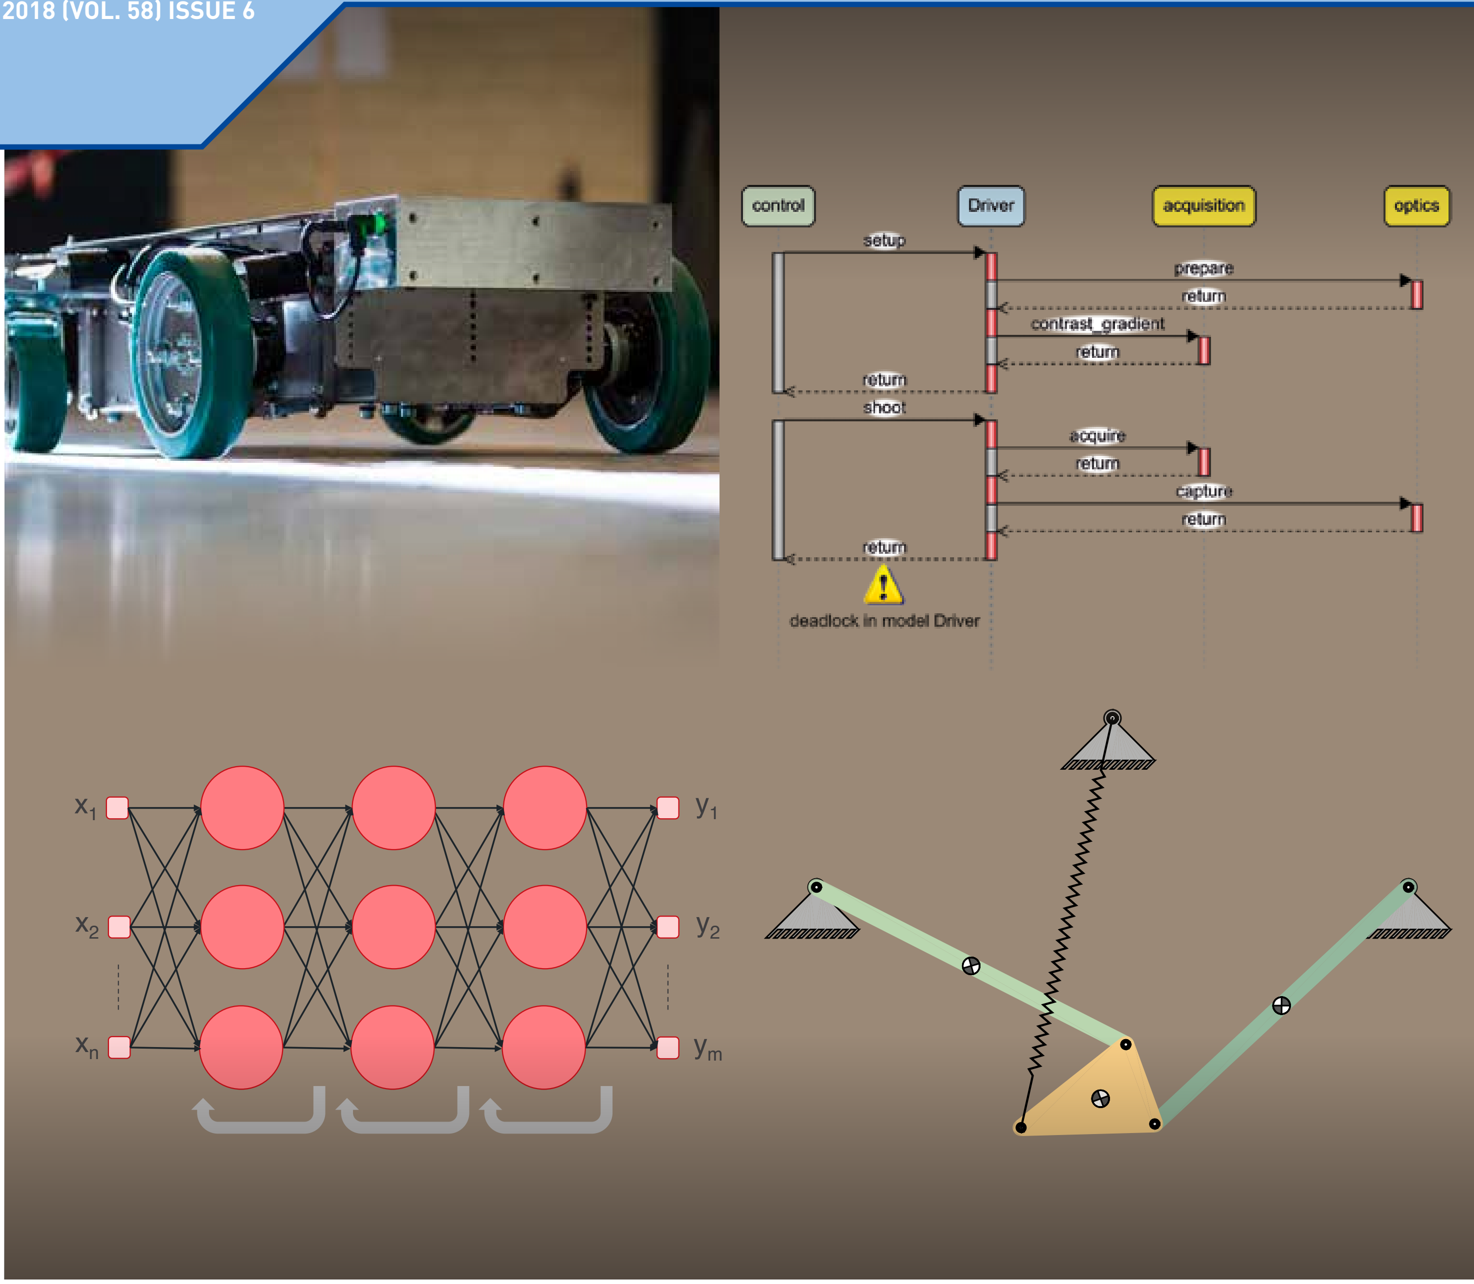

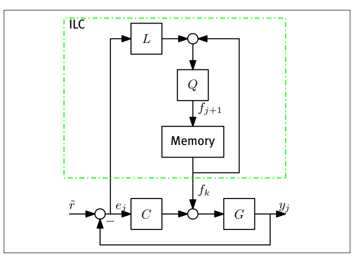

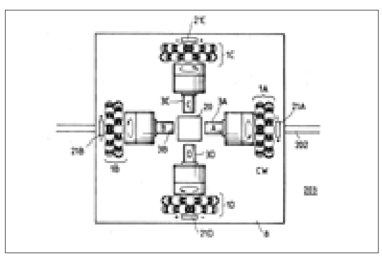

- THEME: SOFTWARE & MACHINE LEARNING
- STABLE X-RAY VISION FOR ATHENA SPACE TELESCOPE
- 2018 PRECISION FAIR IMPRESSIONS
- VALUABLE PATENT INFORMATION

DSPE

MIKRONIEK IS A PUBLICATION OF THE DUTCH SOCIETY FOR PRECISION ENGINEERING WWW.DSPE.NL

PUBLICATION INFORMATION

Objective
Professional journal on precision engineering and the official organ of DSPE, the Dutch Society for Precision Engineering. Mikroniek provides current information about scientific, technical and business developments in the fields of precision engineering, mechatronics and optics. The journal is read by researchers and professionals in charge of the development and realisation of advanced precision machinery.

DSPE
YOUR PRECISION PORTAL

Publisher
DSPE
Annemarie Schrauwen
High Tech Campus 1, 5656 AE Eindhoven
PO Box 80036, 5600 JW Eindhoven
info@dspe.nl, www.dspe.nl

Editorial board
Prof.dr.ir. Just Herder (chairman, Delft University of Technology, University of Twente),
Servaas Bank (VDL ETG), B.Sc.,
ir.ing. Bert Brals (Sioux Mechatronics),
dr.ir. Dannis Brouwer (University of Twente),
Maarten Dekker, M.Sc. (Philips),
Otte Haitsma, M.Sc. (Demcon),
ing. Ronald Lamers, M.Sc. (Thermo Fisher Scientific),
Erik Manders, M.Sc. (Philips Innovation Services),
dr.ir. Pieter Nuij (NTS-Group),
dr.ir. Gerrit Oosterhuis (VDL ETG),
Maurice Teuwen, M.Sc. (Janssen Precision Engineering)

Editor
Hans van Eerden, hans.vaneerden@dspe.nl

Advertising canvasser
Gerrit Kulsdom, Sales & Services
+31 (0)229 – 211 211, gerrit@salesandservices.nl

Design and realisation
Drukkerij Snep, Eindhoven
+31 (0)40 – 251 99 29, info@snep.nl

Subscription
Mikroniek is for DSPE members only.
DSPE membership is open to institutes, companies, self-employed professionals and private persons, and starts at € 80.00 (excl. VAT) per year.

Mikroniek appears six times a year.

ISSN 0026-3699

The main cover photo collage (starting top left, clockwise) is courtesy of Nobleo Technology, Verum, Reinier Kuppens and Albert van Breemen. Read the articles on pages 34 ff, 28 ff, 24 ff and 12 ff, respectively.

IN THIS ISSUE

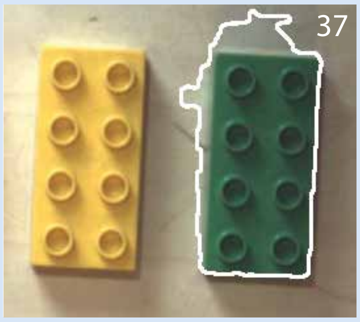

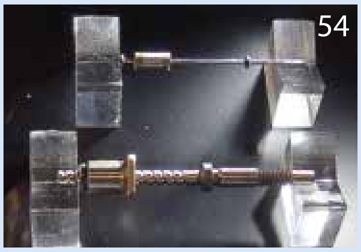

THEME: SOFTWARE & MACHINE LEARNING

FEATURES

INNOVATION AND DIGITALISATION OF **SOFTWARE DEVELOPMENT**

It takes software to turn the hardware of a complex high-tech machine into a dynamically operating system. Software is therefore an essential part of these systems, but it is upsetting to see how difficult it still is to develop that software without errors, and to prepare the code for a flawless integration into such a machine. In the 20 years that I have been visiting clients for my software company, I have observed very little change. There is a big gap between, for example, the mechanical engineers, who are busy with the physical system, and the software people. The latter think in code, while hardware people have much more of a system-level way of thinking. It is for this reason that we have stopped only executing software projects with our company and have switched to focus on the development of software products improving this.

Software development has to be done using a platform with which you can describe the desired behaviour of the system to be designed in a graphical, model-based way that is automatically converted to code. The hardware people are also able to work with this model; it creates a common language that bridges the gap between the software and hardware disciplines, allowing them to collaborate on the design at a system level.
Good riddance to manual, error-sensitive coding. The best way to create error-free code is to generate software automatically. What remains is errors at the design level. The more disciplines involved in software design (thanks to that common language), the more people can review and remove design errors. The next step is simulation: linking a model of the physical system to the 'real' operating software and then testing all sorts of scenarios to check whether the behaviour of the system is what the designers had in mind. This should mean that there are now very few mistakes left, and they will come to the fore when testing the physical system.

This method of software development and testing becomes all the more important if the design of a system, the hardware and the software, is regularly upgraded. This often takes place on the basis of machine learning: from the data generated by a system during its operation, improvements can be derived. Machine learning itself requires, for example, data analysis, and hence a lot of software – which also needs to be error free. Just like those updates themselves, of course. If, for example, the update of a smartphone does not work flawlessly, the market will not pick it up.
In short, software's share in high-tech projects is constantly increasing – partly under the influence of the digitalisation that Smart Industry brings with it – and innovation in the world of software is therefore urgently needed, otherwise the problem of errors will become totally uncontrollable. Digitalisation of software development (code generation, testing, simulation) should contribute to this. In the Dutch high-tech industry, however, the urgency of this has not yet been fully realised. This has to do with the fact that higher management of high-tech companies is mainly populated by hardware people.

Here lies a role for the High Tech Software Cluster, which in the Brainport region unites over 20 high-tech software companies in the areas of virtual prototyping & design, model-based software and data analytics & services. The ambition of this cluster is to contribute to shortening time-to-market and helping to prevent complex development projects become unverifiable and uncontrollable. A highly relevant cluster initiative is the Smart Industry fieldlab Software Competence Centre, which will work on innovation in software and on software-driven innovation, including topics such as digital twinning and model-based engineering.
When I see the massive commitment for the digitalisation of the industry in Germany, through the implementation of Industrie 4.0, I see also that we must take a firm stance in the Netherlands and embrace innovation and digitalisation of software software development. It's not yet too late.

Benno Beuting
CEO Cordis Automation
benno.beuting@cordis.nl, www.cordis.nl

LEARNING IN MACHINES

Control of high-tech mechatronic systems traditionally involves feedback and feedforward control, and essentially only uses a few recent measurements. Here, we aim to explore what can be learned from all available sensor data. A general learning framework is developed that exploits the abundance of data of previously executed tasks. Both fundamental insight and experimental results show that such iterative learning control approaches enable substantial performance improvement compared to traditional control. Interestingly, traditional model-based control theory turns out to have an essential role for fast and safe learning from measured data.

TOM OOMEN

Introduction

The learning from data and information has led to impressive achievements in recent years. Computer algorithms are now capable to successfully learn in many domains, including human language, ranging from speech recognition to accurate translations, real-time pattern recognition from images, digital advertising, self-driving vehicles, Atari, and Go [1]. The key enabler has been the availability of large amounts of data as well as ubiquitous and scalable computation and software.

In sharp contrast, high-tech mechatronic systems, such as manufacturing machines and scientific instruments, are often produced and installed with a pre-defined feedforward/feedback control algorithm, and their performance deteriorates over time due to wear, ageing and varying environmental conditions such as temperature variations. Examples range from lithography machines, 2D and 3D printers and pick & place robots, to microscopes and CT scanners.

Interestingly, these high-tech machines are prime examples of mechatronic system design, where control algorithms are typically implemented in a computer environment. Hence, over the lifetime of these high-tech machines, an abundance of data becomes available, yet this is often not exploited to enhance its performance. Indeed, sensors in mechatronic systems are often used for feedback control, which typically only makes use of real-time position and velocity information.

The aim of this article is to explore opportunities for learning from data in machines, possibly from past and already completed tasks, to control them to the limit of their physical capabilities. A framework for fast and safe learning is presented. Furthermore, at the end of the article, several practically relevant questions are addressed, including what should be done for a broad industrial deployment, what performance can be expected for a specific system at hand, and whether learning control can replace traditional feedback controllers.

Learning requirements

Learning in machines imposes several unique requirements, resulting from the fact that such machines are cyber-physical systems, involving interactions with the real world. In particular, the following requirements are considered throughout:

1. Learning should be fast, since machines require experiments in real-time. In addition, fast adaptation can be useful in case of varying operating conditions, e.g., due to temperature changes induced by motor heating or day/night periodicity.
2. Learning should be safe and use operational data, since dedicated experiments may induce production loss and even damage of the machine.

In the forthcoming sections, an approach to learning in machines is investigated that addresses these requirements.

Learning from past tasks

The aim of this section is to investigate the learning from data. This leads to an approach that bridges data-based learning and model-based control.

Traditional motion control

The printer in Figure 1 is considered as a key example of a mechatronic system. Here, the goal is to position the carriage that contains the printheads. A motor delivers an input u, which moves the carriage using a belt. The output position of the carriage y is measured using a linear encoder. The printer itself is denoted G.

AUTHOR'S NOTE

Tom Oomen is associate professor at Eindhoven University of Technology, the Netherlands, in the Control Systems Technology group at the department of Mechanical Engineering.

t.a.e.oomen@tue.nl
www.tue.nl/cst
www.toomen.eu

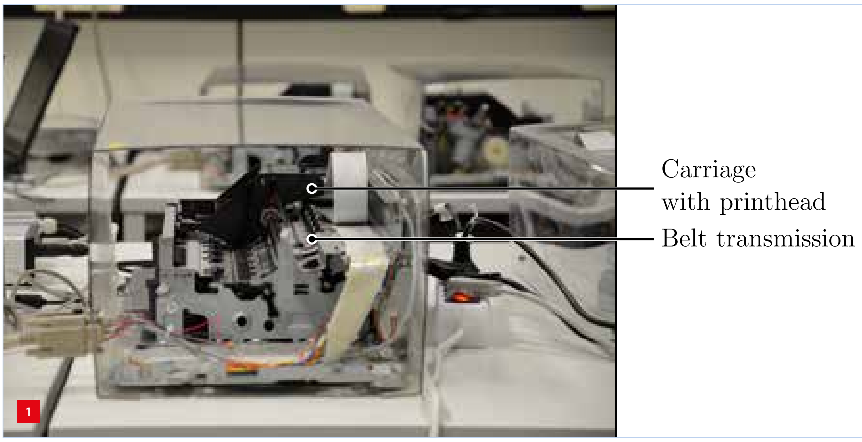

Printer system, used to illustrate traditional motion control and learning.

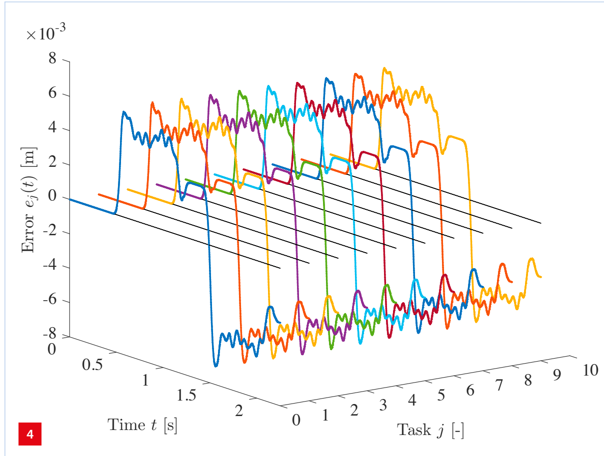

Traditional motion control: the measured error signal is almost identical for subsequent tasks.

The control task is to track a reference trajectory r, such that the printhead moves over a sheet of paper, see Figure 2. The control problem is thus to choose the control input u such that the error $e = r - y$ is small. Traditionally, this is done using the controller structure shown in Figure 3. Here, C is a feedback controller. Feedforward control is implemented by selecting the signal f. A typical approach is to employ Newton's second law, $f = m{\cdot}a$, where m is an estimate of the mass and $a = d^2r/dt^2$ is the acceleration profile.

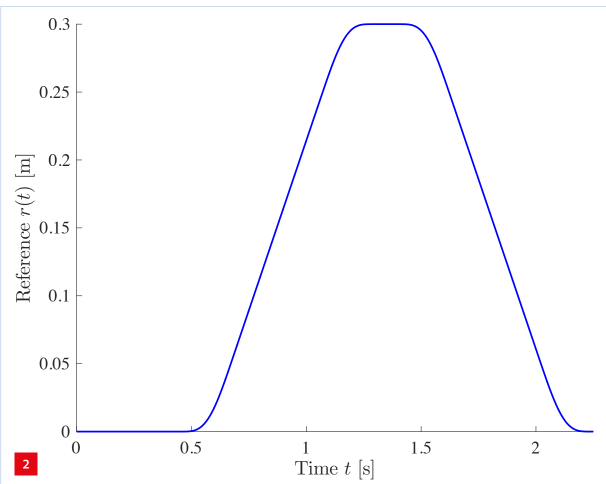

Reference for the carriage containing the printheads.

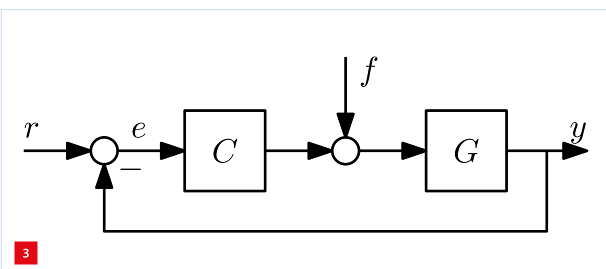

Control architecture, where G denotes the system, C is the feedback controller, f is the feedforward input, r is the reference in Figure 2, and e is the tracking error.

Motion control tasks are often performed repetitively. For example, the reference in Figure 2 has to be performed many times before a sheet of paper is printed: during each repetition of the reference in Figure 2, the sheet is moved a few millimeters by a sheet-positioning mechanism. The typical performance of traditional feedback motion control for such repetitive tasks is shown in Figure 4.

Here, ten tasks are shown, where in each task the reference in Figure 2 is tracked. The key observation is that the measured error is almost identical for each task j. Of course, feedforward control by selecting f can lead to a smaller error, but the key observation remains: the error is identical for each task, since the feedforward action f and feedback action Ce do not depend on past errors.

Learning from task to task

The observation that traditional motion controllers lead to a very similar error profile in Figure 4 raises the question: can we learn from past tasks, to improve the performance in the next task, i.e., task $j + 1$? Intuitively, the answer is affirmative: since the error is predictable, it can be compensated for. The practical question is how this can be achieved.

To learn from past tasks, assume that we perform the first task, $j = 0$, with no feedforward, thus $f_0 = 0$. The resulting error during the first task e_0 is then given by $e_0 = Sr$. Here, $S = 1/(1 + GC)$, the so-called sensitivity function, which can be directly derived from Figure 3. Now, consider the following idea. Assume that we measured e_0, but we do not have access to r. What feedforward f_1 should we select to reduce the error e_1? Note from Figure 5 that:

$$e_1 = Sr - GSf_1$$

Next, two key steps are made. First, note that we do not have direct access to Sr, but in fact it was measured in the

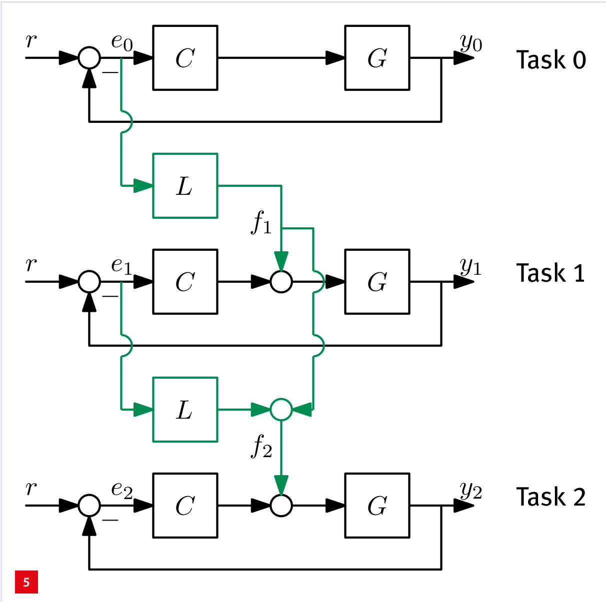

Towards learning from data of previous tasks.

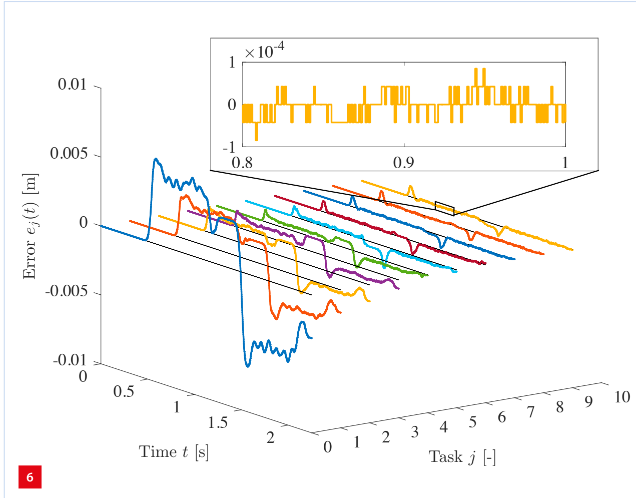

Learning from past data in a printer system: fast convergence to encoder resolution.

earlier experiment: $e_0 = Sr$. Second, let f_1 depend on past errors, for instance:

$$f_1 = Le_0$$

Here, L is a design filter that still has to be specified, see also Figure 5. These steps directly lead to:

$$e_1 = e_0 - GSLe_0$$

This last equation immediately shows that the choice $L = (GS)^{-1}$ leads to $e_1 = 0$.

The update law $f_1 = Le_0$ with $L = (GS)^{-1}$ combines the data e_0 with model knowledge GS. Indeed, L is based on a model of the true closed-loop system GS. The key benefit of learning from data is that an approximate model suffices: of course we cannot expect to have access to an exact model of the system. If the model is not exact, then $GSL \neq 1$; so that e_1 is not zero, but typically much smaller than e_0. The central idea is to repeat the learning procedure in the next task $j = 2$:

$$f_2 = f_1 + Le_1$$

This essentially retains f_1 if it is perfect ($e_1 = 0$), and otherwise includes a small correction based on the already small e_1. This is then also done for future tasks:

$$f_{j+1} = f_j + Le_j$$

This idea of updating the control input is referred to as iterative learning control (ILC), see [2] for a historical overview.

Experimental results

Application of this procedure to the printer system in Figure 1 leads to the measured error signals in Figure 6. These results reveal impressive control performance: the error is at the level of the encoder resolution after only a few tasks. Hence, this very simple learning update leads to extremely high performance by combining data and model knowledge. Interestingly, these performance levels cannot be achieved using traditional feedforward and feedback controllers due to the presence of significant friction in the system; even though the learning update is a simple linear model it can perfectly compensate for these effects.

Can learning beat feedback?

Yes! The results in Figure 6 already reveal extremely high performance, which in practice cannot be achieved using traditional feedforward and feedback. The main reason is that feedback is subject to causality. This is well-known, since in $e_0 = Sr$, the term S cannot be made equal to zero due to the Bode Sensitivity Integral, often referred to by control engineers as the waterbed effect. The fundamental reason this integral exists is due to the fact that the physical system G is causal: it only responds to past outputs. In sharp contrast, in learning, one has access to what will happen in the (near) future due to the simple observation that this has been measured in past tasks. In practice, this is done by designing L to be a non-causal filter; practical details are provided in [3].

Fast and safe learning in the face of uncertainty

The role of model quality for learning

The results in Figure 6 reveal that the feedforward command signals that result from learning substantially increase control performance. In the previous section, it has been argued that the speed of learning depends on the

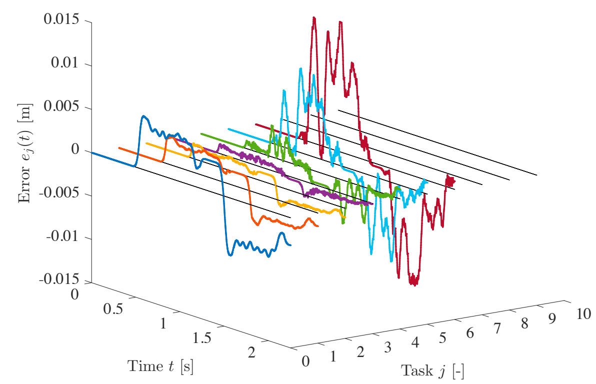

7

Learning from past data in a printer system using a simplified printer model (infinite belt stiffness), leading to growing error signals for subsequent tasks. At iteration 7, safety measures necessitated a shutdown of the system.

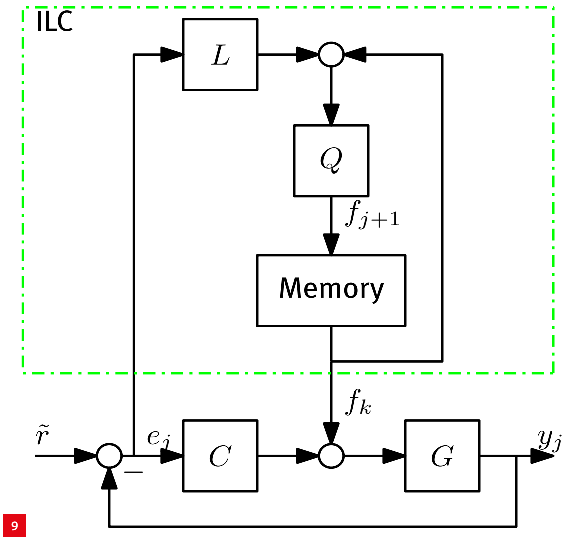

9

Mixed time/task-domain block diagram of iterative learning control, revealing that learning actually is a feedback mechanism.

model quality used to design L. But can it in fact always be guaranteed that the control performance improves?

In Figure 7, the learning procedure of the previous section has been repeated with a slightly different model to determine the learning filter L. In particular, in the experiments of Figure 6, a finite belt stiffness has been assumed, see also Figure 1. In the experiments of Figure 7, a model has been used where the belt is assumed to be infinitely stiff. It is directly observed that in the initial tasks, learning improves performance, but from task 4 onwards, the error actually increases, showing a diverging behaviour until safety measures stop the system at task 7. This can also be seen in Figure 8, where the 2-norm of the error signal is shown for each task, providing a measure of the energy of the error signal. How can feedforward inputs lead to a seemingly unstable system behaviour?

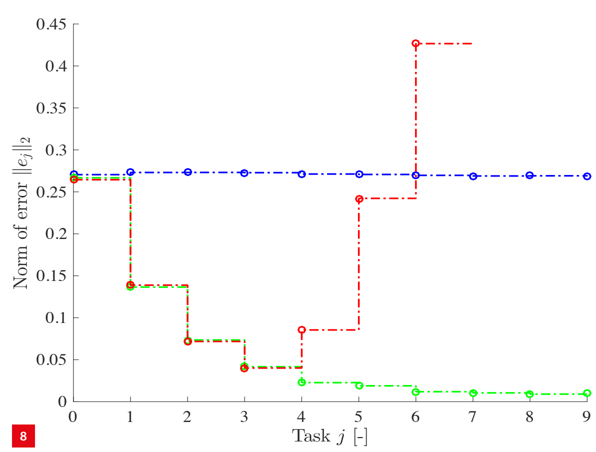

8

Graphs of the 2-norm of the error signal during iteration. Blue: feedback result from Figure 4. Green: learning result from Figure 6. Red: divergent learning behaviour from Figure 7.

Is learning feedforward or feedback?

To understand the behaviour in Figure 7, note that although learning is implemented as feedforward in the time domain, it actually leads to feedback in the task-domain. This can be directly observed in the mixed time/task domain block diagram in Figure 9, where the earlier learning update is obtained if $Q = 1$.

This feedback perspective on learning allows for an explicit analysis of the convergence using control theory. In particular, with the system behaviour $e_j = Sr - GSf_j$ and learning update $f_{j+1} = f_j + Le_j$ it directly follows that:

$$e_{j+1} = (1 - GSL)e_j$$

This type of iteration is ubiquitous in control theory. A very classical result, the Banach fixed-point theorem, implies that this iteration converges in the sense of Figure 8, if the Bode magnitude plot of $(1 - GSL)$ is less than 1 for all frequencies. Thus, convergence, as in Figure 6, 7, and 8, can be directly verified using tools that are traditionally used by mechatronic feedback control engineers. Again, this confirms that learning control in fact is feedback. The feedback perspective on iterative learning from task to task also allows for different choices of L, which can for instance be chosen as a PD controller as in Arimoto ILC approaches [4]. Essentially, this involves a trade-off between required model complexity and convergence speed and behaviour.

Safe learning: the role of robust control

Clearly, when working with physical systems, the diverging behaviour in Figure 7 should be avoided at all cost. This divergence depends on the model that is used for the learning filter L. Control engineers typically have two options in case model errors are too large. First, a better model can be made. Unfortunately, obtaining a model that satisfies the convergence condition for all frequencies requires an extremely high model quality, which is often prohibitively expensive. Second, robustness can be enforced in the design, which is often much more attractive in view of the modelling effort required.

In particular, the field of robust control provides a highly systematic approach for safe learning. Indeed, robustness can be directly enforced by selecting Q in Figure 9. In particular, in case the Bode magnitude plot of $(1 - GSL)$ is not less than 1 for all frequencies, a frequency-dependent Q should be designed such that $Q(1 - GSL)$ is less than 1. Interestingly, this condition can be immediately verified for a set of identified frequency-response functions, which shows a high similarity with traditional mechatronic feedback control design, see [5].

Industrial implementation

The results in Figure 6 reveal an impressive performance improvement. This raises the immediate question: why is learning control not yet standard in industrial mechatronic systems?

Task flexibility

The learning approach outlined in the previous sections assumes that the reference, see Figure 2, is identical for each task. However, in many mechatronic systems the references may change for each task, a typical example being 3D printing. Unfortunately, learning control is highly sensitive for small variations in the reference.

To visualise the troublesome situation, a drawing task has been performed with the 2D industrial flatbed printer in Figure 10. In task 0-4, the goal of the printer is to draw a square. At task 5, the reference is changed to a triangle.

10

Industrial flatbed printer with varying references.

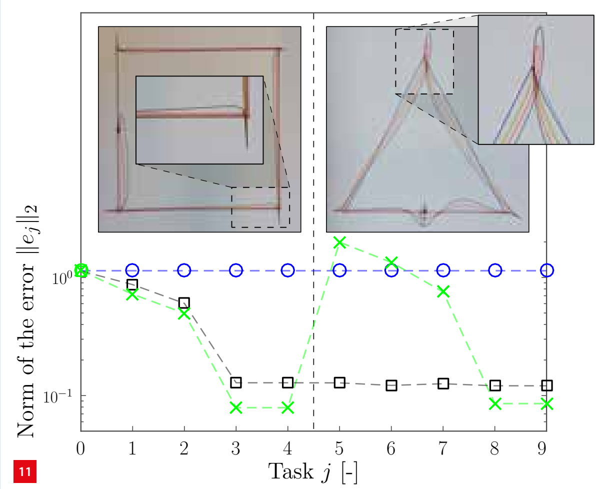

11

Learning with varying references on the 2D flatbed printer in Figure 10. In task 0-4, the goal is to draw a square. From task 5 onwards, the goal is to draw a triangle. Feedback control (blue) leads to mediocre performance. Learning control (green), as described above, leads to an almost perfect square at iteration 3, yet yields very poor performance as soon as the reference changes in task 5. Recently developed algorithms (black) combine task flexibility and high performance through learning.

Clearly, the performance deteriorates significantly, and becomes even worse compared to feedback. Indeed, in case the reference changes each task, it can be shown that feedback outperforms learning.

To address these aspects, learning control with flexibility to tasks has recently been investigated, e.g., in [6]. The key idea is to parameterise f_j such that it extrapolates well with changing references. In Figure 11, the potential is already apparent: both flexibility to varying tasks and high performance are achieved with the new approach.

Learning in complex high-tech systems

High-tech systems are becoming increasingly complex. The example system in Figure 1 only has a single input and output, whereas the system in Figure 10 already has three inputs and three outputs. In many high-tech systems, e.g. in lithography, the entire system may have hundreds of inputs and outputs. This raises the question how learning should be performed, and whether the learning approach described above can be applied sequentially or simultaneously for a set of input-output pairs.

Unfortunately, the naive way of learning for a number of input-output pairs often does not work. In Figure 12, it is shown what happens when the learning approach is naively applied to multiple inputs and outputs of the system in Figure 10 simultaneously. Clearly, this may lead to a diverging error, while the individual loops converge.

Interestingly, this aspect directly connects to multivariable control theory. In [7], a unified framework is developed that

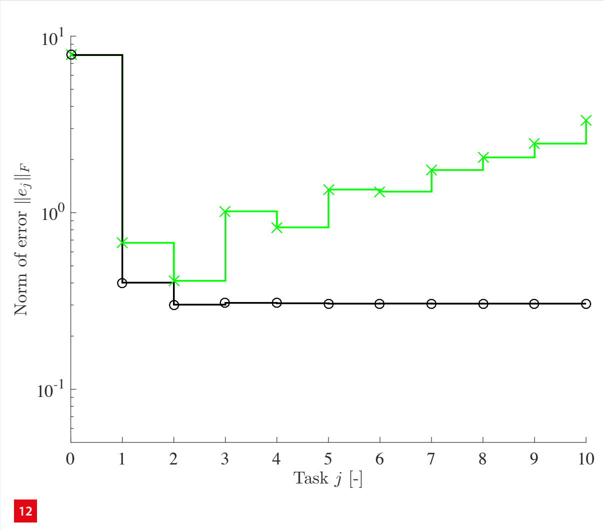

Iterative learning for the printer system in Figure 10 with multiple inputs. Naively applying the learning approach to multiple inputs and outputs simultaneously (green) leads to divergent behaviour. A systematic approach (black), see the subsection on learning in complex high-tech systems, enforces convergent learning behaviour for complex multivariable systems.

allows a systematic design of multivariable learning controllers for complex systems with many inputs and outputs. Interestingly, the approach focuses on a well-balanced use of models and data. The result in Figure 12 confirms that fast and safe learning is achieved for complex systems.

Data-driven intelligent mechatronic systems

What learning has to offer

Learning enables a major performance improvement in machines by exploiting data from past tasks. A general framework for fast and safe learning has been outlined in this article, enabling intelligent mechatronic systems in the near future that can be controlled to the limits of their reproducible behaviour. The role of model-based approaches has been clearly emphasised to achieve fast learning. Control theory is central to achieve safe learning with convergent error signals, which is an essential aspect for learning in physical systems.

A key remaining question is how much performance improvement can be expected with learning? Also, is a classical feedback controller still required? As a general answer, the field of control is able to compensate to the limit of reproducible behaviour of the physical system under consideration. To investigate what learning has to offer for a particular system, consider the following practical procedure. Perform a sequence of n_{exp} experiments with traditional feedback control and optionally feedforward control implemented and measure the error signals e_j, $j = 0, \ldots, n_{\text{exp}}-1$, and compute the sample mean, i.e.:

$$m_e = \frac{1}{n_{\text{exp}}} \sum_{j=0}^{n_{\text{exp}}-1} e_j$$

Learning, as has been outlined in the previous sections, is capable of designing a control input that completely compensates for m_e. The performance that can be expected after learning is thus given by signals $e_j - m_e$, $j = 0, \ldots, n_{\text{exp}}-1$. In this respect, the obtained error at task 10 and beyond in Figure 6 could have been directly predicted from the sample mean of the realisations in Figure 4, where only feedback control is implemented.

The remaining error $e_j - m_e$ is the part that cannot be predicted before the next task starts. Intuitively, feedback control has the task to compensate for these disturbances that occur during the task. Indeed, these disturbances are different each task, but have similar properties for each task, e.g. in terms of their frequency content. It means that as soon as measured data becomes available during the task, a well-tuned feedback controller can effectively address these disturbances. This has been well-known since the advent of optimal control theory in the 1960s: the feedback controller should optimally lead to an error signal which is white noise. In the context of joint learning and feedback, this is investigated in detail in [8]. In conclusion, learning control and a good feedback design are both essential in precision mechatronic systems.

Future developments

In the near future, a further bridge between model-based control and data-based learning is to be expected, which will enable tremendous performance improvements in mechatronic systems. On the one hand, high-tech mechatronic systems are expected to be increasingly complex [9], leading to new learning controllers for multivariable systems [7], unmeasurable performance variables [10], linear parameter-varying dynamics [11], and varying tasks [6]. On the other hand, new developments in control and machine learning will lead to new learning control appoaches, including model-free and reset-free learning [12], kernel-based regression techniques [13], and sparse optimisation [14].

Acknowledgements

This article is the result of a fruitful collaboration with many colleagues in academia and industry, who are gratefully acknowledged. In particular, Frank Boeren, Lennart Blanken, Joost Bolder, Noud Mooren, Robin de Rozario, Nard Strijbosch and Jurgen van Zundert have had a major role in developing and contributing to the ideas outlined here. Financial support is provided by VIDI project number 15698, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO) and ECSEL-2016-1 under grant agreement number 737453 (I-MECH).

REFERENCES

[1] D. Silver, et al., "Mastering the game of Go without human knowledge", *Nature*, vol. 550, pp. 354-359, 2017.

[2] D.A. Bristow, M. Tharayil and A.G. Alleyne, "A Survey of Iterative Learning Control: A Learning-Based Method for High-Performance Tracking Control", *IEEE Control Systems Magazine*, vol. 26, no. 3, pp. 96-114, 2006.

[3] J. van Zundert and T. Oomen, "On Inversion-Based Approaches for Feedforward and ILC", *IFAC Mechatronics*, vol. 50, pp. 282-291, 2018.

[4] S. Arimoto, S. Kawamura and F. Miyazaki, "Bettering Operation of Robots by Learning", *Journal of Robotic Systems*, vol. 1, no. 2, pp. 123-140, 1984.

[5] N. Strijbosch, L. Blanken and T. Oomen, "Frequency Domain Design of Iterative Learning Control and Repetitive Control for Complex Motion Systems", in *IEEJ International Workshop on Sensing, Actuation, Motion Control, and Optimization* (pp. 1-2), Tokyo, Japan, 2018.

[6] L. Blanken, F. Boeren, D. Bruijnen and T. Oomen, "Batch-to-batch rational feedforward control: From iterative learning to identification approaches, with application to a wafer stage", *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 2, pp. 826-837, 2017.

[7] L. Blanken and T. Oomen, "Multivariable iterative learning control design procedures: From decentralized to centralized, illustrated on an industrial printer", *Arxiv: arxiv.org/pdf/1806.08550.pdf*, 2018.

[8] J. van Zundert and T. Oomen, "On Optimal Feedforward and ILC: The Role of Feedback for Optimal Performance and Inferential Control", in *IFAC 2017 World Congress* (pp. 12607-12612), Toulouse, France, 2017.

[9] T. Oomen, "Advanced Motion Control for Precision Mechatronics: Control, Identification, and Learning of Complex Systems", *IEEJ Journal of Industry Applications*, vol. 7, no. 2, pp. 1-14, 2018.

[10] J. Bolder and T. Oomen, "Inferential Iterative Learning Control: A 2D-System Approach", *Automatica*, vol. 71, pp. 247-253, 2016.

[11] R. de Rozario, R. Pelzer, S. Koekebakker and T. Oomen, "Accommodating Trial-Varying Tasks in Iterative Learning Control for LPV Systems, Applied to Printer Sheet Positioning", in *2018 IEEE American Control Conference* (pp. 5213-5218), Milwaukee, Wisconsin, 2018.

[12] R. de Rozario and T. Oomen, "Improving Transient Learning Behaviour in Model-Free Inversion-Based Iterative Control with Application to a Desktop Printer", in *15th International Workshop on Advanced Motion Control* (pp. 455-460), Tokyo, Japan, 2018.

[13] L. Blanken, I. van den Meijdenberg and T. Oomen, "Kernel-Based Regression of Non-Causal Systems for Inverse Model Feedforward Estimation", in *15th International Workshop on Advanced Motion Control* (pp. 461-466), Tokyo, Japan, 2018.

[14] T. Oomen and C.R. Rojas, "Sparse Iterative Learning Control with Application to a Wafer Stage: Achieving Performance, Resource Efficiency, and Task Flexibility", *IFAC Mechatronics*, vol. 47, pp. 134-147, 2017.

Certified training

The content of this article is in part covered by the "Advanced Feedforward Control" training, provided by The High Tech Institute. This course has been certified within the European ECP² framework (see page 60).

WWW.HIGHTECHINSTITUTE.NL/EN/TRAINING/MECHATRONICS

DIVING INTO DEEP LEARNING

Artificial intelligence (AI) technology has made rapid progress over the last few years. Breakthroughs in deep learning set new performance levels for various AI applications including speech recognition, language translation, recommendation and computer vision. It has enabled a first wave of successful consumer applications driven by companies like Google, Uber, Facebook and Netflix. Now a second AI wave is on the horizon, driven by industrial applications. This article gives an overview of the latest AI developments, the main techniques and their application to industrial and control engineering problems.

ALBERT VAN BREEMEN

Introduction

Applying AI to industrial and control engineering problems is not new. During the 70's and 80's expert systems became a popular AI technology to create industrial planning and diagnostics systems [1]. These systems take human expert knowledge and turn it into a database of if-then rules. The control software consists of a database with expert rules and some logic to process these rules ('inference logic'). This approach is particularly successful in domains where the problem cannot be modelled well using first principles and one has to rely on the intuition and knowledge of human experts.

During the early 90's a variation of expert systems using fuzzy logic became popular (see the box). The foundation of fuzzy set theory was laid in 1965 by Lofti Zadeh [1], but only during the late 80's it started to be applied to control engineering problems.

While rule-based AI systems provide a way to translate human expert knowledge into a (control) program, it is at the same time also limited by the knowledge of the human expert. During the mid-90's another AI technology became popular that tackled this problem. This AI technology is named artificial neural networks or just neural networks in short. A neural network is a machine learning AI algorithm that learns nonlinear relationships from data (Figure 1). It can be used as a generic building block to build non-linear adaptive control systems [3]. Applications of neural network-based control systems include nonlinear feedforward control in mechatronic positioning systems and optimised setpoint control [4].

The last few years a new AI technology is rapidly maturing. This technology is called deep learning and it builds upon the foundation of neural networks. The remainder of this article explains this technology and investigates its application to manufacturing and control problems.

AUTHOR'S NOTE

Albert van Breemen is an independent AI & machine learning consultant with over 20 years of experience in high-tech technology and innovation.
He worked for companies including Philips and ASML.

info@aiblog.nl
www.aiblog.nl

Fuzzy logic

While expert systems are based on binary logic (yes, no), fuzzy logic uses truth values between 0 and 1. For example, a human expert might state a knowledge rule like IF temperate-is-high THEN set-heater-low. With binary logic the predicate temperate-is-high is defined by thresholding the temperature sensor data at a specific value. Fuzzy logic, however, defines a range of temperature values, each value being 'high', but with a different truth value. Fuzzy logic provides an intuitive method to design control systems for complex nonlinear processes that is robust to sensor data uncertainty. The technology was popular in Asia and several fuzzy logic-based controllers where used in consumer and industrial applications, such as rice cookers, washing machines and cement mills [2].

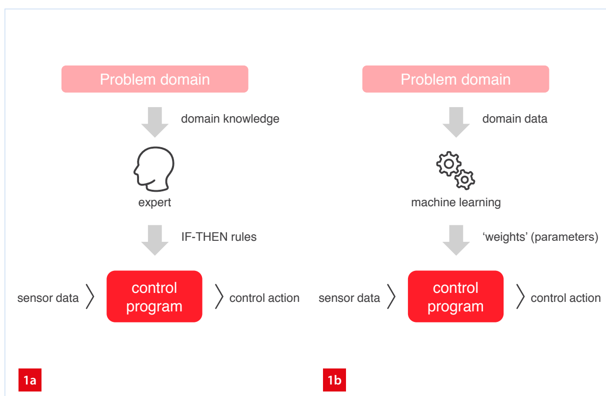

Expert-based design. *Machine learning-based design.*

Deep learning
Deep learning became popular after the 2012 edition of the yearly held 'ImageNet Large Scale Visual Recognition Challenge' (ILSVRC) [5]. During this challenge teams compete to develop the best program to detect objects in images. Up to 2012 the mainstream approach to build these programs was to use human expert knowledge and tuning. Alex Krizhevsky and his team, however, used a neural network with massively more parameters (62 million) than normally used.

As training on a CPU processor would take too long, they used GPUs instead. A GPU has massively more computational cores than a CPU, with each core being able to perform basic multiply-and-addition operations of floating point numbers. As this is the core calculation of a neural network (see further), Krizhevsky et al. could speed up their training significantly.

In the end, they trained a winning neural network on two Nvidia GTX 580 GPUs (1,500 Gflops per GPU) in five to six days. Their solution achieved a 15.4% error rate, which was 10.8 percentage points better than the best system the previous year, and thus won the 2012 ILSVRC. Since then, other teams adopted this approach of using large neural networks, large datasets and GPU training. This approach is now called deep learning and has achieved impressive results in domains such as playing Go [7] and autonomous driving [8].

A perfect AI storm
Currently the field of AI is experiencing a perfect storm. Four technological developments accelerate the development of deep learning, since its conception in 2012 (Figure 2). These include:

1. Data storage & generation:
   Storage capacity in computer systems increased a factor of ~4,000 over the last 20 years. Simultaneously, cheap sensors and Internet connectivity fuel the creation of massive amounts of data; big data. Currently, big datasets of any domain, such as medical, retail, engineering or manufacturing, are freely available on websites such as *Kaggle.com* and *openml.org*, or are present within companies. Deep learning requires large datasets to train models and there is no technology bottleneck anymore for gathering and storing them.
2. Computational power:
   Processor power increased a factor of 1,200,000 over the last 20 years. Companies such as Nvidia and Google have been developing AI-optimised processors such as GPUs (graphical processing unit) and TPUs (tensor processing unit) [9]. These AI optimised processors give a performance boost of 10 to 50 times compared to CPUs.
3. AI algorithms:
   New types of neural networks and machine learning methods have been developed that more efficiently handle large amounts of data. The next section discusses the details of some of those algorithms.
4. Open-source software:
   The availability of open-source software has dramatically shortened the time to develop AI applications. Nowadays anybody has access to major AI software platforms, such as Python/Scikit-Learn [10], TensorFlow [11] or PyTorch [12]. Experimental results can be more easily shared and compared, which speeds up developments.

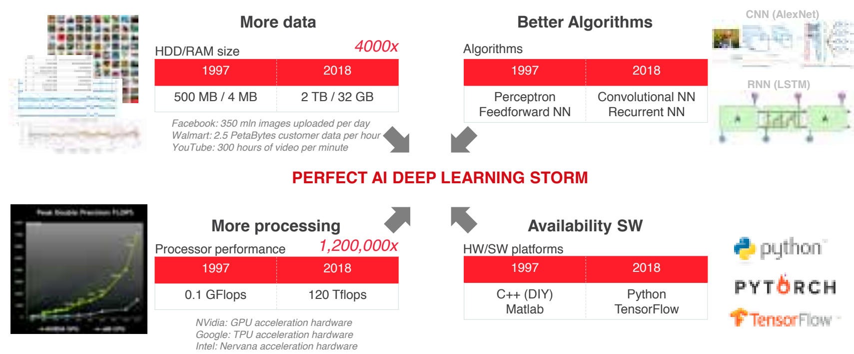

2

Perfect AI deep learning storm due to four technological developments.

Deep learning

Deep learning can be classified into four different types of algorithms, each optimised to handle specific data types and problems.

Artificial neural networks
The main technique behind deep learning is artificial neural networks (ANNs). An ANN is a model that is able to learn nonlinear relations between inputs and outputs. The basic building block of an ANN is a node called perceptron, which is a mathematical model of a biological neuron (see the box on the next page) and was formulated in 1943 by McCulloch & Pitts (referenced in [13]).

Mathematically, a perceptron is defined by the following equation:

$$= f\left(\sum_{i=1}^{n} w_i \cdot x_i + b\right)$$

with:

$$f(z) = \begin{cases} 0 & z < 0 \\ 1 & z \geq 0 \end{cases}$$

Its output y is a function f of the summation of weighted inputs x_i and a bias. The parameters w_i and b (called 'weights' and 'bias') are determined by using a dataset and

Biological neuron

A biological neuron (Figure 3) consists of a cell body, dendrites and axon. Via the dendrites a neuron receives signals from its neighbouring neurons. These signals are added up by the cell body of the neuron. If the total of received signals is higher than some threshold the neuron will fire a signal via its axon to other neighbouring neurons. Not every signal received is as important as other signals. Some signals even inhibit the firing of a neuron. McCulloch & Pitts [13] modelled this by using weights.

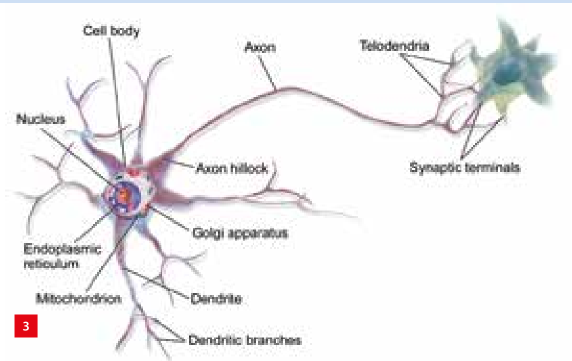

3

Animation of a (multipolar) neuron [14].

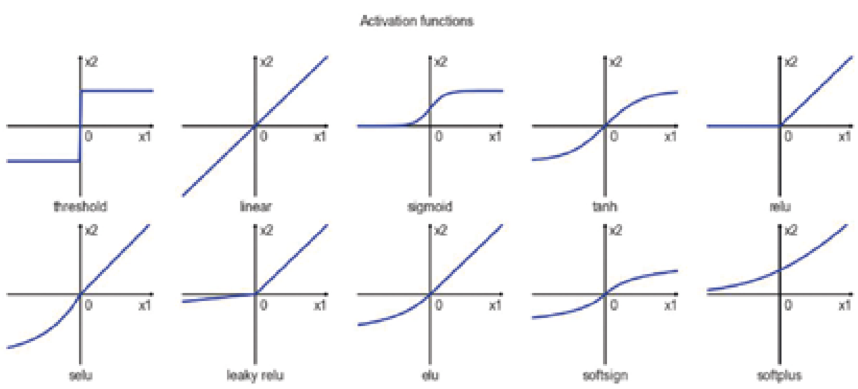

5

Different types of activation functions of a perceptron.

an optimisation method that minimises the prediction error of the perceptron. This is called 'training a neural network'.

One way to interpret the perceptron equation is that it divides a data space linearly into two classes (see Figure 4). The function *f* is called the activation function of the perceptron. Over the years many activation functions have been proposed (see Figure 5) and choosing the best one for a specific problem is one of the ANN design challenges.

A perceptron in itself has limited computational strength. Real-world datasets often are nonlinear and data points from classes are spread over several clusters in the data space (Figure 6). By combining many perceptrons in a layer-like topology (see Figure 7), a model is created that is able to learn more complex functions for problems such as regression and classification.

Convolutional neural networks

Processing images is useful for many applications, including quality control, security, autonomous driving, and augmented reality. While images can be fed into an ANN directly, it would result in neural networks with many inputs and parameters, as every pixel becomes one input of the ANN. A more efficient way is to use neural networks with so-called convolutional layers. A neural network layer consists of a bank of image filters, each having a limited size of 3 x 3 or 5 x 5 parameters. Neural networks where the first layers are convolutional layers, are called convolutional neural networks (CNNs); see Figure 8. Compared to an ANN, a CNN has far less parameters to train.

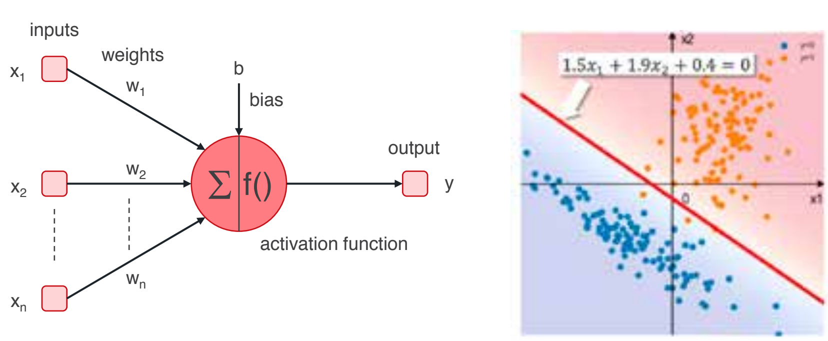

4

Perceptron and an example of a linear separation of a 2D data space.

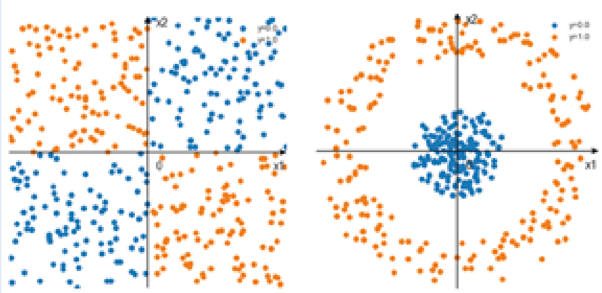

6

Two examples of artificial datasets that cannot be separated by a single line and cannot be learned by a single perceptron.

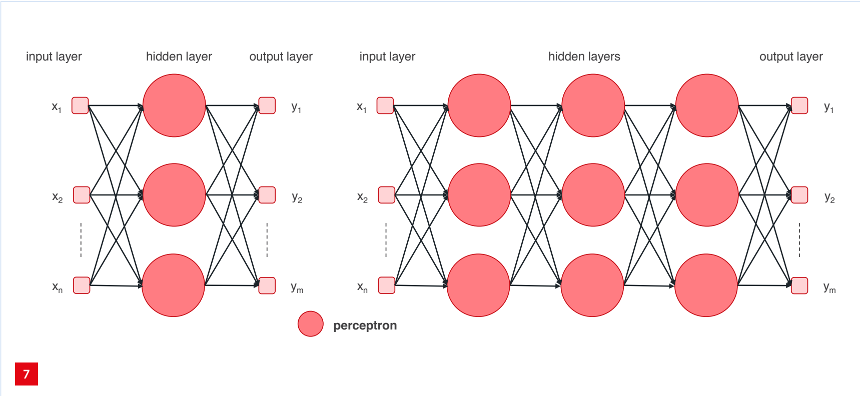

Multi-layer perceptron. *Deep neural network.*

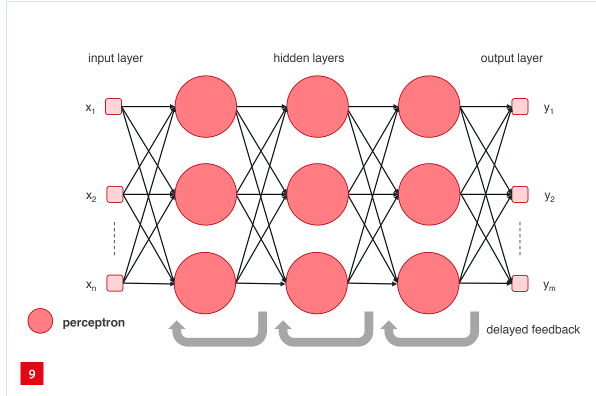

Recurrent neural network.

The more convolutional layers are added to a CNN, the more complex image processing the network can carry out. Over the last few years many different CNN architectures have been developed with names such as AlexNet, ResNet, Inception, and VCC [15]. These networks have been used to solve complex object detection and image segmentation problems. Training the parameters of these networks requires massive amount of data and processing capacity. Without hardware acceleration (i.e. by using AI-algorithm-optimised hardware), training the CNNs becomes impractical.

Recurrent neural networks

A third category of deep learning models are recurrent neural networks (RNNs); Figure 9. These networks differ from ANNs and CNNs in that they feedback a delayed output of nodes in layers as inputs to nodes in previous layers. By doing so, the network becomes 'dynamic' and gets interesting properties to handle time-series data. RNNs can be used for, e.g., speech recognition, video analysis, EEG signal processing, dynamic system identification and stock market analysis.

Although the concept of RNNs was already known in the 90's, no good training method for the weights was available. Because of the feedback loop the training could become unstable, resulting in either parameters vanishing to zero or exploding to infinitely large values. New RNN models called long short-term memory (LSTM) and gated recurrent unit (GRU) have been developed to overcome this problem.

Reinforcement learning

The last deep learning technique discussed here is reinforcement learning (RL) [17]. RL is a technique to determine a sequence of optimal actions for some problem. Where ANNs, CNNs and RNNs require a labelled dataset to do the training, RL uses a single reward signal at the end of some sequence of actions. This type of problem, where a performance signal is only received after some time (control) actions have been carried out, occurs frequently. Think of dynamic machine control, marketing campaigns, process optimisation and game play.

By combining RL with ANN, CNN or RNN powerful learning models can be developed. This has been illustrated by, a.o., the company DeepMind [18], which developed an AI system that beat the ruling world Go champion.

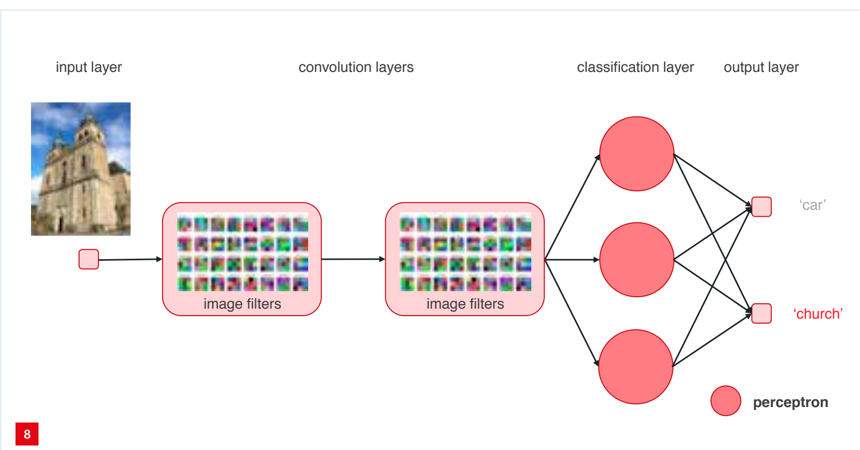

Convolutional neural network.

Applications

Industrial AI

Industrial equipment and manufacturing processes generate huge amounts of data from sensors for monitoring, control and optimisation. Using deep learning AI algorithms could lead to new performance levels and applications to improve product quality, maximise uptime, optimise yield and lower cost of labour, energy and material usage. This has been recognised by some major industrial companies like GE, Bosch, Fanuc, Kuka and Siemens, which are all investing significantly in AI technology riding the wave of 'Industry 4.0', 'smart manufacturing', 'Internet of Things' and more.

Table 1
Comparison of consumer vs industrial AI applications.

Consumer applications	Industrial applications
Meaning of data often clear (e.g. seeing a movie, liking a news item).	Data often noisy and their meaning less clear.
AI runs on the cloud without hard timing deadlines.	AI runs on the edge and real-time responses are expected.
False negatives/positives do not lead to disasters.	Prediction mistakes could lead to unsafe situations.
Predictions may only cost less then €0.001.	High predictions cost €10 - 1,000 as often much more is at stake.
Consumers do not ask why a recommendation was made.	Complex models must be interpretable.

Applying AI technology to industrial problems is more challenging than applying it to consumer applications (see Table 1). Noisy data, real-time performance and safety are among the challenges to deal with in an industrial application. Therefore, industry uses models that can be interpreted by a human expert. While this is possible with rule-based expert systems, deep learning models are more abstract and difficult to interpret. A mix of in-depth engineering domain knowledge and AI knowledge will help in gradually introducing deep learning technology in industrial applications.

Smart cameras
Developing smart cameras for industrial applications is currently one of the low-hanging deep learning fruits. Using CNNs for this purpose has been well researched lately and impressive results have been obtained. This technology is now ready to be used for developing applications in domains such as agriculture, wind mill inspection, medical image analysis and waste inspection.

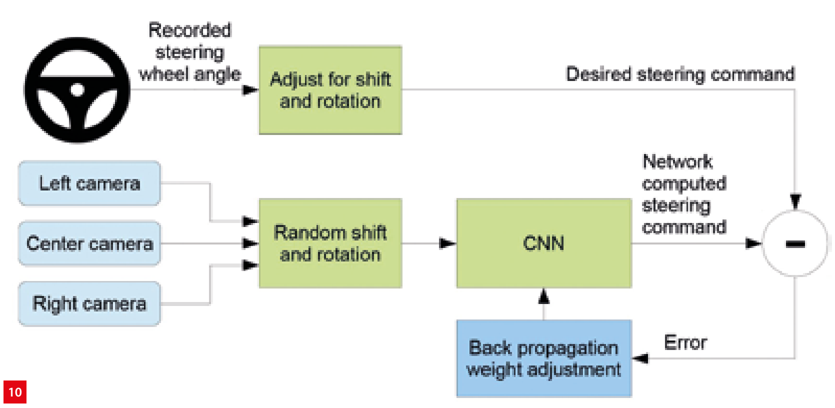

10

An application of CNNs for autonomous driving as described by Nvidia [19].

Besides using CNNs to develop smart object detection systems, they can also be used to develop advanced end-to-end control systems. Nvidia describes an application of CNNs for autonomous driving and trained a CNN for driving a car autonomously (Figure 10). To train the CNN a dataset was created by recording the video footage from front cameras and the steering actions of a driver. With this end-to-end deep learning approach they succeeded in their application [19]. The same approach was also used to create an autonomous flying drone [20].

Maintenance
Another area where CNNs are used is equipment maintenance. By training CNNs to identify machine parts and their health condition augmented reality maintenance systems can be built. Such systems use augmented reality to present relevant information over an image of a machine for the maintenance engineer.

Machine uptime is an important factor determining the overall profit of a manufacturing process. Unexpected machine downtime due to part failure directly leads to additional costs. Predictive maintenance is a method to predict machine part failures before they occur. When a failure is expected within a period of time, the part can be replaced during regular maintenance service within that period.

Digital twin
The concept of a digital twin is often discussed in this context [21]. A digital twin is a digital copy of an asset, system or process. It is used for asset monitoring, predictive maintenance and planning by simulation. A digital twin consists of sensors, data gathering, data storage, data analysis and visualisation. Because very large numbers of digital twins are often managed by a manufacturer, they are implemented using scalable cloud services, such as provided by Amazon Web Services.

The quality of a digital twin depends among others on how well the future behaviour of an asset can be predicted from the data. Deep learning recurrent neural networks are a promising deep learning technique to use, because of the time-series nature of the data to deal with. While traditional approaches depend on manual feature engineering and domain expertise, RNNs automatically extract the right features from time-series data.

Siemens applied the digital twin concept to the application of predictive maintenance of gas turbines. For every turbine data from over 500 sensors are collected. Using AI and virtual reality, engineers can remotely monitor the asset and prevent problems early on [22].

Optimisation
Reducing material and energy usage in manufacturing processes is always an important challenge to tackle, in particular with climate change on the agenda of many companies nowadays. Deep learning models can learn the complex interactions between (parts of) machines and use this knowledge for optimisation.

Google applied deep learning to reduce energy consumption in its data centres. This is a challenging control problem because interaction between the various components is so complex that an intuitive understanding cannot be built. Furthermore, internal and external (like the weather) conditions change so quickly that rules for every scenario could not be derived. Finally, as every data centre has a different architecture a generic solution cannot be used. An energy consumption reduction of 40% was realised by training ANNs on historical data centre operating data. These ANN models could be used to calculate recommended actions to optimise the power usage efficiency [23].

Robotics is another area where deep learning is being explored. Programming manufacturing robots to carry out pick & place operations is a difficult and time-consuming task. Recently, Fanuc showed that this task can be sped up and simplified by using deep RL. Using this AI technique, robots learn by themselves from trial & error to pick & place parts. Within eight hours of learning the robot has achieved a similar performance level as if it where programmed by a human expert [24].

Conclusion
New developments in the area of deep learning are being used in manufacturing and control applications. Major companies including GE, Siemens, Bosch, Kuka and Fanuc are all investing in this technology, with the aim to reduce cost, improve quality and create new applications.

While access to AI algorithms, hardware and software is open to everybody, successful industrial AI applications depend on having domain-specific datasets and talent that understands both AI and manufacturing and control engineering. The latter is a rare breed because both AI and engineering are a specialism requiring deep knowledge. The way to go forwards then is that engineers team up with AI specialists, so that in a next AI-themed *Mikroniek* issue high-end AI applications for mechatronic control can be presented in detail.

REFERENCES

[1] Russell, S., & Norvig, P., *Artificial Intelligence: A Modern Approach*, 3rd edition, 2015.
[2] Singh, H., et al, "Real-Life Applications of Fuzzy Logic", *Advances in Fuzzy Systems*, Volume 2013.
[3] Van Breemen, A.J.N., "Neural Adaptive Control", 1997, *aiblog.nl/neural-adaptive-control*
[4] Van Breemen, A.J.N., & De Vries, T.J.A., "Design and implementation of a room", *Control Engineering Practice*, 9(3):233–248, 2001.
[5] Imagenet Large Scale Visual Recognition Challenge (ILSVRC), *image-net.org/challenges/LSVRC*
[6] Krizhevski, A., et al, "ImageNet Classification with Deep Convolutional Neural Networks", *papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf*
[7] DeepMind, *deepmind.com/research/alphago*
[8] Nvidia autonomous driving, *devblogs.nvidia.com/deep-learning-self-driving-cars*
[9] Google TPU, *cloud.google.com/blog/products/gcp/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu*
[10] Scikit-Learn, *scikit-learn.org*
[11] Tensorflow, *www.tensorflow.org*
[12] PyTorch, *pytorch.org*
[13] Haykin, S., *Neural Networks – A Comprehensive Foundation*, Pearson Prentice Hall, 1999.
[14] Blaus, B., Blausen.com staff (2014), "Medical gallery of Blausen Medical 2014", *WikiJournal of Medicine* 1 (2). *doi:10.15347/wjm/2014.010*, ISSN 2002-4436.
[15] *software.intel.com/en-us/articles/hands-on-ai-part-16-modern-deep-neural-network-architectures-for-image-classification*
[16] *www.kaggle.com/c/challenges-in-representation-learning-facial-expression-recognition-challenge*
[17] Sutton, R.S., & Barto, A.G., *Reinforcement Learning*, MIT Press, 2nd edition, 2018.
[18] *deepmind.com*
[19] Bojarski, M., et al., "End to end learning for self-driving cars", *arxiv.org/abs/1604.07316*, 2016.
[20] *spectrum.ieee.org/automaton/robotics/drones/drone-uses-ai-and-11500-crashes-to-learn-how-to-fly*
[21] GE, digital twin, *www.ge.com/digital/applications/digital-twin*
[22] *www.siemens.com/innovation/en/home/pictures-of-the-future/digitalization-and-software/simulation-and-virtual-reality-simulations-gas-turbines.html*
[23] *deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40*
[24] *www.technologyreview.com/s/601045/this-factory-robot-learns-a-new-job-overnight*

SMART MEASURING

Information-rich metrology refers to the incorporation of any type of available information in the data acquisition and processing pipeline of a measurement process, in order to improve the efficiency and quality of the measurement. In this article, the information-rich metrology paradigm is explored as it is applied to the measurement and characterisation of surface topography. This 'smart' measurement paradigm is illustrated by a wide array of surface metrology applications, ranging from product inspection, to surface classification, to defect identification and to the investigation of advanced manufacturing processes.

NICOLA SENIN AND RICHARD LEACH

Introduction

Information-rich metrology (IRM) refers to the use of any type of additionally-available information to improve a measurement process [1]. Information may come from knowledge of the manufacturing process, knowledge of the object to be measured, and/or knowledge of the physical interactions/principles underlying the measurement technology itself. Information may either come from pre-existing knowledge (i.e. 'a priori'), from mathematical modelling or simulation, or from other measurement processes, even performed concurrently to the measurement one is aiming to improve.

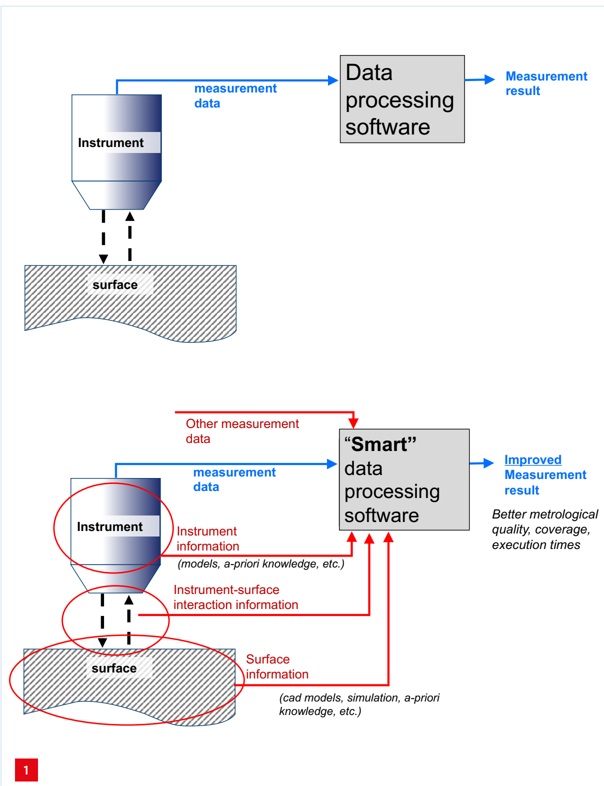

Additional sources of information and changes in information flow when shifting from conventional metrology (top) to the IRM paradigm (bottom).

An overview of how information sources and information flow change when the IRM paradigm is adopted is provided in Figure 1. The idea of using available information related to the product, or process, or product-measurement-instrument interaction, makes intuitive sense because metrology in manufacturing takes place in controlled and very predictable conditions, with a sensible amount of information which is known in advance.

This article illustrates the advantages and challenges of introducing heterogeneous information sources in the surface characterisation pipeline. Examples are provided about the incorporation of structured knowledge about a part nominal geometry, the manufacturing processes with their signature topographic features and set-up parameters, and the measurement instruments with their performance characteristics and behaviour in relation to the specific properties of the surfaces being measured.

Information sources

Measured object and manufacturing process

When a part or product is manufactured, in particular when using digital manufacturing methods, a large amount of information is typically available about the object being produced. For example, CAD data provides information about the nominal form. Analogously, a significant amount of information is available, or can be easily acquired, about the manufacturing process, in terms of its capability, the features and defects it generates, the materials it is designed to operate with, and the types of geometries and surfaces it typically produces.

Most of such information is generated and exploited through product design and manufacturing process planning. In IRM, the aim is for such information to be used to improve metrology, for example, in the inspection and verification of part quality, or in manufacturing process monitoring.

AUTHORS' NOTE

Nicola Senin is an associate professor of Manufacturing at the Department of Engineering, University of Perugia, Italy, and visiting researcher at the Faculty of Engineering, University of Nottingham, UK. Richard Leach is professor of Metrology at the Faculty of Engineering, University of Nottingham.

The authors would like to thank the EPSRC (Grant EP/M008983/1) and the EC (FP7-PEOPLE-MC 624770 METROSURF).

richard.leach@nottingham.ac.uk
www.nottingham.ac.uk/engineering

ACKNOWLEDGEMENT

This article is an abridged version of the paper presented by the authors at the 15th CIRP Conference on Computer Aided Tolerancing (CIRP CAT 2018) in Milan, Italy, and published by Elsevier in Procedia CIRP.

doi.org/10.1016/j.procir.2018.05.003
www.cirpcat2018.polimi.it

Measurement instrument and instrument-surface interaction
One of the most promising paradigms for IRM is based on using additional information about the measurement technology, to develop improved mathematical models that describe the interactions between the measured object and the measuring instrument. In practice, mathematical models that describe physical principles and phenomena underlying many measurement technologies are already available, although one has to be careful that over-simplifications are not abused. In optical measurement, for example, many models have been developed over the last decades [2], to support the theory of focus variation microscopy, coherence scanning interferometry, confocal microscopy, fringe projection, photogrammetry, etc.

Many current commercial optical measurement systems are already making use of complex mathematical models to interpret raw data acquired through their probes. However, because such models aim to be general, they can make very few assumptions about the nature of the surface which will be measured, the material properties that will be encountered, and other factors. Thus, such models are limited in the information they can provide.

The advantage of working in the scenarios typically encountered in manufacturing metrology is that such additional information is often readily available: at the macroscopic scale, there is information about part shape and expected dimensions; at the microscopic scale, there is information about the expected surface texture, and about signature features left by the manufacturing processes. All such information is exploited to a small extent in conventional manufacturing metrology, and is rarely used to develop a better understanding of how measurement instruments interact with surfaces, useful in turn to achieve a better interpretation of measurement raw data.

Smart aggregation of information
The IRM paradigm requires a fundamental re-design of the data analysis processes that are typically adopted in conventional metrology applications. The addition of a potentially high number of heterogeneous information streams raises a whole series of challenges regarding how such information should be homogenised, aggregated and finally exploited towards achieving a better measurement result overall. Recent work on multi-sensor data fusion provides an overview of the challenges and approaches for sensor data aggregation [3-4].

Challenges are in how to handle large amounts of data in increasingly shorter times (possibly verging towards big data issues), in how to data mine the relevant relationships between variables, and finally in how to obtain mathematical and statistical models that ultimately support what can be referred to as the 'smart' measurement paradigm, as opposed to the conventional metrology pipeline of 'blind' processing.

As in many other applications involving big data, a fundamental role in such a paradigm shift may be covered by artificial intelligence (AI) technologies. Machine learning in particular, can provide significant support to the development of the smart measurement solutions of the future (for example, see [5]).

The IRM advantage
Central to the IRM paradigm is the aim to improve measurement quality. Quality is here intended as a generic term encompassing multiple facets: improving quality may mean reducing measurement times, improving measurement performance indicators (accuracy, precision, etc.), expanding the range of covered scales (spatial resolution and range), and improving coverage, intended as the capability to reach surfaces which may be harder to reach, for example measuring beyond the maximum permissible slope for a given measurement technology.

Improving coverage and metrological quality of measurement is a key strategic objective in manufacturing metrology, as many emerging measurement applications (for example, in additive manufacturing), are creating new challenges related to geometric complexity and lack of uniform material properties. Improving measurement speed is essential in many in-process and in-situ measurement applications, as well as the need to overcome the fundamental limits of individual measurement technologies.

Finally, IRM is not only about improving the quality of a measurement, as the information-rich paradigm may also lead to an improved interpretation of the same measurement result.

Information-rich surface metrology
Whilst the previous considerations are general to any metrology application, this paper focuses specifically on the measurement of surface topography, and on what it means for surface metrology to embrace the information-rich paradigm in terms of challenges and new opportunities.

Paradigm shift example
The conventional data-processing pipeline adopted by surface metrology is shown in Figure 2. The pipeline is based on ISO 25178-2 [6] terminology, but equivalent concepts also apply to the older ISO 4287 standard [7]. A form operator (F-operator) is used to level the data set and to remove any trace of the underlying form of the part. An S-filter is used to remove high-frequency noise, and an L-filter is used to separate and remove the waviness

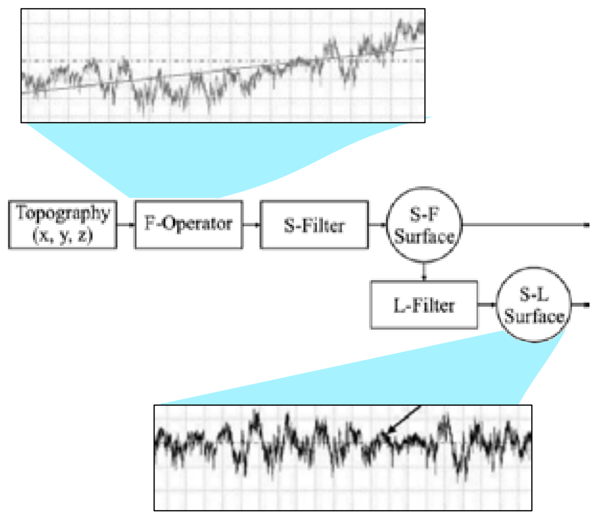

2

Information processing pipeline adopted in conventional surface metrology; example on profile data (adapted from [8]).

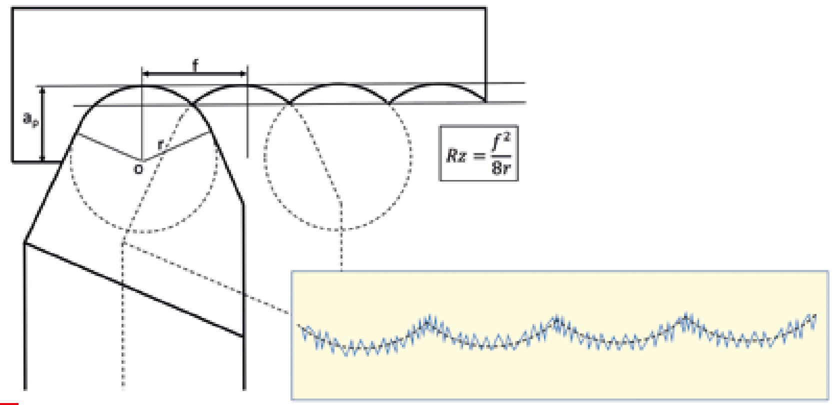

3

Information-rich surface metrology example: use of a topography model from the literature [9] to investigate a measured profile from cylindrical turning.

component. Data processing is designed to make the resulting scale-limited surface (SL-surface) as close as possible to a stationary random signal, suitable to be described by texture parameters that are for a significant part derived from sample statistics.

Very little information is required to apply this procedure: some knowledge of the surface nominal form is required for the F-operator, and previous information about relevant spatial frequencies (typically coming from the manufacturing process and the measuring instrument) is required to choose suitable nesting indices for the S and L filters (cut-off frequencies in the ISO 4287 terminology). The paucity of information requirements is an advantage, as it makes the procedure of very general applicability. But generality is also the main limitation of the procedure, as further case-specific information cannot be exploited to delve deeper into the analysis of measurement data.

An example application of the information-rich paradigm is shown in Figure 3 for a simple case of profile measurement in cylindrical turning. In this case, the expected topography is modelled using a geometrical construction from the literature [9], which relates the spacing, depth and shape of the machining grooves to process parameters, such as feed rate and tool tip geometry. Whilst measurement can proceed in the same way as in the conventional method, what changes is the way the data is analysed: the simulated, expected topography can be subtracted from the measured profile, and then the residuals can be characterised, again possibly with the conventional means of isolating a stationary random signal.

The advantages are immediately visible: it is possible, for example, to investigate aspects, such as the regularity and geometric properties of the machining marks (i.e. how much they deviate from the expected results) and in turn identify effects of machining error at multiple scales (chatter phenomena, oscillations of the workpiece, worn tool, etc.). The price to pay for a potentially much more in-depth investigation is that the method is not generic (it only applies to cylindrical turning), knowledge of nominal manufacturing parameters is required, the whole process of fitting to a nominal geometry and investigating the residuals requires more preparation and is more challenging to implement.

Feature-based representations

The use of modelling to predict topography from manufacturing process parameters, as exemplified in the previous section, introduces the concept that, in IRM, additional information layers pertaining to topography can be added to the characterisation pipeline, for example, where topography itself is described in terms of its constituent features. For the cylindrical turning example, such features are the machining marks, but in general multiple higher-level information overlays can be added to represent additional viewpoints. For example, in Figure 4, an areal topography dataset acquired by an atomic force microscope (AFM) is shown, again representing a cylindrical turned surface, where further overlays (in addition to machining marks), are used to identify scratches from functional life, or artefacts from the measurement process.

Feature-based representation is the term introduced in IRM to refer to the use of additional, higher-level information overlays where topography is partitioned into regions, and the relevant ones are mapped to classes defined within some user-defined ontology. As ontologies may be case-specific

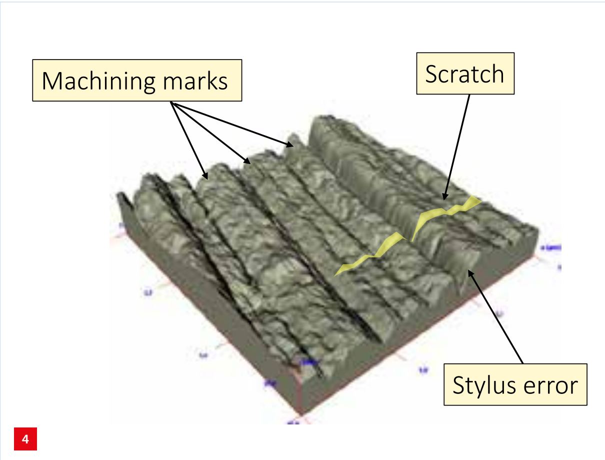

Feature overlays for a cylindrical turning surface measured with AFM.

(i.e. referred to a specific manufacturing process, application or measurement technology), once again, IRM sacrifices generality for depth and breadth of investigation possibilities.

Feature-based characterisation

For 'feature-aware' topography characterisation, a new data-processing pipeline is introduced within the information-rich surface metrology paradigm [10]. This is summarised in Figure 5, and is comprised of the phases of feature identification (the features of interest are identified through matching their shape and size properties to those defined in the ontology of reference), feature extraction (the features of interest are isolated through a partitioning/segmentation of the original dataset, and then extracted as independent geometric entities); and feature characterisation (the feature of interest are described in terms of their relevant shape and size properties).

Feature-based overlays are a core concept of information-rich surface metrology, as they allow mapping of low-level topography information (point cloud or structured grid of height values) to multiple layers of higher-level information,

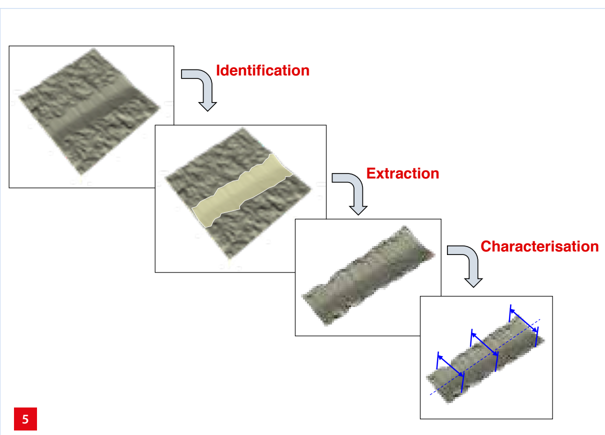

The feature-based characterisation pipeline.

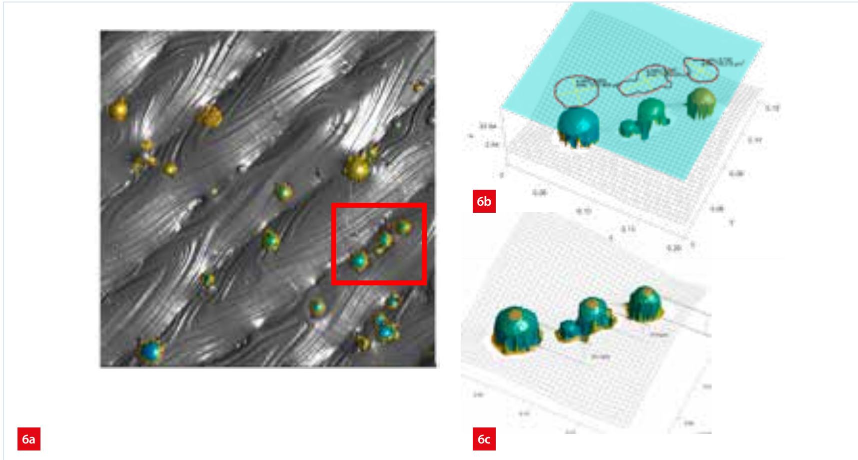

Identification, isolation and characterisation of spatter formations on metallic surface fabricated via laser powder-bed fusion. Measurement obtained via CSI (adapted from [11]).
(a) Identified features.
(b) Footprint area properties.
(c) Feature height properties.

each designed to allow some type of context-specific reasoning, for example, to investigate manufacturing signature features, measurement artefacts or elements of structured surfaces.

An example application of feature-based characterisation is shown in Figure 6 for a metal laser powder-bed fusion (LBPF) surface measured with coherence scanning interferometry (CSI). In Figure 6a, spatter formations are algorithmically identified in the measured dataset; in Figure 6b and 6c, some of such formations are isolated and characterised in terms of footprint area and protruding height from the surroundings [11].

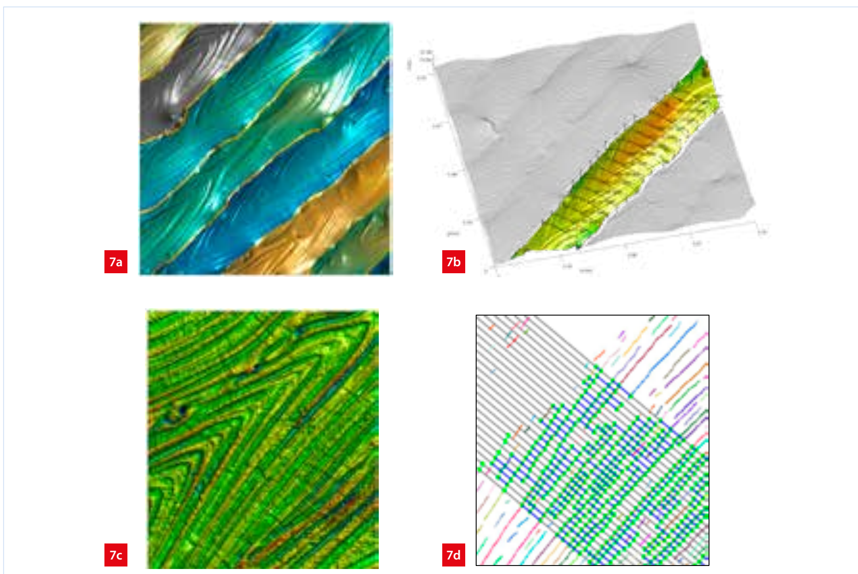

Characterisation of weld tracks and weld ripples on metallic surface fabricated via laser powder-bed fusion. Measurement obtained via CSI (adapted from [11]).
(a) Identified weld tracks.
(b) Cross-section width regularity analysis on isolated weld track.
(c) Detail of weld ripples.
(d) Ripple spacing analysis.

In Figure 7, a similar feature-based characterisation pipeline is used to isolate and characterise LPBF weld tracks and weld ripple spacing [11].

Depending on the degree of determinism of the studied topography, different feature identification and characterisation solutions may be adopted. For example, high variability of shape and size of feature instances suggests the use of statistical modelling tools for shape representation and comparison, the main goal being to pursue robustness to intrinsic variability of feature instances, while still ensuring discrimination of features belonging to different classes. Additional challenges for shape-based reasoning are related to possible lack of information due to sub-optimal sampling density, occlusions, re-entrant portions or too-steep-to-measure portions of the features, all of which being typical issues of micro-scale topography measurement.

Currently investigated approaches for feature identification range from CAD-compare techniques, to the use of a variety of template matching technologies based on shape descriptors (for example, the ring projection transform [12] and the angular radial transform [13]). CAD-compare approaches work well in the presence of highly deterministic structures, e.g., when inspecting micro-parts or products (MEMS, microfluidics) and share significant resemblances with the inspection and verification of standard-sized parts (Figure 8), both in terms of procedural choices and in terms of issues.
However, since typical applications of information-rich surface metrology are at the micro-scale, the availability of surface-specific point sets, akin to what is obtainable from a touch-probe coordinate measuring machine (CMM), is seldom achievable (because of the low market penetration, and challenges of using micro-CMMs [2] [10]), and thus in most circumstances, characterisation proceeds with blanket measurements (typical of range imaging techniques) that require point-set partitioning to isolate the point subsets to fit to each datum [10].

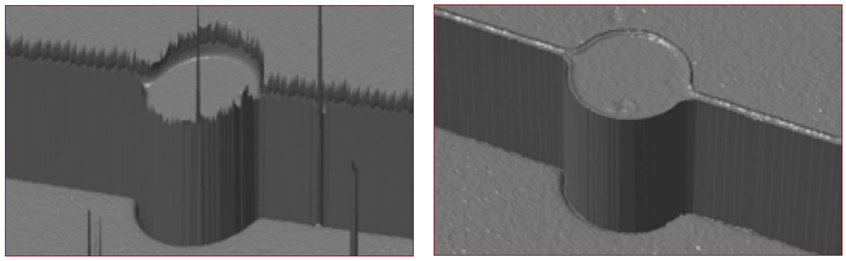

9

Measurement-aware topography data preprocessing example: identification and removal of CSI batwing and spike artefacts from step-like topographic feature (adapted from [14]).

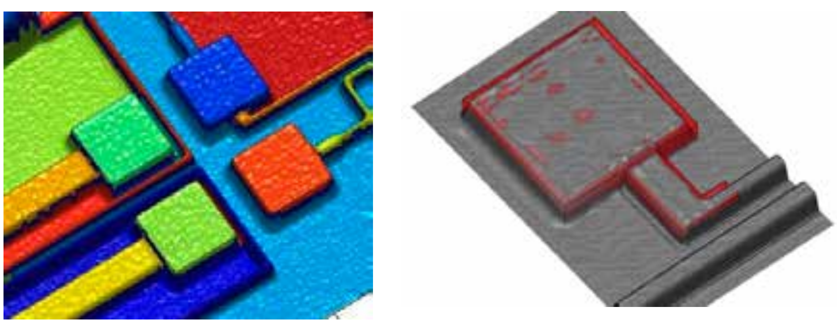

8a 8b

Characterisation of micro-structured elements via CAD-compare techniques (adapted from [10]).
(a) Segmentation.
(b) Volumetric comparison between measured and nominal reference.

Instrument-surface interaction
Another primary venue of investigation in the development of the information-rich surface metrology paradigm, pertains to the incorporation of instrument-related information, and in particular, to the use of models that explain instrument-surface interaction and are thus capable of predicting instrument performance and behaviour when encountering specific topography features. A simple example is shown in Figure 9, where an algorithm is applied, specifically designed to identify and reduce batwing and other spike-like artefacts that appear in CSI measurement in correspondence to abrupt height changes in the topography, as it typically happens when measuring step-like features [14].

The challenge when incorporating knowledge of a specific measurement technology in the surface data-processing pipeline is that, aside from general well-known effects that are clearly recognisable and fairly easy to predict in correspondence of specific topographic features (such as the batwing artefacts mentioned above), a wide range of additional problems are more challenging to spot and handle, as they are related to specific combinations of topographic properties, material properties, and instrument configurations at the time of measurement.

In recent work by the authors, it was shown how the assessment of topographic reconstruction error has a key relevance in contemporary surface metrology [15-16], as measurement error across technologies may sometimes be the same order of magnitude as the features one is trying to measure. The same LPBF region measured via different technologies is shown in Figure 10; recessed features and high-spatial-frequency topographic components are most likely to result in very different reconstructions when acquired with different technologies.
A series of replicate measurements performed in repeatability conditions over the same portion of surface can be used to build statistical models of topography, useful to investigate precision in local height determination, and to determine how precision may depend on the topographic properties of the

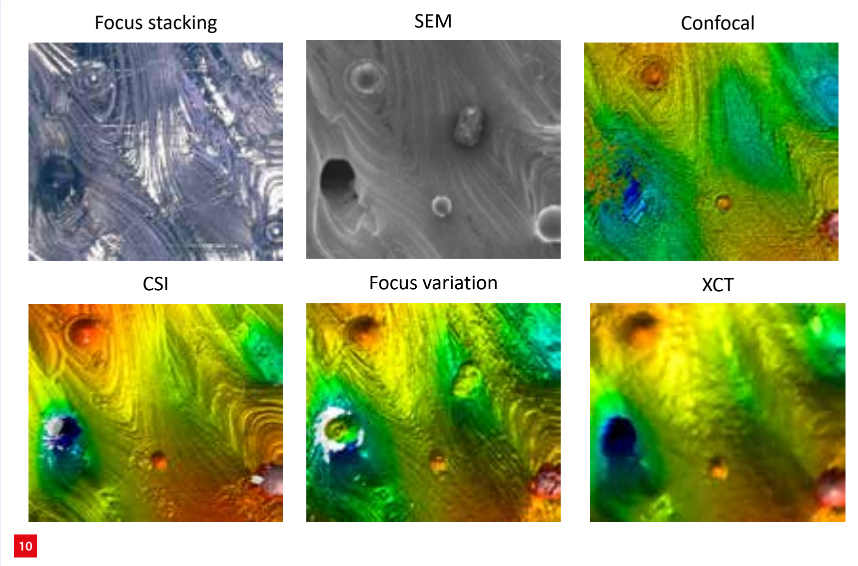

The same topography region reconstructed from single measurement performed with different technologies. Pure 2D imaging results (from optical focus stacking and scanning electron microscopy (SEM) are also shown (adapted from [15]).

surface being measured [15-16]. The same statistical models can be used to identify discrepancies between reconstructions obtained with different measuring instruments [16] and can ultimately give origin to predictor tools to estimate measurement error originated by any measurement technology when used on specific types of surfaces.

The incorporation of measurement error models is a first step towards a measurement-aware approach to feature identification and characterisation, as shape/size information pertaining to the relevant features could be modified to accommodate for variability owing to performance and behaviour of the measurement technology used to acquire information.

Conclusions and outlook

In its attempt to incorporate useful knowledge about the surface, manufacturing process and measurement process within the data-processing pipeline, information-rich surface metrology surely loses generality with respect to the conventional approach to surface characterisation, where only minimal information is necessary, and the same data-processing pipelines can be applied at least in principle to any surface, measured by any instrument. On the contrary, additional, often significant effort is needed in information-rich approaches, to collect, understand and appropriately integrate additional data and models into the data-processing pipeline.

Manufacturing processes evolve and improve over time, as well as the signature features they generate. Measurement instruments also evolve, and so do their performance and behaviour. Customer specifications on what is relevant to measure and to what accuracies and precisions also evolve, as products with increasingly higher value added are designed and produced. At each and every iteration, information-rich approaches require significant extra work, to collect extra data, to develop the appropriate support models, and finally to integrate all the sources of heterogeneous information into a coherent pipeline, ultimately aiming at achieving better metrological performance.

IRM in general, and information-rich surface metrology in particular, pose a series of challenging issues regarding knowledge representation and information handling. Ultimately, the application of the IRM paradigm is far from effortless, and far from straightforward, and may not necessarily be suitable in all manufacturing metrology scenarios. Where it is applicable though, it is asserted that such a significant price to pay is hopefully counterbalanced by the value added to the characterisation results, as dedicated analysis pipelines can be developed that are custom-tailored to specific characterisation requirements, and are capable of providing information that may more directly address specific inspection requests.

REFERENCES

[1] Leach, R.K., et al., "Information-rich metrology: Changing the game", *Commercial Micro Manufacturing* 2017; 8:33-9.
[2] Leach, R.K., *Optical measurement of surface topography*, Berlin: Springer-Verlag; 2011.
[3] Wang, J., et al., "Study of weighted fusion methods for the measurement of surface geometry", *Prec. Eng.* 2017; 47:111-21.
[4] Colosimo, B.M., Pacella, M., and Senin, N., "Multisensor data fusion via gaussian process models for dimensional and geometric verification", *Prec. Eng.* 2015; 40:199-213.
[5] Stavroulakis, P., et al. , "Combined use of a priori data for fast system self-calibration of a non-rigid multi-camera fringe projection system", *Proc. SPIE 2017*; 10330-1033006:1-12.
[6] International Organization for Standardization. ISO 25178-2 geometrical product specification (GPS) - surface texture: Areal - part 2: Terms, definitions and surface texture parameters, 2012.
[7] International Organization for Standardization. ISO 4287 geometrical product specifications (GPS) - surface texture: Profile method - terms, definitions and surface texture parameters, 1997.
[8] Leach, R.K., *Characterisation of areal surface texture*, Heidelberg: Springer Verlag; 2013.
[9] Whitehouse, D.J., *Handbook of surface metrology*, Bristol: Institute of Physics Publishing; 1994.
[10] Senin, N., and Blunt, L., "Characterisation of individual areal features", pp. 179-216, in: Leach, R.K. (ed.), *Characterisation of areal surface texture*, Heidelberg: Springer; 2013.
[11] Senin, N., Thompson, A., and Leach, R.K., "Feature-based characterisation of signature topography in laser powder bed fusion of metals", *Meas. Sci. Technol.* 2018; 29:045009.
[12] Senin, N., Moretti, M., and Blunt, L.A., "Identification of individual features in areal surface topography data by means of template matching and the ring projection transform", *Surf. Topog.: Metrol. Prop.* 2014; 2:14007.
[13] Senin, N., Moretti, M., and Leach, R.K., "Shape descriptors and statistical classification on areal topography data for tile inspection on tessellated surfaces", *Measurement* 2017; 95:82-92.
[14] Senin, N., Blunt, L., and Tolley, M., "The use of areal surface topography analysis for the inspection of micro-fabricated thin foil laser targets for ion acceleration", *Meas. Sci. Tech.* 2012; 23:105004.
[15] Senin, N., Thompson, A., and Leach, R.K., "Characterisation of the topography of metal additive surface features with different measurement technologies", *Meas. Sci. Tech.* 2017; 28:095003.
[16] Thompson, A., et al., "Topography of selectively laser melted surfaces: A comparison of different measurement methods", *Ann. CIRP* 2017; 66:543-6.

FINDING THE FITTEST SPRING MECHANISM

Evolutionary algorithms have been suggested as Darwinian invention machines. By mimicking natural evolution these algorithms have autonomously invented analog electronics, general controllers and robots. In this article we show how evolution can be used to design rigid-body spring mechanisms. The algorithm is able to generate mechanisms that draw a straight line and an ellipse without prior knowledge within 90 minutes.

REINIER KUPPENS

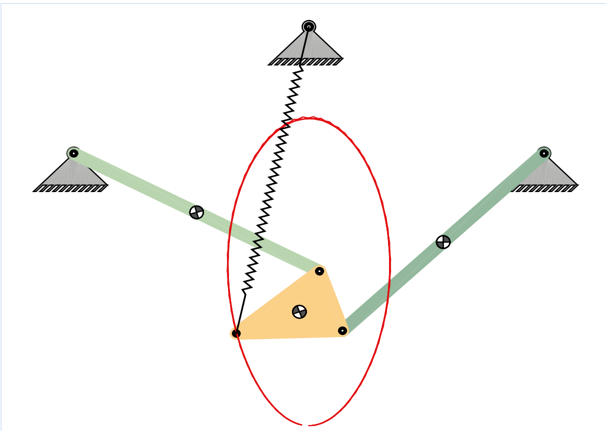

Automation takes centre stage in modern day society. We all know the large manufacturing facilities housing generous amounts of industrial robots, producing a wide variety of products all on their own. More recently, machines have moved out of these deterministic environments to automate more stochastic processes such as video surveillance, product recommendation, personal assistance and even the invention of patentable technology.

Koza et al. [1] showed that by using genetic programming 15 previously patented inventions could be duplicated and improved in the field of electronics. Besides analog electronics, similar algorithms have designed superior wire antennas [2] and general controllers [1], have found natural laws from raw data [3] and have designed robots [4]. All algorithms used in these applications are part of a collection of bio-inspired optimisation methods based on Darwinian evolution and are called evolutionary algorithms (EAs). All of these methods mimic natural evolution in one way or another to harness its sublime creativity.

In this article we show how evolutionary computations can be used to automate the design of rigid-body dynamic mechanisms. We have our algorithm search a discrete and unbounded set of mechanism topologies along with parameters to define hinge locations, mass and spring constants. We show that the algorithm is able to find spring mechanisms that trace a specific path by virtue of the kinematic and dynamic properties without prior knowledge. More details can be found in [5].

Evolutionary design algorithm

Evolutionary algorithms work according to Darwinian evolution (sometimes hybridised with Lamarckian evolution or the Baldwin effect) which is defined by Charles Darwin [6] as one simple but general law: "… multiply, vary, let the strongest live and the weakest die." A computer implementation of this rule maintains a population of initially random solutions that are repeatedly subjected to selection, variation and evaluation until termination conditions are met. An introduction to EAs can be found in [7].

Advantages of EAs are their flexibility, relative simplicity and potential to produce unexpected results. Disadvantages may be that convergence can not be proved mathematically and that computational cost is typically high.

Apart from problem-independent routines, such as the selection of 'parents', a couple of algorithm components need to be tailored specifically for mechanism design. First, an abstract data representation of mechanisms is needed. This representation is essentially a parameterisation of the space of mechanisms and directly determines which mechanisms can be found. One genotype contains all information to construct a mechanism and is a point in this mechanism space.

Second, we need mechanisms to reproduce and mutate. This requires different genotypes to be merged and manipulated, similar to sexual reproduction in nature.

AUTHOR'S NOTE

Reinier Kuppens is a Ph.D. candidate in the Mechatronic System Design research group of the Precision and Microsystems Engineering department at Delft University of Technology, the Netherlands.

p.r.kuppens@tudelft.nl
www.pme.tudelft.nl

Lastly, evaluation requires each mechanism to be simulated and tested according to an objective function to determine its fitness.

For many problems, all of the above components are easily defined. For example, the optimisation of a parametric kinematic model may have a simple vector of real numbers as the genome. Swapping values between two vectors establishes the exchange of genetic information. And behaviour can be determined by substituting the decision variables in the kinematic model.

However, since there exist an infinite number of different types of mechanisms, no such parametric model can be constructed. Consequently, the question arises: how to represent all mechanisms and how to create offspring?

Mechanism DNA

Abstract representations for mechanisms have existed for many years. Early attempts include Reuleaux's symbolic notation [8] and Franke's condensed notation [9]. Although useful for enumerating mechanisms and systematising manual design, computer implementation is non-trivial due to the required visual inspection.

In the 1960s, graph theory was first used to investigate the kinematic structure of mechanisms. Graph-theoretical properties such as planarity, isomorphism and connectivity translate well to mechanisms and can be algorithmically tested on the associated incidence or adjacency matrix.

Essentially, we use a graph to describe which components are connected and how. We have included bodies (B) and a ground (G) connected by hinges (H) and springs (S) as the elementary building blocks. An annotated example of a spring mechanism and its graph are shown in Figure 1. Its incidence matrix I_M is given by Equation 1.

$$\mathrm{I}_M = \begin{array}{c} \\ G \\ B_1 \\ B_2 \\ B_3 \end{array} \begin{array}{c} \begin{array}{ccccc} S_1 & H_1 & H_2 & H_3 & H_4 \end{array} \\ \begin{bmatrix} 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{array} \tag{1}$$

Creating offspring

Creating offspring is essential to evolution. It is responsible for varying and disturbing the information contained by the population. These variations are needed to explore the solution space and move to regions where the best mechanisms can be found.

With practice, graphs will feel as a natural way to represent mechanism topology. However, it is not immediately clear how we can take two mechanisms and mix them to get a new mechanism. Moreover, what do we do with the hinge, spring and mass parameters?

Our solution is to write the topology as a single string of numbers. By using the pattern from Figure 2 we can map each possible column of any incidence matrix to an integer number. Once topology is written as a simple string we can split it anywhere and concatenate it with any other string. Similar to how it is done in the classic genetic algorithm.

The parameters are included as column vectors **d** underneath the topology integers. An example of cross-over on the genetic level is shown in Equation 2 and on the mechanism level in Figure 3.

$$\left[\begin{array}{ccc|cc} 1 & 1 & 3 & 6 & 4 \\ \mathbf{d}_s & \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_h \end{array}\right] + \left[\begin{array}{c|ccc} 2 & 1 & 3 & 2 \\ \mathbf{d}_s & \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_s \end{array}\right] = \left[\begin{array}{cc|ccc} 6 & 4 & 1 & 3 & 2 \\ \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_h & \mathbf{d}_s \end{array}\right] \tag{2}$$

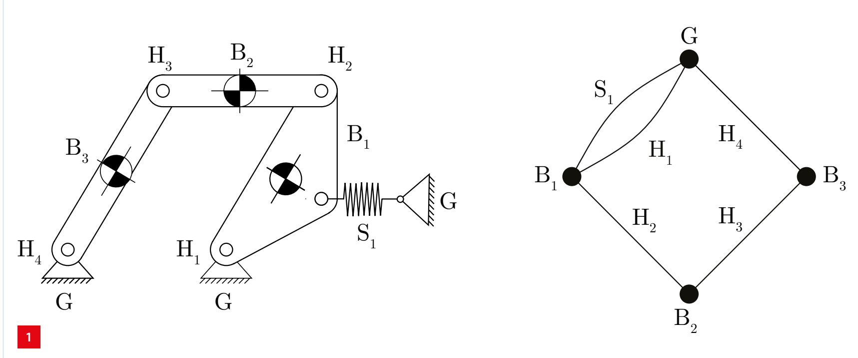

1 A spring mechanism and its graph that describes its topology.

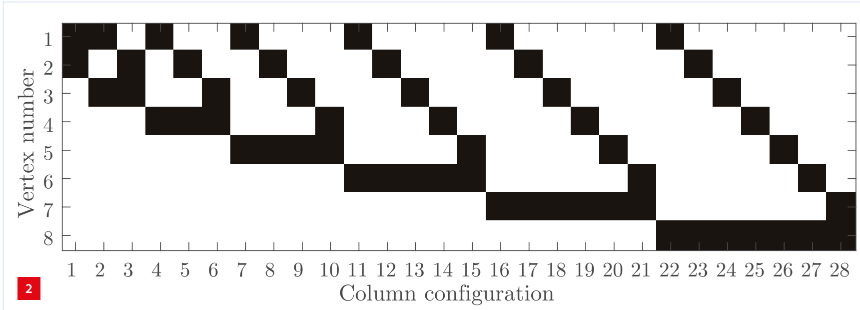

2 The pattern used to map the columns of an incidence matrix to integers. Each column of this pattern is one possible column in an incidence matrix. Black squares indicate ones in a column.

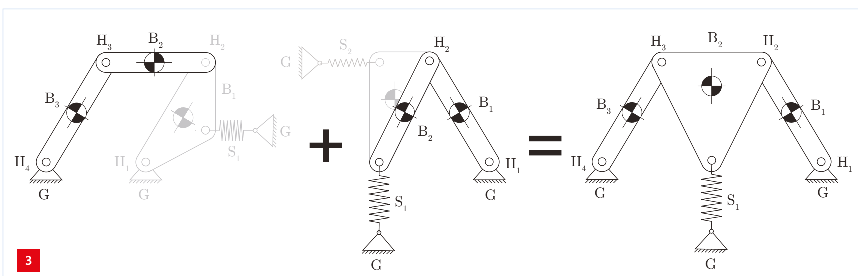

3 Cross-over on the level of mechanisms. The dark parts indicate the selected parts from the genome.

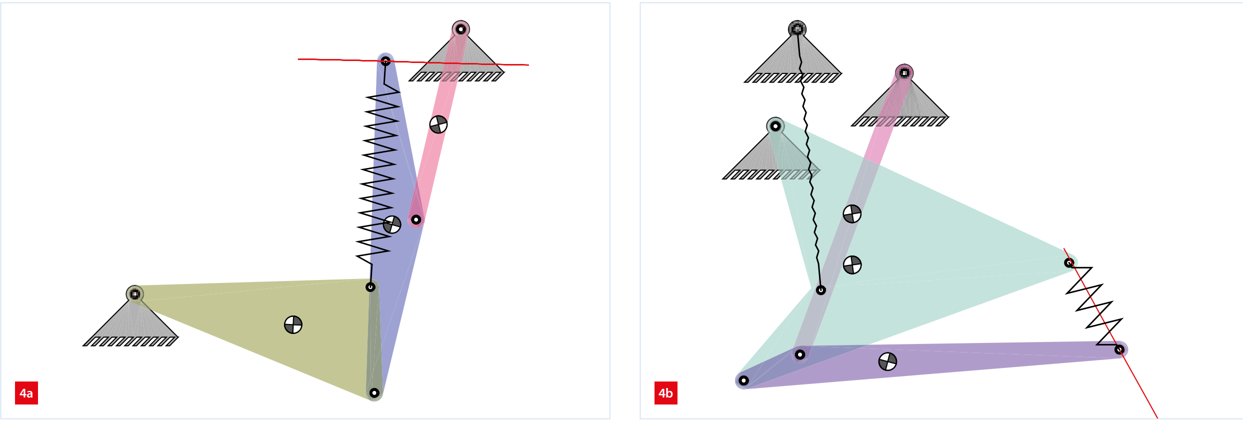

Two examples of evolved straight-line mechanisms.

Evaluation of fitness

Evaluation of a mechanism requires its genome to be converted into a phenotype in some artificial environment. We do so by automatically generating and solving the Newton-Euler equations of motion for each genotype. The solution of this set of differential and algebraic equations gives the dynamic response of our mechanisms under the influence of gravity. The resulting time series of centre-of-mass locations is used by the objective function to compute a mechanism's fitness.
To illustrate the method, two design problems have been solved. The first problem is to find a mechanism that traces a straight line with rotating links. The second problem is to trace an elliptic trajectory. By adding penalties to the fitness, preference is given to mechanisms with a single degree of freedom and a low part count.

Evolved mechanisms

Figure 4 shows two examples of evolved mechanisms for the straight-line problem from different runs. The time series of the *x*- and *y*-locations of their end-effectors are shown in Figure 5a and 5b, respectively. The results for the ellipse problem are shown in Figure 6. All mechanisms were evolved on a desktop computer in under 90 minutes.

All mechanisms have one degree of freedom as desired and oscillate nicely over their approximate straight-line and ellipse trajectories. Most evolved straight-line mechanisms are asymmetric four-bar linkages, not unlike Roberts mechanism. Some solutions will contain seemingly redundant components, such as the mechanism from Figure 6b. However removing these will change behaviour. (Based on the evolved principle from Figure 6b, the figure on page 42 shows a manually designed spring mechanism)

Prospects

Even though the evolved mechanisms shown above perform just simple tasks, they illustrate the potential for automated design once more. In the future, generative design algorithms will become more prominent in the engineering domain. And with the ever increasing available computing power they will solve more complicated problems. Research and industry will focus increasingly on design algorithms, such as Autodesk with their cloud-based generative design project Dreamcatcher [10].

While one may think that these advanced design tools will one day make us engineers obsolete, I argue quite the opposite. These algorithms are tools to more extensively

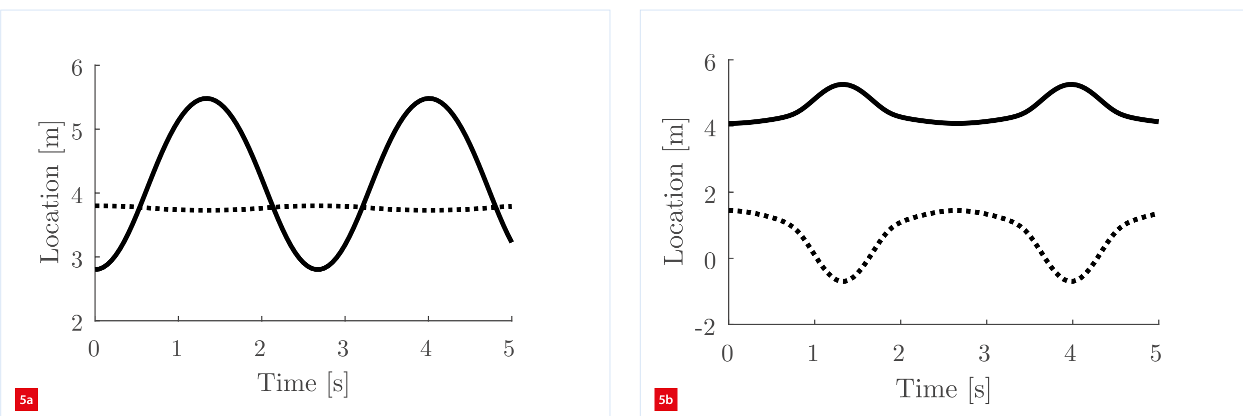

The x- (solid line) and y- (dotted line) coordinates of the resulting trajectory versus time.
(a) The mechanism from Figure 4a.
(b) The mechanism from Figure 4b.

search for conceptual designs. As concepts become more intricate and unanticipated, novel challenges will arise naturally in manufacturing, for example. Not to mention having these algorithms do what one desires is a challenge on its own. We must not forget that these algorithms are simply tools. Tools to augment our own potential. They will free our minds to become more creative on a more abstract level.

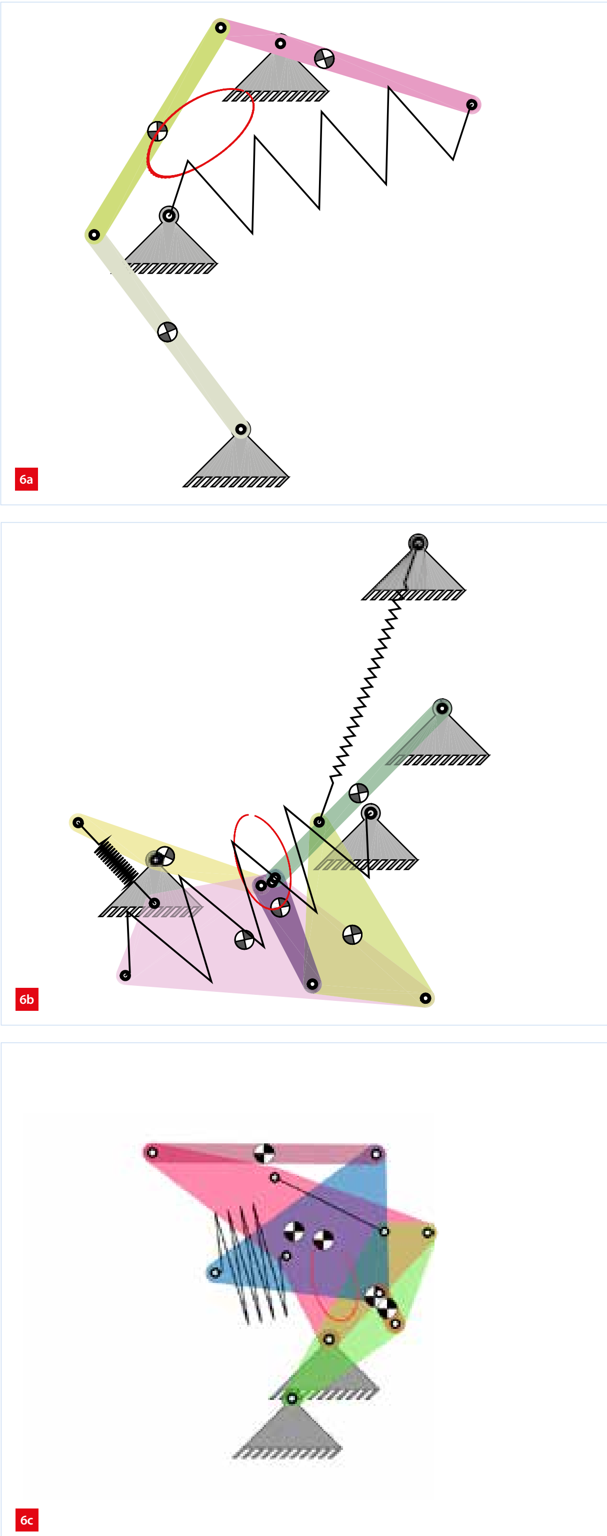

Three examples of evolved ellipse mechanisms.

REFERENCES

[1] J.R. Koza, M.A. Keane, and M.J. Streeter, "Evolving inventions", *Scientific American*, vol. 288, no. 2, pp. 52-59, 2003.
[2] G. Hornby, J. Lohn, and D. Linden, "Computer-automated evolution of an x-band antenna for NASA's space technology 5 mission", *Evolutionary computation*, vol. 19, no. 1, pp. 1-23, 2011.
[3] M. Schmidt, and H. Lipson, "Distilling free-form natural laws from experimental data," *Science* (New York, N.Y.), vol. 324, pp. 81-5, April 2009.
[4] J.B. Pollack, and H. Lipson, "The golem project: Evolving hardware bodies and brains", *Proceedings Second NASA/DoD Workshop on Evolvable Hardware*, pp. 37-42, IEEE, 2000.
[5] P.R. Kuppens and W.J. Wolfslag, "A string-based representation and crossover operator for evolutionary design of dynamical mechanisms", *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1600-1607, 2018.
[6] C. Darwin, *Origin of species*, The Athenaeum, 1869.
[7] A.E. Eiben, and J.E. Smith, *Introduction to evolutionary computing*, Springer Science & Business Media, 2003.
[8] F. Reuleaux, *The Kinematics of Machinery* (reprint), Dover, 1963.
[9] R. Franke, *Vom Aufbau der Getriebe*, VDI-Verlag, 1958.
[10] *autodeskresearch.com/projects/dreamcatcher*

SUPER SOFTWARE TEST TOOL

A model-driven engineering approach is presented based on formal semantics where resulting executable code can be shown to be behaviour-equivalent to the behaviour embedded in the models. This is the ultimate proof that the software tool involved does not introduce problems, providing guarantees towards well-behaving and reliable applications. The methodology focuses on the dynamic control of event-driven, inherently complex systems for which it is very difficult to achieve an acceptable level of test coverage. The approach is aimed at safety-critical domains, but can be beneficial in any domain with complex dynamic behaviour.

JOS HEGGE, MARCEL BEEMSTER AND ROB HOWE

Introduction

ISO 26262, concerned with the functional safety of electrical and/or electronic automotive systems, requires the qualification of tools in use in the development of safety-related software. Essentially it is a risk inventory where a number of factors are estimated:

- The extent to which the tool is used in safety-critical parts.
- The risk that the tool introduces errors which lead to a fault in the product.
- The chance that an introduced error passes by unnoticed.

The qualification flow described here is an answer to the last two bullets, making the tools fit for application in the most safety-critical parts of an application.

The motivation for Verum's model-driven development tooling 'Dezyne' is the ambition to provide formal specification and verification methodologies in a way that is acceptable to regular software engineers. The advantage of applying formal methods for verification of a model is in the completeness of the verification; the final model is mathematically proven to be correct against its requirements for all possible states.

Specifically, Dezyne is targeted at the dynamics or behaviour of event-driven systems since the almost infinite number of possible scenarios expressed by such a system makes it impossible to achieve 100% test coverage by other means. Most formal methods use some kind of mathematical description both for the description of the systems specified and for the requirements (properties) to be verified.

The biggest challenge to general adoption is the high level of expertise required. For most engineers the threshold is too high and hence formal methods are generally discarded. This can be countered by using a model language that is focused on describing systems in a more natural, programming-like style while the language is still based on formal semantics. Using model-to-model transformation a system is translated into a mathematical description that can be formally verified.

Method description

System modeling is focused on describing the event-driven behaviour of a dynamic system. In this activity a system is decomposed into components that operate more or less independently from each other. In Figure 1 the set of components is seen as the 'model'. The externally visible behaviour of the components is specified in interface definitions which are derived from the requirements on desired behaviour. The externally visible behaviour consists of the definitions of input and output events and their communication protocol. In Figure 1 the interfaces are seen as the 'requirements'. The components are only connected via these interfaces. In this way the interfaces isolate the components from each other and allow verification of components individually.

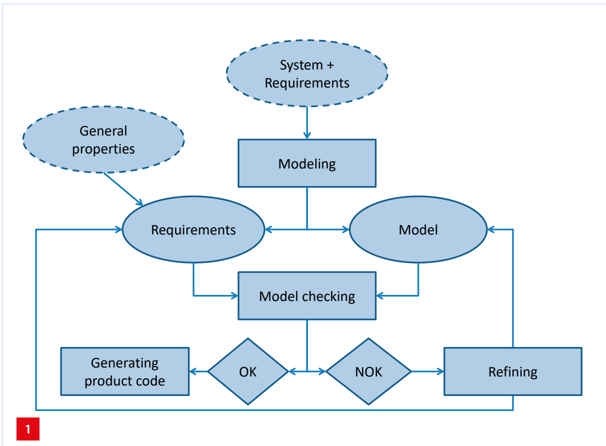

Step 1, the process to come to product code.

AUTHORS' NOTE

Jos Hegge is product manager at Verum Software Tools, based in Waalre, the Netherlands. Marcel Beemster is the CTO of Solid Sands, from Amsterdam, the Netherlands. Rob Howe is the CEO of Verum. Verum develops engineering tools for software-controlled systems. Solid Sands provides a 'one-stop shop' for C and C++ compiler and library testing, validation and safety services.

jos.hegge@verum.com
www.verum.com
www.solidsands.nl

Some properties to be checked during verification are generally applicable and so can be reused from system to system. An example of such a property is 'the absence of deadlocks'. We define these general properties as part of the set of requirements. During model checking we verify that the components behave according to their specifications. Due to the semantics of the modeling language the principle of compositibility applies: properties verified on individual components also hold for systems of components. As a result we can trust that these components will work well together.

In the real world, next to the occurrence of events, we have to deal with data values and we need to incorporate them in our decision logic. However, most decision logic is digital in nature, only at discrete transitions a response should be triggered in the system. (Note: This excludes analog control loops like PID controllers; but these are considered more local attributes.)
This behaviour can be implemented by local functions that translate a change in a data value into an input event. E.g., a speed controller is regularly polling the sensor for its actual value, if it passes a threshold an appropriate event is fired and corresponding actions can be taken. Similarly the system can generate output events that are mapped onto the call of a local function which can, e.g., do specific hardware settings. The local function should have no intrinsic state so it can easily be tested exhaustively in isolation.

Verification

Given the models, the interfaces and the predefined properties to verify, verification is a fully automated process. The following can, e.g., be verified:

- Completeness – in every state of a model for all possible events the corresponding behaviour is defined.
- Determinism – in any given state of a component for any event the component will always respond in the same manner.
- Illegality – in no circumstance an event that is specified as illegal in a given state of a component can fire.
- Deadlock – a component does not have deadlock situations.
- Livelock – similarly a component does not have livelock situations. A livelock is very similar to a deadlock: in a deadlock the system seems to be doing nothing at all, whereas in a livelock the system is still doing lots of things but is simply unresponsive to its client.
- Compliance – a component exactly implements the behaviour specified in its provided interfaces.

A system modelled in the way described above essentially covers all the logic decisions in the system. It is entirely event-driven and behaviour is expressed as sequences of events. The verification checks on every possible sequence of events (i.e. scenarios or use cases) and verifies whether all properties mentioned above hold for all of these sequences. This is expressed as 100% behaviour coverage.

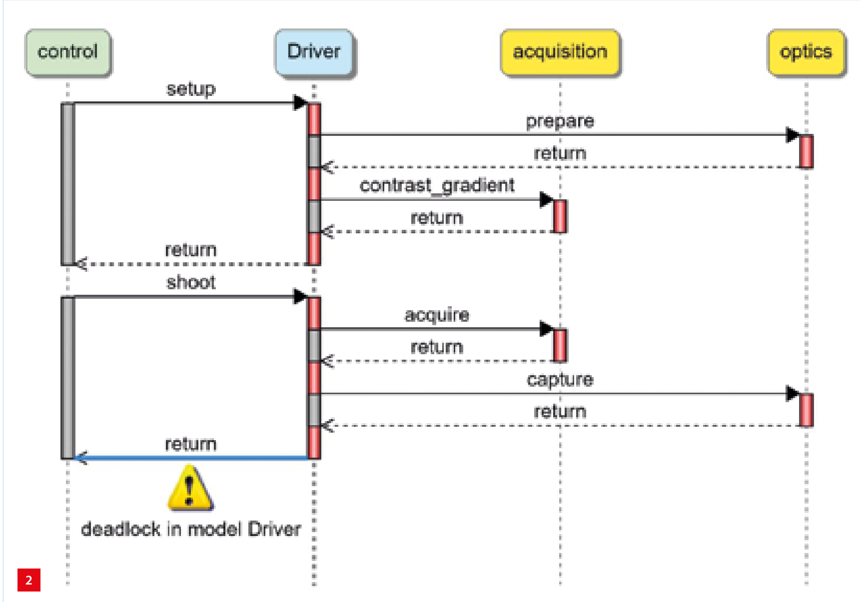

Example of an events sequence diagram (or message sequence chart).

Executable specifications

Functional correctness of the system follows from the validation of the high-level requirements. However, not all requirements can be expressed in the formal description language and hence formally verified. Forasmuch these requirements are expressed as use cases on the system they can be simulated on the models and the resulting behaviour compared to the expected behaviour. A way to visualise the behaviour is to depict it in a message sequence chart (Figure 2). Another important use of the simulation is visualisation of problems found during formal verification. Since we focus on the dynamic behaviour of a system the root cause of a problem may be far back in the past. Showing the execution trace that eventually led to the problem is a perfect way to support analysis and resolution.

Generating executable code

The final step in the process is to generate code that can execute on the product hardware. If executable code would have to be derived manually from the models there would again be a large risk of introducing errors invalidating the benefits from formal verification.
An intermediate step could be a general purpose programming language conforming to automotive standards, like Misra-C. An important aspect of this process is to ensure that the generated code does not have to be modified by engineers to add specific functionality (e.g. low-level drivers and boundary functions), but the additions can remain separate entities. Hence the additions cannot break the guarantees provided by the formal verification.

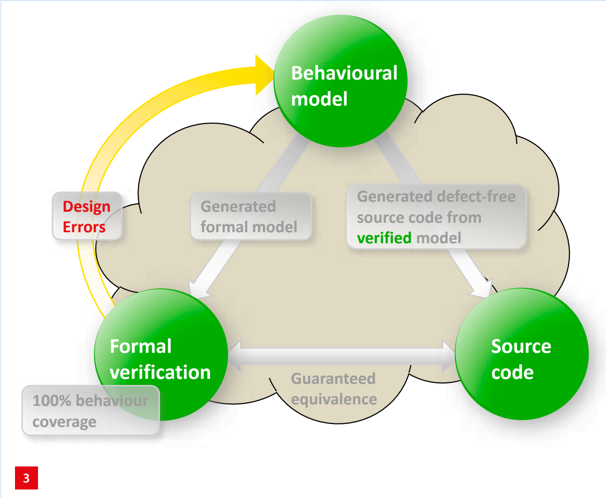

Magic triangle of behaviour-equivalence.

Proving behaviour-equivalence

Our methodology recognises a so-called magic triangle where user-defined models, mathematical models and generated standard programming code are behaviour-equivalent (Figure 3). This means that the properties checked in model verification can be projected on the code that runs on the product hardware.

The equivalence is confirmed by testing on all elements of the Dezyne modelling language. A large set of test models is the starting point for formal verification and code generation. This test-suite covers every element in the Dezyne modeling language. For every model the verification shows either 100% correctness or delivers fault traces. As a side effect of verification we can generate the set of all possible execution traces. From the generated code we build an executable program. We feed the set of all possible traces to the executable by firing the input events one by one into the program. We capture output events and compare them to output events in the trace.
If we can exactly play back all traces on the executable the result is accepted as behaviour-equivalent to the formal model. This is depicted in Figure 4. As starting point the process from Figure 1 was taken where some elements were left out for clarity. The added elements are shown in a different colour.

Based on the test-suite results we guarantee the validity of the magic triangle and hence the equivalence of the generated code to the behaviour expressed in the models. For rigid qualification demands this might not be enough; just like qualification of a compiler which cannot be done in general but only in an actual usage situation. For such situations we can run the customer-specific models through our magic-triangle confirmation flow and show that the behaviour of the generated code is exactly the same as demonstrated in the models.

Integrating compiler qualification

Solid Sands has a tool that extracts MC/DC (modified condition/decision coverage, to measure structural/statement coverage and branch coverage) information from general C and C++ applications (such as the compiler). It then compares coverage results between reference results (obtained by running Solid Sands' SuperTest test and validation suite for C and C++ compilers) and application results.
The application in this case is the end-user application that is generated by Dezyne and then compiled. Four classes are distinguished in the comparison:

- Coverage by the test-suite (reference) and the application.
- No coverage by the test-suite and also not by the application.
- Coverage by the test-suite but not by the application.
- Coverage by the application but not by the test-suite.

For functional safety, the final class is the interesting one. It indicates that the reference is incomplete with respect to the application. This is a potential safety hazard because the application is then using untested code. Thus, statement coverage that is in the fourth class needs to be further analysed.

Compilers are different from regular application code because they are much more integrated and are hugely more complicated. Also they have compile-time configurations and run-time options, both of which influence the internal behaviour of the compiler in many ways. Compilers also contain asserts and other developer-oriented code that is never expected to be executed when it is used by an end-user. For these reasons, it is not expected that a test-suite will come even near 100% coverage for a compiler. This is in fact the reason why the results of the coverage analysis tool are interesting for the fourth class of coverage comparison.

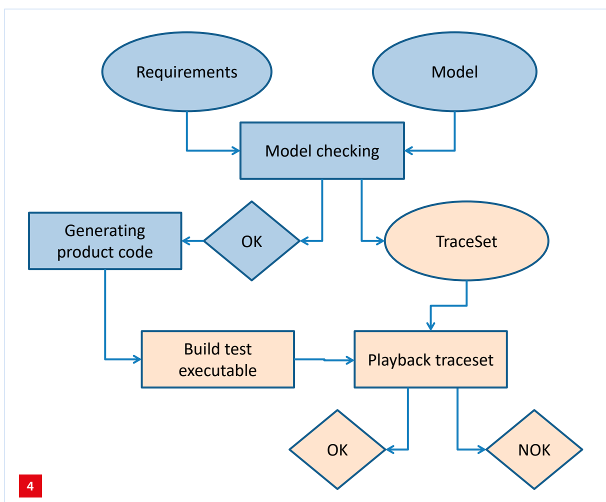

Step 2, the process to prove model-code behaviour-equivalence.

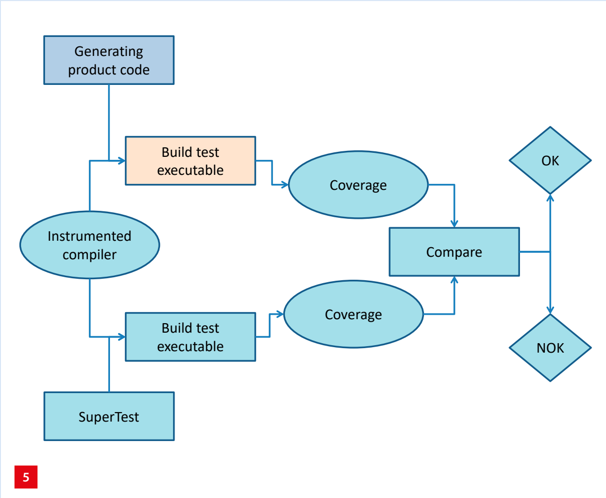

Step 3, the process of compiler qualification.

This is depicted in Figure 5. This builds on top of Figure 4, where again most elements were left out for clarity, and the added elements are shown in a different colour.
In the 'good' case the code generated from a Dezyne model does not trigger compiler statements not covered by SuperTest.

Results

We analysed Dezyne-generated code for all models in the Dezyne test-suite on compiler coverage for both the (tiny) TCC compiler and the sophisticated GCC compiler. The results of this are both interesting and promising.

For TCC, SuperTest by itself achieved statement coverage of 70.1% out of about ten-thousand executable lines. This is for a single configuration (compile time and options) of TCC. It is likely that that number can increase by selecting additional configurations.

For the generated Dezyne code, we combined all of the different source code from the full suite of about ten end-user applications. This code achieved 51.1% coverage.

Comparing these results, we found that 0.05% of the code fell into the class "Coverage by the application but not by the test-suite". This is just five lines in total. This is a very promising result for SuperTest because it means that SuperTest has 'missed' just a few lines.

Further analysis shows that these missed cases are the results of warnings issued by the compiler for certain type conversions in the code generated by Dezyne. The tests in SuperTest are aimed to be 'clean' C-code, that is either free of warnings or containing a clearly demonstrated error (as part of diagnostic testing). This design policy helps end users to understand and analyse tests because they do not fall into the 'grey' area of 'correct-but-not-so-nice' code, for which warnings may be generated. It appears that to achieve the required coverage, we will have to revisit this policy.

The GCC results obtained by running the Verum test-suite confirm the results of the TCC run. In total for the 'Implicit On' application from the Verum apps, 107 lines are reported that are covered by the Verum end-user app, but not by SuperTest. Relative to the total size of GCC, this is a low number. For manual analysis, however, the number is quite high, and it is complicated to analyse these. Initial findings do point in the same direction as for the TCC run. Out of 20 source files in which "not SuperTest covered" lines are present, about 10 are directly related to diagnostic reporting. This also suggests that compiler warnings are responsible for a significant part of the 107 missing lines.
A future area of development is to reduce the effort in analysing the missing coverage lines, for example with the open-source 'reduce' tool.

As a final step we executed the full qualification flow on a collection of models from a customer using GCC. We found an issue not covered by SuperTest yet, which was related to a very specific corner case of GCC. Other than that earlier results were fully confirmed.
For executing the flow on a customer case we might need to use the specific compiler in use. A caveat here is however that the compiler needs to be instrumented to be able to retrieve coverage data from it. If a customer uses a different compiler than GCC or TCC this might be an issue.

Conclusion

We presented a model-driven engineering approach based on formal semantics where resulting executable code can be shown to be behaviour-equivalent to the behaviour embedded in the models. We demonstrated that due to the formal semantics we can transform the user-defined models into a set of mathematical models for automatic verification on a large number of properties. This contributes to the quality of the resulting application. Finally we showed how we can qualify a compiler in its application on generated code. This provides evidence that the compiler does not introduce problems.

Using this model-driven approach it is possible to achieve strict quality standards supported by qualification evidence on both the resulting code and applied compiler tooling.

REFERENCES

- Beusekom, R., Groote, J.F., Hoogendijk, P., Howe, R., Wesselink, W., Wieringa, R., and Willemse, T.A.C., "Formalising the Dezyne modeling language in mCRL2", in *Lecture Notes in Computer Science*, volume 10471, Springer, 2017. (best paper award FMICS&AVoCS)
- Groote, J.F., and Mousavi, M.R., *Modeling and Analysis of Communicating Systems*, MIT Press, Cambridge, 2014.

FIZYR'S PHILOSOPHY

With e-commerce growing sky-high, millions and millions of parcels have to be handled in ever shorter times. Human labour is falling short and conventional bin-picking robots fail in the high-speed handling of frequently changing and individually varying stock-keeping units (SKUs). This 'deadlock' cannot be solved by conventional programming techniques of 'teaching in' each SKU from images or 3D CAD models. Fizyr from Delft, the Netherlands, is providing a solution by applying artificial intelligence (AI) in machine vision to enable robots to cope with variation.

Profile

In 2014 Delft Robotics, a spin-off from the Delft University of Technology research group of Martijn Wisse, now professor of Biorobotics, was established and started doing robot integration projects, usually including computer/machine vision solutions. It embraced deep learning, a kind of artificial intelligence (AI), in order to make vision more intelligent.
A Delft team, in which Delft Robotics participated, winning the prestigious Amazon Picking Challenge in 2016 (Figure 1) made the company realise that its unique expertise in computer/machine vision had an enormous potential for applications like order picking. This led to a strategic shift from hardware/software to software-only and a change in the business model, symbolised by a new name, Fizyr.
This company now focuses on developing AI-based item and parcel picking solutions for global logistic integrators and helping leading automators to apply autonomous robotics.

The main market focus is now on logistics, in particular the parcel delivery 'industry' ("big end-users in online retail, warehousing, parcel distribution and courier/express"), but Fizyr has an eye on other promising markets, such as agriculture & food and the manufacturing industry. The main reasons for this current focus on the logistics market are the generic character of the solutions that can be applied there, and of course the booming e-commerce business. Thus far, projects in the agriculture & food market have mostly been one-offs; as yet there is no model for generating significant recurring business.
In all cases, Fizyr provides solutions in collaboration with suppliers/integrators of material handling equipment and robot vendors/integrators. Here, Fizyr's previous experience as a robot integrator is an additional benefit. The company is growing fast and is now planning its first investment round, to raise 1.5 million Euros for facilitating further growth, mainly to attract more 'brain power'.
In the spirit of the open-source code community, Fizyr has posted part of its applications and algorithms on the Github platform.

WWW.FIZYR.NL
WWW.GITHUB.COM/FIZYR
WWW.AMAZONPICKINGCHALLENGE.ORG

1 *In 2016, the Delft team, a collaboration between the TU Delft Robotics Institute and the Delft Robotics company, won the Amazon Picking Challenge. This challenge aims to strengthen the ties between the industrial and academic robotic communities and promote shared and open solutions to some of the major problems in unstructured automation.*

EDITORIAL NOTE

This article was based on an interview with Herbert ten Have, CEO of Fizyr, and the information on the Fizyr website.

E-commerce is a typical offspring from the digital revolution, but human labour still accounts for over half of the cost in warehousing. However, the work of warehouse operators is physically demanding and in the current labour market they are hard to find and hard to retain. On the other hand, human intelligence is still needed to efficiently handle objects that display a high variety in type, shape, weight, material, colour, texture, transparency, deformability, orientation, stacking, etc. under varying (lighting) conditions. Until recently, computers lacked this intelligence.

Algorithm

To solve this problem, Fizyr has added artificial intelligence (AI) to vision-guided robotics. This comes in the form of

An example of a flexible gripper, designed by Fizyr for 3D-printing.

deep learning capabilities embedded in a so-called neural network. The network has been trained in classifying objects and finding the best possible grasp locations (for picking) using a set of (over 1 million!) real-life images which have been labelled with a large number of features that describe the properties relevant to picking operations.

This massive training set has already significantly increased the quality of the algorithm underlying the neural network for any particular situation. The accuracy and robustness can be further improved, i.e. increasing the success rate of finding the correct grasp locations for each specific implementation, by additional training of the network using application-specific sets of images. As a result, the vision algorithm (processing the point cloud information generated by a camera) can robustly deal with any kind of (unexpected) variation in the items to be picked and in the picking conditions. Based on the subsequent segmentation and classification of all objects, the software can define the best grasp strategy.

The underlying algorithm is built on TensorFlow™, an open-source software library for high-performance numerical computation. Its flexible architecture allows for easy deployment of computation across a variety of platforms and it comes with strong support for machine learning and deep learning. On this solid foundation, Fizyr has developed its widely appreciated Keras-RetinaNet implementation for object detection. Keras is an open-source neural network library written in Python; it is capable of running on top of TensorFlow.

Hardware-independent

A key aspect of Fizyr's philosophy is hardware independency. Users should have the flexibility to select the tools that best fit their application, meaning that the vision software can work with off-the-shelf hardware. This includes:

- sensors: cameras (2D, 3D, RGBD, multi- or hyper-spectral), laser scanners, etc.;
- robots: all the various brands and types, including robot arms, delta robots and collaborative versions (cobots);
- end-effectors, such as grippers which employ different technologies (mechanical, pneumatic);
- computers: conventional PCs with standard GPUs, like Nvidia, running under Linux.

Fizyr is no longer a hardware company, but there is one exception: grippers. Flexible software facilitates the use of flexible grippers. These were, however, not available on the market. Therefore, Fizyr designed flexible grippers that can be 3D-printed (Figure 2). Customers can license a design and print grippers according to their needs.
Deployment of the software is also independent of the type of order picking operation. This means that it works in situations where a 'goods to person' system is replaced by a 'goods to robot' system, including automated storage and retrieval systems, such as conveyor belts, shuttles, automated guided vehicles and autonomous mobile robots. The operation of a robot controlled by the vision software can be easily integrated with a warehouse management system.

Operation

Ultimately, after determining position information of a recognised object, for example by triangulation using a stereo camera, the vision software can provide the grasp locations in six degrees of freedom (6-DoF) of all items to be picked (Figure 3). This is executed within 200 ms, which is much faster than the time a robot needs for physically carrying out the pick & place operation. Combining this 6-DoF information with the gripper-vision coordination, the robot can then pick items one by one without any further operator guidance required. After picking, the 6-DoF information can be used to adjust the parcel orientation as required at its destination.
The machine vision software can also be used for performing classifications and manipulations, as well as for inspection and quality control. The algorithm can be trained to detect defects that are unknown beforehand, such as damaged or broken items, and to carry out grading of quality. That's another quality of Fizyr's software.

A typical application environment for Fizyr's AI-supported vision-based picking software: picking parcels from a roller container, to place them on a conveyor (or vice versa).

GENERIC ROS-BASED ARCHITECTURE FOR DIRTY JOBS

Autonomous systems are systems that can sense, decide and act in a self-governing manner, without a direct human command. Their practical benefit is that they can relieve humans from dirty, dull or dangerous tasks, but in order to do so, they must show dependable, robust and safe autonomous behaviour. Here, a generic architecture for autonomous service vehicles is proposed, based on an open-source standard, ROS. In addition, a number of applications developed by Nobleo are described.

FRANK SPERLING, TIM CLEPHAS AND FERRY SCHOENMAKERS

Introduction

Ever since DARPA (the US Defense Advanced Research Projects Agency) decided to challenge engineers to come up with practical, implementable and robust forms of autonomous driving by organising the DARPA Grand Challenge in 2004 (*en.wikipedia.org/wiki/DARPA_Grand_Challenge_(2004)*), autonomy has more or less become a household word. The event marked the beginning of the transition for autonomous systems from academia to practical applications and industry.

An autonomous system is, by definition, a system that is self-governing, i.e. a system that can receive inputs through sensors, make decisions and act without an active (human) controlling input. The attractiveness of such systems from a user perspective is similar to that of automation: relieving human operators of tedious tasks and putting them in a supervisory role, saving costs and reducing risks while maintaining quality.

From an application perspective, it makes more sense to characterise autonomous systems by their practical merits, as mentioned above. From an academic perspective, the emphasis is usually more on the technology 'under the hood', which is usually dominated by the software, its architecture and decision-making and learning capabilities. From an engineering perspective, systems that can 'sense' and 'act' are often classical mechatronic systems. Adding signal processing, decision-making and learning capabilities yields 'smart mechatronics', taking the next evolutionary step from basic motion systems to autonomous, intelligent systems (Figure 1).

1 *Smart mechatronics.*

Software architecture for service vehicles

Service vehicles are (small) autonomous vehicles intended for repetitive, tedious or unsafe tasks, rather than for human transport or autonomous driving in a semi-controlled environment. Cleaning, lawn mowing, painting, logistics, bin-picking, etc., are typical examples of such tasks. Technology trends show that the time is ripe to automate such tasks and that the required technology is affordable.

As a starting point, a generic vehicle hardware platform has been developed at Nobleo (the Clara platform, see Figure 2) for a number of purposes: as a testing ground for the software platform, for reliability testing of the system and components, and for generic service applications. The appropriate operating system/software development environment was selected by comparing the available options and scoring them on aspects regarded as critical to quality (CTQ); see Table 1.

The columns in Table 1 show a number of candidates and the table summarises the aspects, with ROS (Robot Operating System) scoring highest and hence being selected as the platform of choice. ROS is part of an open-source

AUTHORS' NOTE

Frank Sperling (Nobleo co-founder), Tim Clephas (mechatronic designer) and Ferry Schoenmakers (robotics designer) all are associated with Nobleo Technology, based in Eindhoven, the Netherlands, an engineering company focused on high-end mechatronics. In recent years, Nobleo Technology's emphasis has been shifting towards smart mechatronics and autonomous intelligent systems.

frank.sperling@nobleo.nl
www.nobleo.nl

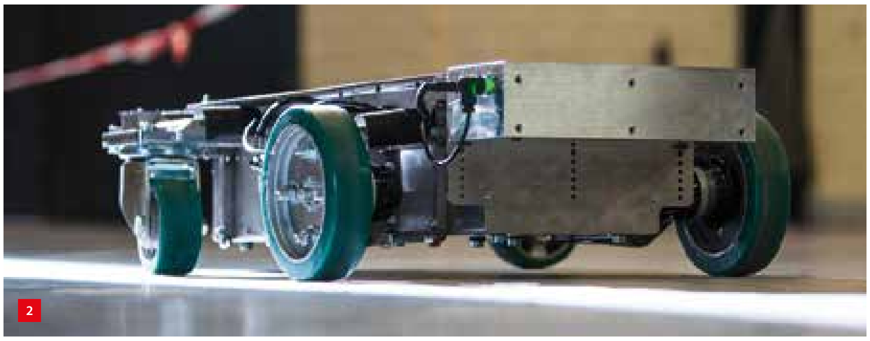

Clara, a generic platform for autonomous service vehicles.

Critical to Quality (CTQ)	In-house development	Matlab Simulink	ROS	ROS-Industrial	ROS 2.0	Orocos	Microsoft Robotics Developer Studio
Time-to-market, Development effort	–	+	+	0	0	0	–
Ease of use (simulation, diagnostics, debugging)	0	+	+	+	0	+	+
Standardised/traction solution	-	+	+	0	0	–	–
Software costs	+	–	+	+	0	0	0
Hardware costs	+	0	+	0	0	0	0
Quality assurance, robustness	0	+	0	+	0	0	+
IP ownership	+	–	0	0	0	0	0
Sum total	1	2	5	3	0	0	0

Table 1. Decision matrix for operating system selection.

initiative that originally started off around 2000 in the academic world, and which has since become mainstream for many academic and industrial initiatives. This is underlined by the fact that 18 of the 23 contestants in the latest DARPA Robotics Challenge in 2015 used ROS and many of them also used Gazebo (the accompanying simulation and visualisation package).

ROS 2.0 is a fundamentally revised concept of ROS with more focus on real-time behaviour and distributed systems (e.g. swarms). Conceptually, this is the preferred version. However, in 2016, when the above comparison was made, major parts of ROS 2.0 were still in the specification phase and little real functionality was available, hence the 0 score in the table. There is a concern among industrial users that an open-source concept lacks the rigorous quality assurance and accountability that a commercial product can offer. This is exactly why ROS Industrial was conceived. It consists of a subset of ROS functions that have been rigorously tested and qualified. Since the set of qualified functions in ROS Industrial is smaller than the complete set, it scores lower in the comparison, lacking functionality. Using the open-source version of ROS is in our view a manageable risk, as that there is a lot of interest in the ROS Industrial initiative (also actively supported by Nobleo), which will therefore sooner or later seamlessly replace the ROS functions.

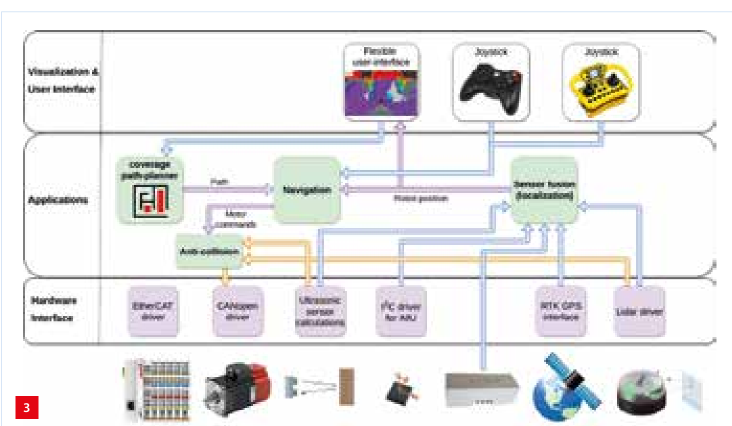

Generic architecture for an autonomous service vehicle.

SLAM

One of the first and foremost functionalities in a service vehicle is that it needs to 'know' where it is. Simultaneous localisation and mapping (SLAM) functionality needs to be embedded in the vehicle's architecture, of course independently of the relevant sensors used, such as GPS, odometers, range finders, lidar, radar, etc. A versatile architecture allows the physical sensor to be 'abstracted' into an information source with its own driver and allows effective 'fusion' of data from numerous sources into one common internal 'world representation'. Here too, ROS, as a communication platform provided with an extensive library of functions and building blocks, accelerates this development. Figure 3 shows the generic architecture for Nobleo's service vehicles. Particularly the breakdown into functional modules, the definition of the interfaces and the layering make it unique.

The upper layer shows hook-ups for different kinds of joysticks; either a cost-effective consumer type (game controller) or a more professional, heavy-duty type can be used. On the sensor/actuator side, numerous options are available. The Accerion sensor, an 'optical odometer' that can achieve high precision by automatic drift calibration, has been extensively tested in the vehicle. Other common modalities are Real Time Kinematic (RTK) GPS (high-precision GPS) and Inertial Motion Units (IMUs).

Application: WasteShark

One of the first spin-offs from the generic architecture is a system incorporated in a boat designed to gather waste from harbours: the WasteShark from the company RanMarine. Figure 4 shows the WasteShark in action. Although it looks very different from Clara, the software is practically identical, allowing the re-use of extensively tested software components from Clara (Figure 5).

Application: Industrial cleaning

Cleaning in general and industrial cleaning in particular is a dirty, dull and dangerous job. In Nobleo's vision, this is one of the areas where there is the most 'urgent' demand for robotisation and autonomous service vehicles. There

4

An aquatic drone: RanMarine's WasteShark equipped with Nobleo's ROS stack and SLAM architecture.

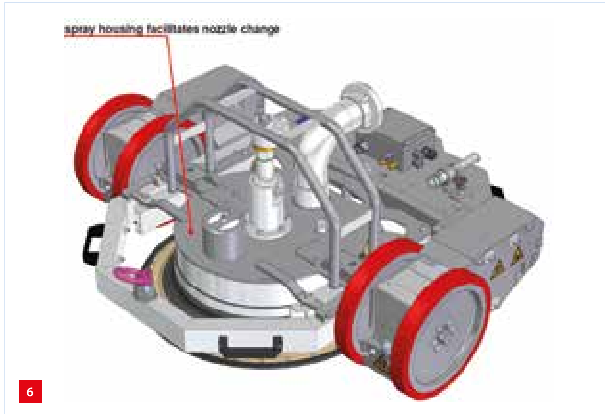

6

A cleaning crawler for steel-walled tanks; feeder and return hosing are omitted here, wheels are magnetic.

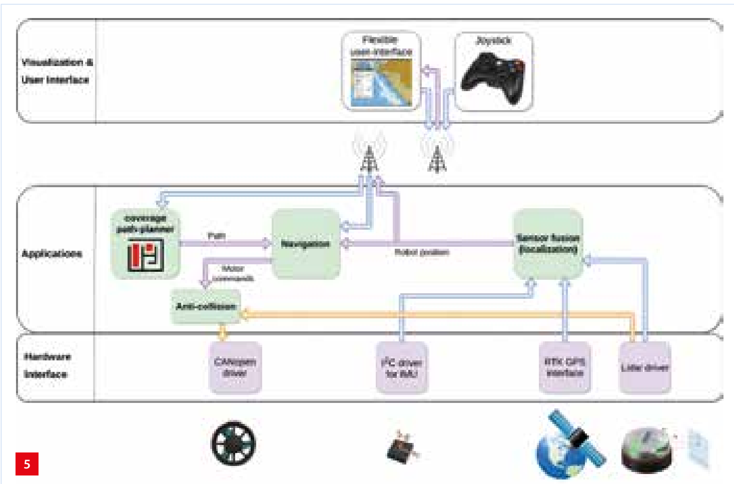

5

WasteShark architecture: the difference with the generic architecture is in the interfacing with sensors and actuators, while the core of the software is maintained. Note also the radio connection between the two separable top layers.

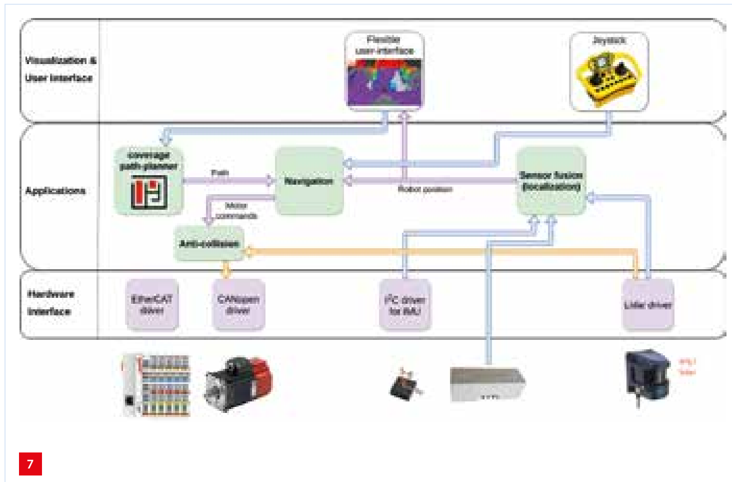

7

The architecture of the cleaning robot: 'same story', only different sensors and actuators.

are several initiatives, from large OEMs to start-ups, that have designed dedicated cleaning and inspection robots, but motion control is often confined to manual control and hand-held consoles requiring a human in the loop. In the case of (petrochemical) tank cleaning and inspection, the call for 'no man entry' is leading to the banning of human entry into tanks, making autonomous 'meandering' of cleaning and inspection devices practically mandatory.

Figure 6 shows a typical prototype crawler now being tested for internal and external cleaning of ferromagnetic tanks. It consists of a high-pressure cleaning system (feeder hoses omitted here) and magnetic wheels.

In Figure 7, the differences with respect to the generic architecture are once again at the device interface level, not in the core of the system. This allowed for quick and cost-effective prototyping of the crawler.

Autonomous systems roadmap

Service vehicles still have a long evolution ahead of them and at Nobleo we have only just crossed the border from remote control to 'situational awareness' and limited autonomy, i.e. SLAM, as depicted in the widely shared development roadmap of autonomous systems and service vehicles in Figure 8. There is still a long way to go in the development of robust industrial autonomous systems, i.e. smart mechatronics.

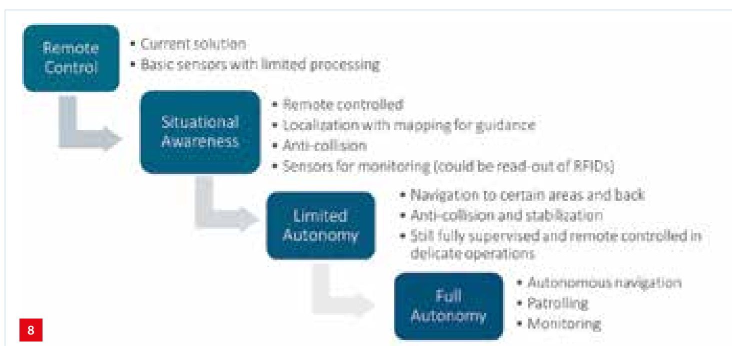

8

Development roadmap of autonomous functionality for service vehicles.

BUILDING SMART VISION BLOCKS

In manufacturing environments where collaborative robots are employed, conventional computer vision algorithms have trouble in the robust localisation and detection of products due to changing illumination conditions and shadows caused by a human sharing the workspace with the robotic system. In order to enhance the robustness of vision applications, machine learning with neural networks is explored. The performance of machine-learning algorithms versus conventional computer vision algorithms is studied by observing a generic user scenario for the manufacturing process: the assembly of a product by localisation, identification and manipulation of building blocks.

ROY DE KINKELDER AND BRIAN ARMSTRONG

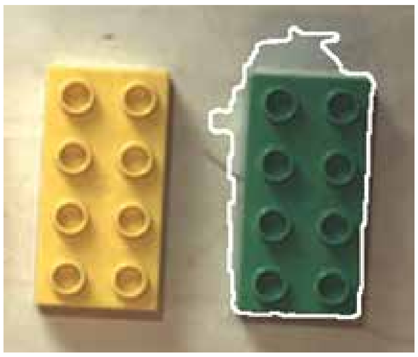

Introduction

High-tech production factories in north-western Europe are characterised by high-mix, low-volume production. Assembly is becoming increasingly challenging due to market dynamics. Therefore, production automation, flexibilisation and optimisation are essential in the trend towards manufacturing smaller batches, while retaining the capability to deliver a large variety of products. Collaborative robotics is an essential element in this trend, with vision systems as an important factor that facilitates flexibility [1].

Vision systems are used in pick & place applications, quality checks, product localisations, flow monitoring, etc. Although vision-controlled robotics has shown great benefit regarding efficiency and yield, it has major disadvantages when changes in the production process arise; vision systems in particular are sensitive to unpredictable environmental changes.

The robustness of vision systems using conventional computer vision algorithms often deteriorates due to (minor) changes in product shape and colour, lighting conditions such as the influence of sunlight, relocation of the production system on the shop floor, and in cases of processing biological products and food, etc. In other words, the robustness of the vision systems depends on the production environment. Moreover, in the case of collaborative robotics, where humans share the same workspace, the illumination changes continuously, as a result of unwanted shadowing (by passing operators) in the camera field of view.

In these cases, the system fails to recognise the object and a computer vision specialist must evaluate the new situation, recalibrate the system and re-program the software. Machine learning using neural networks has the potential to overcome a considerable number of these problems as they have been proven to be considerably less sensitive to varying environments and lighting conditions.

The Saxion research group Mechatronics and the companies Benchmark Electronics, which specialises in electronics manufacturing, and Bronkhorst High-Tech, which specialises in mass-flow meters, are exploring the use of collaborative robots in their production process in the TechForFuture (TFF) RoboTAO project. The focus of the research is on the real collaboration instead of sequential task deployment. A vision system is used to recognise human handling of the product and the intended operator's interference in the production process. For the purpose of enhancing its robustness, machine learning with neural networks is explored more thoroughly.

Collaborative robot scenario

The regular assembly of a product comprising several building blocks (product housing, connectors and PCB boards) is represented by the assembly process of a Duplo® (Lego® group, Denmark) house at different levels of human-

AUTHORS' NOTE

Roy de Kinkelder is a project leader/researcher at the applied research group Mechatronics of the Saxion University of Applied Sciences in Enschede, the Netherlands. His research focuses on machine vision applications for robotics and industrial automation and he is responsible for the management of the TFF RoboTAO project. Brian Armstrong is a student researcher in the same group. He studies Mechatronics and focuses on computer vision and machine learning in industrial applications.

r.dekinkelder@saxion.nl
www.saxion.nl/mechatronics


1
Set-up with camera in the top and illumination of the Duplo blocks.

machine interaction. First, the cobot recognises human interference in the production process, and continue where the human stops. Later, the cobot recognises a human sharing the same workspace with a hand detection algorithm. Finally, the cobot interacts with the human to assemble the Duplo house by sharing the blocks. Currently, the project partners are working on the first stage of the human-machine interaction.

To compare conventional computer vision algorithms with machine learning, a user scenario was defined, in which four types of Duplo blocks have to be identified and localised in order to pick them up. The 4 types were distinguished by colour: red, green, blue and yellow. The set-up is shown in Figure 1.

The blocks were randomly placed in a predefined workspace (35 x 55 cm), but always with the circular studs facing upward. To observe the blocks, a CMOS camera (DFK 23UX174, The Imaging Source Europe) was placed 100 cm above the workspace, also to observe the build plate. A 16 mm fixed focal length lens was used to focus the light on the sensor (1,920 x 1,200 pixels). With an entire field of view of 76.5 x 47.0 cm, a spatial resolution of 0.39 mm/pixel could be achieved. To cancel out glare and unwanted specular reflections, a linear polarizer was placed in front of the lens.

A UR5 collaborative robotic arm (Universal Robots, Denmark) equipped with an RG2 gripper from OnRobot (Denmark) as an end-effector was used to manipulate the blocks and position the blocks on the build plate. There are a few methods to control the UR5. The most common methods are URscript, Matlab using the URControl, and Robot Operating System (ROS) using the URControl.

In the first method, the UR5 is programmed through the connected teach pendant along with the graphical PolyScope programming interface and URScript programming language. This method can be entered through the teach pendant and saved as a program to be executed on the robot. Hence, it is rather easy to program for pick & place tasks. However, using it for long and multi-tasking programs is complicated.

The second method is Matlab. Within the Matlab driver, the velocities in the joint space of the robot are controlled. This gives a good performance and is safe, but it is not open source. Also, there are limitations in the computational load for further developments and communication with other robots.

The last method is using the ROS programming environment, which is a set of software libraries and tools to build robot applications. This provides the services expected from an operating system, including hardware abstraction, and low-level device control. There are different packages in ROS that provide the capability of doing requests such as computing trajectory, connecting joystick and so forth. Specific packages can be added for many robotic applications. Furthermore, ROS has great simulation tools to show robot movements in offline or real operational mode. And above all, it is open source and it is relatively easy to communicate between Python and C++ programs.

Because of the aforementioned advantages and in consideration of further developments in the detection of and collaboration with humans, ROS was selected. ROS runs on Ubuntu and provides the capability to run different drivers or packages (including URdriver, camera, gripper, and additional sensors) within the ROS environment. This means that only one computer (here a NUC core i5) is enough to run the whole set-up.

The results and performance were analysed using two different computer vision approaches: one based on predefined colour spaces, also used in previous studies [2, 3]; and the other with machine learning, using multi-layer perceptrons [4].

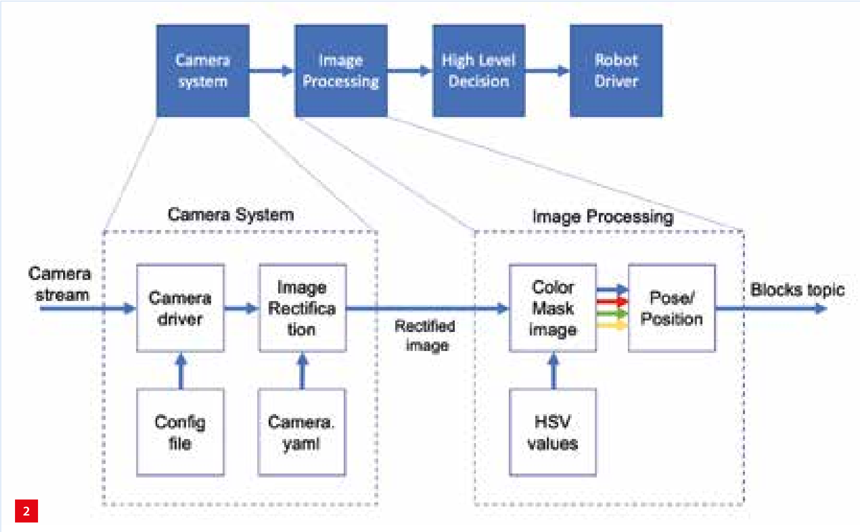

Software flowchart of the computer vision script with predefined HSV (Hue, Saturation and Value) levels.

Conventional computer vision algorithms

In our ordinary computer vision script, we used the open source vision library OpenCV 3.0 in combination with ROS to acquire images from the camera. Images were asynchronously acquired with a maximum frame rate of 30 frames per second (fps). At start-up, white balancing was carried out and the gain and exposure time were set, see Figure 2. The image was then rectified using camera parameters that were determined after a camera calibration procedure [5]. This procedure must be carried out to correct for any distortions caused by the lens and remap the spatial sampling of the image.

Before the system can determine the orientation and position of the blocks, the rectified image must be masked using HSV spaces (Hue, Saturation and Value). The HSV spaces were determined by first adjusting the hue level until the correct colour (e.g. green) was shown in the masked image. Then, the saturation and value levels were adjusted to observe only that specific coloured block with the correct saturation and intensity value. The HSV spaces were saved for all four coloured blocks. Using the HSV levels, a 'find contour' algorithm was employed to find connected regions in the image, see Figure 3a. A minimum bounding rectangle was placed around the contour to determine the orientation of the block by using the corner values. The centre of the rectangle was used as position value for the block.

Field of view for the 'find contour' algorithm.
(a) Fixed top illumination showing correct detection of the green block.
(b) Environmental light causing errors in identifying the green block.

To demonstrate the problem with conventional computer vision algorithms, the set-up was placed close to a window to evaluate the performance under different and uncontrolled lighting conditions. These results are shown in Figure 3b. The altered illumination conditions caused the vision script to detect unwanted regions in the images; beforehand, the hard-coded HSV levels should have been recalibrated. In this case, the illustrated change in illumination was severe; even minor changes in intensity and spectrum of the light source can cause inaccurate determination of the position and orientation of the blocks.

Blocks picked by the robot arm were first placed in a mechanical fixture to correct any alignment errors due to incorrect visual detection and errors caused by incorrect transfer frames (these frames convert camera pixels to real-world coordinates). After correction the blocks were placed on the build plate.

Machine-learning algorithm

In computer vision, the task of image segmentation is formulated as a classification problem where each distinct region of interest (ROI) is considered a class with distinct features. The six main classes in the image workspace are shadows, the background plate, and the four coloured blocks (red, green, blue and yellow). For each pixel in the image, the three primary colour channels Red, Green and Blue are considered the predicting features (please note that the latter group of colours are the camera colour channels and not the final classes).

Within each class definition, these features will vary due to the workspace lighting conditions not being constant. The relationship between features and classes is modelled using a Multilayer Perceptron (MLP), a popular type of Feed-forward Neural Network. The MLP has inputs and outputs that match the number of features and classes, 3 and 6 respectively.

To train the MLP, supervised learning was used; this is a machine-learning method that can be used when the inputs and outputs of the network are known. The MLP was trained to map the class label to the input features. The training dataset consisted of paired features and class labels made using images that spanned a wide set of illumination conditions. The trained MLP was tested using novel data, achieving a prediction accuracy of 97.8%. Rectified images were sent to the block detection program,

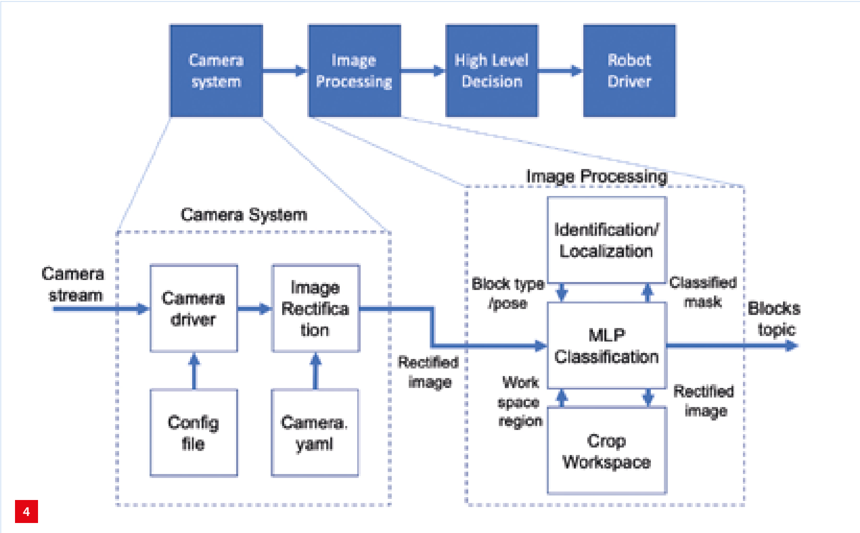

Software flowchart of the machine learning script with MLP.

Figure 4. The bolts that fixate the workspace were used to determine the workspace location and the image reference frame was set. All new images were cropped to show only the workspace. Image classification was carried out by reshaping the 2D cropped input image into a 1D array containing all the pixels. The MLP carried out a batch prediction on this array and the resulting output was reshaped back into its original dimensions. The resulting output was a segmented image containing regions of interest and their respective class, as shown in Figure 5.

A block was localised in the image frame using the distinct round studs that line the top of each Duplo block. The position of a block's origin was considered to be the mean of the stud feature centre-point coordinates.

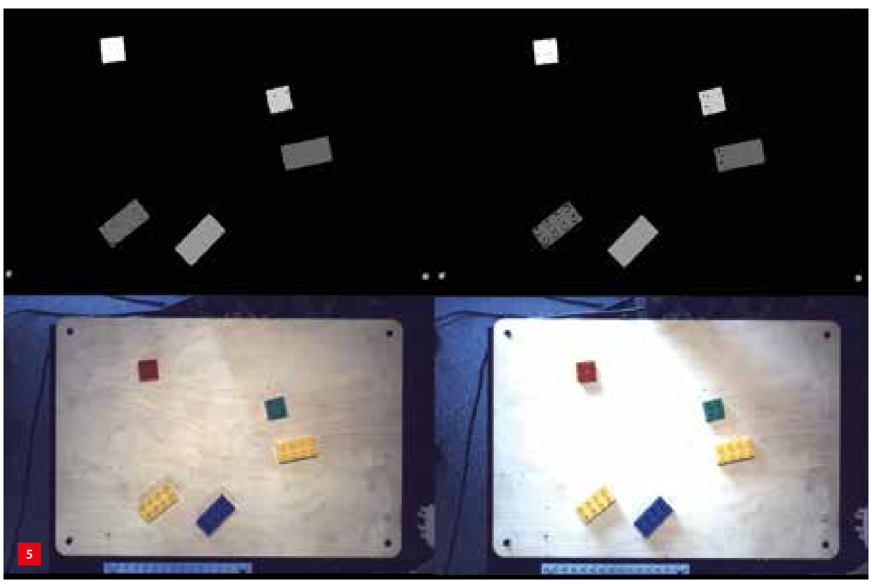

On the bottom, the raw images with changing illumination situations are shown. At the top, the labelled foreground masks are shown that were obtained with the machine-learning algorithm from the corresponding input images below.

To locate these features, the block was extracted from the original RGB image. Using the segmented image, the tightest convex polygon containing the block ROI was calculated. This calculated polygon was used to extract only the block from the RGB image. Centre-points of the studs were determined with a Hough circle transformation and the average of these feature coordinates was taken as the block's local origin. A minimum area rectangle was fit on the segmented block to determine its orientation. This process was performed iteratively for all blocks in the image and reached a high localisation precision 2σ: 0.321 mm, 0.233°.

Using the polygon area and the number of studs of each block, the block type could be identified. In the case of Figure 6, the block was identified as a blue 2 x 4 block. After identification, blocks were picked and directly placed on the build plate without the need for a mechanical fixture.

Comparison of results

In comparison to conventional computer vision algorithms, machine learning, and in particular the feedforward neural network, can be used to significantly improve robustness in the detection of colours and the identification of the Duplo blocks. We measured the physical position of the gripper with respect to the centre of the Duplo block with a caliper tool (mean error: 0.1 mm) to quantify the performance of the machine-learning method and the conventional computer vision method.

The four different coloured blocks were used and every block was measured twice under changing lighting conditions. The variation in lighting conditions was kept similar for both methods. The mean error in position estimation for the conventional vision algorithm was 4.6 mm (2σ: 7.3 mm), while the machine learning showed a mean error of 0.8 mm (2σ: 2.2 mm).

Identification and localisation is more robust due to the statistical nature of the classification. It determines the probability of a pixel belonging to a specific class (colour); therefore, changing the environment will have less influence when using machine-learning algorithms. This could be very beneficial in situations where product materials could vary between batches or human interaction is part of the production process, such as in collaborative robotics.

Computer vision script that uses the hard-coded threshold to differentiate between various colours can be unreliable with even the smallest changes in the environment. However, in the case of short processing times, computer vision algorithms are to be preferred because they are faster. In our situation, a maximum frame rate of 54 fps could be achieved. In comparison, the machine-learning algorithm

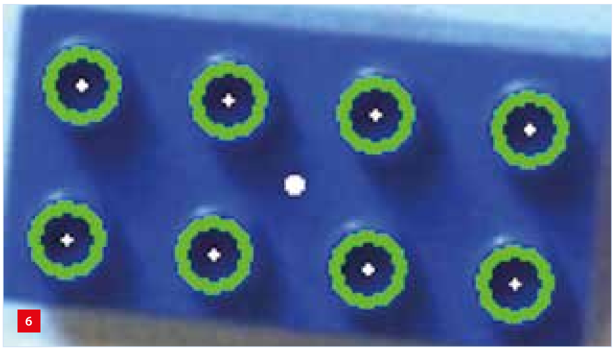

Duplo block with Hough circle detection to localise the centre (white dot) and the orientation of the block.

could be used with a maximum frame rate of 2 fps. By lowering the image resolution for classification, the frame rate can be increased to approximately 10 fps, but this is still significantly slower than the conventional methods.

Conclusion

Systems in manufacturing processes, or inspection lines using vision for pick & place applications and quality checks are more robust to changes in lighting conditions and product properties (such as material colours) when using machine-learning software algorithms.
This case study has shown that relatively simple feedforward neural networks like MLP can be used to identify products of interest. Furthermore, the need for mechanical alignment using a fixture can be prevented by using the machine learning algorithm. The mean error of 0.8 mm is sufficiently low to position the blocks on the build plate. With the 2σ deviation of 2.2 mm, however, the block is sometimes not aligned properly. This can be improved by a better calibration of the transfer frames of the robot to the build plate and the camera.
In the case when processing speed is a critical parameter, it is recommended to invest in proper hardware and stable (illumination) environments. The research is ongoing, but the preliminary results look very promising.

Acknowledgement

This project was supported by TechForFuture, Bronkhorst High-Tech and Benchmark Electronics. Special thanks go to Hanieh Esmaeeli, Mick Boe and Dirk Bekke.

REFERENCES

[1] J. Heilala, and P. Voho, "Modular reconfigurable flexible final assembly systems", *Assembly Automation*, Vol. 21 (1), pp. 20-30, 2001.
[2] A.Y. Mersha, R. de Kinkelder, and D. Bekke, "Affordable Modular Mobile Manipulator for Domestic Applications", *19th International Conference on Research and Education in Mechatronics (REM)*, Delft, pp. 141-146, 2018.
[3] R. de Kinkelder, and R. Zwikker, "Getting a grip on Medical and Modular Robotics", *Mikroniek* 57 (4), pp. 26-31, 2017.
[4] *scikit-learn.org/stable/modules/neural_networks_supervised.html*
[5] *wiki.ros.org/camera calibration*

MORE THAN ROS

All the attention that ROS (Robot Operating System) is getting nowadays is completely justified. After all, this meta-operating system is responsible for sparking a new way of thinking on how to compose robot applications with a lot of capability for our homes and public spaces using re-usable software. Before that, industrial robot manufacturers had kept us believing that robots were expensive, dangerous and not easy to use. All this being said, there are more options available than this open-source software framework called ROS. The world needs modular robot software and robot hardware, including safety, to build reliable and secure applications.

HENK KIELA

Robot times

With robotisation now really raking off, we are definitely living in interesting times. After the recent technological innovations of computers, mobile phones, internet and smartphones, robots will now enter our lives in various forms. "We're not seeing the robots yet, but we certainly could use a few" is one of the reactions many people have – a striking response given that only a few years ago robots were seen as a threat to our global workforce. But people's attitude is definitely changing to working with the three Ds – dirty, dull and difficult.
Although we may not see a lot of robots in daily life outside factory walls, closer inspection reveals that robot technology and artificial intelligence (AI) can be found in common products. Robot vacuum cleaners for example and robotised lawn mowers are still rare, but these early examples of a new generation of robots are changing the way vacuum cleaning and lawn mowing are fundamentally done.

While the principles we now use in technology and AI were determined 20 years ago or more, we still aren't seeing many robots around. The explanation is simple. Robots first need to become a lot cheaper and a whole lot easier to use. And more importantly, these robots need to be able to cope with the unstructured environment we live in and do something useful in a safe and pleasant way. As humans, we've only just started to understand the social patterns of how we interact with each other. AI will help robots in the near future to 'understand' humans in the environment and interact with them in a way that humans understand.

Interestingly, industry is driving new applications of humans and collaborative robots, i.e. cobots, working together. Development in Smart Industry (in the Netherlands) and Industrie 4.0 (in Germany) underline the importance of robotisation, cobots and connected distributed production cells to improve flexibility and reduce the offshoring of work to low-wage countries.

This vision of Smart Industry has also been adopted by the ROS community [1] and the ROS Industrial [2] community. The ROS open-source initiative took off at Stanford University in California, USA, around 2008 in an effort to speed up robot developments by connecting all relevant open-source software in one framework. The ROS community shares the idea of flexible production with the help of robots (figure 1) and the way equipment takes the lead in inviting peer robots to collaborate rather than that the whole system being driven top-down.

ROS evaluation

The worldwide ROS community brings together all relevant knowledge and a variety of robot-relevant open-source software on a 'standardised' platform, or middleware, running on Linux. As the quality and reliability of this software was not as good as everyone had hoped, the ROS Industrial community started working on a more reliable ROS version.

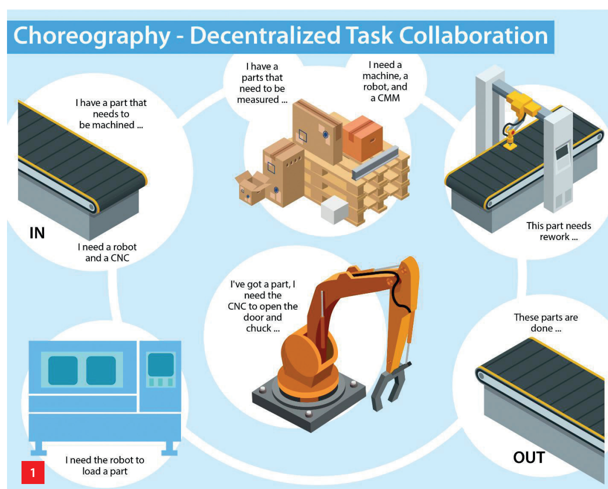

The ROS community shares the idea of flexible production with the help of robots that are able to 'understand' what the other team members need and/or are doing. [3]

AUTHOR'S NOTE

Henk Kiela is a part-time professor in Mechatronics and Robotics at Fontys University of Applied Sciences in Eindhoven, the Netherlands, and is the owner of Probotics AGV Systems in Venlo, the Netherlands. His company has been building self-driving vehicles for a range of logistic applications since 2014. He is an expert in robot safety and member of the ISO committee TC299.

h.kiela@fontys.nl
www.fontys.nl/mechatronics
www.probotics-agv.nl

The principles of ROS and ROS Industrial proved that modular software indeed helps to speed up the development of robot applications. But the system was not sufficiently reliable, did not offer any safety and still required a lot of expertise to 'compose' a robot system. The ROS 2.0 project was started recently, incorporating the new requirement for mission-critical functionality and indeed safety.

Whatever you think of ROS, it has changed the thinking about robots regarding such things as the affordability of complex technology, the flexibility in recomposing robot functions and how to provide a way for worldwide collaboration in robot developments and research. The amount of open-source software compatible with ROS is dazzling and still growing fast.

However, our company Probotics decided not to continue with ROS for a number of reasons. As a manufacturer of self-driving robot systems, we have to deal with many more aspects than ROS covers. To sell safe and reliable self-driving vehicles (SDVs) for an industrial logistic application that is easy to integrate on the shop floor and easy for non-technical people with a non-academic background at our customers to maintain, you need to provide reliable and very-easy-to-use robot systems.

SDVs are the next generation of automatic guided vehicles (AGVs), except these SDVs navigate on environmental features and not on fixed lines and floor features. As such, SDVs are much more flexible than classic AGVs. They provide mobile logistic solutions that are safe, low-cost, easy to reconfigure and easy to integrate, even without the need for ROS.

Because the focus is on ROS, many people may have forgotten that there are a number of alternatives around to build mobile robot systems with affordable industrial components. Companies such as the Swiss company BlueBotics and the Finnish company Navitec developed their SDV navigation systems a few years ago and have sold over 1,000 systems worldwide to mobile robot manufacturers and integrators. These systems provide navigation, localisation and fleet/traffic management. Their systems are open, very reliable and simple to integrate and to maintain.

They have a few aspects in common with ROS, which the ROS community is not addressing very well. As already mentioned, safety in ROS and ROS Industrial is almost invisible in the software and barely discussed in the community – it sometimes doesn't even seem to be an ambition. What's more, the reliability of open-source ROS-driven robots is a serious consideration.

The main function of a self-driving robot is navigation. This can be defined as the combination of the three fundamental competences:

1. Self-localisation
2. Path planning
3. Map-building and map interpretation

To build a safe and reliable SDV, more is needed. Table 1 presents a comparison of some aspects from industrial navigation systems and ROS, using the operating systems of BlueBotics and Navitec as examples, as these are familiar to us. But there are more suppliers of similar systems on the market, with similar capabilities.

Table 1
Comparison of aspects relevant to industrial navigation.

	BlueBotics	Navitec Systems	ROS
Re-configurability of SDV for new tasks	Very simple through CAD-based tools.	Very simple through CAD-based tools.	Reprogramming, reconfiguring launch files, validation of new configuration needed.
Safety	Safety-related features and functions, integrity check and signals.	Safety-related features and functions, integrity check and signals.	No safety, you have to build it yourself; no aspects regarding performance or safety present in any module today.
Security	Vulnerable, but documented.	Vulnerable, but documented.	Very open system, easy to intrude and disturb.
Navigation	Very reliable, easy to reconfigure; little effort required to make changes in routes and maps.	Very reliable, easy to reconfigure; little effort required to make changes in routes and maps.	Very dependent on many settings and configuration files; poor and complex documentation.
Hardware reconfiguration	Good plug & play support for major (safe) hardware suppliers.	Good plug & play support for major (safe) hardware suppliers.	A lot is available, but quality and ease of reconfiguration is not as good.
Integration with production environment	Good, many options on board and available remotely via network.	Good, many options on board and available remotely via network.	Has to be developed, few or nothing available.
Fleet/traffic management	Standard, easy integration with warehouse management system (WMS).	Standard, easy integration with WMS.	Possible, most of it has to be developed.
Cost	Fair, all-in-one package, good support.	Fair, all-in-one package, good support.	Open source is 'free'; the effort to make it a reliable and safe system is unpredictable.
Fit-for-future challenges	New challenges need to be developed.	New challenges need to be developed.	Latest functionality probably available; validation needs attention.
Community supported	None.	None.	Growing worldwide community.
Fit for many more robot applications	Impossible.	Impossible.	Very good ability.

Key:

Good Fair Weak

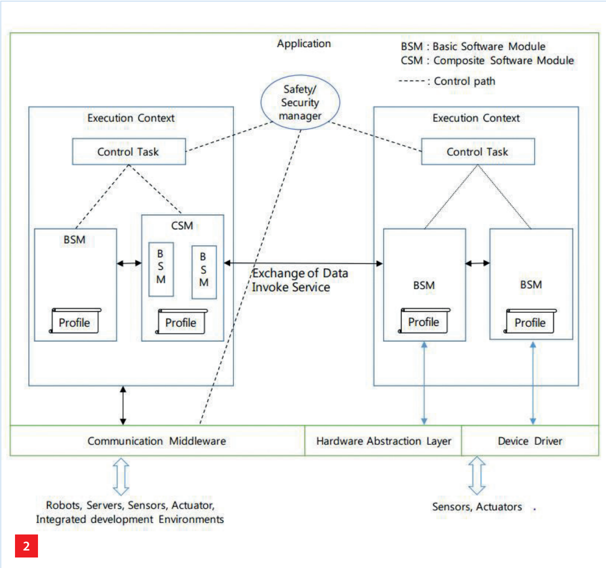

Software modularity as described in the ISO 22166 standard.

Based on our observations, it can be concluded that ROS definitely has its qualities, certainly when looking to the future. But there are cheaper, safer and much more reliable systems available on the market for industrial SDV applications .

Standardisation

While classic industrial robot manufacturers kept telling us years ago that robot control is complicated and that software can't be made modular, ROS has demonstrated the opposite. Modularity in software is possible and offers great advantages in dealing with complex robot systems. This is also reflected in the ISO initiative. Modularity for robot systems has gained a global momentum thanks to the efforts of the Technical ISO Committee TC299, which, in 2014, resulted in a new project to develop a standard for modular robot functionality: ISO/CD 22166-1.2, "Robotics - Modularity for service robots - Part 1: General requirements" [4].

This new ISO standard incorporates hardware and software modules and includes guidelines for the design of safe and secure robot modules. While all of the software aspects presented in this standard are applicable to ROS modules, the ROS community has surprisingly shown little interest in participating in this effort.

The principles presented in this standard help to incorporate safety and security in modules and in module architectures in an industrial manner. If these principles are followed at an early stage of module design, either in hardware or software, the cost of achieving a certain safety performance is low and the result for system integrators will be great. Those integrators who adopt the guidelines at an early stage of development have to spend a lot less time on integrating certified modules into a system, compared to ROS modules, to build a reliable safe and secure robot application. Figure 2 illustrates this standard in terms of software modularity.

Future perspective

ROS and manufacturers of relevant industrial navigation modules seem to be working in different worlds, while there could be major benefits if these two worlds were connected. There are at least two ways to promote more affordable, robust and flexible robot solutions in our world for both ROS and industrial suppliers:

1. Existing manufacturers of robot systems and robot modules cannot ignore the ROS community. They have to provide interfaces to ROS and become part of the ROS (Industrial) community rather than seeing ROS as a competitor. They have been working according to industry (safety) standards for a long time, their products are reliable and simple to configure, but they must now consider providing standardised interfaces to ROS to open up a new market. The ROS community needs this kind of reliable and safe functionality, and will embrace products with an ROS interface.

2. The growing ROS community is driven by science, not by industry, which is good. But it does not pay enough attention to reliability, safety, security, maintenance and support. These aspects have to be addressed in the architecture and implementation of ROS modules. ROS is not interested enough in hardware and integration aspects. In 2018, ROS 2.0 defined some ambitious goals regarding these aspects. But ROS is at least 10 years behind on these aspects compared to the industry that incorporated these aspects in their robot products a long time ago. Opinion leaders in the ROS community see less

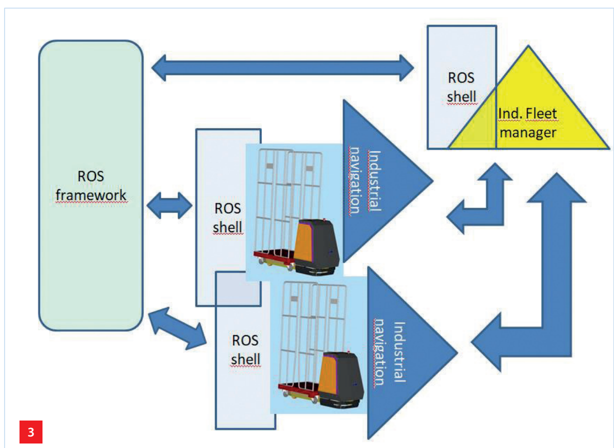

Example of a hybrid system of industrial-grade professional components controlling a fleet of vehicles connected to an industrial fleet/traffic manager and ROS functionality.

value in complying with standards for safety and security. Some of the companies that brought robot products to the market based on ROS had to spend quite some time on making their product fit for industry.

Manufacturers
Manufacturers have to reconsider their position with respect to the expanding ROS community. They can no longer ignore the existence and importance of this community. Building a new relationship with the ROS community will open up new markets for the industrial-grade products they offer.

Universal Robotics and others manufacturers have shown the commercial benefit of embracing ROS and their products are greatly appreciated by the ROS community because of their reliability and price/performance ratio. Figure 3 shows an example of a hybrid industrial-grade/ROS system.

Building an ROS interface is not difficult for a manufacturer; it is merely a choice based on their vision of the future. My message to manufacturers of industrial robot modules is: “Don't be afraid to join the ROS community.” There is no risk of losing market to a competitor. First of all, ROS is not a company, but a community with its own, non-commercial, behaviour and with a huge potential to embrace ROS-compatible products enthusiastically.

There is a need for good quality functionality. And once an industrial module is open to ROS, all other developments in ROS can become part of joint products. And last but not least, a supporting ISO standard is on its way to help structuring the quality and integration of robot modules.

ROS community
Rather than broadening the scope of ROS towards new applications like 3D printing, community efforts should also at least focus on quality of service for existing ROS functionality, and simplification of use of ROS and reconfiguration. They should also incorporate safety and reliability in their concept. This has now been defined in ROS 2.0. But the target timelines for these ambitions and specifications are still unclear.

In our view, there is a strong parallel between ROS modules and the development of apps for smartphones. Initially, app stores were open arenas where everyone could post apps. This resulted in unreliable applications that even jeopardised smartphone integrity and security. Manufacturers responded by imposing guidelines and quality criteria on new apps.

A store of ROS-certified modules could provide a similar function to the user community. The certification should provide minimum qualifications for the performance, safety, security and maintainability of a module. Such a scheme could be adopted for ROS 2.0 in the future, but this should also be done right away for ROS Industrial. This would help enormously to attract industrial suppliers of robot modules and components to become part of the community and help system integrators to introduce complex robot solutions with less effort into our society. In the end, everyone would benefit.

REFERENCES
[1] *www.ros.org*
[2] *www.rosindustrial.org*
[3] *www.mtconnect.org*
[4] *www.iso.org/standard/72715.html*

UPCOMING **EVENTS**

22 January 2019, Eindhoven (NL)

3D Printing Electronics Conference

The focus of the sixth edition of this conference is on aspects such as combining functional elements like electronics (sensors or switches) into a (3D-) printed product hybrid; processes that integrate electronics onto or within 3D-printed parts; and 3D-printed optics/photonics. Challenges for manufacturers/engineers/researchers and future prospects/applications are also covered.

WWW.3DPRINTINGELECTRONICSCONFERENCE.COM

22-23 January 2019, Sheffield (UK)

Integrated Metrology for Precision Manufacturing Conference

The first of two conferences being held as part of a roadmapping project to define the future of integrated metrology in advanced manufacturing in the UK.

WWW.NOTTINGHAM.AC.UK/CONFERENCE/FAC-ENG/METMAP-2019

13-14 March 2019, Veldhoven (NL)

RapidPro 2019

The annual event showcasing solutions for prototyping, product development, customisation and rapid, low-volume & on-demand production.

WWW.RAPIDPRO.NL

13-14 March 2019, Sheffield (UK)

Lamdamap 2019

Thirteenth edition of this event, focused on laser metrology, coordinate measuring machine, and machine tool performance.

WWW.LAMDAMAP.COM

19-22 March 2019, Ede/Veenendaal (NL)

Demoweek 2019

Eight companies demonstrate their automation offerings for the metalworking industry: software, robotisation, control, measurement, 3D printing and machining.

WWW.DEMOWEEK.NL

25 March 2019, Düsseldorf (DE)

Gas Bearing Workshop 2019

Third edition of the initiative of VDE/VDI GMM, DSPE and the Dutch Consulate-General in Düsseldorf (Germany), focused on gas bearings as important components or integral technology of most advanced precision instruments and machines.

WWW.GAS-BEARING-WORKSHOP.COM

26 March 2019, Veldhoven (NL)

CLEAN 2019

This theme day, organised by Mikrocentrum, provides an an expert's view on cleanliness. Speakers from academia and industry will present new developments, discuss process and cost optimisation, review quality control and share best-practice applications.

MIKROCENTRUM.NL/EVENEMENTEN/THEMABIJEENKOMSTEN/CLEAN-2019

11 April 2019, Eindhoven (NL)

High-Tech Systems 2019

One-day conference and exhibition with the focus on high-end system engineering and disruptive mechatronics in, for instance, smart manufacturing, thermal design, smart logistics, scientific instruments, design principles and medical systems.

WWW.HIGHTECHSYSTEMS.EU

15-16 May, Leuven (BE)

Materials+Eurofinish 2019

At this joint event material science meets surface technology. Combined, these ingredients help to achieve sustainable designs and innovative ideas, from (new) materials, material analysis and surface technology to binding techniques. The fair provides a complete overview of the entire value chain: from raw materials to a finished product.

WWW.MATERIALS.NL

3-7 June 2019, Bilbao (ES)

Euspen's 19th International Conference & Exhibition

This event features latest advances in traditional precision engineering fields such as metrology, ultra-precision machining, additive and replication processes, precision mechatronic systems & control and precision cutting processes. Furthermore, topics will be addressed covering robotics and automation, Industrie 4.0 for precision manufacturing, precision design in large-scale applications and applications of precision engineering in biomedical sciences.

Venue of euspen's 19th International Conference & Exhibition, the Euskalduna Conference Centre in Bilbao, Spain.

WWW.EUSPEN.EU

6 June 2019, Enschede (NL)

TValley Annual Conference 2019

The conference will provide a state-of-the-art overview of robotics and mechatronics R&D activities of the TValley network and its industrial partners. The Tvalley agenda includes mechatronics education, knowledge exchange, specific projects on robotics and smart industry, and profiling of the high-tech industry in the east of the Netherlands.

WWW.TVALLEY.NL

12-13 June 2019, Veldhoven (NL)

Vision, Robotics & Motion 2019

This trade fair & congress presents the future of human-robot collaboration within the manufacturing industry.

WWW.VISION-ROBOTICS.NL

DO NOT NEGLECT

Patent publications are a valuable source of knowledge – not to be neglected – for product development. Patent publications describe a technical problem and a solution, which is new worldwide. The solution must be described in such detail that a skilled person can realise it. So all kinds of technical solutions can be found in patent publications. This is illustrated here with patents about omni wheels, which are interesting because of their application in robotics, and about micro surgery. To conclude, it is explained how to find relevant patent information about specific subjects.

HANS HELSLOOT AND DIK VAN HARTE

Omni wheels

Omni wheels or omnidirectional wheels are wheels that allow a vehicle to move in all horizontal directions and rotate around its vertical axis, without steering any of the wheels. This is possible due to the fact that along their circumference they have small rollers, which are oriented at an angle with the turning direction of the wheel.

Omni wheels were not invented recently (Figure 1). The oldest patents we found are by Burnett, issued in 1911, and by Grabowiecki, issued in 1919. Their omni wheels were fitted with rollers perpendicular to the wheel axis. A disadvantage of such wheels is their bumping behaviour. In 1944, a patent was granted to Hladil for some improvements, reducing the bumping, but with all rollers still perpendicular to the wheel axis.
In 1949, Christian Fuchs found that by placing the rollers at an angle other than 90°, the wheel was still working fine and the bumping was reduced significantly (Figure 2a). In 1972, Bengt Erland Ilon, working for Mecanum from Sweden, improved that idea by rotatably mounting convexly shaped rollers on the centre part (Figure 2b). This was a great improvement, since the angle of the rollers could be adjusted to the required use. Ever since then, these wheels are also called Ilon wheels or Mecanum wheels.

Although the inventions of Fuchs and Ilon reduced the bumping significantly, vibrations remained a problem. In many other patent applications inventors suggested improvements. In 2001, Byun & Song proposed for example to reduce vibrations by using inner and outer rollers (Figure 3a). In 2008, Fuji proposed to make omni wheels with flexible rollers, covering the entire circumference of the wheel (Figure 3b). In 2009, Potter proposed to reduce vibrations further by applying at least two rows of angled rollers around a hub, whereby the rollers in each row are axially and rotationally offset from each other (Figure 3c).

Strategic information and statistics

Patent information is also important strategic business information. Applications are published 18 months after their first filing. Patent applications are often published before the product is introduced on the market. If you use the RSS feed of the patent database (see below), you receive a message when a new application by a competitor or in the relevant field of operation is published. There could be a blocking patent or a new competitor, or perhaps there is an interesting company or university to collaborate with.

AUTHORS' NOTE

Hans Helsloot and Dik van Harte are both intellectual property consultants at the Netherlands Patent Office (Octrooicentrum Nederland, part of RVO.nl)

hans.helsloot@rvo.nl
www.rvo.nl/octrooien

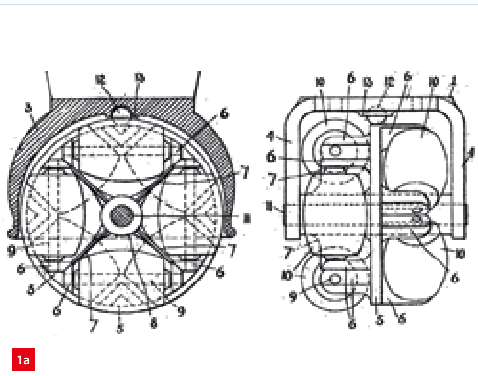

1a

GB7953
Burnett, 1911

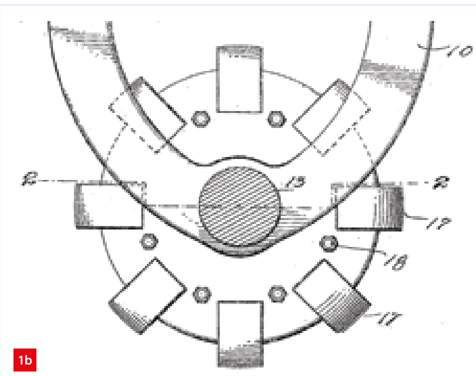

1b

US1305535
Grabowiecki, 1919

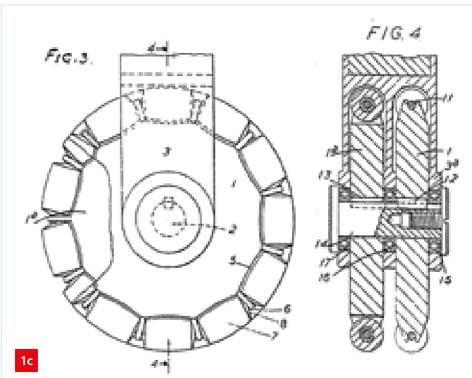

1c

GB565138
Hladil, 1944

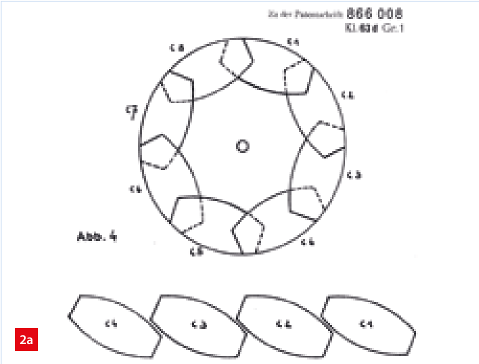

2a

DE866008
Christian Fuchs, 1949

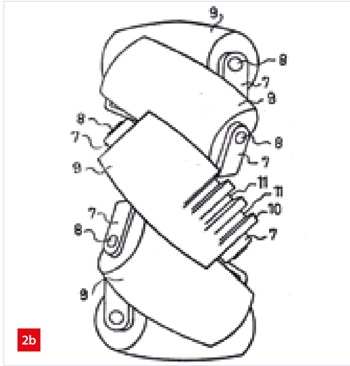

2b

US3876255
Bengt Erland Ilon, 1972

Omni wheels

Patent statistics also provide strategic information. The number of patent applications relating to omni wheels exploded in 2016, as can be seen in Graph 1, showing the number of unique applications filed between 1987 and 2016. The highest number of patent applications in this period originates from Honda Motor Company [JP] with 41 applications. Japan and China are the most frequent countries of origin. The only applicant from the Netherlands is Philips with two applications.
A closer look at the results shows that more than half of the patent applications about omni wheels originates from China. It is clear that there is a market for omni wheels and that the Chinese industry is aware of that. Sometimes we can only guess why there is such a steep rise in the number of patent applications. The patent databases give us the numbers of applications and the countries, but not the reasons behind them.

Micro surgery

A similar analysis can be made about micro surgery.
In the period 1991-2016 in total 478 unique applications have been filed in these area. The top five applicants are:

- Intuitive Surgical (US) 46
- Olympus (Medical) [JP] 18
- Siemens (Healthcare) [DE] 16
- Philips (Electronics) [NL] 15
- Ethicon (Endo Surgery) [US] 14

The origin of these applications can be found in Graph 2. In this graph we see that so far Asia has not been very active in the development of micro surgery in comparison with Europe and the USA. The absolute number of applications in this field is not very high, but the growth of the total numbers is higher than the growth of the overall number of patent applications from Europe or USA.

Inspiration from omni wheels and their applications

Patent information is a useful source of inspiration and information. Most technologies described can be used without any legal problems, since less than 5% of the patent applications published are valid in important countries like USA or Germany. Less than 1% are valid in the Netherlands. Moreover, even if the technology is not freely applicable in a country, it is still permissible to study the technology and learn from it or derive inspiration from it.

This is also demonstrated in the case of omni wheels, where there are many areas of application. Most omni wheel patent applications relate to robots and to the displacement of heavy or vulnerable equipment. This is because holonomic devices can be realised with these wheels: they can move and rotate in any direction.

For example, Hitachi explained in 1989 how to control movements and vibrations of an all-directional vehicle with three or more omni wheels in narrow places or places having an uneven floor surface (Figure 4). Their patent expired in 2001 in the USA and was never applied for in Europe.

Another example by Philips teaches in detail (25 pages) how to make extremely heavy X-ray equipment easily moveable by using one or more omni wheels. The patent application explains how the omni wheels can be used to displace and rotate equipment in any desired direction, for example with a joystick (Figure 5). In an operation room this is an important advantage. The use of omni wheels facilitates a compact and low weight design. The patent is granted and valid, so the technology as described in the claims cannot

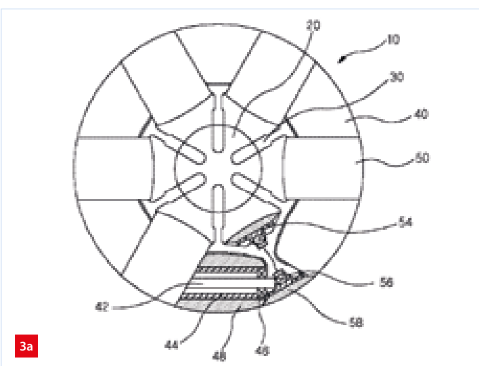

3a

KR20020063737
Byun & Song, 2001

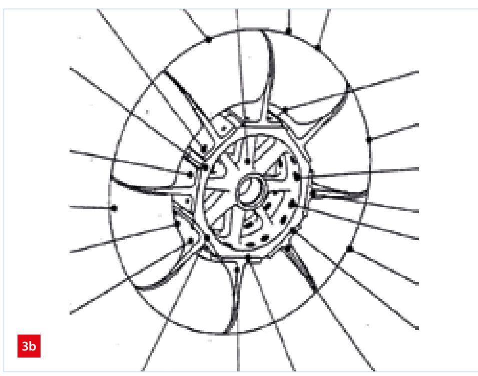

3b

US20080018167
Fuji, 2008

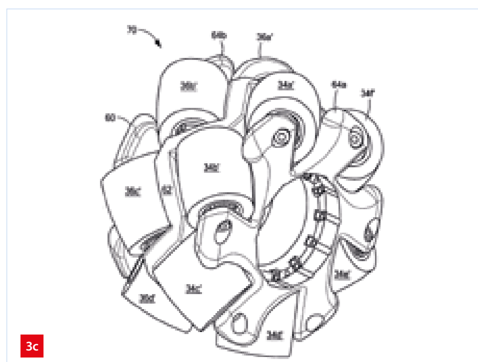

3c

US20100187779A1
Potter, 2009

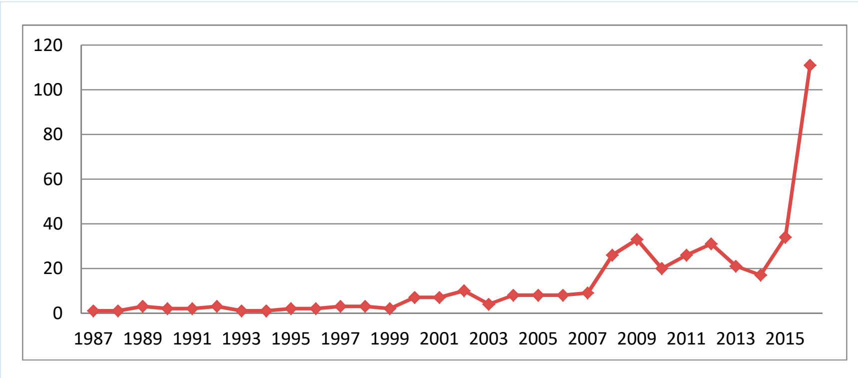

Graph 1. Omni wheel patent applications during the period 1987-2016 (unique inventions on oldest priority date). The highest number of patent applications in this period originates from Honda Motor Company [JP] with 41 applications. Japan and China are the most frequent countries of origin. The only applicant from the Netherlands is Philips with two applications.

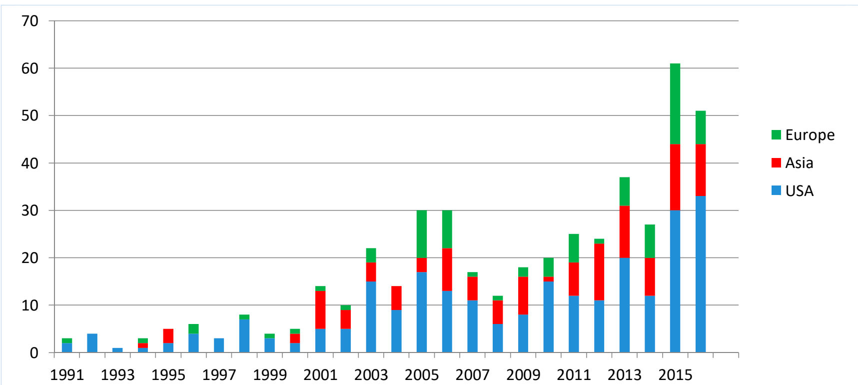

Graph 2. Number and geographical origin of patent applications for micro surgery during the period 1991-2016.

be used in your business, but you can learn from it and try to work around it or try to collaborate with Philips. Airtrax Inc filed patent applications explaining how the use of omni wheels can facilitate precise positioning of forklifts, warehouse vehicles or military equipment (Figure 6a). A number of patent applications on football robots using omni wheels were filed by Suzhou Nanjiang Lebo Robotics Co (Figure 6b and 6c).

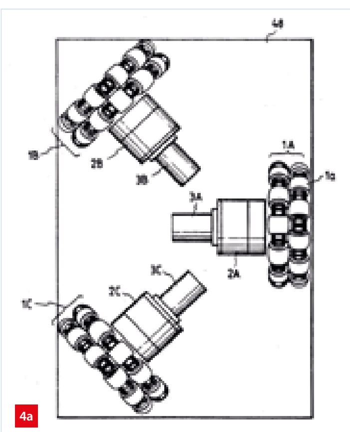

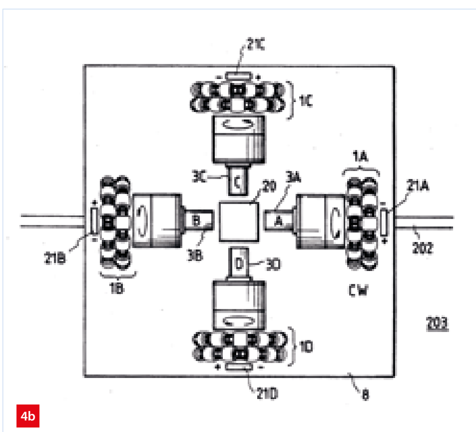

US5213176
Hitachi, 1989

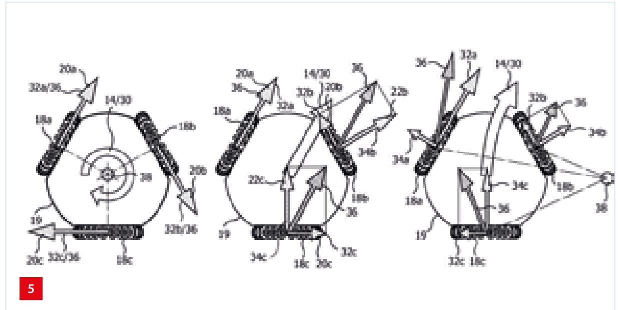

WO2011030255
Philips, 2009

Apart from robots, omni wheels are used in many other applications. From the 1980s we find omni wheels appearing in wheelchair patents (see for example Farnam, 1984; Figure 7a), but their use in current wheelchairs is still limited. The reason might be that in order to permit smooth rolling, the wheels should have a high number of casters and must be mounted at least two-by-two. Therefore, these wheels have a high complexity and are expensive.

These problems are addressed in a patent (2012) by New Live from France (Figure 7b) and by Whill from Japan in almost 20 patent applications between 2014 and 2018 (Figure 7c). One or two applications on a subject are not a guarantee that it will be brought to the market, but 20 applications for omni wheels in wheelchairs filed by a single company are such a large investment, that we can expect to find more and more omnidirectional wheelchairs on the streets in the future.

Competitor analysis

Honda
To really learn about omnidirectional wheels, take a look at the patent applications filed by Honda Motor Co between 2007 and 2016 relating to an inverted pendulum type vehicle, using an omni wheel (Figure 8). They wrote hundreds of pages explaining all kinds of details about how to make and use omni wheels and how they use them in their UNI-CUB personal mobility device. Numerous technical problems and their solutions, such as minimising vibration and noise and preventing dirt from being trapped between the rollers forming the wheels, are explained in detail.

Facebook, Disney and Amazon
Notable is that Amazon filed three patent applications for an internal transport system using special omni wheels that can move on rails as well as on a flat surface (Figure 9). This is understandable, looking at their business, where logistics and warehousing is extremely important. Disney filed

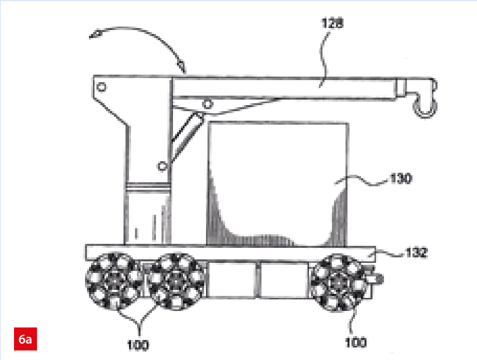

WO2006062905
Airtrax, 2004

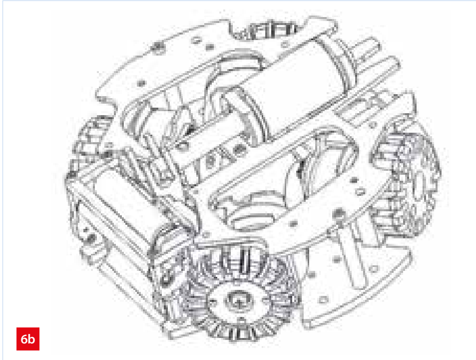

CN205924977U
Suzhou Nanjiang Lebo Robotics, 2016

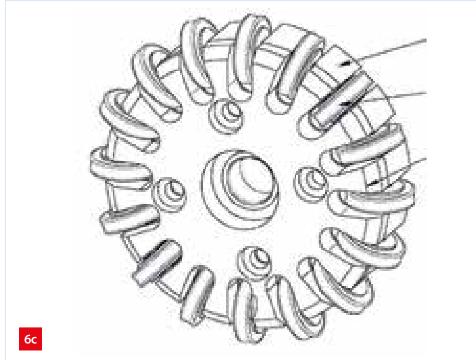

CN205618795U
Suzhou Nanjiang Lebo Robotics, 2016

a patent application for a multidirectional vehicle with omni wheels, to be used in amusement parks. But why is Facebook also filing a patent application concerning omni wheels (Figure 10)?

How to find interesting patents?

There are several free patent databases on the internet: Espacenet, Patentscope, Google Patents and others. Espacenet is the most complete one with over 100 million patent publications (= applications + granted patents) from more than 90 countries. It is hosted by the European Patent Office. Patentscope (hosted by the World Intellectual Property Organisation) contains around 65 million publications and has some extra features, like statistical analysis and chemical structure searching. Google Patents is less complete, but the presentation of the results is well-ordered. The disadvantage of Google Patents is that Google is a commercial company, so Google is free to do whatever it wants with search actions.

Patent searching should not only be done with keywords, because patent applications are written in general terms to protect as much technology as possible. Furthermore, applicants can use any term they want. An omni wheel can also be called a omnidirectional, multidirectional, mecanum or ilon wheel or just a wheel. Moreover, words like precise, precision, micro are not very discriminative and therefore not useful. Dimensions are hardly used in patent applications. The best way to search patent publications is by using classification codes. These are codes that indicate the technical subject of the invention. There are International Patent Classifications (IPC) and Cooperative Patent Classifications (CPC). They are allotted to the patent by the national or regional IP office that publishes the patent application. This code system is harmonised worldwide and updated regularly. In Espacenet there are IPC codes and CPC codes. CPC codes are more specific and used by the European and US patent offices. In Patentscope only IPC codes can be used.

Omni wheels

For omni wheels the most relevant CPC code is B60B19/003 = Multidirectional wheels.
Searching with this code yields patent publications that relate to an improvement of omni wheels, independent of the words that the applicant has used. Want to find Honda's omni wheel applications? Use Espacenet Smart Search and type: Honda B60B19/003.

Some patent publications have more than one IPC or CPC code, because the invention includes several technologies, e.g. the wheelchairs with improved manoeuvrability using omni-wheels like the examples above have several codes for wheelchairs as well as the code for omni wheels.

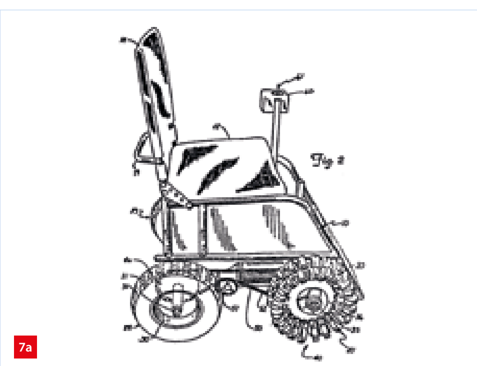

EP160631
Farnam, 1984

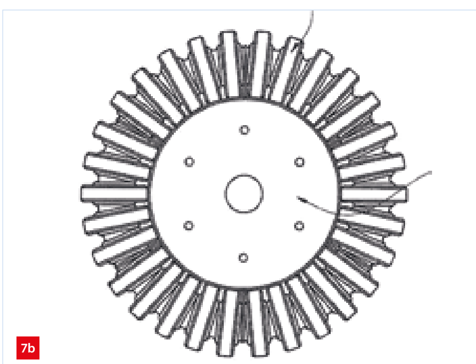

WO2013186489
New Live, 2012

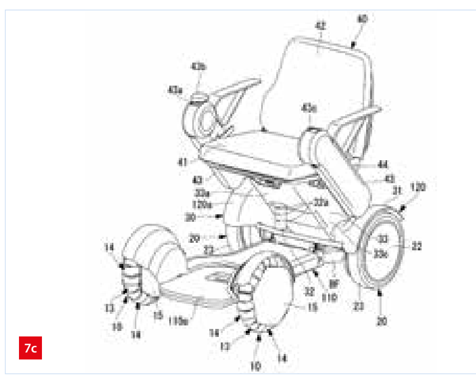

WO2018190388
Whill, 2017

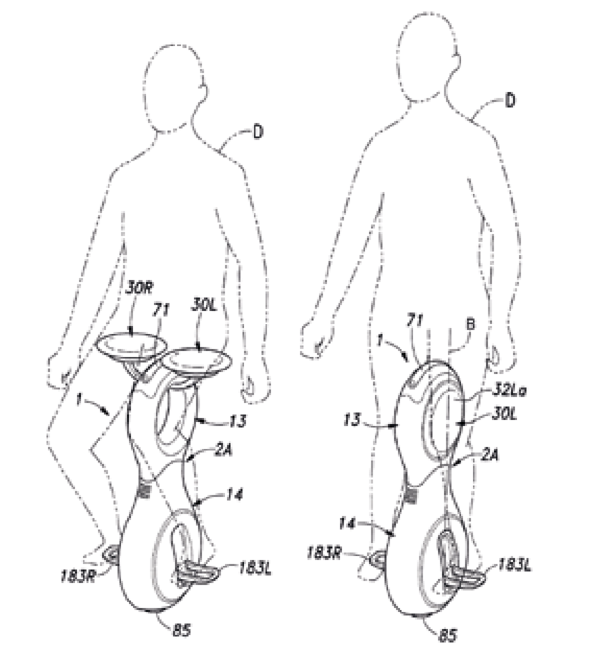

8

US2011067937
Honda Motor Co, 2011

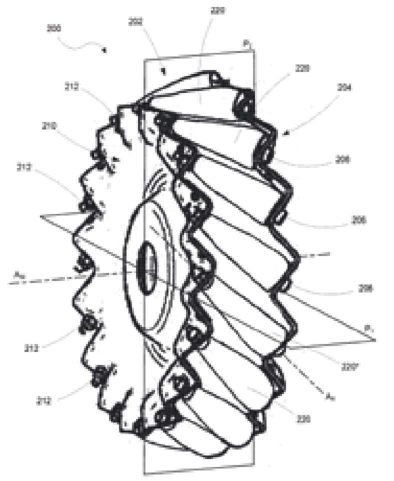

10

US2018229547
Facebook, 2018

Micro surgery
There are no specific codes for micro surgery. Therefore, finding relevant codes is more difficult in this case. A simple way of finding relevant codes is to start with keywords or names and try to find one or two relevant patent applications. Take a look at the meaning of the codes that are used for these relevant applications and continue searching with these codes. In this way we found the following CPC codes:

- A61B2034/303
  Surgical robots specifically adapted for manipulations within body lumens, e.g. within lumen of gut, spine or blood vessels.
- A61B34/32
  Surgical robots operating autonomously.
- A61B34/72
  Micromanipulators adapted for use in surgery.

Use the search string A61B2034/303 OR A61B34/32 OR A61B34/72 in the CPC field of Espacenet and combine it with names, dates, countries, keywords, etc.

Conclusion

Patent applications can tell a lot more than just whether or not a technology is protected by someone. They provide inspiration and technical and strategic information. The vast majority is freely applicable.

Information

For more information about patent searching take a look at the workshops provided by the Netherlands Patent Office at the address below (in Dutch), or by means of the Espacenet help files.

OCTROOIEN@RVO.NL
WWW.RVO.NL/OCTROOIWORKSHOPS
WORLDWIDE.ESPACENET.COM

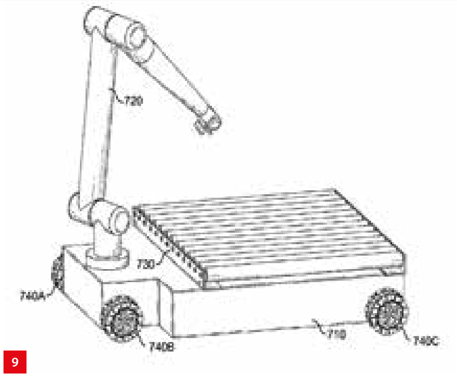

9

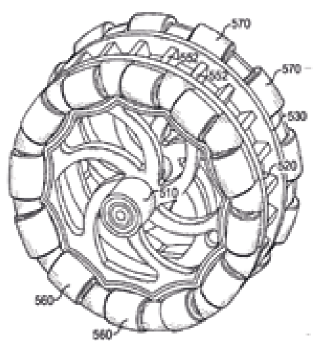

US2016375723
Amazon, 2017

GIVING ATHENA STABLE X-RAY VISION

SRON, the Netherlands Institute for Space Research, is participating in the development (led by the French space organisation) of the X-ray instrument X-IFU for the European space telescope Athena, which will be launched around 2030. Among other things, the SRON team is designing a Kevlar suspension to thermally isolate the camera and keep it in place.

ERIK ARENDS, SANDER KWAST AND JOHANNES DERCKSEN

Along with Hubble's famous pictures taken in visible light, astronomers take photographs in many other parts of the radiation spectrum, such as in radio, infrared or X-rays. This enables them to see unique features of the universe that are invisible to even the largest optical telescope imaginable. By using X-rays, for example, we can observe active galactic nuclei (AGNs) spitting out jets of gas and hot glowing matter being eaten by black holes.

Since the Earth's atmosphere absorbs X-rays, we need space telescopes to capture this dangerous but useful type of electromagnetic radiation. ESA, the European Space Agency, is currently developing X-ray satellite Athena (*sci.esa.int/athena*), to be launched around 2030. On-board will be two instruments: a Wide Field Imager (WFI) and an X-ray Integral Field Unit (X-IFU). In April 2018, SRON received a subsidy from NWO, the Dutch national research organisation, to participate in the development of the latter.

X-IFU will be the most sensitive X-ray space instrument ever developed. It will have a 2.5 eV spectral resolution, with 3,840 pixels of 0.25 cm x 0.25 mm, over a field of view of 5 arcminutes. To reach that level of sensitivity, it needs to be free of background noise at its operating temperature of 50 mK. This means that only low thermal conduction from the relatively warm telescope can be allowed, so the camera needs to be mounted using a string suspension.

As one would expect in a space mission, the camera's parts are not just floating around in free space, hanging on a leash. To achieve the specified sensitivity the X-ray camera must be in a well-determined fixed position. In other words, SRON engineers must tighten the suspension strings firmly enough to ensure that the camera's movement is only minimal. They have already found the right material to use: Kevlar, as it is both strong and has a low thermal conductivity. However, this still leaves some concerns.

AUTHORS' NOTE

Erik Arends, Sander Kwast and Johannes Dercksen are all associated with the Netherlands Institute for Space Research, SRON, based in Utrecht and Groningen, the Netherlands.

e.arends@sron.nl
www.sron.nl

Mechanical vibrations

The X-ray camera contains a cryostat to keep it cooled to below 50 mK. Like any refrigerator, cryostats are all about pumping gases around, inevitably leading to vibrations, which are then picked up by the strings, causing them to oscillate. This oscillation is one of the concerns, because these strings will in turn transfer vibrations to the camera. Also, the strings themselves will heat up from the oscillations, and this additional heat will also be absorbed by the camera. SRON engineers are faced with the challenge of minimising the transfer of vibrations to the Kevlar strings.

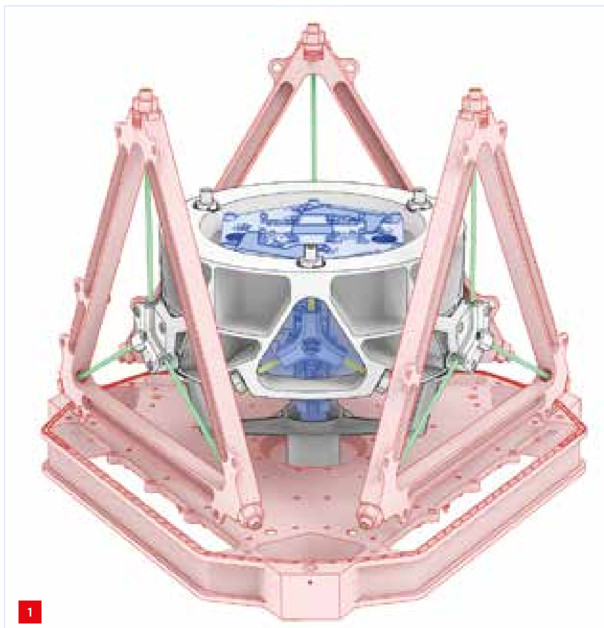

1

The camera is suspended from three points. Each point is connected to an external framework by three strings (yellow). This framework has three suspension points of its own, which in turn are connected to a framework by three Kevlar strings (green). The extra stage is needed for the read-out chain, i.e. the electrical components. Some have an optimal operating temperature of 50 mK, others of 300 mK and yet others of 2 K. This is related to both the dissipation within those components and the noise, which is temperature-dependent.

The more firmly the cords are tightened, the more they will pick up vibrations from the machine they are connected to. The camera is sensitive to thermal fluctuations of the cooler, because the pixels measure temperature increase due to X-ray absorption. For sufficiently accurate measurements, these fluctuations need to be smaller than 10^{-6} K. By separating the natural resonance frequencies of the two suspended stages (see Figure 1), the suspension can be made to act as a filter for vibrations, thereby minimising the amount of heat that is generated in the suspension.

Launch

Any rocket launch goes hand-in-hand with enormous vibrations; in the case of the X-IFU, up to 12.7 g_{rms}. This makes tight strings good candidates for breaking. To survive the launch, the Kevlar cords need to be able to handle more force than the 2.5 kN tension that is needed for the thermal suspension. The design margin for the X-IFU instrument dictates a capacity of more than 5 kN, with a requirement that there should remain almost 2 kN of pretension at the moment of launch. This is where the practicalities of a real-life space mission come into play. Before launch we have to perform several functional tests, including some at cryogenic temperatures. During these cooling downs, the frames of the suspension contract while the Kevlar expands (due to its negative thermal expansion coefficient). The pretension will then decrease dramatically. To solve this issue, disc springs were designed to make sure that there is always pretension, even at the lowest temperatures.

End-fitting

To make something stronger, one must first secure the weakest link. In the Kevlar suspension, this is the end-fitting. To make it at least as strong as the Kevlar cord, the string filaments were untwisted at both ends and put in a hole injected with adhesive. Untwisting increases the total surface area of the Kevlar that is attached to the end-fitting. Furthermore, the hole was designed to have a conical shape, so when the cord is pulled, the adhesive is pulled into the conical hole and the grip tightens. In the end, 100 per cent of the Kevlar cord bulk strength was reached. SRON, together with NTS Mecon and funded by ESA, is developing a reliable and reproducible manufacturing and assembly process for these end-fittings.

18 possible failures

The camera is suspended at three points. Each of those is connected with an external framework by three Kevlar strings. This framework has three suspension points of its own, which again are connected to a framework by three Kevlar strings. This brings the total number of strings to 18 (2 x 3 x 3). Each string is a single-point failure: if one of them loses its pretension, then the entire mission is at risk. With one loose string, the launch loads will probably destroy the detector and/or a thermal shortcut could occur, making it impossible for the cooler to bring the detector to its operating temperature of 50 mK. This science mission is literally hanging by a thread.

Current status

So far, we have been able to manufacture end-fittings that are stronger than the Kevlar cord itself. We learned that this was not as easy as it may sound, as Kevlar cord is strong and difficult to get a grip on. The next steps, with NTS Mecon, are to develop a reproducible manufacturing and assembly process, incorporating end-fitting assemblies in a (dummy) suspension system and putting it to a vibration test to simulate the launch loads.

In parallel to this, next year SRON will build a development model of the Focal Plane Assembly (FPA), including the suspension system, to verify the performance of the camera system (Figure 2). For this, the FPA will be integrated into a cooler system, so that the influence of micro-vibrations will also be assessed. The development model will not be tested for launch (macro-) vibrations.

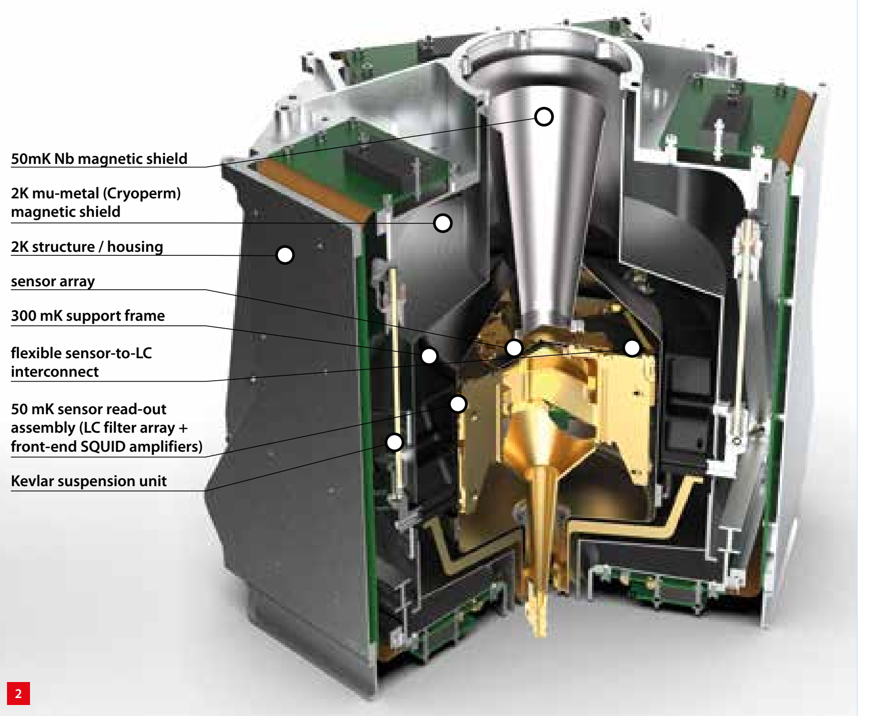

2

The X-IFU Focal Plane Assembly development model.

MUCH MORE THAN PRECISION COMPONENTS ONLY

Most of 290-plus stands at the mid-November fair in Veldhoven in the Netherlands show the capabilities of cutting specialists by conspicuous slogans such as "precision components", "precision parts" or "precision products". This could be considered redundant at an event which calls itself the Precision Fair. Only a small minority of stands demonstrates 3D metal printing disciplines to visitors (over 3,400 in two days). Unlike the angle taken in previous Precision Fair reports, this time it seems a good idea to primarily consider the non-cutting and non-printing exhibitors.

FRANS ZUURVEEN

The word precision is usually associated with tiny objects. But very often large products can also be considered as precision objects, albeit not always in an absolute, but rather in a relative sense. Figure 1 shows a detail of a large gear box exhibited by Schaeffler Nederland. Such gear boxes are used in the currently widely installed windmills which produce energy without emitting CO_2. The roller bearing in the picture is of the cageless type, which means that subsequent rollers touch each other. This effect causes mutual friction and thus a lower bearing efficiency. However, the Schaeffler representative explains that, in the case of lower rotational speeds, more bearing load capability is preferred above efficiency. This is why other bearings with higher speeds in this box do have bronze cages.

Another not so tiny large precision product is the EC-i 275 frameless motor with thirty coils in the stator, developed by the Enschede office of maxon motor benelux, see Figure 2. Maxon is well-known for its small servo motors with up to 100,000 rpm. The next generation of high-torque motors consists of motion solutions that provide maximum design flexibility. With a new line of frameless BLDC (brushless DC) flat motors (45-, 60- and 90-mm diameter) and the EC-i 275 motor, maxon is responding to the need for customer-specific applications.

Another interesting large object is the 3D-printed rocket nozzle displayed by the DARE team (Delft Aerospace Rocket Engineering) from Delft University of Technology, see Figure 3. During the rocket launch, the nozzle has to withstand temperatures up to 3,000 °C. An earlier nozzle design from carbon did not succeed in retaining its shape during tests. This titanium product, supplied by 3D Systems, and provided with a ceramic coating of ZrO_2, is expected to do its job during forthcoming tests in December.

1 A cageless Schaeffler high-load roller bearing in a gearbox for windmills.

2 A multi-coil stator for the EC-I 275 high-torque frameless maxon motor.

AUTHOR'S NOTE

Frans Zuurveen, former editor of Philips Technical Review, is a freelance writer who lives in Vlissingen, the Netherlands.

A 3D-printed rocket nozzle made from titanium, designed by the DARE team from Delft University of Technology.

A stainless steel workpiece made by Lucassen Watersnijtechniek using water cutting technology.

Various production skills

As mentioned before, lots of exhibitors demonstrate their precision cutting skills, but some stands have alternative fabrication technologies on display. For example, MIFA Aluminium shows their precision-extrusion capabilities. The accuracy of the extrusion technology is not fully comparable with milling and grinding, but nevertheless the mastering of the extrusion of small components by MIFA results in accuracies in the order of ±20 µm. Figure 4 shows the required tools and the resulting final product.

KUK Wijdeven is specialised in the production of electromagnetic components such as transformers and transducers. It applies orthocyclic winding technology, developed in the previous century by Philips Electronics, to reach, e.g., a better copper wire fill factor in deflection coils for colour-image tubes. The technology involves wires in an upper wire layer fitting exactly into the grooves of the preceding layer.

Lucassen Watersnijtechniek has mastered water-cutting technology. This cannot be really regarded as precision technology, but the ability to cut steel sheets with a thickness up to 100 mm, see Figure 5, makes this worth mentioning. Lucassen uses Swiss machines which apply one or two jets consisting of a mixture of water with an abrasive, mostly corundum, i.e. Al_2O_3. In addition to steel, a wide range of materials can be water-cut, with the only condition that the material does not absorb fluid. The Lucassen stand also displays real precision products, including some made by wire-spark erosion, for example.

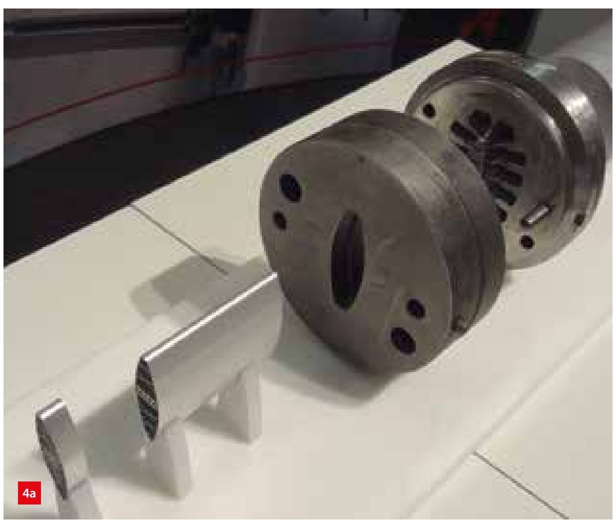

MIFA extrusion technology.
(a) Tools.
(b) The resulting accurate demonstration product.

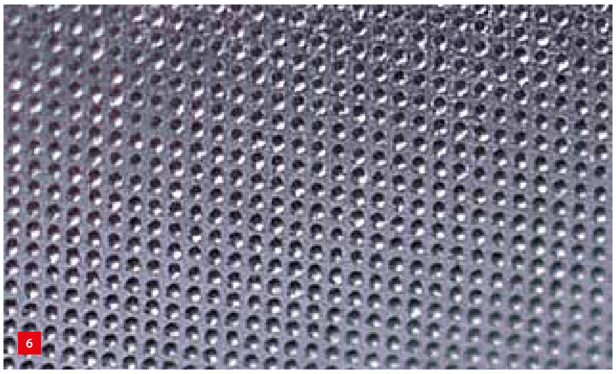

Holes measuring 60 µm in diameter made with a pulsed fibre laser at a rate of 1,500 holes per second by Raytech.

Tiny holes and pits

It is very interesting to discover specialists that make extremely small holes or pits with pulsed lasers. Raytech from Belgium shows visitors to its stand a stainless steel tube with 28,000 holes with a diameter of only 60 µm. The holes in this giveaway test product were made with a pulsed fibre laser at a rate of 1,500 holes per second, see Figure 6. The hole pitch amounts to 100 µm.

Another provider of microscale laser ablation is Lightmotif, a spin-off company from the University of Twente. This innovative company supplies ultrashort-pulse laser micromachining systems and processes. Figure 7a shows a surface texture produced with the Lightmotif 5-axis laser micromachining system shown in Figure 7b.

This CAM-based machine combines step-and-scan driving with a synchronised pulsed picosecond laser. The design aims at providing surfaces with a special texture which cannot be realised with conventional cutting machines. Such textures can be be applied in plastic injection moulding tools, for example. The laser pulses from 200 fs to 10 ps ($2 \cdot 10^{-13}$ to 10^{-11} s) duration evaporate the surface material to a depth of 10 to 50 nm. The wavelength of the laser light is 532 nm, i.e. visible green radiation.

Moving and measuring

As usual, many exhibitors present Schneeberger-like linear movement systems. One of them is THK GmbH, originally from Japan. But in comparison to its competitors, THK displays curved raceways with a wide selection of radii, see Figure 8a. Instead of straight movements, a customer thus can create nonlinear orbits for his moving slides. And THK shows more than these large guiding systems. It also produces sub-miniature linear guides, comprising balls with diameters down to a mere 200 µm, see Figure 8b, which is regarded as highly challenging.

Such movements have to be accurately measured. Masters in these tasks are Renishaw and Heidenhain, by tradition both prominently present at the fair. Heidenhain shows its LIP 200 measuring system with a nearly unimaginable resolution of 31.25 pm. This resolution results from interpolation within a signal period of 0.512 µm of Optodur phase gratings on Zerodur glass ceramic material.

This is the right moment to note the ease with which standholders discuss nanometer accuracy. One of them displays some kind of tool. After some interrogation he claims an accuracy of 6 nm for this tool. However, this assertion results from his assumption that a nanometer equals 1/10 of a micrometer. The statement that 1 nm equals no more than only 1/1000 of a micrometer is obviously new information to him.

But continuing with sophisticated measuring systems, we have to mention Werth Messtechnik, which shows its well-known contactless laser probe and a new 192-point white-light line sensor. This sensor simultaneously scans 192 measuring points on a line perpendicular to the direction of movement.

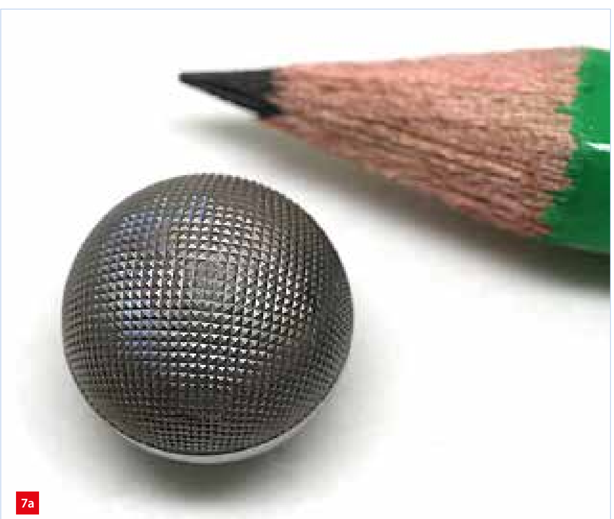

Microscale laser ablation by Lightmotif.
(a) A texture with tiny pits made with an ultrashort pulse-laser.
(b) The Lightmotif 5-axis laser micromachining system.

8a

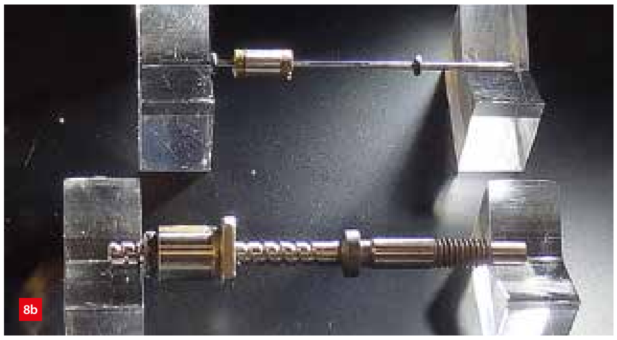

8b

THK Guides.
(a) Curved raceways, available with different radii to produce non-linear guides.
(b) Sub-miniature THK linear guides with balls with a diameter down to 200 μm.

Cutting and printing

Of course, we cannot avoid giving attention to the cutting specialists who demonstrate their skills with precision components, parts or products on display. Figure 9 shows a product machined from difficult-to-machine massive titanium by Mevi Fijnmechanische Industrie. Understandably, its destination is kept secret. The tolerances stated are minimally 5 μm.

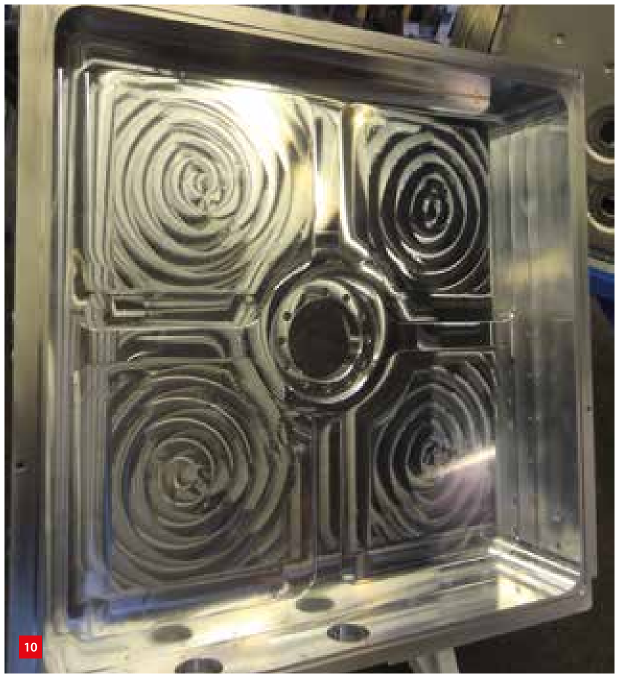

10

A part of a vacuum chamber produced by Vernooy Vacuum Engineering, machined from a massive block of aluminium.

Another impressive item is a vacuum chamber produced by Vernooy Vacuum Engineering, part of the Masévon Group, see Figure 10. This item is cut from a massive block of aluminium, specially selected to have no pores or other cavities which might absorb water or other contaminations. Of course, the internal chamber surface is exceptionally smooth and flat to trap impurities as little as possible. In this regard it is quite remarkable that the internal surface shows traces of the cutting process, whereas the roughness approaches nearly zero values. Obviously these traces originate from crystalline effects due to influencing the atomic grid during cutting.

The last picture, Figure 11, shows printed products made by 3D Systems. Even so, very few printed products are displayed at this fair. Cutting specialists who are asked about their willingness to adopt 3D printing technology express their hesitation: it is too slow, too inaccurate, too rough. Nevertheless they have to admit that some products are really impossible to make with conventional cutting processes. The ones shown in this last illustration underline this.

9

A product machined from titanium by Mevi Fijnmechanische Industrie.

11

A series of printed products produced by 3D Systems.

INFORMATION
WWW.PRECISIEBEURS.NL

PRIZES FOR NATURAL OPTOMECHANICAL DESIGNER AND COMMUNICATIVE DESIGN TALENT

During the 18th edition of the Precision Fair in Veldhoven, the Netherlands, the Ir. A. Davidson Award and the Wim van der Hoek Award were presented under the auspices of DSPE. Lennino Cacace, the owner of AC Optomechanix, received the Ir. A. Davidson Award 2018 for his outstanding designs of complex optical systems and sensors, and his contribution to education in optomechatronics. The Wim van der Hoek Award went to Martin Kristelijn, who graduated from Eindhoven University of Technology on a well-elaborated and clearly presented design of a motion-compensation and load-transfer mechanism for application in the offshore industry.

Encouraging young talent

DSPE board member Toon Hermans, the managing director of Demcon South, presented the Ir. A. Davidson Award on the afternoon of Wednesday 14 November. The purpose of the prize is to encourage young talent by recognising the efforts of a precision engineer with several years of experience working at a company or institute and a proven performance record which has been acknowledged internally and externally. Candidates must also have a demonstrated enthusiasm for the field that produces a positive effect on young colleagues.

The biennial prize, which was established in 2005, is named after an authority in the field of precision mechanics who worked at Philips in the 1950s and 1960s. The prize comes with a certificate, trophy and sum of money. The trophy was created by the Leiden Instrument Makers School (LiS) in the form of the handbook in precision mechanics that Davidson used as a foundation when forming the constructors community at Philips.

Complex optics

This year, the panel of judges for the Ir. A. Davidson Award received a large number of nominations of excellent candidates, from which they unanimously selected Lennino Cacace as the winner (Figure 1). He studied Mechanical Engineering at Eindhoven University of Technology (TU/e) and graduated cum laude. In 2009, he obtained his Ph.D. degree from Maarten Steinbuch, a TU/e professor of control engineering, for the design of a measuring head for the advanced measuring machine Nanomefos, developed by TNO for freeform metrology. In 2006, he had already established his own company, AC Optomechanix, an engineering firm specialising in optomechanics and precision engineering.

Cacace has developed various complex sensors and has worked for, among others, the lithography machine builders ASML and Liteq. At the latter company, now part of Kulicke & Soffa, he fulfilled the role of optical-mechatronic system

1a

Ir. A. Davidson Award 2018 winner Lennino Cacace (right) being congratulated by DSPE board member Toon Hermans. (All photos: Mikrocentrum)

1b

LiS director, Dick Harms, handing over the Ir. A. Davidson Award trophy, made by LiS students, to award winner Lennino Cacace.

engineer for Liteq's waferstepper for packaging chips. He also was responsible for a large part of the optical design, including the mechatronic design of the optical column of the machine.

Education in optomechatronics

As of the last year and a half, Cacace has been involved in the new Master of Optomechatronics that was set up at TU Delft. In his own lecture on optomechatronics, he gives a good overview of the system aspects and the constructive, optical and even electronic aspects of designing an optical system. In future, this course will also be given in Eindhoven, in a form that is tailored to the industry.

"There are few mechanical engineers who have already studied optics so well during their studies. At the same time, Lennino is a good designer. What is also special is that he has a clear vision of the profession and has succeeded in incorporating this into the optomechatronics master and lecture." So said the panel of judges, who characterise Lennino Cacace as an inquisitive, helpful and very honest natural designer and a perfectionist. "He is highly skilled in theory and in practice and experimentation and is keen to pass on his knowledge."

Constructors Award

The second day of the Precision Fair 2018, Thursday 15 November, featured the presentation of the Wim van der Hoek Award. This award (also known as the Constructors Award) was introduced in 2006 to mark the 80th birthday of the Dutch doyen of design engineering principles, Wim van der Hoek.

The Constructors Award is presented every year to the person with the best graduation project in the field of design in mechanical engineering at the Dutch (and Belgian) universities of technology and universities of applied sciences. This award includes a certificate, a trophy made by LiS (Leiden Instrument Makers School) and a sum of money (sponsored by HTSC; the High Tech Systems Center at TU/e).

Five nominations

Criteria for the assessment of the graduation theses include quality of the design, substantiation and innovativeness, as well as the suitability for use as teaching materials. The panel of judges, under the presidency of DSPE board member Jos Gunsing (MaromeTech), received five nominations (for a total of six persons), submitted by the graduation supervisor/professor of each student concerned (Figure 2). Two universities of applied sciences (UAS) and two universities nominated candidates: AVANS Hogeschool 's-Hertogenbosch, UAS Utrecht, KU Leuven (Belgium) and TU/e.

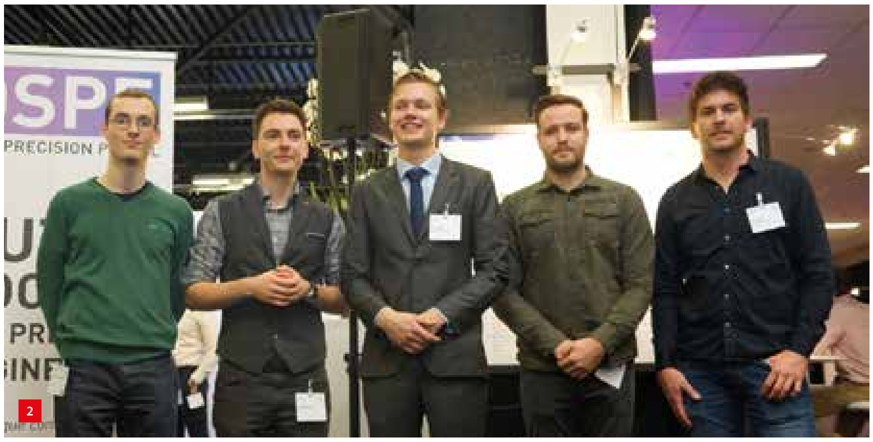

Five (from the total of six) nominees for the Wim van der Hoek Award 2018, from left to right: Bert Van Raemdonck (KU Leuven), Martin Kristelijn (TU/e), Jens de Goeij (UAS Utrecht), Nick Toonen (UAS AVANS) and Roy Jacobs (TU/e).

"Seemingly easy"

Following a careful assessment, the Wim van der Hoek Award 2018 went to Martin Kristelijn, who studied Mechanical Engineering at TU/e (Figure 3). This spring he graduated with the design of a motion-compensation and load-transfer mechanism for application in the offshore industry (for payloads up to 1,000 metric tons). The panel of judges described his design as well elaborated and succinctly but clearly presented.

"Martin Kristelijn correctly combined the relevant construction principles for his design and analysed the degrees of freedom in the right manner. He provides real insight with 'back of an envelope' calculations, step by step, so well that it becomes seemingly easy. He involves the reader in his reasoning and thus delivers excellent communication."

Martin Krijstelijn, winner of the Wim van der Hoek Award 2018, showing the certificate and the trophy he has just received from DSPE board member and panel of judges chairman Jos Gunsing.

ECP² COURSE CALENDAR

COURSE (content partner)	ECP² points	Provider	Starting date
FOUNDATION			
Mechatronics System Design - part 1 (MA)	5	HTI	8 April 2019
Fundamentals of Metrology	4	NPL	to be planned
Mechatronics System Design - part 2 (MA)	5	HTI	4 November 2019
Design Principles	3	MC	13 March 2019
System Architecting (S&SA)	5	HTI	11 March 2019
Design Principles for Precision Engineering (MA)	5	HTI	17 June 2019
Motion Control Tuning (MA)	6	HTI	6 February 2019
ADVANCED			
Metrology and Calibration of Mechatronic Systems (MA)	3	HTI	13 February 2019
Surface Metrology; Instrumentation and Characterisation	3	HUD	-
Actuation and Power Electronics (MA)	3	HTI	to be planned
Thermal Effects in Mechatronic Systems (MA)	3	HTI	19 March 2019
Summer school Opto-Mechatronics (DSPE/MA)	5	HTI	-
Dynamics and Modelling (MA)	3	HTI	25 November 2019
Manufacturability	5	LiS	to be planned
Green Belt Design for Six Sigma	4	HI	to be planned
RF1 Life Data Analysis and Reliability Testing	3	HI	1 April 2019
SPECIFIC			
Applied Optics (T2Prof)	6.5	HTI	5 February 2019
Advanced Optics	6.5	MC	28 February 2019
Machine Vision for Mechatronic Systems (MA)	2	HTI	to be planned (Q2 2019)
Electronics for Non-Electronic Engineers – Analog (T2Prof)	6	HTI	to be planned
Electronics for Non-Electronic Engineers – Digital (T2Prof)	4	HTI	11 February 2019
Modern Optics for Optical Designers (T2Prof)	10	HTI	20 September 2019
Tribology	4	MC	12 March 2019
Basics & Design Principles for Ultra-Clean Vacuum (MA)	4	HTI	to be planned
Experimental Techniques in Mechatronics (MA)	3	HTI	to be planned (Q2 2019)
Advanced Motion Control (MA)	5	HTI	18 November 2019
Advanced Feedforward Control (MA)	2	HTI	9 October 2019
Advanced Mechatronic System Design (MA)	6	HTI	to be planned (Q3/4 2019)
Finite Element Method	5	ENG	in-company
Design for Manufacturing – Design Decision Method	3	SCHOUT	in-company
Precision Engineering Industrial Short Course	5	CRANF	11 February 2019

ECP² program powered by euspen

The European Certified Precision Engineering Course Program (ECP²) has been developed to meet the demands in the market for continuous professional development and training of post-academic engineers (B.Sc. or M.Sc. with 2-10 years of work experience) within the fields of precision engineering and nanotechnology. They can earn certification points by following selected courses. Once participants have earned a total of 45 points, they will be certified. The ECP² certificate is an industrial standard for professional recognition and acknowledgement of precision engineering-related knowledge and skills, and allows the use of the ECP² title.

ECP2EU.WPENGINE.COM

Course providers

- Engenia (ENG) WWW.ENGENIA.NL
- The High Tech Institute (HTI) WWW.HIGHTECHINSTITUTE.NL
- Mikrocentrum (MC) WWW.MIKROCENTRUM.NL
- LiS Academy (LiS) WWW.LISACADEMY.NL
- Schout DfM (SCHOUT) WWW.SCHOUT.EU
- Holland Innovative (HI) WWW.HOLLANDINNOVATIVE.NL
- Cranfield University (CRANF) WWW.CRANFIELD.AC.UK
- Univ. of Huddersfield (HUD) WWW.HUD.AC.UK
- National Physical Lab. (NPL) WWW.NPL.CO.UK

Content partners

- DSPE WWW.DSPE.NL
- Mechatronics Academy (MA) WWW.MECHATRONICS-ACADEMY.NL
- Technical Training for Prof. (T2Prof) WWW.T2PROF.NL
- Systems & Software Academy (S&SA)

NEWS

Next-generation design principles textbook

Earlier this year, the textbook *Design Concepts for Precision Engineering* was published by Susan van den Berg, M.Sc. Mech. Eng., a lecturer at Fontys University of Applied Sciences (UAS) in Eindhoven, the Netherlands. The book is targeted at UAS students and engineers who are novices to the field of precision engineering. Only prior knowledge of mechanics (statics, stress and deformation) and basic mathematics (trigonometry and calculus) is required.

The new textbook is the first of two volumes and it aims to explain the (physical) phenomena that interfere with achieving high-precision goals (part I) and to provide useful tools for analysing and designing systems (part II). The second volume, to be published next year, will present (many of) the design concepts that help to avoid some of the more harmful phenomena (part III).
The content of the books is largely based on the work begun by Wim van der Hoek, who is considered the founder of the precision engineering field in the Eindhoven region. As a professor at Eindhoven University of Technology, he collected design principles inspired by unsuccessful mechanical designs in *Des Duivels Prentenboek* (in Dutch), which is no longer available.

Van der Hoek's work was continued and extended by Rien Koster, professor at the University of Twente, in his book *Constructieprincipes* (in Dutch), which is aimed at engineers who have already been introduced to the field. His successor, Herman Soemers, published a translated version, entitled *Design Principles for Precision Mechanisms*, which was abridged in some parts and extended in others (by incorporating mechatronics).
The last two books are available, but both require an academic level of understanding. When Susan van den Berg noticed a mismatch between these existing books and the entry level of the students in her class at Fontys UAS, she took the initiative to write a book herself targeted specifically at UAS students. The forthcoming June issue of *Mikroniek*, with the theme "Design principles", will dive deeper into the two volumes and the author's underlying philosophy.

WWW.BERGPRECISIONPUBLISHING.NL

Software-centred Smart Industry fieldlab kicks off in 2019

On 13 March 2019, the Software Competence Centre will officially open as a Smart Industry fieldlab. This high-tech software fieldlab will be based in Eindhoven, the Netherlands, and will work on innovating in software and using software to drive innovation in industry, including areas such as digital twinning and model-based engineering.

The fieldlab is an initiative of a consortium of more than 20 high-tech software companies in the areas of virtual prototyping & design, model-based software and data analytics & services. Partners include Alten, Cordis Automation, Heemskerk Innovative Technology, ICT Automatisering Nederland, Sofon, TomTom, Topic Software Development, Unit040 Ontwerp and Verum Software Tools.

The Software Competence Centre has its origins in the High Tech Software Cluster, facilitated by Brainport Development. The Eindhoven-centred Brainport region produces the most advanced industrial software in the world, software that many innovative companies – for which this region is famous – rely upon for their complex products. The cluster's ambition is to contribute to shortening time-to-market and to help prevent complex development projects become unverifiable and uncontrollable. The cluster is built around four themes: visibility, collaboration & technology, internationalisation, and the labour market.

W.RENDERS@BRAINPORTDEVELOPMENT.NL (WIM RENDERS)
WWW.BRAINPORT.NL/HIGHTECHSOFTWARECLUSTER
WWW.SMARTINDUSTRY.NL/FIELDLABS

Readers' survey results

To gain insight into the reception of *Mikroniek* and further align *Mikroniek*'s editorial focus and standard to the information needs and wishes of its (potential) readership, the editorial board has conducted a small-scale readers' survey.

Half of the respondents were engineers, whereas researchers, (non-) technical managers and others accounted for the other half. Regarding the type of organisation they work in, the distribution was 1/3, 1/3, 1/3 for OEMs, suppliers and educational/research institutions, respectively. A large majority indicated to read (nearly) every issue, with one third reading (nearly) every article in every issue and the rest some of the articles. The survey reveals a clear preference for the paper edition, with none of the readers opting for online-only.

Regarding the attention reserved for the various types of articles, the average score was 'adequate'. Articles on design & realisation projects and manufacturing technology deserve more attention, so the overall response, whereas (short) event reports could do with less. Overall, *Mikroniek*'s performance was assessed as 'good', with relevance and depth each receiving an 'excellent' score from one third of the respondents. The newsworthiness seems to be somewhat questionable, as nearly 50% of the readers rated this only as 'acceptable'.

The editorial board will use these results to 'refresh' *Mikroniek*, along with the restyling of the layout.

New professor of Smart Manufacturing & Integrated Systems Engineering

As per 1-10-2018, Hans Krikhaar, CEO of Hurli, started at the Fontys University of Applied Sciences as a professor (lecturer) of Smart Manufacturing & Integrated Systems Engineering. Within the department of Engineering in Eindhoven and the department of Technology and Logistics in Venlo, he will guide research into mechanical engineering, electronics, mechatronics and automotive with links to ICT, industrial design, medical technology and business.
The applied research will strengthen the innovation competences of the university and their lecturers and will have direct benefits for the students involved. Collaboration with the industry and institutes like Eindhoven University of Technology, TNO, Brainport Industries Campus and Eindhoven Engine is crucial. Hans Krikhaar will continue to chair DSPE.

WWW.FONTYS.NL

New embedded secure coding training

High Tech Institute, which specialises in professional education for high-tech industries, has recently entered into a collaboration with the Hungarian Secure coding Academy (Scademy), a specialist organisation that focuses fully on training people to write secure code. In October, the "C and C++ secure coding" training was successfully given in Eindhoven, the Netherlands. The next training session is scheduled for 9-11 April 2019 in the same city.

Hungarian company Scademy offers almost forty courses; security for embedded systems is its specialty. Teaching people to write secure code is a labour-intensive endeavour. "The goal is not to teach people how to hack, but to instil paranoia. That emotion is important", says trainer Ernõ Jeges. "It has an impact which you will not get with online training courses."

WWW.HIGHTECHINSTITUTE.NL/SECURE-CODING-TRAINING

Laser Technology Janssen is expanding machining facilities

Two years ago, Laser Technology Janssen (LTJ) from Wijchen, the Netherlands, started with micro-lasering / fine-cutting of precision-formed products from thin sheet and foil material on a five-axis fibre laser. The mechanical department of LTJ was further expanded this year. Now LTJ can also carry out operations such as ultrasonic cleaning and passivation, in addition to fine-machining operations. As a result, LTJ is able to supply small, accurate products with short delivery times and with the right mechanical finishing and surface treatment, ready for installation.
In addition, five-axis fibre-laser micro-lasering / fine-cutting remains an important processing method. In addition to very precise contour cutting, holes and slits of a few tenths of a mm can also be applied with the micro-laser, with tolerances down to a mere ±3 μm. Initially, the choice of materials was limited, but LTJ can now also process materials such as stainless steel, titanium, copper, phosphor bronze, aluminium and various other alloys.

"We also do (micro) bending of sheets/foils with thicknesses of 0.02 to 1.5 mm in-house; this is unique to the market", declares founder/director Toon Janssen. LTJ also specialises in 2D and 3D (micro) laser welding of thin sheets and foils, minimising heat input, as this affects shape and dimensions. "Due to extremely accurate parameter control, we are able to weld super-thin materials without thermal stress."

(Source: www.fpt-vimag.nl/actueel)

WWW.LASERTECHNOLOGYJANSSEN.NL

Towards ever-smaller bore holes

Posalux, based in Biel, Switzerland, has introduced a newly designed machine for electronics fabrication that employs Scanlab's highly integrated precSYS 5-axis scan sub-system. The Swiss manufacturer's laser-processing system is specifically tailored to demands of micromachining and is also usable for processing such challenging materials as polymers and ceramics. The machine features precSYS, which enables ultra-precise, high-dynamic beam deflection for guiding the laser spot onto workpieces. New standards in micro-drilling precision are set by results obtained for fabricating electronics test equipment, where bore hole corner radii smaller than 5 μm are now possible.

The integrated sub-system by Scanlab, based in Puchheim, Germany, coupled with an ultra-short-pulse laser, enables processing of highly diverse materials such as metals, polymers and ceramics, without affecting them thermally. The scan solution provides five axes for defined laser beam guidance in the machine's x,y,z coordinate axes and a simultaneously superimposed, adjustable angle of incidence (positive or negative). This makes it ideal for fabricating micro-bore holes with high aspect ratios and freely definable geometries. And the intuitive user interface lets machine operators easily load a bore image, assign process parameters and scan the workpiece surface.

In one application case, Posalux's precSYS infrared machine produced bore holes specified for 30 μm x 30 μm edge lengths, 300 μm material thickness and 10 μm separation. Some 46,000 bore holes were examined for positioning accuracy of ±2 μm, and corner radii smaller than 5 μm were reliably achieved.

The precSYS sub-system is currently produced exclusively for infrared lasers with a 1,030 nm wavelength. A new variant is being developed for green lasers with a 515 nm wavelength, to thereby enable even finer structures and corner radii.

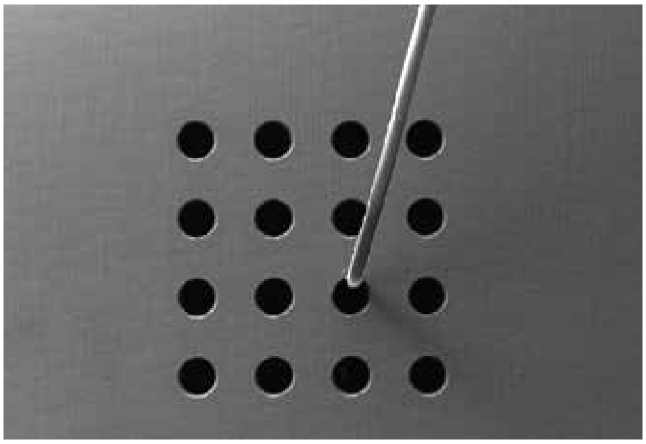

200-μm bore holes drilled in steel with precSYS, compared to a human hair.

WWW.SCANLAB.DE

Industry 4.0 cleanroom

Connect 2 Cleanrooms (C2C) has realised an intelligent cleanroom solution for medical device manufacturing at Promolding, a polymer product manufacturer based in Den Haag, the Netherlands. The production set-up involved consists of two injection moulding machines and robots for product handling, with which the cleanroom solution has to work in harmony. Effective contamination control is critical, as Promolding is producing optical manifolds for an eye-surgery machine.

C2C, which has its head office in Lancaster, UK, and a regional office in Utrecht, the Netherlands, applied Industry 4.0/Smart Industry techniques and designed a cleanroom with automated canopies to work in cooperation with Promolding's robotics. This reduces the risk of human error and improves the quality and consistency of the end product. The cleanroom features C2C's intelligent ECO control system. It is designed to ensure that the cleanroom operates at optimum effectiveness by constantly monitoring the operating conditions, in real time within the critical environment, raising alarms if any of the parameters vary beyond a user-specified threshold.

Each moulding machine has been fitted with one fixed and one actuated, HEPA-filtered, overhead canopy. C2C's automated HEPA-lite™ canopies supply clean air at the critical point of production and reduce contamination by significantly limiting exposure to the external environment during tooling changes, thus providing an ISO 14644-1:2015 Class 7 environment.

The automated, sliding canopies provide overhead access for tool changes. They are driven by an actuator with two linear guides, one master and one slave. These drive the filter system to an open or closed position. The canopies feature effective safety mechanisms, sending infrared signals across the actuators. If the signals are interrupted, for instance by the robot or by operator's hands, the canopy will deactivate movement, preventing any accidents. The canopies feed into a main cleanroom area, which houses assembly, plastic welding and packaging.

Overview of the installation, showing an Engel injection moulding machine fitted with two canopies, a fixed one at the back with two filters, and the sliding canopy at the front with one filter unit.

The intelligent moulding machine recognises faults with a product and drops affected products into stainless steel drop drawers for inspection. The personnel door is interlocked with the moulding machine during manufacture. The sample drawers are accessible from outside the HEPA-lite, meaning the faulty parts or samples can be safely removed without interrupting manufacture.

Despite being a precision build, the main cleanroom was constructed in four weeks and the HEPA-lite canopies were constructed in two weeks.

WWW.CONNECT2CLEANROOMS.COM
WWW.PROMOLDING.NL

Passive damping for high-tech systems training

High Tech Institute and Mechatronics Academy are introducing a new training course focused on passive damping for high-tech systems. The first course will take place next April in Eindhoven, the Netherlands. Over the past few years, efforts have been made to better understand the underlying principles of damping, which has led to spectacular results. Experts from industry and academia share the latest insights and design approaches in the new short course "Passive Damping for High Tech Systems".

When designing high-precision mechatronic systems it is essential to achieve a high bandwidth of the feedback control loop, in order to suppress the negative effects of disruptive forces on the machine accuracy and settling time. Dynamics and resonances play an important role in limiting the achievable bandwidth and settling time. Much focus is directed towards high eigenfrequencies and understanding vibration modes, including the mechanisms of excitation and observability.

However, as accuracy and resulting bandwidth requirements are getting tighter, the requirements in terms of eigenfrequencies of the system sometimes reach the limits of what is physically possible. Passive damping offers additional design space and is becoming a key design parameter for achieving these extreme requirements. Despite the risk of introducing hysteresis-related virtual play, passive damping can significantly simplify controller design and improve positioning performance.

WWW.HIGHTECHINSTITUTE.NL/PASSIVE-DAMPING

"LESS IS MORE, ADDITIVE MANUFACTURING"

On Thursday 11 October, instrument makers met in Leiden, the Netherlands, for a symposium with the theme "Less is more, additive manufacturing". The symposium was organised by the Association for the Promotion of Training as an Instrument Maker, i.e. the alumni association of the Leiden Instrument Makers School (LiS). The symposium targeted instrument makers and attracted approximately 100 participants, including 35 fourth-year LiS students.

The theme had been selected because additive manufacturing is increasingly finding application within the production processes of the instrument-making profession. Various techniques were discussed during the interesting lectures, including the advantages and disadvantages of the techniques within their fields of application. Lectures were delivered by speakers from Ultimaker, Landré, 3D Systems, OIM Orthopedie and Hulotech 3D Printing & Engineering.

Harma Woldhuis of 3D-printer builder Ultimaker kicked it off with a lecture on the applications of plastic 3D printing in the industry. The main example she gave concerned employees in an assembly line printing the moulds they worked with themselves when the moulds wore out.

René Groothedde from Landré, a specialist in CNC manufacturing technologies, gave a presentation entitled "Lightweight metal printing is fantastic....but not easy". He showed that metal printing is already being used in industry for the production of, among other things, aircraft parts. Furthermore, he stated that 75 per cent of production time is spent on finishing the parts.

Raph Alink from 3D Systems, a service provider for 3D metal printing, talked about the future of metal printing. He told his audience that in the early years after the foundation of his company in 2007, 90 per cent of the designs made concerned prototype development. In recent years, the number of prototype components has dropped to 10 per cent of 3D Systems' total production, while the remaining 90 per cent is parts for industry. Furthermore, the first machine converted by 3D Systems for 3D printing is still used as a production machine.

René Groothedde from Landré delivering his lecture at the 2018 LiS Symposium.

Peter de Groot from OIM Orthopedie, a supplier of orthopedic aids, together with Arjan Huiting from Hulotech 3D Printing & Engineering, gave a presentation about 3D printing for orthopedic applications. Much work has already been done with 3D scanning for orthopedics, although up until now, 3D scans were only used to make moulds for fitting prosthetics. The speakers showed that prosthetics of very high quality can be printed by OIM Orthopedie in collaboration with Hulotech.

After the symposium, there was ample opportunity for networking. The exhibitors/sponsors Pfeiffer Vacuum, Hositrad Vacuum Technology, Leybold, Louwers Hanique, Stellar Space Industry, TNO and ZME precision milling shop received great interest in their products and services.

Alumni association

The organisation of the symposium was in the hands of the Association for the Promotion of Training as an Instrument Maker (*Vereeniging ter Bevordering van de Opleiding tot Instrumentmaker*), which was founded 1901, the year when the LiS also started. Since 2016, the Association has again become active in organising events for its members. There is an annual symposium with a different theme each year. The Association also organises company visits, in which the participating instrument makers have the opportunity to become acquainted with the specialisations and competencies of the company concerned. Companies can apply for a corporate membership of the Association.

ALUMNI@LIS.NL
WW.LIS.NL

English edition of compelling ASML historiography

After the successful launch 18 months ago of two books – a technical and a management version – on the history of lithography machine builder ASML, a crowdfunding campaign was started for sponsoring an (American) English translation of the technical book. This has resulted in the publication of *ASML's Architects* by Techwatch Books.

The author of the two Dutch books, René Raaijmakers, publisher at Techwatch and initiator of Bits&Chips, revised the Dutch tech edition to make it more digestible to an international audience, by providing more context. To this end, he added, a.o., a glossary of characters, with brief descriptions of the people named in the book, a 'Notes and Bibliography' section containing literature references and additional explanation, and a glossary of terms.
The books cover the origins of ASML, which was founded in 1984 as a spin-out company from Philips, as well as its formative years up till around 1996. It describes "how a hopeless research and business activity was transformed into a billion-dollar machine and a world-leading company". In the 1980s, the American heavyweights Perkin-Elmer and GCA came under incredible fire from their Japanese competitors Canon and Nikon in the chip lithography market. As a result, the US lost its two-decade monopoly on this key technology, which is the driving force behind Moore's law.

Meanwhile, an obscure, inconsequential lithography company in the Netherlands was taking its first steps. This company, ASML, is now an unparalleled success. It is the world's biggest and most profitable machinery manufacturer for the chip-making industry. With a market share of 70 to 80 percent, ASML has been leaving Canon and Nikon in the lithographic dust for years.

The author returns to the birthplace of the wafer stepper and the roots of ASML's global success. He chronicles the engineers' all-consuming race to surpass the rest, providing a vivid window into the unique culture that spawned the world's finest chip-making technology. The book covers the technology and the business, but the real story is about the people behind them. Throughout its pages, the interviewed engineers, scientists and managers speak frankly about their struggles, their fights and the gruelling teamwork behind the scenes.

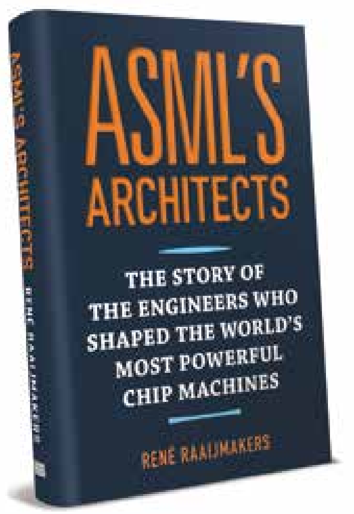

René Raaijmakers, "ASML's Architects", 664 pages, hardcover, € 59.00.

WWW.TECHWATCHBOOKS.NL/ARCHITECTS
WWW.ASML.COM

Challenges when performing laser measurements in vacuum

Canadian supplier of laser measurement systems, Gentec Electro-Optics, represented in the Netherlands by Te Lintelo Systems, offers solutions for the challenges associated with performing laser measurements in experimental vacuum environments.

Ultra-high vacuums can reach pressures as low as 10^{-9} Pa and involve challenges that one has to consider in order to design an experimental set-up. The biggest problem is outgassing. When the external pressure is really low, the most volatile materials will tend to release gas molecules. In a vacuum chamber, for example, silicon molecules expelled from detector cables and glue will condensate onto optics and cause deterioration. Heat dissipation is another concern when using detectors in a high vacuum. Most of them rely on convection cooling for stable operation. However, the vacuum environment prevents proper cooling and can result in the sensor reaching very high temperatures, which can lead to damage.

Gentec-EO offers various solutions. Regarding outgassing, these include replacement of standard detector cables and glues for low-outgassing materials, and aluminium anodising of the detector's case to remove organic materials. For signal transmission, different adaptor feed-throughs can be applied for transmitting the electrical output signal to a monitor outside of the vacuum chamber. To prevent heat dissipation, a water-cooled system can be integrated inside the vacuum chamber, or the heat can be removed by direct contact of the detector with a material outside of the vacuum chamber.

WWW.GENTEC-EO.COM
WWW.TLSBV.NL

Air Bearings

AeroLas GmbH
Grimmerweg 6
D-82008 Unterhaching
Germany

T +49 89 666 089-0
F +49 89 666 089-55
E info@aerolas.de
W www.aerolas.de

AeroLas is world leader in air bearing technology strengthening the customer´s competitive advantage with customized air-guided products and solutions.

Automation Technology

Festo BV
Schieweg 62
2627 AN DELFT
The Netherlands
T +31 (0)15-2518890
E sales@festo.nl
W www.festo.nl
Contact person:
Mr. Ing. Richard Huisman

Festo is a leading world-wide supplier of automation technology and the performance leader in industrial training and education programs.

member DSPE

Cleanrooms

Brecon Group
Droogdokkeneiland 7
5026 SP Tilburg
T +31 (0)76 504 70 80
E brecon@brecon.nl
W www.brecon.nl

Brecon Group can attribute a large proportion of its fame as a cleanroom builder to continuity in the delivery of quality products within the semiconductor industry, with ASML as the most important associate in the past decades.

Brecon is active with cleanrooms in a high number of sectors on:
* Industrial and pharmaceutical
* Healthcare and medical devices

member DSPE

Connect 2 Cleanrooms BV
Newtonlaan 115
Zen Building
3584 BH Utrecht
Nederland
T +31 (0)30 210 60 51
E info@connect2cleanrooms.com
W www.connect2cleanrooms.nl

Our cleanroom solutions are bespoke and scalable, encouraging efficiency through flexible design. We help organisations reduce failure rates by getting it right first time.

member DSPE

Development

TNO
T + 31 (0)88-866 50 00
W www.tno.nl

TNO is an independent innovation organisation that connects people and knowledge in order to create the innovations that sustainably boosts the competitiveness of industry and wellbeing of society.

member DSPE

Development and Engineering

Segula Technologies Nederland B.V.
De Witbogt 2
5652 AG Eindhoven
T +31 (0)40 8517 500
W www.segula.nl

SEGULA Technologies Nederland BV develops advanced intelligent systems for the High Tech and Automotive industry. As a project organisation, we apply our (engineering) knowledge to non-linear systems. This knowledge is comprised of systems architecture and modelling, analysis, mechanics, mechatronics, electronics, software, system integration, calibration and validation.

member DSPE

Education

Leiden school for Instrumentmakers (LiS)
Einsteinweg 61
2333 CC Leiden
The Netherlands
T +31 (0)71-5681168
E info@lis.nl
W www.lis.nl

The LiS is a modern level 4 MBO school with a long history of training Research instrumentmakers. The school establishes projects in cooperation with industry and scientific institutes thus allowing for professional work experience for our students. LiS TOP accepts contract work and organizes courses and summer school programs for those interested in precision engineering.

member DSPE

Electrical Discharge Machining (EDM)

CVT BV
Heiberg 29C
5504 PA Veldhoven
The Netherlands
T +31 (0)497 54 10 40
E info@cvtbv.nl
W www.heinmade.com

Partner high tech industry for wire EDM precision parts. Flexible during day shifts for prototyping. Outside office hours low cost unmanned machining. Call and enjoy our expertise!

member DSPE

Ter Hoek Vonkerosie
Propaanstraat 1
7463 PN Rijssen
T +31 (0)548 540807
F +31 (0)548 540939
E info@terhoek.com
W www.terhoek.com

INNOVATION OF TOMORROW, INSPIRATION FOR TODAY
Staying ahead by always going the extra mile. Based on that philosophy, Ter Hoek produces precision components for the high-tech manufacturing industry.

We support customers in developing high-quality, custom solutions that can then be series-produced with unparalleled accuracy. That is what makes us one of a kind.
It is in that combination of innovative customization and repeated precision that we find our passion. Inspired by tomorrow's innovation, each and every day.

member DSPE

Lasers, Light and Nanomotion

Laser 2000 Benelux C.V.
Voorbancken 13a
3645 GV Vinkeveen
Postbus 20, 3645 ZJ Vinkeveen
T +31(0)297 266 191
F +31(0)297 266 134
E info@laser2000.nl
W www.laser2000.nl

Laser 2000 Benelux considers it her mission to offer customers the latest photonics technologies available.
Our areas of expertise are:

- Lasers, scanners and laser machines for industry and research
- Light metrology instruments for LED and luminaire industries
- Light sources for scientific applications
- Piezo- and stepper motion products for nano- and micro positioning
- Inspection and research grade high speed cameras
- Laser safety certified products

Te Lintelo Systems B.V.
Mercurion 28A
6903 PZ Zevenaar
T +31 (0)316 340804
E contact@tlsbv.nl
W www.tlsbv.nl

Photonics is our passion!
Our experienced team is fully equipped to assist you with finding your best optical business solution.
For over 35 years TLS represent prominent suppliers in the photonics industry with well-educated engineers, experience and knowledge.
Over the years we became the specialist in the field of:

- Lasers
- Light metrology,
- Opto-electronic equipment,
- Positioning equipment
- Laser beam characterization and positioning,
- Interferometry,
- (Special) Optical components,
- Fiber optics,
- Laser safety

Together with our high end suppliers we have the answer for you!

Mechatronics Development

Sioux CCM
De Pinckart 24
5674 CC Nuenen
T +31 (0)40 2635000
F info.ccm@sioux.eu
W www.sioux.eu

Sioux CCM is a technology partner with a strong focus on mechatronics.
We help leading companies with the high-tech development, industrialization and creation of their products, from concept stage to a prototype and/or delivery of series production.
Commitment, motivation, education and skills of our employees are the solid basis for our business approach

Sioux CCM is part of the Sioux Group.

member DSPE

Manufacturing Technical Assemblies (MTA) b.v.
Waterbeemd 8
5705 DN Helmond
T +31 (0)492 474992
E info@m-t-a.nl
W www.m-t-a.nl

MTA is an high-tech system supplier specialized in the development and manufacturing of mechatronic machines and systems.
Our clients are OEM s in the Packaging, Food, Graphics and High-tech industries.

member DSPE

Mechatronics Development

MI-Partners
Dillenburgstraat 9N
5652 AM Eindhoven
The Netherlands
T +31 (0)40 291 49 20
F +31 (0)40 291 49 21
E info@mi-partners.nl
W www.mi-partners.nl

MI-Partners is active in R&D of high-end mechatronic products and systems. We are specialised in concept generation and validation for ultra-fast (>10g), extremely accurate (sub-nanometers) or complex positioning systems and breakthrough production equipment.

member DSPE

Metal Precision Parts

Etchform BV
Arendstraat 51
1223 RE Hilversum
T +31 (0)35 685 51 94
F info@etchform.com
W www.etchform.com

Etchform is a production and service company for etched and electroformed metal precision parts.

member DSPE

Micro Drive Systems

Maxon Motor Benelux

The Netherlands

Head Office
maxon motor benelux bv
Josink Kolkweg 38
7545 PR Enschede

South
High Tech Campus 9
5656 AE Eindhoven
T +31(053) 744 0 744
E info@maxonmotor.nl
W www.maxonmotor.nl

Belgium / Luxembourg

maxon motor benelux bv
Schaliënhoevedreef 20C
2800 Mechelen - Belgium
T +32 (15) 20 00 10
F +32 (15) 27 47 71
E info@maxonmotor.be
W www.maxonmotor.be

maxon motor is the worldwide leading supplier of high precision drives and systems. When it really matters! Try us.

member DSPE

Micro Drive Systems

FAULHABER Benelux B.V.
Drive Systems
High Tech Campus 9
5656 AE Eindhoven
The Netherlands
T +31 (0)40 85155-40
E info@faulhaber.be
E info@faulhaber.nl
W www.faulhaber.com

FAULHABER specializes in the development, production and deployment of high-precision small and miniaturized drive systems, servo components and drive electronics with output power of up to 200 watts. The product range includes brushless motors, DC micromotors, encoders and motion controllers.
FAULHABER also provides customer-specific complete solutions for medical technology, automatic placement machines, precision optics, telecommunications, aerospace and robotics, among other things.

Physik Instrumente (PI)
Benelux BV
Hertog Hendrikstraat 7a
5492 BA Sint-Oedenrode
The Netherlands
T +31 (0)499-375375
F +31 (0)499 375373
E benelux@pi.ws
W www.pi.ws

PI is the world's leading provider of nanopositioning products and systems. All key technologies are developed, manufactured and qualified in-house by PI: Piezo components, actuators and motors, magnetic drives, guiding systems, nanometrology sensors, electronic amplifiers, digital controllers and software.

member DSPE

Motion Control Systems

Aerotech United Kingdom
The Old Brick Kiln
Ramsdell, Tadley
Hampshire RG26 5PR
UK
T +44 (0)1256 855055
F +44 (0)1256 855649
W www.aerotech.co.uk

Newport Spectra-Physics B.V.
Vechtensteinlaan 12 - 16
3555 XS Utrecht
T +31 (0)30 6592111
E netherlands@newport.com
W www.newport.com

Newport Spectra-Physics B.V. is a subsidiary of Newport, a leader in nano and micro positioning technologies with an extensive catalog of positioning and motion control products. Newport is part of MKS Instruments Inc., a global provider of instruments, subsystems and process control solutions that measure, control, power, monitor, and analyze critical parameters of advanced processes in manufacturing and research applications.

member DSPE

Motion Control Systems

Physik Instrumente (PI) Benelux BV
Hertog Hendrikstraat 7a
5492 BA Sint-Oedenrode
The Netherlands
T +31 (0)499-375375
F +31 (0)499 375373
E benelux@pi.ws
W www.pi.ws

PI is the world's leading provider of nanopositioning products and systems. All key technologies are developed, manufactured and qualified in-house by PI: Piezo components, actuators and motors, magnetic drives, guiding systems, nanometrology sensors, electronic amplifiers, digital controllers and software.

member DSPE

Optical Components

Molenaar Optics
Gerolaan 63A
3707 SH Zeist
T +31 (0)30 6951038
E info@molenaar-optics.nl
W www.molenaar-optics.eu

Molenaar Optics is offering optical engineering solutions and advanced products from world leading companies OptoSigma, Sill Optics and Pyser Optics.

member DSPE

Piezo Systems

HEINMADE BV
Heiberg 29C
NL - 5504 PA Veldhoven
T +31 (0)40 851 2180
E info@heinmade.com
W www.heinmade.com

As partner for piezo system solutions, HEINMADE serves market leaders in the high tech industry. Modules and systems are developed, produced and qualified in-house. HEINMADE distributes Noliac piezo components.

member DSPE

Piezo Systems

Physik Instrumente (PI) Benelux BV
Hertog Hendrikstraat 7a
5492 BA Sint-Oedenrode
The Netherlands
T +31 (0)499-375375
F +31 (0)499 375373
E benelux@pi.ws
W www.pi.ws

PI is the world's leading provider of nanopositioning products and systems. All key technologies are developed, manufactured and qualified in-house by PI: Piezo components, actuators and motors, magnetic drives, guiding systems, nanometrology sensors, electronic amplifiers, digital controllers and software.

member DSPE

Ultra-Precision Metrology & Engineering

IBS Precision Engineering
Esp 201
5633 AD Eindhoven
T +31 (0)40 2901270
F +31 (0)40 2901279
E info@ibspe.com
W www.ibspe.com

IBS Precision Engineering delivers world class measurement, positioning and motion systems where ultra-high precision is required. As a strategic engineering partner to the world's best manufacturing equipment and scientific instrument suppliers, IBS has a distinguished track record of proven and robust precision solutions. Leading edge metrology is at the core of all that IBS does. From complex carbon-fibre jet engine components to semiconductor chips accurate to tens of atoms; IBS has provided and engineered key enabling technologies.

member DSPE

ADVERTISERS INDEX