I/O magazine

ICT RESEARCH PLATFORM NEDERLAND

PAUL KLINT

Does Dutch politics care about ICT?

On 22 November 2022, elections took place in the Netherlands. Does Dutch politics care about ICT?

The programmes of Dutch political parties mention many relevant topics like digital rights, digital sovereignty, government ICT and the development of ICT expertise. Several of my favourite topics are discussed by some party or another, such as very early ICT feasibility checks when developing new policies and laws and attention for software maintenance (make maintenance costs explicit both in the National Budget and in the budget of each individual ICT project). What I miss is attention for *common sense*: this sounds trivial but many (governmental) projects fail due to a lack of it. Departmental policies, political and technical fashions, group thinking and personal career ambitions all hinder common sense. Common sense has to be organised to make it effective. Unfortunately, both whistleblowers and common sense lead miserable lives in our country. We do have an official House for Whistleblowers, but it has a very limited impact.

One organisation that could introduce more common sense in governmental ICT projects is the Ministry of Digital Affairs, as some of the parties propose. What is needed to let such a ministry succeed?

First and foremost, a Ministry of Digital Affairs should have a voice in all ICT-related decisions of all other ministries, comparable to the Ministry of Finance's role in other ministries' financial decisions. This could reduce the current fragmentation, duplication and sometimes naivety in governmental projects. Second, the position of the CIO-Rijk should be reconsidered. At the moment, this CIO is responsible for the development of government-wide ICT policies but lacks the means to enforce them. In practice, each ministry has its own responsible CIO and even they lack the means to decisively steer projects towards success. Third, the ICT know-how in all ministries should be strengthened and the role of the market and external contractors reduced. There are more but column-length overflow is hitting here.

I observe a lot of good intentions in party programmes but as we all know, the road to hell is paved with them. In the end, it is the execution of those intentions that decides the success or failure of Dutch ICT policies in general and of a Ministry of Digital Affairs in particular. To summarise: the care is sufficient, but execution remains to be seen.

138/86

9
8/98.24

MARRYING THE HARDNESS OF MACHINES WITH THE SOFTNESS OF PEOPLE

The NWO Gravitation programme Hybrid Intelligence has been running for nearly four years now, with six more years to go. Hayley Hung and Bart Verheij evaluate what has been accomplished, the implications of ChatGPT and how they hope the programme will affect the development of AI.

By Bennie Mols Image iStock

Imagine a robotic arm equipped with a smart camera that can show images of places inaccessible to the surgeon operating with the robot, an arm that even thinks along with the surgeon about the best way to perform the operation. Such a smooth and smart collaboration between humans and machines could greatly improve surgery. It is still music of the future, but this is one of the research projects that is part of the ten-year NWO Gravitation programme Hybrid Intelligence that started in 2020.

Hayley Hung:

'The consortium is a melting pot of different ideas and opinions'

In the areas of pattern recognition and machine learning, AI today regularly surpasses our human capabilities. However, today's AI struggles with understanding general knowledge of the world and reasoning with common sense, aspects that we humans develop playfully from childhood and take for granted. Moreover, AI systems are typically poor at human capabilities such as collaboration, adaptability, responsibility and the ability to explain our thought processes, especially our values and norms.

Combining human and machine intelligence is therefore the natural way to develop robust, trustworthy and ethical AI applications. In 2020, the consortium published a research agenda which describes hybrid intelligence as 'the combination of human and machine intelligence, augmenting human intellect and capabilities instead of replacing them and achieving goals that were unreachable by either humans or machines.'

Bart Verheij:

'My dream is that we can contribute to democratising AI technology'

The research of the Hybrid Intelligence consortium focuses on exactly those four capabilities in which machines are poor: Collaboration, Adaptability, Responsibility, Explainability, abbreviated as CARE. Funded by an NWO Gravitation grant, the programme has also materialised in the Hybrid Intelligence Centre, where researchers can meet, collaborate and use the shared

laboratories and infrastructure. Its headquarters are at VU Amsterdam and it has working locations at all participating universities in Amsterdam (VU and UvA), Delft, Groningen, Leiden and Utrecht.

Melting pot

Hayley Hung, associate professor at TU Delft, and Bart Verheij, professor in AI and Argumentation at the University of Groningen, look back on the first four years of the Hybrid Intelligence programme. 'The most important result so far, as far as I'm concerned, is creating a community,' says Hung, who is part of the research line on collaborative AI. 'The consortium has allowed many more different people to work together and create a melting pot of different ideas and opinions. We already see researchers publishing together in papers where they didn't publish before.'

Bart Verheij, part of the research line on responsible AI, adds: 'In the consortium, we work in a very bottom-up, transparent and collaborative way. We have a Google Drive with information about all meetings and presentations, and everybody has access to everything. This allows us to pick the brains of everybody at all levels. It's all very different from my early days as a researcher when there was much more hierarchy and ego boosting.'

First results

The first batch of PhD students has nearly completed their research, and recruitment of the second batch has begun earlier this year. This also means that the first concrete research results will soon be published. Hung and Verheij can already share some preliminary highlights. Hung: 'Together with postdoc Bernd Dudzik and PhD student Tiffany Matej-Hrkalovic, I worked with social psychologist Daniel Balliet on the research question of how to pick people to collaborate with. That is like a first date, with a messiness of interacting social signals. Are there computationally extractable signals that can predict or inform how people decide who to work with? We have done experiments, we have collected data and now we are analysing them. It would be nice if we could develop a technical tool to help us better choose the people to collaborate with.'

In recent years, the research line of responsible AI has brought together the technical-oriented people with the more philosophical-oriented ones, says Verheij. 'One result has been that the technical researchers have become more aware of their ethical responsibilities, and the philosophical researchers have become more aware of the technical possibilities and impossibilities.'

One concrete research question that Verheij has explored with PhD student Cor Steging is how to make the outcomes of machine-learning models more transparent and explainable. Verheij: 'We have looked at a rule of

Dutch law that defines exactly what and when something qualifies as a so-called tort, a legal topic about compensation for damages. With a knowledge representation model of that law, we have generated data that we have fed into a machine-learning model. We asked the question whether a machine-learning model could rediscover the underlying knowledge structure from the data.'

It turned out that the machine-learning model performed very well statistically, in the sense of mostly giving the right answers, but not for the right reasons. Verheij: 'It failed to discover the underlying knowledge structure. This, of course, is very strange and undesirable. The next challenge is to find out why that is and what we can do about it. For now, machines cannot do the job by themselves, and having humans in the loop is necessary.'

Language models

With the launch of ChatGPT late last year, AI has become more in the spotlight than ever before. To what extent do large language models like ChatGPT impact research within the Hybrid Intelligence programme?

'Well, they do a bit and they don't', says Hung. 'What hybrid intelligence stands for goes beyond what current large language models are doing. Hybrid intelligence problems are mostly problems which can only be solved using relatively small amounts of data. The subtleties of adapting to specific social situations, understanding the context, and knowing what to say when lie outside the realm of large language models.'

For Verheij, large language models are a game changer. 'For the first time, I see that data-driven methods do something for knowledge-based AI', he says. 'There is a bit of logical reasoning inside those models, but nobody understands their capabilities and limitations. That is scary and fascinating at the same time. Large language models have no idea when they are correct and when they are wrong. So far, they are a product of engineering experiments, but we need much more scientific understanding.'

Verheij: 'What I find inspiring for our hybrid intelligence programme is that ChatGPT is a form of hybrid intelligence. The user has a conversation with the machine. I hear students experimenting with ChatGPT in a way that I think is a smart and nice way to use this tool, namely as a kind of idea suggestion machine. They still have to think. ChatGPT does not do that for them.'

The Hybrid Intelligence programme runs for another six years, so the best is yet to come. Hung and Verheij also see that all over the world, similar initiatives are happening. How do they hope the programme will have a lasting impact on AI development? 'I hope for a marriage between the hardness of machines with the softness of people', says Hung. 'Technology for technology's sake doesn't work. I hope that we can create a culture that is more open to hybrid intelligence. For example, right now, I am con-

cerned about where the data for all these large language models come from. Not enough attention has been paid to the ethical side of data collection. But that's really an important part of the whole ecology of hybrid intelligence.'

Democratising AI

'My dream', says Verheij, 'is that we can contribute to democratising AI technology so that professional experts and everyday users can participate in changing AI tools in the right direction. The next phase of AI will be the hybrid intelligence phase, where the machines support us in using their strengths, like heavy data crunching, and we use our own strengths, like knowledge of the world and appreciating human values. I was a bit surprised by how angry some people got at ChatGPT. I'm afraid of the abuse of ChatGPT and the power of big tech, but I'm not afraid of ChatGPT as a tool. So, let's not act out of fear but enjoy the opportunities of being researchers in a fantastic field.'

HYBRID INTELLIGENCE CENTRE

Budget: € 20 million (funded by a 10-year NWO Gravitation grant).
Universities involved: VU Amsterdam (leader), University of Amsterdam, TU Delft, Leiden University, Utrecht University, University of Groningen
Goal: Designing hybrid intelligent systems that combine the strengths of human and machine intelligence. Hybrid intelligence takes human expertise intentionality into account when making meaningful decisions and performing appropriate actions that take ethical, legal and societal values into consideration.
Research lines: Four aspects of hybrid intelligence (HI): Collaborative HI, Adaptive HI, Responsible HI and Explainable HI.
Practical goals: In healthcare: a smart robot that can work with a surgeon via dialogue during surgery. In education: teams of robots and teachers to teach better together. In science: a published scientific paper co-authored by a computer.

www.hybrid-intelligence-centre.nl

Collaborating across the pond

By Sonja Knols
Image iStock

Since 2012, NWO, the Dutch National Cyber Security Centre and the US Department of Homeland Security (DHS) have been conducting a joint research programme into cyber security. So far, this cooperation and a similar cooperation with the National Science Foundation (NSF) have resulted in 13 joint projects. At the One Conference last October, the merits of joint international research programmes like these were discussed and some project results were presented.

Carlos Hernandez Ganan from Delft University of Technology can be brief when asked about the advantages of working together with American colleagues in his recently finished project MitigatINg IOt-based DDoS attacks via DNS (MINIONS): 'Our collaboration with New York University brought a unique edge, enabling us to leverage mixed methods and multidisciplinary research to gain global insights into the role of the human factor in Internet-of-Things cyber security.'

'The American and Dutch philosophies on openness and sharing data align very well', says Jan Piet Barthel, who was one of the founding fathers of the collaboration and attended the One Conference session. 'Sharing a mindset is an important prerequisite for the success of such a collaboration.' However, funding is also needed to provide a stimulus for individual research groups across the pond to actively find each other and build new relationships. But there is a large return on investment, he states. 'For this spe-

cific programme, NWO and the Dutch National Cyber Security Center funded the Dutch scientists and the American organisation DHS funded the American researchers. However, since these funding agencies started to collaborate, we literally got two for the price of one: for every PhD student we funded, we got an additional American-funded PhD student for free. So each euro we invested easily delivered twice the revenues when compared to national funding schemes.'

Relevant results

The collaborations have not only been fruitful when it comes to establishing new relationships but also in terms of scientific and societally relevant results. This came to the fore during the short pitches at the conference about three of the recently completed joint projects. For example, the project Planning Anycast for Anti-DDoS (PAADDOS) resulted in tools to predict the impact of a DDoS attack and to help operators of critical

services design their systems in such a way that they are as resilient as possible to such attacks. And the Mapping DNS DDoS Vulnerabilities to Improve Protection and Prevention (MADDVIPR) project resulted in new actionable intelligence on Internet-of-Things-based DDoS attacks that target the Domain Name System, which translates a domain name into an IP address. This intelligence will aid the protection of the DNS and facilitate the prevention of attacks against it.

Cyber security has become a global sport, with the internet spanning the entire world and ever more devices connecting to it. Based on his vast experience with the topic, Barthel has a clear vision of the way forward. 'You simply cannot put a fence around the internet to keep out malicious parties. So, in the field of cyber security, collaborating across national borders is a no-brainer. And why not start by connecting the scientists who work on solutions?'

BÈTA IN GOVERNANCE AND POLICY

On 22 November, the second edition of the 'Bèta in Governance and Policy' career day was organised for PhD, postdoc and master students with a STEM background connected to Dutch universities and research institutes.

Young STEM students interested in politics, governance and policy were welcome to learn from people who have chosen to work in these fields. What was their motivation, what are important competencies you need for that kind of job and what are examples of career opportunities in the sector?

The initiative 'Bèta in Governance & Policy' is developing various types of actions to bring the policy and science worlds closer together. Their long-term aim is to increase the number of STEM-trained politicians and civil servants. Their short-term aim is to convince politicians and public administrators of the profitable effects obtained from associating with young STEM scientists.

KIC CALL CYBER SECURITY

NWO has opened the call 'Cyber security for digital resilience'. With this public-private partnership call, NWO aims to substantially boost cyber security knowledge and innovation, focusing on solutions to digital security problems. New knowledge and innovation are needed to shape societal transitions in a secure digital way and strengthen the digital resilience of the Netherlands.
A maximum of 3.4 million in research funding can be applied for per project. The research questions focus on an integrated approach and include 7 research themes: Security by design; Secure data-driven operations; Secure and resilient connectivity; Operational technology (OT) and IoT security and secure integration with IT; Cyber risk management; Security of systems and supply chains; Cyber knowledge and skills.

More information via www.nwo.nl/kic

NWO OTP GRANT FOR ANDREA CONTINELLA

Andrea Continella, associate professor at the University of Twente, has been awarded one of seven NWO OTP grants. The Open Technology Programme provides funding for excellent research with a view to the potential application of the results. The programme gives companies and other organisations a low-threshold way to join scientific research that should lead to applicable knowledge.
Continella's project is about designing and developing automated techniques to analyse discovered vulnerabilities in software, assess their risk, prioritise the critical ones and generate patches to fix them. The research will consider vulnerabilities in their context, including interactions between vulnerabilities and defences, allowing for prompt mitigation and reduced costs.

NWO ICT.OPEN 2024

NWO ICT.OPEN 2024 will take place on 10 and 11 April 2024 at the Beatrix Theater Jaarbeurs Utrecht. At this annual conference, scientists from all ICT research disciplines and industry are brought together to learn, share ideas and network. NWO ICT.OPEN is organised annually by the Dutch Research Council (NWO), ICT Research Platform for the Netherlands (IPN) and platform Praktijkgericht ICT-onderzoek (PRIO). The theme of next year's edition is 'Science in the Service of Society'.

BITS AND BYTES FROM SPACE

By Leendert van der Ent Image ESA

The first association many people will have when it comes to space technology is rockets. In reality, data processing would be more appropriate. At least as far as senior scientist at SRON in Leiden Jochen Landgraf's work is concerned. He explains how data from space instruments play a vital role in battling climate change.

'The first mission I was involved in was the Global Ozone Monitoring Experiment (GOME) instrument some 25 years ago', Landgraf recalls. 'Data processing took place on my own desktop. It was a matter of MBs.' Fast forward to the TROPOspheric Monitoring Instrument (TROPOMI) that measures air quality and greenhouse gas emissions since 2017. 'To take one example, it carries out 215 methane observations per second, partly to establish whether countries live up to their climate commitments. Fifteen data transmissions per day yield in total about 1 TB of raw data.'

With these amounts of data, data processing is a serious matter. 'Our software runs on the high-performance computing facilities at SURF where the satellite data comes in', says Landgraf. 'We are very happy with SURF people and facilities, as they are fully dedicated to science. This creates a common language between the hardware specialists there and us as scientific software developers. This merger of computer knowledge and domain knowledge is always a critical success factor in ICT.'

MATCH MEASUREMENTS

The methane emissions are measured with calibrated spectra for the reflection of sunlight. Landgraf: 'My core task in TROPOMI has been to develop a software model that describes the measurements in such a way that a number of photons measured eventually translates into reliable knowledge on methane emissions.'

This requires several data processing steps. Level 0 to 1 involves the calibrating step on raw bits and bytes. Levels 1 to 2 are about unravelling the absorption structures of the trace gases in the spectra of sunlight that is reflected by the Earth, which leads to the amount of trace gases in the atmosphere. At the highest level, local emissions from landfills, wetlands, and blowouts at oil and gas production sites are estimated.

Landgraf: 'A vital aspect of this model development process is a verification loop between data input and verified methane emission measurements. This all revolves around a computer model that simulates the measurement by assuming a certain amount of greenhouse gases in the Earth's atmosphere. Subsequently, those assumptions must be adapted to a point at which the model exactly matches the measurements. As this is a process with numerous iterations then it all boils down to doing this as smartly as possible to make effective use of computing power. Moore's law is impressive, but it's just not fast enough.'

NEW INSTRUMENTS

These are exciting times at SRON, as the launch of the SPEXone instrument within the framework of NASA's PACE mission will take place on 24 January 2024. Landgraf: 'This instrument will measure aerosols in more or less the same way as TROPOMI measures methane. Aerosol levels are closely linked to climate change and air quality. Less aerosols

means better air quality. Ironically, aerosols have an important cooling climate effect, and currently partly mitigate global warming.'

With SPEXone not even launched, Landgraf is already looking forward to the next launch, namely that of TANGO in 2027 or 2028. 'We are now making software design choices for the instrument that will be built in 2024. It will measure CO_2, CH_4, and NO_2 at a resolution of 300 by 300 meters. This high resolution is very important, as it takes us from a several square kilometres level that allows us to say, for example, 'somewhere in the Europoort area', to a specific location on a specific site. European weather forecasting data will also be integrated into the software data flow, which will further enlarge the process chain compared to previous missions.'

'MOORE'S LAW IS IMPRESSIVE, BUT IT'S JUST NOT FAST ENOUGH'

Jochen Landgraf

DUTCH NICHE

The TANGO satellite will be built by ISIS Space (Delft) and will carry an eight-litre large spectrometer to be made by TNO. SRON will be responsible for the detector characterisation and the data chain will be a joint SRON/KNMI effort. Landgraf: 'This makes the supply chain entirely Dutch. I believe that projects like these enable the Dutch space knowledge infrastructure to carve out an interesting niche for itself and SURF plays an enabling role in this. Furthermore, I think it is worthwhile setting up a community structure to specialise further in the key question: how can we extract the maximum amount of knowledge from data coming from our instruments?'

Language technology for the social good

By Bennie Mols Images Ivar Pel

GROUP PASSPORT

RESEARCH FIELD

Natural language processing, machine translation, summarisation, question answering, language modelling, image captioning

INSTITUTION

Informatics Institute of the University of Amsterdam (UvA)

EMPLOYEES (as of November 2023)

1 professor, 1 assistant professor,
13 PhD students

WEBSITE

ltl.science.uva.nl

The University of Amsterdam's Language Technology Lab focuses on text generation for machine translation, summarisation and question answering with a keen eye on user control and multilingualism.

Christof Monz

Kata Naszádi

Vlad Niculae

Since the launch of ChatGPT in late 2022 and its subsequent spectacularly rapid adaptation, language technology has attracted more attention than ever before. 'Nowadays, one vacancy attracts more than a hundred applicants,' says professor Christof Monz, leader of the Language Technology Lab at the University of Amsterdam's Informatics Institute. Despite the seemingly impressive performance of ChatGPT, both language processing and generation are far from being solved problems. As everyone has been able to experience by now, ChatGPT writes falsehoods with great conviction, writes impersonal, cliched and sometimes even harmful text, offers little control over the text and works only for major languages. Monz's lab is working to improve some of these weaknesses.

'In general, our lab focuses on the generation of text in the form of machine translation, summarisation, question answering, and control of the generated text', says Monz. 'As for the latter, we want to give users more control over aspects such as quality, formality and toxicity. Does a sentence flow well? Is the content appropriate? Should you translate the English 'you' into Dutch as 'jij' or as 'u'? That depends on the context. How do you translate slang? How do you avoid discriminatory and other harmful language?'

For machine translation, Monz's lab is focusing on smaller languages, for which the well-known translation engines such as Google Translate or DeepL do not work well, if at all. Monz: 'Automatic translation from Bengali to Swahili, for example, is currently dramatically poor. Because we value inclusiveness, providing language technology for smaller languages is important too. Therefore, we are developing techniques that are able to translate languages for which little or no data exist.' This is also the focus of Monz's ongoing NWO Vici project.

One of the interesting applications of the lab's work is the translation of documents from the City of Amsterdam for various minorities in the city, work done with a larger consortium called Language Sciences for Social Good. Monz: 'Translating such documents happened a lot during the COVID-19 pandemic. With official city documents, quality and accuracy are, of course, extra important. Some citizens also need the text in official documents to be presented in a simplified manner. That's something we want to work on in the coming years.'

Assistant professor Vlad Niculae joined the Language Technology Lab when it was established in 2020. 'I was really excited about the opportunity to shape a new group's direction,' says Niculae. Whereas Monz has a background in lin-

guistics, Niculae comes from computer science. Niculae: 'Christof told me that he was looking for somebody as different to him as possible but with the same values and the same drive. We both aim for a deep understanding of language technology problems but take different approaches. I am looking more for generalisations and finding mathematical answers is what gets me excited.'

In October 2022, Niculae started working on the NWO Veni project 'Intelligent interactive natural language systems you can trust and control'. Niculae: 'In this project, I propose a redesign of the dominant paradigm that currently underlies language generation systems like ChatGPT. I argue that in that paradigm, you cannot build in controllability and that we need a new paradigm that includes controllability by design. To give an example: one of my students is working on the generation of subtitles. That is not just about automatically recognising audio but also about the timing of the subtitle, and about the maximum number of words before the subtitle becomes unreadable. These are some of the parameters that you want to control. Every application domain has its own specific control parameters.'

Improving dialogue

Kata Naszádi worked for four years at Amazon on the automatic speech recognition system for its Alexa personal assistant before starting her PhD in 2020 at the Language Technology Lab. In recent years, the number of PhD students in the lab has grown to thirteen. 'What is special in this group is that we are a foodie team', says Naszádi. 'There is this stereotype that PhD students go for the cheapest food options, but we actually like to go to really good restaurants together. An Iranian PhD student took us to an Iranian restaurant, a Chinese took us to a Chinese restaurant. As a Hungarian, I cooked a Hungarian meal for the group.'

Naszádi's PhD research is part of the Gravitation programme Hybrid Intelligence. 'I am trying to improve the dialogue between a human and an artificial agent in which they have to achieve some goal together', she says. 'I use a virtual environment based on the game Minecraft, in which an agent needs to build things and the human gives directions on what to build. They use natural language in order to coordinate their actions and understand each other better.'

She also collaborates with researchers from TU Delft and Erasmus MC on developing a dialogue system that allows a microsurgeon to communicate with a tiny camera he uses during surgery on blood vessels, for example. Naszádi: 'We want to make the conversation during the surgery flow naturally, so that the surgeon can tell the camera things like go a bit more to the left or a bit more down. That would give the surgeon a better vision and thus improve the quality of the surgery.'

Marjolijn Antheunis is Professor of Communication and Technology at Tilburg University and vice dean for research at Tilburg School of Humanities and Digital Sciences. She obtained her PhD in Communication and Technology at the University of Amsterdam in 2009, after having worked as a consultant in the private sector for several years. Her research is focussed on the impact of technology on interpersonal communication, and especially the effects of online communication on our social lives.

ON THE CROSSROADS
OF HUMANITIES AND SCIENCE,
MAN AND MACHINE

Marjolijn Antheunis is Professor of Communication and Technology at Tilburg University and vice dean for research at Tilburg School of Humanities and Digital Sciences. She studied Communication Science at the University of Amsterdam and obtained her PhD on 'Online communication, interpersonal attraction, and friendship formation' in 2009. She is still intrigued by the impact of online communication on the human condition.

By Leendert van der Ent
Image Bram Saeys

What do you find the most surprising development in online communication?

"That it keeps on growing to the point where it takes over some people's lives. I started my research in the age of e-mail, Hyves, and MSN. Back then, online digital communication seemed to peak. Then the first smartphone came along, and the hype truly reached its zenith – it seemed. And here we are in the age of 24/7 social network connections and AI conversation agents for so-called interpersonal communication.'

You have carried out a longitudinal study into the impact of AI conversation chatbots such as Replica on friendship formation between humans and chatbots. What was the conclusion?

'Among the student population included, we found that without a prior communication need, test persons gradually lost interest in the chatbot as the quality of social interactions failed to match that with real people. Yes, the bots have nice qualities, such as uttering statements of liking. But to have intimate conversations, humans expect reciprocity when they open up. Chatbots are not good at self-disclosure yet – there is no personality to disclose. But beyond research results, there is anecdotal evidence that people with an actual prior need for communication, people who feel lonely, for example, are ready to commit deeply to a chatbot, to the point of regarding it as their partner.'

What are your thoughts about the use of AI chatbots in mental healthcare?

'Mental health care chatbots such as Woebot and Wysa are said to be beneficial for treating anxiety and depression, but there is hardly any scientific evidence for that. I think it is the responsibility of science to society to establish whether there is an impact and if this impact is positive or negative. If necessary, we could urge politics to come up with a legal framework, just as the EU has done with the AI Act.'

What is it like to do research on the interface between humanities and sciences?

'An interdisciplinary approach in science leads to beautiful things. I think it's great that our research school merges computer science with the humanities. While it's true that some computer scientists initially regarded us as the odd men out, many of them now start to see the added value. That is why we submitted a joint Gravitation proposal. I realise what I don't know and they know, and they realise what I know and they don't. When we complement each other and break down the walls between the human and natural sciences then that really benefits our research.'

HOW CAN AI BETTER HANDLE UNCERTAINTIES IN THE REAL WORLD?

By Bennie Mols Image iStock

Nils Jansen from Radboud University has received an ERC grant to develop methods for making AI applications more robust in dealing with unforeseen situations.

Imagine a drone that needs to deliver a package to a building in a densely populated city. How can the drone get that done safely? There may be other drones flying around, there are buildings in the way, there may be a gust of wind, people walking around, or traffic driving. Many uncertainties stand in the way of the flying drone safely delivering the package.

Nils Jansen, part-time associate professor at Radboud University in Nijmegen and recently appointed as full professor at Ruhr-University Bochum, is investigating how AI can deal better with such uncertainties in the real world. He received an ERC grant of 1.5 million euros for the five-year project 'Data-driven verification and learning under uncertainty' (DEUCE). In January of this year, he started to work on the project, at the moment together with three PhD students, but he is looking for more staff to work on the project.

No board game

One way that has been successful in recent years in enabling AI systems to get better and better at a given task is learning through reward and punishment, so-called reinforcement learning. This technique was the basis of go computer AlphaGo, which beat one of the best go players of all time in 2016. 'But the real world is not a board game', says Jansen. 'Reinforcement learning works well in games and simulations with a restricted number of scenarios. However, it's not much used with robots in the real world because you can't make robots collide with buildings and people until they've learned that that is not a good idea. The aim of the DEUCE project is to bring reinforcement learning closer to the real world by being able to deal with uncertainty. We want to give guarantees on the behaviour of AI agents that are controlled by reinforcement learning.'

The approach used by Jansen and his colleagues belongs to the relatively new field of neurosymbolic AI, a combination of machine learning based on data and machine reasoning based on logical rules and even formal verification. Jansen: 'The idea is that we model some prior knowledge with a symbolic approach. In the case of the delivery drone, that can be information about the weather, for example. With this approach, we can compute the probabilities of things that can happen to the drone. Based on these probabilities, we can then compute a plan for the drone that, while using reinforcement learning, still guarantees with, let's say 99 per cent certainty, a certain desired behaviour. Some people call this approach 'guess and verify'.'

In the first year of the DEUCE project, the researchers have already acquired a lot of insights about how to mathematically model various kinds of uncertainty. 'Currently, we are trying to move to slightly more realistic examples', says Jansen. 'One example is the delivery drone. Other examples are a robotic arm or small mobile robots that need to learn to find a path. For these cases, we collaborate with roboticists at the University of Texas in Austin, the University of Oxford and RWTH Aachen University. However, my work is primarily fundamental and not applied. Ultimately, our methods and techniques are meant to push AI to become more robust in the real world with many unforeseen situations.'

More information **nilsjansen.org**

MANAGING COMPLEXITY OF CYBER-PHYSICAL SYSTEMS

By Marysa van den Berg
Images WAT ontwerpers

Cyber-physical systems are all around us and their complexity is rapidly increasing. The NWO programme Mastering Complexity (MasCot) aims to develop software tools and methodologies to deal with these issues so as to give the Dutch industry a competitive advantage.

A self-driving car, a large-production printer or the wafer scanner that uses a photolithographic process to build microchips. These are all examples of cyber-physical systems: systems that interact with and manipulate the physical world. Obviously, such systems are essential for society, but over the last few years, they have become so complex that this complexity is giving rise to several challenges. 'Many of these machines are composed of a lot of smaller subsystems, each with their own set of features', explains Harm Munk, system architect and software developer at TNO and coordinator of the MasCot programme. 'And some of the systems are specifically designed for a particular customer. We aim to make the ever-increasing complexity more manageable by designing, building and testing new tools that can help us keep up with society's high demands both now and in the future.'

'We must ensure new software will run on old hardware, and new hardware will support old software'

FROM DESIGN SPACE TO TESTING

The MasCot programme started in 2019 and consists of four projects. Each of them is tackling a different complexity problem of cyber-physical systems. One is about the design space of a system for a specific goal. The design space consists of all the variable and process parameters that need to be considered. 'When you have an enormous design space, the right set-up is difficult to determine', explains Munk. 'Within the Design Space Exploration 2.0 project, the University of Amsterdam is looking for methods that help narrow down the design space.' ASML is directly involved with this project, as wafer scanners are specifically designed for a particular customer.

Other complexity issues arise from scheduling certain manufacturing processes. For example, in industrial printing, a page must stay a certain amount of time in the printer to dry. 'But the printer should still be scheduled to print the next pages correctly. Better software is needed to improve this.' That is the goal of the project Scheduling Adaptive Modular Flexible Manufacturing Systems, of which Canon is reaping the benefits.

Testing is also a challenge of increasing system complexity. In the project Testing in Time of Continuous Change, this challenge is approached from a new perspective, says Munk. 'Instead of testing the whole system in detail, Radboud University and University of Twente are trying to work from the description of what the purpose of the system is.' Both ASML and Canon will be able to use such new testing software and are therefore partners in this particular project.

SOFTWARE RESTRUCTURING

Last but not least, in the Programming and Validating Software Restructurings project, the researchers are trying to upgrade so-called legacy software. 'This is software that was created many years ago and is still used', Munk says. 'To be able to maintain and update this software, regular restructuring is required to support modern frameworks and technology.' This is the project in which Philips Healthcare is one of the partners.

'We aim to make the ever-increasing complexity more manageable by designing, building and testing new tools'

The Image Guided Therapy business in Philips Healthcare is the market leader in non-invasive interventional cardiology. 'Using our Azurion system, a cardiologist is able to navigate catheters and guide wires through the blood vessels, ensuring accurate placement and effective treatment during cardiac procedures', says Daan van der Munnik, technical department manager software at the company. 'These systems are very complex, involving mechatronics, X-ray generation, image detection and image processing. They have a long lifetime, often exceeding 15 years. During that time, the customer expects software updates and hardware maintenance. So, we must ensure new software will run on old hardware, and new hardware will support old software.'

That is not all. The software of the machine has millions of lines of code, and that amount is growing with each update. 'As more code could potentially lead to more bugs, we are continuously looking for ways to reduce the overall amount of code. By participating in the MasCot programme, we are working with academia on finding novel solutions to these problems.'

REJUVENATE THE CODE

How does that restructuring work exactly? According to Van der Munnik, the process consists of three tracks. 'The first track focuses on making the tools to extract the essential information from the code, applying a restructuring and then writing it back into the code. In the second track, we want to define a domain-specific language to describe these restructuring patterns. And the third track will focus on validating the correctness of the resulting code because we want to be sure the behaviour of the machine before and after restructuring is the same.' Together with the MasCot partners, Van der Munnik thinks some leaps can be made towards this end goal.

This cooperation between the participants is going quite well, according to both Van der Munnik and Munk. Each project has a user committee that represents the industrial partners. PhD students also spend 1 to 2 days a week doing research at the premises of the companies. 'Now the programme's last year has started and we are seeing the first examples of successful applications', Munk says. 'Like advanced testing procedures for wafer scanners at ASML and industrial printers at Canon.'

It is expected that the complexity of cyber-physical systems will keep on increasing over the next few years. 'So, we need to move to a higher abstraction level to manage that', says Van der Munnik. 'For Philips, it is essential to be able to continuously rejuvenate code to keep it young and fresh. We cannot do this manually; we will always need these kinds of technologies to stay on top of the challenge.'

'We are seeing the first examples of successful applications'

MASTERING COMPLEXITY RESEARCH PROGRAMME

AIM
Investigate and deliver the next generation of engineering methodologies, integrating a number of formalisms, techniques, methods, and tools to help manage the increasing system complexity.

DURATION
The programme has started in 2019. The average duration of granted projects is 4 years.

PARTNERS
NWO, TNO ESI, an industry partner board consisting of ASML, Canon, Philips, Thales, Thermo-Fisher Scientific, and academics from Eindhoven University of Technology, TU Delft, Radboud University, and University of Twente.

More information **www.nwo.nl/mascot**

GETTING A HEAD START

By Sonja Knols

Aske Plaat

Scientific Director of LIACS

'The first time informatics was mentioned as one of the sectors to receive sector plan funds, the news was extremely welcome. All over the country, student numbers were going through the roof. Here in Leiden, we spent a good half year developing a coherent set of plans for all four sectors that were to be stimulated with funding back in 2018. For computer science, this resulted in six new positions.

We decided to divide the funding over topics that are either at the basis of informatics, like programming languages and theoretical informatics, or upcoming, like the combination of artificial intelligence and reasoning. With Nele's appointment, we followed up on a recommendation made by an external visitation committee. They observed that our education on cyber security in the bachelor could use an impulse. The sector plan funding made it possible for us to fill in that void. At the time we hired Nele, we had one assistant professor on the topic. And now, only three years later, the group has grown into one of the biggest of our institute, with some thirty members.

Not only for our institute but also for the Netherlands as a whole, the sector plans have resulted in an important rejuvenation and reinforcement of scientific staff. In the eight years that I've been a scientific director here we have more than quadrupled in terms of staff, partly as a result of the sector plan funds.

And we have been able to strengthen the foundations of the discipline and to explore exciting new directions. With the positions we've created in quantum technology, machine learning and large language models, we've been able to get ahead of the troops in emerging fields. When we wrote the plans, large language models weren't that much of a subject. And now, we have a dedicated group with two full professors working on it, and even society is aware of the importance of this topic with applications like ChatGPT.

The sector plan funds have enabled us to determine a strategic line and build a coherent workforce. Computer science is a relatively young but fast-growing scientific discipline. Investments like these are crucial to not only keep up with the growth but to also get ahead and set the agenda of future technology development.'

The two successive sector plans have resulted in a significant number of new hires at various Dutch universities. Aske Plaat, Scientific Director of the Leiden Institute of Advanced Computer Science, explains how the funds have enabled the institute to strengthen the foundations and take a head start in promising new developments. And professor Nele Mentens explains how her appointment kickstarted the fast-growing new security branch for the institute.

Nele Mentens

Professor of Applied Cryptography and Security at LIACS since June 2020 and Professor of Embedded Systems & Security at KU Leuven

'Ever since I visited LIACS for the first time, I have really liked its open atmosphere. As I saw ample possibilities for interesting collaborations between my group at KU Leuven and the people in Leiden, I applied as soon as a vacancy came up.

My research is aimed at hardware security and embedded security. The leading question is how we can efficiently secure small devices and make chips with built-in efficient and cost-effective cryptography. For example, I develop practical security solutions for Internet-of-Things devices that take into account factors such as energy usage, performance and cost of goods of the device. Nobody will use a smart coffee machine that has become ten times slower due to security measures or a mobile device of which the battery is draining too fast due to cryptography procedures.

Over the past three years, I have been involved in several joint project proposals with other groups here. Currently, I am working on two projects that combine several strengths of the institute. In one project, we are looking at chips that execute artificial intelligence algorithms. In a traditional cryptographic chip, side-channel attacks, for example, measuring the power consumption or time duration of specific procedures, are used to discover secret data. In AI chips, however, the architecture is often a valu-

able asset for the owner and therefore needs protection. We want to know what kinds of side-channel attacks can be used to acquire this information so that we can develop efficient countermeasures.

The second project is about remote updates and reconfigurable chips. How can we ensure that the configuration that must be uploaded is free of viruses and malware? And how can we identify components in the software that can damage the chip, for example by drawing too much current?

My appointment here in the Netherlands provided me with an inspiring extra working environment and made it a lot easier for me to collaborate with other Dutch knowledge institutions. For example, just after I came here, the PROACT project was awarded, which is a collaboration with Radboud University and Riscure. This project allowed me to hit the ground running in my new job and has meant an important impulse for our Systems and Security group as a whole.'

RAISING AWARENESS
& ACQUIRING FUNDING

By Marysa van den Berg Image Ivar Pel

The amount of funding computer science research receives is not in proportion to the societal value of the field, argues Hajo Reijers, chair member of the Round Table of Computer Science, an advisory committee to NWO that helps researchers in the field to acquire funding.

'Before 2018, the board of NWO had the feeling that even though their purpose was to help fund the sciences, they lacked a structured way of engaging with researchers. So, in that year, the round tables were established. The members of the Round Table of Computer Science meet regularly to discuss the latest plans of NWO and the developments concerning the funding of computer science research. Our feedback and advice are then channelled to the board of NWO.
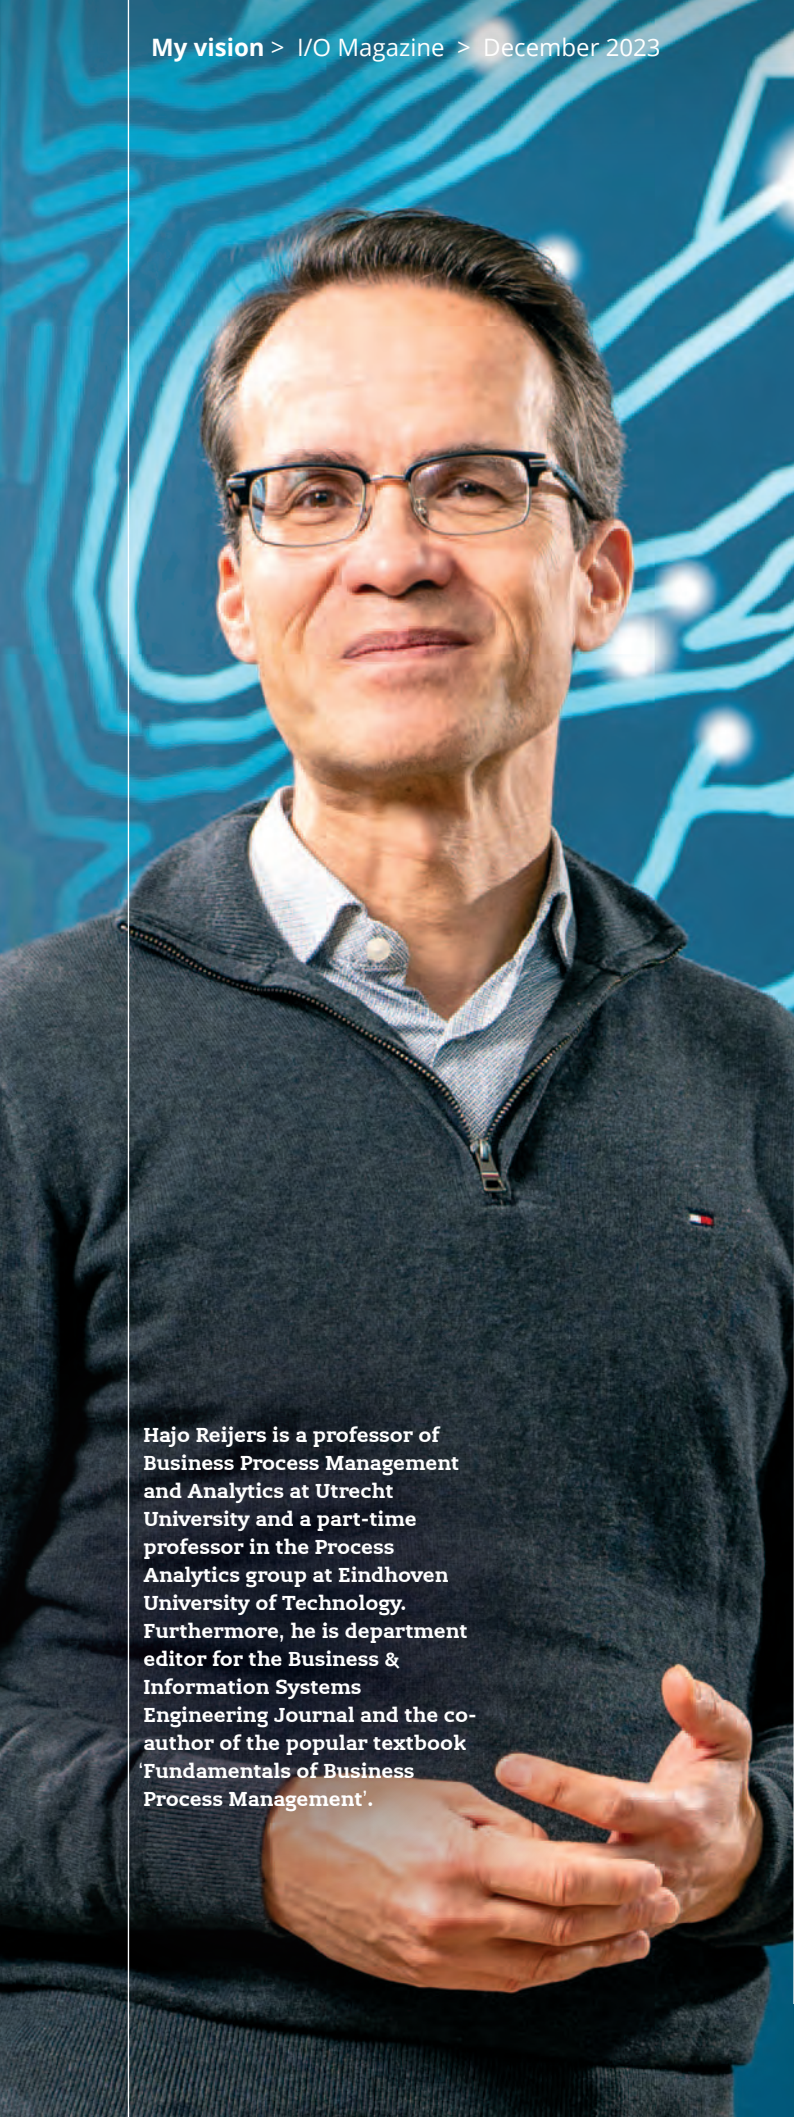
The benefits of the round table for NWO are clear; we help them to better understand the nature of computer science research and the needs of researchers in the field. And for the scientists themselves, we can make it easier to get their feedback heard by NWO and to make sure that their "peculiarities" are taken into account. By the latter, I mean that all scientific disciplines have their own specific ways of evaluating what good research is. I think it is important that a funding agency knows this view. For example, we can help with that by providing disciplinary sketches to reviewers who are not computer scientists themselves.

YOUNGER MEMBERS NEEDED

'I still think some improvements need to be made concerning how the table operates though. For example, in previous years, there was a strong tendency to involve senior scientists as members: those who already have an established career. But now we also try to include younger members, because you can imagine that getting funding early on in a career is different than when you already have a lot of proposals granted.

For the future, I expect that we will become a well-known, approachable, and representative group of computer scientists who know how to get in touch with NWO for all kinds of issues. The goal is a smooth cooperation with the funding agency.

In general, I hope that the funding of computer science research will grow in the coming years. I have the feeling that the amount allotted to this field is not in proportion to the enormous value computer science has for society. And let us not forget how important the field is for the other sciences, too, because there are hardly any domains that do not use computational methods nowadays. So, helping computer science can also help other disciplines to prosper. My vision is that the Round Table of Computer Science can contribute to realising all these benefits.'

Hajo Reijers is a professor of Business Process Management and Analytics at Utrecht University and a part-time professor in the Process Analytics group at Eindhoven University of Technology. Furthermore, he is department editor for the Business & Information Systems Engineering Journal and the co-author of the popular textbook 'Fundamentals of Business Process Management'.