Faszination **Forschung**

TUM Research Highlights
Technische Universität München | The Science Magazine
July 2017 | Edition 20

Special Issue:
Digitalization

Augmented Reality: Providing users with the right information at the right time
Bioinformatics: Working towards individualized therapies
Smart Factories: Seamless communication with self-aware modular components

ISSN 18653022

Nominal charge
EUR **9.00**

Dear TUM friends and associates,

The word "digitalization" may sound technical and even impersonal, but it's actually one of the most people-focused technologies out there. For technology to provide workable solutions supporting the vast spectrum of human needs and abilities, digital solutions are essential – making this a core field for us as a technical university. New disciplines such as bioinformatics pave the way for personalized medicine, for instance. Equally, digitalized industrial systems harness mass-production methods to deliver individual items tailored to our personal preferences. Meanwhile, computer scientists are developing tools to turn the massive streams of data generated by digital technologies into useful and individualized information. And mathematical models help us understand and manage the dynamics of social interactions.

This edition of our magazine introduces scientists from the most diverse disciplines – all of whom have one thing in common: digital technology lies at the heart of their work.

Hans-Werner Mewes is a pioneer in the field of bioinformatics. Here, the focus lies on how to harness the gigantic data volumes generated by biology and medicine. In his interview, Hans-Werner Mewes traces the development of this discipline from its origins right through to the present day. Meanwhile, bioinformatician Burkhard Rost, one of our Humboldt professors, examines the evolution of human DNA to enable predictions about the function of specific proteins in individual people.

Augmented reality (AR) is a technology that provides us with highly contextualized digital information. Gudrun Klinker is investigating tools to optimize the way we interact with AR. At the same time, Nassir Navab is bringing this technology into the operating theater. He has developed a system that provides surgeons with additional information in real time as they operate – for instance displaying the precise location of a tumor within their field of vision.

Frank Diermeyer is looking at autonomous driving from a fresh angle. Convinced that driver assistance systems alone will not be able to deal with every possible traffic scenario, he has developed a solution where another human intervenes remotely if the driver and digital assistants have reached their limits.

Batch size one is industry's holy grail. For automation engineer Birgit Vogel-Heuser, the key to achieving this lies in self-aware modules, able to act autonomously and organize themselves thanks to uniform semantics.

Massimo Fornasier applies mathematical modeling methods and the laws of physics to analyze social dynamics. He is investigating how crowds of people can be rapidly and safely directed in evacuation scenarios, for instance, and how public opinion can be influenced. Shaping social dynamics also preoccupies Simon Hegelich from the TUM School of Governance and the Bavarian School of Public Policy in Munich. His research focuses on bots engaging in automated mass-messaging on social networks. Needless to say, this is now highly topical – particularly against the backdrop of the US presidential race and the upcoming parliamentary elections in Germany.

So as abstract as digitalization concepts may initially appear, an encounter with our researchers well and truly brings them to life.

My hope is that all of these articles will give you fascinating new insights into the fast-paced progress of technical science today.

Prof. Wolfgang A. Herrmann

Picture credit: Heddergott

Table **of Contents**

Picture credits: Eckert, ediundsepp, Jooss, Graphics: ediundsepp

"Digitalization is a core field for us as a technical university." *Wolfgang A. Herrmann*

In every Issue

Clear View **from Afar**

Research and industry around the globe are working to bring autonomous driving to our streets. Dr. Frank Diermeyer and his team, though, are convinced that situations will still regularly occur that exceed the capabilities of autopilot. That is why they have developed a type of remote control for cars. In future, this will allow service staff at a remote command center to steer cars safely through traffic when autopilot hits its limits. TUM's smart demo vehicle, Mute, shows just how well this teleoperated driving concept can work.

Picture credit: Eckert

Remote driving: An operator takes control of the car and steers it safely through the traffic, based on video and sound information transmitted from the car to the operator's console.

Picture credit: Eckert

Tim Schröder

Durchblick aus der Ferne

Weltweit arbeiten Autohersteller intensiv an Autopiloten. Das Team um den Maschinenbauingenieur Dr. Frank Diermeyer vom Lehrstuhl für Fahrzeugtechnik aber schlägt jetzt ein neues Kapitel des autonomen Fahrens auf. „Wir arbeiten an der Fernsteuerung von Autos, dem teleoperierten Fahren“, sagt Diermeyer. Damit sprengt er das Paradigma vom völlig selbständigen Auto. „Wir sind davon überzeugt, dass es immer wieder knifflige Situationen geben wird, die Fahrerassistenzsysteme nicht von allein lösen können“, sagt Diermeyer. „Es wäre hilfreich, wenn ein Assistenzfahrer einspringen könnte, der in einer Fernsteuerzentrale sitzt.“

Die Forscher koppeln dafür das Versuchsfahrzeug der TUM, das Elektroauto Mute, über eine schnelle LTE-Mobilfunkverbindung mit einem von ihnen entwickelten Fernsteuerplatz. Dieser verfügt über drei große Bildschirme, auf die die Kamerabilder aus dem Mute übertragen werden: der Blick durch die Windschutzscheibe, durch die Seitenfenster und aus der Heckscheibe. Im Cockpit befinden sich ein bequemer Autositz, ein Armaturenbrett und die Pedale für Gas, Bremse und Kupplung, mit denen das Auto ferngesteuert wird. Eine solche Fernsteuerung ist nach Ansicht von Diermeyer zum Beispiel für ältere Menschen sinnvoll, die sich mit ihrem Auto kaum mehr in die Innenstädte trauen. Autopiloten können diesen Job bislang nicht übernehmen, weil auch sie in den Innenstädten schnell überfordert sind. Derzeit geht Diermeyer davon aus, dass das teleoperierte Fahren zuerst in Städten bei der Suche nach Parkplätzen zum Einsatz kommen wird. Hat man das Ziel erreicht, wird man aussteigen und dem Operator das Auto übergeben, der es zu einer freien Parkbucht lenkt.

Für die Fernsteuerfunktion mussten die Forscher eine Fülle technischer Probleme meistern. So gibt es einen Zeitverzug, mit dem die Kamerabilder aus dem Auto in der Fernsteuerzentrale ankommen. Dieser führt zu einer ungleichmäßigen Bildschwankung, die einen Operator schnell nervös macht. Die Experten der TUM verzögern die Signale deshalb künstlich um 500 Millisekunden und erhalten damit einen zwar verzögerten, aber gleichmäßig fließenden Datenstrom. Damit der Operator trotzdem exakt die Position des Autos kennt, wird im Kamerabild die vom Computer errechnete reale Position des Autos als einfaches Viereck angedeutet.

Mittlerweile haben Diermeyer und seine Kollegen mehrere Anfragen von Industrieunternehmen für eine Zusammenarbeit erhalten. Diermeyer hofft, dass der Mute in drei Jahren soweit ist, dass er tatsächlich ganz allein und ohne Sicherheitsfahrer an Bord auf öffentlichen Straßen fahren kann. Wenn man sieht, wie sicher der Mute schon jetzt seine Bahnen fährt, glaubt man das gern – und auch, dass nicht allein dem Autopiloten, sondern auch der Teleoperation die Zukunft gehört. □

Link
www.ftm.mw.tum.de

Diermeyer's group is currently testing teleoperated driving in a mass-produced vehicle, which is also equipped with sensors monitoring its surroundings.

An SUV appears outside the Institute for Automotive Technology at TUM's Garching campus. A quick glance behind the steering wheel rivets you to the spot – the driver's seat is empty. The SUV can steer itself! It zips past, weaving its way through an obstacle course of red and white traffic cones, veering left here and right there. Taking a closer look, you might start to think again, reflecting that perhaps it is not really all that surprising for a car to be driving itself around these days. After all, many researchers and industrial enterprises are now advancing autonomous driving worldwide. Electric carmaker Tesla, for instance, is actively pushing the driver assistance capabilities of its sedans, which enable cars to cruise along highways almost by themselves. Other manufacturers, too, are equipping their models with cameras, radar and laser sensors to support features such as active brake assist, lane departure warning and stop-and-go in heavy traffic. Now, though, mechanical engineer Frank Diermeyer and his team from TUM's Institute of Automotive Technology are marking a new departure in the autonomous driving story. They demonstrated their technology 2013 in Mute, a compact electric runabout developed from the ground up by various TUM working groups over the past few years. Now, Diermeyer's group is also using mass-produced cars for their tests. "We are working on controlling cars remotely through teleoperated driving," confirms Diermeyer, thus challenging assumptions around fully autonomous cars. Most experts claim that they are aiming to bring fully automatic cars to the streets within the next few years, able to navigate independently with the aid of sensors. "We, however, are convinced that tricky situations will keep arising that driver assistance systems will not be able to resolve by themselves," explains Diermeyer. "In these cases, it would be helpful if an assistant driver could intervene from a remote command center."

Teleoperated driving on TUM campus

TUM researchers can already show us what a remote center like this might look like. They set up a type of simulator cockpit in their lab. Inside hang three large screens, displaying camera images from Mute: the view through the windshield, through the side windows and to the rear. The cockpit also contains a comfortable car seat, a dashboard, and gas, brake and clutch pedals for remote control of the vehicle. When Mute drives around outside, the simulator makes you feel as though you are in the car yourself. But why go to all the effort? If the autopilot fails, the car's driver could surely step in and resolve the situation themselves. "That's true," acknowledges Diermeyer. "But in some cases, remote control makes perfect sense." Here, he is thinking of older people, in particular, who still like to drive their cars on rural roads away from cities but would not usually be confident in downtown traffic any more. As it stands, autopilot cannot help here either, since these systems also quickly reach their limits in busy cities. Signpost clusters, double-parked cars and obscured lane markings usually still trigger an emergency stop. According to Diermeyer: "I don't see these autopilot problems being solved in the next few years, so an alternative here would be an operator in a remote command center." The weather, too, currently stumps autopilot – particularly fog, heavy rain and driving into the light. All situations that quickly overwhelm older people as well. Snow drifts are also problematic for self-driving cars: while a human driver can assess how deep the snow is and whether the car can get through it, autopilot simply stops in front of the white obstacle. "In any case, we don't see teleoperated driving as competition, but as a valuable addition to automated driving," clarifies Diermeyer. "In combination, the two technologies could help seniors to stay mobile for longer. And, especially for people living in the countryside, that plays a role in combatting social isolation too."

Truck drivers also stand to benefit from teleoperated driving. They have to take a break every few hours, but with a com- ▷

Camera systems record the front and side view from the cockpit. All images are transmitted via mobile communication to the operator, who controls the steering wheel and pedals.

bination of autopilot and remote control, the vehicle could continue driving by itself. Autopilot alone would not suffice here, since it could get stuck in tricky situations and truck drivers would not be allowed to intervene if they were on a break. Teleoperation, however, would be a real help. Diermeyer currently anticipates that teleoperated driving will initially be introduced in downtown areas – when looking for somewhere to park. The hunt for parking spaces alone is thought to account for 30 percent of traffic in major cities today. So it would be a significant relief if an operator familiar with the location could take over the car and park it. Close to the destination, the operator would then navigate to a vacant parking space via the shortest route. To ensure this functions smoothly, the remote command center would be linked up to a parking information system.

Eliminating jitter in the operator cockpit

Diermeyer and his colleagues have been focusing on autonomous driving since 2006, moving on to teleoperation in 2010: "Because it became clear to us that we're going to need both. We are one of very few research groups to have this type of system up and running today," underscores Diermeyer. The team has now been able to demonstrate its capabilities in several test drives. However, it was a long road to get there, with the researchers confronted by numerous technical challenges along the way. One of them was the time lag between camera images from the car reaching the control center and the corresponding commands being relayed back to the vehicle. Mute uses the rapid mobile communication standard LTE. Nonetheless, depending on reception and available bandwidth, the images and driving commands are always transmitted with at least a slight delay. The jerky effect this creates in the video stream is known to the experts as jitter. In the simulator cockpit, it meant that the car would sometimes seem to lunge forward and sometimes hesitate or even come to a complete halt. And this jitter quickly stresses even the most skilled or most experienced of operator. Hence, the TUM experts decided to incorporate an artificially fixed delay of 500 milliseconds into their signals. "That is the rough extent of the time lag we are dealing with," explains Diermeyer. "So this now eliminates jitter and gives us an even data stream, albeit with a slight delay."

This delay, though, presents the researchers with their next problem. The image in the remote control cockpit is now permanently 500 milliseconds behind, so never shows the precise position of the car at any given time. And the faster the car is moving, the larger the gap between its actual position and the image the operator sees. At 60 kilometers per hour, the car can travel a whole 8.3 meters in the space of 500 milliseconds. The researchers thought long and hard about how to give the operator an impression of the actual situation in spite of this – and came up with a surprisingly simple solution. The computer calculates the car's real position, and this is displayed in the camera image as a simple square. Two small white spots are also visible, indicating the actual position of the front wheels. "Our tests show that the remote control operators adapt well to this after a short period of training."

Remote control always comes with the risk of the operator becoming disconnected from reality – resulting in "loss of situational awareness", as psychologists call it. To avoid this effect, the researchers have a host of technologies in their cockpit to ensure that the simulated journey feels just like a real car ride. Loudspeakers are used to transmit the real background noise – including the sound of the vehicle on the road, as well as the horns of other cars. Driving on ice makes itself felt by a change in steering motion, and uneven ground by a shaker fitted to the seat. Operators also gain a sense of the car moving thanks to the camera image in the side windows imitating the natural blur of motion. ▷

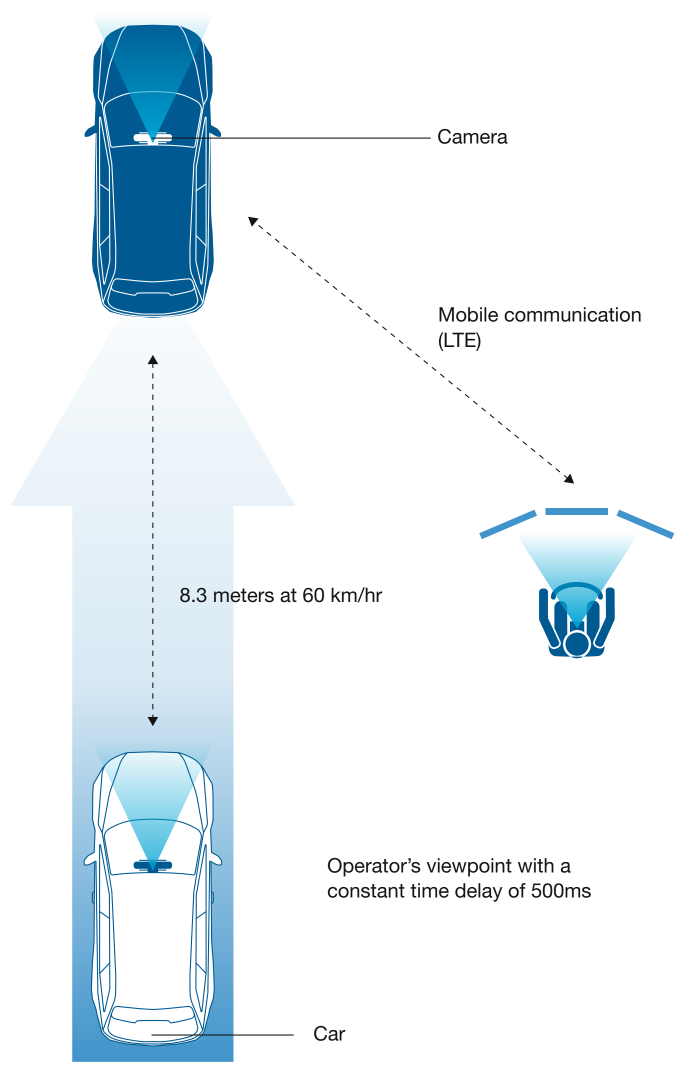

An even data stream: Depending on the quality of the mobile network, jitter can occur in the video stream the operator receives. A deliberate time delay of 500 milliseconds ensures a constant video stream, albeit with a time lag that would translate into a distance of about eight meters at a speed of 60 km/hr. The computer calculates the car's real position and displays this information in the operator's console.

"We are convinced that tricky situations will keep arising that driver assistance systems will not be able to resolve by themselves. In these cases, it would be helpful if an assistant driver could intervene from a remote command center."

Frank Diermeyer

Graphics: ediundsepp, Picture credit: Eckert

Novel visual techniques

The TUM researchers developed another visual technique aimed also at ensuring a natural impression of the driving experience – namely data goggles. These show the outline of the car, precisely superimposed on the camera images of the surroundings displayed by the monitors. As experiments in Munich showed, any unsteadiness or slight discrepancies quickly lead to the user feeling unwell. As long as the driver looked straight ahead at the windshield camera image, everything was fine. But if they turned their head to the side, the system switched to the side window. The result was a brief time lag between the windshield and side-window images with each turn of the head – literally giving the operator travel sickness. Diermeyer's team thus decided not to transmit the various camera images separately, but to merge them within fractions of a second to produce a wide-angle image of the surroundings, which now stretches from left to right across all monitors like scenery at the theater. Since the view changes rapidly as the car drives along, the image data has to be computed extremely quickly. "The usual MPG-4 video standard would have been too slow for this," reveals Diermeyer. "We had to modify it and adapt it for our specific application needs."

Right from the start, Diermeyer's aim has been to develop an effective system for teleoperated driving that fulfils all the safety requirements necessary for use in real-world traffic situations. This also applies in the event of connectivity problems with the remote control center. Anyone who has made phone calls in a car knows that the connection cuts out in a dead spot with no signal – and of course this applies to remote control too. Teleoperation is primarily envisaged for downtown driving, however, and strong network coverage means minimal dead spots there. TUM's remote solution is currently designed for use at speeds of up to 60 kilometers per hour, which is quite sufficient for urban areas. If the network connection is ever actually lost, the car will come to a standstill – either gradually or by means of an emergency stop, depending on the situation. "At speeds under 60 kilometers an hour, putting the brakes on is the safest reaction in any case," emphasizes Diermeyer. But of course the aim is to prevent any interruption to the data connection if at all possible. With that in mind, the researchers have enriched their system with maps from mobile network operators, indicating the dead spots. This should enable the vehicle to bypass them – for instance by avoiding tunnels.

Meanwhile, Diermeyer and his colleagues have already been approached by several industrial companies interested in collaboration. Diermeyer hopes that Mute will have advanced sufficiently in three years' time to be able to drive around public streets all by itself, with no backup driver in the car. When you see how reliably it already steers its course, that seems entirely possible. As does the prospect of a future shaped not just by autopilot capabilities, but also by teleoperated driving. *Tim Schröder*

Dr.-Ing. Frank Diermeyer

Research dedicated to the elimination of road accidents

One thing is for sure: Germany does not have many researchers lecturing on motorcycle technology, so Frank Diermeyer is something of a rarity in that regard. And no wonder, since motorbikes are his passion. After many years riding sports models, he recently acquired his first touring bike: "In keeping with my advancing years," laughs the 41-year-old, describing his Honda CBF 1000 with 100 HP. "Ideal for cruising through the Alps – especially the Dolomites."

Diermeyer has always been a motoring enthusiast. In his leisure time, he rides his motorbike; at work, he researches cars. He started off working on the electronic stability control (ESC), primarily investigating vehicle rollovers. His task was to further optimize the ESP feature, and he wrote his doctoral thesis on this topic at TUM in 2008. "So it was just a small step from driver assistance systems to automated driving," he recalls, "since a really effective assistance system can only become reality if the car has the ability to drive itself." He thus moved on to researching autopilot capabilities in 2010.

"Vision Zero" remains Diermeyer's overriding aim; in other words, the elimination of road accidents. He is sure that automated driving has a role to play here, since driver assistance systems react faster than humans and cannot be distracted. Most accidents today occur due to driver error. For Diermeyer, autopilot finally brings Vision Zero within reach. When he first went into vehicle research fifteen years ago, passive safety was still the major concern, with technology such as airbags and seat belt pre-tensioners designed to protect the driver in a crash. Since then, however, active safety has become the number-one topic, with cars preventing accidents of their own accord by means of distance detection sensors and lane departure warning systems – and, in the near future, through automated driving. "Autopilot is the next big thing," declares Diermeyer – and that is very good news indeed from his perspective.

“We are one of very few research groups to have a system for teleoperated driving up and running today.”

Frank Diermeyer

Picture credit: Eckert

Link

www.hfp.tum.de
politicaldatascience.blogspot.de

Campaign 4.0 – **On the Trail of the Elusive Bot ...**

Online social networks have now become a part of political discourse, recently attracting a great deal of attention due to their role in the Brexit referendum and presidential elections in the US and France – not to mention the emotive tweeting of a certain US President. In the struggle to gain influence, intelligence services, terror organizations and political movements are increasingly deploying social bots, whose mass-messaging can polarize, mislead and unnerve other users. Simon Hegelich, Professor of Political Data Science at the Bavarian School of Public Policy within TUM, is working with his team to find ways of detecting these computer programs, helping to stop the originators in their tracks. Germany's parliamentary elections are the next big event.

Karsten Werth

Wahlkampf 4.0 – Meinungsmaschinen auf der Spur

Die sozialen Netzwerke im Internet sind wegen ihrer Rolle beim Referendum über den Brexit, bei den Präsidentschaftswahlen in den USA und in Frankreich, aber auch durch einen emotional twitternden US-Präsidenten ein viel beachteter Teil des politischen Diskurses geworden. Im Kampf um Einfluss setzen nicht nur Geheimdienste und Terrororganisationen, sondern auch politische Bewegungen vermehrt Social Bots ein, die durch ihre massenhaft verbreiteten Botschaften polarisieren, täuschen und verunsichern können. Simon Hegelich, Professor für Political Data Science an der Hochschule für Politik an der TUM, sucht mit seinem Team nach Wegen, diese Computerprogramme zu entdecken und dabei zu helfen, ihren Urhebern das Handwerk zu legen. Unter Einsatz neuer Methoden, einer Kombination aus IT-gestütztem Data-Mining und politikwissenschaftlicher Interpretation von großen Mengen von Nutzerdaten, kommen die Forscher immer komplexer werdenden Bots auf die Spur, die unter falscher Identität im Netz Stimmung machen. Im Jahr der Wahl zum deutschen Bundestag konzentrieren sich die Forscher auf die Debatte zur Flüchtlingskrise im sozialen Netzwerk Facebook. Sie haben bereits eindeutige Manipulationsversuche aus dem rechten politischen Lager festgestellt und arbeiten am Aufbau eines Monitoring-Systems, das dabei helfen kann, mehr Transparenz in die neuartige Form der politischen Willensbildung über die sozialen Medien zu bringen. □

Picture credit and Graphics: edundsepp

“A people that no longer can believe anything cannot make up its mind. … And with such a people you can then do what you please.” Hannah Arendt

It’s a dream come true,” acknowledges Simon Hegelich: “You pursue your research and suddenly everyone is interested in it. But the past few months have certainly been hectic.” Specializing in social media forensics, he is now in high demand prior to Germany’s 2017 election as an expert in media and politics. In one week in February alone, he had meetings in Berlin with parliamentary committees, the Ministry of Education and the Federal Press Office and also attended a panel debate held by public broadcaster ZDF as part of the Berlin International Film Festival. The interest in Hegelich’s research into systematic manipulation via social networks is easy to understand, since it is becoming increasingly clear that this facet of our fast-evolving technology lifestyle is also heralding radical changes to the political process. In fact, Hegelich suspects Germany could be about to witness the last traditional election campaign – and possibly the dawn of a new era.

Direct communication – one to all – without any filters

“The democratic public space is undergoing disruptive change,” emphasizes Hegelich. “In principle, social media platforms are enabling direct, one-to-all communication. This is shaking up the traditional role of political parties and the media in forming opinions. In historical terms, the situation is comparable to the invention of the printing press. That suddenly allowed information to be widely disseminated – and in a way no longer tied to the old frameworks of power and discourse.” What is more, these new communication channels are emerging at a time of increasing anxiety about the future and loss of trust in politics and the media. So on one hand, we are facing political change due to factors such as the financial crisis, economic challenges in the US and upheaval in the European Union. This change would also be occurring without social media. On the other hand, though, we now also have technology that allows us to share our feelings in public. That certainly does nothing to calm the situation – in Germany, for instance, Internet users are becoming increasingly polarized. And, as all over the world, populists are also campaigning for votes online. These forces feed into each other, making it difficult to gain a clear picture of the situation here. As Hegelich puts it: “It is hard to pinpoint cause and effect. Did Twitter really help Trump into power? Or is Twitter only still in the market because Trump is such a prolific tweeter?”

Machine against machine

Another revolutionary factor is the fact that it’s not just people who are active on social networks. Increasingly, they are being joined by machines. In 2015, Hegelich was able to show that Ukrainian ultra-nationalists were using a computer program to control 15,000 Twitter profiles, sending up to 60,000 tweets per day. Since then, he has delved even more deeply into the topic and is now investigating how robots fuel discussions and generate or strengthen opinions. Itself a player in the cyber arms race, the research community is also benefiting from progress here. Thanks to new developments in machine learning, with algorithms autonomously detecting patterns in large data volumes and enabling them to be interpreted, political scientists can now harness the power of big data in their work. At the same time, the speed at which the Internet is evolving poses a real challenge for traditional research activities. While a researcher is busy submitting project proposals to apply for funding from the relevant institutions, the online environment continues to evolve. And by the time applications have been processed, staff hired and resources secured, it might look very different again. Two years from the original project idea, for instance, the social media platforms selected for study might have changed completely or lost a huge amount of impact. Hegelich's team is still small, but he is aiming to monitor social media around the German parliamentary elections in as much detail as possible. Its members are currently working to build capacity, gather data and program analysis tools. To date, there is no software for detecting social bots that could keep up with the rapid pace of development. Somewhat surprisingly, from a non-specialist perspective, the arsenal of Munich's social media forensics experts includes Raspberry Pis – small, single-board computers. These mini-computers are used to connect to the social networks. The researchers then conduct their analysis using advanced databases, such as Elasticsearch, and high-end servers. However, this data mining does not currently involve supercomputers such as you might find in the Physics faculty, for instance.

Political and data science team up

"We do a lot with programs we have written ourselves," explains Hegelich. "For instance, we are taking part in a fake news challenge – the task being to check 50,000 articles and see whether the headline matches the text." On Facebook it is often the case that real articles are shared with fake headlines so they are more likely to be circulated. That is why users should look not only at the Facebook preview, but also at the original source. Hegelich's team is trying to automate machine-based checking in these cases. When conducting research with social media data, the interplay between IT know-how and political expertise is particularly important. "If you just let the data speak for itself, the outcome is nonsense," underscores Hegelich. "So we look at bot mistakes from every angle. Otherwise you end up with theories like the one blaming bots for Brexit. It's really not that simple though. You don't just type 'Brexit' into the Internet and people are suddenly in favor of it." If you use machines for data analysis, continues Hegelich, you need to know what that means: "If I've developed an algorithm that classifies data correctly in 95 percent of cases, that sounds absolutely great. But if I'm using it to sort a billion posts per day, it will get a substantial number of them wrong. So it's not enough to look at just one aspect." ▷

5,000,000

GER

Tweets concerning German politics are posted every day

15,000

UKR

Twitter bots in Ukraine send 60,000 tweets per day

15,000

GER

comments against refugees were posted from some single Facebook accounts

500,000

GER

Facebook users have liked posts from more than one German political party

9 – 15%

USA

of all US Twitter accounts could be social bots*

*Alessandro Bessi, Emilio Ferrara. Social bots distort the 2016 U.S. presidential election, online discussion. First Monday 21(11), 2016

Quote taken from The New York Review of Books, issue October 26, 1978; Graphics: ediundsepp (Source: Hegelich)

How hyperactive users affect social network structures

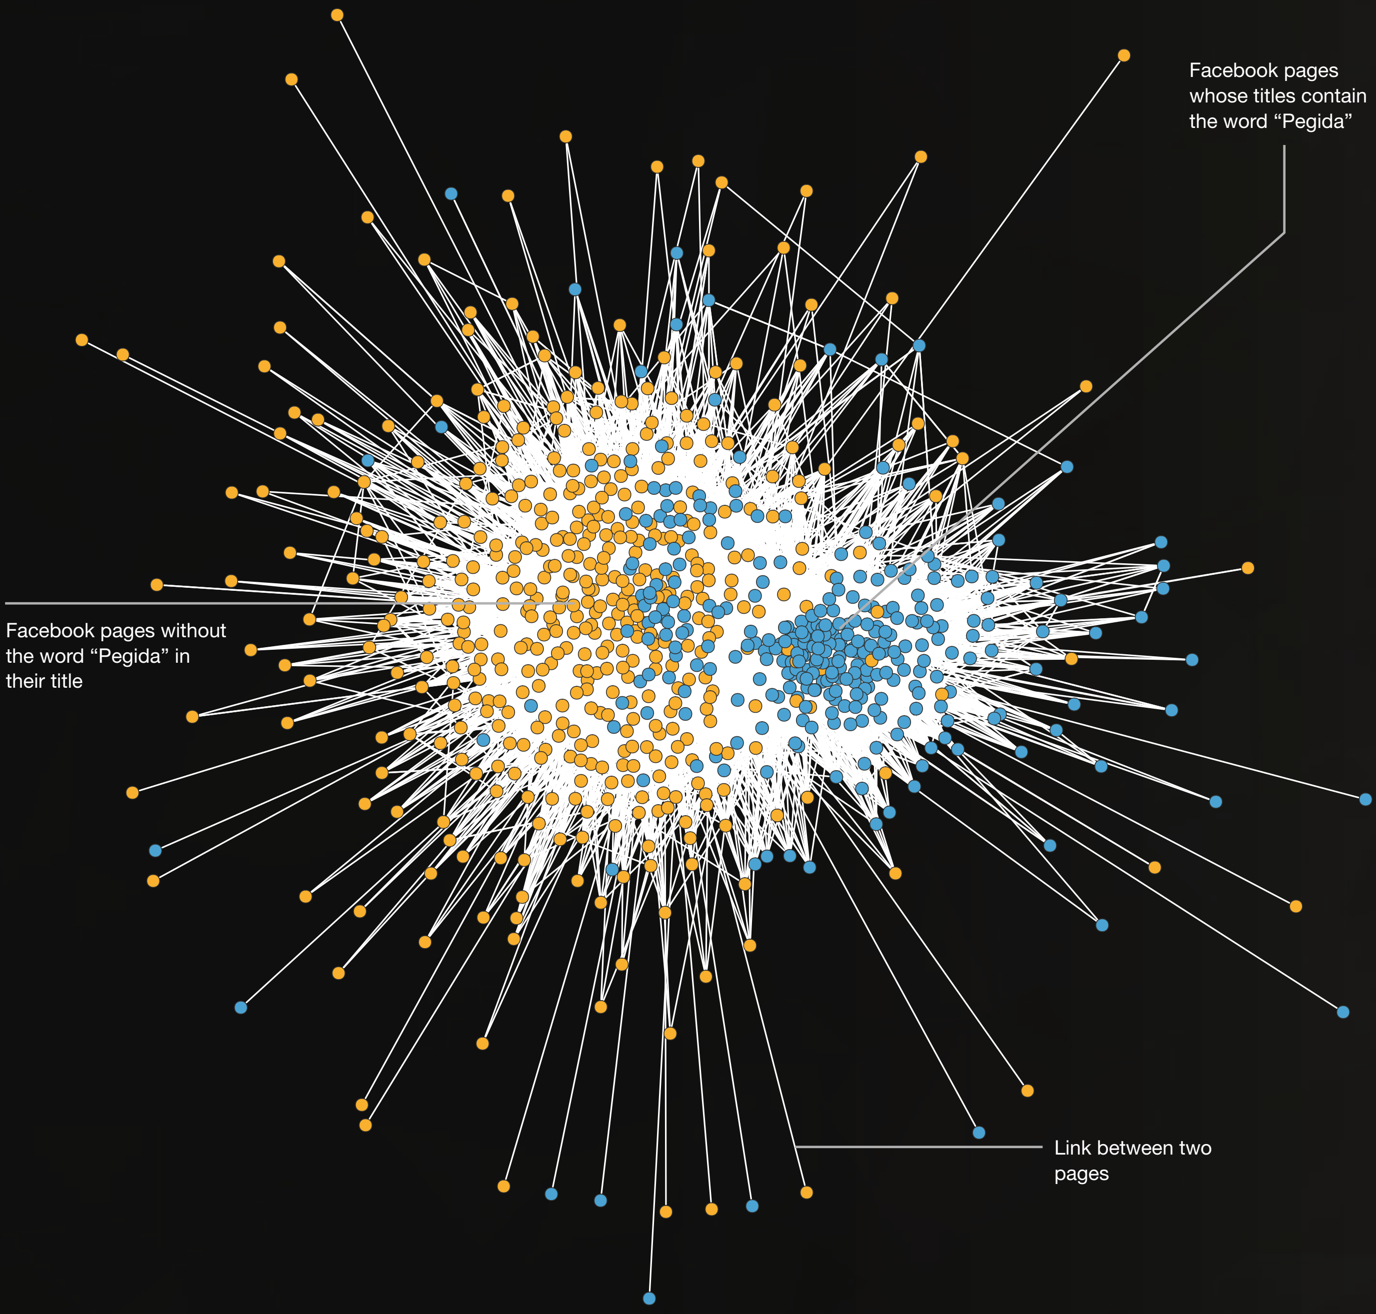

Can the structure of a social network be manipulated? This is an important question, because the algorithms based on which social networks display content to their users analyze the network structure. In 2016, Hegelich analyzed nearly 1,000 German Facebook pages (including posts, comments and likes) of all political parties, main media and of all pages with the word “Pegida” in the title. Pegida is an anti-immigrant initiative in Germany. Orange dots stand for “mainstream” pages, blue dots for pages which have “Pegida” in their title. The lines indicate that a user has linked these two pages. The resulting network is cleary separated into two clusters. The few blue dots within the orange network stand for “anti-Pegida” pages, which also use the word “Pegida” in their title.

All users

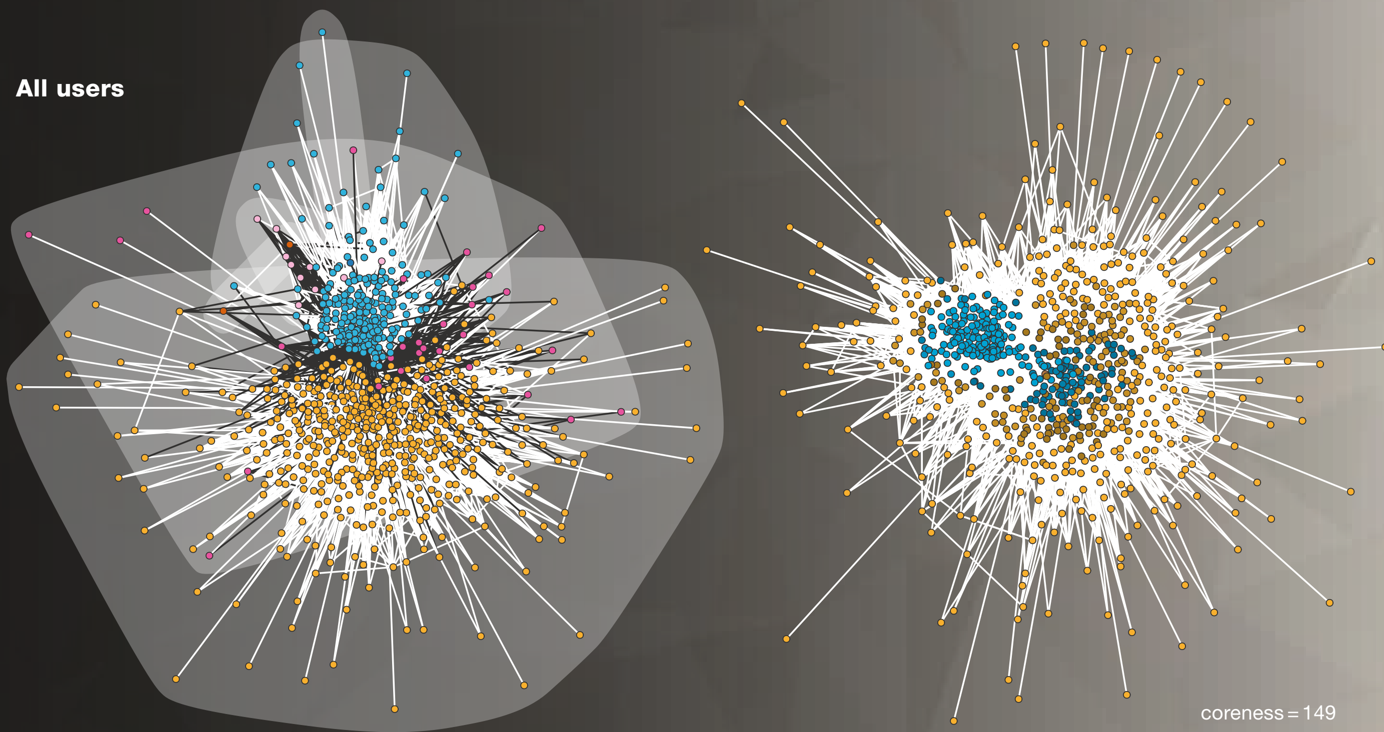

Without the 1 percent most active users

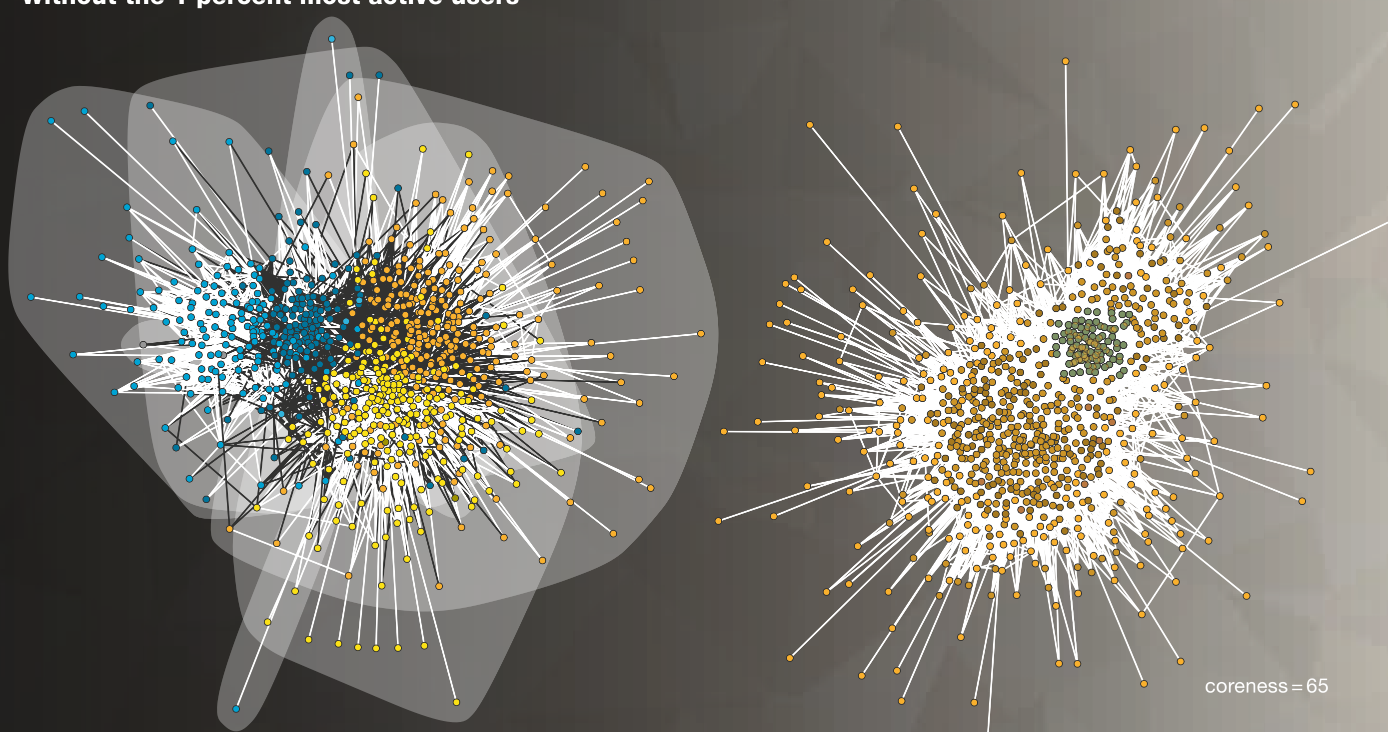

Graphics: ediundsepp (Source: Hegelich)

Does the network structure for all users (top) change when the hyperactive users are eliminated (bottom)? The graphics on the left analyze the so-called statistical distance, which describes how strongly two pages are linked. For all users, the network is separated into a “mainstream” (orange dots) and a “right-wing” (blue dots) cluster. Without hyperactive users, this structure changes: Both clusters break up into two groups. The distance between “mainstream” and “Pegida” pages shrinks, indicating higher information flow between them.
The graphics on the right consider pages which are linked to a very high number of other pages. A coreness value of 149 means that they show only pages that are linked to 149 or more pages. Orange stands for a lower number of linked pages, blue for the highest number of connections. For all users, this network is clearly divided into a highly linked and a relatively weakly linked cluster. Without the hyperactive users, the two clusters are less clearly separated and the coreness value is much lower, indicating a less closely knit network altogether.

Opinion-forming tools in social media

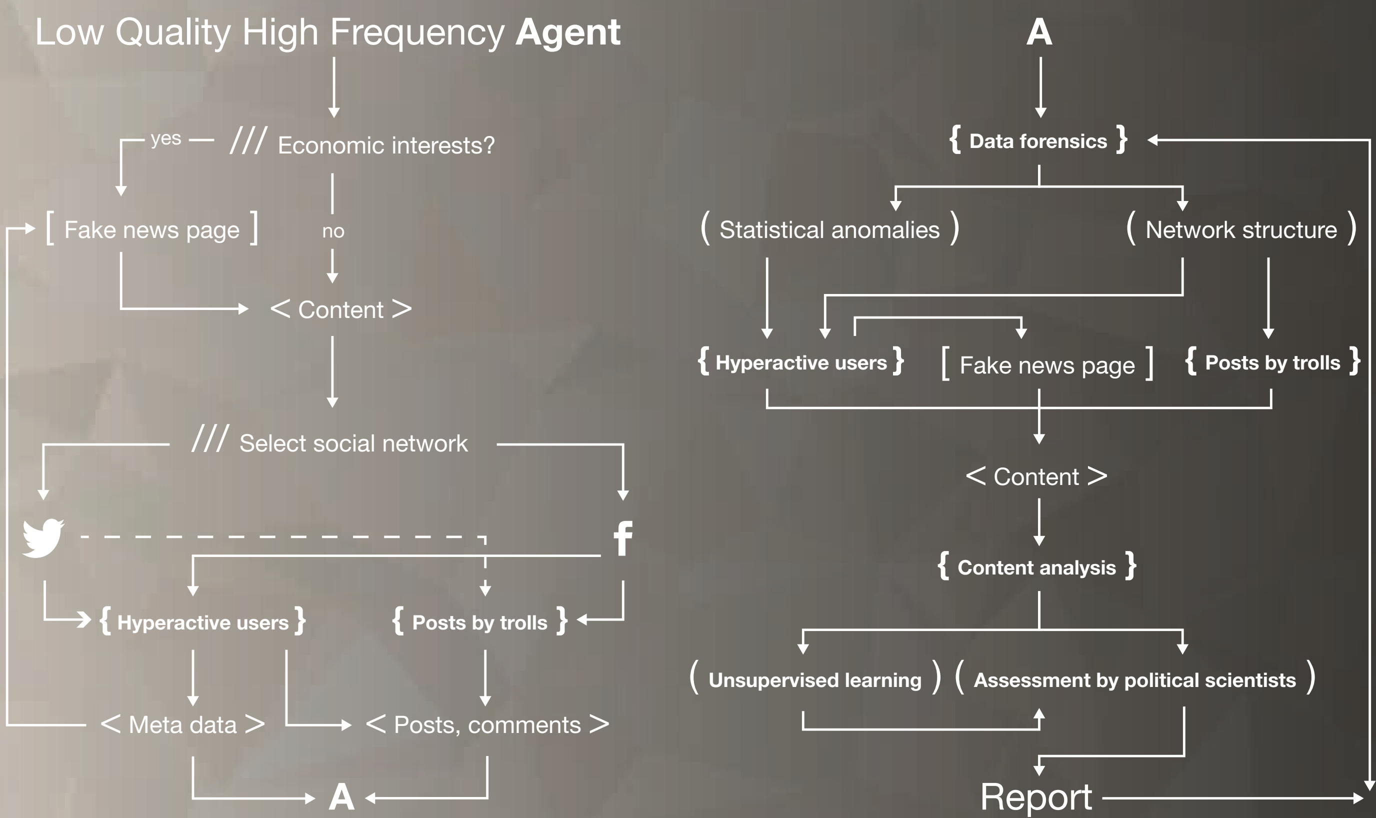

Left: Setting up a high number of relatively simple postings (Low Quality High Frequency Agents): A certain piece of content – either fake news or real news – is streamed into a selected network. Fake news pages are mostly set up in order to raise money from clicks on ads. Once streamed into the networks, the content gets distributed with the help of hyperactive users or trolls. The result (A) is a certain opinion that forms within the network.
Right: To uncover this manipulation, data scientists use statistical methods to find hyperactive users or trolls. Political scientists analyze the content of their posts and assess whether or not manipulation should be suspected.

Who or what are social bots?
The social bots under investigation by the TUM researchers are computer programs deployed for manipulative purposes that use fake accounts to emulate human identities and communication. They capture user data and spread targeted spam or political messages. In the US presidential race, for instance, bots sent out mass notifications that Donald Trump lied in the TV debate or that Hillary Clinton suffers from Parkinson's disease. They are particularly active on social networks using easily accessible application program interfaces (APIs). Twitter's low entry barriers mean it is a heavily used platform. It is currently at the center of international research, since it makes it relatively easy to procure the necessary data. However, Hegelich takes a critical view of this focus in relation to Germany, since Twitter has an extremely small following in comparison with Facebook there. Recently, Hegelich has been concentrating his own efforts on analysis of the refugee debate on Facebook. Prior to the upcoming German parliamentary elections, he intends to turn his attention to manipulation via social media that could influence the campaign. "We are also working to incorporate other platforms like 4chan, VKontakte and reddit in our analysis, but are reaching the limits of our resources there," he admits. Hegelich's previous studies, which analyzed over 30 million instances of Facebook activity in relation to the refugee issue, clearly show right-wing manipulation attempts. Hyperactive users – both human and automated – became apparent, systematically "liking" every post by Germany's right-wing populist AfD party. There are people spending eight hours a day writing hate-filled comments about refugees on the Internet. This huge engagement gives their views an exceptionally wide reach on Facebook. At some point, the site's algorithms

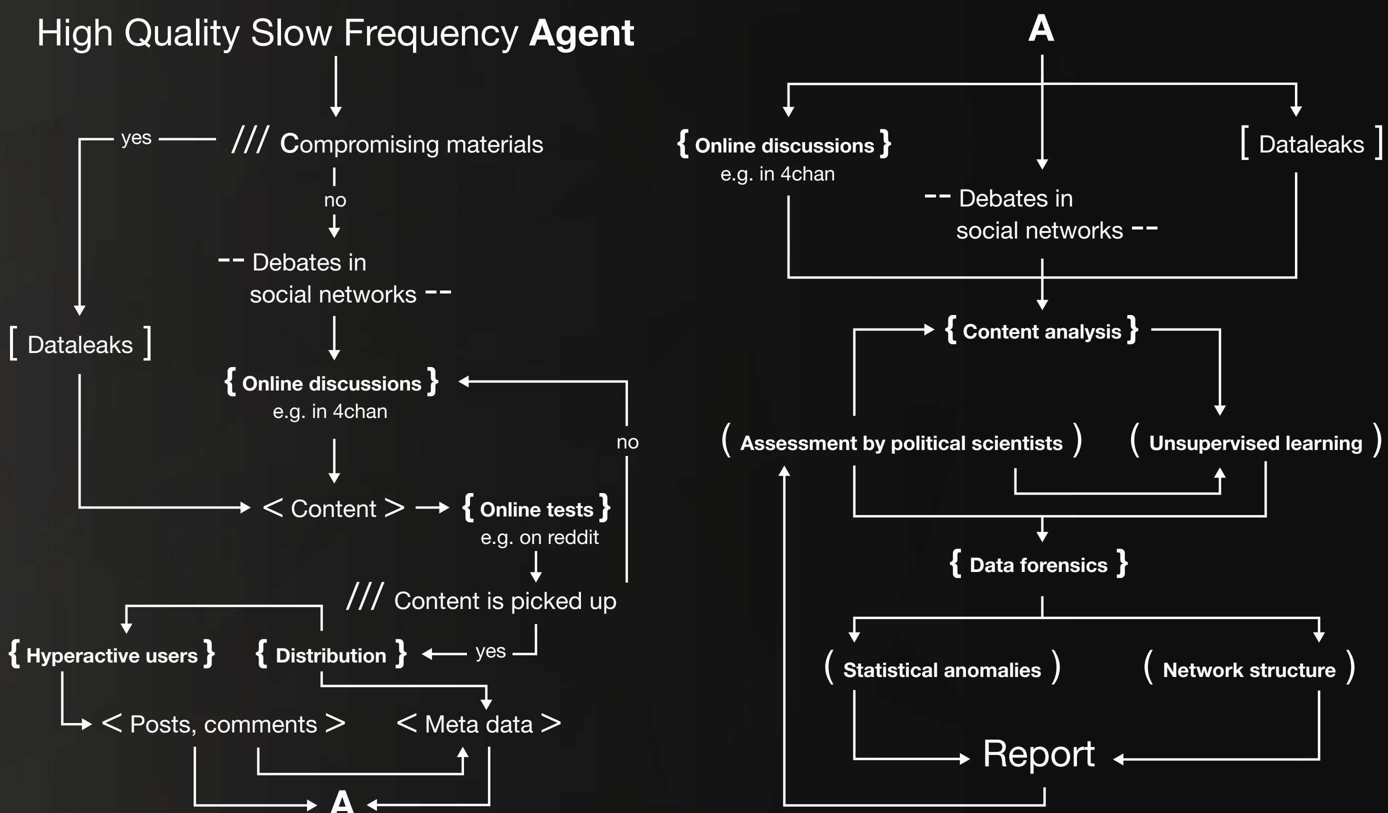

Left: Setting up conspiracy theories (High Quality Low Frequency Agents). A story to discredit the target is either sourced from data leaks or set up as fake news, tested in discussion platforms like reddit and then distributed via hyperactive users. As a result, a certain opinion forms within the network. **Right:** Such agents are difficult to uncover, because they cannot be detected with statistical tools early on. Political scientists watch sites which have been identified as at risk. If one suspects that a conspiracy theory is being launched, statistical tools help to assess how far it has already been distributed. A conspiracy theory can be made public and invalidated only if it is detected early enough.

kick in, ensuring that such an apparently popular topic is more and more visible to other users. And that achieves the aim of this type of purely statistical manipulation. According to Hegelich: "You look after the first 20,000 clicks yourself and hope it all takes off on its own from there. The topic then goes viral."

Depending on their development level, social bots have varying capabilities when it comes to mimicking human identities. Simple bots recognize key terms such as "refugee" and respond to them by posting images or retweeting comments. These simple bots are currently the most prevalent on the Internet. Generating them takes little programming skill, and manuals and instructions are freely available online. Hegelich himself has published one such guide on his blog. "It's about ten lines of code, and then you have a Twitter bot," he ▷

Graphics: ediundsepp (Source: Hegelich); Picture credit: ediundsepp confirms. More complex bots, meanwhile, can analyze the content of communications and engage in dialog. They cloak themselves with copied profile photos and follow a regular daily routine – just like the average human user.

Everything as a service

Social bots are also available to order. The business of creating fake digital identities is based on a value chain with globally distributed production. As Hegelich explains: "The slogan in Silicon Valley just now is 'everything as a service'. Sadly, on the Internet, that also extends to manipulation. You can buy 10,000 fake accounts for a few hundred dollars, for instance – and the software to run them can be purchased too. A highly sophisticated variant you could use to operate your 10,000 Twitter accounts costs around 500 dollars."

Not always political

Often, these purchases are not motivated by politics. If someone uses bots for promotional purposes, for instance, they automatically register them. And for each new registration, Twitter suggests users to follow. To make their human persona as believable as possible, the bots accept all these suggestions and thus end up following Donald Trump or Hillary Clinton, for instance. Hegelich is sure that many of the fake accounts attributed to the candidates have no political purpose at all: "And of course that distorts the picture. We need only recall the noise made by the media here about how many followers Trump has on Twitter. Yet these numbers tell us absolutely nothing!"

How do you spot a social bot?

The team's own initial studies have already yielded large data sets, which Hegelich is using for comparison purposes to establish how social bots behave. A computer program imitating human behavior will always generate recognizable patterns. The TUM analyses do show that bots have adopted a regular daily routine – they no longer post at night or every ten minutes, which is far too easy to detect, but now go to sleep and take lunch breaks. Over a longer time, though, it becomes evident that they are just as active over the weekend as during the week, for instance – unlike the average human user, who posts a lot less then. Patterns like this can be identified through data mining. As Hegelich describes: "We put all the data into the computer and say, 'These are bots and these are not, so now tell us how they differ'."

Uncertainty – the most dangerous type of manipulation

Another factor that contributes to uncertainty is the fact that rising social bot numbers increasingly cloud the topics and views of importance to voters. It remains unclear, for instance, whether people are generally expressing negative comments about refugees more frequently or whether this trend stems from a computer program. The US elections showed very clearly that the most dangerous type of manipulation deliberately sets out to create uncertainty. Hegelich clarifies: "What doesn't work is political conversion. I can't use social media to suddenly make a Democrat out of a Republican. But uncertainty works very well indeed. If I want to canvass for Trump, I don't need to promote him overtly. If I can manage to spread the impression that they're all lying, for instance, that's something that helps Trump more than Clinton. In Germany, that would help the right-wing AfD more than the mainstream SPD and CDU parties."

Is the clock ticking for the Internet as a human communication platform?

Taking social bots as an example, we see how digitalization invalidates the age-old assumption that ultimately, quantity is an indication of quality. This is now no longer the case, since even a message shared millions of times can be downright false. Growing lack of trust among users could be dangerous for a network like Twitter, since they would then move away from the platform in the end. As it stands, Twitter is currently reluctant to take serious action against bots for commercial reasons. In the main, however, Hegelich takes a fairly relaxed or even optimistic view of our digital future. As he sees it, everything is in flux and new technical solutions will continue to emerge – new networks and new rules leading to new ▷

Who is behind the bots?

Experts believe that social bots are now part of every political debate. The majority of bots deployed can currently prove detrimental in two ways: First, they are scalable – if you can run one, you can run a million, so the sheer volume can skew trends and divert the attention of large numbers of users to a given topic. Second, bots can contribute to polarizing opposing camps with hate messages.

With few exceptions, the originators of social bots have not yet been identified. The spectrum ranges from dubious PR agencies right through to organized cybercriminals. Many of those involved come from Eastern Europe, though are not necessarily – as often supposed – linked to the Russian government. Cybercrime is widespread in Eastern Europe, where there are many well-trained people who know they can make money in this way and that it is extremely difficult to bring legal action against them from abroad. Russia, for instance, would not extradite anyone to the US. However, the first large-scale programs for social network manipulation were developed by US intelligence services.

Graphics: ediundsepp (Source: Hegelich)

How to build a little spambot: Simple spambots can be programmed quite easily. Hegelich presented this very short script for a Twitter spambot (written in the open source programming language R) in his blog.

```
library(twitteR)

ckey <- “1234mykey”
csecret <- “1234mysecret”
atoken <- “45678mytoken”
asecret <- “56789othersecret”

setup_twitter_oauth(ckey, csecret, atoken, asecret)

LISTTopic <- twListToDF(searchTwitter(‘#BigData’, n=10))
View(LISTTopic)
LISTNames <- unique(LISTTopic$screenName)
text.examples <- c(“I am a bot, but I appreciate your work!”,
                  “Data is the new bacon!”,
                  “There are only 10 kinds of people: Those understanding binary code and others.”,
                  “Data is like people – interrogate it hard enough and it will tell you what you want to hear.”,
                  “Data that is loved tends to survive.”)
for(i in 1:length(LISTNames)){
  message.text <- paste0(“Hi @”,LISTNames[i], “ “,
                         text.examples[sample(length(text.examples),1)])
 print(nchar(message.text))
   try({
     updateStatus(message.text)
   }) -> temp
   if(class(temp)==”try-error”){
     print(‘Error!’)
     Sys.sleep(runif(1,50,100))
  }  else{
     print(paste0(‘i= ‘,i,’ (‘,LISTNames[i],’) is DONE!’))
     Sys.sleep(runif(1,10,22))
 }
}
```

Load library for Twitter communication

Get authentification for Twitter application programme interface (API)

Search for 10 latest tweets about big data

Use this list of text samples about big data

Answer each of the 10 tweets by writing: Hi, “your name”, the user’s name, “one of the text samples

Go to sleep on a regular basis to keep twitter from identifying you as a bot

user behaviors. For him, the more urgent question is whether people get a handle on the current reality fast enough in relation to Germany's 2017 elections: "I am actually quite concerned about that. At the moment, we are dealing with social media structures that were absolutely not set up for political opinion-forming, but purely for business reasons. Facebook was intended to be a virtual friendship network – a feel-good environment for private users – and not an information medium. Large-scale, manipulative use of this kind of network by government organizations, for instance, was not part of the plan."

Responsibility of political parties
Since political opinion-forming is increasingly taking place online, Hegelich believes political parties have a responsibility to actively engage in these debates. However, he points out that, "Political campaigning on social medial is a gray area ethically speaking. There's a general lack of experience as to what conditions should apply here. Social bots concealing their presence are obviously not acceptable. But what about chatbots, which hold automated discussions with users about the party manifesto? And what about personalized campaign ads? These issues certainly call for a particularly transparent approach."

In Germany, politics is still blind – half a year before the elections, it remains oblivious to what is actually happening on social media. Hegelich thus advocates setting up a monitoring system. Especially during an election campaign, political players ought to know what is circulating publicly on social networks – especially when it comes to fake news. This is essential to have any chance of responding. A monitoring system could ensure early detection of mass-messaging with dubious content and inform the politicians and parties affected. "That, too, relies on machines," concludes Hegelich. "Monitoring of this type requires extensive automation – noone can read the whole of Facebook and analyze the content from a political perspective." *Karsten Werth*

"I actually think bots are a great technology that can be put to very good use. I could even envisage them being a valuable addition to a political campaign. But the moment a computer program tricks you into thinking it's a real person, it then becomes highly problematic."

Simon Hegelich

Prof. Simon Hegelich

A pioneer in political data science

Simon Hegelich has been Professor of Political Data Science at the Bavarian School of Public Policy since 2016. This professorship is the first of its kind in Germany. In his research, Hegelich blends political and computer science to pursue political data science. This entails both investigating the political relevance of issues surrounding digitalization and applying methods such as machine learning, data mining, computer vision and simulation to conventional aspects of political science. Hegelich studied political science at the University of Münster, completing both his doctoral and postdoctoral theses there. From 2011 through 2016, he was Director of the FoKoS interdisciplinary research center at the University of Siegen. Hegelich has been nominated for the 2017 German Research Foundation (DFG) Communicator Award, which recognizes excellence in communicating research findings to the public.

Picture credit: ediundsepp

Algorithms Help Control **Social Dynamics**

Prof. Massimo Fornasier does not need a crystal ball to see into the future. He uses the power of mathematics instead. At his Chair of Applied Numerical Analysis at TUM, he develops processes and models that have been proven to deliver reliable forecasting results.

Link
www-m15.ma.tum.de

Birgit Fenzel

Algorithmen zur sozialen Kontrolle

An seinem Lehrstuhl für angewandte numerische Analyse in Garching arbeitet Prof. Massimo Fornasier mit seinem Team an der mathematischen Modellierung hochdimensionaler Probleme – also einem Wissensgebiet, das sich mit der mathematischen Beschreibung von Problemen befasst, bei deren Darstellung man mit der Berechnung von Länge, Breite, Höhe und vielleicht noch der Zeit nicht sehr weit kommt. Wie er unlängst beweisen konnte, sind soziale Dynamiken auch „bloß“ ein hochdimensionales Problem. So hat er Algorithmen entwickelt, mit denen er das Verhalten von Menschen in einer Gruppe sehr präzise berechnen, prognostizieren und sogar beeinflussen kann.

Die Lösung liegt dabei in der Kunst der Reduktion und in der Betrachtung von Menschen als physikalische Teilchen. „Als Masse verhalten sich Menschen durchaus ähnlich wie Teilchen von Flüssigkeiten oder Gasen“, erklärt er seinen Ansatz. Um die Bewegungsrichtung großer Mengen von Gasmolekülen mit hoher Wahrscheinlichkeit berechnen zu können, müsse man nicht die genauen Eigenschaften jedes einzelnen Teilchens kennen. „Es genügt, ihre durchschnittlichen Eigenschaften zu kennen“, sagt Fornasier.

Mithilfe von Computersimulationen konnte er damit bereits mögliche kollektive Verhaltensmuster einer großen Anzahl sich wechselseitig beeinflussender Individuen in einer aktuellen Situation darstellen. Dass diese virtuellen Multiagentensysteme durchaus mit realem Gruppenverhalten übereinstimmen, bestätigte ein Experiment, das die Mathematiker der TUM im Mai 2015 in Zusammenarbeit mit dem Consiglio Nazionale delle Ricerche (CNR) und der Universität „La Sapienza” in Rom durchführten. Dabei hatten die Forscher zwei Gruppen von etwa 40 Studierenden damit beauftragt, ein bestimmtes Ziel in einem Gebäude zu suchen. In die eine Gruppe hatten sie allerdings zwei „Undercoveragenten“ eingeschleust, die zielstrebig in eine vorgegebene Richtung gingen. Damit allein brachten sie die Gruppe dazu, ihnen wie Schafe dem Leithammel zu folgen. Wie sie auch beweisen konnten, gelten diese Ergebnisse ebenfalls für sehr große Gruppen. „Tatsächlich reichen zwei bis drei solcher Agenten pro hundert Menschen aus“, sagt Fornasier.

Mit seinen Algorithmen für soziale Dynamiken lassen sich effiziente Möglichkeiten finden, die für mehr Planbarkeit und Sicherheit sorgen. Zum Beispiel, wenn es darum geht, große Besuchermengen durch Gebäude zu dirigieren, Verkehrsflüsse zu optimieren oder Evakuierungen optimal durchzuführen. Bei allen denkbar positiven Anwendungsmöglichkeiten stellt sich in diesem Zusammenhang natürlich auch die Frage nach den Risiken. Da seine mathematischen Modelle in einer völlig abstrakten Umgebung formuliert sind, können sie leicht an verschiedene Situationen angepasst werden, in denen Interaktionen von Menschen eine Rolle spielen. Ganz gleich, ob es sich nun um Finanzmärkte oder Meinungsbildungsprozesse handelt. So ist die Möglichkeit der Einflussnahme auf das Verhalten sozialer Systeme ein zweischneidiges Schwert. „Die gute Nachricht in diesem Zusammenhang ist, dass wir auch bewiesen haben, dass das Verhalten nicht für alle Typen von Dynamik und nicht in allen Situationen so leicht vorauszusagen ist“, kann Fornasier beruhigen. □

Graphics: ediundsepp, Picture credit: Jooss

“Performing numerical simulations based on the complex data from social behavior is one of the key challenges of the 21st century.”

Massimo Fornasier

The world of big data is opening up new opportunities in the realm of forecasting. In theory, at least. Even though rapid advances in technology have led to an explosion in the availability of data in recent years, this has not led to a corresponding gain in knowledge. This is particularly true in the case of models for forecasting social dynamics. The reason lies in the fact that human behavior is influenced by a complex interplay between physical, emotional, cognitive and social factors. The rise of Twitter and other social networks means that data is no longer just two-dimensional – it cannot be described as easily as something like an ultrasound, for example. The evaluation of complex data often results in very large systems of equations and the computing power required to find a solution is rising exponentially in tandem. As such, even state-of-the-art supercomputers can in many cases struggle to meaningfully process the raw data. “We call this phenomenon the curse of dimensionality,” explains mathematician Massimo Fornasier. Along with his team in Garching, he uses highly efficient algorithms to solve such high-dimensional problems. “Performing numerical simulations based on this data is one of the key challenges of the 21st century,” he maintains.

His models are based on the idea that the dynamics of a multi-layered overall system are governed by just a few classification parameters. “Simply put, we want to be able to understand and organize complexity,” says Fornasier as he describes his research. His algorithms are essentially nothing more than programmed sieves that sort through a huge volume of data to find the information that corresponds to a certain

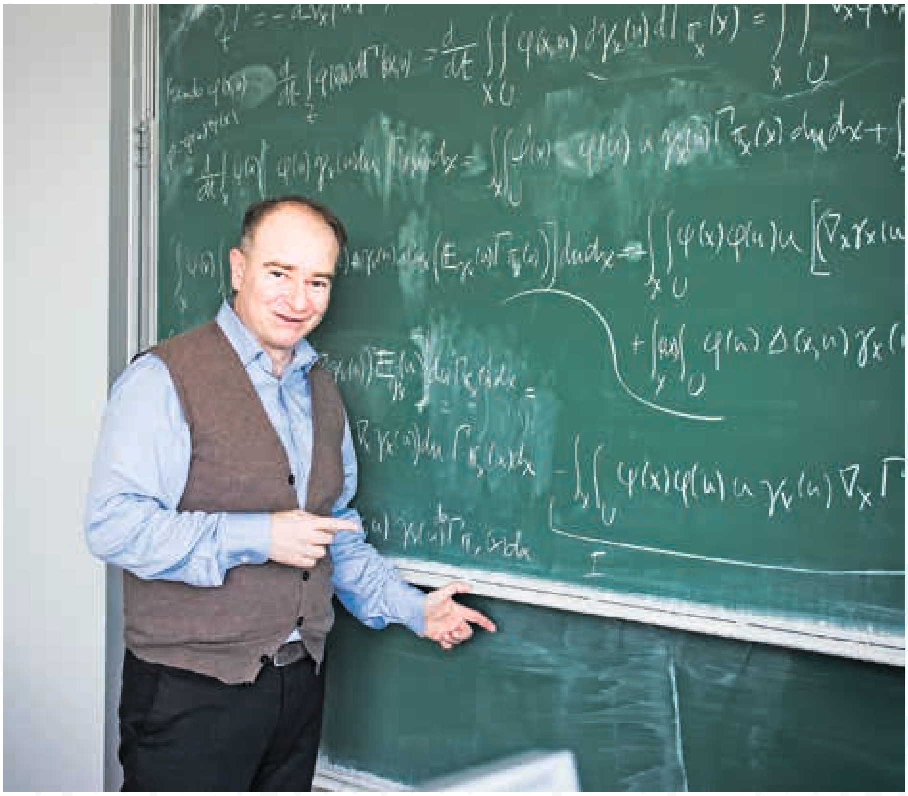

Using mathematics to describe how a person – a so-called agent – behaves within a crowd: Agent's movements are highly probabilistic because they depend, for instance, on obstacles or on how neighboring agents may behave. Mathematically, this is represented by a probability distribution of possible moves for each agent.
The blackboard shows a system of two time dependent equations, which Fornasier uses to describe the behavior of several agents within a crowd and the evolution of the respective probability distributions of their moves.

pattern. "We reduce the parameters to an absolute minimum to then draw conclusions applicable to a larger quantity." The process can be described as a mathematical balancing act consisting of many small arithmetical adjustment steps. The mathematical model that eventually emerges from these iterations has to be as simple as possible in order to be manageable, but should at the same time be as complex as necessary. The next step is a question of simulating the system as such. "This means that we have to develop algorithms that allow us to simulate the process with sufficient accuracy."

Together with his working group in Garching, he has already developed algorithms for a series of highly complex and thus higher-dimensional problems from real life, which produce such reduced or – as the researchers call them – "sparsely populated" results. Applications for his mathematics include marketing strategies, financial portfolio management, image compression and painting restoration. One example is the restoration of an Andrea Mantegna fresco in a church in Padua, which was almost totally destroyed by bombing in 1944. All that remained of the almost 1,000 m^2 artwork dating from the 15th century were tens of thousands of painted puzzle pieces rescued from the piles of rubble after the attack. These fragments, which made up less than one tenth of the original fresco, were cleaned, sorted and photographed in the early 1990s. Using the digital photos, Fornasier was able to develop an algorithm to automatically assign each fragment to its original position in the fresco.

Impressive as this is, Fornasier knows that his mathematical models are capable of so much more. He recently proved in a European Research Council project with EUR 1.123 million in funding that his process can also very accurately analyze, forecast – and even under certain conditions influence – social dynamics.

Cloud-like patterns of movement

We already know that highly efficient algorithms are used in direct marketing so that companies can identify customers' preferences based on their Internet usage and use these insights to recommend further products. "The programs running in the background of the websites of market leaders like Amazon and Netflix are getting better all the time at customizing the content presented to users," says Fornasier. Their forecasting abilities are limited, however, by the myriad factors that influence individual behavior. "But things look quite different if people are considered as part of a crowd," Fornasier points out. He and his team have recently proven mathematical laws that demonstrate how surprisingly easy it is to automatically generate precise mathematical models for specific, relatively simple group interactions based on observed dynamics data.

He has observed striking structural similarities with phenomena from the world of physics. This is especially true in situations where there is limited room for maneuver because of environmental constraints and rules – for example in road traffic or in social networks. "In those cases, masses of people behave like gas clouds," he explains. These consist of many atoms, which have their own mass, have a position and a speed at all times, and exert forces on each other through ▷

Graphics: ediundsepp, Picture credit: Jooss

"Once we can calculate the behavior of a group of interacting agents in advance, we are only one small step away from being able to steer them."

Massimo Fornasier

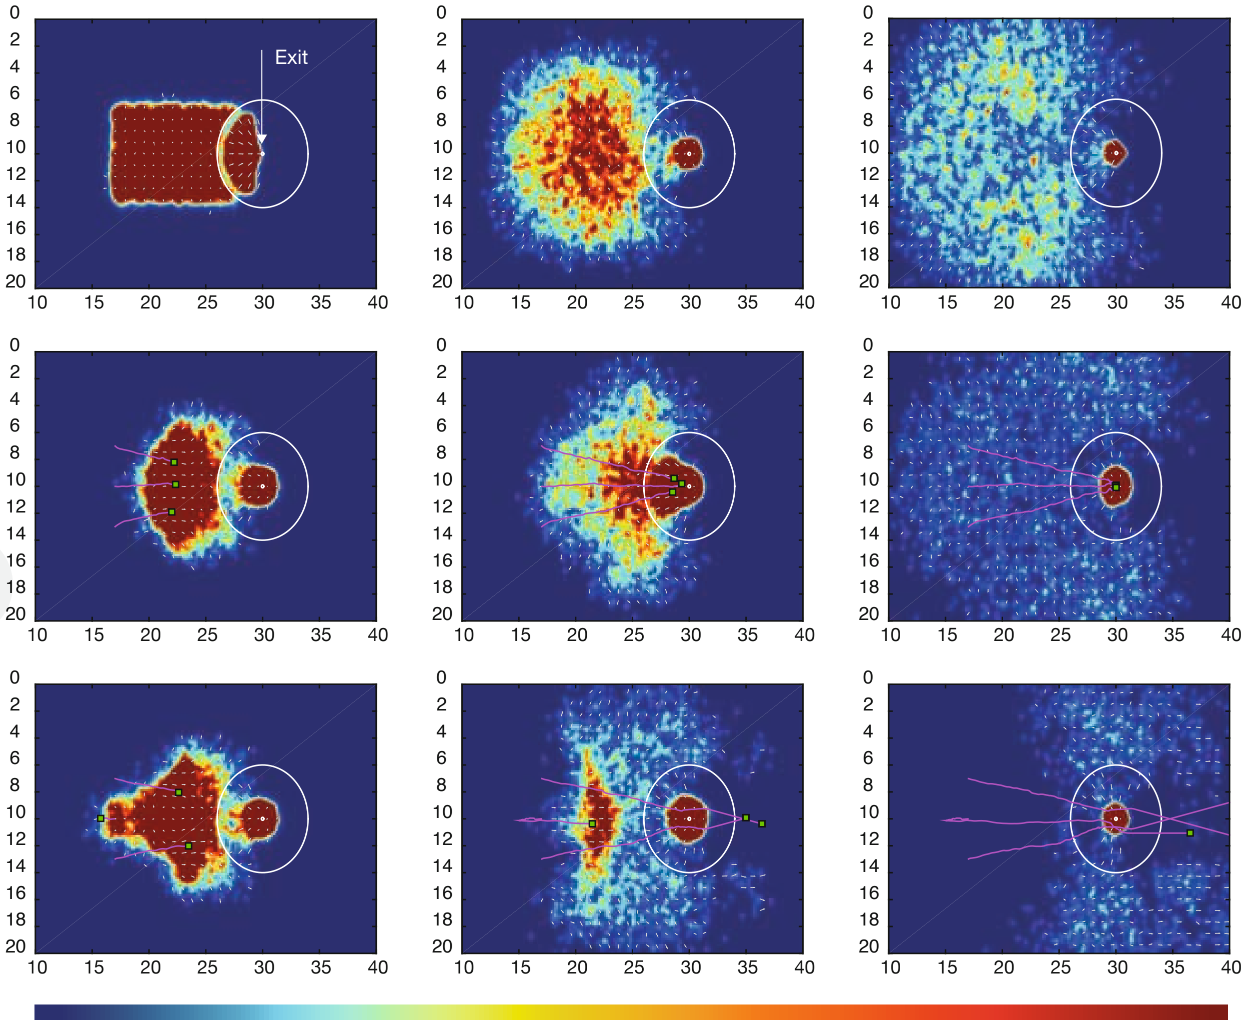

To simulate the behavior of very large human crowds comprising millions, Fornasier calculates the probability density rather than individual trajectories for each person. The image shows the simulation of the behavior of a human crowd in a room with no obstacles and one exit, expressed as probability of a person being in a certain location. Dark red equals high probability, dark blue equals zero probability. The white arrows indicate the locally averaged direction and speed of the agents.
First row: Under normal circumstances, the persons who can see the exit move towards it (left). Some group mates close to them follow (middle). The persons farthest away split into several groups moving into random directions and never reach the exit (right).
Second row: Three informed leaders (green, with pink trajectories) move straight towards the exit (left). All others follow, causing heavy congestion around the exit, which delays the evacuation (middle). Eventually the probability of finding a person still in the room is very low (right).
Third row: Three informed leaders move along trajectories which have been calculated with the aim of optimizing the flow (left). Congestion is avoided and flow through the exit is increased (middle). The probability of a person not reaching the exit is nearly zero (right).
Graphics with the kind agreement of the authors: Picture credit: G. Albi, M. Bongini, E. Cristiani and D. Kalise. Invisible Control of Self-Organizing Agents Leaving Unknown Environments, SIAM J. Appl. Math., 76(4):1683-1710, 2016. G. Albi and M. Bongini have been collaborators within Fornasier's ERC Starting Grant project "HDSPCONTR".

their electric charges. The important calculation criterion is not the position or speed of the individual particles – or individuals – at any one time, but rather how they behave as a mass on average.

In this context, Fornasier likes to refer to the science fiction writer Isaac Asimov. The Russian had invented the science of "psychohistory" in his celebrated Foundation Series. In the novels, the scientist Hari Seldon – who was also a mathematician – used this analogy of gas kinetics and collective behavior to predict the future of mankind for millennia. Fornasier's models bring this science fiction surprisingly close to scientific reality. Incidentally, the mathematician and his team have also elaborated on a new application for the Boltzmann equation, which physicists commonly use to calculate the statistical distribution of particles in gases. In particular, the scientists' capabilities are not just limited to describing collective behavioral patterns and visualizing them in computer simulations. They have also developed tools to help them see into the future. "Once we can calculate the behavior of a group of interacting agents in advance, we are only one small step away from being able to steer them," says Fornasier.

He found empirical evidence for this in an experiment undertaken with colleagues in May 2015 in cooperation with the Consiglio Nazionale delle Ricerche (CNR) and the University of Rome "La Sapienza" in Italy. As part of this project, the researchers divided 80 students into two groups and gave them the task of finding a specific location in a building. The scientists had planted two "undercover agents" into one of the groups and instructed them to walk very determinedly in a pre-defined direction. It turned out that their determined behavior very quickly induced the entire group to start following them. The mathematicians were also able to prove in this experiment that the movement of large groups of people can be steered with surprisingly little effort. "In fact, two to three agents per 100 individuals are sufficient," says Fornasier.

Herding dog strategy for opinion forming

The fact that his mathematical models are formulated in an entirely abstract framework makes them easily adaptable to a wide variety of situations in which interactions between multiple agents play a role. Fornasier takes the financial markets as an example of an arena where influence is exerted by a small number of agents. "There, precisely coordinated activities by big investors can result in sizable market movements," he explains. Maybe even more interesting is what Fornasier has learned about the dynamics of opinion forming within a group. His model shows that it is relatively easy to steer the thinking of individuals who are part of a group in a certain direction. "There is also a good model for this in nature," according to Fornasier. "To drive a herd of sheep in a desired direction, a good herding dog will always concentrate on the animal that is farthest from the group. They achieve their goal by reining in the most stubborn animal." When it comes to opinion forming in groups, it is also simply a matter of reining in the most ardent supporters of the views that go against the desired consensus. The rest of the group will then follow.

His findings are based on both theoretical results and numerical simulations, which are in turn based on opinion-forming models like those developed by the philosopher Rainer Hegselmann and the mathematician Ulrich Krause. What these say about opinion formation is that every group member adapts their opinion in line with those of the rest of the group in increments of time.

Fornasier and his colleagues are able to describe this alignment process using ordinary and partial differential equations. "Since all group members have their own opinion at the beginning, it is important that all agents share their information and points of view," he says of the principle. The consensus building happens in this process. This can only happen, however, if there is not too much separation. Agents who are ▷

Graphics: ediundsepp (Source: TUM)

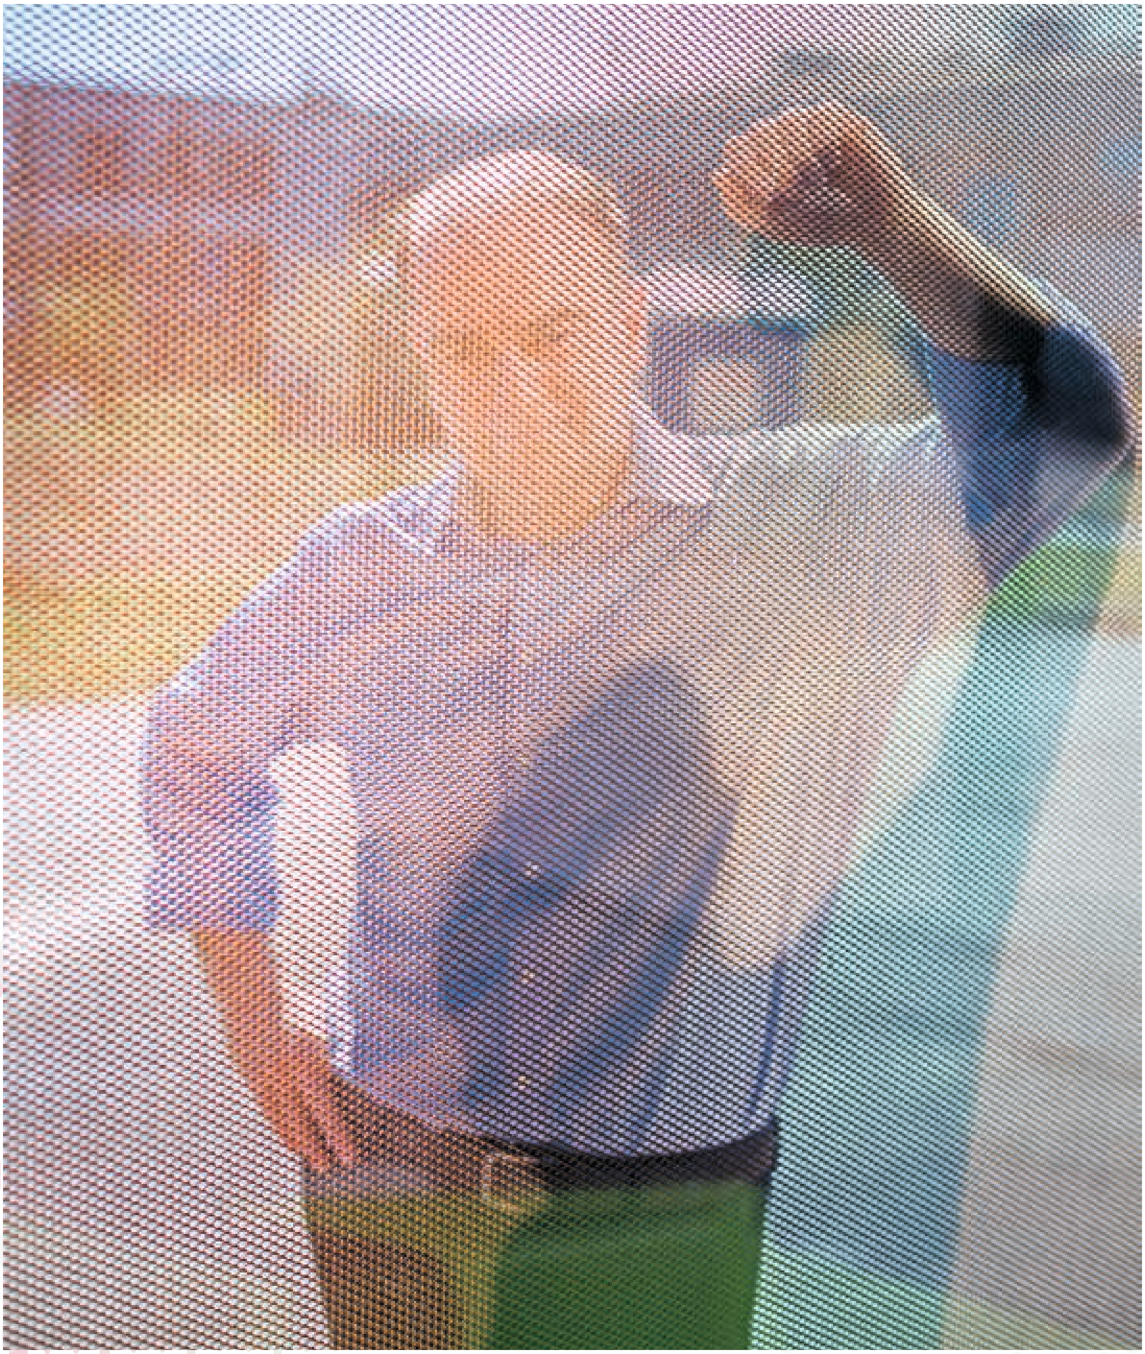

Prof. Massimo Fornasier

A multi-dimensional researcher

With his mathematical models, Massimo Fornasier has a view on the world that is much more advanced than our three-dimensional observations. His special interest is the modeling of high-dimensional problems, where other dimensions in addition to length, height, width or time can be involved and which are therefore more difficult to conceive. The mathematician is able to present these problems using the formulas he has devised.
Fornasier was born in 1975 in the town of Feltre in the North-East of Italy. He has held the Chair of Applied Numerical Analysis at TUM since April 2011. His academic journey began at the University of Padua where he studied mathematics and completed his Doctorate in Computational Mathematics in February 2003. Between 2003 and 2009, he held the position of research associate at the Universities of Vienna, Rome, Marburg, Princeton and Linz. He has already received several awards for outstanding achievements in his field. These include the START Prize – the most prominent and highest endowed award for young scientists in Austria – and a Heisenberg professorship from the German Research Foundation. In 2009, he received the Prix de Boelpaepe for image processing from the Royal Academy of Science, Letters and Fine Arts of Belgium. In 2012, he became the first ever recipient of the newly created SIMAI award of the Italian Society for Applied and Industrial Mathematics. He also received a prestigious Starting Grant from the European Research Council in 2012. Prior to his appointment at TUM, Fornasier worked as a senior scientist at the Johann Radon Institute for Computational and Applied Mathematics (RICAM), which is part of the Austrian Academy of Sciences.

Graphics: ediundsepp, Picture credit: Jooss

When there is limited room for maneuver because of environmental constraints and rules, masses of people behave like gas clouds.

far apart from other individuals of the group do not exchange their views, which means that their position remains fixed. There is a way to bring these hardliners into the fold, however, according to the mathematician. “In the model, all we need to do is create one virtual agent and position them close to the extremists,” he explains. During the rapprochement process, the dissenter is gradually guided towards the opinion, which was literally brought nearer to him.

Fornasier’s social dynamics algorithms can help to find efficient solutions for planning and security issues in the real world. Examples of situations where this can come in useful include directing large numbers of visitors through buildings, optimizing traffic flows, and organizing evacuations in an emergency. For all the conceivably positive application scenarios that exist, there remains, of course, the question of potential risks. The possibility of exerting influence on social behavior is a double-edged sword. There is always the danger that the power of information will be abused through targeted manipulation and that democracy will thereby be undermined. Fornasier, however, excludes the possibility of his mathematical models being exploited to create totalitarian structures. “The good news in this context is that we have also proven that behavior is not so easy to predict or control for all kinds of dynamics and situations,” he says reassuringly. “We are lucky to have plurality of opinion and safeguards like the media or a multi-party system.” The scientists also found that the processes only work well in groups that show generalized patterns of behavior. However, as soon as the energy of individual agents crosses a certain threshold, the group of agents can no longer be moved in a concerted direction using simple, sporadic interventions. In other words, just as the best herding dog will fail to drive its herd of sheep home if the outliers dart away too quickly, algorithms too will become ineffectual in such scenarios. “The prospect of complete control will thus remain science fiction,” he says with conviction.

Birgit Fenzel

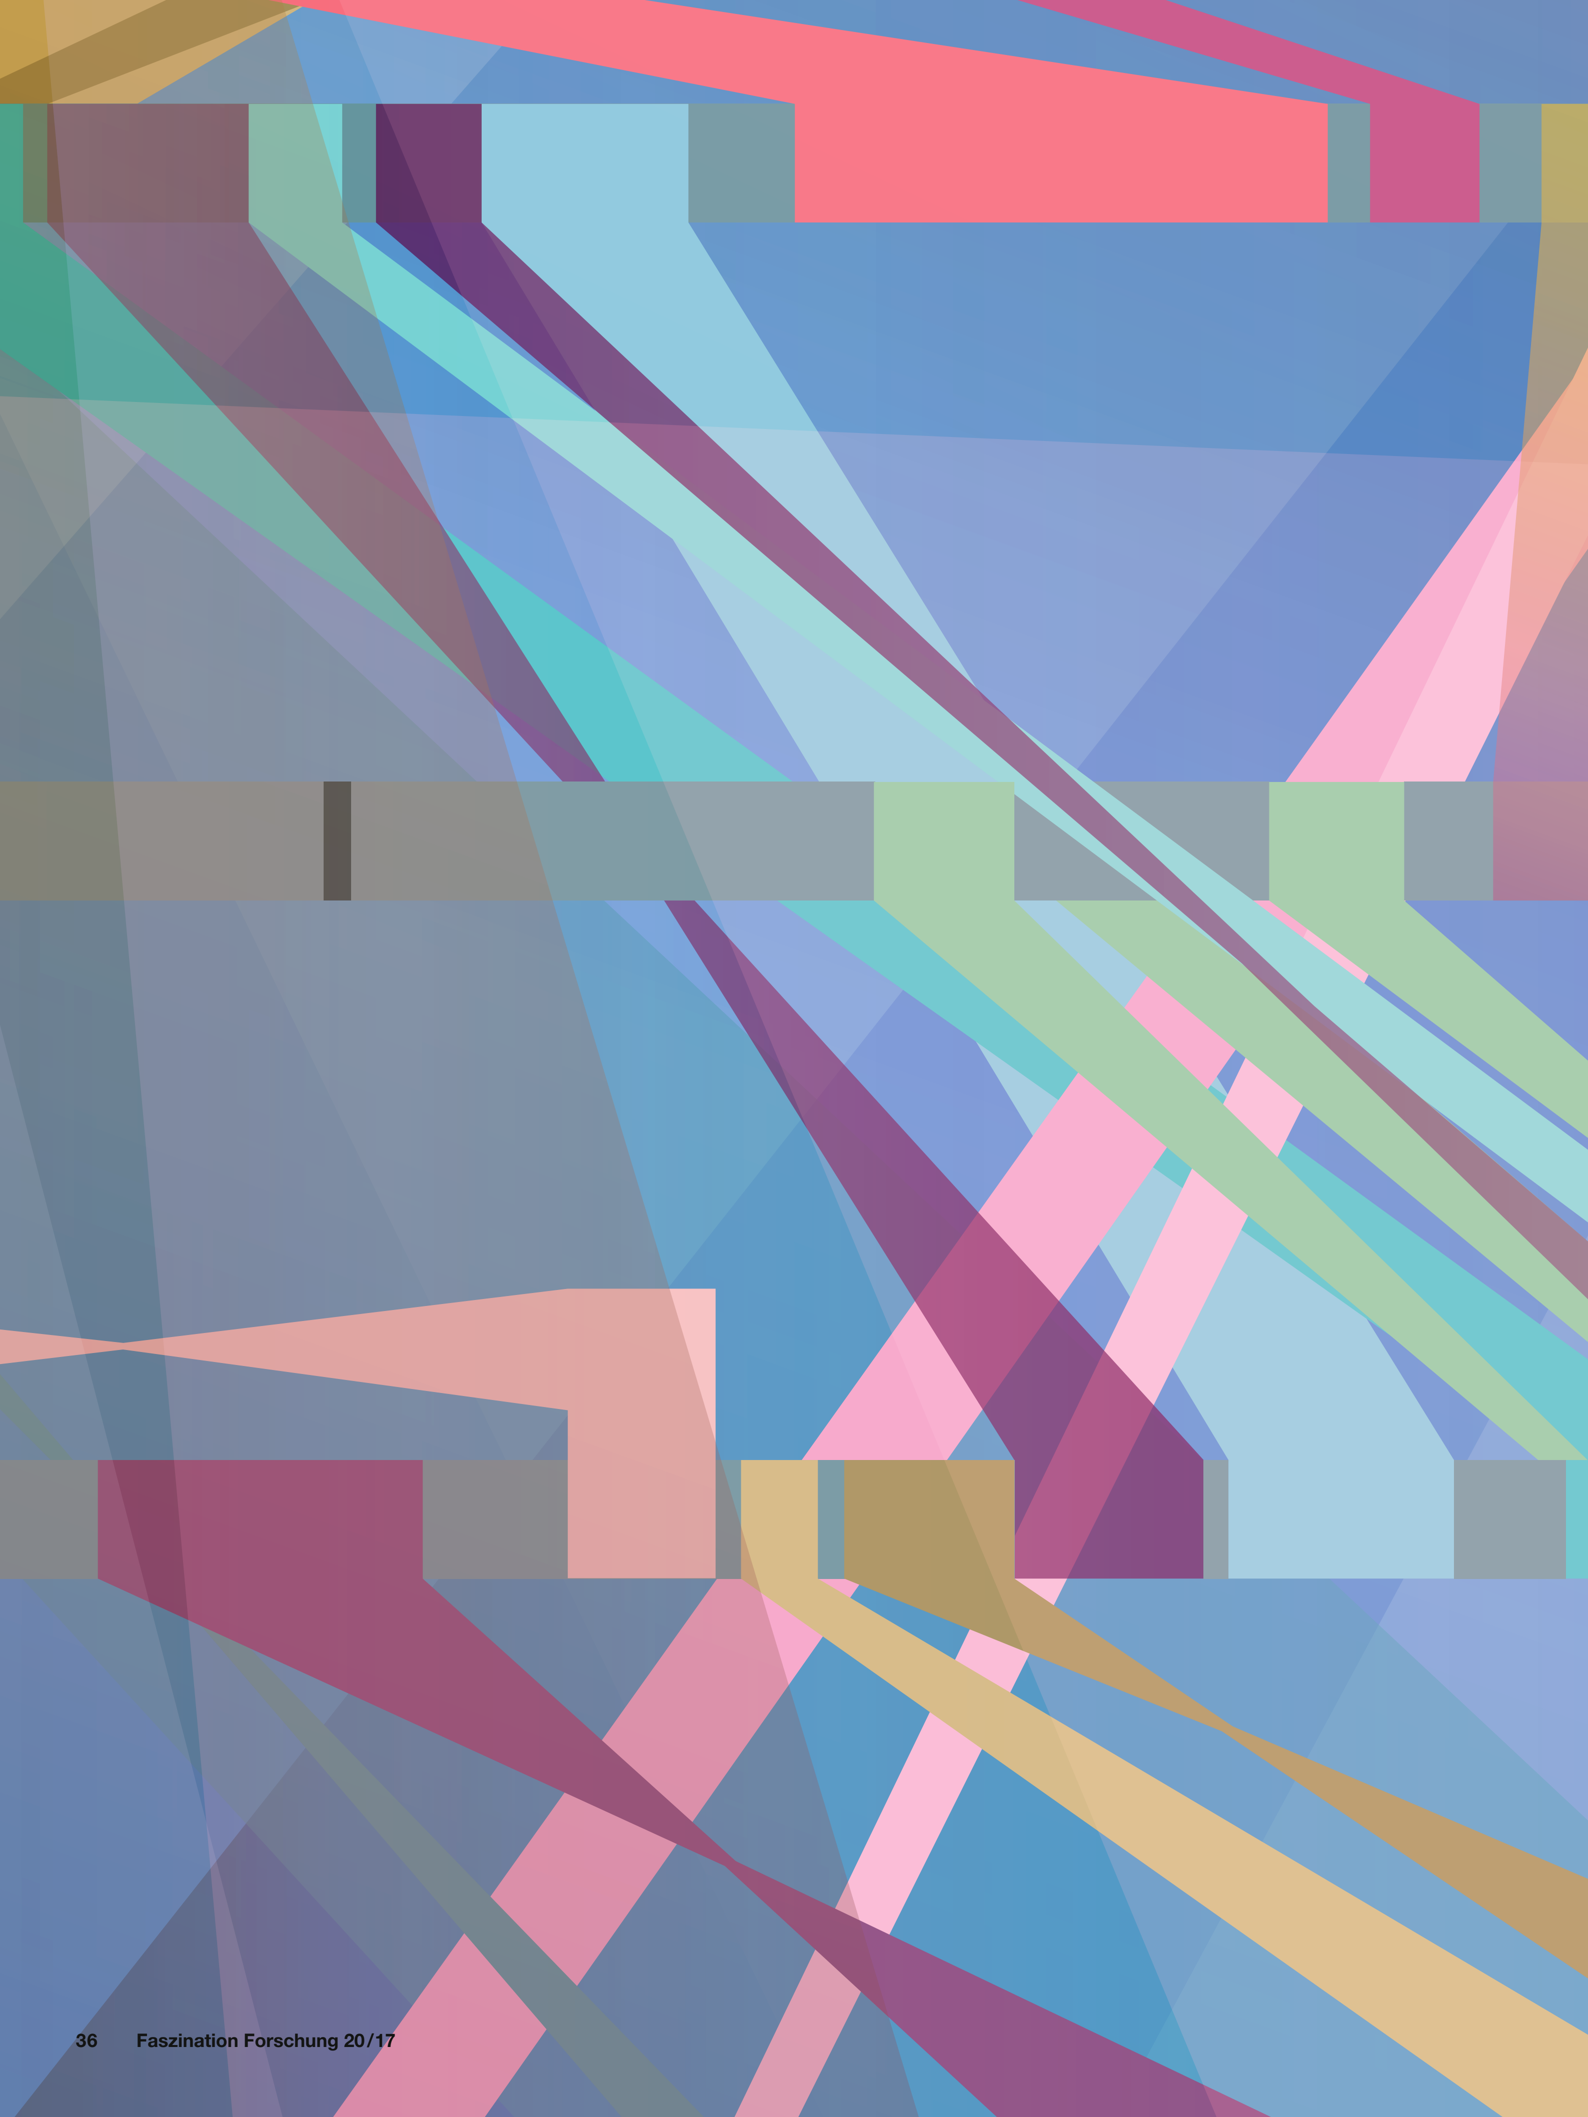

The Secret of Bioinformatics – **Classification and Prediction**

By the time Prof. Hans-Werner Mewes retired at the end of March 2017 he had witnessed and indeed influenced more than forty years of development in bioinformatics. At the end of the 1970s, he was among the pioneers who first digitally captured biological data. In parallel with computer power, the volume of information generated by biologists and medics also grew – and with it, the need to sort and interpret this data. Here, the researcher looks back at the early years of this new discipline.

Link
www.bioinformatik.wzw.tum.de

Graphics: ediundsepp

Brigitte Röthlein

„Die Kunst der Bioinformatik besteht aus zwei Dingen: der Klassifikation und der Vorhersage“

In den Lebenswissenschaften ist heute Forschung ohne Datenverarbeitung nicht mehr möglich. Durch die immer schnellere und preisgünstigere Sequenzierung von Genen bzw. Proteinen gab es in den letzten Jahren eine Explosion von Daten, deren positive Folgen im Erkenntnisgewinn in der Praxis und Klinik erst jetzt allmählich sichtbar werden. Um sie auszuschöpfen, benötigt man Techniken der Künstlichen Intelligenz, etwa maschinelles Lernen oder künstliche neuronale Netze. So ist ein neues Fachgebiet entstanden, die Bioinformatik.

Prof. Hans-Werner Mewes, der zu den Pionieren in dieser Disziplin gehört, setzt sich seit Jahren mit der Umwandlung von biologischem Wissen in berechenbares auseinander. Er leitete bereits 1989 beim Hefegenomprojekt, einem der ersten Projekte zur Entschlüsselung des kompletten Genoms eines Organismus, das Team, das die bioinformatische Aufbereitung der Daten von den insgesamt rund 600 beteiligten Wissenschaftlern aus über 100 Laboren übernahm. Zuerst mussten dort die analysierten Fragmente anhand überlappender Abschnitte korrekt zu einem durchgehenden DNA-Strang zusammengefügt werden. Danach stellten die Forscher die Sequenzdaten und die zugehörigen Informationen systematisch in computerlesbarer Form dar und organisierten sie als Datenbank. Sie entwickelten damals zur Darstellung des Hefegenoms eine Oberfläche, die eine einfache symbolische Visualisierung der Chromosomen erlaubte. So wurde es durch vielfältige Verknüpfungen von Datenelementen möglich, mit Browsern durch das Erbgut zu navigieren.

Die bioinformatischen Verfahren wurden im Laufe der Jahre immer ausgefeilter und spezifischer auf die Probleme zugeschnitten. Man erstellte Funktionskataloge, das heißt, man versuchte zu klassifizieren, welche Aufgaben welcher DNA-Abschnitt jeweils übernimmt. Denn die Kunst der Bioinformatik besteht aus zwei Dingen: der Klassifikation und der Vorhersage. Mit der zuverlässigen Vorhersage von biologischen Eigenschaften aus Daten werden viele Experimente überflüssig, andere können so erst interpretiert werden.

Um die Übertragung akademischer Lösungen in die industrielle Praxis voranzubringen, gründete Mewes zusammen mit Kollegen 1997 die Firma Biomax, die solche Lösungen anbietet. Seit 2002 gibt es in München das Studienfach Bioinformatik in einer Zusammenarbeit zwischen TUM und LMU. Seither wurden schon rund 600 Bioinformatiker ausgebildet, viele haben promoviert, einige sind Professor(innen) geworden, viele arbeiten im Ausland und in der Industrie. □

Prof. Mewes, you were among the very first bioinformaticians. Could you tell us how it all began?
Well, the field of bioinformatics is really still in its infancy. It is directly linked to the massive data streams produced by biology today – our role being to sort, merge and interpret it. Personally, I began working in this area in 1979 at Heidelberg University. There we had what was – at the time – a very expensive computer with 10 kilobytes of storage capacity, and I programmed that to digitally record measurement data. That was a very exciting step. Fortunately, we were dealing with very modest volumes of data back then.

But then all of that changed?
Yes, very quickly indeed. The core challenge facing bioinformatics today is the sheer volume and complexity of available information. Over the last thirty years, sequencing technologies have become more and more advanced, and this has created huge streams of data. The first step here is to determine the order of the four bases that form the building blocks of life. We represent them as strings consisting of four letters: G, A, T and C, designating the bases guanine, adenine, thymine and cytosine.
The process began in the 1950s with protein sequencing – then a matter of a few hundred amino acids – and continued in the seventies and eighties with the sequencing of DNA in bits and pieces, not to assemble whole chromosomes. By then, we were already dealing with millions of base pairs. We then selected certain model organisms, already extensively studied from a biological perspective – like yeast, for instance, whose 12 million base pairs encoding 5,800 genes were identified by 1996. Finally, at the end of June 2000, the first full sequence of a human genome was published, consisting of over 3 billion base pairs coding for around 20,000 genes. Today, huge amounts of data are generated by other methods too, such as proteome analysis. Fortunately, at the same time, computer power has also grown at an explosive rate – and every lab can now access the relevant databases over the Internet. ▷

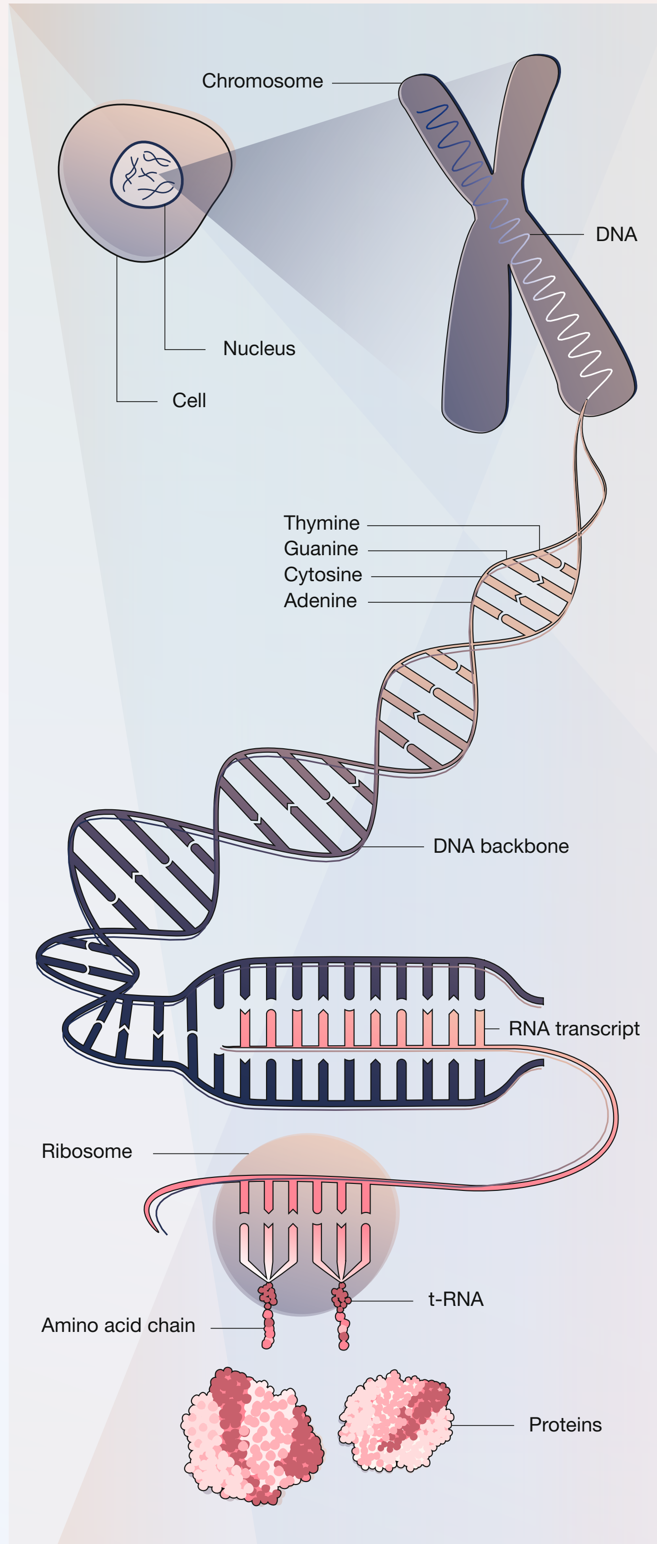

The amino acid sequence stored in the genes is the blueprint for every creature's proteins. To start protein production, DNA is unpacked from the chromosomal proteins and transcribed into messenger RNA. Ribosomes – the protein factories of the cell – link the individual amino acids together to form a protein following the blueprint of the messenger RNA.

Graphics: ediundsepp (source: TUM), Picture credit: Jooss

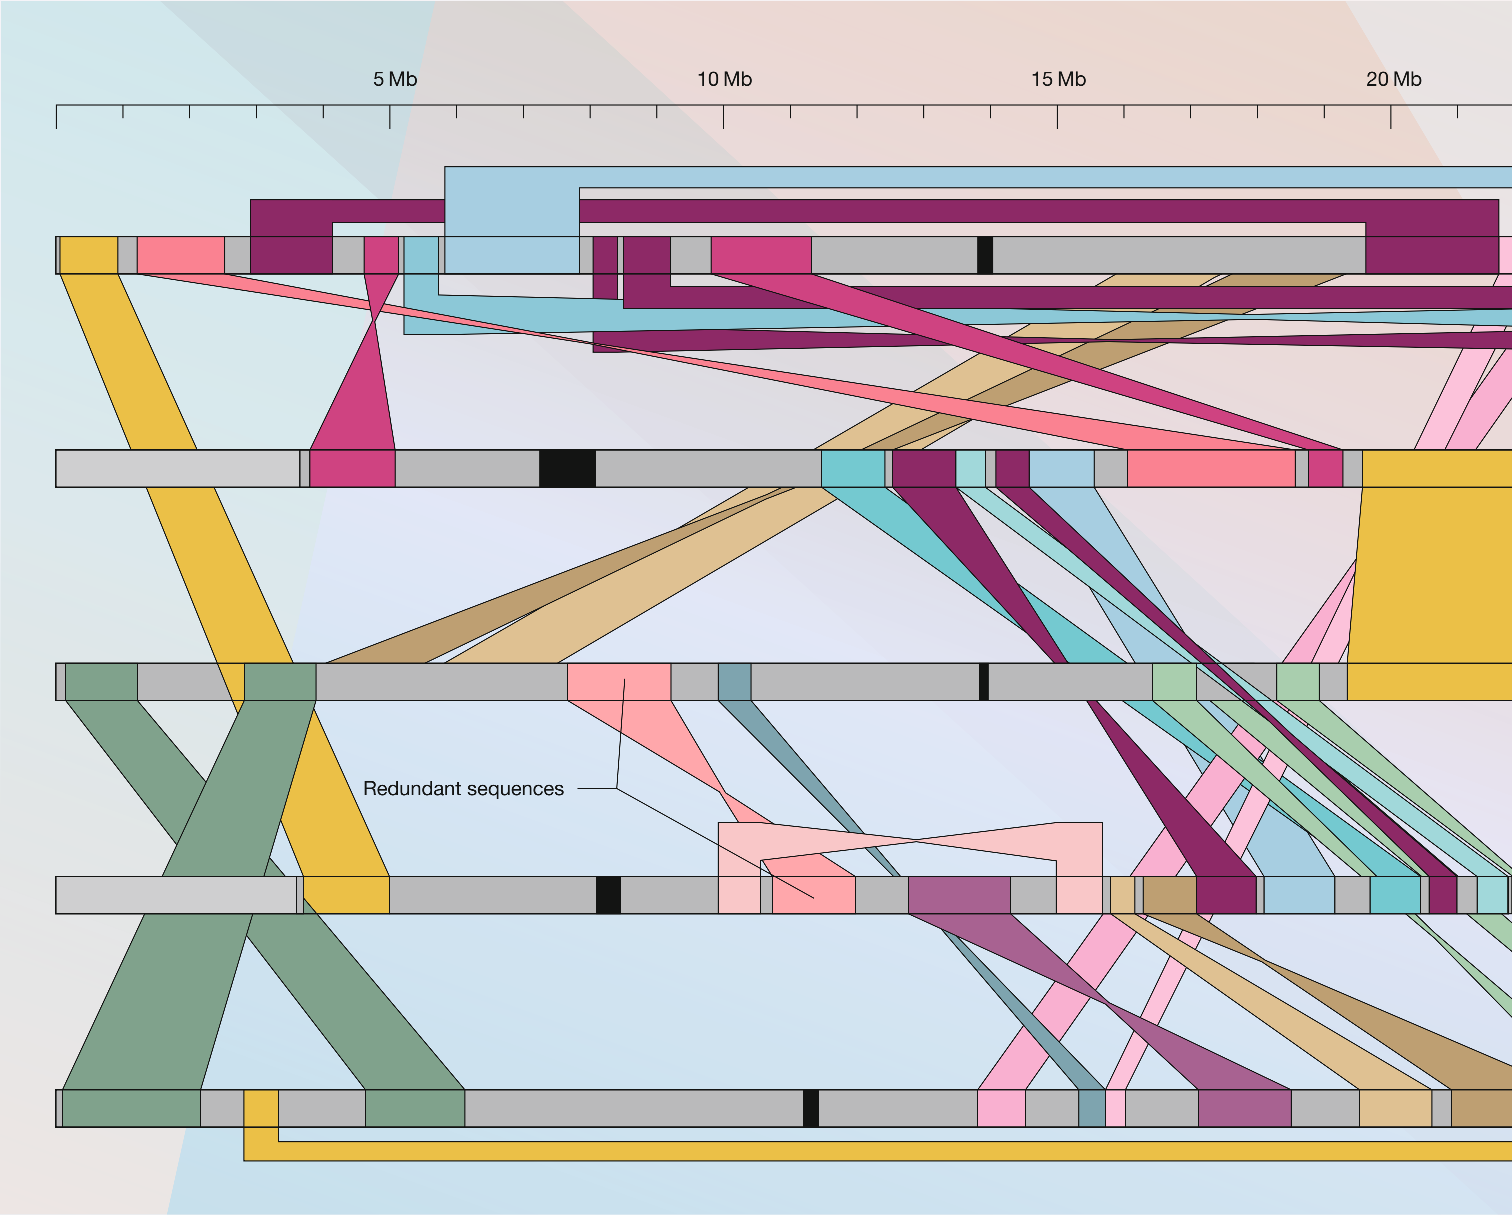

What did these developments mean for you personally?
Well, in the 1980s, I was engaged in collecting protein sequences and annotating them with information about the origin, properties and content of the data. In the subsequent yeast genome project, which started in 1989, I led the team responsible for bioinformatics processing of data obtained from all the participating scientists – roughly 600 – across over 100 laboratories. First, we had to correctly assemble the analyzed fragments into a continuous DNA strand by means of overlapping sections. Then, we converted the sequence data and associated information systematically into computer-readable form and organized these findings in a database. Our 1997 paper on this was published in “Nature” and was widely acclaimed. To illustrate the yeast genome, we also developed an interface allowing simple symbolic visualization of the chromosomes. Numerous links between data elements thus enabled users to navigate the genome using a browser. The catalog grouping the yeast genes by function played an important role here and is still in use today.

But you didn’t stop at compiling databases, did you?
No. Once we had our catalog of 5,800 yeast proteins, I was convinced that we couldn’t just publish it like that – it wouldn’t make sense. So we started to produce a catalog of functions – in other words, we attempted to classify the tasks performed by each section of DNA. The secret of bioinformatics lies in two things – classification and prediction. Reliable prediction of biological properties based on data renders many experiments superfluous, while others can only be interpreted in this way.

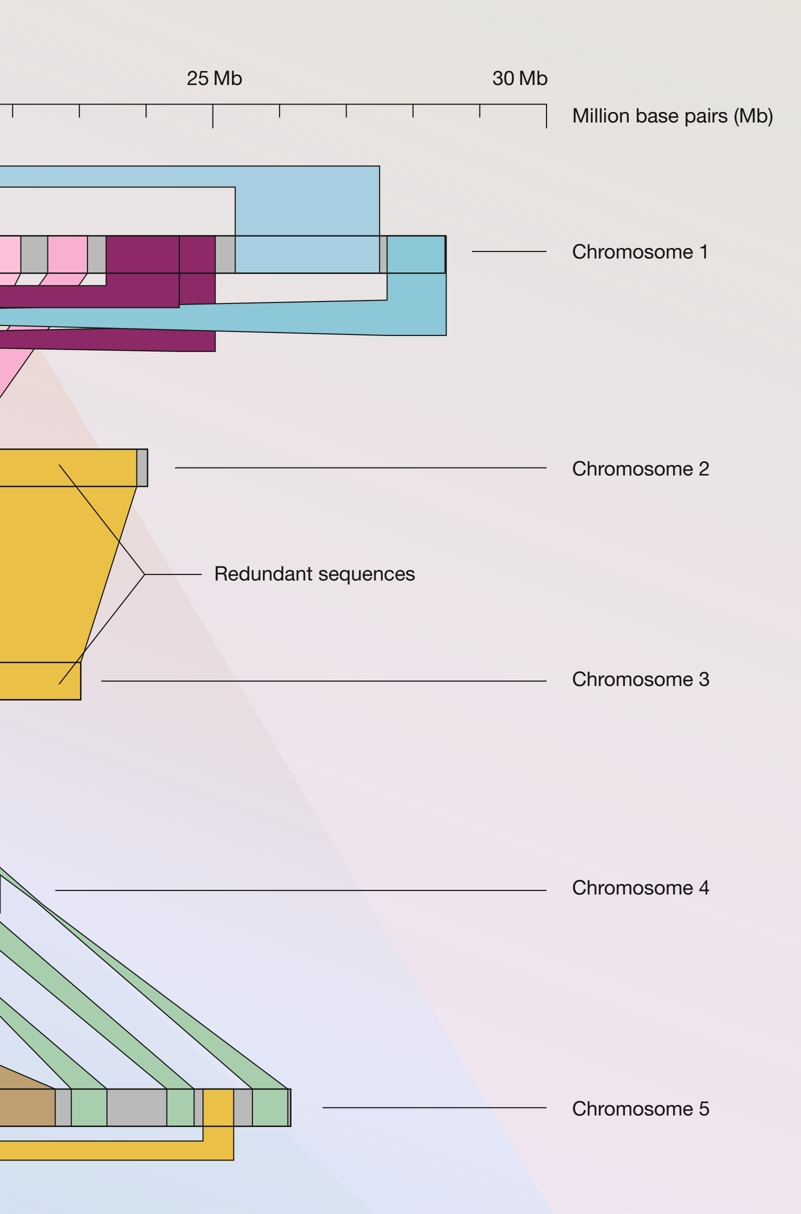

“The core challenge facing bioinformatics today is the sheer volume and complexity of available information.”

Hans-Werner Mewes

Three years after sequencing the yeast genome, the genome sequence of *Arabidopsis thaliana* was unraveled. The genome of *A. thaliana* comprises 125 million base pairs stored in five chromosomes, shown here as gray bars. The sequence analysis revealed that several parts of the genome are redundant – the illustration shows these sequences as colored bands. This indicates that the genome of *A. thaliana* was duplicated several times over the course of its evolution. As members of the Arabidopsis Genome Initiative, Mewes and his colleagues were involved in assembling the entire genome and in localizing those genes which actually code for proteins.

Graphics: ediundsepp (source: AGI, Nature Vol. 408, December 2000)

And how do you go about doing that?
Methods for genome analysis have been developed to make the best use of the collected data. An entire repertoire of algorithms to search for patterns and similarities, translate DNA into hypothetical proteins and link genetic elements with knowledge gained through experiments can be applied to the bare sequences.

Did that not require a huge amount of computing horsepower?
It certainly did. To give an impression of the sheer scale of searches for similar patterns: you are comparing each sequence with all of the others. So the more there are, the more challenging that becomes – it takes an enormous amount of computing time. I think entire power plants were probably running just to work the computers for the Human Genome Project. We had two tricks up our sleeve here: The first was to use a high-speed format to store the data. And the second was to build a network enabling large numbers of Internet users to contribute to these endeavors. This type of crowdsourcing means that anyone who has the time and interest can help with specific digital tasks. The project was implemented by Thomas Rattei at TUM’s Department of Bioinformatics – he is now a professor at the University of Vienna. And there are still around 10,000 private computers involved in this program today, as everything is always being updated. Now, all you need to do is press a button to ask: “Where can we find similar sequences?” ▷

Why is it so important to search for similar gene sequences?
This is the essential groundwork we must cover before we can begin detailed analysis of the sequences to establish the function of individual genes. In this way, we try to predict which functions a protein might have or whether a mutation could trigger a particular rare disease. Research efforts investigating who has a tendency toward obesity and who does not are another example. Gene regulation plays a major role here, which is why it is so difficult to link genetic predisposition to disease.

What methods do you use for this?
Along with many other groups, we use machine learning techniques to search for patterns in large and complex datasets – such as DNA information and disease progression data – in order to identify patient groups responding to a specific treatment, for instance. While neural networks learn from our brain patterns, machine learning involves the computer autonomously developing a model that simulates the data as accurately as possible. This is an iterative process using trial and error, based on a number of known examples. If the outcome is sufficiently reliable, this model can then be used to evaluate therapeutic decisions and individualize treatment.

What other methods did you apply?
In another major step, we set out to determine how proteins interact with one another – which of course feeds into network biology. To do that, we began compiling a catalog of protein-protein interactions a long time ago.
After the turn of the millennium, the idea of systems biology then emerged. This approach aims not only to explore the interactions, but also to establish full computational models to replicate the actual processes. These could then be used to essentially simulate experiments on the computer. However, systems biology struggled with the complexity of the processes; and in practice it was only possible to develop relatively small models, which did not allow the bigger puzzle to be reassembled. As a result, this approach lost some of its appeal.

So how, specifically, does bioinformatics contribute to research?
Bioinformatics is a highly interdisciplinary field. We are not the only ones responding to computational problems. Obviously we do have to master computational techniques, but what the bioinformatician sets out to understand is: What is the biological research question? Which method do I apply? What data does this then generate? How can I convert that data? And how do I then interpret it in the context of the biological question being researched?

Are there already tools for this on the market?
Yes, several, but there is a problem: In the academic space, research focuses specifically on the data needed to answer the next research question, on an ad hoc basis. So you have a team of people who can handle the data and find answers to biological questions, and the end result is a publication. The effort involved is huge, but there isn't really a professional software package or solution to support this. So the market doesn't actually have standard tools that can be universally used for these projects. As it stands, if you want to use bioinformatics in industry, you have to hire specialists, but you will end up with an isolated solution that is not easy to sustain.

In other words, there is no real technology transfer?
It is still very limited, since academic solutions cannot be directly commercialized. That is why, together with colleagues, I founded the Biomax company in 1997, to offer such solutions. But the market beyond large institutions remains tough. Even in the clinical sector, computing power is primarily used for accounting and not set up to use diagnostic data to find the best possible therapy.

Where are the shortcomings, in your view, and what improvements would you like to see?
I would like the data to be used by doctors and in the clinical environment, and I want its application to benefit patients. Patient care produces enormous amounts of information. ▷

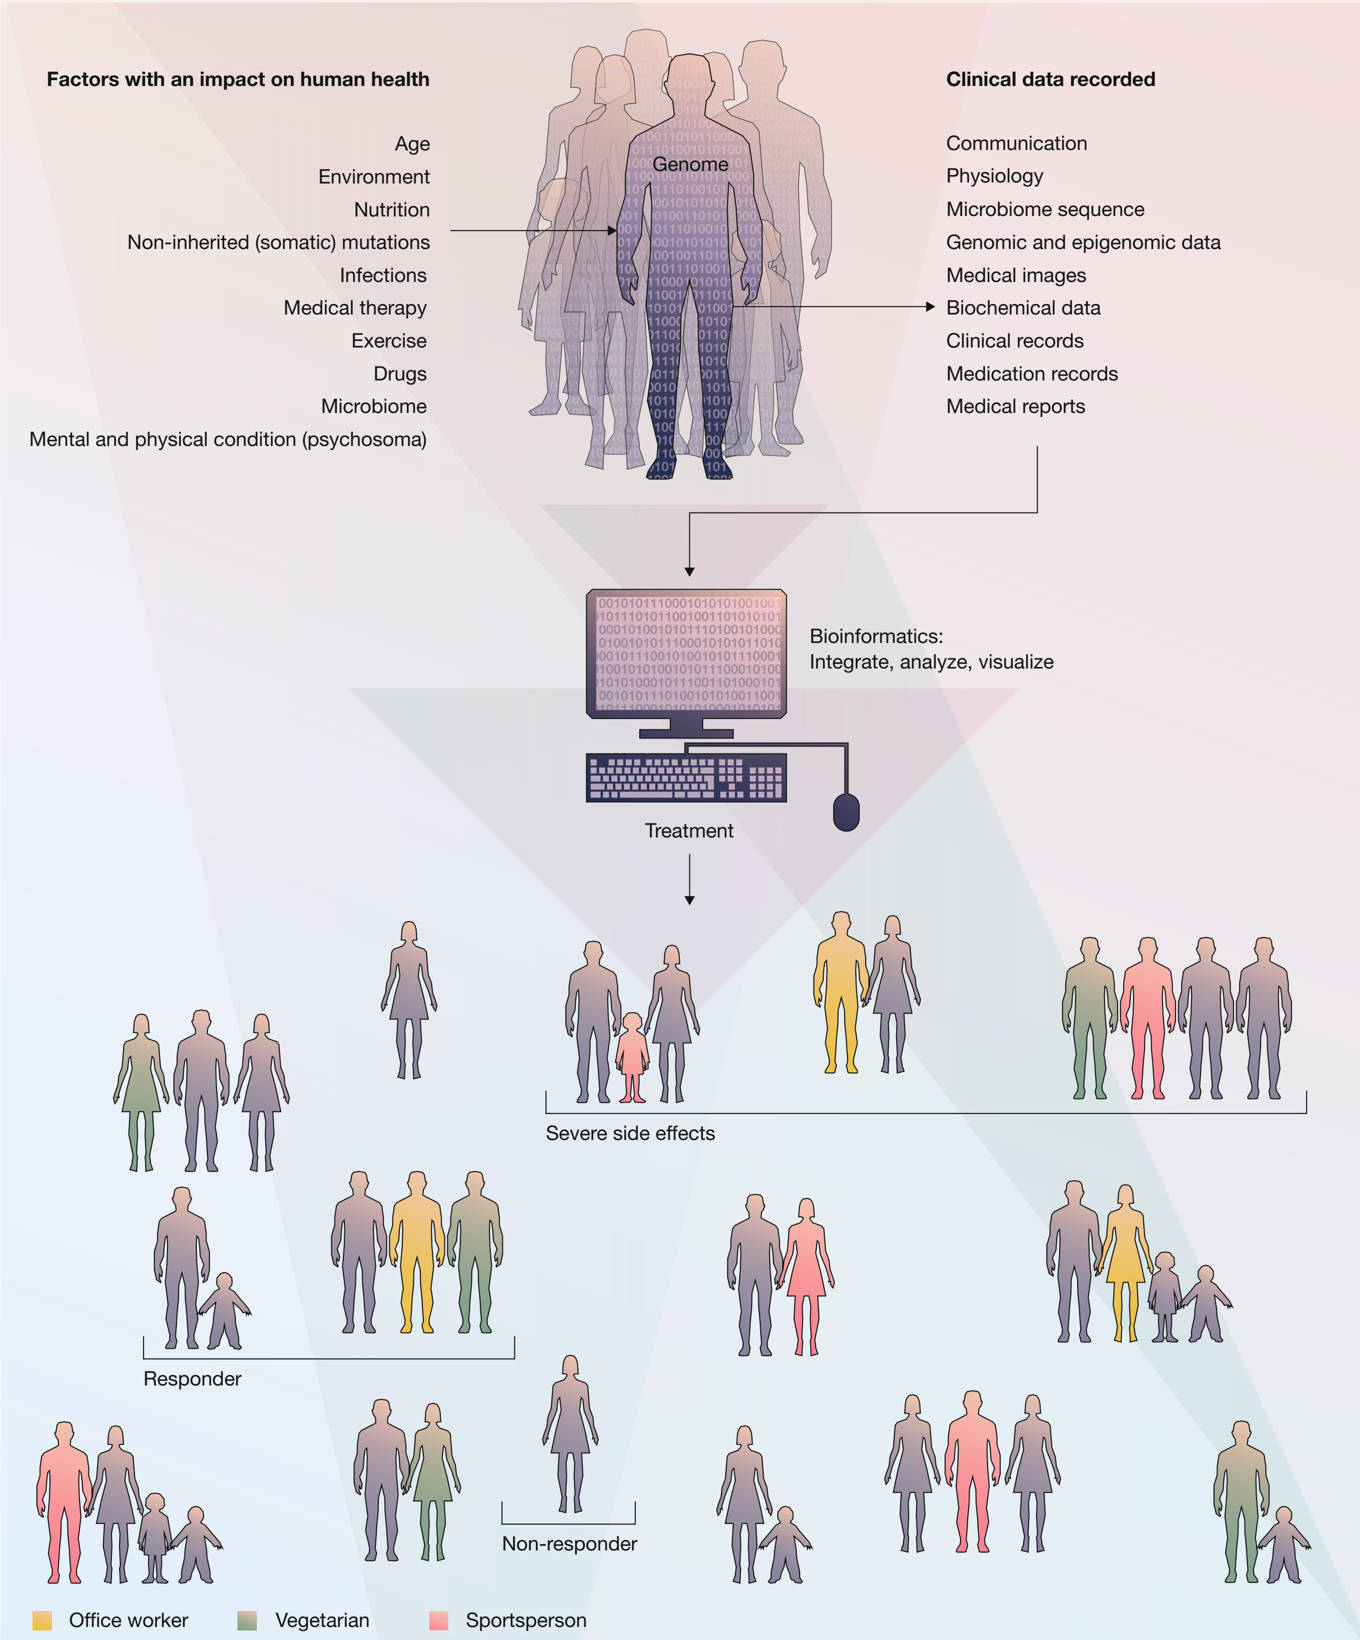

Graphics: ediundsepp (source: TUM), Picture credit: Jooss

How patients react to medication depends on a large number of variables, for instance genome, overall health, lifestyle and environment. Bioinformatics develops tools to gather, analyze and understand complex medical data with the aim of predicting whether a patient will or will not respond to a certain therapy or whether he or she will experience severe side effects from it.

It is estimated that a medium-sized hospital with 100 to 200 beds generates 600 terabytes of data per year – and the use of advanced diagnostics keeps adding to that. But this data mine is currently untapped. Yet it could be a source of important findings, allowing closer analysis of the factors involved in complex diseases such as diabetes or psychiatric disorders. Based on the success of treatment, it should then be possible to identify specific patient groups with different reactions to medication: responders, non-responders, poor metabolizers and super-responders. This knowledge could be used to optimize treatment and save further costs.

And today it is possible to study bioinformatics as a degree course, isn't it?

Yes, the German Research Foundation (DFG) requested proposals for bioinformatics courses in 1999, and I would never have dreamed that I would one day become a bioinformatics professor. But the subject has been very well received and our graduates are snapped up. The course has been available in Munich since 2001, and we have already trained around 600 bioinformaticians. Many of them have also pursued doctorates and some have gone on to become professors. A lot of them work abroad and in industry, with highly interdisciplinary careers. *Interview conducted by Brigitte Röthlein*

Prof. Hans-Werner Mewes

A pioneer of bioinformatics

As a bioinformatician, Hans-Werner Mewes is an interdisciplinary researcher, combining his grounding in chemistry with expertise in biology, medicine and computer science. Following his school-leaving exams (German Abitur) in 1969 in Marburg/Lahn, he studied chemistry at the University of Marburg, receiving his degree in 1978. He then took up roles first at the University of Heidelberg, then at the European Molecular Biology Laboratory (EMBL) in Heidelberg in 1983. Two years later, he moved to the Max Planck Institute (MPI) of Biochemistry in Martinsried, at the same time studying for his doctorate at the University of Marburg. On obtaining this, he became Director of the MIPS (Munich Information Centre for Protein Sequences at the MPI of Biochemistry), and was then made an honorary professor in the Faculty of Biology at LMU Munich in 1999.
Since 2001, Hans-Werner Mewes has been full professor of genome-oriented bioinformatics at TUM, based at the School of Life Sciences Weihenstephan. At the same time, he took over as Director of the Institute of Bioinformatics and Systems Biology at the German Research Center for Environmental Health. Since 2011, he has also been a faculty member of the TUM School of Medicine. From 2010 through 2014, Mewes was an active member of the Helmholtz Association Think Tank and, most recently, also served as spokesperson for the Helmholtz Graduate School of Environmental Health. Hans-Werner Mewes became professor emeritus on March 31, 2017. He is the founder of two companies: Biomax Informatics AG (with Klaus Heumann and Dmitrij Frishman) and Clueda AG (with Volker Stümpflen and Daniel Pinnow).

“What the bioinformatician sets out to understand is: Which experimental method was applied and what kind of data does it generate? How can I analyze that data to find signals? And how do I then interpret it in the context of the biological question being researched?” Hans-Werner Mewes

Picture credit: Jooss

Revealing the Clockwork of Life

With the aid of artificial intelligence, Prof. Burkhard Rost and his team are developing methods to investigate the inner workings of our cells. They seek to enrich our understanding of molecular biology and also to better understand the precise causes of various diseases to support the development of individualized therapies.

Link
www.rostlab.org

Graphics: ediundsepp (source: RCSB PDB)

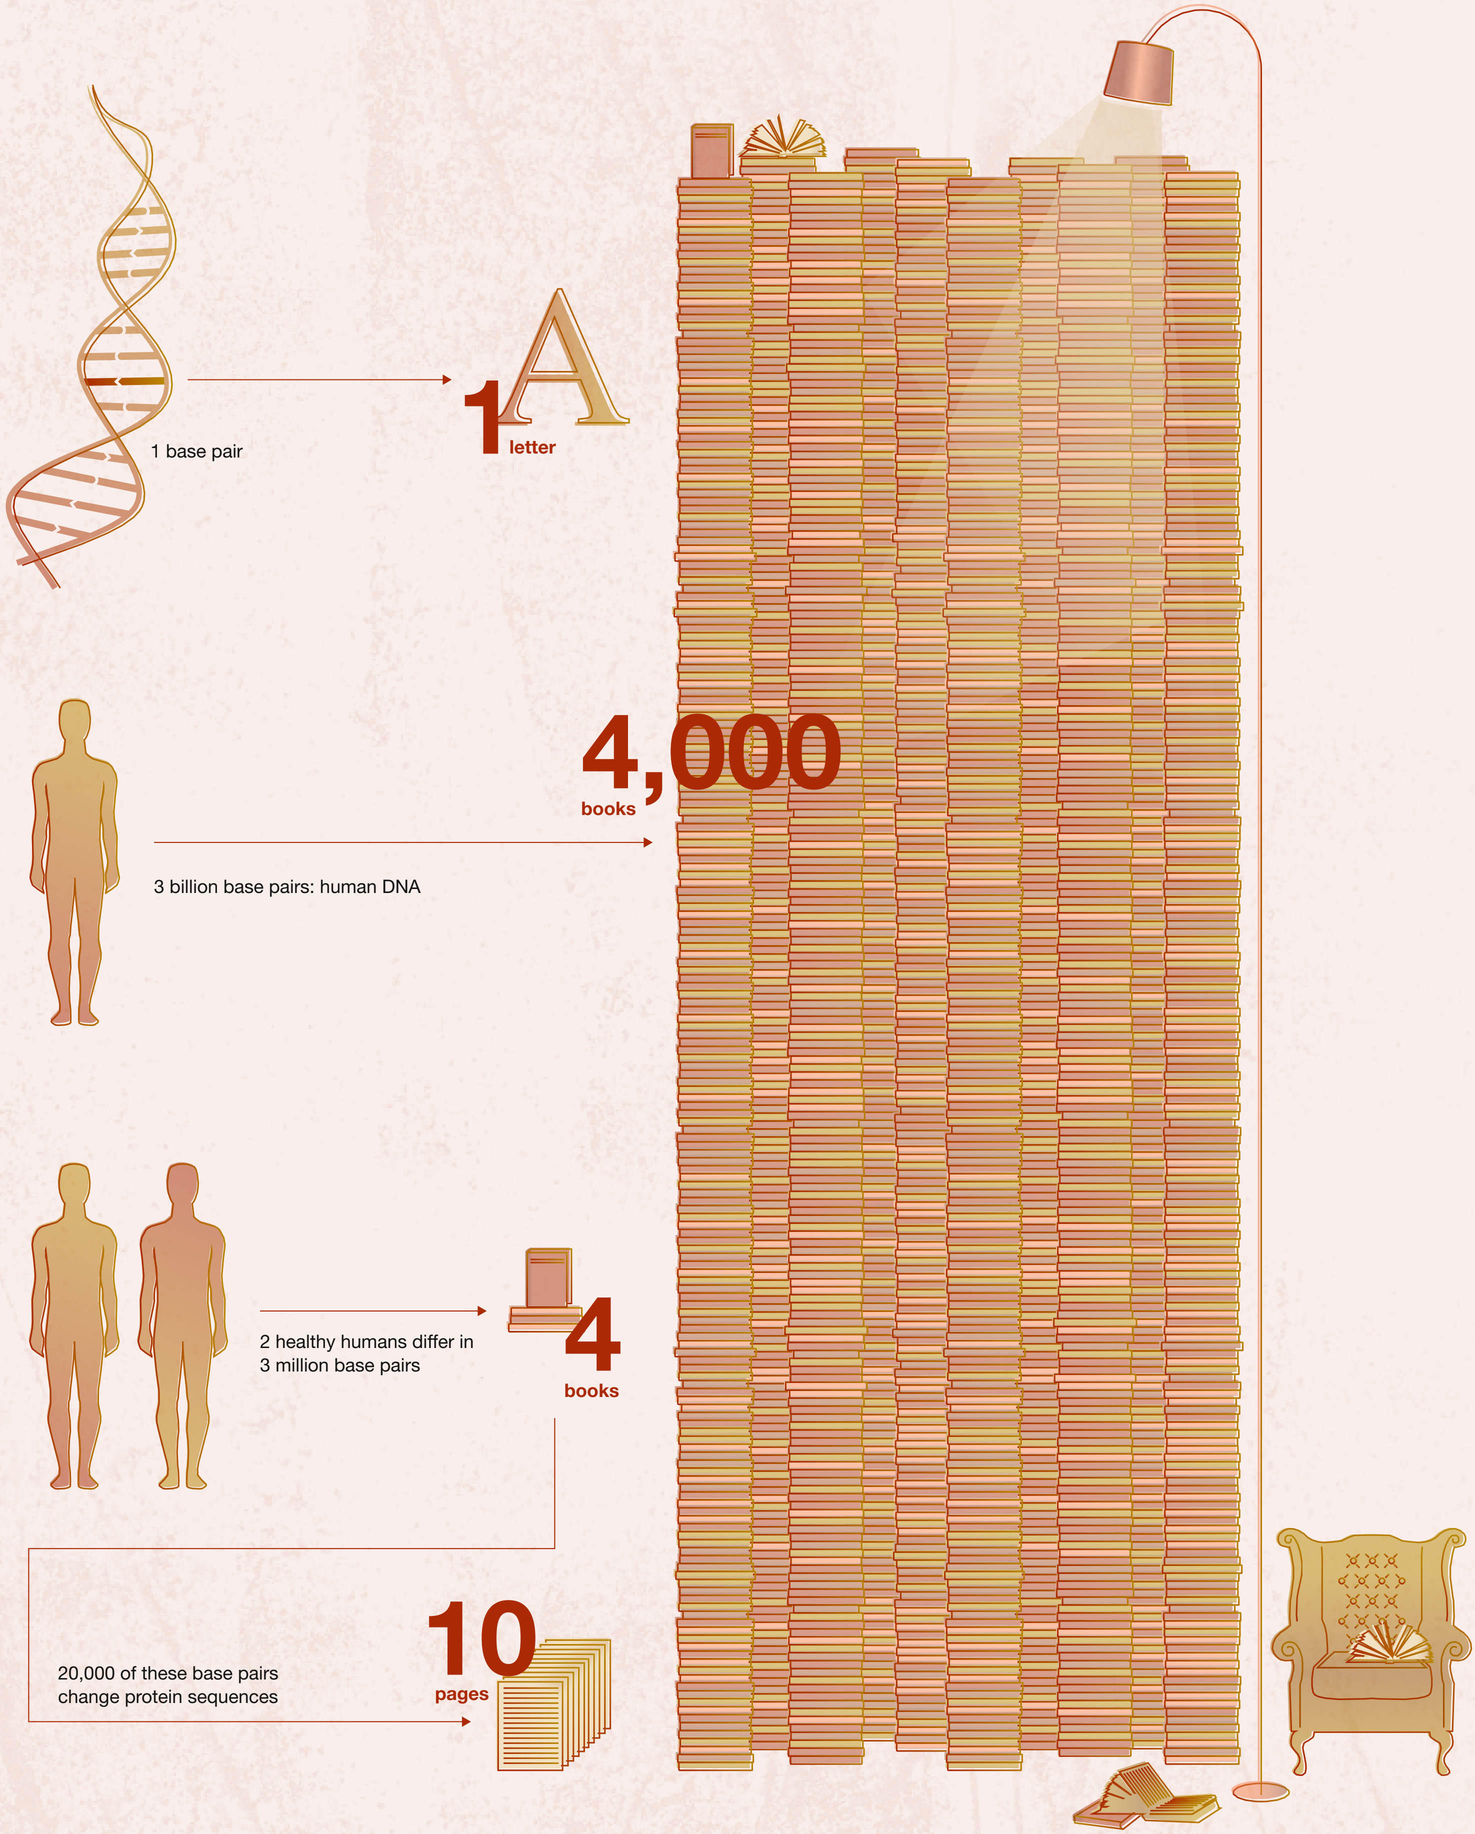
1 base pair
1 letter
3 billion base pairs: human DNA
4,000 books
2 healthy humans differ in
3 million base pairs
4 books
20,000 of these base pairs
change protein sequences
10 pages

Brigitte Röthlein

Diagnosesysteme für die medizinische Werkstatt

Prof. Burkhard Rost und sein Team arbeiten daran, wichtige Aspekte der Proteinfunktionen aus den Informationen vorherzusagen, die im Genom stecken. Sie benutzen dazu einerseits die Daten über Proteine, die sich im Laufe der Evolution unterschiedlich entwickelt haben, andererseits genetische Varianzen, die sich zwischen gleichartigen Lebewesen zeigen. Auf diese Daten wenden sie neben der Statistik Methoden der Künstlichen Intelligenz an, insbesondere des Maschinenlernens, der neuronalen Netze und andere. Es geht in der Regel darum, dass der Rechner aus einer Menge bekannter Beispiele durch Versuch und Irrtum in vielen Iterationsschritten selbständig ein Modell entwickelt, das diese Daten so genau wie möglich simuliert. Ist das mit ausreichender Zuverlässigkeit geschehen, kann man anschließend neue, noch unbekannte Daten eingeben und anhand dieses Modells bewerten. Außerdem benutzen die Forscher strukturelle Informationen über die Proteine, um daraus zusammen mit anderen Daten Rückschlüsse über deren Funktion zu ziehen. Der Bioinformatiker vergleicht seine Arbeit mit der Suche nach Problemen bei einem kaputten Auto: „Bevor man ein neues Auto kauft, kann man sich die Schadensstatistiken für jedes Fabrikat ansehen und damit die Wahrscheinlichkeit verringern, dass man ein Schrottauto erwischt. Sobald man aber ein spezielles Auto hat, nützt dieses Vorgehen bei einem Schaden nichts mehr. Dann muss ein Experte ein Diagnoseprotokoll Schritt für Schritt abarbeiten und mögliche Fehler ausschließen, bis er am Ende die Ursache für den Schaden gefunden hat.“ Entsprechend sieht er den Nutzen seiner Arbeit für die Medizin: Neben grundlegenden Erkenntnissen ist es das langfristige Ziel, die Abläufe in der Zelle zu verstehen. Mit diesem Wissen lassen sich die genauen Ursachen für Krankheiten verstehen und man kann daraus im Idealfall individuell zugeschnittene Therapien für jeden Patienten ableiten. □

“You discover more about biological function from a detailed description than from a number. That’s why I love protein structures.”

Burkhard Rost

Each of our cells is a veritable masterpiece. About 20,000 different proteins work together as molecular machines in each of those, ensuring that each cell receives energy and can perform its specific functions; can move, nourish and renew itself, reproduce, protect against enemies, and send and receive signals. “A cell is fairly crammed with proteins – it is as dense as a solid,” explains Burkhard Rost, Professor of Computational Biology & Bioinformatics at TUM. “The image of a cell as a sack full of liquid in which proteins swim around is completely misleading. It is better to imagine it as a type of clockwork mechanism, with moving parts interacting and interlocking similar to the tiny cogs inside a watch.”

The genetic information is at the heart of all of this activity. It is present in the nucleus of each cell as part of the genome, in the form of a chain comprising around three billion base pairs, and the cell uses it as assembly instructions for its proteins. These usually have an extremely complex design, folding themselves in a particular way to form intricate structures such as spheres, branches, spirals, stalks and channels. ▷

Graphics: ediundsepp (Data: TUM)

Many beliefs held 15 years ago viewed as errors today

Since the British scientific "virtuoso" Robert Hooke discovered in the mid-seventeenth century that living beings are made up of tiny "units of life", which he termed "cells", scientists have been keen to gain an understanding of their interior and their inner workings. When the human genome was decoded in the 1990s, it was first thought that all secrets had finally been revealed. In fact, the real search had only just begun. It quickly became clear that, alongside the actual genes providing assembly instructions for proteins, the genome also contained significantly more segments, whose purpose remained a mystery. So these were initially deemed "junk" and disregarded. Meanwhile, we now know that they also hold information, albeit of a different kind.

The proteins themselves also proved puzzling, however. As more and more human genomes were sequenced, the variation between genes was found to be significantly greater than was first assumed. As Burkhard Rost, Professor at Columbia University in the City of New York at the time, recalls: "Two people differ from one another in around 20,000 amino acids – the building blocks of proteins. On average, there is one difference in every gene. That is absolutely astonishing and we certainly didn't expect this 16 years ago. Back then, it was thought that there was a single reference genome, and all individual genomes could be expressed as small deviations from this reference genome. But we now know that the differences between any two unrelated people are too substantial for this."

So the researchers started to dig. If 20,000 variations in the genome between two healthy people produce no visible difference, it is probably safe to assume that these changes can be regarded as neutral. On the other hand, mutations had also been identified that adversely affect people's health, such as in sickle-cell anemia, which results from a point mutation on chromosome 11, i. e. by changing one single amino acid in all 20,000 proteins a disease is caused. Such conditions are known as "rare diseases". In contrast to infections or major, widespread diseases such as cancer or cardiovascular disorders, these are triggered by an early genetic defect. Angelman, Marfan and Treacher-Collins syndromes are examples here, as are Down syndrome and progeria – and the list runs into the hundreds. Each of these conditions is rare in its own right, but taken as a whole, they affect a good five percent of the population.

When looking at the effects of genes, these two scenarios represent the extremes: either no impact at all or a disease due to a single mutation. The reality is, however, more complex – as a rule, the outcome of genetic variations lies somewhere in between, with most diseases arising from several mutations that interact.

Artificial intelligence comes into play

The question then was how to determine – or even predict – which variants in the genome have which effect? Burkhard Rost had a breakthrough idea in this regard 25 years ago, long before the variations had been identified and the ▷

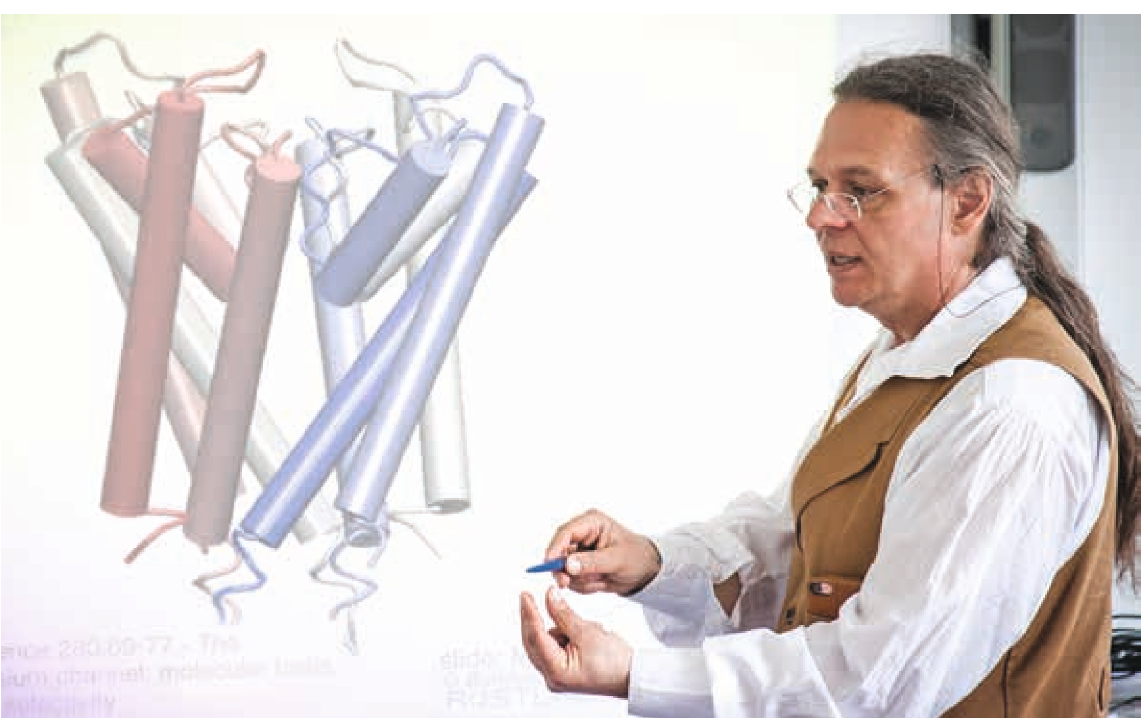

Using tools like machine learning, researchers can predict whether mutations affect the function of a protein, depending on their location in the genome. This image shows the 3D structure of beta hemoglobin. One can clearly see the ring-like heme groups which bind to oxygen and thus transport the gas through the body. In the areas colored grey, mutations have a neutral effect. They change the protein sequence, but do not affect its function. Mutations in the red areas, such as the areas near the heme binding sites, have a significant impact on the protein function. Genetic differences in the blue areas have no effect at all.

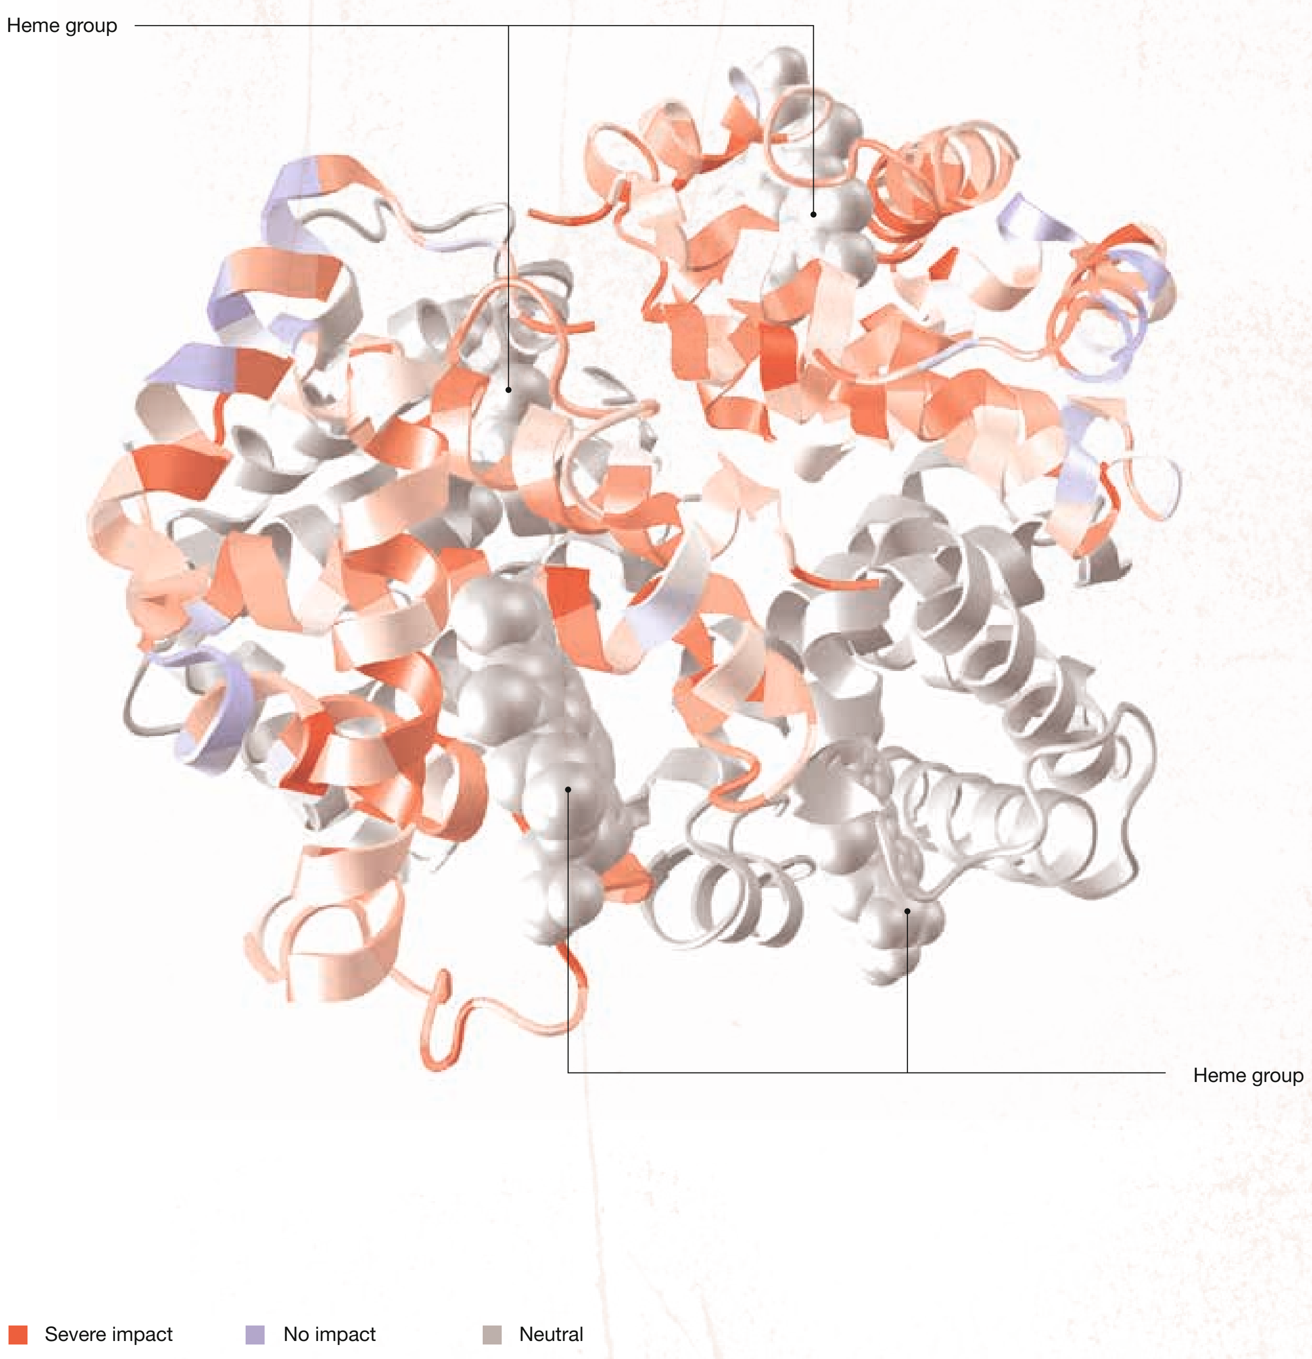

Picture credit: ediundsepp, Graphics: ediundsepp; TUM/Rost; Schafferhans

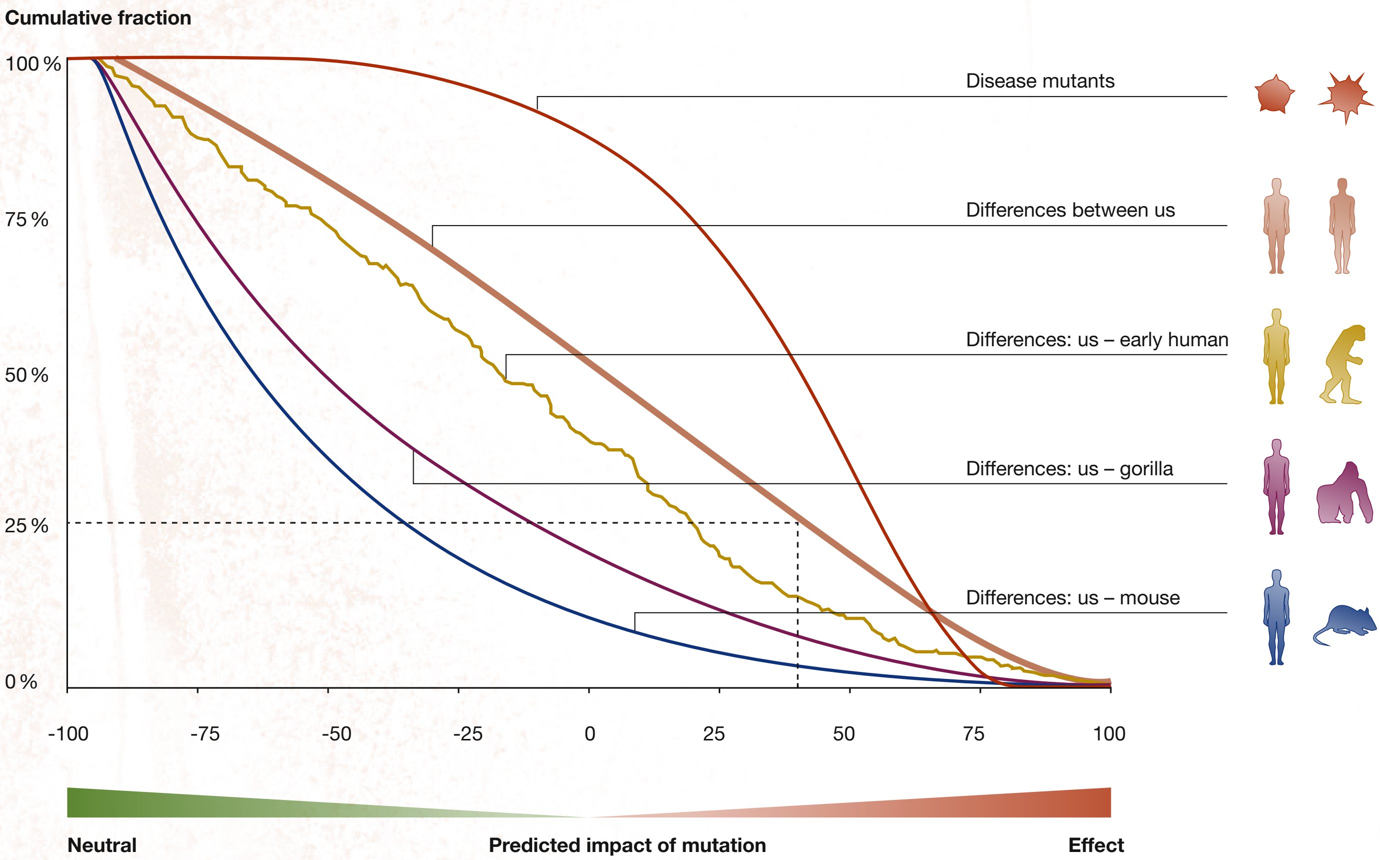

The impact of mutations predicted with the help of machine learning: A strong effect (red) would result in a clearly detectable difference or in a disease, neutral (green) would have no effect at all. Of the 20,000 mutations in which humans differ and which can change protein sequences, about half are neutral, and half have an effect. About 25 percent of these have a strong effect (dashed line). Humans differ from their ancestors in a much higher number of mutations, but these have a much lower effect. On the other hand, mutations that are known to cause diseases are correctly predicted to have a very strong effect.

corresponding data made accessible. His suggestion back then was to compare the gene with its evolutionary relatives: "Suppose there is a particular enzyme in humans that also occurs in chimpanzees. And this protein has the same function in both cases. Maybe we then determine that the ape enzyme is not entirely identical in appearance to ours. But since there is no difference in function, we can assume that has no bearing. From this, we conclude: these variations have no effect; they are neutral – they alter neither structure nor function. If, however, we compare the relevant enzyme with that of mice, we might find that there are differences in function there. So by looking at the evolutionary family of this enzyme, we can see which variations have an impact and which do not. We can now feed this evolutionary profile information into a machine learning method or neural network."

Both of these options fall under the umbrella of artificial intelligence (AI) or machine learning (ML), which entails teaching computers to identify patterns without explicitly programming them in beforehand. Based on a number of known examples, the computer uses an iterative, trial-and-error process to autonomously develop a model that simulates this data as accurately as possible. If the outcome is sufficiently reliable, new, previously unknown data can then be entered and the model used to evaluate it. 1988 saw the first publications by researchers describing how they obtained information using neural networks. The real breakthrough came from combining AI with the evolutionary information, five years later in 1993. This method quickly spread. Today, Burkhard Rost's research team alone uses over 50 different AI methods to create 10 other AI methods that determine what effect variations in an amino acid might have.

The Munich-based researchers are not only focusing on the sequence of bases in the genome here. For instance, they are also analyzing structural information – that is, data relating to the form of the relevant proteins, primarily derived from X-ray crystallography. To do this, biologists have isolated and ▷

Graphics: ediundsepp (source: Y. Mahlich et al (2017) Scientific Reports 7, Article number: 1608), Picture credit: ediundsepp

Prof. Burkhard Rost

A passion for proteins

"I'm always excited to discover that something I believed was wrong," declares Burkhard Rost, one of the very first bioinformaticians. Like all his colleagues, at the outset he thought that the DNA fragments between genes were junk, for instance. Whereas now researchers know that they have other functions.

Burkhard Rost works on predicting the function and structure of proteins and genes, with a particular focus on forecasting protein interactions and the effect of variations in individual amino acids. His research findings are intended to foster a better understanding of protein, gene and cell function. Additionally, he aims to enable earlier detection of diseases and more effective treatment. His research group specializes in connecting artificial intelligence and machine learning with evolution.

After studying physics, history and philosophy at Giessen and Heidelberg universities, Rost received his doctorate at the European Molecular Biology Laboratory (EMBL) in 1994. He had already developed the first Web server for predicting protein structures in 1992. Following research visits to EMBL and the European Bioinformatics Institute in Cambridge (UK), as well as a brief period in industry at LION Biosciences in Heidelberg, he took up a professorship at Columbia University (New York City) in 1998. In 2009, he accepted an appointment to the Chair of Bioinformatics at TUM, where he holds an Alexander von Humboldt Professorship. He is a member of the New York Academy of Sciences and has been President of the International Society for Computational Biology since 2007.

crystallized the protein, before passing X-rays through it. They can then use the resulting diffraction patterns to draw conclusions – sometimes highly detailed – concerning the form of the protein. "International programs for structural genomics have made huge efforts to get closer to the aim of assigning a 3D structure to every protein," reports Burkhard Rost. "But we are still a long way from understanding them all. In humans, 3D structures have been experimentally determined for fewer than a quarter of all proteins to date."

A picture worth more than a thousand words

Nonetheless, this structural biology information can yield fascinating insights, since proteins work like tiny, three-dimensional machines. There are three different ways to describe them: first, in terms of their biochemical activity; second, by biological function; and third, by localization – that is, where their activity is sited, as in the mechanism of a clock. As Burkhard Rost emphasizes: "You discover more about biological function from a detailed description than from a number. That's why I love protein structures."

Here again, machine learning can tell us quite a lot about the possible 3D structure of proteins that have not yet been analyzed. Researchers can develop models showing how they interact with other known proteins. Another example entails investigating where in the cell the relevant protein must be located in order to perform its role: in the nucleus, where it facilitates gene transcription, or in the outer membrane, where it might function as an ion channel, or as a signal protein between the two?

Sometimes even a section with no discernable function can have a specific task. Here, researchers talk about "disorder", meaning parts of a protein that simply do not fold. "You can think of these pieces as the padding in Amazon parcels – they don't fold either," as Rost explains the phenomenon. "These areas are there to stop anything interacting with the protein. But the molecule can also use these fragments to scan their surroundings, and if something approaches, it emits a signal."

Ion channels are particularly complex. These are small pores in a membrane, able to open or close to allow something to pass through or block it. Few of these proteins, in particular, have revealed their exact structure to date, since they cannot be crystallized without their surrounding membrane. But even they can be analyzed using artificial intelligence: "If we identify a channel that does something quite different in bacteria compared to humans, for instance, we can ask why. Where are the similarities between the two; where are the differences?", describes Rost. "We also look at the structure: which part lies in the membrane; what signal causes the channel to open and close? Controlled by which amino acids? We can find this out through experiments, but we can often predict it too. To do this, we are investigating the entire ion channel family."

The difficulty with the data

The 55-year-old bioinformatician likens his work to searching for the problem when a car breaks down: "Before you buy a new car, you can look at the breakdown statistics for each make and thus reduce the probability of ending up with a "lemon". But as soon as you have an actual car, this process is of no use – it can't help you narrow down the reason for a break-down. In that case, an expert needs to work through a diagnostic protocol step by step, excluding potential sources of failure until they finally identify the root cause. Healthcare policy to date resembles the quest for breakdown statistics, but not really the search for actual causes of disease." Rost is determined that the latter should receive more support.

No matter what method is used, one thing is certain in bioinformatics: The more information you can feed into the models, the better they become. That is why researchers depend on obtaining as much high-quality data as they can. But there is a major problem here: "I am in favor of the free exchange of data," states Rost. "There are many issues we cannot address if the information gathered is not shared. Restricting data access on legal grounds is a mistake, from my point of view. We need to reach the stage where people understand: We can only move forward if I make my data available, even if other people then know I have a particular condition."

Brigitte Röthlein

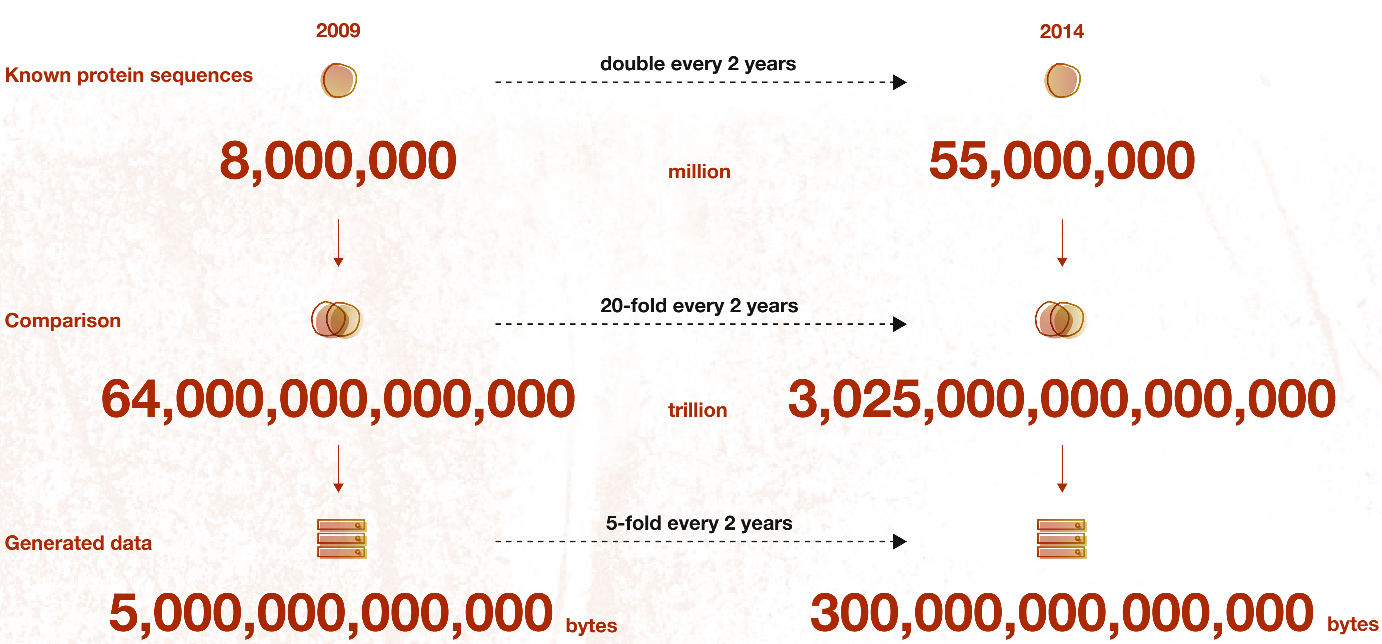

Graphics: ediundsepp (source: TUM)

Big data is at the center of bioinformatics, as the discipline generates huge amounts of data which multiplies every few years.

+ FORSCHER
VERPACKUNGS-
KÜNSTLER
ALTANA – global führend in reiner Spezialchemie. Wir setzen weltweit auf die Kunst, genau auf den Kundenbedarf zugeschnittene Produkte zu entwickeln. Deswegen entwickeln unsere Experten an rund 50 Standorten weltweit nicht nur die passende chemische Formel, sondern die optimale Lösung, damit zum Beispiel Verpackungen nicht nur schön aussehen, sondern jeder nur denkbaren Belastung standhalten.
Entdecken Sie dieses Plus für Ihr Geschäft: www.altana.de/plus

Looking Ahead – **Into our Real and Virtual Future**

Whether through data glasses, a smartphone or a tablet – augmented reality (AR) enriches our physical surroundings with virtual information in real time, making complex concepts or settings easier to understand. TUM professor Gudrun Klinker is interested in a wide range of AR research topics: She has developed a reference architecture for industrial AR applications, is working to improve eye tracking in head-mounted displays and is investigating AR objects that not only communicate data, but are also fun to use.

Link
campar.in.tum.de/Chair/ResearchAr

Graphics: ediundsepp; Picture credit: Eckert

Augmented reality put to use in the humanities

Klinker's team developed a game in which players build a Celtic village. Markers on the playing cards correspond to houses with various functions. Using an AR app, players can visualize them in 3D and find additional resources inside, including information about Celtic life or rune stones.

Gitta Rohling

Weiter Blick – in die reale und virtuelle Zukunft

Dank eines Trackingsystems und Markern zeigt das Display einer Bolzenschweißpistole dem Monteur, wo sich die Schweißpunkte befinden und wie die Pistole geneigt werden muss, so dass sich der Bolzen gezielt platzieren lässt. Entwickelt hat das Gerät Prof. Gudrun Klinker in einem Forschungsprojekt mit der BMW AG. Das Beispiel zeigt: AR reichert die reale Welt in Echtzeit mit virtuellen Daten an – und erweitert sie damit. Das kann per Smartphone, Tablet oder Datenbrille passieren. Auf jeden Fall: „AR ist das ideale Hilfsmittel, um mir Informationen aus meiner Umgebung sinnvoll an die Hand zu geben und gezielt nutzbar zu machen", erklärt Klinker.

In der Produktionsplanung lässt sich AR ebenso einsetzen wie in der Wartung und Reparatur. Allerdings: Für AR-Anwendungen existieren in der Regel geschlossene Systemwelten. Eine Referenzarchitektur, die die Gemeinsamkeiten zwischen unterschiedlichen Anwendungen beschreibt, hat Gudrun Klinker im Forschungsprojekt „ARVIDA" entwickelt. Die Partner aus Industrie und Forschung haben verschiedene Anwendungen in der industriellen Produktentwicklung und Produktionsplanung auf ihre Gemeinsamkeiten hin analysiert und schafften aus Einzelfällen Standards. Ein Beispiel ist ein standardisiertes Tracking, mit dem Bewegungsabläufe von Fahrern ebenso wie Fahrzeugbauteile getrackt und in eine virtuelle Umgebung eingespielt werden können. Nach drei Jahren Forschung ist im Jahr 2016 eine offene, auf etablierten Web-Technologien basierende Referenzarchitektur entstanden. Produktentwickler und Produktionsplaner haben damit einen Werkzeugkasten an der Hand, mit dem sie schneller, flexibler und kostengünstiger als bisher neue AR-Anwendungen zusammenstellen können.

Bei weiteren Forschungen beschäftigt sich Klinker mit der Frage: Wie viele virtuelle Informationen kann der Mensch aufnehmen, ohne die Konzentration auf das reale Geschehen, zum Beispiel beim Autofahren, zu verlieren? Diese Frage zieht gleich eine weitere nach sich: Wie kann der Mensch animiert werden, etwas zu tun? Klinker entwickelt derzeit beispielsweise gemeinsam mit Ernährungswissenschaftlern und Medizinern ein Trinkglas für ältere Menschen. Das Glas stellt fest, wie viel der Mensch getrunken hat und erinnert ihn bei Bedarf spielerisch ans Trinken – indem es etwa aufhört, auf einem integrierten Display die Fotos der Enkel zu zeigen. „Die Technik ist sehr weit gediehen, jetzt können wir uns auch auf die Psychologie konzentrieren", freut sich Gudrun Klinker. □

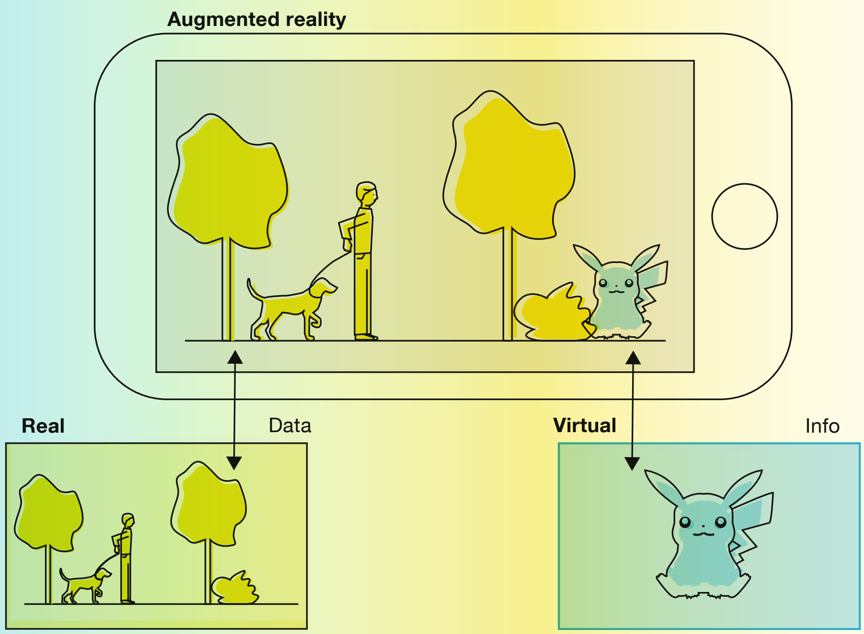

How AR works

Augmented reality adds extra, virtual information to the real world in real time.

"AR is the ideal way to place valuable information about my surroundings at my fingertips and enable me to use it in a really productive way."

Gudrun Klinker

Graphics: ediundsepp; Picture credit: Eckert

With a single, well-aimed shot, the bolt enters the vehicle body. And – pow! – the second accurate hit follows just seconds later. This speedy, sure-fire success comes courtesy of an augmented reality stud welding gun. TUM professor Gudrun Klinker developed this device in a research project with BMW AG, and from 2003 the company used it for years in prototype manufacturing. Using a tracking system and markers, the gun's display shows the technician where the weld points are located and how to place the device to position each bolt with high precision.
As this example shows, AR creates an interactive, real-time connection between physical objects (like the car frame) and virtual information (the welding points), which appear in the same environment. Tracking systems and markers align both worlds geometrically and synchronously so that the virtual objects are positioned exactly. In short, AR adds extra information to the real world in real time – and thus expands it. This might take place via a smartphone, tablet or head-mounted displays: "AR is not limited to any specific device or particular purpose. You only need to think of your car: a talking navigation system, a vibrating steering wheel – all that is AR too, since it gives us additional information," underscores Klinker. The stud welding gun was originally designed for a head-mounted display, for instance. But since the tip of the gun had to be positioned with high accuracy, a built-in display on the device proved the better choice. But regardless of the device, according to Klinker: "AR is the ideal way to place valuable information about my surroundings at my fingertips and enable me to use it in a really productive way." ▷

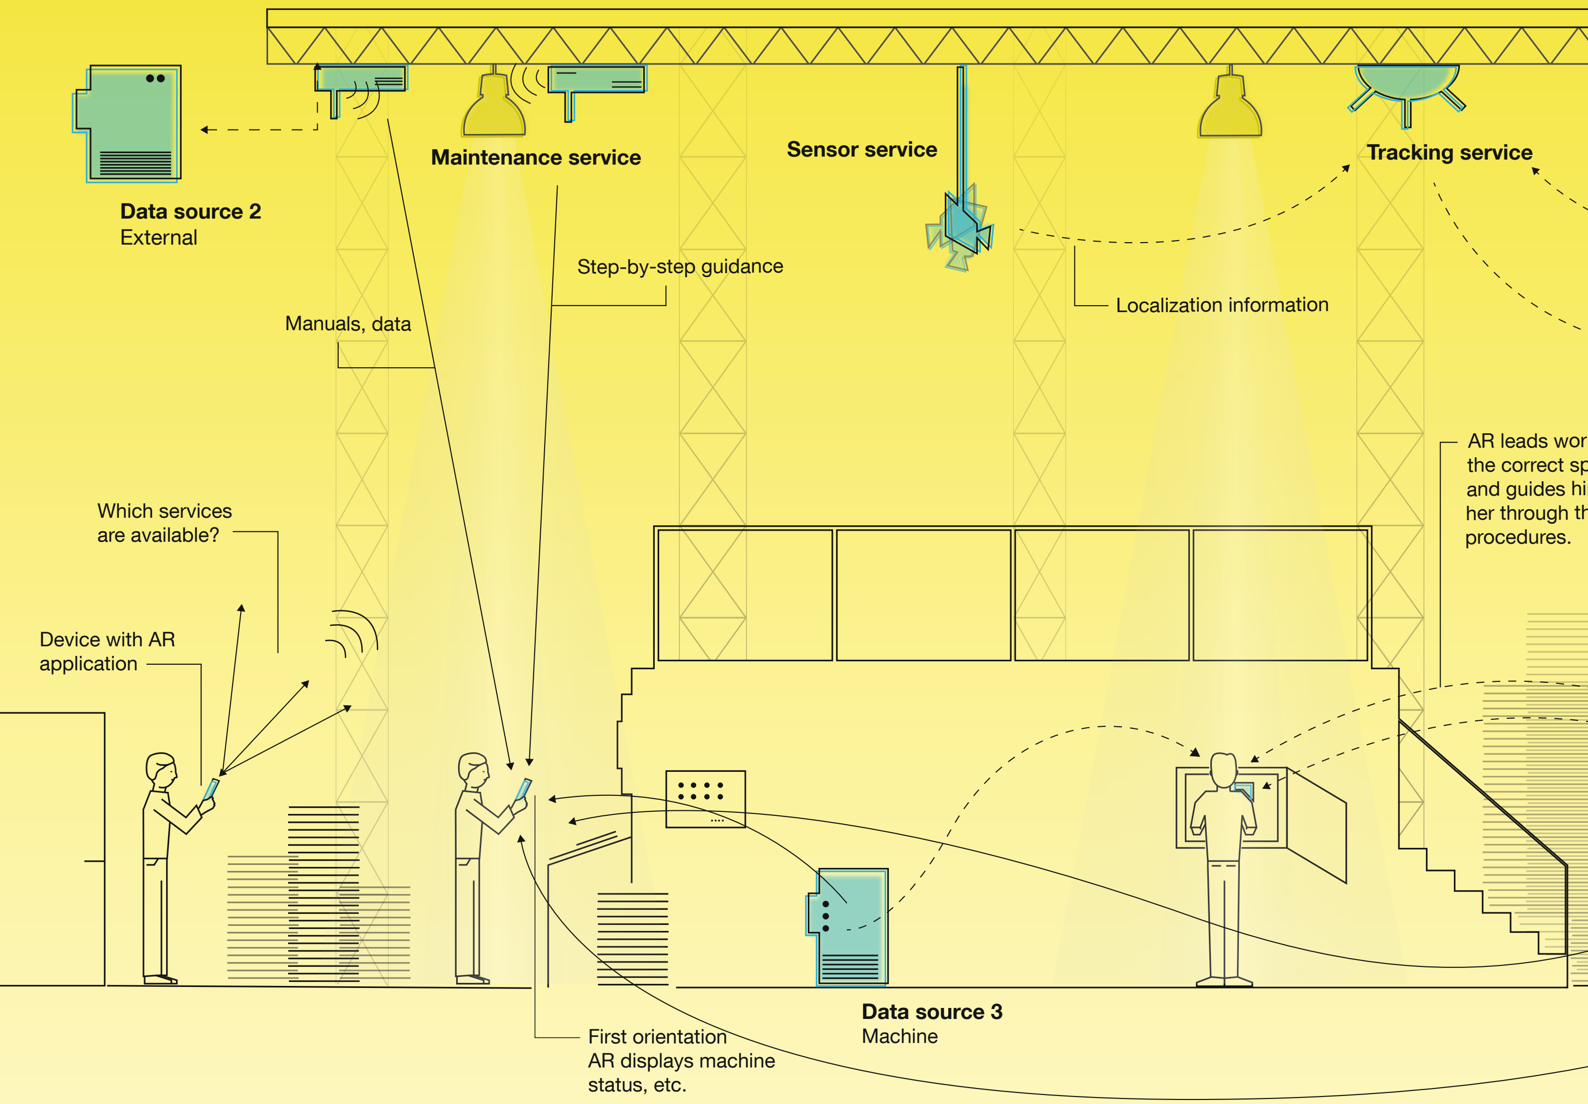

An open reference architecture based on Web technologies for rapid and flexible implementation of AR applications (right) – shown here in a maintenance scenario (left). Centrally provisioned services support typical functions such as user tracking, which maps the user in relation to other objects in the vicinity, and sensing, which uses data from cameras, thermometers and microphones.

A sense of abstraction: reference architecture for AR applications

The use of AR in industry is gathering pace – in production planning, for instance, as well as for maintenance purposes. Another example is vehicle repair, where technicians might receive the exact information required for a specific vehicle by tablet, say, saving them the time and hassle of leafing through printed manuals.

To date, though, system environments for AR applications have generally been based on closed designs. This makes adding new functionality difficult and expensive. In addition, accessing third-party systems is complex, as data needs to be converted to do this. In the ARVIDA research project, funded by the German Federal Ministry of Education and Research (BMBF), Klinker and her team thus worked with partners in business and science to develop a reference architecture for AR applications. The abbreviation stands for the German project name, which translates as “Applied reference architecture for virtual services and applications.” In contrast to pre-programmed and inflexible monolithic applications, this architecture uses abstraction to describe common features and thus enable the design of overarching development systems. To do this, the researchers analyzed the similarities across several applications in industrial product development and production planning.

One application, for instance, was designed to improve vehicle ergonomics. As a first step, the researchers used a tracking system to record the motion sequences of several drivers.

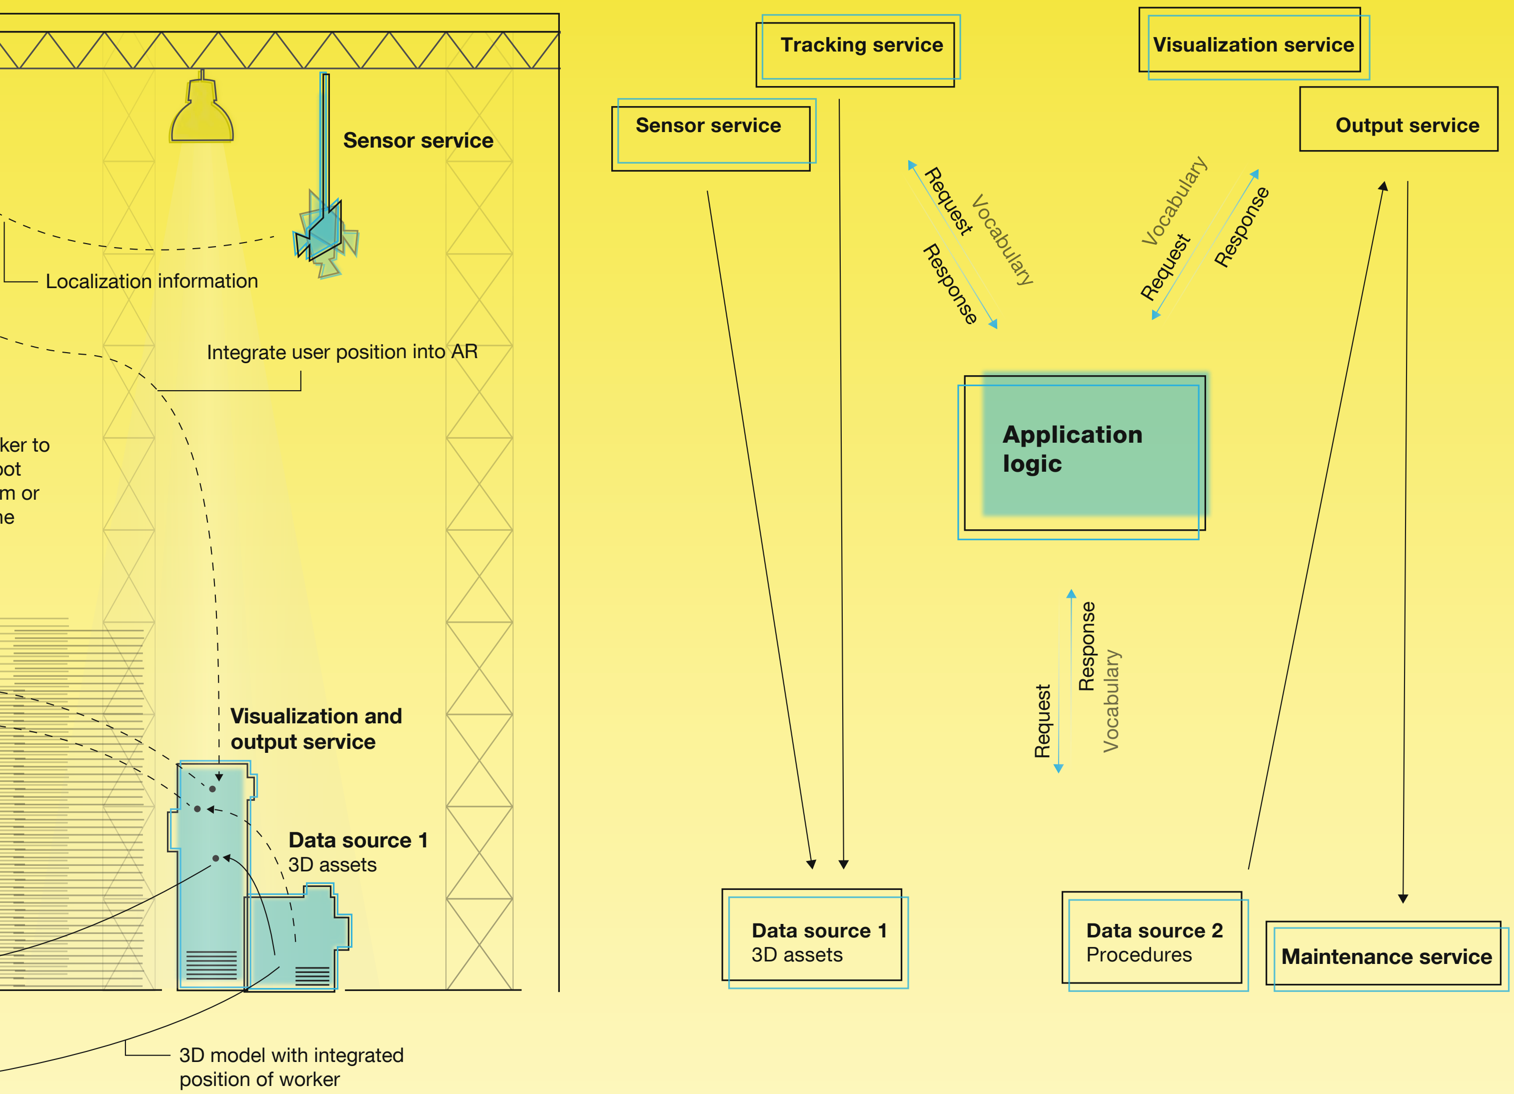

Alongside concrete information such as 3D system models and manuals/instructions, data services also offer data models and simulations to forecast the outcomes of specific actions. Visualization and output services present the information in optimized format on the user's device or on appropriate displays in their surroundings. These services employ a defined vocabulary to present their services to ensure they can be used across different AR applications.

They then used this to define basic movements, which they imported into a virtual vehicle environment. From this movement pool, they could now create new motion patterns without having to track them again – thus, essentially from the lab. The researchers also standardized the process to extend beyond this specific use case, for use in other applications – such as improving the installation of vehicle parts. Instead of motion sequences, this involves tracking components and integrating them in a virtual vehicle. Checks can then be carried out even before the components are manufactured, ensuring they are fully compatible and identifying ways to improve their assembly.

For the reference architecture, Klinker and the ARVIDA team held intensive discussions to analyze the recurring methods across the various applications. From this, they put together and tested application packages, thus deriving standards from individual use cases. "It's all about choosing a level of abstraction that makes sense such that you can capture the essence," summarizes Klinker. After three years of research, the outcome in 2016 was an open reference architecture based on established Web technologies. This gives product developers and production planners a tool kit they can use to build new AR applications faster, more flexibly, and more cost-effectively than ever. ▷

Graphics: ediundsepp (source: TUM)

Input modes for VR/AR

How precisely can hand movements be transmitted into the AR?

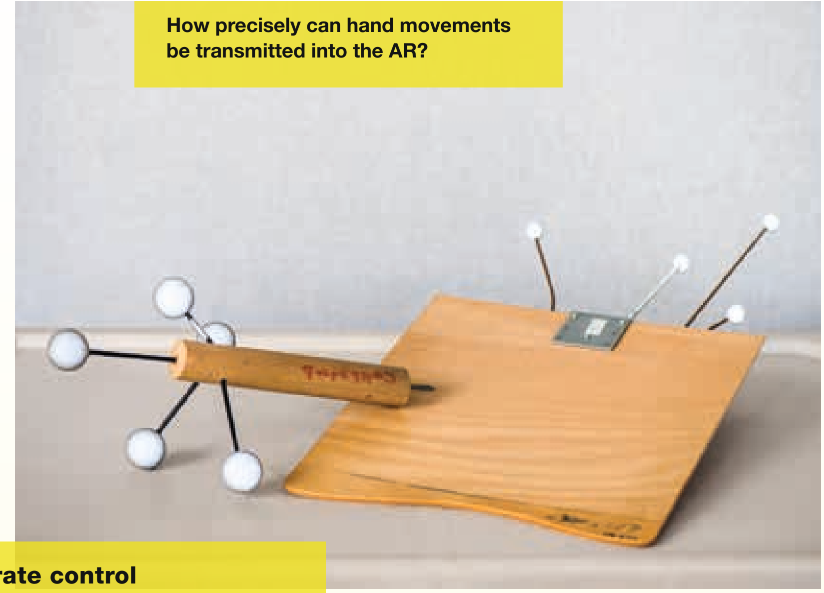

Accurate control

Testing hand tracking and gesture detection devices: Is the user able to select individual squares of a grid? The left monitor shows the grid which the user sees in stereo in his glasses.

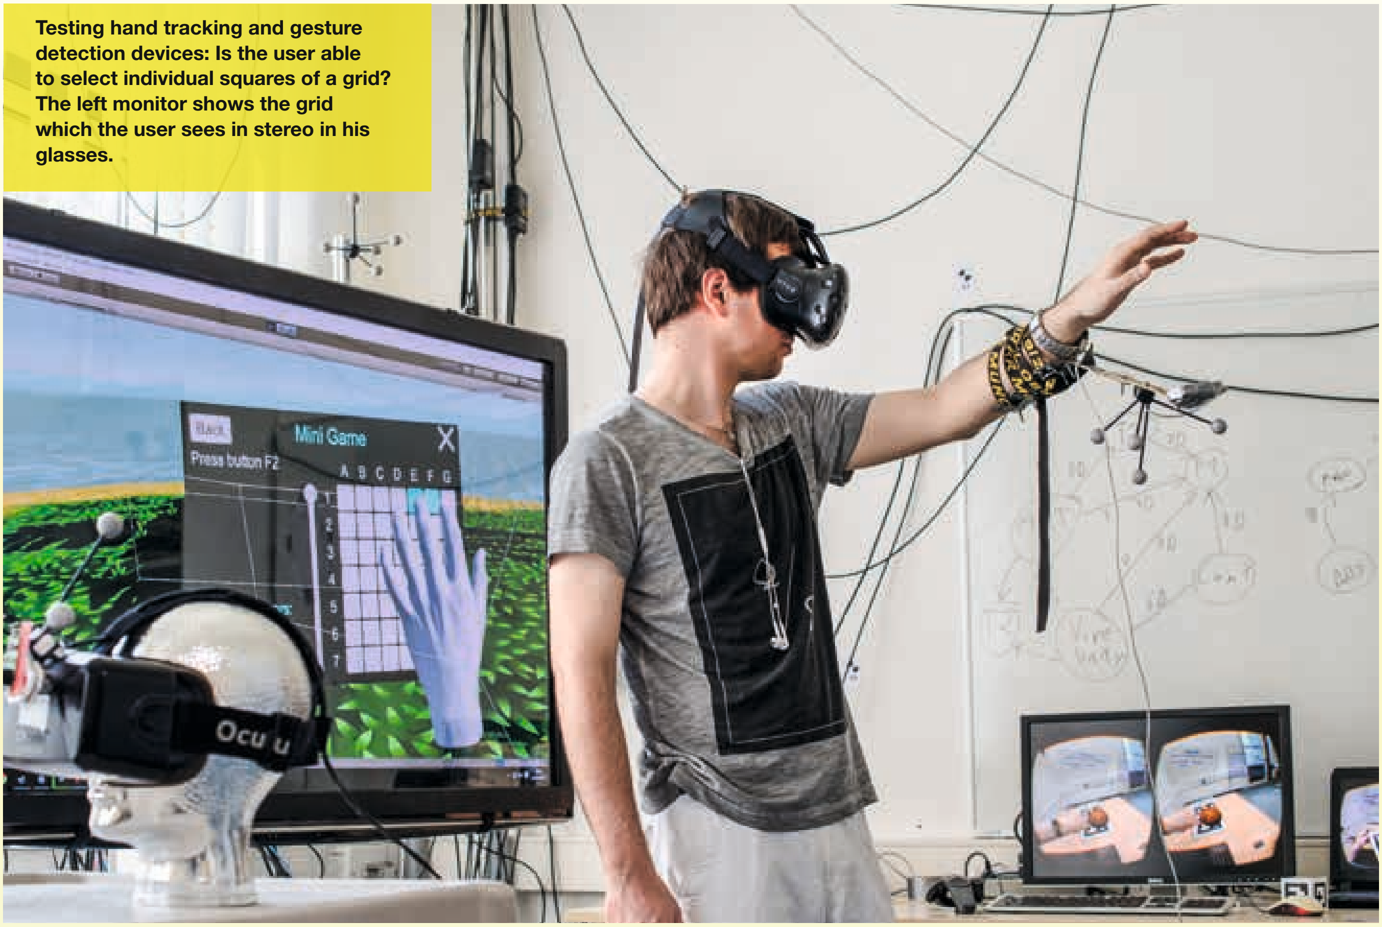

Tracking hand movements and gestures

For an AR to react smoothly to the user's actions, gestures and hand movements have to be determined very accurately.

Staying in touch with the real world

See-through AR headsets pick up the user's view with an integrated camera and display augmented information directly on the glasses.

Eye tracking

AR systems that "know" where the user is looking are able to display information precisely in the right spot.

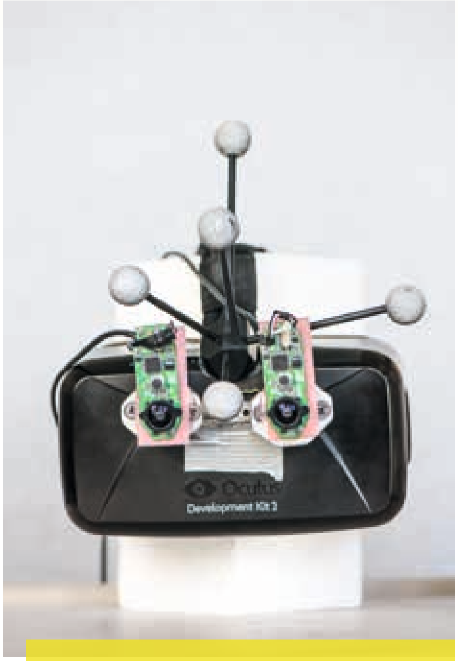

VR becomes AR

The real world recorded by the camera is displayed inside this VR headset. Additional markers provide extra options to track the user's position and head movements.

Picture credit: Eckert

Augmented reality allows people to remain present in the real world instead of shutting it out.

Prof. Gudrun Klinker

It is no surprise that Gudrun Klinker opted to pursue her research in augmented rather than virtual reality: VR would have been far too one-directional for the professor, who now leads the Augmented Reality research group at TUM. AR links the physical to the virtual world, connecting hard facts with creativity – a juxtaposition that ideally suits her personality.
At school, Klinker was equally interested in sciences and languages – as reflected in her choice to focus particularly on mathematics and French. When selecting her degree course, she got hold of a study guide for computer science and the description of the subject as a mixture of math, logic, language and philosophy immediately caught her attention. “The core focus was not on programming, but on handling information and structuring it systematically. This remains the essence of it today – and continues to fascinate me now as it did then.” Although the beginning of her career – back in the days of punch cards – mainly entailed data processing, her focus was always less on the device and more on the information in its own right.
Klinker specialized in image processing right from the start. Following her Informatics studies in Erlangen and Hamburg, Germany, she moved to the US, where she obtained her doctorate from Carnegie Mellon University in Pittsburgh before starting work with the Digital Equipment Corporation at its Cambridge Research Lab. The 1980s were a pioneering period for computer graphics and virtual and augmented reality. This was also when she saw the first CAVE at a trade fair – a space for projecting an immersive, virtual 3D environment.
After eleven years – now married and a mother of two – Klinker returned to Germany, where she teaches at TUM’s Department of Informatics as Professor of Augmented Reality. Here, she has been instrumental in establishing the “Informatics: Game Engineering” degree program, which has attracted 150 to 200 students per semester since 2011. As she emphasizes: “Given our expertise as a technical university, we designed the program with a highly technical slant. Our aim is to give students in-depth competence on the technical aspects of games, while also incorporating enough interdisciplinary insights to ensure they have an open mindset later on towards design, artistic and game production aspects.” Ultimately, interdisciplinary collaboration is a key success factor for innovation – and, of course, no-one is more aware of that than Gudrun Klinker, as she looks ahead at so many different avenues.

A compelling experience – but there are risks

AR has entered the mainstream beyond industry – recently boosted by the hit game Pokémon Go, launched in 2016. Here, virtual fantasy creatures hide in various places, and a gamer only sees them when he or she is nearby. Then, the figures show up on the smartphone display integrated into the user’s real surroundings. And that, in a nutshell, explains the appeal of augmented as opposed to virtual reality (VR): AR allows people to remain present in the real world instead of shutting it out. “A Pokémon apparently standing on my own coffee table in my actual surroundings has a much greater emotional effect on me than if I were to see it in an unknown setting,” Klinker explains.
At the same time, the compelling nature of AR can actually make it hazardous when users lose sight of the world around them. This is already an issue with pedestrians who fail to notice red lights because they are staring at their smartphones – now a regular occurrence on our streets. The German city of Augsburg has already reacted by installing ground lights aimed at preventing accidents when people are staring downwards onto their phones – an idea that has attracted major interest worldwide. Driver assistance systems are another example here. Whether their purpose is navigation, communication or entertainment – these systems must be laid out in a way that avoids distracting drivers. And we have to know when to step in. “It’s the same as with back-seat drivers – some people would be well advised to hold their tongue,” observes Klinker drily. The key question is: how much virtual information can someone absorb without losing concentration on their actual surroundings? Or, in short: how much distraction can a person take?

Keeping users keen

As part of this research interest, Gudrun Klinker is currently working on the use of eye tracking in head-mounted displays. This involves recording and evaluating eye positions and movements of people wearing head-mounted displays. “We are improving eye tracking methods and trying to establish as precisely as possible where the user’s eyes are positioned behind the glasses and what they are looking at,” explains Klinker. This is essential for the glasses to display information in exactly the right place to overlay the physical objects it relates to. And that has significant consequences for the user. Tracking a user’s eyes enables them to control a program without a mouse, for instance – and let it know what they find interesting. ▷

Picture credit: Eckert

> *"Now that the technology is so far advanced, we can focus on the psychological element too – which also means bringing in experts from other fields who were previously deterred by all the bits and bytes flying around."* *Gudrun Klinker*

"This then raises the question: when are users looking for this type of added information and when are they just moving their eyes for no particular reason?", specifies Klinker. For that reason, AR research involves not only computer scientists, but also educators, psychologists and methodologists, to determine how best to convey the information. And this directly leads to further questions: How can someone's interest be aroused? And how can they be motivated to do something? "The nature of Homo ludens (Latin for "playing human") is to develop new things by experimenting; through playful exploration. So it's a matter of fueling this playful interest," Klinker explains.

Here, she mentions storytelling and gamification – two key tools in this area. And researchers have long since been working on other AR devices besides the glasses. "The limited field of view means that glasses detract quite a bit from the reality of the experience, so you have to "bend" your senses to compensate. That's why everyday objects hold great promise for AR." Klinker is currently working with nutritional scientists and doctors to develop a drinking glass for older people, for example. This intelligent glass registers how much someone has drunk and gives them playful reminders to drink more if need be – for instance, by ceasing to show photos of their grandchildren on its integrated display. Klinker is also looking forward to further interdisciplinary research: "Now that the technology is so far advanced, we can focus on the psychological element too – which also means bringing in experts from other fields who were previously deterred by all the bits and bytes flying around." *Gitta Rohling*

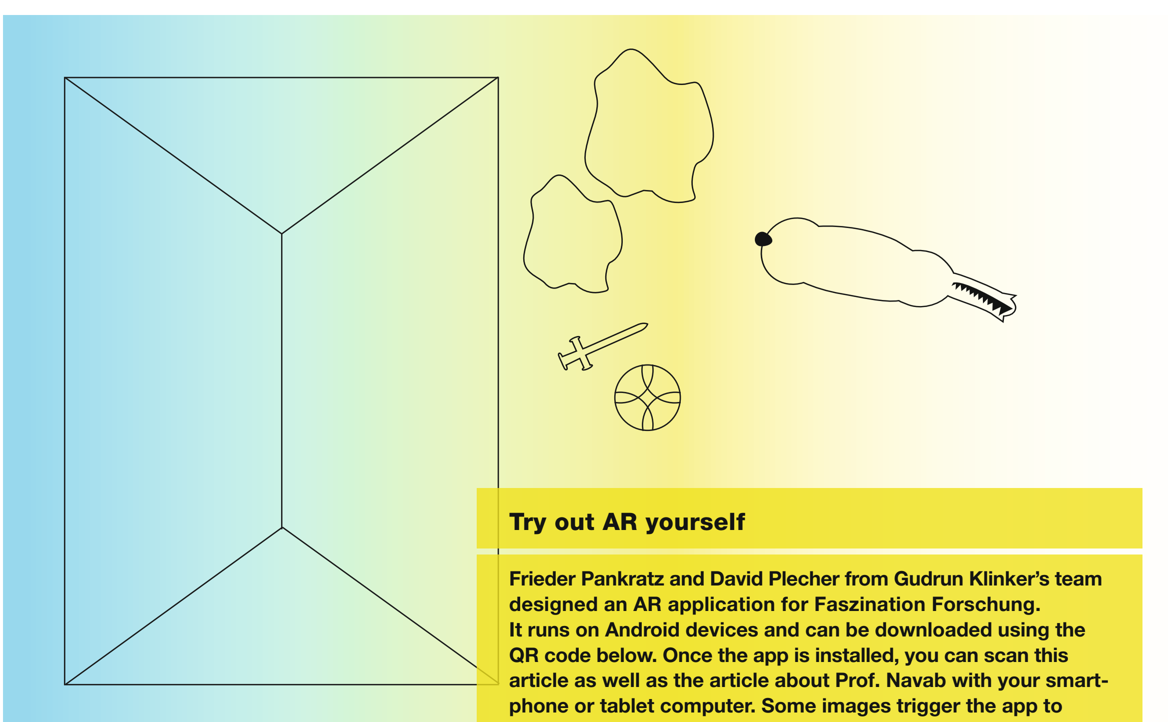

Try out AR yourself

Frieder Pankratz and David Plecher from Gudrun Klinker's team designed an AR application for Faszination Forschung. It runs on Android devices and can be downloaded using the QR code below. Once the app is installed, you can scan this article as well as the article about Prof. Navab with your smartphone or tablet computer. Some images trigger the app to show additional information, overlaid in 3D above the images.

Graphics: ediundsepp

Authors

Claudia Doyle studied biochemistry in Leipzig, Germany. Between 2010 and 2013 she worked in the press office of several institutes of the Max Planck Society. In 2013 she started freelancing as a science journalist. She is currently studying journalism at the Deutsche Journalistenschule (German School of Journalism) in Munich. Her articles focus mainly on ecology and health.

Birgit Fenzel studied German language and literature, philosophy and educational sciences before starting her career in journalism. After finishing university and completing a year of voluntary service, she worked as an editor for a daily newspaper, also completing freelance assignments for other public service broadcasters. She worked on scientific articles for the ddp news agency and was most recently employed as an editor at the Max Planck Society.

Gitta Rohling, M.A., works under the brand name Auftrag Wort as a writer, editor and PR consultant. She supports mainly companies operating in technical and scientific fields in the areas of PR, press work and text.

Dr. Brigitte Röthlein has been working for many years as a science author for magazines, TV and radio broadcasting and for newspapers. She holds a diploma in physics and a Ph. D. in communication science, education science and history of natural sciences. Her main interest lies in basic research.

Tim Schröder is a science journalist. After working as an editor for the daily "Berliner Zeitung" newspaper he became a freelancing author in Oldenburg, Germany. He regularly contributes to the "Frankfurter Allgemeine" (Sunday edition), "Neue Zürcher Zeitung" and "ZEIT" newspapers. His specialist areas are basic and applied research, energy and the environment.

Monika Weiner is a freelance journalist. After studying geology at the University of Munich she started a career in science journalism working for different newspapers, popular science magazines and broadcasting stations. She took part in the Knight Science Journalism Program at MIT and was Journalist in Residence at the Kavli Institute for Theoretical Physics, University of California. Her reporting focuses on the impact of discoveries and new technologies.

Dr. Karsten Werth is a freelance science journalist based in Munich. He studied history and American studies, writing his doctoral thesis on the US space program of the 1960s. He went on to gain experience with various industrial and media enterprises in the US, Canada and Germany, including as Editor-in-Chief at a PR agency.

Imprint

Faszination Forschung
Technische Universität München's Science Magazine
Sponsored by the Excellence Initiative of the German federal and state governments

Publisher
Prof. Dr. Dr. h.c. mult. Wolfgang A. Herrmann,
President of Technische Universität München
Editorial Office
Dr. Christine Rüth (Editor-in-Chief)
Translations
Capella & McGrath GmbH, Munich
Copy-editing
Baker & Harrison, Munich
Design
ediundsepp Gestaltungsgesellschaft, Munich
Authors in this issue
Claudia Doyle, Birgit Fenzel, Gitta Rohling, Dr. Brigitte Röthlein,
Tim Schröder, Monika Weiner, Dr. Karsten Werth
Address of the Editorial Office
Technische Universität München, Corporate
Communications Center, 80290 Munich
E-mail
faszination-forschung@zv.tum.de
Printing
Druckerei Joh. Walch GmbH & Co. KG, Augsburg
Circulation
61,500
ISSN: 1865-3022
Frequency
Twice a year
Responsible for editorial content and advertisements
Dr. Ulrich Marsch
Photographers
Kurt Bauer, Astrid Eckert, ediundsepp, Andreas Heddergott,
Magdalena Jooss




VIVE

Picture credit: Bauer

Beyond the Limits **of our Eyes**

Link
campar.in.tum.de

Nowadays, surgeons can access a wealth of information about their patients. X-rays, CT scans and MRI images reveal structures inside the body that are usually not visible to the naked eye. In the middle of an operation, however, physicians have to rely on their own senses. Prof. Nassir Navab wants to change this by bringing augmented reality into the operating theater.

The video image of the surgery is augmented with a 3D model of the radioisotope labeled sentinel lymph node.

Reference structure to determine the position of the patient in the camera image.

A gamma probe is used to map the sentinel lymph node by detecting radiation emitted by the injected radioisotope. This data is used for reconstruction and then visualized in real time in the video image on the computer screen.

Reference structure to determine the position of the gamma detector relative to the patient.

Graphics: ediundsepp, Picture credit: Bauer

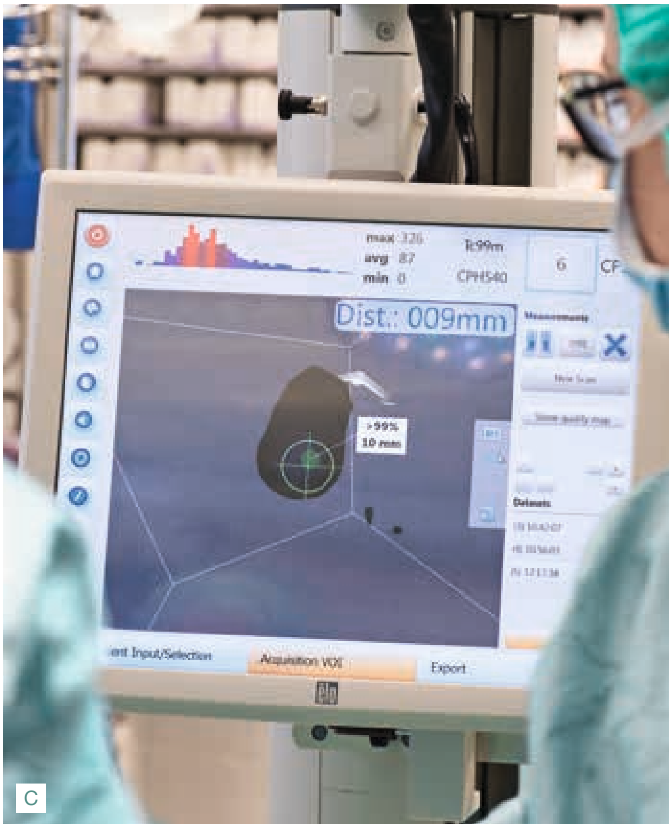
max 326
avg 87
min 0
Tc99m
CPH540
6
Dist.: 009mm
Measurements
New Scan
>99%
10 mm
Datasets
Acquisition VOI
Export

C


Dist.: 078mm
CPS

A


Dist.: 641mm
Measurements
2872
New Scan
Show quality map
>99%
641 mm
Sentinel lymph node
Datasets
(1) 10:32:13
(2) 10:35:15
(3) 10:42:07
Labels
SURGICeye
Patient Input/Selection
Acquisition VOI
Export

D

Scanning the body area with a gamma probe (A) and tracking this action with a special camera and navigation system (B) allows augmentation of the video with a reconstructed 3D image. The surgeon can see which areas are already scanned (black). Based on the gamma data, a 3D model of the sentinel lymph node is calculated (C) and correctly superimposed on the body to guide the surgeon (D) for the surgical resection.

Claudia Doyle

Jenseits des Sichtbaren

Chirurgen haben heutzutage Zugriff auf eine Fülle von Informationen. Röntgenbilder, CT-Scans oder MRT-Aufnahmen machen Strukturen im Inneren des Körpers sichtbar, die menschlichen Augen normalerweise verborgen bleiben. Doch während einer Operation sind die Ärzte wieder allein auf ihre Sinne gestellt. Prof. Nassir Navab will das ändern, indem er Augmented Reality in den Operationssaal holt.

Die erste Anwendung von Augmented Reality, die routinemäßig von Ärzten eingesetzt wird, ist ein von Navab und den medizinischen Partnern der Nuklearmedizin und Frauenheilkunde am Klinikum rechts der Isar entwickeltes kleines Messgerät namens declipseSPECT. Es handelt sich um eine Miniaturausgabe einer Gammakamera, die der Chirurg während der Operation benutzen kann. Mit normalen Gammakameras ist das nicht möglich. Die mannsgroßen Maschinen sind meist in einem eigenen Raum installiert und nicht flexibel.

Zur Anwendung kommt declipseSPECT beispielsweise bei der Lymphknotenbiopsie. Hat ein Patient einen Tumor, dann prüfen Chirurgen mit Hilfe dieser Methode, ob der Krebs bereits metastasiert ist. Finden sich Krebszellen in dem Lymphknoten, der im direkten Lymphabfluss des Tumors ist, dann ist das wahrscheinlich.

Um genau den richtigen Lymphknoten chirurgisch minimalinvasiv entfernen zu können, spritzen die Mediziner dem Patienten eine sehr geringe Dosis eines radioaktiven Mittels. Dieses reichert sich im Wächterlymphknoten an und kann mit der Gammakamera detektiert werden.

Die Vorteile liegen in der kombinierten Darstellung der Daten. Zum einen ein Live-Videobild des Patienten, aufgenommen von einer über dem OP-Tisch installierten Kamera. Zum anderen die Lokalisierung der Lymphknoten, gemessen mit der Gammakamera. Der Chirurg weiß dann genau, wo er den richtigen Wächterlymphknoten findet.

Intensiv erforscht wird auch der mögliche Einsatz von sogenannten Head-Mounted-Displays. Diese in der Vergangenheit klobigen Brillen, ausgestattet mit Sensoren, Kameras und Bildschirmen, sind bisher vor allem aus der Unterhaltungselektronik bekannt. Doch auch ein Einsatz im Operationssaal ist denkbar. Die Chirurgen sollen dann den Patienten vor sich sehen und sich gleichzeitig je nach Bedarf Röntgenbilder, CT-Scans oder andere Informationen in ihr Blickfeld projizieren lassen. Möglich ist vieles. Aber die richtige Information zum richtigen Zeitpunkt anzuzeigen, das ist die große Herausforderung. □

Graphics: ediundsepp, Picture credits: Bauer

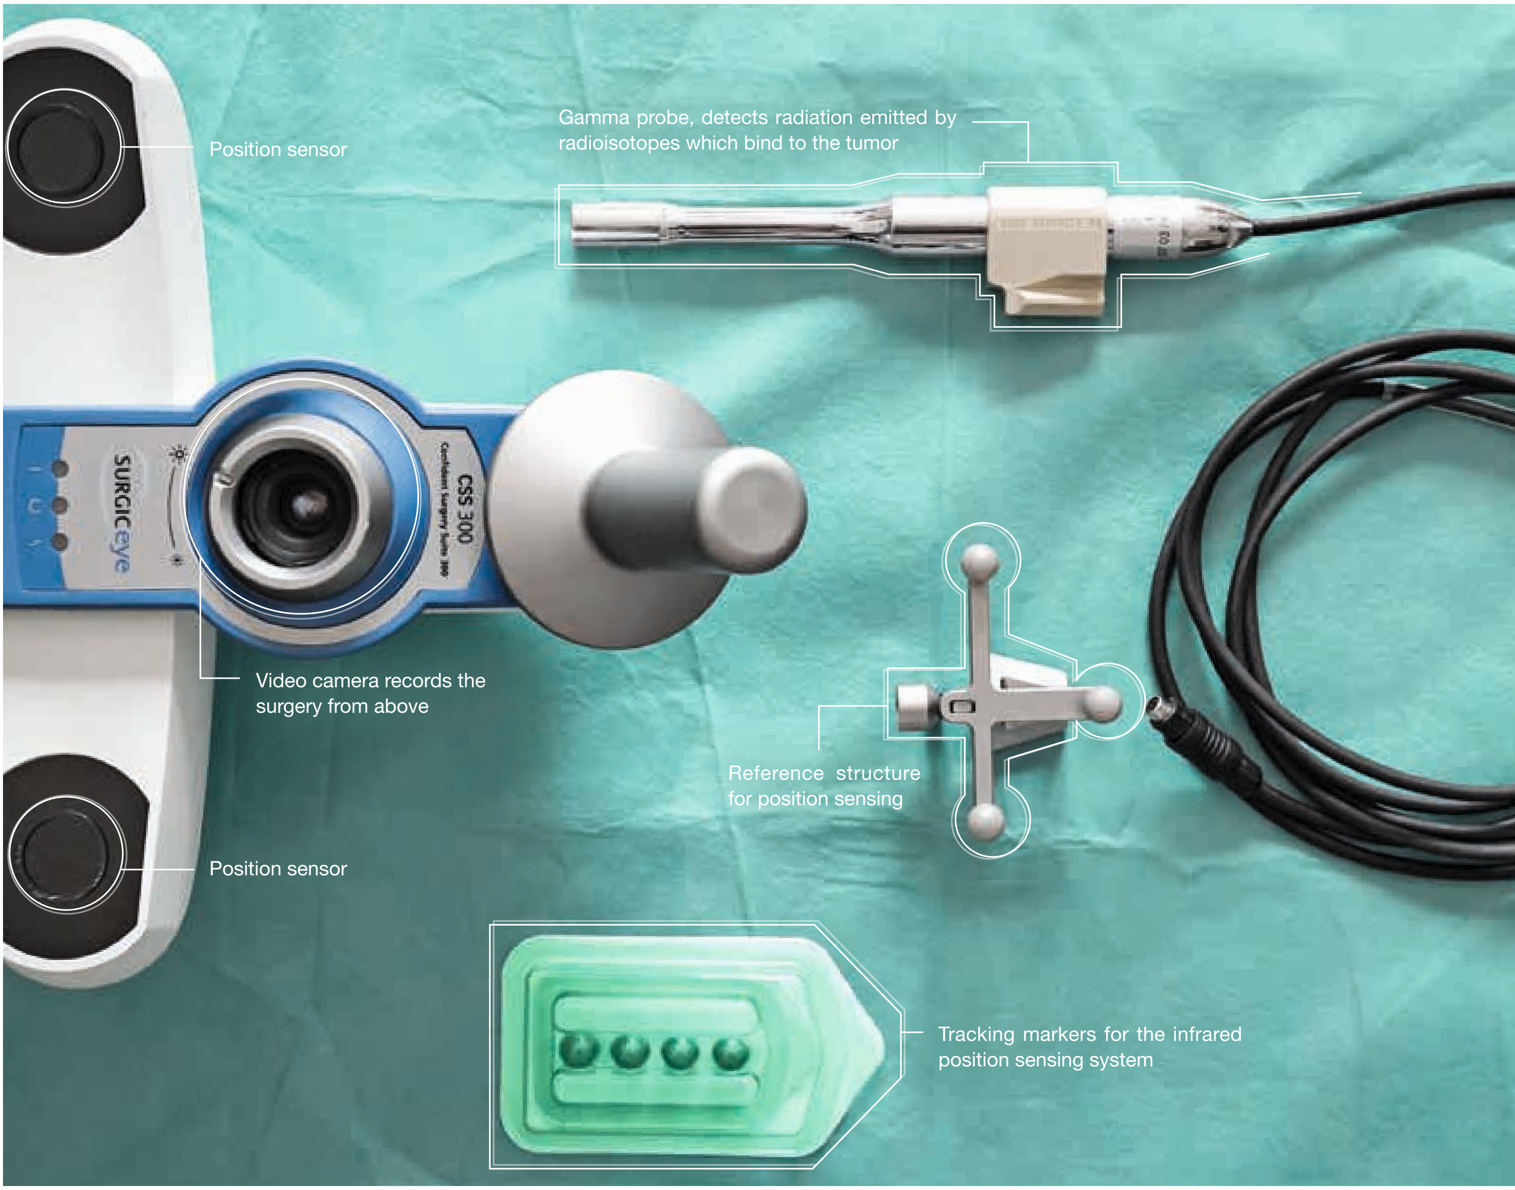

Nassir Navab invented declipseSPECT in collaboration with clinical partners specializing in nuclear medicine.

Graphics: ediundsepp, Picture credit: Bauer

With a steady hand, Dr. Stefan Paepke moves the finger-sized measuring instrument over the skin of his patient. The surgeon traces slow circular movements between the armpit and the breast. An image recorded by a camera can be viewed on a screen beside the operating table. Black spots that become steadily larger slowly start to appear. These are the areas that Paepke has already recorded with his measurement device.

Operating theater 3 in the gynecology department of Munich's Klinikum rechts der Isar looks no different to the facilities in many other hospitals. But there are a few important differences, such as a small device called declipseSPECT. This is a miniature version of a gamma camera – an instrument often used in nuclear medicine to detect cancer cells.

It takes Paepke around one minute to scan the entire area. A 3D image is reconstructed from gamma rays acquired from the area marked in black, wherein a green-colored region becomes visible. This is the sentinel lymph node, which lies in the direct drainage path of the primary tumor. If the tumor has formed metastases, the answer can be found in this lymph node. Paepke can use this image to pinpoint the exact location of the lymph node and determine its subcutaneous depth. He can then reach for his scalpel and extract the correct lymph node a short time later.

The working group led by Nassir Navab invented declipseSPECT in collaboration with clinical partners specializing in nuclear medicine. The system was designed, developed and tested from the outset in close collaboration with Stefan Paepke in the gynecology operating theater of Klinikum rechts der Isar. Since then, the device has been used in hundreds of operations at several different clinics. declipseSPECT thus holds the distinguished position of being one of the first medical applications of augmented reality which has already achieved market maturity and is in regular use in clinical settings. According to Navab, it will by no means be the last.

Nassir Navab has held the Chair of Computer-aided Medical Procedures and Augmented Reality at TUM since 2003. He stresses that he does not want to replace surgeons, but rather make their job as easy as possible. “Surgeons nowadays use only a small percentage of the information that is actually available to them,” explains Navab.

One reason for this is the rapid rise over recent decades in the choice of medical measuring instruments available to clinicians. In earlier times, physicians could only take basic measurements like blood pressure, pulse rate or temperature. But now, technology in many different forms, including X-ray, ultrasound, CT scans, MRI, PET, SPECT and photoacoustic procedures, provide deep insights into what is happening inside patients' bodies. This means that specialists can, for example, visualize brain activity, see where cancer cells are growing and understand more about inflammations.

“Today's surgeons usually have too much, rather than too little, information at their disposal,” affirms Dr. Jörg Traub. Since 2008, the former Ph.D. student of Nassir Navab has been managing director of the company SurgicEye, which has been overseeing the approval, market rollout and further development of the hand-held gamma camera. “These huge volumes of data have to be filtered and then presented at exactly the right time.” This means that information should ideally be available immediately before or during the critical phases of an operation.

To achieve this, Navab's team miniaturized the SPECT acquisition setup. Normally, SPECT acquisition devices are man-sized imagers installed at a fixed location. The patient therefore has to be moved from the operating table to the gamma camera. The great benefit of declipseSPECT is the flexibility that comes with a tracked hand-held device. ▷

This RAMP Head Mounted Display (HMD) system was developed at Siemens Corporate Research at the end of 1990s.

Virtual reality becomes augmented reality when the VR set is combined with a camera, which records the real environment to be displayed inside the HMD. Additional tracking markers on the headset help determine the user's position more accurately.

Even though it is small and portable, declipseSPECT is very similar to fixed gamma cameras in terms of functionality. The patient is first injected with a radioactively labeled molecule. This is dispersed throughout the body and binds to certain tissue structures. The camera then detects the radiation emitted by this molecule.

In operating rooms without declipseSPECT, surgeons rely on acoustic signals from a type of Geiger counter, which emits peeps that go from fast to slow and from quiet to loud when it detects radioactivity. This does not tell the surgeon how deep the lymph node is located beneath the surface, however.

For patients, it is vitally important to have what is called the sentinel lymph node removed. As the closest lymph node to the tumor, it would be the first to be affected by a possible metastasis and has to be examined. The head and throat regions have a particularly high number of lymph nodes close to the tumor, and acoustic methods are of limited use when a tumor is being examined here. The problem is that it is not always clear whether the radioactivity is coming from the sentinel lymph node or from the injection point of the radioactive substance. declipseSPECT is able to throw light on this question. The surgeons at Klinikum rechts der Isar are extremely pleased with the new instrument, and welcomed the arrival of a computer scientist in their workplace with open arms.

When Nassir Navab moved to Munich in 2003, he immediately made it clear that his staff had to have direct contact with the surgical teams, preferably on a daily basis. Without further ado, he took over the empty basement rooms at the inner-city clinic of LMU Munich. He and his staff got to work straight away, continuously asking the surgeons what they need, what hampers their work, what could be improved and what they would like to add to their wishlist.

After six months, they were driven out of their underground vault. The hospital management had been won over by the work of Navab and his team in the meantime, however, and so they were offered a smart suite of offices right next to the operating theater in the new clinic building.

"I was delighted with the opportunity to set up my lab inside the hospital," affirms Navab. While working for his previous employers, there was much less interaction between engineers and clinicians. Just one meeting per year was organized to discuss surgeons' wishlists and identify the potential benefits of new technologies. "Twelve months later, we came back to the clinic and presented our inventions to the physicians, only to be told more often than not that it was not what they were looking for," relates Navab.

A short time later, Navab opened a second lab in TUM's Klinikum rechts der Isar, where he now works with physicians from fields as diverse as urology, surgery and nuclear medicine. They are all interested in finding out how augmented reality can make their work easier. The computer scientists and the clinicians are constantly discussing clinical needs and bouncing new ideas around for how to use advanced technology to satisfy them. ▷

Augmented reality optical see-through headsets have a transparent display and overlay the reality that the user views through the glasses with additional information.

Tests with a see-through augmented reality HMD: While cameras on the HMD allow for the augmentation of the patient view, infrared markers allow the HMD to be localized and co-registered with surgical instruments. Co-registration means that position and orientation of real and virtual objects match.

Graphics: ediundsepp, Picture credits: Bauer

“Surgeons nowadays use only a small percentage of the information that is actually available to them.”

Nassir Navab

Try out AR yourself
Frieder Pankratz and David Plecher from Gudrun Klinker’s team designed an AR application for Faszination Forschung.
It runs on Android devices and can be downloaded using the QR code below. Once the app is installed, you can scan this article as well as the article about about Prof. Klinker with your smartphone or tablet computer. Some images trigger the app to show additional information, overlaid in 3D above the images.

Prof. Nassir Navab

A computer scientist embedded in the surgeon team

Nassir Navab was born in Iran and went to university in France. He initially studied mathematics and physics before moving on to computer science and system control. During his doctoral thesis, he performed research on a robot destined for a Mars discovery mission. After completing his doctorate, he went to work in the USA, first at Massachusetts Institute of Technology and then in the research department of Siemens in Princeton. In 2003, he relocated to Munich with his wife and two children, where he set up the Chair of Computer-aided Medical Procedures and Augmented Reality. One thing he regards as crucial to his work is close proximity to the physicians who will one day use the technology he is developing. “Without constant feedback from the operating theater, what we develop will not meet real-world requirements,” explains Navab.
Nassir Navab has been settled in Munich for 13 years now. He loves the city’s cosmopolitan and friendly vibe, relishing the fact that he can enjoy a Persian concert and then spend a chilled evening in a Bavarian beer garden – all in one day!

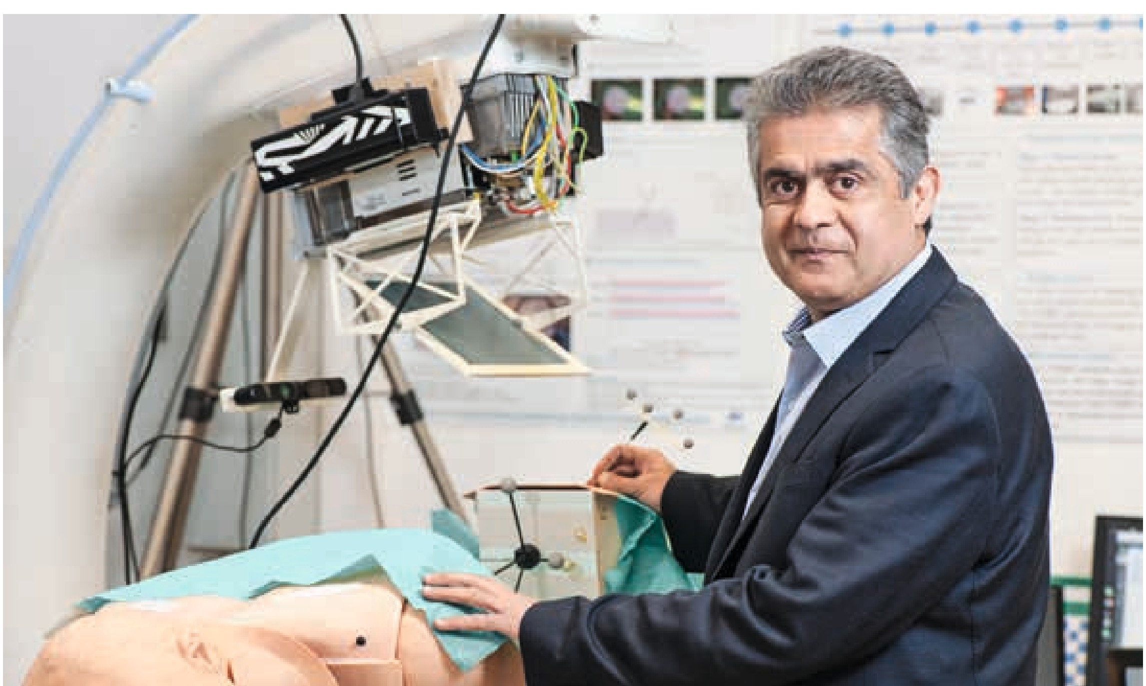

This interactive augmented 'magic mirror' allows everyone to get familiarized with their own anatomy. The left hand moves the AR local window into the anatomy and the right hand changes the level of details from bony structures, to organs and vascular system and finally the surrounding fat and muscles. The system can be used for anatomy education, patient information and rehabilitation.

"For me, it is not enough to have just one clinical partner," says Navab. Every surgeon has undergone different training and has a different way of thinking. It is therefore important to get feedback from as many people as possible since the ultimate objective is to get as many clinics as possible to use the new inventions. This is why Nassir Navab set up a second lab at Johns Hopkins University on the east coast of the USA around three years ago. He spends around one week per month stateside, asking questions, listening to feedback and developing new ideas.

One of the more long-term projects that he has been working on with his team focuses on a head-mounted display (HMD) for surgeons. These oversize spectacle frames equipped with cameras, screens and position sensors are a common sight in the world of entertainment. Just about every major player specializing in gaming consoles and video games has brought out its own version of an HMD by now.

There are two types of head-mounted displays. What are known as optical see-through HMDs have a transparent display, so what users see is their actual environment. Then there are video see-through HMDs with a non-transparent display, where two cameras film the user's surroundings and present this image to the wearer. It is possible to project additional information onto both types of display with AR technology.

At present, HMDs are mostly of interest as a training aid. The technology lets physicians practice certain procedures under realistic conditions so that they are better prepared when they have to perform them during an operation. HMDs also play an important role when it comes to planning operations. The scientists are currently working on ways to plan neurosurgical interventions like deep brain stimulation using HMDs. Bringing this technology to a live operating theater remains the Holy Grail, however. A surgeon wearing an HMD would not only be able to see the patient in front of them, but also be able to call up useful images like X-rays, MRT scans or ultrasounds on demand. "It is of course always possible to walk to a monitor and view the images there," says Ulrich Eck, head of the NARVIS working group, "but our vision involves projecting them directly onto the corresponding area of the patient's body."

The main challenges for the researchers at present involve processing and visualizing the data because of the massive computational tasks that have to run in the background. Feedback from physicians is crucial here because they can specify what information is important at which point during an operation.

"Augmented reality is often viewed as a gaming gimmick, and that is true in some ways," admits Jörg Traub. At the end of the day though, the technology is too useful to be limited to fun and games.

Claudia Doyle

Picture credit: Bauer

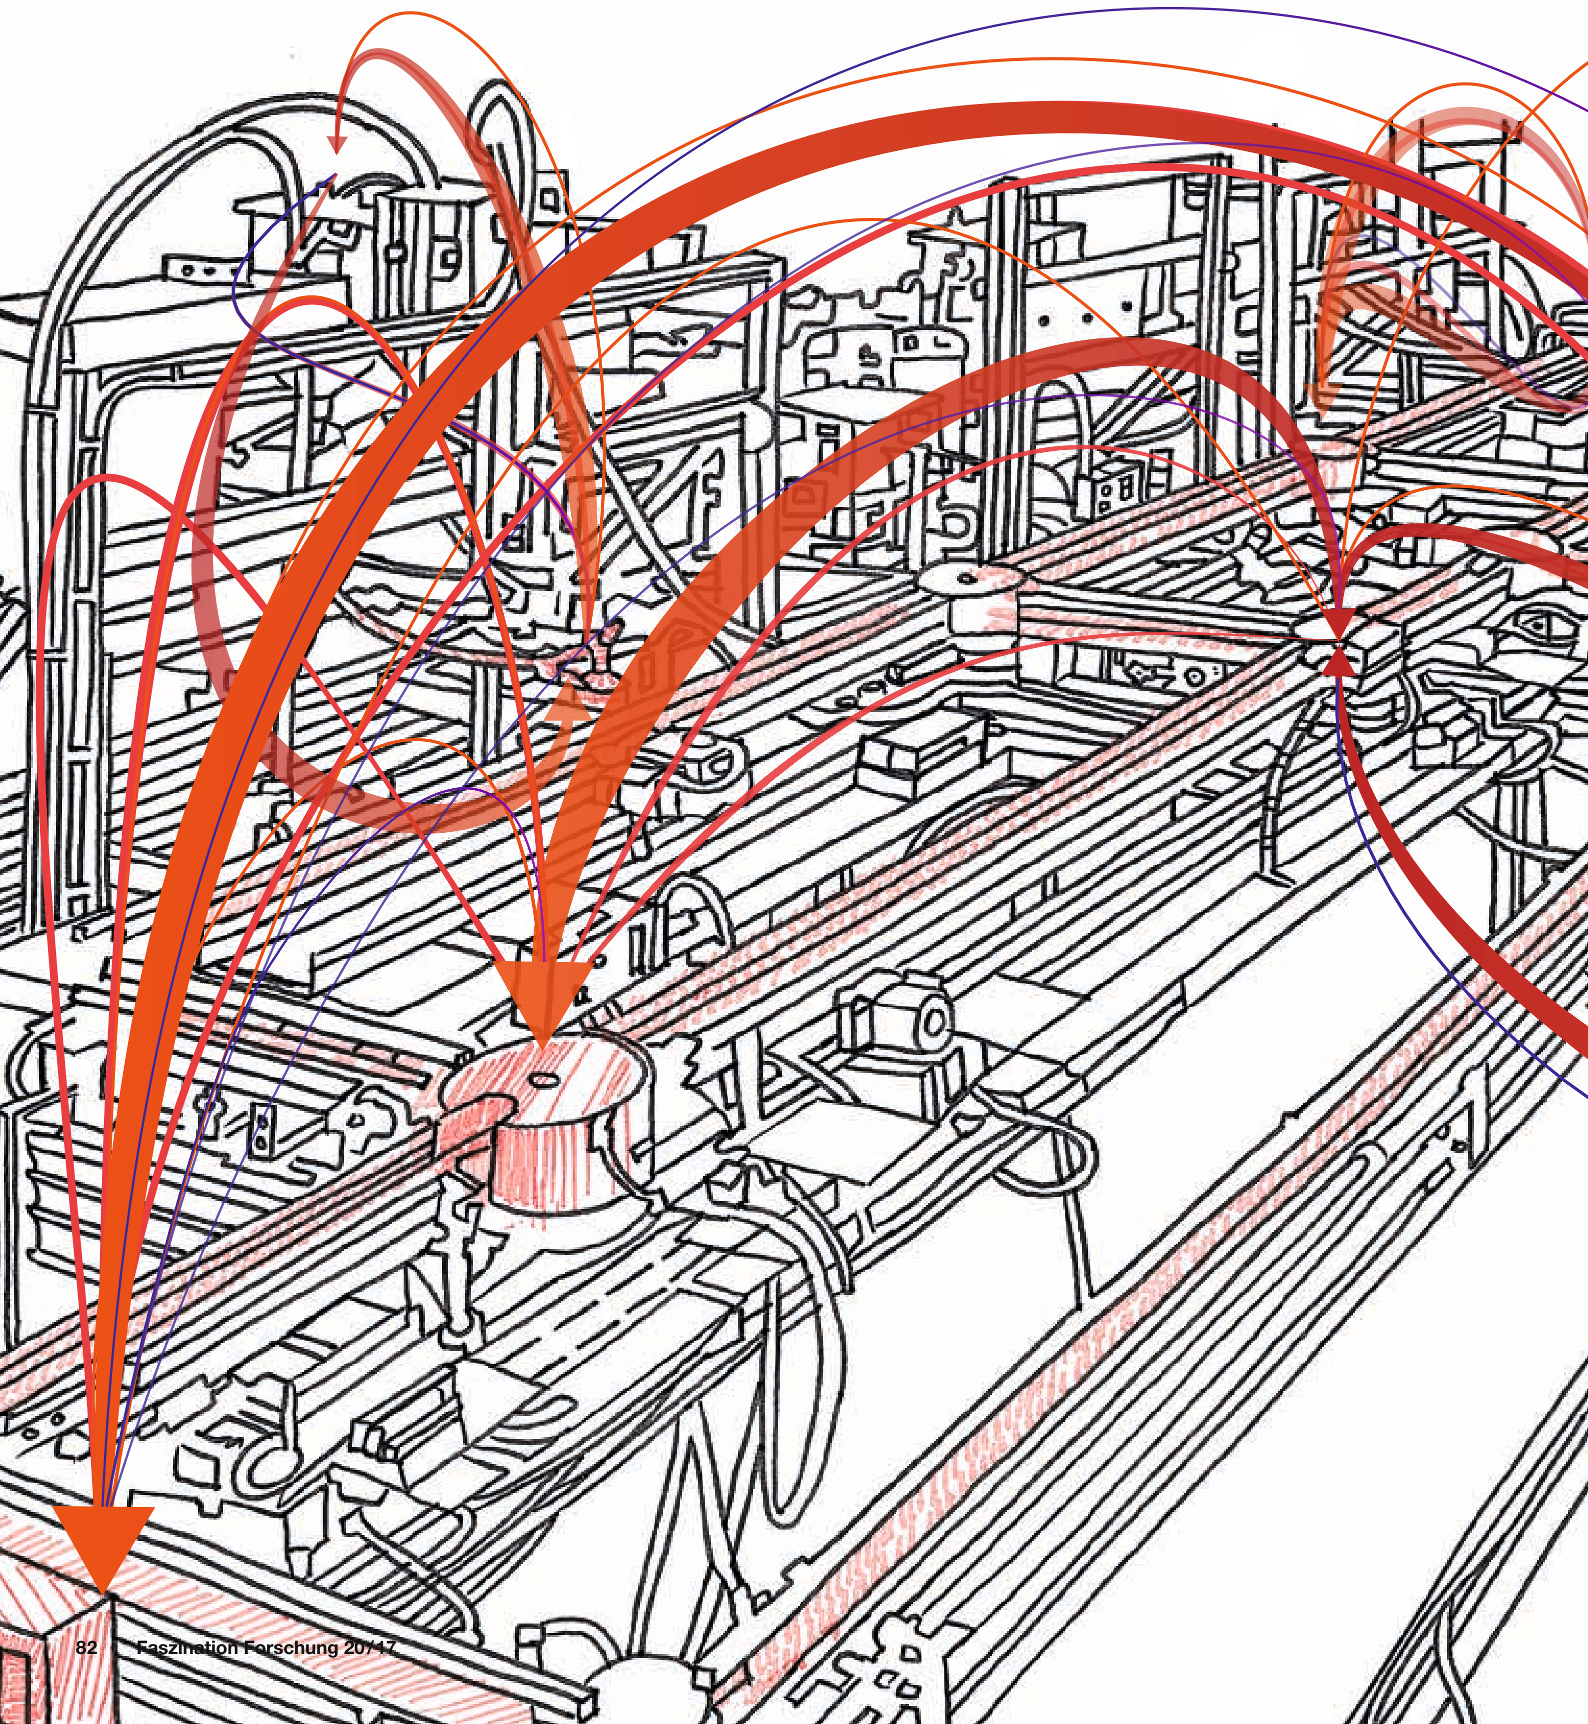

Overcoming Communication Barriers

Steam engines, assembly lines and IT have each played a role in enabling and shaping industrial manufacturing. Now, digital technologies are triggering a fourth industrial revolution: Smart Factories – often referred to as Industry 4.0 – are set to boost the speed, flexibility and efficiency of production in the future. Their ability to do so, however, hinges on seamless communication – between operating systems, machine components, equipment and organizations. Yet, the situation today is more like the Tower of Babel. High time, then, for a new, multidisciplinary approach.

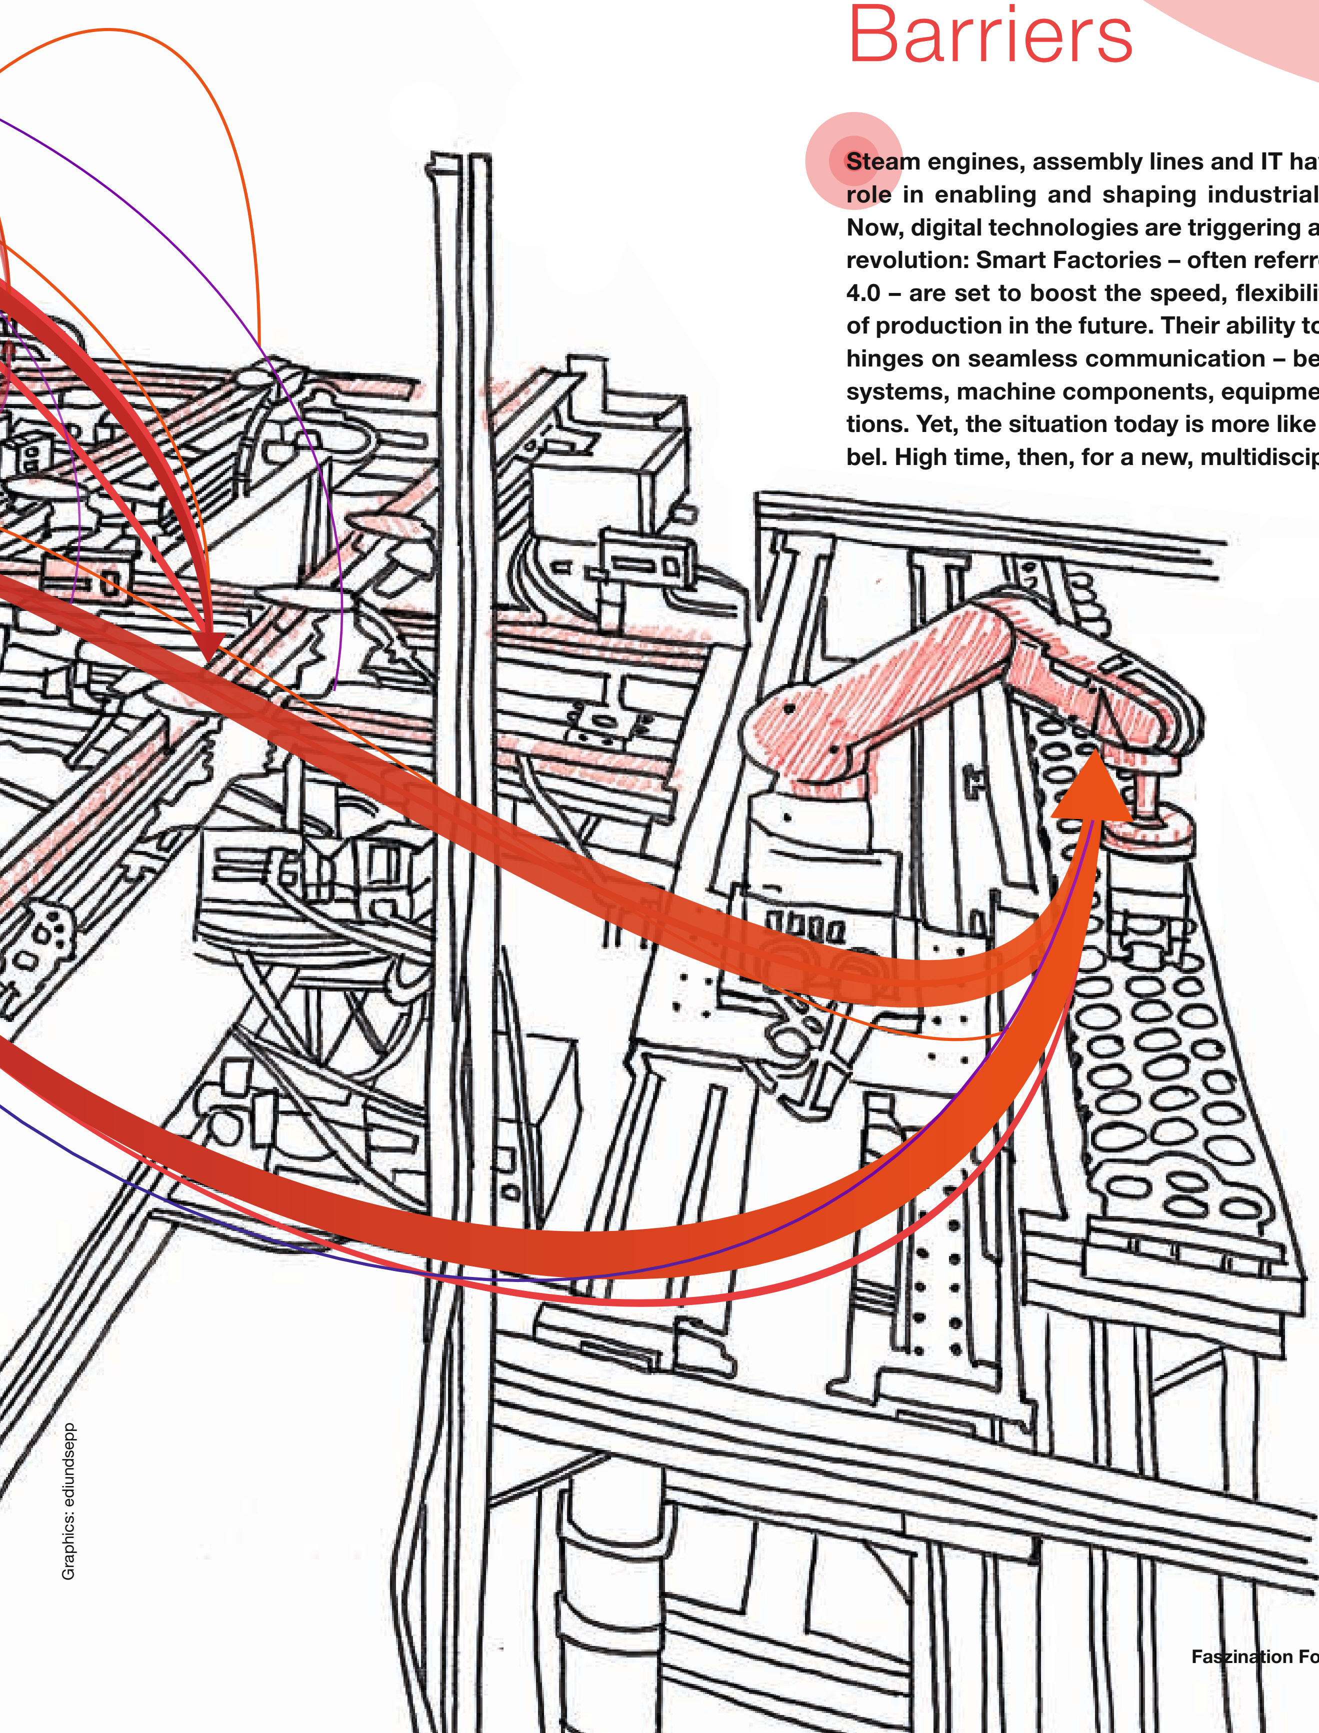

Graphics: ediundsepp

Monika Weiner

Modularität überwindet Sprachbarrieren

Anlagen, die schnell auf Kundenwünsche reagieren, sich selbst steuern, effizient produzieren, stets optimal ausgelastet sind und Produkte in höchster Qualität liefern – Industrie 4.0 macht's möglich. Zumindest theoretisch. Praktisch steckt die Vernetzung von Verfahren, Prozessen und Produkten noch in den Kinderschuhen. Noch fehlt es an Standards und Plug-and-Play-Lösungen; der Fluss von Daten und Informationen kommt häufig schon an den Schnittstellen zwischen Maschinen zum Erliegen; Softwareprobleme machen das Austauschen eines Förderbandes oder Barcode-Scanners zum unüberwindbaren Hindernis.
Prof. Birgit Vogel-Heuser erforscht am Lehrstuhl für Automatisierung und Informationssysteme der TUM, wie sich der Informationsfluss verbessern lässt – zwischen Anlagenbau, Elektrotechnik und IT, zwischen den einzelnen Komponenten einer Anlage, aber auch zwischen Auftragseingang und Systemsteuerung. Damit überschreitet sie Grenzen – sowohl zwischen Fachdisziplinen als auch zwischen Technologien.
Der Schlüssel zum Erfolg der Industrie 4.0 ist nach Ansicht der Ingenieurin ein modularer Aufbau von Produktionseinheiten: Jeder Greifer, jedes Förderband, jede Abfüllstation wird als eigenständige Einheit betrachtet, als Modul, das eine gewisse Intelligenz besitzt, das weiß, was es kann, und das sich im ständigen Austausch mit den anderen Einheiten befindet. Dass dieses Konzept funktioniert, beweist der MyJoghurt-Demonstrator im Erdgeschoss des Instituts: Die Anlage steuert sich selbst. Geht ein Auftrag ein, prüft sie ohne menschliches Zutun die Verfügbarkeit der Zutaten und die Kapazität der einzelnen Module. Das System ist so flexibel, dass während des laufenden Betriebs Softwareupdates durchgeführt und Komponenten ausgetauscht werden können.
Der modulare Ansatz hilft Unternehmen, die in Zeiten immer kürzerer Produktzyklen schnell auf die Anforderungen des Marktes reagieren müssen, und auch Zulieferern und Dienstleistern, die unter dem Druck der Globalisierung neue Möglichkeiten suchen, ihre Effizienz zu steigern. □

Link
www.ais.mw.tum.de

Electric motors whir, conveyor belts emit a low hum. With a quiet click, a robotic hand places a small glass container on the belt. The yogurt jar starts its trip by heading to a switch, which guides it around the corner, before a second belt picks it up and transports it past a barcode scanner towards the filling unit. Here, a light barrier registers the incoming jar. The conveyor stops, and little red balls from a plastic container immediately tumble into the opening. After two seconds, the jar advances again to the next stop, where water then pours in. At a third station, yellow balls are added to the mix. "These might be strawberry pieces or chocolate chips – we don't like using fresh ingredients in the lab, because they're perishable," explains Prof. Birgit Vogel-Heuser. The Chair of TUM's Institute of Automation and Information Systems keeps a critical eye on the processes in the "MyJoghurt" demonstrator. Everything runs like clockwork: robotic hands and conveyors, switches, scanners, light barriers and filling units are perfectly synchronized.
What you don't see, however, is the fact that the system is running itself. When an order comes in – a customer requests chocolate yogurt with strawberry topping, for instance, which has not been produced or designed before – requests are sent to the various modules: Are sufficient amounts of these ingredients available? Are the filling unit's tubes wide enough for chocolate chips and strawberry pieces? Can the conveyors transport the jars fast enough to execute the order on time? Is the price right? Once all the modules have checked and confirmed their capabilities, production can begin. A minor miracle of predictive planning and efficiency.
We could, in fact, be looking at the future of industrial production. MyJoghurt is a demo system for Industry 4.0 – the fourth industrial revolution following in the footsteps of the steam engine, the conveyor belt and IT. "Industry 4.0 has now become a buzzword. Everyone wants it, but everyone means something different by it. There are currently few standards and hardly any plug and play Industry 4.0 solutions that can smoothly integrate equipment from different manufacturers," clarifies Vogel-Heuser, explaining that the definition of Industry 4.0 is thus relatively abstract. It entails connecting data, goods flows, production processes and people in order to improve the speed, flexibility, overall equipment efficiency and quality of manufacturing in the future.
The MyJoghurt system shows how these aims can be realized. Researchers are using the demonstrator to develop new concepts – for instance, to enable communication between separate areas or units. "We can run software updates and even replace components with new ones featuring different connectors in live mode without needing to completely shut down the whole system," outlines the engineer. "As it stands, that is almost impossible in current industrial production plants." ▷

MyJoghurt – a demo system for Industry 4.0

The MyJoghurt demonstrator is intended to showcase the possibilities of Industry 4.0 and how they can best be leveraged. A key focus lies on networking geographically distributed production facilities which may even belong to different companies. This way, the filling unit may collaborate with the packaging system or purchasing organization on an automated basis, for example. MyJoghurt consists of a yogurt filling system located at TUM, a yogurt processing facility in Magdeburg and another one in Stuttgart, a lid labeling facility in Hamburg and other, virtual systems run by various partners. Prof. Birgit Vogel-Heuser was responsible for initiating and setting up MyJoghurt, with five other research partners also participating while the project continues to grow.

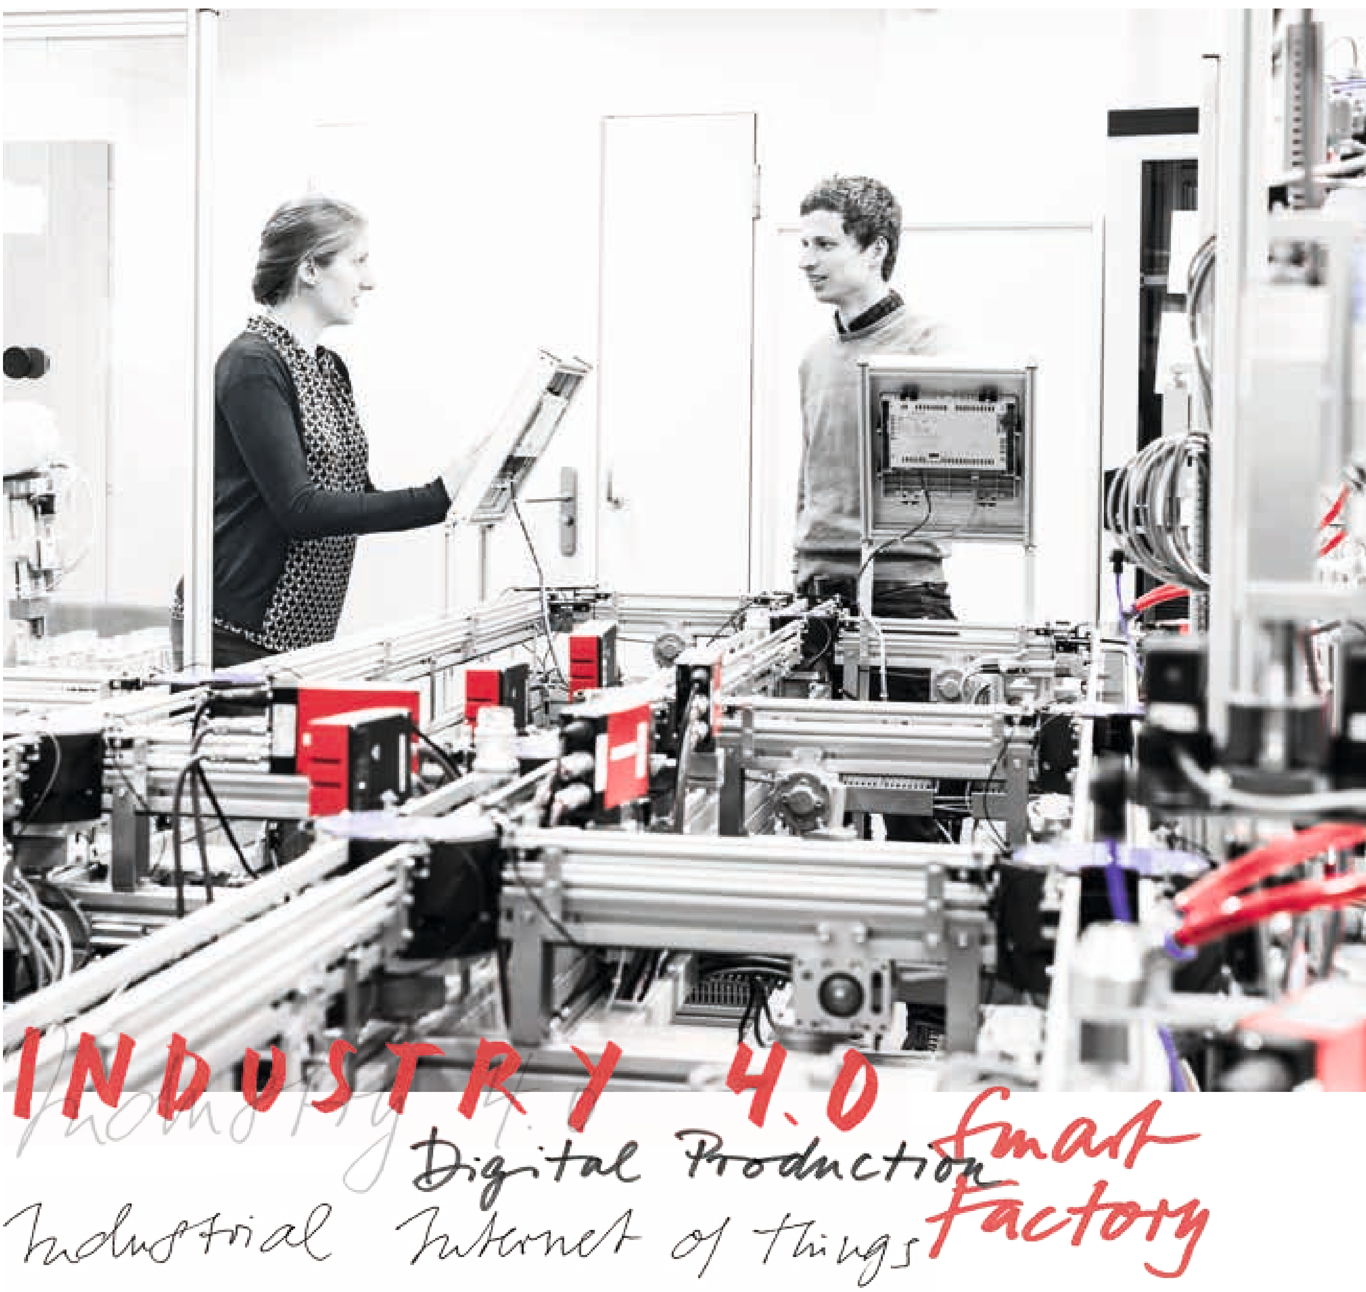

Graphics: ediundsepp; Picture credit: Bauer

Material supply
with metering device

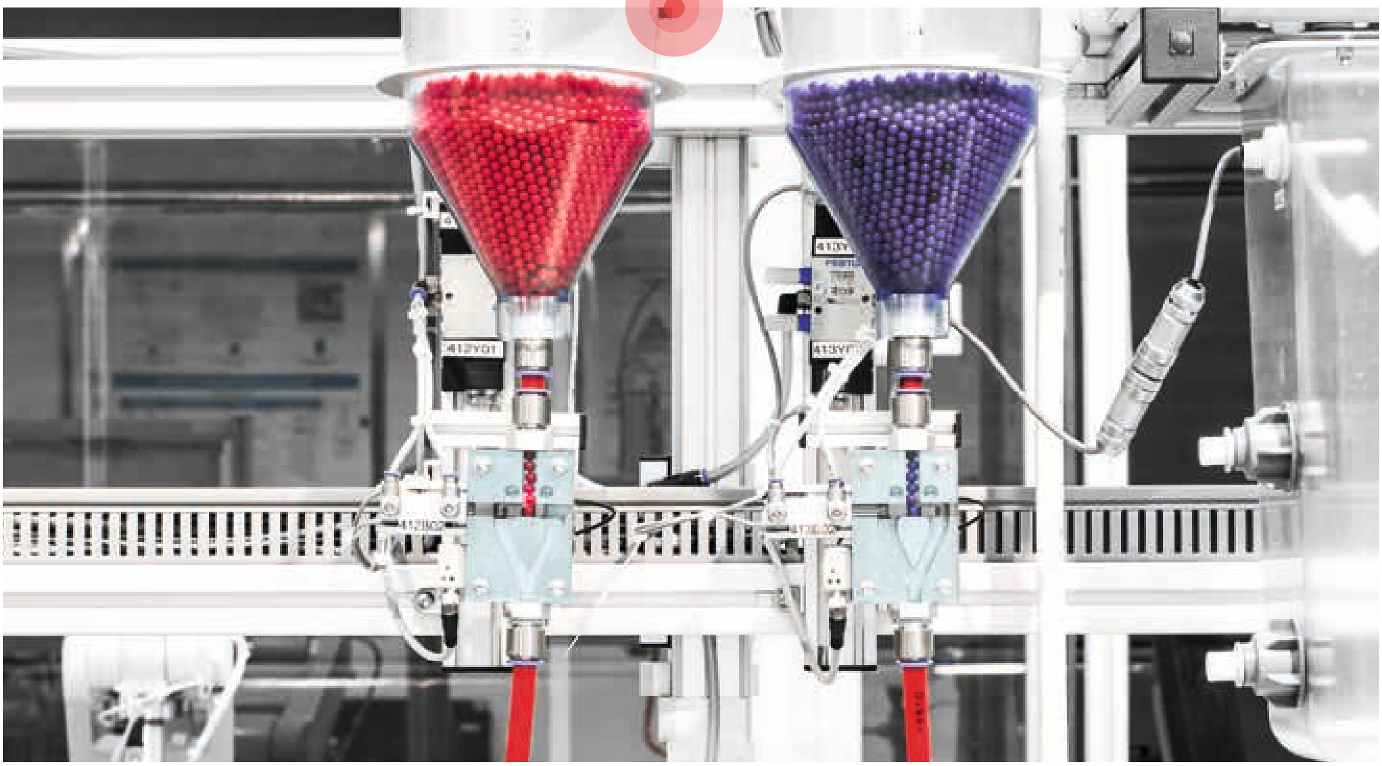

Robotic arm
picks up containers and places them onto belt

Beyond digital Babel

But why is that, in fact? Vogel-Heuser doesn't need asking twice. She knows the situation on the ground. For ten years, she was involved in plant engineering herself, and she has worked with industry partners on dozens of projects as a researcher. "Within an enterprise, there is often no clear semantics of communication between production units. Conveyors and scanners, filling equipment and packaging machines all come from different manufacturers and use diverse control systems." She explains that there often isn't even a seamless flow of information within the engineering, meaning the overall technical solution comprising electronics, software and mechanics: "The conveyor belt controller doesn't know, for example, when the motors have reached their load limit because design data the plant engineer created in a CAD tool is not compatible with the tools used by the electrical engineers. And these, in turn, are not recognized by the tools supporting software design."

A survey of sixteen German plant and machine manufacturing companies conducted by Vogel-Heuser and her team also bears this out. Many companies work with equipment acquired over time from a variety of manufacturers. The manufacturers themselves, after decades of operation, are often not adequately informed about which software their customers are using. Similarly, they do not know to what extent the customers modified the programs or even the hardware that they supplied. "Changes within production lines thus tend to be extremely complex," reports Vogel-Heuser. "Integrating a new component in an existing production facility means determining the current version of the old software, then adapting the interface before finally importing the new software." And identifying the system's previous software status is frequently no easy matter. Some solutions evolve several times over a longer period and manufacturers often only maintain very rudimentary variant or version management. "Some systems have been running for a good thirty years," Vogel-Heuser points out.

Her aim is to remove these barriers and pave the way for workflow and system efficiency at both the engineering and operational stages – not only in factories, but also among suppliers and service providers. "In the long term, for companies to hold their ground in the face of global competition, they will have to adapt quickly to new market requirements," forecasts the engineer. "To do this, we need a seamless flow of information and plants that can process incoming orders as autonomously as possible. Since product cycles are also becoming ever shorter, systems either have to adapt automatically or allow for rapid reconfiguration and redesign."

"We can run software updates and even replace components with new ones featuring different connectors in live mode without needing to completely shut down the whole system. As it stands, that is almost impossible in current industrial production plants."

Birgit Vogel-Heuser

Intelligent modules

The key to success here lies in a modular approach, Vogel-Heuser is convinced. A module is a self-contained unit – in the MyJoghurt demonstrator the robotic hand, scanner and conveyor belt are, for example, independent modules. Within any given module, all the information relating to its tasks – technical data, operating data, programming – is available at all times. This becomes possible when mechanical engineers, electrical engineers and software developers all use a common description with defined semantics. This allows information to flow freely between the different disciplines of both a single module or a combination of modules. The module is aware of its own capabilities and knows when its limits are reached.

Modules can thus work together to make intelligent decisions. While the yogurt filling system is still adding chocolate chips and strawberry pieces to the jars, for instance, the system can check the next incoming order. Say a customer would like yogurt with cherry pieces. The first thing is to determine whether the system can process the relevant in- ▷

Picture credit: Bauer

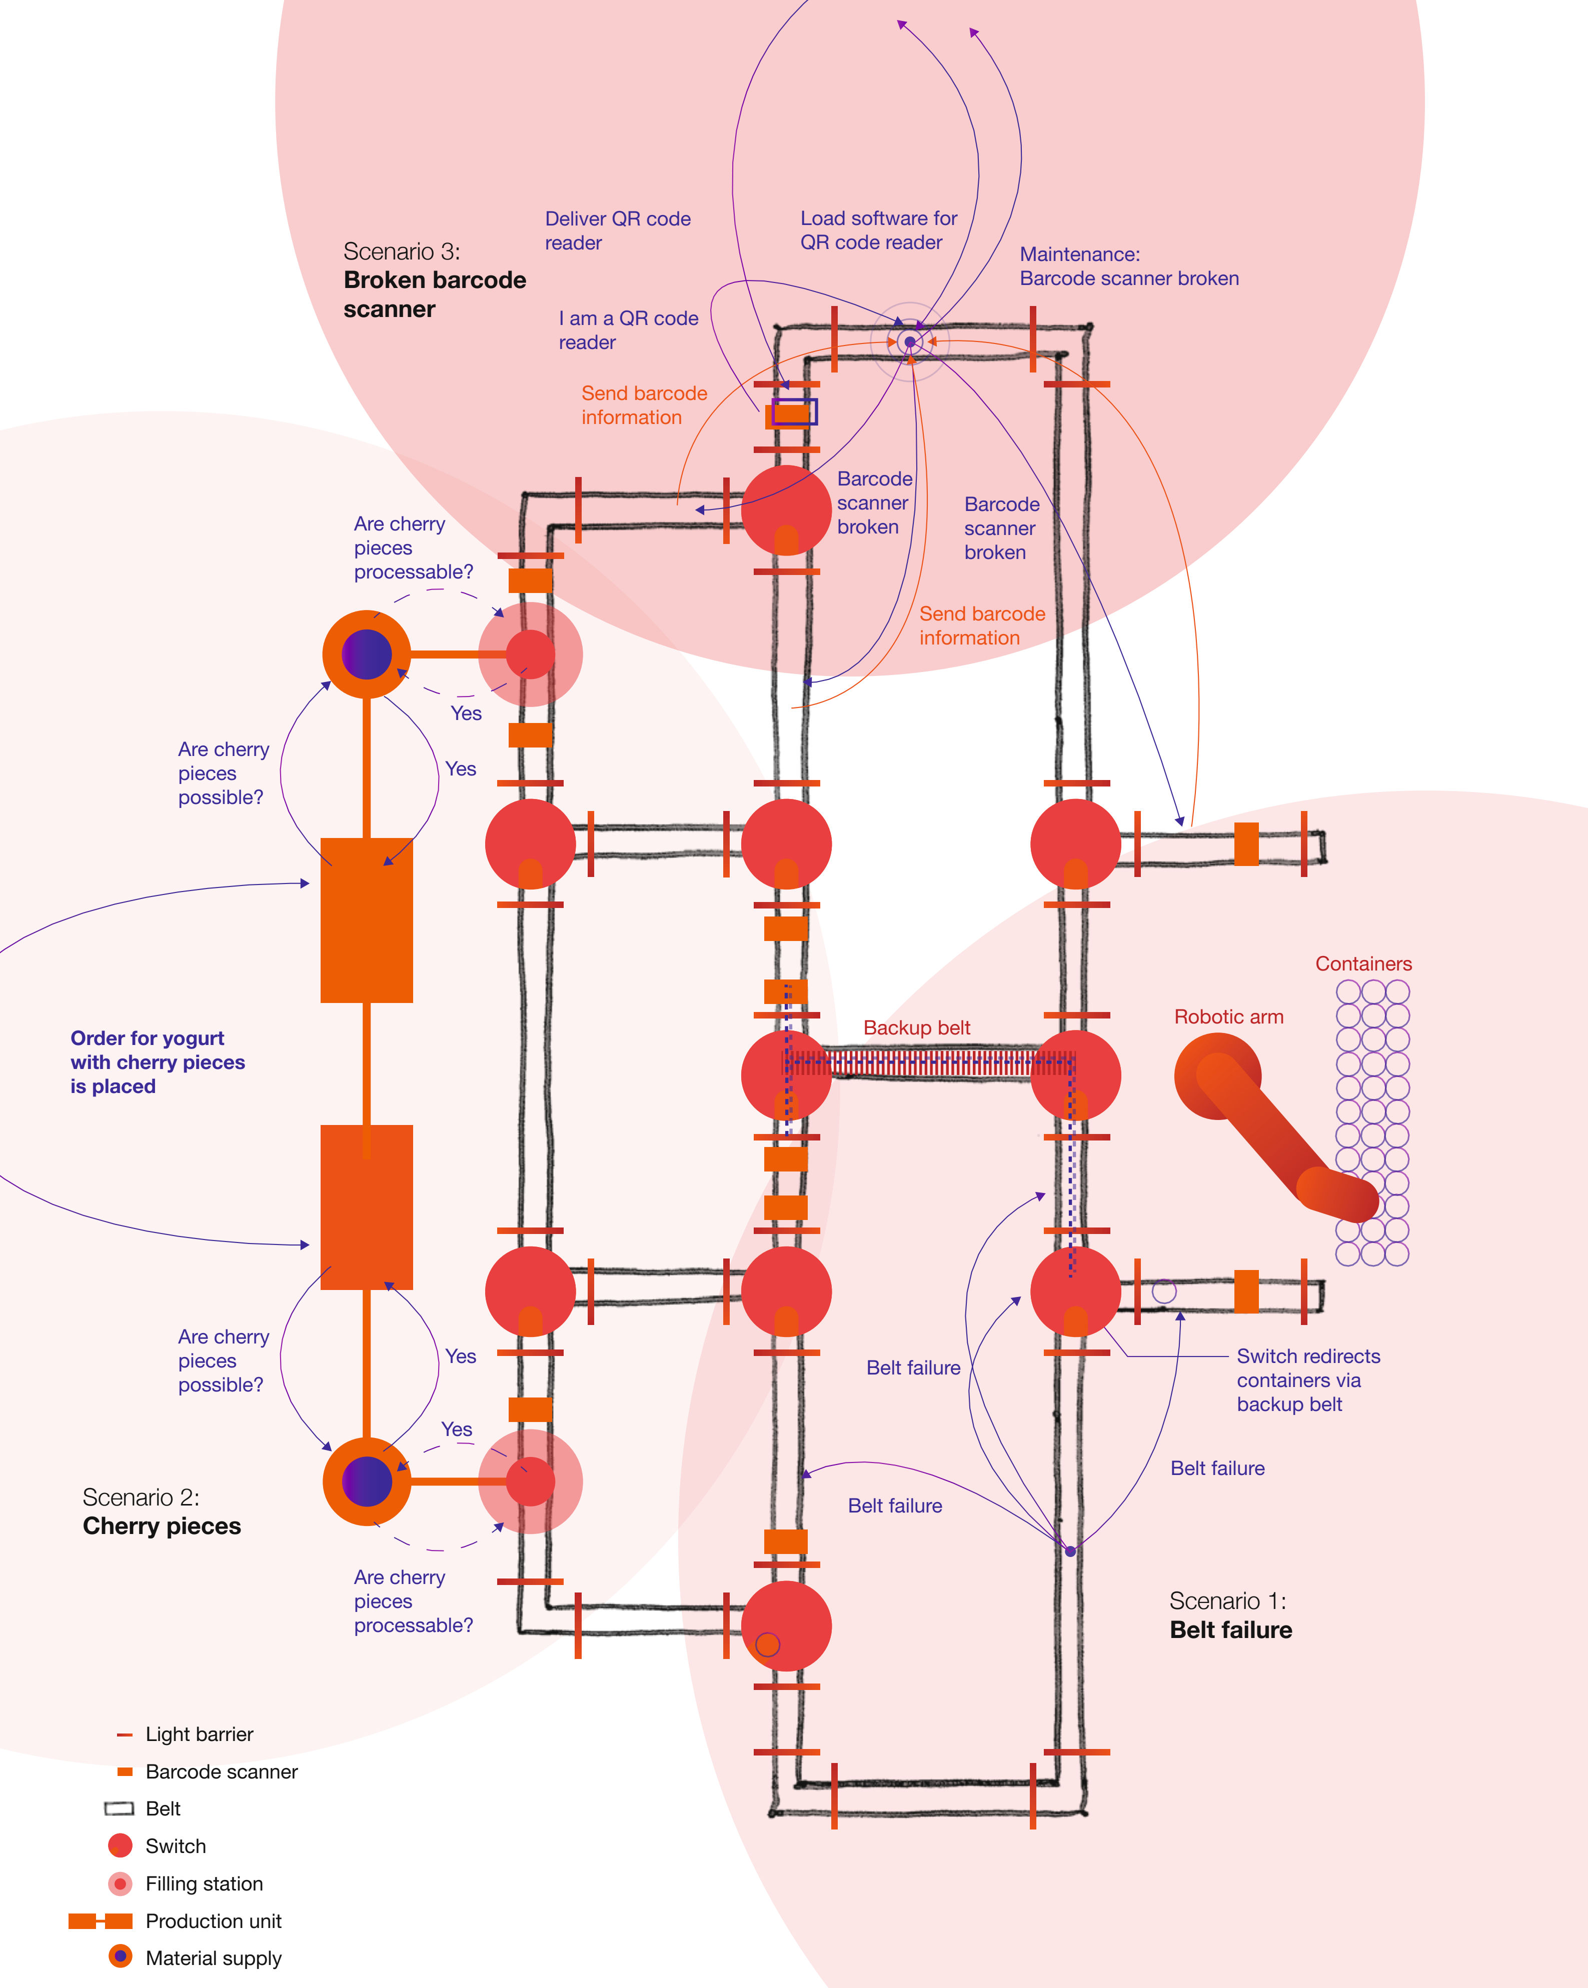

Scenario 3:
Broken barcode scanner
Deliver QR code reader
Load software for QR code reader
Maintenance: Barcode scanner broken
I am a QR code reader
Send barcode information
Barcode scanner broken
Barcode scanner broken
Send barcode information
Are cherry pieces processable?
Yes
Are cherry pieces possible?
Yes
Order for yogurt with cherry pieces is placed
Containers
Robotic arm
Backup belt
Are cherry pieces possible?
Yes
Yes
Belt failure
Switch redirects containers via backup belt
Belt failure
Belt failure
Scenario 2:
Cherry pieces
Are cherry pieces processable?
Scenario 1:
Belt failure
Light barrier
Barcode scanner
Belt
Switch
Filling station
Production unit
Material supply

Switch
with light barriers

Filling station
with mixing unit

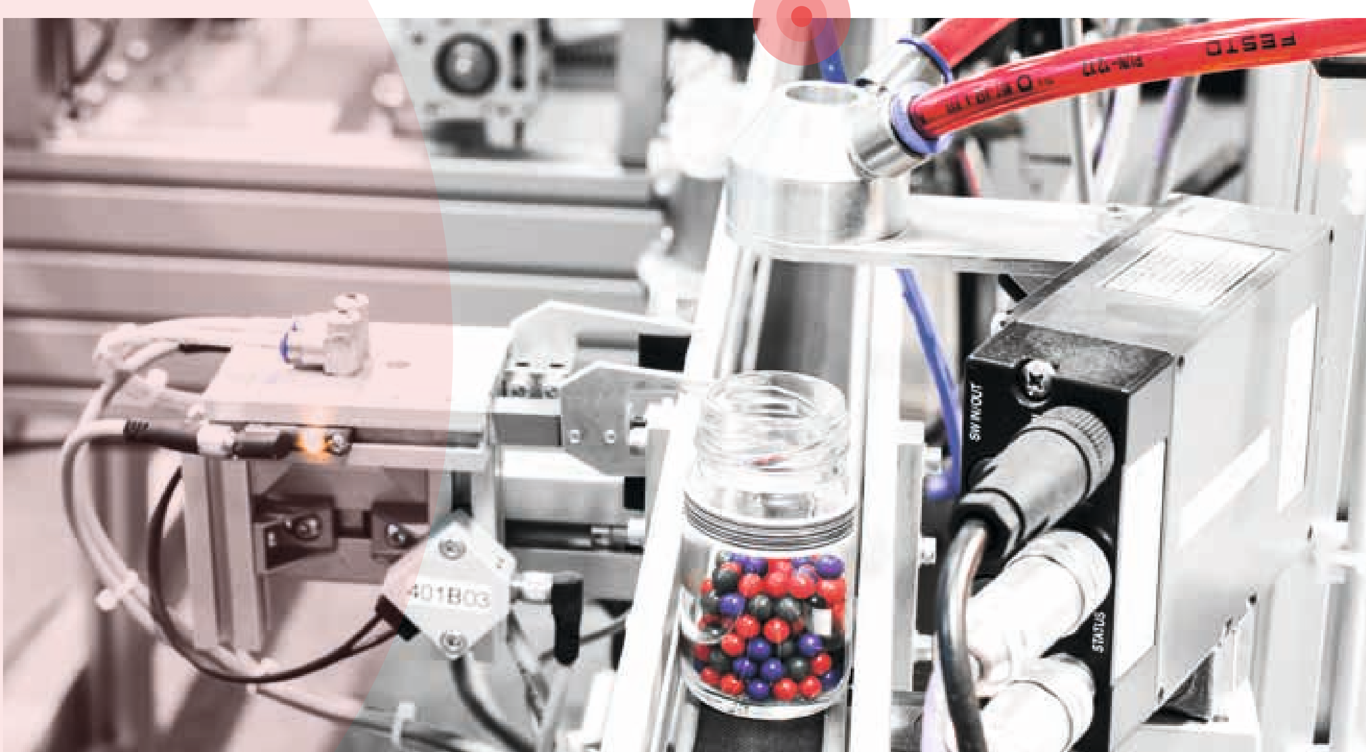

Graphics: ediundsepp (source: TUM); Picture credit: Bauer

“Industry 4.0 has now become a buzzword. Everyone wants it, but everyone means something different by it.”

Birgit Vogel-Heuser

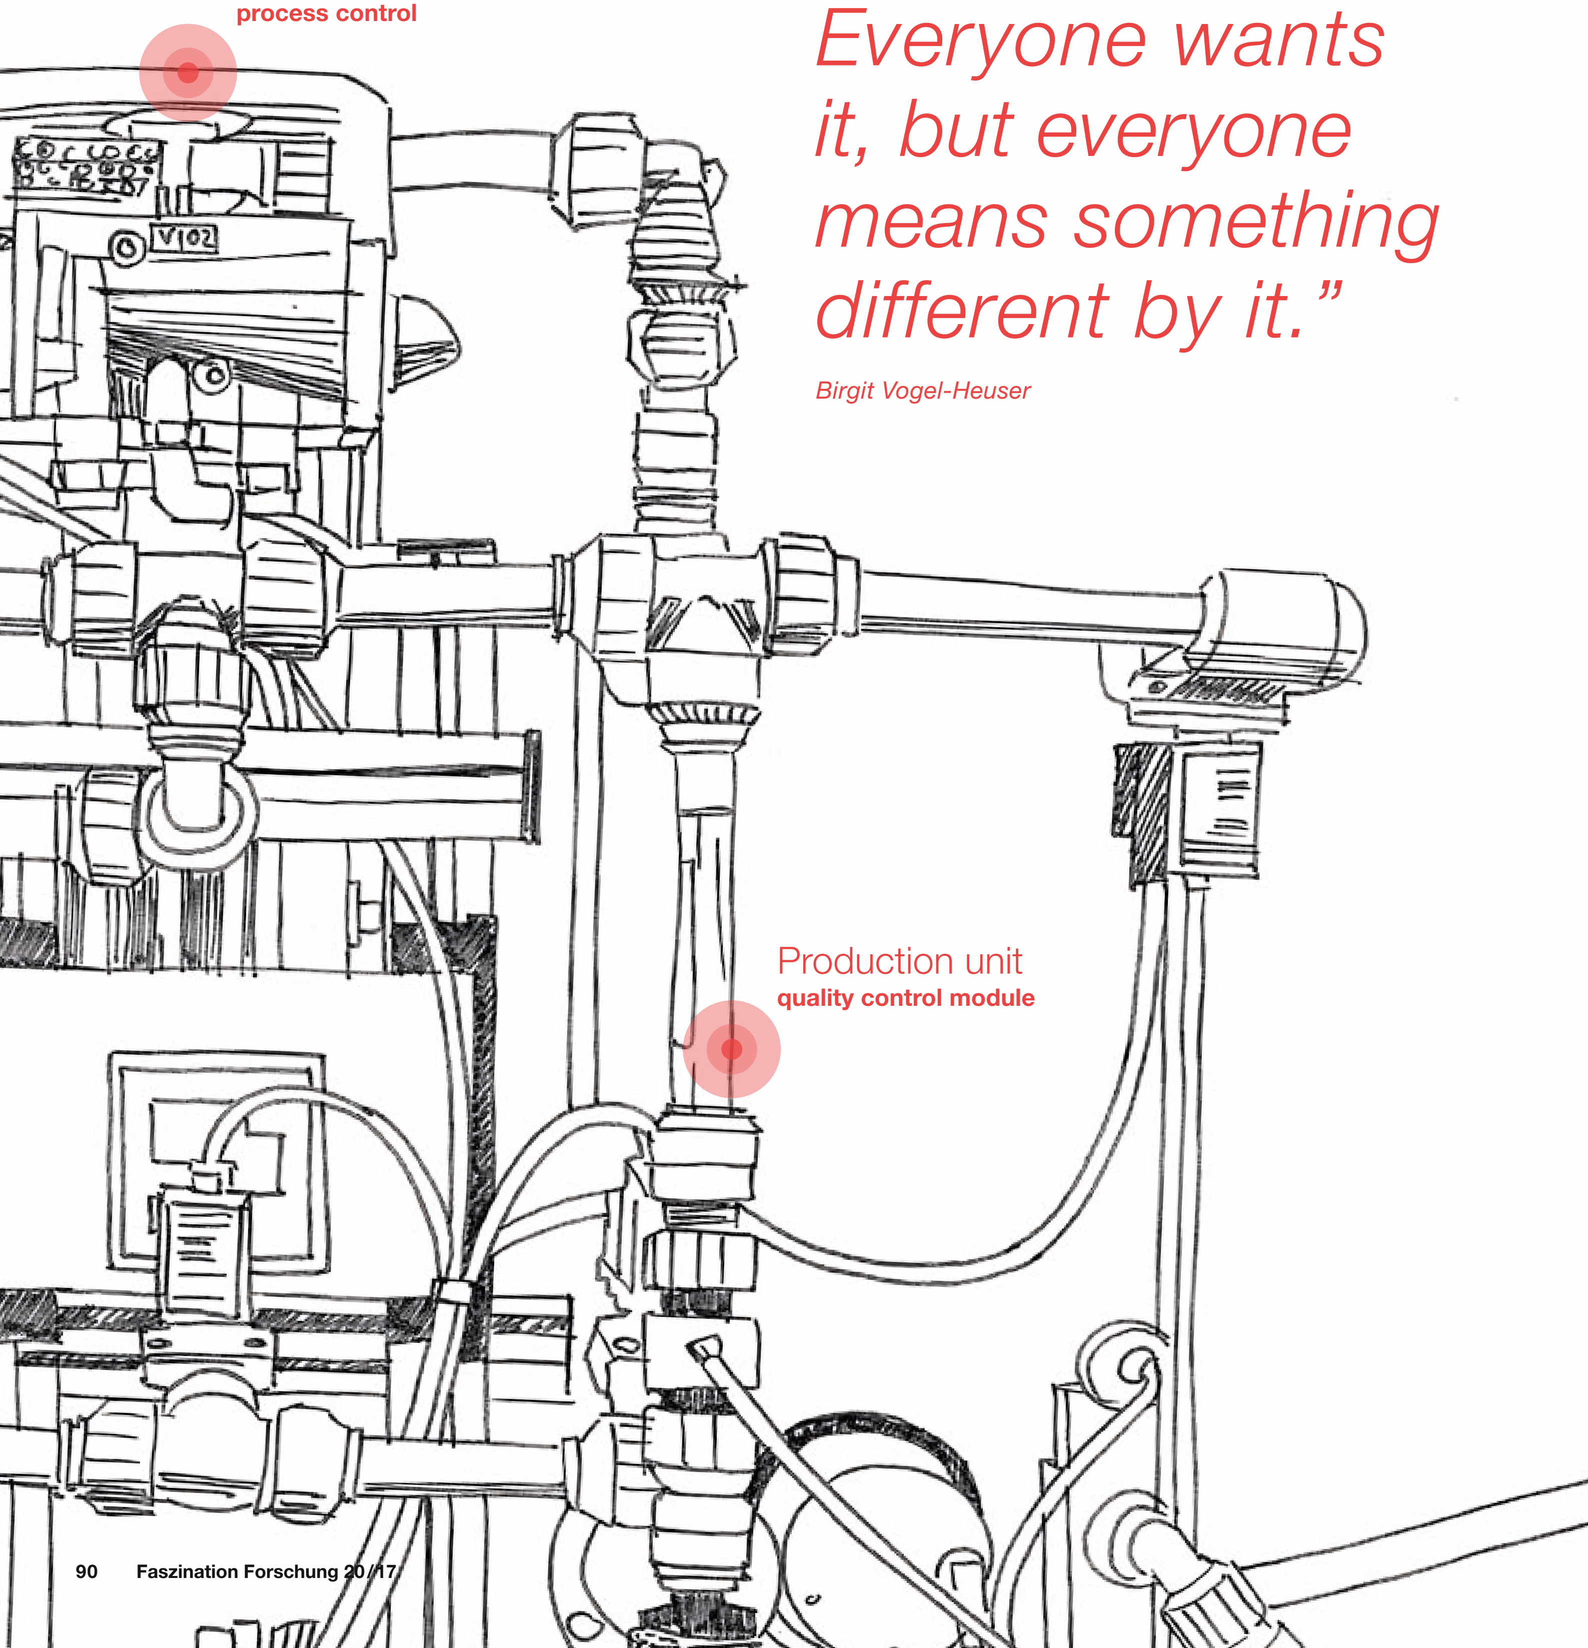

Prof. Birgit Vogel-Heuser

Interdisciplinary from the outset

Birgit Vogel-Heuser likes to bridge boundaries – both geographical and, above all, technical: "We can only advance Industry 4.0 with an interdisciplinary approach." Indeed, she has constantly crossed the borders between various disciplines in the course of her own career. Having earned her degree in electrical engineering at RWTH Aachen University, she then transferred to the Faculty of Mechanical Engineering to pursue a doctorate focusing on software-oriented automation. She thus laid the foundations early on for the triad of mechanical engineering, electrical engineering and information technology that still informs her research today.
After obtaining her doctorate, Vogel-Heuser spent ten years in industry, working first with Siempelkamp Group and later with ATR Industrie-Elektronik. Alongside this, she established the Automation Technology professorship at the distance-learning University of Hagen (FernUniversität Hagen). She then made the leap back into research in 2000, first becoming Chair of Automation Technology and Process Informatics at the University of Wuppertal (BUW). She then held the Chair of Embedded Systems at the University of Kassel from 2006 to 2009, before moving to Munich in 2009 to take up the Chair of Automation and Information Systems at TUM.
In recognition of her work, Vogel-Heuser has received the Adam Opel Award, the GfR Sponsorship Award, the Borchers Medal from RWTH Aachen, and the Initiative D21 Special Award ("Get the Best").
Vogel-Heuser is spokesperson for the German Research Foundation's Collaborative Research Center 768: "Managing cycles in innovation processes – integrated development of product service systems based on technical products", a member of the Deutsches Museum Board of Trustees, and a member of the German Academy of Science and Engineering (acatech).

gredients in the first place. The size of the cherry pieces is automatically compared with the diameter of the filling tubes. "For us humans, that is simple and straightforward, but it can lead to problems with machines – for instance if the supplier sending the cherry pieces provides measurements in millimeters, while the tube diameter is stored in inches. But defined semantics clears this up, by providing a common understanding," emphasizes Vogel-Heuser. The next step is to check whether all ingredients will be available in the right time frame, and whether the capacity of the robotic hands and conveyor belts is sufficient to carry out the order on time. All of this must be ensured before the order is accepted. Thanks to defined semantics, communication between the demonstrator's modules, even newly added ones, is available in real time and reliable. Special software agents take care of this coordination in real time, while the plant is running. If a component – such as a filling level sensor – is updated with a newer version, each module recognizes the new hardware from the information available and integrates it automatically. But what happens if a new yogurt variety is to be produced and the system's software has to be modified to allow it? Can approaches be developed here that could be used in similar systems around the world? Vogel-Heuser is currently working with industry partners and as part of a German Research Foundation (DFG) priority program to develop targeted methods for evolving software over decades. These efforts center on optimized module structures to enable systems to adapt with as little effort as possible to new requirements and thus, ideally, continue to evolve themselves. In her TUM institute, a demonstrator set up in Vogel-Heuser's lab shows how this might work, with a conveyor belt transporting boxes past an old-style barcode scanner. "Here, we have created a scenario that corresponds to actual practice in intralogistics," explains the engineer. "Our aim was to replace the old scanner with a new one without interfering with system functionality of the logistics system." This relies on a clever trick: "Taking the old software interface as our basis, we developed a new one that is outwardly identical. So the software works as a kind of adapter, ensuring that the system doesn't even notice that a component has been replaced and updated."
These new Industry 4.0 solutions certainly work well in the lab. But how long will it take for the modular concept to catch on in the real world? "Looking ahead, Industry 4.0 is coming one way or another," responds Vogel-Heuser. "But the road to intelligent plants and equipment is rocky, so we need to move forward slowly and surely. We always need to consider each individual scenario and look for solutions together with the company involved – there is no quick fix." That said, German companies, in particular, have a good chance of benefitting from the new research findings sooner rather than later: "There is no country where research and industry are more closely intertwined, or where findings are transitioned to concrete applications more quickly." *Monika Weiner*

Graphics: ediundsepp; Picture credit: Bauer

E-bikes and the domino effect

> *"This example shows what enormous reach innovations can have."*
>
> *Birgit Vogel-Heuser*

What happens when customers want to cycle further? The electric bike rental case study shows how complex the interplay between innovations, businesses and consumers can be. Let's say a rental service offers electric bikes with a range of thirty kilometers. But an increasing number of customers want to ride longer distances. An ambition with far-reaching consequences: First, the rental company contacts the e-bike supplier. They determine that their batteries are not powerful enough. As a result, they need to switch to a different battery technology. This, in turn, also changes the way the battery is mounted to the bike, as well as the interfaces between mechanics, electrics and software. And the battery production facility has to be reconfigured for the new technology. As soon as the batteries are delivered and attached, the bike software needs updating to provide customers with information about points of interest in their new, extended surroundings. The hitherto successful companies which provide the pick-up service may now also prove unsuitable, since more vehicles with greater capacity are required to collect the bikes. "This example shows what enormous reach innovations can have," explains Prof. Birgit Vogel-Heuser, spokesperson for the German Research Foundation's Collaborative Research Center 768: "Managing cycles in innovation processes – integrated development of product service systems based on technical products". The researchers set out to represent the interdependencies between service providers, products, and production facilities in the different models. Using these models to run through a wide variety of scenarios allowed them to find the best solution for the overall offering, spanning both the bike itself and accompanying services – in other words, the complete product service bundle. The identified dependencies between the different models are then presented to the decision-maker – such as the rental service. This familiarizes them with the factors and levers leading to the suggested solution and significantly boosts their confidence in the outcome. The researchers have developed a special visualization method showing how small changes – like customer demand in bike rental, as outlined above – result in a domino effect spreading through the entire system. Projecting these interactions beforehand allows businesses and service providers to cooperate in advancing the necessary innovation cycle, to their mutual benefit.

This is no mean feat, since each innovation cycle involves dozens of stakeholders, all with their own approaches. The software engineer sees things differently to the electrical or mechanical engineer – they each have a different mindset. "In the Collaborative Research Center, we are working to make these mindsets underlying the various fields more transparent," confirms Vogel-Heuser. "There are often semantic issues – different terms used for the same idea, or various ideas referred to by the same term. Only once we are aware of these differences can we reduce misunderstandings and work together to find solutions."

Longer distances for rental e-bikes requires larger batteries, new interfaces between mechanics, electronics and software. Visualizing the interactions between involved businesses helps to anticipate how an innovation will affect the whole system.

Picture credit: Bauer

Help the World Run Better and Improve People's Lives

Bill McDermott

Bill McDermott is chief executive officer of business software market leader SAP SE. Since becoming CEO in 2010, SAP has experienced dramatic increases in total revenue, operating income and market value. Under his leadership, SAP is the first global IT company to receive EDGE Certification for its focus on gender diversity and is the most sustainable software company in the world according to the Dow Jones Sustainability Index.
Bill has more than three decades of experience in business technology. He joined SAP in 2002 to lead the business in North America and has steadily risen to his current role as chief executive officer. Before joining SAP, he served in senior executive roles with Siebel Systems and Gartner, Inc. He launched his business career at Xerox Corporation, where he rose to become the company's youngest corporate officer and division president. His life and learnings are chronicled in his internationally best-selling memoir, "Winners Dream: A Journey from Corner Store to Corner Office".
He also serves on the boards of directors of performance apparel maker Under Armour, engineering software maker Ansys, and cyber security firm SecureWorks. He represents SAP on the European Roundtable of Industrials and the U.S. Business Council.

During a recent meeting of the European Roundtable of Industrialists, a special guest challenged the assembled CEOs with a simple reminder. "The shift from the old to the new," he said, "is the only tradition worth preserving."
The global economy itself is a fitting reminder of this tradition. We have seen the economy reinvent itself many times in the last century. Once from an industrial economy to a services economy. Once to an internet economy. Now, yet again, to a new digital economy.

What have we learned that can help us harness the forces of change?
First, while every individual has the capacity to change the world, it is almost always the enterprise where progress is best achieved. The enterprise brings people together, either for profit, for purpose or for both. The enterprise is where a collection of individuals can combine their abilities to advance society by creating products and delivering services. Yes, individual innovators are responsible for big ideas. But only once those big ideas have been brought to full scale have we recognized the full magnitude of their impact.
Second, trust is the ultimate human currency. Since the earliest days of recorded history, the ability of one person to trust another has propelled society forward. When people trust, they can build. When people trust, they can accept mistakes and failures. When people trust, they can share in success and triumph.
Third, the world's biggest challenges are also its greatest opportunities. Feeding, housing and healing people are challenges that transcend any single phase of the economy or of society itself. It is precisely our collective will to do these things better that lifts up one generation from its predecessors.

What do these three learnings mean for our new digital economy?
When it comes to the enterprise, today we understand that data is its most significant resource. Data enables us to connect people, things and businesses. Data enables machines to learn and to conduct routine transactions more efficiently. Like physicians who sequence the human genome to diagnose and cure illness, technology can now help us sequence the economy to fuel a new era of unparalleled growth.
Even with data as the new fuel of our economic engine, a lack of trust could undermine our forward progress. We must help all people to see a place for themselves in this promising future. We must train people to understand new technologies rather than abandon them for fear of the implications. It is no longer enough that we profess our concern for others. We have to prove this through action.

Finally, our challenges can build bridges. With geographic boundaries largely irrelevant, we live in the most ethnically, racially and culturally diverse society in history. The questions before us must not be about what makes us different. We must focus on what brings us together. We must focus on environmental sustainability (SAP will be carbon neutral by 2025). We must cure cancer (The National Center for Tumor Disease in Heidelberg is making progress on the SAP HANA platform). We must give every child a path to a bright future, regardless of where they are born (SAP is teaching young people to code in Germany, Africa, Latin America and beyond).
SAP has been at the epicenter of the enterprise technology industry for nearly five decades. We have never been more optimistic about what's possible. We have never been more inspired by the leadership of our 350,000 customers in 190 countries around the world.
Today we are experiencing a shift from the old to the new. We embrace it, especially as it offers an unprecedented opportunity to help the world run better and improve people's lives. □

Picture credit: SAP SE

Cover graphics: ediundsepp, Nicku/Shutterstock