botenstoff

SPECIAL

Artificial intelligence in health

Developments, applications, trends

human.
technology.
styria.

Contents

Pictures created with AI

DALL·E 2 is an artificial intelligence system that can create realistic images and art from prompts written in natural language.
For example, the prompt used to generate the cover was, "health industries cooperate with AI, photorealistic". For this issue, we have generated several images in different styles using DALL·E 2. The prompts used for each image are shown in the image credits.

Foto: Sissi Furgler Fotografie

Foreword

Live digital, work medical

We are Human.technology Styria; technology is part of our name. Our partners here in the region are highly technology driven: they think digital, they breathe digital, they live digital. In particular, our community here in Styria is working at the cutting edge of many applications in artificial intelligence: discovering covid variants, improving quality management in pharma production, interpreting medical images and rolling out telemedicine.

AI is often talked about as the disruptive technology of the 21st century; the kind of thing where sometimes we are in the driver's seat and sometimes technology drives us. As Dr. Siegfried Meryn said in a keynote talk in Graz: "In the future it's not about what medical treatment will be possible with your mobile phone, it will be about what will not be possible".

AI will change the face of how many medical tasks are done and how many procedures and conditions are managed. These developments also throw up many novel challenges such as trust, explainability, data availability, usability, and ethical questions. And all of this will change how we educate and train doctors and other healthcare professionals.

These matters are not not self-contained within our local R&D region, they are things we are dealing with in networks stretched around the world, working with partners in centres of innovation such as the Mayo Clinic. Our international activities led us to the Twin Cities this year, and are bringing us to Helsinki & Turku in Finland next year. As surely as the southeast foothills of the Alps are our geographical location, these contacts and collaborations are where our thoughts and developments are at home.

It is my pleasure that my first "botenstoff special" as CEO of our human technology cluster is able to present some of the exciting projects happening here, and I'm looking forward to many more to come.

Lejla Pock
lejla.pock@human.technology.at

„Modern AI is a performance amplifier for human cognitive abilities. It can work much faster and look at much larger data sets."

Wolfgang Kienreich

AI – history, state of the art and future challenges

Artificial intelligence (AI) deals with the simulation of human cognitive functions by machines. Whether you're playing a board game, booking a trip, finding your way to a venue, or writing an article like this: Modern AI can perform such tasks largely without human assistance. At the Know-Center in Graz, we have been following and shaping the development of AI for more than 20 years. Here we provide an overview of the history, state of the art and future challenges of this fascinating topic.

Early visions of AI emerged in the mid-20th century and described a general AI that would simulate human intelligence in every aspect. This included the physical manifestation in the form of humanoid robots, a fantasy that lives on in our attitudes to AI. But initial research showed that the development of general AI was a remote and likely impossible goal. The rule-based methods used at the time often failed due to the complexity of most real-world tasks. It's like when you try to cook something from a recipe and find that in the reality of your kitchen, a lot of details arise that it doesn't deal with. As an alternative, inspired by the workings of the human brain, neural networks were invented. These are computing systems made up of large numbers of interconnected, simple decision-making units that can be trained to do a task using examples. For example, if we show such a network a large set of X-ray images together with the correct diagnoses for each case, it learns the relationship between the image and the diagnosis and then it is able to correctly diagnose images it hasn't seen before.

But when we consider that an X-ray image in diagnostically meaningful resolution consists of millions of pixels, and that hundreds of different diagnoses are conceivable – and then, that for each of these variants a large number of images are needed to train the neural network – then you can see that this is a task that takes some serious computing power. Only in the last few years that have we had sufficient computing power and memory to use neural networks meaningfully for many real-world tasks.

In the meantime, AI spent a long time in what some called a winter period. And yet, that time was not wasted. It was used to develop many methods that are still at the heart of many problem-solving strategies in AI today. For example, in image processing, "convolution neuronal networks" (CNNs) were established as a tool whose function is intuitively understandable. You would certainly not inspect millions of pixels individually to recognize a holiday picture. But if you compress the image to just a few pixels, and what remains is blue above yellow, you are probably looking at a beautiful beach shot. Other technologies such as LSTMs (long-term short-term memory models) provided similar breakthroughs for processing speech. Most of this progress was barely noticed by the public.

At the beginning of the 21st century, the development of the Internet made very large amounts of data on most topics immediately available. Handling and using these datasets in a meaningful way has been enabled by Big Data technologies. At the same time, massively parallel computing power became affordable outside of dedicated data centres. Finally, the way was clear for practical application of neural networks. AI awoke from its slumber. But AI methods had come a long way from the early visions. Modern AI is data-driven and problem-solving oriented: It processes large amounts of data describing how specific problems were solved in the past and computes a so-called model. This model is then able to solve new, previously unknown problems of the same type. For example, a pharmaceutical AI can learn a model of the relationships between active ingredients and their effects from a large set of known formulations – and then it can predict the effects of a new formulation, or even propose formulations with certain effects itself. In this way AI shows itself to be highly intelligent within the context of the problem and the data.

The latest advances in AI are particularly impressive in human language processing. Modern language models are able to thematically categorize and summarize arbitrary texts and can even compose new texts based on a given topic. If an AI can summarize thousands of descriptions of side effects by patients in a meaningful and grammatically correct way, then the historical dream of the artificial brain seems within reach. But this impression is deceptive: the AI always remains within the space defined by the underlying data. But what it does do is arrange these with enormous speed in ever new patterns, which are often helpful.

It's probably best to see modern AI as a performance amplifier for human cognitive abilities. AI can work much faster and look at much larger data sets than human experts. This shifts the working relationship between humans and machines. At the Know-Center, we have recently been able to successfully transfer routine tasks to AIs in hundreds of use cases in industries like production, logistics, energy, medicine, and pharmaceuticals. But when tasks go beyond the routine, beyond model boundaries, and special cases need to be considered – then we still need human understanding and general knowledge.

This lack of creativity isn't really the most central problem for us at the moment. The things we're grappling with in practical AI applications are quite different and arise from their data-oriented nature. In many industrial applications, relevant data is still not available in digital form. For example, digital twins would be useful for simulation of production processes, for instance to predict quality fluctuations before they occur. But this requires complete digitization of the production environment, which is still far off. Processing large volumes of patient data by an AI has a lot of potential in medical research, for example in developing biomarkers. But data protection concerns often preclude such applications of personal data. And in recommender systems, we are increasingly finding that the commercially driven data on which many models are based leads to disadvantages for minority subgroups of users.

It seems there are many reasons not to trust modern, data-driven AI, even leaving aside the clichéd notion of the malevolent artificial brain. Efforts in current AI research have a lot to do with creating trustworthy artificial intelligence that ensures data security, is based on balanced data, respects human privacy, acts transparently, and is resistant to manipulation. Promising approaches to these issues are being developed in Europe in particular: for example, it has recently become technically possible to let AI learn and decide based on encrypted data. Trustworthy AI is needed in order to enable the application of AI in sensitive areas and to increase its acceptance. It will be at the centre of the work we are doing at the Know-Center in Graz in the coming years, and we are already finding first practical answers.

Wolfgang Kienreich
Know-Center

Best Practise Applications

© Innophore

INNOPHORE

The global variant oracle

Best practice example in the field of AI in health: Innophore and Amazon Web Services have jointly developed an AI that can predict how dangerous new Covid 19 variants may be. This new technology can be called a global variant oracle. Knowing the infectiousness and potential severity of new variants gives researchers a head start when it comes to developing or adapting drugs or vaccines against Covid-19.

To make reliable predictions, Innophore researchers began with a detailed study of the virus' structural features. Because the spike proteins exposed on the surface of Covid-19 play a central role in the mechanism of infection, the team quickly turned its attention to these proteins. With sequence data from around the world, Innophore created a library of models containing every variant recorded. This massive dataset is used to train the AI's algorithm so that it can make predictions about the assertiveness or severity of emerging variants. Strong binding affinity between human receptors and the viral spike protein is a solid indicator of a variant of concern and is one of the key features calculated and used in Innophore's prediction model.

Using AI to make predictions about Covid-19 is only the first step. By choosing an appropriate target protein, the method can be adapted to a wide range of infectious diseases such as HIV, RSV, influenza and others. At Innophore, the expansion of our portfolio of detectable diseases has just begun.

About Innophore

Innophore, based in Graz, is a high-tech spin-off of the University of Graz and the Austrian Centre of Industrial Biotechnology (acib) and specialises in the fields of digital drug discovery and enzyme search using 3D point clouds, AI and Deep Learning.

About AWS

Amazon Web Services (AWS), the cloud computing service of Amazon.com, is the largest cloud service provider globally. From its data centres, the business provides over 200 fully featured services including computing, storage, and databases. AWS currently has 26 regions and 84 availability zones in operation.

LEFTSHIFT ONE

Artificial intelligence in the pharmaceutical industry

Graz-based startup Leftshift One has developed an AI solution that is now being used in quality control for drug production by Fresenius Kabi, the international pharmaceutical group. The new approach was initiated through a pilot project at the pharmaceutical company's site in Graz – where 100 million doses of drugs such as anaesthetics, analgesics and antibiotics are produced every year. In pharmaceutical production, there must be no deviations in planned processes, and if they do occur, they must be checked closely and, if necessary, an entire batch must be destroyed. Good documentation is crucial, and until now this has been done exclusively by employees on site. Now they are getting help in form of artificial intelligence.

Symbiotic cooperation between man and machine

"With the help of AI, we want to ensure that the information for analyzing a deviation is sufficient and complete – regardless of parameters such as the employee documenting or the time of day," explains Fresenius Kabi Austria innovation manager Christoph Köth. The AI builds on the documentation of the employee: "Our technology is able to analyze and classify the textual description of the employee. This converts previously unstructured texts into comparable data sets," explains Patrick Ratheiser, CEO of Leftshift One. The results are recommendations that should make it easier for the quality experts to decide how to deal with the deviations.

The Fresenius Kabi case also uses AIOS, Leftshift One's trusted AI operating system, which enables the rapid implementation of AI applications. "With the help of our technology, we can offer initial solutions to companies across industries within a few weeks. What is important is to combine the respective domain know-how with our data expertise and to keep people in the loop. This is particularly true in this project, as the preliminary classification of AI has to be comprehensible, transparent and explainable for the user," emphasizes Ratheiser. After all, only the symbiosis of human and machine makes the application of AI meaningful.

© Leftshift One

AIT

Artificial intelligence for telemedicine

How AI will multiply the capacity of telehealth centers - Telemed5000

Telemedicine has arrived in standard care and is becoming increasingly important in the treatment of chronic diseases. One best practice example in this field is HerzMobil, a disease management programme that provides telehealth support for patients with heart failure in Tyrol, Styria and Carinthia. For these high-risk patients, telehealth solutions can significantly reduce both the mortality rate and the rehospitalization rate; a recent data analysis by Dr. Gerhard Pölzl, Medical Director of HerzMobil Tirol, showed a 61% risk reduction for death within 12 months and a 46% risk reduction for readmission or death in 6 months, compared to standard care. This form of monitoring also has big benefits for patients in terms of security and independent living. This is just one example of how telemedicine brings the care of patients with chronic diseases into the age of digitalization.

This kind of care is so successful that it really should be available to everyone with this condition. But individual telemedicine centres, such as the one at the famous Charité hospital in Berlin, cannot care for more than 750 patients remotely. And the demand is much greater than this, and growing all the time. The prevalence of heart failure is increasing: firstly because in an ageing society, the number of people experiencing heart failure is rising, and secondly because more and more people are surviving cardiovascular events, but with persistent damage to their heart, and need careful monitoring. In Germany alone, there are around 2.5 million people living with chronic heart failure, with around 300,000 new cases every year. The current capacities of telemedicine centres are not enough to provide adequate care to these large patient groups.

TELEMED5000: raising capacity from 750 to 5000 people cared for

Since the beginning of 2020, the Charité Berlin has been working together with the Digital Health Group of the AIT Center for Health and Bioresources in the TELEMED5000 project, to develop an intelligent system for the telemedical care of large collectives of cardiological risk patients. The mission is to use AI, such as deep learning, federated learning and the Internet of things to create a system solution that makes it technically possible to manage large numbers of patients. The goal is to be able to manage up to 5000 cardiac risk patients from a single telemedicine centre, using an intelligent system. The project builds on the five-year Fontane Study conducted at the Charité by Prof. Friedrich Koehler, head of the Center for Cardiovascular Telemedicine, which was the first demonstration that telemedical co-care can prolong the lives of cardiac patients and that they spend fewer days in hospital. It is equally suitable for patients in rural areas and in urban centres.

An important milestone for telemedicine in Germany was the decision of the Federal Joint Committee (G-BA) on March 30, 2021, to reimburse telemedicine for patients with heart failure as a standard benefit of the statutory health insurance funds. The decision means that approximately 200,000 patients in Germany will be entitled to telemedical care in the future.

In Austria, the clinical information system of the public health system, ELGA, has for the first time created a "telemonitoring episode report" as a format and infrastructure for handling patient-submitted heath parameters, such as blood pressure.

AIT is a pioneer in research and care for Digital Health in Austria with over 20 years of expertise in telemonitoring and therapy management, and with over 2 million telemo- nitoring days of experience. AIT has developed and tested a telehealth technology platform to support disease management programs and integrated care models through national and international research projects. telbiomed Medizintechnik und IT Service Gmbhwas founded as a spin-off from AIT for national and international commercialization of the telehealth solution.

Artificial intelligence in the Telemed5000 project

The development of a decision support system (EUS) for prioritizing patients includes relevant physiological parameters, including physical activity.

As an Austrian partner in the bilateral project, AIT has the task of evaluating and integrating suitable sensors for monitoring physical activity and subsequently refining the predictive analytics based on the collected parameters.

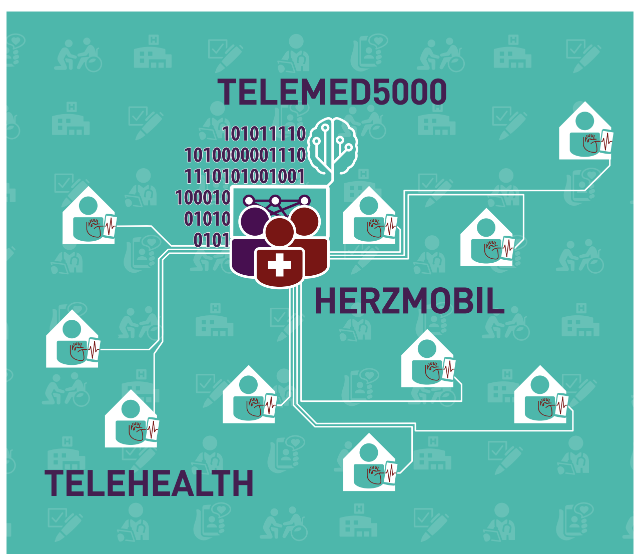

© AIT

© © DALL•E 2 | prompt: "mathematical knowledge, woodcut"

FRAUNHOFER

Mathematics for better health

Dr. Eva Eggeling

A lot of progress has been made in recent years in terms of computer power and the availability of large amounts of digital data. As a result, mathematical methods, both old and new, are experiencing a golden age. They allow us to find solutions quicker, and they help us to recognize recurring patterns in the data, and to quickly identify correlations, when manually sifting through the data would take weeks or months. In medicine in particular, AI and data analytics can mean a huge step forward. Mathematics provides a data-based recommendation, but the final decision is ultimately made by the physician.

"Data Driven Design" encapsulates all these aspects, and that is why I chose it as the name for the Fraunhofer Austria Center, which I lead and which has one of its locations in Graz. Our motto is: We make more out of data.

We have already started collaborations with medical institutions, and we have been able to apply our mathematical knowledge in the health sector. One of our current projects - "Curin' "– is dedicated to the detection of urinary tract infections. In order to reduce the workload in culture and to detect as many negative samples as possible in advance, it is already common practice to perform a screening. Current techniques, however, only use two or three of 40 available parameters for diagnosis, achieving a reduction of 45% of workload, while maintaining a sensitivity of 95%. We have now received data from 15,000 urine samples from our cooperation partner, a clinic in Italy, and are now in the process of training an AI to recognize further correlations and thus improve this method even further. A very different approach, but no less exciting for public health, was our statistical analysis of Google queries during the pandemic. Again, we used mathematical methods and we were able to show that searches for typical Covid 19 symptoms such as headache or chest pain increased in each wave of infections.

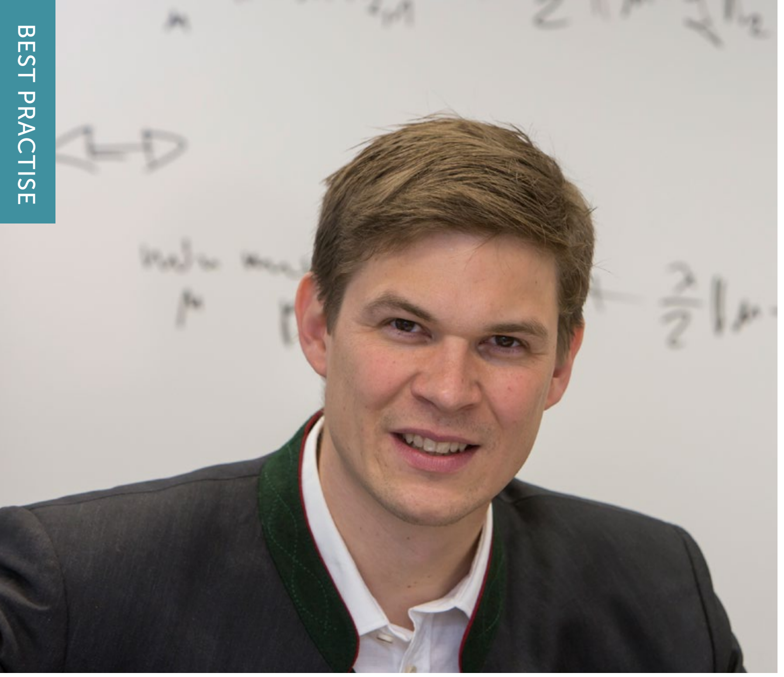

Thomas Pock, Institute of Computer Graphics and Vision, TU Graz

TU GRAZ

AI in medical imaging

Artificial intelligence (AI)-based algorithms are increasingly being used in medical imaging.

AI algorithms can reconstruct high-quality images even from incomplete and noisy measurement data. As a result, they have become widely used in medical imaging to support diagnostic and therapeutic decisions. The use of these algorithms brings great practical advantages, as Thomas Pock from the Institute of Machine Vision and Display at Graz University of Technology explains: "In computer tomography, for example, the dose of contrast agent can be reduced, and in MR imaging, the acquisition time can be significantly shortened." AI algorithms have already been implemented in the latest generations of scanners, so they are beginning to be used in routine clinical practice. However, the algorithms used so far, known as "discriminative AI algorithms", have a major drawback: They are inflexible in terms of the imaging protocol. While they are good at distinguishing between noise and real signal if the imaging protocol is kept constant, a growing number of experiments show that the quality decreases rapidly as soon as the imaging protocol is changed.

Based on fundamental principles of probability theory, Thomas Pock and his team at TU Graz are developing AI-based generative models. These generative AI algorithms, in contrast to discriminative algorithms, are able to generate new images and, as a by-product, calculate the probability that the image is a real medical image. The image used here for illustration shows an example of an artificially generated MR image of a knee. "A key advantage of generative models is that they can be learned without having to rely on a specific image acquisition protocol. So they are much more flexible," explains Thomas Pock. In the reconstruction process, the generative model is then used to generate the most likely image that matches the given measurement data. "In our latest results, we show that the reconstruction quality of generative models is on a par with discriminative algorithms, but also allows for arbitrary modification of the acquisition protocol," Pock said.

Martina Aubel with her team (from left) Thomas Truskaller, Philip Stampfer and Julian Gutheil

JOANNEUM RESEARCH

Who keeps track? AI!

Digitalisation offers great opportunities in the health sector. Health data contain an enormous wealth of experience, but the sheer volume makes it impossible for humans alone to analyze everything. Collecting structured data and using AI to learn from it promises ways to put the information to good use and sustainably improve medical care. The "Clinical Decision Support" competence group at the HEALTH Institute of JOANNEUM RESEARCH is working on this.

Martina Aubel leads the highly motivated group of researchers that develops digital technologies to support healthcare professionals in diagnostics and therapy. "One of our focus areas is geriatrics. This often presents a challenge in hospitals," explains Aubel. Older people naturally become more frail and their probability of developing single or multiple illnesses increases. Medical treatment is therefore usually complex, as many different things have to be taken into account. In the process, some risks may not be perceived. Recognising individual factors too late can lead to serious complications and thus have far-reaching consequences for those affected. "In order to create a real improvement for patients and also for healthcare professionals, we use artificial intelligence: it keeps track for us by recognising potential risks at an early stage and warning us. Doctors can then initiate preventive and therapeutic measures in time to avoid possible complications," Aubel says.

Of course, even an artificial intelligence cannot do this just like that – it has to be specifically trained for its task. For this purpose, anonymous electronic health data are used. The AI practices and learns to recognise risk patterns in data independently.

PARKSIDE

The UX success factor in the use of artificial intelligence

Artificial intelligence (AI) holds enormous potential for automation, for the creation of new products and services, for increasing efficiency and for data processing. However, AI can also create a sense of scepticism and discomfort among users or patients, especially when used in the human technology sector. In addition to the technical challenges, human-system interaction is also an important success factor in the use of AI, because it does not replace the individual, but rather interacts with the users and supports them in their daily work. Even in highly automated processes, the human being is the ultimate controller, and so it is crucial to also address the psychological aspects of this interaction. Especially in the environment of medical devices, a sufficient risk analysis and subsequent usability tests are of crucial importance, as even the smallest uncertainties and errors in use can have serious consequences for patients.

To ensure that AI-based applications are also used profitably, it is crucial to involve UX (user experience) and UI (user interface) designers in the development process in addition to data scientists and engineers. They ensure that the human factor is taken into account and that technically perfect solutions also achieve the necessary usability and acceptance.

In AI applications, the UX focuses on the following topics, among others:

- the motivation of the user to use an AI-based tool
- the trust in the results of the AI, the data security and its explicability and comprehensibility
- the user's expectation management of the AI application and
- a well thought-out interaction according to common UX criteria, which offers the AI application the possibility to learn.

Learn more about the most important focus areas for UX in artificial intelligence systems at: www.parkside.at/ux4ai

Artificial intelligence as an ethical challenge

AI, at least as a buzzword, is well known to the general public. How AI works, what potentials and possible dangers it harbours, and how to deal with it in a meaningful way, are not so well understood. People's emotional reactions to AI include everything from flat-out rejection to massive existential fears to ideas of paradise. A similar spectrum of attitudes is evident even among experts. To arrive at a sense of how we as a society want to handle the use of AI, and what measures should be taken, we need to consider and discuss the ethics of its development and application.

Ethics is primarily concerned with the question of what is good and with the evaluation of actions and decisions as to whether they are good or bad. The potentials, but also the risks and dangers of AI require us to be clear about what is good and right for us. There are now a number of proposals to establish ethical orientation principles in dealing with AI at all levels, especially among decision-makers. These include, as basic principles, autonomy, non-harm, "do good"/ welfare, and justice/fairness. Other normative orientations would be to protect groups that are particularly worthy of protection, clearly regulate the distribution of power and information, and thus also create clear responsibilities. Issues that need answers are social impacts, prohibition of discrimination and preservation of diversity, special protection of privacy (confidentiality, integrity) and security, but above all transparency and explainability and the necessity of stakeholder analyses. In case of conflict, it is always necessary to justify which principles and values are to be rated higher than others. Not all of these orientations can always be realized at the same time.

However, an ethical framework is needed at all levels, starting with the initial ideas and extending to the developers, the developments (ethics-by-design) and companies, all the way to concrete applications and political regulations. Sensitization and adequate value attitudes must be incorporated into personal behaviour everywhere and must not simply be delegated to clients or to higher authorities. Each individual has to take responsibility for their own actions. Transparency is a key tool here, enabling us to accompany developments with critical reflection at an early stage and to stimulate a social discourse about what we want and what is better not to do. Since not every end justifies the means, it must also be possible to forego certain opportunities in order to safeguard fundamental values and principles. This requires impact assessments, clearly defined areas of application, small implementation steps, sufficient safety nets, permanent evaluations and clear responsibilities in order to construct future AI in a "human-compatible" and "human-sensitive" way.

At the moment it is hardly possible to predict where the development of future AI systems will lead. But we can develop our ethical discussions now and by preparing the relevant ethical principles, we can be ready to exploit these potentials in a way that we want to without having to suffer the worst of the risks. What matters is that we can only do this together.

Univ.Lektor Dr. Andreas Klein
Private lecturer at the University of Vienna, and independent health care ethicist

> „We need to construct future AI in a 'human-compatible' and 'human-sensitive' way"
>
> **Andreas Klein**

AI in education for health professions: radiography

Explainable AI offers significant opportunities especially in the diagnosis and treatment of rare diseases. Like the discovery of X-rays, there is, on the one hand, outspoken enthusiasm for the new technology and, on the other hand, a justifiably critical attitude. From our point of view, an open/academic discourse between scepticism and euphoria is necessary in the interest of education in health care and in the interest of students. It is the professional ethical responsibility of teaching to explore new procedures in the radiological processes such as X-ray diagnostics, nuclear medicine and radiation therapy. Although the field for the application of artificial intelligence is a very complex domain that requires a lot of mathematical knowledge to understand fully, we should nevertheless be prepared to face this educational challenge in the academic environment of the university of applied sciences. AI topics need to be covered both in the curriculum of the undergraduate bachelor's degree program in radiography/radiologic technology and in the specialisms offered in master's degree programs. In the discussion, it is often stated that the professional group of radiology technologists should stick to their core competencies. But the fact is that core competencies of every profession must adapt in the course of social changes, driven by technological developments. An example of this is the bookbinder's trade, which has been completely replaced by automated technology (disruptive technologies). In radiology, AI enables support in decision-making or evaluation procedures as well as an emergence of interventional radiology.

Radiography may be under pressure, if we want to maintain it as a proper academic and professional discipline, to face these complex developments and either expand into new territory or to retreat in the face of disruptive technologies. I think it is clear that in interplay with many medical and non-medical professions, radiography should take the chance to develop further and integrate the new advances in medical data processing in order to support the overall health system in safety and effectiveness.

FH-Prof. Dr. Helmut Ritschl MA MSc
Head of the Institute of Radiologic Technology,
FH Joanneum University of Applied Sciences, Graz

„Despite all justified expectations, it is particularly the limits of artificial intelligence that are often completely misunderstood in public discourse."

Wolfgang Granigg

No way around it

Just like big data, cloud computing or data science, the term artificial intelligence is almost omnipresent. AI has found its way into numerous areas and disciplines and is regarded as a great hope for the future. AI is also increasingly being used in medicine, especially in the context of medical imaging.

The history of artificial intelligence (AI) is quite interesting. Initially, attempts were made to artificially represent intelligence by means of a manually defined set of rules, usually in connection with a knowledge database, to solve various problems. Although some progress and success were achieved, the high expectations were not fulfilled in this approach. In the meantime, AI has taken a different path. Instead of manually mapping a set of fixed rules, modern AI systems 'calculate' and 'learn' rules and patterns themselves from a large set of data (cf. 'computational intelligence' or 'machine learning') and store them implicitly, for example in an artificial neural network.

What does the omnipresence of AI mean for education – especially in medicine? Well, from my point of view, there is probably no way around dealing with AI as part of education in almost any discipline. Even if not everyone can or wants to use such models and algorithms themselves, it is at least useful to have a reasonably accurate idea of the possibilities, but also the limits of AI. In my view, this should be part of almost every curriculum – including medicine. Despite all justified expectations, it is particularly the limits of AI that are often completely misunderstood in public discourse, because there is one thing that artificial intelligence can currently not do – directly transfer rules and patterns learned in one domain or in one application to another domain or another application. Such intelligent transfers are currently (at least as far as we know) solely reserved for primates and especially humans.

MMMMag. DDr. Wolfgang Granigg
Head of Degree Programme 'Data Science and Artificial Intelligence' | Institute of Business Informatics and Data Science | FH Joanneum University of Applied Sciences, Graz

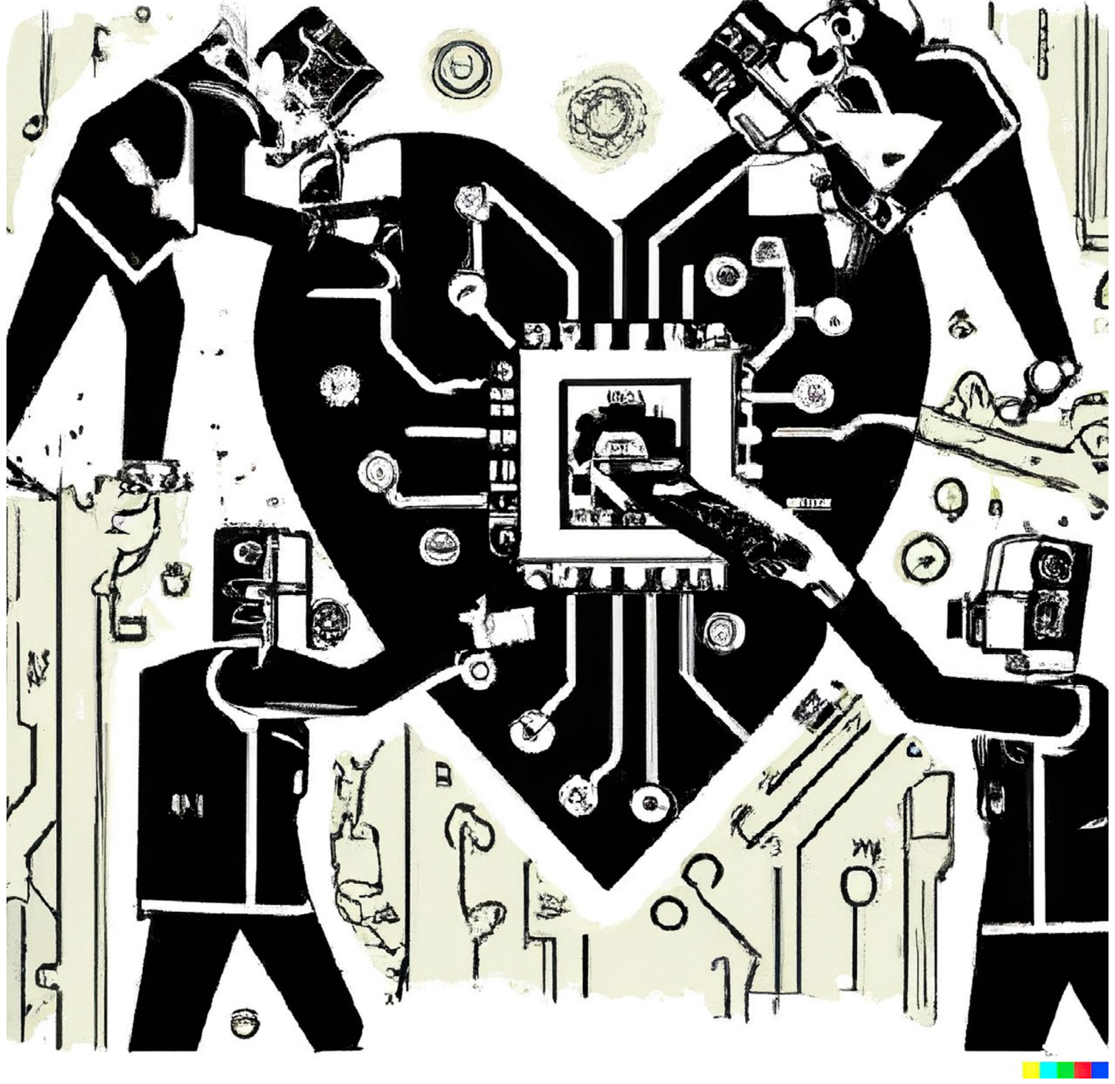

How can AI and ML solutions be meaningfully integrated into clinical care?

My job as a clinician informaticist is to serve as a bridge between clinical practice and cutting-edge technologies, such as Artificial Intelligence. My work as both a hospitalist at the Mayo Clinic and as a Medical Director at the Mayo Clinic Platform allows me to focus on how AI and ML solutions can be meaningfully integrated into clinical care. Healthcare is at an interesting juncture, where the speed of innovation from an AI perspective is outpacing our ability to effectively incorporate these technologies into practice in ways that improve outcomes, promote fairness, reduce bias, and ameliorate practice inefficiencies. As we innovate, we must remember that clinicians everywhere face enormous day-to-day challenges in meeting high patient demands and complexity, strained resources, and burdensome administrative obligations. When we think of AI/ML solutions, we need to pay special attention to these circumstances and pressures of day-to-day practice. Healthcare workers need solutions that solve real problems without complicating workflows or increasing administrative demands. And, most importantly, such solutions should improve patient outcomes while promoting safety and increasing fairness.

To effectively deliver upon these fundamental aims, we must first understand some of the limitations of AI/ML in today's healthcare landscape. Barriers to effective and appropriate AI are many, but we can broadly group many challenges into the categories of 'technical', 'process', and 'ethical'.

- **1. Technical:** From data storage systems, to computing power, to model development, there are many technical considerations that must be addressed when developing an ML algorithm – most of which I will not elaborate on here. It is worth noting a few key issues that occur in the AI/ML development phase that have significant impacts on algorithm efficacy and adoption. The first is data. Ideally, all algorithms would be trained on high-quality, complete, and accurate data that is diverse and free from bias. But this is seldom the case. Data quality is difficult to quantitatively and qualitatively assess in a comprehensive way. Many data repositories represent populations of a specific geography; this may limit the diversity of the datapoints represented. Further, data often represents the various societal and practice biases that exist on a systemic level. For example, if physicians happen to be implicitly biased in under-diagnosing a certain subgroup of patients with depression, the dataset that contains this subset of patients will not accurately represent the true incidence of depression in that population. The algorithms will be trained on biased data and may perpetuate those same biases. This can be extraordinarily difficult to address, quantify, and mitigate.

- **2. Processes:** Processes must also exist to incorporate algorithms efficiently into clinical practice. Once an algorithm is deemed appropriate, effective, and safe, it must then be integrated into clinical work flows in intelligent ways. If we have 100 high quality algorithms but all of them fire pop-up alerts right when a patient's chart is opened, the value of the algorithms' predictions is unlikely to be fully realized. Predictions and recommendations must be presented for the right patient at the right time and in the right way.

- **3. Ethical:** Algorithm development and implementation also face various ethical challenges that hinder widespread adoption. Complex mathematics and pattern generation underly algorithm generation. Humans are unable to perform and often unable to fully understand these computations, leading to the "black box" nature of AI. It is difficult for clinicians to fully trust a tool that lacks transparency, particularly when recommendations affect clinical management and treatment plans. Algorithm developers need to assure transparency of models so that end-users have sufficient information to assess and audit model performance. Further, models should be as explainable as possible – a task that can be very difficult to achieve. Ideally, information should be published that provides explicit evidence and/or explanation for why a model benefits users, why it should be accepted or trusted, and how it adheres to regulatory compliance requirements. Sufficient details surrounding the technical production of the model should be published, and "knowledge limits" should be established to provide conditional guard rails of when and where the model is appropriate for use. Finally, models need appropriate governance structures to facilitate auditability and accountability – all of which have yet to be standardized across healthcare systems.

These limitations, of course, show the way to a potentially very bright future of AI/ML in healthcare. We are seeing more cooperative data networks, where data owners can share their data within a federated system while retaining ownership of their data. Using advanced cryptographic methods, data can remain behind the firewall of an organization, preserving data privacy and security, while permitting innovators to create algorithms on that data from behind their own firewall, thereby preserving their intellectual property. Soon, we will have parameters we can use to measure algorithm fairness and bias, allowing greater transparency to clinicians and patients. We can also look forward to intelligent systems that deliver relevant predictions for the appropriate patient at the appropriate time in a manner that simplifies care, not complicates it. Finally, we will see a workforce transformation that will empower clinicians and healthcare workers to better understand the benefits and limitations of AI solutions so that they can more effectively implement them in care.

Sonya Makhni, M.D., M.S.
Senior Associate Consultant,
Department of Hospital Internal Medicine
Mayo Clinic Rochester MN

Imprint
responsible for content:
Human.technology Styria GmbH
Neue Stiftingtalstraße 2 | Eingang B
1. Stock 8010 Graz | Austria
Lejla Pock

Editors: Mag.a Eva Bucht und Pascal Mülner | HTS
Grafik: cardamom
Cover: DALL·E 2 prompt: “health industries cooperate with AI, photorealistic”
Printing: Medienfabrik Graz

Circulation: 500 print copies, 1650 e-Paper subscribers
Frequency of publication: twice yearly. Articles by named authors do not necessarily reflect the opinion of the editors or publisher. Subject to errors in setting and printing.

This information valid as of December 2022

QR Code
Read the botenstoff as an ePaper on your smartphone

human.
technology.
styria.

HUMAN.TECHNOLOGY STYRIA
CONNECTING COMPANIES
SUPPORTING BUSINESS
ENABLING GROWTH

human.
technology.
styria.

„HUMAN.TECHNOLOGY.STYRIA is the hub of a cluster of around 140 Styrian companies that work for human health. We act as a hub for contacts, know-how and information for this community. Our goal is to create economic added value for the cluster community, to further develop the strengths of the region in a targeted manner and thus to improve the international visibility of the location.“

humantechnology.at

humantechnologystyria

humantechnologystyria

European Innovation Partnership on Active and Healthy Ageing

SFG
NEUES DENKEN. NEUES FÖRDERN.