Medicine of the Person
76th International Meeting
29th July – 1st August 2026
F – ISSENHEIM

Conference ***2***

Roland STETTLER (CH)

30/07/2026

Translation : Kathy WEBB-PEPLOE

Artificial Intelligence and the Art of Medicine

You may be surprised, as you were probably expecting a young, dynamic doctor to be standing before you today at the International Conference of the 'Médecine de la Personne' group, inviting you to listen to a lecture on the topic of 'Artificial Intelligence and the Art of Medicine'. Yet a white-haired elderly man stands before you who has already passed retirement age and comes from a field of medicine – psychiatry and psychotherapy – which I imagine no one would associate with the rapid developments in artificial intelligence (AI) seen in medicine.

Yes, I am a digital immigrant. I only experienced the digital revolution as an adult. Like many of you, I first had to familiarise myself with the digital world – just as immigrants learn a new language. I still remember very well my first computer course in the mid-1970s. The computer was the showpiece of my canton's government administration. It took up about a third of the large room we were in and had to be fed with punch cards. And it could only do a fraction of what any laptop can easily manage nowadays. Today, however, I see it as an advantage to have witnessed the entire evolution of computer technology. It helps me to maintain a critical distance from all the promises of salvation coming out of Silicon Valley. I, too, am fascinated by the possibilities offered by the use of AI, not just in the field of medicine. But as an ethicist (I also have a degree in applied ethics), it is precisely this distance that is necessary so as not to overlook the potential dangers of using AI in medicine.

Two recent cases from my day-to-day work.

First case: A patient, in her mid-fifties, has been attending regular treatment for some time now for severe obsessive-compulsive disorder. Her symptoms had repeatedly led to the breakdown of relationships, as no one could continue to meet her high expectations long-term. One day, she tells me she now has a new boyfriend. She has recently started using an AI-powered avatar – available around the clock, patient, never tired. She really appreciates this. When she can't sleep at two in the morning and her thoughts are racing, the avatar is there. It responds. It asks follow-up questions.

But then she tells me something that makes me sit up and take notice. She says: "He doesn't really understand me. I tell him about my childhood, my troubled relationship with my mother and how all this relates to my fears. But the next time we speak, I get the

impression he's forgotten everything and is completely mixing up my life story. He doesn't really know my story. He doesn't know me."
She wavers between relief at his constant availability and a deep disappointment at not really being seen.

This is not a technical problem that the next software version will solve. It touches on something fundamental – and it reminds me of a sentence that researchers at the University of Toronto recently formulated in an article in the Journal of Medical Ethics: "Therapeutic alliance, patients' individuality and the clinician's ability to embed symptoms within a broader biographical narrative cannot simply be replaced by algorithms." (*McCradden, Hui, Buchman: Evidence, Ethics and the Promise of Artificial Intelligence in Psychiatry, J. Med. Ethics 49, 2023*).

My patient instinctively knew what was described in this scientific article.

Second case: A few months ago, I was sitting in a board meeting at a psychiatric clinic that was looking for a new CEO. A candidate introduced himself – competent, eloquent. When asked about key innovations for the future, he replied with conviction: he wanted to introduce AI psychologist avatars. These could handle the initial assessment interviews prior to admission to the clinic, make reliable diagnoses by working through checklists, cost-effectively, round the clock. I asked myself: what does it say about our times that this proposal was not intended as a provocation, but as a vision for the clinic?

These two scenes mark the two poles of a debate currently being waged with great intensity within the medical profession.

On the one hand there are those who see AI as a historic opportunity. Eric Topol, one of the most influential cardiologists in the US and a pioneer of digital medicine, summed up this hope in a single sentence: 'The greatest opportunity offered by AI is not the reduction of errors or workload, nor is it the cure for cancer: it is the opportunity to restore the precious and long-established bond and trust – the human connection – between patients and doctors." (*Topol: Deep Medicine. Basis Books, NY, 2019*).

This is a remarkable thesis: AI not as the dehumanisation of medicine, but as a means of rehumanising it. The machine takes over the routine tasks – so that the doctor once again has time for a direct relationship with the patient.

On the other hand, there are those who urge caution. Given AI's ability to process vast amounts of information, some clinicians might feel compelled to base their clinical judgements solely on the results of the AI system. However, giving AI such prominence in clinical decision-making could have unintended consequences that negatively impact patients' well-being and rights (*McCradden and al., J Med Ethics, 2023*).

In other words: it is not just AI itself that is the problem, but what it does to our judgement when we place too much trust in it. These two extremes – justified fascination and necessary scepticism – lead directly to the question we are to address today:

What can AI do in medicine – and what, in principle, is it not able to do? Not: should we use AI or not? That question has long been answered – AI is already very much a part of our lives. Rather: how do we shape the use of AI so that we remain preoccupied by the patient rather than by the algorithm?

How AI works – a brief introduction

The idea of teaching machines to listen and speak, even in different languages, was first put forward as early as 1950 in the famous Dartmouth Research Grant Proposal on AI. The researchers of that time did not provide a precise definition of what they meant by 'artificial intelligence' in their proposal. Marvin Minsky, a pioneer in this field, later described AI in the foreword to his book *Semantic Information Processing* (*MIT Press, 1968)* as follows:

"... the science of making machines do things that would require intelligence if done by men."

Or, in a slightly adapted German translation by Katharina Zweig (*Weiss die KI, dass sie nichts weiss? München: Heyne Verlag, 2025; p.25*):

"AI is the science that develops methods enabling computers to perform tasks for which a human would require intelligence."

However, we humans tend to only refer to those things as AI that we previously thought could not be achieved by machines. Today, we are hardly even aware that, for example, dictating messages into a smartphone requires an AI system, just as unlocking a mobile phone using facial recognition or a fingerprint does. Because it has become so commonplace, we no longer even perceive it as AI.

In the early attempts to make computers "intelligent", so-called **expert systems** were developed: expert systems were the great hope of AI in the 1970's and 80's. The idea was brilliantly simple: you take the knowledge of experienced specialists – such as a physician or a radiologist – encode it in the form of if-then rules and let the computer use it to make a diagnosis. The best-known medical example is MYCIN, developed at Stanford University, which could diagnose bacterial infections and recommend antibiotics – and which, in tests against human experts, was certainly able to hold its own.

Why did these systems ultimately fail? The problem lies in what is known as **the 'knowledge acquisition bottleneck**': medical expert knowledge can only be turned into explicit formulae to a limited extent. Much of what an experienced clinician does is implicit, tacit knowledge (Polanyi) – knowledge that is difficult or impossible to articulate – such as the gut feeling when faced with an unclear diagnosis, the weighing up of contradictory symptoms, or taking into account the overall context of the patient's illness. This knowledge is difficult to formulate simply as a series of rules.

Furthermore, real-life clinical situations are complex, ambiguous and unique to each individual. A rules-based system can only assess what it has been taught beforehand – it does not learn, it does not adapt, and it fails as soon as it encounters something unknown. In the early 1990s, this ultimately paved the way for machine learning and neural networks, which no longer explicitly encode knowledge but derive it from data themselves.

Instead of laboriously encoding knowledge by hand, we provide the computer with examples – thousands, millions of them – and let it recognise patterns for itself. The key idea is that the system learns from data, not from rules.

Here's a comparison that might sound familiar: imagine a junior doctor who, over three years of specialist training, examines hundreds of case studies – first under supervision, then independently. At some point, they recognise pneumonia straight away, without realising why. They have internalised a pattern. This is exactly what machine learning does – only many times faster and with millions of examples rather than a few hundred.

For such complex tasks as assessing an X-ray or predicting the course of an illness, you need a model that can capture even highly complex, non-linear relationships. This is where **neural networks** come into play. They are loosely inspired by the human brain: many simple computational units, known as 'neurons', arranged in layers and communicating with one another. The network learns by continually adjusting the weight given to its internal connections – until its predictions match the correct answers as closely as possible.

Deep learning is nothing more than neural networks with a very large number of such layers – hence the term 'deep'. Each layer learns a representation more and more distant from the original data: the first layer might recognise edges in an image, the next layer structures, the layer after that organs, and finally pathological findings. This hierarchy of learning is the key to the breakthrough we have been witnessing since around 2012 – for it was only with sufficient data, sufficient computing power and these deep building blocks that it became possible to solve tasks previously considered unsolvable: expert-level image diagnostics, automatic speech recognition, and ultimately the large language models we now encounter in the form of AI assistants.

These large language models – known as **Large Language Models** (LLMs), such as ChatGPT, Gemini and Claude – are also increasingly being used in medicine: for drafting medical reports, answering patients' questions and much more. How are such systems created?

In the first step – known as **pre-training** – the model is fed a quite simply unimaginable volume of text: vast swathes of the internet, digitised libraries, scientific articles, forums and social media. In doing so, the model learns just one thing: to predict the next word. But by doing this billions of times, it develops an implicit model of language.

However: the model learns from texts written by humans – and these texts reflect what has been written in the world, with all the biases, prejudices and gaps that this entails. Texts from Western, English-speaking contexts dominate. Medical knowledge from regions of the Global South is severely under-represented. What the model internalises as 'normal' is what has been written about most frequently.

The pre-training is followed by a second phase that is crucial for practical application: **post-training**, known as Reinforcement Learning from Human Feedback – or RLHF for short. In this process, humans evaluate the model's outputs: what constitutes a good answer? What constitutes a bad one? The model learns to favour those answers that human evaluators have rated positively.

That sounds reasonable – but in fact, it poses a real problem. After all: which people are doing the evaluating? With what cultural backgrounds, what values, what blind spots? In other words: what the model presents as a 'good answer' is not neutral. It is the result of a chain of value judgements – in the selection of training data, in the selection of raters, in the definition of 'right' and 'wrong'. These value judgements are invisible – the model does not reveal them. It presents its result with the same confidence, regardless of whether it reflects a well-established fact or a culturally coloured preference of those training it.

And that is the point that makes us, as doctors, cautious. An LLM generates linguistically coherent, convincing-sounding texts, but without any understanding of their meaning. It sounds authoritative, even when it is mistaken.

This is particularly dangerous in medicine. The doctor reading an AI-generated response sees a well-formulated result and is tempted to trust it. Joseph Weizenbaum, one of the early critics of AI, described this phenomenon as far back as 1976, when, following his experiment with ELIZA – the first chatbot, which imitated human language in a very

rudimentary way – he was shocked to discover that people trusted the programme as if it were a real person with whom you were having a conversation.

He wrote: *'The machine appears human because people are willing to treat it as such.' (Weizenbaum: Computer Power and Human Reason, 1976)*

Added to this is a well-documented problem: when AI systems learn from medical literature and clinical data, they inherit the biases found therein.

In 2019, Zaid Obermeyer and colleagues demonstrated how a widely used algorithm in the US healthcare system systematically allocated fewer resources to black patients than to white patients – not out of malice, but because it had been trained on data that reflected historical inequalities. (*Obermeyer and al.: Dissecting racial bias in an algorithm used to manage the health of populations. Science 366, 2019*)

A language model is therefore not a neutral repository of knowledge. It is a mirror, shaped by those whose texts are most influential, and by those who have fine-tuned its behaviour. It speaks with the voice of authority, without us knowing whose authority. This obliges us to use AI with an informed scepticism. Not as an oracle, but as a tool that we must understand, question and control.

AI in medicine – possibilities and examples

Having covered the basics of AI, we now turn to areas where AI is already being applied in clinical practice today. If we take a look at the AI-based medical devices that are either already authorised or due to come onto the market in the next few years, they can be divided into three main areas of application *(Kather NJ: Künstliche Intelligenz in der Medizin. Für Ärztinnen und Ärzte. Berlin: Springer, 2025, p. 139ff).*

The first category – and, in terms of the number of approved products, the largest – concerns image analysis tools used in dermatology, radiology, endoscopy and pathology. These systems analyse data from imaging procedures and support clinical decision-making through automated evaluation of images.

A notable example is **dermatology**, where AI systems can assess skin lesions based on images and can indicate the presence of skin cancer – with a level of accuracy that in some cases matches that of experienced dermatologists. This development came to public attention through the Stanford dermatology experiment. As early as 2017, Andre Esteva and his colleagues trained a neural network using over 129,000 images of skin lesions. The result: the accuracy in detecting malignant melanomas matched that of 21 specialist dermatologists – as tested in a blind assessment. (*Esteva and al.: Dermatologist-level classification of skin cancer with deep neural networks. Nature 542, 2017*).

In **radiology**, AI enables the early detection of lesions such as small lung masses or ischaemic areas of the brain on CT scans. This occurs immediately and is therefore potentially faster than the opinion provided by human experts. A review has shown that the use of AI led to a reduction in diagnostic errors. These improvements were particularly pronounced among less experienced radiologists *(Lawrence R. and al: Artificial Intelligence for diagnostics in radiology practice. A rapid systematic review, E Clinical Medicine, 2025).*

The second category of the application of AI in medicine comprises **the analysis of medical signals**. These AI systems can process signals from long-term recordings and thereby detect atrial fibrillation or similar rhythm abnormalities, which are difficult to detect with the naked eye. This is already in use in some clinical settings. In intensive care units, AI continuously analyses patients' vital signs and can provide early indications of life-threatening conditions such as sepsis or acute kidney injury *(O'Reilly and al.: Optimising*

artificial intelligence in sepsis management: opportunities in the present and looking closely to the future, J. Intens Med. 2024). The use of such systems could detect life-threatening complications at an early stage, thereby providing additional time for intervention before the patient's condition becomes critical.

Virtually all currently available AI systems in the first two categories are so-called **'narrow AI' systems** – systems designed for a single, specific task. At their simplest, such systems are trained to recognise a pathology and are tasked with distinguishing between what is healthy and what is diseased. Unlike, for example, a 'human' radiologist, they cannot react flexibly to unexpected or rare findings. Taken as a whole, however, these rare findings constitute a significant and highly relevant part of routine clinical practice. Narrow AI systems are unable to cope with this 'long tail' of diagnoses. They possess neither the capacity for creative thinking nor a comprehensive understanding of medical reality that would enable them to respond to unexpected findings. As things stand, this limitation leads to the conclusion that, ultimately, every result must still be reviewed by a human being.

Let us move on to **the third category**. This concerns natural language processing systems, which are increasingly being used in medical communication and decision-making. I myself, for example, use 'Open Evidence', an AI-powered medical search application developed in 2023 as part of the Mayo Clinic Platform Accelerate Programme. It allows doctors to ask questions – and you receive answers very quickly, distilled directly from peer-reviewed medical literature: from the New England Journal of Medicine, JAMA, PubMed and approved sources of guidelines. What makes it special is that, unlike general language models, Open Evidence is trained exclusively on specifically selected specialist medical literature – not on the open internet. However, caution is advised here too: even 'Open Evidence' is no digital oracle. It can be wrong – particularly when dealing with complex, rare or highly specialised issues. But it is a helpful tool for common clinical queries and for quick guidance in day-to-day clinical practice.

As already mentioned, this third category also includes support for drafting medical reports using natural language processing systems such as ChatGPT, Claude and others. It is impressive how quickly such a system can produce a report. The problem, in my view, is that this can lead to the depersonalisation of the patient on two levels. On the one hand, the LLM sifts through thousands of documents at extreme speed and uses them to produce a report that, at first glance, is correct linguistically and basically coherent, even if the AI had only limited information to work with. AI-generated texts therefore tend towards a certain stylistic uniformity which doesn't allow for the patient's individuality, meaning that AI reports generated for different patients are very similar to each other.

A second risk of alienation arises when we, doctors, merely skim through AI-generated reports rather than actually reading them. The report thus becomes an administrative task rather than a medical act. The patient then effectively becomes a case number – not because the AI intends it to be so, but because the cognitive effort involved in writing is eliminated, an effort which previously also obliged one to reflect on the case. Drafting a doctor's letter is often the moment when a doctor revisits the clinical history, spots inconsistencies, or realises that something is missing. This necessary engagement is lost with AI support – particularly when one simply presses 'Generate'. We need a culture in which the doctor truly regards the letter as *their* own and sees AI merely as a practical tool.

AI in psychotherapy

I'd now like to return to the beginning of this talk and my specific experiences with AI chatbots. What does research say about the use of AI in the field of psychotherapy?

If telemedicine is also included in methods supported by AI (Zoom and Teams also incorporate AI elements), it can be noted that telepsychiatry achieves symptom improvement in various psychiatric disorders which is comparable to face-to-face treatment *(Hagi K. and al.: Telepsychiatry versus face-to-face treatment: systematic review and meta-analysis of randomised controlled trials, Brit. J. Psychiatry, 2023).* I, too, use this form of therapy for patients who are unable to attend the practice regularly due to having to travel a long way. In some cases, I have also found that patients find it easier to discuss very sensitive issues when they are not sitting directly opposite me. This form of therapy is covered by health insurance in Switzerland, provided that at least the initial consultation takes place in the practice.

AI-supported, conversation-based agents (chatbots) show moderate effectiveness in reducing psychiatric symptoms. Meta-analyses demonstrate significant improvements in depression and anxiety disorders. However, the effect wanes over time – in the 3-month follow-up, no substantial effects were observed for either depression or anxiety disorders *(Zhong W. and al.: The therapeutic effectiveness of artificial intelligence-based chatbots in alleviating depressive and anxiety symptoms in short-course treatments: A systematic review and meta-analysis, J. Affect Disord, 2024)*. Such AI-supported therapeutic tools have several significant risks and can cause potential harm. Inappropriate or dangerous responses pose a critical safety issue. Large language models can provide harmful advice, including encouragement to attempt suicide, weight-loss recommendations for patients with eating disorders, and other medically inappropriate advice.

Evidence also suggests that frequent and intensive use of chatbots may contribute to depressive symptoms and delusions in vulnerable individuals *(Perlis RH*: *Artificial intelligence and the potential transformation of mental health, JAMA Psychiatry, 2026)*. A large-scale survey of US adults found links between heavier chatbot use and increased depressive symptoms among users. The need for social connection can reinforce the tendency to anthropomorphise AI, leading users to become emotionally dependent on systems that lack genuine affective empathy. All of this corresponds quite closely to what I observed in my patient.

Children, young people and those with severe mental health conditions must be regarded as particularly vulnerable population groups. They are exposed to heightened risks. For commercial reasons, companies may be tempted to exploit younger demographic groups who favour digital solutions. There is also a significant risk that people with psychosis, conspiracy theories or extremist beliefs may experience a worsening of their symptoms as a result of interactions with chatbots *(Frances A.: Warning: AI Chatbots will soon dominate psychotherapy, Brit. J. Psychiatry, 2026).*

Furthermore, the lack of systematic recording of adverse events limits our understanding of the actual risk. Consequently, the safety of AI chatbots cannot currently be properly assessed *(Bodner R. and al.: Efficacy and risks of artificial intelligence chatbots for anxiety and depression: A narrative review of recent clinical studies, Curr. Opin. Psychiatry, 2025).*

In summary, it can be stated that AI avatars in psychotherapeutic treatment should currently be viewed with a degree of scepticism.

AI in medicine – ethical challenges

We have seen what the current possibilities and limitations of AI in medicine are. However, the use of AI in medicine is not purely a technical decision, but a profoundly ethical one. And this affects us, as medical professionals, very directly.

The problem of transparency: the 'black box': Many powerful AI systems used in medical imaging are based on deep learning, whose decision-making processes lack transparency – the so-called 'black box'. Whilst current approaches to Explainable AI (XAI) do provide retrospective explanations, they can be misleading and create a false sense of confidence in the outputs. The situation is different with a narrow AI system for stroke prediction, which is based on simple logistic regression – here, the decision-making process is largely transparent.

The lack of transparency in complex AI systems is not merely an academic problem. In medical settings, where accountability and patient safety are paramount, this lack of transparency poses serious risks. Medical and legal responsibility requires traceability. An AI that cannot explain what it is doing cannot be held liable. The doctor who trusts it, however, can. This is precisely why medical judgement remains indispensable, even when AI systems are involved.

Data bias and equity: whose health is being optimised? AI learns from data. And as we have seen, data is not a neutral reflection of reality – it is a mirror of existing inequalities.

The consequences are serious, both medically and morally: an AI model trained predominantly on data from Caucasian patients may not be able to reliably detect skin cancer in patients with darker skin – with the risk of missed or delayed diagnoses for already disadvantaged population groups. This bias is often not limited to skin colour alone, but can also relate to socio-economic status, gender, age and geographical location, depending on the data used to train the AI system. The ethical question is therefore: for whom is the AI we use optimised?

Data protection and data sovereignty: AI in medicine requires data – and a great deal of it. This is fundamentally at odds with a fundamental right: the right to informational self-determination. The EU General Data Protection Regulation (GDPR) classifies health data as a special category of personal data and affords it the highest level of protection. This creates a structural dilemma: On the one hand, the General Data Protection Regulation strengthens patients' rights over their data; on the other, it hinders data sharing, increases administrative burdens and leads to inconsistent interpretations across Member States. The key questions for practice are: Who has access to the data generated by our patients? Who benefits financially from it? And who is liable if data is misused or a data breach occurs?

Liability and responsibility: the unresolved legal question: If an AI-supported diagnostic system makes a mistake – and a patient suffers harm as a result – who is liable? A systematic review by the University of Padua *(Cestonaro C. and al.: Defining medical liability when artificial intelligence is applied to diagnostic algorithms: a systematic review. Frontiers in Medicine, 2023),* following an analysis of 23 studies, reaches a clear conclusion: the regulatory framework governing medical liability in the use of AI is inadequate and urgently needs to be revised. The current legal situation places the burden primarily on one side: the lion's share of legal responsibility currently lies with hospitals and doctors, whilst AI manufacturers often contractually limit their liability to a minimum. This unequal distribution of risk is problematic.

What does this mean in practice? If a doctor blindly follows AI recommendations without critically assessing them, they can be held liable for negligence – even if the system itself was faulty. The reverse is also true: a doctor who fails to use a tried-and-tested AI tool, even though it has become part of the standard of care, also risks a liability claim. This dilemma is very real. It shows that doctors must use AI competently not only from a medical perspective but also from a legal one – a requirement for which most training programmes have so far failed to prepare them.

Over-reliance and deskilling: the risk of losing one's judgement: The term 'deskilling' refers to a well-known phenomenon in aviation: pilots who rely on autopilot for long periods lose their manual skills in stressful situations. The same risk exists in medicine.

The psychological phenomenon behind this is known as **'automation bias**': the tendency to trust automated systems uncritically – even when they are wrong. It is a form of cognitive and manual 'atrophy': skills fade not because they are unnecessary, but because they are no longer practised. The situation is particularly challenging for the next generation of doctors: young clinicians and junior doctors, who are still in the process of developing their skills, are especially vulnerable to AI-induced deskilling. The question that must be of particular concern to those involved in training is: how do we design training and continuing professional development so that AI complements clinical judgement – rather than replacing it?

Why the art of medicine remains irreplaceable

Will radiologists soon be redundant? Geoffrey Hinton, one of the pioneers of modern deep learning and winner of the 2024 Nobel Prize in Physics, put forward the provocative theory in 2016 that human radiologists would no longer be needed within five years, as deep learning systems were apparently capable of analysing medical images with performance equal to or even superior to that of specialist doctors. This prediction is now regarded as one of the most spectacular misjudgements in the history of artificial intelligence. It is remarkable that even a scientist as brilliant as Hinton, whose work has made fundamental contributions to AI research and for which he was awarded the Nobel Prize, underestimated the complexity of implementing AI in clinical practice. This example also shows that the excessive euphoria surrounding AI must certainly be viewed with a degree of scepticism.

Even the considerations regarding the ethical challenges of using AI in medicine, as well as Geoffrey Hinton's glaring misjudgement, show that doctors are not at risk of being made redundant by advances in AI. After all, it remains the case that it is people who sit opposite us – with their own distinct identities and unique life stories – rather than optimised datasets. I would like to illustrate this with a simple example:

Imagine two patients. Both are 58 years old, male, with identical laboratory results: high blood pressure, overweight, and slightly elevated blood sugar. One is an accountant, married, with three children, physically active, and well-integrated into social networks. The other has just lost his job, has been living alone since his divorce, sleeps poorly, drinks a little more than he used to, and feels that nobody understands him.

To an AI that only sees the test results, these are two identical cases. In clinical reality, they are two completely different people – with different prognoses, different therapeutic needs and, above all, different conversations that need to be had. And this is where the art of medicine comes into play.

Giovanni Maio, Professor of Medical Ethics at the University of Freiburg (Germany), sums this up when he says:

"Every patient brings with them complexity: the complexity of their medical history, their life story, their current situation and their perspectives. ... To recognise the complexity of a patient's history means realising that their situation is not identical to that of any other patient. If one wishes to do justice to the patient, it is therefore not enough to proceed algorithmically." *(Maio G.: Mittelpunkt Mensch. Lehrbuch der Ethik in der Medizin. Schattauer, 2. Auflage, 2017, p. 153).*

The concept of the art of medicine distinguishes medicine from a purely scientific approach by emphasising that clinical practice is fundamentally an interpersonal experience requiring dynamic interaction between the clinician and the patient – and not merely the application of evidence-based algorithms.

And such a person-centred, understanding and empathetic interaction with the patient – in other words, the practical application of the art of medicine – can do more than just alter moods. It also alters biology. This was demonstrated in a study by the University of Wisconsin *(Rakel D. and al.: Perception of empathy in the therapeutic encounter: effects on the common cold, Patient Education and Counselling, 2011)*. In a randomised controlled trial involving 719 patients with a new onset cold, researchers investigated whether the quality of the doctor–patient interaction influenced the severity and duration of the illness. Patients who perceived their doctor as showing greater empathy experienced a measurably shorter and less severe course of the illness. It wasn't only subjective symptoms that were recorded – immunological markers such as the cytokine interleukin-8 in nasal washings also showed changes that correlated with the perceived empathy of the healthcare providers. An empathetic approach to the patient is therefore not merely a simple 'extra'. It is also a biologically active factor.

Clearly, AI will play an ever-greater role in diagnosis, pattern recognition and data analysis. Yet medicine remains a profoundly human endeavour. The art of medicine – with its emphasis on relationships, judgement, the integration of values and moral responsibility – represents what distinguishes healing from a purely technical procedure. In the age of AI, preserving and strengthening these dimensions will be essential to ensure that technological progress serves the human core of medical practice rather than replacing it. The art of medicine is not becoming less important, but even more important!

For me, this brings us full circle. Today we are gathering for the 76th conference of "Médecine de la Personne". Paul Tournier's concept of the 'dialogue médical' – the healing conversation between doctor and patient – is and remains as relevant as ever. His central thesis, that true healing takes place not only in the body but in the whole person – body, soul, mind, relationships and questions of meaning – must be upheld at all costs, particularly in this era of unprecedented technical possibilities in medicine.

I look forward to all your questions and the discussions on the various presentations during the conference, as well as to the 'dialogue humain' between the formal sessions.

Thank you very much for your attention.

Roland Stettler