Assistive Systems: A Paradigm Shift in Technology and Medical Care

Michael Meyer, Arthur F. Pease, and Michael Lang

Siemens AG Germany, Healthcare Sector
Karlheinz-Kaske-Str. 2, 91052 Erlangen, Germany
michael-meyer@siemens.com

Abstract. Medical diagnostic procedures are overwhelming physicians with data. In response, doctors are turning to computer programs that help them assess and interpret results. Newly-developed software is now providing fast, accurate decision support. An army of algorithms is being developed. Built on expert knowledge and capable of learning from experience, these systems are pointing out anomalies in radiology exams, providing decision support in a range of fields, and optimizing split-second decision making in high-speed industrial processes. As such systems harvest knowledge, there will be no limit to what we will learn from them.

1 Introduction

For most industrialized countries, the demographic, structural and social changes are resulting in a higher proportion of elderly people within the population. The social and financial impact of this change cannot be met by today's medical system and the public institutional care systems. Hence the transition of elderly people into medical care centers and institutional care has to be reduced or delayed, and work intensive tasks within these facilities have to be facilitated. Assistive systems are therefore demanded at the elderly people's home as well as in institutional medical care.

An Ambient Assisted Living (AAL) environment is an integration of stand-alone assistive technologies, with elements of smart homes, and telehealth services. AAL refers to technical systems that support people in need of assistance in their everyday lives, to promote personal independence and to improve the availability and quality of health and social care services in the home. New market potentials are created at the interface between traditional sectors such as health services, care, medical technology, housing and the IT industry on the other. The use of ICT and the establishment of networks between everyday objects, social players and the welfare system mean high requirements have to be met regarding the reliability, user acceptance and serviceability of the IT based solutions.

One important application is a continuous monitoring of activities and behavior of the elderly user based on diverse sensor modalities. This enables the system to recognize and assess the current situation, to provide assistance and to call external support if needed. In this context robots can be of use in order to provide additional sensing capabilities. In addition, it is expected that future intelligent robots will provide

D. Tavangarian et al. (Eds.): IMC 2009, CCIS 53, pp. 1–11, 2009.

physical assistance tasks in home and care facilities. They will also feature a close integration into intelligent environments.

2 Harvest without End[1]

Patterns previously invisible to machines and humans are today providing insights that make medical treatments increasingly personalized and effective, production more customized and efficient, and intelligence — whether in a security camera or a picture archiving system — more distributed and flexible. Across the board — from health care to energy management, and from finance to security and sales — information is being mined from machines, processes and experts, and crystallized into machine knowledge used by algorithms. These algorithms, which range from systems that can interrogate cardiac data for anomalies to the analysis of sales information to predict a customer's probability of consummating an order, are becoming our invisible assistants.

2.1 Experts Inside

Regardless of the class of problems they are engineered to solve, assistants provide support in an area humans are ill equipped to deal with: discovering trends in huge databases. In the medical area, for instance, this process begins with data mining. "We are taking various patient data sources, mining them to build predictive models, and embedding the results in applications that allow physicians to dynamically interact with the information in a computer aided detection (CAD) environment," says Alok Gupta, PhD, vice president of the CAD and Knowledge Solutions Group at Siemens Medical Solutions (SMS) in Malvern, Pennsylvania.

For SMS, the spot where this avalanche of data converges is a comprehensive knowledge platform for medical decision support called the Remind (Reliable Extraction and Meaningful Inference from Nonstructured Data) platform. The ultimate invisible assistant, "Remind will make it possible to dynamically integrate medical images, in-vitro diagnostic information, and genetic information into a patient's profile, providing personalized decision support based on analysis of data from large numbers of patients with similar conditions," explains Bharat Rao, PhD, senior director of Knowledge Solutions for Healthcare Providers at Siemens Medical Solutions in Malvern and inventor of the Remind platform. Remind adds up to a diagnostic crossroads for Siemens' imaging-related businesses and its more recently acquired in-vitro businesses, now known as Siemens Diagnostics (for more, see Pictures of the Future, Spring 2007, p. 54). "The vision is to integrate the information from imaging and laboratory tests into a single database, and eventually a single patient record," says Gupta.

On the long road to realizing the Remind vision, Siemens is developing an army of invisible assistants designed to support physicians as "second readers."

The idea is that once a specialist has examined a scan, he or she can run the appropriate assistant to increase the probability that nothing has been missed. Known as

[1] This article by Arthur F. Pease was originally published in the Siemens brochure "Pictures of the Future", pages 89-91, Spring 2008.

knowledgedriven products, these assistants (which plug into Siemens' syngouser interface) offer computer aided detection of lung nodules, colon polyps, breast lesions, and much more.

Other assistants support physicians in accelerating the process of accurately quantifying functions such as cardiac ejection fraction and vessel flow abnormalities, and in providingcomparative analysis of images produced at different times and from different imaging modalities.

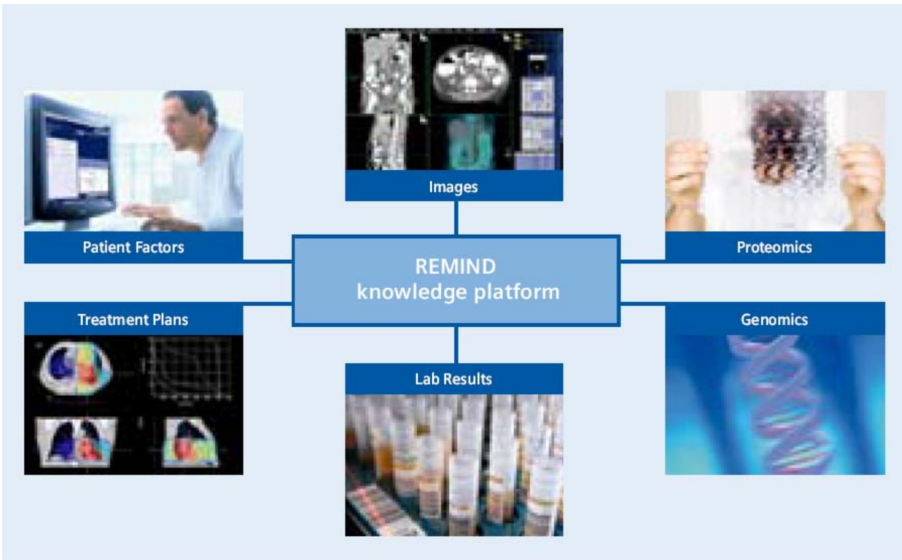

Fig. 1. By combining different sources of medical information in a single database, the Remind platform will support the creation of new, specialized decision-support assistants

Among the many assistants heading for commercialization is one that extracts a 4D model (3D over time) of the aortic valve from ultrasound data "that will allow physicians to interrogate it regarding a variety of real-time, quantitative functions," says Helene Houle, a senior sonographer with Siemens Ultrasound in Mountain View, California, who worked closely with Siemens Corporate Research in Princeton on its development. Another assistant now under joint development will create a 3D interactive model of the heart from computer tomography (CT) data. The model, now in prototype, will display the outlines of the beating heart and provide information regarding anomalies in the volume of blood pumped by the atria.

But such assistants are just the beginning. "We are looking at what it would mean to add genetic information to the imaging data in these products," says Gupta. With this in mind, Siemens is working with an expanding group of medical centers in the context of the EUfunded Health-e-Child program (see Pictures of the Future, Spring 2007, p. 72). The program, which is coordinated by SCR and the CAD group, is developing an integrated health-care platform for pediatric information designed to provide seamless integration of traditional sources of biomedical information, as well as emerging sources, such as genetic and proteomic data.

2.2 Voice Command

As medical assistants multiply and their underlying databases expand, new systems of addressing this cornucopia of information will be needed. One solution that is approaching market introduction in 2008 is Automatic Localization and Parsing of Human Anatomy (ALPHA). Trained on a huge anatomical database and capable of learning with each exam, ALPHA recognizes landmarks throughout the body, thus opening the door to voicebased interaction. “Questions such as ‘show me the lower left lobe of the patient’s lung and compare it with the previous two exams,’ will become routine,” says Arun Krishnan, PhD, head of CAD research and development at SMS in Malvern. “This will accelerate throughput, because it will no longer be necessary to search through image sets to find a desired anatomical slice. The target will appear automatically in response to a voice command,” he says. Compatible with hospital picture archiving and communication systems, ALPHA will provide a quantum leap in terms of the rapid accessibility of CT, MR, PET, and other imaging modalities and their content.

Understanding the complex meanings and information locked in images is a topic that is also being examined by Theseus, a German Federal Ministry of Education and Research project led by Siemens. “A big part of the Theseus vision is to automatically recognize image data in order to transform it from an unstructured to a structured state, so that it can be used in the semantic Web for retrieval,” say Dr. Hartmut Raffler, coordinator of Theseus and head of the Information and Communications Division of Siemens Corporate Technology (CT). Adds Dr. Volker Tresp, who is responsible for day-to-day management of Theseus and is a specialist in data mining, machine learning and decision support at CT, “This is a vast area because it opens up the entire field of picture, video, multimedia and content archives for deep exploration as they relate to security, robotics, entertainment, environmental sciences, and much more.”

Specifically, a research area within Theseus known as “Medico” is building an intelligent, scalable picture archiving and search system (that could be supported by ALPHA) capable of retrieving images by content. Suppose, for in stance, that a cardiologist is examining MR images of a patient with a pulmonary valve deficiency. “To help determine whether the deficiency warrants surgery, he might ask Theseus to show him images of pulmonary valves that look similar to the one he is looking at in terms of morphology and function before and after surgery,” says Dr. Dorin Comaniciu, head of the Integrated Data Systems department at Siemens Corporate Research and one of the initiators of Theseus Medico.

2.3 Communicative Cameras

But the areas of application for this kind of search engine extend well beyond medical uses. Says Ramesh Visvanathan, PhD, head of the Real-time Vision and Modeling Department at SCR, “In the context of the Theseus project, our Vision Center of Competence in Munich is defining metadata languages for the automatic identification of video content. In terms of security applications, for instance, this will mean that cameras will be able to track a target of interest by describing it in a standardized language and passing the information from one camera to another." The technology would thus make it possible to follow an intruder as he or she leaves one camera's field of view and enters the area monitored by another camera.

And what about the quality of the images that intelligent systems select? Regardless of whether an image originates in a surveillance camera or a medical database, the highest quality must be guaranteed if the evaluation of its content is to be reliable. Image retrieval systems therefore need a way of ensuring selection of the best available images. Work now in the pipeline at Beijing's Tsinghua University that is sponsored in part by Siemens may provide an answer. "The idea is to develop an assistant that will select the best and most relevant images for doctors from large data sets," says Comaniciu. Trained by using the criteria doctors themselves use for selecting images, the assistant may even be able to enhance images that are less than perfect.

2.4 Algorithms and Automation

Just as intelligent assistants are rapidly reproducing in the health-care universe, they are also beginning to populate other areas — particularly in industry. In steel production, for instance, the trend toward total automation is leading to increasing use of decentralized intelligence.

"Depending on the grade of steel, the manufacturing components involved may have individual strategies for monitoring and managing each step while taking a collective view ofthe process," says Dr. Michael Metzger, a specialist in steel industry solutions at Corporate Technology in Munich. He explains that this boils down to the use of "algorithms stationed near associated actuators working together to solve a control problem within a community of machines." Such systems must, furthermore, be able to learn at lightning speed. "In order to accomplish this," says Metzger, "these systems are based on control and optimization process models that are themselves based on relationships derived from physics and expert knowledge. But they must also be able to learn from the huge amount of data produced by an automation system, thus enabling the control system to respond optimally in real time to variables such as rolling force and temperature," he explains.

As in health care, a process of customization is in full swing here. This begins with expert knowledge and data mining, which discover key parameters, such as deformation historyand cooling rate for a given grade of steel. Then, to optimize results for a particular order, the entire production process is simulated — including neural networks and learning algorithms. Once optimized in the virtual world, the information is transferred to the rolling mill and put to the test. Values for each process step are taken and compared against the simulated (and thus optimized) values. "As a result," says Metzger, "the models learn how to improve themselves based on this comparison. Ultimately," he adds, "such systems will provide decision support and finally decision automation."

2.5 Digital Repairmen

Not only do learning systems keep track of what works best under a complex variety of circumstances. They also keep an eye on the long-term factors that cause machine wear and tear, and predict when service should be performed with a view to minimizing downtime. With this in mind, in 2007 Siemens established a strategic program called the Machine Monitoring Initiative. ”The project will tap basic research throughout the organization in data mining, learning systems and decision support,” says Claus Neubauer, a data integration specialist at SCR. The results will be used to automate the prediction and scheduling of maintenance on everything from power, rail and communication networks to MR scanners and windmill gearboxes.

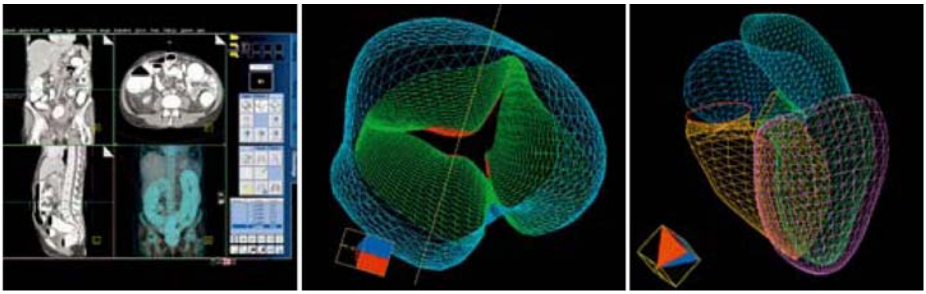

Fig. 2. Medical assistants recognize anomalies in the intestine (left), in the function of the aortic valve (center), and in the amount of blood pumped by the atria of the heart over time

Predicting when machines will need maintenance and which parts will need to be replaced may sound like a tall order, but what about predicting whether a customer will actually purchase a wind park or an MR scanner? Surprisingly, agents are already zeroing in on this kind of information as well. Research conducted at SCR has come up with an agent technology that is “70 to 80 percent accurate,” says Amit Chakraborty, who leads the Modeling and Optimization program at SCR. “In providing this decision support, the agent takes many factors into account, including customer reliability, competitors, and sales force information,” he adds.

2.6 Intelligence Everywhere

Naturally, given the fact that they are weightless, not particularly expensive to produce, and capable of incrementally increasing the productivity of hardware, invisible agents will eventually pop up just about everywhere. The trend toward decentralized intelligence in highly automated production facilities will have its counterparts in traffic and rail management, building and home automation, safety and security technology, power generation and distribution, and of course health care. The implications of these invisible entities for entertainment, information accessibility, security, environmental protection, and the way humans communicate, organize, work and live could be profound.

“We should keep in mind that this is all about solutions that support human activities,” says CT’s Raffler. “Based on this, agents will understand what we are looking for, present results more intelligently than is now possible, answer questions, deal with large bodies of unstructured data, compose services, and propose new processes for solving problems.”

Information, something we produce more of with every passing second, will become increasingly valuable as we learn to mine it, combine its streams, and refine its messages. What lies ahead, in short, is a harvest without end.

3 Digital Decision Support[2]

The Maastro Clinic is a leading cancer treatment facility located in the vicinity of Maastricht University in the Netherlands. The clinic's Radio Therapy section has a friendly reception area where patients referred from numerous other Dutch hospitals await cancer screenings, follow-up treatments, and treatment simulations. To provide the best possible treatment for these patients, and improve cancer research, the facility houses an interdisciplinary team of radiation therapy specialists, biologists, physicists, and computer scientists, as well as experts from Siemens Healthcare, all of whom have access to high-tech medical equipment and state-of-the-art software.

A key member of the team is Professor Philippe Lambin, a radiation oncologist who is medical director of the Maastro Clinic. "We're conducting research on a computer-aided decision-support system for personalized treatment of patients with lung cancer," Lambin explains.

"We're doing this because a study carried out by Maastricht University revealed that most doctors are unable to reliably assess how well their treatments are working, and therefore have difficulty in choosing the right treatment. We plan to improve predictions of the effectiveness of radiation therapies with the help of sophisticated software." The software Lambin is referring to is based on Remind, a data mining comtool from Siemens (Pictures of the Future, Spring 2006, p. 9).

Remind (Reliable Extraction and Meaningful Inference from Nonstructured Data) statistically analyzes all types of medical information, including everything from physicians' letters to medical images and laboratory diagnoses, and then identifies specific patterns. A research prototype system tested at the Maastro clinic is able to predict the two year survival rate of lung cancer patients with high accuracy. The twoyear survival rate is used by doctors to assess the success of individual radiation treatments. At the moment, 47 percent of all lung cancer patients survive for the first two years after diagnosis, if their cancer is detected at an early stage.

The first commercial application of Remind is Soarian Quality Measures software, which can measure quality of care from patient records based on established standards. At the Maastro Clinic, however, Remind is being optimized for cancer research in a research project that calls for Siemens experts to work with clinic specialists on site.

The system requires as much medically relevant patient data as possible to issue statistically meaningful prognoses. Such data includes sociological information on the individual in question, measurements taken with imaging methods, and biological data such as cell division biomarker analyses. Remind analyzes and links more than 100 of these parameters. In this research project, Remind then computes the likelihood of two year survival, and the risk of side effects, for various treatment options

[2] This article by Michael Lang was originally published in the Siemens brochure "Pictures of the Future", pages 92-94, Spring 2008.

for a patient. The intention is to help physicians select the optimal treatment for each individual patient. capability and radiation sensitivity, which can be determined through gene and blood biomarker analyses. Remind analyzes and links more than 100 of these parameters.

In this research project, Remind then computes the likelihood of two year survival, and the risk of side effects, for various treatment options for a patient. The intention is to help physicians select the optimal treatment for each individual patient.

3.1 Combining Diagnostics and Treatment

Maastro physicians can use state-of-the-art Siemens technology for their diagnoses and treatments. For instance, a combined positron emission tomography (PET) and computer tomography (CT) scanner makes it possible to obtain 3D images of the lungs in spite of breathing- related movement — a must in the case of lung cancer patients. The PET unit uses a lowradiation marker substance to provide crosssectional images that depict biochemical and physiological processes, while the CT details the anatomy and location of the tissue being studied.

The combination of both technologies provides doctors with information on the type of tumor they're dealing with, as well as its precise shape and position. When it comes to treatment, one of Lambin's preferences is adaptive radiotherapy. This Siemens solution provides oncologists with a 3D data set of the patient, which allows them to optimally adapt a radiation procedure to the position and size of the tumor in question.

Here too, Remind supports physicians with prognoses and treatment planning assistance by assessing the results from a database of post-treatment examinations. Lambin describes this combination of diagnostics and treatment therapy as "computer-aided theragnostics."

Another aspect of this research project is to configure Remind to predict the relative probability of typical radiation therapy side effects, such as esophagitis (perforation of the esophagus). This is achieved on the basis of parameters such as radiation dose, treatment time, concomitant chemotherapy, and the concentration of white blood cells. The intention is that such a tool helps doctors to recognize early signs of esophagitis, thus avoiding premature discontinuation of therapeutic treatment.

The next step in the research project will be to include the costs of potential complications associated with the therapy in question. Lambin's primary goal for 2008 is to broaden the system's database. "To make a fairly accurate prediction of the survival rate after a specific therapy, we need to have at least 500 to 1,000 patients in our database," he says. "We also need an external dataset to validate the predictions — that's the bottleneck."

The Maastro research database now contains data on approximately 1,000 patients, 500 of whom were diagnosed with lung cancer. In order to expand this database, the Maastro Clinic has plans to establish a digital network link with hospitals in Leuven and Liège in Belgium, and in Groningen in the Netherlands. Due to data security considerations, however, Maastro's Remind system will only be given anonymous parameters via the link; the data itself will remain in other clinics. The resulting broader research database will make a completely new type of clinical research possible. This is because specialists in Maastricht plan to use the data to simulate clinical studies, much in the same way the pharmaceutical industry uses machine-learning-based software to simulate experiments.

3.2 Digital Radiology

Dr. Marco Das works in the Department of Diagnostic Radiology at Aachen University Hospital, which is located around 40 kilometers from the Maastro Clinic. The focus of his work is the detection of growths in the lungs, such as cancers, metastases, and benign tumors. The clinical routine here involves using CT to create 3D data sets of the lung, after which Das searches for suspicious-looking structures in digital images. Das examines 30 to 40 patients this way every day, meaning that he only has a couple of minutes for each diagnosis.

To raise the probability that no tumor is overlooked, a second radiologist double-checks all of his findings. Das also utilizes CAD (computer- aided detection) software that may eliminate the need for a second radiologist, as is already the case at many hospitals. CAD is a technology based on pattern recognition and not on artificial intelligence. CAD systems for lungs analyze differences in thickness in lung tissues and compare these with stored images of typical lung tumor patterns. They are therefore able to recognize such patterns in other CT images as well.

3.3 Tumor Marker

“All of this functions very well in practice,” says Das, who uses syngo Lung- CAD software from Siemens. The system examines lungs for tumors even before a radiologist has finished making his or her assessment. It takes the software only around four minutes to check up to 700 image slices, each of which is one millimeter thick — and it works even faster with thicker layers and a correspondingly lower number of images.

After Das completes his diagnosis, he analyzes the results produced by the software, which means there’s no waiting time in between. The software automatically marks suspicious areas with red circles. “All studies to date show that CAD software has had a positive effect on the accuracy of radiologist diagnoses,” says Das. The system does make errors, however. These take the form of false-positive diagnoses, which, according to Das, don’t cause any major problems, since they can quickly be spotted by an experienced radiologist.

“CAD programs are very good as second readers, but they’ll never replace radiologist diagnoses because a doctor’s experience is the key to evaluating results,” says Das. An additional advantage offered by the new syngo CT Oncology software — which includes syngo LungCAD functionality — is that it helps to accelerate diagnostic decision making, according to Das. For instance, doctors need to measure changes in tumor size in order to determine whether a treatment is working. Until recently this was done by manually calculating a tumor’s diameter onscreen. Such measurements are extremely imprecise, however, and can vary from doctor to doctor. Syngo CT Oncology, on the other hand, improves measurement accuracy by automatically calculating the volume of all different types of tumors. It also enables

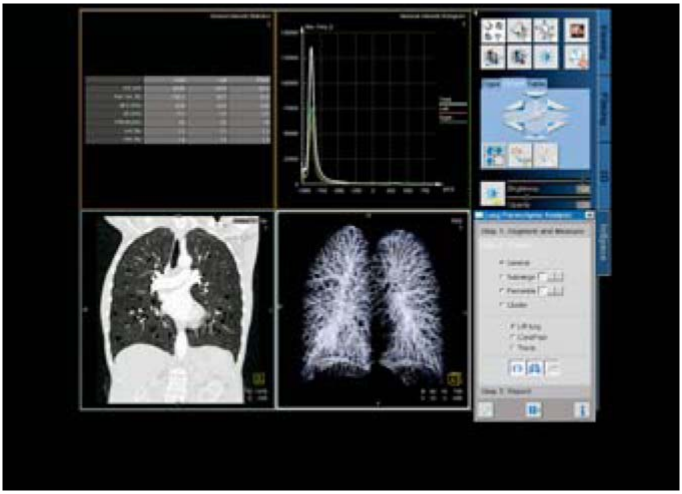

Fig. 3. Siemens software supports accurate diagnostic decision making regarding lung tumor characteristics

doctors to determine tissue density — a measurement that cannot be performed manually. Tissue density, in turn, provides an initial indication of whether or not a tumor is malignant.

Such measurements are also often used on patients with emphysema, a disease usually caused by smoking that destroys the alveoli in the lungs. Here, syngo InSpace4D Lung Parenchyma Analysis software from Siemens measures density distribution throughout the entire lung, whereby a diseased lung will, due to its burst alveoli, have more free air in its tissue (and will therefore be less dense) than a healthy lung. “This software solution makes it possible for the first time ever to quantify the early stages of emphysema and thus to effectively monitor treatment,” says Das. “This used to be an extremely difficult process requiring several indirect tests.”

3.4 Virtual Colonoscopy

Colon cancer screenings are another application where computeraided detection is very helpful. Dr. Anno Graser from the Institute for Clinical Radiology at Munich University Hospital uses syngo Colonography with PEV (Polyp Enhanced Viewing) software to review the results of virtual colonoscopies. Unlike physicians in Aachen, Graser does not have a second radiologist and therefore relies on PEV software for a second opinion.

“The program, which can be used by any doctor, delivers very good results, as long as the colon has been properly cleansed beforehand,” says Graser, who also tested the software in several studies. He’s not only satisfied with the program’s accuracy, but also happy that “the software simplifies and accelerates the entire process.” In Graser’s institute, it only takes four minutes, in fact, for the program to calculate the PEV results — about as long as it takes a gastroenterologist.

Graser has been screening one or two patients per day with the system since he concluded his clinical studies of the software. But there is still some resistance to the new technology. “Health insurance companies in Germany only pay for conventional colonoscopies, unless you have a situation where an intestinal infection or obstruction would not allow for such a procedure,” he explains.

Nevertheless, patients prefer the virtual procedure because it’s much shorter than the conventional one. Another major benefit is that its polyp detection software is extremely sensitive, thus improving the chances of early detection. “These benefits are going to help the system achieve a major breakthrough in terms of acceptance,” says Graser.