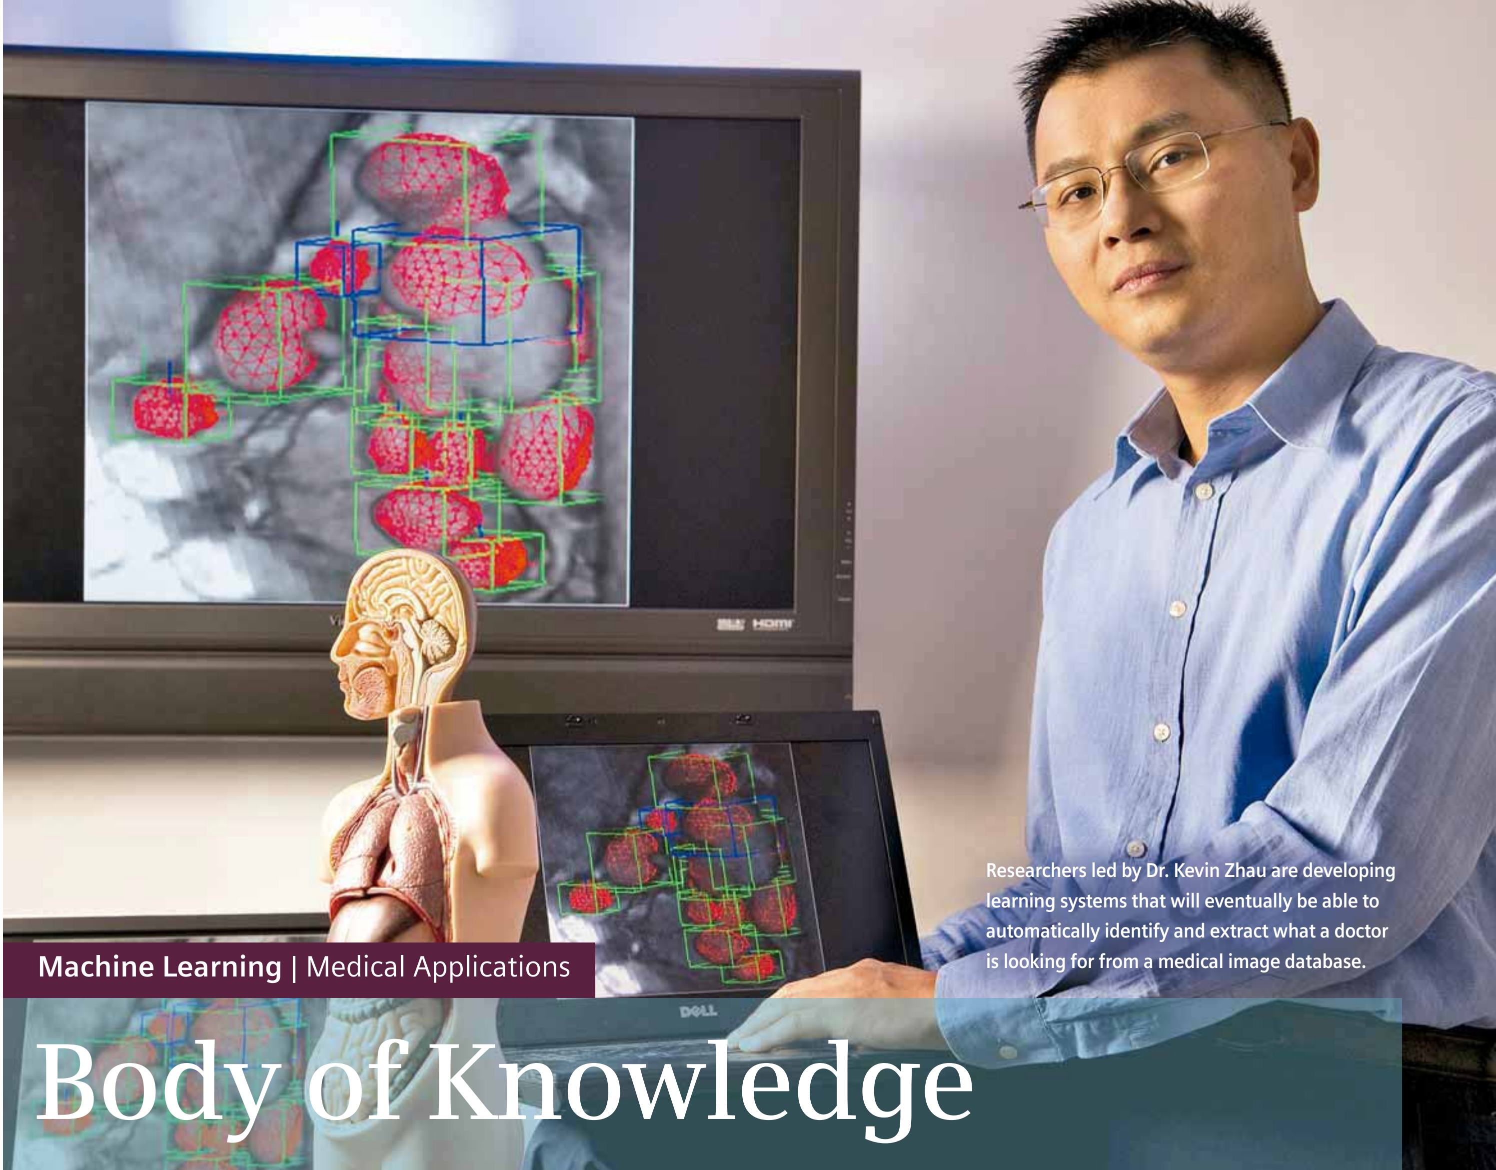

Researchers led by Dr. Kevin Zhau are developing learning systems that will eventually be able to automatically identify and extract what a doctor is looking for from a medical image database.

Machine Learning | Medical Applications

Body of Knowledge

Trained on thousands of annotated images, software systems have learned to identify organs and even recognize cancer stages in pathology slides. Such capabilities are opening the door to a world of diagnostics and treatment in which anatomy and physiology become semantically addressable.

Imagine if you could have a medical scan so complete that the location and function of every cell in your body would be stored. This would make it possible, for instance, to instantly visualize all cardiac cells or all prostate cells, thus presenting an unobstructed three-dimensional view of an organ from any desired angle, and allowing you to zoom in on any part of that organ — or any element of its function — down to almost any level of detail by simply moving a joystick or typing in a request. Although such a vision is perhaps 20 years from realization, scientists are already approaching some of these functions in limited areas of the body on the level of voxels — in other words, 3D pixels — each of which represents roughly 100,000 cells.

"The end result of our efforts should be the ability to automatically label every voxel in a scan," says Dr. S. Kevin Zhou, who heads a program that focuses on whole body image analytics at Siemens Corporate Technology (CT US) in Princeton, New Jersey. "From that, we will eventually be able to develop services such as semantic searching, which will make it possible for a doctor to simply mention, for instance, a liver tumor, and the system will pull up images of that tumor from a patient's most recent exams, measure its relative size in each image, and thus illustrate how it has responded over time to treatment. It will all add up to a faster, more accurate, and more efficient workflow."

Before an image analysis system can determine whether that liver it's looking for might be in a given image, however, it must first get its bearings. To do so, such systems start out by looking for anatomical landmarks. In the thorax, for instance, these include locations such as the top of the lung and the lower end of the aorta. "Landmarks keep an image analysis system from getting mixed up, and allow it to orient itself," explains Zhou.

How Machines Memorize. Behind the growing ability of machine learning systems to identify landmarks and zero in on objects of interest is the development of software that can learn to identify the content of an image based on vast numbers of "classifiers" or characteristics that are common to all examples of a target object.

Once trained on thousands of images of, say, the liver, each one of which has been annotated by experts, such software has essentially memorized the three-dimensional shape of a human liver and can therefore generalize

Thanks to machine learning, the ability to distinguish the precise outlines of organs and their constituent anatomies regardless of occlusions, angle of view, imaging modality, or pathology is being automated, thus opening the door to faster, more precise diagnostics.

to the extent that it can identify and segment (separate from its surroundings) a liver in any medical image, regardless of occlusions, angle of view, imaging modality, or pathology. And the same is true for a rapidly-growing number of anatomical entities throughout the body, from organs and bones, to the outlines of a fetus or a lesion.

Flat Ribs. Once a system has learned to automatically identify part of the anatomy, an amazing world of possibilities begins to unfold. Take, for instance, what happens after a routine whole-body computed tomography scan. Today, in many countries, radiologists are required by law — regardless of the reason for the scan — to examine all the major organs in the image set and the entire rib cage, including inside surfaces, to determine if there are any signs of disease. "Examining the ribs is a particularly time-consuming activity because it is difficult to navigate all those curved surfaces," says Zhou.

But software now being developed by Siemens Corporate Technology in cooperation with the company's Computed Radiology business unit could one day make it possible to automatically segment the rib cage from the rest of an image and flatten the ribs, thus substantially accelerating the process. "The program will use machine learning to find each rib and locate its center line. This, in turn, will make it possible to apply a simple program that would then flatten each rib," adds Zhou. "

Fusing X-Ray and Ultrasound Information. For over 50 years, patients with many of the most serious cardiac conditions have had to undergo the trauma of open heart surgery. But today, thanks to steadily-improving imaging techniques and the advent of systems that can learn to identify and automatically track everything from valves and chambers to catheters and stents, a growing number of patients can be treated using nothing more invasive than a specialized catheter. For instance, one year ago, as reported in *Pictures of the Future,* Fall 2010, page 79, Siemens introduced a new X-ray-based visualization and guidance technology to facilitate the implantation of a replacement aortic valve. Now, thanks to machine learning algorithms that automatically identify the same anatomical landmarks in different modalities, procedures such as aortic valve replacement are set to become more precise.

Machine learning is helping computed tomography systems to identify calcified tissues.

"We call this new technique 'model-based fusion,'" says Dr. Razvan Ionasec at Corporate Technology in Princeton, New Jersey. "Three-dimensional, X-ray-based angiography is great for seeing the location of a catheter, but not optimal for visualizing tissues. Ultrasound, on the other hand, is exactly the opposite. So the idea is to combine the two."

With this in mind, researchers led by Dr. Terrence Chen, also at Corporate Technology US in Princeton, are developing a learning-based detection and tracking technology that will help to automatically optimize the registration of angiography images with images produced by a miniature intravascular ultrasound (IVUS) device. Such devices are often used to determine the quantity of plaque in the coronary arteries. Here, the learning process focuses on automatically recognizing the ultrasound transducer and an associated guiding catheter in X-ray-generated angiography images as they move through blood vessels. "This helps to determine the exact location of a plaque deposit and thus supports treatment planning," says Chen.

Automatic Identification of Calcifications. Working along related lines, a team of researchers led by Corporate Technology Visual-and-Solid-Modeling Program Manager Dr. Tong Fang has developed a technology called dynamic tissue contrast enhancement (DTCE) that identifies human anatomy in ultrasound images and then "optimizes image quality using advanced noise reduction and structure enhancement technologies," according to Fang. Based on off-line training in which annotated sample images were used for learning, the software provided "superior image quality and clinical diagnosis benefit," in a pilot study, says Fang.

Researchers are also using machine learning technology to train computed tomography systems to identify calcified tissues in images of the heart. "Calcification is the main reason for aortic valve replacement and a key factor in coronary artery disease," explains Ionasec. "Computed tomography images already provide outstanding anatomical detail. But in the future, with software that is now in the pipeline, we expect to develop a system that will, for the first time, help clinicians quantify the extent of calcium deposits on an aortic valve and in the thoracic aorta. This information will help them to predict the chances of success for a replacement valve, decide which type of valve to use, and the amount of pressure to apply through a balloon when fixing the new valve in place."

Further down the road, researchers hope that machine learning will help them to detect the differences between normal plaque, which remains anchored to the surfaces it occupies, and so-called "unstable" plaque, which can break away from the surface and potentially cause a heart attack or stroke — a major risk factor in many interventional treatments. "We can see different kinds of plaque in computed tomography and magnetic resonance scans," says Dr. Gareth Funka-Lea, a specialist in cardiovascular diseases at Corporate Technology US in Princeton, "but we still do not know how to differentiate unstable plaque. It is possible, however, that we will eventually find the answer by harnessing machine learning and massive data mining."

Semantic Heart. Siemens' teams of cardiovascular and machine learning experts have also expanded their focus from the aortic valve to cover virtually the entire human heart. "As part of a major Siemens R&D project called the 'Semantic Heart,' we are now using machine learning to automatically identify all four valves and are integrating this information with our models of the cardiac chambers to produce a full model of the heart," says Ionasec. The idea is that, eventually, clinicians will be able to model and compare the effects of different forms of cardiac intervention — from insertion of a stent or repair of an aneurysm to replacement or repair of a valve — on the dynamics of a patient's entire heart before the treatment is performed.

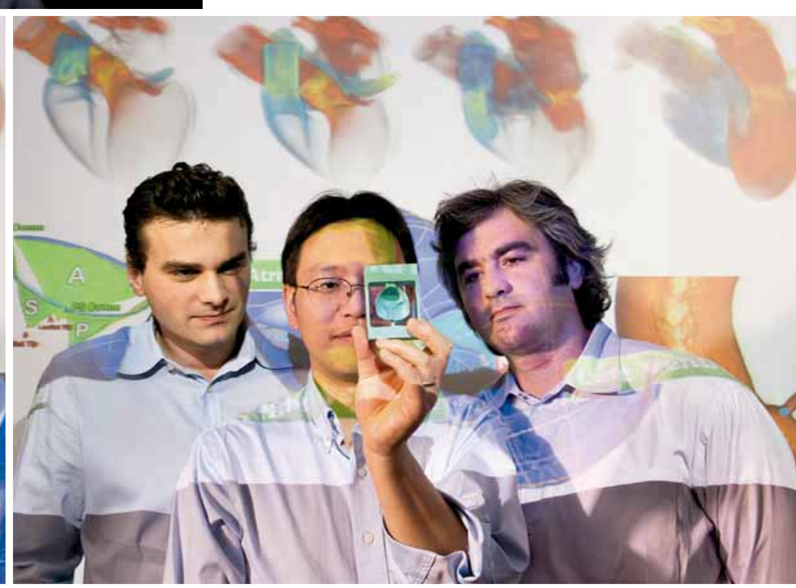

Using machine learning, researchers involved in Siemens' Semantic Heart project are setting the stage for producing a fully-functional model of each patient's heart...

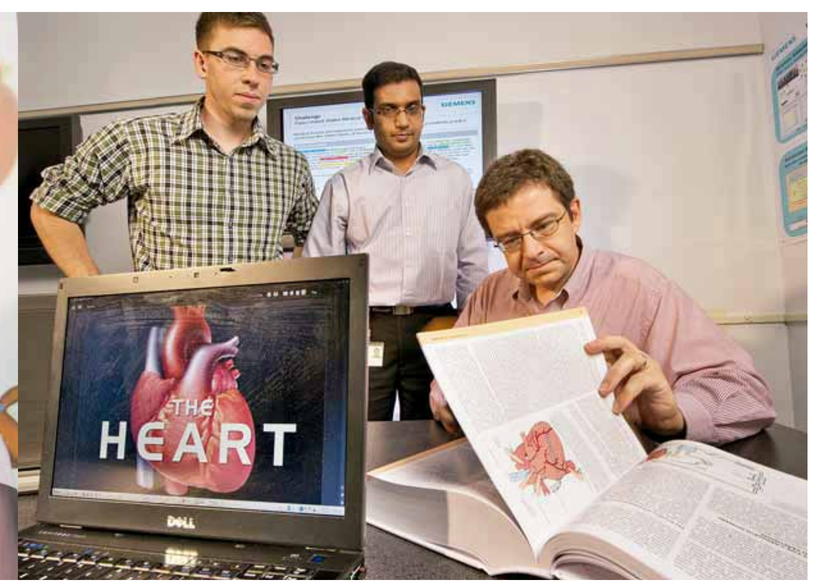

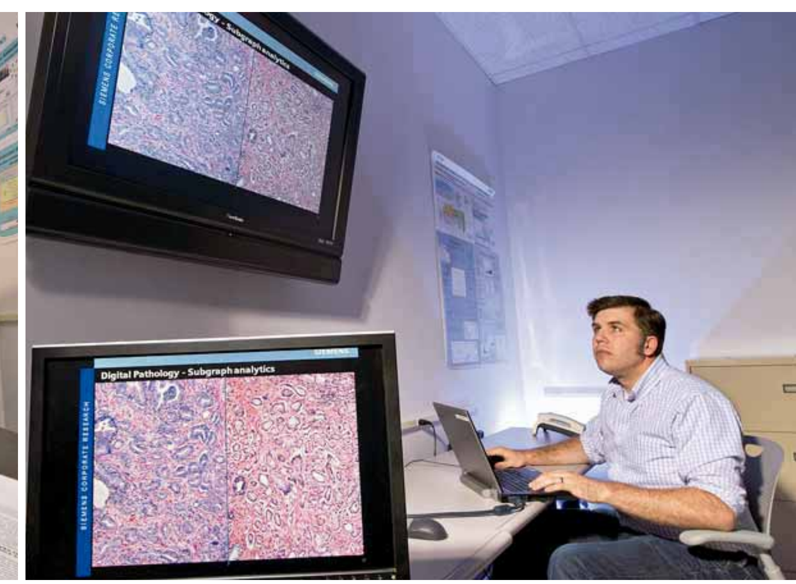

...While others develop a medical reasoning engine (see insert page 60) to support physicians. Learning systems identify cancer stages in prostate biopsies (right).

One of the most far-reaching results of the Semantic Heart project is the rapidly-evolving ability to model the mitral valve, which controls blood flow from the left atrium to the left ventricle. Far more complex than the aortic valve, the mitral valve is kept in check by a network of string-like tendons that keep its two flaps from reversing direction into the left atrium. But the tendons can snap in response to overexertion or disease — with consequences that can range from minor to life-threatening. The condition can be repaired by means of a trans-catheter procedure that involves clipping the flap with the snapped tendon to the remaining healthy flap. "But attaching a minuscule clip to two moving flaps by means of a catheter using only fluoroscopy to see what you're doing is not easy," says Ionasec.

In view of this challenge, Ionasec's research team is therefore developing an approach that combines the sub-millimeter resolution from pre-operative ultrasound images generated by a transducer in the esophagus, with intra-operative X-ray images acquired with syngo DynaCT Cardiac on a Siemens Artis zee angiography system. The approach, which is based on algorithms trained on thousands of patient images, uses machine learning to automatically recognize and track the anatomy and movements of the flaps, as well as to fuse the X-ray and ultrasound images. The new procedure is expected to enter clinical trials in Germany late in 2011.

Reading the Hidden Language of Cells

One day, there will be a device called a digital diagnostic pathology scanner. It will process thousands of pathology slides per hour, each loaded with a paper-thin slice of tissue suspected of harboring disease, and will deliver highly-accurate analyses at minimal cost. Its output will be combined with each patient's results from other areas, such as genetics, physiology, anatomy and demographics. And of course it will learn from each and every slide, thus constantly refining the accuracy of its results. In fact, such machines will probably be networked, allowing them to learn from each other.

Although such a machine may seem to be a distant vision, researchers are today compiling the basic knowledge that will eventually drive such a device. In Princeton, New Jersey, for instance, a team of researchers led by Leo Grady, PhD, a specialist in biomedical image analytics at Siemens Corporate Technology, is using machine learning to predict the cancer stage of samples from prostate biopsies.

Using slides previously marked by expert pathologists as belonging to one of the four cancer stages, "the system tries to identify features such as cell structure and arrangement that are consistently associated with a stage," explains Grady. "For each 100 graded slides, the system is trained on ninety, and tested on the remaining ten. Then — always randomly — another 90 slides are selected for training and another ten for testing."

The process is repeated until performance is good across the board, at which point the system has learned to generalize from experience — sometimes with very surprising results. For instance, not only has the system, as expected, learned to identify what different kinds of cells look like — thus opening the door to automated counting — but it has discovered something the researchers were not even aware of.

"The system extracted the fact that although there are loop-shaped patterns of cancer cells and normal cells in each of the images, the length of the loop and the number of cells in it is sufficient to predict the cancer stage," says Grady. "That was a surprise to us. But when we discussed this with a pathologist he said that yes, this structure is something specialists look for to determine cancer stage. In this case, however, the system discovered this on its own."

Arthur F. Pease

A Reasoning Engine for Tomorrow's Physicians

Envisioning a system that will one day support physicians in answering complex medical questions, researchers at Siemens Corporate Technology in Princeton are developing a deep reasoning machine that learns from large quantities of data. The simplified example shown below illustrates four steps in the deep reasoning process:
(1) acquire patient history and physical examination data,
(2) determine differential diagnoses,
(3) recommend diagnostic tests to cover existing knowledge gaps: e.g. perform ECG to detect ST-Elevation (i.e. occlusion of a coronary artery) and Q-Waves (i.e. local electrical dysfunction of heart muscle cells), in order to
(4) select the most likely diagnoses.
"The system," explains Project Leader Mathaeus Dejori, PhD, "reflects the decision-making process in medical practice. Physicians typically receive lists of patient values and are expected to make hard decisions." Adds Vinay Shet, PhD, who is also involved in the project: "Our system avoids the complexity of dealing with language directly. Instead, it works from semantic concepts such as 'coronary occlusion' and 'acute chest pain'. The deep reasoning machine has an understanding of these concepts and uses medical knowledge to draw conclusions." Dejori, Shet, and co-researcher Dan Tecuci, PhD, envision the technology as an intelligent assistant that will help doctors easily make use of rapidly-growing reservoirs of digital information.