How Would a Computer Diagnose Arthritis on a Radiograph?

● KENNETH URISH, MD, PHD, FAAOS, AND ALAN M. REZNIK, MD, MBA, FAAOS

The first article in this series ("Understanding the Impact of Artificial Intelligence on Orthopaedic Surgery," *AAOS Now*, September 2018) provided an overview of the history of artificial intelligence (AI) and a few general concepts regarding how AI problems are approached. Inevitably, the question becomes: Is it possible to replace doctors with AI machines? As an example, this article discusses how, using basic AI techniques, an AI machine can diagnose osteoarthritis (OA) and its severity on radiographs.

For a human, diagnosing OA requires just three features absent deformity: osteophytes, joint space narrowing, and sclerotic bone. A computer may detect OA by different textural features. It may look for loss of signal variation around the subchondral bone that represents loss of joint space or sclerotic bone, and it must do it by brute force of mathematical calculations. Understanding the difference in the approach is key to seeing how AI works when looking at radiographs. Therefore, diagnosing OA on radiographs is a multistep problem involving computer vision.

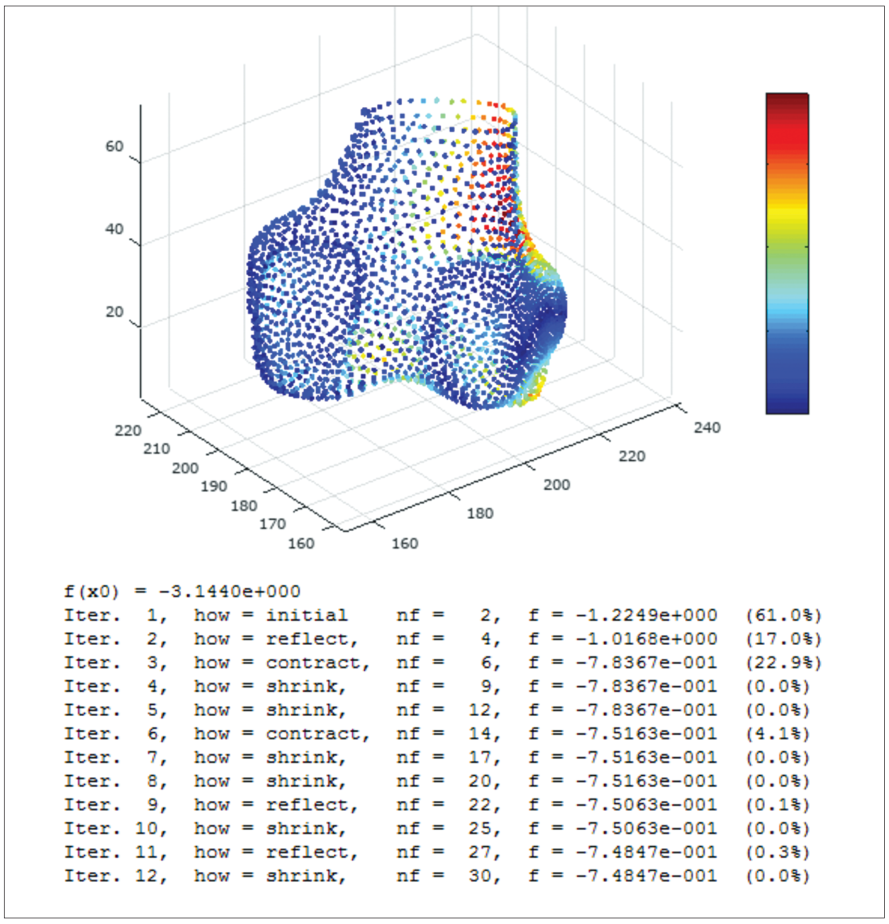

Fig. 1 An anterior-posterior radiograph of the knee segmented by a computer vision algorithm
COURTESY OF ENHATCH

Using AI to diagnose OA

The goal is for the machine to appreciate the structures on the image, extract features for analysis, and develop pattern recognition to make the diagnosis. There are many variations and other more complex approaches that can be used to eliminate the need for computer vision, including automated feature extraction; but for illustrative purposes, this article will simplify the problem by using this initial approach.

Still, even this is far easier said than done. The human mind is exceptional at recognizing global patterns and being able to see an entire image as a single object. This is a challenge for computers. The foundation of computer vision is based on approaching images pixel by pixel. Groups of pixels are known as neighborhoods. Humans are excellent at seeing the entire image from a global perspective, whereas computers must build an image pixel by pixel.

> **Editor's note:** This is the second in a series of articles about artificial intelligence and how it may impact orthopaedic surgeons' treatment plans, protocols, algorithms, and "IQ." The first article was featured in the September issue of *AAOS Now*; visit www.aaos.org to read more.

AI step-by-step analysis of a radiograph

The first step in diagnosing OA on radiographs is computer vision—recognizing the location of the important anatomic structures of cartilage and bone. This is also known as image segmentation, where anatomic structures can then be annotated. An easy initial approach is in determining the location of bone and assuming cartilage is in the space between the bone. This is a very popular approach, as there is high contrast between bone and soft tissue, making it easier to recognize the structures. Essentially, the edge of the interface between bone and soft tissue needs to be recognized (Fig. 1).

The second step is extracting features from the image to help make a diagnosis. For humans, feature extraction involves recognizing osteophytes, joint space narrowing, and areas of sclerotic bone around the joint. We look for a minimal number of features to recognize the pattern and make the diagnosis. In the AI algorithm, it is the exact opposite—the data are collected pixel by pixel and turned into a large number (hundreds or thousands) of apparent abstract and meaningless features. This can be thought of as looking at the texture of an image pixel by pixel. It is important to appreciate that a radiograph is really an image with thousands of pixels. In an infinite number of neighborhoods around each pixel, each neighborhood will have a mean value, a variation of the surrounding pixels, and other mathematical measurements that can define the texture (Fig. 2).

The final step in a basic AI routine becomes recognition and classification of the pattern. Most AI algorithms recognize a pattern by separating the populations by a mathematical expression. Machine learning involves feature reduction, where the thousands of collected features are reduced to only the most important ones, and then a mathematical expression is used to separate the two populations of healthy and OA radiographs.

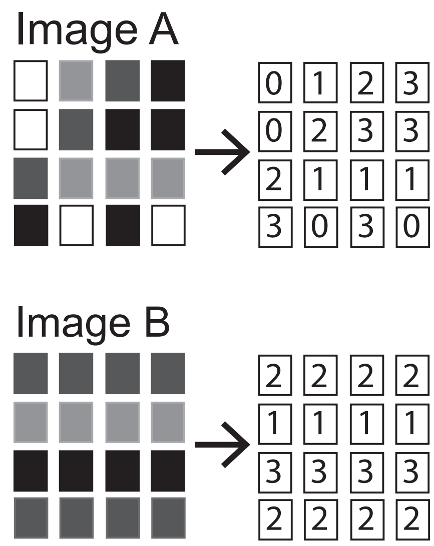

Fig. 2 An image is really a neighborhood of pixels where each pixel intensity is represented by a numeric value. In image A, there is a heterogeneous texture with increased variation in each pixel row and a lower mean compared to image B.
COURTESY OF ENHATCH

AI computer models and 'fit'

"Fitting" the data with a good model is key to the best results. Here is where the concept of over-fitting and under-fitting comes into play. In mathematics, if we fit a complex relationship with a straight line, we may not have a good model. This is an example of under-fitting. At the same time, if we have 500 data points and fit the relationship with a 500-variable equation, the model does not predict anything but the data we have now. This is an example of over-fitting. To avoid this problem in AI, unnecessary features that are not needed for a final pattern are discarded. In other words, if we use all of the thousands of features in an image for the model, the dataset would be over-fit. The key becomes recognizing the most important features. The process of using only the more important features and discarding the rest is known as feature reduction. It

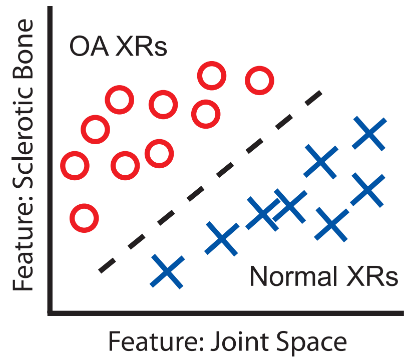

Fig. 3 In our simplified model, two features can be used to separate normal and osteoarthritis radiographs. Two features selected by the computer that roughly correspond to sclerotic bone and joint space separate the two populations, allowing the computer to mathematically define the pattern.
COURTESY OF ENHATCH

has become clear that many AI problems would become intractable without feature reduction.

In our case, the model needs to create a good fit for bone sclerosis and cartilage loss but not tumors or stress fractures. If the AI model was built with two features—pixel variation around articular bone (sclerosis) and the area of mean pixel intensity in the joint space (joint space loss)—then a linear model could be used to separate the two populations (Fig. 3). A model is then built from a training set that defines the two different populations. Deep learning involves more complex, higher-dimensional situations where a nonlinear model, such as a neural network, is needed. However, the principle of separating these two classes is the same (Fig. 2).

You can write artificial intelligence (AI) code to detect patterns and diagnose osteoarthritis via radiographs, but you cannot write an AI program to help a patient decide or sign an informed consent.

How is this applied to radiographs?
In application, when a radiograph is presented to this AI algorithm, the computer will segment the image for the anatomic location of the bone, extract the relevant features related to sclerotic bone and joint space loss, and then compare the features to its known arthritis model. This would allow the computer to determine whether the patient has or does not have arthritis. Other features could be added to the model from the medical record, like Western Ontario and McMaster Universities Osteoarthritis Index scores, to determine the presence of symptomatic arthritis.

Conversations about AI in medicine invariably lead to concerns about physicians being replaced. AI is ideal for repetitive tasks that do not require human interaction and creativity, such as determining whether a radiograph demonstrates OA. Orthopaedic surgery is more than pattern recognition. The value of orthopaedic surgeons is not that they can recognize OA or perform different procedures such as steroid injections or total knee replacements. The value of orthopaedic surgeons is appreciating many alternatives; adjusting plans; blending in the patient's psychosocial factors, quality of life, and physical demands; and sharing decision-making and treatments. You can write AI code to detect patterns and diagnose OA via radiographs, but you cannot write an AI program to help a patient decide or sign an informed consent.

Future articles in this series will explore natural language interpretation for chart reading and diagnosis and how, for humans, there may be a requirement for AI to explain what it is doing—so-called explainable AI or XAI.

Kenneth Urish, MD, PhD, FAAOS, specializes in arthroplasty and serves on the AAOS Patient Safety Committee. He is an assistant professor in the Department of Orthopaedic Surgery and associate medical director at the Magee Bone and Joint Center at the University of Pittsburgh. He has an adjunct appointment in the Department of Biomedical Engineering at Carnegie Mellon University.

Alan M. Reznik, MD, MBA, FAAOS, specializes in sports medicine and arthroscopic surgery and serves on the *AAOS Now* Editorial Board, AAOS Communications Cabinet, and Committee on Research and Quality. Dr. Reznik is chief medical officer of Connecticut Orthopaedic Specialists, associate professor of orthopaedics at Yale University School of Medicine, and a consultant.