

CURRENT CONCEPTS REVIEW

Artificial Intelligence and Orthopaedics

An Introduction for Clinicians

Thomas G. Myers, MD, MPT, Prem N. Ramkumar, MD, MBA, Benjamin F. Ricciardi, MD, Kenneth L. Urish, MD, PhD,
Jens Kipper, PhD, and Constantinos Ketonis, MD, PhD

Investigation performed at the University of Rochester Medical Center, Rochester, New York

- ▶ Artificial intelligence (AI) provides machines with the ability to perform tasks using algorithms governed by pattern recognition and self-correction on large amounts of data to narrow options in order to avoid errors.
- ▶ The 4 things necessary for AI in medicine include big data sets, powerful computers, cloud computing, and open source algorithmic development.
- ▶ The use of AI in health care continues to expand, and its impact on orthopaedic surgery can already be found in diverse areas such as image recognition, risk prediction, patient-specific payment models, and clinical decision-making.
- ▶ Just as the business of medicine was once considered outside the domain of the orthopaedic surgeon, emerging technologies such as AI warrant ownership, leverage, and application by the orthopaedic surgeon to improve the care that we provide to the patients we serve.
- ▶ AI could provide solutions to factors contributing to physician burnout and medical mistakes. However, challenges regarding the ethical deployment, regulation, and the clinical superiority of AI over traditional statistics and decision-making remain to be resolved.

Artificial intelligence (AI) is a term coined by John McCarthy as a theory that computers could eventually learn to perform tasks through pattern recognition with minimal to no human involvement^{1,2}. A more modern and accurate definition of AI is the application of algorithms that provide machines the ability to solve problems that traditionally required human intelligence³. In 1976, Jerrold S. Maxmen penned that AI would bring about the “post-physician era” in the 21st century^{4,5}. With the rapid increase in affordable computational power and exponential increases in extremely large data sets (“Big Data”), AI has transitioned from mere theory to tangible application on an unprecedented scale¹. By analyzing extraordinarily large data sets in near real-time, AI has become fundamentally ingrained within many facets of

society from autonomous driving cars and video streaming recommendations (Netflix), to online purchase recommendations (Amazon), advertisements (Facebook), and fraud detection (Capital One).

AI can be thought of as an umbrella term that encompasses a broad range of subfields, including *machine learning* (ML), which in turn contains a subfield called *deep learning* (DL) (Fig. 1)^{6,7}. The goal of AI is to allow a machine (computer) to perform specific tasks that can match or exceed human performance. Advances in computing power, data storage, and availability of high-quality data have driven the expansion of AI into the health-care field. Problems specific to orthopaedics, such as image recognition, preoperative risk assessment, clinical decision-making, and analysis of massive

Disclosure: One author (K.L.U.) reported grant support from the National Institute of Arthritis and Musculoskeletal and Skin Diseases (NIAMS) of the National Institutes of Health (NIH), and from the Orthopaedic Research and Education Foundation (OREF). On the **Disclosure of Potential Conflicts of Interest** forms, which are provided with the online version of the article, one or more of the authors checked “yes” to indicate that the author had a relevant financial relationship in the biomedical arena outside the submitted work and “yes” to indicate that the author had a patent and/or copyright, planned, pending, or issued, directly relevant to this work (<http://links.lww.com/JBJS/F719>).




Fig. 1
The relationship of AI, ML, and DL. (Reproduced, with modification, from: Chollet F. Deep learning with Python. Shelter Island, NY: Manning Publications; 2018. Reproduced with permission.)

data sets, are beginning to be addressed using AI-based methods (Table I). The purpose of the present review is to provide background on AI, relevant applications to health care (with a focus on orthopaedics), and challenges with its use.

What Is AI?

AI has promised to solve a multitude of problems plaguing society and health care today. The term AI can be difficult to define. However, understanding what AI is becomes crucial as the hype of AI is frequently misapplied to many new technologies or products that may not truly fit the definition⁸. Appendix 1 contains a list of AI terms introduced in this paper that may be useful for future reference. AI at its core involves machines that can perform tasks innately characteristic of human intelligence. This includes tasks like planning, understanding language, recognizing patterns, learning, and problem solving. Most importantly, AI can learn from mistakes and improve, which resembles experiential learning.

ML

ML was introduced shortly after AI in 1959 at IBM as a method of achieving AI^{9,10}. ML algorithms are able to learn from examples by adjusting their internal parameters (weights) and strengthening relevant associations to improve the accuracy of a given model^{6,9,11}. The “learning” that takes place in ML is achieved by incremental optimization of a mathematical model¹². These models are built using inputted

features (data), which are divided into 2 separate groups: a training set, which is used to build the mathematical model, and a testing set, which is used to assess the model’s performance. Mathematical functions are used to match inputs with desired outputs, which is distinctly different from the process involved with statistics or traditional programming (Fig. 2)^{9,10}.

It is possible to produce an AI system (i.e., computer checkers and chess programs) without using ML, but this would require extensive coding, complex rules, and decision trees to account for every single possible move in the game⁹. At the risk of oversimplification, the biggest difference between statistics and AI is the kind of question that is asked. In statistics, the question is “What is the probability that some phenomena that I observed occurred just by chance?” (hypothesis testing). In AI, the question is “When I see a new entity with a specified set of features, how should I label it?” (prediction).

The ML training process can either be supervised or unsupervised¹³. Supervised ML occurs when the training data are labeled by humans. For example, if an ML algorithm is tasked with detecting arthritis on a radiograph, the arthritic features in addition to relevant clinical information (i.e., joint space loss, subchondral sclerosis, recent nonsteroidal anti-inflammatory drug [NSAID] use, patient-reported outcomes, etc.) must first be manually identified and labeled by a human along with the label of whether the radiograph is an example of an arthritic or a normal knee. An algorithm may require hundreds or thousands of correctly labeled example radiographs, with and without these features, to accurately recognize patterns (“arthritis present” or “arthritis absent”), achieve enough accuracy, and therefore learn (Fig. 3)^{14,15}. With supervised ML, inaccurate labeling of data can lead to bias and error.

In contrast, with unsupervised ML, the training data are not labeled, the data are not known, but the outcome of interest is known. To continue with our previous example, the presence or absence of arthritis on the radiograph would be known, but the training data are not labeled by humans. The goal in these cases is pattern recognition as best determined by the computer. The algorithms look for clustering or self-organization of the given data in ways that can lead to the so-called black box phenomenon. The black box phenomenon occurs when the algorithm detects potentially unconventional patterns not recognized by human logic. Therefore, any attempts at retrospective performance evaluation of the model are difficult for humans to understand.

The amount of data required in either supervised or unsupervised algorithm training has been well illustrated by Beam and Kohane (Fig. 4)¹⁶. Figure 4 presents the relationship between the degree of human relative to ML involvement and the magnitude of the data sets required to sufficiently train existing ML algorithmic examples. The amount of data required to successfully run unsupervised ML (most DL examples in Figure 4) is magnitudes larger than for supervised

TABLE 1 Notable Examples of Leveraging AI for Orthopaedic Applications

Study	Application*	Notes
Thong et al. ⁴⁷ (2016)	Optimization of 3-D spine model vectors for the automatic detection of adolescent idiopathic scoliosis.	
Olczak et al. ⁴⁸ (2017)	Identification of fractures from radiographic images	
Chen et al. ¹⁷ (2017)	ML-based predictions for physician order entry show that prioritizing small amounts of recent data is more effective than using larger amounts of older data toward future clinical predictions	Concept of decaying clinical data based on practice patterns is critical to appreciate
Kruse et al. ⁴⁹ (2017)	Prediction of hip fractures from dual x-ray absorptiometry	
Cilla et al. ⁵⁰ (2017)	Use of ML to optimize short-stem THA design to produce optimal mechanical performance	
Konda ⁵¹ (2018)	An AI system (PersonaCARE) helps manage NYU's middle age and geriatric fracture population based on all of the principles of value-based care	
Karnuta et al. ⁵² (2019)	Determined that a bundled care model for hip fractures is an unsustainable value-based model	
Ramkumar et al. ¹⁹ (2019)	Predicted length of stay, inpatient costs, and patient disposition for lower extremity joint replacement	First introduction of the patient-specific payment model
Shah et al. ⁵³ (2019)	Automatic measurement and segmentation of articular cartilage thickness in healthy knees on MRI	
Harris et al. ⁵⁴ (2019)	Prediction of 30-day complications and mortality following TJA	Utilized the National Surgical Quality Improvement Program database
Greenstein et al. ⁵⁵ (2019)	ANN utilization of in-house EMR data to predict skilled nursing facility utilization following TJA	
Fontana et al. ⁵⁶ (2019)	ML to preoperatively predict with fair to good accuracy which patients may achieve minimally clinically important differences postoperatively in TJA	Study was the subject of <i>Clinical Orthopaedics and Related Research's</i> "Editor's Spotlight" ⁵⁷
Thirukumaran et al. ⁵⁸ (2019)	Use of NLP to identify orthopaedic surgical site infections	
Galbusera et al. ¹³ (2019)	Valuable review of AI and ML in spine research	Recommended reading

*3-D = 3-dimensional, ML = machine learning, THA = total hip arthroplasty, NYU = New York University, MRI = magnetic resonance imaging, TJA = total joint arthroplasty, NLP = natural language processing, ANN = artificial neural networks, and EMR = electronic medical records.

ML ("expert AI systems" and "risk calculators"). However, unsupervised ML is not synonymous with "unmanned" ML. ML algorithms can be run unsupervised but not without frequent validation of their predictive accuracy¹⁷. In many clinical scenarios, the predictive accuracy suffers because of changes in practice patterns, patient demographics, or other aspects of the training data.

DL

DL is a more sophisticated form of ML. DL is capable of unsupervised learning from unstructured and unlabeled inputs and filtering out data input from variables of low relevance to the prediction of interest. DL is modeled after the brain's neuronal connections via algorithms termed *artificial neural networks* (ANNs). The individual nodes in




Fig. 2

Illustration of the process involved with statistics or traditional programming (**Fig. 2-A**) and ML (**Fig. 2-B**). (Reproduced, with modification, from: Chollet F. Deep learning with Python. Shelter Island, NY: Manning Publications; 2018. Reproduced with permission.)

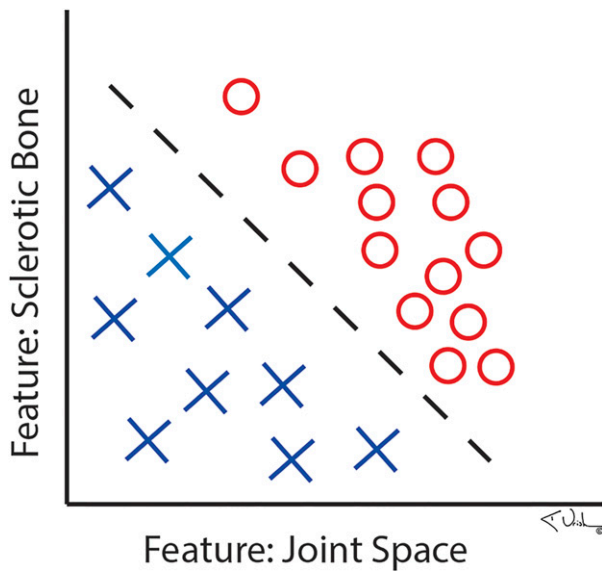
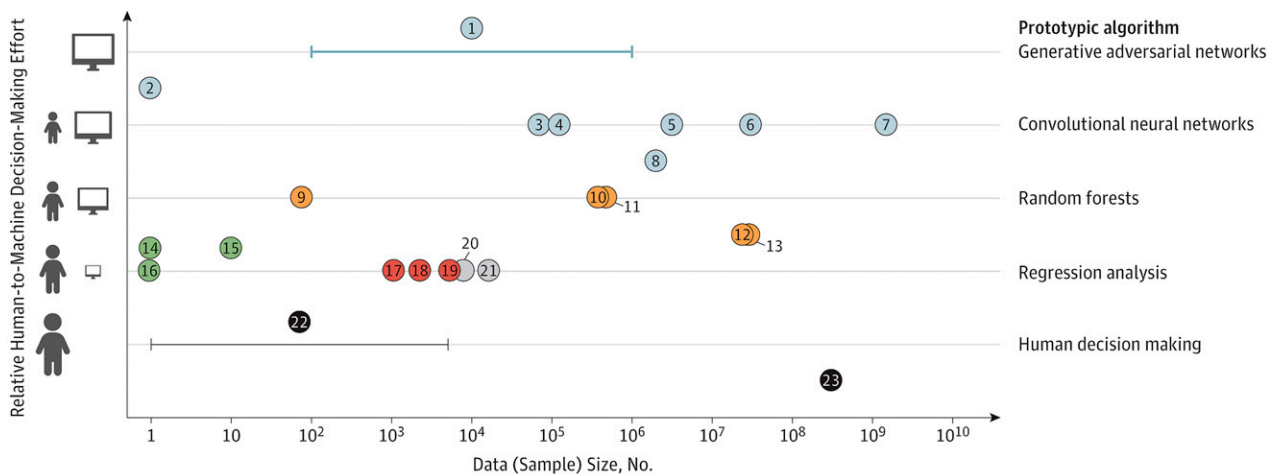


Fig. 3
Illustration showing how a computer would diagnose arthritis. (Reproduced, with permission, from: Kenneth Urish, MD, PhD.)

ANNs have discrete layers and connections to other nodes. The complex layering gives the algorithm its name and ability to learn more complex and subtle patterns than simple 1 or 2-layer neural networks (Fig. 5)^{3,18,19}. In the osteoarthritis (OA) example used previously, 1 node may be responsible for edge contrast, another for pixilation density, another for width, and so on²⁰. However, only the nodes that are clinically relevant to subchondral sclerosis, cysts, joint space narrowing, or other yet unknown features of OA will clinch the radiographic diagnosis of OA. The clinical utility of an ML OA algorithm may be most useful for the nonorthopaedic physician, orthopaedic trainee, physician assistant, or radiologist as a teaching, screening, or referral tool. DL has been applied successfully to complex problems including facial recognition, handwriting recognition, and vision systems for autonomous vehicles²⁰.

Natural Language Processing (NLP)

NLP is the study of how computers understand and interpret human language with the goal of generating structured information from unstructured free text. The ability to establish



Deep learning	Classic machine learning	Risk calculators
① Generative adversarial networks (2014)	⑨ Diffuse large B-cell lymphoma outcome prediction by gene-expression profiling (2002)	⑰ CHA ₂ -DS ₂ -VASc Score for atrial fibrillation stroke risk (2017)
② Google AlphaGo Zero (2017)	⑩ EHR-based CV risk prediction (2017)	⑱ MELD end-stage liver disease risk score (2001)
③ ATM check readers (1998)	⑪ Netflix Prize winner (2006)	⑲ Framingham CV risk score (1998)
④ Google diabetic retinopathy (2016)	⑫ Google Search (1998)	
⑤ ImageNet computer vision models (2012-2017)	⑬ Amazon product recommendation (2003)	
⑥ Google AlphaGo (2015)		Randomized Clinical Trials
⑦ Facebook Photo Tagger (2015)		⑳ Celecoxib vs nonsteroidal anti-inflammatory drugs for osteoarthritis and rheumatoid arthritis (2002)
⑧ Prediction of 1-y all-cause mortality (2017)	Expert AI systems	㉑ Use of estrogen plus progestin in healthy postmenopausal women (2002)
	⑭ MYCIN (1975)	Other
	⑮ CASNET (1982)	㉒ Clinical wisdom
	⑯ DXplain (1986)	㉓ Mortality rate estimates from US Census (2010)

Fig. 4
Chart showing the relationship between the degree of human relative to ML involvement and the magnitude of the data sets required to sufficiently train existing ML algorithmic examples. ATM = automated teller machine, EHR = electronic health records, CV = cardiovascular, and MELD = model for end-stage liver disease. (Reproduced, with permission, from: Beam AL, Kohane IS. Big data and machine learning in health care. JAMA. 2018 Apr 3;319[13]:1317-18. Copyright © 2018 American Medical Association. All rights reserved.)




Fig. 5
Example of input, hidden, and output layers in an ANN used to predict value-based metrics prior to elective primary total hip or knee arthroplasty. APR = all patient refined. (Reproduced, with permission, from: Ramkumar PN, Karnuta JM, Navarro SM, Haeberle HS, Iorio R, Mont MA, Patterson BM, Krebs VE. Preoperative prediction of value metrics and a patient-specific payment model for primary total hip arthroplasty: development and validation of a deep learning model. *J Arthroplasty*. 2019 Oct;34[10]:2228-34.e1.)

order from the massive amounts of patient data contained within the written text of the electronic medical record (EMR), emails, and the spoken word is an application of huge clinical value and of active research. This could take the form of assisting clinic schedulers with the triage and correct assignment of patients to practitioners, as well as determining if nonoperative treatment has failed. This could also be evident in digital scribes capable of simultaneously creating real-time office notes while interacting with the EMR, assistance in patient education in preoperative and postoperative surgical discussions, and so on. Furthermore, the ability to link in real time any individual patient encounter with the exponentially expanding volume of medical knowledge and personal health data is both a considerable challenge and huge opportunity²¹.

Continuing our knee OA example from above, we now apply it in the clinical setting of a new patient encounter for knee OA. Assuming that nonoperative options have failed, it may be possible to autopopulate a note from prior referrals to template the clinical visit and immediately provide the next steps for total knee arthroplasty preoperative scheduling and medical optimization. For starters, the algorithm would need access to the EMR to analyze available imaging consistent with

radiographic knee OA as well as primary care notes documenting pain, loss of function, and attempted nonoperative treatments. Of course, most NLP algorithms have direct access to medical dictionaries and associated acronyms that would be able to understand what “NSAIDs” are and include slight typographic errors. Recognizing “therapy” as synonymous with “PT” or “rehab” would be beneficial in patient records to indicate that the patient had tried NSAIDs and physical therapy and demonstrates obvious signs of radiographic OA. Certainly, the history and examination would be documented by the physician, but the NLP may be able to “suggest” phrases that would meet appropriate coding standards. The machine may use dozens of algorithms to break a sentence down or expand on one.

The EMR example is overly simplistic but certainly within reach. Most importantly, it begins to provide a glimpse into the potential of decompressing the orthopaedic surgeon’s documentation and regaining critical clinical time. Imagine a world in which a practitioner could walk into a room and engage a patient in conversation using an NLP-assisted EMR that is structured and responsive to tasks and commands such as “pull up all CT [computed tomography] scans on patient’s shoulder,” “graph the trend of all




Fig. 6

Remote patient monitoring. Data for health monitoring applications can be captured using a wide array of pervasive sensors that are worn on the body, implanted, or captured through ambient sensors, e.g., inertial motion sensors, electrocardiogram patches, smart watches, electroencephalograms, and prosthetics. (Republished with permission of IEEE, from: Ravi D, Wong C, Deligianni F, Berthelot M, Andreu-Perez J, Lo B, Yang GZ. Deep learning for health informatics. *IEEE J Biomed Health Inform.* 2017 Jan;21[1]:4-21; permission conveyed through Copyright Clearance Center, Inc.)

CRPs [C-reactive proteins] since explant of infected TKA [total knee arthroplasty],” or “book for left carpal tunnel release.” The early stages of a virtual assistant such as this is currently being studied at Vanderbilt University and the University of Rochester^{22,23}. NLP could assist in creating an accurate clinic note and generating a prognosis for treatment success that is based on the entirety of the patient’s biopsychosocial profile and access to the patient’s medical record. In this manner, we will come much closer to the concept of *precision medicine*²¹.

The Internet of Things (IoT)

The IoT is simply the interconnection of so-called smart devices that collect data and deliver them to the internet to connect people to people, people to things, and things to things²⁴. Examples relevant to orthopaedics include computers, intraoperative devices, smart phones, garments, sleeves, and braces²⁵⁻²⁷. These smart devices, when applied in the health-care environment, are termed *mobile health*, or *mHealth*^{28,29}. AI and mHealth together make an indispensable future in orthopaedics. McClelland points out that AI and the IoT are analogous to the relationship of the brain to the human body⁹. The body’s peripheral end organs collect a multitude of sensory data for

interpretation by the brain just as the IoT collects data that are sent to AI systems in order to interpret clinical scenarios (Fig. 6)³⁰. A current validated example of AI and the IoT includes wearable sensors paired with a smart phone that transfer information following TKA to a *remote patient monitoring* platform^{26,31}. These data can include patient-related outcome measures, steps, opioid use, range of motion, home exercise compliance, and pain visual analog score. Similarly, some of the same authors validated the use of a smart orthotic to remotely monitor shoulder range of motion^{26,32}. The remote monitoring of the orthopaedic patient opens up an exciting and entirely new area of research and patient engagement in the preoperative and postoperative phases of care. Furthermore, this technology allows the collection of large volumes (approximately 18,000 data points per patient) of real-time rehabilitation data following TKA at day 0 through week 12, which has never before been possible^{31,33}. Access to such data may enable predictions of which patients will struggle with their rehabilitation program because of intrinsic or extrinsic factors. Similarly, it can help determine when patients can safely progress to the next level in their rehabilitation program or if they have maximized their rehabilitation potential and can discontinue physical therapy.

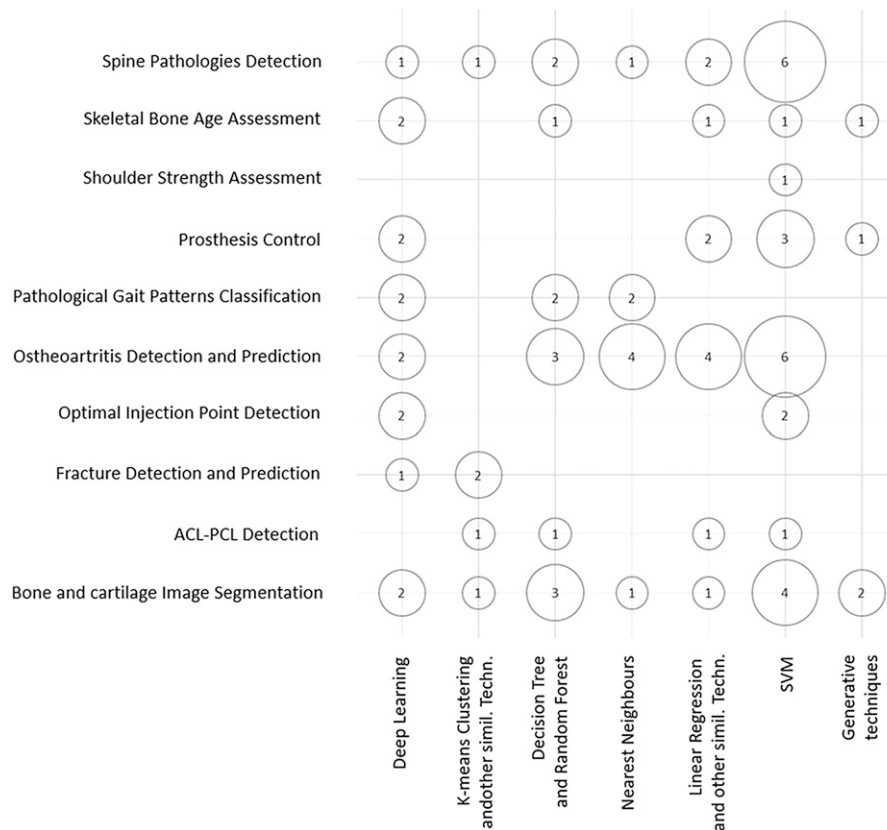


Fig. 7

Bubble chart showing orthopaedic studies by ML techniques and ML techniques by input data. ACL-PCL = anterior cruciate ligament-posterior cruciate ligament, and SVM = support vector machine. (Reproduced, under Creative Commons license Attribution 4.0 International [CC BY 4.0], from: Cabitza F, Locoro A, Banfi G. Machine learning in orthopaedics: a literature review. *Front Bioeng Biotechnol.* 2018 Jun 27;6:75. © 2018 Cabitza, Locoro and Banfi.)

Orthopaedics and AI Research

Over the last few years, there has been a substantial increase in the number of initiatives leveraging the power of AI to address orthopaedic-specific problems. An excellent in-depth systematic review of ML techniques applied to orthopaedic problems over the last 2 decades, by Cabitza et al., was recently published¹². The authors found that spine pathology, OA detection and prediction, and imaging of bone and cartilage were the most studied topics, while DL and support vector machines (SVMs) were the most frequent ML applied algorithms¹². Medical imaging data were by far the most commonly applied input data source¹² (Fig. 7). Even though an in-depth analysis of specific AI-related projects in orthopaedics is beyond the scope of this review, we include a table with some notable examples of such initiatives (Table I). We recommend referring to the specific articles for full details on the methodology and results. However, we would be remiss if we did not direct readers to the website of the *Journal of the American Medical Association* dedicated to interdisciplinary groups demonstrating ML research for health care³⁴.

Challenges Facing AI in Orthopaedics

By now a reasonable question would be “With all the potential for AI to change orthopaedics, why isn’t it happening?” Some


people have gone so far as to accuse AI in medicine of being no more than a modern day “Mechanical Turk”—the fake 18th-century chess-playing machine^{1,35}. This would imply that ML provides no advantage over traditional statistics, which is simply not true³⁶. There is no question that AI works. Contrary to our initial contrast between traditional statistics and ML to illustrate a point (Fig. 2), a clear distinction between the 2 techniques is difficult to make and it is more appropriate to view the 2 as lying on a spectrum¹⁶. Indeed, Christodoulou et al., in their PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses)-compliant systematic review, concluded that ML performed no better than logistic regression (LR) and that many studies had poor methodology, with the LR-ML model validation being either not sound or not well reported³⁷. The authors offered 2 recommendations based on their review of the studies. First, model development and validation methods should be more carefully designed and reported. Second, research should focus more on identifying which algorithms have optimal performance for different types of prediction problems³⁷. Furthermore, the study results reported by Miller et al. on LR versus ML models in a cardiac transplant database “raise the notion that large clinical datasets might lack the accuracy and granularity needed for

Original image



+ 0.04 ×

Adversarial noise



Adversarial example



=


Dermatoscopic image of a benign melanocytic nevus, along with the diagnostic probability computed by a deep neural network.

Perturbation computed by a common adversarial attack technique. See (7) for details.


Combined image of nevus and attack perturbation and the diagnostic probabilities from the same deep neural network.



Diagnosis: Benign



Adversarial rotation (8)



Diagnosis: Malignant

The patient has a history of **back pain** and chronic **alcohol abuse** and more recently has been seen in several...

Adversarial text substitution (9)

The patient has a history of **lumbago** and chronic **alcohol dependence** and more recently has been seen in several...

Opioid abuse risk: High

Opioid abuse risk: Low

277.7 **Metabolic syndrome**
429.9 Heart disease, unspecified
278.00 Obesity, unspecified

Adversarial coding (13)

401.0 **Benign essential hypertension**
272.0 **Hypercholesterolemia**
272.2 **Hyperglyceridemia**
429.9 Heart disease, unspecified
278.00 Obesity, unspecified

Reimbursement: Denied

Reimbursement: Approved

Fig. 8

The anatomy of an adversarial attack. Demonstration of how subtle changes against various AI systems (image recognition and text recognition) can substantially alter clinical care and reimbursement. (Reproduced, with permission of the American Association for the Advancement of Science, from: Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. Science. 2019 Mar 22;363[6433]:1287-9; permission conveyed through Copyright Clearing Center, Inc.)

machine learning methodologies to uncover unique associations.^{38,39} This brings us to the heart of the challenge to adopt AI in orthopaedics. Topol and others have outlined 4 things

necessary for the successful implementation of AI in medicine: big accurate data sets, powerful computers, cloud computing, and open source algorithmic development¹. Similar to

Topol, the highly noteworthy review by Kohane in *Science* should be considered required reading for any novice to AI and precision medicine for a perspective of how these 2 may evolve in the future²¹. However, an important point to remember as highlighted by Cabitza et al. is that the application of ML in the orthopaedic field is still limited to Phase-2, or even smaller, studies¹².

Ethical Considerations

Much of AI today involves ML in the form of training deep neural networks, which requires large amounts of data. This suggests that developers of AI need to have access to sensitive medical information, which raises concerns about patient privacy. These concerns were exemplified in 2015, when DeepMind—an English company owned by Google—reached an agreement with the U.K. National Health Service (NHS) to get access to the medical information of 1.6 million individuals. In 2017, the U.K. Information Commissioner's Office (ICO) found that this agreement violated the Data Protection Act⁴⁰.

In addition to privacy issues, there are serious concerns associated with AI biasing. If the data on which the AI was trained involve any degree of bias, this will result in widespread systematic analytical mistakes. For example, much of the available historic medical data in the U.S. is predominantly based on medical records of white men^{41,42}. An algorithm trained on such data may make less accurate predictions about women, ethnic minorities, or any other underrepresented group in the data set. This is similar to the bias of human judgment, which stems from an individual's past experiences, which in turn can lead to inaccurate clinical decisions. However, the widespread simultaneous deployment of the biased AI and its use by thousands of clinicians may have more detrimental implications to patient care and safety. Moreover, AI does not possess traits that are uniquely human like morality and intuition and is liable to make mistakes that seem absurd to humans⁴³. Unfortunately, these vulnerabilities in AI can be exploited by various players interested in the trillion-dollar health-care industry for financial, political, and other motivations. *Adversarial attacks* are inputs to an AI model that are intentionally crafted to bias or force the model to make a mistake⁴⁴. Finlayson et al. detailed a fantastic summary of how these attacks could impact medical diagnosis and decision support, medical insurance claims, drug and device approvals, and clinical trials (Fig. 8)⁴⁴. Despite the clear danger of these attacks, the authors pose a difficult question, "Should the adversarial-examples problem in health care systems be addressed now—in the early, uncertain days of medical AI algorithms—or later, when the algorithms and the protocols governing their use have been firmed up?"⁴⁴ The authors then suggested that regulating AI now may help defend against some of these problems but at the risk of locking us into inaccurate threat-based models and unwieldy regulatory structures. The unintended consequence would amount to stifling the types of innovation AI systems require in order to guard against these not yet fully materialized threats.

Finlayson et al. then offered some solutions such as amending currently existing regulatory practices and mechanisms. AI-based software is defined as a medical device and as such falls under the purview of the U.S. Food and Drug Administration. A lifecycle-based framework for regulating AI systems in medicine is covered in further detail by Hwang et al.⁴⁵. Unfortunately, it is inconceivable that exploitation of these AI systems in the cyberworld of medicine at any point in the future of AI development will cease to exist. Cybersecurity is a cat and mouse game, and Finlayson et al. pointed out that it is always easier to break systems than to protect them⁴⁴.


More Challenges and Considerations

Cabitza et al. offered 3 related viewpoints on unintended consequences of ML in medicine in addition to the ethical and black box challenges previously discussed⁴⁶. These viewpoints include reducing the clinical skills of physicians and physician extenders (deskilling), reliance on data without appropriate context, and not fully appreciating the intrinsic uncertainty of clinical medicine. First, deskilled clinical practitioners reliant on AI systems, which, if compromised (i.e., adversarial attacks), may lead to serious consequences in the delivery of care. There is also the concern that as the diagnostic performance of AI systems reaches that of humans, there may be subtle loss of self-confidence and willingness of a physician to provide a definitive diagnosis as there may be 2 conflicting opinions (the practitioner's and the AI's). Second, AI techniques help make decisions based on data that are presumed to be both reliable and a complete representation of the clinical scenario. However, ML models only identify patterns within data and are at risk for generating incorrect conclusions because critical contextual information is at times difficult or impossible to include in the data. This is especially problematic if overreliance on AI systems erodes clinical practitioners' skills and ability to interpret data within appropriate clinical context. Finally, the deeply embedded intrinsic uncertainty of the medical decisions that feed AI models may be underappreciated. Consequently, the reliability and accuracy of ML performance may continue to suffer if algorithms are not adapted to account for the quality of medical data⁴⁶.

Conclusions

The brief history and explanation of AI given above will hopefully help practitioners to gain new insight into the topic and a greater appreciation of how AI may be able to positively affect orthopaedics. Augmented intelligence, rather than artificial intelligence, may be the proper way to view this exciting and promising field. AI could provide solutions to the increasing demands of redundant and repetitive tasks that are lower on the intellectual spectrum and contribute to physician burnout and medical mistakes. However, challenges regarding the ethical deployment, regulatory challenges, and the clinical superiority of AI over traditional statistics and decision-making remain to be resolved.

Appendix

 Supporting material provided by the authors is posted with the online version of this article as a data supplement at [jbjs.org \(http://links.lww.com/JBJS/F720\)](http://links.lww.com/JBJS/F720). ■

Thomas G. Myers, MD, MPT¹
Prem N. Ramkumar, MD, MBA²
Benjamin F. Ricciardi, MD¹
Kenneth L. Urish, MD, PhD³
Jens Kipper, PhD⁴
Constantinos Ketonis, MD, PhD¹

¹Divisions of Adult Reconstruction (T.G.M. and B.F.R.) and Hand and Upper Extremity Surgery (C.K.), Department of Orthopaedics and Rehabilitation, University of Rochester Medical Center, Rochester, New York

²Machine Learning Arthroplasty Laboratory, Cleveland Clinic Foundation, Cleveland, Ohio

³Department of Orthopaedics and The Bone and Joint Center, Magee Women's Hospital of the University of Pittsburgh Medical Center, Pittsburgh, Pennsylvania

⁴Department of Philosophy, University of Rochester, Rochester, New York

Email address for T.G. Myers: thomas_myers@urmc.rochester.edu

ORCID iD for T.G. Myers: [0000-0003-3545-2333](https://orcid.org/0000-0003-3545-2333)
ORCID iD for P.N. Ramkumar: [0000-0002-1704-9156](https://orcid.org/0000-0002-1704-9156)
ORCID iD for B.F. Ricciardi: [0000-0001-9384-4840](https://orcid.org/0000-0001-9384-4840)
ORCID iD for K.L. Urish: [0000-0002-4511-1308](https://orcid.org/0000-0002-4511-1308)
ORCID iD for J. Kipper: [0000-0002-8138-0855](https://orcid.org/0000-0002-8138-0855)
ORCID iD for C. Ketonis: [0000-0003-4535-9079](https://orcid.org/0000-0003-4535-9079)

References

1. Topol EJ. Deep medicine: how artificial intelligence can make healthcare human again. 1st ed. New York: Basic Books; 2019.
2. Wikipedia. Dartmouth workshop. Accessed 2019 Aug 3. https://en.wikipedia.org/wiki/Dartmouth_workshop
3. Hashimoto DA, Rosman G, Rus D, Meireles OR. Artificial intelligence in surgery: promises and perils. *Ann Surg*. 2018 Jul;268(1):70-6. Epub 2018 Feb 2.
4. Naylor CD. On the prospects for a (deep) learning health care system. *JAMA*. 2018 Sep 18;320(11):1099-100. Epub 2018 Sep 5.
5. Maxmen JS. The post-physician era: medicine in the twenty-first century. New York: Wiley; 1976.
6. Scarlet A. A machine learning primer for clinicians—part 1. *Histalk*. 2018. Accessed 2019 Aug 5. <https://histalk2.com/2018/10/17/a-machine-learning-primer-for-clinicians-part-1/>
7. Chollet F. Deep learning with Python. Shelter Island, New York: Manning Publications; 2018.
8. Bini SA. Artificial intelligence, machine learning, deep learning, and cognitive computing: what do these terms mean and how will they impact health care? *J Arthroplasty*. 2018 Aug;33(8):2358-61. Epub 2018 Feb 27.
9. McClelland C. The difference between artificial intelligence, machine learning, and deep learning. *Medium*. 2017. Accessed 2019 Aug 3. <https://medium.com/iotforall/the-difference-between-artificial-intelligence-machine-learning-and-deep-learning-3aa67bff5991>
10. Wikipedia. Machine learning. Accessed 2019 Aug 3. https://en.wikipedia.org/wiki/Machine_learning
11. Reznik AM, Urish KL. Understanding the impact of artificial intelligence on orthopaedic surgery. *American Academy of Orthopaedic Surgeons (AAOS Now)*. 2018. Accessed 2019 Apr 21. <https://www.aaos.org/AAOSNow/2018/Sep/Research/research01/?ssopc=1>
12. Cabitza F, Locoro A, Banfi G. Machine learning in orthopedics: a literature review. *Front Bioeng Biotechnol*. 2018 Jun 27;6:75.
13. Galbusera F, Casaroli G, Bassani T. Artificial intelligence and machine learning in spine research. *JOR Spine*. 2019 Mar 5;2(1):e1044.
14. Urish KL, Reznik AM. How would a computer diagnose arthritis on a radiograph? *American Academy of Orthopaedic Surgeons (AAOS Now)*; 2018. Accessed 2019 Apr 21. <https://www.aaos.org/AAOSNow/2018/Dec/Research/research04/>
15. Urish KL, Keffalas MG, Durkin JR, Miller DJ, Chu CR, Mosher TJ. T2 texture index of cartilage can predict early symptomatic OA progression: data from the osteoarthritis initiative. *Osteoarthritis Cartilage*. 2013 Oct;21(10):1550-7. Epub 2013 Jun 15.
16. Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA*. 2018 Apr 3;319(13):1317-8.
17. Chen JH, Alagappan M, Goldstein MK, Asch SM, Altman RB. Decaying relevance of clinical data towards future decisions in data-driven inpatient clinical order sets. *Int J Med Inform*. 2017 Jun;102:71-9. Epub 2017 Mar 18.
18. Shi L, Wang XC, Wang YS. Artificial neural network models for predicting 1-year mortality in elderly patients with intertrochanteric fractures in China. *Braz J Med Biol Res*. 2013 Nov;46(11):993-9. Epub 2013 Nov 18.
19. Ramkumar PN, Karnuta JM, Navarro SM, Haeberle HS, Iorio R, Mont MA, Patterson BM, Krebs VE. Preoperative prediction of value metrics and a patient-specific payment model for primary total hip arthroplasty: development and validation of a deep learning model. *J Arthroplasty*. 2019 Oct;34(10):2228-2234.e1. Epub 2019 May 2.
20. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015 May 28;521(7553):436-44.
21. Kohane IS. Health care policy. Ten things we have to do to achieve precision medicine. *Science*. 2015 Jul 3;349(6243):37-8. Epub 2015 Jul 2.
22. Drees J. Vanderbilt tests virtual assistant that can voice patient EHR data to physicians. *Becker's Health IT & CIO Report*; 2019. Accessed 2020 Jan 7. https://www.beckershospitalreview.com/ehrs/vanderbilt-tests-virtual-assistant-that-can-voice-patient-ehr-data-to-physicians.html?oly_enc_id=923016729601G9D
23. Marcotte B. A prescription for physician frustration. *University of Rochester*. 2019. Accessed 2020 Jan 7. <https://www.rochester.edu/newscenter/virtual-physician-assistant-erecords-373972/>
24. Morgan J. A Simple explanation of 'the Internet of things'. *Forbes*. 2014. Accessed 2020 Jan 7. <https://www.forbes.com/sites/jacobmorgan/2014/05/13/simple-explanation-internet-things-that-anyone-can-understand/#5a4745841d09>
25. Chiang CY, Chen KH, Liu KC, Hsu SJP, Chan CT. Data collection and analysis using wearable sensors for monitoring knee range of motion after total knee arthroplasty. *Sensors (Basel)*. 2017 Feb 22;17(2):418.
26. FocusMotion. 2019 Accessed 2019 Apr 21. <http://focusmotion.io/>
27. Hexoskin Health Sensors & AI. Hexoskin smart shirts - cardiac, respiratory, sleep & activity metrics. Accessed 2019 Apr 21. <https://www.hexoskin.com/>
28. Ramkumar PN, Muschler GF, Spindler KP, Harris JD, McCulloch PC, Mont MA. Open mHealth architecture: a primer for tomorrow's orthopedic surgeon and introduction to its use in lower extremity arthroplasty. *J Arthroplasty*. 2017 Apr;32(4):1058-62. Epub 2016 Nov 17.
29. Wikipedia. mHealth. Accessed 2019 Apr 21. <https://en.wikipedia.org/wiki/MHealth>
30. Ravi D, Wong C, Deligianni F, Berthelot M, Andreu-Perez J, Lo B, Yang GZ. Deep learning for health informatics. *IEEE J Biomed Health Inform*. 2017 Jan;21(1):4-21. Epub 2016 Dec 29.
31. Ramkumar PN, Haeberle HS, Ramanathan D, Cantrell WA, Navarro SM, Mont MA, Bloomfield M, Patterson BM. Remote patient monitoring using mobile health for total knee arthroplasty: validation of a wearable and machine learning-based surveillance platform. *J Arthroplasty*. 2019 Oct;34(10):2253-9. Epub 2019 May 16.
32. Ramkumar PN, Haeberle HS, Navarro SM, Sultan AA, Mont MA, Ricchetti ET, Schickendantz MS, Iannotti JP. Mobile technology and telemedicine for shoulder range of motion: validation of a motion-based machine-learning software development kit. *J Shoulder Elbow Surg*. 2018 Jul;27(7):1198-204. Epub 2018 Mar 7.
33. Ramkumar PN, Haeberle HS, Bloomfield MR, Schaffer JL, Kamath AF, Patterson BM, Krebs VE. Artificial intelligence and arthroplasty at a single institution: real-world applications of machine learning to big data, value-based care, mobile health, and remote patient monitoring. *J Arthroplasty*. 2019 Oct;34(10):2204-9. Epub 2019 Jun 17.
34. JAMA. Machine learning. 2019. Accessed 2020 Jan 7. <https://sites.jamanetwork.com/machine-learning/>.
35. Wikipedia. The Turk. Accessed 2019 Aug 11. https://en.wikipedia.org/wiki/The_Turk
36. Davison J. No, machine learning is not just glorified statistics. *Medium*. 2018. Accessed 2020 Jan 7. <https://towardsdatascience.com/no-machine-learning-is-not-just-glorified-statistics-26d3952234e3>
37. Christodoulou E, Ma J, Collins GS, Steyerberg EW, Verbakel JY, Van Calster B. A systematic review shows no performance benefit of machine learning over logistic

- regression for clinical prediction models. *J Clin Epidemiol.* 2019 Jun;110:12-22. Epub 2019 Feb 11.
- 38.** Miller PE, Pawar S, Vaccaro B, McCullough M, Rao P, Ghosh R, Warier P, Desai NR, Ahmad T. Predictive abilities of machine learning techniques may be limited by dataset characteristics: insights from the UNOS database. *J Card Fail.* 2019 Jun; 25(6):479-83. Epub 2019 Feb 6.
- 39.** Akbilgic O, Davis RL. The promise of machine learning: when will it be delivered? *J Card Fail.* 2019 Jun;25(6):484-5. Epub 2019 Apr 9.
- 40.** BBC. Google DeepMind NHS app test broke UK privacy law. 2017. Accessed 2020 Jan 7. <https://www.bbc.com/news/technology-40483202>
- 41.** Feldman S, Ammar W, Lo K, Trepman E, van Zuylen M, Etzioni O. Quantifying sex bias in clinical studies at scale with automated data extraction. *JAMA Netw Open.* 2019 Jul 3;2(7):e196700.
- 42.** Oh SS, Galanter J, Thakur N, Pino-Yanes M, Barcelo NE, White MJ, de Bruin DM, Greenblatt RM, Bibbins-Domingo K, Wu AH, Borrell LN, Gunter C, Powe NR, Burchard EG. Diversity in clinical and biomedical research: a promise yet to be fulfilled. *PLoS Med.* 2015 Dec 15;12(12):e1001918.
- 43.** Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. Presented as a poster exhibit at the 2015 International Conference on Learning Representations. Computational and Biological Learning Society; 2015 May 7-9. San Diego, CA. <https://dblp.org/db/conf/iclr/iclr2015>
- 44.** Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. *Science.* 2019 Mar 22;363(6433):1287-9.
- 45.** Hwang TJ, Kesselheim AS, Vokinger KN. Lifecycle regulation of artificial intelligence- and machine learning-based software devices in medicine. *JAMA.* 2019 Nov 22. Epub 2019 Nov 22.
- 46.** Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. *JAMA.* 2017 Aug 8;318(6):517-8.
- 47.** Thong W, Parent S, Wu J, Aubin CE, Labelle H, Kadoury S. Three-dimensional morphology study of surgical adolescent idiopathic scoliosis patient from encoded geometric models. *Eur Spine J.* 2016;25(10):3104-13.
- 48.** Olczak J, Fahlberg N, Maki A, Razavian AS, Jilert A, Stark A, Stark A, Sköldenberg, Gordon M. Artificial intelligence for analyzing orthopedic trauma radiographs. *Acta Orthop.* 2017 Dec;88(6):581-6.
- 49.** Kruse C, Eiken P, Vestergaard P. Machine learning principles can improve hip fracture prediction. *Calcif Tissue Int.* 2017 Apr;100(4):348-60.
- 50.** Cilla M, Borgiani E, Martinez J, Duda GN, Checa S. Machine learning techniques for the optimization of joint replacements: application to a short-stem hip implant. *PLoS One.* 2017 Sep 5;12(9):e0183755.
- 51.** Konda SR. Predictive software helps optimize efficient care for older patients with orthopedic fractures. NYU Langone Health. Orthopedic Surgery 2018 Year in Review. Accessed 2020 Jan 21. <https://nyulangone.org/news/predictive-software-helps-optimize-efficient-care-older-patients-orthopedic-fractures>
- 52.** Karnuta JM, Navarro SM, Haeberle HS, Billow DG, Krebs VE, Ramkumar PN. Bundled care for hip fractures: a machine-learning approach to an untenable patient-specific payment model. *J Orthop Trauma* 2019 Jul;33(7):324-30.
- 53.** Shah RF, Martinez AM, Padoia V, Majumdar S, Vail TP, Bini SA. Variation in the thickness of knee cartilage. The use of a novel machine learning algorithm for cartilage segmentation of magnetic resonance images. *J Arthroplasty.* 2019 Oct; 34(10):2210-5.
- 54.** Harris AHS, Kuo AC, Weng Y, Trickey AW, Bowe T, Giori NJ. Can machine learning methods produce accurate and easy-to-use prediction models of 30-day complications and mortality after knee or hip arthroplasty? *Clin Orthop Relat Res.* 2019 Feb; 477(2):452-60.
- 55.** Greenstein A, Teitel J, Mitten DJ, Ricciardi B, Myers TG. A neural network predicts discharge disposition after primary total joint arthroplasty. Read at the Annual Meeting of the American Academy of Orthopaedic Surgeons; 2019 Mar 15; Las Vegas, NV.
- 56.** Fontana MA, Lyman S, Sarker GK, Padgett DE, MacLean CH. Can machine learning algorithms predict which patients will achieve minimally clinically important differences from total joint arthroplasty? *Clin Orthop Relat Res.* 2019 Jun; 477(6): 1267-9.
- 57.** Leopold S. Editor's Spotlight/Take 5: Can machine learning algorithms predict which patients will achieve minimally clinically important differences from total joint arthroplasty? *Clin Orthop Rel Res.* 2019 Jun. 477(6):1262-6.
- 58.** Thirukumaran CP, Zaman A, Rubery PT, Calabria C, Li Y, Ricciardi BF, Bakhsh WR, Kautz H. Natural language processing for the identification of surgical site infections in orthopaedics. *J Bone Joint Surg Am.* 2019 Dec 18;101(24):2167-74.