

Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.

# Artificial intelligence in medical devices and clinical decision support systems

Elliot B. Sloane<sup>a,b</sup>, Ricardo J. Silva<sup>b,c</sup>

<sup>a</sup>Foundation for Living, Wellness, and Health, Osprey, FL, United States, <sup>b</sup>Computing Sciences, Villanova University, Villanova, PA, United States, <sup>c</sup>Foundation for Living, Wellness, and Health, Orlando, FL, United States

Artificial Intelligence (AI) is a broad conceptual category of computer software that is designed to mimic, emulate, or improve human decision-making. There are two broad categories of AI: rules-based expert systems (ES), and machine-learning (ML) systems.

Rules-based ES are designed to replicate the interpretation and decision-making of a subject matter expert. Some simple examples include

- car antilock braking systems, which automatically apply and release (pump) the brakes like an expert driver would do in order to maintain maximum control of the auto;
- college selection and recommendation system based on student grades, aptitudes, and preferences; and
- Medication selection or recommendations derived from clinical best practices rules and based on the input of basic data about patient, symptoms, side effects, and costs.

ML AI is based on a different premise and design, in which the computer software either is trained to recognize or infer desired patterns, or trial-and-error experimentation to obtain or attain desired results. Examples include

- digital camera software, which detects people's eyes and faces, and uses that information to adjust focus and flash intensity;
- modern antivirus software that is based on detection of anomalous system behaviors or parameters; and
- Cardiac arrhythmia detection algorithms used in automated electronic defibrillators (AEDs).

Both ES and ML system can be paired with one or more human operators to improve the quality, performance, or safety of the overall system. In many cases, human operators also take part in training the ML system, and the process is much like training a pet. The operator provides examples

and positive reinforcement for correct decisions and examples and negative reinforcement for poor decisions. For example, the computer can be taught or can infer patterns from a sufficiently large set of matching (correct) or defective (incorrect) examples that were previously scored or provided by one or more experts. An automated blood laboratory system, or breast cancer screening software, might be based on ML. In other cases, the human operator can take control if and when the computer can no longer make a successful interpretation of the data (e.g., a self-driving car might return control to the human driver if confused, overwhelmed, or specified safety parameters are exceeded, or a bank's credit card supervisor might personally interview an applicant if the software cannot make a credit approval decision).

In addition, ES and ML tools can be combined into a single product or suite of products, such as the features offered by most contemporary antivirus software.

## History and background

The earliest AI systems were analog mechanical devices, such as wind-vane self-steering for boats, centrifugal speed controls for engines, or automatic chokes and transmissions for automobiles. Mainframe computer programs in the 1960s were developed to automate manufacturing planning, based on organized bills of materials for components and subassemblies, and/or for early military attack and targeting. The advent of small, low-cost microprocessors in the 1970s allowed computerized implementation of prior analog AI, such as self-steering boat systems (Francis West, 1979), following ES rules and algorithms to assert positive, intelligent control of many manufacturing, transportation, and related tasks.

By the mid-1980s, real-time minicomputers coupled with advancing analog-to-digital conversion chips allowed the development of early ML-based experiments and products. A classic early ML experiment involved a computer learning to balance a broomstick on an artificial hand (Widrow, 1987). The general concept has the words embedded themselves: methods for machines to “learn” on their own, without a priori writing of programs and algorithms. There are two broad categories of ML that have evolved: machines that “learn” using teaching examples or subsets, and machines that explore or discover facts and solutions on their own.

An example of the first ML category might be an industrial robot that learns how to see if a component is properly assembled and located on a circuit board. In this case, training the computer would entail showing it examples of the target component, from many angles that its camera might be presented. In addition, the operator might show images of incorrect components. “Correct” examples might include labeling the component, shape, color, and size. “Incorrect” examples might include outdated and incorrect or inferior substitution components. Further training might include images of correct positioning and alignment with respect to other nearby components, and “incorrect” examples might illustrate common or potential erroneous assembly.

“Deep Learning” systems are capable of receiving, recalling, and analyzing large sets of correct and incorrect examples, and can deduce patterns (Weiss, 1990), such as “object is of a class of bird, because not like fish, reptile, or land mammal, but is not among a set of known living birds. Therefore, perhaps it is a Dodo, extinct pigeon, or dinosaur.”

An example of the second ML category might be a program that can try many possible configurations and combinations of assembly, and either testing the success of the assembly, or receiving a score or feedback from a human trainer or a separate testing subsystem. In this category of ML, the computer must “explore,” which infers that the computer is capable of inventing different assembly alternatives. That is, either the computer has programs and actuators that can randomly or methodically try different positions and parts, or an associated subsystem—or operator—creates the iterative or random or iterative combinations.

Computer games (and real war-fighting tools) further advanced the field, adding intelligent computer responses and counterattacks to increase the complexity of single- and multiplayer games.

## Current state

Today, large numbers of smartphone apps, home appliances, and personal products like automobiles include ES and ML capabilities, enabling tasks like voice recognition, complex Global Positioning System (GPS)-based routing

recommendations, and a growing number of Food and Drug Administration (FDA)-approved medical devices (Jiang, 2017).

In the case of rule-based ES, AI is somewhat easy to visualize using simple logic statements, like “if daylight is bright enough to trigger a sensor, then turn off the porch light,” or, “if the child has already received the maximum safe daily dose of acetaminophen, and the fever is still above 102°, and the child has become listless, and the child is not taking fluids or urinating, then call the family physician immediately.”

Underlying all computer-based ML programs and systems is a binary digital computer, which only has a limited—but fast—ability to add and subtract “1s and 0s.” Therefore, despite the apparent “ability” of a ML program to successfully identify a chicken or balance a broom, the actual digital computer is ultimately limited to rudimentary statistical or numeric calculations.

To simplify: the computer ultimately must compute a numeric representation for the object it is examining. That representation may be in multiple dimensions, such as size, shape, color, and orientation. That numeric representation is ultimately compared to the template or a set of alternatives. The simple digital computer method is a simple subtraction of the first objects numeric representation from the second object’s numeric representation. If the objects are the same, the subtraction would result in a zero remainder. The larger the difference between numerical representations of the two objects’ dimensions, the further the resulting subtraction result will be from zero.

Similarly, groups of similar objects, such as pieces of candy, will share statistically similar average values (and/or the min, max, median, or mode values). In order to decide if a new object is the same as the original group of objects, the digital computer is limited to subtracting the numeric representation of the new object from the average values of the original group of objects. Again, the closer the result is to zero, the stronger the evidence is that the objects are identical. To compare two groups of samples, such as stained pathology samples of cancer tissues, the digital computer “knows” they are identical only by subtracting the numeric representations from each other.

The underlying statistical tools can be rather sophisticated, representing, for example, multifactor covariance calculations. The clinical user no longer needs to understand deeply the underlying calculations or statistics, because those tools have become broadly, and often freely, available. What the user needs to understand, or provide the interpretation about, is cofactors that are likely related (such as obesity, diet, exercise, education, and diabetes), vs ones that are less likely associated, such as eye color, native language, or shoe size.

One other related computational area or methodology is fuzzy logic. In simple terms, fuzzy logic algorithms are

designed to handle vague similarities, or shades of gray. For example, there are many unique hybrid chickens and parrots. A new bird sample can be segregated into chicken or parrot categories based on similarities that are more, or less, like the matching template. Fuzzy logic software can help by providing mathematical representations and categorizations such as “more like a chicken, less like a parrot, or not much like either.”

Many ML approaches and algorithms are in commercial use, and even more are in experimental trials. Recent examples that are well-publicized include self-driving cars (Google, Tesla, and Uber), voice-activated home assistance devices (Cortana, Alexa, Google Voice, and Siri), and autonomous robots that are being used in industrial, home, and healthcare settings (Bickmore, 2018).

## Natural language processing (NLP)

There may be no area of AI with as much history, hope, and hype as NLP. On the surface, the newcomer to the topic may well wonder why a process that very young babies master so easily. Many breeds of domesticated animals even seem to master rudimentary commands easily. Further, a significant number of “everyday people” who have lived near multilingual national borders for centuries, like most of Europe, have been able to acquire and master multiple languages with apparent ease. Some countries like Switzerland officially incorporate multiple languages, and since World War II, many, perhaps most, countries have accepted English as de facto second language for business, science, and international trade.

The challenge of NLP for computers begins with the binary nature of the technology, because all concepts must be reduced to strings of 1s and 0s, and the computer’s only way to “compute” similarity or difference is simple addition or subtraction. That is, two words or sentences are identical if the computer can subtract the two binary representations and get a result of zero.

Natural human languages, however, contain very rich and complex nuances, some of which are only understood by the context of the message, or the peculiarities of the local language itself.

Some speech recognition examples may help illustrate:

*‘Time flies like an arrow’ is a phrase that we easily interpret as ‘time goes very swiftly, like an arrow flies’. However, the computer is stuck with the task of grouping the words together in different combinations, trying to ‘make sense’ of the possible combinations. For example, do ‘time flies’ ‘like’ ‘arrows’ in the same sense that ‘fruit flies’ ‘like’ ‘bananas?’ Or, maybe, ‘like’ infers affection, rather than food or scent! Further, do these ‘time flies’ like only ‘an arrow’ singly, or do ‘time flies’ like all arrows? Or, perhaps, ‘time’ has wings and ‘flies like an arrow’. And, if so, is ‘Time’ or are ‘time*

*flies’ truly winged insects like fruit flies, or some other form of life such as a bat, a bird, or even a pig that simply shares the characteristic of ‘liking of an arrow’.*

*Or, as another health-oriented example, an emergency 911 caller might proclaim: ‘My baby is burning up!’ Is ‘My baby’ an infant, a child, sibling, a bride, a husband, or, perhaps another cherished and beloved friend? Is ‘My baby’ really ‘burning up’, and on fire, or perhaps running a very high fever, badly sun-burned, or perhaps injured with boiling water or oil, or, even, perhaps, suffering a life-threatening reaction to something that was eaten, swallowed, or touched. Last, perhaps ‘My baby’ is burning with anger, embarrassment, or resentment about something said, heard, or seen!*

Now, add to the above examples the complexity of personal and/or community dialects, and the problem becomes even worse. Even without possible effects of alcohol, medicine, injury, or speech impediment, “time,” for example, may be pronounced “tiime,” or “taahm,” (Southern US drawl) or even “toyme” (British Cockney) or “tohm” (formal British).

And, lastly, consider the variations of idiosyncratic uses of words in different national variants of, for example, English. In the United States, a car’s front hood is called the bonnet in countries based on British dialects; trunk becomes boot, and windshield becomes windscreen. Similarly, pants become knickers, and telephone becomes telly. In health care, the common US medicine named acetaminophen is known more widely worldwide as paracetamol, and a US “stretcher” to transport sick or injured patients is a “trolley” in British parlance.

A different class NLP problem is a bit easier to tackle: voice synthesis from computerized text or numbers. This is an easier task for the computer because the context can be provided by the human listener. The computer can either have a library of digitized words and phrases that it can select, or it can piece together portions of the word as a child does, syllable by syllable. For example, early computerized weather reporting could be handled with limited libraries of typical words, such as “wind” “one” “two” “knots” “north” “heading” “one” “two” “zero.”

In the late 1990s and early 2000s, the earliest personal and professional voice translation and transcription software, such as Dragon, or professional dictation systems sold by companies like Dictaphone, began to make it possible for ever-growing clusters of users such as lawyers or radiologists to correctly dictate reports or observations directly into a computer (Devine, 2000). Assistive technologies began to appear in personal computer (PC)-based products, too, that could read out loud for vision-limited users or pronounce street names and directions in a GPS designed for busy drivers.

In the late 2000s, in fact, low-cost voice synthesis chips and low-cost analog-to-digital chips brought constrained

voice recognition and speech capabilities to games and toys, mass-commercializing products that could entertain, and ultimately “communicate” simple concepts. Even a simple baby doll that could cry, say Mama, and utter phrases like “I love you” were surprisingly well received. Also, web-based tools like Google Search and Microsoft Bing raised awareness of the power that computer-assisted text searching, and similarity ranking can bring to every-day life and business. Suddenly, AI tools like Google Search could begin to overcome many common errors, such as misspelling Mississippi, Spanakopita, or Sjogren’s Syndrome.

In the latter 2000s, the arrival of smartphones, tablets, and apps introduced a much richer array of consumer-friendly NLP products like Apple’s Siri, which set the stage for the huge upswelling of smart home automation and Internet of Things appliances like Amazon’s Alexa and Google Assistant.

Voice-activated automated attendant software is now common in many, if not most, banking, finance, travel, and other businesses. Alexa and Google Assistant are apparently soon to be joined by updated and enhanced versions of Apple’s Siri, which will add voice recognition for many everyday tasks and activities.

The two large branches of AI, ES, and ML, could become much more accessible and “natural” for users if the keyboard and mouse can be eliminated with NLP (Jiang, 2017). This may be especially helpful for aging citizens with deteriorating vision, hearing, or hand-eye coordination limitations. An ES example might be assisting a patient to select a nearby medical specialist for, say, cardiology, solely by voice dialog. ML could also be enhanced with NLP by allowing users to train the system to their personal preferences. That might look like a next-generation NEST (<https://nest.com/>) thermostat which the user could train to increase or set-back heating or air-conditioning to meet their personal preferences, and/or to override the default settings if staying at home on a weekday with a sick child or parent.

## Healthcare AI

The applications of ML probably were first widely seen in diagnostic healthcare products, such as tumor detection or medication recommendation and safety systems. Large-scale systems ML-based systems first emerged in laboratory and pathology, where expert pathologists could manage the training and validation of the results. This kept the “learned intermediary” (i.e., the licensed pathologist or a certified technician) “in the loop” with the labs’ diagnostic products and services. The quality and safety regulations embedded in the 1988 CMS Clinical Laboratory Improvement Amendments (CLIA) programs were developed to ensure that competent human operators followed well-documented quality assurance practices to ensure the

automated laboratory devices produced accurate, trustworthy results (Centers for Disease Control and Prevention (CDC) et al., 2003). CLIA processes are intended to ensure that the human expert remains the authoritative expert in control of the diagnostic systems (Ehrmeyer, 2004).

Many emerging and production ML-based imaging analysis systems depend on an expert human operator’s control and decisions for quality and nuanced diagnosis. Although the ML-based tumor detection algorithms may be “near perfect,” there is such wide diversity in human physiology and disease morphology that computers cannot always make definitive decisions (ter Voert, 2018). On the other hand, the human expert can become fatigued and/or distracted, and the ML-system can help double-check the human operator’s work.

As ML-based diagnostic systems have become more accepted, therapeutic applications have emerged. The use of Smart Infusion Pump Systems has become the preferred method of ensuring IV medication safety (Giuliano, 2018). Most of those systems are based on expert-system AI, rather than ML, but the robustness and reliability of such devices have led to more complete ML-based applications like implanted insulin pumps and emerging closed-loop artificial pancreas devices.

The distinction of closed-loop ML devices for life-critical medical care cannot be overstated. In a closed-loop implanted defibrillator, pacemaker, or pancreas, the “learned intermediary,” or expert operator is being gradually removed from the process. It has taken decades of development and regulatory refinements to allow such products to come to market.

In the absence of any major or catastrophic setbacks, one can reasonably expect that ML-based “intelligence” being developed for life-critical consumer and commercial products like self-driving cars will eventually reinforce society’s willingness to allow more ML-based, autonomous healthcare diagnostic and therapeutic devices. The positive potential for improving safety, flexibility, availability, and affordability for ever-growing, ever-aging, and largely dispersed populations of citizens and patients will need every possible automation and safety enhancement possible. Several healthcare-related examples have been cited above, and applications in health care are multiplying rapidly thanks to apps and ubiquitous memory and cloud computing resources. Broad classes of Healthcare AI (HAI) include (Jiang, 2017; Shah, 2018)

- vision and pattern detection (tumors, breaks, foreign objects, gait irregularities);
- speech and natural language detection and response (voice-controlled robots, self-service patient kiosks, and neurological deficit detection);
- expert clinician advice to assist other caregivers or patients (stroke diagnosis and treatment, best clinical practices and guidelines, and medication recommendations);

- autonomous surgical instruments or enhancements (beating heart surgery, robotic surgery sensory augmentation); and
- closed-loop medical devices (artificial pancreas, AED)

AI has been introduced into most electronic medical record systems for a wide variety of tasks. Most are supplemental tools to either accelerate medical decisions, reduce or eliminate errors, and/or improve healthcare quality, compliance to standards, cost-effectiveness, or satisfaction. Most smart infusion pump systems, for example, are now designed with modules and ancillary devices like bar code and radio frequency identification (RFID) readers to help enforce “Five Rights:” right patient, right drug, right dose, and right route, right time (Podaima, 2018).

### NLP and medical technology

There are great hopes surrounding potential improvements in healthcare delivery, efficiency, and efficacy based on enhanced NLP application. One application, for example, is to add NLP-based “syndromic reporting” to population health data. The CDC describes syndromic reporting as an outgrowth of the September 11 terrorist attacks and subsequent anthrax attacks (Henning, 2004). The goal of syndromic reporting was originally to obtain earlier, sentinel notice about potential local, regional, or national public health problems from real-time multiple data streams instead of waiting for retrospective, reactive reporting from physicians and hospitals. Since the original surge of interest in 2001, many additional nonterrorist possible benefits have been identified, including the unexpected transmission of highly contagious pandemic diseases like severe acute respiratory syndrome (SARS), Ebola, Zika, or avian flu.

Families, parents, and individuals usually discuss and describe their illnesses using informal descriptors. In fact, few would ever think to use formal international classification of diseases (ICD)-10 codes to explain their illnesses, let alone the abbreviated descriptors in the ICD-10 system. Laypersons will usually describe their symptoms, maladies, or syndromes with phrases or words like “tummy-ache,” “bones hurt,” “on fire,” “sleepless,” or “sore throat.” Though vaguely descriptive, such terms are imprecise, and they also lack context. If the tummy-ache happened immediately after eating week-old left-overs, the probable meaning is quite different than a kid who just binged on Halloween candy or a tourist who just returned from an Ebola-stricken area of Africa! Also, such reports can be combined and analyzed in conjunction with similar reports created by international agencies e.g., (World Health Organization, 2017).

Google, the CDC, and researchers have also done some interesting work correlating Pharmacy prescriptions, Google Search queries, and CDC Flu data trying to identify reliable, earlier flu detection and location information

in a more timely fashion than the weekly retrospective Flu reports posted by the CDC (Centers for Disease Control and Prevention, National Center for Immunization and Respiratory Diseases (NCIRD), 2019).

Another hotly debated aspect of eHealth is the potential, correct, and incorrect uses of free-text reporting and annotations in electronic health record (EHR) and personal health record (PHR) systems. Proponents have argued that requiring and constraining physician and nurse reporting to codes like ICD-10, CPT (<https://www.ama-assn.org/practice-management/cpt>), SNOMED (<http://www.snomed.org/>), LOINC (<https://loinc.org/>), and others have two significant downsides: (1) the effort required to find and post codes is so demanding that only minimal detail is likely to be recorded; (2) the physician and nursing natural language notes are more likely to contain rich, important, and nuanced details that are not readily captured by code-only enforcement; and (3) highly trusted and successful NLP-based tools like Google Search seem to do a very impressive and effective job organizing data from the world wide web without imposing arcane standards or cumbersome codes on every web site and page.

Perhaps the “Holy Grail” of NLP applications in health care would be replacement or enhancement of electronic medical record systems so that physicians, nurses, and patients could reliably use free text or voice input instead of using tedious codes like ICD-10, SNOMED, RxNorm, LOINC, and others. Technically, though, this goal is very, very difficult to achieve, especially for data input (compared to data retrieval and reporting). The ancient computing axiom of Garbage In, Garbage Out (GIGO) is always at work: GIGO. For example, if a specific tumor, like a variety of melanoma is not initially correctly entered into the electronic medical record, analysis, or retrieval of melanomas of that type may be impossible to achieve. Yes, a more general report, analysis, or query might capture the data in a broader category of “cancer,” but little more specificity may be possible. Similarly, medicine names are often so closely similar that typographic—or faulty memory—errors could cause life-threatening risks. Many antibiotics have similar-sounding names, but the benefits, side-effects, and allergic complications can be very, very different. Similarly, many narcotic pain killers use terms that sound like “codeine,” but some are 10× stronger than others, and could cause life-threatening complications from overdose or allergic reactions!

Hopefully at some future date, advanced AI software might reliably infer and even correct information as it is entered. For example, a patient with chronic or reoccurring disease and treatments, the computer might predict the correct terminology during data input. Such a system might, for example, alert the physician or nurse that “According to the EHR records, Mrs. Jones’s osteoarthritis pain has in the recent past been treated with 100mg of generic Celebrex,

or celecoxib. Is that the medicine you want to prescribe, or do you want to continue with this prescription for 500 mg of ibuprofen?” Or, a system might suggest that “Mrs. Jones has had a diagnosis of rheumatoid arthritis since 2009. Do you want to add a new diagnosis of osteoarthritis?”

For query, reporting or analytical purposes, the software might easily search for all potential variants of arthritis, or melanomas, or antibiotics. Retrieval, analysis, or reporting software cannot add specificity or details that are not in the original data, however, nor can errors be corrected. However, broad-based data analysis can take place, e.g., summary statistics of all patients with skin cancers, including all varieties of melanomas, are possible.

## Clinical decision support systems (CDSS)

CDSS are software-based AI tools that can assist physicians, nurses, patients, or other care-givers to make better decisions. A common use of CDSS is to analyze past, current, and new patient data and identify or suggest gaps, errors, safety concerns, or care pathway improvements to the user.

CDSS can be based on either of the two classes of AI: ES or ML. To implement the first category of CDSS, ES-based CDSS, many EHR systems offer sophisticated rule-management subsystems. Some physicians may specialize in specific types of surgery, such as knee replacement, and they may create customized “Order sets” that orchestrate virtually every part of the patient’s care to ensure optimal results. For example, the order set might include presurgical lab tests, prophylactic medication, and patient instructions, customized preparation of trays and sets of surgical instruments, implants, sutures and related supplies, video enhancement systems, and surgical devices, postsurgical patient recovery, intensive care, pain management medication, physical therapy, hospital discharge, and, finally, outpatient and home-care procedures. Each aspect of the order set may include optional or alternative steps to deal with expected care variations and situations.

The ES-based CDSS is relatively simple to set up and validate because the rules often use binary decisions. The rule may be based on a physiologic parameter, such as “If the patient weighs greater or equal to 300 pounds then follow Procedure 1; If the patient weighs less than 100 pounds, follow Procedure 2; Else follow Procedure 3.”

ML-based CDSS may be based on statistical inferences, such as “If this patient’s physiological and medical conditions are similar to the majority of other patients with the same category of disease, then the most common treatment will probably be A, B, and C.” IBM’s Watson Health is an example of such a system. For example, International Business Machines (IBM) teams with highly respected clinical teams, like Memorial Sloan Kettering Hospital for cancer, or Cleveland Clinic for Heart or Orthopedic Surgery.

IBM and the medical experts “train” an instance (i.e., a dedicated Watson Health system) using that hospital’s medical records to create a trained CDSS tool that can readily match that hospital’s optimized practices (Malin, 2013).

Rather than the physician experts going through the tedious ES-system rule documentation, however, the Watson Health ML system uses statistical matching algorithms to group similar and dissimilar patients, treatments, and outcomes automatically. If the hospital’s patient data is accurate and complete, ideally the Watson Health system will infer the basic rules described above (e.g., “If the patient weighs greater than or equal to 300 pounds then...”) Of course, the ML system is completely dependent on correct and complete training. If, for example, Cleveland Clinic follows more flexible decisions for the first weight category, say patients above 280 pounds are usually grouped with the 300 pound patients, then the ML system may select the lower limit, or the ML system may create a fuzzy set of rules, which might roughly be stated as “if the patient is close to or above 300 pounds...”.

One of the biggest challenges of ML systems can be validation and verification of the ML system rules that have been inferred. Some ML systems could self-disclose the “logic” and data underlying the rule to the user, so the user can evaluate whether the recommendation is correct. For example, the ML system might be able to show a distribution of patient weights for source data that led to the 300-pound rule, such as 96% of patients who were 300 pounds or more received Procedure 1, and 3% of patients 280 pounds and above also received Procedure 1.

## Why CDSS is useful

CDSSs can assist clinicians, patients, or other users by serving an advisory or watchdog role. In the advisory role, the CDSS might suggest a best practice for postsurgical patient discharge that could include recommended medications and doses, exercise regimens, and periodic follow-up check-ups and tests to assure optimal results. CDSS can help patients decide on alternate treatment or rehabilitation choices, and/or assist with most-appropriate escalation and treatment of emergent complications in the safest and most cost-effective way possible.

CDSS can also play an important role in reducing medical errors, including patient errors. Many EHR software systems have competent medication error subsystems to protect patients. These CDSSs will check prescribed drugs against known patient allergies, potential drug–drug interactions with existing medication regimens, and/or call out better or less expensive medication alternatives for consideration.

CDSS can offer many advantages to clinicians and to hospital and practice managers. Physicians and nurses may, for example, add or modify rules to assist them in implementing new procedures and rules.

For example, on November 13, 2017, the American College of Cardiology and the American Heart Association reduced the threshold for determining high blood pressure (hypertension) to 130/80 from 140/90 (Vongpatanasin, 2018). Though these blood pressure levels seem very, very close to one another, the new, lower limit may well shift nearly half of the US adult population into the high blood pressure category (Cunningham, 2017)! For a physician or hospital, this new guideline may need to trigger earlier and more aggressive interventions, including modifications of diet and exercise, and medication. In addition, this definition will have implications for health insurance availability and cost, because many more adults will suddenly find they have a new “preexisting condition” even though they did not have that medical condition in their medical records prior to that date!

Presuming the scientific and medical evidence continues to support these new hypertension guidelines, the EHR’s CDSS can help ensure every patient gets the best-recommended treatment even if the physician does not notice a patient’s elevated blood pressure.

Another very volatile recent patient safety issue relates to the prescription of opioid pain medication. Because of the scourge of addictions and deaths, many physicians and hospitals—and government agencies—are considering limiting prescriptions of these narcotics to 7-day doses. A CDSS can help enforce the transition to these new, safer limits by reminding physicians, and/or enforcing strict constraints and reviews on physicians’ attempts to override these new limits. The governor of Florida even proposed limiting prescriptions to 3 days (Lardieri, 2017).

CDSS can also provide very valuable assistance to nurses and physicians by maintaining immediate access to large repositories of uncommon diseases, treatments, and complications. When the Ebola crisis (2014–15) occurred, CDSS which were configured to ask and analyze patients’ recent international travel patterns could add Ebola to the list of potential diseases and treatments if a patient suddenly arrived at the emergency department exhibiting fever, severe headache, and stomach pain.

A separate category of CDSS in health care is found in certain medical devices such as mammography tumor screening, blood pathology analysis systems, and some implanted devices like implanted defibrillators, and a few free-standing devices like AEDs.

CDSS can also augment preventive medicine, too, by offering guidelines and advice before physiological measurements trigger disease warnings. For example, Body Mass Index or fasting blood sugar levels can trigger diet or exercise recommendations, and ongoing alcohol or tobacco use can trigger remediation program guidance.

Ultimately, CDSS can be very useful because such systems are consistent and ever-vigilant. They are not susceptible to negligence, fatigue, training lapses, data overload, and/or distractions.

## CDSS regulations

CDSS technologies represent a challenging regulatory issue for the FDA because the FDA’s regulatory authority is significantly limited to manufacturer’s published or verbal claims. Because of FDA’s limited regulatory authority, virtually every EHR vendor is careful to document that its CDSS products are intended to be used by a “learned intermediary,” i.e., a licensed medical professional or under the direction of same. The FDA—and the vendors—use this distinction to clarify that the software is not intended or approved to deliver medical diagnosis and/or therapy without the supervision and control of an appropriate licensed professional. Since the FDA does not nominally operate, inspect, or impose itself directly in a physician’s office, clinic, or hospitals, the daily use of devices is not nominally visible to it.

If a CDSS is part of a system or device that is intended to make a direct, independent diagnosis or deliver a therapeutic treatment, then the FDA categorizes it as a medical device, which then is subject to strict design, registration, safety, and quality measures. The level of liability assumed by the device and vendor is quite a bit higher, too, since it is not necessarily controlled by a “learned intermediary.”

Some products sit on the edge, or middle, of such distinctions. For example, a mammography breast tumor detection system could be used as a backstop for physician diagnosis. The benefit is the “learned intermediary” is doing the actual diagnosis, and the CDSS is double-checking his/her work. Because mammography reviewers are known to encounter periods of mistakes, they often spot-check each other’s work, which is a role that the CDSS can potentially do more effectively and consistently.

The cost and complexity of obtaining FDA approval for a CDSS-based medical device is heavily based on the potential risk to patients if the device is faulty or ineffective. If the underlying scientific and medical facts are crystal clear for the CDSS rules, and the rules have been extensively tested and approved in similar nonautomated applications, the patient risks are low, and competent human oversight is likely, the vendor may only need to document the intrinsic performance and reliability of the software and device. For example, a portable digital chest X-ray device that self-limits the radiation exposure to assure high-quality image functions solely within very well-known laws of physics. The voltage and duration of the radiation-emitting device can be automatically controlled by a radiation sensor and safety system at a level of precision that may well exceed human measurement capability.

If, however, the device introduces novel capabilities and is meant to function relatively autonomously, like an implanted defibrillator or AED, then carefully planned and documented clinical trials may be required by the FDA before the device can be sold.

## CDSS in practice

CDSS have become very common in many clinical settings. Most of the above examples illustrate applications where the “learned intermediary,” typically a physician or nurse, operates the device in a safe and effective manner.

The largest challenge to CDSS deployment and management is the long-term maintenance, verification, and validation to assure the CDSS is functioning properly. Physicians have gotten themselves in trouble—and injured patients—when the erroneous customization of their order sets prevented the CDSS from safe operation. Because many departments, specialists, and clinical situations require customized rules, and because the rule systems (whether ES or ML) are complex, large, and difficult to assess, verification that they are operating as intended, and/or validation that their functions are indeed safe and optimal for current medical care is a growing challenge.

Take, as one example, the new hypertension rules, which define a lower blood pressure for hypertension diagnosis and treatment. In principle, every physician, hospital, medical device, electronic medical record system, and pharmacy information system may need every affected CDSS system to be updated, verified, and validates. Though cumbersome for an ES-based system, the method of retraining an ML-based CDSS may be much more complex, because the system may need to “unlearn” prior examples, or it may need to be retrained or tweaked in some manner.

An emerging class of smartphone- and app-based medical devices are becoming more common is devices that are nominally patient-controlled, like self-administered blood-glucose monitoring and insulin injection. The FDA has approved a growing number of “blood glucose management systems” that transfer the data from the blood glucose monitor to a physician or nurse management team, allowing earlier detection of problems and intervention prior to expensive and dangerous emergency crises.

The entire infrastructure of physician practice electronic medical records (EMRs), hospital electronic health record (EHRs), regional hyperimmunoglobulin E syndrome (HIEs), and medical devices with bi-directional data communication capabilities may eventually open the way to more closed-loop patient health management systems. One only must look at the apparent rapid pace of self-driving car research and on-road testing to see that society is becoming more tolerant and interested in expanding its trust of technology to protect human life. Only time will tell if, when, and how quickly such innovations will make their way into health care.

Until then, practitioners can expect to see more and more “augmentive” CDSS technologies that are intended to empower clinicians, patients, and caregivers.

## Intersection of AI and big data

Big Data is a term or concept that tries to capture several related, but very different, aspects of clinical information.

1. In one dimension of Big Data, we can consider what is being described as the “quantified self.” In other words, each person is the source of a large and growing accumulation of health-related data that spans, literally, from cradle to grave. An extraordinarily healthy person’s data may largely consist of vital signs and limited laboratory tests from annual check-ups, while the unfortunate patient born with multiple genetic defects and diseases may generate a veritable private encyclopedia of observations, test results, experiments, and outcomes. This patient’s health records may be influenced by many personally unique environment factors, too, including education, financial situation, nutrition, medication, climate, culture, and more.
2. In another dimension, we are gathering incredible volumes of data at the very smallest molecular biological level. This data is coming from fields that we label genomics (the DNA of the person), proteomics (the molecular effluent of tumors), epigenomics (temporary and/or permanent modifications of the DNA caused by diseases, injuries, stress, pollution, or other effects).
3. In yet another dimension, we are gathering local and regional community health data in electronic and/or PHR, which can be used to discern emergence or arrival of diseases like Ebola, anthrax, influenza, and/or environmental effects like lead in the water supply.
4. In a larger dimension, we can obtain, aggregate, and analyze national and global trends such as obesity, diabetes, heart disease, and also environmental data including increasing baseline outside temperatures, expansion of brackish water intrusion into water and food supplies, and/or antibiotics, steroids, and illegal drugs in the water supply.
5. Plus, we cannot forget the temporal/time dimension, which is allowing us to accumulate data across the spectrum of the fine-grained moment to moment heartbeats through weekly, monthly, and annual sampling of all the above data; and
6. In a composite dimension, much of the above data is now available from mobile sources, including wearable sensors and monitors that can incorporate not only human physiologic data, tagged with geospatial locations and synchronized microclimatic and local pollutant measurements.
7. To all the above descriptive and numeric data, we must now add multimedia formats (pictures, video, and

sound) that are accumulating everywhere, from smartphones, pervasive surveillance cameras, dental computed tomography (CT) images, hospital images from virtually all specialties, and more.

In a very real sense, human health data is now available as if we were zooming in from a satellite in layers from the overall population to communities, the individual, and the individual's molecules.

This abundance of “Big Data” is creating opportunities that never existed before, but the volume of data has outstripped manually guided data analysis and display. Even 10 years ago, the CDC or the World Health Organization (WHO) were limited to relatively static maps and map layers that were able to display retrospective data that was rarely more current than a few weeks in the past.

Thanks to higher-density AI and graphic display tools developed for military net-centric warfare and commercial logistics applications (e.g., [Amazon.com](https://www.amazon.com)), we are beginning to see very innovative applications emerge for health care.

### Challenges of data acquisition

Certainly, we cannot overlook one of the axioms of all data and information analytics: GIGO. If we do not control the quality and consistency of data that is being captured, stored, and analyzed, we cannot add precision or quality afterwards without running big risks of data distortion.

PHR illustrate one sort of risk. In principle, mHealth device companies have been arguing that capturing every nuanced detail of a person's health record can assist a physician in more precise diagnosis or detection of problems. If only, the argument goes, we could track a diabetic patient's every blood sugar measurement, every episode of exercise, every hour of restful sleep, and every detail of nourishment, then we might be able to pinpoint risks and problems precisely. However, we also know that patients get very resentful of being overmonitored and managed—and dealing with very unpleasant side effects—they forget details, and, not surprisingly, sometimes they simply mix in white lies—errors of omission and commission—to help them regain control and privacy over their lives. Also, patients go through cycles of grief, including anger and denial, which affects their willingness and ability to comply with rigid regimens of care.

On top of these human foibles, the personal health monitoring devices that patient can buy from Amazon, eBay, or Walmart are not necessarily subject to routine calibration, which can introduce even more inaccuracies into the data stream.

Also, clinical providers have only a finite—and dwindling—amount of time, energy, and focus they can devote to reviewing patient data. The average physician is therefore

not enthusiastically encouraging pouring moment-by-moment PHR data into their carefully and scientifically controlled EHR.

If we then add the legal liability that providers face if they make a mistake, we can quickly understand why many view PHR data as a risky and burdensome source of erroneous data, and the quantity of PHR data may be orders of magnitude more voluminous than the trusted EHR data, too.

Researchers must also consider that there may be two or more Data Acquisition challenges: (1) Source data to teach the Machine Learning (ML) AI tools may be incomplete or biased, and (2) System users may introduce their own distortions or biases.

Acquisition and analysis of source data are the way an ML system learns, or is taught, “Truth,” or “best practices.” An ML AI system could identify disease and case management practices from (EMR). If studying congestive heart failure (CHF) patients, for example, an ML system would likely notice that CHF patients often have second or third disease diagnoses, and that most of the patients would be put on similar low-sodium diets and, depending on coexisting diseases, certain patterns of medications would emerge. The electronic records would also likely indicate that CHF patients wind up being transported to hospitals when they start showing one or more significant clinical distress signs, including shortness of breath, high and/or erratic heart rate, swollen joints and difficulty moving, and confusion caused by lower blood oxygen levels. Many of these patients would also show significant water-weight gain upon arrival at emergency departments, and they are likely to receive strong doses of Lasix, a medication that stimulates the kidneys to accelerate urination to eliminate the accumulated water.

Consider, though, the inferences that an ML might make if the patient was part of an active home-care regimen in which food, exercise, and medication were constantly adjusted to minimize these emergency situations.

Consider, too, a third ML system that tries to understand patterns in a large rural population of unemployed, uninsured, and uneducated patients who have little access to primary care and are forced to rely on ER treatment.

If three different ML systems are trained with these disparate sets of data, each ML system will “correctly” infer the common, and perhaps most successful, clinical care pathways for each situation, but none of the three will necessarily have suitable information for each other's circumstances and patient population.

This is, in many ways, a computerized version of the parable about the four wise men and the elephant, in which each individually describes an ear, a trunk, a tail, or a leg as their “truth” about the nature of an elephant.

The user may misapply an ML system too, by asking the “wrong questions,” and the user may also misunderstand the ML's guidance if it is provided without the necessary

context. For example, a caregiver in a rural setting may try to follow the best practice of sending a patient to their primary care physician for follow-up and medication management, or by ordering state-of-the-art home care support. But if neither of those options exists for the patient, due to the patient's circumstances, the guidance will simply be ignored.

- Benefits of Big Data and AI
- Challenges of Processing Big Data using AI
- Current Technology Approaches

Despite the limitations mentioned above, AI is beginning to make contributions for healthcare improvement using Big Data healthcare sets. AI tools are tireless, and they are scalable across thousands of computers, too. AI tools can be used to separate sets of data on an as-needed basis, and multiple sets of AI tools can be deployed to identify either anomalous or simultaneous patterns and display them to human operators for interpretation and action.

One interesting application being developed and deployed at Stanford is using NLP tools to analyze retrospective pharmaceutical drug research and clinical records to identify new drug application patterns and opportunities. In simple terms, if Drug A and B successfully treat Diseases 1 and 2, and Disease 3 is much like Diseases 1 and 2, Drugs A and B might not be beneficial for Disease 3? What if there are some small numbers of cases where, indeed, Disease 3 was improved by Drug A and/or Drug B, but those had never been formally part of a clinical trial because the cases and research happened on different continents and was reported in different languages?

Another very interesting and actively discussed body of research is emerging from IBM's Watson Health programs. IBM has partnered with many clinical centers of excellence to explore ML AI systems trained by "the experts." Clinical partners have been chosen in fields like cancer (Memorial Sloan Kettering), cardiology (Cleveland Clinic), and primary care medicine (Kaiser Permanente and Mayo Clinic). In the early stages, such tools are expected to help each group of experts learn more about their own practices and to possibly help discover new opportunities they have not yet noticed independently. Watson Health's tools can also assist those expert teams to retrospective research studies by looking more carefully for patterns in past patient care that may have been overlooked.

Watson health also represents a possible business opportunity for IBM and each clinical partner, because each trained and validated ML system could be sold or rented to other clinical users who want to leverage the embedded expert knowledge. A challenge, of course, is the matching of the new client's patients, environment, and documentation standards to the particular Watson Health model they wish to use. If, for example, a national health system would like to implement Cleveland Clinic's ML model,

but has not integrated, well-developed, or roughly similar hospital, home-care, or emergency care programs, the client may not find the ML system initially valuable or use of the model may initially increase cost and workload for the nation.

In time, however, such experiments have the potential of training growing numbers of hospitals, health systems, and even national healthcare programs how to replicate or customize Cleveland Clinics expertise to improve the safety, effectiveness, and cost of prior programs.

A latent challenge in Big Data AI systems is discovering and implementing systems to retrain the knowledge tools to incorporate new best practices. For example, in 2017 the American College of Cardiology and the American Heart Association revised the guidelines for high blood pressure diagnosis and treatment. Patients who previously had blood pressures of 130/80 were considered "OK," but under the new guidelines, such patients should be diagnosed and treated for high blood pressure (Perry, 2017)!

The pattern of significant periodic revision of scientific and medical knowledge, and best practices has a long history in health care. In fact, the latency in healthcare adoption of new, better practices has been quite long (e.g., it took many, many decades for penicillin to be recognized as the standard of care to treat infections.) When new tools are developed to retrain ML AI systems, perhaps the delays can be significantly reduced, so that medical care and outcomes can be optimized more quickly.

In addition, the growing availability of Big Data from electronic medical record-keeping is allowing retrospective review of patient care patterns to more readily discern benefits of potential paradigm changes like reducing the high blood pressure limit for diagnosis and treatment. Large population retrospective analysis can identify clusters of potential improvements that bear more careful research and challenge prior assumptions.

## Potential benefits/challenges of AI

As with many technologies, there is a two-edged sword effect. Beneficial features are often offset by unanticipated side effects, unintended consequences, and/or "emergent behaviors." An emergent behavior is often observed when people interact with technologies in unexpected ways. Cell phones, for instance, provided unheralded mobility, but texting while driving was/is a dangerous emergent behavior that was not planned or anticipated.

For providers, future AI/ML tools could help support or protect patients who may need assistance before emergency conditions occur. For example, patients who forget to take their medicine, or who take them at the wrong time or in undesirable combinations, could be prioritized for nurse visits or calls, or for counseling and education during routine visits (Hsu, 2016; Sloane, 2016).

Similarly, patients with CHF can be given an electronic “safety net” to enable safer, more affordable, and potentially better quality of life using remote AI/ML monitoring systems (Guidi, 2015; Guidi, 2016). Also, in 2017, the first prescription-only app was approved by the US FDA (Software as a Medical Device) to support patients with substance abuse management (Waltz, 2017). This is an interesting app. Lastly, another 2017 FDA approval was for a medication/device combination that allows monitoring when a medication has been ingested (U.S. Food and Drug Administration, 2017).

Such innovations may have much higher adoption potential in the future, as US Accountable Care Organization (ACO) reimbursement models provide financial incentives for improved patient outcomes and satisfaction. To the extent that such tools allow patients’ diseases to be successfully remotely monitored and managed outside of the expensive hospital environment, providers may have significant motivation to adopt and deploy them. Also, if patient-focused compliance incentives or penalties catch on, patients may actively seek them out.

There is a hidden challenge in each of these technologies, including matching appropriate physicians, patients, and devices to assure beneficial outcomes without adding inadvertent risks and costs. For all mobile monitoring and support systems, for example, there is threshold infrastructure expectation that must be met, including reliable and secure communication and cloud resource availability. Remote rural areas, or highly congested urban areas, may introduce unacceptable risks to patient care.

Emergent behaviors by patients—or caregivers—may harm the success, safety, or reliability of a remote app to support substance abuse app or medication ingestion monitoring system. Patients who wish to fool the system, or other honesty or cognition limitations of the patient, could harm, delay, or derail widespread adoption and deployment. If the patient experiments with ways to fool such systems, by either lying to the app or feeding the self-reporting pill to a pet, the intended systemic benefits may fail.

Provider incentive misalignment can provoke unintended consequences, too. In an ACO, for example, the community of providers must share the fixed per-patient reimbursement. This, in turn, can result in overloading lower-cost home-care providers with supervisory responsibilities they cannot or will not fulfill correctly. To make matters worse, many current patient discharge quality monitoring schemes focus on eliminating re-admission or -treatment within 30 days. A dishonest or incompetent provider partner could inappropriately hold patient care off until the 30-day period has expired despite patient discomfort or deterioration (Wadhera et al., 2018).

Possibly the most widely heralded potential future benefits fall under the rhetoric of “precision medicine,” or medical care that can optimally tease out or fine-tune the

best pathways of care based on unique personal differences. Each patient often has unique side effects and limitations of benefits for every available medicine, and for every combination medicine. The differences could be related to genetic make-up but may also be related to novel allergies and the interaction of multiple simultaneous illnesses. There is also a growing body of knowledge about human chronobiologic variations of drug responses throughout the 24-h day (McEachron, 2012). In theory, at least, an AI/ML-based patient care optimization tool could help track and identify personally optimized treatments offering the most cost-, quality-, risk-, and efficacy-focused care (De Sonis, 2017).

### Potential benefits/challenges for payers

In a growing number of situations, payers are being pressed to manage the costs and outcomes for every-larger patient populations based on fixed per-patient revenues. Broadly speaking the aggregation of risk and costs across large patient pools would be expected to allow cost averaging. In other words, though some patients are more expensive, others are less expensive, so the costs and savings might balance each other out. However, large-scale population trends like aging, obesity, and diabetes can move the average costs higher year by year. In addition, many older patients develop multiple simultaneous diseases (multimorbidity) as they age, which can add significant complexity, risk, and cost to the overall population.

On the one hand, AI- and ML-enhanced technologies hold the promise of improving—or at least managing—growing patient care complexity and cost. By 2050, however, more than half of the population will allegedly be older than 65, and continued life-sustaining success is expected to shift the average US age higher and higher.

The payers are in the unenviable role of trying to drive down short-term costs without simultaneously driving up long-term costs and risks. There is no historic precursor data to draw on to predict the ultimate costs and complexity of rapidly aging communities, though, so much future effort may necessarily be reactive instead of proactive.

An emergent behavior by patients is being witnessed by payers, too: healthier, more active patients are actively pursuing activities like adventure vacations and active sports, both of which can increase overall rates of injuries, complications, and costs!

### Potential benefits/challenges for consumers/patients

Consumers, patients, and family/community caregivers are faced with ever-growing out-of-pocket expenses. Rising insurance rates, changes in deductibles, copays, and pharmaceutical benefits restrictions are among the causes, along with the ever-increasing age and multimorbidity of

patients. Consumers and patients find themselves deferring medical visits, using over-the-shelf alternatives, and simply going without. AI/ML-based clinical decision tools could benefit the consumer by helping clarify their needs and assess treatment alternatives more effectively. Though there are many fact-based online sources available, like PubMed, consumer sites provided by leading health centers like Cleveland Clinic, Johns Hopkins, Mayo Clinic, and others, and advertising-funded sites like WebMD and [RxList.com](http://RxList.com), there is no effective way to integrate the information and/or match it to individual needs.

## Future of AI

Given the progress of autonomous devices like industrial robots, airplanes, drones, and self-driving cars, it would seem likely that more autonomous diagnostic and therapeutic measurements, products, and services are likely to emerge. Some could be simple, such as self-service eyeglass or immunization screening and dispensing, drone and/or robotic delivery of medicines and med/surg supplies like needles and bandages, and self-managed clinical care ([Waltz, 2017](#)).

However, AI observers should always be aware of the classic hierarchy of data, which progresses from raw data to information to knowledge, and then wisdom. Modern AI seems to be encroaching more and more into areas of “knowledge” by leverage ever-increasing computing and memory capacity. Imparting or deriving wisdom into or from AI, however, lies beyond the horizon for now. For example, an AI-based babysitter may be able to “see” a child putting a fork or spoon in its mouth, and “know” that is safe. It may even “know” that a pocket knife does not belong in a child’s hands, let alone near its mouth or eyes. A wise parent, however, will anticipate that a child may try to put virtually anything into its mouth, including toxic chemicals, nearby animals, or virtually any mobile object, and will add a razor, broken piece of pottery, or nearby worm into the threat category without delay. That same wise parent will apply the same rules to a grandparent with dementia without much prompting, too. Current AI programs are hampered by several constraints, including limitation of its learning to past events and rules, dependence on direct success and failure experiences, and an absence of “awareness.” When humans say, “I could see that child was in danger,” or even more abstractly, “I could see that child was upset,” we include a massive amount of personal and collective knowledge into our wisdom, judgment, and awareness. To date, AI has not demonstrated that capacity, nor the senses of compassion and fear that are part of our biologic heritage.

Adding to the philosophical and scientific debate about potential AI risks is the concern that computed AI could erroneously identify a person as a risk to him/herself or the system, and constrain or harm a person in destructive or dangerous ways. Taking that electronic babysitter, for

example, a “perfect AI babysitter” might reasonably put a child in a box or cage to keep it from harming itself or the environment. Although that action might be logical and even at some level justified, human parents and caretakers, know that children learn from falling, scraping their knees, bumping their heads, cutting their fingers, and numerous other mistakes. Teenagers and adults, unfortunately, sometimes must learn from much, much more disastrous mistakes. Creativity and innovation come along with significant measures of risk. Great leaders and thinkers like Elon Musk and Stephen Hawking have spent a good deal of time and energy analyzing and presenting such concerns about AI’s risks and limitations to modern scientists and policymakers, hopefully opening the dialog to better future AI.

Another considerable challenge to AI is that medical, scientific, technical, and engineering knowledge do not remain constant. In fact, facts and knowledge change as time, experience, and circumstances evolve. At one time, for example, antibiotics were mixed in soaps and product surfaces to eradicate germs, but we have come to realize that excessive or incomplete use of antibiotics has given rise to dangerous generations of antibiotic-resistant germs. Even rules of thumb, like the IV Therapy’s “Five Rights” may inevitably be reconsidered as we realize we have overlooked, or might improve outcomes if we add more rights ([Federico, 2015](#))!

Perhaps last, but certainly not least, there are significant ethical and societal issues being exposed in the development and deployment of AI systems in all applications (health care included). One of the simplest medical ethics issue is the right of the patient to choose to accept, decline, or cooperate with treatment. How, for example, should an AI healthcare system deal with a patient with liver disease who cannot/will not stop drinking alcohol or injecting illegal narcotics? Or consider the conundrum that faces the self-driving car industry: no-win choices and decisions provoked by circumstances. For example, if a car must choose from harming a child on one side of the street, a group of adults crossing the street, a pregnant mother on the other side of the street, or risk harming the passenger by running into an oncoming bus, what is the “best” decision ([Silva, 2019](#))? Is this dilemma much different than the choices facing a late-stage cancer patient, for whom chemotherapy, radiation therapy, surgery, and/or morphine all have complex risks, consequences, costs, suffering, and ultimately death?

## References

- Bickmore, T.W., 2018. Patient and consumer safety risks when using conversational assistants for medical information: an observational study of Siri, Alexa, and Google Assistant. *J. Med. Internet Res.* 20, e11510.
- Centers for Disease Control and Prevention (CDC), Centers for Medicare & Medicaid Services (CMS), HHS, 2003. Medicare, Medicaid, and CLIA programs; laboratory requirements relating to quality systems and certain personnel qualifications. Final rule. Bethesda MD, Fed. Regist. 68, 3639–3714.

- Centers for Disease Control and Prevention, National Center for Immunization and Respiratory Diseases (NCIRD), 2019. 2018-2019 Influenza season week 2 ending January 12, 2019. In: Weekly U.S. Influenza Surveillance Report. Retrieved from, <https://www.cdc.gov/flu/weekly/index.htm>.
- Cunningham, A., 2017. New Blood Pressure Guidelines Put Half of U.S. Adults in Unhealthy Range. Science News. Retrieved from, <https://www.sciencenews.org/article/new-blood-pressure-guidelines>.
- De Sonis, A.M., 2017. Pain chronobiology and epigenetic regulation nutrition, nutraceuticals, acupuncture. *J. Clin. Epigenet.* 36, 1–3.
- Devine, E.G., 2000. Comparative evaluation of three continuous speech recognition software packages in the generation of medical reports. *J. Am. Med. Inform. Assoc.* 7, 462–468.
- Ehrmeyer, S.S., 2004. Has compliance with CLIA requirements really improved quality in US clinical laboratories? *Clin. Chim. Acta* 346, 37–43.
- Federico, F., 2015. The Five Rights of Medication Administration. Institute for Healthcare Improvement. Retrieved from, <http://www.ihp.org/resources/Pages/ImprovementStories/FiveRightsofMedicationAdministration.aspx>.
- Francis West, J., 1979. USA Patent No. US4165704A.
- Giuliano, K.K.-T., 2018. Intravenous smart pump drug library compliance: a descriptive study of 44 hospitals. *J. Patient Saf.* 14, e76–e82.
- Guidi, G.P., 2015. A multi-layer monitoring system for clinical management of Congestive Heart Failure. *BMC Med. Inform. Decis. Mak.* 15, S5.
- Guidi, G.R., 2016. Case study: IBM Watson analytics cloud platform as analytics-as-a-service system for heart failure early detection. *Future Internet* 8, 32.
- Henning, K.J., 2004. What is syndromic surveillance? *Morb. Mortal. Wkly. Rep.* 53 (Suppl.), 5–11. <https://www.cdc.gov/mmwr/preview/mmwrhtml/su5301a3.htm>.
- Hsu, W.L., 2016. Utilization of a cloud-based diabetes management program for insulin initiation and titration enables collaborative decision making between healthcare providers and patients. *Diabetes Technol. Ther.* 18, 59–67.
- Jiang, F.J., 2017. Artificial intelligence in healthcare: past, present and future. *Stroke Vasc. Neurol.* 2 (4), 230–243. <https://doi.org/10.1136/svn-2017-000101>.
- Lardieri, A., 2017. Fla. Gov. Announces Proposed 3-Day Limit on Opioid Prescriptions. U.S. News. Retrieved from, <https://www.usnews.com/news/national-news/articles/2017-09-26/fl-gov-announces-3-day-limit-on-opioid-prescriptions>.
- Malin, J.L., 2013. Envisioning Watson as a rapid-learning system for oncology. *J. Oncol. Pract.* 9, 155–157.
- McEachron, D., 2012. Chronobioengineering: Introduction to Biological Rhythms With Applications. vol. 1 Morgan & Claypool Publishers, Williston, VT.
- Perry, J., 2017. New Guidelines for High Blood Pressure From AHA/ACC Will Effect Millions In U.S. Medicine Matters. Retrieved from, <https://news.med.virginia.edu/medicinematters/uvasomrestriction/public/2017/11/17/new-guidelines-for-high-blood-pressure-released-by-aha-and-acc/>.
- Podaima, B.W., 2018. A review of emerging smart RFID in healthcare. In: CMBES Proceedings.
- Shah, R.A., 2018. IOT and AI in healthcare: a systematic literature review. *Issues Inf. Syst.* 19, 33–41.
- Silva, R.A., 2019. Ethical implications of social internet of vehicles systems. *IEEE Internet Things J.* 6, 517–531.
- Sloane, E.B., 2016. Cloud-based diabetes coaching platform for diabetes management. In: 2016 IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI). IEEE-EMBS, Las Vegas, pp. 610–611.
- ter Voert, E.E., 2018. Quantitative performance and optimal regularization parameter in block sequential regularized expectation maximization reconstructions in clinical 68 Ga-PSMA PET/MR. *EJNMMI Res.* 8, 70.
- U.S. Food and Drug Administration, 2017. FDA Approves Pill With Sensor That Digitally Tracks If Patients Have Ingested Their Medication. FDA News Release. Retrieved from, <https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm584933.htm>.
- Vongpatanasin, W.C., 2018. Diagnostic thresholds for blood pressure measured at home in the context of the 2017 hypertension guideline: analysis from 2 US cohorts. *Hypertension* 72, 1312–1319.
- Wadhwa, R.K., Joynt Maddox, K.E., Wasfy, J.H., et al., 2018. Association of the hospital readmissions reduction program with mortality among medicare beneficiaries hospitalized for heart failure, acute myocardial infarction, and pneumonia. *JAMA* 320 (24), 2542–2552.
- Waltz, E., 2017. First Prescription App for Substance Abuse Approved by FDA. IEEE Spectrum. Retrieved from, <https://spectrum.ieee.org/the-human-os/biomedical/devices/first-prescription-app-for-substance-abuse-approved-by-fda>.
- Weiss, S.M., 1990. An empirical comparison of pattern recognition, neural nets and machine learning classification methods. In: *Readings in Machine Learning*, pp. 177–183. <https://www.amazon.com/Readings-Machine-Learning-Morgan-Kaufmann/dp/1558601430>.
- Widrow, B., 1987. The original adaptive neural net broom balancer. In: *Proceedings of the IEEE Symposium on Circuits and Systems*, pp. 351–357.
- World Health Organization, 2017. Pacific Syndromic Surveillance Report. Week 35, Ending 4 September 2017. Retrieved from, [http://www.wpro.who.int/southpacific/programmes/communicable\\_diseases/disease\\_surveillance\\_response/PSS-4-September-2017/en/](http://www.wpro.who.int/southpacific/programmes/communicable_diseases/disease_surveillance_response/PSS-4-September-2017/en/).