From “If-Then” to Artificial Intelligence…

18 Uses of AI for population Health

June 2018

AGATE
CONSULTING

What you will find in this report…

Executive Summary

Despite several high-profile fits and starts, Artificial Intelligence (AI) represents one of the most promising technologies to improve population health. Last year, health executives identified AI as the most potentially disruptive technology to the industry. Goldman Sachs estimates that the technology can eliminate $28bn annually in payer and provider inefficiencies by 2025. Accenture is even more optimistic. Even if some of this hype is potentially excessive, AI will be a meaningful component of all population health improvement efforts going forward.

AI refers to computers having the ability to perform cognitive functions typically associated with humans. This includes perceiving, reasoning, learning, and problem-solving. Machine Learning refers to the computational techniques that enable AI. AI is profoundly different than the standard computational approach in making predictions and performing tasks. Traditionally, a computer program would use multiple “if-then” rules to assess a system and decide what to do. In this traditional method, the computer is not learning. Rather it is automating the knowledge of its human computer programmers. While useful, there are many inherent limitations to such a deterministic method. AI addresses many of these. Like humans, AI-based algorithms can use unstructured, incomplete, imprecise, and/or inconsistent data to make predictions about a system. These predictions, like human thought, have various levels of certainty in their accuracy. As the algorithms are trained on more and more data, their sensitivity and specificity can be improved. Hence why “big data” sets are so intertwined with AI.

Understanding the complexities of the many types of Machine Learning algorithms is, at least today, the job of the data scientists. However, all health executives charged with improving population health should at least be familiar with the potential of AI. To help with this, this report outlines many AI-based capabilities that population managers could explore, and even some that they may unwittingly already be using. This does not mean to say that all health systems need to do all these things. Rather, the list intends to shed light on this important technology while spurring creativity.

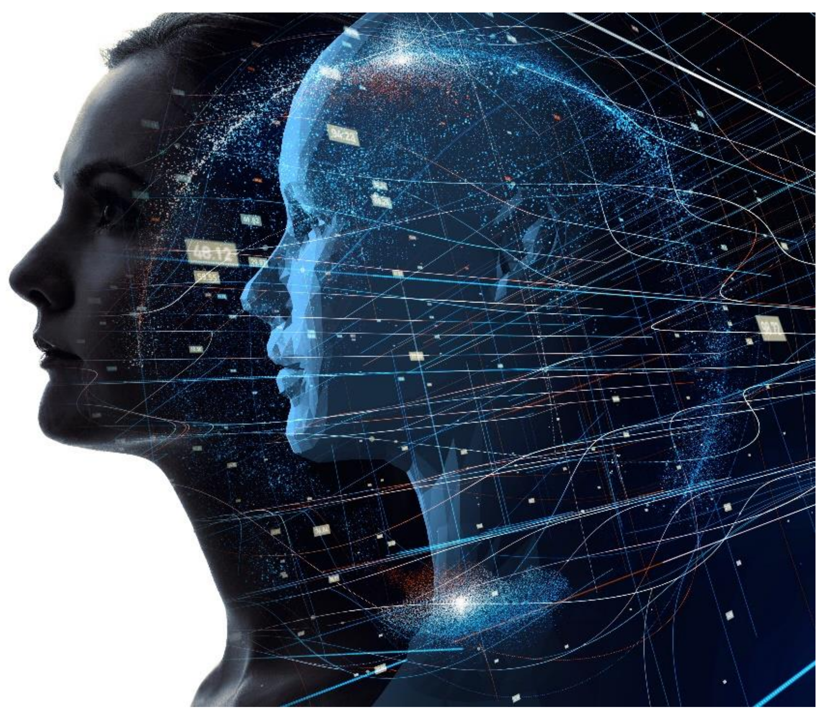

Why AI is needed in population health

Recent announcements by leading AI vendors have caused a some healthcare executives to fear AI is another passing fad.[ii] Others executives are still just trying to understand what exactly AI means to healthcare. While the current woes of some leading AI-players serve as a valuable cautionary tale about the challenges of deploying new health technologies, AI is already a feature of the healthcare IT landscape. It will grow dramatically. It must. Traditional deterministic "if-then" based analysis already fails to meet all of the health industry's needs. Its shortcomings will only get worse over time.

Fortunately, incumbent organizations do not need to reinvent all of healthcare today to leverage AI. In fact, as with other health technologies, unbridled enthusiasm around AI is likely counterproductive.[iii] Rather, organizations that start their AI journey to address tactical, solvable, high-value problems can get immediate results, while building a foundation for future expansion. This report lists several AI-based capabilities for health systems to consider experimenting with. While none of these may lead to a transformation in cancer care on day one, or the rapid creation of a full-fledged robotic doctor, they may help improve care quality, safety, cost-efficiency, and stakeholder satisfaction.

AI by the numbers

Last year, healthcare executives rated AI as the industry's most potentially disruptive technology. It beat out the Internet of Things (IoT), 3D Printing, Robotics, Augmented Reality, and the often-hyped Blockchain.[iv] Goldman Sachs called AI "the apex technology of the information age," and "a needle-moving technology for the global economy."[v] In healthcare alone, Goldman estimates that AI can create $54bn in annual cost savings by 2025. This savings is split almost evenly between lower costs to discover new medications ($26bn) and driving efficiencies in healthcare delivery ($28bn).[vi] Accenture's forecasts are even more optimistic. They project that AI applications will create $150bn in annual savings for US healthcare alone by 2026.[vii] Even if current assessments of AI's potential are somewhat overstated, the opportunity is too large to be ignored by all healthcare incumbents—large or small, tech-savvy or naïve, data-rich or poor, and even wealthy or cash-strapped. Once demystified, much of the AI toolkit can be embraced by almost all healthcare organizations. Due to the given complexity of healthcare and the failures of traditional analytics, it will need to be.

A working definition of AI

AI can be defined as the ability of a machine to perform cognitive functions typically associated with human minds. These include perceiving, reasoning, learning, and problem-solving. Tactically, this means having computer algorithms draw probabilistic insights from unstructured and imperfect information. This is just what humans do every day.

Machine learning refers to the computation techniques that achieve AI. Instead of hard-coding software routines with specific instructions to accomplish a task, machine learning is a way of "training" algorithms to learn how to do something on their own. To illustrate, consider a computer algorithm predicting which book an online shopper might want to buy next. One way is to have a programmer write a lot of "if-then" statements. For example, "if the person is male and bought other war histories outside traditional holiday shopping periods in the last year, suggest the new history of the Gulf War." Alternatively, a programmer could code the rule: "if the reader is female and has purchased dramatic fiction in the past, suggest the soon-to-be-released new JK Rowling's book.[viii] This approach of writing thousands of hard-coded "if-then" statements would neither scale, nor likely be that helpful. Rather, it would be better to train a computer algorithm to make recommendations using all the data it has from past sales. Such a "recommendation engine" could start with a simple observation: "people who bought books by Jane Austen on average also seem to buy books by Margret Atwood. I see you bought a copy of Austen's Emma last month, so I will suggest Atwood's A Handmaid's Tale." By making such guesses, and then getting constant feedback on whether the guess was good, the algorithms can get smarter over time. This "training" requires feeding huge amounts of data to the algorithm. Thus, "big data" sets become the foundation on which machine learning algorithms are developed. It is worth noting in the example that the algorithm did not need to know anything about the works of Jane Austen or Margret Atwood, e.g., that they both feature strong female protagonists. All the algorithm needed to know was that people who liked one seem to like the other; correlation was enough in this case. No causative reason was required. This capability of a computer learning over time versus blindly automating the already learned insights of the programmer is the power of AI.

Some basics types of machine learning

There are many different methods computers employ to learn. Three common types are supervised learning, unsupervised learning, and reinforcement learning.

Supervised machine learning. In supervised machine learning, a computer model is trying to predict a certain result based on various inputs. Such results could predict a number, e.g., the probability of getting cancer next year, or a discrete result, such as "has cancer" or "does not have cancer." If a model predicts a number, it is called a predictive regression model. If the model predicts options among fixed choices, it is a classifying predictive model. To understand regression, consider the example of a computer model that predicts future health expenses of patients using past claims data. To build such a model, a computer algorithm is fed a training data set. This includes total patient medical expenses—the outcome being sought—along with various patient features, e.g., age, chronic illnesses, past use of care, and associated primary care clinician. Using a variety of mathematical techniques, the model identifies which factors predict costs to what degree. This learning process is called "training." The model can then be fed a new data set to make predictions on future costs. This model is considered supervised because the output variable during the training, the total medical expense, was known to the model developer, as were the potential patient features that could influence it. Classification is like regression, with the exception that the output is not a continuous variable (e.g., next year's spending) but multiple choice. For example, a model that takes in several different physiological parameters to assess if a hospitalized patient is septic or not.

Unsupervised machine learning. In the past two supervised learning examples, the model developer knew the correct answer in the training data. In the first, each patient's total medical expenses were known. In the second, it was known if the patient was septic or not. In unsupervised learning, the model builder does not know what the "right" answer is, even in the training data. For example, imagine a health-system is trying to understand how various groups engage in their care. To do this, a data set is compiled with demographic information, office visit use, medication compliance experience, ED visits, hospital admission data, social history, health system, web portal use, etc. Mathematical techniques (e.g., K-means clustering, hierarchical clustering) are then used to identify similar patient groups. This analysis may yield a cluster of high cost, young, seemingly very engaged chronically ill patients, and other groups of recently operated-on patients on certain medications. Such a grouping of similar patients, called "clusters," can then help drive additional research on patient engagement. This type of clustering analysis can also be used to group newly published medical research into common themes. It can also analyze vast genomic data sets to find similar patients for better research targeting. What all these examples have in common is that the model builder does not know in advance what clusters will emerge. The algorithms automatically organize the data.

Reinforcement learning. Reinforcement learning can be thought of as a hybrid of supervised and unsupervised learning. An algorithm may start unsupervised, making predictions on whether something is true or not based on several variables; e.g., patients with these conditions tend to do better on this drug vs. that. The model's prediction is then compared to real-world data. When the prediction is correct, the algorithm is reinforced. When it is incorrect, it is penalized, and the model may change its weighting of various parameters to try and do better. Over time, with timely data feedback, the model accuracy continually improves. The book recommendation engine discussed above is an example of reinforcement learning.

Even within these three types of machine learning, there are many different sub-methods and mathematical techniques. These details are the realm of the data scientists and clinical informaticians. However, executives focused on improving population health will need at least a basic understanding of the power of machine learning. To help in this journey, below are some examples of how the technology can be applied.

18 AI-Based Population Health Management Capabilities

AI will be key to helping health systems and payers achieve the overarching goal of population health – have each patient get the right care, at the right time, from the right provider, across the care continuum. The examples below are foundational AI-based tools that population health managers can employ today to varying degrees. They run the gamut of machine learning types described above.

1. Patient and provider identity matching. Population health data analysis first and foremost requires matching the same patient or provider across all data sets. Consider all the data on a sample patient Jennifer Smith. Her records span across many systems and many years. Pulling this all together, which seems trivial to laymen, can be a major undertaking. Even within one IT system, patients and providers may not be consistently represented. Patient Smith could sometimes be "Jennifer," sometimes "Jen," or "Jenn." Of course, she was Jennifer Jones until 2014, when she got married. While most systems have her correct birthdate of 7/3/1975, some flipped it to 3/7/1975. Most of her providers use a common Medical Record Number, but some do not. Her member ID from her insurance claims is entirely different. Basically, the data is messy. Probabilistic identity matching models based on AI are commonly employed to make a "best guess" on whether two patients or providers refer to the same entity. Hard coded "if this field and that field and that field are all the same, then the two patients are the same" logic would not get it done. Such a probabilistic approach is already quite common.[ix] However, much additional work is required. 86% of clinicians and health IT professions surveyed in 2016 by Ponemon Institute® said they knew of a medical error that was a result of patient misidentification.[x] Better algorithms are needed.

2. Patient record deidentification. Another key population health data management capability is deidentifying patient information. Deidentifying data protect patient privacy and allows information to be more freely shared with other users and systems. It is done on claims and other structured data sets as well as unstructured textual documents. Deidentifying structured information is somewhat straightforward, e.g., delete patient names, month & day of birth, address fields. Deidentification is much harder to do on medical notes. Natural Language Processing (NLP) is the form of AI where machines learn to understand human language. In NLP, the computer makes a prediction of what the text means. For example, a computer can read the sentence in a medical record, "Mrs. Turner has Turner's Syndrome," and replace it automatically with "Patient XYZ has Turner's Syndrome." Hence, the new data can be manipulated without knowing the identity of Patient XYZ, and then reidentified at the end of the process as needed. As an aside, the author tried this very example at the Google Booth at HIMSS 2018. Google was demo-ing their patient deidentification API. The demo failed, returning "[Patient 1] has [Patient 1]'s Syndrome." No doubt Google can do this analysis properly, the demo was just not fully set up. However, it was a good lesson. Just because something is said to have "Artificial Intelligence," do not assume that it is truly "intelligent" out-of-the-box.

3. Health event rate prediction. Population health managers often need to know the probability of a specific health event occurring in a specific patient within a specific time-frame. Some of these tools are already employed today in clinical practice to help decide optimal treatment. For example, the Framingham Index® computes the probability of a given patient having an MI or death in the next 10 years. Such an overall risk assessment can influence the decision to treat or not. A more avant-garde example is The Oncotype DX® risk assessment tool, which uses tumor genomic data and other information to predict the recurrence of breast cancer. Other models predict which patients with Chronic Kidney Disease will progress to ESRD. Health event prediction could include operational variables as well. The renal care provider Fresenius® uses model to predict which renal patients will miss their dialysis appointments, including such variables as snowfall on day of treatment.[xi] There are many health event prediction models being used across all specialties today. As digitized data sets grow, models mature, and just as importantly providers and patients learn how to best use these models, the number of such models will increase.

4. Utilization/cost prediction. The ability to predict the average total consumption, either in units or dollars, of a patient population for a given set of services over a given period of time is a core component of population health management. For example, what is the total medical claims cost that this group of people will likely incur next year, or how many Emergency Department visits will another group likely make next month. Predicting total claims cost of a group is what the common predictive models, e.g., DxCG®, ACG®, or HCCs, are mostly used for. These models are also often used for patient stratification for care management as well; e.g., nurses will focus on those patients with the highest likely future expense. Further, some models like Verisk® and Milliman® bifurcate the risk into addressable or not-addressable issues. Newer, more heavily AI-based technologies, e.g., Cardinal Analytx Solutions®, markedly expand on past solutions in sensitivity, specificity, and actionability.

5. Maintaining active problem lists. Creating an active problem list for a patient is useful in treatment decision making, overall risk prediction, and resource planning. For any single patient, diagnoses are listed in many sources, e.g., claims and notes, albeit not always correctly. Several tools today distill claims data to create such problem list. Many of these tools employ deterministic (non-AI) rules like "to be considered a diabetic a patient must have either two independent outpatient claims for diabetes or one claim and be on one anti-diabetic medication." AI-based solutions can greatly improve on this, both in terms of employing other data sets, e.g., labs and notes, as well as the accuracy of the prediction with the data already captured. Companies like Apixio® have employed this approach very successfully to help Medicare Advantage plans capture under-documented diagnoses. This is critical to Medicare plans as the patient's documented disease burden is the driver of revenue for the patient.

6. Coding & documentation quality management. Accurate, complete, easily-readable, and efficiently-created medical documentation helps ensure that all patients gets the care they need, that quality metrics are accurately reported, and that payments are correct. To help improve medical documentation, here again are NLP and AI. In particular, AI is being used to power "computer assisted coding" solutions. Specifically, computer algorithms are trained to read through medical documents and already existing code sets, to: (i) suggest the appropriate codes to submit, (ii) to ensure that the codes submitted were the correct ones, (iii) to find things done that were not coded at all, and/or (iv) to ensure that the underlying documentation supports the submitted codes. The Computer Assisted Coding market, including 3M®, MModal®, Precyse®, and others, have revenues of $2.5bn in 2016, growing at 13%.[xii] AI will continue to shape this market.

7. Defining episodes-of-care. AI can be used to aggregate discrete care services into a bundle of related activities. A bundle could include everything related to a joint replacement surgery or treatment of pneumonia. For example, a bundle could include hospital costs, physician fees, medications, and post-acute services. Bundles are commonplace now—Medicare uses them for select payments under the Bundled Payments for Care Improvement (BPCI) programs. Prometheus®, the analytics effort of HCI3®, publishes definitions for 97 such bundles, from Urinary Tract Infections to Acute MI to Colorectal Cancer. AI offers the potential to expand the scope and specificity of bundled payment analytics dramatically. For example, clustering analysis could identify different sub-bundles within a broader category and/or find outliers cases. On the cutting edge, companies like Clarify Health® are using AI to both better create bundles, as well as identify high-leverage points to improve each patient's journey along their episode of care.

8. Patient safety alerting. Healthcare is dangerous. Errors of all types are commonplace.[xiii] There are errors of commission. For example, an ordered medication could adversely interact with the patient's current medication regimen or their existing medical conditions. Alternatively, there are errors of omission, e.g., patients with diabetes not started on an ACE inhibitor, surgical teams failing to administer antibiotic prophylaxis pre-op, or women over 50 years old without an order for a mammogram. Today each one of these example errors can generate (multiple) alerts in an EMR. Most of these alerts are based on limited data and do not include clinical context. More generally, they, in many situations, are all but useless.[xiv] Worse than the annoying and unhelpful popups in the EMR is the cacophony of alarms on every hospital floor. Monitoring devices, IVs, and other equipment are constantly beeping. Literally, nurses and doctors cannot find a meaningful signal amongst all the physical noise. All these alerts are based on simplistic "if this and this then that" logic. Smarter, AI-based ones, are imperative. While there are many non-technical barriers holding these up, most notably liability, they are already coming into being.

9. Symptom triaging and care routing. A key element of population health improvement is getting patients to the right type of care when needed. Today, many patients with an acute illness self-assess if they should treat the issue themselves, schedule a physician office visit, start a virtual tele-visit, make a trip to the Urgent Care Center, or call an ambulance. Other patients may use a nurse-call-line that will walk the patient through a branching-logic-based (deterministic) triaging tool. Finally, other patients will reach out to their on-call primary care doctor, who will use her intuition to assess the patient's next step. These non-computer-based algorithms and simple triaging tools can be supported with AI. The issue of care routing exists in routine care as well as acute care—for example, whether to refer a patient to a specialist. The technology company Kyruss® found that 20M referrals each year were "clinically inappropriate."[xv] Even in high-acuity care settings, triaging and routing relies, rightly and wrongly, on the art of medicine as well as the science. For example, many patients presenting in the ED with chest pain can be safely discharged after an algorithmic assessment of vitals, ECG, and negative cardiac enzymes.[xvi] While AI-based algorithms may not be ready to make definitive triaging or care-routing recommendations yet, they can help with the process.

10. Care deterioration alerting. Today, numerous technologies aggregate multiple real-time data streams on patients, regardless if the patient is in a hospital bed, at home, or any place in between. Unfortunately, interpreting this data is both time intensive and difficult for clinicians. Algorithmic surveillance with AI can greatly help. Sepsis can be detected sooner in the ICU with tools like CLEW Medical®. The Baby Scripts® platform can detect abnormal changes in blood pressure, weight, or blood sugar levels of a pregnant woman at home—prompting early treatment. The data need not be objective. Amazon®'s Alexa® can simply ask a patient at home, "how are you feeling today?" and include that data in the algorithm as well. In each of these examples, it is not simply "if data point on the patient is above X or below Y, then do Z." Rather, it is an algorithm that is constantly scoring the health of the patient to probabilistically assess risk.

11. Protocol and/or clinical trial matching. As more providers attempt to standardize care delivery, they are developing an ever-broadening set of clinical guidelines and protocols. Initially, simple deterministic rules can be used to match patients with the appropriate guideline, e.g., "if a patient presents with a C. Diff. infection, apply the hospital's C. Diff. infection control guideline," or "if a child is diagnosed with Type 1 diabetes initiate the Type 1 Diabetes new onset protocol." However, scaling this becomes quite challenging, especially for polychronic patients, for complex guidelines which leverage multiple histologic and genomic inputs such as those in oncology, and for atypical manifestations of diseases. In these cases, an AI-based probabilistic approach to matching patients with guidelines may help. A specific use of such probabilistic patient/guidelines matching is finding appropriate patients for the more than 110,000 registered clinical trials in the US alone.[xvii]

12. Understanding variations in treatments and guideline application. A cornerstone of any quality improvement program is detecting unwanted variations in treatments of similar patients for similar problems, and comparing treatment to established care standards. Today, clinical decision support tools can do this deterministically. For example, assuming the data is clean and structured, a computer can assess if the correct statin was prescribed for all patients who are (i) diabetic, (ii) between the ages of 40 to 75, (iii) have an LDL between 70 and 190 mg/dl, and (iv) have no contraindications.[xviii] Even for just enforcing this deterministic rule, AI can help structure the data. For example, AI can pull the LDL values out of a text note or find contraindications. More importantly, AI can help identify clusters/causes of guideline variations among patients and providers to help address the root causes of quality issues.

13. Defining and refining clinical guidelines. In addition to helping identify drivers of guideline non-compliance, AI can help improve and expand clinical guidelines. Using the same above example of statin use for hypercholesterolemia, there are several additional factors clinicians need to consider for lower-risk patients, including family history, likelihood to modify lifestyle, coronary artery calcification, ankle-brachial indices, and C-reactive protein levels. AI offers more sophisticated population health managers the opportunity ever to evolve their guidelines and treatment pathways while enforcing the ones they have in place.

14. Personalizing treatments and recommendations. Standardizing care delivery does not imply giving each patient the exact same treatment. Rather, it is reducing unwanted variation. It is just as important to personalize treatments to each patient's specific needs and genotype. Such a "personalized medicine" approach is common in cancer care. Here therapies can be targeted based on the patient's and tumor's genomics. AI is essential to providing such complex care. However, personalized medicine is not the exclusive purview of oncology. Companies like Nutrino Health® use AI to optimize the diets of diabetics by analyzing their food logs and Continuous Glucose Monitoring (CGM) data.

15. Medical knowledge digestion. In 1950, the universe of medical knowledge doubled every 50 years. By 2020, it will do so every 73 days.[xix] Finding up-to-date, relevant information is not easy. Fortunately, AI has already helped a lot. Search engines use AI to continually predict which piece of content the searcher will find most helpful. There is much more to do here though. For example, medical literature search can be integrated into the EMR workflow to push the most pertinent pieces of knowledge to the point-of-care when it is needed. Better periodic digests of key content can be created for each clinician—leveraging the knowledge of the physician's specific patient panel, relevant areas of research, and her unique practice patterns. Keeping up with published literature will still be hard, but AI can make it easier.

16. Visual pattern recognition and change detection. In addition to better understanding structured data and gleaning insight from human language, AI can also see. Already AI's ability to assist radiologists with image interpretation is well established. While a fully computerized virtual radiologist is far off in the future, elements of "Dr. AI" are already here. Earlier this year, the FDA approved IDx-DR®, an AI algorithm that detects diabetes retinopathy from images of the eye.[xx] Unlike past tools, IDx-DR is approved to determine more than mild diabetic retinopathy autonomously, not merely suggest it to a radiologist. Further, like humans, computers can interpret images in motion, not just static pictures. Companies like CaptureProof® use traditional cameras to assess the range of limb motion in post-op patients. Clearly, teaching computers to interpret complex radiology studies or video streams is beyond the scope of most health systems or payers. Employing the technology is not. There are also less sophisticated applications of "visual learning," including basic wound tracking, port/catheter monitoring, and even some basic dermatology screening, that can be experimented with.

17. Patient behavioral segmentation. Population health management requires as much consumer marketing capabilities (a.k.a., "patient engagement") as clinical skills. A patient's ultimate outcome is driven in large part by their social determinants of health. These include education, economic stability, physical environment, and community interaction. Understanding these social determinants of health are key to cost-effectively driving consumer action, improving patient compliance with care plans, and ensuring patients make better lifestyle choices. Furthermore, different people respond differently to different stimuli. Some may be easily swayed to change behaviors with direct messaging, others may require financial motivations. Outside healthcare, this process of figuring out who needs what, and what will make them buy, is called customer segmentation. Such segmentation outside healthcare is already heavily AI-based. Healthcare will follow.

18. Sentiment/frustration identification. Beyond segmentation and communication personalization, changing patient behavior requires empathy. Just as children learn at an early age to sense anger, frustration, fear, and joy in their parents' voices, computers can learn this too. Companies like Cogito® have algorithms that listen to conversations between nurses and patients. The tool gives the clinician feedback on how well the discussions are going and how the patient is reacting. While Cogito started using AI to help patients and clinicians better interact, they have now expanded outside healthcare as well. In general, AI offers caregivers entirely new ways to measure and improve the efficacy of provider-patient communications.

Going forward—AI is here, no ifs, ands, or buts...

AI is being used in various ways currently to improve healthcare delivery. The above list is merely a smattering to whet the appetite. This is the tip of the iceberg. At its core, healthcare is characterized by "decision making under uncertainty." Providers of every type make decisions every day with incomplete, inaccurate, and, often, contradictory information. As do patients and insurers. Informatic tools predicated on rigid, deterministic logic will never be able to fully meet their needs. The frustration with the thousands of useless EMR pop-up messages are a testament to that. AI-based analytics, combined with redesigned processes and clinician and patient training, represent a lynchpin of any healthcare organization's IT strategy going forward.

Healthcare's challenges are infinitely more complex than any board game, quiz show, or even autonomously driving car. There is 100% certainty that AI will have a meaningful role in addressing them. Healthcare is far too complicated not to have humans and AI working in concert.

Acknowledgments

This paper was inspired by the panel discussion *AI for Population Health* at the 2018 American Telemedicine Association conference in Chicago. The author is indebted to the panelists Kamal Jethwani (Partners HealthCare®), Alexandra Pelletier (Deloitte®), and Jean Drouin (Clarify Health®) for their outstanding insights. Of course, the author is solely responsible all the content herein, and takes full responsibility for any errors or omissions.

Financial & Conflict-of-Interest Disclosures

This report was funded entirely by Agate Consulting, Inc. It was conducted solely for Agate Consulting, Inc.'s internal purposes.

About the author & Agate Consulting, Inc.

This report was authored by Kevin Agatstein, Founder & CEO of Agate Consulting.

As founder of Agate Consulting, and a former member of both McKinsey & Company's Payer/Provider Practice and Arthur Andersen's Healthcare Practice, Kevin Agatstein brings clients more than 15 years of US and overseas healthcare consulting experience. With these firms, Kevin has served leading providers, payers, healthcare IT companies, life sciences organizations, and healthcare venture capital and private equity firms. His areas of expertise include clinical analytics, telehealth, population health, and a payer-provider collaboration.

Before founding Agate, Kevin designed a consumer-oriented mobile pharmacy management platform, now marketed as emWellics® by the company RxEOB®. Prior to that, Kevin led operations, corporate development, and marketing for CareKey, Inc., a leading medical management application suite employed today by some of the largest payers in the US. CareKey was sold to its long-time channel partner, The TriZetto Group, in December of 2005 for $100M.

Kevin holds a degree in Chemical Engineering from MIT. At MIT, he was also a researcher and project manager at the MIT Sloan School of Management. Additionally, Kevin completed McKinsey & Company's mini-MBA program and is a published author in strategy, marketing, and organizational design.

For more information:

Kevin Agatstein
Agate Consulting, Inc.
77 Alderwood Road
Newton, MA 02459
+1.617.901.8161
kevin@agate-inc.com

References

[i] For more information on Agate Consulting, Inc. see http://www.agate-inc.com.
[ii] Gist Healthcare. Big Blue finds success in its Watson business not so elementary. Jun 1, 2018.
[iii] Bresnick J. Success of AI in Healthcare Relies on User Trust in Data, Algorithms. Health Analytics. May 2018.
[iv] Williams B. Enabling better healthcare with artificial intelligence. PwC. Aug 28 2017.
[v] Terry HP et al. Artificial Intelligence, AI, machine learning, and data fuel the future of productivity. Goldman Sachs Equity Research, Nov 14, 2016. P. 2.
[vi] Ibid. p. 59.
[vii] Collier M et al. Artificial intelligence: Healthcare's new nervous system. Accenture. 2017.
[viii] The examples above are purely for illustration. I have included gender as an input variable as it seems to be meaningfully predictive of book preference. The Huffington Post (Nov 1, 2017) described a 40,000-person survey by Goodreads® that found profound differences in book preferences among genders.
[ix] Bresnick J. Data Integrity Strategies for Patient Matching, Identification. HealthIT Analytics.
[x] Ponemon Institute. 2016 National Patient Misidentification Report. Dec 2016.
[xi] Chan KE et. all. J Am Soc Nephrol. 2014 Nov; 25(11): 2642–2648.
[xii] Marketwatch. Computer-Assisted Coding Market 2018 Global Analysis, Opportunities and Forecast To 2023. Apr 2018.
[xiii] National Patient Safety Foundation. Free from Harm: Accelerating Patient Safety Improvement Fifteen Years after To Err Is Human. National Patient Safety Foundation, Boston, MA; 2015.
[xiv] Wachter R. The Digital Doctor. McGraw-Hill. 2015.
[xv] Pennic J. 19.7M "Clinically Inappropriate" Physician Referrals Occur Each Year. HIT Consultant. Nov 2014.
[xvi] Colwell. Which chest pain patients can go home? ACP Hospitalist. Oct 2015.
[xvii] Number of registered clinical trials in "US Only" and "Both US and non-US" as of May 31, 2018 on clinicaltrials.gov.
[xviii] ACC/AHA Release Updated Guideline on the Treatment of Blood Cholesterol to Reduce ASCVD Risk, Am Fam Physician. 2014 Aug 15;90(4):260-265.
[xix] Poorman E. Staying Current in Medicine: Advice for New Doctors. NEJM Knowledge+ Blog. Nov 2016.
[xx] Taylor NP. FDA approves diabetic retinopathy-detecting AI algorithm. Fierce Biotech, Apr 2018.

