A Value Proposition for AI-Enabled Population Health Management

Chris DeRienzo, MD, Cardinal Analytx; Michael Barr, MD, NCQA; Nigam Shah, MBBS, and Arnold Milstein, MD, Clinical Excellence Research Center and the Department of Medicine, Stanford University

Contact: Arnold Milstein, amilstein@stanford.edu

Abstract

What is the message?

This paper proposes a unifying value proposition for the application of modern data science tools to population health management. The authors believe artificial intelligence (AI) will deliver value to health plans, large employers, providers, and patients primarily in two ways – first, through linear improvements in traditional health plan functions like reserve setting and fraud prevention, and second through geometric improvements in pairing predictions with targeted screenings and upstream clinical and social-determinant interventions.

What is the evidence?

While AI may be hitting "peak hype" in healthcare, Americans are in truth surrounded by machine learning. Proof points exist across industries on the power of AI to perform markedly better at sorting tasks than either humans or baseline statistical models, and the health plan world will be no different. Furthermore, AI has proven able to redefine entire technology platforms in areas like speech and image recognition. Applied to population health management, the same power will create entirely new health platforms based on increasingly creative benefit design.

Submitted: February 22, 2019; accepted after review: April 15, 2019

Introduction

In a widely published 1994 interview with *Rolling Stone Magazine,* Steve Jobs said "Technology is nothing. What's important is that you have a faith in people, that they're basically good and smart, and if you give them tools, they'll do wonderful things with them." When it comes today's buzziest technology – artificial intelligence (AI) – people are only just beginning to figure out how to use their newest tools to transform entire industries. Though some argue we've reached peak AI hype, the authors believe significant potential remains for AI-powered solutions to transform population health management in the service of healthier people. What's lacking, however, is a common framework — a value proposition that captures the wide range of possible applications of modern data science tools from managing core health plan operations to reducing members' costs and improving clinical outcomes.

In this paper, we present a unifying construct that we believe forms the base of driving value in AI-enabled population management – namely, that AI will ***simultaneously drive a digital metamorphosis of traditional health plan functions while creating new platforms to fundamentally improve people's health.***

The machine learning metamorphosis of traditional health plan processes: Supervised and reinforcement learning

Incumbents in the health plan space today grapple with real business problems that AI can help resolve. These are extant, functional challenges that baseline technology solutions can address but will prove inferior to AI-enabled algorithms. At present, the most promising AI tools with real application to these business problems are grounded in two primary modes – supervised machine learning and reinforcement learning.

At its core, supervised learning is the combination of curated data sets containing the data (for predictions) and the answer. In these kinds of applications, machine learning models "win" by using data to predict the best answer. While this sounds like a traditional approach to data analytics, unlike conventional regression models a machine learning model is not constrained by assumptions about the best way to analyze the data or about what variables (individually or in combination) may be most important in yielding an optimized prediction.

Reinforcement learning is a little more complex as the data themselves are not curated in advance. Instead, reinforcement learning models are built using feedback to help impute "correct" or "incorrect" answers. These algorithms learn through interaction with the system. For example, if a reinforced learning speech algorithm recognizes words incorrectly, then subsequent user behavior, such as repeating the word, provides feedback that the algorithm made a mistake the first time.

Of course, supervised and reinforcement learning can also work together. Consider training your favorite speech-enabled smartphone assistant with the sounds of 100,000 people saying "tomato" (supervised learning). Then imagine setting her loose into the wild, unleashing her algorithm's existing "tomato" recognition on new consumers who provide new data to learn from and further refine her "tomato" algorithm (reinforcement learning).

These two machine learning constructs – both of which mirror aspects of human learning within a targeted subject area – are applicable to any number of traditional population health management functions. Here are just three examples.

Example 1: Automation of manual processes

As early as 2003, the United States Patent and Trademark Office reported patent applications for models that "automatically classify health insurance claims using classification models that are trained to predict whether a health insurance claim will be accepted or rejected by a target payer, analyze why the claim will be rejected, and then target the intervention(s) needed to appropriately handle the claim."[1] In late 2017, IBM and Singapore insurer NTUC Income announced a partnership that capitalizes on patents like these to automate the processing of 14,000+ pre/post hospitalization[2] claims each month. When JP Morgan announced a similar

transition to using a machine learning algorithm for commercial loan agreements it reported taking only seconds to perform tasks that historically took lawyers 360,000 hours to do by hand.[3] Such efficiency gains rival those of the industrial revolution, and promise to spread across other back office functions that traditionally demanded endless hours of human tedium.

Example 2: Actuarial reserve-setting

In 2016, the Society of Actuaries spent barely a page of its 90 page report on the *Accuracy of Claims-Based Risk Scoring Models* to machine learning.[4] Barely six months later, the same group hosted a 75 minute "Primer" webcast titled "Insurance Analytics with Machine Learning." With training steeped in the need to quickly identify and integrate emerging trends into their work, it should be no surprise that actuaries rapidly recognized the power of machine learning to drive improved individual and population-level risk analyses. The sheer mechanics of a machine learning model's statistics enables an actuary to identify hidden relationships in massive datasets without pre-specifying their importance. Like the transition from SD to HD video, these novel connections promise a less pixelated view of a population's actuarial risk, thereby also yielding more precise estimates of risk, better understandings of markets and market segments, more accurate bidding models, and optimized plan design.

Most importantly, however, is their potential impact on the work of human actuaries. Despite their power, AI-enabled risk models still need real people to interpret their outputs. In an ideal example of AI designed to augment rather than overpower human intelligence, risk-based AI-enabled solutions will allow an actuary to practice at the top of her license, focusing her limited time and resource on making her best-informed, most accurate risk assessment.

Example 3: Early-identification and prevention of fraud

Fraud prevention is a third obvious candidate for the use of an AI tool in the health plan world,[5] in much the same manner as in the credit card industry,[6] finance,[7] and the rest of insurance market.[8] By training on a massive universe of labeled compliant and fraudulent claims, even a basic machine learning algorithm can be applied upstream in the claim payment pathway, signaling to a human when the characteristics of a claim warrant deeper investigation. Instead

of weeding through 4 billion U.S. health insurance claims each year by hand, AI-powered fraud analysis can point humans to the very small, but overall hugely expensive, number of fraudulent ones that demand the most human attention. Moreover, in staying one step ahead of fraudsters, AI may also support a defensive strategy against increasingly creative approaches to fraud developed using novel machine learning-based strategies.

Create novel platforms to fundamentally transform population health

While machine learning provides immediate value through automation of traditional health plan processes, AI's major long-term benefit will be the broader transformation of population health management. Entities on the leading edge of the technology adoption curve – including forward-thinking incumbent health plans, tech-native new market entrants, providers with population health business models, and self-insured employers – have already begun adopting AI-enabled platforms to address both costs and clinical outcomes. Here are three examples of how we believe AI can geometrically change the population health management value equation.

Example 1: Augment human-based care / Move clinical interventions upstream

According to the Centers for Disease Control (CDC), the United States spends nearly $3 trillion each year – or 90% of all health expenditures – to manage Americans' chronic challenges with physical and mental health.[9] In Medicare alone, the Commonwealth Fund reported in 2017 that the average frail elderly patient – who made up only 4% of Medicare's total population – generated over $6,500 per capita in excess preventable healthcare spending each year.[10]

While identifying and managing high-cost patients is nothing new to health plans, existing technologies have failed to support the identification of the 60% of new patients who join the "high cost" cohort any given year. AI promises to change that, enabling plans to identify members whose health will worsen before it happens, figure out why, and offer clinicians a directed opportunity to intervene upstream similar to the opportunity offered by clinical "early warning systems" in use at hospitals around the nation.

A recent article published by AM and NS shows this is more than just theory.[11] Using Denmark's National Health Service and Civil Registration System as data source, they analyzed the

effectiveness of a machine learning model's ability to predict "cost blooms" for the entire population of Western Denmark from 2004 to the end of 2011. Importantly, these "blooms" weren't already high-cost high-utilizers – rather, these Danes were seemingly fine one year and markedly sicker and more expensive the next. Compared to traditional diagnosis-based models, the AI-enabled model was 30% better at finding people who would become sick.

The implication is clear, if a little *Minority Report* like. With novel models of high-touch and concierge primary care exploding across America, imagine being able to target the most comprehensive — and highest-cost — care coordination teams not only to today's highest-acuity patients but to those members whose cost and clinical outcome curves stand to bend the earliest and the furthest. While complete delegation of intervention targeting to machine learning models may risk exacerbating silent dataset biases,[12] we believe this could be overcome in two major ways: first, through purposeful attention to bias during model creation, curation, and validation, and second, by again designing AI-enabled solutions to augment rather than entirely replace human-to-human care management.

Example 2: Align elective patient needs with ideal providers

While they may disagree on how "value" is defined, it's no secret that health plans and providers alike recognize a spectrum of value exists within any healthcare market. In work funded by the Peterson Center on Healthcare, Stanford University's Clinical Excellence Research Center (CERC) has worked to systematically identify "Bright Spots" providers operating at the frontier of value, defined by both high quality of care and low total annual healthcare spending. Were the major features of Bright Spots providers in primary care alone to be adopted nationwide, CERC conservatively projects a potential savings of more than $300 billion.[13]

That said, the current era of value-based care suffers from a fatal flaw. Outside of full-risk capitation models, the majority of provider revenue still driven by fee-for-service medicine. As a result, any value-based reimbursement model that subsequently decreases volume cannibalizes one revenue stream for another. This is where AI-powered choice architecture can make a real difference.

A recent New York Times article[14] highlighted how the treasure trove of Amazon's customer

purchase data is being used to hyper-target advertising for local companies. Instead of a baby formula company targeting consumers who recently purchased bottles, imagine instead a tool like this being used to help shape a patient's health care experience, driving patients to higher-value in-network Bright Spots providers. Instead of targeting a purchase of a product, these approaches could also target purchases of services, food choices or health behaviors.

Consider the potential effectiveness of a program to encourage utilization of Bright Spots providers, delivered when people are first contemplating utilization — or even before they are consciously considering such a purchase — and the benefit of such a program for patients, high-value providers, plans and employers. We see profound changes possible for markets driven by these kinds of programs, perhaps even providing the final bridge for providers to accept more risk in return for market share.

Example 3: Dramatically improve benefit design

On the consumer side, AI already drives many of the purchasing decisions we make without our knowing it. As noted above, Amazon is not only pushing us recommendations based on our own experience and that of millions of other consumers but selling the ability to target us as members of a definable cohort. These same kinds of cohorts exist in health: men aged 35-49 with rising risk of heart disease, women aged 18-34 with rising risk of diabetes, and so on. People in these cohorts may not bloom in cost this year or next year, but they're statistically more likely to bloom at some point than people in other, healthier cohorts.

Even relatively rare cohorts stand to experience significant benefit. Consider people with

familial hypercholesterolemia (FH), which affects about 1 in 250 global citizens.[15] FH is caused by mutations in genes coding for the LDL receptor which reduce the receptor's ability to recycle so-called "bad cholesterol." The result yields circulating LDLs of up to six times higher and a risk of coronary heart disease five times higher than the general population. Yet in the U.S., fewer than 10 percent of those with FH know they have the disease.

While pharmacogenetic research has made it possible to both diagnose and treat FH, screening or treating the entire world population remains wildly cost prohibitive. Data science offers a solution. Working with the FH Foundation, NS and a Stanford colleague trained an algorithm on the electronic health records of known FH cases to identify individuals at risk of FH. In a

validation study, the algorithm was correct 8 out of 10 times when it flagged a patient as high-risk for FH. Imagine using such an approach at scale, offering near instantaneous economic optimization of screening/treatment combinations for hundreds if not thousands of types of high-risk patients.

With the massively increased effectiveness of AI-powered targeting, we can expect similar improvement in benefit utilization. Those payers with the longest time-horizons – think big single-payer nationalized systems, Medicare, and perhaps some very large employers – stand to gain the most economically from such an advantage. Imagine identifying a member at age 67 who in the next ten years has a modifiable trajectory in clinical outcomes and cost ranging from $75,000 to $750,000. With radically redesigned benefit structures to address not just clinical determinants of health but personal-choice driven determinants of health (food choice, activity, smoking, injury prevention), the human health and economic impacts of such an approach could both prove enormous.

Looking forward

In its landmark 2017 report assessing the potential impact of AI in healthcare, the elite JASON group[16] noted that "Unlike previous eras of excitement over AI, the potential of AI applications in health may make this era different because the confluence of the following three forces has primed our society to embrace new health centric approaches that may be enabled by advances in AI: 1) frustration with the legacy medical system, 2) ubiquity of networked smart devices in our society, 3) acclimation to convenience and at-home services like those provided through Amazon and others." We wholeheartedly agree. As we've outlined here, we believe the value proposition AI-enabled population management will follow a logical flow:

- Massive data processing and analysis capabilities to generate...
- New insights from clinical, financial, actuarial data...
- That will disrupt health plan operations and care delivery systems...
- Thus improving efficiency, effectiveness, safety and quality of care...
- While distributing resources to maximally and positively affect the health of people served by population health managers

With American consumers not only prepared for but expecting intelligent services and increasingly accustomed to trading better price, convenience, quality, and experience in exchange for access to their data – the time for AI-powered health plans is now. Whether captured by tech-native newcomers, attentive incumbents, or vertically aligned combinations of the two, AI's power will unquestionably drive new value in population health management.

References

1. Rao B, Landi W, Rucker D, Inventors. Systems and methods for automated classification of health insurance claims to predict claim outcome. 2005-06-23, 2005.
2. NTUC Income becomes the first insurer in Singapore to use IBM Watson Explorer to improve the efficiency of pre- and post-IncomeShield claims – IBM Fintech portal [press release]. https://www.ibm.com/think/fintech/asia/, August 15 2017.
3. Brynjolfsson E, McAfee A. The Business of Artificial Intelligence. *Harvard Business Review*. 2017(July).

4. Hileman G, Steele S. *Accuracy of Claims-Based Risk Scoring Models.* 2016.
5. Shin H, Park H, Lee J, Jhee WC. A scoring model to detect abusive billing patterns in health insurance claims. *Expert Systems with Applications.* 2012;39(8):7441-7450.
6. Ryman-Tubb N, Krause P, Garn W. How Artificial Intelligence and machine learning research impacts payment card fraud detection: A survey and industry benchmark. *Engineering Applications of Artificial Intelligence.* 2018;76:130-157.
7. Castellanos S, Nash K. Bank of America Confronts AI's 'Black Box' With Fraud Detection Effort. *Wall Street Journal.* 2018. Accessed January 30, 2019.
8. Melendez S. Insurers turn to artificial intelligence in war on fraud. *Fast Company.* 2018. Published June 26, 2018. Accessed January 30, 2019.
9. CDC. Health and Economic Costs of Chronic Disease. https://www.cdc.gov/chronicdisease/about/costs/index.htm. Published 2018. Updated November 2018. Accessed January 30, 2019, 2019.
10. Jha A. Preventable Spending High-Cost Medicare Patients. The Commonwealth Fund. https://www.commonwealthfund.org/publications/journal-article/2017/oct/concentration-potentially-preventable-spending-among-high. Published 2017. Accessed January 30, 2019.
11. Tamang S, Milstein A, Sørensen HT, et al. Predicting patient 'cost blooms' in Denmark: a longitudinal population-based study. *BMJ Open.* 2017;7(1):e011580.
12. Zou J, Schiebinger L. AI can be sexist and racist — it's time to make it fair. *Nature.* 2018;559(7714):324.
13. Bright Spots Research. Clinical Excellence Research Center – Stanford Medicine. https://med.stanford.edu/cerc/research/bright-spots-research.html. Published 2019. Accessed January 30, 2019.
14. Weise K. Amazon Knows What You Buy. And It's Building a Big Ad Business From It. *The New York Times.* January 20, 2019.
15. Goldberg A, Gidding S. Knowing the Prevalence of Familial Hypercholesterolemia Matters. *Circulation.* 2016;133(11):1054-1057.
16. JASON. *Artificial Intelligence for Health and Health Care.* The MITRE Corporation; December 2017.