Raising the Digital Quotient of Healthcare

Dale Sanders
Health Catalyst
May 2018
New England Regional HIMSS Chapter Meeting

My Background

- **15 years in military industrial complex where conflict leads to profits**
- **21 years in the healthcare industrial complex where illness leads to profits**

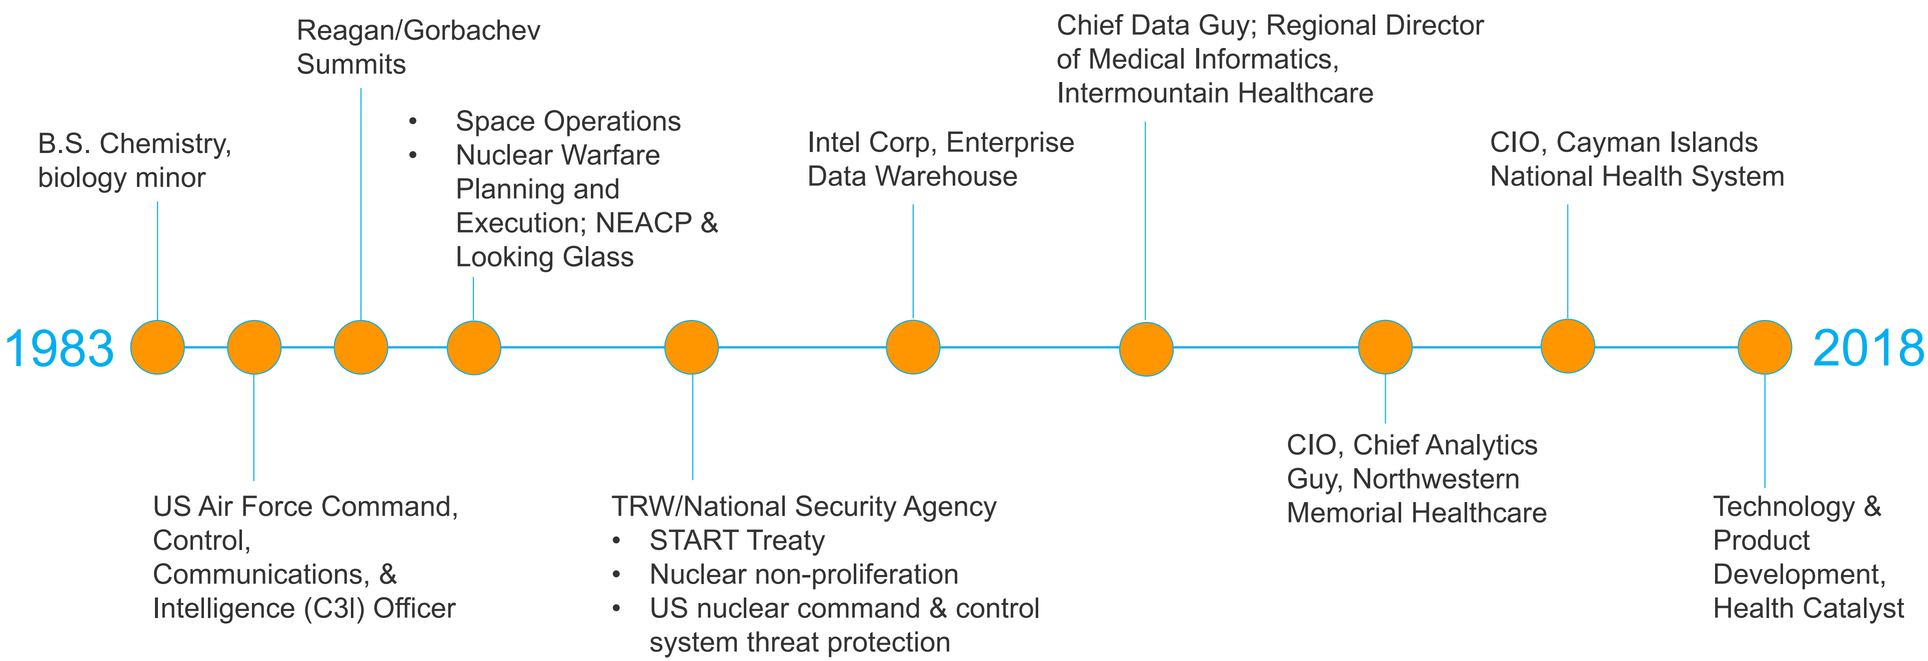

Proof to the young that having no career plan can still turn out ok ☺

Career Advice for Your Digital Future

- It's a great skill to have, but you don't have to be a coder
- The consistent value in my career seems to be…
- Know the pros and cons and what's probable and improbable with the technology
- Take an idea over here and move it over there

Graphic courtesy of Financial Express

Speaking of Skills…

- As measured by McKinsey, healthcare's DQ is not good
- Healthcare executives need to raise their Digital Quotient
- It's up to everyone in this room to help and be role models

"Healthcare CEO, what is your organization's Digital Quotient?"

Healthcare is one of the least digital sectors, and it shows in profit margin growth.

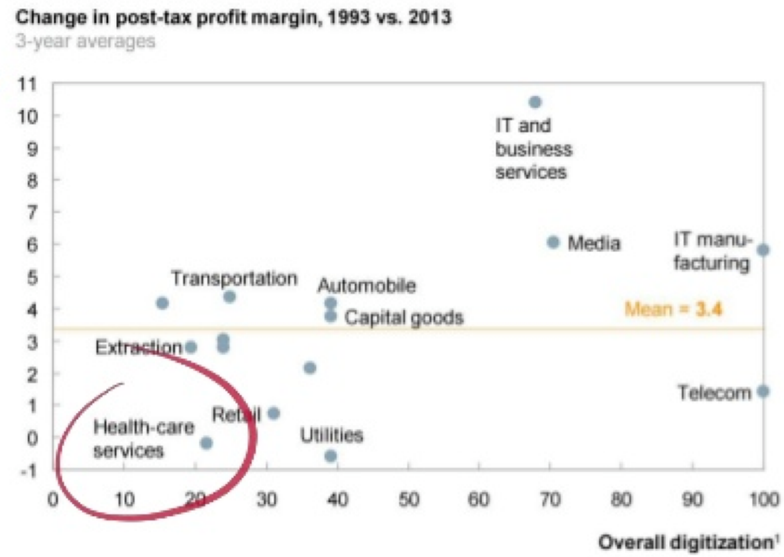

DQ = Data Assets x Data Usage x Data Skilled Labor

Source: McKinsey Corporate Performance Analysis Tool

Today's Story

- Assertions, observations, and futures about data and digitization in healthcare
- Attributes of a modern digital platform
- Thoughts on AI and precision medicine

Mastery, Autonomy, Purpose

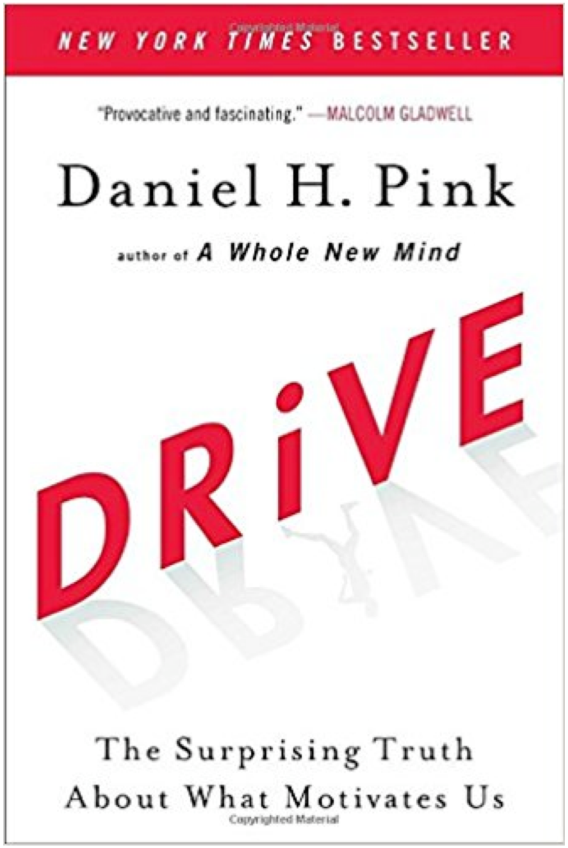

Our current "data-driven" strategy in healthcare is sucking the life out of physicians' sense of Mastery, Autonomy, and Purpose

Our National Data Strategy is a Train Wreck

We've lost our physicians


The NEW ENGLAND JOURNAL of MEDICINE

Annals of Internal Medicine®

LATEST ISSUES CHANNELS CME/MOC IN THE CLINIC JOURNAL CLUB WEB EXCLUSIVES AUTHOR INFO

IDEAS AND OPINIONS | 8 MAY 2018

Physician Burnout in the Electronic Health Record Era: Are We Ignoring the Real Cause?

N. Lance Downing, MD; David W. Bates, MD, MSc; Christopher A. Longhurst, MD, MS

Article, Author, and Disclosure Information

FULL TEXT

MORE

Physician burnout is reaching crisis proportions in the United States (1). Studies have noted a rising prevalence of emotional fatigue. One study suggested that more than half of physicians in some disciplines are burned out and that this proportion is increasing. The number of clinicians leaving the workforce represents a major concern to health care professionals and to the health of the nation. Many factors contribute, but the physician's interaction with electronic health records (EHRs) is especially important now that EHRs have been broadly adopted across the country.

Although EHRs have great potential to improve care, they may also have perverse effects. Some studies suggest that U.S. physicians now spend as much time on "desktop medicine" (interacting with the computer) as they do face to face with patients (2, 3). Providers must divide their attention between patients and the EHR, and many believe that this compromises patient–physician relationships (4). Although few physicians support reverting to paper, there is a growing sense within the medical community that the EHR is driving professional dissatisfaction and burnout.

Time Out — Charting ... Measurement

Catherine H. MacLean, M.D., Ph.D., E...

Performance mea... care system has ... past 30 years. T... Clearinghouse now list... measures. These measures used in various quality-repor... accountability, and payment ... grams sponsored by comme... payers, government agencies, independent quality-assess... organizations. The Centers ... Medicare and Medicaid Ser... (CMS) aims to base 90% of ... care fee-for-service paymen... clinicians on "value" by the ... of 2018 by using perform... scores.

Although most physicians ... the delivery of high-quality ... as a professional imperative,[1] performance-measurement activities face increasing resistance from physicians and some policymakers who believe that current mea- (see box). Using a modified version of the method developed at RAND and UCLA for evaluating the benefits and harms of a medical intervention, we applied the QPP measures list, we identified and rated the validity of 86 that the committee considered relevant to ambulatory general internal medicine. Among these, 32

N ENGL J MED 378;19 NEJM.ORG MAY 10, 2018 1757

- 271 measures in QPP
- 86 related to General Internal Medicine
- 37% invalid, 28% questionable validity
- Highest suicide rate of any profession
  - American Psychiatric Association (APA) 2018. Abstract 1-227, presented May 5, 2018
- >50% burnout in all specialties

- **Quantitative predictability is the metric of scientific precision**
- **The progression of any body of science is measured by its predictability**

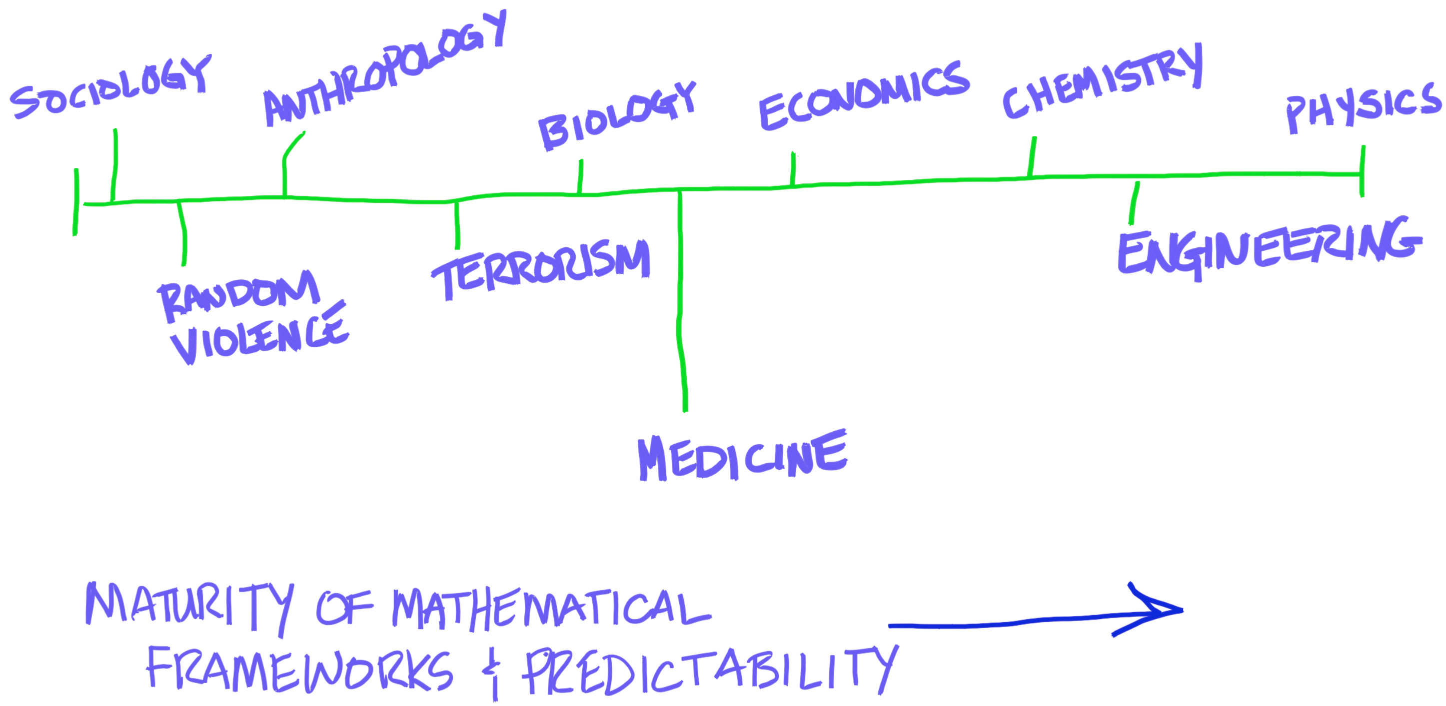

Enabling the Digital Healthcare Conversation

Between a physician and their patient

"I can make a health optimization recommendation for you, informed not only by the latest clinical trials, but also by local and regional data about patients like you; the real-world health outcomes over time of every patient like you; and the level of your interest and ability to engage in your own care. In turn, I can tell you within a specified range of confidence, which treatment or health management plan is best suited for a patient specifically like you and how much that will cost."*

Outcomes and cost data, predictive analytics, machine learning, recommendation engines

***—Inspired by the Learning Health Community**

What's Required to become "Digitized?"

Creating the Digital Twin

1. **Digitize the assets** you are trying to manage and optimize

Airplanes

Patients

2. **Digitize your production process** for managing the assets you are trying to understand and optimize

Air traffic control, baggage handling, ticketing, maintenance, manufacturing

Registration, scheduling, encounters, diagnosis, orders, billing, claims

We haven't digitized the patient, and we've only digitized a clinical encounter to drop a bill.

At Best, EHRs Hold 8% of the Data We Need

- Only 20% of factors affecting health outcomes fall inside traditional healthcare delivery
- On average, patients have 3 healthcare encounters per year
- We are missing data for the other 362 days of the year
- Healthy patients represent our ideal AI training set… but we have no data on healthy patients

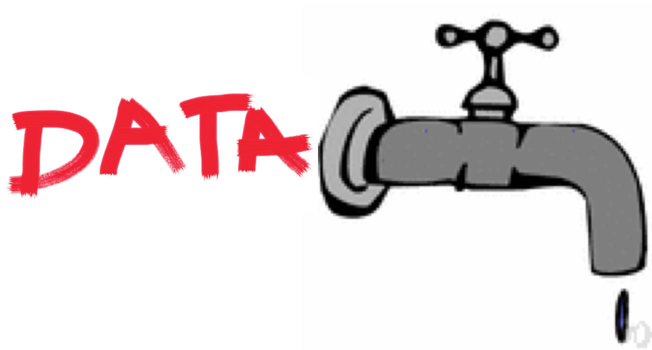

The Simple, Ultimate Analytic Goals

Answer these two questions…

1. What was the **Cost per Unit of Health Outcome Achieved** for this patient?

2. What are the **precise, personal interventions** that will maximize the outcome and minimize the cost?

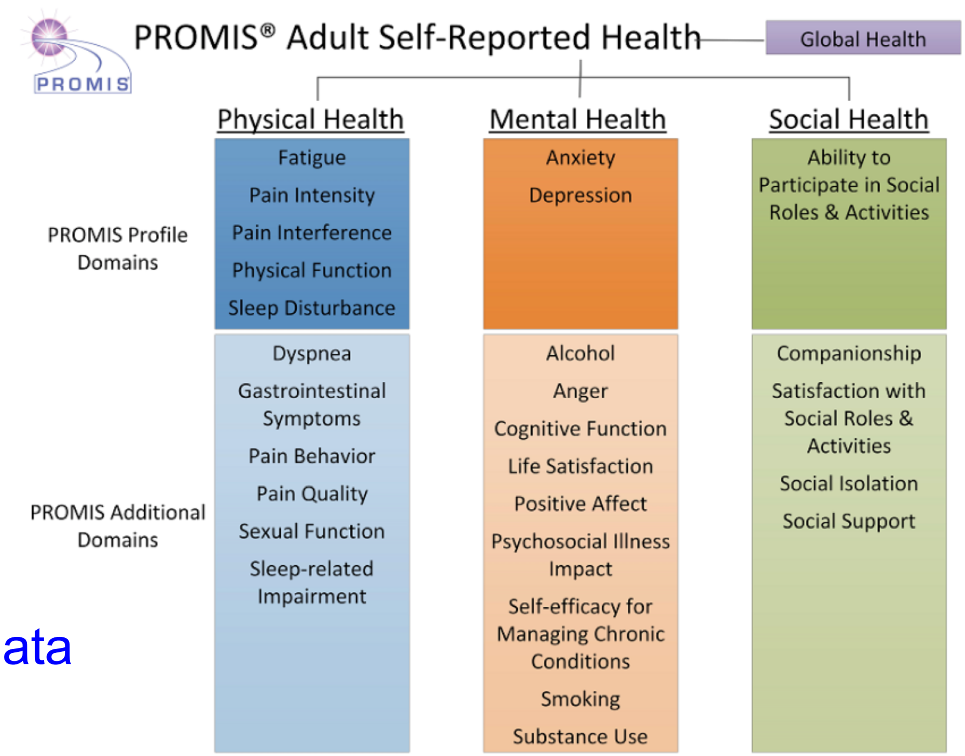

We need outcomes and precise cost data

My Observation About Patient Engagement

- About two-thirds of patients don't want or cannot be "engaged"
- What they really want: When they are sick, they want to be treated safely, affordably, personally, efficiently, and precisely
- Keep that in mind as we lay out a strategy and priorities for digital health

OXFORD
ACADEMIC

Family Practice

Issues Advance articles Submit ▼ Purchase Alerts About ▼ All Family Practice

Change in frequency of patient requests for diagnostic screening and interventions during primary care encounters from 1985 to 2014

Jenny van den Broek, Kees van Boven, Hans Bor, Annemarie A Uijen

Family Practice, cmy031, https://doi.org/10.1093/fampra/cmy031
Published: 26 April 2018

Cite Permissions Share ▼

Abstract

Background
The reason why patients contact a care provider, the reason for encounter (RFE), reflects patients' personal needs and expectations regarding medical care. RFEs can be symptoms or complaints, but can also be requests for diagnostic or therapeutic interventions.

Objectives
Over the past 30 years, we aim to analyse the frequency with which patients consult a GP to request an intervention, and to analyse the impact of these requests on th

Requests for blood tests: 2x increase
Requests for urine tests: 26x increase
Requests for radiology/imaging: 2.4x increase
Requests for medication prescription: 1.2x increase

Patients Owning Their Care

- Netherlands study
- Rate of patient requests for a specific therapeutic or diagnostic intervention
- From 1985-2014
- Significant increase in requests by patients
- Significant increase in compliance by GPs

Future of Diagnosis and Treatment

- Enabled by bio-integrated sensors, patients hold more data about themselves than the healthcare system
- Their data is constantly being updated and uploaded to cloud-based AI algorithms
- Those algorithms diagnose the patient's condition, calculate a composite health risk score, and recommend options for treatment or maintaining health
- The algorithm suggests options for a "best fit" care provider and the ability to socially interact with other patients like them

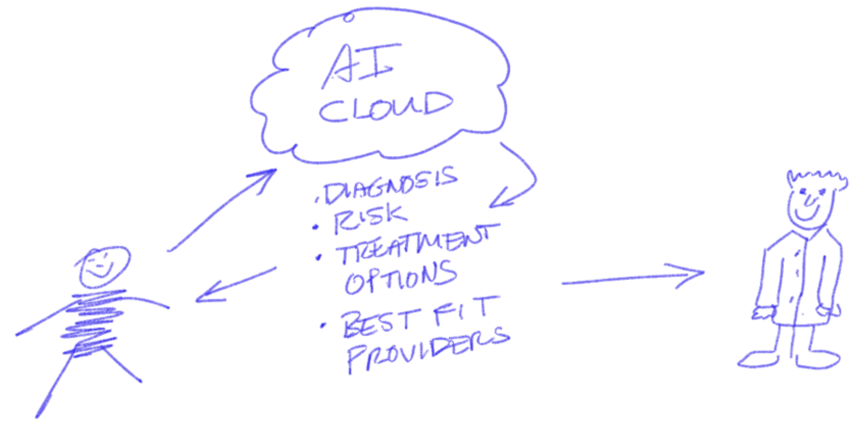

- The patient engages with the care provider, enabled with the output of the AI algorithms

But AI Needs Breadth & Depth of Data in the Domain

This is our strategic data acquisition roadmap

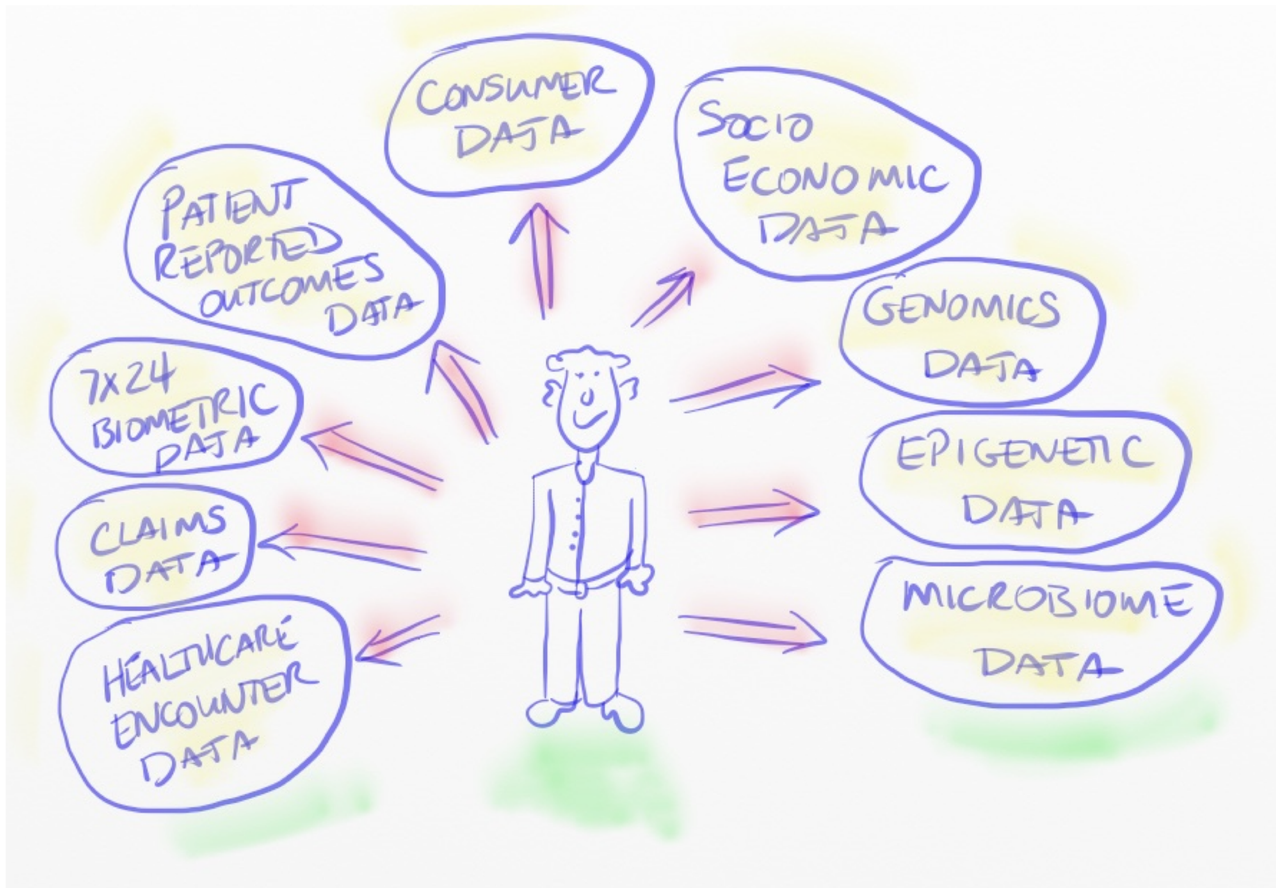

This is my life

This is healthcare's digital view of my life

Digital Accuracy $\propto$ Digital Sampling

We can't possibly provide personal health or precision medicine with only three patient data samples per year

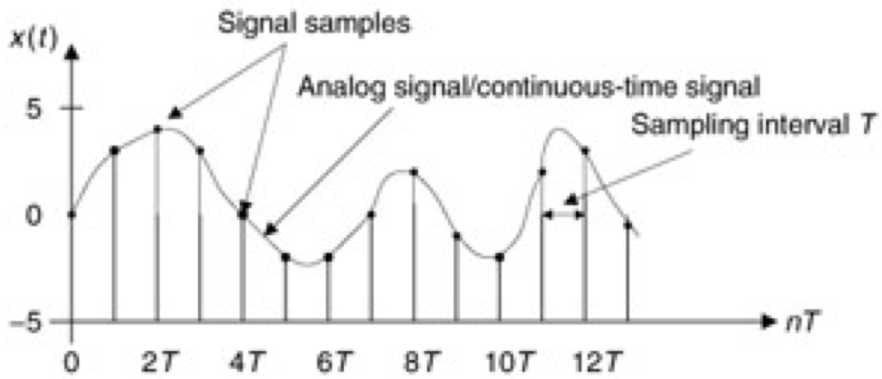

We are not "Big Data" in healthcare yet

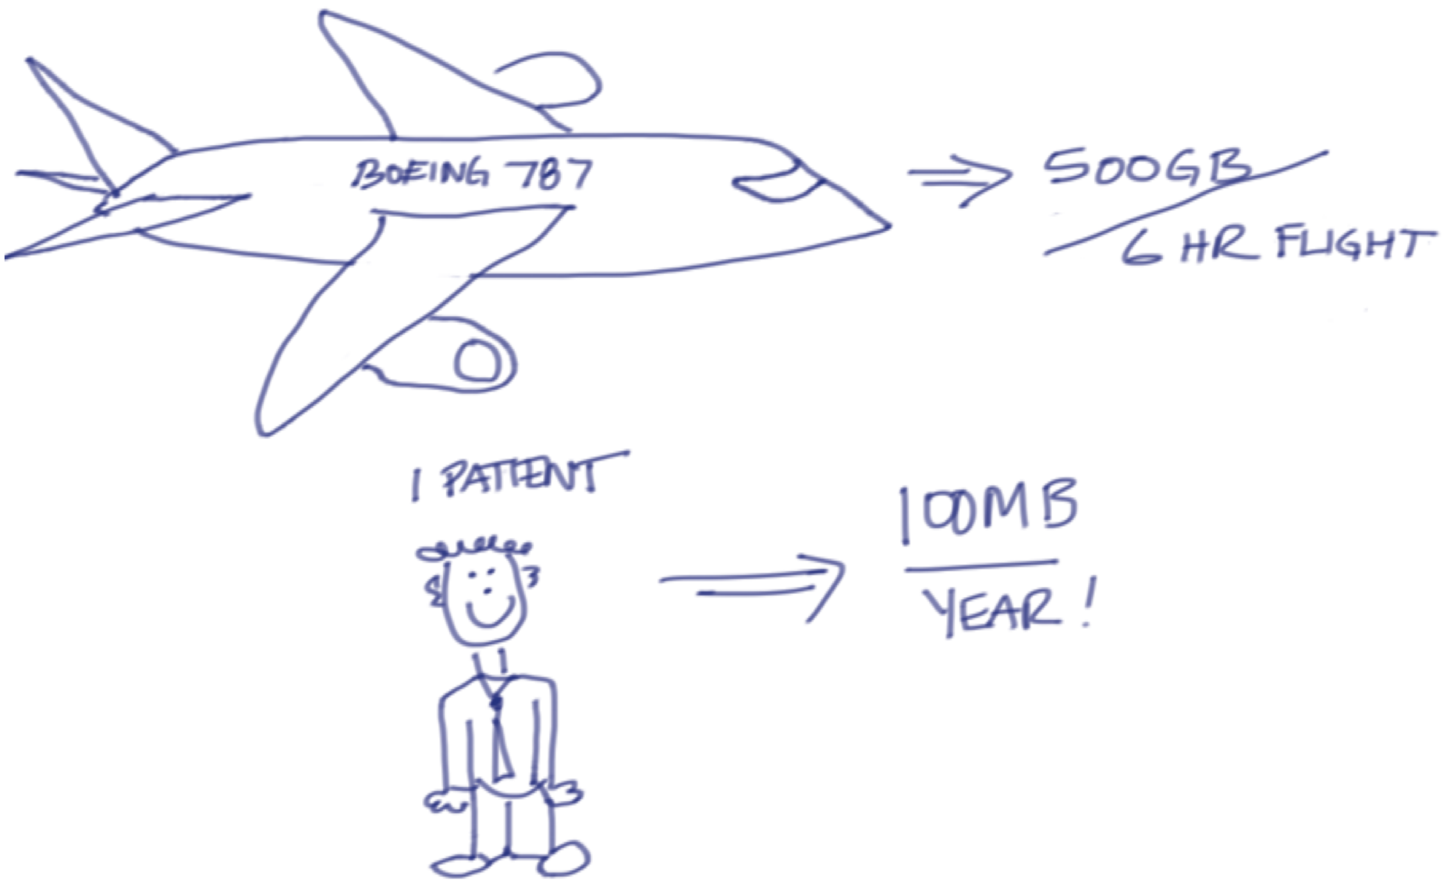

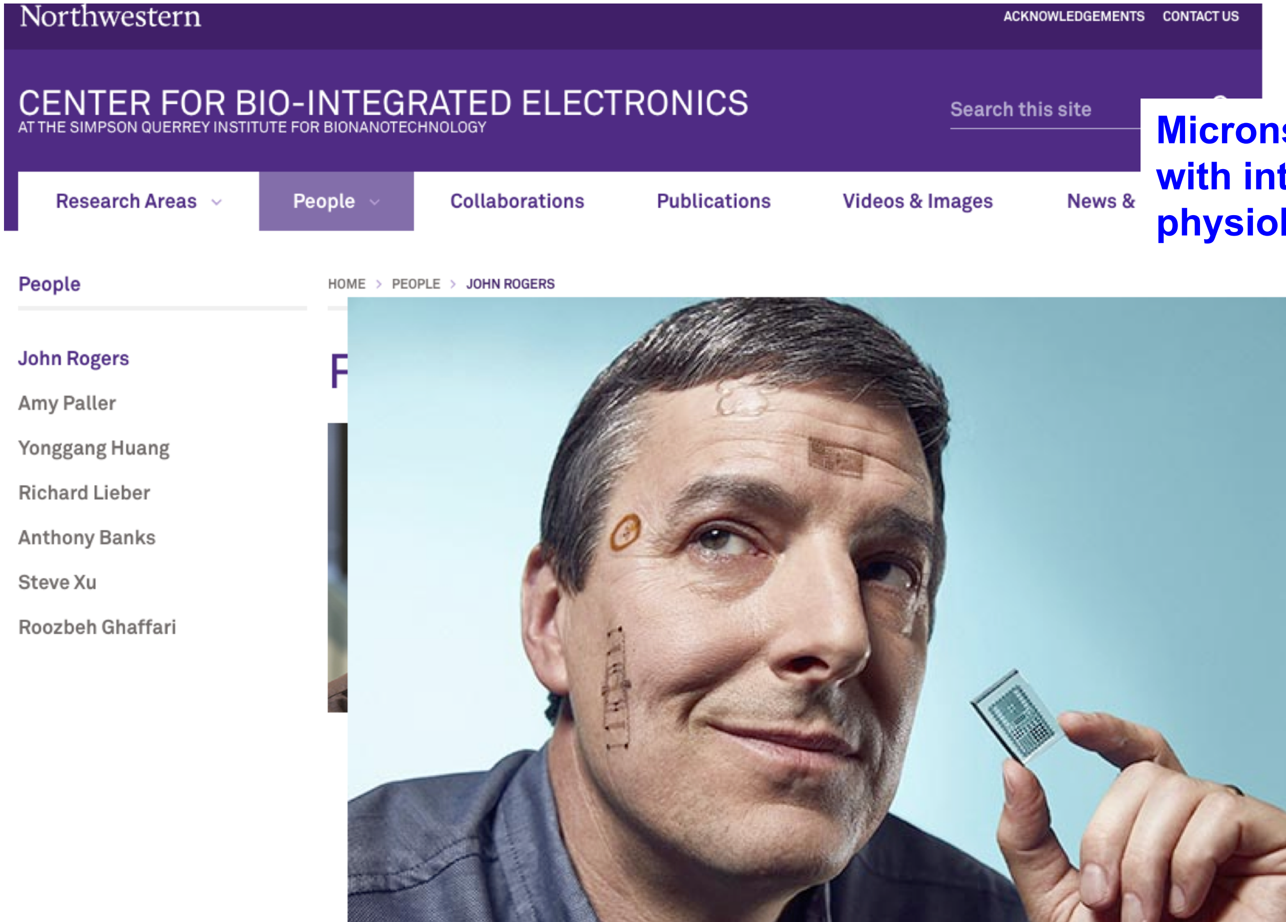

Microns-thin, one-inch skin-pliable sensors with integrated Bluetooth antenna, CPU, physiologic monitors, and wireless power

Rise of The Digitician and Patient Data Profiles

- **Different patient types** have **different data profiles** required for the active management of their outcomes and health
- I'm **not** talking about **quality measures**
- I'm talking about **telemetry**, diagnostics and functional status **about the state of the patient**, not the state of healthcare processes

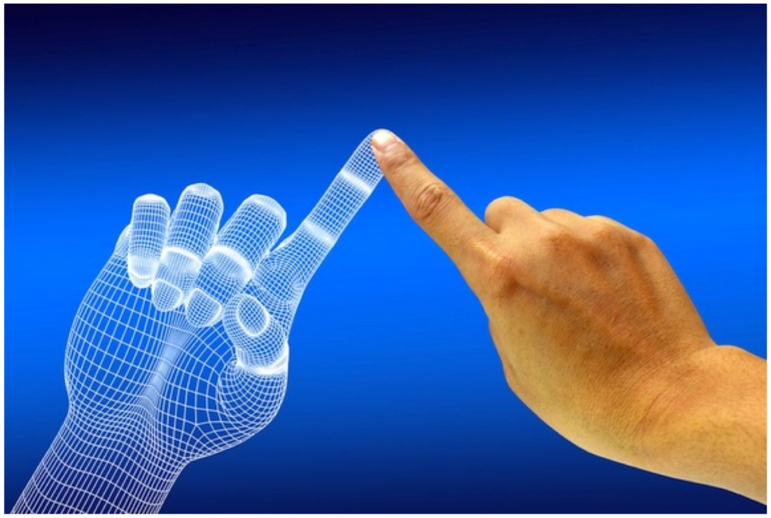

- It's the **Digitician's job** to proactively collect this data for patients in their panel, and feed the analytics of that to the care team and patient

Attributes of a Modern Digital Platform

As computer scientists, we overlooked the last and critically important layer in the technology stack…

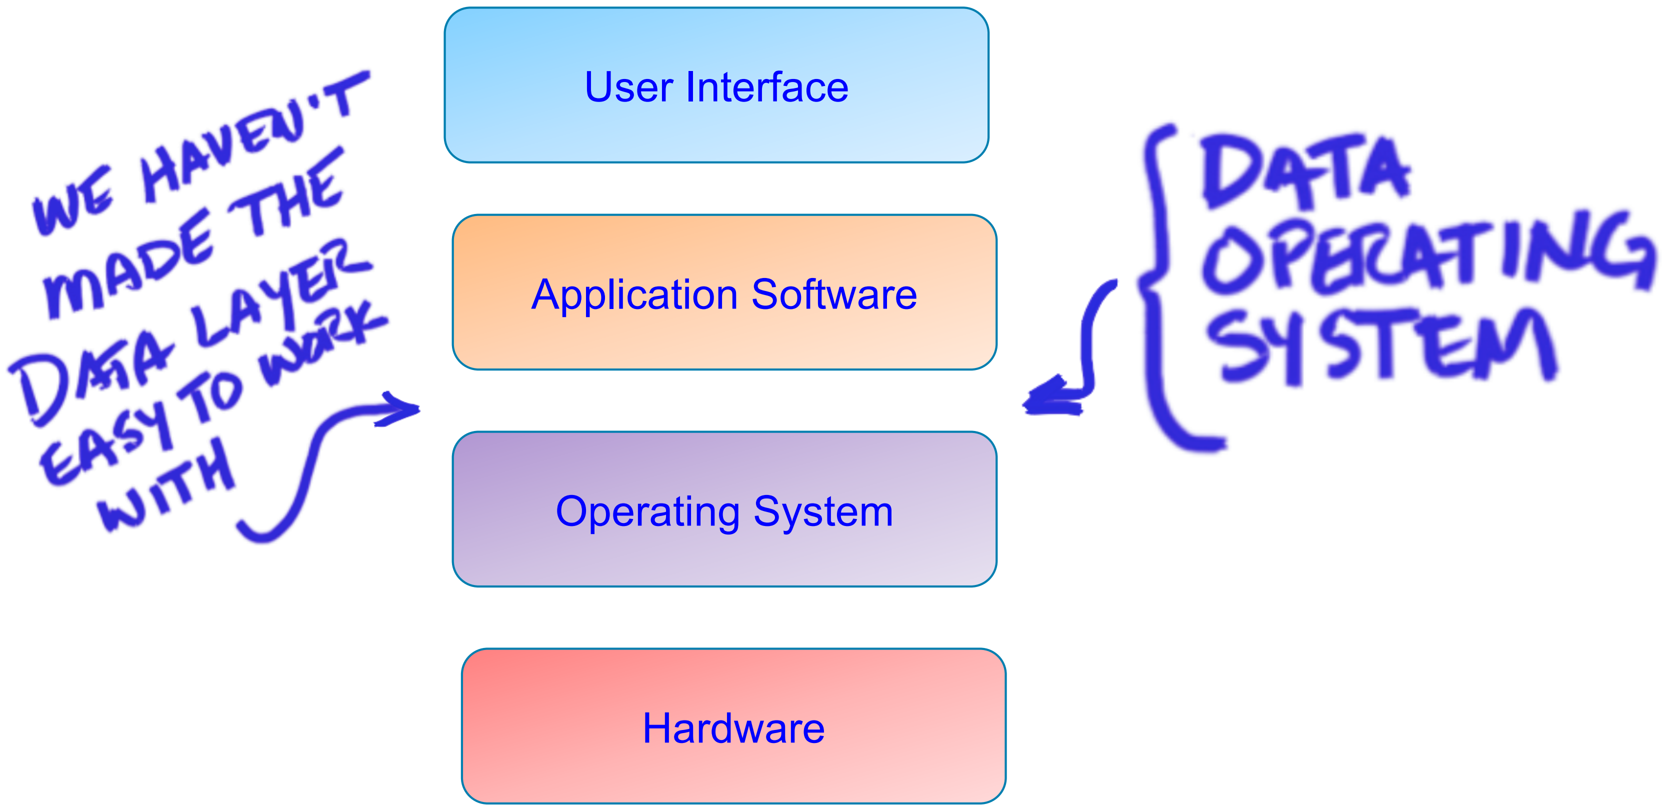

The Evolution of Data Modeling in Analytics

"We know all the use cases, *a priori*"

Monolithic, enterprise data model

"We know some of the use cases, *a priori*"

Intermediate data models

- Harmonized vocabulary
- Comprehensive and persistent agreement about binding logic, e.g., CMS value sets

"We know none of the use cases, *a priori*"

Late binding data models, aka, schema on read

Binding Data in the Pipeline, Not the Data Model

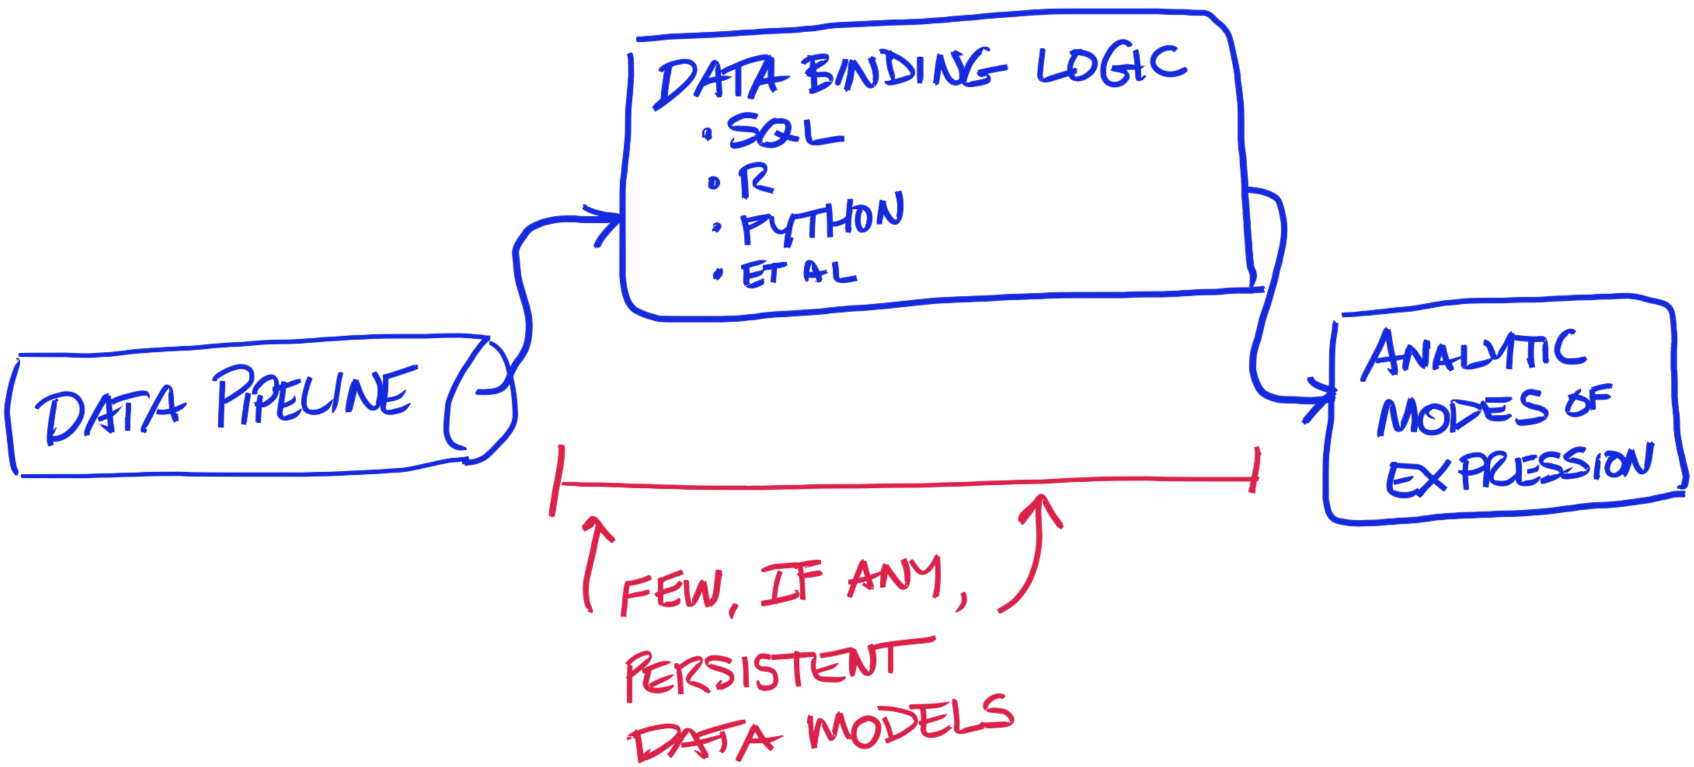

7 Attributes of a Modern Digital Platform

1. **Reusable clinical and business logic:** Registries, value sets, and other data logic lies on top of the raw data and can be accessed, reused, and updated through open APIs, enabling third-party application development.

2. **Single data stream feeds analytics and workflow applications:** Near- or real-time data streaming from the source all the way to the expression of that data through the platform that can support transaction-level exchange of data or analytic processing.

3. **Integrates structured and unstructured data:** Integrates text, images, and discrete structured data in the same environment.

4. **Closed-loop capability:** The methods for expressing the knowledge in the platform, include delivering that knowledge at the point of decision making, for example, back into the workflow of source systems, such as an EHR.

5. **Microservices architecture:** In addition to abstracted data logic, open microservices APIs exist for platform operations such as authorization, identity management, data pipeline management, and DevOps telemetry. These microservices also enable third-party applications to be built on the platform, and constant delivery of software updates, rather than massive, major updates.

6. **AI/Machine learning:** Natively runs AI and machine learning models, and enables rapid development and utilization of ML models, embedded in all applications.

7. **Agnostic data lake:** The platform can be deployed over the top of any healthcare data lake. The reusable forms of logic must support different computation engines; e.g., SQL, Spark SQL, SQL on Hadoop, et al.

The Health Catalyst Data Operating System Architecture

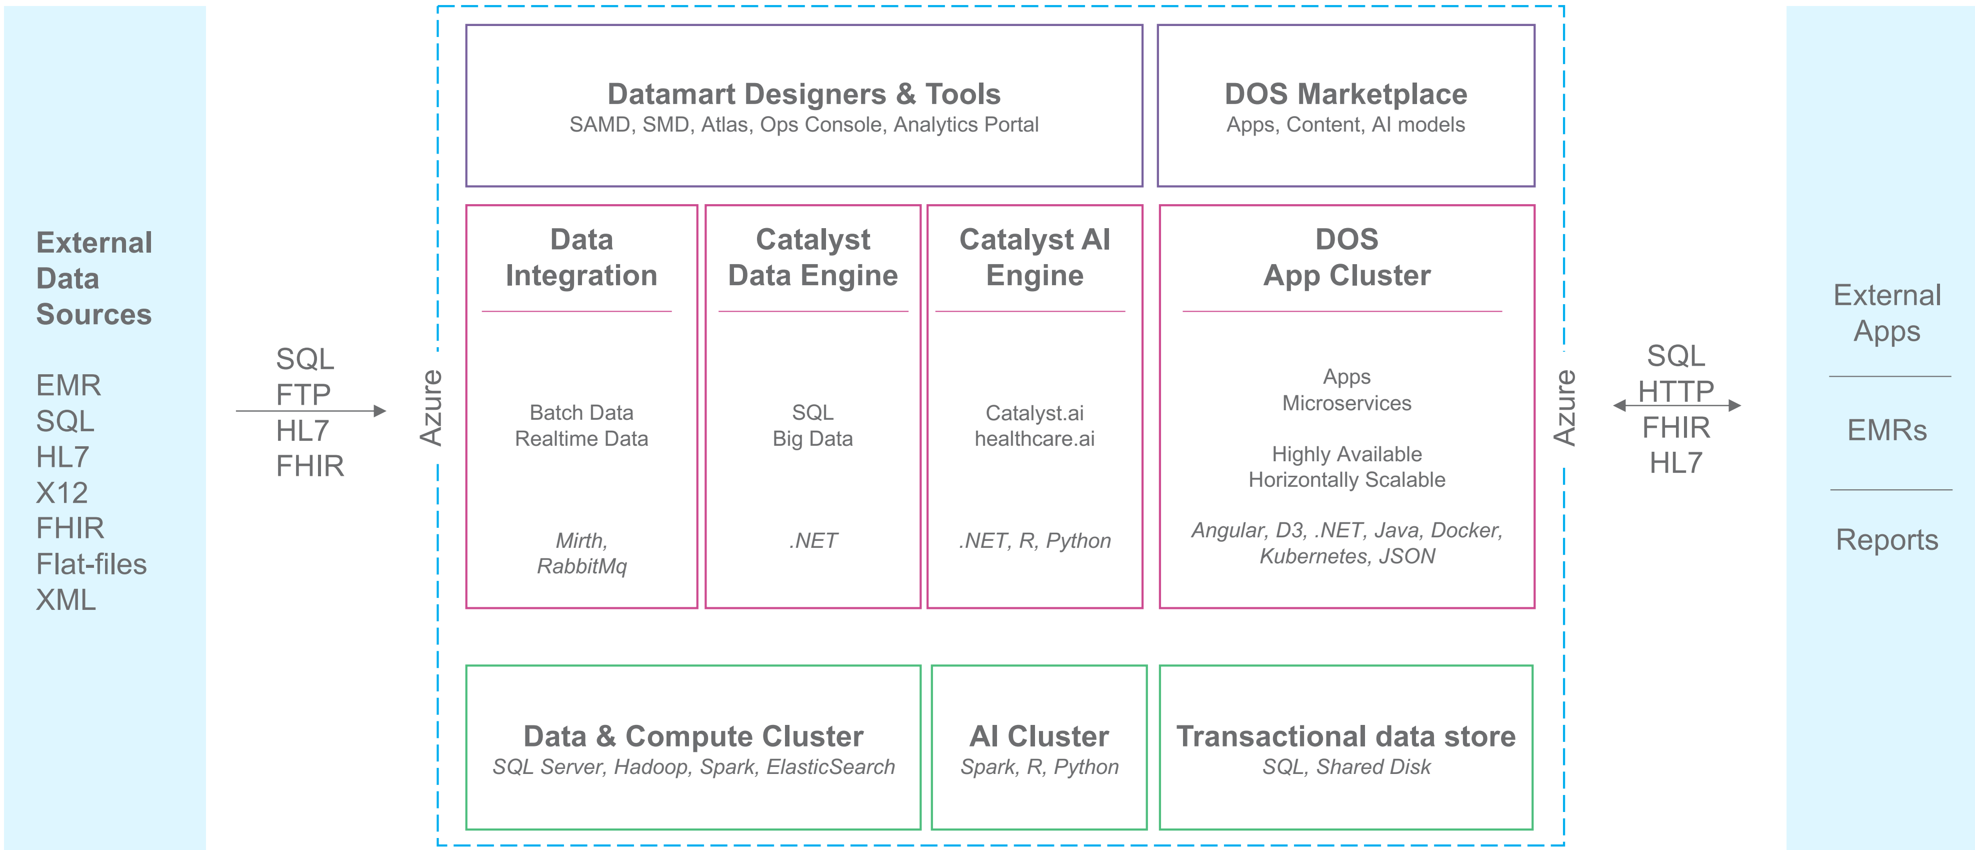

Thoughts on AI and Precision Medicine

Predictive Risk Fatigue

Predictions of risk, without a plan or the ability to intervene, are a liability to the decision maker, not an asset

Discriminative neural networks mimic the human pattern recognition & classification process… "Those are people."

Generative Adversarial Networks (GANs) mimic the opposite human process… "This is what people look like."

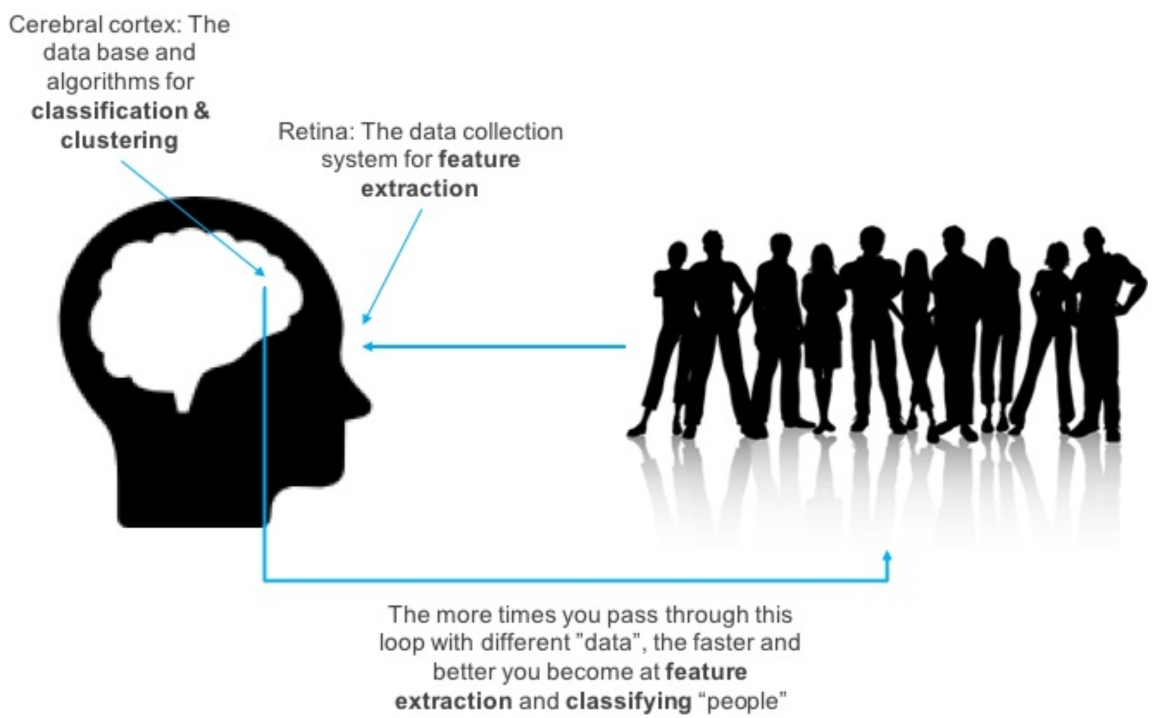

Machine learning boils down to pattern recognition then doing something based on that pattern

Combining three fundamental patterns that will disrupt traditional clinical trials and evidence based care

Patients like this
[pattern 1]

Who were treated like this
[pattern 2]

Had outcomes like this
[pattern 3]

THE LANCET
Diabetes & Endocrinology

Online First Current Issue All Issues Special Issues Multimedia About the Journal Advisory Bo

All Content Search Advanced Search

< Previous Article Volume 6, No. 5, p361–369, May 2018 Next Article >

Articles

Novel subgroups of adult-onset diabetes and their association with outcomes: a data-driven cluster analysis of six variables

Emma Ahlqvist, PhD, Petter Storm, PhD, Annemari Käräjämäki, MD†, Mats Martinell, MD†, Mozhgan Dorkhan, PhD, Annelie Carlsson, PhD, Petter Vikman, PhD, Rashmi B Prasad, PhD, Dina Mansour Aly, MSc, Peter Almgren, MSc, Ylva Wessman, MSc, Nael Shaat, PhD, Peter Spégel, PhD, Prof Hindrik Mulder, PhD, Eero Lindholm, PhD, Prof Olle Melander, PhD, Ola Hansson, PhD, Ulf Malmqvist, PhD, Prof Åke Lernmark, PhD, Kaj Lahti, MD, Tom Forsén, PhD, Tiinamaija Tuomi, PhD, Anders H Rosengren, PhD, Prof Leif Groop, PhD

† Contributed equally

Published: 01 March 2018

PlumX Metrics

DOI: https://doi.org/10.1016/S2213-8587(18)30051-2 | CrossMark

Article Info

Summary | Full Text | Tables and Figures | References | Supplementary Material

Summary

Background

Diabetes is presently classified into two main forms, type 1 and type 2 diabetes, but type 2 diabetes in particular is highly heterogeneous. A refined classification could provide a powerful tool to individualise treatment regimens and identify individuals with increased risk of complications at diagnosis.

Methods

We did data-driven cluster analysis (k-means and hierarchical clustering) in patients with newly diagnosed diabetes (n=8980) from the Swedish All New Diabetics in Scania cohort. Clusters were based on six variables (glutamate decarboxylase antibodies, age at diagnosis, BMI, HbA_{1c}, and

Rules-Based Registries Have Flaws

- Lund University, Sweden
- Using k-means and hierarchical clustering
- Five distinct subtypes of adult-onset diabetes

Topographical Data Analysis & Diabetes Subgroups

- Mt Sinai study
- Visualizes and explores clusters of patients grouped together by algorithms
- 2,551 Type 2 diabetic patients clustered on 73 clinical variables
- A rules-based approach would not find these subgroups

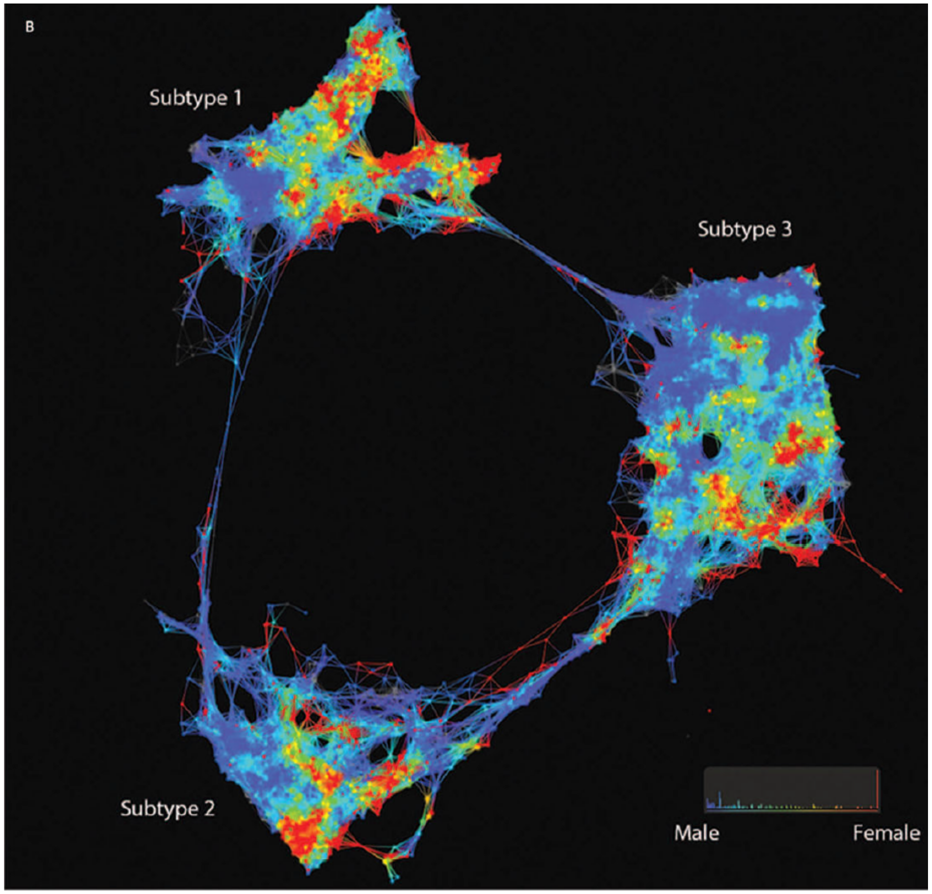

Li L, Cheng W-Y, Glicksberg BS, et al. Identification of type 2 diabetes subgroups through topological analysis of patient similarity. *Science translational medicine*. 2015;7(311):311ra174. doi:10.1126/scitranslmed.aaa9364.

Data Volume vs. AI Model Sophistication

"The Unreasonable Effectiveness of Data", March 2009, IEEE Computer Society; Alon Halevy, Peter Norvig, and Fernando Pereira, of Google

"Invariably, simple models and a lot of data trump more elaborate models based on less data."

EXPERT OPINION

Contact Editor: **Brian Brannon,** bbrannon@computer.org

The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

Eugene Wigner's article "The Unreasonable Effectiveness of Mathematics in the Natural Sciences"[1] examines why so much of physics can be neatly explained with simple mathematical formulas such as $f = ma$ or $e = mc^2$. Meanwhile, sciences that involve human beings rather than elementary particles have proven more resistant to elegant mathematics. Economists suffer from physics envy over their inability to neatly model human behavior. An informal, incomplete grammar of the English language runs over 1,700 pages.[2] Perhaps when it comes to natural language processing and related fields, we're doomed to complex theories that will never have the elegance of physics equations. But if that's so, we should stop acting as if our goal is to author extremely elegant theories, and instead embrace complexity and make use of the best ally we have: the unreasonable effectiveness of data.

One of us, as an undergraduate at Brown University, remembers the excitement of having access to the Brown Corpus, containing one million English words.[3] Since then, our field has seen several notable corpora that are about 100 times larger, and in 2006, Google released a trillion-word corpus with frequency counts for all sequences up to five words long.[4] In some ways this corpus is a step backwards from the Brown Corpus: it's taken from unfiltered Web pages and thus contains incomplete sentences, spelling errors, grammatical errors, and all sorts of other errors. It's not annotated with carefully hand-corrected part-of-speech tags. But the fact that it's a million times larger than the Brown Corpus outweighs these drawbacks. A trillion-word corpus—along with other Web-derived corpora of millions, billions, or trillions of links, videos, images, tables, and user interactions—captures even very rare aspects of human behavior. So, this corpus could serve as the basis of a complete model for certain tasks—if only we knew how to extract the model from the data.

Learning from Text at Web Scale

The biggest successes in natural-language-related machine learning have been statistical speech recognition and statistical machine translation. The reason for these successes is not that these tasks are easier than other tasks; they are in fact much harder than tasks such as document classification that extract just a few bits of information from each document. The reason is that translation is a natural task routinely done every day for a real human need (think of the operations of the European Union or of news agencies). The same is true of speech transcription (think of closed-caption broadcasts). In other words, a large training set of the input-output behavior that we seek to automate is available to us *in the wild*. In contrast, traditional natural language processing problems such as document classification, part-of-speech tagging, named-entity recognition, or parsing are not routine tasks, so they have no large corpus available in the wild. Instead, a corpus for these tasks requires skilled human annotation. Such annotation is not only slow and expensive to acquire but also difficult for experts to agree on, being bedeviled by many of the difficulties we discuss later in relation to the Semantic Web. The first lesson of Web-scale learning is to use available large-scale data rather than hoping for annotated data that isn't available. For instance, we find that useful semantic relationships can be automatically learned from the statistics of search queries and the corresponding results[5] or from the accumulated evidence of Web-based text patterns and formatted tables,[6] in both cases without needing any manually annotated data.

1541-1672/09/$25.00 © 2009 IEEE
Published by the IEEE Computer Society

AI Algorithms are Commodities, Digital Platforms and Infrastructure are Not

Neural Information Processing Systems (NIPS)
Advances in Neural Information Processing Systems 28 (NIPS 2015)

Hidden Technical Debt in Machine Learning Systems

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips
{dsculley,gholt,dgg,edavydov,toddphillips}@google.com
Google, Inc.

Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, Dan Dennison
{ebner,vchaudhary,mwyoung,jfcrespo,dennison}@google.com
Google, Inc.

Abstract

Machine learning offers a fantastically powerful toolkit for building useful complex prediction systems quickly. This paper argues it is dangerous to think of these quick wins as coming for free. Using the software engineering framework of *technical debt*, we find it is common to incur massive ongoing maintenance costs in real-world ML systems. We explore several ML-specific risk factors to account for in system design. These include boundary erosion, entanglement, hidden feedback loops, undeclared consumers, data dependencies, configuration issues, changes in the external world, and a variety of system-level anti-patterns.

"…it is dangerous to think of these quick wins as coming for free. Using the software engineering framework of technical debt, we find it is common to incur massive ongoing maintenance costs in real-world ML systems."

The machine learning code, in the black box, is a small fraction of the ML ecosystem

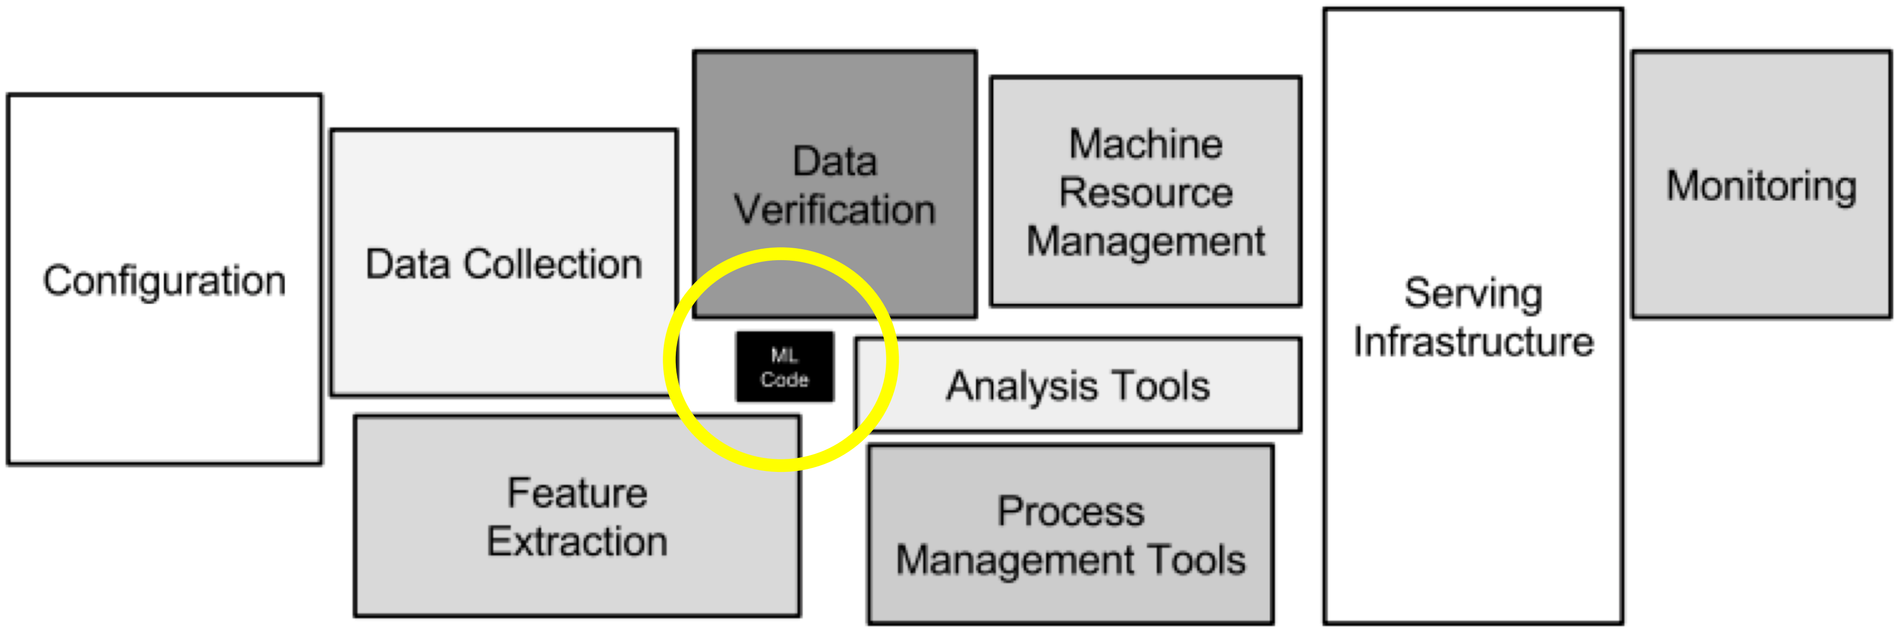

This is not the land of small, niche startups or home grown systems

The Siren's Temptation of Home Grown Digital Platforms

Plug your sailors' ears, Odysseus ☺

- Public cloud makes the infrastructure an incredibly appealing and affordable commodity

The hard part is…

- The collection, curation, and management of data and the logic associated with that data
- The development of APIs and applications

Remember when we were all building our own PCs?

1868, Firmin Girard

In Closing…

- **Drive**: Our digital strategy must enhance Mastery, Autonomy, and Purpose
- **Freud**: Our data isn't as big as we like to think it is in healthcare
- **Platform**: It's overdue in healthcare by 10 years
- **Debt**: AI will disrupt healthcare, no doubt about that, but it's not a slam dunk
- **Tom Brady**: Come on, I had to mention him somewhere ☺