Towards Online Universal Quality Healthcare through AI

Xavier Amatriain (@xamat)
xavier@curai.com

Curai

Being a physician is hard(er)

- Doctors have ~15 minutes to capture information* about a patient, diagnose, + recommend treatment
- *Information
    - Patient's history
    - Patient's symptoms
    - Medical knowledge
        - Learned years ago
        - Latest research findings (70+ journal articles per day)
        - Different demographics
- Data is growing over time, so is complexity
- Manual adaptation is challenging

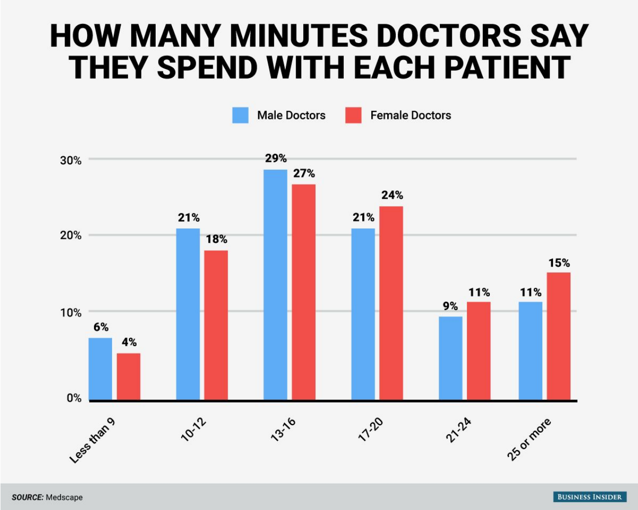

Opinions

When medical care is delivered in 15-minute doses, there's not much time for caring

The Washington Post

Democracy Dies in Darkness

Being a physician is hard(er) (and tech hasn't helped!)

CURAI

Study: EHRs Contribute to Family Physician Stress, Burnout

January 16, 2019 12:06 pm Michael Devitt – In theory, health information technology is supposed to improve communication among health care professionals, make it easier to access and review patient data, cut through the billing and insurance bureaucracy, and enhance the overall health experience for physicians and patients alike.

But as most family physicians know, what sounds good in theory doesn't always play out that way in the real world. In fact, increasing evidence suggests that use of electronic health records (EHRs) (www.annfammed.org) can take up a significant amount of a family physician's workday, making it more

AAFP

Cost of medical errors

CURAI

- 12M misdiagnoses/y
- Errors cause 400k deaths & 4M serious health events
    - Compare to 500k deaths from cancer or 40k from vehicle accidents
- Almost half of those are preventable

Preventability Rationale	Percentage of Events*
Preventable Events (n=133)	
Error was related to medical judgment, skill, or patient management	58%
Appropriate treatment was provided in a substandard way	46%
The patient's progress was not adequately monitored	38%
The patient's health status was not adequately assessed	23%
Necessary treatment was not provided	17%
Event rarely happens when proper precautions and procedures are followed**	14%
Communication between caregivers was poor**	8%
Facility's patient safety systems and policies were inadequate or flawed**	3%
Breakdown in hospital environment occurred (equipment failure, etc.)**	2%
Nonpreventable Events (n=155)	
Event occurred despite proper assessment and procedures followed	62%
Patient was highly susceptible to event because of health status	50%
Care provider could not have anticipated event given information available	35%
Patient's diagnosis was unusual or complex, making care difficult	29%
Harm was anticipated but risk considered acceptable given alternatives**	14%

Online search and/or Healthcare access?

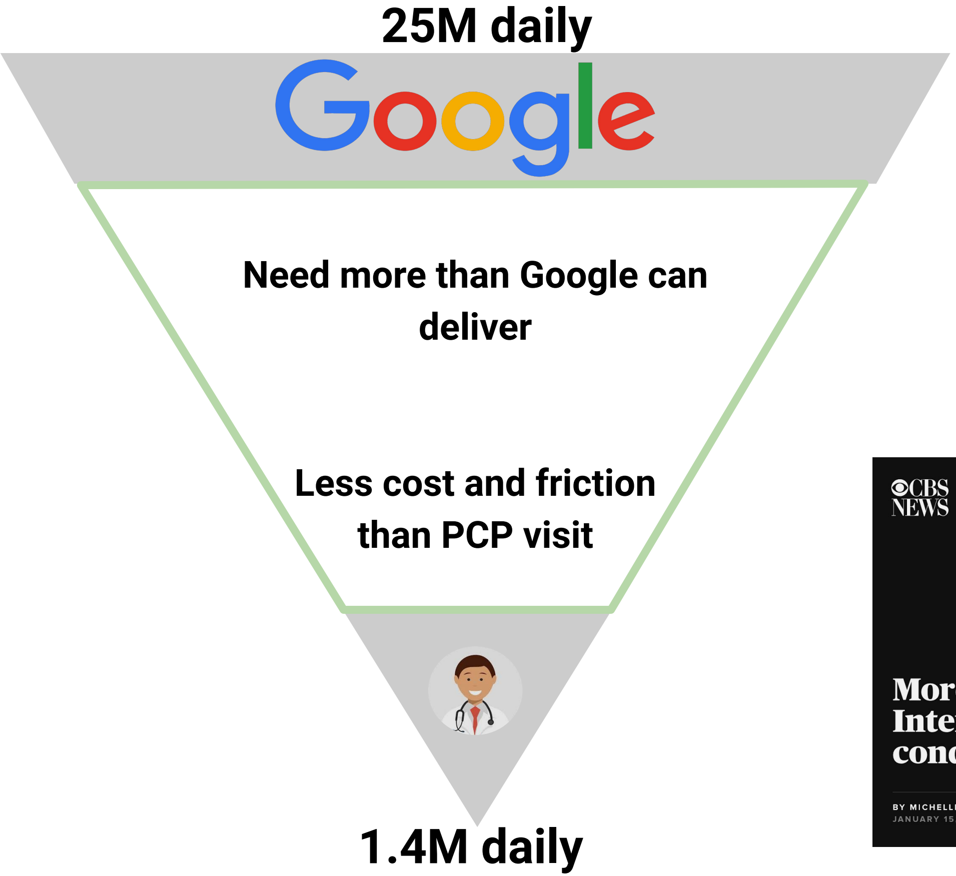

“72% of internet users say they looked online for health information within the past year” [Pew Research]

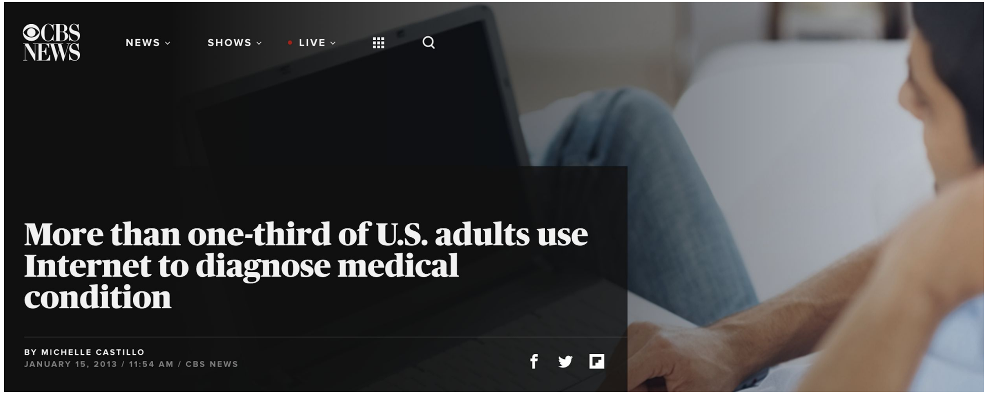

We can do better than Google + Webmd

“Roughly 1 percent of searches on Google are symptom-related”
[Google]

We have an opportunity to reimagine healthcare

We have an ~~opportunity~~ obligation to reimagine healthcare

Towards Online Universal Quality Healthcare

- **Online**= mobile, always on
- **Universal** = scalable, low cost, easy to access/use
- **Quality** = as good as the best doctors, based on best-known practices and scientific evidence, replicable
- **Healthcare** != URL. Healthcare = actionable recommendations + directed access to resources including human professionals when needed

How?

- Online + physicians in the loop = telehealth
- However, standard telehealth is (1) hard to scale, (2) not better quality

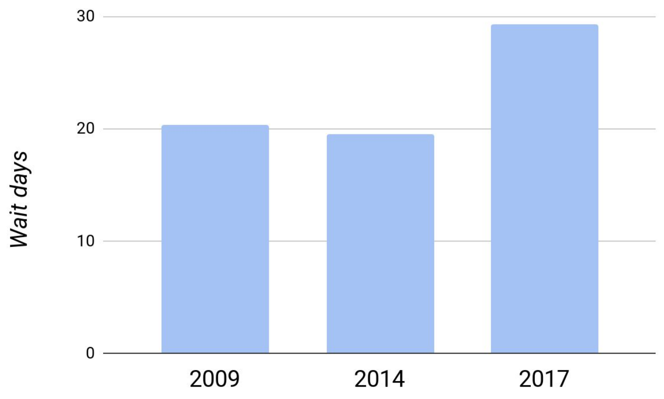

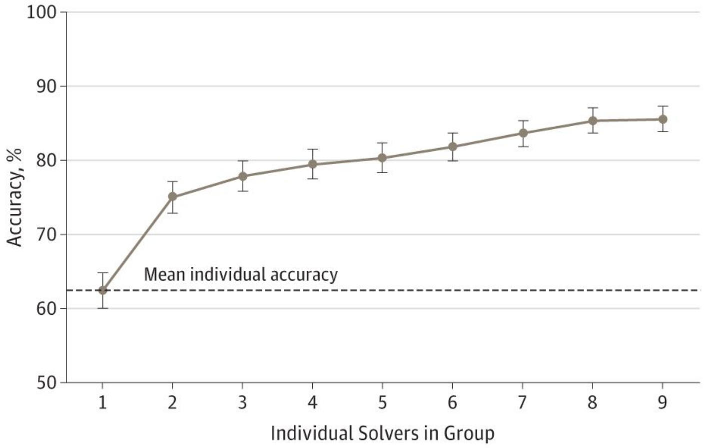

shortage of **120,000 physicians** by 2030

[Merrit Hawkins, 2017 survey]

Barnett et.al. Comparative Accuracy of Diagnosis by Collective Intelligence of Multiple Physicians vs Individual Physicians JAMA, 2018

How?

Scale + Quality =

Automation (aka AI)

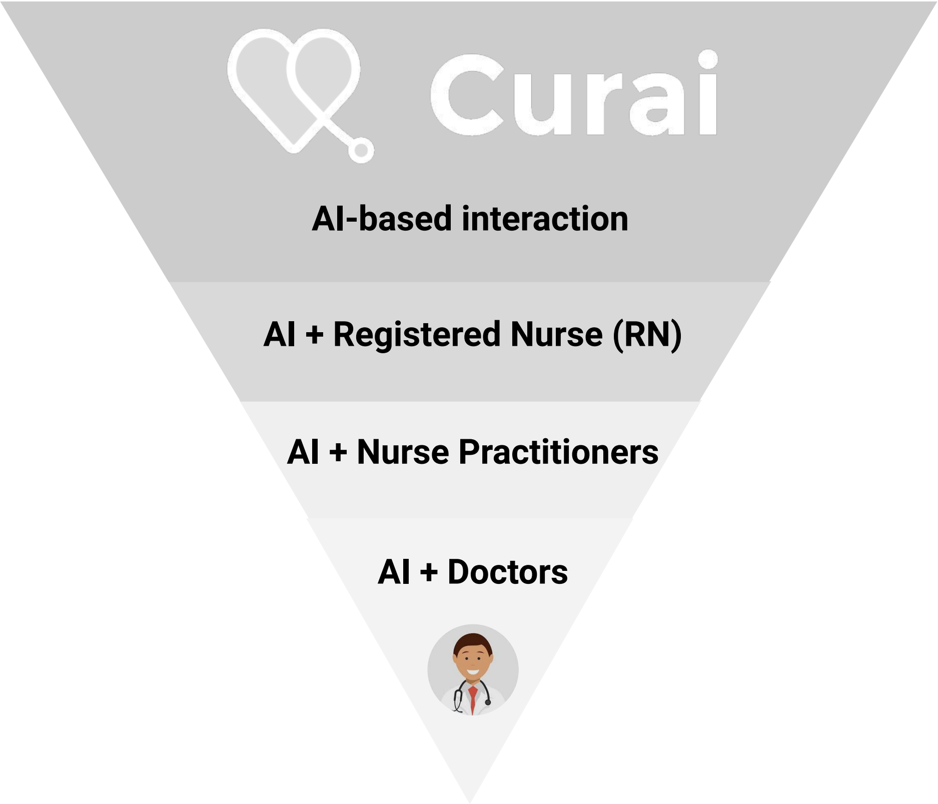

Part II.
Medical AI for online Healthcare

Medical AI for Online Healthcare

CURAI

1. Medical knowledge extraction/representation
2. Conversational Healthcare Systems
3. Automated Triage/Diagnosis/Treatment
4. Multimodal inputs

Part II.

1. Knowledge Extraction & Representation

Medical Knowledge extraction

- Medical ontologies
- Electronic access to medical research
- Access to Electronic Health Records

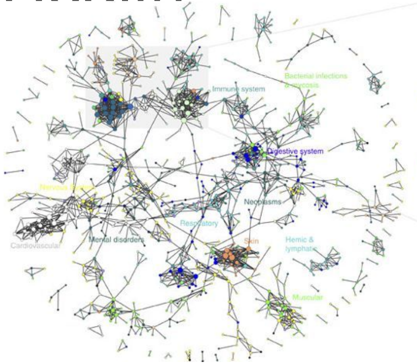


128
© 2008 IMIA and Schattauer GmbH

Extracting Information from Textual Documents in the Electronic Health Record: A Review of Recent Research

S. M. Meystre[1], G. K. Savova[2], K. C. Kipper-Schuler[2], J. F. Hurdle[1]
[1] Department of Biomedical Informatics, University of Utah School of Medicine, Salt Lake City, Utah, USA
[2] Biomedical Informatics Research, Mayo Clinic College of Medicine, Rochester, Minnesota, USA

Summary
Objectives: We examine recent published research on the extraction of information from textual documents in the Electronic Health Record (EHR).
Methods: Literature review of the research published after 1995,

Introduction
In the biomedical domain, the rapid adoption of Electronic Health Records (EHR) with the parallel growth of narrative data in electronic form, along

rules or based on statistical methods and machine learning. The information extracted can then be linked to concepts in standard terminologies and used for coding. The information can also be

Generating Multi-label Discrete Patient Records using Generative Adversarial Networks

Edward Choi[1] MP2893@GATECH.EDU
Siddharth Biswal[1] SBISWAL7@GATECH.EDU
Bradley Malin[2] BRADLEY.MALIN@VANDERBILT.EDU
Jon Duke[1] JON.DUKE@GATECH.EDU
Walter F. Stewart[3] STEWARWF@SUTTERHEALTH.ORG
Jimeng Sun[1] JSUN@CC.GATECH.EDU

[1] Georgia Institute of Technology [2] Vanderbilt University [3] Sutter Health

Abstract

Access to electronic health record (EHR) data has motivated computational advances in medical research. However, various concerns, particularly over privacy, can limit access to and collaborative use of EHR data. Sharing synthetic EHR data could mitigate risk.

In this paper, we propose a new approach, medical Generative Adversarial Network (medGAN), to generate realistic synthetic patient records. Based on input real patient records, medGAN can generate high-dimensional discrete variables (e.g., binary and count features) via a combination of an autoencoder and generative adver-

Knowledge Representation: Medical ontologies

- Snomed Clinical Terms
  - collection of medical terms used in clinical documentation and reporting.
  - clinical findings, symptoms, diagnoses, procedures, body structures, organisms substances, pharmaceuticals, devices…
- UMLS
  - Compendium of many controlled vocabularies
  - Enables translating between terminology systems
- ICD-10
  - International Statistical **C**lassification of **D**iseases and Related Health Problems

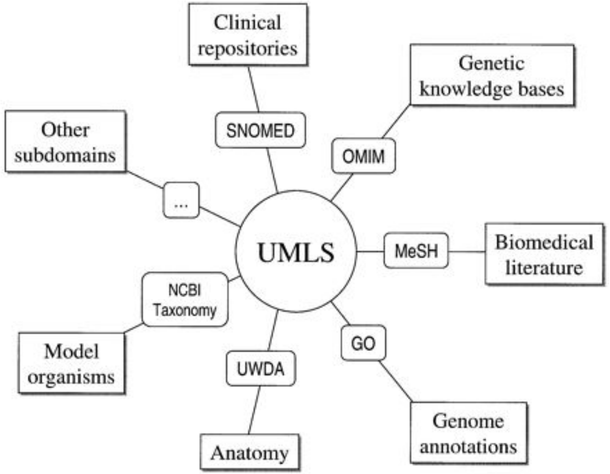

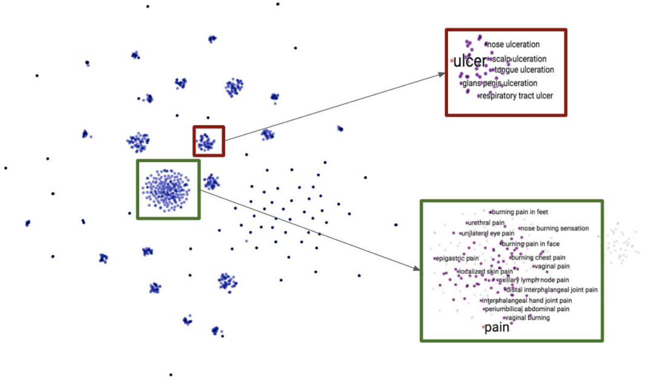

Health knowledge graphs

SCIENTIFIC DATA

OPEN

SUBJECT CATEGORIES
» Data mining
» Diagnosis
» Epidemiology
» Outcomes research

Building the graph of medicine from millions of clinical narratives

Samuel G. Finlayson, Paea LePendu & Nigam H. Shah

Electronic health records (EHR) represent a rich and relatively untapped resource for characterizing the true nature of clinical practice and for quantifying the degree of inter-relatedness of medical entities such as

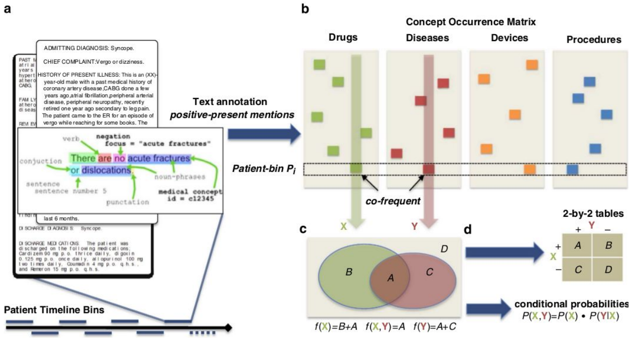

Figure 1. Workflow Architecture. The architecture of our workflow starts with (a) patient notes that are grouped together based on their nearness in time. Given the patient timeline bins, clinical terms are recognized from the notes and recorded into (b) the clinical concept occurrence matrix, which is scanned for (c) counting pairwise the frequency and co-frequency of concepts. This data can be used to calculate (d) contingency tables and Bayesian probability estimates. For example, the concept X has a frequency of *f*(X) and is pairwise co-frequent with concept Y exactly *f*(X,Y) times.

SCIENTIFIC REPORTS

OPEN

Learning a Health Knowledge Graph from Electronic Medical Records

Received: 3 March 2017
Accepted: 1 June 2017
Published online: 20 July 2017

Maya Rotmensch[1], Yoni Halpern[2], Abdulhakim Tlimat[3], Steven Horng[3,4] & David Sontag[5,6]

Demand for clinical decision support systems in medicine and self-diagnostic symptom checkers has substantially increased in recent years. Existing platforms rely on knowledge bases manually compiled through a labor-intensive process or automatically derived using simple pairwise statistics. This study explored an automated process to learn high quality knowledge bases linking diseases and symptoms directly from electronic medical records. Medical concepts were extracted from 273,174 de-identified patient records and maximum likelihood estimation of three probabilistic models was used to automatically construct knowledge graphs: logistic regression, naive Bayes classifier and a Bayesian network using noisy OR gates. A graph of disease-symptom relationships was elicited from the learned parameters and the constructed knowledge graphs were evaluated and validated, with permission, against Google's manually-constructed knowledge graph and against expert physician opinions. Our study shows that direct and automated construction of high quality health knowledge graphs from medical records using rudimentary concept extraction is feasible. The noisy OR model produces a high quality knowledge graph reaching precision of 0.85 for a recall of 0.6 in the clinical evaluation. Noisy OR significantly outperforms all tested models across evaluation frameworks ($p < 0.01$).

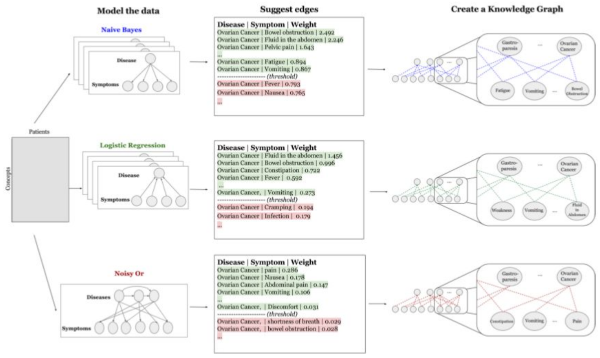

Figure 2. Workflow of modeling the relationship between diseases and symptoms and knowledge graph construction, for each of our 3 models (naive Bayes, logistic regression and noisy OR).

Part II.
2. Conversational Healthcare

Understanding patients and doctors

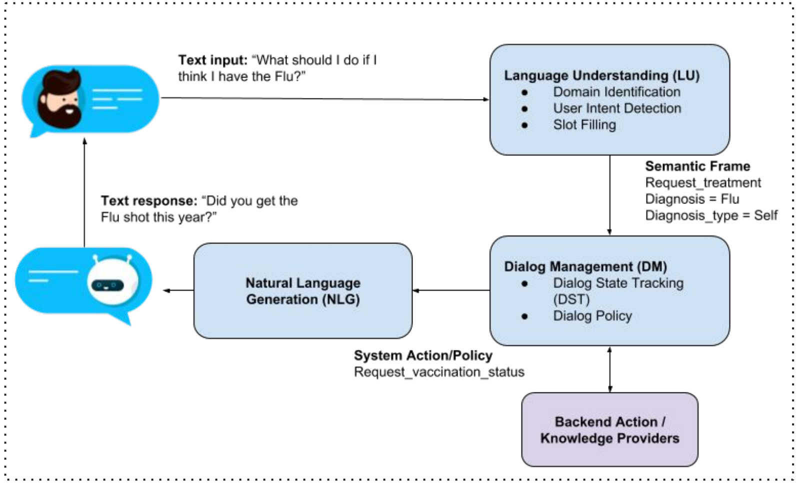

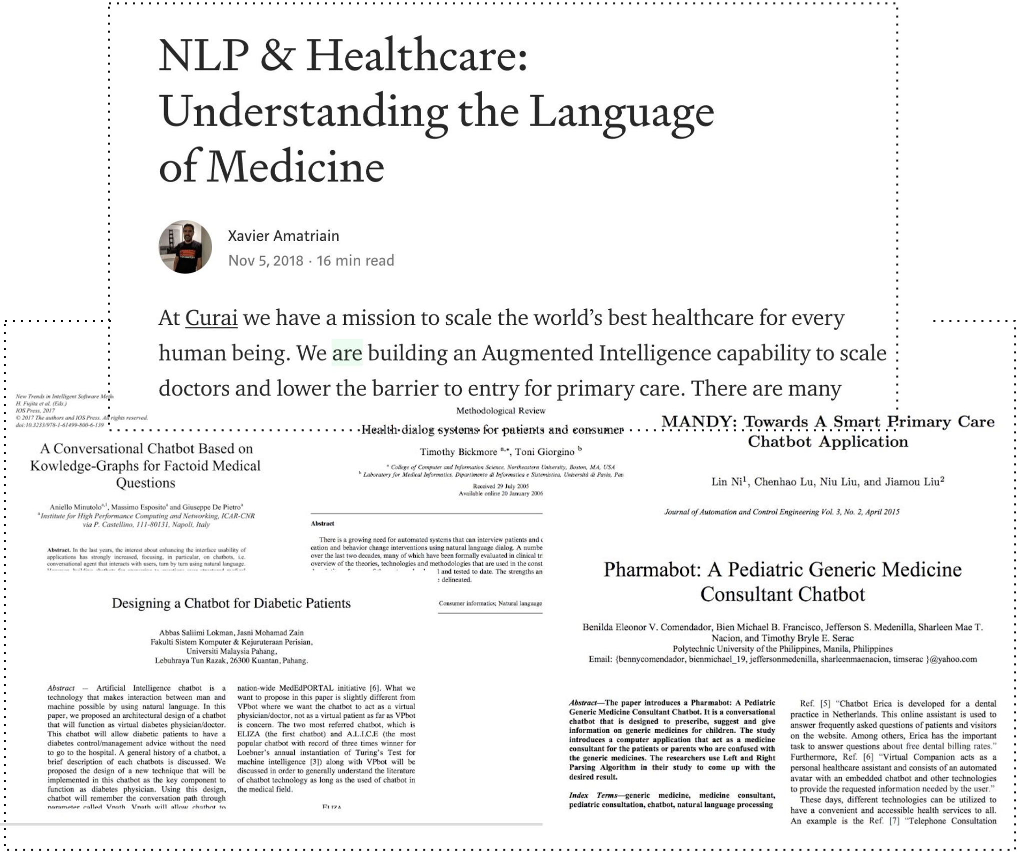

Learning a medically-aware dialog system

CURAI

```
User   : Right now my stomach hurts.
User   : It feels like I need to do a clean out. If you know what I mean
Dr   : Sorry for the abdominal pain Laura. When did you have last bowel movement?
User   : It was yesterday
Dr   : What was the consistency of stool was it soft well-formed or was it hard?
User   : Right now I just went and it is watery and very loosely
User   :
User   : That was causing the problem with my stomach hurts
Dr   : Any blood or mucus with stools? Was it foul smelling?
User   : Nope for all three
Dr   : Any fever?
User   : P
User   : Nope
Dr   : I asked as blood or mucus in stool fever can be due to an underlying infection.
Dr   : Any nausea/ vomiting?
User   : Nope
User   : Why does this happen to me
User   : Is it something that I have ate
Dr   : Diarrhea can often be due to indigestion or an infection. Did you eat outside food or any packaged foo

User   : Yes
```

Part II.
3. Automated(Assisted) Triage/Dx/Rx

Medical Diagnosis

- Diagnosis (*R.A. Miller 1990*):
  - Mapping from patient's data (history, examination, lab exams…) to a possible condition.
  - It depends on ability to:
    - **Evoke history**
    - **Surface symptoms and findings**
    - **Generate hypotheses** that suggest how to **refine** or **pursue** different hypothesis
  - In a **compassionate**, **cost-effective** manner

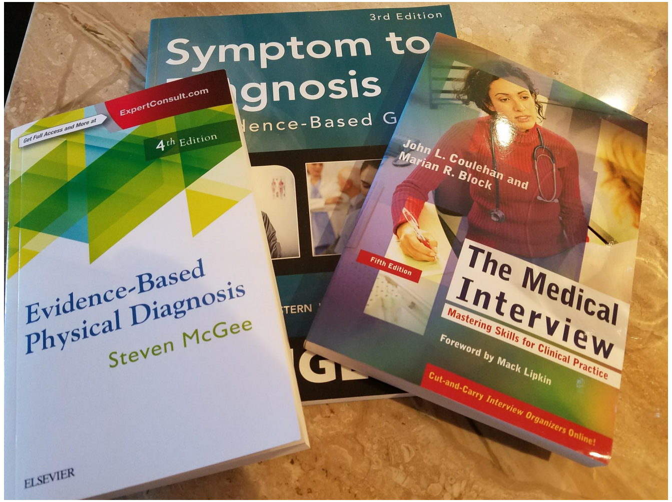

An example: Internist-1/QMR/Vddx

Table 4 A Sample Manifestations List*

Reproduced with permission. From Miller RA, Pople HE Jr, Myers JD. INTERNIST-1, An Experimental Computer-based Diagnostic Consultant for General Internal Medicine. N Engl J Med 1982;307:468-76.

DISPLAY WHICH MANIFESTATION LIST?
ALCOHOLIC HEPATITIS

AGE 16 TO 25	0 1	CHOLESTEROL BLOOD DECREASED	2 2
AGE 26 TO 55	0 3	KETONURIA	1 2
AGE GTR THAN 55	0 2	PROTEINURIA	1 2
ALCOHOL INGESTION RECENT HX	2 4	SGOT 120TO 400	2 3
ALCOHOLISM CHRONIC HX	2 4	SGOT 40TO 119	2 3
SEX FEMALE	0 2	SGOT GTR THAN 400	1 2
SEX MALE	0 4	UREA NITROGEN BLOOD LESS THAN 8	2 2
URINE DARK HX	1 3	UROBILINOGEN URINE ABSENT	1 1
WEIGHT LOSS GTR THAN 10 PERCENT	0 3	UROBILINOGEN URINE INCREASED	2 4
ABDOMEN PAIN ACUTE	1 2	WBC 14000 TO 30000	0 3
ABDOMEN PAIN COLICKY	1 1	WBC 4000 TO 13900 PERCENT NEUTROPHIL(S) INCREASED	0 3
ABDOMEN PAIN EPIGASTRIUM	1 2	WBC LESS THAN 4000	1 1
ABDOMEN PAIN NON COLICKY	1 2	ACTIVATED PARTIAL THROMBOPLASTIN TIME INCREASED	1 3
		ANTIBODY MITOCHONDRIAL	1 1
		ANTIBODY SMOOTH MUSCLE	2 3
		BSP RETENTION INCREASED	1 5

- Internist (1971) - Dr. Jack Myers
- Process for adding a disease requires 2-4 weeks of full-time effort and doctors reading 50 to 250 relevant publications

Evoking Strength	Interpretation
0	Nonspecific—manifestation occurs too commonly to be used to construct a differential diagnosis
1	Diagnosis is a rare or unusual cause of listed manifestation
2	Diagnosis causes a substantial minority of instances of listed manifestation
3	Diagnosis is the most common but not the overwhelming cause of listed manifestation
4	Diagnosis is the overwhelming cause of listed manifestation
5	Listed manifestation is pathognomic for the diagnosis

Frequency	Interpretation
1	Listed manifestation occurs rarely in the disease
2	Listed manifestation occurs in a substantial minority of cases of the disease
3	Listed manifestation occurs in roughly half the cases
4	Listed manifestation occurs in the substantial majority of cases
5	Listed manifestation occurs in essentially all cases—i.e., it is a prerequisite for the diagnosis

ML/AI Approaches to Diagnosis

CURAI

- Early DDSS based on Bayesian reasoning (60s-70s)
- Bayesian networks (80s-90s)
- Neural networks (lately)

Proceedings of Machine Learning for Healthcare 2016 JMLR W&C Track Volume 56

Doctor AI: Predicting Clinical Events via Recurrent Neural Networks

Edward Choi, Mohammad Taha Bahadori MP2893,BAHADORI@GATECH.EDU
College of Computing
Georgia Institute of Technology
Atlanta, GA, USA

Andy Schuetz, Walter F. Stewart SCHUETA1,STEWARWF@SUTTERHEALTH.ORG
Research Development & Dissemination
Sutter Health
Walnut Creek, CA, USA

Jimeng Sun JSUN@CC.GATECH.EDU
College of Computing
Georgia Institute of Technology
Atlanta, GA, USA

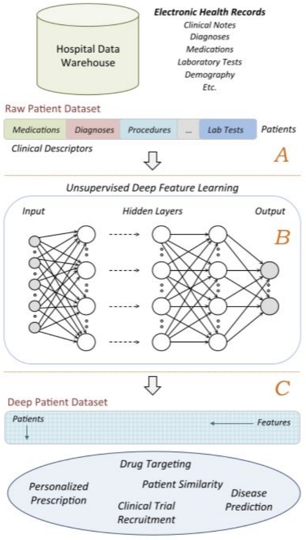

Figure 1. Conceptual framework used to derive the deep patient representation through unsupervised deep learning of a large EHR data warehouse. (**A**) Pre-processing stage to obtain raw patient representations from the EHRs. (**B**) The raw representations are modeled by the unsupervised deep architecture leading to a set of general and robust features. (**C**) The deep features are applied to the entire hospital database to derive patient representations that can be applied to a number of clinical tasks.

SCIENTIFIC REPORTS

OPEN

Received: 28 January 2016
Accepted: 27 April 2016
Published: 17 May 2016

Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records

Riccardo Miotto[1,2,3], Li Li[1,2,3], Brian A. Kidd[1,2,3], Joel T. Dudley[1,2,3]

Secondary use of electronic health records (EHRs) promises to advance clinical research and better inform clinical decision making. Challenges in summarizing and representing patient data prevent widespread practice of predictive modeling using EHRs. Here we present a novel unsupervised deep feature learning method to derive a general-purpose patient representation from EHR data that facilitates clinical predictive modeling. In particular, a three-layer stack of denoising autoencoders was used to capture hierarchical regularities and dependencies in the aggregated EHRs of about 700,000 patients from the Mount Sinai data warehouse. The result is a representation we name "deep patient". We evaluated this representation as broadly predictive of health states by assessing the probability of patients to develop various diseases. We performed evaluation using 76,214 test patients comprising 78 diseases from diverse clinical domains and temporal windows. Our results significantly outperformed those achieved using representations based on raw EHR data and alternative feature learning strategies. Prediction performance for severe diabetes, schizophrenia, and various cancers were among the top performing. These findings indicate that deep learning applied to EHRs can derive patient representations that offer improved clinical predictions, and could provide a machine learning framework for augmenting clinical decision systems.

Our approach: Expert systems as Prior

Knowledge base

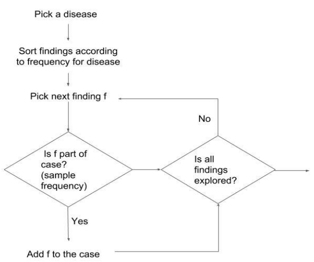

Clinical case simulator

Hepatitis acute viral

anorexia
jaundice
abdomen pain epigastrium
hepatomegaly present
liver enlarged moderate
liver tender on palpation
feces light colored
hands palmar erythema
skin spider angiomas
abdomen pain acute
abdomen pain not colicky
vomiting recent
constipation
vomiting coffee ground

Example of simulated case

Sign

The science of assisting medical diagnosis: From Expert systems to Machine-learned models

Anitha Kannan Follow
Apr 15 · 9 min read

Curai's mission is to scale the world's best healthcare for every human being. We are building an augmented intelligence system to help scale physicians' abilities as well as to lower users' barrier to entry to care. There are many

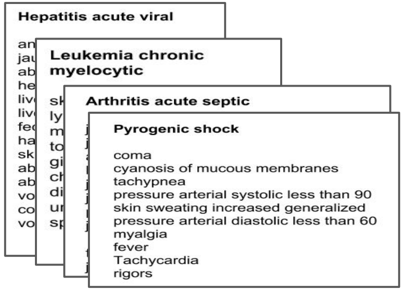

simulated cases

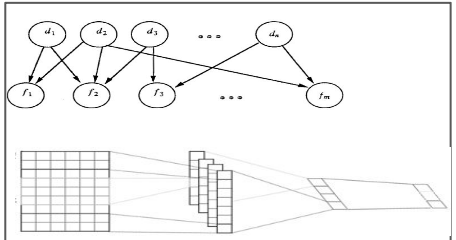

Machine learned models

Our approach: Incorporating data from EHR

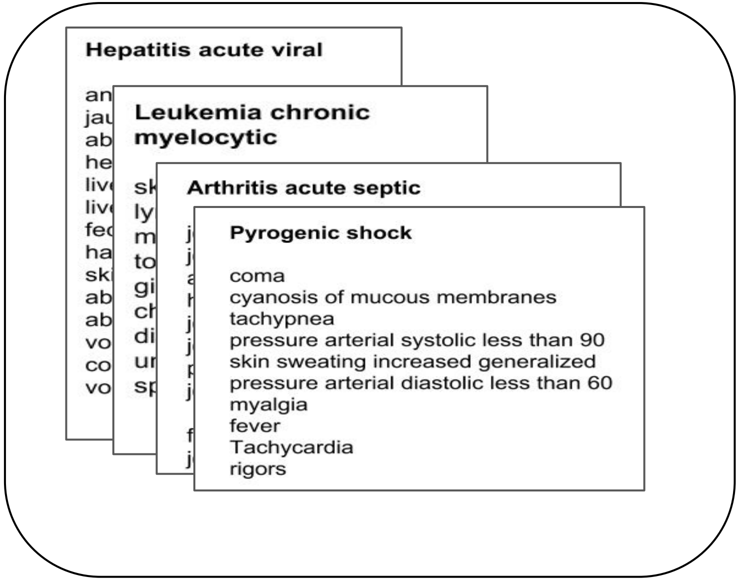

clinical cases simulated from expert systems

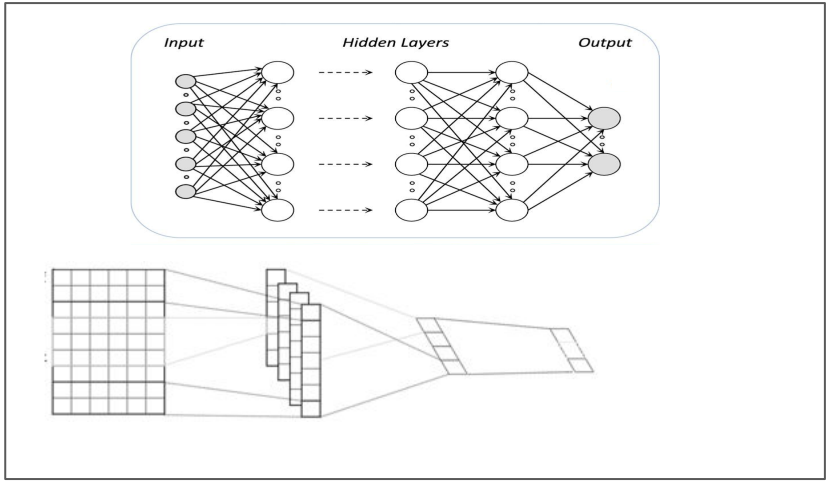

ML classification models for differential diagnosis

clinical cases other sources eg. electronic health records

Part II.
4. Multimodality

Multimodal Inputs

CURAI

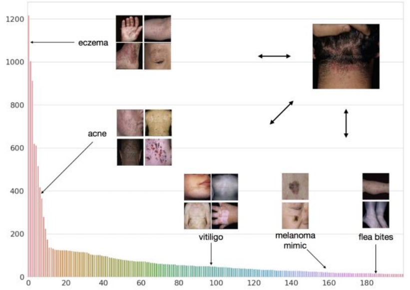

Figure 1: Long-tailed class distribution of Dermnet (shown here for the top-200 classes). Also shown are nearest neighbors to four of the many prototypes learned for select classes using the proposed Prototypical Clustering Network approach. This is illustrative of the huge intra-class variability in the data. For a novel test image, shown at the upper right corner, the model predicts the correct class by measuring weighted similarity to per-class clusters in the embedding space learned through a deep convolutional neural network.

Prototypical Clustering Networks for Dermatological Disease Diagnosis

Viraj Prabhu*[,1] Anitha Kannan[3] Murali Ravuri[3]
Manish Chablani[3] David Sontag[2] Xavier Amatriain[3]
[1]Georgia Tech [2]MIT [3]Curai
virajp@gatech.edu dsontag@csail.mit.edu {anitha, murali, manish, xavier}@curai.com

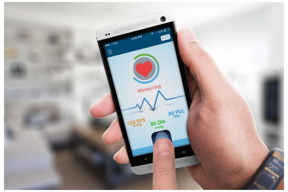

RSNA Radiological Society of North America

Members | Trainees | International | Companies | Media

Search...

Annual Meeting | Membership | Science | Education | Journals | Informatics | R&E Foundation

RSNA News

VISIT US AT RSNA 2017 Booth #2529 That's Radiological
©2017 Bayer Learn more

Archive 2017

Archive 2016

Look Ahead - Machine Learning in Radiology

Noted radiologist Nick Bryan, M.D., Ph.D., discusses the role of Machine Learning (ML) in the radiology practice of tomorrow. BY R. NICK BRYAN, M.D., PH .D.

Bryan

May 1, 2016 | Share | Tweet | Share | Email

Given the many scientific papers, newspaper and magazine articles, and

Insight Data Follow
Insight Fellows Program—Your bridge to a career in Data Science and Data Engineering
Mar 15, 2016 · 6 min read

Detecting heart arrhythmias using machine learning and Apple Watch data

Yancheng is an Insight alumnus from the first Health Data Science session and is now a data scientist at AthenaHealth. While at Insight, he partnered with the UCSF Health eHeart study to detect atrial fibrillation patients using Apple Watch heart rate data. This content originally appeared on his personal website.

Stanford | News

Search Stanford news...

Home | Find Stories | For Journalists | Contact

JANUARY 25, 2017

Deep learning algorithm does as well as dermatologists in identifying skin cancer

In hopes of creating better access to medical care, Stanford researchers have trained an algorithm to diagnose skin cancer.

Conclusions

Recap

CURAI

- More people access healthcare through search than PCPs
- Doctors don’t have time and make mistakes
- We should be able to offer a better online experience than Google + Webmd
- Online Universal Quality Healthcare
    - **Online**, mobile, always on
    - **Universal** = scalable, low cost, easy to access/use (i.e conversational)
    - **Quality** = as good as the best doctors, based on best-known practices and scientific evidence, replicable
    - **Healthcare** = not only a website you need to parse and figure out on your own, but directed access to healthcare resources including human professionals when needed
- How?
    - Online + physicians in the loop = telehealth
    - Standard telehealth is (1) hard to scale, (2) not better quality
    - Scale + Quality = automation (aka AI)

References

References

- “*Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base*” . Shwe et al. 1991.
- “*Computer-assisted diagnostic decision support: history, challenges, and possible paths forward*” Miller. 2009.
- “*Mining Biomedical Ontologies and Data Using RDF Hypergraphs*” Liu et al. 2013.
- “*A ‘Green Button’ For Using Aggregate Patient Data At The Point Of Care*” Longhurst et al. 2014.
- “*Building the graph of medicine from millions of clinical narratives*” Finlayson et al. 2014.
- “*Comparison of Physician and Computer Diagnostic Accuracy*” Semigran et al. 2016.
- “*Clinical Tagging with Joint Probabilistic Models*” . Halpern et al. 2016.
- “*Deep Patient: An Unsupervised Representation to Predict the Future of Patients from EHR*”. Miotto et al. 2016.
- “*Learning a Health Knowledge Graph from Electronic Medical Records*” Rotmensch et al. 2017.
- “*Clustering Patients with Tensor Decomposition*”. Ruffini et al. 2017.
- “*Patient Similarity Using Population Statistics and Multiple Kernel Learning*”. Conroy et al. 2017.
- “*Diagnostic Inferencing via Clinical Concept Extraction with Deep Reinforcement Learning*”. Ling et al. 2017.
- “*Generating Multi-label Discrete Patient Records using Generative Adversarial Networks*” Choi et al. 2017
- *“The Use of Autoencoders for Discovering Patient Phenotypes”*. arXiv.org. Suresh, H., et al. 2017.
- “*Comparative Accuracy of Diagnosis by Collective Intelligence of Multiple Physicians vs Individual Physicians*” Barnett et.al. 2018
- *“Scalable and accurate deep learning with electronic health records”.* Rajkomar et al. 2018
- ***“Learning from the experts: From expert systems to machine-learned diagnosis models”*. Ravuri et al. 2018**
- ***“Prototypical Clustering Networks for Dermatological Disease Diagnosis”*. Prabhu et al. 2018**

CURAI

Curai