Curai health

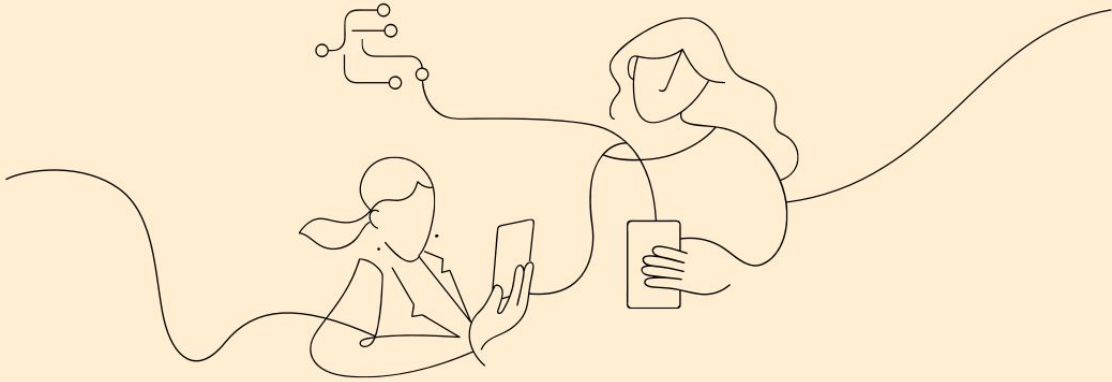

AI for COVID-19: An online virtual care approach


X. Amatriain (@xamat), PhD


Geoff Tso, MD


Anitha Kannan, PhD

http://www.curai.com

COVID-19 and AI Virtual Conference
Stanford 04/01/2020

Curai health

AI-powered virtual care for everyone

Healthcare access and scalability

- >50% world with no access to essential health services
  - ~30% of US adults under-insured
- ~15 min. to capture information, diagnose, recommend treatment
- ⅓ of Americans self-diagnose online

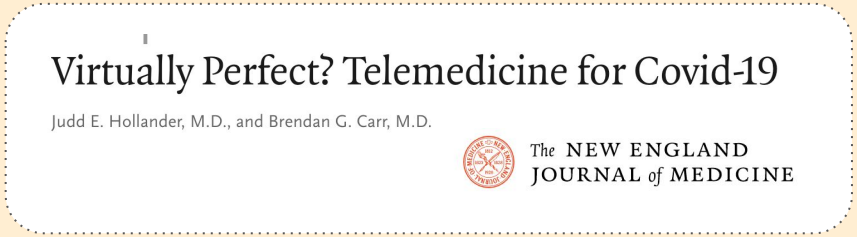

Towards AI powered scalable health systems

- **Mobile-First Care,** always on, accessible, affordable
- **AI + human providers in the loop** for quality care
- Always-**Learning** system
- AI to operate **in-the-wild**

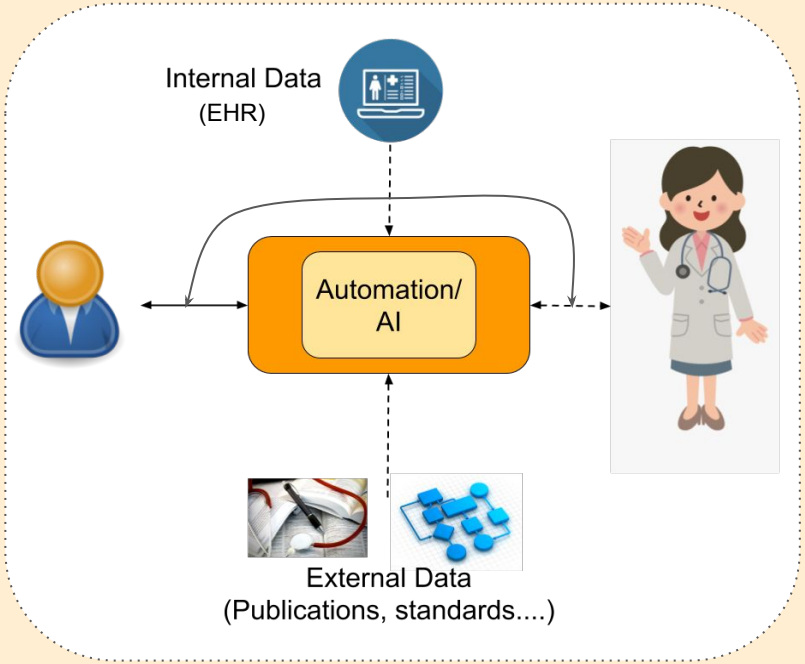

How does this look like?

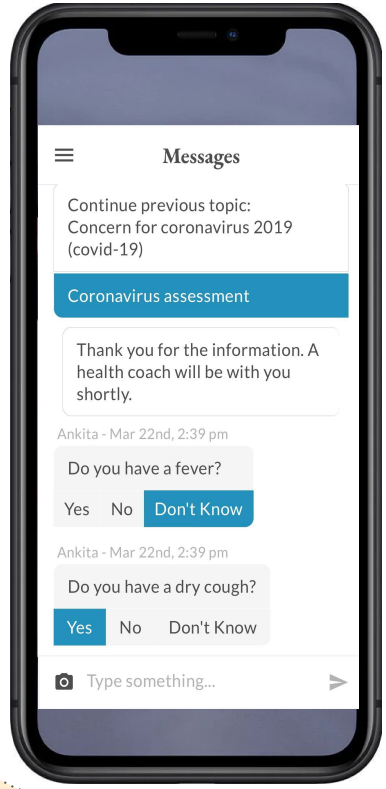

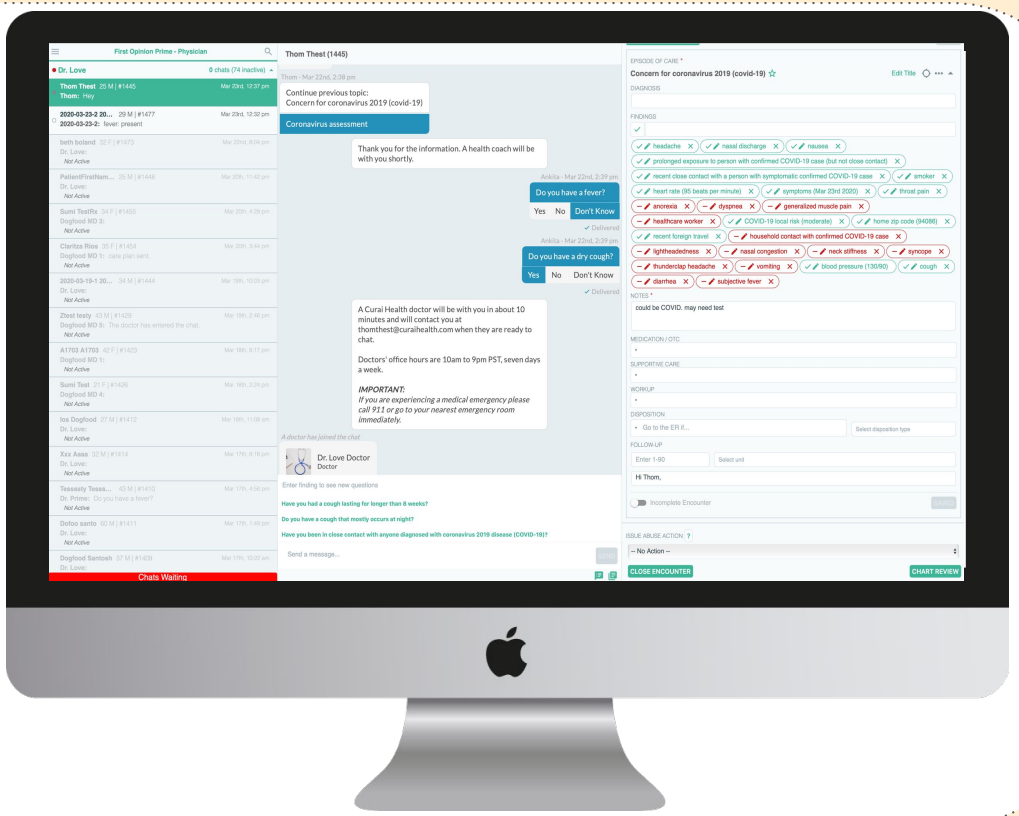

Curai health

Our approach to COVID-19

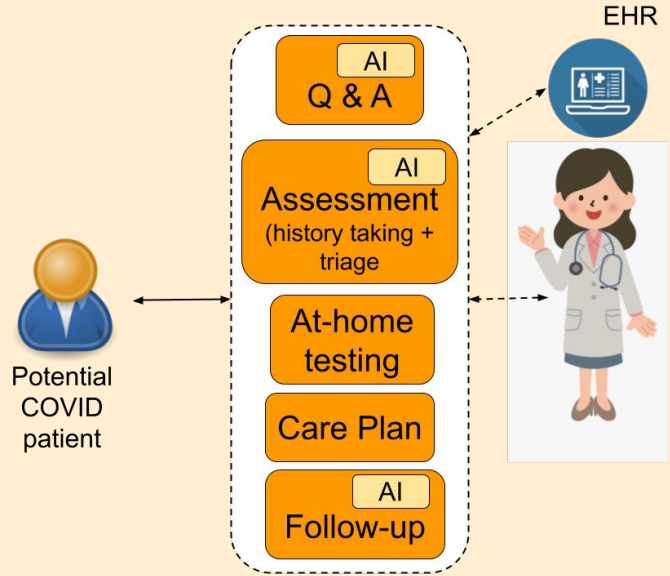

Personalized Diagnostic Assessment

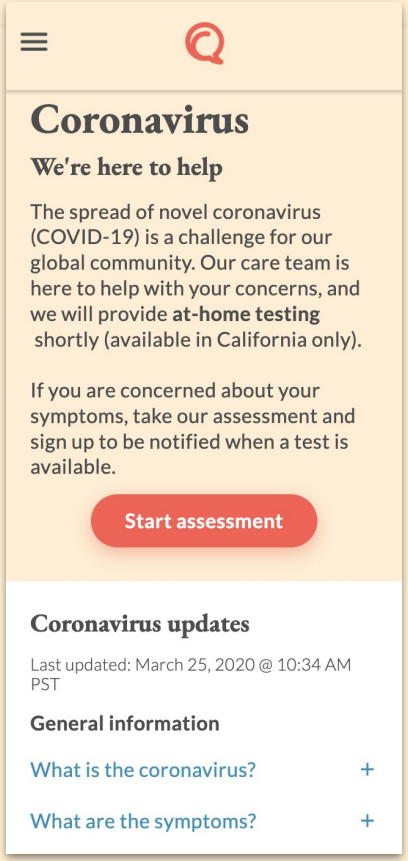

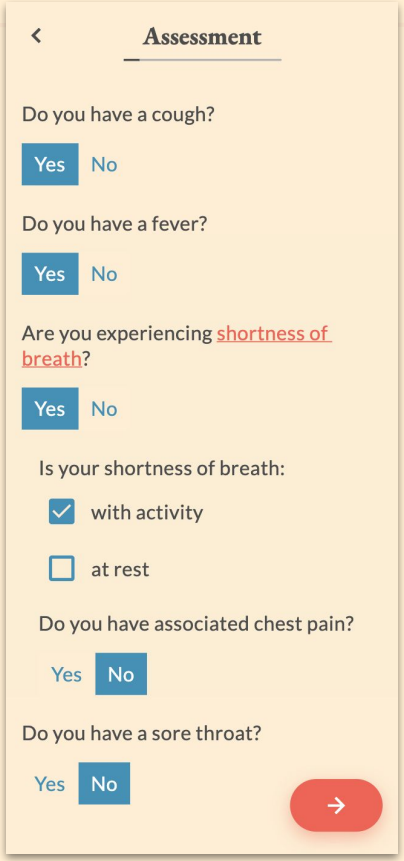

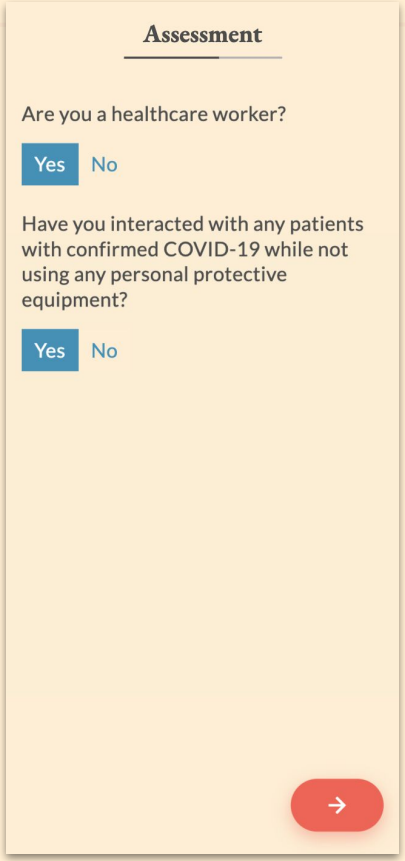

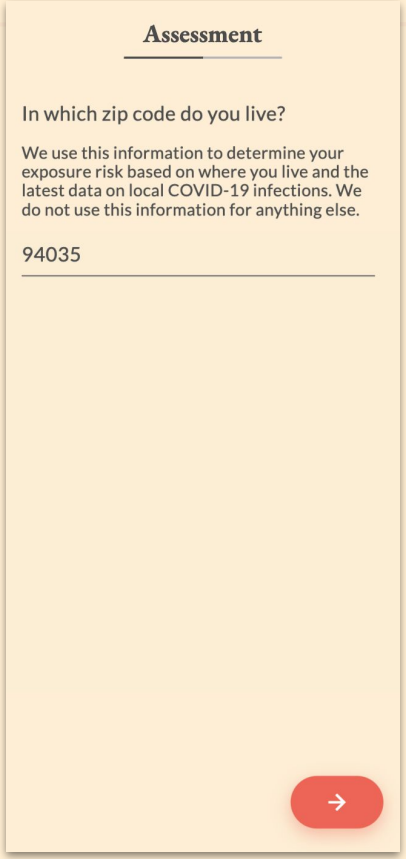

AI + Providers in the loop + Daily follow up

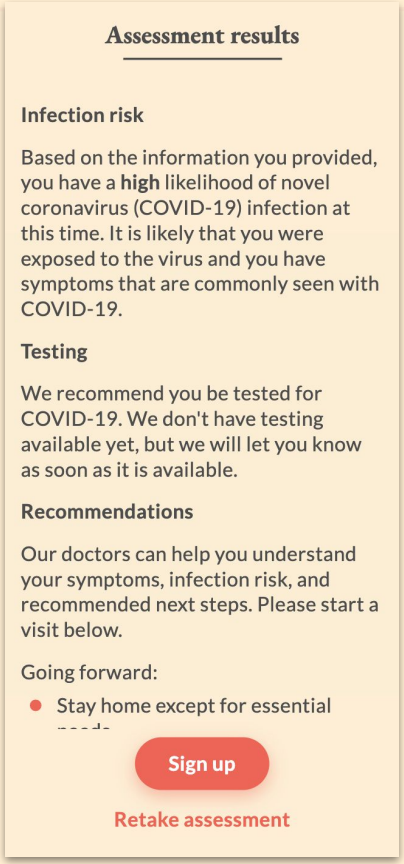

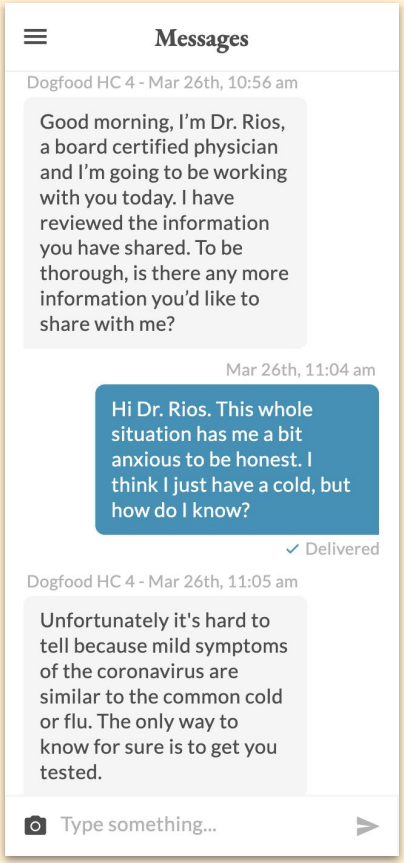

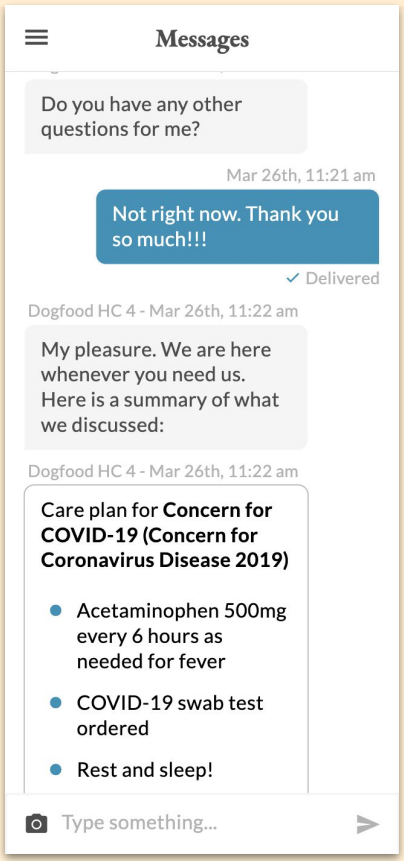

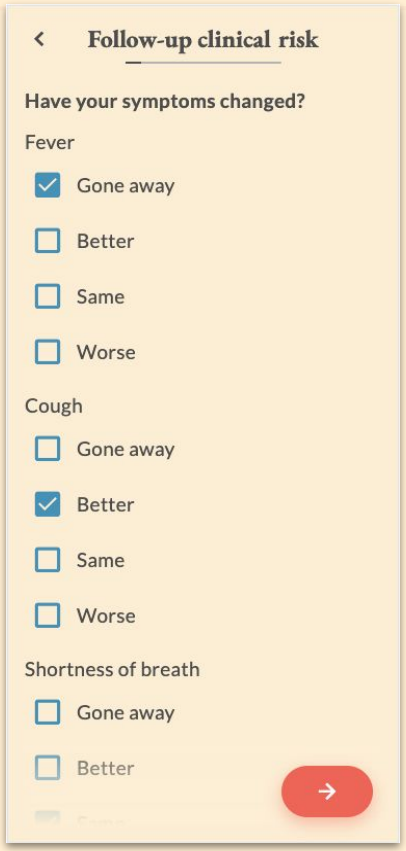

Working on...

- Added at-home testing capabilities to our platform, waiting for FDA to approve
- In discussions with existing health systems who may use platform to scale and improve access

UnitedHealth Group Study Clears Path for Self-Administered COVID-19 Test, Improving Testing Efficiency and Protecting Health Care Workers

- *Simple, scalable patient-collected COVID-19 testing process is as effective as clinician-administered test, according to study of 500 OptumCare patients*
- *Less-invasive testing is better tolerated by patients, reduces exposure for health care workers, and preserves personal protective equipment*
- *FDA has updated its guidance based on UnitedHealth Group data to allow patients to self-administer swab tests for COVID-19*

Start-Ups Jump the Gun on Home Kits for Coronavirus Testing

After a federal warning, companies have stopped marketing kits that let consumers collect their own saliva or throat swabs and send them to labs.

Peer-reviewed research at Curai

Open Set Medical Diagnosis

Viraj Prabhu*,[1] Anitha Kannan[3] Geoffrey J. Tso[3] Namit Katariya[3]
Manish Chablani[3] David Sontag†,[2] Xavier Amatriain[3]
[1]Georgia Tech [2]MIT [3]Curai

Abstract

Machine-learned diagnosis models have shown promise as medical aides but are trained under a closed-set assumption, i.e. that models will only encounter conditions on which they have been trained. However, it is practically infeasible to obtain sufficient training data for every human condition, and once deployed such models will invariably face previously unseen conditions. We frame machine-learned diagnosis as an *open-set* learning problem, and study how state-of-the-art approaches compare. Further, we extend our study to a setting where training data is distributed across several healthcare sites that do not allow data pooling, and experiment with different strategies of building open-set diagnostic ensembles. Across both settings, we observe consistent gains from explicitly modeling unseen conditions, but find the optimal training strategy to vary across settings.

Learning from the experts: From expert systems to machine-learned diagnosis models

Murali Ravuri MURALI@CURAI.COM
Anitha Kannan ANITHA@CURAI.COM
Geoffrey J. Tso GEOFF@CURAI.COM
Xavier Amatriain XAVIER@CURAI.COM

Abstract

Expert diagnostic support systems have been extensively studied. The practical applications of these systems in real-world scenarios have been somewhat limited due to well-understood shortcomings, such as lack of extensibility. More recently, machine-learned models for medical diagnosis have gained momentum, since they can learn and generalize patterns found in very large datasets like electronic health records. These models also have shortcomings - in particular, there is no easy way to incorporate prior knowledge from existing literature or experts. In this paper, we present a method to merge both approaches by using expert systems as generative models that create simulated data on which models can be learned. We demonstrate that such a learned model not only preserves the original properties of the expert systems but also addresses some of their limitations. Furthermore, we show how this approach can also be used as the starting point to combine expert knowledge with knowledge extracted from other data sources, such as electronic health records.

Classification as Decoder: Trading Flexibility for Control in Medical Dialogue

Sam Shleifer[1] Manish Chablani[2]
Anitha Kannan[2] Namit Katariya[2] Xavier Amatriain[2]
[1]Stanford [2]Curai

Abstract

Generative seq2seq dialogue systems are trained to predict the next word in dialogues that have already occurred. They can learn from large unlabeled conversation datasets, build a deeper understanding of conversational context, and generate a wide variety of responses. This flexibility comes at the cost of control, a concerning tradeoff in doctor/patient interactions. Inaccuracies, typos, or undesirable content in the training data will be reproduced by the model at inference time. We trade a small amount of labeling effort and some loss of response variety in exchange for quality control. More specifically, a pretrained language model encodes the conversational context, and we finetune a classification head to map an encoded conversational context to a response class, where each class is a noisily labeled group of interchangeable responses. Experts can update these exemplar responses over time as best practices change without retraining the classifier or invalidating old training data. Expert evaluation of 775 unseen doctor/patient conversations shows that only 12% of the discriminative model's responses are worse than the what the doctor ended up writing, compared to 18% for the generative model.

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

Abstract

We consider the problem of image classification for the purpose of aiding doctors in dermatological diagnosis. Dermatological diagnosis poses two major challenges for standard off-the-shelf techniques: First, the data distribution is typically extremely long tailed. Second, intra-class variability is often large. To address the first issue, we formulate the problem as low-shot learning, where once de-

surge in online services and telemedicine for closing the gap of healthcare access, these services also have similar problems [20]. The need to find effective solutions to *aid* doctors in accurate diagnosis motivates this work.

Why is diagnosis of skin conditions hard for doctors? One important factor is the sheer number of dermatological conditions. The International Classification of Disease 10 (ICD 10) classification of human disease[1] enumerates more than 1000 skin or skin-related illnesses. However, most

Domain-Relevant Embeddings for Medical Question Similarity

Clara H. McCreery[1] Namit Katariya[2] Anitha Kannan[2] Manish Chablani[2]
Xavier Amatriain[2]
[1]Stanford [2]Curai

Abstract

The rate at which medical questions are asked online far exceeds the capacity of qualified people to answer them, and many of these questions are not unique. Identifying same-question pairs could enable questions to be answered more effectively. While many research efforts have focused on the problem of general question similarity for non-medical applications, these approaches do not generalize well to the medical domain, where medical expertise is often required to determine semantic similarity. In this paper, we show how a semi-supervised approach of pre-training a neural network on *medical question-answer pairs* is a particularly useful intermediate task for the ultimate goal of determining medical question similarity. While other pre-training tasks yield an accuracy below 78.7% on this task, our model achieves an accuracy of 82.6% with the same number of training examples, and an accuracy of 80.0% with a much smaller training set.

The accuracy vs. coverage trade-off in patient-facing diagnosis models

Anitha Kannan, PhD[1], Jason Alan Fries, PhD[2], Eric Kramer, MD[1],
Jen Jen Chen, MD[1], Nigam Shah, MBBS, PhD[2] and Xavier Amatriain, PhD[1]
[1]Curai, Palo Alto, CA, USA [2]Stanford University, Palo Alto, CA, USA

Abstract *A third of adults in America use the Internet to diagnose medical concerns, and online symptom checkers are increasingly part of this process. These tools are powered by diagnosis models similar to clinical decision support systems, with the primary difference being the coverage of symptoms and diagnoses. To be useful to patients and physicians, these models must have high accuracy while covering a meaningful space of symptoms and diagnoses. To the best of our knowledge, this paper is the first in studying the trade-off between the coverage of the model and its performance for diagnosis. To this end, we learn diagnosis models with different coverage from EHR data. We find a 1% drop in top-3 accuracy for every 10 diseases added to the coverage. We also observe that complexity for these models does not affect performance, with linear models performing as well as neural networks.*

Question similarity

- Transfer learning
- Double-finetune BERT model
  - Handle data sparsity
  - Medical domain knowledge through an intermediate QA binary task
- Out-of-the-box applied to COVID specific Q/A

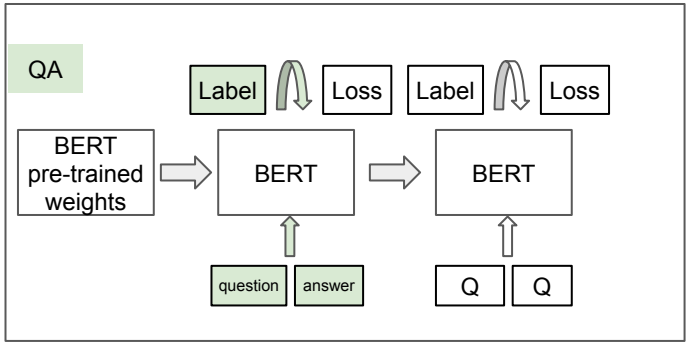

Effective Transfer Learning for Identifying Similar Medical Questions

Clara H. McCreery
mccreery@cs.stanford.edu

Namit Katariya
namit@curai.com

Anitha Kannan
anitha@curai.com

Manish Chablani
manish@curai.com

Xavier Amatriain
xavier@curai.com

Automated Question/Answering

- Automated question answering via question similarity
- Practitioners' vetted answers with ability to follow up

User Questions	Matching Questions in our FAQ (Answer not shown here for brevity)
When do COVID symptoms start after exposure?	How long is it between when a person is exposed to the virus and when they start showing symptoms?
I am asymptomatic and have been social distancing/self-isolating for the past x days. Can I still transmit the infection?	How can someone pass along coronavirus when asymptomatic? If not sneezing or coughing, how can they infect others?
Currently I'm experiencing a **cough** and **slight chest pain**. Should I just stay at home? At what point will I know I have to go to the ER?	When should you go to the emergency room? 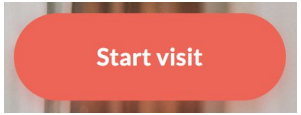

ML + Expert systems for Dx models

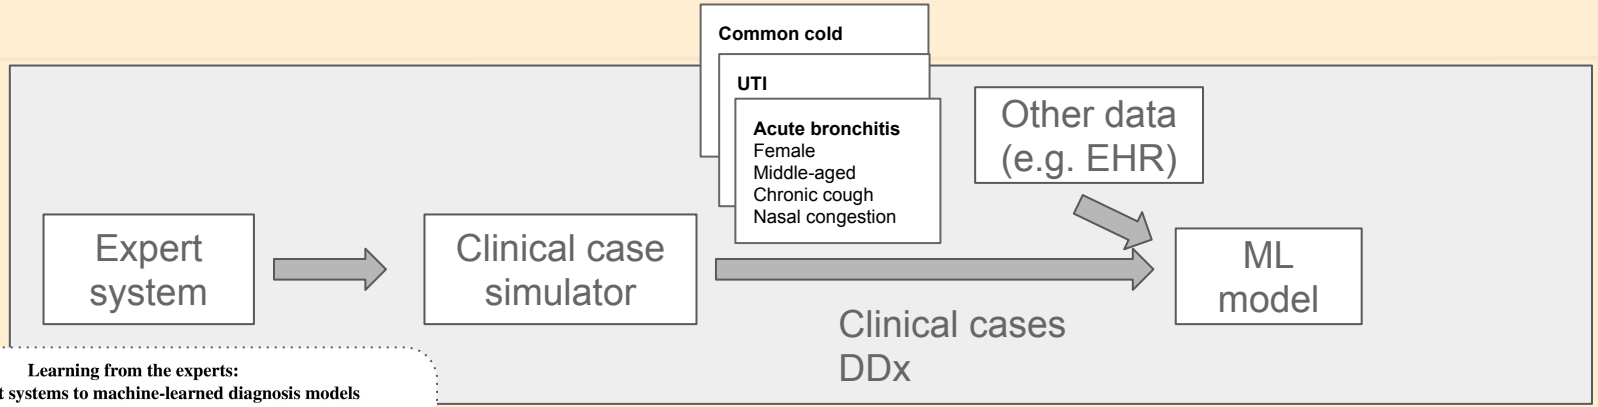

Learning from the experts:
From expert systems to machine-learned diagnosis models

Murali Ravuri MURALI@CURAI.COM
Anitha Kannan ANITHA@CURAI.COM
Geoffrey J. Tso GEOFF@CURAI.COM
Xavier Amatriain XAVIER@CURAI.COM

Inputs

female
middle aged
fever
cough

DDx with expert systems

Influenza 16.9
bacterial pneumonia 16.9
acute sinusitis 10.9
asthma 10.9
common cold 10.9

DDx with ML model

influenza 0.753
bacterial pneumonia 0.205
asthma 0.017
acute sinusitis 0.008
pulmonary tuberculosis 0.007

Feedback loop

Inputs	*DDx before COVID*	*DDx after COVID*
female middle aged fever cough nasal congestion	influenza 0.634 adenovirus infection 0.159 bacterial pneumonia 0.114 acute sinusitis 0.05 asthma 0.019	influenza 0.512 **COVID-19** 0.256 adenovirus infection 0.106 bacterial pneumonia 0.069 acute sinusitis 0.026
female middle aged fever cough healthcare worker	influenza 0.753 bacterial pneumonia 0.205 asthma 0.017 acute sinusitis 0.008 pulmonary tuberculosis 0.007	**COVID-19** 0.913 influenza 0.048 bacterial pneumonia 0.024 pulmonary tuberculosis 0.004 adenovirus infection 0.003

Conclusions

- Healthcare needs to scale quickly, and this has become obvious in a global pandemic like the one we are facing
- The only way to scale healthcare while improving quality and accessibility is through technology and AI
- AI cannot be simply “dropped” in the middle of old workflows and approaches
    - It needs to be integrated in end-to-end medical care benefitting both patients and providers

Curai health