Artificial Intelligenec and its role in clinical contexts

Samira A.Rahimi Eng, Ph. D.
Assistant Professor
Department of Family Medicine

McGill | Department of Family Medicine

Outline

Rahimi,S
McGill

What is AI?

What has been done in clinical contexts

Future

- **Recommendations**

What is Artificial Intelligence?

Rahimi,S
McGill

What is Artificial Intelligence?

Thinking/ Reasoning

Thinking Humanly

Thinking Ra

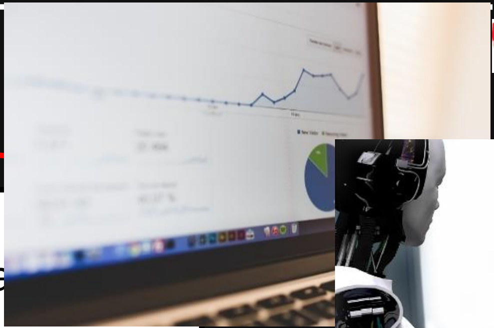

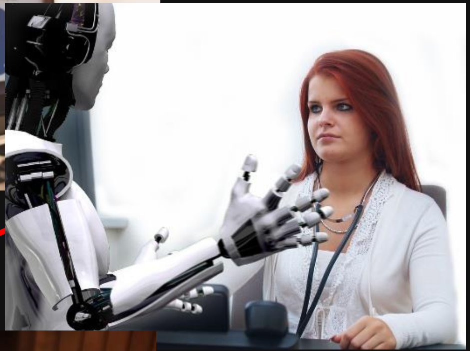

Acting/ Behaving

Acting Humanly

Acting Ratio

Russell & Norvig, 2010

Rahimi,S

McGill

What is Artificial Intelligence?

Thinking/

Behaving

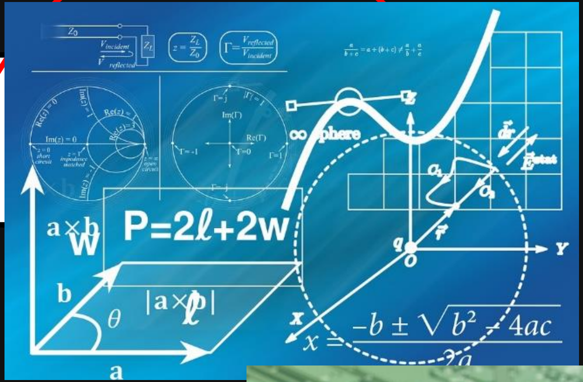

Acting Rat

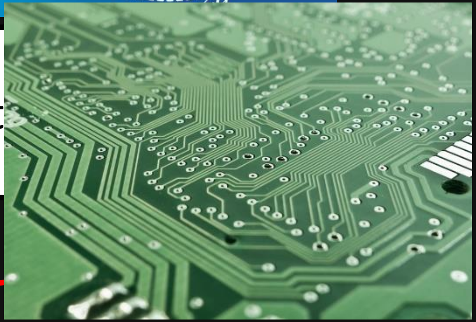

Russell & Norvig, 2010

Rahimi,S
McGill

What is Artificial Intelligence?

Thinking/ Reasoning

Thinking Humanly

"The exc
mal
think...ma
(Hau

Thinking Rationally

"The study of mental
ugh the use of
nal models"
McDermott,
85)

Acting/ Behaving

Actin

"The study o
that perform functions that require intelligence when performed by people" (Kurzweil, 1990)

g Rationally

oncerned with intelligent behavior in artifacts" (Nilsson, 1998)

Russell & Norvig, 2010

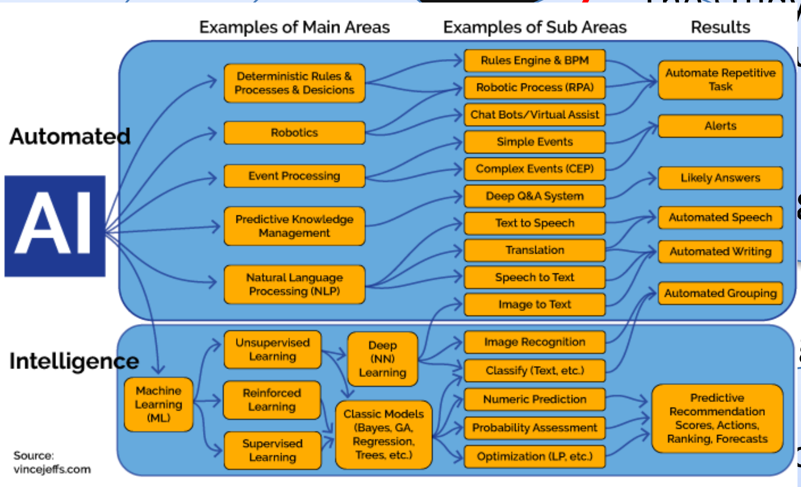

Brief history of AI:

A LOGICAL CALCULUS OF THE IDEAS IMMANENT IN NERVOUS ACTIVITY

WARREN S. MCCULLOCH and WALTER H. PITTS

A. M. Turing (1950) Computing Machinery and Intelligence. ***Mind 49:*** **433-460.**

COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous, If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

1943

1949

1950

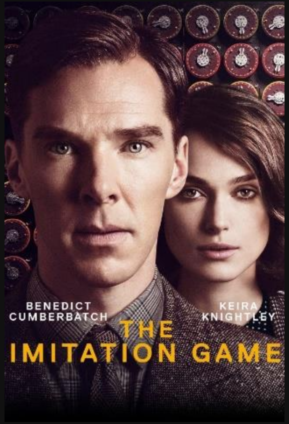

Rahimi,S
McGill

The birth of AI (1956):

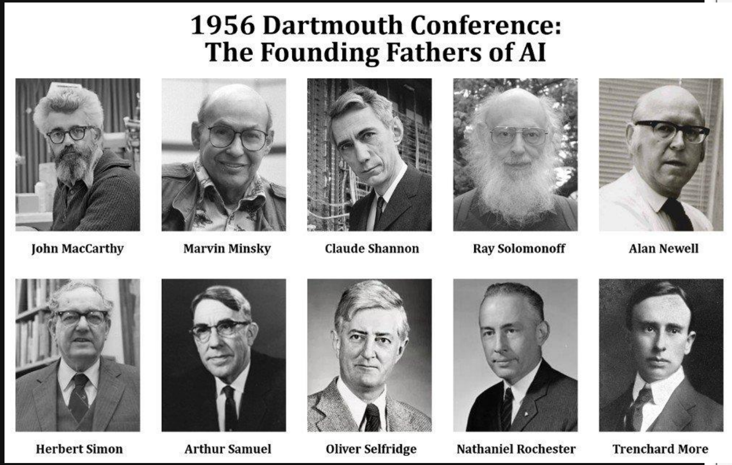

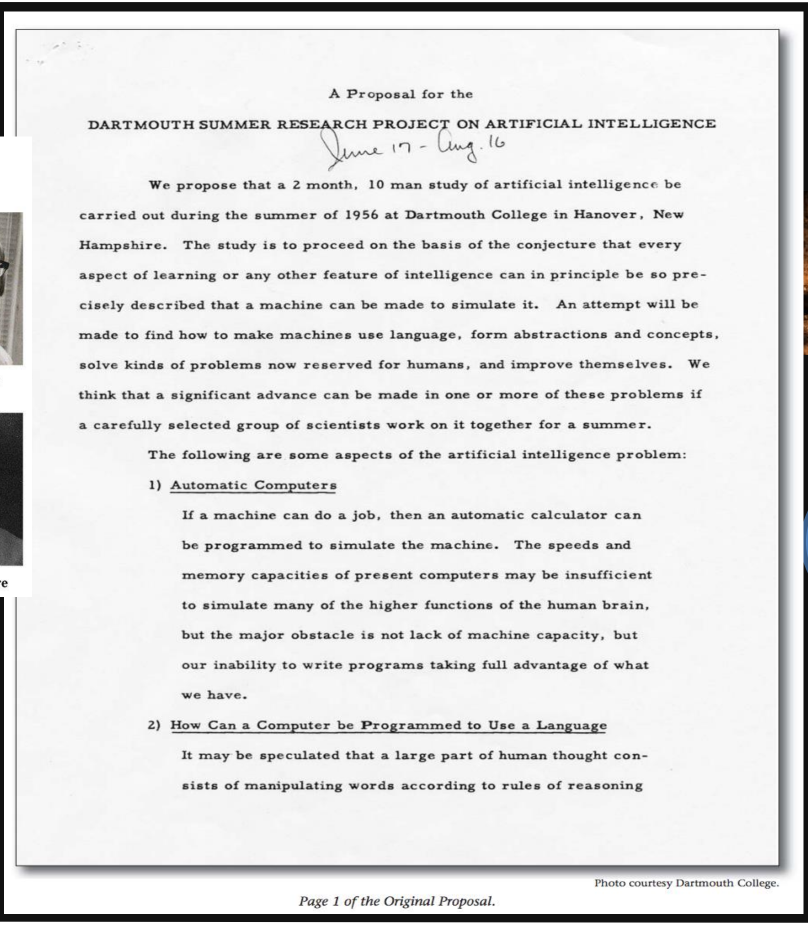

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

Photo courtesy Dartmouth College.

Page 1 of the Original Proposal.

1956

Rahimi,S

McGill

Very large data sets:

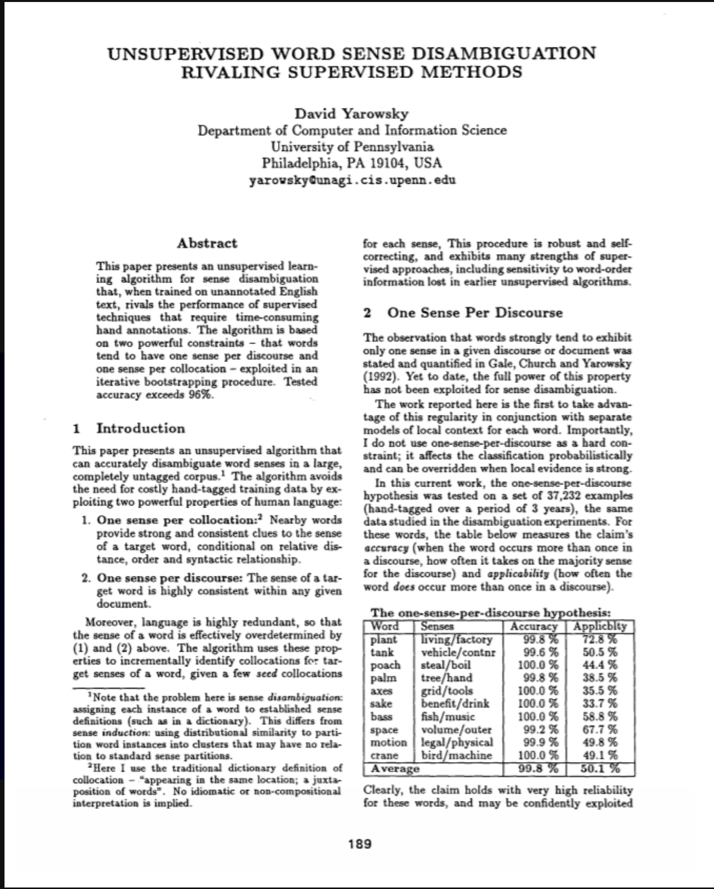

UNSUPERVISED WORD SENSE DISAMBIGUATION RIVALING SUPERVISED METHODS

David Yarowsky
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104, USA
yarowsky@unagi.cis.upenn.edu

Abstract

This paper presents an unsupervised learning algorithm for sense disambiguation that, when trained on unannotated English text, rivals the performance of supervised techniques that require time-consuming hand annotations. The algorithm is based on two powerful constraints – that words tend to have one sense per discourse and one sense per collocation – exploited in an iterative bootstrapping procedure. Tested accuracy exceeds 96%.

1 Introduction

This paper presents an unsupervised algorithm that can accurately disambiguate word senses in a large, completely untagged corpus.[1] The algorithm avoids the need for costly hand-tagged training data by exploiting two powerful properties of human language:

1. **One sense per collocation:**[2] Nearby words provide strong and consistent clues to the sense of a target word, conditional on relative distance, order and syntactic relationship.
2. **One sense per discourse:** The sense of a target word is highly consistent within any given document.

Moreover, language is highly redundant, so that the sense of a word is effectively overdetermined by (1) and (2) above. The algorithm uses these properties to incrementally identify collocations for target senses of a word, given a few *seed* collocations for each sense, This procedure is robust and self-correcting, and exhibits many strengths of supervised approaches, including sensitivity to word-order information lost in earlier unsupervised algorithms.

2 One Sense Per Discourse

The observation that words strongly tend to exhibit only one sense in a given discourse or document was stated and quantified in Gale, Church and Yarowsky (1992). Yet to date, the full power of this property has not been exploited for sense disambiguation.

The work reported here is the first to take advantage of this regularity in conjunction with separate models of local context for each word. Importantly, I do not use one-sense-per-discourse as a hard constraint; it affects the classification probabilistically and can be overridden when local evidence is strong.

In this current work, the one-sense-per-discourse hypothesis was tested on a set of 37,232 examples (hand-tagged over a period of 3 years), the same data studied in the disambiguation experiments. For these words, the table below measures the claim's *accuracy* (when the word occurs more than once in a discourse, how often it takes on the majority sense for the discourse) and *applicability* (how often the word *does* occur more than once in a discourse).

The one-sense-per-discourse hypothesis:

Word	Senses	Accuracy	Applicblty
plant	living/factory	99.8 %	72.8 %
tank	vehicle/contnr	99.6 %	50.5 %
poach	steal/boil	100.0 %	44.4 %
palm	tree/hand	99.8 %	38.5 %
axes	grid/tools	100.0 %	35.5 %
sake	benefit/drink	100.0 %	33.7 %
bass	fish/music	100.0 %	58.8 %
space	volume/outer	99.2 %	67.7 %
motion	legal/physical	99.9 %	49.8 %
crane	bird/machine	100.0 %	49.1 %
Average		99.8 %	50.1 %

Clearly, the claim holds with very high reliability for these words, and may be confidently exploited

[1] Note that the problem here is sense *disambiguation*: assigning each instance of a word to established sense definitions (such as in a dictionary). This differs from sense *induction*: using distributional similarity to partition word instances into clusters that may have no relation to standard sense partitions.

[2] Here I use the traditional dictionary definition of collocation – "appearing in the same location; a juxtaposition of words". No idiomatic or non-compositional interpretation is implied.

1995

START

What is AI?

What has been done in clinical contexts

Future

Q & A

What can AI do?

Robatics

Game playing

Speech recognition

Logistic planning

...

Autonomous planning and scheduling

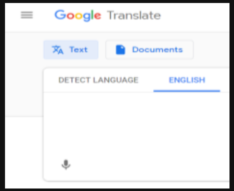

Machine Translation

Rahimi,S
McGill

Outline

What is AI?

What has been done in clinical contexts

Future

- **Recommendations**

Rahimi,S
McGill

Clinical contexts

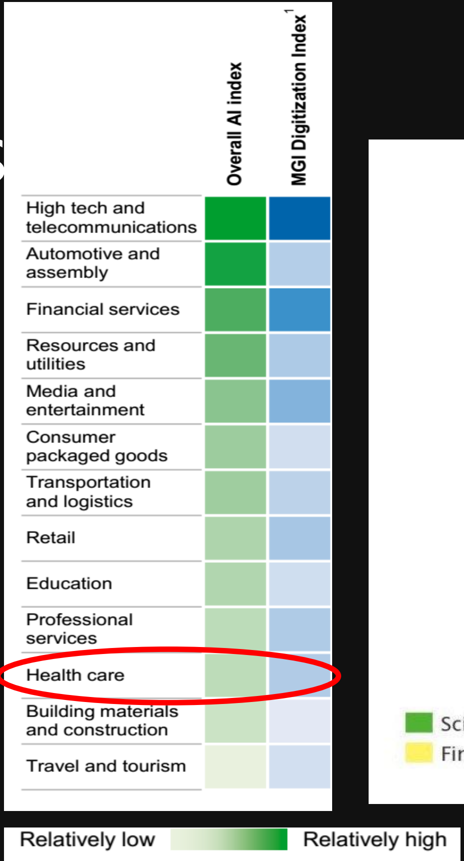

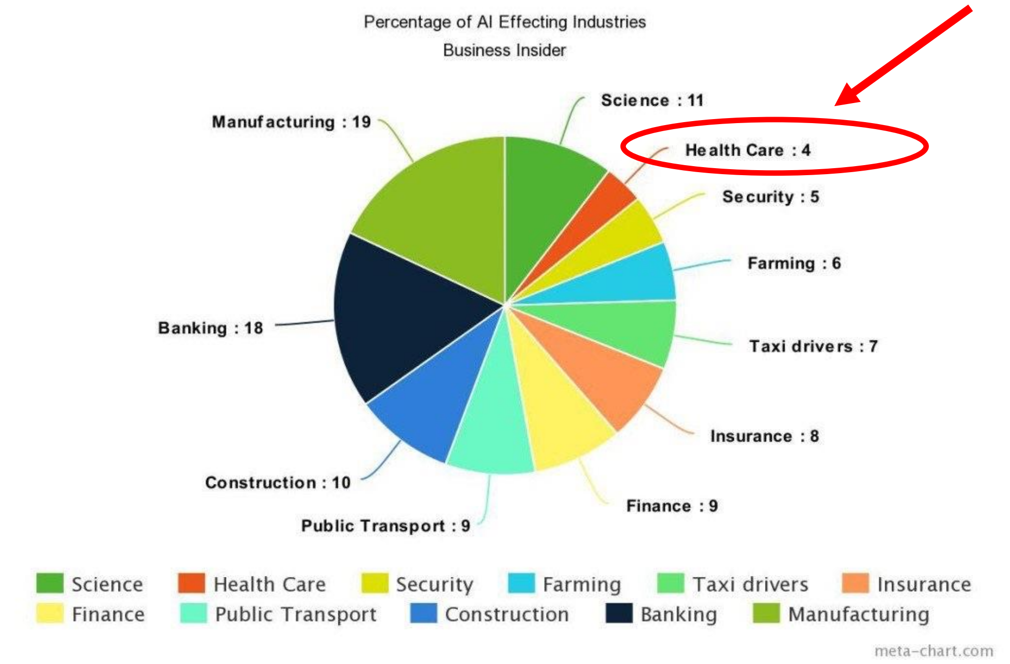

McKinsey Global Institute (MGI), 2017

Rahimi,S
McGill

What AI can do in Clinical contexts?

Improving the quality of training and Patient-doctor relationship

Address health & wellbeing gap

1.
2.

Ad... ga...

1. F...
2. F...

Health care	▪ **$300 billion possible savings** in the United States using machine learning tools for population health forecasting ▪ **£3.3 billion possible savings** in the United Kingdom using AI to provide preventive care and reduce nonelective hospital admissions	▪ **30–50% productivity improvement** for nurses supported by AI tools ▪ **Up to 2% GDP savings** for operational efficiencies in developed countries

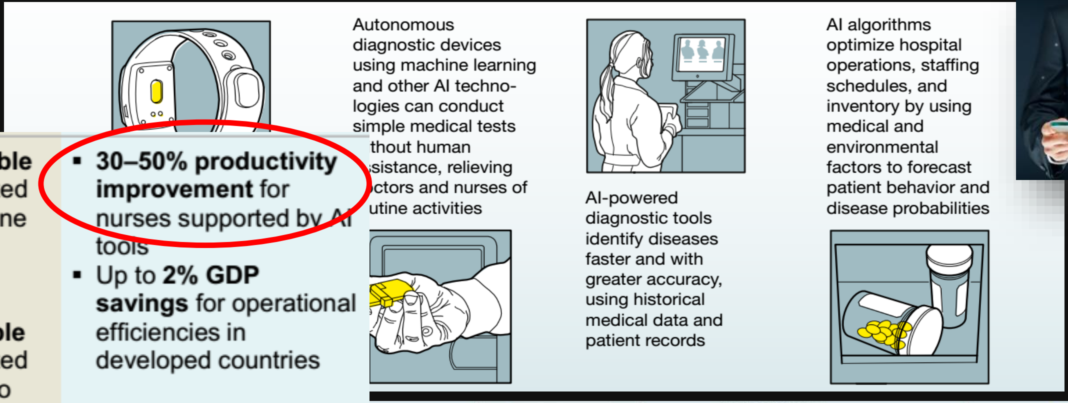

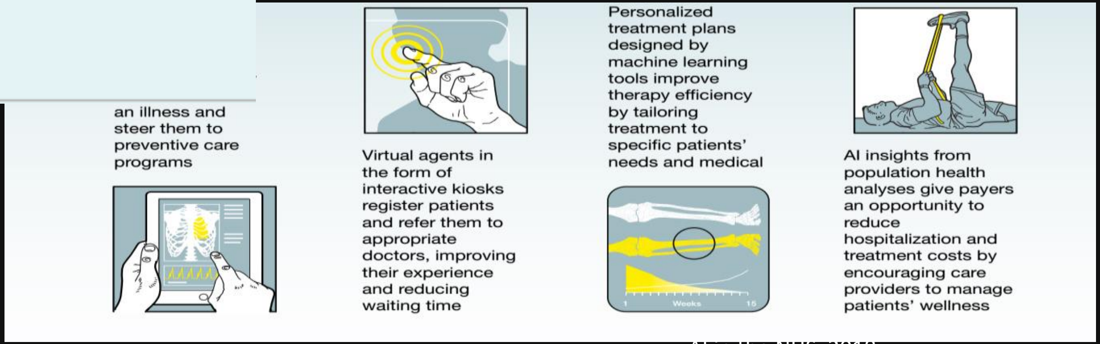

Address the care & quality gap

1. Improved diagnostics (e.g. 30x faster & accurate in mammograms)
2. Personalized Treatment

McKinsey Global Institute (MGI), 2017

AI in the NHS, 2018
Patel et al, 2017, Cancer: 123 (1)

Dermatology:

Rahimi,S

McGill

13000

LETTER

doi:10.1038/nature21056

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva[1*], Brett Kuprel[1*], Roberto A. Novoa[2,3], Justin Ko[2], Susan M. Swetter[2,4], Helen M. Blau[5] & Sebastian Thrun[6]

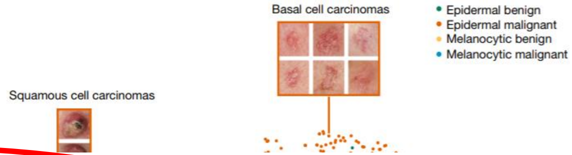

Extended Data Table 2 | General validation results

a.

Classifier	Three-way accuracy
Dermatologist 1	65.6%
Dermatologist 2	66.0%
CNN	69.4 ± 0.8%
CNN - PA	**72.1 ± 0.9%**

b.

Classifier	Nine-way accuracy
Dermatologist 1	53.3%
Dermatologist 2	55.0%
CNN	48.9 ± 1.9%
CNN - PA	**55.4 ± 1.7%**

c. **Disease classes: three-way classification**

0. Benign single lesions
1. Malignant single lesions
2. Non-neoplastic lesions

d. **Disease classes: nine-way classification**

0. Cutaneous lymphoma and lymphoid infiltrates
1. Benign dermal tumors, cysts, sinuses
2. Malignant dermal tumor
3. Benign epidermal tumors, hamartomas, milia, and growths
4. Malignant and premalignant epidermal tumors
5. Genodermatoses and supernumerary growths
6. Inflammatory conditions
7. Benign melanocytic lesions
8. Malignant Melanoma

ıds represent the different disease ithm clusters the diseases. Insets show points. Images reprinted with permission ırary (https://licensing.eri.ed.ac.uk/i/ html).

94% & 96%

Accuracy rate

Cardiology:

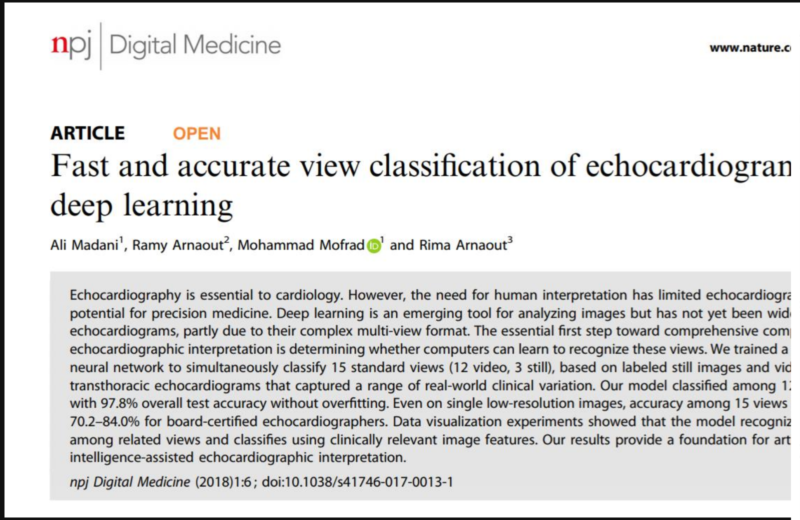
npj | Digital Medicine

www.nature.c

ARTICLE OPEN

Fast and accurate view classification of echocardiogran deep learning

Ali Madani[1], Ramy Arnaout[2], Mohammad Mofrad[1] and Rima Arnaout[3]

Echocardiography is essential to cardiology. However, the need for human interpretation has limited echocardiogra potential for precision medicine. Deep learning is an emerging tool for analyzing images but has not yet been wid echocardiograms, partly due to their complex multi-view format. The essential first step toward comprehensive com echocardiographic interpretation is determining whether computers can learn to recognize these views. We trained a neural network to simultaneously classify 15 standard views (12 video, 3 still), based on labeled still images and vid transthoracic echocardiograms that captured a range of real-world clinical variation. Our model classified among 1 with 97.8% overall test accuracy without overfitting. Even on single low-resolution images, accuracy among 15 views 70.2–84.0% for board-certified echocardiographers. Data visualization experiments showed that the model recogniz among related views and classifies using clinically relevant image features. Our results provide a foundation for art intelligence-assisted echocardiographic interpretation.

npj Digital Medicine (2018)1:6 ; doi:10.1038/s41746-017-0013-1

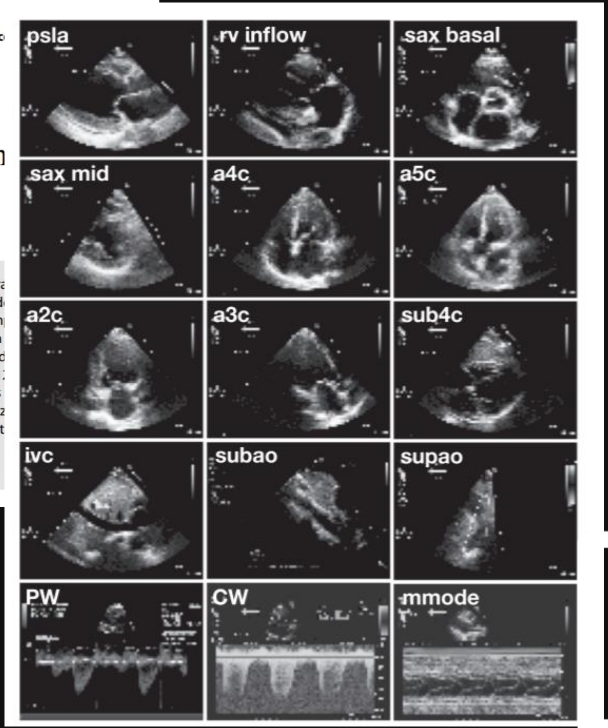

91% vs 70-84%

Accuracy rate (comparing)

Rahimi,S

McGill

Rheumatology :
DSS for Diagnosis of Rheumatoid Arthrit

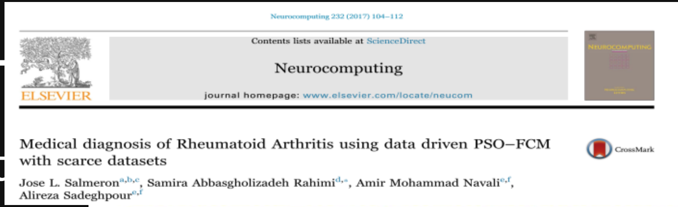
Neurocomputing 232 (2017) 104–112
Contents lists available at ScienceDirect
Neurocomputing
journal homepage: www.elsevier.com/locate/neucom

Medical diagnosis of Rheumatoid Arthritis using data driven PSO–FCM with scarce datasets

Jose L. Salmeron[a,b,c], Samira Abbasgholizadeh Rahimi[d,*], Amir Mohammad Navali[e,f], Alireza Sadeghpour[e,f]

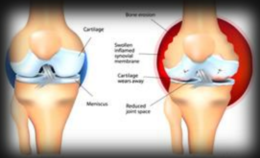

OUR TEAM

- Rheumatologists
- Orthopedic surgeons
- System Engineer
- Software experts

Three countries

- Canada
- Spain
- Iran

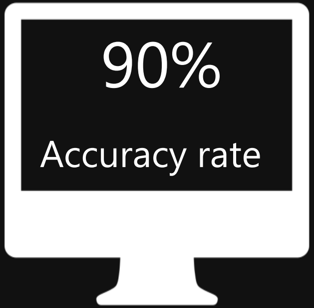

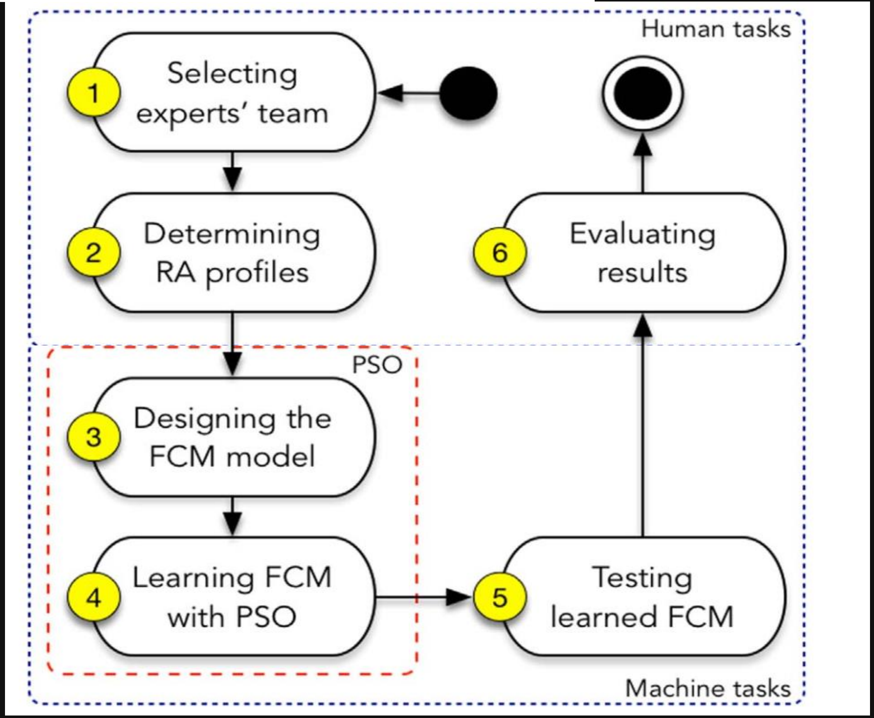

Rahimi,S
McGill

FDA approved AI tools (2017-2018):

Company	FDA Approval	Indication
Apple	September 2018	Atrial fibrillation detection
Aidoc	August 2018	CT brain bleed diagnosis
iCAD	August 2018	Breast density via mammography
Zebra Medical	July 2018	Coronary calcium scoring
Bay Labs	June 2018	Echocardiogram EF determination
Neural Analytics	May 2018	Device for paramedic stroke diagnosis
IDx	April 2018	Diabetic retinopathy diagnosis
Icometrix	April 2018	MRI brain interpretation
Imagen	March 2018	X-ray wrist fracture diagnosis
Viz.ai	February 2018	CT stroke diagnosis
Arterys	February 2018	Liver and lung cancer (MRI, CT) diagnosis
MaxQ-AI	January 2018	CT brain bleed diagnosis
Alivecor	November 2017	Atrial fibrillation detection via Apple Watch
Arterys	January 2017	MRI heart interpretation

X-ray

CT

MRI

Topol, E., 2019

Outline

What is AI?

What has been done in clinical contexts

Future
- **Recommendations**

Rahimi,S
McGill

Recommendations

Recommendation 1 (Data):

Rahimi,S
McGill

Recommendation 2 (Cross-deciplinary collaborations):

Rahimi,S

Rahimi,S
McGill

Recommendation 3

(Education, training and awareness building):

- Culture of learning
- Support educators
- Education of current HCPs
- Education of future HCPs (university level)

Rahimi,S
McGill

Recommendation 4

(Education, training and awareness building):

Recommendation 5

(Patient and public engagement):

Rahimi,S
McGill

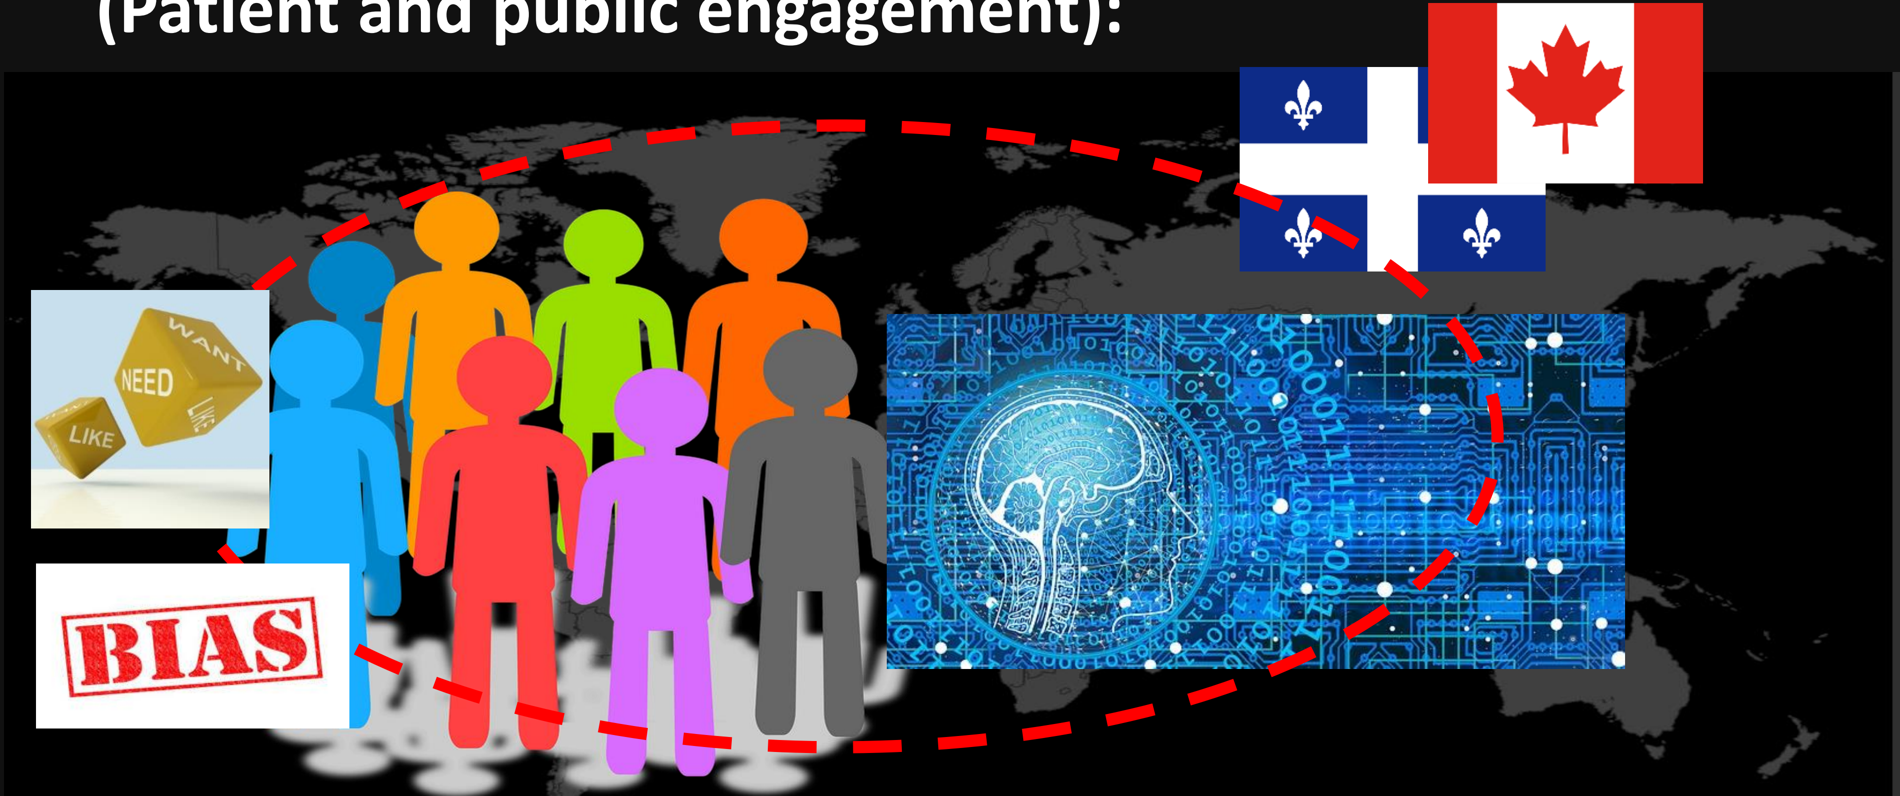

Merci!!

Email: samira.rahimi@mcgill.ca