Introduction to Artificial Intelligence (AI) Applications for Healthcare

Sambath Parthasarathy PhD

sambath.narayanan@dataeverconsulting.com

www.dataeverconsulting.com


76 BPM
10:09
24sec
It helps to rest your arms on a table or your legs.

Apple Watch


76 BPM
10:09
24sec
It helps to rest your arms on a table or your legs.

Apple Watch

Approved by the US FDA in September of last year, the ECG technology on the Apple Watch Series

This technology is the first consumer-available product that allows users to take an ECG from their own wrist,

and can provide critical data to physicians.

AFib is the most common heart arrhythmia experienced in patients

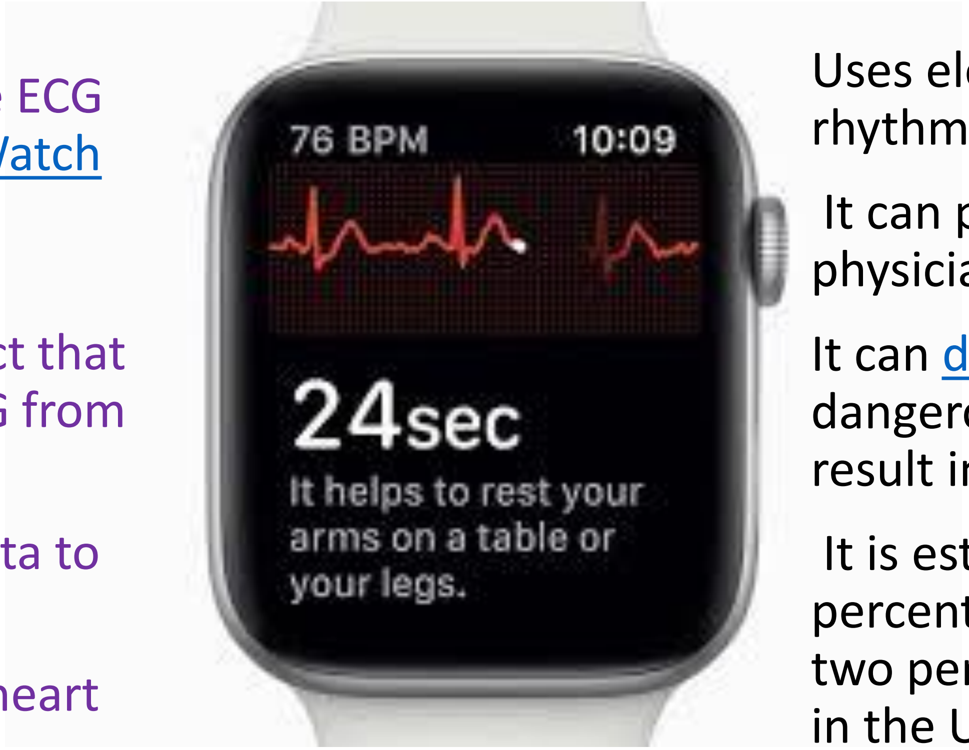

Uses electrodes to capture heart rhythm irregularities.

It can provide critical data to physicians.

It can detect atrial fibrillation, a dangerous arrhythmia that can result in stroke if left untreated.

It is estimated to affect up to nine percent of those over the age of 65, two percent of younger populations in the U.S.

AI/ML Deep Learning applications

- Computer Vision
  - Object Detection
  - Classification
  - Segmentation
  - Visual Q & A
- Speech and Audio
  - Automatic Speech Recognition
  - Generation
  - Processing
  - Denoising
- Natural Language Processing(NLP)
  - Sentiment Analysis
  - Search
  - Recommender Systems
  - Q & A
  - Language Translation

Healthcare industry challenges

- **Radiologist** is a medical practitioner professionally trained in analyzing MRI or CT scan, X Ray and using the analysis for diagnosing medical conditions

- The time it takes to analyze these scans, combined with the sheer number of scans that healthcare professionals have to go through (over 1,000 a day at one location alone), can lead to lengthy delays between scan and treatment – even when someone needs urgent care.

AI Healthcare market size and trends

“Public and private sector investment in healthcare AI is expected to reach $6.6 billion by 2021”

“Top AI applications may result in annual savings of $150 billion by 2026 “

–Accenture

“…35% to 40% of hospital operating rooms in the United States will be integrated with AI and virtual reality applications by 2022.“

– Frost&Sullivan

https://www.forbes.com/sites/insights-intelai/2019/02/11/ai-and-healthcare-a-giant-opportunity/#30c1c7014c68
https://go.frost.com/NA_PR_TH_MFernandez_MDD8_Operating_Jul19

AI Healthcare companies

- Using ECG-gated CT scans, the Zebra Medical Vision platform calculates the degree of calcification in a patient's coronary artery. AI-algorithms that can detect bone density, fat in the liver, and emphysema in the lungs from images.
- Caption Health's technology was found to have less variability in analyzing left ventricular ejection fraction (EF) than most cardiologists in a study conducted with the Minneapolis Heart Institute.

aidoc

Analyzes CT scan for a suspected Hammorahage then prioritizes on Radiologists worklists for Immediate attention

zebra MEDICAL VISION

Zebra is empowering radiologists with its AI1 offering

Caption Health

The AI platform automatically reviews digital video clips from a patient's echocardiogram and selects the best ones for calculating EF.

aidence
human sense in artificial intelligence

HEALTHCARE | March 26, 2020

27 Companies Changing Health Outcomes Through AI

Ejection Fraction (HF-rEF)

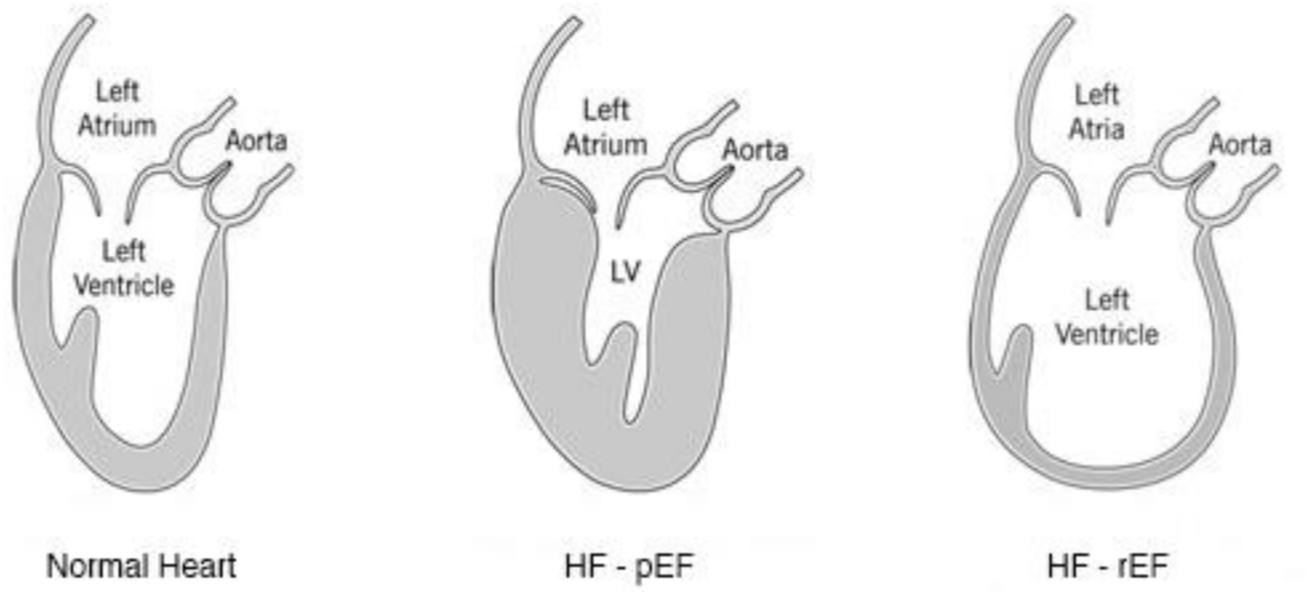

Potential Use Cases of AI in Healthcare

- Improve diagnostic accuracy
- Make use of EHR more effective
- Develop drugs and vaccine in shorter period
- Virtual Nursing Assistant
- Chatbot based Telemedicine

- Analyze blood samples. Track glucose levels in diabetic patients.
- Detect heart problems like Ejection Fraction
- Using image analysis to detect tumors.
- Detecting cancerous cells and diagnosing cancer using DNA sequencing/analysis
- Medical Imaging - detecting osteoarthritis from an MRI scan before the damage has begun.

https://mobisoftinfotech.com/resources/blog/data-science-in-healthcare-use-cases/

Source : Blog by Sinhasane, March 2019

Remote patient care management

- The Solution
- combining software from , vivifyhealth ading remote care platform, and handheld tabs, which provided a seamless **mobile solution for patients to track their medical conditions.**

- Each day, remote care patients use **Bluetooth-enabled medical devices** — such as **weight scales, blood pressure cuffs and oximeter/pulse readers — to capture biometric data on the Samsung tablets and transmit it to AHC's Command Center.**

- Patients also answer a series of daily questions, based on their conditions and care plans. When AHC medical teams notice troubling trends in the data or spot potential problems, they can immediately reach

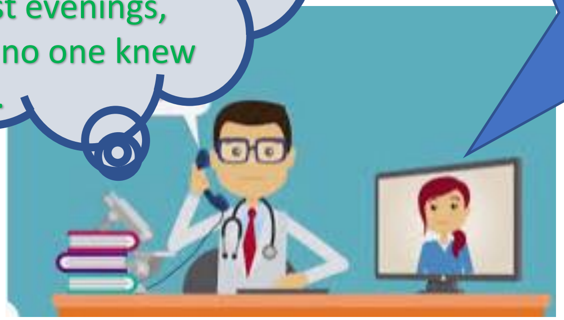

Blood pressure was lower than normal, heart rate was elevated, importantly, the nurse could see the patient was very pale on the video …

Nurse realized the patient might have gastrointestinal bleed ?

Courtesy: Cartoonstock, DoodleVideoProduction.com

High throughput sequencing(HTS)

Assessing and scoring chronic disease risk

- Use AI as part of cancer treatment program to understand and fight different types of cancer.
- Analyze the DNA of individual tumors and evaluate which specific mutations might be making a cancer grow.
- With Biorepository, a central data source of biological samples and linked clinical data, scientists can determine new biomarkers and other discoveries to yield new treatment options.
- Genetic samples are sequenced and analyzed over a 16-day period by a combination of **AI and a virtual molecular tumor board** that interprets the results and recommends targeted treatment.
- For more than 80 percent of patients with advanced-stage cancer, this technique is able to identify actionable mutations and offer additional treatment options to target the mutations.

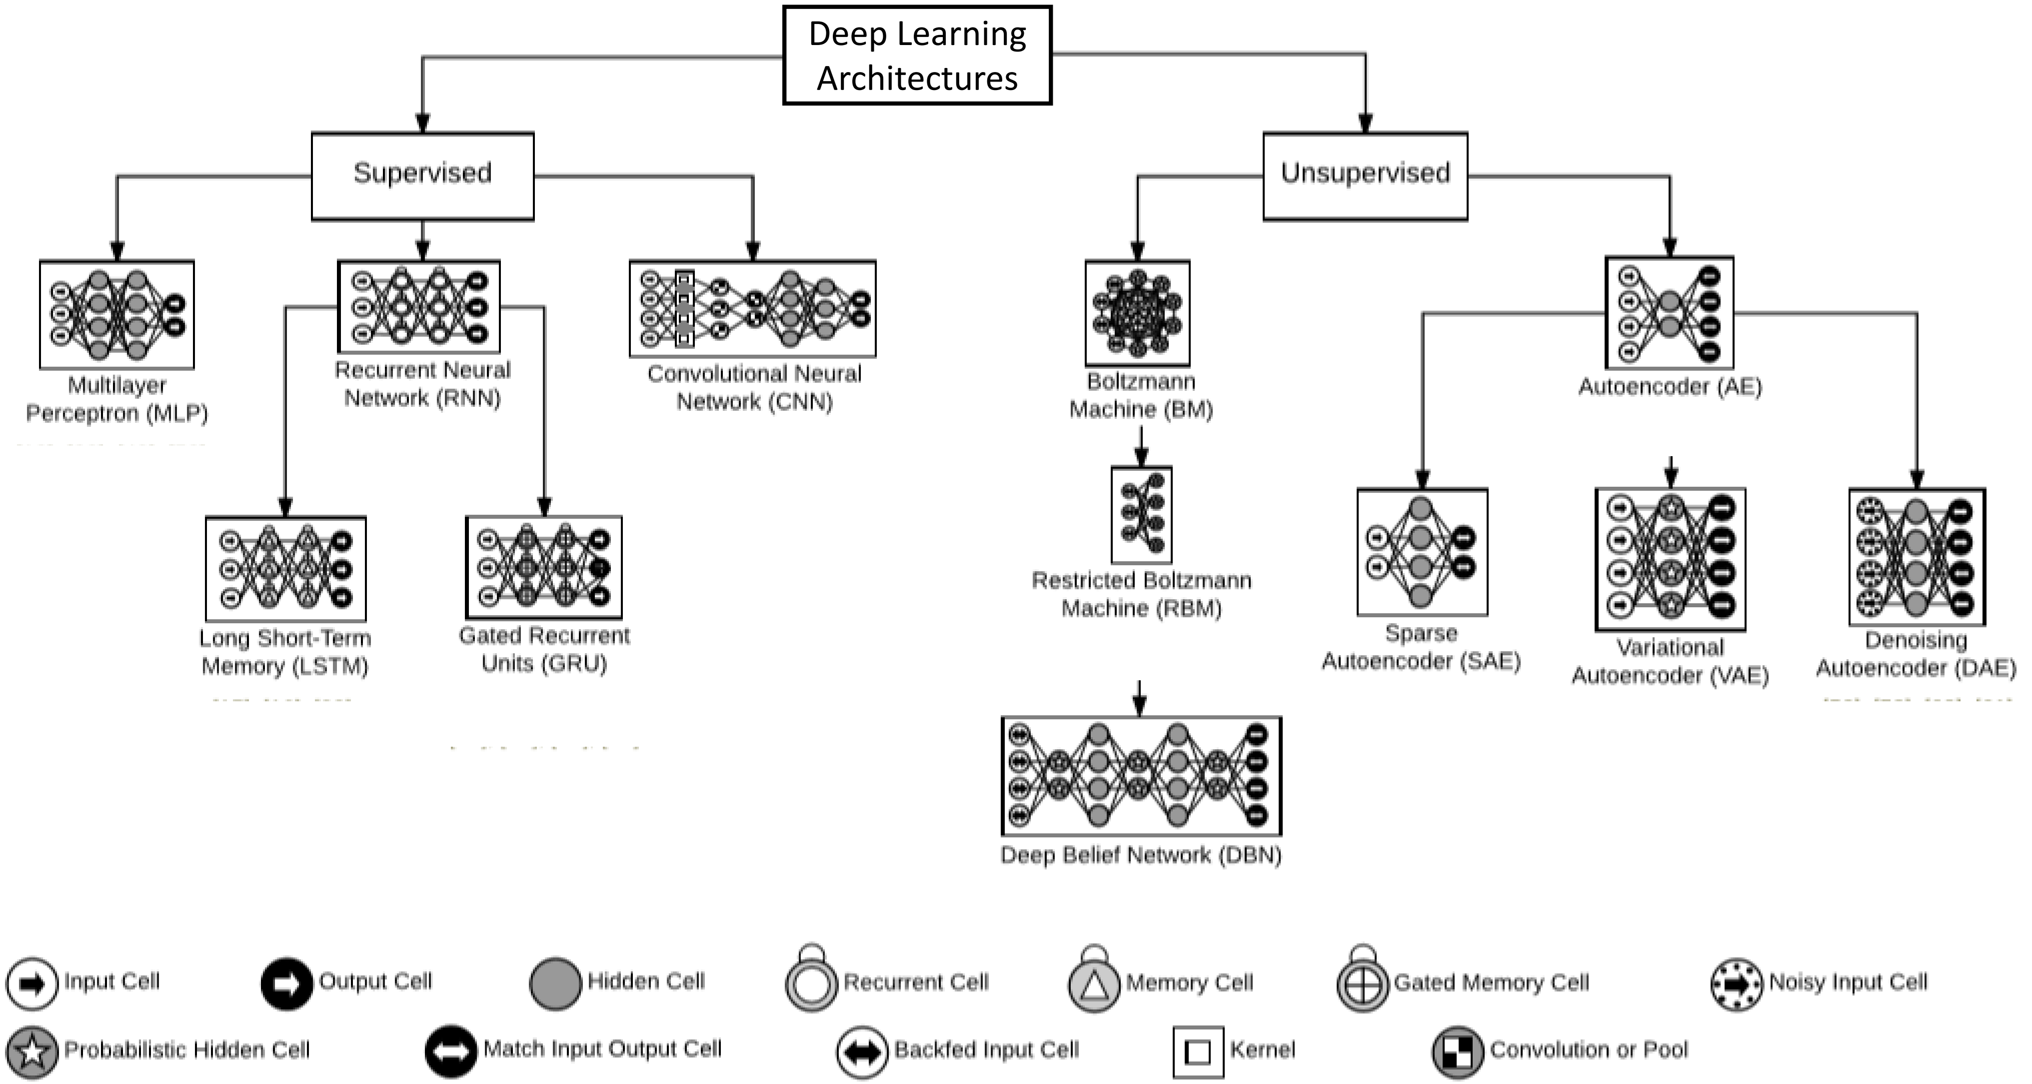

Deep Learning Architectures
Supervised
Unsupervised
Multilayer Perceptron (MLP)
Recurrent Neural Network (RNN)
Convolutional Neural Network (CNN)
Long Short-Term Memory (LSTM)
Gated Recurrent Units (GRU)
Boltzmann Machine (BM)
Restricted Boltzmann Machine (RBM)
Deep Belief Network (DBN)
Autoencoder (AE)
Sparse Autoencoder (SAE)
Variational Autoencoder (VAE)
Denoising Autoencoder (DAE)
Input Cell
Output Cell
Hidden Cell
Recurrent Cell
Memory Cell
Gated Memory Cell
Noisy Input Cell
Probabilistic Hidden Cell
Match Input Output Cell
Backfed Input Cell
Kernel
Convolution or Pool

The Dense Net is a primary architecture for very high capacity networks

- The DenseNet architecture is highly efficient, both in terms of parameter use and computation time.
- On the ImageNet ILSVRC-2012 dataset , a 201-layer DenseNet achieves roughly the same top-1 classification error as a 101-layer Residual Network (ResNet) , while using half as many parameters (20M vs 44M) and half as many floating point operations (80B/image vs 155B/image).
- **Each DenseNet layer is explicitly connected to all previous layers within a pooling region, rather than only receiving information from the most recent layer.**
- These connections promote feature reuse, as early-layer features can be utilized by all other layers. Because features accumulate, the final classification layer has access to a large and diverse feature representation.
- #a 161-layer DenseNet(with k = 48 features per layer and 29M parameters) achieves a top-1 single-crop error of 22.2% on the ImageNet ILSVRC classification dataset.
- It is reasonable to expect that larger networks would perform even better. However, with most existing DenseNet implementations, model size is currently limited by GPU memory.

Huang et al.[9]

Memory bottleneck challenges

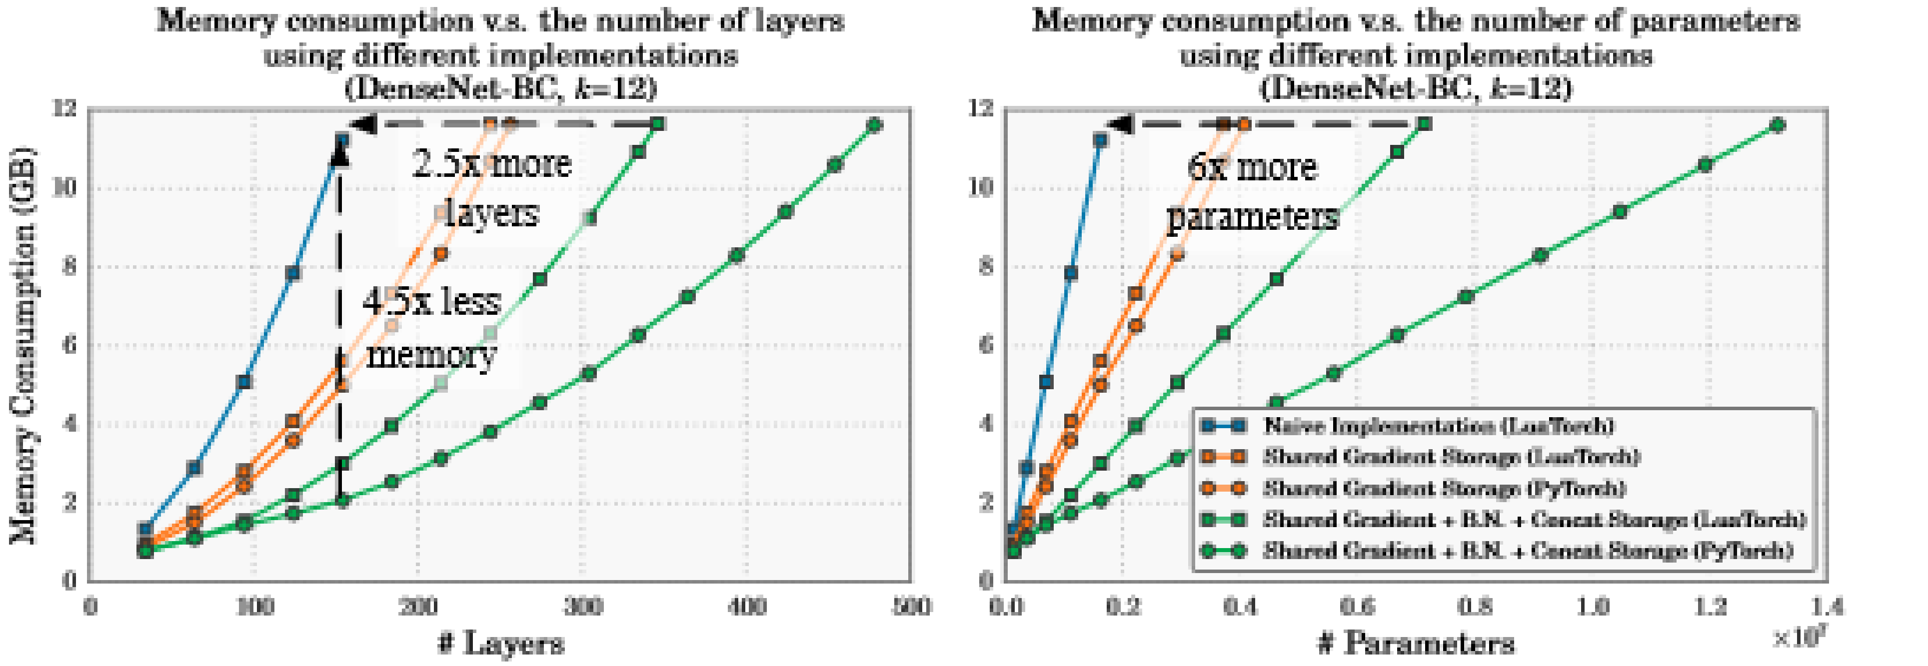

Figure 4: GPU memory consumption as a function of network depth. Each model is a bottlenecked Densenet (Densenet-BC) with $k = 12$ features added per layer. The efficient implementations enables training much deeper models with less memory.

Autoencoders

- Autoencoders are an unsupervised learning technique in which we leverage neural networks for the task of **representation learning**. Specifically, we'll design a neural network architecture such that we *impose a bottleneck in the network which forces a **compressed** knowledge representation of the original input.*
  - Encoder – compression, auto – NN
- https://www.jeremyjordan.me/autoencoders/
  - Intro to autoencoders

Electronic Health Record and DeepLearning

- Electronic Health Record (EHR) systems store patient data, required personal information, medical history records, and lab results.
- EHR systems improve the rate of correct diagnosis and the time it takes to reach a prognosis, via the use of deep learning algorithms.
- DL algorithms use data stored in EHR systems to detect patterns in health trends and risk factors and draw conclusions based on the patterns they identify.
- Researchers can use data in EHR to create deep learning models that will predict the likelihood of certain health-related outcomes
  - *E.g likelihood of a person contracting a disease or medical condition*

EHR information extraction

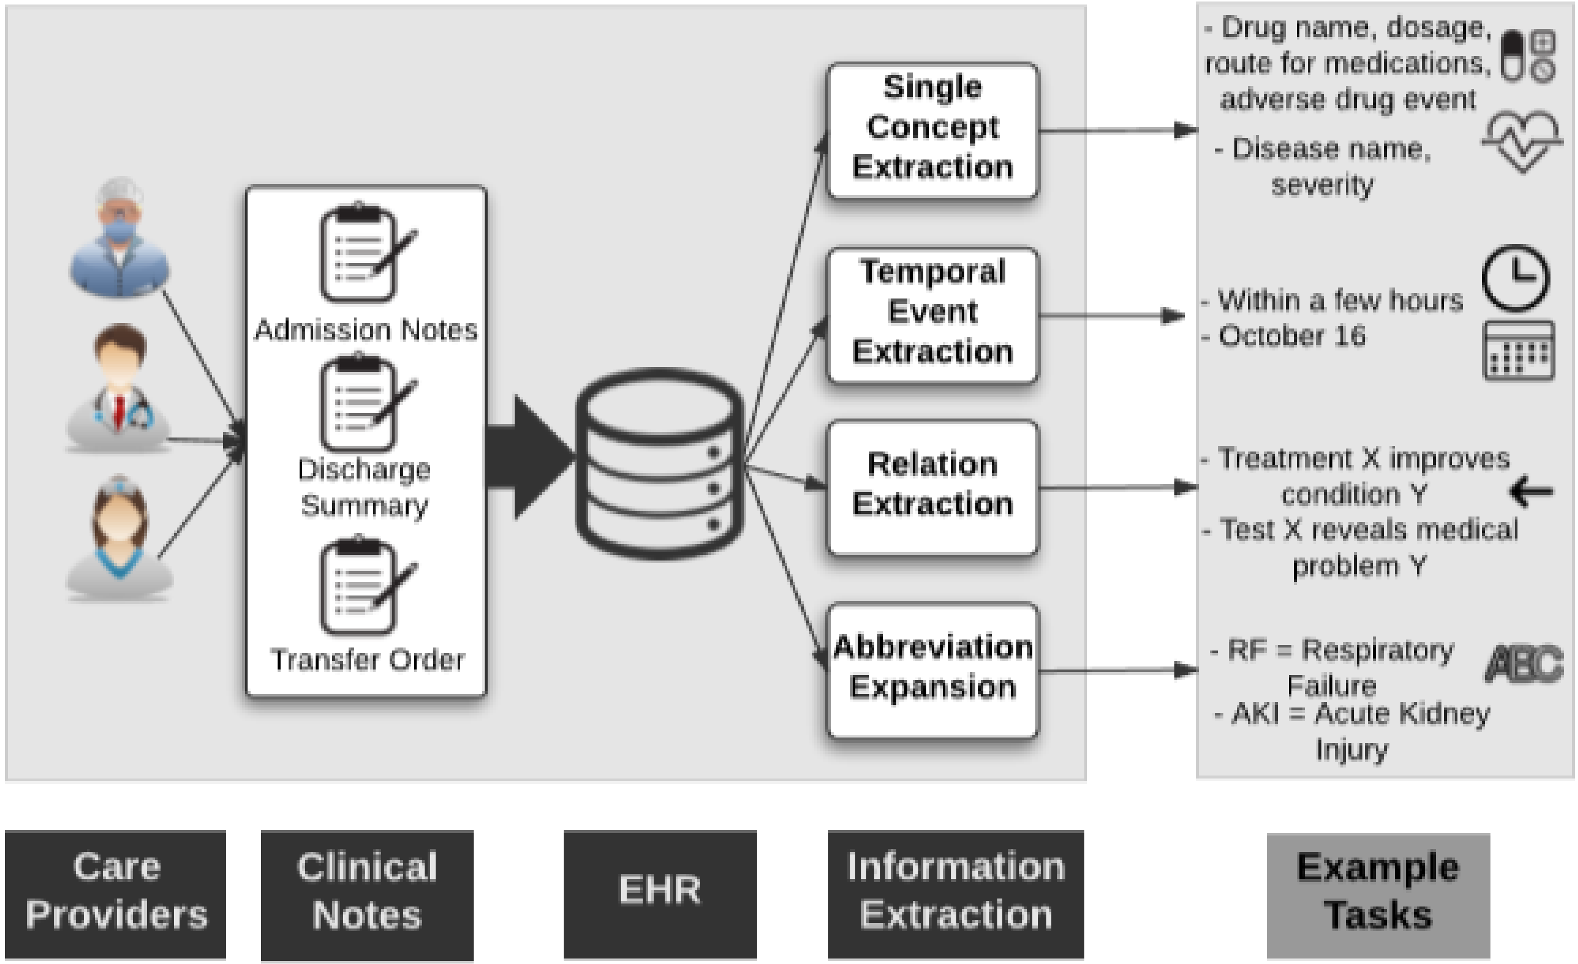

Summary of EHR Deep Learning processes

Task	Subtasks	Input Data	Models
Information Extraction	(1) Single Concept Extraction (2) Temporal Event Extraction (3) Relation Extraction (4) Abbreviation Expansion	Clinical Notes	LSTM, Bi-LSTM, GRU, CNN RNN + Word Embedding AE Custom Word Embedding
Representation Learning	(1) Concept Representation (2) Patient Representation	Medical Codes	RBM, Skip-gram, AE, LSTM RBM, Skip-gram, GRU, CNN, AE
Outcome Prediction	(1) Static Prediction (2) Temporal Prediction	Mixed	AE, LSTM, RBM, DBN LSTM
Phenotyping	(1) New Phenotype Discovery (2) Improving Existing Definitions	Mixed	AE, LSTM, RBM, DBN LSTM
De-identification	Clinical text de-identification	Clinical Notes	Bi-LSTM, RNN + Word Embedding

AE – auto encoder
GRU – gated recurrent unit
RBM – restricted Boltzman machine
DBN – Deep Belief Network

Relation extraction

- Relation Extraction: relation extraction deals with structured relationships between medical concepts in free text,
  - *treatment X improves/worsens/causes condition Y, or test X reveals medical problem Y. Lv et al. [36]*
- use standard text pre-processing methods and UMLS-based word-to-concept mappings in conjunction with sparse autoencoders to generate features for input to a CRF classifier, greatly outperforming the state of the art in EHR relation extraction.

Information Extraction

- **Identification and extraction of relevant facts and relationships from unstructured text**; the process **of making structured data from unstructured and semi-structured text**.
- **Single concept extraction** - Grouping of words and phrases into semantically similar groups.
- Temporal event extraction - **to extract medical events from the clinical text.**
  - a medical event is anything that is clinically important and that can also be mapped to a timeline.
  - E.g Seizure, hemorrhage,

https://www.sciencedirect.com/topics/mathematics/concept-extraction

Patient Representation Construction

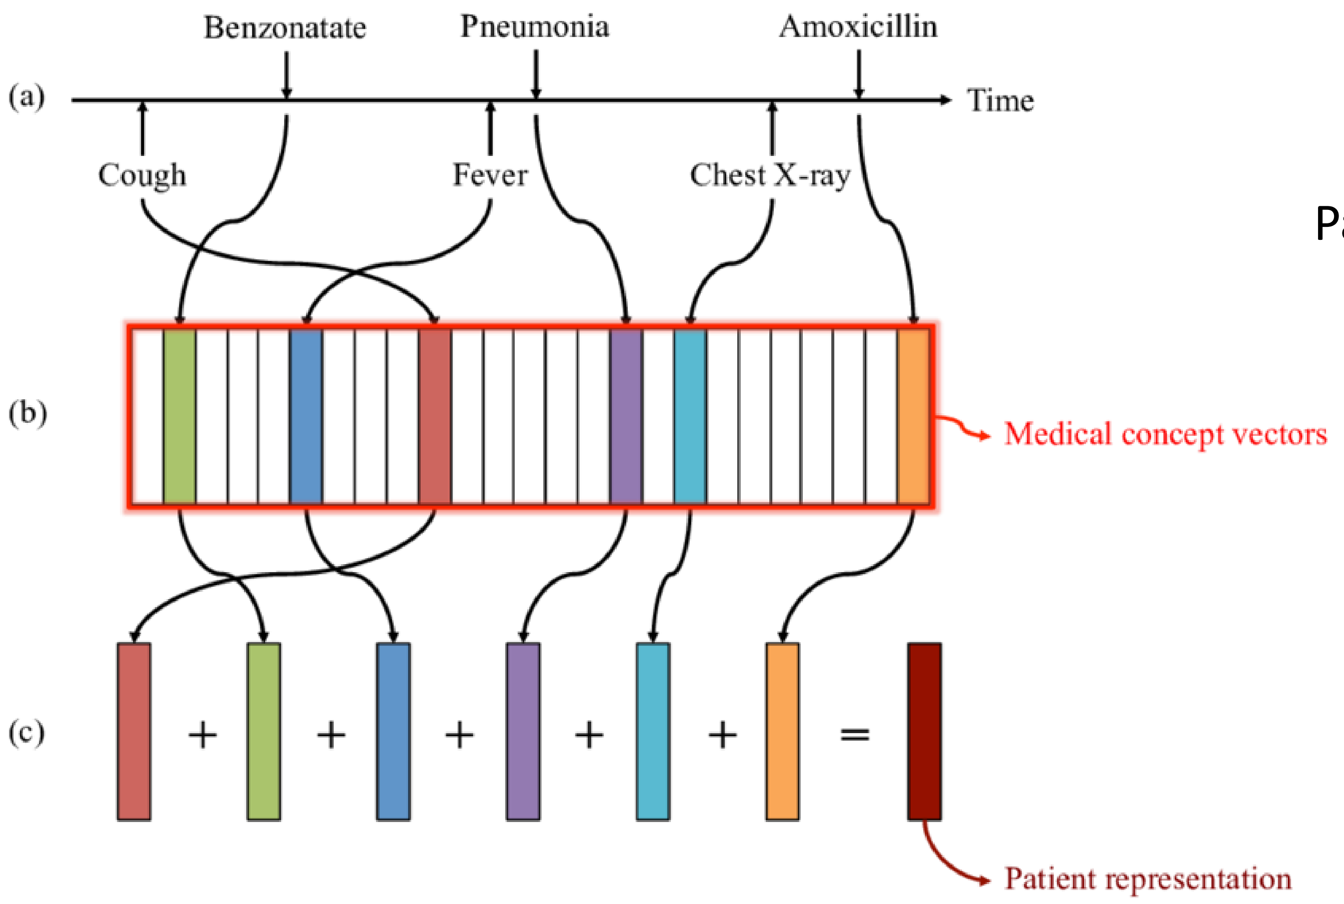

Patient representation construction

(a) represents a medical record of a patient on a timeline

(b) The medical concepts are represented as vectors using the trained medical concept vectors

(c) The patient is represented as a vector by summing all medical concept vectors.

Concept Representation

- Derive vector representations from sparse medical codes such that similar concepts are nearby in the lower-dimensional vector space. Once such vectors are obtained, codes of heterogeneous source types (such as diagnoses and medications) can be clustered
- techniques such as t-SNE wordcloud visualizations of discriminative clinical codes , or code similarity heatmaps .
- Word2Vec is an unsupervised ANN framework for obtaining vector representations. skip-gram, a model in word2vec .

Autoencoders help transform extremely sparse patient vectors into a more compact representation.

- **Medical coding** is the transformation of healthcare diagnosis, procedures, medical services, and equipment into universal medical alphanumeric codes.

- Medical codes are represented as binary categorical features, raw patient vectors can have dimensions in the thousands.

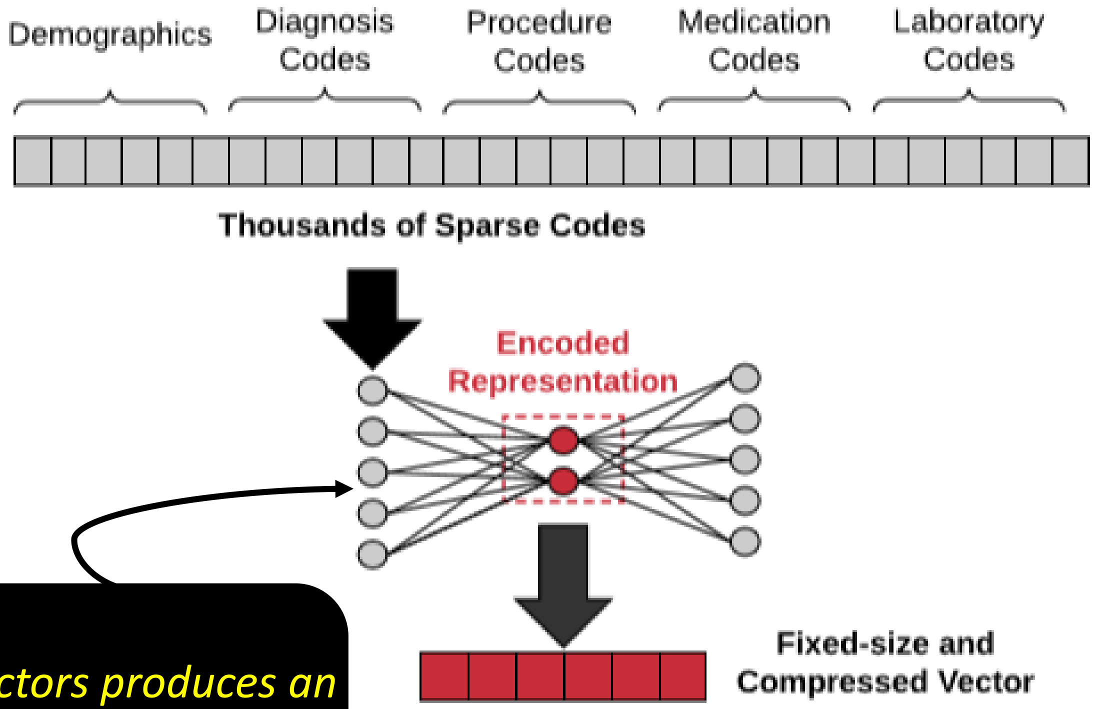

Training an autoencoder on these vectors produces an encoding function to transform any given vector into it's distributed and dimensionality-reduced representation/compressed.

Unsupervised autoencoder discovers latent structure in the raw clinical data without any human input.

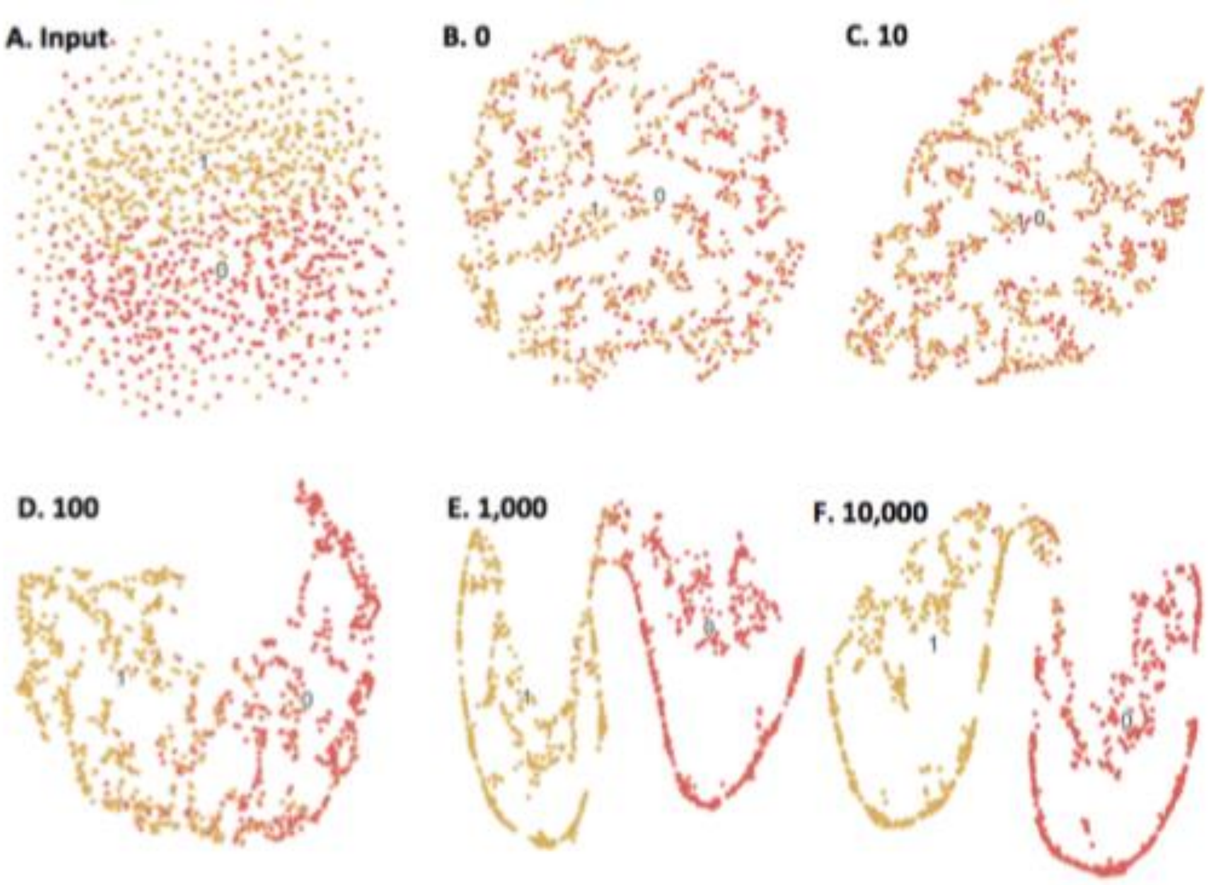

Phenotype stratification

autoencoder-based phenotype stratification for case (1) vs. control (0) diagnoses, illustrated with t-SNE.

(A) shows clustering based on raw clinical descriptors, where there is little separable structure.

(B-F) show the resulting clusters following 0-10,000 training epochs of the single-layer autoencoder.

As the autoencoder is trained, there are clear boundaries between the two labels,

Methods of evaluation for outcome prediction

- Most of the methods of evaluation for outcome prediction using deep learning techniques make use of standard classification metrics such as AUC[@]
  - AUC is Area Under Curve
  - Thumb rule AUC of 0.7 is required for practical application of model
  - AUC > 0.9 is outstanding model

[@] 5.3.9 ROC and AUC, Machine Learninng using Python, Manaranjan Pradhan, U Dinesh Kumar

CNN vs Dermatologists

“ These findings show that deep learning CNN are capable of outperforming dermatologists, in the task of detecting melanomas "
- Holger Haenssle, MD, senior managing physician, University of Heidelberg, Germany

abnormal growth or appearance compared to the skin around it

Dermatologists

- Having additional information improved the dermatologists' performance where the mean sensitivity was 88.9%; the mean specificity 75.7%, and the mean ROC AUC 0.82 (p<.01).

CNN performance

- on the same reader study level-II test showed a sensitivity of 95%, a specificity of 90%, and an ROC AUC of 0.95.

Annals of Oncology, May 2020
ROC - Receiver Operating Characteristic

Challenges - deep learning in medical imaging

- Challenges associated to memory and compute consumption when using CNNs with higher-dimensional image data
- Other important challenges are related to data, trust, interpretability, workflow integration, and regulations

https://www.sciencedirect.com/science/article/pii/S0939388918301181

Deep learning, medical imaging and MRI

AI, approach for electrocardiogram (ECG) analysis

- January 2019 edition of Nature Medicine showing expert-level detection of cardiac arrhythmias using a new deep learning, approach for electrocardiogram (ECG) analysis across a variety of diagnostic classes
- The term "arrhythmia" means an irregular heartbeat. It is used to describe manifestations ranging from benign, harmless conditions to severe, life-threatening disturbances of the heart's rhythm
- **the algorithm, a 34-layer Deep Neural Network, which learned from 91,232 ECG reports first time an artificial intelligence model has been able to successfully detect and label these many arrhythmias**
- ECG records are collected from 53,549 unique patients using **the Zio by iRhythm ambulatory continuous cardiac monitoring device.**

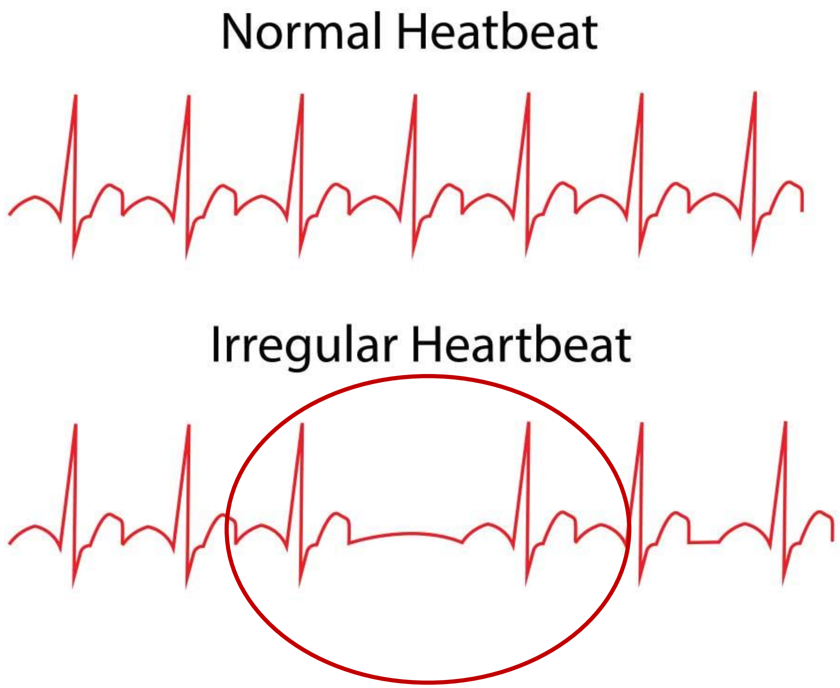

Courtesy:
https://www.diversalertnetwork.org/health/heart/arrhythmias-overview
https://www.irhythmtech.com/company/press-releases/research-published-in-nature-medicine-demonstrates-the-promise-of-algorithm-based-ambulatory-cardiac-monitoring

COVID-Net, a deep convolutional neural network

- Makes predictions using an explainability method in an attempt to gain deeper insights into critical factors associated with COVID cases, which can aid doctors in improved screening, but also audit COVID-Net in a responsible and transparent manner to validate that it is making decisions based on relevant information from the Chest X-RAY images.
- COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-Ray Images, Linda Wang and others, https://arxiv.org/pdf/2003.09871.pdf

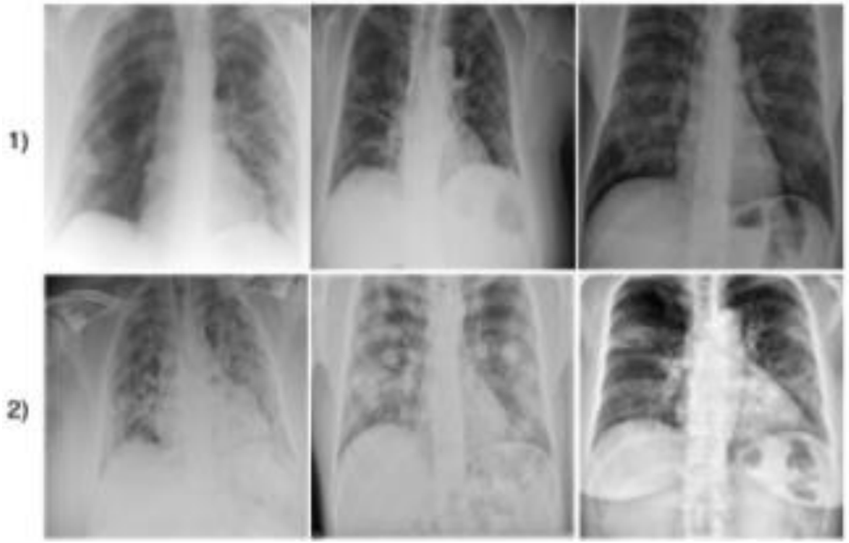

COVID-Net

Projection-expansion-projection-extension (PEPX) design

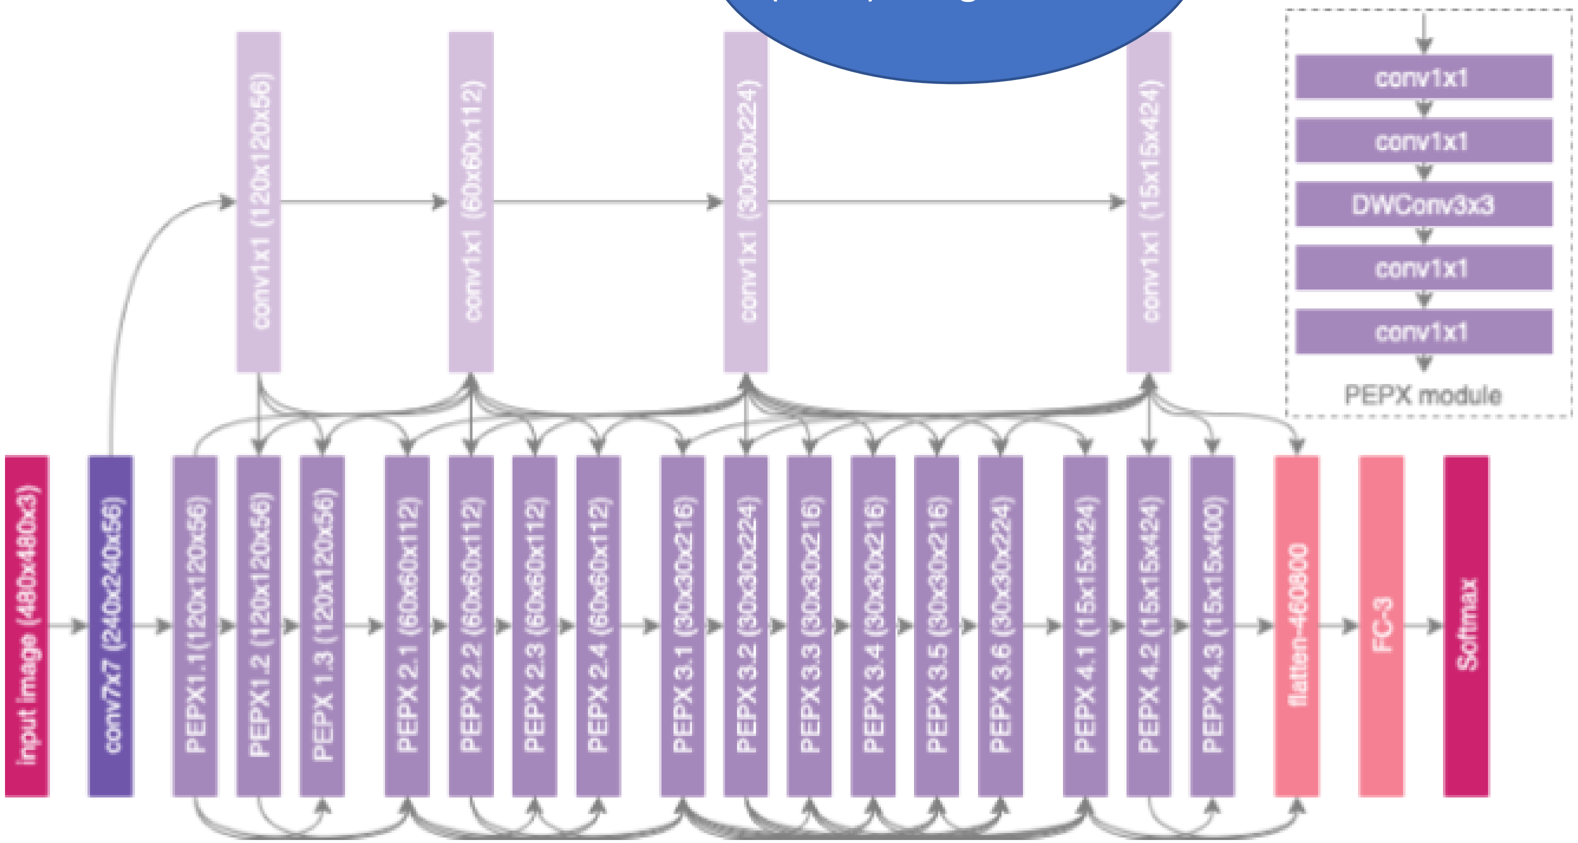

- COVID-Net was pretrained on the ImageNet36 dataset ,COVIDx datasets.
- Adam optimizer
- The hyperparameters used for training are: learning rate=2e-4, number of epochs=22, batch size=64, factor=0.7, patience=5.
- The initial COVID-Net prototype was built and evaluated using the Keras deep learning library with a TensorFlow backend.

Deep learning in the MR signal processing chain, from image acquisition and image reconstruction, to image restoration and image registration.

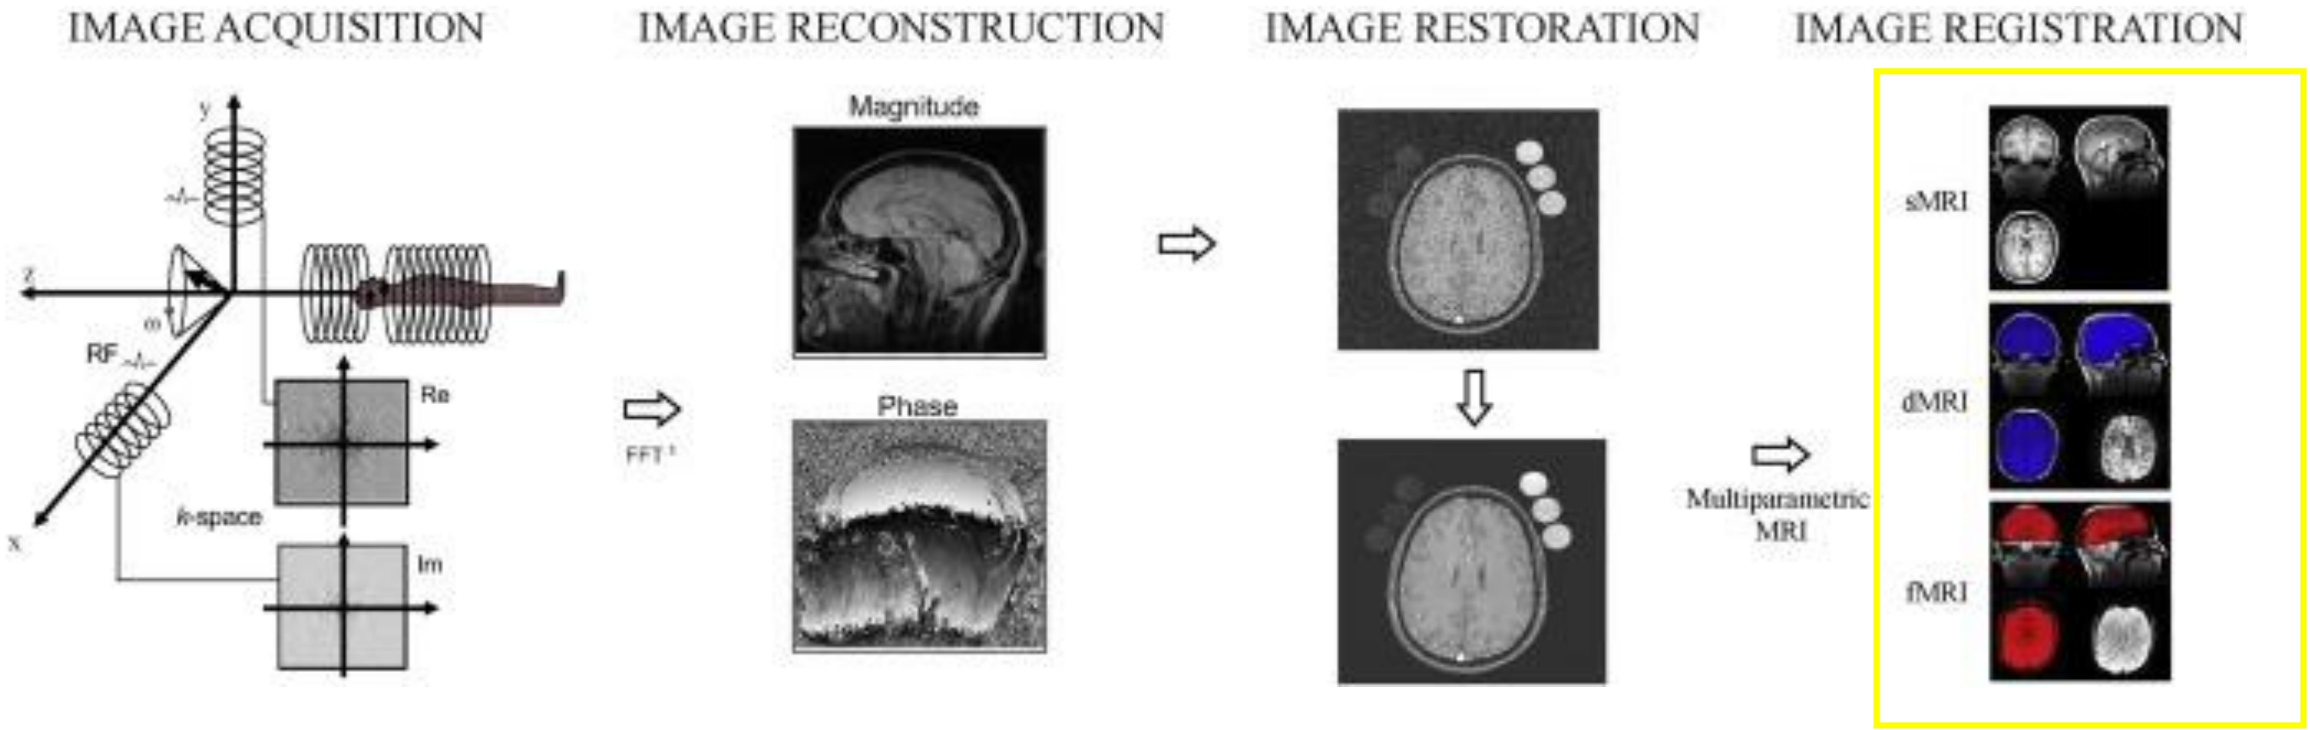

Yellow Box - illustrates multimodal brain MRI. sMRI = structural 3D T1-weighted MRI, dMRI = diffusion weighted MRI (stack of slices in blue superimposed on sMRI), fMRI = functional BOLD MRI (in red).

Dataever Consulting

NiftyNet

- NiftyNet is a TensorFlow-based open-source CNN platform for research in medical image analysis and image-guided therapy. NiftyNet's modular structure is designed for sharing networks and pre-trained models
- Using NiftyNet's modular structure
  - you can get started with established pre-trained networks using built-in tools
  - Quickly build new solutions to your own image analysis problems
- https://niftynet.readthedocs.io/en/dev/index.html
- The code is available via GitHub.
- Also you can quickly get started with the released versions in the Python Package Index.

Imaging of other human organs

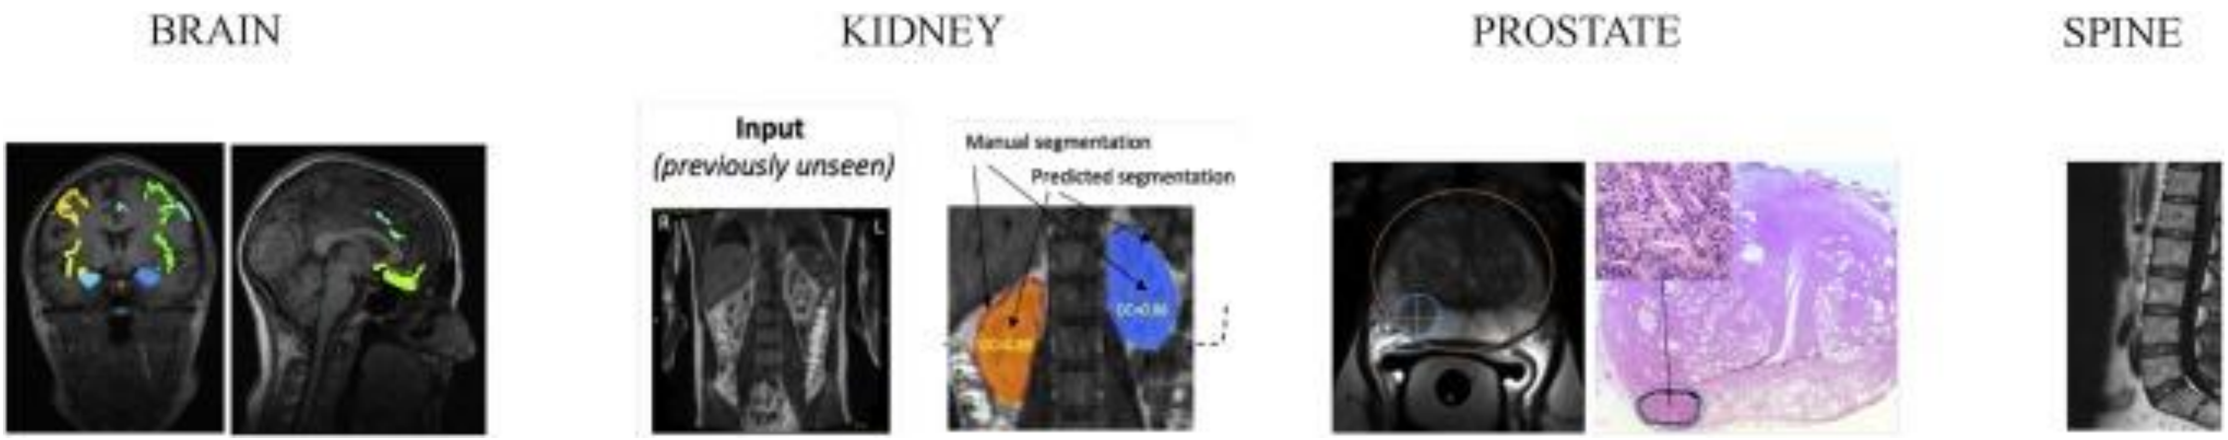

Dataever Consulting

Medical imaging competitions

- Automatically locate lung opacities on chest radiographs
  - https://www.kaggle.com/c/rsna-pneumonia-detection-challenge
- Grand challenges in biomedical image analysis. Hosts and lists a large number of competitions
  - https://grand-challenge.org/

https://www.sciencedirect.com/science/article/pii/S0939388918301181

Medical imaging dataset repositories

- Open Neuro - An open platform for sharing neuroimaging data under the public domain license. Contains brain images from 168 studies (4,718 participants) with various imaging modalities and acquisition protocols.
  - https://openneuro.org
- TCIA – The cancer imaging archive hosts a large archive of medical images of cancer accessible for public download. Currently contains images from 14,355 patients across 77 collections.
  - http://www.cancerimagingarchive.net

Health care industry is an early adopter of AI

- The AI/ML applications for Healthcare has been a successful journey so far.
- Open source approach has been a driving force behind this success
  - Most of the main new ideas and methods are posted to the arXiv preprint server, and the accompanying code shared on the GitHub platform. The data sets used are often openly available through various repositories.
- For a beginner, this along with numerous online educational resources, make it suitable to get started in the field.

Backup slides

Assignment: Folding@home

- Folding@home is a distributed computing project for simulating protein dynamics, including the process of protein folding and the movements of proteins implicated in a variety of diseases.
- It brings together citizen scientists who volunteer to run simulations of protein dynamics on their computers.
- Insights from this data are helping scientists to better understand biology, and providing new opportunities for developing therapeutics.
- https://ngc.nvidia.com/catalog/containers/hpc:foldingathome:fah-gpu

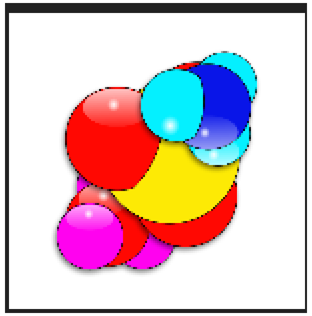

Outcome Prediction

- https://www.youtube.com/watch?v=6FblFg7fYWQ

AutoML

- Automated machine learning is the process of automating the process of applying machine learning to real-world problems.
- AutoML covers the complete pipeline from the raw dataset to the deployable machine learning model.
- Automated machine learning, or **AutoML**, aims to reduce or eliminate the need for skilled data scientists to build machine learning and deep learning models. Instead, an **AutoML** system allows you to provide the labeled training data as input and receive an optimized model as output.

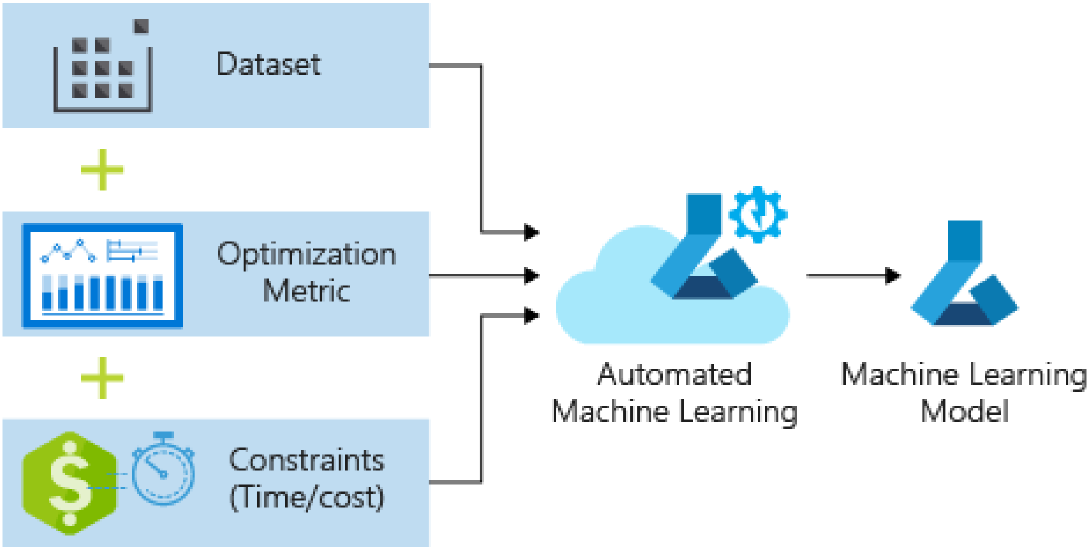

Collaborative decision making

Courtesy Philips Healthcare

Family history information extraction via Deep Joint Learning

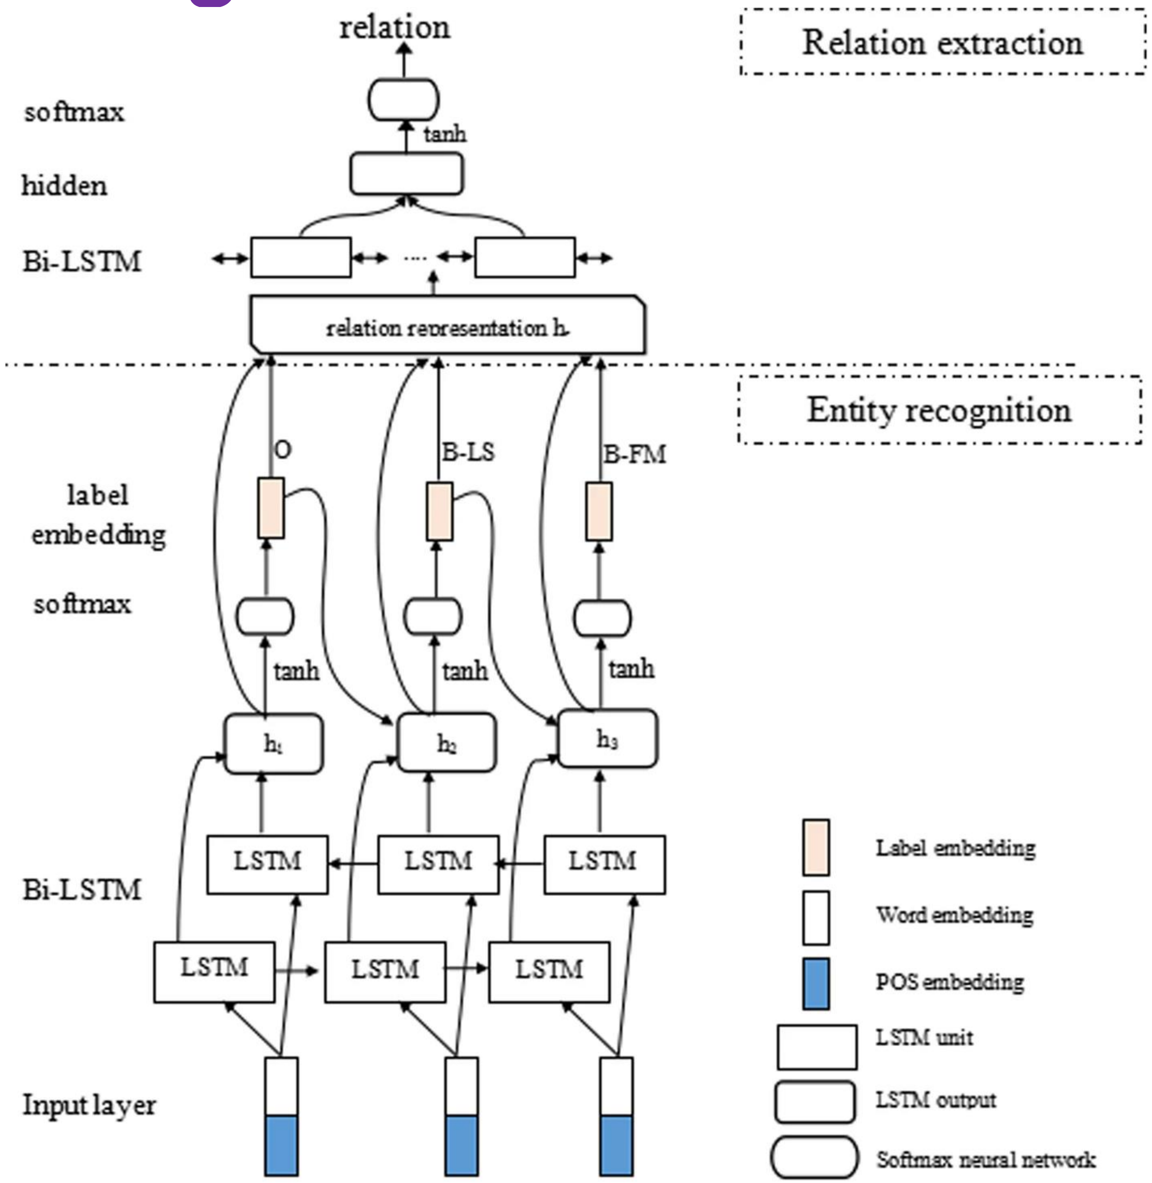

ML Techniques for analyzing EHR data

- Techniques for analyzing EHR data were based on traditional machine learning and statistical techniques such as logistic regression, support vector machines (SVM), and random forests .

DL Techniques for analyzing EHR data

- deep learning techniques have achieved success in many domains through deep hierarchical feature construction and capturing long-range dependencies in data

Phenotypes

- The term used in genetics for the composite observable characteristics or traits of an organism.

The shells of species show diverse coloration and patterning in their phenotypes.