PHILIPS

www.philips.com

The role of Artificial Intelligence in medical imaging: from research to clinical routine

Philips Research
Nicolas Villain
15-01-2018
France

innovation + you

2019
12th
Paris Hepatology Conference

The Need for AI in Healthcare – Big Picture

Growing demand

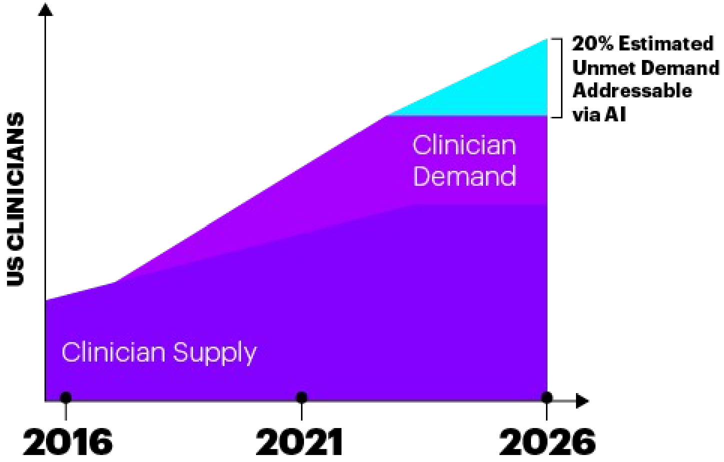

Demand for clinicians is predicted to outgrow supply and offer opportunities for AI enabled solutions

Impending crisis

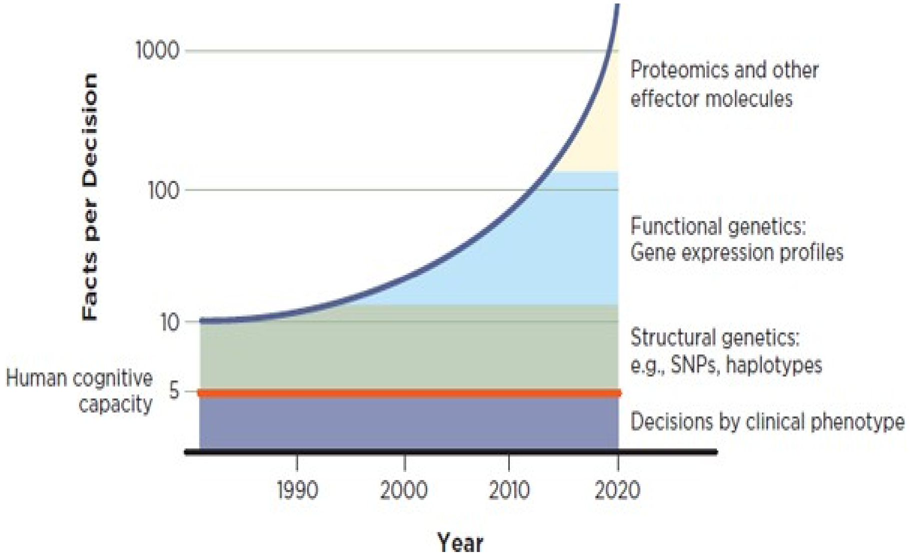

Clinical data outstrips human cognitive capacity by orders of magnitude

Source: Accenture analysis. Graph is not to scale and is illustrative.
https://www.nap.edu/read/13444/chapter/6

The AI Hype ‘Roller Coaster’

What is Driving The Roller Coaster

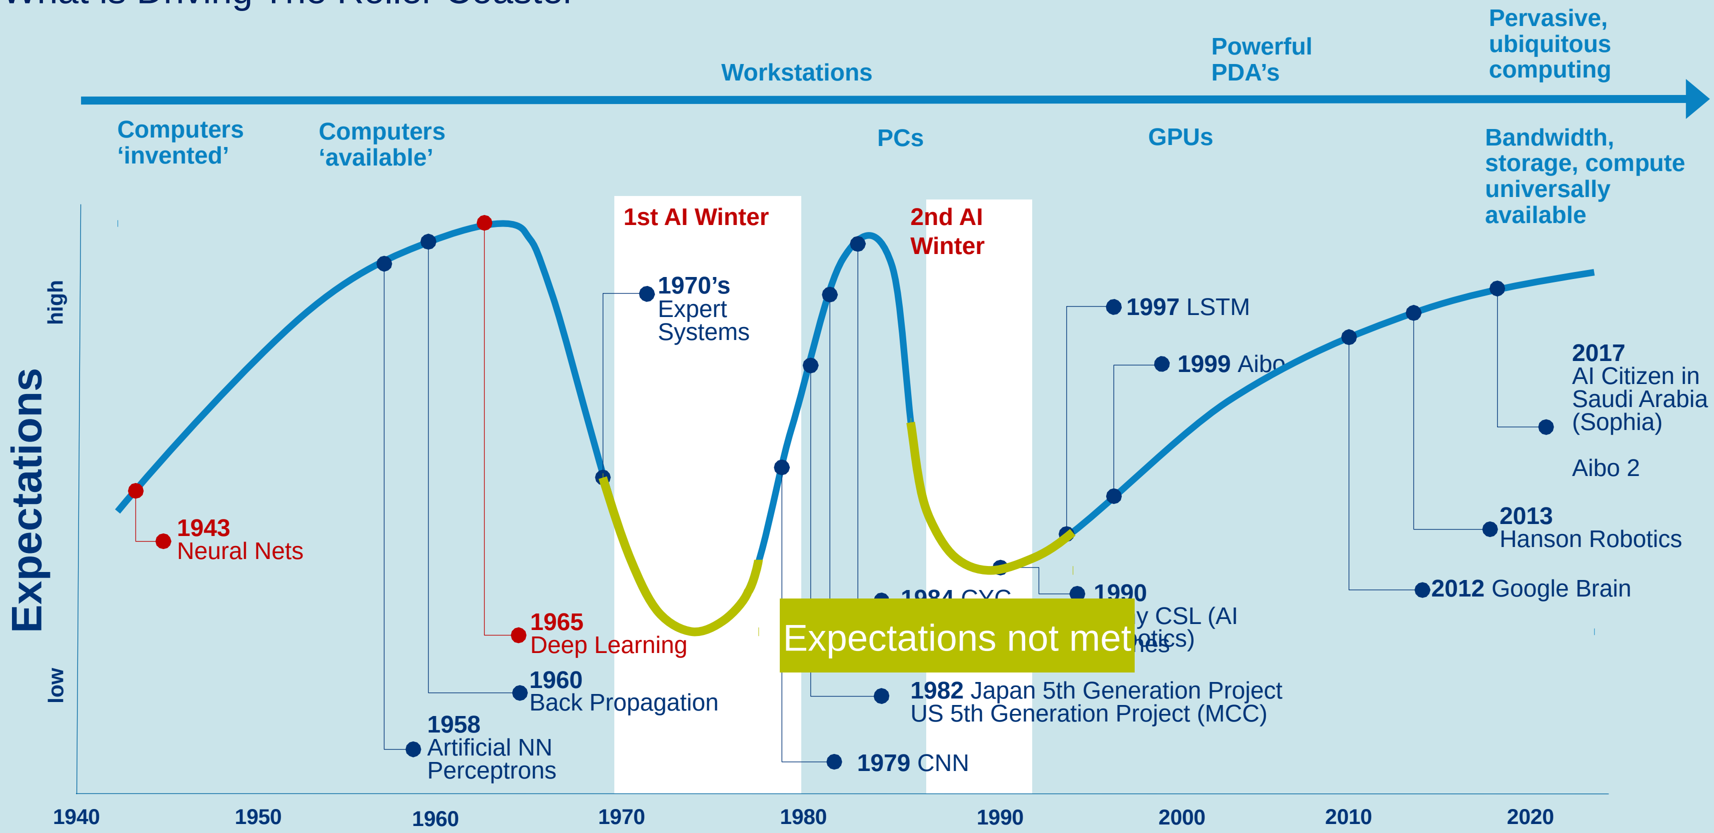

- Artificial Intelligence:
- the many successes of Deep Learning

- Games:
  - Go
  - Chess & many others
- Language
  - Natural language processing
  - Speech recognition
  - Automated translation
- Computer vision:
  - Object detection and recognition
  - Self-driving cars
  - Robotics
  - Medical image processing

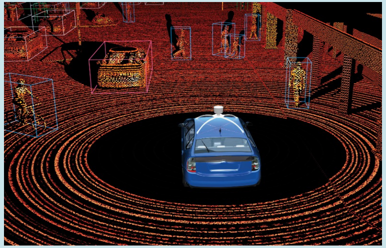

“Traditional” Computer Vision

Labeled Data

“dog”

“cat”

Software Engineer

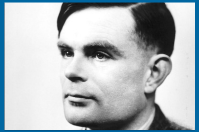

Programming

Prediction

Unlabeled Data

Algorithm

Prediction

“dog”

A Computer Vision expert writes a computer program / algorithm to detect dogs & cats.

Deep Learning

A **Lot** of Labeled

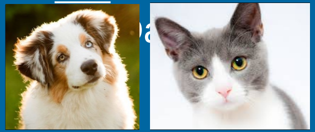

"dog" "cat"

Learning Algorithm

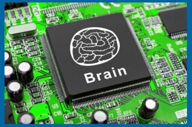

Training

Prediction

Unlabeled Data

Learned Model

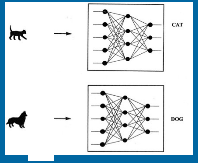

Prediction

"dog"

Deep Learning is a technique that provides computers with the ability to **learn** without being explicitly programmed

Deep Learning

A **Lot** of Labeled

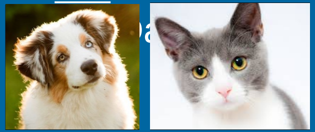

"dog" "cat"

Learning Algorithm

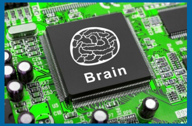

Training

Prediction

Unlabeled Data

Learned Model

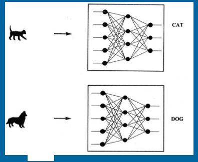

Prediction

"?"

Deep Learning is a technique that provides computers with the ability to **learn** without being explicitly programmed

Looking Inside of the Black Box

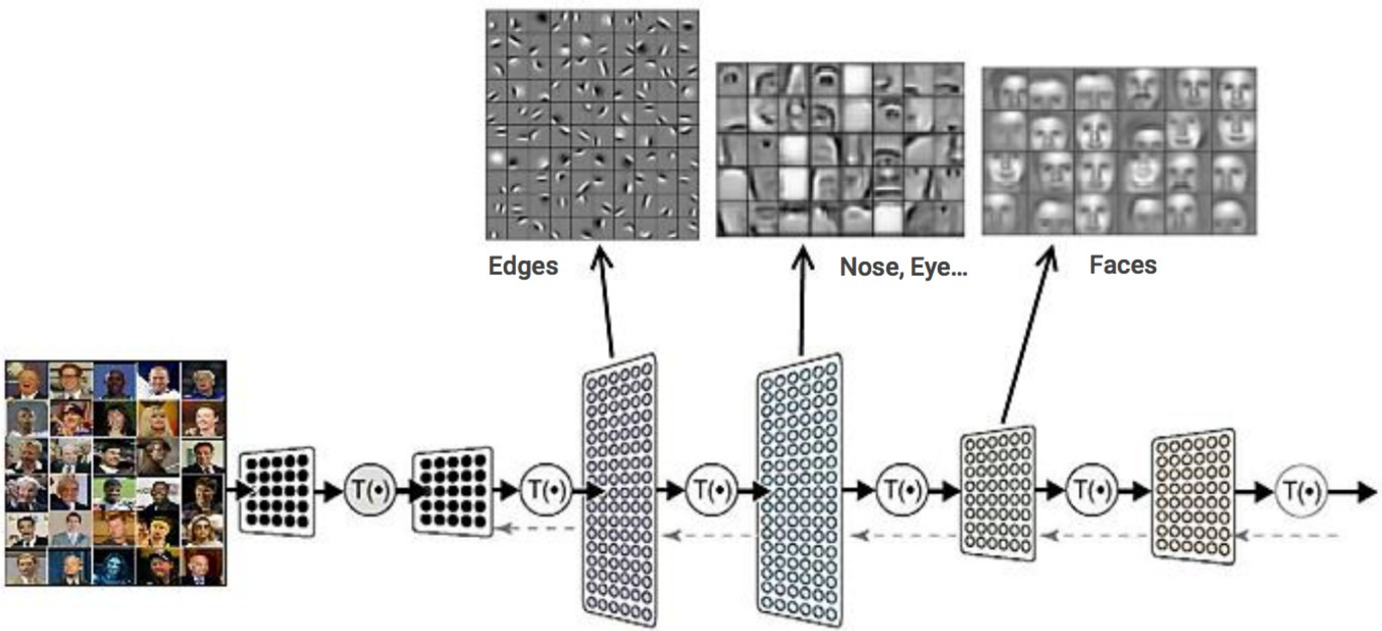

nose -> face). The output layer combines those features to make predictions.

Deep Neural Nets interpret images like radiologists

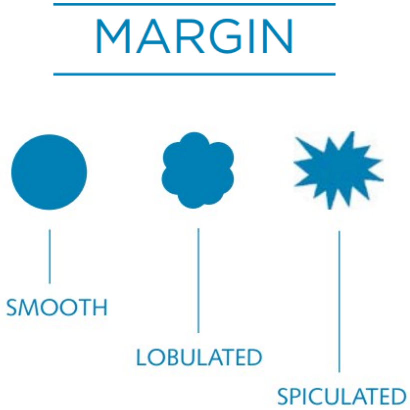

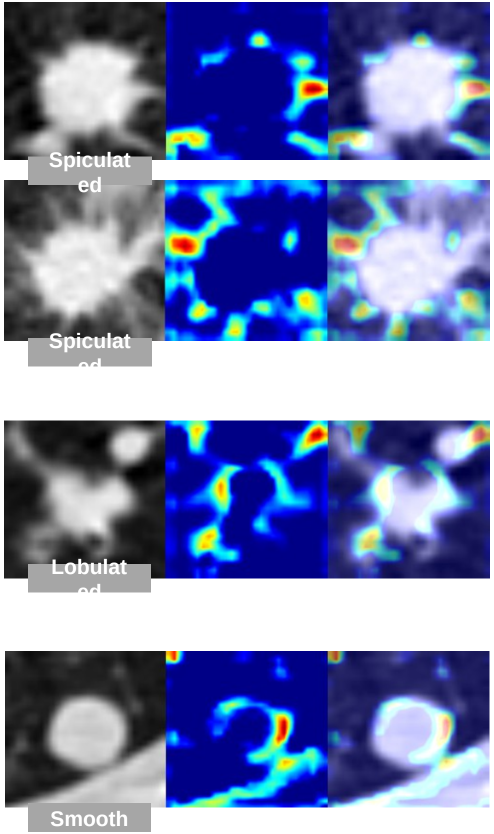

prediction = 0.73
correct class = 2

prediction = 0.81
correct class = 2

prediction = 0.37
correct class = 1

prediction = 0.23
correct class = 0

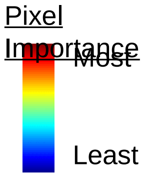

Prédire les accidents cardiovasculaires

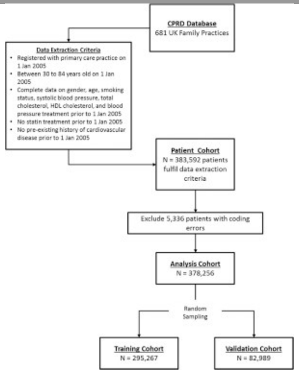

Prospective cohort study using routine clinical data of 378,256 patients from UK family practices, free from cardiovascular disease at outset.

Four machine-learning algorithms were compared to an established algorithm (**American College of Cardiology guidelines**) to predict first cardiovascular event over 10-years.

The best one—neural networks—correctly predicted 7.6% more events than the ACC/AHA method, and it raised 1.6% fewer false alarms. In the test sample of about 83,000 records, **that amounts to 355 additional patients whose lives could have been saved**.

Prédire les accidents cardiovasculaires

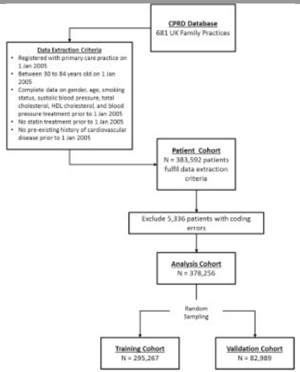

Prospective cohort study using routine clinical data of 378,256 patients from UK family practices, free from cardiovascular disease at outset.

Four machine-learning algorithms were compared to an established algorithm (**American College of Cardiology guidelines**) to predict first cardiovascular event over 10-years.

The best one—neural networks—correctly predicted 7.6% more events than the ACC/AHA method, and it raised 1.6% fewer false alarms. In the test sample of about 83,000 records, **that amounts to 355 additional patients whose lives could have been saved**.

ACC/AHA Algorithm		
Men	**Women**	**ML: Neural Networks**
Age	Age	Atrial Fibrillation
Total Cholesterol	HDL Cholesterol	Ethnicity
HDL Cholesterol	Total Cholesterol	Oral Corticosteroid Prescribed
Smoking	Smoking	Age
Age x Total Cholesterol	Age x *HDL Cholesterol*	Severe Mental Illness
Treated Systolic Blood Pressure	Age x Total Cholesterol	SES: Townsend Deprivation Index
Age x Smoking	Treated Systolic Blood Pressure	Chronic Kidney Disease
Age x *HDL Cholesterol*	Untreated Systolic Blood Pressure	*BMI missing*
Untreated Systolic Blood Pressure	Age x Smoking	Smoking
Diabetes	Diabetes	Gender

Italics: Protective Factors

https://doi.org/10.1371/journal.pone.0174944.t003

Several of the risk factors that the machine-learning algorithms identified as the **strongest predictors are not included in the ACC/AHA guidelines**, such as severe mental illness and taking oral corticosteroids.

Meanwhile, **none of the algorithms considered diabetes**, which is on the **ACC/AHA list**, to be among the top 10 predictors.

2019 12th Paris Hepatology Conference

PHILIPS

IA helps transform large amounts of data into actionable insights to augment clinicians and empower patients with their own health.

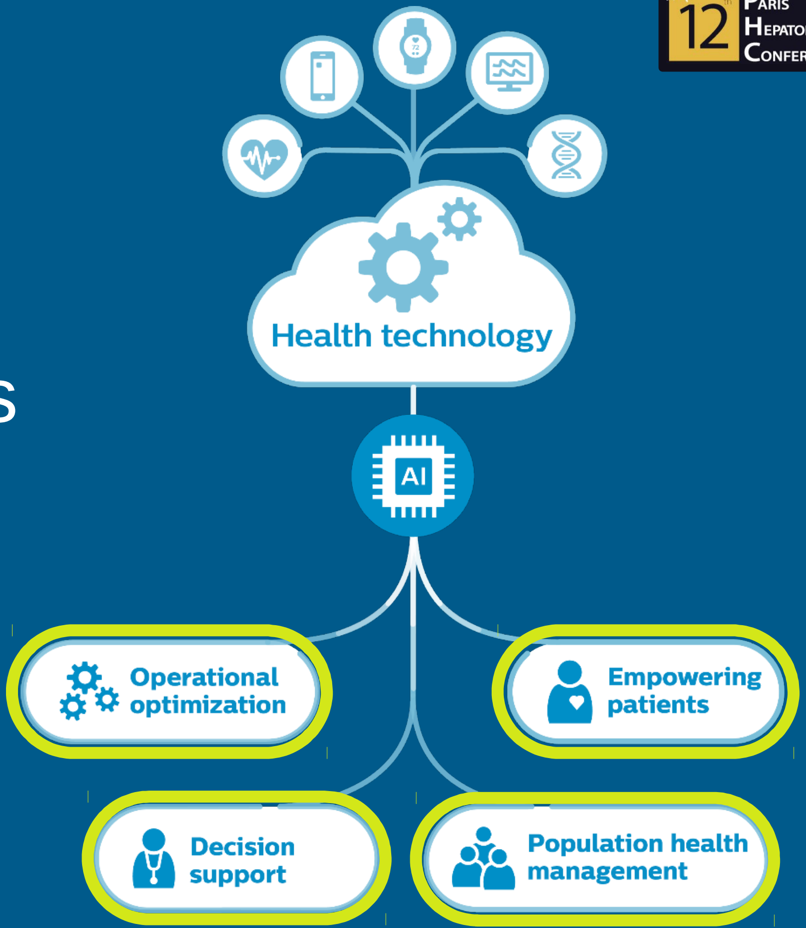

- The AI Challenges in Healthcare

Large quantities of data with appropriate patient privacy protection

Curated clean databases

Trust & transparency

Combining data and knowledge driven learning

The last mile: understanding of the (local) clinical context

Review cycles with clinicians, prove outcomes, Q&R

New ways of working

Fragmented Healthcare IT infrastructure

Training the Best Model is not the Main Problem…

Establishing the context to leverage AI methods is far more complex

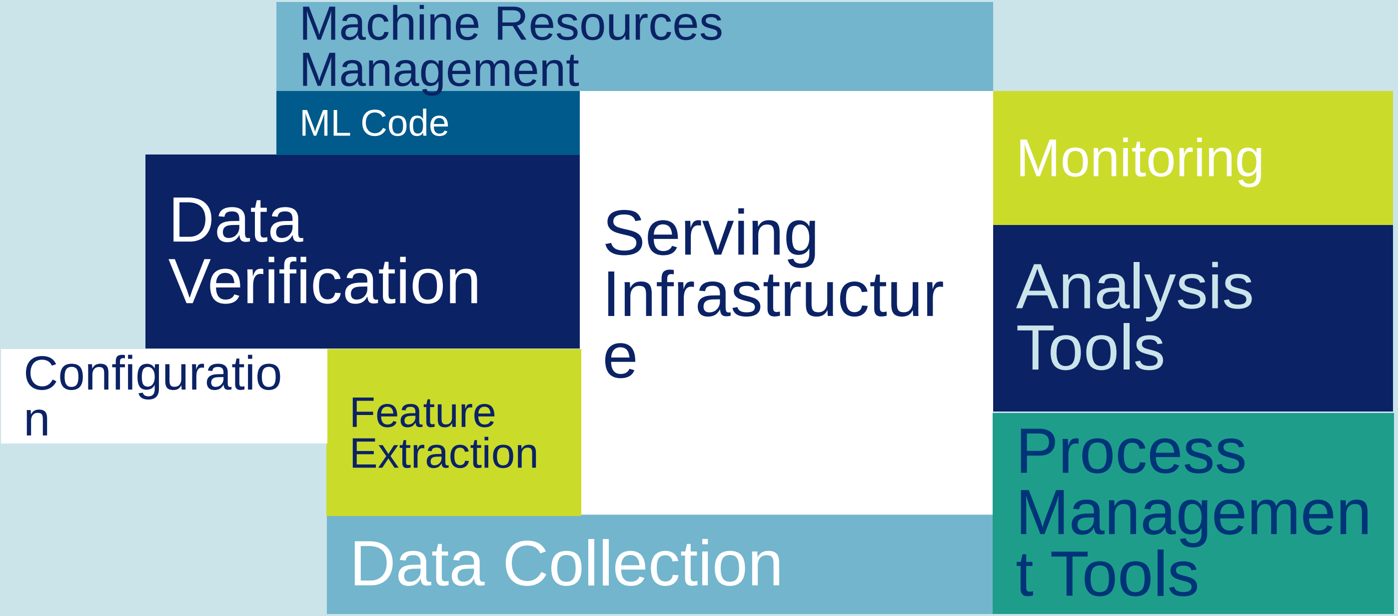

From Google article: “Hidden Technical Debt in Machine Learning Systems” D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips

How to Create and Deploy an AI solution

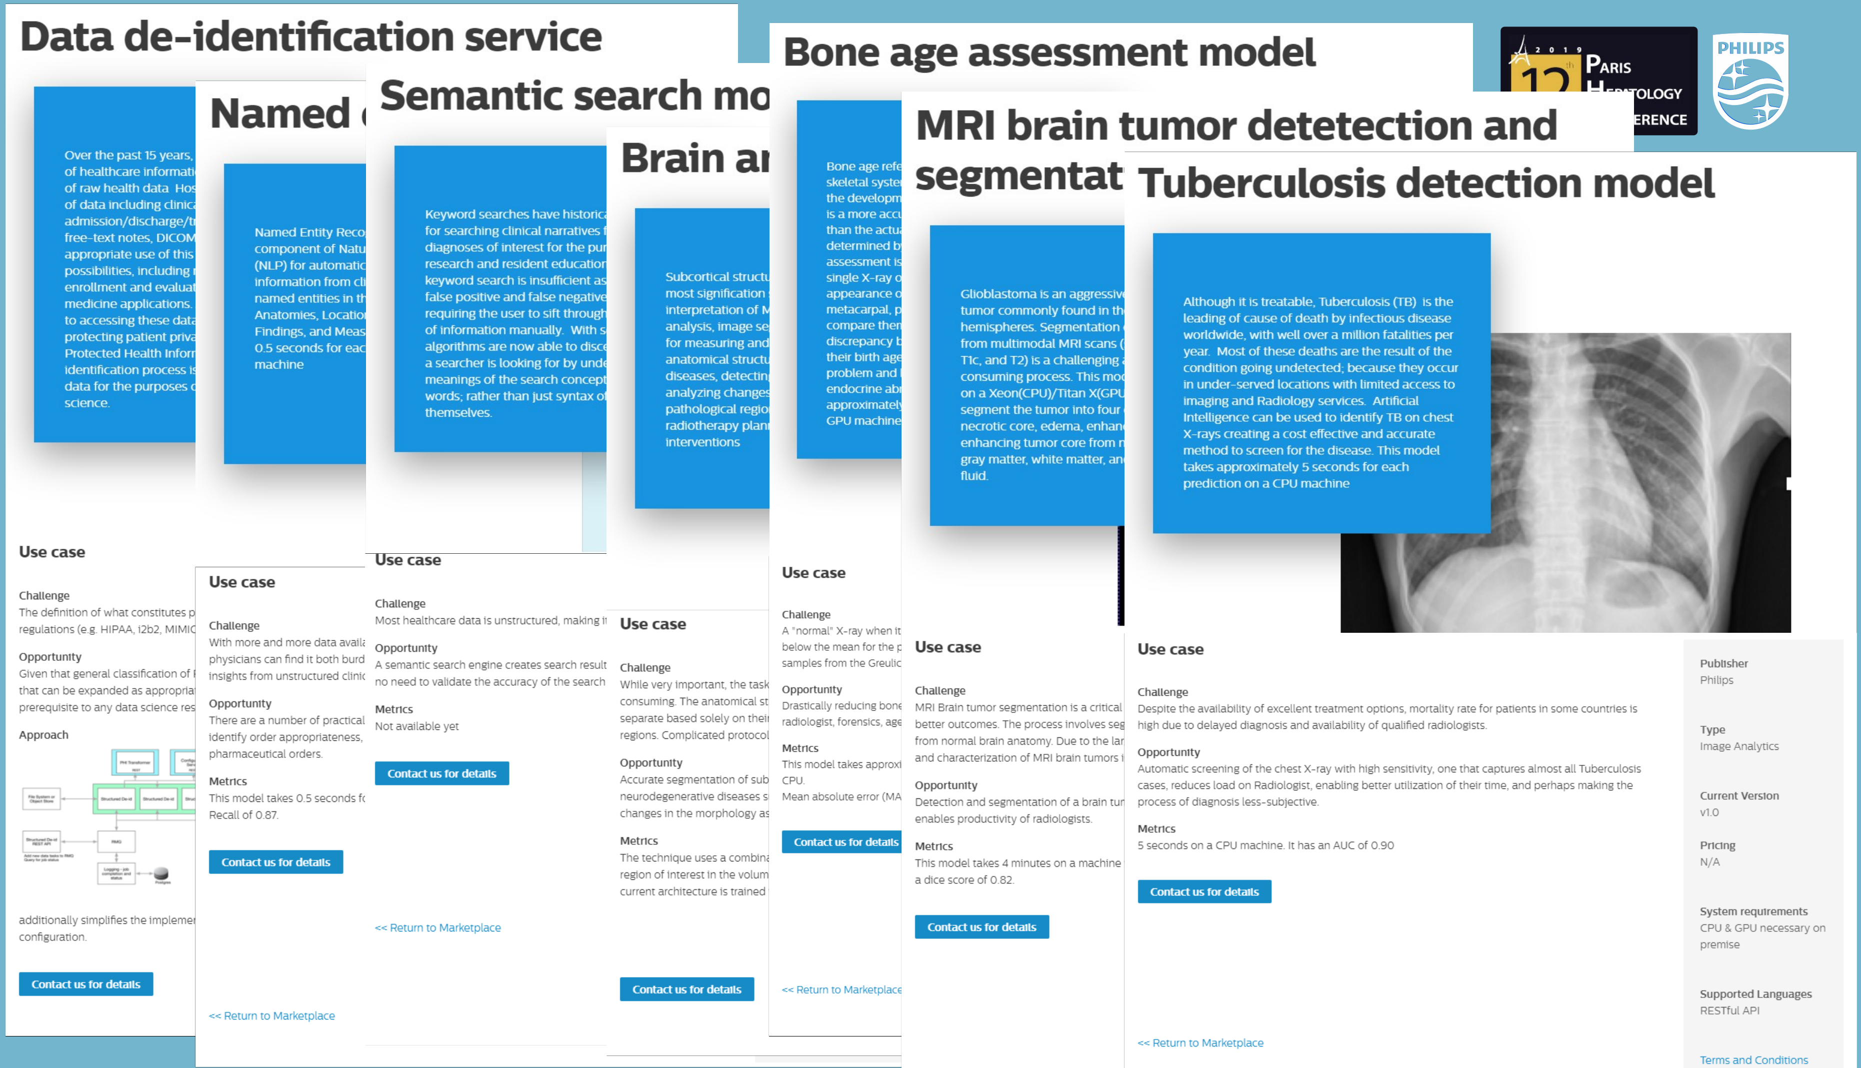

Data de-identification service
Over the past 15 years,
of healthcare informati
of raw health data Hos
of data including clinica
admission/discharge/t
free-text notes, DICOM
appropriate use of this
possibilities, including
enrollment and evaluat
medicine applications.
to accessing these data
protecting patient priva
Protected Health Inform
identification process is
data for the purposes o
science.
Use case
Challenge
The definition of what constitutes p
regulations (e.g. HIPAA, i2b2, MIMIC
Opportunity
Given that general classification of
that can be expanded as appropria
prerequisite to any data science res
Approach
additionally simplifies the implemen
configuration.
Contact us for details
Named
Named Entity Reco
component of Natu
(NLP) for automatic
information from cl
named entities in th
Anatomies, Locatio
Findings, and Meas
0.5 seconds for eac
machine
Use case
Challenge
With more and more data avail
physicians can find it both burd
insights from unstructured clinic
Opportunity
There are a number of practical
identify order appropriateness,
pharmaceutical orders.
Metrics
This model takes 0.5 seconds f
Recall of 0.87.
Contact us for details
<< Return to Marketplace
Semantic search mo
Keyword searches have historica
for searching clinical narratives f
diagnoses of interest for the pur
research and resident education
keyword search is insufficient as
false positive and false negative
requiring the user to sift through
of information manually. With s
algorithms are now able to disce
a searcher is looking for by unde
meanings of the search concept
words; rather than just syntax of
themselves.
Use case
Challenge
Most healthcare data is unstructured, making i
Opportunity
A semantic search engine creates search result
no need to validate the accuracy of the search
Metrics
Not available yet
Contact us for details
<< Return to Marketplace
Brain ar
Subcortical structu
most signification
interpretation of M
analysis, image se
for measuring and
anatomical structu
diseases, detectin
analyzing changes
pathological regio
radiotherapy plan
interventions
Use case
Challenge
While very important, the task
consuming. The anatomical st
separate based solely on their
regions. Complicated protocol
Opportunity
Accurate segmentation of sub
neurodegenerative diseases s
changes in the morphology as
Metrics
The technique uses a combina
region of interest in the volum
current architecture is trained
Contact us for details
Bone age assessment model
Bone age refe
skeletal syste
the developm
is a more accu
than the actu
determined b
assessment is
single X-ray o
appearance o
metacarpal, p
compare them
discrepancy b
their birth age
problem and
endocrine abn
approximately
GPU machine
Use case
Challenge
A "normal" X-ray when it
below the mean for the p
samples from the Greulic
Opportunity
Drastically reducing bone
radiologist, forensics, age
Metrics
This model takes approxi
CPU.
Mean absolute error (MA
Contact us for details
<< Return to Marketplace
MRI brain tumor detetection and segmentat
Glioblastoma is an aggressive
tumor commonly found in th
hemispheres. Segmentation
from multimodal MRI scans (
T1c, and T2) is a challenging
consuming process. This mod
on a Xeon(CPU)/Titan X(GPU
segment the tumor into four
necrotic core, edema, enhan
enhancing tumor core from n
gray matter, white matter, an
fluid.
Use case
Challenge
MRI Brain tumor segmentation is a critical
better outcomes. The process involves seg
from normal brain anatomy. Due to the lar
and characterization of MRI brain tumors i
Opportunity
Detection and segmentation of a brain tu
enables productivity of radiologists.
Metrics
This model takes 4 minutes on a machine
a dice score of 0.82.
Contact us for details
Tuberculosis detection model
Although it is treatable, Tuberculosis (TB) is the leading of cause of death by infectious disease worldwide, with well over a million fatalities per year. Most of these deaths are the result of the condition going undetected; because they occur in under-served locations with limited access to imaging and Radiology services. Artificial Intelligence can be used to identify TB on chest X-rays creating a cost effective and accurate method to screen for the disease. This model takes approximately 5 seconds for each prediction on a CPU machine
Use case
Challenge
Despite the availability of excellent treatment options, mortality rate for patients in some countries is high due to delayed diagnosis and availability of qualified radiologists.
Opportunity
Automatic screening of the chest X-ray with high sensitivity, one that captures almost all Tuberculosis cases, reduces load on Radiologist, enabling better utilization of their time, and perhaps making the process of diagnosis less-subjective.
Metrics
5 seconds on a CPU machine. It has an AUC of 0.90
Contact us for details
<< Return to Marketplace
Publisher
Philips
Type
Image Analytics
Current Version
v1.0
Pricing
N/A
System requirements
CPU & GPU necessary on premise
Supported Languages
RESTful API
Terms and Conditions
2019 12th Paris Hematology Conference
PHILIPS

- A Typical Image Classification Example

- **Tuberculosis** (TB) has been a major global health challenge, especially in developing nations.

- Despite the availability of excellent treatment options, mortality rate for patient countries is high due to delayed diagnosis and availability of qualified radiologists.

- Acute need for an **automatic screening solution** based on Deep Learning algorithms with a good compromise between sensitivity (capture almost all TB cases) and specificity (low false positives)

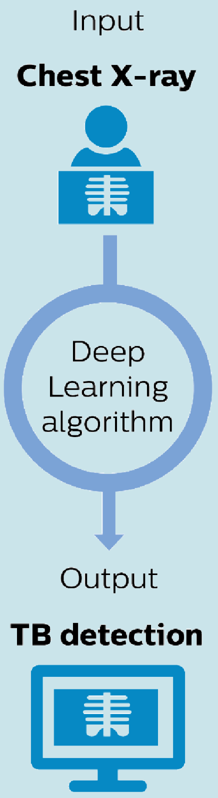

- Medical Context and Data Access are Key for DL

- Sidra Medical and Research Center is an all-digital academic medical center in Doha, Qatar that is collaborating with Philips to develop an automatic screening solution for TB based on Deep Learning algorithms
- Sidra provided
  – Vast set of data (**nearly 800,000 chest X-ray images**) with accurate labeling of the images indicating whether they are normal or abnormal
  – **Infrastructure** (16 GPU cluster) to create the model
  – **Validation data sets** to verify the algorithms

Training and Deployment
Collaborative approach

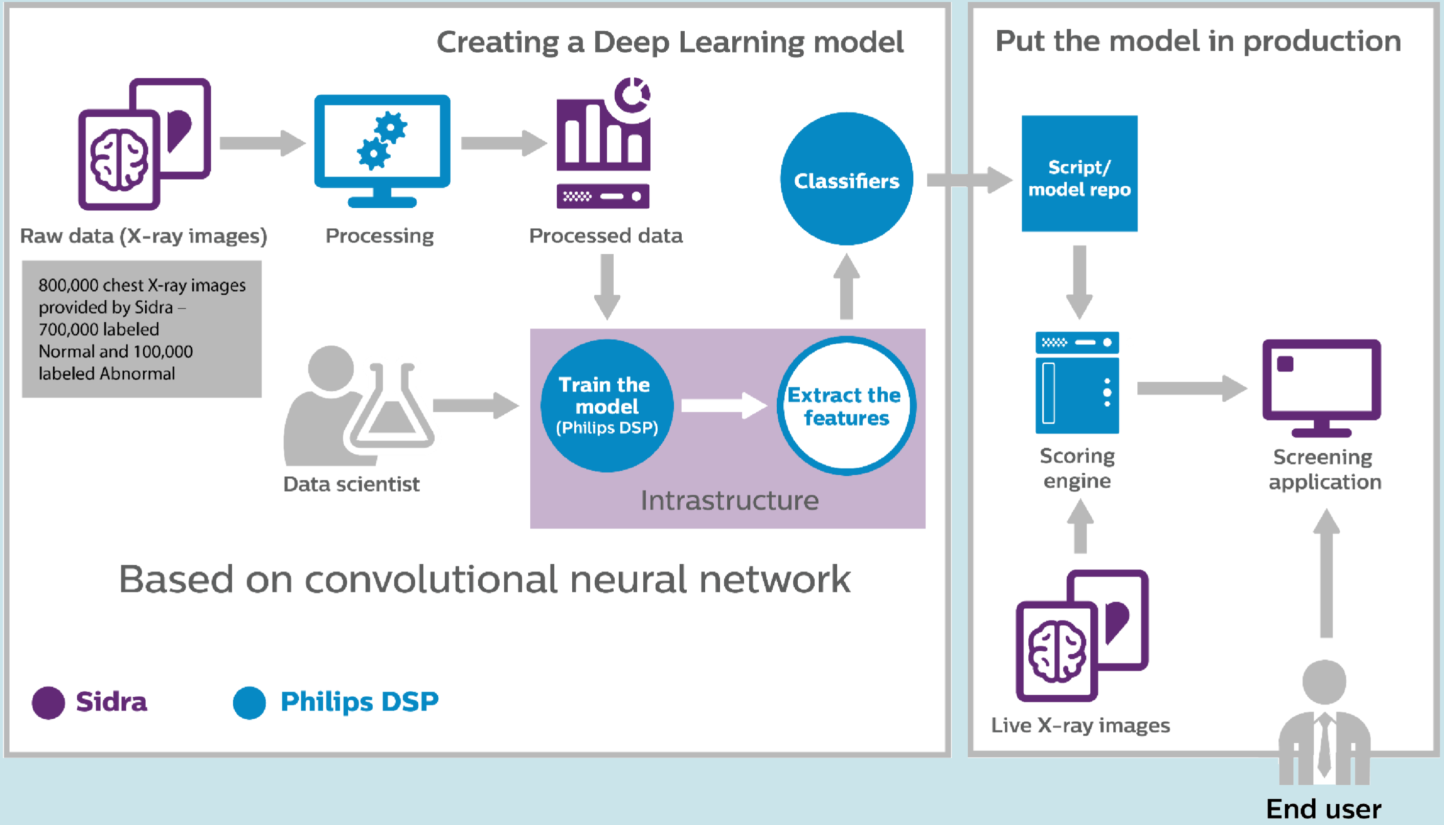

Results

Validation performed on data sets provided by Sidra as well as publicly available data sets

- Blinded data set: 16311 Normal + 3457 Abnormal
- Public data sets: 322 Normal + 312 Abnormal

Overall performance: **Specificity : 1-(FP/Normal) = 0.90**

2 seconds/image **Sensitivity : 1-(FN/Cases) = 0.90**

Results	Normal		TB case	
	Sidra	Public	Sidra	Public
Predicted condition +ve	13711	289	304	30
Predicted condition -ve	2600	33	3153	282

DL-based liver Contour extraction

CNN are very good
for segmentation tasks

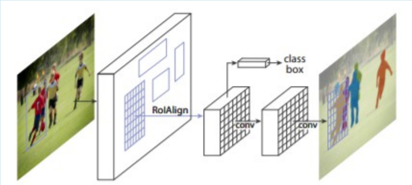

even with complex
anatomy or pathology
(polycystic liver)

Image from Pr Valette
HCL, Lyon, France

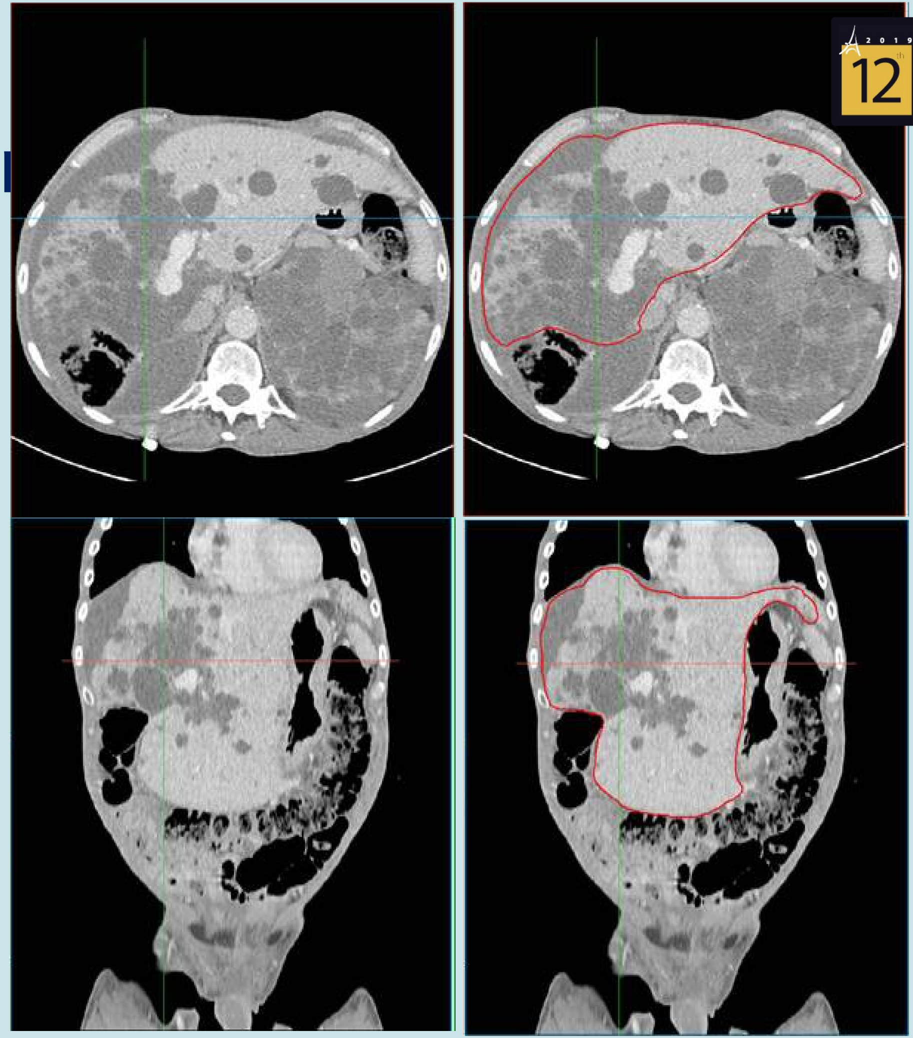

DL-based liver Contour extraction

CNN are very good for segmentation tasks

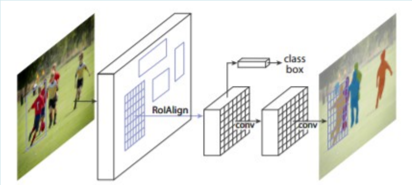

even with complex anatomy or pathology (metastatic liver)

Image from Pr Vilgrain
Hôpital Beaujon, Clichy, France

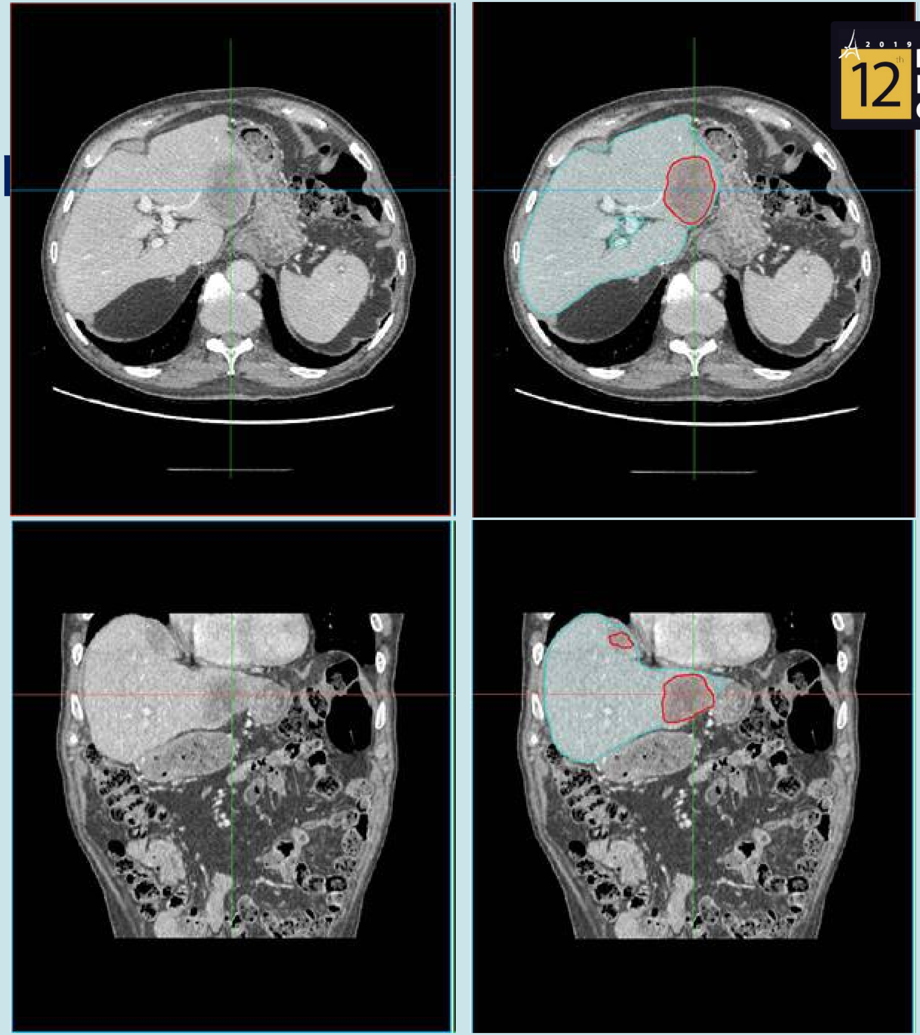

Digital, Computational Oncology

Data integration to optimise patient care

Patient Health

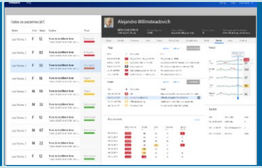

Family, environment, clinical data

Radiolog y

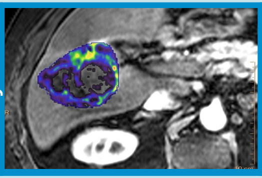

Analysis of organs, vessels, tumors, metabolism

Digital Pathology

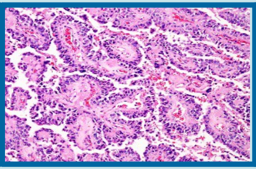

Tissue and cell analysis

Genomics

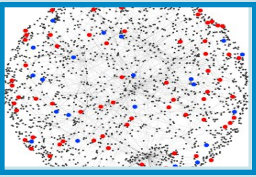

Full genome, transcriptomics, proteomics

Multidisciplinary Team (MDT)

Definitive Diagnosis

Personalized treatment

2019 12th Paris Hepatology Conference

PHILIPS

Our Objective: The Quadruple Aim

Patient Experience, Population Health, Reducing Costs, Care Team Well-Being

Artificial intelligence along the health continuum
Solutions to optimize the Quadruple Aim


Personalized health programs

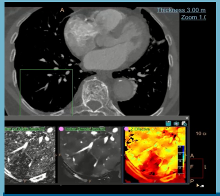


Advanced visualization

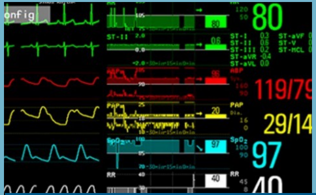
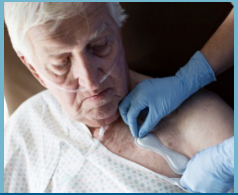

Predictive monitoring

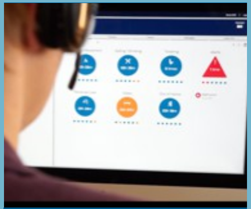

Télémédecine

Minimally invasive treatment

360° Patient

Quantification

Genomics

Adaptive Interfaces

PHILIPS