


University of Luxembourg

Multilingual. Personalised. Connected.

Machine Learning in a Nutshell

Brief Presentation at ISED 2019



Dr. Tegawendé F. BISSYANDE

Research Scientist, [SnT at University of Luxembourg](#)

bissyande [@Google Scholar](#) [@DBLP](#) [@Orbi.lu](#)

Research interests:

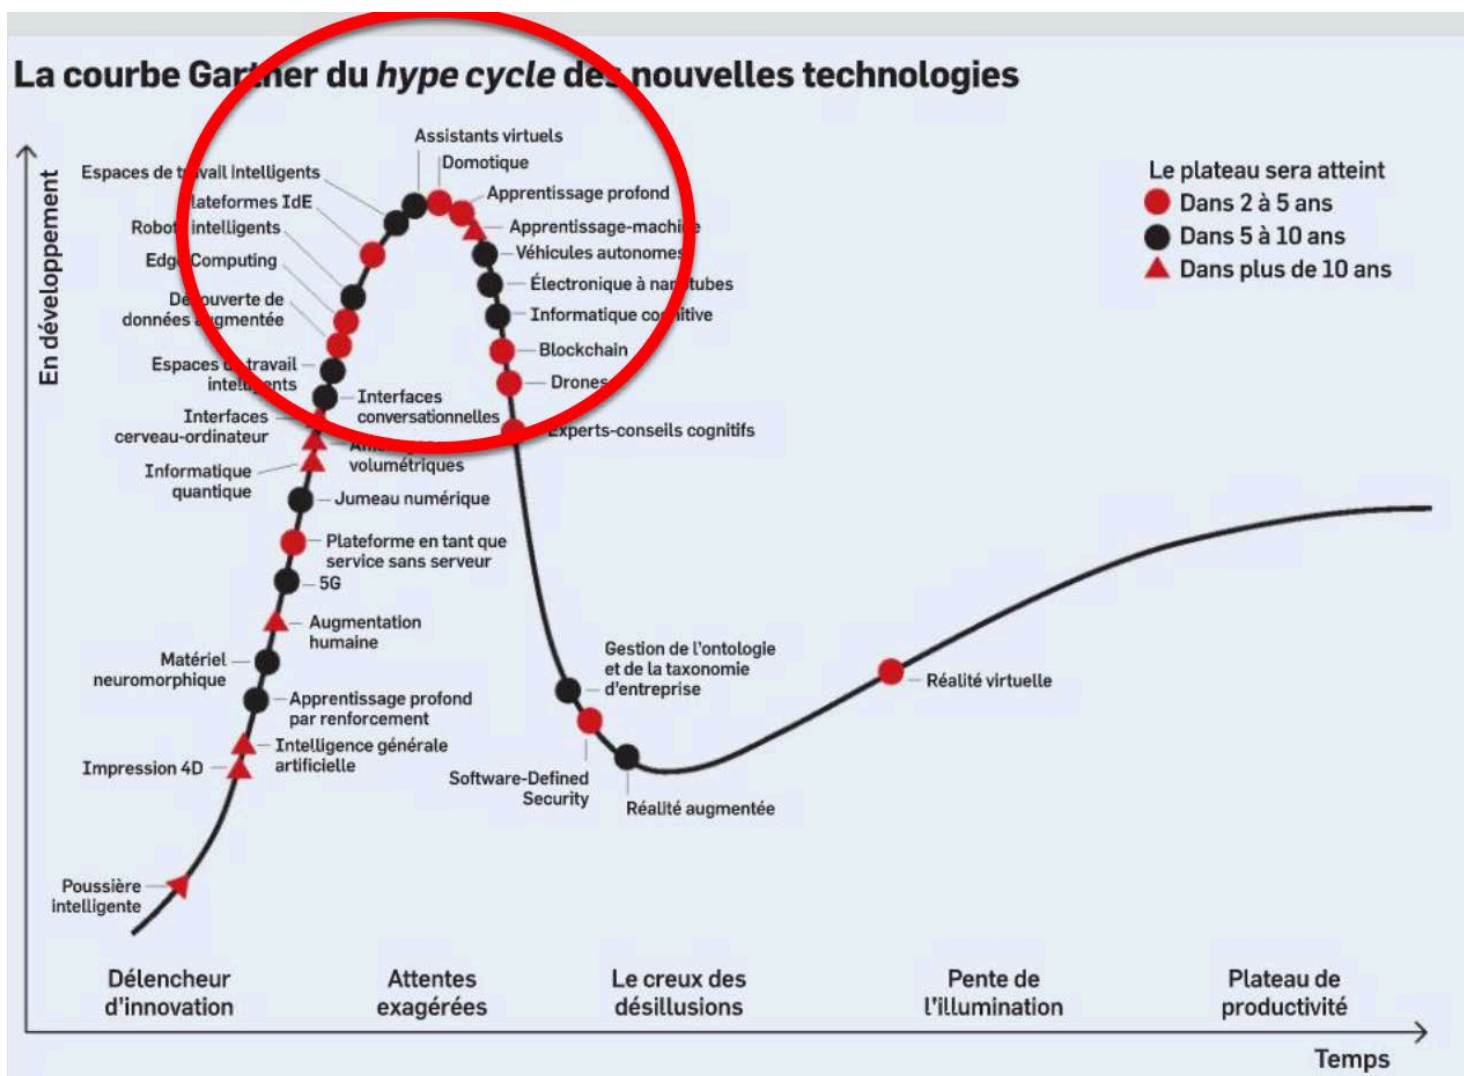
- Software Debugging (especially bug localization and **program repair**),
- Software Security (especially **malware detection** and analysis),
- **Code Search** (both free-form and semantic code-to-code),
- and more broadly software engineering and cyber-security.



I use and develop Machine Learning methods
Applied to Security and Software.

→ We are now investigating **industrial applications** for FinTech, RegTech, Cold management

The Hype




AI – What are we talking about?



- COMPANION
- FRIEND
- HELP

AI – What are we talking about?




- EVIL
- POWER
- DANGER

AI – What are we talking about?



Extremely Good at a Specific Task!

AI – What are we talking about?




Détection d'images et voitures autonomes :



2015 : Les performances obtenues sont équivalentes à celles des humains.

2012-2015 : percée en vision

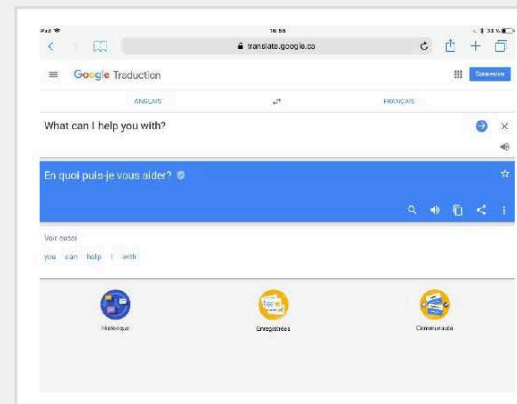


Reconnaissance vocale :



Garner prédit que d'ici 2020, 75% des ménages américains auront un assistant virtuel chez eux.

2010 à 2012 : percée importante



La traduction automatique :



La traduction neuronale est aussi le modèle utilisé par Microsoft et Systran.

Malgré des percées importantes, la traduction neuronale est capable d'erreurs importantes qu'un traducteur humain ne ferait jamais.

AI – Myths and Facts

Myth:

Superintelligence by 2100 is inevitable

Mon	Tue	Wed	Thur	Fri	Sat	Sun
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	✓ 21	22	23	24	25
26	27	28	29	30		

Fact:

It may happen in decades, centuries or never: AI experts disagree & we simply don't know



Myth:

Superintelligence by 2100 is impossible

Myth:

Only Luddites worry about AI



Fact:

Many top AI researchers are concerned




Mythical worry:

AI turning evil



Actual worry:

AI turning competent, with goals misaligned with ours




Mythical worry:

AI turning conscious

Myth:

Robots are the main concern



Fact:

Misaligned intelligence is the main concern: it needs no body, only an internet connection



Myth:

AI can't control humans



Fact:

Intelligence enables control: we control tigers by being smarter



Myth:

Machines can't have goals



Fact:

A heat-seeking missile has a goal



Mythical worry:

Superintelligence is just years away

PANIC!




Actual worry:

It's at least decades away, but it may take that long to make it safe

PLAN AHEAD!



AI – Tension / We don't know / We do not agree!



AI – Most people do not understand it...

 **Darren Cunningham** @dcunni · 25 Jul 2017
Zuckerberg blasts @elonmusk warnings against artificial intelligence as 'pretty irresponsible' bizjournals.com/sanjose/news/2... @svbizjournal #ai



Zuckerberg blasts Musk warnings against AI as 'pretty irresponsible'
"People who are naysayers and try to drum up these doomsday scenarios — I just, I don't understand it," the Facebook CEO said. "It's really negative
bizjournals.com

102 717 1.5K

 **Elon Musk** ✓
@elonmusk [Follow](#)

Replying to @dcunni @svbizjournal


I've talked to Mark about this. His understanding of the subject is limited.

12:07 AM - 25 Jul 2017

8,907 Retweets 34,808 Likes

1.2K 8.9K 35K

AI – Wait, Let's Look at the definition



Artificial intelligence is intelligence exhibited by machines, rather than humans or other animals. The field of AI research defines itself as the study of "intelligent agents": any device that perceives its environment and takes actions that maximize its chance of success at some goal

AI – Are we talking about the same thing?



Artificial Narrow Intelligence (ANI): Machine intelligence that equals or exceeds human intelligence or efficiency **at a specific task**.



Artificial General Intelligence (AGI): A machine with the ability to **apply intelligence to any problem**, rather than just one specific problem (*human-level intelligence*).



Artificial Superintelligence (ASI): An **intellect that is much smarter than the best human brains** in practically every field, including scientific creativity, general wisdom and social skills.

AI – By the way, What is intelligence?

Learn

Create a system of knowledge, and be able to integrate new knowledge

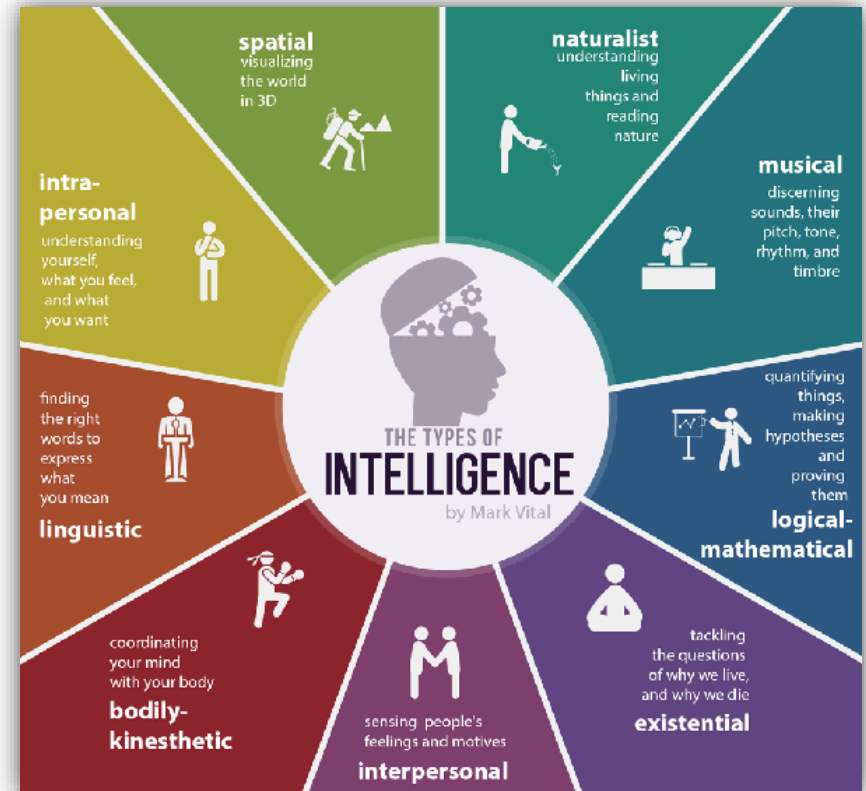
Reason, Deduce, Anticipate

From the system of knowledge, and based experience data, be able to produce new knowledge


Possess a history

Possess a consciousness

Possess sentiments



AI – Like many modern tech, it is boosted by WAR




Modern AI is initiated by **Alan Turing**

- WWII: How to crack Enigma? With a “learning” Bombe machine
- John McCarthy coined the term in 1956

Imitation game: “A machine that would converse with humans, without the humans knowing that it is a machine”



Turing, Alan (c. 1941). "Paper on Statistics of Repetitions".
The National Archives (United Kingdom): HW 25/38.




Machine Learning is NOT Deep Learning NOR AI or AGI

Two Major AI Techniques

- **Logic and Rules-Based Approach**



- **Machine Learning (Pattern-Based Approach)**





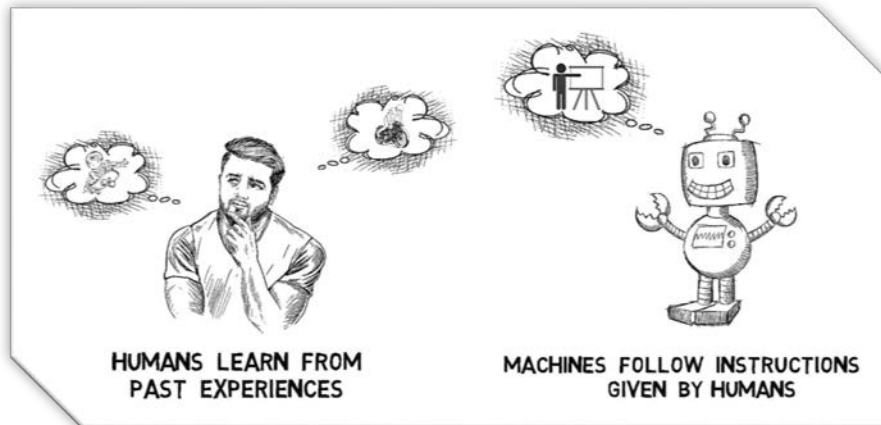
Why learn?



**HUMANS LEARN FROM
PAST EXPERIENCES**

**MACHINES FOLLOW INSTRUCTIONS
GIVEN BY HUMANS**

Why learn?



Traditional Programming



Machine Learning



When do we use machine learning?

Navigation on Mars



Human expertise
does not exist

Speech recognition




Human can't explain
their expertise

Personalized medicine



Models must be
customized

Genomics



Models are based
on huge data

→ Learning isn't always useful
Would you « learn » to calculate payroll?

A classical example of task that requires machine learning

It is very hard to say what makes a 2



“Learning is any process by which a system **improves** performance from **experience**.”

- *Herbert Simon*


Definition by Tom Mitchell (1998):

Machine Learning is the study of algorithms that



- improve their performance P
- at some task T
- with experience E .

A well-defined learning task is given by $\langle P, T, E \rangle$.

So we need a lot of data...



... from various perspectives!



Images and movies taken from Sebastian Thrun's multimedia website.

Motivating Example: Learning to Filter Spam

Spam - is all email the user does not want to receive and has not asked to receive


T → Identify Spam Emails

P → % of spam emails that were filtered / % of ham (non-spam) emails that were incorrectly filtered-out

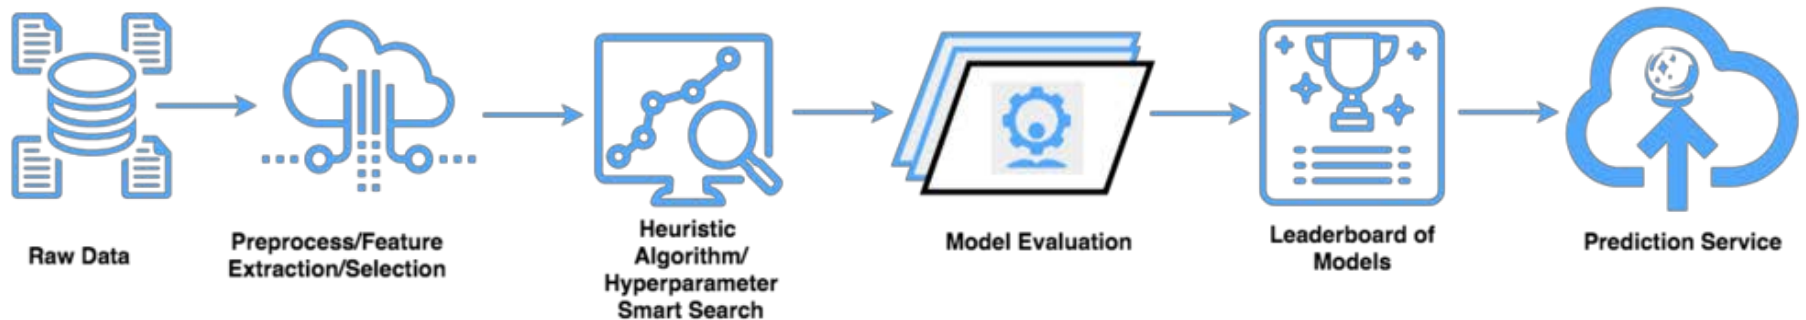
E → a database of emails that were labelled by users



Motivating Example: Learning to Filter Spam

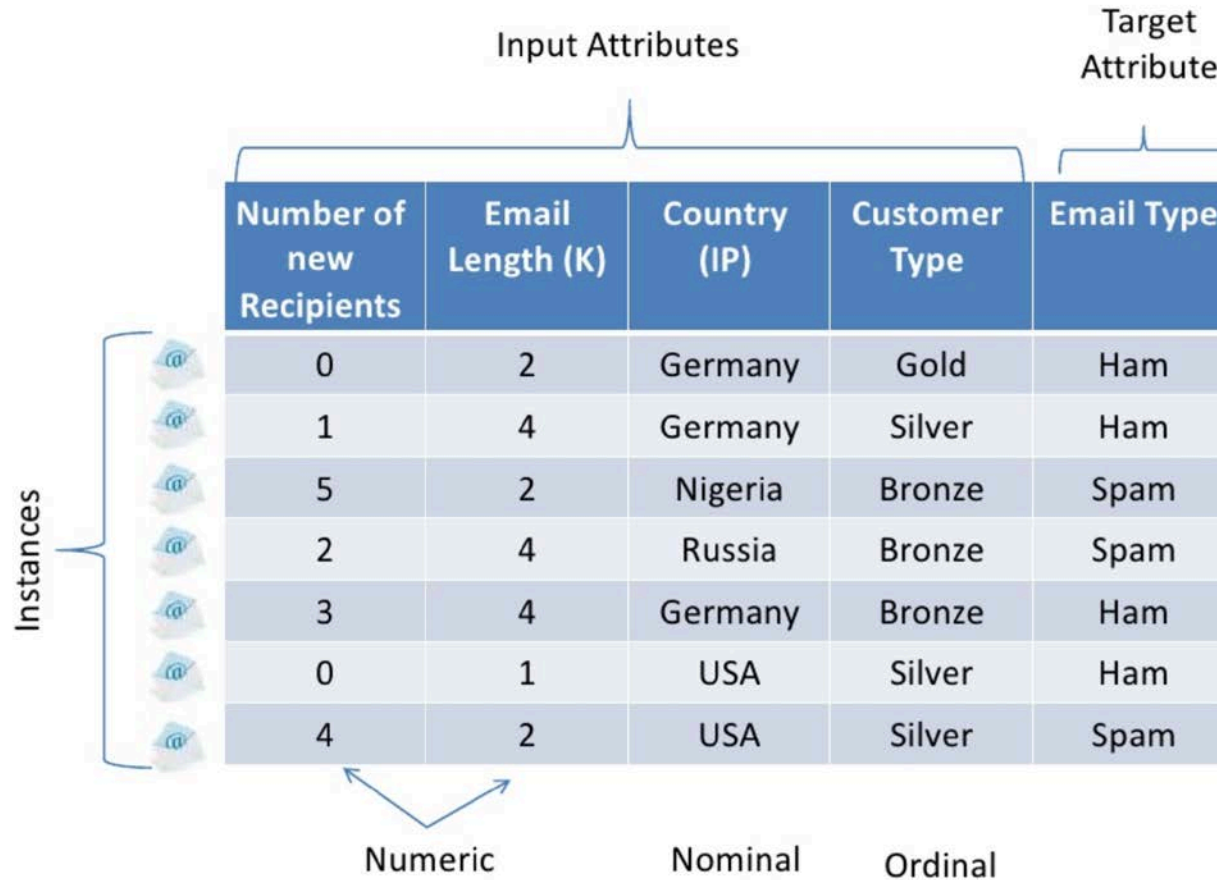


Motivating Example: Learning to Filter Spam



- # of recipients
- Email length
- New sender?
- Sender IP (country)

Motivating Example: Learning to Filter Spam




Input Attributes				Target Attribute
Number of new Recipients	Email Length (K)	Country (IP)	Customer Type	Email Type
0	2	Germany	Gold	Ham
1	4	Germany	Silver	Ham
5	2	Nigeria	Bronze	Spam
2	4	Russia	Bronze	Spam
3	4	Germany	Bronze	Ham
0	1	USA	Silver	Ham
4	2	USA	Silver	Spam

Numeric Nominal Ordinal


Motivating Example: Learning to Filter Spam

Non Spam (Ham)




Motivating Example: Learning to Filter Spam

Non Spam (Ham)




Motivating Example: Learning to Filter Spam

Non Spam (Ham)




Motivating Example: Learning to Filter Spam

Non Spam (Ham)




Motivating Example: Learning to Filter Spam

Non Spam (Ham)




Motivating Example: Learning to Filter Spam

Non Spam (Ham)





Motivating Example: Learning to Filter Spam




Email Length

Motivating Example: Learning to Filter Spam



Discriminative power of features is key




Types of Learning

- **Supervised (inductive) learning**
 - Given: training data + desired outputs (labels)
- **Unsupervised learning**
 - Given: training data (without desired outputs)
- **Semi-supervised learning**
 - Given: training data + a few desired outputs
- **Reinforcement learning**
 - Rewards from sequence of actions

Supervised Learning: Regression


- Given $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$
- Learn a function $f(x)$ to predict y given x
 - y is real-valued == regression



data from G. Witt, Journal of Statistics Education, Volume 21, Number 1 (2013)


Unsupervised Learning

- Given x_1, x_2, \dots, x_n (without labels)
- Output hidden structure behind the x 's
 - E.g., clustering



Supervised Learning: Classification

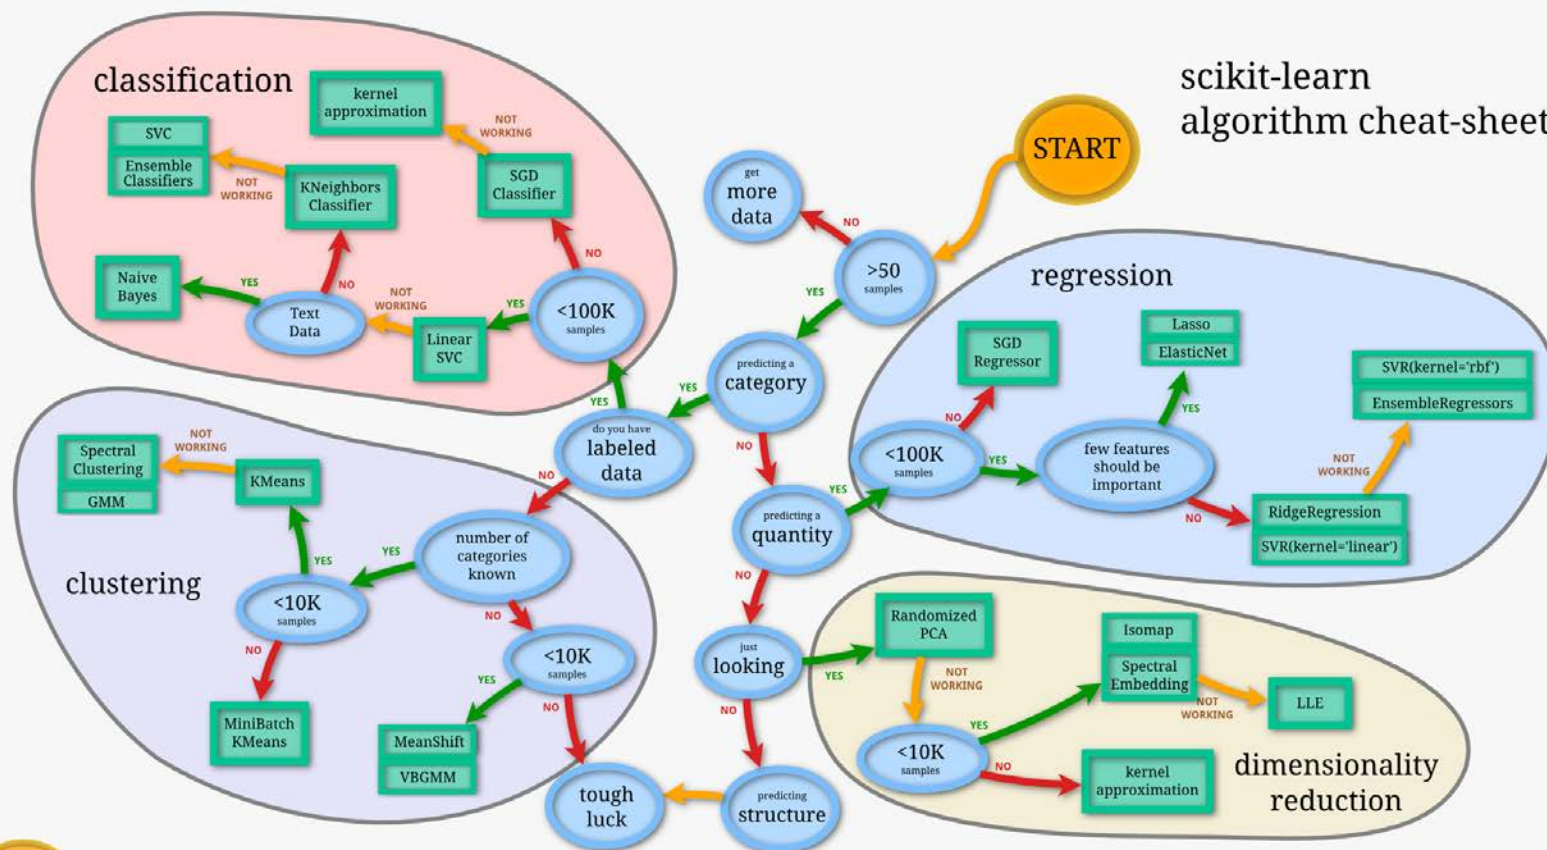
- Given $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$
- Learn a function $f(x)$ to predict y given x
 - y is categorical == classification




Based on example by Andrew Ng

Types of Traditional Learning

scikit-learn
algorithm cheat-sheet





Deep Learning

The Brain: An Amazingly Efficient "Computer"

10^{11} neurons, approximately

10^4 synapses per neuron

10 "spikes" go through each synapse per second on average

10^{16} "operations" per second

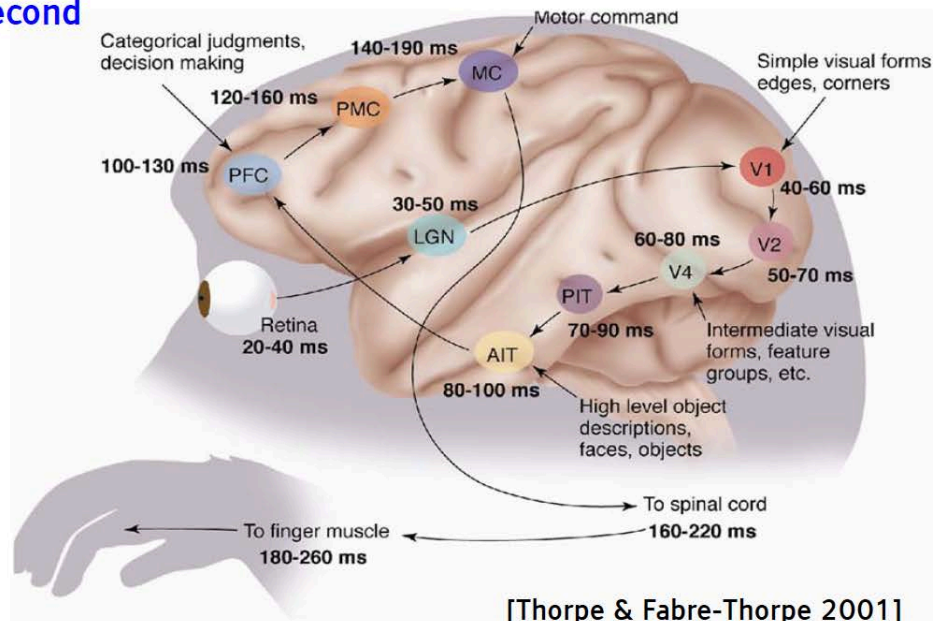
25 Watts

► Very efficient

1.4 kg, 1.7 liters

2500 cm²



► Unfolded cortex



The Brain: An Amazingly Efficient "Computer"


Intel Xeon Phi CPU

- ▶ 2×10^{12} operations/second
- ▶ 240 Watts
- ▶ 60 (large) cores
- ▶ \$3000



NVIDIA Titan-Z GPU

- ▶ 8×10^{12} operations/second
- ▶ 500 Watts
- ▶ 5760 (small) cores
- ▶ \$3000



Are we only a factor of 10,000 away from the power of the human brain?

- ▶ Probably more like 1 million: synapses are complicated
- ▶ A factor of 1 million is 30 years of Moore's Law
- ▶ 2045?

Can we build AI systems by copying the brain?

Are computers only a factor of 10,000 away from the power of the human brain?

- ▶ Probably more like 1 million: synapses are complicated
- ▶ A factor of 1 million is **30 years of Moore's Law**

Will computers be as intelligent as human by 2045?

- ▶ Compute power is not the whole story
- ▶ Moore's Law may not continue for that long
- ▶ We need to understand the **principles** of learning and intelligence



Getting inspiration from biology is a good thing

But blindly copying biology without understanding the underlying principles is doomed to failure

- ▶ Airplanes were inspired by birds
- ▶ They use the same basic principles for flight
- ▶ But airplanes don't flap their wings & don't have feathers



Can we build AI systems by copying the brain?

It's nice imitate Nature,

But we also need to **understand**

- ▶ How do we know which details are important?
- ▶ Which details are merely the result of evolution, and the constraints of biochemistry?

For airplanes, we developed aerodynamics and compressible fluid dynamics.

- ▶ We figured that feathers and wing flapping weren't crucial

QUESTION: What is the equivalent of aerodynamics for understanding intelligence?



L'Avion III de Clément Ader, 1897





(Musée du CNAM, Paris)

His "Eole" took off from the ground in 1890, 13 years before the Wright Brothers, but you probably never heard of it (unless you are french).

1957: The Perceptron (the first learning machine)

A simple simulated neuron with **adaptive** “synaptic weights”

- ▶ Computes a weighted sum of inputs
- ▶ Output is +1 if the weighted sum is above a threshold, -1 otherwise.



The Perceptron: a Trainable Classifier [Rosenblatt, 1957]


■ Example: classifying letters "A" from "B"

■ Learning: find the weight values that produce +1 for A and -1 for B

■ Training set: $(X^1, Y^1), (X^2, Y^2), \dots, (X^P, Y^P)$

■ Example: $(A, +1), (B, -1), (A, +1), (B, -1), (A, +1), (B, -1), \dots$

Supervised
Learning



The traditional model of pattern recognition

■ The traditional model of pattern recognition (since the late 50's)




- ▶ Fixed/engineered features (or fixed kernel) + trainable classifier



hand-crafted
Feature Extractor


"Simple" Trainable
Classifier

■ Perceptron (Cornell University, 1957)




Traditionnel vs deep models of pattern recognition

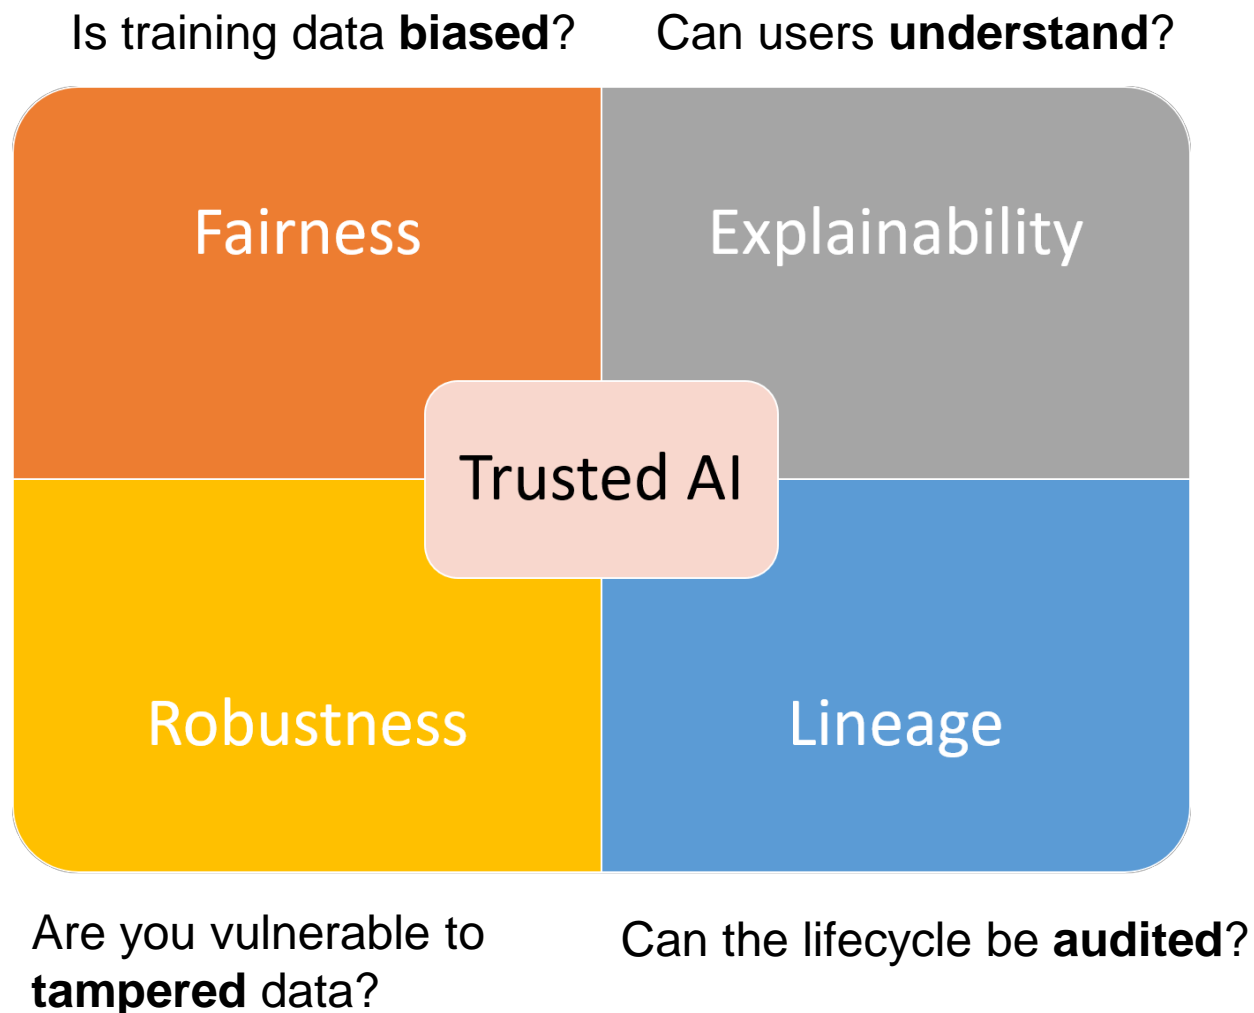
Machine Learning



Deep Learning



Before we leave... Beware of the TRUST problem



Stay connected with us!

