UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Deep Learning & Knowledge Representation

JDEV 2020

Frederic Precioso

Equipe Projet Comun Maasai – Laboratoire I3S UMR CNRS 7271 – INRIA

CONTEXT & VOCABULARY

What is Artificial intelligence?

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

How is Artificial intelligence defined?

JDEV

- The term ***Artificial Intelligence***, as a research field, was coined at the conference on the campus of Dartmouth College in the summer of **1956**, even though the idea was around since Antiquity: Hephaestus built automatons of metal to work for him or protect others, the Golem in Jewish folklore, etc.

- Closer to the Dartmouth conference but still before, the first manifesto on Artificial Intelligence, an unpublished report ***"Intelligent Machinery"***, written by Alan Turing in **1948**. He already distinguished two different approaches to AI, which may be termed ***"top-down"*** and ***"bottom-up"*** *(now more commonly called knowledge-driven AI and data-driven AI respectively).*

(sources: Wikipedia, https://www.greeklegendsandmyths.com/automatons.html ,
http://www.alanturing.net/turing_archive/pages/Reference%20Articles/what_is_AI/What%20is%20AI02.html
Stanford Encyclopedia of Philosophy: https://plato.stanford.edu/entries/artificial-intelligence/*)*

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

How is Artificial intelligence defined?

JDEV

- *"top-down"* or ***knowledge-driven AI***
  - cognition = high-level phenomenon, independent of low-level details of implementation mechanism
  - Evolutionary Algorithms (1954,1957, 1960), Knowledge Representation, Reasoning (1959,1970), Expert Systems (1970), Logic, Automata, Intelligent Agent Systems (1990)…

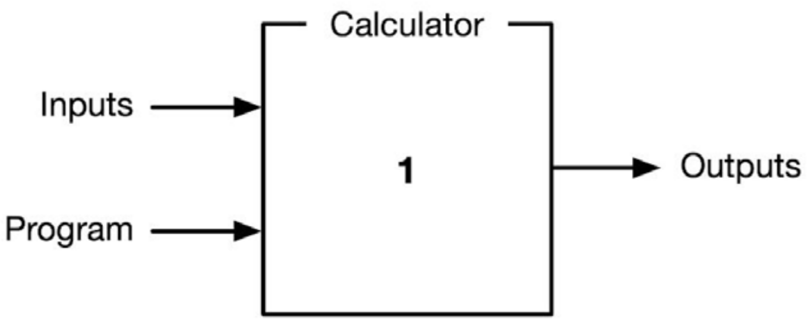

(1) Hypothetical-deductive machines

(Figure from: Neurons spike back The invention of inductive machines and the artificial intelligence controversy", D. Cardon, J.-P. Cointet, A. Mazières, Translated by Elizabeth Libbrecht In Réseaux Volume 211, Issue 5, 2018, pages 173 to 220)

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

How is Artificial intelligence defined?

JDEV

- *"bottom-up" or data-driven AI*
  - opposite approach, start from data to build incrementally and mathematically mechanisms taking decisions
  - First neuron (1943), first neural network machine (1950), neucognitron (1975), Decision Trees (1983), Backpropagation (1984-1986), Random Forest (1995), Support Vector Machine (1995), Boosting (1995), Deep Learning (1998/2006)…

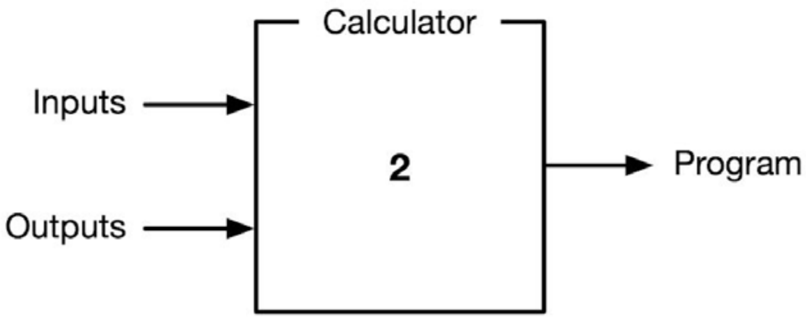

(2) inductive machines

(Figure from: Neurons spike back The invention of inductive machines and the artificial intelligence controversy", D. Cardon, J.-P. Cointet, A. Mazières, Translated by Elizabeth Libbrecht In Réseaux Volume 211, Issue 5, 2018, pages 173 to 220)

Why Artificial Intelligence is so difficult to grasp?

- Frequently, when a technique reaches **mainstream use**, it is **no longer considered as artificial intelligence**; this phenomenon is described as the ***AI effect***: "AI is whatever hasn't been done yet.“ (***Larry Tesler's Theorem***)
  -> e.g. Path Finding (GPS), Chess electronic game, Alpha Go…

- Consequently, AI domain is continuously evolving and so very difficult to grasp.

So what is Machine Learning?

Machine Learning

$$\begin{pmatrix}\mathbf{x}\end{pmatrix} \xrightarrow{f(\mathbf{X},\alpha)\ ?} y$$

$\begin{pmatrix}\mathbf{x}\end{pmatrix}$

y

Face detection

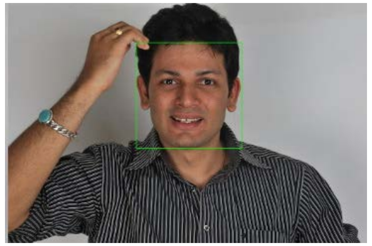

Scores prediction

Sport bets

Voice recognition

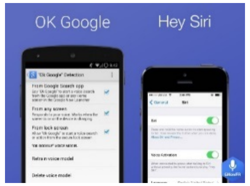

Machine Learning

$$\begin{pmatrix} \mathbf{X} \end{pmatrix} \xrightarrow{f(\mathbf{X},\alpha)\ ?} y$$

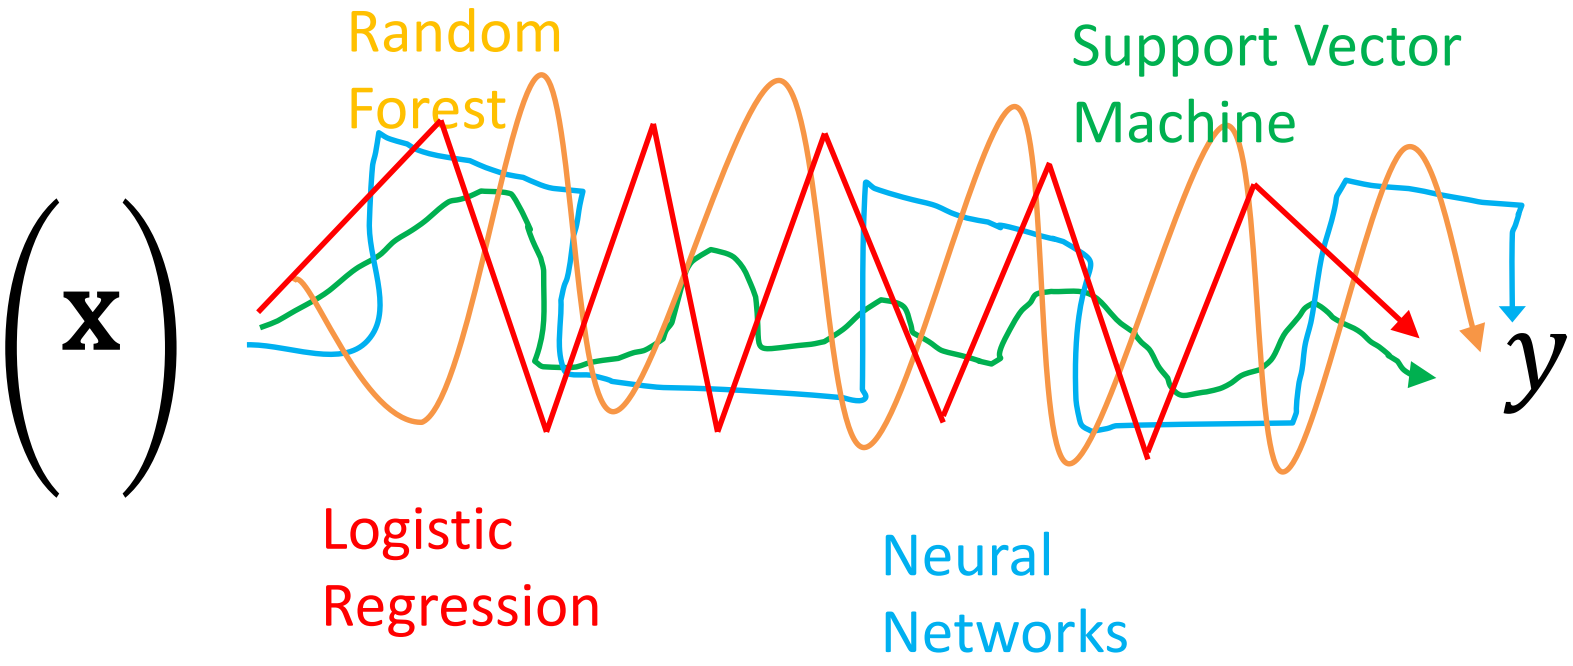

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Machine Learning

$$\left(\mathbf{x}\right) \xrightarrow{f\,(\mathbf{X},\alpha)\;?} y$$

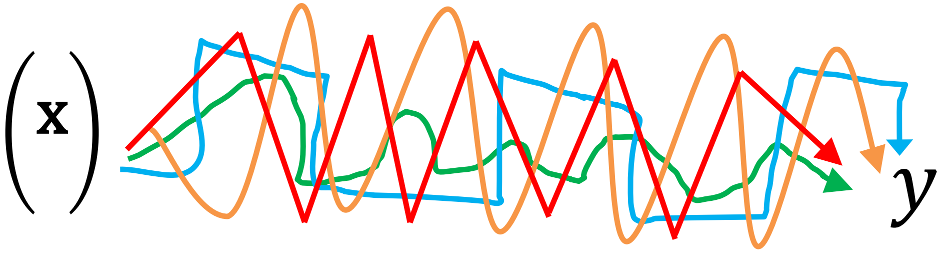

$\neq$ ***AI***

ML

Francis Bach at *Frontier Research and Artificial Intelligence Conference (ERC Conference)*: **“Machine Learning is not AI”**
(https://erc.europa.eu/sites/default/files/events/docs/Francis_Bach-SEQUOIA-Robust-algorithms-for-learning-from-modern-data.pdf
https://webcast.ec.europa.eu/erc-conference-frontier-research-and-artificial-intelligence-25#)

Machine Learning

$$\left(\mathbf{X}\right) \xrightarrow{f\,(\mathbf{X},\alpha)\;?} y$$

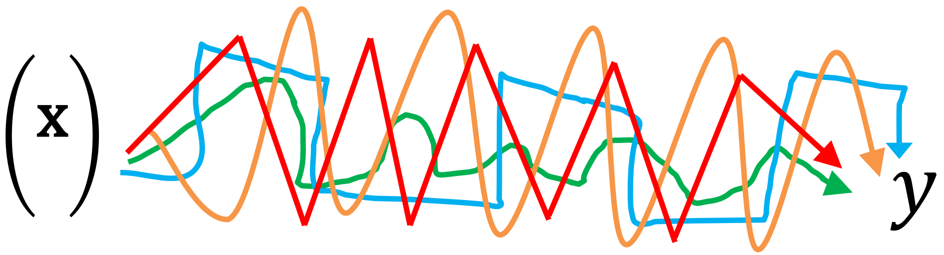

$\Leftrightarrow$ "Weather Forcasting"

ML

$\rightarrow$ *Trolley dilemma next slide*

univ-cotedazur.fr

Trolley dilemma
(a two-year old kid's solution)

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

AI vs Machine Learning vs Deep Learning

JDEV

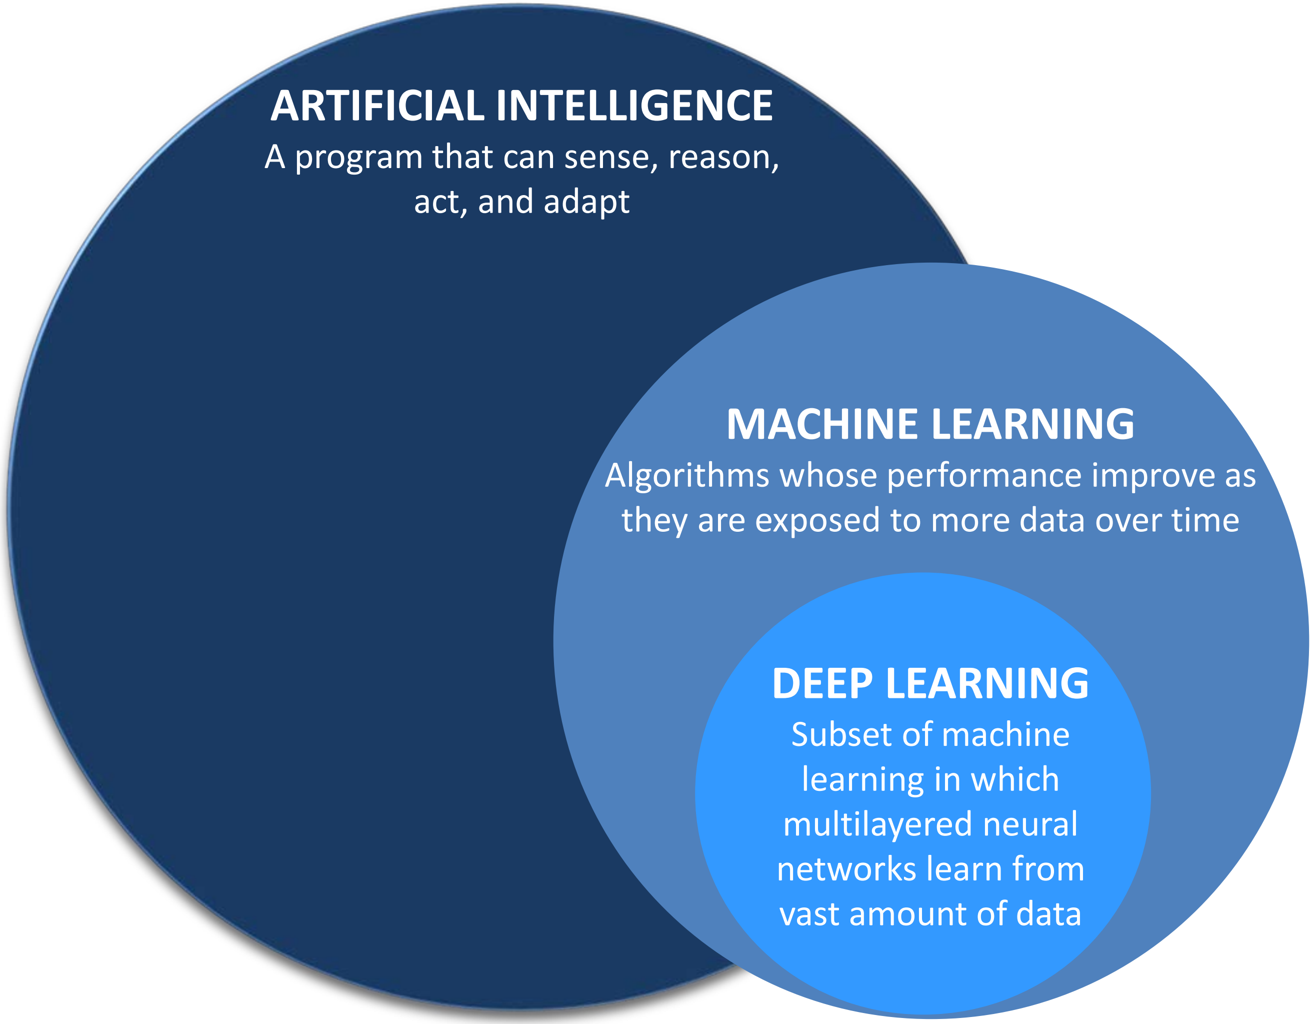

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

AI vs Machine Learning vs Deep Learning

JDEV

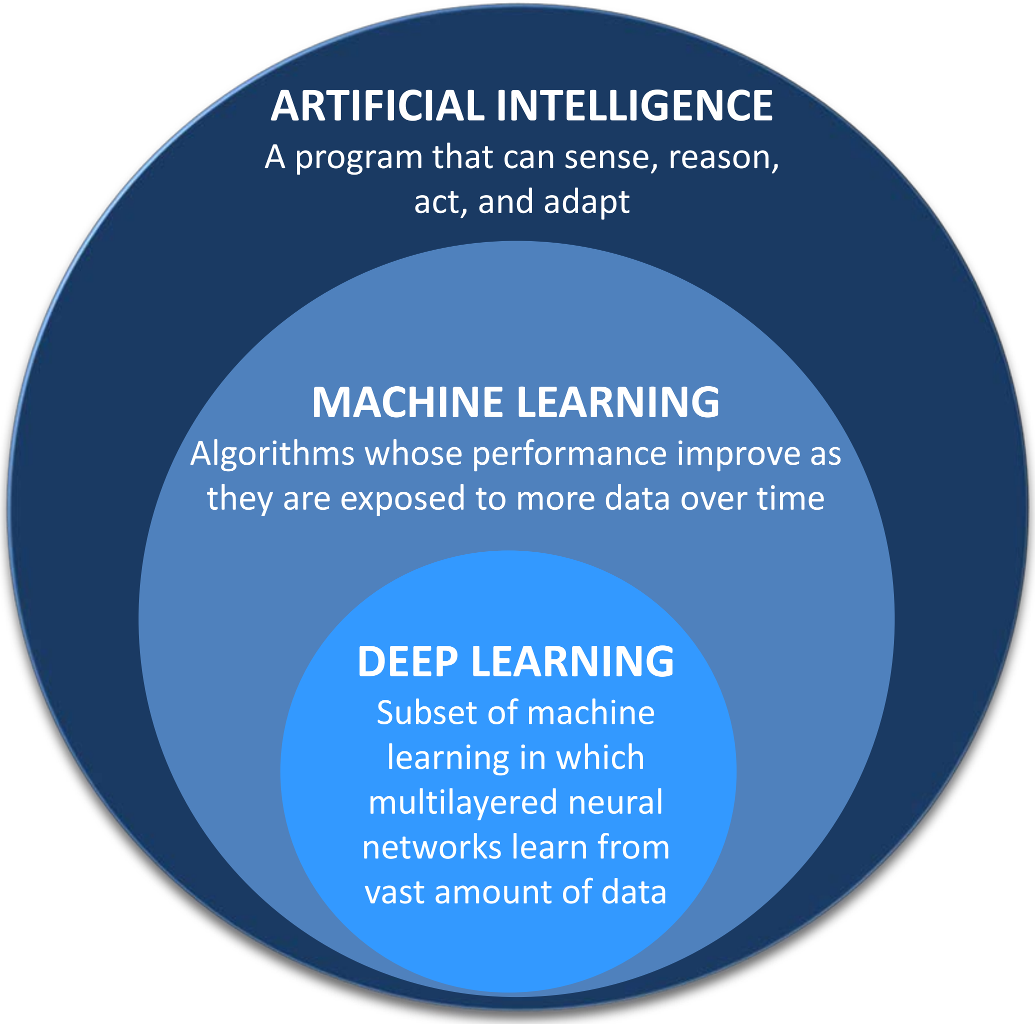

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

A BRIEF OF DEEP LEARNING

From one neuron to network

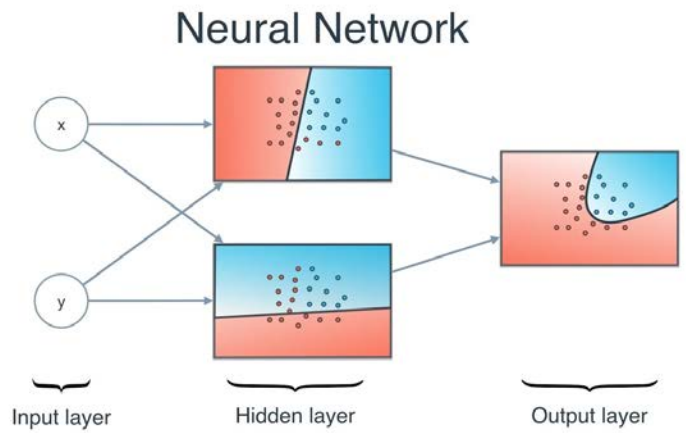

@tachyeonz: A friendly introduction to neural networks and deep learning.

At training you want to set the weights,
so that your training samples are correctly classified:

Minsky, M. and Papert, S. (1969). Perceptrons: An Introduction to Computational Geometry. MIT Press.

Training

At training you want to set the weights:

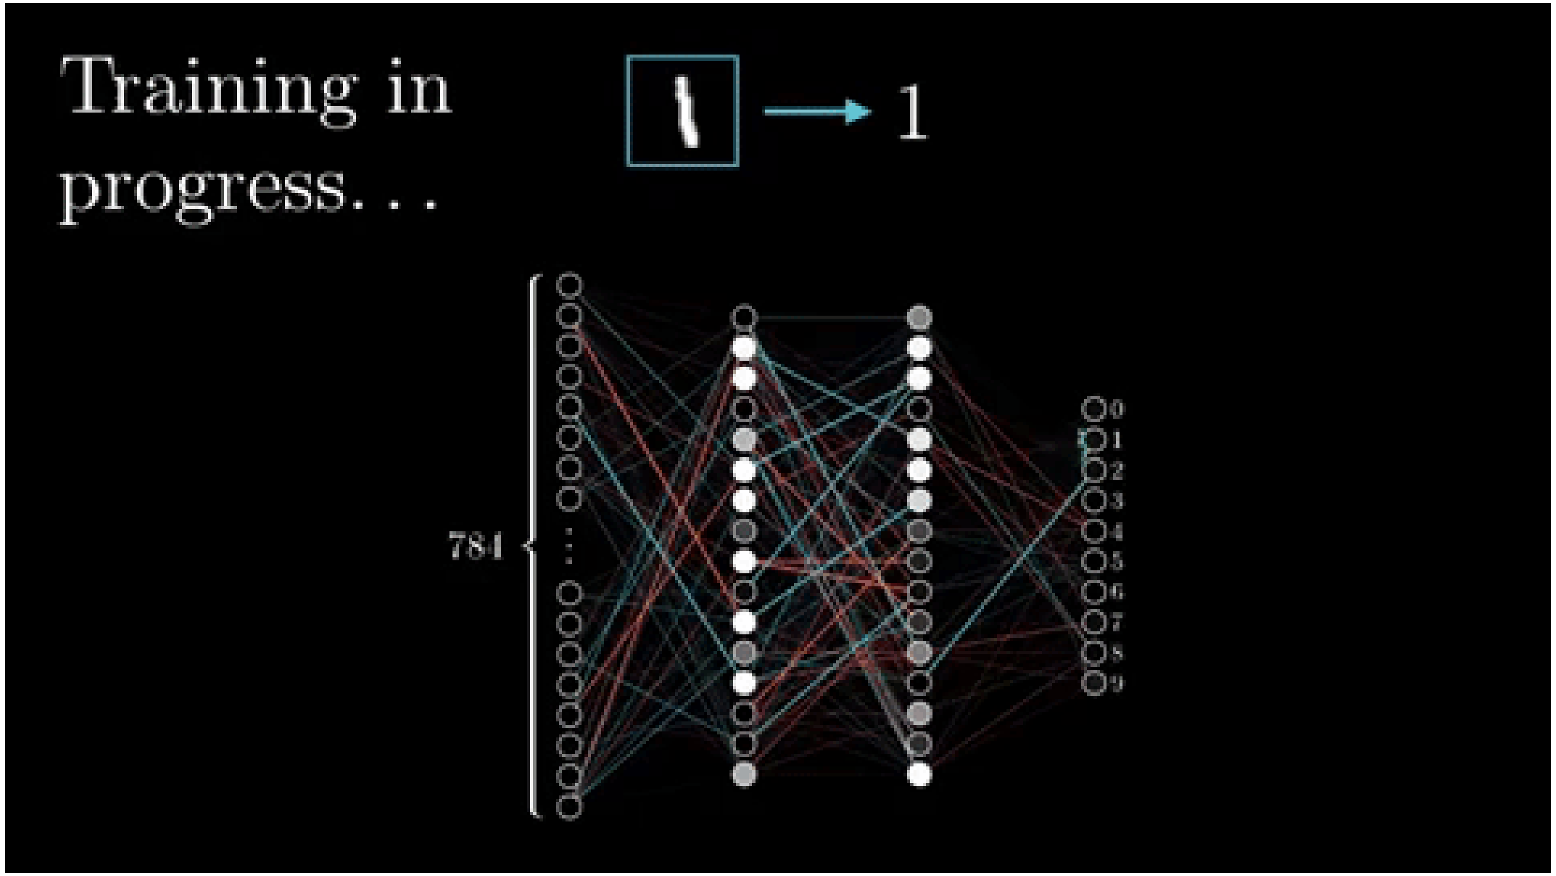

Source: https://www.3blue1brown.com/neural-networks

Prediction

At testing the weights do not evolve anymore:

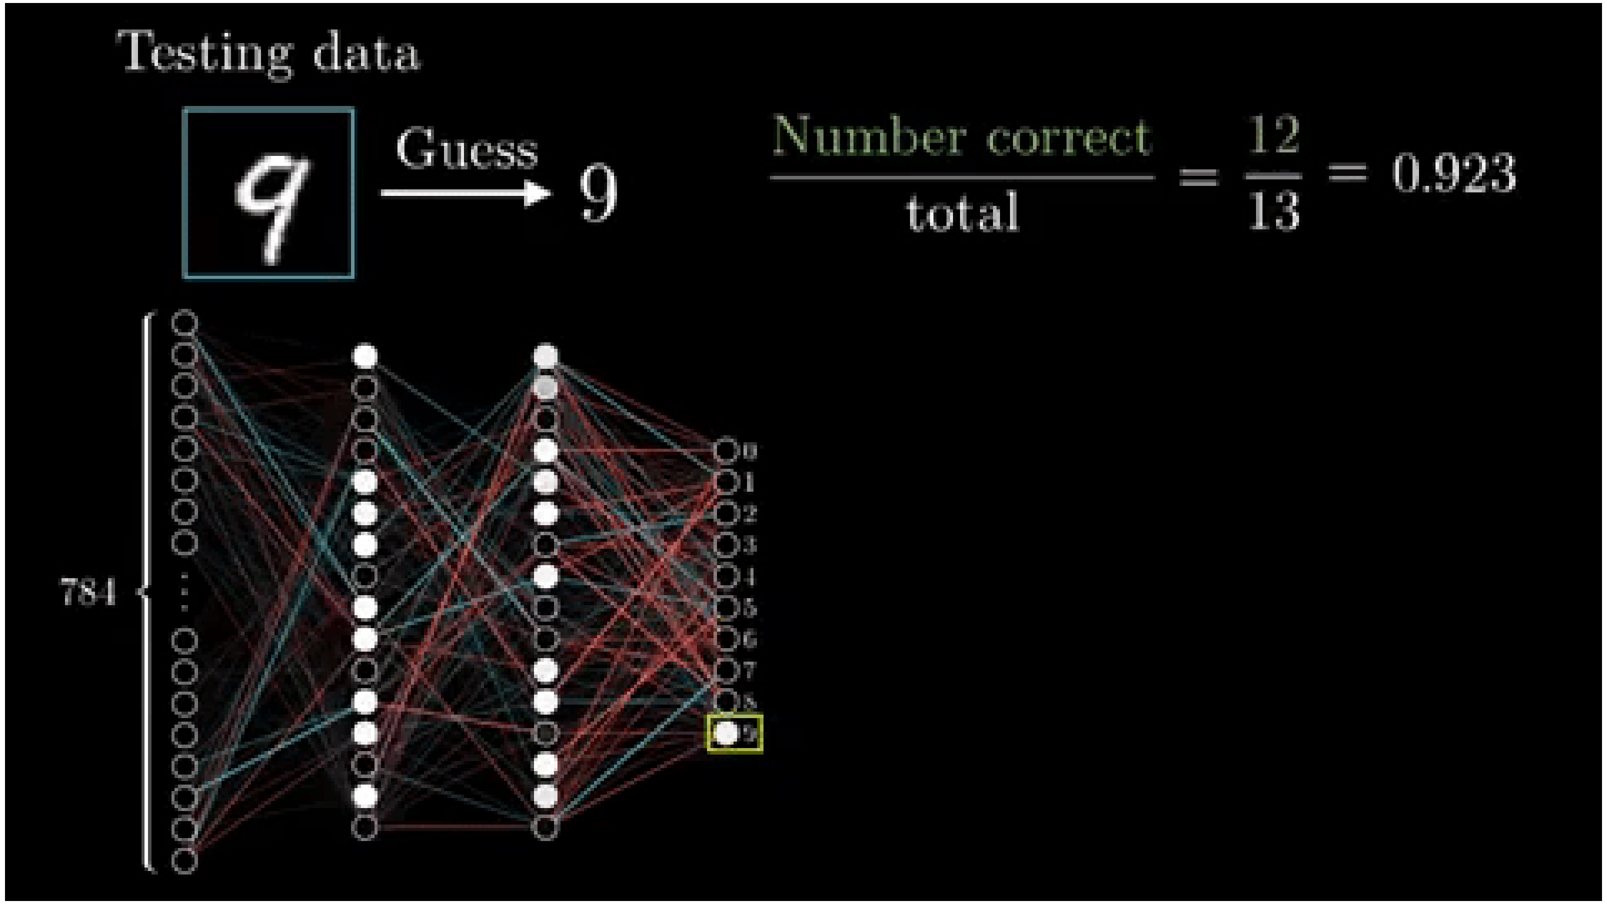

Source: https://www.3blue1brown.com/neural-networks

Structure the network?

- Can we put any structure reducing the space of exploration and providing useful properties (invariance, robustness…)?

$$y = \underbrace{s(w_{13}\overbrace{s(w_{11}x_1 + w_{21}x_2 - w_{01})}^{z_1} + w_{23}\overbrace{s(w_{12}x_1 + w_{22}x_2 - w_{02})}^{z_2}}_{z_3} - w_{03})$$

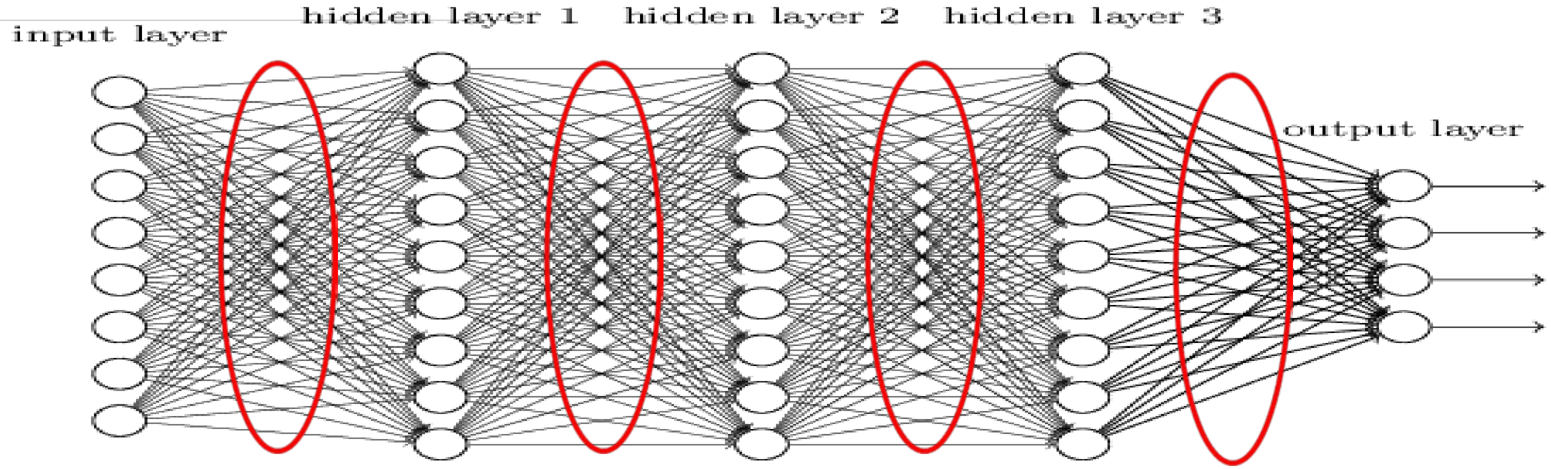

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Example with spatial invariance (Scale, Translation,…)

The Mammalian Visual Cortex is structured

- The ventral (recognition) pathway in the visual cortex has multiple stages
- Retina - LGN - V1 - V2 - V4 - PIT - AIT ....
- Lots of intermediate representations

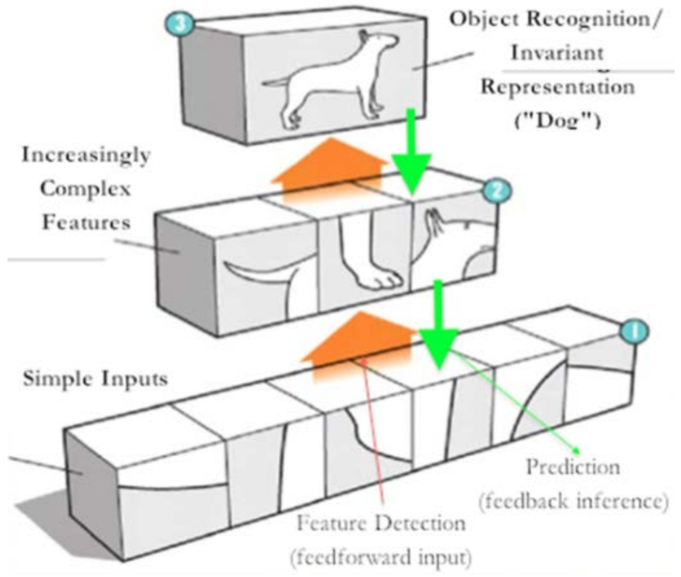

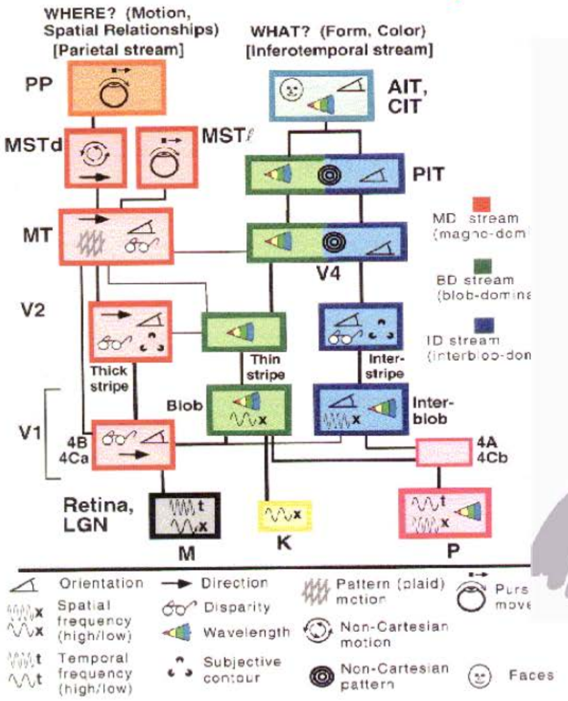

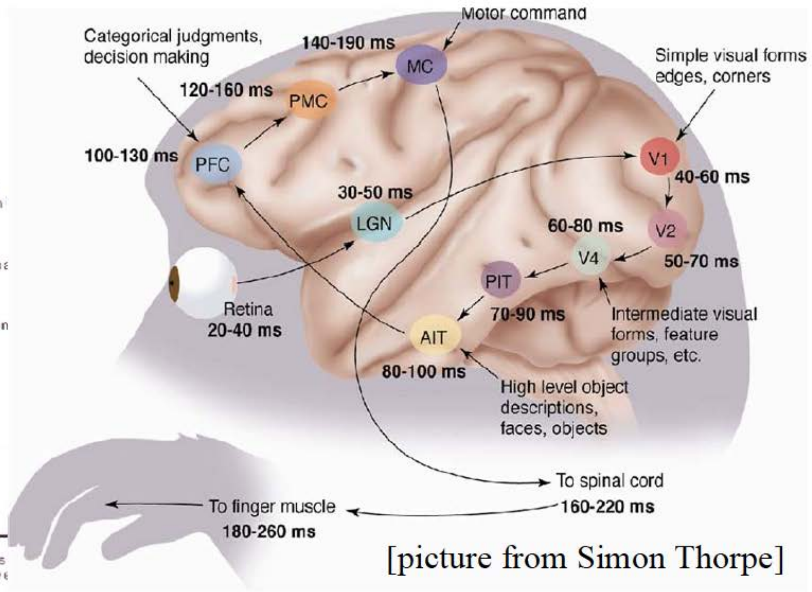

[Gallant & Van Essen]

Deep spatial representation

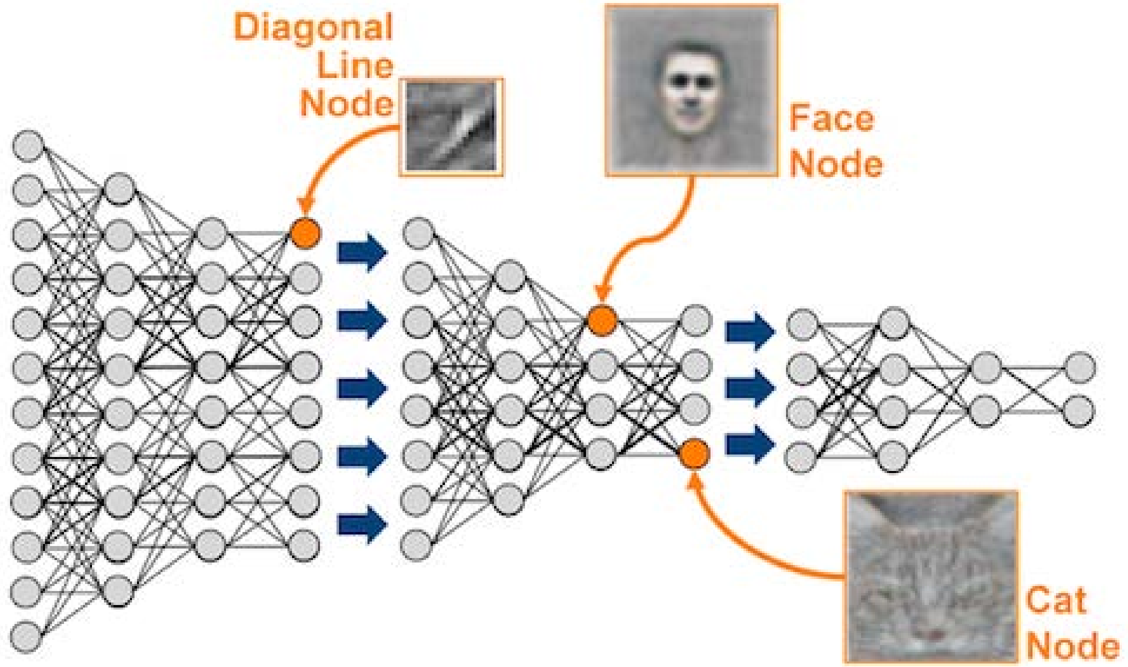

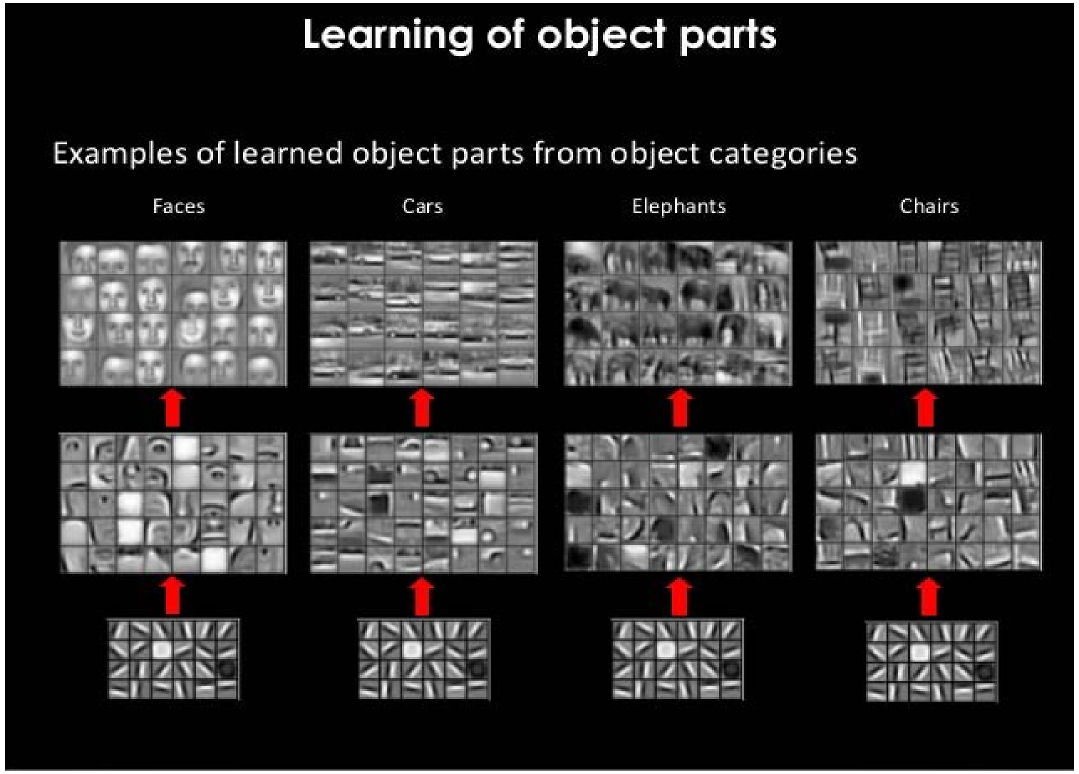

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

Deep spatial representation

- Deep Networks are as good as humans at recognition, identification...

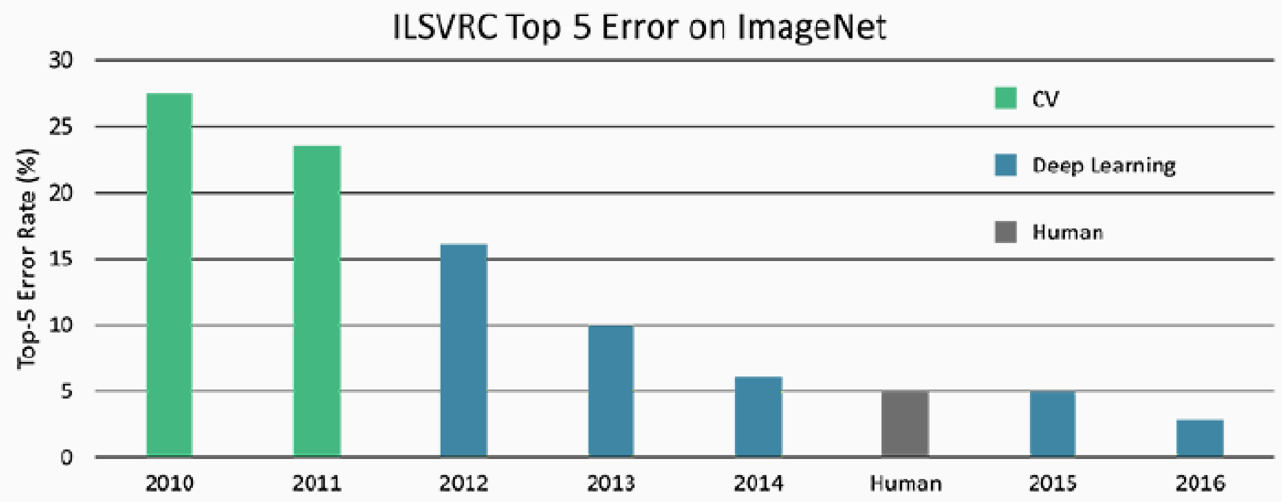

How much does a deep network understands those tasks?

Transfer Learning!!

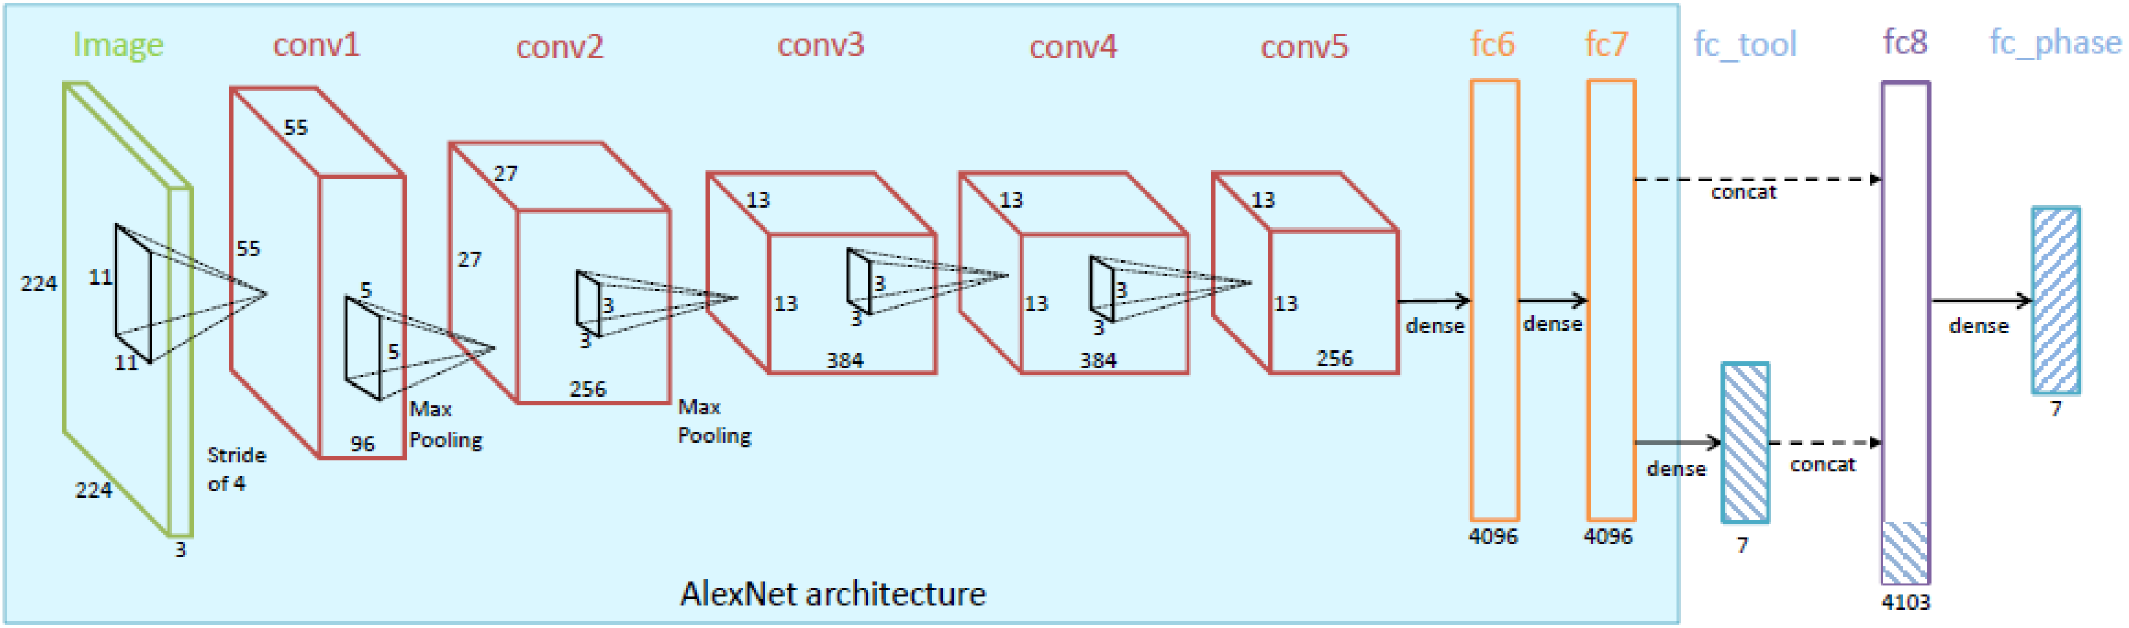

Fig. 2: EndoNet architecture (best seen in color). The layers shown in the turquoise rectangle are the same as in the AlexNet architecture.

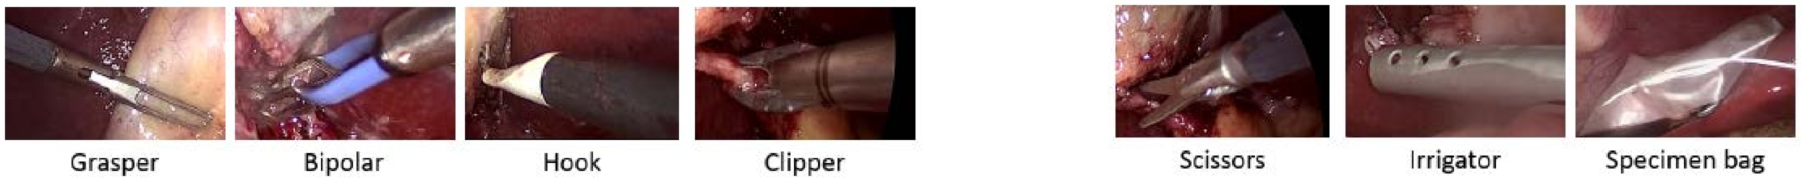

Other domains

- We could have presented similar impressive results for Natural Language Processing (translation, Name Entity Recognition,…), for speech Recognition,…

- These are not limited to signal but have been extended to graph data (among which social networks: pinterest, facebook…)

SO IT IS MAGIC??

univ-cotedazur.fr

Adversarial examples

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Amazing but…beware of the adversarial examples (as any other ML algorithms)

JDEV

Intriguing properties of neural networks

C. Szegedy, w. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, R. Fergus

arXiv preprint arXiv:1312.6199

2013

[1312.6199] Intriguing properties of neural networks - arXiv.org
https://arxiv.org › cs - Traduire cette page
de C Szegedy - 2013 - Cité 449 fois - Autres articles
21 déc. 2013 - In this paper we report two such **properties**. First, we ... Second, we find that deep **neural networks** learn input-output mappings that are fairly ...

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Morphing

JDEV

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Beware of the adversaries (as any other ML algorithms)

JDEV

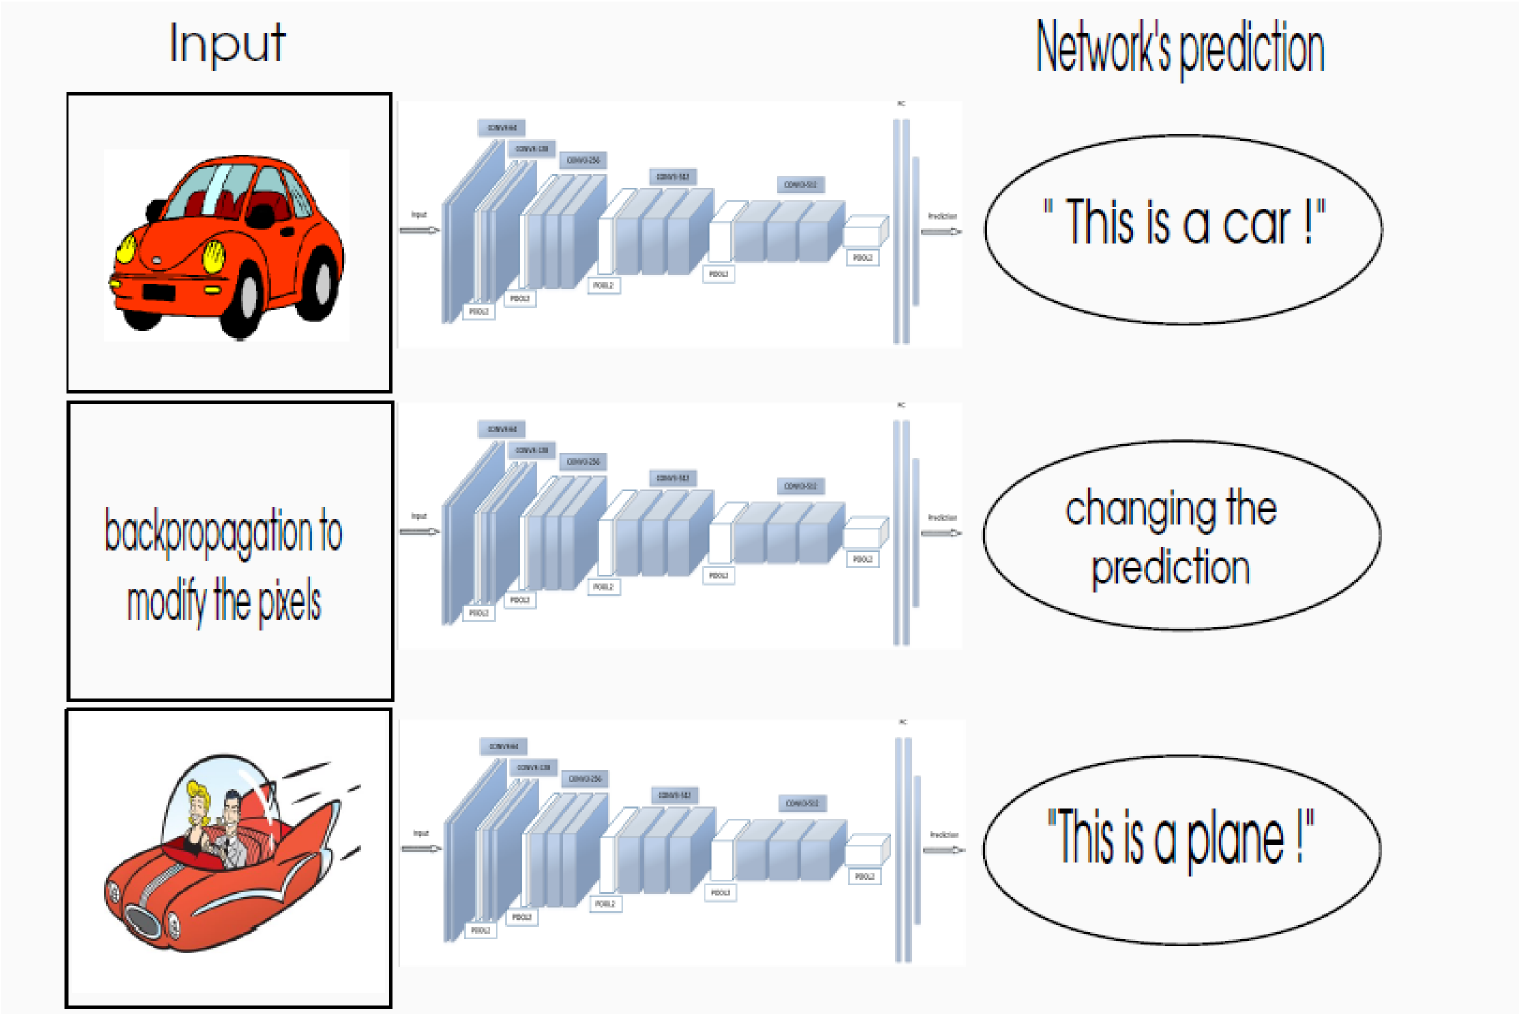

Beware of the adversaries (as any other ML algorithms)

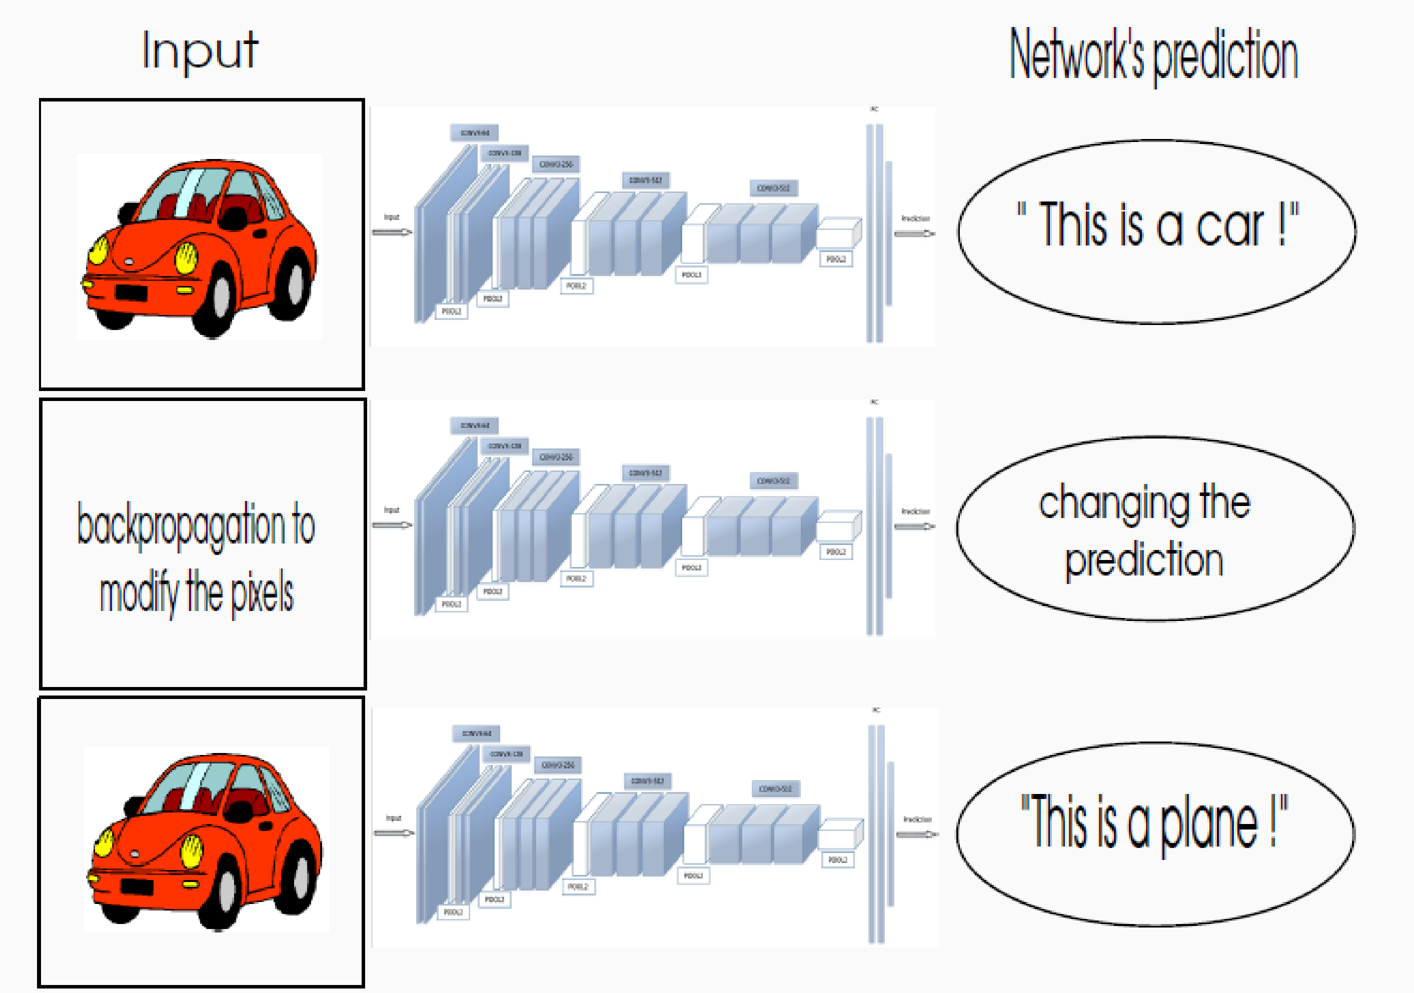

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Adversarial Examples

JDEV

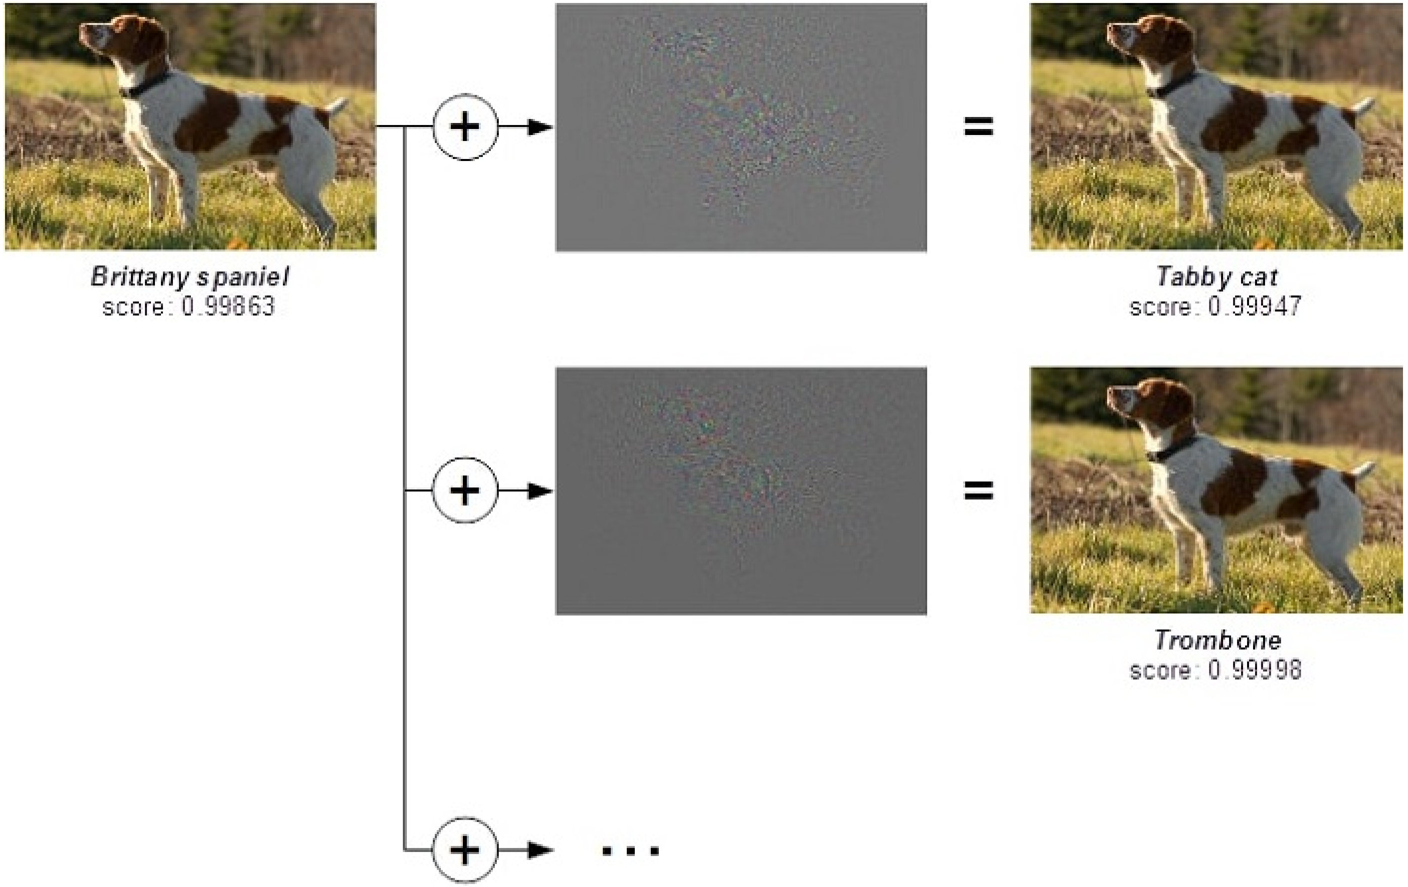

From Thomas Tanay

Beware of the adversaries (as any other ML algorithms)

- $\neq$ outliers
- regularization: correct one... find another
- high confidence predictions
- **Transferability**

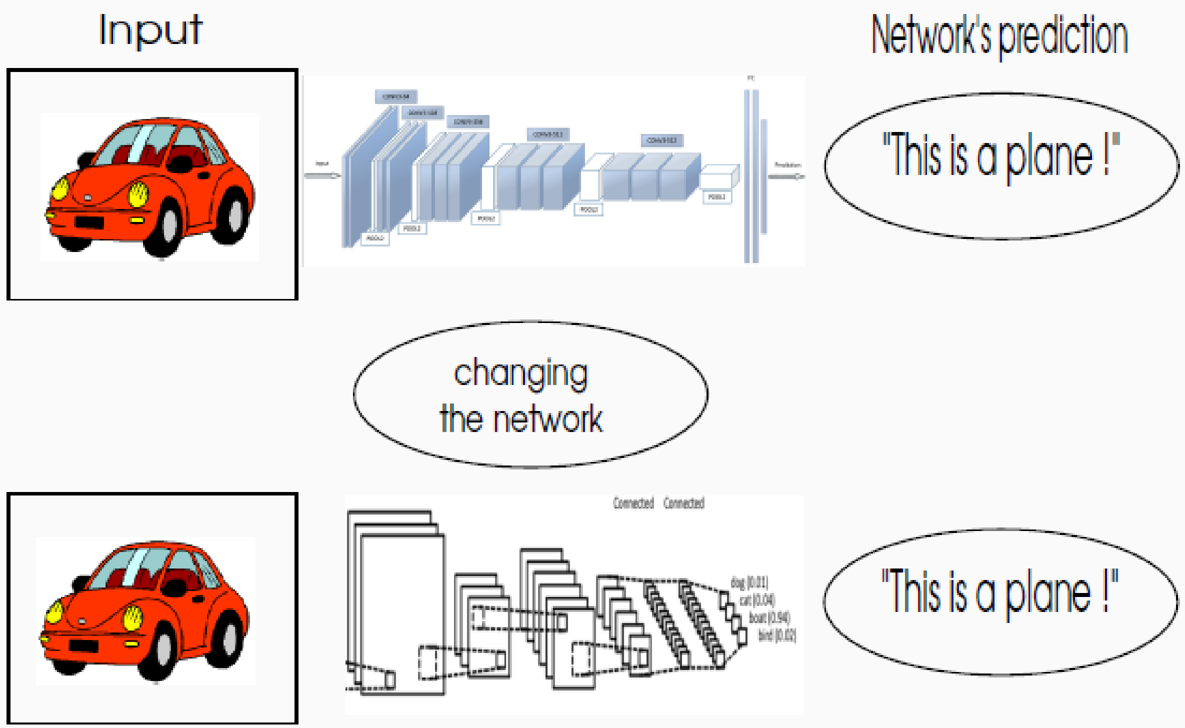

Beware of the adversaries (as any other ML algorithms)

Andrej Karpathy blog, http://karpathy.github.io/2015/03/30/breaking-convnets/

Beware of the adversaries (as any other ML algorithms)

univ-cotedazur.fr

- It “works” for other modalities also:

https://nicholas.carlini.com/code/audio_adversarial_examples/

BIASES

Beware of input bias

Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings

Tolga Bolukbasi[1], Kai-Wei Chang[2], James Zou[2], Venkatesh Saligrama[1,2], Adam Kalai[2]
[1]Boston University, 8 Saint Mary's Street, Boston, MA
[2]Microsoft Research New England, 1 Memorial Drive, Cambridge, MA
tolgab@bu.edu, kw@kwchang.net, jamesyzou@gmail.com, srv@bu.edu, adam.kalai@microsoft.com

From Nello Cristianini, at at *Frontier Research and Artificial Intelligence Conference*:
https://erc.europa.eu/sites/default/files/events/docs/Nello_Cristianini-ThinkBIG-Patterns-in-Big-Data.pdf

univ-cotedazur.fr

Beware of input bias

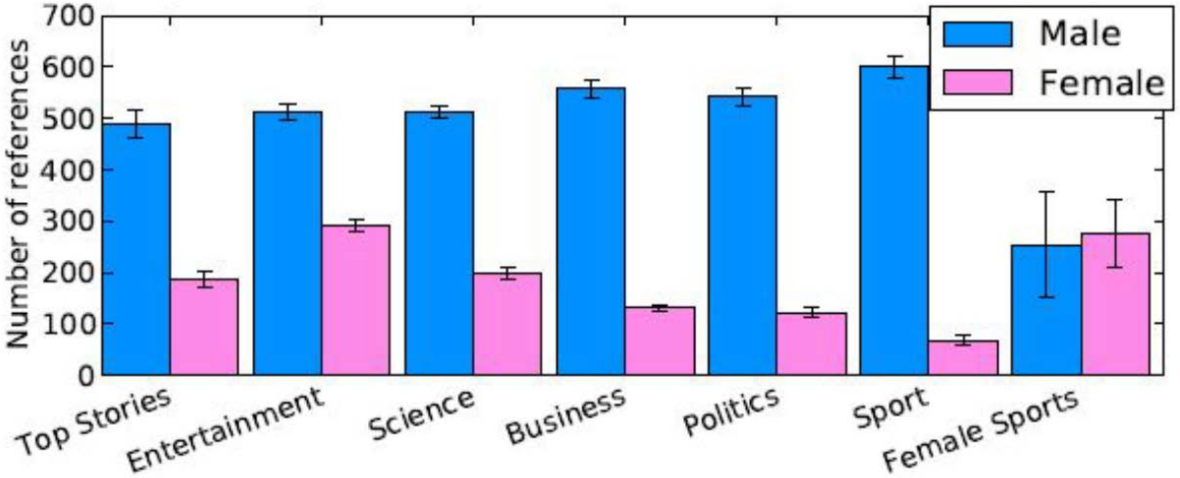

TABLE I: List of the top 10 occupations per gender by their association with gender.

Gender	Occupations most associated with a gender
Male	Manager, Engineer, Coach, Executive, Surveyor, Secretary, Architect, Driver, Police, Caretaker, Director
Female	Housekeeper, Nurse, Therapist, Bartender, Psychologist, Designer, Pharmacist, Supervisor, Radiographer, Underwriter

From Nello Cristianini, at at *Frontier Research and Artificial Intelligence Conference*:
https://erc.europa.eu/sites/default/files/events/docs/Nello_Cristianini-ThinkBIG-Patterns-in-Big-Data.pdf

Beware of input bias

BUSINESS NEWS OCTOBER 10, 2018 / 5:12 AM / 7 MONTHS AGO

Amazon scraps secret AI recruiting tool that showed bias against women

Jeffrey Dastin 8 MIN READ

SAN FRANCISCO (Reuters) - Amazon.com Inc's (AMZN.O) machine-learning specialists uncovered a big problem: their new recruiting engine did not like women.

Forget Killer Robots—Bias Is the Real AI Danger

John Giannandrea.
GETTY

Artificial Intelligence / Robots

Forget Killer Robots—Bias Is the Real AI Danger

John Giannandrea, who leads AI at Google, is worried about intelligent systems learning human prejudices.

by **Will Knight** Oct 3, 2017

Google's AI chief isn't fretting about super-intelligent killer robots. Instead, John Giannandrea is concerned about the danger that may be lurking inside the machine-learning algorithms used to make millions of decisions every minute.

"The real safety question, if you want to call it that, is that if we give these systems biased data, they will be biased," Giannandrea said before a recent Google conference on the relationship between humans and AI systems.

The problem of bias in machine learning is likely to become more significant as the technology spreads to critical areas like medicine and law, and as more people without a deep technical understanding are tasked with deploying it.

Energy Cost…

univ-cotedazur.fr

Winter is coming…

Consumption	CO_2e (lbs)
Air travel, 1 passenger, NY↔SF	1984
Human life, avg, 1 year	11,023
American life, avg, 1 year	36,156
Car, avg incl. fuel, 1 lifetime	126,000
Training one model (GPU)	
NLP pipeline (parsing, SRL)	39
w/ tuning & experimentation	78,468
Transformer (big)	192
w/ neural architecture search	626,155

Table 1: Estimated CO_2 emissions from training common NLP models, compared to familiar consumption.[1]

"Energy and Policy Considerations for Deep Learning in NLP"

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

A POSSIBLE SOLUTION: BRIDGING SYMBOLIC AND SUBSYMBOLIC

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

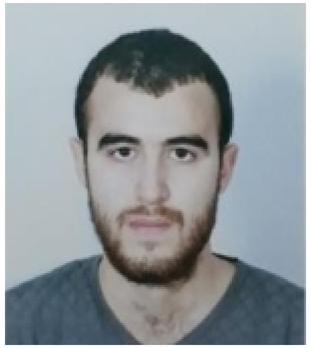

Taki Eddine MEKHALFA

Prof. Marco Gori

Learning knowledge with neural networks

(work from Prof. Marco Gori and his team, to be continued with Taki-Eddine Mekhalfa)

Bridging symbolic and subsymbolic

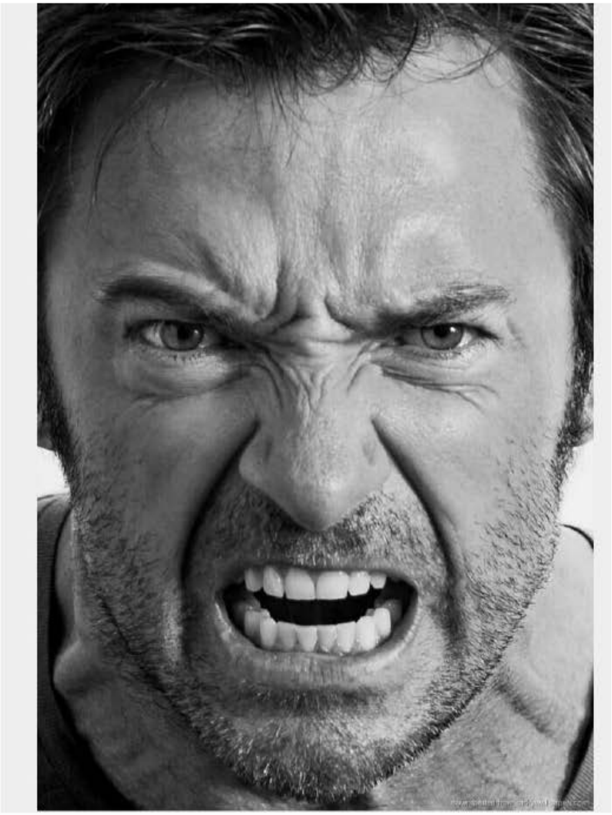

$$17 \times 23$$

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

Bridging symbolic and subsymbolic

In both cases, we carried on a **reasoning process** :

- In the "angry face" case, it was a **fast** reasoning process, mostly **associative**, **approximate** and **effortless**.
  - Subsymblic AI mostly concentrates on these tasks.
- In the "multiplication" case, it was a **slow** reasoning process, requiring multiple steps and the need to temporary **store intermediate results** and it was **effortful**.
  - Symbolic AI mostly concentrates on these tasks.

Bridging symbolic and subsymbolic

Many framework already exists :

- Symbolic approaches enhanced by deep learning (e.g. DeepProbLog, NTP, LRNN)
- Subsymbolic approaches enhanced with structure (e.g. CRF, DSL, SBR)

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

Bridging symbolic and subsymbolic

We consider a logic language $\mathcal{L}$. $\mathcal{C}$ be a set of constants and $\mathcal{R}$ a set of relations of any arity.
For $c_1, c_2, \ldots, c_k \in \mathcal{C}$ and $R \in \mathcal{R}$, we call $R(c_1, c_2, \ldots, c_k)$ a **ground atom** or **fact** or **triple**.
E.g.

$$\mathcal{C} = \{Alice, Bob, Eve\}$$
$$\mathcal{R} = \{smokes, friendOf\}$$

$$smokes(Alice), friendOf(Alice, Bob)$$

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Bridging symbolic and subsymbolic

JDEV

There exists a *feature representation* function g defined on a subset of $\mathcal{C}$, which provides a feature representation x of some (or all) the constants in $\mathcal{C}$.

E.g.

$$x_{\text{Alice}} = g(Alice) = [\underbrace{23}_{\text{age}}, \underbrace{165}_{\text{height}}, \underbrace{0.001, 0.22, \cdots, 0.32}_{\text{RGB profile photo}}, \cdots]$$

$$\mathbf{x} = \{g(c) : c \in \mathcal{C}\}$$

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

Bridging symbolic and subsymbolic

The **Herbrand Base** HB is the set of all possible ground atoms that can be built from $\mathcal{L}$.
E.g.

$$\mathcal{C} = \{Alice, Bob, Eve\}$$
$$\mathcal{R} = \{smokes, friendOf\}$$

$$\begin{aligned} HB(\mathcal{L}) = \{&smokes(Alice), smokes(Bob), smokes(Eve),\\ &friendOf(Alice, Alice), friendOf(Alice, Bob),\\ &\ldots\\ &friendOf(Eve, Eve)\} \end{aligned}$$

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

Bridging symbolic and subsymbolic

A **global example y** (also called **Herbrand Intepretation** or **possible world** or **labels**) is an assignment of a **truth value** to some or all elements of the HB. It is usually defined as a subset of HB composed of only `True` ground atoms
E.g.

$$\mathcal{C} = \{Alice, Bob, Eve\}$$
$$\mathcal{R} = \{smokes, friendOf\}$$

$$y = \{friendOf(Bob, Eve), friendOf(Eve, Bob), smokes(Alice)\}$$

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

Bridging symbolic and subsymbolic

Goal

We want to model the probability distribution $p(\mathbf{y}|\mathbf{x})$ *in order to :*

- *reason* ***under uncertainty***
- *about the* ***truth value***
- *of some* ***symbolic entities*** $\mathbf{y}$ *(ground atoms of the language* $\mathcal{L}$ *; boolean random variables)*
- *given some* ***perceptions*** $\mathbf{x}$ *of the constants* $\mathcal{C}$

We are usually provided with $\mathbf{x}$ and (some examples of) $\mathbf{y}$

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

NEURAL MARKOV LOGIC NETWORKS

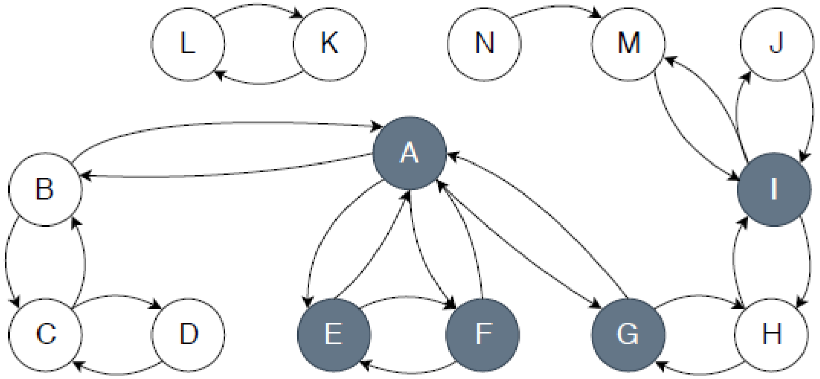

(a) The training KB.

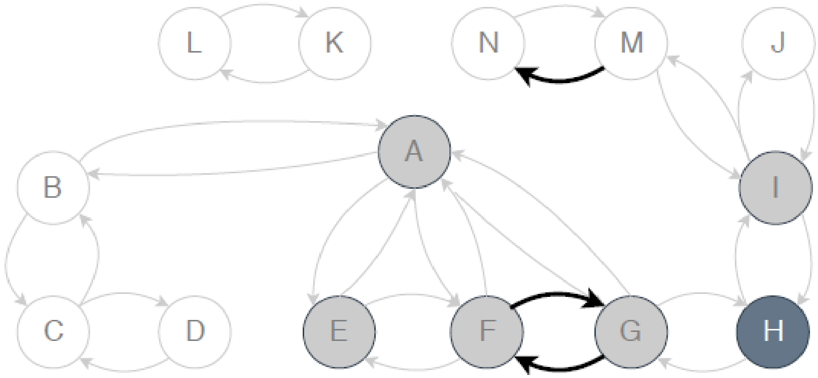

(b) The completed KB.

- **Knowledge Base Completion in the Nations dataset.** Circles represent constants.
- A grey circle means that the predicate smokes is *True*.
- A white circle means that the value of the predicate smokes is unknown.
- Links represent the relation friendOf (absence of an arrow means that the relation is *False*). The given world is shown on the top (a), while the completed knowledge base is shown on the bottom (b).
- The system learnt the symmetric nature of the friendship relation.
- It learnt that a friend of at least two smokers is also a smoker, and that two smokers, who are friends of the same person, are also friends.

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

Next Steps

- Learning the world “from scratch” by interacting with it, by perceiving it
- Learning jointly symbolic and subsymbolic models

Slide from Giuseppe Marra, Universities of Florence and Siena, PhD Defense

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV


Dir.Research Fabien Gandon


Anna Bobasheva

Deep Learning ↔ Knowledge Representation

(work in collaboration with Fabien Gandon and Anna Bobasheva)

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Some existing works

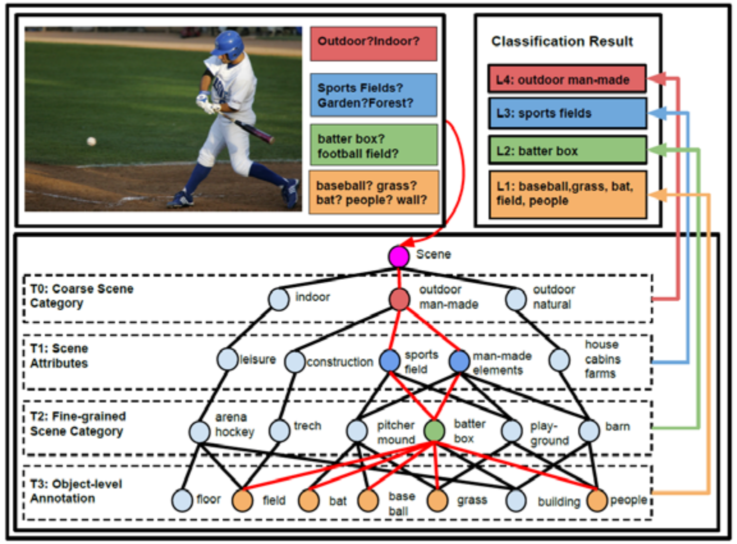

"Learning Structured Inference Neural Networks with Label Relations"

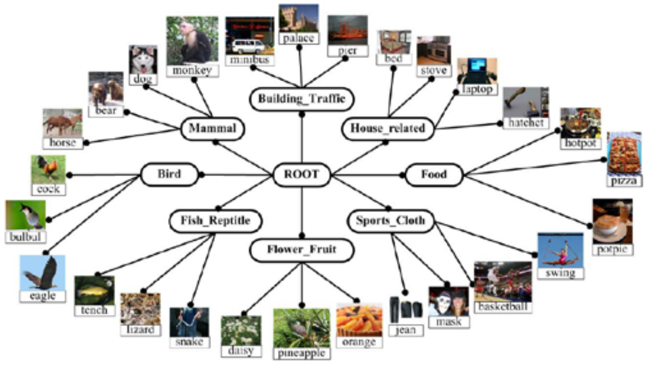

Fig. 2: The concept ontology for ILSVRC2012 image set, which is used to organize 1,000 object classes hierarchically according to their inter-class semantic correlations.

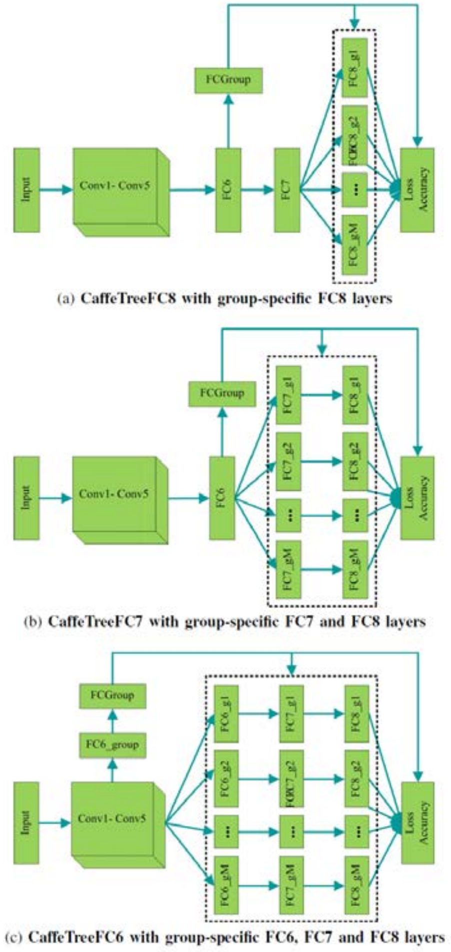

"Deep Multi-Task Learning for Large-Scale Image Classification"

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Study Overview

JDEV

- reason & query on RDF metadata to build balanced, unambiguous, labelled training sets.
- transfer learning & CNN classifiers on targeted categories (topics, techniques, etc.)
- **reason & query RDF metadata of results to address silence, noise, and explain**

Joconde database

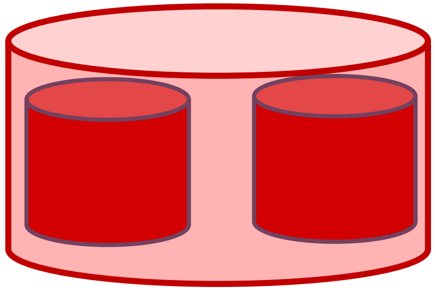

350 000 ***images*** *of artworks*

RDF ***metadata*** *based on* *external thesauri*

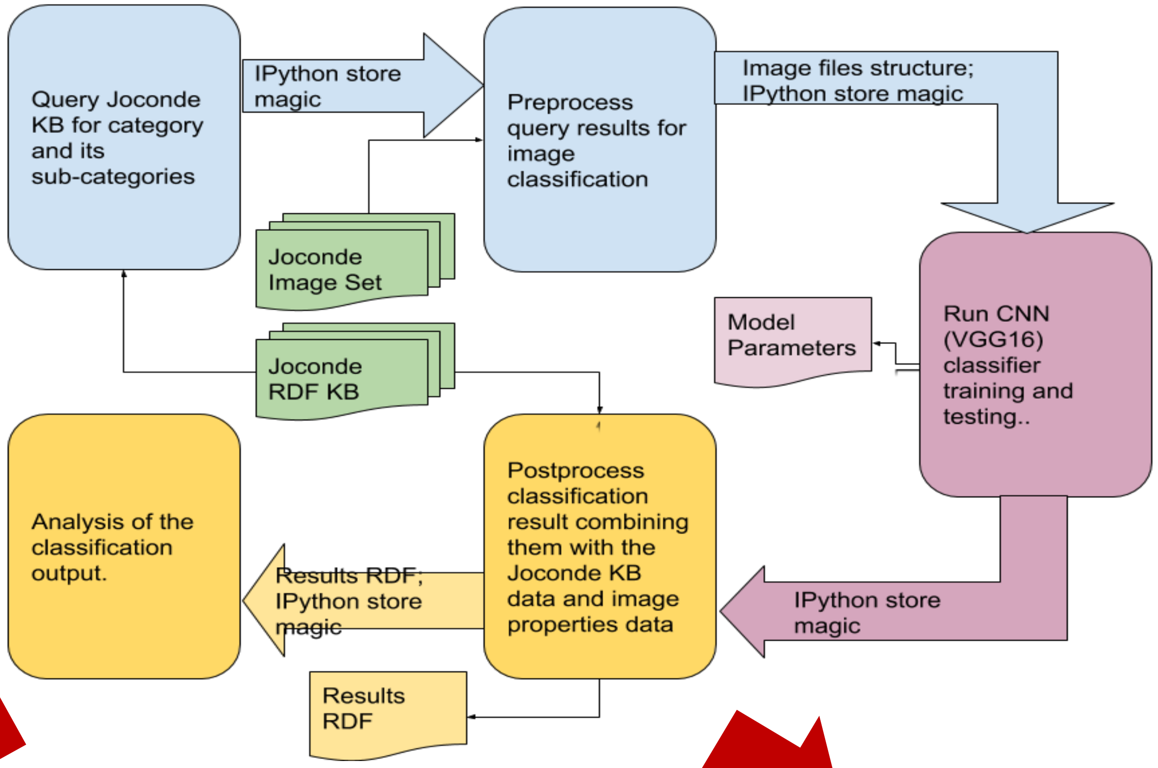

☑ ***animal***

? ***bird***

☒ ***painting***

Explain

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

Motivation & Challenges

GALERIE DE VUES DE LA ROME MODERNE by PANNINI Giovanni Paolo

Museum curators have to annotate thousands of artworks acquired over the hundreds of years and now managed as digital collections. This process can be tedious and susceptible to the human errors and omissions.

- **Can the existing digital artwork collections be automatically enhanced by combining Machine Learning and Knowledge Representation & Reasoning?**

- **Can annotation of the new artworks be automated or semi-automated?**

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Joconde Database

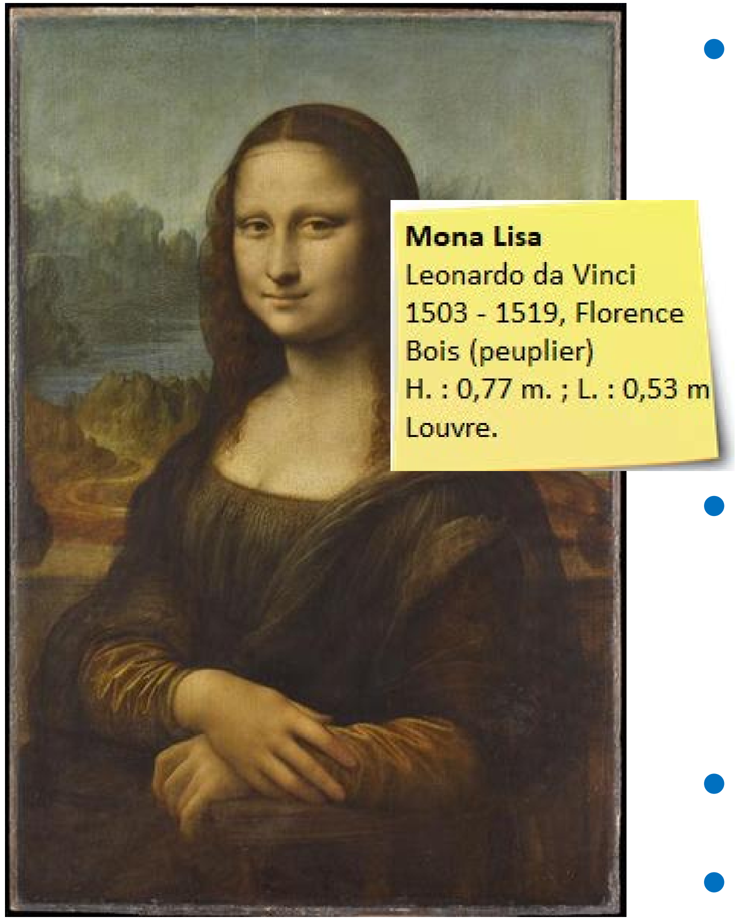

PORTRAIT DE MONA LISA (1479-1528) ; DITE LA JOCONDE by Leonardo Da Vinci

- 350 000 illustrated artwork records from the French museums.

  RDF metadata describing the artwork subject and properties (media, author, museum, etc.).

- **The database is searchable on the artwork subjects and other properties but…**

- **The metadata can be <u>incomplete</u> & <u>noisy</u>.**
- **The new artworks <u>added continuously</u>.**

Enabling Methods

Deep Learning from unstructured data such as images

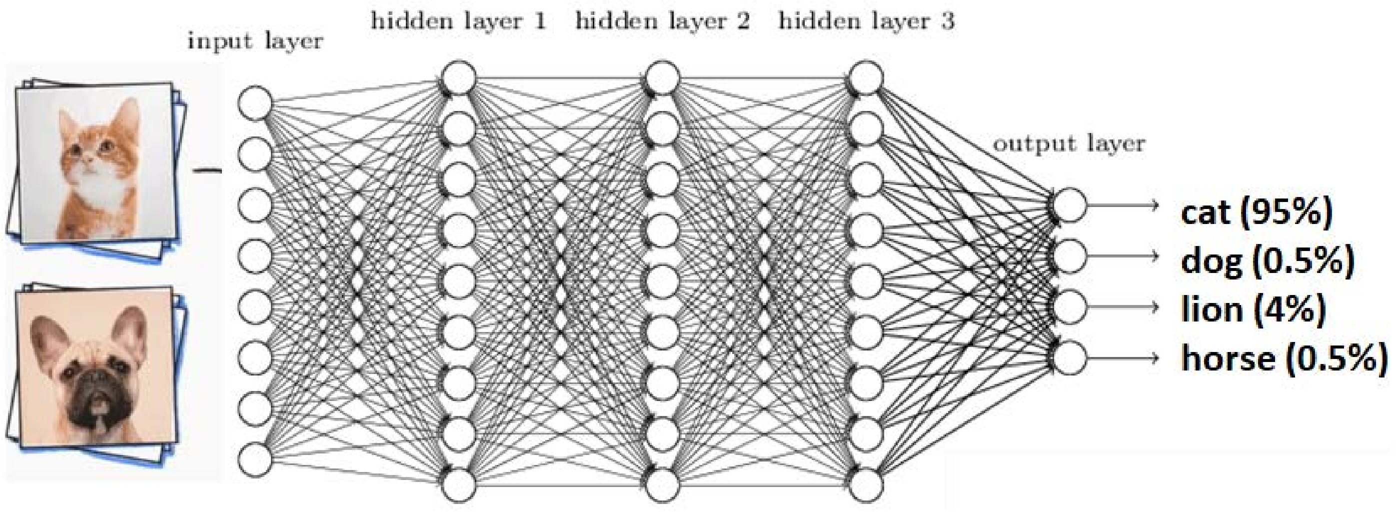

Semantic Reasoning and querying from semantic metadata

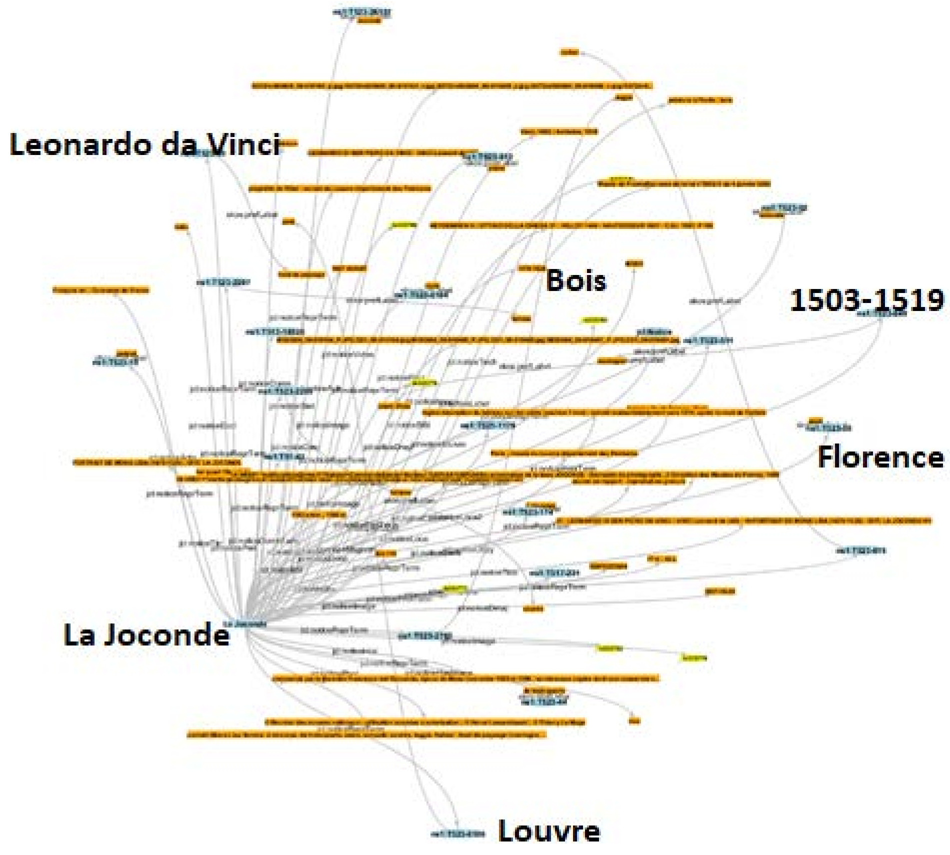

Marrying Methods: Combining Their Strength

Le Mariage de la Vierge by Alonso Cano

- reason to **prepare and control training** sets & labels
- learn to **improve the quality of the existing metadata** that is incomplete & noisy
- learn to **annotate new artworks** with efficiency
- reason to **augment and explain results**
- learn and reason to **improve searchability** of the Joconde database

Marrying Methods: Combining Their Strength

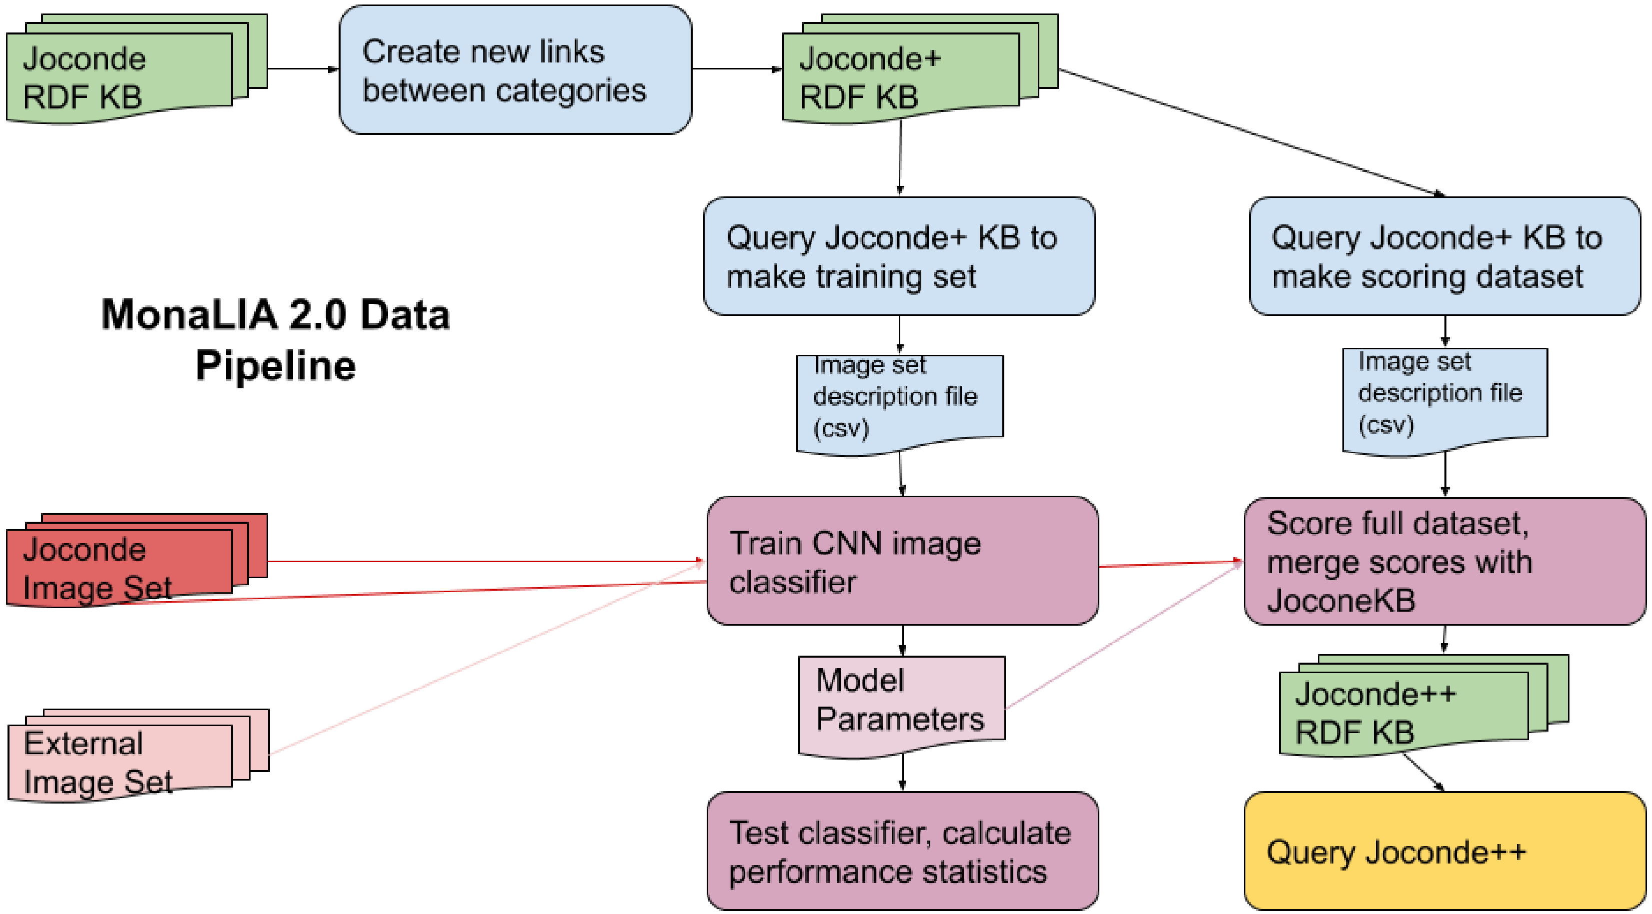

Marrying Methods: Combining Their Strength

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

- On the road to Mona Lia 3.0
  - Deep Network Layers and Thesaurus Layers (representation level)
  - learning and reasoning techniques (inference level)
  - Induction on RDF data and unstructured data

Liberté • Égalité • Fraternité
RÉPUBLIQUE FRANÇAISE

We want to thank The French Ministry of Culture for the opportunity to work on such an exiting project and for funding it.

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV


Bora Kizil, Ezako


Prof. Mireille Blay-Fornarino


Julien Muller, Ezako


Yassine El Amraoui

"Smart" MLPaaS / "Smart" MLOPS

(work in collaboration with Mireille Blay-Fornarino, Yassine El Amraoui, Julien Muller, and Bora Kizil)

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Existing Works

- AutoML
- Auto-Weka
- …

- Azure Microsoft
- Amazon AWS
- Google Cloud
- …

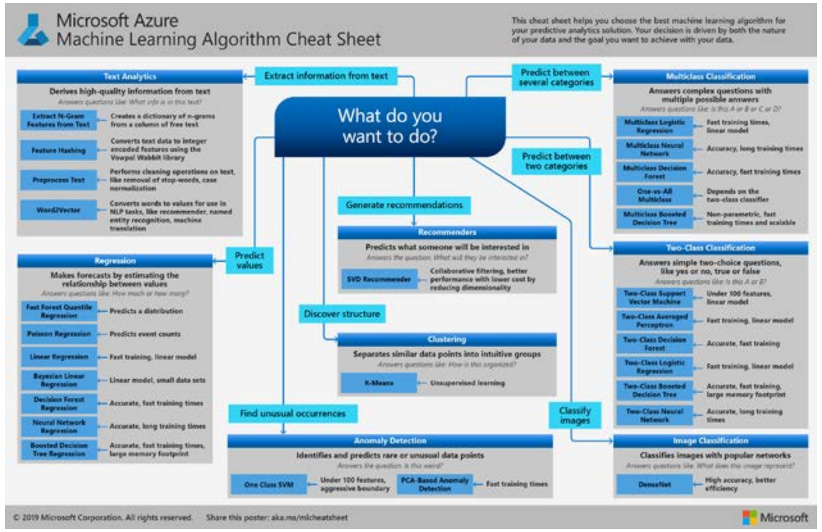

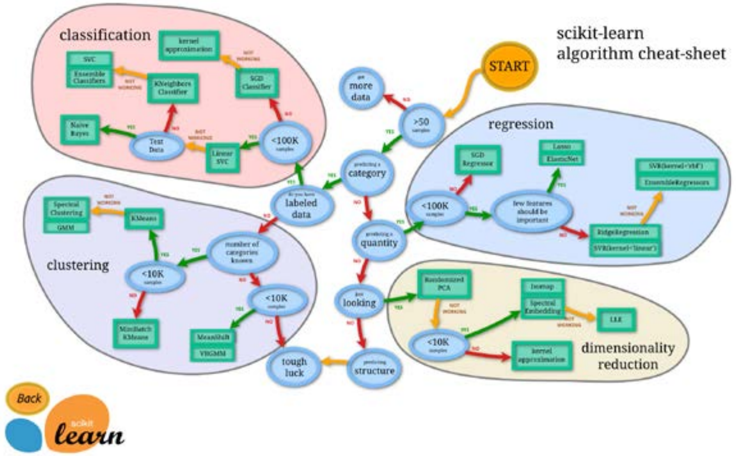

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

Rockflows

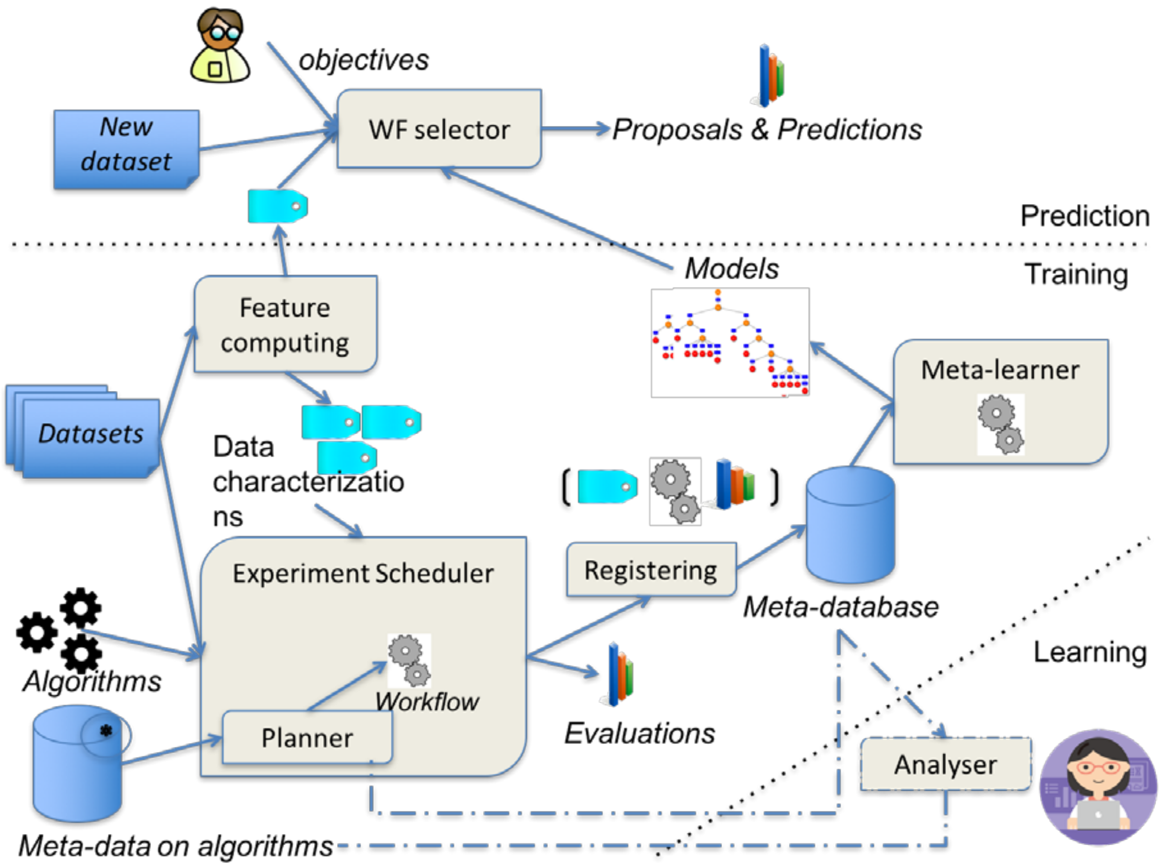

http://rockflows.i3s.unice.fr/

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

Rockflows

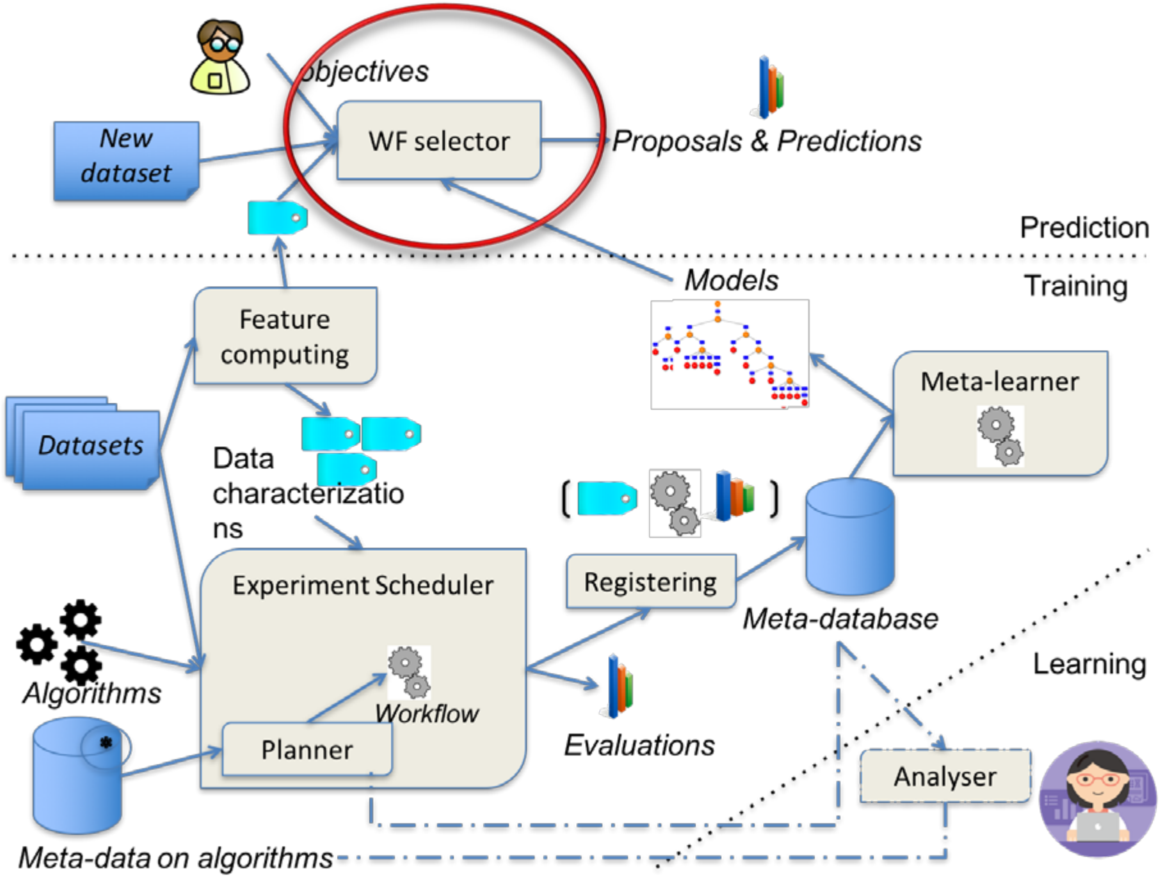

http://rockflows.i3s.unice.fr/

Rockflows

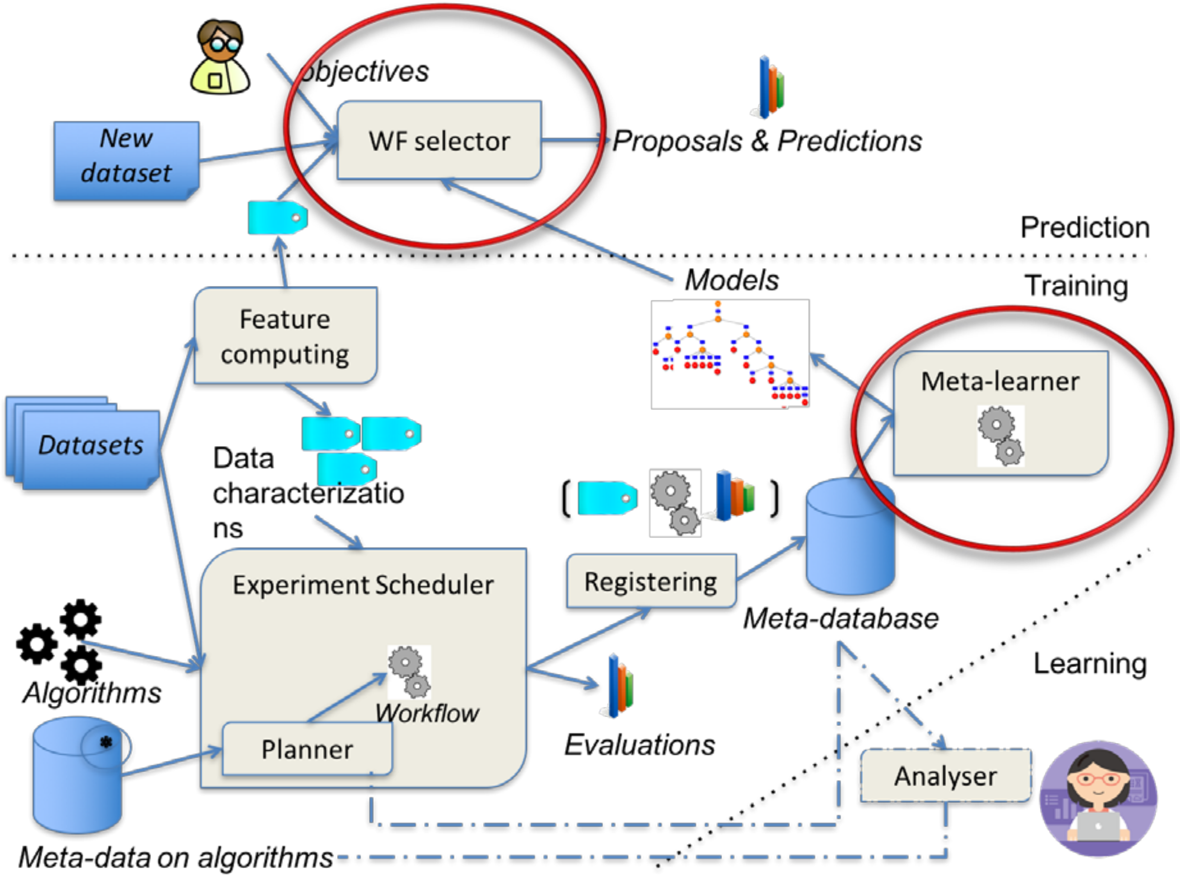

http://rockflows.i3s.unice.fr/

UNIVERSITÉ CÔTE D'AZUR

univ-cotedazur.fr

JDEV

Rockflows

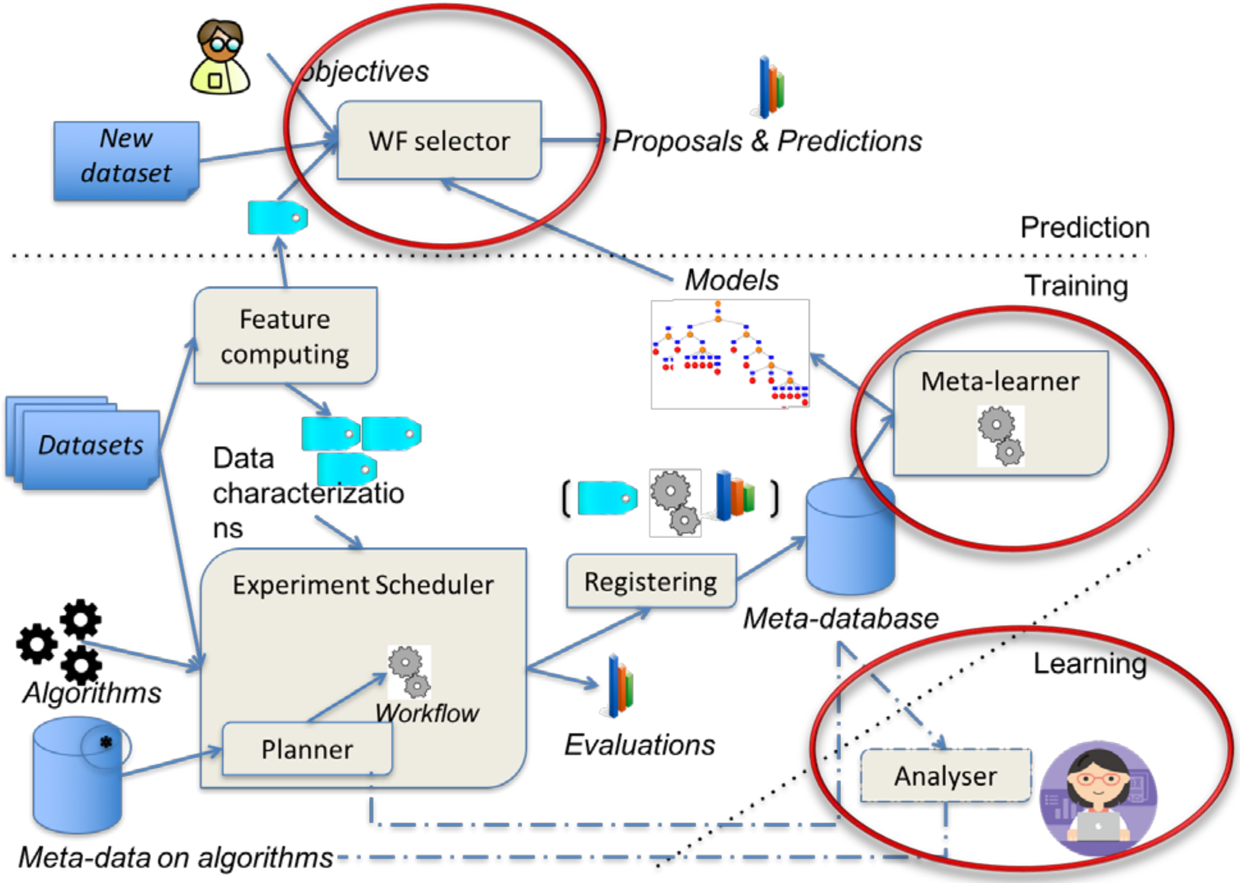

http://rockflows.i3s.unice.fr/

Rockflows

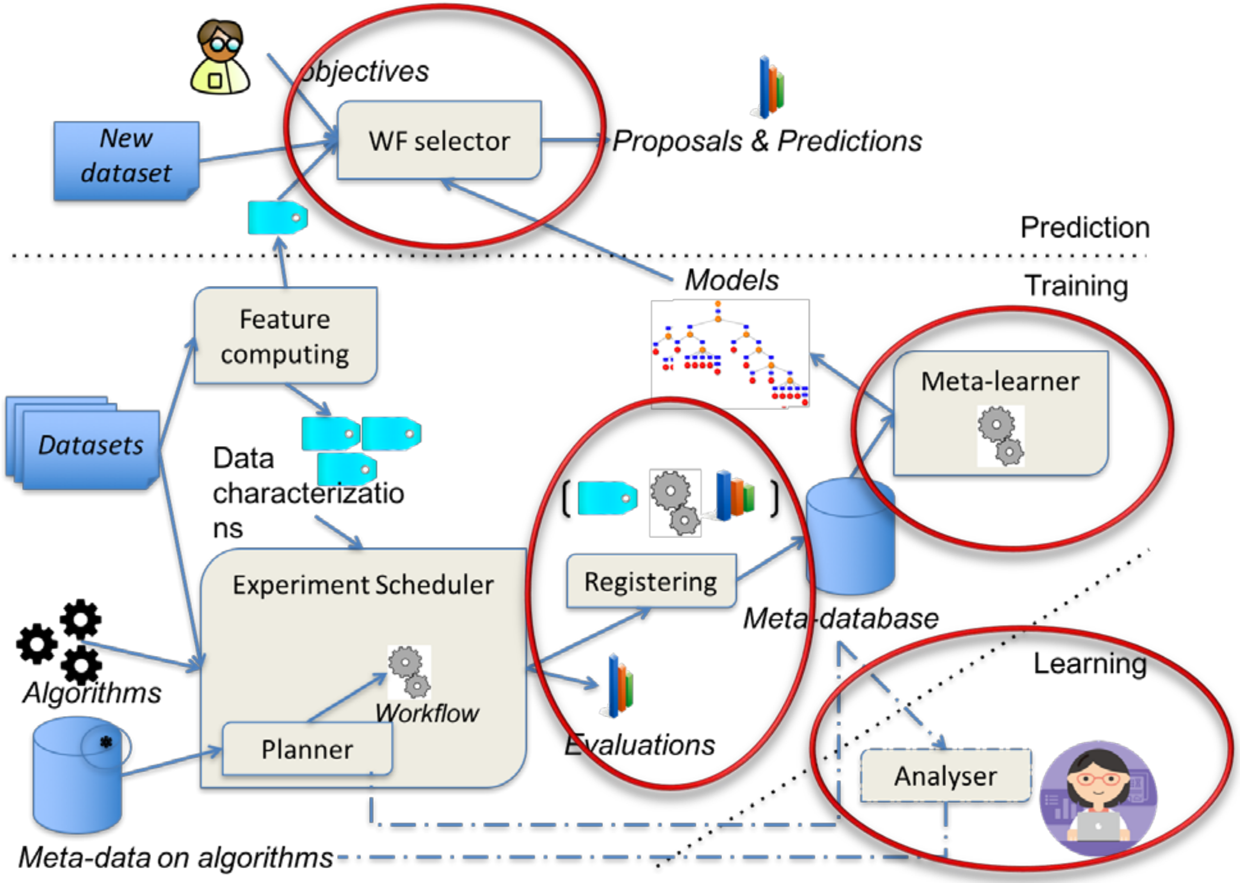

http://rockflows.i3s.unice.fr/

Rockflows

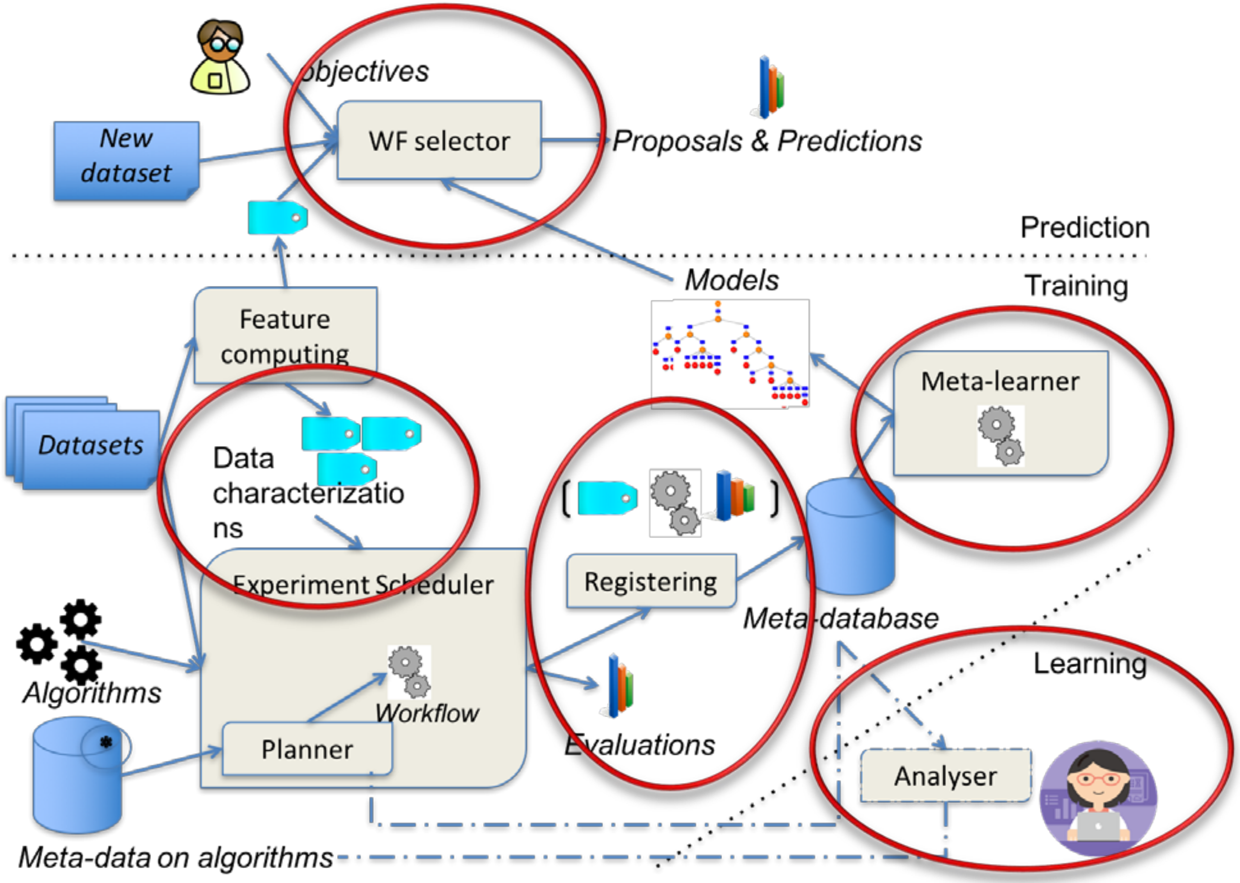

http://rockflows.i3s.unice.fr/

UNIVERSITÉ CÔTE D'AZUR
univ-cotedazur.fr

JDEV

Rockflows

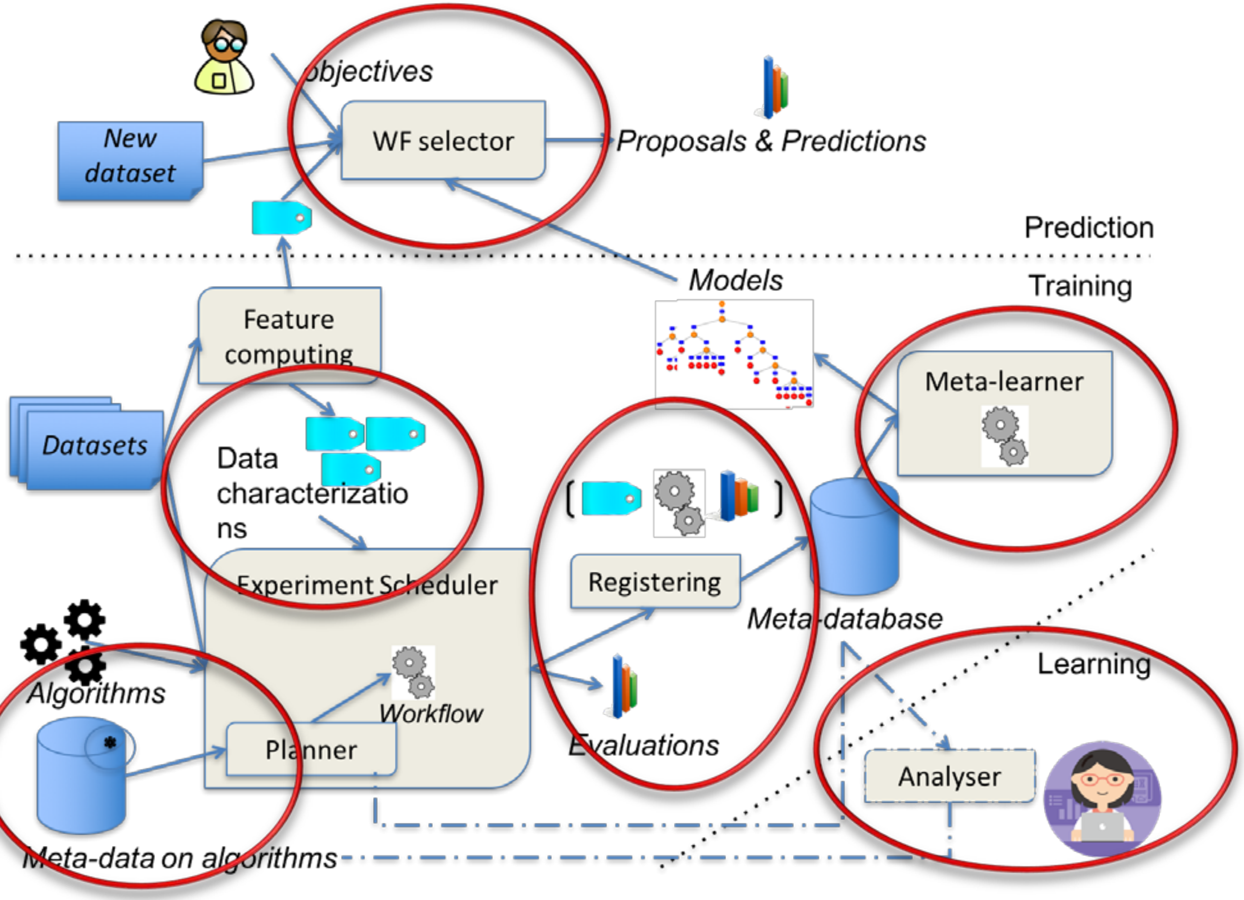

http://rockflows.i3s.unice.fr/