Inductive Learning:

Antoine Cornuéjols

AgroParisTech – INRA MIA 518

EKINOCS research group

Inductive Learning:

a Risky Business

Antoine Cornuéjols

AgroParisTech – INRA MIA 518

EKINOCS research group

Inductive learning: what **it does**

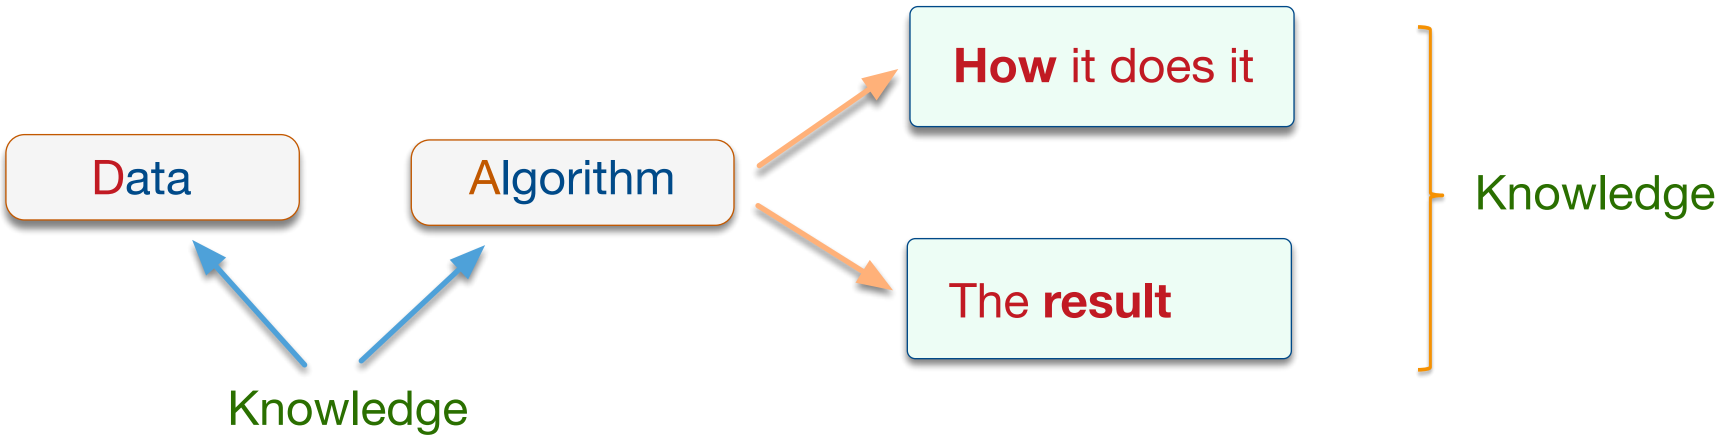

...

Inductive learning: what **it does**

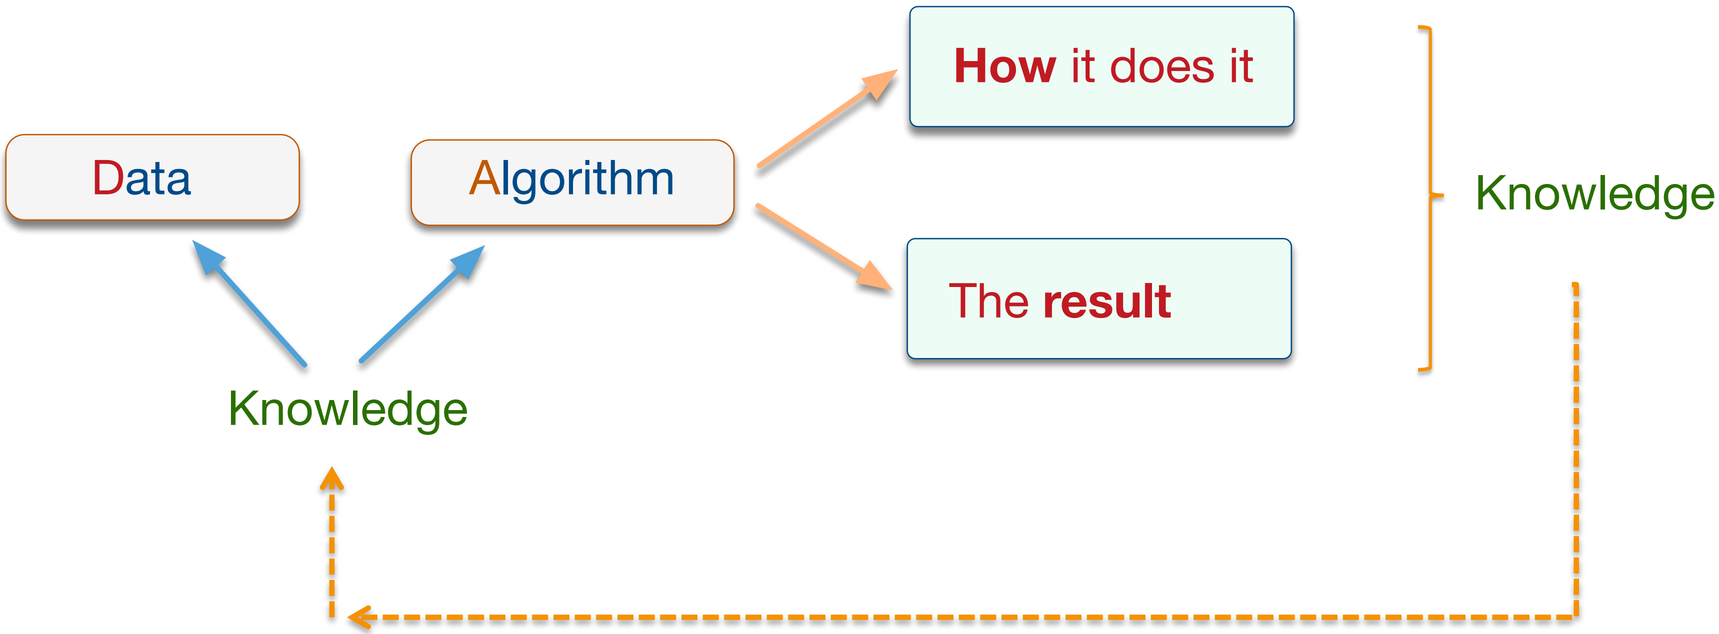

...

Supervised Induction

Descriptors

Labels

	#Cycles	Masse	PH	Carboxyle	Activité
M1	1	faible	<5	non	nulle
M2	2	moyen	<5	oui	toxique
M3	0	moyen	>8	oui	toxique
M4	0	moyen	<5	non	nulle
M5	1	lourd	~7	non	nulle
M6	2	lourd	>8	non	toxique
M7	1	lourd	>8	non	toxique
M8	0	faible	<5	oui	toxique

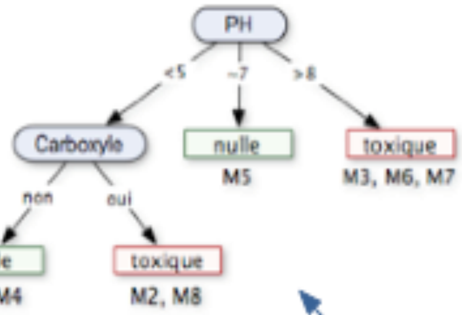

R1: Si pH > 8
alors toxique
R2 : Si pH < 5 & Carboxyle = oui
alors toxique
....

Example

Values of the descriptors

Concept
or
hypothesis

What do you expect?

- Better **understand** your data

What do you expect?

- Better **understand** your data

- Be able to make **prediction**

What do you expect?

- Better **understand** your data

- Be able to make **prediction**

- Be able to make **prescription**

What do you expect?

- Better **understand** your data

- Be able to make **prediction**

- Be able to make **prescription**

(1) **Understand** your data

- **Re-express** it
  - In a concise way
  - To be interpretable by an expert of the domain

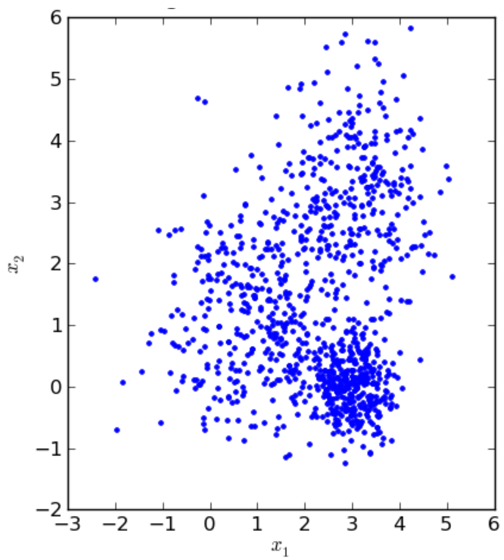

Original data

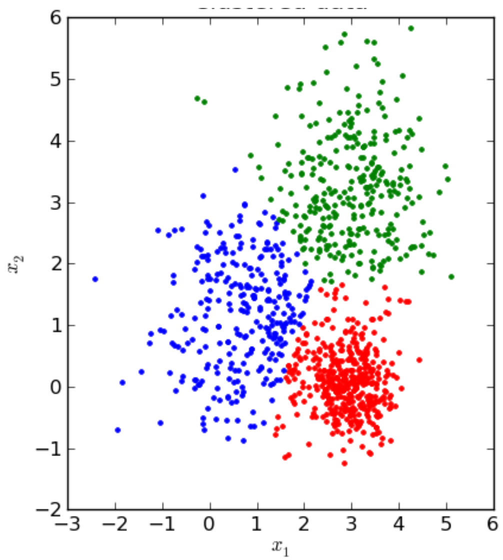

Clustered data

Three groups of customers with such and such **characteristics** ...

(2) Make **predictions**

- Extrapolate your data to find **predictive correlations**
  - From a **training** set $S = \{(x_1, y_1), (x_2, y_2), \ldots, (x_i, y_i), \ldots, (x_m, y_m)\}$

(2) Make **predictions**

- Extrapolate your data to find predictive correlations
  - From a **training** set $S = \{(x_1, y_1), (x_2, y_2), \ldots, (x_i, y_i), \ldots, (x_m, y_m)\}$

$$S = \{(x_1, y_1), (x_2, y_2), \ldots, (x_i \xrightarrow{f} y_i), \ldots, (x_m, y_m)\}$$

(2) Make **predictions**

- Extrapolate your data to find predictive correlations
  - From a **training** set $S = \{(x_1, y_1), (x_2, y_2), \dots , (x_i , y_i), \dots , (x_m, y_m)\}$

$$S = \{(x_1, y_1), (x_2, y_2), \dots , (\overset{\overset{f}{\curvearrowright}}{\underset{\underset{h}{\curvearrowright}}{x_i , y_i}}), \dots , (x_m, y_m)\}$$

(2) Make **predictions**

- Extrapolate your data to find predictive correlations
  - From a **training** set $S = \{(x_1, y_1), (x_2, y_2), \ldots , (x_i , y_i), \ldots , (x_m, y_m)\}$

$$S = \{(x_1, y_1), (x_2, y_2), \ldots , (x_i \overset{f}{\underset{h}{\rightrightarrows}} y_i), \ldots , (x_m, y_m)\}$$

New x --> y ?

(2) Make **predictions**

- Extrapolate your data to find predictive correlations
  - From a **training** set $S = \{(x_1, y_1), (x_2, y_2), \dots, (x_i, y_i), \dots, (x_m, y_m)\}$

$$S = \{(x_1, y_1), (x_2, y_2), \dots, (x_i, y_i), \dots, (x_m, y_m)\}$$

f

h

$$x - h \to y$$

(2) Machine Learning as …

… Learning a function from an **input** space X to an **output** space Y

Cats vs. **dogs**

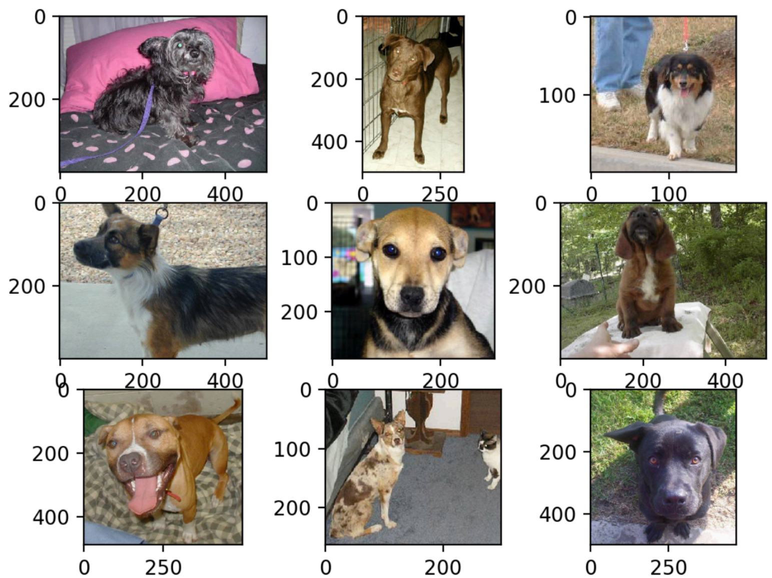

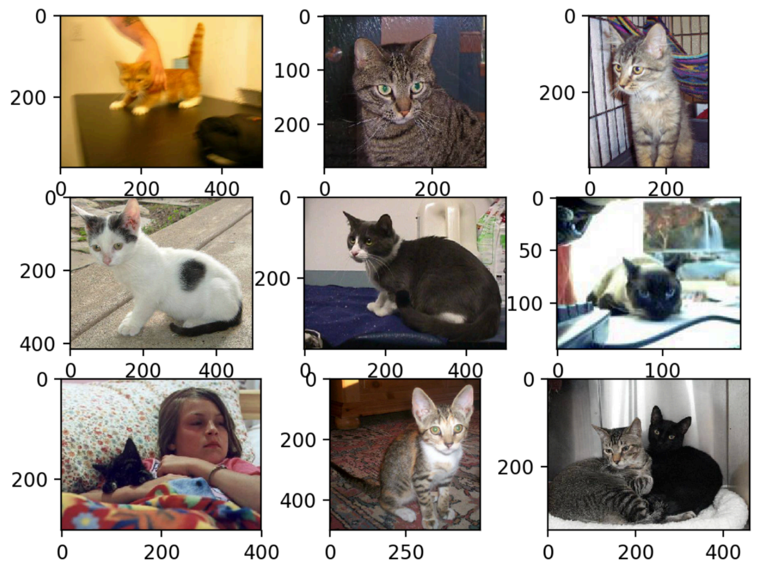

Supervised learning

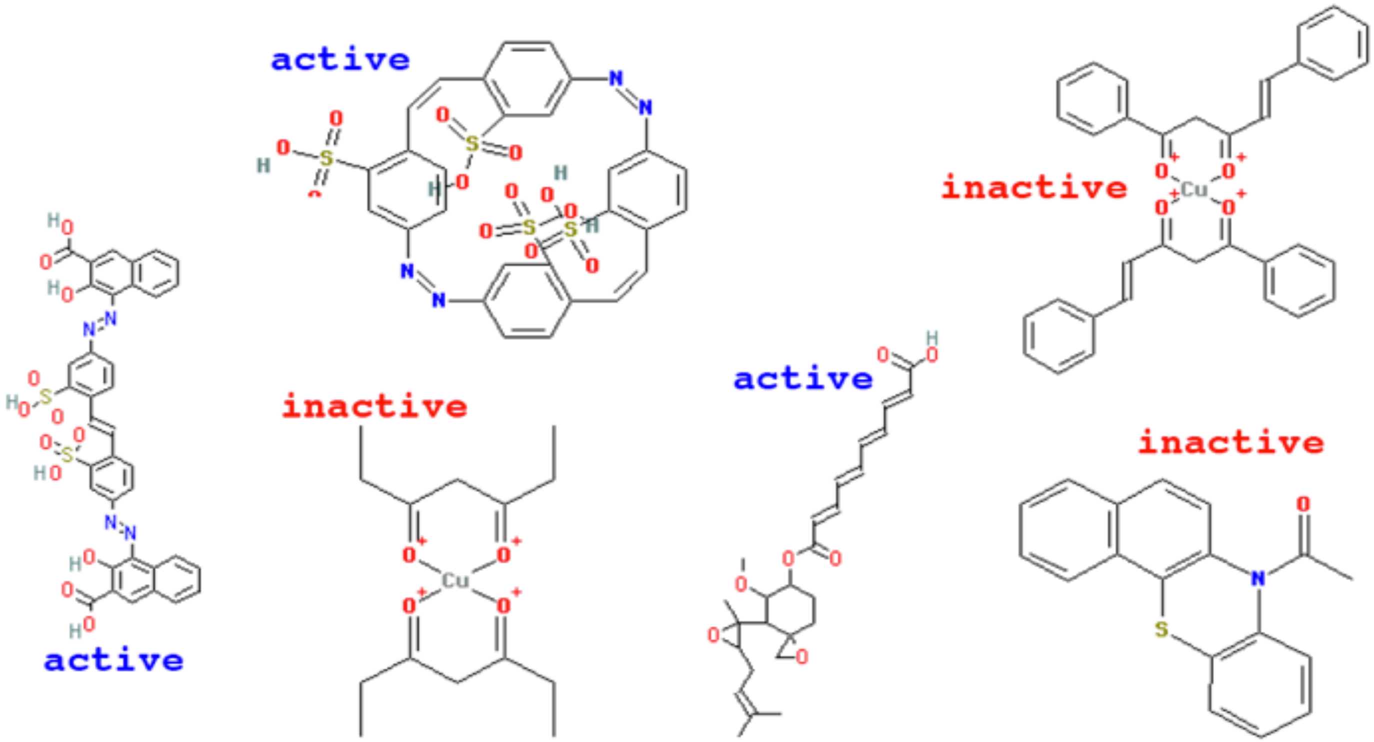

NCI AIDS screen results (from http://cactus.nci.nih.gov).

(3) Make **prescriptions**

- Learning causal relationships
  - The barometer **allows the prediction** of tomorrow's weather
  - But tampering with its needle **will not** change the weather

Correlation **is not** causality

Discovering causal relationships

(generally) requires knowing **more than the data**

Looking for causal relationships

- What **causes appetite** for protein dishes?

  - **Hunger** ?
  - The **hour** in the day?
  - **Gender**?
  - The **visual** aspect?
  - The **olfactory** aspect?
  - The high protein content of **previous meals**?
  - …

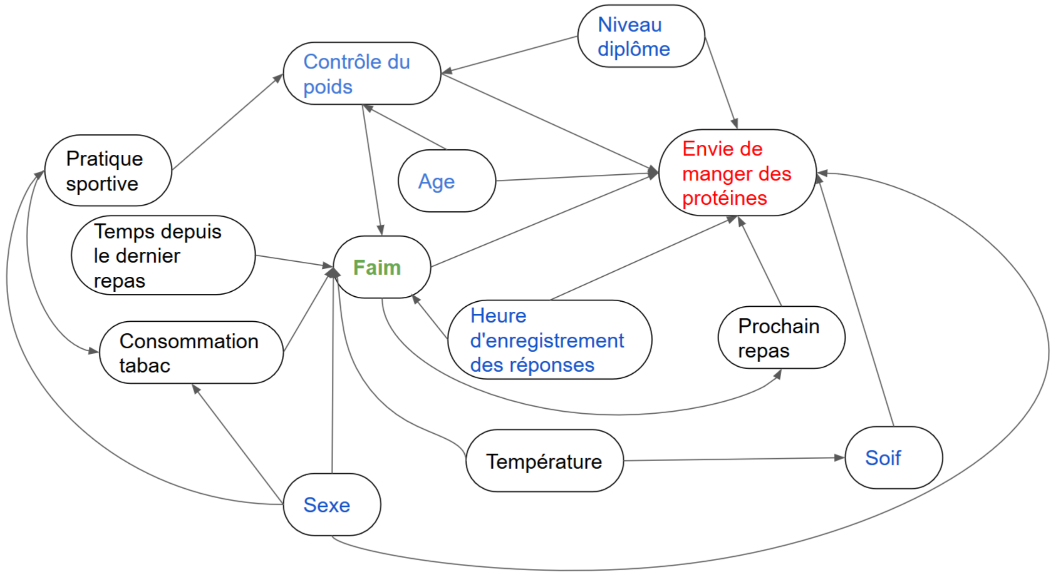

How do you **evaluate** the **results?**

How do you **evaluate** the **results?**

- **Descriptive** learning

How do you **evaluate** the **results?**

- **Descriptive** learning
  - **Validation** by the **expert**

How do you **evaluate** the **results?**

- **Descriptive** learning
  - **Validation** by the **expert**
  - The expert can be wrong
    - **Wants to see things** that are not really there
    - **Blind** to interesting but out of the blue patterns

How do you **evaluate** the **results?**

- **Descriptive** learning
  - **Validation** by the **expert**
  - The expert can be wrong
    - **Wants to see things** that are not really there
    - **Blind** to interesting but out of the blue patterns

Descriptive learning usually takes place in an **exploratory phase**

→ Be **very careful**

How do you **evaluate** the **results?**

- **Predictive** learning

How do you **evaluate** the **results?**

- **Predictive** learning
  - **Predictive** performance (on a test set)
    - E.g. error rate

How do you **evaluate** the **results?**

- **Predictive** learning
  - **Predictive** performance (on a test set)
    - E.g. error rate
  - But, this is not all there is to it
    - **Interpretability** of the **model**
    - Explanation/**justification** of the **prediction**
    - **Fruitfulness** wrt. the **domain theory**

How do you **evaluate** the **results?**

- **Predictive** learning
  - **Predictive** performance (on a test set)
    - E.g. error rate
  - But, this is not all there is to it
    - **Interpretability** of the **model**
    - Explanation/**justification** of the **prediction**
    - **Fruitfulness** wrt. the **domain theory**

Often, we are **not** interested in prediction alone,

but in **understanding** the **prediction** and/or the **predictive model**

How do you **evaluate** the **results?**

- **Predictive** learning
  - **Predictive** performance (on a test set)
    - E.g. error rate
  - But, this is not all there is to it. We want also
    - **Interpretability** of the **results**
    - **Interpretability** of the **model** and the **process**
    - Gaining a **better understanding of the world**
      when including the learned decision function in an existing theory

Often, we are **not** interested in prediction alone,
but in **understanding** the **prediction** and/or the **predictive model**

We are **thinking** beings

A **basic** principle

- Machine Learning “just” **reformulates** what has been given as input
- A conservation theorem:
  - **No information is “added**
  - Data + prior knowledge

A **basic** principle

- Machine Learning "just" reformulates what has been given as input
- A conservation theorem:
  - **No information is "added"**
  - Data + prior knowledge

Little data + **lots** of prior knowledge
Big data + **less** prior knowledge

A basic principle

- Machine Learning "just" reformulates what has been given as input
- A conservation theorem:
  - **No information is "added"**
  - Data + prior knowledge

Little data + **lots** of prior knowledge
Big data + **less** prior knowledge

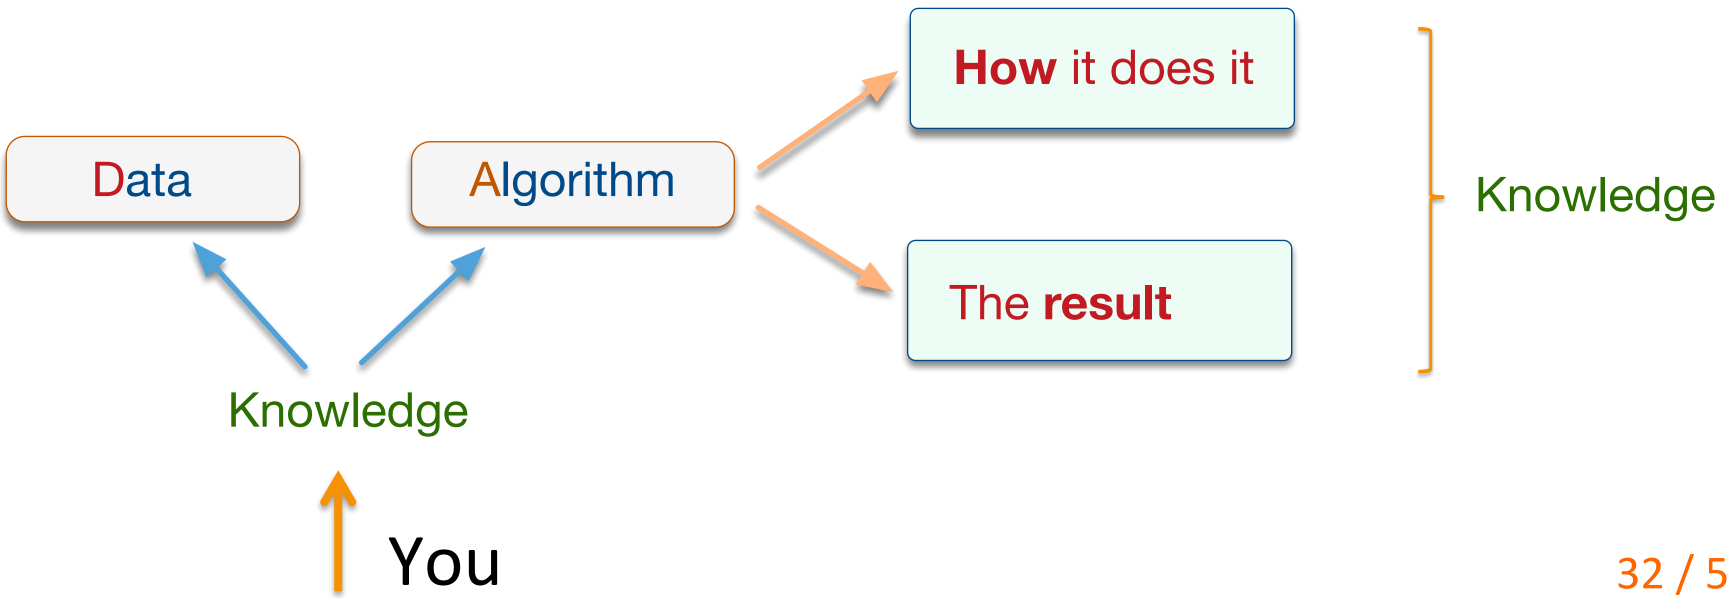

A basic principle

- Machine Learning "just" reformulates what has been given as input
- A conservation theorem:
  - **No information is "added"**
  - Data + prior knowledge

Little data + **lots** of prior knowledge
Big data + **less** prior knowledge

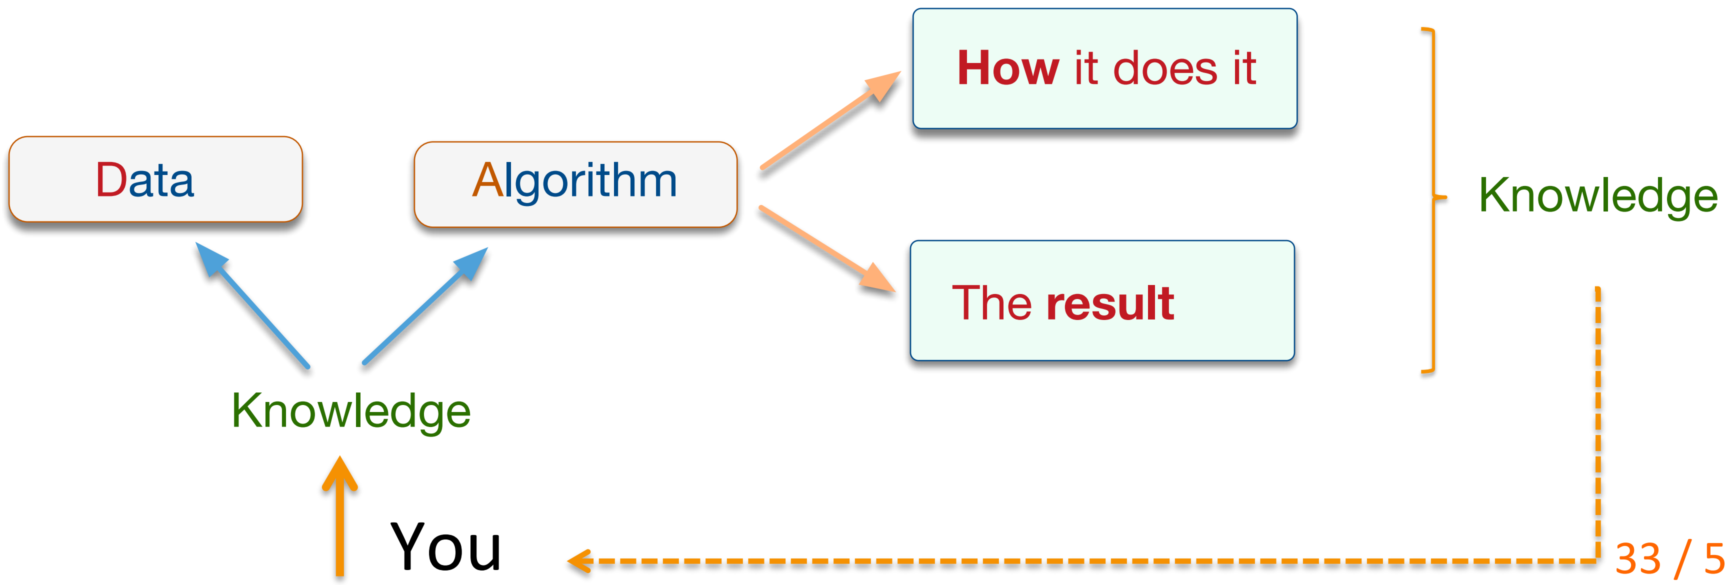

Prior knowledge

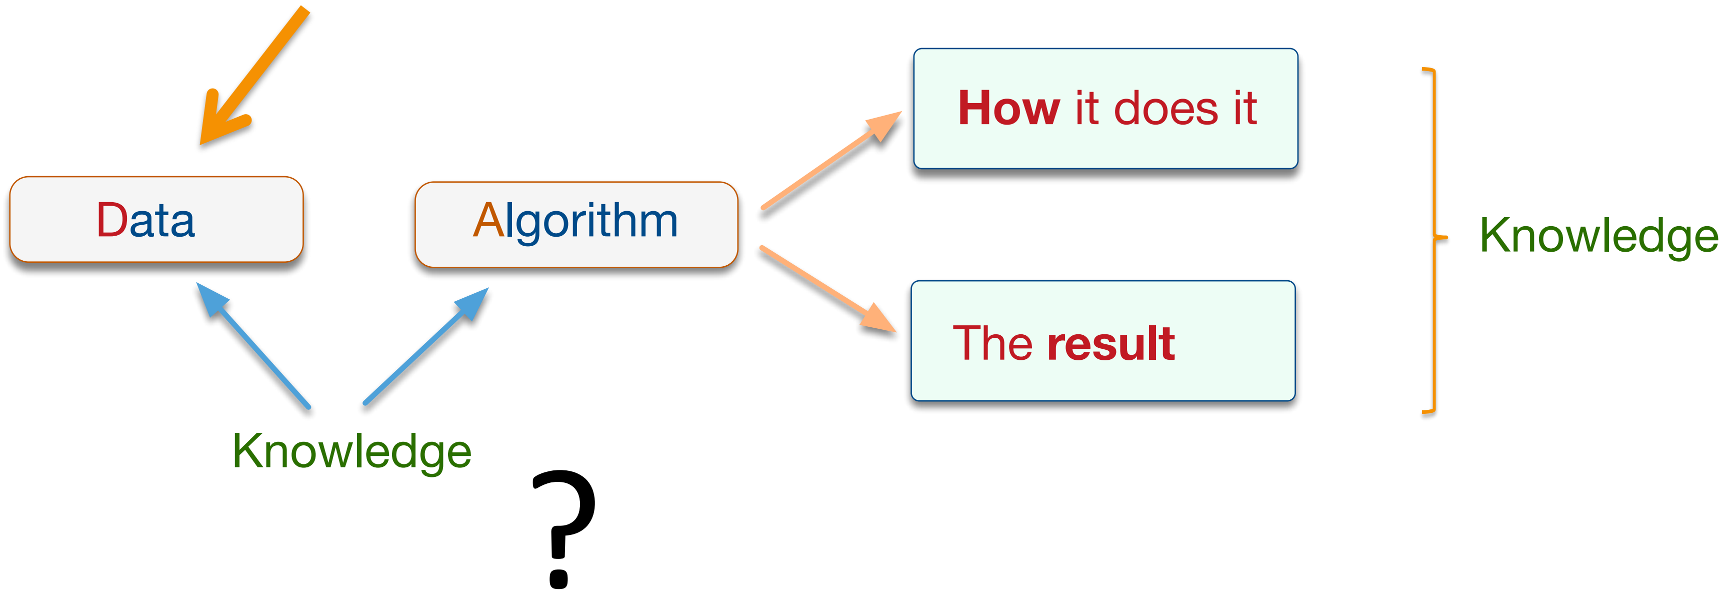

...

Knowledge as **input** to ML

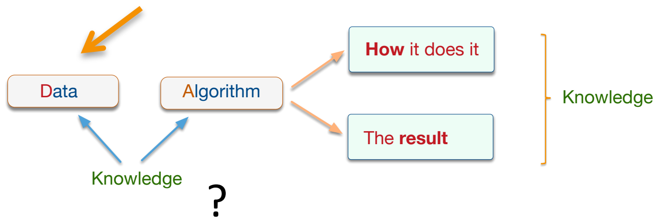

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

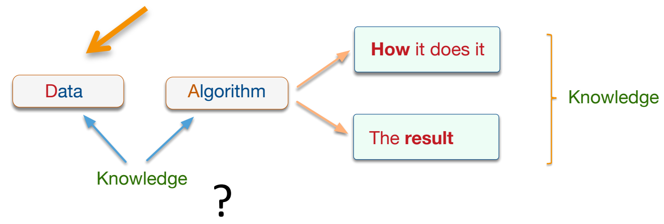

- Knowledge in the data
  - The **experimental apparatus** → With its own imperfection and biases
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

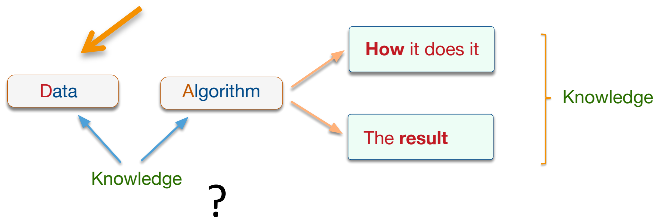

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features) → Necessarily biased
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

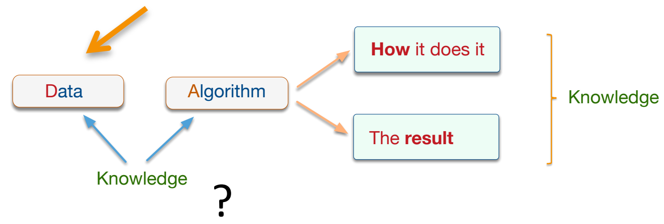

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies → Prior (and biased) knowledge
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

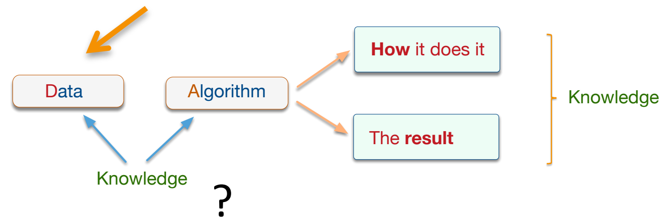

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values → No perfect normalization
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

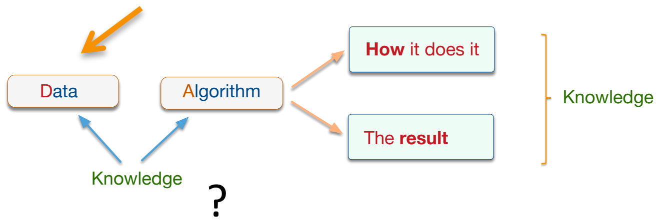

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values → What choice of imputation method?
  - Possibly **added data point**
    - With invariances in mind
  - …

Knowledge as **input** to ML

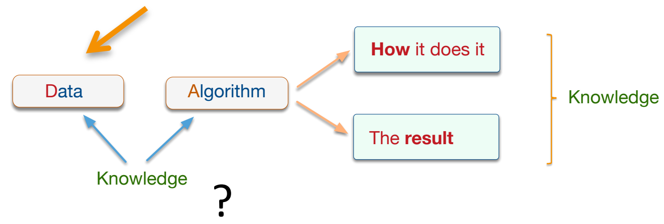

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

How do you add points?

Choice of invariance (prior assumptions)

Knowledge as **input** to ML

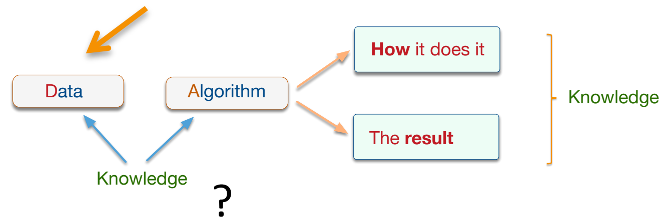

- Knowledge in the data
  - The **experimental apparatus**
  - Choice of **the descriptors** (the features)
  - **Enrichment** using ontologies
  - **Normalization** of the values
  - **Missing** values
  - Possibly **added data point**
    - With invariances in mind
  - …

Prior assumptions everywhere

Knowledge as **input** to ML

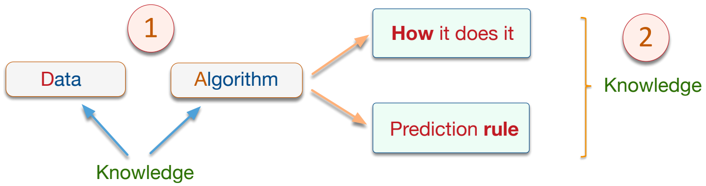

- Knowledge in the data

And sometime, there is not even that much **available** data!

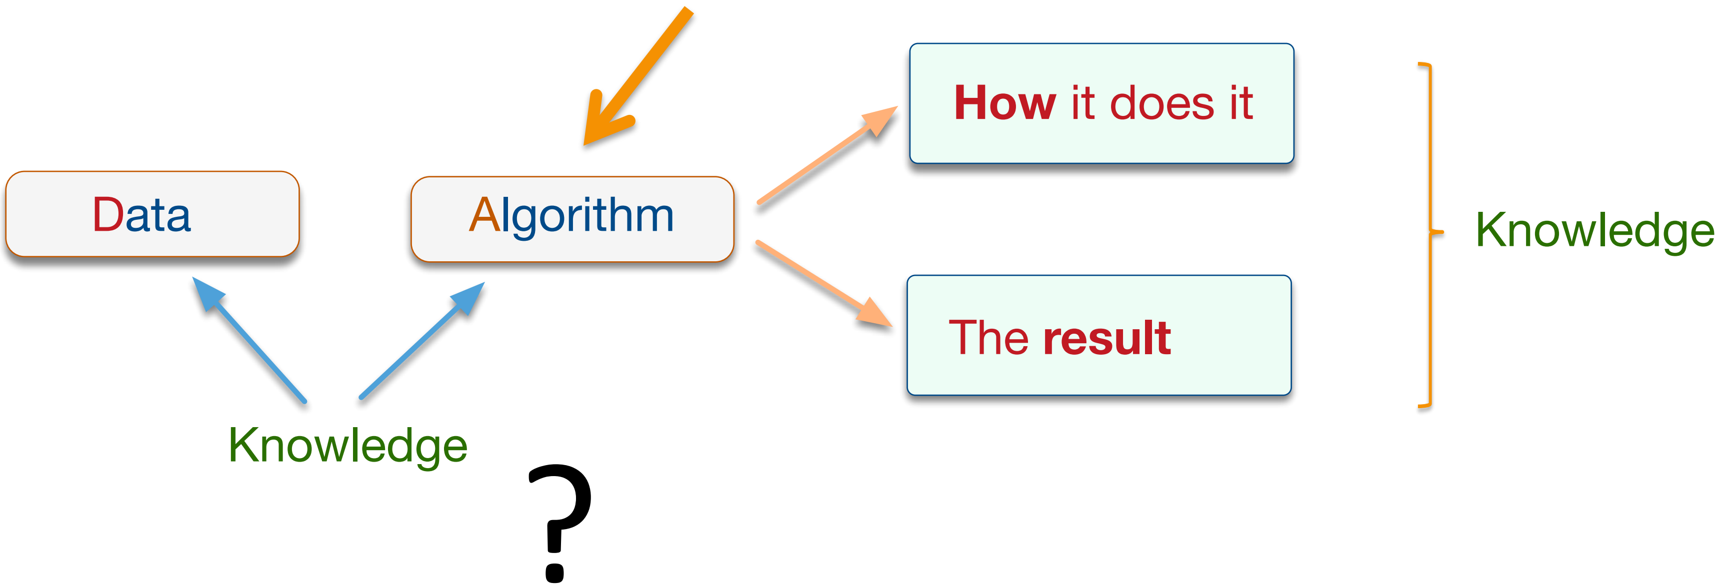
Data
Algorithm
How it does it
The result
Knowledge
Knowledge
?

...

Knowledge as **input** to ML

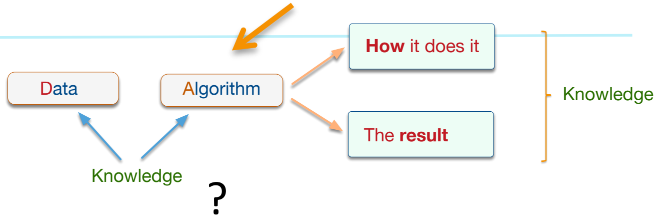

- Knowledge in the learning algorithm
  - Constraints on the hypothesis space: **representation bias**

$$h^\star \; = \; \operatorname*{ArgMin}_{h \in \mathcal{H}} \left[R_{\mathrm{Emp}}(h) \; + \lambda \, \mathrm{reg}(h) \right]$$

Looking for **sparse linear hypotheses**

$$h^\star \; = \; \operatorname*{ArgMin}_{h \in \mathcal{H}} \left[\frac{1}{m} \sum_{i=1}^{m} \ell(h(\mathbf{x}_i), y_i) \; + \lambda \, ||h||_1 \right]$$

Favors hypotheses with few non null coefficients

Knowledge as **input** to ML

- Convolutional Neural Networks
  - Knowledge embedded in the architecture of the network

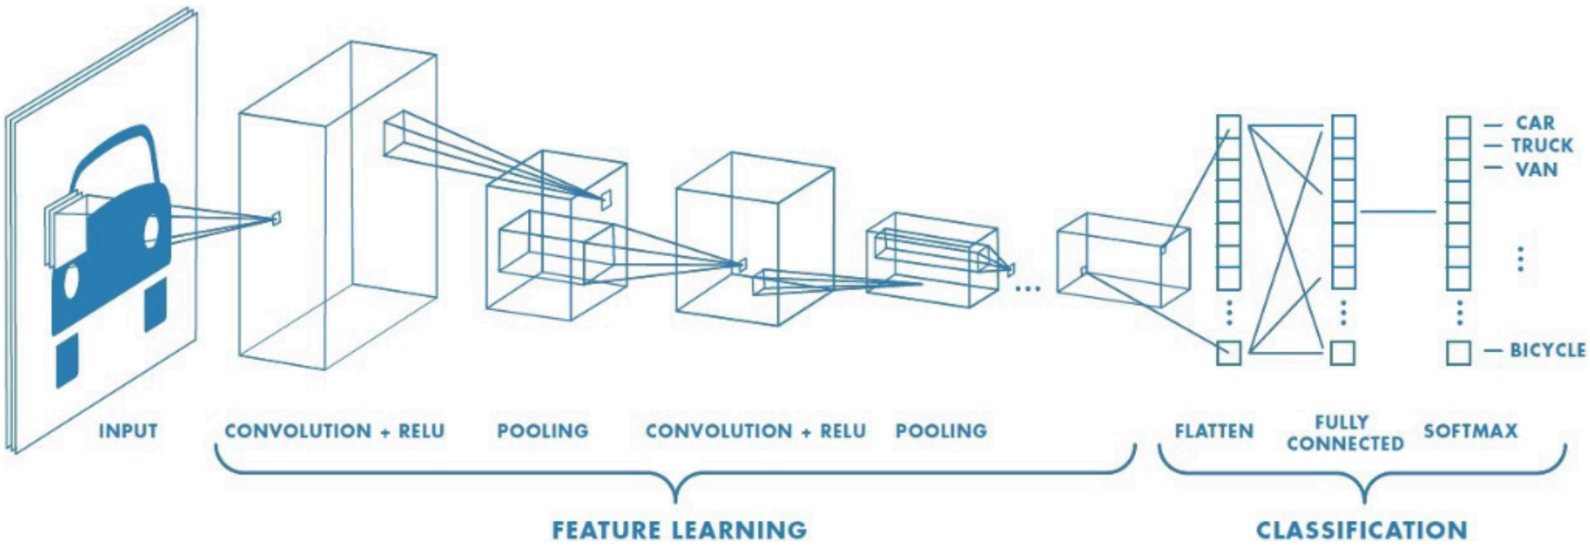

From https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53

What current Inductive Learning is **good at**

1. Identifying patterns in data (DESCRIPTIVE learning)
   - But **no guarantees** about the value of their value

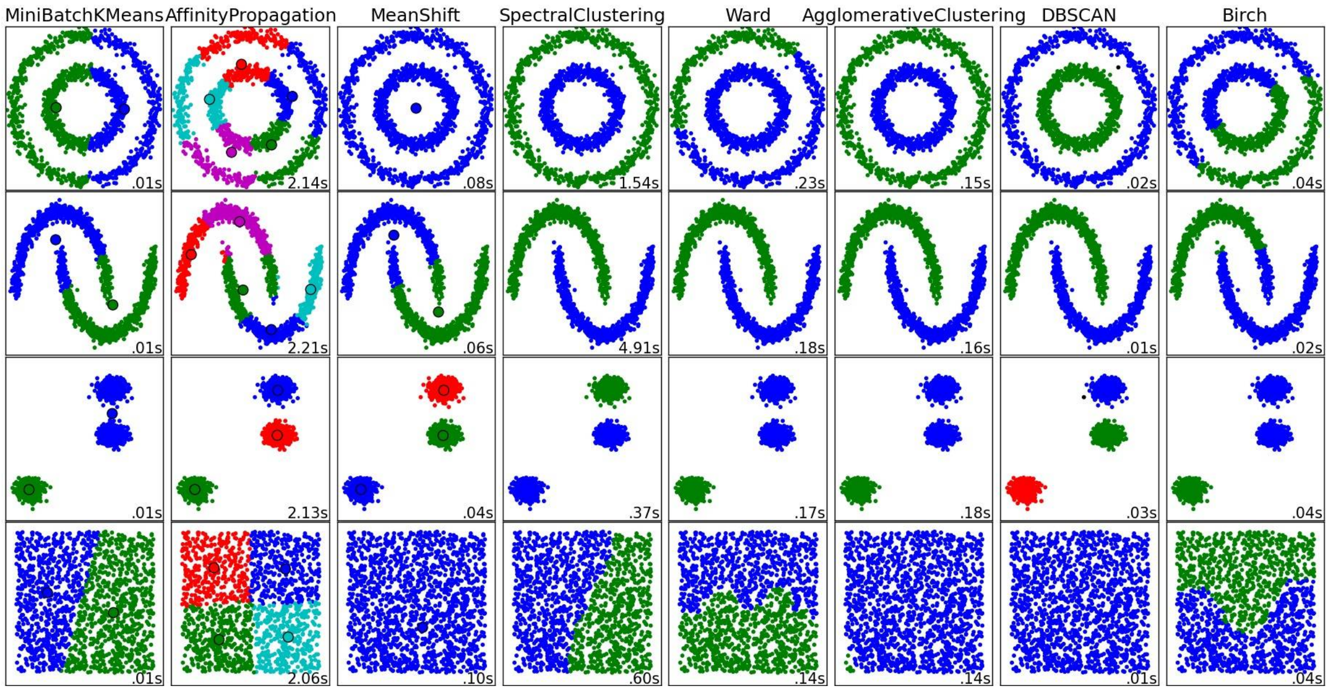

What current Inductive Learning is **good at**

2. Discover prediction rules based on statistical correlations (PREDICTIVE learning)
   - Geared towards **minimizing prediction errors**
   - In **stationary** environments
   - **Statistical** correlations: needs lots of data and ...

Is this less of a car
because the context is wrong?

Some current **challenges**

- Non stationary environments
- Heterogeneous multiple sources
- Interpretability of the results
- Integration in larger reasoning systems (including people)

Induction is a **risky** business

1. You have to **invest a lot**

2. And **be very careful** about the yield

Machine Learning DOES NOT produce absolute truths

Do not give up your **critical sense** at every step!

Induction is a **risky** business

1. You have to **invest a lot**

2. And **be very careful** about the yield

But do not abandon **hope**

Machine Learning is a useful tool ... in good hands