Machine Learning and Knowledge Engineering

Knut Hinkelmann

Knowledge Engineering

Knowledge

internal knowledge

documented knowledge

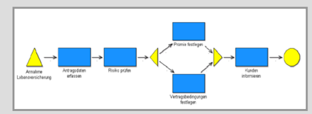

$E = m\,c^2$

automated knowledge

knowledge bases

Knowledge Engineering: Human-Created Knowledge Base

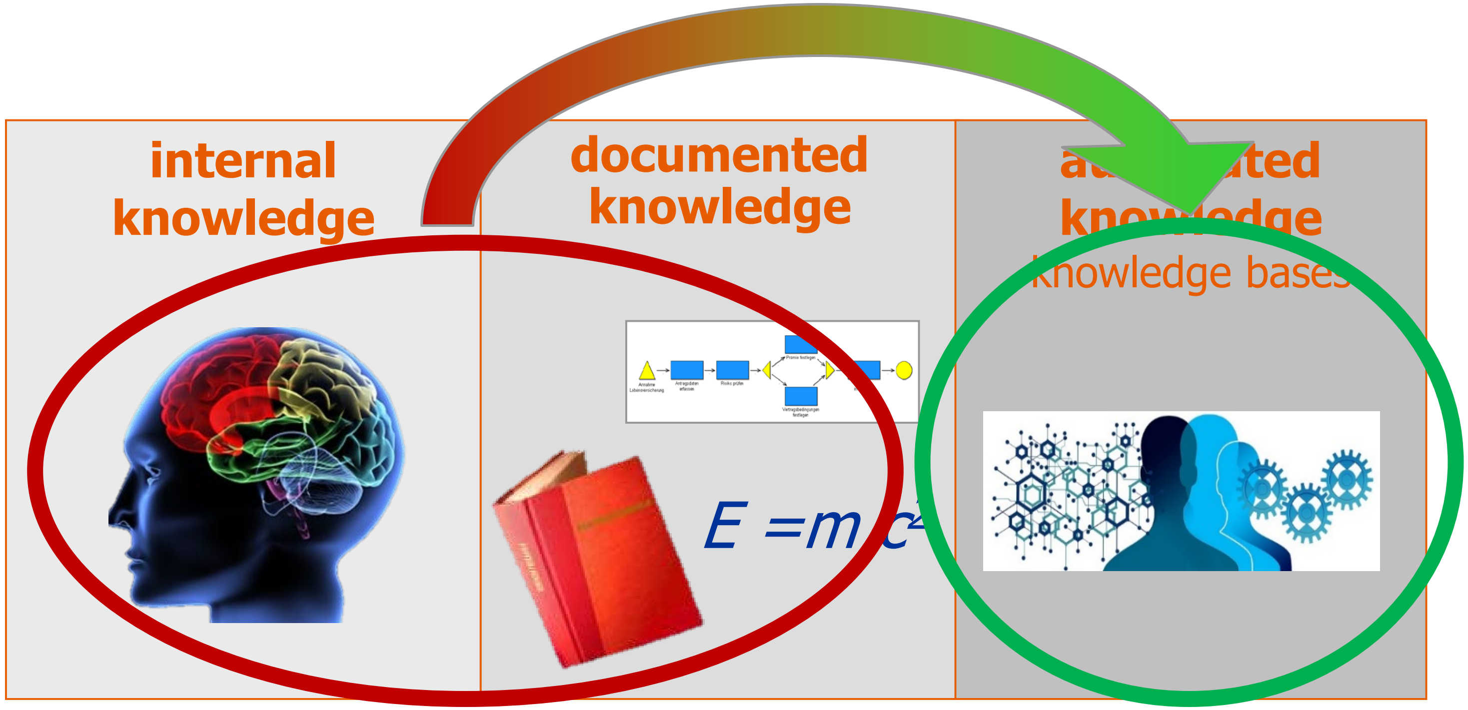

Knowledge-Based Systems (Expert Systems)

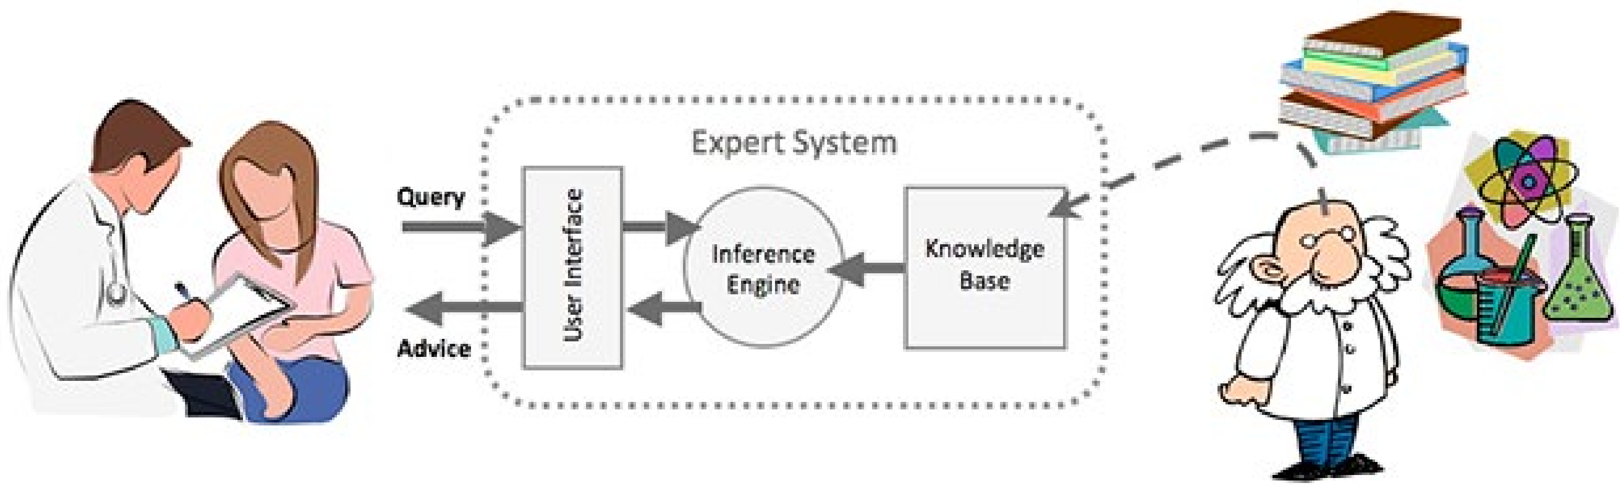

Knowledge-Based Systems (Rules & Facts)

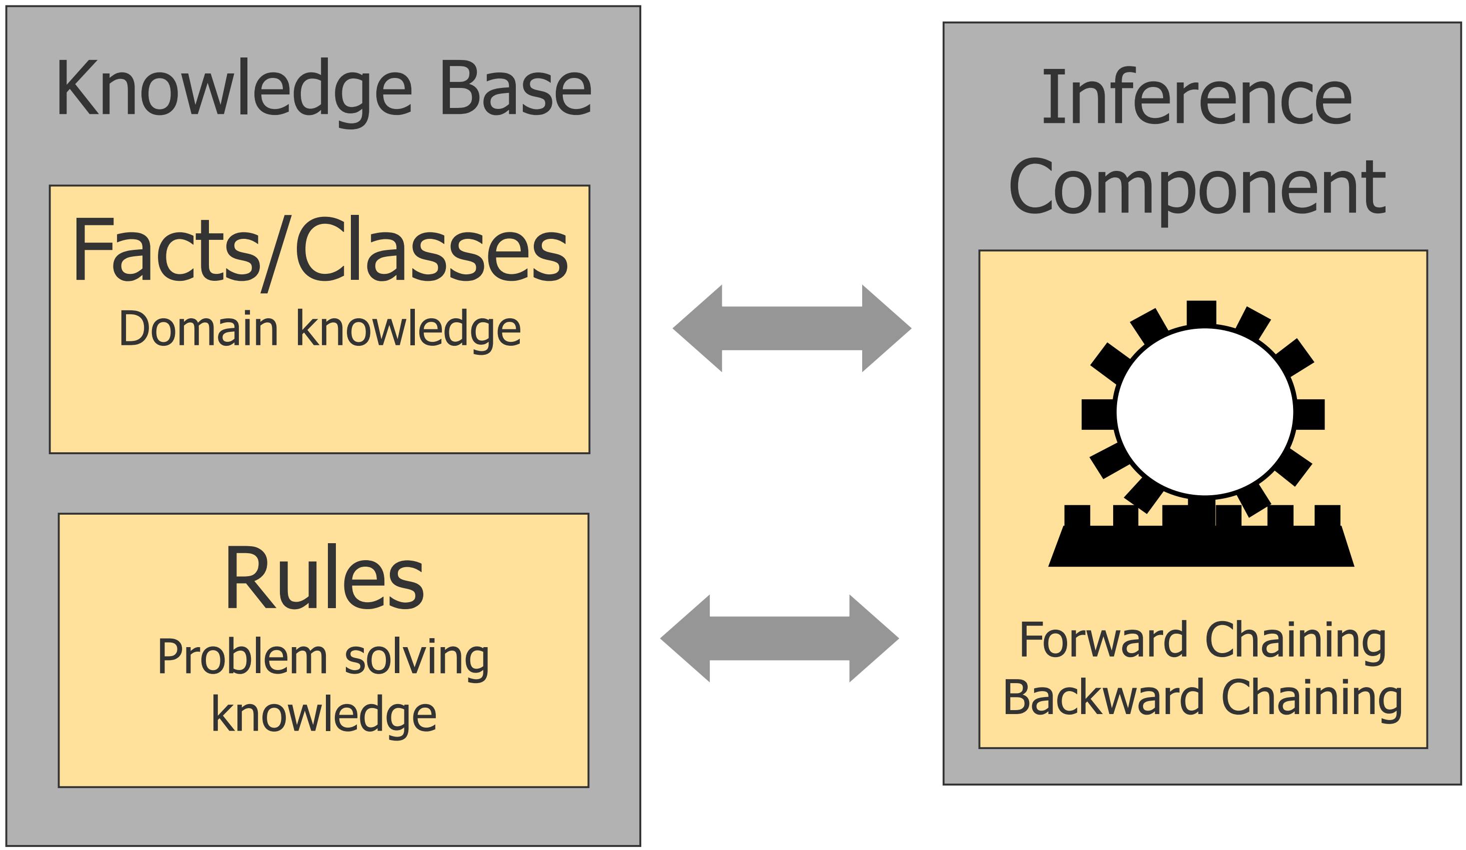

Knowledge Bases and Inference

- *Inference:* Making implicit knowledge explicit by generating new facts from knowledge
- Example:

  From the knowledge that ***socrates is a human*** *and that* ***all humans are mortal****, one can conclude that* ***socrates is mortal***

Knowledge Base

Facts: human(socrates)

Rules: human(X) → mortal(X)

—inference→ mortal(socrates)

Inference

What knowledge can be derived from this knowledge base?

Knowledge Base

Facts:

father(peter,mary)
father(peter,john)
mother(mary,mark)
mother(jane,mary)

Rules:

father(X,Y) → parent(X,Y)
mother(X,Y) → parent(X,Y)
father(X,Y) AND parent(Y,Z) → grandfather(X,Z)
mother(X,Y) AND parent(Y,Z) → grandmother(X,Z)
parent(X,Y) AND parent(X,Z) → sibling(Y,Z)

Knowledge-Based Systems (Classes and Instances)

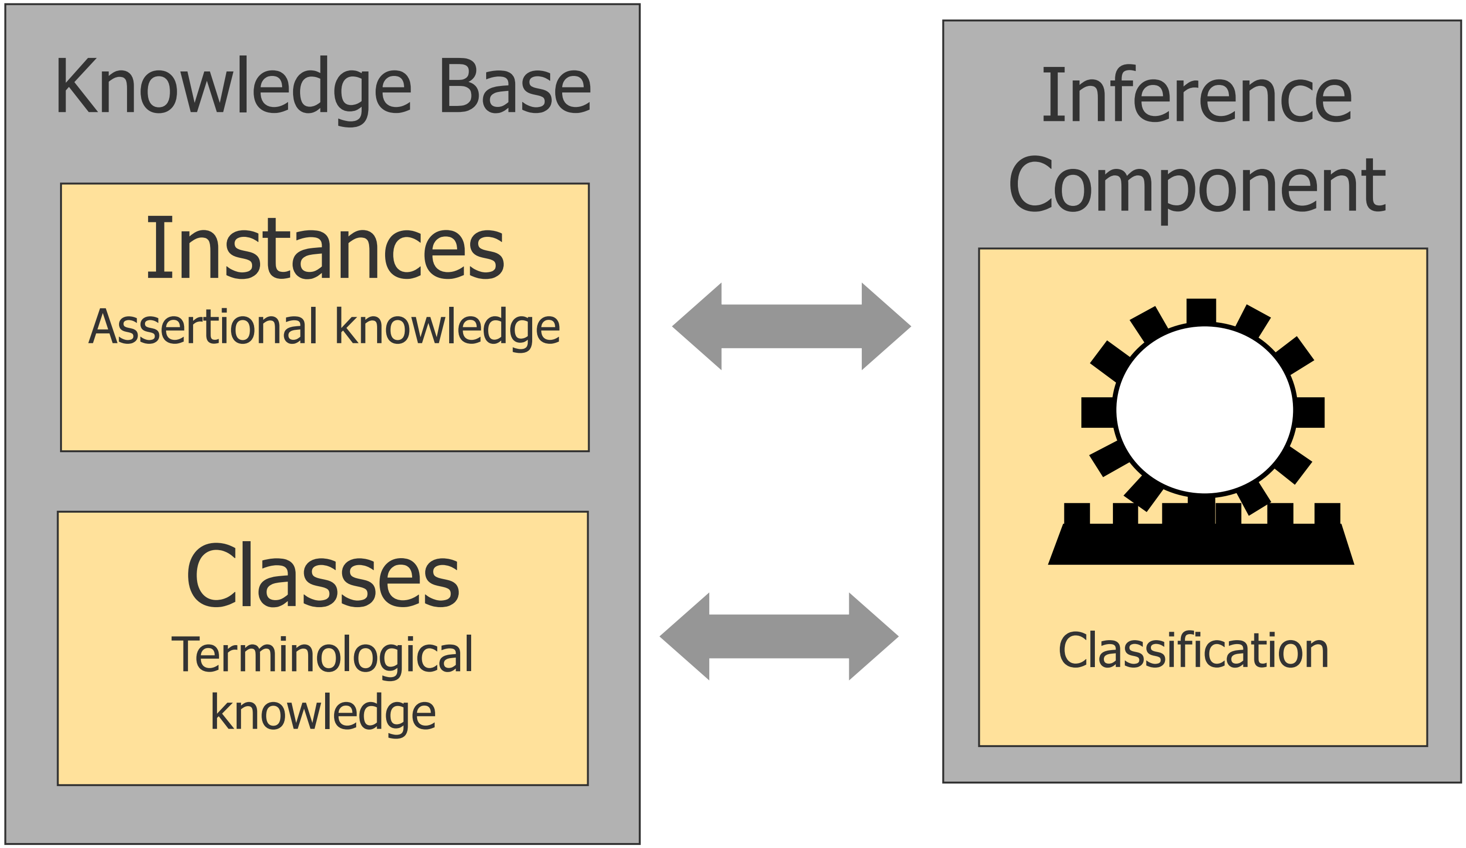

Domain Concepts, Instances and Relations

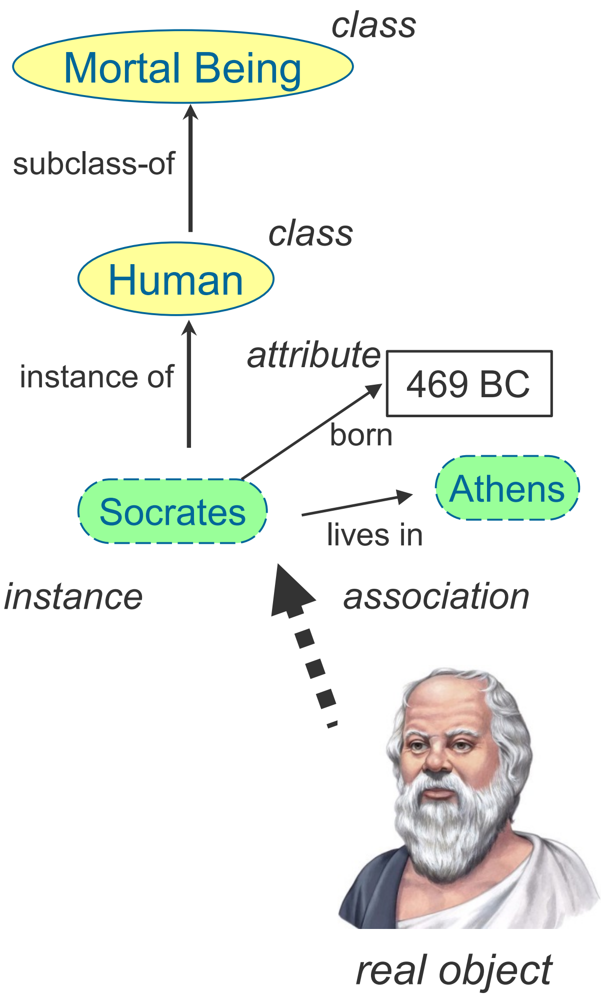

- There are two kinds of concepts:
  - **classes**
  - **instances (facts)**
- There are different kinds of relations
  - **generalisation** ("is a")
    - between classes **(subclass of)**
    - between instance and class **(instance of)**
  - **associations/properties**
    - any other kind of relationship
- Attributes can be regarded as properties whose value is not an instance but is of a primitive type (number, string).

Machine Learning

Self-driving Cars

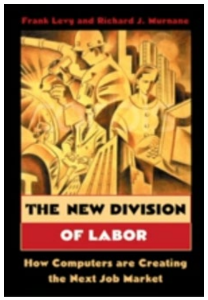

"… it is hard to imagine discovering the set of rules that can replicate the driver's behavior."

(Levy & Murnane 2006)

Face Recognition

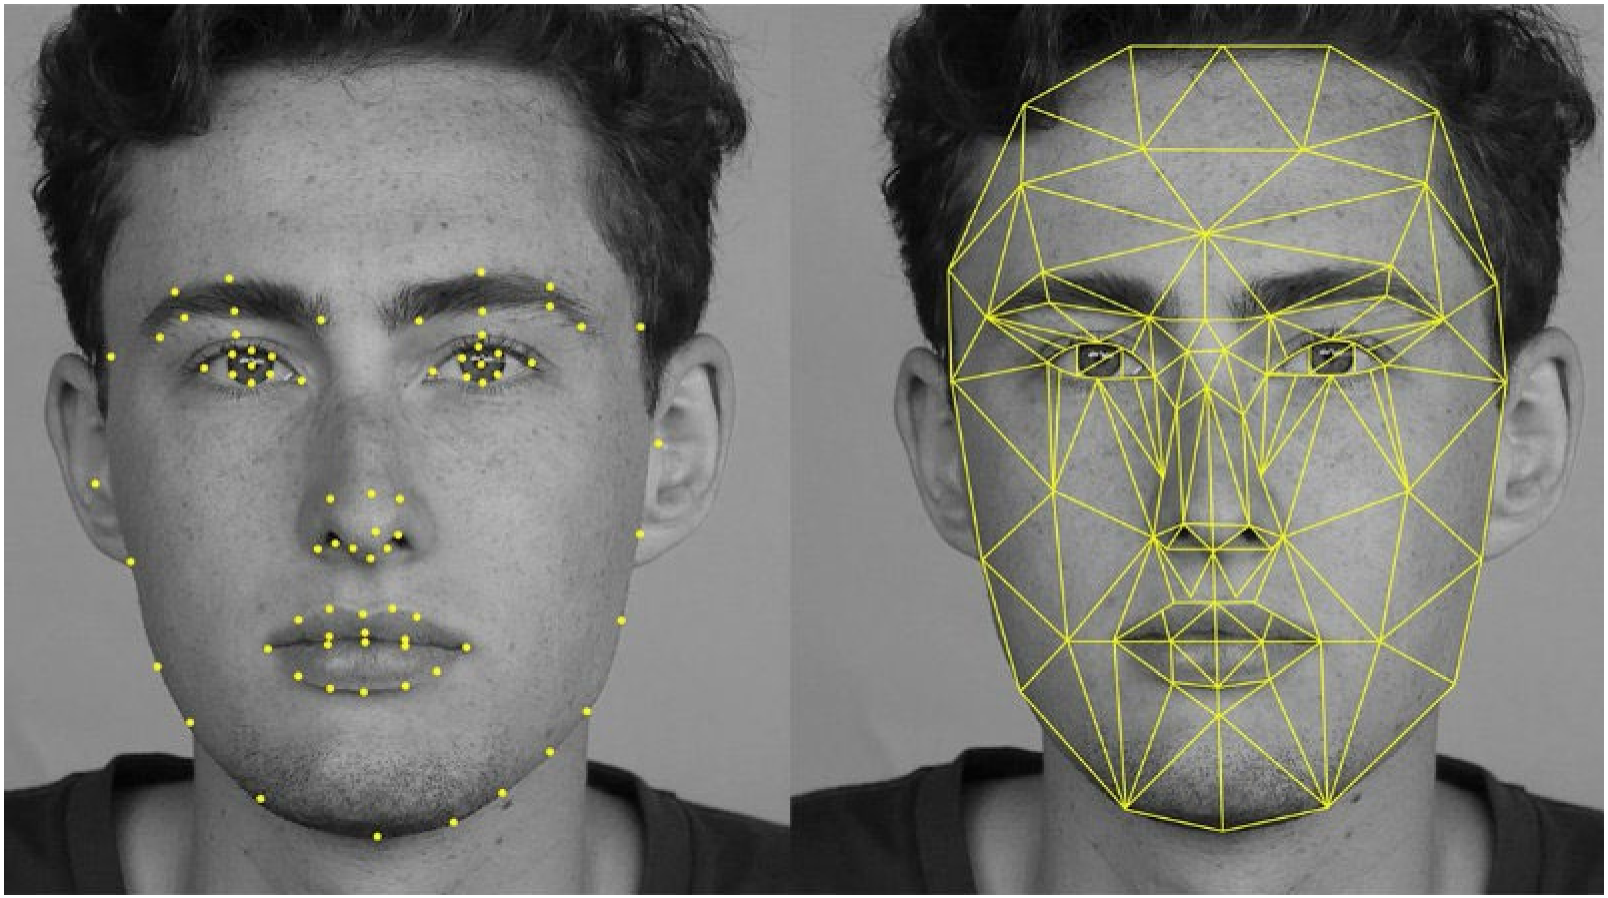

Spam Filter

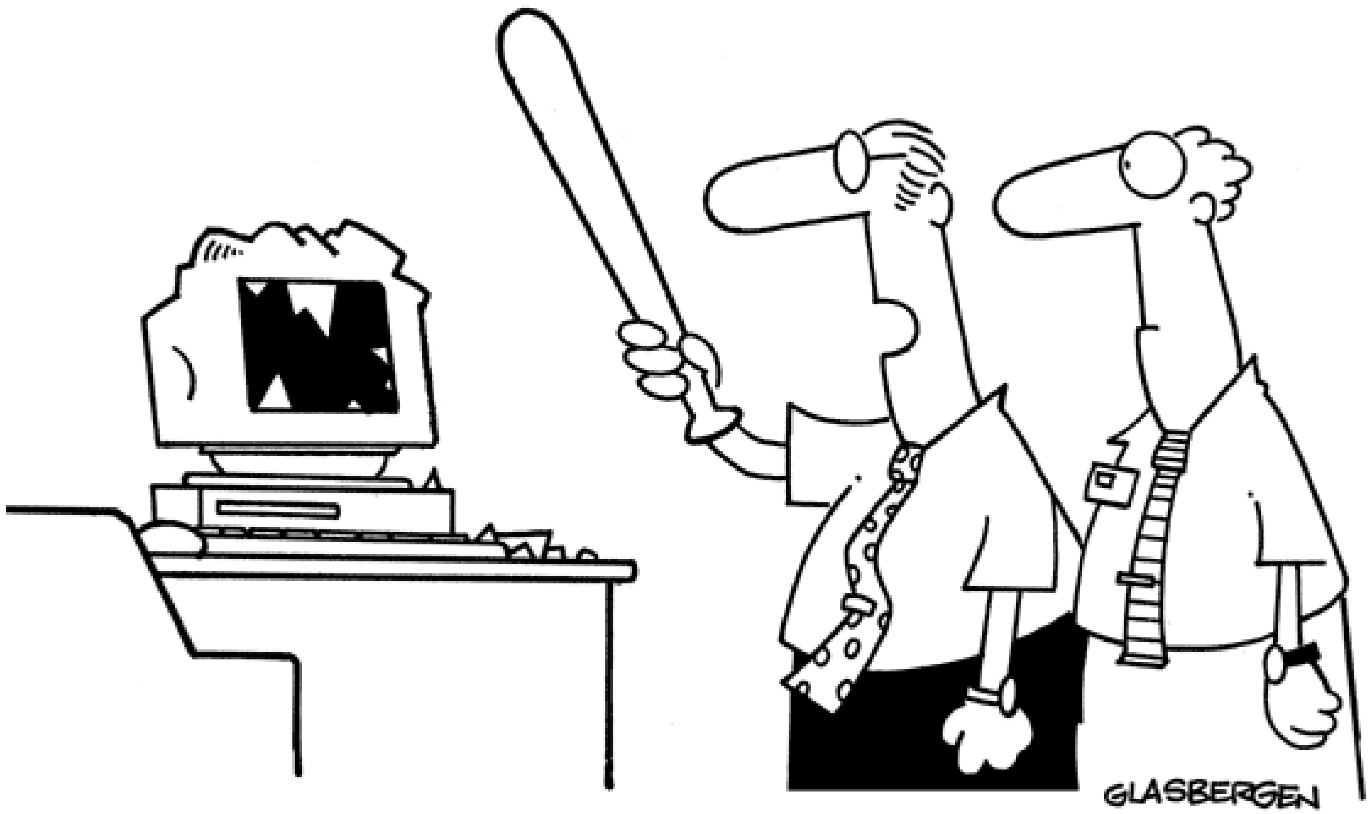

"It's not the most sophisticated Spam blocker I've tried, but it's the only one that works!"

Translation

"The spirit was willing, but the flesh was weak" ... *"The vodka was good, but the meat was rotten"*

"Out of sight, out of mind." ... *"Blind and insane."*

https://www.snopes.com/language/misxlate/machine.asp

Knowledge

<table>
<tr><th colspan="2">internal knowledge</th><th colspan="2">external knowledge</th></tr>
<tr>
<td>tacit knowledge
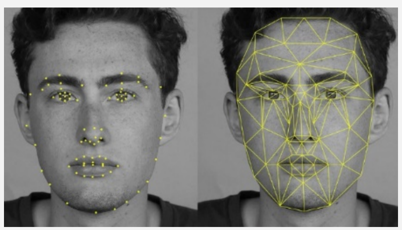</td>
<td>self-aware knowledge
</td>
<td>documented knowledge

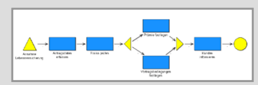
$E = m c^2$</td>
<td>automated knowledge
knowledge bases
</td>
</tr>
</table>

Machine Learning: Learning (Tacit) Knowledge from Data

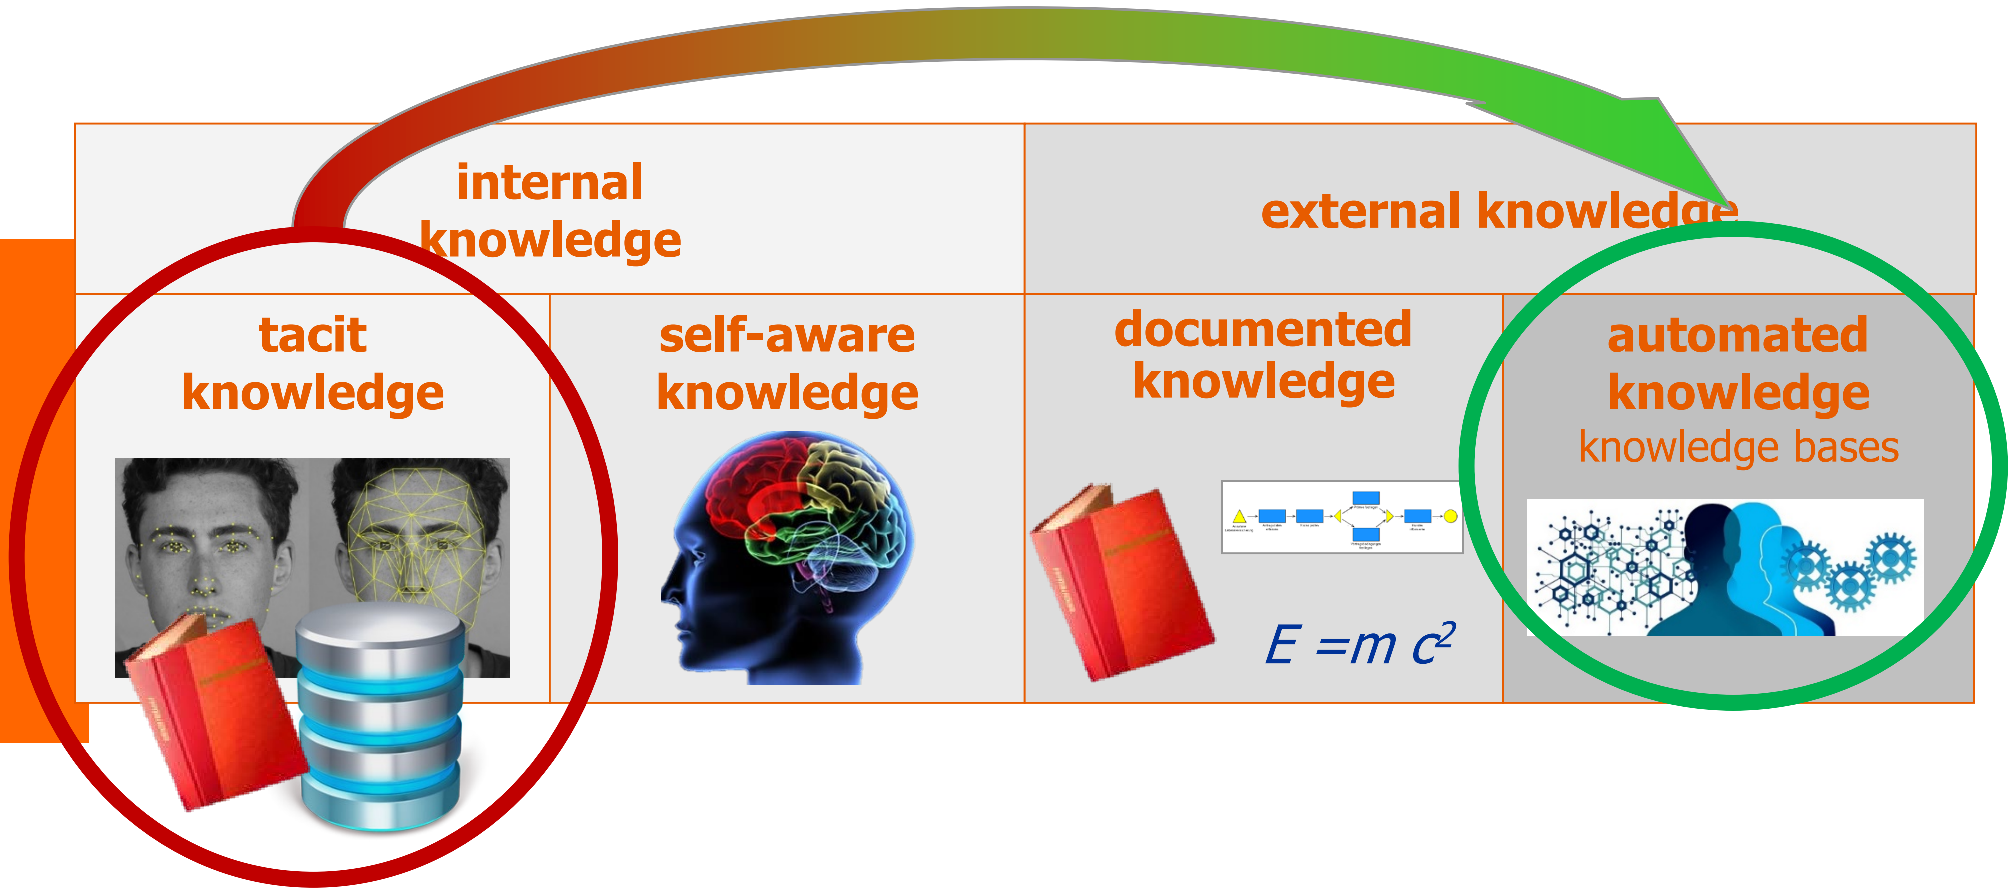

Machine Learning: General Idea

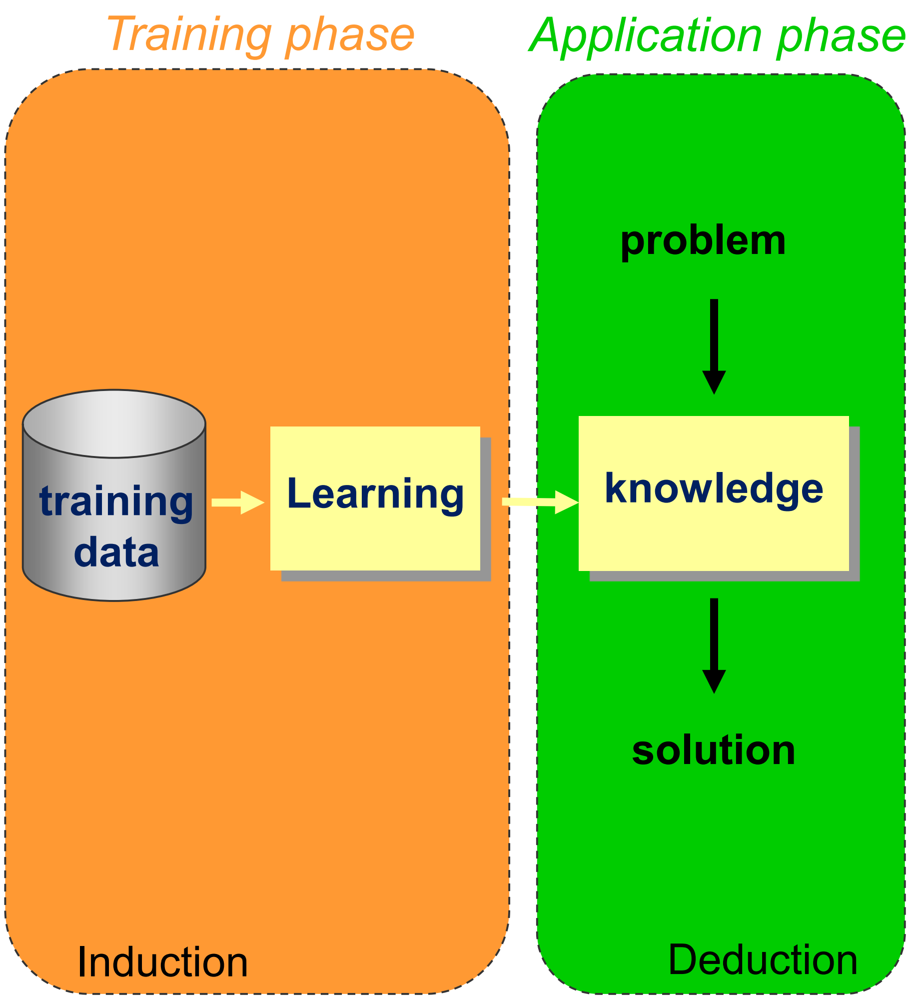

- **Training**
  - Collect data for the problem
  - Use the data to learn how to solve the type of problem
  - Result: Knowledge base
- **Application**
  - Use the learned knowledge for new problems

Machine Learning vs. Knowledge-based Systems

Knowledge-based System

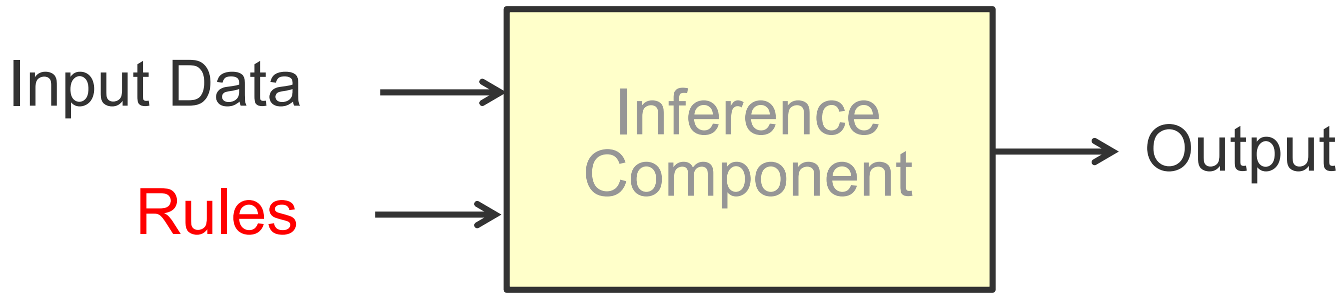

Machine Learning

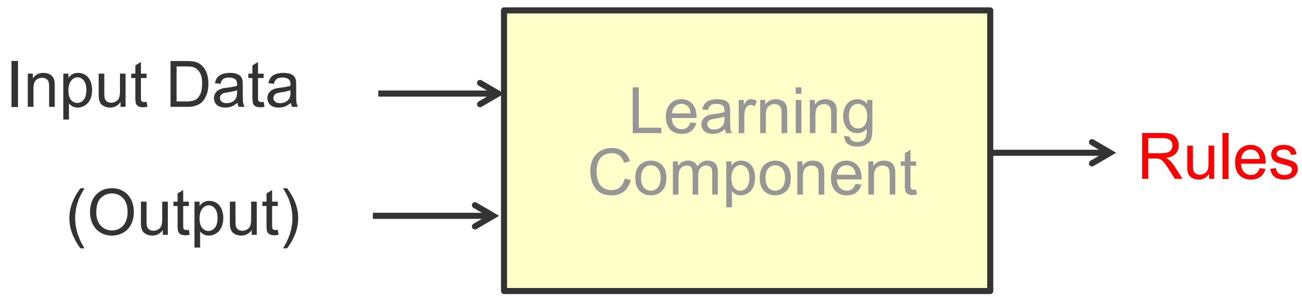

Source: Vibhav Gogate, UT Dallas

Machine Learning in Context

- Machine Learning (Data Mining) is a step to discover knowledge in data

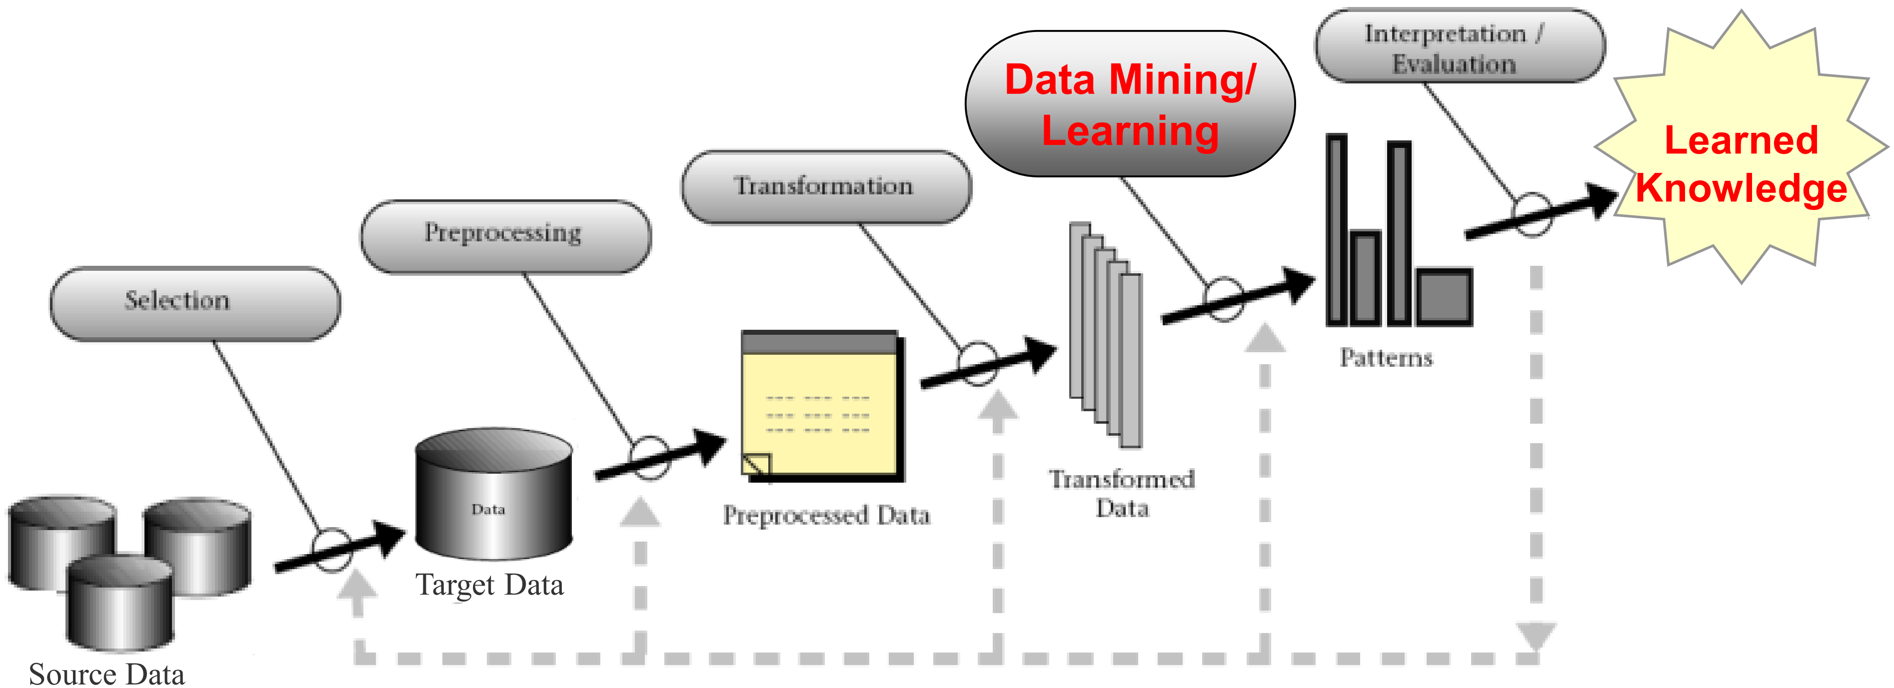

Knowledge is then used in processes and applications.

(Fayyad et al., 1996)

Types of Learning

- The learning method depends on the kind of data that we have at our disposal
  - The data contains sets of inputs and corresponding outputs: (i,o)
  - No prior knowledge: The data contains only the inputs i: output has to be determined
  - The data contains sets of inputs without corresponding «correct» output, but we can get some measure of the quality of an output o for input i.
    Rewards for good output quality.

Supervised Learning

Unsupervised Learning

Reinforcement Learning

(Lison, 2012)

Supervised Learning: Application Examples

	Input ***i***	Output **o**
Spam filtering	An email	{spam, non-spam}
Face recognition	An image	Identified faces
Machine translation	A sentence in language A	A sentence in language B
Speech recognition	A speech signal	A (text) sentence
Fraud detection	A financial transaction	{fraud, non-fraud}
Robot motion	Sensory data	Motor control

(Lison, 2012)

Supervised Learning

Example: Classification

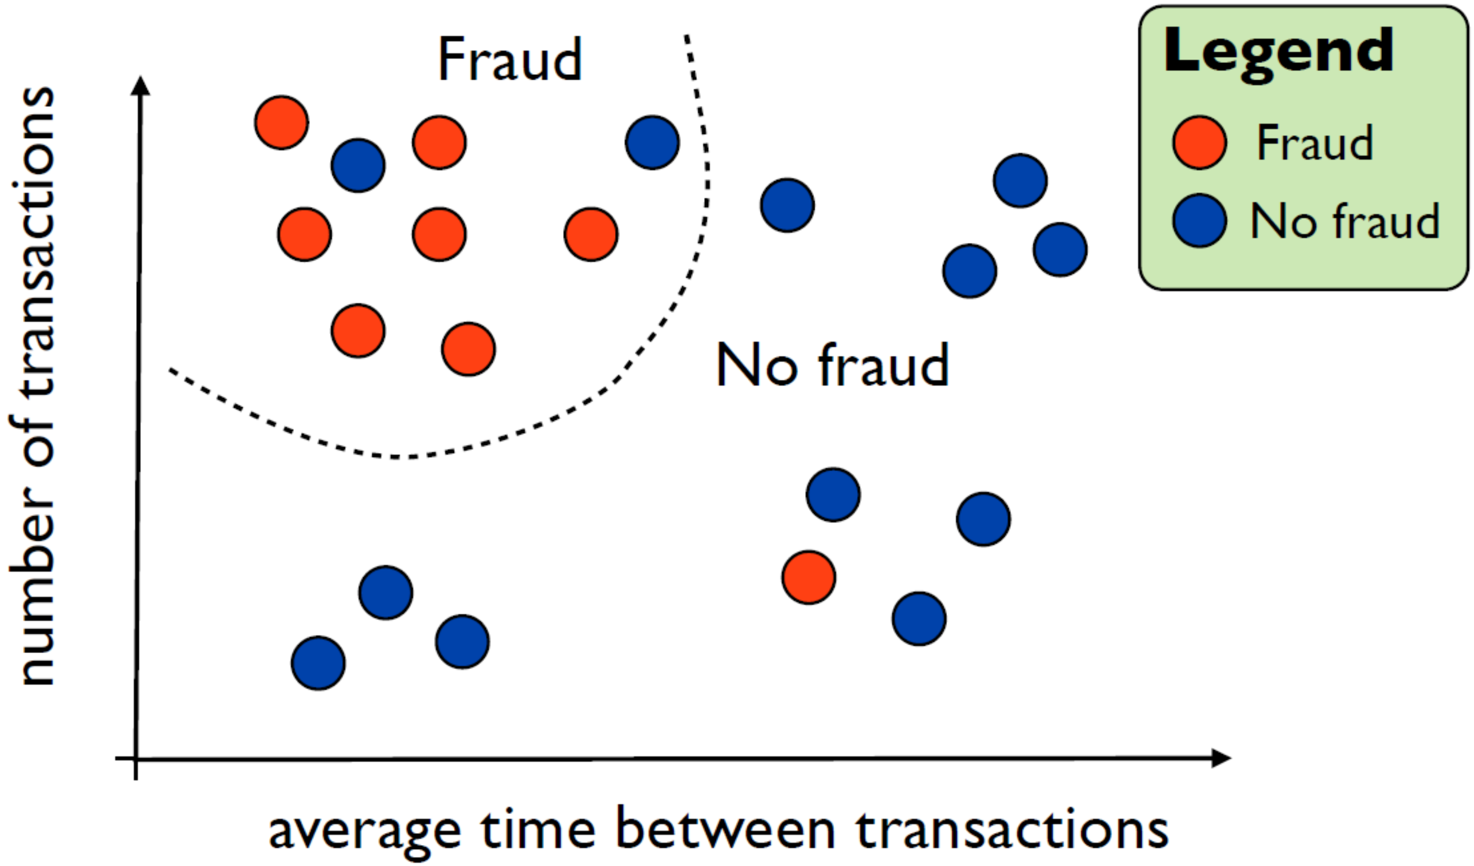

(Lison, 2012)

Classification

- Assign objects (input) to known classes (output)
- Examples:
  - credit assessment
    Input: customers of a bank
    Classes: credit worthy
    not credit worthy
  - Spam filtering
    Input: email
    Classes: spam
    non-spam
  - optical character recognition (OCR)
    Input: scanned pixel image
    Classes: ASCII characters

Training and Application Phase

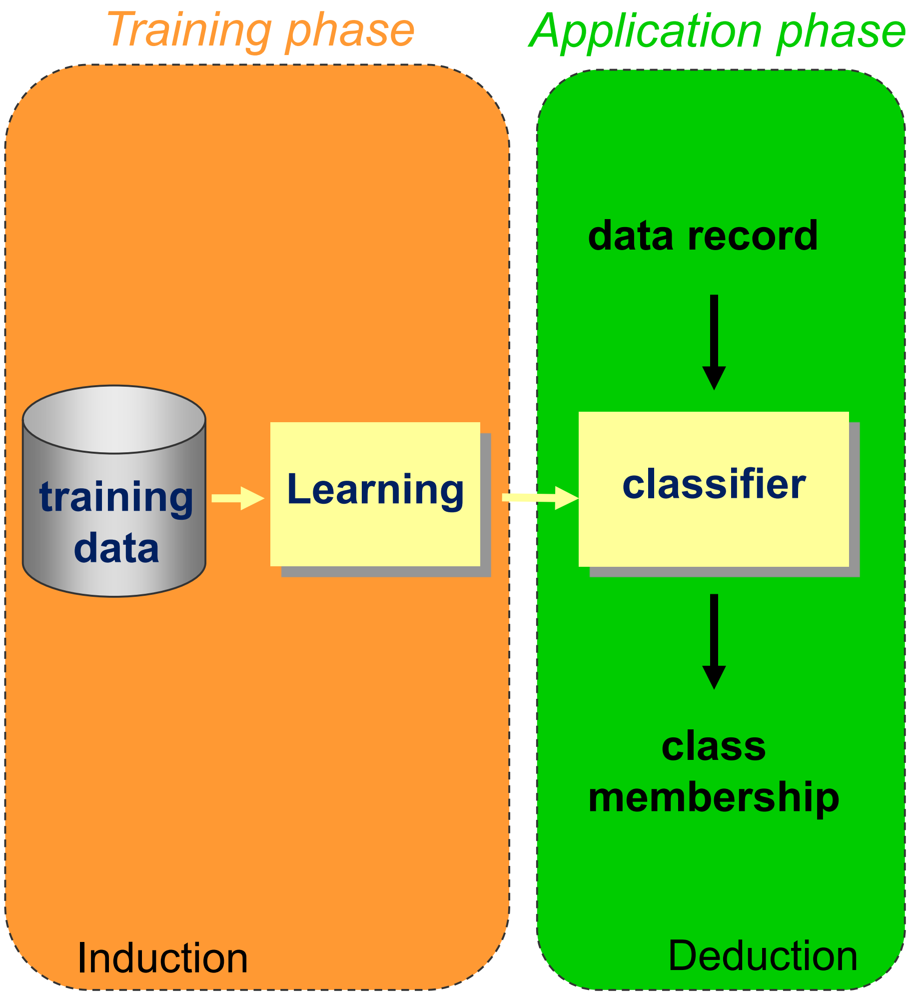

- **Training:** Learning the classification criteria
  - Given: sample set of training data records
  - Result: Decision logic to determine class from values of input attributes (decision tree, rules, model)
- **Application**: Classification
  - Goal: assign a class to previously unseen records of input data as accurately as possible

Supervised Learning: Classification Criteria

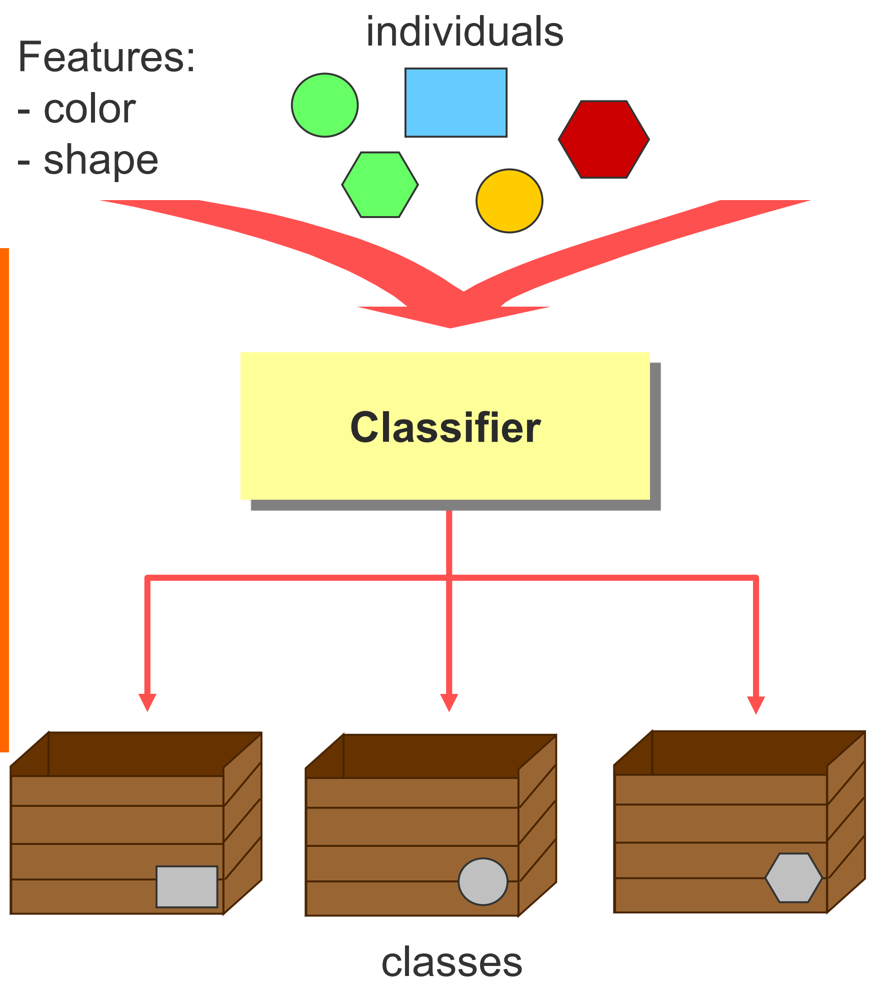

- The classifier decides, which individual belongs to which class
- Problem:
  - Input has different features
  - The criteria for the decision are not always obvious
- Supervised Learning:
  - Learn the classification criteria from known examples
  - Criteria = relevant features and their valures

Example for Supervised Learning

Problem: When to give credit

Tid	Employed	Marital Status	Taxable Income	accept
1	No	Single	125K	No
2	Yes	Married	160K	Yes
3	Yes	Single	70K	No
4	No	Married	120K	No
5	Yes	Divorced	95K	Yes
6	Yes	Married	60K	No
7	No	Divorced	220K	No
8	Yes	Single	85K	Yes
9	Yes	Married	95K	No
10	Yes	Single	90K	Yes

Training Data

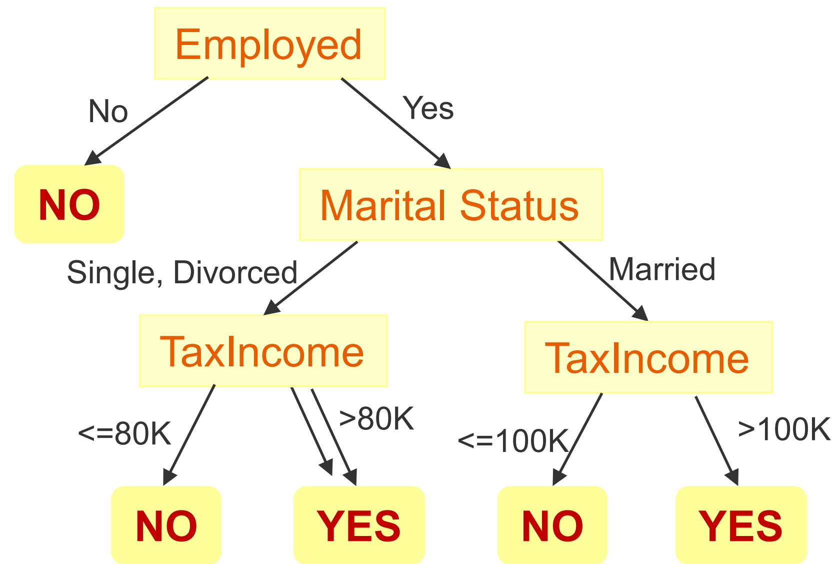

Credit Worthiness

	Employed	Marital Status	Taxable Income	Accept
	Yes, No	*Single, Divorced, Married*	*Integer*	*Yes, No*
1	No			No
2	Yes	Single	> 80K	Yes
3	Yes	Divorced	> 80K	Yes
4	Yes	Single	≤ 80K	No
5	Yes	Divorced	≤ 80K	No
6	Yes	Married	> 100K	Yes
7	Yes	Married	≤ 100K	No

Knowledge Base: Decision Tree , Decision Table

Classification Methods

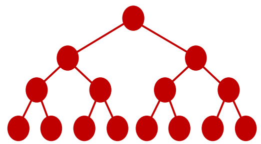

Decision Trees

criteria	class
… … …	…
… … …	…

Decision Table

IF …
THEN …

Rules

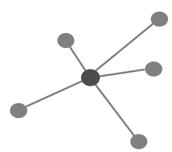

k-Nearest Neighbor

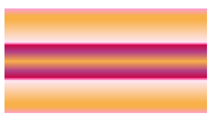

Genetic Algorithms

symbolic

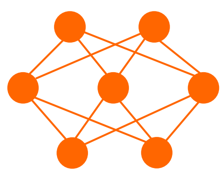

Neuronale Netze

subsymbolic

Unsupervised Learning

- Sometimes, we don't have access to any output value **o**, we simply have a collection of input examples **i**
- Input: data sets without corresponding output values.
- Objective: learn the underlying patterns of our data
  - Are there any *correlations* between features?
  - Can we *cluster* our data set in groups which behave similarly?

(Lison, 2012)

Unsupervised Learning

Example: Clustering (= identify new classes)

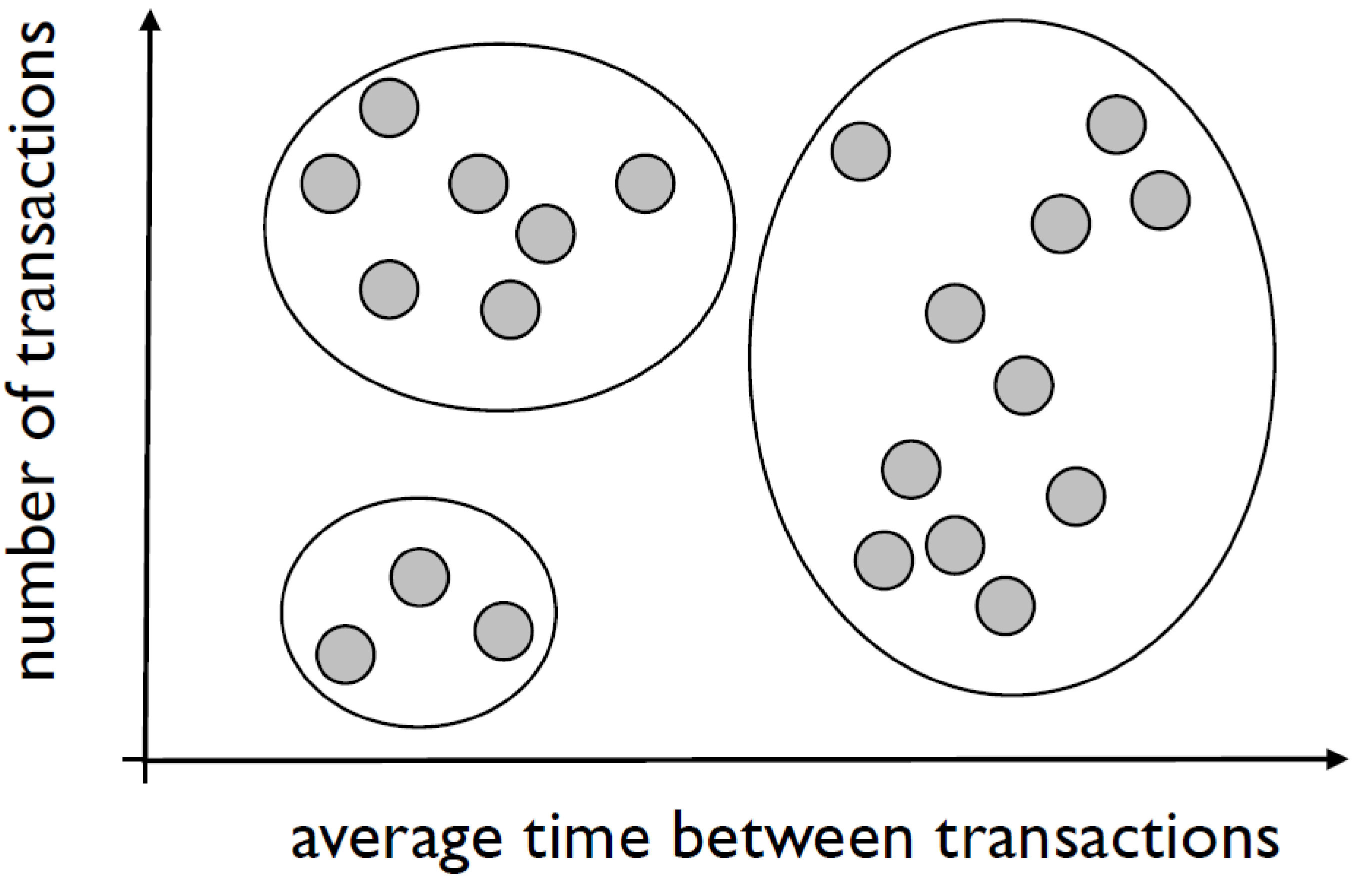

(Lison, 2012)

Example: Recommender Systems

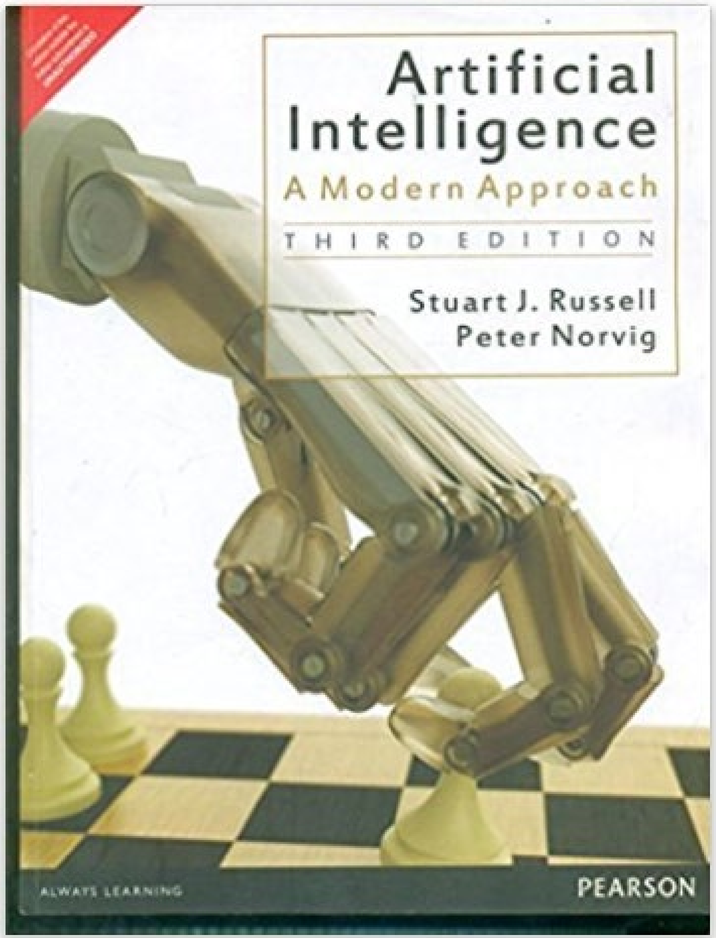

Reinforcement Learning

- Sometime we don’t have direct access to «the» correct output ***o*** for an input ***i***
- But we can get a measure of «how good/bad» an output is
  - Often called the *reward* (can be negative or positive)
- The goal of the agent is to learn the behaviour that maximises its expected cumulative reward over time
  - To learn how to flip pancakes, the reward could for instance be +3 if the pancake is flipped, -1 if the pancake stays in the pan, and -5 if it falls

Combining Machine Learning and Knowledge Engineering

Knowledge in Intelligent Systems

Machine Learning

Knowledge Engineering

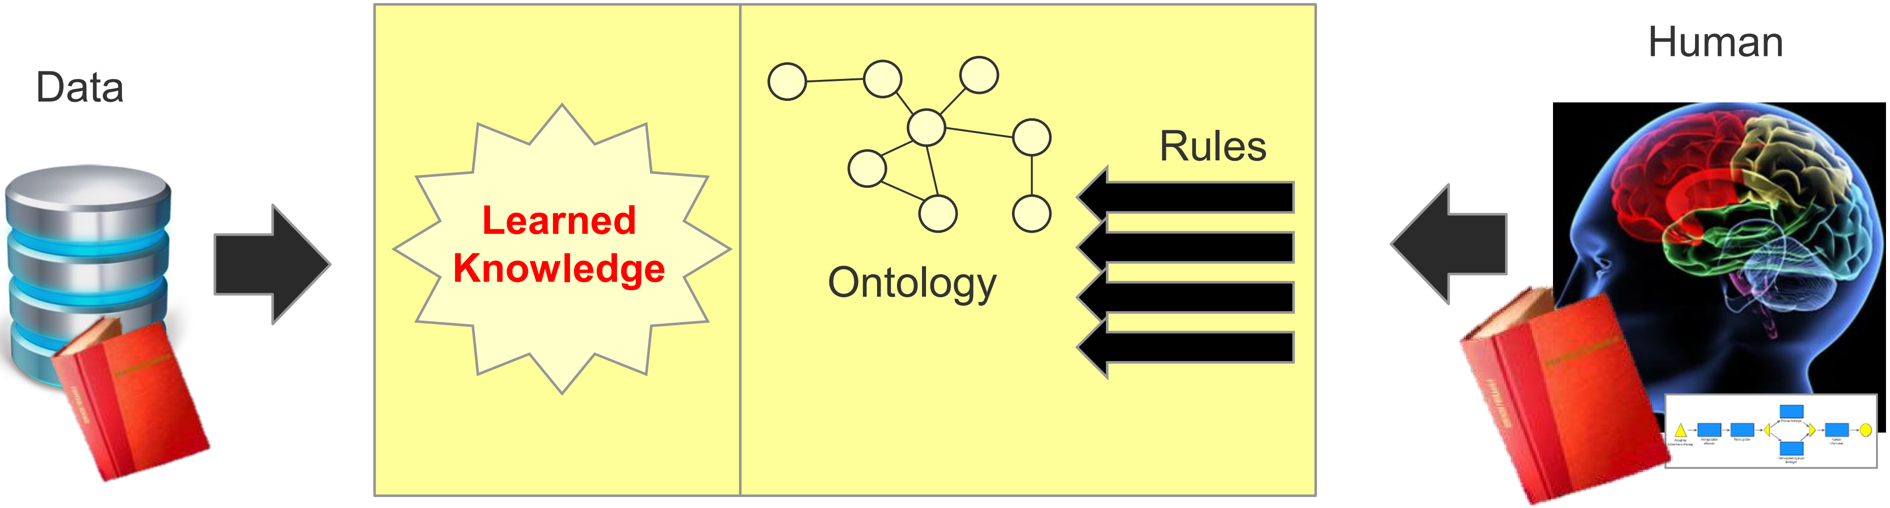

- Tacit or unknown knowledge
- Adaptable to new situations

- Knowledge we are aware of
- Knowledge that must be correct

Autonomous Driving

- Machine Learning:
  Driving Behaviour

- Knowledge Engineering:
  Traffic Rules

Eligibility Decision

Example: Insurance

Accept yes/no

Combining Machine Learning and Knowledge Engineering for Eligibility Decisions (1/2)

- Example: Application of health insurance
  - Machine Learning: data records about risks of clients

Age	surgery	docvisit	allergy	med	disieases	bmi	class
20	0	2	no	no	cholesterol	28	low
21	0	4	no	no	no	23	low
49	2	12	yes	yes	heart	34	high
22	0	3	no	no	no	23	low
51	2	2	yes	yes	diabetes	26	high
52	2	8	no	no	heart	31	high
52	0	3	yes	no	no	22	low
52	2	12	yes	yes	diabetes	27	high
52	0	11	yes	no	cholesterol	29	high
23	0	3	no	no	no	23	low

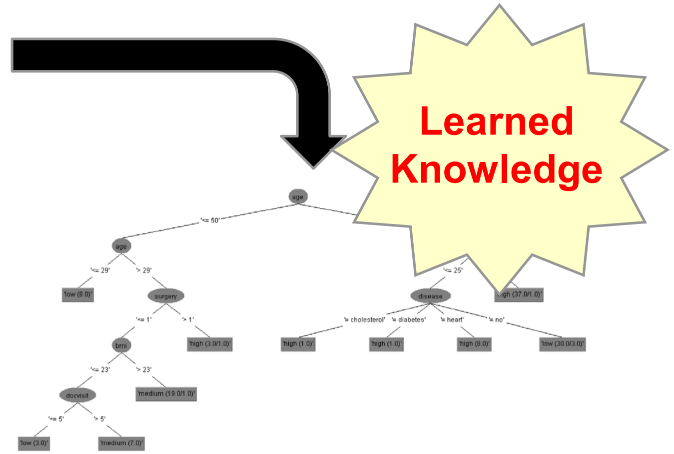

  - Engineered knowledge: eligibility and compliance

Applicants from Switzerland are eligible.
A person younger than 21 year is not able to apply
…

Engineered Knowledge

Combining Machine Learning and Knowledge Engineering for Eligibility Decisions (2/2)

Examples of learned rules:

risk (Person, high) :- age(Person,A), A > 50,
bmi(Person, Bmi), Bmi =<25,
disease(Person, diabetes).
risk (Person, low) :- age(Person,A), A =< 29.

Learned Knowledge

Examples of engineered rules:

eligible(Person, no) :- age(Person,A), A =< 21.
eligible(Person,no) :- country(Person,C), C =/= switzerland.

Engineered Knowledge

Combining engineered and learned rules:

accept(Person, yes) :- eligible(Person, yes), risk(Person, low).
accept(Person, yes) :- eligible(Person, yes), risk(Person, medium).
accept(Person, no) :- eligible(Person, no).
accept(Person, no) :- risk(Person, high)

Engineered Knowledge

Summary: Creating Knowledge Bases

- **Knowledge Engineering:** Human experts build knowledge base
  - For knowledge we are aware of
  - For knowledge that must be correct (e.g. compliance rules)
  - Inferences are explainable (trust)
- **Machine Learning:** automatic creation of knowledge from example data
  - Can solve complex tasks for which
    - knowledge is not known
    - knowledge is tacit
  - Reliance on real-world data instead of pure intuition
  - Requires large sets of data
  - Can adapt to new situations (collect more data)