Artificial Intelligence and Chatbots

Knut Hinkelmann

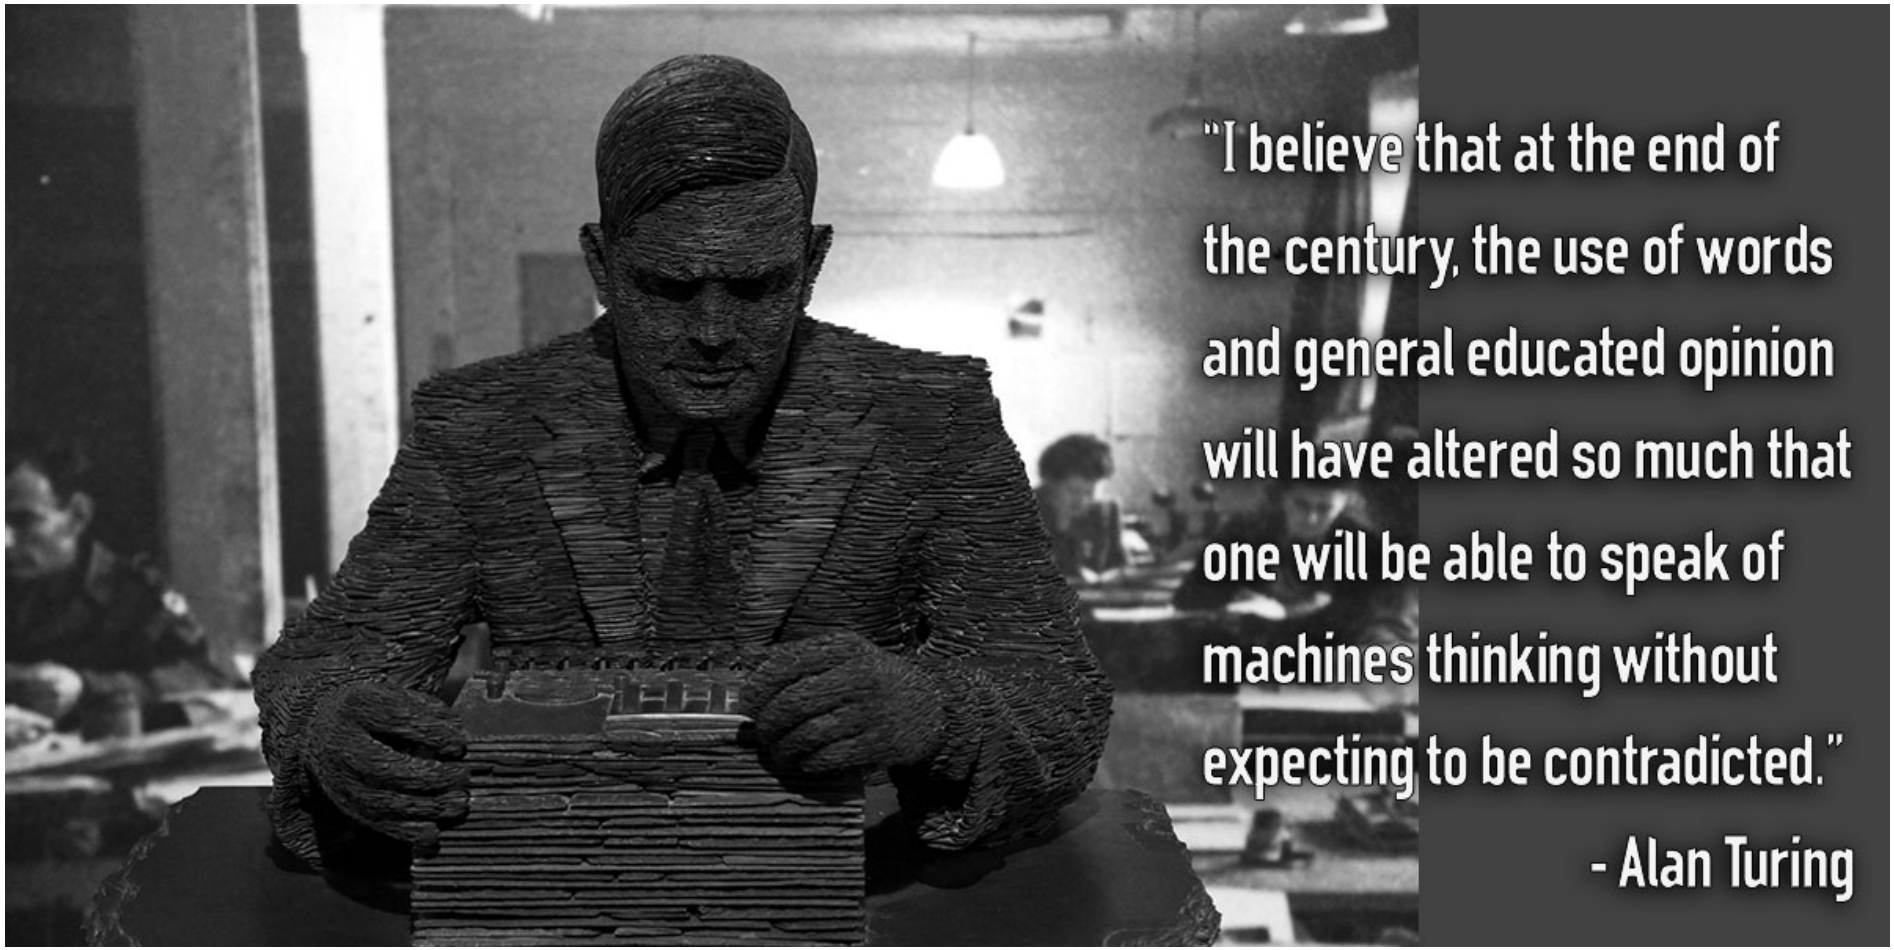
"I believe that at the end of the century, the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted."
- Alan Turing

Turing Test

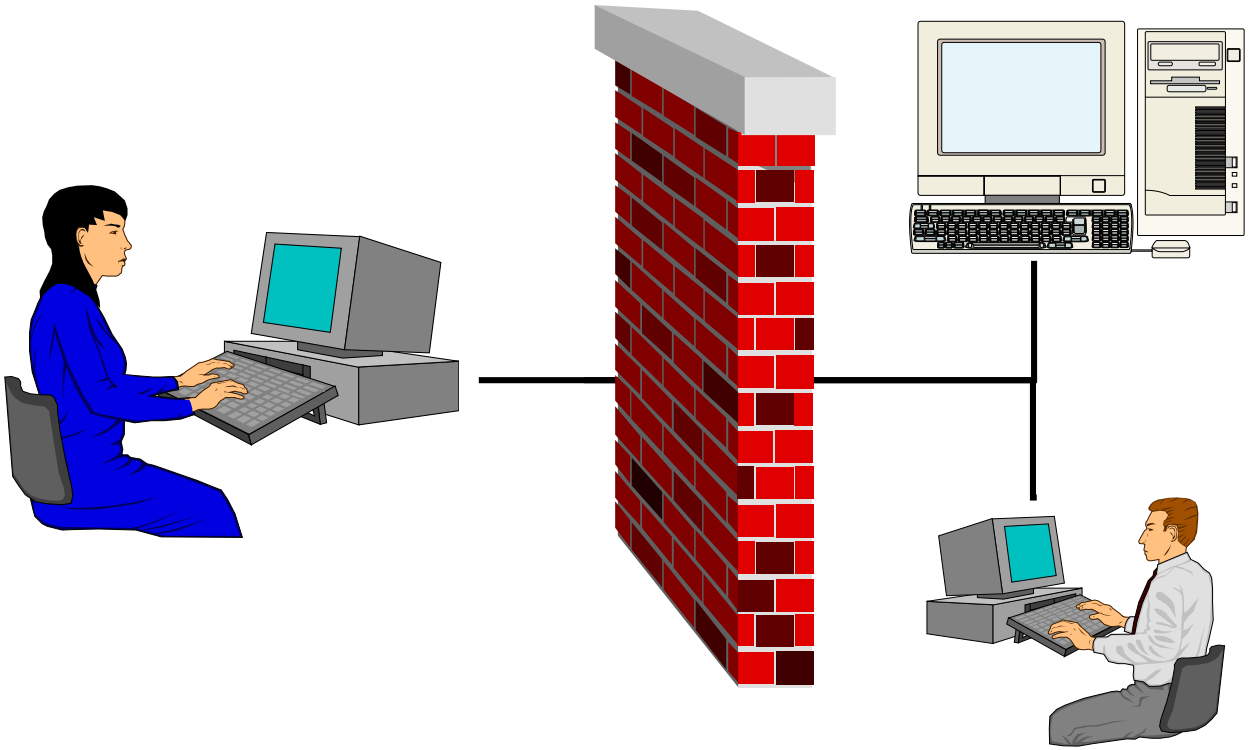

Applications and Methods of AI

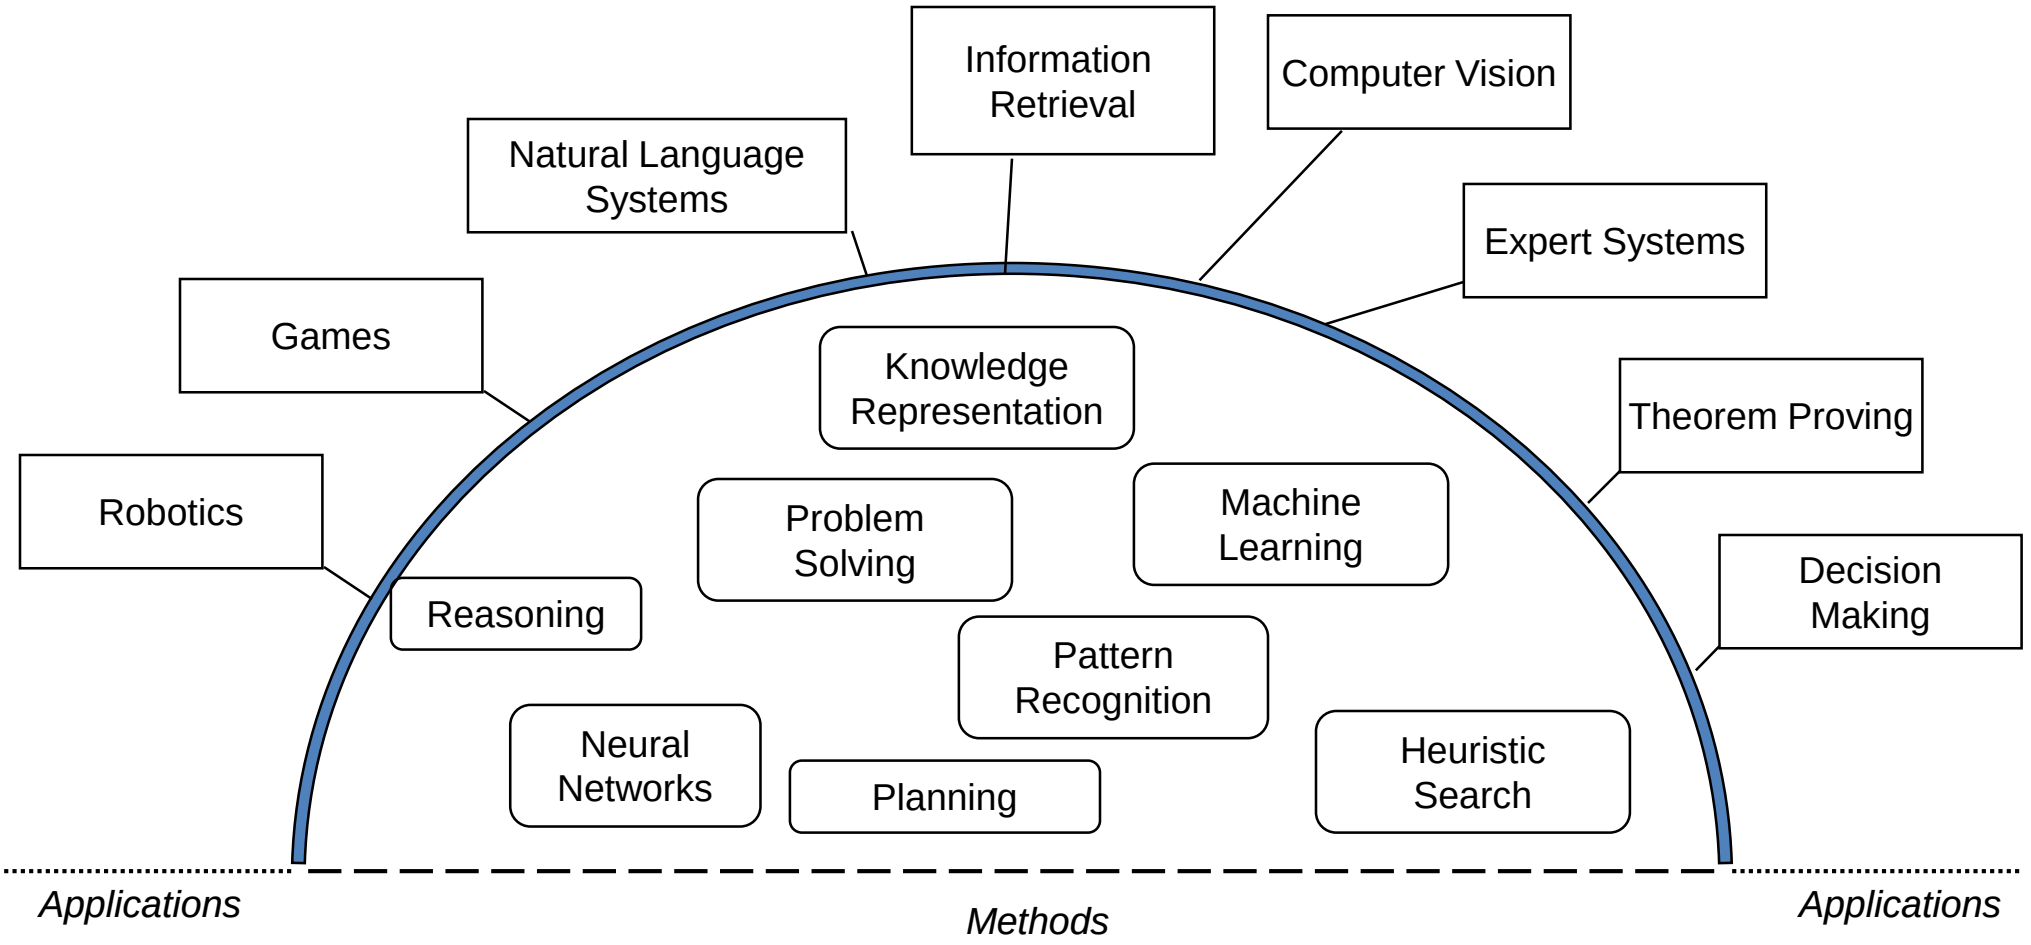

Symbolic AI

Knowledge-Based Systems (Expert Systems)

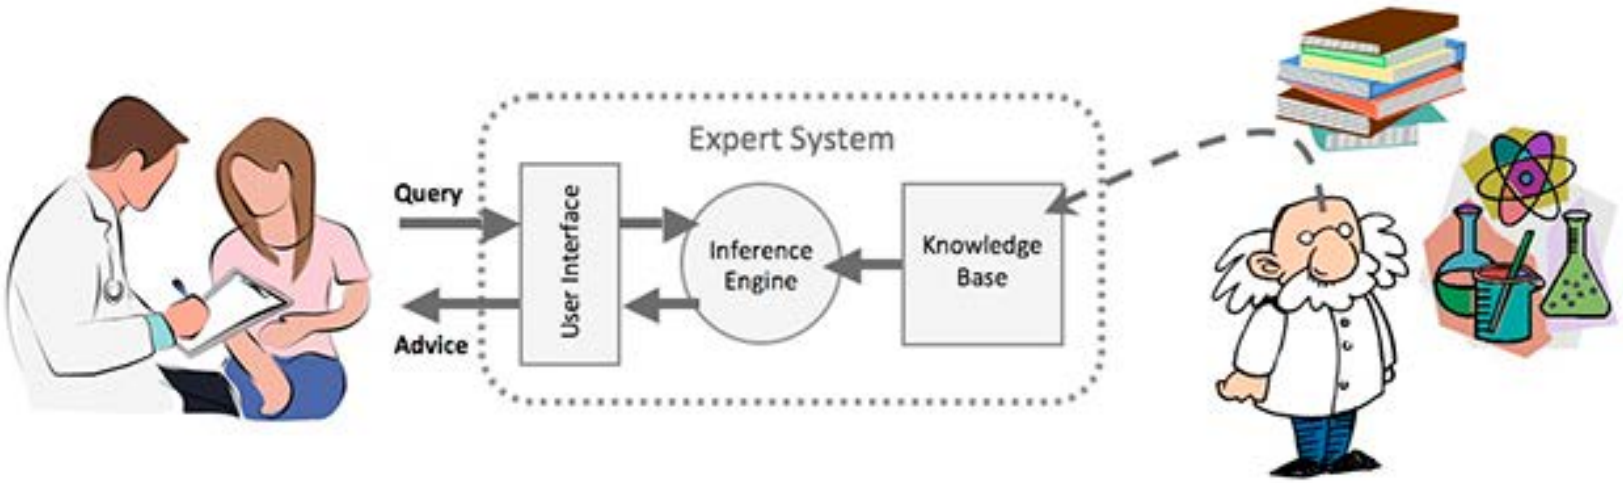

Knowledge Sources in a Knowledge Base

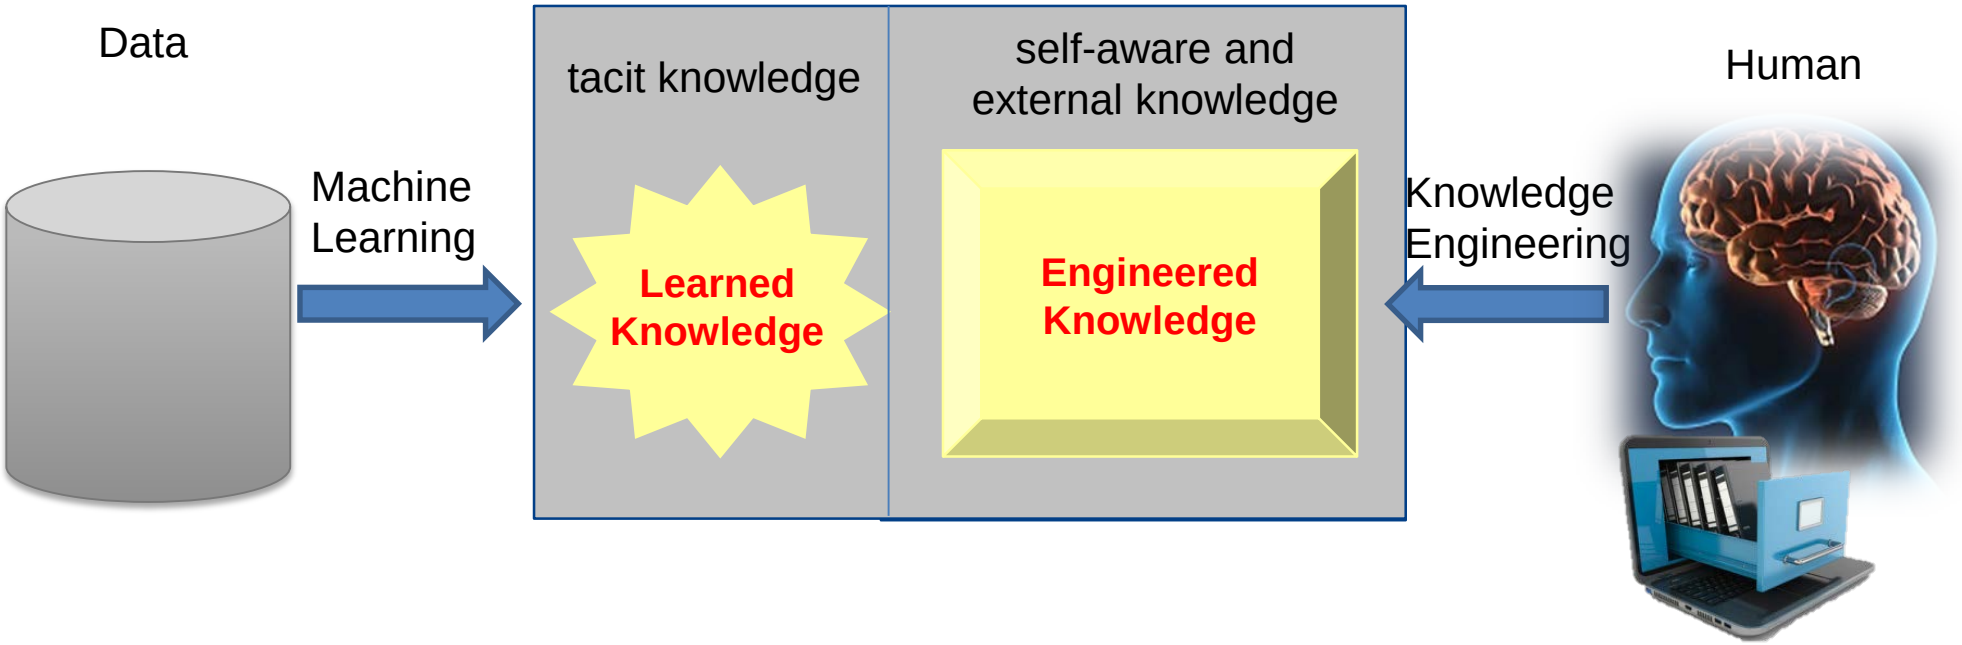

Knowledge

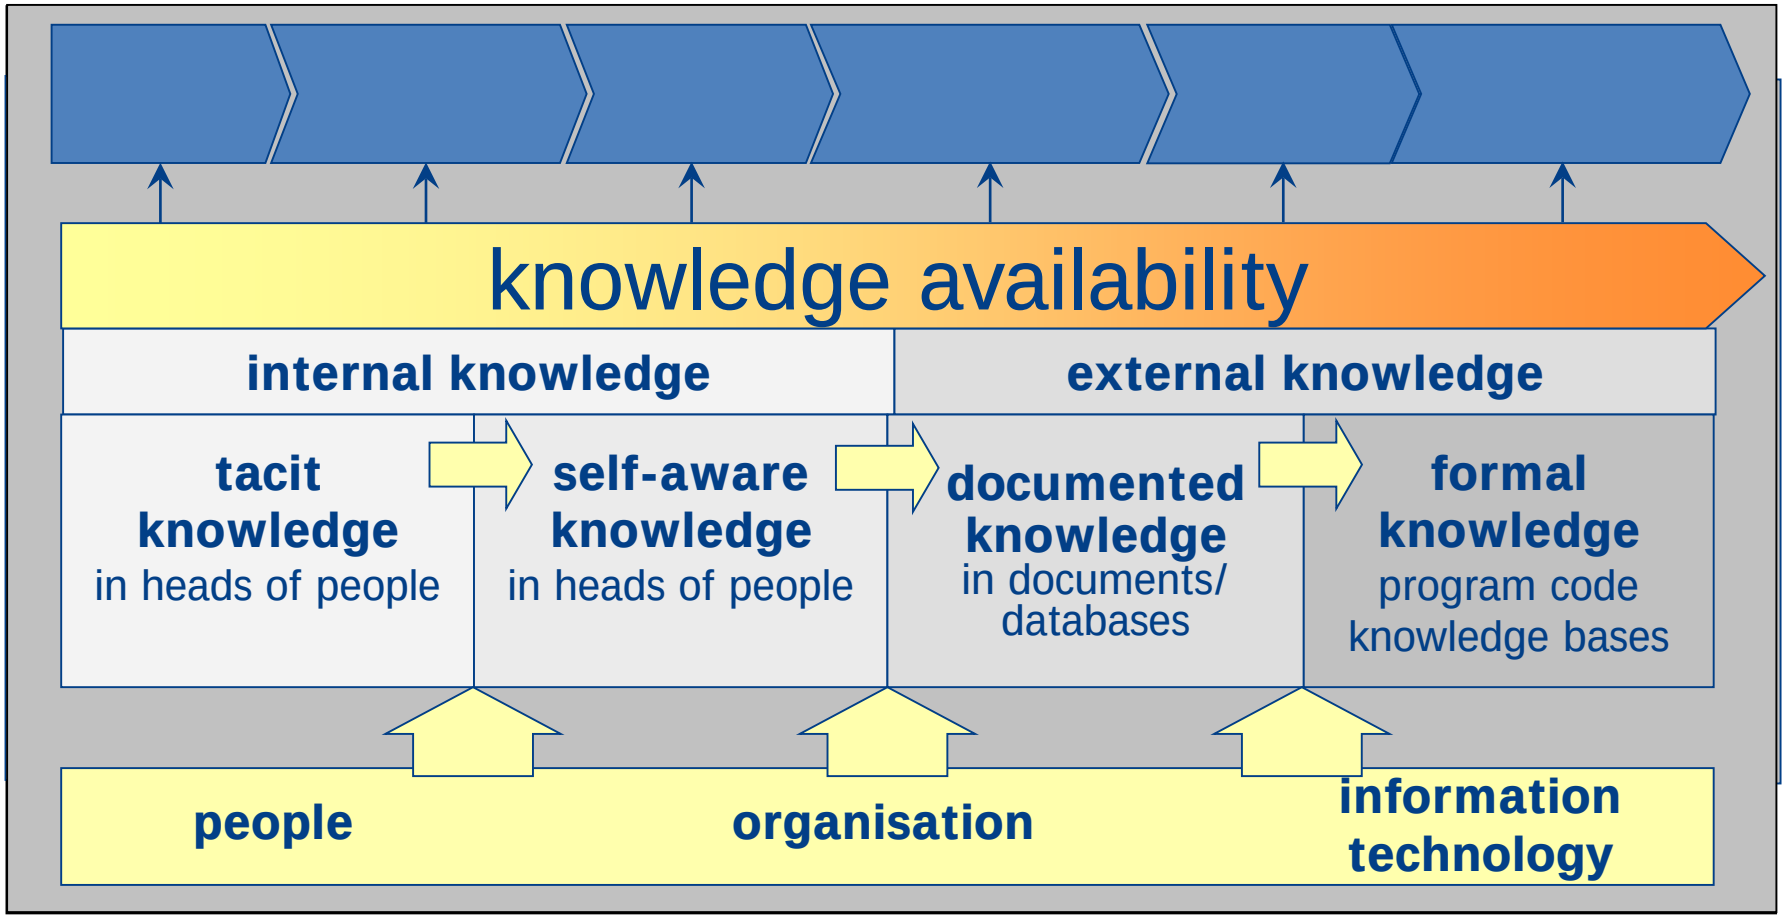
knowledge availability
internal knowledge
external knowledge
tacit knowledge
in heads of people
self-aware knowledge
in heads of people
documented knowledge
in documents/ databases
formal knowledge
program code
knowledge bases
people
organisation
information technology

Knowledge Engineering

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***

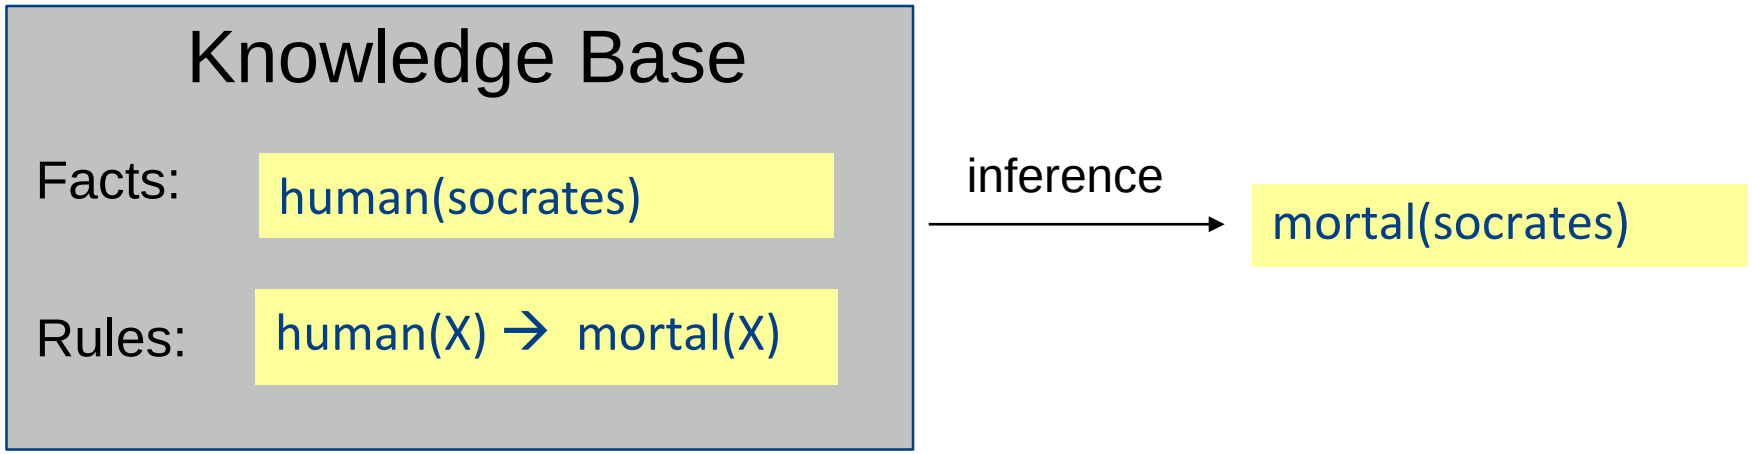

Inference

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)

Domain Concepts, Instances and Relations

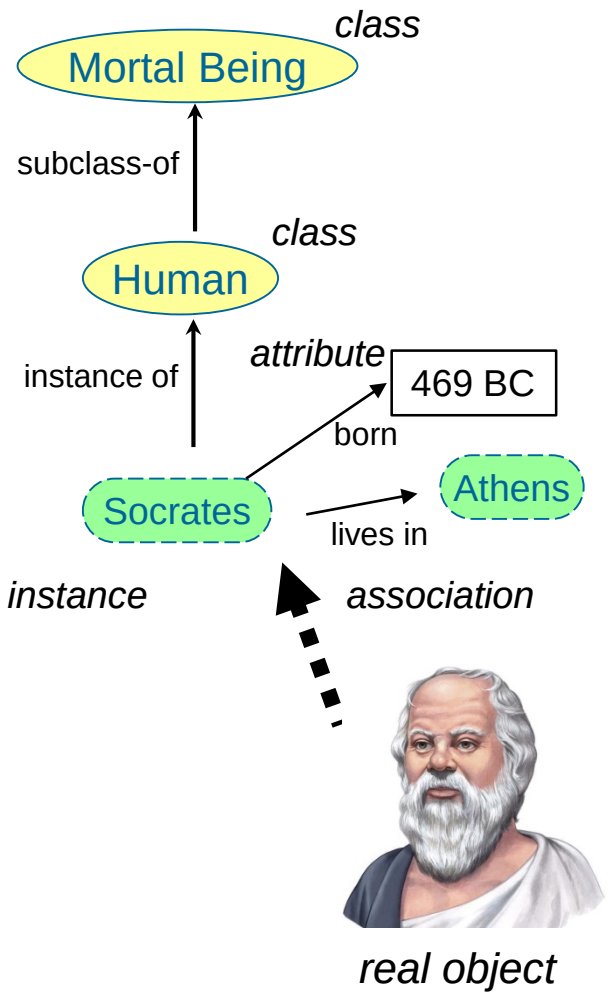

- There are two kinds of concepts:
  - **classes**
  - **instances (facts)**
- There are different kinds of relations
  - **generalisation** ("is a")
    - between classes **(subclass of)**
    - between instance and class **(instance of)**
  - **associations/properties**
    - any other kind of relationship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

Machine Learning:
Make Knowledge explicit with the Use of Data

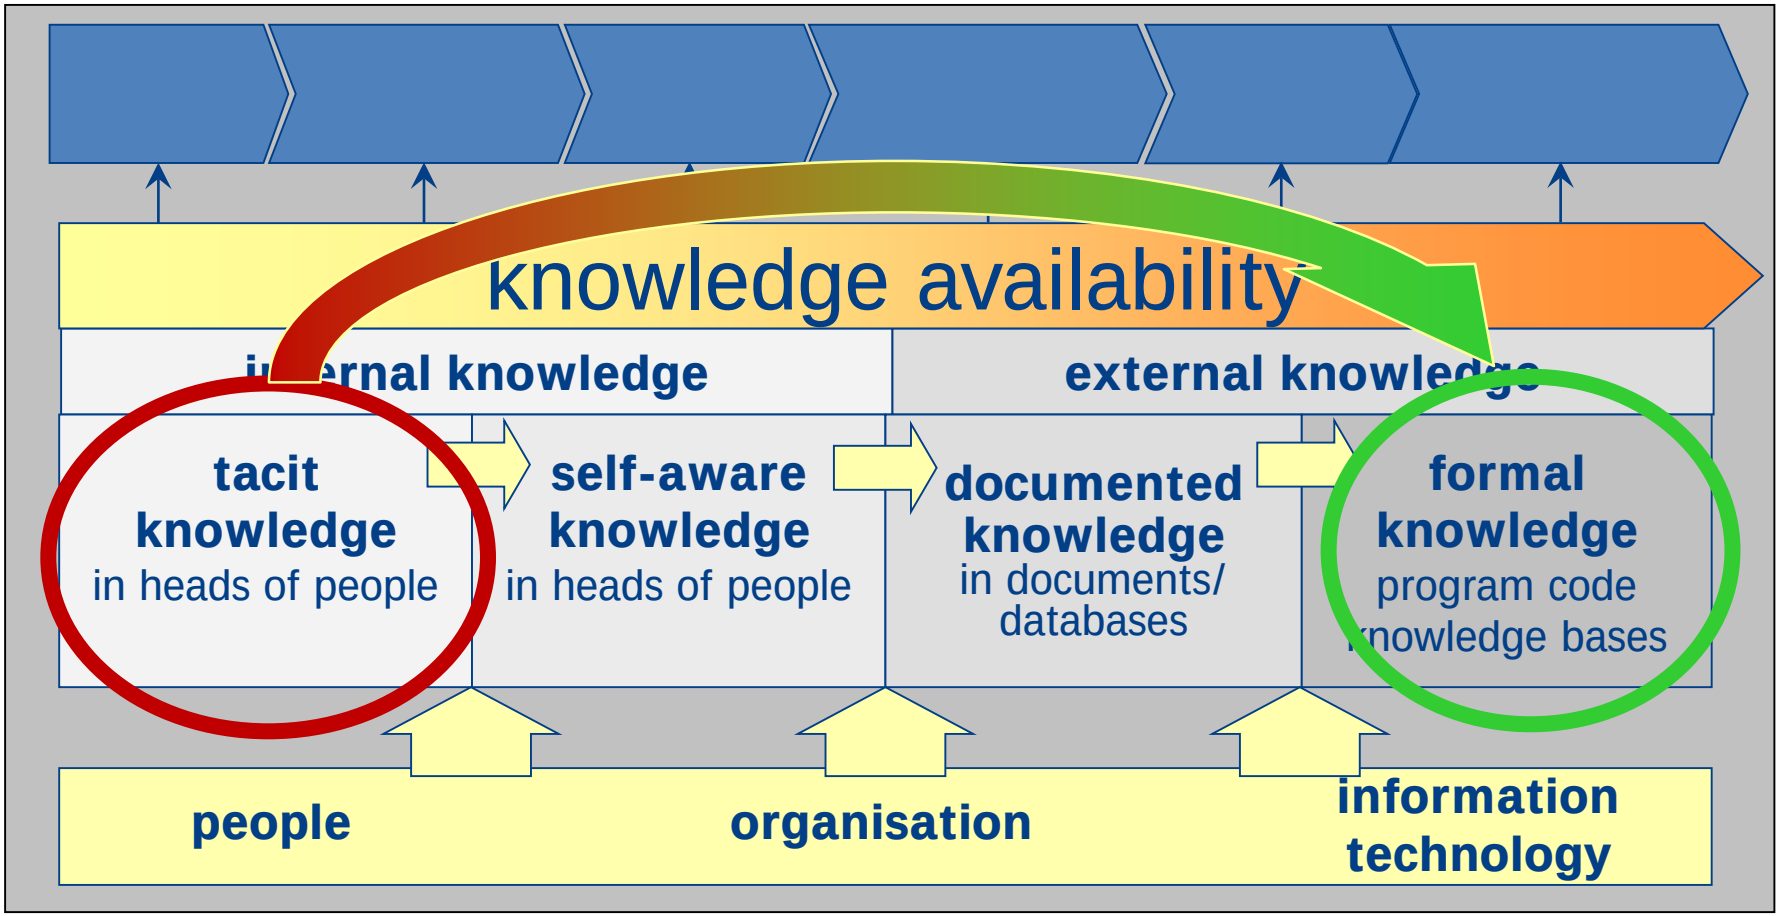



Self-driving Cars

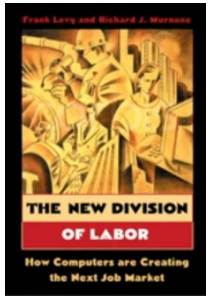

"… it is hard to imagine discovering the set of rules that can replicate the driver's behavior."

(Levy & Murnane 2006)

Face Recognition

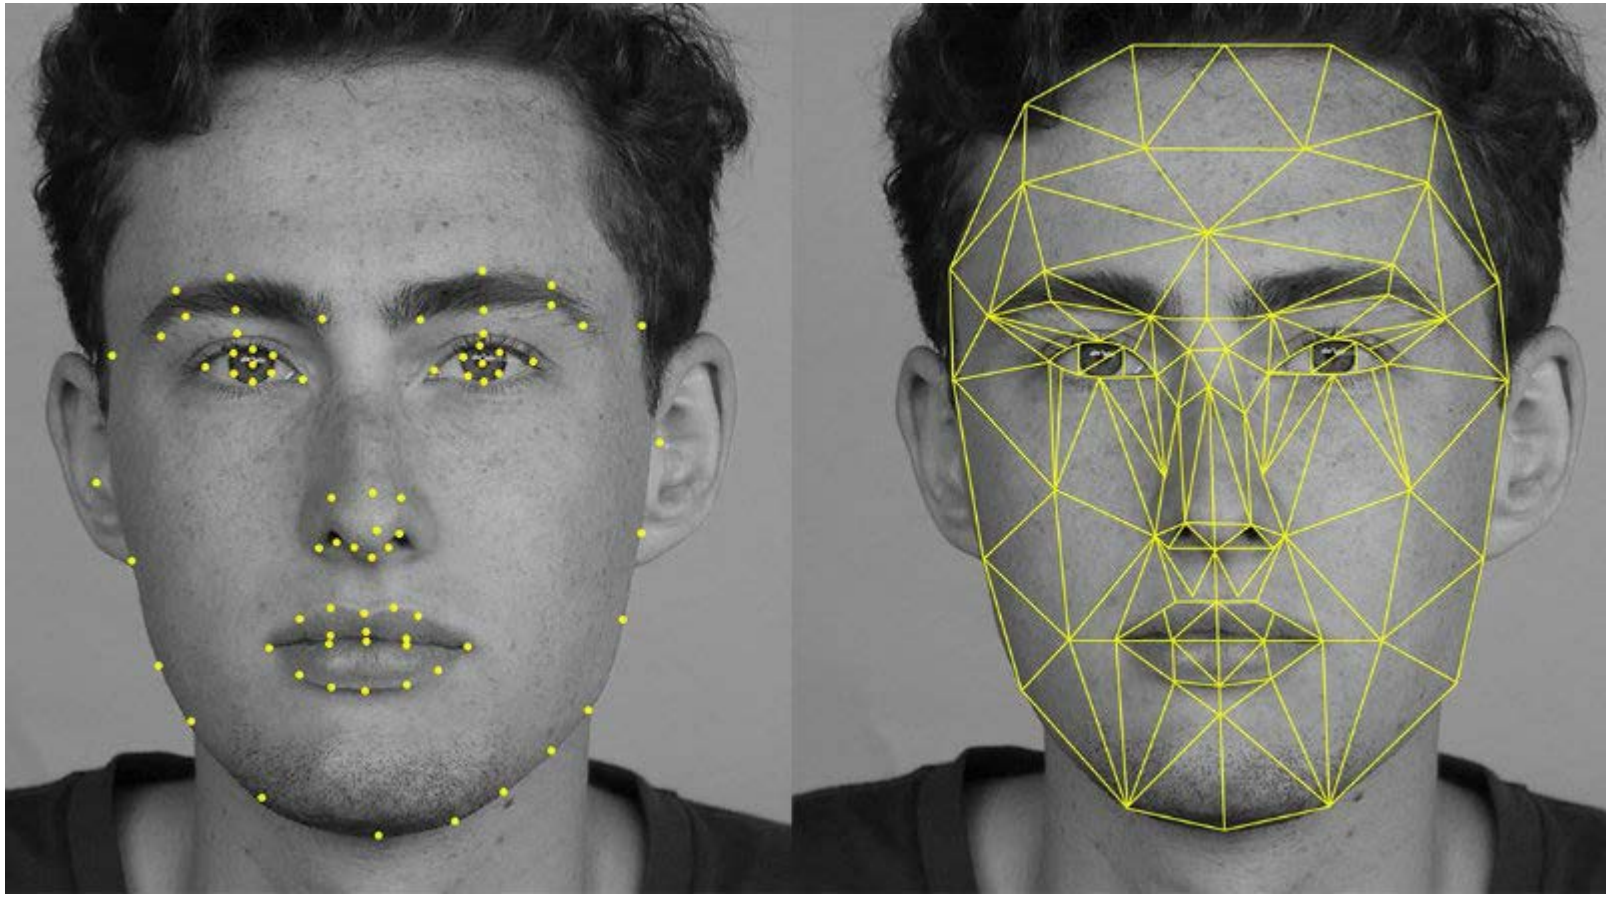

Spam Filter

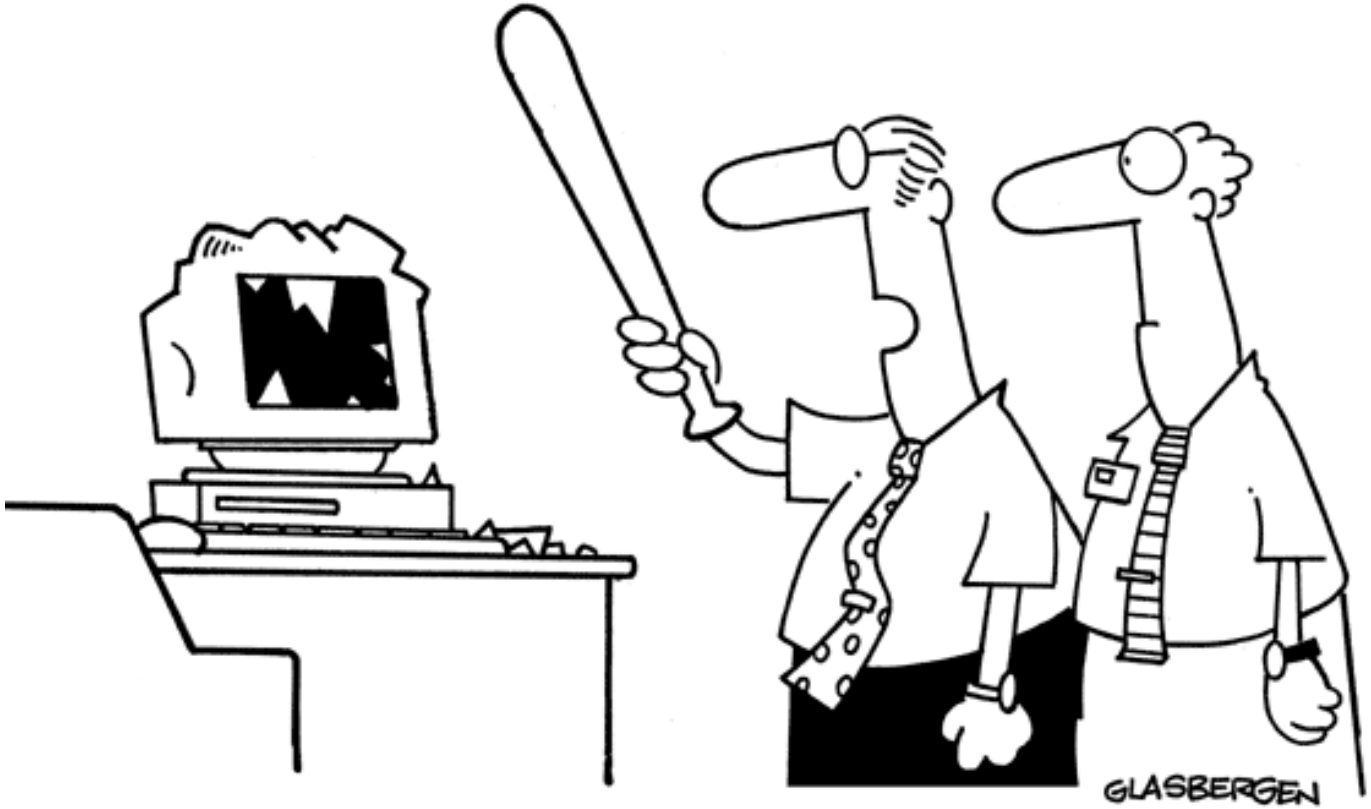

"It's not the most sophisticated Spam blocker
I've tried, but it's the only one that works!"

Machine Learning: General Idea

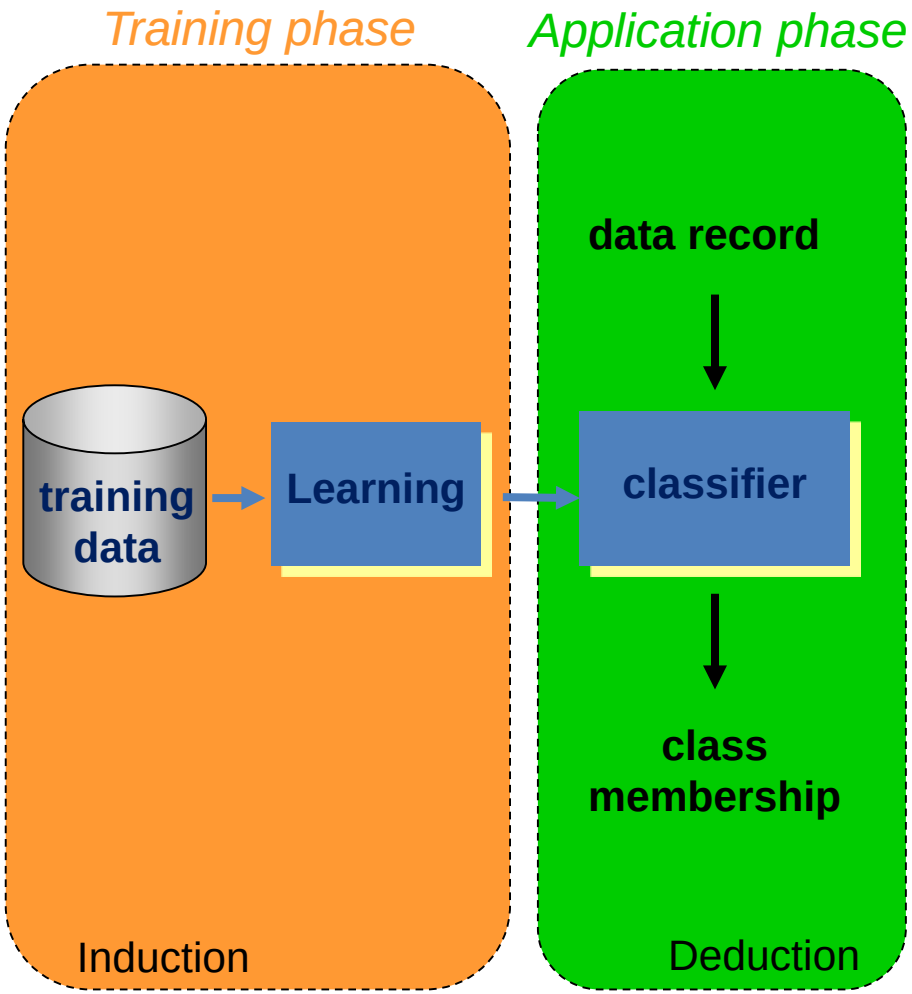

- **Training**
  - Collect data for the problem
  - Use the data to learn how to solve the type of problem
  - Result: Knowledge base
- **Application**
  - Use the learned knowledge for new problems

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

Example: Classification

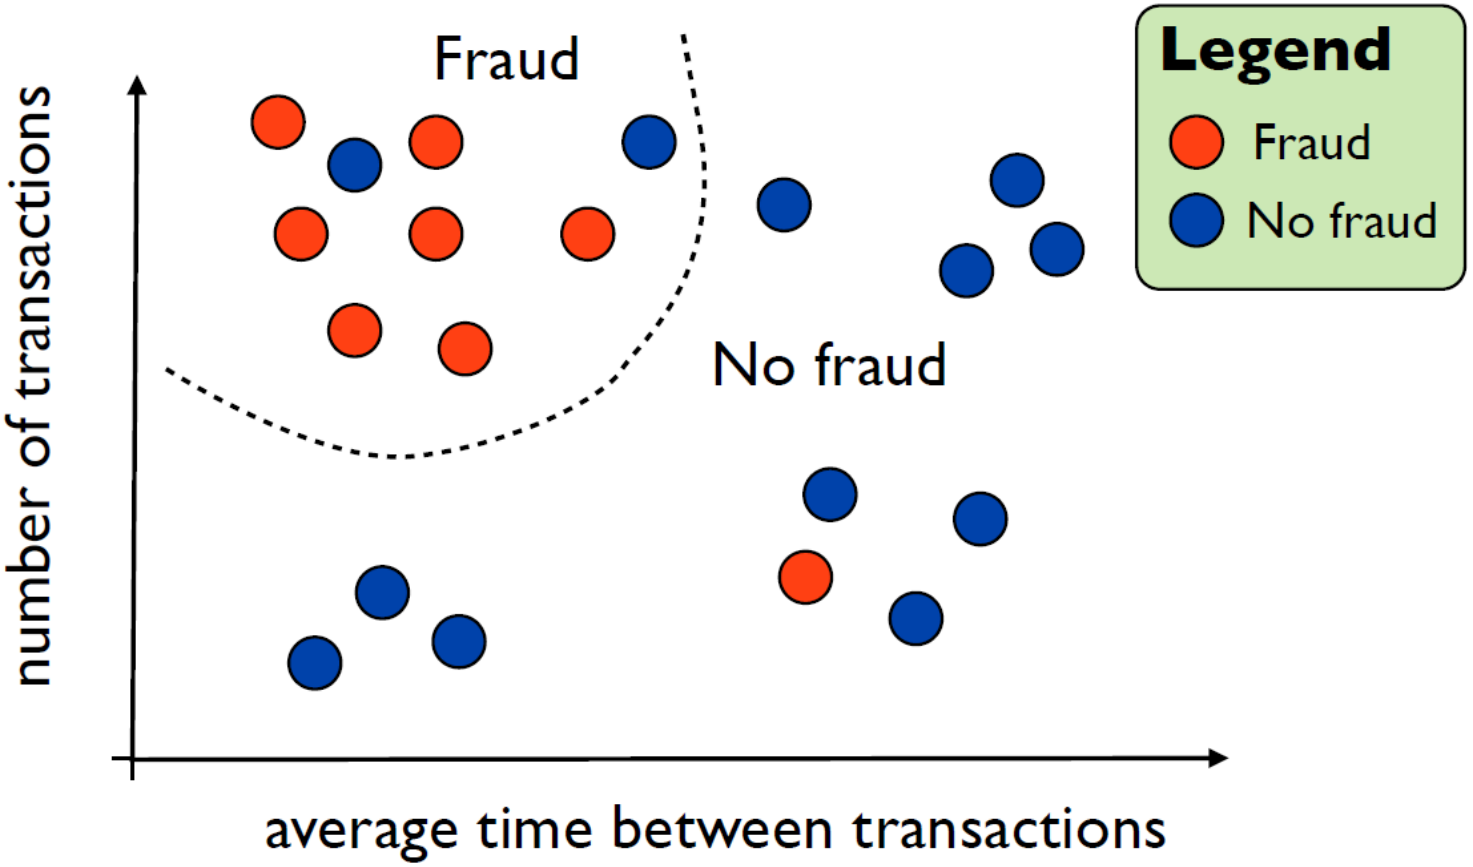

(Lison, 2012)

Input and Output for Machine Learning

	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)

Classification Criteria

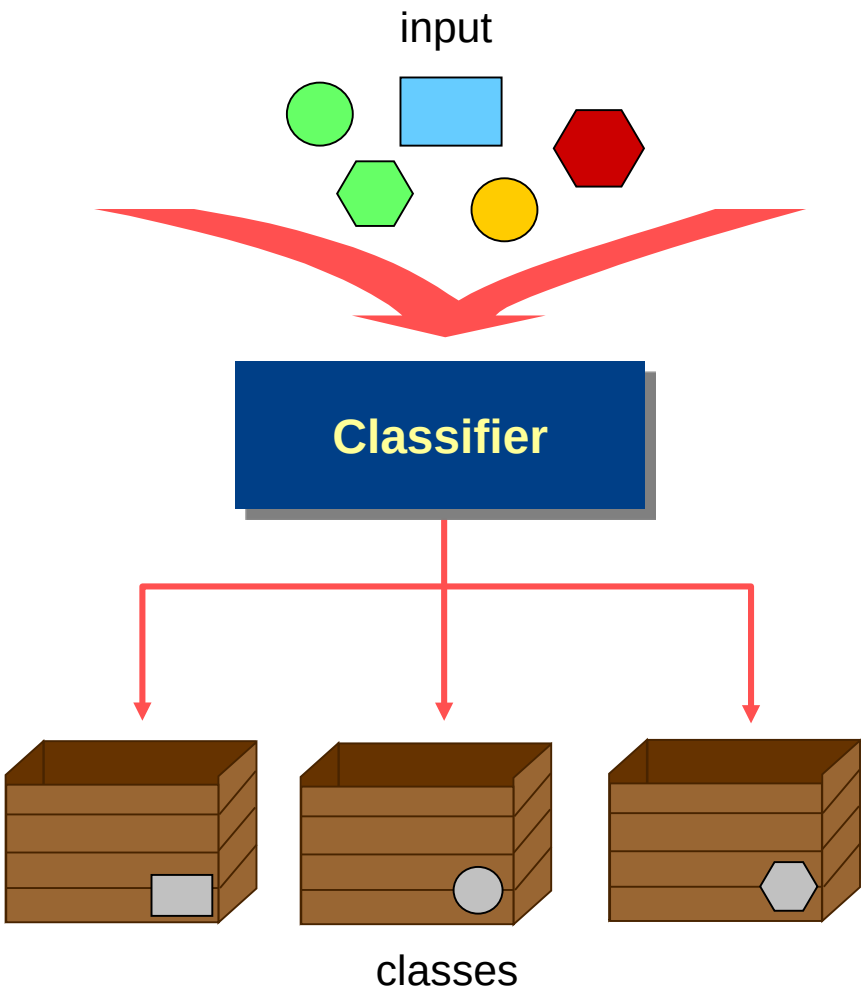

- The classifier decides, which individual belongs to which class
- Problem:
  - The criteria for the decision are not always obvious
- Supervised Learning:
  - Learn the classification criteria from known examples

Machine Learning: General Idea

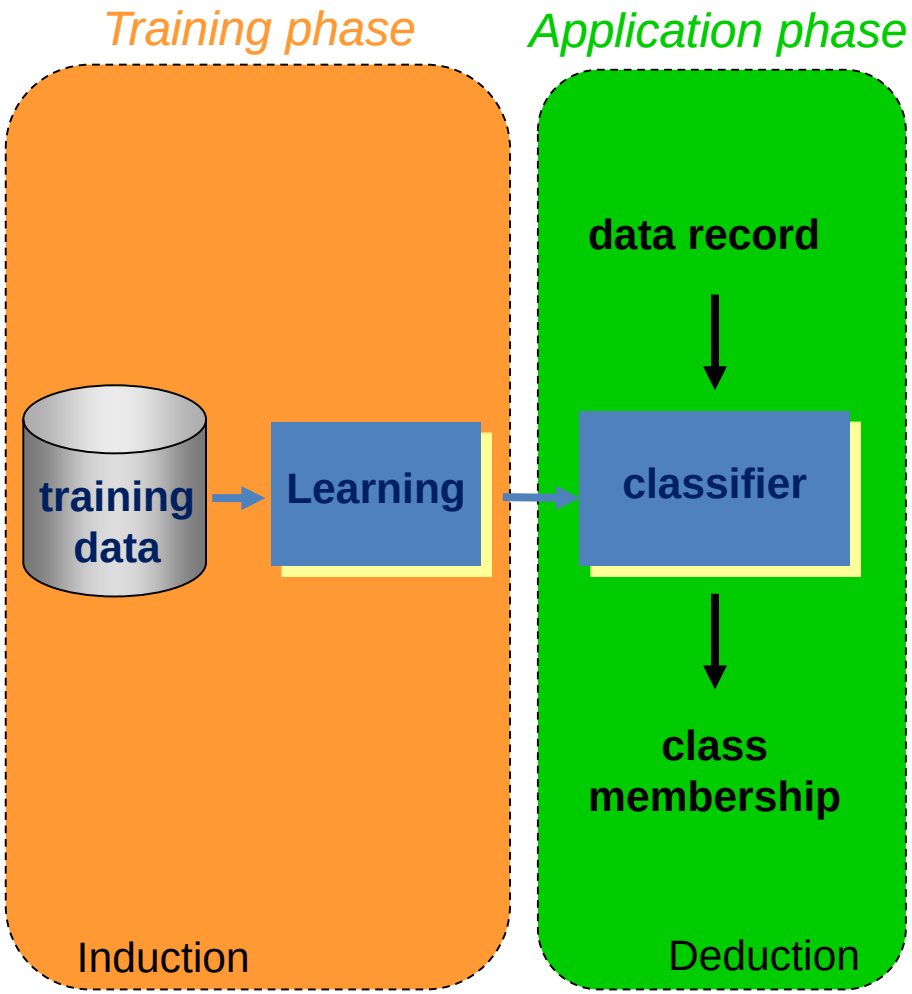

- **Training:** Learning the classification criteria
  - Given: sample set of training data records
  - Result: Decision logic to determine class from values of input attributes (decision tree, rules, model)
- **Application:** Classification
  - Goal: assign a class to previously unseen records of input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

Classification Methods

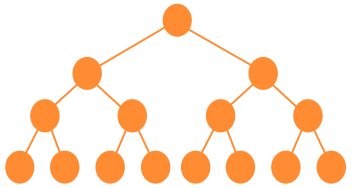

Decision Trees

IF ...
THEN ...

Rules

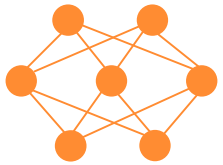

Neural Network

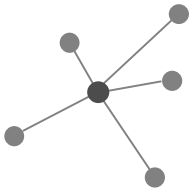

k-Nearest Neighbor

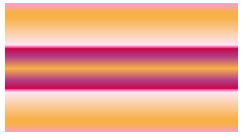

Genetic Algorithms

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

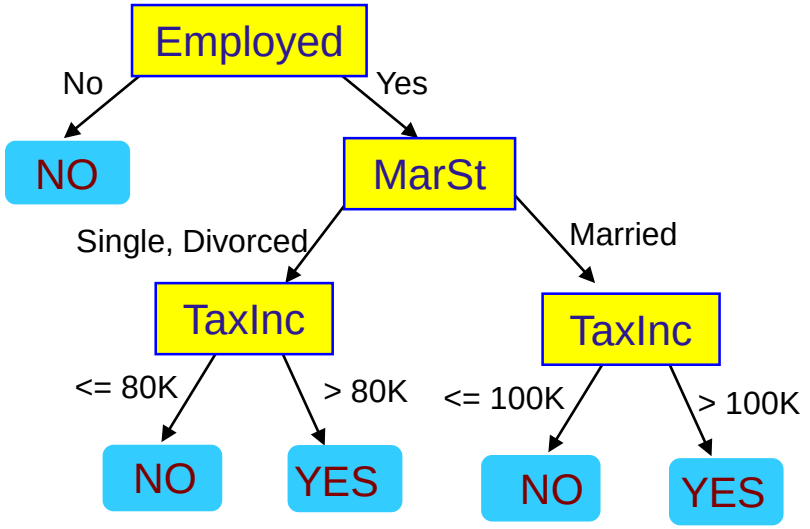

Credit Worthiness	Employed	Marital Status	Taxable Income	Accept
	Yes, No	*Single, Divorced, Married*	*Integer*	*Yes, No*
1	No			No
2	Yes	Single	> 80K	Yes
3	Yes	Divorced	> 80K	Yes
4	Yes	Single	≤ 80K	No
5	Yes	Divorced	≤ 80K	No
6	Yes	Married	> 100K	Yes
7	Yes	Married	≤ 100K	No

Model: Decision Tree/Table

Unsupervised Learning

Example: Clustering (= identify new classes)

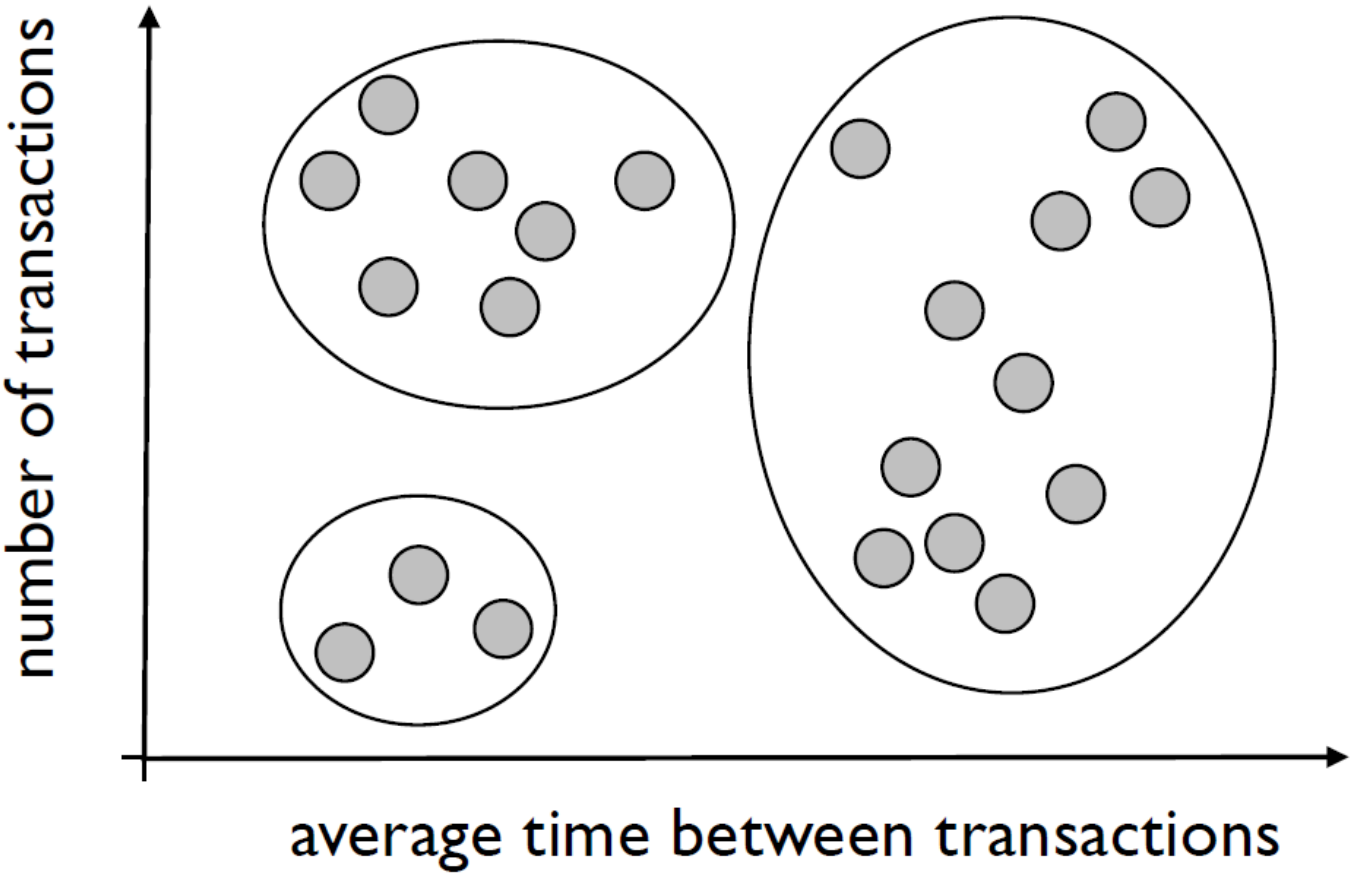

(Lison, 2012)

Unsupervised Learning

- Sometimes, we don't have access to any output value **o**, we simply have a collection of input examples **i**
- In this case, what we try to do is to learn the underlying patterns of our data
  - are there any *correlations* between features?
  - can we *cluster* our data set in a few groups which behave similarly, and detect *outliers*?

(Lison, 2012)

Example: Recommender Systems

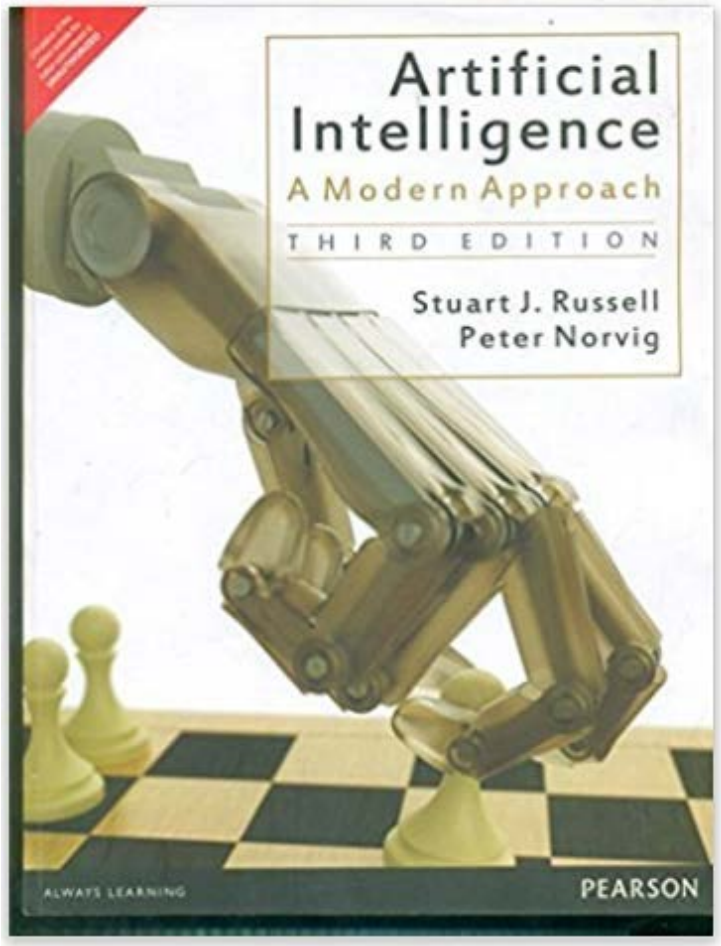

Customers who bought this item also bought

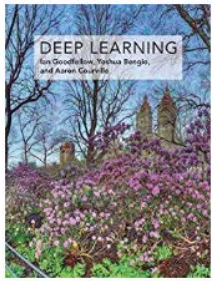

Deep Learning (Adaptive Computation and Machine Learning series)
› Ian Goodfellow

Hands-On Machine Learning with Scikit-Learn and TensorFlow:...
› Aurélien Géron

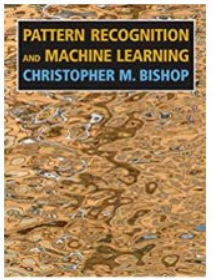

Pattern Recognition and Machine Learning (Information Science...
› Christopher M. Bishop

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

Subsymbolic AI

Biological Neuron

The basic computational unit of the brain is a **neuron**. Approximately 86 billion neurons can be found in the human nervous system and they are connected with approximately 10^{14}—10^{15} **synapses**.

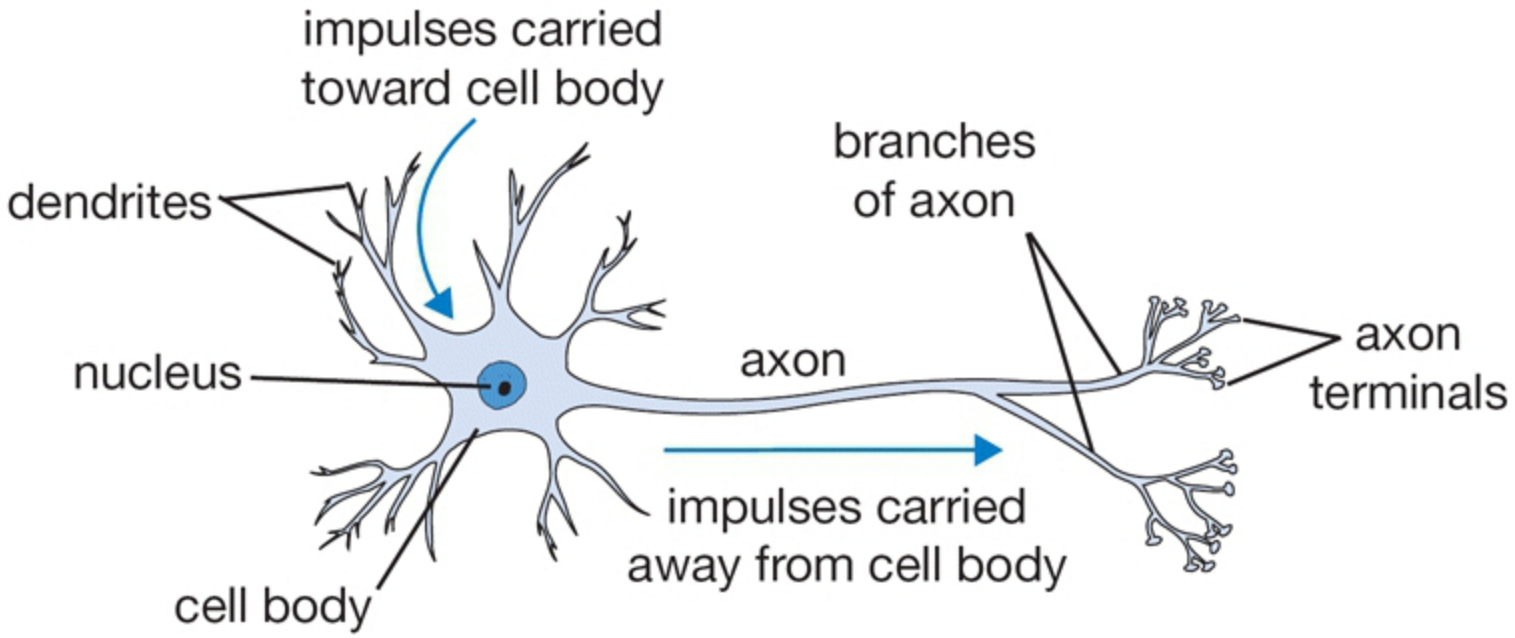

Source: David Fumo: https://towardsdatascience.com/a-gentle-introduction-to-neural-networks-series-part-1-2b90b877

Mathematical Model of a Neuron

The basic unit of computation in a neural network is the neuron (often called a node or unit). It receives input from some other nodes, or from an external source and computes an output. Each input has an associated weight (w), which is assigned on the basis of its relative importance to other inputs. The node applies a function to the weighted sum of its inputs.

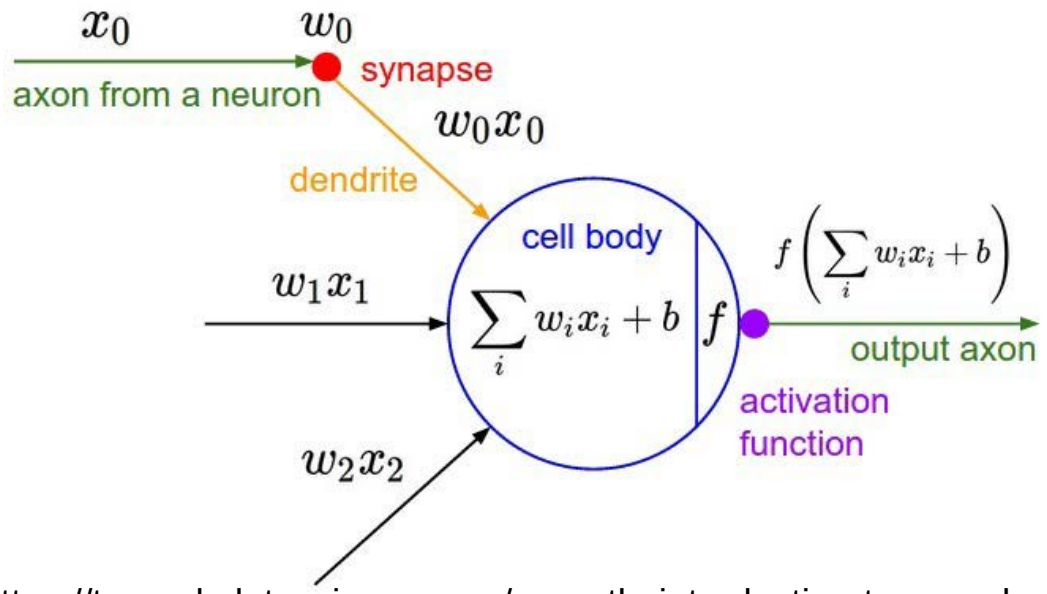

Source: David Fumo: https://towardsdatascience.com/a-gentle-introduction-to-neural-networks-series-part-1-2b90b877

“Backprops” (Neural Networks)

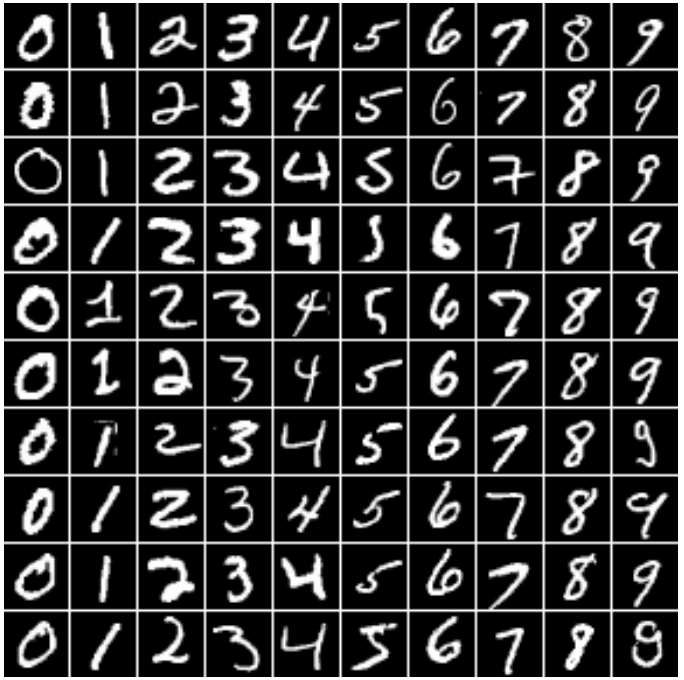

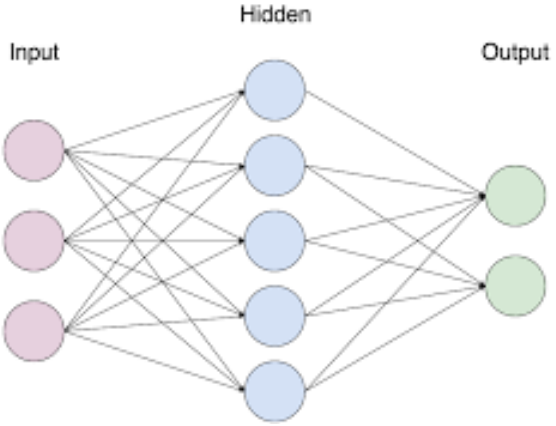

- Classify the numbers

Recognizing Numbers

- It is very hard to specify what makes a «2»

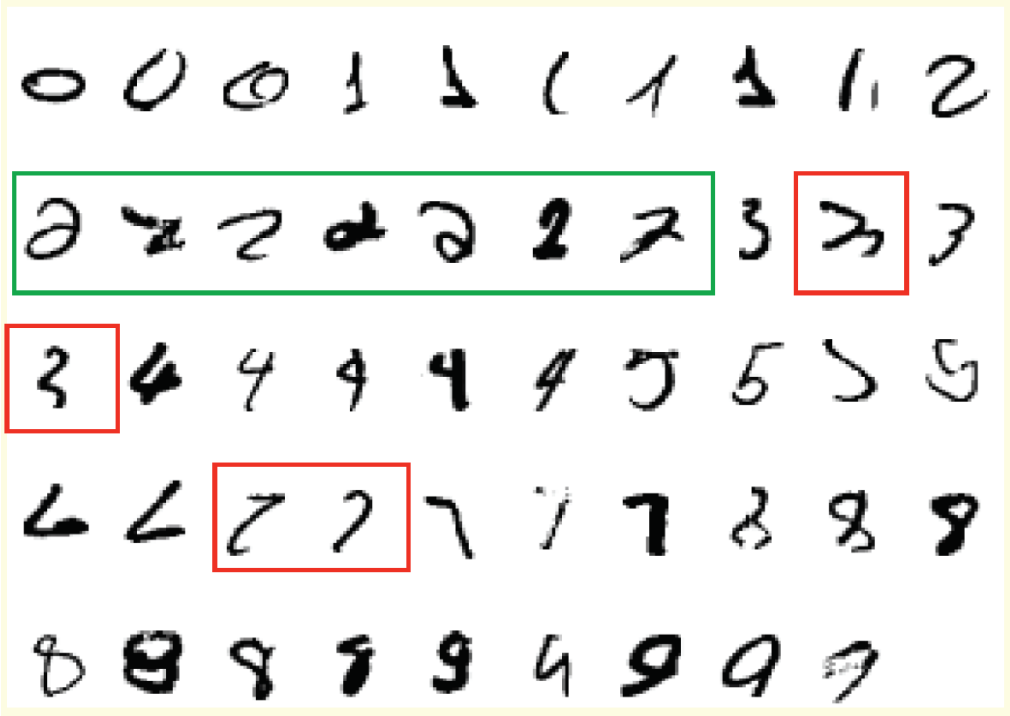

Source: Geoffrey Hinton, https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec1.pdf

Multi-layer Perceptrion

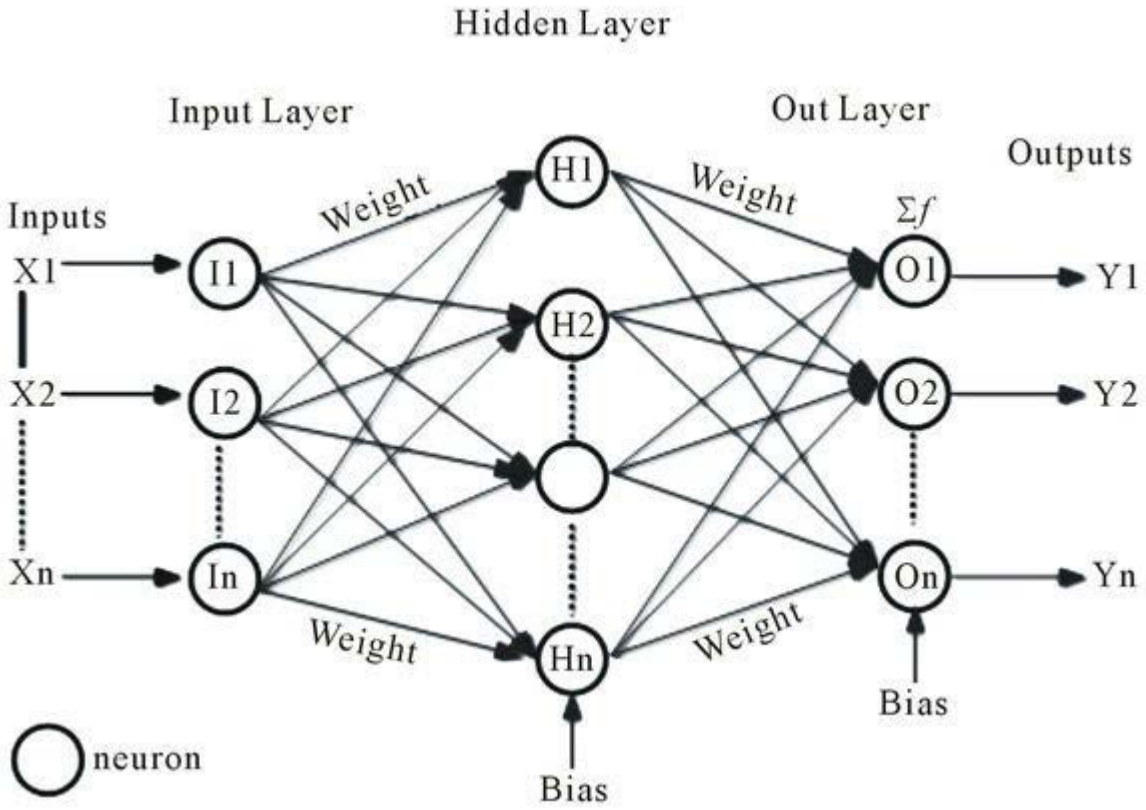

Source: David Fumo: https://towardsdatascience.com/a-gentle-introduction-to-neural-networks-series-part-1-2b90b877

Learning: Backpropagation of Errors

- The supervisor ***corrects*** the network whenever it makes mistakes (supervised learning)

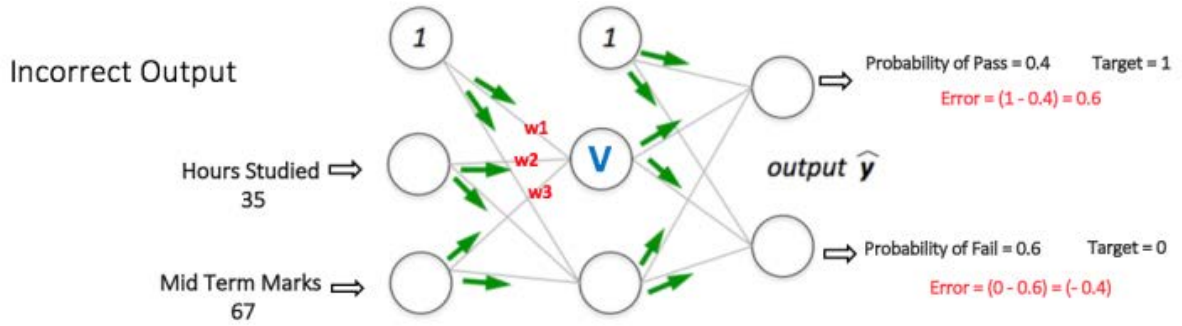

- Total error at the output nodes are propagate back through the network leading to new weights

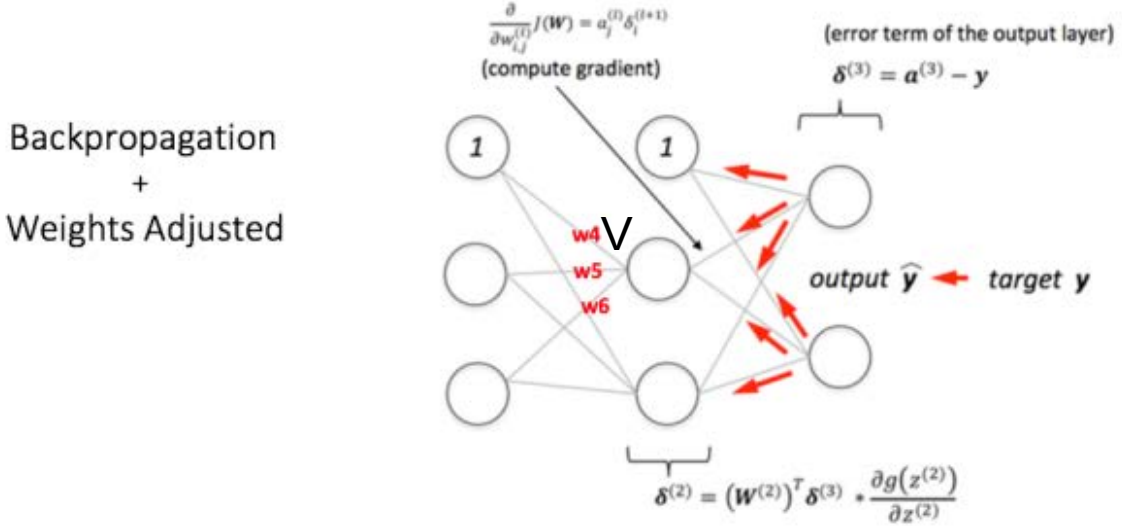

https://ujjwalkarn.me/2016/08/09/quick-intro-neural-networks/

Deep Neural Networks

"Non-deep" feedforward neural network

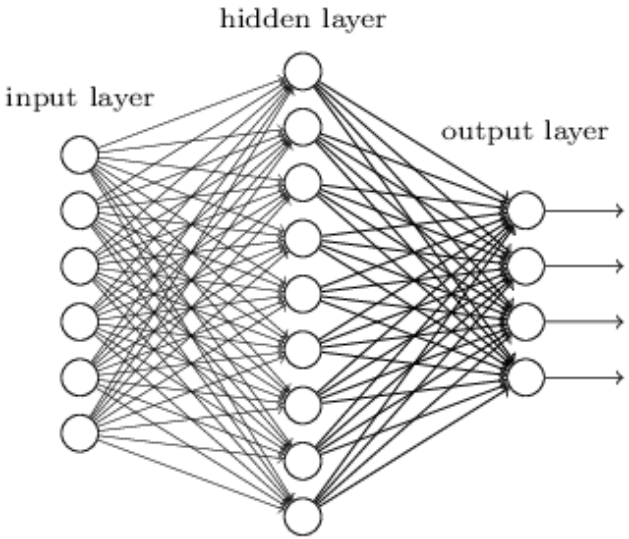

Deep neural network

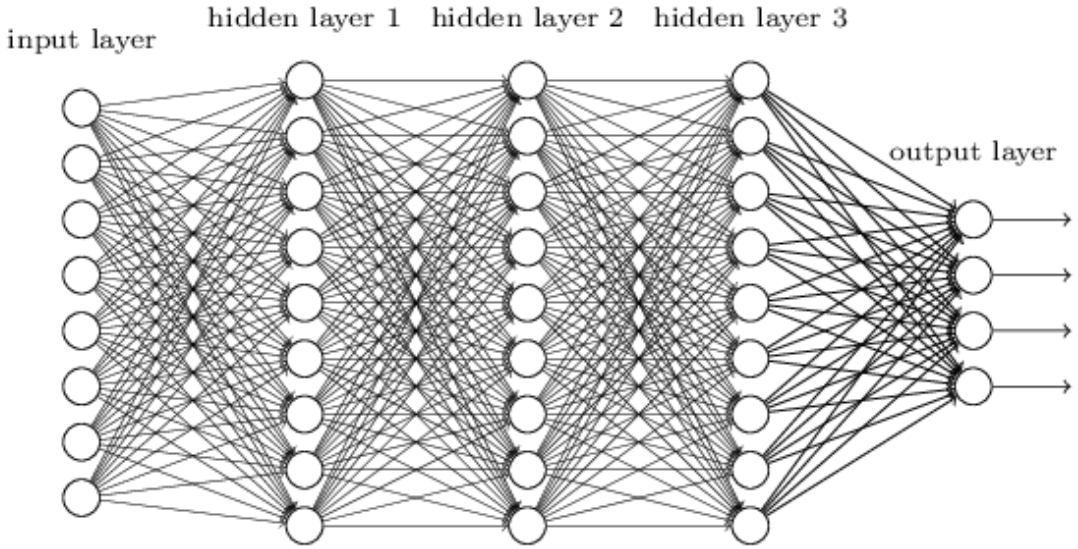

Applications of Deep Learners

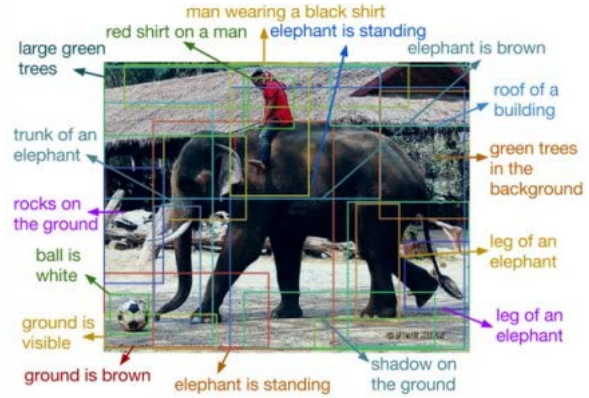

Describing photos

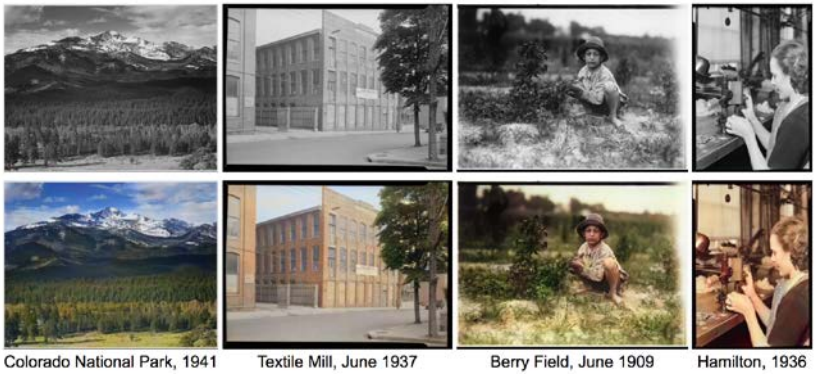

Picture coloring

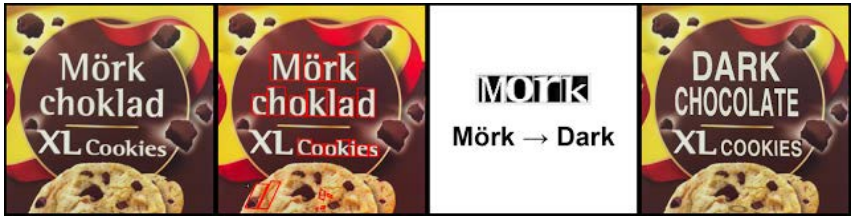

Translation

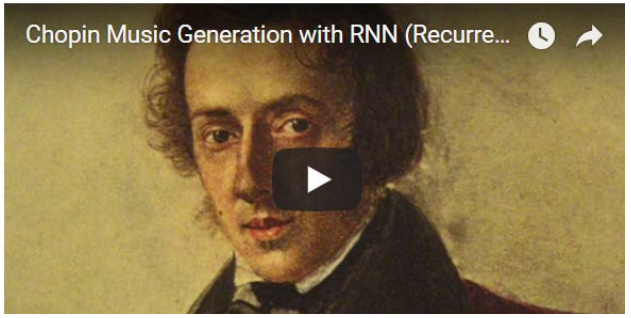

Musik composition

Self-driving Cars

© http://www.yaronhadad.com/deep-learning-most-amazing-applications/ [28.02.2018]

Neural Networks vs. Decision Trees/Rules

- Classification with Neural Networks is very good
- Decisions with neural networks are not comprehensible
- Decision Trees and Rules are often more trustworthy
  - Decisions with trees and rules are comprehensible and explainable
  - One can see which rules are applied to make a decision
- In applications, in which trust in the decision or explainability is important, people prefer decision trees or rules

Summary: Creating Knowledge Bases

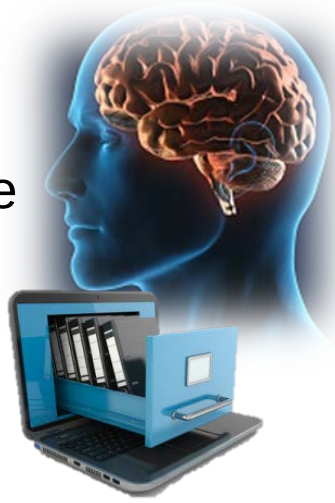

- **Knowledge Engineering:** Human experts build knowledge base
  - For knowledge we are aware of
  - For knowledge we need to be sure of (e.g. compliance rules)
  - Inferences are explainable (trust)

- **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)

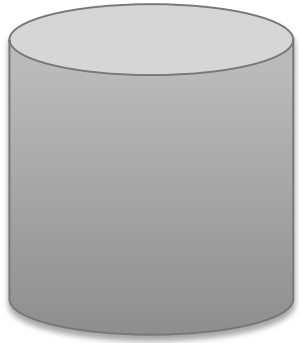

Chatbots

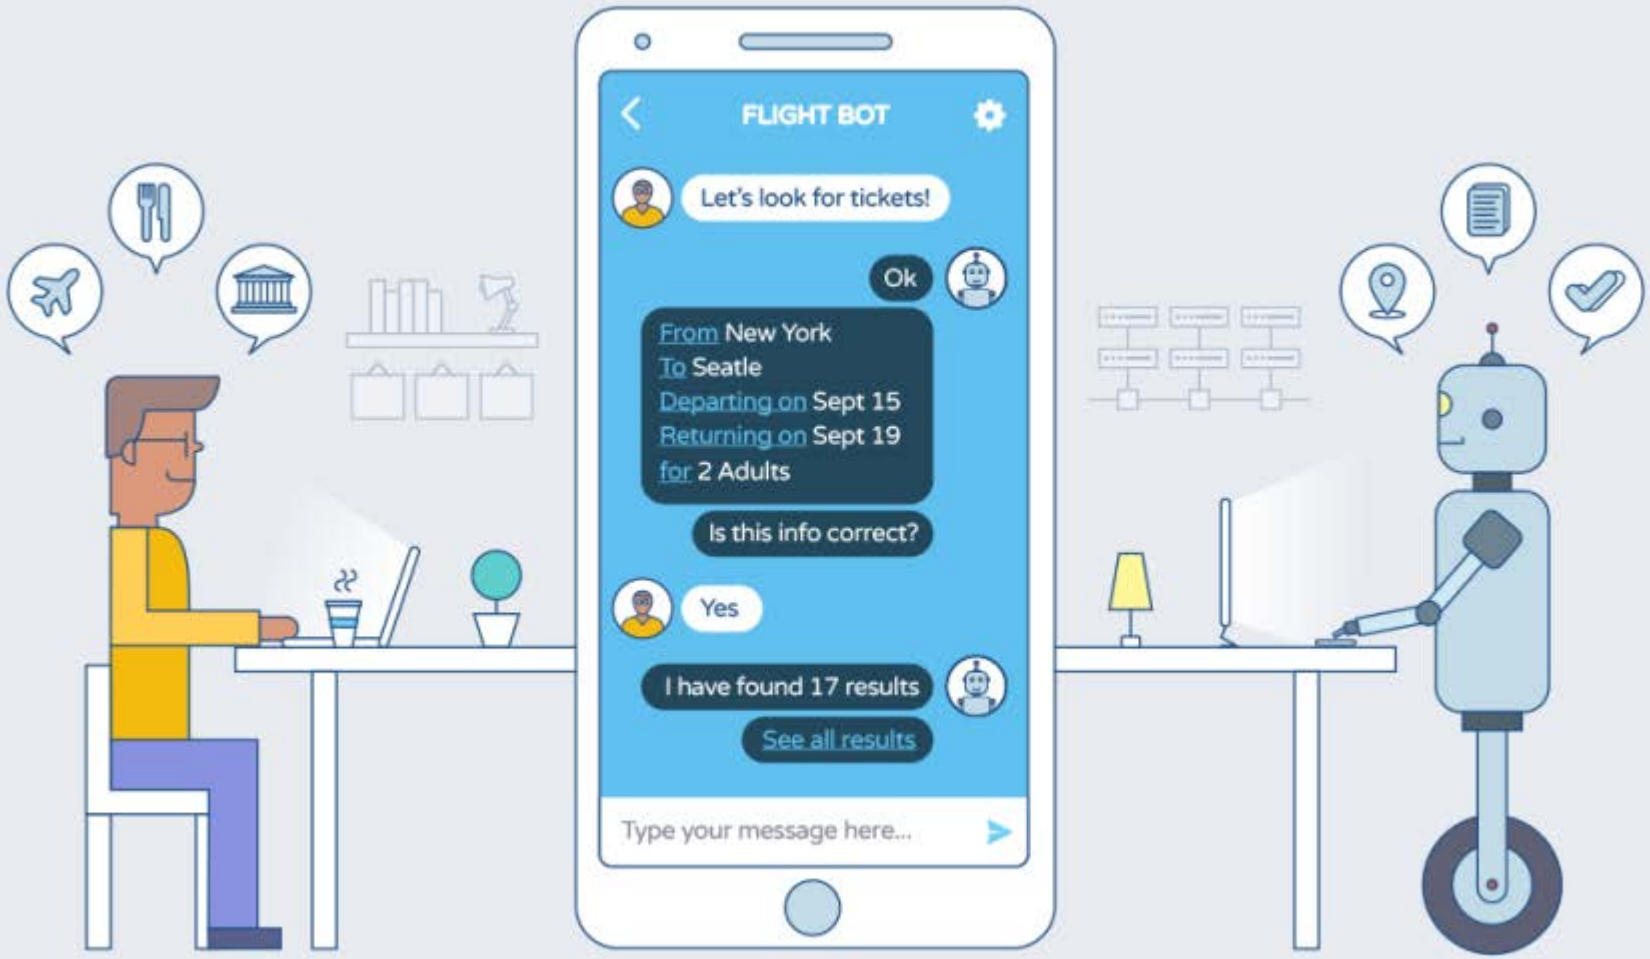
FLIGHT BOT
Let's look for tickets!
Ok
From New York
To Seatle
Departing on Sept 15
Returning on Sept 19
for 2 Adults
Is this info correct?
Yes
I have found 17 results
See all results
Type your message here...

A simple Dialog

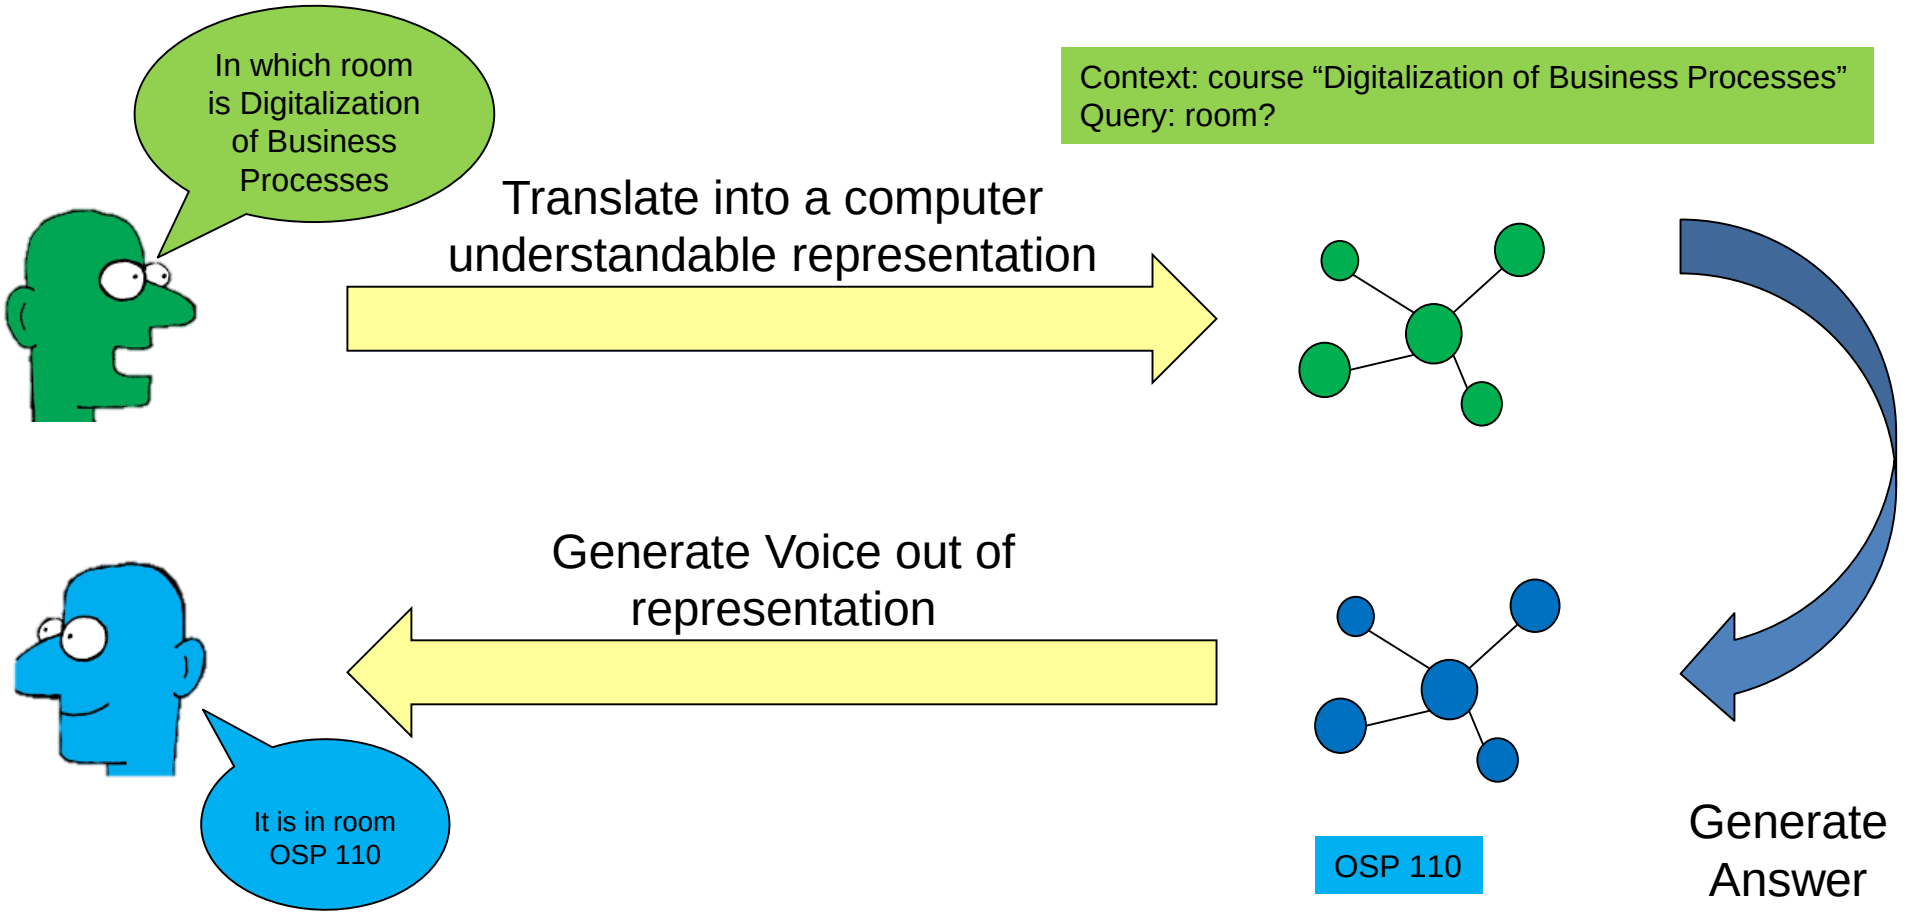

Knowledge Sources for Generating the Answer

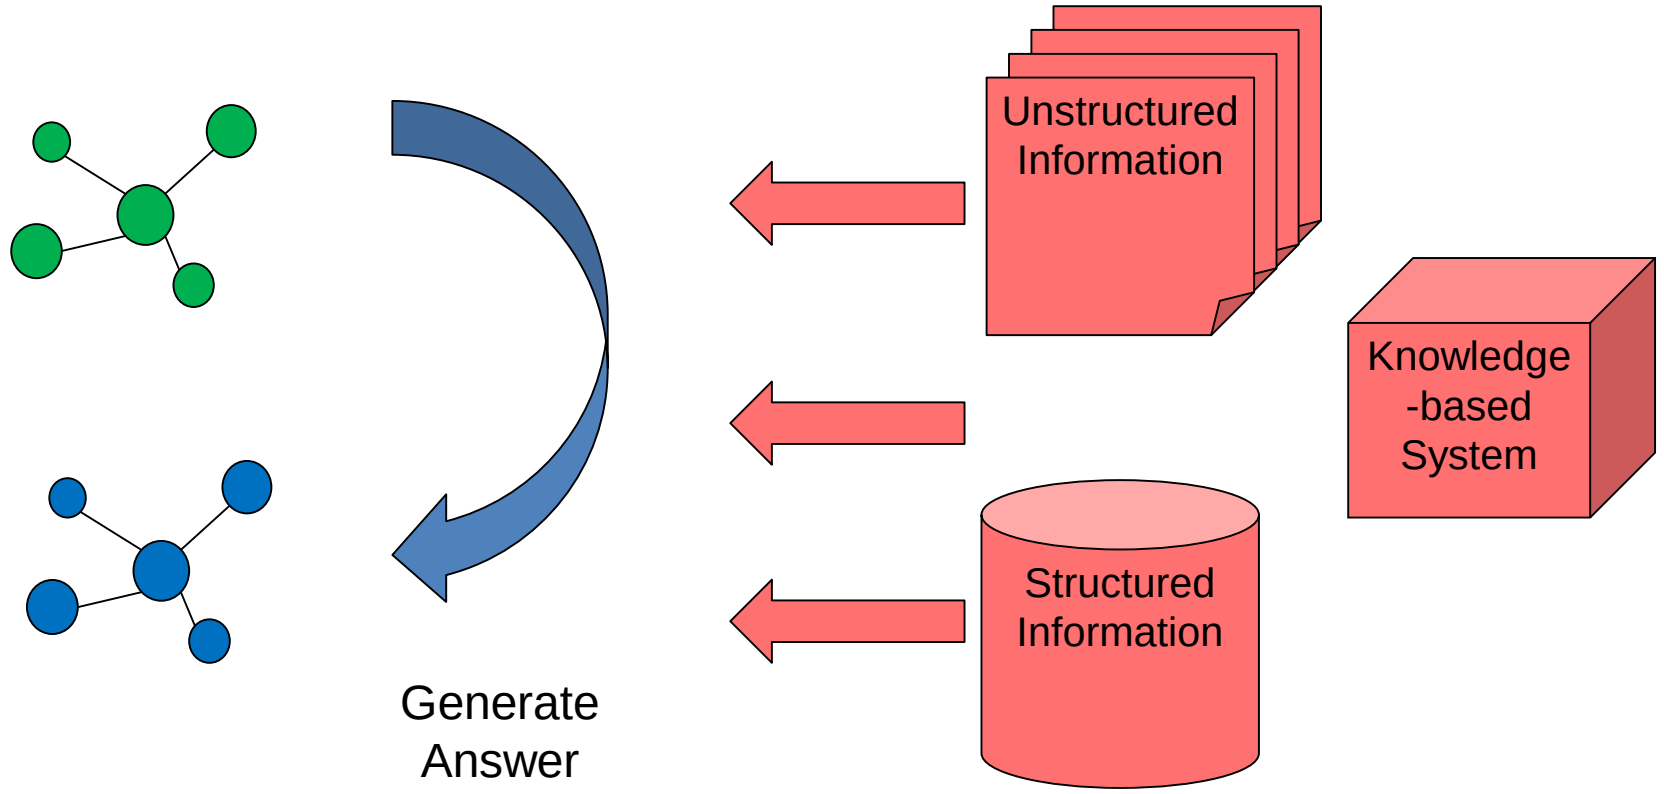

Understanding: the ”easy but stupid way”

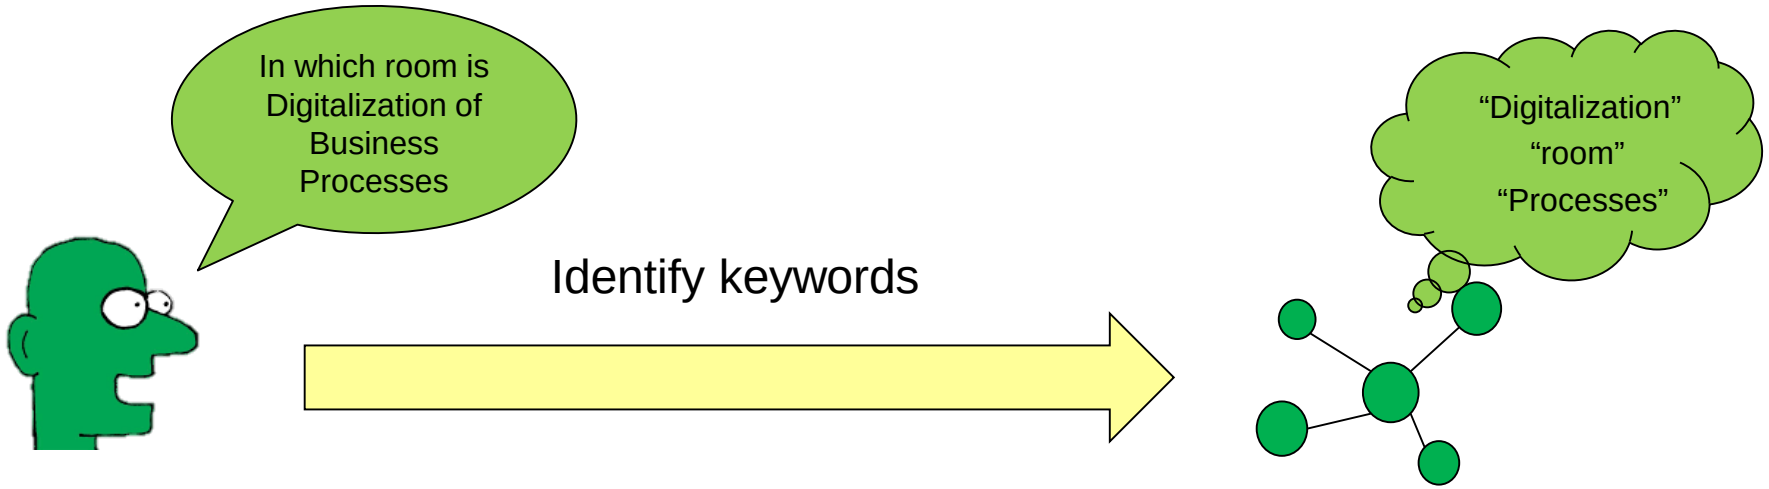

Understanding: the ”complicated way”

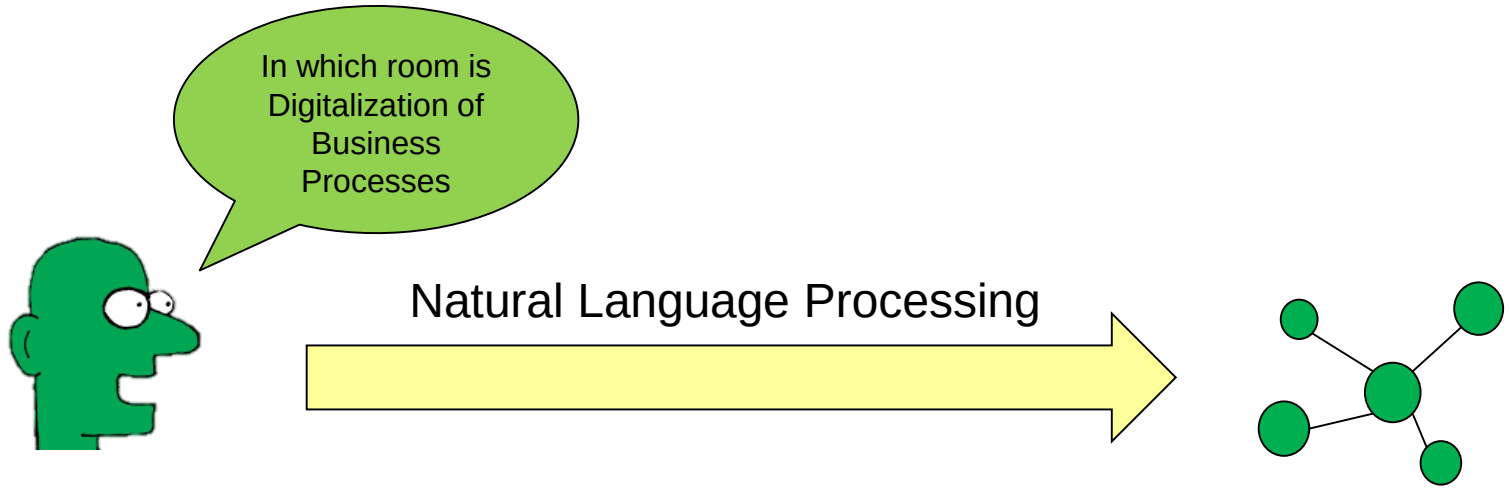

What is Natural Language Processing (NLP)

- Conversion of human language input into a useful form of representation.
- Objective: getting computers to perform useful and interesting tasks, e.g.
  - User Interface
  - Translation
  - Question Answering
- Components
  - Natural Language Understanding
  - Natural Language Generation

Knowledge of Language

- **Semantics** – meaning of words and how these meanings combine to meaning of sentences
  - Problem: Words can have different meaning

    "I saw you at the bank"

- **Syntax** – how to form correct sentences, what structural role each word plays in the sentence and what phrases are subparts of other phrases
  - Problem: Different ways to parse a sentence

    "Time flies like an arrow. Fruit flies like a banana."

- **Pragmatics** – how sentences are used in different situations and how use affects the interpretation of the sentence.
  - Problem: Sentences are meant differently

    "Do you know what time it is?"
    "Yes"

Knowledge of Language (cont.)

- **Discourse** – immediately preceding sentences affect the interpretation of the next sentence.
  - Problem: Words can have different meaning

    "Mary gave a book to Kelly. She didn't like it."

- **World Knowledge** – includes general knowledge about the world. What each user must know about the other's beliefs and goals.
- **Phonology** – how words are related to the sounds that realize them.
- **Morphology** – how words are constructed from more basic units called morphemes.

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