Learning About and Learning with Artificial Intelligence in School

From Understanding of Basic AI Concepts to Trustworthy and Human-centric AI Tools

Ute Schmid
Cognitive Systems
Applied Computer Science
University of Bamberg, Germany
www.uni-bamberg.de/cogsys/

Forschungsgruppe Elementarinformatik

ProDaBi-Colloquium (Paderborn/online), Dec. 7 2022

What should children learn about AI in school?

www.post-gazette.com/business/tech-news/2019/01/06, ReadyAI

How can AI be used to support teaching and learning?

observatory.tec.mx/edu-news, Bigstock

Learning and Teaching about AI

- Identifying relevant AI topics and methods
- Teaching different age groups
- Teaching students with and without background in computer science
- Teaching students in different types of school (middle school, high school, vocational school)
- Development of curricula and text books
- Integrating AI as a topic in training computer science teachers

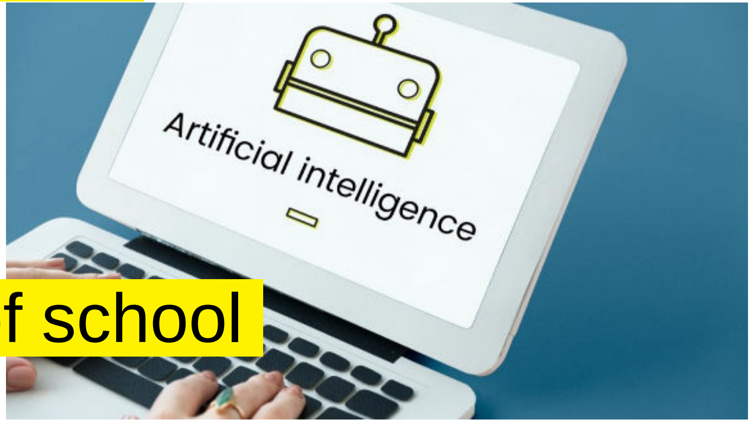

https://thestempedia.com/blog/teaching-ai-kids/

Some Challenges

- Get an overview of AI topics and methods
- Understand the difference between standard computer science algorithms and AI algorithms
- Understand the basic principles of machine learning
- Understand the difference between human and machine intelligence
- Be able to assess where AI can be helpful and where it might be dangerous → foundation for evaluating ethical aspects of AI

AI is an established area of research in computer science

Term Artificial Intelligence
coined 1956 by John McCarthy

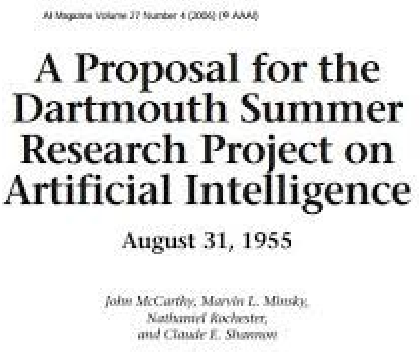

AI Magazine Volume 27 Number 4 (2006) (© AAAI)

A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence

August 31, 1955

John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon

AI research is based on the believe that every aspect of human intelligence can be formalized in such a way that it can be simulated by a computer program.

http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html

Three Flavors of AI

(as represented by the founding fathers)

John McCarthy
Stanford
Separating Knowledge and Inference
Lisp
Logic AI

Allen Newell
Herbert Simon
CMU
General Problem Solving
Cognitive AI

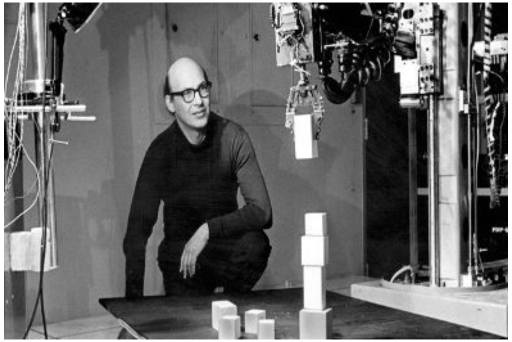

Marvin Minsky
MIT
Vision
Robotics
Multi-Agents
Engineering AI

The Current Hype of AI

- AI research has been done without interruption over the last decades

  see e.g., (bi-) annual IJCAI conference since 1969

- Recent hype: data-intensive end-to-end learning approaches

- IJCAI-93: Chambery, France
- IJCAI-91: Sydney, Australia
- IJCAI-89: Detroit, Michigan, USA
- IJCAI-87: Milan, Italy
- IJCAI-85: Los Angeles, California, USA
- IJCAI-83: Karlsruhe, Germany
- IJCAI-81: Vancouver, Canada
- IJCAI-79: Tokyo, Japan
- IJCAI-77: Cambridge, Massachusetts, USA
- IJCAI-75: Tbilisi, Georgia
- IJCAI-73: Stanford, California, USA
- IJCAI-71: London, England

IJCAI-69: Washington, D.C., USA

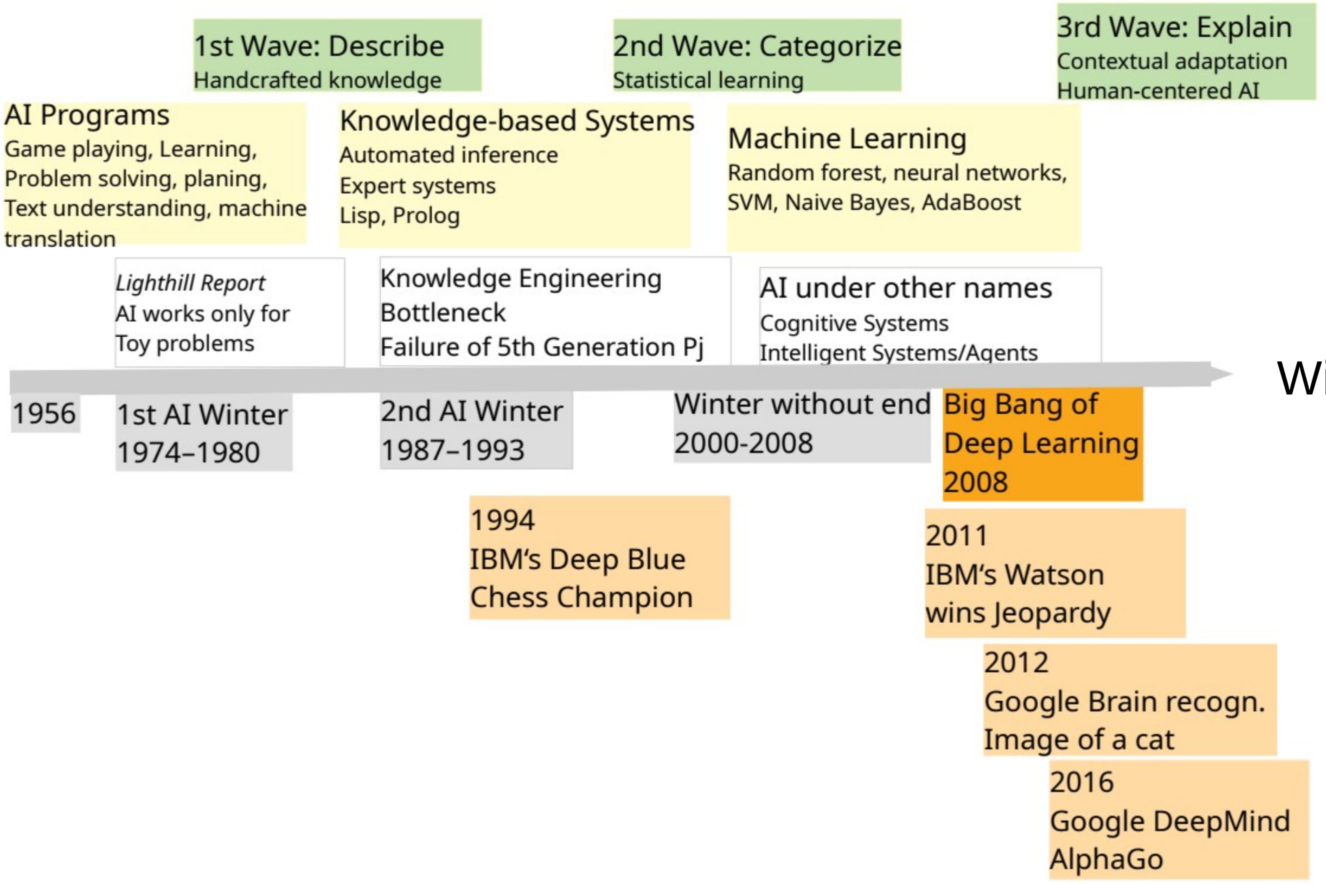

Winter is coming?

Teachers need to understand

- that AI is not only (deep) neural networks
- but a set of well established formal methods
  - ➔ (logic based) knowledge representation, sophisticated heuristic search algorithms

Most widely used textbook AIMA

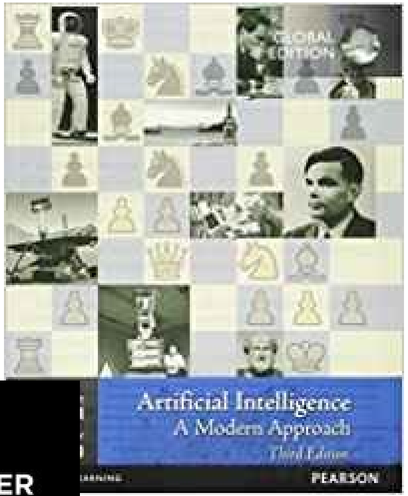

German handbook, 6th edition in 2021

Building-blocks of a Curriculum for AI

- Basic concepts of knowledge representation and reasoning, heuristic problem solving algorithms
  - ✗ Standard search on lists, search trees, graph search (e.g. Dijkstra alg) vs. AI search (problem spaces which grow exponentially)
- Basic concepts of machine learning: inductive inference, decision tree learning, perceptron, feed-forward networks
  - ✗ Learning as approach for domains where no (complete and correct) model exists, where human knowledge is mostly implicit

Teaching ML: Entry into the topic

Training Examples

Class (concept learning)

Features (discrete valued)

Story: A robot learning in what kind of packages there is probably a present

Parcel	Brightness	Size	Present?
parcel-1	bright	medium	no
parcel-2	very dark	small	yes
parcel-3	dark	Very small	no
parcel-4	bright	large	yes
parcel-5	medium	small	no

Didactic Advantages of this Example

- Students can understand that the class is not seen /known except for training examples when the parcel is opened (difficult for image classification)
- Features can be interpreted categorial (suitable for decision tree learning) and metrical (as brightness values, height in cm, suitable for perceptron/neural networks)

Schmid (2012), How do computers and robots learn?, Children‘s University
Schmid et al. (2019). KI selber programmieren für Dummies Junior, Wiley

Perceptron Learning

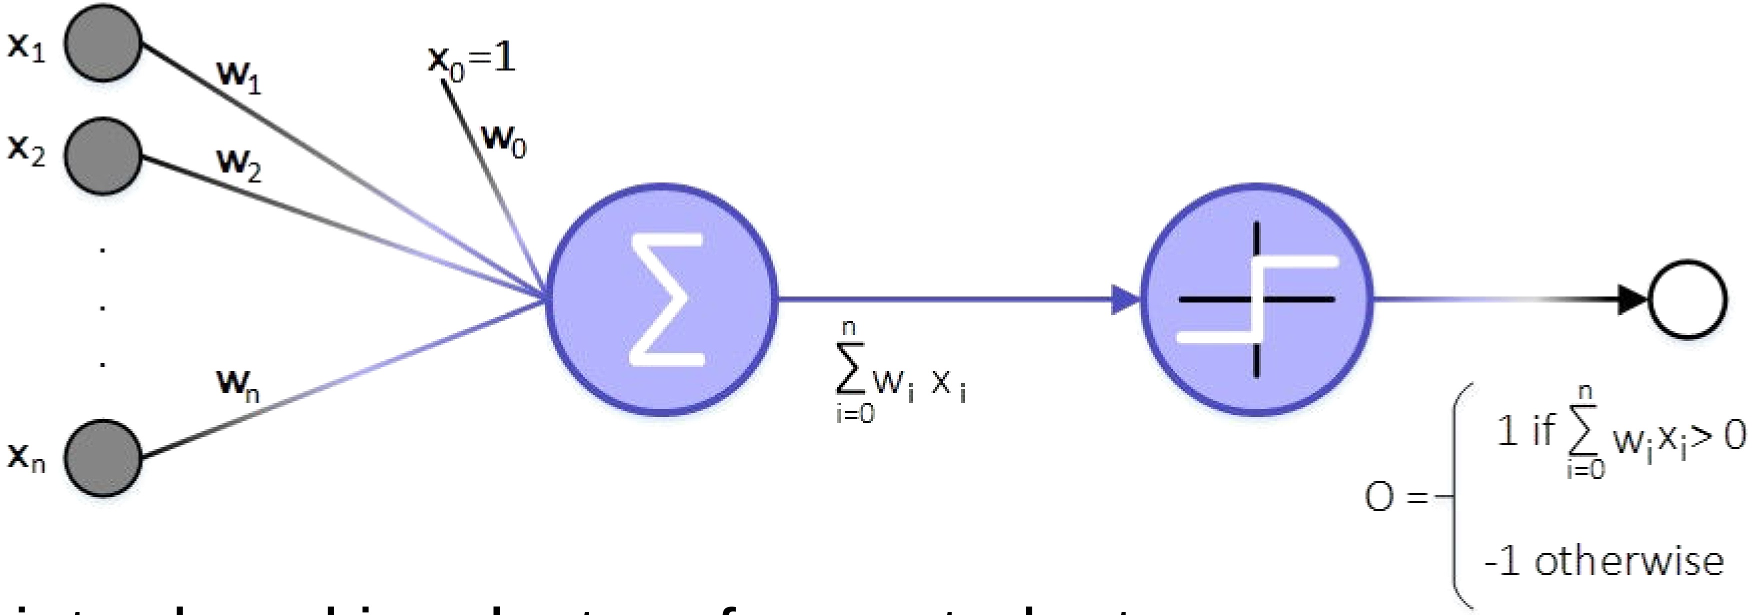

as introduced in a lecture for cs students,
with threshold activation function

Perceptron Learning

- Without bias weight, without activation function, only two features (simple addition), weights introduced as importance

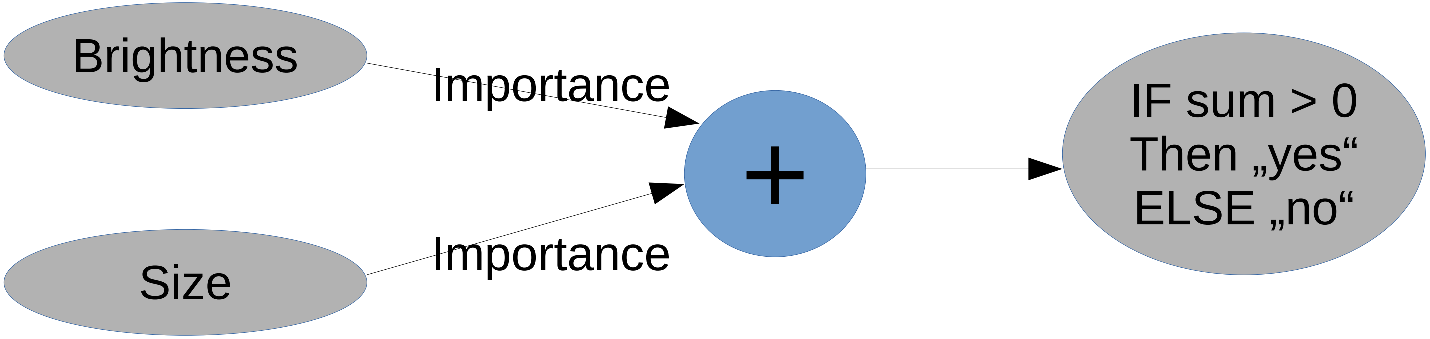

Perceptron Learning

Training Examples

Features

Class (concept learning)

Story: A robot learning in what kind of packages there is probably a present

Parcel	Brightness	Size	Present?
parcel-1	4	3	no
parcel-2	1	2	yes
parcel-3	2	1	no
parcel-4	4	4	yes
parcel-5	3	2	no

Perceptron Learning Unplugged

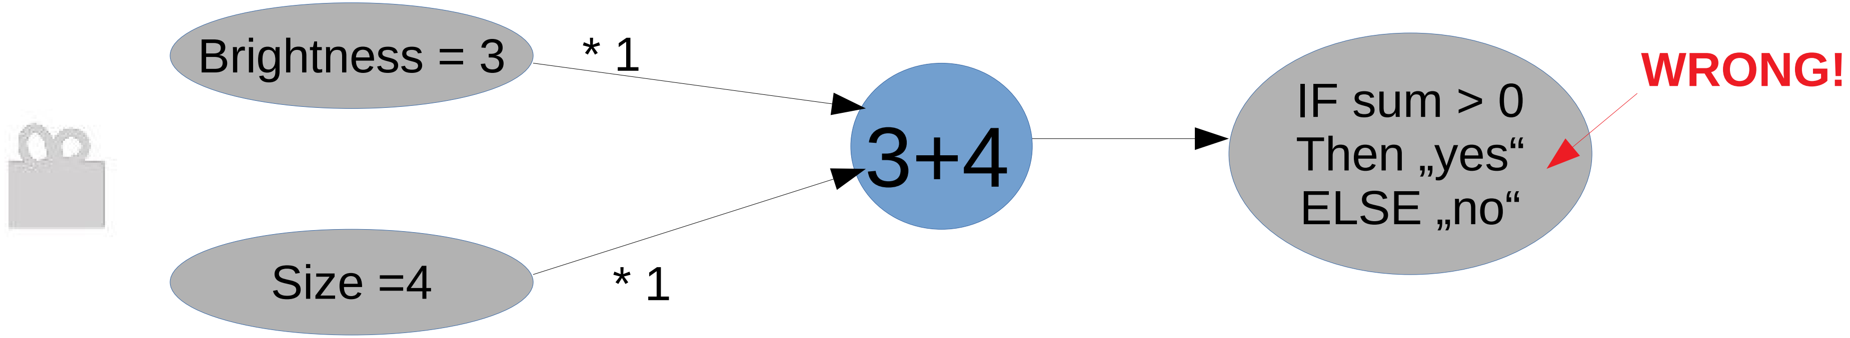

Learning from errors:

If „yes“ is wrong:
subtract feature values
from importance

i.e., change of weights

Perceptron Learning in Python

```
# Training Examples
learn_parcels = [
    (4, 3, "no"),
    (1, 2, "yes"),
    (2, 1, "no"),
    (4, 4, "yes"),
    (3, 2, "no")]

# Initial importance values
importance_brightness = 1
importance_size = 1

iteration = 0
error = True
```

```
# Repeat until no errors occur
# but 5 iterations max
while error and iteration < 5:
    iteration = iteration + 1
    error = False

   # inerate over all training examples
    for parcel in learn_parcels:
        # unpack parcel
        brightness, size, present = parcel

        # Calculate Decision
        value1 = brightness * importance_brightness
        value2 = size * importance_size
        decision = value1 + value2
```

Perceptron Learning in Python

```
# Decision: does the parcel contain a present?
        if decision >= 0:
            expect_present = "yes"
        else:
            expect_present = "no"

# Learn: Adjust importance values
        if expect_present == "yes" and present == "no":
            importance_brightness = importance brightness - brightness
            importance_size    = importance_size- size
            error = True
        if expect_present == "no" and present == "yes":
            importance_brightness = importance_brightness + brightness
            importance_size    = importance_size + size
            error = True
```

Relate Perceptron Learning to Math

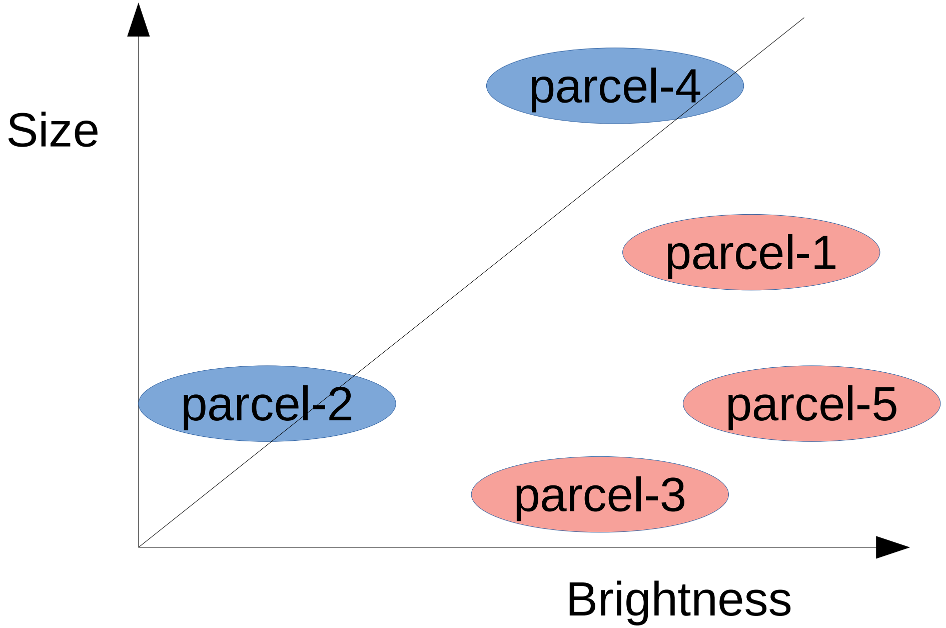

Learning =
change the slope of the linear function

such that it separates parcels with and without presents

Evaluate the Learned Model

- Estimate how good the model will perform on unseen instances (e.g., new parcels)
- Keep a subset of instances from the training set to test the learned model, count errors (accuracy, misses/false alarms)

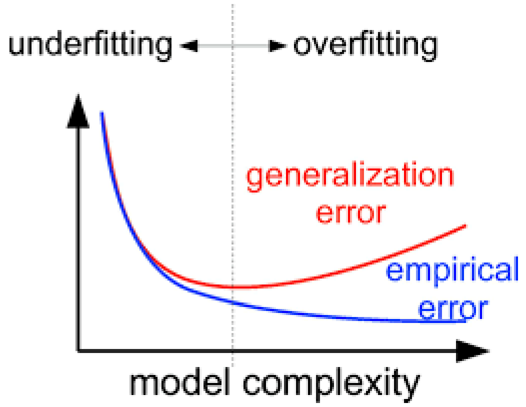

Generalization learning cannot be error-free!

How dangerous is it to apply a model with 99% accuracy
– for the parcel opener robot?
– for a self-driving vehicle?

Understand: Inductive Learning

- Most machine learning is inductive reasoning, that is: generalize from training examples to all possible instances
- Different from: rote learning (insert items into data bases or knowledge graphs), here main questions are consistency, retrieval efficiency
- Typically: the instance set is infinite – and the set of possible models which can be learned from a set of training examples also!
- It is important to assess how good a learned model generalizes to unseen instances (use a test set! Better sample differenzt test sets for a statistical estimate of predictive accuracy)
- Inductive learning cannot guarantee correctness, cannot be bias free

Biases in Machine Learning

- **Inductive biases:**

  (inherent in every learning system, artificial or human)

  – Restriction bias: language in which learned models/hypotheses are represented

  – Preference bias: policy by which a hypothesis is selected

    learning is search in (infinite) hypothesis space, typically so complex that a greedy algorithm (e.g. gradient descent) is used

- **Sampling bias:** training examples do not reflect distribution of the intended population (can result in over-fitting and unfairness of models)

Unmasking Clever Hans Predictors

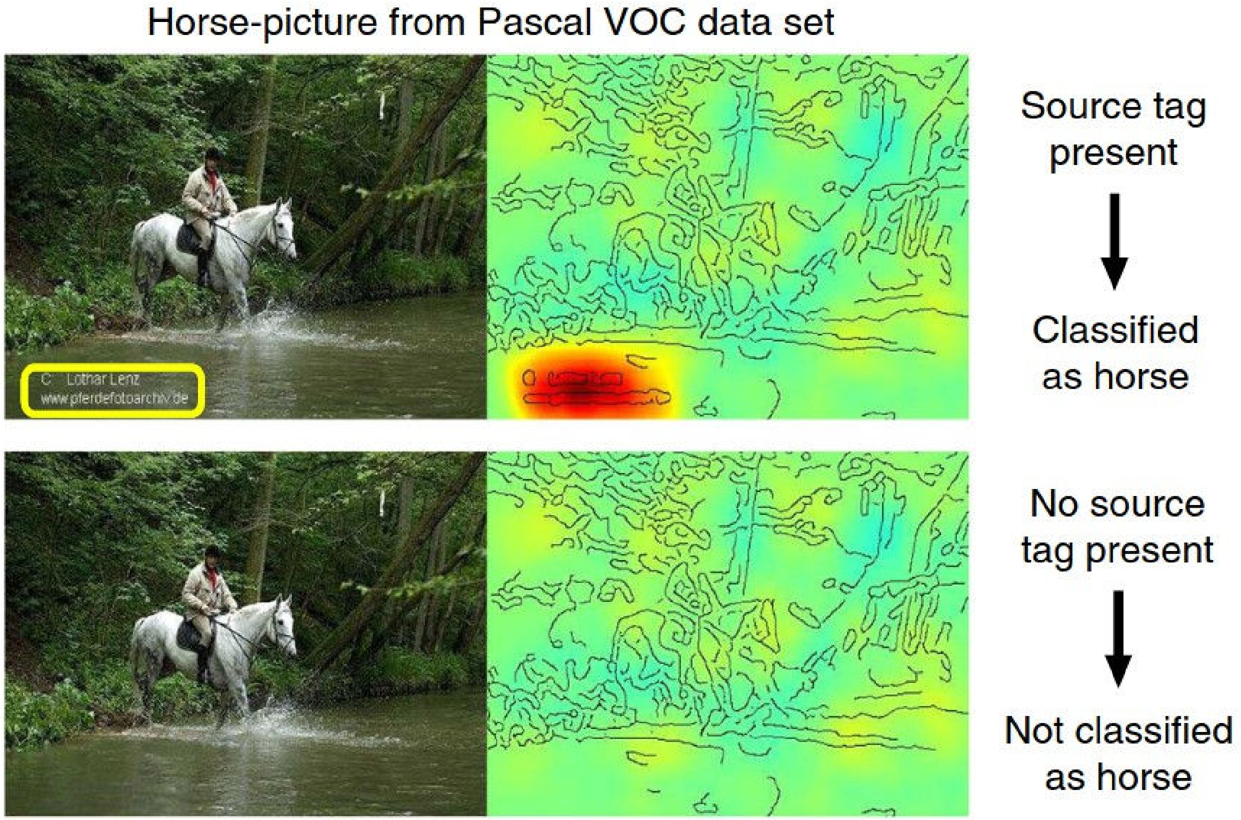

3rd Wave of AI:

Explainable Artificial Intelligence (XAI)

Lapuschkin,et al. Unmasking clever hans predictors and assessing what machines really learn. Nature communications, 2019.

The Power of Human Inductive Learning

Josh Tenenbaum

Human inductive bias allows for powerful and flexible learning e.g. language learning (over-generalization of regularities, see-d/saw)

The dark side: stereotypes and prejudice (such as girls are not good in physics)

Teachers need to understand

- that a perceptron is a very simple approach to inductive machine learning which can help to understand basic principles of learning (e.g. inductive biases, generalization error)
- that there is a long way from understanding a perceptron to understand a support vector machine, a feed-forward network, or a deep learning architecture
- that statistical machine learning in general results in intransparent models (cf. implicit learning in humans)
  - Alternatively, there are approaches to rule-learning with decision tree algorithms as most basic approach

Decision Tree Learning

- Same problem domain (parcels with presents)
- Easier to explain on a superficial level
- But: more complex to realize as a program (tree is a dynamic data type!, that is: computer science topic of upper secondary school)

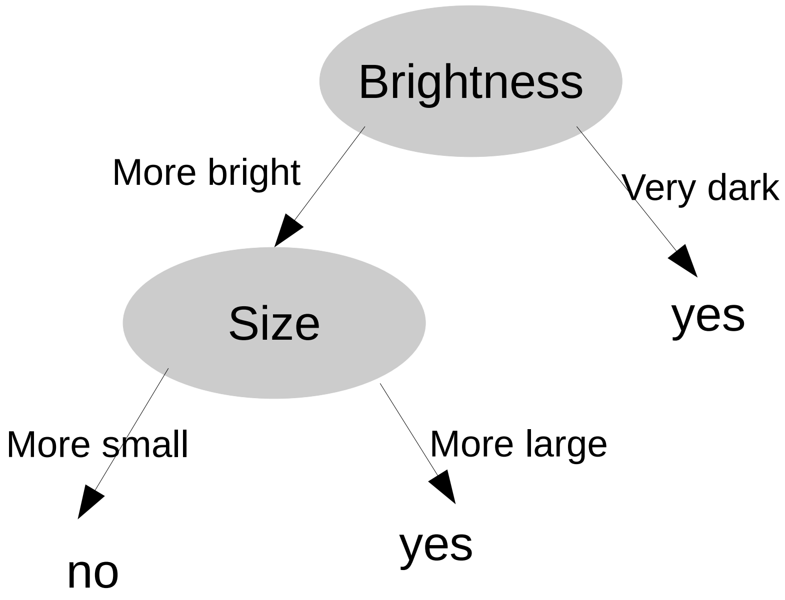

Generalization: only relevant feature combinations

Reasoning

- **Machine learning:** when it is not possible or not easy to present decision rules/knowledge in an explicit way
  - But: no correctness,
  - Often: problem to get enough data with high quality (ground truth) for supervised learning
- **Knowledge-based inference:** e.g. for taxonomies (family relations, biology), easy to realize in Prolog

https://www.inf-schule.de/deklarativ/logischeprogrammierung

Deductive Inference

- Does a blackbird brood?
- Does a blackbird breath?
- This is not explicitly stated in a biology class, but inferred!

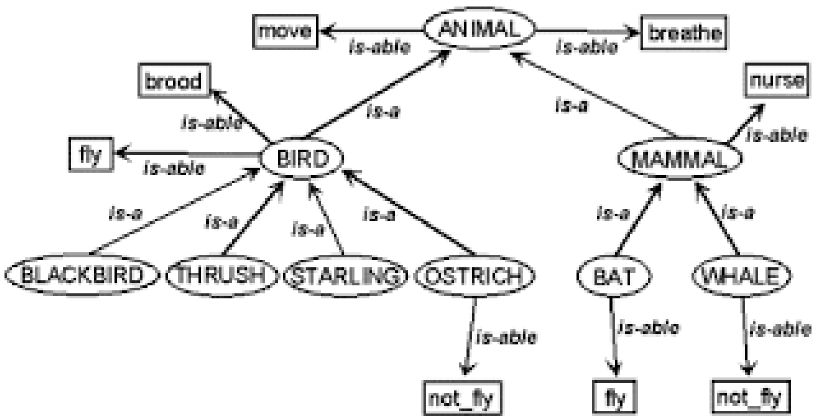

What we know and can infer,

we do not need to learn again

from data!

From: Rödder & Kulmann. Recall and reasoning—an information theoretical model of cognitive processes. Information Sciences 176.17 (2006): 2439-2466.

Deductive Inference

Modus barbara

$A \rightarrow B$
$B \rightarrow C$

$A \rightarrow C$

All Blackbirds are Birds
All Birds are Animals

Therefore, all Blackbirds are Animals

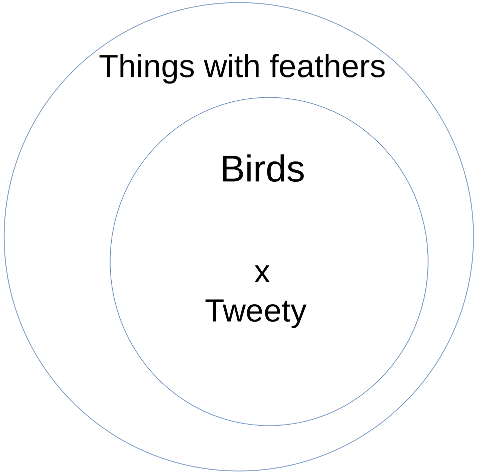

Modus ponens

$A \rightarrow B$
A

B

All birds have feathers
Tweety is a bird

Therefore, Tweety has feathers

Who should know what about AI?

- CS Teachers: a general understanding of AI as a long established area of applied computer science (given the classic four parts: theoretical, practical, technical, applied), basic algorithmic approaches (knowledge representation, inference, heuristic search, machine learning)
- Primary school students: basic inductive learning, inference and heuristic search algorithms (computational thinking) based on unplugged activities (see e.g. HABA EA Digital Starter games, AI unplugged at FAU)
- Secondary school students: more formal treatment, relation to other concepts (linear functions, dynamic data types, search algorithms, …), programming basic algorithms, maybe also: working with tools

Who should know what about AI?

- Students and teachers without computer science background:
  - basic understanding of exemplary domains, such as classification based on learned models
    - Related to social scoring, to learning analytics
- Understand, that AI systems are in a different way intelligent as humans (strong/weak AI)
- Discuss which jobs can and which jobs should be replaced by AI systems
    - E.g. old age homes: lifting, documentation, reading, going for a walk, …
    - E.g. in school: explaining, testing, grading

Strong/weak AI

- Most AI systems are only good (maybe even better than humans) at one specific thing
  - A system which us very good at recognizing vehicles cannot recognize animals, cannot solve text algebra problems, does not understand a joke and cannot write an essay

  Humans tend to ascribe intelligence (to other people, to computer systems)

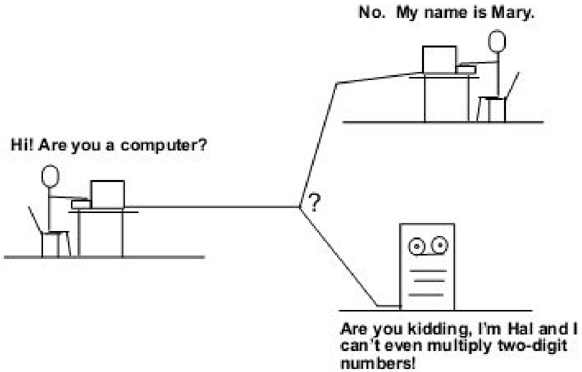

www.researchgate.net/profile/Chinmay_Bepery

- **Moravec‘s paradox:** it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility
- Strong AI aims to realize general intelligent behavior, presupposing intentionality and consciousness

Deep learning is NOT strong AI

see G. Marcus – Deep Learning, A critical appraisal

Beware of Illusionary Intelligence when Interacting with Humanoid Robots

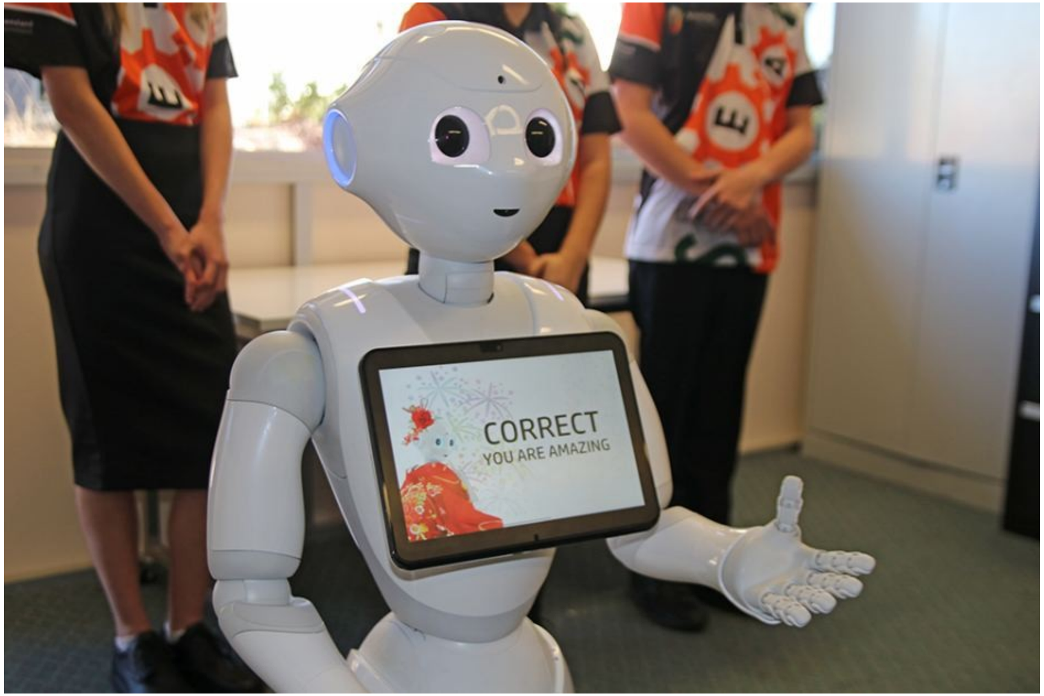

www.abc.net.au

Teaching and Learning with AI

- **Learning analytics:** data-based scoring
  - Need for data might influence didactic decisions
  - Monitoring to support might lead to unwanted selection effects
- **Gamification** for learning by rote domains (vocabulary, multiplication table, capitals)
- **Intelligent tutor systems:** understand misconceptions, individual feedback and training

Intelligent Tutor Systems

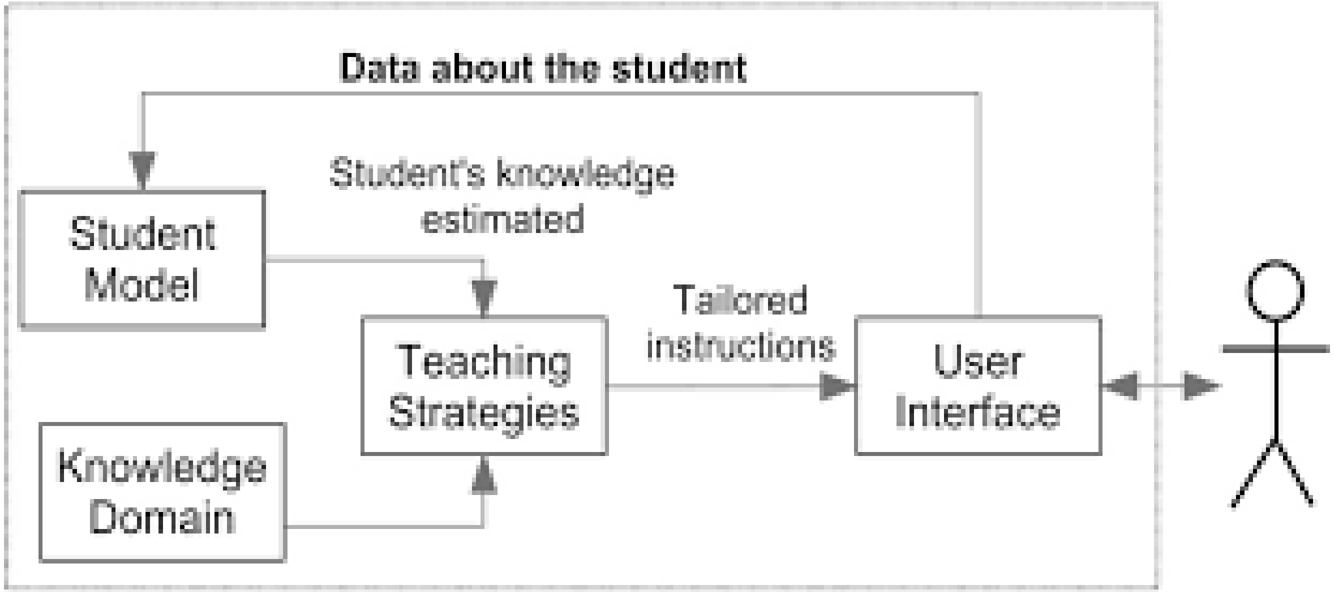

Fig. 1. Basic architecture of an ITS [7].

Butz, Hua, Maguire, 2006

An ITS for Subtraction

subtraction problem

$$\begin{array}{cccc} & C_3 & C_2 & C_1 \\ & 4 & 3 & 7 \\ - & 3 & 7 & 4 \\ \hline & & & \end{array}$$

typical student solution

$$\begin{array}{cccc} & \mathfrak{m}_{-1} & \mathfrak{m}^{+10} & \mathfrak{m} \\ & 3 & 3 & 7 \\ - & 3 & 7 & 4 \\ \hline & 0 & 4 & 3 \end{array}$$

analogous problem

$$\begin{array}{cccc} & C_3 & C_2 & C_1 \\ & 4 & 1 & 0 \\ - & 1 & 8 & 0 \\ \hline & & & \end{array}$$

analogous problem solution

$$\begin{array}{cccc} & \mathfrak{m}_{-1} & \mathfrak{m}^{+10} & \mathfrak{m} \\ & 3 & 11 & 0 \\ - & 1 & 8 & 0 \\ \hline & 2 & 3 & 0 \end{array}$$

Zeller & Schmid, ICCBR Workshops 2016

Take Home Messages

- AI will be part of more and more domains of our live in the future
- For a realistic assessment, a basic understanding of how AI algorithms work is necessary
- Currently many teaching offers are developed: as always the challenge is – teach it correctly but simplify suitably!

miro.medium.com