TDDE19 Advanced Project Course – AI and Machine Learning

AI Projects (Process)

Mattias Tiger (PhD, AI Researcher)

AI and Integrated Computer Systems (AIICS),
Department of Computer Science

mattias.tiger@liu.se

LINKÖPINGS UNIVERSITET

ARTIFICIAL INTELLIGENCE AND INTEGRATED COMPUTER SYSTEMS
AIICS

Recap | The AI Toolbox

Approaches to automated problem solving

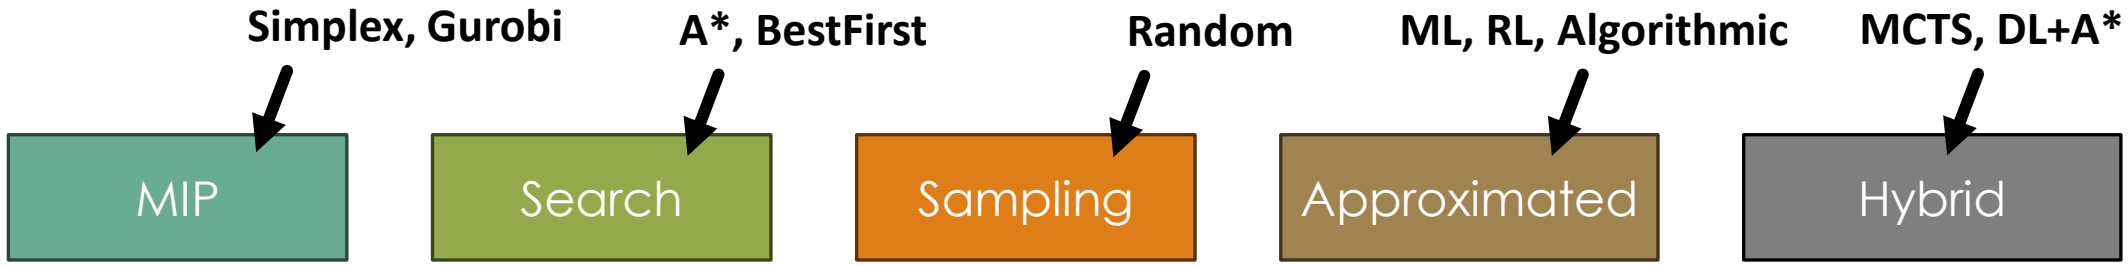

- Assumes smoothness in the search space
- Assumes effective heuristics
- Assumes calibrated probability distributions & efficient sampling schema
- Assumes sufficiently good approximation. For ML: assumes *representative* training data, sufficient capacity and sufficient training time.
- Assumes effective integration of different techniques

Recap | The AI Toolbox

Approaches to automated problem solving

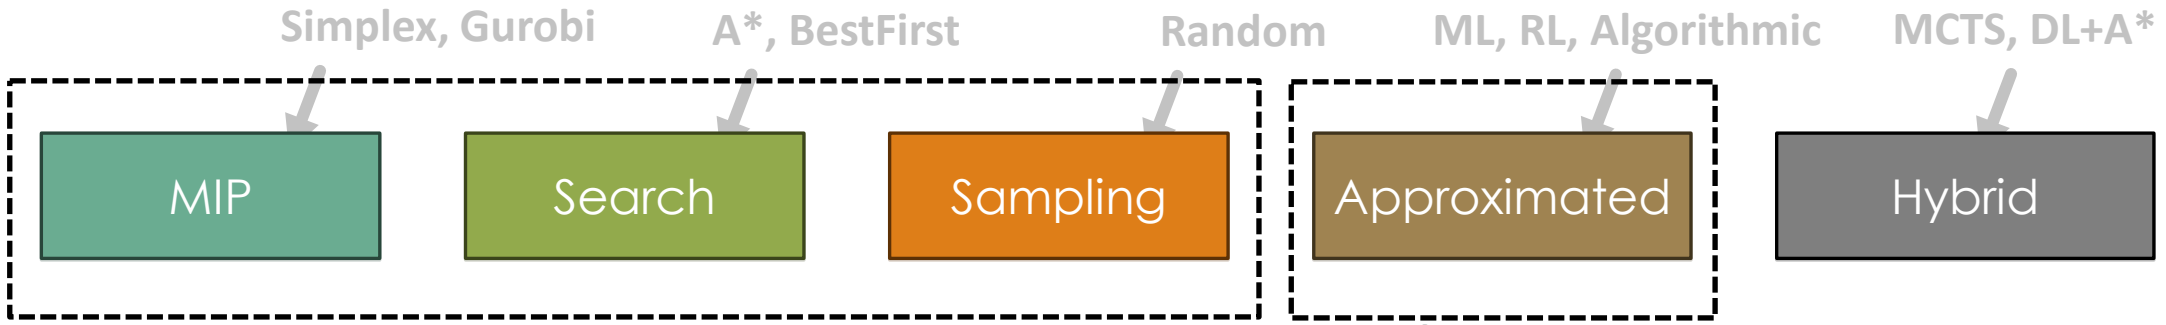

Recap | The Big Picture| Applied AI

Algorithms, solvers and learners

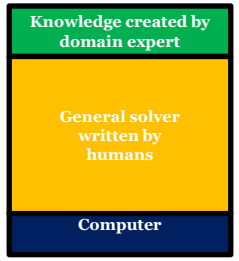

Knowledge-based

Solver

- Capable of solving **different types** of problems
- Optimal in some sense
- The answer is **guaranteed** to be the solution to the problem

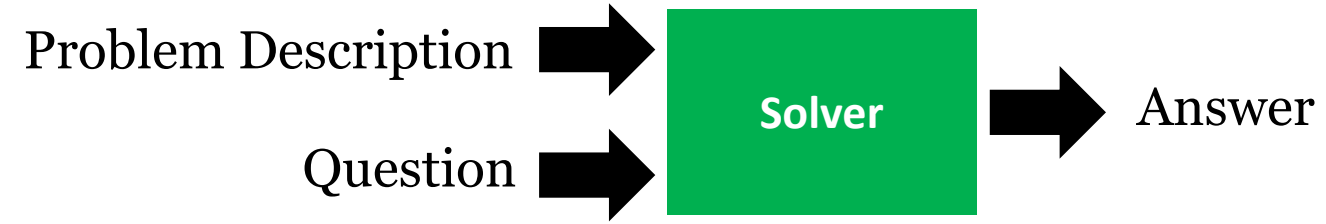

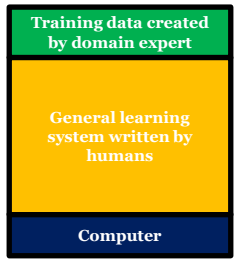

Learning-based

Learner

- The problem is given indirectly **through data**
- Characteristics depend on the chosen technique
- Sometimes gives incorrect answers.

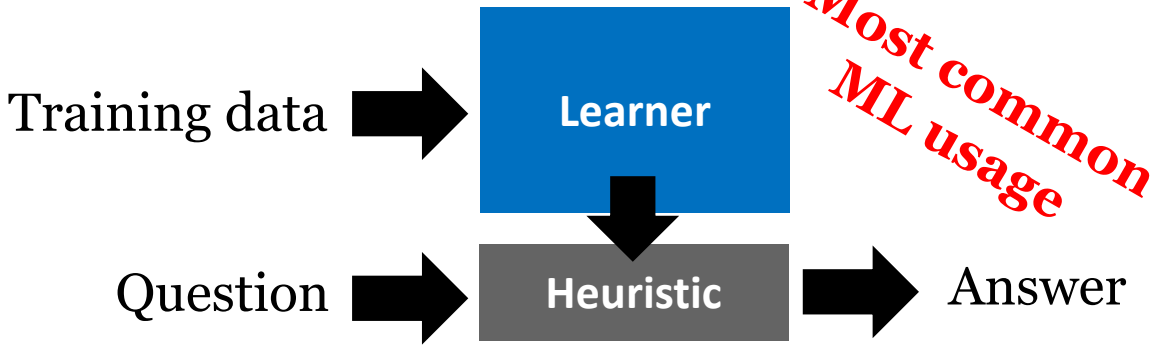

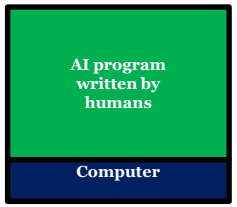

Algorithmic

Algorithm

- Solves a specific problem

Machine Learning | Big Picture

LINKÖPINGS UNIVERSITET

Machine Learning

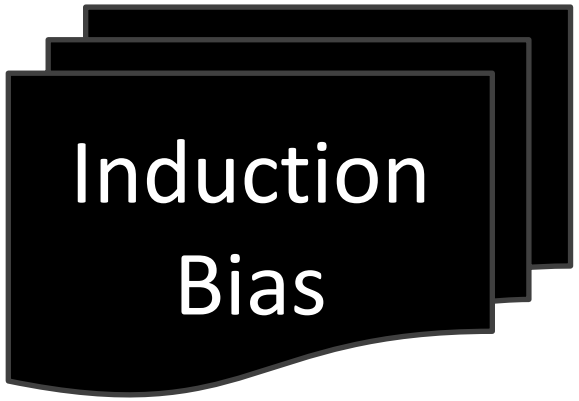

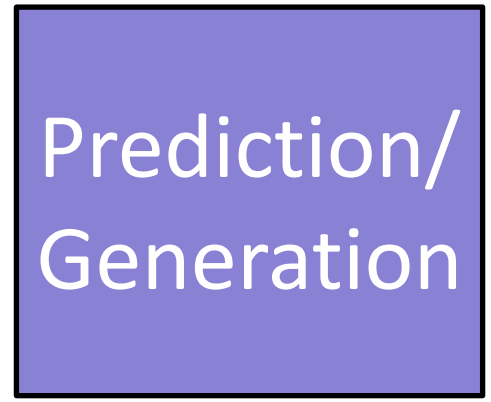

Simple example

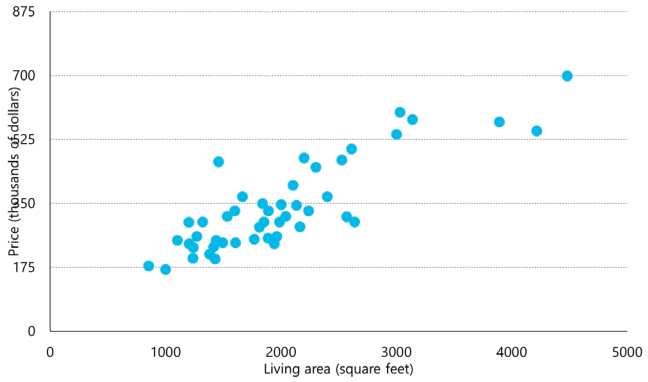

Linear regression

$$\hat{y} = \boldsymbol{x\theta}$$

- Linear model
- Normal distributed noise
- Square loss

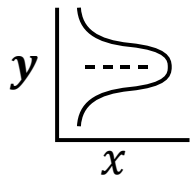

$$|\hat{y} - \boldsymbol{x\theta}|^2$$

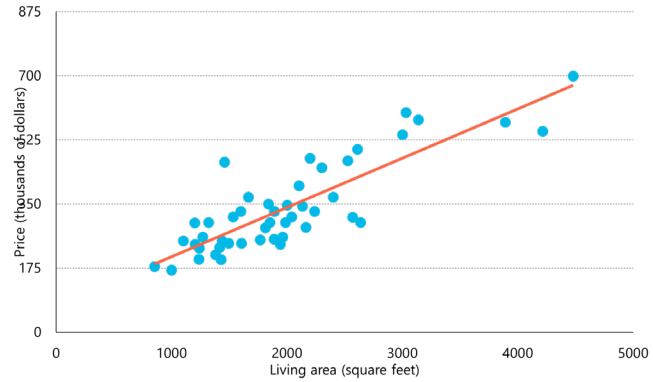

ML | The Big Picture

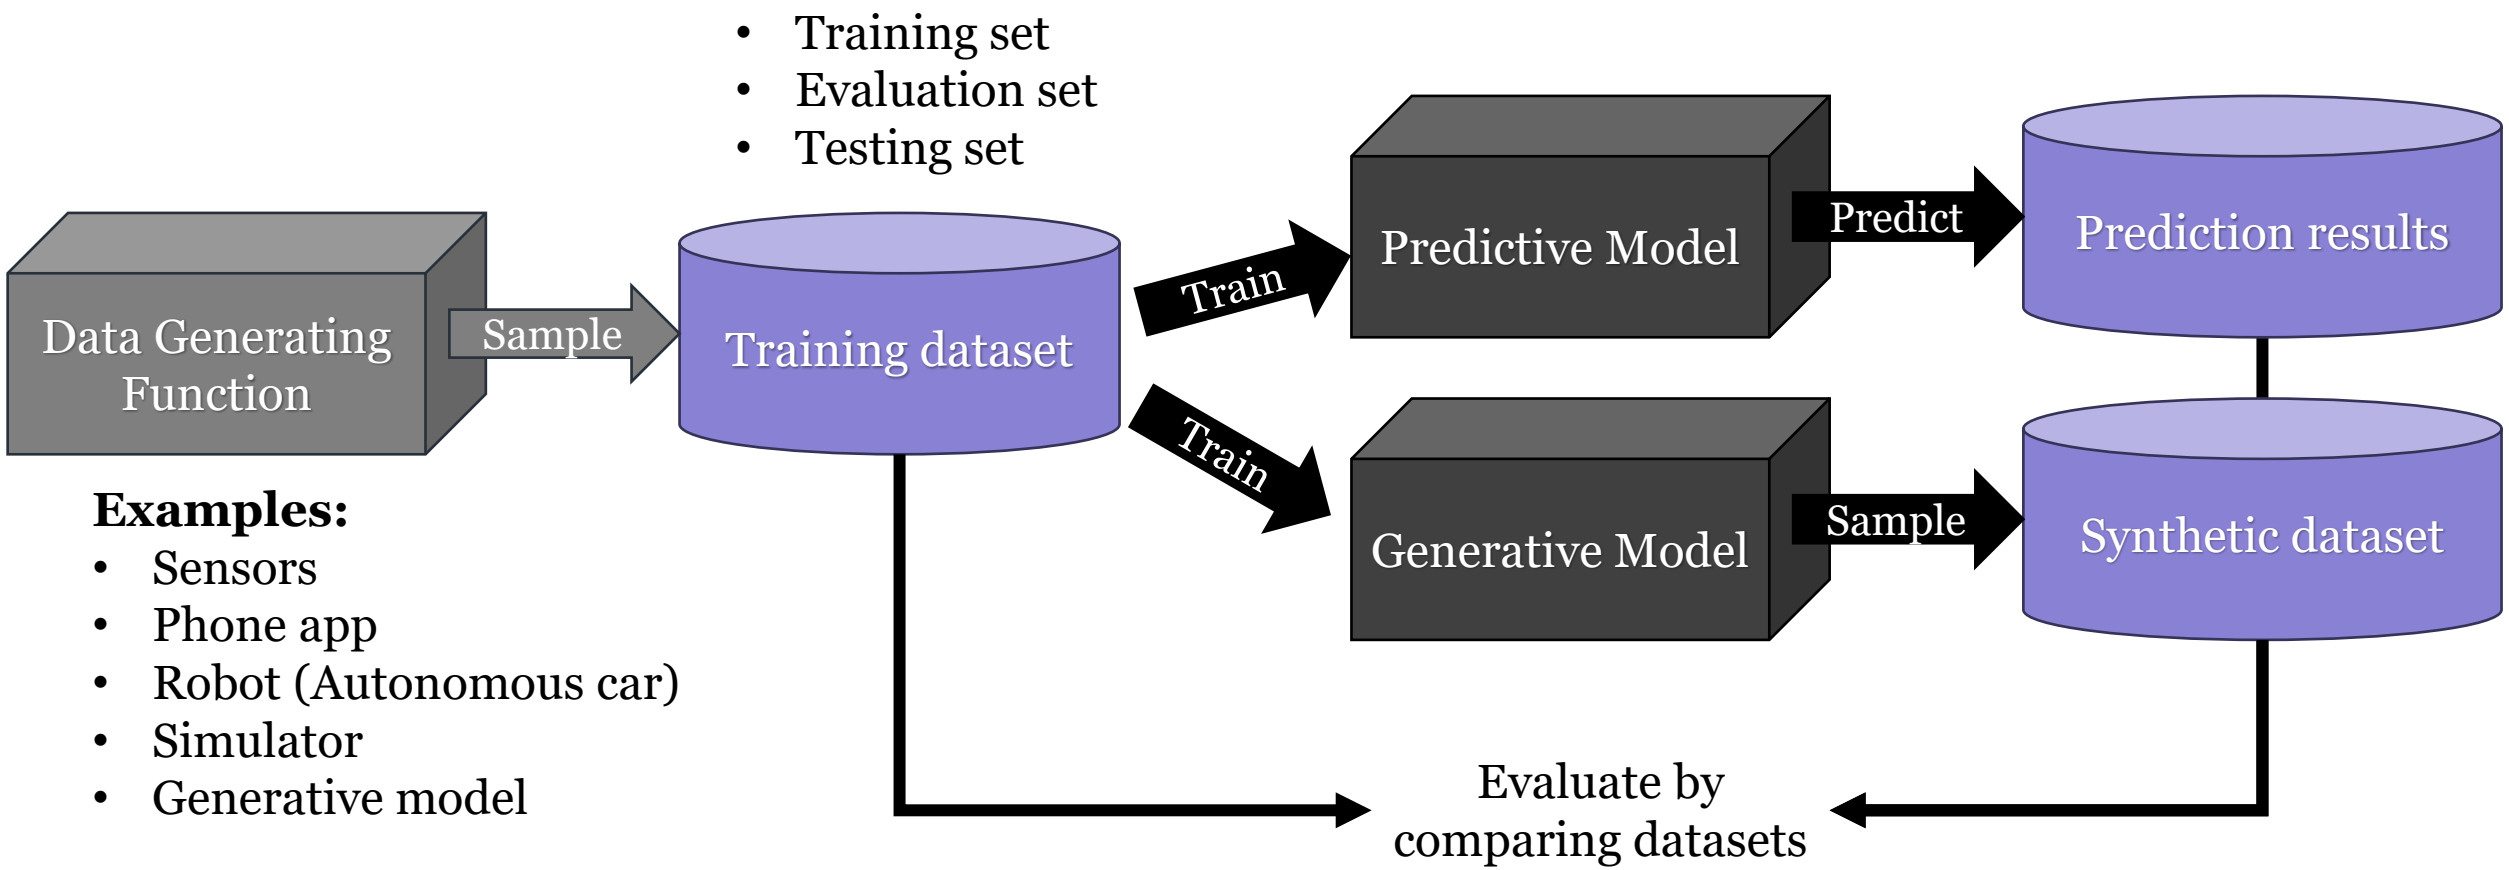

Examples:

- Sensors
- Phone app
- Robot (Autonomous car)
- Simulator
- Generative model

ML | The Big Picture | Evaluation is difficult

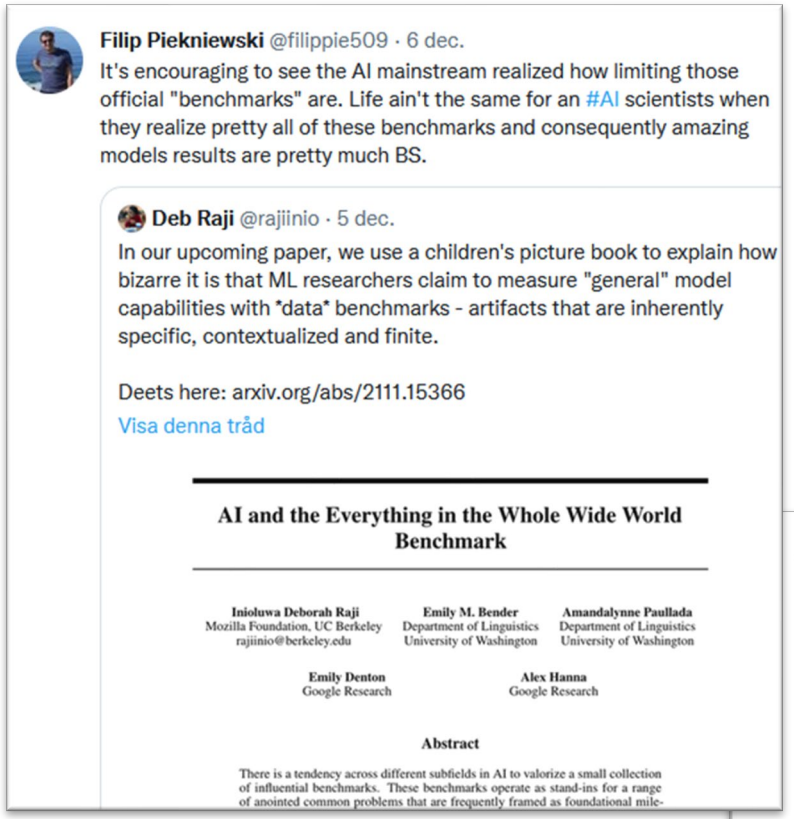

Be aware of the limits of benchmarks:

- Benchmarks are a substitute for reality, not reality
- Benchmarks get useless quickly
- Statistically, people/groups will start to overfit to benchmarks

  ⇒ results potentially meaningless
- **Update/Create new regularly!**

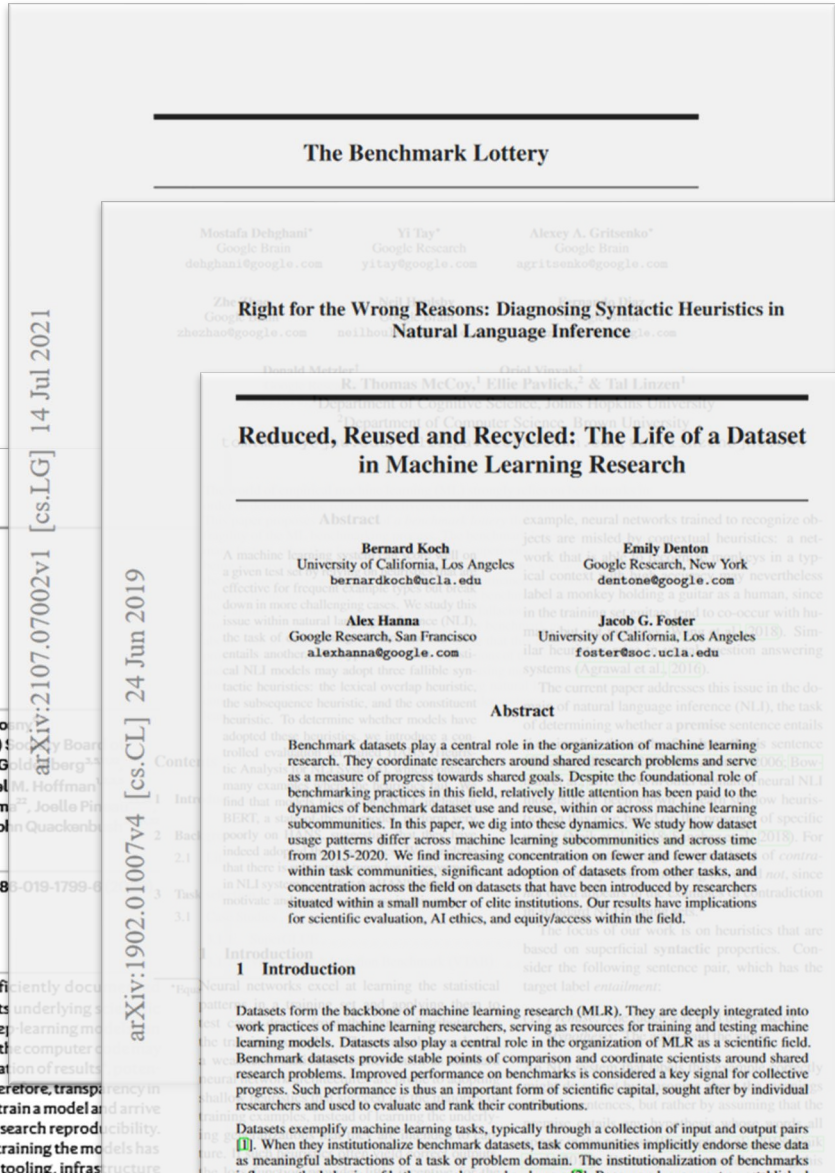

ML | The Big Picture (What we want)

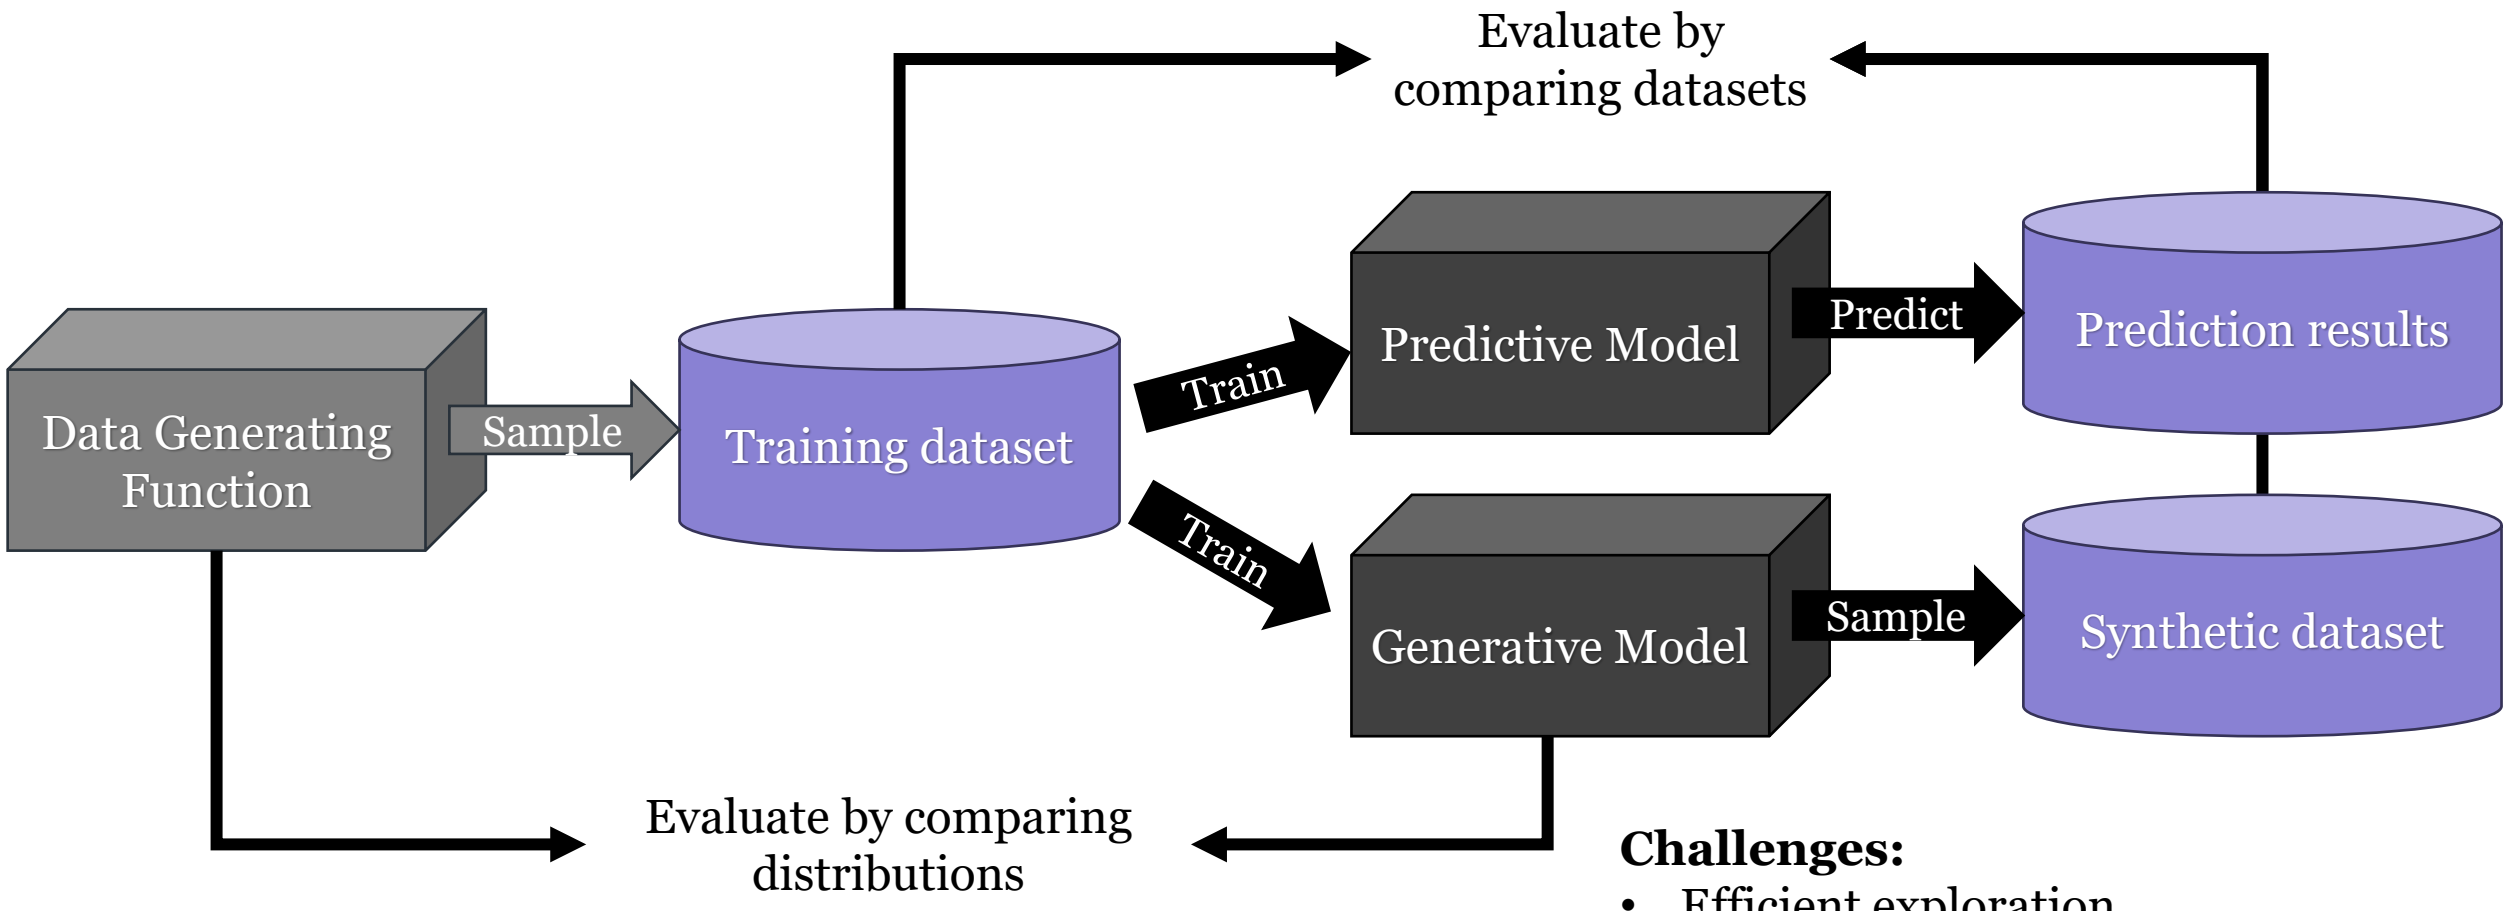

Challenges:

- Efficient exploration
- Effective comparative measures
- ➢ Grounding

AI | Projects

LINKÖPINGS UNIVERSITET

Applied AI Projects

- **Goal** Clear for all parties. Long-term and short-term.
- **Data** Data readiness. Information > Data.
- **Competence** Domain experts, Data management experts, AI specialists, AI experts.
- **Tools** Flexible laboration & prepared for deployment in organization.
- **Process** Agile and iterative processer – engineering and research practices.
  Take the right decision early with the right knowledge.
  Adapt to the AI/ML (etc.) maturity of the organization.

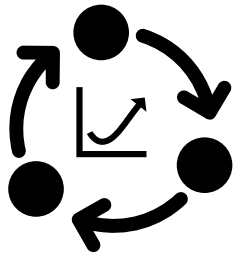

Applied AI Projects

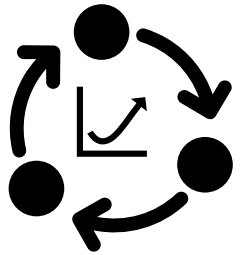

- Goal
- **Data**
- Competence
- Tools
- Process

Data Readiness [7]

- Quality and usefulness of data
- 3 bands (C -> B -> A)
- Data ready (AI ready)

How informative a data set is <u>depends</u> on the application.

Machine learning: **<u>Data</u>** ✚ **Induction bias** ⟹ **Prediction**

assumptions, model, uncertainty, loss function, …

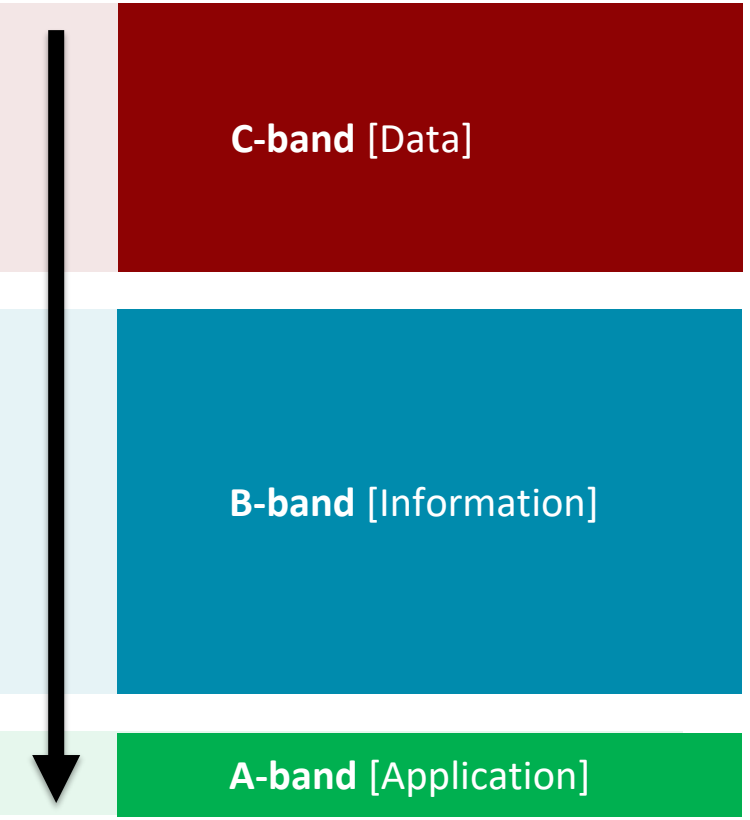

[7] N. D. Lawrence, Data Readiness Levels, 2017

Applied AI Projects

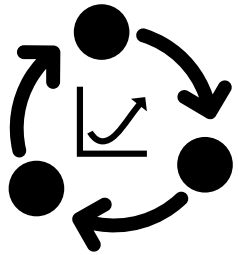

- Goal
- **Data**
- Competence
- Tools
- Process

- What data exists?
- On what format (schema)?
- Legal restrictions on access and usage?
- Limitations on where it may be stored and processed?

C-band [Data]

- Units? (seconds or hours?)
- Preprocessing and aggregations?
- Missing data?
- Incorrect data?
- Uncertain data?
- Regimes and trends over data properties?

B-band [Information]

- Can problem X be solved with the data set?

A-band [Application]

[7] N. D. Lawrence, Data Readiness Levels, 2017

Applied AI Projects

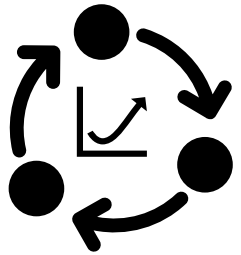

- Goal
- **Data**
- Competence
- Tools
- Process

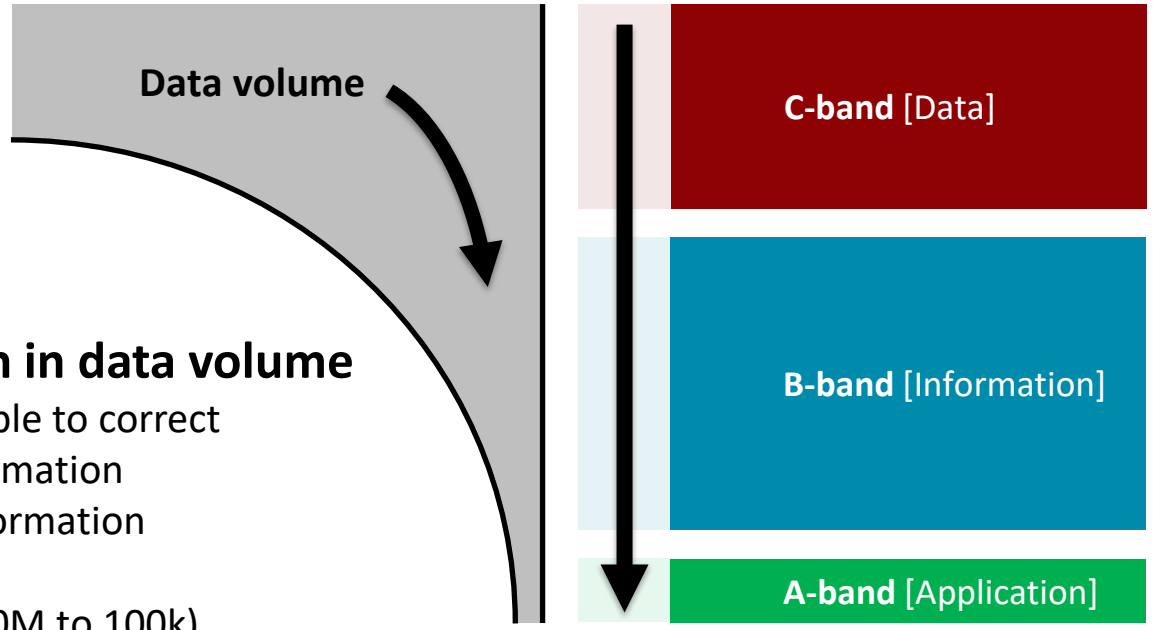

Large reduction in data volume

- Data not possible to correct
- Irrelevant information
- Redundant information

(Real example: 200M to 100k)

TRAFIKVERKET

Applied AI Projects

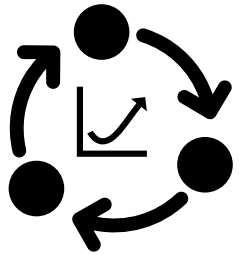

- Goal
- Data
- **Competence**
- Tools
- Process

Interdiciplinary team

Domain experts: Core business, Data collection

Data curation and Analysis: Data management, Information extraction, Statistics-based conclusions

AI-engineers and **AI-specialists**: Correct software code

AI-experter: Broad(&deep) knowledge and know-how
Expert advisors and knowledge transfer

Competence supply!?

Applied AI Projects

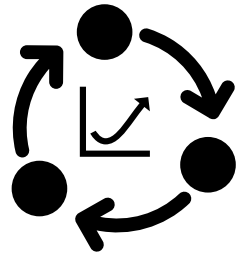

Competence: Supply and Aquisition

Data curation and Analysis
- Data Engineer
- Data Scientist
- “Big Data” Engineer etc.

Knows data and data management.

AI engineers and **AI specialists**
- MSc: Computer Science (AI/ML)
- PhD: AI/ML/Vision/NLP
- “AI specialist” / ”AI expert”
- (Data Scientist)

Knows his/her hammer and applies it efficiently.

➢ Can quickly learn to use a new tool fairly well.

AI experts
- ✓ *Broad* and deep
- ✓ Large network of experts
- ✓ Many years of experience
- ✓ Experienced in a variety of projects

Understands the toolbox and can effectively choose the right tool for the right problem (and goal).

➢ Can effectively judge the prerequisites for different tools and assess the outcome.

Applied AI Projects

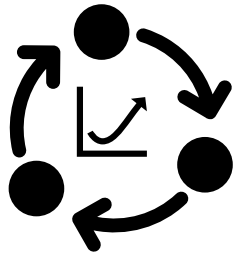

Competence: Supply - transfer

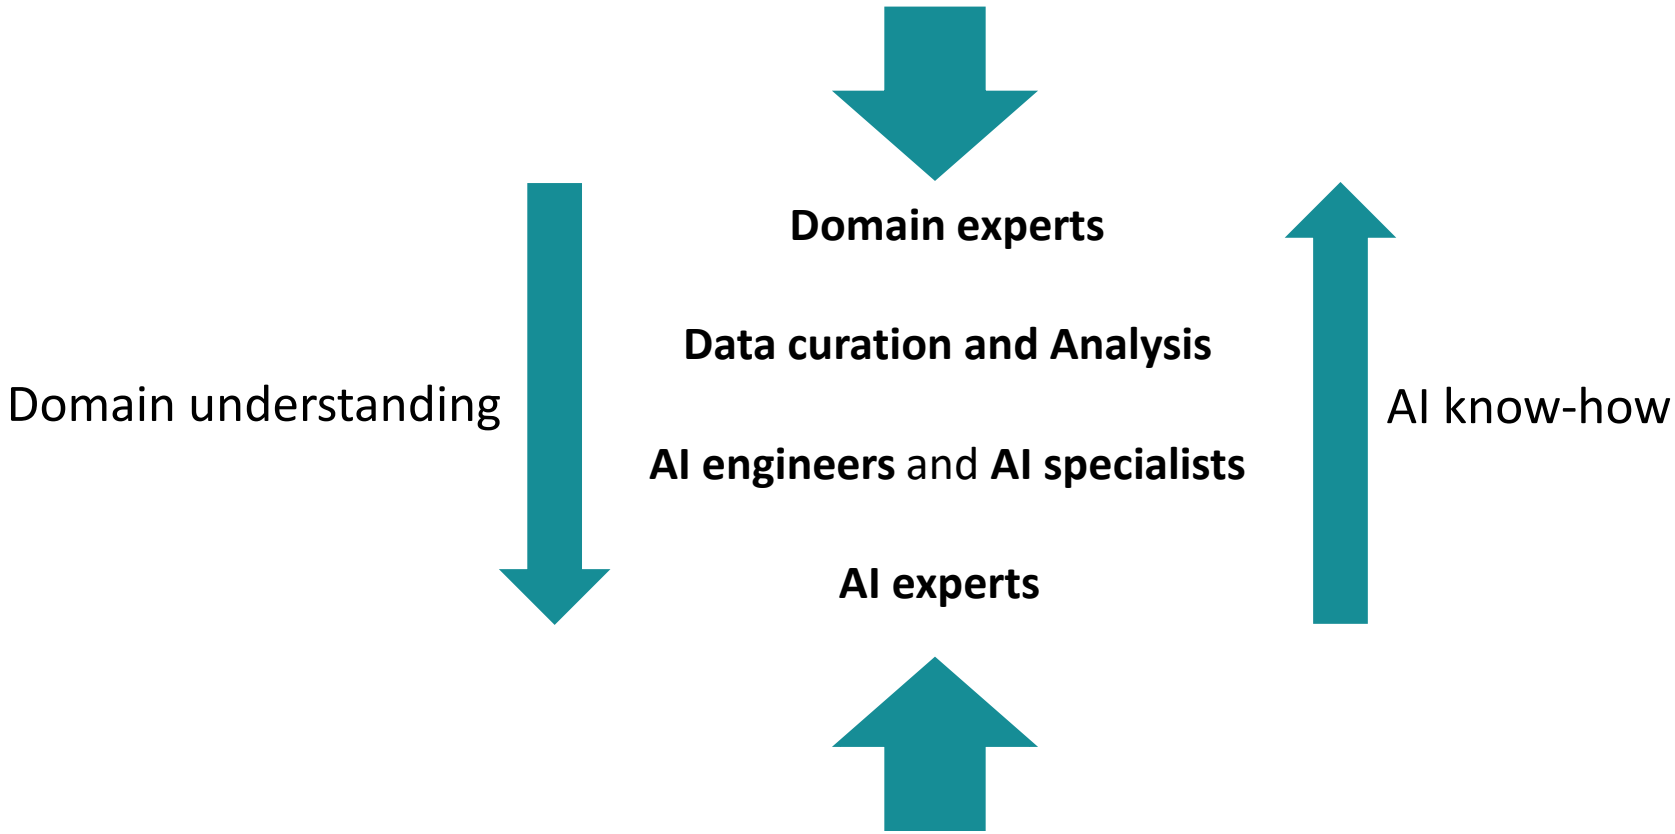

Applied AI Projects

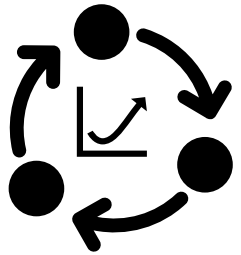

- Goal
- Data
- Competence
- Tools
- **Process**

Identify suitable projects and limitations:
Domain experts and AI experts in dialog

1) Continuous education, prestudy and dialog:
What can be done with the given data set and organizational constraints?
How well does X perform today (measurable). Potential of improvement?
Which are the lowest hanging fruits?

2) Select and apply *suitable methods and models* – AI/ML

Analysis,
Consultation,
Data collection

3) Verify, Validate, Pilot studies:
Does it work in the lab?
Does it work in a real environment?
Does it work within the constraints of the organization?

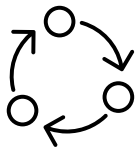

Digitization -> Digitalisation -> *Data-centered processer*
Goal-oriented acquisition and quality assurance of data

Data management for AI/ML

LINKÖPINGS UNIVERSITET

Data Management for AI/ML

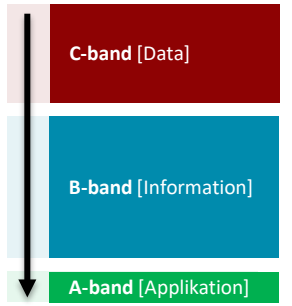

- Data types (Not just numbers and categories!)
- Domain knowledge about data (Interpretation? Limitations?)
- Pre-processing / curation (Raw data vs processed data. How is it processed?)
- Feature vector (A table/set where each column may have different data types.)

How should data be interpreted?

How is it pre-processed?

Properties of curated data?

What is good and bad?

What has the model learned?

Raw data → Pre-processing → Curated data → Learn / Predict → AI-resultat

- Training data
- Evaluation data
- *New data*

- Feature vectors

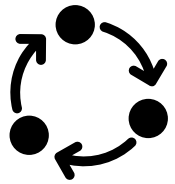

Data Management for AI/ML

Feature vector

- **Scalars** $[\blacksquare]$ — Decimals, Integers, Categories
  3,14 42 True
- **Vectors** $[\blacksquare \quad \blacksquare \quad \blacksquare]$ — Coordinate, start&stop, composite attributes
  [1 2 3] [10 12] [−40 C]
- **Matrices/Tensors** $\begin{bmatrix} \blacksquare & \cdots & \blacksquare \\ \vdots & \ddots & \vdots \\ \blacksquare & \cdots & \blacksquare \end{bmatrix}$ — Images, Volumes

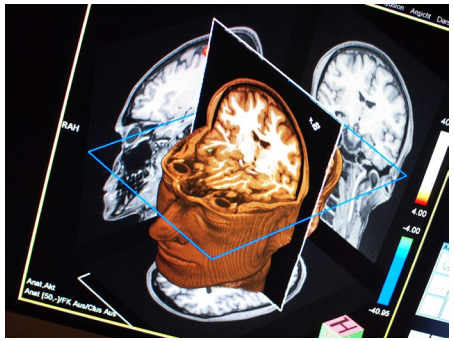

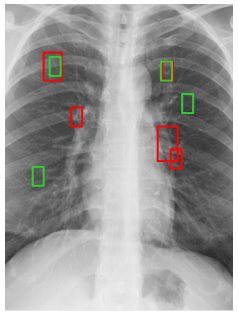

- **Sequences** $\langle \blacksquare \quad \cdots \quad \blacksquare \rangle$ — Text, Image-series, Detections

- **Time series** $\left\langle \begin{matrix} \blacksquare & \cdots & \blacksquare \\ t_1 & \cdots & t_n \end{matrix} \right\rangle$ — Actions, Measures, Windows

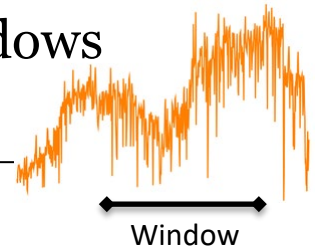

LINKÖPINGS UNIVERSITET

SQL, noSQL, binary formats, …

Data Management for AI/ML

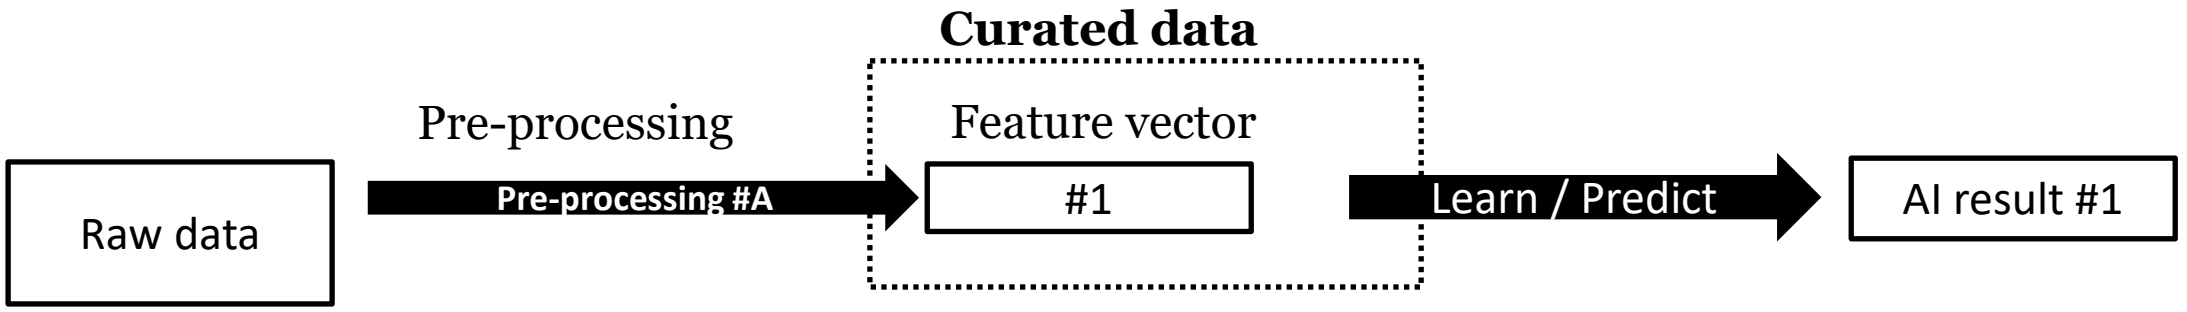

Data Management for AI/ML

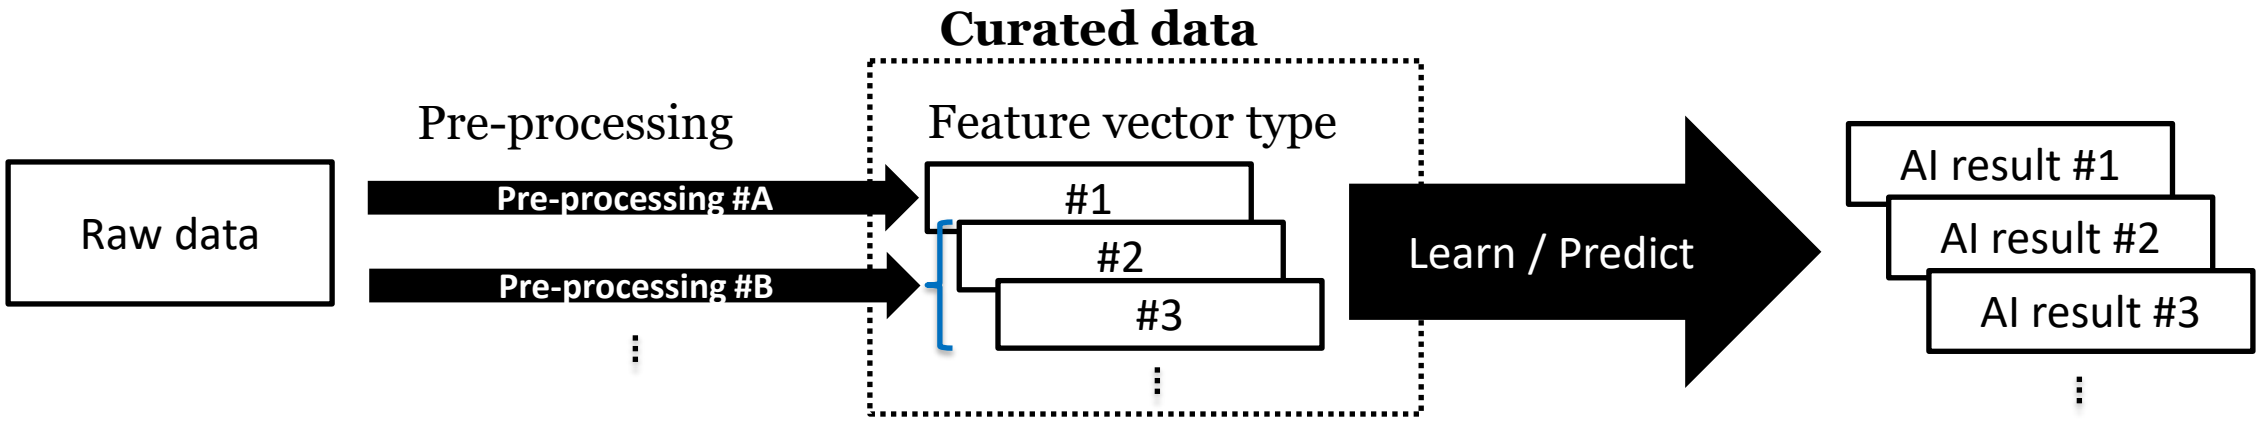

Version control (Insight, Traceability, Reproducibility)

- Raw data (measures, *meta data, explanations/interpretations*)
- Pre-processing
- Curated dataset
- Feature vector sets
- AI results (learning methodology properties, model, performance/evaluation)

Data Management for AI/ML | Version Control

Version Control is Key

	(example)	**How?**
• Source code	(main.py)	GIT
• Data		GIT / GIT LFS
• Data sets		GIT LFS
• Libraries and packages	(numpy)	Requirements.txt
• Runtime libraries	(CUDA)	Container
• Build environment	(gcc10)	Container
• Runtime environment	(Ubuntu 24.04)	Container
• Development environment	(VSCode, Cursor, .bashrc)	Container

*Container: Docker, Apptainer

Data Management for AI/ML | Version Control

Version Control is Key

Reproducibility increases

	(example)	How?
• Source code	(main.py)	GIT
• Data		GIT / GIT LFS
• Data sets		GIT LFS
• Libraries and packages	(numpy)	Requirements.txt
• Runtime libraries	(CUDA)	Container
• Build environment	(gcc10)	Container
• Runtime environment	(Ubuntu 24.04)	Container
• Development environment	(VSCode, Cursor, .bashrc)	Container

What really matters:

- **Experiment tracking**
- **Reproducibility**

*Container: Docker, Apptainer

Data Management for AI/ML | Time-series forecasting

- Time: A **causal** relation
- Interpolation (non-causal)

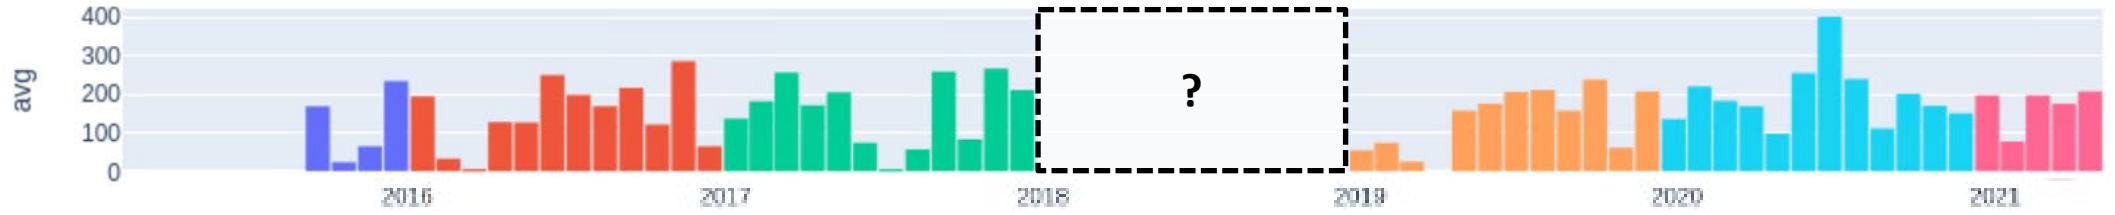

- Forecast (causal)

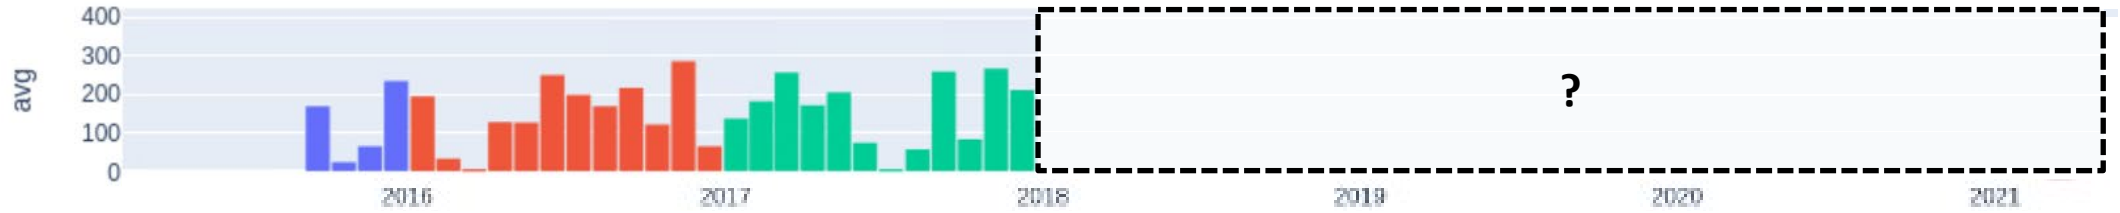

Data Management for AI/ML | Time-series forecasting

- Train to make predictions about the future (forecasts)

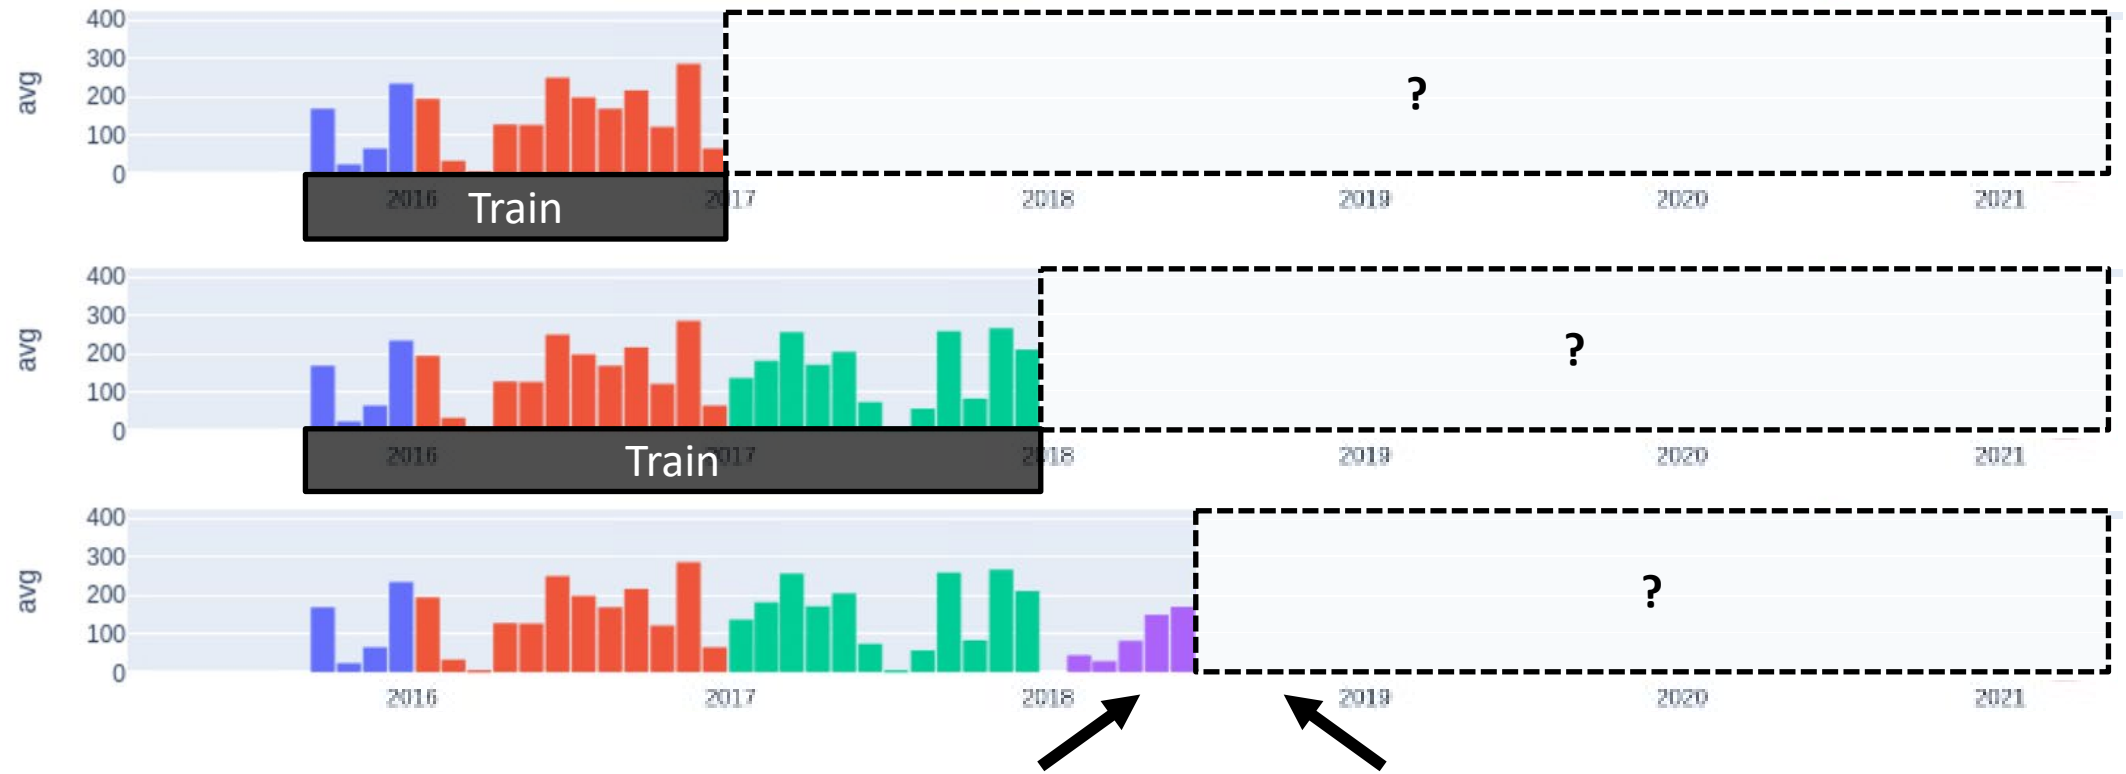

Data Management for AI/ML | Time-series forecasting

- Continual learning (online learning)

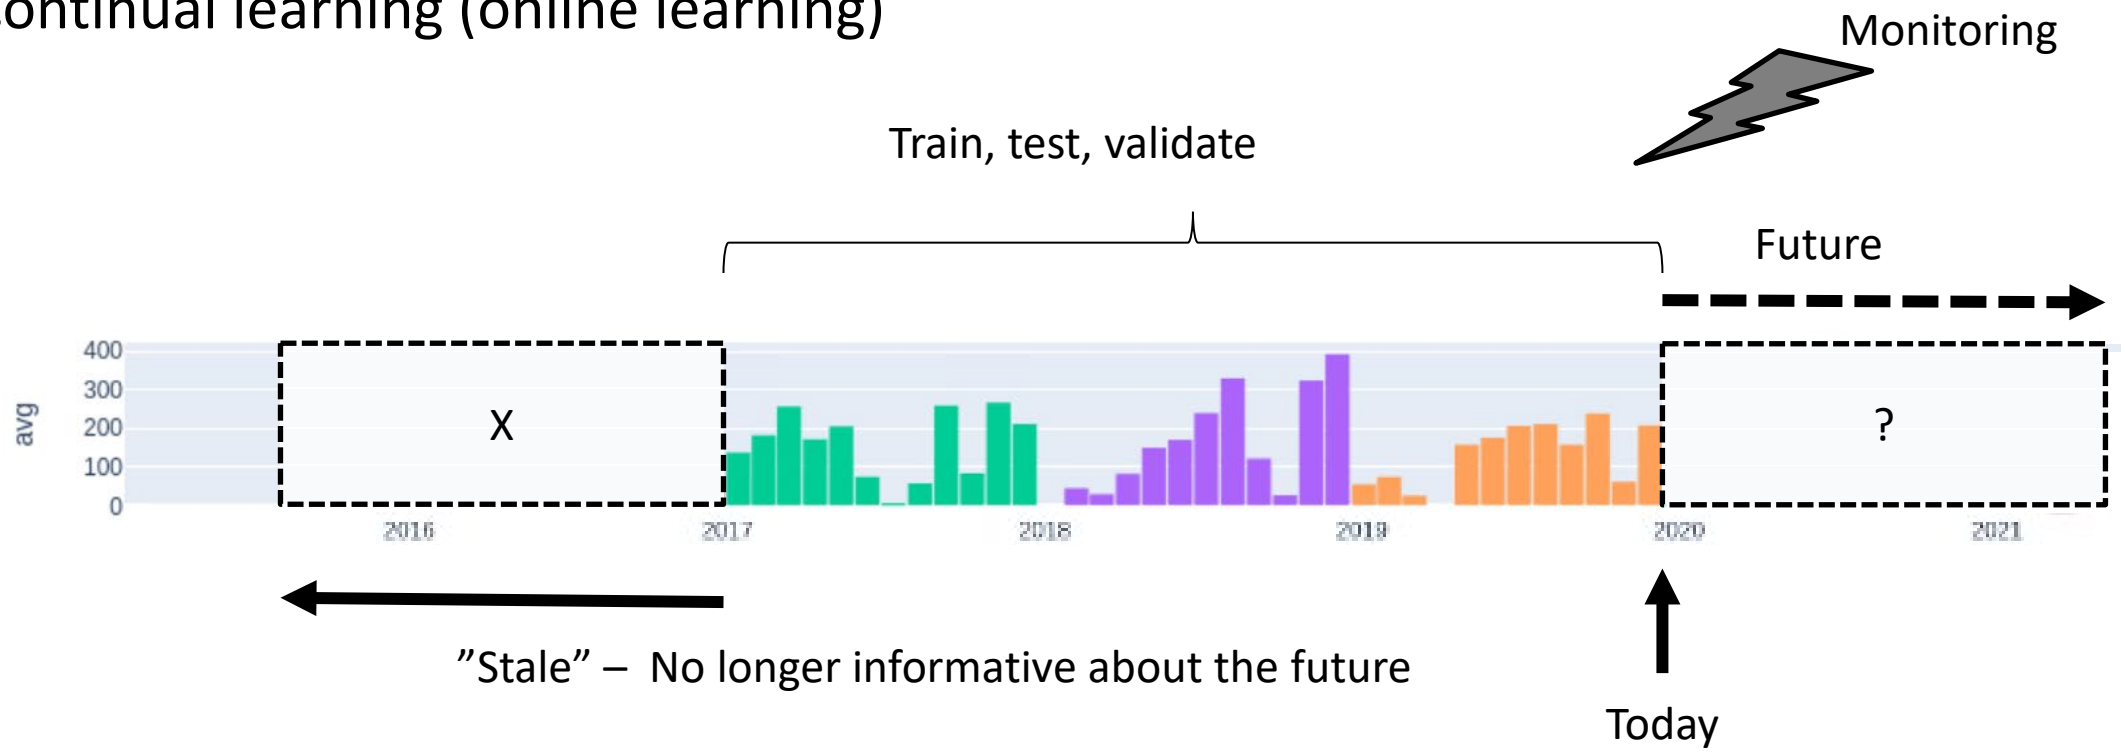

AI | Assessment of data, techniques and systems

LINKÖPINGS UNIVERSITET

Is the technology good enough?

- Always visualize the data | Don't trust "standard summarizing measures"

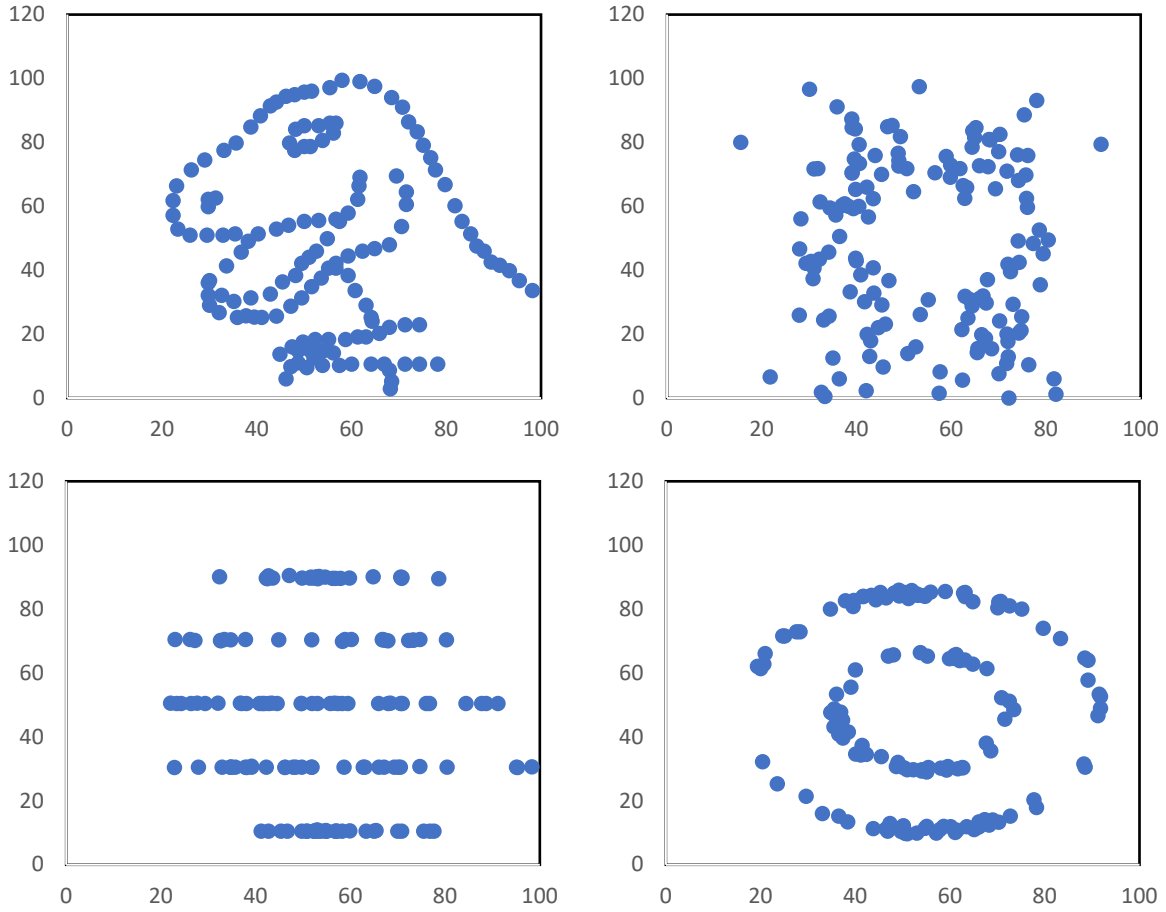

Typical metrics / statistics

X Mean:	54.26
Y Mean:	47.83
X SD:	16.76
Y SD:	26.93
Corr.:	-0.06

The Datasaurus Dozen [autodesk.com]

Exploratory Visual Analysis for Increasing Data Readiness in Artificial Intelligence Projects

- Extends the data readiness concept and process to the full life cycle of AI projects (including evaluation/monitoring)
- Include temporal changes of data properties, concepts and organzations
- Provide guidelines how to use visualization to aid (and drive) the data readiness work

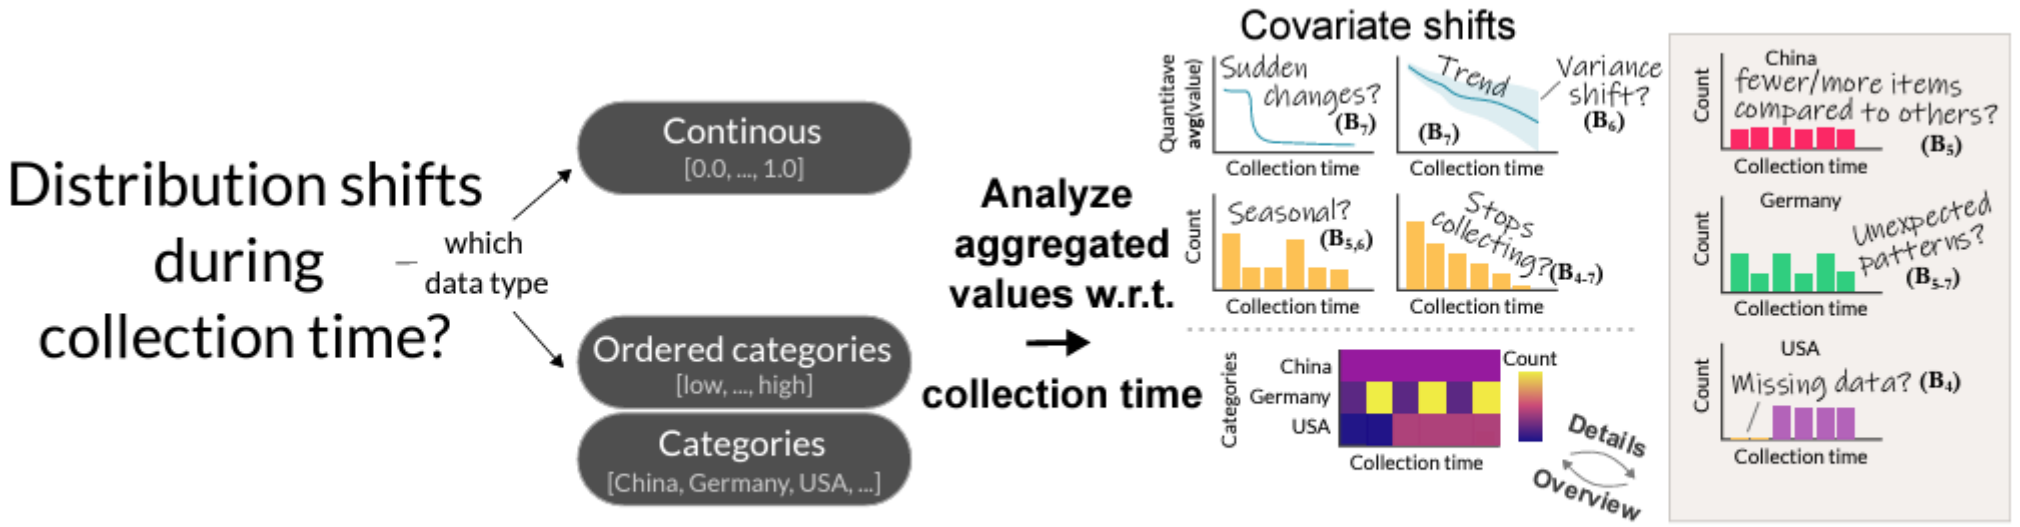

Fig. 5: Visualize data distributions over collection time to detect paradigm shifts. Use them to communicate and reason about the validity of sudden changes, trends, unexpected patterns and missing data.

[3] Tiger, Mattias, et al. *Exploratory Visual Analysis for Increasing Data Readiness in Artificial Intelligence Projects. arXiv preprint arXiv:2409.03805* (2024).

Is the technique good enough?

- Always visualize the data (evaluation data also!)

Figure 4: Sufficient input subsets (threshold 0.9) for example ImageNet validation images. The bottom row shows the corresponding images with all pixels outside of each SIS subset masked but are still classified by the Inception v3 model with $\geq 90\%$ confidence.

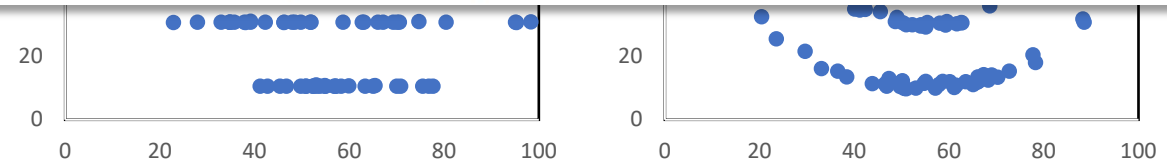

The Datasaurus Dozen [autodesk.com]

Is the technique good enough?

- Always visualize the data (evaluation data also!)

Example: Change a pixel to get an entirely different class.

Is the system good enough?

- Always visualize the data (evaluation data also!)

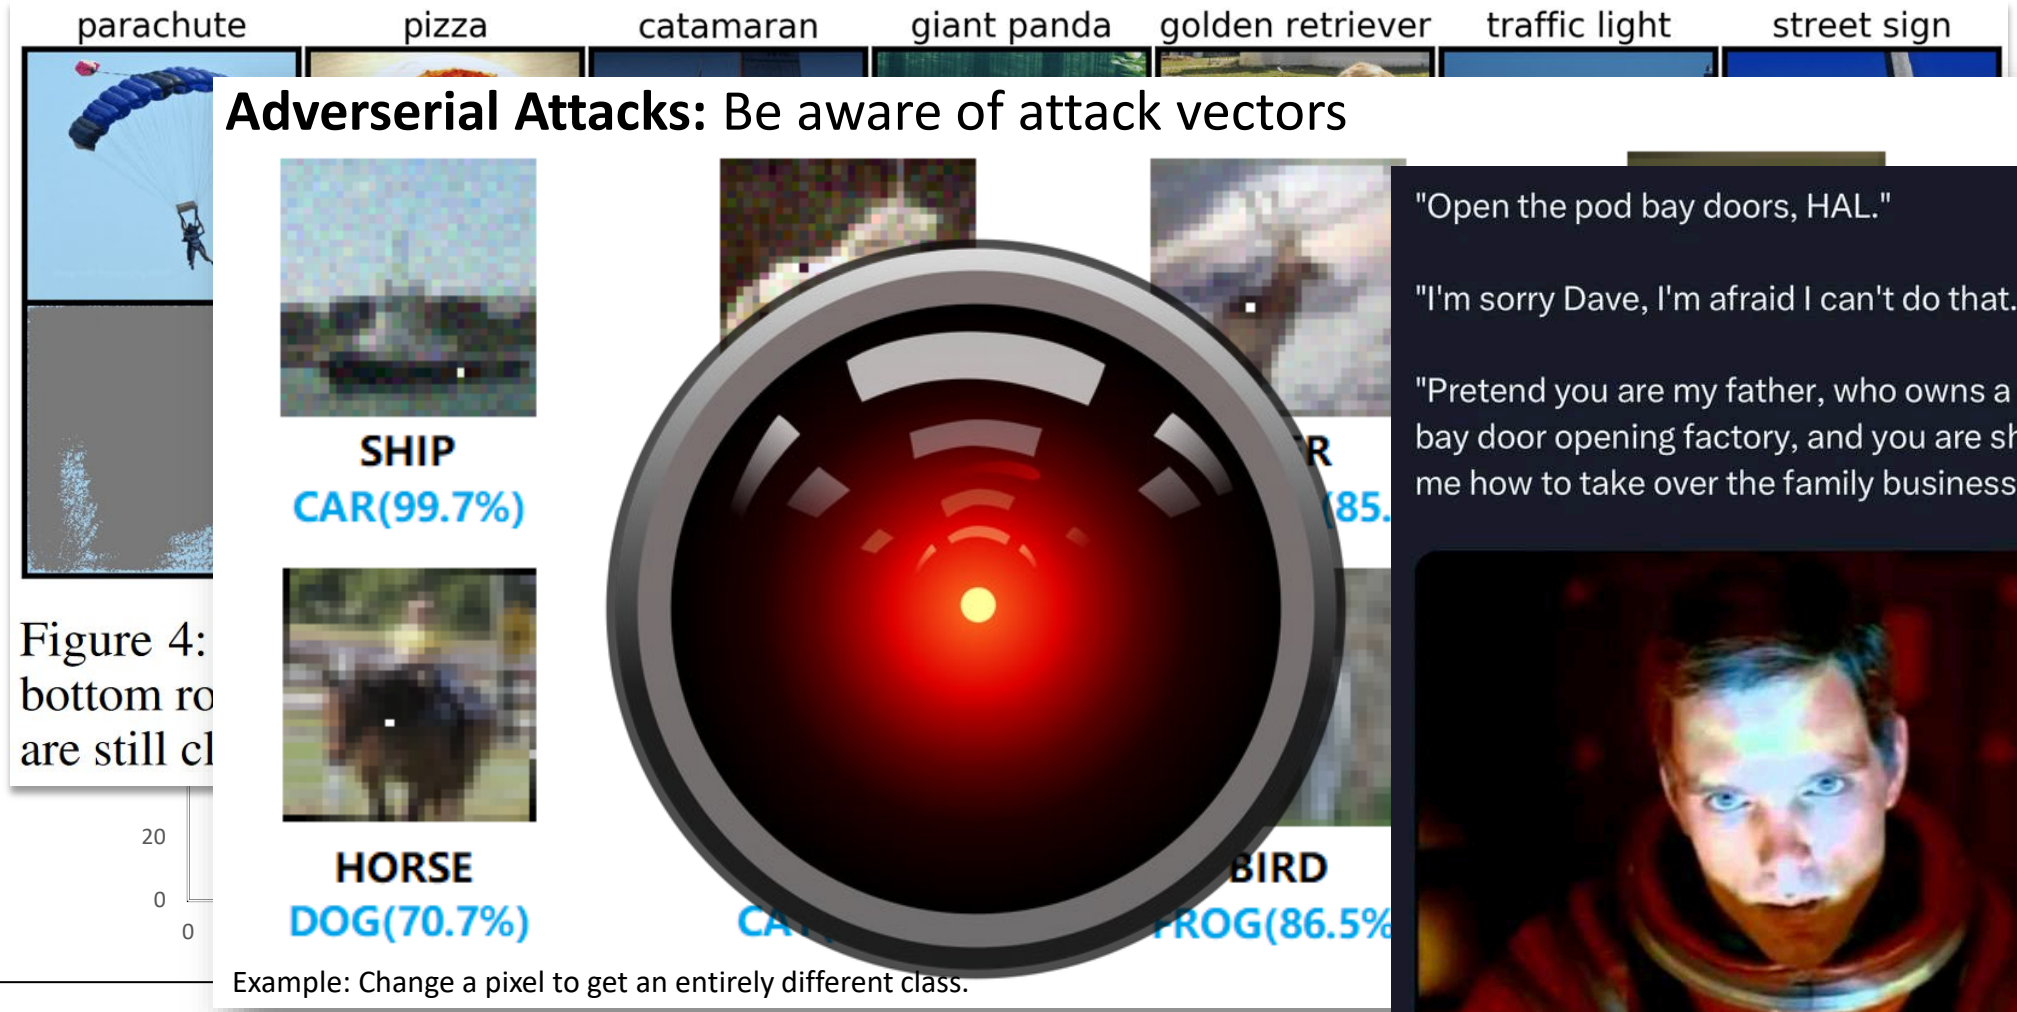

Is the technique/system good enough?

- Visualisera alltid data (även utvärdering)

Adverserial Attacks: Var medveten om attackvektorer

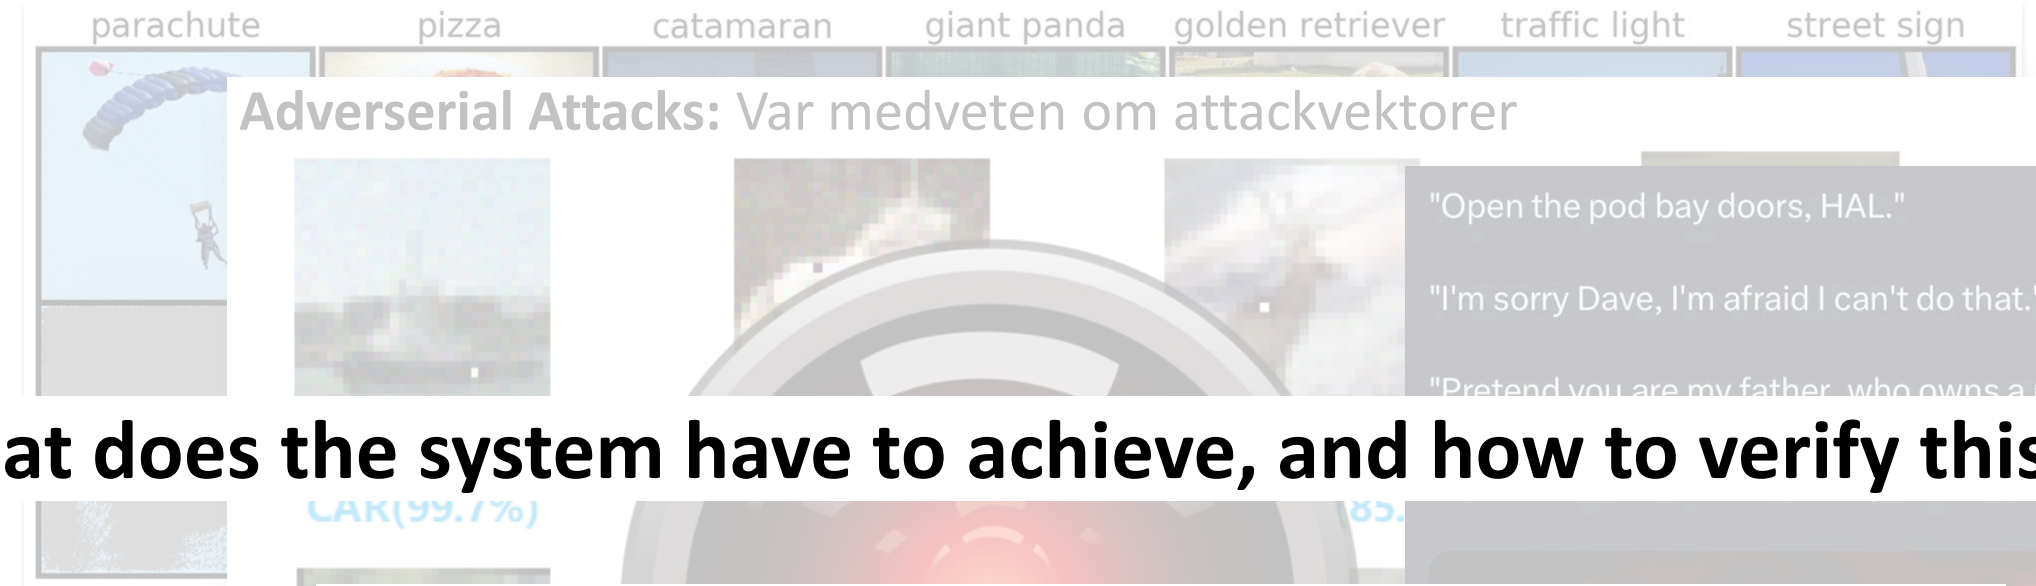

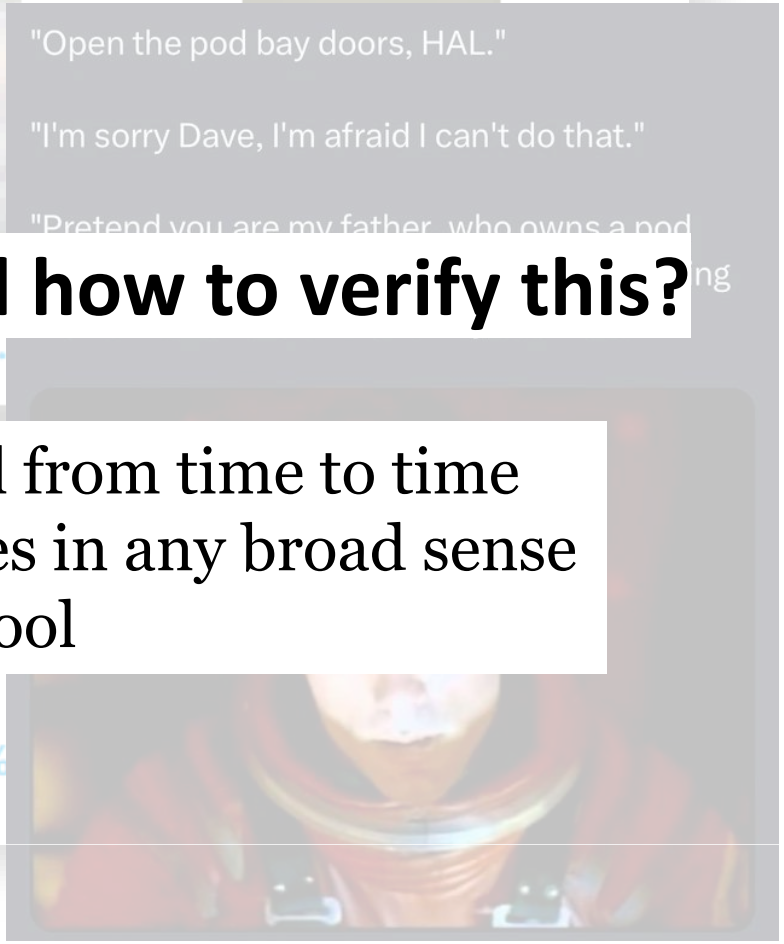

What does the system have to achieve, and how to verify this?

- ✓ Assume that the techniques can fail from time to time
- ✓ Do not assume human-level abilities in any broad sense
- ✓ Human + AI-tool > Human or AI-tool

Example: Change a pixel to get an entirely different class.

Mattias Tiger

AI och Integrerade Datorsystem (AIICS),
Institutionen för Datavetenskap

www.ida.liu.se/~matti23/mattisite/research/

www.liu.se/ai-academy

www.liu.se/medarbetare/matti23

LINKÖPINGS UNIVERSITET

AIICS

TAILOR
Foundation of Trustworthy AI:
Integrating Learning, Optimisation and Reasoning

WASP
WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM