What AI Can Do for You

Kerstin Kleese van Dam
Brookhaven National Laboratory

Is it really AI?

- Often Machine Learning is sold as AI
- Machine Learning is learning from data
- AI is taking decisions and actions based on data
- AI often includes Machine Learning, but also many other components

The Smiling Girl, Andrew W. Mellon Collection, (c) The National Gallery of Art, Washington, D.C.

What is Machine Learning?

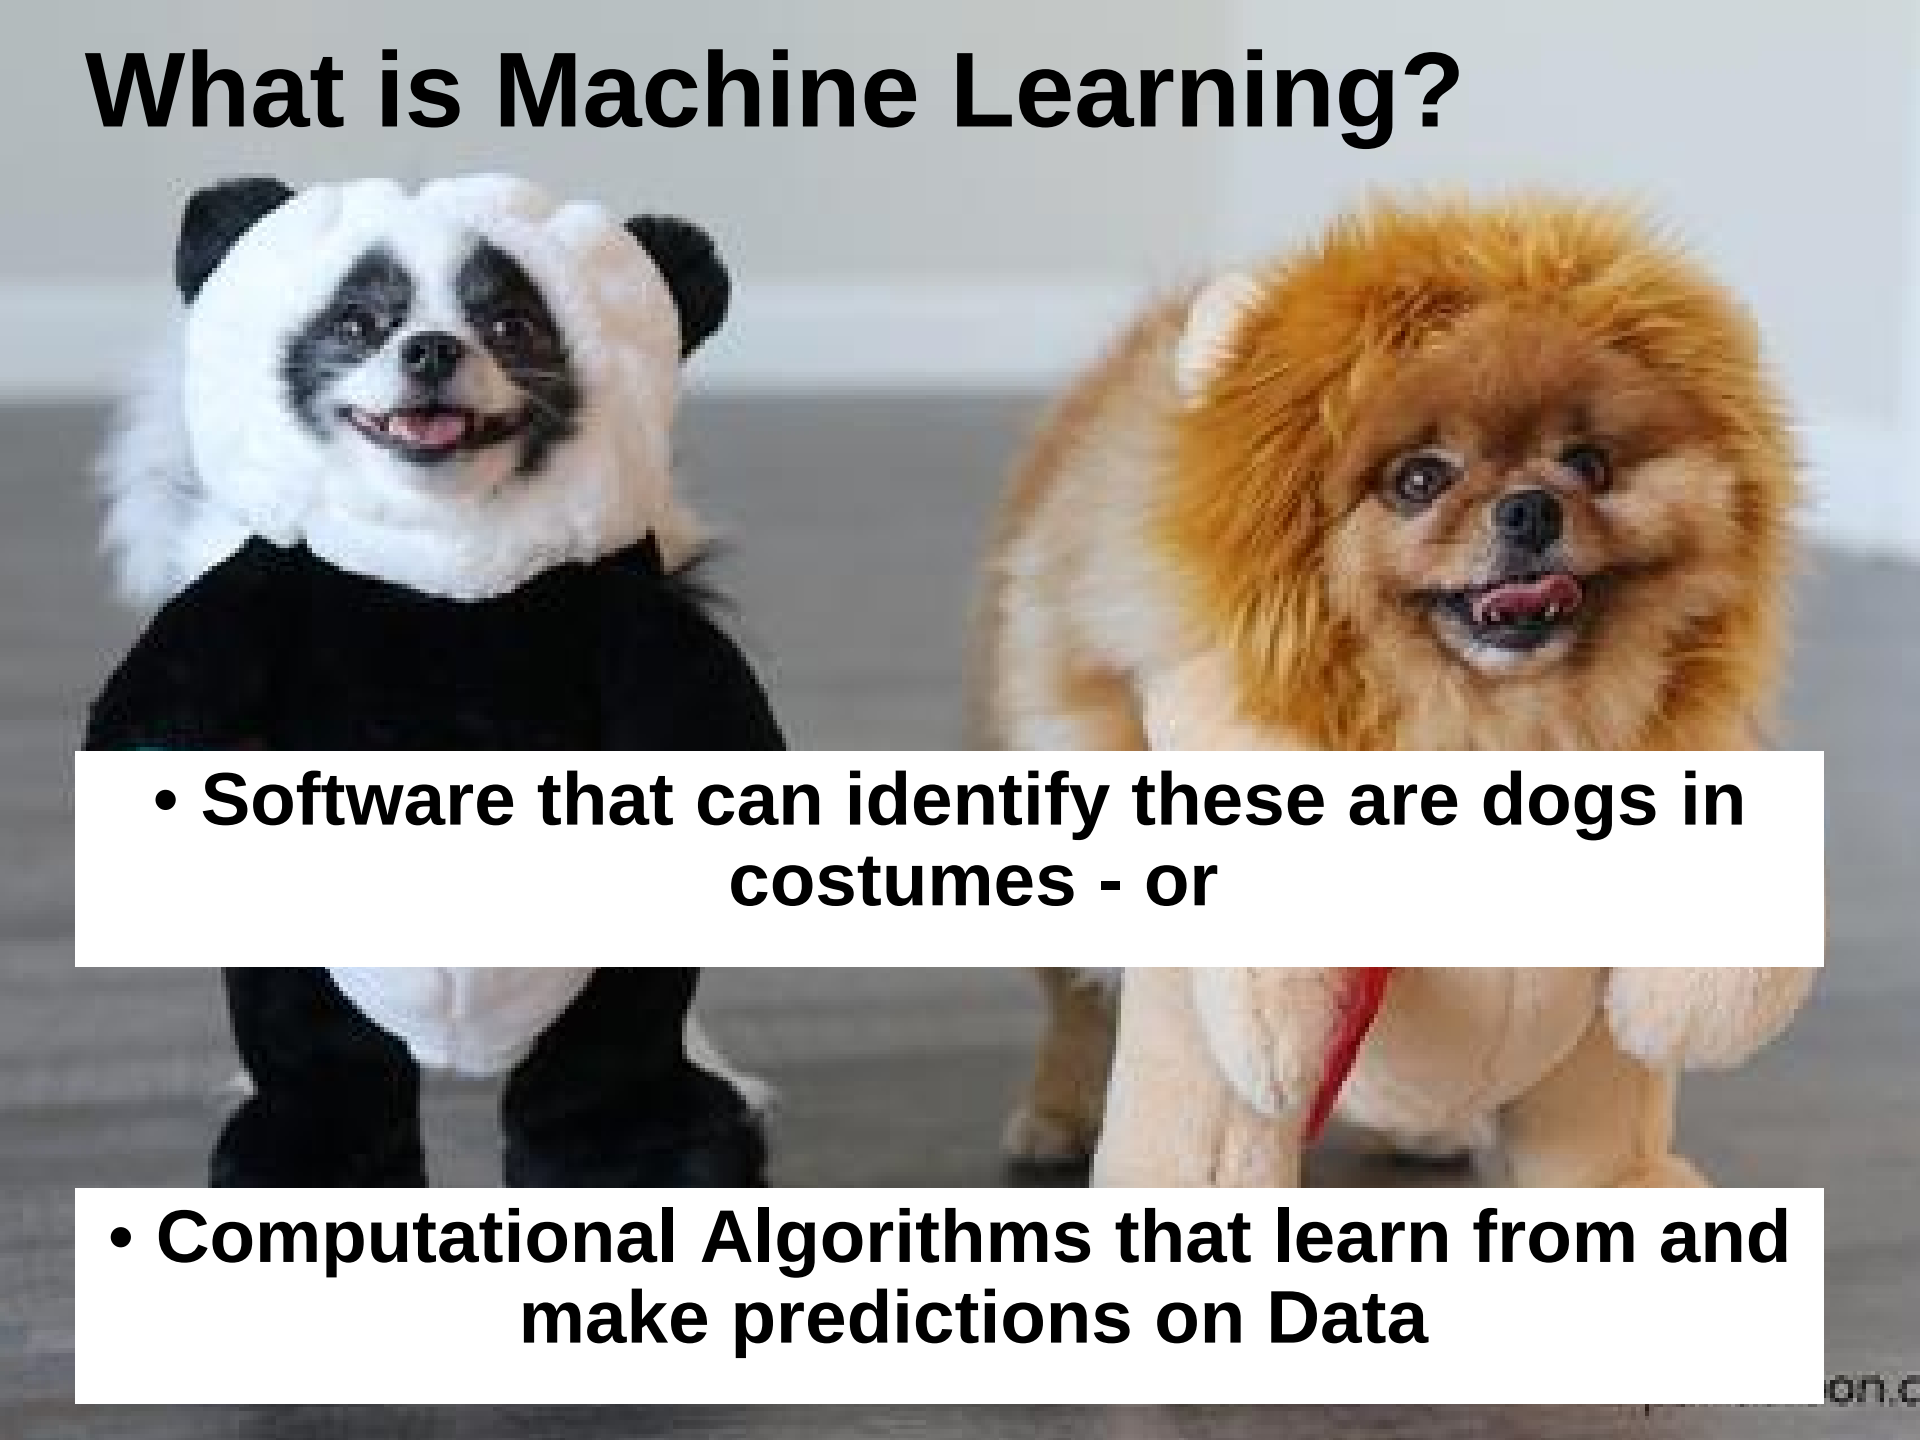

- **Software that can identify these are dogs in costumes - or**

- **Computational Algorithms that learn from and make predictions on Data**

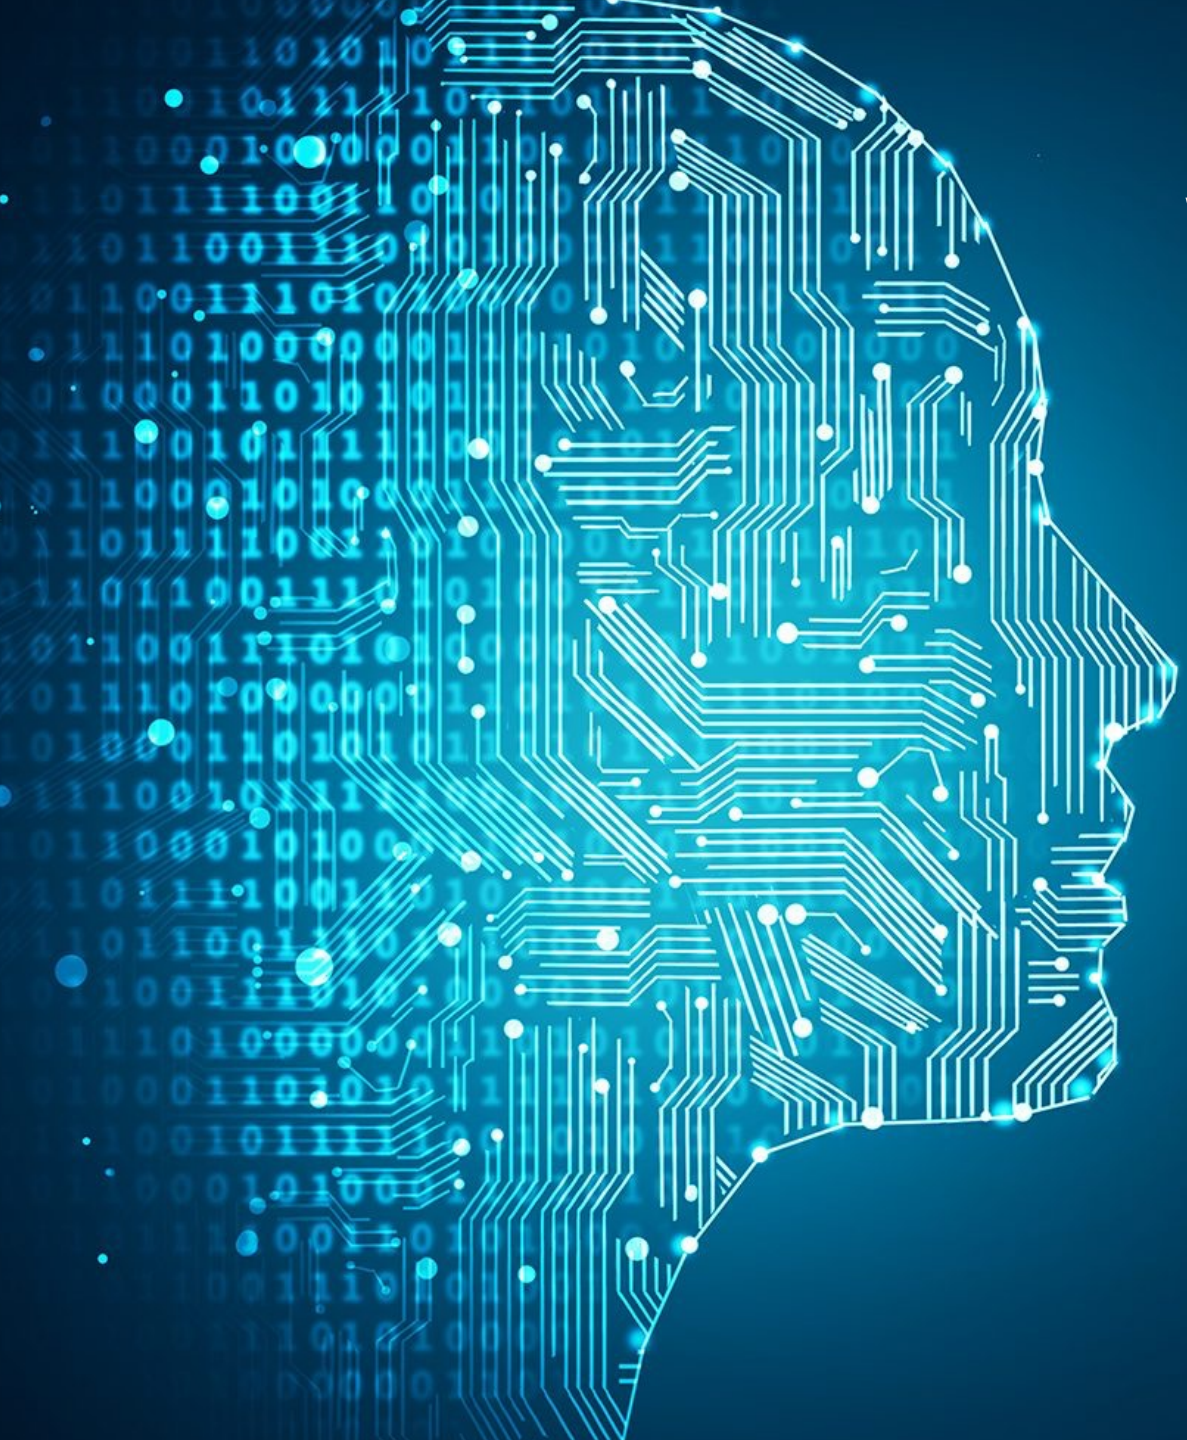

What can Machine Learning do?

Clustering Machine Learning

- Automatically Identifying things and concepts that belong together, ability to distinguish between similar items
- e.g. Data Points, Images, Documents, Sentiments
- Aggregate Heat Kernels, Laplace Beltrami Normalization or Local Density Affinity Transformation are example methods

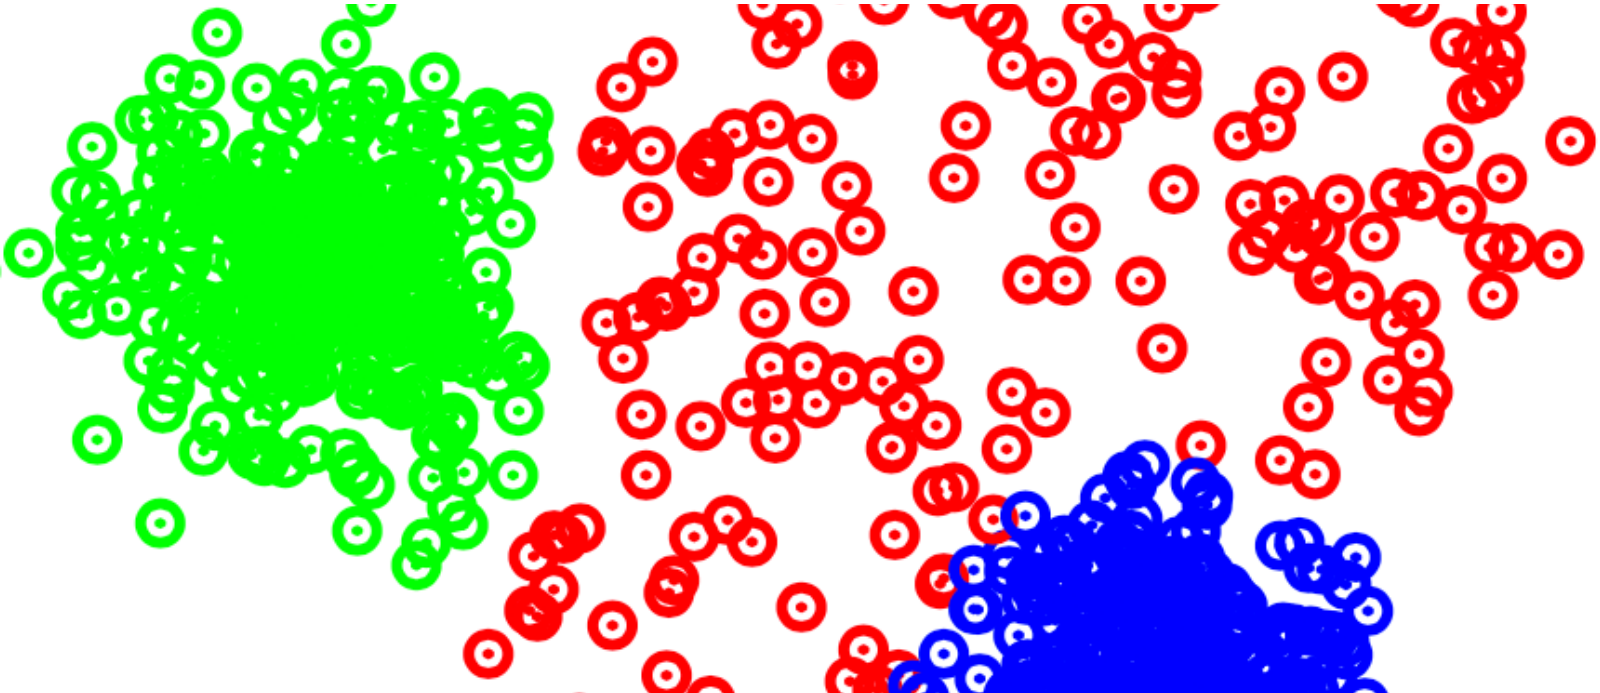

Feature Detection

- Identifying specific features - simple or complex, many or few - in a data collection.
- Often applied in images, finding the same person in different images - a bank robber that may have visited you before as customer and was captured on CCTV
- Multi Cluster / Class Feature Selection (MCFS) or Feature Selection on Data Streams (FSDS) are example methods

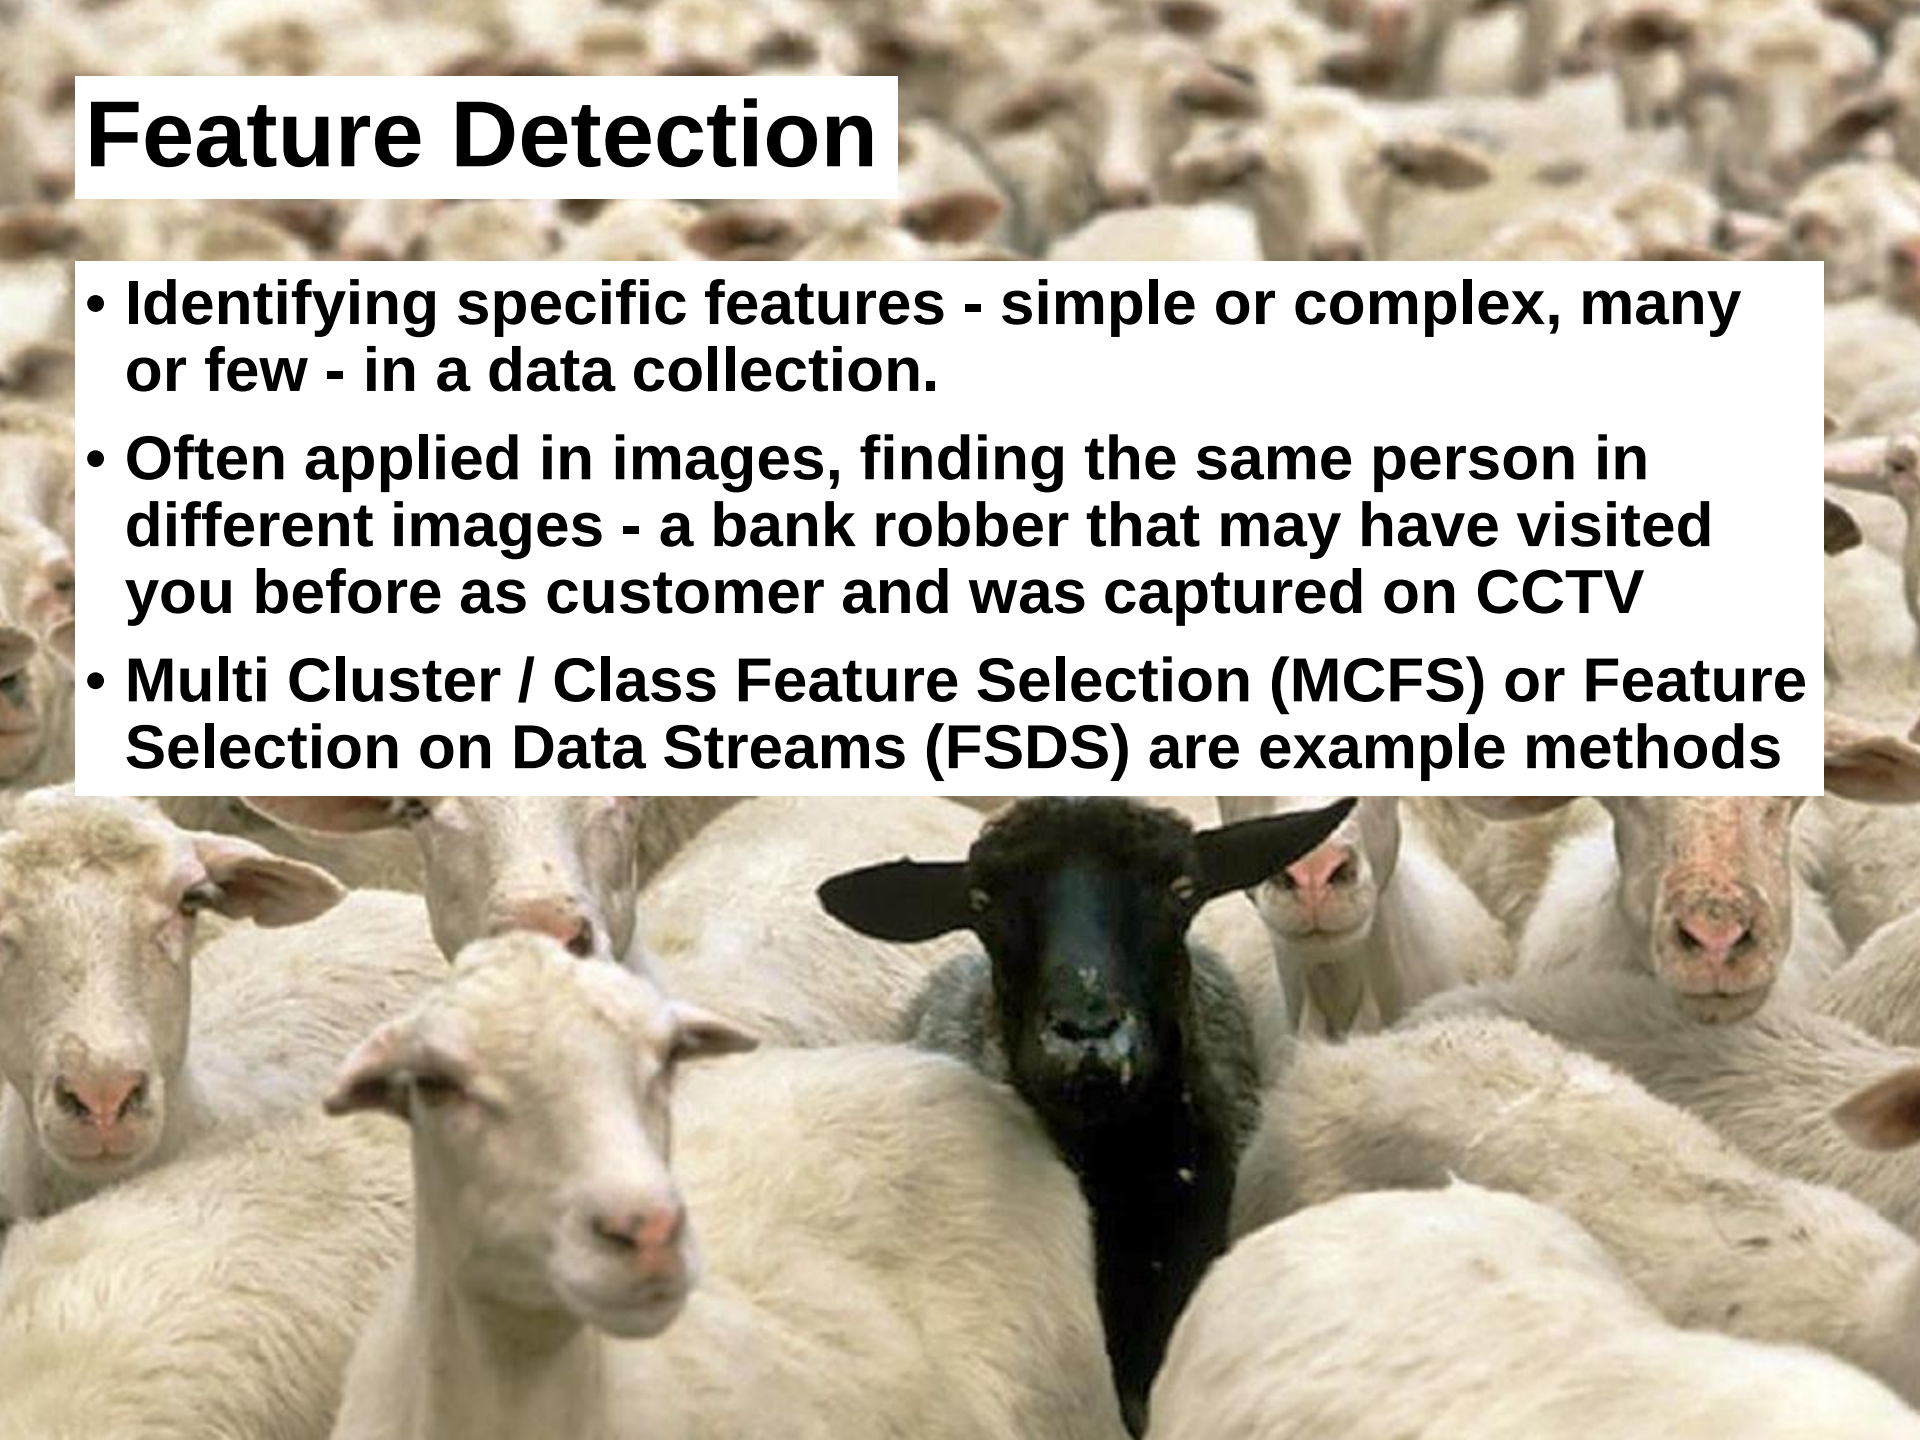

Outlier Detection

- Detecting items that are not like the others, either completely or subtle so.
- Fraud detection is a main application in your area
- Local Outlier Factor or Fermi Density Descriptor are method examples

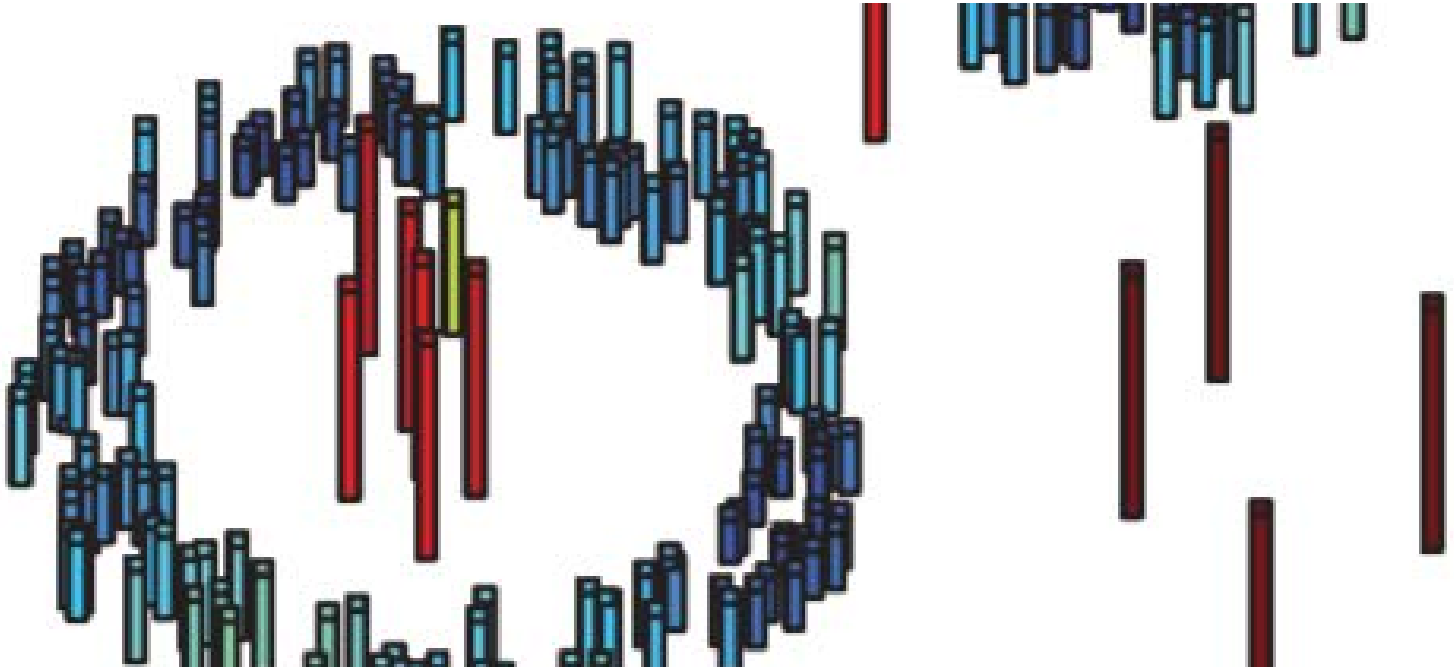

Causal Analysis

- **Identify the reason why the anomaly may have occurred. Can we find signs in the data that would indicate the event could happen? Can we use this knowledge to avoid future occurrences?**
- **Computer fails -> CPU reported overheating -> Other components reported overheating -> likely cause cooling fan broken — Check other system logs to see if others report similar early warning signs.**
- **Lasso or Conditional Granger Causal Inference with two step prior incorporation are example methods**

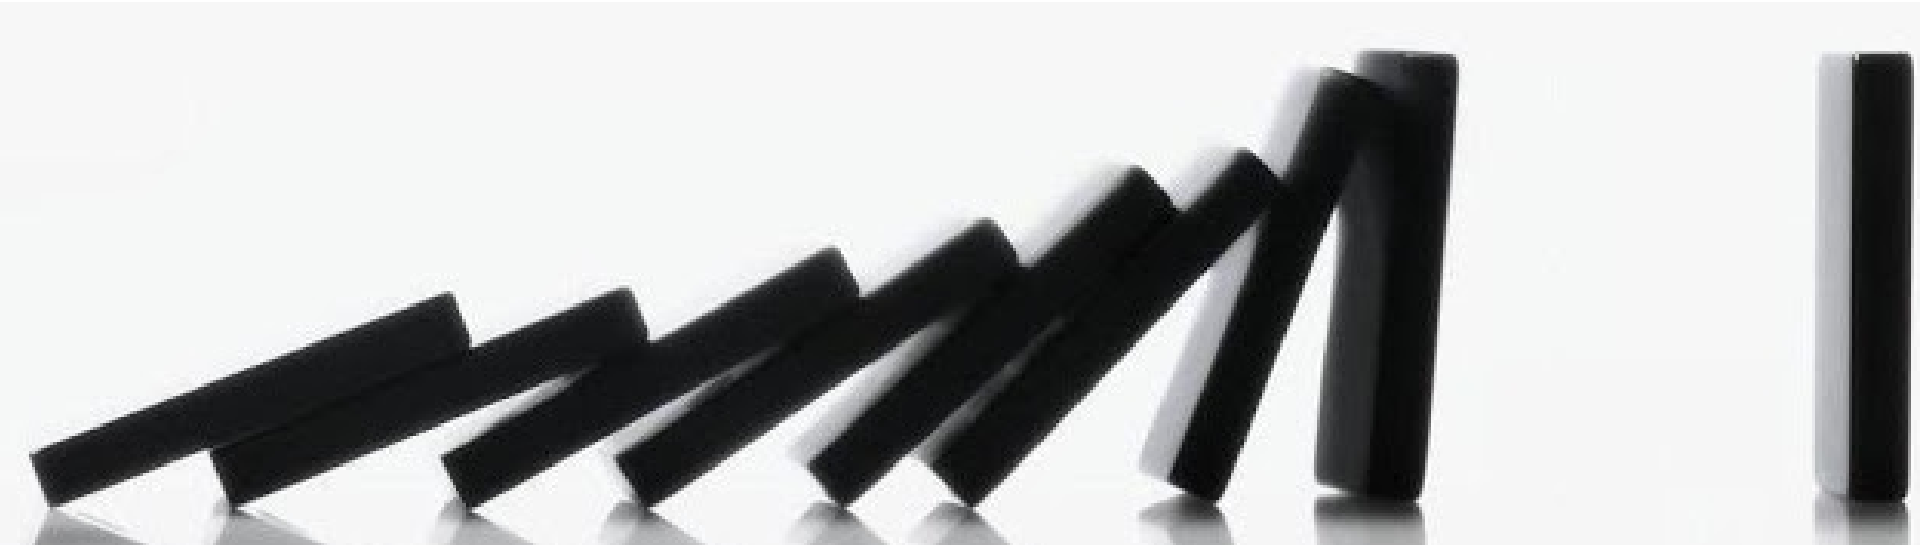

Dimension Reduction

- You have complex data with many features, and want to reduce the number of those features, while retaining as much information as possible.
- Linear Dimensionality Reduction (LDA) or Manifold Learning algorithms such as Diverse Power Iteration Embedding (DPIE) are example methods.

Multi-Source Data Integration and Prediction

- We often have to consider many different types of data to make a prediction, ML can help to integrate and find pattern across data.
- Alzheimers - Algorithms were trained on magnetic resonance imaging, brain phenotyping and patient health records. We achieved nearly 100% Alzheimer's detection accuracy and 83% prediction accuracy for early-stage Alzheimer's.
- Machine learning analytics extracted key features from the data to classify subjects with Alzheimer's disease. The combination of the different data sets significantly improved the prediction accuracy.

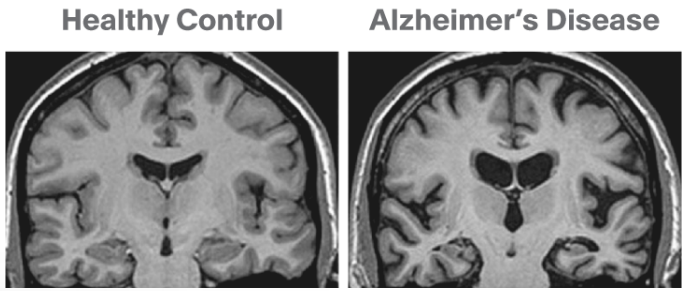

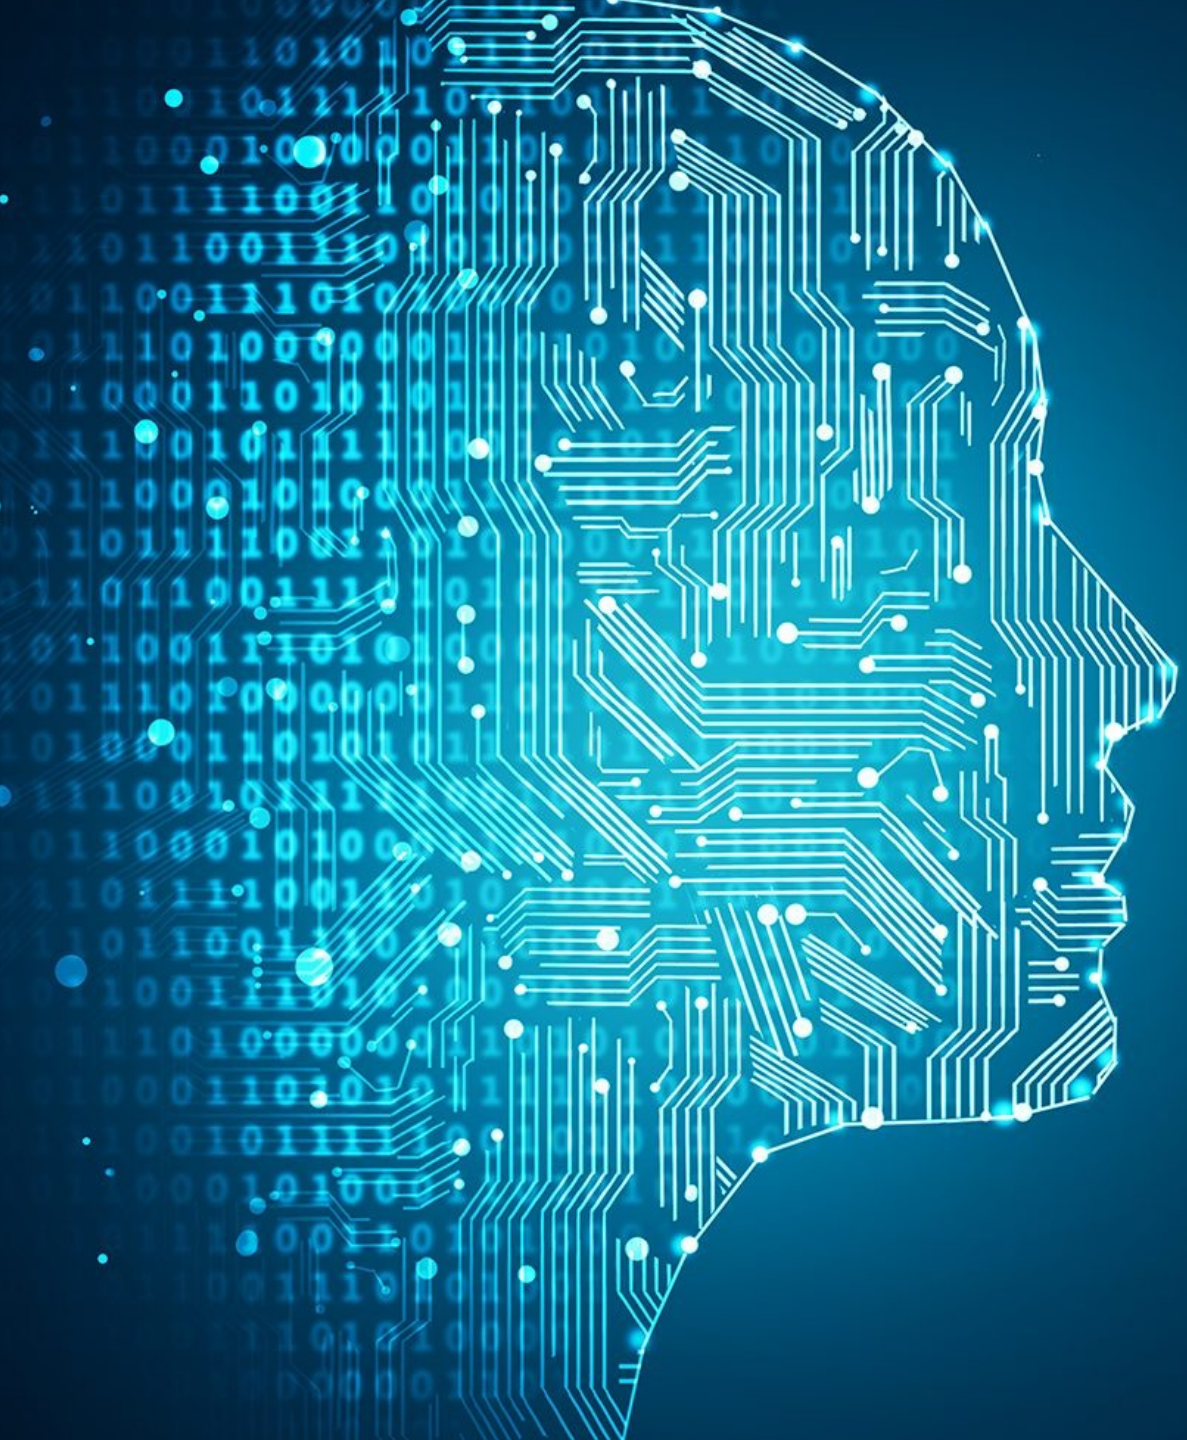

How can machines learn?

Supervised Learning

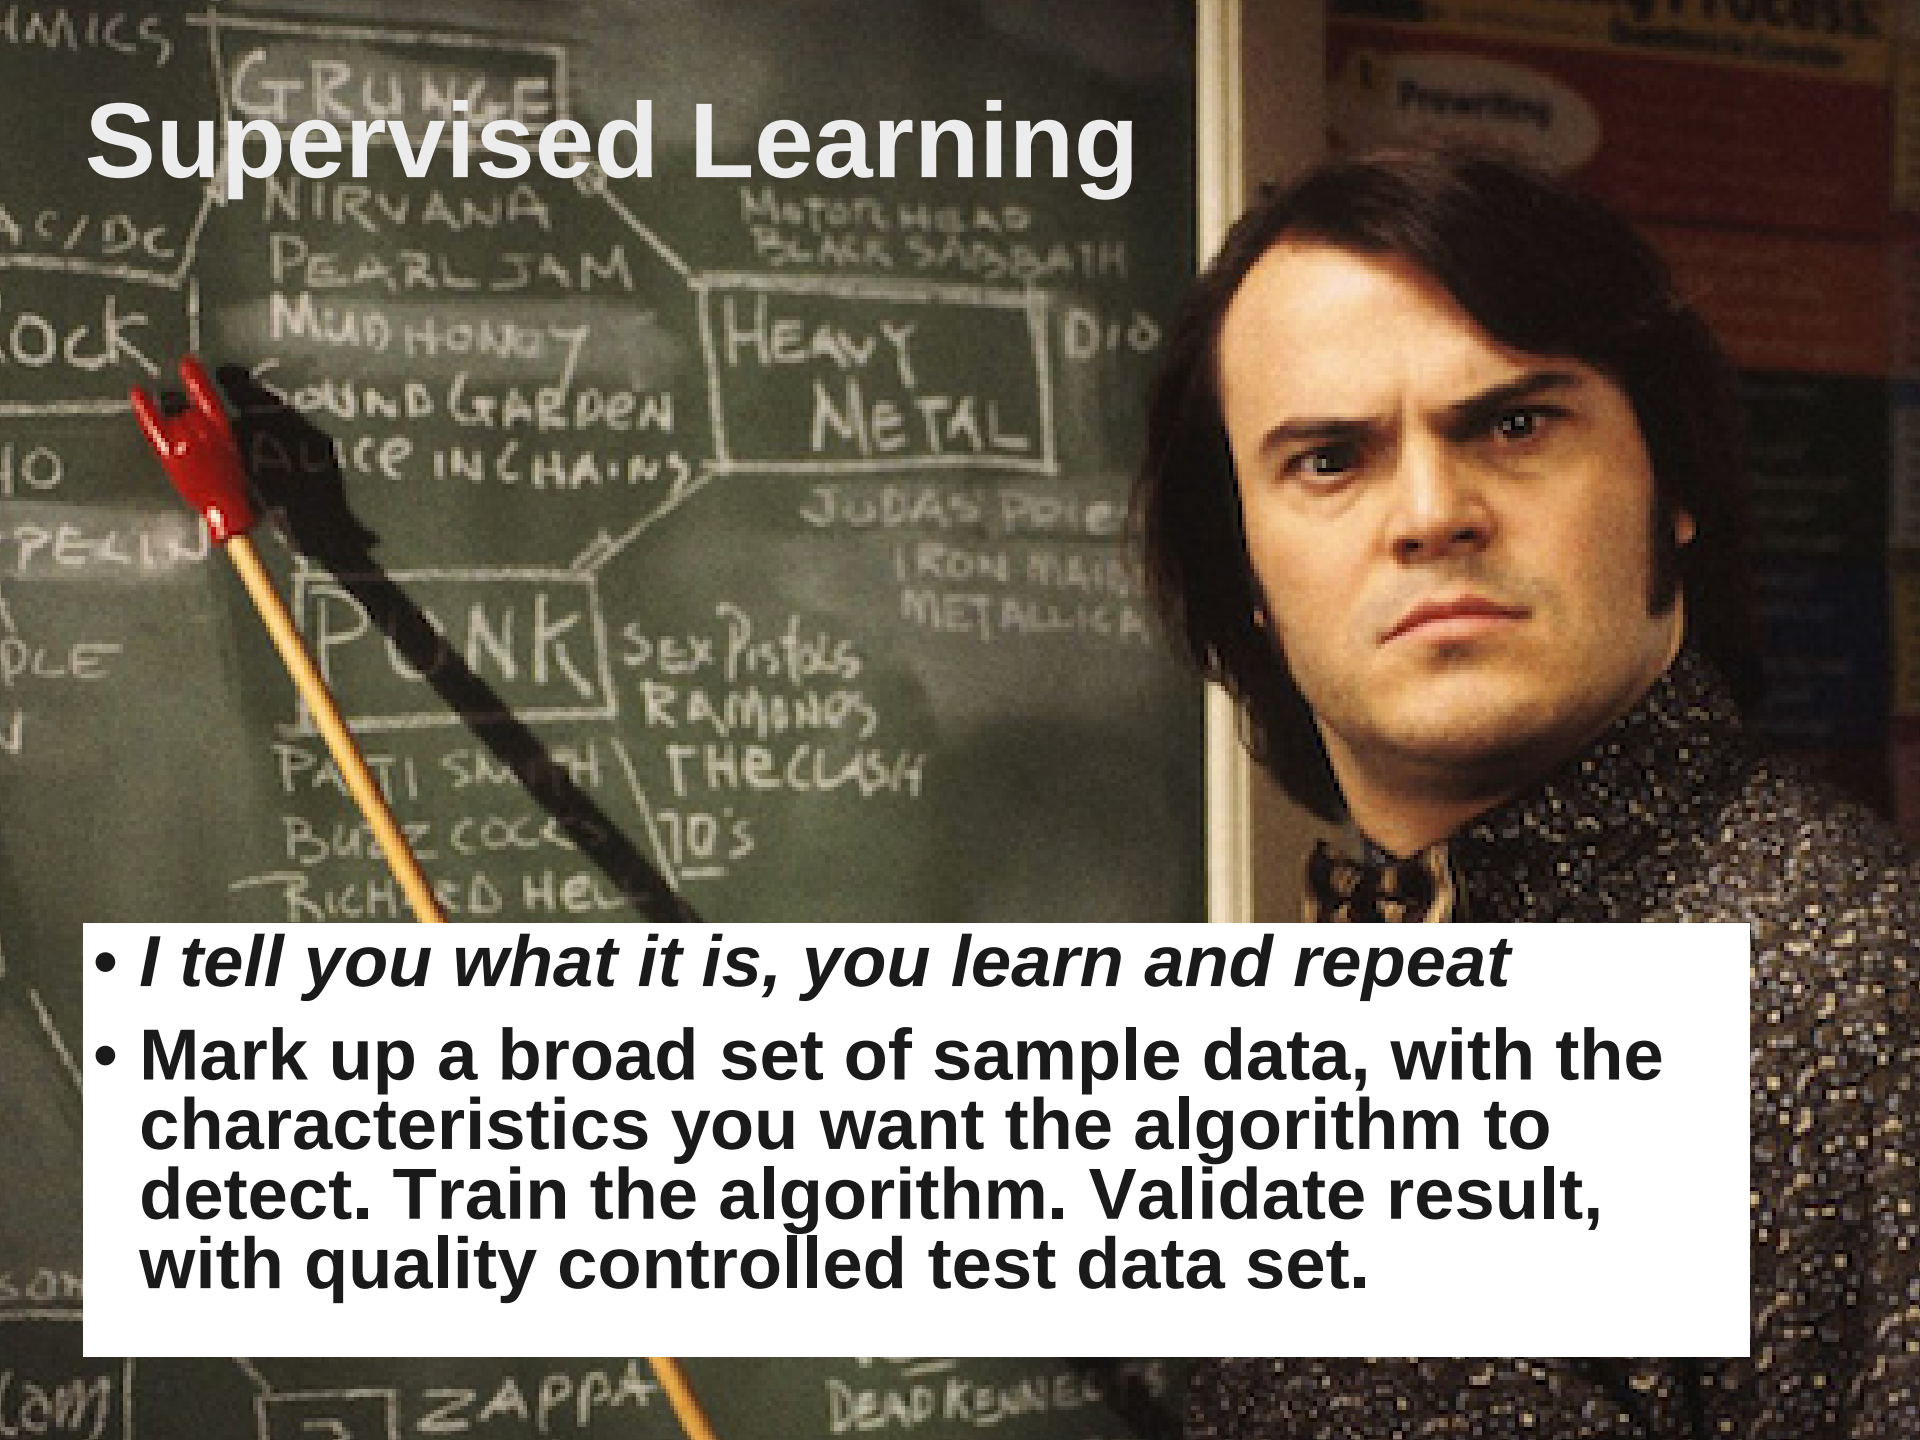

- *I tell you what it is, you learn and repeat*
- **Mark up a broad set of sample data, with the characteristics you want the algorithm to detect. Train the algorithm. Validate result, with quality controlled test data set.**

Unsupervised Learning

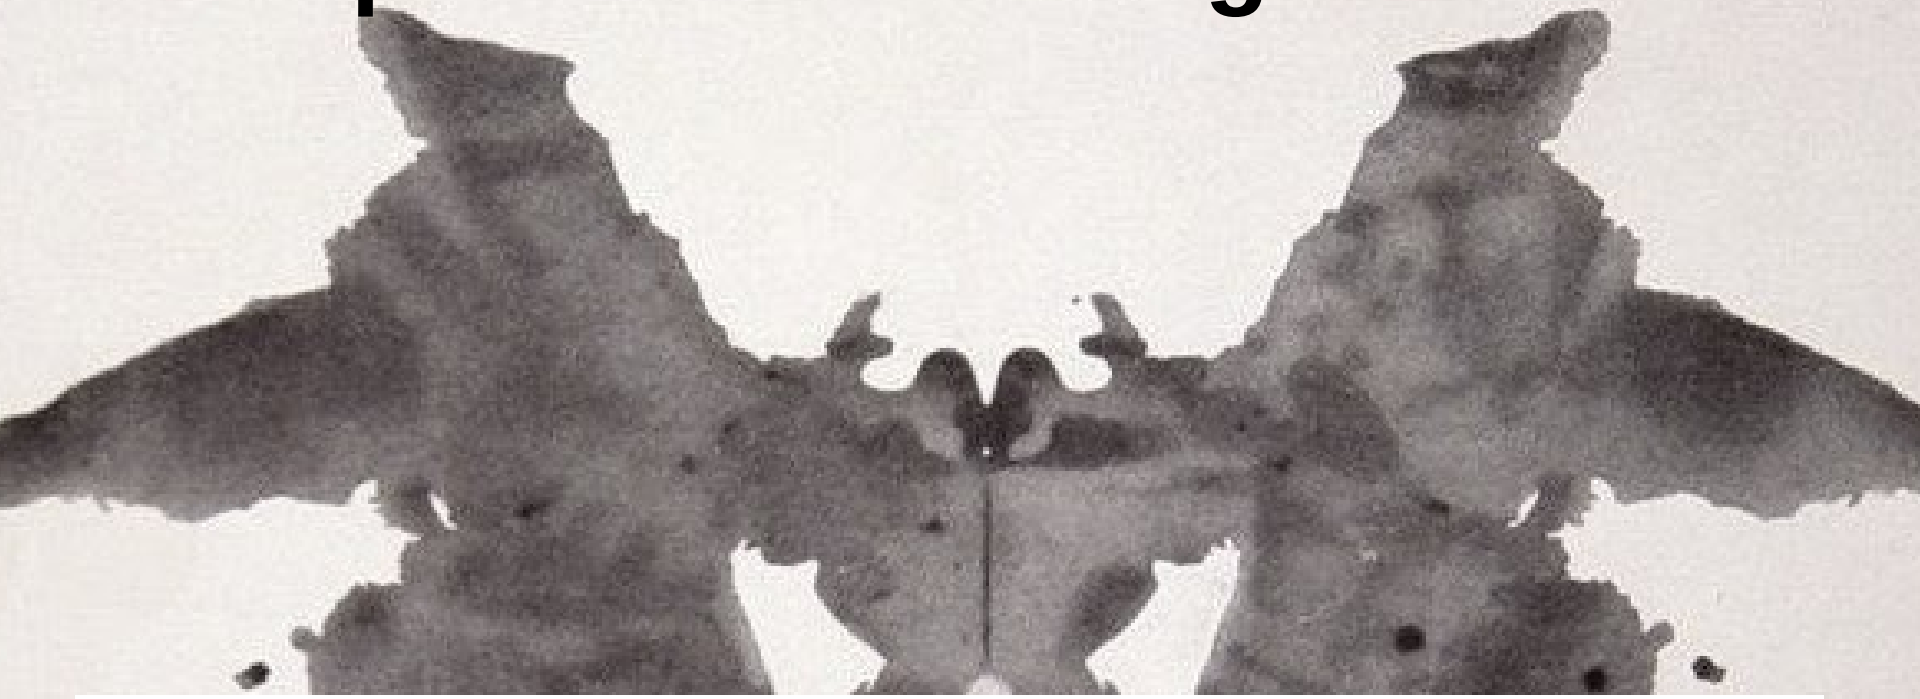

- ***Tell me what you see***
- **Data that has not been marked up, is given to the machine learning software, to see what patterns it may find, and what inferences it may be able to draw**

Reinforcement Learning

- ***Tell me what you think and I tell you if you are right***
- **Provide suitable actions to reward specific behavior in the algorithm, in certain situations, to encourage a desired path to solution or reaction.**

How to train for Unicorns?

- Most algorithms, need 100,000's of labeled data, that cover a broad spectrum of possible data the algorithm may see in operation - think colors of the rainbow
- Unfortunately we do not always have sufficient data - rare cases - what can we do?
  - Transfer Learning - We train on data that is similar first and than refine with the rare data
  - Simulation - In cases that follow clear rules, we can write programs that create synthetic data, that looks like the real data and follows the same principles
  - Create Synthetic Data - In more complex cases you can use machine learning to generate the synthetic data - via Generative Adversarial Networks (GANs)

How sure are you?

- Machine Learning algorithm can be very good in their capability, however the cases they work on can be tricky and even they may not be 100% sure what the right answer is - just as humans sometimes are not a 100% sure.
- In particular in cases where the answer has broad implications: Breast cancer yes/no rather than is this a Panda - One needs to also provide the certainty or rather the uncertainty of the result provided.
- This is still a research area and so today the results of **machine learning** algorithms should always be used with caution - it **can guide, but one should not necessarily 100% rely on them**.

What is in the Box?

- Machine Learning is often seen as a black box.
- In normal algorithms every decision, every rule has to be programmed and can therefore be verified by mathematical or physical models.
- Machine learning learns from the data on its own, potentially with little or no supervision.
- So how do we know the answer is not only correct, but arrived at in the right way?

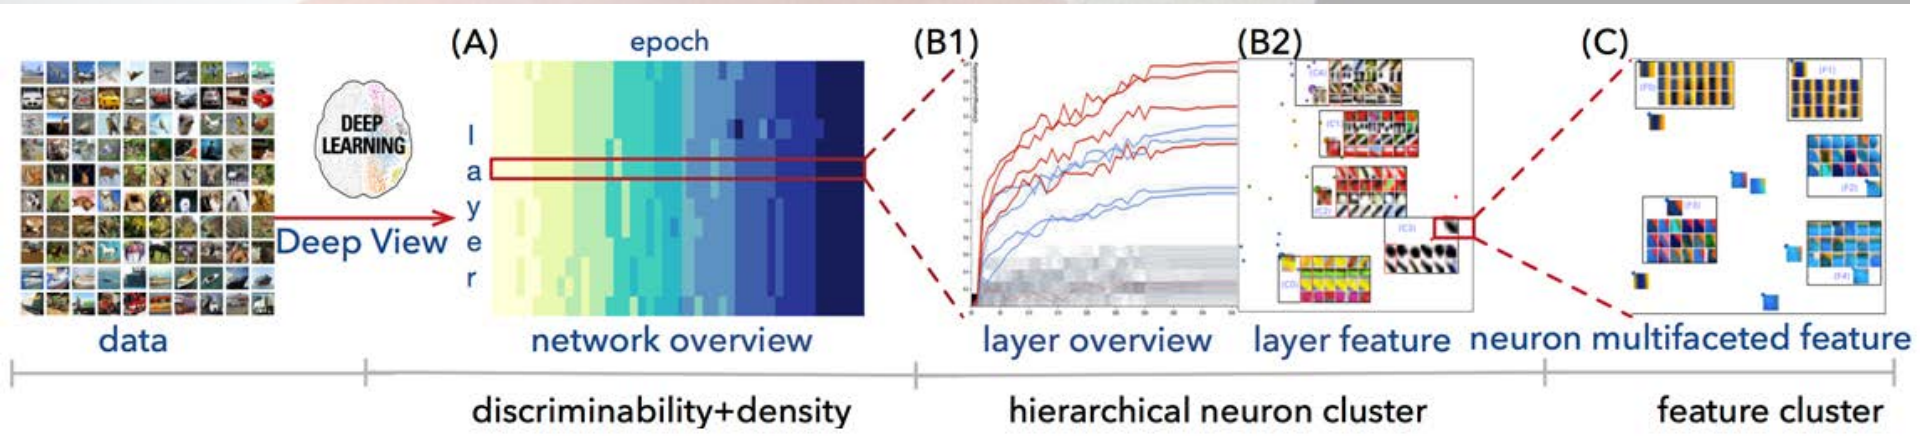

Batch versus Streaming

- Traditionally Machine Learning was applied to a large collection of Static Data
- All data is available at the start of the analysis
- There is time to consider all data before producing a result or making a prediction
- Example: Risk classification of customer

- Data is arriving while the algorithm is working
- Results are expected before the last data arrives
- In some cases you only get to look at each data point once (YOLO)
- Example: Fraud Detection in Transaction Streams

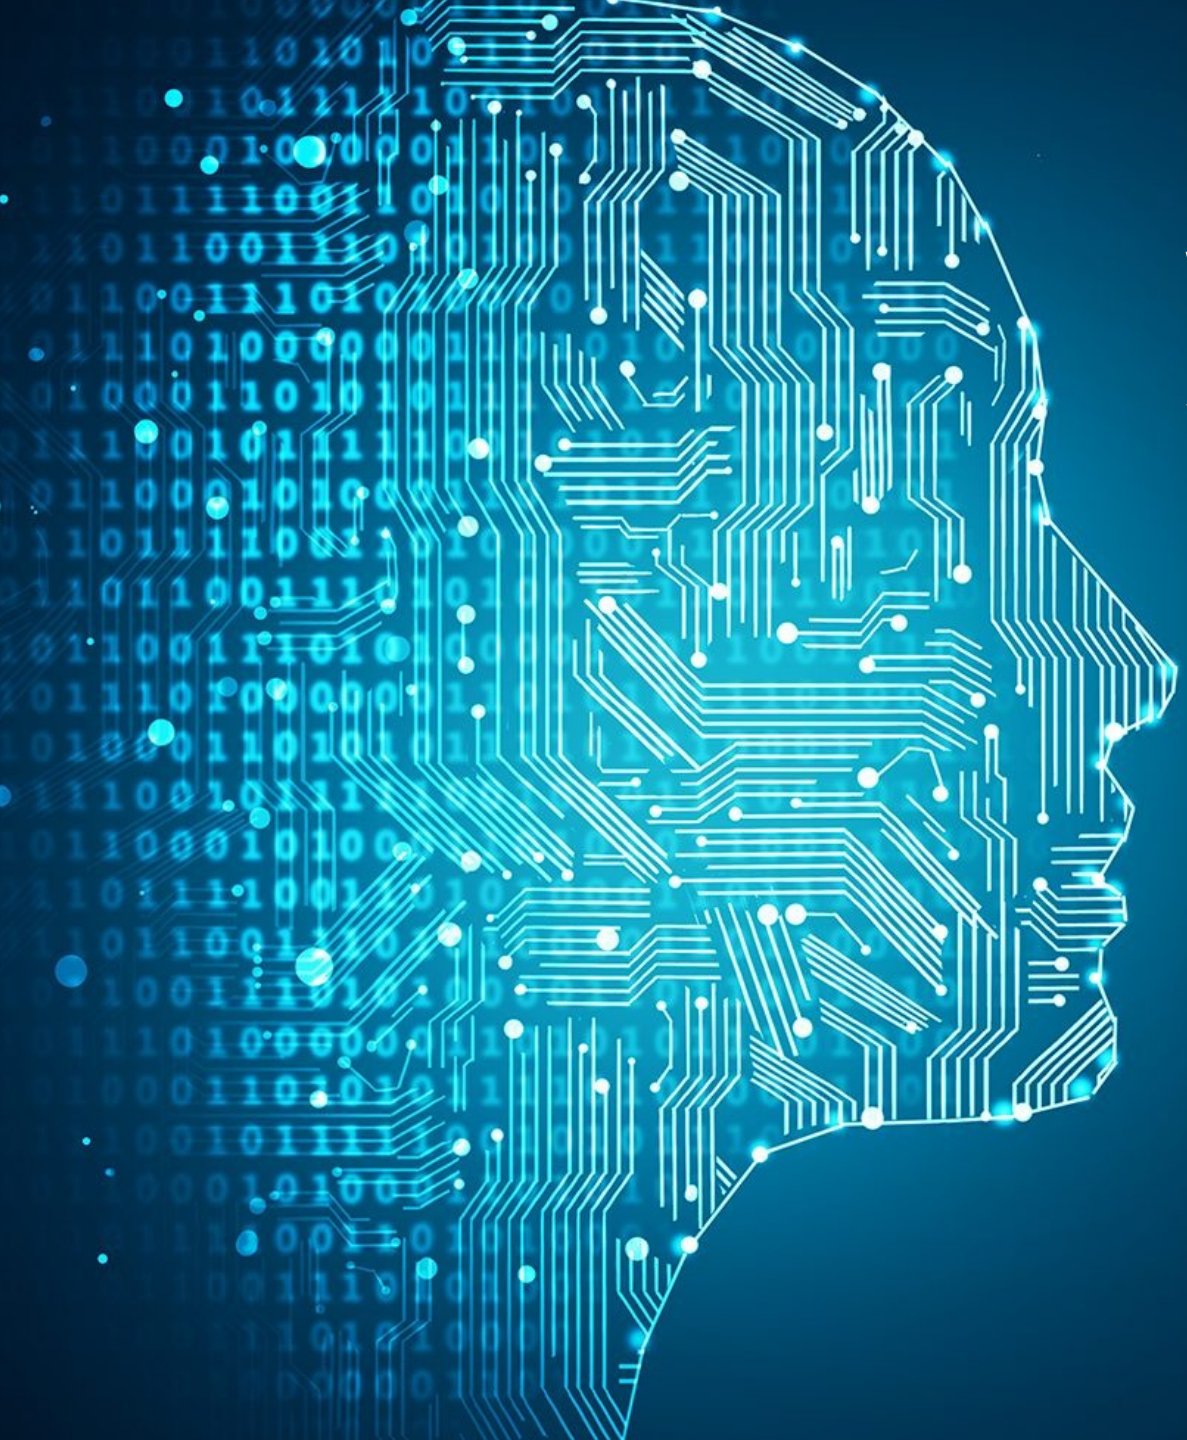

Where can you find AI today?

Imitating Human Creativity

- New AI solutions today are starting to mimic human creativity - so far they are largely imitating specific artists, rather than creating on their own.
- Examples are:
  - Painting - A new Rembrandt, through analysis of several hundred original works of the master, studying techniques, features and composition
  - Music - Several solutions exist to compose music that sounds like e.g. the Beatles
  - Acting - AI solutions have written comedy sketches

Amazon Alexa

- Build on voice recognition machine learning algorithms to understand your commands
- AI is used to generate suitable responses - can become rather independent - Have you ever argued with your Alexa?
- AI is used to take actions - switch off lights, order and item online
- **Less well known, thousands of humans listen to your voice recordings to improve Alexa's voice recognition - labeling data, supervised learning**

Self-Driving Car

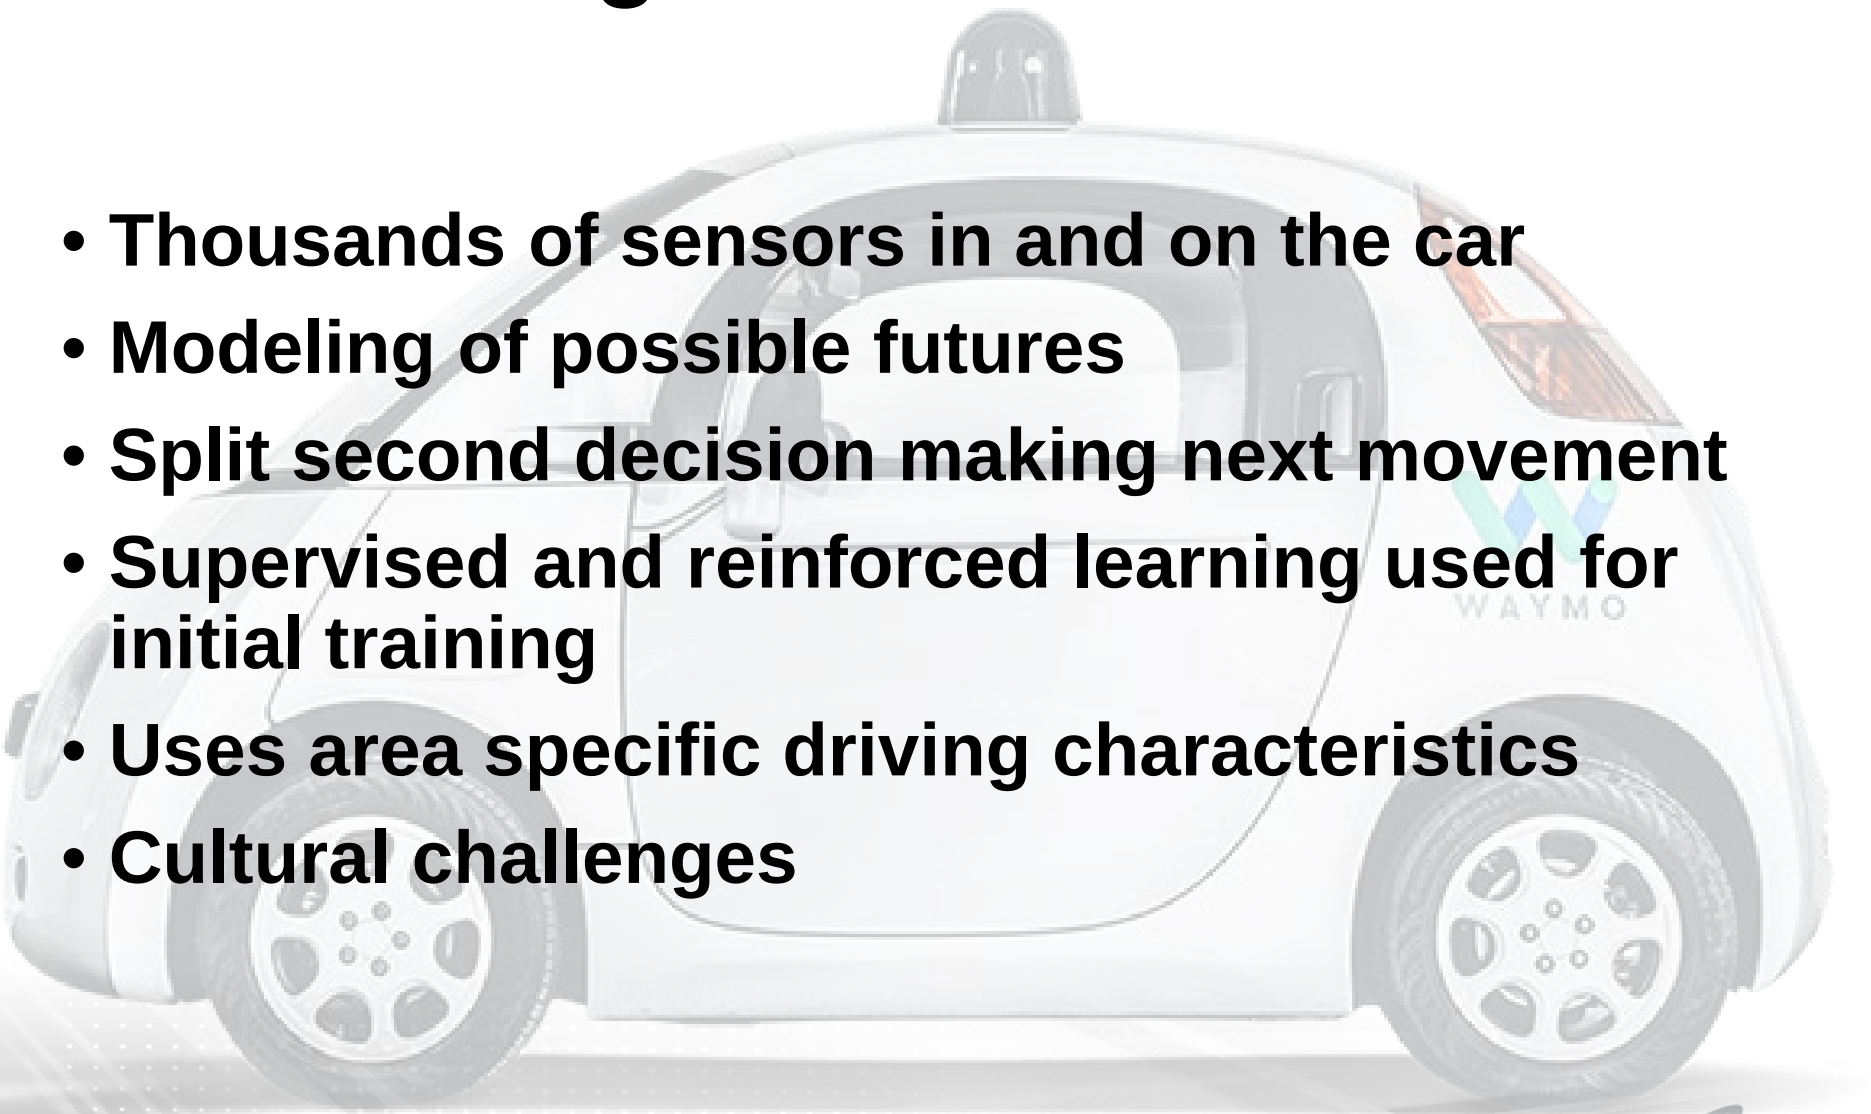

- Thousands of sensors in and on the car
- Modeling of possible futures
- Split second decision making next movement
- Supervised and reinforced learning used for initial training
- Uses area specific driving characteristics
- Cultural challenges

Artificial Intelligence Application Characteristics

- Data arrives at high velocity and volume
- Data is often extremely varied
- Tacit knowledge is needed for accurate interpretation of data
- Critical decisions need to be taken while data is still arriving
- Critical events are usually rare

What does it take to build an AI application?

- Identify emerging phenomena in high velocity streaming data **- Streaming Statistics, Data Mining, Machine Learning**
- Determine what is of interest and impact, generate candidate explanations **– Streaming Deductive Reasoning**
- Human-Computer collaboration to jointly adjust data collection, reasoning and insights **– Science of Interaction, Cognitive Depletion Detection, Hypothesis Exchange, Adaptive Algorithms and Workflows**
- Evaluate the impact of possible decisions **- On Demand Prediction**
- Document which decisions were taken during the analysis process to explain the results **- Provenance**

Challenge: Human Computer Collaboration

- Cognitive Limitations
- Different Levels of Ambiguity Tolerance
- Potentially Different Mental Models
- Unexpected Behavior
- Cultural Differences
- Interaction Speed
- Trust
- Ethics

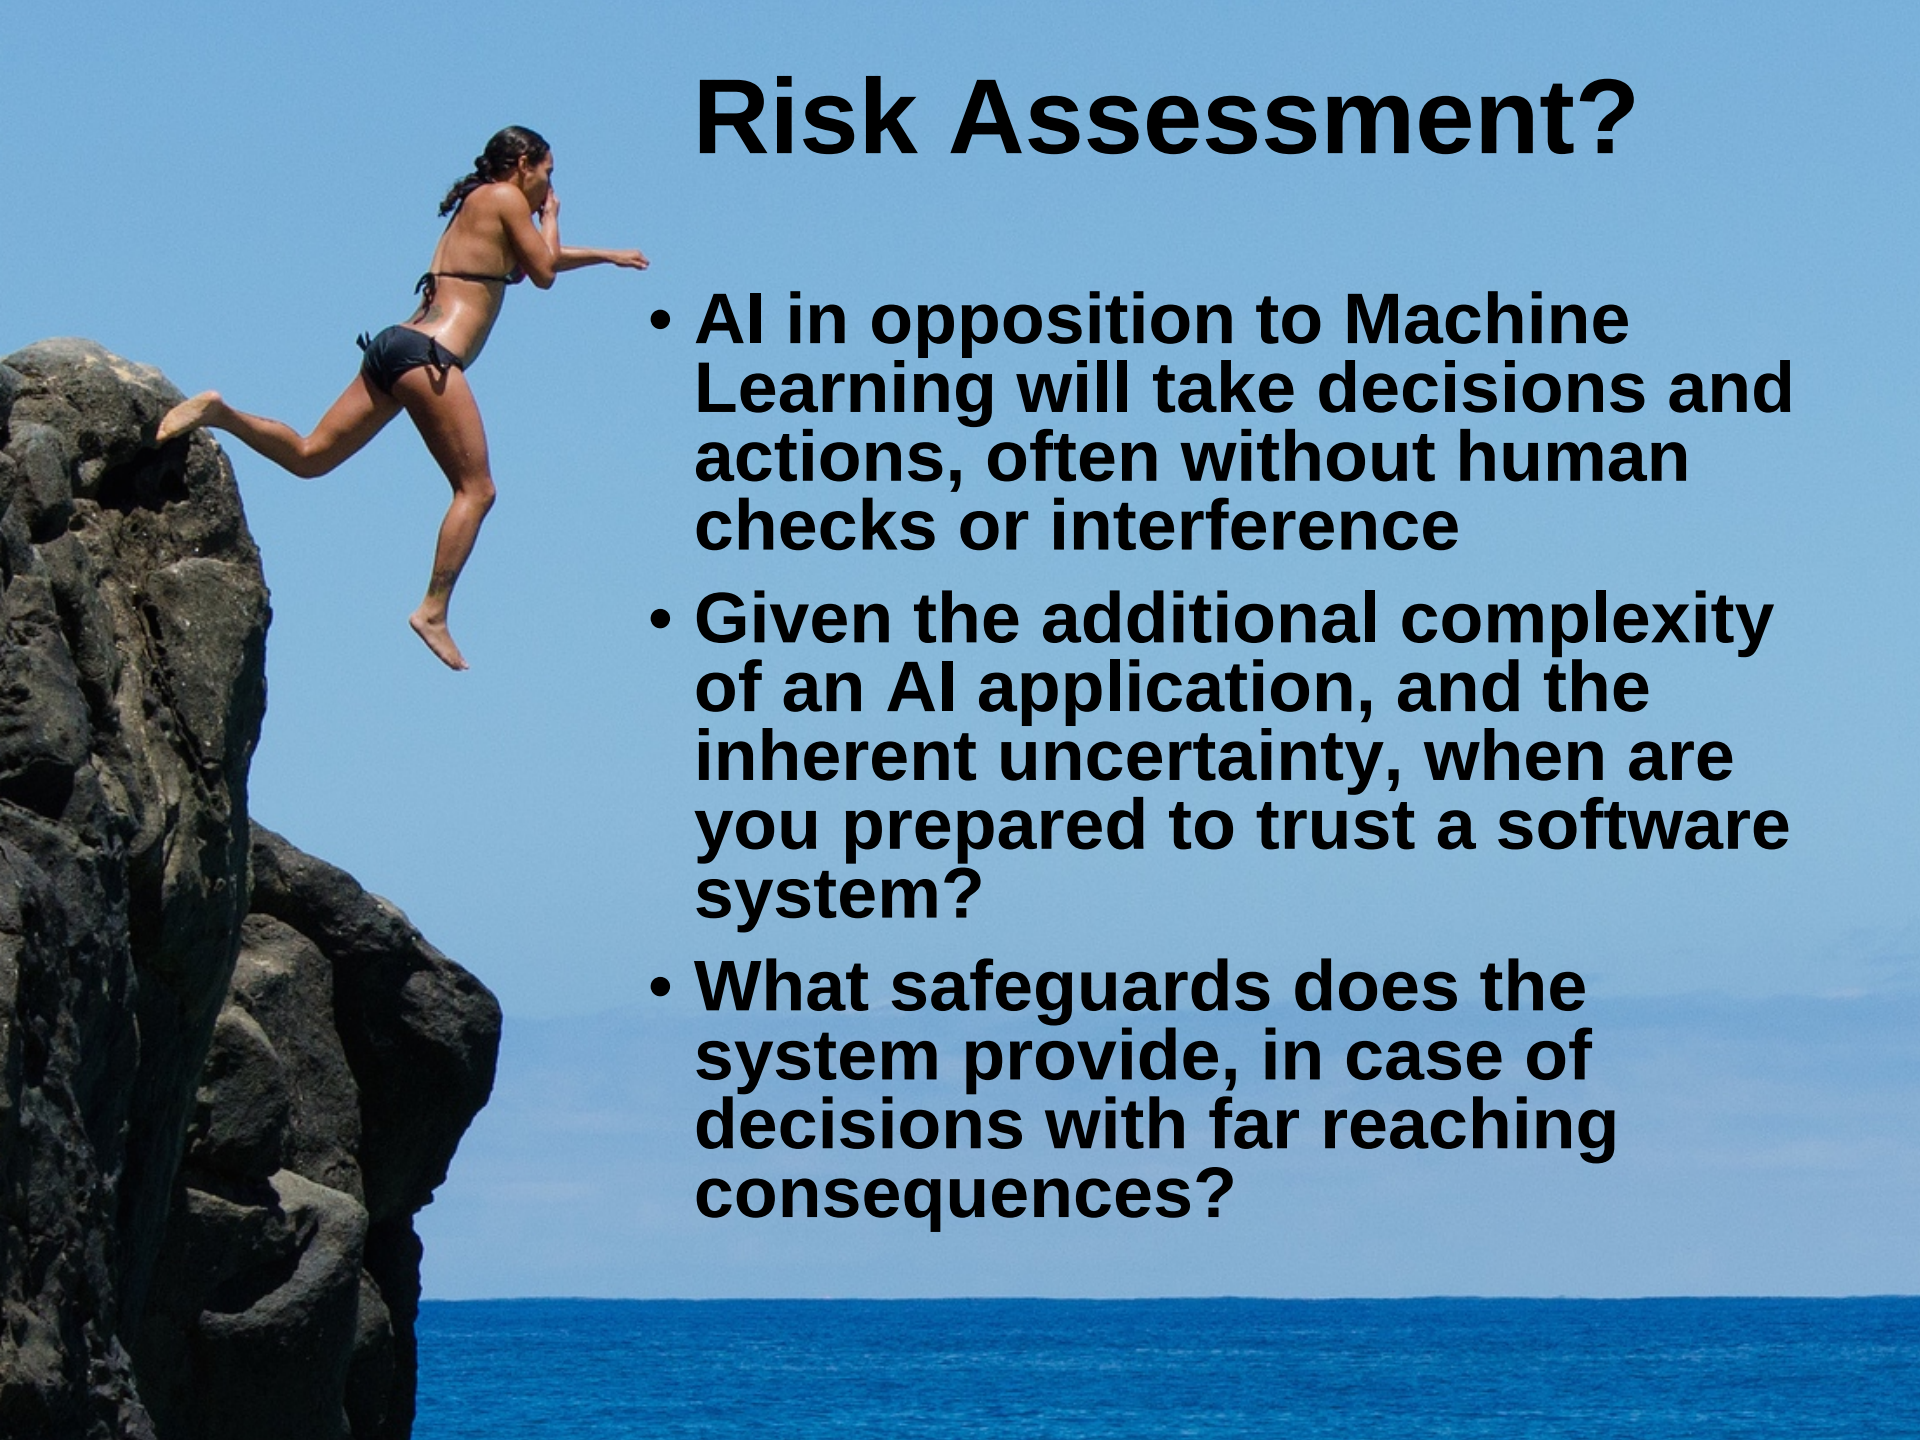

Risk Assessment?

- AI in opposition to Machine Learning will take decisions and actions, often without human checks or interference
- Given the additional complexity of an AI application, and the inherent uncertainty, when are you prepared to trust a software system?
- What safeguards does the system provide, in case of decisions with far reaching consequences?

We are only at the beginning of exploring what AI can do for us. There are many possible applications in banking, healthcare, manufacturing, transport or science, but we must proceed with care, being aware of the limitations and pitfalls.


BNL
Brookhaven
National Laboratory
OPERATED BY BROOKHAVEN SCIENCE ASSOCIATES UNDER CONTRACT WITH THE UNITED STATES DEPARTMENT OF ENERGY

Thank you!