BRIDGING THE GAP BETWEEN AI IMPLEMENTATION AND PRODUCTION

DataBeacon

MONDAIS

AI for safety-critical systems

(SUPERVISED) MACHINE LEARNING BASICS

Features

Label

Rows

Size	Beds	Baths	Zip	Price
1100	1	1	64576	1.29
1900	3	1.5	78321	2.14
2800	3	3	98712	3.10
3400	4	3.5	25721	3.75

Columns

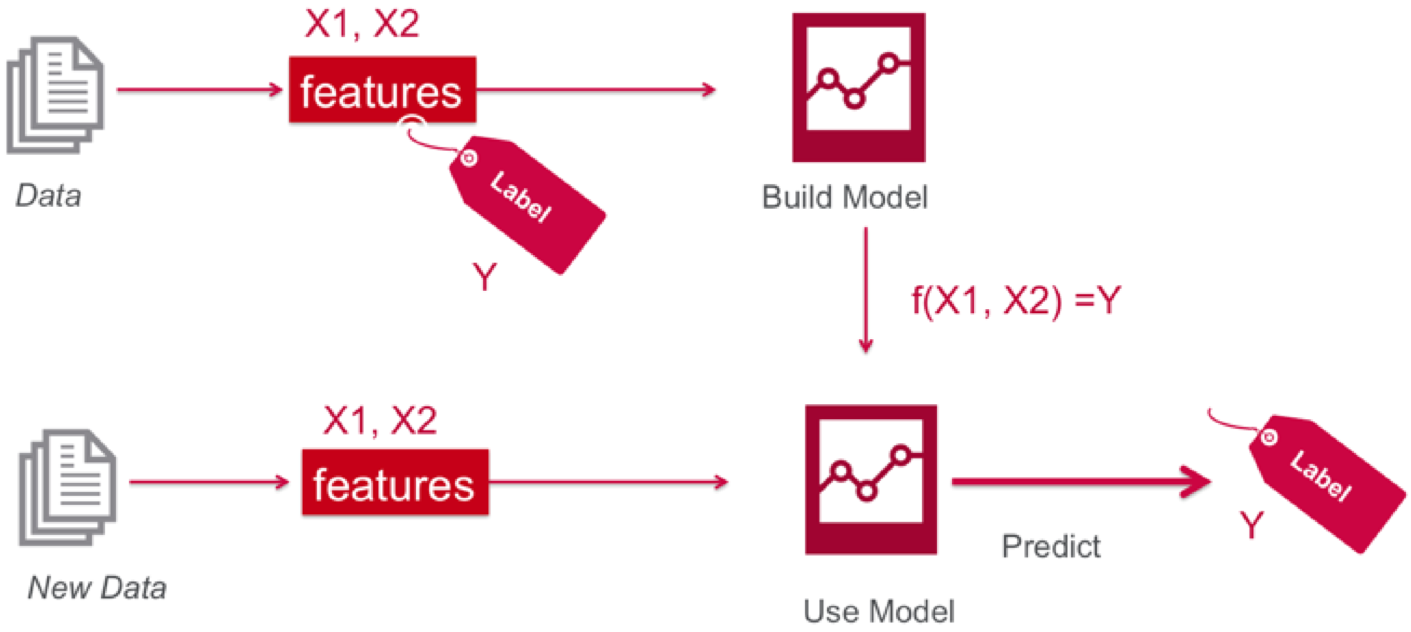

MACHINE LEARNING VS TRADITIONAL PROGRAMMING

TRADITIONAL PROGRAMMING

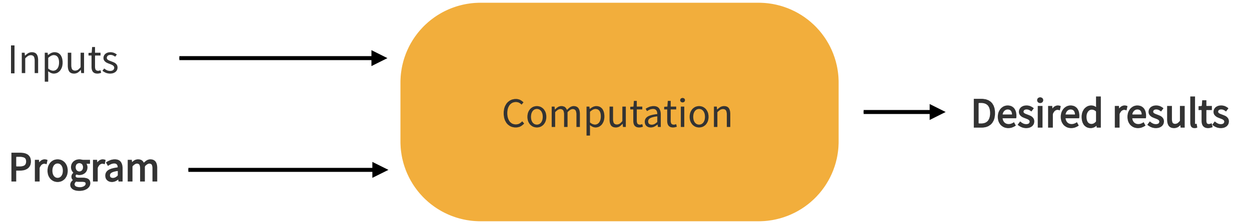

MACHINE LEARNING

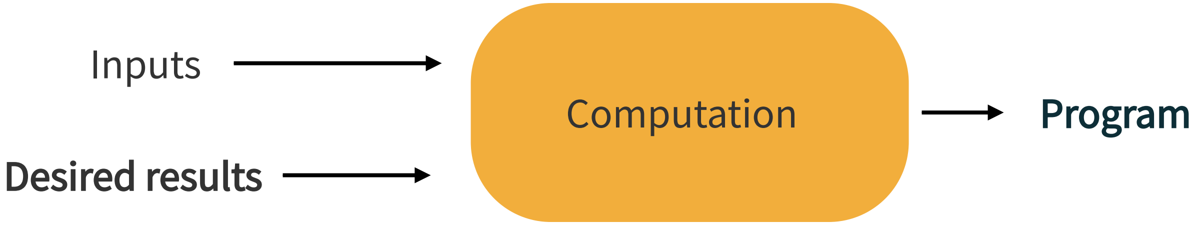

DataBeacon

CODED RULES VS ESTIMATED WEIGHTS

TRADITIONAL PROGRAMMING

- Manually formulate/code rules.
- Hundred of thousands of lines of code.
- Written by **humans**.
- To match the **desired and known output.**

MACHINE LEARNING

- Extract the rules from **historical data**.
- Hundred of weights and mathematical relations.
- Written **automatically** by an **optimization process**.
- To optimise an **objective function.**

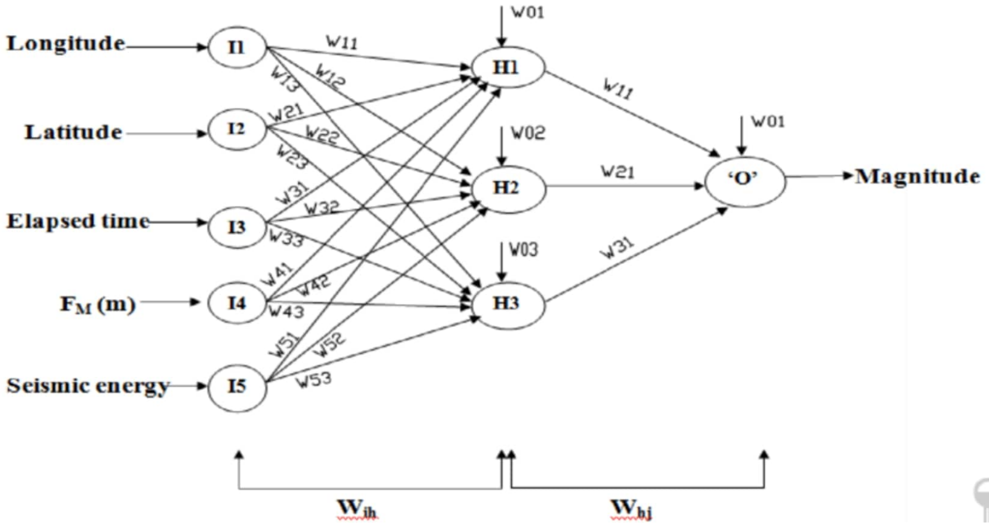

DataBeacon

AN EXAMPLE: ACTIVITY RECOGNITION

TRADITIONAL PROGRAMMING

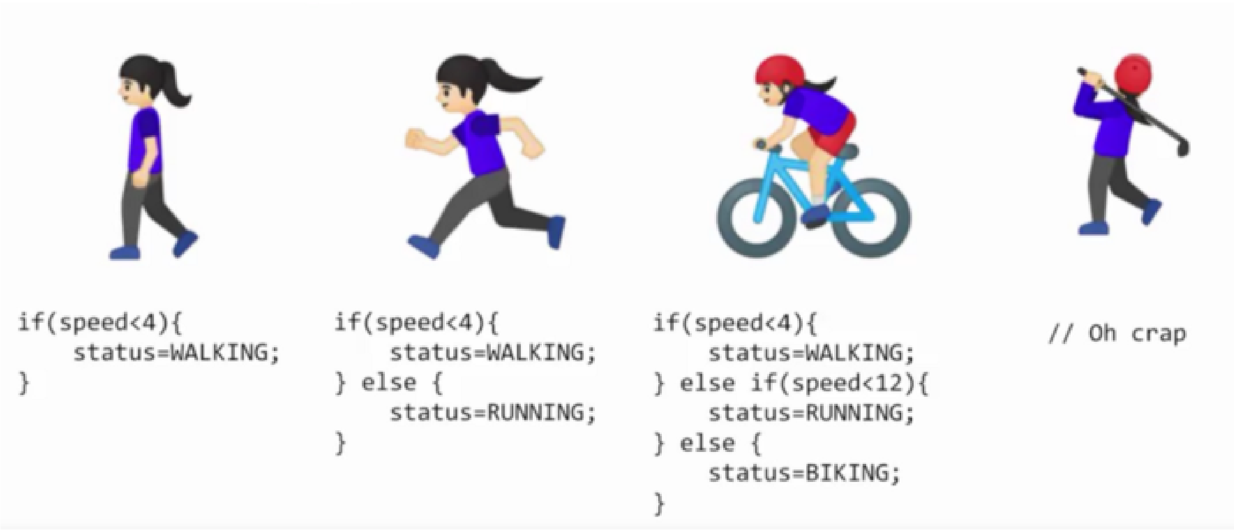

Rules specified on a highly abstracted but known representation.

MACHINE LEARNING

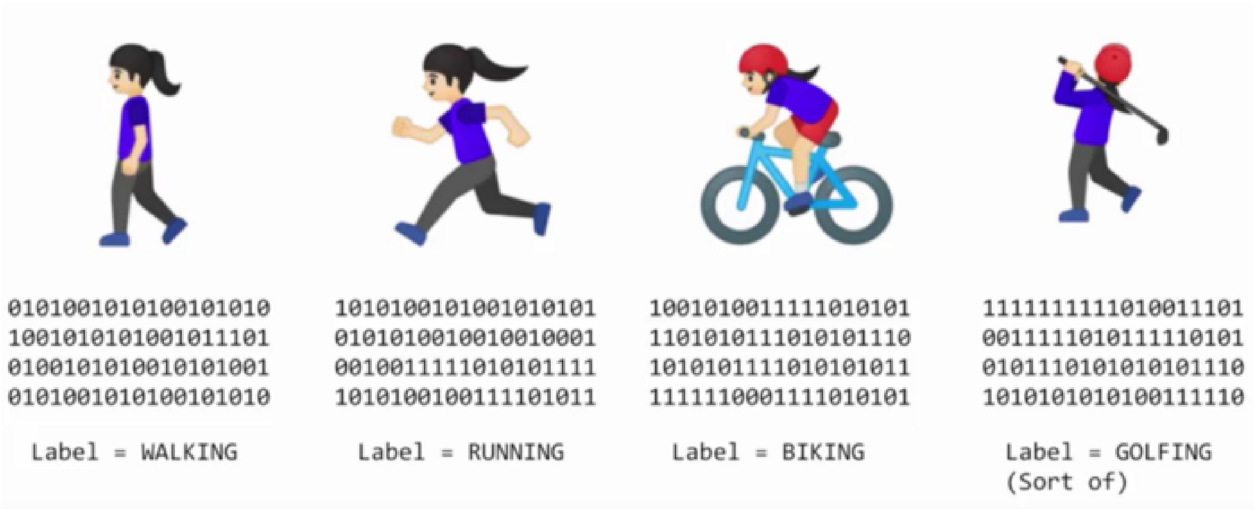

Just because a Neural network says so, based on a lot of examples.

Image source: https://rstefanus16.medium.com/

THE CURRENT LANDSCAPE...

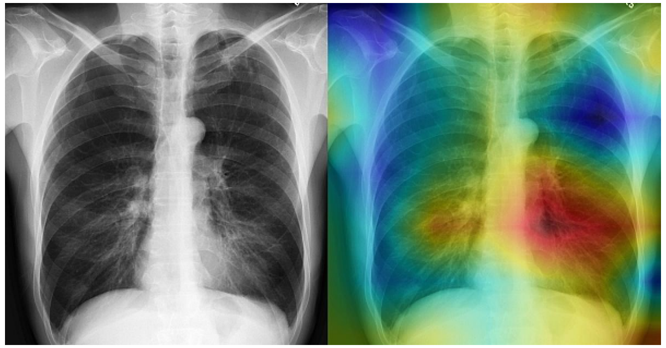

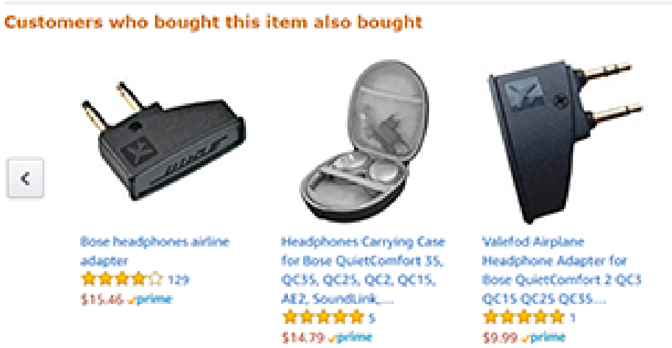

DataBeacon

SAFETY CRITICAL SYSTEMS IN AVIATION

- **Failure** would cause a significant increase in the safety risk for the **people** or **environment involved**.
- How safe human-controlled transportation systems are?
  - cars: 7 deaths every 10^8 miles
  - buses and trains: 0.1-0.4 deaths every 10^8 miles
  - **aircrafts: 0.07 deaths every 10^8 miles**

  Source: I.Savage, "Comparing fatality risks in United States transportation across modes and over time", *Research in Transportation Economics,* 43:9-23. 2013)

- How this is achieved in the aviation industry?
  - Strong **certification process** and high requirements.
  - **Model-Based** Design
  - **Fleet-wide** analysis
  - Very **structured** operating environments
  - **Well trained** personnel (pilots, FAs)

DataBeacon

THE "EASY" CHALLENGES OF ML IN AVIATION

- **Failure rates**
  - Numbers that are acceptable in other industries (e.g. banking or energy) are too high for Aviation standards.
- **Poor metrics** for certification
  - "classification error" is not enough for a lawyer...
- Data requirements...
  - often **unknown**...
  - but always large (millions of operations).
- Size of errors in relation to the **amount of training data.**
- How do you catch **outliers** and **corner cases**?
  - In Aviation they are very very very rare...
- Individual outputs are not relevant
  - ... but **system behaviour** is.
- Lack of data providers and data sharing.

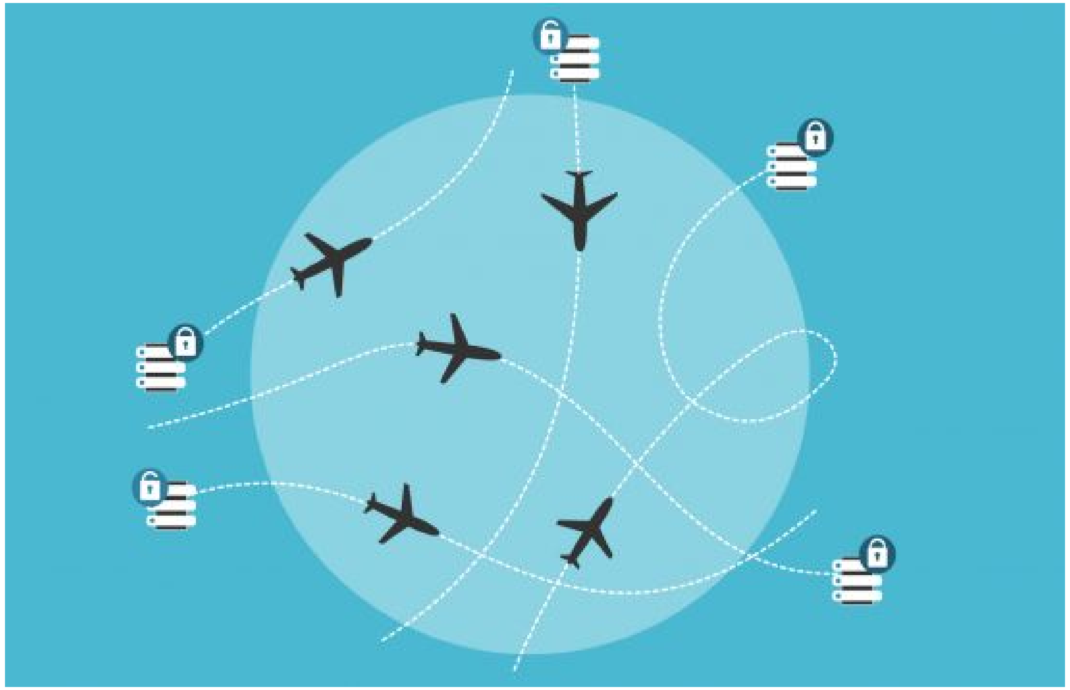

DataBeacon

THE HARD CHALLENGES OF ML IN AVIATION

- We are way past the peak of the hype cycle...

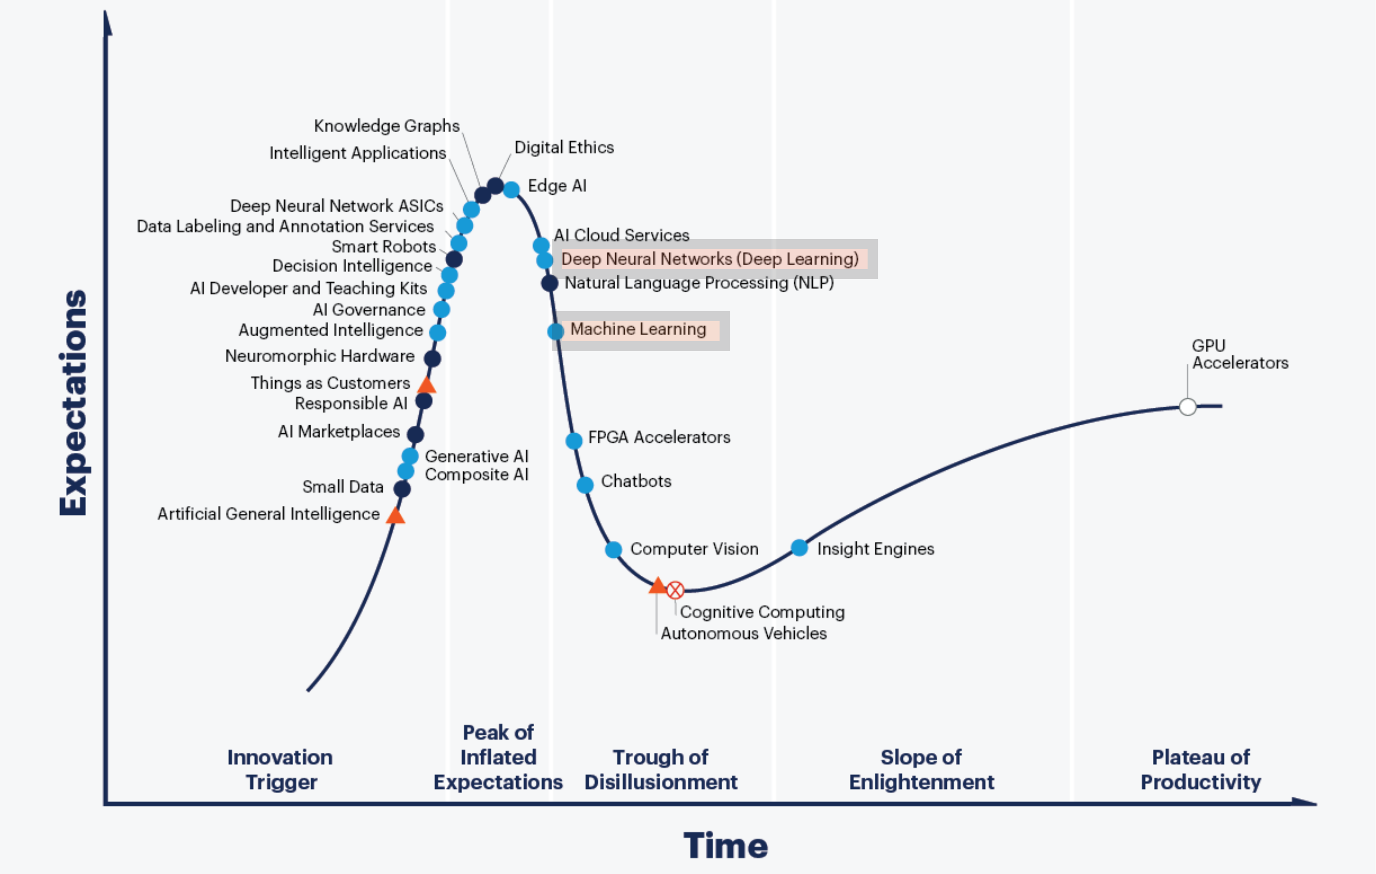

- The **development timeline** is much longer than what many expect...
- **Fundamental issues** for adoption remain **unsolved**.
- A tyrannical effort by industry, academia and **regulatory bodies** is required.

DataBeacon

THE PAST, PRESENT AND FUTURE OF AI: DATASETS

Amount of lost sleep over...

Andrej Karpathy
Formerly PhD Student at Stanford. Now at Tesla

PhD

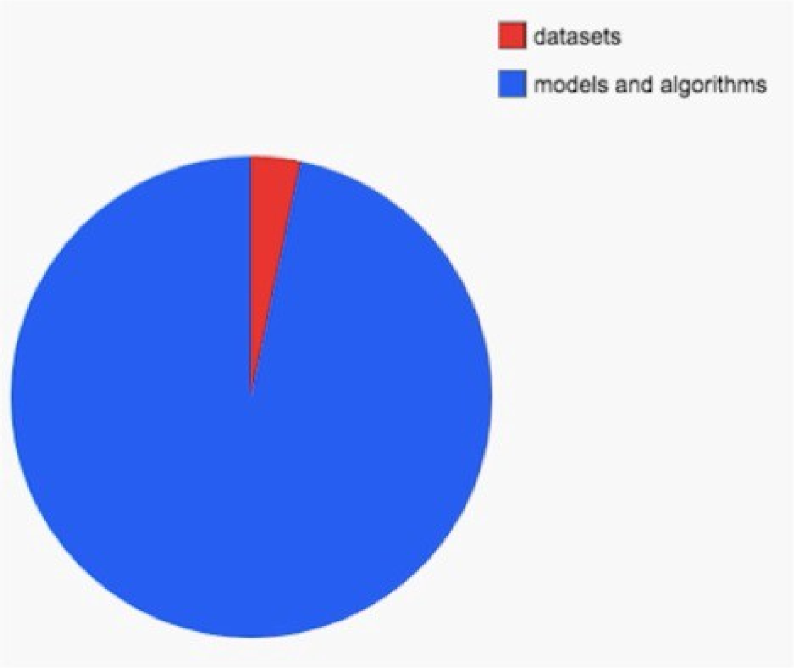

Tesla

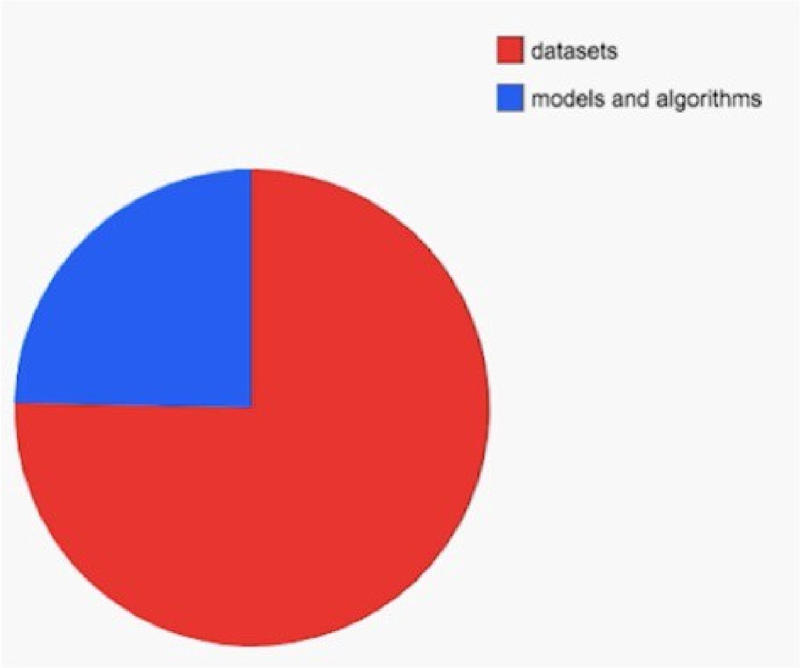

DataBeacon

"BECOMING ONE WITH DATA"

- Data is the **single point** of success or failure of an AI product.
- A **good understanding of domain knowledge** is required.
- Understand the training set:
  - What makes certain instances truly interesting from the rest?
  - What is the importance of certain feature?
  - Do your features cover every aspect of the problem?
- Know how your data behaves over time:
  - Will your data change? It is stationary?
  - Have your data changed? (hello, Covid-19)
- Assess if your data is complete:
  - Do you have enough examples?
  - Do your historical data cover every possibility?
  - Do you need synthetic data?

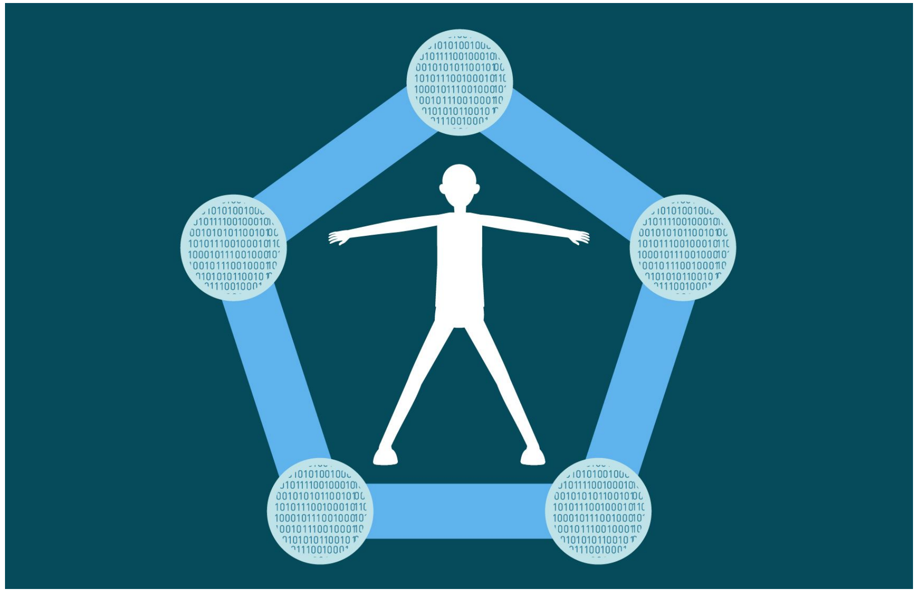

DataBeacon

THE DATA LABELLING CHALLENGE

HIGHLY INTERACTIVE AND NON-TRIVIAL

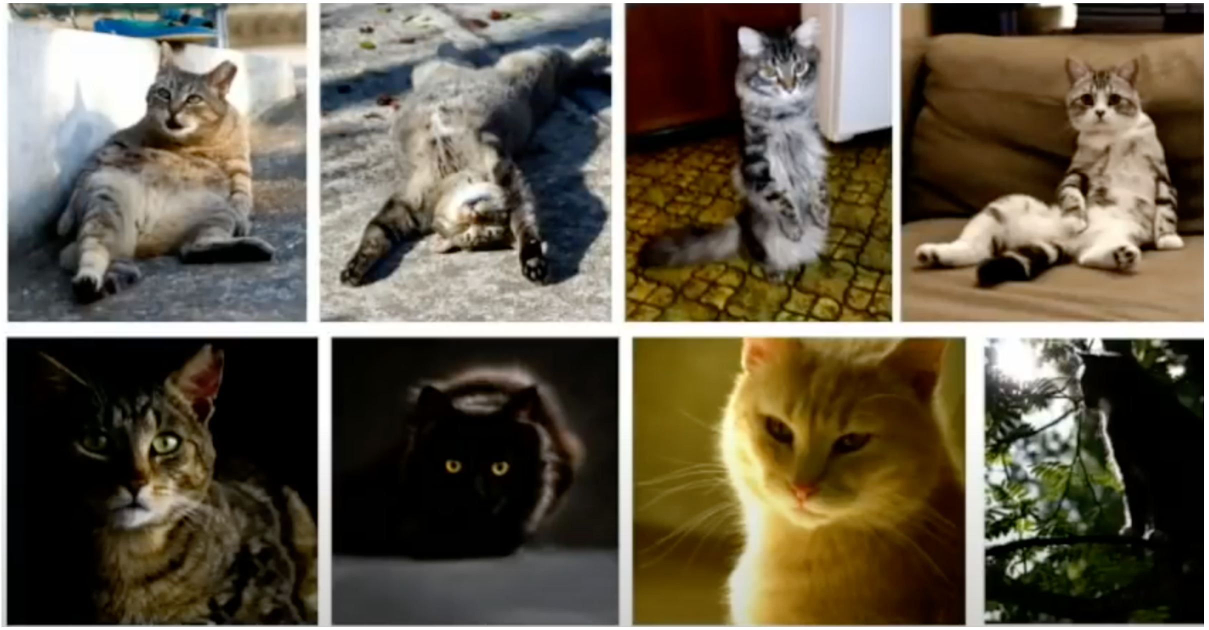

CARE ABOUT LABEL IMBALANCES…

THEY MATTER IN **SAFETY CRITICAL SYSTEMS**.

DataBeacon

ACHIEVING REALISTIC DATASETS

- We need to train with the most **realistic dataset** possible.
- This sometimes means to **tweak** the data.
  - Make it hard to learn...
- Test with "**worse case scenarios**".
- Look for:
  - High label/examples imbalance: Discard 95% of the label of a class.
  - Noisy labels: Perturb a % of the labels.
  - Multi-Task solutions: Add 10 any other labels.
  - Incompleteness; add 1M of unlabeled cases.

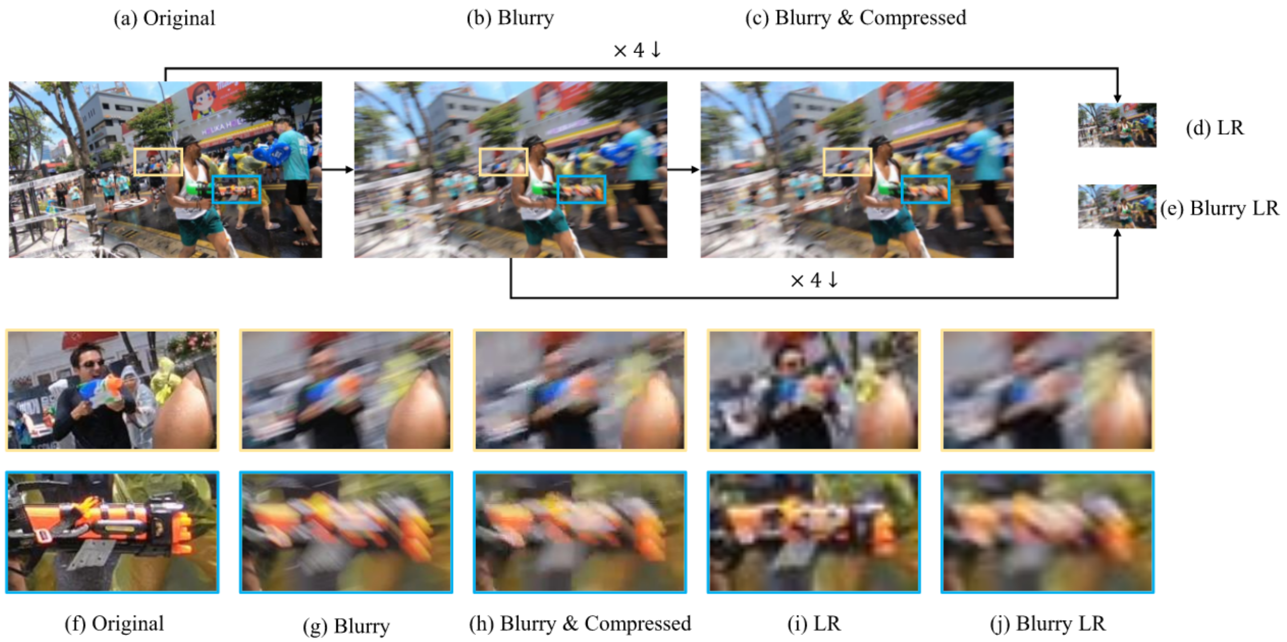

DataBeacon

IF MACHINE LEARNING "CODE" IS AN AUTOMATED OPTIMISATION THEN...

WHAT ARE ALL THE HUMANS DOING?

- Label unseen examples.
- Visualize datasets and labels.
- Create/edit labels based on feedback from experts.
- Study mislabeled examples.
- Suggest new examples to label.
- Flag labeller disagreements.
- Engineer and Validate features
- Model supervision

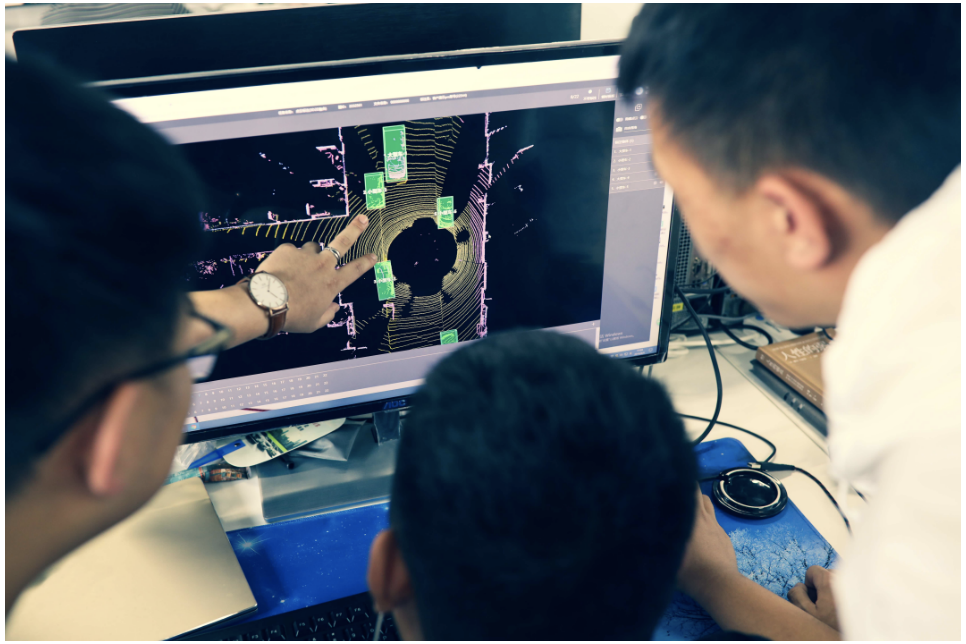

Source: Testin’s Beijing data labeling office

HUMAN-IN-THE-LOOP (HITL)

DataBeacon

THE POWER OF ACTIVE LEARNING

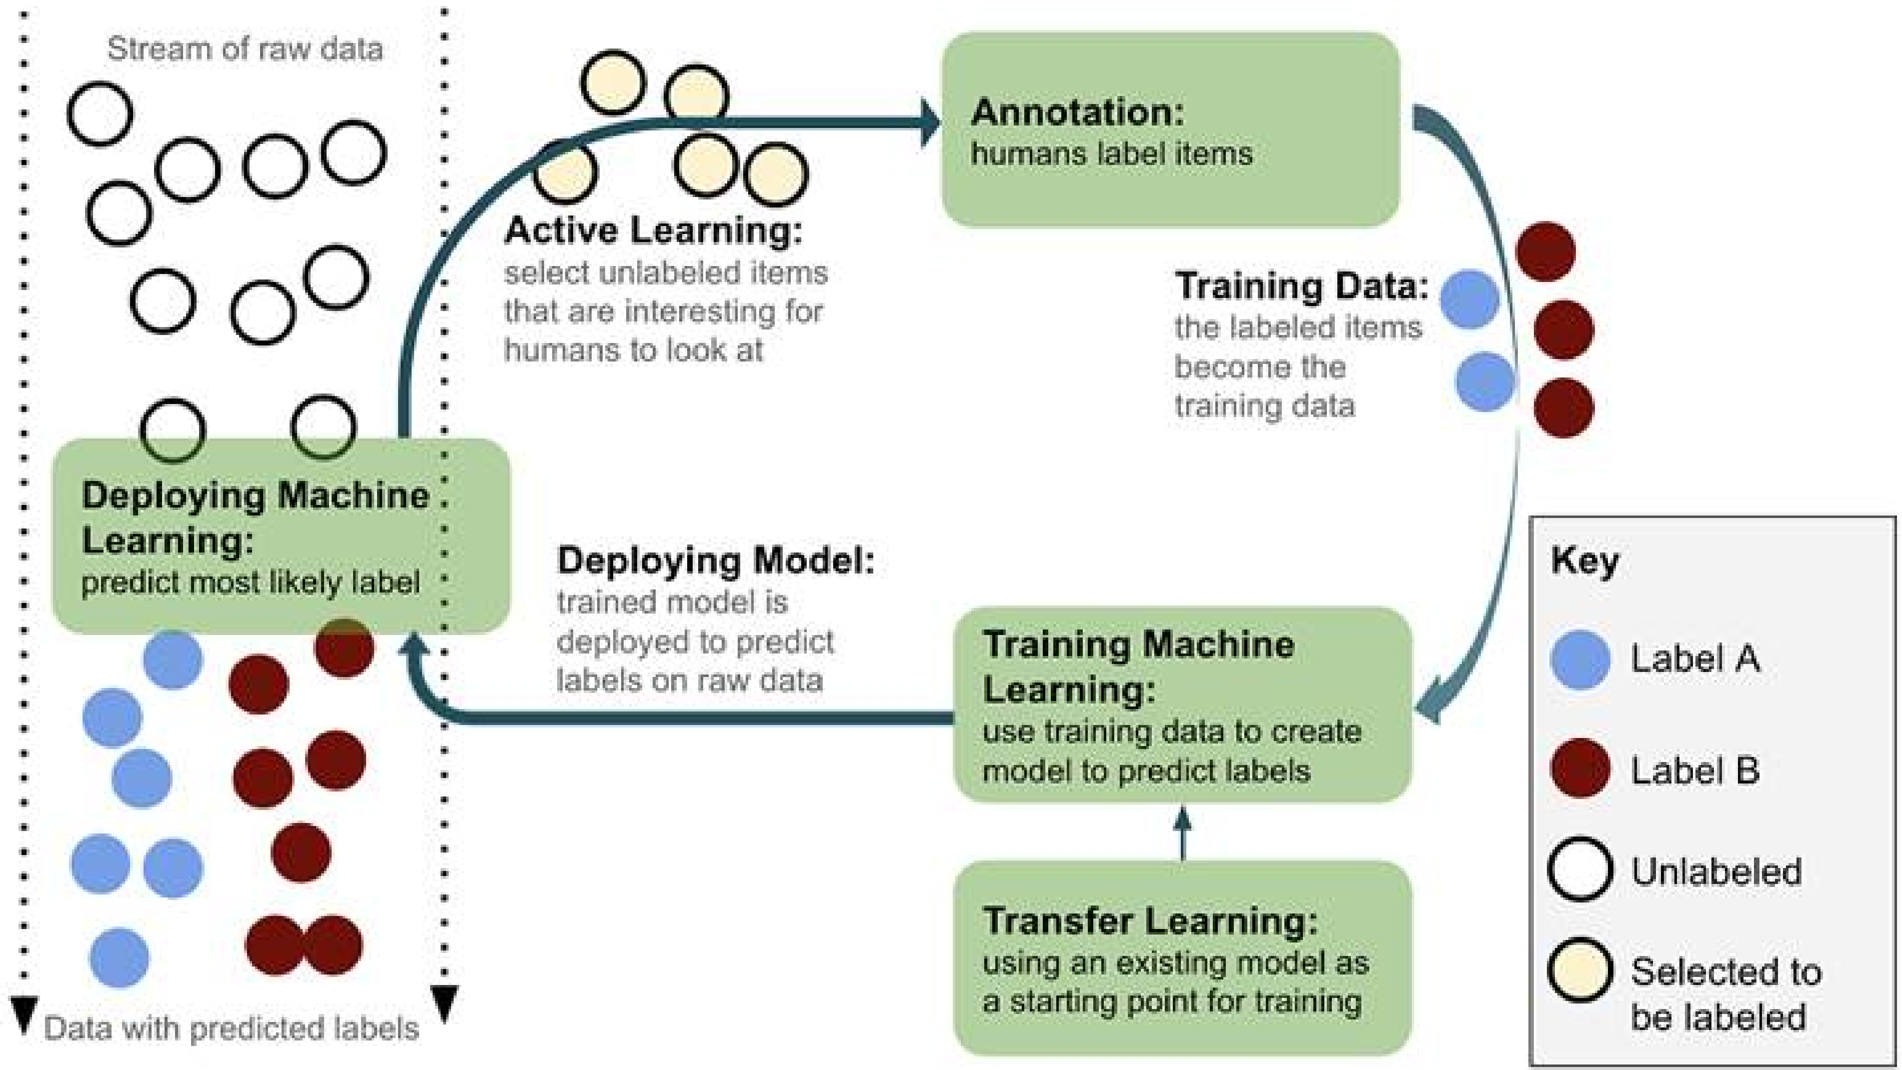

Source: Munro Monarch, Robert (2021). Human-in-the-Loop Machine Learning: Active learning and annotation for human-centered AI (1st ed.). Manning.

DataBeacon

HUMAN ANNOTATIONS: THE UGLY PART

- Humans **will** provide erroneous labels.
- Labels depend totally on the **previous experience**, **field knowledge**, personality, etc... of your workforce.
- **Disagreements** are expected, after all, these are dubious cases.
- This scales very badly with highly technical problems; e.g. aviation problems.
- Overcoming this challenge requires surprisingly **sophisticated statistics**.

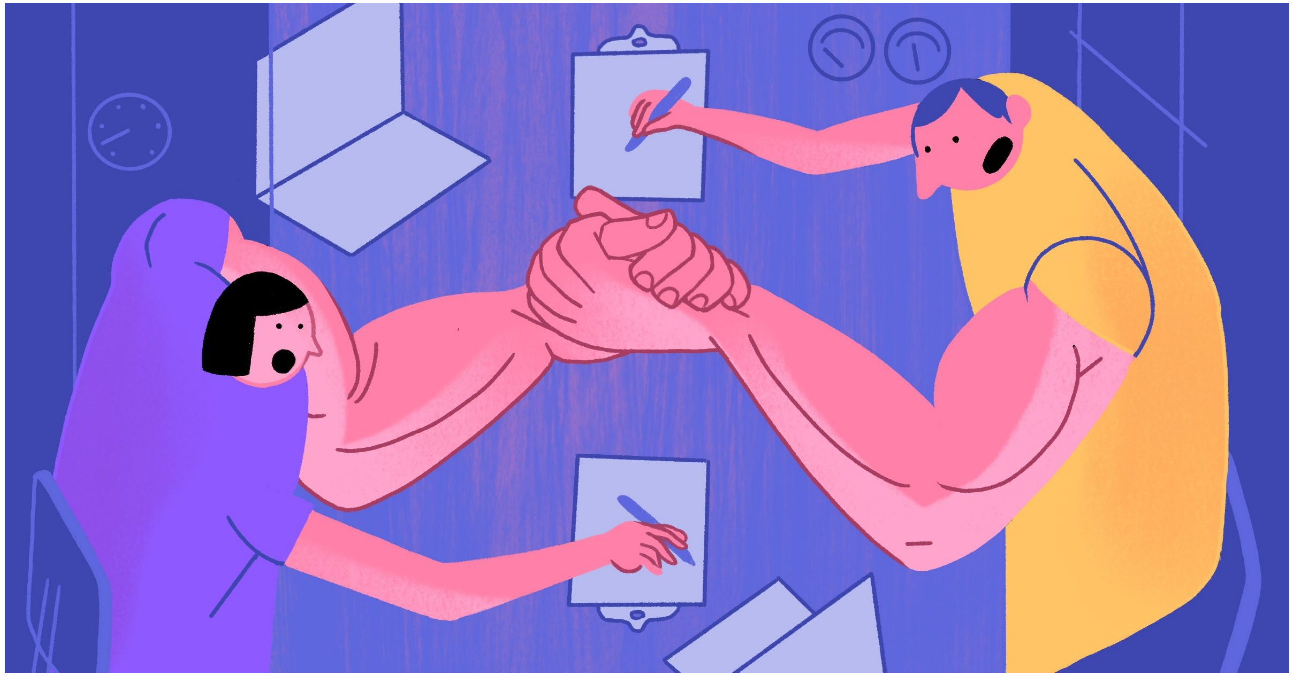

DataBeacon

THE "BLACK BOX" PARADIGM

Geoffrey Hinton @geoffreyhinton · 20 feb. 2020

Suppose you have cancer and you have to choose between a black box AI surgeon that cannot explain how it works but has a 90% cure rate and a human surgeon with an 80% cure rate. Do you want the AI surgeon to be illegal?

810 1,7 mil 5,1 mil

DataBeacon

ACCURACY VS INTERPRETABILITY

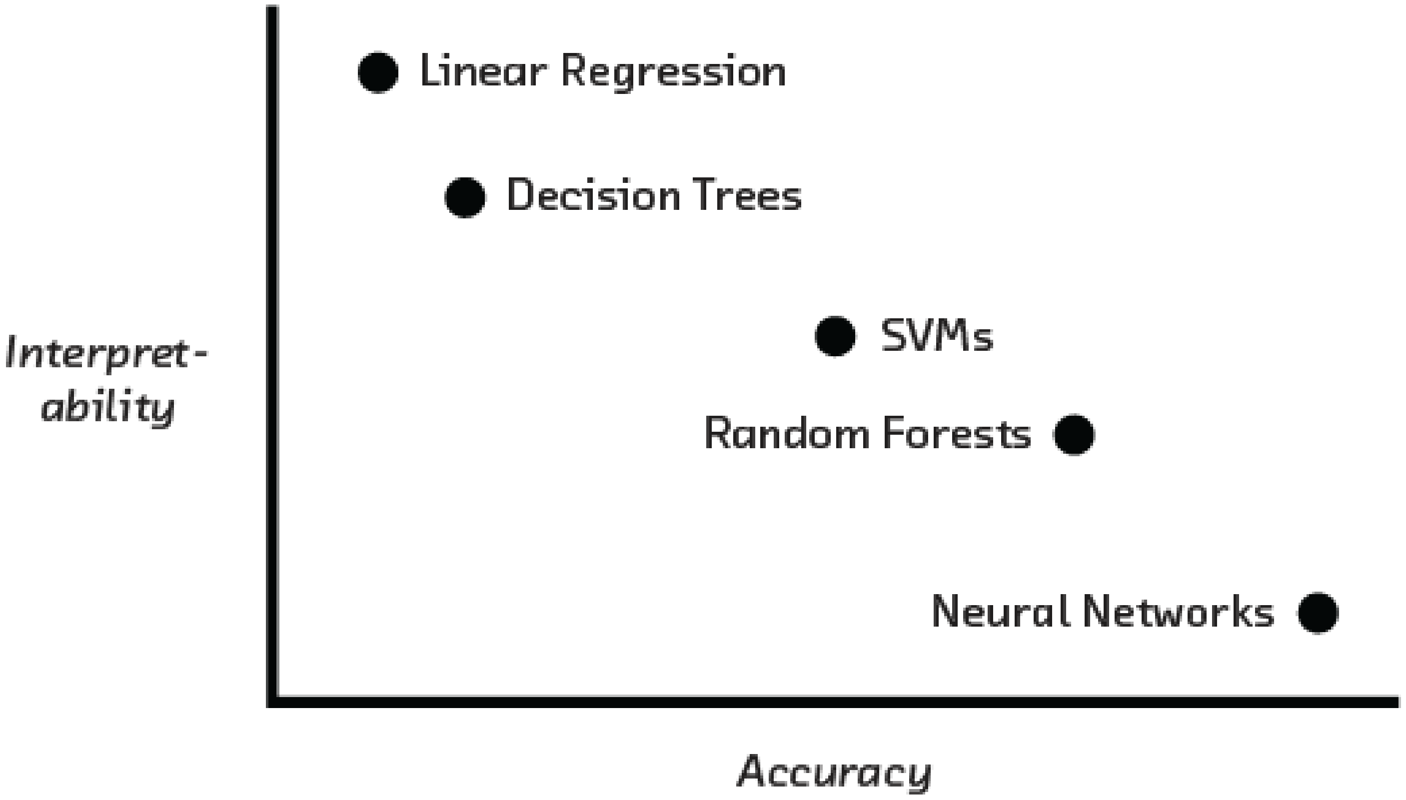

Source: https://towardsdatascience.com/guide-to-interpretable-machine-learning-d40e8a64b6cf

DEEP LEARNING AND "THE FUTURE"

- February 2020: Microsoft's Turing NLG (17 billion parameters)

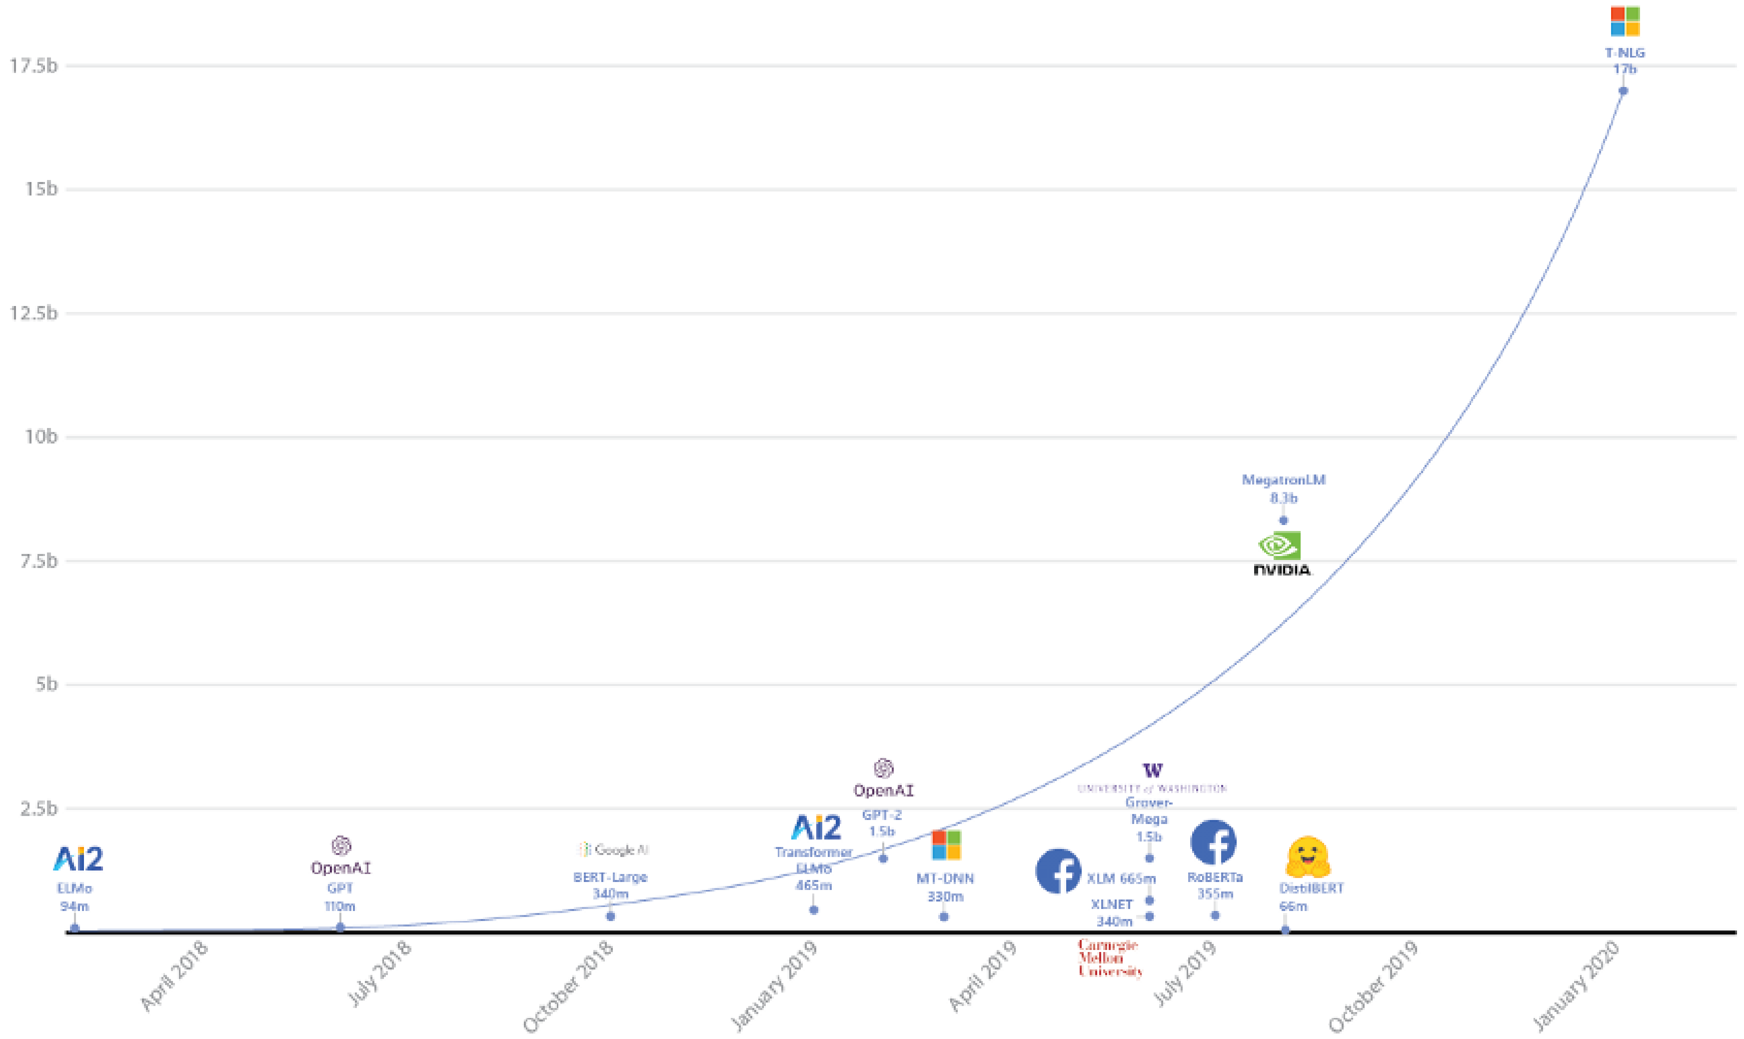

Source: https://www.microsoft.com/en-us/research/blog/turing-nlg-a-17-billion-parameter-language-model-by-microsoft/

DataBeacon

OPENING THE BLACK BOX

- **Visualizations** are the key.
  - Training set
  - Labels
  - Predicted values
  - Features

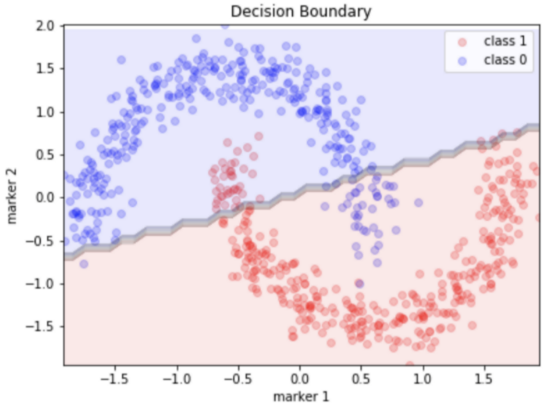

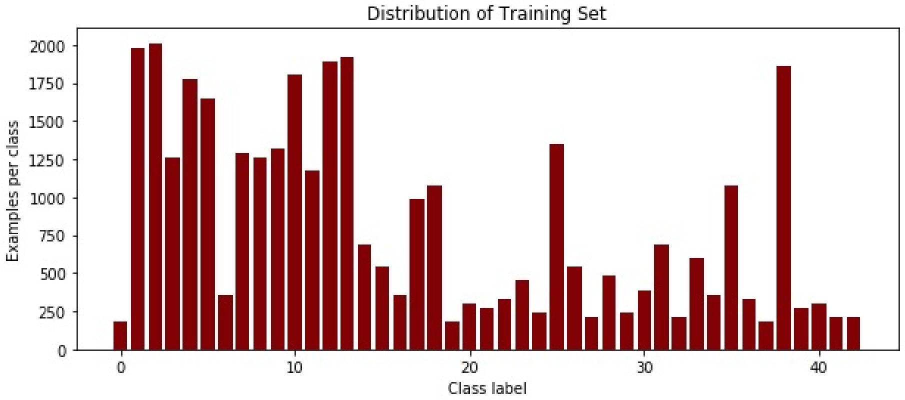

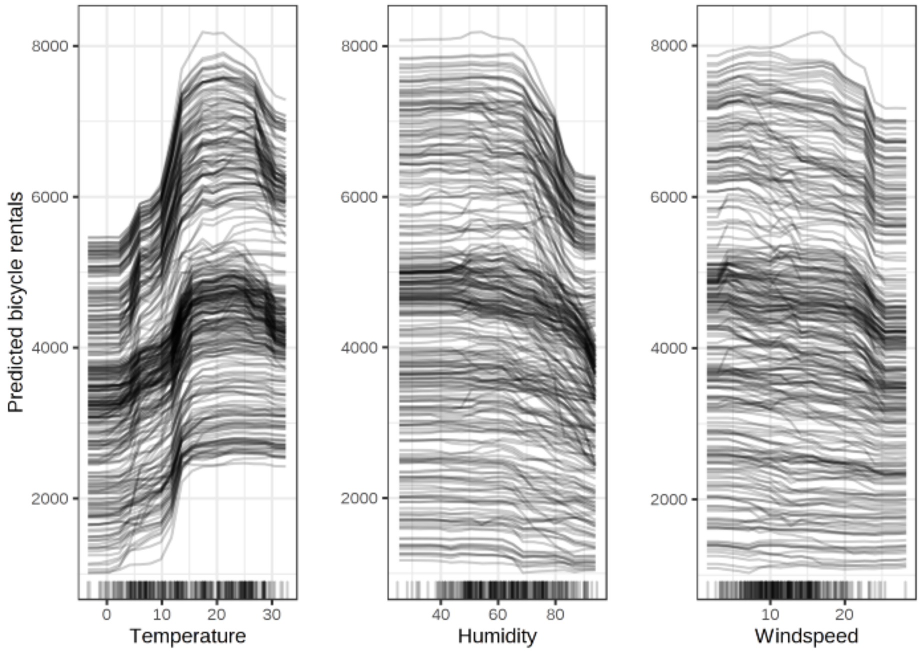

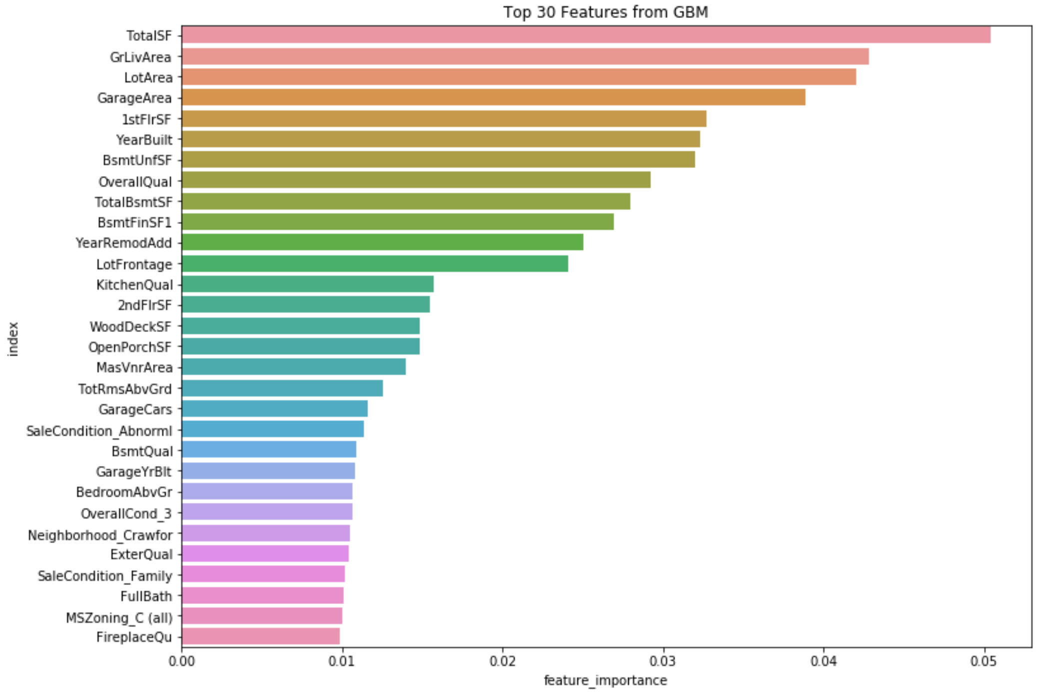

DataBeacon

OPENING THE BLACK BOX

- Local techniques
  - Study a small portion of the ML model
  - For example, an exact layer of a neural network
- Global techniques
  - Study the model as a whole.
  - For example, in/out of the complete neural network

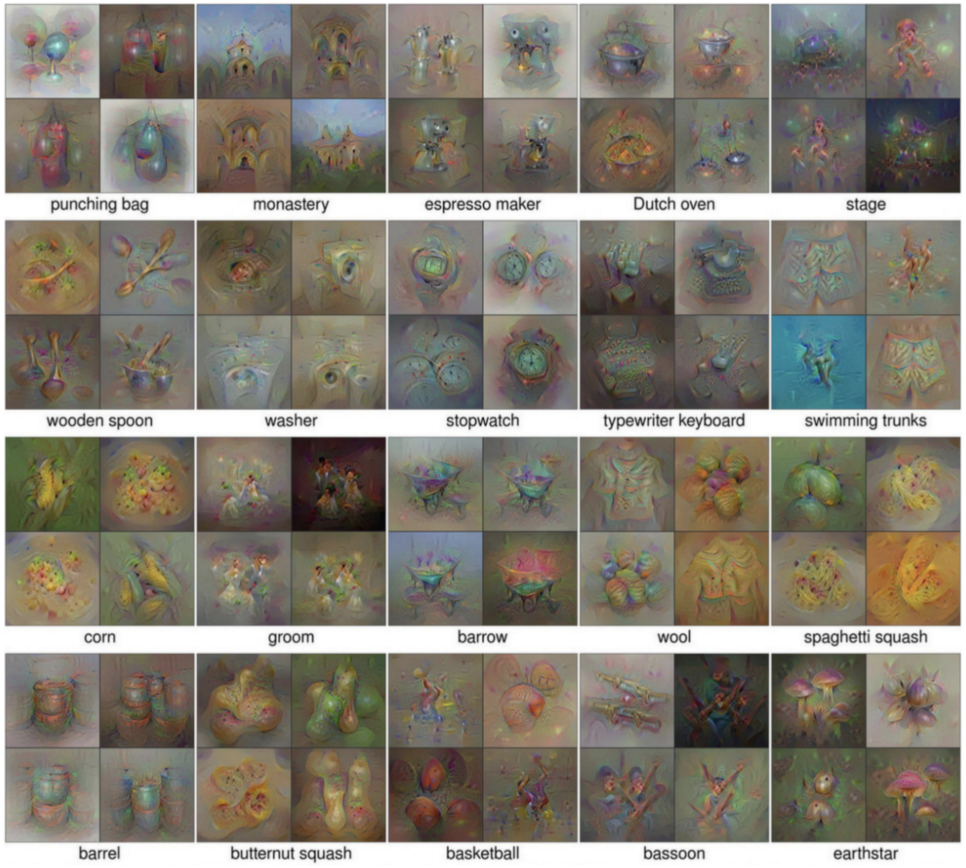

DataBeacon

WE HAVE THE TOOLS TO INTERPRET ML

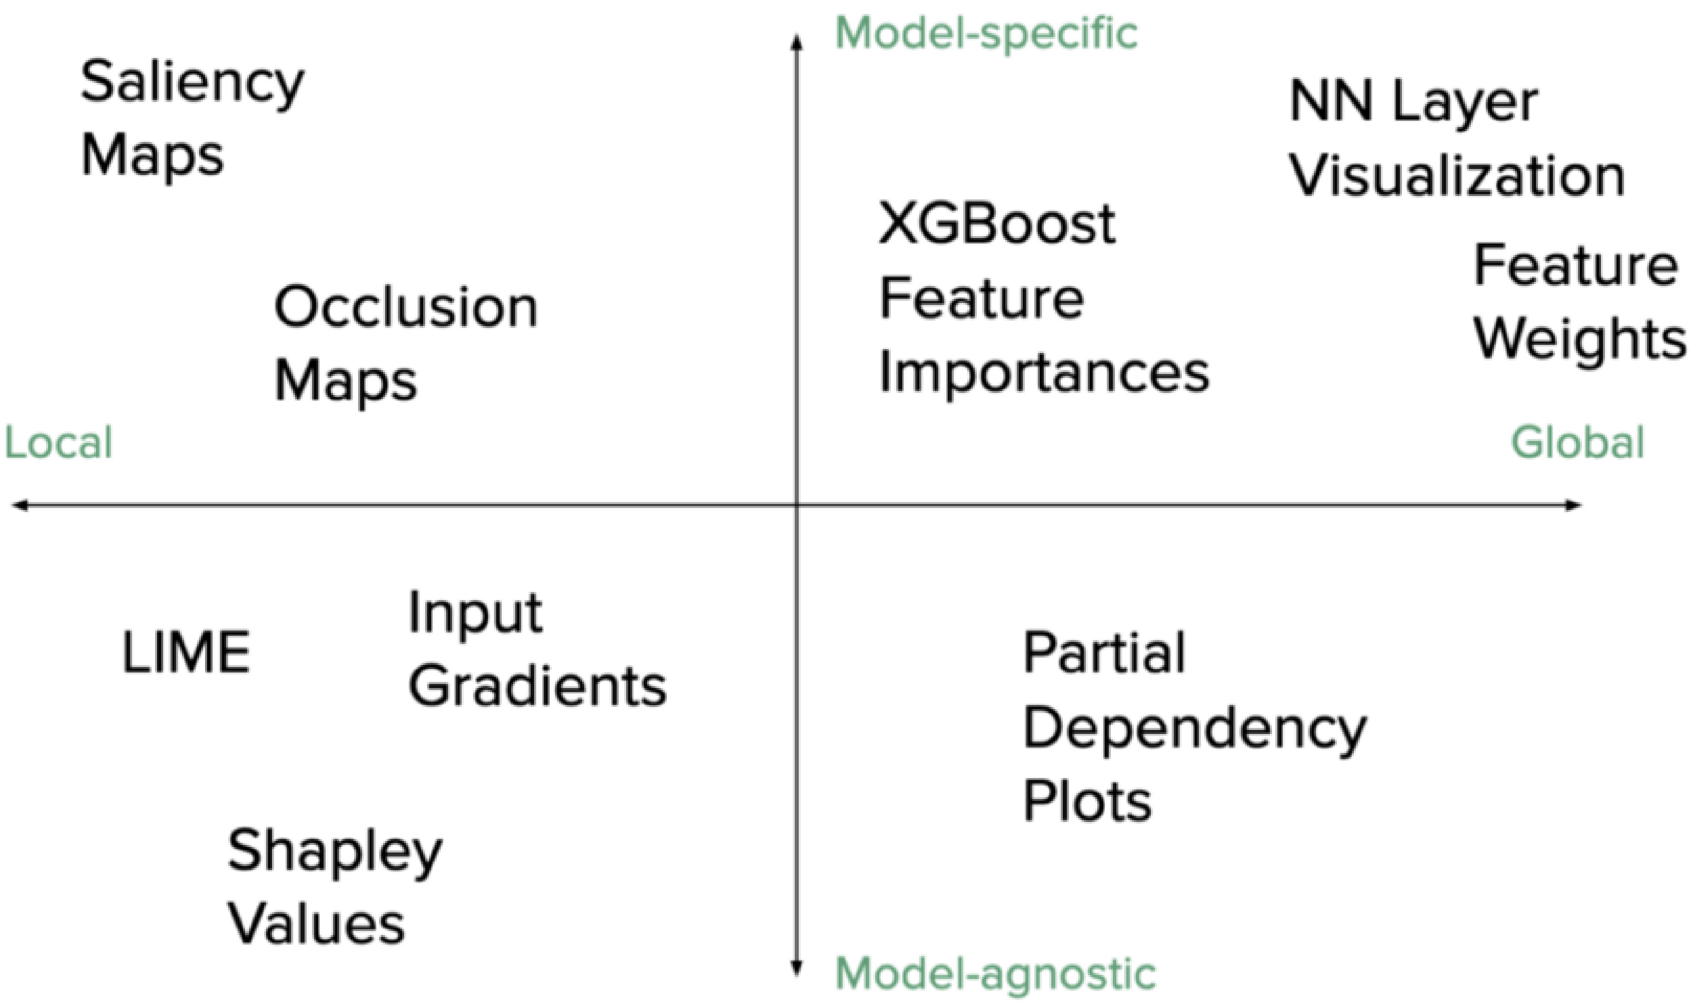

DataBeacon

DataBeacon

www.DataBeacon.aero
info@databeacon.aero