It is about Augmented Intelligence, not Artificial Intelligence

Vilnius, 29th of November 2018

Let’s define the Use Case

Challenges for Fraud Detection

Low Detection Rate

ONLY ~40%

of fraud cases are detected

Many false positives

99.5%

of cases are not fraud related

High Fraud Loss

Tens of Millions

€ lost each month

Fast evolving fraud sophistication

Ambitions for Fraud Project

- ✔ The Bank advanced analytics blueprint
- ✔ Reduce false-positives & Enhance fraud detection rate
- ✔ Blue-Print approach to real time scoring of transactions

teradata.

Fraud Types – Customer Initiated

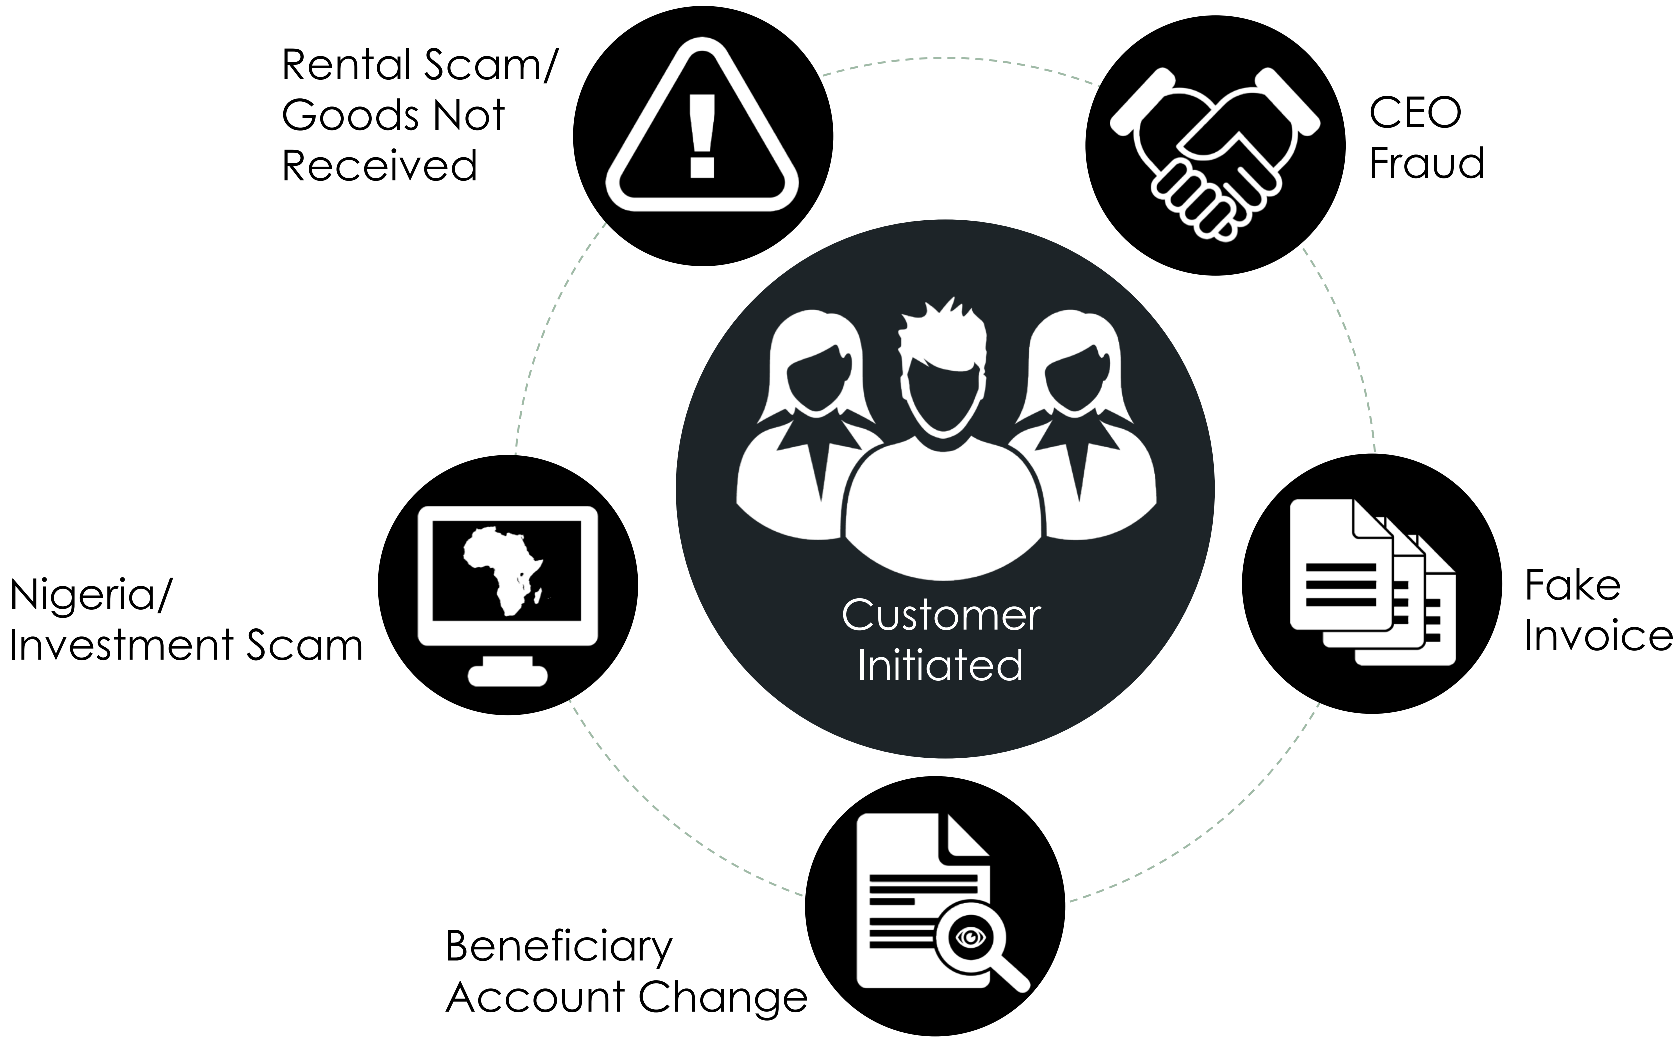

Fraud Types – Fraudster Initiated

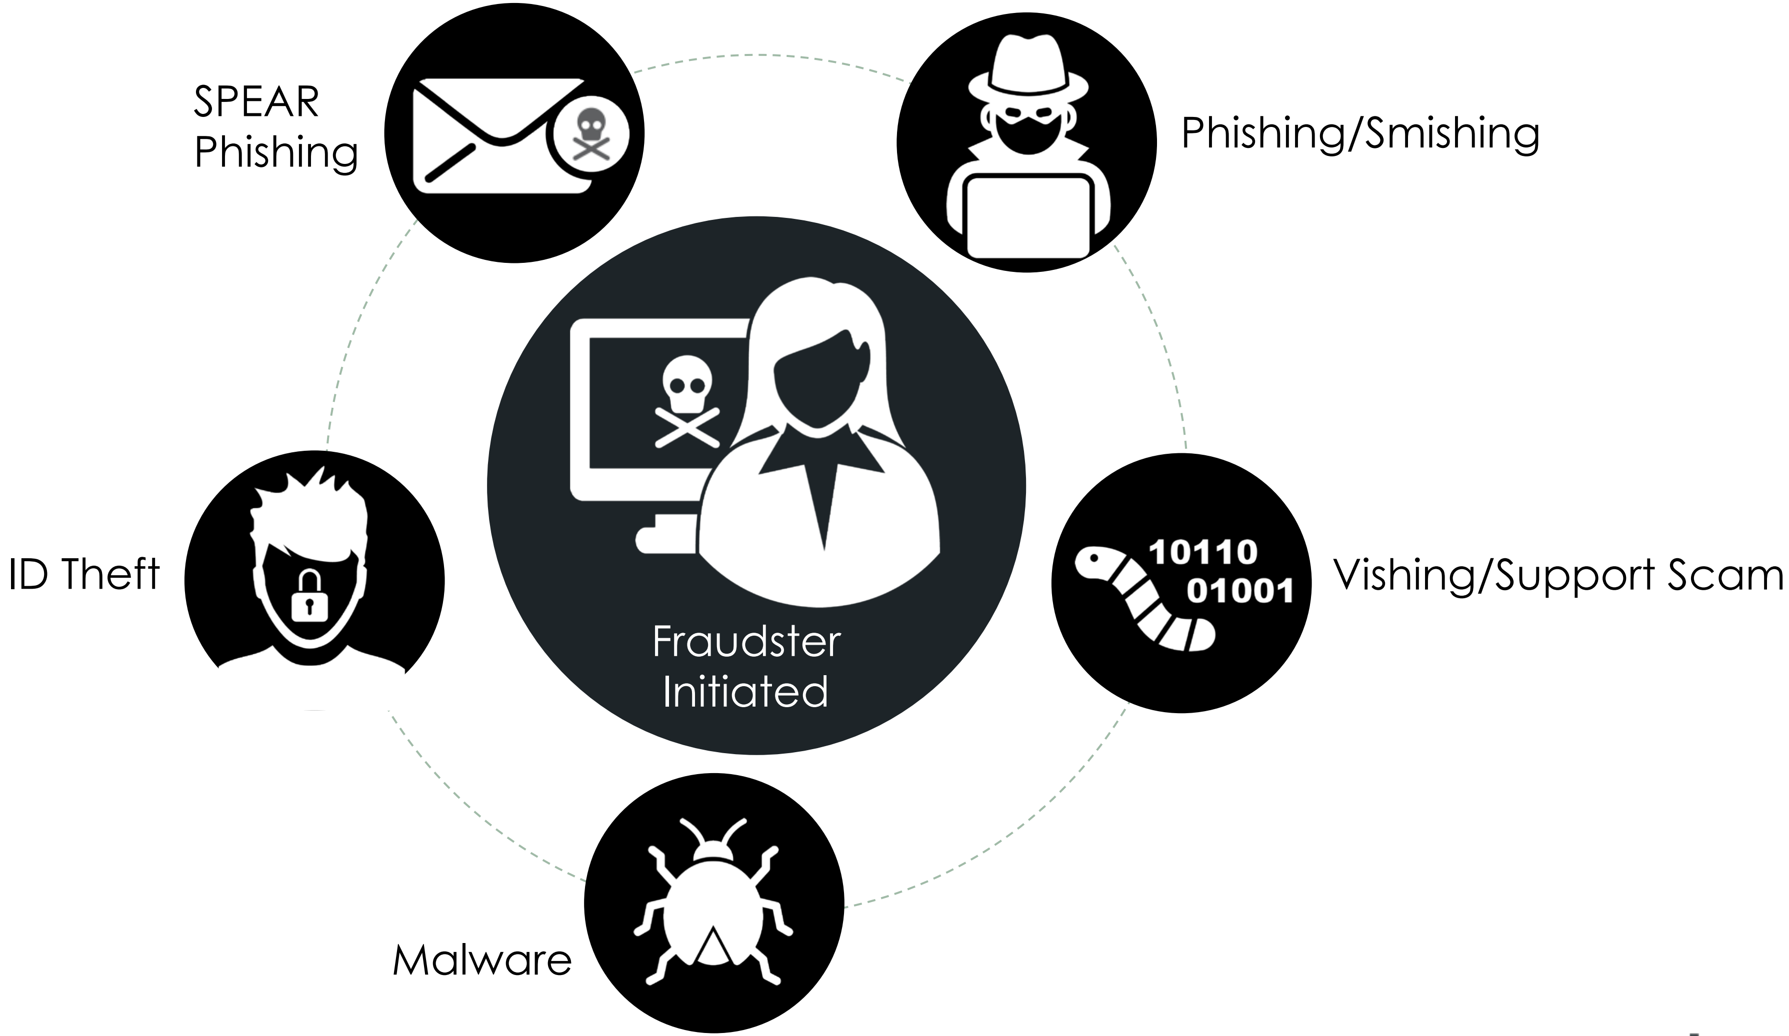

Modeling Challenges

- **Class imbalance**
  (100,000:1 non-fraud vs. fraud)
- Assigning fraud labels from historic data
- Fraud is ambiguous
- Not all features available in real-time
- Most machine learning sees transactions atomically

teradata.

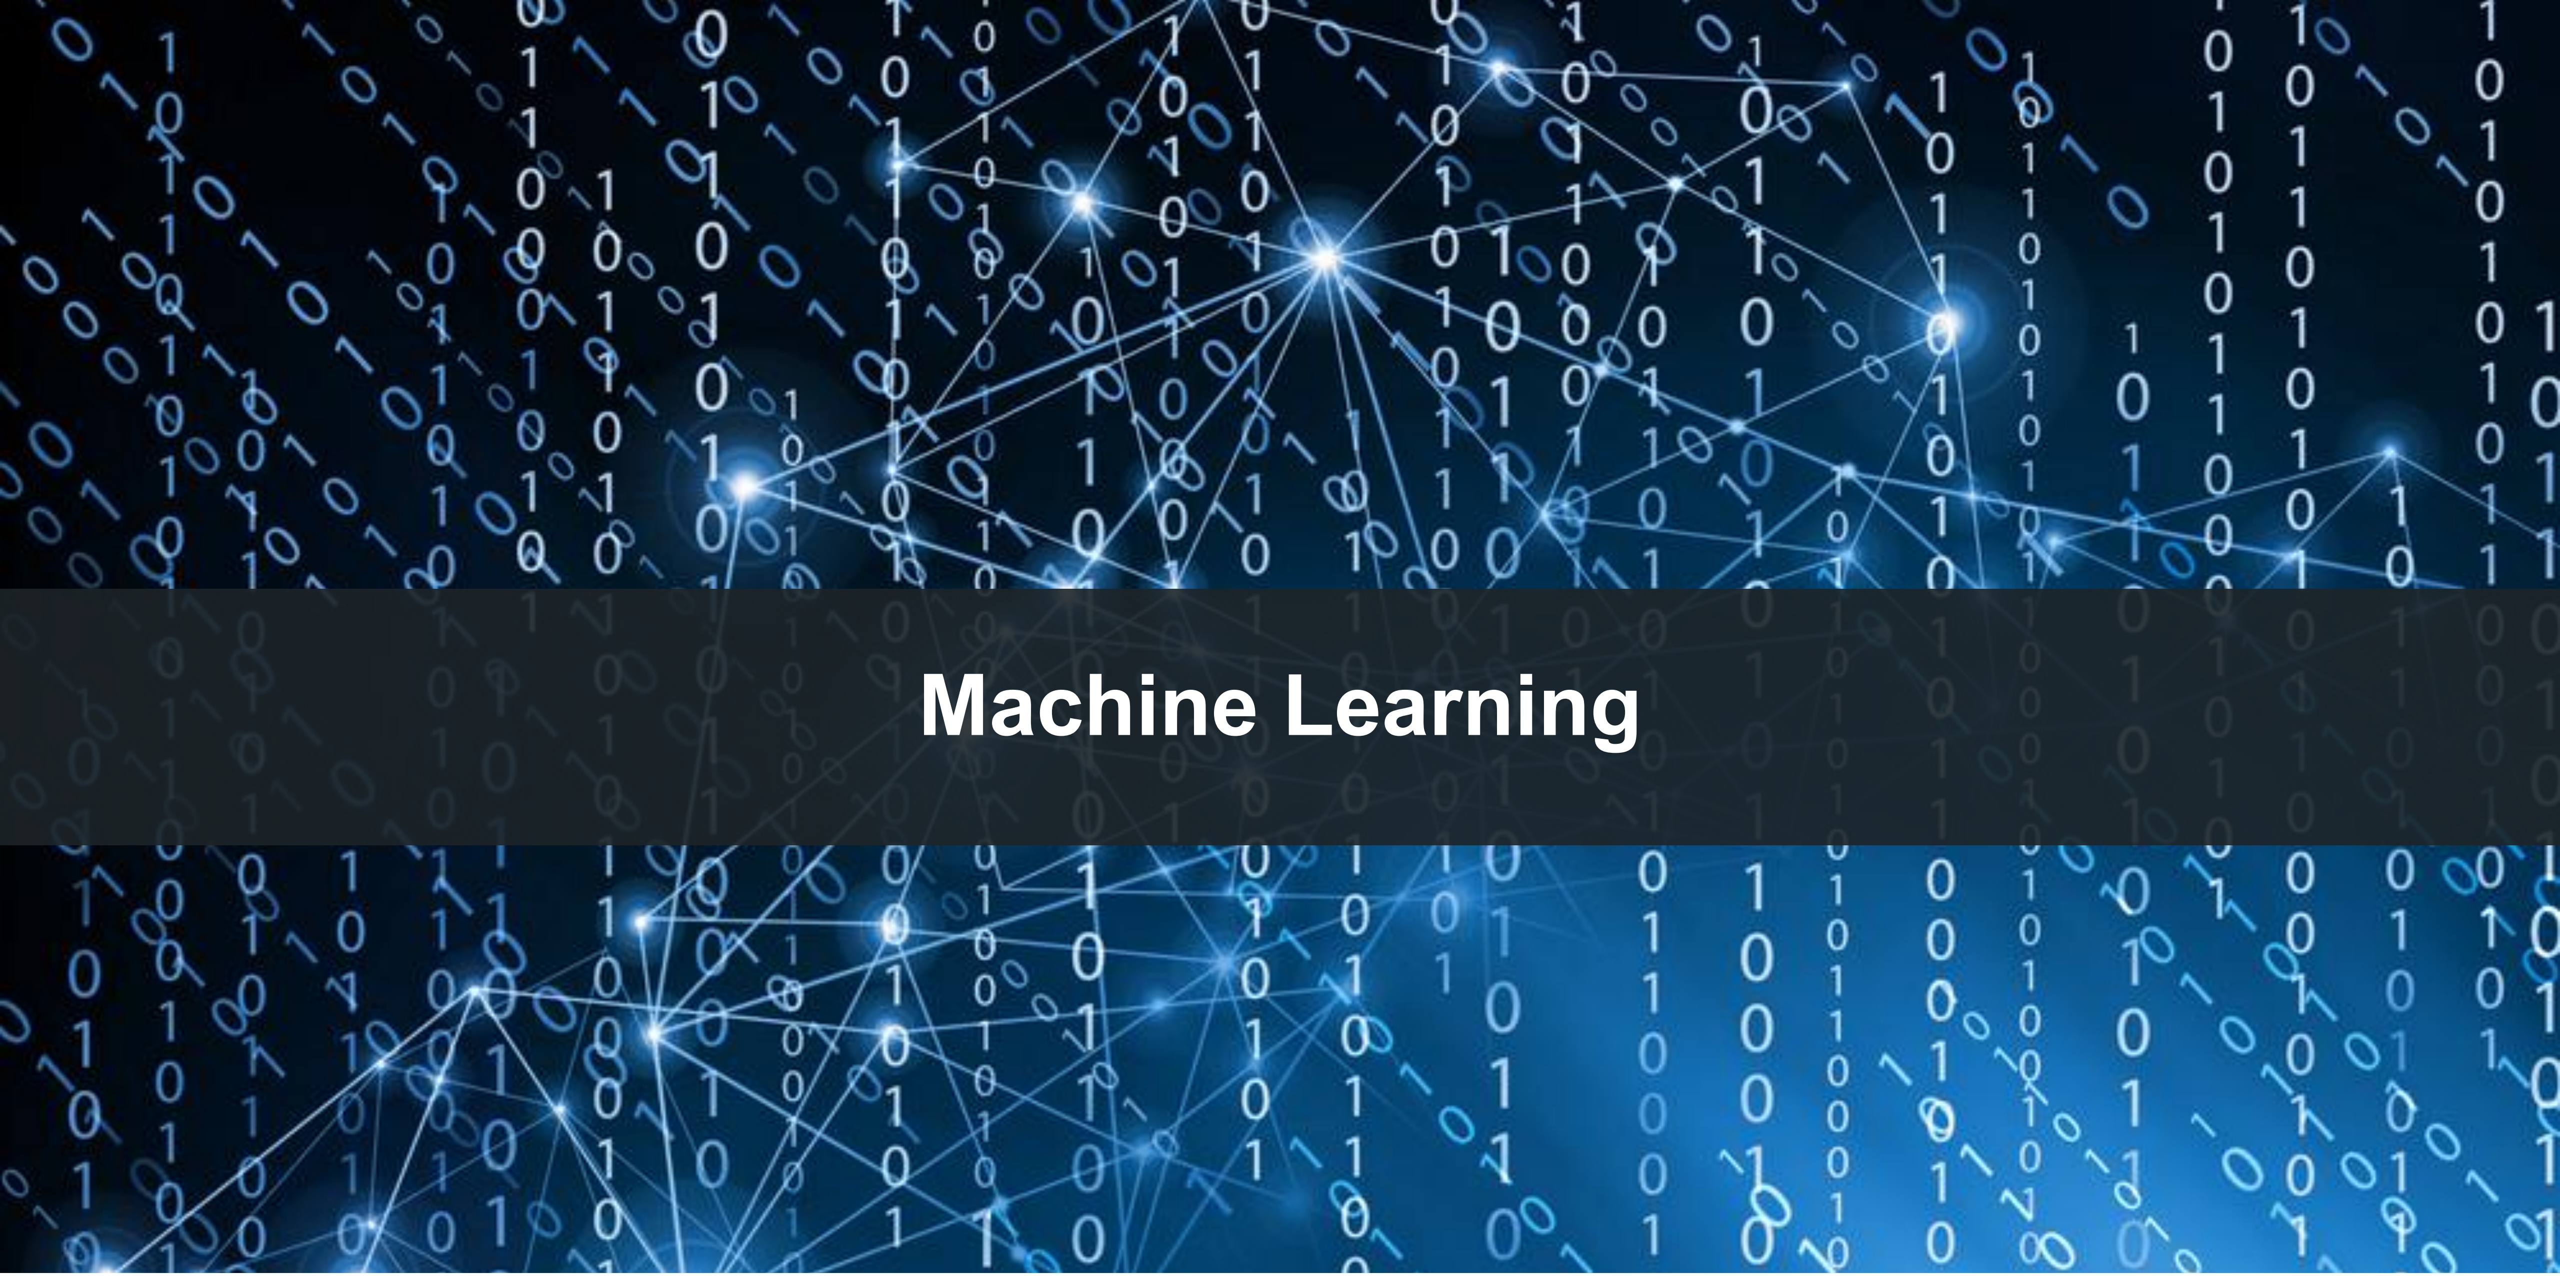

Machine Learning

teradata.

Blue Print vs. Black Boxes

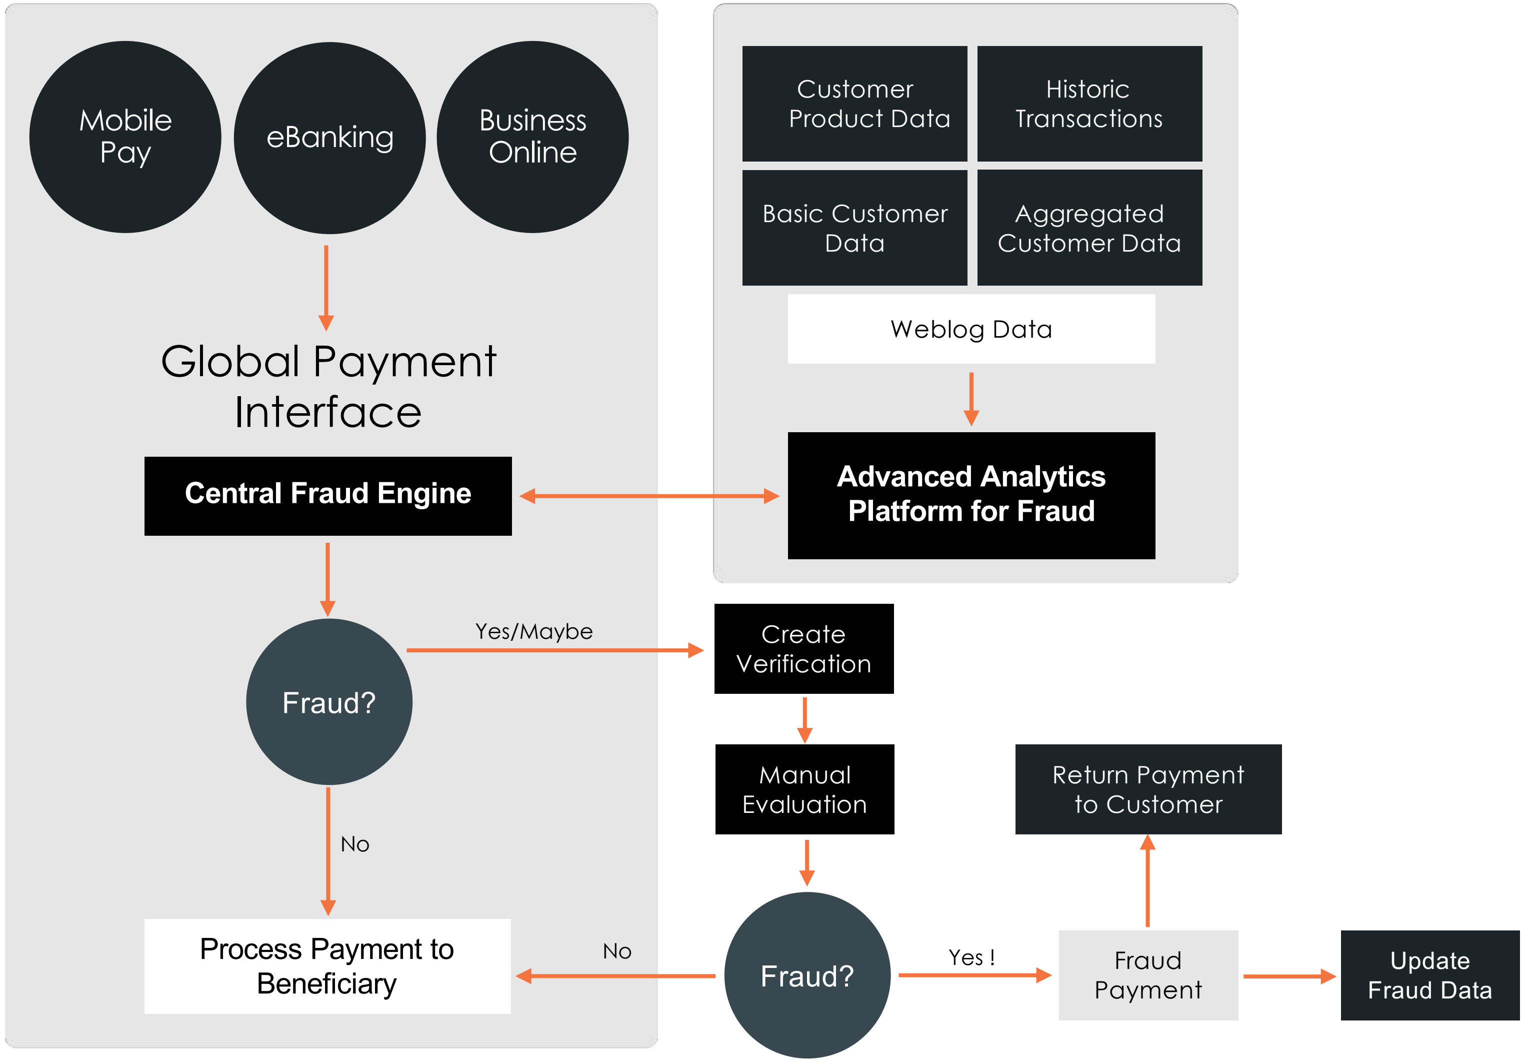

Advanced Platform for Fraud: Data-Driven Approach

Advanced Analytics Platform Architecture - Components

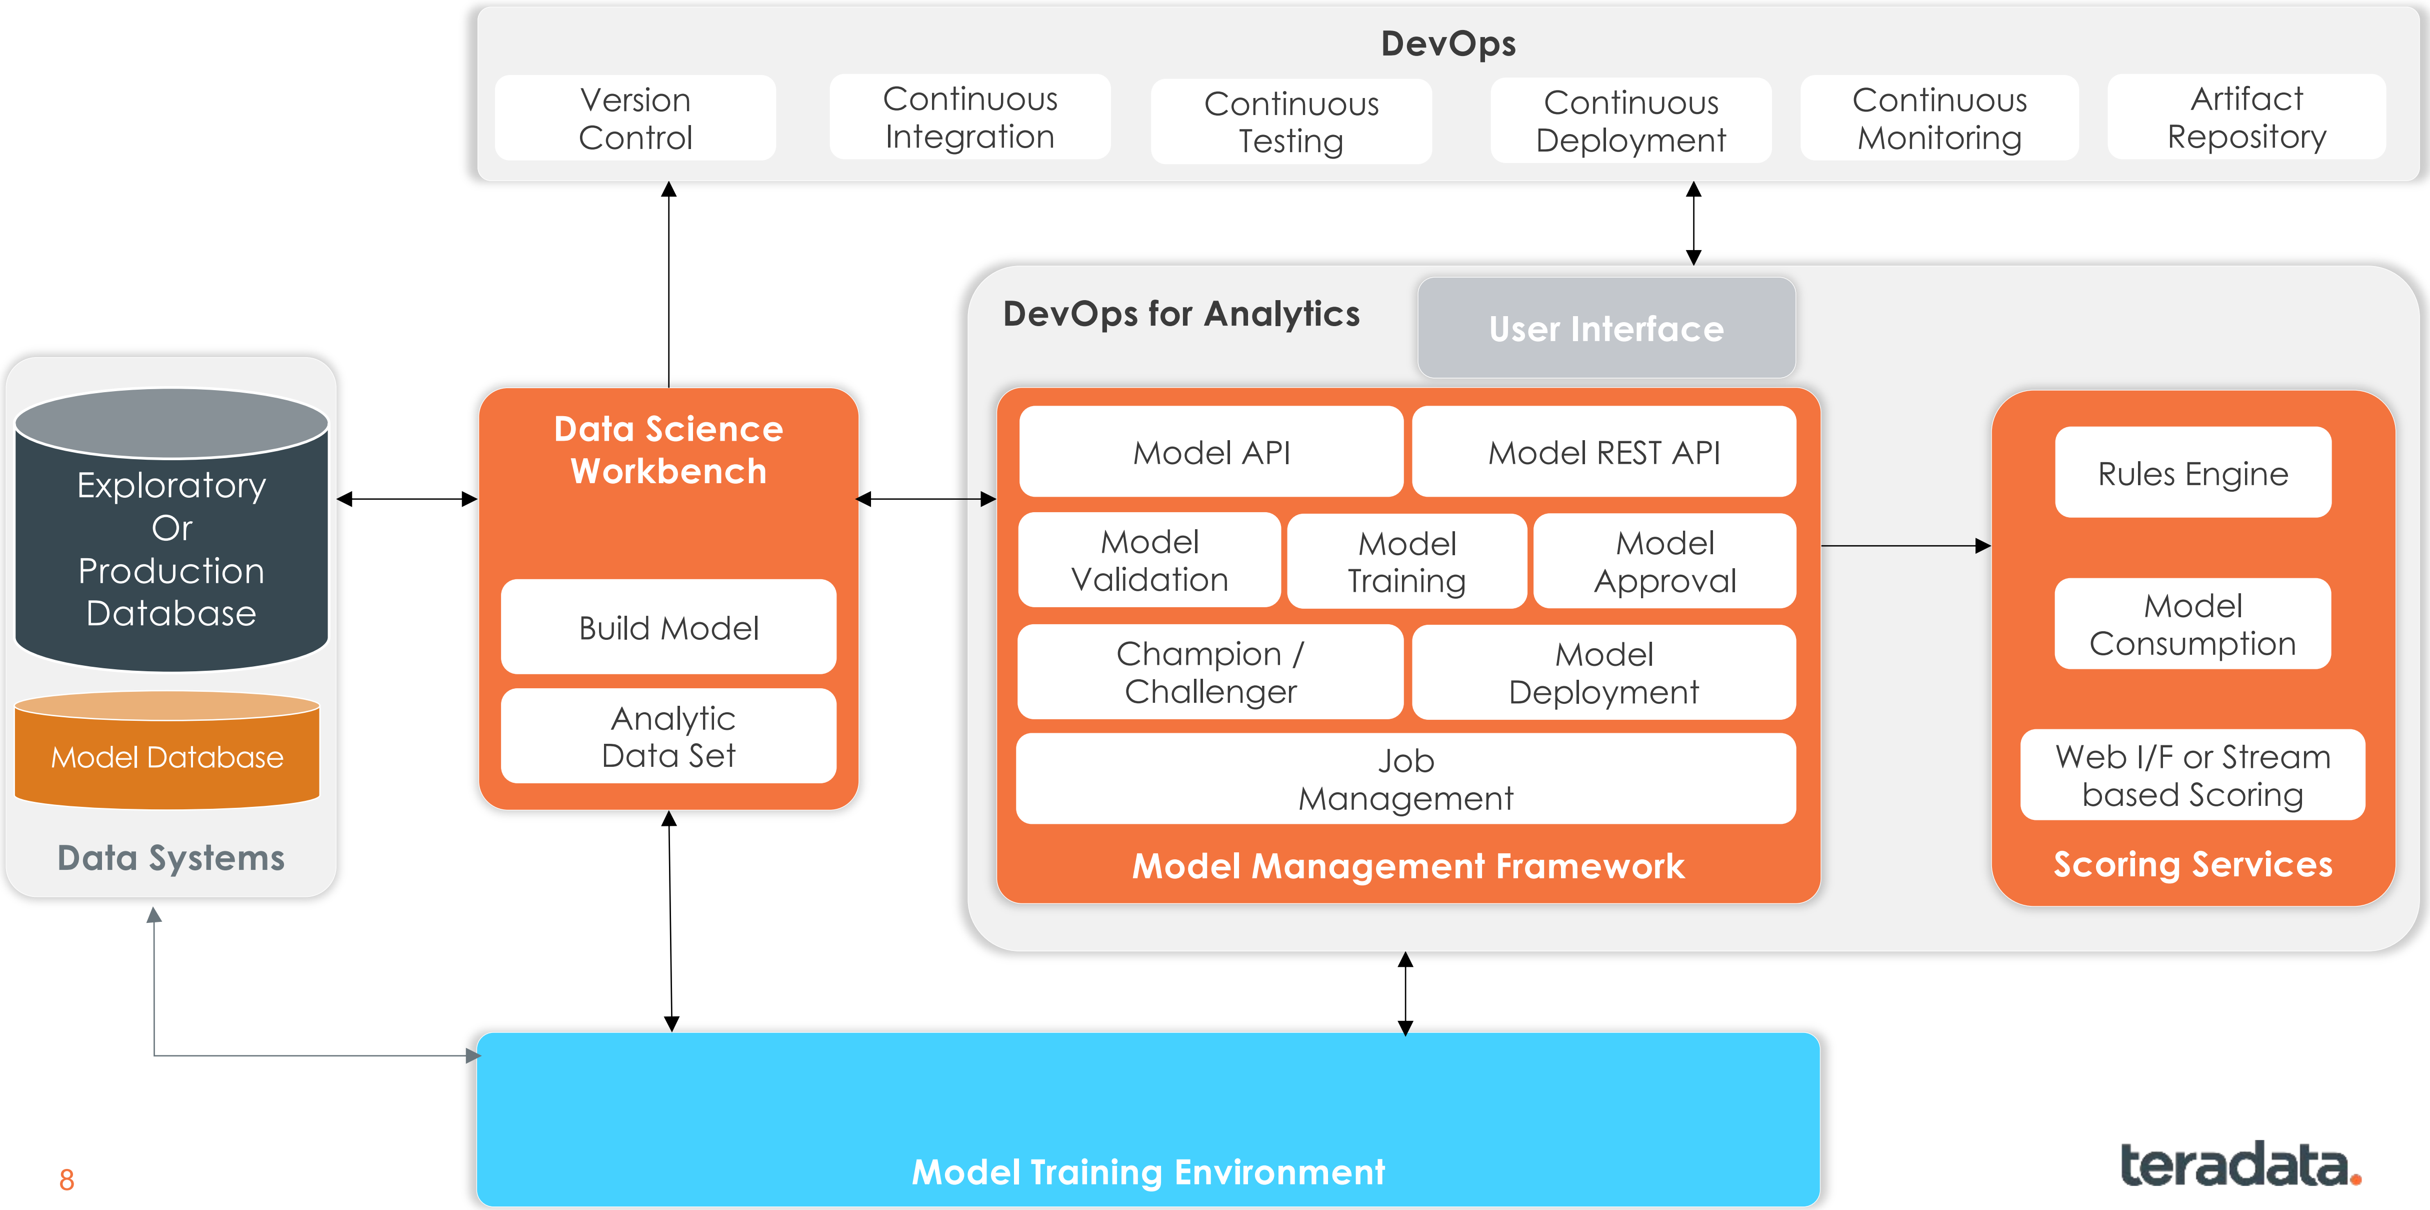

Advanced Analytics Platform Architecture - Components

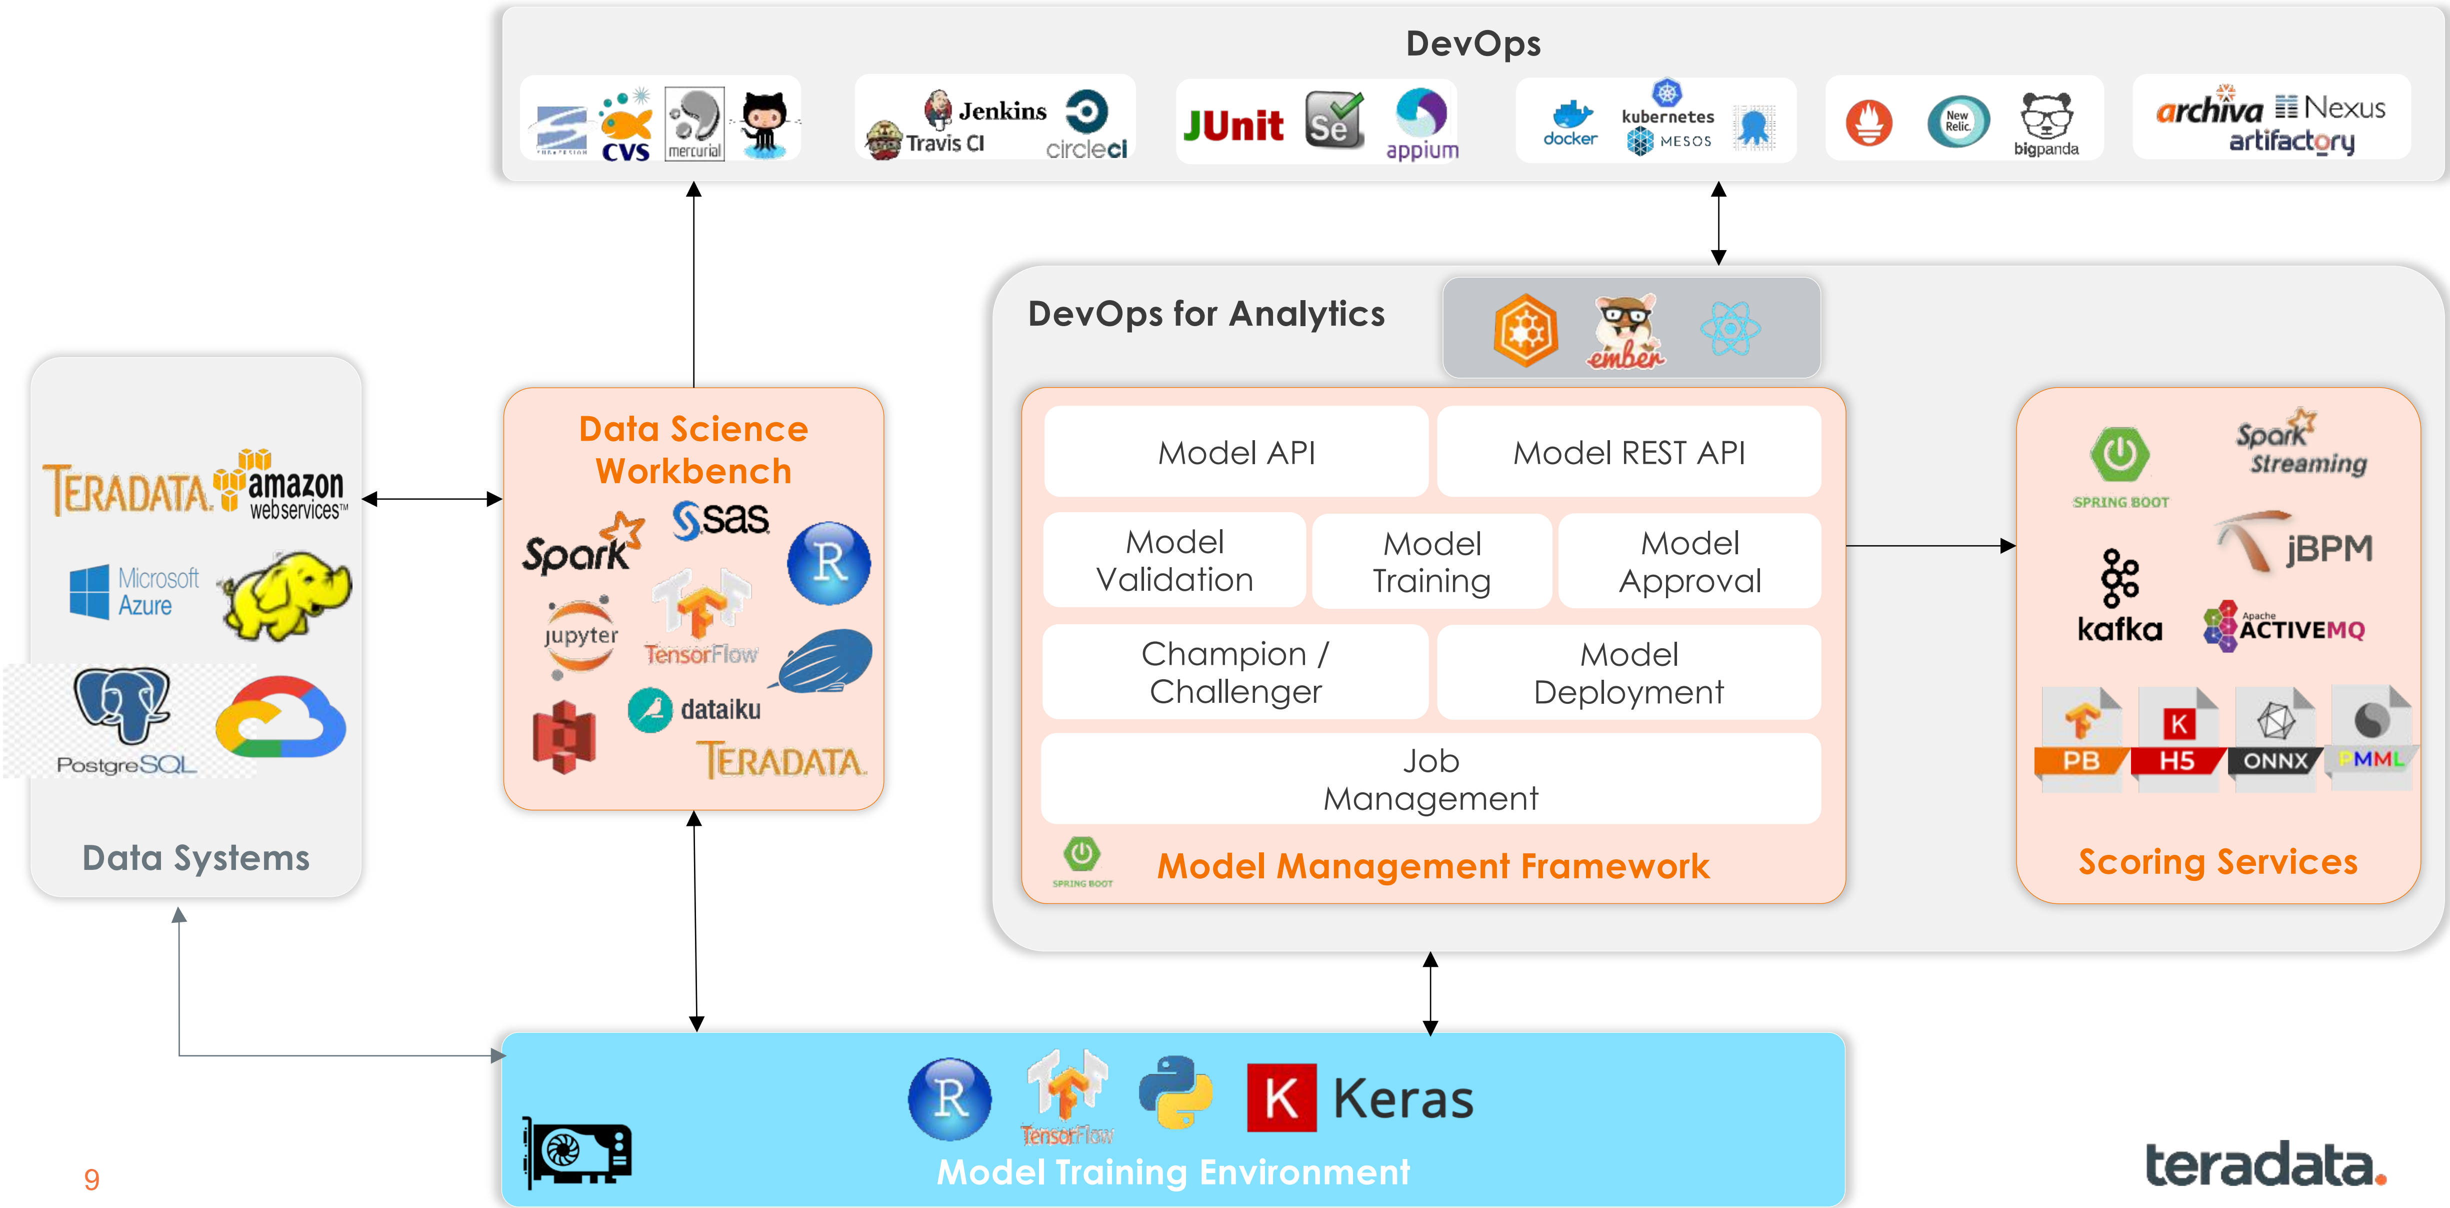

teradata.

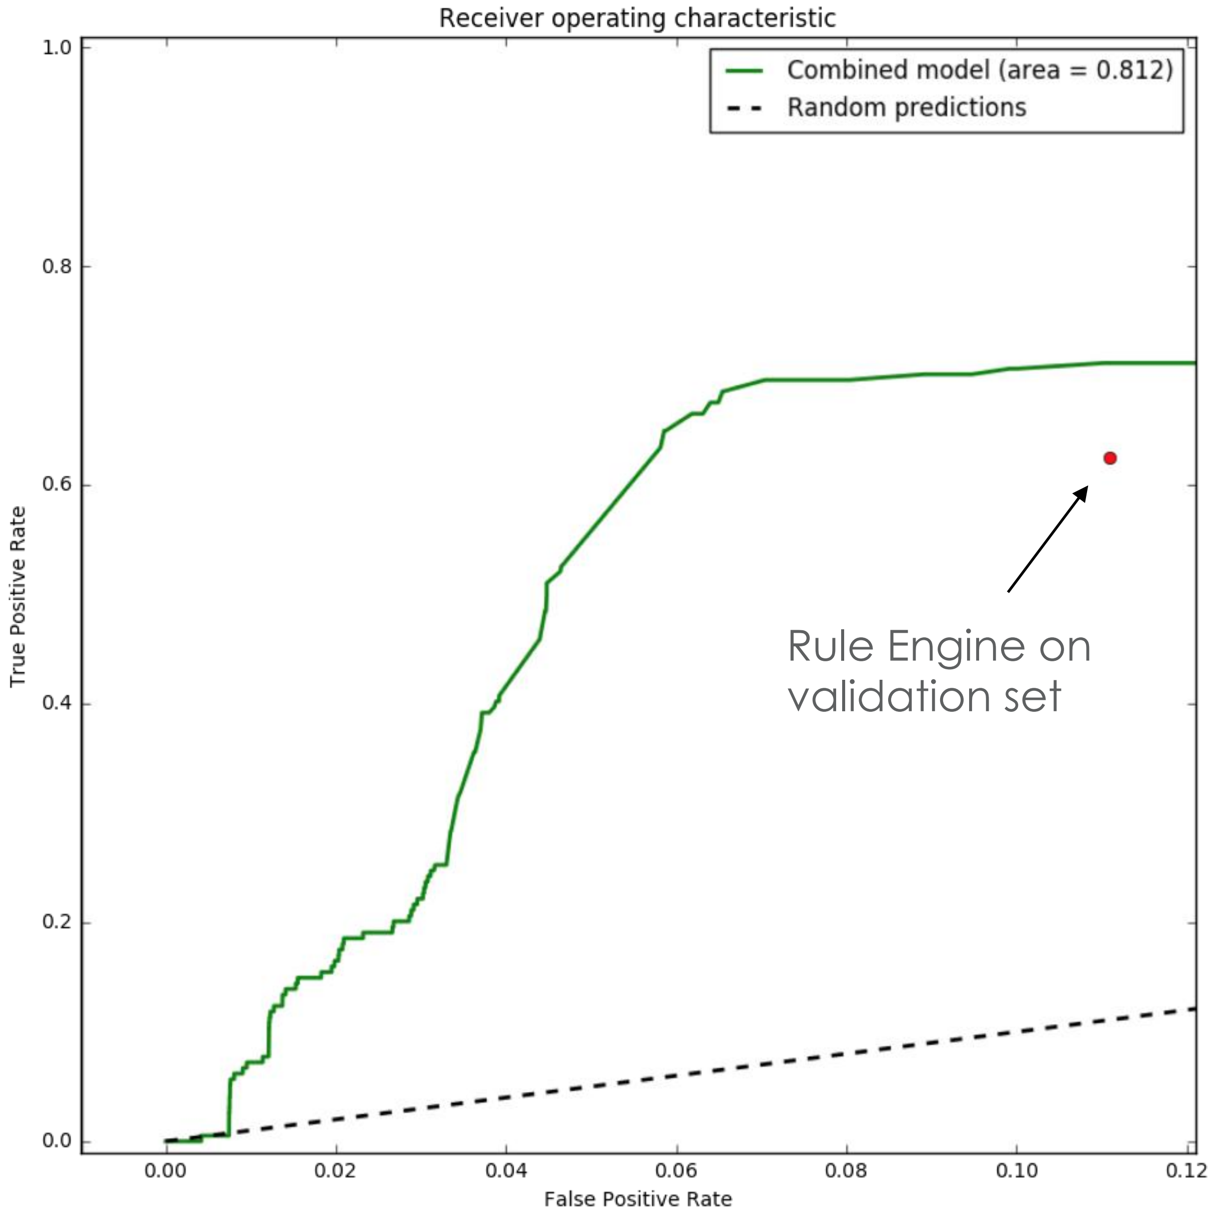

Machine Learning Results

(Live System: 60 transactions/sec.)

Ensemble of boosted decision trees and logistic regression.

From online validation of the model:

- 25-30% false positive reduction, with over 35% increase in detection rate
- Opportunity to expand model with additional features, retrain on recent data and add additional models to the ensemble.
- Models can be expanded to additional channels

Can you build trust based in accuracy?

Wolf or Husky?

You created a snow detector

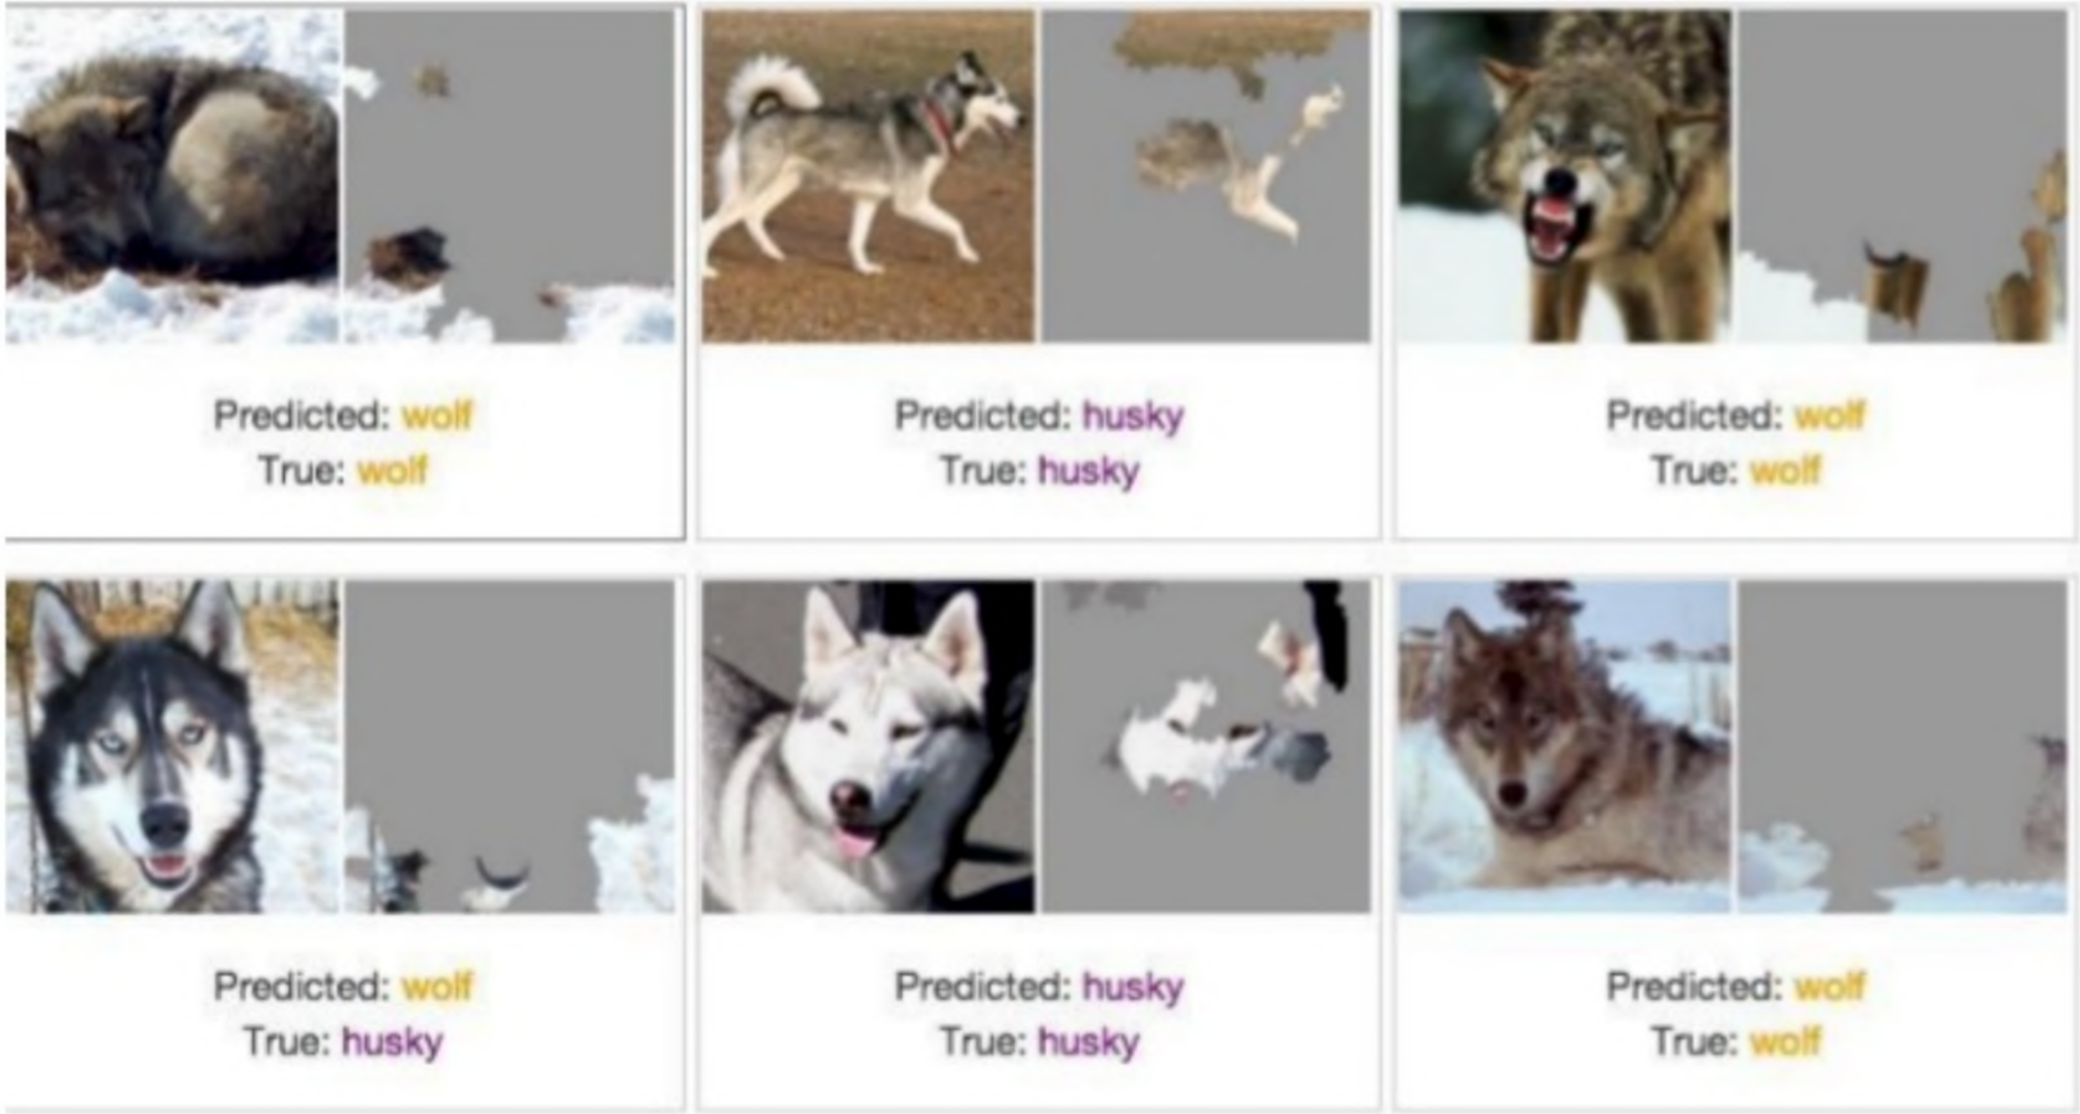

Transparency and Interpretability

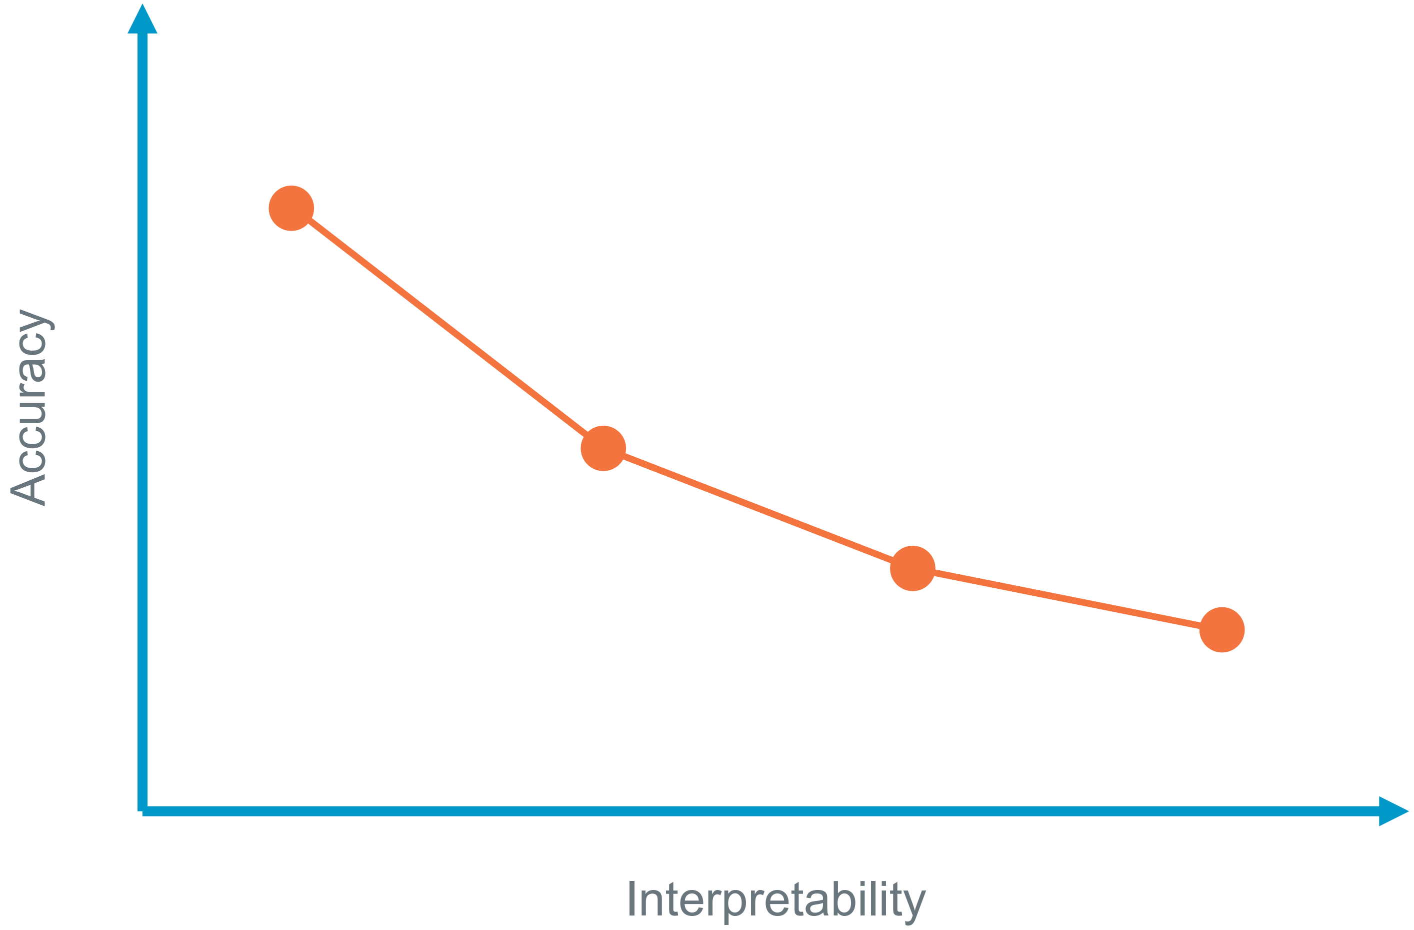

Key Requirement for Black-Box Models: Model Interpretation

- We have deployed LIME (Locally Interpretable Model Explanation) for customers
  - Improves trust
  - Compliance with EU's General Data Protection Regulation (GDPR)

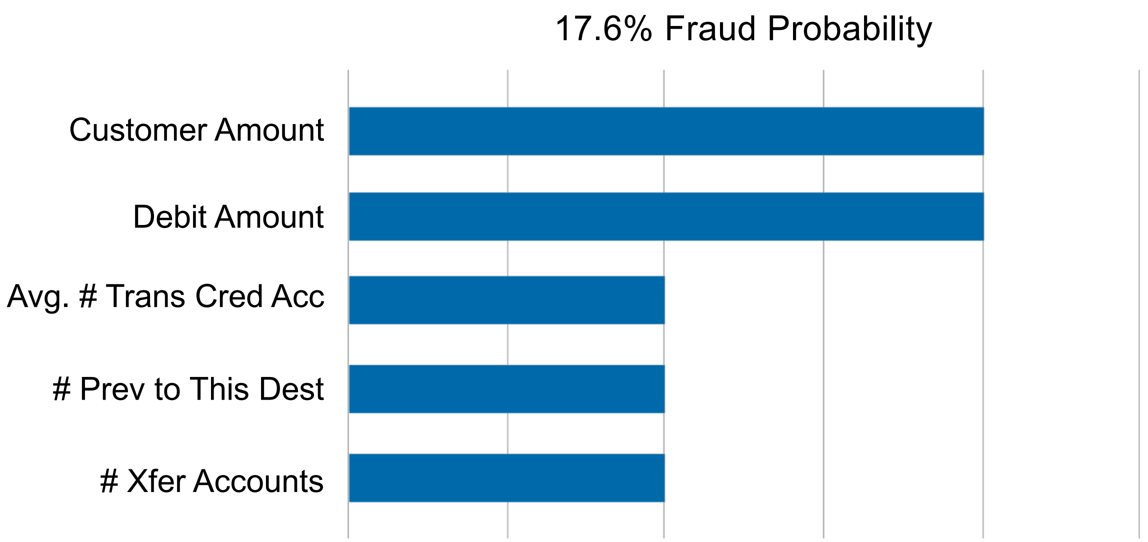

What features are most important to this decision?

```
X% score due to:
+ transfer amount
+ destination country
+ last year monthly spend
```

Deep Learning

teradata.

Deep Learning Opportunity

Current models can only catch ~70% of all fraud cases

Traditional ML models view transactions atomically

Often missed fraud transactions are part of a series

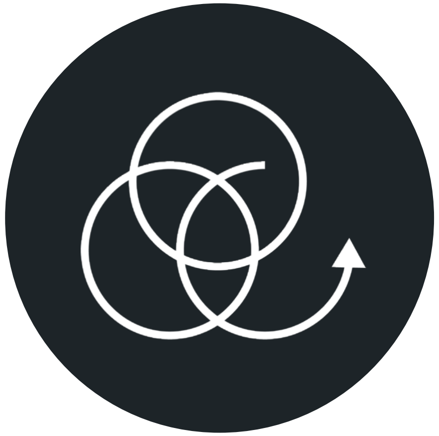

Capturing correlation across many features

teradata.

Three Deep Learning Architectures to Deliver Value

ConvNet

- Designed for spatial correlated features, but by transforming transactions into a 2D image, we can learn temporal correlated features.
- Deeper ConvNet allows learning more complex & general features.

Goal: Learn kernels from temporal & static features to gain insight into the characteristics of fraud.

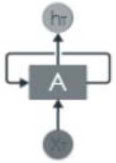

LSTM

- Learn temporal information and classify if the sequence of transactions contains fraud.
- Shares knowledge across learning time.

Goal: Learn transaction patterns within a window. Two solutions can be tested: flag fraud or predict next transaction and define an error.

Auto-Encoders

- Anomaly detection: Learn how to generate normal transactions, potentially large volumes of non-fraud data.
- AE provide a low level representation of the data.

Goal: Build a model that learns how to generate non-fraud data. To detect fraud, define a reconstruction error rate for the fraud cases

How Can We Create an Image From Bank Transactions?

Input

Raw Features

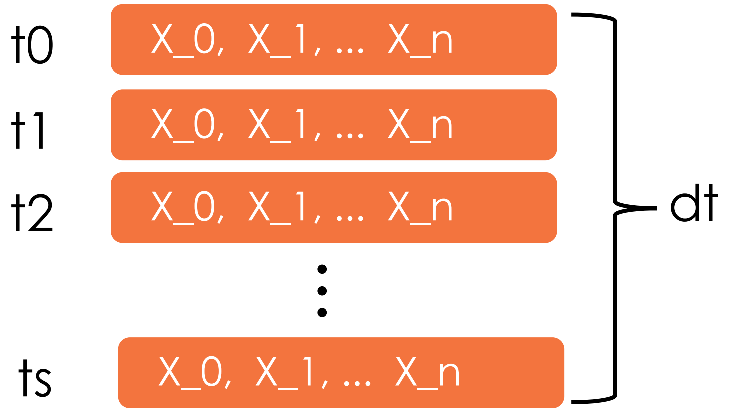

Top k Features Correlation

```
{0: [41, 5, 30, 29, 31, 10, 37, 3],
 1: [42, 40, 32, 15, 35, 2, 16, 31],
 2: [3, 15, 4, 1, 28, 40, 31, 49],
 3: [15, 41, 29, 16, 0, 2, 6, 14],
```

Output

Add correlated features in a clock-wise manner

Image size is: **[10 x 3, 50 x 3, 1]**

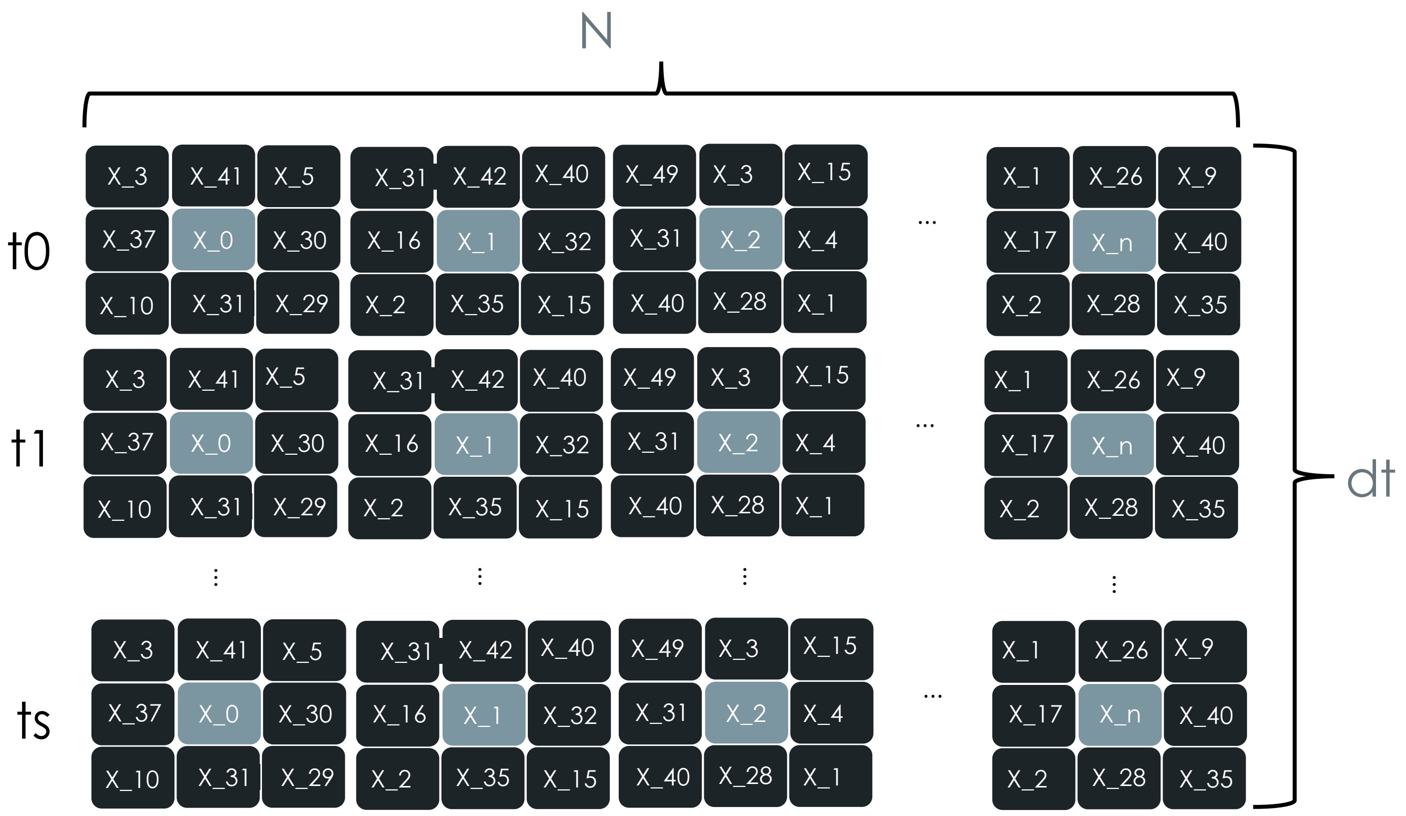

teradata.

Convolutional Layers for Trans2D

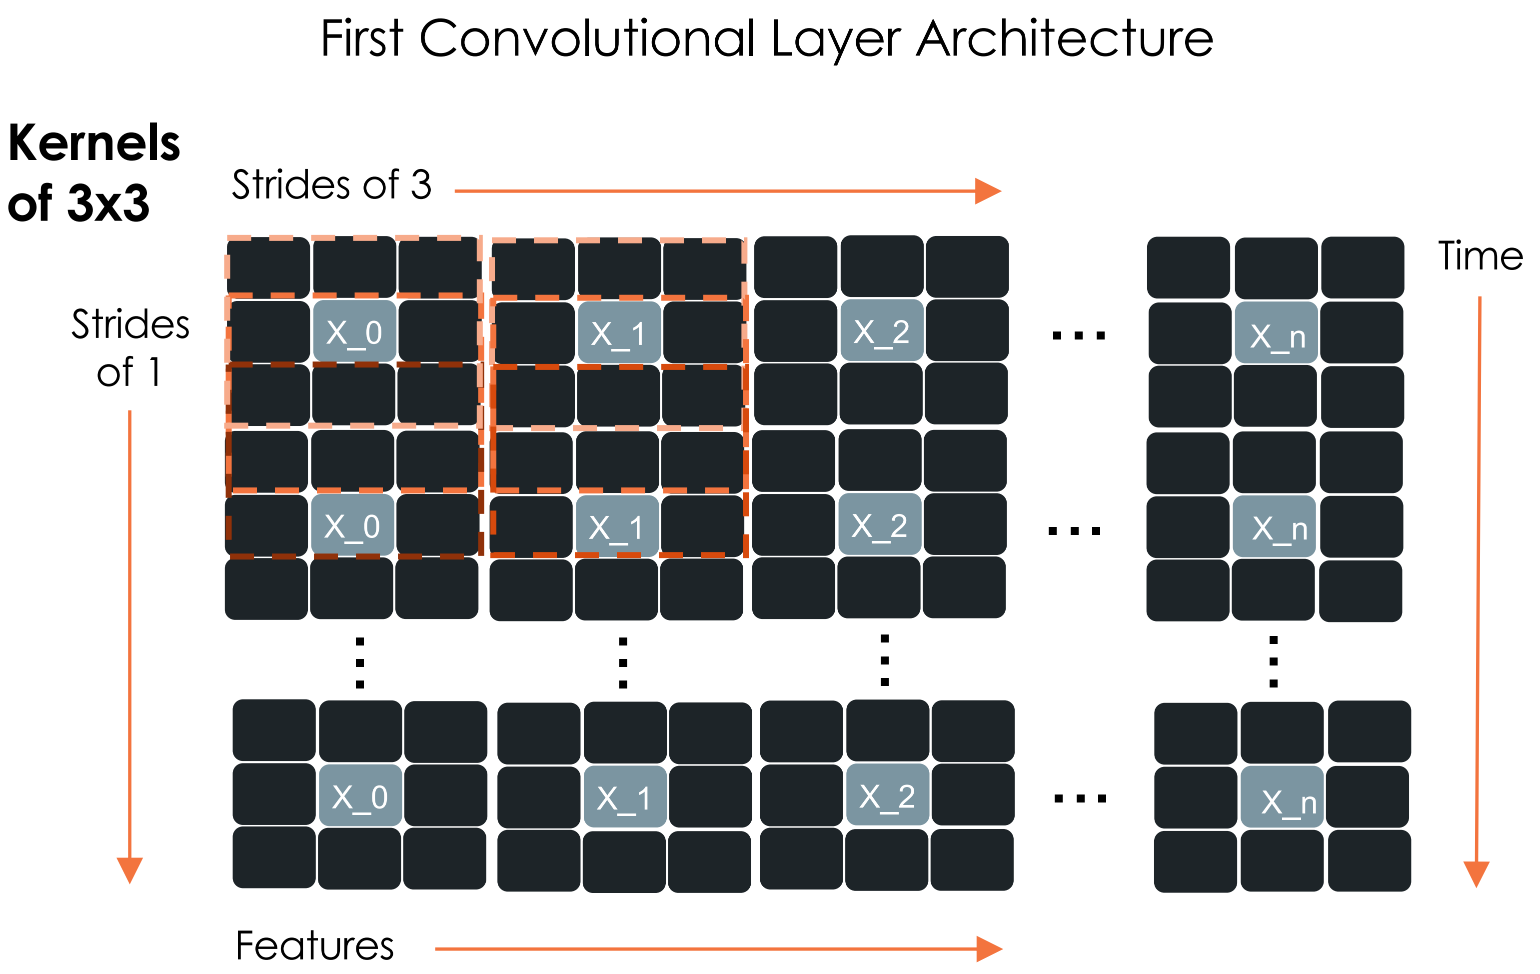

teradata.

2D Transaction Image Example

Non-fraud Transaction Image

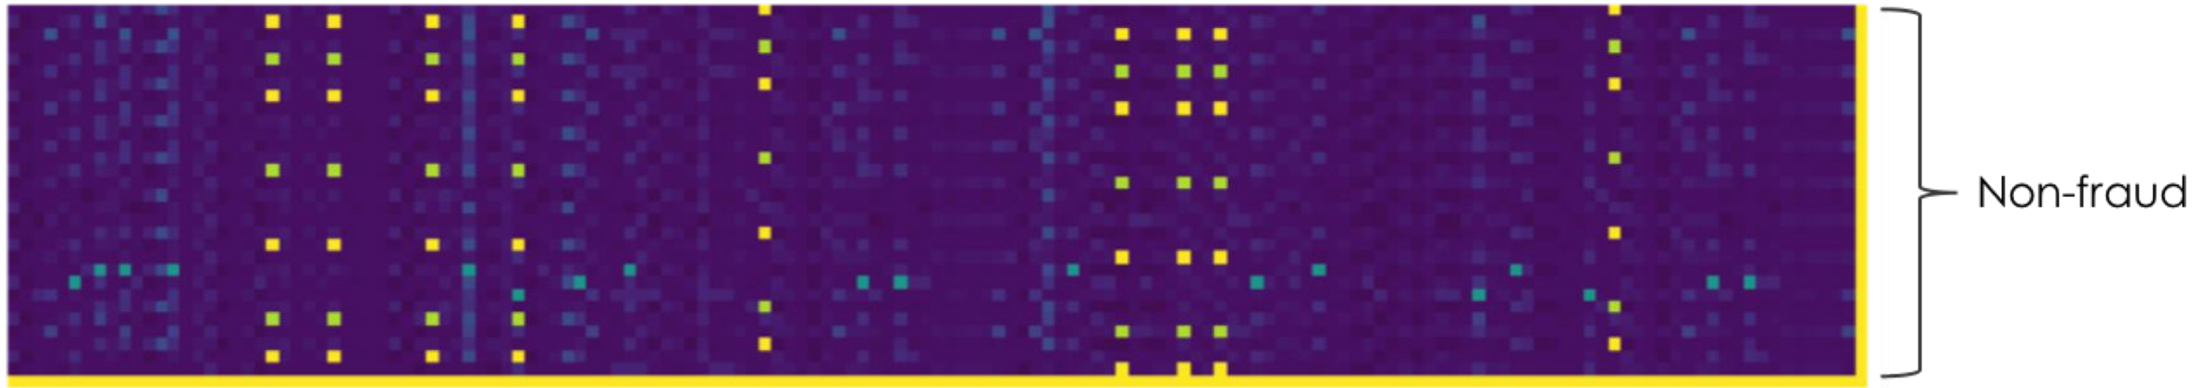

Fraud Transaction Image

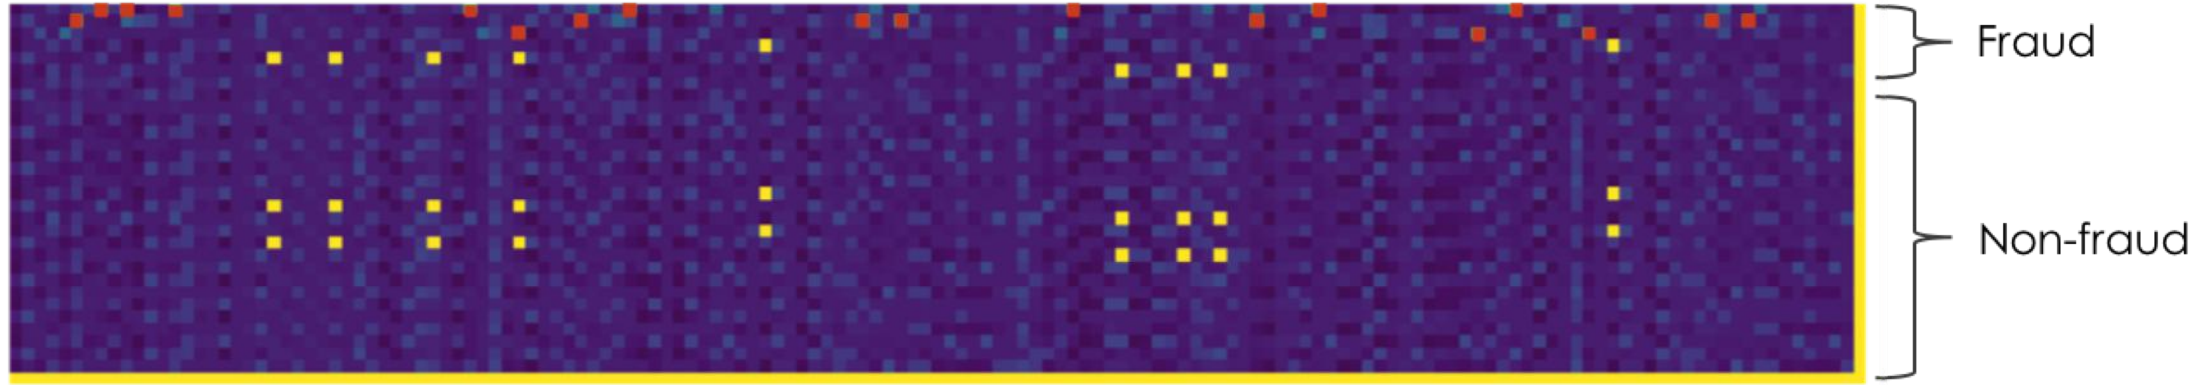

X-axis: features, Y-axis: time

Inside the ResNet model

64 Filters
Activations
After the CNN
Residual Blocks

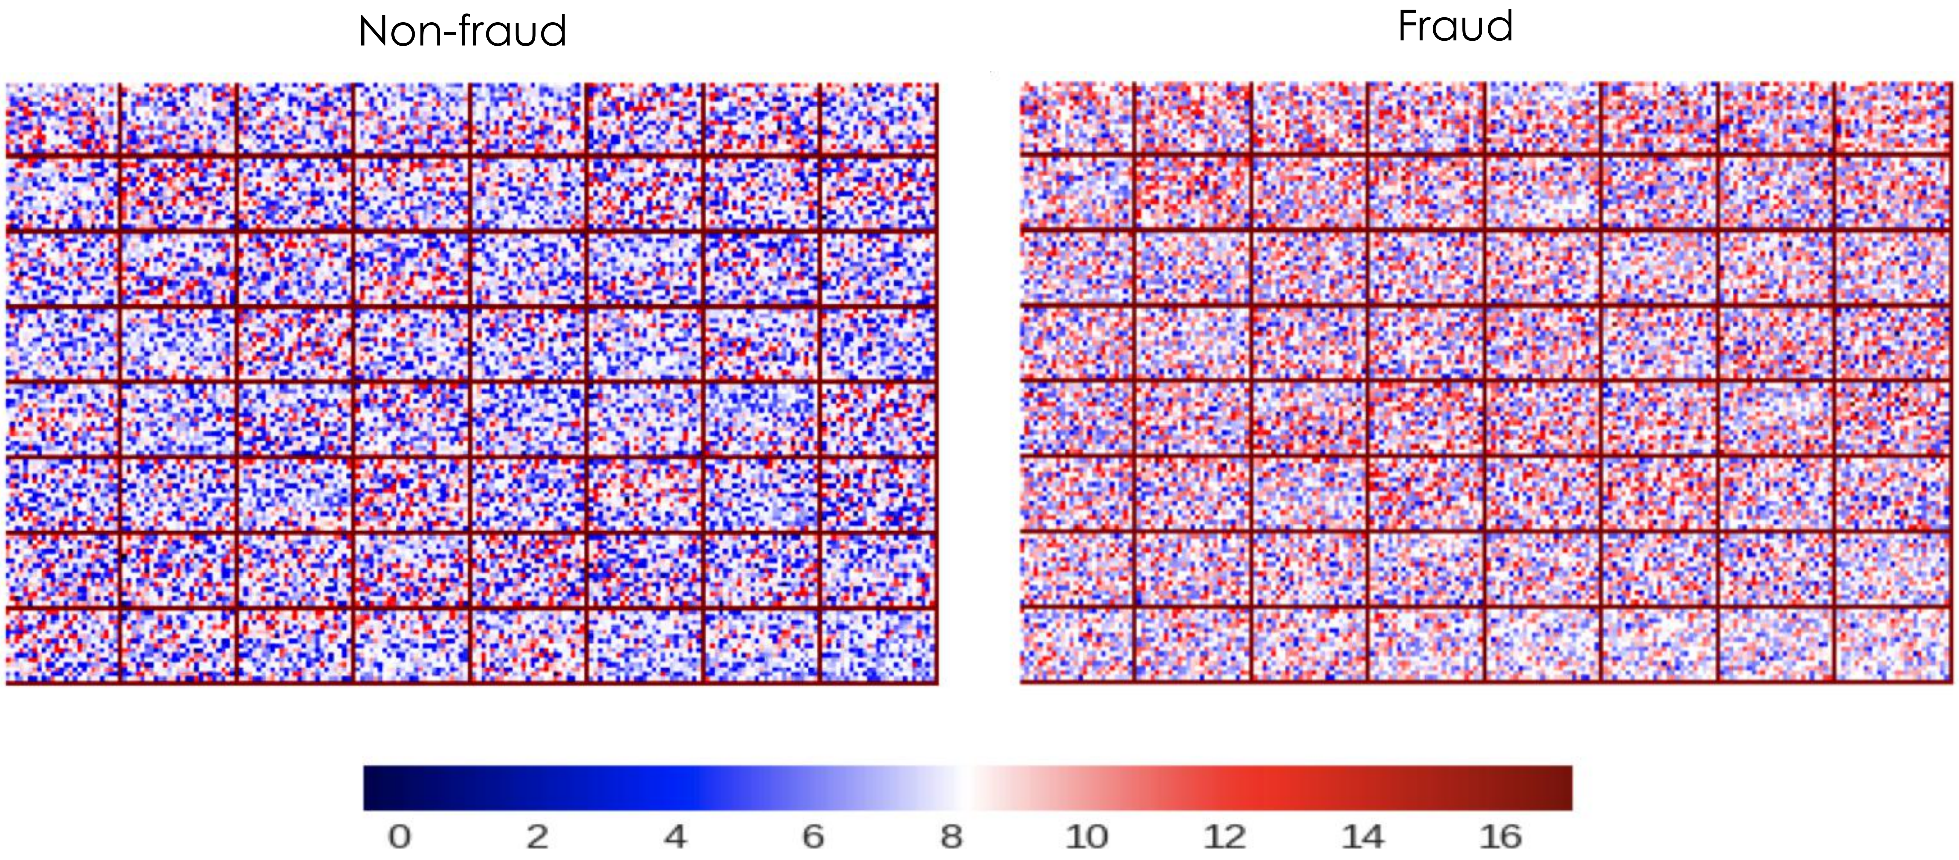

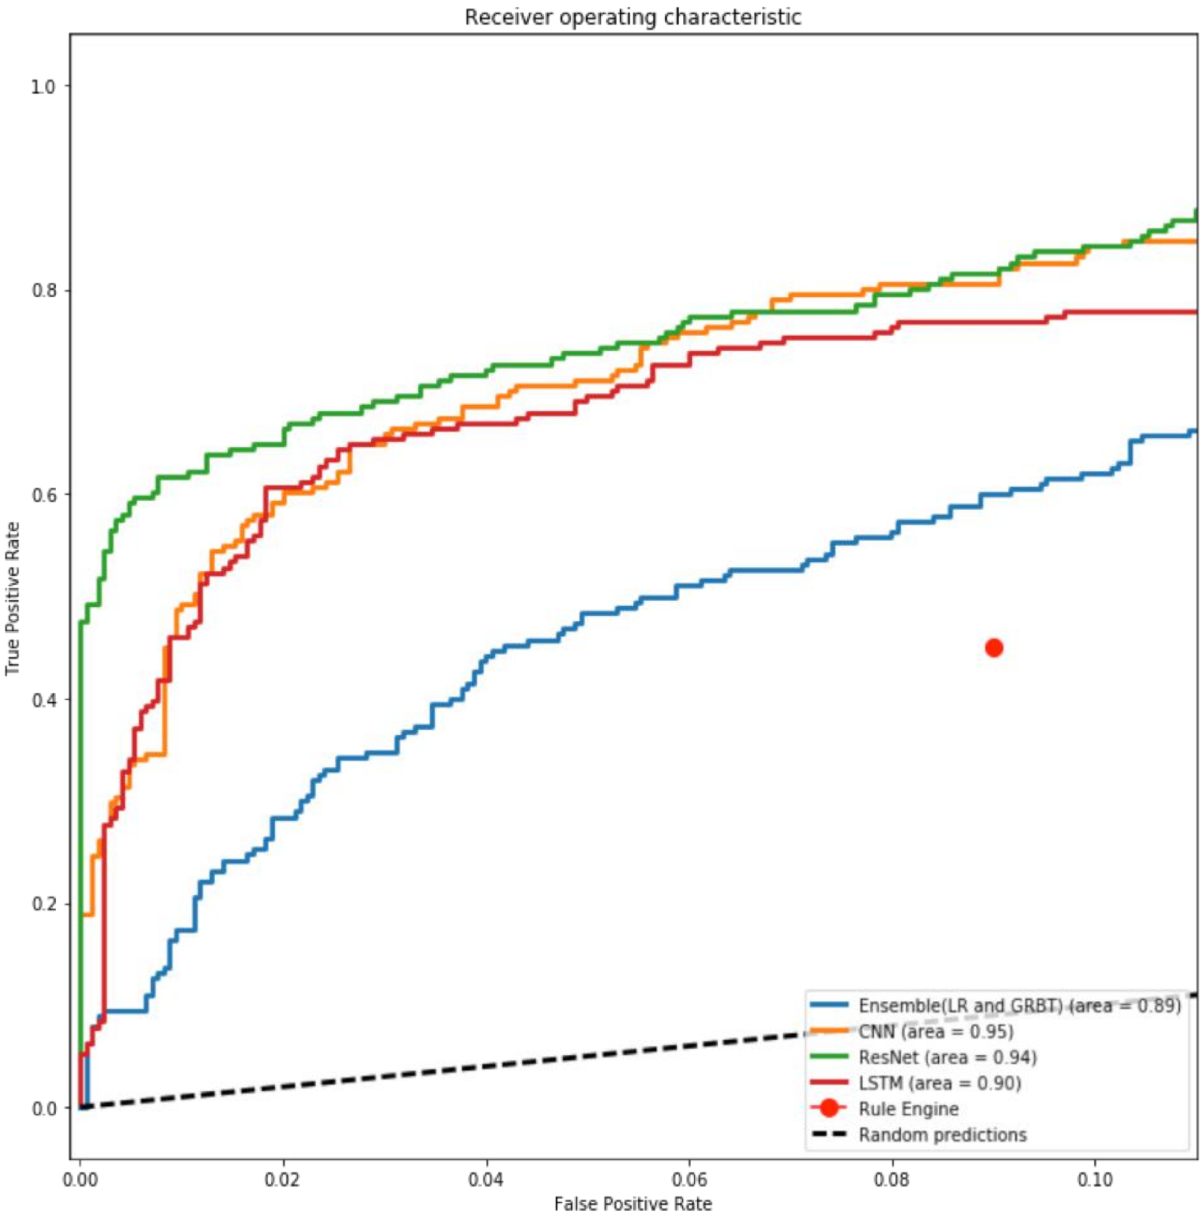

Deep Learning First Results

on the fraud verification dataset

Comparison of the three deep learning models andthe traditional machine learning ensemble model.

- Ensemble model (AUC 0.89)
- ConvNets (AUC 0.95)
- LSTM (AUC 0.90 – looking at 10% of transaction)
- ResNet (AUC 0.94)

Lessons Learned: Takeaways From this Example

Deep learning adoption from pictures to financial transactions

Enhancement of data quality & cluster capabilities with **data ingestion**

Building **Analytics Ops** capabilities to support business units

Leveraging experience from Fraud advanced analytics to deliver extra use cases

teradata.

Big Data Team – Lessons Learned

Success: From PowerPoint to production in 8 sprints

Team effort: Thorough collaboration across IMD, GFU and Think Big

Synergy: Successfully spearheaded innovation in all involved systems

Inspiration: Incorporation of an advanced analytics blueprint sets a generic scene for combatting new challenges in advanced analytics

Agile influence: Using an agile approach we were able to quickly deliver within the challenging timeframe.

teradata.

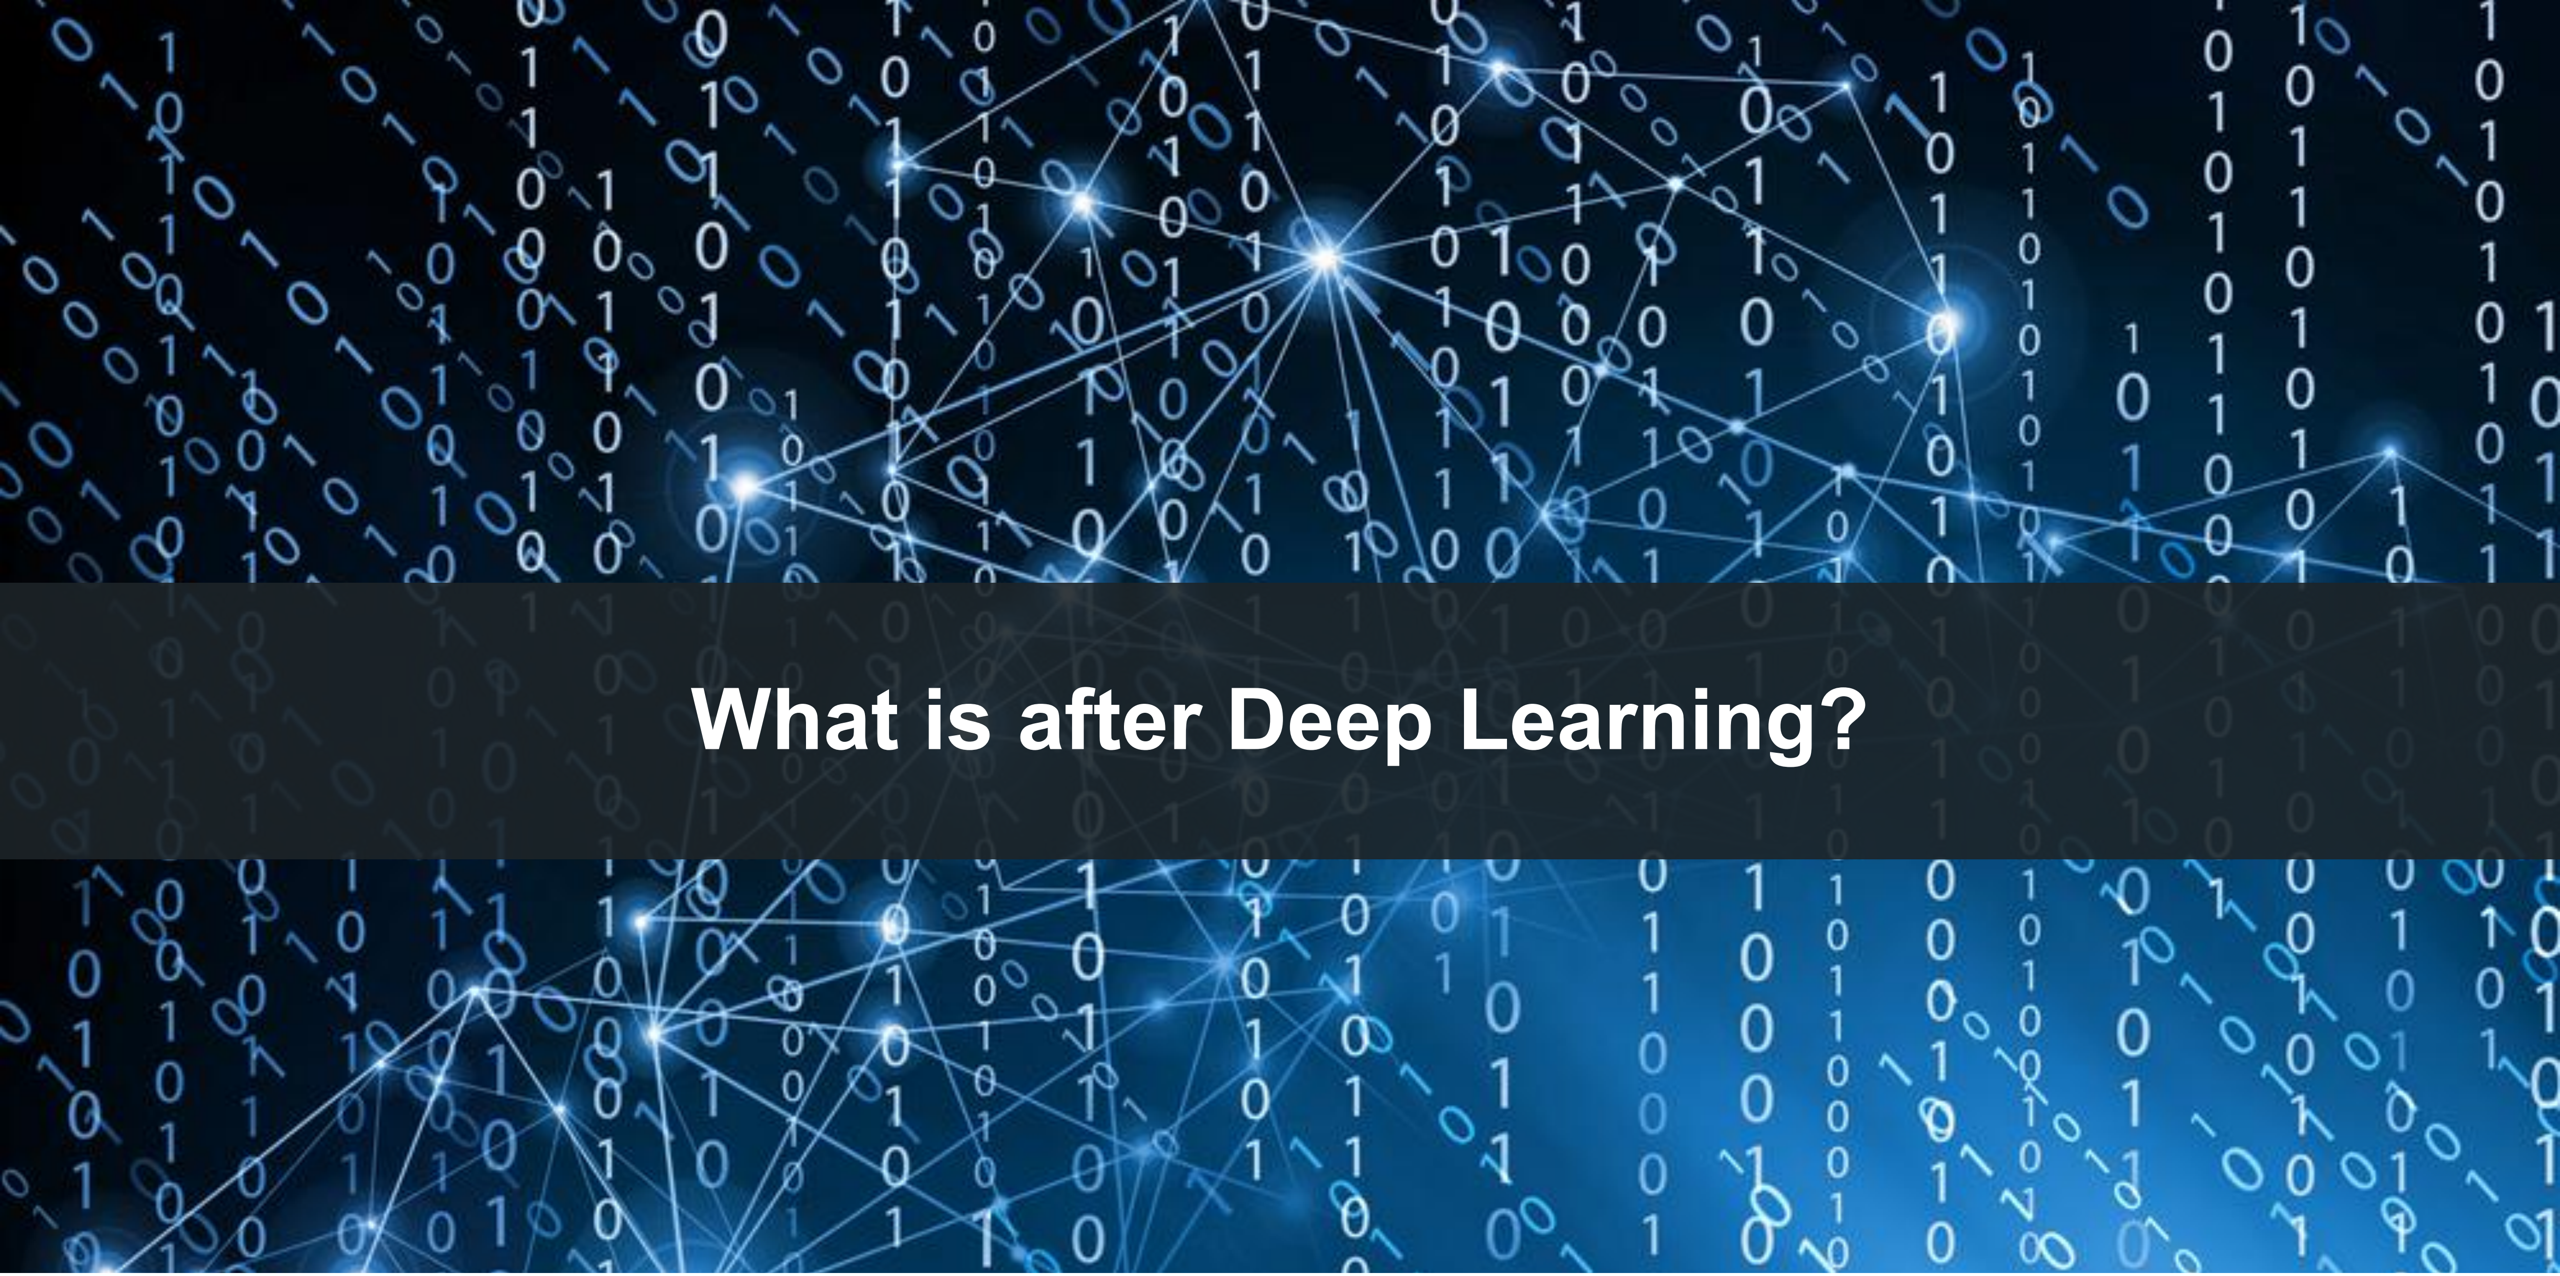

What is after Deep Learning?

teradata.

- I pounded a nail on the wall.
- I pounded a nail on the floor.

Intelligence is the ability to model the world and act on it!

Deep neural networks are easily fooled:

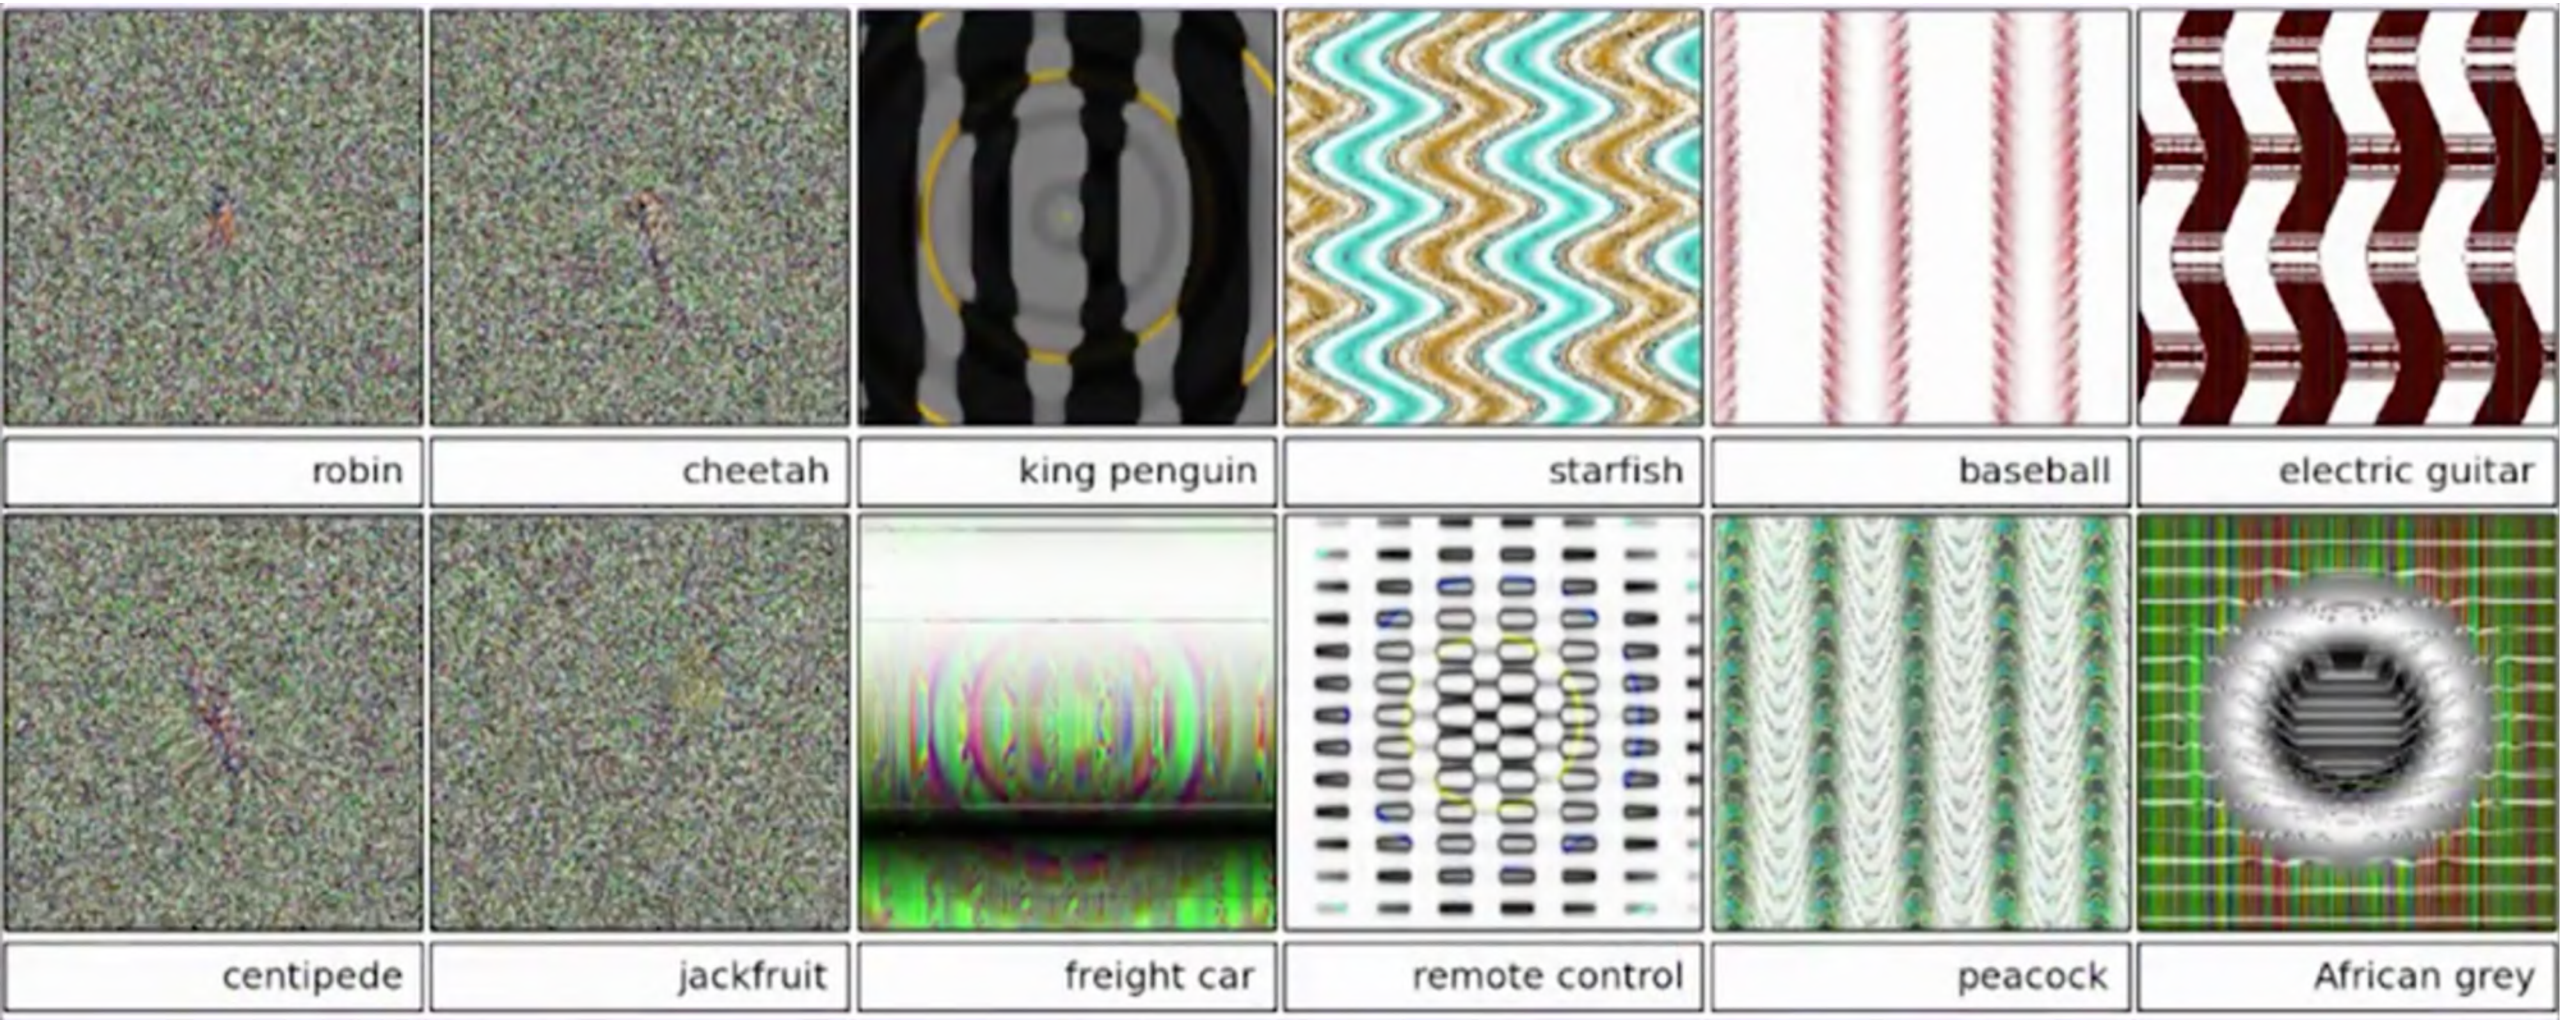

teradata.

teradata.

Trained application-specific network with limited ability for imagination and reasoning

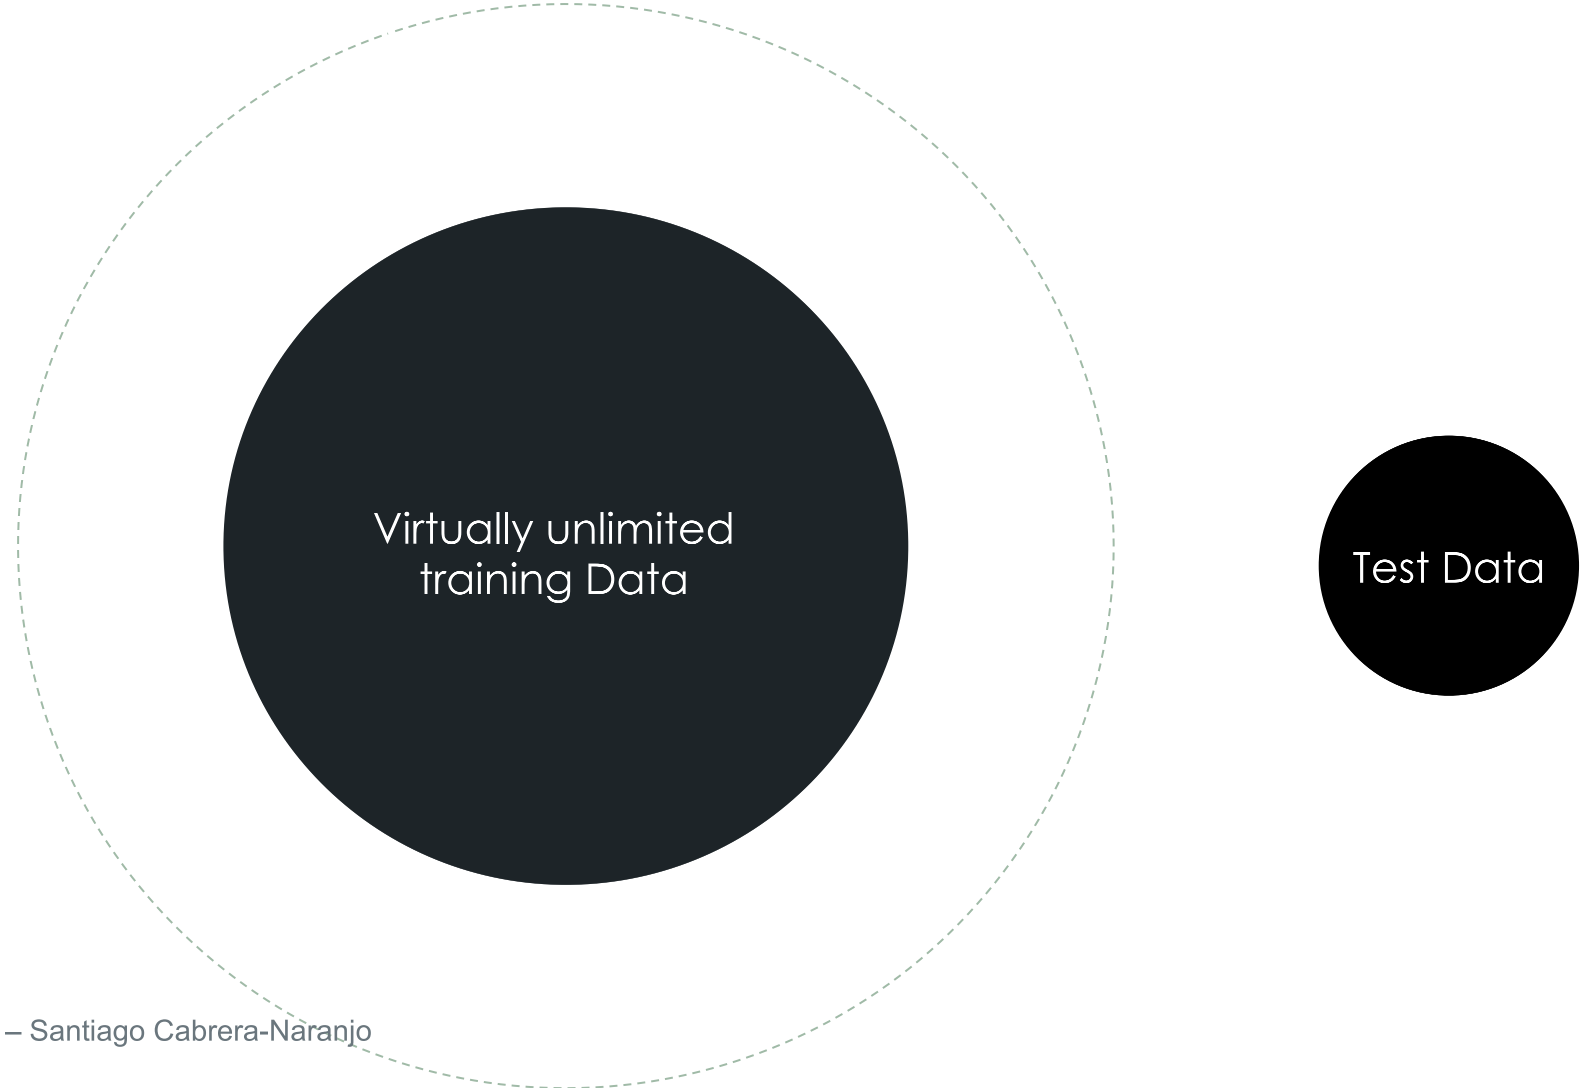

General AI

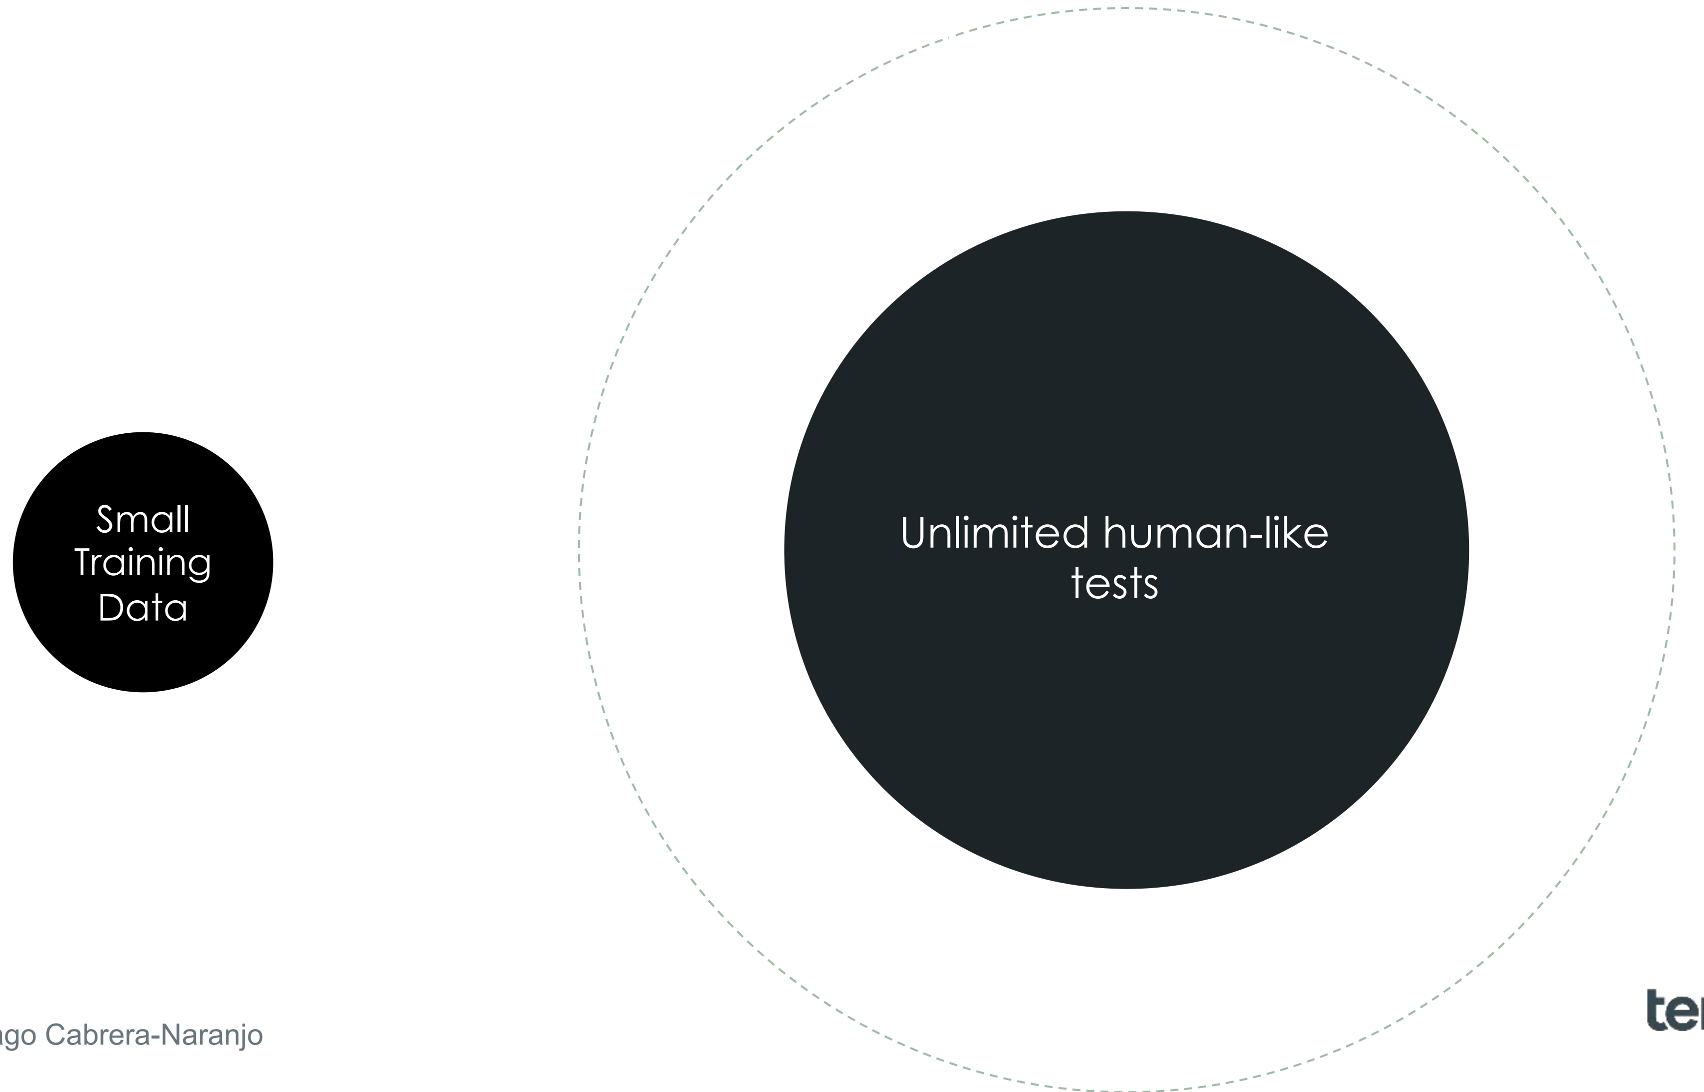

The common sense problem:

The ability to explain cause and effect presumptions about the type and essence of ordinary situations is an important aspect of explainable AI.