Enterprise-wide AI-enabled Digital Transformation

Mehdi Maasoumy, PhD

Lead Data Scientist

April 16th 2019 | ISPD19

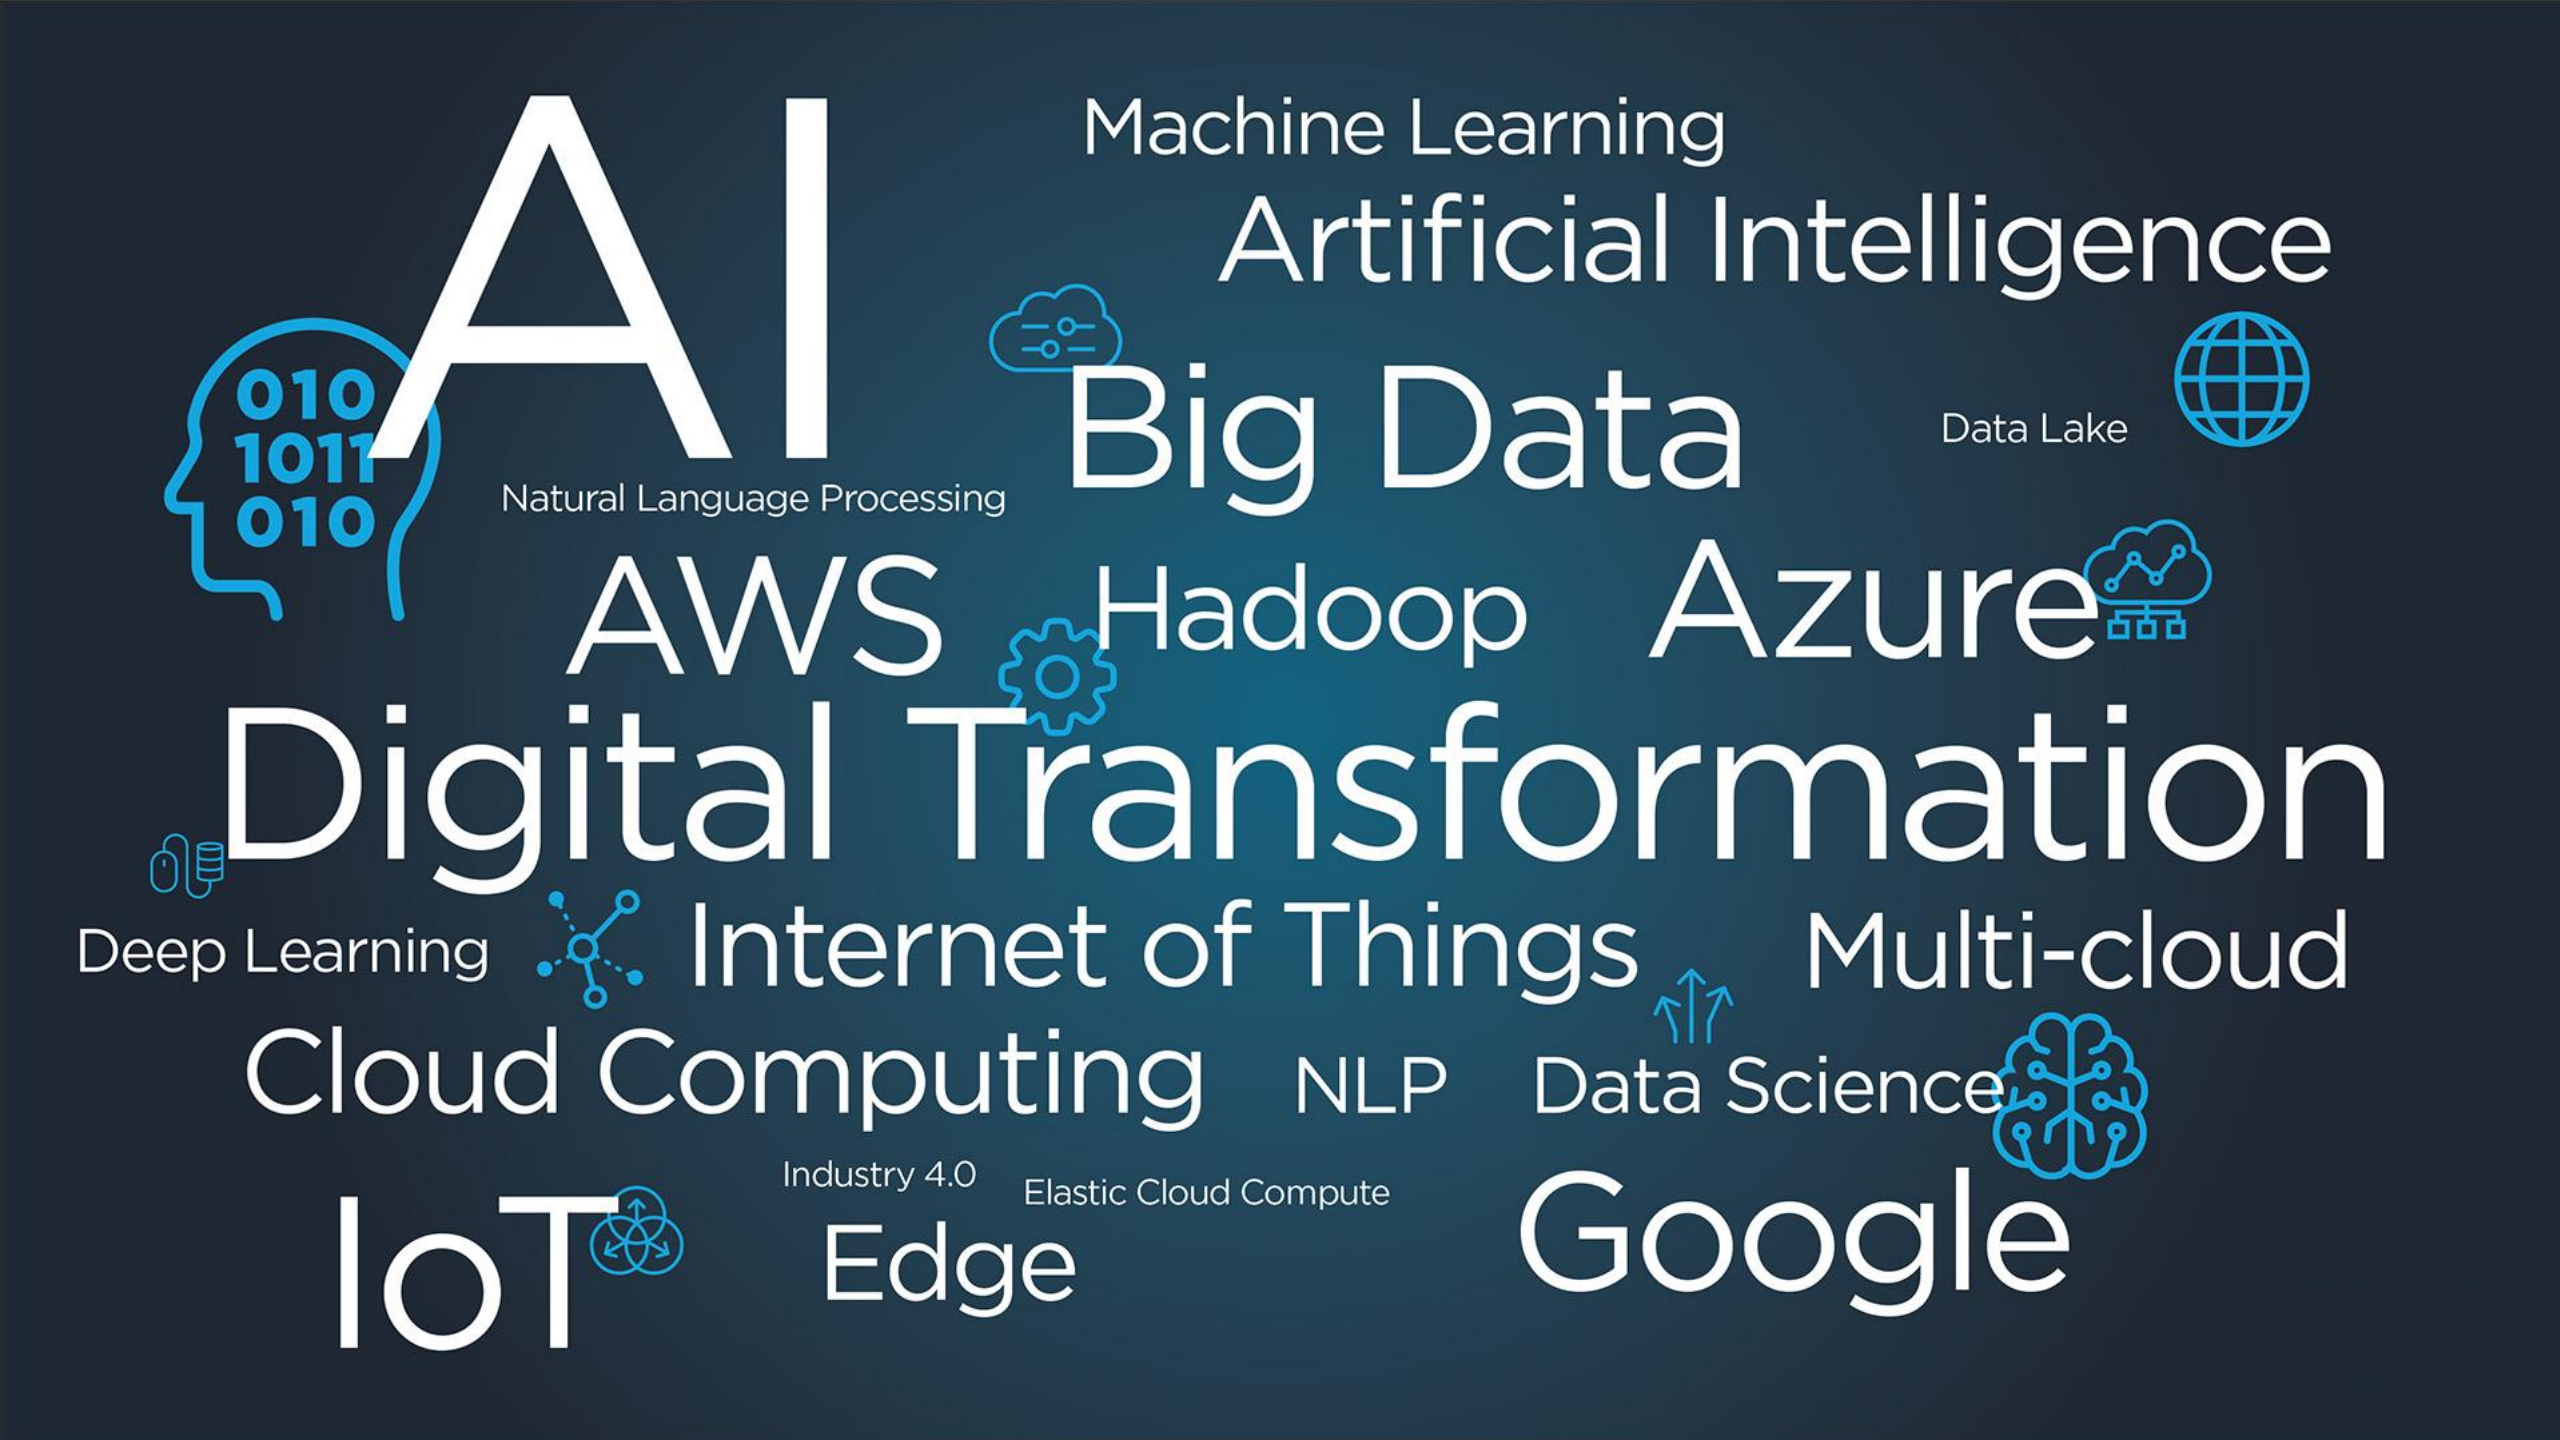
AI
Machine Learning
Artificial Intelligence
010
1011
010
Big Data
Data Lake
Natural Language Processing
AWS
Hadoop
Azure
Digital Transformation
Deep Learning
Internet of Things
Multi-cloud
Cloud Computing
NLP
Data Science
Industry 4.0
Elastic Cloud Compute
IoT
Edge
Google

AI

Future Factory

Production Schedule

Autonomous Vehicles

Route Optimization

Inventory Optimization

Supply Chain Optimization

Demand Forecasting

Predictive Maintenance

AI Augmented Intelligence

Precision Health

Order to Promise

Yield Optimization

Anti-Money Laundering

Banking

Hospitality

Automotive

Smart Cities

Financial Services

Utilities

Healthcare

Aerospace

Retail

Travel & Transportation

Oil & Gas

Government

Defense

Entertainment

Telecommunications

Predictive Maintenance

Digital Oil Field

Smart Grid

Smart Cities

Health Care Analytics

Pharmaceuticals

CRM AI

Aerospace

Supply Chain Optimization

Next-Generation Retail

Financial Services

In enterprise ML systems, the algorithm is <5 % of the code.

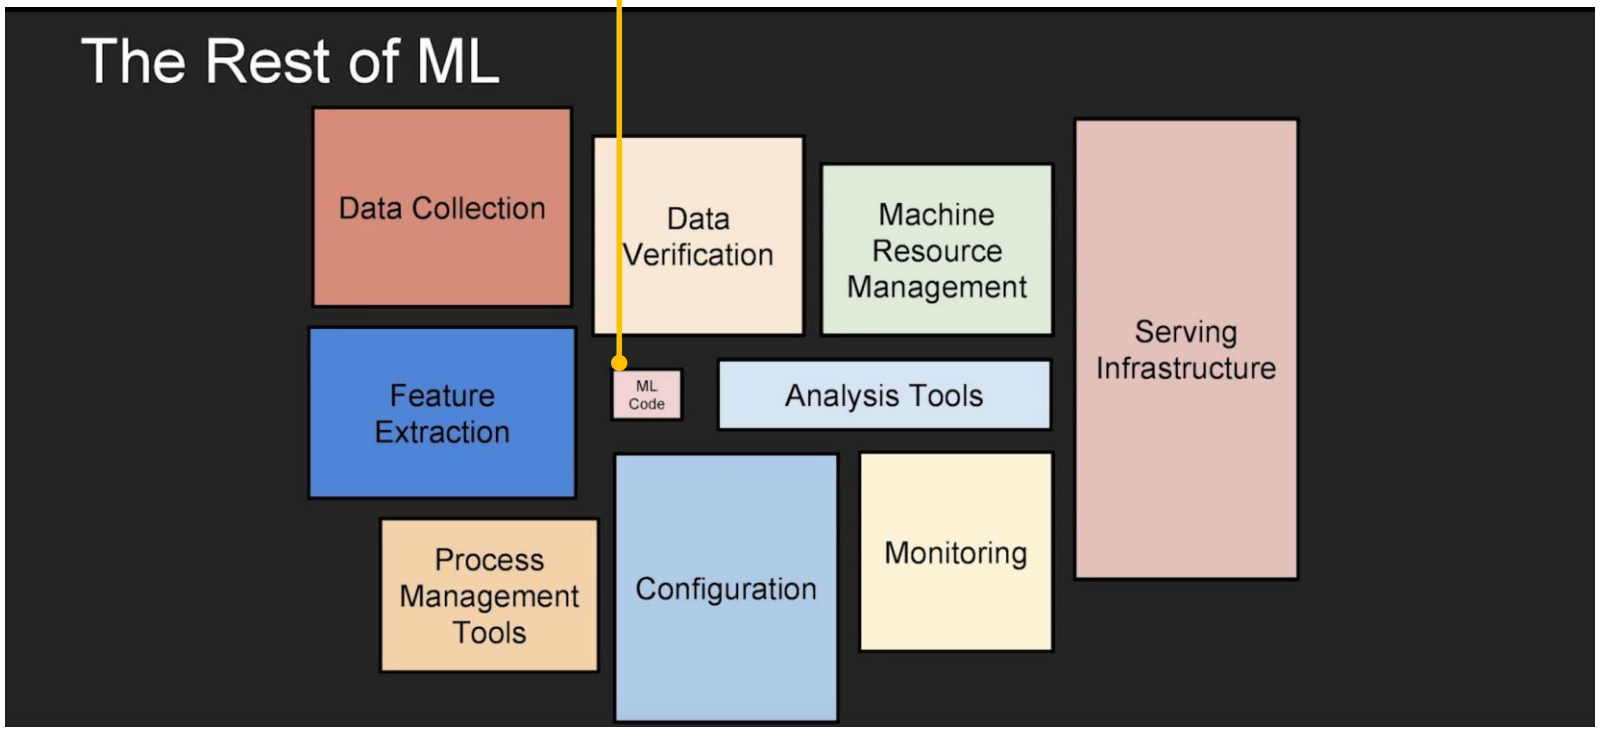

-- D. Sculley, et. al., Hidden Technical Debt in Machine Learning Systems, NIPS 2015[1]

Only a small fraction of real-world ML systems is composed of the ML code, as shown by the small box above in the middle. The required surrounding infrastructure is vast and complex.

1. *http://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf*

C3.ai

Simplified View of “AI Apps”

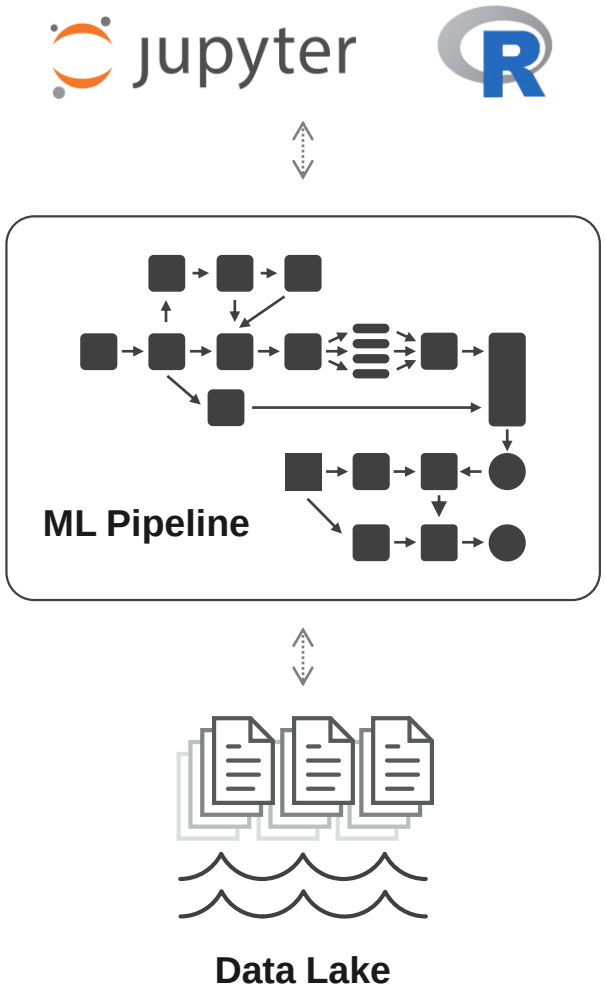

Simplified “AI Apps” Don’t Scale Across Enterprise

- Consistent View of Business
- Reusable services across enterprise
- Real-time, closed loop machine learning
- Common framework for developers and data scientists
- Ease of replacing technical components

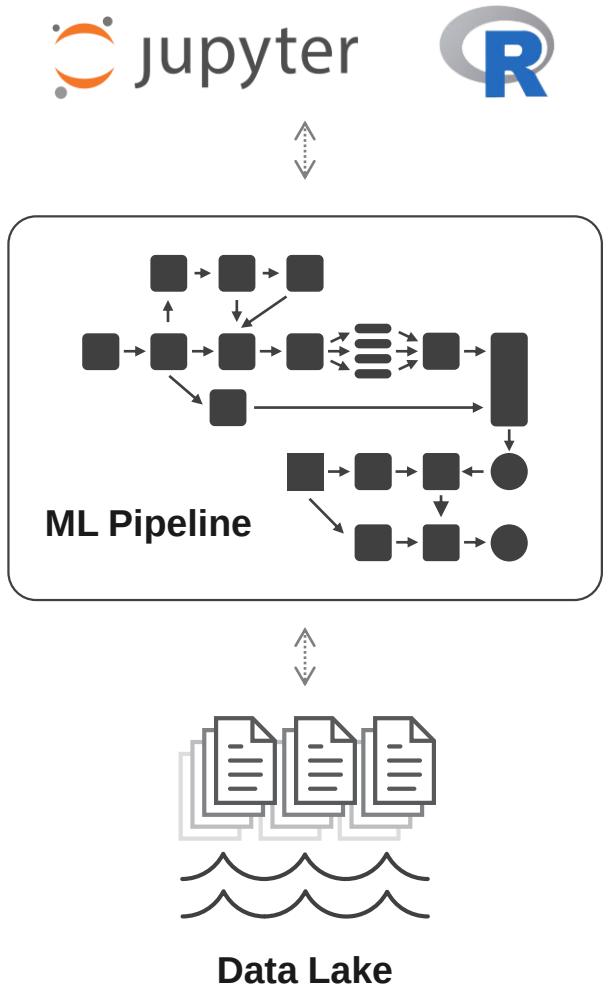

- Future proof of investments as underlying technology evolves
- Enterprise grade security
- End to end model management
- Integration from multiple data sources (persisted or virtual)
- Models replicated for each use case

C3.ai

AI Software Stack

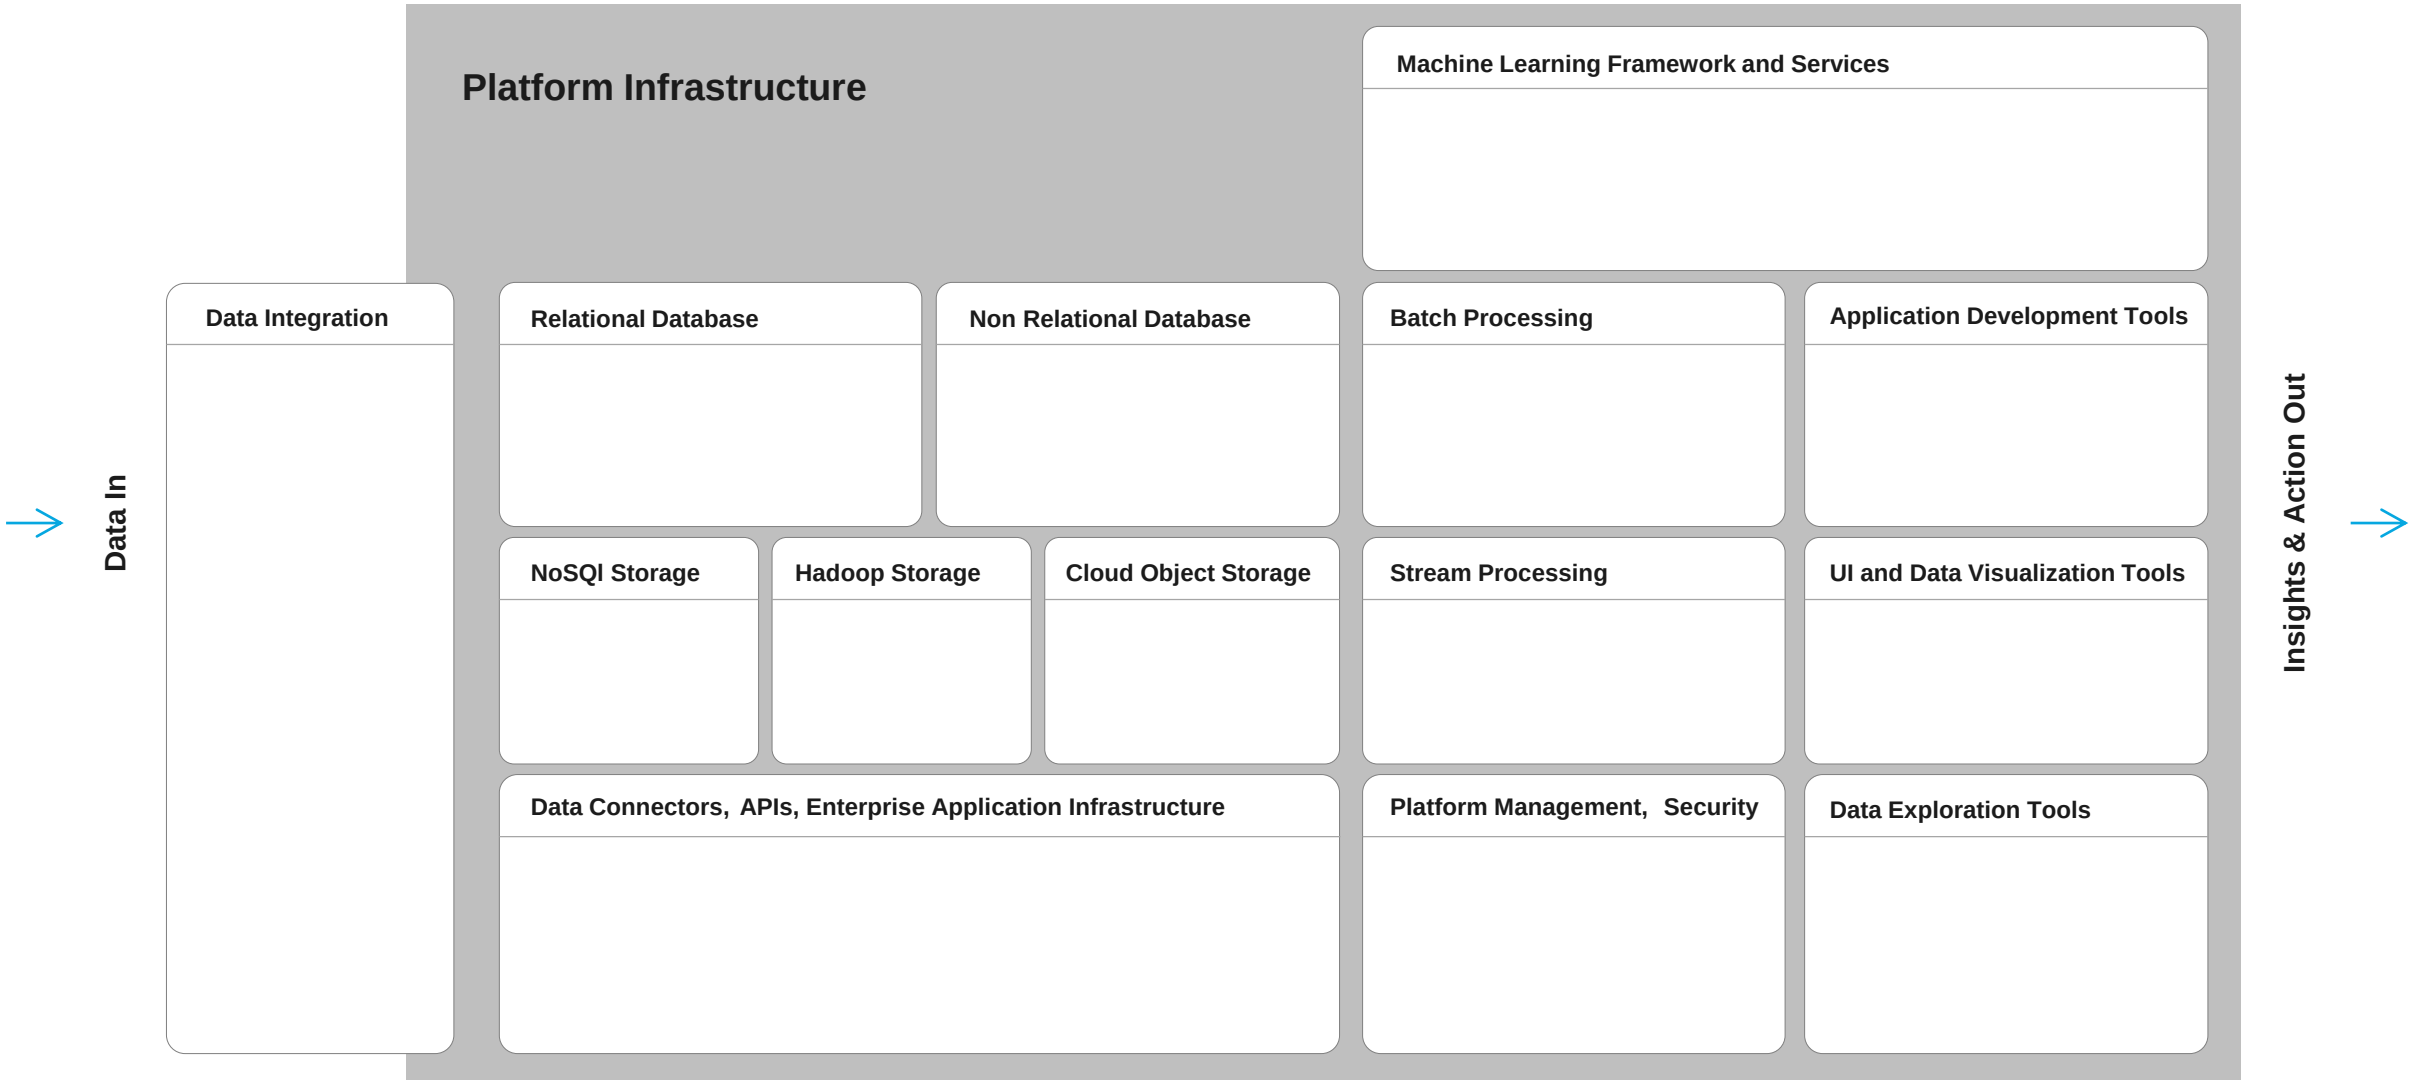

AWS IoT Events

C3.ai

C3.ai

AI Software Stack

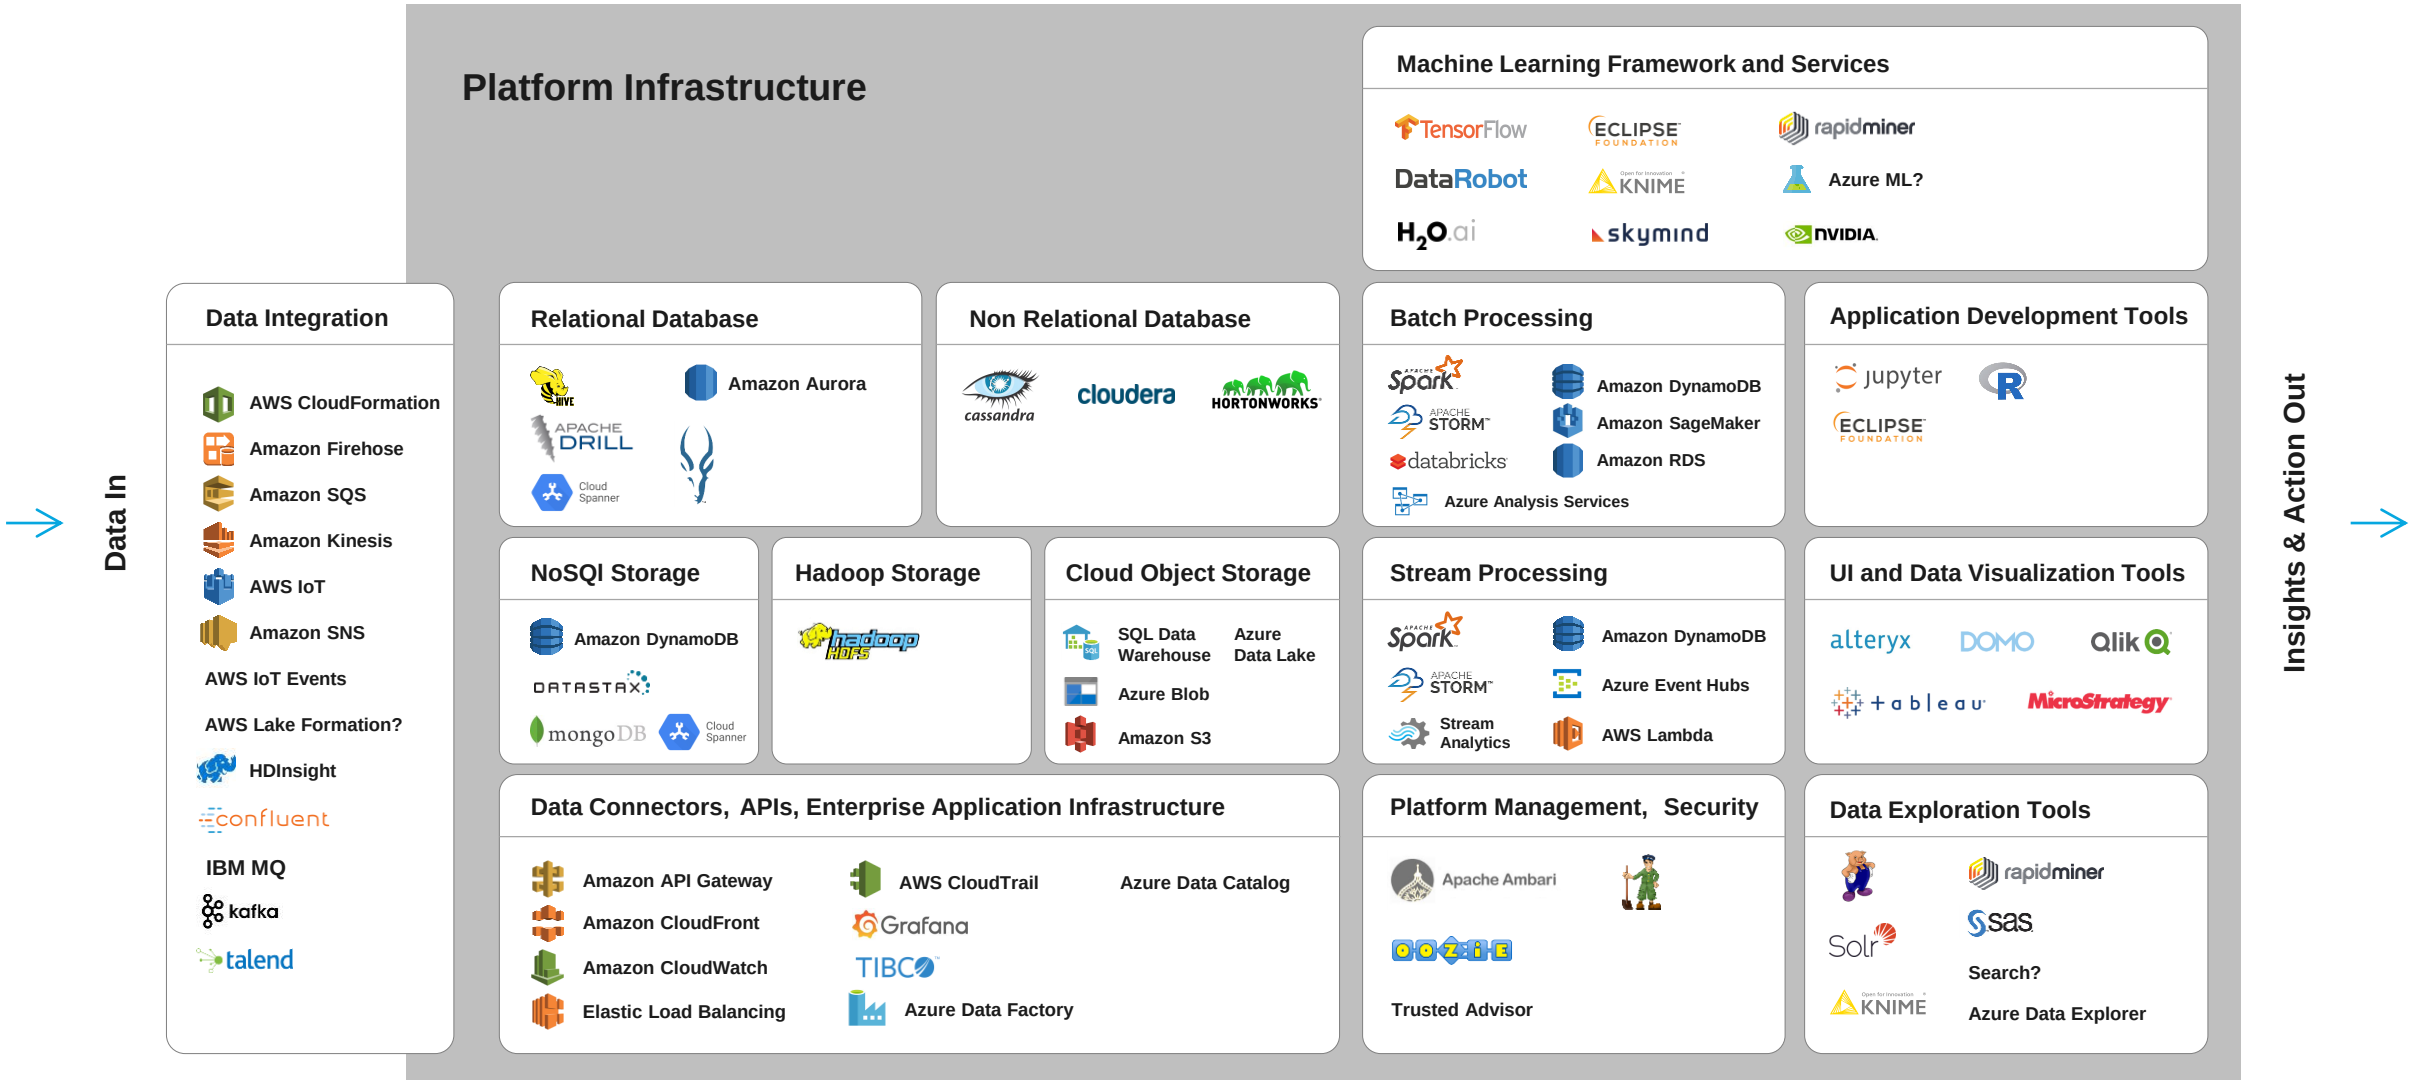

C3.ai

AI Software Stack

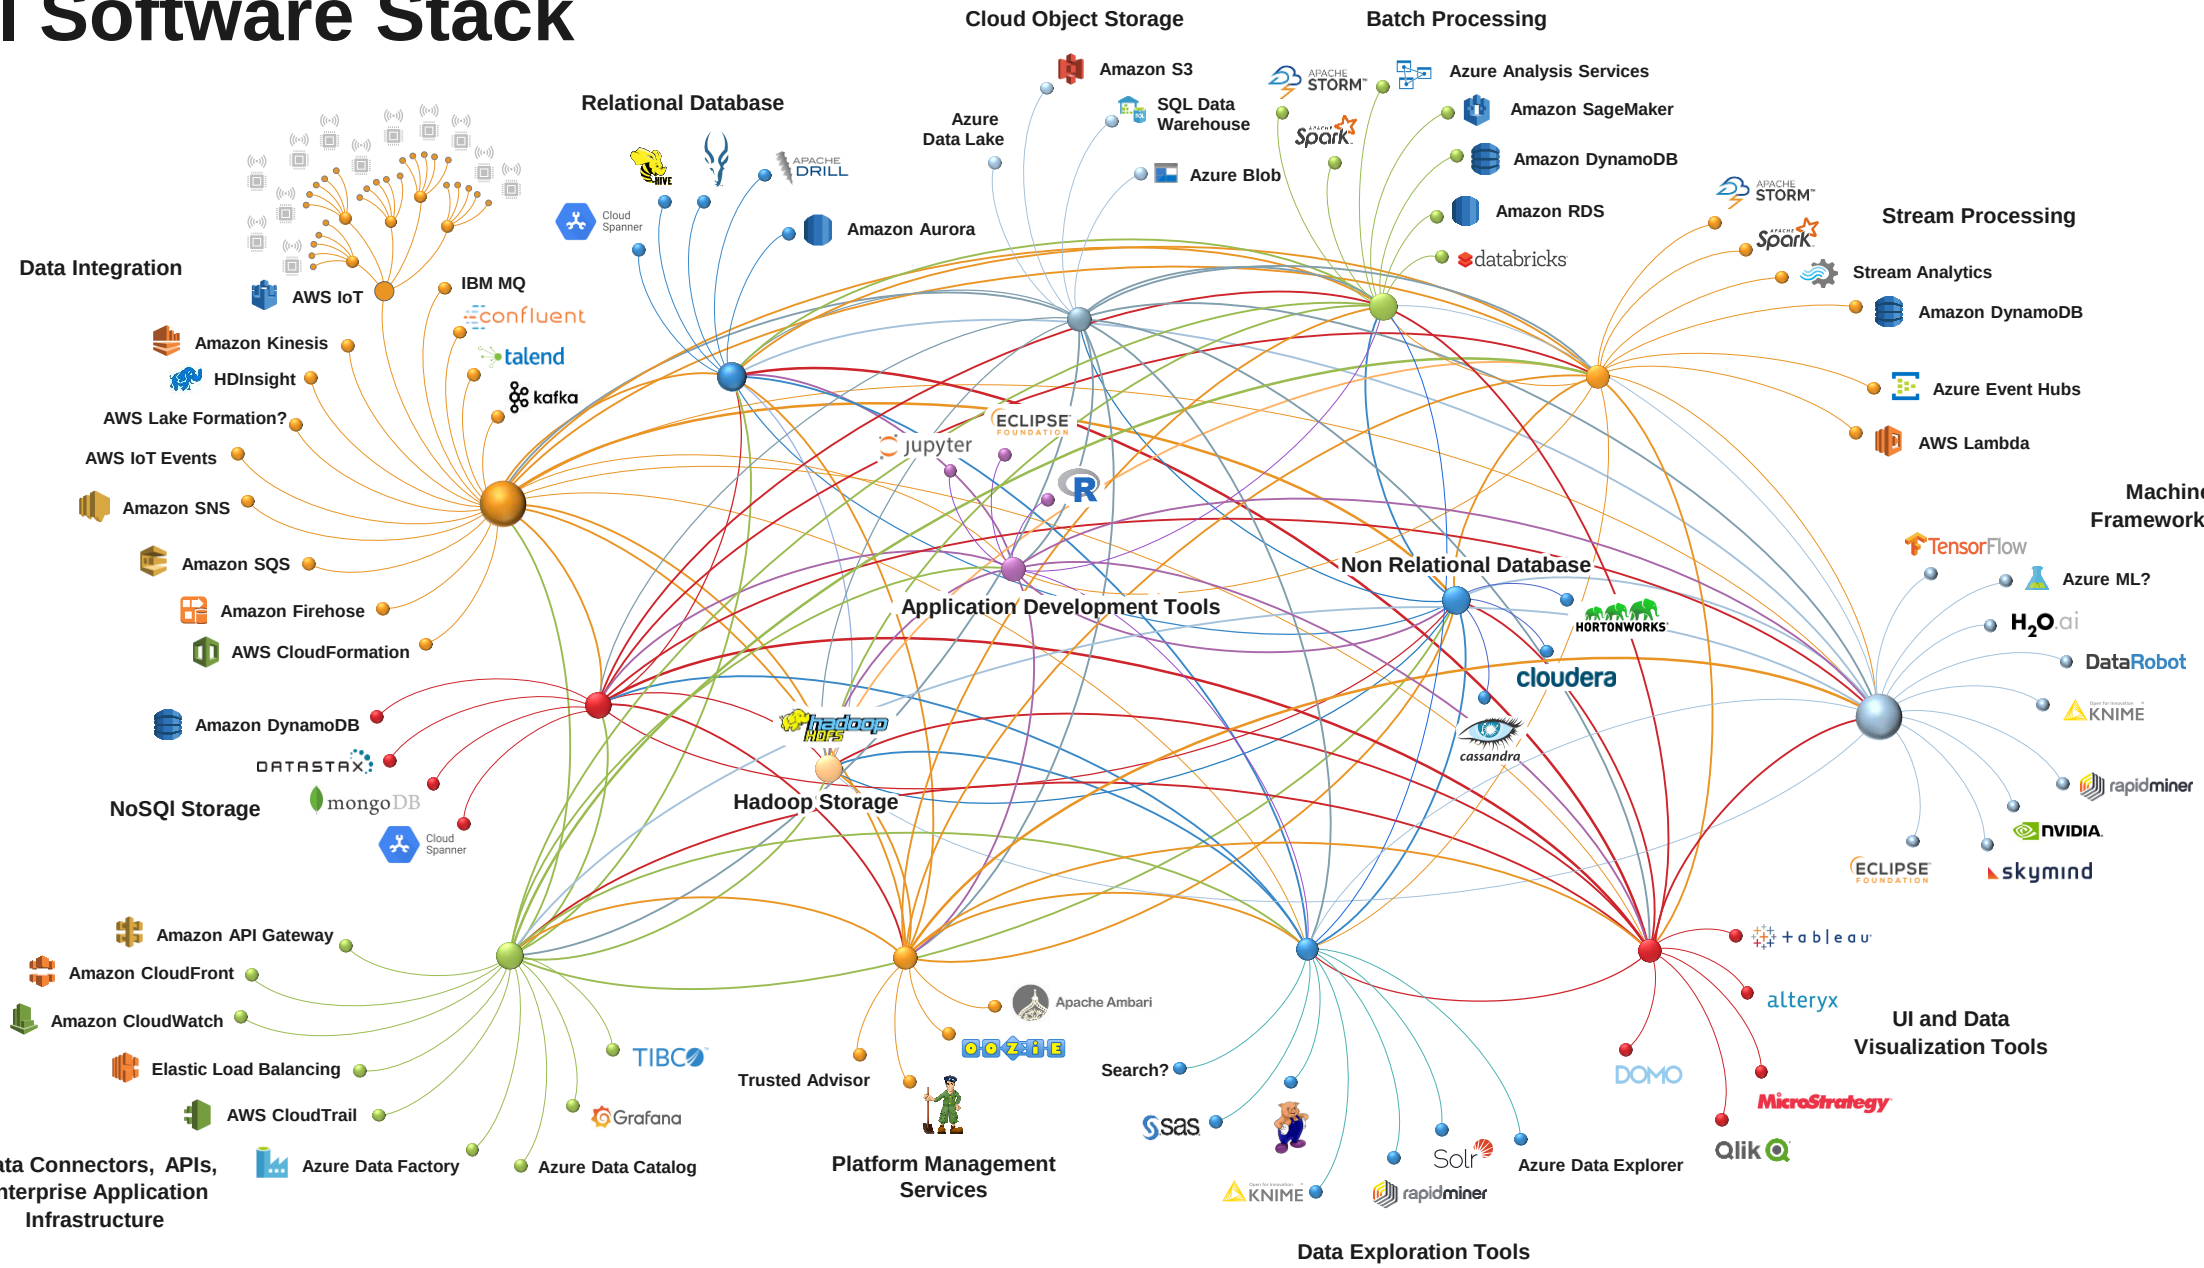

C3 AI Suite for Digital Transformation

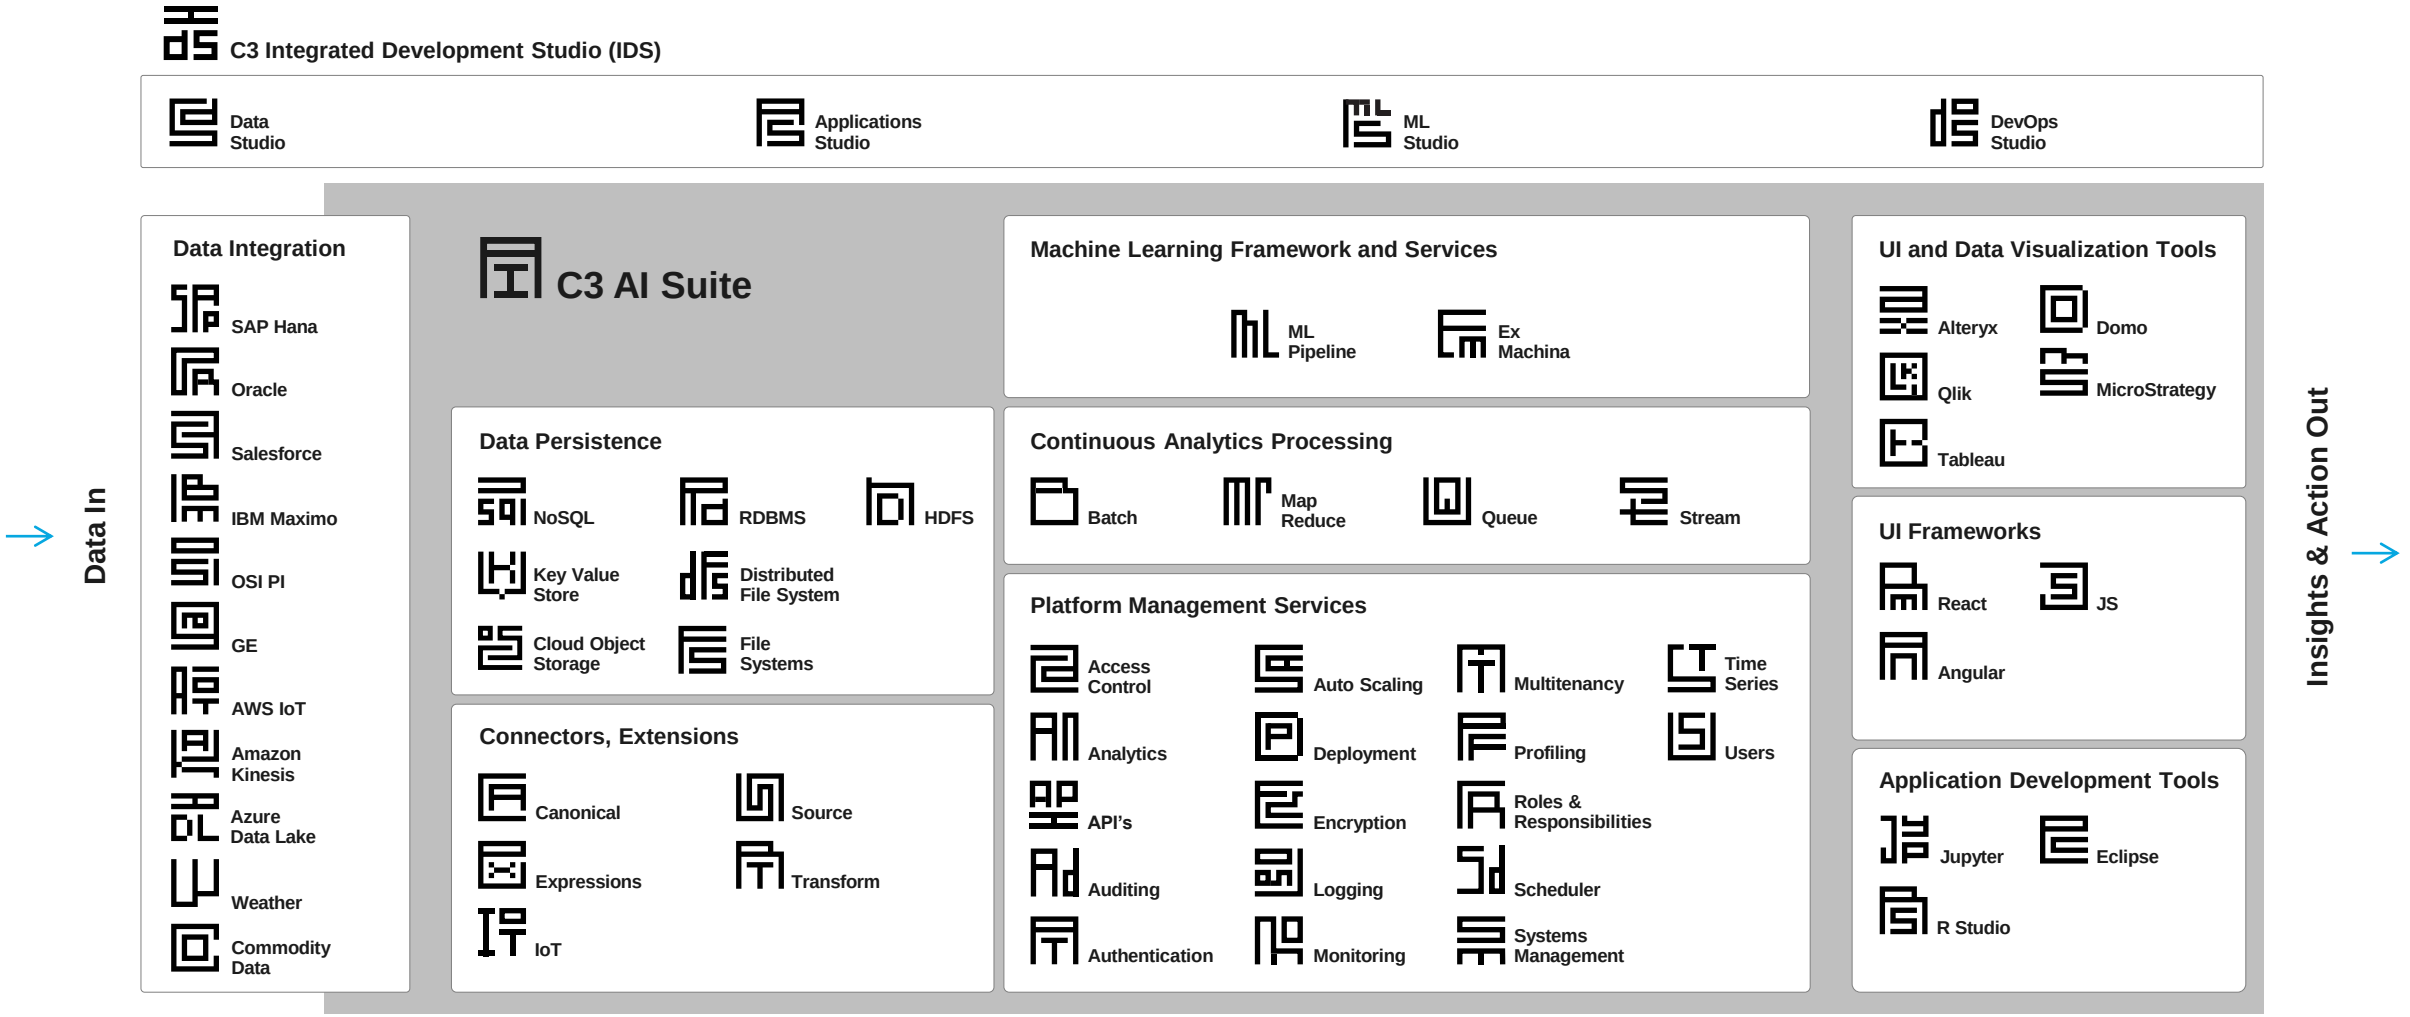

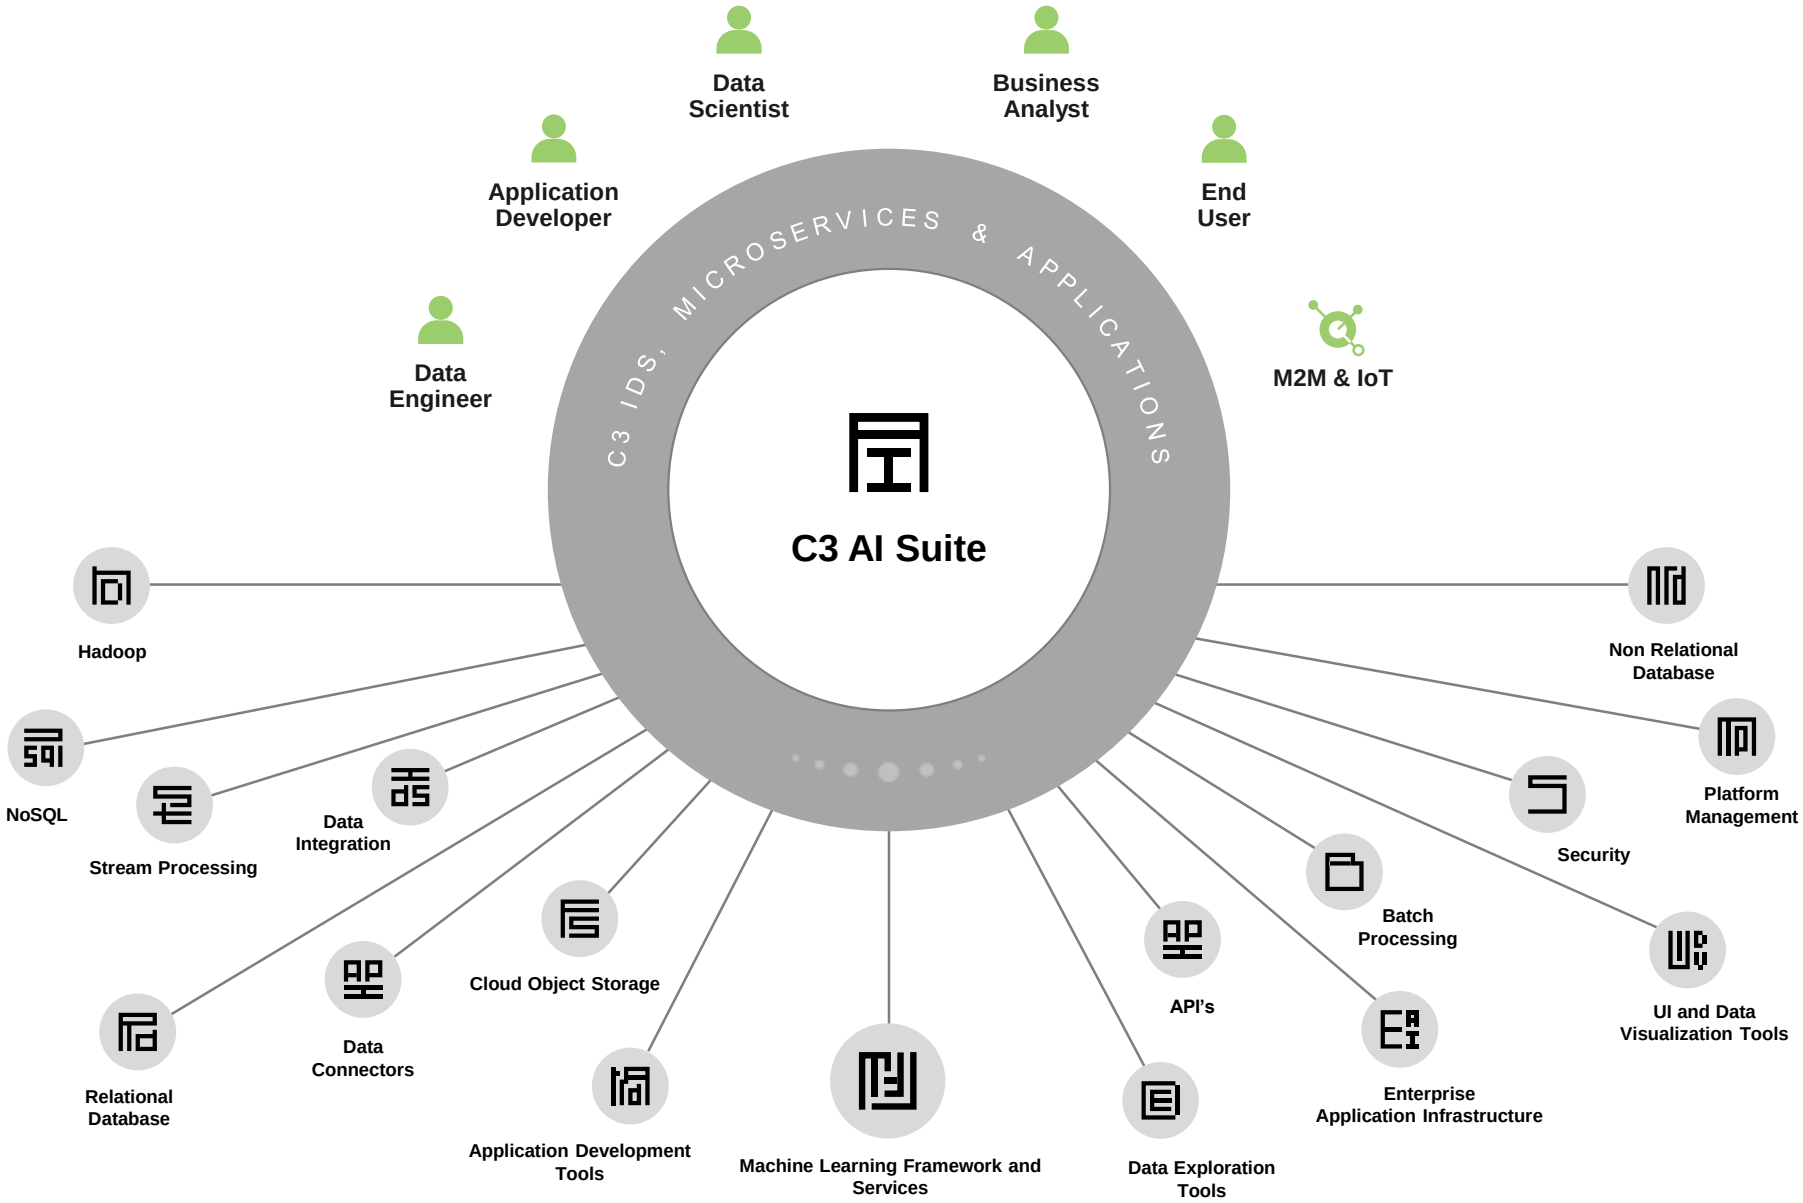

Enterprise-scale machine learning reaches beyond model training, includes data integration and full model life cycle management

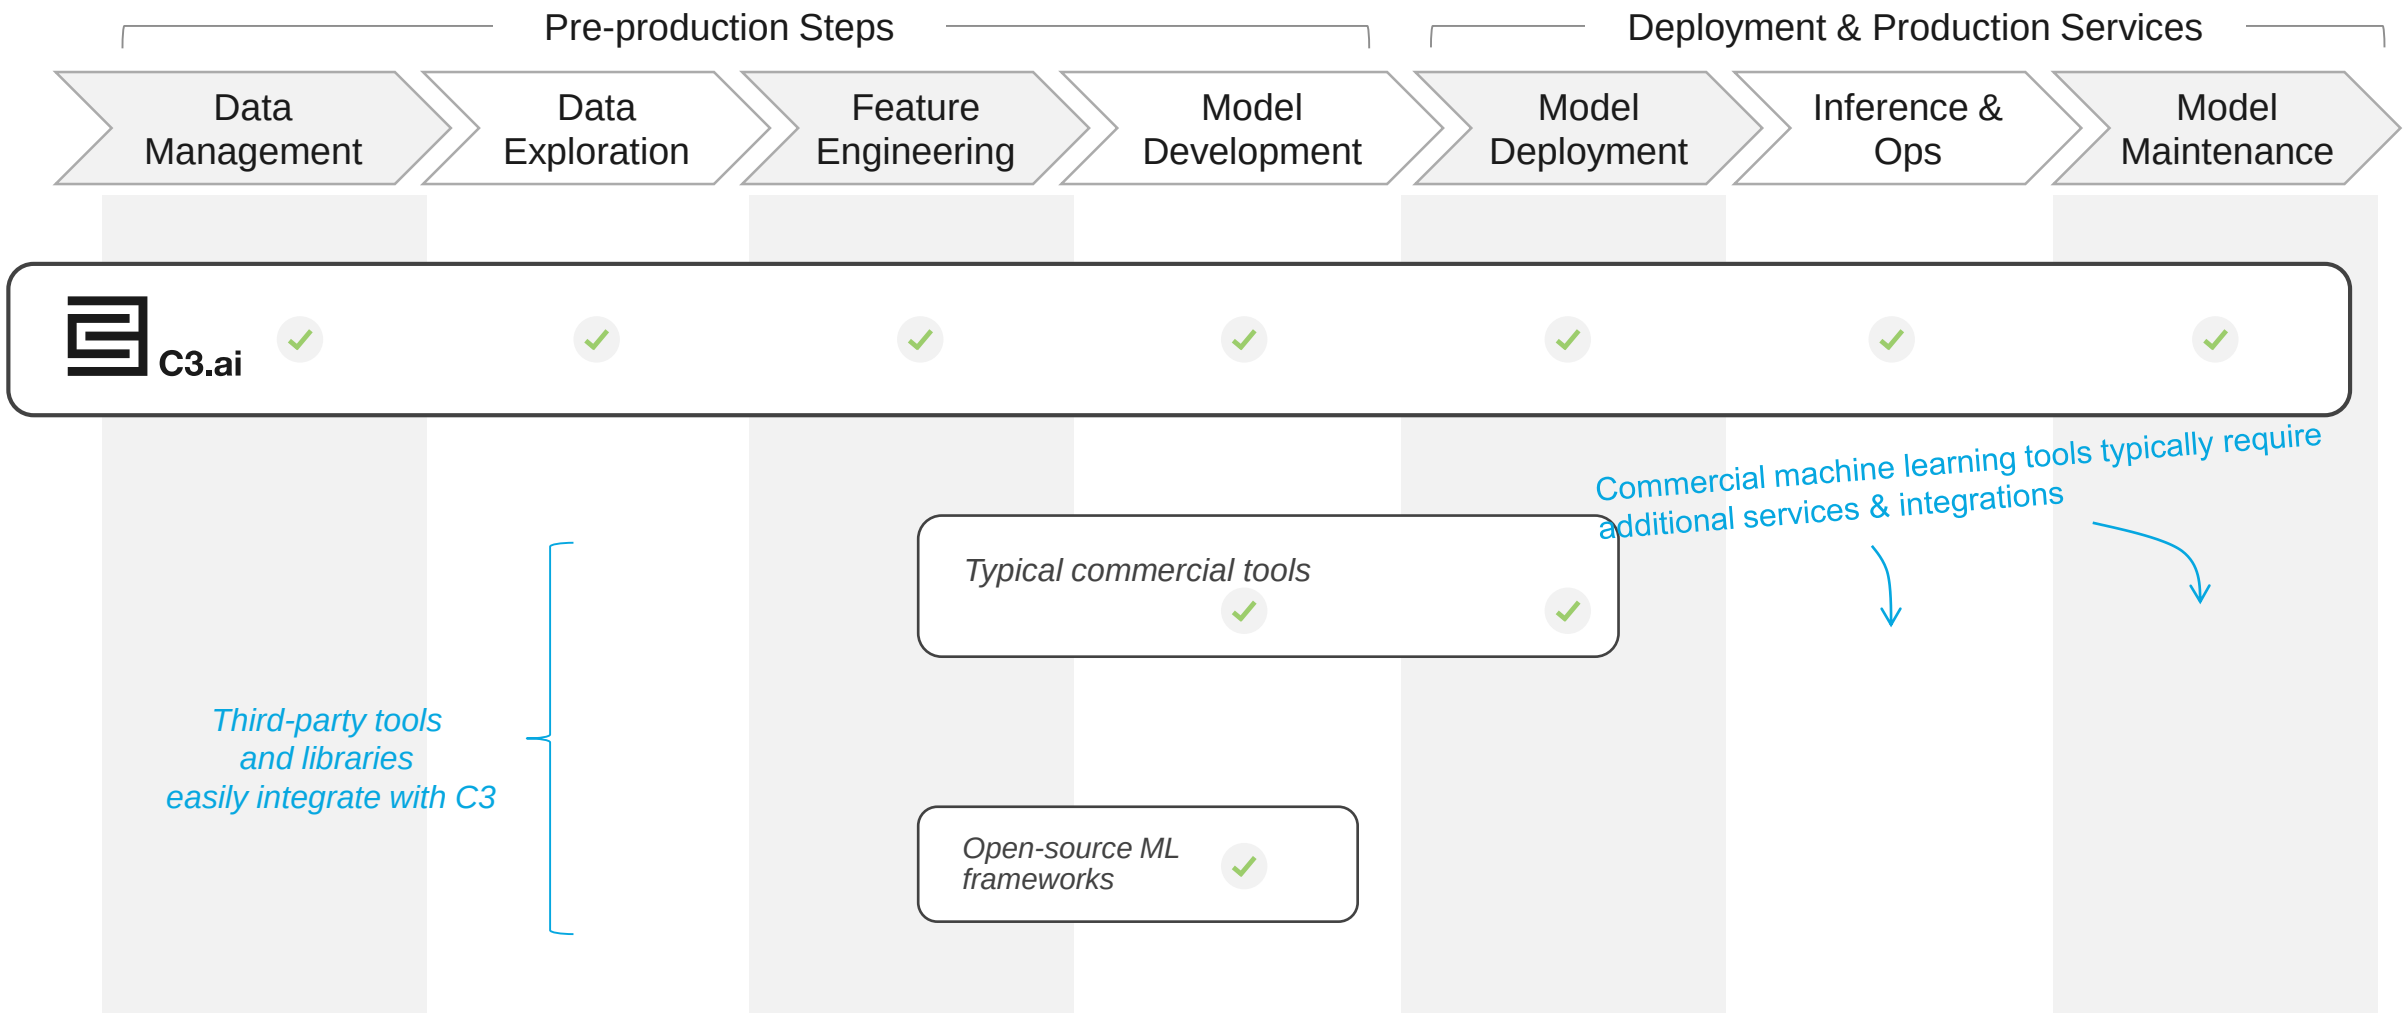

C3.ai

A typical data science process for enterprise-scale deployments traverses many functional areas, all available and configurable in C3

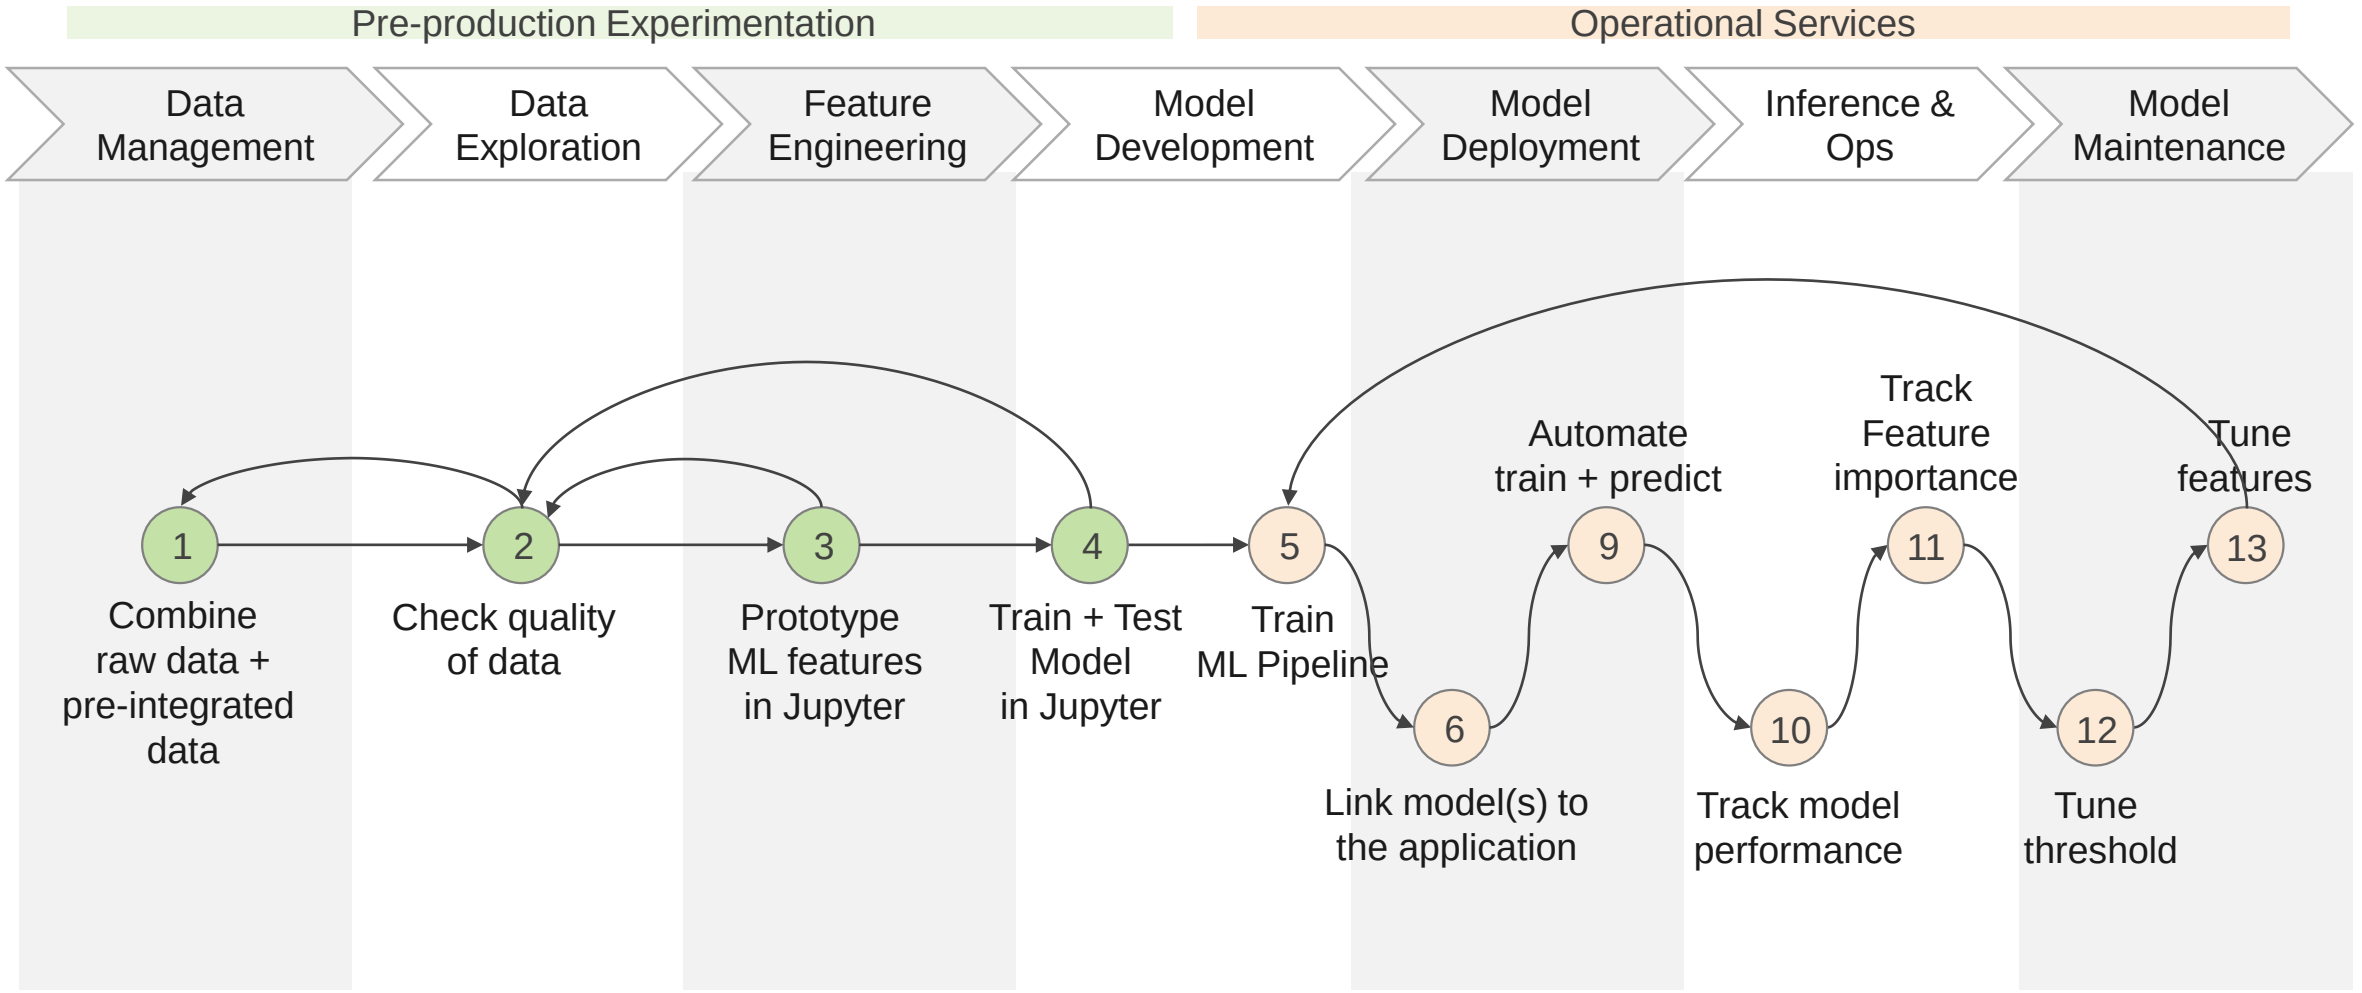

C3.ai

Use Cases

GLOBAL 500
FINANCIAL
SERVICES COMPANY

CASE STUDY

Anti-Money Laundering

C3.ai

GLOBAL 500
FINANCIAL
SERVICES COMPANY

Multinational Investment Bank
and Financial Services Company

$38B Annual Revenue

15M Customers

61,000 Employees

$2.8T Assets Under Management

C3.ai

GLOBAL 500 FINANCIAL SERVICES COMPANY

Challenge ▸ Pinpoint patterns in transactions and event data streams that identify illegal actions and signal money laundering activities

Project Details ▸

12	6	2	250+
Weeks Project Completion	Source Systems	Years' History	SAR Labels

Process and correlate unstructured data from wires with structured transaction and account data. Apply advanced analytics and machine learning to reduce false positives associated with the transaction monitoring system

Results ▸

85% Reduction in false positives using machine learning prioritization

200% Increase in recall of SAR cases identified through transaction data alone

5,000 Time-based expressions modeling the complexity of data interaction across clients

C3.ai

GLOBAL
DISCRETE
MANUFACTURER

CASE STUDY

Optimize Inventory Levels for a $30B Global Discrete Manufacturer

C3.ai

GLOBAL DISCRETE MANUFACTURER

Discrete Manufacturer with Broad Range of Industrial Equipment

$30B Annual Revenue

180 Years in Operation

60,000 Employees

C3.ai

GLOBAL DISCRETE MANUFACTURER

Challenge ▸ Build application to optimize inventory levels of one product line with over 40k unique parts

Project Details ▸

12	42	9M	3.2M
Weeks Project Completion	Files	Rows of Data	Material Movement Events

Highly customized, made-to-order product with a complex bill of materials

Results ▸

28-52% Savings in inventory holding costs

$100-$200M Annual economic value

Scaling application to 40+ global factories across over 1000 product lines

C3.ai

Largest Petrochemical Company in Latin America

PETROCHEMICAL COMAPANY

CASE STUDY

Optimize Yield by Predicting Product Quality in Real Time for a $15B Plastics Producer

C3.ai

Largest Petrochemical Company in Latin America

32B LBS
Thermoplastic resins and petrochemicals produced per year

€15B
Annual Revenue

36
Industrial Plants across 3 continents

7,700
Employees

C3.ai

CHALLENGE

Develop machine learning models to predict the melt flow and xylene soluble content of the material being produced in the reactor, in near real-time, to enable yield optimization

PROJECT DETAILS

10
Weeks Project Completion

8
Data Sources

2B
Rows of Operational Data

38k
Historical lab tests to train and test against

RESULTS

<3%
Error in melt flow and xylene predictions relative to tests performed in lab

8.1 M
Pounds of product saved per line per year

>$30M
Increase in operating margin globally

50+%
Reduction in lab tests to determine product quality

C3.ai

QGC

Predict PCP Gas Well Failures and Failure Modes

C3.ai

Challenge ▸ Build predictive maintenance application to predict three types of failures across 3000 Progressive Cavity Pump natural gas wells

Project Details ▸

16	73	10B	7	1
Weeks	Distinct File Types	Rows of raw data	Machine learning models	Day to go from trial to daily prediction application

Results ▸

600M+ Estimated Savings

~86% ROC performance in identifying all types of well failures based on initial labels

84-87% ROC performance in identifying individual failure types based on initial labels

C3.ai

CASE STUDY

Largest Production Deployment of AI & IoT Applications

C3.ai

Global Fortune 100

96GW Generation Capacity

€71B Annual Revenue

73M Customers Globally

63,000 Employees

TOP 5 Fortune’s “Change The World” List

C3.ai

Challenge ▸ Enterprise-wide digital transformation involving the largest deployment of AI and IoT applications in the world.

Project Details ▸

5+	**25+**	**50TB**	**10M**
Year Strategic Partnership	Data Sources	Unified Federated Cloud Image	Node In-memory Graph Network

Results ▸

4	Production applications deployed at scale
50M+	Sensors monitored
€200M+	Annual recurring economic value
€6.72B	Potential per year in economic benefit once fully implemented

C3.ai

MANAGED
HEALTHCARE
COMPANY

CASE STUDY

Predict Risk of Opiate Dependency for $200B Managed Healthcare Company

MANAGED HEALTHCARE COMPANY

CHALLENGE ▸ Develop AI application to predict risk of opiate dependency

PROJECT DETAILS ▸

9	6	2.5B	120M
Weeks Project Completion	Data Sources	Rows of Medical Data	Insured Members

RESULTS ▸

80% Accuracy of opioid risk identification

40x Improvement in development productivity

5x Reduction in machine learning algorithm code on the C3 IoT Platform

Opioid Dependency Risk Application

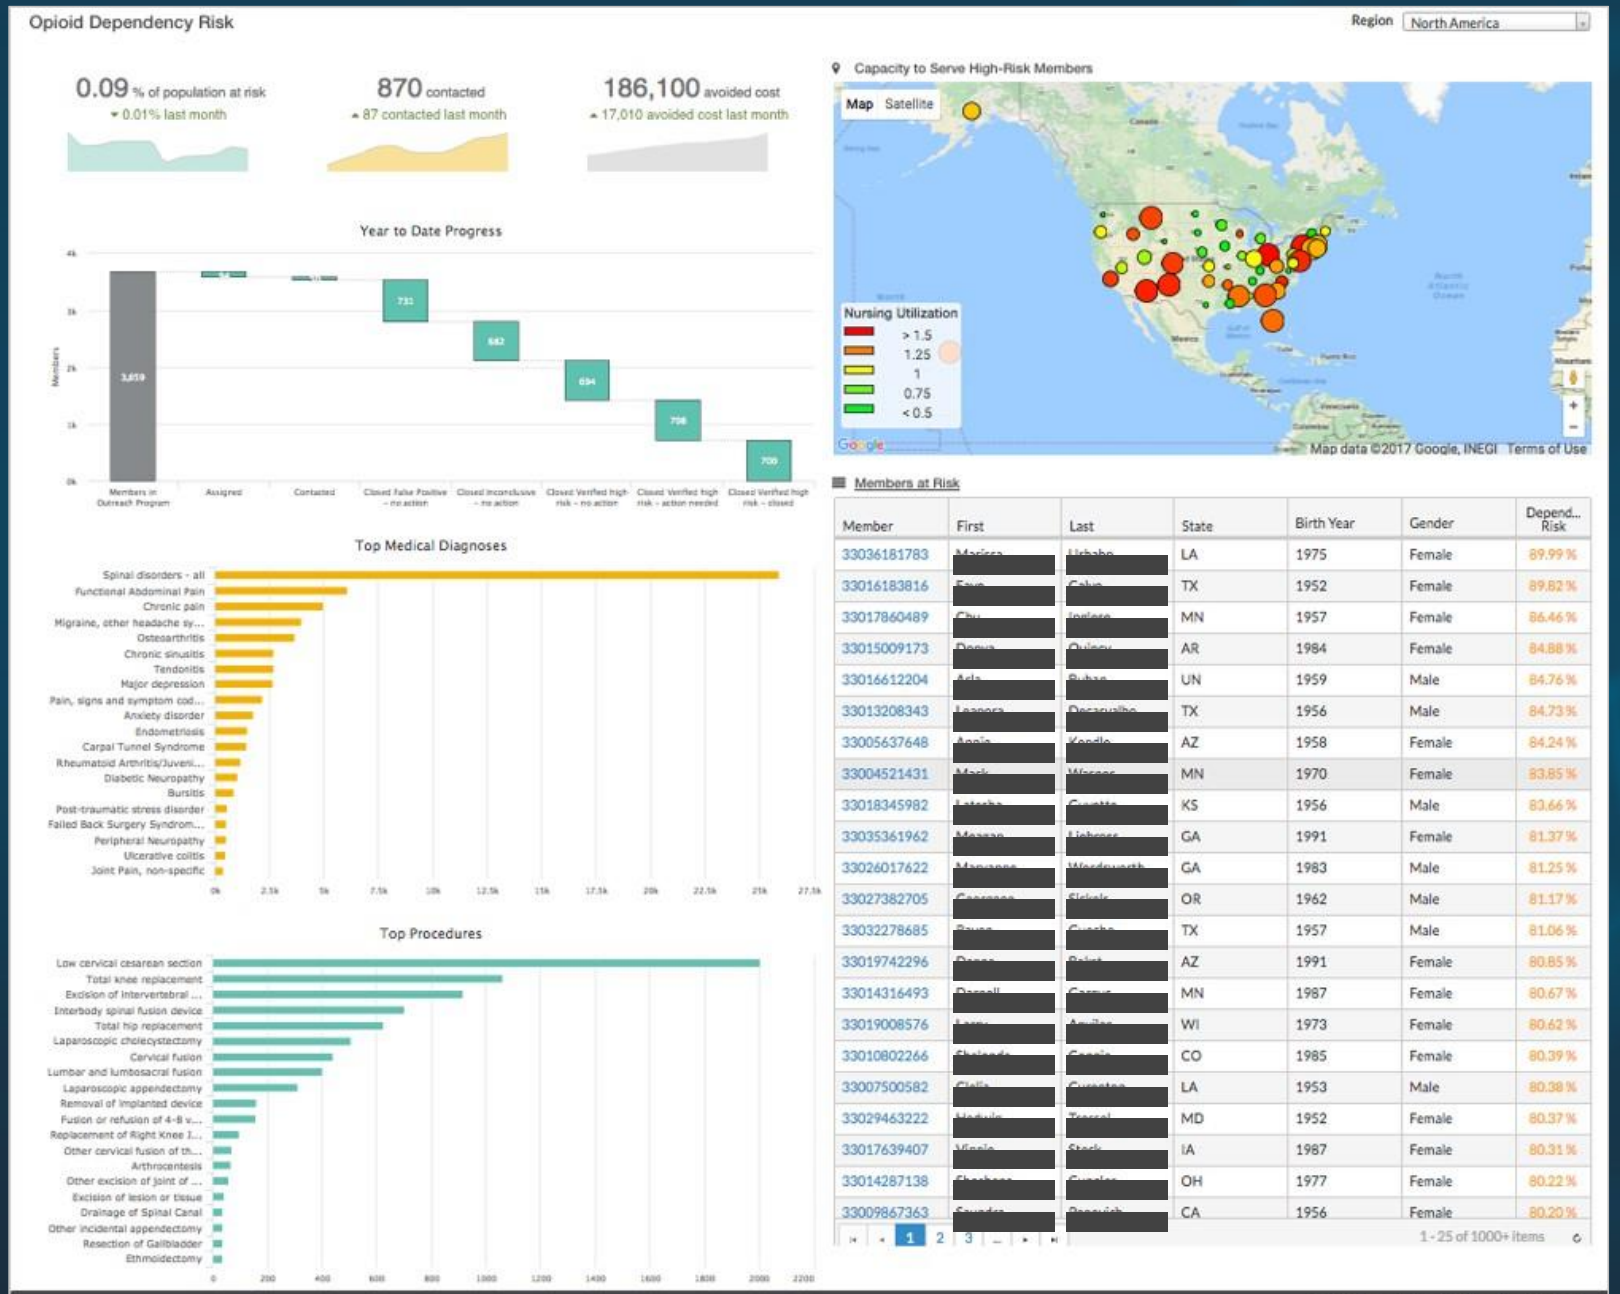

U.S. AIR FORCE

CASE STUDY

Predictive Maintenance for Aircraft Sub-systems

C3 IoT

C3 Industry Applications

	Utilities	Oil & Gas / Mining	Industrial / Manufacturing	Consumer Packaged Goods	Aerospace & Defense	Financial Services	Healthcare	Retail	Telecom	Public Sector / Cities
C3 CRM™	●	●	●	●	●	●	●	●	●	●
C3 Predictive Maintenance™	●	●	●	●	●	●	●	●	●	●
C3 Energy Management™	●	●	●	●	●	●	●	●	●	●
C3 Sensor Health™	●	●	●	●	●	●	●	●	●	●
C3 Fraud Detection™	●	●	●	●	●	●	●	●	●	●
C3 Supply Network™	●	●	●	●	●	●	●	●	●	●
C3 Inventory Optimization™	●	●	●	●	●	●	●	●	●	●

C3 AI Suite™

Sensor Network (edge) | External (social media, weather, traffic, prices) | Product Design | Engineering | Manufacturing | Supply Chain | Billing | Financials | HR | Industry Specific

DATA SOURCES ACROSS VALUE CHAIN

C3.ai

AI-Based Stochastic Supply Chain Optimization

Companies Often Hold Large Quantities of Excess Inventory

Large difference between Actual Inventory and Consumption

Example: Part A

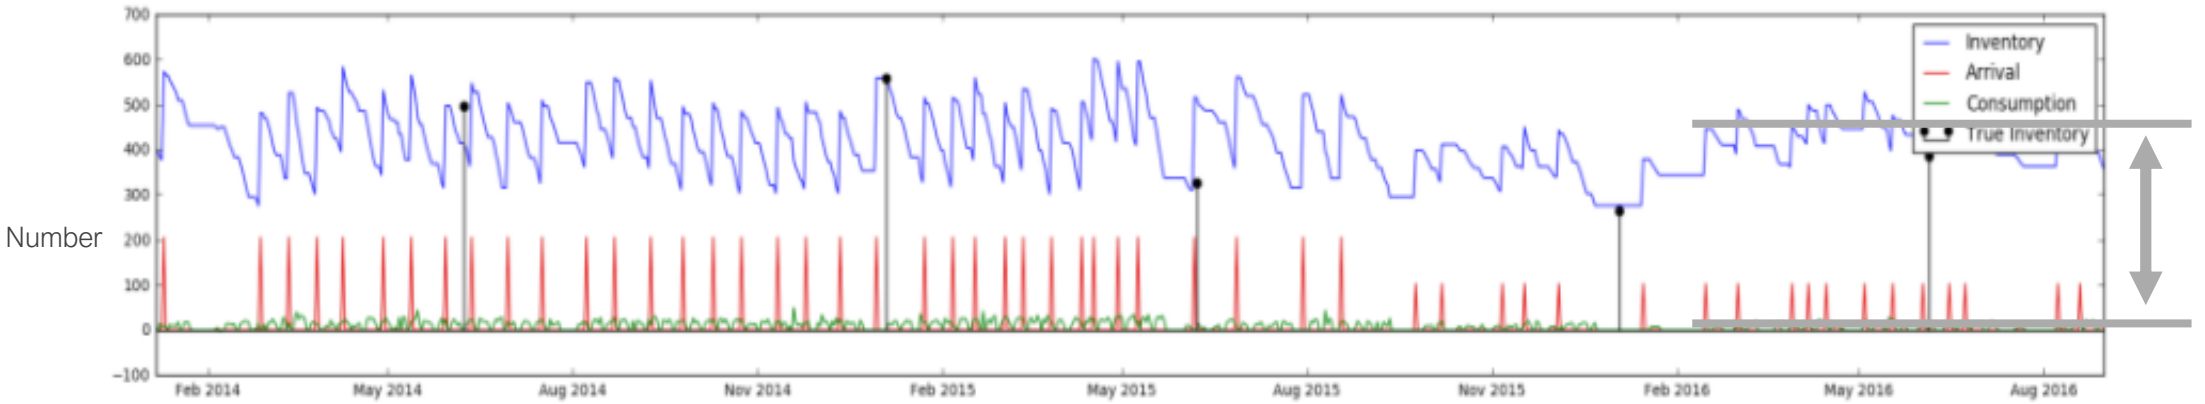

Large difference between inventory levels and part consumption levels – indicating opportunity to reduce inventory

At the Same Time, Supplier Delays Cause Cascading Downstream Issues

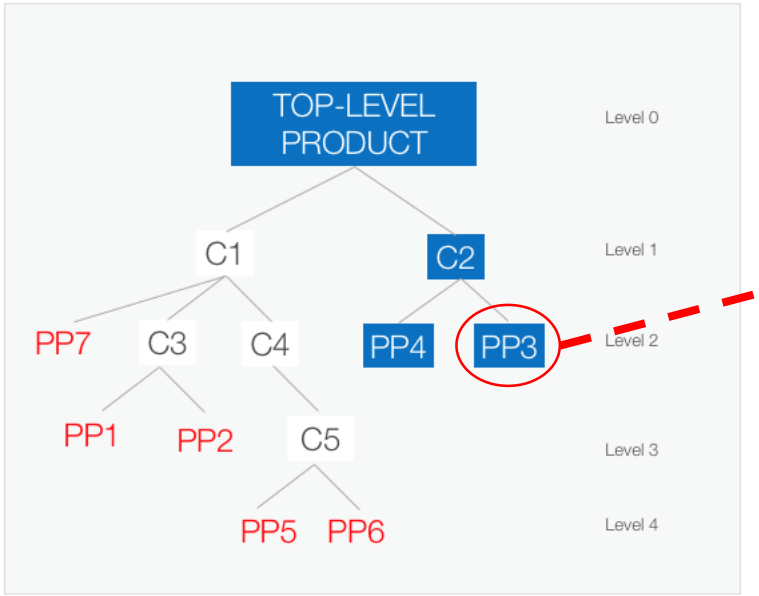

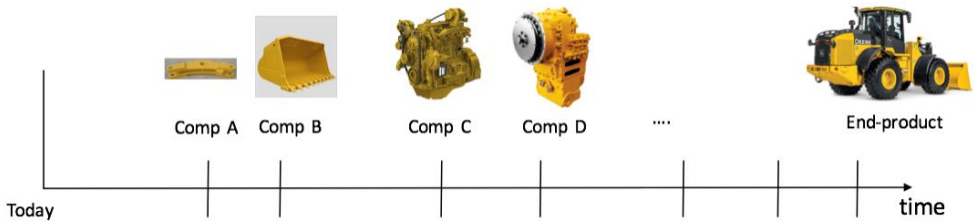

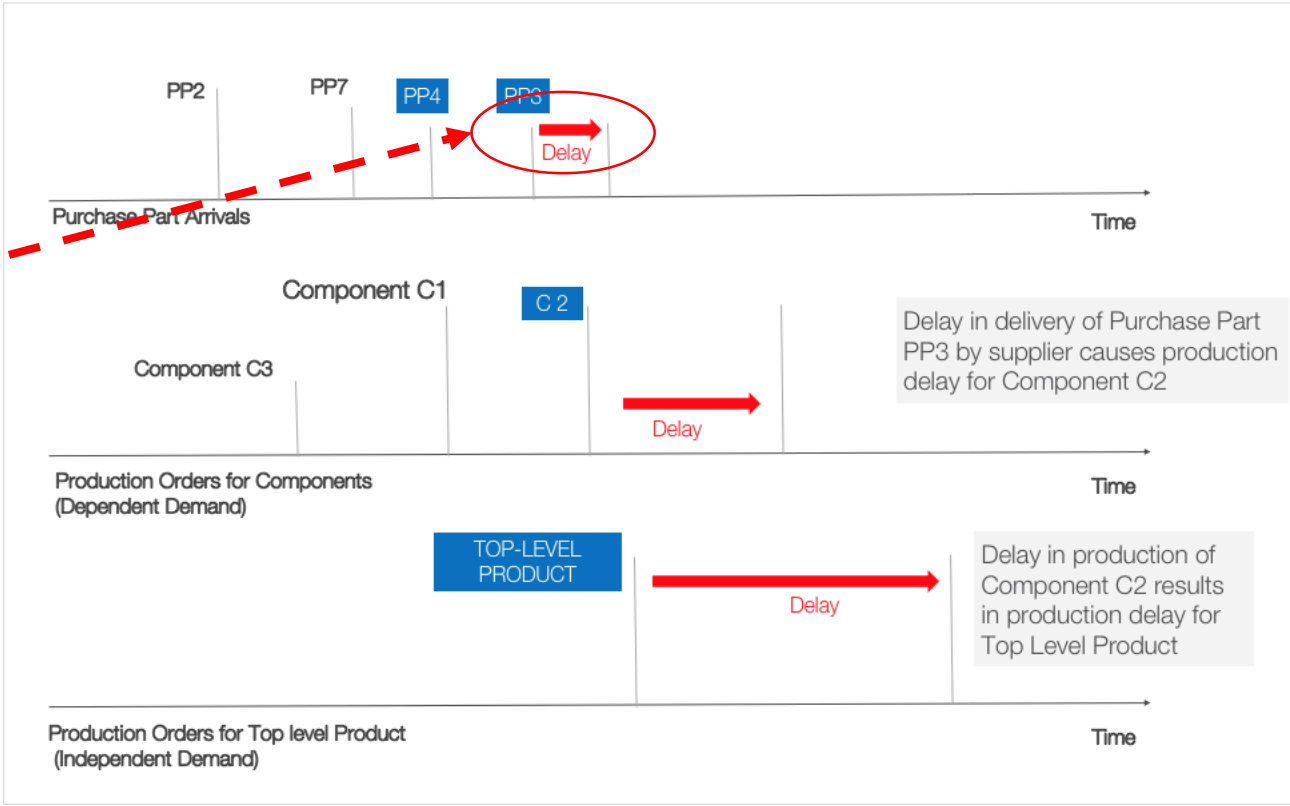

Purchase part delays lead to Trapped Inventory and Increased Inventory Holding Costs

Cascading Delays and Trapped Inventory. Delay in PP3:

- Delays C2 production
- Increases inventory holding costs of other purchased products

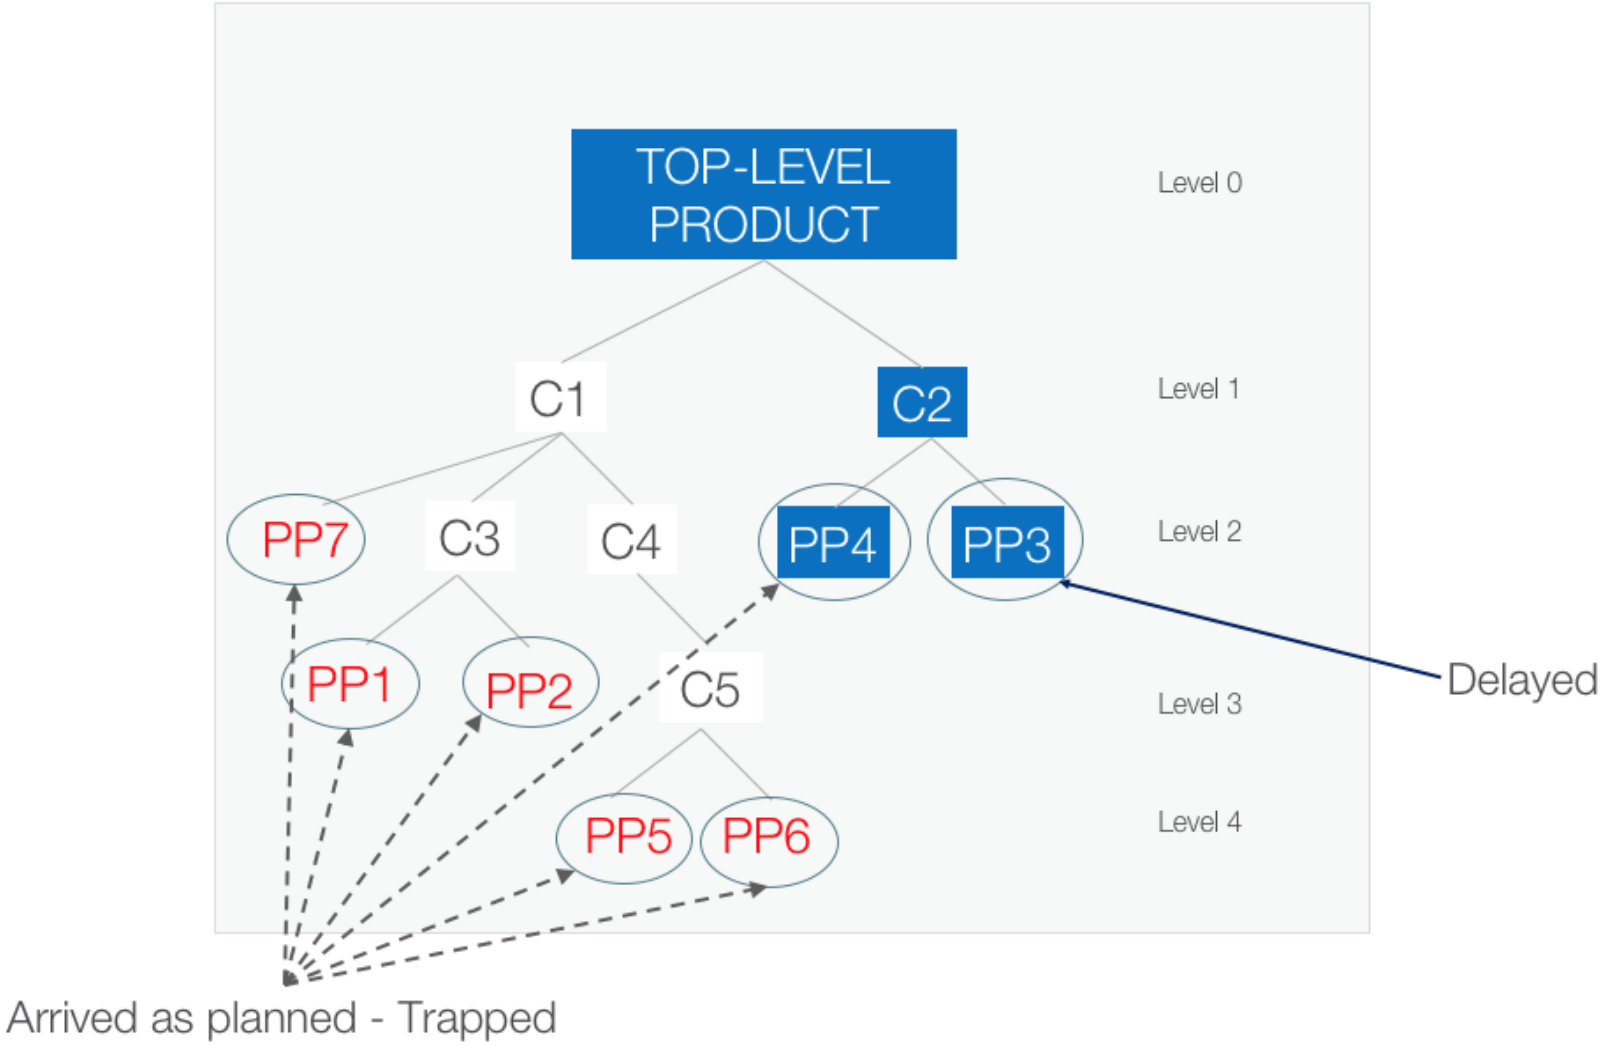

Production Often Needs to Be Shifted Disrupting Manufacturing Plans

Example Part

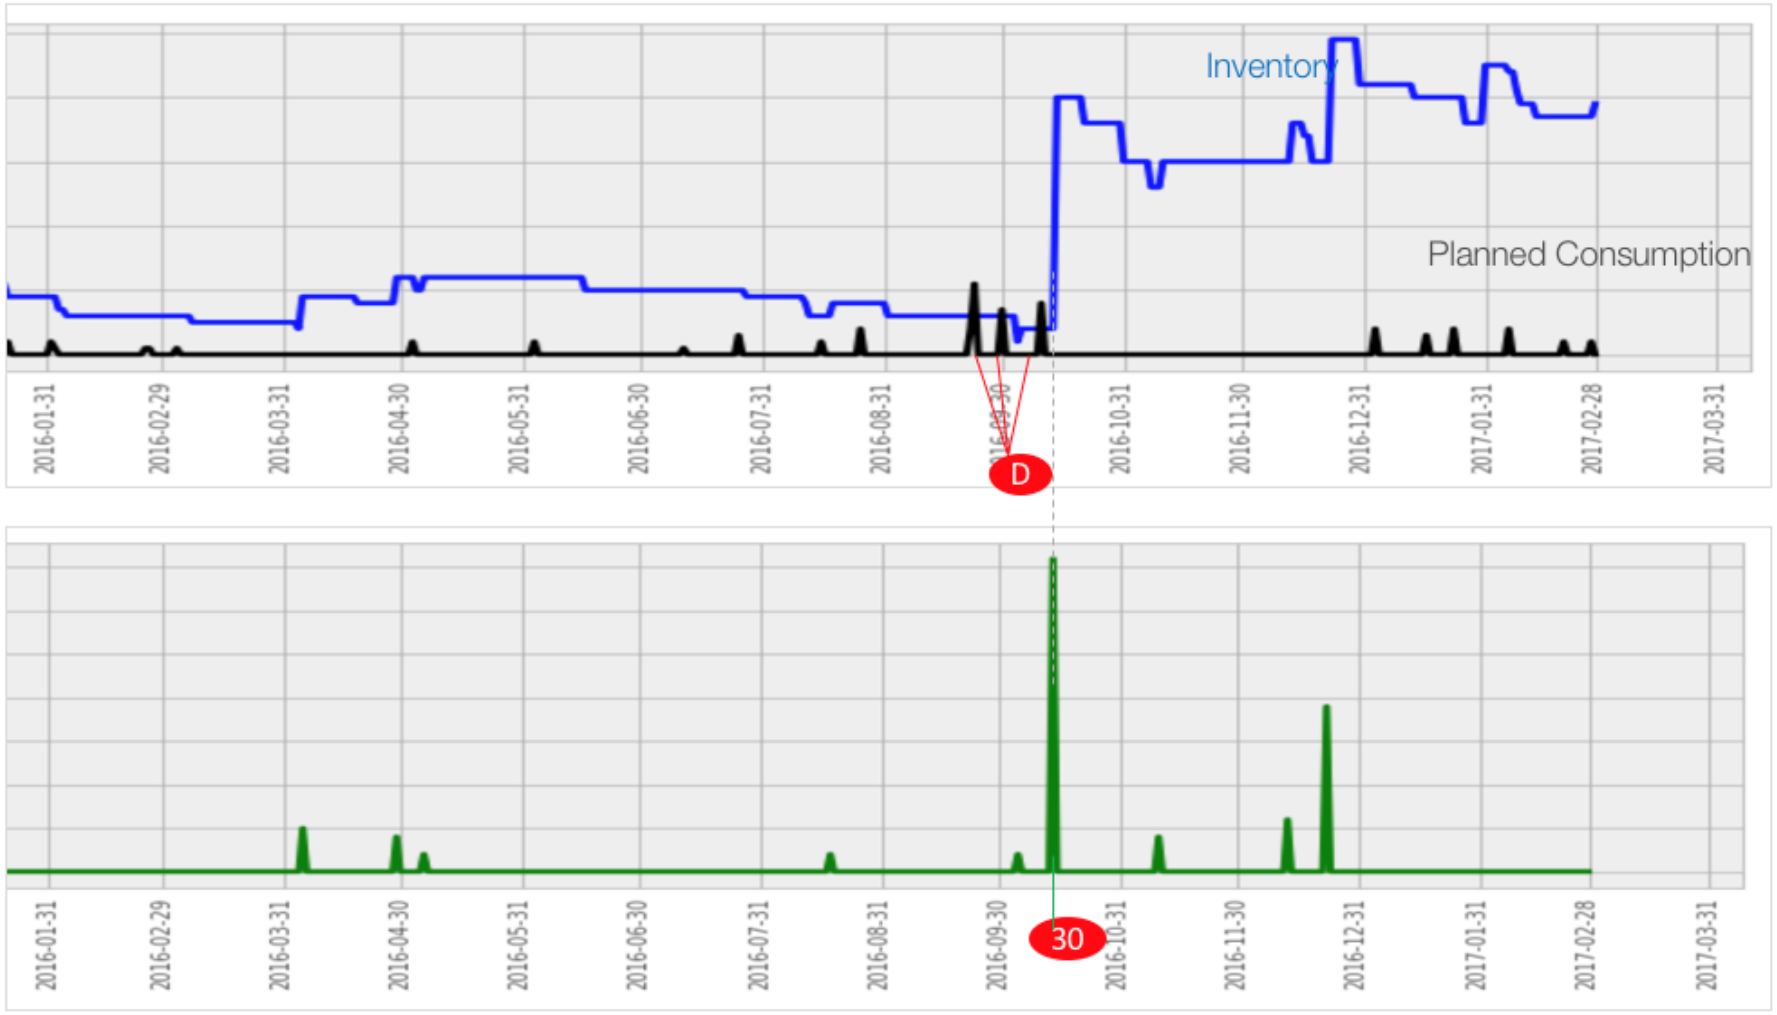

D — Planned consumption was greater than available inventory. Delay occurred on 3 separate occasions

30 — Insufficient inventory was caused by a 30 day delay in arrival of Purchase Part

C3 Inventory Optimization™ Enables Dynamic Optimization To Reduce Costs, Increase Service Levels, and Manage Suppliers

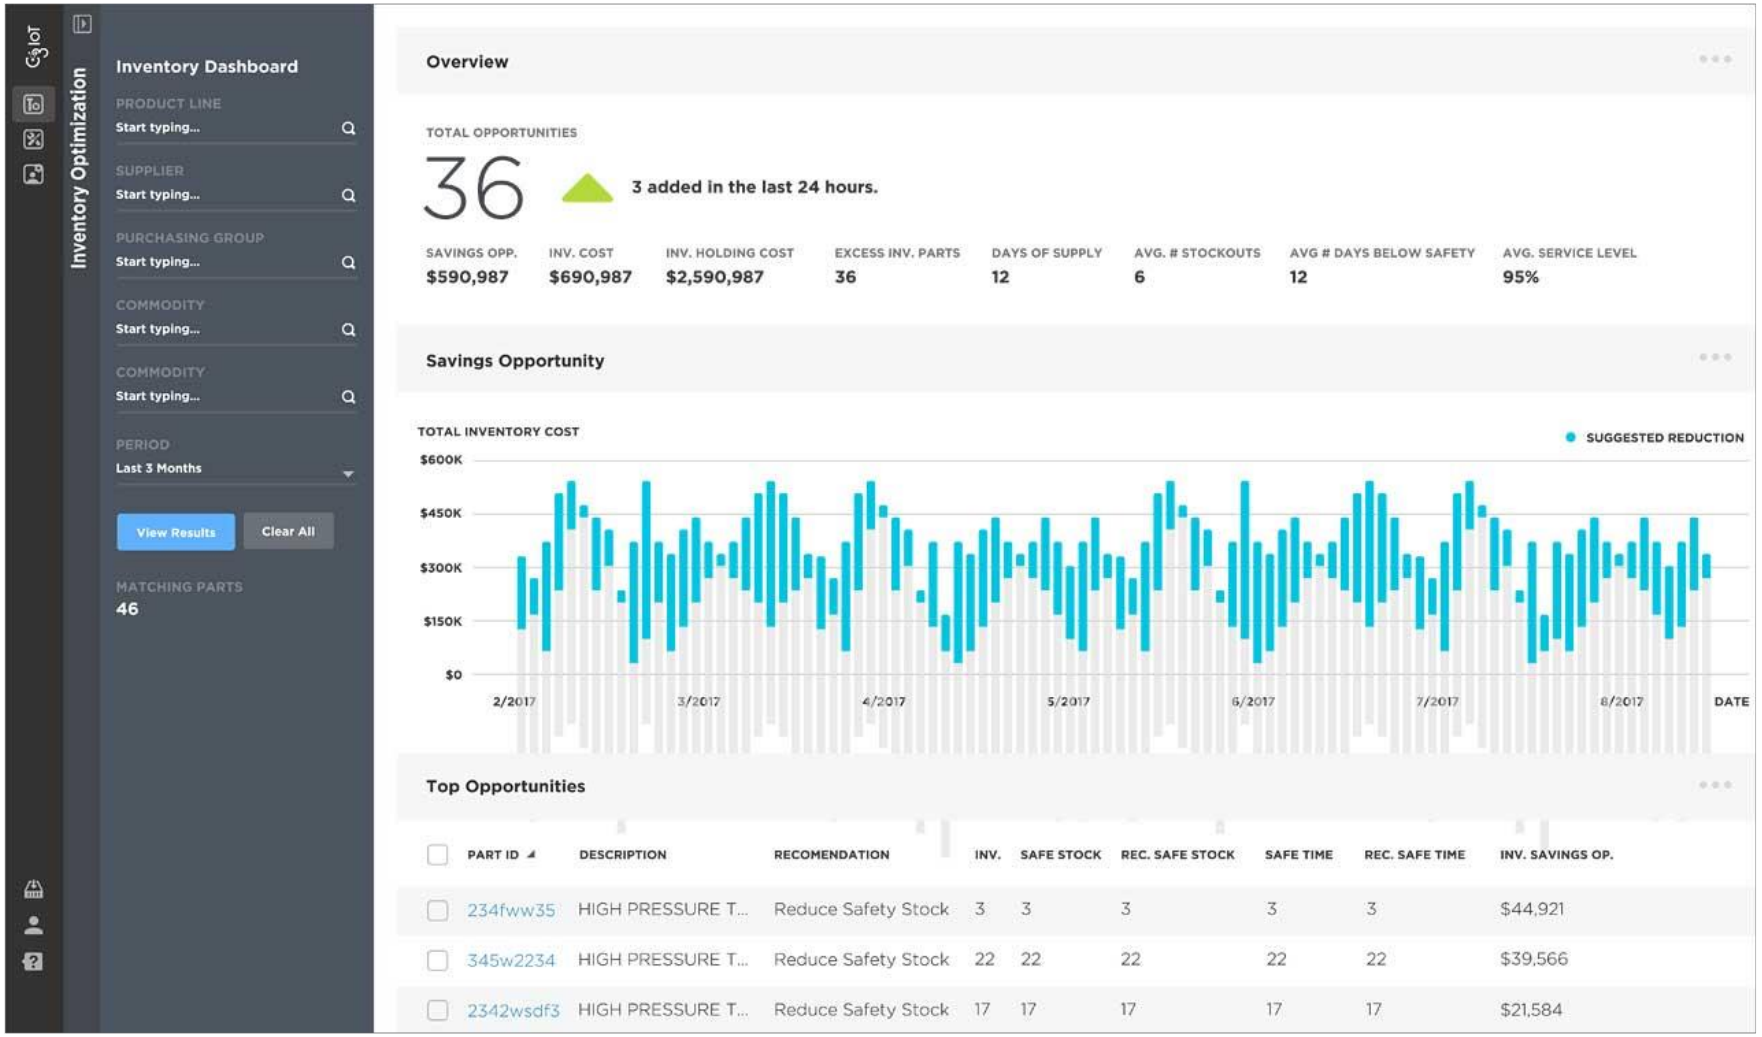

C3.ai

C3 Inventory Optimization

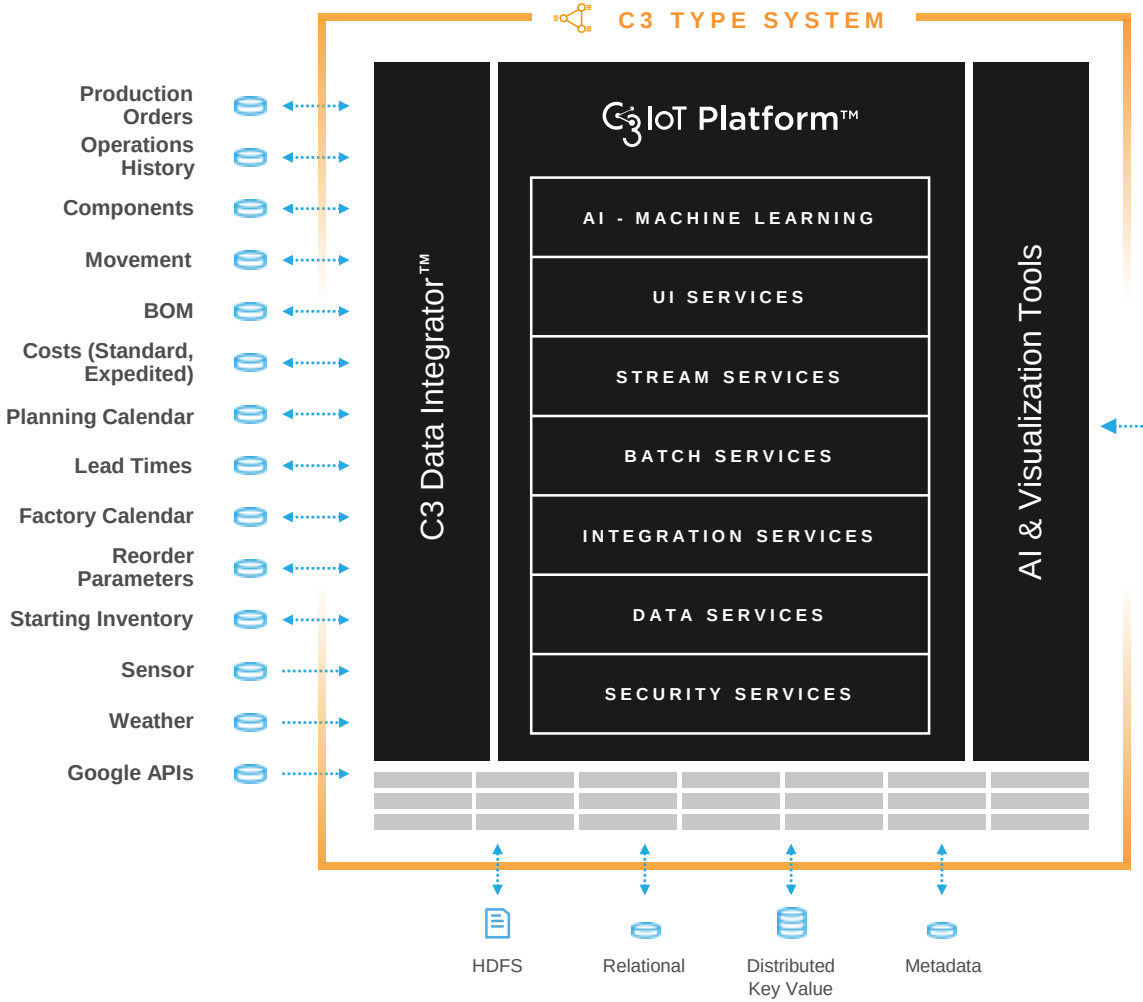

C3 Inventory Optimization™

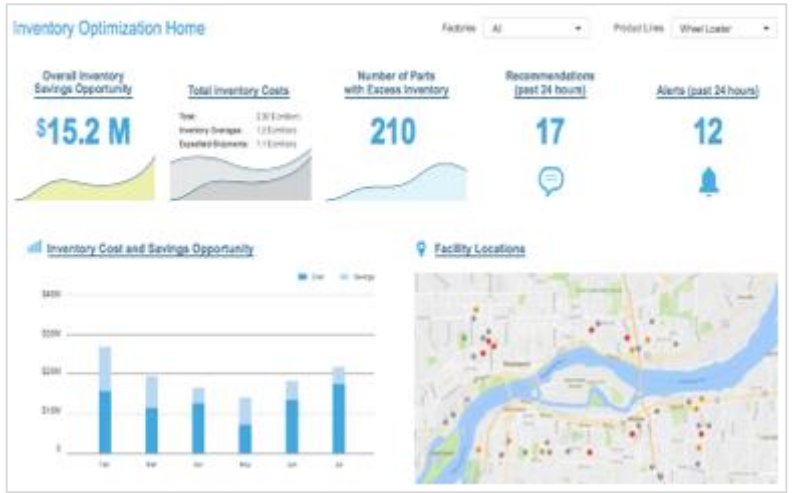

- **Dynamically optimize re-order parameters** (e.g., safety stock, safety time) while maintaining inventory at a certain confidence level
- **Improve service levels** and avoid production re-scheduling
- **Predict supplier delays/ decommits** at a part-level

C3.ai

What Makes this Problem Difficult

- **Companies often have complex supply chains**
- **Large numbers of products/ SKUs with complex Bills of Materials (BOMs)**
- **Significant uncertainties –** both on the supply and demand-side
- **Need to work within the constraints of legacy MRP systems** and business processes
- **Embed optimization recommendations within human workflows**

MACHINE LEARNING FOR INVENTORY OPTIMIZATION

Objective: Optimize safety stock given historical uncertainty in actual orders while maintaining confidence the customer will not run out of parts

Inputs	Optimization	Savings

Actual Inventory

Actual consumption

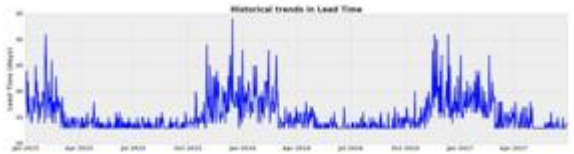

Planned consumption (from MRP)

Leverage natural stochasticity from historical variation in consumption

Optimization

What Safety Stock Minimizes Total Cost

Total Cost

=

Inventory Holding Cost

+

Shipping Cost

Subject to constraints (specifically either never run out of parts; or run out of parts on <x% of days)

Actual Inventory ($)

–

Simulated Optimized Inventory ($)

=

Historical Savings Opportunity ($)

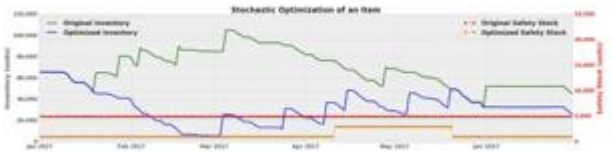

C3.ai

For Inventory Optimization, we use machine learning to model supply chain and demand uncertainties, which are used as inputs to optimize re-order parameters

Data Sources

- Production Orders
- Operations History
- BOM
- Costs
- Lead Times
- Reorder Parameters

Machine Learning Model

Machine Learning Model

Quantity Uncertainty:

- **Material Movement Uncertainty:** Uncertainty due to errors in factory operations (e.g., received parts were not in appropriate location or part of manufacturing process)
- **Demand Uncertainty:** Uncertainty due to differences between predicted demand forecast and actual consumption
- **Supplier Quantity Uncertainty:** Quantity uncertainty due to differences between each suppliers' ability to satisfy orders appropriately (e.g., supplier provided 90 parts on a 100-part order)

Time Uncertainty:

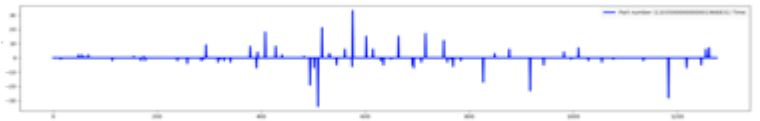

- **Supplier Time Uncertainty:** Time uncertainty due to differences between each suppliers' ability to satisfy orders appropriately (e.g., supplier is often X days late in having orders received)

Re-order Parameters

Optimized Safety Stock and Safety Time

AI/Machine learning and optimization details

Quantity uncertainty due to material movement and demand forecast

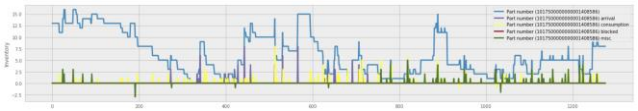

+

Quantity uncertainty due to supplier

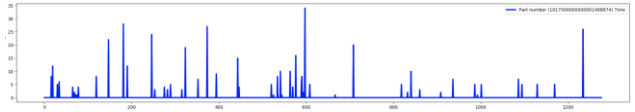

Time uncertainty due to supplier

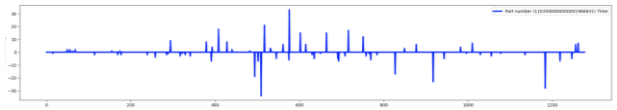

GLM for mixture models

One model for each part, facility

Distribution of time uncertainty

Quantile Regression Model

One model for each part, facility, and quantile value

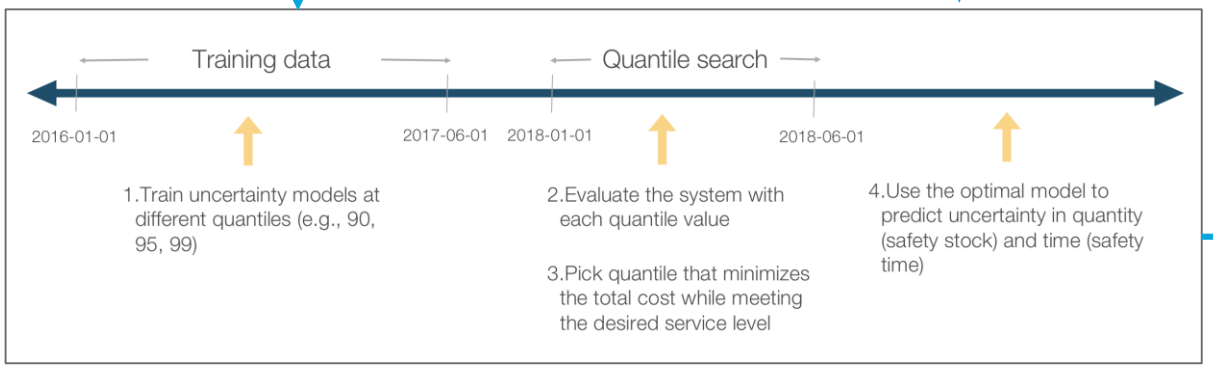

procurement process and inventory optimization

Optimizer

$$\underset{x^k, A^k, o^k, y^k, z^i}{\text{minimize}} \quad \sum_{k=1}^{N} \left[(c_o^k)^T A^k + (c_h^k)^T x^k \right] + \sum_{i=1}^{m} (c_p^i)^T z^i$$

$$\begin{aligned}
\text{subject to} \quad & \Delta x^k = e_1 x_0^k + \tilde{A}^k + A^k - D s^k \quad k = 1, \ldots, N \\
& x^k + Z s^k \geq 0 \quad k = 1, \ldots, N \\
& A^k - (MO^k + RV^k o^k) + M(1 - y^k) \geq 0 \quad k = 1, \ldots, N \\
& A^k - (MO^k + RV^k o^k) \leq 0 \quad k = 1, \ldots, N \\
& A^k - M y^k \leq 0 \quad k = 1, \ldots, N \\
& A^k \geq 0 \quad k = 1, \ldots, N \\
& o^k \geq 0 \quad k = 1, \ldots, N \\
& Z \leq Z_{SLA} \\
& \Delta Z \leq D \\
& Z \geq 0
\end{aligned}$$

Optimized scheduled Arrivals

Legacy compatible replenishment

Inventory

BOM → BOM Engine

Product Demand Forecast

Safety stock and safety time

Previous Scheduled Arrivals

Production Planning → MRP Engine

Scheduled Arrivals

EXAMPLE RESULTS FROM OPTIMIZATION ALGORITHM

Optimization goal: Minimize Safety Stock
Constraint: Ensure inventory level > 0 over the historical period
Example: Part: ‘A’
Current value of Safety Stock: 45
Optimized Safety Stock: 0

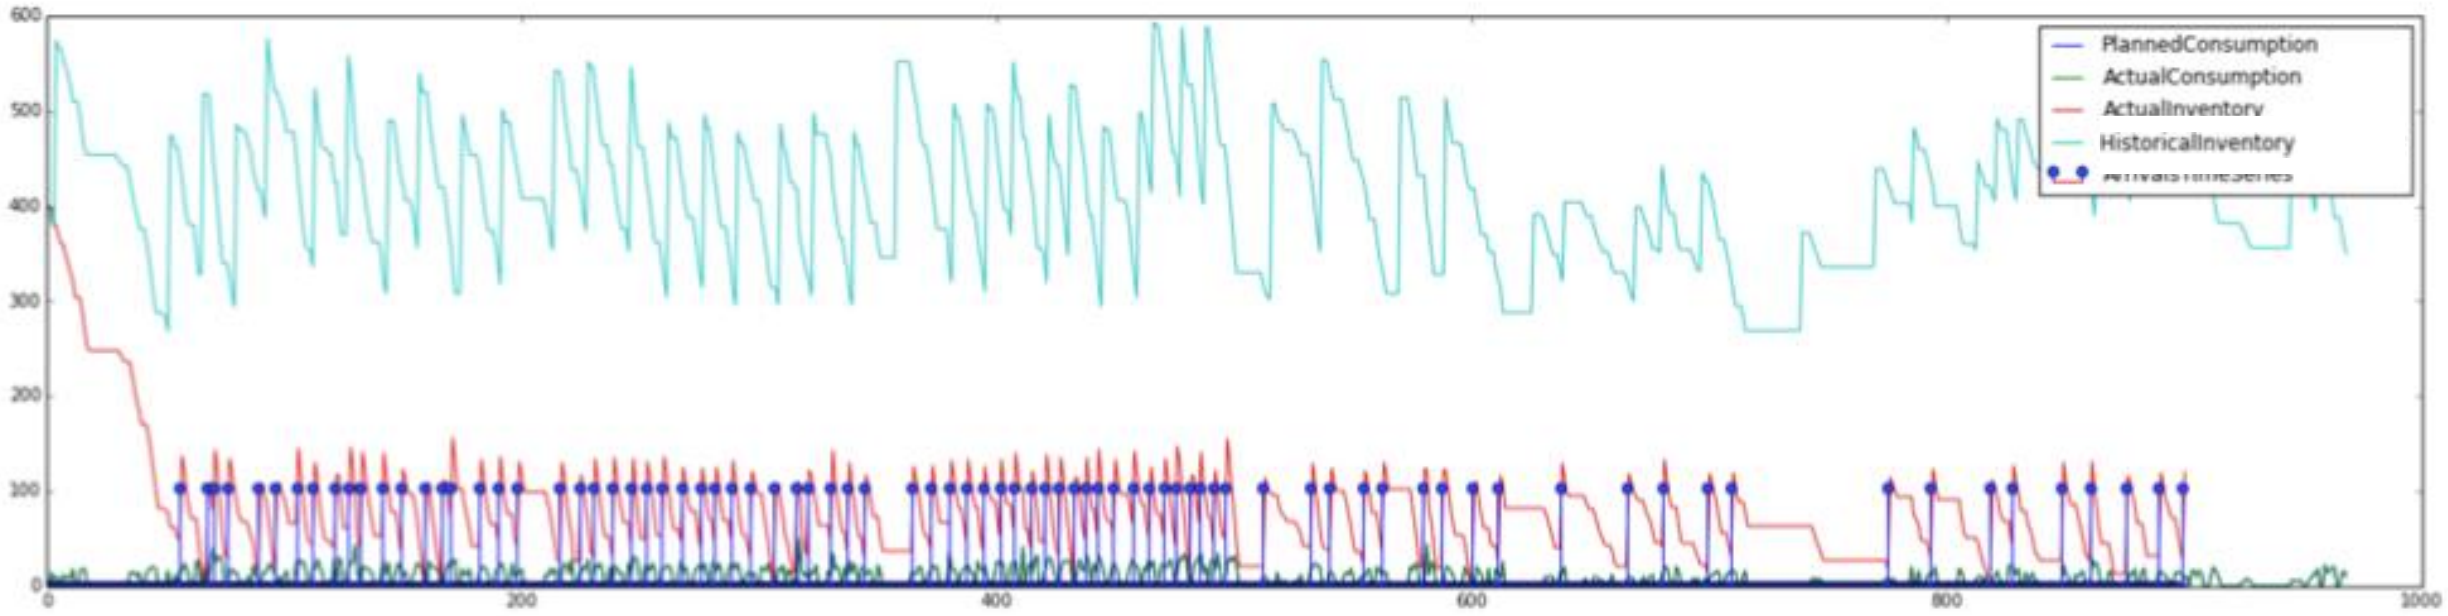

Inventory levels can be significantly reduced; safety stock can be reduced to 0

C3 Inventory Optimization™

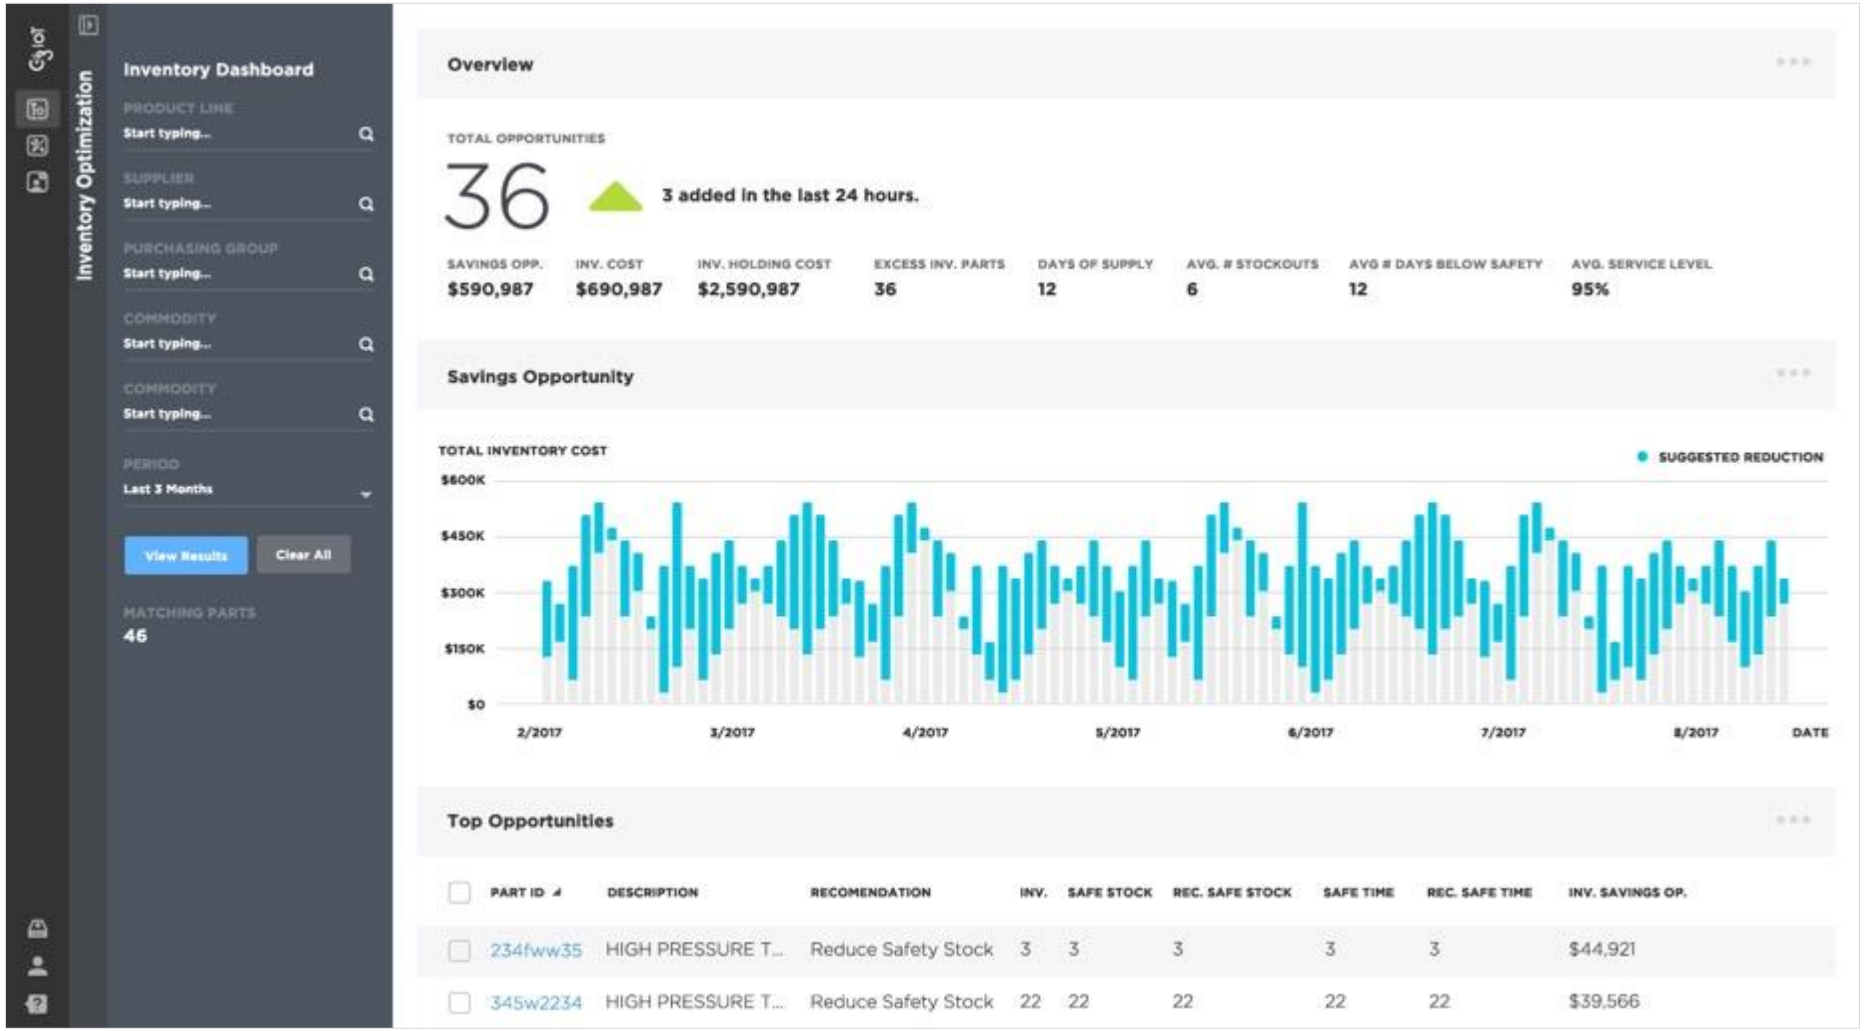

Typical Results with C3 Inventory Optimization

30-50% Savings in Inventory holding costs

70-80% Precision in identifying supplier delays

20-30% Increase in service levels

Billions of $ In economic value for most large enterprises

C3.ai

Fraud Detection Using Smart Meters

The Smart Grid: $2 Trillion Upgrade this Decade

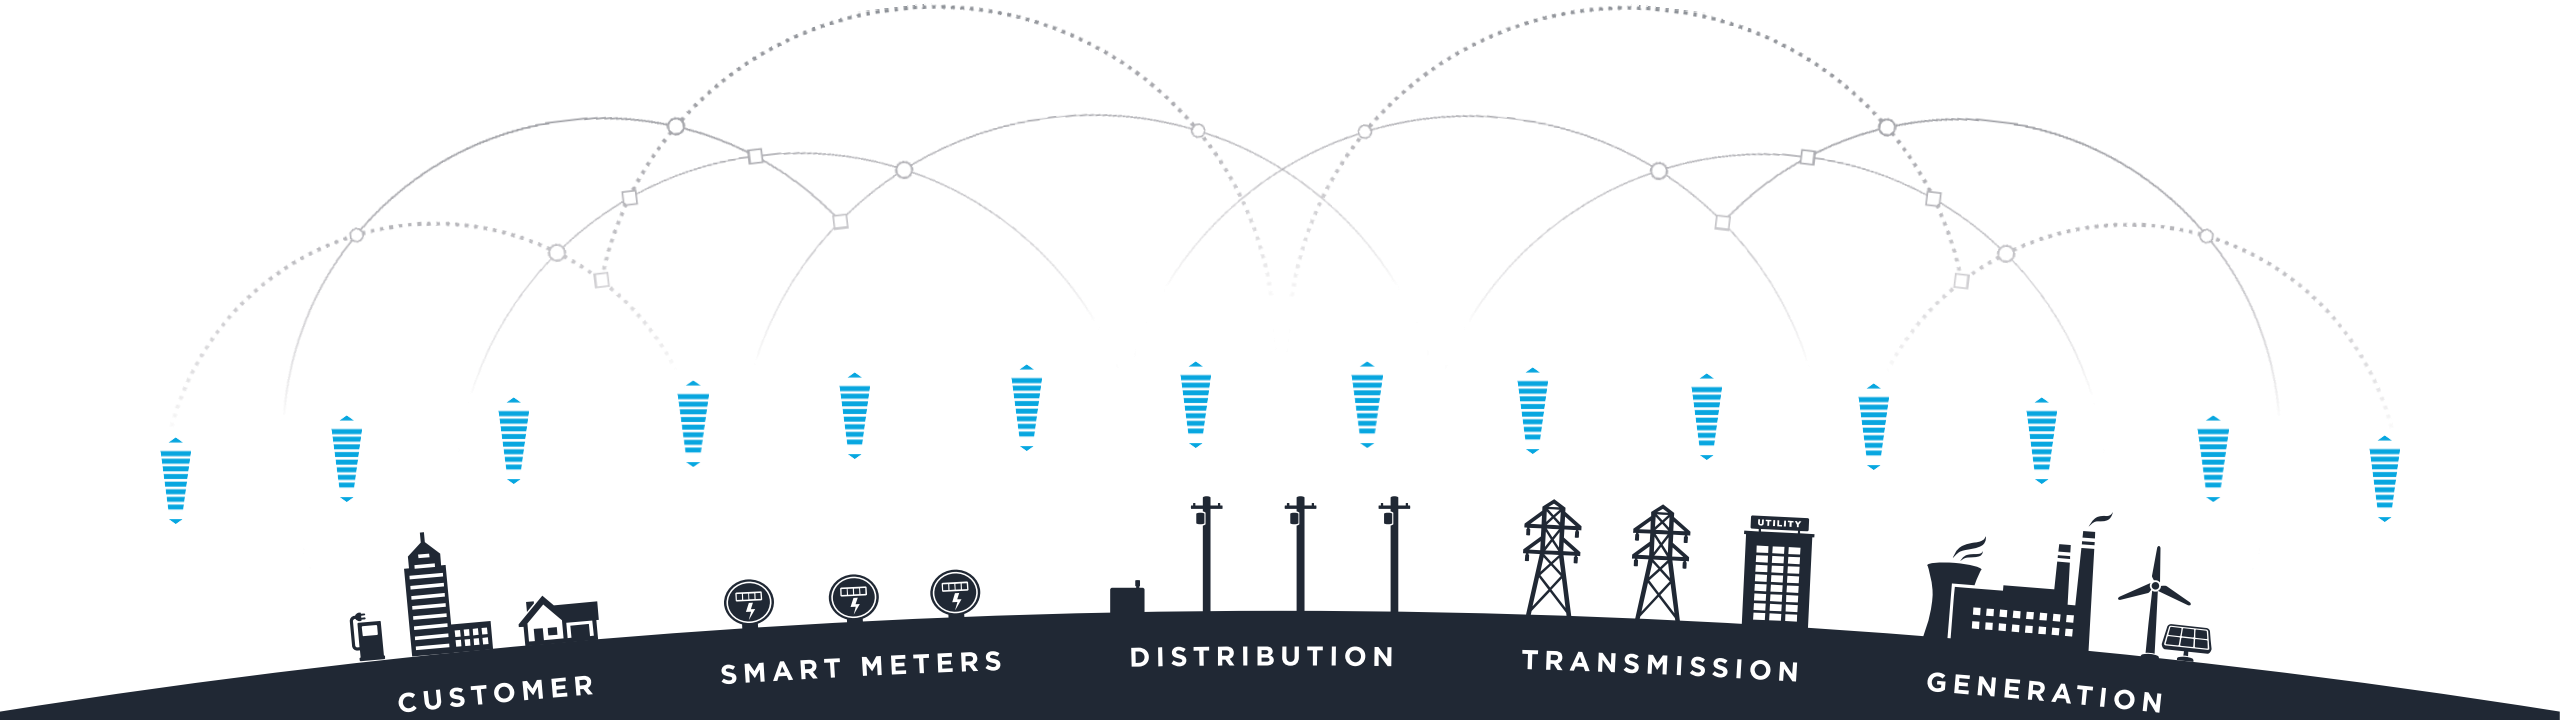

- CIS / Billing
- CRM
- Call Center

- MDM
- Head End
- Asset Mgmt

- OMS
- DMS
- WMS

- EMS
- SCADA
- GIS

- Market Ops
- Generation Ops
- TRM

Largest Production Deployment of AI & IoT Applications

96GW Generation Capacity

€71 Billion Annual Revenue

65M Customers Globally

73,000 Employees

TOP 5 Fortune's "Change The World" List

4 Continents

40 Countries

Detecting Electricity Fraud

- Customers bypass meter to get *free electricity*
- Surprisingly common: ~3% of electricity in Italy is estimated to be stolen
- Train a classifier based upon historical fraud investigations (~100,000 cases)
- Input data includes 86 signals sent from meter, aggregated daily

Challenges in Meter Fraud

- Vastly different methods of theft
- Ability to recover payment across different regions
- We only get credit for a fraud prediction if utility is able to recover payment (“recover energy”)
- Limited resources and budgets across regions, each office wants highest probability, highest value cases in their own region

Our path to the final model

- Compared different Deep Learning techniques:
  - Multi-layer Perceptron (MLP),
  - Convolutional Neural Networks (CNNs),
  - Recurrent Neural Network (RNN)

- **Winner:**
  - RNN (stacked BiLSTM with attention) has the best results – preliminary results exceed conventional machine learning

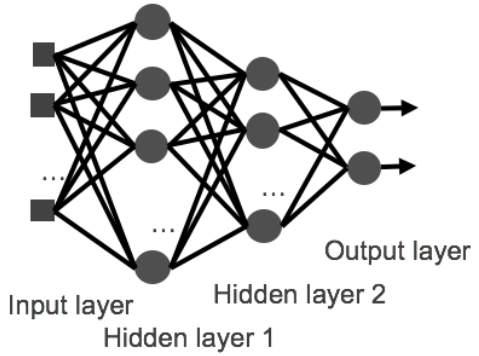

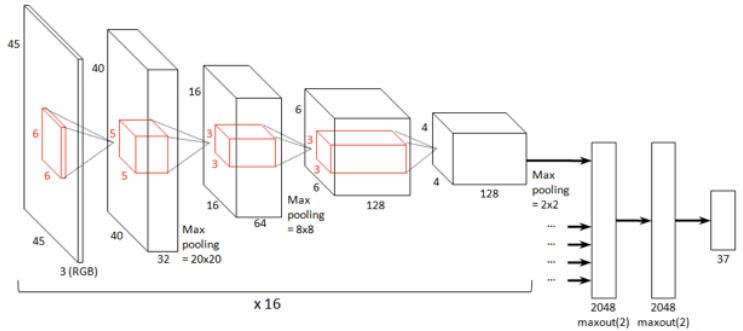

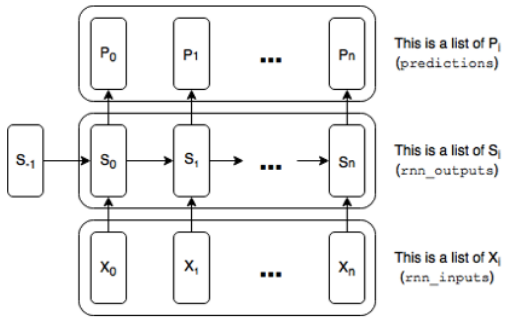

Revenue Protection – Convolutional Neural Net

- Deep learning has been proven superior for image and audio applications and is nowadays as good as humans in most image recognition tasks.
- Since humans can look at graphs of consumption, serviceActive, workOrders, events etc., and detect fraud and anomalies, we think that deep learning will be able to learn from graphs to detect fraud/anomalies.

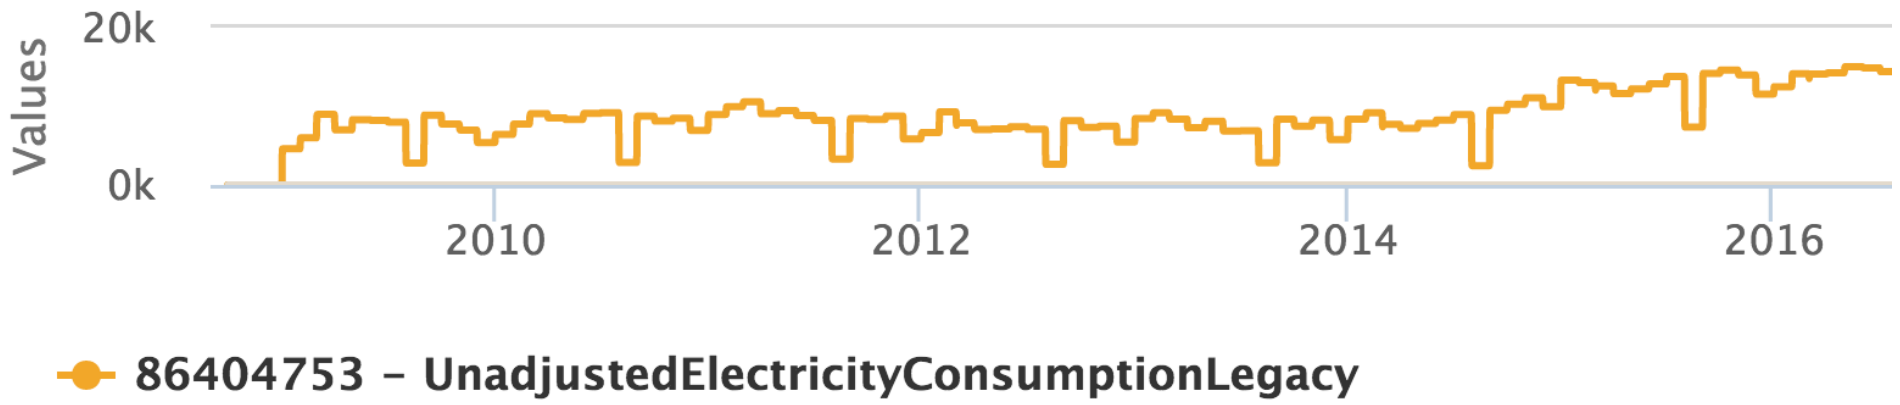

Revenue Protection – Convolutional Neural Net

- Convolutional neural nets (CNN) have been proven effective for images. To better handle time series data, we modified a CNN architecture.

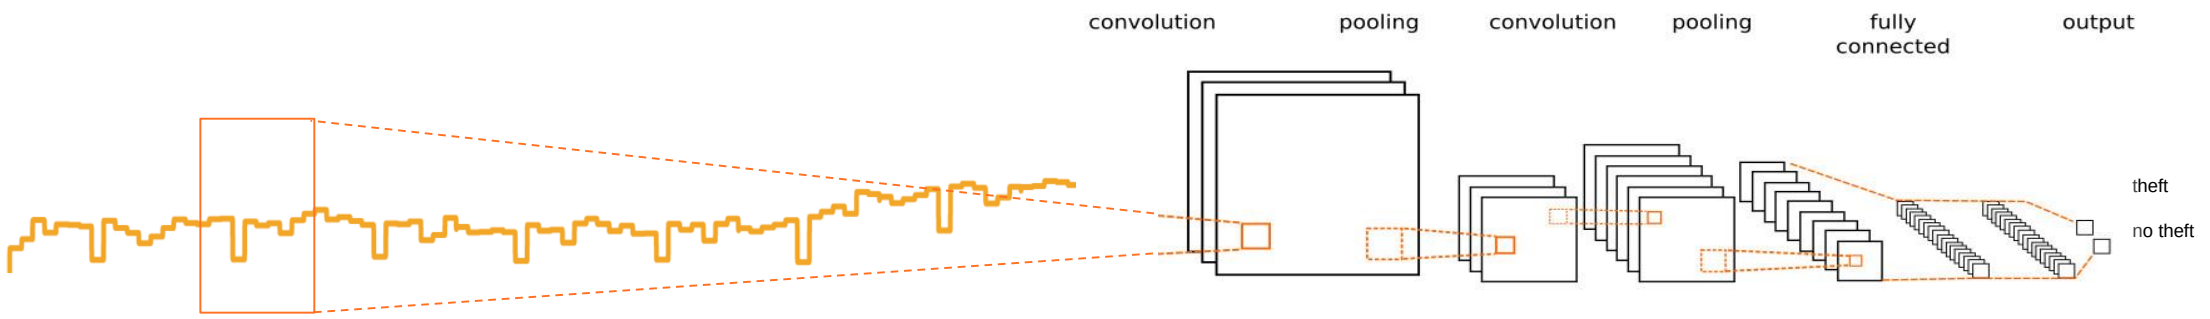

CNN architecture adapted to this problem for handling time series data

- The prediction results of CNN are comparable to the baseline (current classifier, traditional machine learning).

Final Machine Learning Approach

- **Approach:** stacked BiLSTM with attention, predict probability of theft from raw time series
- Separate classifier to predict probability of recovering energy, and third analytical rule to estimate recovered energy
- Train system end-to-end to maximize tradeoff between total recovered energy cases and amount of recovered energy
- Ranking-based loss function

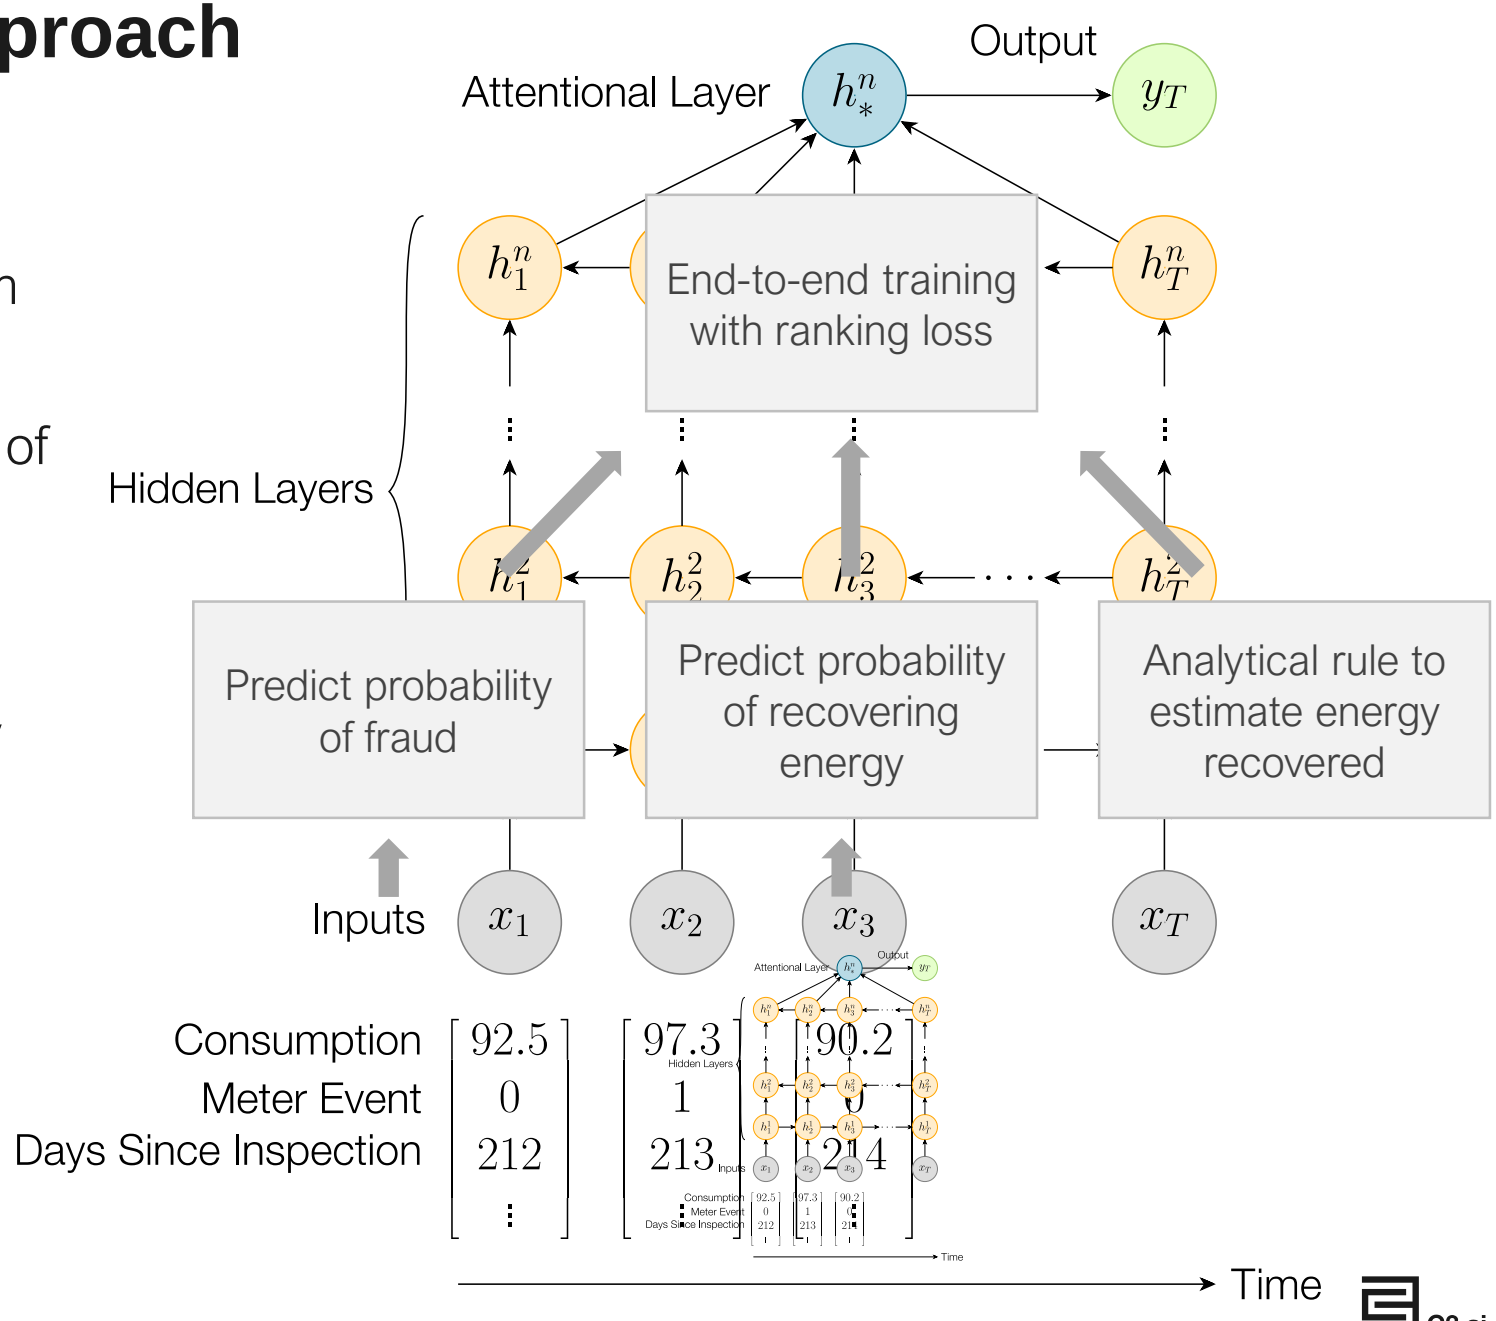

Ranking-based cost function

- We used the P-Norm Push*: A Simple Convex Ranking Algorithm that Concentrates at the Top of the List
- Perform especially well near the top of the ranked list
- Perform sufficiently well on the rest of the list

Height of this example = 4

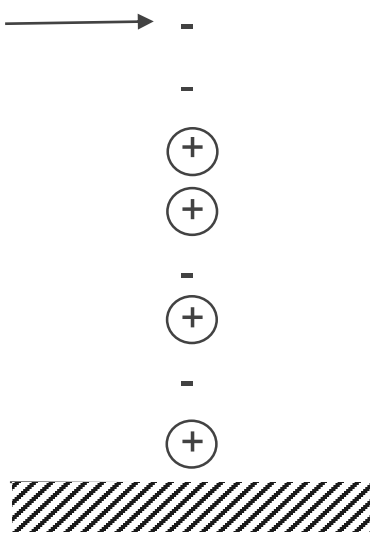

$$\text{Height}(k) := \sum_{i=1}^{I} \mathbf{1}_{[f(\mathbf{x}_i) \leq f(\tilde{\mathbf{x}}_k)]}$$

- The objective, i.e. sum of the prices for the negative examples, is:

$$R_{g,1}(f) := \sum_{k=1}^{K} g\left(\sum_{i=1}^{I} \mathbf{1}_{[f(\mathbf{x}_i) \leq f(\tilde{\mathbf{x}}_k)]}\right) \qquad \text{where} \qquad g(r) = r^p$$

* Rudin, C., 2009. The p-norm push: A simple convex ranking algorithm that concentrates at the top of the list. *Journal of Machine Learning Research*, *10*(Oct), pp.2233-2271.

Results

- Trained a BiLSTM network with ranking loss
- Training period 2015-12-01 to 2016-12-01
- Testing period 2017-01-01 to 2017-04-01
- Backtest methodology:
  - Obtain score distribution on a random set of 50K
  - Count the number of TPE and FPs from the test set in each percentile and calculate the hit-rate
- Potential improvements at the top of the list were observed

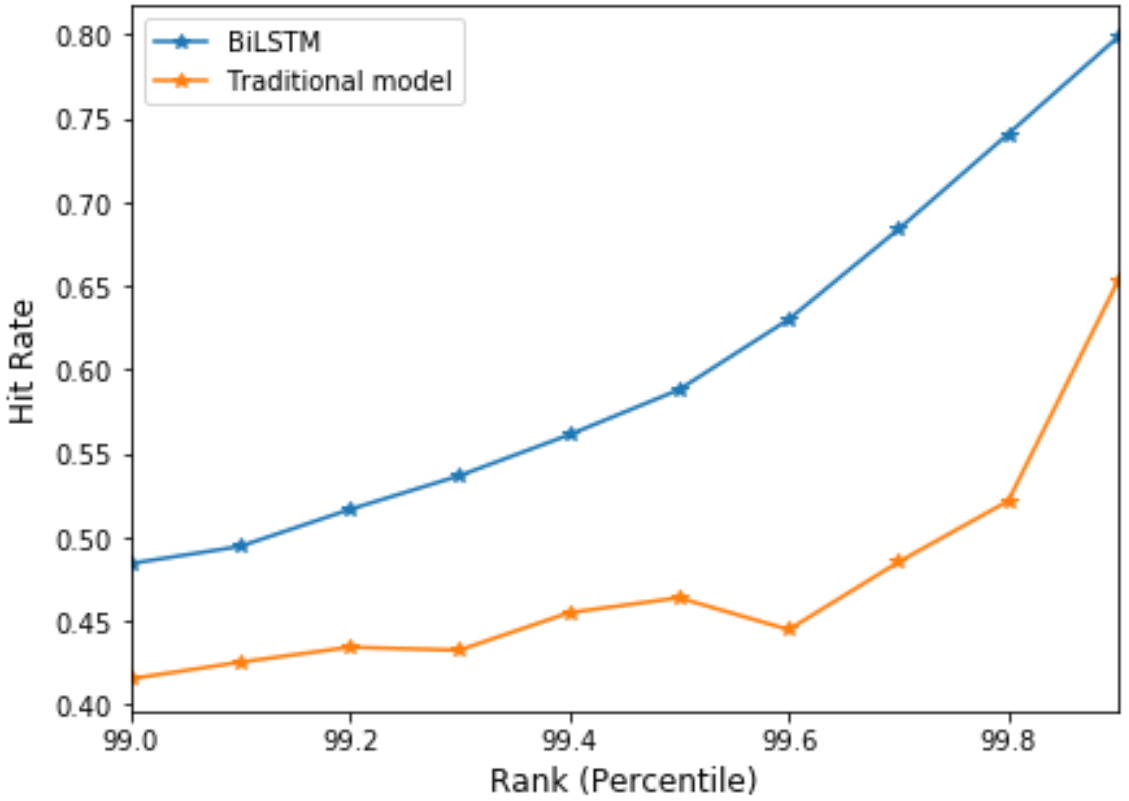

C3.ai

Results: Meter Fraud Detection

	HUMAN EXPERTS	TRADITIONAL ML APPROACH	DEEP LEARNING
Number of raw features	??	575	86
Development time spent	~20 years	2 years	2 months
Hit-rate (precision) at available resources	~25%	44%	49%

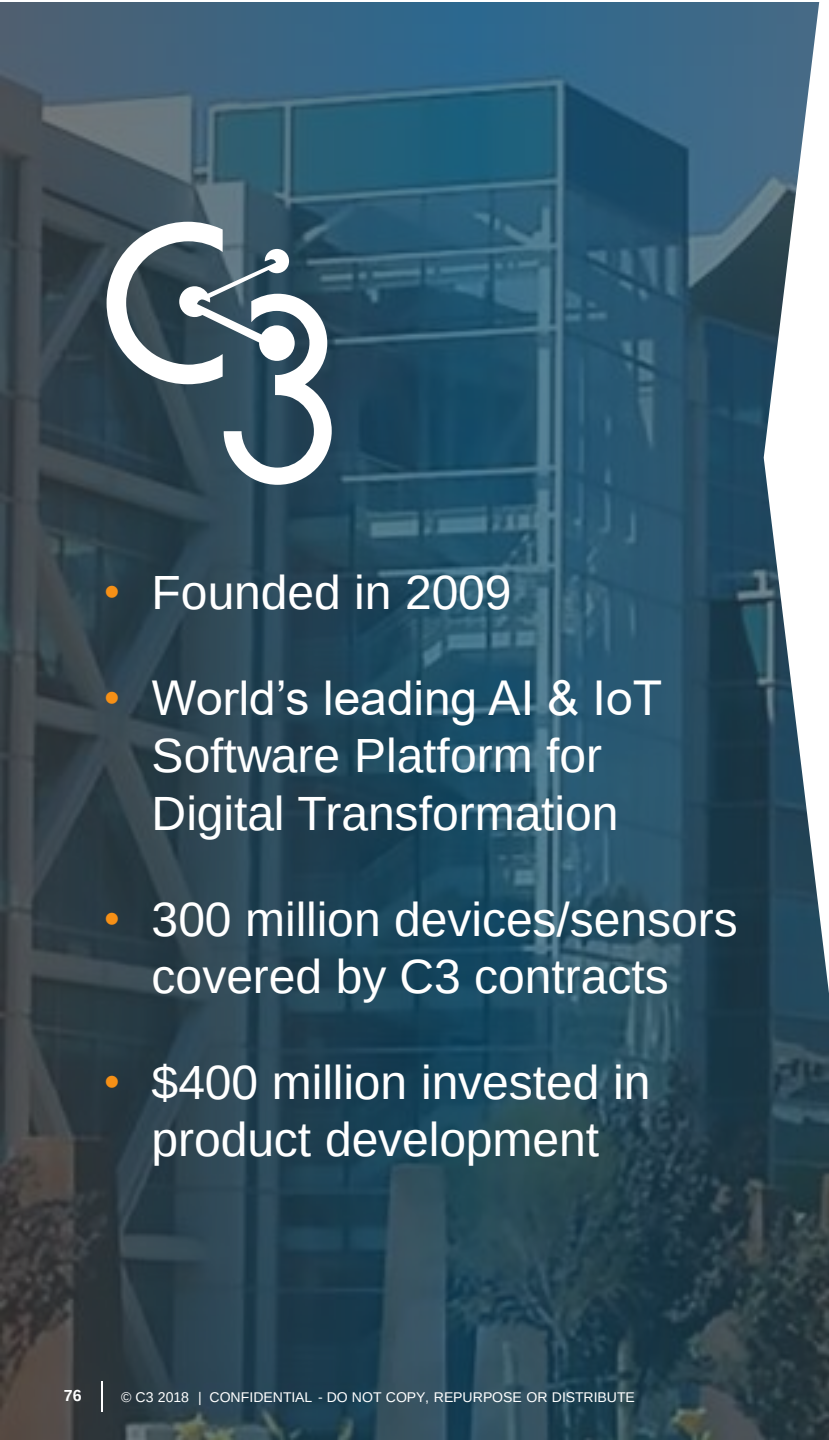

- Founded in 2009
- World's leading AI & IoT Software Platform for Digital Transformation
- 300 million devices/sensors covered by C3 contracts
- $400 million invested in product development

CNBC DISRUPTOR 50 2018

glassdoor 2018 TOP CEOs EMPLOYEES' CHOICE

glassdoor 2017 BEST PLACES to work Employees' Choice

Goldman Sachs 100 MOST INTRIGUING ENTREPRENEURS

JMP HOT 100 SOFTWARE COMPANIES 2018

FORRESTER® WAVE LEADER 2018 Industrial IoT Software Platforms

Forbes 2017 THE Cloud100 WORLD'S BEST CLOUD COMPANIES

Forbes THE 2018 Cloud100 THE WORLD'S BEST CLOUD COMPANIES

datanami DATA SCIENCE • AI • ADVANCED ANALYTICS 2018 EDITORS' CHOICE AWARDS

EY Entrepreneur Of The Year® 2017 Award Winner

constellation RESEARCH

A NextGen Platform for NextGen Apps

C3.ai

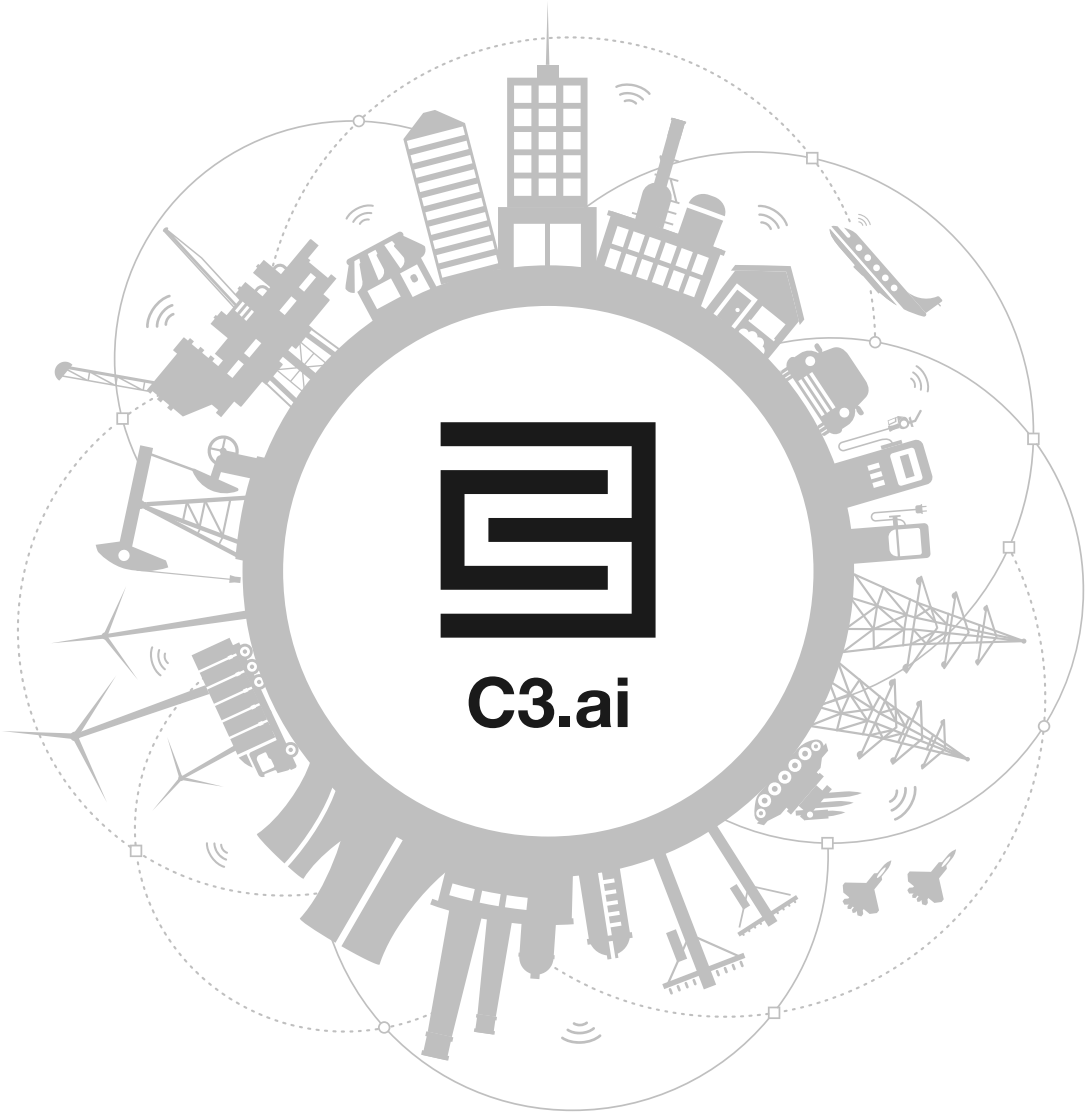
C3.ai

Appendix

Text Sentiment Analysis Using Deep Learning

Machine Vision Use Cases

Deep Learning

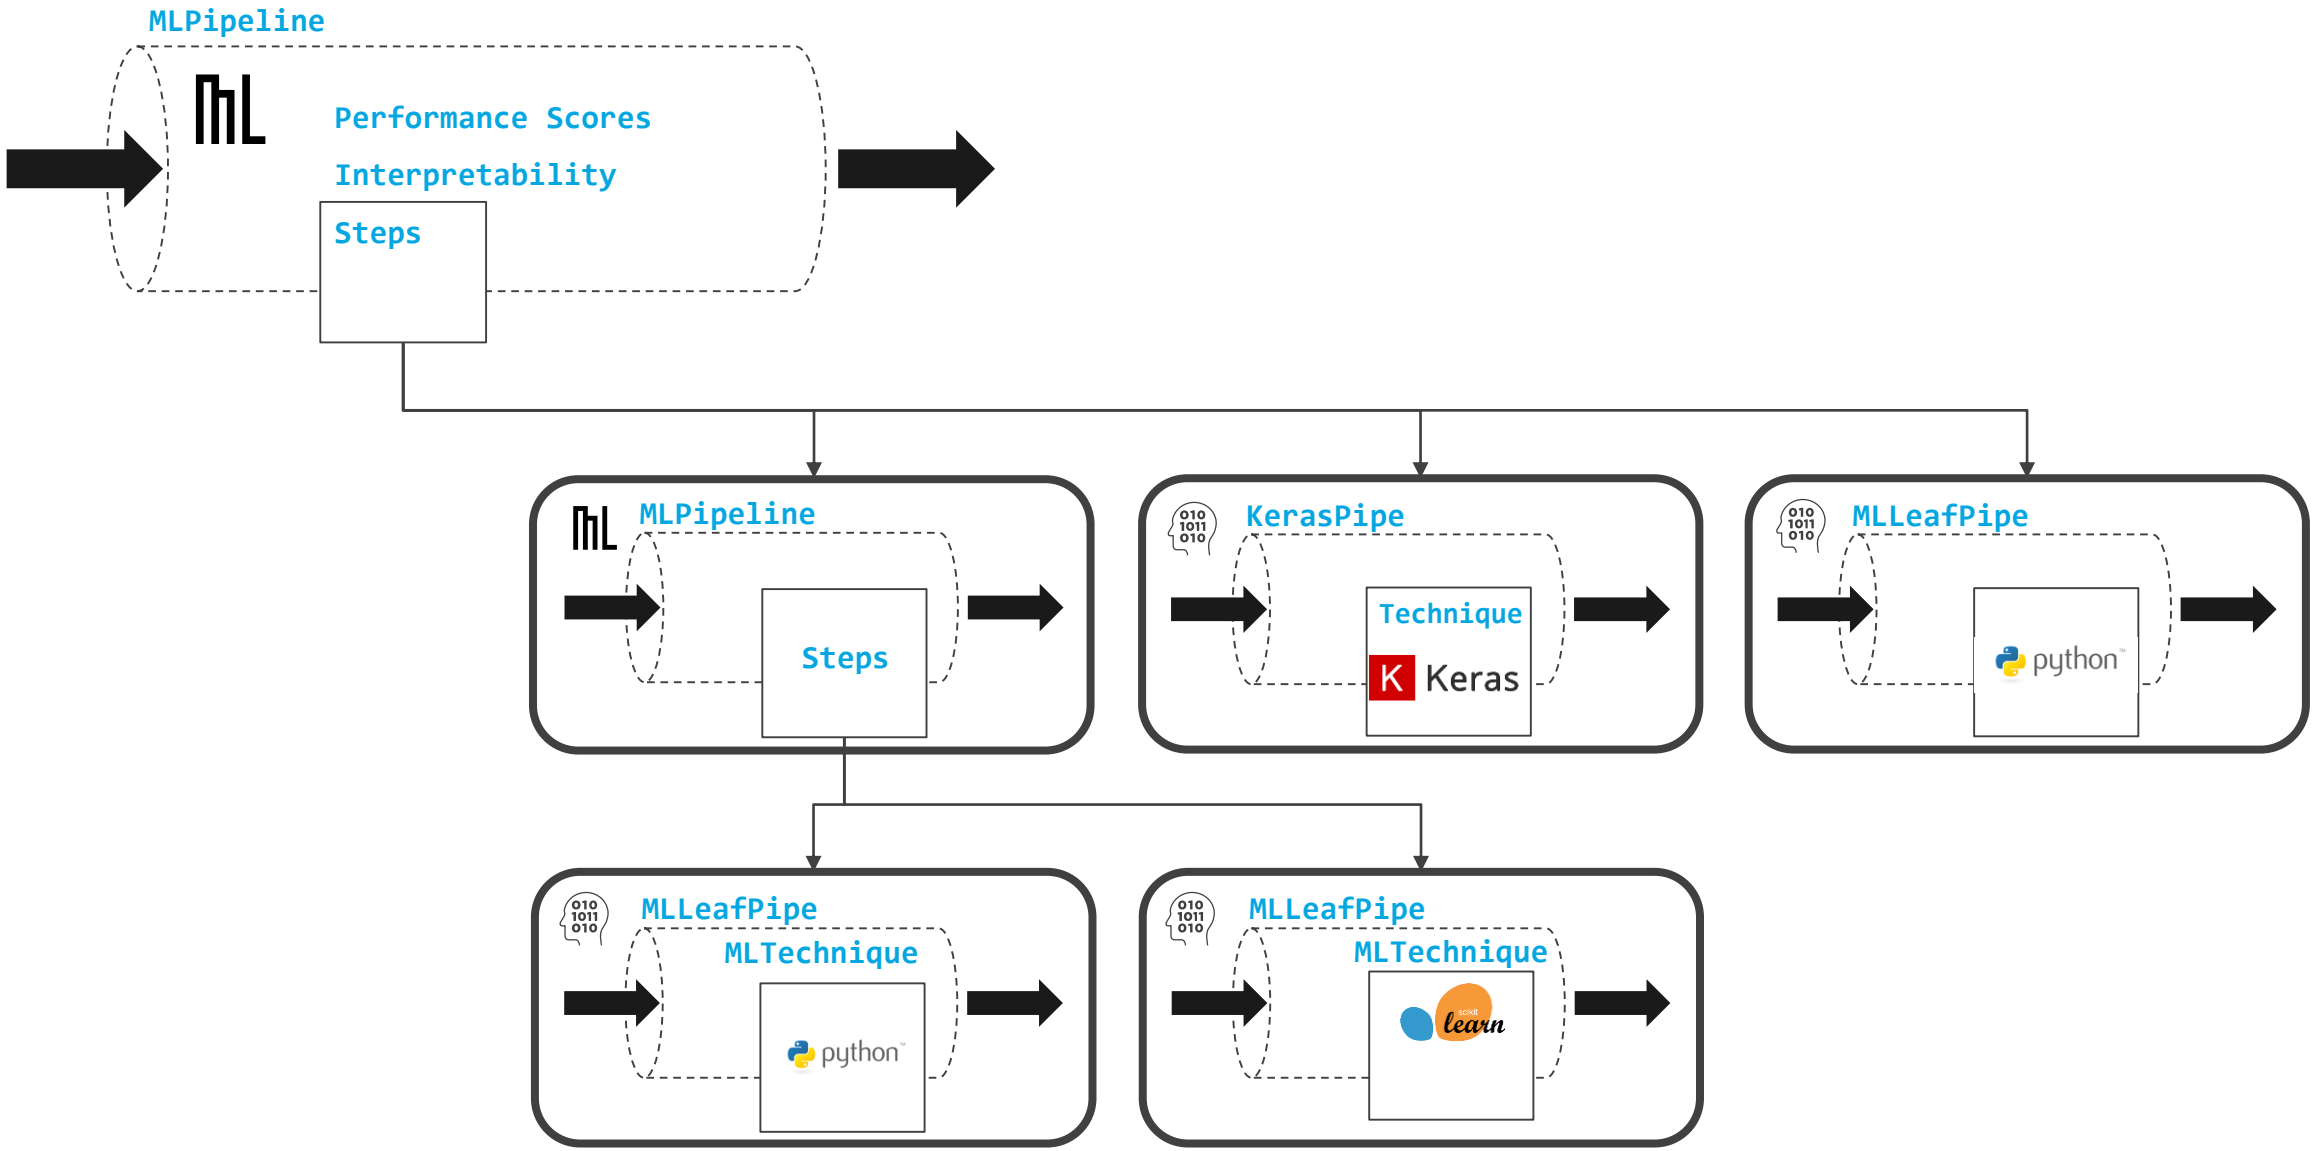

C3.ai

Deep Learning - Machine Vision

Object Detection, Classification, and Recognition

Detection

Recognition (and classification)


```
Predicted class : Audi TT with probability score : 1.0
```

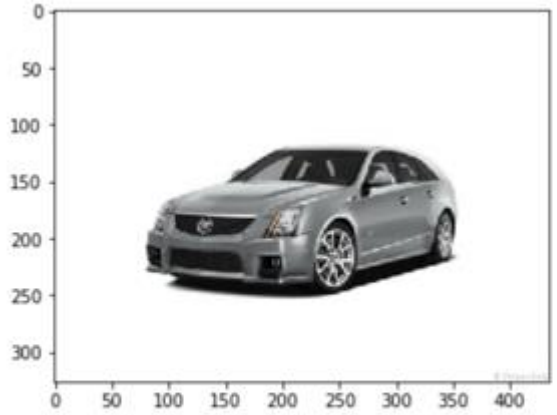


```
Predicted class : Cadillac cts-v with probability score : 0.81
```


C3.ai

Ariel Image Analysis

- Detect and count cars
- Size of objects in satellite images is ~10px
- A CNN network inspired by RetinaNet was trained
- Non-max suppression for cars in overlapping regions
- Trained model achieved a state-of-the-art performance with F1 score of 0.91
- A few false positives include garbage bins (looks similar to cars from overhead)

Feature Pyramid Network

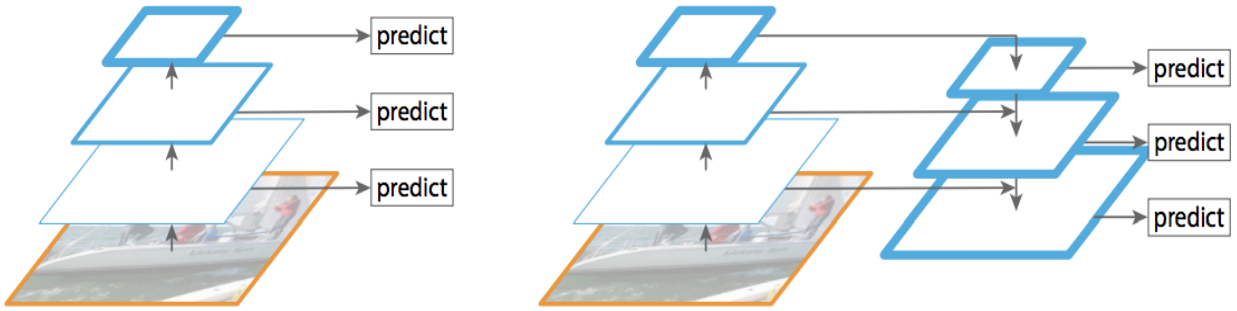

Localized cars from satellite image on Toronto city

Safety Adherence Detection Model

- Detect and classify helmets and croissants
- Upsampled the dataset to balance it and heavily used data augmentation techniques like rotation, zoom, flip,..
- Used transfer learning on RetinaNet architecture with frozen ResNet50 backbone and weights trained on COCO
- The best performing model achieved an mAP of 0.74

Predictions from Object Detection Model trained on C3 ML Pipelines

Vehicle Model Recognition

- Recognize make and model of car from security surveillance cameras
- A custom dataset is prepared from various data sources
- Heavily used data augmentation techniques to add more variance and enhance performance of the model
- Used differential learning rate for convolutional layers in the InceptionV3 network
- The best performing model uses transfer learning on InceptionV3 network trained on ImageNet
- Best performing model achieves accuracy of 94% on balanced test set


```
Predicted class : Volvo XC90 with probability score : 0.97
```


```
Predicted class : Lamborghini Gallardo with probability score : 1.0
```

Natural Language Use Cases

Sentiment Analysis of Images of Tweets

Goal: Demonstrate the Natural Language Processing (NLP) and deep learning capabilities of the C3 AI Suite.

Use case: Parse tweets from PDF files and classify them as having positive/negative sentiment.

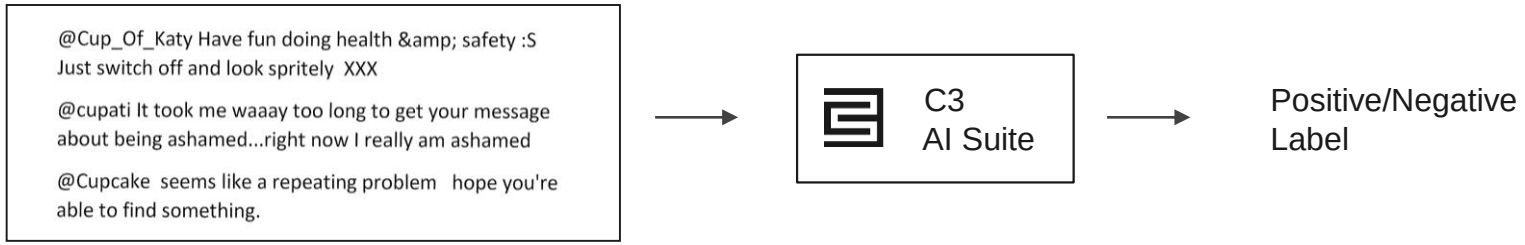

Technologies integrated:

OCR

Minecart/Tesseract Python library for text extraction from PDF files.

Regex

Regular expressions for preprocessing text.

spaCy

SpaCy for advanced NLP (NER, POS, word embeddings)

TensorFlow

TensorFlow for deep learning

C3.ai

Sentiment Analysis of Images of Tweets -- Training Pipeline

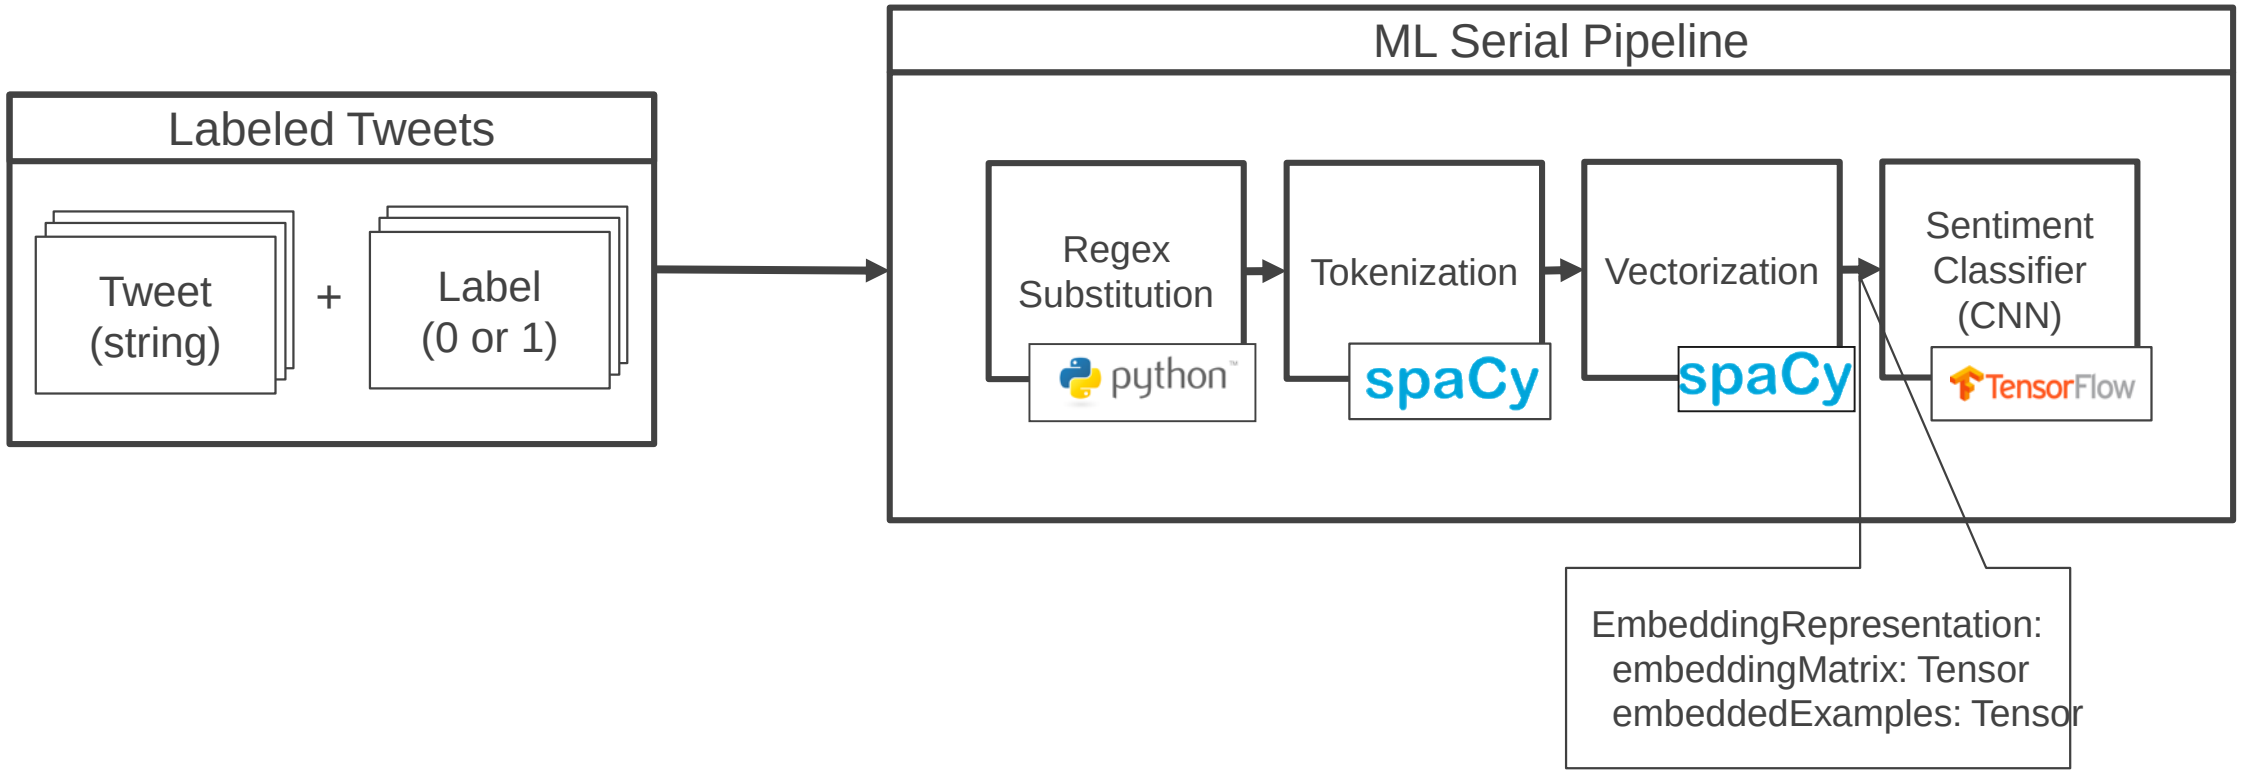

C3.ai

Sentiment Analysis of Images of Tweets -- Inference Pipeline

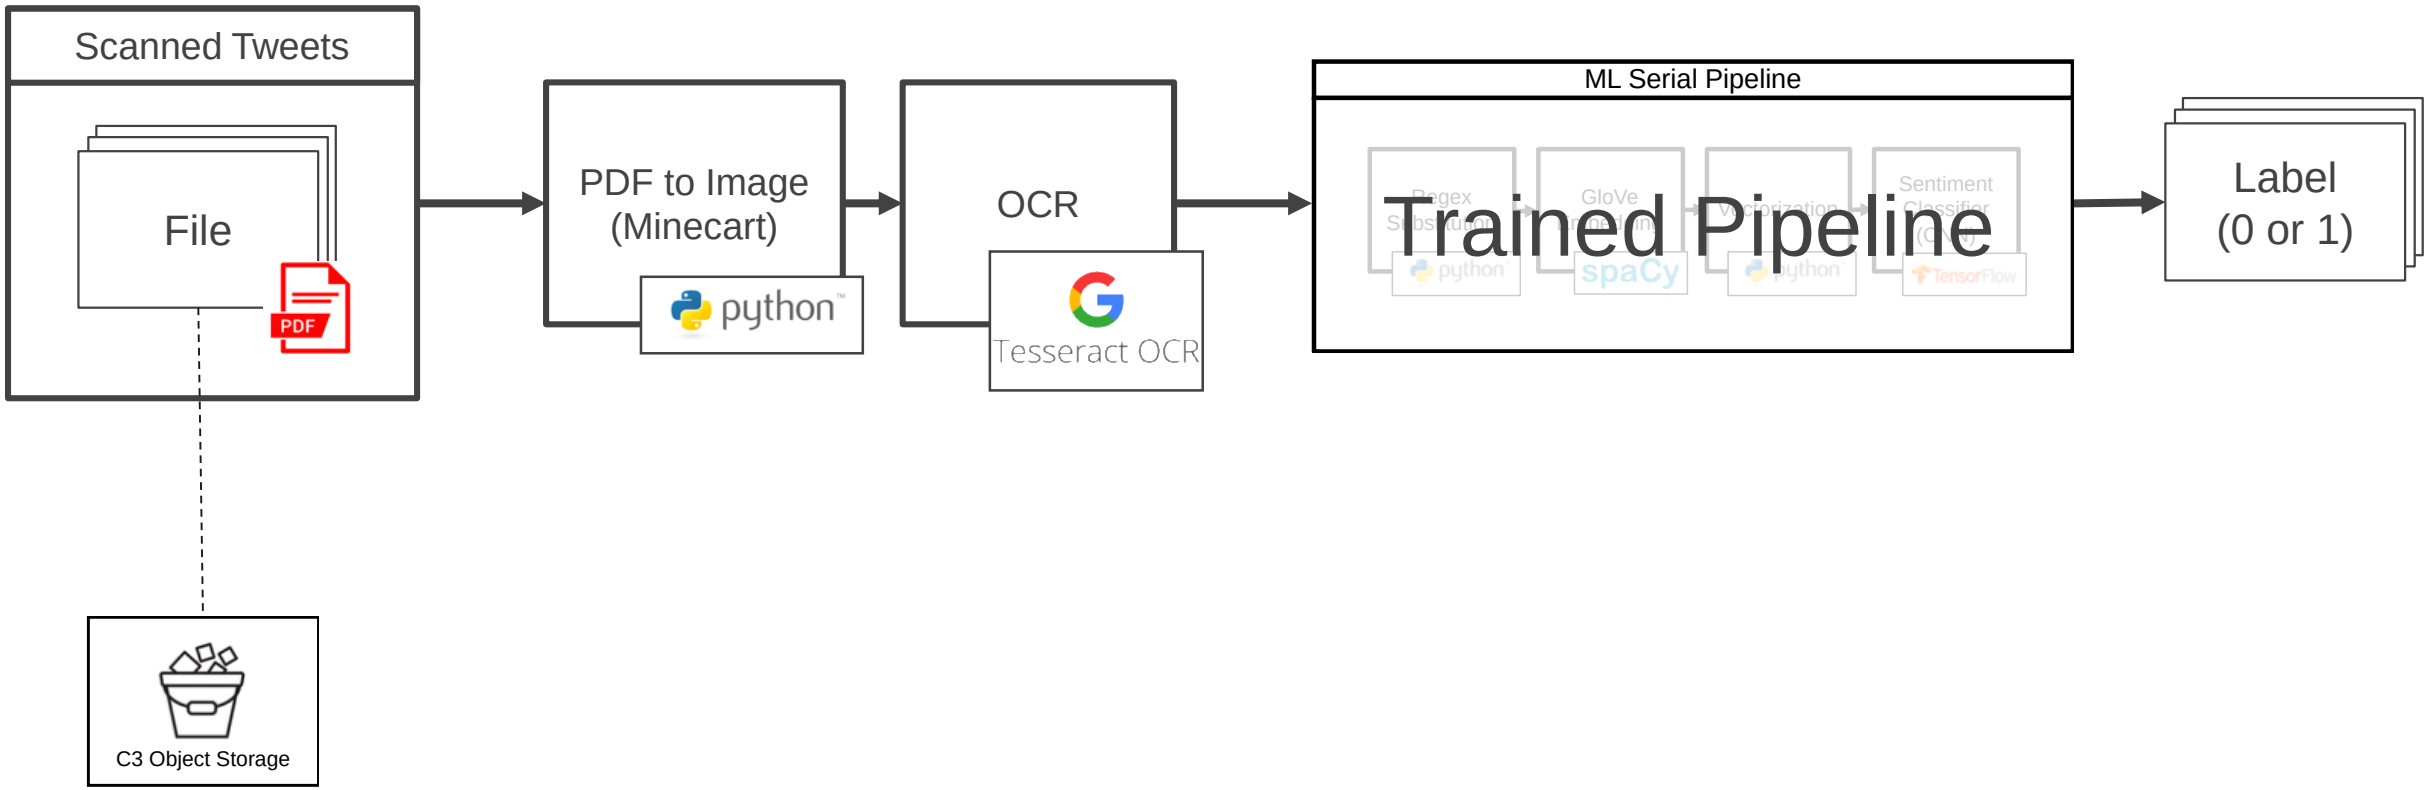

C3.ai

Deep Learning – Word Embedding

- Instead of having one-hot encoding of tokens (potentially with dimension 1e6), map words to low dimensional dense vectors (e.g. dimension 300) with the following constraint:
    - Words with similar meaning (sentiment) should be relatively close in this low dimensional space
- For training words $w_1, w_2, \dots, w_T$ try to maximize:

$$\sum_t \sum_{-c \le j \le c} \log P(w_{t+j}|w_t) \,/T, \text{ where}$$

$$P(w_o|w_I) = \exp(v_o^T v_I) \,/ \sum_w \exp(v_w^T v_I) \text{ (softmax)}$$

- More efficient approaches
    - Hierarchical softmax
    - Negative sampling
- There are some APIs (Word2Vec) and pertained models (Spacy) that we are going to use
- Demo:
    - Cosine of angle between ‘two’ and ‘three’: 0.956
    - Cosine of angle between ‘three’ and ‘tree’: 0.263
    - Angle between ‘man’ and ‘king’ ~ angle between ‘woman’ and ‘queen’

Sentiment Analysis of Images of Tweets -- CNN for NLP

- Chose a CNN approach over a traditional RNN approach due to easier parallelization and computation
  - Uses 1D convolutional layers with filters across sequence dimension
- Similarly to WaveNet architecture, used strides and dilation > 1 in earlier layers to increase the receptive field size of
- each neuron without making the network too deep.
- Limitation – requires fixed input size (needs padding/truncation)

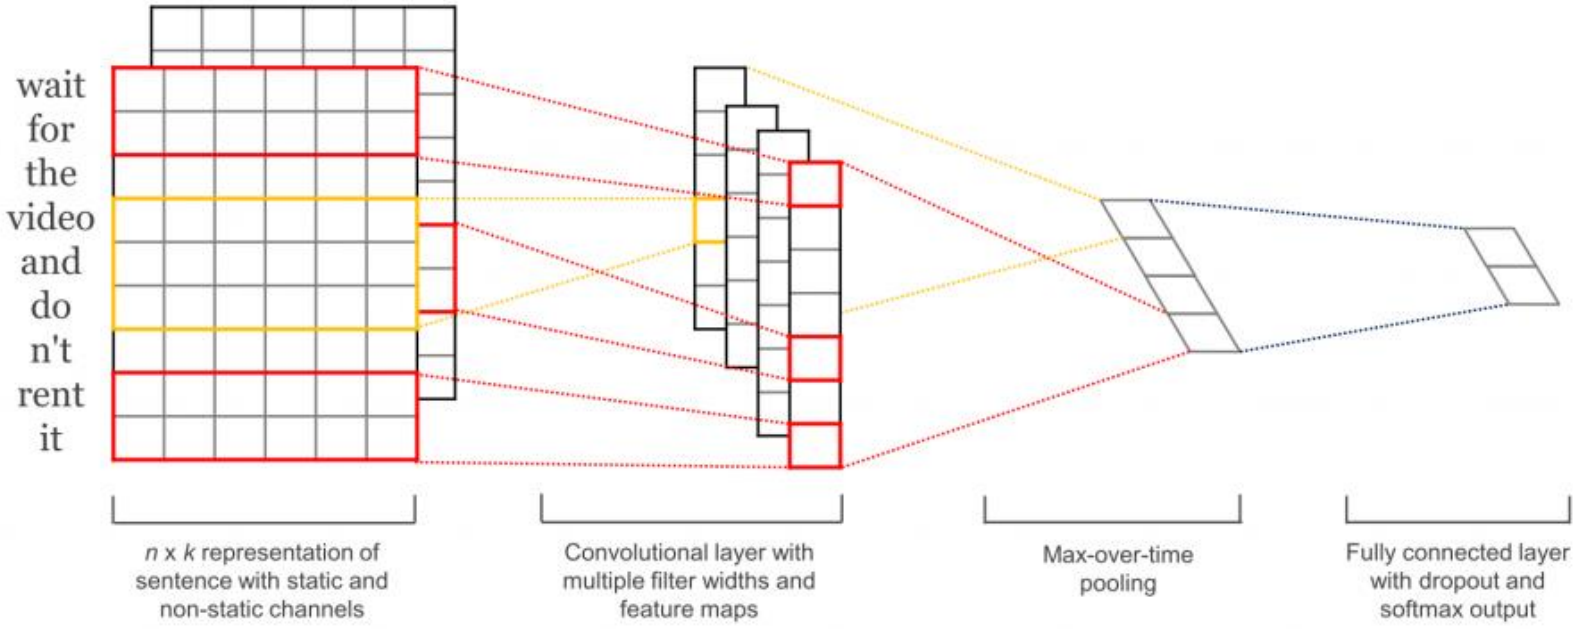

Deep Learning – CNN

- Convert texts to images using their word embedding

kay pollack (the man behind this movie) is a real great man who tries to share his life philosophy in different ways . he has written a bunch of good and well written books about how to control your senses and keep your soul happy . the message in most of his books and this movie , is about that your thoughts in fact is what causes your problems and that the reason of your anger hardly ever is caused of what you think of . the main message is that you can choose to be happy , but hardly ever do that . to watch this movie and learn something very important on life , you have to keep your mind very open and l i s t e n to all the " hidden messages " (or guidelines to get through life) which most of the parts in this movie contains if you listen and watch . watch it with your ears . you wo n't learn the meaning of life , but you 'll learn how to live and get the most out of it ... so , while watching , please keep in mind : " the mind is like a parachute , it does n't work unless it 's open ! "

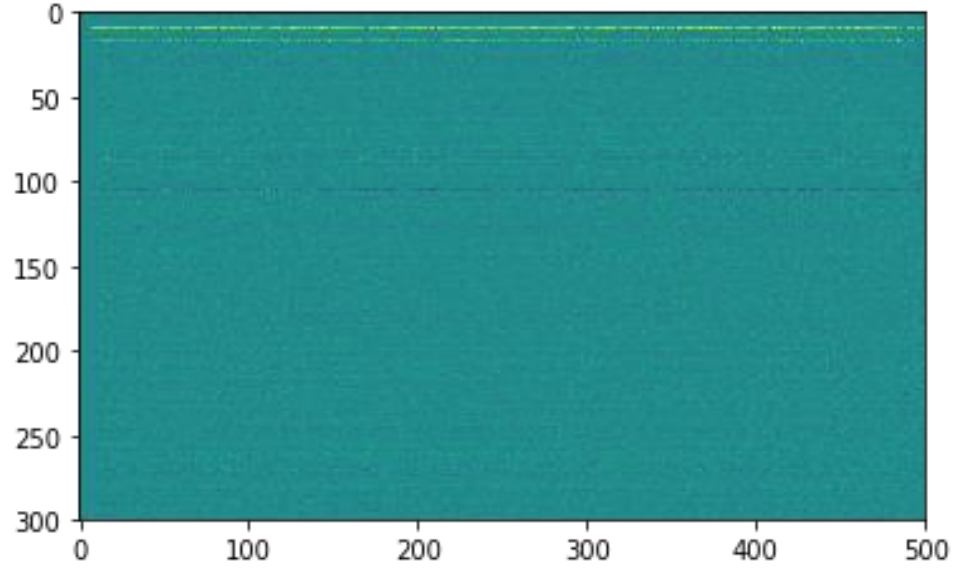

- Applied CNN to the image where the width of the kernel is the same as the image (1-d convolution)
- Model structure:
  - Two layers of Convolution followed by Two layers of dense networks
  - Rectified linear units as activation function
  - Dropout and early stopping to reduce overfitting
  - Cross entropy as the loss function
  - Adam optimizer with default learning rate

CNN – Results (1/2)

- Loss-Accuracy

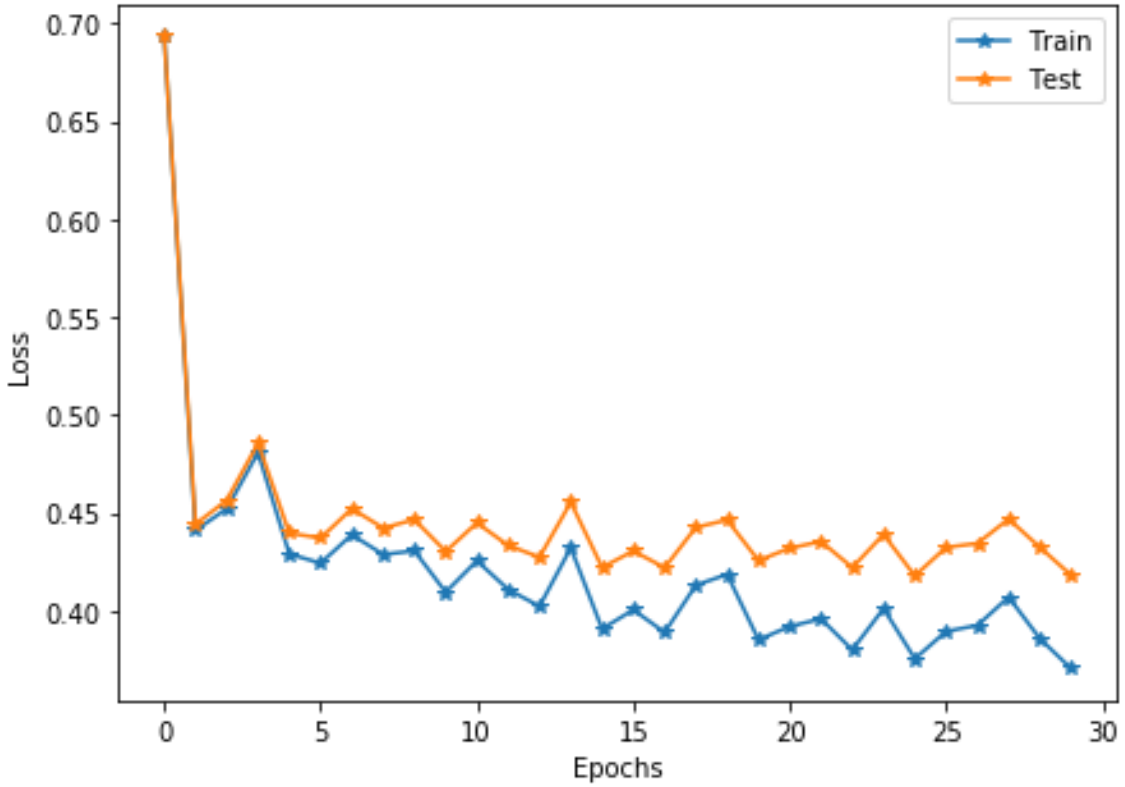

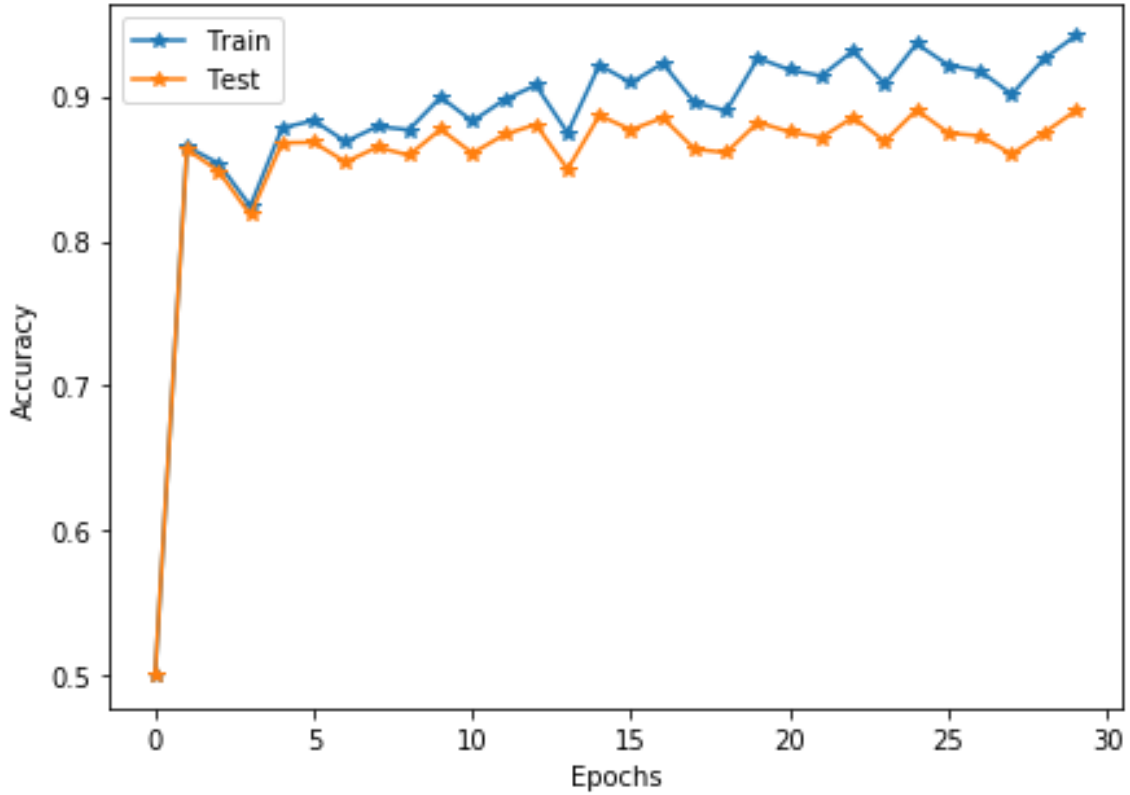

CNN – Results (2/2)

- ROC, Precision-Recall for the selected model

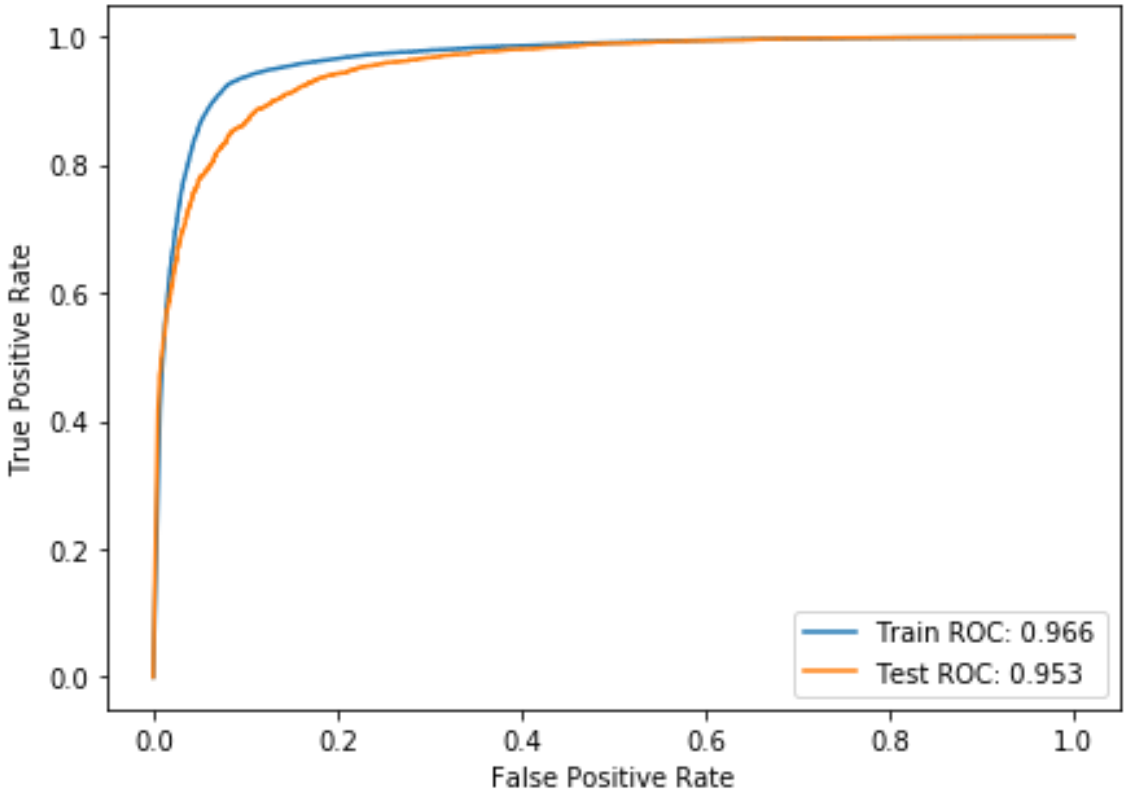

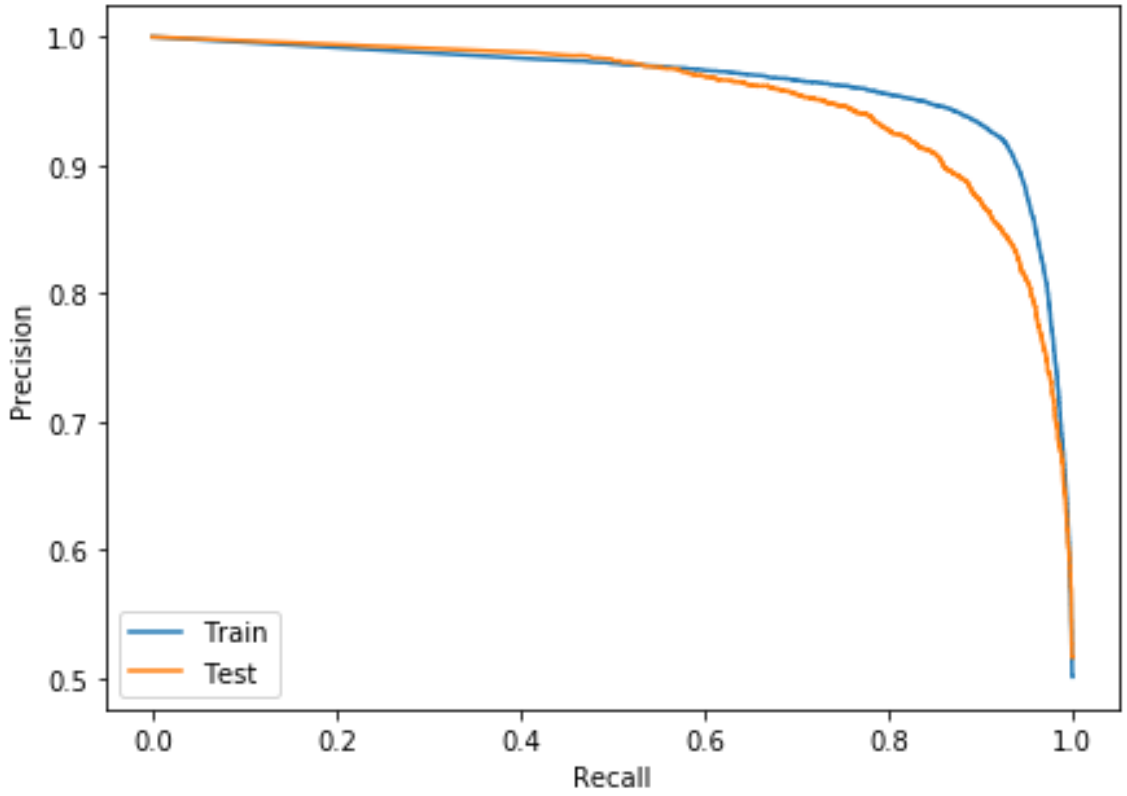

AI for Banking

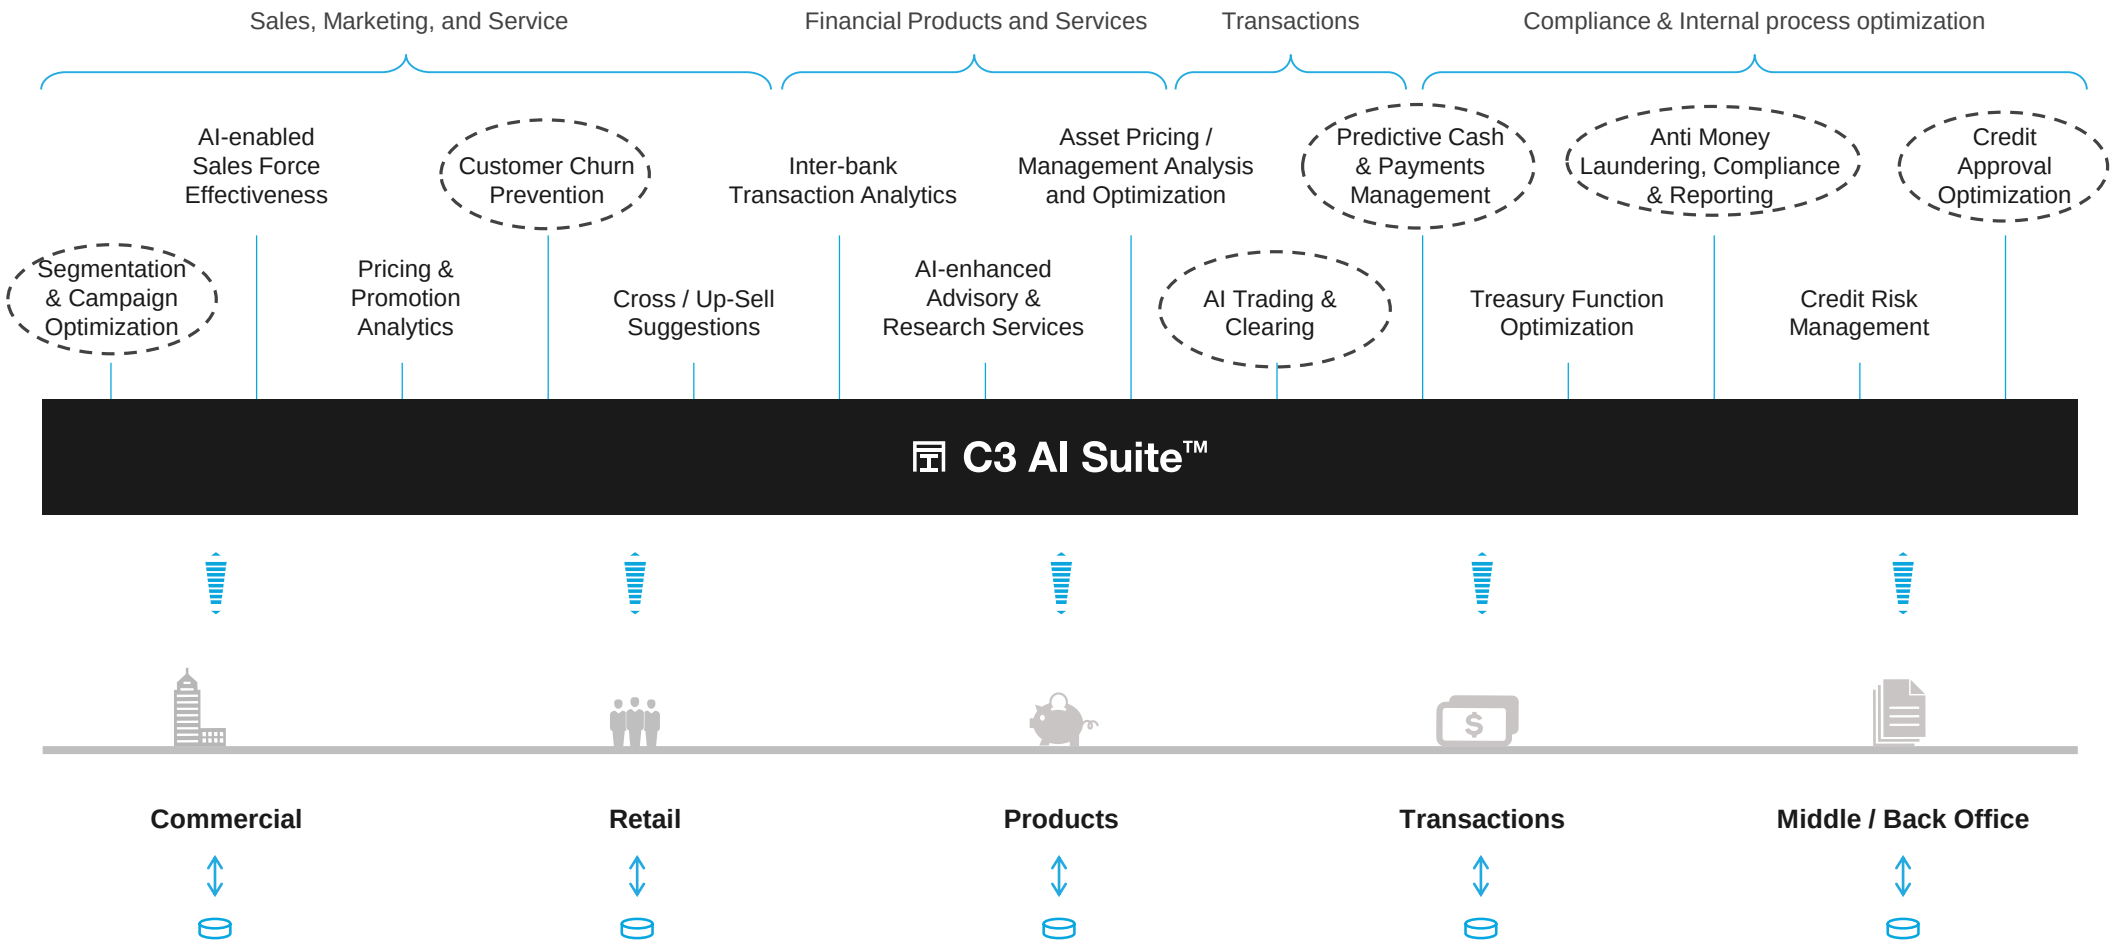

C3.ai

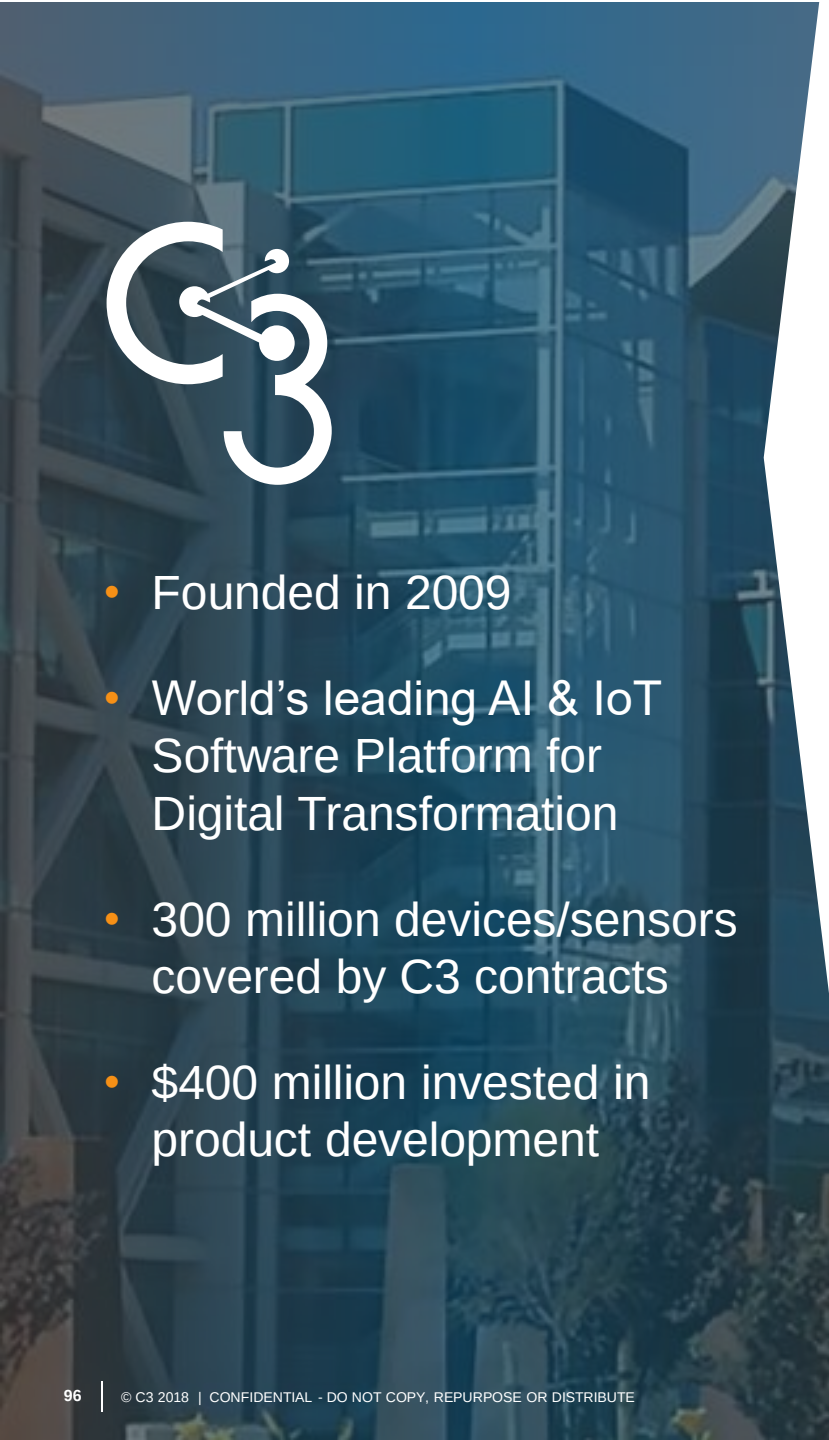

- Founded in 2009
- World's leading AI & IoT Software Platform for Digital Transformation
- 300 million devices/sensors covered by C3 contracts
- $400 million invested in product development

CNBC DISRUPTOR 50 2018

glassdoor 2018 TOP CEOs EMPLOYEES' CHOICE

glassdoor 2017 BEST PLACES to work Employees' Choice

Goldman Sachs 100 MOST INTRIGUING ENTREPRENEURS

JMP HOT 100 SOFTWARE COMPANIES 2018

FORRESTER® WAVE LEADER 2018 Industrial IoT Software Platforms

Forbes 2017 THE Cloud100 WORLD'S BEST CLOUD COMPANIES

Forbes THE 2018 Cloud100 THE WORLD'S BEST CLOUD COMPANIES

datanami DATA SCIENCE • AI • ADVANCED ANALYTICS 2018 EDITORS' CHOICE AWARDS

EY Entrepreneur Of The Year® 2017 Award Winner

constellation RESEARCH

A NextGen Platform for NextGen Apps

C3.ai

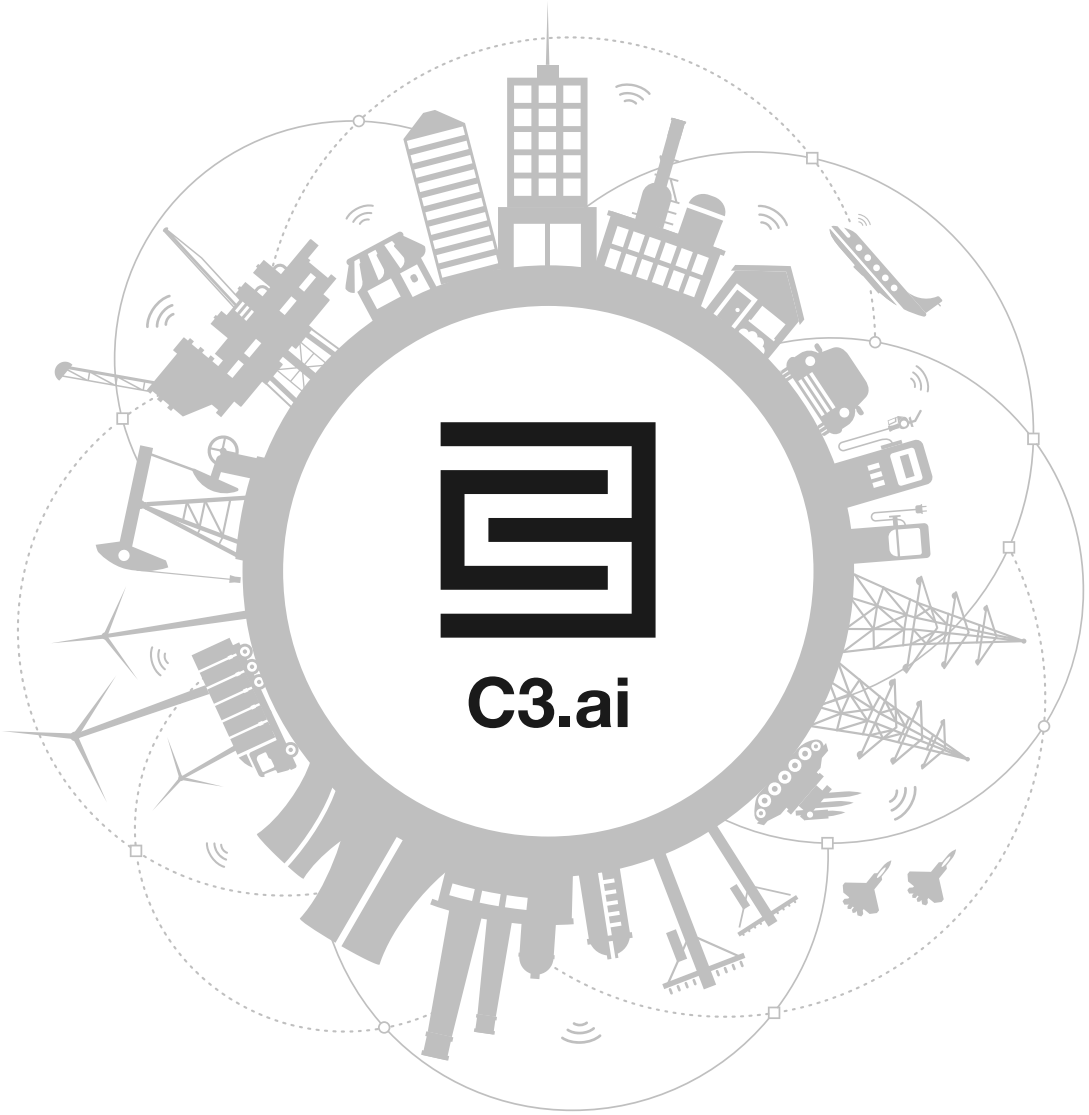
C3.ai