Machine Learning and Deep Learning for IIOT

Chanchal Chatterjee, Dell EMC

Reston, March 22 2016

Goals of the Meeting

- Provide insights on methods and systems for machine learning and deep learning.
- Provide machine/deep learning use cases for IIOT.
- Provide architectures and frameworks for machine/deep learning for IIOT.

Machine Learning & Deep Learning – Confusing, Eh!

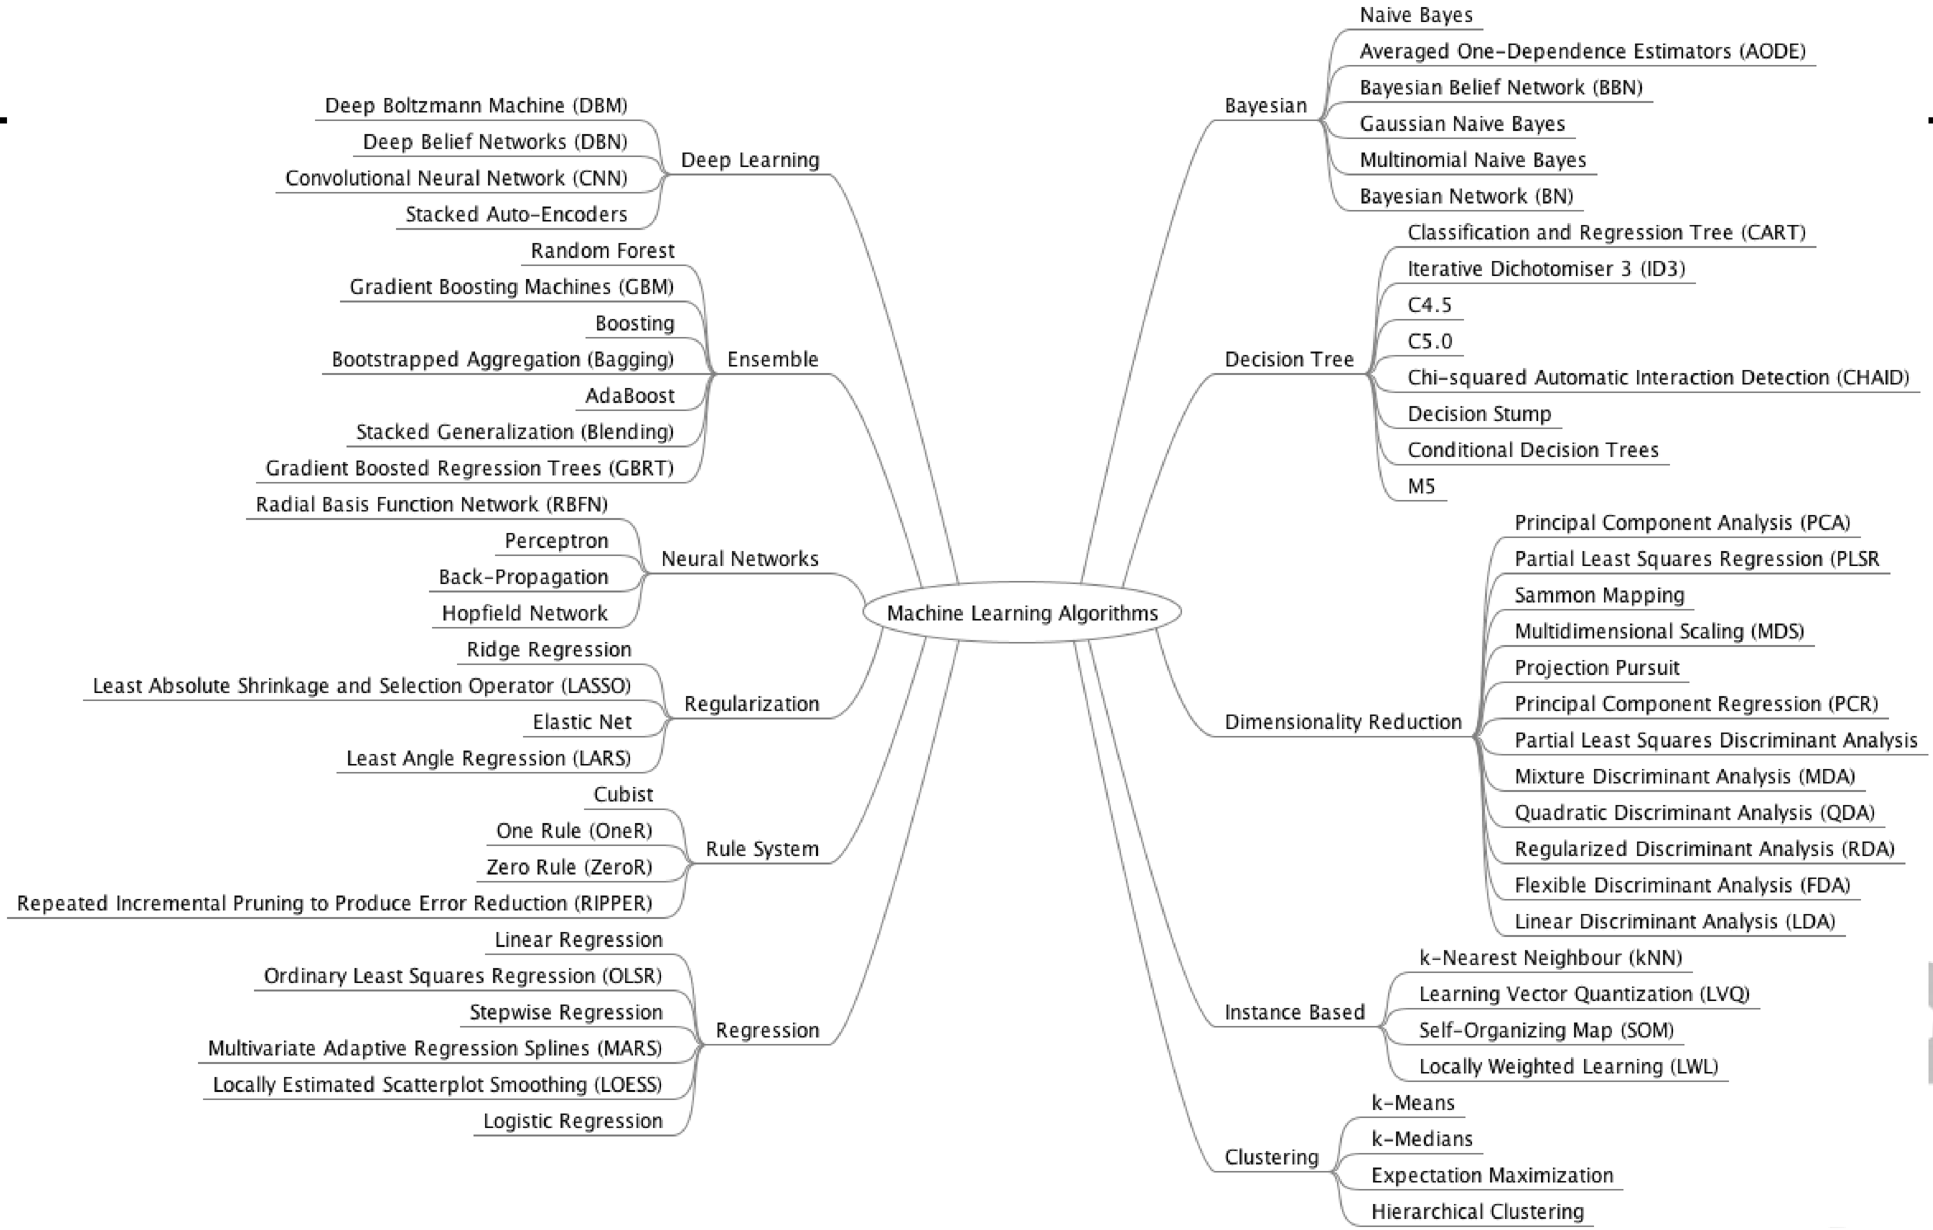

From [Machine Learning Mastery](http://machinelearningmastery.com/) (http://machinelearningmastery.com/)

Machine Learning and Deep Learning Dependencies

- **Types of Data**

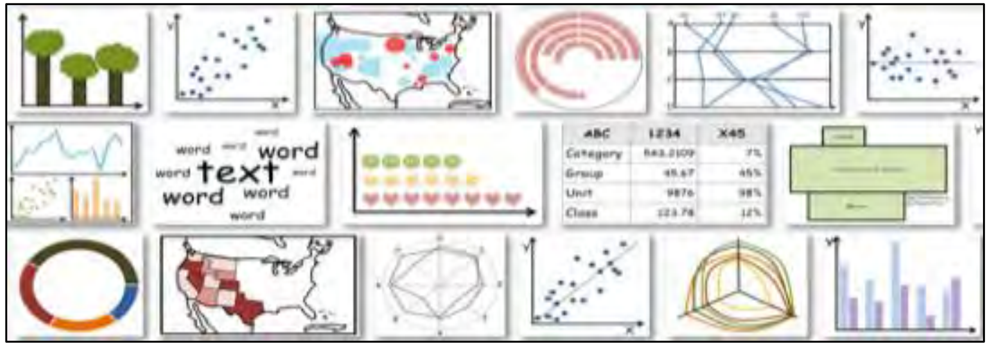

- **Types of Learning**

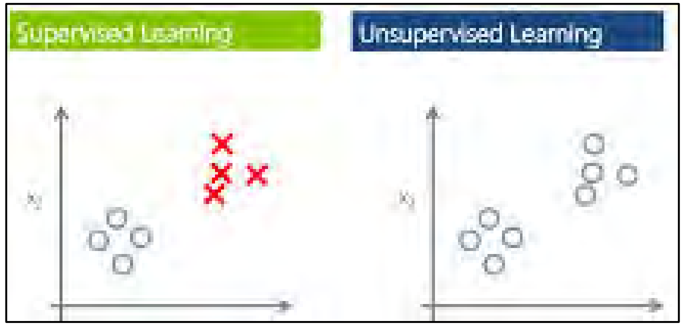

- **Types of Algorithms**

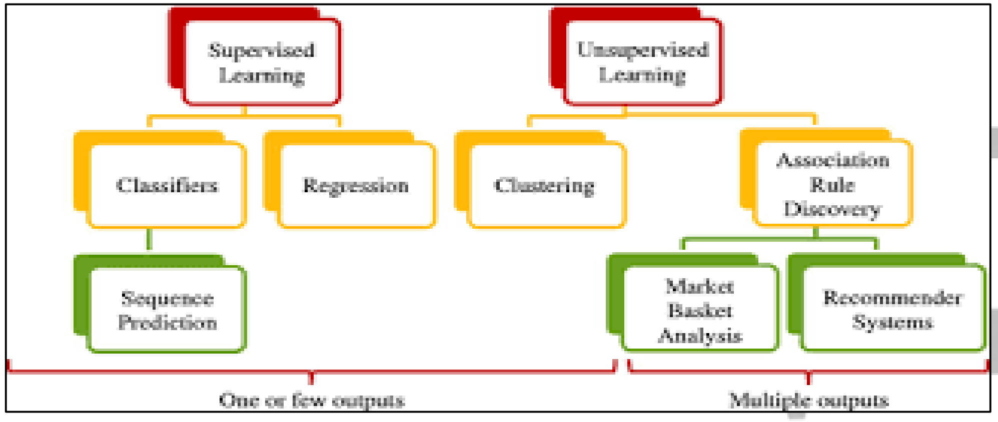

Types of Data

- **Structured Data**
  - Time Series
  - Events
  - Graph
- **Unstructured Data**
  - Video/Images
  - Voice
  - Text

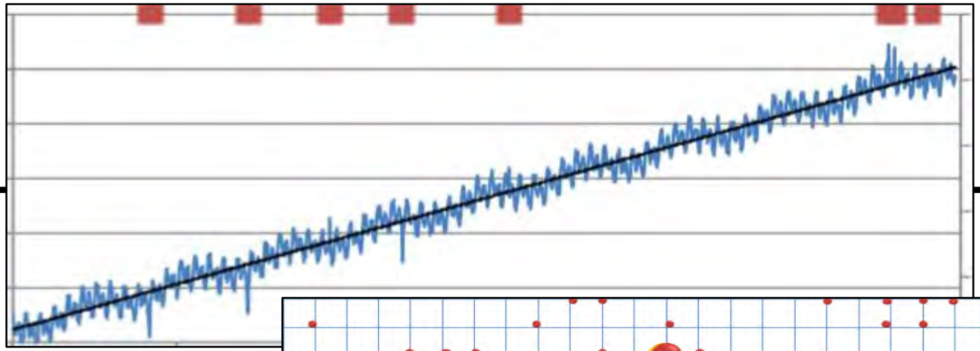
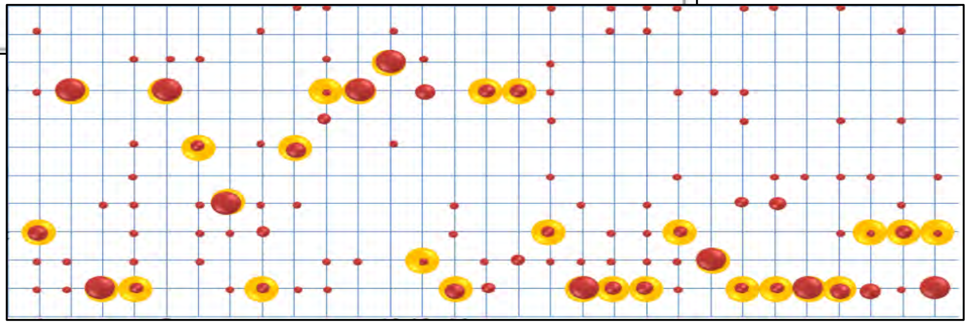
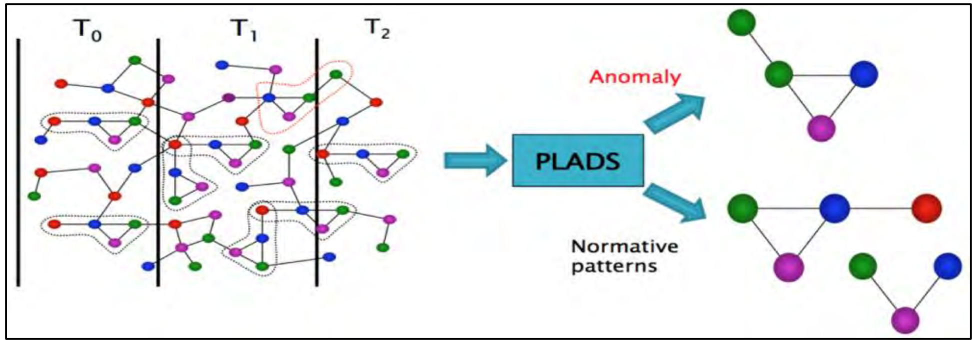

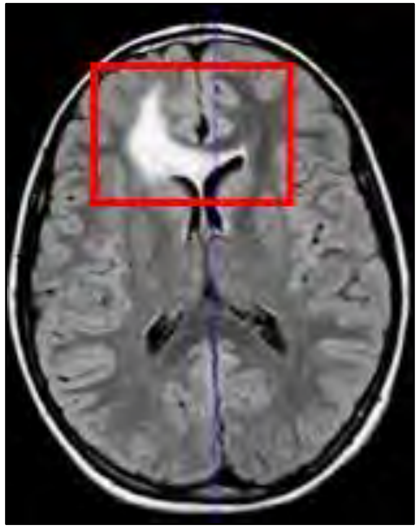
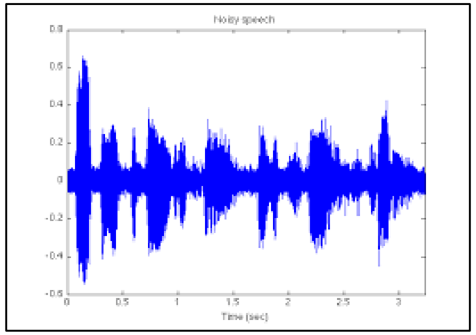

Types of Learning

- **Un-Supervised**
  - Do not require training data
  - Assume normal instances far more frequent than anomalies

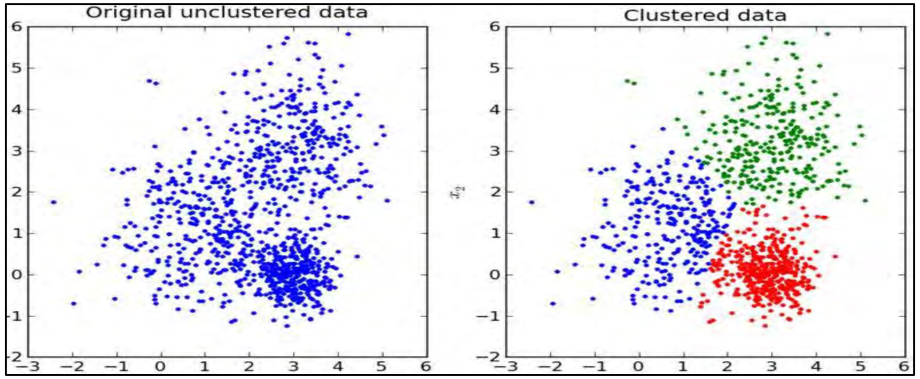

- **Semi-Supervised**
  - Training data has labeled instances for only the normal class
  - Assume normal instances far more frequent than anomalies

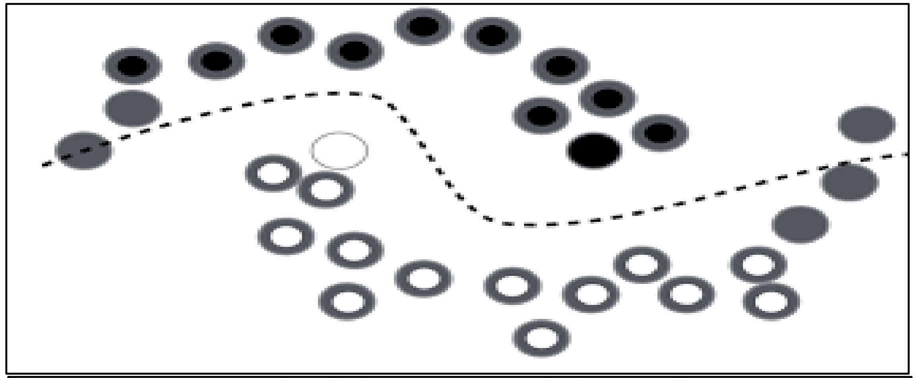

- **Supervised**

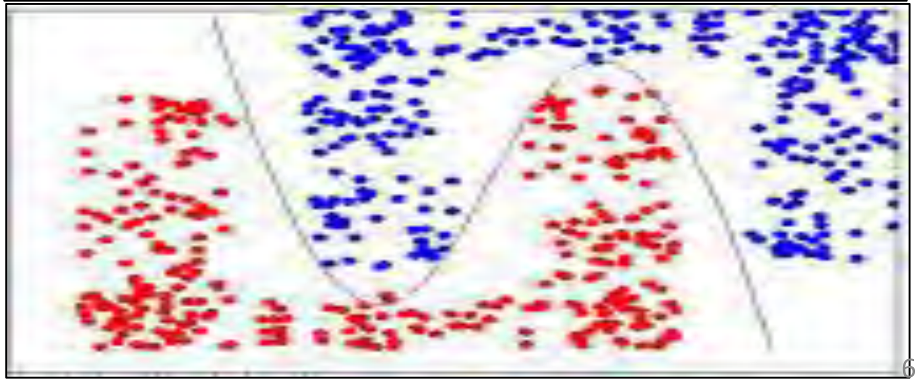

Types of Algorithms

- **ML: Machine Learning**
  - Anomaly Detection
  - Trends, Predictions & Forecasting
  - Association & Grouping
- **DL: Deep Learning**
  - Ladder Network
  - Convolutional Neural Network
  - Recurrent Neural Network
  - Deep Belief Networks

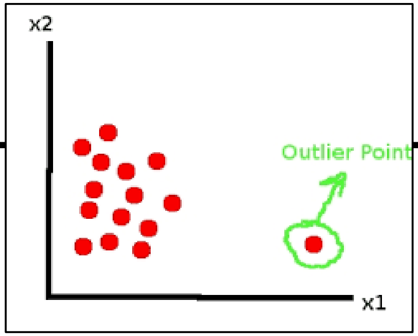

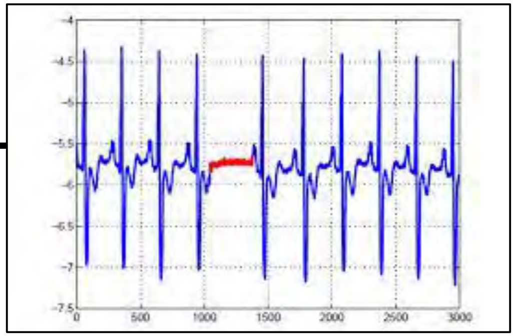
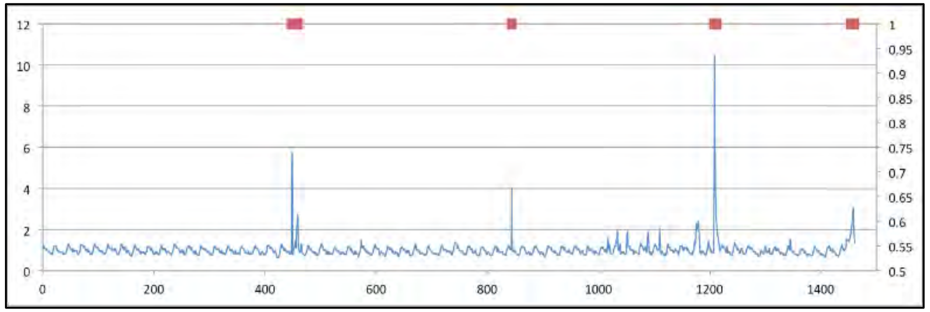
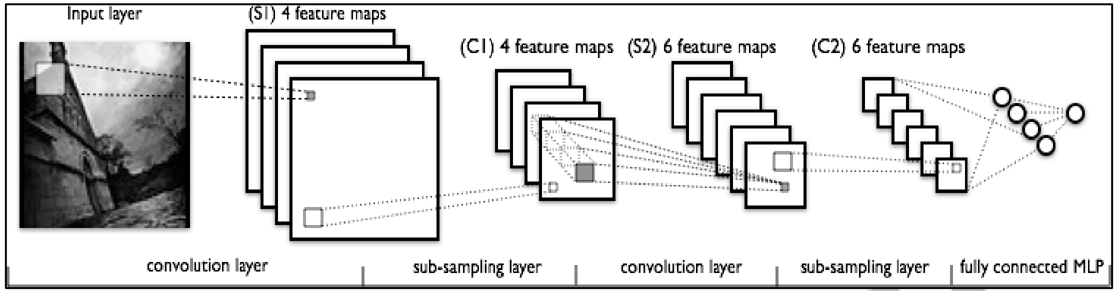

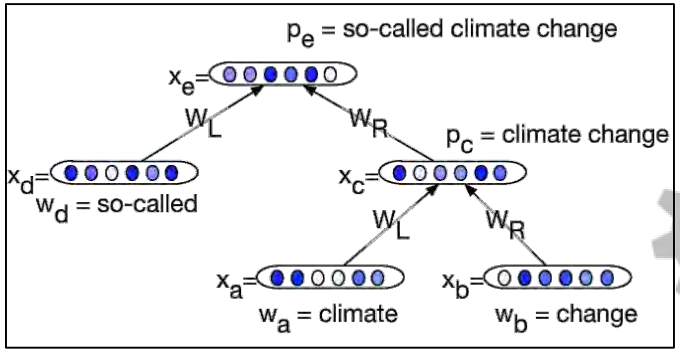

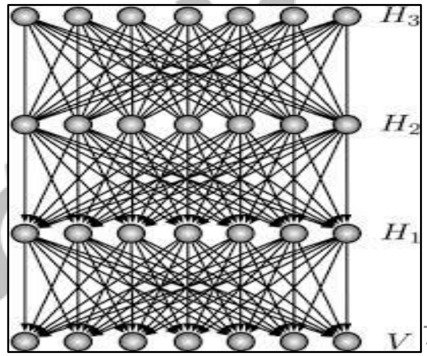

Some Details

Machine Learning

- **Anomaly Detection**
  - Point Anomaly
  - Contextual Anomaly
  - Collective Anomaly
  - Graph Anomaly
- **Trends, Predictions & Forecasting**
- **Associations & Grouping**

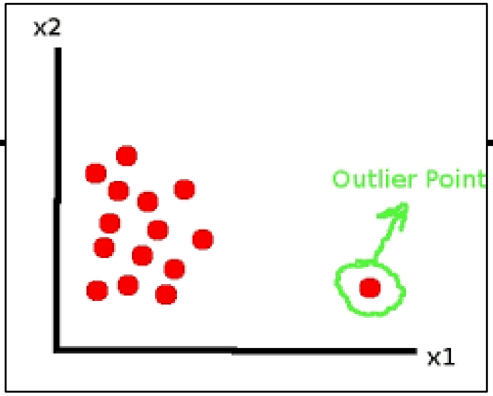

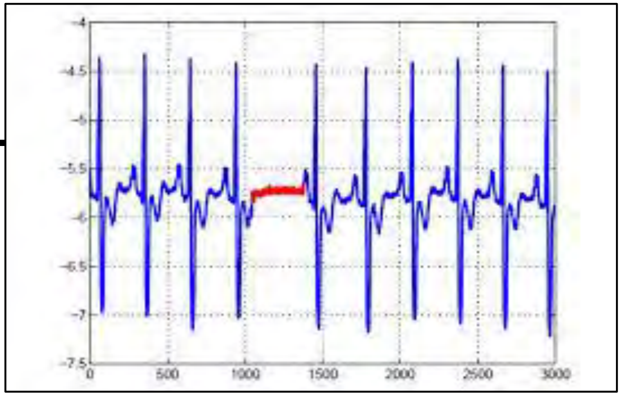

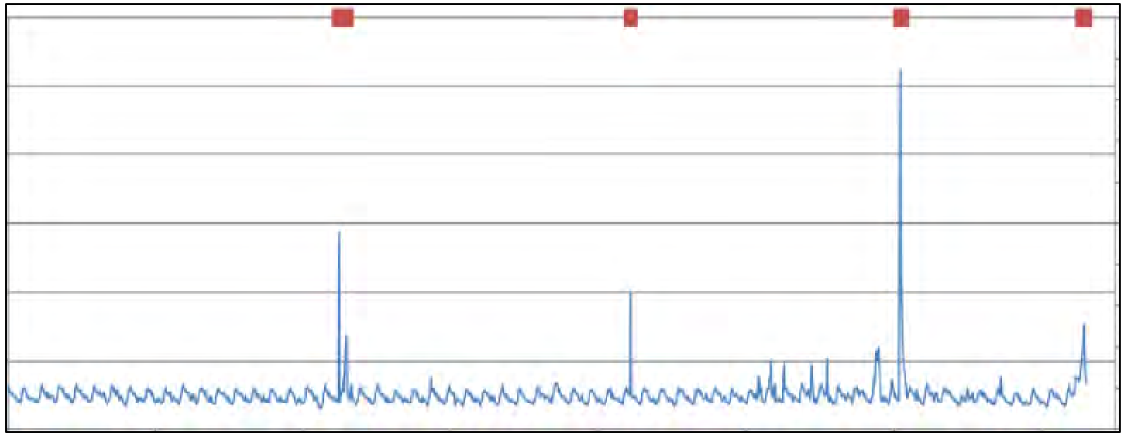

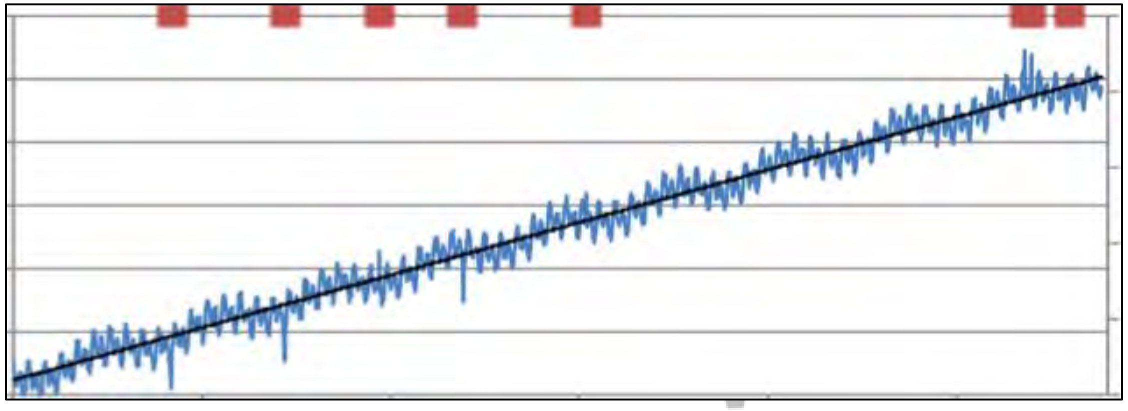

Deep Learning

- **Ladder Network**
- **Convolutional NN (CNN)**
- **Recurrent NN (RNN)**
  - Recurrent Recursive NN (R^2NN)
  - Long Short Term Memory (LSTM)
- **Deep Belief Networks (DBM)**
  - Restricted Boltzmann Machine(RBM)

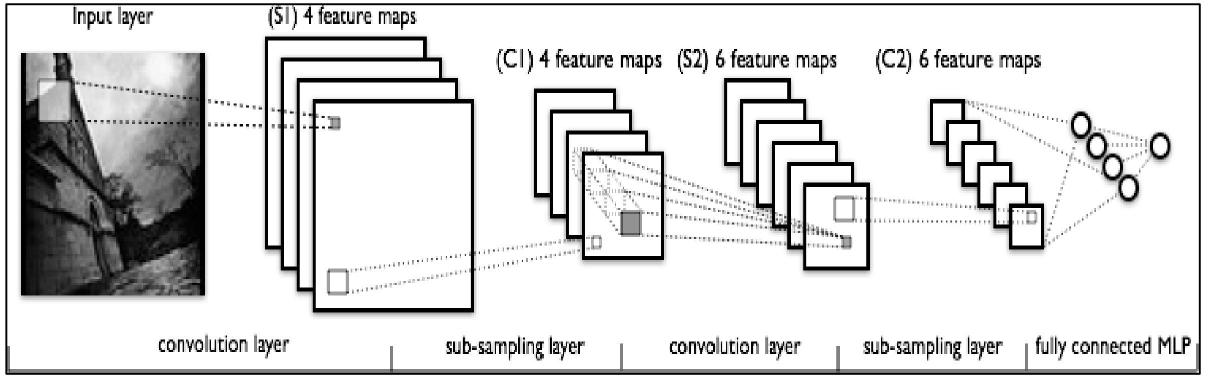

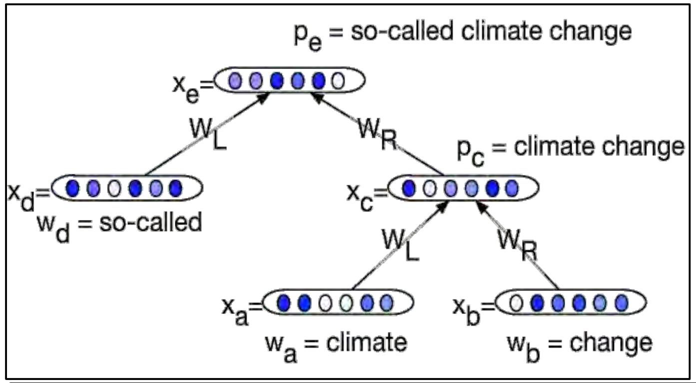

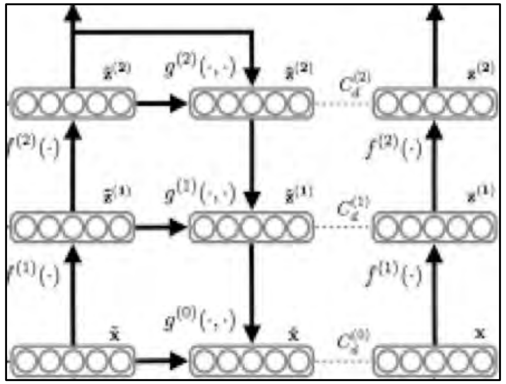

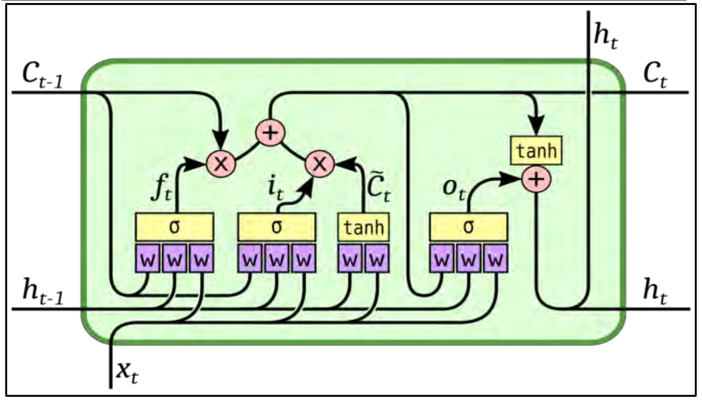

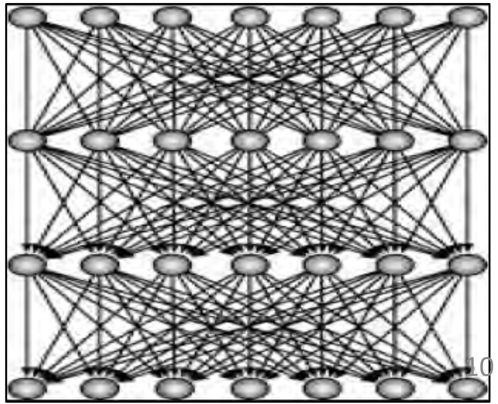

Deep Learning Networks

Source: http://www.asimovinstitute.org/neural-network-zoo/

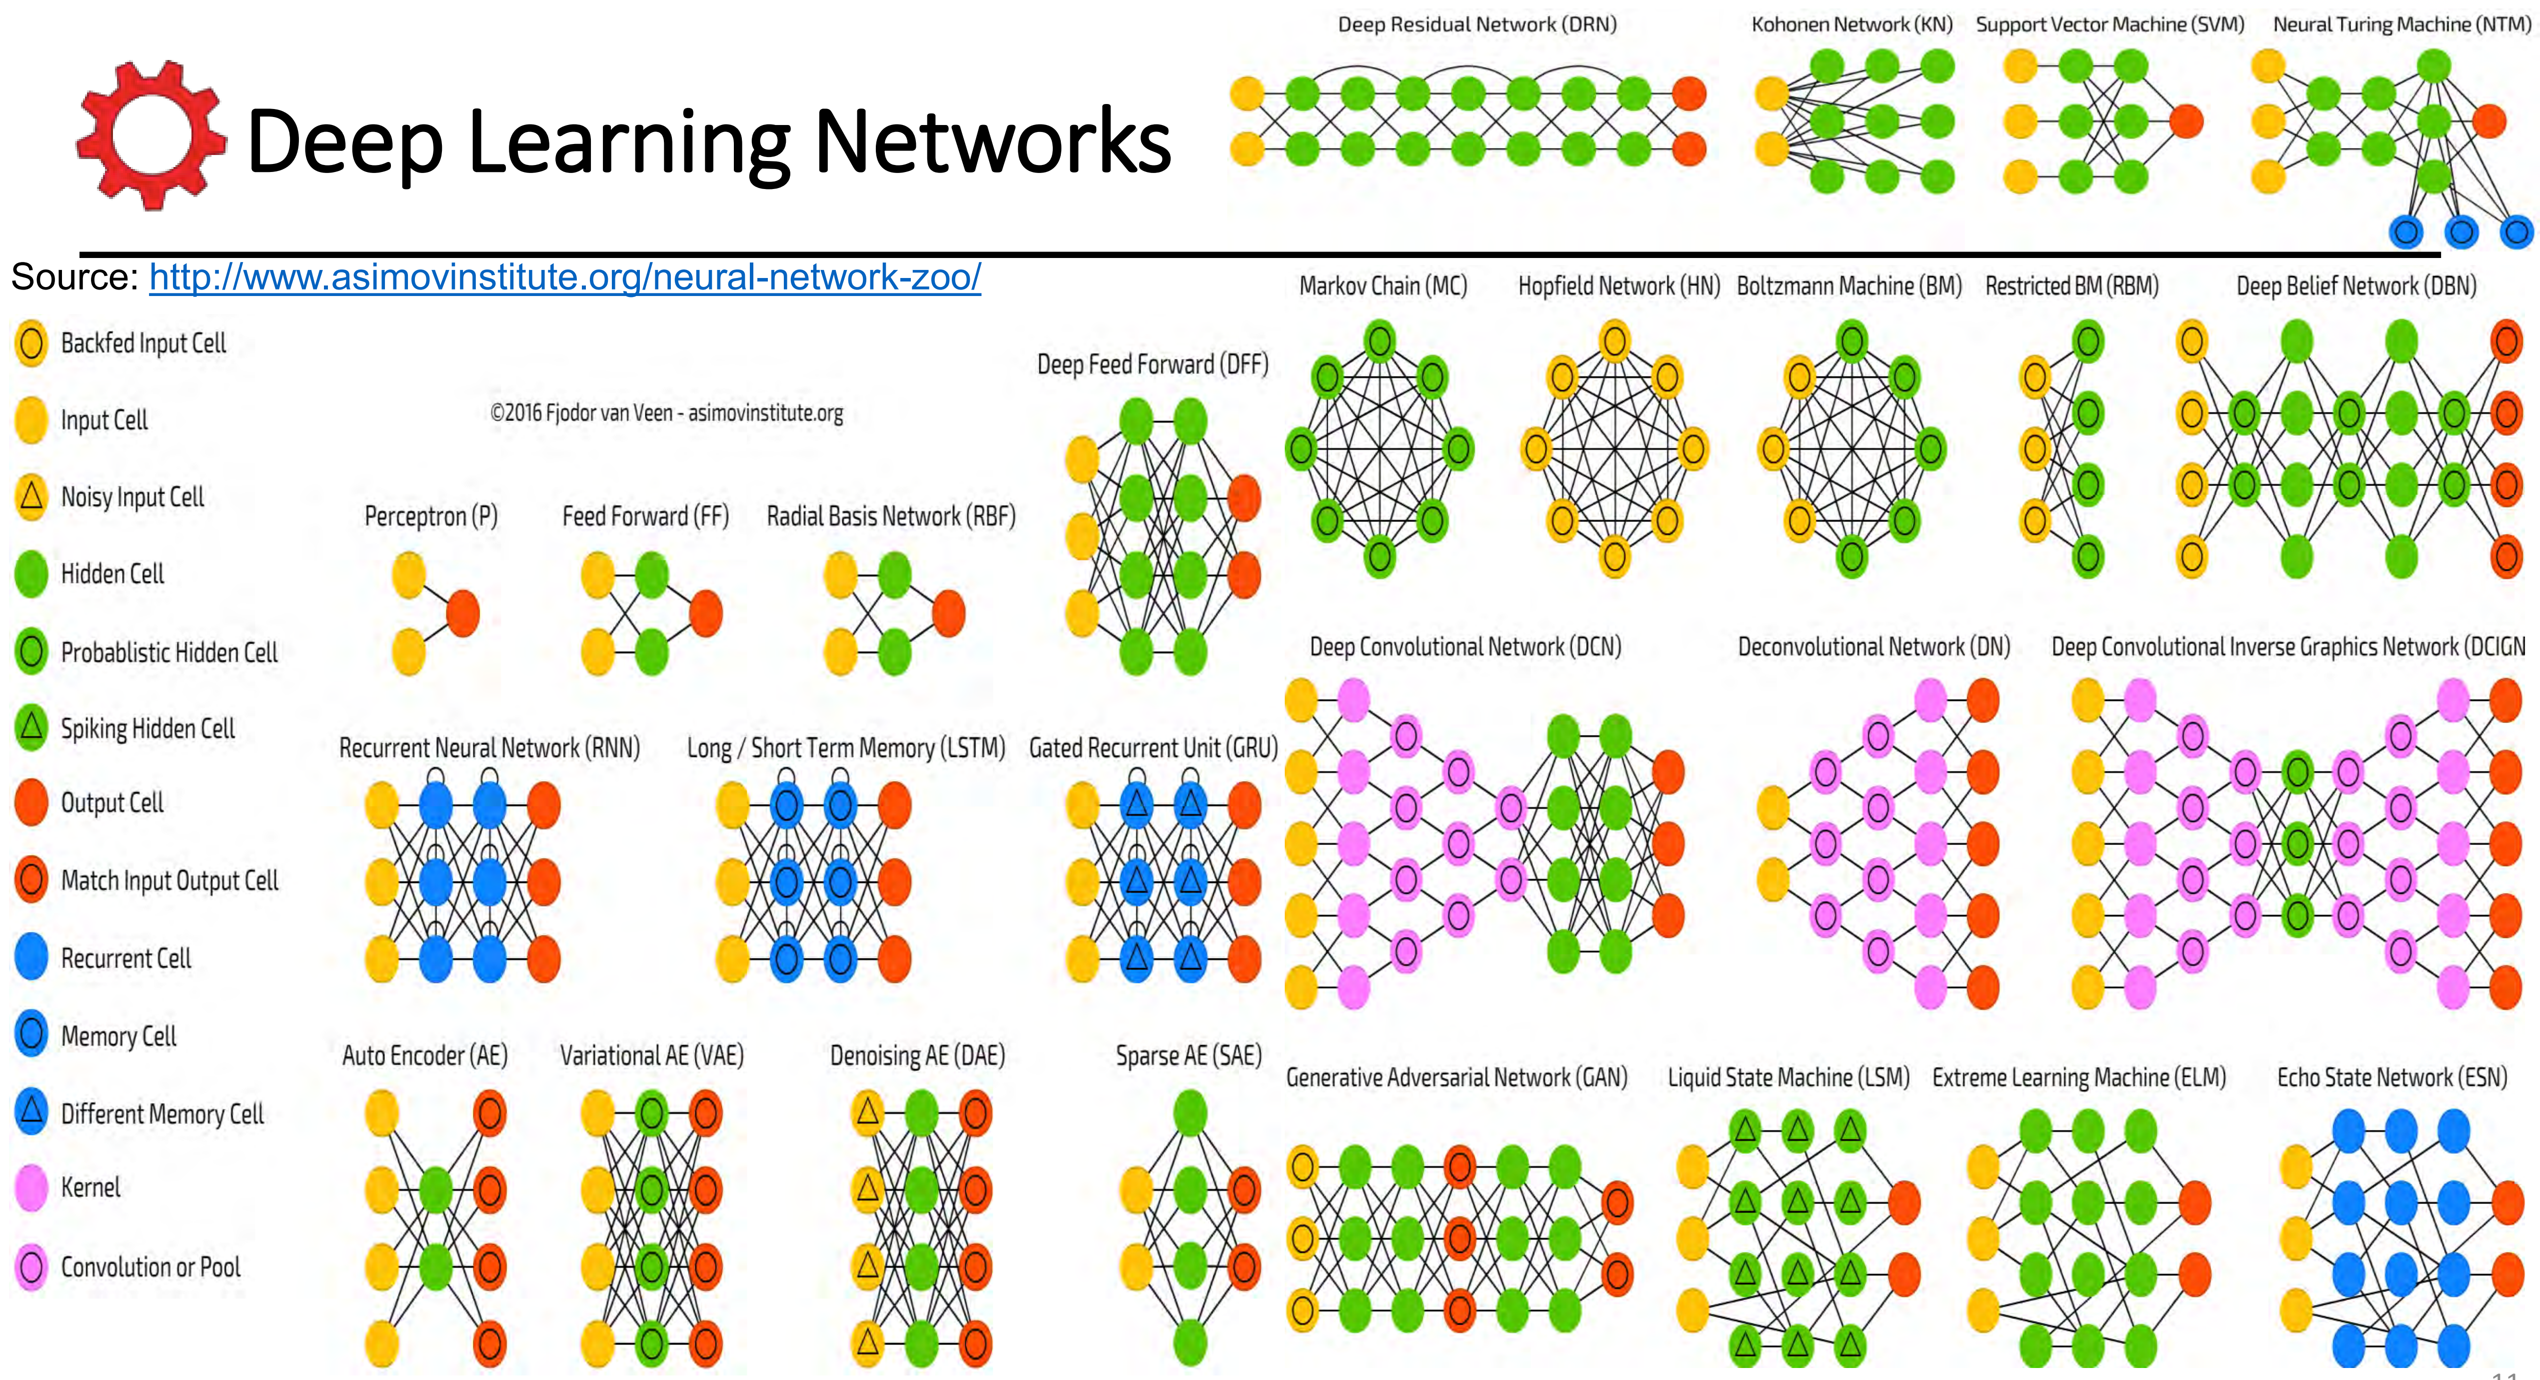

Small Fraction of Real-World ML Systems Have ML Code

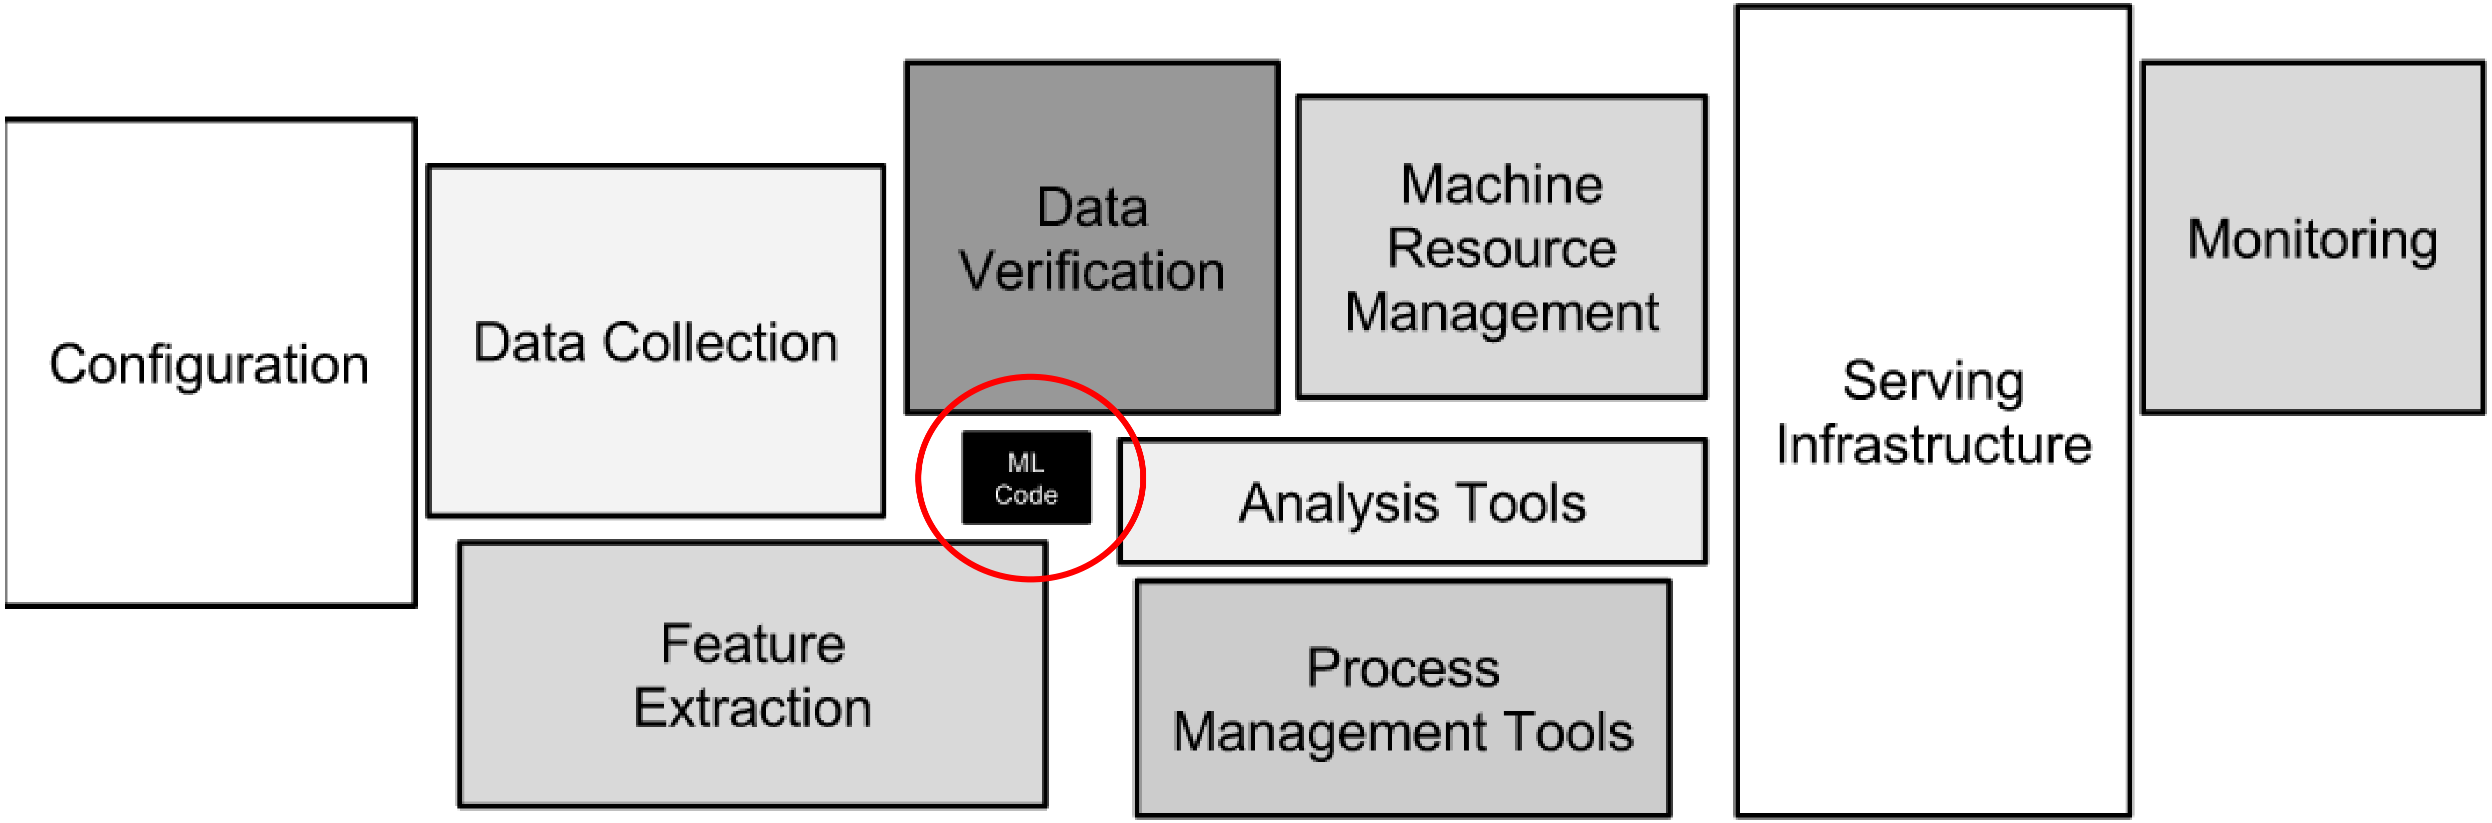

From Hidden Technical Debt in Machine Learning Systems
Scully et al, NIPS 2016

Manufacturing IoT Use Cases

Use Cases	Drivers	Challenges
• Predictive maintenance • Process optimization • Warehouse/supply chain optimization • Remote asset maintenance and control • Product lifecycle monitoring • Integrated plant management • Product-as-a-service	• Increase yield/asset utilization • New revenue streams • Operational efficiencies • Increased worker satisfaction/safety • Eco-sustainability	• Low latencies • Data security • Interoperability between diverse sets of equipment (typically with their own proprietary control system and data interchange standard) • Rapid interpretation of large volumes of data • Reliable indoor/outdoor coverage in harsh environments • Connectivity across different access technologies

Source: SDx Central IoT Infrastructure Report 2017

IIoT Architectures

Three Tier Deployment Model

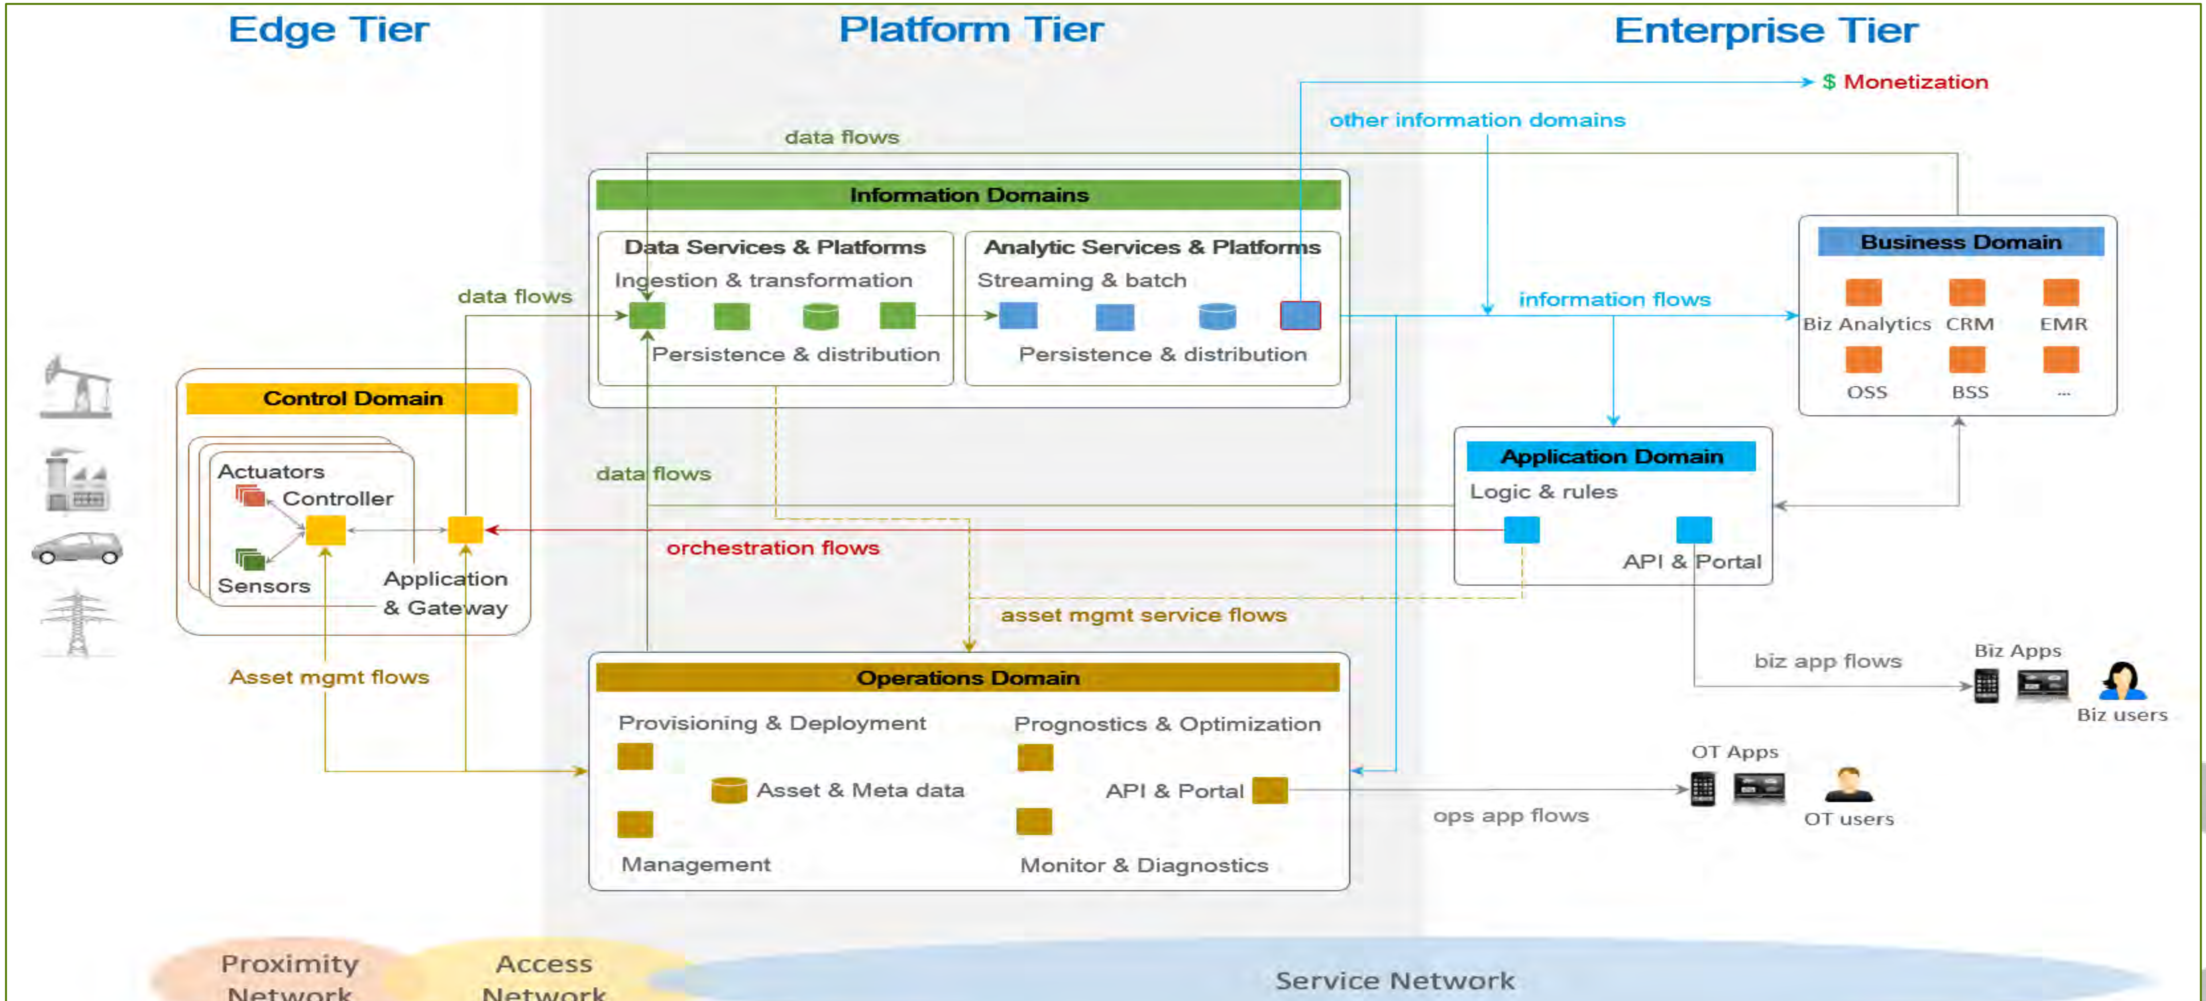

Taken from Industrial Internet Reference Architecture v1.8

Edge2Core Crosscutting Functions

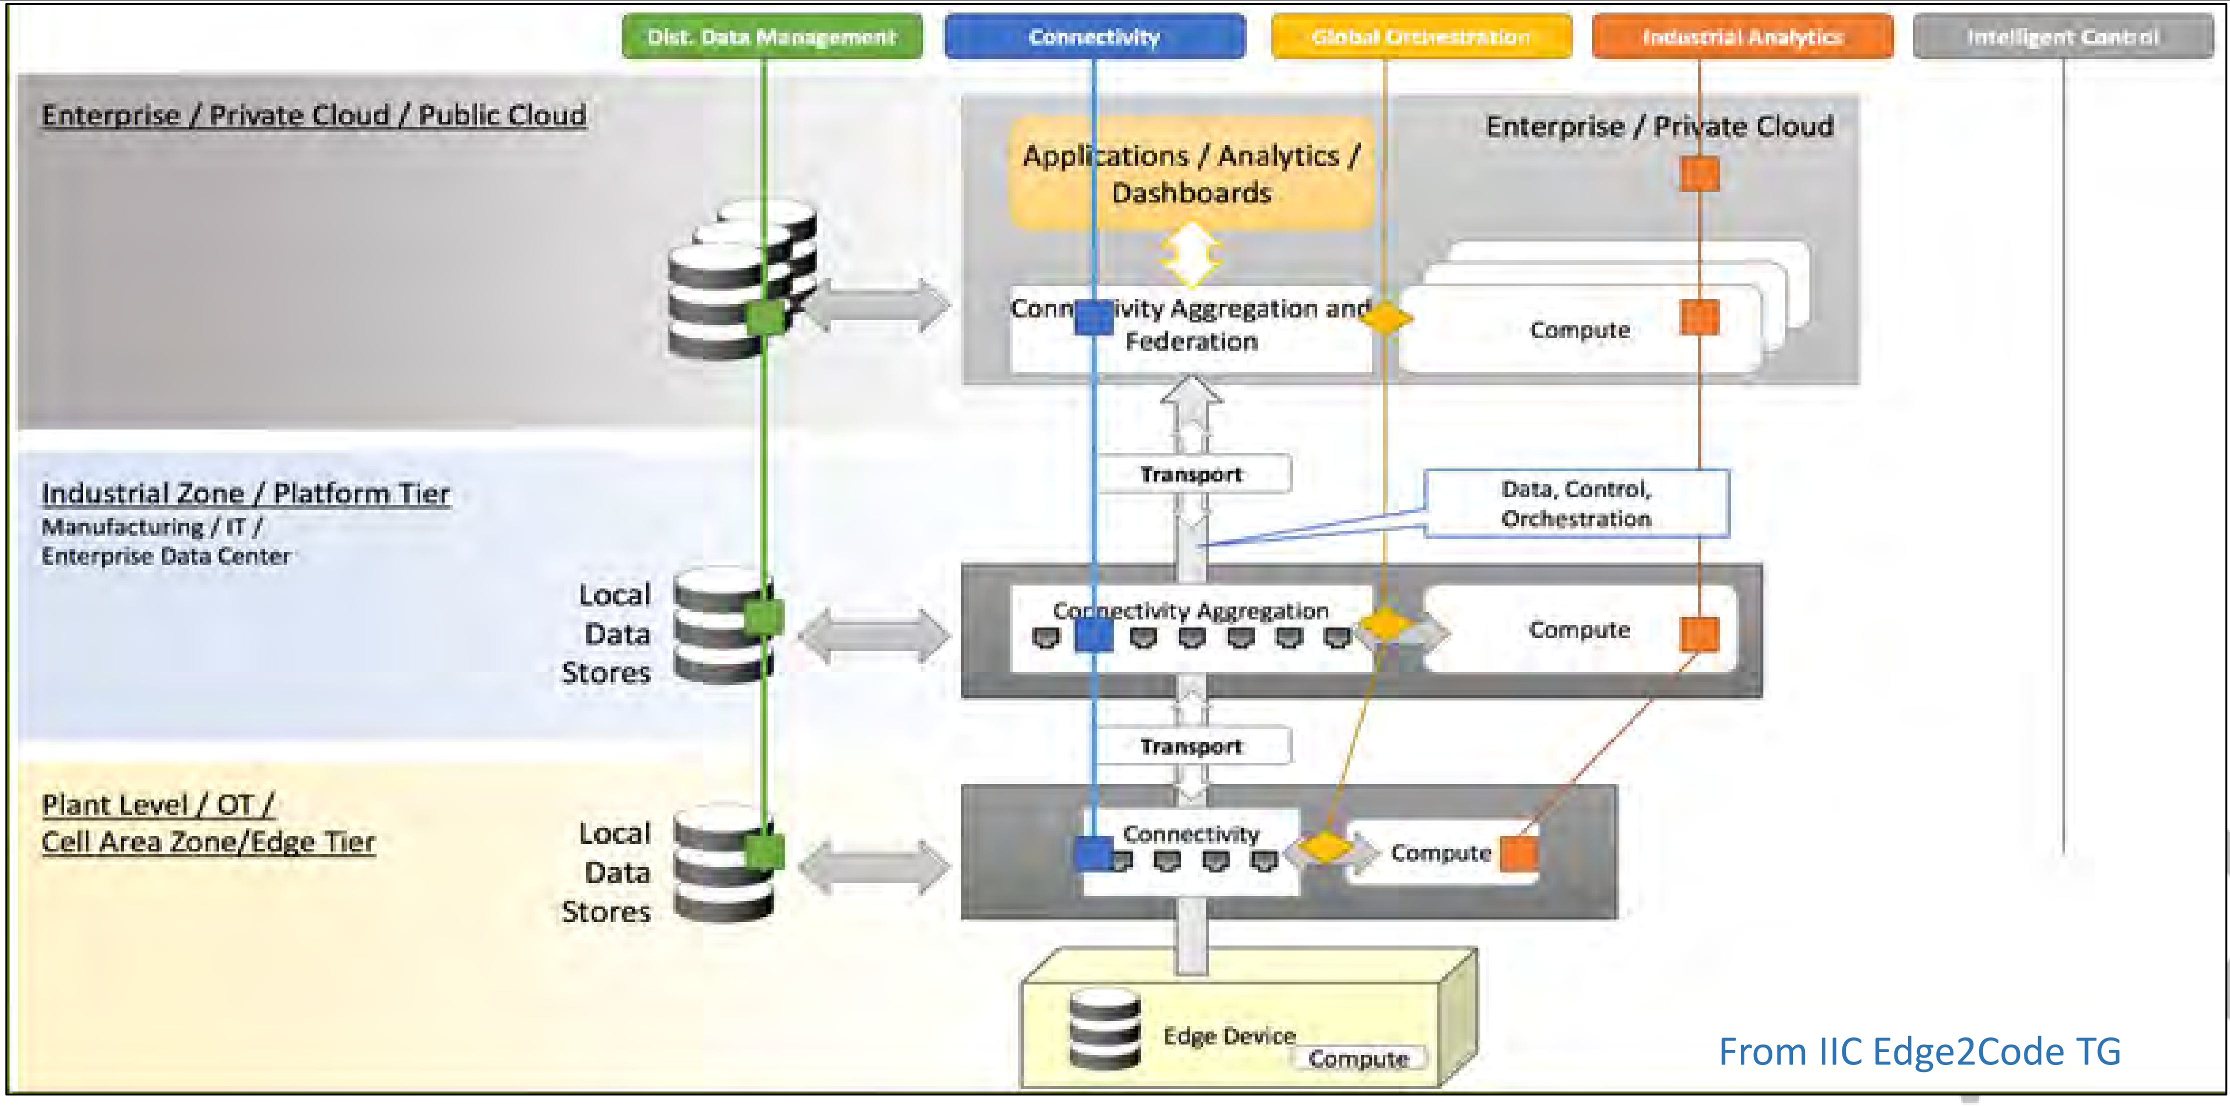

From IIC Edge2Code TG

IIOT Analytics Implementation Challenges

- How to Combine Streaming and Batch Processing Frameworks

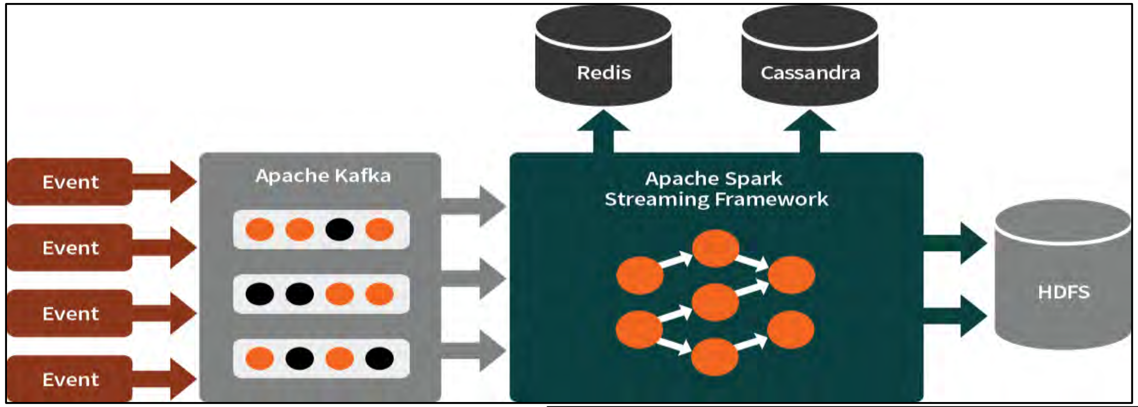

- How to Introduce Human Domain Knowledge
  - NLP, Reinforced Learn, …

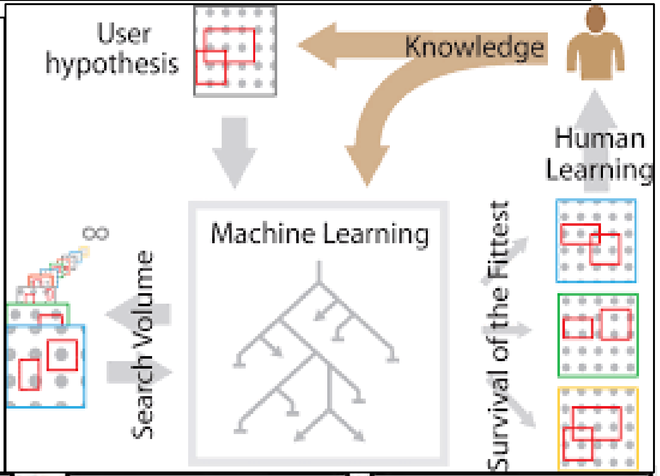

- How to Distribute Processing and Data at the Tiers

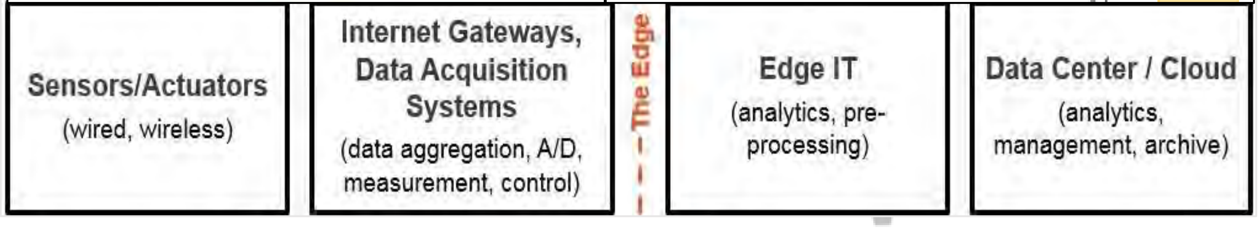

IIOT Analytics Implementation Challenges

- How to train DL with Unlabeled data using Algorithms & Domain knowledge

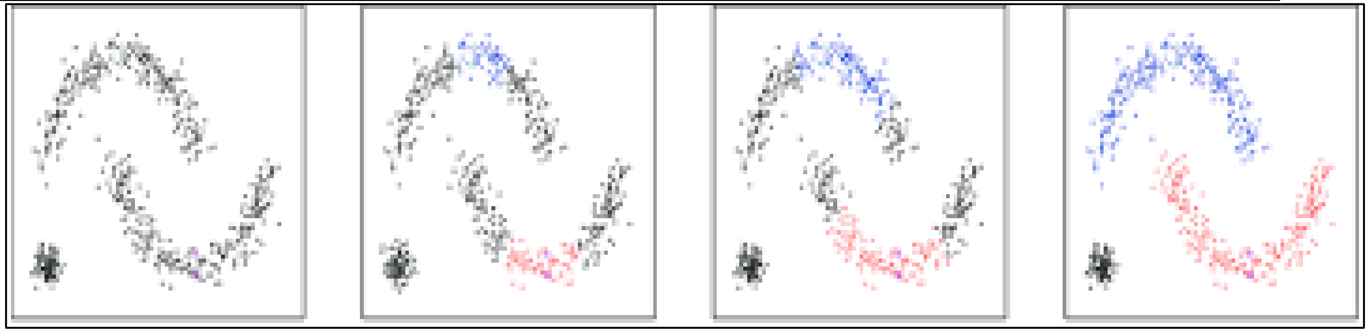

- How to scale DL into multiple nodes

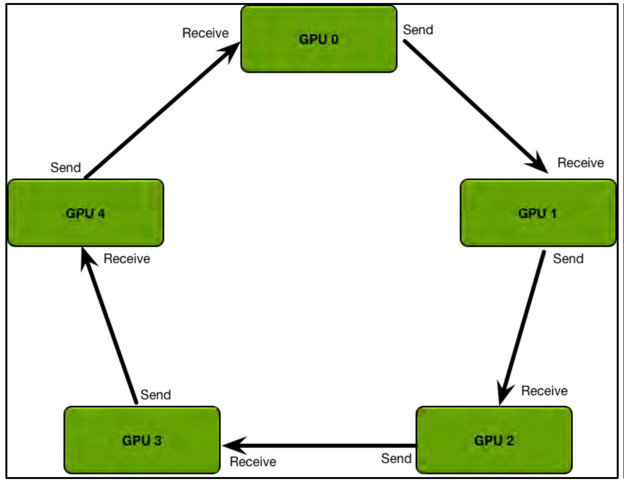

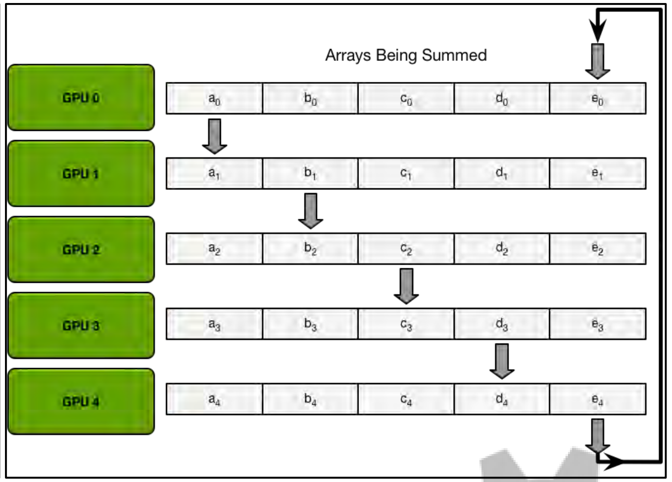

- How to tune DL Nws
    - architecture, parameters

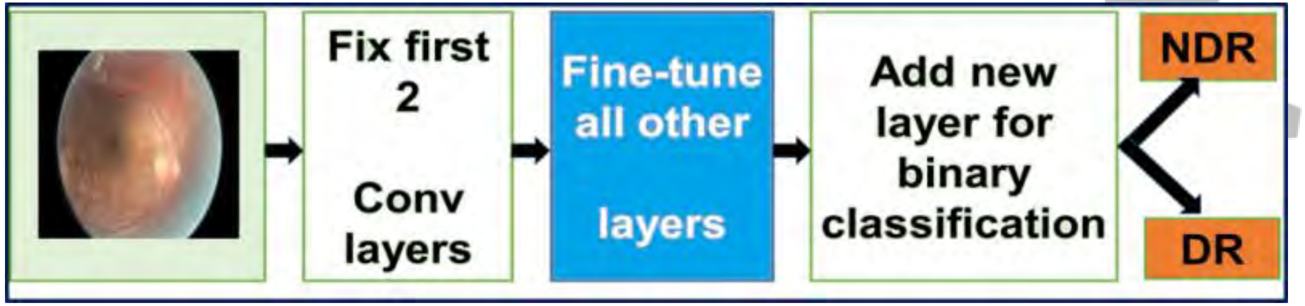

IIOT End to End Frameworks

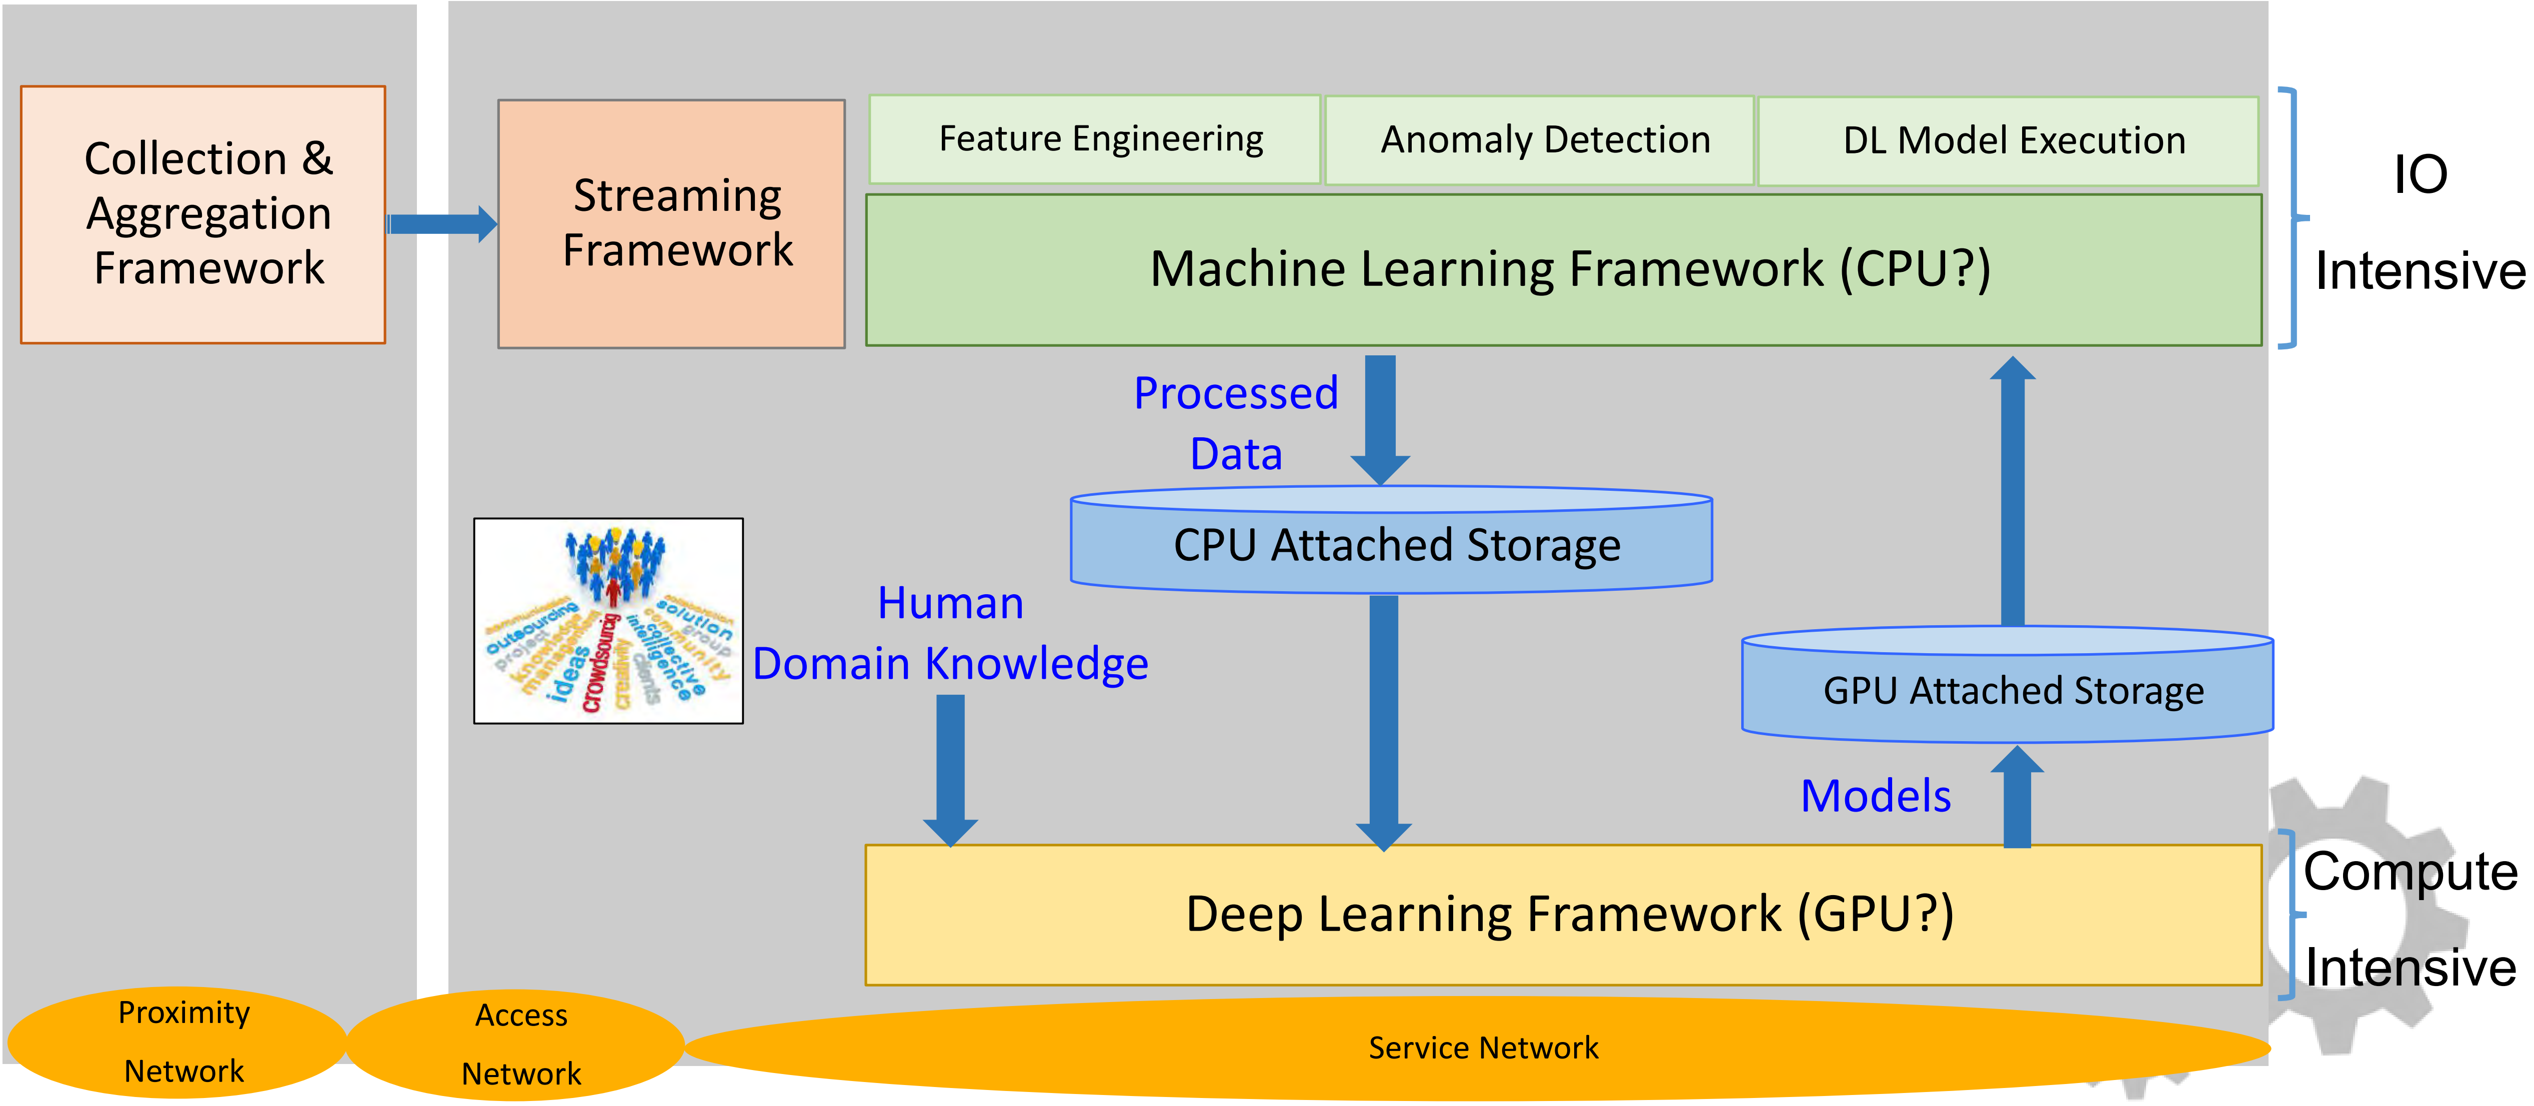

Open Source Frameworks for ML and DL

Deep Learning Frameworks

- Apache SINGA
- Brainstorm
- Caffe
- Chainer
- CNTK (Microsoft)
- DL4J
- DMLC
- Fbcunn (Facebook)
- Lasagne

- Minerva
- Mocha.jl (Julia)
- MXnet
- Neon (Nervana)
- Purine
- Tensorflow (Google)
- Theano
- Torch
- Warp-CTC (Baidu)

- Brain (Javascript)
- Cudamat
- Deep Learning Framework (Intel)
- Deepnet
- Hebel
- Infer.NET
- Keras
- Leaf

- MLPNeuralNet
- Neural Network Toolbox (MatLab)
- Neuraltalk
- Neurolab
- OpenDeep
- PyBrain
- Swift-AI
- VELES (Samsung)

Each differ on – Licensing, Language implemented, OpenMP Support, OpenCL support, CUDA support, Various networks implemented, Pretrained model support and parallel implementations

Deep Learning Frameworks

Software	Software license	Platform	OpenMP support	OpenCL support	CUDA support	Recurrent nets	Convolutional nets	RBM/DBNs
TensorFlow	Apache 2.0	Linux, Mac OS X, Windows	No	On roadmap	Yes	Yes	Yes	Yes
Caffe	BSD 2-Clause License	Linux, Mac OS X, unoffl Android, Windows	No	3rd party implementation	Yes	Yes	Yes	No
Keras	MIT license	Linux, Mac OS X, Windows	Only if Theano backend	Under dev for Theano backend	Yes	Yes	Yes	Yes
Deeplearning4j	Apache 2.0	Linux, Mac OS X, Windows, Android (Cross-platform)	Yes	On roadmap	Yes	Yes	Yes	Yes
MXNet	Apache 2.0	Linux, Mac OS X, Windows, AWS, Android iOS, JavaScript	Yes	On roadmap	Yes	Yes	Yes	Yes
Theano	BSD license	Cross-platform	Yes	Under development	Yes	Yes	Yes	Yes
Facebook Torch	BSD License	Linux, Mac OS X, Windows, Android, iOS	Yes	3rd party implementations	Yes	Yes	Yes	Yes

Started from Wikipedia

TensorFlow for Deep Learning

- Open source library for Machine Learning and Deep Learning by Google.
- Supports CUDA, CNN, RNN and DBN. Distributed TensorFlow offers flexibility to scale up to hundreds of GPUs, train models with a huge number of parameters.
- Has a well documented Python API, less documented C++ and Java APIs.
- Uses XLA, JIT, AOT, and other compilation techniques to minimize execution time and maximize computing resources.
- TensorBoard – Visualize TensorFlow graphs, monitor training performance, and explore how models represent data.
- Layers, Estimators, and Canned Estimators for defining models.
- Keras DL framework can be used in Tensorflow. DeepMind also uses TensorFlow.
- TensorFlow models can be deployed in iOS and Android apps, and Raspberry Pi.
- TensorFlow Serving, a flexible, high-performance ML serving system designed for production environments.
- TensorFlow has a toolkit of ML algorithms.

Apache Spark for Streaming and Machine Learning

- Open source library for SQL, Streaming, ML and Graph in a distributed cluster.
- Provides APIs for Scala, Java, Python and R.
- DAG execution engine supports acyclic data flow and in-memory computing.
- Runs on Hadoop, Mesos, standalone, or in the cloud. It can access diverse data sources including HDFS, Cassandra, HBase, Hive, and S3.
- Supports standalone (native Spark cluster), Hadoop YARN, or Apache Mesos.
- Spark Streaming has support built-in to consume from Kafka, Flume, Twitter, ZeroMQ, Kinesis, and TCP/IP sockets.
- Spark MLib simplifies large scale machine learning pipelines, including:
  - Summary statistics, correlations, stratified sampling, hypothesis testing, random data generation[16]
  - Classification and regression: support vector machines, logistic regression, linear regression, decision trees, naive Bayes
  - Collaborative filtering techniques including alternating least squares (ALS)
  - Cluster analysis methods including k-means, and Latent Dirichlet Allocation (LDA)
  - Dimensionality reduction techniques: singular value decomposition (SVD), and principal component analysis (PCA)
  - Feature extraction and transformation functions
  - Optimization algorithms such as stochastic gradient descent, limited-memory BFGS (L-BFGS)
- GraphX is a distributed graph processing framework.

THANK YOU