Machine and Deep Learning with MATLAB

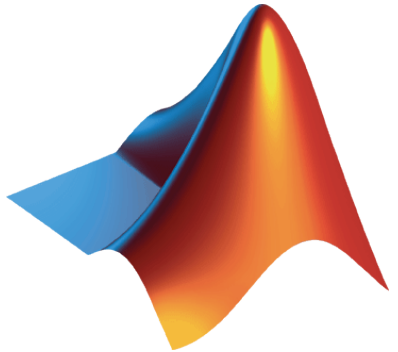

Alexander Diethert, Application Engineering

May, 24th 2018, London


RICE HOUSE

Agenda

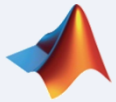	**Artificial Intelligence enabled by Machine and Deep Learning**
	Machine Learning
	Deep Learning
	Outlook: Integration in Production Systems

Artificial Intelligence Is in Early Adoption

Percentage of Respondents

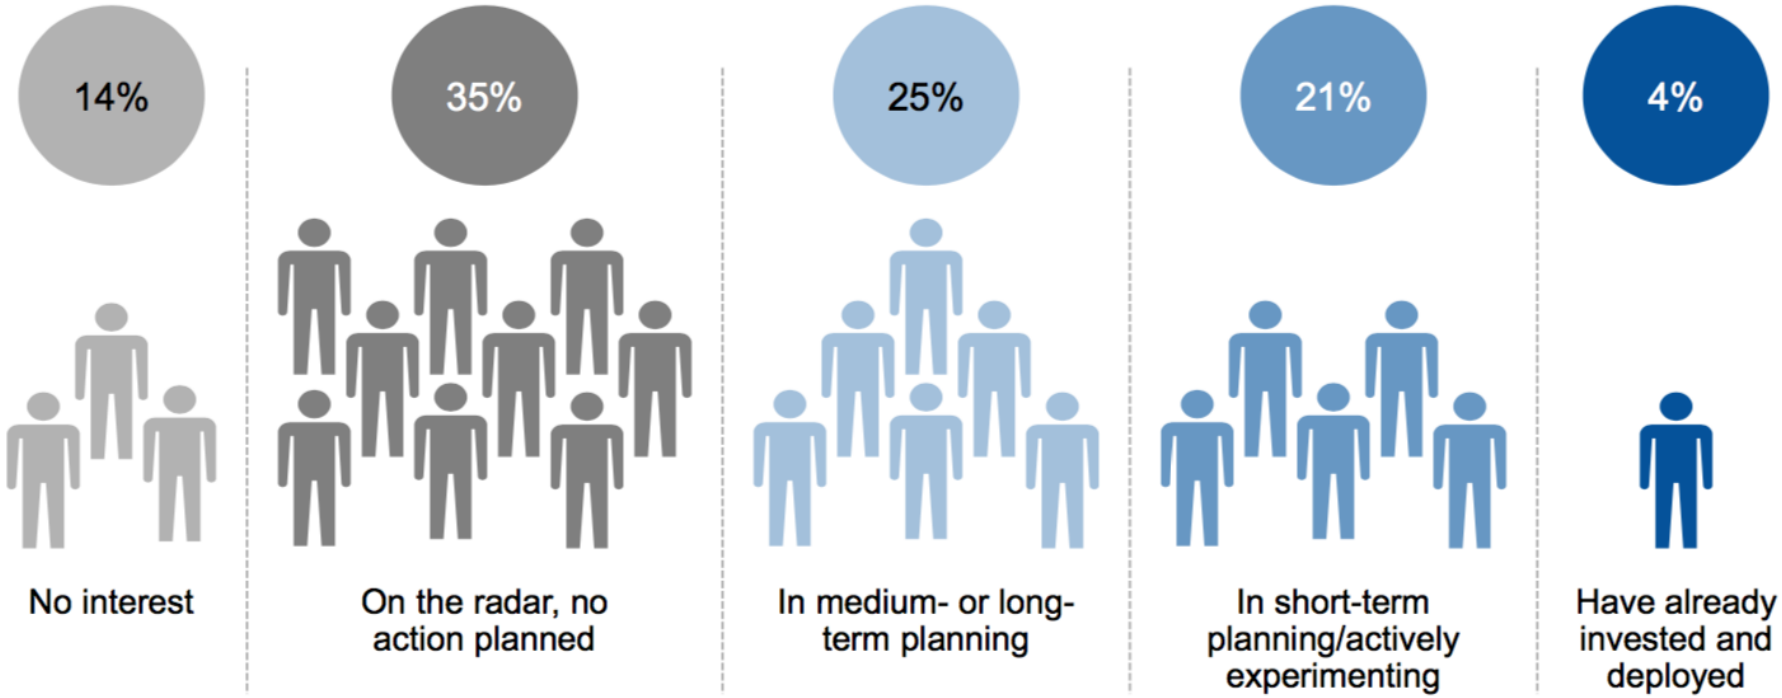

Q: What are your organization's plans in terms of artificial intelligence?
Base: All Answering, n = 3.138
Source: Gartner 2018 CIO Survey

Big Data

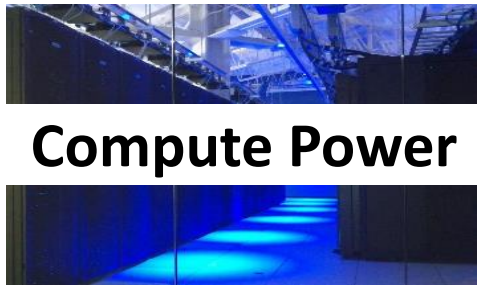

Compute Power

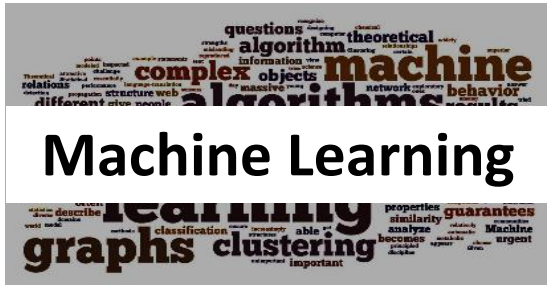

Machine Learning

Analytics are pervasive – *Why Now?*

We have data

- **Engineering**
- **Business**
- **Transactional**

We have compute

- **Desktop**
  Multicore, GPU
- **Clusters**
- **Cloud computing**
- **Hadoop with Spark**

We know how

- **Neural Networks**
- **Classification**
- **Clustering**
- **Regression**
- **…and much more…**


amazon


Luke Skywalker
94.5%


ВЫХОД В ГОРОД
ACCESS TO CITY

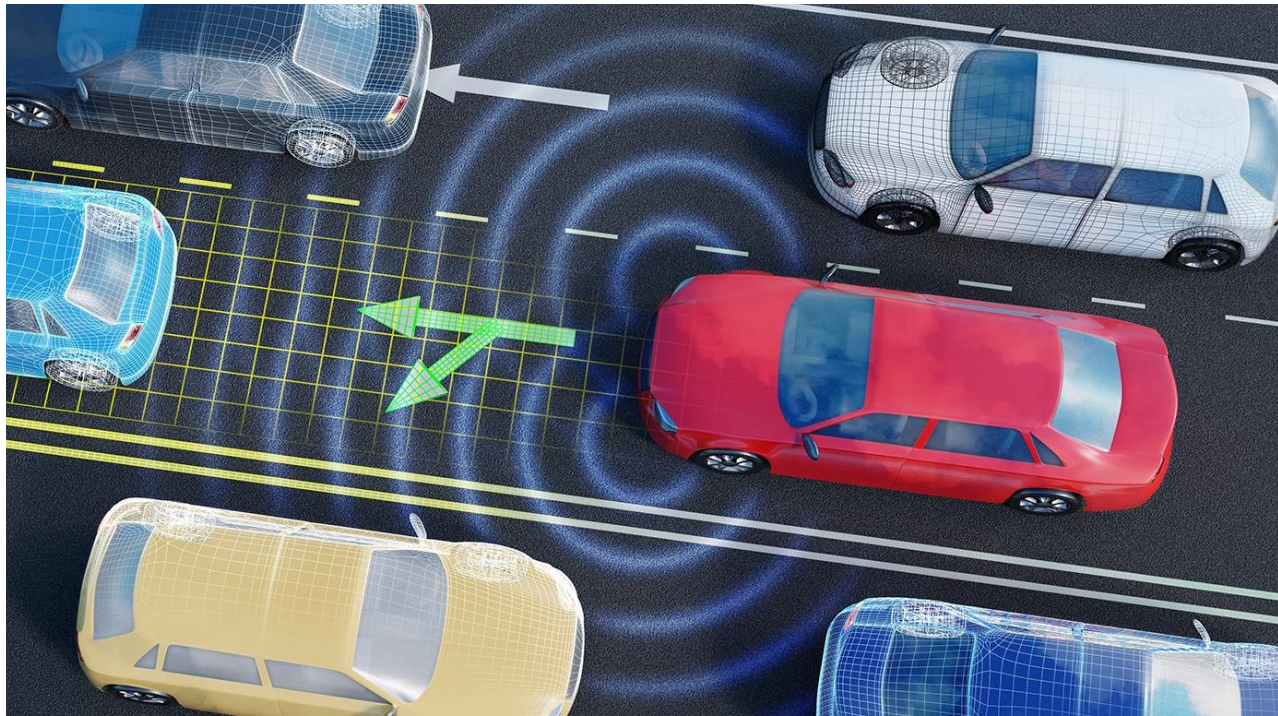

There are two ways to get a computer to do what you want

Traditional Programming

There are two ways to get a computer to do what you want

Machine Learning

There are two ways to get a computer to do what you want

Machine Learning

Artificial Intelligence

Machine Learning

AI, Machine Learning, and Deep Learning

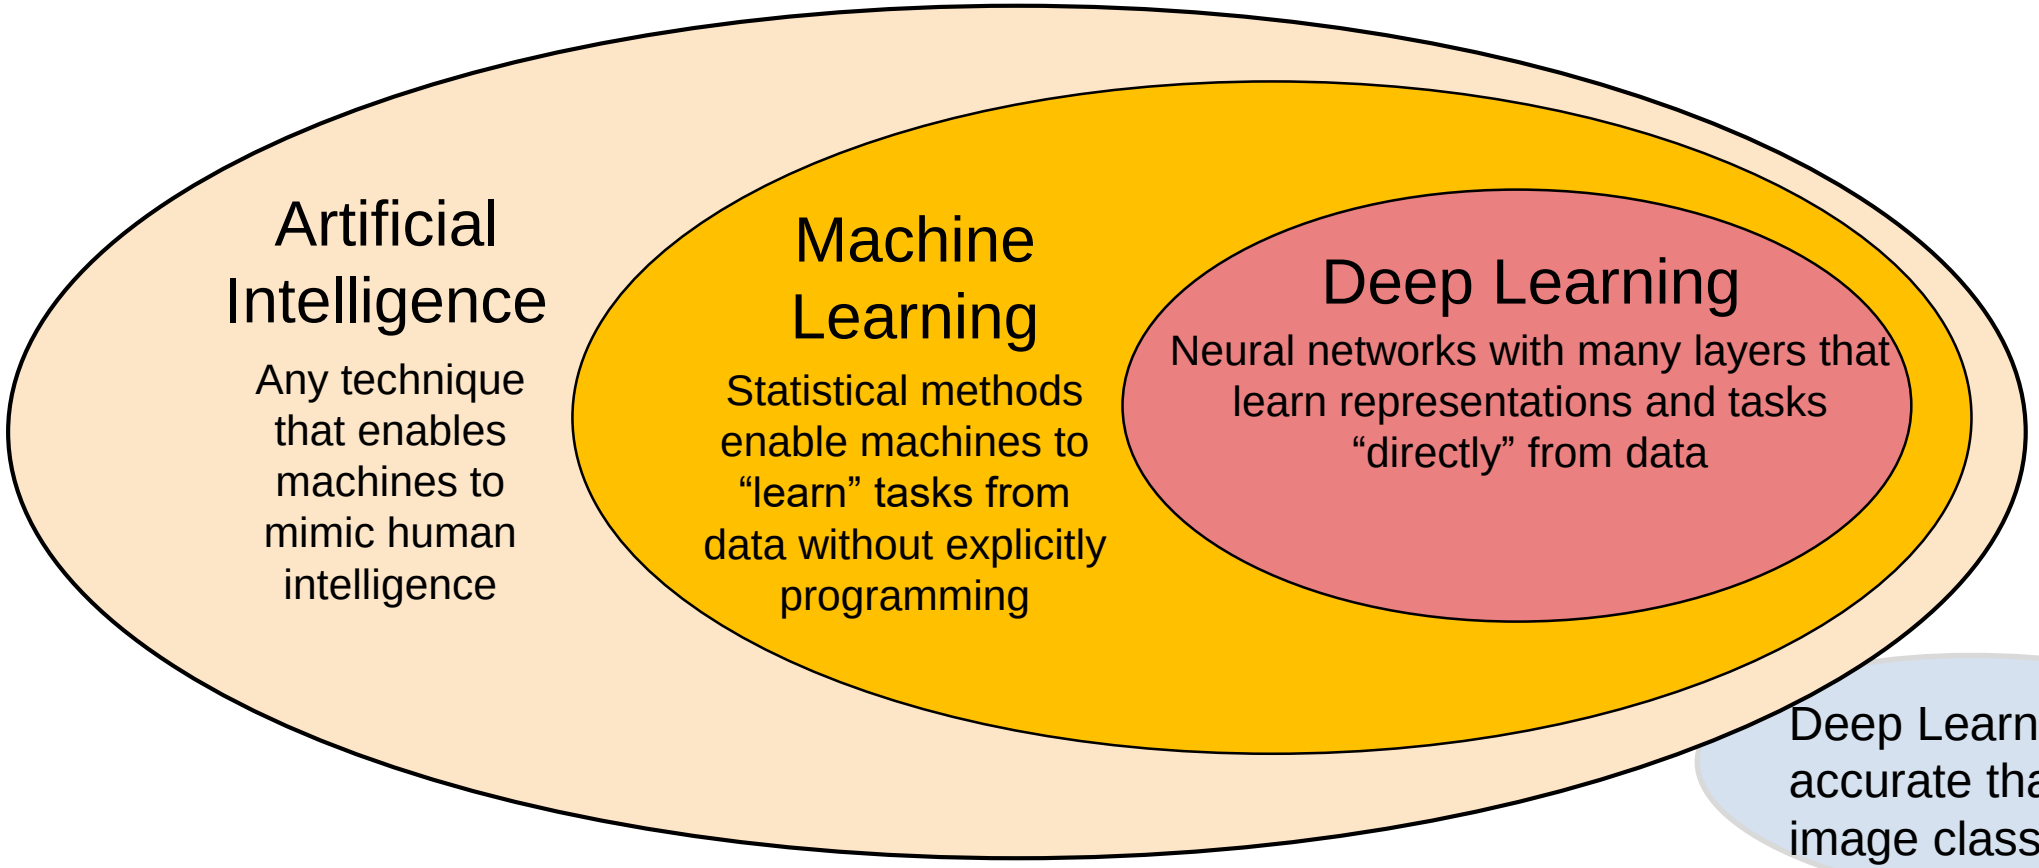

Deep Learning more accurate than humans on image classification

1950s — 1980s — 2015

FLOPS — *Thousand* — *Million* — *Quadrillion*

What can Machine and Deep Learning do?

Machine learning

Machine learning deals with the problem of extracting ***features*** from data so as to solve many different ***predictive*** tasks:

- ❑Forecasting (*e.g. Energy demand prediction, sales*)
- ❑Imputing missing data (*e.g. Netflix recommendations*)
- ❑Detecting anomalies (*e.g. Intruders, virus mutations*)
- ❑Classifying (*e.g. Credit risk assessment, cancer diagnosis*)
- ❑Ranking (*e.g. Google search, personalization*)
- ❑Summarizing (*e.g. News zeitgeist, social media sentiment*)
- ❑Decision making (*e.g. AI, robotics, compiler tuning, trading*)

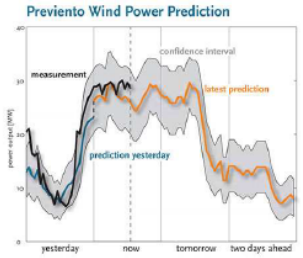

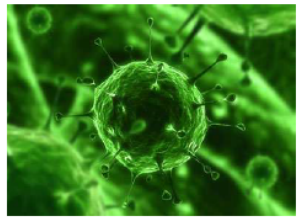

When to apply machine learning

- ❑Human expertise is absent (*e.g. Navigating on Mars*)
- ❑Humans are unable to explain their expertise (*e.g. Speech recognition, vision, language*)
- ❑Solution changes with time (*e.g. Tracking, temperature control, preferences*)
- ❑Solution needs to be adapted to particular cases (*e.g. Biometrics, personalization*)
- ❑The problem size is to vast for our limited reasoning capabilities (*e.g. Calculating webpage ranks, matching ads to facebook pages*)

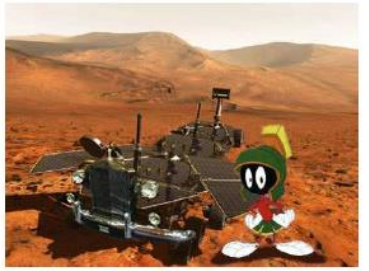

http://www.cs.ubc.ca/~nando/340-2012/lectures/l1.pdf

Example: Predictive Analytics in e-commerce

Engineering Data

Images

Social profile

Geolocation

Keystroke logs

Business Data

Transactions

Use Image Processing

to add image data to the model,

improving performance

IMPROVED Predictive Model

Improved Offer to Customer

Applications of Machine Learning and Deep Learning in Finance

Algorithmic Trading

Sentiment Analysis

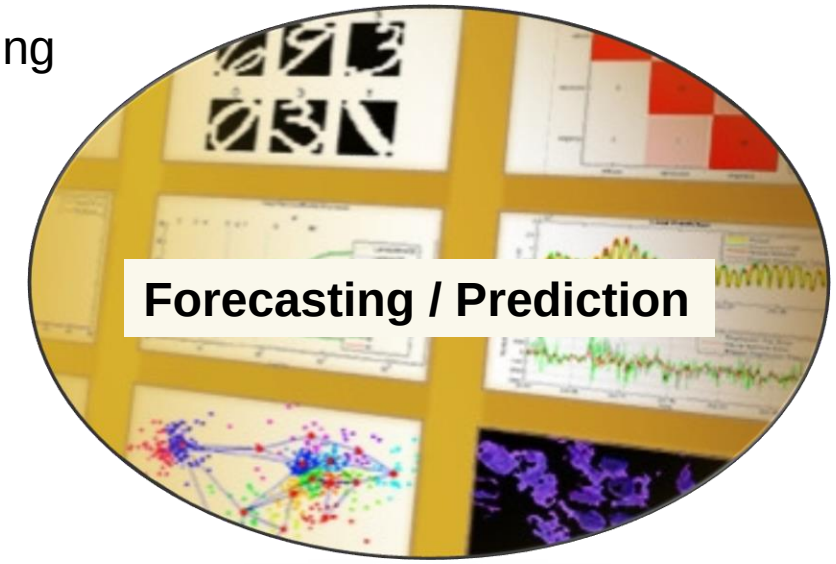

Fraud Detection

Credit Decision Making

Financial Planning

Agenda

	Artificial Intelligence enabled by Machine and Deep Learning
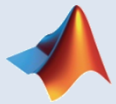	**Machine Learning**
	Deep Learning
	Outlook: Integration in Production Systems

Customer References

MATLAB Used to Predict Financial Crises in Emerging Markets

"Because MATLAB is both powerful and easy to use, I felt confident that the Bank of Indonesia would be able to implement the MATLAB programs and use them as an early warning system for financial distress."

— *Dr. Paul McNelis, Georgetown University*

Machine Learning and Visualisation in the Context of a Large Enterprise
Arjun Viswanathan, Citi

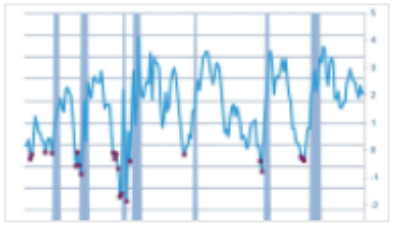

Portfolio Allocation with Machine Learning and MATLAB Distributed Computing Server on Microsoft Azure Cloud
Emilio Llorente-Cano and James Mann, Aberdeen Asset Management

Gas Natural Fenosa Predicts Energy Supply and Demand

"Because we need to rapidly respond to shifting production constraints and changing demands, we cannot depend on closed or proprietary solutions. With MathWorks tools we get more accurate results—and we have the flexibility to develop, update, and optimize our models in response to changing needs."

— *Angel Caballero, Gas Natural Fenosa*

Portomouros hydroelectric dam.

Machine Learning and Applications in Finance
Christian Hesse, Deutsche Bank and University College London

Example: Machine Learning for Risk Managers

Machine learning is enabling better models for complex problems

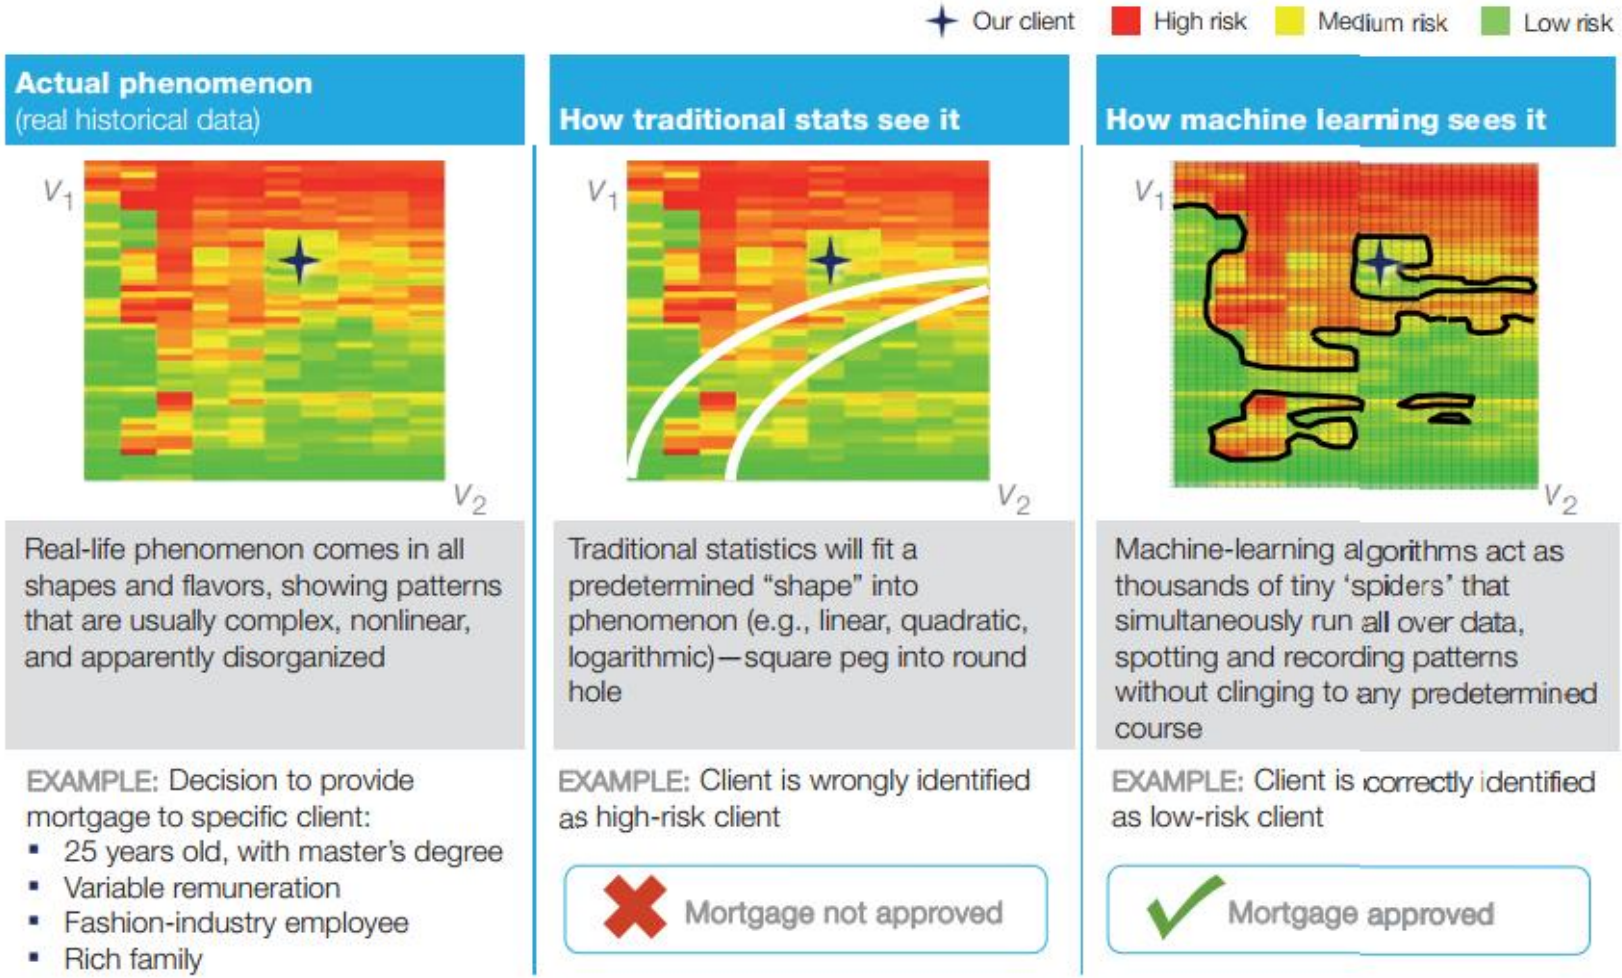

https://www.mckinsey.com/~/media/mckinsey/dotcom/client_service/risk/pdfs/the_future_of_bank_risk_management.ashx

SOURCE: McKinsey analysis

Machine Learning Workflow

1 Access and Explore Data

Files

Databases

Sensors

2 Preprocess Data

Working with Messy Data

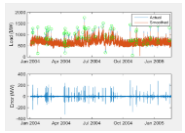

Data Reduction/ Transformation

Feature Extraction

3 Develop Predictive Models

Model Creation e.g. Machine Learning

Parameter Optimization

Model Validation

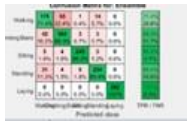

4 Integrate with Production Systems

Desktop Apps

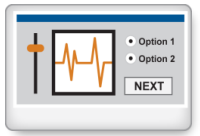

Enterprise Scale Systems

kafka

AWS Kinesis

OSIsoft.

Embedded Devices and Hardware

5 Visualize Results

3rd party dashboards

Spotfire® TIBCO Software

+ableau

Qlik

Power BI

Web apps

Types of Machine Learning

Type of Learning **Categories of Algorithms**

Machine Learning

Unsupervised Learning

Group & **interpret** data based only on **input data**

Clustering

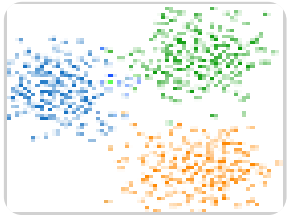

Output is the # of groups formed from similar data. Find natural groups and patterns from input data only

Supervised Learning

Develop **predictive model** based on both **input and output** data

Classification

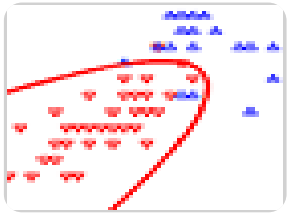

Output is a choice between classes (True, False) (Red, Blue, Green)

Regression

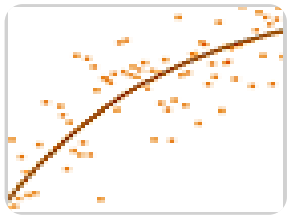

Output is a prediction of the future state

Workflows of Machine Learning

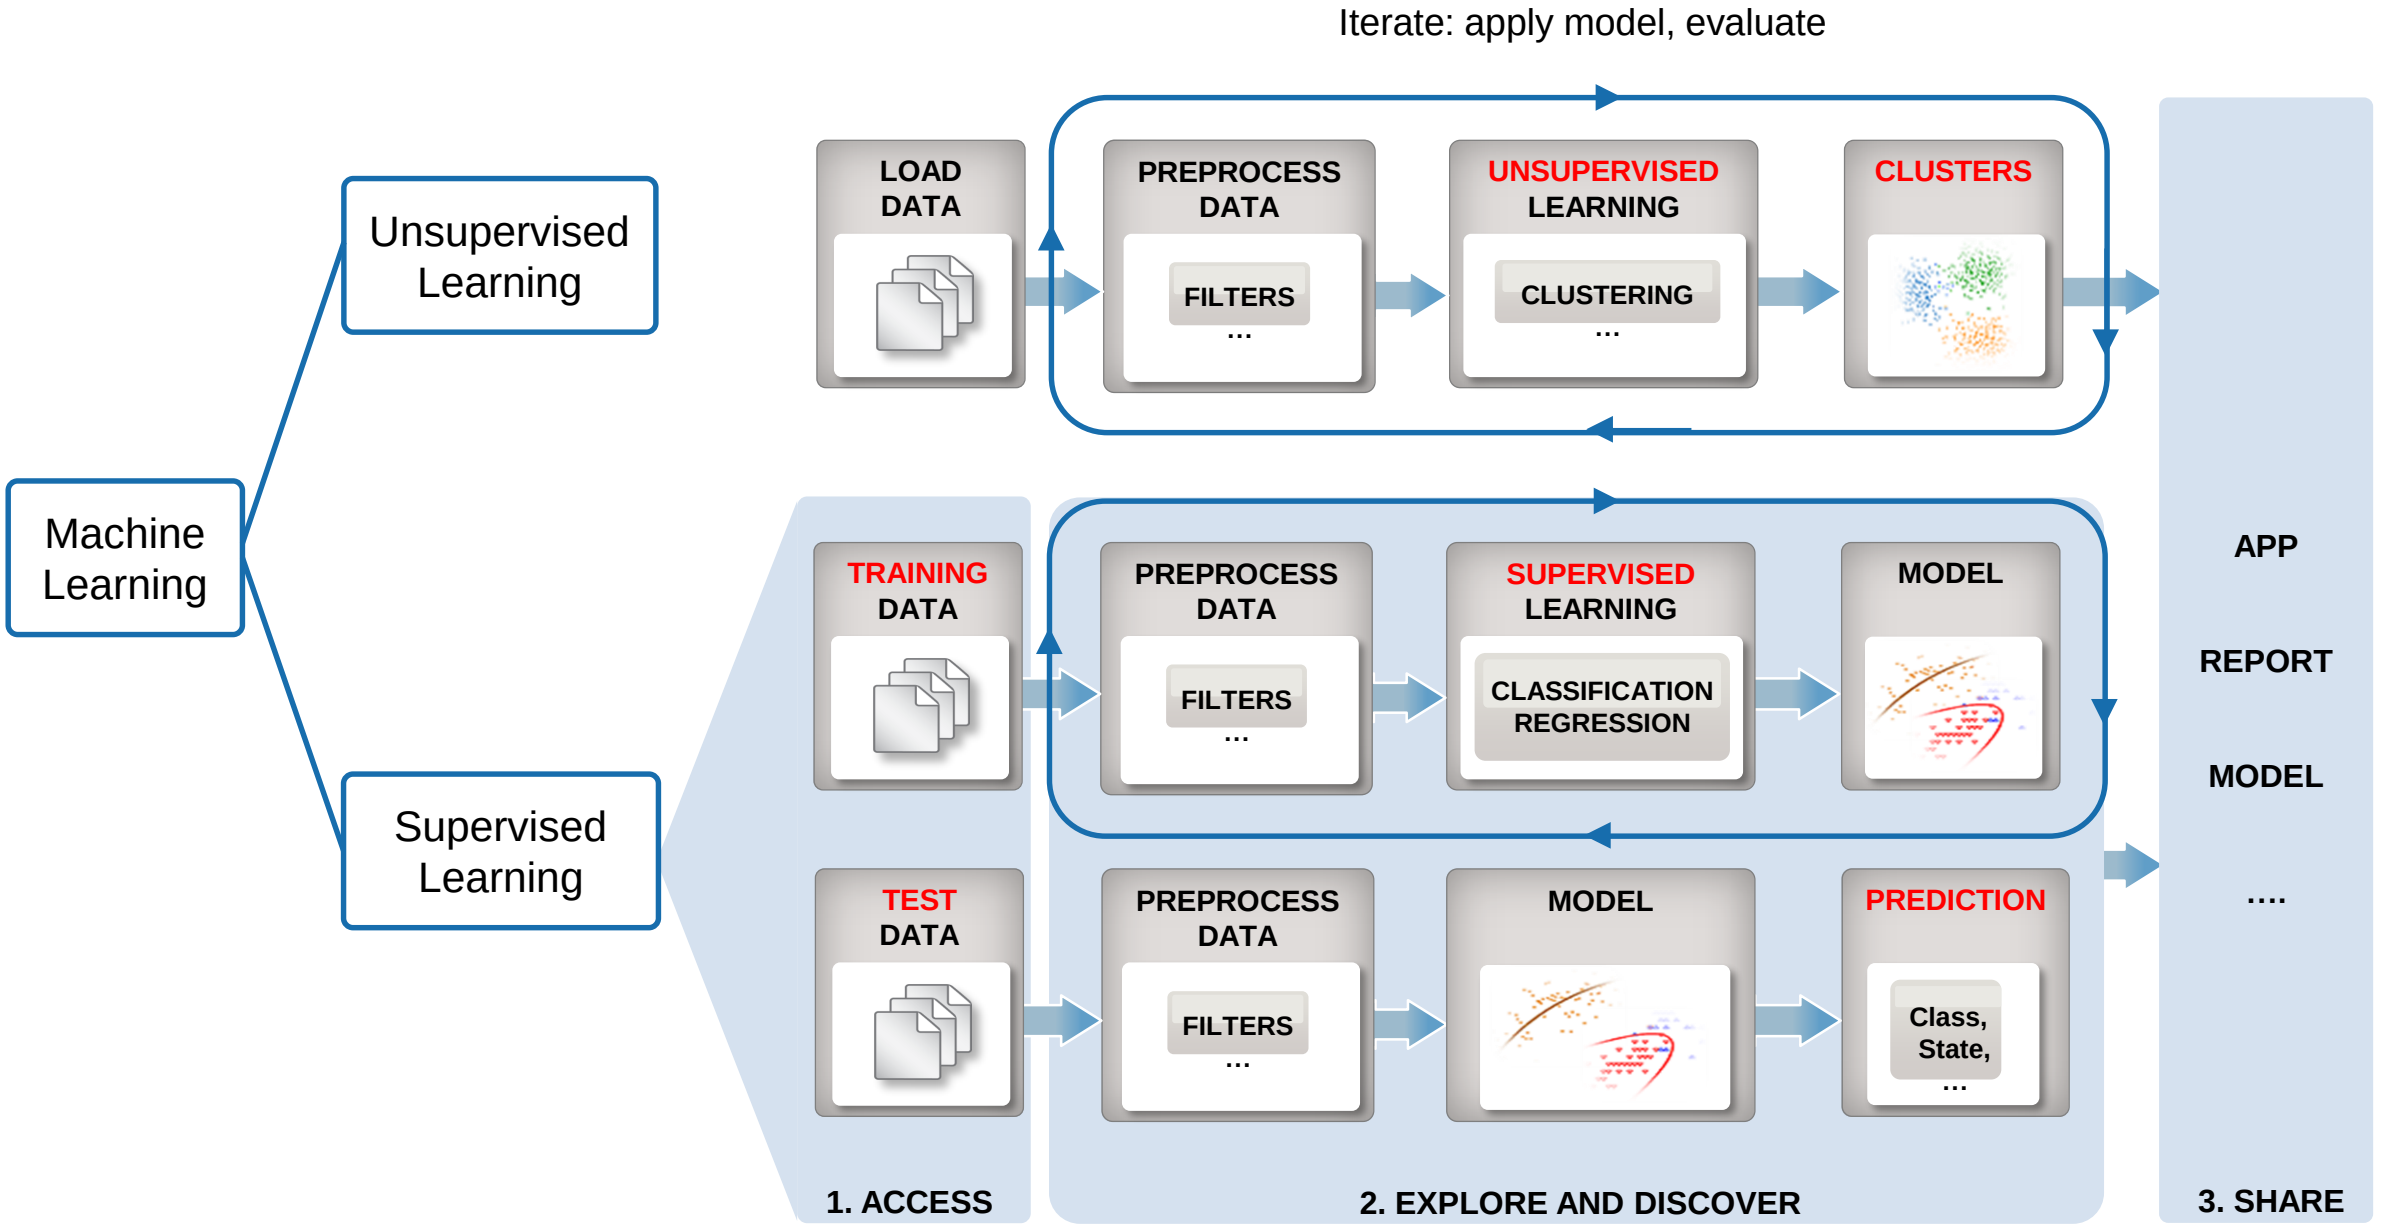

Demo: Classification Learner App

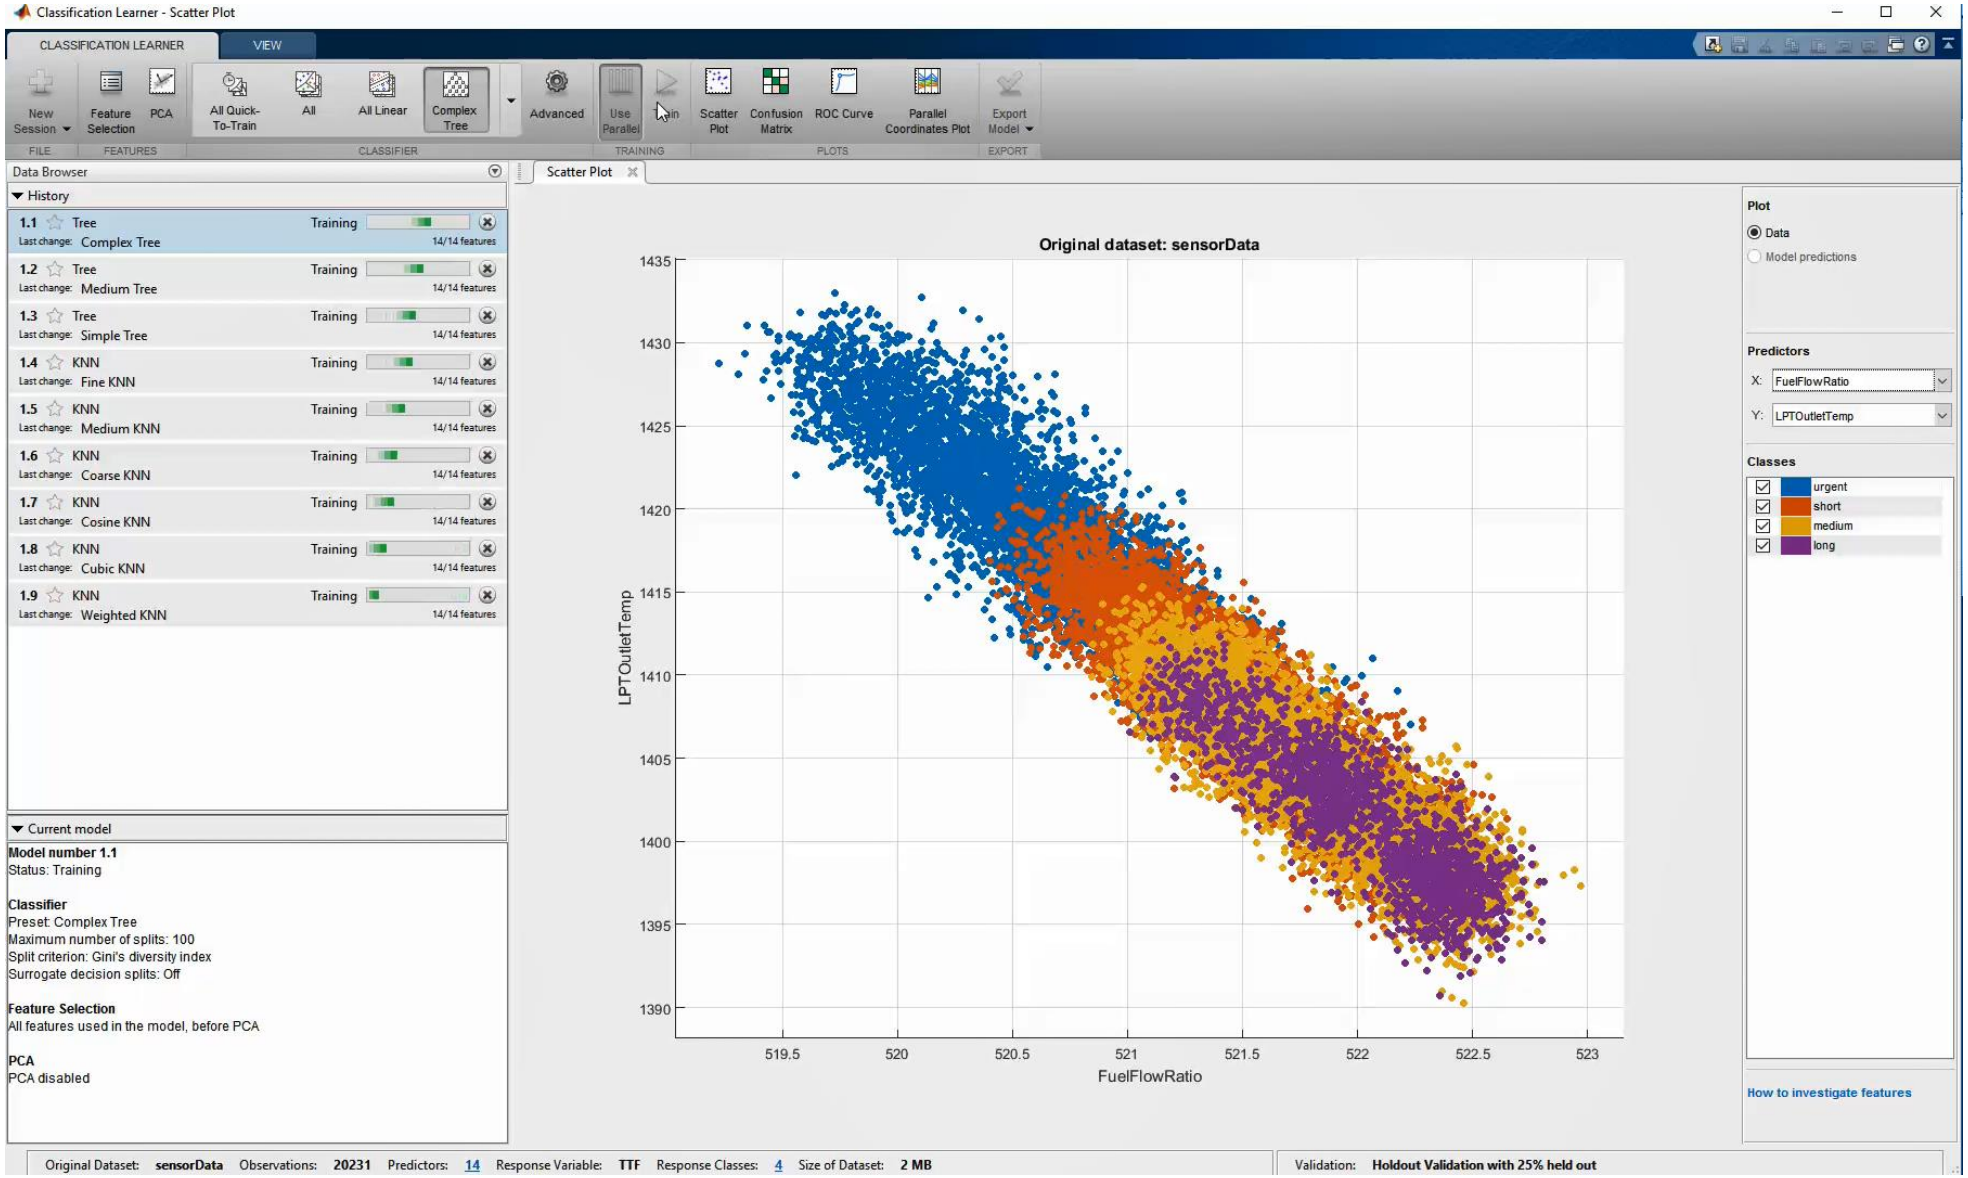

Machine Learning Apps for Classification and Regression

- Point and click interface – no coding required
- Quickly evaluate, compare and select regression models
- Export and share MATLAB code or trained models

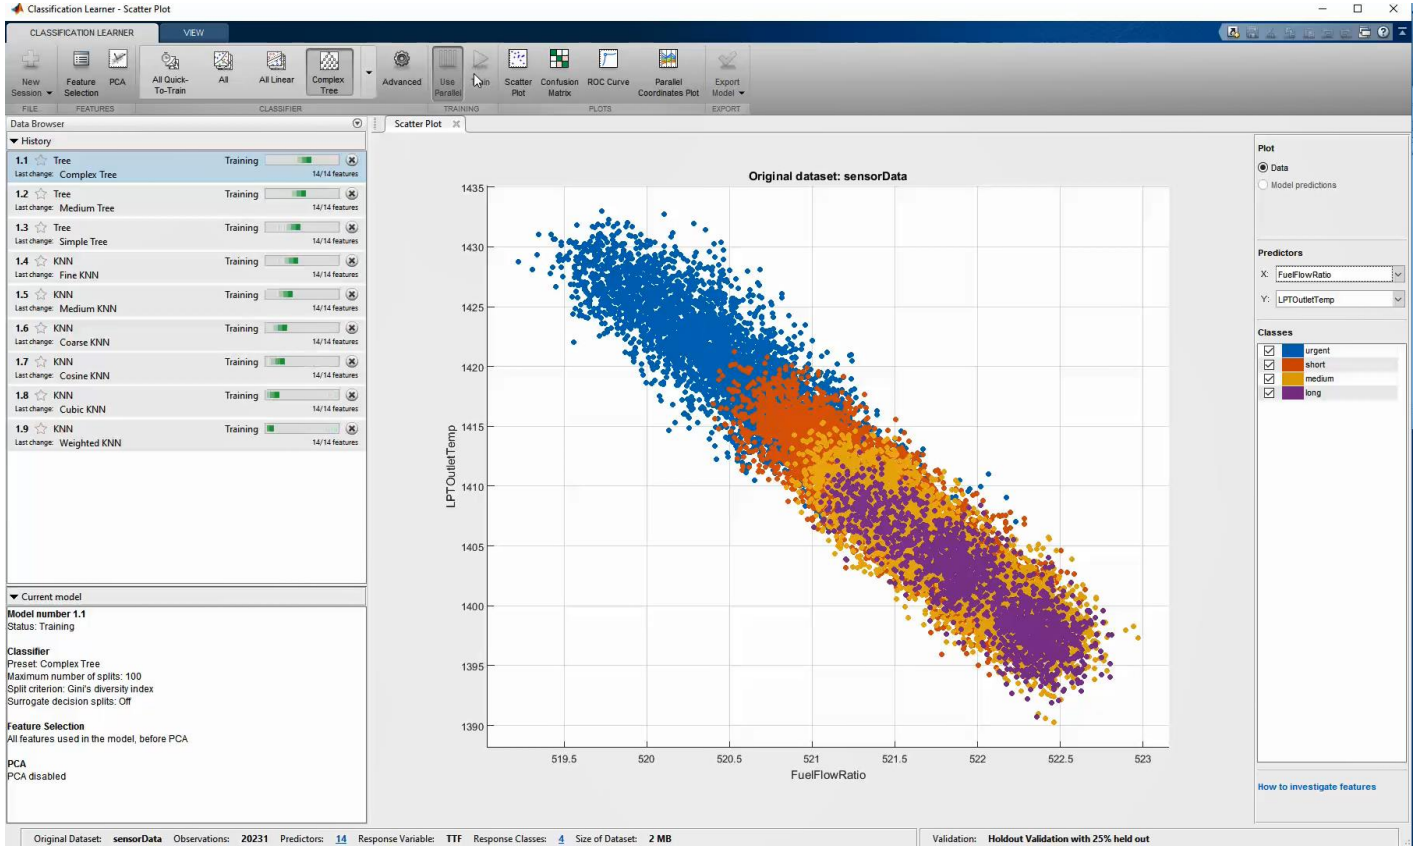

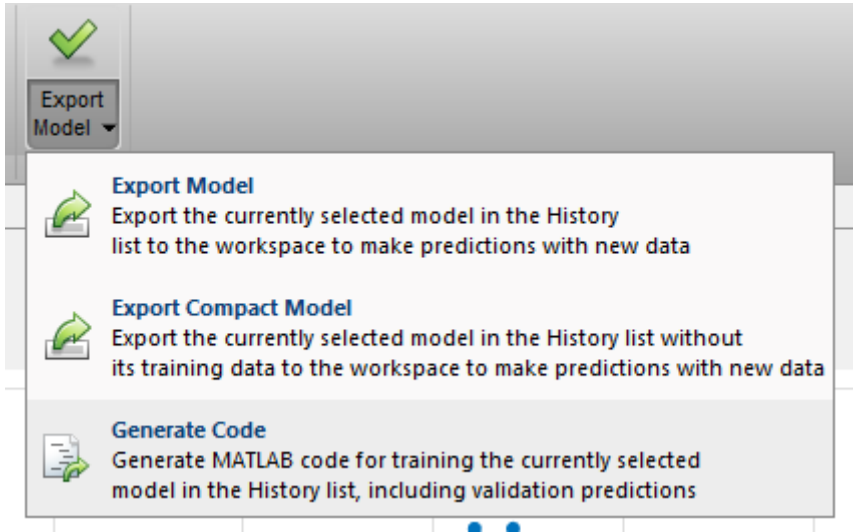

Fine-tuning Model Parameters

Why?
- Manual parameter selection is tedious and may result in suboptimal performance

When?
- When training a model with one or more parameters that influence the fit

Capabilities
- Efficient comparted to standard optimization techniques or grid search
- Tightly integrated with fit function API with pre-defined optimization problem (e.g. bounds)

Hyperparameter Tuning with Bayesian Optimization

```
template = templateSVM(...
    'KernelFunction', 'linear', ...
    'PolynomialOrder', [], ...
    'KernelScale', 0.25, ...
    'BoxConstraint', 0.1, ...
    'Standardize', true);
m = fitcecoc( T, 'Species', 'Learners', template )
```

Previously tuning these parameters was a manual process

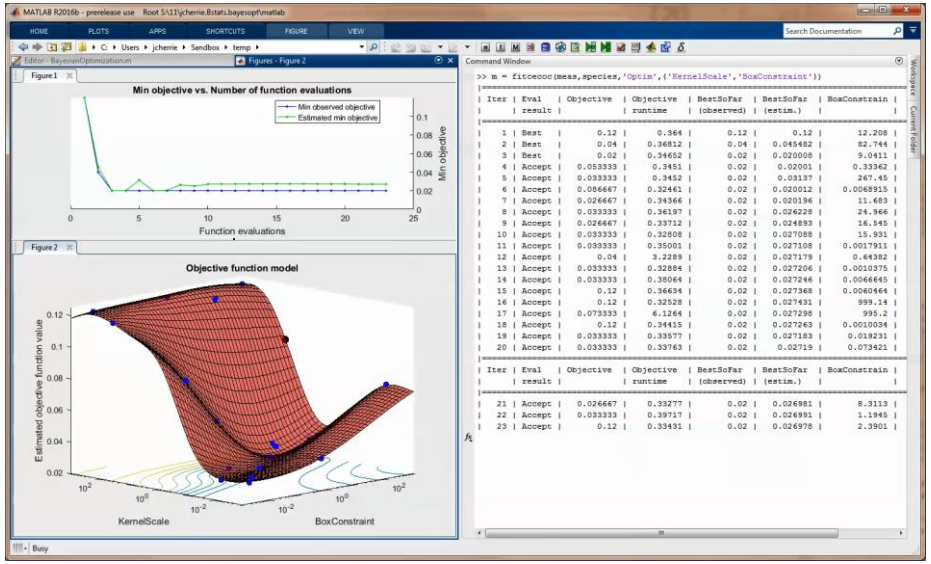

R2016b

Building out your Machine Learning Tool

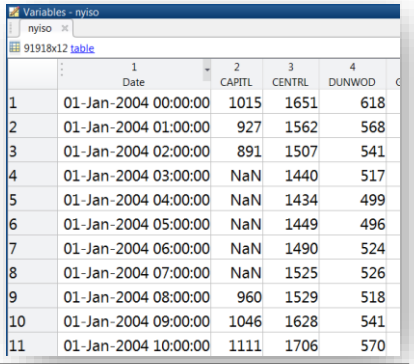

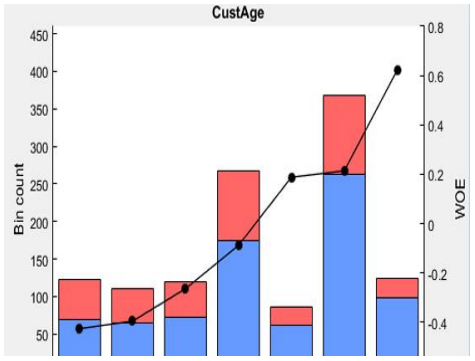

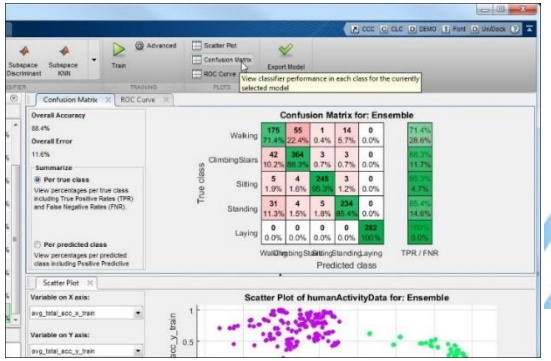

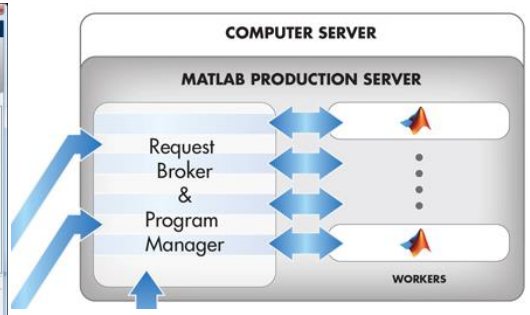

Access and Explore Data

Process Data and Create Feature

Build and Validate Models

Deploy Model
Review Model

Agenda

	Artificial Intelligence enabled by Machine and Deep Learning
	Machine Learning
	**Deep Learning**
	Outlook: Integration in Production Systems

Machine learning vs deep learning

Deep learning performs **end-to-end learning** by learning **features, representations and tasks** directly from **images, text and sound**

Deep learning algorithms also **scale with data –** traditional machine learning **saturates**

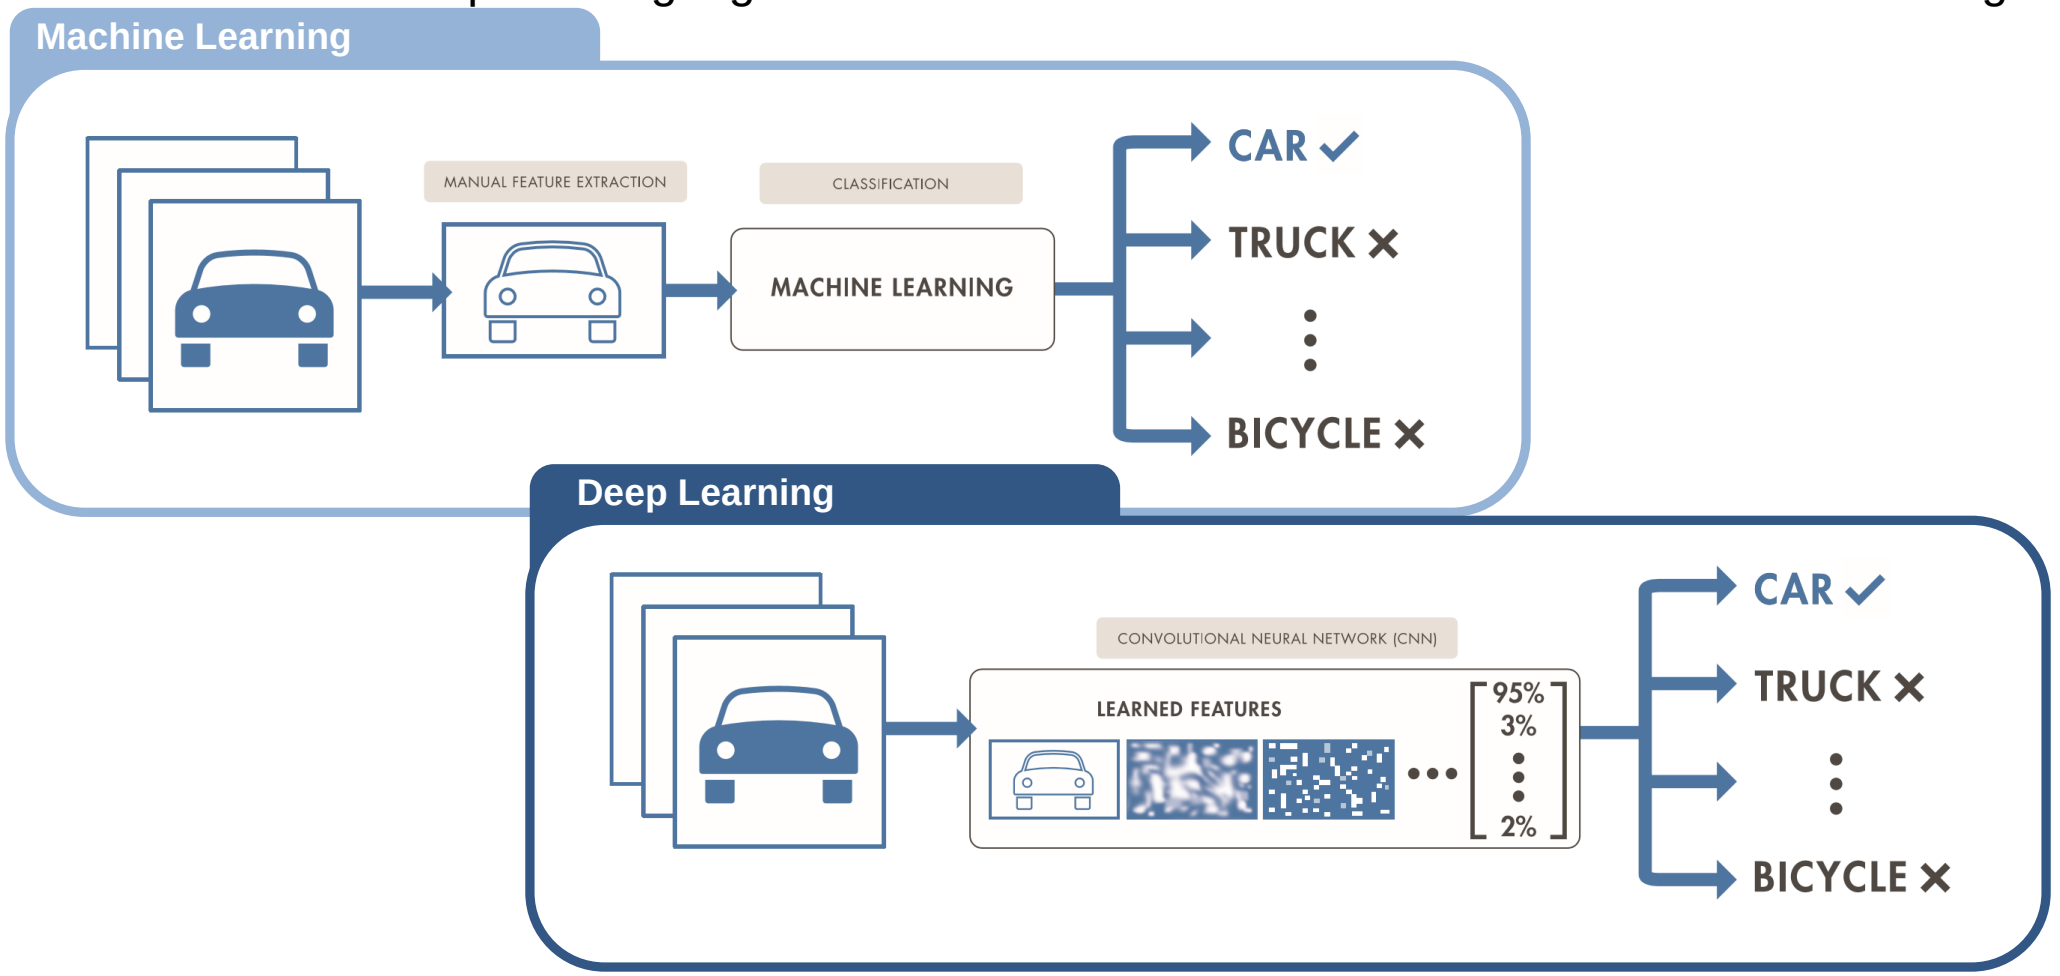

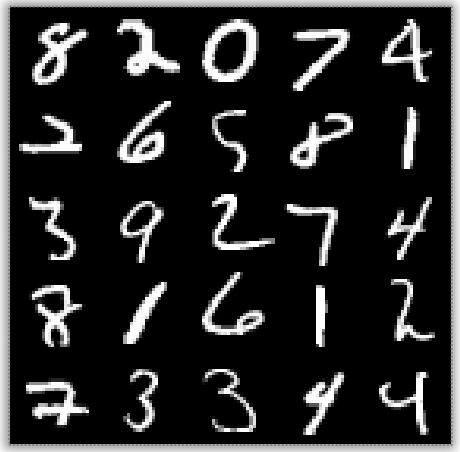

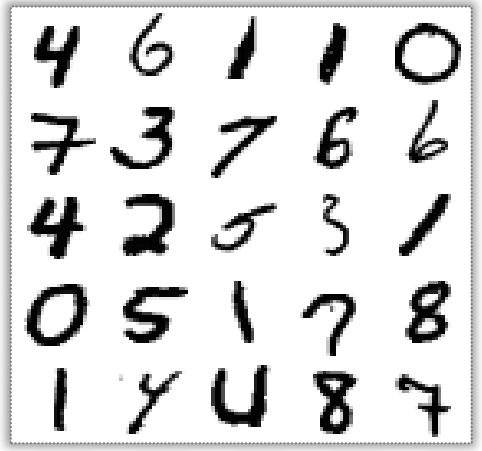

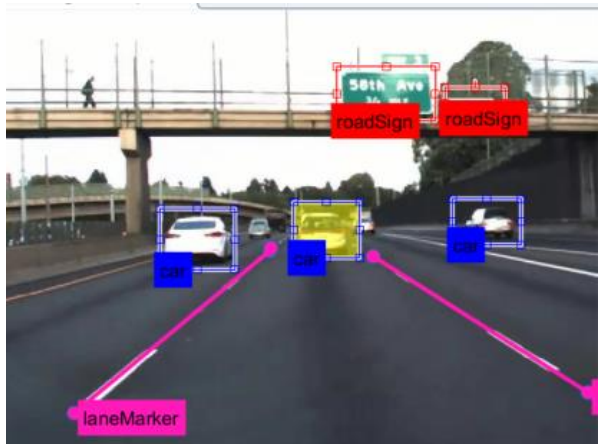

What is Deep Learning?

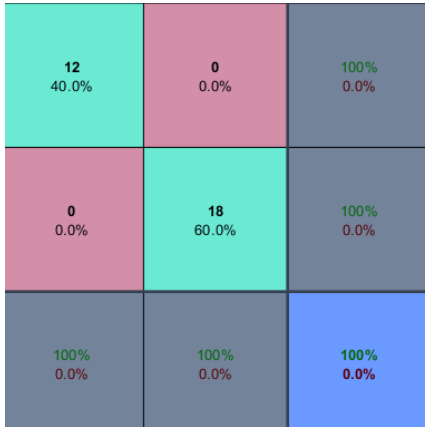

Data Types for Deep Learning

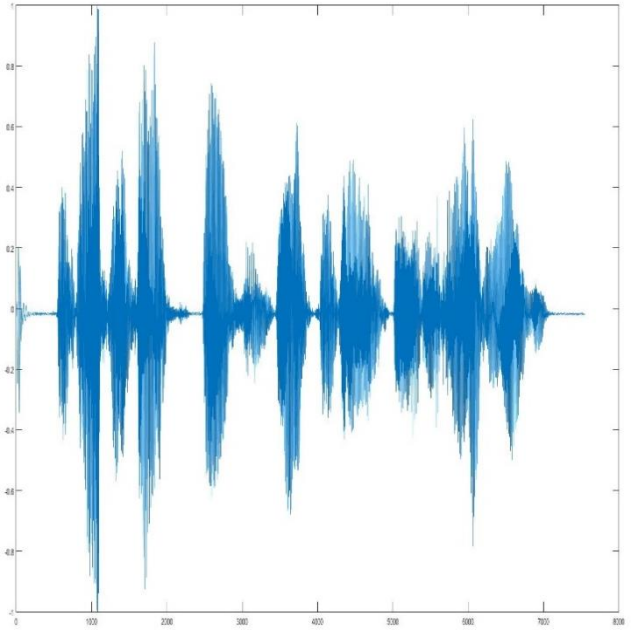

Signal

Text

Image

Deep learning and neural networks

- Deep learning == neural networks; Data flows through network in **layers**
- **Layers** provide transformation of data

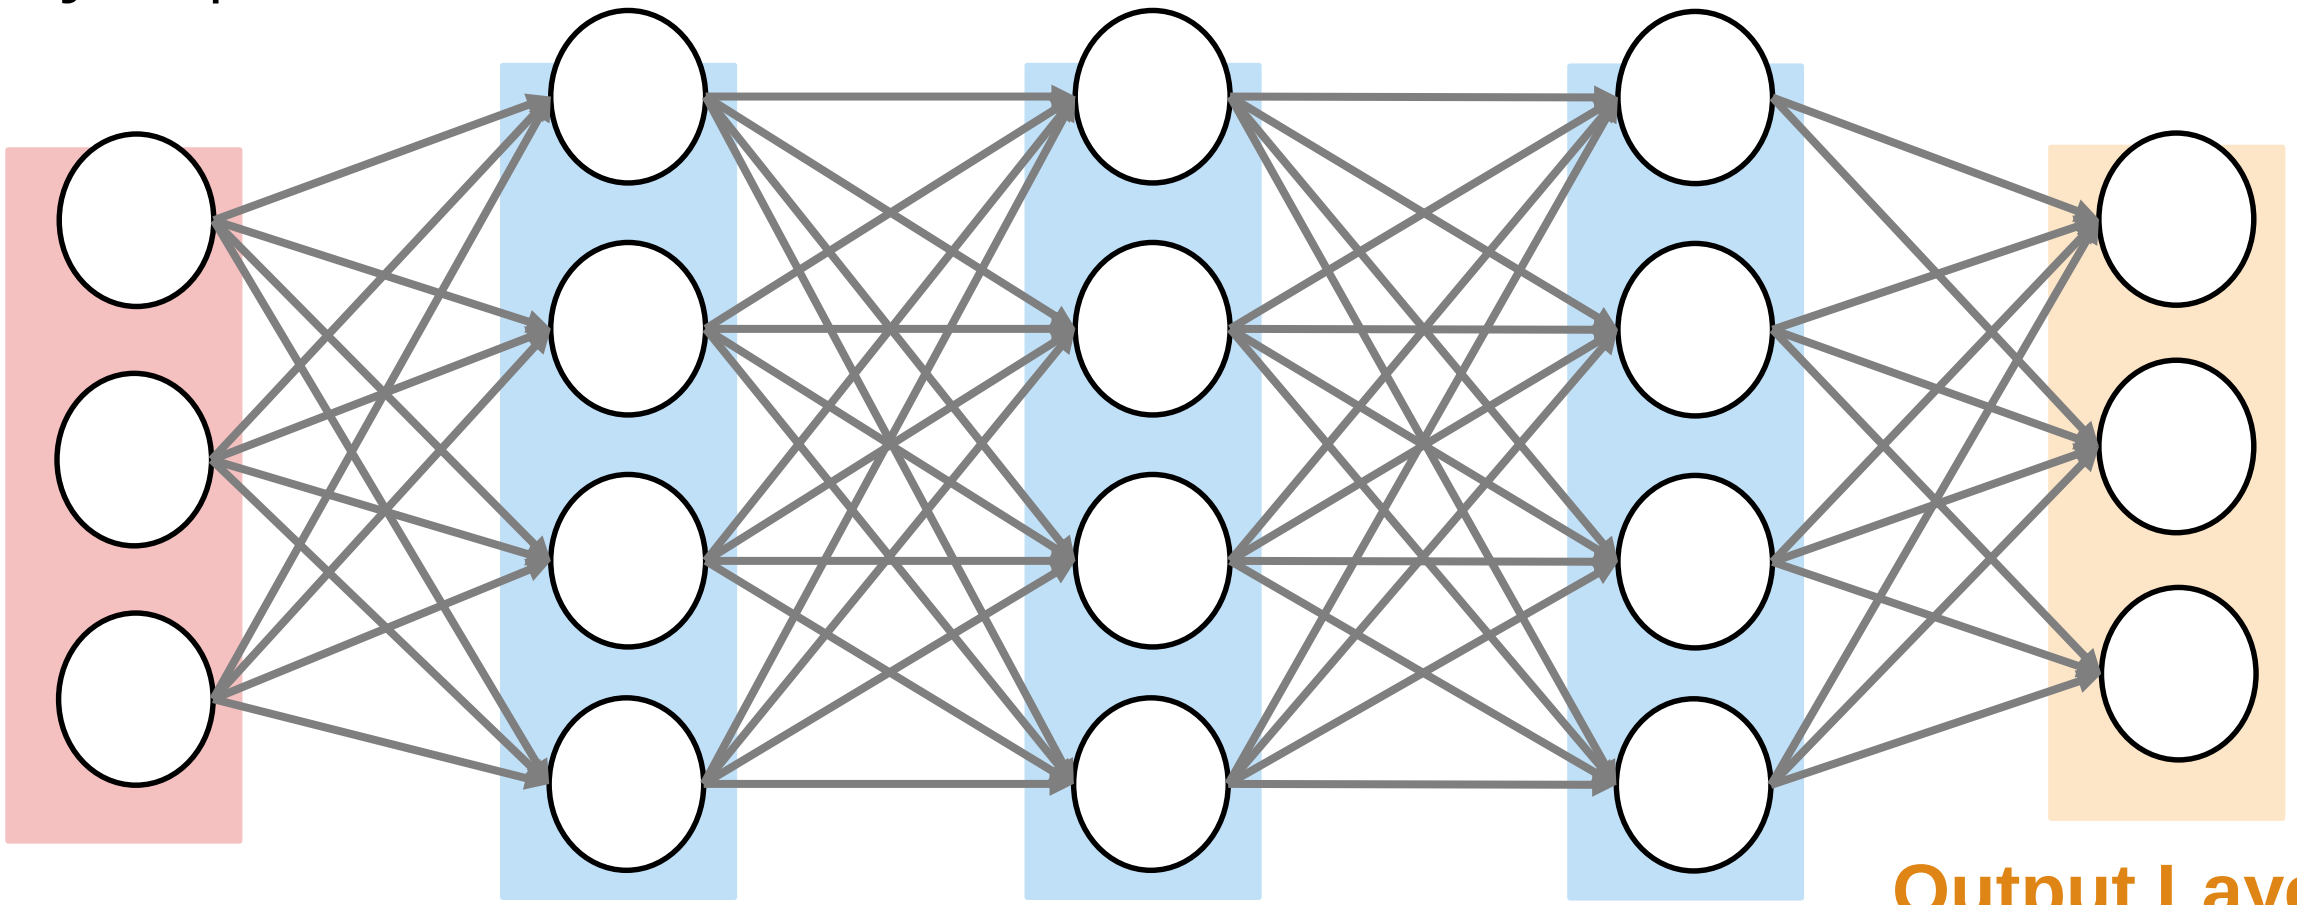

Thinking about Layers

- Layers are like blocks
  - Stack on top of each other
  - Replace one block with a different one
- Each hidden layer processes the information from the previous layer

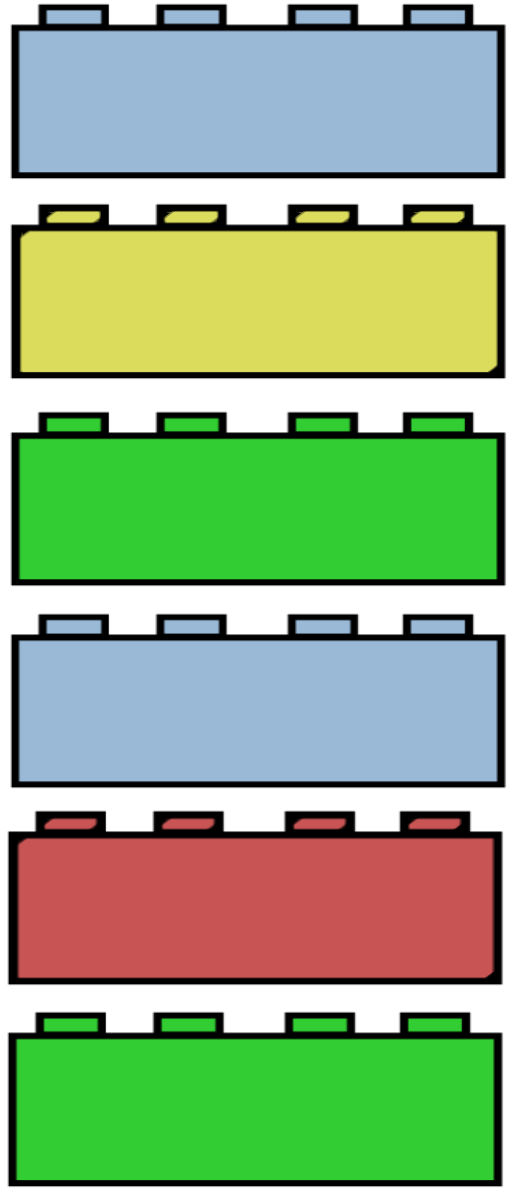

Thinking about Layers

- Layers are like blocks
  - Stack them on top of each other
  - Replace one block with a different one
- Each hidden layer processes the information from the previous layer
- Layers can be ordered in different ways

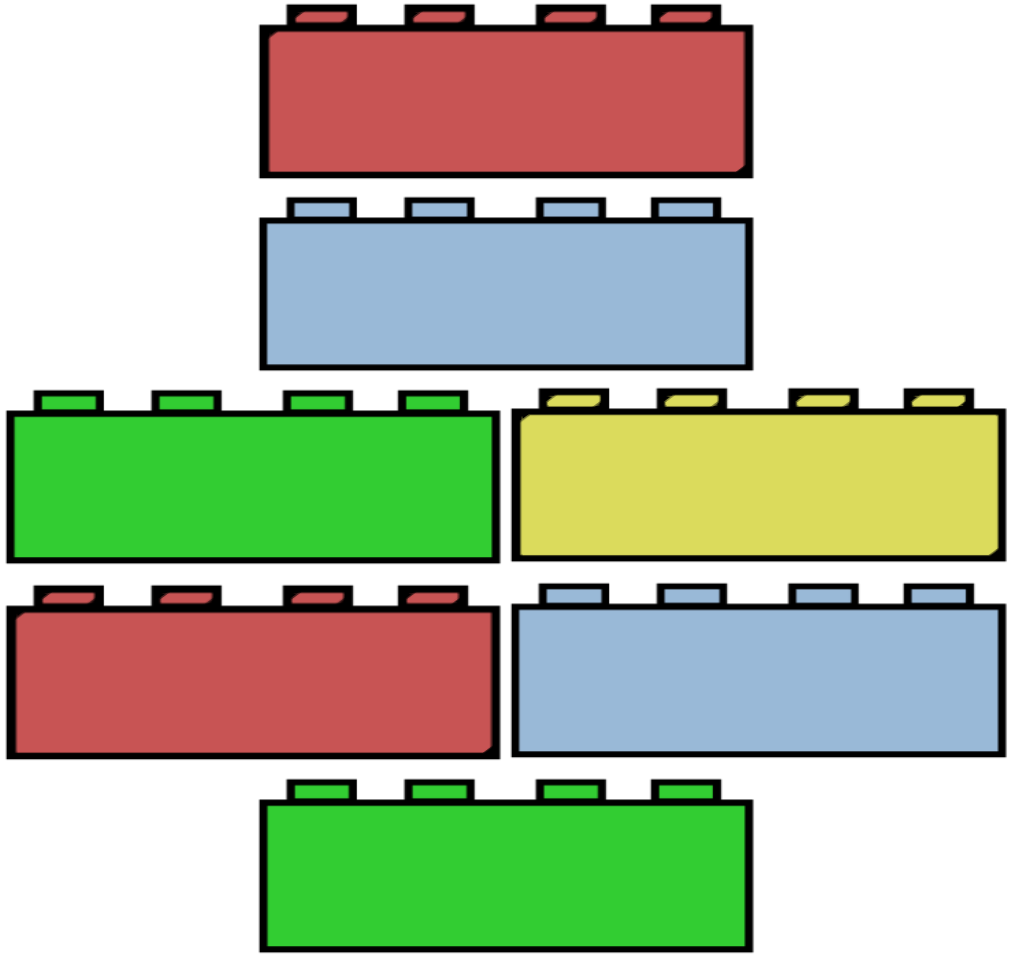

Convolutional neural networks

- Train “deep” neural networks on structured data (e.g. images, signals, text)
- Implements Feature Learning: Eliminates need for “hand crafted” features
- Training using GPUs for performance

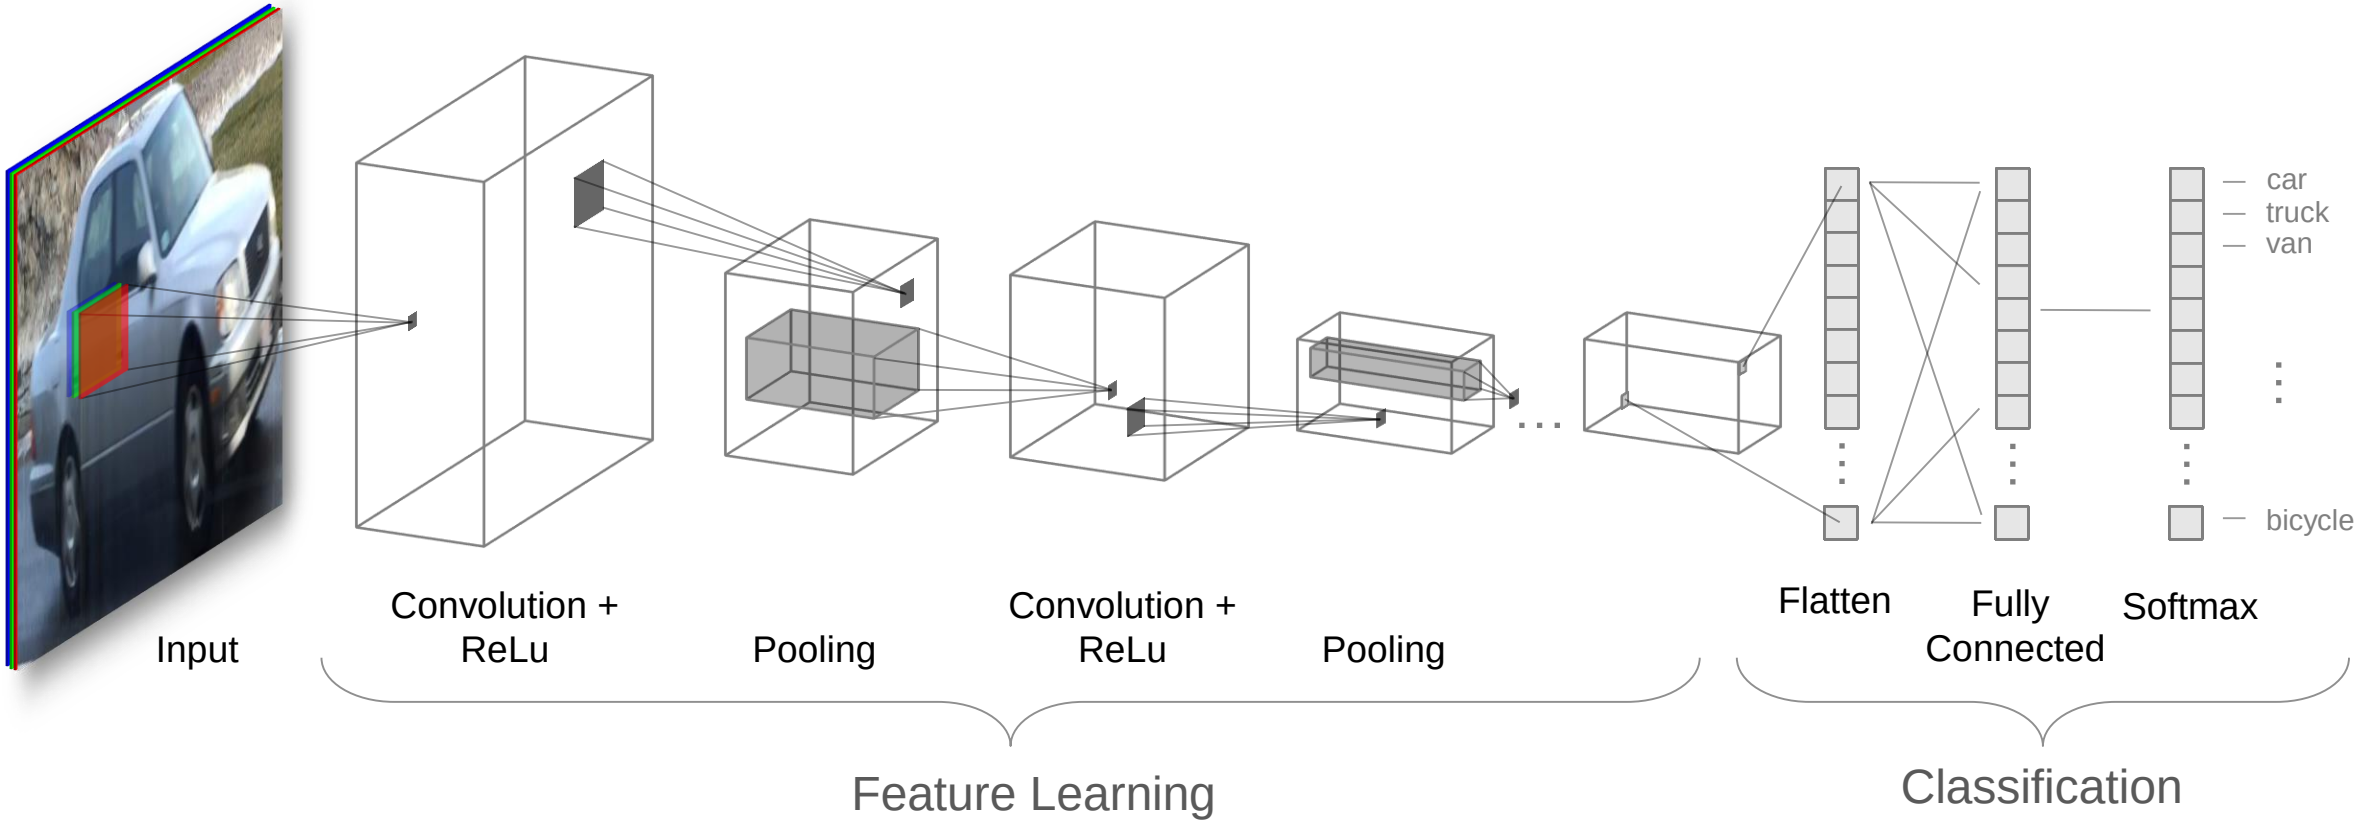

Convolutional Neural Networks (CNN)

- CNN take a fixed size input and generate fixed-size outputs.
- Convolution puts the input images through a set of convolutional filters, each of which activates certain features from the input data.

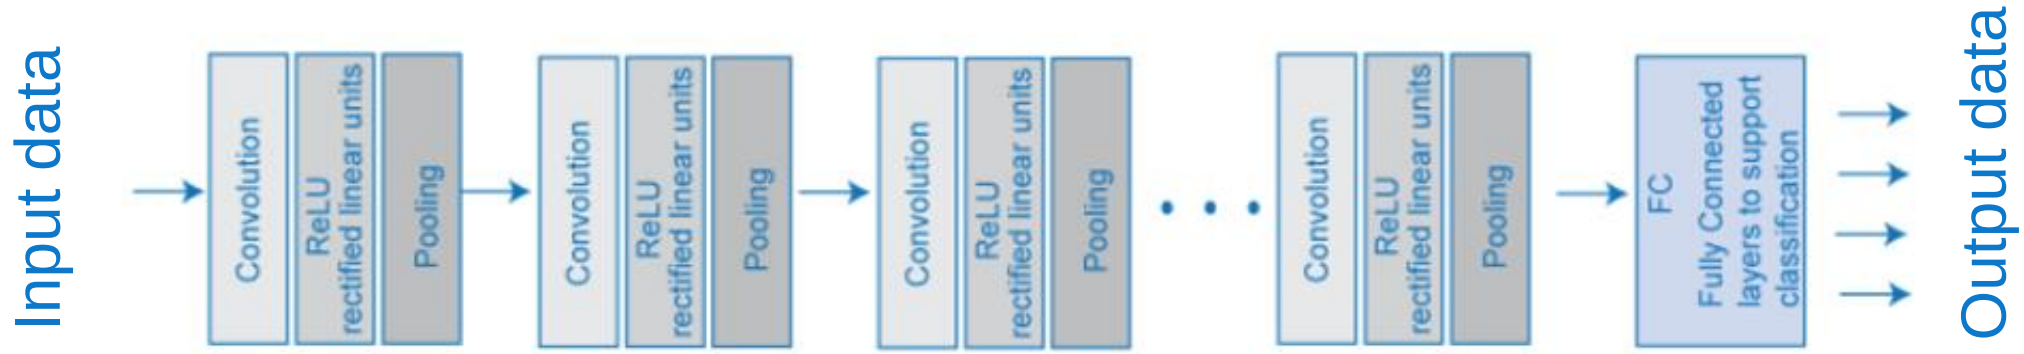

Another Network for Signals - LSTM

- LSTM = Long Short Term Memory (Networks)
  - Signal, text, time-series data
  - Use previous data to predict new information
- I live in France. I speak _____.

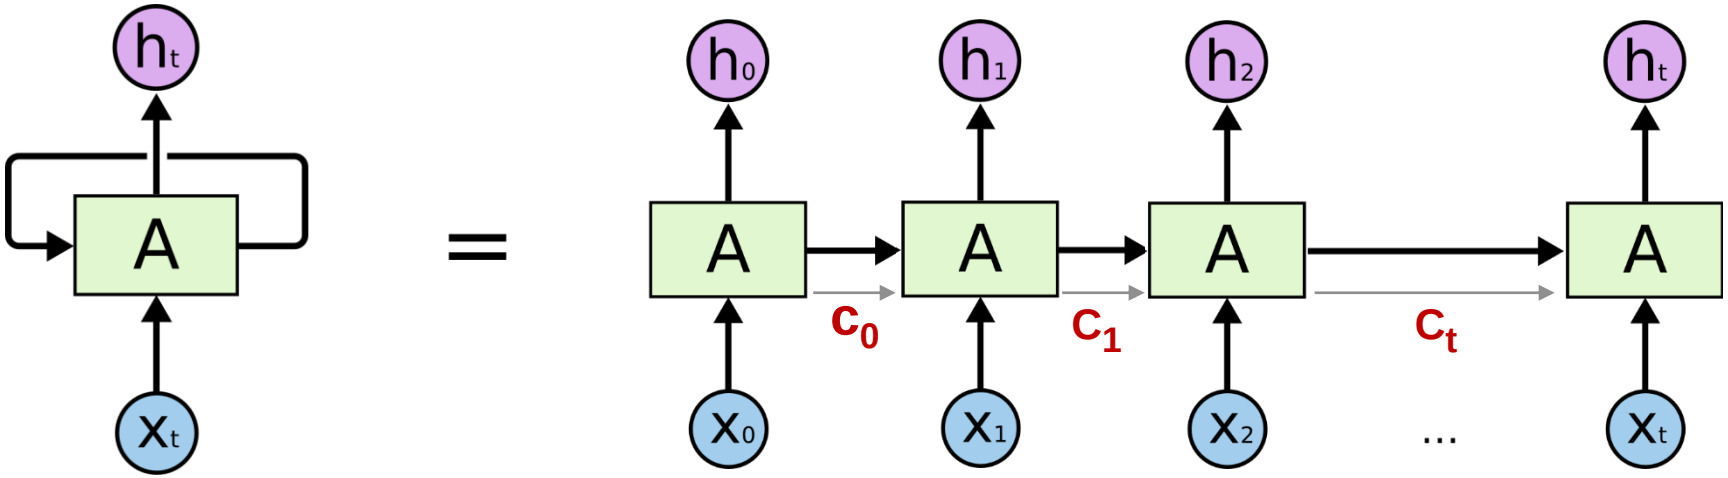

Long Short-Term Memory (LSTM)

- LSTM are an extension of Recurrent Neural Networks.
- RNN can handle arbitrary input/output lengths.
- They have the capability to use the dependencies among inputs.
- LSTMs just like every other RNN **connect through time.** They are capable of preserving the long-term and short-term dependencies that occur within data.

Example: Algorithmic Trading

Run LSTM networks

Define LSTM layers

```
layers = [sequenceInputLayer(nFeatures)
    lstmLayer(256, 'OutputMode', 'last')
    fullyConnectedLayer(1)
    regressionLayer]
```

Set training options

```
miniBatchSize = 250;
opts = trainingOptions('sgdm',...
    'Plots','training-progress', ...
    'LearnRateSchedule', 'piecewise', ...
    'shuffle', 'never',...
    'InitialLearnRate', 0.01,...
    'MiniBatchSize', miniBatchSize,...
    'MaxEpochs', 20);
```

Train networks

```
net = trainNetwork(XTrain, YTrain, layers, opts);
```

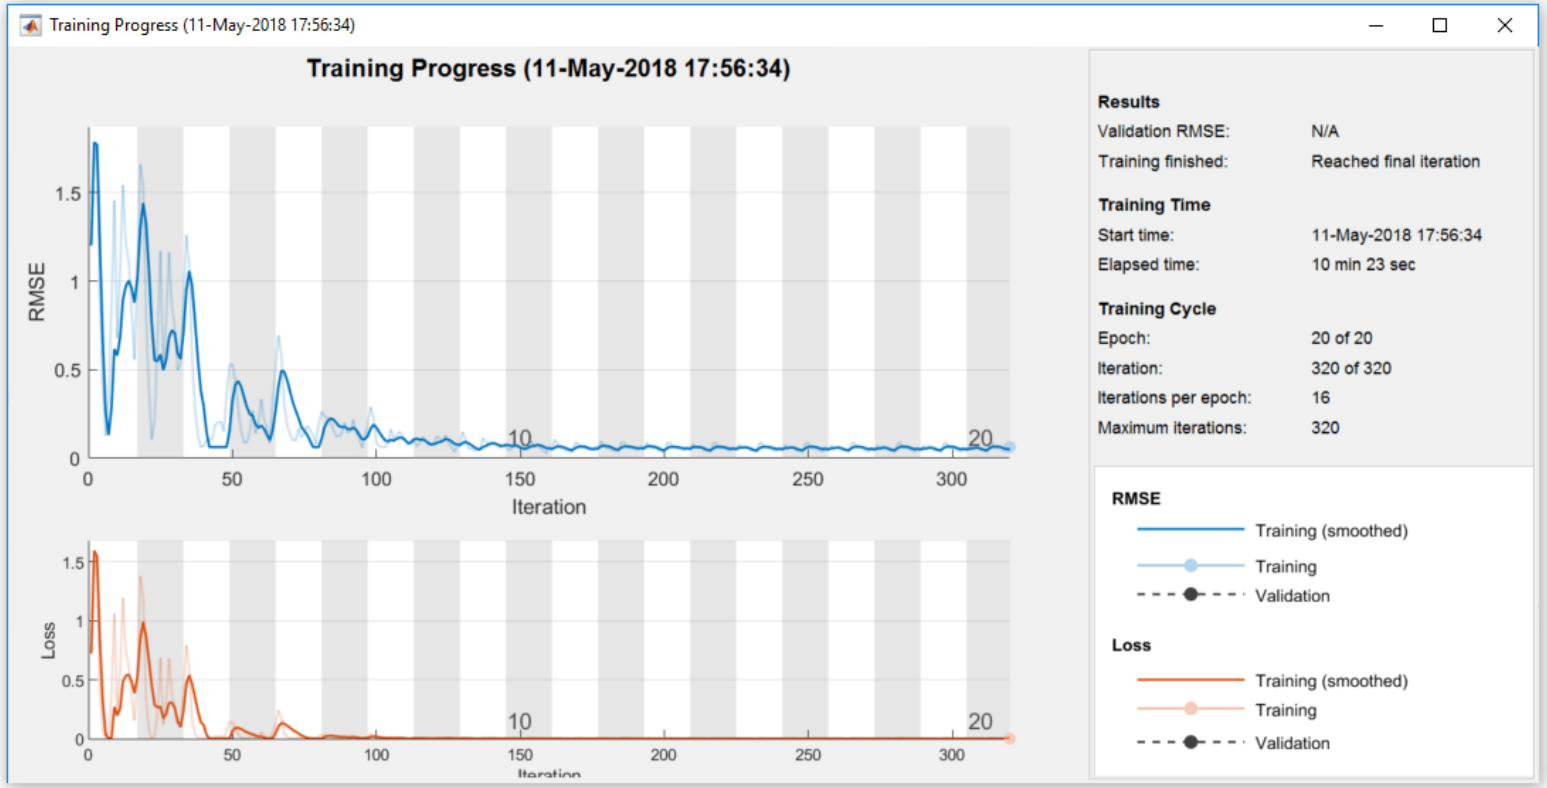

Another Application: Sentiment Analysis with Twitter Data

Access Tweets → **Preprocess Tweets** → **Develop Model** → **Predict Sentiment**

Preprocess Tweets: **Clean-up Text** → **Convert to Numeric**

Apple's iPhone 8 to Drive 9.1% Increase in Shipments Per IDC https://t.co/n085F65upk $AAPL $GRMN $GOOG

apples iphone drive increase shipments per idc

	apple	iphone	increase	sell
tweet1	1	1	1	0
tweet2	1	0	0	1

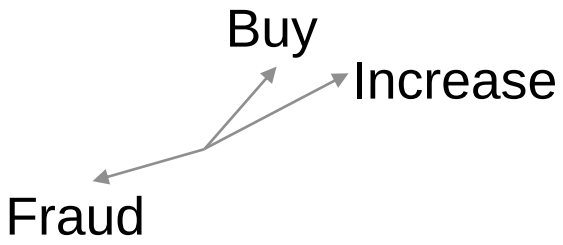

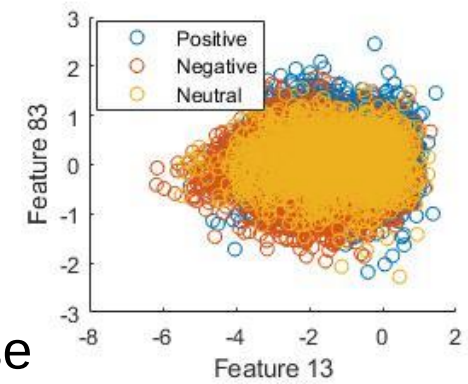

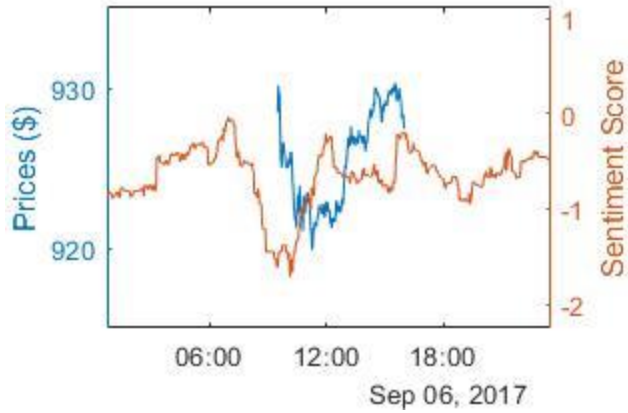

	Track 1
13:30	New Financial Modelling Capabilities in MATLAB *Kevin Shea and Stuart Kozola, MathWorks*
14:15	Building an Online Quantitative Lecture in MATLAB *Dr. Ioannis Kyriakou, Cass Business School*
14:45	Break
15:15	Using MATLAB for Sentiment Analysis and Text Analytics *Liliana Agapito de Sousa Medina, MathWorks*

Deep Learning on CPU, GPU, Multi-GPU and Clusters

Single CPU

Single CPU
Single GPU

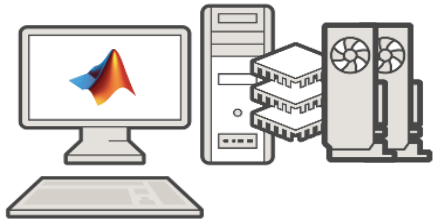

Single CPU, Multiple GPUs

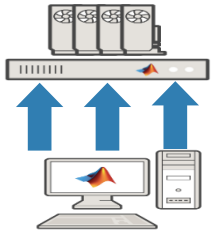

On-prem server with GPUs

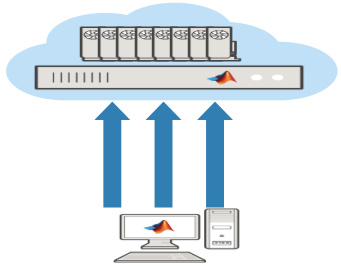

Cloud GPUs (AWS)

How to target?

```
opts = trainingOptions('sgdm', ...
    'MaxEpochs', 100, ...
    'MiniBatchSize', 250, ...
    'InitialLearnRate', 0.00005, ...
'ExecutionEnvironment', 'auto' );
```

```
opts = trainingOptions('sgdm', ...
    'MaxEpochs', 100, ...
    'MiniBatchSize', 250, ...
    'InitialLearnRate', 0.00005, ...
'ExecutionEnvironment', 'multi-gpu' );
```

```
opts = trainingOptions('sgdm', ...
    'MaxEpochs', 100, ...
    'MiniBatchSize', 250, ...
    'InitialLearnRate', 0.00005, ...
'ExecutionEnvironment', 'parallel' );
```

GPU Coder

- Automatically generates **CUDA** Code from MATLAB Code
  - can be used on NVIDIA GPUs

- CUDA extends C/C++ code with constructs for parallel computing

Agenda

	Artificial Intelligence enabled by Machine and Deep Learning
	Machine Learning
	Deep Learning
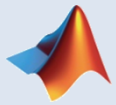	**Outlook: Integration in Production Systems**

Integrate with Production Systems

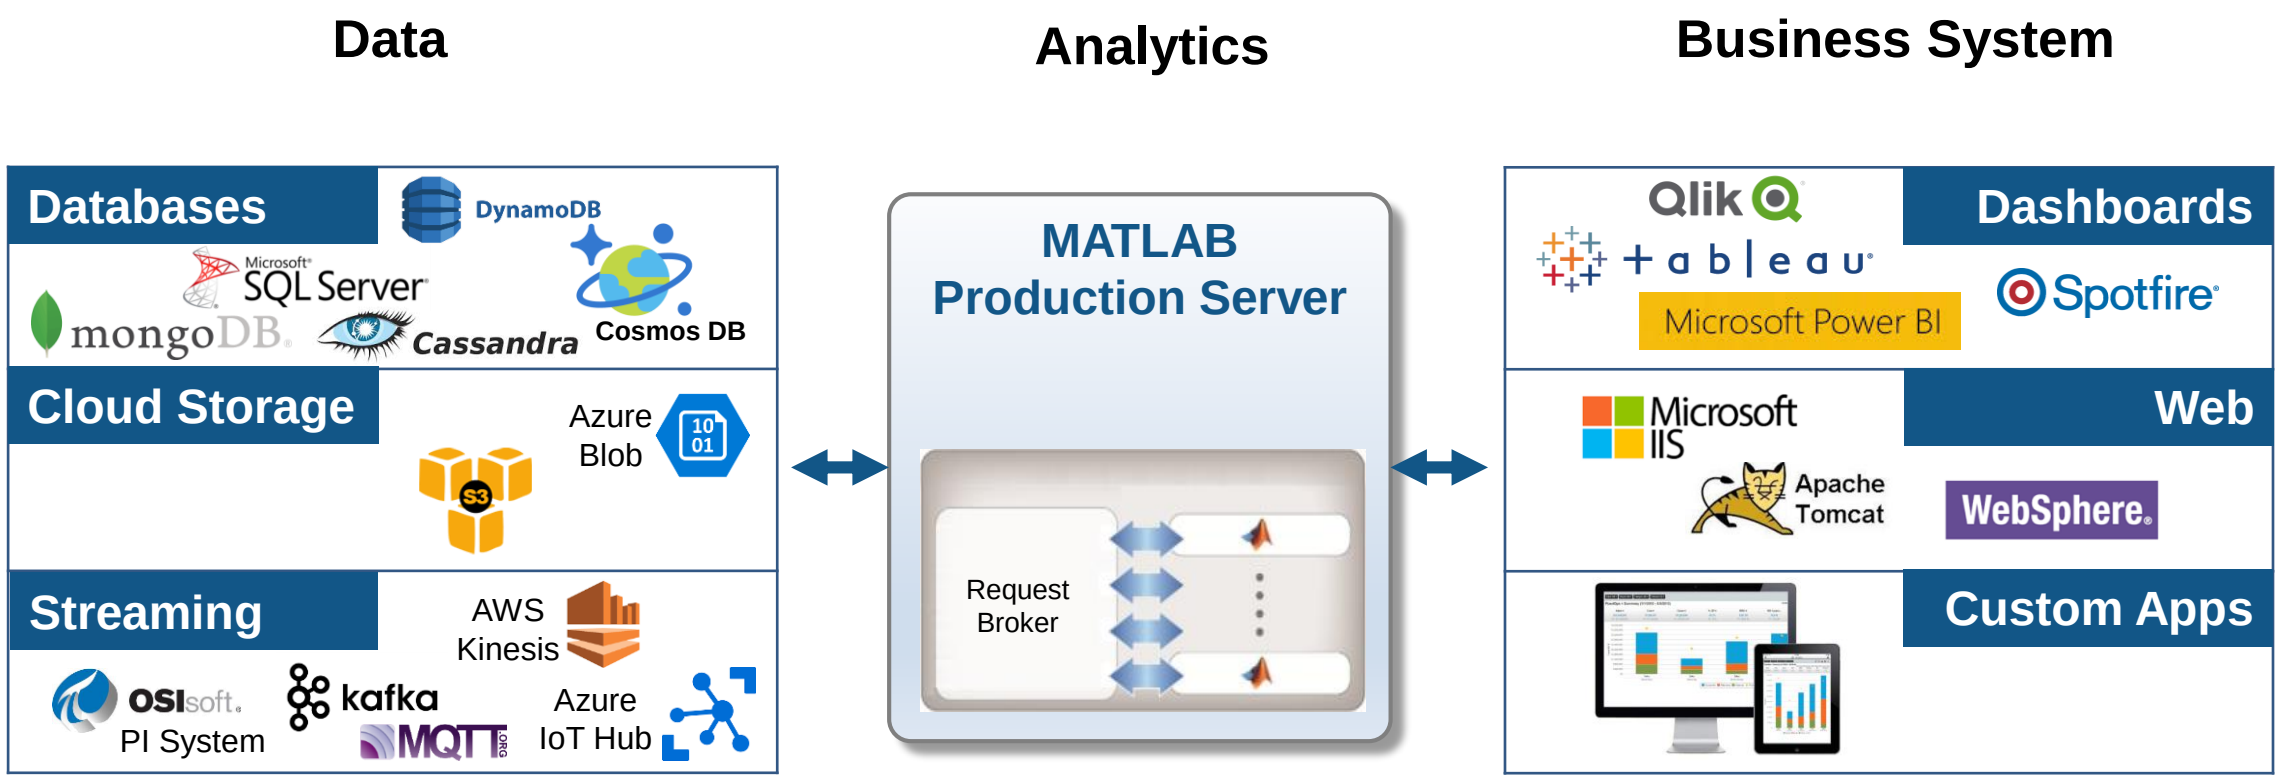

Thank you for your attention