HPE POINTNEXT

AI thoughts

Athanassios Liakopoulos
Global Practice for Network, Workplace & Edge

GRNOG #8

Hewlett Packard Enterprise

Who am I?

— I am not an AI expert!

— Just sharing my thoughts

— Not assigned to any AI role in my organization

- Leading solutions for Enterprise IoT

Outline

- AI use cases
- Ethics
- Machine Learning
  - Basic terms, phases, reality, etc
- Reinforcement learning - Lets play Atari!
- CNN - Where is the cat?
- ML in networking / PdM

“The illiterate of the 21st century will not be those who cannot read and write, but those who cannot learn, unlearn, and relearn.”

Alvin Toffler

Backwards Brain Bicycle

Knowledge vs Understanding

https://ed.ted.com/featured/bf2mRAfC

AI examples

AI use cases

- Deep Blue vs Garry Kasparov (1996-1997)
- Autonomous cars (Tesla Autopilot 2014)
- Google DeepMind plays Atari Breakout (2015)
- AlphaGo (2016)
- AI vs Doctors
  - Skin cancer (2017), Cardiography (2017) , MRI (2019)
- “My personal IoT experience” (2017)
- Uber accident (2018)

AI ethics

Ethics challenges

– Self-driving Uber taxi accident

– Accept credit based on customer (vs financial) profile

– Admission to the university based on non-academic (vs academic) profile

– Get private insurance based on your gene profile

– Get employed based on …

– Fake news - Stop Misinformation and False News

  – Responsibility fall in Intenet giants, e.g. Facebook, etc

Big Data – a tool for inclusion or exclusion?

Areas of concern

- Protected classification
  - Race, ethnicity, national origin, gender, religion, etc
- Biased data
  - Data input is biased
  - Use of irrelevant data (e.g. zip code)
- Tension between *Individual Rights* and *Organisational Rights*
  - The AI algorithm is an intellectual property – Need to be protected
- Negligence laws
  - Who's fault is the Uber taxi accident?
    - AI programmer? Car provider? Service provider? Taxi driver? Other?
- Customer analytics
  - Consumer Reporting Agencies (CRAs)
    - Implement reasonable procedures to insure a maximum possible accuracy of consumer reports
    - Provide consumers with access to their information
    - Provide consumers with the ability to correct errors
  - Unfair and deceptive practices can arise when companies are maintaining large amounts of sensitive data about individuals and are not securing it from misuse

- General data protection regulation (GDPR)
  - Privacy is a fundamental right for people in Europe
- AI & GDPR
  - Big data analytics need to be fair
  - Permission to process
  - Purpose limitation
  - Holding on to data
  - Accuracy
  - Individual rights and access to data
  - Security measures and risks
  - Accountability
  - Controllers and processes

Basic terms

Artificial intelligence

Predict and anticipate possible future events

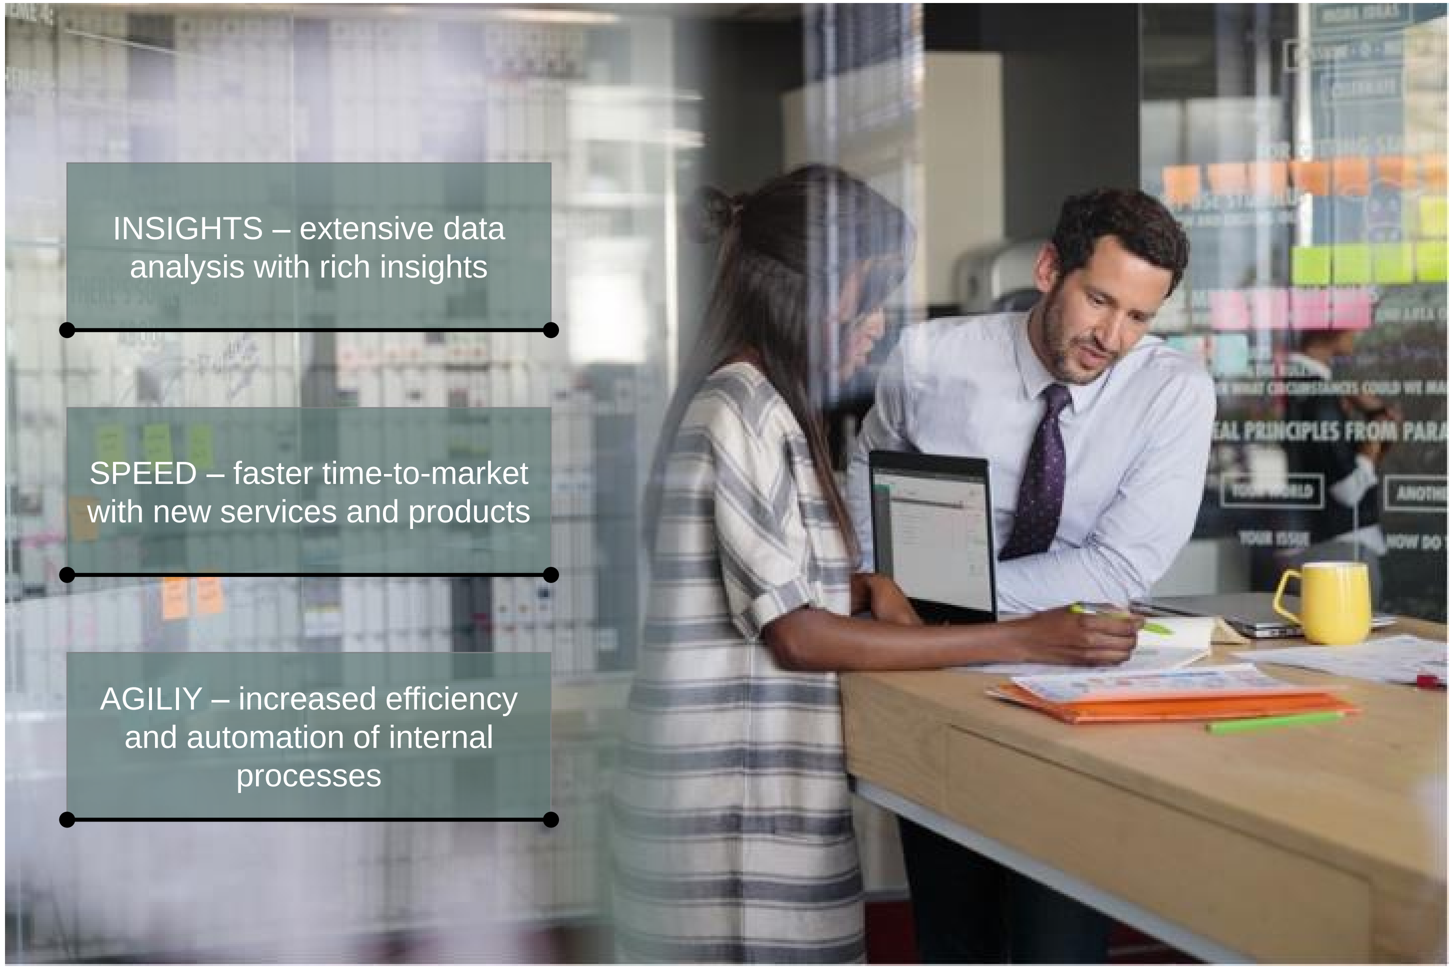

INSIGHTS – extensive data analysis with rich insights

SPEED – faster time-to-market with new services and products

AGILIY – increased efficiency and automation of internal processes

Artificial intelligence is a technology that appears to emulate human performance typically by learning, coming to its own conclusions, appearing to understand complex content, and engaging in natural human-like actions.

Let's get grounded...what is AI?

What makes a machine intelligent?

Artificial intelligence (AI)

Mimics human behavior. Any technique that enables machines to solve a task in a way like humans do.

Example:
Siri

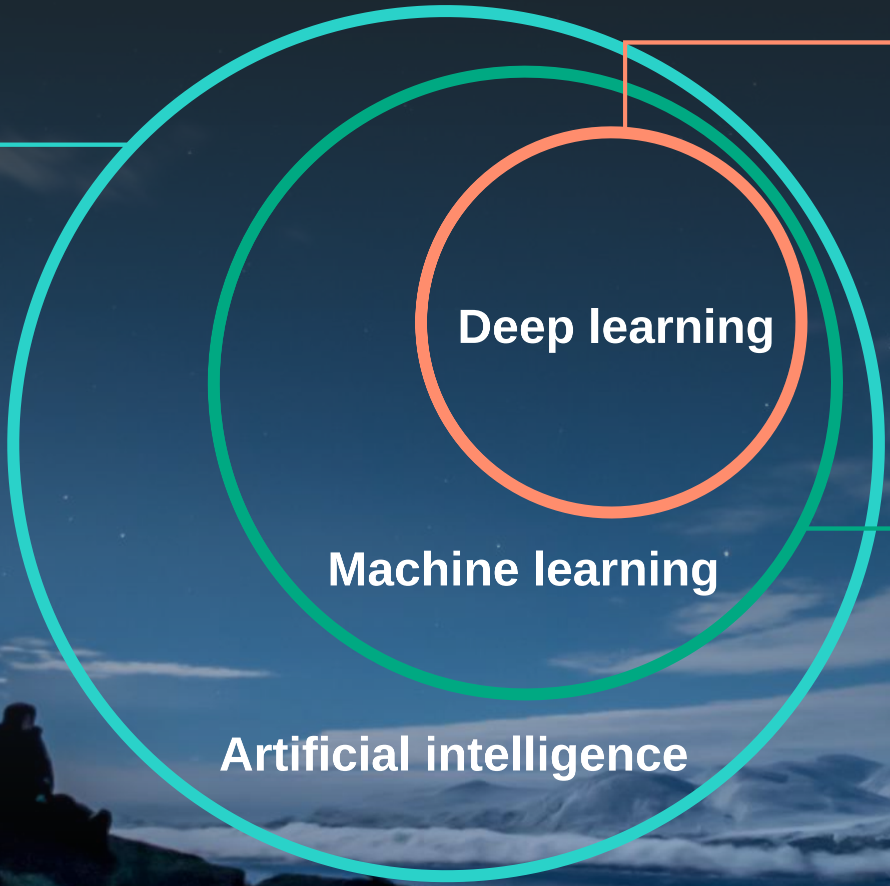

Deep learning (DL)

Subset of ML, using deep artificial neural networks as models, inspired by the structure and function of the human brain.

Example:
Self-driving car

Machine learning (ML)

Algorithms that allow computers to learn from examples without being explicitly programmed.

Example:
Google Maps

Why should you be interested in AI / advanced analytics?

Our customers want AI and advanced analytics….

AI and advanced analytics is one of the
top 3 priorities
for CIOs

AI and advanced analytics infrastructure could constitute
15-20%
of the market by 2021[1]

….but face many challenges

Use cases

New roles, skill gaps

Culture and change

Data preparation

Legacy Infrastructure

[1] IDC. Goldman Sachs. HPE Corporate Strategy.2018

Data explosion and need for advanced analytics

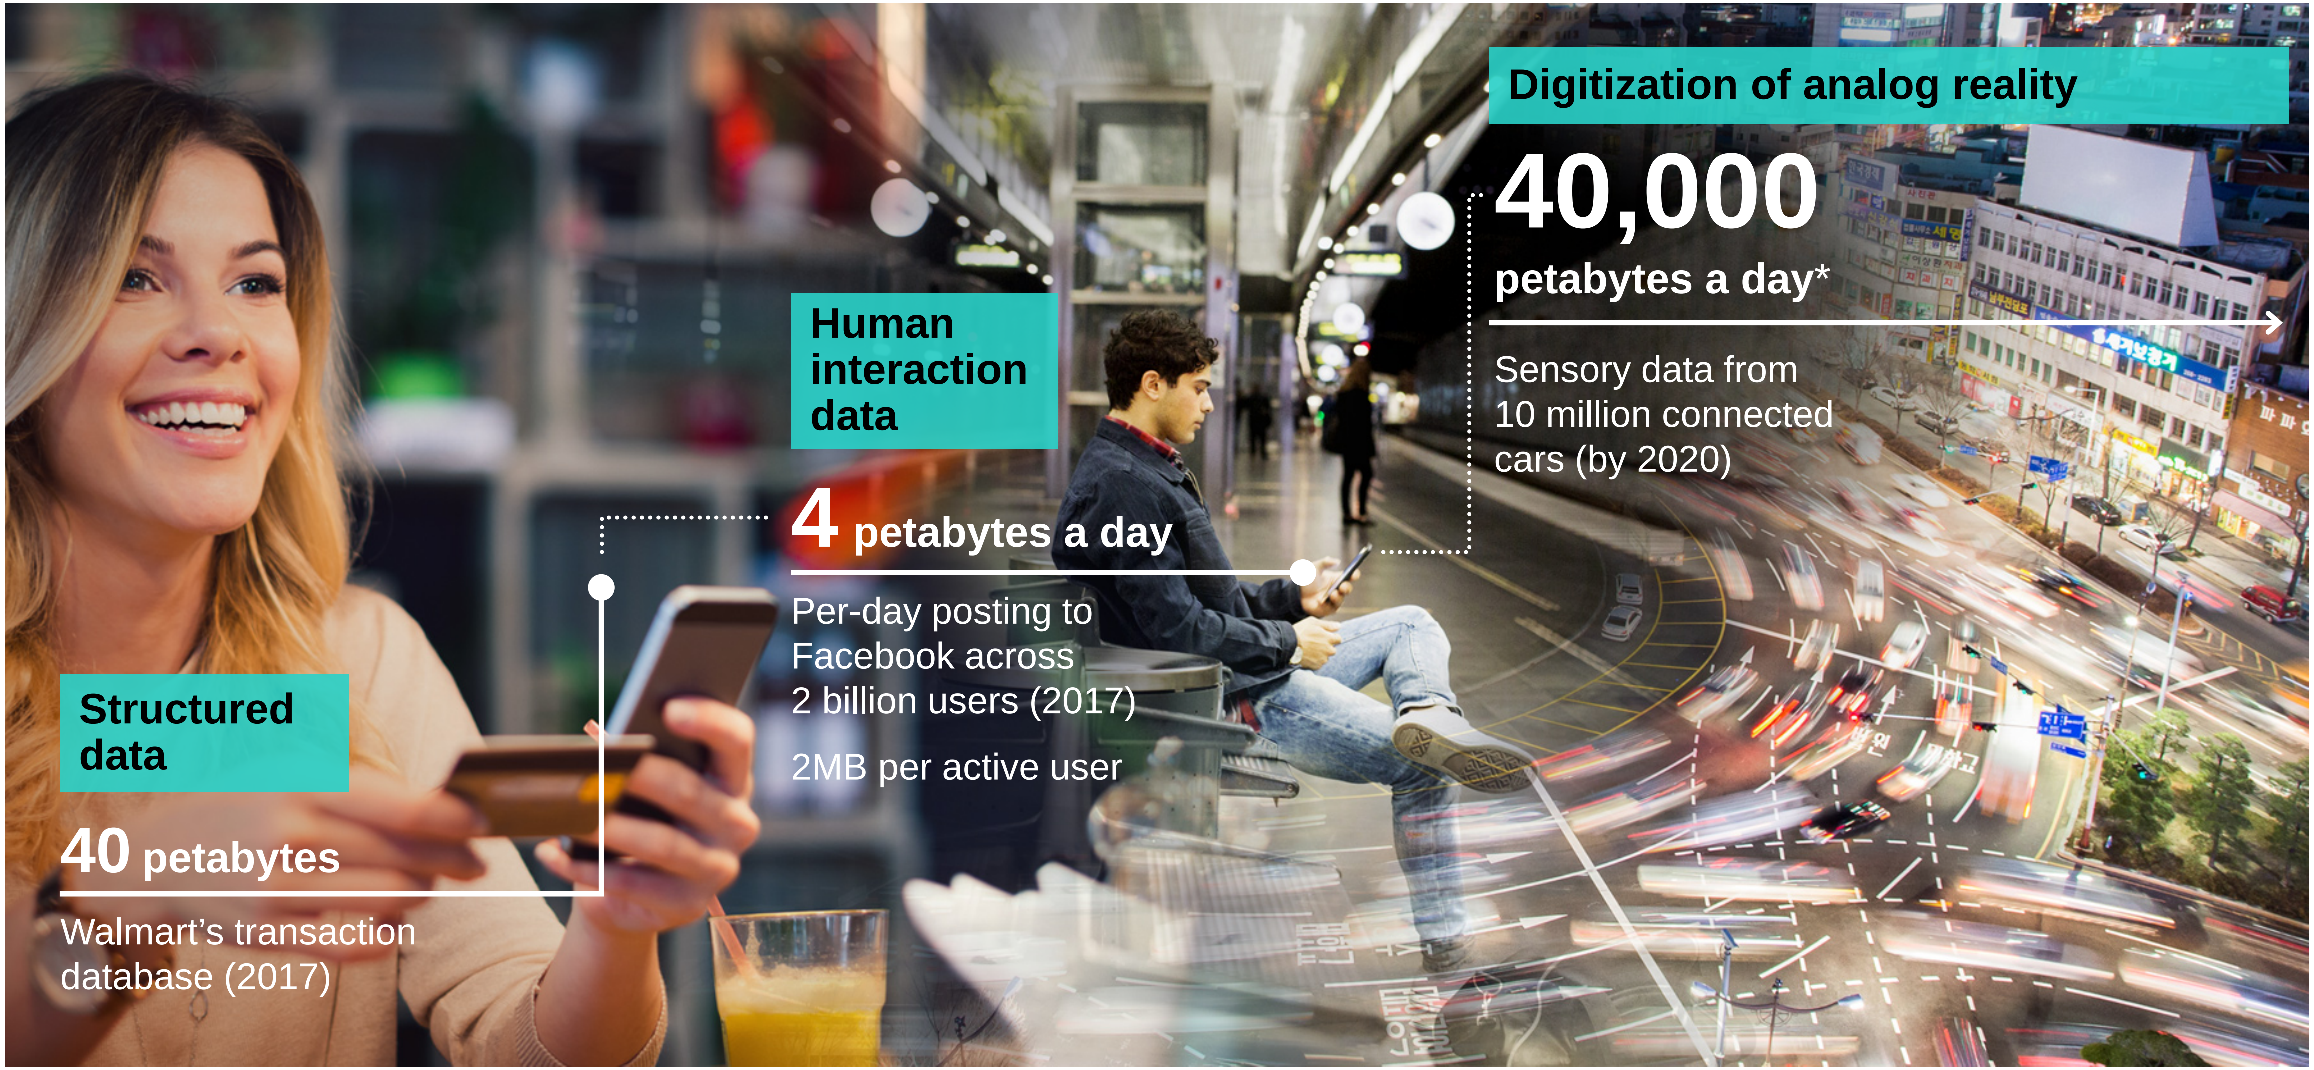

Structured data

40 petabytes

Walmart’s transaction database (2017)

Human interaction data

4 petabytes a day

Per-day posting to Facebook across 2 billion users (2017)

2MB per active user

Digitization of analog reality

40,000 petabytes a day*

Sensory data from 10 million connected cars (by 2020)

* Driver assistance systems only

Taking a holistic view to artificial intelligence

AI Technologies

Machine, Deep Learning
Natural-Language Processing
Voice, Face Recognition
Robots and Sensors

AI Use Cases

Customer Service
Internet of Things
Security and Fraud
Digital Workplace

AI Applications

Bots, Chatbots
Virtual Assistants
Predictive Analytics
Smart Objects, Sensors

AI People

Data Scientists, Data Officers
Developers
Enterprise Architects
Businesses

Key questions remain

What opportunities does AI bring to your business? What are the major use cases?

How do you get started with gaining intelligence with your data?

What is the best way to prepare your company culture for data centric and AI future?

How do you integrate your AI and data ecosystem for advanced analytics?

How do you modernize and prepare your EDW or Hadoop big data foundation for AI?

Machine Learning

HPE
POINTNEXT

ML Analogy

How long am I going to live?

– Do I have a balanced diet?

– Do I have the right BMI?

– Do I exercise?

– Did my parents lived beyond 90?

– Do I use seat belt when I drive?

– Do I know to swim?

– Do I smoke?

– Do I take drugs?

– Do I climb to Everest?

– Am I F1 driver?

– Do I travel to Syria?

ML basics

Supervised Machine Learning

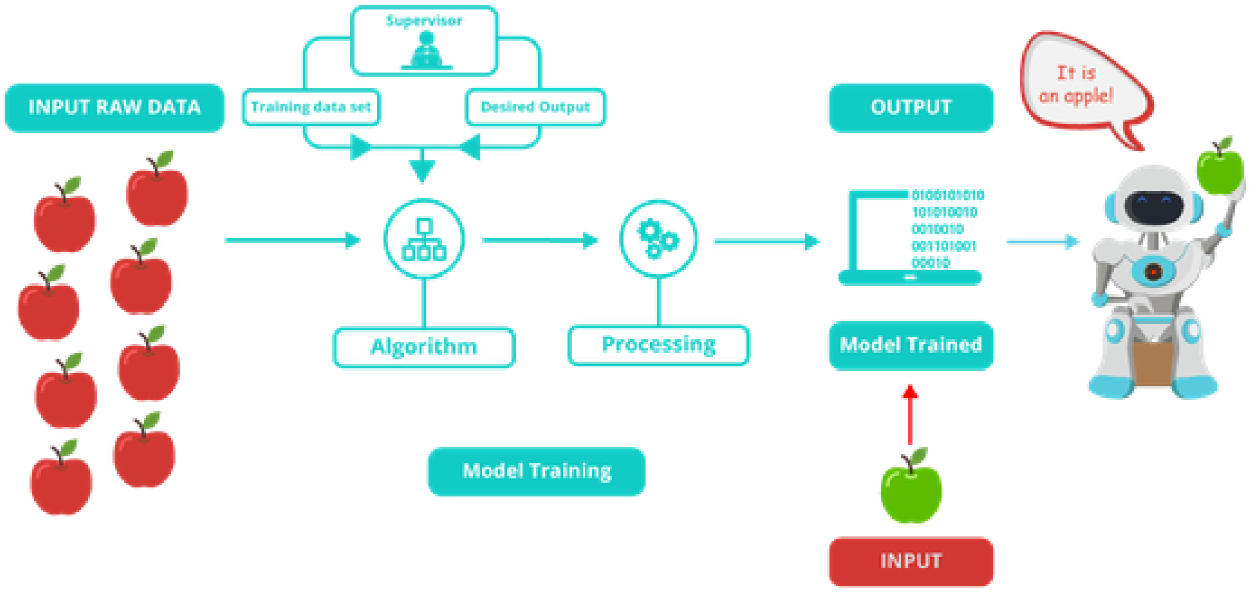

Unsupervised Machine Learning

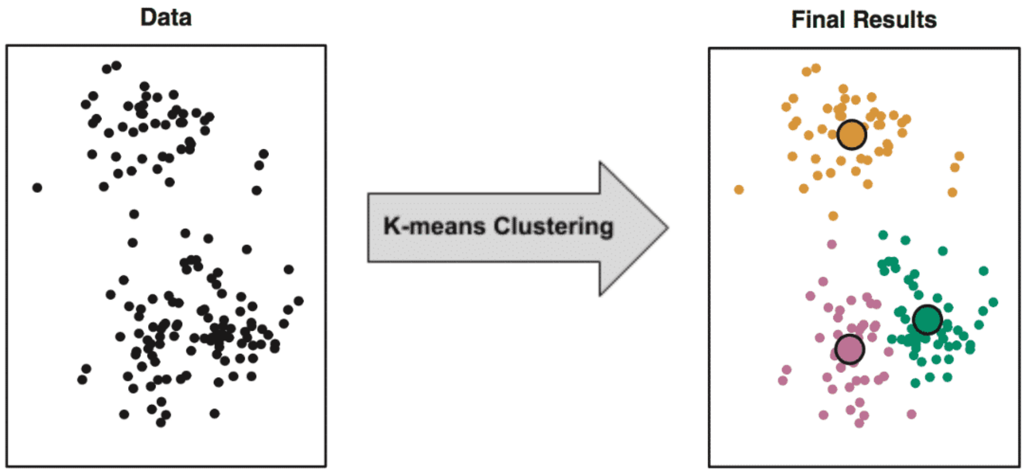

ML basics

Classification

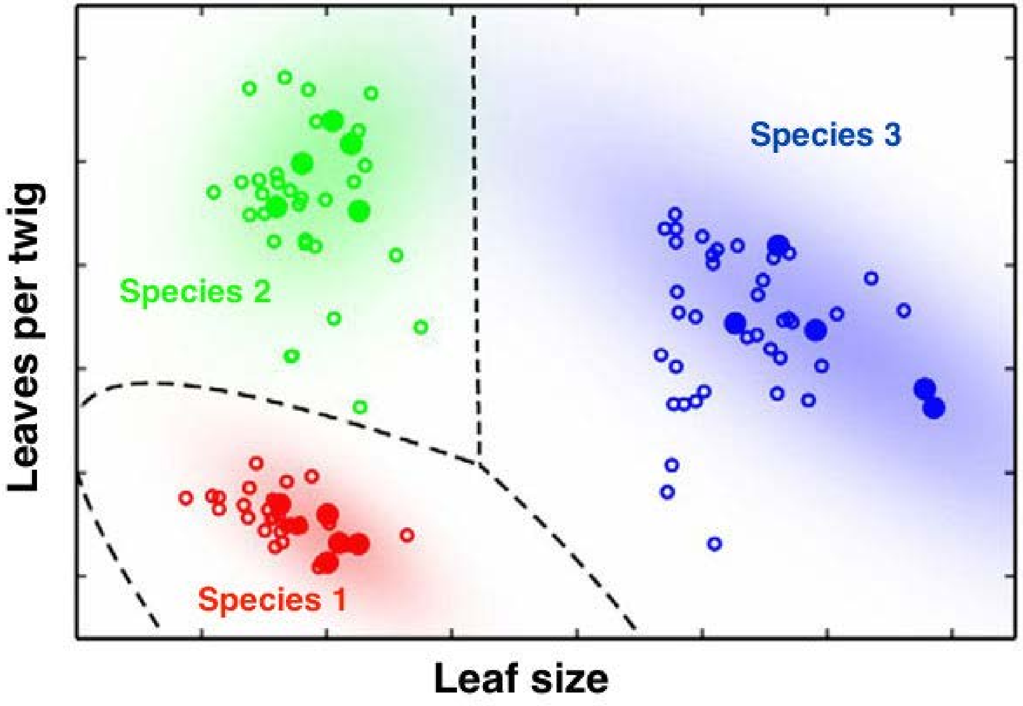

Regression

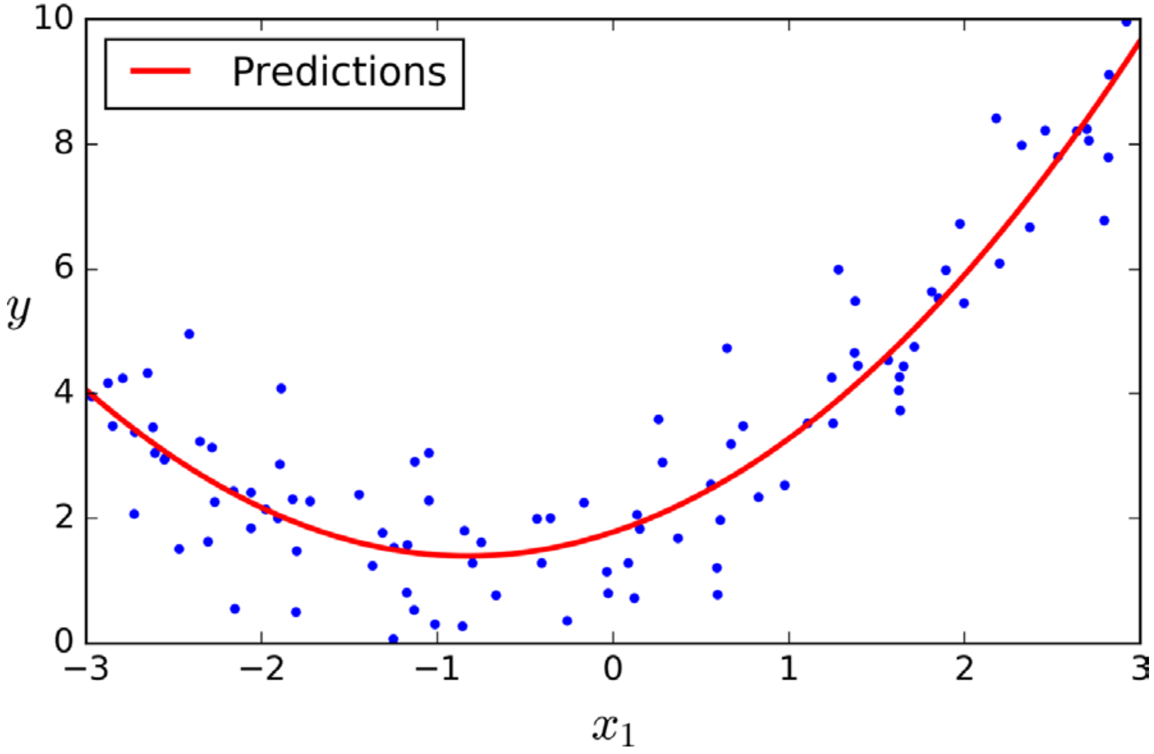

Overfitting

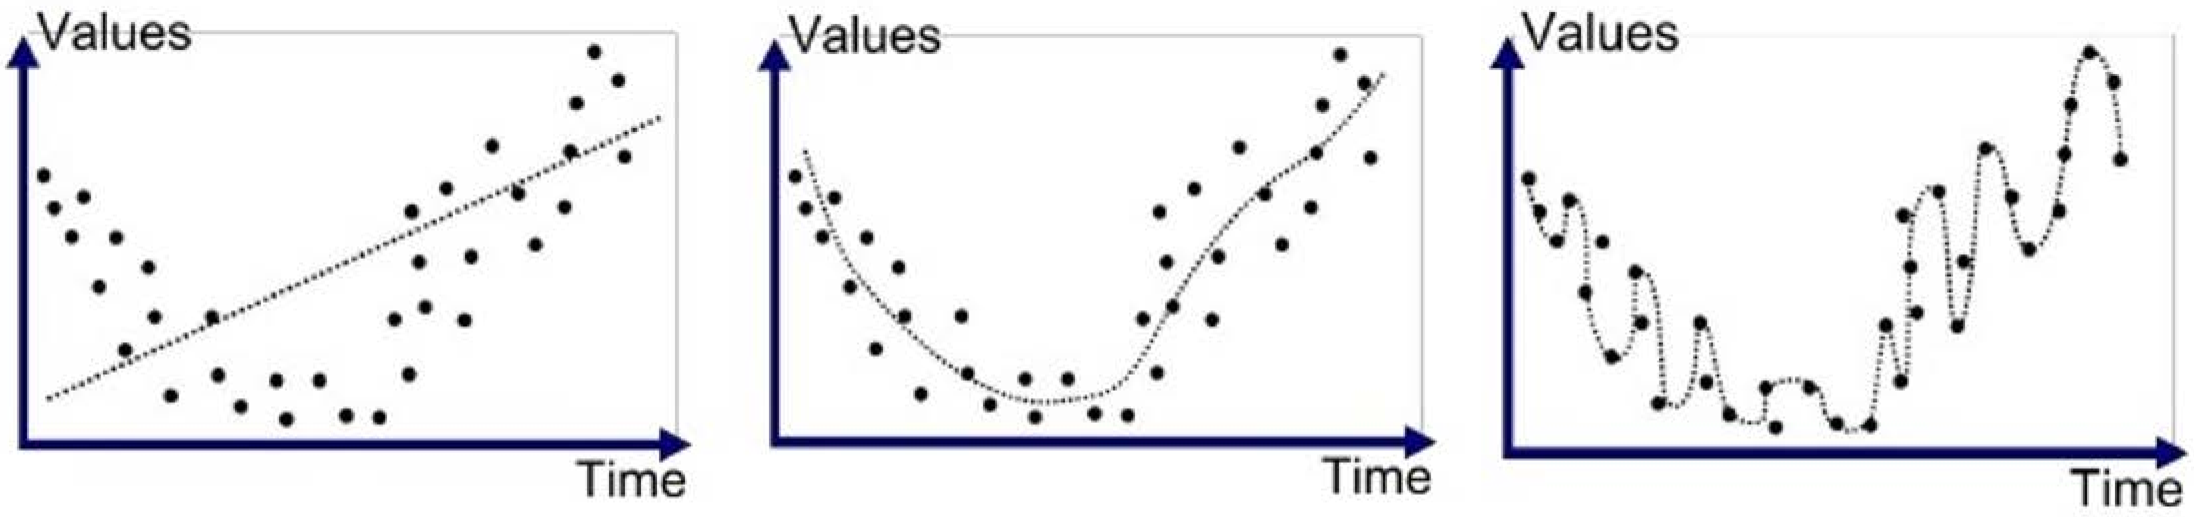

Underfitted | Good Fit/Robust | Overfitted

Overfitting & Model complexity

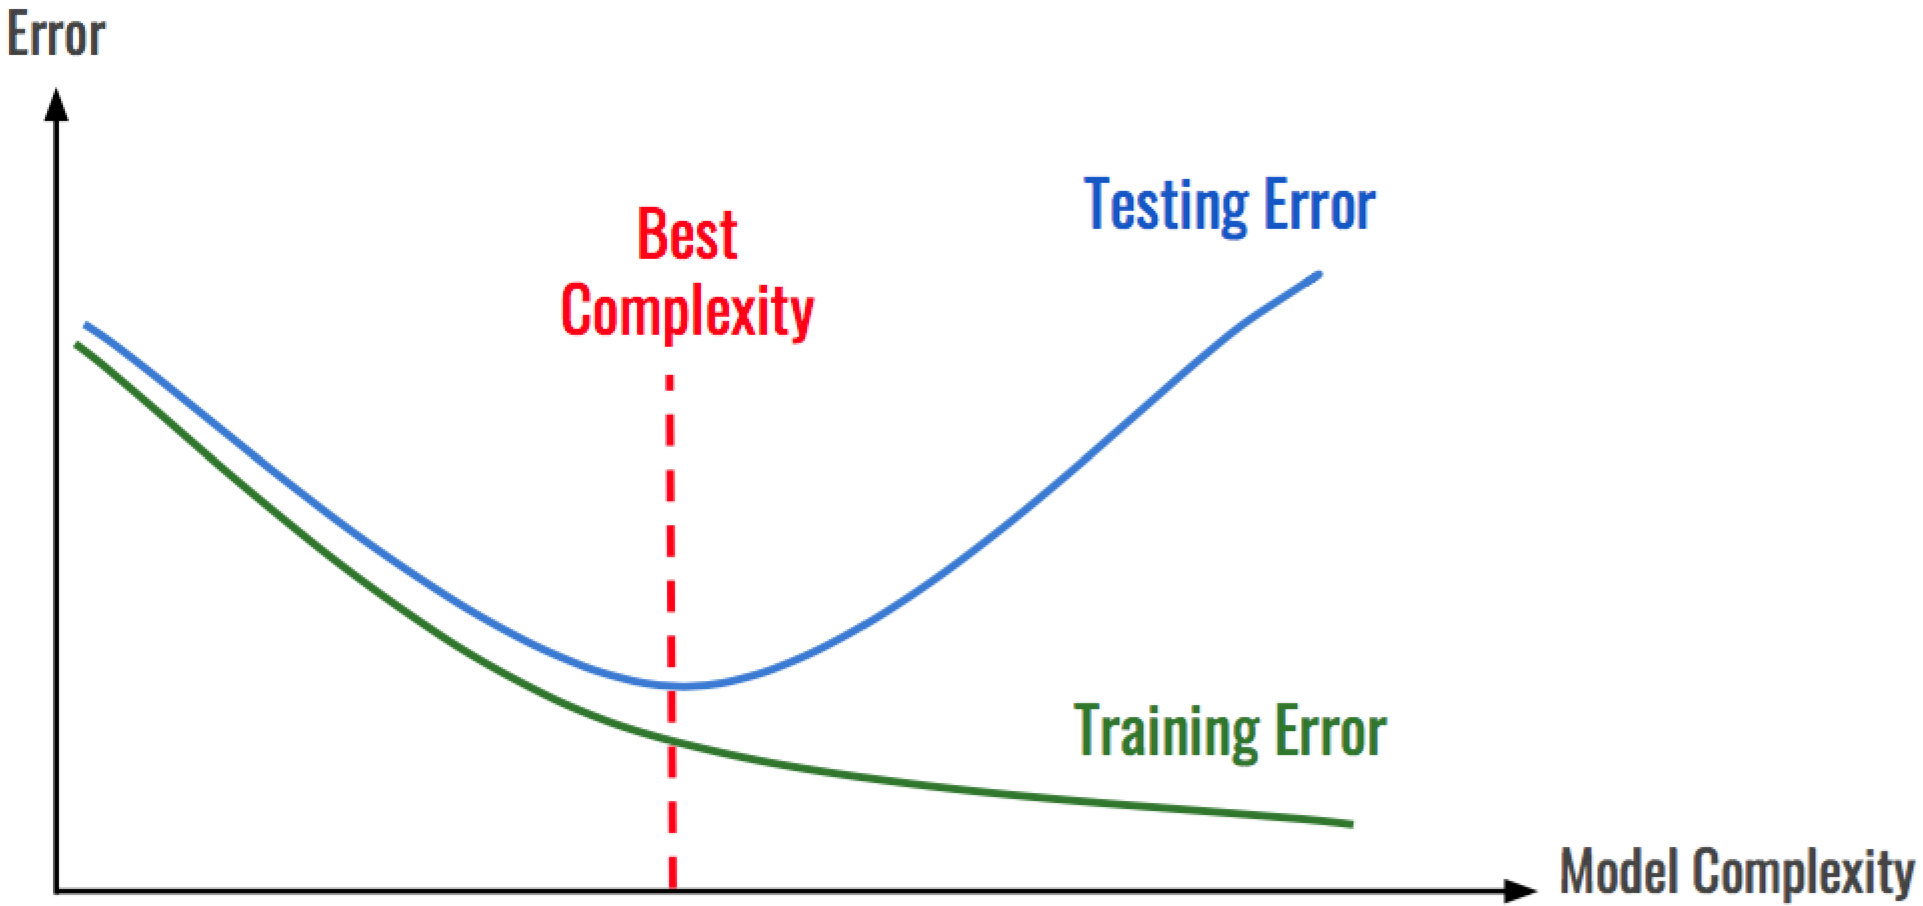

Data (example)

	Daily Time Spent on Site	Age	Area Income	Daily Internet Usage	Ad Topic Line	City	Male	Country	Timestamp	Clicked on Ad	LABEL
0	68.95	35	61833.90	256.09	Cloned 5thgeneration orchestration	Wrightburgh	0	Tunisia	2016-03-27 00:53:11	0	L1
1	80.23	31	68441.85	193.77	Monitored national standardization	West Jodi	1	Nauru	2016-04-04 01:39:02	0	L2
2	69.47	26	59785.94	236.50	Organic bottom-line service-desk	Davidton	0	San Marino	2016-03-13 20:35:42	0	L2
3	74.15	29	54806.18	245.89	Triple-buffered reciprocal time-frame	West Terrifurt	1	Italy	2016-01-10 02:31:19	0	L1
4	68.37	35	73889.99	225.58	Robust logistical utilization	South Manuel	0	Iceland	2016-06-03 03:36:18	0	L3
5	59.99	23	59761.56	226.74	Sharable client-driven software	Jamieberg	1	Norway	2016-05-19 14:30:17	0	L1
6	88.91	33	53852.85	208.36	Enhanced dedicated support	Brandonstad	0	Myanmar	2016-01-28 20:59:32	0	L1
7	66.00	48	24593.33	131.76	Reactive local challenge	Port Jefferybury	1	Australia	2016-03-07 01:40:15	1	L3
8	74.53	30	68862.00	221.51	Configurable coherent function	West Colin	1	Grenada	2016-04-18 09:33:42	0	L2
9	69.88	20	55642.32	183.82	Mandatory homogeneous architecture	Ramirezton	1	Ghana	2016-07-11 01:42:51	0	L2

AI Algorithms

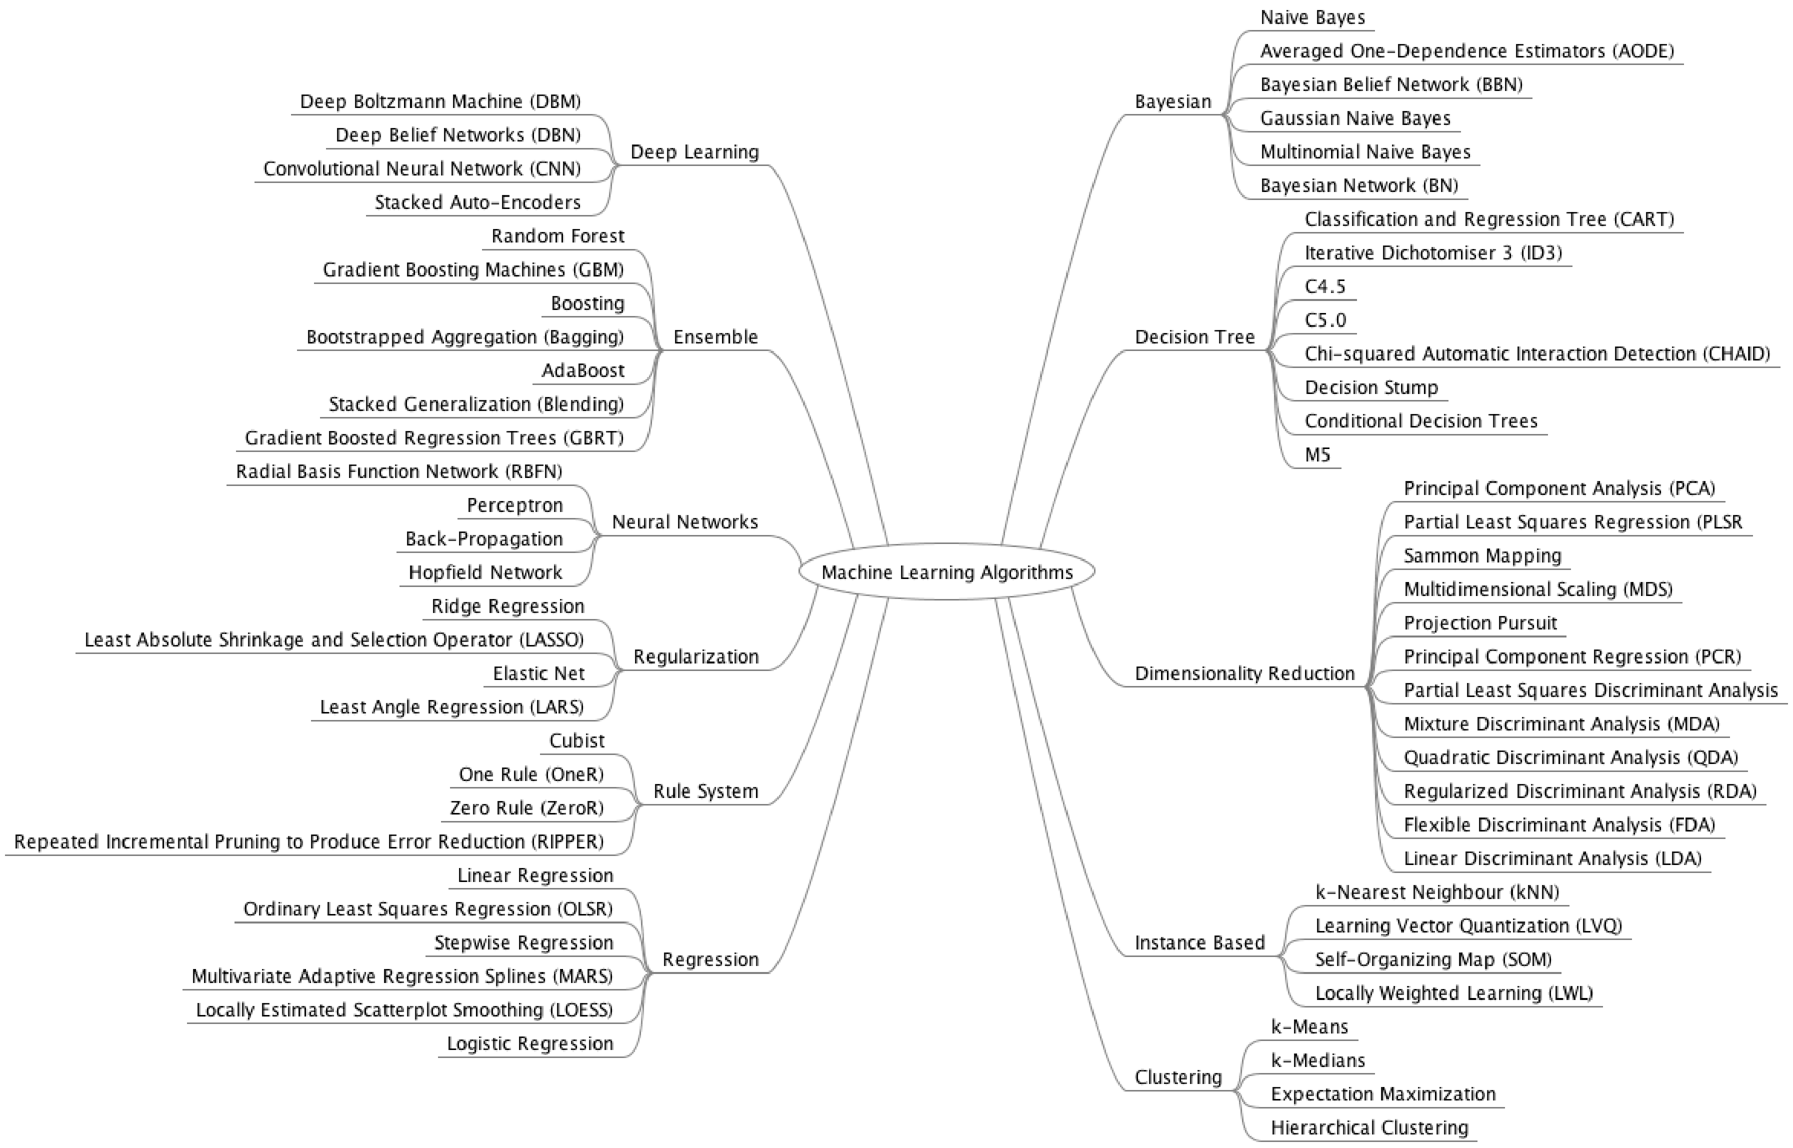

HPE POINTNEXT

ML phases

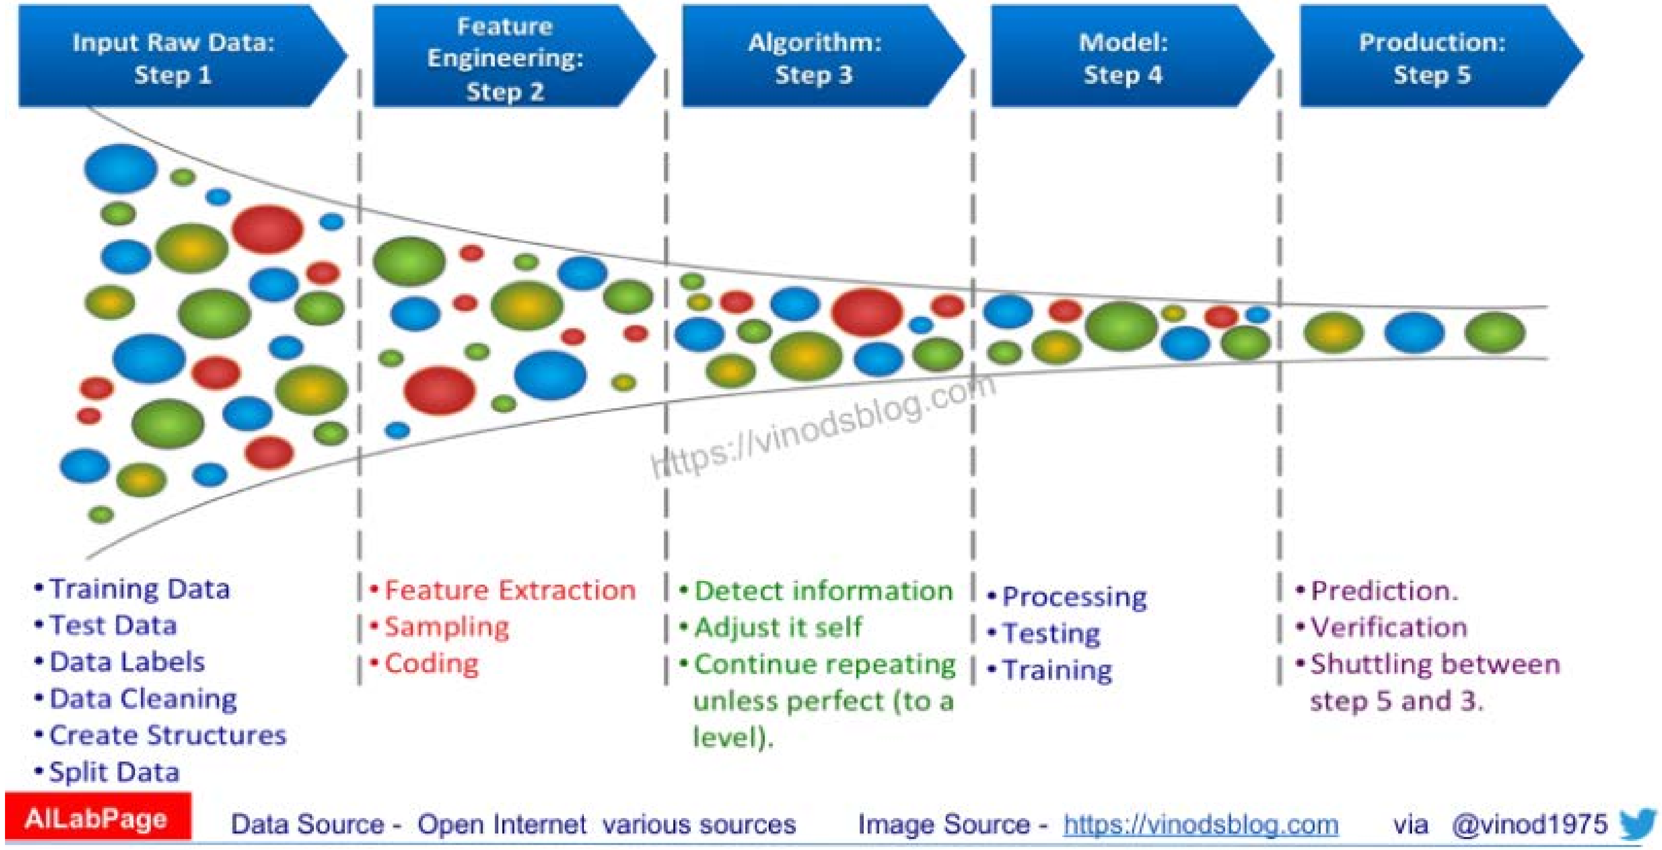

- Data preparation,
- Feature engineering
- Curse of dimensionality
- Models are two lines

It takes more than having tomatoes to build a Greek salad

Neural Networks

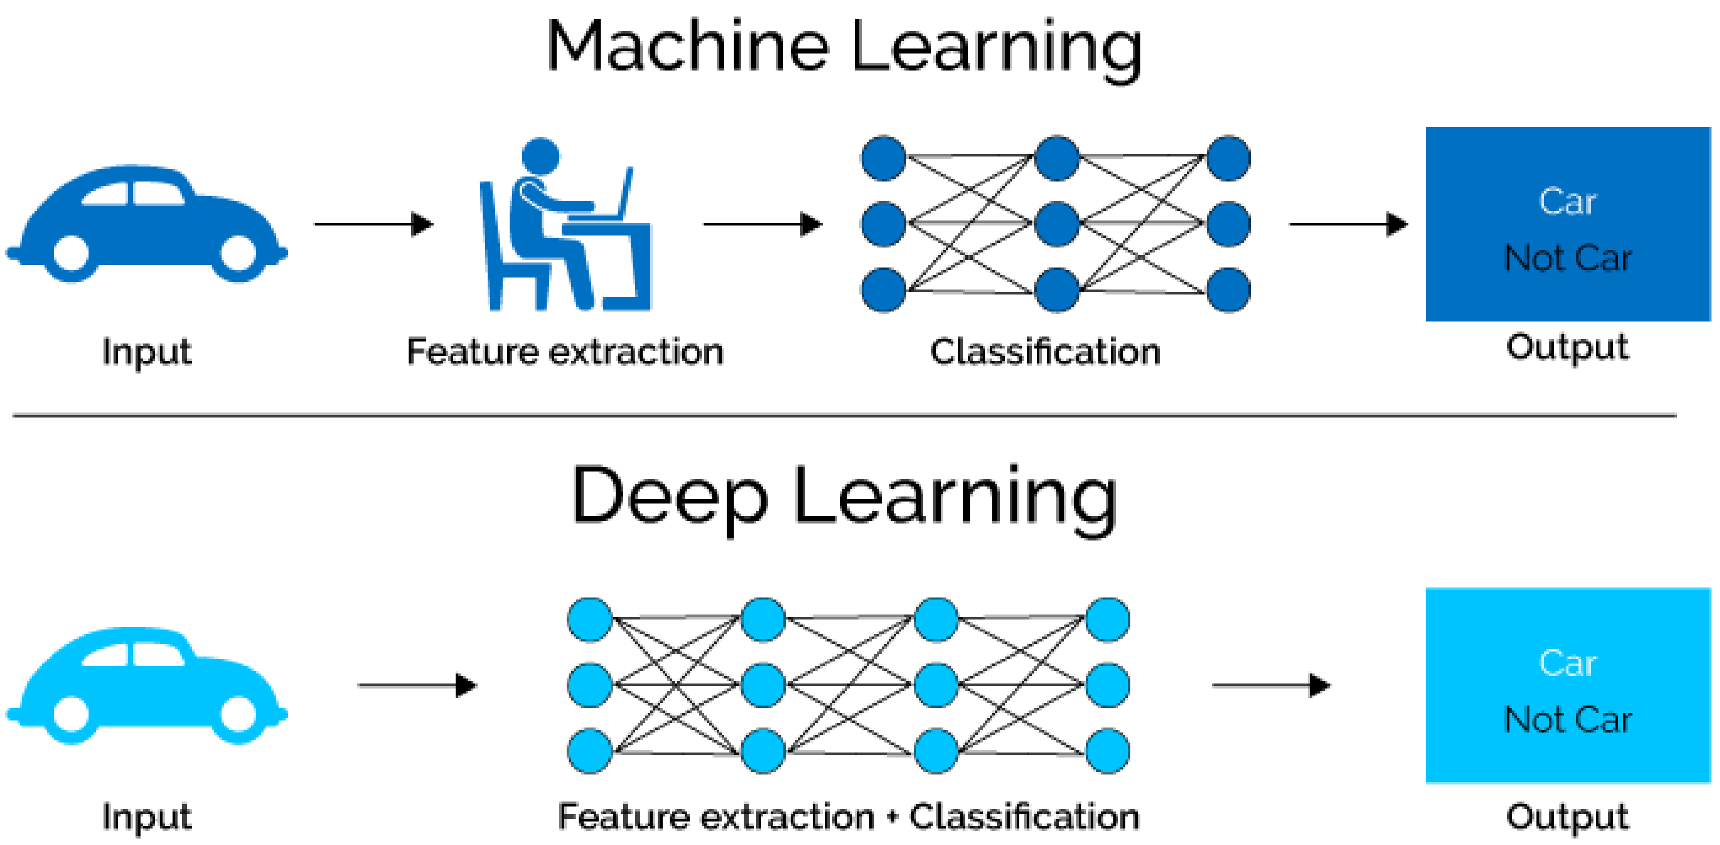

Machine Learning
Car
Not Car
Input
Feature extraction
Classification
Output
Deep Learning
Car
Not Car
Input
Feature extraction + Classification
Output

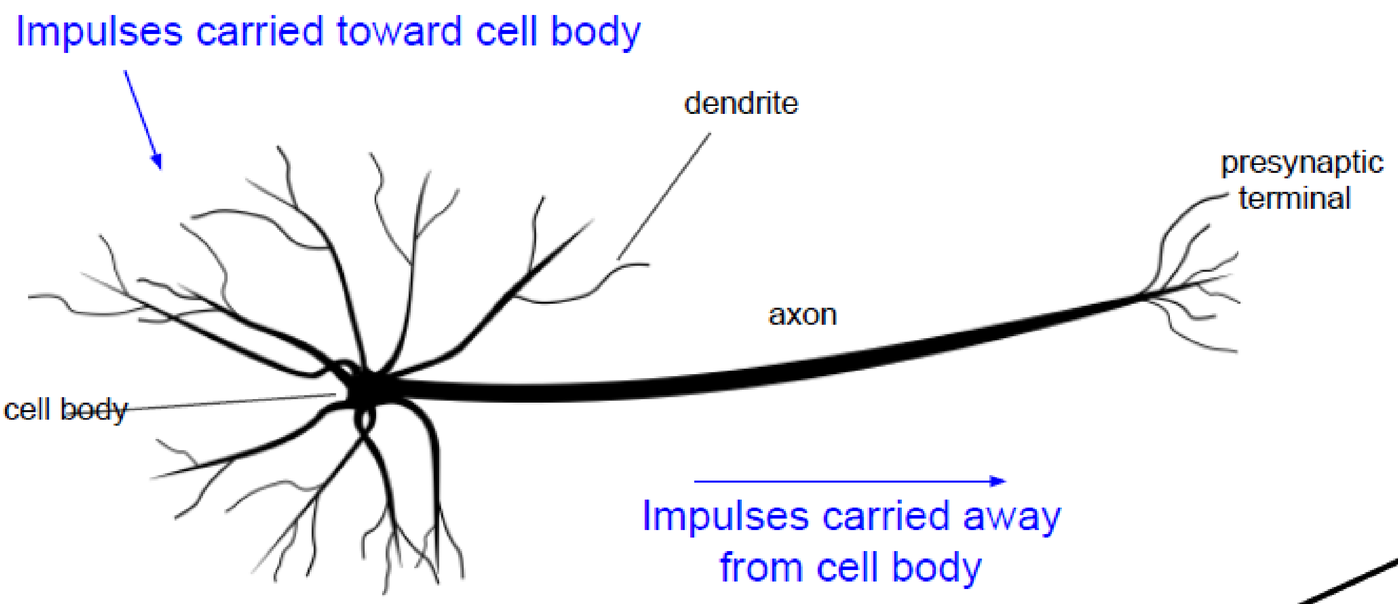
Impulses carried toward cell body
dendrite
presynaptic terminal
axon
cell body
Impulses carried away from cell body

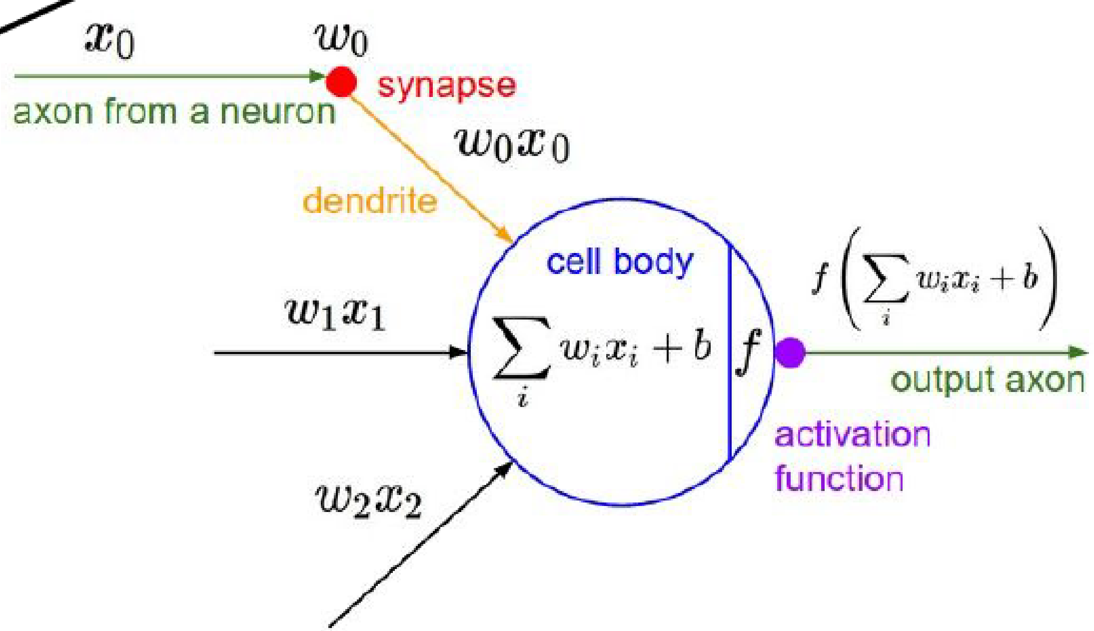
x_0
w_0
synapse
axon from a neuron
$w_0 x_0$
dendrite
cell body
$f\left(\sum_i w_i x_i + b\right)$
$w_1 x_1$
$\sum_i w_i x_i + b$
f
output axon
activation function
$w_2 x_2$

Neural Network

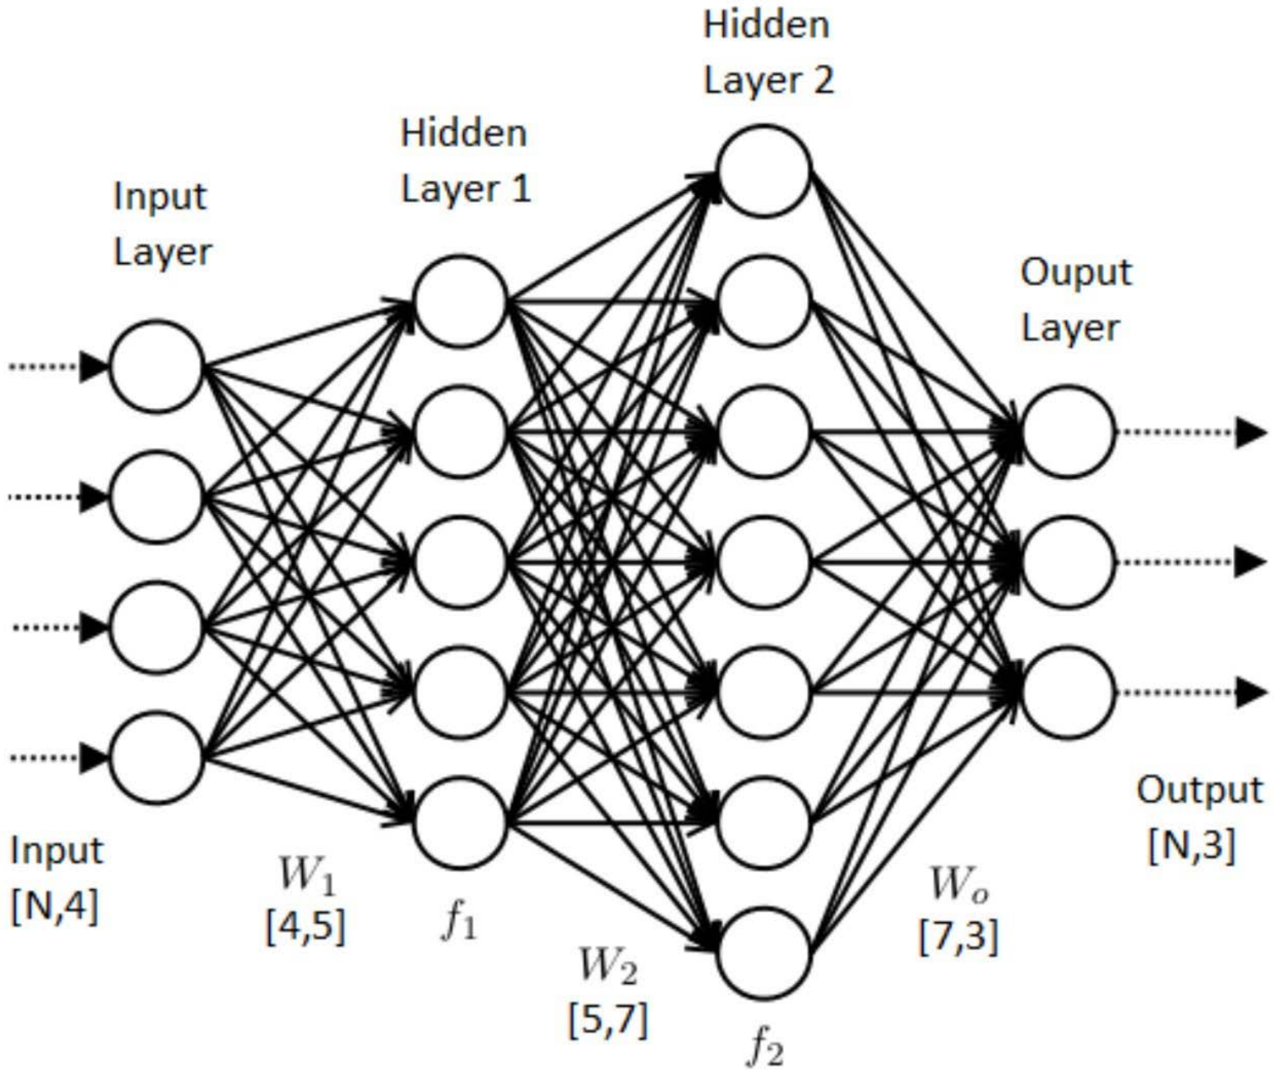

Backpropagation

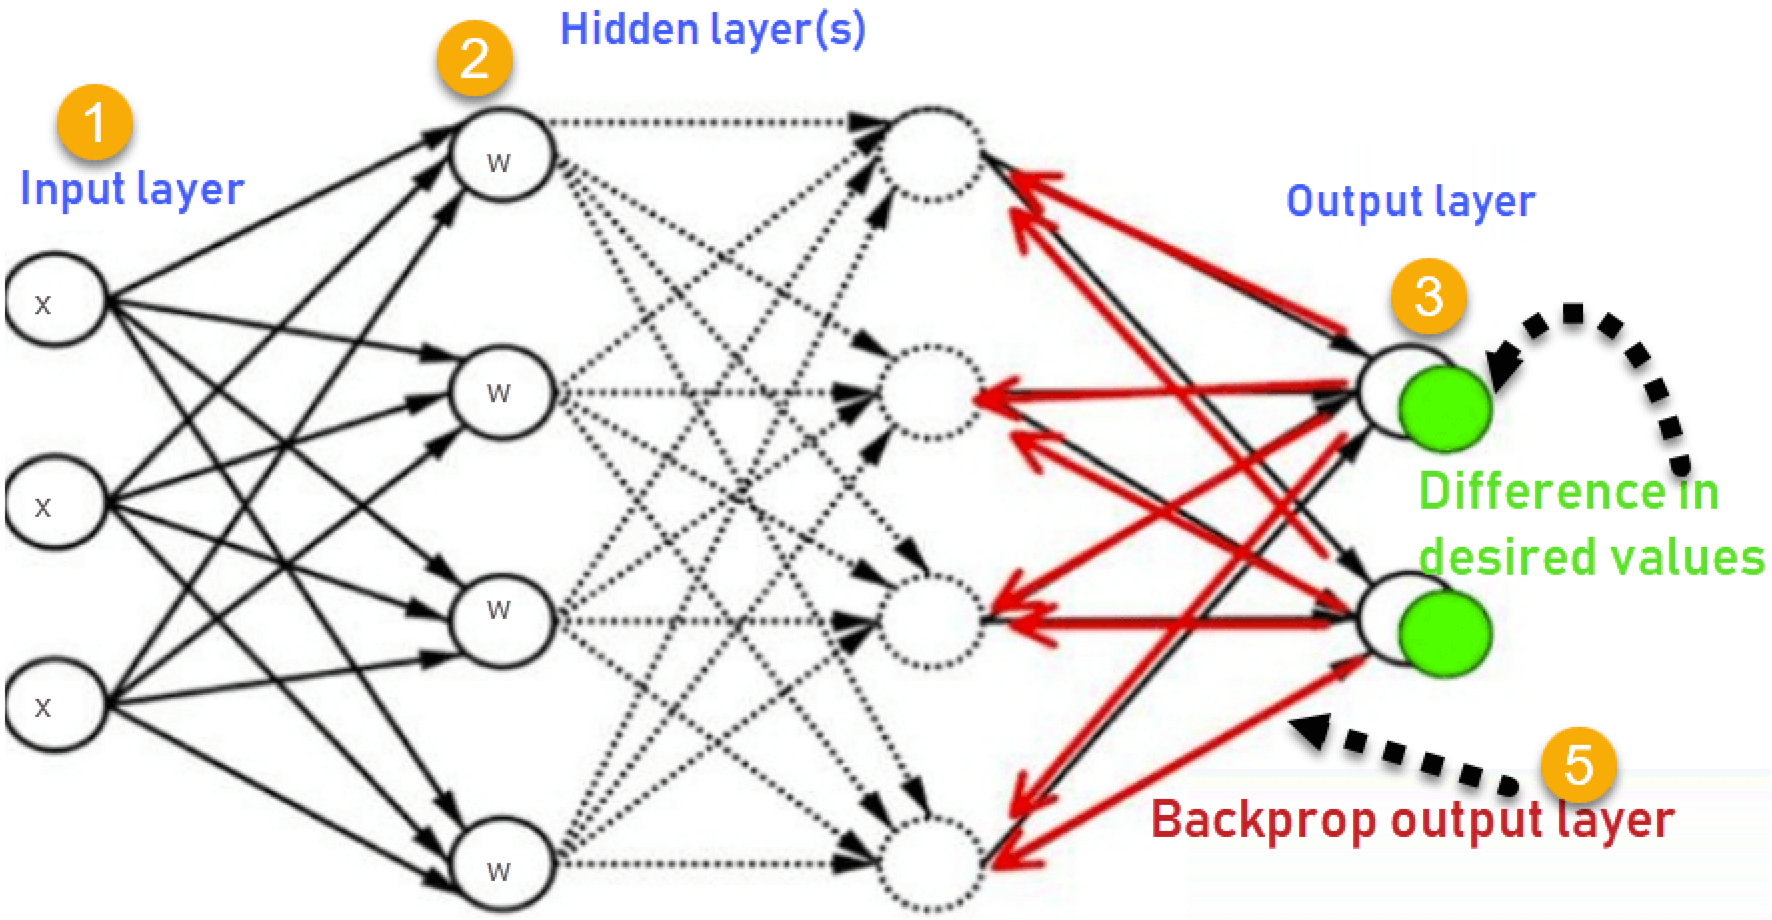

Reinforcement learning

HPE POINTNEXT

Reinforcement Learning

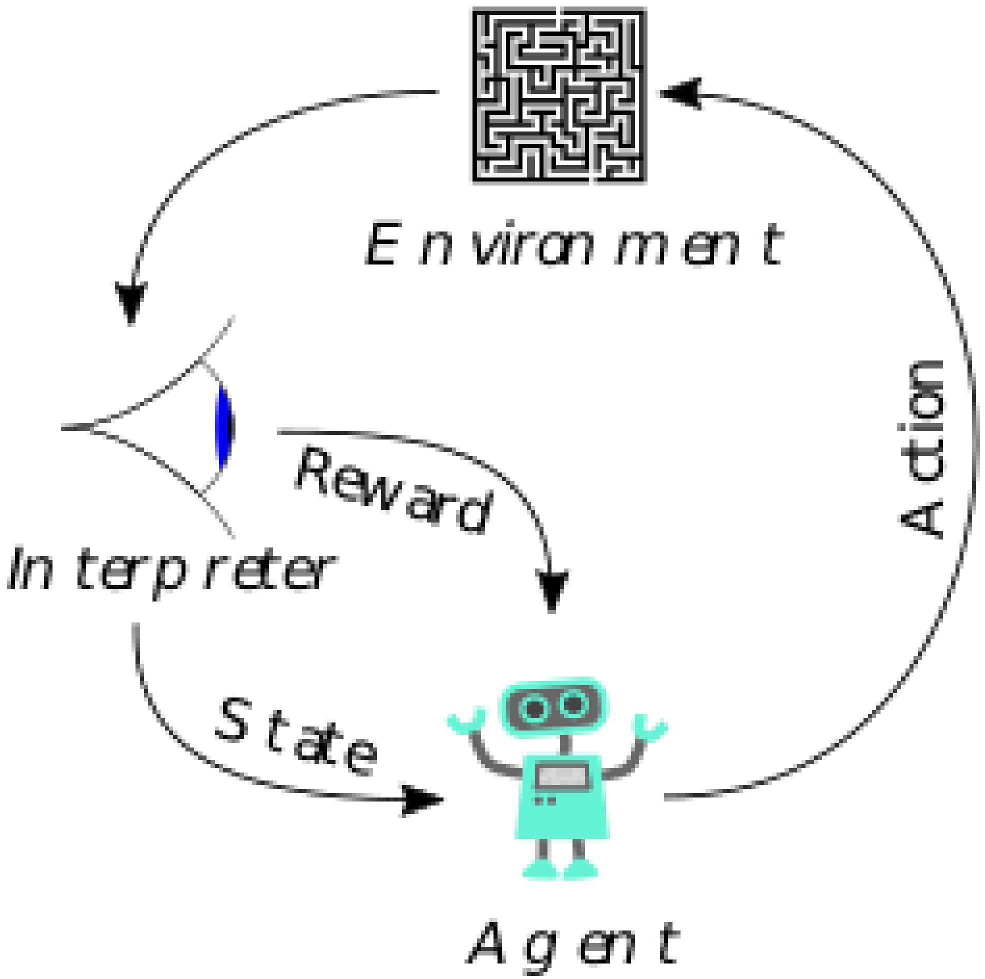

Reinforcement Learning

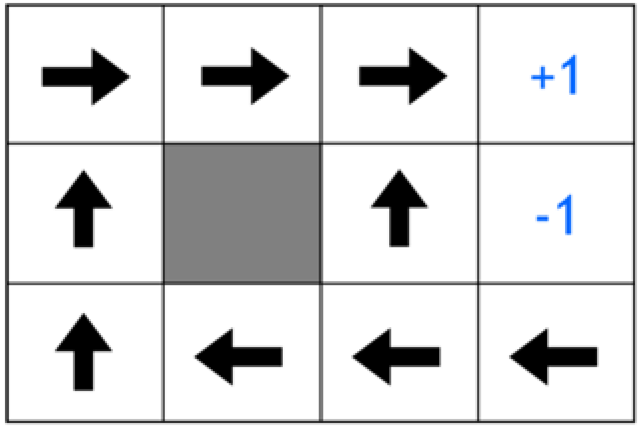

Q-learning: Learn function $Q : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$

Require:
  Sates $\mathcal{X} = \{1, \ldots, n_x\}$
  Actions $\mathcal{A} = \{1, \ldots, n_a\}$, $\quad A : \mathcal{X} \Rightarrow \mathcal{A}$
  Reward function $R : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$
  Black-box (probabilistic) transition function $T : \mathcal{X} \times \mathcal{A} \to \mathcal{X}$
  Learning rate $\alpha \in [0, 1]$, typically $\alpha = 0.1$
  Discounting factor $\gamma \in [0, 1]$
  procedure QLEARNING($\mathcal{X}$, A, R, T, α, γ)
    Initialize $Q : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$ arbitrarily
    while Q is not converged **do**
      Start in state $s \in \mathcal{X}$
      while s is not terminal **do**
        Calculate π according to Q and exploration strategy (e.g. $\pi(x) \leftarrow \arg\max_a Q(x, a)$)
        $a \leftarrow \pi(s)$
        $r \leftarrow R(s, a)$ ▷ Receive the reward
        $s' \leftarrow T(s, a)$ ▷ Receive the new state
        $Q(s', a) \leftarrow (1 - \alpha) \cdot Q(s, a) + \alpha \cdot (r + \gamma \cdot \max_{a'} Q(s', a'))$
        $s \leftarrow s'$
    return Q

Let play Atari

https://www.youtube.com/watch?v=V1eYniJ0Rnk

HPE POINTNEXT

Convolutional Neural Networks

HPE POINTNEXT

CNN

– Visual recognition problems
  – Face detection, object detection, action classification, … “fun”!
  – Image classification subproblems: view point variation, illumination, background clutter, deformation, occlusion, etc
– Hello World – Image recognition
– Filters examples (eyes)

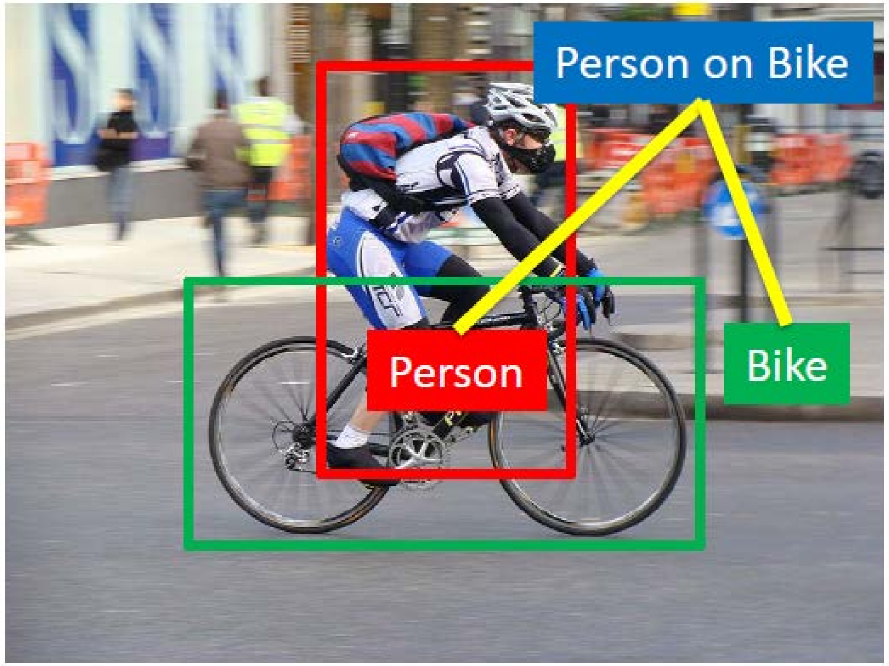

Our data

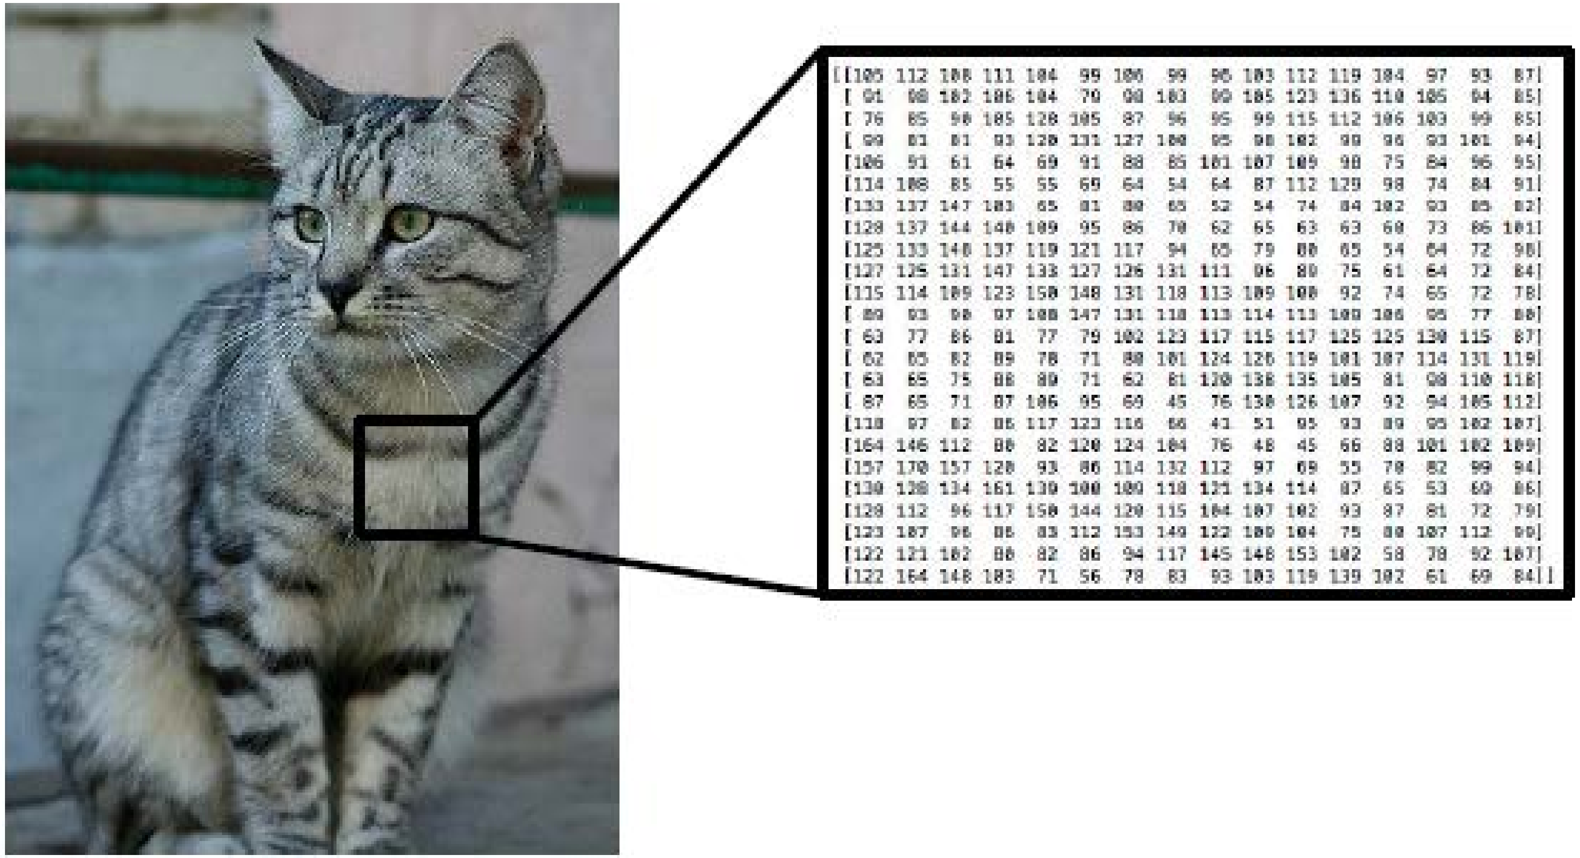
[[105 112 108 111 104 99 106 99 96 103 112 119 104 97 93 87]
[91 98 102 106 104 79 98 103 99 105 123 136 110 105 94 85]
[76 85 90 105 128 105 87 96 95 99 115 112 106 103 99 85]
[99 81 81 93 120 131 127 100 95 98 102 99 96 93 101 94]
[106 91 61 64 69 91 88 85 101 107 109 98 75 84 96 95]
[114 108 85 55 55 69 64 54 64 87 112 129 98 74 84 91]
[133 137 147 103 65 81 80 65 52 54 74 84 102 93 85 82]
[128 137 144 140 109 95 86 70 62 65 63 63 60 73 86 101]
[125 133 148 137 119 121 117 94 65 79 80 65 54 64 72 98]
[127 125 131 147 133 127 126 131 111 96 89 75 61 64 72 84]
[115 114 109 123 150 148 131 118 113 109 100 92 74 65 72 78]
[89 93 90 97 108 147 131 118 113 114 113 109 106 95 77 80]
[63 77 86 81 77 79 102 123 117 115 117 125 125 130 115 87]
[62 65 82 89 78 71 80 101 124 126 119 101 107 114 131 119]
[63 65 75 88 89 71 62 81 120 138 135 105 81 98 110 118]
[87 65 71 87 106 95 69 45 76 130 126 107 92 94 105 112]
[118 97 82 86 117 123 116 66 41 51 95 93 89 95 102 107]
[164 146 112 80 82 120 124 104 76 48 45 66 88 101 102 109]
[157 170 157 120 93 86 114 132 112 97 69 55 70 82 99 94]
[130 128 134 161 139 100 109 118 121 134 114 87 65 53 69 86]
[128 112 96 117 150 144 120 115 104 107 102 93 87 81 72 79]
[123 107 96 86 83 112 153 149 122 109 104 75 80 107 112 99]
[122 121 102 80 82 86 94 117 145 148 153 102 58 78 92 107]
[122 164 148 103 71 56 78 83 93 103 119 139 102 61 69 84]]

Face detection

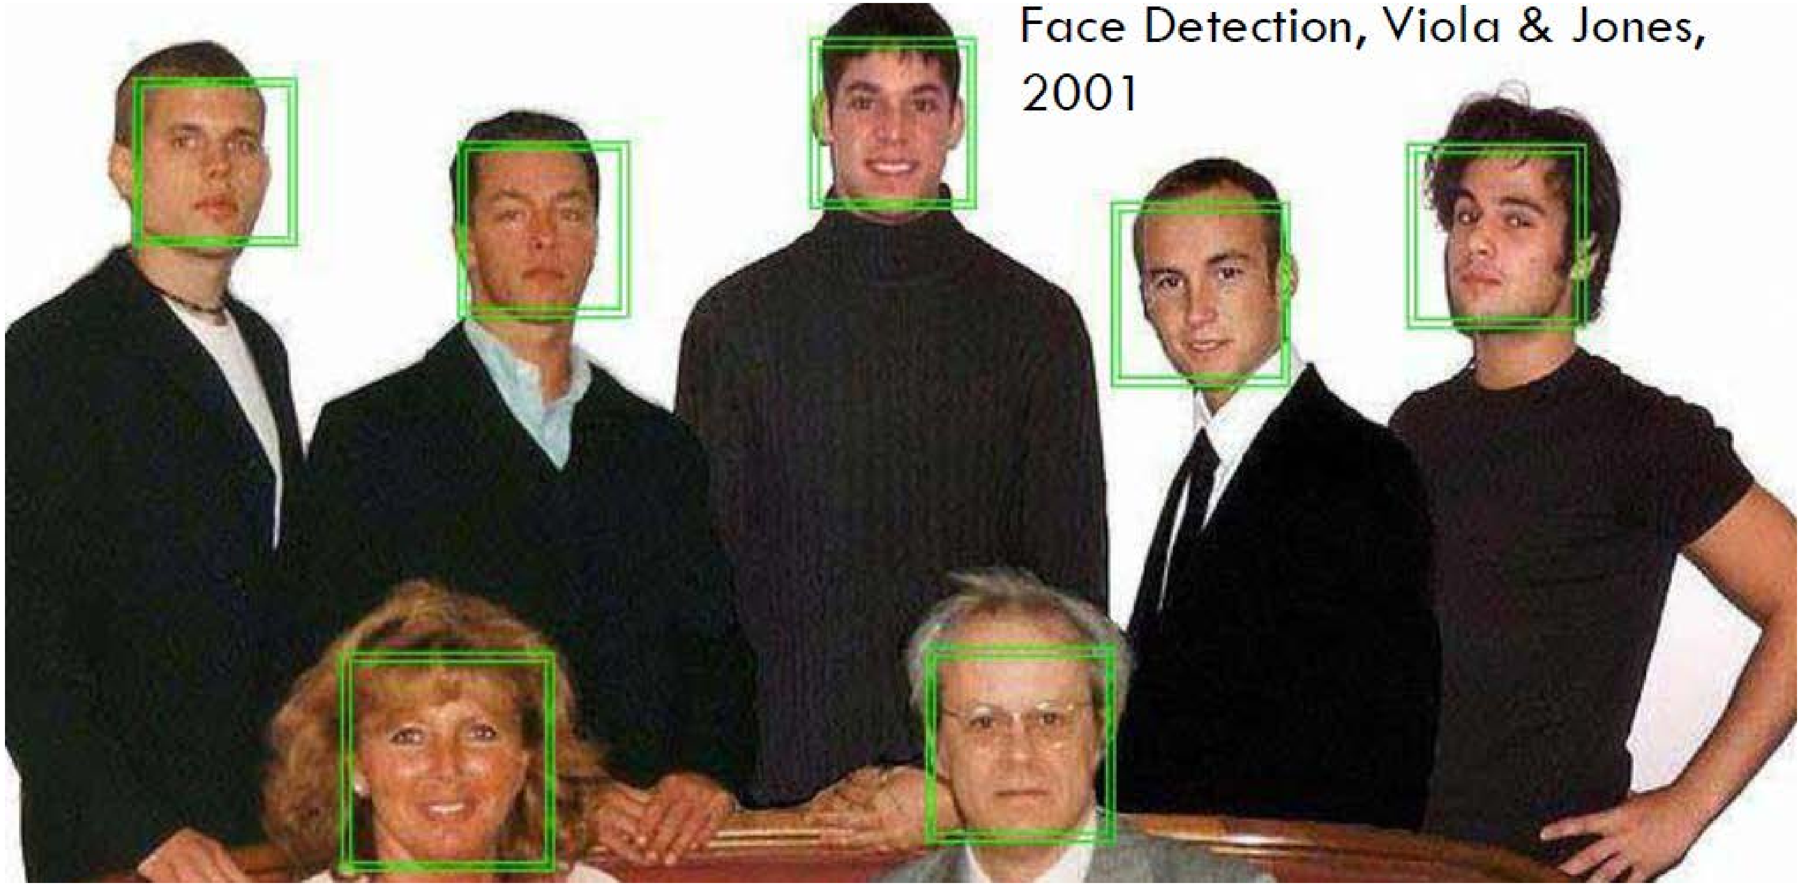

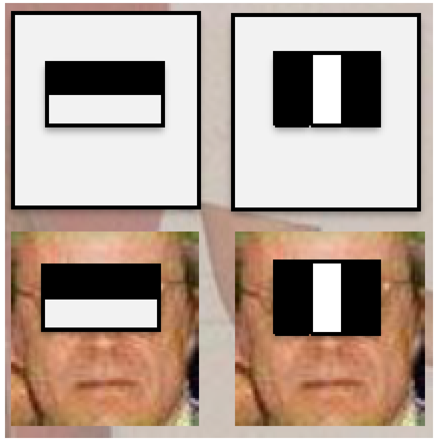

Example with an image with 4 pixels, and 3 classes (cat/dog/ship)

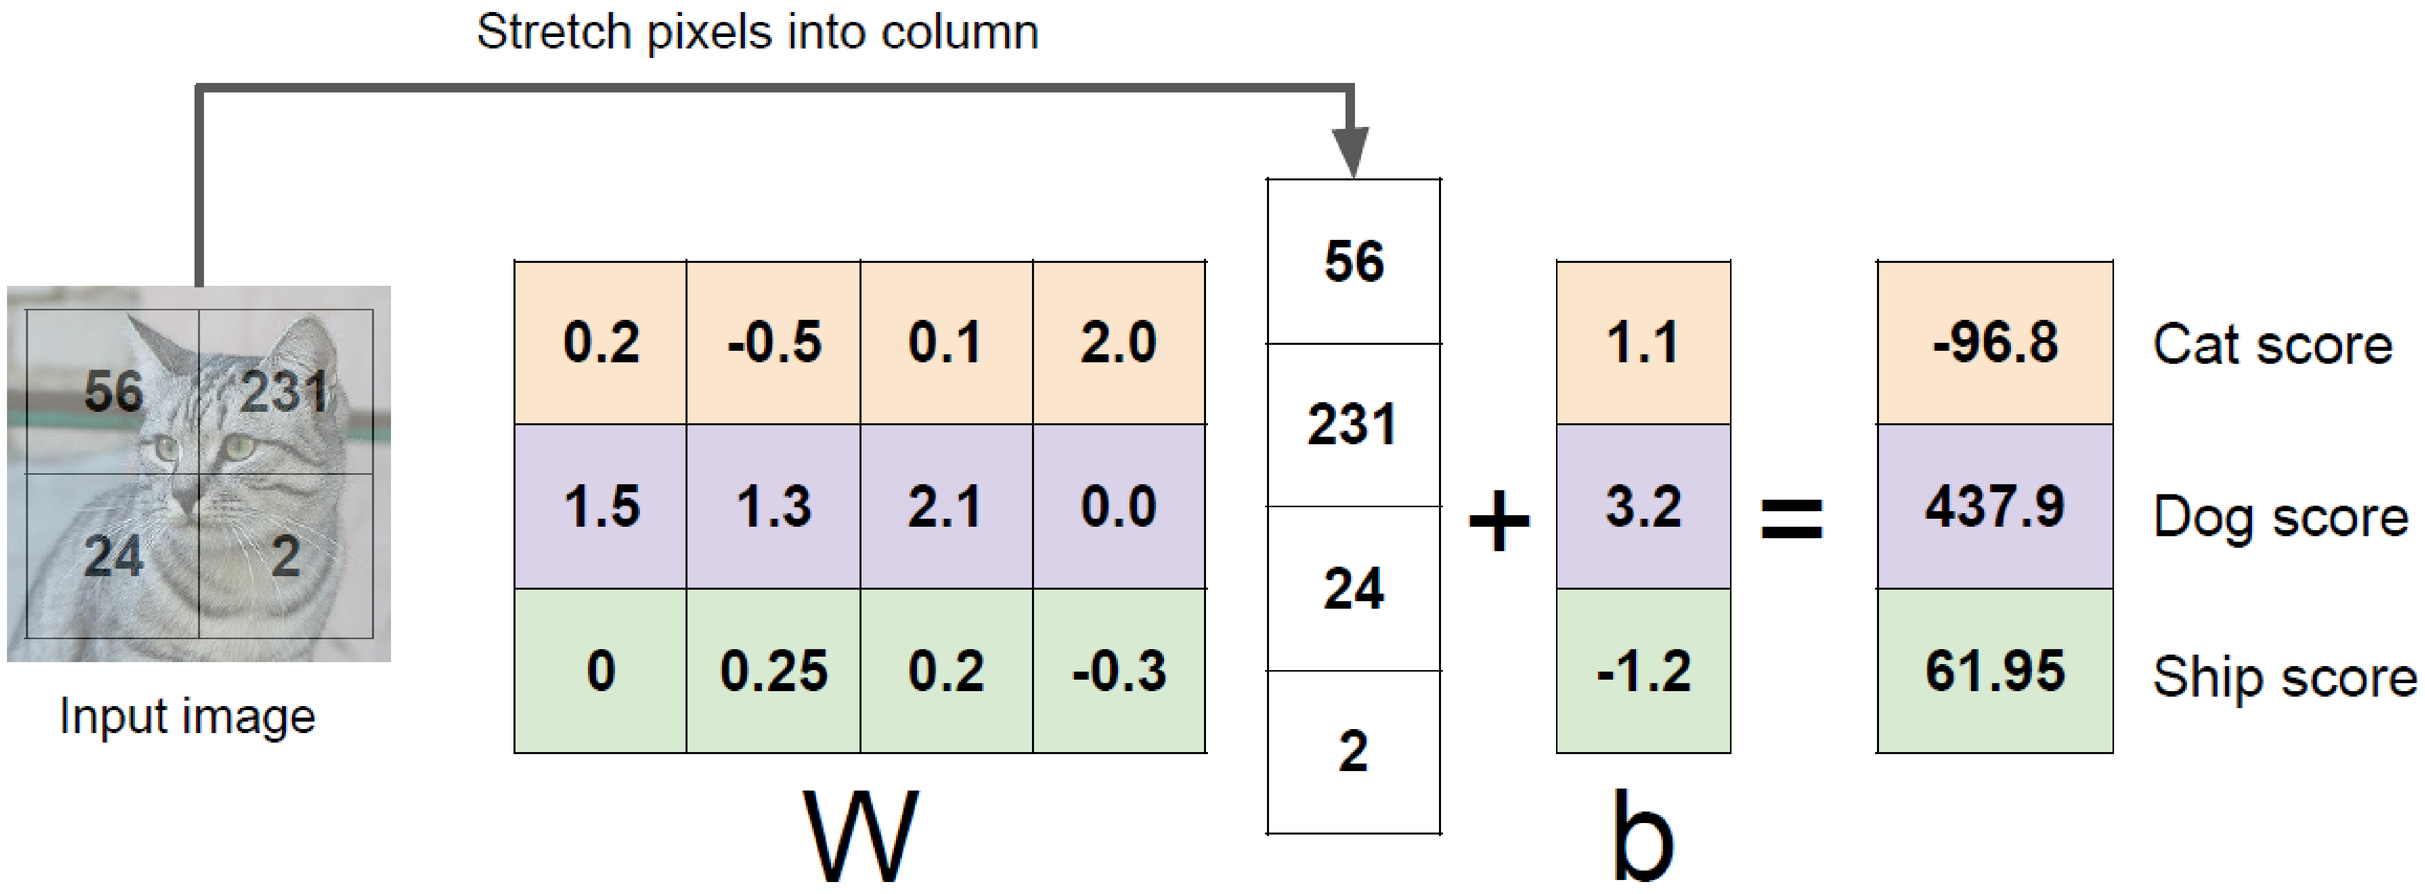

1998

LeCun et al.

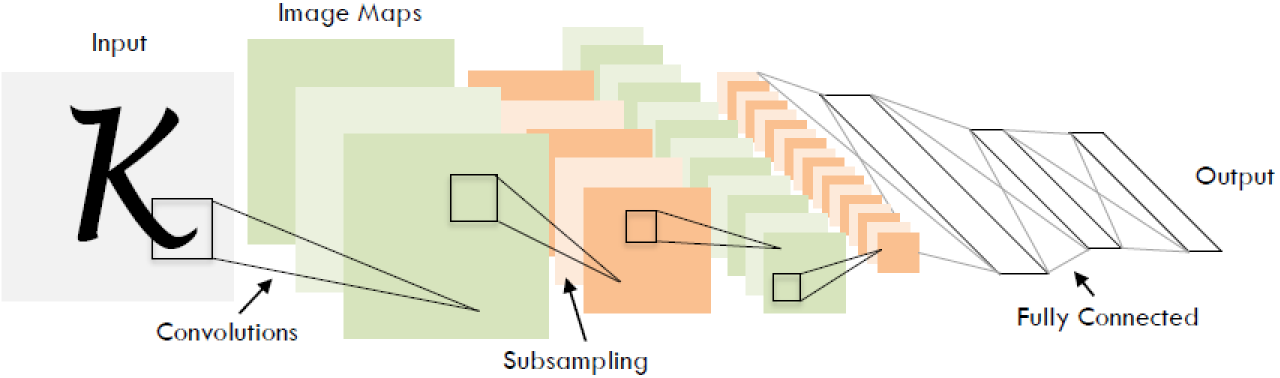

# of transistors

10^6

# of pixels used in training

10^7 NIST

2012

Krizhevsky et al.

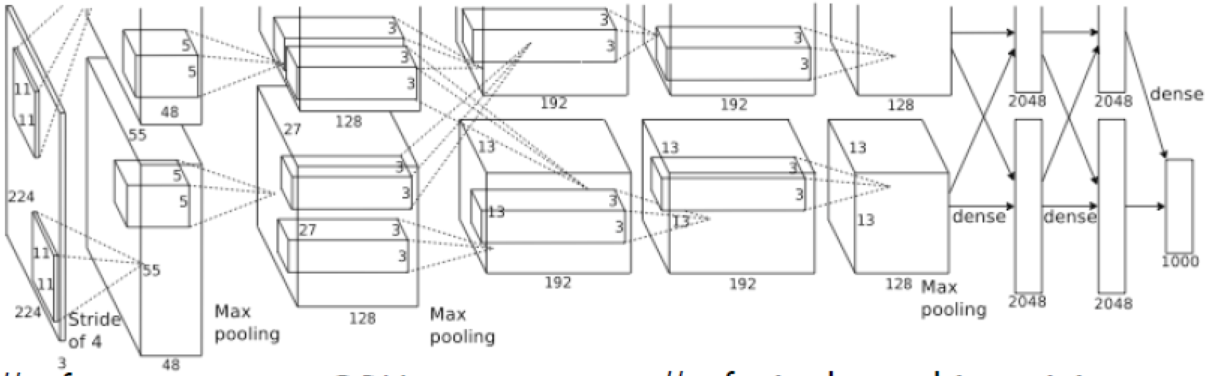

# of transistors

10^9

GPUs

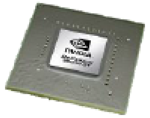

# of pixels used in training

10^{14} IMAGENET

Figure copyright Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012. Reproduced with permission.

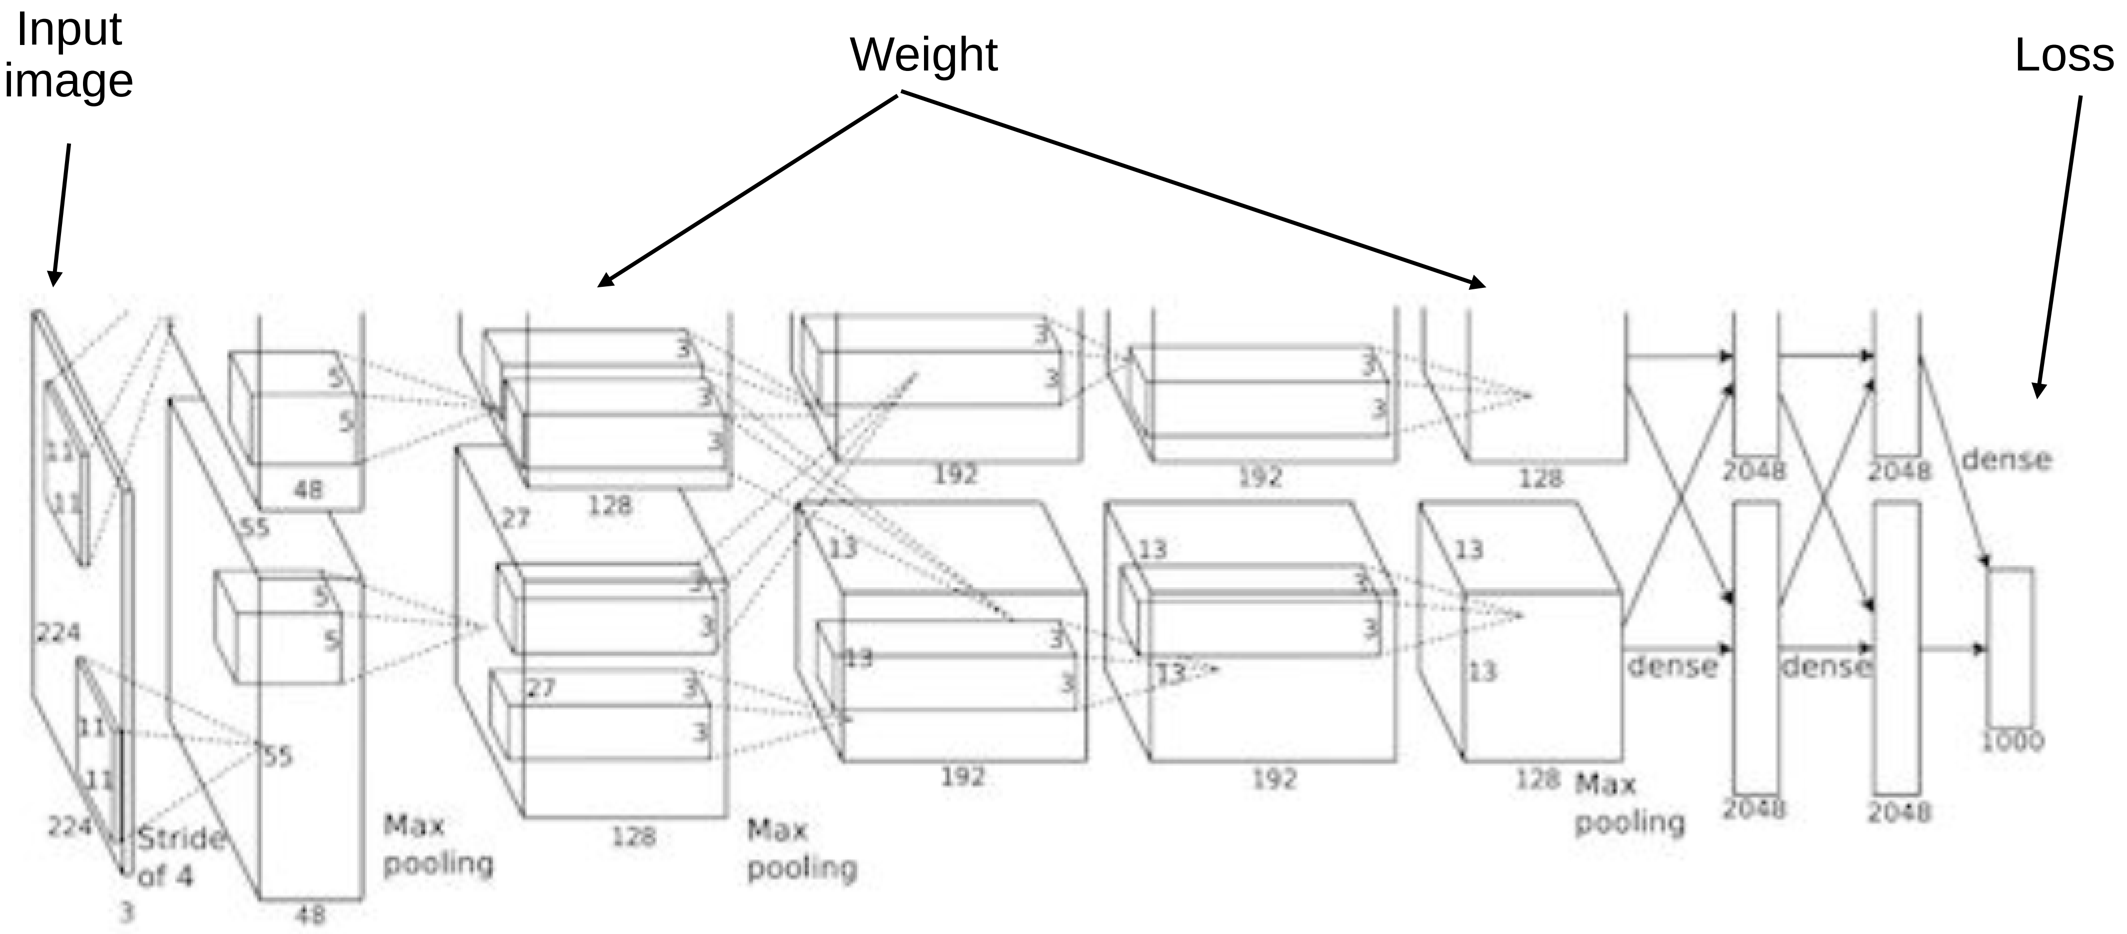
Input image
Weight
Loss
224
11
11
224
Stride of 4
3
55
5
5
48
48
Max pooling
27
128
3
3
128
Max pooling
13
192
192
192
192
128
128
Max pooling
dense
dense
dense
2048
2048
2048
2048
1000

AlexNet

Case Study: AlexNet

[Krizhevsky et al. 2012]

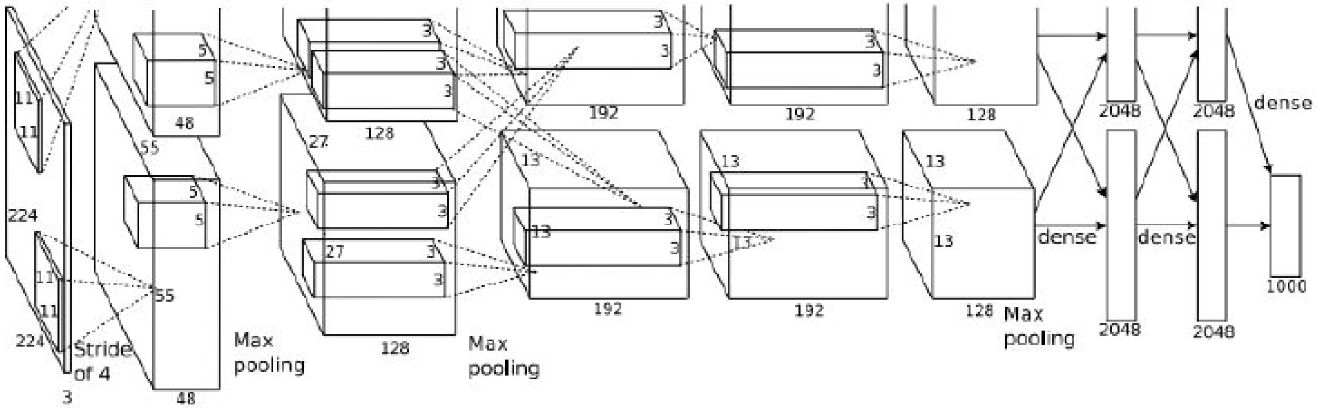

Full (simplified) AlexNet architecture:
[227x227x3] INPUT
[55x55x96] CONV1: 96 11x11 filters at stride 4, pad 0
[27x27x96] MAX POOL1: 3x3 filters at stride 2
[27x27x96] NORM1: Normalization layer
[27x27x256] CONV2: 256 5x5 filters at stride 1, pad 2
[13x13x256] MAX POOL2: 3x3 filters at stride 2
[13x13x256] NORM2: Normalization layer
[13x13x384] CONV3: 384 3x3 filters at stride 1, pad 1
[13x13x384] CONV4: 384 3x3 filters at stride 1, pad 1
[13x13x256] CONV5: 256 3x3 filters at stride 1, pad 1
[6x6x256] MAX POOL3: 3x3 filters at stride 2
[4096] FC6: 4096 neurons
[4096] FC7: 4096 neurons
[1000] FC8: 1000 neurons (class scores)

ML in networking (example)

IntroSpect User and Entity Behavior Analytics (UEBA)

HPE POINTNEXT

Behavior – Many Different Dimensions

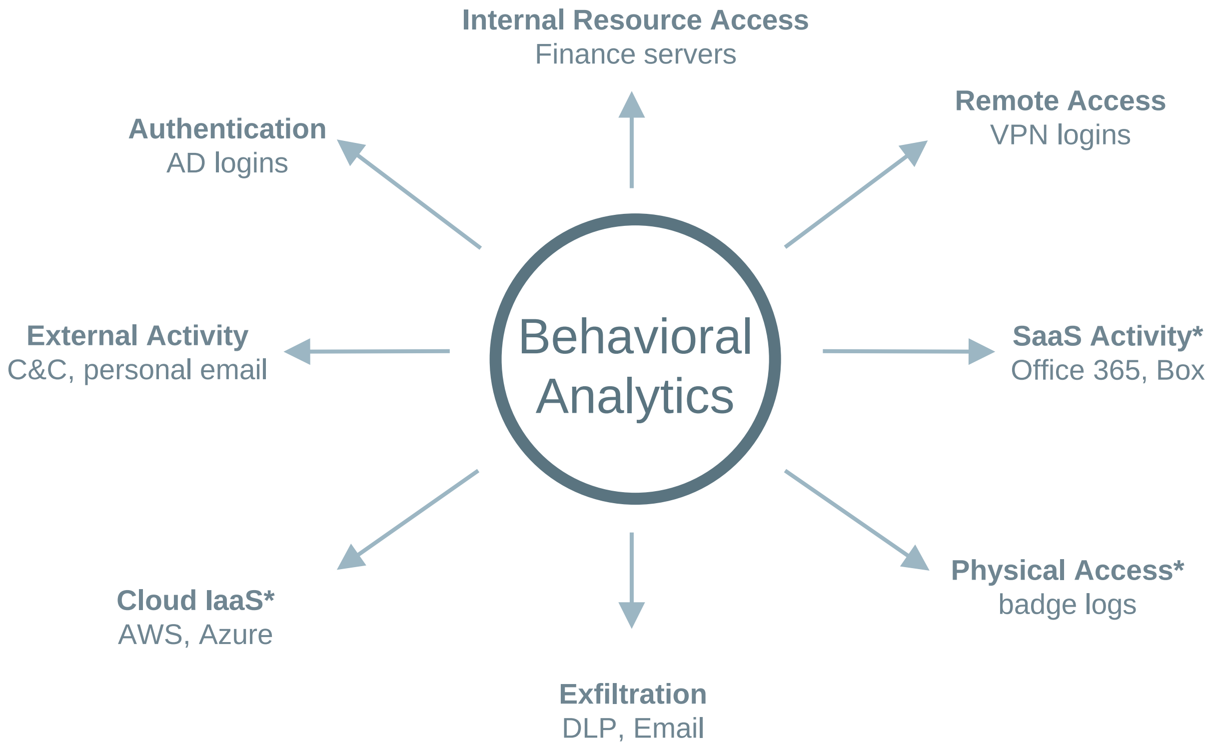

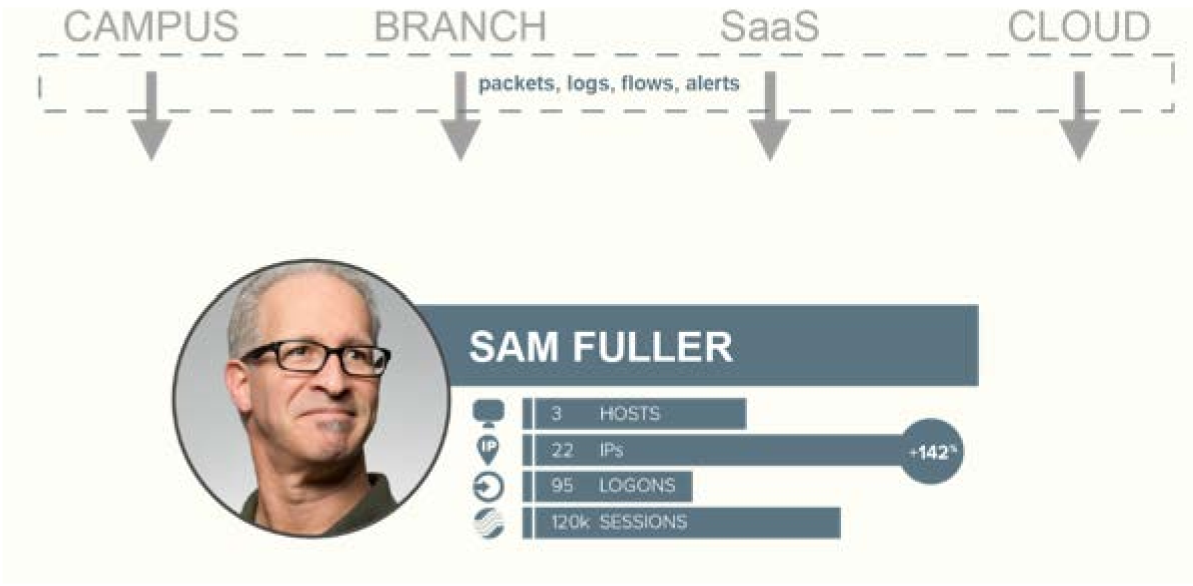

IntroSpect UEBA

https://www.arubanetworks.com/products/security/ueba/

Peer Baseline Anomaly

Context

Entities	Entity Risk Score		Alert Severity	Confidence Score	Attack Stage	Conversations
mjohnson	76 When Jan 19, 2016 10:00:44 AM	83 Detected Jan 20, 2016 12:53:10 PM	80	89	Internal Activity	

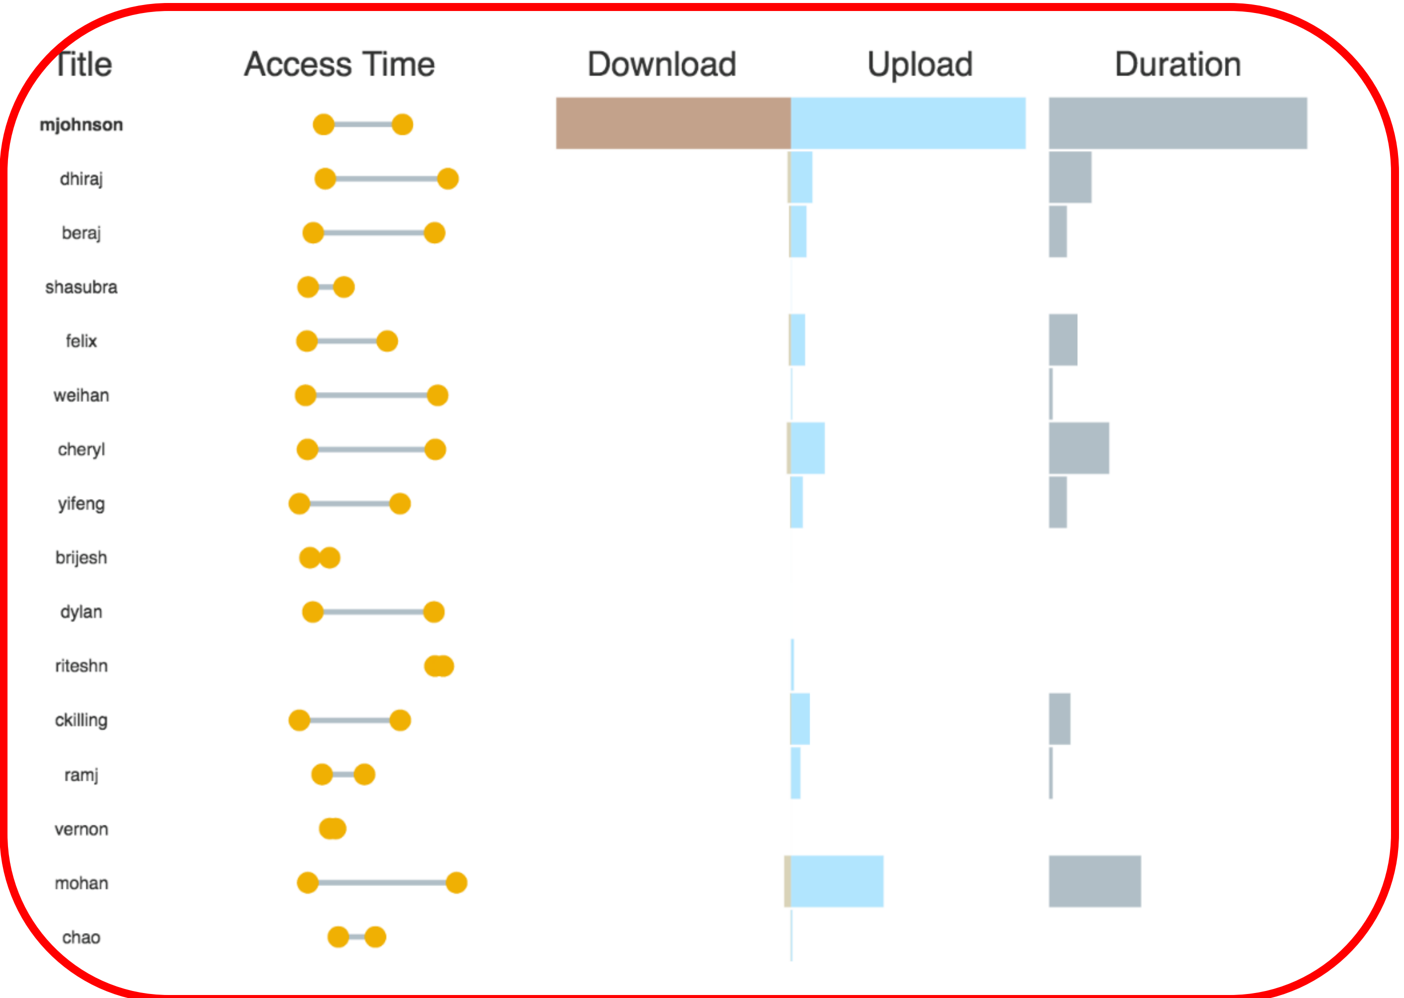

	Daily Time Spent on Site	Age	Area Income	Daily Internet Usage	Ad Topic Line	City	Male	Country	Timestamp	Clicked on Ad
0	68.95	35	61833.90	256.09	Cloned 5thgeneration orchestration	Wrightburgh	0	Tunisia	2016-03-27 00:53:11	0
1	80.23	31	68441.85	193.77	Monitored national standardization	West Jodi	1	Nauru	2016-04-04 01:39:02	0
2	69.47	26	59785.94	236.50	Organic bottom-line service-desk	Davidton	0	San Marino	2016-03-13 20:35:42	0
3	74.15	29	54806.18	245.89	Triple-buffered reciprocal time-frame	West Terrifurt	1	Italy	2016-01-10 02:31:19	0
4	68.37	35	73889.99	225.58	Robust logistical utilization	South Manuel	0	Iceland	2016-06-03 03:36:18	0
5	59.99	23	59761.56	226.74	Sharable client-driven software	Jamieberg	1	Norway	2016-05-19 14:30:17	0
6	88.91	33	53852.85	208.36	Enhanced dedicated support	Brandonstad	0	Myanmar	2016-01-28 20:59:32	0
7	66.00	48	24593.33	131.76	Reactive local challenge	Port Jefferybury	1	Australia	2016-03-07 01:40:15	1
8	74.53	30	68862.00	221.51	Configurable coherent function	West Colin	1	Grenada	2016-04-18 09:33:42	0
9	69.88	20	55642.32	183.82	Mandatory homogeneous architecture	Ramirezton	1	Ghana	2016-07-11 01:42:51	0

Finding the Malicious in the Anomalous

BUSINESS CONTEXT
High Value Assets
High Value Actors

MACHINE LEARNING

SUPERVISED
UNSUPERVISED

THIRD PARTY ALERTS

DLP
Sandbox
Firewalls
STIX
Rules
Etc.

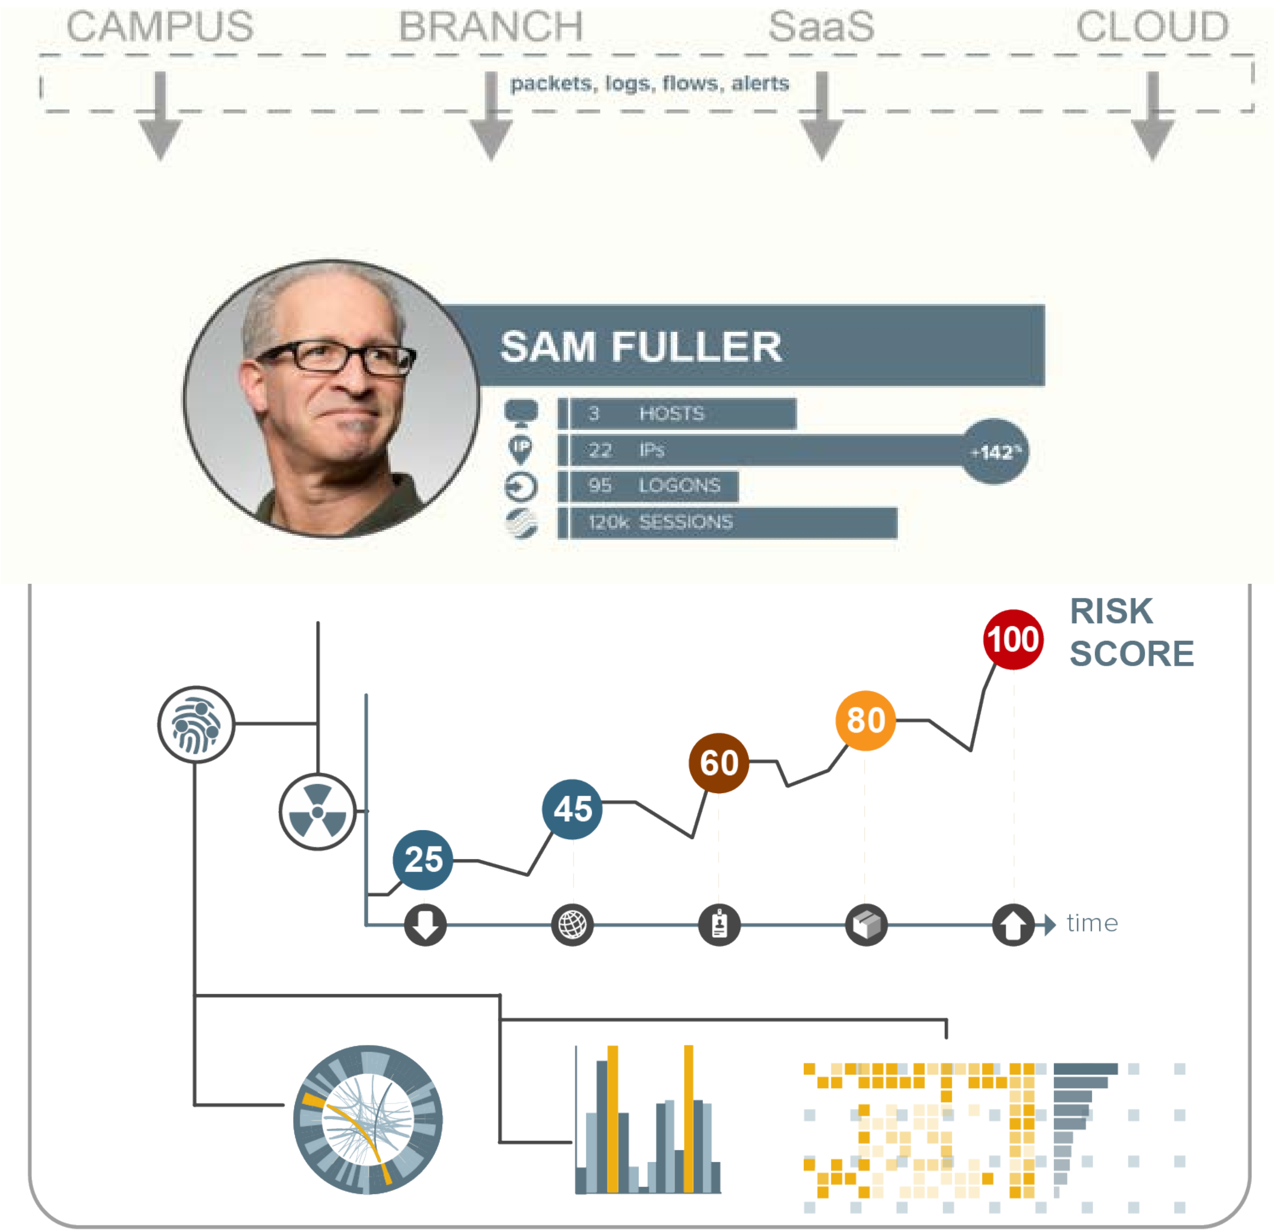

AI in production

HPE Prescriptive Maintenance

HPE POINTNEXT

HPE Digital Prescriptive Maintenance Use Case

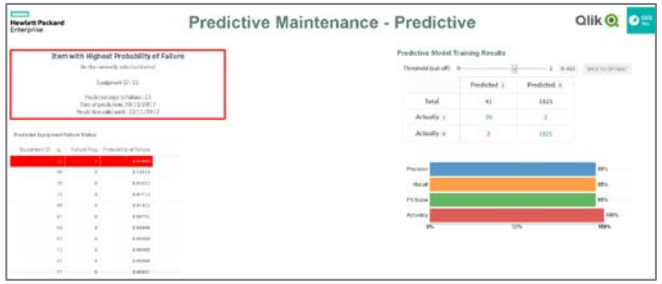

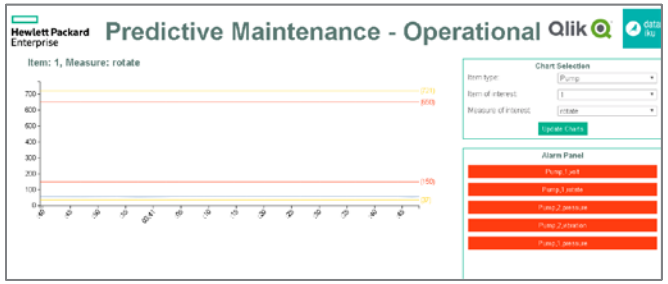

HPE POINTNEXT

Workshop	Proof-of Value	Implementation

Analytics and ML Software

Data lakes, IoT, Infrastructure

HPE EdgeLine

HPE Apollo

HPE Servers

- Prescribes best maintenance decisions and automates actions
- Open end-to-end solution framework including data ingestion, visualization analysis, prediction and automation
- Integrates with various structured and unstructured sources at the Edge and Hybrid Cloud
- Works with leading maintenance technology partners (Dataiku, PTC/thingsworx or GE/Predix)
- Delivered through a workshop, POV, implementation services

Last remarks

HPE POINTNEXT

Last remarks

- Fascinating world!
- Don’t start if you are afraid maths
- Learn python (?)
- It is more important to understand and analyse the data rather than to write code
- AI ethics are important
- AI is already here!

Further reading

- Deep learning & machine learning solutions
  - https://www.hpe.com/us/en/solutions/hpc-high-performance-computing/deep-learning.html
- Enterprise artificial intelligence
  - https://www.hpe.com/us/en/solutions/artificial-intelligence.html
- How to prepare your business for AI
  - https://www.hpe.com/us/en/insights/articles/how-to-prepare-your-business-for-ai-1711.html

HPE is making AI enterprise-grade

Using AI to perfect race strategy

Converging AI, IT and OT
to boost output and quality

Accelerating Alzheimer’s research
100x using in-memory analytics

HPE
POINTNEXT

Thank you

Athanassios.Liakopoulos@hpe.com

Hewlett Packard
Enterprise