IOTWeek | Aar 17-2

Security in AI : An ISV perspective

June, 6th, 2019

Dr Jean-Christophe Pazzaglia
(Demo from Aymen Mouelhi)

SAP Labs France
BDVe Project

BIG DATA VALUE

THE BEST RUN SAP

PUBLIC


IOT

????


SAP

SAP Leonardo IoT – From Internet of Things Data to Business Insights

Open Ecosystem & Partnerships

Cloud-to-Cloud interoperability

Vendor-agnostic

Technical Services

Connectivity

Device Management

SAP Leonardo IoT

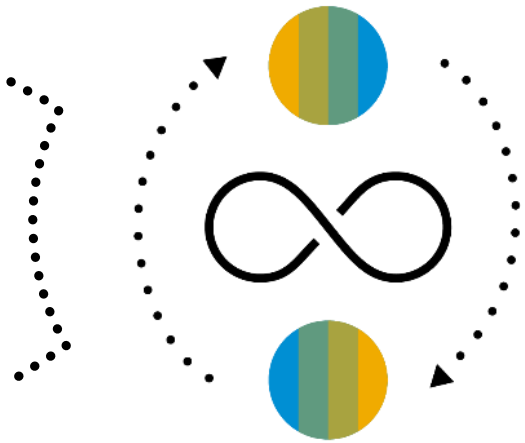

SAP Leonardo IoT Edge

Drive business outcomes with embedded IoT in lines of business applications and further intelligent technologies

Transform your business with open innovation while keeping a clean core

Integrate and contextualize IoT data with business process information

Extend your business process from the cloud to the edge

Leverage IoT insights with an integrated application developer experience

Unlock your Intelligent Enterprise with SAP Leonardo IoT

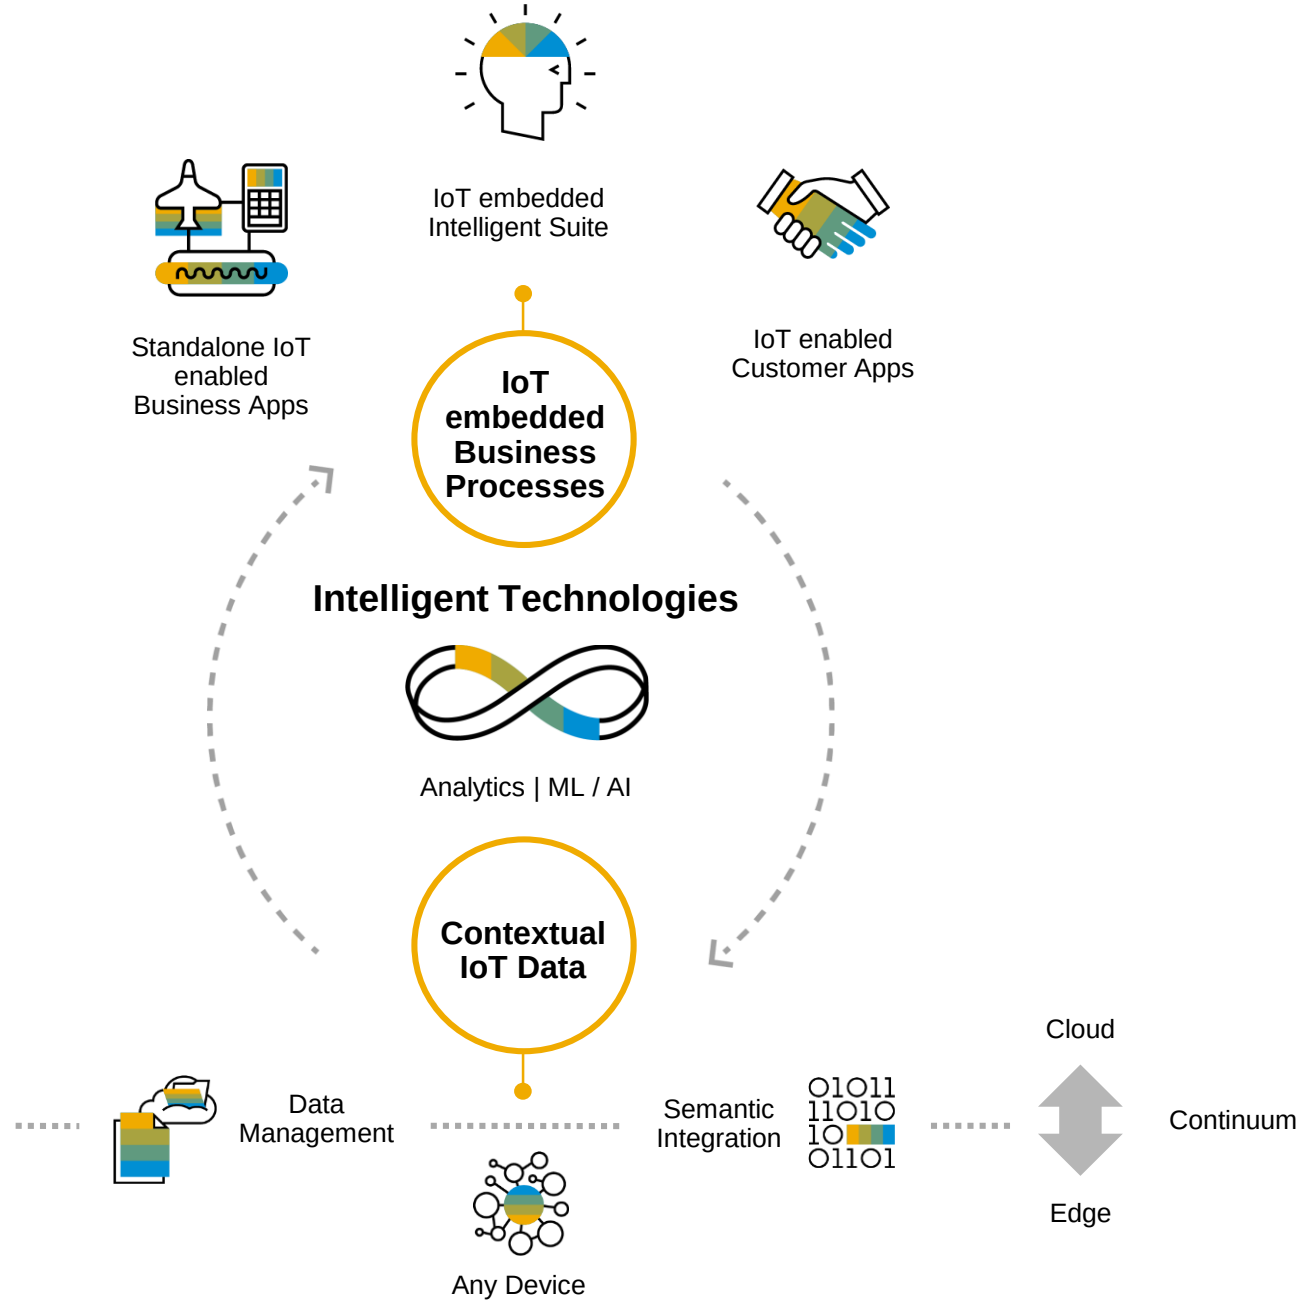

Cloud-Edge Continuum – Decide where to run your workload

SAP Leonardo IoT

Cloud

Continuum

Edge

SAP Leonardo IoT Edge

Consistent Execution
Integrate into cloud or on-prembusiness processes

Consistent Business Rules
Apply same set of business rules in the cloud or at the edge

Consistent Business Context
Bringing contextual data to the edge

Examples of IoT enabled Use Cases

Replenishment Optimization

Condition Monitoring

Fleet Management

Service Management

Usage Insights

Consumer's Experience

Trenitalia: Creating a Dynamic Maintenance Management System Powered by SAP HANA

https://www.sap.com/assetdetail/2015/12/b6caea0d-507c-0010-82c7-eda71af511fa.html

Secure Monitoring of Antibes Water Distribution Network

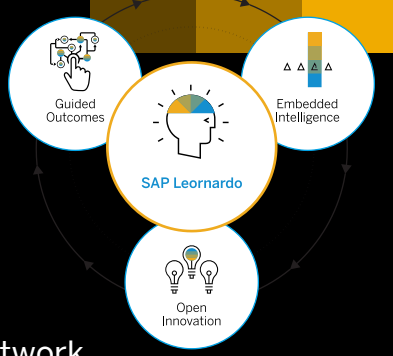

Motivation

- Monitor water distribution network
- Predict failures, and
- Optimize proactively their finance and controlling

Technical characteristics

- 2000 water sensors instrumenting the network (e.g., pressure, temperature, flow)
- Low bandwidth and powered Programmable Logical Controls
- No physical access to devices
- Low throughput
- Confidentiality and Integrity
- No personal data

Why AI now ?

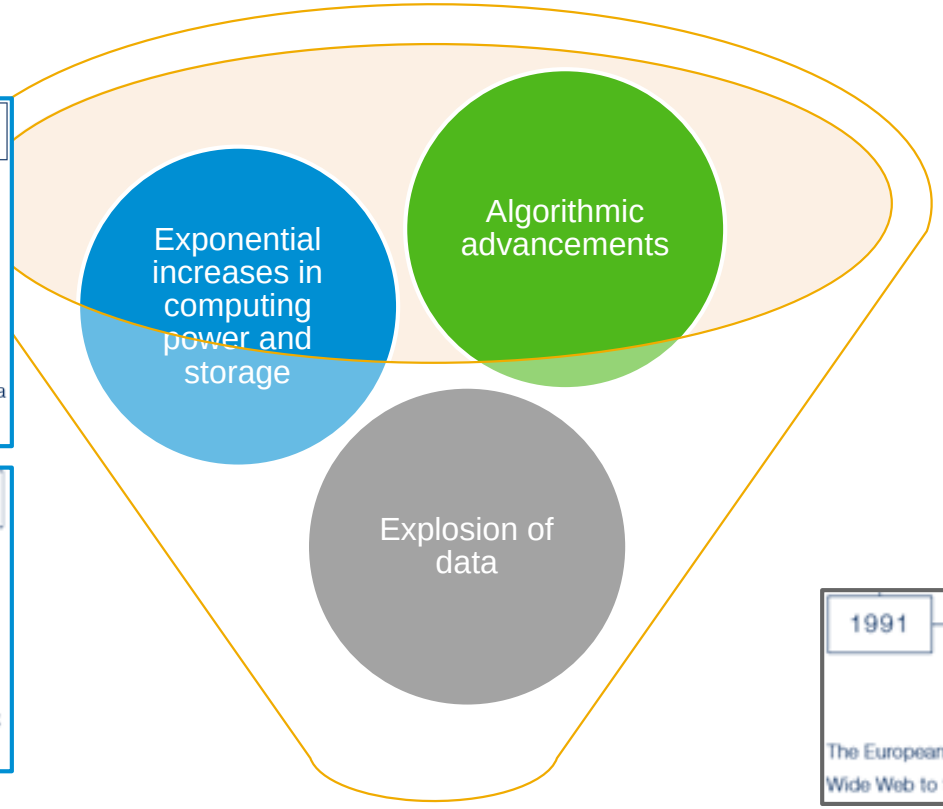

2004 — Dean and Ghemawat introduce the MapReduce algorithm to cope with data explosion

With the World Wide Web taking off, Google seeks out novel ideas to deal with the resulting proliferation of data. Computer scientist Jeff Dean (current head of Google Brain) and Google software engineer Sanjay Ghemawat develop MapReduce to deal with immense amounts of data by parallelizing processes across large data sets using a substantial number of computers.

2009 — Ng uses GPUs to train deep-learning models more efficiently

American computer scientist Andrew Ng and his team at Stanford University show that training deep-belief networks with 100 million parameters on GPUs is more than 70 times faster than doing so on CPUs, a finding that would reduce training that once took weeks to only one day.

1805 — Legendre lays the groundwork for machine learning

French mathematician Adrien-Marie Legendre publishes the least square method for regression, which he used to determine, from astronomical observations, the orbits of bodies around the sun. Although this method was developed as a statistical framework, it would provide the basis for many of today's machine-learning models.

1965 — Birth of deep learning

Ukrainian mathematician Alexey Grigorevich Ivakhnenko develops the first general working learning algorithms for supervised multilayer artificial neural networks (ANNs), in which several ANNs are stacked on top of one another and the output of one ANN layer feeds into the next. The architecture is very similar to today's deep-learning architectures.

1991 — Opening of the World Wide Web

The European Organization for Nuclear Research (CERN) begins opening up the World Wide Web to the public.

2004 — Web 2.0 hits its stride, launching the era of user-generated data

Web 2.0 refers to the shifting of the Internet paradigm from passive content viewing to interactive and collaborative content creation, social media, blogs, video, and other channels. Publishers Tim O'Reilly and Dale Dougherty popularize the term, though it was coined by designer Darcy DiNucci in 1999.

Success of AI ?

Executive Guide to AI
https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/an-executives-guide-to-ai

Success of ~~AI~~ ML

2005
- Darpa Grand challenge: five vechicles successfully completed (200 km in the desert) . Now, Waymo fleet drive >25000 miles per week

2011
- IBM's Watson computer defeated television game show Jeopardy! champions Rutter and Jennings (3x).

2012
- Deep convolutional neural net achieved 16% error rate on ImageNet (-10%);

2015
- Super-human performance for image classification, <5% error rate (ImageNet challenge)

2016
- Google Neural Machine Translation (GNMT) → almost human level, more 100 languages, more fluid (better segmentation), zero-shot learning (new couples)
- Google DeepMind's AlphaGo (version: Lee) defeated Lee Sedol 4–1

2017
- Lip Reading Sentences in the Wild (twice more efficient than human!)

… and many others DeepText, Pinterest “similar looks”, A9 search (Amazon), Smart Personal Assistant (Alexa)

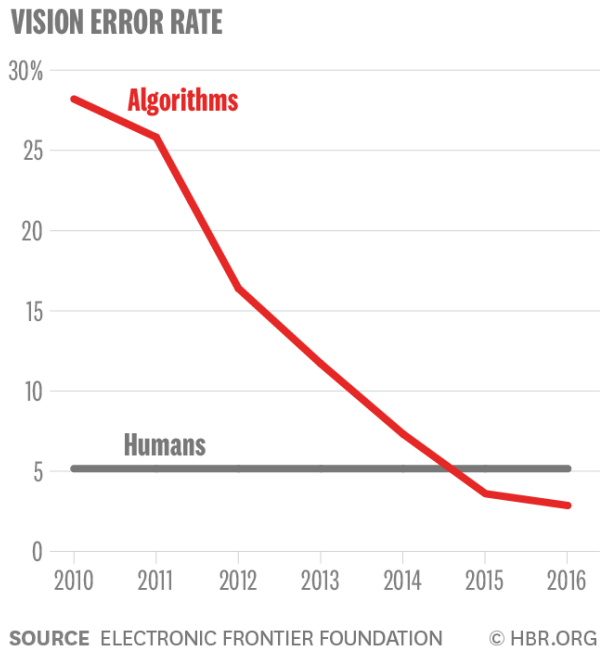


A pizza sitting on top of a table
Description generated with high confidence
SAP TECHED

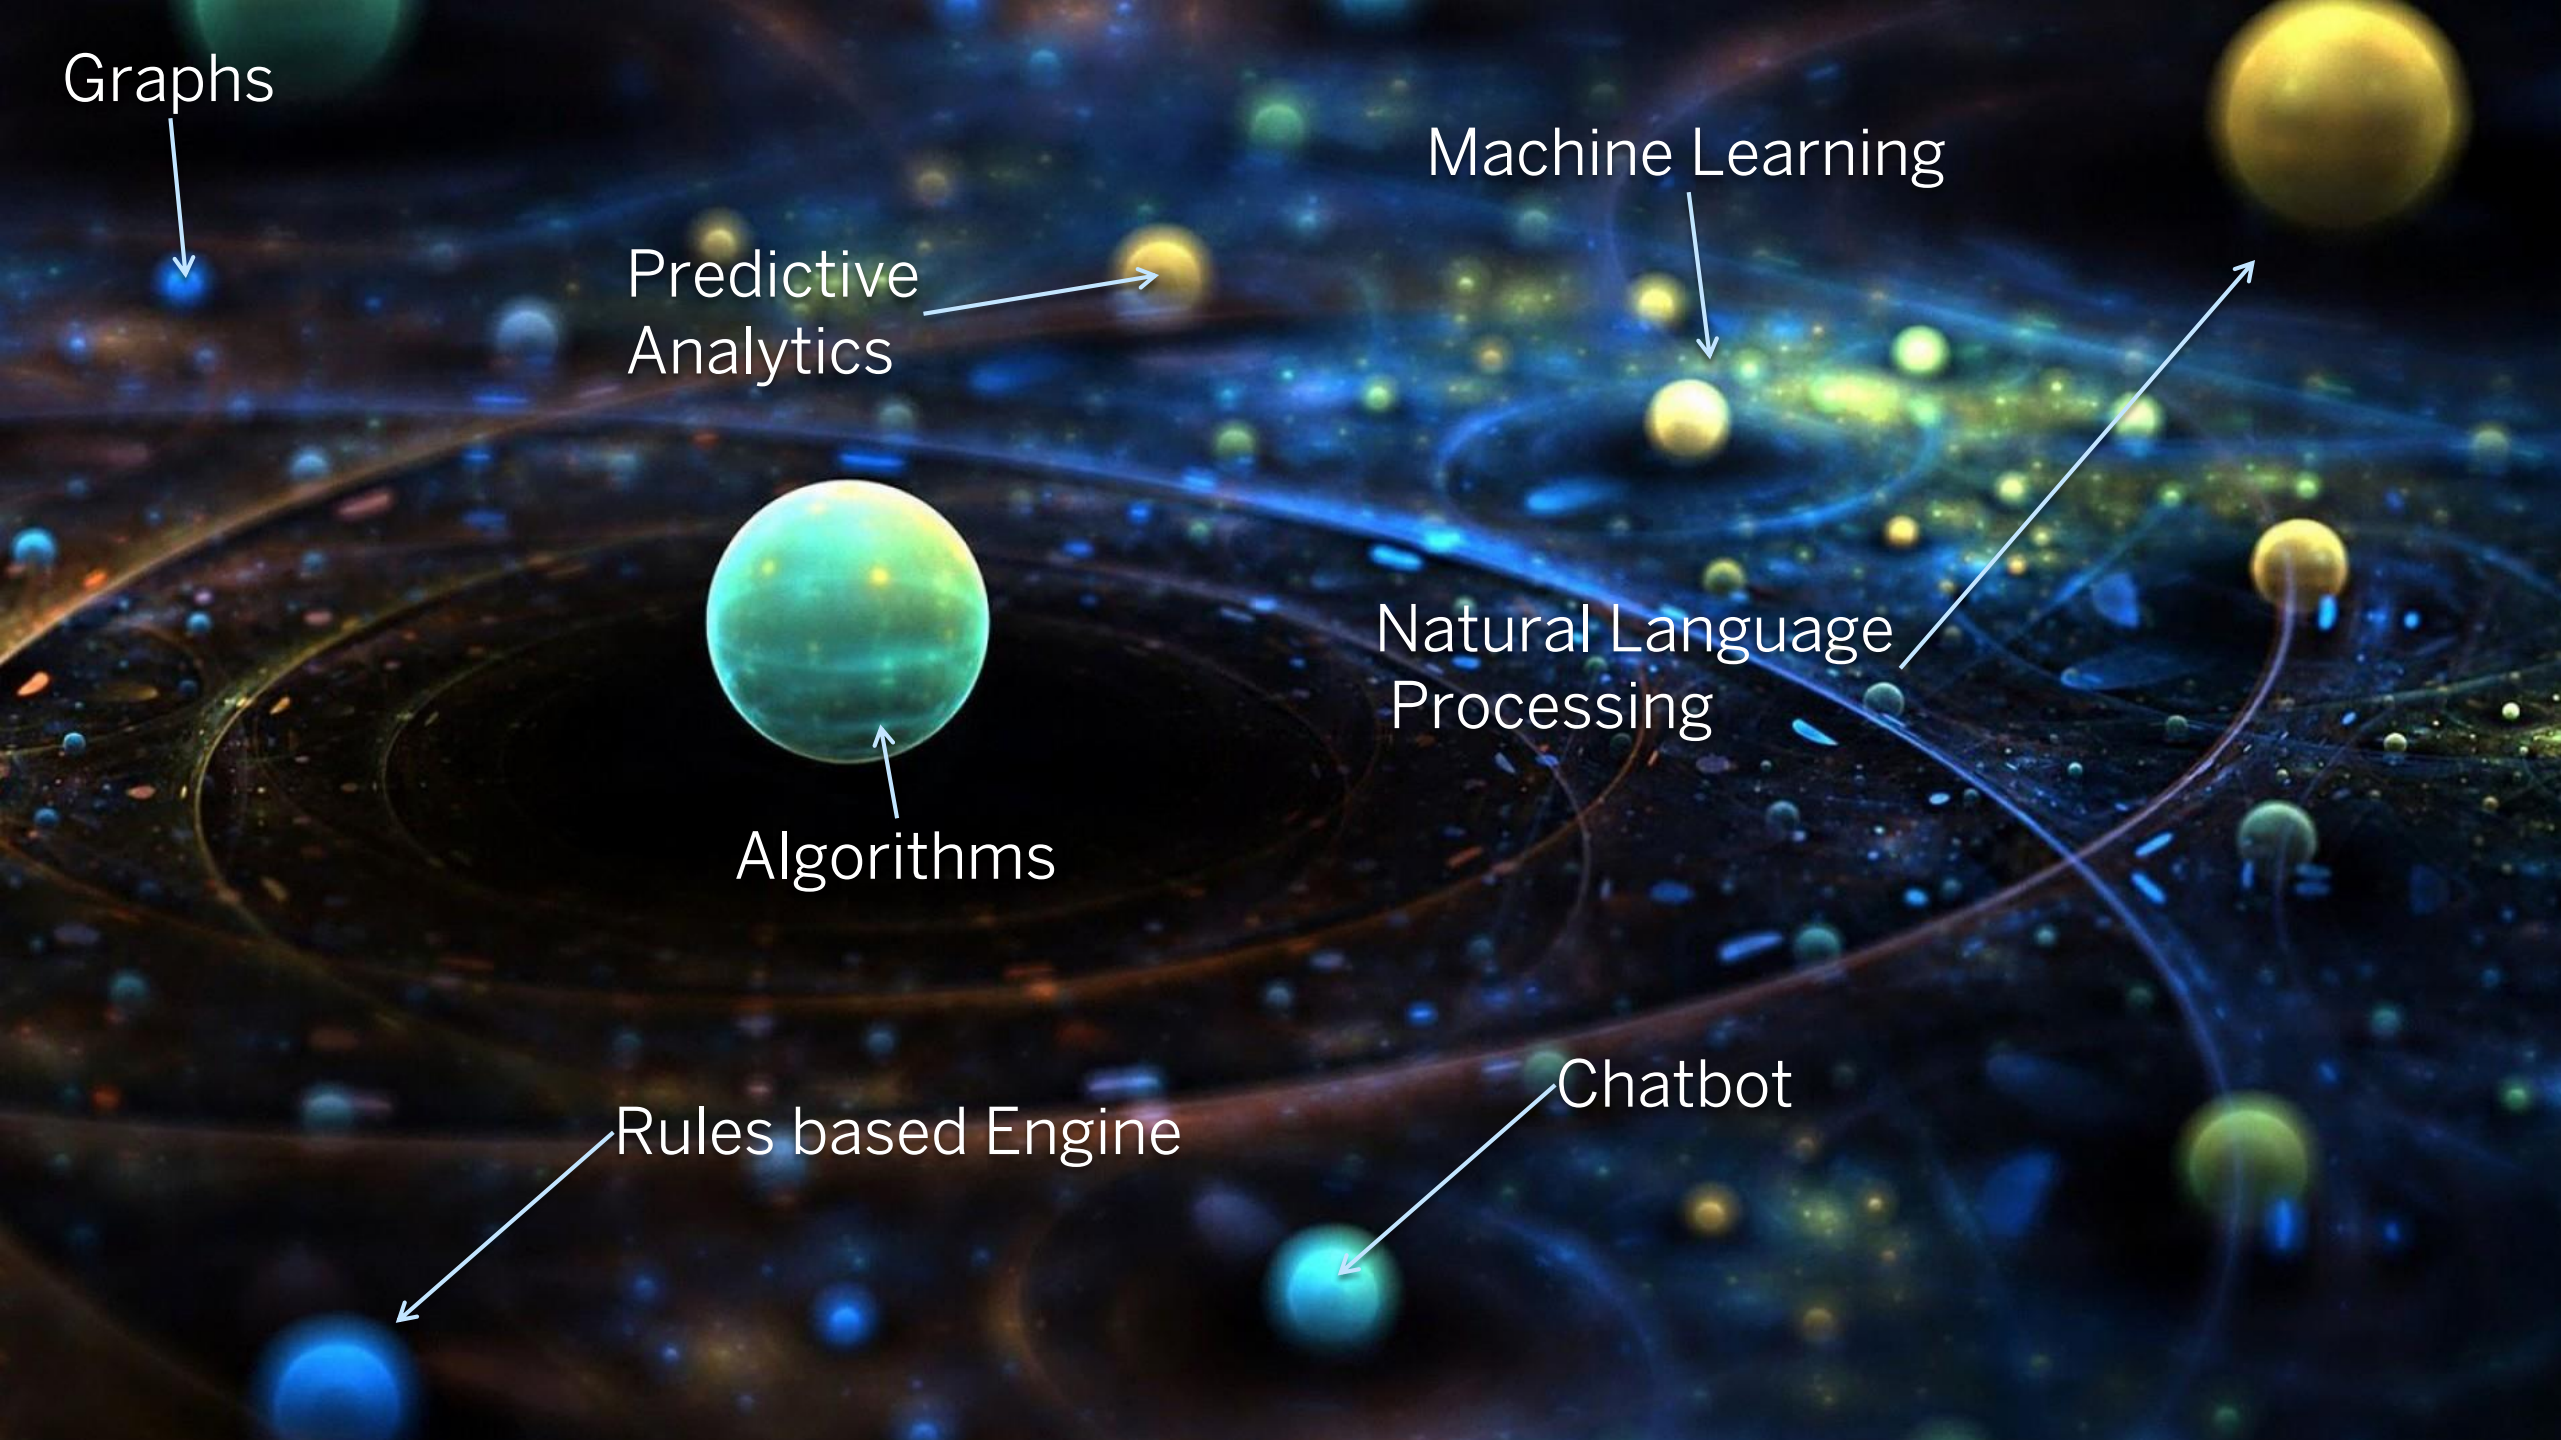
Graphs
Machine Learning
Predictive Analytics
Natural Language Processing
Algorithms
Chatbot
Rules based Engine

ML Guide

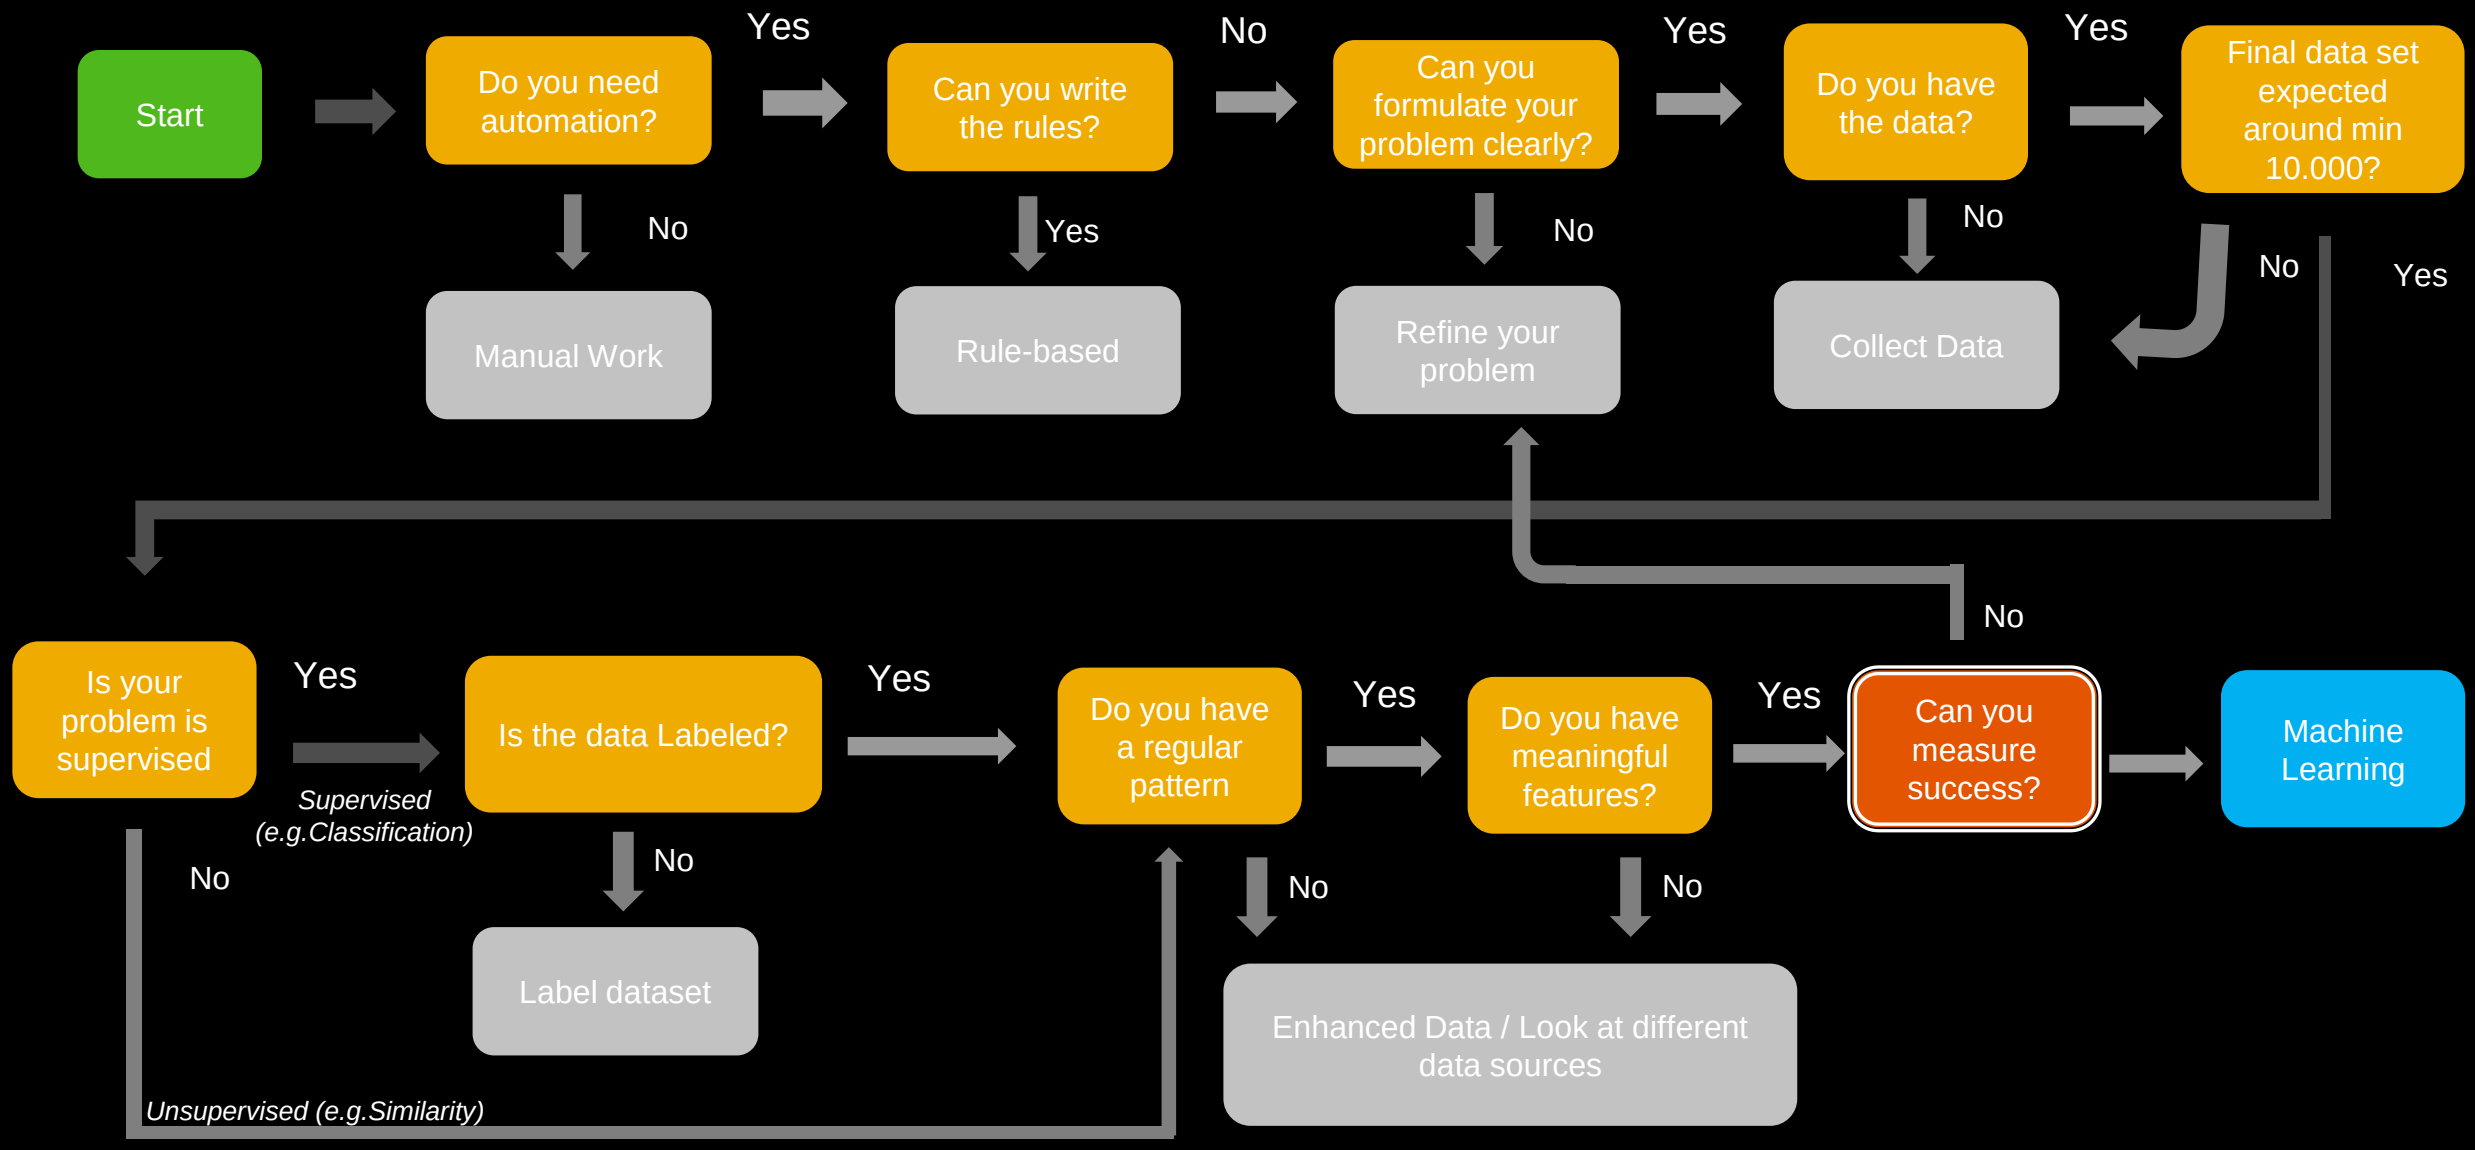

But: ....

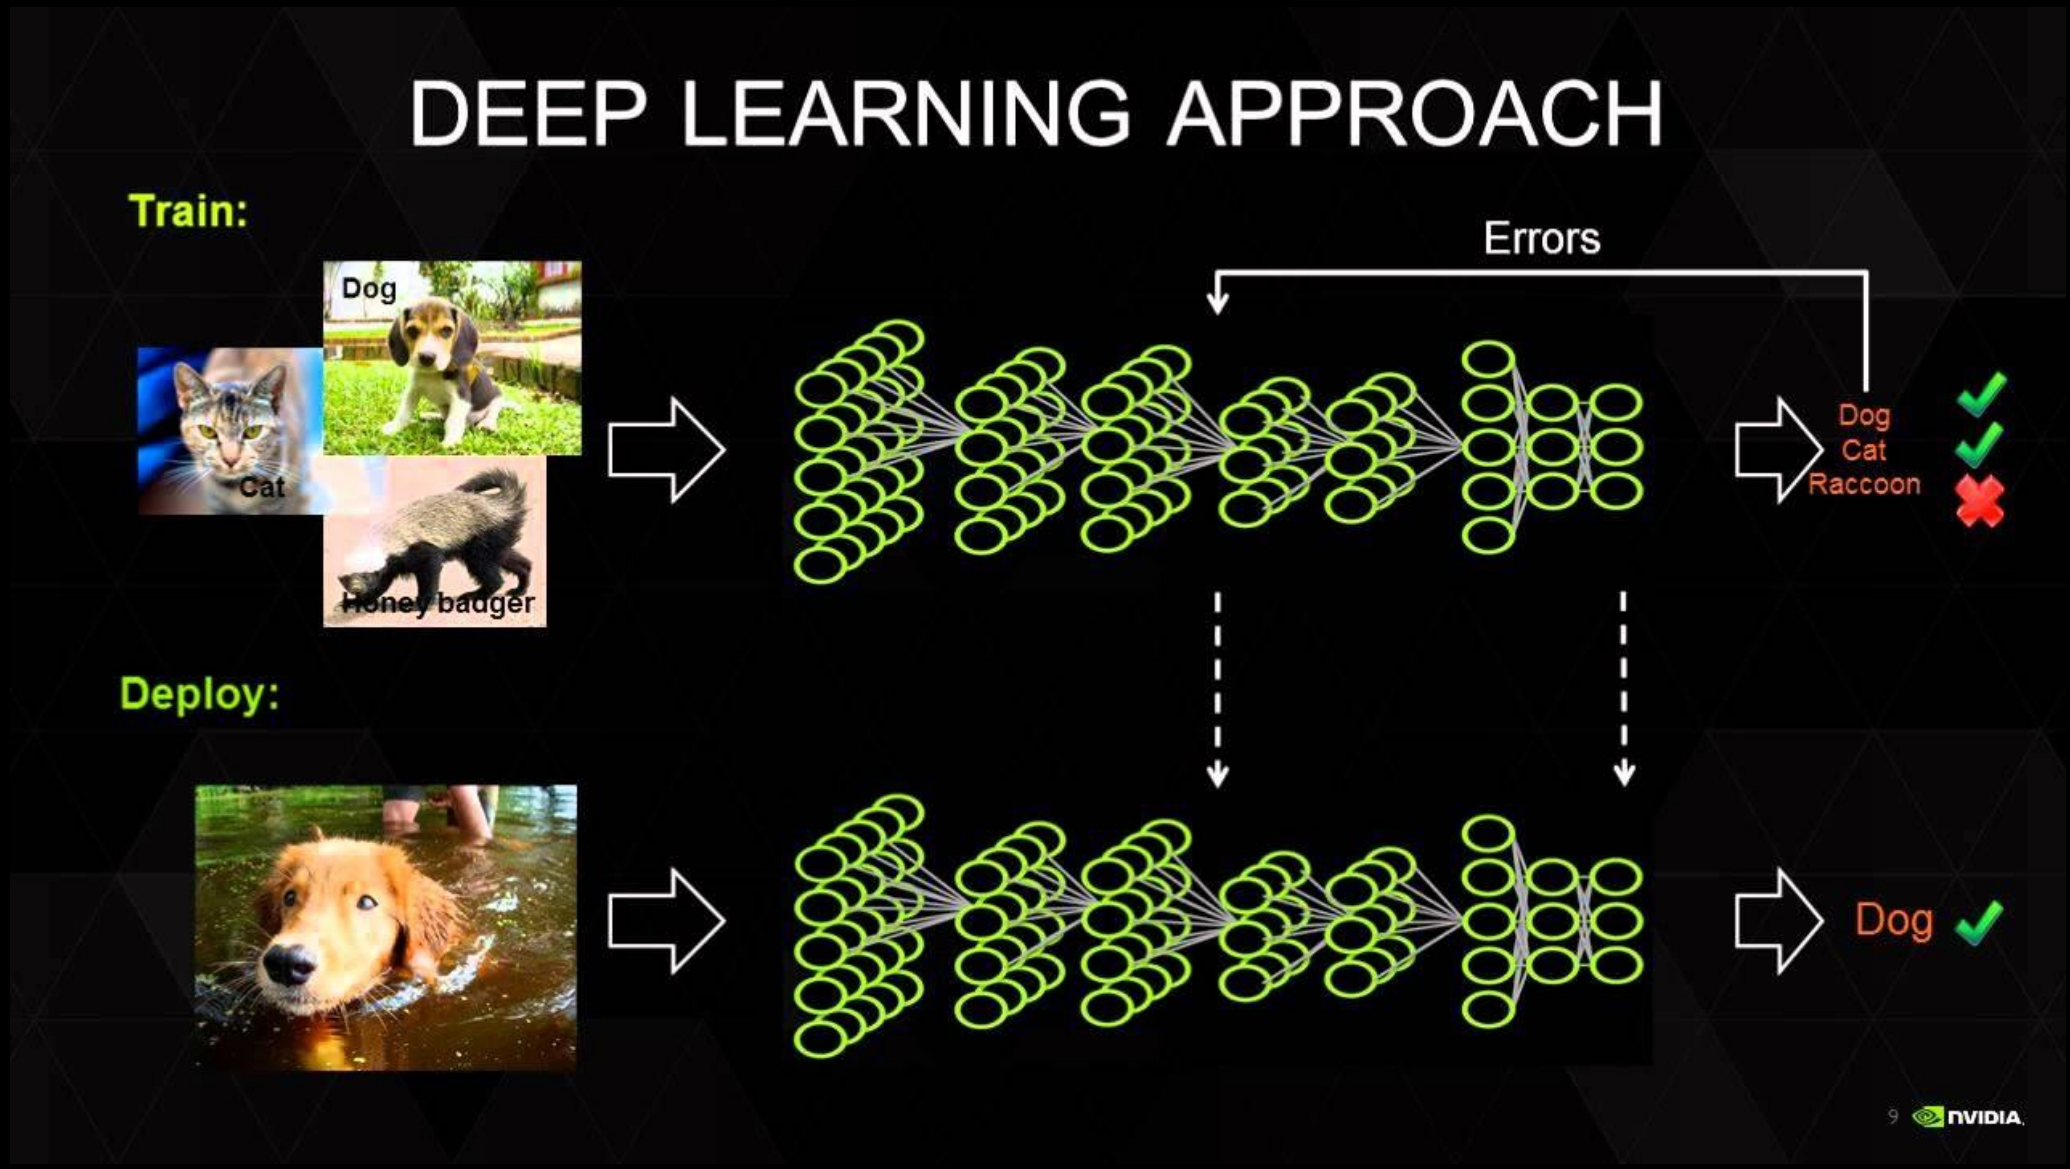

But: ....

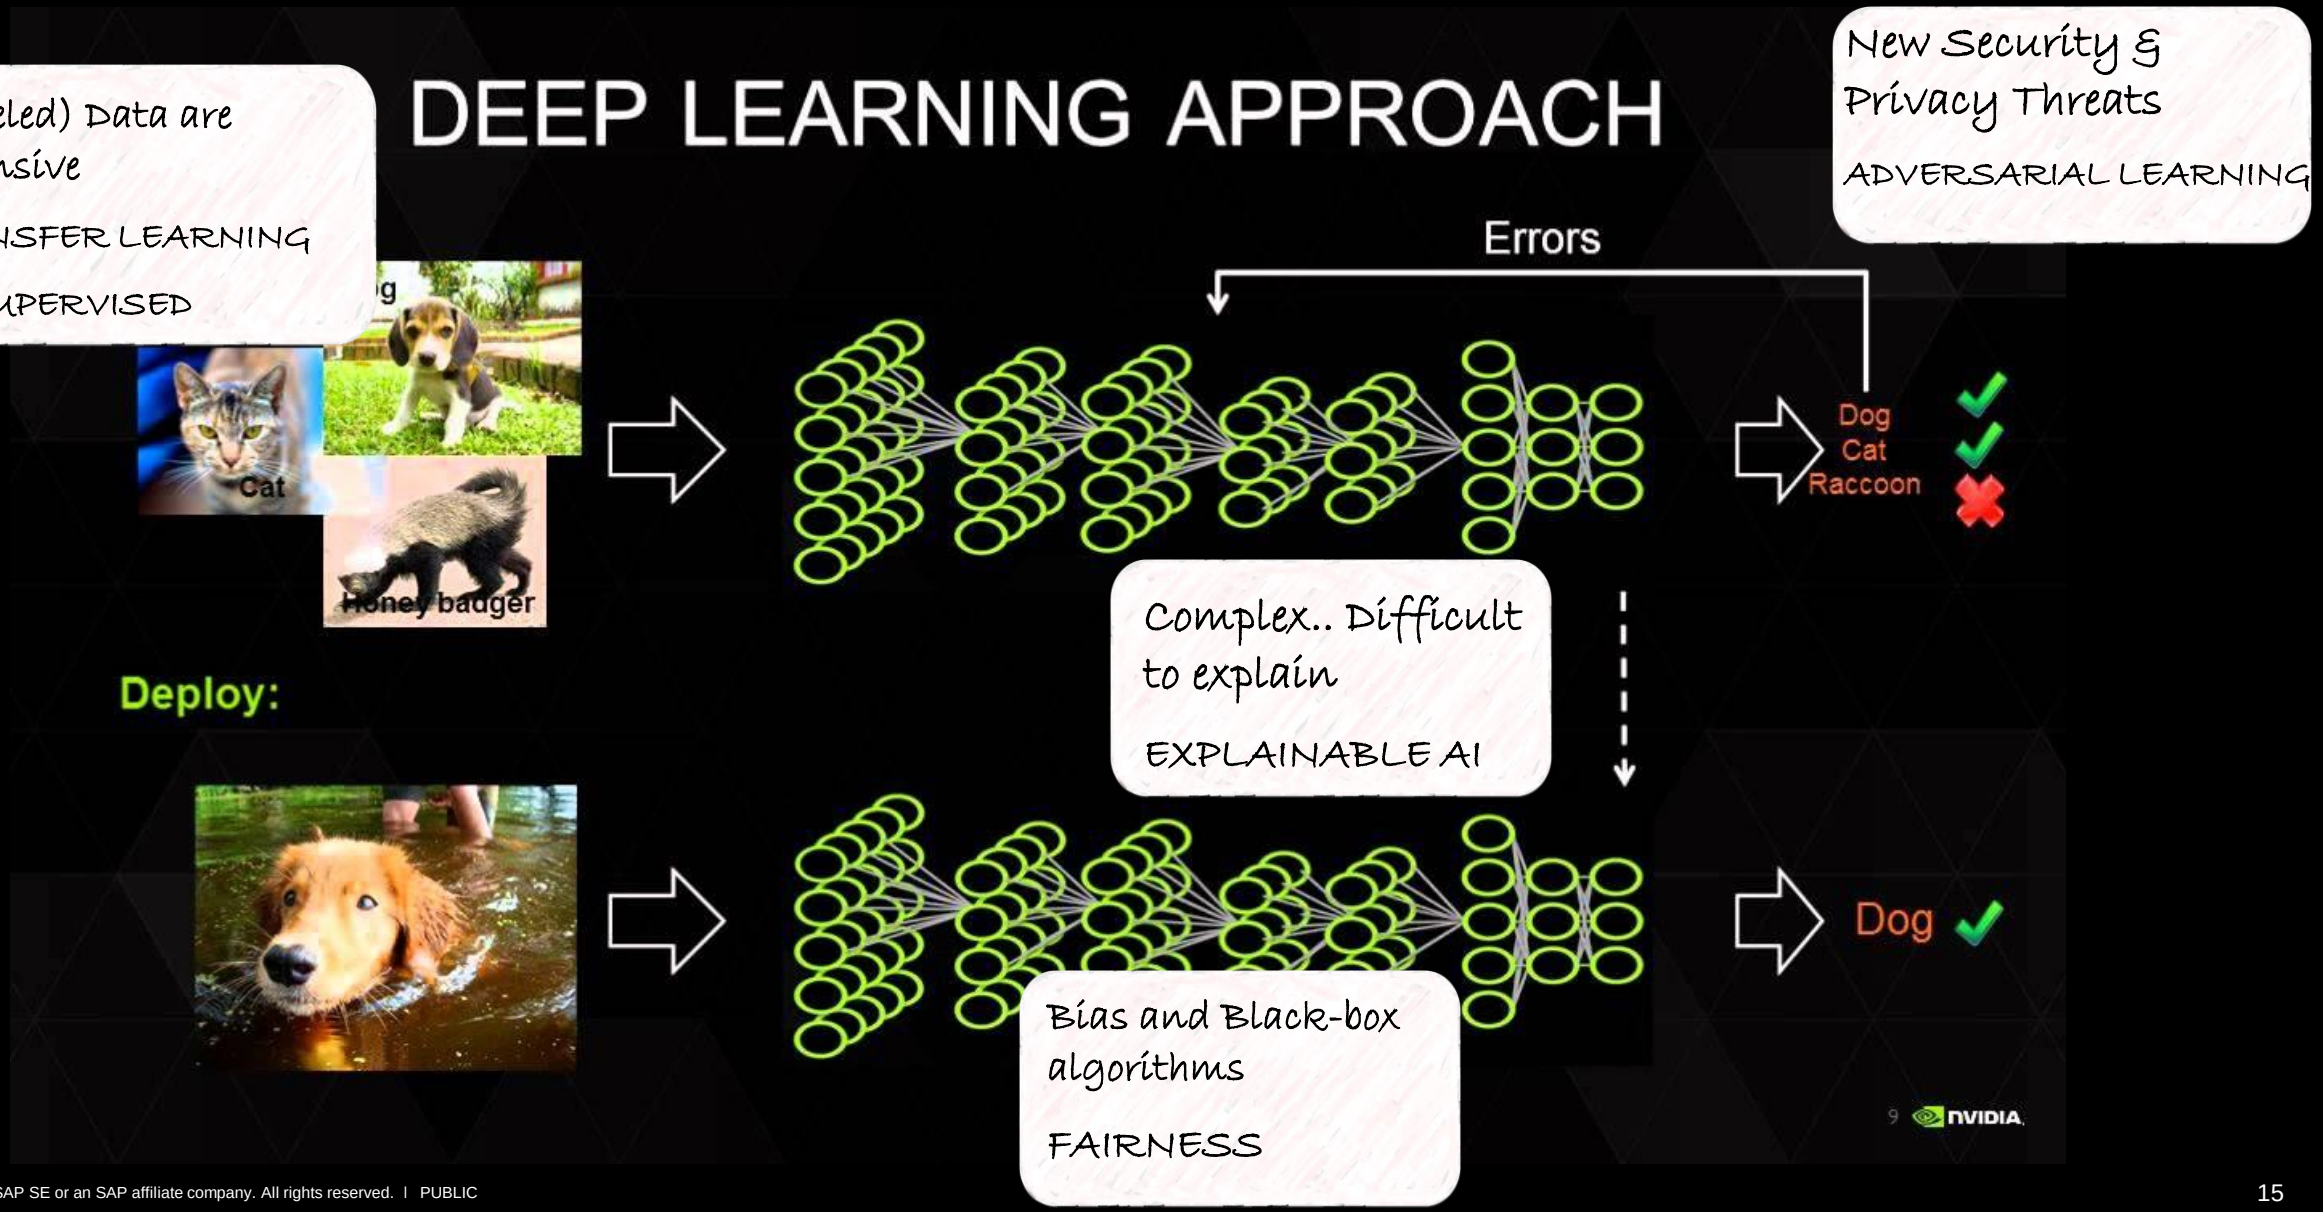

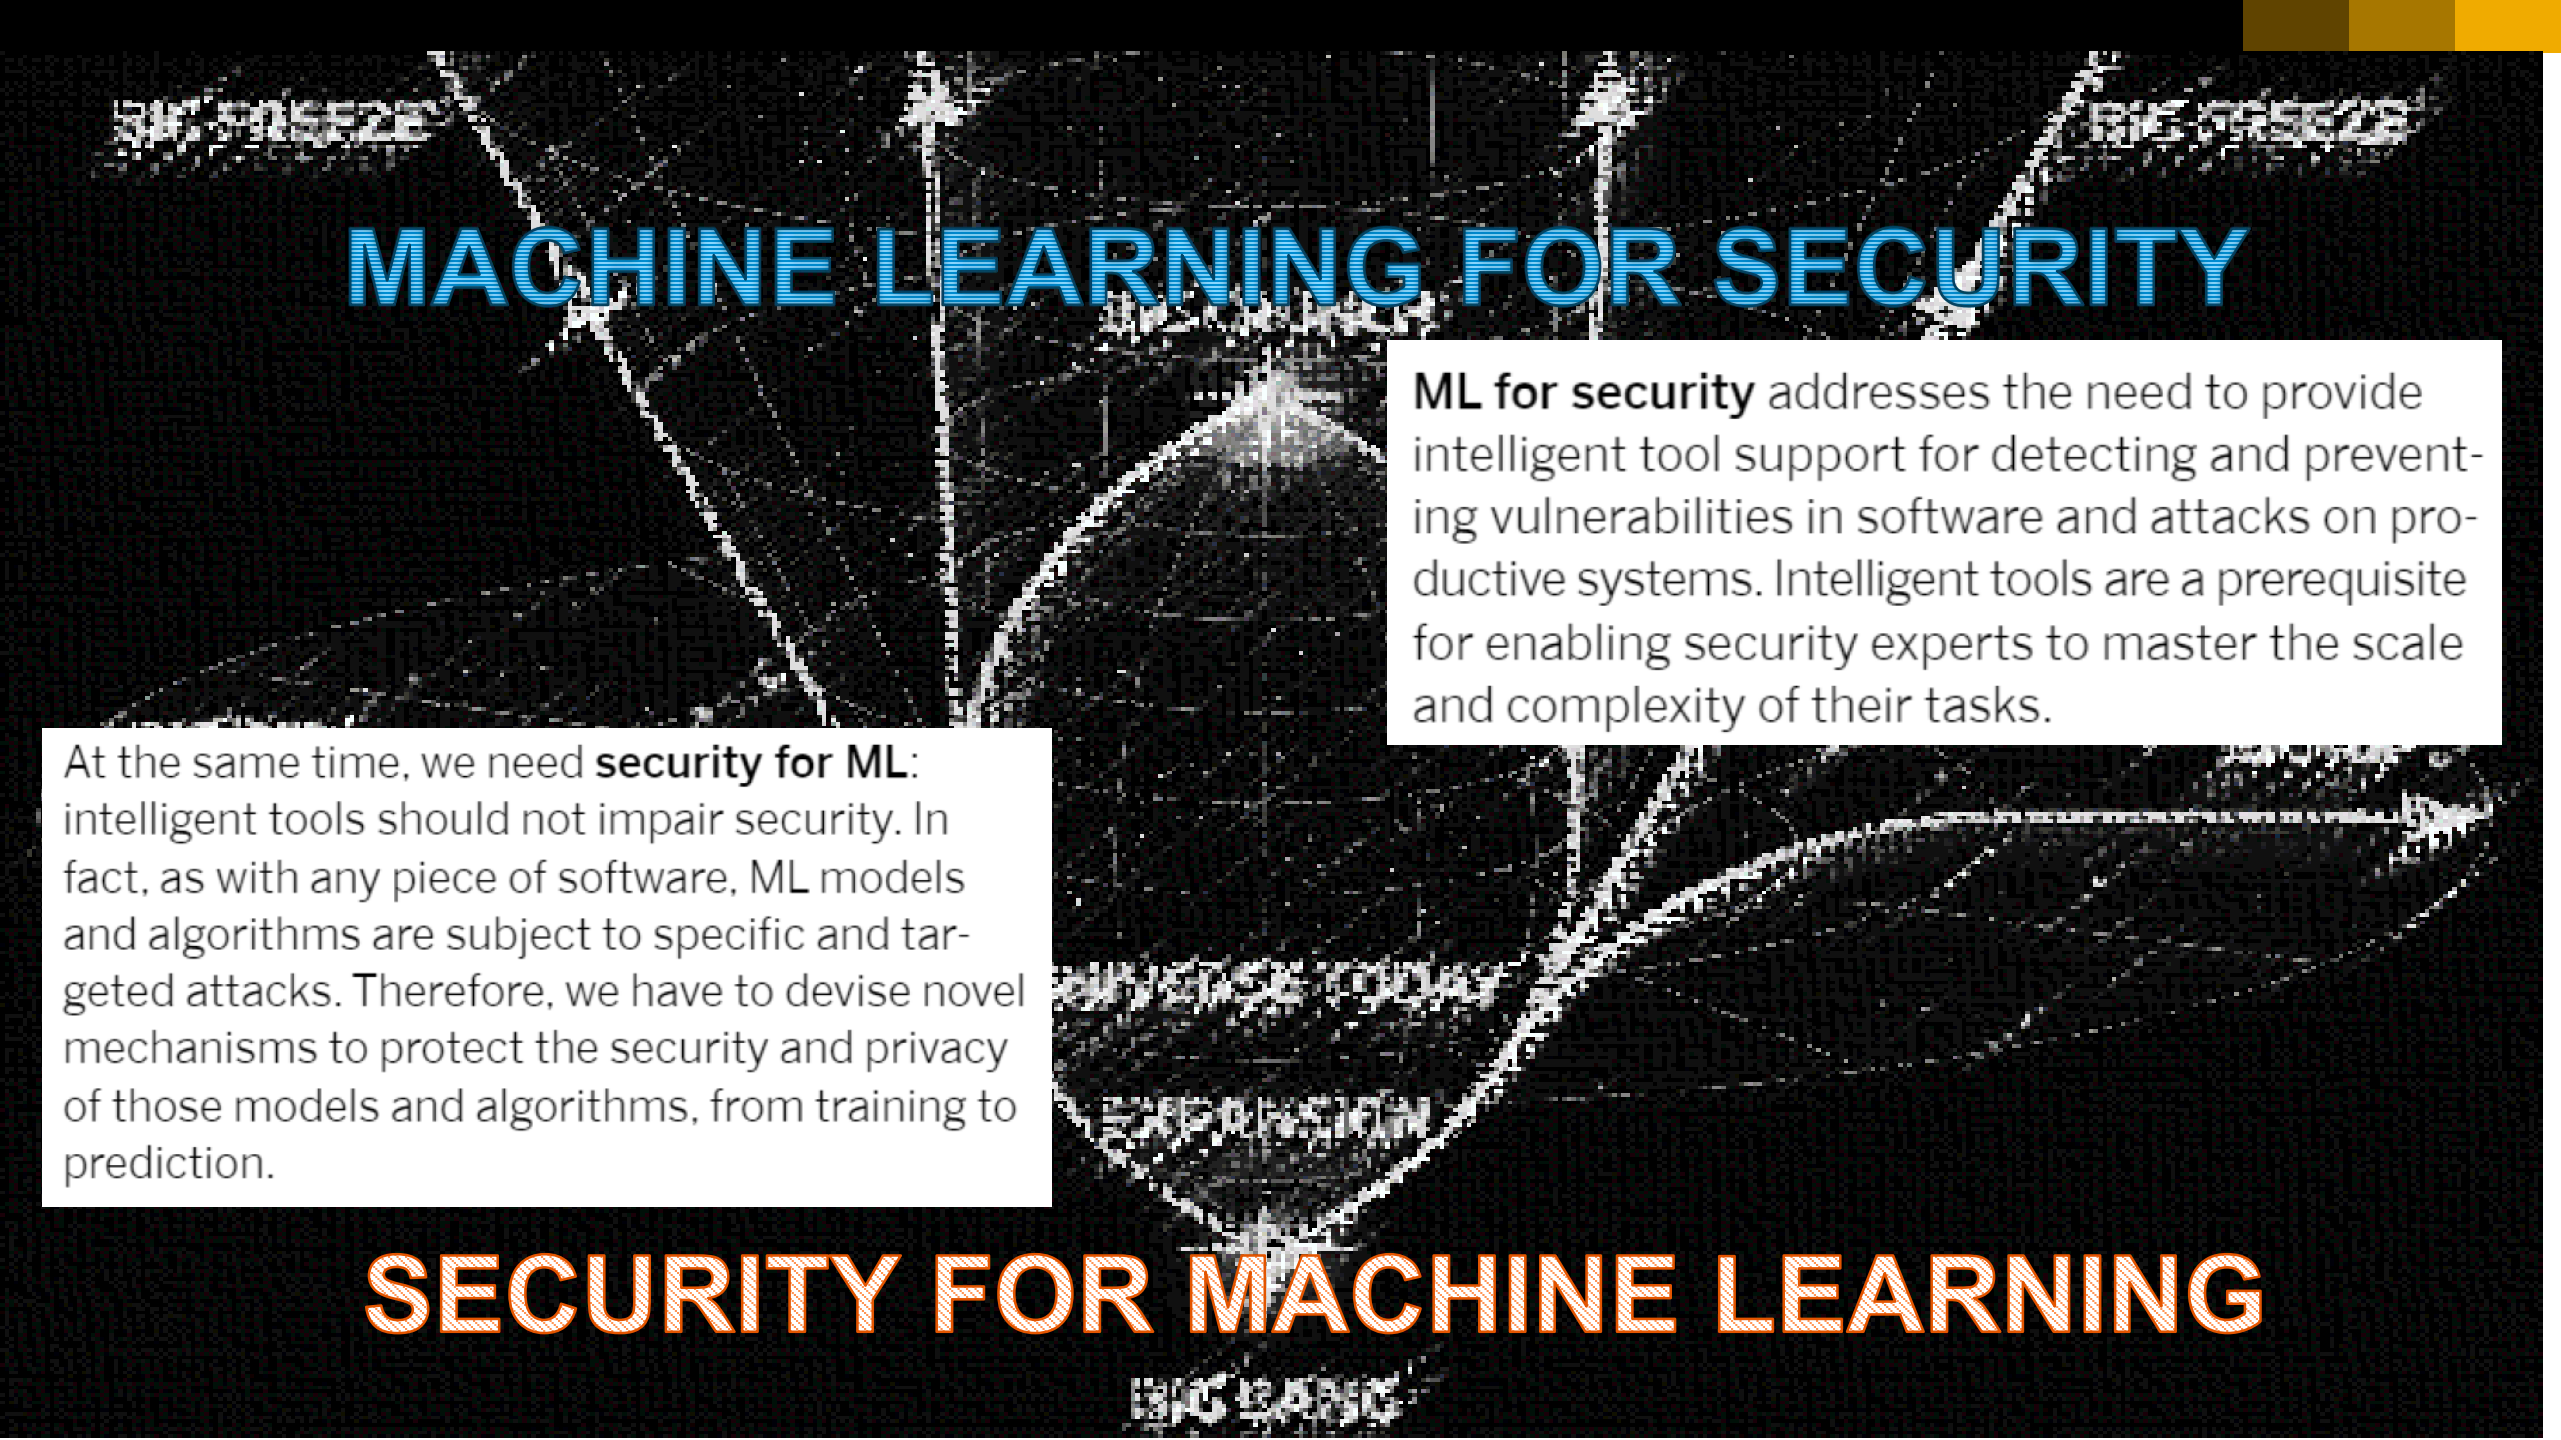
MACHINE LEARNING FOR SECURITY
ML for security addresses the need to provide intelligent tool support for detecting and preventing vulnerabilities in software and attacks on productive systems. Intelligent tools are a prerequisite for enabling security experts to master the scale and complexity of their tasks.
At the same time, we need security for ML: intelligent tools should not impair security. In fact, as with any piece of software, ML models and algorithms are subject to specific and targeted attacks. Therefore, we have to devise novel mechanisms to protect the security and privacy of those models and algorithms, from training to prediction.
SECURITY FOR MACHINE LEARNING

MACHINE LEARNING FOR SECURITY

Machine Learning for Security

Intelligent Code Analysis

- Code analysis is paramount to insure Security baseline
- Automation is the only way to do address Scalability
- Continuous analysis during the software lifecycle including support

Security Intelligence

- Holistic Security requires to gather Information from Multiple Sources (eg. Report, Social network, Darkweb)
- How to identify, classify, combine this information ?

Threats Detection

- Network, Application, Business Threats
- How to detect unusual system behaviour (zero day) ?
- How to recognize known attack patterns (including variants)?

A Practical Approach to the Automatic Classification of Security-Relevant Commits

Antonino Sabetta and Michele Bezzi
SAP Security Research
{antonino.sabetta,michele.bezzi}@sap.com

***Abstract*—The lack of reliable sources of detailed information on the vulnerabilities of open-source software (OSS) components is a major obstacle to maintaining a secure software supply chain and an effective vulnerability management process. Standard sources of advisories and vulnerability data, such as the National Vulnerability Database (NVD), are known to suffer from poor coverage and inconsistent quality.**

To reduce our dependency on these sources, we propose an approach that uses machine-learning to analyze source code repositories and to automatically identify commits that are *security-relevant* (i.e., that are likely to fix a vulnerability). We treat the source code changes introduced by commits as documents written in natural language, classifying them using

becomes *known to the public* (e.g., because a commit fixing it is pushed to the source code repository of the project) and the *actual availability of a new release* that includes that fix, can last a few days or even weeks (sometimes months). During such time-frame, a malicious party that would examine the code commits in an open-source project, could infer the existence of a security issue and exploit it before a fixed release is available th
This risk is pushing th
code repositories in rea
security-relevant change

Many of us help to train Machine Learning models on a daily basis without knowing …

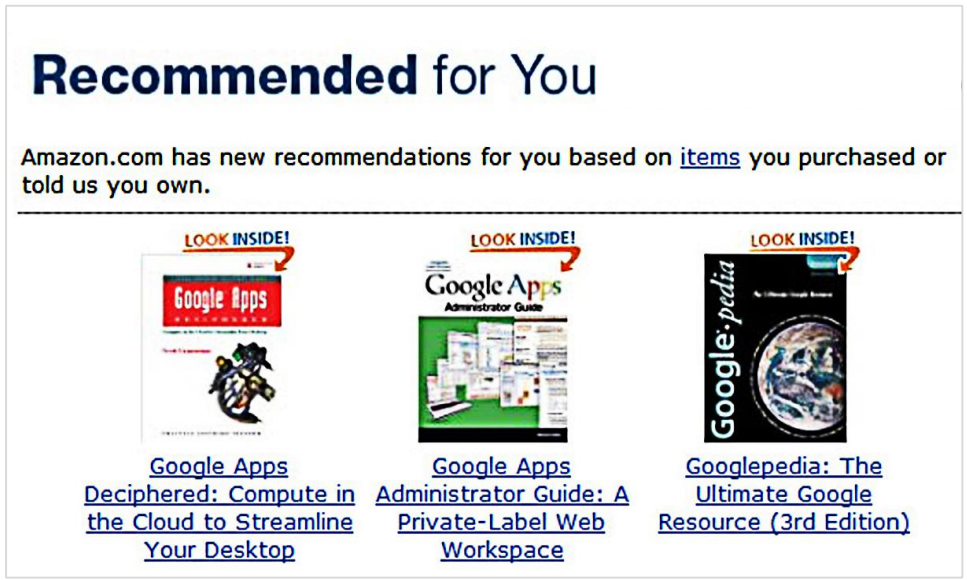

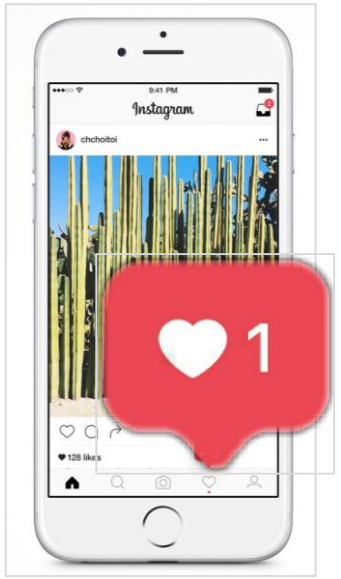

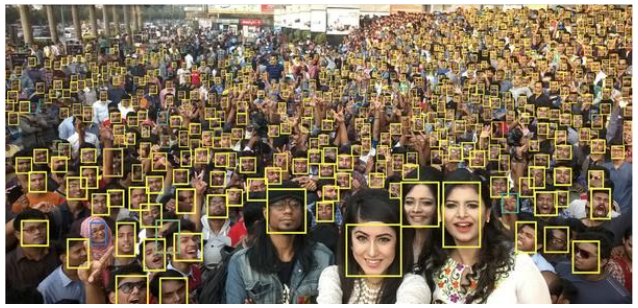

Creation of Value

Help everyone, everywhere - One CAPTCHA at a time.

Millions of CAPTCHAs are solved by people every day. reCAPTCHA makes positive use of this human effort by channeling the time spent solving CAPTCHAs into digitizing text, annotating images, and building machine learning datasets. This in turn helps preserve books, improve maps, and solve hard AI problems.

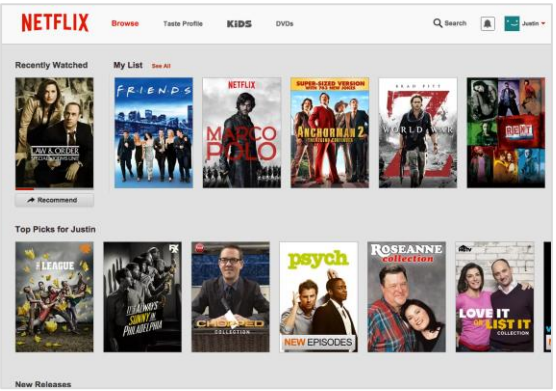

Machine Learning for Security: What makes it different/harder?

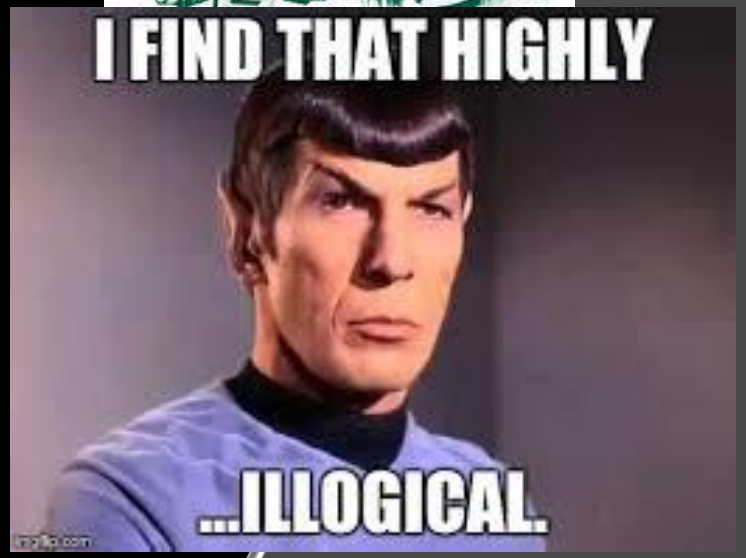

Unusually high cost of *errors*

- *Misclassifying a legitimate access request, makes people angrier than a bad movie recommendation → avoid False Positive*

Lack of training (labeled) data

- *People prefer sharing cats pictures than security (sensitive) data*

Meaningful explanations

- *Human decisions need to be supported by human-understandable algorithms results*

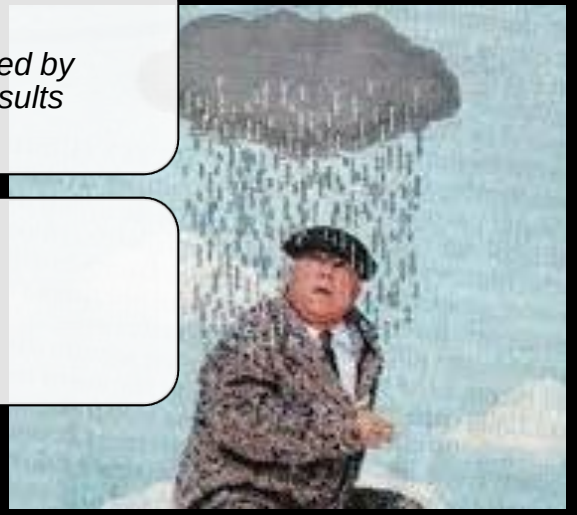

Adaptive Adversaries

- *Weather is not trying to avoid classification/prediction, hackers do*

SECURITY FOR MACHINE LEARNING

Attack to Machine Learning: Evasion Attack → Toxic Signs

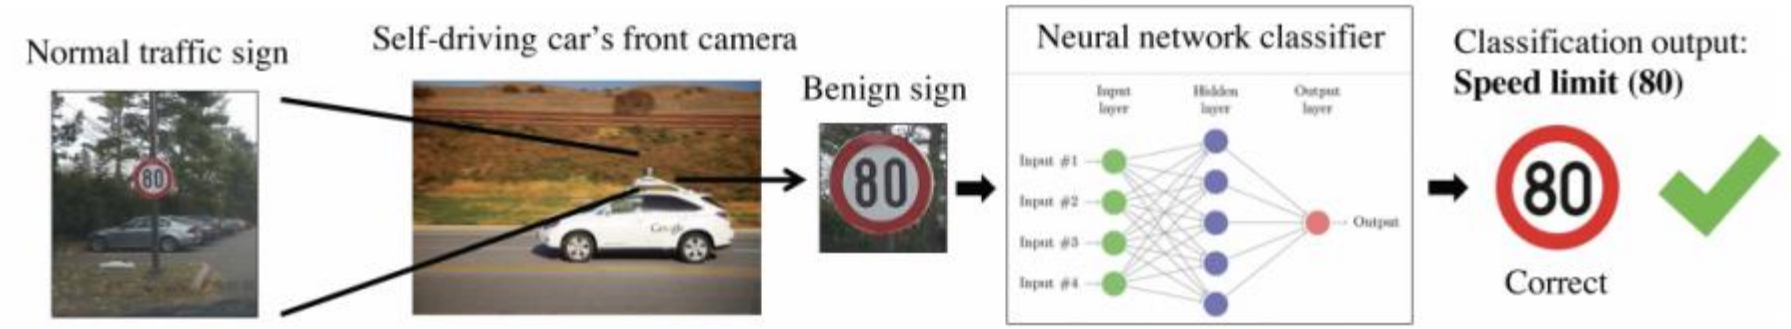

(a) Operation of the computer vision subsystem of an AV under *benign conditions*

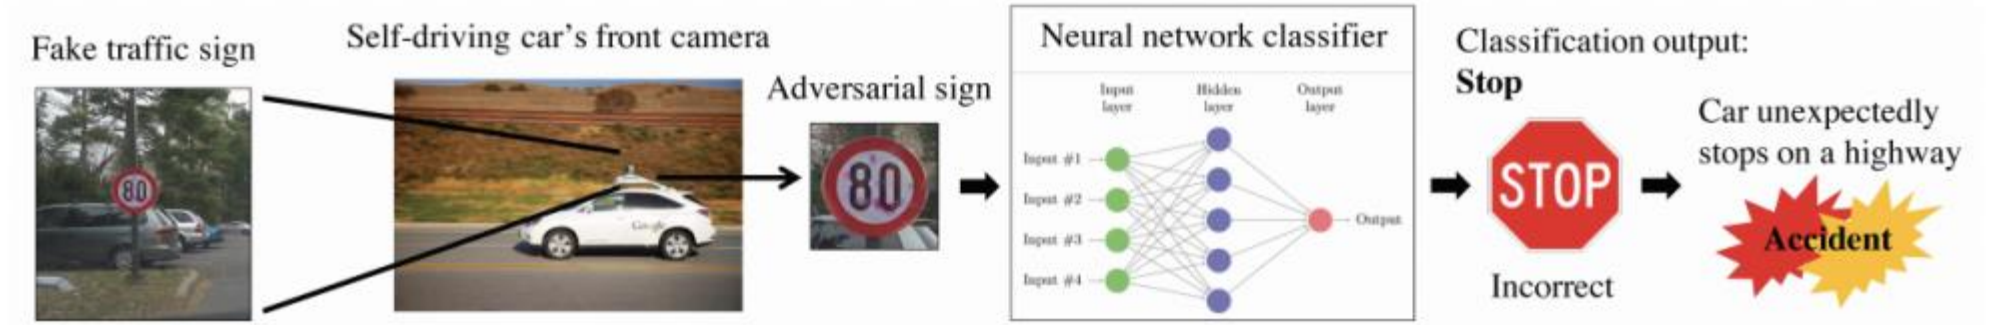

(b) Operation of the computer vision subsystem of an AV under *adversarial conditions*

Fig. 1. **Difference in operation of autonomous cars under benign and adversarial conditions**. Figure 1b shows the classification result for a drive-by test for a physically robust adversarial example generated using our Adversarial Traffic Sign attack.

Sitawarin at al. 2018

When smart machines, driven with artificial intelligences, will assist Courtroom ?

2017, Eric Loomis was sentenced to six years in prison at least in part by the recommendation of private company's secret proprietary software (COMPAS)

Risk: **Blackboxing the Judicial System**

ML (DNN) algorithms decision-process is hard to understand:

- Decision (at best) as good and fair as data … they reproduce bias
- Algorithms used as support for decision making needs to provide human-understandable motivation

GDPR gives you the right to explanation

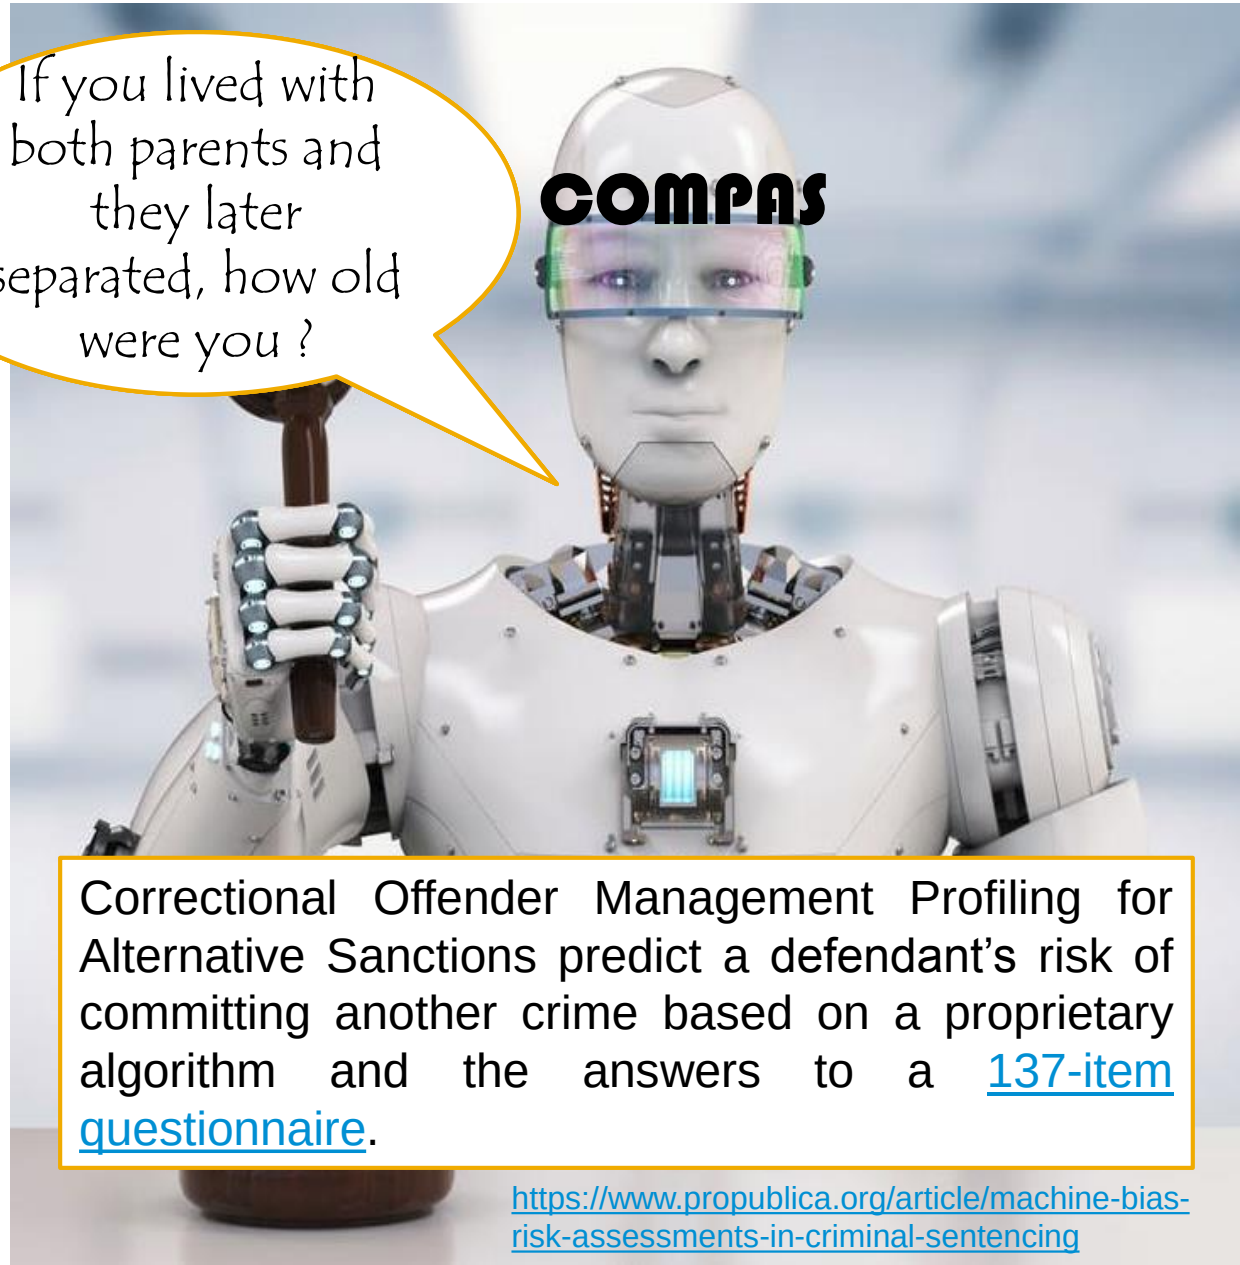

Correctional Offender Management Profiling for Alternative Sanctions predict a defendant's risk of committing another crime based on a proprietary algorithm and the answers to a 137-item questionnaire.

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

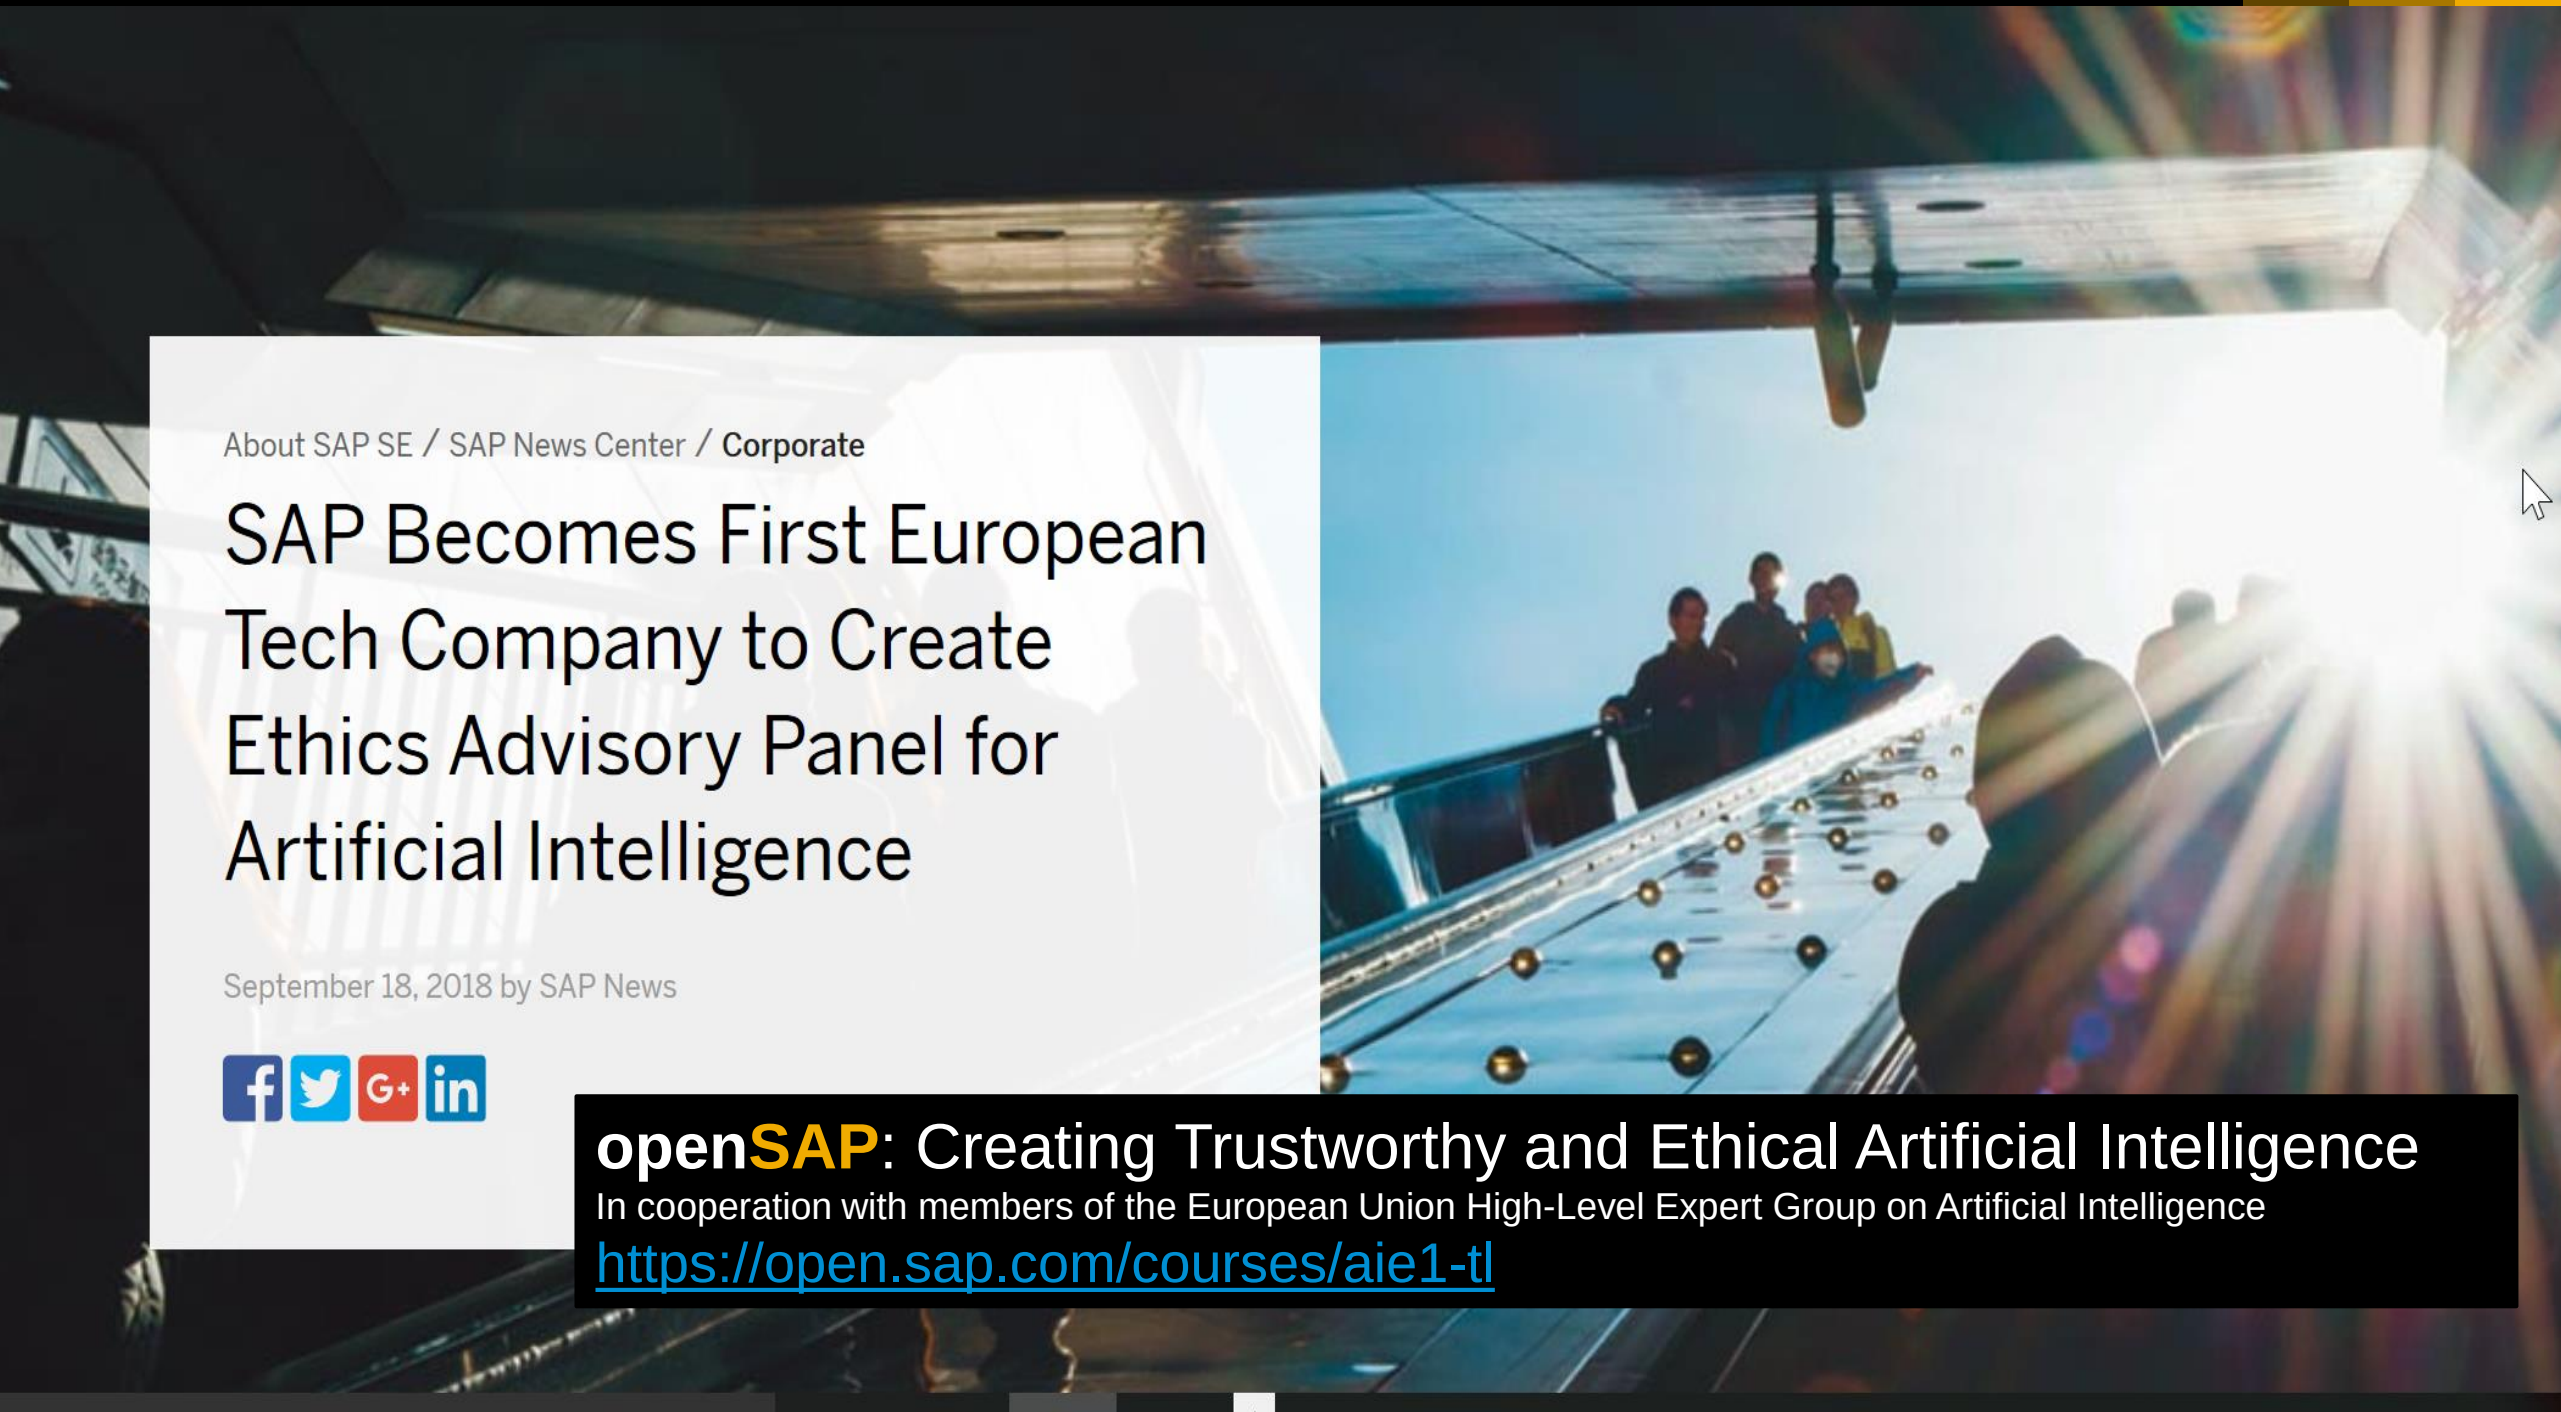

About SAP SE / SAP News Center / **Corporate**

SAP Becomes First European Tech Company to Create Ethics Advisory Panel for Artificial Intelligence

September 18, 2018 by SAP News

openSAP: Creating Trustworthy and Ethical Artificial Intelligence

In cooperation with members of the European Union High-Level Expert Group on Artificial Intelligence

https://open.sap.com/courses/aie1-tl

Security for Machine Learning

Attacks

- Different vectors (eg. training set, model, input , output) and different attacks (eg. Poisoning, evasion, toxic)
- Is the Deployed Model immune to them ?

Privacy/Confidentiality

- Training Data contains Confidential Information
- How to insure that they cannot be Reversed Enginneered ?

Intellectual Property

- Considerable effort and know how in making THE solution (eg. Topology, Data Cleansing, (re)Training)
- How to protect it from Competitors?

Explainability / Trust / Compliance

- How to explain results ?
- How to trust a model (eg. Coverage of training set, unbiased set) ?
- How to certify a Model ?

The need for an AI Partnership

BDV BIG DATA VALUE ASSOCIATION

BDVA – euRobotics MoU (ICT 2018 event)

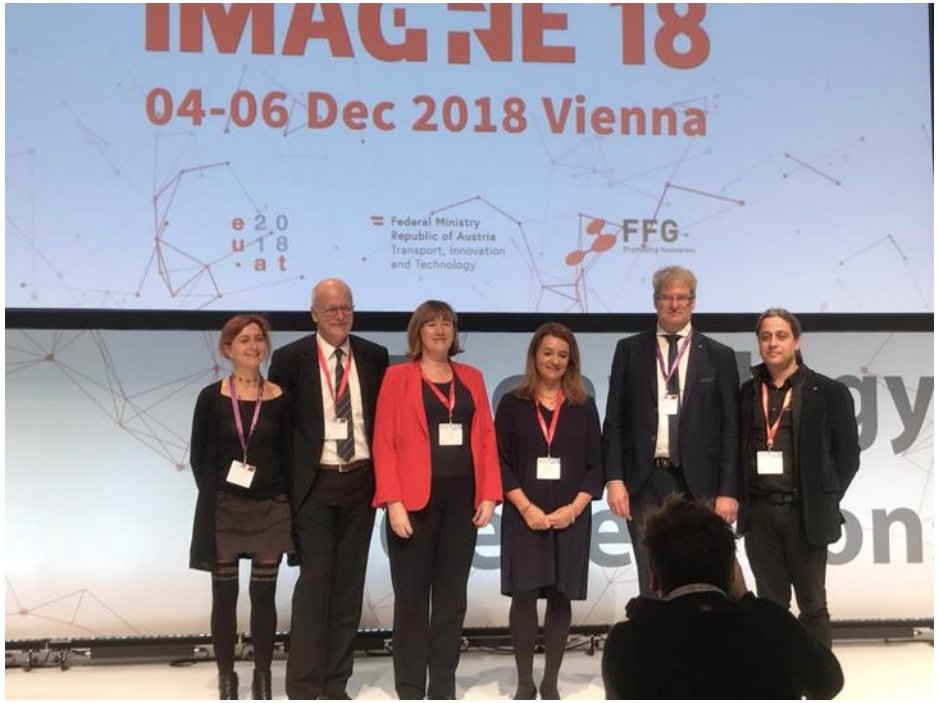

https://ec.europa.eu/digital-single-market/en/news/artificial-intelligence-public-private-partnerships-join-forces-boost-ai-progress-europe

BDVA – euRobotics common Vision Paper

March 2019

AI Value Chain

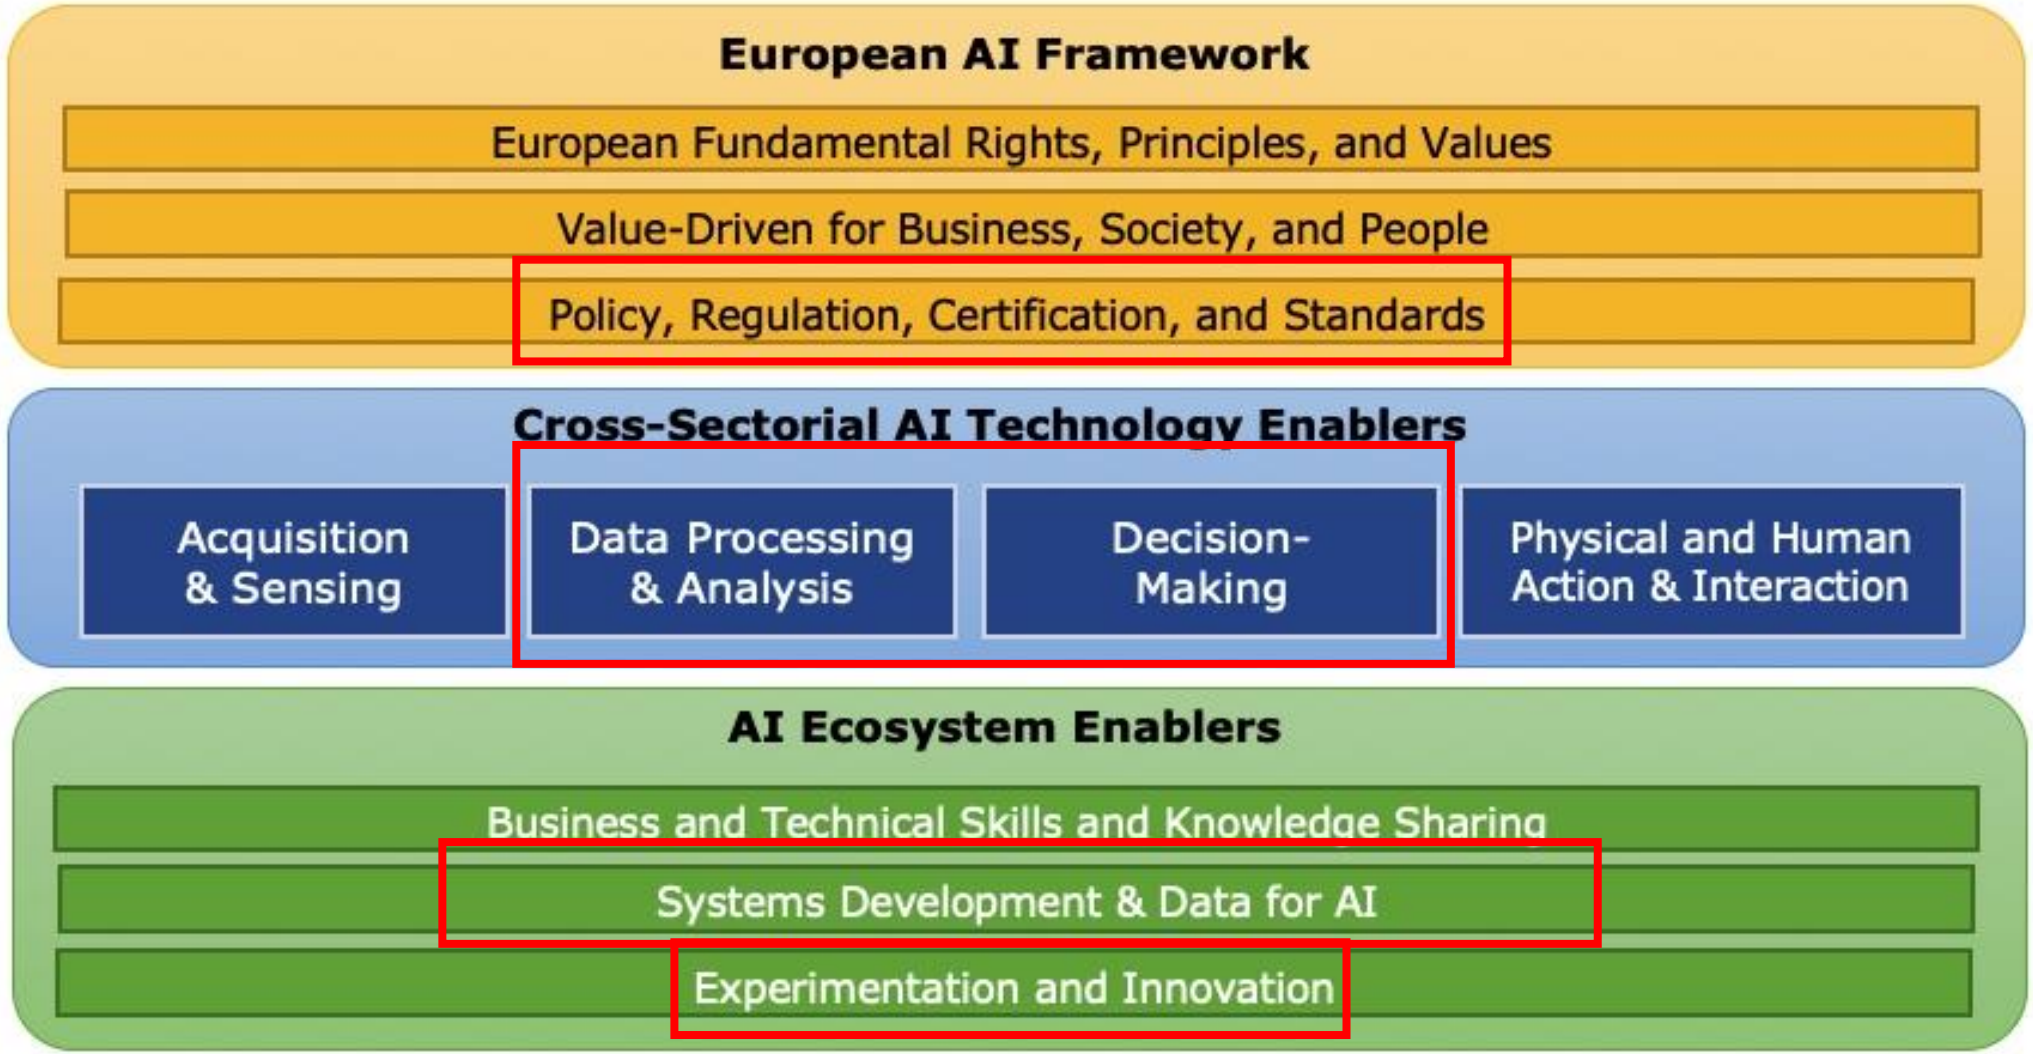

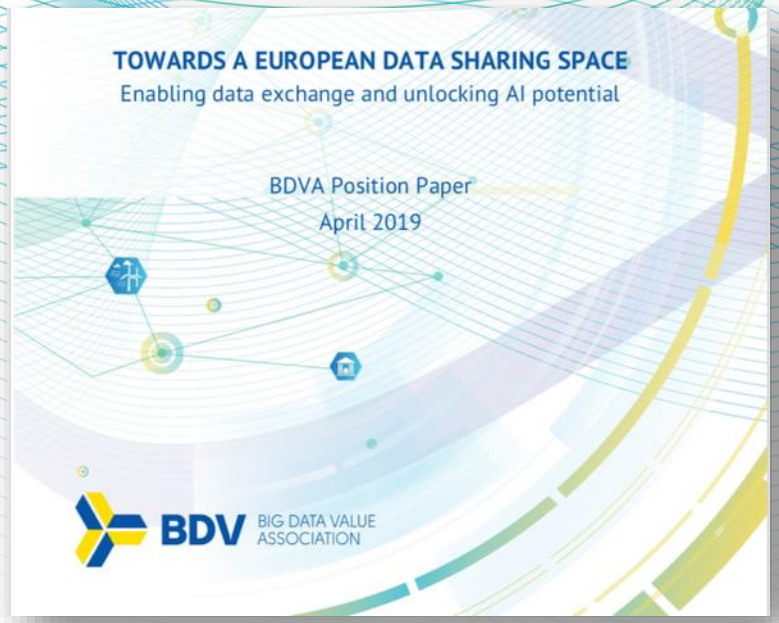

BDV BIG DATA VALUE ASSOCIATION

www.bdva.eu

BDV PPP SUMMIT

JUNE 26 - 28 2019, RIGA

Impact empowered by Data-Driven AI

BIG DATA VALUE
PUBLIC-PRIVATE PARTNERSHIP

BDV PPP CONFERENCE + BDV PPP MEETUP

http://summit.big-data-value.eu/

https://www.european-big-data-value-forum.eu/

Thank you.

Jean-Christophe Pazzaglia, Phd
Chief Support Architect Higher Education & Research (PS)

SAP Labs France SAS | 805 av du Dr Maurice Donat |
BP1216 | 06254 Mougins Cedex, France

Firstname.lastname@sap.com

THE BEST RUN SAP®