The long and winding road to Self-driving networks

IFIP Networking

e-Paris, June 2020

Dario Rossi
Chief Expert, Network AI
Director, DataCom* Paris Lab
dario.rossi@huawei.com

(*) Data Communication Network Algorithm & Measurement Technology Laboratory

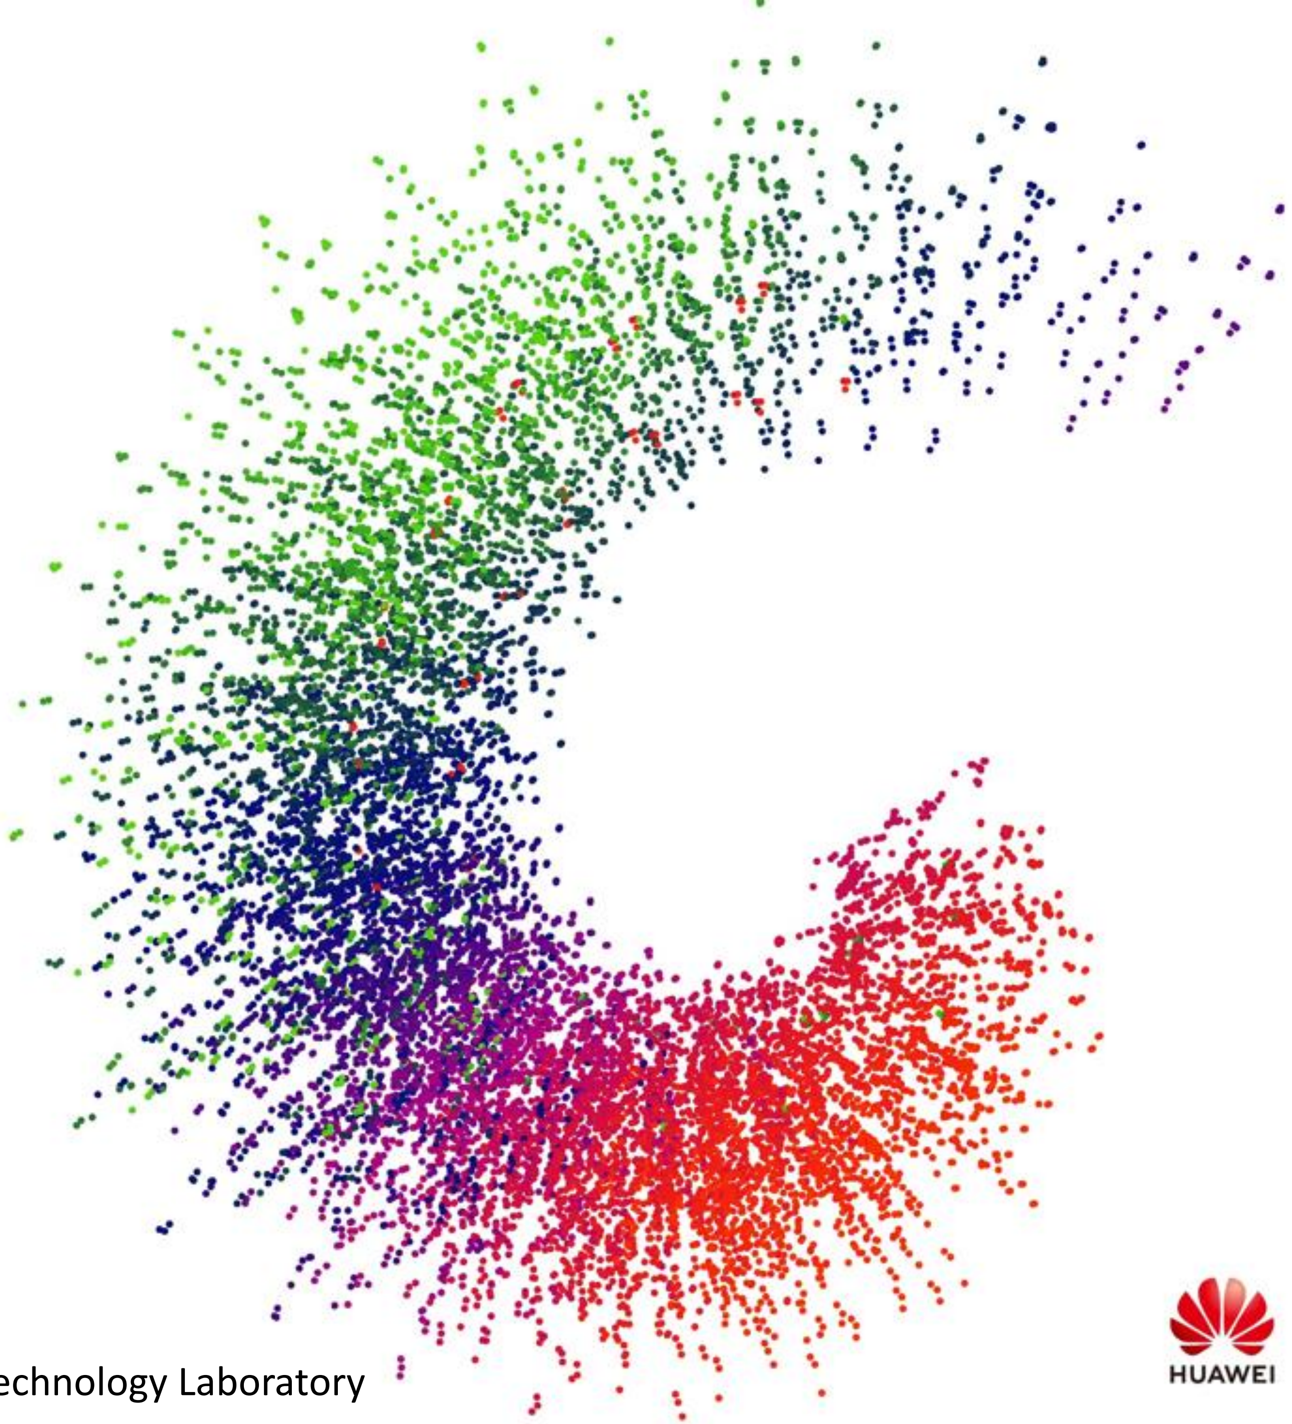

HUAWEI

Absence
of information

Encryption
operational obscurity

Excess
of information

Data deluge
operational overload

Opportunity
for AI & ML

Ascend
Unified AI chip architecture

Tackle
operational obscurity & operational overload

Huawei's IDN in a nutshell

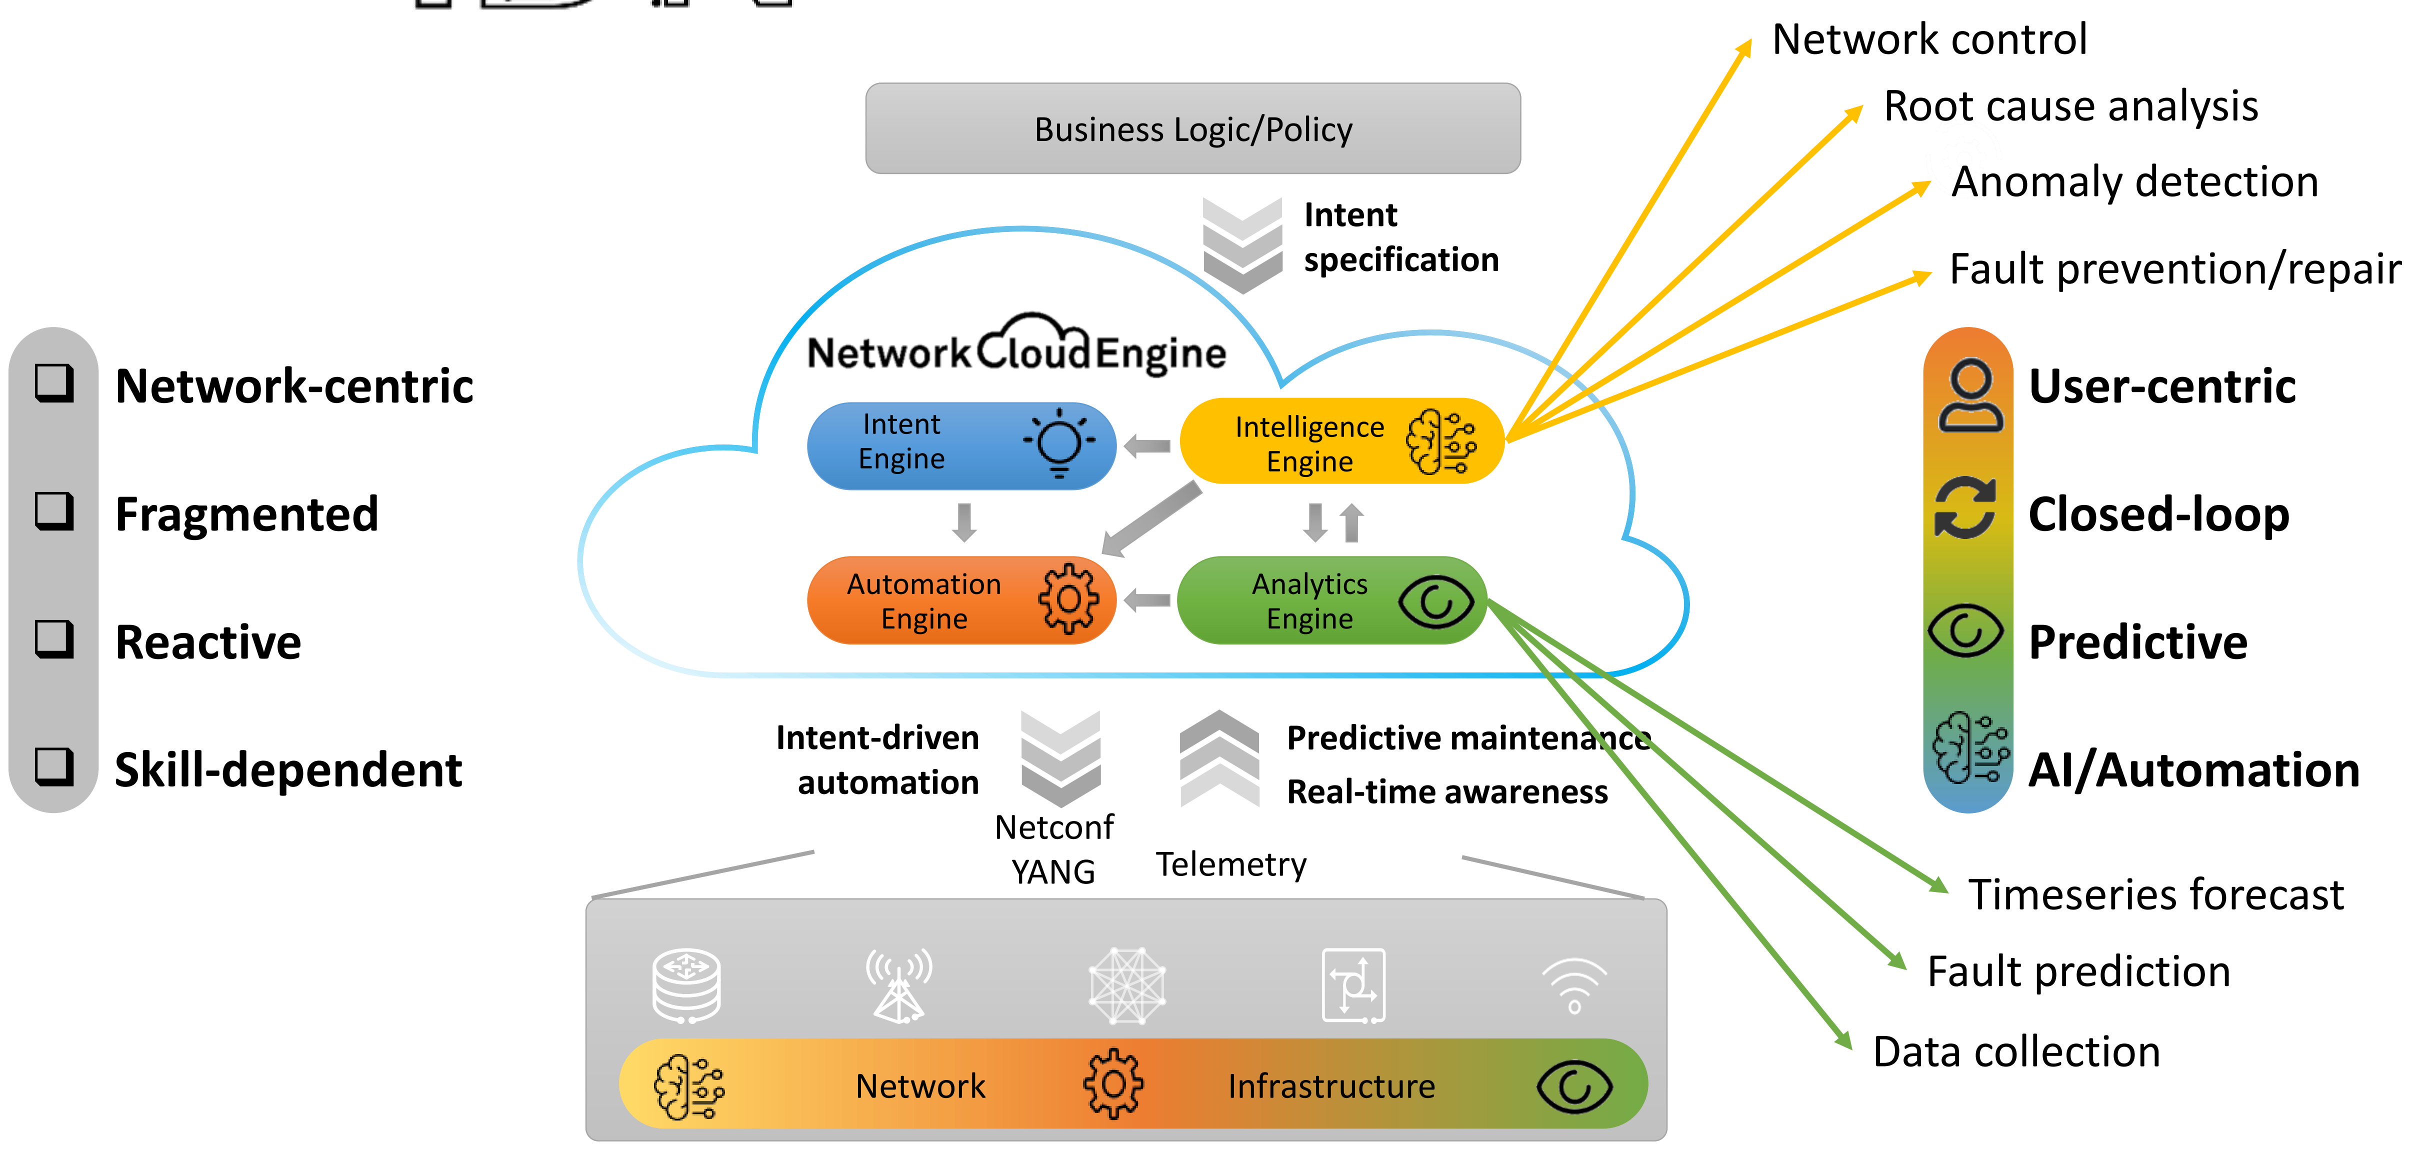

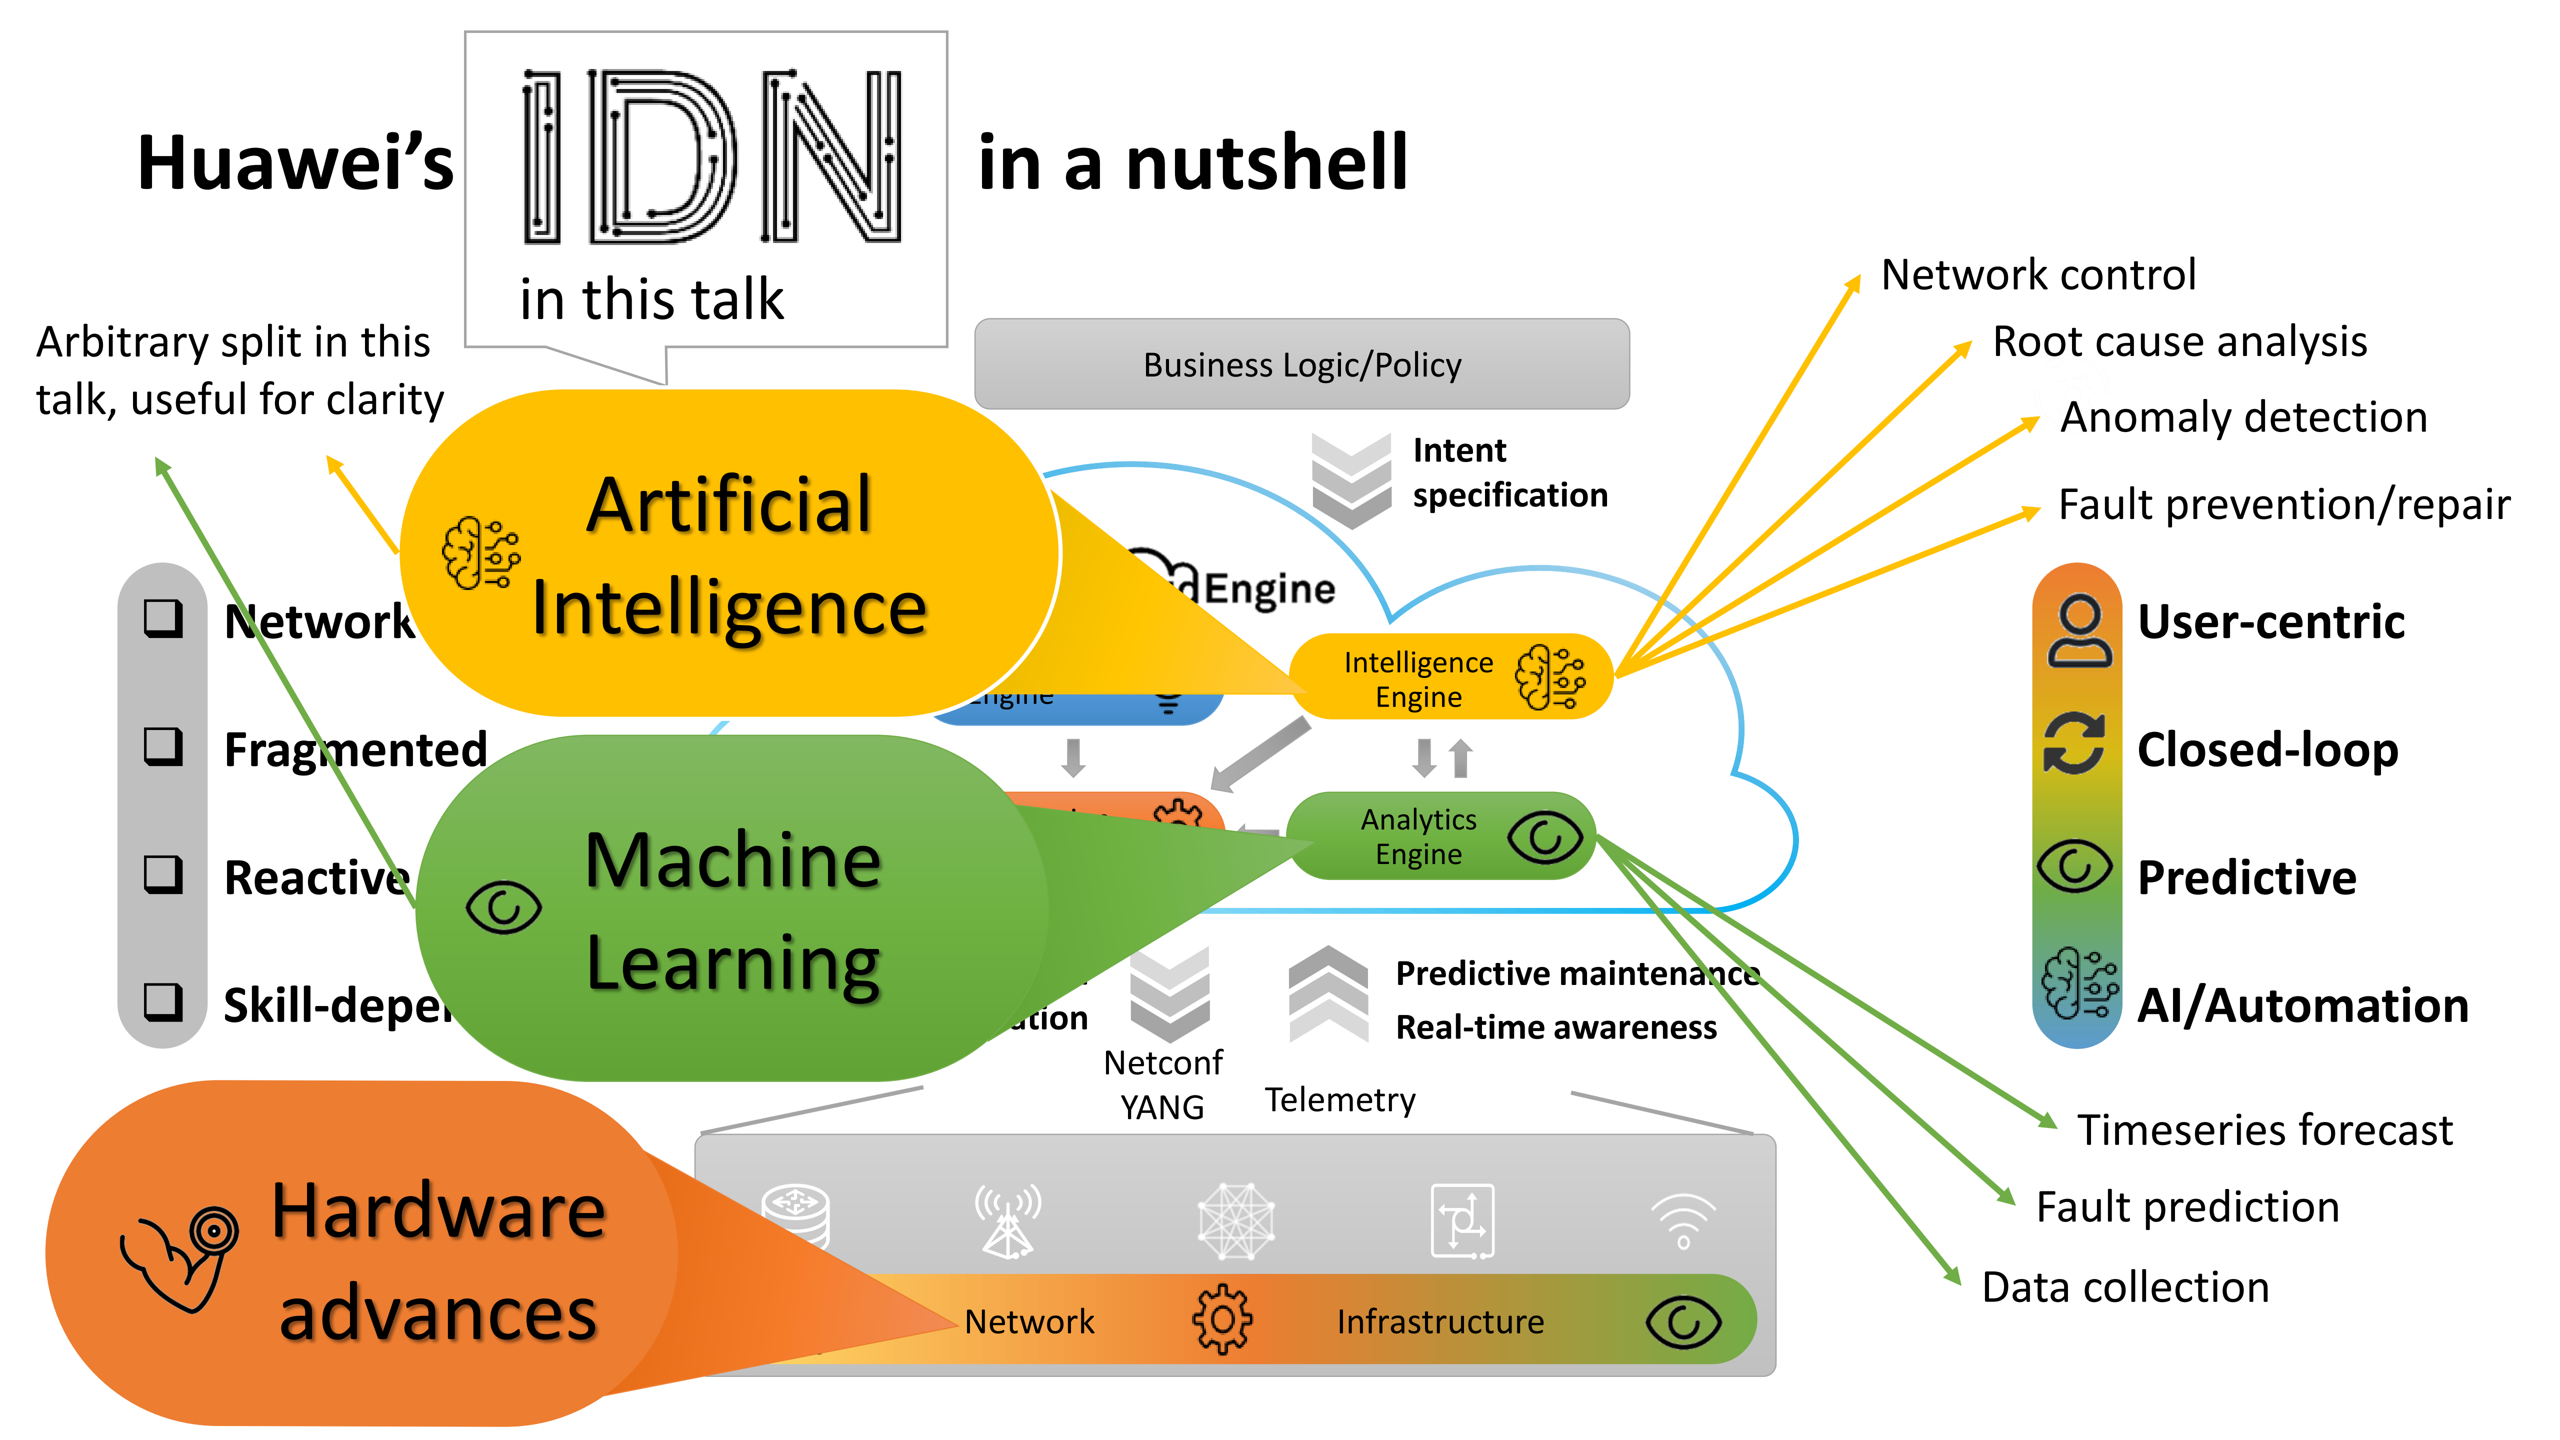

Huawei's IDN in a nutshell
IDN in this talk
Arbitrary split in this talk, useful for clarity
Business Logic/Policy
Intent specification
Artificial Intelligence
Machine Learning
Hardware advances
Network
Fragmented
Reactive
Skill-depe
Engine
Intelligence Engine
Analytics Engine
Network control
Root cause analysis
Anomaly detection
Fault prevention/repair
User-centric
Closed-loop
Predictive
AI/Automation
Predictive maintenance
Real-time awareness
Netconf YANG
Telemetry
Timeseries forecast
Fault prediction
Data collection
Network
Infrastructure

Agenda

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

Network data

So much data,

so few labels

Understand the network

- ❑ Explicability
- ❑ Evolution
- ❑ Security

Control the network

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Aim of this talk

Tips to avoid bumps in the road to network AI

+ Flash few examples out of our activities

Agenda

Hardware advances | Network data | Understand the network | Control the network

- ❑ History
- ❑ Trends
- ❑ AI chips

So much data,

so few labels

- ❑ Explicability
- ❑ Evolution
- ❑ Security

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Aim of this talk

Tips to avoid bumps in the road to network AI

+ Flash few examples out of our activities

Hardware advances

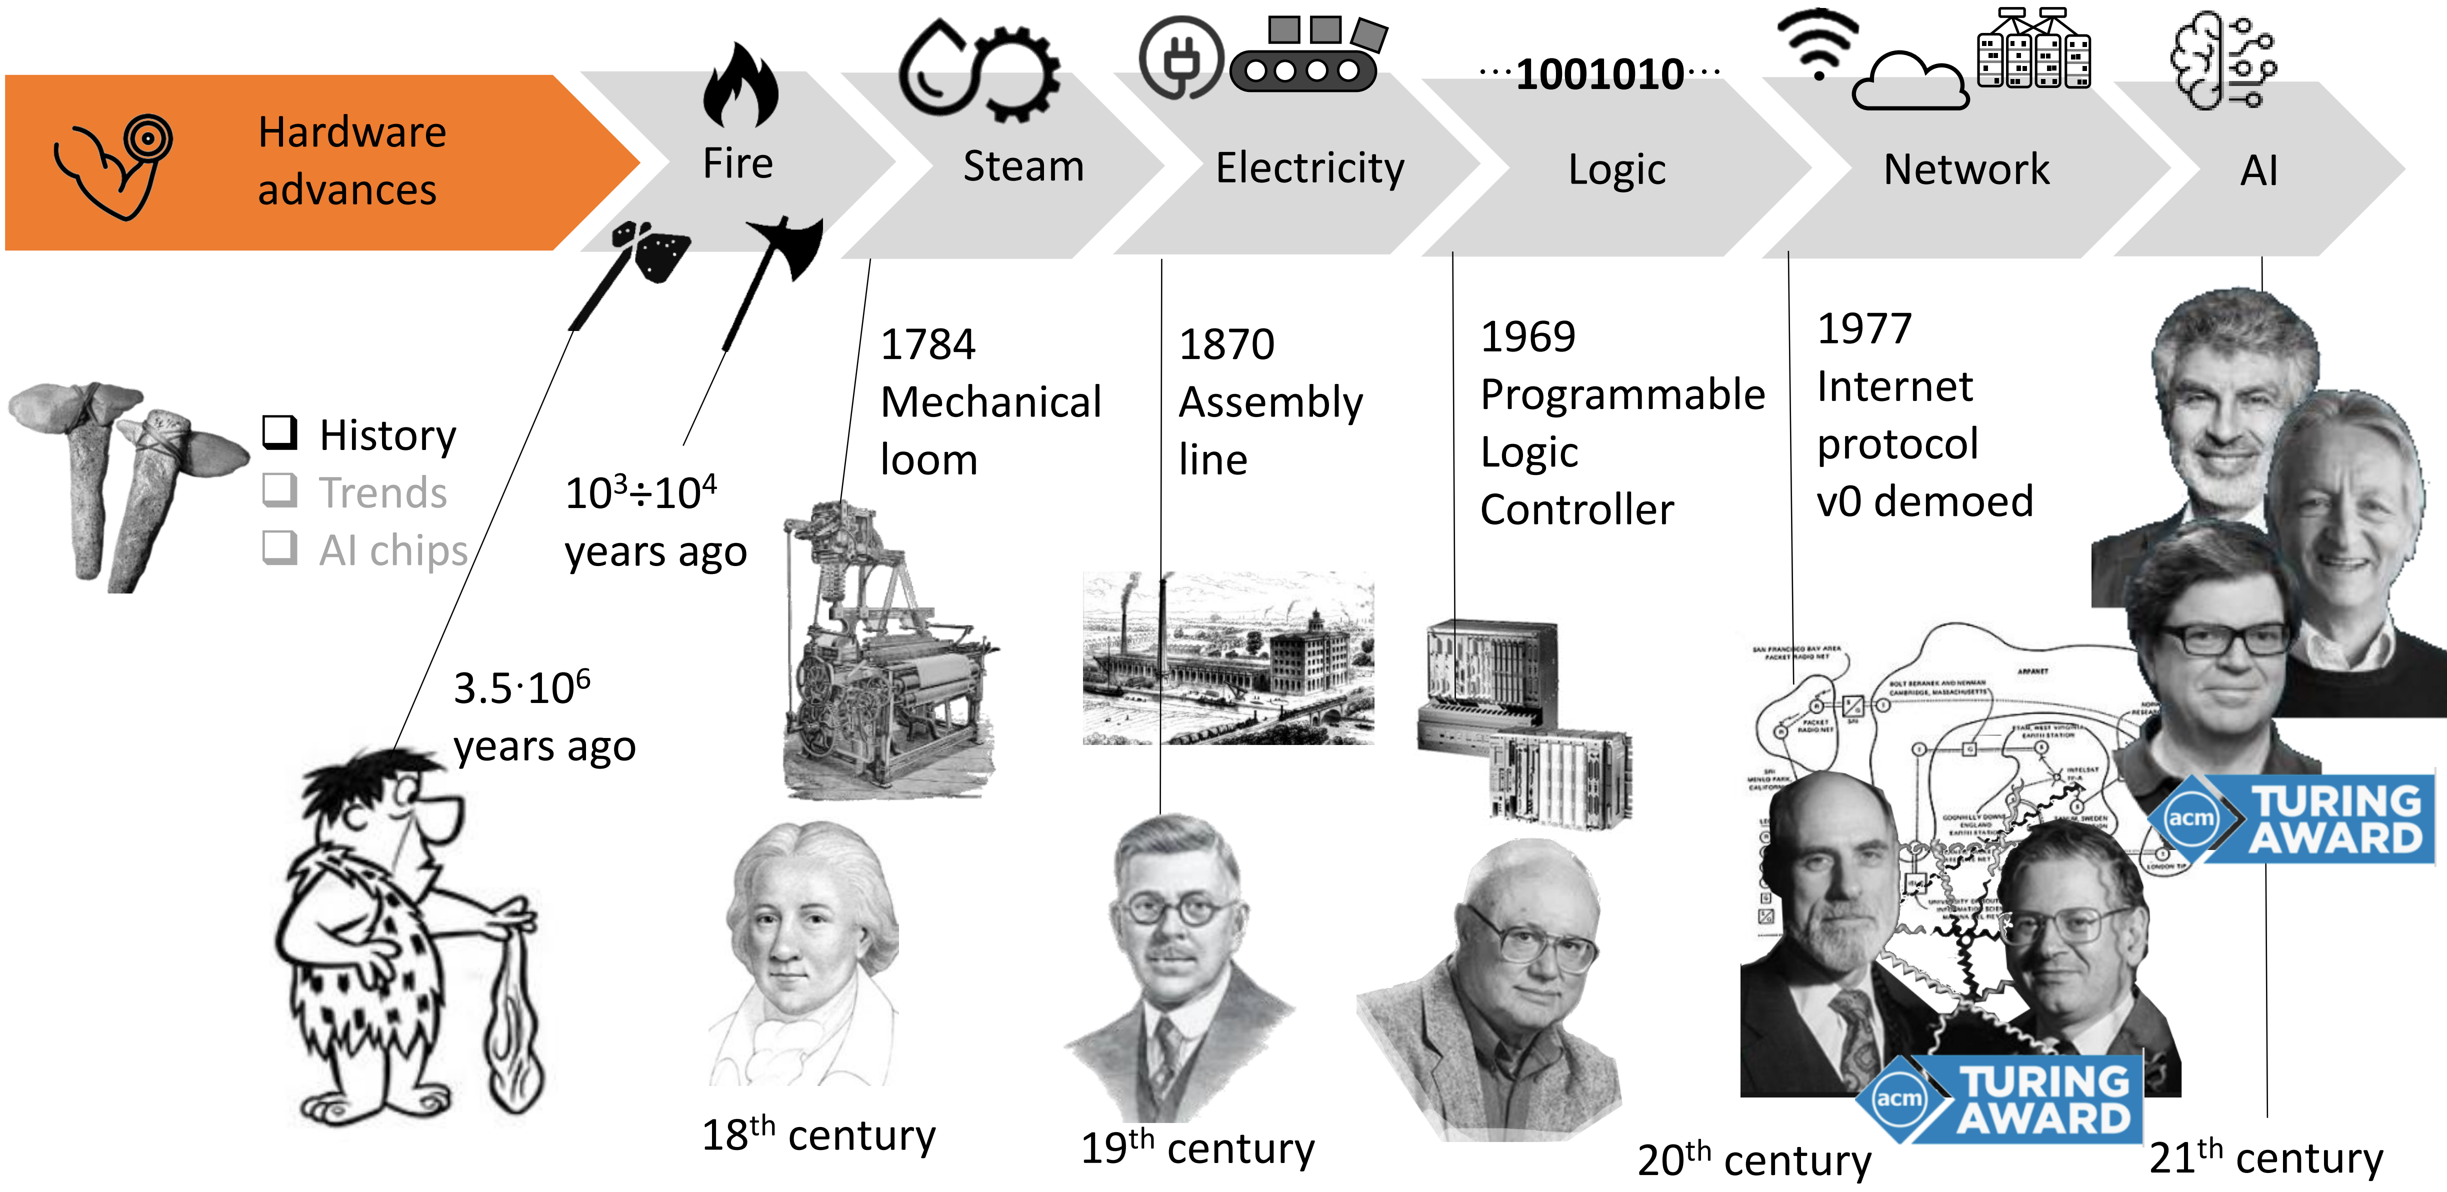

Hardware advances, but not only

Hardware advances → Fire → Steam → Electricity → ···1001010··· Logic → Network → AI

- History
- Trends
- AI chips

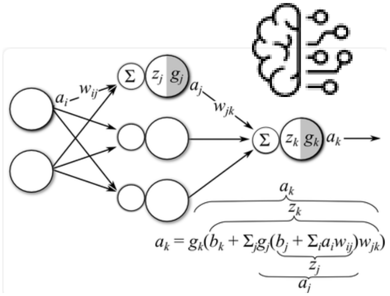

Theoretical advances

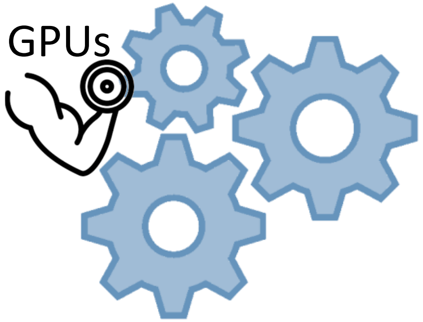

Massive amount of computational power

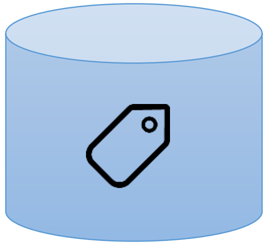

Massive volume of labeled data

Keys of success

21st century

Deep neural networks trend

“Natural” neural network
~20 W

Hardware advances

- History
- Trends
- AI chips

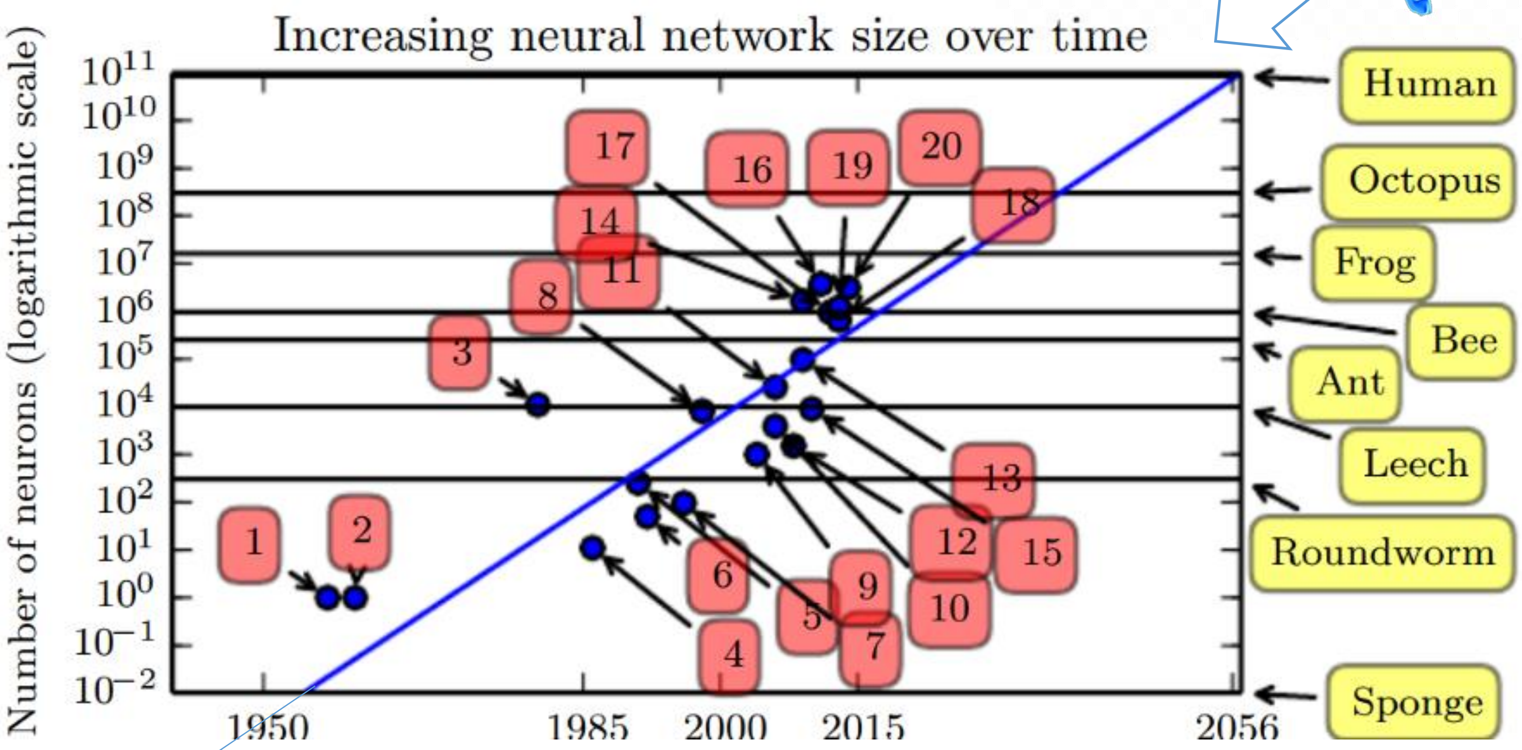

1.
Numbers of neurons increases faster than the number of transistors

Ian Goodfellow and Yoshua Bengio and Aaron Courville, Deep learning, MIT Press https://deeplearningbook.org

Hardware advances for general purpose computing

Hardware advances

- History
- Trends
- AI chips

2003, Moore law starts fading

2018, 15x gap

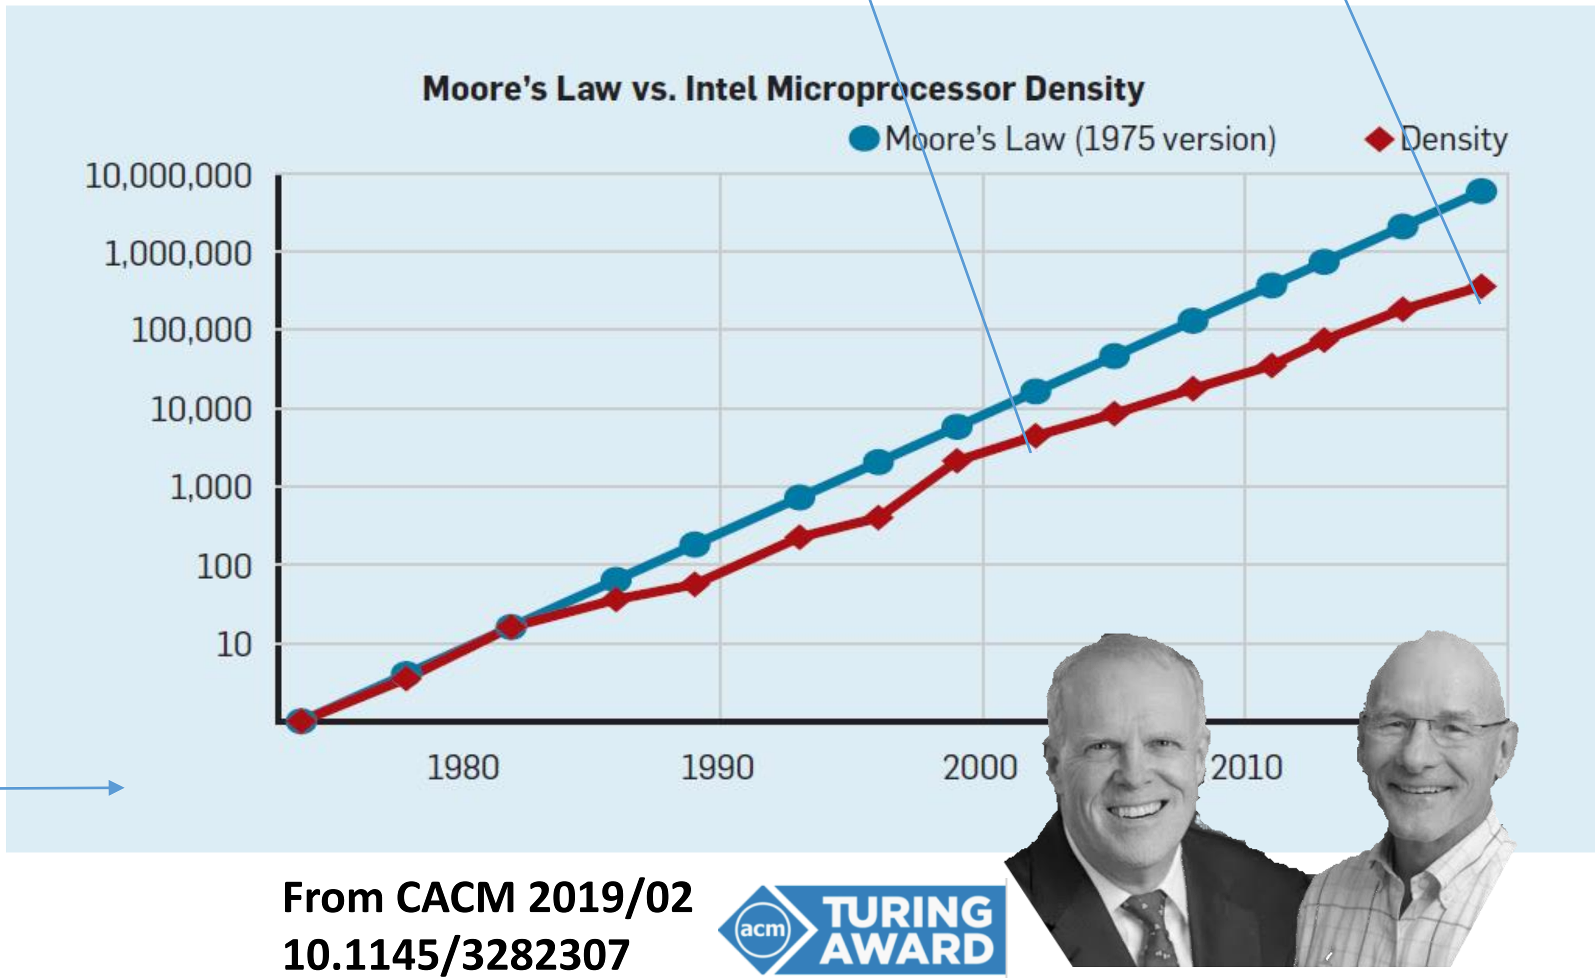

2.
Moore law will come to a stop eventually (the gap is already big)

From CACM 2019/02
10.1145/3282307

TURING AWARD

Hardware advances for general purpose computing

Hardware advances

- History
- Trends
- AI chips

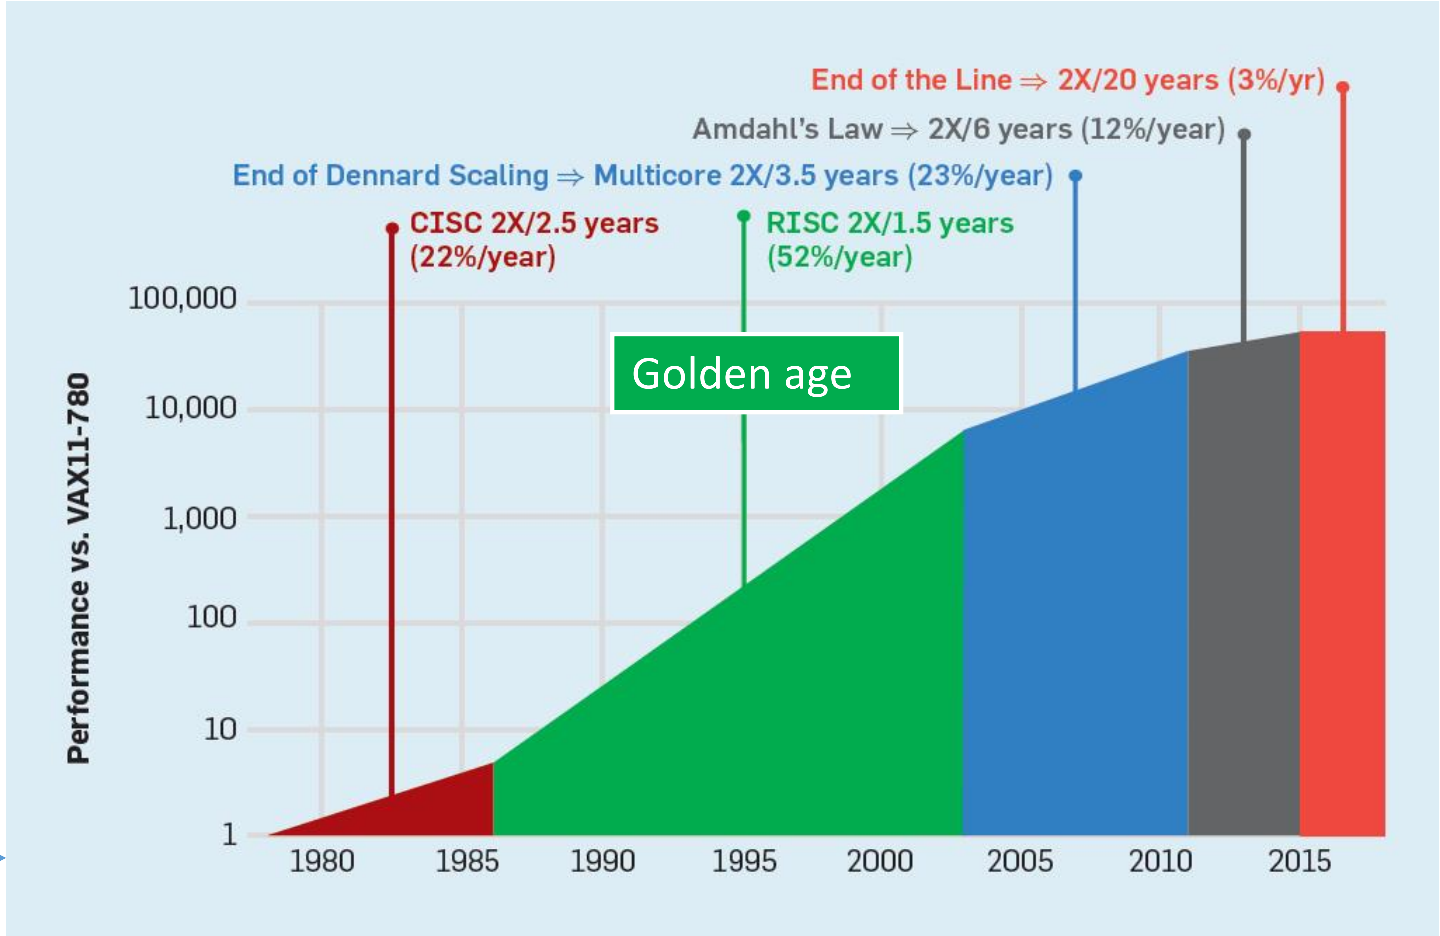

2b.
Computing performance increase is slowing down (it's not just Moore law…)

TURING AWARD

From CACM 2019/02
10.1145/3282307

Hardware advances for general purpose computing

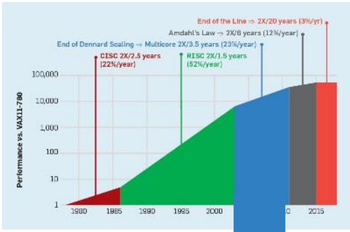

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

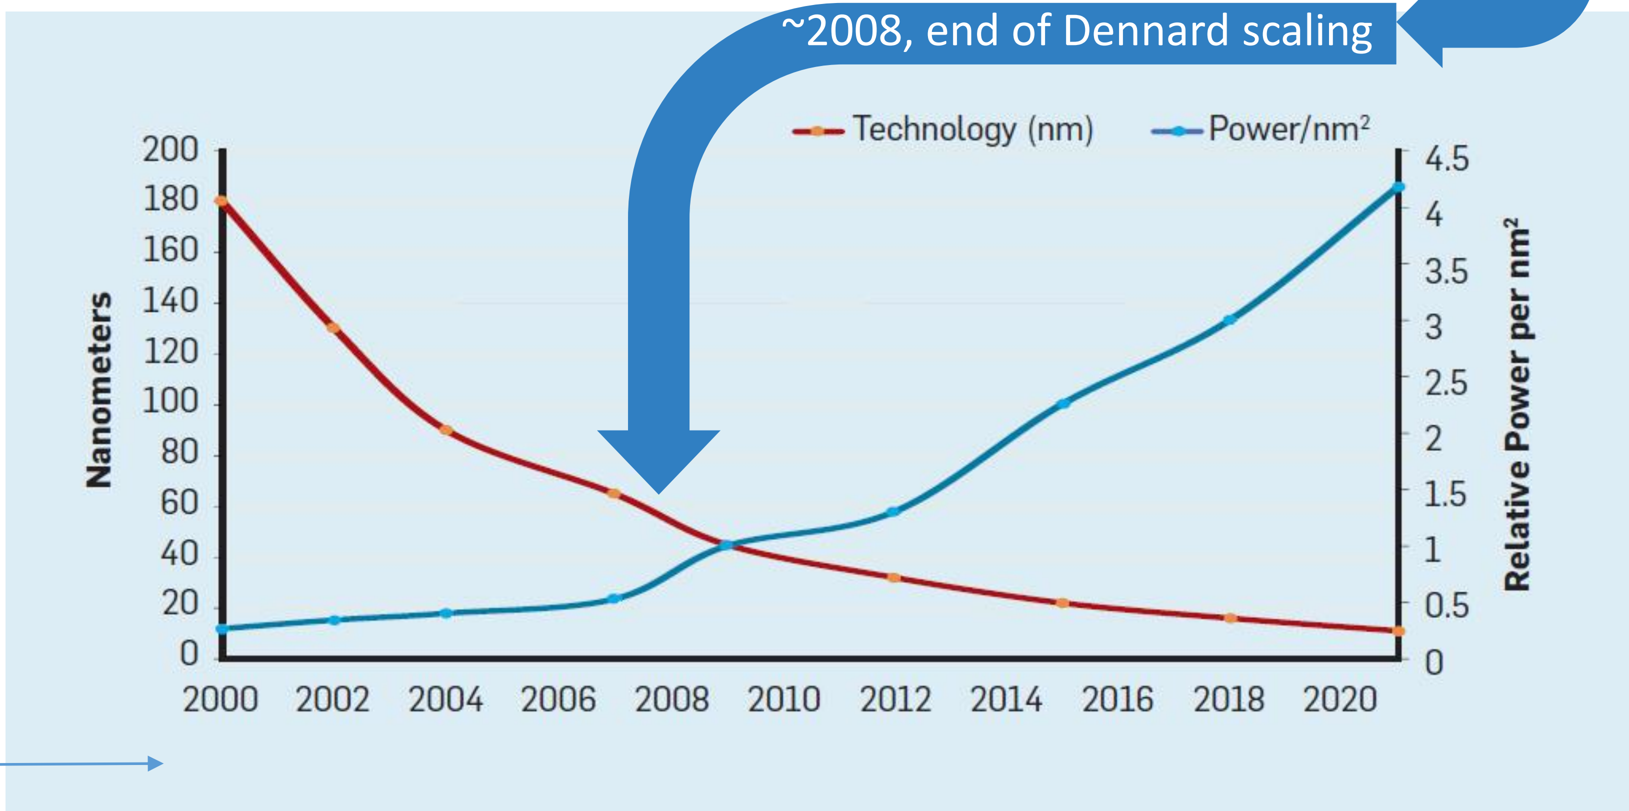

2c.
Dennard scaling also practically stopped,
(⇨ multicore)

From CACM 2019/02
10.1145/3282307

Hardware advances for general purpose computing

Hardware advances

- [ ] History
- [ ] Trends
- [ ] AI chips

$$\lim_{s\to\infty} S_{\text{latency}}(s) = \frac{1}{1-p}.$$

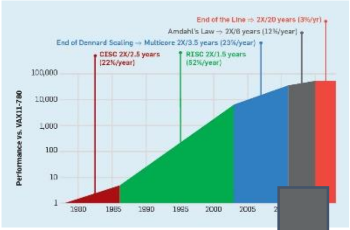

~2010, Ahmdal law

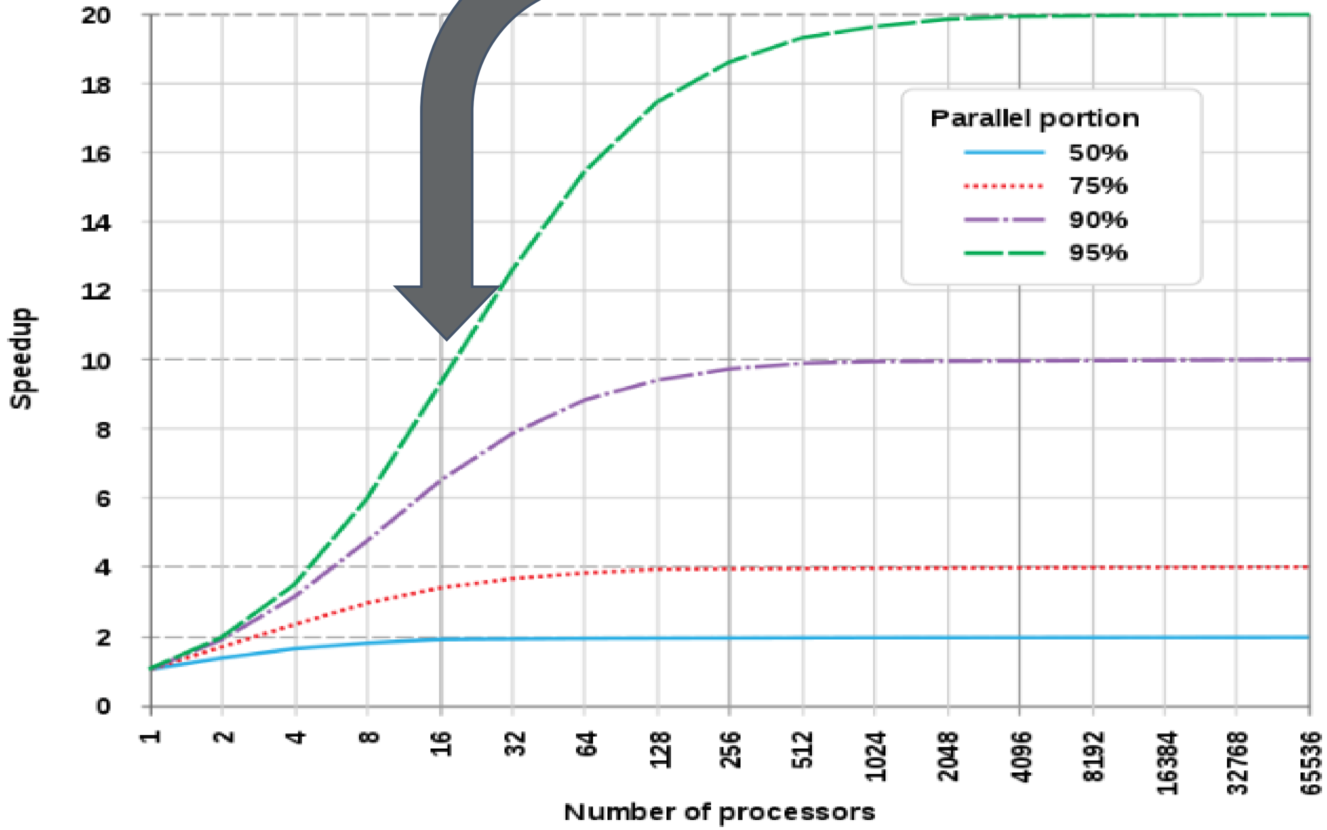

2d.
However Ahmdal's law limit the practical appeal for multicore CPUs in many cases

Hardware power consumption ?

Hardware advances

- ❑ History
- ❑ **Trends**
- ❑ AI chips

3.
General purpose designs hitting a power wall

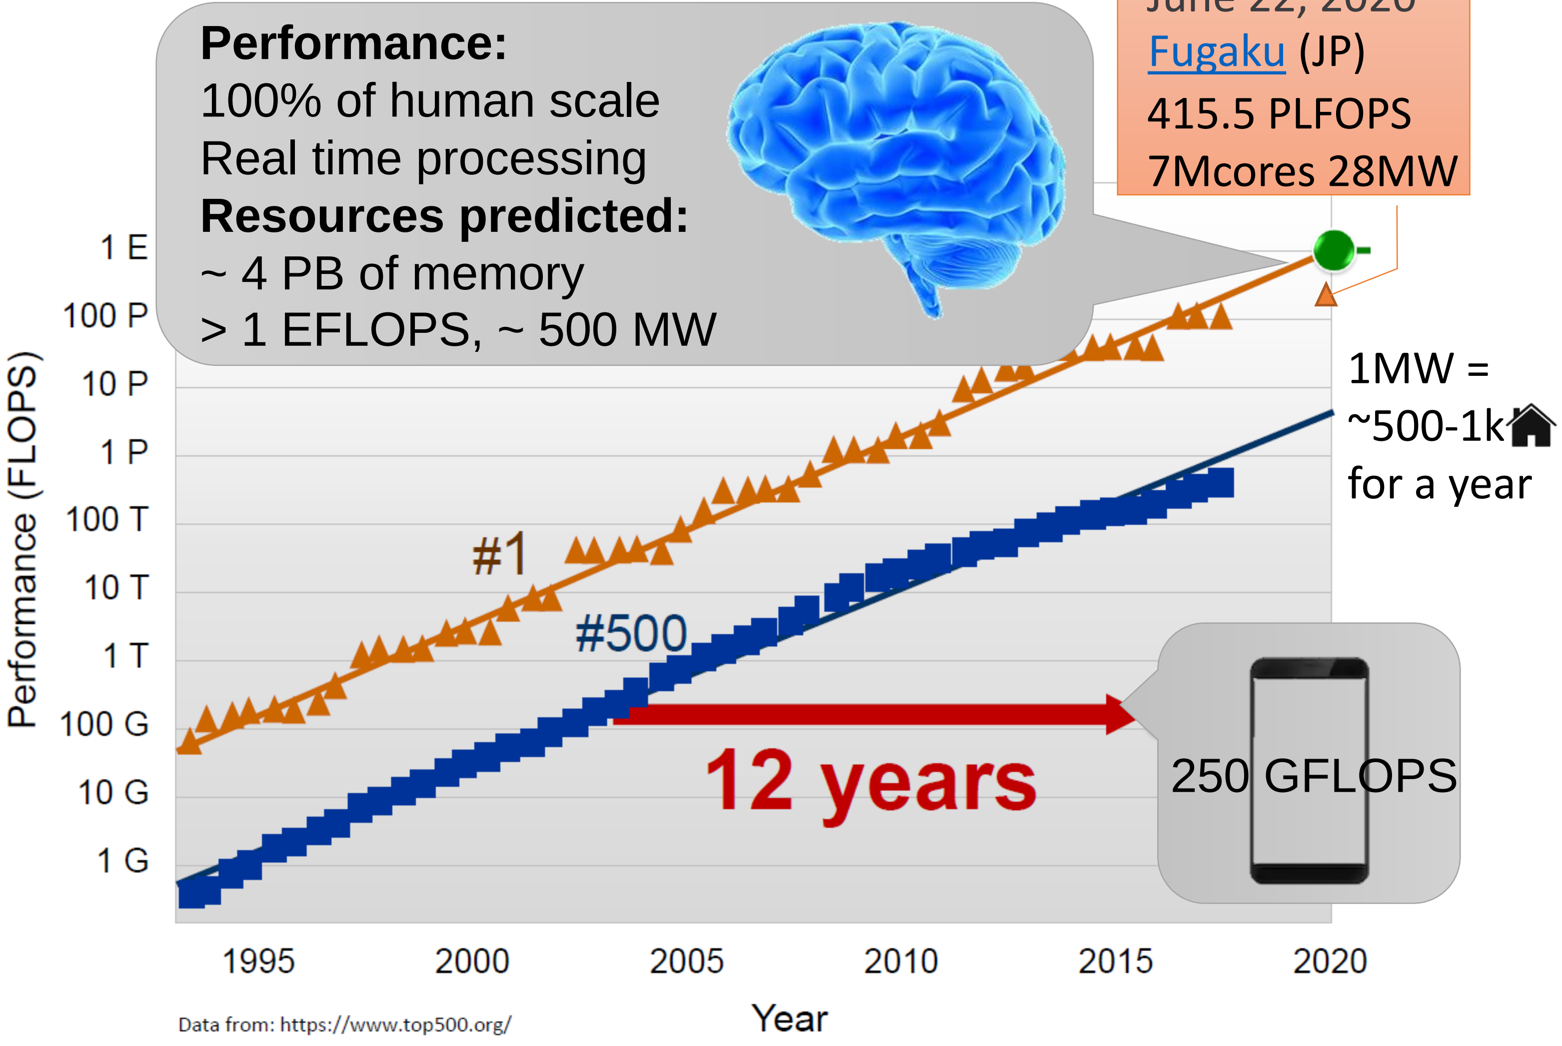

Courtesy H.S. Philip Wong (黃漢森), Stanford & TSMC

Hardware bottleneck for AI processing ?

Hardware advances

- History
- Trends
- AI chips

4.
General purpose designs hit a memory wall for AI

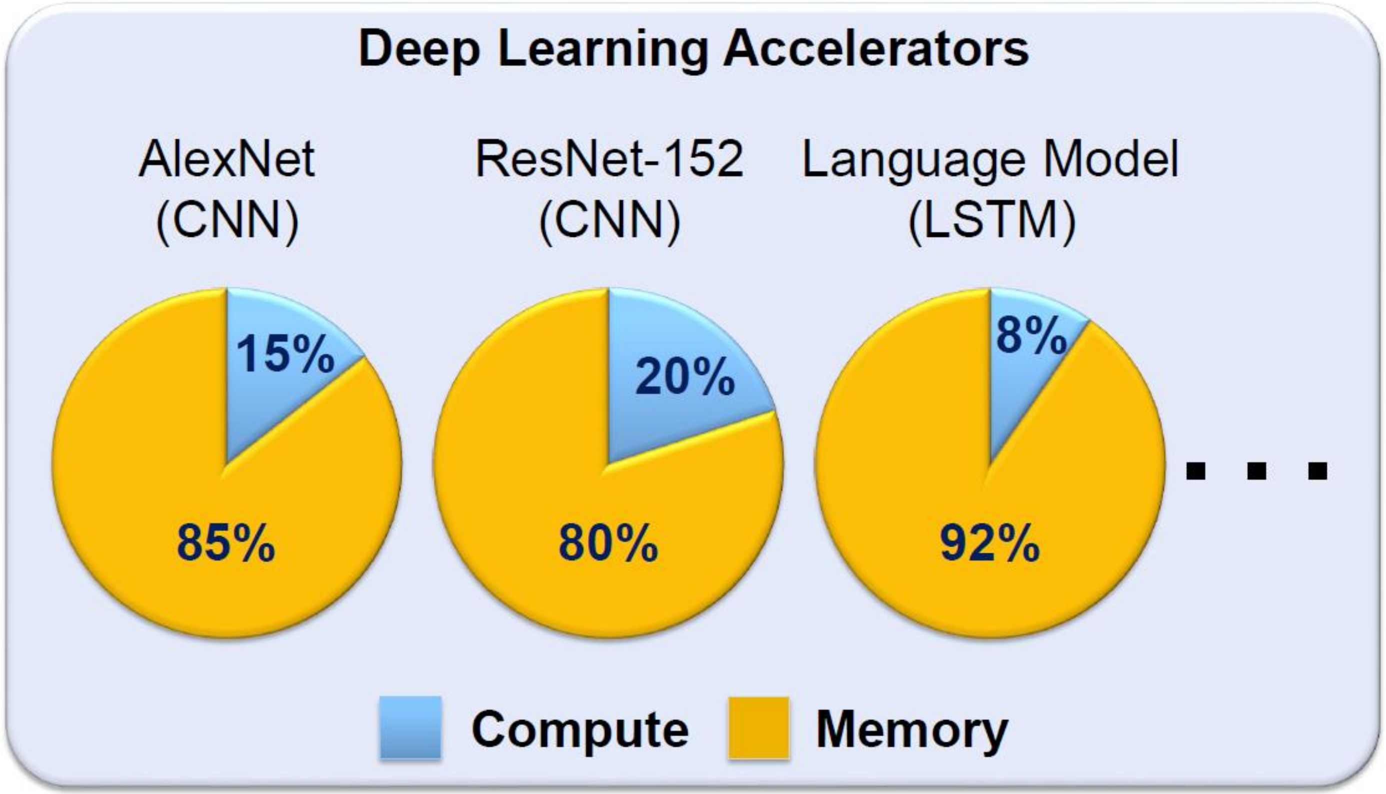

Intel performance counter monitors 2 CPUs, 8-cores/CPU + 128GB DRAM

Source: S. Mitra (Stanford)

H.S. Philip Wong (黃漢森),
Stanford & TSMC

Hardware design trends

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

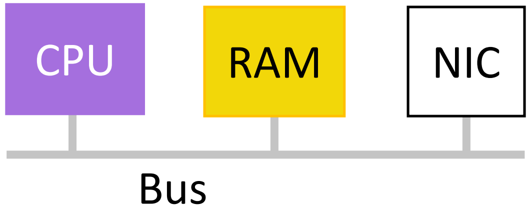

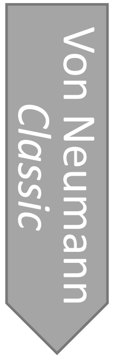

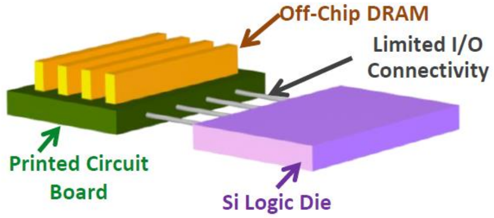

⇨

Go beyond classic
Von Neumann architectures

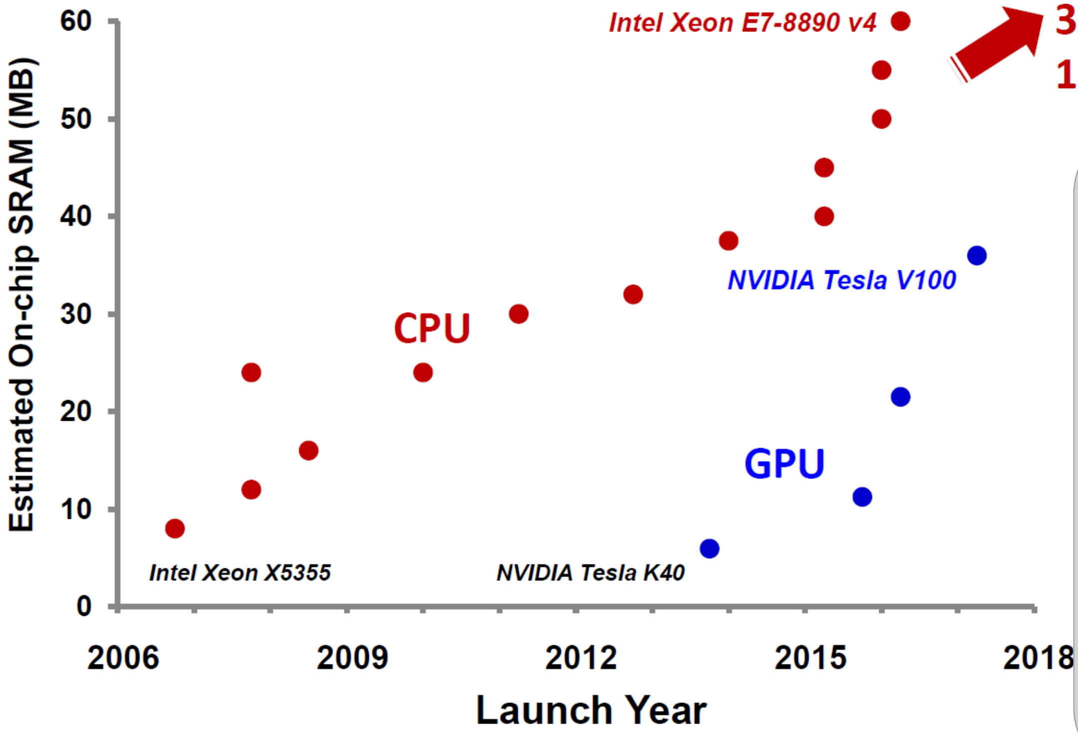

3.8 GBytes @ 1.4 nm node

Source: W. Hwang, Prof. S. Mitra (Stanford)

Recall
~ 4 PB memory

10^{11} neurons each connected to 10^4 synapses

H.S. Philip Wong (黃漢森),
Stanford & TSMC

Hardware design trends

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

⇨

Go beyond classic
Von Neumann architectures
(⇨ memory-compute integration)

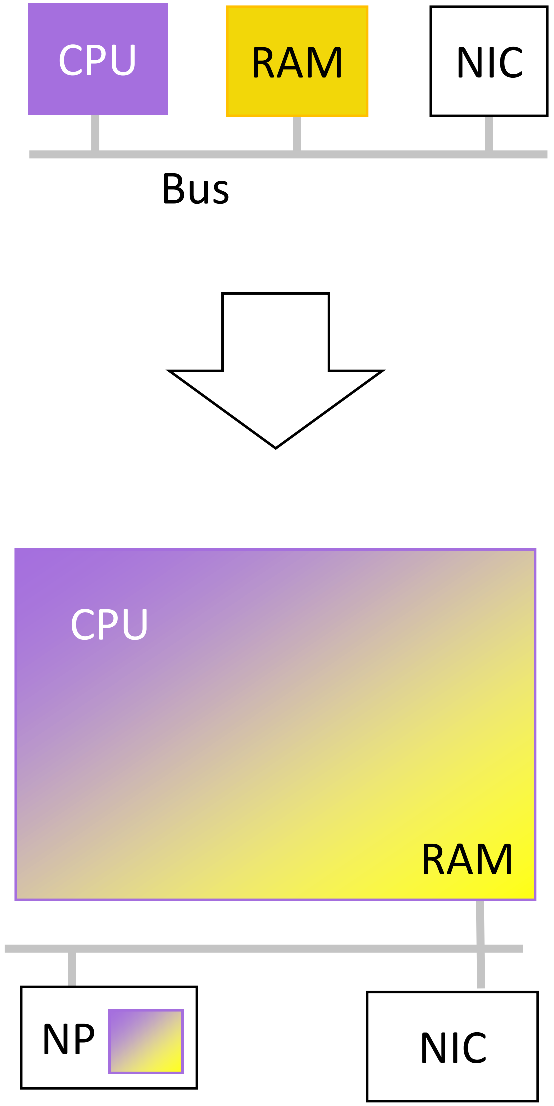

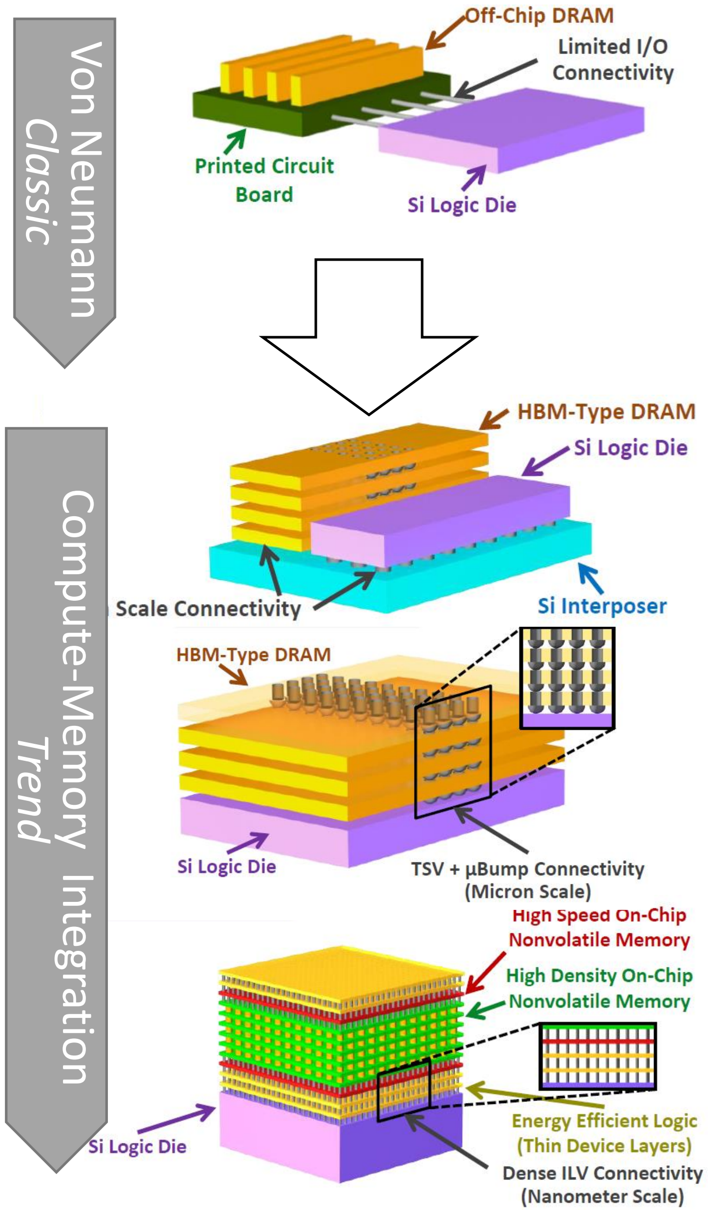

Hardware design trends

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

⇨

Go beyond classic
Von Neumann architectures
(⇨ design tailored for CNNs)

Huawei Ascend

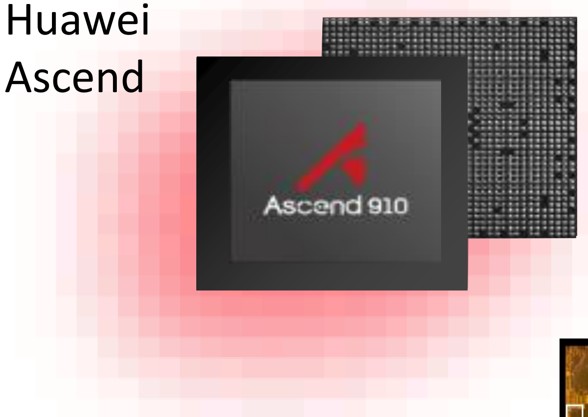

Coral.ai

Tesla FSD

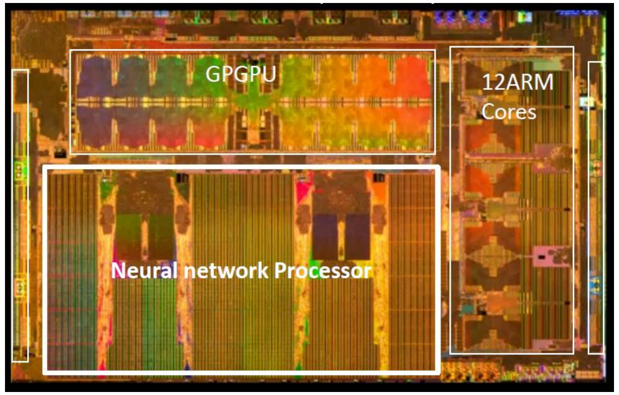

Google TPU v3.0

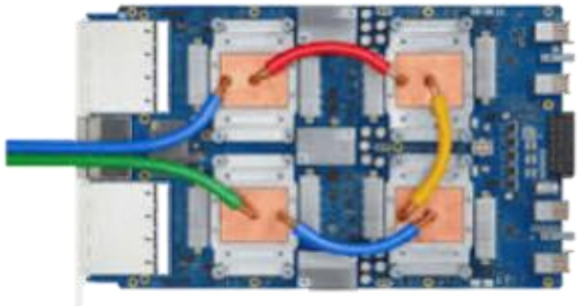

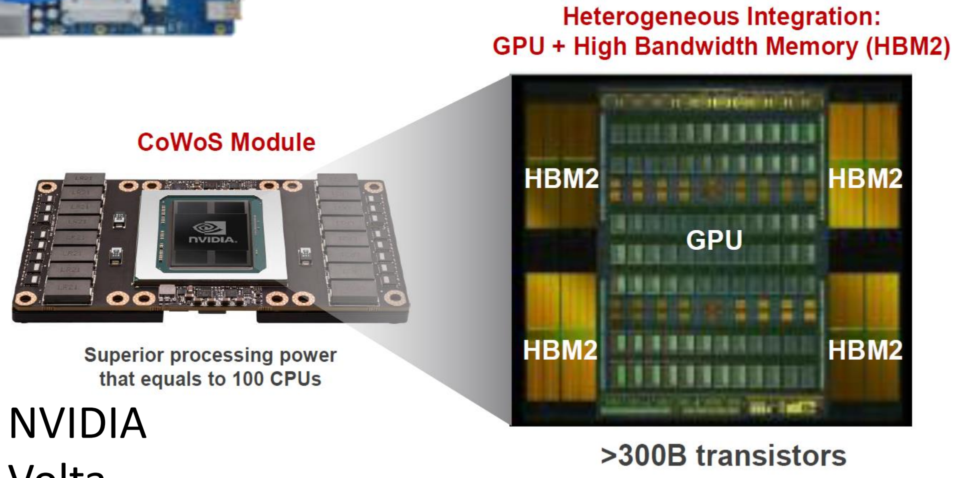

NVIDIA Volta

Hardware design trends

Hardware advances

- History
- Trends
- AI chips

➩

Go beyond classic
Von Neumann architectures
(➩ flexible design, edge intelligence)

Huawei Ascend

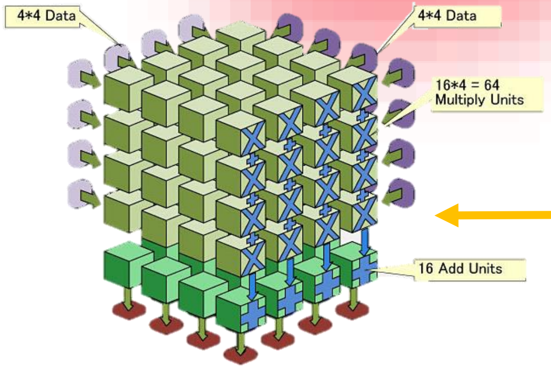

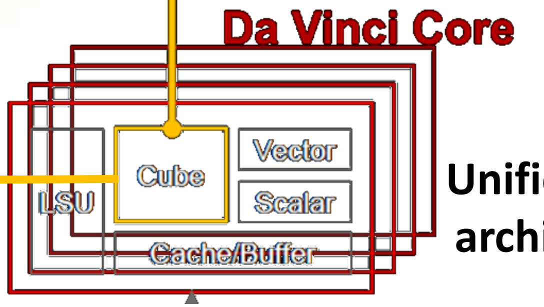

***DaVinci* Unified chip architecture**

Ascend310 (Mini)
FP16：8 TFLOPS
INT8：16 TOPS

Power: 8W
Process: 12nm

Ascend910 (Max)
FP16: 256 TFLOPS
INT8: 512 TOPS

Power: 310W
Process: 7+ nm

DaVinci chips → DaVinci Server → DaVinci Cluster

Hardware design trends

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

⇨

Go beyond classic
Von Neumann architectures
(⇨ flexible design, cloud)

Huawei
Ascend

8x Ascend 910

Ascend board
DaVinci node

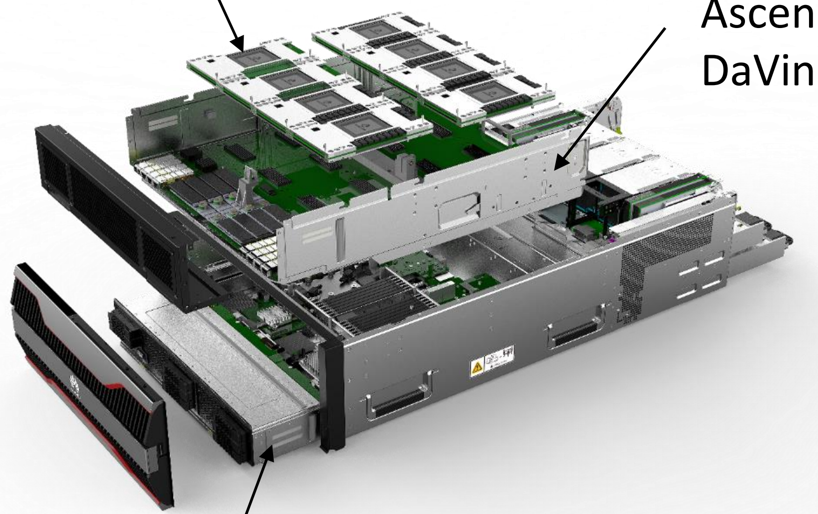

X86 Node

DaVinci chips → DaVinci Server → DaVinci Cluster

Hardware design trends

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

➪

Go beyond classic
Von Neumann architectures
(➪ flexible design, hyperscale)

Huawei Ascend

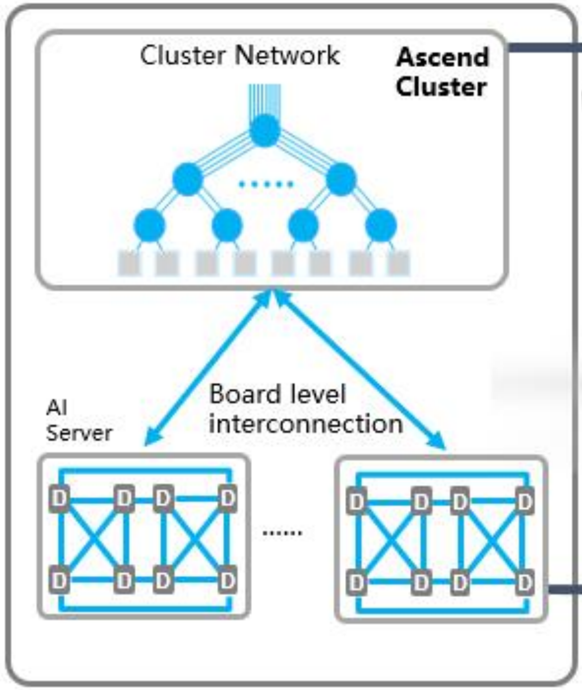

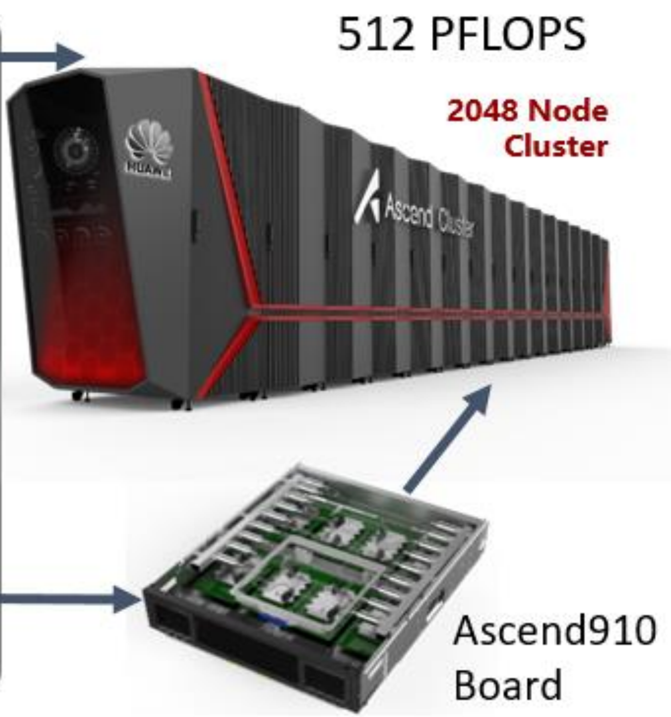

Ascend910 Board

DaVinci chips ➤ DaVinci Server ➤ DaVinci Cluster

Hardware desing trends

Hardware advances

- History
- Trends
- AI chips

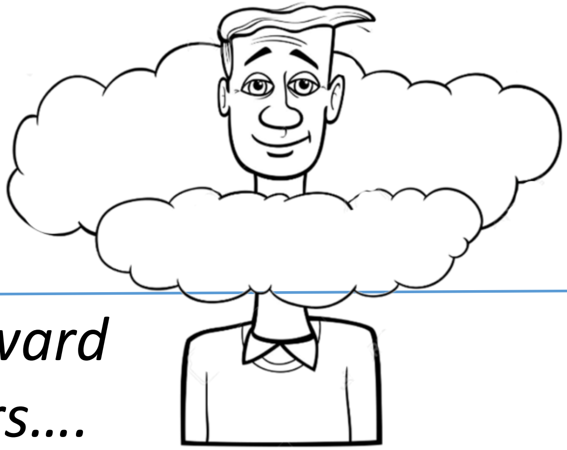

Fast forward
~10 years….

“Artificial neural networks” (synchronous)

310W
~16GB memory
FP16: 256 TFLOPS
(INT8: 512 TOPS)

“Natural neural networks” (asynchronous)

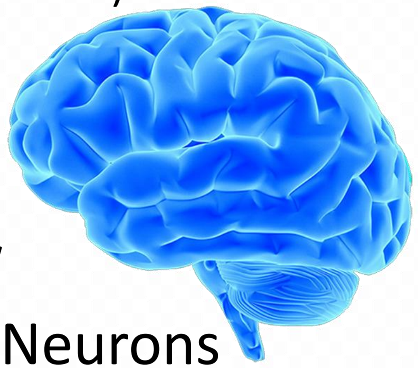

20W
10^{11} Neurons
~ 4 PB of memory
> 1 EFLOPS

Spiking neural networks & neuromorphic chips (asynchronous)

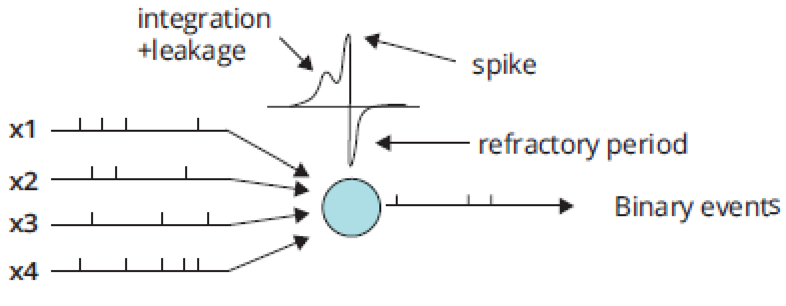

Figure 8-1 A simple Demonstration of Leaky Integrate-and-Fire Algorithm

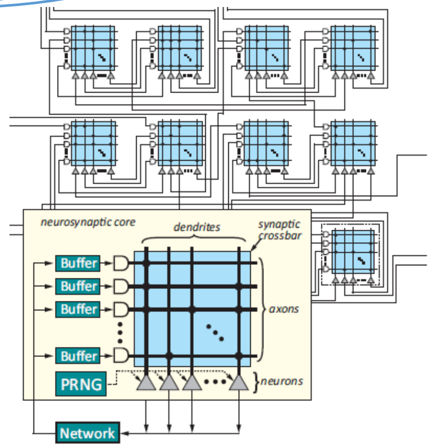

Figure 8-3 Architecture of Neuromorphic Chip

Tsinghua “AI Chips” whitepaper (2018)

Hardware is key, but software needed to exploit it!

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

A bit extreme example, but valid point!

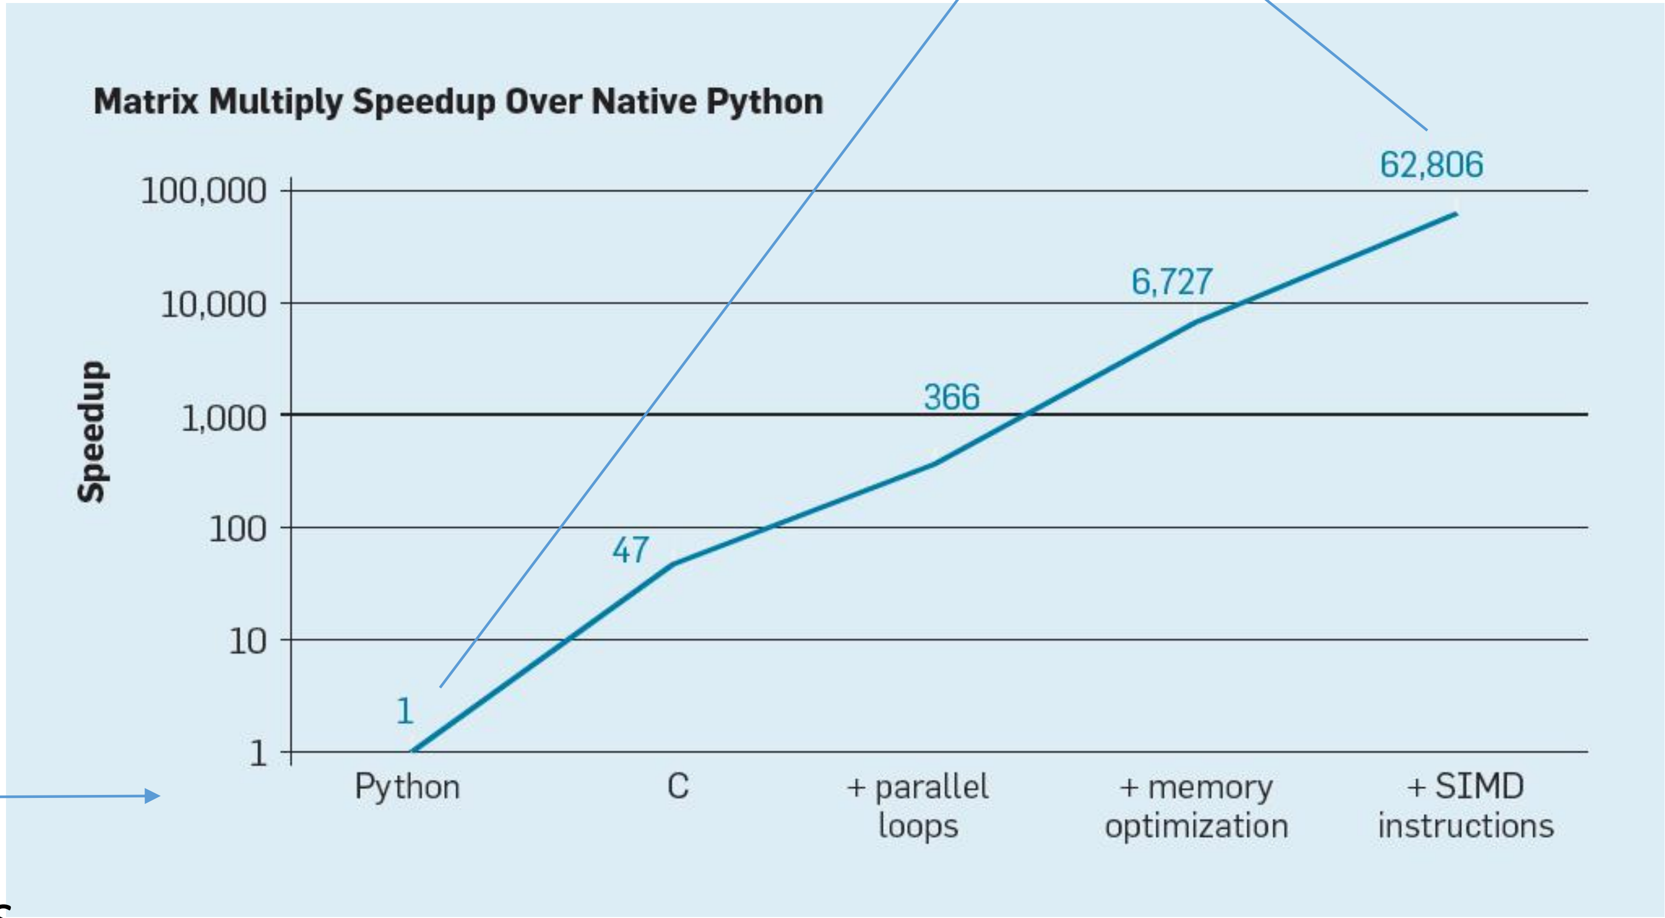

⇨

Go beyond classic
Von Neumann architectures
(⇨ software still matters)

Ex. from Leiserson. C, "There plenty of room at the top"
Illustration from CACM 2019/02 10.1145/3282307

Hardware is key, but software needed to exploit it!

Hardware advances

- ❑ History
- ❑ Trends
- ❑ AI chips

Don't expect the L3 cross-compiler to just do *all* the magic

The more you know, the better *your* program

Ascend software stack

No free lunch…

Level 3 Library (written by novice programmer)		TBE LIB			
Level 3 Compiler (mathematical programming model)	TVM/XLA	TBE	Halide	Halide	
Level 2 Library (written by skilled programmer)	CudaNN/ CuBLAS	TIK LIB			
Level 2 Compiler (parallel/kernel programming model)	Cuda/OpenCL	TIK			
Level 1 Library (written by expert)		CCE Lib			Eigen/numpy
Low Level 1 Compiler (Intrinsic C) (Architecture defined programming)		CCE C	Intrinsic C	Intrinsic C	C/C++ LLVM
Instruction Set Architecture	GPU	NPU	Vision Processor	DSP	CPU + SVE

Software

Hardware

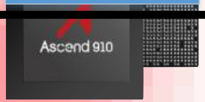

The long and winding road

To self-driving networks

Network

AI

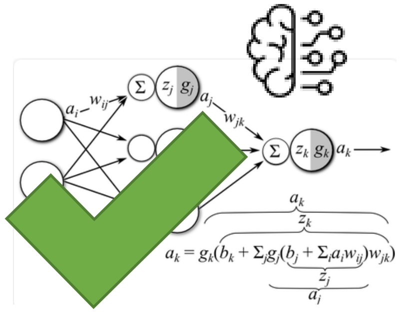

Theoretical advances

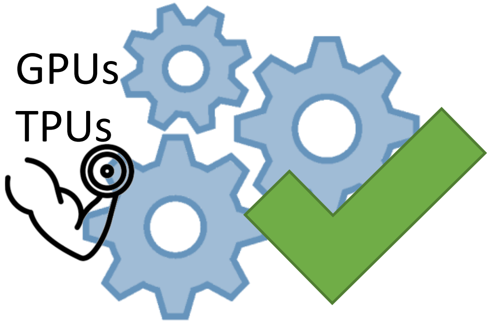

Massive amount of computational power

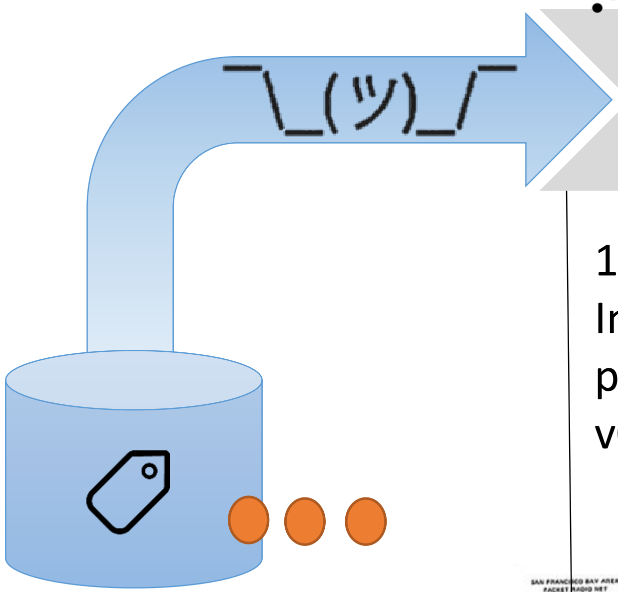

Massive volume of labeled data

Keys of success

1977
Internet
protocol
v0 demoed

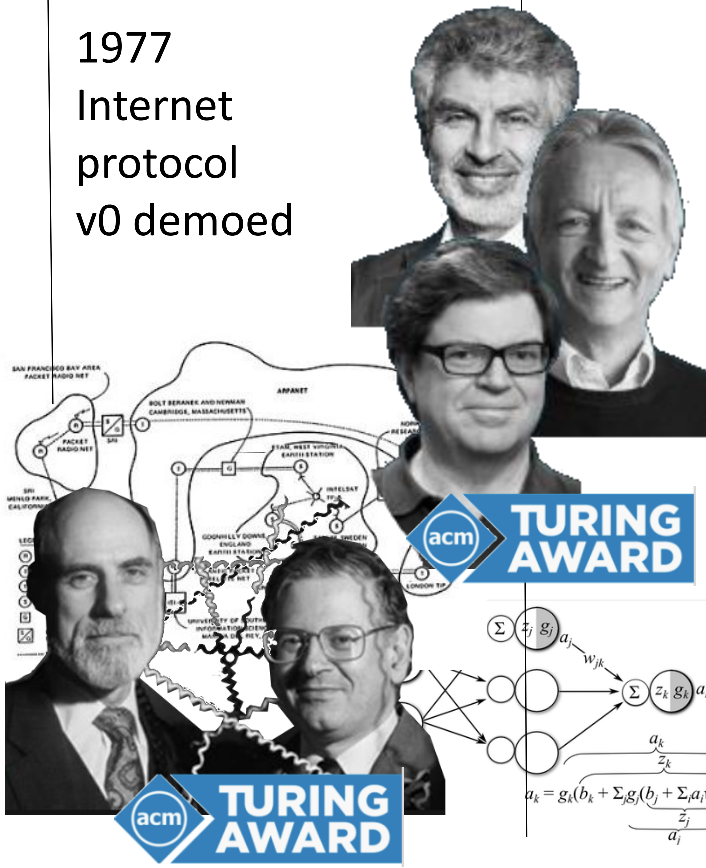

The long and winding road

Data preparation & management essential for AI in products

To self-driving networks

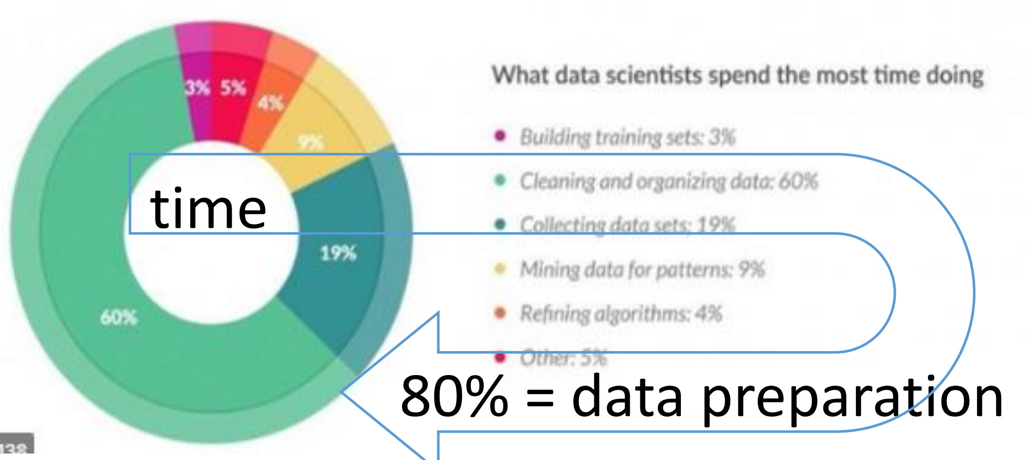

time

80% = data preparation

"Data is a key asset for AI system"

Andrew Ng (co-founder of Google Brain and former Vice President and Chief Scientist at Baidu)

"Amount of time on Algorithm / Data : PHD = 90% / 10% Tesla = 20% / 80% "

Andrej Karpathy (director of Artificial Intelligence & Autopilot Vision at Tesla)

Data management

Agenda

Hardware advances

Network data

Understand the network

Control the network

- ❑ History
- ❑ Trends
- ❑ AI chips

So much data,

so few labels

- ❑ Explicability
- ❑ Evolution
- ❑ Security

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Aim of this talk

Tips to avoid bumps in the road to network AI

+ Flash few examples out of our activities

Networking data for ML / AI

Network
data

So much data,

so few labels

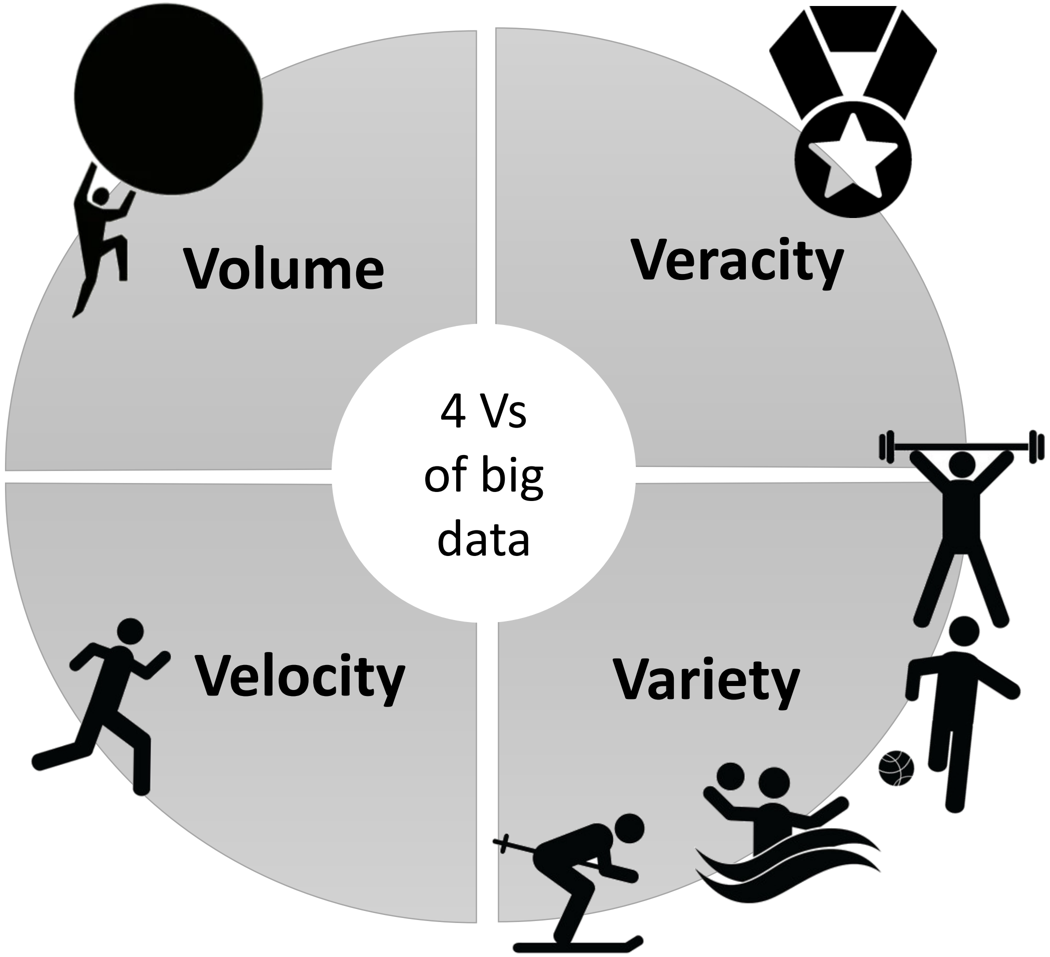

Networking data for ML / AI

Network data

So much data,

so few labels

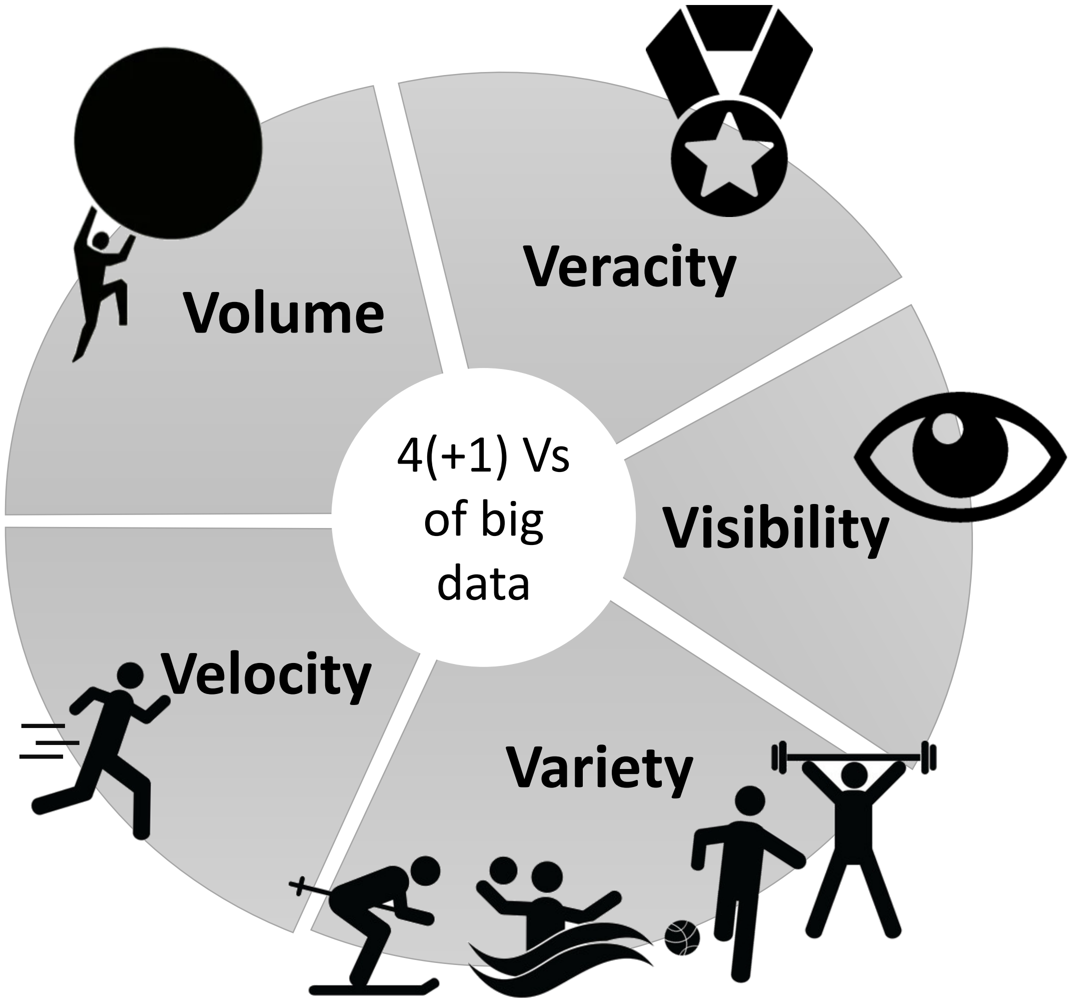

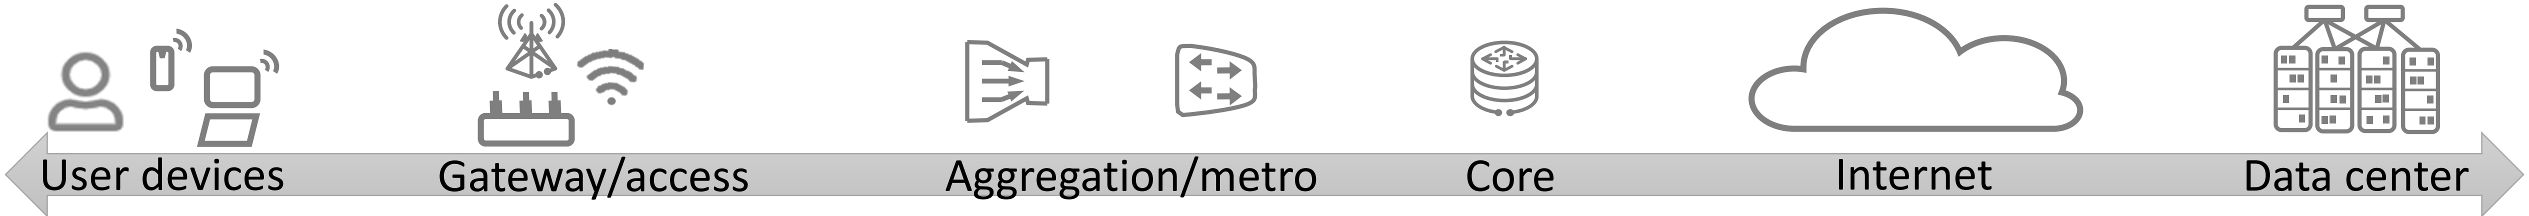

Networking data for ML / AI

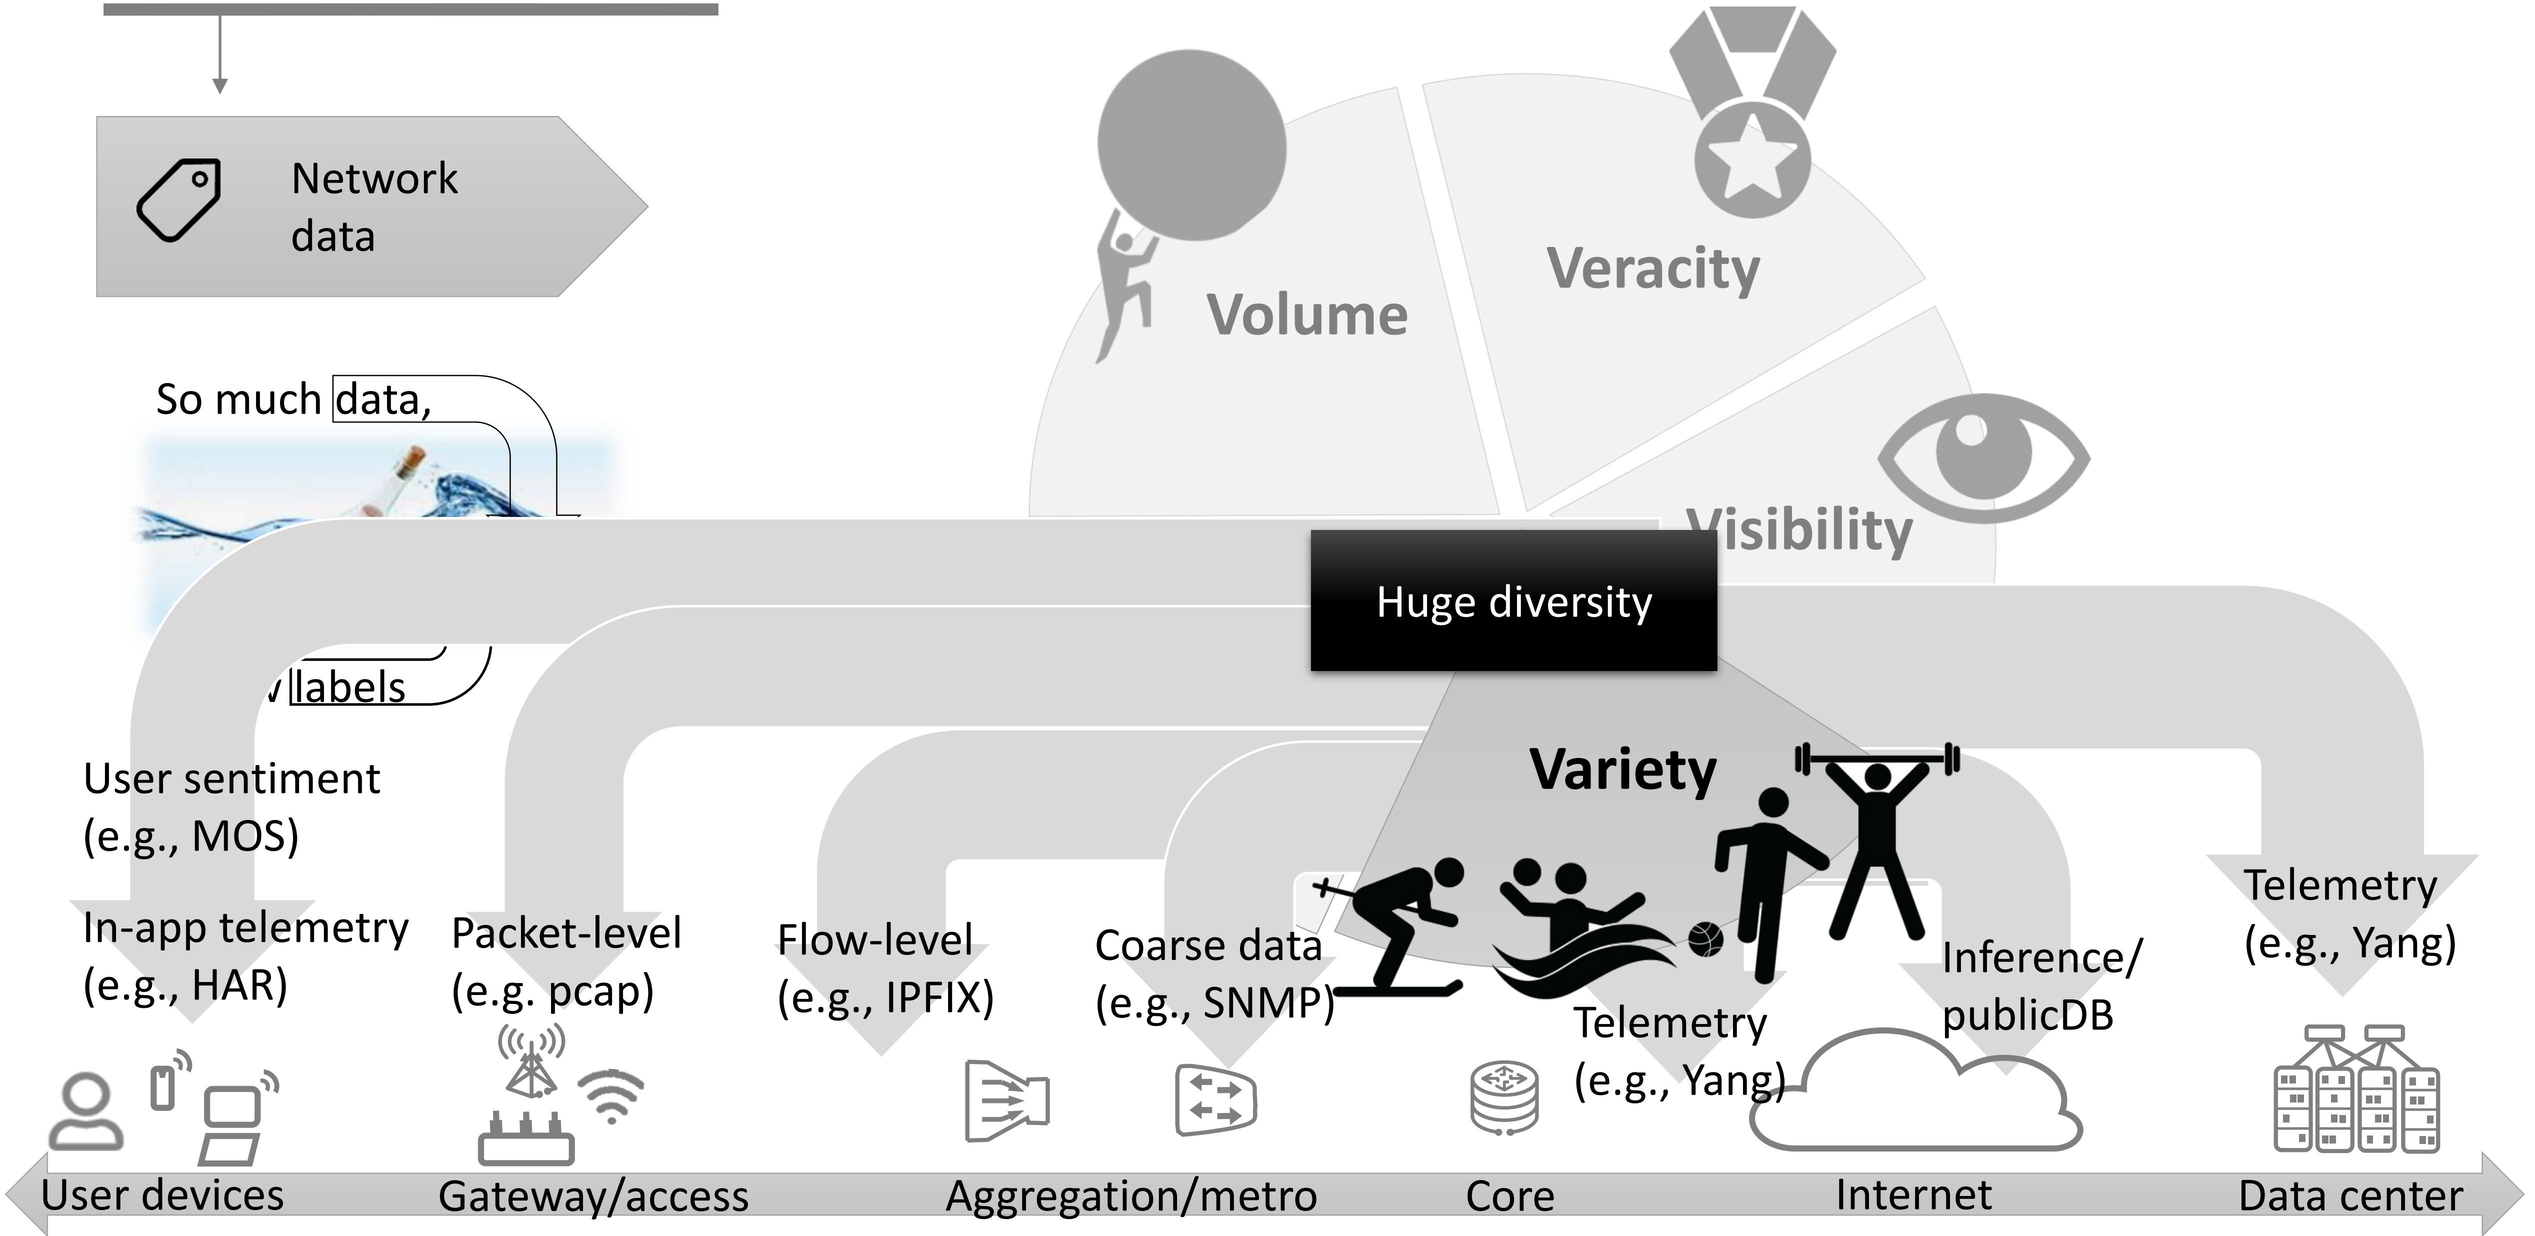

https://www6.slac.stanford.edu/research/scientific-computing
https://home.cern/resources/faqs/facts-and-figures-about-lhc

Networking data for ML / AI

Network data

So much data,

so few labels

Linac Coherent Light Source
10^{10} bps, 4 PB

Large Synoptic Survey Telescope
10^{8} pixel/sec

Large Hadron Collider (LHC) 10^{12} bps

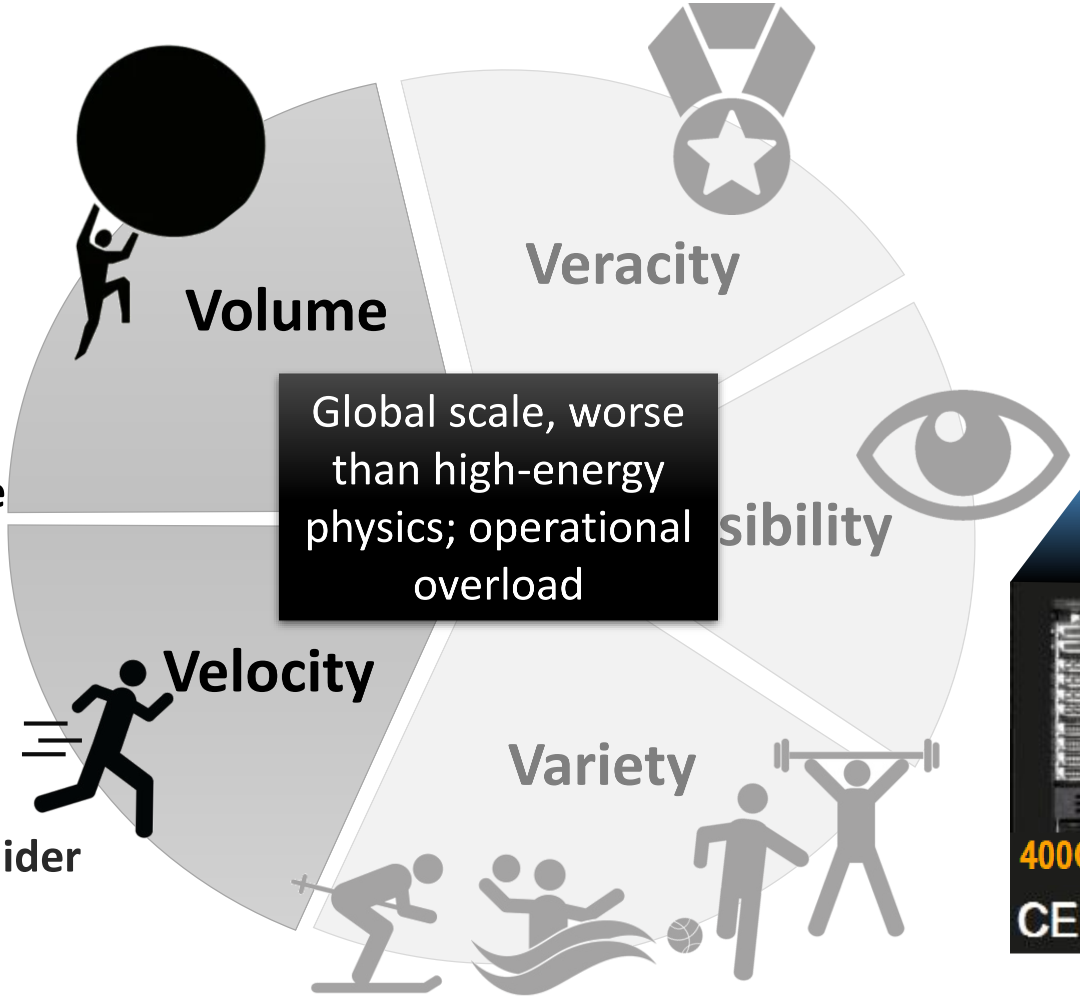

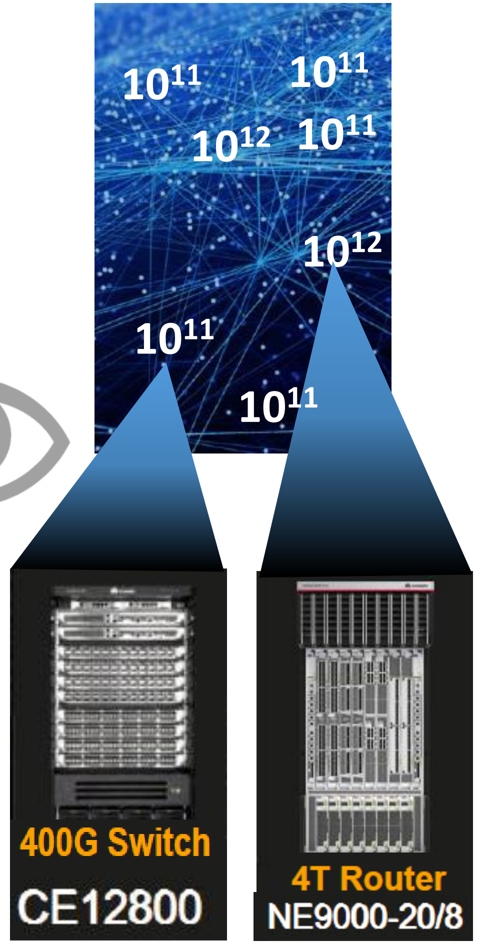

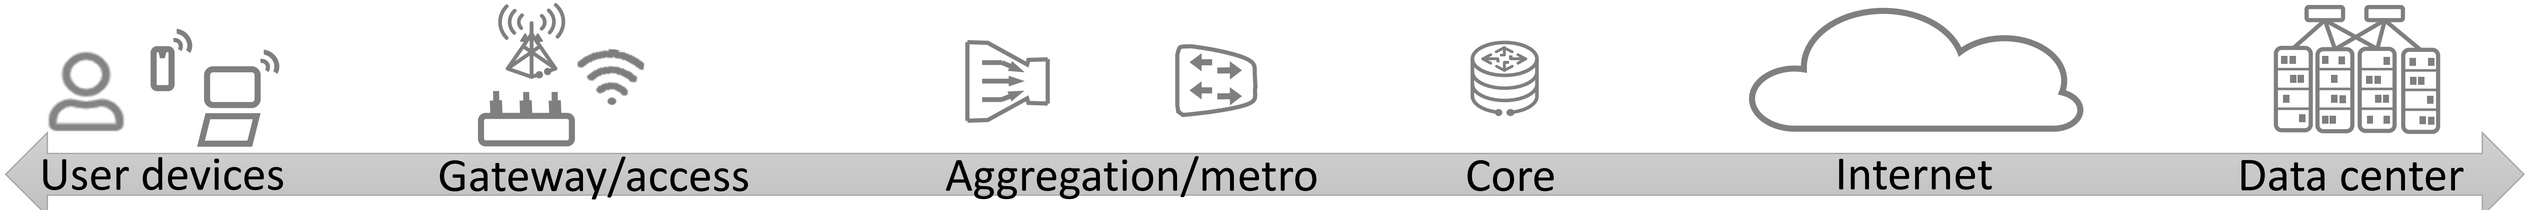

Networking data for ML / AI

Network data

So much data,

so few labels

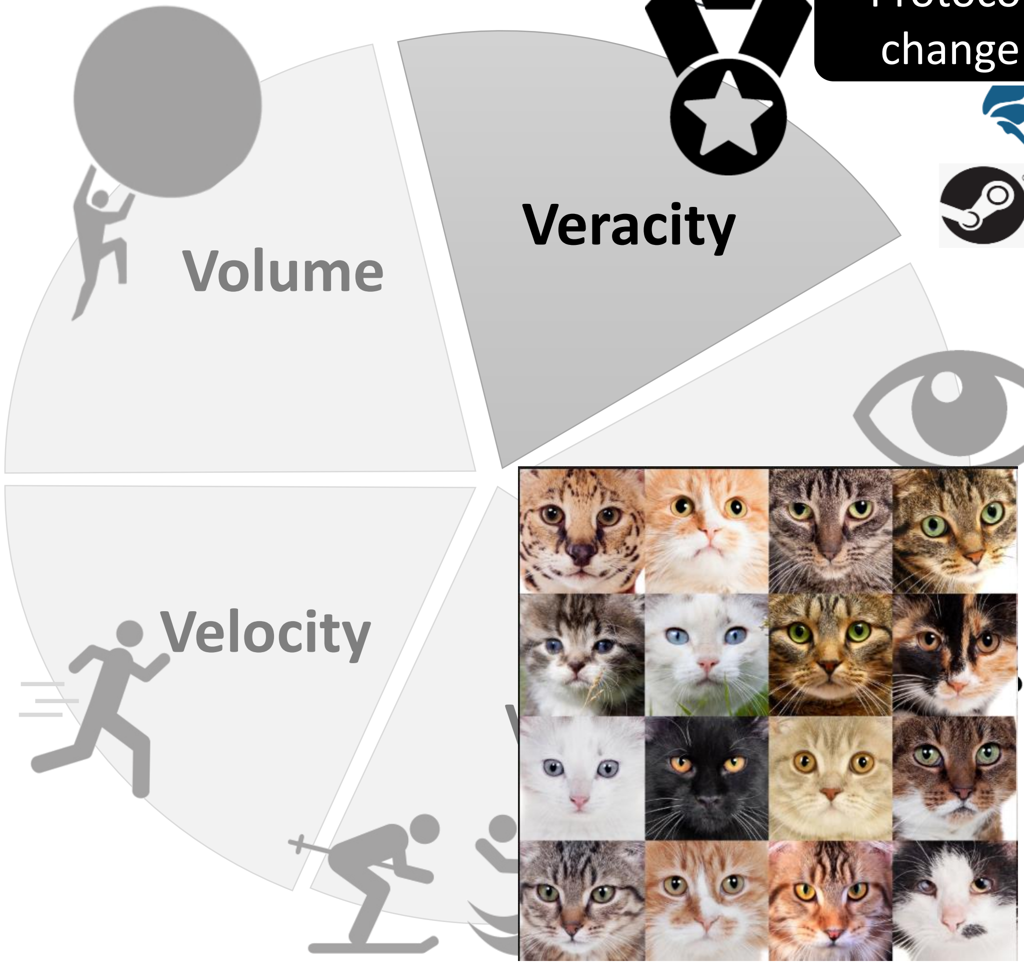

Protocol continuosuly evolve, change and die. So do *labels*

Cats are cats since 10^6 years

IMAGENET

$1.5 \cdot 10^7$ labeled images

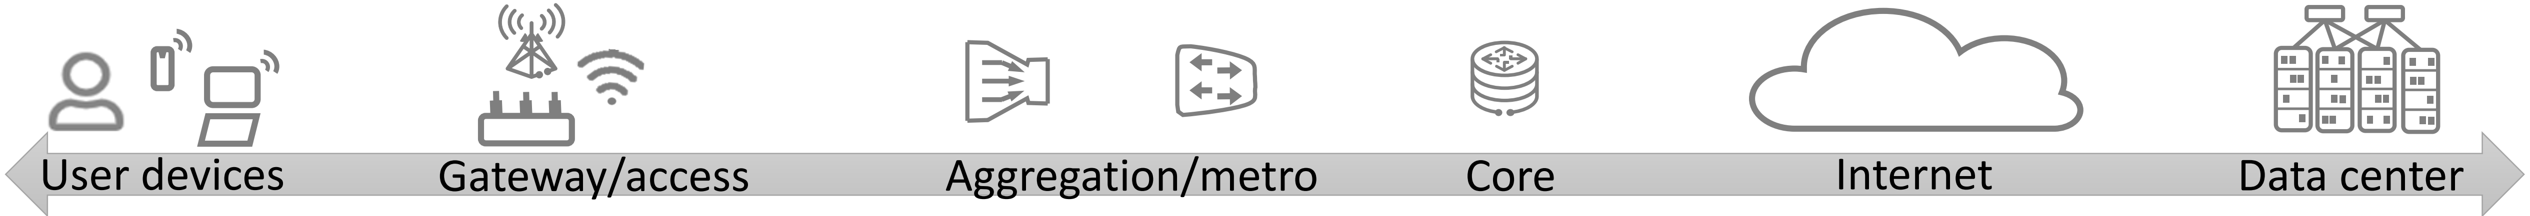

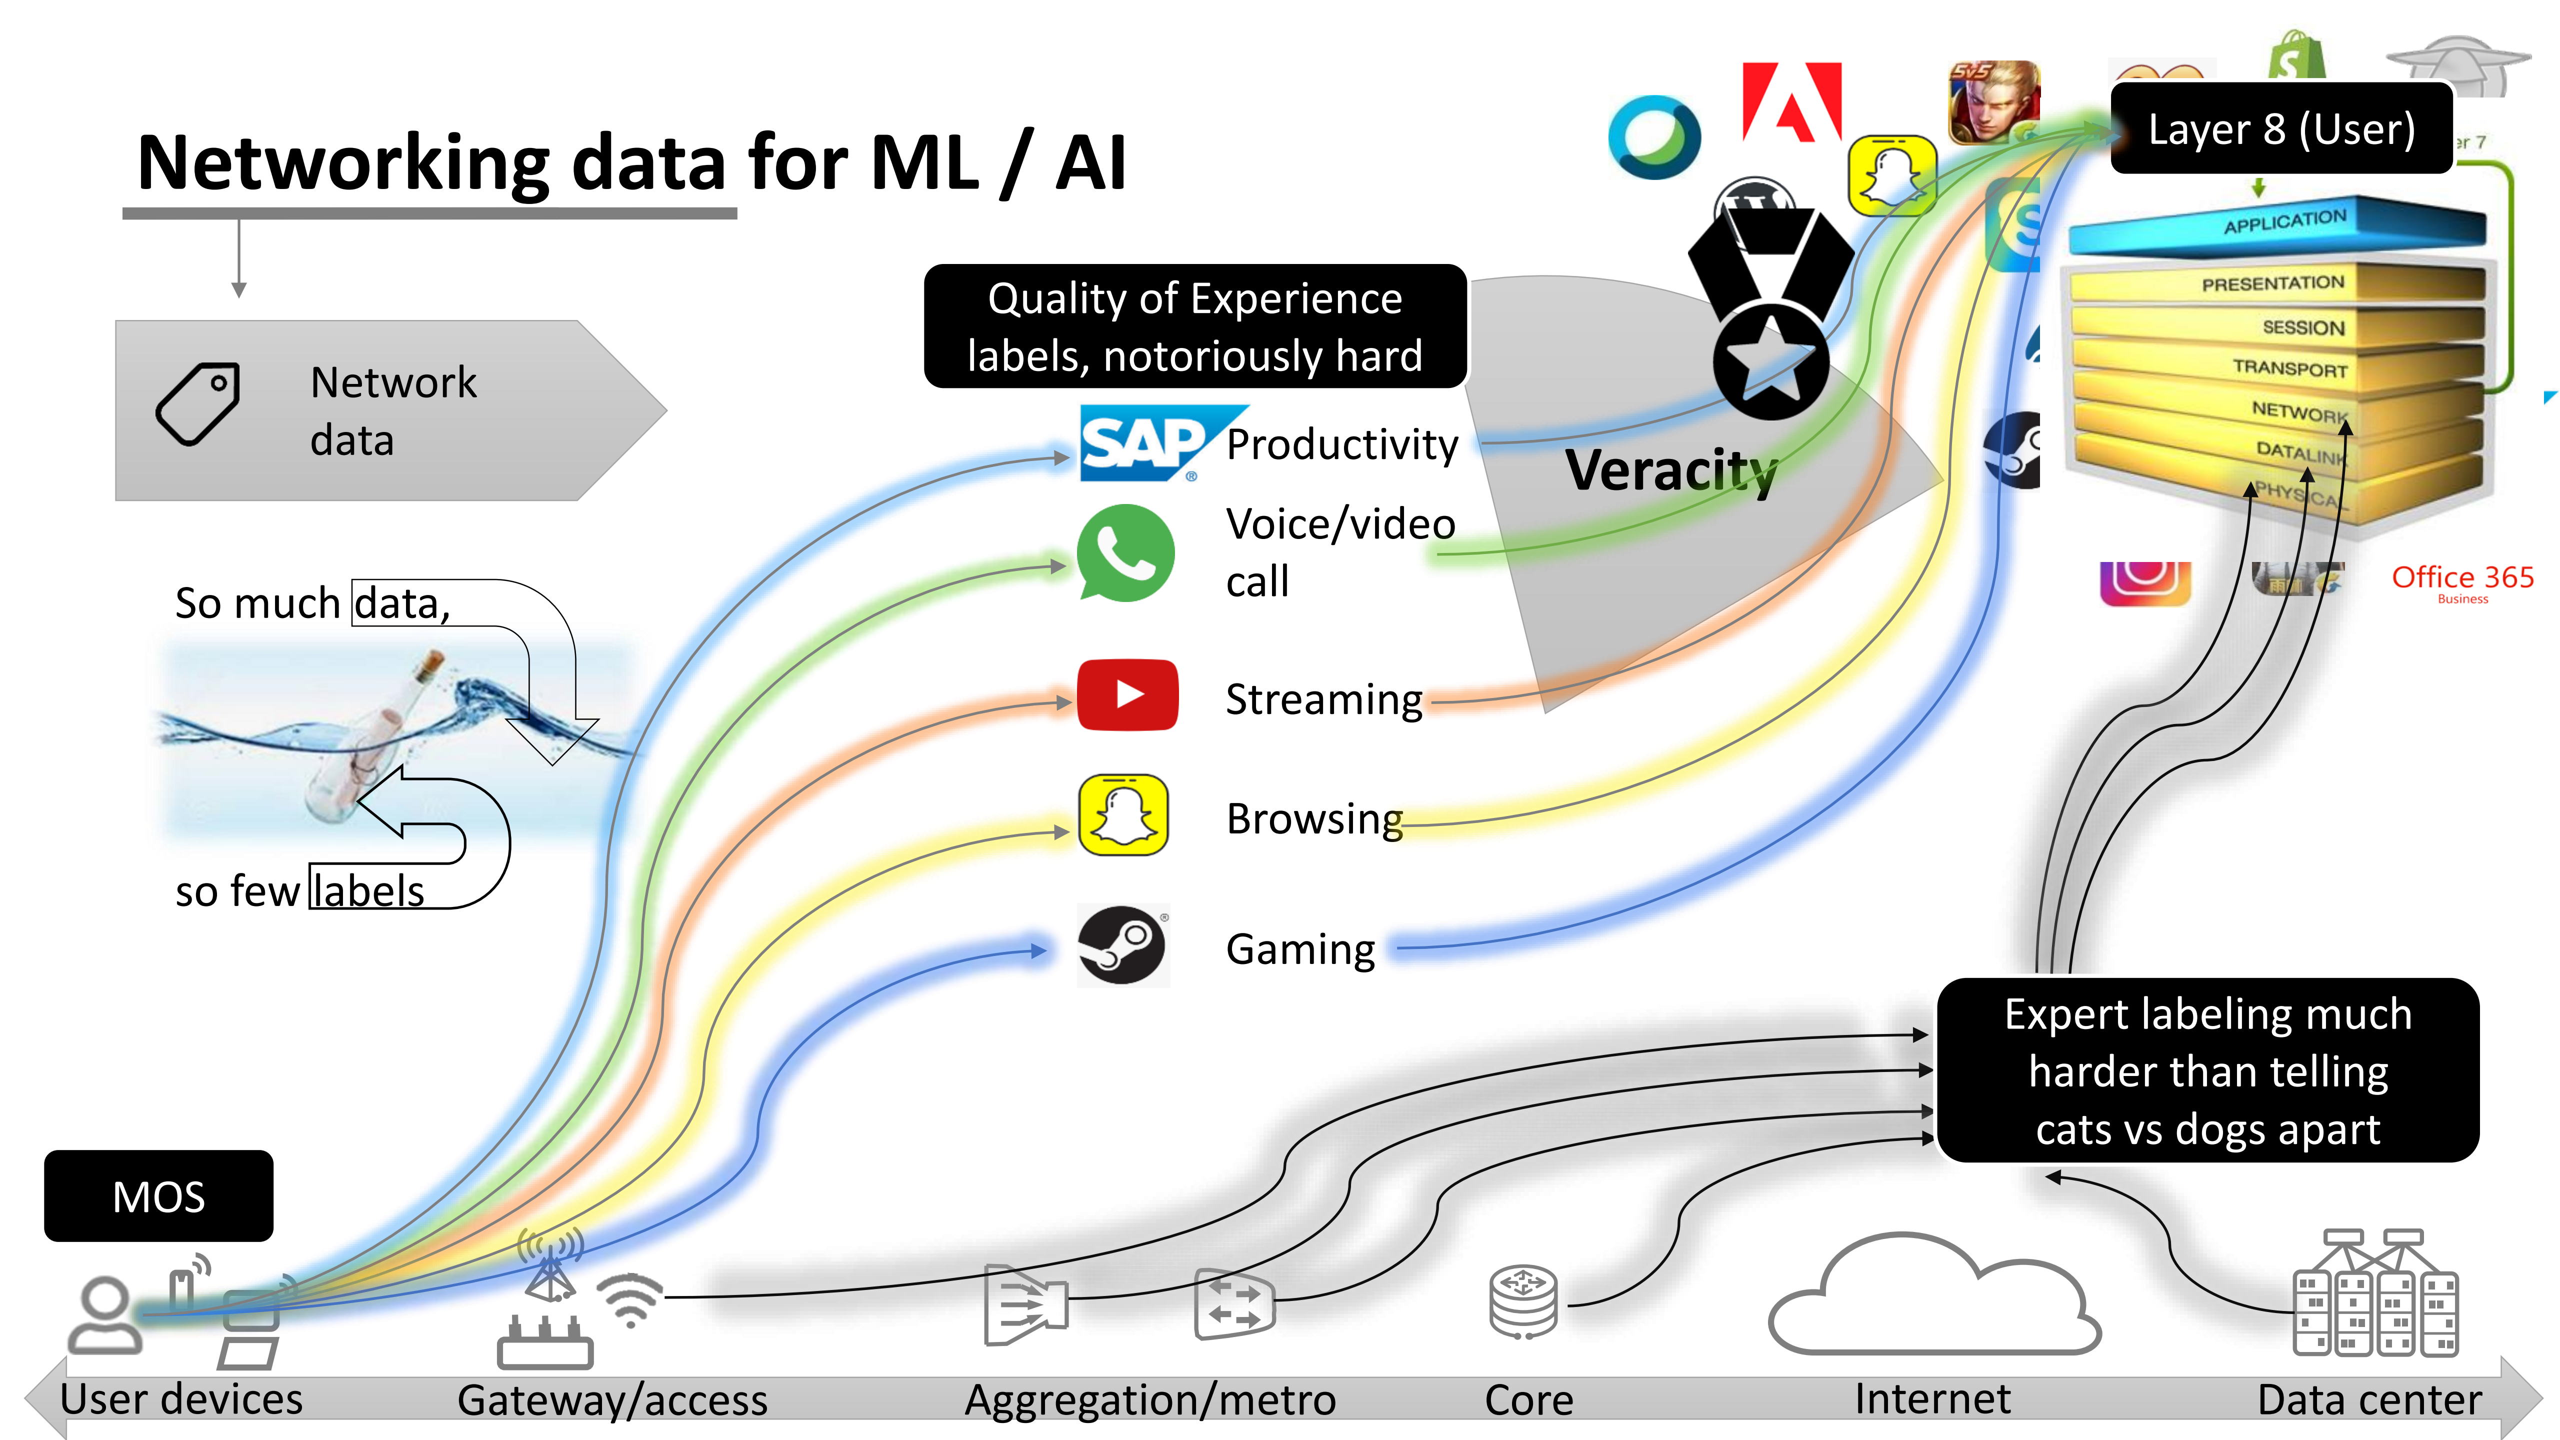

Networking data for ML / AI
Network data
So much data,
so few labels
Quality of Experience labels, notoriously hard
Productivity
Voice/video call
Streaming
Browsing
Gaming
Veracity
Layer 8 (User)
APPLICATION
PRESENTATION
SESSION
TRANSPORT
NETWORK
DATALINK
PHYSICAL
Office 365
Business
Expert labeling much harder than telling cats vs dogs apart
MOS
User devices
Gateway/access
Aggregation/metro
Core
Internet
Data center

Networking data for ML / AI

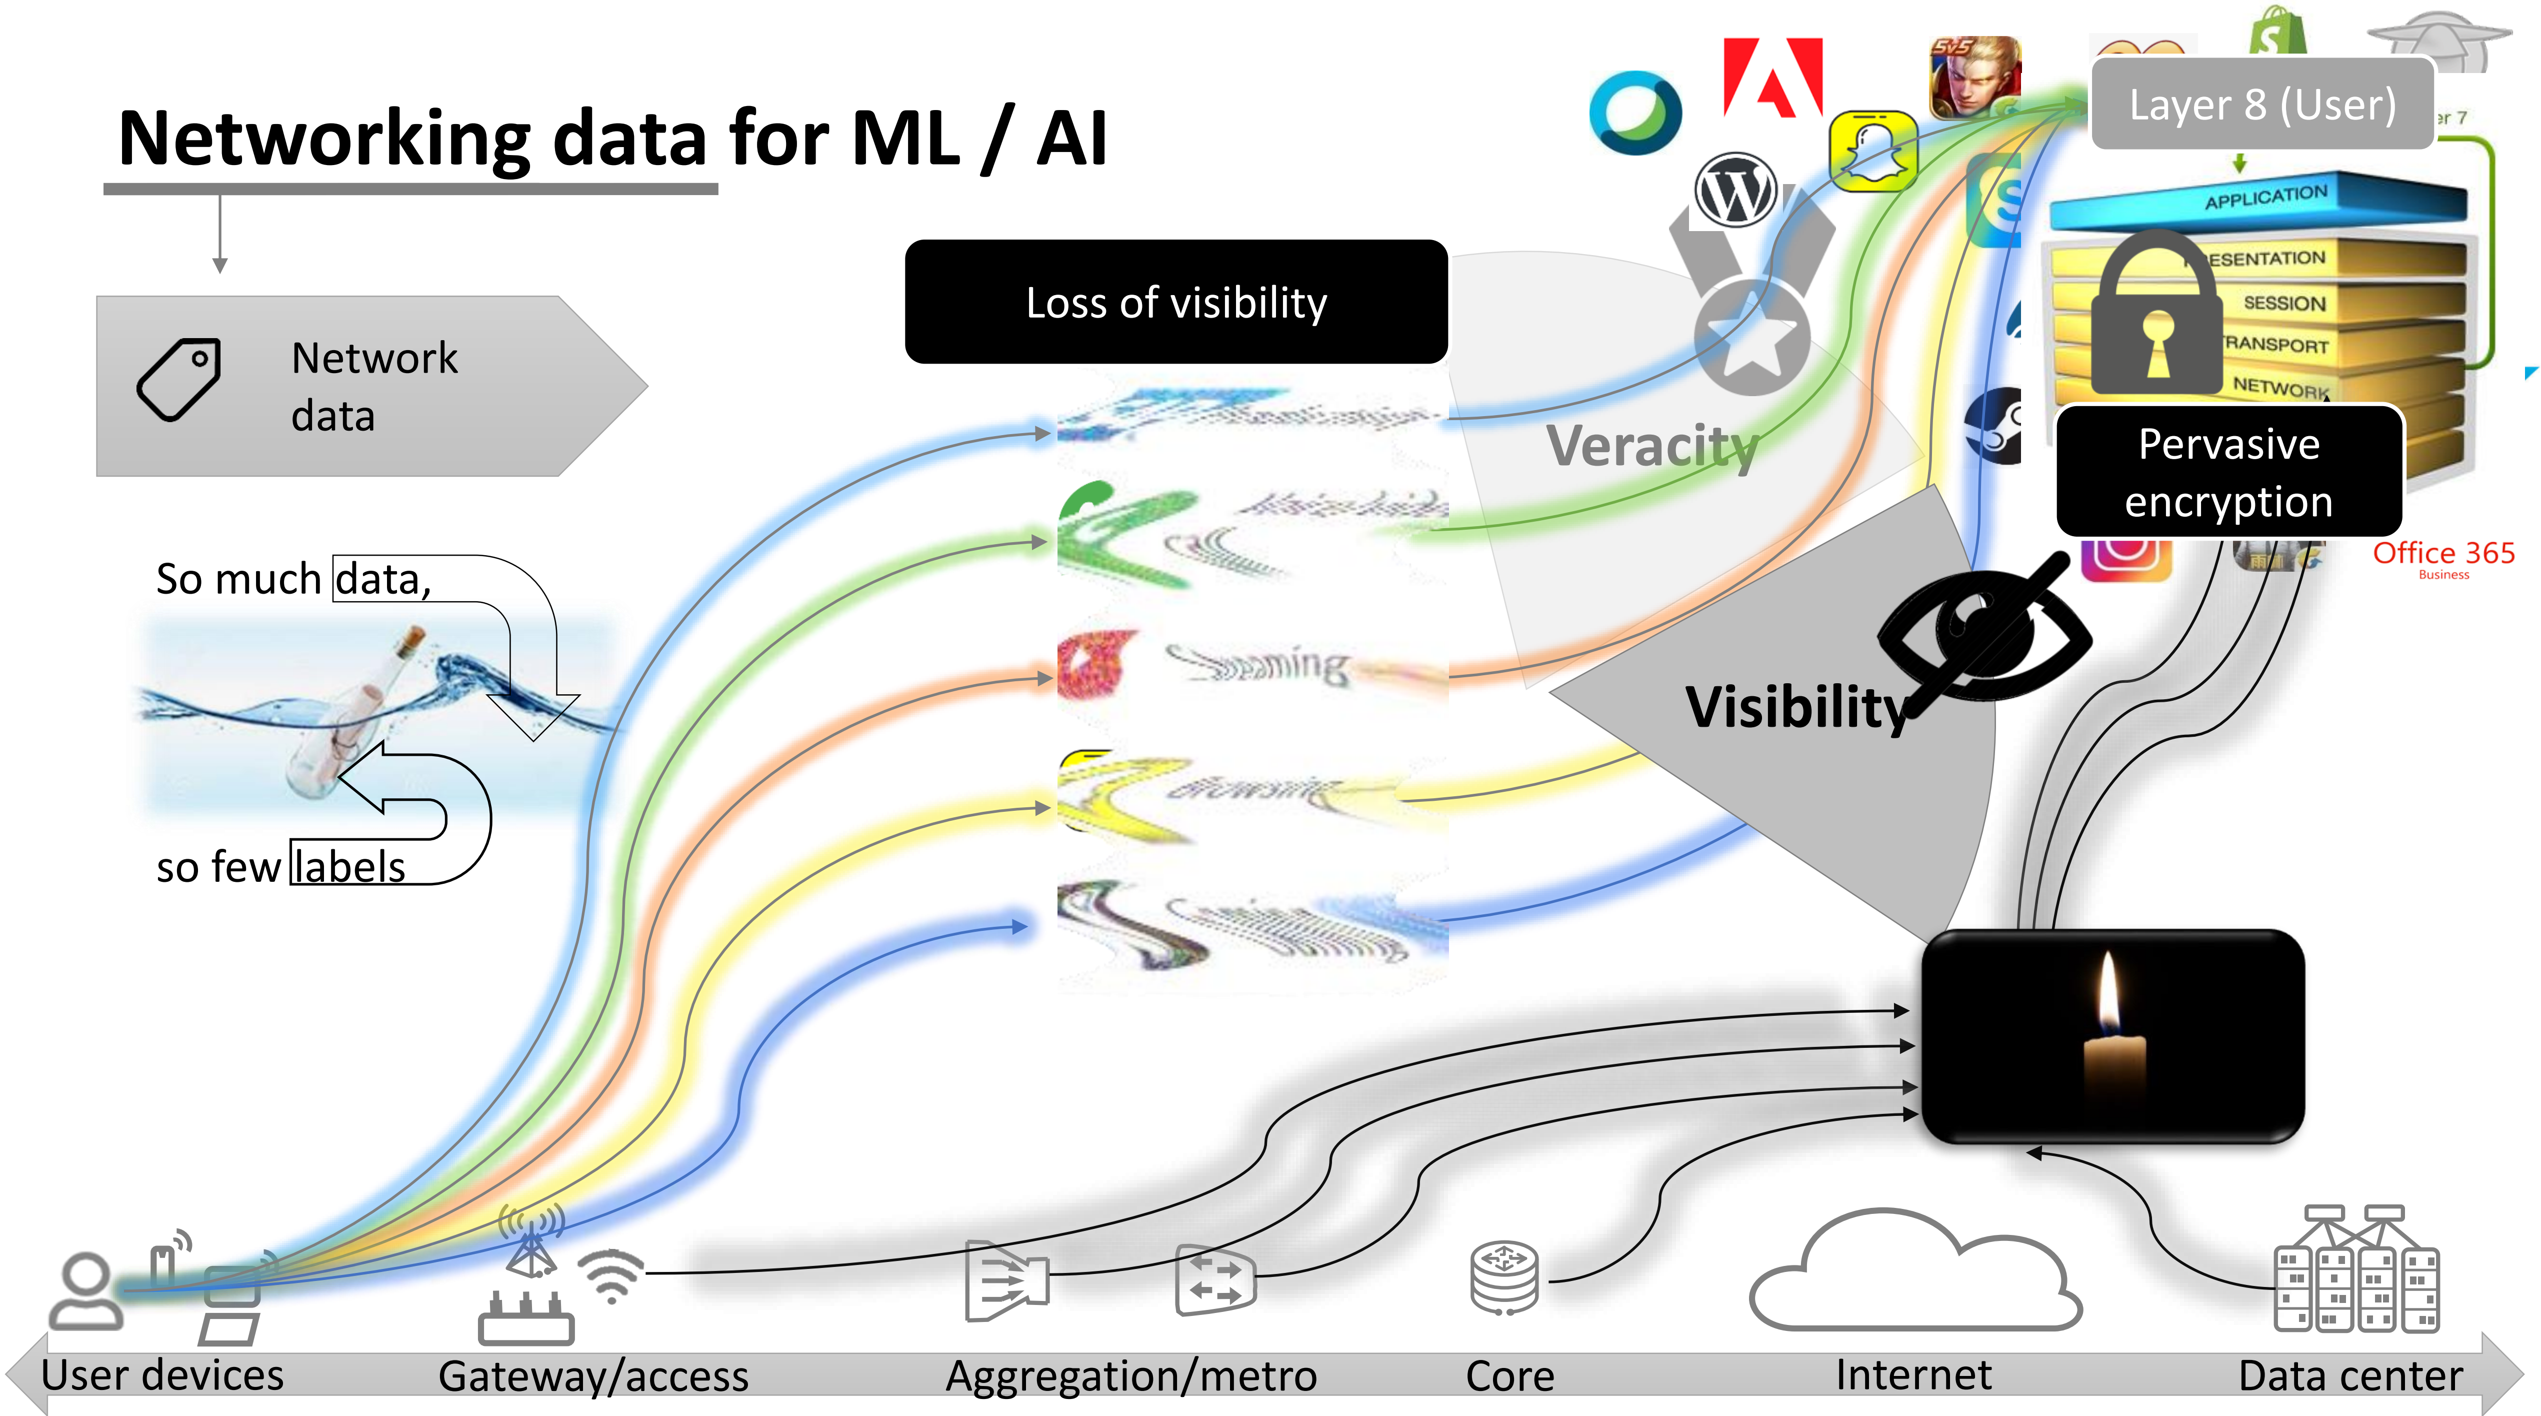

Networking data : added ML / AI value

It's optimal! (increase efficiency, same budget)
It's automated! (decrease human effort, save money)

Example: Automated Application Recognition

Application packets

Data: 1 2 3 4 5 6 ...

New traffic flows

Inputs

Labels: "Ground truth"

Labeled instances of applications of interest ... used for training

Algorithm / system

Expert models

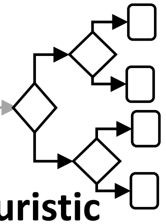

Reverse engineering & heuristic

Machine learning

Mean packet size, flow rate, timing

Feature extraction

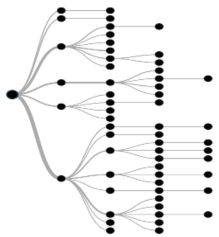

Classification

Deep Neural Networks

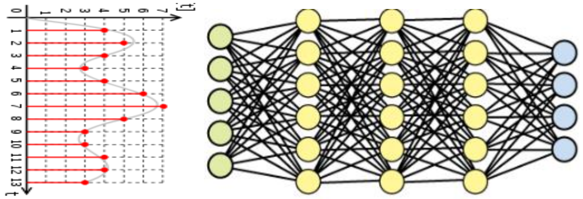

Feature extraction + Classification

Output

Prediction for each flow

- **Expert model**: manual effort, difficult to maintain
- **Machine learning**: algorithms to automatically learn optimal separation boundaries from *engineered* data
- **Deep Neural Nets**: algorithms to automatically learn non-linear functions from *raw data*

User devices | Gateway/access | Aggregation/metro | Core | Internet | Data center

Agenda

Hardware advances

Network data

Understand the network

Control the network

- ❑ History
- ❑ Trends
- ❑ AI chips

So much data,

so few labels

- ❑ Explicability
- ❑ Evolution
- ❑ Security

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Aim of this talk

Tips to avoid bumps in the road to network AI

+ Flash few examples out of our activities

ML-powered networks

Care about interpretability, not just performance as a black-box

Understand the network

Some jobs will be lost, but humans operators will remain even with self-driving networks

- ❑ Explicability
- ❑ Evolution
- ❑ Security

Several techniques inherently *as efficient as obscure*

- ❑ Convolutional Neural Networks
  - ➔ weights of densely connected neurons?
- ❑ Support Vector Machines
  - ➔ representative examples of each class?

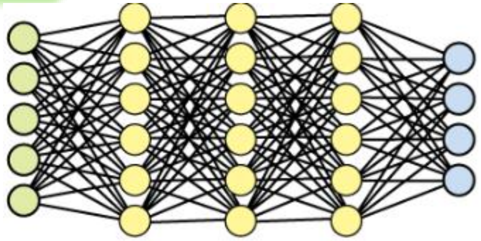

Often difficult to explain results to a *domain expert*

- ❑ Dimensionality reduction (PCA / tSNE)
  - ➔ very compact, but how to interpret?
- ❑ Outlier detection
  - ➔ along which of the many dimension?

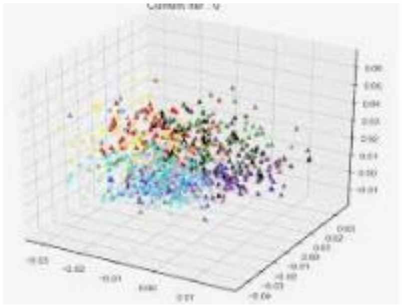

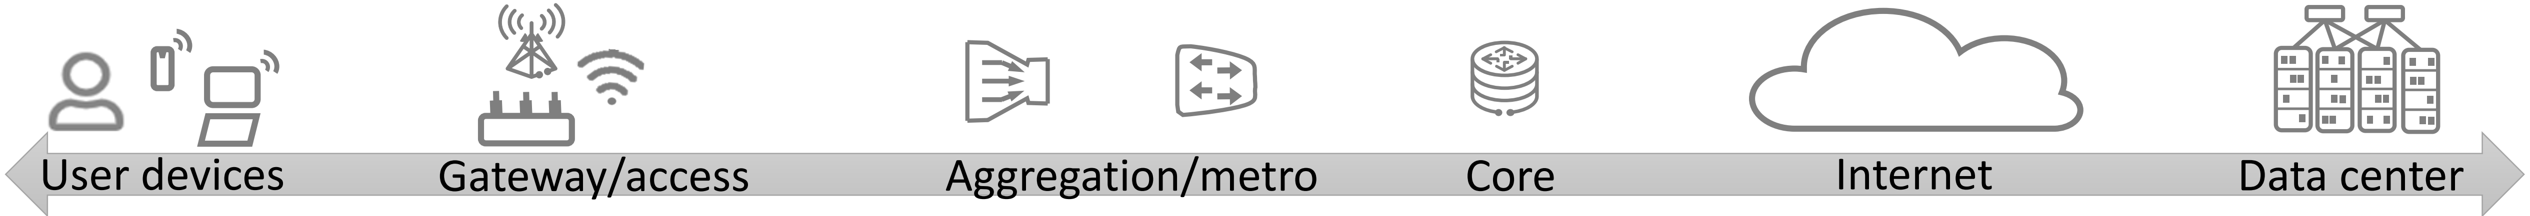

Example #1 Human-readable anomaly detection

Like Baidu for network anomalies

异常

Give to the human operator an ordered list of likely causes of anomalous behavior, in decreasing order of algorithmic importance

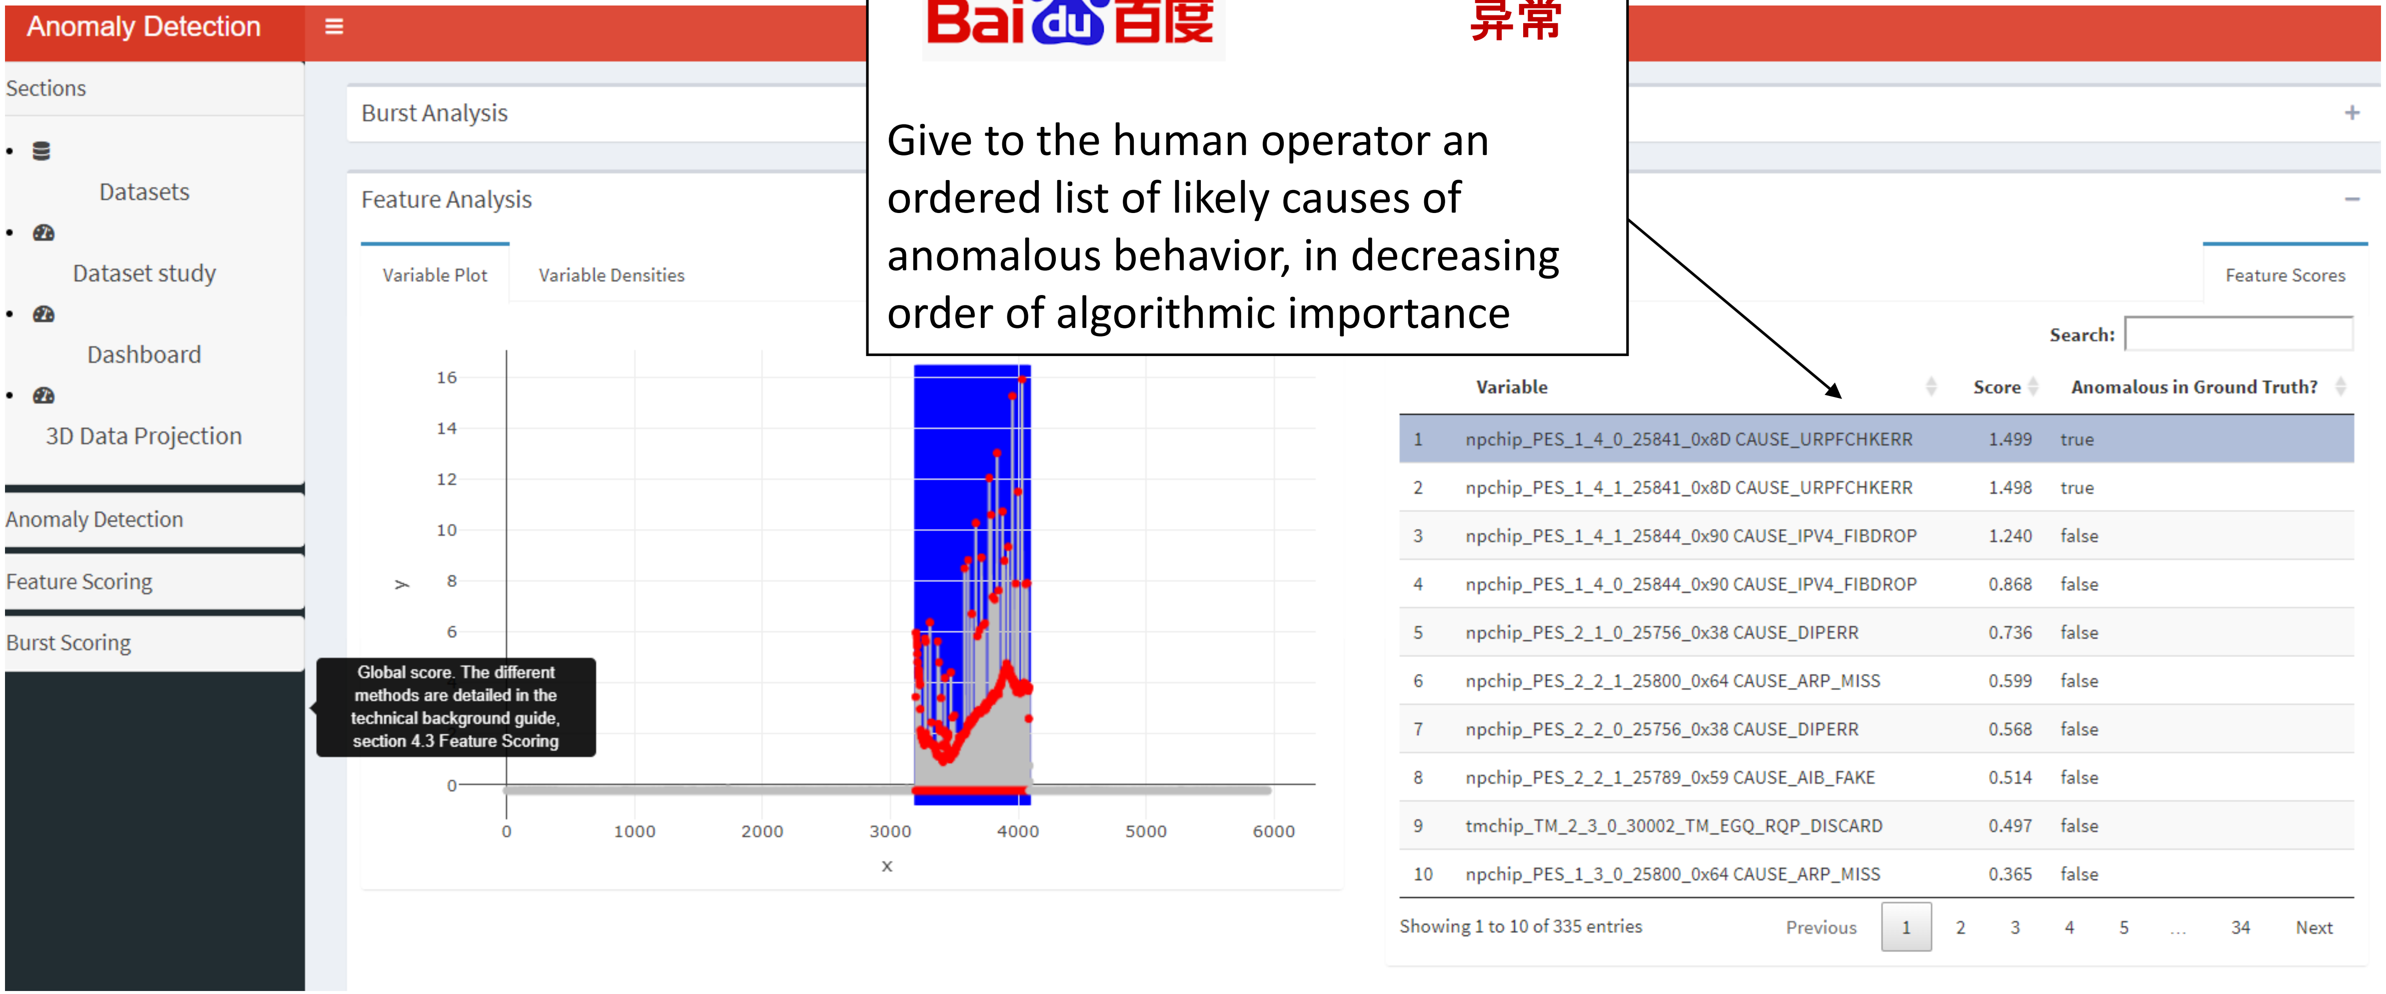

	Variable	Score	Anomalous in Ground Truth?
1	npchip_PES_1_4_0_25841_0x8D CAUSE_URPFCHKERR	1.499	true
2	npchip_PES_1_4_1_25841_0x8D CAUSE_URPFCHKERR	1.498	true
3	npchip_PES_1_4_1_25844_0x90 CAUSE_IPV4_FIBDROP	1.240	false
4	npchip_PES_1_4_0_25844_0x90 CAUSE_IPV4_FIBDROP	0.868	false
5	npchip_PES_2_1_0_25756_0x38 CAUSE_DIPERR	0.736	false
6	npchip_PES_2_2_1_25800_0x64 CAUSE_ARP_MISS	0.599	false
7	npchip_PES_2_2_0_25756_0x38 CAUSE_DIPERR	0.568	false
8	npchip_PES_2_2_1_25789_0x59 CAUSE_AIB_FAKE	0.514	false
9	tmchip_TM_2_3_0_30002_TM_EGQ_RQP_DISCARD	0.497	false
10	npchip_PES_1_3_0_25800_0x64 CAUSE_ARP_MISS	0.365	false

[INFOCOM'20] J.M.Navarro et al. HURRA! Human-Readable Router Anomaly Detection IEEE Infocom, Demo session

ML-powered networks

Understand the network

- ❑ Explicability
- ❑ **Evolution**
- ❑ Security

Online/streaming ML algorithms

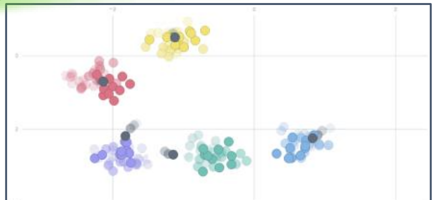

Unsupervised

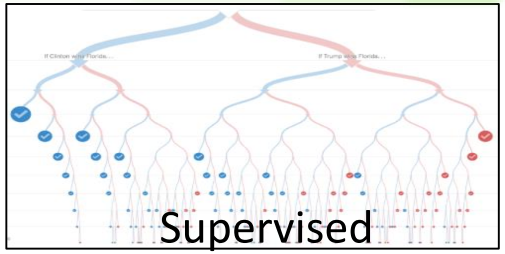

Supervised

- ❑ *Network evolves, so should your models*
- Clustering (e.g, Dgrid, DenStream, CluStream)
- Trees (e.g., Hoeffding tree, Adaptive Random Forest)

Model fusion

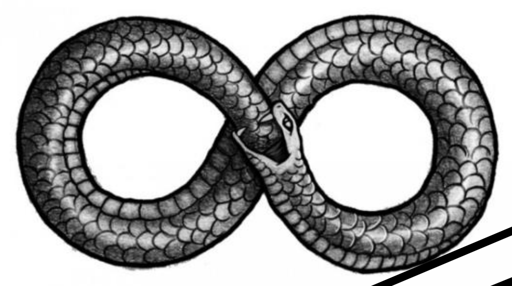

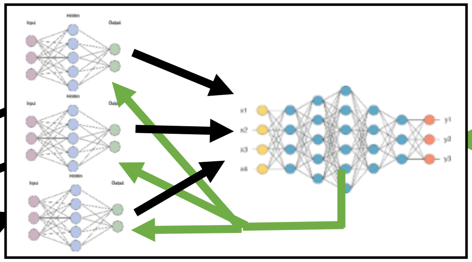

Federated learning

- ❑ *Networks have a large set of sensors, fusing this models better than exchanging data*
- Federated Learning (at the edge)
- Transfer Learning (more general concept)

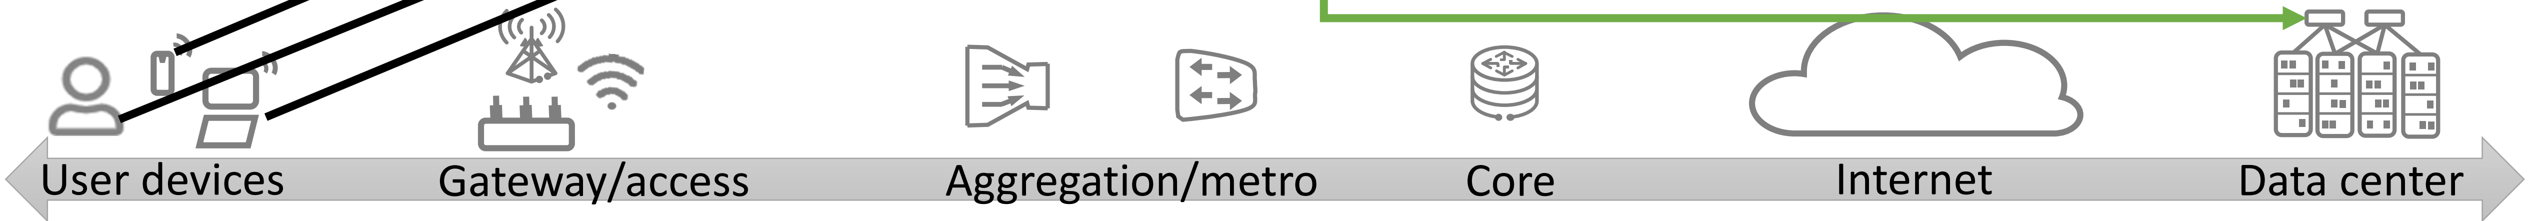

ML-powered networks

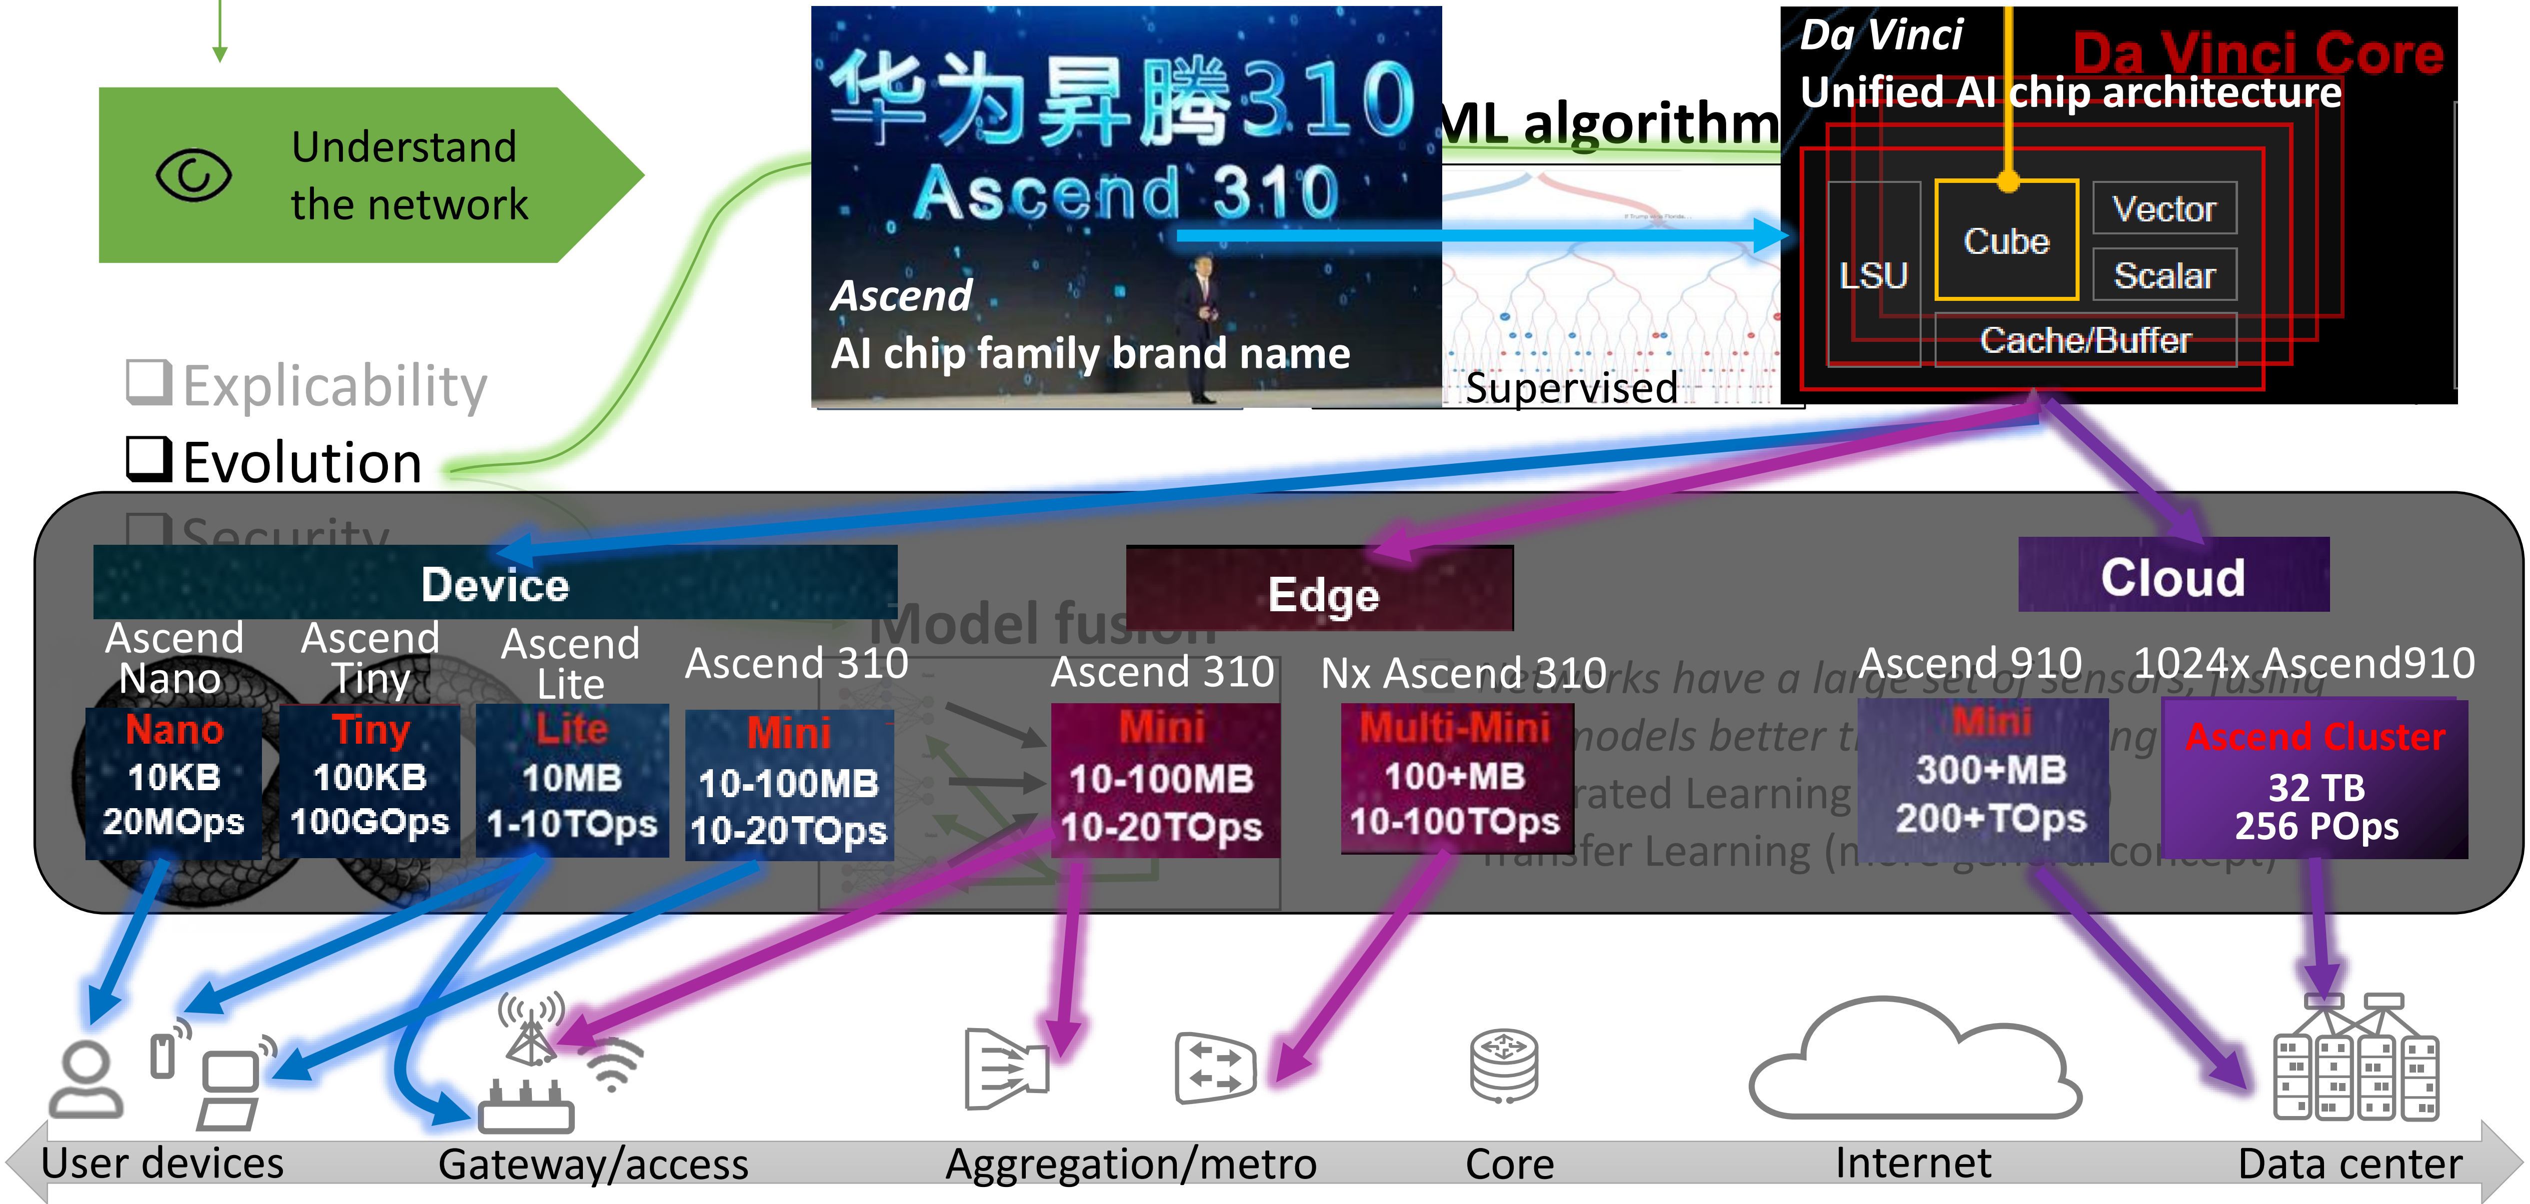

ML-powered networks

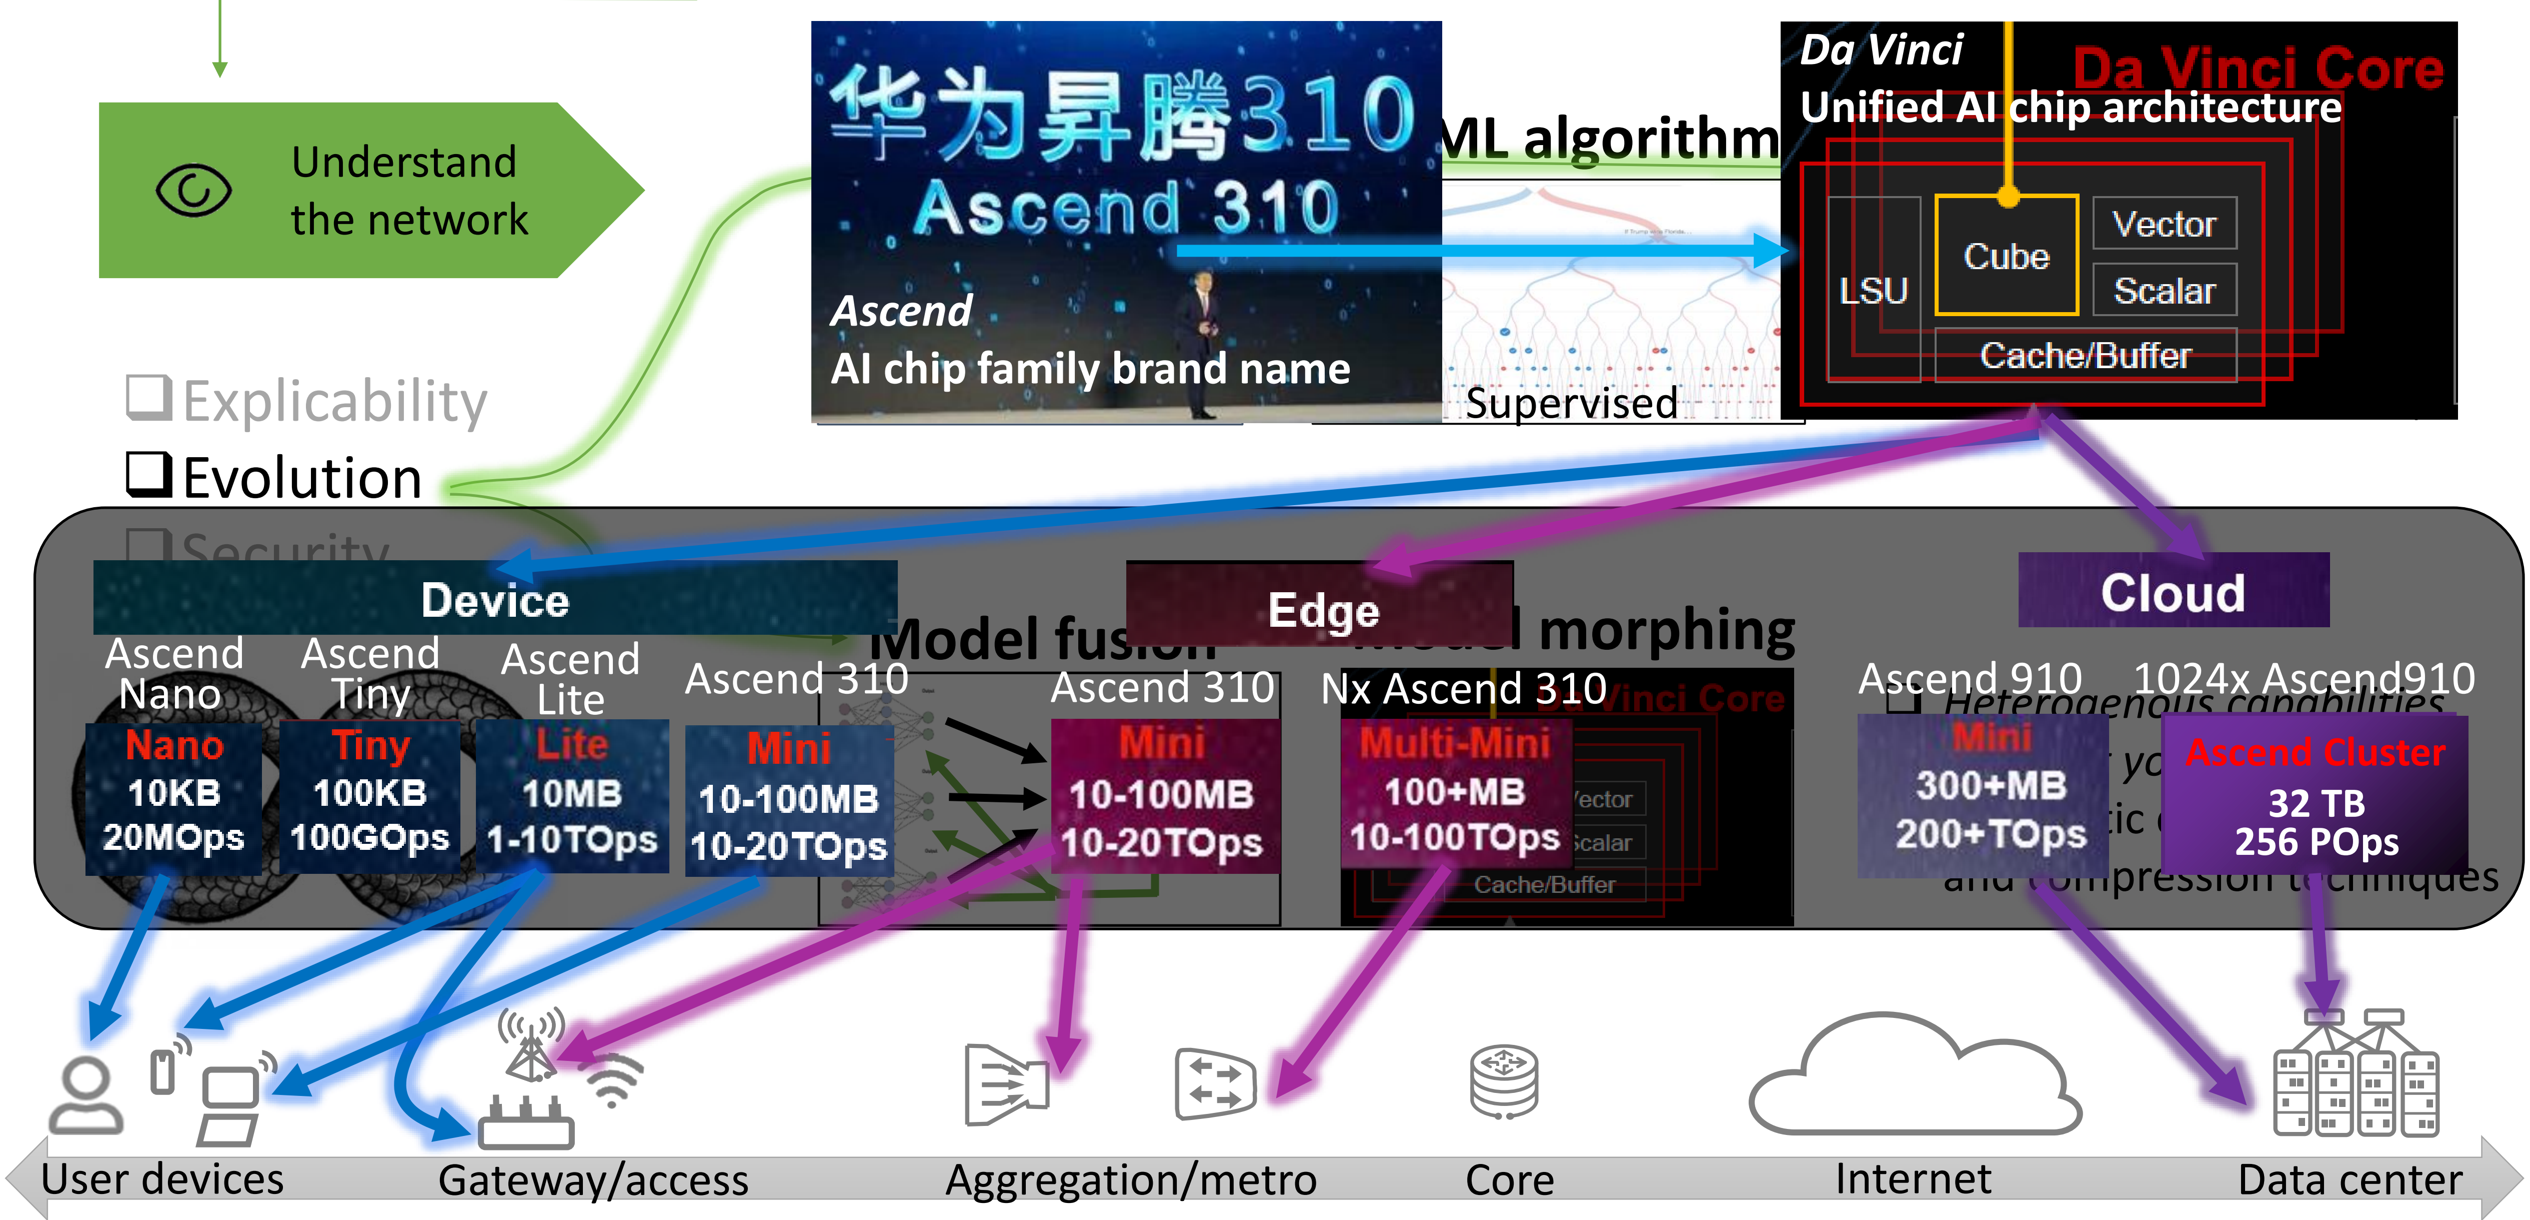

ML-powered networks

In ML, the journey matters more than the destination

Understand the network

- ❑ Explicability
- ❑ **Evolution**
- ❑ Security

Online/streaming ML algorithms

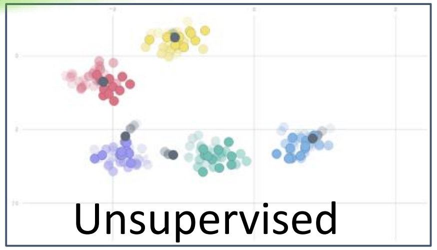

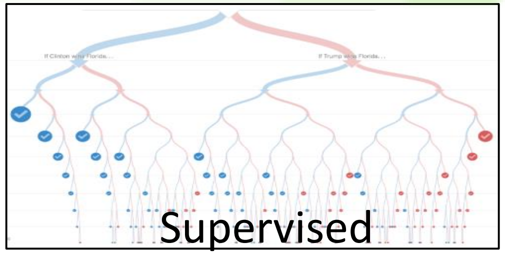

- ❑ *Network evolves, so should your models*
- ▪ Clustering (e.g, Dgrid, DenStream, CluStream)
- ▪ Trees (e.g., Hoeffding tree, Adaptive Random Forest)

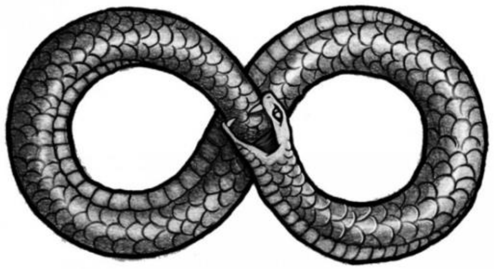

Model fusion

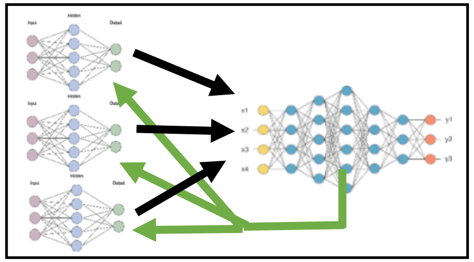

Federated learning

+Model morphing

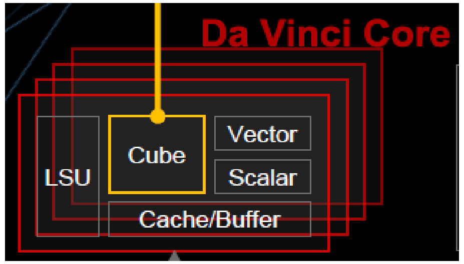

+ Embrace AIOps

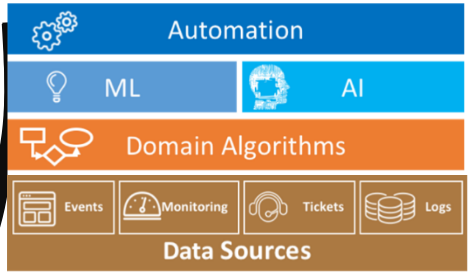

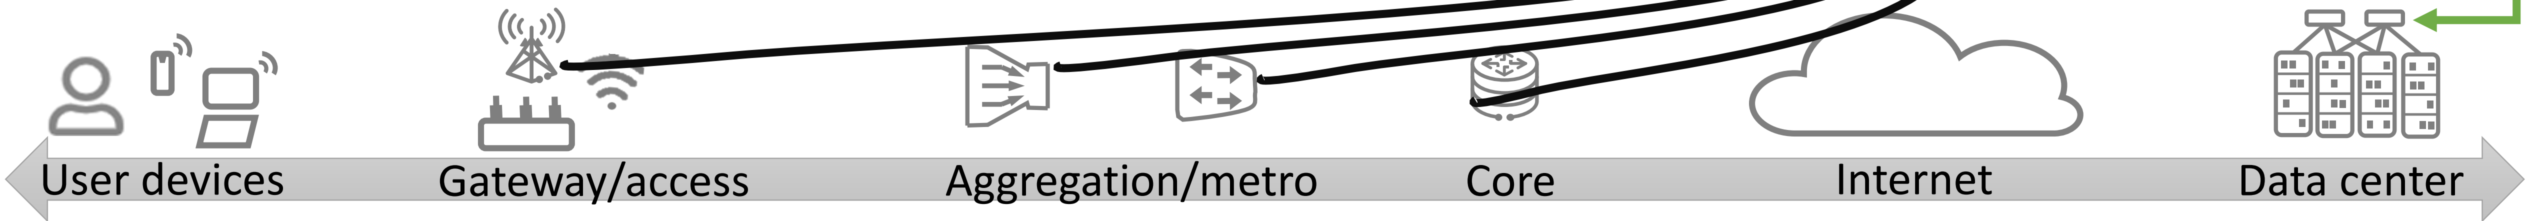

Example #2 Encrypted & unknown traffic classification

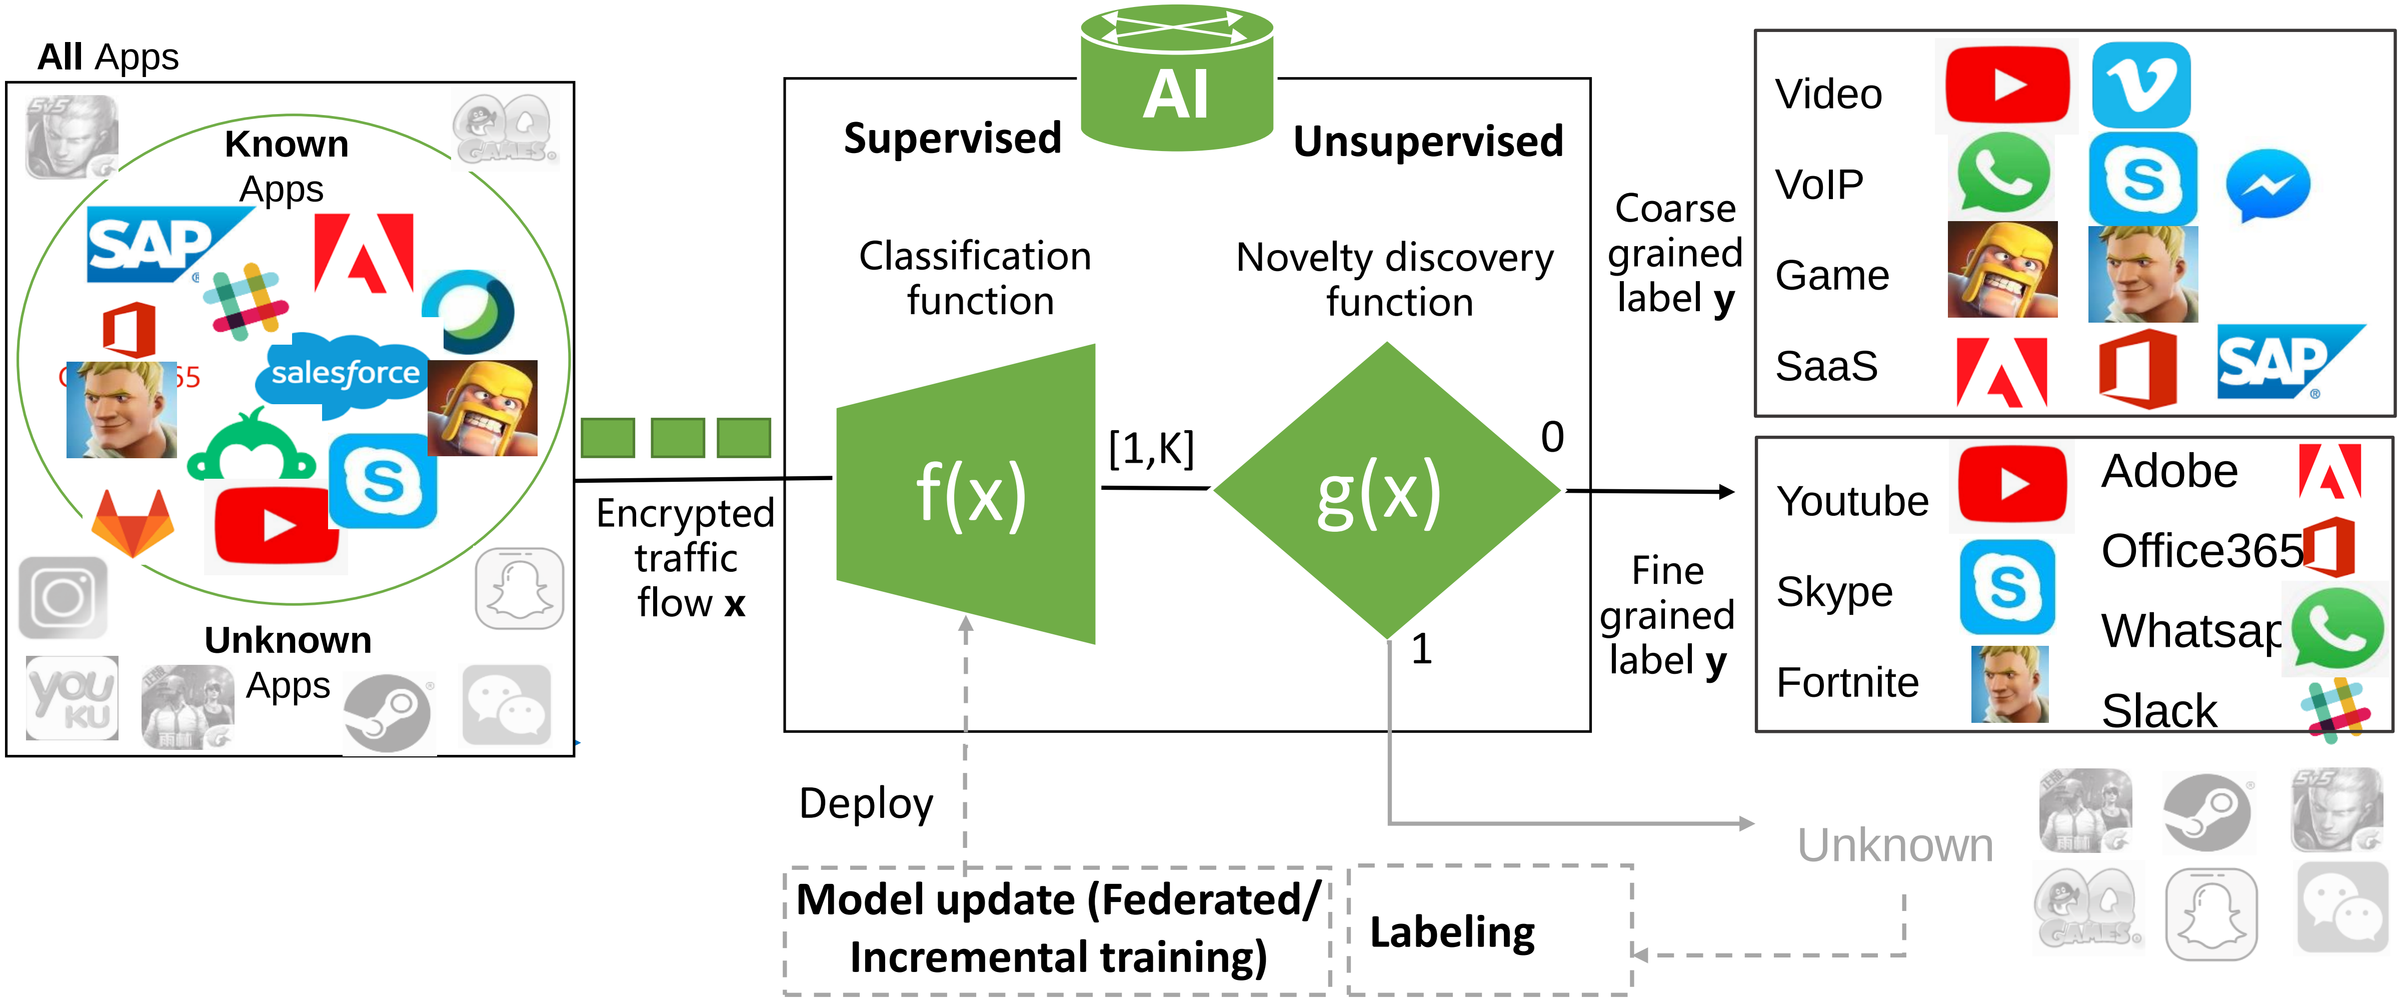

[IJCAI'20] L. Yang et al. Heterogeneous Data-Aware Federated Learning, International Joint Conference on Artificial Intelligence, FL workshop

[INFOCOM'20] C. Beliard et al. Opening the Deep Pandora Box: Explainable Traffic Classification IEEE Infocom, Demo session

Example #2 Encrypted & unknown traffic classification

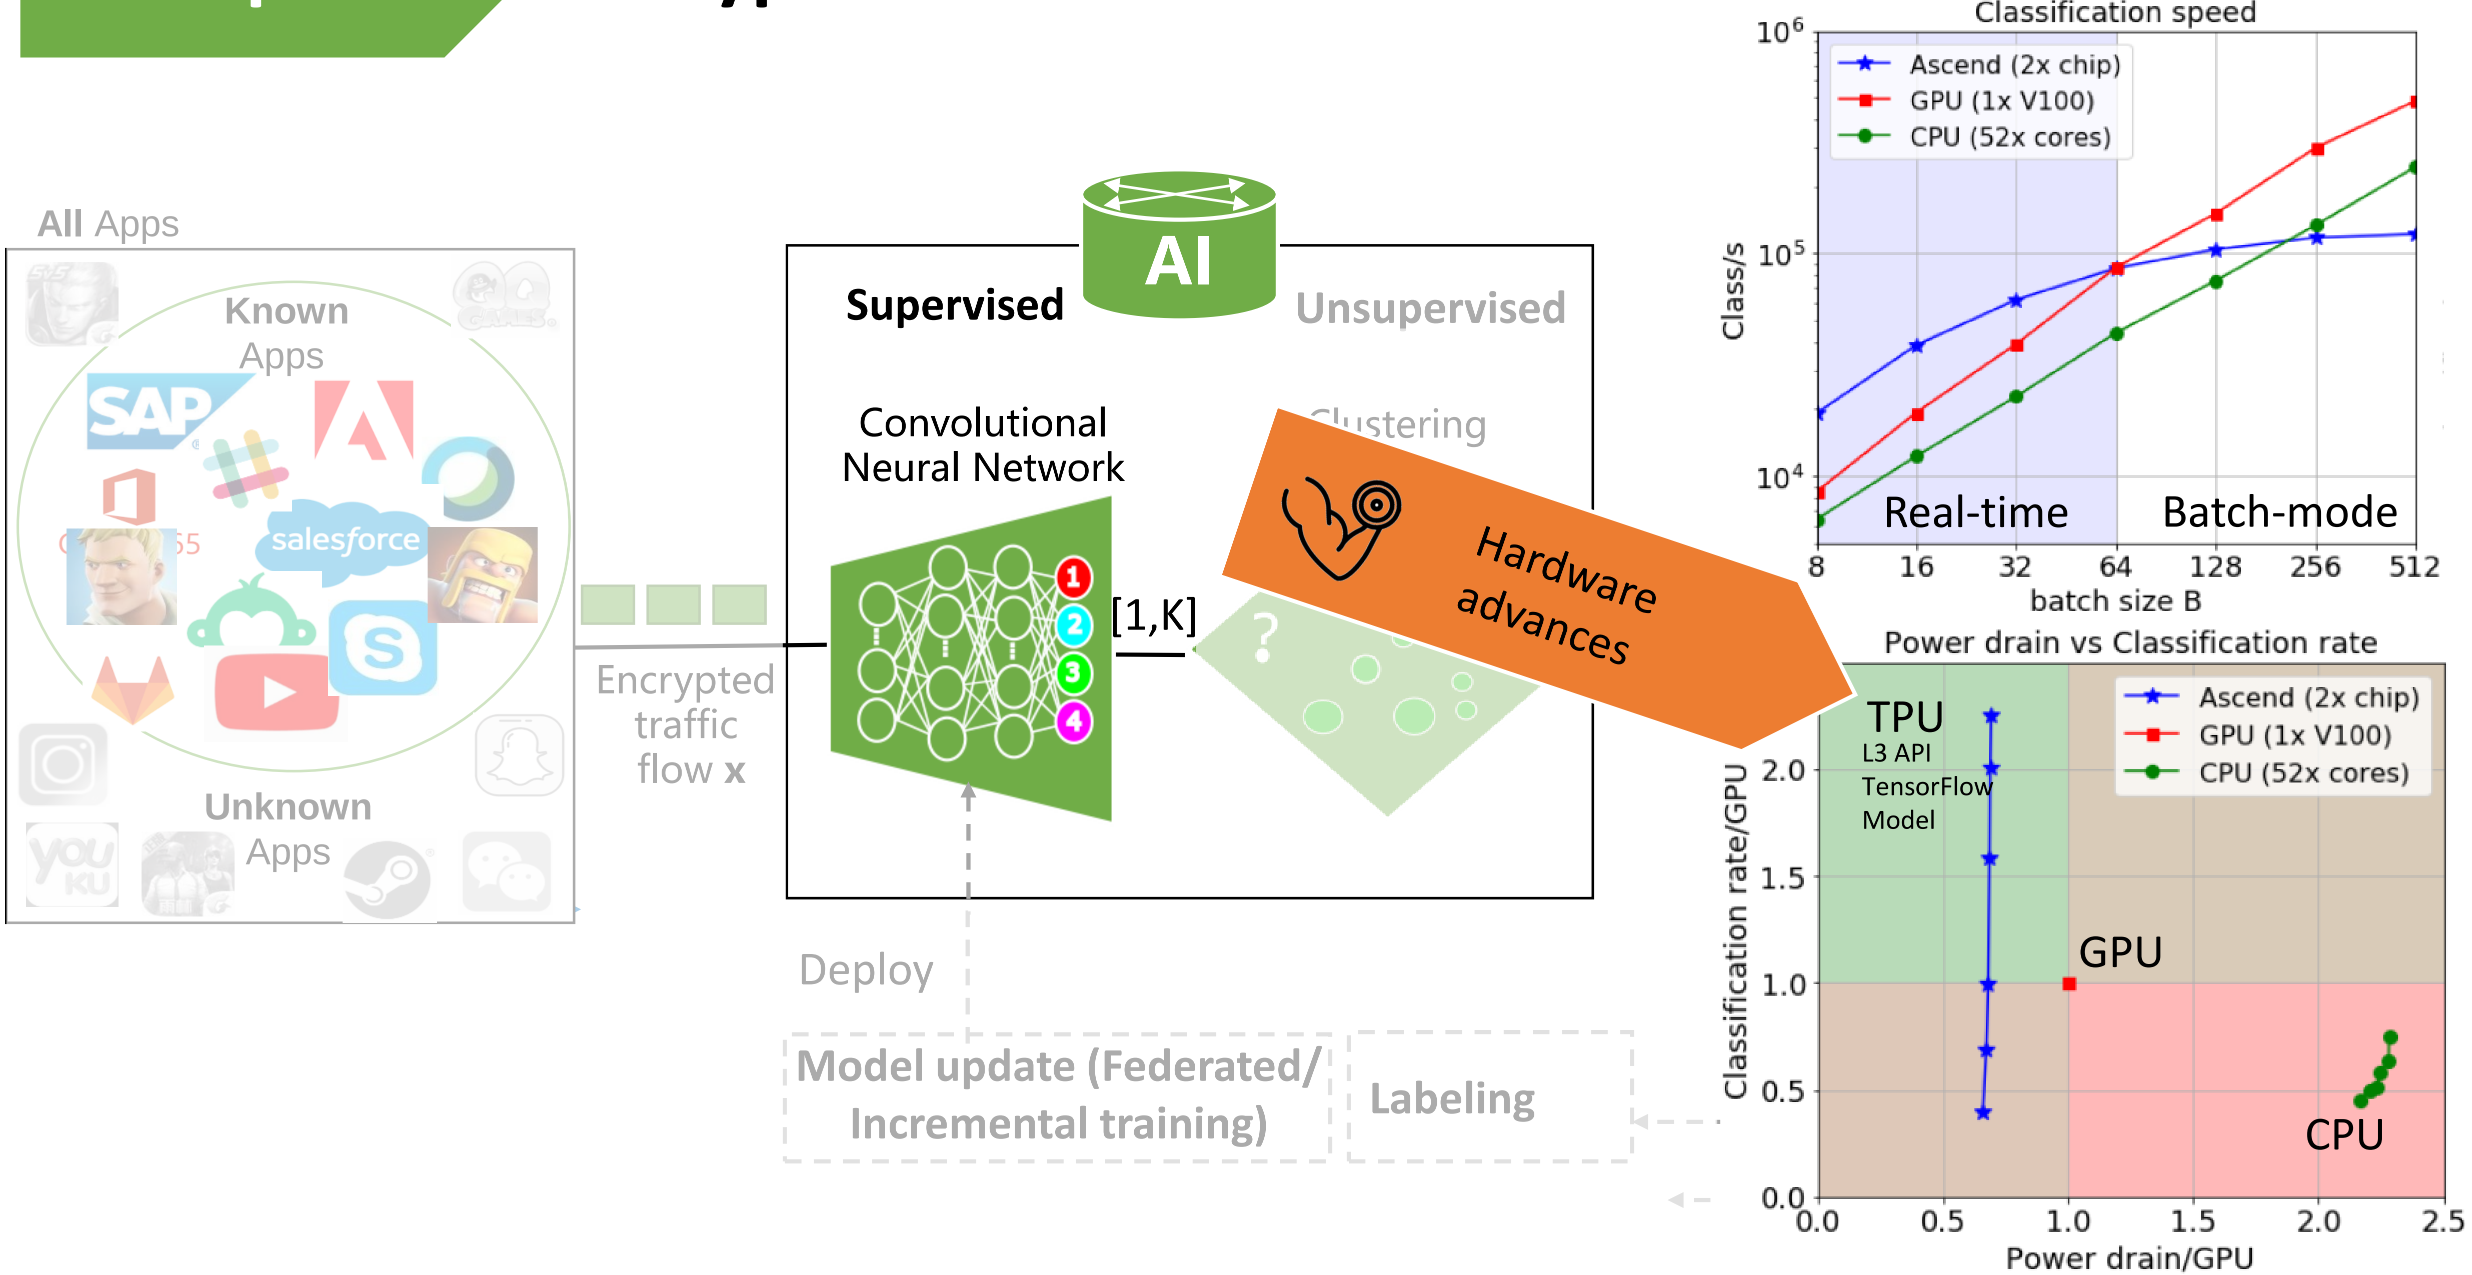

Just as network protocols,
ML can (& will) be hacked

ML-powered networks

Understand the network

- ❑ Explicability
- ❑ Evolution
- ❑ **Security**

ML Evasion

- ❑ Can happen locally, when a model is deployed
- ❑ *E.g.*, Adversary circumvents/alters traffic classification results by purposely altering its own features

Adversarial ML

- ❑ Can happen for streaming techniques, during the learning phase
- ❑ Adversary alters the ML training process by purposedly mislabeling data, affects all systems

Panda Perturbation Gibbon

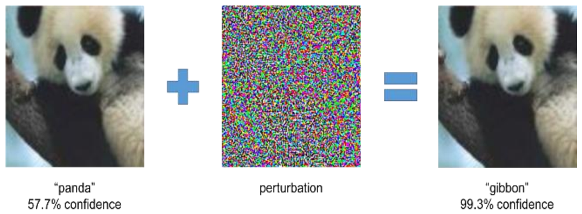

Leak of sensitive information

- ❑ E.g, adversary extracts information from shared/accessible ML models

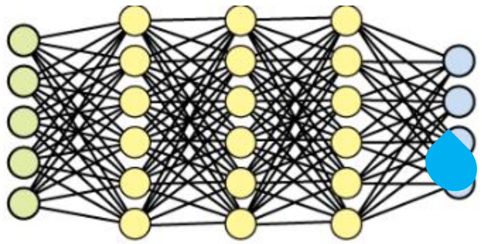

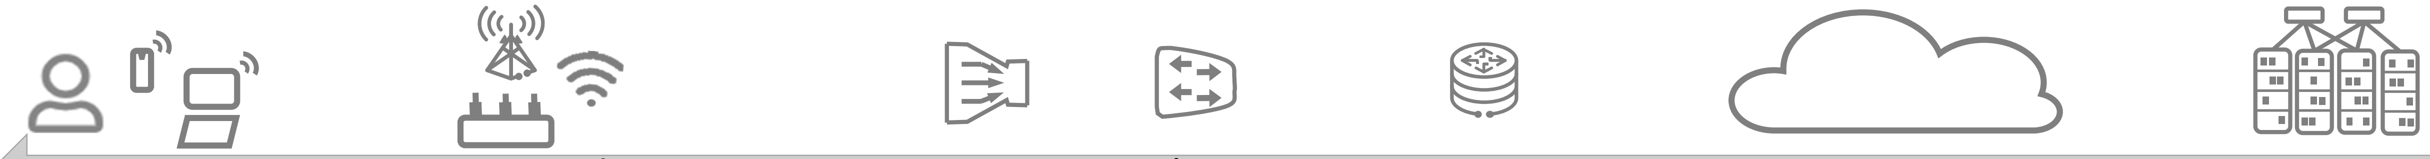

User devices | Gateway/access | Aggregation/metro | Core | Internet | Data center

Agenda

Hardware advances

Network data

Understand the network

Control the network

- ❑ History
- ❑ Trends
- ❑ AI chips

So much data,

so few labels

- ❑ Explicability
- ❑ Evolution
- ❑ Security

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Aim of this talk

Tips to avoid bumps in the road to network AI

+ Flash few examples out of our activities

AI-powered networks

When closing the loop, mind the gap!

Control the network

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

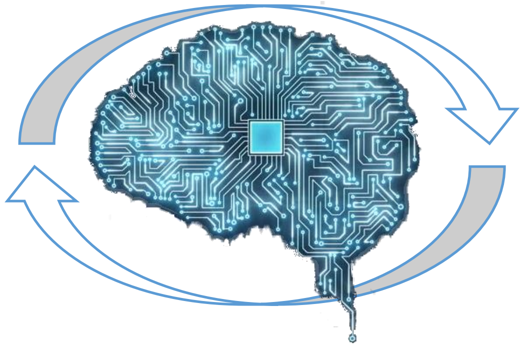

Games (Go state space ~ 10^{100})

- ❑ AlphaGo (10,000s of human amateur and professional games, 3 days training, 1920 CPUs, 280 GPUs, elo rating 3.16)
- ❑ AlphaGo Zero (simply plays against itself, 4 TPUs, 40 days to beat AlphaGo, achieving elo rating 5.16)/AlphaZero/MuZero
- ❑ Portability? Add one row ▣ to the board !! Add a ● player !?

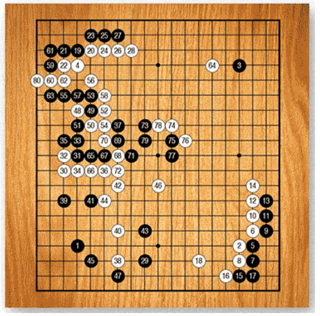

Networks (state space $\mathbb{R}^N$, with N>>100)

- ❑ Portability is essential: you cannot sell an AI product that will make performance *worse* for over a month !
- ❑ Results coupled with delay of telemetry, and delay to actuate actions in the controller
- ❑ Convergence speed matters ! for any techniques (Reinforcement learning, Deep reinforcement learning, Stochastic optimization, etc.)

Example #3 WLAN traffic optimization

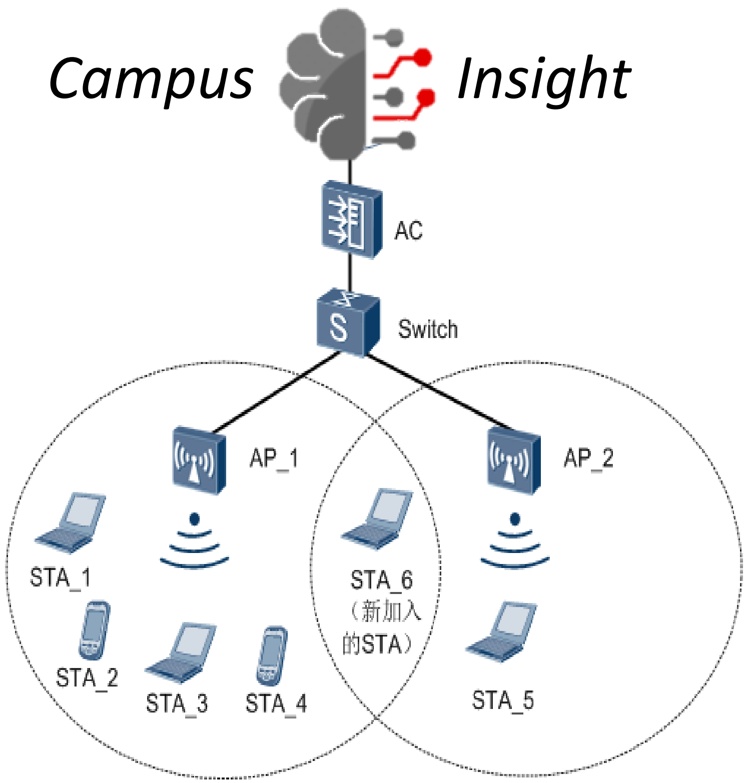

(Deep) reinforcement learning

Reward= f(T, Δ, QoE, I, RSSI, …)

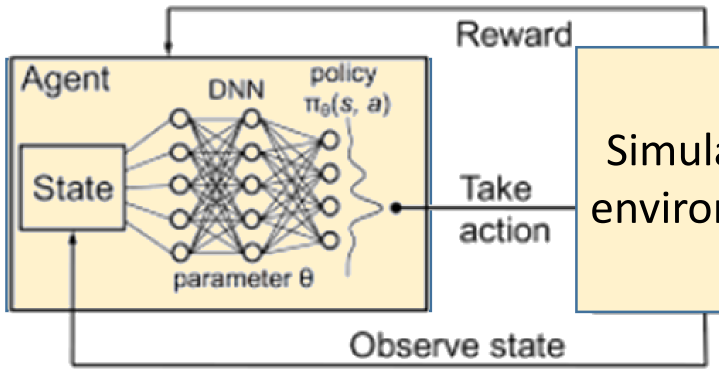

Simulation environment	Emulation environment	Real world

Speedup state exploration

Combine multiple environments

Simulation	Emulation	Real world

Example #3 WLAN traffic optimization

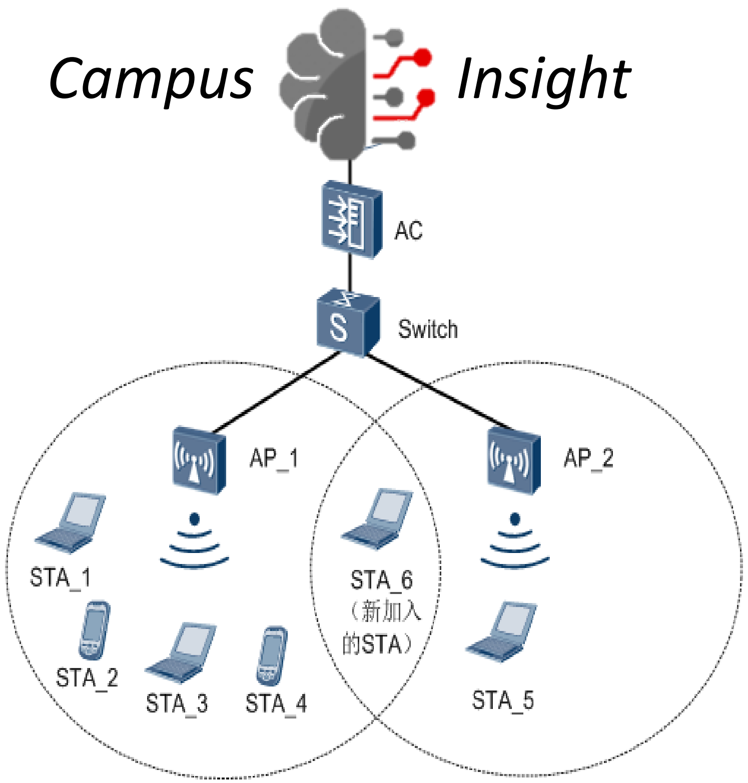

(Deep) reinforcement learning

Reward= f(T, Δ, QoE, I, RSSI, …)

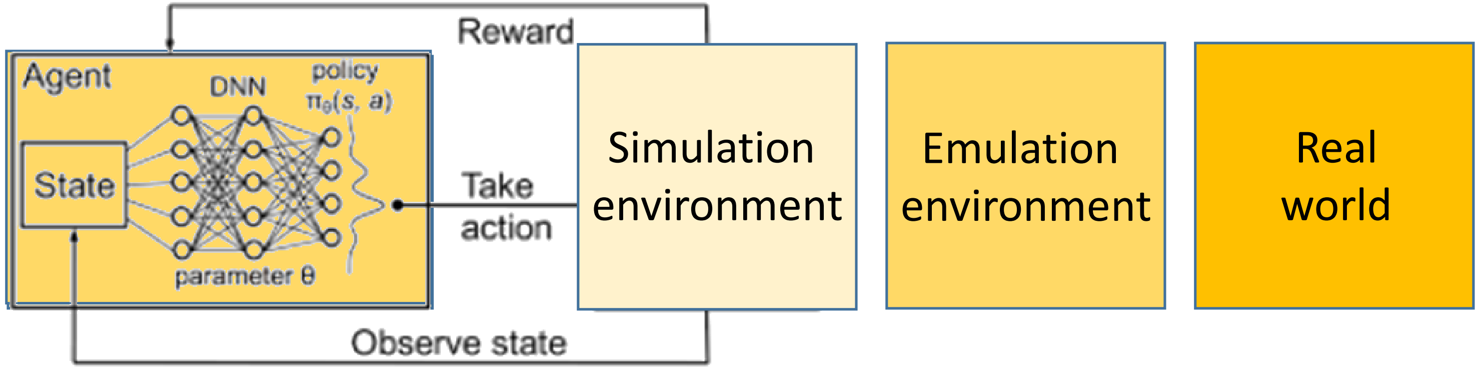

Speedup state exploration

Combine multiple environments

Simulation | Emulation | Real world

Example #3 WLAN traffic optimization

Campus *Insight*

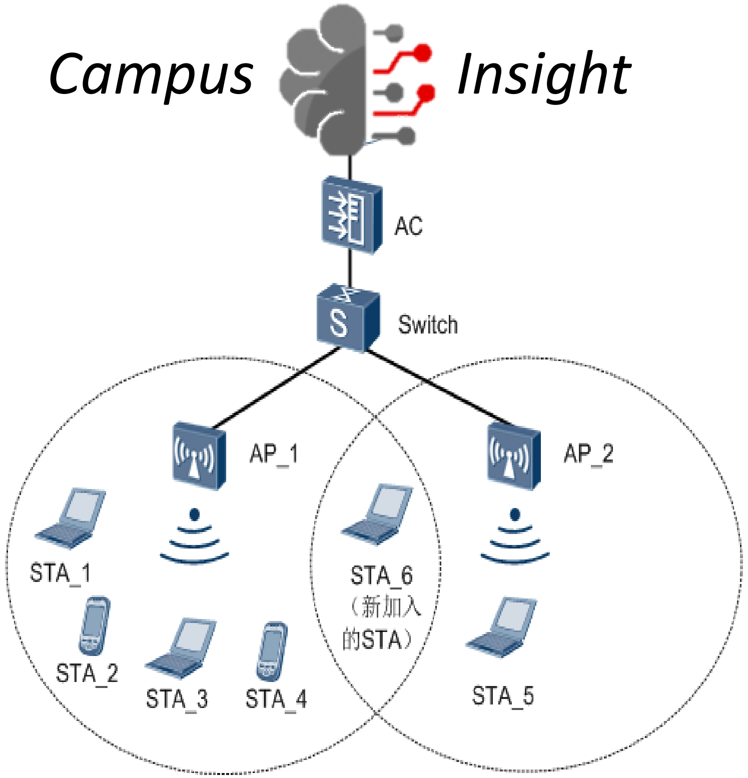

(Deep) reinforcement learning

Reward= f(T, Δ, QoE, I, RSSI, …)

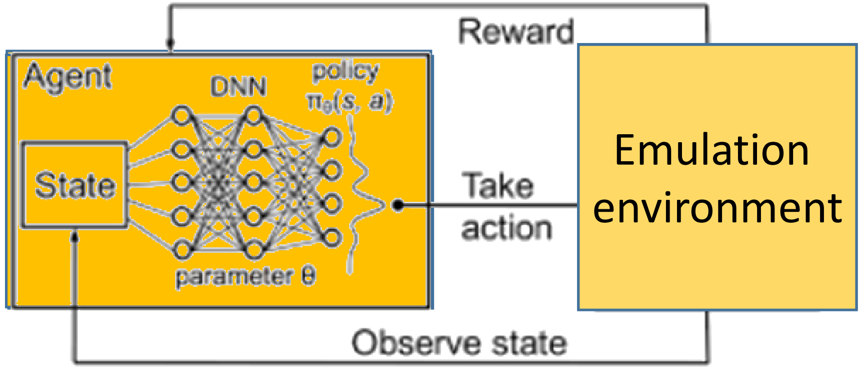

Real
world

Speedup state exploration

Combine multiple environments

Simulation | Emulation | Real world

AI-powered networks

Keep humans in the (slow) loop,
facilitate their interaction with AI

Control
the network

- ❑ Closing the loop
- ❑ **Humans & the loop**
- ❑ System aspects

QoE driven network management

In most cases, *users* in the end-to-end loop

- ❑ Must avoid humans in the *fast* loop (else it breaks the autonomic paradigm)
- ❑ Useful to keep humans in the *slow* loop (e.g. involve end-users to ensure AI controlled networks works better than before!)

Human-resilient AI

In most cases, *human operators* will not have a clue (or anyway will not be experts) of AI technologies

- ❑ AI should be resilient in spite of poor/adversarial training, bad calibration, overfitting, unfairness, …
- ❑ Artificial intelligence must use techniques to be robust and survive in spite of human stupidity….

Example #4 Web QoE

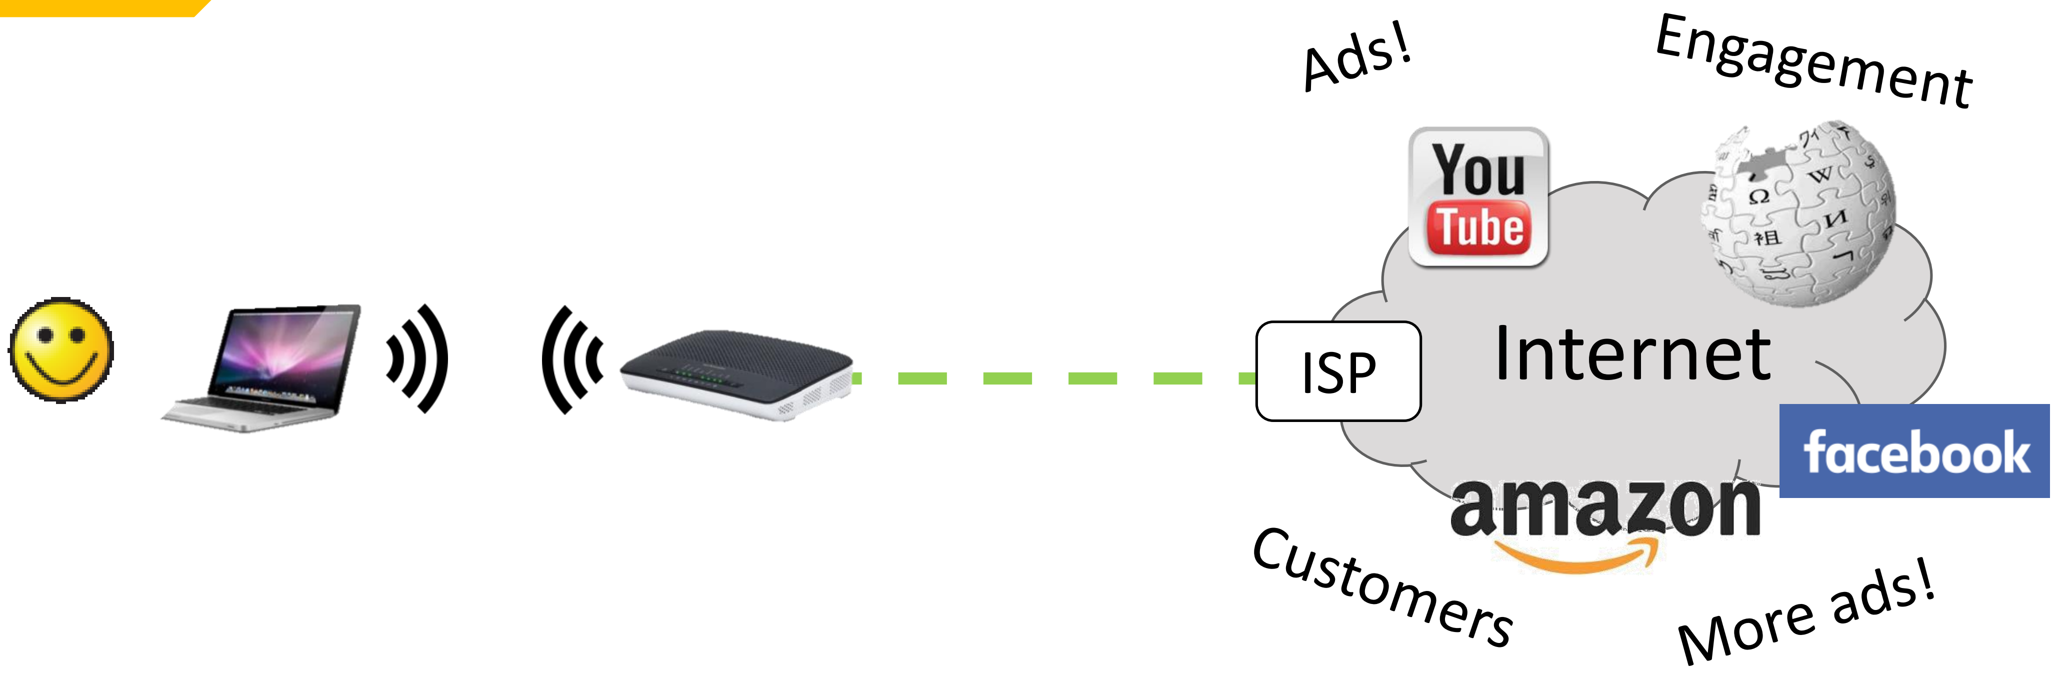

Offering G😊😊d user Q😊E is a common goal

Example #4 Web Quality of Experience

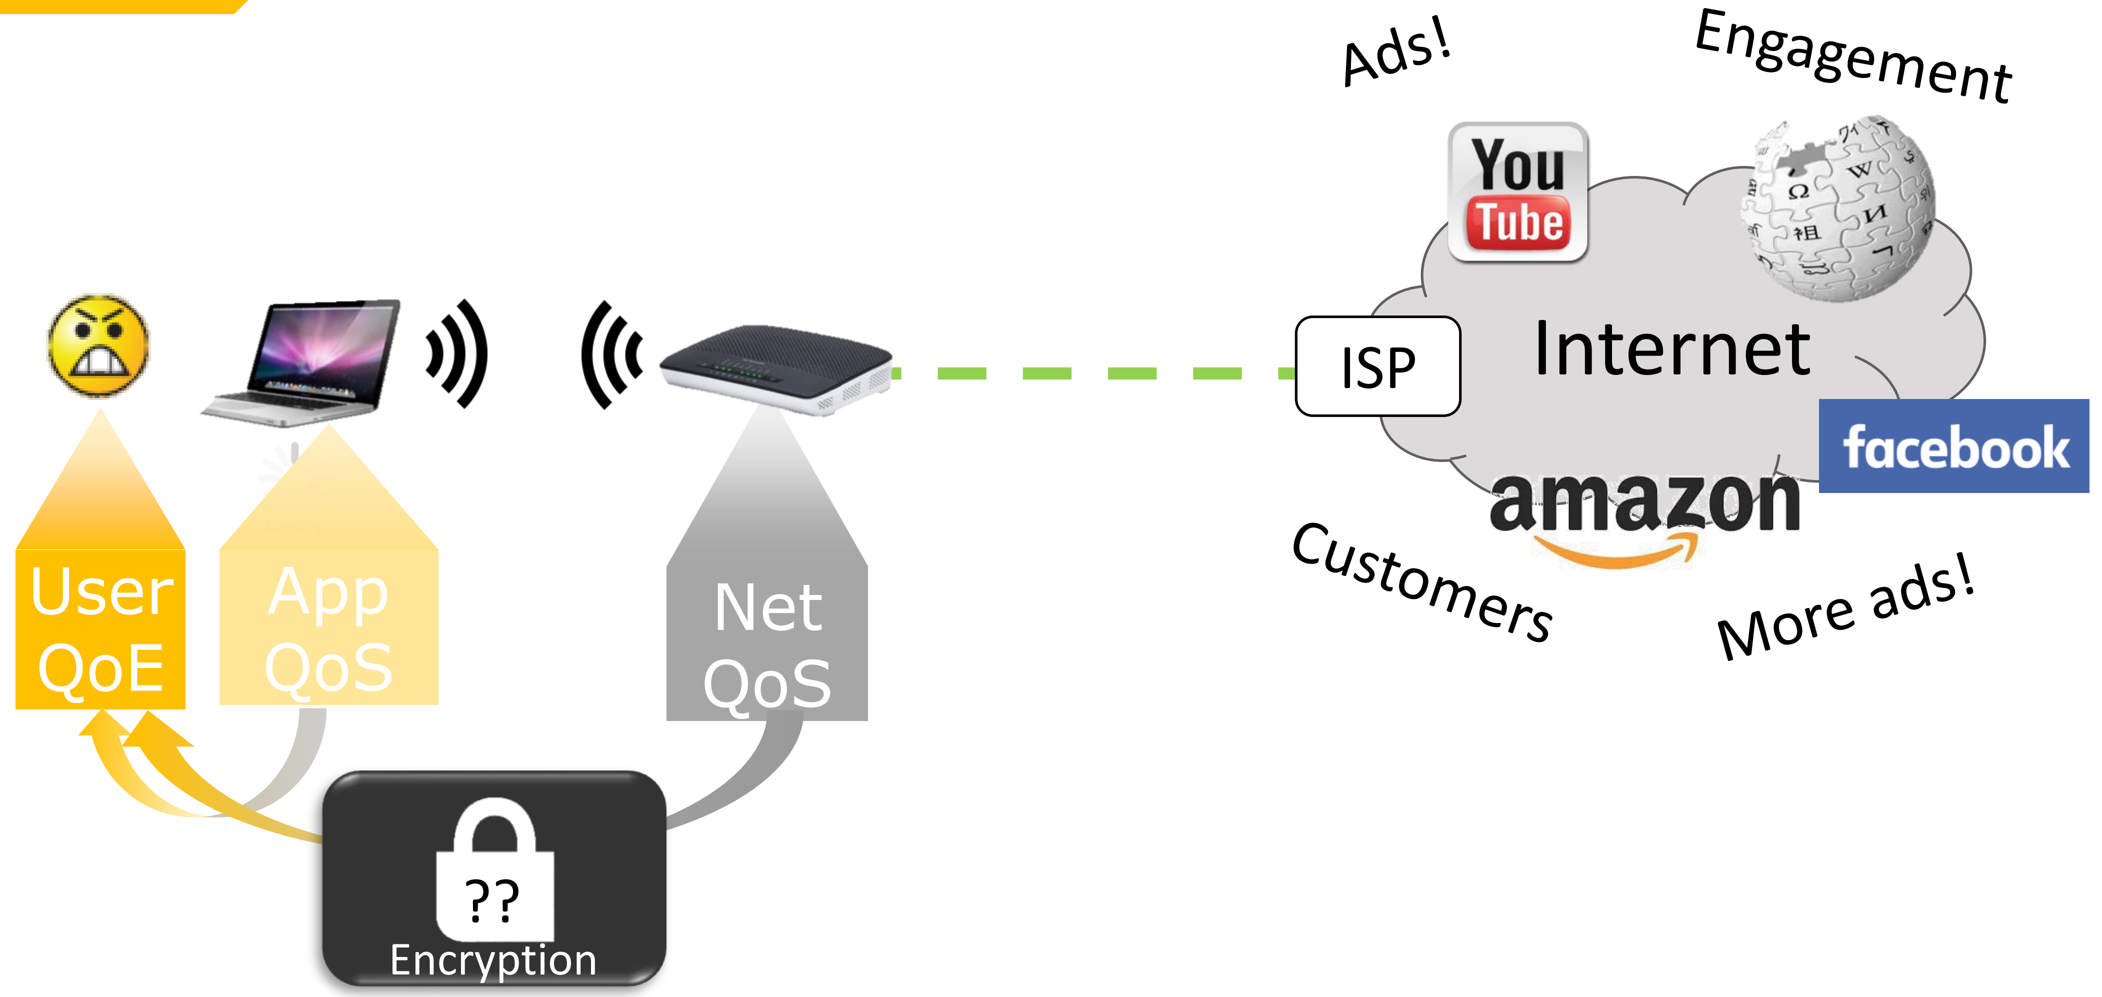

Detecting/preventing user QoE degradation is important!

Example #4 Web QoE

Play

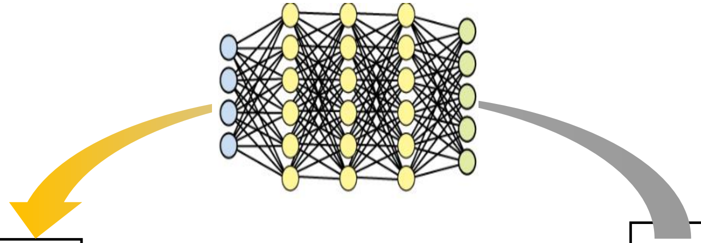

Webpage rendering

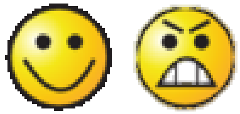

User

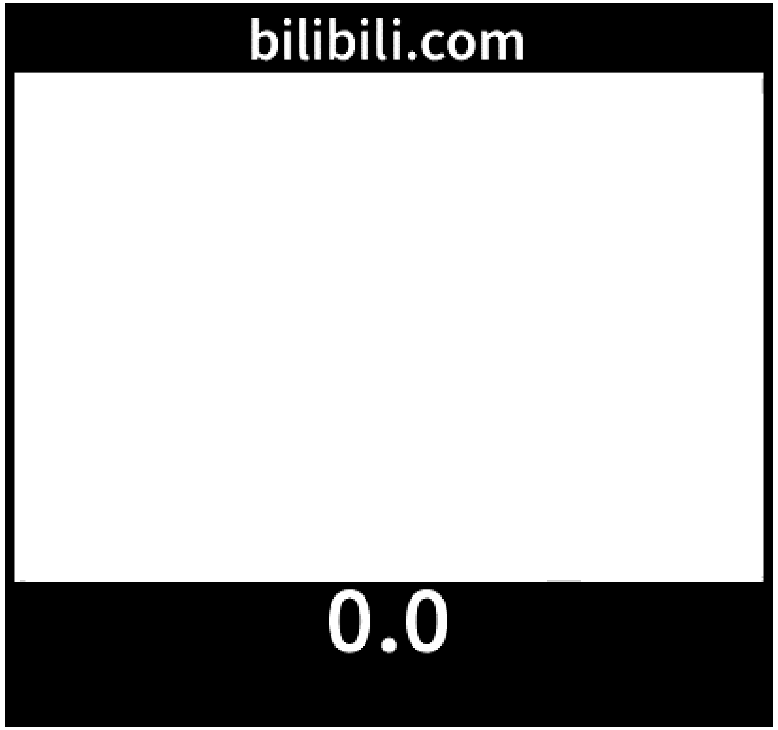

1 burst = 1 object
1 color = 1 domain

Brows

1 burst = 1 packet
1 color = 1 IP server

L3)

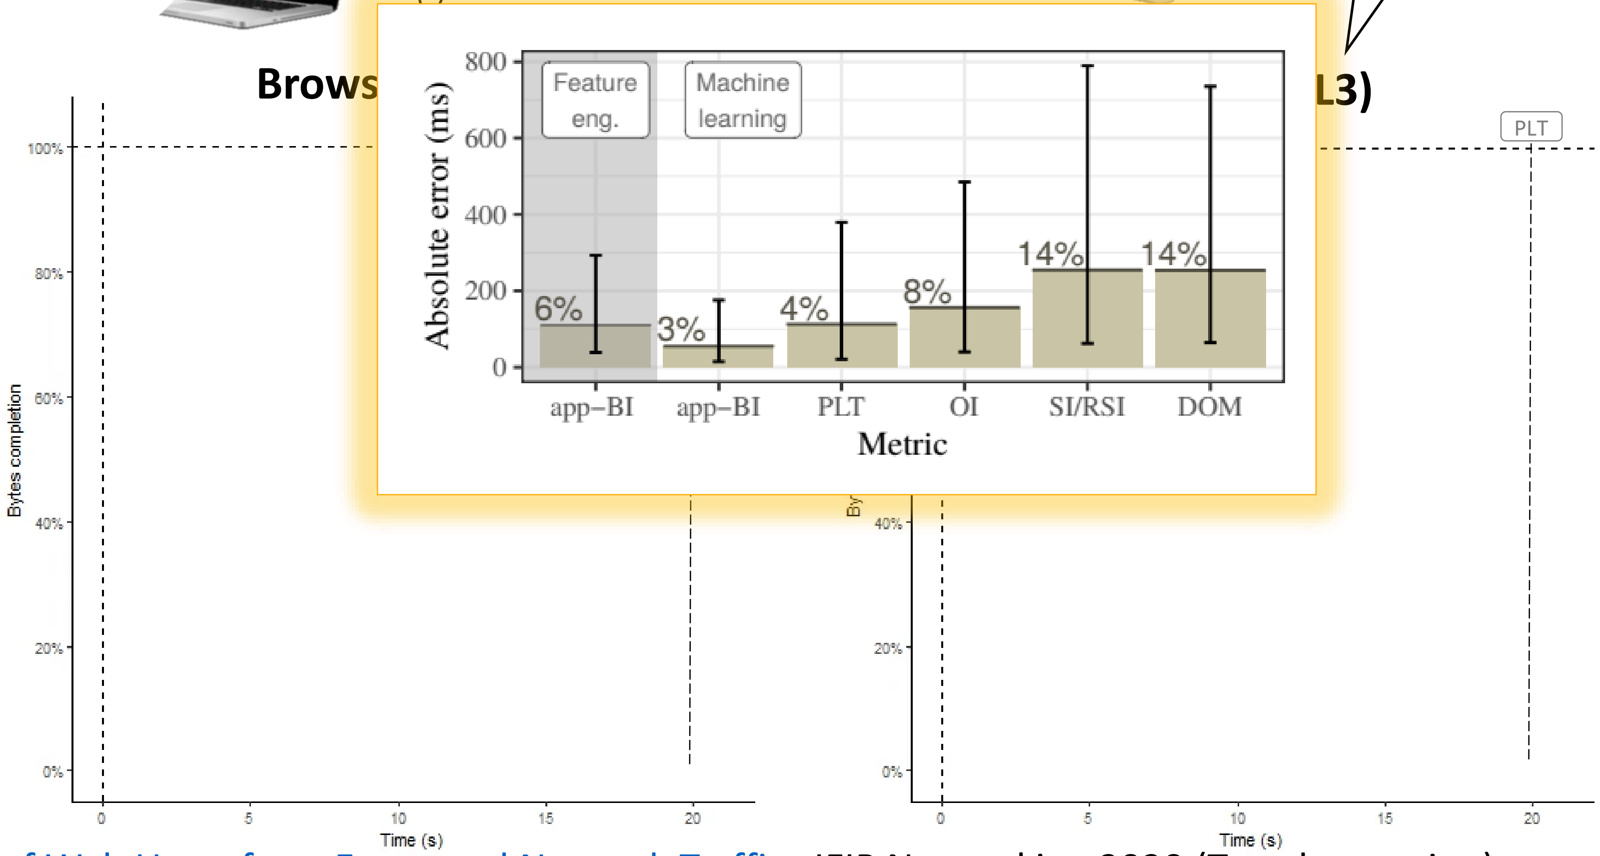

[Networking 2020] A. Huet et al. [Revealing QoE of Web Users from Encrypted Network Traffic](#) IFIP Networking 2020 (Tuesday session)

Example #4 Game QoE

We add network latency

We record AIs score

We let trained AIs play

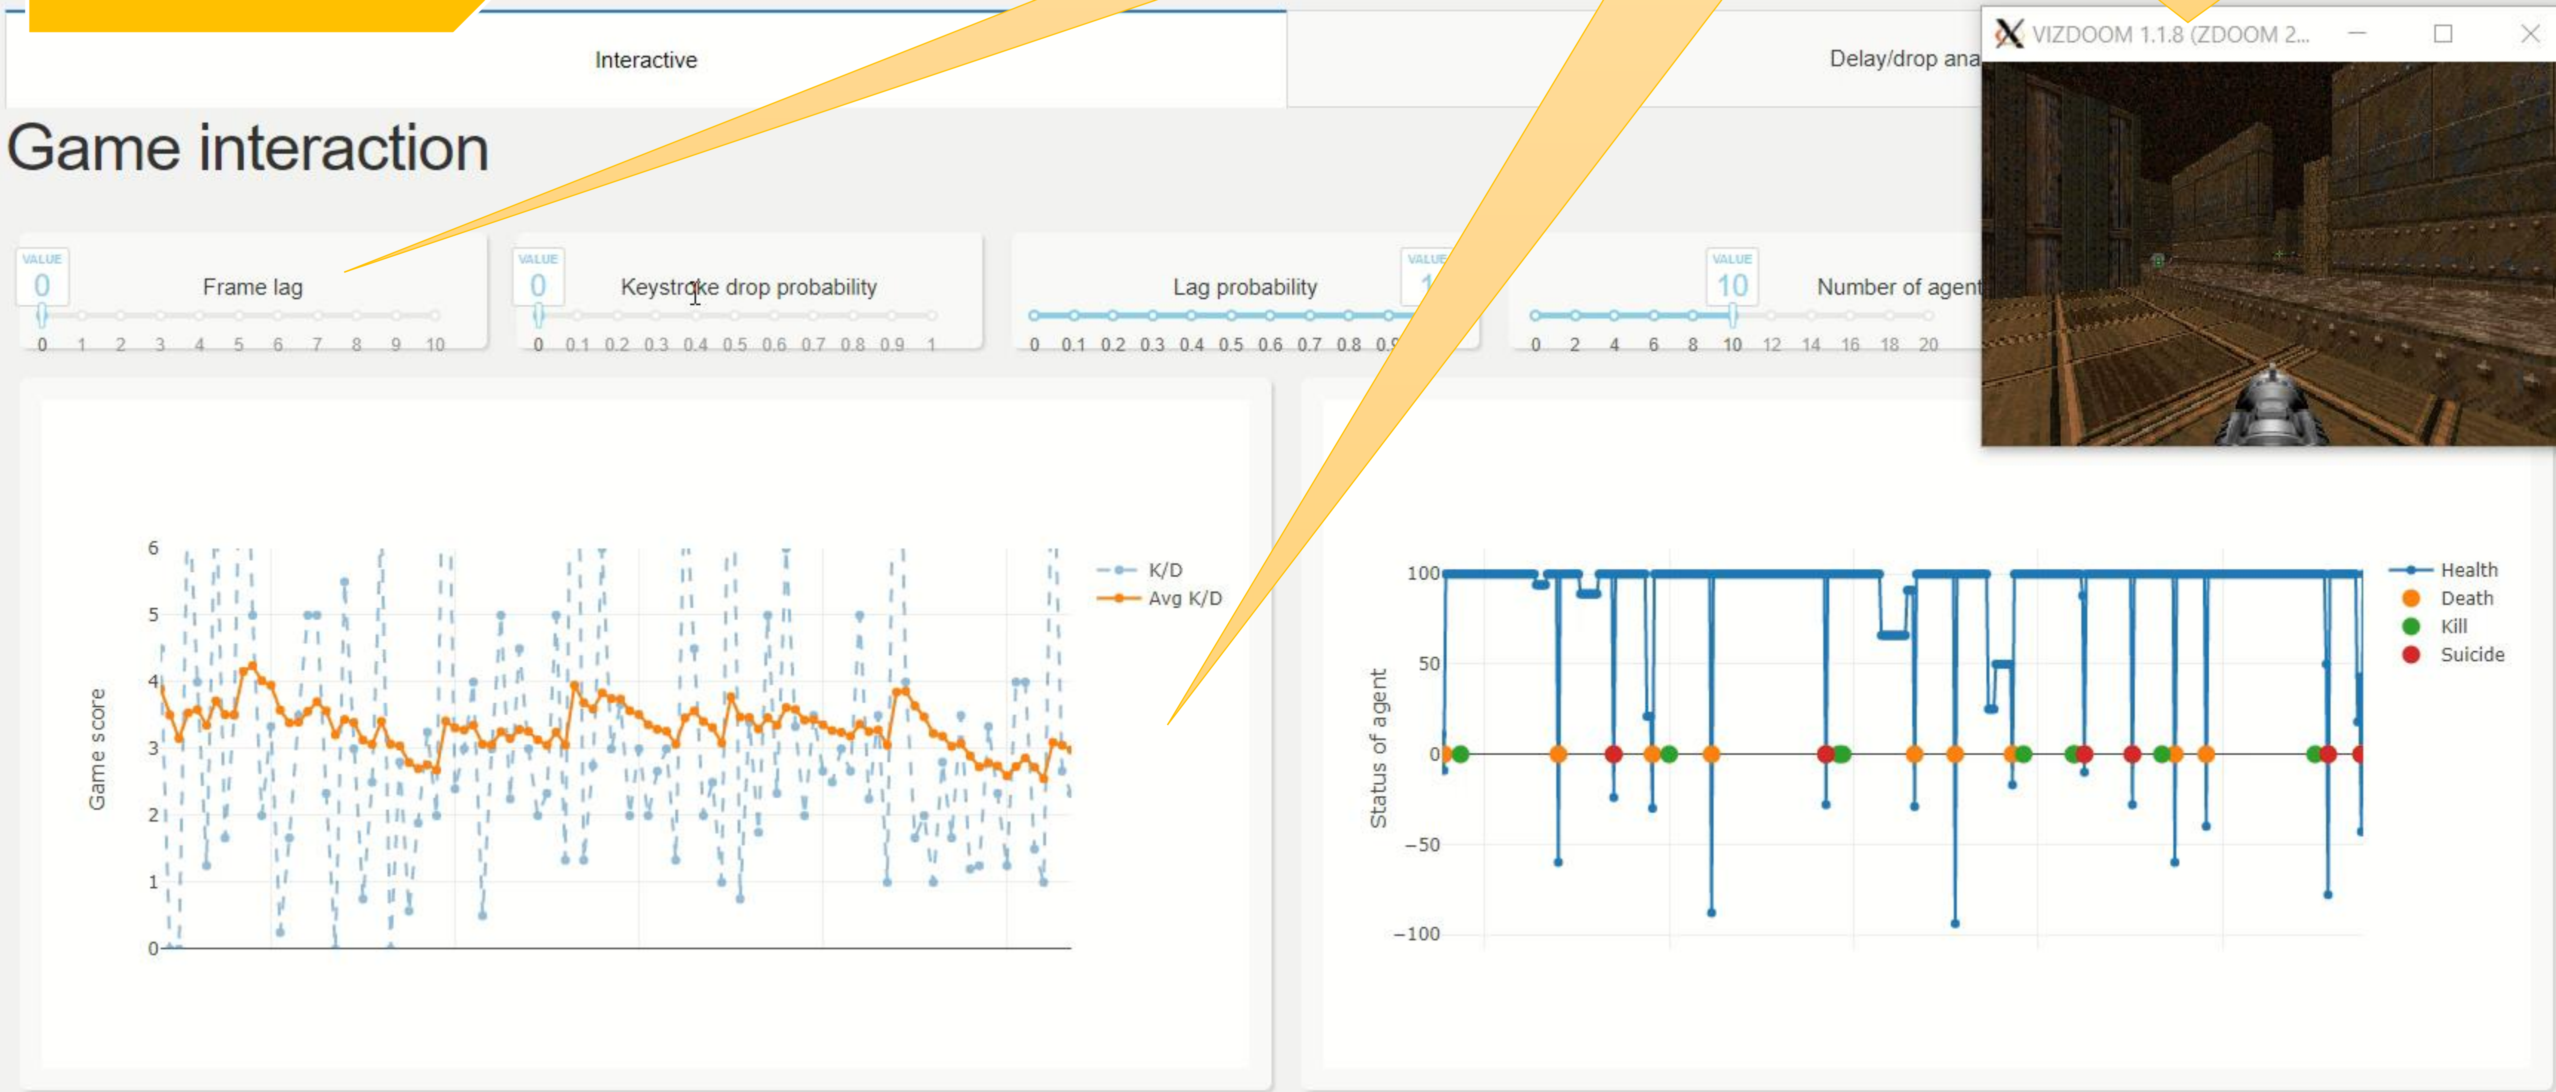

[INFOCOM'20] G. Sviridov et al., [Removing human players from the loop: AI-assisted assessment of Gaming QoE](#) IEEE INFOCOM Workshop + Demo

AI-powered networks

Statistical approach not a silver bullet. AI resource allocation !

Control the network

- ❑ Closing the loop
- ❑ Humans & the loop
- ❑ System aspects

Need for deterministic algorithms

- ❑ Machine learning is not a *silver bullet:*
  - ML accuracy 99.9% (dream model) 100,000 configuration lines = 100 errors
  - Ops, the problem just got a worse nightmare
- ❑ Autonomus configuration must use formal models for rigorous and deterministic guarantees

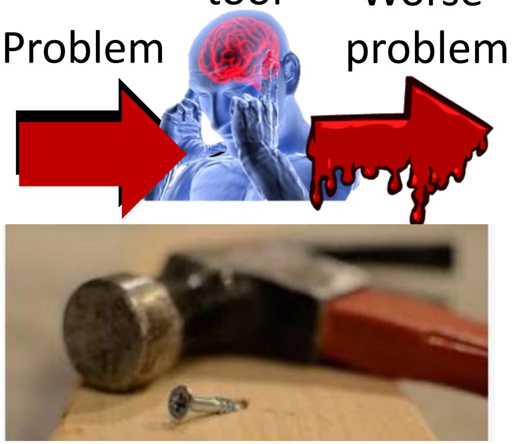

AI-resource allocation

AI powered chips extend the NFV resource allocation problem to a new dimension: the chip memory/processing resources!

- ❑ New tradeoff: chip memory/operations vs bandwidth
- ❑ New problems: how to split the in-network processing?

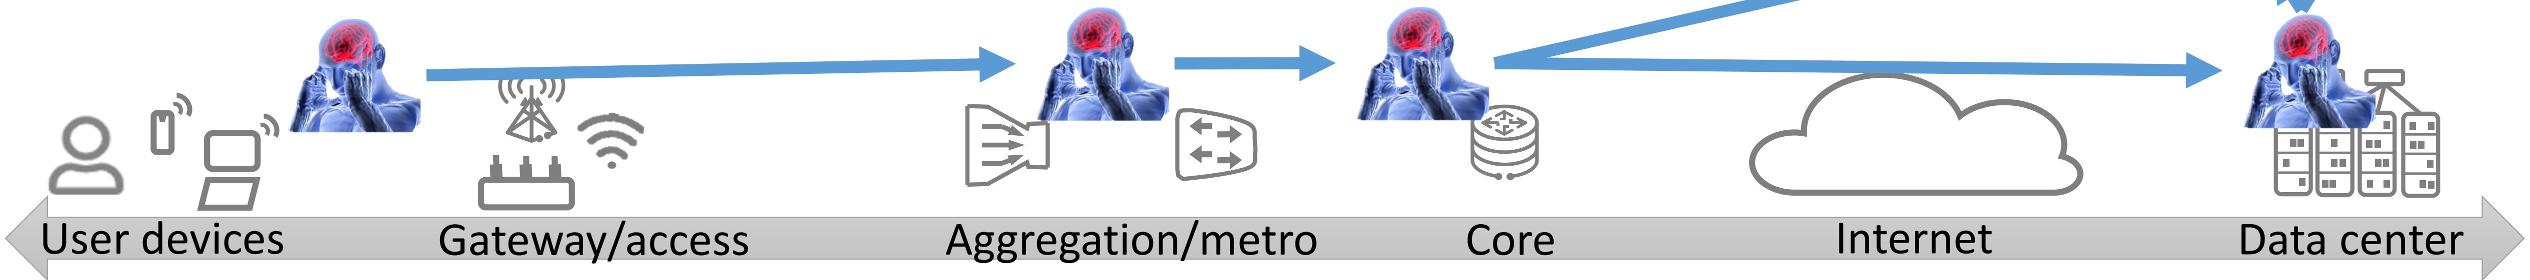

Takeway messages
for the road

Hardware advances

Recent hardware advances true enablers of “edge intelligence”

Network data

Heterogeneous, asynchronous, evolving unlabeled massive data

Understand the network

In ML, the journey matters more than the destination

Care about interpretability, not just performance as a black-box

Just as network protocols, ML can (& will) be hacked

Control the network

Keep humans in the (slow) loop, facilitate interaction with AI

When closing the loop, mind the gap!

Statistical approach not a silver bullet. AI resource allocation !

Good Practices

IO data pipeline essential for AI in products

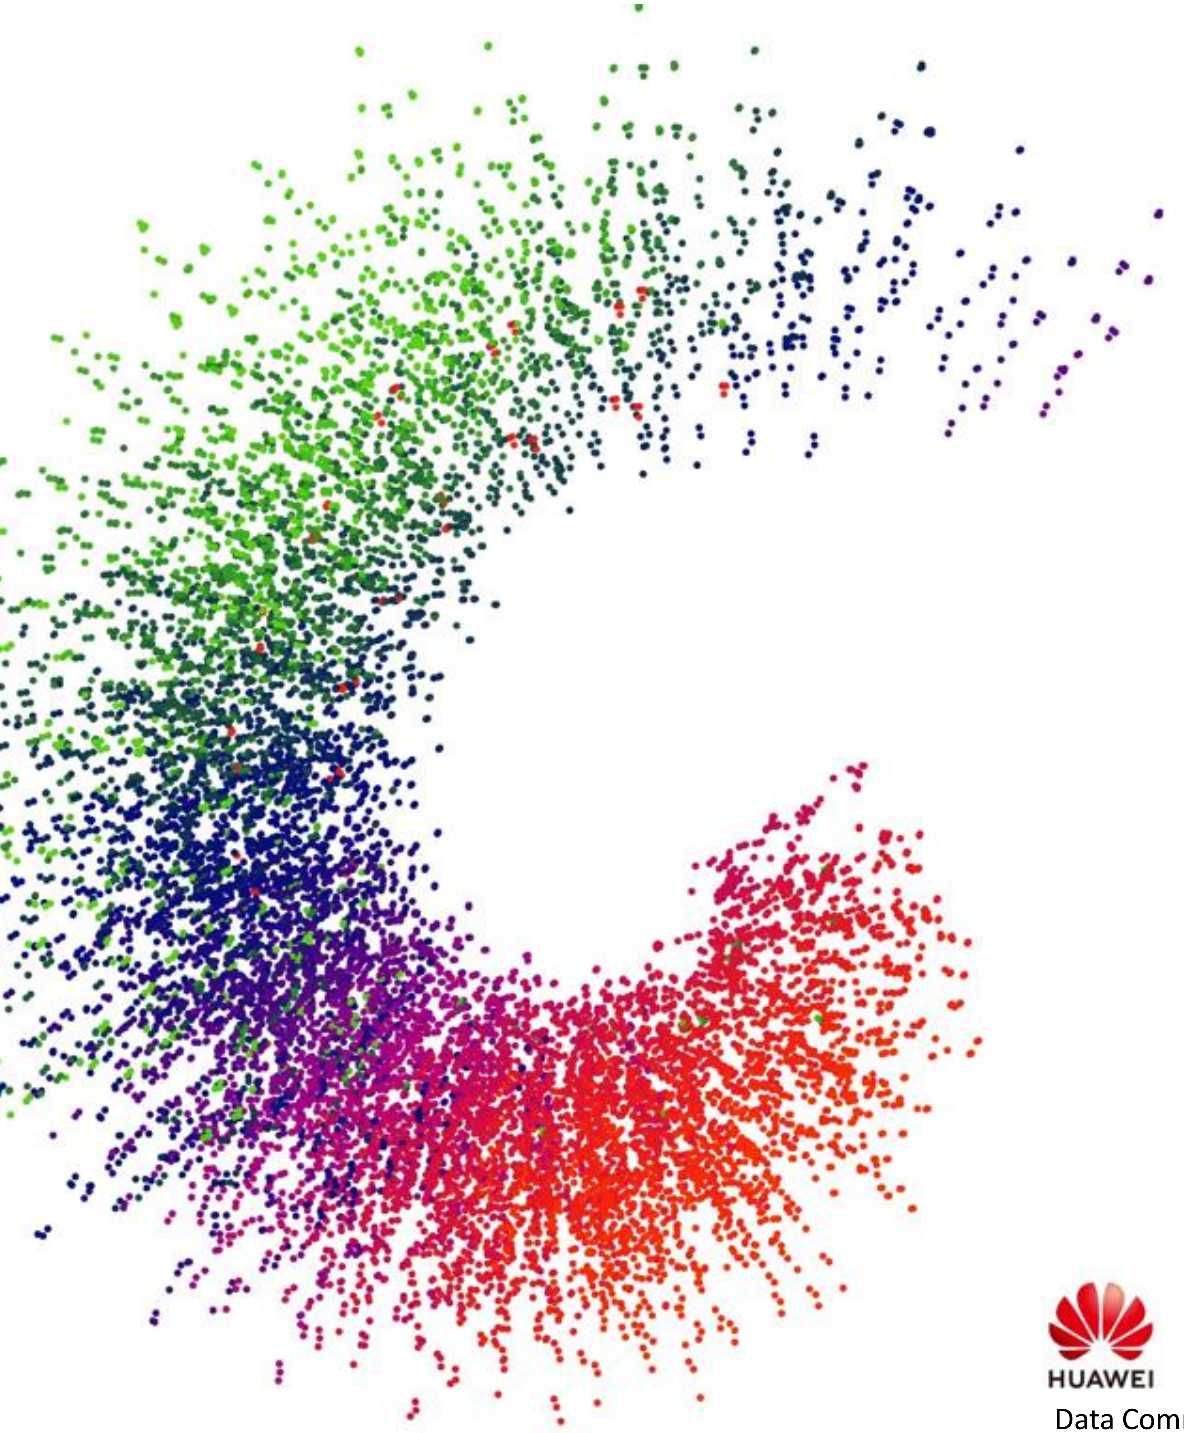

Thanks

Dario Rossi

dario.rossi@huawei.com

https://nonsns.github.io

Chief Expert, Network AI

Director, DataCom Paris Lab

Data Communication Network Algorithm and Measurement Technology Laboratory