AUSTRIAN INSTITUTE OF TECHNOLOGY

Photonic Information Processing

All set for maximizing efficiency?

Photonics 4

Sustainable Environment

May 18 2021

Dr. Bernhard Schrenk

AIT Austrian Institute of Technology

erc

Today‘s Presentation

Addressing signal processing,
as a main challenge in ICT:

data generation and processing shows a CAGR of 60%

collect – fuse – process everywhere

Photonic **Information** **Processing**

Mimicking biological information processing – without the need for virtualizing real-world information structures at high energy detriment.

- Accelerators for HPC
- AI hardware

- Enable high information rates
- Enable an ultra-low time-of-flight latency
- Lower the required energy and break the energy brick-wall of microprocessors

The Internet Explosion

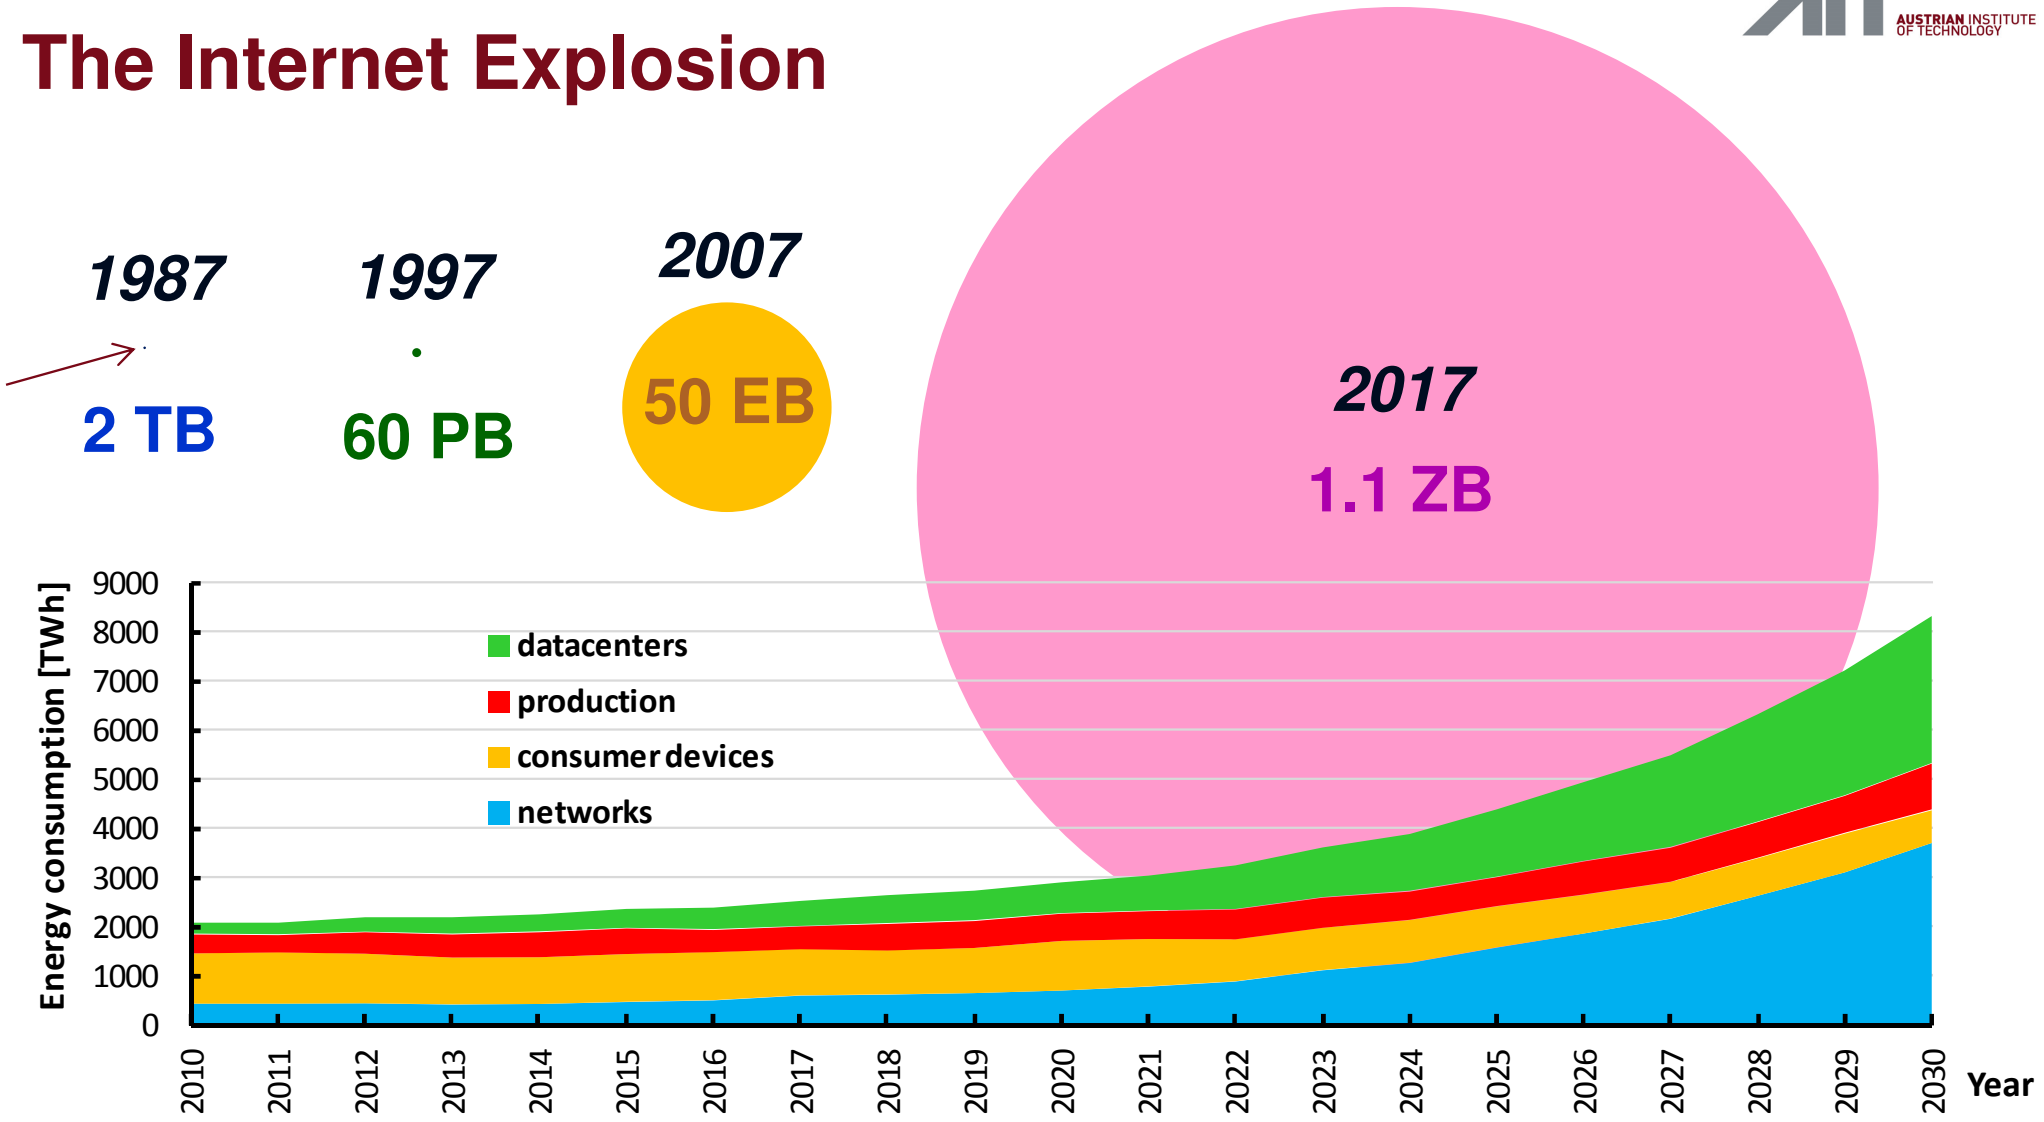

N. Jones, Nature 561, 163 (2018)

Share of ICT in Global Energy Consumption

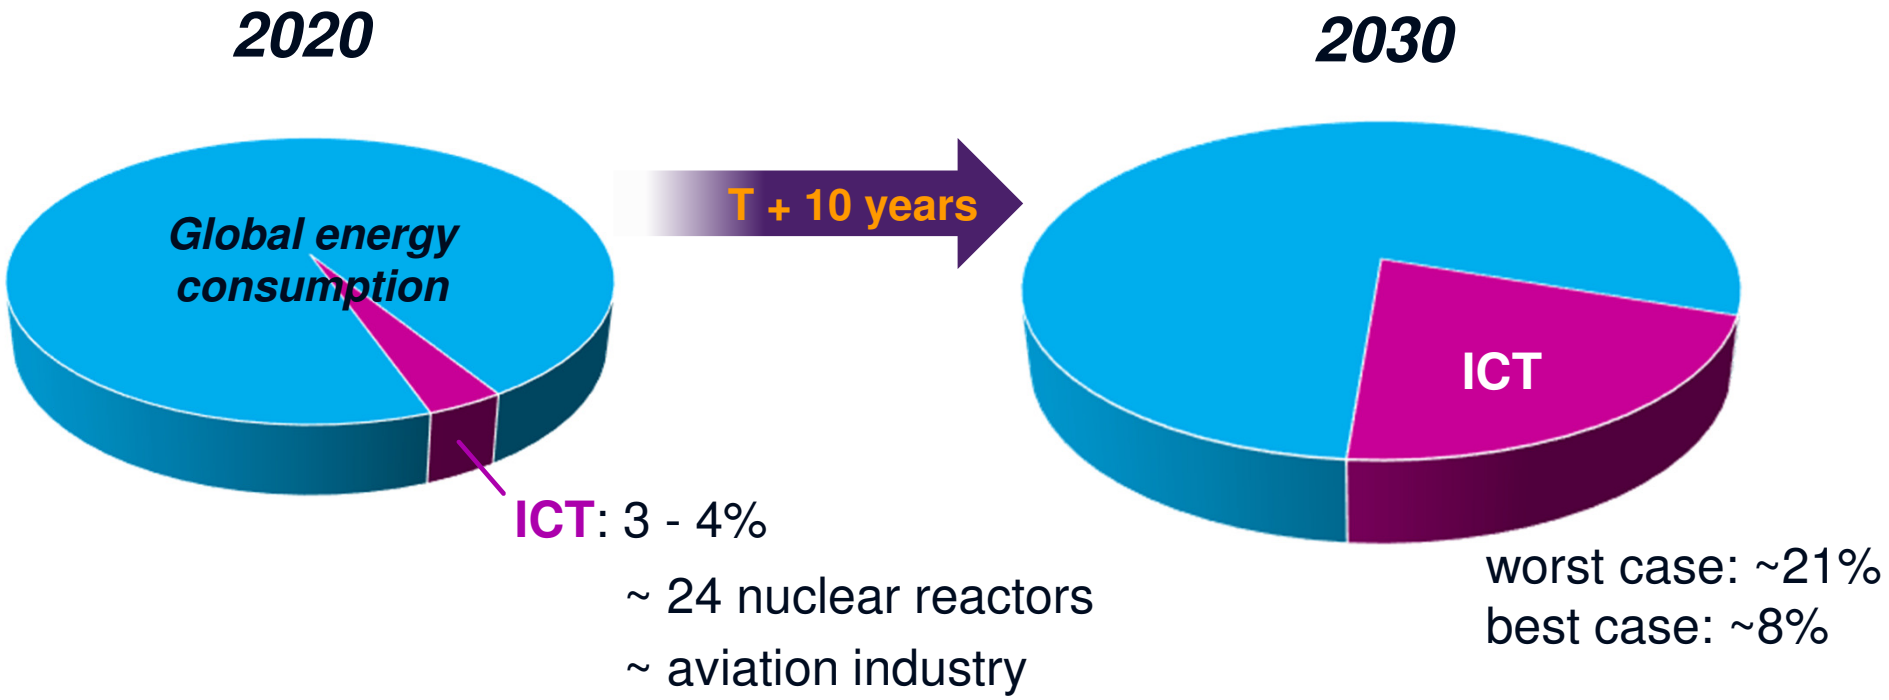

N. Jones, Nature 561, 163 (2018)

Bitcoin Mining

- 80% happening in China
  - 40% thereof fuelled by coal
- equivalent to 0.6% of world's electricity production
- similar footprint as Italy or Saudi Arabia

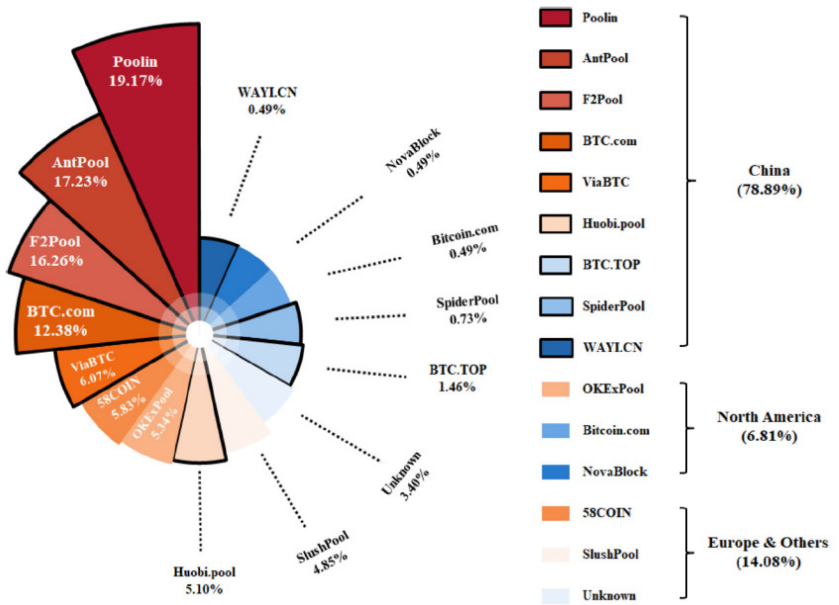

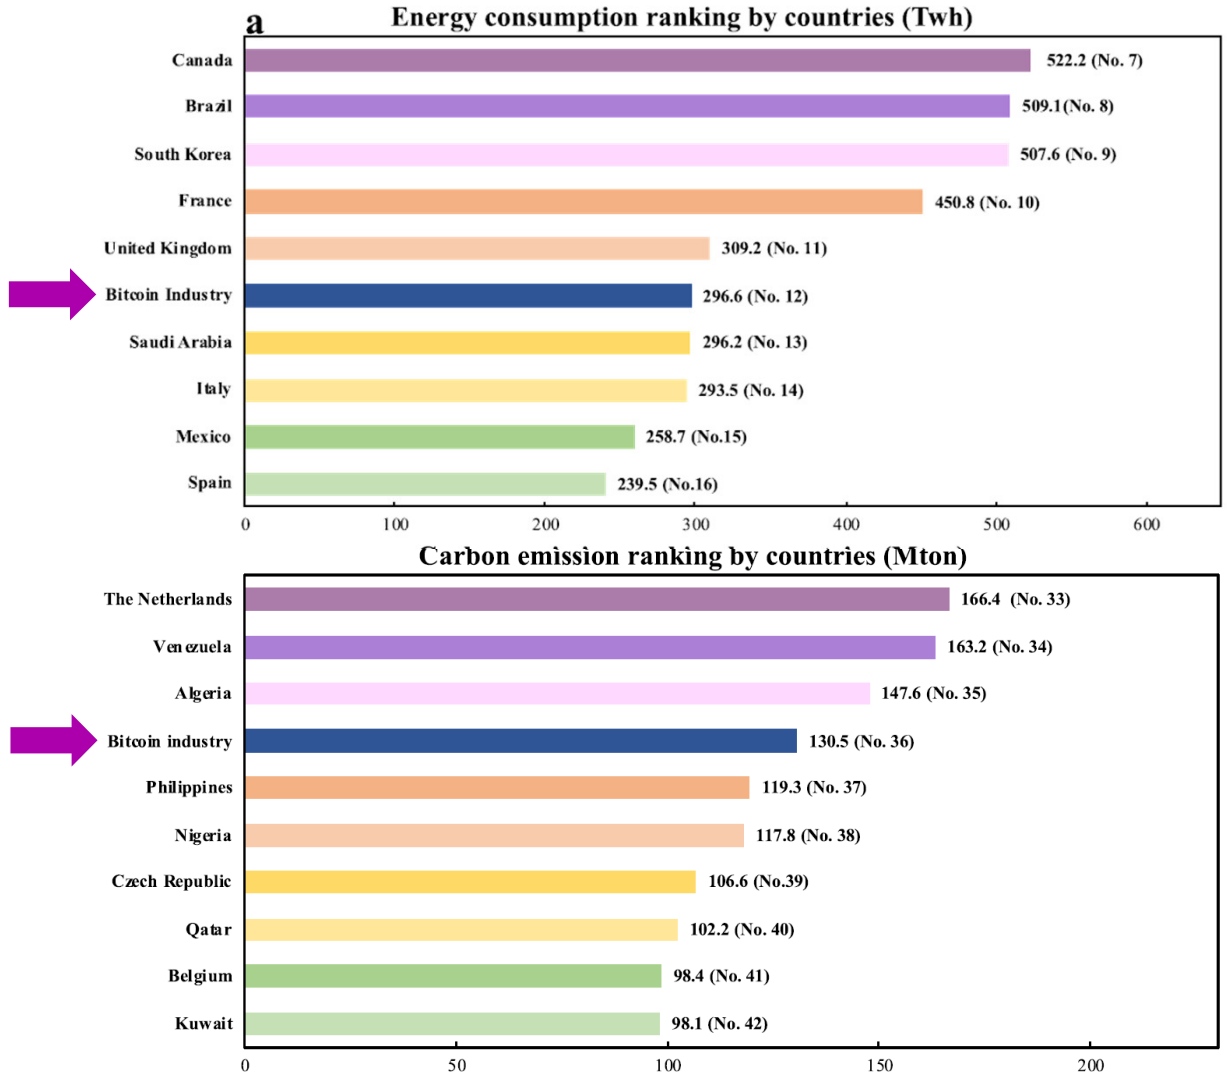

S. Jiang, Nature Comm. 12, 1938 (2021)

Inside the Information Factory

Cloud Datacenter

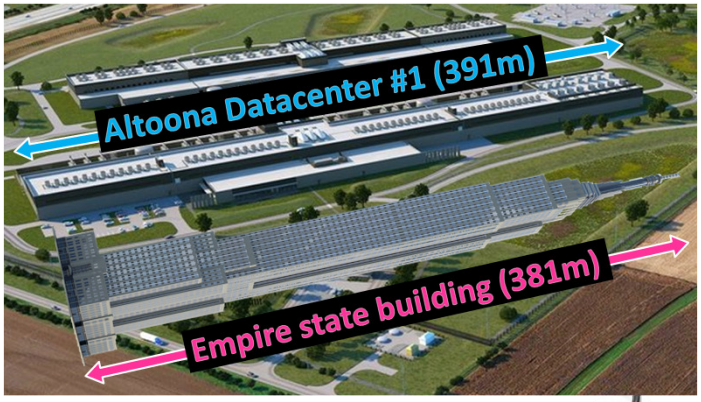

~20 000 servers
~20 MW

HPC

537 212 TFlop/s
29.9 MW

Human

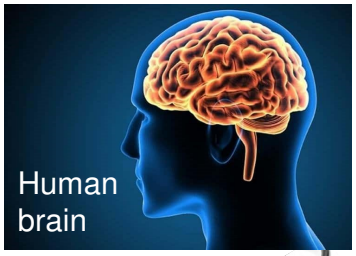

~2 000 TFlop/s
0.000020 MW

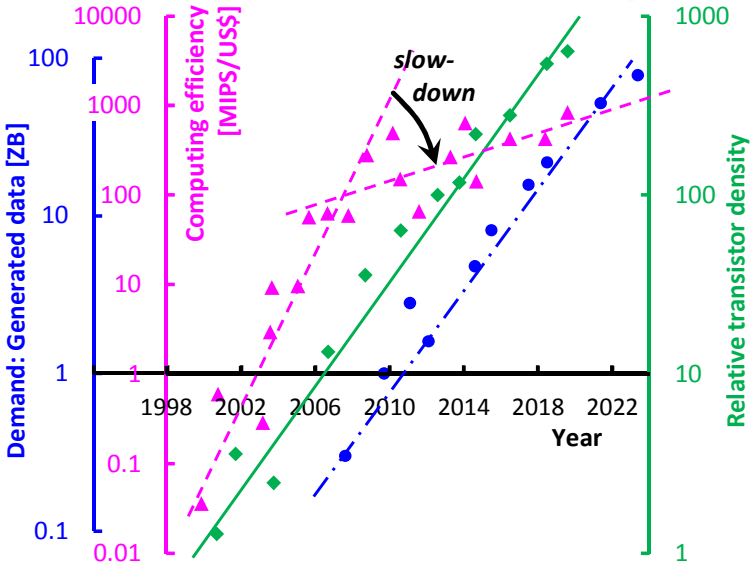

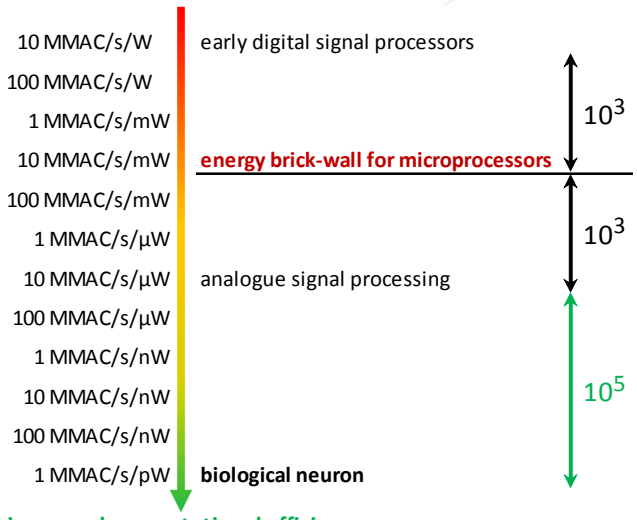

J. Kendall, Appl. Phys. Rev. 7, 011305 (2020)

The Rise of AI

- Natural language processing
  - Siri, Alexa
- Game playing (Go, chess)
  - 2016:
    Lee Sedol vs. AlphaGo
    1 - 4
- Face recognition and situation classification
- Autonomous vehicles (and UAVs)
- Control and optimization

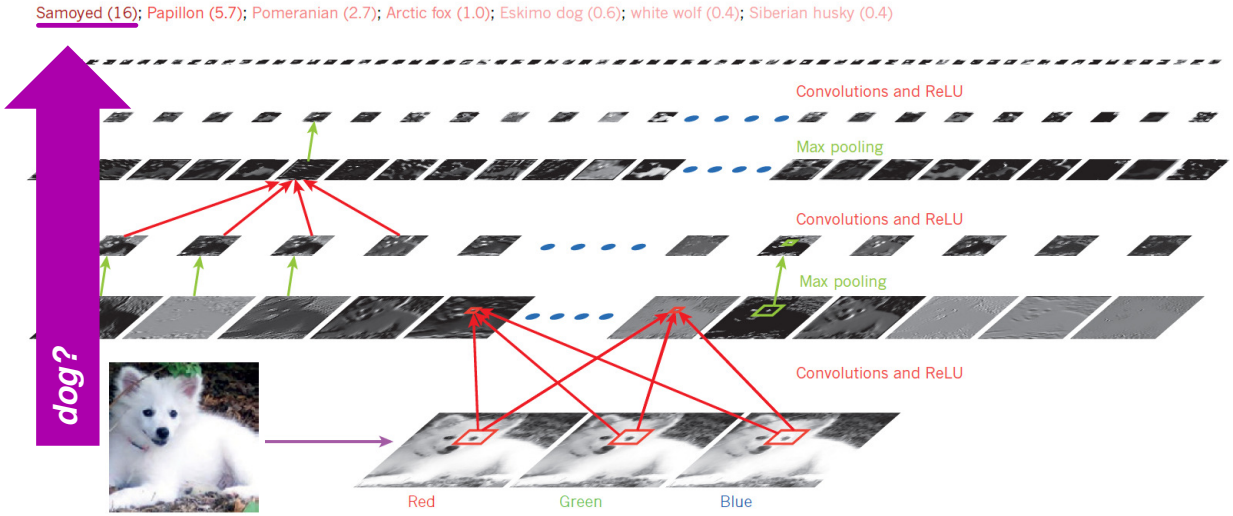


A woman is throwing a **frisbee** in a park.


A **dog** is standing on a hardwood floor.


A little **girl** sitting on a bed with a teddy bear.


A group of **people** sitting on a boat in the water.

Y. LeCun, Nature 521, 436 (2015)
M. Waldrop, PNAS 116, 1047 (2019)

Pattern Recognition

Msot plpeoe wlil hvae no peormbls radneig tihs txet, alothguh the oderr of leterts is rndaom (wtih the epeixoctn of the frist and the lsat leettr).

- There is only 1 correct solution and ~ 121 885 070 000 000 000 000 000 possibilities.
- We compute on-the-fly as we read over the text – a fantastic example of pattern recognition.
- Hardly any child could do this.

Artificial Intelligence

- Multi-layered, deep **neural network**
  - accomodates many **neurons**
    Human brain: 10^{11}
    Intel Loihi: 130,000
- Weighted synaptic **interconnect**
  - dense vector-matrix multiplications
  - routing becomes challenging when scaling up the data movement
    Human brain: 10^4 inputs/neuron
- Each layer needs to be **trained** …
- … to yield **time-of-flight inference**

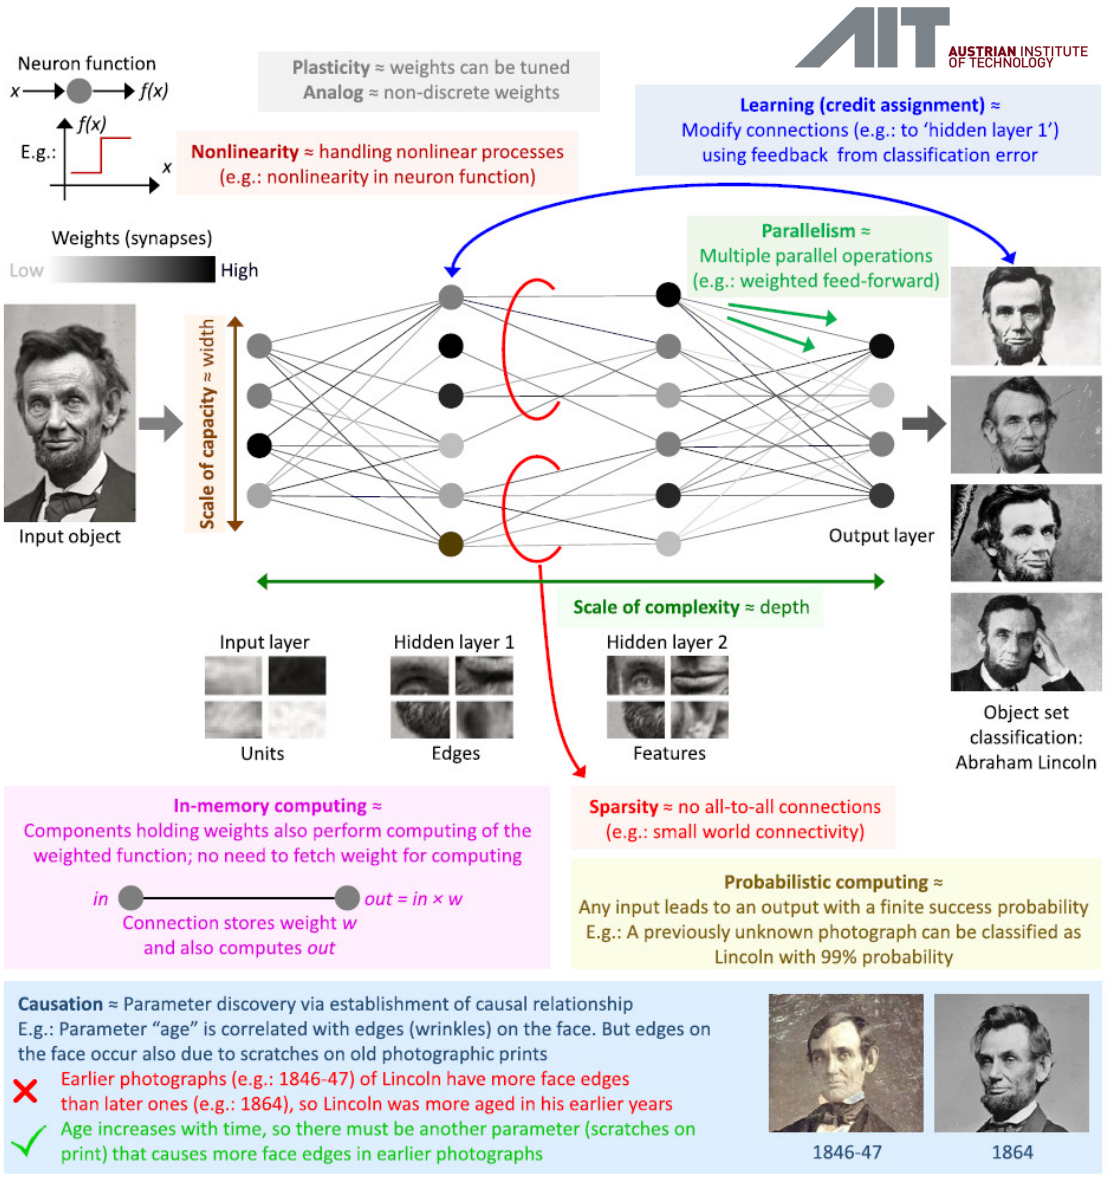

J. Kendall, Appl. Phys. Rev. 7, 011305 (2020)

The Cost of Training AI

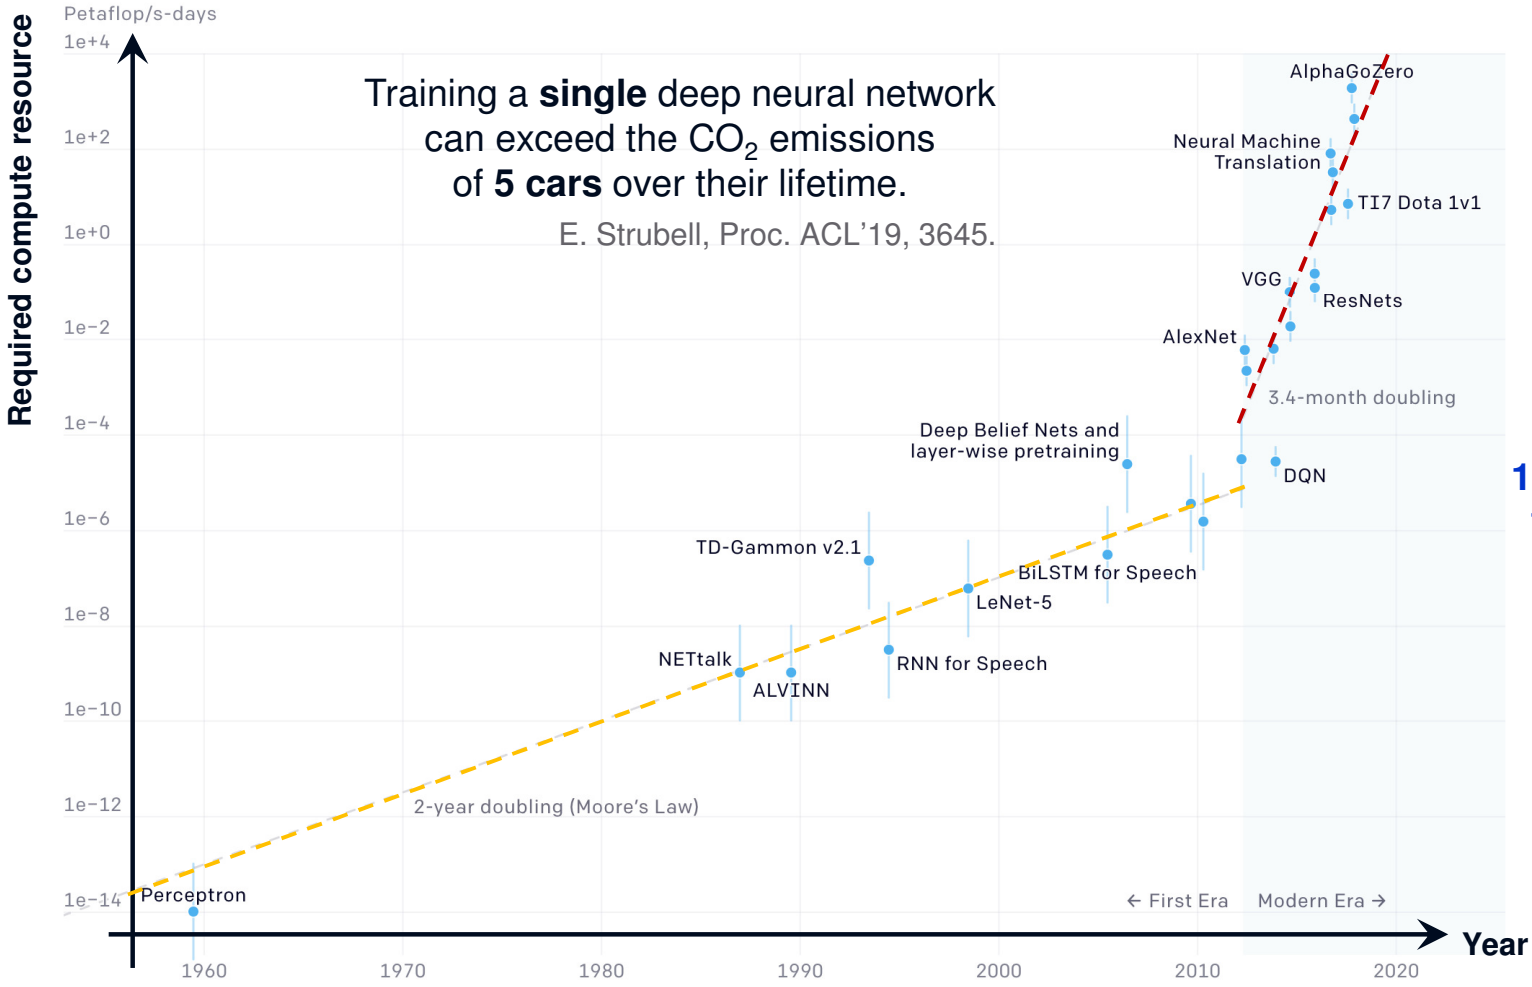

Go 2016:
AlphaGo vs Lee Sedol

1202 CPUs
176 GPUs vs **1 human brain**

1 MW vs **20 W**

Source: https://openai.com/blog/ai-and-compute/

The Neuron Goes "Light"

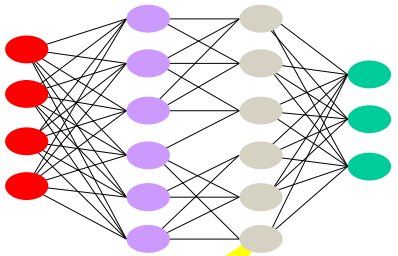

- Neuron model: Leaky integrate-and-fire
  - Linear operations:
    - (1) receive multiple inputs
    - (2) weigh each input
    - (3) sum all inputs

$$\sum_k v^{(k)} w_k$$

$$w_k \in [-1, +1]$$

multiply & accumulate

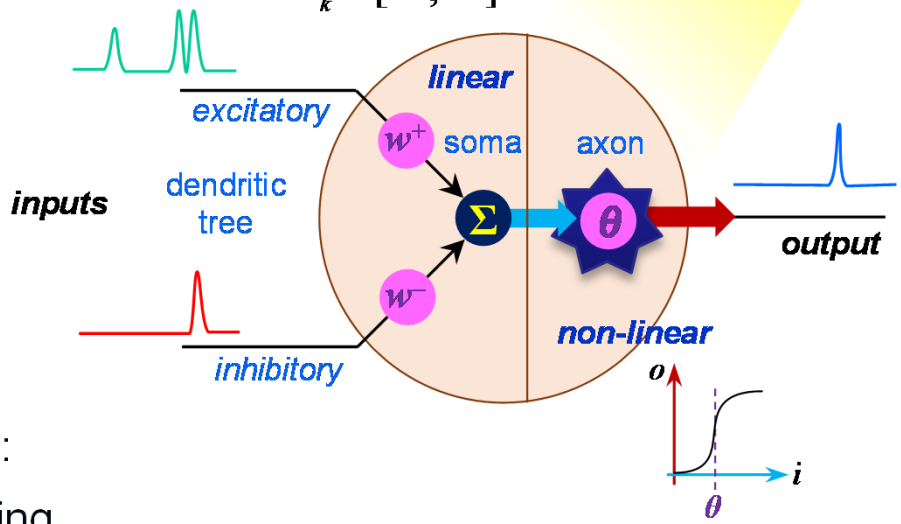

  - Non-linear operations:
    - (4) perform thresholding

Why photonics?

- Synaptic interconnect
  - can exploit wavelengths at fan-in and fan-out:

    single bus hosts a plethora signals
- Operation at high speed
  - spike train at 1 GHz with 100-ps spikes
  - becomes energy efficient

How can Photonics contribute?

Is there a Moore's Law for PICs?

- …and does it scale according to application needs?

e.g., predictive control for object at flight

- 24 neurons,

  convergence time of 10 ns

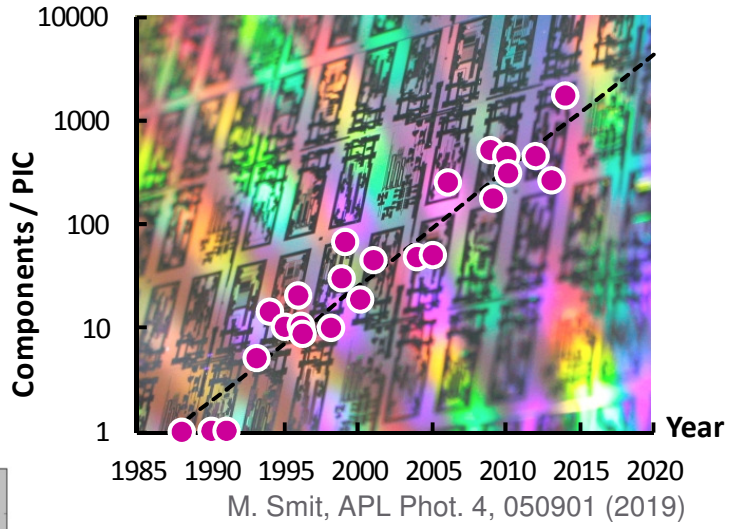

M. Smit, APL Phot. 4, 050901 (2019)

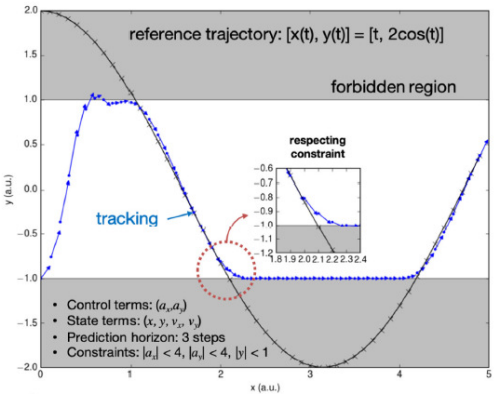

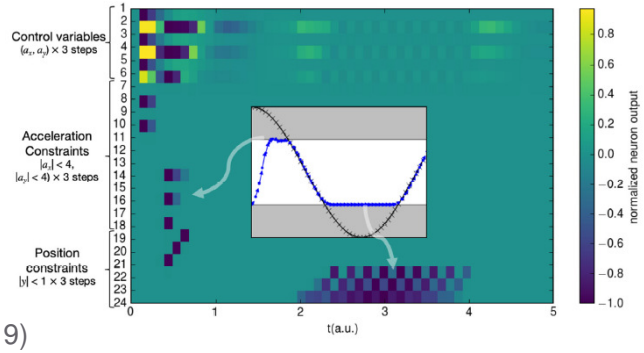

T. Ferreira de Lima, JLT 37, 1515 (2019)

What about energy efficiency?

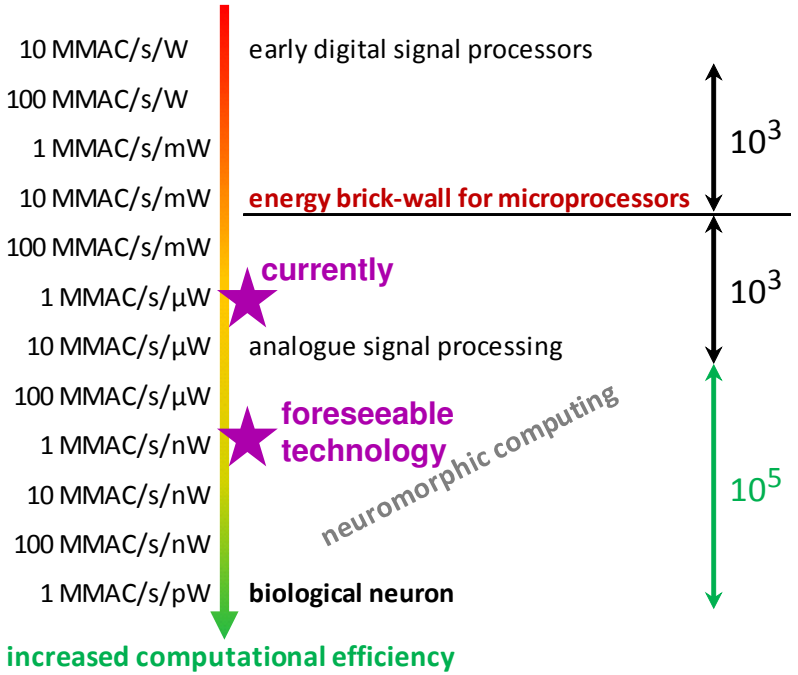

- foreseeable:

  1.1 fJ/MAC

K. Nozaki, OFC'20 (2020)

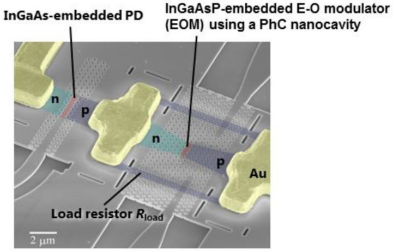

Reducing the CO_2 Footprint: Photonics at Power-Play

Information Transmission

- communication among people
- communication between machines

→ 2 pJ/bit optical
vs. 18 pJ/bit electrically wired
vs. ~1000 pJ/bit wireless

Information at Rest

- short-term caching
- long-term storage

→ SSD/HDD requires 30-60 nJ/bit per year
+ migration every 2-3 years

optical data memory: 1 TB / 100 cm^2,
written/read at Gb/s with 1 nJ/bit

→ energy / resource
reduction by factor 100

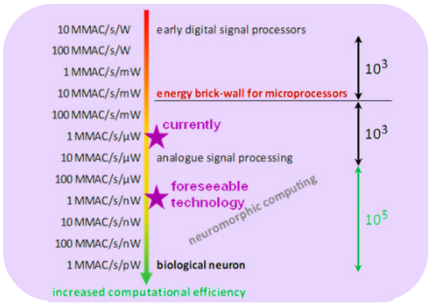

Information Transformation

- Exascale computing
- low-latency signal processing

"You can't stop the waves, but you can learn to surf."

Jon Kabat-Zinn

Dr. Bernhard Schrenk

bernhard.schrenk@ait.ac.at

Senior Scientist

Center for Digital Safety & Security

AIT Austrian Institute of Technology

Part of the work related to the presented results have received support by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 804769).