THE FOURTH INDUSTRIAL REVOLUTION, DEEP LEARNING AND ITS NEW COMPUTING PLATFORMS

Dr. Eng. Didier El Baz
CDA Team
LAAS-CNRS, Toulouse France

March 7, 2017

Outline

- 1. When?
- 2. What?
- 3. How?
- 4. Where?
- 5. CDA & AI
- 6. Conclusions and Perspectives

1. When?

- A new industrial revolution has started:
- 4th revolution.
- AI comeback: AI today → Autonomy.

Fig.1.1 Google car, USA

1.1 The first industrial revolution

- The Industrial revolution (1760 - around 1840);
- from hand production to use of machines and rise of factory system;
- the Age of Steam (1760 - around 1914).

Fig.1.2 Cugnot self-propelled vehicle, France 1770 (4 km/h)

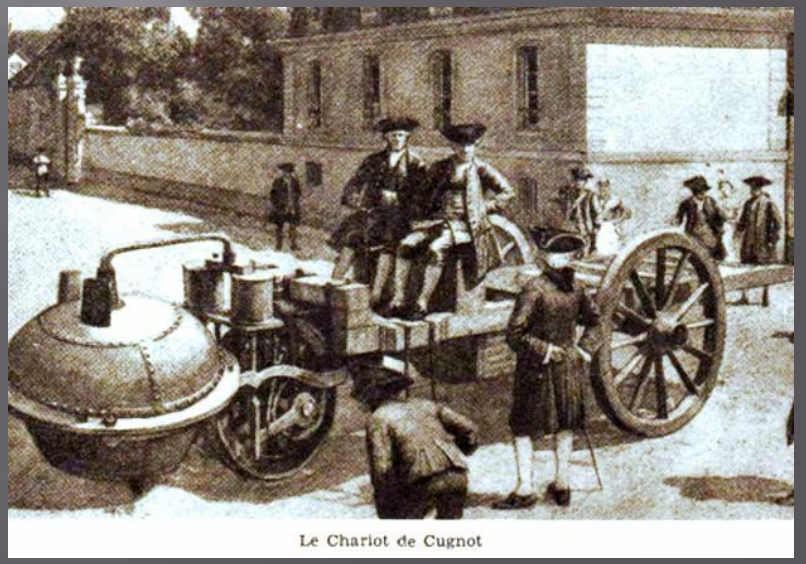

1.2 The second industrial revolution

- Electrification (1860 - 1950)
  - Factory, household, city & railway electrification.

Fig.1.3 Arc lamps at Avenue de l'Opera, Paris France in 1878.

1.3 The third industrial revolution

- Computer & communication age (1950 – 2010)
- Processors everywhere;
- Internet has changed the way people communicate and exchange data.
- Messaging applications

(WeChat, WhatsApp, …)

- The Digital Age (1960 -)

Fig.1.4 *5 MB* hard drive being loaded in an airplane by IBM, 1956

1.4 The fourth revolution

- Fusion of physical, digital world and the Internet (2010 -).
- Autonomy, collaboration, factory of the future.

Fig.1.5 Collaboration of man & robot AUDI, Ingolstadt, Germany

2. What?


中国乌镇
DEEP LEARNING APPROACH
Train:
Errors
Dog
Cat
Raccoon
Deploy:
Dog
WE WILL SHARE THIS KNOW-HOW
WITH ENTIRE WORLD

2.1 Amazing AI

Fig. 2.1 The next Rembrandt
TU Delft, Microsoft, ING, Mauritshuis 2016.
Cloud platform: Microsoft Azure VM
Parallelism (up to 1000 servers).

2.1 Amazing AI

Fig. 2.1 The next Rembrandt

Limited interest

2.1 Amazing AI

Fig. 2.1 The next Rembrandt

Fig. 2.2 The next Hollande

2.1 Amazing AI

Fig. 2.1 The next Rembrandt

Fig. 2.2 The next Hollande Deepfake

2.1 Amazing AI

Fig. 2.1 The next Rembrandt

Fig. 2.2 The next

Fig. 2.3 The next GO game champion

2.2 AI for sciences

- Counting Adélie penguin via AI (751,527 pairs)
- Deep Neural Network (DetectNet) to analyze photo collage & counting penguin nests.

Fig. 2.4 Penguin nests in Danger Island, Antartica February 2018

2.2 AI for sciences

- Looking for planet 9 via Data Mining, AI, HPC.
- Passing all detections through a machine-learning system trained to catch and reject artifacts: satellite trails, hot pixels, cosmic rays.
- Cori supercomputer,
  Cray XC40 rank 8 of Top500
  14 Pflops, Xeon Phi.

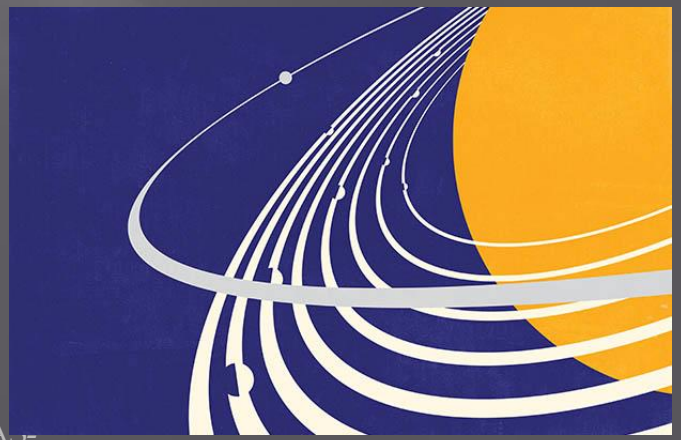

Fig. 2.5 Looking for planet 9

2.3 Opinions

- Popular topic today
- « Artificial intelligence is the future …
  it comes with colossal opportunities but also threats that are difficult to predict.»

  V. Poutine
  Talk to the students
  4 Septembre 2017

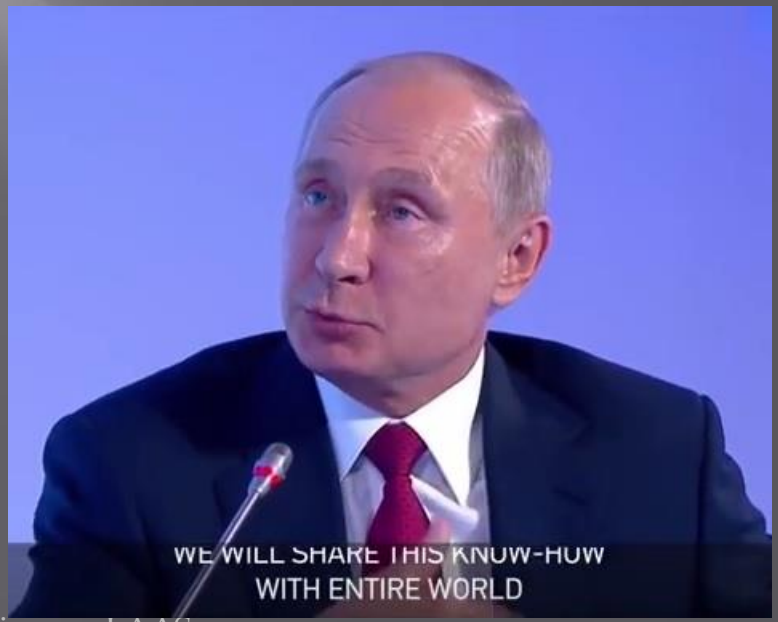

2.3 Opinions

- « L'Europe aura besoin d'investir de façon importante dans ses infrastructures de calcul, domaine sur lequel elle est très en retard par rapport à la Chine et aux Etats-Unis, ainsi que dans le hardware, l'industrie du semi-conducteur. Ce sera une grosse, grosse affaire. »

  C. Villani
  L'Obs
  28 Février 2018

3. How?

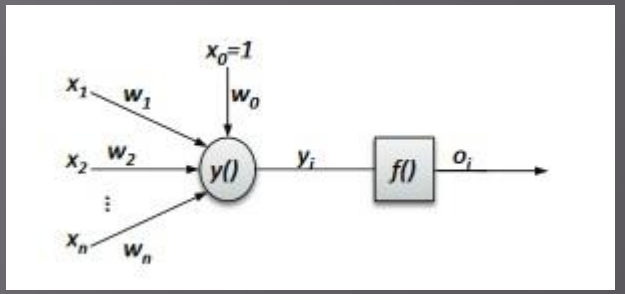

- Mid 40s to 2018.
- From Computational models for neural networks to parallel deep Learning methods.
- Artificial Neural Networks (ANN), Perceptron.

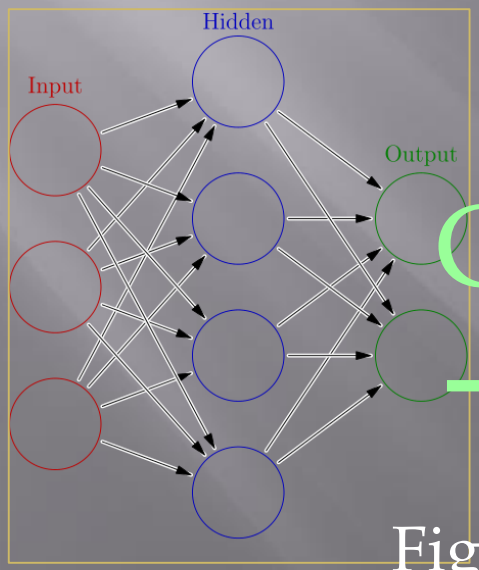

Fig 3.1 ANN

GPUs →

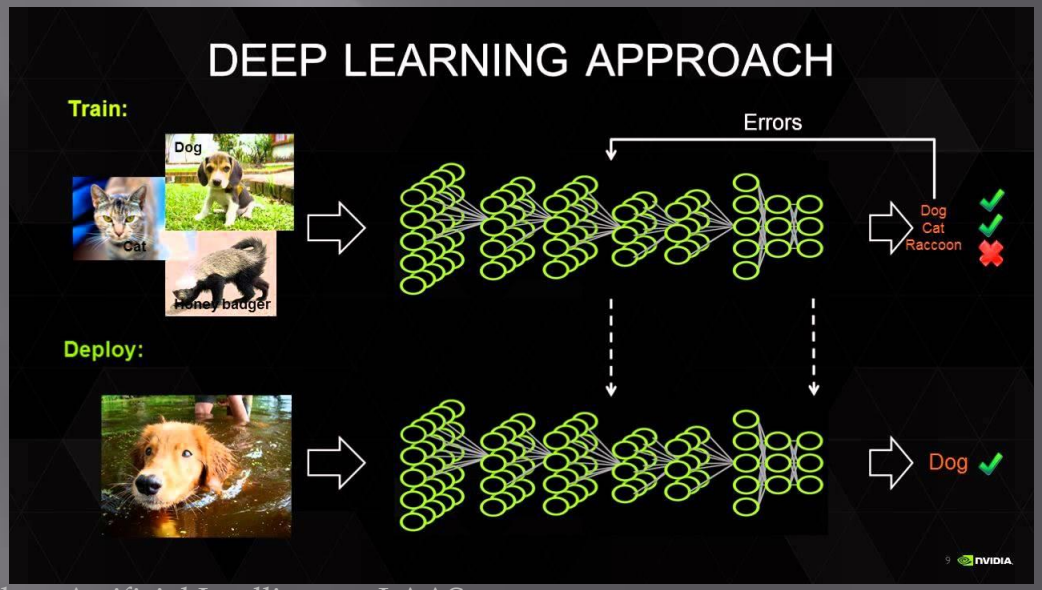

3.1 Algorithms and Data

- Algorithm that reproduces human / animal decision making.
- Heuristics or metaheuristics like
  - genetic algorithm;
    genetic algorithms → distributed computing.
  - ant colonies, swarms, flocks, fish school;
  - neural networks, deep learning.
  - Huge computations → parallelism; GPUs.
    Training ANN.

4. Where?

- Where is intelligence?
- Embedded intelligence (in the device);
  cost.
- Distributed intelligence (in the network),
  e.g., modular cyber-physical systems.
  resilience, volume, security issues.
- Hosted intelligence (deported on a server).
  miniaturization, data mining, security issues.
- In a supercomputer.
  solving difficult combinatorial optimization

4.1 New computing platforms and AI

- Use massive parallelism of

  Graphics Processing Units (GPU)

e.g. NVIDIA Jetson TX2 (embedded intelligence)
256 CUDA cores ;

< 10 Watts ;

< $400 ;

> 1 Tera flop (simple précision).
Up to six cameras.

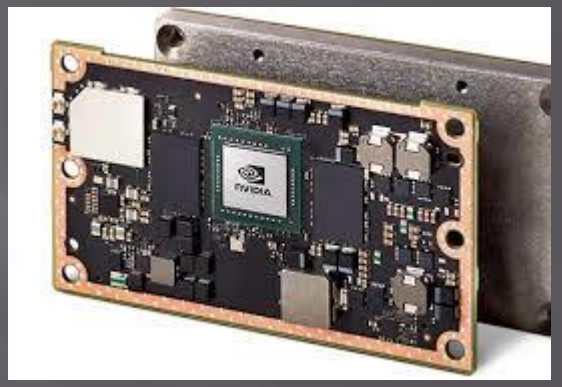

NVIDIA workshop, LAAS, March 22, 2016.

Club des Affiliés.

Fig. 4.2 Jetson TX2

4.2 AI supercomputers

- AI supercomputer: NVIDIA DGX-1
- Eight GPUs: P100 or V100.

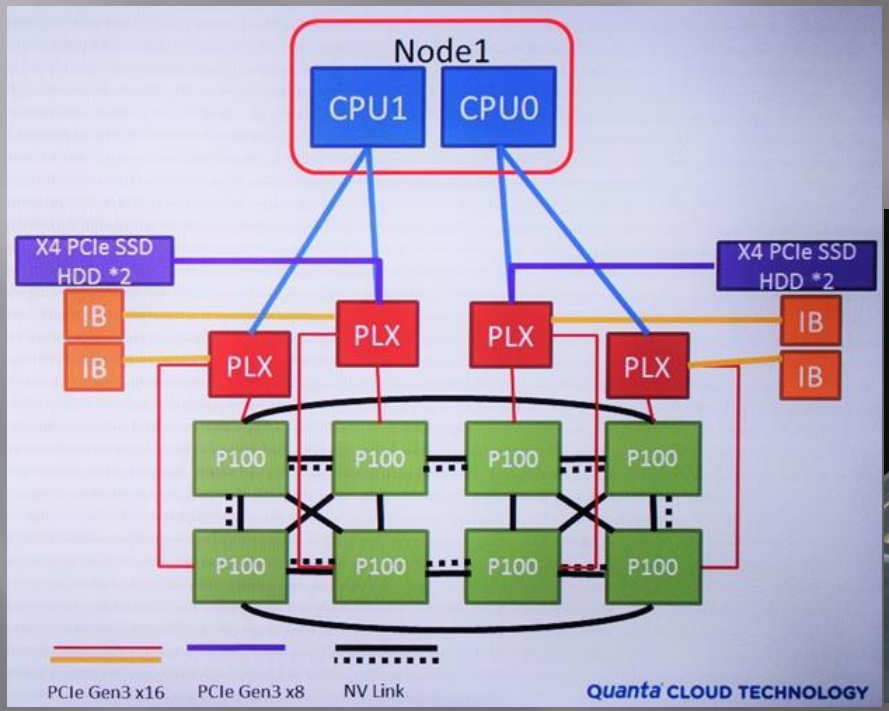

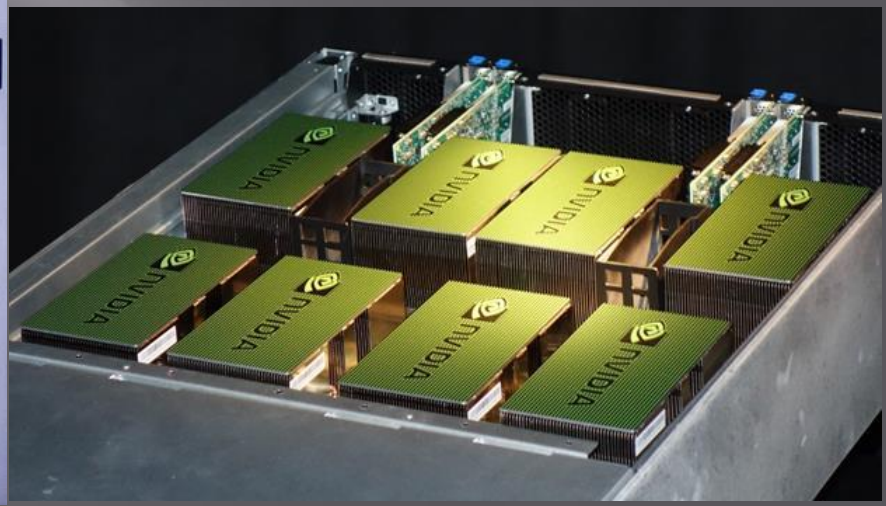

Fig. 4.2 NVIDIA DGX-1 supercomputer

4.2 AI supercomputers

- AI supercomputers

- From 170 TFLOPS up to 1 PFLOPS.
- From 28,672 up to 40,960 CUDA cores.
- Tensorflow.
- 96 x Faster Training than with Dual Xeon E5 -2699 (SIMT model).

5. CDA & AI

- Parallel or distributed metaheuristics.
- Training many neural networks in parallel on GPUs.
- Reconfigurable distributed smart conveyors.

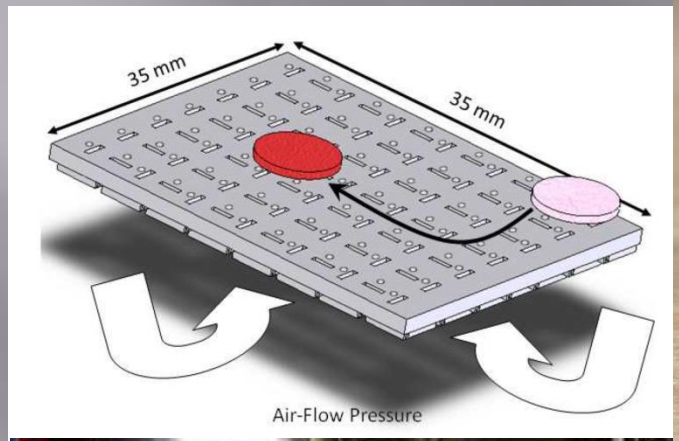

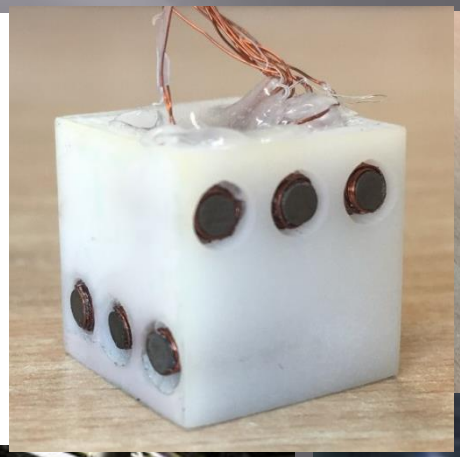

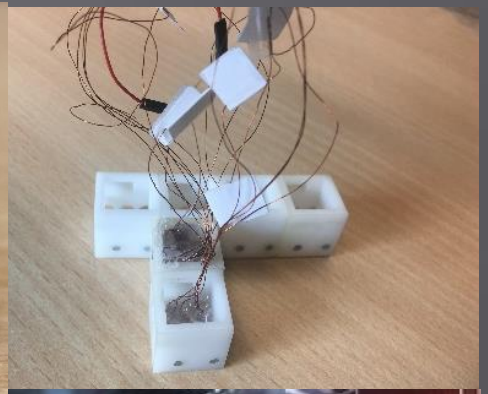

5.1 Smart Surface

- The Smart Surface conveyor in manufacturing industry.
- ANR 06 ROBO 0009, 2007 – 2010.
- FEMTO-ST, LAAS, LIMMS.
- Distributed part differentiation.

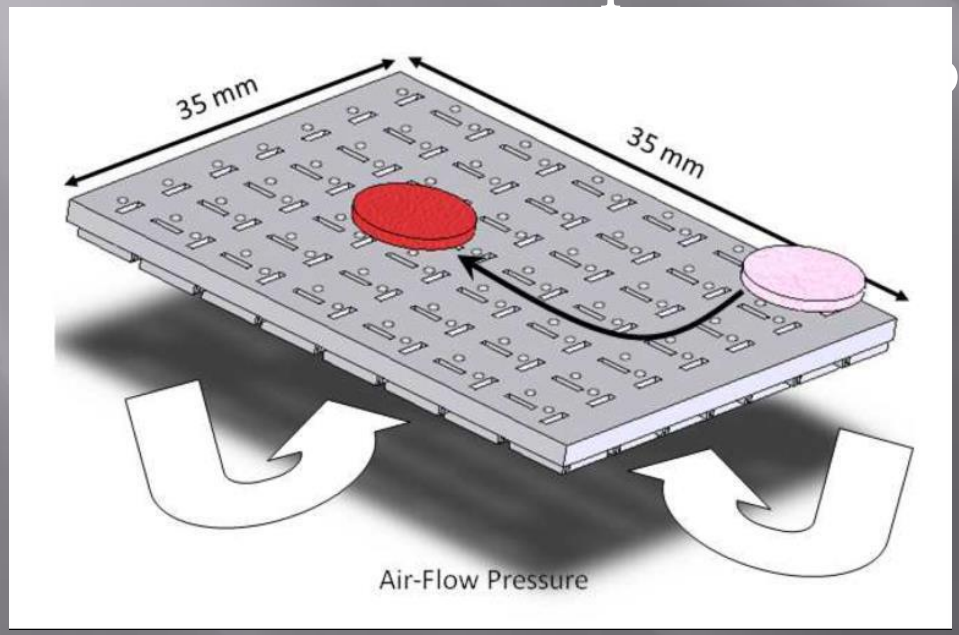

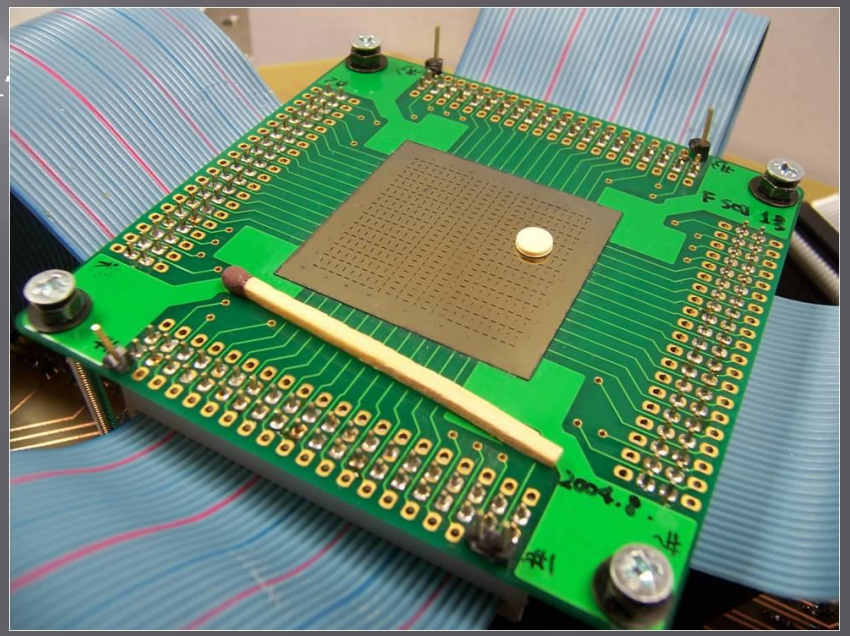

Fig. 6.1 Smart Surface

5.2 Smart Blocks (CDA)

Distributed autonomous modular system
Reconfigurable conveyor
Cyber-physical systems.
ANR-2011-BS03-005, 2011 – 2015

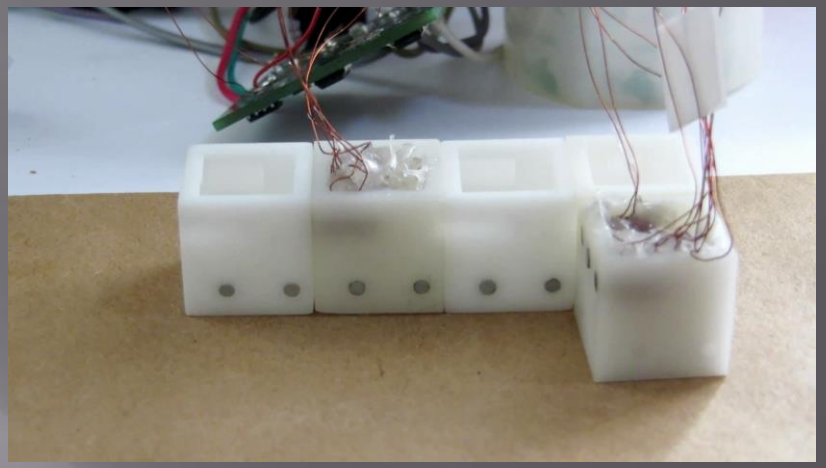

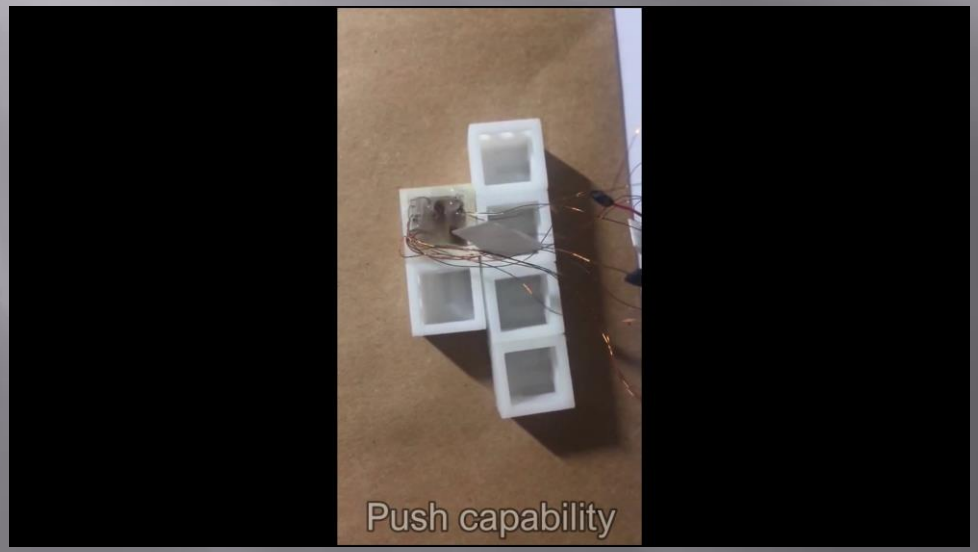

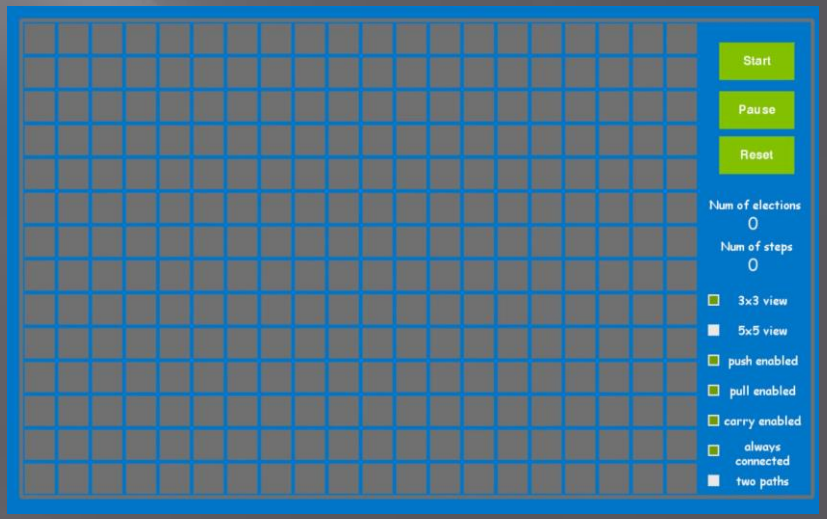

Li Zhu's Ph.D. February 2018

5.3 Scheduling and Parallel Metaheuristics

- Scheduling problems
- Energy Efficient Dynamic Flexible Flow Shop Scheduling.
- GPU-based Parallel Genetic Algorithm; K40 GPU.
- Jia Luo Ph.D. student.
- IEEE Workshop PDCO in conjunction with IEEE IPDPS

5.4 ANN Training and GPUs

- Product demand of a brewery company (real data).
- Training many ANN in parallel via back-propagation, K20 & K40 GPUs.
- J. Cruz Lopez, Mexican Master student visiting.

5.5 Intelligent Flying Machines

Deep Learning

Decision Making

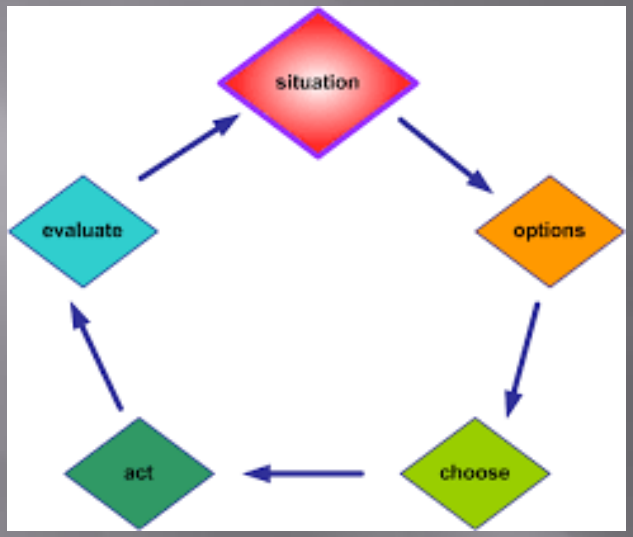

NonMonotonic Reasoning for Uncertain Situations

AI is driving UAV Intelligence

UAV testbed at Luminy, JL Vilchis Ph.D. student

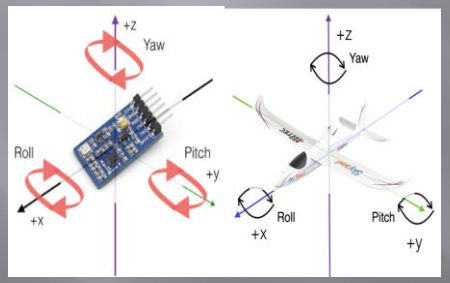

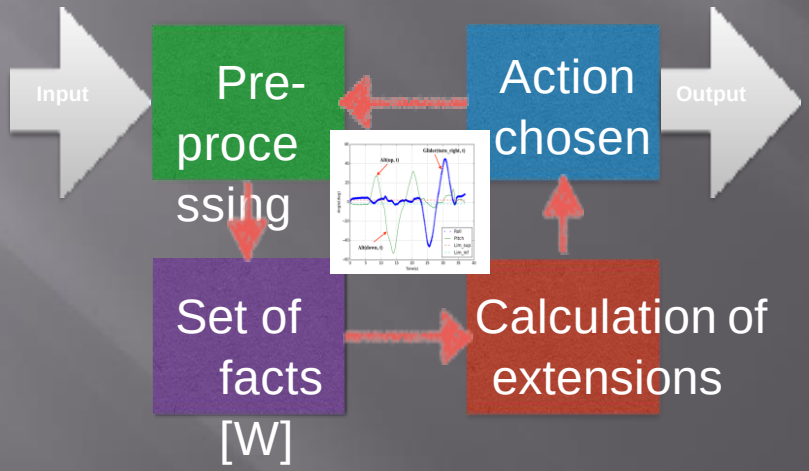

6. Conclusions

- AI & 4th Industrial Revolution.
- Convergence of many domains like

  AI, Data Mining, Non linear Optimization, HPC.
- Design of parallel or distributed AI algorithms (HPC) is a hot topic.

6. Conclusions

- AI: Tremendous opportunities in transport and manufacturing industry.
- Smart systems, smart cities, smart world
  - Huge impact on society.
  - Less urbanization?
- IEEE SmartWorld Congress 2016 Toulouse.
  UIC, ATC, ScalCom, CBDCOM, IOP
  2017 → San Francisco
  2018→ Guangzhou, October 2018.
  participation of neOCampus

6. Conclusions

- AI may lead to important advances in the way to exploit country resources.
- AI may lead to important discoveries in sciences and technology when combined with HPC & data mining, with application to
  - Astronomy & Astrophysics;
  - Physics;
  - Geology;
  - Biology & Medicine;
  - Language Models;
  - and many more.

6. Conclusions

- "Two things are infinite: the universe and human stupidity; and I'm not sure about the universe."

  A. Einstein

- Thankfully or unfortunately: there is nothing that AI can do that has not already been done by man.

Acknowlegments

- Andrei Doncescu,
- O. Llopis,
- Li Zhu,
- Lundl Chen,
- J. Luo,

- NVIDIA,
- TC on Smart World of IEEE CIS.

Publications

- D. El Baz et al. Distributed part differentiation in a smart surface, Mechatronics, Vol. 22, Issue 5, 2012, p. 522-530.
- V. Boyer, D. El Baz, M. A. Salazar-Aguilar, GPU Computing Applied to Linear and Mixed Integer Programming, Chapter 10 in Advances in GPU, Research and Practice, H. Sarbazi-Azad editor, Morgan Kaufmann, Elsevier, Amsterdam Boston, ISBN 978-0-12-803738-6, 2017, p. 247 - 271.
- D. El Baz, Keynote Speaker : Cyber-physical systems and various computer science issues in smart distributed autonomous robots, Third International Scientific School Incident Management and Countering Targeted Cyber-Physical Attacks in distributed large-scale critical systems (IM & CTCPA 2017), Saint-Pétersbourg Russie, 19 Décembre 2017
- D. El Baz, Keynote Speaker : High performance issues related to Internet of Things and Smart Earth, International Conference Parallel Computing, ParCo2017, Bologne Italie, 12-15 Septembre 2017.

Publications

- D. El Baz, Keynote Speaker : Challenges in Computing Accelerators and Heterogeneous Computing, 25th International Conference on Parallel Distributed and networked based Processing (PDP 2017) Saint Petersburg Russie, 6 au 8 Mars 2017.
- D. El Baz, Keynote Speaker : IoT and the Need for High Performance Computing, in Proceedings of the International Conference on Identification, Information and Knowledge in The Internet of Things (IIKI2014), 17-18 Octobre 2014, Pékin Chine, p. 1-6, IEEE CPS.
- J. Cruz-Lopez, V. Boyer, D. El Baz, Training Many Neural Networks in Parallel via Back-Propagation in Proceedings of the 27th IEEE Symposium IPDPSW 2017 / PDCO'17, Orlando USA, 29 May 2 June 2017.
- Didier El Baz, Mhand Hifi, Lei Wu, Xiaochuan Shi, A Parallel Ant Colony Optimization for the Maximum-Weight Clique Problem in Proceedings of the 30th IEEE Symposium IPDPSW 2016 / PCO 2016, Chicago 2016, 23-27 May 2016, p. 796 - 800.

Publications

- J.L. Velchis, P. Siegel, A. Doncescu, " Autonomous Aerial Vehicle Based on Non-Monotonic Logic", VEHITS 2017, Porto, Portugal, 22-24 Avril, 2017.