Some thoughts from ETS*:
- Part 1: Lessons learned with RISC-V;
- Part 2: AI for test/diagnosis

CLAUDE THIBEAULT

PROFESSOR, EE DEPT., ÉCOLE DE TECHNOLOGIE SUPÉRIEURE. MONTREAL

* WITH HELP FROM POLYTECHNIQUE MONTREAL FOR PART 1

Part 1: *Lessons learned with RISC-V*

CLAUDE THIBEAULT, ÉTS
MICKAËL FIORENTINO, POLYTECHNIQUE MONTREAL
YVON SAVARIA, POLYTECHNIQUE MONTREAL

Agenda

- **Part 1: *Lessons learned with RISC-V***
  - RISC-V Initiative
  - Research context
  - Lessons learned (so far)
- **Part 2: *AI for Test/Diagnosis***

Part 1: *Lessons learned with RISC-V*

- RISC-V initiative:
  - RISC-V:
    - Open, free Instruction Set Architecture (ISA) + tool chain
    - University of California at Berkeley
    - Krste Asanović and David A. Patterson
    - Initial development: 2010 to 2014
  - RISC-V Foundation
    - Started in 2015
    - 100+ members including: Google; Microsemi; NVIDIA; NXP.
    - Workshop (each fall and spring)

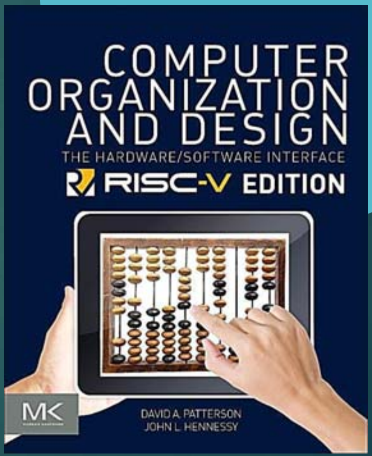

Part 1: *Lessons learned with RISC-V*

- For more information:
  - Web site: www.riscv.org
  - Technical report: "Instruction Sets Should Be Free: The Case For RISC-V": https://www2.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-146.pdf
  - First RISC-V Workshop opening session : https://riscv.org/wp-content/uploads/2015/01/riscv-intro-workshop-jan2015.pdf

Part 1: *Lessons learned with RISC-V*

- Research context:
  - In collaboration with Octasic, development of a highly testable low power asynchronous version of an ARM processor
  - Working on a design flow:
    - Verification through emulation on FPGA + ASIC implementation
    - First proof of concept using a Mini-MIPS processor
    - Needed a bigger benchmark circuit → We chose RISC-V because:
      - Base Integer ISA easy to implement
      - Open-source & versatile toolchain (including good documentation)
      - **Growing and dynamic community** → good impact on research

Part 1: *Lessons learned with RISC-V*

- Lessons learned (so far):
  - Two students attended 2016 Workshop in Boston:
    - Occasion to meet the people from the community
    - Topics : ISA, software ecosystem, microarchitectures, hardware implementation
    - A practical session on soft-core integration on FPGA using the Chisel HDL language

Part 1: *Lessons learned with RISC-V*

- Lessons learned (so far):
  - Developed a VHDL synchronous version of the RISC-V on FPGA
    - Multiple open-source RISC-V implementations (Verilog, VHDL, Chisel…) available
      - We chose to start from scratch in VHDL:
        - RV32IM (32 bit integer + multiply/divide) ISA
        - 5-stage pipeline, basic branch prediction, basic interrupt
        - Adapting the toolchain to fit the design

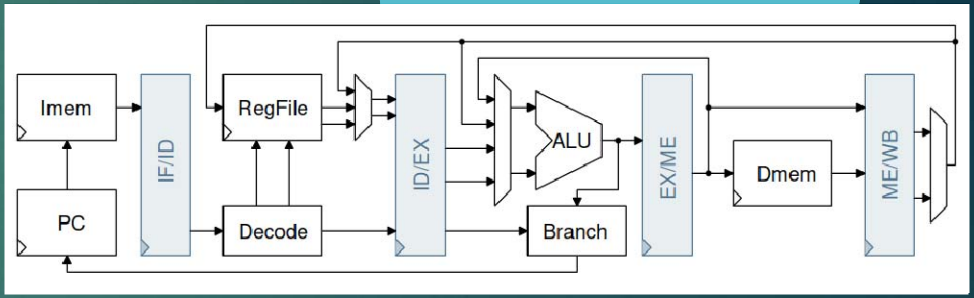

    - Overall, it took about 2 months to one (very good) PhD student complete and verify the VHDL model

Part 1: *Lessons learned with RISC-V*

- Conclusions:
  - So far, so good; we are about to start working on the asynchronous version.
  - RISC-V is a good vehicle for microarchitectural development :
    - Versatile ISA : a wide variety of implementations
    - Complete open-source software ecosystem (compilers, OS, libraries…)
  - We recommend it
    - Easy to learn → Enables multi-disciplinary projects
    - Dynamic and growing research community

Part 2: AI for test/diagnosis

CLAUDE THIBEAULT, ÉTS

Agenda

- Part 1: *Lessons learned with RISC-V*
- **Part 2: *AI for Test/Diagnosis***
  - Taxonomy
  - A bit of history
  - Where are we?
  - My wish list…

Part 2: *AI for test/diagnosis: Taxonomy*

Artificial Intelligence

Machine-Learning

Deep Learning

- "A technique for implementing machine learning ," [4]
- " ...one of many approaches to machine learning," [3]
- With respect to other approaches, automatic feature extraction [
- Good when features unknown/hard to derive + enough data [1]
- "Deep ANNs... with more than one hidden layers ..." [5]

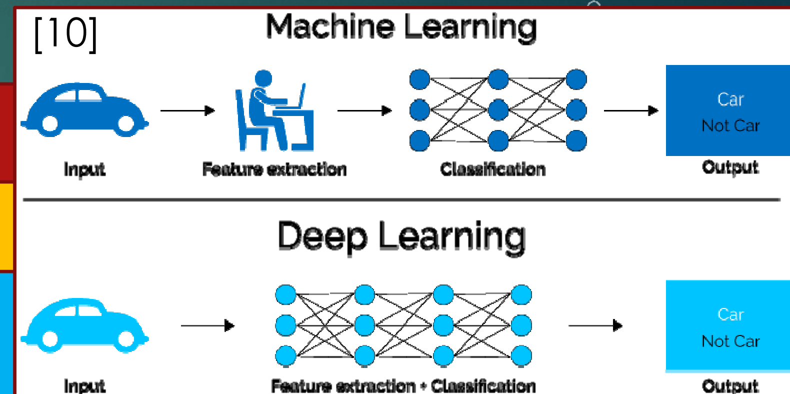

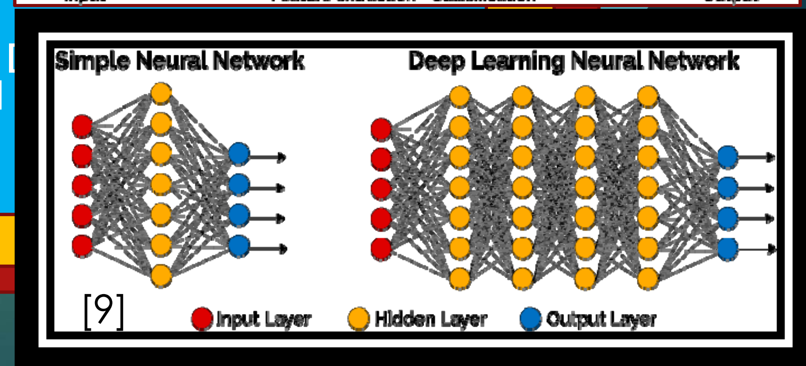

Part 2: *AI for test/diagnosis: a bit of history*

- When the terms appeared:
  - Artificial Intelligence: 1956
  - Machine Learning: 1959
  - Deep Learning: 1986
- So, why this hype now?
  - Processing power: Moore's Law + GPUs (speedup training by 100X)
  - Tons of data for training: IOT, data centers, etc.
  - Driving application(s): (IOT, data centers), autonomous cars, etc.

Part 2: *AI for test/diagnosis: a bit of history*

- What about test/diagnosis?
- No surprise here, it started a long time ago:

IEEE TRANSACTIONS ON SEMICONDUCTOR MANUFACTURING, VOL. 2, NO. 2, MAY 1989 47

A Machine-Learning Classification Approach for IC Manufacturing Control Based on Test Structure Measurements

M. E. ZAGHLOUL, SENIOR MEMBER, IEEE, D. KHERA, MEMBER, IEEE, L. W. LINHOLM, MEMBER, IEEE, AND C. P. REEVE

Part 2: *AI for test/diagnosis: a bit of history*

- Test/diagnosis : A lot of test data ...

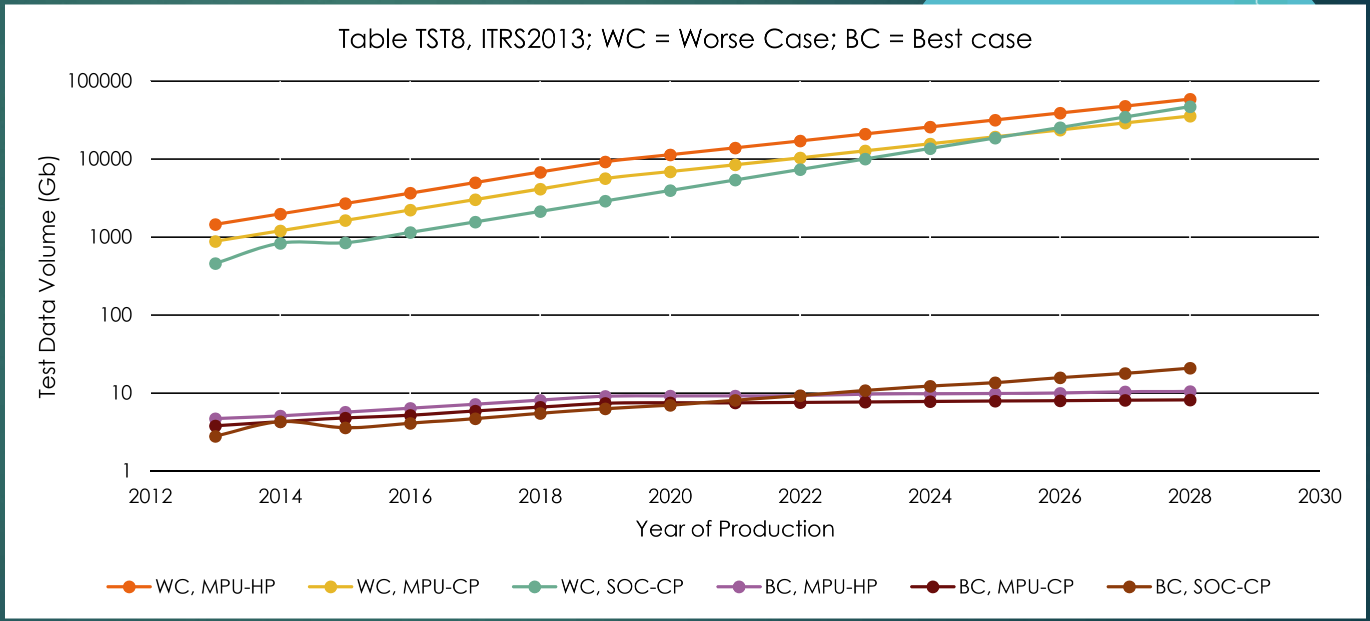

Part 2: *AI for test/diagnosis: a bit of history*

- Test/diagnosis: A lot of test data ...
  - Smart phones: In 2017, Samsung shipped approximately 310 million units [6]
  - 310M/year = 849k shipped units/day
  - Assuming a (optimistic) 80% yield, 849k/(day * 80%) = 1M tested units /day
  - Assuming 10Gb/unit, daily test data volume = 10^{16} bits = 10Petabits = 10000Terabits; daily diagnosis data volume 2000 Terabits (Big Data!)
- A high test accuracy: assuming 200DPM, 99.98%; not as high for diagnosis

Part 2: *AI for test/diagnosis: a bit of history*

- To deal with this data volume:

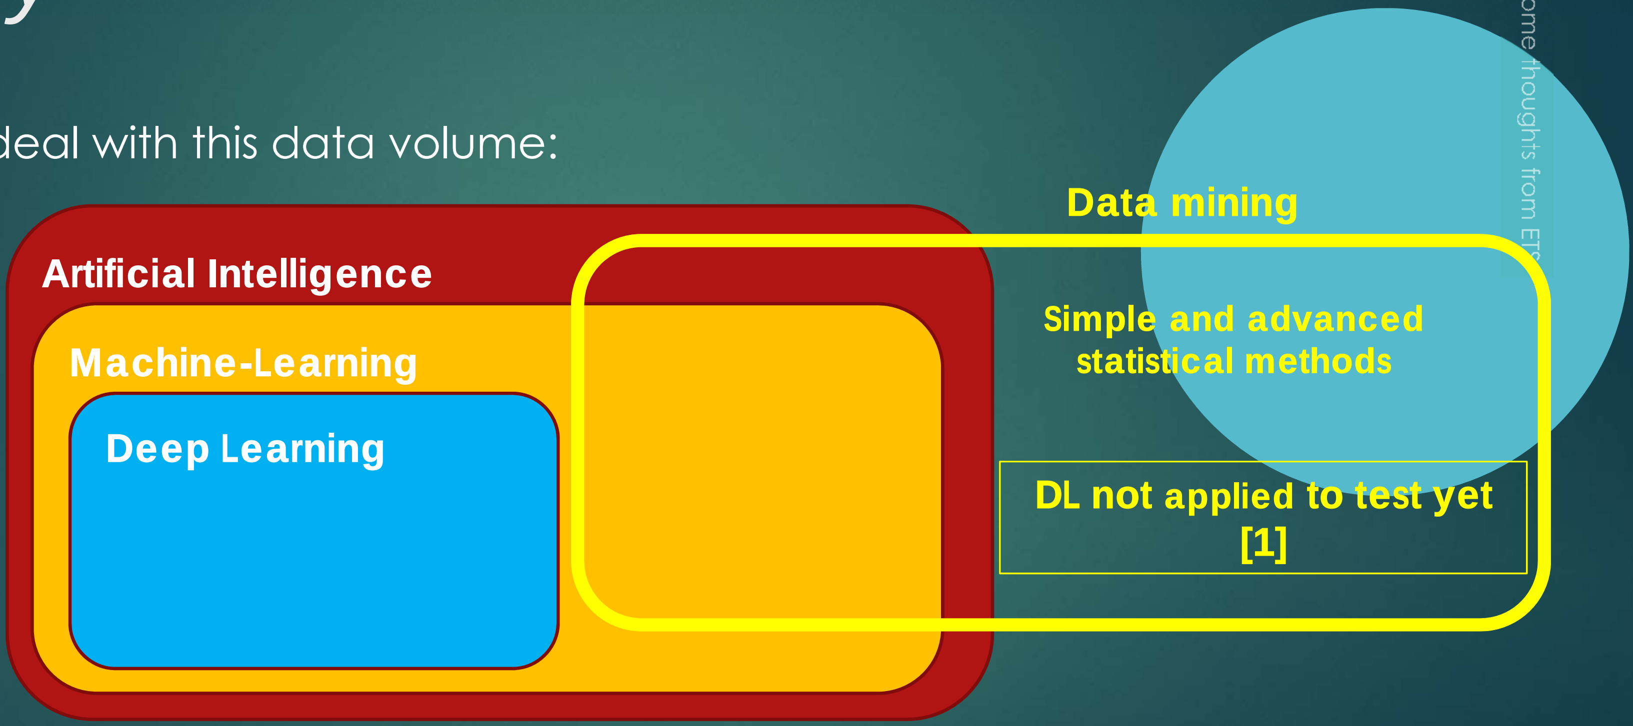

Part 2: *AI for test/diagnosis: a bit of history*

- Over the years, different ML approaches explored:
  - Bayesian Inference:
    - Diagnosis: Z. Zhang et al., "Board-Level Fault Diagnosis using Bayesian Inference," VTS2010.
    - Yield learning: W.-T. Cheng et al., "Volume Diagnosis Data Mining", ETS2017

P{A | B} = P{B | A} * P{A} / P{B} = Prob. of hypothesis A given observed evidence B

Where:
P{B | A} = Prob. of evidence B, given hypothesis A
P{A} = Prob. of hypothesis A , before observing evidence B
P{B} = Prob. of evidence B, under all hypotheses A

Part 2: *AI for test/diagnosis: a bit of history*

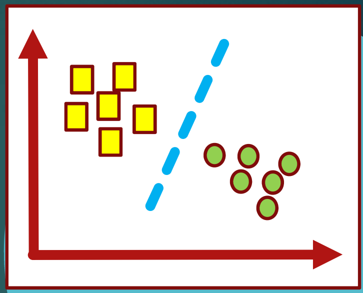

- Over the years, different ML approaches explored:
  - Support Vector Machine:
    - Diagnosis
      - F. Ye et al., “Board-Level Functional Fault Diagnosis Using Multikernel Support Vector Machines and Incremental Learning,” TCAD2014.
    - Yield learning
      - J. Tikkanen et al., “Yield Optimization Using Advanced Statistical Methods”, ITC2014
  - See Youtube video: https://www.youtube.com/watch?v=3liCbRZPrZA

Part 2: *AI for test/diagnosis: where are we?*

- Coming soon:
  - VTS2018 program (April 22nd to 25th):
    - Morning Tutorial – Machine Learning and Its Applications in Test
      - In this tutorial, we will start by covering the basics of machine learning. We will proceed to give a brief overview of the new and exciting field of deep learning. We will show how easy it is to try using machine learning and **deep learning**, thanks to powerful, free libraries. After offering the required background in machine learning, we will review several important papers in the field of DFT, diagnosis, yield learning, and root cause analysis, which use machine learning algorithms for solving various problems. Finally, we will propose future research directions in the area of testing, where we think machine learning (especially **deep learning**) can make a big impact.

Part 2: *AI for test/diagnosis: where are we?*

- Coming soon:
  - VTS2018 program (11 presentations on Machine Learning, including 3 on deep learning):
    - **Exploiting Deep Learning System-level Vulnerabilities from the Intelligent Supply Chain**
      Presenter: Wujie Wen, Florida International University
    - **Overcoming the Challenges of Hotspot Detection using Deep Learning**
      Kareem Madkour, Mentor, A Siemens Business
    - **Machine Learning in Semiconductor Test: Can Deep Learning Save The Day?**
      Presenter: Yiorgos Makris, UT Dallas

Part 2: *AI for test/diagnosis: My wish list*

- As a ML user:
  - Processing power
  - Tools & tutorials
  - (A lot of good) Data
- As a (potential) ML developer:
  - FPGA-based acceleration: Hardware platform
- Last thought: We should take advantage of the local ML (and DL) expertise

Part 2: *AI for test/diagnosis: conclusions*

- Conclusions:
  - About to enter the DL era in test/diagnosis
  - We need tools + data + processing power to use it
  - We need hardware platform to improve it
  - We should leverage local expertise on DL

Thank you!

Part 2: *AI for test/diagnosis: sources*

Some thoughts from ETS

[1] Y. Huang, “Machine Learning for Test,” International Test Conference, Nov. 1, 2017, Special Session 4.1.

[2] B. Marr, “What is the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning?”, Forbes, DEC 6, 2016, https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/#5fc0c1b12742

[3] C. McClelland,” The Difference Between Artificial Intelligence, Machine Learning, and Deep Learning,” Medium, Dec. 4, 2017, https://medium.com/iotforall/the-difference-between-artificial-intelligence-machine-learning-and-deep-learning-3aa67bff5991

[4] M. Copeland,” What’s the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning?”, Nvidia, July 29, 2016, https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

[5] DL4J,” What’s the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning?”, https://deeplearning4j.org/ai-machinelearning-deeplearning

[6] https://mobilesyrup.com/2018/02/14/samsung-apple-market-share/

[7] ITRS 2013: www.itrs2.net

[8] VTS 2018 program: http://tttc-vts.org/public_html/new/2018/technical-program/

[9] https://becominghuman.ai/deep-learning-made-easy-with-deep-cognition-403fbe445351

[10] https://medium.com/swlh/ill-tell-you-why-deep-learning-is-so-popular-and-in-demand-5aca72628780