Using ANN for power consumption estimation

Georges Da Costa

SEPIA Team

Jounée Deep Learning @ IRIT

26 juin 2018

Improving energy-efficiency of ICT

Knowledge is Power

Power consumption knowledge is crucial

- Datacenter operators
  - Power capping
  - Billing in colocation centers
  - Improving energy efficiency using autonomous systems
- Users and developers
  - Increasing their awareness
  - Choosing the best application or library

Monitoring power consumption

Two techniques:

- Power meters
  - Costly

- Power models
  - Analytical models based on DataSheets
  - Regression models
  - Neural Networks

Multi-scale models

Power consumption through the ages

- Linear classical model: error ~ 10-15%
  - $Power = P_{min} + Load \times (P_{max} - P_{min})$
- Adding frequency and voltage: error ~ 5-9%
  - $Power = P_{min} + Load \times \alpha\ Voltage^2\ Frequency$
- Adding temperature: error ~ 4-7%
  - $Power = P_{min} + Load \times \alpha\ Voltage^2\ Frequency + \lambda\ Temperature$
- Adding PSU bias, cooling, …: error ~ 2-3%
  - $Power_{DC} = \omega_0 + \omega_1\ Power_{AC} + \omega_2\ Power_{AC}$
- Learning method
  - Neural network: error ~ 2%

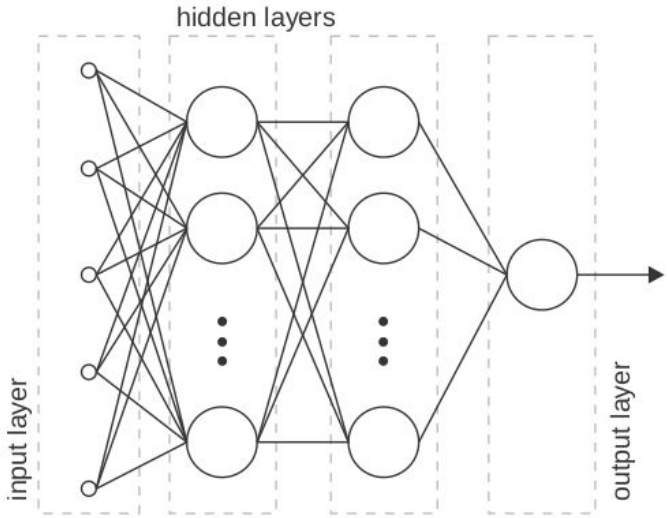

1. Acquiring data
2. Tuning the network
3. Results

Data preprocessing

Clapperboard are useful

Direct power measurement is difficult

- Raw power data are noisy
- Time and value precision are low
- High overhead of measures

Several methods to cope with imprecision

- Reaching steady states
- Reduce measure points
- Use movie-like Clap

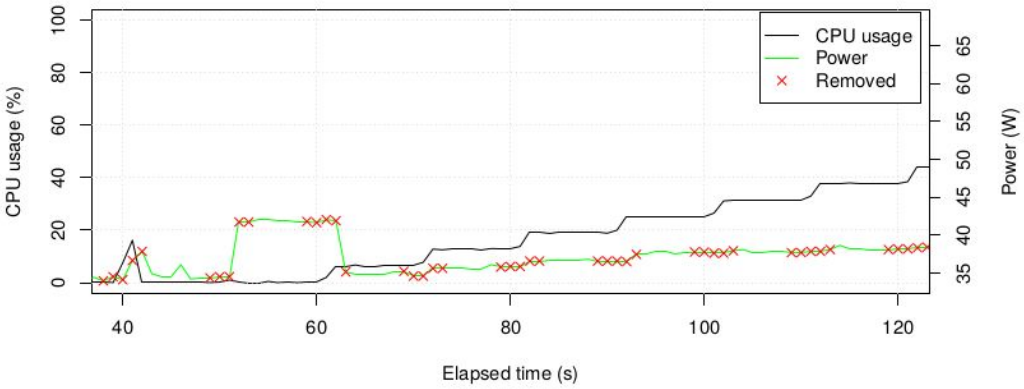

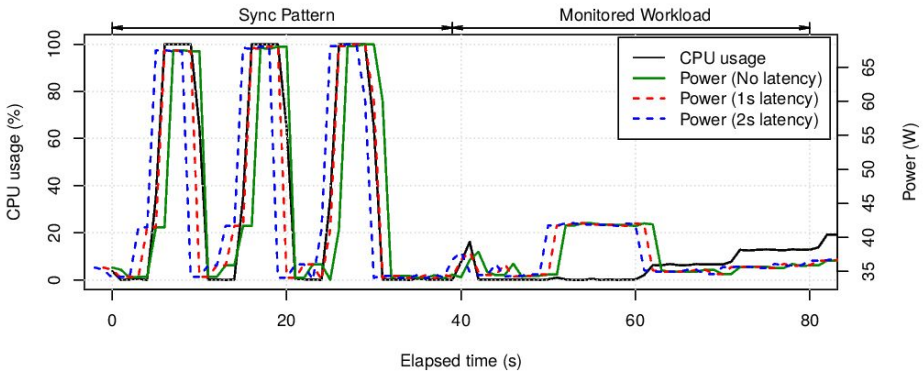

Network topology

Multi-layer perceptron (MLP) with one or two hidden layers

- Number of weights and biases to optimize:
  - $wb = (i + 1) * l_1 + (l_1 + 1) * l_2 + (l_2 + 1) * o$
  - i: inputs; o: outputs
  - l_1 and l_2 neurons on hidden levels

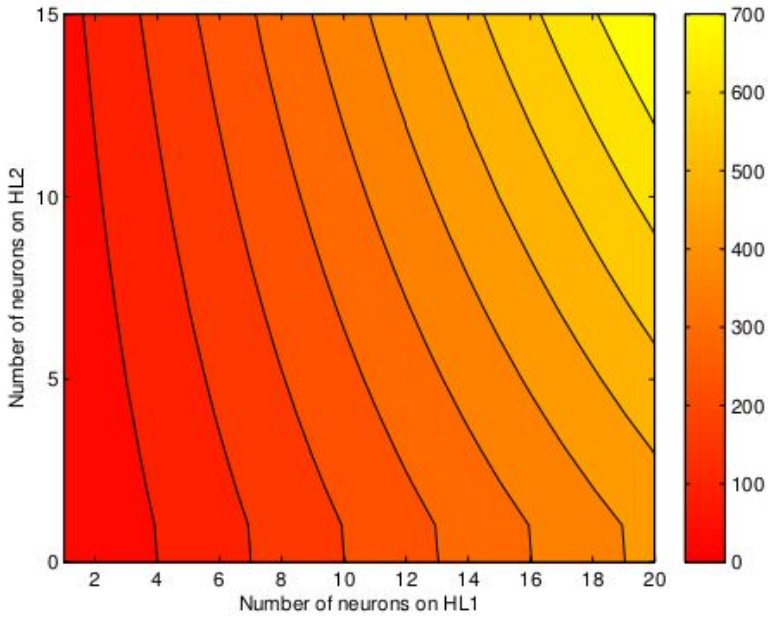

(a) 20 dimensions

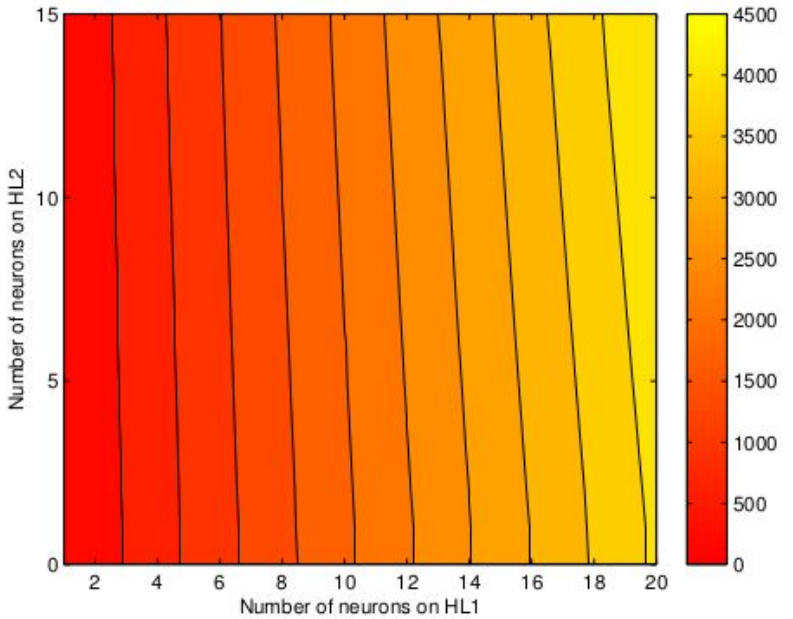

(b) 200 dimensions

Variable reduction

Input from more than 250 system-level sensors

One experiment : 1.3 Go/node

Multiple methods for variable reduction

- Increasingly add the most correlated variable
- Increasingly add the variable correlated with residual
- Increasingly add the best (using ANN) variable

Need of a stop criteria

- Adding a variable stops improving the ANN

Validation

Coverage problem

- ANN cannot guess
- Problem of overfitting

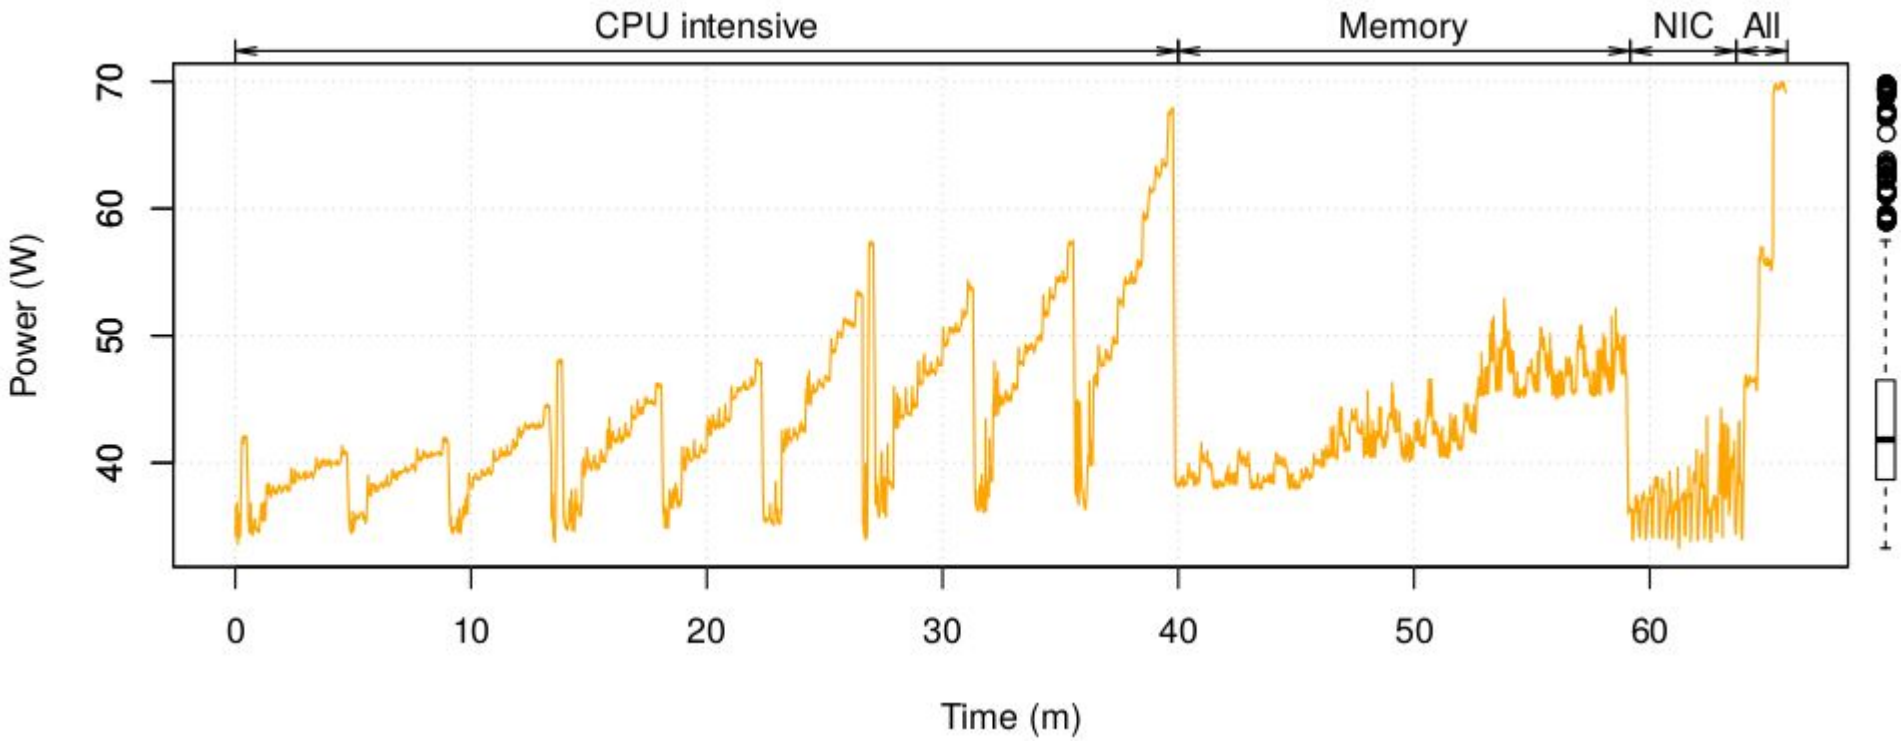

Results

- Validation with similar inputs
- With bias removal
- With optimized parameters

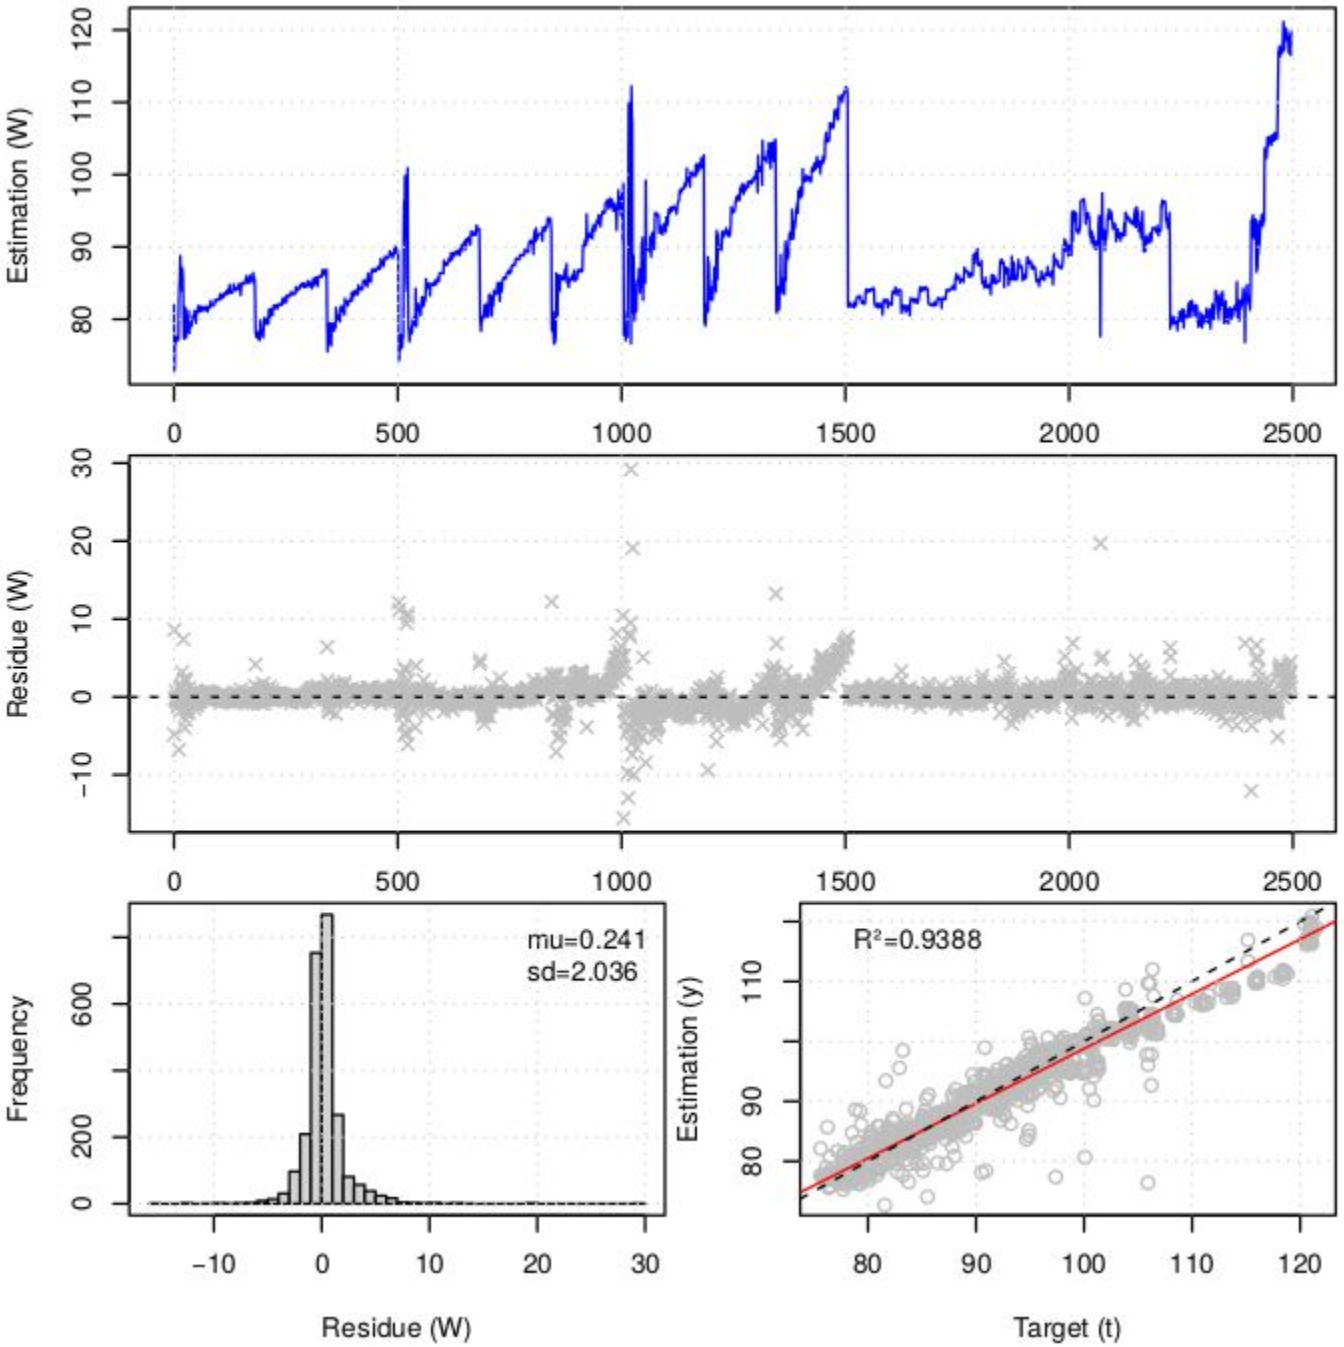

Performance evaluation

Using workload re-execution

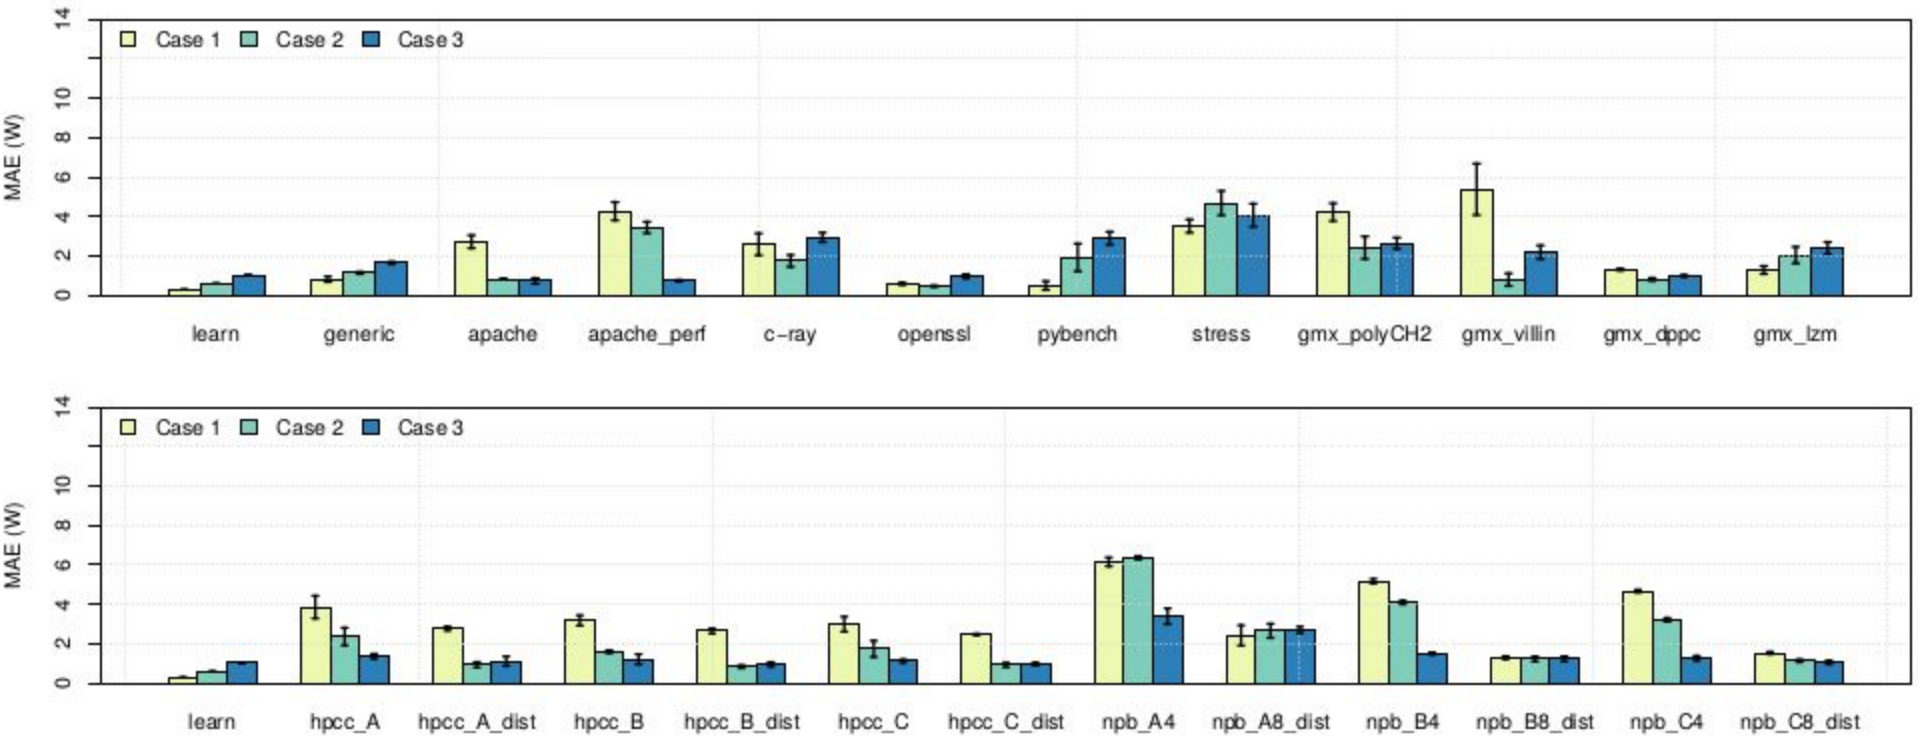

Perspective for SEPIA team

- Advanced **users**
  - Tune models to reduce their impact
  - Use ANN as $\mathbb{R}^n \Rightarrow \mathbb{R}$
- Usecases
  - Power models
  - Prediction models
    - Power
    - Time
  - Crisis management: similarities detection for rapid-flooding
    - Short term: real-time alert
    - Long term: improve area resilience
  - Smart cities and Datacenter optimization