Direct Load Control of Thermostatically Controlled Loads Based on Sparse Observations Using Deep Reinforcement Learning

Frederik Ruelens*, Bert J. Claessens†, Peter Vrancx‡, Fred Spiessens*, and Geert Deconinck*

*KU Leuven/EnergyVille, Leuven, Belgium † REstore, Antwerp, Belgium ‡VUB, Brussels, Belgium

Motivation

- Significant growth of renewable energy sources
  - **Limited controllability**
  - **Stochastic character**

Challenge: Causes problems with keeping the balance between electricity supply and demand

- Enabling technology is demand response using thermostatically controlled loads
  - Heat pump for space heating
  - Electric water heater
- These loads can use their thermal inertia, e.g. a water buffer or building envelope, as a thermal battery to store energy and shift energy consumption

Demand response objective

- Dynamic pricing objective: $\rho_k = \lambda_k * u^{\text{phys}_k}$ with λ_k and u_k^{phys} the current electricity price and electricity consumption.

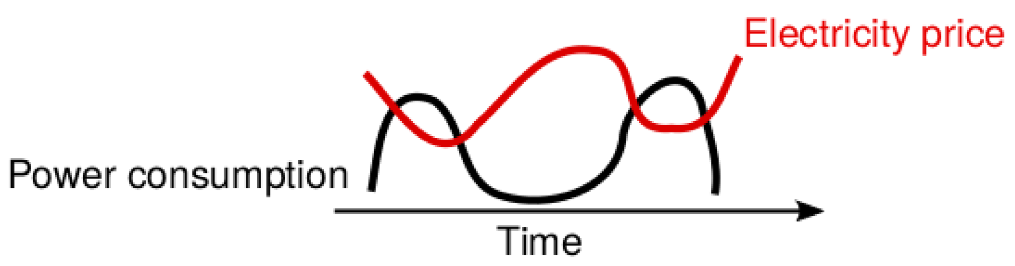

Problem statement

In most complex real-world problems, such as demand response, an agent cannot measure the exact full state of its environment, but only a partial observation of the state. Depending on how good this partial observation can be used to model future interactions, using partial state information may result in sub-optimal policies.

Partial observability

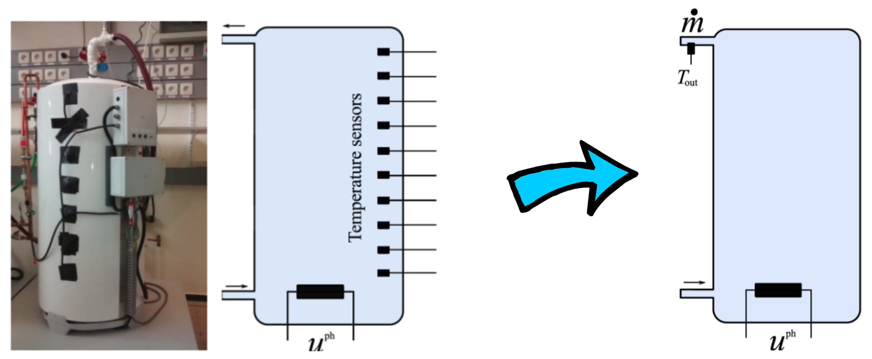

- Electric wather heater
  - Power consumption
  - Temperature and flow rate of water exiting the buffer
- Heat pump for space heating
  - Power consumption
  - Outside temperature

Partial state information

$$x_k^{\text{aug}} = (x_k^{\text{time}}, [u_{k-1}, \dots u_{k-h}], [u_{k-1}^{\text{phys}}, \dots u_{k-h}^{\text{phys}}], [\dot{m}_k, \dots \dot{m}_{k-h+1}], [T_k^{|\mathcal{L}|}, \dots, T_{k-h+1}^{|\mathcal{L}|}]),$$

- with:
  - x_k^{time}: timing information: current quarter in the day
  - u_k: control action
  - u_k^{phys}: electricity consumption
  - $\dot{m}_k$: flow rate of water exiting buffer vessel
  - $T_k^{|\mathcal{L}|}$: temperature of water exiting buffer vessel
  - h: history length

Recurrent neural network

- Unfolded recurrent neural network

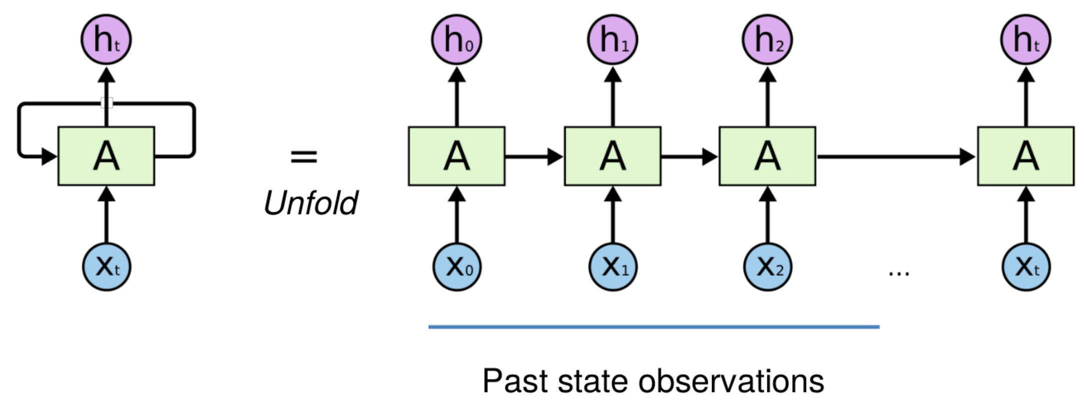

`https://colah.github.io/posts/2015-08-Understanding-LSTMs`

Approximator architecture using LSTMs

- Q-function approximator used within fitted Q-iteration (Ernst *et al.*)

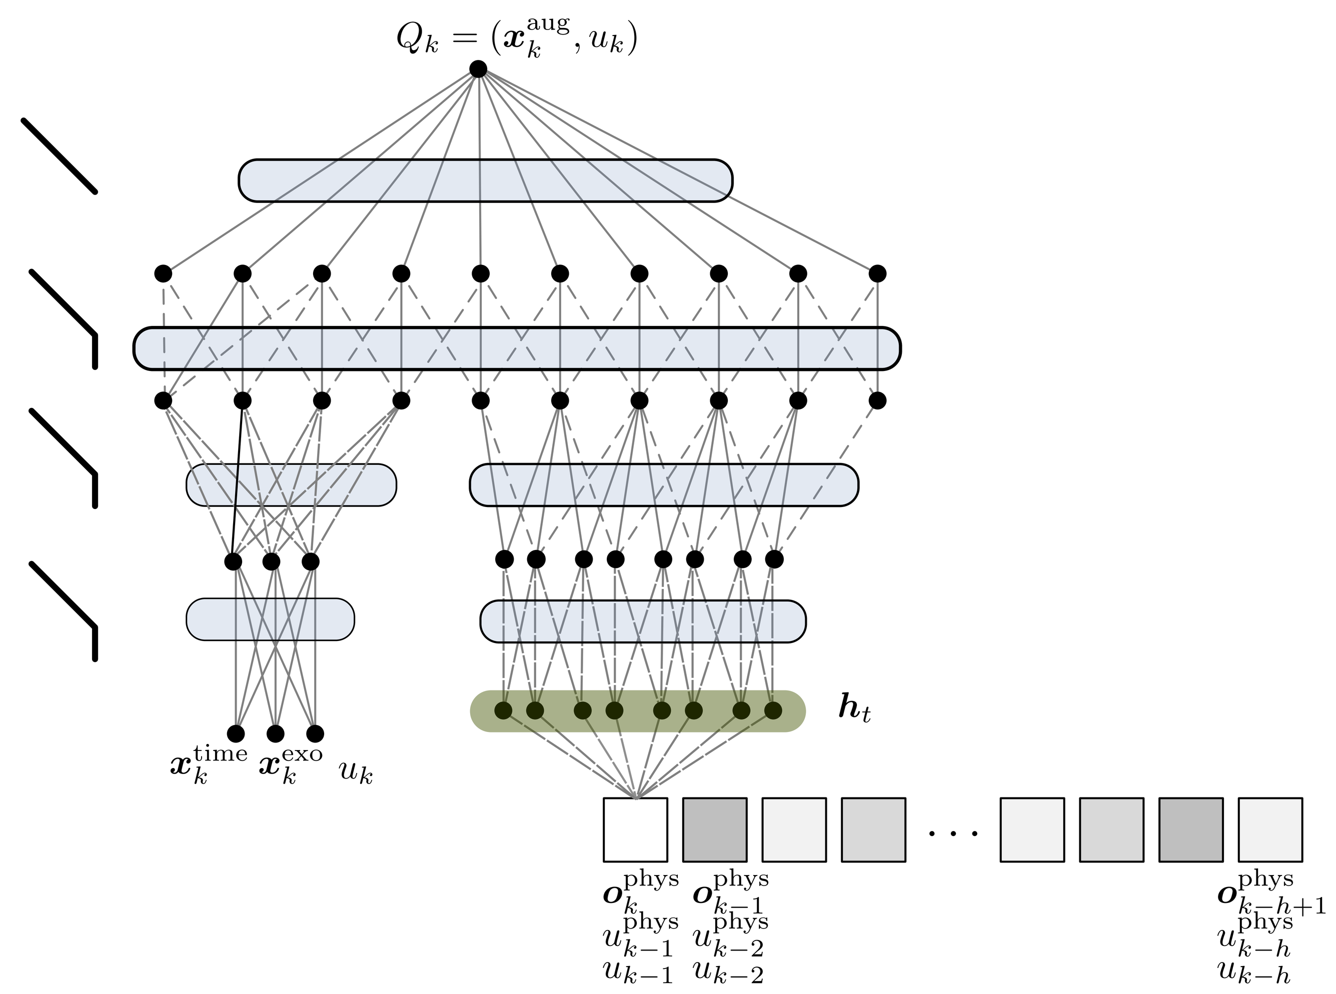

Data Science Summer School, Paris 2017

Simulation Results

- Electric water heater (grey: electricity price, blue: power consumption, green: state of charge)

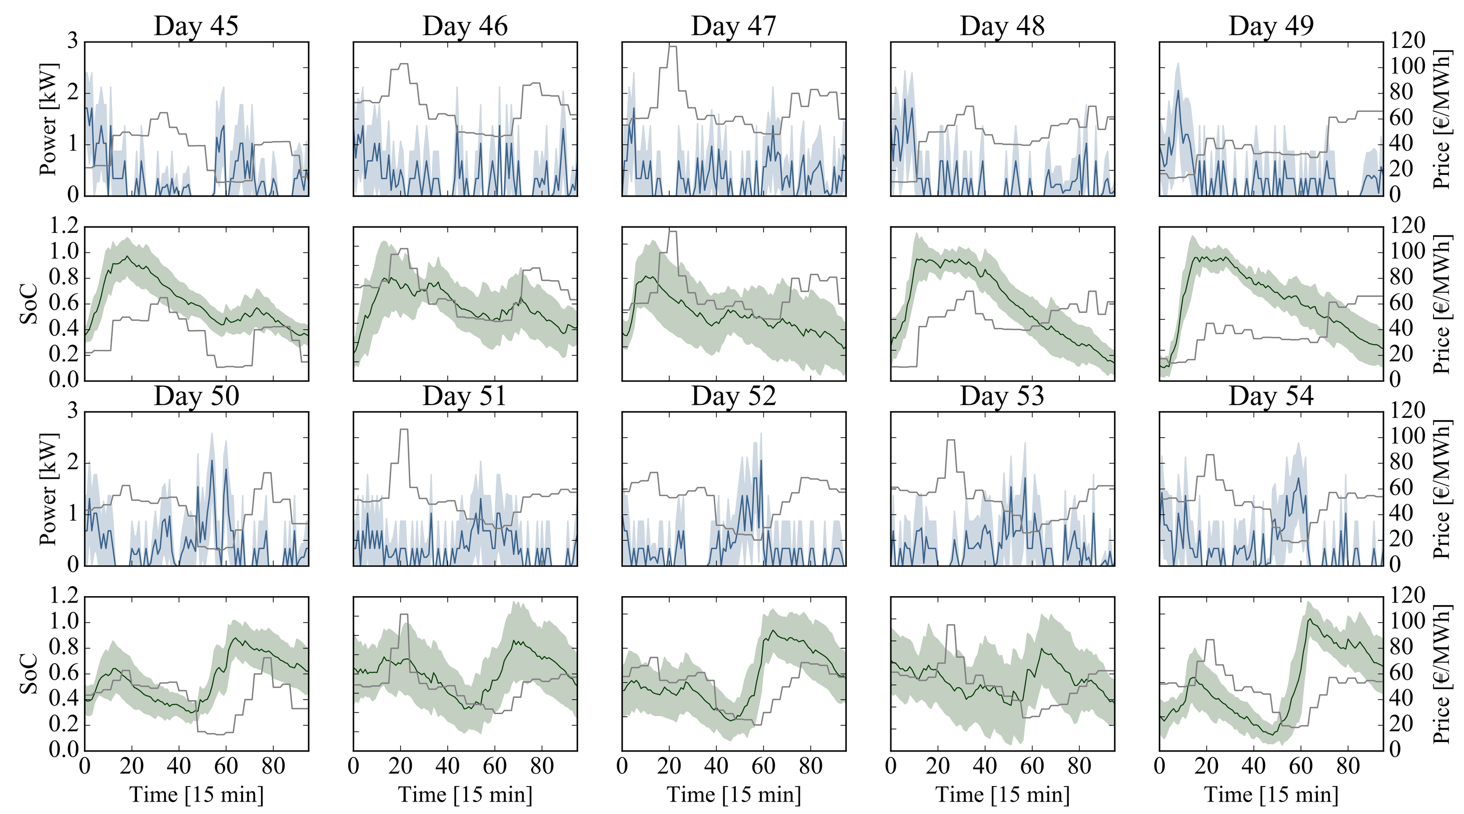

- Heat pump (green: indoor air temperature)

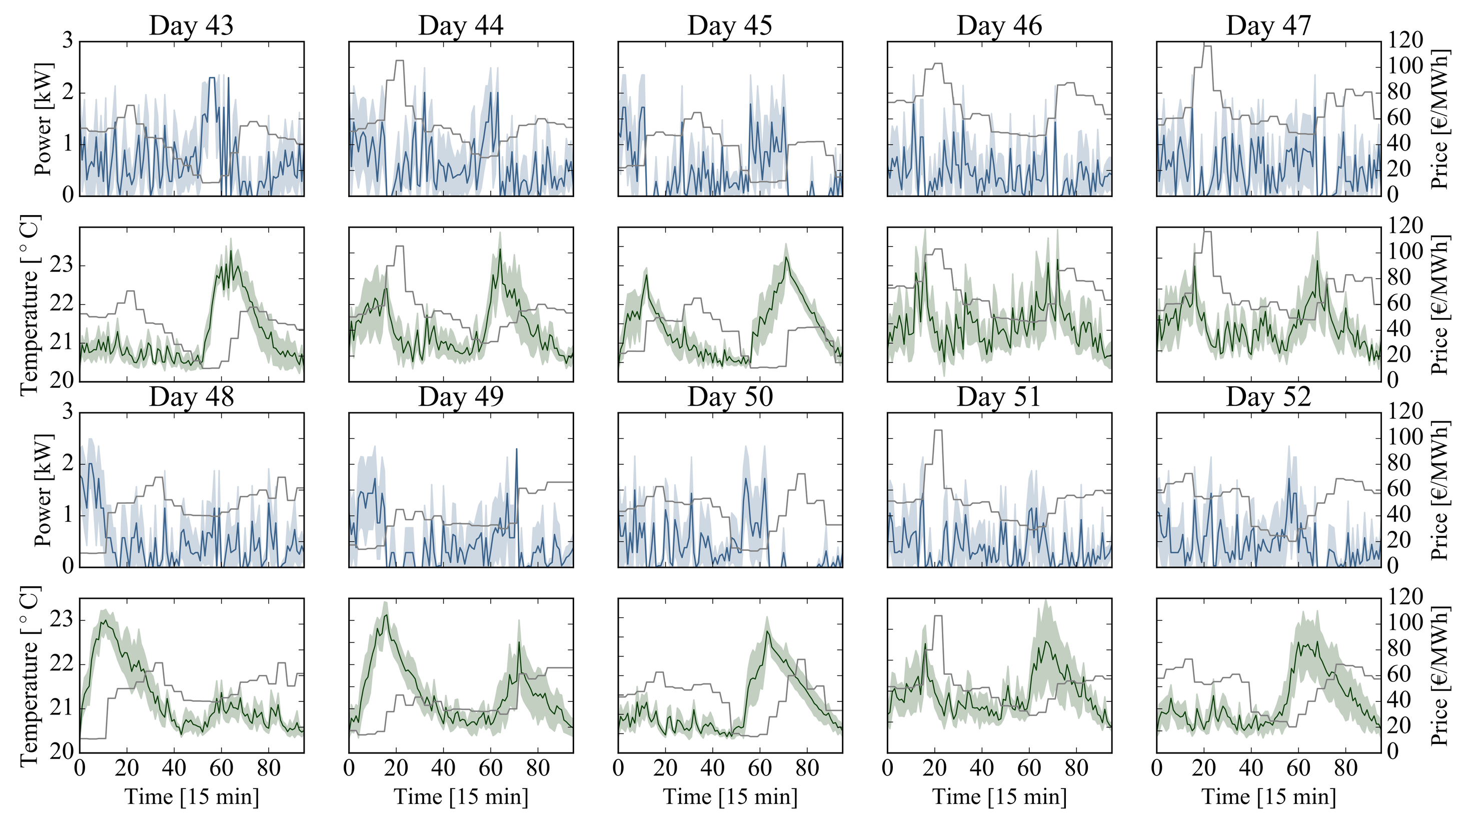

Conclusions

1. Using an LSTM network performed better than other techniques such as a neural network, convolutional neural network when sparse observations are used.
2. Its internal memory cell allows the LSTM network to process sequences of sparse observations and extract relevant features from it that can represent the underlying state of charge of the application.

Acknowledgements

This research is supported by IWT-SBO-SMILE-IT, funded by the Flemish Agency for Innovation, promoting Strategic Basic Research.

Paper @ `https://arxiv.org/pdf/1707.08553.pdf`