# Better Supply-Side and Demand-Side Energy Efficiency via Reinforcement Learning

Jesse Read (ack. to Alban Puech and Simo Alami)

29 July 2023
Seminar on Sustainable Development @ LIX

# Motivation: More Renewable Production & Less Demand

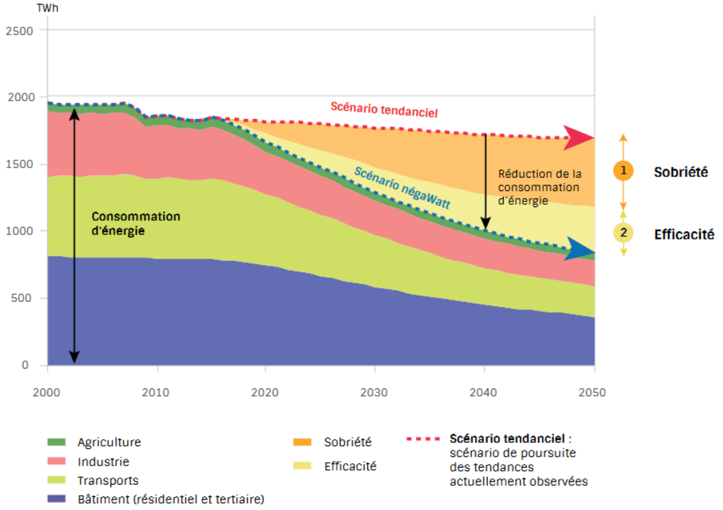


i.e., goal: being more efficient; and using less.

# Outline

# Introduction: Reinforcement Learning

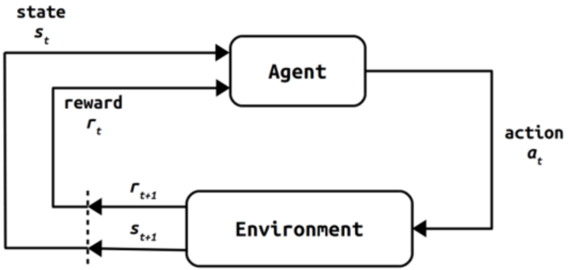


Gain (with discount factor $\gamma$): $G_t = \sum_{t=0}^{\infty} \gamma^t R_t$

Objective (to be maximized, wrt $a_t$):

$$\mathbb{E}[G_t \mid s_t, a_t]$$

where $s', r' \sim p(s', r' \mid s, a)$ (interaction with the Environment).

# Introduction: Wind Turbines

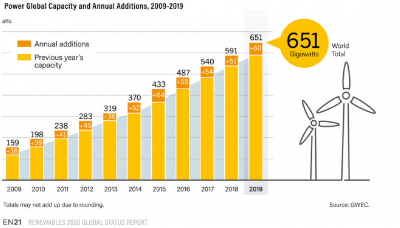


- Promising renewable energy source
- Can produce over 5MW each (per turbine)
- For a given turbine, energy output depends on wind speed, but also its operation

# Yaw Misalignment

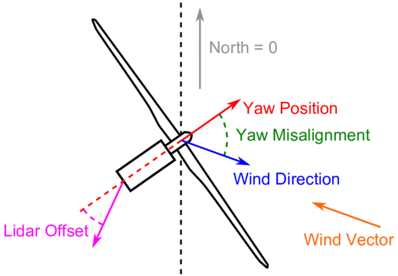


*Top View*
*All angles are drawn positive*

Yaw misalignment ($\epsilon$): difference between wind direction and nacelle (shown above)

Consequences on the power output, safety, durability.

$$P_\epsilon = P_* \cos^3(\epsilon)$$

# Existing Yaw Control Strategies

Conventional Yaw Control Algorithm:

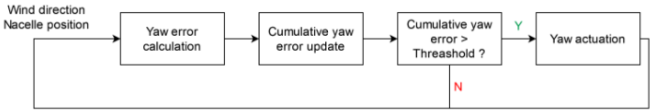


RL-solutions already exist but are currently focussed on long-term control or based on overly-synthetic simulations.

# Our Approach

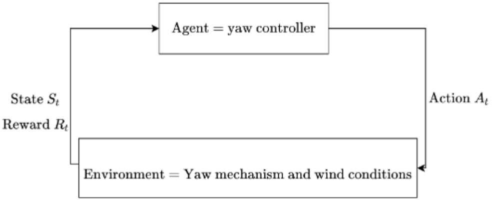


**State** $s_t = \{a_{\tau-1}, \epsilon_\tau, \phi_\tau, v_\tau \mid \tau = t, \ldots, t-j\}$
where yaw misalignment and wind speed and direction

**Action** $a_t \in \{\circlearrowright, \circlearrowleft, 0\}$ (turn or do nothing)

**Reward**

$$R_t = w \prod_{i=t-k+1}^{t} 1\{a_i = 0\} - (1-w)\epsilon_{t+1}^2 v_{t+1}^3$$

where $w$ the trade-off of preserving the yaw mechanism (left term), vs seeking better alignment (right term).

We use *Proximal Policy Optimization* (PPO).

# Results

- decreased the yaw misalignment by 5.5% to 11.2% compared to the conventional active yaw control algorithm.
- average net energy gain 0.31% to 0.33% algorithm
- this amounts to a 1.5k to 2.5k EUR per annum *per turbine*
- based on real world wind logs obtained from a RE power MM82 2MW turbine

Technology

**Software update for world's wind farms could power millions more homes**

An AI that predicts wind changes could boost wind turbine efficiency by 0.3 per cent, which globally would amount to enough extra electricity to keep a country running

By Matthew Sparkes

21 May 2023

A small efficiency gain for every wind turbine would lead to a big boost globally

Justin Paget/Digital Vision/Getty Images

Current algorithms track wind patterns and adjust the turbine blades to anticipate changes, but Alban Puech and Jesse Read at the Polytechnic Institute of Paris believe artificial intelligence can do better. They have trained a reinforcement-learning algorithm to monitor wind patterns and develop its own strategy to maintain the turbine at the correct angle. And because

(New Scientist magazine)

And [3] Puech and Read, ECML-PKDD, 2022.

# Outline

Introduction and Motivation: Complex Energy Systems

Part I (Being More Efficient): An Improved Yaw Control Algorithm for Wind Turbines

Part II (Using Less) Understanding Consumers, and Reducing Demand

# Problem: Energy Systems are Complex

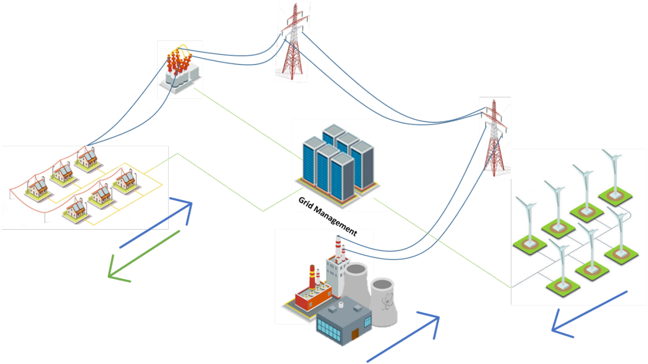


- Increasing use of renewable energy = more variability
- But demand and supply must be matched

We cannot easily control supply (we cannot control the weather).

Can we influence demand?

# Inverse Reinforcement Learning

- Reinforcement Learning: Find agent/policy $\pi : S_t \mapsto A_t$ to maximize [expected sum of] rewards $R_t = r(S_t)$; i.e., $\mathbb{E}[G_t]$.
- Inverse Reinforcement learning: Find the reward function $r : S \mapsto \mathbb{R}$ which agent/policy $\pi$ is maximizing (via $G$)

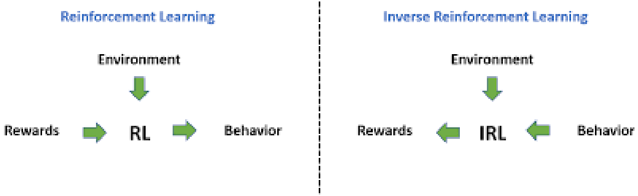


Having $r$, we might understand how an agent (consumer) $\pi$ will react (or *not* react) in a given environment $p(s' \mid s, a)$.

We could change the electricity network in a way that does not affect their comfort; or even 'nudge' them to act differently.

# First Challenge

- We can assume that people ($\pi$) *already* consume ($a_t = \pi(s_t)$) in order to optimize (max) their comfort ($r_t$).
- Challenge: learn user habits without infringing their privacy. Can only know the overall consumption of a given household/building.

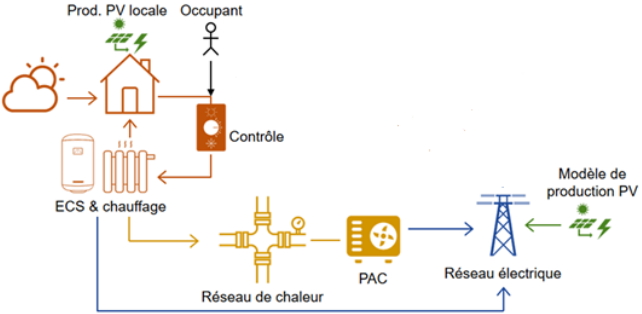

# Source Separation and Non Invasive Load Monitoring (NILM)

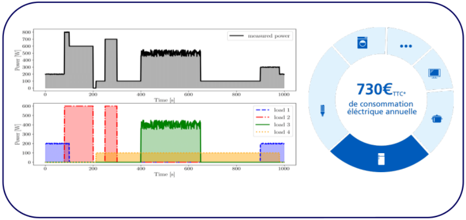


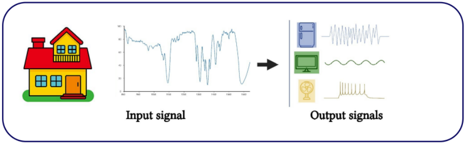


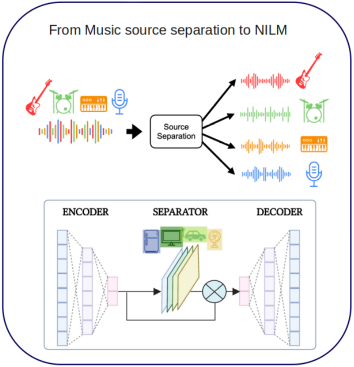

# Conv-NILM-net

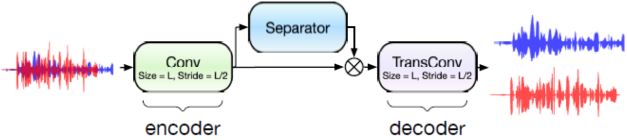


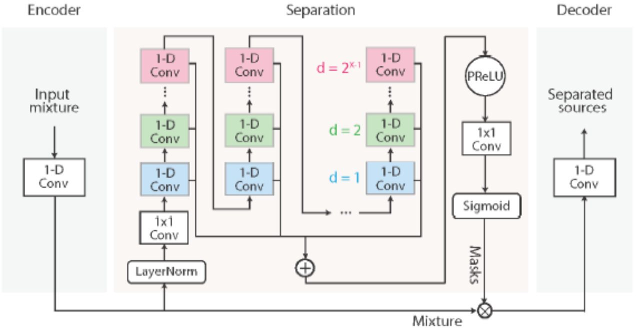


[2] Simo Alami et al., MLBEM 2022: ECML-PKDD Workshop on Machine Learning for Buildings Energy Management.

# Results

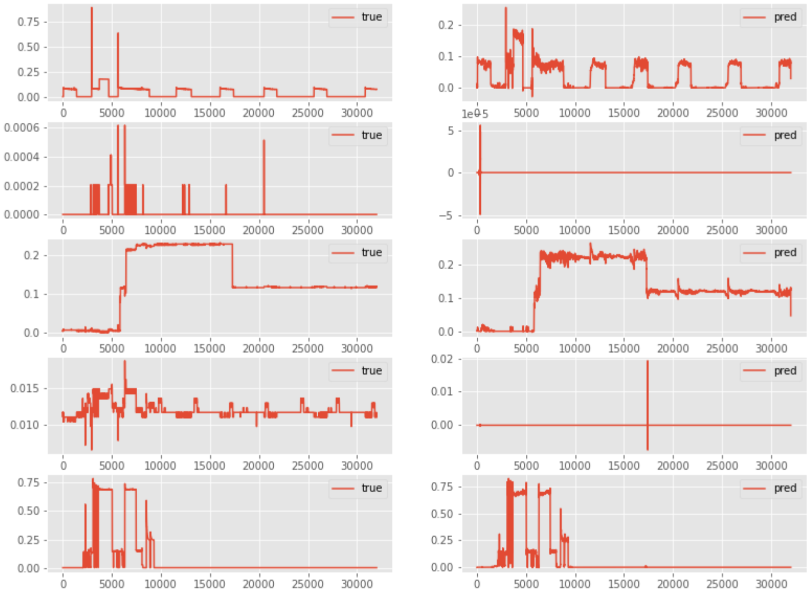

# Summary and Next Steps

So far we can learn habits in a non-intrusive way [2]; and consider a generalizable reward/metric $r$ [1].

Next steps: alter the environment $(s' \mid s, a)$ in a way that will not impact $r$; but at the same time, minimize 'global' costs: energy use, expense, renewable sources, . . . . We may even try to influence $r$ directly (find the easiest ways to nudge consumers).

# Outline

# Thank You. Questions?


Simo Alami Chehboune, Rim Kaddah, and Jesse Read.
Transferable deep metric learning for clustering.
In *IDA 2023: Advances in Intelligent Data Analysis XXI, 21st International Symposium*, pages 15–28, April 2023.

Simo Alami Chehbourne, Jérémie Decock, Rim Kaddah, and Jesse Read.
Conv-NILM-Net, a causal and multi-appliance model for energy source separation.
In *MLBEM 2022: ECML-PKDD Workshop on Machine Learning for Buildings Energy Management*, 2022.

Alban Puech and Jesse Read.
An improved yaw control algorithm for wind turbines via reinforcement learning.
In *ECML-PKDD 2022: 33rd European Conference on Machine Learning*, pages 614–630. Springer Nature Switzerland, 2023.
ADS Track.


And some images used in these slides sourced from: `https://blog-isige.minesparis.psl.eu/2018/07/24/le-scenario-negawatt-pour-la-transition-energetique/`
`https://www.evwind.es/2020/07/05/wind-energy-expanded-19-in-2019-with-around-60-gw-of-new-capacity/75563`
`https://towardsdatascience.com/inverse-reinforcement-learning-6453b7cdc90d`