AMLD EPFL

Wind Farm Dynamic Yield Optimization using Reinforcement Learning

Dr. Giorgio Cortiana
Head of Energy Intelligence,
Global Advanced Analytics and AI @ E.ON

e·on

Lausanne – Switzerland, January 27th 2020

Wind Turbines Basics

Main operational parameters

- Yaw angle (orientation with respect to incoming wind direction)
- Blade pitch angle (to modulate the amount of energy extracted from the incoming wind)

Turbine Wakes

- Downstream wind flow structures characterized by reduced wind speed, higher turbulence. Negatively impact the power output of a wind park. Can be mitigated by:
  - layout optimization
  - coordinated wind park control (wake management: pitch/yaw)

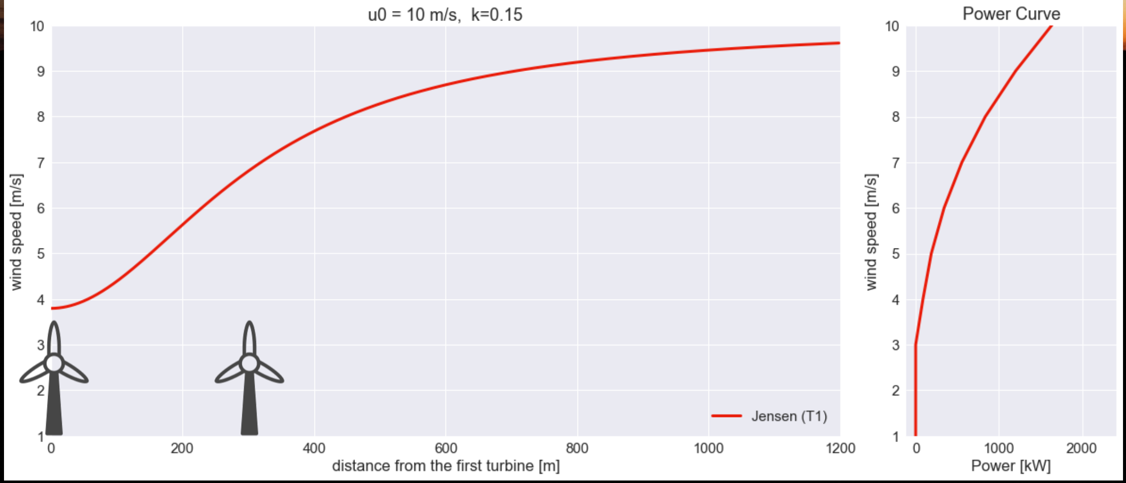

Wind Turbines Basics

Main operational parameters

- Yaw angle (orientation with respect to incoming wind direction)
- Blade pitch angle (to modulate the amount of energy extracted from the incoming wind)

Turbine Wakes

- Downstream wind flow structures characterized by reduced wind speed, higher turbulence. Negatively impact the power output of a wind park. Can be mitigated by:
  - layout optimization
  - coordinated wind park control (wake management: pitch/yaw)

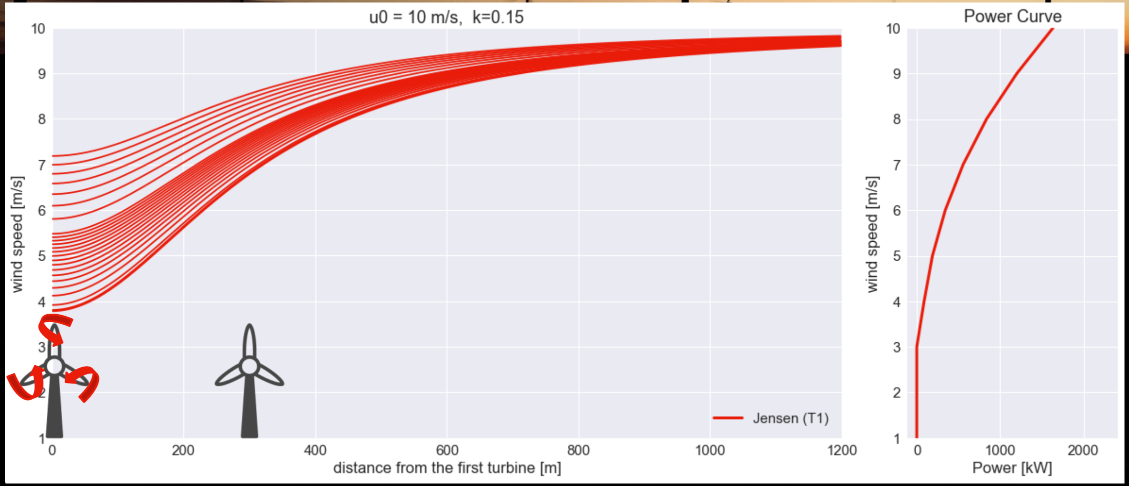

Wind Turbines Basics

Main operational parameters

- Yaw angle (orientation with respect to incoming wind direction)
- Blade pitch angle (to modulate the amount of energy extracted from the incoming wind)

Turbine Wakes

- Downstream wind flow structures characterized by reduced wind speed, higher turbulence. Negatively impact the power output of a wind park. Can be mitigated by:
  - layout optimization
  - coordinated wind park control (wake management: pitch/yaw)

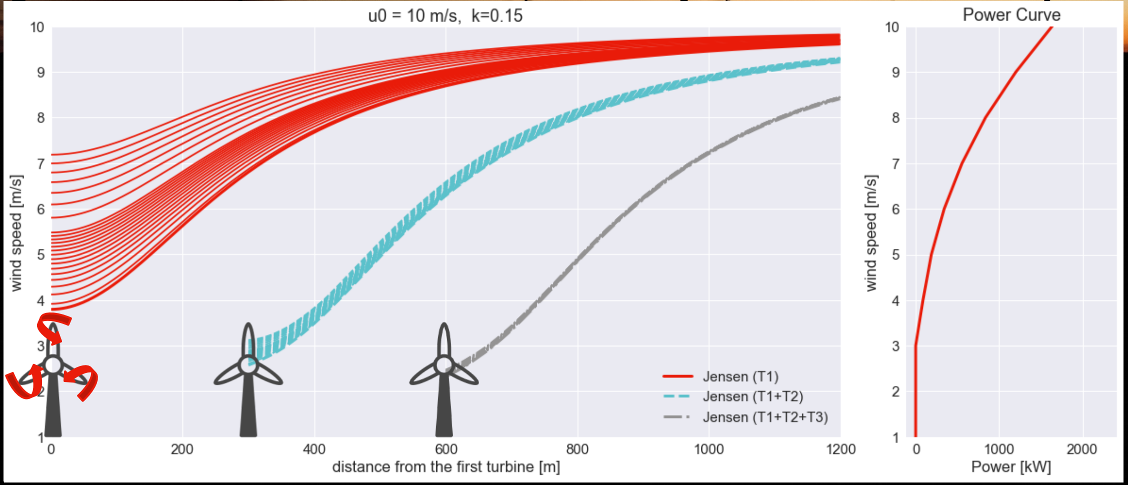

Reinforcement learning (RL) for wake management

RL is specially designed to interact with the its environment and learn (as fast as possible) from the interactions with it.

Could Computational Fluid Dynamics be used instead?

Yes… BUT

- Computationally very demanding and costly, not viable for online control
- Wake model dependent
- Wind park conditions, neighborhoods and environment are constantly changing

We needed a data-driven, self-learning and adaptive approach, and that could run on commodity hardware

Successful applications

- Computer/board games (go)
- Robot control
- Machine tuning
- Other optimization problems (NN setups, complex cooling systems…)

RL Architecture

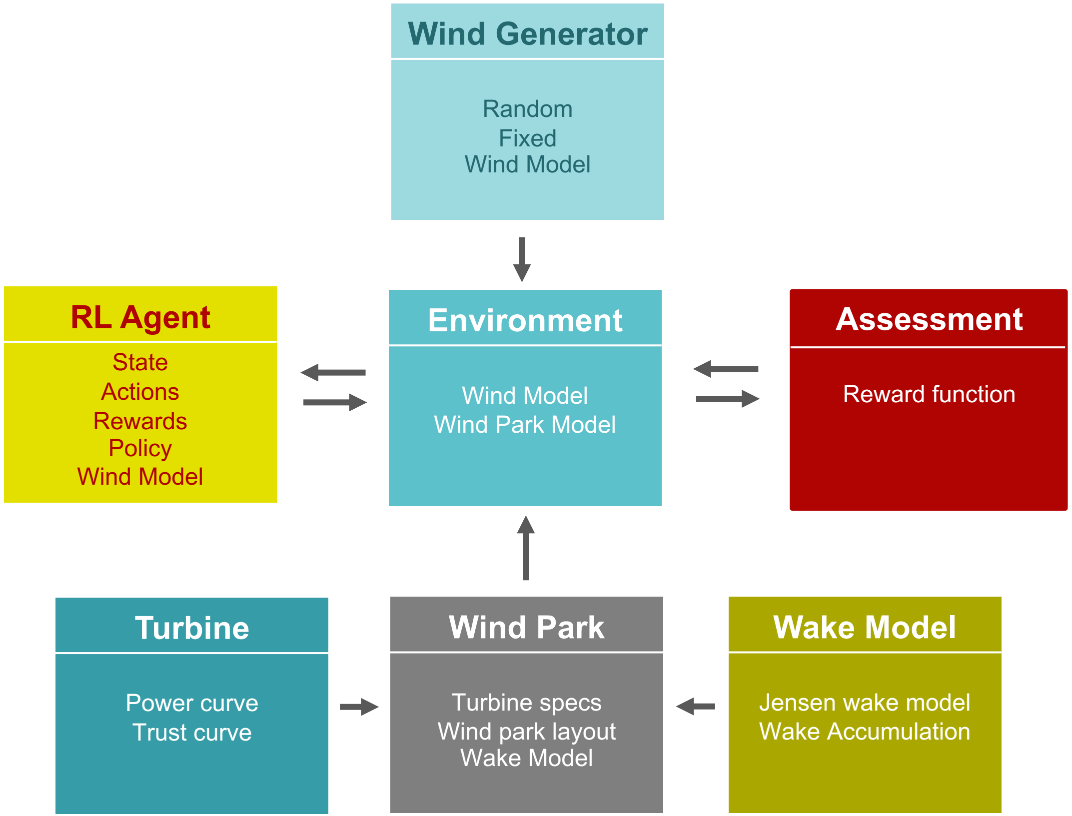

For RL algorithm development and test, a complete wind park simulation suite has been developed

RL Architecture

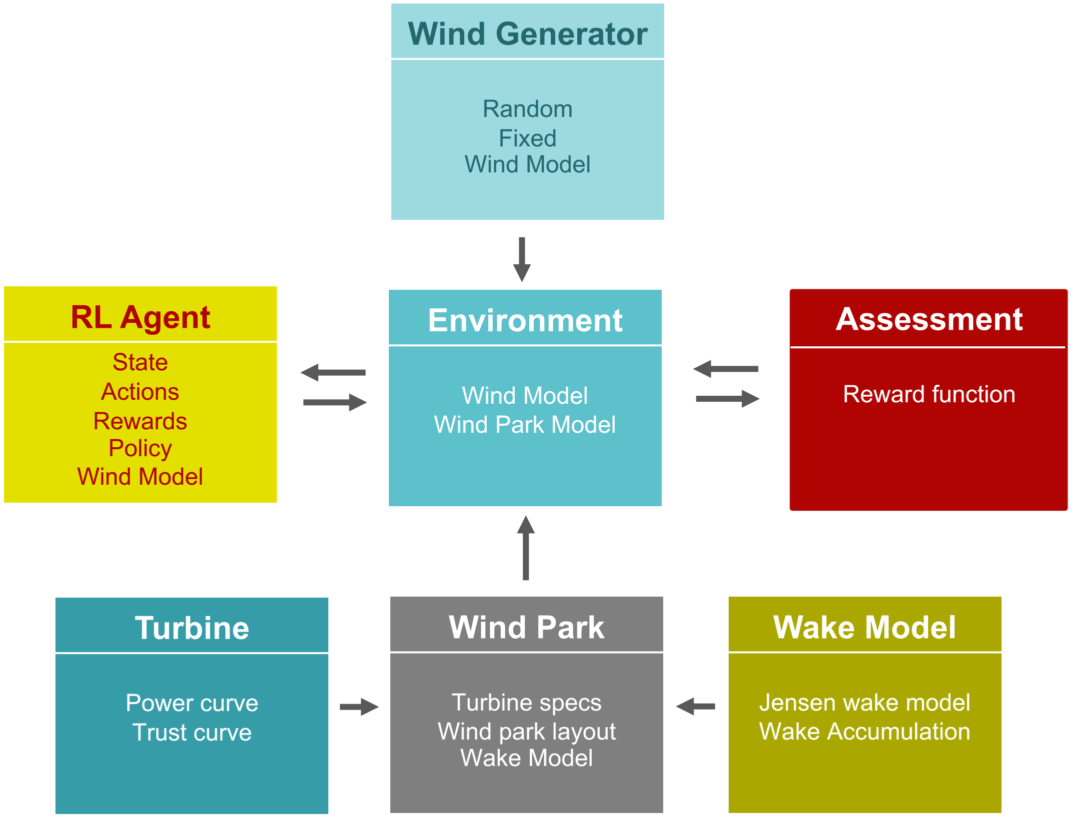

The goal is to find the best operational settings of the wind turbines (policy), depending on the wind conditions

The RL agent explores the environment and updates the policy based on the experience gained (reward function approx.)

The state is defined by the combination of wind-direction and wind-speed, turbulence intensity, curtailment level (blade pitch angle settings)

The reward is the gain/loss in power with respect nominal operation (no curtailment) for the same wind conditions*

- Taking into account the wake propagation time

Pilot implementation

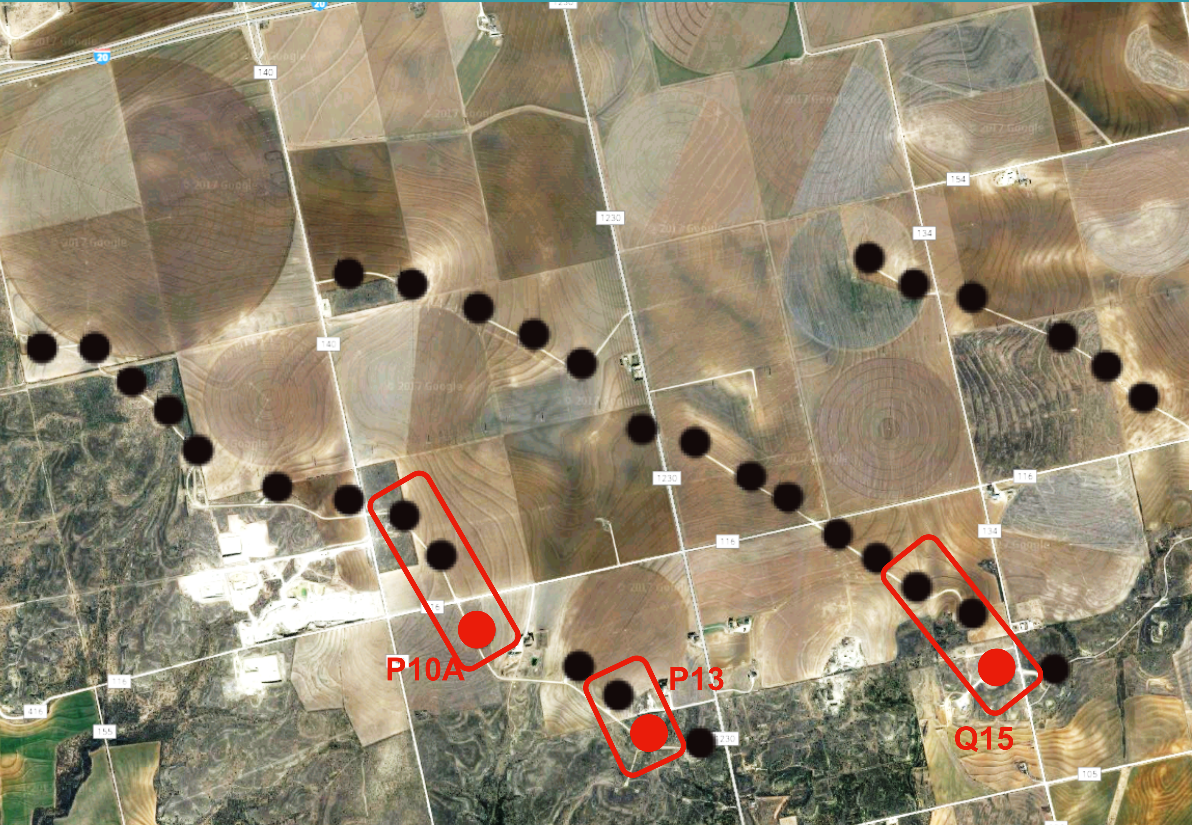

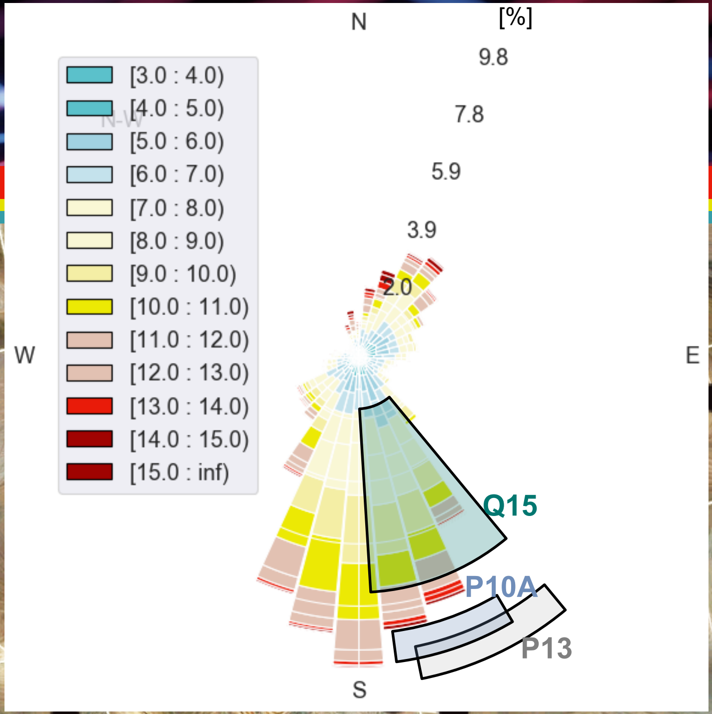

Pilot implementation

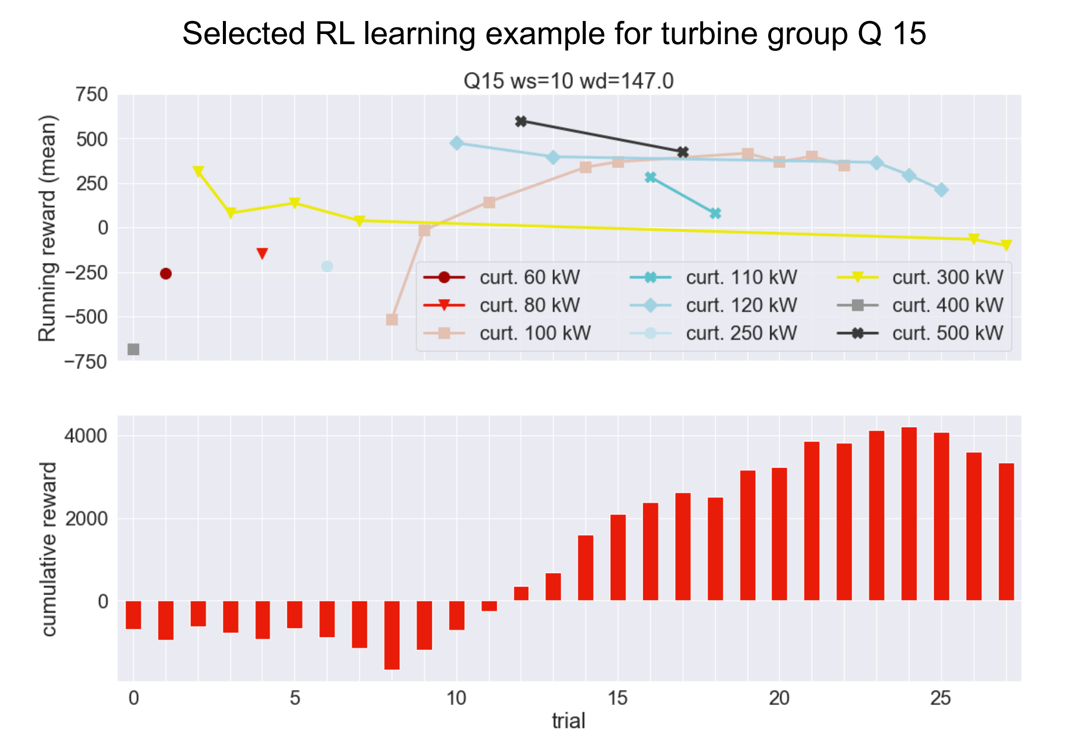

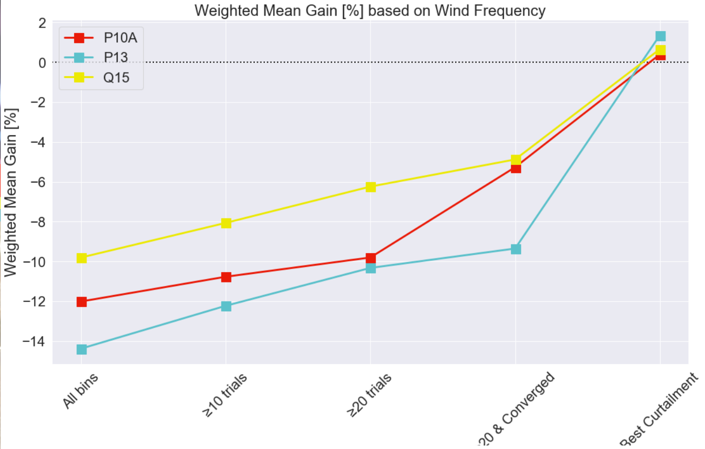

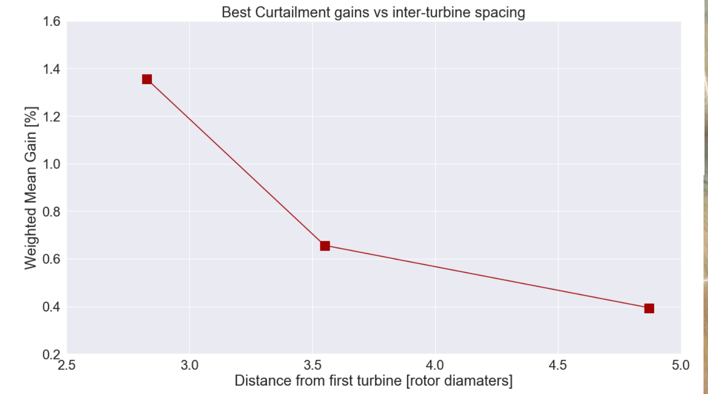

We exploited Reinforcement Learning (RL) to dynamically optimize the energy yield of a wind farm

The developed RL framework has been successfully applied a real and uncontrolled environment, demonstrating its potential to solve complex real-life problems

Observed energy gains are small, and further algorithm improvements (yaw angle control, and better handles on syst. unc.) have the potential to improve the performance and open the doors to full scale roll-out

Pilot results prove that it is possible to influence, and control wind turbines wakes based on a full data driven approach, using state of the art AI technology

Thank you for your attention

Acknowledgments: The project here described was realized as a joint collaboration between the Energy Intelligence team at the E.ON Advanced Analytics and AI unit and the colleagues from the E.ON renewables business.