Learning To Run a Power Network Competition

Benjamin Donnot[†] Isabelle Guyon[‡ •]

Antoine Marot [†]

‡ UPSud and Inria TAU, Université Paris-Saclay, France.

• ChaLearn, Berkeley, California. † RTE France.

September 16, 2019

Abstract

We present the results of the first edition as well as some perspective for a next potential edition of the "Learning To Run a Power Network" (L2RPN) competition to test the potential of Reinforcement Learning to solve a real-world problem of great practical importance: controlling power transportation in power grids while keeping people and equipment safe.

1 Context

Power grids transporting electricity across states, countries, or continents, are vital components of modern societies, playing a central economical and societal role, by supplying power reliably to industry, services, and consumers. Blackouts may lead to significant losses and delay in public services and strategical industries, *de facto* freezing society. Grid operators are in charge of ensuring that a reliable supply of electricity is provided everywhere, at all times.

In Western Europe or Northern America, powergrid has been built decades ago. They were mainly developed to ensure the power production (often centralized and controllable) could be securely transmitted to de-centralized loads such as heavy industries, hospitals or final consumers. When the powergrid was too weak in some area e.g. due to new industries installing near a city or a growth in the number of inhabitants, the preferred choice used to be to build new *hardware* (e.g. new transportation lines or transformers).

Today, power grids are facing new challenges. The total demand for energy stops increasing and the public acceptance for this *heavy* and *costly* infrastructures has become very low. In this context, it is really difficult to build new powerlines. Yet, at the same time, power production becomes more and more decentralized with the emergence of renewables energies (wind and solar power production), new consumers behaviors (self-consumption, installation of solar panel on their roof), etc. All these factors combined make the power grid less and less fitted to these new emerging usages while it is difficult to develop it.

A possible solution to continue ensuring a high quality of service in such context could be to invest in software e.g. having smart automatons (or *controllers*) that could leverage the data available and take decisions to ensure the security of the powergrid [1]. In the community of power systems, optimal control methods [2] have been explored but with limited success until now. With the latest breakthroughs from Alpha Zero at various board games [3] and Libratus at poker [4], Reinforcement Learning (RL) seems a promising new avenue to develop an artificial agent able to operate a complex power system in real-time. On the other hand, RL community focuses more and more on problems of the real world [6] so we believe the modeling of this critical importance problem, as well as the methods developed to solve it, could also benefit this community. We propose a series of competitions "Learning to run a Power Network" using a challenging game-like environment we have designed, to enable the RL community to join forces with the power system community in order to tackle this burning smart grid issue at larger scales.

2 Feedbacks and Future Competitions

This competition is inspired by the "Learning to run" competition at NIPS 2017[1] and has been presented in the CiML Neurips 2018 workshop [6].

The challenge was part of the "International Joint Conference on Neural Networks" https://www.ijcnn.org/ competitions. Inspired by the "Learning to run" competition, this first L2RPN competition, based on the pypownet[2] open-sourced platform relying on OpenGym RL framework, ran over 6 weeks starting on May 2019 over the Codalab challenge platform. 100 participant signed in, 15 of which were particularly active with many submissions every week. Among them, 7 managed to outperformed the baselines most of them using either reinforcement learning, expert system or a combination of both. For this first edition, the managed powergrid was limited in size, yet the behavior of the winning agents showed similar properties to what humans are doing today, such as acting only on some specific nodes of the system and only at some specific time steps.

These results are encouraging for the next competition held in 2020. Contestants will be asked to manage a bigger more realistic powergrid based on an improved version of the environment and more realistic data generation processes. We will pursue our effort into developing a complete environment for the challenge that will allow to integrate planned maintenance (obstacles that make the control task harder) and a wider "action space" allowing the controller to perform not only discrete actions (as it was the case in the first edition) but also continuous actions (such as changing the amount of power produce by a thermal power plant) thus making this real world problem even more interesting for the RL community.

More challenging and more realistic the design of this (these) new competition(s) will be made possible thanks to the help from experts of both "machine learning" and "power system" communities made possible by the organization and the communication that has taken place around this first edition.

[1] `https://www.crowdai.org/challenges/nips-2017-learning-to-run`

[2] Available on GitHub at `https://github.com/MarvinLer/pypownet`

Acknowledgments

This work would not have been possible without the support of RTE (French Power Grid Operator), INRIA as well as Chalearn, and without the contribution of many advisors and students, including Patrick Panciatici, Marc Schoenauer, Louis Wehenkel, Olivier Pietquin, Joao Araujo, Marvin Lerousseau, and Kimang Khun.

References

[1] Benjamin Donnot, Isabelle Guyon, Marc Shoenauer, Patrick Panciatici, Antoine Marot, Introducing machine learning for power system operation support, IREP Conference 2017

[2] Á. Lorca & al, Multistage adaptive robust optimization for the unit commitment problem, Operations Research, 2016.

[3] Silver, David, et al. "Mastering the game of Go without human knowledge." Nature 550.7676 (2017): 354.

[4] N. Brown and T. Sandholm, Safe and Nested Subgame Solving for Imperfect-Information Games, NIPS 2017, https://arxiv.org/abs/1705.02955.

[5] A. Marot, D. Donon, I. Guyon, B. Donnot, Learning To Run a Power Network Competition, CiML Workshop @NeurIPS 2018 available at `https://docs.google.com/a/chalearn.org/viewer?a=v&pid=sites&srcid=Y2hhbGVhcm4ub3JnfHdvcmtzaG9wfGd4OjM4OWI`

[6] T. H. G. Dulac, D. Mankowitz. Challenges for real-world reinforcement learning, `http://arxiv.org/abs/1904.12901`