

Fully funded Ph.D. Position:

Integrated Machine Learning and Optimization for Dynamic Decision-Making in Continuously Monitored Systems

The [chair of Risk and Resilience of Complex Systems \(RRCS\)](#), Industrial Engineering Laboratory (LGI) within the French engineering school [CentraleSupélec](#) at the [University of Paris Saclay](#) is looking for a doctoral student working on integrated machine learning and optimization approaches for dynamic decision-making problems in continuously monitored systems, applied to operation and maintenance planning of systems such as fleets of wind farms, high-speed trains, and telecommunication antenna/data-centers.

The chair RRCS is funded by three prominent actors of the French industry: EDF (Power supply provider), Orange (Telecom), and SNCF (Railway). The chair has two main missions: i) to ensure a level of scientific excellence and ii) to promote the transfer of knowledge and technology. The objective for the partners is to share common concerns, contribute to the development of pooled models, and exchange use cases.

The Ph.D. candidate will be directed by Associate Professor [Yiping Fang](#), Professor [Anne Barros](#), and Associate Professor [Adam F. Abdin](#). She/he will have the opportunity to build a rich network by interacting with the three partners at the top management level and the operational ones. She/he will also have great opportunities to stay in several other prestigious research institutes or universities in France or abroad, including with other laboratories at CentraleSupélec (especially the Laboratory of Signal and Systems and the Mathematics and Informatics Lab), INRIA Paris, École Normale Supérieure Paris, Delft University (Netherlands), Rutgers University (USA), University of New South Wales (Australia), and the University of Edinburgh Business School (UK) within the collaboration network of the Chair project.

Research problem statement

Numerous private and public organizations heavily rely on physical assets to produce goods or deliver services. Examples of such assets include power plants (nuclear, wind turbines, etc.) in energy systems, base stations and data centers in telecommunication, and fleets of trains in railway transport, among many others. The operations of those systems inevitably suffer from uncertainties such as random shocks and deterioration which not only affects system productivity but also increases the risk of unexpected failures, resulting in costly downtimes.

Recent advances in data collection (through advanced sensor techniques) and algorithmic capabilities in data processing for decision making have resulted in renewed potential to streamline the operation and maintenance activities of industrial systems, a careful planning of which is effective to mitigate the severe consequences of asset deterioration, ensuring the reliability, resilience, and productivity of those systems. Various technical issues arise during the development of optimal operation and maintenance policies for modern complex industrial systems. The two most commonly encountered issues are *partial information* and *high dimensionality*. Firstly, information obtained from condition monitoring sensors or field human inspection is inherently noisy, i.e., the systems are *partially observable*. Also, how an asset degrades is inherently stochastic, and the degradation dynamic can often not be accurately estimated due to the lack of historical data or heterogeneous operational environments. This means that decision-makers must plan maintenance interventions under uncertainty about the true operational state and degradation dynamic of the system. Secondly, a high dimensional state (and action) representa-

tion is necessary for modeling real-world systems, especially those composed of multiple heterogeneous subsystems (e.g., a set of geographically distributed wind farms, each composed of tens of wind turbines), adding another layer of complexity to maintenance decision-making.

The focal dynamic operation and maintenance planning problems are essentially about the sequential allocation of the limited resources, corresponding to the classical *restless multi-armed bandit problem* (RMABP) which is often studied in the framework of Partially Observable Markov Decision Processes (POMDP) [1]. Conventional heuristics, such as Whittle's index policies, are generally suboptimal for the RMABP, as this is known to be PSPACE-hard [2].

Furthermore, *economic* and *functional dependencies* often exist in practice among the maintenance interventions of different units of a system, e.g., induced by the fixed intervention/deployment cost [3] and system-level functionality/reliability/availability constraints [4]. The dependencies make methods like the linear programming relaxation of decomposable Markov decision processes (MDP) [5] likely not suitable for our problems, and an alternative development must be pursued.

This thesis aims to develop credible models and efficient algorithms for data-driven operation and maintenance optimization of large-scale multi-unit industrial systems. From the modeling perspective, we will focus on strongly coupled RMABP by considering system-level reliability/availability constraints. Also, the idea of integrating online learning of the degradation dynamics (based on online-monitoring data) and maintenance planning, such as the framework of "smart predict-then-optimize" [6], could be in the scope. From the algorithmic perspective, recently, the operations research community started to focus on systematic ways to analyze and solve optimization problems of combinatorial and dynamic natures with the help of machine learning [7], and several promising architectures [8-9] have been proposed. Thus, the idea of exploiting cutting-edge machine learning architectures for solving large-scale RMABP problems will also be emphasized.

Keywords: *Data-driven sequential decision-making; restless multi-armed bandit problems; machine learning for optimization; condition monitoring; operation and maintenance optimization; smart predict-then-optimize*

Your profile

You will work within an interdisciplinary project, at the interface of operations research, machine learning, and industrial engineering, for which we expect that:

- You are proactive and highly motivated, having an MSc degree (or equivalence) in a quantitative discipline (such as applied mathematics, computer science, or engineering) from a recognized university.
- You have a strong background in quantitative modeling (statistics, mathematical programming and optimization, and machine learning) and computer programming.
- Professional command of English (both written and spoken) is mandatory. French is not mandatory, but surely a plus.
- You will need to enjoy working in a dynamic and international environment with other doctoral students and postdocs.

How to apply?

We look forward to receiving your application with the following documents compressed in **one PDF file** and sending it as an attachment to phd.candidate.rrcs@gmail.com with the title "[Your name]-RRCS 2023 PhD Application":

- Curriculum Vitae, max. two pages

- Motivational Letter, max. two pages
- Transcript of academic records (including detail of grades and rankings when available)
- Two reference letters with contact details

Please note that the position will start as soon as possible, and we exclusively accept applications submitted through phd.candidate.rrcs@gmail.com before **June 15, 2023**.

Further information about the RRSC team and LGI Laboratory can be found on the website. Questions regarding the position should be directed to Dr. Yiping FANG (yiping.fang@centralesupelec.fr), Dr. Adam F. Abdin (adam.abdin@centralesupelec.fr), and Professor Anne Barros (anne.barros@centralesupelec.fr) (**no applications**).

References

- [1] Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A), 287-298.
- [2] Papadimitriou, C. H., & Tsitsiklis, J. N. (1994, June). The complexity of optimal queueing network control. In *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory* (pp. 318-322). IEEE.
- [3] Keizer, M. C. O., Flapper, S. D. P., & Teunter, R. H. (2017). Condition-based maintenance policies for systems with multiple dependent components: A review. *European Journal of Operational Research*, 261(2), 405-420.
- [4] Do, P., Vu, H. C., Barros, A., & Bérenguer, C. (2015). Maintenance grouping for multi-component systems with availability constraints and limited maintenance teams. *Reliability Engineering & System Safety*, 142, 56-67.
- [5] Bertsimas, D., & Mišić, V. V. (2016). Decomposable Markov decision processes: A fluid optimization approach. *Operations Research*, 64(6), 1537-1555.
- [6] Elmachtoub, A. N., & Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1), 9-26.
- [7] Bengio, Y., Lodi, A., & Prouvost, A. (2021). Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2), 405-421.
- [8] Bertsimas, D., & Stellato, B. (2022). Online mixed-integer optimization in milliseconds. *INFORMS Journal on Computing*, 34(4), 2229-2248.
- [9] Khalil, E., Dai, H., Zhang, Y., Dilkina, B., & Song, L. (2017). Learning combinatorial optimization algorithms over graphs. *Advances in neural information processing systems*, 30.