

Potential and challenges of AI-powered decision support for short-term system operations

Viebahn, Jan; Naglic, Matija ; Marot, Antoine; Donnot, Benjamin; Tindemans, Simon H.

Publication date

2022

Document Version

Final published version

Citation (APA)

Viebahn, J., Naglic, M., Marot, A., Donnot, B., & Tindemans, S. H. (2022). *Potential and challenges of AI-powered decision support for short-term system operations*. Paper presented at CIGRE paris Session 2022, Paris, France.

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

**Potential and challenges of AI-powered decision support
for short-term system operations**

**Jan VIEBAHN^{1*}, Matija NAGLIC¹, Antoine MAROT²,
Benjamin DONNOT², Simon H. TINDEMANS³**

¹TenneT TSO, The Netherlands

jan.viebahn@tennet.eu, matija.naglic@tennet.eu

²RTE TSO, France

antoine.marot@rte-france.com, benjamin.donnot@rte-france.com

³TU Delft, The Netherlands

S.H.Tindemans@tudelft.nl

SUMMARY

Given the increasing need to meet the new operational requirements of power systems and prepare for the future, adaptation of cutting-edge Artificial Intelligence (AI) technologies in the operational processes is paramount to timely meet the challenges. The focus of this paper is on applying AI in power system operations, in particular for the development of decision support tools. First, the paper elaborates on the decision-making process of the power system operators and presents a mirroring digital framework consisting of AI and control theory to mimic sequential decision making of the operators. Next, a demonstrating example in the field of congestion management is presented by a real-world AI use-case at TenneT TSO. The paper continues with state-of-the-art on sequential decision making applied to congestion management and elaborates on research challenges when applying AI to the power systems problems. Finally, the paper elaborates on the enabling capabilities with focus on people, data, and platform pillars an organisation needs for mastering the development as well as maintenance of AI solutions, and proposes a cyclic (agile) process approach to decrease time from development to actual deployment and cooperation between research and industry organisations.

KEYWORDS

Artificial Intelligence - Control Centre - Decision-Making - Decision-Support - R&D Ecosystem - From Research to Deployment - Congestion Management - L2RPN Competitions.

1 INTRODUCTION

The European electricity system and associated energy market systems are transforming at a rapid pace due to the decarbonisation of the electricity system to meet climate goals and the facilitation of cross-border energy markets. As a consequence, power system operation is becoming increasingly complex, resulting in a growing need for reliable real-time decision support tools to assist the power system operators. In particular, the need for time-continuous decision support that offers effective and persistent decisions over a time horizon is becoming more urgent. It is expected that conventional EMS/SCADA functionalities of control centres, which largely focus on individual snapshots of the power system state instead of sequences of snapshots, will become insufficient for reliable and affordable power system operation in the future. Nevertheless, the ongoing wide-spread digitalisation in the power system enables opportunities to advance the control room systems and tools by combining the best of human and computer intelligence [1-3].

Recently, the field of artificial intelligence (AI) went through a deep learning revolution, which drastically enlarges its potential for real-world application. Deep-learning AI systems are able to digest large volumes of information, memorize historical datasets, and learn to quickly infer effective actions in context by taking time horizons and forecast uncertainties into account. In particular, the recent advances in deep reinforcement learning (RL) [4-5] demonstrate that it's often possible to achieve performance that is comparable to or exceeds that of humans for sequential decision making in complex systems. RL has been applied with great success in various sectors, for example, robot open-walk [6], self-driving cars [7], autonomous navigation of stratospheric balloons [8], data-center cooling [9]. RL has also been considered for decision support in power systems [10], but this application has remained as a research niche without substantial real-world impact and deployment to date. Yet, these cutting-edge approaches create unprecedented opportunities for advancing of the existing EMS/SCADA functionalities and designing advanced decision support tools for power system operators with the aim to improve the efficiency and safeguard the reliability, security, and resilience of power systems.

With this paper, the authors inform the power system community about potential and challenges of AI-based decision support for power system operations with a focus on sequential decision making. More specifically, in section 2 we give a more precise notion of AI and indicate more precisely why AI is a relevant technology for the evolution of control centres. In section 3 we share as a motivating example results from a real-world use case at TenneT TSO. Subsequently, in section 4 we summarize the state-of-the-art research on sequential decision making applied to congestion management, followed by section 5 in which research gaps for the application of AI tools in system operations are identified. In order to overcome the research gaps, we propose in section 6 a procedural embedding of AI development in the business practice of system operators as a crucial step for the eventual successful deployment of AI tools in the system operation. Section 7 provides a short summary.

2 The operator's decision-making and AI

Typical control centres are core places of the power system, providing to groups of human operators the necessary working environment to remotely monitor and operate the power system in real time. The operators interact with the power system, on the one hand, by observing the continuously changing power system state, and, on the other hand, by manually performing a broad range of control actions on the grid. The actions can be discrete like line switching and substation reconfiguration, or continuous like adjusting voltage setpoints or power dispatches of generators, and many more.

Each of the control actions, performed by power system operators, usually not only affects the current state of the power system but also the future state and availability of future control actions, that is, short-term actions can have long-term consequences. As a result, the decision problem of power system operators is typically a **sequential decision problem** in which the current decision can affect all future decisions. Moreover, due to possible nondeterministic changes of the power system state (e.g., due to




Figure 1: (Left) In black the typical high-level framework of the operator’s decision-making process. In blue an additional decision support interface offering the operator a set of functionalities that enable shared knowledge representations and situational awareness over relevant time-horizons and allow for contextual collaborative decision-makings through adapted interactions. AI is well-suited to provide some of these functionalities. (Right) In black the general framework of AI and control theory consisting of an agent or controller that executes actions in a complex environment and receives feedback from the environment via observations [13]. In green the additional elements present in RL in which the agent also receives a reward signal that provides an indication of the quality of its behaviour. An RL agent is represented by a policy that maps observations to actions, and an RL algorithm uses sequences of observations and rewards in order to iteratively find an optimal policy that maximizes the sum of rewards [4-5]. The right figure is a modified version of [14].

unplanned outages or the intermittent behaviour of renewable energy sources) and different sources of error (e.g., measurement errors, state estimation errors, flawed judgement) the operators need to **handle uncertainty** in their decisions. Finally, operational decisions must often be made quickly, under **hard time constraints**. The energy transition increases both the relevance of this challenge as well as the complexity of time horizons, uncertainty, and time constraints in control centres [11].

Control centres need to adapt to the rapid changes in operational requirements of the power system. When improving the functionality of the control centre systems, it is crucial to consider soft frameworks of the human decision-making processes in order to develop efficient and user-friendly tools for power system operators [2, 12]. As visualised in Figure 1 (left, black), these frameworks typically consist of four high-level steps, namely, (i) perception of elements in the current situation; (ii) processing of observations or comprehension of their meanings and relations (information inference); (iii) action exploration or projection of future states with the given knowledge; (iv) implementation of actions in real-world environment. Different models typically vary in the amount of detail shown for each step and in the number of feedback loops or shortcuts between the different steps [2,12].

AI represents a technology that is well suited to support the evolution of control centres since AI in large parts inherently deals with sequential decision making under uncertainty. Furthermore, the standard model of AI [13], as shown in Figure 1 (right, black), is neatly in line with frameworks for decision-making by operators, as shown in Figure 1 (left, black). More precisely, AI is concerned with building effective agents¹ that interact with complex environments in order to effectively perform tasks, according to one or more objectives. That is, AI is not just concerned with perceiving, understanding, or predicting a complex environment but also with building intelligent entities – machines that can compute how to act effectively and safely in a wide variety of unprecedented situations [13]. Moreover, AI agents can be designed as decision support tools, as indicated in Figure 1 (left, blue), providing to power system operators (possibly in an interactive manner) a set of actions to be potentially used in the

¹Note that a plain agent is just something that perceives and acts in an environment.





Figure 2: (Left) Part of the 110 kV grid in the North of the Netherlands. Red circles indicate substations and power lines are indicated in black. The use case is focussed on the sub-grid to the East where the superimposed green circles indicate the relative amount of solar in-feed per substation in a near-future scenario. (Right) Maximum line loading across the sub-grid as a result of various control strategies during a day of a near-future scenario. The control strategies are: fixed reference topology (red), greedy agent (green), RL agent proposing two substation reconfigurations (cyan), RL agent proposing three substation reconfigurations (blue).

given situation, with the operator still being in charge of deciding which actions are acceptable and should be implemented [2,12]. Several design principles are possible for human-in-the-loop AI [3, 15] which represents a growing research field (see also section 5).

We note that AI is not limited to a specific technique. AI solutions typically employ a suite of different approaches ranging from efficient brute force search and expert rules/systems to probabilistic reasoning and machine learning (ML) [13]. In particular, reinforcement learning (RL) represents an approach that is specifically tailored to sequential decision making under uncertainty (see also Figure 1, right-green) and receives a lot of attention recently [4-5]. Finally, we emphasize that the concept of a controller in control theory (especially for stochastic optimal control, well-known in power systems research) has an identical role as that of an agent in AI [13]. Hence, the challenge of bringing the two fields of control theory and AI closer together [16] is less a conceptual problem but a practical/organizational problem, as we further detail below in sections 4 and 6.

3 A motivating example

As a concrete real-world example, we share results of a use case at TenneT TSO on congestion management, which is concerned with the task of fully mitigating overloads in the power grid. The use case is focussed on a part of the 110 kV grid in the Netherlands that has limited transport capacity and redispatch options (Figure 2, left) [17]. Moreover, the land is relatively cheap leading to a sharp increase in requested customer connections of solar parks. Consequently, in the near future, operators will be confronted with new (more generation-dominated) flow patterns and they will need to operate the grid with an increasing number of interventions and even closer to its limits. Figure 2 (right) illustrates that large overloads are expected in a near-future scenario when operating the grid in the reference topology (all busbar couplers closed).

In the use case a first step towards AI-based decision support is developed to increase grid utilization and to unlock unexploited flexibilities. Due to the limited redispatch options in this region, the question is: Are topological actions capable of fully mitigating these overloads? Dynamic grid topology reconfiguration is an interesting option for system operators since it is a cost-efficient and flexible solution for congestion management that uses existing infrastructure. But it is still beyond the state-of-the-art to optimally control the grid topology “at scale” due to the problem’s nonlinear and discrete combinatorial nature leading to a large search/optimization space [18,19]. The traditional optimization formulation of identifying optimal topology at a given snapshot is a large-scale non-convex mixed-integer non-linear programming problem [20]. This is a computationally intensive optimization problem

to solve even for commercial solvers. Moreover, the real-world problem is not a single snapshot problem but rather the optimal topology must be designed considering the variation of load and generator injections over a time horizon (several time steps). This significantly increases the computational complexity of the problem. However, there is a need for real-time/fast optimal topology recommendation systems. No existing optimal power flow solver can yet tackle this problem [19].

In the use case at TenneT the amount of possible substation configurations is kept relatively small by choosing a subset of only 7 controllable substations. This leads to about 200 different valid unitary substation configurations (i.e., topologies with only one substation being reconfigured). However, even for this relatively small number of substations and configurations the combinatorial explosion is significant. For example, allowing for two substations being reconfigured leads to about 15k possible topologies and reconfiguring 3 substations allows for about 600k topologies [18]. Assuming that a load flow calculation takes a second, it turns out that brute force load flow computation for all these topologies already takes days to weeks. For larger grid sizes the number of topologies is practically infinite.

A simple approach to this problem is the so-called greedy agent (or controller) in which no time horizon is taken into account. That is, each timestamp is optimized independently except that the current topology is determined by the sequence of the previous actions. We also employ the constraint that only one substation (or no substation at all) can be reconfigured per timestamp. In this case at each timestamp the effects of all possible unitary substation reconfigurations (i.e., about 200 in the use case) are computed and the grid configuration that is optimal for that timestamp is chosen. In Figure 2 the green curve depicts the result of the greedy agent. Obviously, it is able to reduce the amount of overload but it is not able to fully mitigate all overloads. Consequently, it is not a trivial task to fully mitigate the overloads.

In order to solve this task, topology controllers that provide optimal control actions over a time horizon are needed. In the use case a simple RL-based agent [18] is used and the result is shown in Figure 2. The RL-agent is able to fully mitigate the overloads. Since the RL approach is probabilistic it is even able to offer several successful solutions with either two (cyan curve) or three (blue curve) substations being reconfigured. We note that a single substation reconfiguration is not sufficient to mitigate the overloads.

4 State-of-the-art research on sequential decision making applied to congestion management

Research on sequential decision making applied to real-time power network operations is still in its infancy. The opportunities for scientists to work collaboratively at scale on the problem were limited by a lack of commonly usable environments, baselines, data, networks and simulators. However, the ongoing energy transition forces industry and academia to invest significant resources in this topic. Recently, RTE TSO developed the open-source GridAlive ecosystem² to facilitate the development and evaluation of controllers (or agents) that act on power grids. With the Grid2Op framework³ at its core, any type of control algorithm in interaction with simulators of one's choice can be used such that gaps between research communities can be overcome.

Moreover, based on the GridAlive ecosystem, RTE TSO and collaborators launched a series of competitions, the so-called *Learning to Run a Power Network* (L2RPN) challenge. In each competition the participants need to develop controllers that control a power network to maintain a supply of electricity to consumers on the network over a given time horizon by avoiding a blackout. The controllers are exposed to realistic (stochastic) production and consumption scenarios, and the remedial actions are subject to real-world network constraints [21]. The aim of L2RPN is to foster faster progress in the field by creating the first large open-benchmark for solutions to the real-world problem of complex continuous-time network operations, building on previous advances in AI such as the ImageNet

²<https://github.com/rte-france/gridAlive>

³<https://github.com/rte-france/Grid2Op>

benchmark for computer vision [22]. On the one hand, the competitions present an opportunity for the AI community to demonstrate recent breakthroughs which could successfully find applications in real-world problems [23-25]. On the other hand, the aim is to raise awareness among the power system community about the innovations and potential of AI algorithms to solve network challenges and to embrace the application of different approaches to traditional problems [2].

While initial L2RPN competitions tested the feasibility of developing realistic power network environments [26] and the applicability of RL agents [18-19, 27-28], the 2020 L2RPN competition at NeurIPS [21, 29-30] and the 2021 *L2RPN with trust* competition [31] had increased complexity. These competitions came with realistically sized network environments, implying very large discrete and combinatorial action space dimensions due to the topological flexibilities of the network. Moreover, three real-world network operation challenges are addressed, namely, *robustness*, *adaptability*, and *trustworthiness*. More precisely, in the robustness track controllers had to operate the network maintaining supply to consumers and avoiding overloads while targeted unforeseeable line disconnections create N-1 situations by disconnecting one of the most loaded lines at random times [32]. In the adaptability track controllers need to cope with unseen energy mix distributions at test time which are different from the energy mix distributions provided during training. Finally, in the 2021 *L2RPN with trust* competition controllers additionally had to provide a confidence level of their actions. That is, controllers were not only evaluated based on their operational performance but also based on how trustworthy a controller will be for human operators.

Throughout the competitions, ML approaches showcased continuous robust and adaptable behaviours over long time horizons. This behaviour was not previously exhibited by the expert systems [33], or by optimization methods that are limited by computation time [34-36]. Participation and activity were steady with entries from all over the world, and corresponding research is emerging [18-19, 27-28, 30]. It is worth noting so far that the best teams such as Baidu and Huawei actually came with little power system knowledge but great AI expertise, confirming the benefits that the AI community could bring the power system community. The winning solutions employ a combination of expert rules, brute force simulation for action validation, and on top of that different RL approaches to increase planning abilities and get a final boost in operational performance [21, 31]. The confidence levels are still mainly determined via rule-based approaches. We also note that until now even the best agents still fail over 30% of the L2RPN test scenarios and sending alarms based on confidence levels is equally successful. Consequently, the competitions so far demonstrate high potential but also still indicate a lot of room for improvement.

5 Research challenges when applying AI to power systems

Although impressive proof-of-concept results have been obtained there remains a large gap to deployment, compared to other sectors like the automotive or the biotechnology industries [37]. This relates both to methodological challenges that are specific to the application domain, and to organizational challenges stemming from the way that system operators and the surrounding ecosystem are currently organised. Methodological challenges are addressed below and organizational challenges are discussed in the next section.

Power systems exhibit several properties that require domain-specific research efforts. More precisely, the power system is based on (I) a large-scale real-world system with (II) unique (cyber-)physical properties and dynamics; it is a (III) critical infrastructure that is (IV) subject to significant uncertainty due to its open nature (e.g., international borders, market integration) and a variety of evolving objectives (regulation and the energy transition); finally, (V) humans in the form of operators play a central role in maintaining adequate grid security margins. These properties imply the following research challenges:

(I) A large-scale real-world physical system is not physically reproducible, so it cannot be put in a laboratory (similar to the climate system, for example). Moreover, as a critical infrastructure (see also III below) the power system can never be taken out-of-service to be tested as a whole. To study the

entire system under a broad range of conditions, researchers need to work with well-developed models and simulators. In the AI community, *Sim2Real* refers to the concept (which emerged in robotics) of transferring capabilities learned in simulation to the real system [38-39]. *Sim2Real* draws its appeal from the fact that it is cheaper, safer (see also III below) and more configurable to perform experiments in simulation than in the real world. However, it has the obvious pitfall that discrepancies may exist between the simulation model and the real-world environment. One approach to address this is to explicitly account for (uncertainties due to) model-errors when designing or training agents/controllers [37].

Furthermore, as also shown by the L2RPN competitions (see section 4), the power network gives rise to large observation and action spaces. These large state and action spaces can present serious issues for traditional RL algorithms [37]. So far, these spaces are mainly reduced by using brute force approaches. Efficient search techniques are needed in order to exploit all available flexibilities. Moreover, decisions span a range of time scales; [40] have proposed a proxy-based method to efficiently deal with the complexity resulting of time scale hierarchies.

Finally, large amounts of data are important for agent training and validation, especially in the presence of changing distributions, but real-world data sources are limited. Consequently, extra emphasis should be placed on real-world data collection (see also section 6). On the other hand, techniques for creating realistic synthetic data sets are a crucial supplement [41-42].

(II) The power system has unique physical properties (electrical, topological and dynamical) and it is closely interlinked with embedded measurement and control systems. To start with the latter, one can view this as *cyber layer* that is superimposed on the physical network layer: operators can only access the physical network via specific IT infrastructure, and automata are deployed on the network. These specific structures need to be taken into account when building (decision support via) agents/controllers [43].

Focusing on the physical properties of the system, there is an opportunity to enhance ML models using the physical equations governing the system, which gives rise to the new research field of *physics-informed ML* [44].

Along similar lines, the bus-line representation of power systems suggests a natural fit for the use of *graph neural networks* (GNNs), a class of deep learning methods designed to perform inference on data described by graphs that receives a lot of attention recently [45-46]. For example, an obvious example of binary graph classification via GNNs applied to power systems is diagnosing whether a power network state satisfies the N-1 principle [47]. Also, faster load flow solvers have been developed using GNNs [48-49]. However, the unique properties of the power network hamper easy transfer of AI methods developed for other sectors even if the methods are already tailored to graph-like structures. For example, the actions of busbar splitting effectively split a node in a graph in two nodes (or conversely merge two nodes into one), which is not common in other networks. The unique structure of substations requires adequately adapted approaches in order to incorporate specific network constraints and exploit specific network flexibilities [48-49]. Equally, the effective handling of long-range dependencies (i.e., non-local effects) and the fast propagation speed of electricity represent power system specific challenges. In other words, the performance of GNNs on power grids (in particular, across different topologies and different power grids) still needs to be thoroughly researched [50]. The same holds for the question of how agents/controllers developed for a specific network can be generalized, that is, effectively transferred to other networks (transfer learning).

(III) As a critical infrastructure, the power system is designed and operated with high availability in mind. Ensuring its security through control actions therefore amounts to the analysis and mitigation of low-probability, but potentially high-impact, events [51]. It is generally desired that the agent performs robustly for all task instances and not just in expectation. Therefore, its performance cannot be summarized by a single scalar describing cumulative reward, but must consider the full distribution of behaviours both during training and testing. Operators will usually be averse to high-impact scenarios,

especially so because these rare events are often associated with large uncertainties (e.g., due to model mismatch, measurement errors or data scarcity).

The broad area of risk-averse reinforcement learning is known as *Safe RL* [52]. A typical approach to include risk-aversion is to use a Conditional Value at Risk (CVaR) objective [53], which looks at the expected *tail* reward (beyond a given percentile), rather than the expected reward. Chow et al. show that by optimizing reward CVaR, the agent is able to improve upon its worst-case performance, a property central in operating such safety critical infrastructures. Alternatively, the CVaR can be added used in constraints for policies with other reward objectives [54]. Safe RL methods should be further developed to deal with strict constraints on high-impact low probability events [37]. This also necessitates dealing with the potential for catastrophic model errors in a methodologically consistent manner.

(IV) The fourth set of challenges is related to the wider environment in which the power system must operate. First, instead of a single objective, the operator may desire to optimize a range of objectives: operational costs, long-term costs and asset wear, social welfare, environmental factors, and regulatory risks. When these objectives cannot be combined into a single objective using weights that are known in advance, specific *multi-objective RL* methods are required. There are approaches to learning the pareto-optimal reward function [55], and recent attempts to develop such methods for the deep reinforcement learning setting [56].

Second, many systems are electrically connected to neighbouring systems. This either necessitates tight integration with regional control centres (and possibly a centralised or hierarchical AI solution), or development of less tightly coupled *multi-agent RL* schemes.

(V) And last but not least, real systems are owned and operated by humans who need to be reassured about the controllers' intentions and require insights regarding failure cases [2, 12]. Explainability of proposed actions is important in this setting, and even more so in the context of a highly regulated industry, where human operators must be able to justify control actions. Explainability takes the form of *a priori* explanations of proposed actions, especially when proposed actions are unexpected, *a posteriori* explanations of actions that were later considered inappropriate. Moving beyond an advisory approach where AI agents suggest control actions to the human operator, the authors envision a hybrid intelligence approach for system operations, where operators and intelligent agents cooperate to control the grid. Hybrid intelligence or human-centred AI is a young and active research area [15, 57-58], where further development is required.

6 Embedding of AI development in system operators

To facilitate the transition of AI solutions for power systems from research to business-as-usual deployment, a number of institutional challenges must be faced. This section firstly identifies typical key enabling capabilities for the development of AI solutions, and secondly a vision of the production cycle of AI solutions in power system operation [59].

In order to enable and ultimately master the design, development, validation, training and deployment of AI solutions, each organisation needs to develop sufficient capabilities in the *people*, *data*, and *platform* pillars [60]:

- (A) The *people pillar* represents human resources including skill, culture, and organisation required for the successful launching of AI initiatives, their realisation, as well as their long-term maintenance. In general, a broad range of different human expertise and skills are necessary to identify, design, build, validate, and maintain AI solutions. Especially when starting-up, timely setting of way-of-working strategies, putting data collection and governance practices in place, as well as the right infrastructure to support the development as well final implementation are of crucial importance to gain momentum and swiftly demonstrate the added value of AI to the power system operators.

For system operators, it is essential to gain in-house knowledge of AI methods and broader digital competencies, partially through training for existing staff and, for new staff, by embedding these subjects in curricula of further and higher education. Another important aspect is to facilitate knowledge sharing and building through in-house AI communities, like the TenneT AI Center of Excellence. Of course, the purpose is not only to grow knowledge, connect and retain talent, but also to educate others about AI possibilities and promote AI developments internally in an organisation and connect and co-develop with similar AI communities externally. Finally, AI communities should define best practices including ethic guidelines when developing, deploying and maintaining AI solutions.

- (B) The *data pillar* is about organising data in terms of data accessibility, governance, and quality. In the majority of cases, availability of both historical and real-time process data is of key enabling importance to swiftly start with an AI initiative. Typically, access to data set dumps is sufficient for the development phase of an AI solution, while for production phase, data must be regularly available. Within the TSO/DSO organisation, it is often worth considering setting up a dedicated database to store structured and non-structured data for long-term, and to consolidate and democratize data sources, improve data quality and security via central data governance, and simplify data access. Nevertheless, it is also worth considering at the earliest stage to save various operational and non-operational data like SCADA and PMU measurements, power system models, grid security analysis results, energy market outcomes, and weather data, to have such data sets available when needed for future applications. In other words, an ambitious digitalization step is crucial within companies since quality and rich data is an essential raw material: it should now be considered as a company product/asset.

In particular, it is crucial for AI development and testing to thoroughly define *AI learn- and testbeds*. AI solutions need to be trained, tested, and benchmarked against other solutions. In the last decade fast AI improvements in computer vision [61], natural language processing [62], or biology [63] for example have indeed been largely driven big open benchmarks [64-66]. For this purpose, datasets of various power system operation scenarios (i.e. grid models, injections, constraints) that embody the past, present and future conditions of power network operations are necessary, and testbeds need to be developed according to the specific needs of system operators. Moreover, while test datasets should be as close as possible to real-world scenarios to best evaluate any solution, training datasets/environments should be rich and large enough for AI developments, possibly somewhat different from scarce and complex real data. In practice, this leads to a range of testbeds that increase in complexity from abstract model systems with model-generated data to full-fledged digital twins.

- (C) The *platform pillar* is about tools and technology used to develop, test and implement AI solutions. For development, platforms should be developer-friendly and supported by set of data analytics, visualisation, and programming tools, database pipelines, scalable compute and storage resources, in order to quickly develop prototypes and demonstrate value. In addition, the platform should be flexible enough to enable seamless yet secure cooperation with external parties when co-developing AI solutions. For a production environment, scalability and reliability are key aspects. Crucial is the ability to support continuous integration and deployment of new improved versions of the AI solution for maintenance purpose without affecting the rest of the system.

The above three pillars mostly focus on required enablers for AI solution realisation within a company. Nevertheless, for a seamless and rapid production cycle from design to implementation of AI solutions, the company needs to be organised to facilitate collaboration between Research, Development, and Deployment innovation phases. However, the process of technology innovation often takes place as schematically indicated in Figure 3. Hereby, each of the phases include Data, Code, People, and Testing as required resources/capabilities per project.

The research phase is often performed in a closed environment within small research teams, by using benchmark power system models, and without direct involvement from system operations experts. The main goal of this phase is to innovate and develop new AI techniques applicable to power system problems in order to show value.

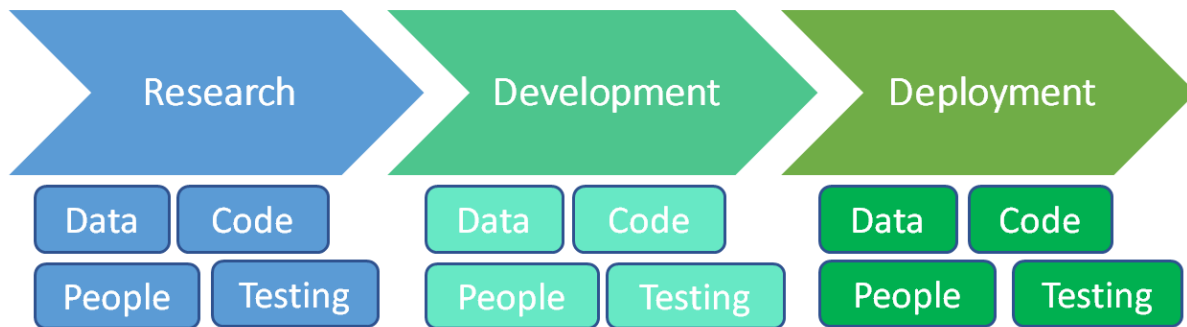


Figure 3: Schematic of the linear and compartmentalised process of technology innovation. Resources, practices, and capitalization are largely isolated and unshared between research, development, and deployment phases. This can lead to reduced alignment and slow progress.

The development phase is often performed by an external expert company, by using actual realistic power system models and historical datasets, and with time-to-time involvement of system operations experts. The main aim of this phase is to take the outcomes of the Research phase, further innovate, and develop a minimum-viable product and later final solution that reliably works for a specific real-world use case.

The deployment phase is then performed by utilities themselves or the EMS/SCADA system suppliers, by using real-time data streams, and with direct involvement of the system operations experts. The main aim of this phase is to take the final product of the Development phase, and implement it in the operator training simulator and production environments of EMS/SCADA systems to be used by the power system operators.

As can be observed in Figure 3, the innovation process of largely independent phases prevents the virtuous cycle of rapid development/testing/deployment that other sectors have used to good effect, in particular in AI-driven companies. To gain the most in the shortest time when developing AI solutions, it is crucial firstly that each company integrates the phases of the innovation process, and secondly that the *integration of system operation and the global AI R&D ecosystem* is increased, as schematically indicated in Figure 4. Near real-world testing environments, providing realistic scenarios and validation criteria, need to be continuously defined and refined back and forth from research to deployment. The use of open data, models and code with open access papers should be encouraged, much like within the LF Energy Initiative (<https://www.lfenergy.org/>), to enlarge the ecosystem for AI tool development. First steps in this direction are currently taking place, as exemplified by the L2RPN competitions and the European AI on Demand Platform (<https://www.ai4europe.eu/>).

7 Summary

Given the increasing need to meet the new operational requirements of power systems and prepare for the future, adaptation of cutting-edge Artificial Intelligence (AI) technologies in the operational processes is paramount to timely meet the challenges. The focus of this paper is on applying AI in power system operations, in particular for the development of decision support tools. First, the paper elaborates on the decision-making process of the power system operators and presents a mirroring digital framework consisting of AI and control theory to mimic sequential decision making of the operators. Next, a demonstrating example in the field of congestion management is presented by a real-world AI use-case at TenneT TSO. The paper continues with state-of-the-art on sequential decision making applied to congestion management and elaborates on research challenges when applying AI to the power systems problems. Finally, the paper elaborates on the enabling capabilities with focus on people, data, and platform pillars an organisation needs for mastering the development as well as maintenance of AI solutions, and proposes a cyclic (agile) process approach to decrease time from development to actual deployment and cooperation between research and industry organisations.




Figure 4: Two schematics for fast AI production. (Left) Schematic of the cyclic and largely shared process of technology innovation. It enables fast iteration and synchronization leading more efficiently to products suitable for real-world application. It applies to both the organization within a company as well as the collaboration between companies. (Right) Schematic of an R&D ecosystem for efficient and real-world applicable AI innovation. The dashed boxes indicate resources that can be largely shared between organizations. System operators have a unique position in the ecosystem since they are the only ones that can access and provide feedback from the real system. Consequently, system operators have a unique responsibility in forming the testing environment.

BIBLIOGRAPHY

- [1] Donnot, B. et al., "Introducing machine learning for power system operation support." arXiv preprint arXiv:1709.09527 (2017).
- [2] Prostejovsky, A. et al., "The future role of human operators in highly automated electric power systems." *Electric Power Systems Research*, 175:105883 (2019).
- [3] Marot, A. et al., "Towards an AI assistant for human grid operators." arXiv preprint arXiv:2012.02026 (2020).
- [4] R.S. Sutton and A.G. Barto. "Introduction to Reinforcement Learning", MIT Press (2018).
- [5] Kelly, A. et al., "Reinforcement Learning for Electricity Network Operation." arXiv preprint arXiv:2003.07339 (2020).
- [6] Li, Z. et al., "Reinforcement Learning for Robust Parameterized Locomotion Control of Bipedal Robots," 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 2811-

- 2817), (2021). doi: 10.1109/ICRA48506.2021.9560769. arXiv preprint arXiv:2103.14295 (2021).
- [7] Cohen, J., “How Google’s Self-Driving Cars Work” (2020). <https://heartbeat.fritz.ai/how-googles-self-driving-cars-work-c77e4126f6e7>
- [8] Bellemare, M. G. et al., “Autonomous navigation of stratospheric balloons using reinforcement learning.” *Nature*, 588(7836), 77-82 (2020).
- [9] Parfitt, N., “Towards the better Datacenter” (2020). <https://www.intelligentcio.com/wp-content/uploads/sites/12/2020/11/Towards-the-Better-Datacenter-Huawei.pdf>
- [10] Glavic, M. et al., “Reinforcement learning for electric power system decision and control: Past considerations and perspectives.” *IFAC-PapersOnLine*, 50(1):6918–6927 (2017).
- [11] O’Malley, M. et al., "Enabling Power System Transformation Globally: A System Operator Research Agenda for Bulk Power System Issues," in *IEEE Power and Energy Magazine*, vol. 19, no. 6, pp. 45-55, Nov.-Dec. 2021, doi: 10.1109/MPE.2021.3104078.
- [12] Marot, A. et al., “Perspectives on Future Power System Control Centers for Energy Transition,” in *Journal of Modern Power Systems and Clean Energy*, vol. 10, no. 2, pp. 328-344, (2022). doi: 10.35833/MPCE.2021.000673. <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9744623>
- [13] S. Russell and P. Norvig “Artificial Intelligence: A Modern Approach”, Pearson, (2021).
- [14] MathWorks, “Reinforcement Learning for Control Systems Applications” (2022). <https://nl.mathworks.com/help/reinforcement-learning/ug/reinforcement-learning-for-control-systems-applications.html>
- [15] Monarch, R., “Human-in-the-Loop Machine Learning: Active learning and annotation for human-centered AI”, Manning, (2021).
- [16] Shamma, J., “Feedback Control Perspectives on Learning”, invited talk at NeurIPS 2020, https://nips.cc/virtual/2020/public/invited_16168.html (2020).
- [17] Enexis, (2021). <https://www.enexis.nl/zakelijk/duurzaam/beperkte-capaciteit/gebieden-met-schaarste>
- [18] Subramanian, M. et al., “Exploring grid topology reconfiguration using a simple deep reinforcement learning approach,” arXiv preprint arXiv:2011.13465 (2020).
- [19] Matavalam, A. et al., “Curriculum Based Reinforcement Learning of Grid Topology Controllers to Prevent Thermal Cascading.” arXiv preprint arXiv:2112.09996, (2021).
- [20] Pourahmadi, F. et al., “Dynamic uncertainty set characterization for bulk power grid flexibility assessment,” *IEEE Systems Journal*, (2019).
- [21] Marot, A. et al., “Learning to run a power network challenge: a retrospective analysis.” arXiv preprint arXiv:2103.03104 (2021).
- [22] Deng, J. et al., “Imagenet: A large-scale hierarchical image database.” *IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee (2009).
- [23] Silver, D. et al., “Mastering chess and shogi by self-play with a general reinforcement learning algorithm”. arXiv preprint arXiv:1712.01815, (2017).
- [24] Vinyals, O. et al., “Grandmaster level in starcraft ii using multi-agent reinforcement learning”, *Nature*, 575 (7782):350–354, (2019).
- [25] Badia, A. et al., “Agent57: Outperforming the Atari human benchmark” in *International Conference on Machine Learning*, pages 507–517. PMLR, (2020).
- [26] Marot, A. et al., “Learning to run a power network challenge for training topology controllers.” *Electric Power Systems Research* 189:106635 (2020).
- [27] Lan, T. et al., “AI-Based Autonomous Line Flow Control via Topology Adjustment for Maximizing Time-Series ATCs”, 2020 IEEE Power & Energy Society General Meeting (PESGM), pages 1-5, IEEE, (2020). arXiv preprint arXiv:1911.04263 (2019).
- [28] Yoon, D. et al., “WINNING THE L2RPN CHALLENGE: POWER GRID MANAGEMENT VIA SEMI-MARKOV AFTERSTATE ACTOR-CRITIC”, *International Conference on Learning Representations (ICLR)*, 5, (2021). <https://openreview.net/pdf?id=LmUJqB1Cz8>
- [29] Marot, A. et al., “L2RPN: Learning to Run a Power Network in a Sustainable World NeurIPS2020 challenge design”. <https://www.public.asu.edu/~ywgeng2/Tutorial5/pdf/111.pdf> (2020).
- [30] Zhou, B. et al., “Action Set Based Policy Optimization for Safe Power Grid Management”, *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 168–181. Springer, 2021. arXiv preprint arXiv:2106.15200 (2021).

- [31] Marot, A. et al., “Learning to run a power network with trust”, arXiv preprint arXiv:2110.12908 (2021).
- [32] Omnes, L. et al., “Adversarial Training for a Continuous Robustness Control Problem in Power Systems”, arXiv preprint arXiv:2012.11390, (2021).
- [33] Marot, A. et al., “Expert system for topological remedial action discovery in smart grids” IET Digital Library, (2018)
- [34] Ruiz, P. A. et al., “Security-constrained transmission topology control MILP formulation using sensitivity factors,” IEEE Transactions on Power Systems, vol. 32, no. 2, pp. 1597-1605 (2017).
- [35] Little, E. et al., “Optimal Transmission Topology for Facilitating the Growth of Renewable Power Generation.” arXiv preprint arXiv:2103.15677 (2021).
- [36] Heidarifar, M. et al., “An optimal transmission line switching and bus splitting heuristic incorporating AC and N-1 contingency constraints”, International Journal of Electrical Power & Energy Systems, Volume 133, (2021). <https://doi.org/10.1016/j.ijepes.2021.107278>.
- [37] Dulac-Arnold, G. et al., “Challenges of real-world reinforcement learning.” arXiv preprint arXiv:1904.12901 (2019).
- [38] Fragkiadaki, K., course entitled “Deep Reinforcement Learning and Control”, (2019). <https://www.andrew.cmu.edu/course/10-403/slides/S19sim2real.pdf>
- [39] Kaspar, M. et al., “Sim2Real Transfer for Reinforcement Learning without Dynamics Randomization” arXiv preprint arXiv:2002.11635 (2020).
- [40] Dalal, G. et al., “Hierarchical decision making in electricity grid management” in International Conference on Machine Learning (pp. 2197-2206), PMLR, (2016).
- [41] Zheng, X. et al., "PSML: A Multi-scale Time-series Dataset for Machine Learning in Decarbonized Energy Grids." arXiv preprint arXiv:2110.06324 (2021).
- [42] Wang, S. et al., “Generating Multivariate Load States Using a Conditional Variational Autoencoder,” arXiv preprint, arXiv:2110.11435 (2021).
- [43] Allgöwer, F. et al., “Position paper on the challenges posed by modern applications to cyber-physical systems theory”, Nonlinear Analysis: Hybrid Systems, 34, 147-165 (2019).
- [44] Raissi, P. et al., “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”, Journal of Computational Physics 378, 686-707 (2019).
- [45] Zhou, J. et al., “Graph Neural Networks: A Review of Methods and Applications”, AI Open 2021, arXiv preprint arXiv:1812.08434, (2021).
- [46] Liao, W. et al., “A Review of Graph Neural Networks and Their Applications in Power Systems" in Journal of Modern Power Systems and Clean Energy, (2021). <https://ieeexplore.ieee.org/abstract/document/9520300>
- [47] Fullhase, S., “Testing the n-1 principle with Graph Neural Networks”, master thesis, (2020). https://www.researchgate.net/profile/Jacco-Heres/publication/343099809_Testing_the_n

[1_principle_with_Graph_Neural_Networks/links/5f16d65e299bf1720d56ad33/Testing-the-n-1-principle-with-Graph-Neural-Networks.pdf](https://hal.archives-ouvertes.fr/hal-02175989/document)

- [48] Donon, B. et al., “Graph Neural Solver for Power Systems”, International Joint Conference on Neural Networks, (2019). <https://hal.archives-ouvertes.fr/hal-02175989/document>
- [49] Donon, B. et al., “Neural networks for power flow: Graph neural solver”, Electric Power Systems Research, Volume 189, (2020).
- [50] Ringsquandl, M. et al., “Power to the Relational Inductive Bias: Graph Neural Networks in Electrical Power Grids”, arXiv preprint arXiv:2109.03604 (2021).
- [51] GARPUR project, “Generally Accepted Reliability Principle with Uncertainty modelling and through probabilistic Risk assessment” (2017). <https://www.sintef.no/projectweb/garpur/>
- [52] Garcia, J. et al., “A comprehensive survey on safe reinforcement learning”, Journal of Machine Learning Research, 16(1), 1437-1480, (2015).
- [53] Chow, Y. et al., “Risk-sensitive and robust decision-making: a cvar optimization approach”. arXiv preprint arXiv:1506.02188, (2015).
- [54] Yang, et al., “WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning”, Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, 35(12), 10639-10646, (2021). <https://ojs.aaai.org/index.php/AAAI/article/view/17272>
- [55] Roijers, et al., “A survey of multi-objective sequential decision-making”, Journal of Artificial Intelligence Research. Oct 18;48:67-113. (2013).
- [56] Nguyen, et al., “A multi-objective deep reinforcement learning framework”, Engineering Applications of Artificial Intelligence. Nov 1;96:103915, (2020).
- [57] Heer, J., “Agency plus automation: Designing artificial intelligence into interactive systems.”, Proceedings of the National Academy of Sciences, 116(6), 1844-1850, (2019).
- [58] Schmidt, A et al., “Artificial Intelligence for Humankind: A Panel on How to Create Truly Interactive and Human-Centered AI for the Benefit of Individuals and Society” in IFIP Conference on Human-Computer Interaction (pp. 335-339). Springer, (2021).
- [59] Paleyes, A. et al., “Challenges in Deploying Machine Learning: a Survey of Case Studies.” arXiv preprint arXiv:2011.09926 (2020).
- [60] GoDataDriven 2021, The AI Maturity Journey. Go Data Driven, 2021, available: <https://cdn2.hubspot.net/hubfs/697348/GDD%20General%20files/WP-AI-Maturity-Journeys/AI%20Maturity%20Journey-Whitepaper-GoDataDriven.pdf>
- [61] Kolesnikov, A. et al., “Big transfer (BiT): General visual representation learning.” arXiv preprint arXiv:1912.11370, (2019).
- [62] Brown, T. B. et al., “Language models are few-shot learners.” arXiv preprint arXiv:2005.14165 (2020).
- [63] Jumper, J. et al., "High accuracy protein structure prediction using deep learning." Fourteenth Critical Assessment of Techniques for Protein Structure Prediction (Abstract Book), 22, 24 (2020).
- [64] MOULT, J. et al., “A decade of CASP: progress, bottlenecks and prognosis in protein structure prediction.” Current opinion in structural biology, vol. 15, no 3, p. 285-289 (2005).
- [65] Bentivogli, L. et al., "The Fifth PASCAL Recognizing Textual Entailment Challenge." In TAC (2009).
- [66] Wang, A. et al., “GLUE: A multi-task benchmark and analysis platform for natural language understanding.” arXiv preprint arXiv:1804.07461 (2018).

