



HAL
open science

The challenge of controlling microgrids in the presence of rare events with Deep Reinforcement Learning

Tanguy Levent, Philippe Preux, Gonzague Henri, Réda Alami, Philippe Cordier, Yvan Bonnassieux

► **To cite this version:**

Tanguy Levent, Philippe Preux, Gonzague Henri, Réda Alami, Philippe Cordier, et al.. The challenge of controlling microgrids in the presence of rare events with Deep Reinforcement Learning. IET Smart Grid, In press, 10.1049/stg2.12003 . hal-02971554

HAL Id: hal-02971554

<https://hal.science/hal-02971554>


Submitted on 21 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ORIGINAL RESEARCH PAPER

The challenge of controlling microgrids in the presence of rare events with deep reinforcement learning

Tanguy Levent¹  | Philippe Preux² | Gonzague Henri³ | Réda Alami⁴ |
Philippe Cordier⁴ | Yvan Bonnassieux¹

¹Ecole Polytechnique, CNRS, IP Paris, Palaiseau, France

²Université de Lille, CRNS, Lille, France

³Total EP R&T, Houston, USA

⁴Total, Palaiseau, France

Correspondence

T. Levent, Ecole Polytechnique, CNRS, IP Paris, Palaiseau, France.

Email: tanguy.levent@polytechnique.edu

Abstract

The increased penetration of renewable energies and the need to decarbonise the grid come with a lot of challenges. Microgrids, power grids that can operate independently from the main system, are seen as a promising solution. They range from a small building to a neighbourhood or a village. As they co-locate generation, storage and consumption, microgrids are often built with renewable energies. At the same time, because they can be disconnected from the main grid, they can be more resilient and less dependent on central generation. Due to their diversity and distributed nature, advanced metering and control will be necessary to maximise their potential. This paper presents a reinforcement learning algorithm to tackle the energy management of an off-grid microgrid, represented as a Markov Decision Process. The main objective function of the proposed algorithm is to minimise the global operating cost. By nature, rare events occur in physical systems. One of the main contribution of this paper is to demonstrate how to train agents in the presence of rare events. Merging the combined experience replay method with novel methods called ‘Memory Counter’ unstucks the agent during its learning phase. Compared to baselines, an extended version of double deep Q-network with a priority list of actions into the decision making strategy process lowers significantly the operating cost. Experiments are conducted using 2 years of real-world data from Ecole Polytechnique in France.

NOMENCLATURE

a_t	Action taken by the agent at t	J_{obj}	Cost/Objective function
a_a	Action taken by the DDQN-EMS _a	k	Episode number
a_p	Action taken by the DDQN-EMS _p	m	Batteries operational cost
\mathcal{A}	Action Space of a MDP	$P_B(t)$	Batteries power delivered at t
\mathcal{A}_a	Action Space of the DDQN-EMS _a	$P_G(t)$	Diesel generator power delivered at t
\mathcal{A}_p	Action Space of the DDQN-EMS _p	$P_C(t)$	Load curtailment power at t
c	Load curtailment cost	$P_{PV}(t)$	PV power at t
\mathcal{C}	Memory Counter	$P_L(t)$	Load power at t
\mathcal{D}	Replay memory	$P_{Net}(t)$	Net demand at t
$E_{Bcap}(t)$	Batteries SoC at t	P_{Bmin}	Minimum power delivered by the batteries
E_{Bmin}	Batteries SoC minimum limit	P_{Bmax}	Maximum power delivered by the batteries
E_{Bmax}	Batteries SoC maximum limit	P_{Gmin}	Minimum power delivered by the diesel generator
\exp_t	Experience store in the replay memory at t	P_{Gmax}	Maximum power delivered by the diesel generator
g_z	Activation Function of the neuron z	q	Diesel generator operational cost
\mathbf{G}	Discounted Return	$Q(s_t, a_t)$	State Action value function at t
G	Memory Counter Capacity		

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *IET Smart Grid* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.

$Q^*(s_t, a_t)$	Optimal State Action value function at t
r_t	Immediate reward at t
\mathbf{R}	Transition function of a MDP
R	Replay Memory Capacity
s_t	State of the environment at t received by the agent
\mathcal{S}	State Space of a MDP
\mathbf{T}	Reward function of a MDP
t	Time step of an episode
T	Length of a training episode
w	Weights of the policy Q network
w^-	Weights of the target Q network
\mathcal{W}	Weights Space
x	Input of a function
\mathcal{X}	Input Space of a function
y	Desired output value of a neural network
\hat{y}	Estimated output value of a neural network
γ	Discounted factor
π	Policy
π^*	Optimal Policy
α	Learning rate
ϵ	Exploitation rate
δ	Frequency of time step t in \mathcal{C}
ρ	Threshold of the memory counter for δ
ϕ	New exploration rate value for MemC strategy
τ	Frequency of updating the Q target
Δt	Time interval between two time steps
Δw	Weights update
∇_w	Weights derivative

1 | INTRODUCTION

One of the main challenges of the twenty-first century is to reduce greenhouse gases emissions to comply with the 2015 Paris Agreement [1]. To tackle this challenge, there has been a global increase in investments for renewable energy projects [2] and for distributed energy resources (DER), as demonstrated in Referenes [3,4]. As a result, utilities must adapt the grid infrastructure to handle stochastic resources such as solar or wind energy in order to maintain grid reliability and stability. A microgrid is a small scale power system that consists of renewable energy sources (wind turbine or photovoltaic panels), traditional generators (diesel generators), batteries, loads, and an energy management system (EMS). The principal definitions and foundations of a microgrid have been developed and explained in References [5–7]. A microgrid may operate either connected to the main grid, or disconnected from it, in the islanded mode. A microgrid may also be completely disconnected from the main grid (off-grid). Furthermore, developing renewable energy sources (RES) capacity will impact electricity markets with non-dispatchable resources. Microgrids are considered as an important technology for the energy transition and integrating more renewable while increasing resiliency. In Reference [8], the author

proposes that the grid could evolve from a monolithic system centrally operated to a system of microgrids.

In this paper, we focus on the EMS; more specifically, we are interested in designing an algorithm that is able to manage the operations of a microgrid. Often, when making decisions, the EMS needs to consider uncertain future states: the demand and production depend on both the human activity and the weather. In such a setting, we find the reinforcement learning (RL) framework to be a strong candidate to tackle this problem [9]. In this paper, we present a Deep Reinforcement Learning (DRL) approach for the EMS of a hybrid off-grid microgrid based on PV panels. An EMS has to deal with rare events, which are situations that are less likely to occur which significantly affect the performance, as explained in References [10,11]. In our study, rare events have a probability of occurrence smaller than 5%, which are considered as significant rare events (SRE). According to Reference [12], conventional RL methods are not robust when SRE occur. Indeed, they fail when there are rare events that affect significantly the expected performance. There is no published research to date regarding SRE in the case of controlling a microgrid with RL methods.

One of the main considerations made in this paper is the absence of forecasters; this makes our work significantly different from the state of the art on managing microgrids. In the literature, a forecaster provides a prediction of the next 24 h of the PV power output and the load consumption; an optimisation method is used to manage the different units of the system using this forecast (planning). Developing a forecaster for every microgrid would be difficult, as this requires a weather forecaster, load measurement infrastructures, and eventually a sufficient amount of historical data [13]. Therefore, we focus on the design of a controller without any forecasting capability, reacting to changing and unplanned conditions. The purpose of this study is to propose a novel approach using a DRL algorithm that minimises the operational cost of an off-grid microgrid. This algorithm is reactive; it does not use any forecaster; and it deals with SRE. We analyse the efficiency and robustness of this method by experimenting over a large variety of daily conditions, covering almost 2 years of solar and load conditions. We created a dataset with 2 years of activity of our experimental microgrid. A part of this dataset (training dataset) is used to train a DRL agent; then, its performance is assessed on the remaining part of the dataset (testing dataset). As EMS agent, we propose two versions of DRL agents that stem from the family of double deep Q-network (DDQN) algorithms. These two DRL agents are benchmarked against other methods: a RL decision tree [14] (RL-DT), a rule-based algorithm and dynamic programming (DP).

1.1 | Reinforcement learning applications for power system

RL has been proposed for several power system applications. In Reference [15], the authors use Q-Learning and an ensemble neural network for operation and maintenance in power systems with degrading elements and equipped with prognostics and health management capabilities. Deep Q-network (DQN) and

deep policy gradient have been studied for the scheduling of electricity consuming devices in residential buildings [16]. The authors in Reference [17] have used Q-learning to propose a dynamic pricing for demand response in a electricity market over two days of simulations. Authors in Reference [18] propose an EMS integrated by a fitted Q-Iteration algorithm to sell/buy the surplus/deficit electricity power output of smart homes. Regarding recent works relative to the energy management of microgrids using RL algorithms, a 2-step ahead Q-Learning algorithm was proposed in Reference [19] to manage the energy storage device of a microgrid to optimise its utilisation rate. A battery energy management in a microgrid of a two-month period has used a batch RL algorithm in Ref. [20]. The study in Reference [21] proposes an adaptive DP and RL framework to learn a control policy in order to optimise the critical load operation into a microgrid. A DRL algorithm has been studied in Reference [22] for operating a hydrogen storage device into an islanded microgrid. Finally, a deep review of RL algorithms applied into the electric power system domain was published in Reference [23], focusing on the past considerations and the new perspectives.

1.2 | Contributions and outline

The main contributions of this paper are:

- Designing a novel approach accelerates the learning phase of the agents and deal appropriately with significant rare events.
- Proposal of two RL agents solving the economic dispatch problem of the microgrid. This involves an approach based on a novel priority list of actions.
- Development of an off-grid hybrid microgrid simulator (*MGSimulator*) specifically tailored for RL.

This paper is organised as follows: Section 2 introduces briefly Markov Decision Processes (MDPs) and then the RL framework which provides a family of algorithms to solve MDPs. The microgrid management problem we are tackling is presented in Section 3. Section 4 introduces its modeling as an MDP. In Section 5, we present the algorithms we have designed to manage the problem of rare events. Section 6 serves two purposes: first to demonstrate by experiments the robustness of the new capability added to the learning agent to tackle rare events and then to demonstrate the performance of the DRL agent over a wide range of data (PV power and consumption). Then, we compare the results with other methods. Finally, we conclude in Section 7 by an overview of the current work and by suggesting future directions.

2 | REINFORCEMENT LEARNING BACKGROUND

The problem of managing operations can be seen as a sequential decision-making problem in a stochastic environment. RL is a sub-domain of machine learning, where an agent learns to complete a task by interacting with its environment.

To reach this goal, the agent learns to optimise a certain objective function that is defined by the consequences of its decisions. The most common way to model a RL problem is as a MDP. We briefly introduce this notion in the next section.

2.1 | Markov decision processes

A MDP is a discrete time framework for modeling sequential decision making problems. To apply an RL algorithm, the problem must be expressed in this formalism. In our case, the MDP of interest is defined as a 5-tuple $(\mathcal{S}, \mathcal{A}, \mathbf{T}, \mathbf{R}, \gamma)$ such that:

- \mathcal{S} denotes the finite set of the states,
- \mathcal{A} denotes the finite set of the possible actions,
- $\mathbf{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ denotes the transition function:
- $\mathbf{T}(s_{t+1}, a_t, s_t) = \mathbb{P}[s_{t+1} | s_t, a_t]$ is the probability that the emission of action $a_t \in \mathcal{A}$ in state $s_t \in \mathcal{S}$ at time t will lead to state $s_{t+1} \in \mathcal{S}$ at time $t + one$
- $\mathbf{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denotes the reward function and models the consequence of actions. $r_t = \mathbf{R}(s_t, a_t)$ is the immediate reward received by the agent after performing action a_t in state s_t . A reward value can be positive or negative to account for good and bad consequences.
- It is essential to distinguish the immediate reward from the optimisation of the objective function. The goal is to optimise rewards over a span of time, not immediate rewards: to optimise the objective function, one may have to perform actions that have bad immediate rewards but are necessary to reach the best long term behaviour.

In this paper, the objective function is the sum of discounted rewards: $\mathbf{G}(s_t) = \sum_{k=0}^T \gamma^k r_{t+k}$ with $\gamma \in [0, 1)$. γ controls how far we consider the consequences of actions in the future: γ close to 0 makes the agent focus on short-term (immediate or so) rewards, whereas γ close to 1 leads to long term optimisation.

2.2 | Reinforcement learning

RL aims at optimising the agent's behaviour while facing an unknown environment, that is an MDP where \mathbf{T} and \mathbf{R} are unknown. The behaviour of the agent is known as its 'policy' π which is a mapping $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ where $\pi(a_t | s_t)$ denotes the probability of taking action $a_t \in \mathcal{A}$ in state $s_t \in \mathcal{S}$.

The RL agent learns by trial-and-error, continuously interacting with its environment: at each time step t , the agent observes the current state of its environment $s_t \in \mathcal{S}$ and then chooses an action $a_t \in \mathcal{A}$ to perform. The agent performs a_t which leads the environment to $s_{t+1} \in \mathcal{S}$; the agent receives a reward $r_t = \mathbf{R}(s_t, a_t) \in \mathbb{R}$. It is important to note that the environment satisfies the Markov property as the next state s_{t+1} depends only on the current state s_t and the current action a_t : $\mathbf{T}(s_{t+1}, a_t, s_t)$.

Since in early days of 1980s, RL has been successfully used in a variety of problems, such as playing Backgammon [24], in health [25], video games [26] and board games where DRL has

been able to learn to play beyond human performance in games such as Chess and then Go, Reversi and others using only the rules of the game [27]. It should be noted that the recent achievements on board games concern environments that are deterministic: this point simplifies the problem. Moreover, these environments are stationary, that is the terms of the MDP do not change along the time. In our case, the environment is stochastic: the next state cannot be predicted exactly, and the stationarity of the environment is also not assumed.

2.3 | Q-learning

The state-action value $Q(s_t, a_t)$ quantifies how good it is to emit an action $a \in \mathcal{A}$ in a given state $s \in \mathcal{S}$ at time t . It is defined by $Q(s_t, a_t) = \mathbb{E}[G(s_t, a_t)]$. The Q function of the optimal policy is learned by repeated interaction between the agent and its environment and update according to Bellman equation:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha(r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)) \quad (1)$$

where $\alpha \in (0, 1)$ denotes the learning rate.

The Equation (1) converges towards the optimal value function $Q^*(s_t, a_t) = \max_{\pi} Q^{\pi}(s_t, a_t)$ reached for the optimal policy: $\pi^* = \arg \max_{\pi} Q^{\pi}(s_t, a_t)$ [28]. This update equation is the essence of the Q-Learning algorithm [9]. Another basic element of an RL algorithm concerns the exploration-exploitation dilemma. At each time step t , the agent needs to perform an action. Then, the agent has two possibilities: either the agent chooses the action, it has yet observed to be the best one in the current state (exploitation), or the agent chooses an action that seem sub-optimal to gain more information about it (exploration). Exploitation turns out to choose action $\arg \max_{a \in \mathcal{A}} Q(s_t, a)$. To learn, the agent has to explore; to perform optimally while optimising the objective function, the agent has to exploit. Exploitation is the obvious choice once the agent has learned the optimal policy. However, in a stochastic environment, and worse, in a non-stationary environment, the agent can never be certain it has found the optimal policy, so it has to keep on exploring, at least some times to times. As a result, the agent needs a strategy to balance exploration–exploitation. In an environment that does not change too fast (in which case it is difficult to learn anything at all), the agent has to explore with high probability at the beginning of its interactions with its environment; then, progressively, the balance has to shift towards exploitation. The study of the exploration–exploitation dilemma is the field of research known as the multi-armed bandit problem. This is a very active field of research, and many algorithms have been proposed to cope with various settings. Despite all these efforts, in the case of the RL problem, one of the most effective strategy remains a very simple one; it is called ϵ -decreasing greedy, and it is used in this paper. In this strategy, ϵ is the probability that the agent explores at step t ; then, with probability $1 - \epsilon$, it exploits and performs the action currently

estimated the best: this is a greedy choice, hence the name. In ϵ -decreasing greedy, ϵ is slowly decreasing along time, leading the RL agent to explore less and less, and then, to exploit more and more, hence to stick to its goal of optimising the objective function. So, we should subscript ϵ with a t but we drop it for the sake of simplicity of notations. With such a scheme, the agent may also increase ϵ if it detects that it needs to acquire more information; this is typically the case when the environment is non stationary and its dynamics changes along time. A pseudo-code of the Q-Learning algorithm is described in Algorithm.1.

Algorithm 1 Q-Learning

```

Set hyperparameters:  $\alpha \in [0, 1]$ ,  $\epsilon > 0$ .
Initialise  $Q(s, a)$ ,  $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$ 
for each episode do.
  Initialise the agent ( $s_0$ );  $t \leftarrow 0$ .
  While Terminal state not reached do.
    Choose  $a_t$  for  $s_t$  using  $Q$ .
    Emit action  $a_t$ , observe  $r_t, s_{t+1}$ 
    Update  $Q$  using (1)
  end while.
  Decrease the value of  $\epsilon$ 
end for.

```

A key element of the Q-Learning algorithm (or any RL algorithm) is the structure representing the Q function. Q may be seen as a table of real numbers with two indices; one for the state, one for the action. When the cardinal of state space and the cardinal of action space are small, Q is implemented in this way in basic Q-Learning implementations (so-called tabular Q-Learning). In real applications, it is common that the number of states is large, and even infinite (uncountable). In this case, we need a structure that can represent an infinite number of real values. This structure is called a ‘function approximator’. Over the years, various function approximators have been used (decision trees, random forests, support-vector machines and also k nearest neighbor techniques) among which neural networks are important. Due to this fact, we dedicate the next section to a brief recap about neural networks, a field today better known as Deep Learning (DL).

2.4 | Deep learning

DL is a branch of Machine Learning. It is the modern name for ‘artificial neural networks’. In the last 10 years, DL has revolutionised 50 years old research fields such as computer vision, signal processing, and natural language processing [29]. A neural network is made of a sequence of layers of neurons/units connected together by edges, in a feed forward way*, from the input to the output. Each edge is characterised by a weight, that is a real number. A data is input into the network, and a prediction is output. Traditionally, a neuron inputs a d -dimensional

*we do not consider recurrent neural networks here because they remain out of the scope of our work.

real vector $x_j \in \mathcal{X} \subset \mathbb{R}^d$; each element x_j is weighted by $w_j \in \mathbb{R}^d$; the neuron z outputs $\hat{y}_z = g_z(x_j, w_j)$ where g_z usually denotes a non-linear function, that is sigmoid, hyperbolic tangent, rectified linear unit... The data are integrated as the inputs of the network, its attributes being processed in each layer to finally calculate the output \hat{y} . This process is called the forward propagation. The output of the network $\hat{y}(x_i)$ for data x_i is compared to the desired output y_i . Then, the weights of the network are modified in order to reduce the difference between the output value and the expected value. This process is repeated again and again over the whole set of training data, until convergence. A rather recent and comprehensive description of DL is available in Reference [30].

2.5 | Reinforcement learning decision tree ↔ deep Q-network: the deep Q network algorithm

As presented above, Q-Learning is tabular: the Q values are stored in an array. In practice, this approach is not viable because Q-Learning would be restricted to small size state space, and, even worse, it does not generalise. To get around this problem, in a previous study [14], we use a decision tree algorithm after the RL learning phase using a tabular Q-Learning, in order to approximate a function between the states and the best associated action. Though generalising, this method remains limited to small discretised state spaces. Tabular RL has soon been extended to handle these problems by substituting the tabular representation of Q by a function approximation to learn, store, and estimate the state-action value. In DRL, the state vector feeds a neural network and an estimate of the Q-values is output for each action. Weights are updated following the Q-Learning update rule in Equation (1). Because of the combination of the updates of the weights of the neural network and the updates of Q estimates, Equation (1) can be simplified: the learning rate α can be removed as it is already used during the backward propagation phase resulting in:

$$Q(s_t, a_t, w) = r_t + \gamma \max_a Q(s_{t+1}, a, w) \quad (2)$$

Supervised learning methods such as neural networks require a dataset of examples, that is a set of input-output pairs. In Deep Q Network, we create an experience replay memory \mathcal{D} . An experience at each time step is defined by the tuple: $\text{exp}_t = (s_t, a_t, r_t, s_{t+1})$ and is stored in \mathcal{D} . Nevertheless, as each transition exp_t is recorded, we have a problem of correlation between close experiences, which is inconvenient for training the neural network. To avoid that, we select randomly a batch of experiences of N transitions from the pool of stored in \mathcal{D} to stabilise the input dataset. The learning algorithm is called Deep Q Network or DQN [31]. Equation (2) is treated as the Q target, and we update the weights iteratively using $w \leftarrow w - \Delta w$ such that:

$$\Delta w = \alpha \left[r_t + \gamma \max_a Q(s_{t+1}, a, w) - Q(s_t, a_t, w) \right] \quad (3)$$

$$\nabla_w Q(s_t, a_t, w)$$

2.6 | Double deep Q-network

In Equation (3), we need to compute the difference between the Q target and the current estimated Q-values. Both values use the same neural network of weights w . As a result, the target Q-values change when the weights $w \in \mathcal{W}$ are updated during the backward propagation phase, which leads to big oscillations during training the agent. In order to stabilise the algorithm, we use the trick proposed in Reference [31]. This trick consists in separating the network between the target and the current Q-values. Every τ updates, we copy the current Q network weights to the periodically fixed target Q network ones. We note w^- the weights corresponding to the target Q network, and we keep the w notation for the current Q network. A second improvement called Double DQN (DDQN) and introduced in Reference [32] concerns the problem of overestimations of Q-values because we use the max operator to choose the estimated Q-value of the next state in Equation (3). The solution consists in using the Q network to select the best action of the next step and then use the target network to evaluate Q:

$$Q(s_t, a_t) = r_t + \gamma Q(s_{t+1}, \arg \max_a Q(s_{t+1}, a, w), w^-) \quad (4)$$

3 | MODEL

We have designed a hybrid off-grid microgrid consisting of solar PV panels, a diesel generator (genset), batteries, a building and an EMS. Hybrid refers to the fact that the microgrid generates energy with both renewable and fossil resources. These different components are illustrated in Figure 1.

3.1 | Problem formulation

In this paper, we aim at minimising the operational cost of the proposed microgrid while respecting the system constraints over a period of time T . The marginal cost of the PV production is considered to be zero, thus it is not taken into account in the cost function. We assume a fixed marginal cost for the genset and the batteries. As a result, our objective function is the cumulative cost to operate the genset and the set of batteries over T , where the time step Δt is 1 h. Power losses on the feeders are not considered in the operational cost function. Finally, for the sake of simplicity, we assume that the electricity power (kW) at time t will consist of the production (kWh) during $[t, t + \Delta t]$. The EMS cost function is thus formulated as:

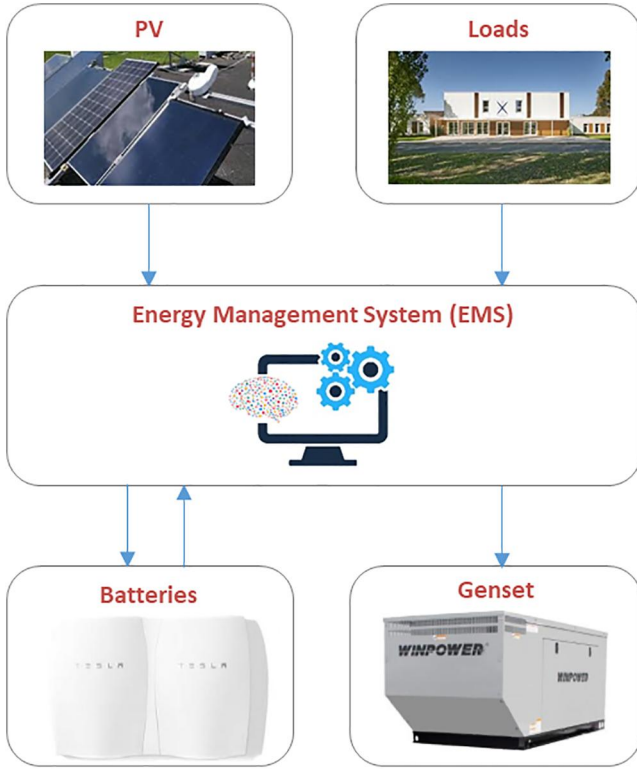


FIGURE 1 The off-grid hybrid microgrid

$$J_{obj} = \sum_{t=0}^T |P_B(t)|m + P_G(t)q + P_C(t)c \quad (5)$$

The variables m , q and c represent respectively the operational cost of the set of batteries, genset and load curtailment, respectively, and P_B , P_G and P_C are the charge or discharge power of the batteries, the genset power output and the power curtailment, respectively.

3.2 | Constraints and hypothesis

The difference between the PV power produced P_{PV} and the load consumption P_L is equal to the amount of electricity to be managed properly. The value can be positive or negative regarding the two components and is defined as a net demand P_{Net} . We formulate this difference of power at time t as:

$$P_{Net}(t) = P_{PV}(t) - P_L(t) \quad (6)$$

The principal constraint concerns the power balance of the microgrid as the generation needs to match the load. This constraint must be satisfied at any time t :

$$P_B(t) + P_G(t) + P_C(t) = P_{PV}(t) - P_L(t) \quad (7)$$

The second constraint is related to the battery energy capacity E_{Bcap} bounds at each time t . If the EMS does not respect the energy storage limits, it will automatically assume a crash:

$$E_{Bmin} \leq E_{Bcap}(t) \leq E_{Bmax} \quad (8)$$

The charging and discharging power limits of the batteries must satisfy:

$$P_{Bmin} \leq P_B(t) \leq P_{Bmax} \quad (9)$$

Furthermore, the dynamical set of batteries is modeled regarding the operation and the capacity as:

$$E_{Bcap}(t) = E_{Bcap}(t-1) - P_B(t)\Delta t \quad (10)$$

We consider a simple battery model with perfect efficiency. In future work, we will investigate more complex battery models. Finally, the diesel generator is also constrained by its own limits in terms of delivered power range:

$$0 \leq P_{Gmin} \leq P_G(t) \leq P_{Gmax}, \quad (11)$$

where $P_G(t)$ is equal to zero when it is turned off mode.

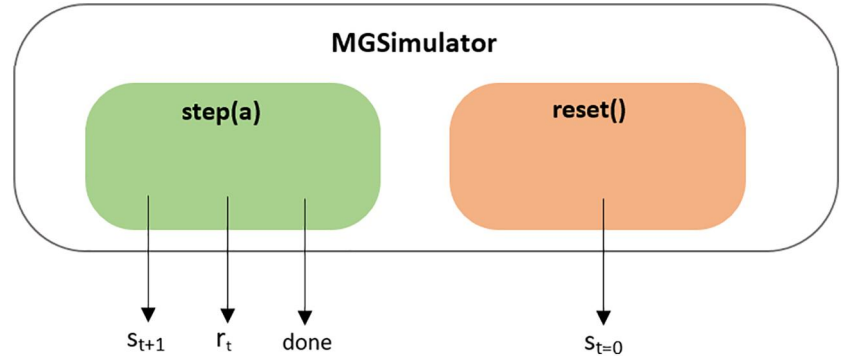
3.3 | MGSimulator

We have created a microgrid simulator called *MGSimulator* and implemented it in Python. The simulator calls a pre-processing module which returns a training and testing net demand dataset. To make *MGSimulator* useful and easily used by the RL community, *MGSimulator* follows OpenAI Gym [33] design and API. This is part of a larger effort to create a generic, easy to use, simulator of microgrids that is available to the communities studying microgrids and RL. The main function of the simulator is named `step(a)`, which takes the agent action $a \in \mathcal{A}$ as input and it returns three elements to the agent: the next state s_{t+1} , the reward r_t , and a Boolean value determining if the terminal state is reached or not (variable name: `done`). A function called `reset()` resets the environment to its initial state. Figure 2 illustrates the different blocks used in the *MGSimulator*.

4 | MICROGRID ENERGY MANAGEMENT SYSTEM AS AN MARKOV DECISION PROCESSES

This section discusses how to transform a microgrid model into an MDP.

FIGURE 2 The MGSimulator architecture. The two main function: step in green and reset in orange are illustrated. Arrows return the results of the functions



4.1 | States

The RL agent uses states to perceive its environment at any given moment. The state space is made of relevant information that the agent can use to take decisions. In our study, two observations compose the state: $s_t = (P_{Net}(t), E_{Bcap}(t)) \in \mathcal{S}$. The net demand $P_{Net}(t)$ defined in Equation (6) is deduced from our PV power output and load time series data. $P_{Net}(t) \in \mathbb{R}$ may have a negative, null, or positive value. The batteries energy $E_{Bcap}(t)$ describes the state of charge (SOC) of the home batteries and follows Equation (10). The constrained values of $E_{Bcap}(t)$ are represented by Equation (8). During an episode, a terminal state is reached when $t = T$ or when the agent takes a bad decision, meaning that constraints are not respected (game over).

4.2 | Actions

To control the microgrid, a set of actions \mathcal{A} is designed. At each time step, the RL agent chooses an action a_t based on s_t . In this study, we propose two sets of actions \mathcal{A}_a and \mathcal{A}_p . Two agents are designed, one called **DDQN-EMS_a** that implements the first action set \mathcal{A}_a , while the second agent **DDQN-EMS_p** implements the priority list set of actions \mathcal{A}_p . The purpose of having two sets of actions is to understand the relationship between the action space and the overall performance.

The first set of actions $a_a \in \mathcal{A}_a$ can only dispatch one generator per time step. In the second case, an action $a_p \in \mathcal{A}_p$ can be a priority list, that is a ranked list of actions, each concerning one generator. Priority lists are a popular algorithm to dispatch generators. It refers to dispatching generators in the order of the list. Once the first generator reaches its maximum power output, the second generator turns on.

Two actions are related to the batteries unit: discharging ($a_a = 0$) or charging ($a_a = 1$ and $a_p = 1$) at full rate of $P_{Net}(t)$. The priority list allows the agent to discharge the battery with the highest priority and then to produce electricity with the genset with the lowest priority ($a_p = 0$). The diesel generator produces electricity ($a_a = 2$ and $a_p = 2$) at full rate of $P_{Net}(t)$. Finally, if the solar energy produced by the PV panels is equivalent to the load consumption in the Equation (6), then the EMS will be in an 'Idle' mode ($a_a = 3$ and $a_p = 3$).

4.3 | Reward function

The main purpose for the DDQN-EMS agents is to optimise the economic cost function J_{obj} of the microgrid system defined in Equation (5). The immediate reward $\mathbf{R}(s, a)$ at time t is associated to the cost of the generators used to meet the net demand $P_{Net}(t)$. In addition, the constraints of the microgrids need to be respected, otherwise either an outage happens, or the load is unwillingly curtailed in Equation (7). As a result, a cost is defined for a constraint-violating decision taken by the agents. This value is not a realistic cost but it is used to penalise such decisions. The reward function is defined as:

$$\mathbf{R}(s, a) = \begin{cases} -m P_{Net}, & \text{if charge batteries} \\ -q P_{Net}, & \text{if power produced by the genset} \\ -c P_{Net}, & \text{if the constraints are not respected} \\ 0, & \text{if do nothing} \end{cases} \quad (12)$$

As the reward function is indexed on the action taken, it is necessary to distinguish two different rewards calculation for **DDQN-EMS_a** and **DDQN-EMS_p** when the action discharge ($a_a = 0$ or $a_p = 0$) is chosen:

- If **DDQN-EMS_a** is used, we add another element in Equation (12): the cost for discharged batteries at full rate of P_{Net} , defined as:

$$\mathbf{R}(s, a) = -m P_{Net} \quad (13)$$

- If **DDQN-EMS_p** is used, the previous reward $\mathbf{R}(s, a)$ of Equation (13) is modified and results in the sum between a cost of discharging batteries denoted by P_{NetBat} and a cost of producing power with the diesel generator to supply the load P_{NetGen} . The reward $\mathbf{R}(s, a)$ is defined as:

$$\mathbf{R}(s, a) = -(mP_{NetBat} + qP_{NetGen}) \quad (14)$$

where

$$P_{Net} = P_{NetBat} + P_{NetGen} \quad (15)$$

4.4 | Transition function

The transition function $\mathbf{T}(s_{t+1}, a_t, s_t)$ is not available, hence unknown to the agents.

5 | METHODS PROPOSED TO HANDLE RARE EVENTS WITH DOUBLE DEEP Q-NETWORK

The two agents described in Section 4 are proposed to tackle the energy management problem of the microgrid, based on the Double Deep Q Network algorithm described in Section 2.

During the learning phase, SRE occurs, which lead the agents to take bad decisions. As a result, the agent must restart the episode at the beginning without the possibility to explore further the episode. We consider the agent ‘stuck’ in such a case, where the agent will always fail at the same time step and thus never learn.

5.1 | Rare events

Rare events are low probability events which significantly affect the expected performance. In our case, they are characterised by a combination of a rare state and a rare action. Figure 3 exhibits rare events (in red) present in our dataset, representing the discrete net demand observed at each time step. We can observe that the two points with a net demand below -20 are not considered as rare events. The reason is that the agent takes the right decision during these two situations and the expected performance is therefore not affected. In our case study, we define rare events when the net demand P_{Net} is

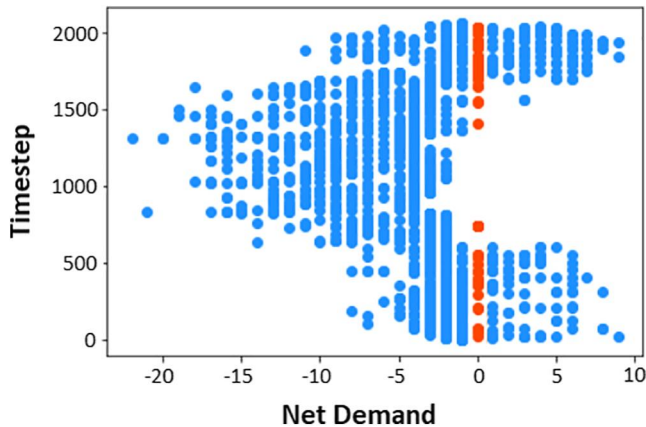


FIGURE 3 Rare events are highlighted by red dots in the dataset. (Blue dots are non rare events.)

equal to zero, that is when the power produced by the PV panels is equal to the power consumption. In Section 4, we have defined action *Idle* ($a = 3$) to effectively manage these situations. The other actions result as bad decisions which cause an outage into the off-grid hybrid microgrid because the power balance constraints in Equation (7) is not respected. The rare events represents 2.66% of our dataset and are classified as SRE.

5.2 | Double deep Q-network-energy management system improvements

In this subsection, we describe the two solutions proposed to unstuck the learning agents during SRE. The first method called **memory counter (MemC)** noted **MemC** is the key to unstuck the DDQN-EMS agent when a rare event occurs. The second method called **combined experience replay** noted **combined experience replay technique (CER)** improves the performance **MemC**. Identifying a rare event is challenging for the agent, and a standard DDQN performs poorly during the learning phase for several reasons:

- training a neural network is a slow process,
- the more time passes, the lesser the opportunity to explore during an episode with the ϵ -decreasing greedy strategy,
- at each episode, the experience replay process requires to sample a small batch of the replay memory \mathcal{D} randomly to train the neural network,
- the size of an episode is long (more than 2000 steps),
- if a bad decision is taken by the agent, it restarts the episode at the beginning.

Altogether, these points raise the necessity of better managing SRE in DDQN. Rare events are under-represented in the replay memory \mathcal{D} because there only represent 2.66% of the dataset. This means that a rare event is less likely to be picked in the batch memory to train the neural network if it has been stored in the replay memory. A solution to unstuck an agent would be to increase its exploration rate, however as epochs pass, the probability to explore decrease. Therefore, as time passes, the agent capability to unstuck itself at a late stage of training becomes null.

To address these problems, we propose two mechanisms that are embedded in the DDQN.

5.2.1 | Double deep Q-network-energy management system with the memory counter capability

We propose to equip DDQN with a new memory called **MemC** of capacity G and denoted by \mathcal{C} (**MemC**). The purpose of \mathcal{C} is to keep track of the maximum step reached during an episode before restarting at the beginning. The capacity of \mathcal{C} is equal to G , meaning that \mathcal{C} consists of the last step numbers of

the last G episodes. For instance, if the agent fails in step $t = 88$, then 88 is added to \mathcal{C} . At each time step t of an episode, a new ratio noted δ_t is calculated. δ_t corresponds to the frequency of the time step ' t ' in \mathcal{C} . As a result, if δ_t is above a certain threshold ρ , the exploration rate ϵ is raised to a higher value $\phi > \epsilon$. ϕ is manually defined by the user. By this way, if the agent is blocked at a certain moment of an episode without the possibility to explore because of a too small ϵ , then this mechanism forces the agent to explore more.

5.2.2 | Double deep Q-network-energy management system with memory counter and combined experience replay technique capabilities

The size of the replay memory \mathcal{D} plays an important role in the performance of the RL agent. We have used the CER proposed in Reference [34] by forcing the last element of \mathcal{D} in the batch experience replay to be sure that if a rare event occurs; then, this experience will be selected to train the neural network. This mechanism always forces the agent to train with the last experience stored in \mathcal{C} . It is a simple but effective trick. We summarise the two capabilities in Algorithm 2.

Algorithm 2 DDQN-EMS with MemC and CER capabilities

Initialise empty replay memory \mathcal{D} to size capacity R .
 Initialise empty counter step memory \mathcal{C} to size capacity G .
 Initialise current Q network with random weights w .
 Initialise target Q network with weights $w^- \leftarrow w$.
for each episode **do**.
 Initialise $t \leftarrow 0$.
 Initialise *MGSimulator*.

while Terminal State not reached **do**.
 if $\delta_t > \rho$ **then**.
 Force $\epsilon \leftarrow \phi$
 end if
 Choose action a_t using ϵ -greedy strategy.
 Emit action a_t in *MGSimulator*, observe r_t, s_{t+1}
 Store transition $\text{exp}_t = (s_t, a_t, r_t, s_{t+1})$ in \mathcal{D}
 $t \leftarrow t + 1$.
end while.
 Sample random batch of $N - 1$ transitions from \mathcal{D}
 Add the latest transition stored in \mathcal{D} to the batch.
if episode ends at s_{t+1} **then**.
 $Q(s_t, a_t) \leftarrow r_t$
Else.
 Set $Q(s_t, a_t)$ according to Equation (4)
end if.
 Perform gradient descent step using Equation (3)
 Every τ steps, update weights $w^- \leftarrow w$ weights.
 Store the time step t in \mathcal{C}
end for.

6 | EXPERIMENTS

6.1 | Experimental settings

The dataset used in this study comes from different sources: the load measurements from an office building, the Drahi X Novation Center, and the PV generation measurements from the SIRTÀ atmospheric laboratory [35]. Figure 4 illustrates the PV power and load consumption profiles in the data. The sensors are located at the same place on the campus of the École Polytechnique in Palaiseau, France. The agent uses historical data and interacts with the simulator *MGSimulator* (cf. Section 3). The code of our study is available in open source:

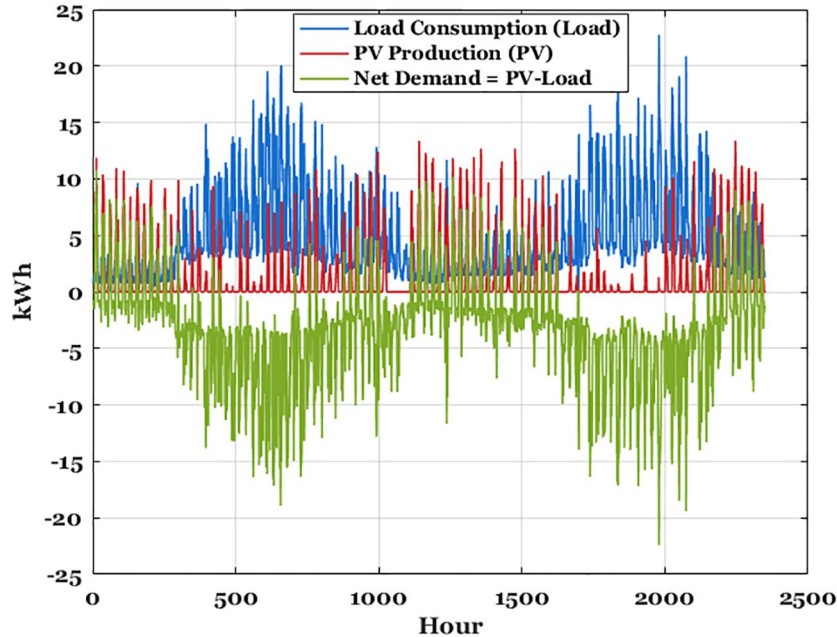


FIGURE 4 The dataset of load consumption, PV power production and net demand

<https://bit.ly/3jmaasD>. First the agent is trained. This period is called the *training phase* or the *learning phase*. When this phase is over, we test the learned agent on new conditions of the microgrid. This is called the *testing phase*. The training phase is conducted on simulating 43 active days of the microgrid, which are randomly selected over 43 successive weeks along the year. A time step of 30 min is taken during the training phase which represents $2 \times 24 \times 43 = 2064$ steps for one episode. We have chosen to set the time step to half an hour to increase the amount of training data, which is enough to have good performance without increasing too much the computing time. The more data we have in an episode, the more time is needed to train our agent. The testing phase is performed on a new series of 52 days, over 52 successive weeks along the year, with a time step of 1 h. Note that the testing dataset represents a large variety of days, including seasonal variations, holidays, etc.

The parameters of the objective function J_{obj} (which are also used in the reward function) are set to $m = 0.5$ €, $q = 1.5$ €, and $c = 10$ €. Table 1 provides the parameters of the simulated microgrid.

The hyperparameters used for the two agents: DDQN-EMS_a and DDQN-EMS_p, are given in Table 2. The selection of these hyperparameters is sensitive because they affect directly the learning performance. DDQN-EMS_p (with a priority list of actions) needs more episodes to converge because the estimation of Q is more difficult to learn. We have noticed that to obtain better performances with DDQN-EMS_p, the size of the

TABLE 1 Specifications for microgrid components

	Energy capacity (kWh)	Rated power (kW)
PVs	-	15
Batteries	90	42
Genset	-	12
Load	-	22

TABLE 2 Specifications for microgrid components

	DDQN-EMS _a	DDQN-EMS _p
Learning rate α	0.001	0.00025
Epsilon max ϵ	1.0	1.0
Epsilon min ϵ	0.001	0.001
Gamma γ	0.5	0.5
Memory size R	500	20,000
Batch size N	32	32
Number of episodes	3500	50,000
$\hat{Q} = Q$ every τ (step)	Episode	5000
Step memory size G	10	10
Ratio limit ρ	0.8	0.8
Epsilon MemC ϕ	0.99	0.99

replay memory \mathcal{D} has to be made larger. The minimum value of ϵ is fixed at 0.001 in order to give the agents the possibility to reach the successful terminal state s_T . Indeed, as the agent restarts at s_0 when a bad decision is made, having a minimum ϵ at 0.1 is too high to consider reaching the 2064th step. The neural networks for the target and Q networks consist of three hidden layers of 100 neurons each, each with a Rectified Linear Unit (ReLU) activation function. We tested different architectures and this one gives the best performance. We train the neural network with Adam, an adaptive learning rate optimisation algorithm [36] widely used in the DRL domain.

Abbreviations: DDQN, double deep Q-network; EMS, energy management system.

6.2 | Memory counter and combined experience replay capabilities experiments

To validate the improvements made in the DDQN-EMS algorithm, we have tested four agents (based on the DDQN-EMS_a settings) with different capabilities, on the same microgrid environment with rare events: a basic DDQN agent (without any improvement), an agent with only the combined experience replay (**CER**), an agent with only the MemC and finally an agent with the both capabilities. The purpose is to compare their ability to learn how to control an off-grid hybrid microgrid effectively when SRE occur. Each agent is tested over 1400 episodes, corresponding to a learning phase. Each episode consists of 2064 steps maximum. Each agent runs 10 learning phases in order to understand the variation in its performances. Figure 5 displays the performance of each agent along training. The y -axis is the average of steps reached over the last 100 episodes and is on a logarithmic scale. The cloud around the average line represents the 10% and 90% percentiles.

The basic DDQN agent without improvements and the agent with only **CER** capability fail early, not reaching a time step above 25–30 in any episode; hence, they can not learn a good strategy. **CER** alone does not help the learning agent and performs equivalently to the basic DDQN-EMS. The red curve shows the performance of DDQN-EMS equipped with **MemC**: this agent enhances significantly the performance of the basic DDQN-EMS agent with 410 steps performed on average at the 1400th episode. Finally, combining both **MemC** and **CER**, DDQN-EMS outperforms **MemC** by 75%, an average of about 720 steps being reached at each episode at the end of the learning/training phase. We conclude that the combine experience replay improves significantly the **MemC** capability. Hence, their combination performs very well.

Finally, Figure 6 illustrates how many times an agent is stuck in the microgrid environment. We consider that an agent is stuck if during the last 10 episodes, the agent fails at the same time step. We show that both the basic DDQN-EMS and the agent with only **CER** are stuck and do not manage rare events. The agent with the **MemC** mechanism is able to get unstuck. With regards to the basic DDQN-EMS agent, we obtain impressive results with an enhancement of 686% and 1042%

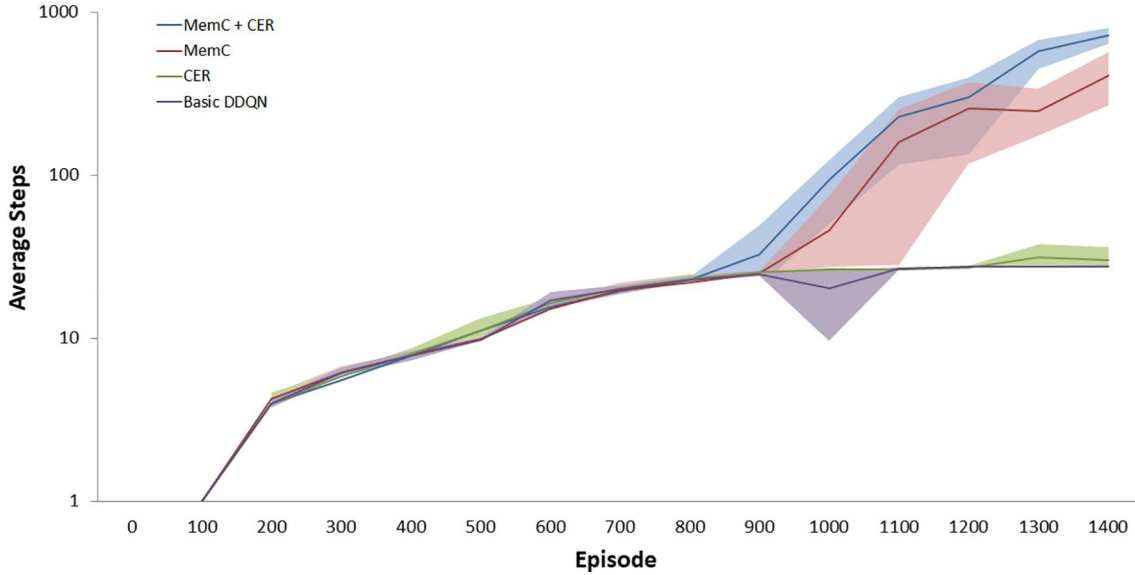


FIGURE 5 Comparison of the learning performance of the basic double deep Q-network (DDQN)-energy management system (EMS) and the three proposed variants of DDQN-EMS_{*a*} with combined experience replay (**CER**), with the memory counter (**MemC**), with both (**MemC + CER**)

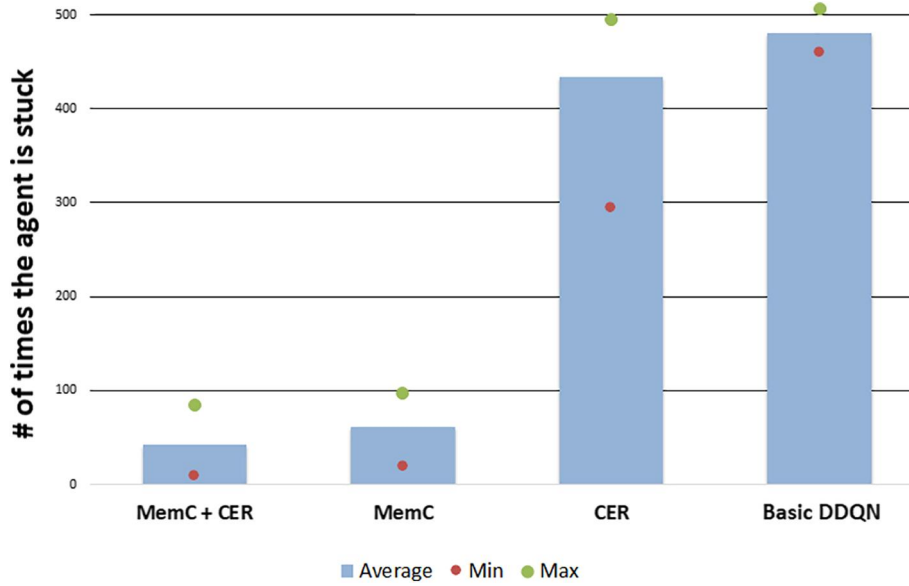


FIGURE 6 Double deep Q-network (DDQN)-energy management system (EMS) performance improvement with regards to the used mechanism: average number of times the learning agent is stuck. The lower, the better. The minima (red squares) and the maxima (green squares) show the variability in the performance of the agent during a learning phase

for the agents equipped with **MemC** and **MemC + CER**, respectively. We notice that the **CER** agent obtains the larger variance during its learning phases, which is greatly reduced with **MemC**.

6.3 | Training phase

Now that we have studied and found mechanisms to have the agent able to learn during whole episodes, we focus on its performance to control a microgrid. First, we consider the

training phase of the DDQN-EMS agents and we compare the two types of action sets. Then in the next section, we compare DDQN-EMS agents with other algorithmic approaches.

Each version of the DDQN-EMS agent (DDQN-EMS_{*a*} without the priority list of actions and DDQN-EMS_{*p*} with it) has a training phase to find a good estimate of Q . Each version is equipped with both **MemC** and **CER** mechanisms. To deal with the exploration-exploitation dilemma, the agent uses the ϵ -decreasing greedy strategy, explained in Section 2. To validate that the agent is learning correctly, we examine the learning curve: Figure 7 shows how the performance of DDQN-EMS_{*a*}

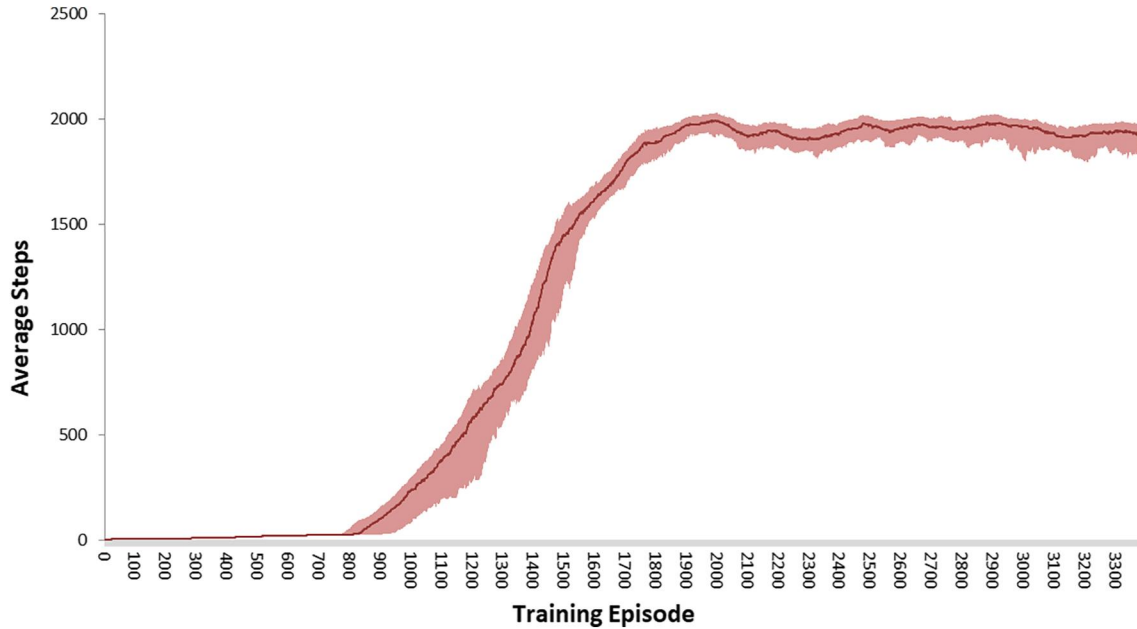


FIGURE 7 Learning curve of the DDQN-EMS_a algorithm

improves along training episodes, averaged over 100 episodes. The computation time (with an Nvidia GeForce GTX 1080) for a training phase of DDQN-EMS_a is around 30 min and more than 7 h for the DDQN-EMS_p. This time gap is relative to the number of episodes. Figure 7 illustrates the average performance and its variability for agent DDQN-EMS_a: 100 trainings are performed, each one giving a learning curve. The cloud around the red line represents the 10% and 90% percentiles. Only the DDQN-EMS_a is illustrated in this work because it requires too much time (more than 700 h) to have the same figure for DDQN-EMS_p. The shape of the learning curve function is common for RL agents and validates their ability to learn: initially, the agent explores its environment and does not perform well at all; then, after a while, there is a rapid improvement of the performance, leading to a phase during which the performance stagnates again. At this point, the agent has learnt a good (if not optimal) policy. By doing ‘test and error’ over a lot of hyper-parameter combinations, we have succeeded to have a low variance around the average. With all the tests we have done for DDQN-EMS_p, we can be pretty sure that we have the same type of learning curve than the one obtained with DDQN-EMS_a.

6.4 | Testing phase

In the testing phase, the trained agents are facing a new dataset, that is new conditions under which to control the microgrid. We measure their performance on these new conditions. The cumulative cost obtained to control our off-grid hybrid microgrid is reported. At test time, the agent does not perform exploration: actions are selected greedily. To assess DDQN-EMS_a and the DDQN-EMS_p agents, we compare their performance with those obtained by other methods. The first of

these methods is a RL Decision Tree (RL-DT) algorithm, proposed in Reference [14]. This method is a combination of a Q-Learning and a CART decision tree. The second method is a hand-crafted rule-based control algorithm. The third method is a DP algorithm called Value Iteration, as proposed in Reference [9]. DP algorithms give an optimal policy given a perfect model of the environment. We provide the DP agent with the real net demand P_{Net} of the next 24 h. The implemented DP algorithm does not have a priority list of actions.

Figure 8 represents the cumulative cost for an episode horizon $T = 1248$ averaged over 10 runs. This plot shows that the two versions of the DDQN-EMS outperform both the RL-DT algorithm (6491.8€) and the hand-crafted rule-based (6498.5€). DDQN-EMS_a obtains an average cost of 6463.5€, while the DP approach obtains a lower average cost of 6452.5 €. Thanks to its priority list of actions, DDQN-EMS_p outperforms all the methods, with an average cumulative cost of 6436.5€. DDQN-EMS_a has a standard deviation of 0.81€ and DDQN-EMS_p 3.9€. The two versions of DDQN-EMS exhibit very small variation of the final performance, hence a high robustness.

Table 3 illustrates the computational time to take each decision during the testing phase. The rule-based, RL-DT and DDQN-EMS algorithms take a decision almost instantly (with best time performance for the rule based), while it takes DP method 1000 more time to select an action. It is worth mentioning that DP is relatively fast in this simple case. However, the method will not scale to a more complex system (more generators and possible actions) at a smaller time step. The advantage of the DDQN-EMS method is once trained, it obtains better results with almost equivalent or less computational time to take actions than the benchmark methods. The second advantage is that even if the microgrid complexity increases by adding further generators,

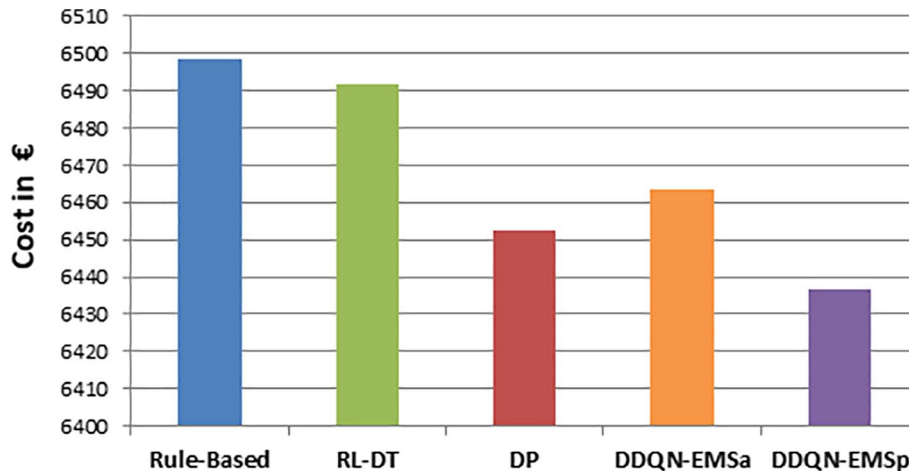


FIGURE 8 Performance of the competing algorithms: see text for details

TABLE 3 Computational performance to select one action

	Time (s)
Rule-based	$8.5 \times e^4$
RL-DT	$1.7 \times e^3$
DP	0.3
DDQN	$1.0 \times e^3$

Abbreviations: DDQN, double deep Q-network; DP, dynamic programming; RL-DT; reinforcement learning decision tree.

the computation time will not change for the DDQN-EMS and RL-DT. In addition, the DDQN-EMS does not use a forecaster and a model to take decisions, whereas the DP method needs predictions and a solver, which means that if the system changes, the DP controller needs to be design again, which is not effective. Finally, it is important to notice that the DDQN-EMS learns its environment in an offline way with simulations and afterward implements its knowledge during a testing phase in an online fashion: the process is simulated in our case but it could be implemented in a real off-grid hybrid microgrid.

7 | CONCLUSION AND FUTURE WORKS

In this paper, we have optimised the operational cost of an off-grid hybrid microgrid. We assume that no forecaster is available. This choice is motivated by the difficulty of the forecasting task. We investigate RL algorithms in this context. The specificity of this study is by nature associated with rare events during an episode. We have noticed that rare events create problems during the learning phase of an agent. As a result, we have proposed a novel approach by merging an existing method called **CER** and a new method called **MemC**. We have experimentally demonstrated that this approach is efficient to unstuck the agent during the learning phase. We

show experimental results for our proposed algorithms and standard algorithms for energy management (including rule-based and dynamic programming). Experiments are based on a simulated microgrid with real data. As usual in machine learning, the data used for training is different from the data used to assess the performance of the algorithms: they have to generalise their knowledge from the training data and hence be able to cope with unseen situations. We have demonstrated that our proposed DDQN-EMS agents succeed to adapt themselves in a new testing year with good performances, outperforming standard methods. The proposed algorithm could be used without modification for an other off-grid hybrid microgrid. Nevertheless, an improvement of this work would be to design different microgrid architectures (islanded and connected to the grid with different dataset), with more or less complexity, in order to study the scalability and the robustness of this method. Regarding the different components of a microgrid, the MDP has to be transformed to match the new power system environment. Another next step would be to validate the performance of the proposed method in a physical test. New perspectives should be explored by adapting the agent action into multiple actions like we have proposed with the priority list mechanism. Comparing DDQN with Policy Gradient methods is an other interesting track for further research.

ACKNOWLEDGEMENTS

This work was supported by the CSE & AI R&D of Total S.A and the TREND-X research program of École polytechnique. The measurements from the SIRTA's (<http://sirta.ipsl.fr>) outdoors photovoltaic test bench are done in a joint collaboration of four laboratories (GeePs, LMD, Limsi and LPICM). Philippe Preux acknowledges the support of Inria and the members of SequeL for the very friendly, fruitful, stimulating, and exciting atmosphere of the group. Tanguy Levent thanks the LPICM laboratory of École Polytechnique and the CSE & AI R&D of Total S.A team for the equipment, the fascinating discussions and their delightful supports.

ORCID

Tangyu Levent  <https://orcid.org/0000-0003-4394-7133>

REFERENCES

- United Nations Treaty Collection.: Enviroment, Chapter XXVII, 7. d Paris Agreement (2015). https://treaties.un.org/pages/ViewDetails.aspx?src=TREATY&mtdsg_no=XXVII-7-d&chapter=27&clang=en
- Frankfurt School-UNEP Collaborating Centre for Climate & Sustainable Energy Finance, Bloomberg New Energy Finance: UN Environment: Global trends in renewable energy investment. (2018). <https://europa.eu/capacity4dev/file/71900/download?token=57xpTJ4W>
- NRG.: The Evolution of Distributed Energy Resources. Microgrid Knowledge (2018)
- Chu, S., Majumdar, A.: Opportunities and challenges for a sustainable energy future. *Nature*. 488(7411), 294–303 (2012). <https://doi.org/10.1038/nature11475>
- Lasseter, R.H., et al.: Integration of Distributed Energy Resources: The Certs Microgrid Concept. CERTS (2003)
- Lasseter R.H.: Microgrids. In: IEEE power Engineering Society Winter meeting Conference Proceedings, vol. 1, pp. 305–308. IEEE (2002)
- Hatzigrygiou, N., et al.: Microgrids. *IEEE Power Energy Mag.* 5(4), 78–94 (2007)
- Lasseter, R.H.: Paigi, P. Microgrid: a conceptual solution. In: 2004 IEEE 35th Annual Power Electronics Specialists Conference (IEEE Cat. No.04CH37551), vol. 6, pp. 4285–4290. IEEE (2004)
- Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction, 2nd ed. The MIT Press (2018). <http://incompleteideas.net/book/the-book-2nd.html>
- Paul, S., Osborne, M.A., Whiteson, S.: Fingerprint policy optimisation for robust reinforcement learning. In: Proceedings of the 36th International Conference on machine learning (ICML), pp. 5082–5091. PMLR 97 (2019). <http://proceedings.mlr.press/v97/paul19a.html>
- Frank, J., Mannor, S., Precup, D.: Reinforcement learning in the presence of rare events. In: Proceedings of the 25th International Conference on Machine Learning (ICML'08), pp. 336–343.(2008). <https://doi.org/10.1145/1390156.1390199>
- Ciosek, K., Whiteson, S.: OFFER: off-environment reinforcement learning. In: The Thirty-first AAAI Conference on artificial Intelligence (AAAI-17) (2017)
- Agüera.Pérez, A., et al.: Weather forecasts for microgrid energy management: review, discussion and recommendations. *Appl Energy*, 228,265–278 (2018)
- Levent, T., et al.: Energy management for microgrids: a reinforcement learning approach. In: 2019 IEEE PES Innovative smart grid Technologies Europe (ISGT-Europe), pp. 1–5. IEEE (2019)
- Rocchetta, R., et al.: A reinforcement learning framework for optimal operation and maintenance of power grids. *Appl Energy*. 241, 291–301 (2019)
- Mocanu, E., et al.: On-line building energy optimization using deep reinforcement learning. *IEEE Trans. Smart Grid.* 10(4), 3698–3708 (2019)
- Lu, R., Hong, S., Zhang, X.: A dynamic pricing demand response algorithm for smart grid: reinforcement learning approach. *Appl Energy*. 220, 220–230 (2018)
- Berlink, H., Costa, A.H.R.: Batch reinforcement learning for smart home energy management. In: Proceedings of the 24th International Conference on Artificial Intelligence. IJCAI'15., pp. 2561–2567. AAAI Press (2015)
- Kuznetsova, E., Li, Y.F., Ruiz, C., Zio, E., Ault, G., Bell, K.: Reinforcement learning for microgrid energy management. *Energy*. 59, 133–146 (2013). <http://www.sciencedirect.com/science/article/pii/S0360544213004817>
- Mbuwir, B.V., et al.: Battery energy management in a microgrid using batch reinforcement learning. *Energies* 10, 1846 (2017)
- Venayagamoorthy, G.K., et al.: Dynamic energy management system for a smart microgrid. *IEEE Trans. Neural Netw. Learn. Syst.* 27(8), 1643–1656 (2016)
- François-Lavet, V., et al.: Deep reinforcement learning solutions for energy microgrids management. In: 2016 European Workshop on Reinforcement Learning (EWRL) (2016)
- Glavic, M., Fonteneau, R., Ernst, D.: Reinforcement learning for electric power system decision and control: past considerations and perspectives. *IFAC-PapersOnLine*. 50(1), 6918–6927 (2017). 20th IFAC World Congress. <http://www.sciencedirect.com/science/article/pii/S2405896317317238>
- Tesauro, G.: Temporal difference learning and td-gammon. *Commun ACM*. 38(3), 58–68 (1995). <https://doi.org/10.1145/203330.203343>
- Yu, C., Liu, J., Nemat, S.: Reinforcement learning in healthcare: a survey. *ArXiv preprint, arXiv: abs/1908.08796* (2019)
- Mnih, V., et al.: Playing atari with deep reinforcement learning. *arXiv: abs/1312.5602* (2013). <http://arxiv.org/abs/1312.5602>
- Silver, D., et al.: Mastering the game of go without human knowledge. *Nature*. 550, 354 (2017). <https://doi.org/10.1038/nature24270>
- Melo, F.S.: Convergence of q-learning: a simple proof. *Institute Of Systems and Robotics, Technical Report* pp. 1–4 (2001)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097–1105. Curran Associates, Inc. (2012)
- Goodfellow, I.J., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge, (2016). <http://www.deeplearningbook.org>
- Mnih, V., et al.: Human-level control through deep reinforcement learning. *Nature*. 518(7540), 529–533 (2015). <https://doi.org/10.1038/nature14236>
- Hasselt, H.V., Lafferty, J.D., Williams, C.I.K.: Shawe: Double q-learning. In: Taylor, J., Zemel, R.S., Culotta, A. (eds.) *Advances in Neural Information Processing Systems*, vol. 23, pp. 2613–2621. Curran Associates, Inc. (2010). <http://papers.nips.cc/paper/3964-double-q-learning.pdf>
- Brockman, G., et al: Openai gym. *ArXiv:1606.01540* (2016). <http://arxiv.org/abs/1606.01540>
- Zhang, S., Sutton, R.S.: A deeper look at experience replay. *ArXiv:abs/1712.01275* (2017)
- Haefelin, M., et al.: Sirta, a ground-based atmospheric observatory for cloud and aerosol research. *Ann Geophys.* 23(2), 253–275 (2005). <https://www.ann-geophys.net/23/253/2005/>
- Kingma, D., Ba, J.: Adam: a method for stochastic optimization. In: *International Conference on Learning Representations* (2014)

How to cite this article: Levent T, Preux P, Henri G, Alami R, Cordier P, Bonnassieux Y. The challenge of controlling microgrids in the presence of rare events with deep reinforcement learning. *IET Smart Grid*. 2021;1–14. <https://doi.org/10.1049/stg2.12003>