energies

Article

Cooperative Wind Farm Optimization Using Policy Search Reinforcement Learning

Yasser Bin Salamah

Department of Electrical Engineering, King Saud University, Riyadh 11451, Saudi Arabia; ybinsalamah@ksu.edu.sa

Abstract

This paper introduces a policy-search-based reinforcement learning algorithm aimed at generating optimal set-points of wind turbines in wind farms. The proposed approach addresses the problem of multivariable optimization in systems where the objective function is unknown or difficult to model. The algorithm is a model-free framework and relies solely on measured performance of the system. Namely, it does not require gradient information of the objective function or an explicit model of the aerodynamic interaction between wind turbines. The proposed scheme utilizes stochastic policy perturbations to explore the search space and update the policy parameters directly based on the observed reward signal. In this way, the algorithm progressively drives the control variables toward optimal operating conditions. The proposed policy-search reinforcement learning framework is analyzed to establish its connection with gradient-free optimization methods. The proposed method is applied to wind farm power optimization, where multiple turbine control variables must be adjusted in the presence of wake interactions cooperatively. The performance of the proposed approach is evaluated through extensive simulations under both steady-state and time-varying wind conditions. The proposed algorithm is compared with an extremum-seeking control method that was previously suggested for the same problem. The results demonstrate that the proposed approach is able to effectively maximize power production in wind farms while maintaining a simple and model-free optimization structure.

Keywords: gradient-free optimization; policy search reinforcement learning; wind farm optimization; renewable energy systems; extremum seeking control

1. Introduction

1.1. Motivation

Wind energy has become one of the most important renewable energy sources in modern power systems. Large-scale wind farms are widely deployed to increase energy production while reducing installation and operating costs. By grouping multiple wind turbines at a single location, wind farms allow for the efficient use of land resources and facilitate the large-scale generation of clean energy. However, despite these advantages, the operation of wind farms presents significant control and optimization challenges. For example, when wind flows through a wind farm, upstream turbines extract kinetic energy from the airflow. As a result, downstream turbines experience reduced wind speed due to wake effects created by upstream machines [1,2]. Figure 1 shows a representation of wind turbines grouped as series-connected turbines considering full wake interaction under different wind directions. These aerodynamic interactions introduce strong coupling among turbines, which significantly affects the overall power production of the wind farm.

Academic Editor: Tonio Sant

Received: 18 March 2026
Revised: 25 April 2026
Accepted: 26 April 2026
Published: 29 April 2026

Furthermore, accurate modeling of wake interactions remains a challenging problem; the wake dynamics depend on several factors, including wind conditions, turbine spacing, and atmospheric turbulence. Although various wake models have been proposed, their accuracy can vary significantly under different operating conditions. These challenges motivate the development of control strategies designed to address two primary objectives:

1. Maximizing total power output cooperatively by accounting for the complex wake interactions and aerodynamic coupling between turbines.
2. Utilizing model-free methods to eliminate the need for complex nonlinear model development, as these phenomena are notoriously difficult to characterize analytically.

Motivated by these considerations, this paper proposes a cooperative policy-search reinforcement learning (RL) framework for the multivariable optimization of wind farms. The proposed approach is inherently model-free, relying solely on the measured power output of the wind farm to drive the optimization process.

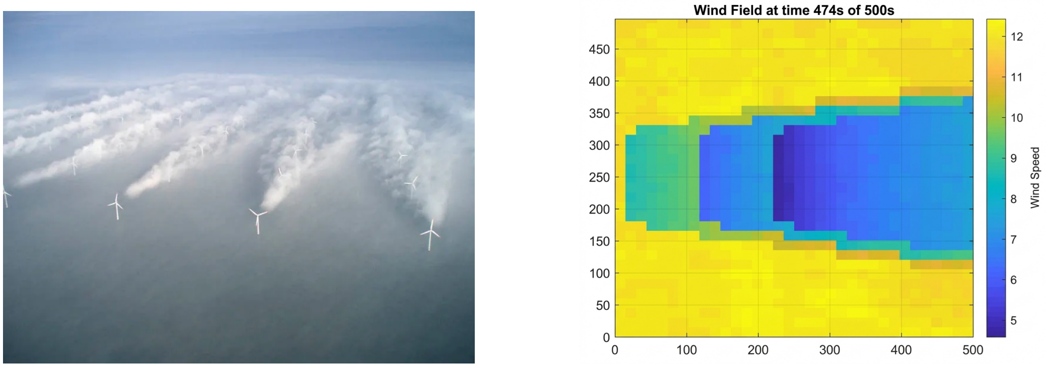

(**a**) (**b**)

Figure 1. Illustration of wind turbine wake effects: (**a**) real offshore wind farm showing wake patterns (Denmark Horns Rev Offshore Wind Farm [3] (credit: Christian Steiness)) and (**b**) simulated wake field of three turbines.

1.2. Related Work

Wind farm optimization has attracted significant attention in recent years due to the increasing deployment of large-scale wind energy systems. In general, optimization problems in wind farms can be classified into two main categories: wind farm layout optimization and operational wind farm control [2,4].

Wind farm layout optimization focuses on determining the optimal placement of turbines within a wind farm in order to maximize energy production while minimizing wake interference among turbines. Various optimization techniques have been proposed for this problem, including evolutionary algorithms, particle swarm optimization, and gradient-based methods [5,6]. These approaches are typically applied during the design stage of a wind farm and are therefore performed offline.

The second category concerns the operational optimization of wind farms. In this case, the objective is to maximize the total power output of the wind farm during operation by adjusting the control variables of individual turbines. In contrast to layout optimization, operational optimization must account for the dynamic aerodynamic interactions between turbines. Several control variables have been considered in the literature for this purpose, including yaw angle control [7,8], pitch control [9], and axial induction control [10]. Among these approaches, induction-based control has received considerable attention since it directly influences the aerodynamic wake interaction among turbines. In this work, the focus is placed on induction-based control strategies for wind farm optimization.

Different optimization and control methods have been proposed to address the wind farm control problem. One class of methods relies on model-based optimization approaches. These approaches use wake models to estimate aerodynamic interactions between turbines and compute optimal operating conditions accordingly [11,12]. Although model-based methods can provide useful insights, their performance strongly depends on the accuracy of the wake models used. In practice, wake dynamics are influenced by several environmental factors such as wind speed variations, atmospheric turbulence, and changing wind directions, which makes accurate modeling challenging.

To overcome these limitations, several model-free optimization techniques have been proposed. Extremum-seeking control (ESC) is one of the most widely studied approaches for real-time optimization of wind farms [13]. ESC methods employ perturbation signals to estimate gradient information of the objective function and iteratively drive the system toward optimal operating conditions. Various ESC-based algorithms have been proposed in the literature, including sliding-mode extremum-seeking schemes that improve convergence properties and robustness [1,14]. In addition, several distributed ESC [15–18] that build on the sinusoidal-perturbation approach have been analyzed. A pure optimization-based approach was proposed in [19], where the authors use ADMM to enable distributed, real-time power optimization of wind farms with guaranteed convergence.

Game-theoretic approaches have also been investigated to model the interaction between wind turbines in a wind farm [20]. In such formulations, each turbine is treated as an individual decision-making agent that seeks to optimize its own objective. However, these approaches often converge to Nash equilibrium solutions, which may not necessarily maximize the total power output of the wind farm. In a related study, a cooperative game-theoretic framework was introduced in [21]. In that work, both the yaw misalignment angles and the axial induction factors were adjusted to enhance the overall power output of the wind farm.

More recently, reinforcement learning (RL) has emerged as a promising tool for solving complex control and optimization problems. Several RL-based approaches have been proposed for wind farm optimization [4]. Early farm-level RL studies primarily relied on value-based or tabular methods, particularly distributed Q-learning for yaw control under wake delays [22]. More recent developments include Q-learning and deep reinforcement learning approaches [23,24], which enable the learning of optimal control policies directly from system measurements without requiring explicit wake models. However, many of these approaches rely on complex function approximators and require extensive training data, which may limit their practical deployment.

In the context of cooperative control, ref. [25] proposed a deep reinforcement learning framework with knowledge-assisted learning for wind farm optimization. Nevertheless, many traditional cooperative RL frameworks implicitly capture inter-agent interactions without explicitly modeling the underlying coupling structure. To address this limitation, recent research has explored graph-based representations. For instance, ref. [26] demonstrates that integrating graph neural networks (GNNs) with multi-agent reinforcement learning (MARL) enables explicit modeling of topological relationships and interaction strengths among agents, thereby improving cooperative performance. In wind farm applications, such graph-based formulations are particularly relevant, as turbines naturally form a spatially coupled system governed by wake interactions.

1.3. Contributions

Motivated by these observations, this paper investigates the use of a policy-search reinforcement learning framework for cooperative wind farm optimization. In contrast to greedy strategies that optimize individual turbines independently, the proposed approach directly optimizes the total power output of the wind farm. Furthermore, the proposed method is compared with a previously developed sliding-mode extremum-seeking controller to highlight the similarities and differences between reinforcement learning and extremum-seeking optimization methods. This work builds upon our previous efforts to develop robust control and optimization strategies for complex industrial systems [1,14]. Namely, we design and analyze an RL framework that bridges the gap between stochastic search and classical optimization through a direct, gradient-free formulation. The key contributions of this paper are summarized as follows:

1. We propose a framework that utilizes stochastic policy perturbations to explore the search space. This enables the update of policy parameters directly based on measured system performance, allowing for the optimization of high-dimensional systems without requiring an explicit mathematical model of the plant or the objective function.
2. We provide a formal analysis of the proposed framework, proving that the stochastic policy update approximates gradient ascent on a smoothed version of the objective function. This result establishes a fundamental mathematical connection between policy-search reinforcement learning and classical gradient-free methods, such as Extremum-Seeking Control.
3. Application to Wind Farm Power Generation: We demonstrate the framework's effectiveness by applying it to a wind farm setup. The controller successfully optimizes multiple turbine control parameters simultaneously, effectively managing the complex aerodynamic coupling and wake interactions—phenomena that typically present significant modeling challenges in renewable energy contexts.
4. We conduct a comparative study using sliding-mode extremum-seeking control (SM-ESC) and the REINFORCE policy-gradient method as benchmark approaches. The results highlight the structural similarities between gradient-free optimization and policy-search reinforcement learning, while demonstrating competitive performance and improved convergence behavior of the proposed PSRL framework under the considered operating conditions.

1.4. Organization

The rest of this paper is organized as follows. Section 2 presents the wind farm model and formulates the associated multivariable optimization problem. Section 3 introduces the proposed policy-search reinforcement learning framework and discusses its convergence properties. Section 4 provides simulation results and evaluates the performance of the proposed approach through a comparison with the sliding-mode extremum-seeking control method. Finally, Section 5 concludes the paper and discusses potential directions for future research.

2. Wind Farm Model and Optimization

This section describes the formulation of the wind farm optimization problem. In a wind farm, wind turbines extract kinetic energy from the incoming airflow. As the air passes through the turbine rotor, part of its energy is converted into electrical power, which results in a reduction of the wind velocity downstream of the turbine. This reduction in wind speed behind a turbine is known as the wake effect. The wake produced by an upstream turbine influences downstream turbines by decreasing the effective wind speed reaching them, which consequently reduces their power generation capability.

Accounting for aerodynamic interactions among turbines is therefore essential when optimizing wind farm performance. Recent studies [27] demonstrate that incorporating wake interactions in the control strategy can lead to energy gains exceeding ten percent compared with the standard controller developed by the National Renewable Energy

Laboratory (NREL). In the following subsections, we first review the wind turbine control framework and then formulate the wind farm optimization problem. For clarity, Table 1 summarizes the notation used throughout the paper.

Table 1. Summary of notation used in this paper.

Symbol	Description
$\theta \in \mathbb{R}^n$	Vector of axial induction factors
θ_i	Axial induction factor of turbine i
$J(\theta)$	Total wind farm power (objective function)
$J_i(\theta)$	Power generated by turbine i
$\mu \in \mathbb{R}^n$	Policy mean (decision variable in PSRL)
σ	Exploration variance
$\epsilon \sim \mathcal{N}(0, I)$	Gaussian perturbation vector
α_k	Learning rate at iteration k
N	Number of sampled candidates per iteration
R_i	Reward corresponding to sampled input θ_i
$C[j, i]$	Wake interaction coefficient between turbines j and i
V_∞	Free-stream wind speed
$V_i(\theta)$	Effective wind speed at turbine i
ρ	Air density
A_i	Rotor swept area of turbine i
$C_p(\theta_i)$	Power coefficient

In this work, the wind farm is modeled as an array of n interconnected wind turbines where the wake interactions among turbines are fully considered. As illustrated in Figure 1, this configuration does not restrict generality, since many wind farm layouts can be rearranged into equivalent series-connected turbine groups depending on the wind direction [10].

2.1. Wind Turbine Control

Several control strategies for wind turbines have been proposed in the literature. The most common approaches include yaw control, pitch control, and generator torque control. Yaw control adjusts the orientation of the rotor relative to the wind direction. However, this action typically operates on a slow time scale and therefore does not influence the wind farm power output as rapidly as other control mechanisms.

In contrast, pitch control and torque control can be coordinated through the *axial induction factor*. The axial induction factor represents a dimensionless measure of the reduction in wind velocity between the free-stream flow and the rotor plane. It is widely used as a high-level control variable for regulating power production in wind turbines (see, for example, ref. [28]).

The axial induction factor is determined by two low-level controllable variables: the blade pitch angle and the tip-speed ratio λ. The tip-speed ratio is defined as

$$\lambda = \frac{\omega R}{V}, \tag{1}$$

where V is the undisturbed wind speed, ω is the rotor angular velocity, and R denotes the rotor radius.

Wind turbine control strategies depend strongly on the operating wind speed. When the wind speed is below the rated value (referred to as Region-2 operation), the turbine typically regulates the tip-speed ratio using generator torque control to maximize energy capture. When the wind speed exceeds the rated value, the blade pitch angle becomes the primary control input to maintain safe operation. Detailed descriptions of turbine operating regions can be found in [29]. These mechanisms collectively determine how the axial induction factor can be adjusted.

Remark 1. *This work focuses on the design of a high-level controller that generates reference set-points for the axial induction factor. These references are assumed to be tracked by the turbine's low-level control loops according to the turbine specifications.*

2.2. Wake Interaction Model

One of the main challenges in wind farm optimization lies in accurately modeling the wind velocity field while accounting for wake interactions generated by multiple turbines. Among the various wake models proposed in the literature, the Park wake model [20,30] is widely adopted due to its simplicity and widespread use in the wind farm control literature as a benchmark model for evaluating optimization and control strategies. It is important to emphasize that the purpose of employing this model is not to provide a high-fidelity representation of wake dynamics, but rather to establish the feasibility and effectiveness of the proposed control framework under a commonly accepted test environment.

In this work, the Park wake model is adopted. Accordingly, for turbine i, the effective wind velocity is described as

$$V_i(\theta) = V_\infty(1 - \delta V_i(\theta)), \tag{2}$$

where V_∞ represents the free-stream wind speed, θ denotes the vector of axial induction factors, and $\delta V_i(\theta)$ represents the velocity deficit experienced by turbine i. The velocity deficit is given by

$$\delta V_i(\theta) = 2\sqrt{\sum_{j<i} (\theta_j C[j,i])^2}. \tag{3}$$

Here, $C \in \mathbb{R}^{n\times n}$ is a matrix describing the wake interaction between turbines and depends on the wind farm layout, including the relative spacing between turbines and the overlap of their wakes. Each element of C is defined as

$$C[j,i] = \left(\frac{D_j}{D_j + 2k(x_i - x_j)}\right)^2 \frac{A_{j\to i}^{overlap}}{A_i}, \tag{4}$$

where D_j denotes the rotor diameter of turbine j, A_i is the rotor swept area of turbine i, and $A_{j\to i}^{overlap}$ represents the overlapping region between the wake generated by turbine j and the rotor area of turbine i.

Remark 2. *The proposed policy-search framework is model-free and does not depend on a specific wake model. The Park model is adopted as a simplified, widely used benchmark for evaluating the feasibility and performance of the proposed approach. The method relies solely on measured power signals and can be directly extended to more advanced wake models or real-world systems.*

2.3. Model of Produced Power

For each turbine i, the axial induction factor θ_i is treated as the primary control variable and takes values within the interval $\left[0, \frac{1}{2}\right]$. It is defined as

$$\theta_i = \frac{V_\infty - V_{rotor}}{V_\infty}, \tag{5}$$

where V_{rotor} represents the wind speed at the rotor plane.

The electrical power generated by turbine i can be expressed as

$$J_i(\theta_1, \theta_2, \ldots, \theta_i) = \frac{1}{2}\rho_{air} A_i C_p(\theta_i) V_i(\theta)^3, \tag{6}$$

where

- ρ_{air} is the air density;
- A_i denotes the rotor swept area of turbine i;
- $C_p(\theta_i)$ is the power coefficient given by

$$C_p(\theta_i) = 4\theta_i(1 - \theta_i)^2. \tag{7}$$

From (7), it can be verified that the power coefficient reaches its maximum value when $\theta_i = \frac{1}{3}$, corresponding to the well-known Betz limit [28].

2.4. Optimization Problem

Wind turbine power optimization problems can generally be classified into two main categories. The first focuses on optimizing the performance of a single turbine. In this case, the objective is to maintain the power coefficient C_p near its optimal value at $\theta_i = \frac{1}{3}$, which corresponds to the well-known Maximum Power Point Tracking (MPPT) problem (see, for example, ref. [31]).

The second category, which is the focus of this work, considers cooperative control among turbines within a wind farm. In this scenario, turbines coordinate their operating points to maximize the total power produced by the entire farm rather than maximizing individual turbine outputs.

Figure 2 illustrates this concept for a simple two-turbine case. It can be observed that operating the upstream turbine exactly at its individual optimal point ($\theta_1 = \frac{1}{3}$) maximizes its own power output. However, this configuration increases the wake effect on the downstream turbine and may reduce the overall farm power. Consequently, operating the upstream turbine slightly below its optimal point can increase the total generated power.

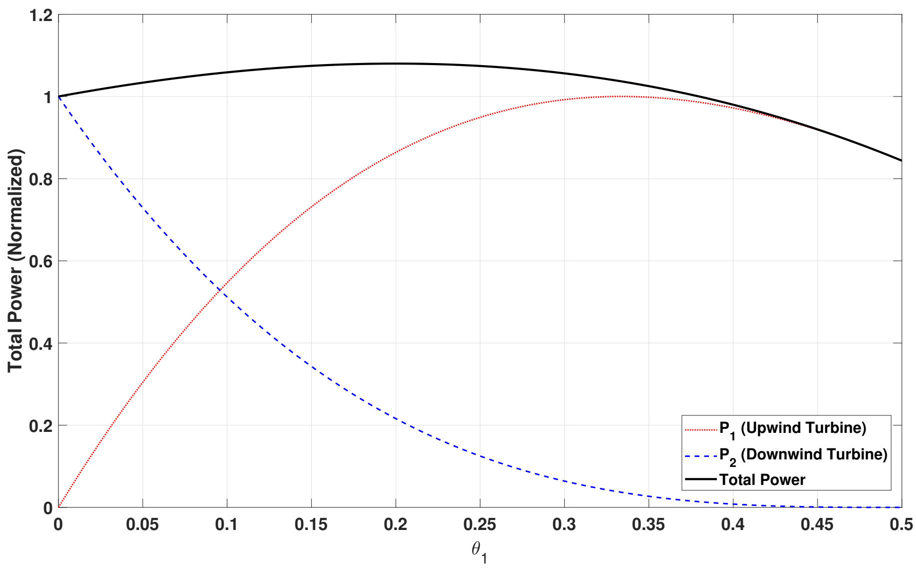

Figure 2. Wake interaction between two turbines. Operating the upstream turbine at $\theta_1 = \frac{1}{3}$ maximizes its individual power (solid red) but may reduce the total farm power (solid black).

In this work, a wind farm consisting of n turbines aligned along the wake direction is considered. Each turbine is associated with a controller and is modeled as an *agent*. Each agent i, for $i = 1, 2, \ldots, n$, is capable of measuring its local power generation $J_i(\theta(t))$. Furthermore, agents can exchange information through a communication network represented by a fixed, time-invariant, undirected, and strongly connected graph $G = (V_G, E_G)$. The vertex set $V_G = \{1, 2, \ldots, n\}$ represents the turbines, while E_G denotes the communication links between them.

To accurately capture the physical interactions among turbines, each agent is assumed to measure its generated power directly rather than estimating it. The overall optimization objective is to maximize the total power generated by the wind farm, defined as

$$P_{Total}(\theta_1, \theta_2, \ldots, \theta_n) = \sum_{i=1}^{n} J_i(\theta_1, \theta_2, \ldots, \theta_i). \tag{8}$$

Finally, each local cost function $J_i(\theta(t))$ is assumed to be concave and to possess a unique maximizer. Under these assumptions, the global optimization problem in (8) admits a unique finite solution. These assumptions are consistent with the properties derived from (6) and (7).

3. Proposed Control

This section presents the proposed control framework based on policy search reinforcement learning (PSRL). The objective of the proposed scheme is to determine optimal control inputs for wind turbines such that the total power production of the wind farm is maximized. The framework operates in a model-free fashion and relies only on measured power signals from the wind farm.

The general structure of the reinforcement learning scheme is illustrated in Figure 3. In our setup, the control input is generated as an agent action. The controller generates candidate control inputs using a stochastic policy. These inputs are applied to the wind farm (the environment) and the resulting power output is measured. The measured performance signal is then used to update the policy parameters in order to improve the operating point of the wind farm. In this way, the proposed framework can be viewed as a gradient-free optimization algorithm in which the gradient information is estimated through stochastic perturbations of the policy parameters.

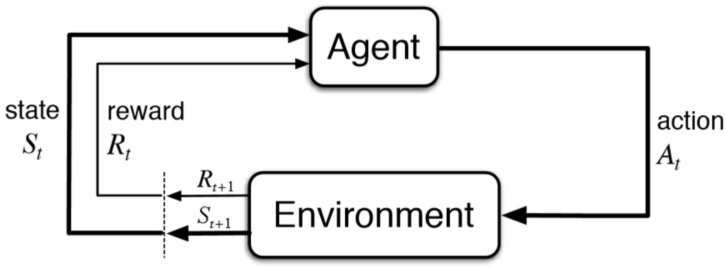

Figure 3. General reinforcement learning framework illustrating the interaction between the agent and the environment. At time step t, the agent observes the current state S_t and selects an action A_t. The environment then returns a reward R_{t+1} and transitions to the next state S_{t+1}.

3.1. Policy Search RL

Consider a wind farm consisting of n wind turbines. Let

$$\theta = [\theta_1, \theta_2, \ldots, \theta_n]^T \tag{9}$$

denote the vector of control variables associated with the turbines. In this work, θ_i represents the axial induction factor of turbine i.

The total power produced by the wind farm is denoted by

$$J(\theta) = P_{\text{tot}}(\theta) \tag{10}$$

where $J(\theta)$ is the performance index to be maximized.

Due to the aerodynamic wake interactions among turbines, the analytical expression of $J(\theta)$ is generally unknown or difficult to obtain. Consequently, conventional gradient-based optimization methods cannot be directly applied. To address this challenge, a stochastic policy is introduced to generate candidate control inputs.

Remark 3. *The proposed formulation does not rely on a full Markov Decision Process (MDP) structure. Instead, the problem is treated as a bandit optimization setting, where policy updates are driven by instantaneous reward measurements.*

3.1.1. Stochastic Policy

The control variables are generated according to the policy

$$\theta = \mu + \sigma\epsilon \tag{11}$$

where $\mu \in \mathbb{R}^n$ denotes the policy mean, $\sigma > 0$ is the exploration variance, and $\epsilon \sim \mathcal{N}(0, I)$ is a Gaussian random vector. The associated probability density is given by

$$\pi(\theta \mid \mu) = \mathcal{N}(\mu, \sigma^2 I). \tag{12}$$

The parameter μ represents the current estimate of the optimal operating point of the wind farm. At iteration k, a set of N candidate control vectors is generated as

$$\theta_i = \mu_k + \sigma\epsilon_i, \qquad i = 1, \ldots, N \tag{13}$$

where $\epsilon_i \sim \mathcal{N}(0, I)$. Each candidate control vector is applied to the wind farm and the corresponding reward is measured as

$$R_i = J(\theta_i). \tag{14}$$

It is worth mentioning that the proposed policy is intentionally parameterized using a low-dimensional Gaussian distribution rather than a neural-network-based function approximator. This choice is motivated by the compact action space and the availability of direct reward measurements, allowing for a lightweight and computationally efficient policy-search framework.

3.1.2. Policy-Gradient Estimate

The objective of the policy search method is to maximize the expected reward

$$\mathcal{J}(\mu) = \mathbb{E}_{\theta \sim \pi(\cdot|\mu)}[J(\theta)]. \tag{15}$$

Using the likelihood-ratio identity, the gradient of $\mathcal{J}(\mu)$ with respect to the policy mean can be written as

$$\nabla_\mu \mathcal{J}(\mu) = \mathbb{E}\big[J(\theta)\nabla_\mu \log \pi(\theta \mid \mu)\big]. \tag{16}$$

For the Gaussian policy in (11),

$$\nabla_\mu \log \pi(\theta \mid \mu) = \frac{\theta - \mu}{\sigma^2}. \tag{17}$$

Hence,

$$\nabla_\mu \mathcal{J}(\mu) = \frac{1}{\sigma^2}\mathbb{E}[J(\theta)(\theta - \mu)]. \tag{18}$$

In practice, the expectation in (18) is replaced by the sample average

$$\hat{\nabla}_\mu \mathcal{J}(\mu_k) = \frac{1}{N\sigma^2} \sum_{i=1}^{N} R_i(\theta_i - \mu_k). \tag{19}$$

The policy mean is then updated according to

$$\mu_{k+1} = \mu_k + \alpha_k \hat{\nabla}_\mu \mathcal{J}(\mu_k) \tag{20}$$

where $\alpha_k > 0$ is the learning rate.

Equation (20) shifts the policy mean toward control inputs that yield larger wind farm power. Therefore, the proposed PSRL controller searches for a cooperative operating point that maximizes the total wind farm power rather than improving the performance of turbines independently.

Remark 4. *The proposed controller is model-free. In particular, the update law in* (20) *does not require an explicit wake model, nor does it require the gradient of the wind farm power with respect to the induction factors.*

Remark 5. *The role of the stochastic perturbation in* (13) *is twofold. First, it allows the controller to explore the search space. Second, it provides the information needed to estimate the ascent direction through the correlation between perturbations and measured rewards.*

The implementation of the proposed PSRL controller is summarized in Algorithm 1. As shown in Figure 3, the wind farm plays the role of the environment, while the policy-update block generates new candidate control vectors and updates the policy mean according to the measured total power.

Algorithm 1 Policy Search Reinforcement Learning for Cooperative Wind Farm Control

Require: Initial policy mean μ_0, exploration variance σ, learning-rate sequence $\{\alpha_k\}$, batch size N

1: **for** $k = 0, 1, 2, \ldots$ **do**

2: **for** $i = 1, \ldots, N$ **do**

3: Sample perturbation $\epsilon_i \sim \mathcal{N}(0, I)$

4: Generate candidate control input:

$$\theta_i = \mu_k + \sigma \epsilon_i$$

5: Apply θ_i to the wind farm

6: Measure the resulting total power:

$$R_i = J(\theta_i)$$

7: **end for**

8: Compute gradient estimate:

$$\hat{\nabla}_\mu \mathcal{J}(\mu_k) = \frac{1}{N\sigma^2} \sum_{i=1}^{N} R_i(\theta_i - \mu_k)$$

9: Update the policy mean:

$$\mu_{k+1} = \mu_k + \alpha_k \hat{\nabla}_\mu \mathcal{J}(\mu_k)$$

10: **end for**

https://doi.org/10.3390/en19092160

3.2. Convergence Analysis

This subsection provides theoretical insight into the convergence properties of the proposed PSRL framework. Before that, the following assumptions are introduced.

Assumption 1. *The objective function $J(\theta)$ is continuously differentiable in the region of interest.*

Assumption 2. *The reward is bounded, that is, there exists a constant $M > 0$ such that*

$$|J(\theta)| \leq M \tag{21}$$

for all admissible θ.

Assumption 3. *The learning-rate sequence satisfies the Robbins–Monro conditions*

$$\sum_{k=0}^{\infty} \alpha_k = \infty, \qquad \sum_{k=0}^{\infty} \alpha_k^2 < \infty. \tag{22}$$

It is worth mentioning that, because the control input is sampled from a Gaussian distribution, the proposed algorithm does not optimize $J(\theta)$ directly. Instead, it optimizes the smoothed objective

$$J_\sigma(\mu) = \mathbb{E}_\epsilon[J(\mu + \sigma\epsilon)]. \tag{23}$$

The quantity $J_\sigma(\mu)$ represents the expected wind farm power when sampling around the current policy mean.

The gradient of $J_\sigma(\mu)$ can be obtained as

$$\nabla_\mu J_\sigma(\mu) = \frac{1}{\sigma^2}\mathbb{E}[J(\theta)(\theta - \mu)]. \tag{24}$$

Comparing (24) and (19), it follows that the gradient estimator used in the proposed controller is an unbiased estimator of the gradient of the smoothed objective.

Remark 6. *The smoothing induced by σ improves robustness to measurement noise and small local irregularities in the objective function. However, a large value of σ may also slow convergence by averaging the objective over a wider neighborhood.*

Theorem 1. *Suppose that Assumptions 1–3 hold. Consider the PSRL update law in (20) with gradient estimator (19). Then the sequence $\{\mu_k\}$ converges with probability one to the set of stationary points of the smoothed objective $J_\sigma(\mu)$, that is,*

$$\lim_{k\to\infty} \|\nabla_\mu J_\sigma(\mu_k)\| = 0 \quad \textit{with probability one.} \tag{25}$$

Proof. From (19) and (24), the update law in (20) can be written as a stochastic approximation recursion of the form

$$\mu_{k+1} = \mu_k + \alpha_k\big(\nabla_\mu J_\sigma(\mu_k) + \xi_k\big) \tag{26}$$

where ξ_k is a zero-mean estimation error satisfying

$$\mathbb{E}[\xi_k \mid \mu_k] = 0. \tag{27}$$

Under Assumptions 1 and 2, the gradient of the smoothed objective is well defined and bounded on the region of interest. Assumption 3 guarantees that the learning-rate sequence satisfies the standard stochastic approximation conditions. Therefore, by classical Robbins–

Monro stochastic approximation results, the parameter sequence μ_k converges almost surely to the set of stationary points of $J_\sigma(\mu)$. $\square$

Remark 7. *Theorem 1 establishes convergence to a stationary point of the smoothed objective J_σ, not necessarily to the global maximizer of the original objective J. In practice, this is a standard feature of policy-search methods and gradient-free optimization schemes.*

Remark 8. *A decreasing exploration variance can be used in practice to improve the approximation between the stationary points of J_σ and those of the original objective J. Nevertheless, for fixed σ, the convergence result in Theorem 1 remains valid for the smoothed objective.*

The above analysis shows that the proposed control law has a clear optimization interpretation. It performs stochastic gradient ascent on a smoothed approximation of the wind farm power function. As a result, the proposed controller provides a simple and cooperative model-free optimization framework for wind farm control.

4. Simulation Results

This section evaluates the performance of the proposed policy search reinforcement learning (PSRL) controller on the wind farm optimization problem. Then, it compares the proposed strategy with a sliding-mode extremum-seeking control (SM-ESC) method and a classical policy-gradient method (REINFORCE). The objective is to maximize the total power generated by the wind farm by adjusting the axial induction factors of the turbines.

4.1. Simulation Setup

A wind farm consisting of three aligned wind turbines is considered, as shown in Figure 4. The turbine diameter is selected as $D = 80$ m, and the free-stream wind speed is set to $V_\infty = 8$ m/s. The air density is chosen as $\rho = 1.225$ kg/m^3. The turbine spacing is defined as $q_1 = 0$, $q_2 = 400$, and $q_3 = 800$, while the wake expansion coefficient is selected as $k = 0.075$.

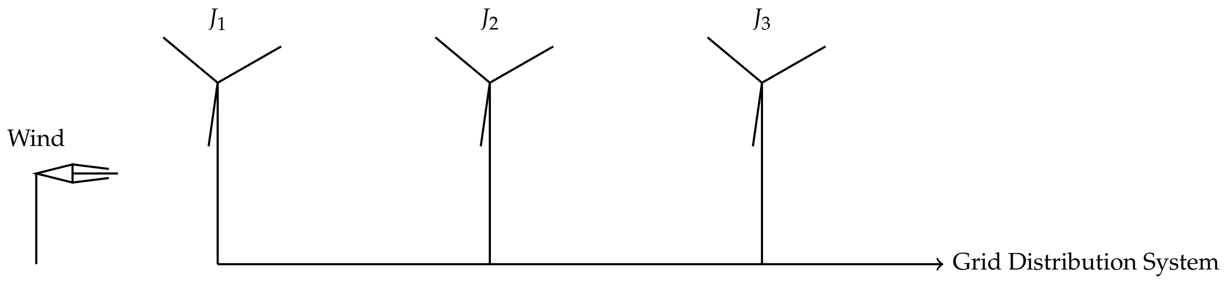

Figure 4. Conceptual configuration of the considered wind farm. Three turbines (J_1, J_2, and J_3) are arranged along the incoming wind direction and connected through a common power collection line that transfers the generated electricity to the grid distribution system.

The total power production of the wind farm is described by

$$\begin{aligned} P_{\text{tot}}(\theta_1, \theta_2, \theta_3) = \frac{1}{2}\rho A \Big(& 4\theta_1(1-\theta_1)^2 V_\infty^3 \\ & + 4\theta_2(1-\theta_2)^2 \big(V_\infty(1 - 2\theta_1 c_{12})\big)^3 \\ & + 4\theta_3(1-\theta_3)^2 \big(V_\infty(1 - 2\sqrt{(\theta_1 c_{13})^2 + (\theta_2 c_{23})^2})\big)^3 \Big) \end{aligned} \tag{28}$$

where θ_i denotes the axial induction factor of turbine i, and the coefficients c_{12}, c_{13}, and c_{23} model the wake interaction between turbines.

The admissible range of the control variables is chosen as

$$0 \le \theta_i \le 0.5, \qquad i = 1, 2, 3. \tag{29}$$

For the proposed PSRL method, the initial policy mean is selected as

$$\mu_0 = [0.4,\ 0.4,\ 0.4]^T. \tag{30}$$

At each iteration, $N = 40$ candidate solutions are sampled from the Gaussian policy. The learning rate is set to $\alpha = 10^{-3}$, and the exploration variance is chosen as $\sigma = 0.01$. In addition, a normalized advantage function is used to reduce the sensitivity of the update law to the absolute scale of the reward signal. Gradient clipping is also employed in order to improve numerical robustness.

For comparison, a sliding-mode extremum-seeking controller is implemented using the total power output of the wind farm as the performance signal. The SM-ESC scheme updates the induction factors through discontinuous control laws derived from the sliding surfaces defined by the difference between the measured power and a monotonic reference signal.

4.2. Performance of the Proposed PSRL Scheme

The first set of simulations illustrates the evolution of the policy mean under the proposed policy search reinforcement learning (PSRL) algorithm. Starting from the initial condition $\mu_0 = [0.4, 0.4, 0.4]^T$, the policy mean gradually moves toward an operating point that yields larger wind farm power. The evolution of the policy parameters μ_1, μ_2, and μ_3 is shown in Figure 5. It can be observed that the parameters converge toward the neighborhood of the cooperative optimal operating point. In particular, the parameters settle close to the reference values $(0.232, 0.208, 0.333)$ obtained from the optimal cooperative solution. This behavior confirms that the proposed policy-search update law successfully drives the operating point toward a cooperative configuration of the wind farm.

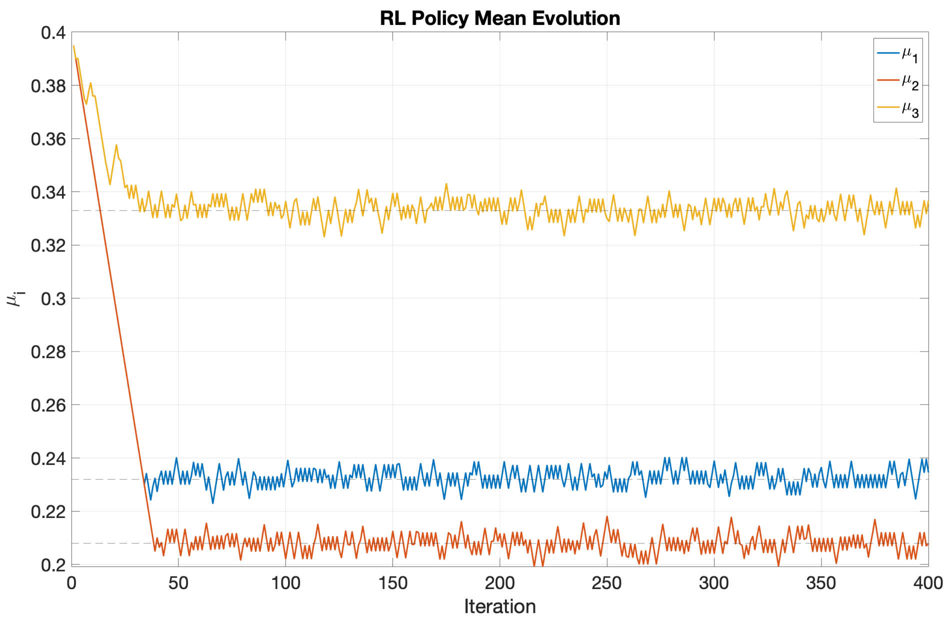

Figure 5. Evolution of the policy mean parameters μ_1, μ_2, and μ_3 under the proposed PSRL algorithm. The dashed lines indicate the reference optimal cooperative operating point.

The evolution of the total reward (i.e., the total power generated by the wind farm) is shown in Figure 6. The reward evaluated at the policy mean gradually increases as the learning process progresses. This demonstrates that the PSRL algorithm effectively improves the operating point of the wind farm over successive iterations.

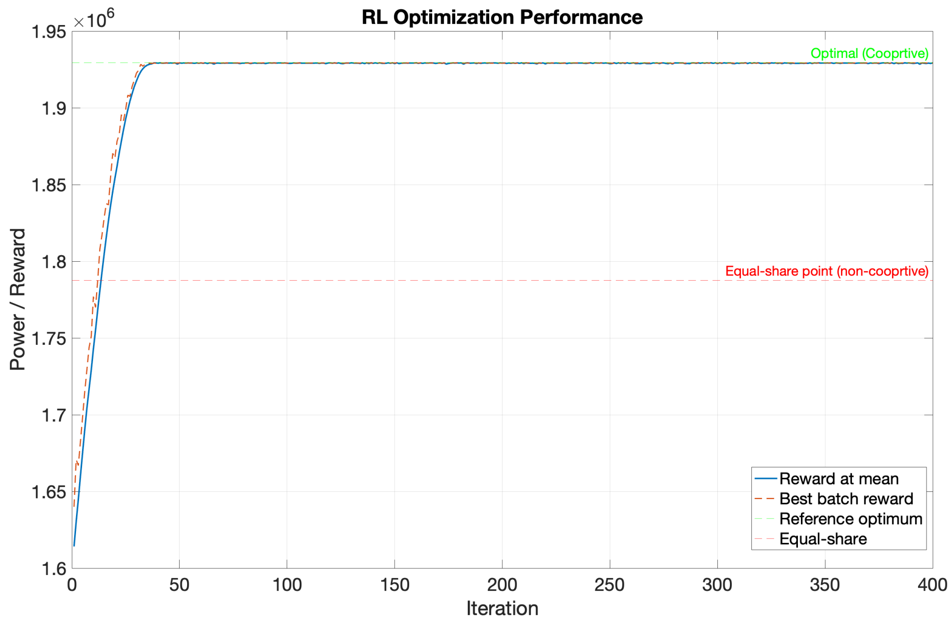

Figure 6. Reward evolution during learning. The reward evaluated at the policy mean converges toward the cooperative optimal value. The dashed curves indicate the reference optimal solution and the equal-share operating point.

Furthermore, the results show that the best sampled reward in each batch remains close to the optimal operating point. This observation indicates that the proposed framework is capable of exploring the search space efficiently while maintaining stable convergence properties. In addition, the learning process converges rapidly within a relatively small number of iterations.

It is also worth noting that the obtained operating point significantly improves the total power generation compared with the equal-share configuration. In the equal-share case, all turbines operate with identical induction factors, which neglects the aerodynamic wake interactions among turbines. In contrast, the proposed PSRL approach naturally accounts for these interactions through direct optimization of the total wind farm power. As a result, the algorithm converges toward a cooperative operating point that maximizes the overall energy capture of the wind farm. To further evaluate the adaptability of the proposed PSRL framework, additional simulations were conducted under time-varying wake interaction conditions. In particular, the overlap between wind turbine wakes was varied in a stepwise manner to emulate different aerodynamic interaction regimes within the wind farm.

4.3. Adaptation Under Time-Varying Wake Interactions

Figure 7 shows the evolution of the total power generated by the wind farm as the overlap level changes during the learning process. The overlap ratio is modified at predefined iterations, which results in a corresponding shift in the optimal operating point. The results demonstrate that the proposed PSRL controller is able to rapidly adapt to these changes and track the new optimal power levels associated with each overlap configuration.

Figure 8 illustrates the evolution of the policy parameters together with the overlap stages applied during the experiment. It can be observed that the policy means smoothly adjust to the new operating conditions whenever the overlap ratio changes. This behavior confirms that the reinforcement learning framework is capable of continuously adapting the turbine induction factors to maintain near-optimal wind farm performance despite variations in the aerodynamic coupling among turbines. Overall, these results highlight the robustness and adaptability of the proposed PSRL-based optimization strategy in dynamic wind farm environments.

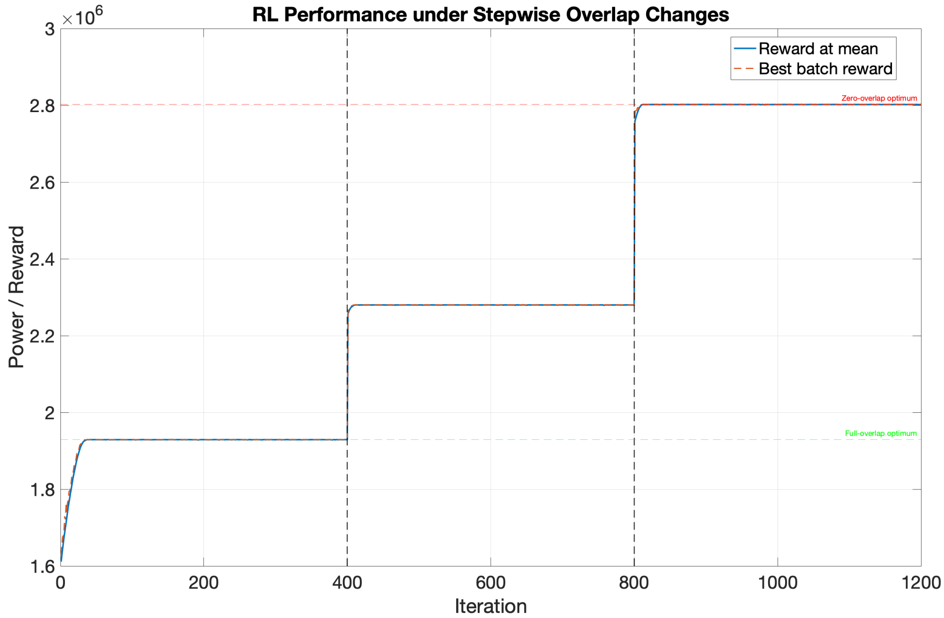

Figure 7. Performance of the proposed PSRL algorithm under stepwise changes in turbine wake overlap. The total wind farm power adapts to the new optimal operating point as the aerodynamic interaction between turbines changes.

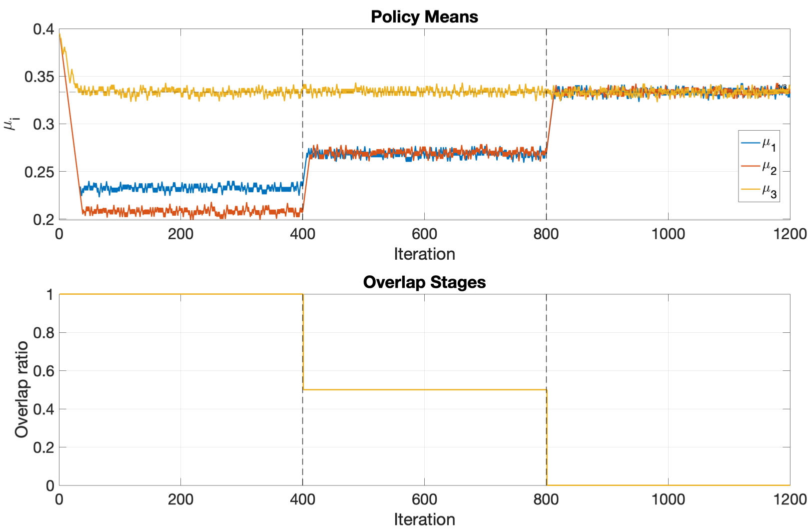

Figure 8. Evolution of the policy mean parameters and the corresponding overlap stages. The controller adapts the induction factors as the aerodynamic coupling between turbines varies over time.

4.4. Parameter Sensitivity Analysis

To evaluate the robustness of the proposed PSRL framework, a sensitivity analysis is conducted with respect to the exploration variance σ and the learning rate α. The impact of key parameters on the convergence behavior of the proposed PSRL algorithm is illustrated in Figure 9. As shown in Figure 9a, the exploration variance σ influences the trade-off between exploration and stability. In contrast, Figure 9b demonstrates the effect of the learning rate α on convergence speed and stability. The results show that increasing σ

improves exploration in the early stages, helping the algorithm avoid suboptimal operating points. However, large values may introduce oscillations and slow convergence, while smaller values lead to smoother but more conservative updates.

The learning rate α primarily affects the convergence speed. Larger values accelerate learning but may cause instability near the optimum, whereas smaller values ensure stable but slower convergence. Overall, the results demonstrate that the proposed method maintains reliable performance over a range of parameter settings, with moderate values of σ and α providing a good balance between exploration and stability.

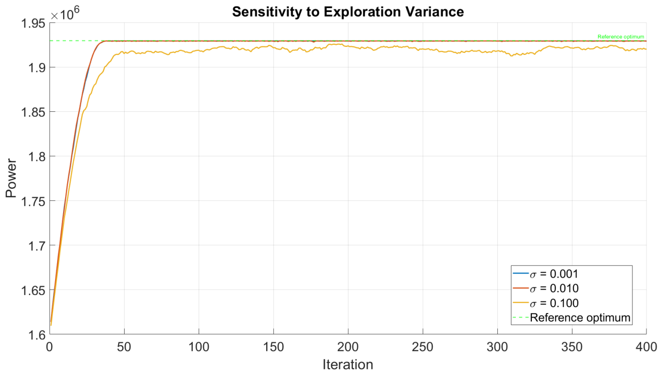

(**a**)

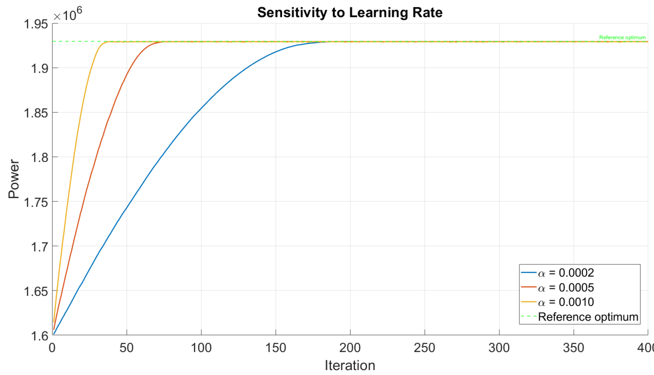

(**b**)

Figure 9. Parameter sensitivity analysis of the proposed PSRL framework. (**a**) Effect of exploration variance σ on convergence. (**b**) Effect of learning rate α on convergence.

4.5. Comparative Evaluation with Gradient-Free and Policy Gradient Methods

To further assess the effectiveness of the proposed approach, the PSRL controller is compared with the sliding-mode extremum-seeking control (SM-ESC), which represents the closest methodological baseline due to its gradient-free and perturbation-based structure. In addition, a comparison with the REINFORCE algorithm is included, as it belongs to the same family of policy-gradient reinforcement learning methods. The comparison is performed using the same wind farm model and identical initial operating conditions. Figure 10 illustrates the evolution of the wind farm power obtained using the PSRL, REINFORCE, and SM-ESC approaches. It can be observed that all methods successfully improve the wind farm power starting from the initial operating point. In particular, the PSRL algorithm gradually increases the total power through stochastic policy updates and converges to the neighborhood of the optimal cooperative solution. The REINFORCE method exhibits a similar convergence behavior, albeit at a slower rate. Furthermore, the SM-ESC algorithm also converges to the optimal solution, although the convergence process is driven by discontinuous switching signals inherent to sliding-mode extremum seeking.

The evolution of the induction factors for the three methods is shown in Figure 11. The PSRL controller produces a smoother parameter evolution due to the stochastic gradient-based policy updates. In contrast, the SM-ESC algorithm exhibits sharper transitions caused by the discontinuous switching structure of the sliding-mode update law. Nevertheless, both approaches converge toward operating points close to the optimal cooperative solution.

The final operating points obtained using the different approaches are summarized in Figure 12a. The figure shows that all three methods—PSRL, REINFORCE, and SM-ESC—converge to operating points that are close to the reference optimal solution. This result confirms that the proposed reinforcement learning framework is capable of identifying cooperative operating conditions for the wind farm.

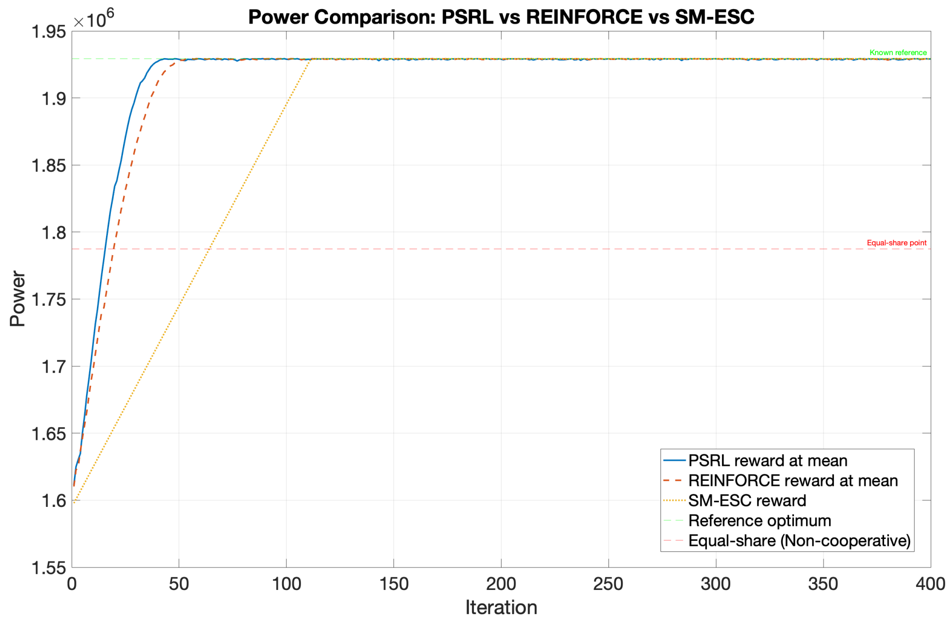

Figure 10. Evolution of wind farm power using PSRL, REINFORCE, and SM-ESC. Both methods converge toward the cooperative optimal solution.

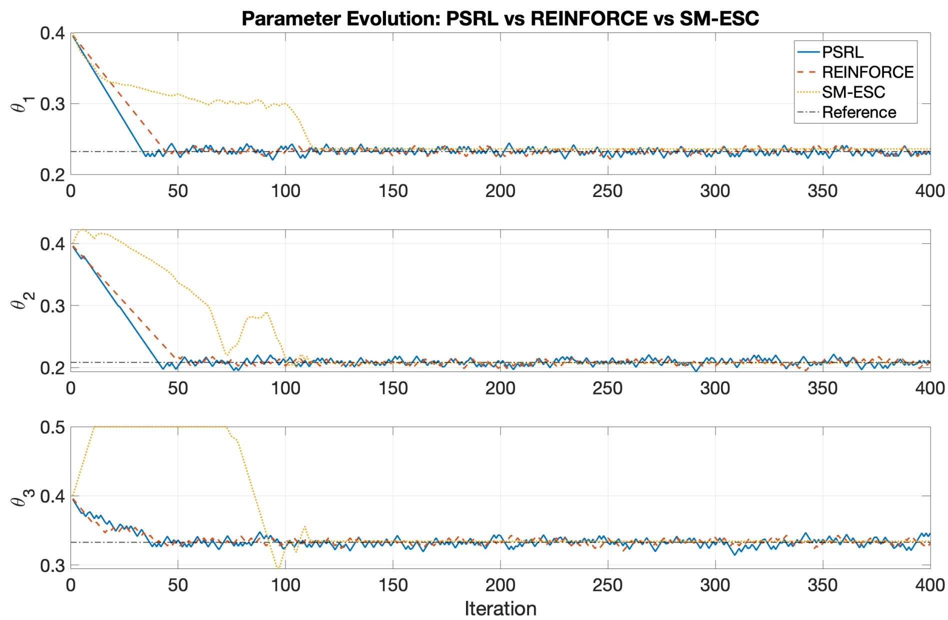

Figure 11. Evolution of the induction factors using PSRL, REINFORCE, and SM-ESC.

Finally, the steady-state power achieved by the different approaches is presented in Figure 12b. The results show that both PSRL and SM-ESC achieve power levels that are very close to the reference optimal solution and significantly higher than the equal-share operating point. This confirms that cooperative optimization of the induction factors can substantially improve the overall energy capture of the wind farm.

Overall, the results demonstrate that the proposed PSRL framework achieves performance comparable to sliding-mode extremum-seeking control (SM-ESC) while providing a smoother and more stable optimization process. In addition, the RL-based formulation offers a flexible framework that can be extended to more complex wind farm control problems and other multivariable optimization settings. The simulation results confirm that the proposed PSRL controller converges to a cooperative operating point that is close to the reference optimal solution while preserving a simple model-free structure.

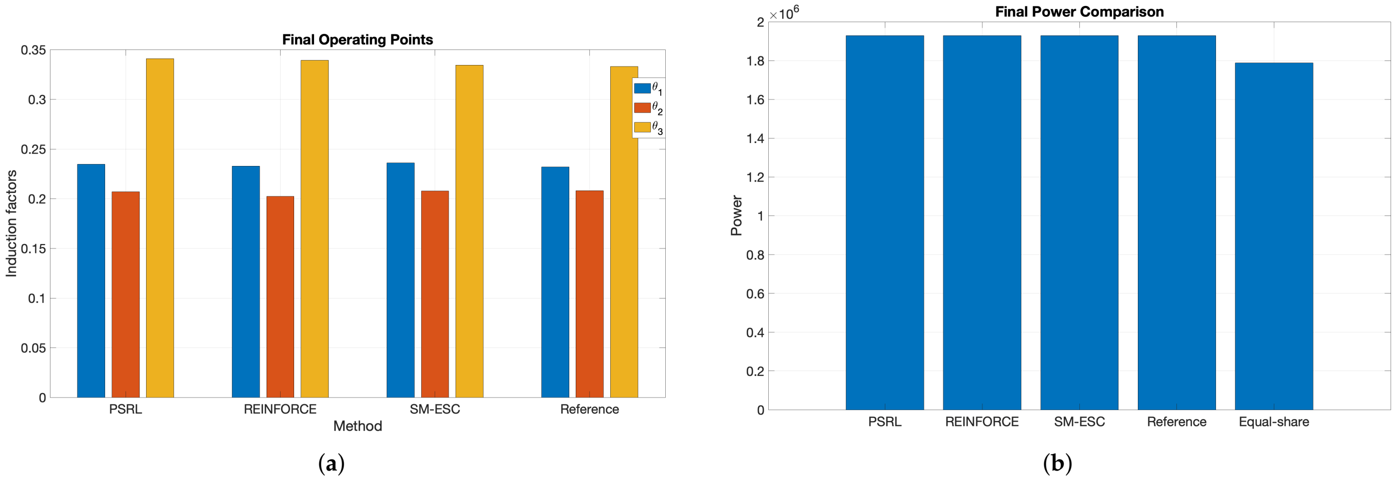

Figure 12. Comparison of the final operating points obtained by PSRL, REINFORCE and SM-ESC. (**a**) Final induction factors obtained using PSRL, REINFORCE, SM-ESC, and the reference optimal solution. (**b**) Final wind farm power achieved by PSRL, REINFORCE, SM-ESC, and the equal-share operating point.

In contrast, the REINFORCE baseline, although capable of reaching similar operating regions, exhibits slower convergence and increased sensitivity to reward scaling, requiring gradient clipping to ensure stable updates under the same conditions. This highlights the advantage of the proposed PSRL framework, where normalized advantage-based updates improve numerical stability and convergence efficiency.

In summary, the presented results highlight two important observations. First, the proposed PSRL framework is capable of solving the wind farm optimization problem without requiring an explicit model of the wake interaction between turbines. Second, the comparison with SM-ESC and REINFORCE shows that the proposed approach achieves competitive performance while maintaining a simple implementation structure and improved convergence behavior. These characteristics make the proposed framework particularly attractive for applications where reliable gradient information is unavailable or difficult to obtain. It is important to note that the presented results are obtained using the Park wake model, which provides a simplified steady-state representation of wind turbine interactions. While this model is widely used in the literature as a benchmark for validating control and optimization strategies, it does not capture all real-world effects such as wake meandering, turbulence, and actuator dynamics. Nevertheless, the proposed PSRL framework is inherently model-free and relies solely on measured power signals rather than model predictions. This reduces the risk of overfitting to specific simulator assumptions and improves robustness to modeling inaccuracies. In addition, the stochastic perturbation mechanism introduces an implicit smoothing effect, which enhances robustness to noise and small disturbances.

4.6. Practical Considerations

For practical implementation, additional considerations such as actuator rate limits, measurement noise filtering, and safety constraints can be incorporated into the policy update. In particular, bounded and rate-limited updates can be enforced to ensure smooth control actions and physical feasibility of the induction factors. Additional operational objectives, such as structural load mitigation, actuator fatigue reduction, and grid-support requirements, can also be incorporated into the optimization framework through constrained or multi-objective formulations.

In addition, while the computational complexity of the proposed policy-search framework scales linearly with the number of control variables, increasing the number of turbines leads to a higher-dimensional search space, which may impact convergence speed and sample efficiency. Additionally, the convergence behavior may be influenced by the strength of aerodynamic coupling among turbines, where strong or ill-conditioned interactions can further slow the learning process.

In the current formulation, the controller is centralized, which may limit scalability for large wind farms. Future work will investigate distributed and hierarchical extensions of the proposed approach to improve scalability. Furthermore, validation using high-fidelity dynamic wake models and experimental data, as well as incorporating actuator dynamics and communication delays, will be considered to better assess real-world applicability.

5. Conclusions

This paper presented a policy-search reinforcement learning framework for cooperative wind farm optimization. The proposed approach develops a model-free optimization strategy for maximizing the total power captured by a wind turbine farm through the adjustment of turbine induction factors. In contrast to conventional control strategies that optimize individual turbine performance in a greedy manner, the proposed scheme directly maximizes the overall power generated by the wind farm, thereby accounting for aerodynamic wake interactions among turbines.

The proposed reinforcement learning formulation treats the wind farm optimization problem as a multivariable policy-search problem. By employing stochastic perturbations and reward-based updates, the controller gradually adjusts the policy parameters toward operating points that yield higher total power. The convergence behavior of the learning process has been analyzed, and the resulting algorithm provides a simple yet effective gradient-free mechanism suitable for online optimization applications.

The performance of the proposed PSRL scheme has been verified through numerical simulations using a wind farm power model. The results demonstrate that the proposed algorithm successfully converges toward the cooperative optimal operating point and significantly improves the power capture compared with the equal-share operating condition. Furthermore, a comparison with the sliding-mode extremum-seeking control method shows that the proposed RL-based approach achieves comparable steady-state performance while providing a smoother optimization trajectory. Moreover, comparison with a standard policy-gradient baseline further confirms improved convergence behavior under identical sampling conditions.

Overall, the presented results highlight the potential of reinforcement learning as a flexible framework for multivariable cooperative optimization in renewable energy systems. Future work may focus on extending the proposed framework to dynamic wind conditions, larger wind farms, and more complex control variables such as yaw and pitch control. In addition, further investigation may consider the influence of measurement noise, time delays, and uncertain wake interactions on the learning performance.

Funding: This research was funded by King Saud University through the Ongoing Research Funding program, (ORF-2026-270), King Saud University, Riyadh, Saudi Arabia.

Data Availability Statement: The data are contained within the article.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Bin Salamah, Y.; Ozguner, U. Sliding Mode Multivariable Extremum Seeking Control with Application to Wind Farm Power Optimization. In *2018 Annual American Control Conference (ACC)*; IEEE: New York, NY, USA, 2018; pp. 5321–5326.
2. Dong, H.; Xie, J.; Zhao, X. Wind farm control technologies: From classical control to reinforcement learning. *Prog. Energy* **2022**, *4*, 032006. [CrossRef]

3. Hasager, C.B.; Rasmussen, L.; Peña, A.; Jensen, L.E.; Réthoré, P.E. Wind farm wake: The Horns Rev photo case. *Energies* **2013**, *6*, 696–716. [CrossRef]
4. Abkar, M.; Zehtabiyan-Rezaie, N.; Iosifidis, A. Reinforcement learning for wind-farm flow control: Current state and future actions. *Theor. Appl. Mech. Lett.* **2023**, *13*, 100475. [CrossRef]
5. Zhang, Z.; Li, J.; Lei, Z.; Zhu, Q.; Cheng, J.; Gao, S. Reinforcement learning-based particle swarm optimization for wind farm layout problems. *Energy* **2024**, *313*, 134050. [CrossRef]
6. Dong, G.; Qin, J.; Wu, C.; Xu, C.; Yang, X. Reinforcement learning-enhanced genetic algorithm for wind farm layout optimization. *Renew. Energy* **2025**, *259*, 125093. [CrossRef]
7. Dar, Z.; Kar, K.; Sahni, O.; Chow, J.H. Windfarm power optimization using yaw angle control. *IEEE Trans. Sustain. Energy* **2016**, *8*, 104–116. [CrossRef]
8. Lucas Frutuoso, T.R.; Castro, R.; Pereira, R.B.S.; Moutinho, A. Advancements in wind farm control: Modelling and multi-objective optimization through yaw-based wake steering. *Energies* **2025**, *18*, 2247. [CrossRef]
9. Routray, A.; Sivakumar, N.; Hur, S.H.; Bang, D.J. A comparative study of optimal individual pitch control methods. *Sustainability* **2023**, *15*, 10933. [CrossRef]
10. Bin Salamah, Y.; Ozguner, U. Distributed extremum-seeking for wind farm power maximization using sliding mode control. *Energies* **2021**, *14*, 828. [CrossRef]
11. Gebraad, P.M.; van Wingerden, J.W. Maximum power-point tracking control for wind farms. *Wind Energy* **2015**, *18*, 429–447. [CrossRef]
12. Heer, F.; Esfahani, P.M.; Kamgarpour, M.; Lygeros, J. Model based power optimisation of wind farms. In *2014 European Control Conference (ECC)*; IEEE: New York, NY, USA, 2014; pp. 1145–1150.
13. Ahmad, M.A.; Hao, M.R.; Ismail, R.M.T.R.; Nasir, A.N.K. Model-free wind farm control based on random search. In *Proceedings of the 2016 IEEE International Conference on Automatic Control and Intelligent Systems (I2CACIS), Selangor, Malaysia, 22–22 October 2016*; IEEE: New York, NY, USA, 2016; pp. 131–134.
14. Bin Salamah, Y.; Fiorentini, L.; Ozguner, U. Cooperative Sliding Mode Extremum Seeking Control for Distributed Optimization. In *Proceedings of the 2018 IEEE 57th Annual Conference on Decision and Control (CDC)*; IEEE: New York, NY, USA, 2018; pp. 1281–1286.
15. Xu, J.M.; Soh, Y.C. Distributed extremum seeking control of networked large-scale systems under constraints. In *Proceedings of the 52nd Annual Conference on Decision and Control*; IEEE: New York, NY, USA, 2013; pp. 2187–2192.
16. Menon, A.; Baras, J.S. Collaborative extremum seeking for welfare optimization. In *Proceedings of the Decision and Control (CDC), 2014 IEEE 53rd Annual Conference, Los Angeles, CA, USA, 15–17 December 2014*; IEEE: New York, NY, USA, 2014; pp. 346–351.
17. Ye, M.; Hu, G. Distributed extremum seeking for constrained networked optimization and its application to energy consumption control in smart grid. *IEEE Trans. Control Syst. Technol.* **2016**, *24*, 2048–2058. [CrossRef]
18. Ebegbulem, J.; Guay, M. Distributed Extremum Seeking Control for Wind Farm Power Maximization. *IFAC-PapersOnLine* **2017**, *50*, 147–152. [CrossRef]
19. Xu, Z.; Chu, B.; Geng, H.; Nian, X. Distributed power optimization of large wind farms using ADMM for real-time control. *IEEE Trans. Power Syst.* **2022**, *37*, 4832–4845. [CrossRef]
20. Marden, J.R.; Ruben, S.D.; Pao, L.Y. A model-free approach to wind farm control using game theoretic methods. *IEEE Trans. Control Syst. Technol.* **2013**, *21*, 1207–1214. [CrossRef]
21. Park, J.; Kwon, S.; Law, K.H. Wind farm power maximization based on a cooperative static game approach. In *Proceedings of the Active and Passive Smart Structures and Integrated Systems 2013*; International Society for Optics and Photonics: Bellingham, WA, USA, 2013; Volume 8688, p. 86880R. [CrossRef]
22. Stanfel, P.; Johnson, K.; Bay, C.J.; King, J. A distributed reinforcement learning yaw control approach for wind farm energy capture maximization. In *Proceedings of the 2020 American Control Conference (ACC)*; IEEE: New York, NY, USA, 2020; pp. 4065–4070.
23. Dong, H.; Zhao, X. Wind-farm power tracking via preview-based robust reinforcement learning. *IEEE Trans. Ind. Inform.* **2021**, *18*, 1706–1715. [CrossRef]
24. Soler, D.; Mariño, O.; Huergo, D.; de Frutos, M.; Ferrer, E. Reinforcement learning to maximize wind turbine energy generation. *Expert Syst. Appl.* **2024**, *249*, 123502. [CrossRef]
25. Zhao, H.; Zhao, J.; Qiu, J.; Liang, G.; Dong, Z.Y. Cooperative wind farm control with deep reinforcement learning and knowledge-assisted learning. *IEEE Trans. Ind. Inform.* **2020**, *16*, 6912–6921. [CrossRef]
26. Luo, F.; Wang, S.; Lv, Y.; Mu, R.; Fo, J.; Zhang, T.; Xu, J.; Wang, C. Domain knowledge-enhanced graph reinforcement learning method for Volt/Var control in distribution networks. *Appl. Energy* **2025**, *398*, 126409. [CrossRef]
27. Xiao, Y.; Li, Y.; Rotea, M.A. CART3 field tests for wind turbine region-2 operation with extremum seeking controllers. *IEEE Trans. Control Syst. Technol.* **2018**, *27*, 1744–1752. [CrossRef]
28. Vali, M.; Petrović, V.; Boersma, S.; van Wingerden, J.W.; Pao, L.Y.; Kühn, M. Adjoint-based model predictive control for optimal energy extraction in waked wind farms. *Control Eng. Pract.* **2019**, *84*, 48–62. [CrossRef]

29. Johnson, K.E.; Pao, L.Y.; Balas, M.J.; Fingersh, L.J. Control of variable-speed wind turbines: Standard and adaptive techniques for maximizing energy capture. *IEEE Control Syst. Mag.* **2006**, *26*, 70–81.
30. Katic, I.; Højstrup, J.; Jensen, N.O. A simple model for cluster efficiency. In *Proceedings of the European Wind Energy Association Conference and Exhibition*; A. Raguzzi: Rome, Italy, 1987.
31. Shen, D.; Khayyer, P.; Izadian, A. Sliding mode extremum seeking control for maximum power point tracking in wind system. In *Proceedings of the 2016 IEEE Power and Energy Conference at Illinois (PECI)*; IEEE: New York, NY, USA, 2016; pp. 1–6.