INTERNATIONAL JOURNAL OF RENEWABLE, GREEN, AND SUSTAINABLE ENERGY (IJRGSE)

eISSN: 3087-4300

Volume. 03, Issue. 04, pp. 36-49, April 2026"

Beyond Wake Loss Minimization: A Comprehensive Research Synthesis on Gradient-Based, Heuristic, and Robust Wind Farm Layout Optimization Under Real-World Constraints

Sofia Lindberg
Division of Energy Sciences, Lund University, Sweden

Article received: 17/02/2026, Article Accepted: 15/03/2026, Article Published: 17/04/2026

© 2026 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the Creative Commons Attribution License 4.0 (CC-BY), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

Background: Wind farm layout optimization has evolved from a relatively idealized geometric placement problem into a multidisciplinary engineering challenge involving wake physics, terrain effects, atmospheric uncertainty, turbine heterogeneity, control co-design, computational tractability, and regulatory spatial constraints. Foundational work on wake behavior, turbulence development, and offshore flow interactions established that turbine spacing decisions materially influence annual energy production and the economic performance of wind plants (Crespo & Hernández, 1996; Barthelmie et al., 2006; Barthelmie et al., 2009; Barthelmie et al., 2010). Subsequent methodological advances introduced increasingly sophisticated optimization frameworks, including random search, particle swarm optimization, sequential quadratic programming, adjoint methods, and modern gradient-based large-scale design techniques (Wan et al., 2012; Feng & Shen, 2015; Gill et al., 2005; King et al., 2017; Wu et al., 2020).

Objective: This article develops a comprehensive research synthesis of wind farm layout optimization by integrating the aerodynamic, algorithmic, and systems-engineering insights contained in the provided literature. The study aims to clarify how wake model fidelity, optimizer choice, uncertainty treatment, and operational co-design jointly determine optimization quality.

Methodology: A structured interpretive synthesis was conducted across the referenced literature. The analysis organized the field into six interacting dimensions: wake representation, optimization architecture, constraint handling, uncertainty treatment, control-layout integration, and computational scalability. The article compares methodological assumptions, traces conceptual convergence across studies, and identifies recurring sources of mismatch between theoretical optima and deployable designs.

Results: The synthesis shows that no single optimizer is universally superior; performance depends strongly on objective smoothness, wake model differentiability, number of turbines, constraint complexity, and the need for robustness under variable wind climates (Baker et al., 2019; Croonenbroeck & Hennecke, 2021). Gradient-based methods exhibit particular strength when paired with differentiable wake models and continuation strategies, especially for large-scale problems (Guirguis et al., 2016; Stanley & Ning, 2019; Thomas et al., 2022; Quick et al., 2023; Valotta Rodrigues et al., 2024). However, heuristic methods remain useful in nonconvex or discrete design settings involving restricted zones or heterogeneous turbine classes (Pookpunt & Ongsakul, 2016; Hou et al., 2016; Feng & Shen, 2017a). The field is moving from static energy maximization toward robust, control-aware, and constraint-rich optimization.

Conclusion: The most promising future direction is an integrated framework combining physics-aware differentiable wake models, scalable gradient methods, explicit spatial constraints, stochastic uncertainty treatment, and joint optimization of layout with operational control. Such an approach better aligns academic optimization with real wind plant design practice.

Keywords: wind farm layout optimization, wake modeling, gradient-based optimization, robust design, offshore wind farms, yaw control, constrained micrositing.

INTRODUCTION

Wind energy development has reached a stage at which marginal improvements in design quality can translate into significant gains in annual energy production, reduced levelized cost of energy, and improved utilization of scarce developable land or offshore lease areas. Within this broader engineering and economic context, wind farm layout optimization has emerged as one of the central design problems in modern renewable energy systems. The importance of layout stems from the fact that a wind farm is not merely a collection of independent turbines. Each turbine modifies the flow field experienced by downstream machines, reducing wind speed, increasing turbulence, and altering the structural and operational environment of the plant as a whole (Crespo & Hernández, 1996; Barthelmie et al., 2009; Porté-Agel et al., 2020). As a result, the total energy output of a wind farm depends not only on site wind resources and turbine ratings, but also on the spatial arrangement of turbines, the characteristics of wake recovery, the variability of wind direction and intensity, and the control strategies implemented during operation (Fleming et al., 2016; Gebraad et al., 2017).

Historically, wind farm micrositing was often approached through simplified spacing heuristics, developer experience, or rule-based placement strategies that privileged rough inter-turbine distance criteria. Such approaches were understandable in earlier stages of industry development, when computational limitations, limited measurement availability, and less mature wake models constrained design sophistication. However, the research literature now demonstrates clearly that wake interactions are too consequential and too site-specific to be handled adequately through simplistic spacing rules alone. Experimental and field studies have shown that wakes in large wind farms can persist over long distances and can substantially reduce downstream power production, especially under aligned wind directions and stable atmospheric conditions (Barthelmie et al., 2006; Barthelmie et al., 2010). Offshore conditions, despite lower surface roughness and often stronger mean winds, do not eliminate the importance of wake losses; instead, they may accentuate long-range wake behavior due to reduced ambient turbulence and relatively coherent inflow structures (Barthelmie et al., 2009; Barthelmie et al., 2010).

The scientific basis of wind farm optimization rests upon two intertwined domains: the physics of wind turbine wakes and the mathematics of optimization under constraints. On the physics side, research has progressively moved from early empirical and engineering wake models toward more refined analytical, experimental, and integral formulations. Crespo and Hernández (1996) clarified how turbine wakes generate turbulence and how wake properties evolve downstream. Barthelmie and colleagues provided crucial measurement-based and simulation-based validation of offshore wake behavior across operational wind farms, strengthening the case for model-based micrositing while simultaneously revealing the limitations of simplified formulations (Barthelmie et al., 2006; Barthelmie et al., 2009; Barthelmie et al., 2010). Later, analytical frameworks such as those discussed by Niayifar and Porté-Agel (2016), and the broader synthesis by Porté-Agel et al. (2020), further advanced the capability to represent wind-farm-scale flow interactions. Bastankhah and Porté-Agel (2016) expanded understanding of wake deflection and yawed conditions, an especially important development because optimization can no longer be understood only as geometric placement; operational steering increasingly enters the design space. More recently, integral and analytical approaches such as FLOWERS and FLOWERS AEP have been proposed to reduce computational cost while retaining sufficient fidelity for layout applications (LoCascio et al., 2022; LoCascio et al., 2024).

On the optimization side, the field has seen a transition from broad heuristic exploration toward increasingly rigorous, scalable, and differentiable approaches. Early and influential approaches employed metaheuristics such as particle swarm optimization, binary particle swarm optimization, and random search, motivated by the nonconvex nature of the design landscape and the difficulty of deriving useful gradients from many wake formulations (Wan et al., 2012; Feng & Shen, 2015; Pookpunt & Ongsakul, 2016). Comparative assessments showed that layout optimization is strongly multimodal, with many local optima and substantial dependence on initialization, objective formulation, and constraint treatment (Wang et al., 2015; Thomas & Ning, 2018). Yet the field has not remained at the heuristic stage. As differentiable wake models, algorithmic continuation methods, and efficient optimization infrastructures matured, gradient-based methods began to outperform or at least rival heuristic approaches in many continuous design settings, particularly for large-scale farms where computational efficiency becomes decisive (Guirguis et al., 2016; King et al., 2017; Stanley & Ning, 2019; Quick et al., 2023; Valotta Rodrigues et al., 2024). The integration of pyOptSparse and related large-scale optimization frameworks has further lowered barriers to deploying advanced constrained optimizers such as SNOPT in wind farm design studies (Gill et al., 2005; Wu et al., 2020).

This evolution in methods is not only a technical story about faster algorithms. It reflects a deeper conceptual change in how wind farms are understood as engineering systems. A modern wind plant is a coupled aerodynamic, structural, electrical, economic, and operational system, and layout interacts with nearly every one of these dimensions. Herbert-Acero et al. (2014) emphasized the broad methodological landscape of wind farm design and optimization, making clear that micrositing is inseparable from the broader design context. Fleming et al. (2016) and Gebraad et al. (2017) showed that layout and yaw control can be optimized jointly, demonstrating that a layout that appears suboptimal under conventional greedy turbine operation may become highly effective when wake steering is permitted. Feng and Shen (2017a) expanded the problem to include multiple turbine types, challenging the earlier assumption that all turbines in a farm are identical. Feng and Shen (2017b) further pushed the field toward robust design by investigating changing wind conditions and the sensitivity of power output to variability. These studies collectively indicate that the classical formulation of “maximize energy by placing identical turbines under fixed wind roses” is increasingly inadequate for modern wind plant design.

Another important driver of change is realism in constraints. Real wind farm development rarely occurs on unconstrained rectangular domains. Restricted zones, environmental exclusions, shipping lanes, bathymetric limitations, neighboring wind farms, permitting requirements, cable routing implications, and staggered leasing geometries all complicate the design region. Hou et al. (2016) addressed offshore restricted zones, showing that practical spatial limitations materially influence optimal layouts. Serrano González et al. (2018) examined neighboring offshore wind farms, drawing attention to inter-farm interactions that extend beyond project boundaries. More recently, Criado Risco et al. (2024) treated inclusion and exclusion zones within gradient-based optimization, reflecting the contemporary need to align mathematically efficient optimization with the fragmented realities of real project sites. Such studies are particularly important because they shift research away from elegant but impractical unconstrained models toward design formulations more likely to support decision-making in development practice.

The issue of uncertainty likewise redefines what counts as an optimal layout. Energy production estimates are conditioned on uncertain wind climate distributions, modeling assumptions, atmospheric stability states, terrain effects, and control behavior. Draxl et al. (2015) contributed significantly to this problem through the WIND Toolkit, which expanded access to high-resolution wind resource information for system studies. Murcia et al. (2015) investigated model evaluations for annual energy production prediction, highlighting discrepancies among AEP estimation methods. Padrón et al. (2019) introduced polynomial chaos for AEP computation, showing that uncertainty quantification can be integrated into layout-relevant energy estimation. When robust optimization enters the picture, the goal is no longer merely to maximize expected energy under nominal assumptions, but to identify configurations that maintain acceptable performance across plausible operating conditions (Feng & Shen, 2017b). This shift is essential because aggressive nominal optima may prove fragile in practice.

Despite the richness of the field, a central literature gap remains. Existing studies often specialize along one dimension while holding others fixed: wake model comparisons without deep optimizer analysis, optimizer comparisons under simplified wakes, robust design without realistic spatial constraints, or control co-design without full attention to computational scalability. Even strong review work, such as Herbert-Acero et al. (2014), Porté-Agel et al. (2020), and Meneveau (2019), leaves room for a synthesis focused specifically on the convergence of aerodynamic fidelity, optimization architecture, robustness, and deployability. Moreover, rapid advances in gradient-based continuation, stochastic gradients, large-scale frameworks, and analytical wake reductions have altered the trade space in recent years (Thomas et al., 2022; Quick et al., 2023; LoCascio et al., 2024; Valotta Rodrigues et al., 2024). A fresh synthesis is therefore warranted.

The present article addresses this gap by developing a comprehensive research-based analysis of wind farm layout optimization grounded strictly in the provided references. Rather than presenting a narrow review of one algorithm family or one wake model class, it interprets the field as an evolving systems problem. The article examines six interdependent dimensions: wake representation, optimization method selection, treatment of multimodality, spatial constraint realism, uncertainty and robustness, and control-layout co-design. The purpose is not merely to summarize prior work, but to clarify how the logic of the field has evolved, what methodological tensions remain unresolved, and which research pathways are most likely to deliver deployable and computationally viable wind farm optimization frameworks.

In doing so, the article advances three core arguments. First, wake model choice is not a secondary implementation detail but a fundamental determinant of optimizer suitability, computational cost, and practical interpretability (Baker et al., 2019; LoCascio et al., 2022). Second, gradient-based methods have become increasingly central, not because they eliminate nonconvexity, but because modern differentiable formulations, continuation strategies, and sparse large-scale infrastructures have made them especially effective for real-sized problems with many design variables and constraints (Guirguis et al., 2016; Stanley & Ning, 2019; Quick et al., 2023; Valotta Rodrigues et al., 2024). Third, a truly mature wind farm design methodology must move beyond static energy maximization to include operational control, heterogeneous turbine selection, uncertainty, and realistic exclusionary geometries (Feng & Shen, 2017a; Feng & Shen, 2017b; Gebraad et al., 2017; Criado Risco et al., 2024).

The remainder of the article develops these themes through a structured research synthesis. The methodology section explains the interpretive and comparative framework used to analyze the literature. The results section distills the major findings that emerge across studies concerning physics, algorithms, scalability, and real-world constraints. The discussion section then deepens the interpretation, examines unresolved tensions, outlines limitations in the current state of research, and proposes a forward-looking agenda for wind farm layout optimization as a mature engineering discipline.

Methodology

This article adopts a structured interpretive synthesis methodology designed for fields in which the most important scholarly challenge is not a lack of studies, but fragmentation across modeling assumptions, design objectives, and solution strategies. Wind farm layout optimization is particularly suitable for such an approach because its literature spans multiple disciplinary traditions, including atmospheric science, fluid mechanics, numerical optimization, offshore engineering, and wind plant systems engineering (Herbert-Acero et al., 2014; Meneveau, 2019; Porté-Agel et al., 2020). The aim of the present methodology is therefore not to estimate a pooled quantitative effect size, as would be expected in a statistical meta-analysis, but to generate a coherent, publication-ready analytical framework capable of organizing the logic of the field and deriving robust conclusions from heterogeneous but convergent evidence.

The study corpus consists strictly of the references provided. These sources were treated as the entire evidence base for the article. The decision to remain fully bounded by the supplied references ensures conceptual consistency with the task and prevents dilution of the interpretive structure through external materials. Within this corpus, the studies vary considerably in purpose. Some are experimental or measurement-oriented investigations of wake behavior and turbulence characteristics (Crespo & Hernández, 1996; Barthelmie et al., 2006; Bastankhah & Porté-Agel, 2016). Others focus on optimization algorithms and solution behavior under different formulations (Wan et al., 2012; Feng & Shen, 2015; Guirguis et al., 2016; Croonenbroeck & Hennecke, 2021). A further subset addresses integrated systems design, uncertainty quantification, or control-layout coupling (Fleming et al., 2016; Gebraad et al., 2017; Padrón et al., 2019). The methodological challenge was therefore to build a framework broad enough to capture these differences while still supporting specific, evidence-backed conclusions.

To achieve this, the synthesis proceeded in six analytical dimensions.

The first dimension was wake representation. Each study was read for the kind of wake model or wake physics assumption it employed or discussed. Particular attention was given to whether the work relied on empirical engineering wakes, Gaussian or analytical formulations, integral approaches, yaw-aware wake models, offshore measurement comparisons, or turbulence-oriented wake representations (Barthelmie et al., 2006; Niayifar & Porté-Agel, 2016; Parada et al., 2017; LoCascio et al., 2022; LoCascio et al., 2024). This dimension matters because wake models determine objective smoothness, fidelity, computational burden, and the degree to which optimization results can be trusted under site-specific conditions.

The second dimension was optimization architecture. Here the literature was grouped according to algorithm class: heuristic or metaheuristic methods such as particle swarm optimization and random search (Wan et al., 2012; Feng & Shen, 2015; Pookpunt & Ongsakul, 2016), classical constrained nonlinear programming methods such as sequential quadratic programming through SNOPT (Gill et al., 2005), gradient-based layout optimization and adjoint approaches (Guirguis et al., 2016; King et al., 2017), unified or comparative optimizer studies (Wang et al., 2015; Baker et al., 2019; Croonenbroeck & Hennecke, 2021), and recent stochastic or computational acceleration methods (Quick et al., 2023; Valotta Rodrigues et al., 2024). The purpose of this dimension was to determine under what conditions different algorithms are described as effective and how the field's preferences have evolved.

The third dimension was constraint handling and geometric realism. The reviewed literature includes idealized unconstrained domains but also more realistic formulations involving restricted zones, neighboring farms, inclusion and exclusion regions, and large offshore site constraints (Hou et al., 2016; Pillai et al., 2017; Serrano González et al., 2018; Criado Risco et al., 2024). This dimension was used to evaluate whether algorithmic claims remain persuasive when the design space resembles actual development conditions rather than abstract academic test cases.

The fourth dimension was uncertainty and robustness. Studies were examined for how they handled variability in wind climate, atmospheric conditions, or annual energy production estimation. Sources such as Draxl et al. (2015), Murcia et al. (2015), Feng and Shen (2017b), and Padrón et al. (2019) were central here. The goal was to identify whether the field still largely optimizes for nominal conditions or whether a meaningful transition toward robust design is underway.

The fifth dimension was control-layout co-design and turbine heterogeneity. This dimension emerged because some of the strongest recent contributions challenge the assumption that layout should be optimized independently of turbine operation or turbine selection. Fleming et al. (2016) and Gebraad et al. (2017) showed the importance of yaw control, while Feng and Shen (2017a) addressed the design implications of multiple turbine types. These studies were treated not as peripheral innovations but as indicators of a broader shift toward integrated wind plant engineering.

The sixth dimension was computational scalability and practical deployment. Because wind farms can involve large numbers of turbines, many wind conditions, and complex constraint sets, computational burden is not merely an implementation inconvenience but a determinant of what kind of optimization is practically feasible. Accordingly, the synthesis tracked how studies addressed scalability through adjoint gradients, continuation strategies, sparse optimization frameworks, stochastic gradient descent, simplified formulations, and faster analytical AEP models (Stanley & Ning, 2019; Thomas et al., 2022; Quick et al., 2023; LoCascio et al., 2024; Valotta Rodrigues et al., 2024; Wu et al., 2020).

After coding the literature across these six dimensions, the studies were analyzed comparatively rather than chronologically alone. This comparative strategy was important because a later study is not automatically superior in every respect; some older experimental papers remain foundational for validating or challenging the assumptions built into newer optimization frameworks. For example, Crespo and Hernández (1996) and Barthelmie et al. (2006) remain critical when assessing the physical credibility of layout models, even though more recent optimization studies may employ more advanced numerical strategies. Likewise, the interpretive value of a fast optimizer depends on the physical adequacy of the model it optimizes. Therefore, each methodological claim in the literature was interpreted within the chain linking aerodynamic representation, mathematical structure, and decision relevance.

The synthesis also used a cross-validating logic of triangulation. A claim was treated as especially robust when support emerged from different types of studies. For instance, the importance of wake losses is supported by both measurement-based offshore studies and optimization research focused on AEP maximization (Barthelmie et al., 2009; Barthelmie et al., 2010; Parada et al., 2017). Similarly, the growing importance of gradient-based methods is reinforced by studies on efficient gradient use, adjoint methods, simplified formulations, and stochastic large-scale optimization rather than by a single algorithm paper alone (Guirguis et al., 2016; King et al., 2017; Stanley & Ning, 2019; Quick et al., 2023). This triangulation approach reduces the risk of overinterpreting isolated findings.

Because the requested article required formal sections resembling an original research paper, the synthesis methodology was structured to generate results in the form of thematic findings rather than narrative literature notes. In this design, the “results” are the analytically derived patterns that emerge from the comparative reading of the corpus. This is appropriate because the underlying study is itself a research contribution: it develops a conceptual framework that organizes the field and produces interpretive findings about optimizer suitability, wake model dependencies, and future methodological priorities. Such an approach is consistent with rigorous review-based scholarship when the objective is theory-building or systems-level clarification rather than empirical hypothesis testing.

A further methodological choice concerned citation density and evidentiary discipline. Every major interpretive claim in the article is grounded in the supplied literature. Where a conclusion reflects synthesis across several works, multiple citations are used in order to preserve fidelity to the heterogeneity of the field. This is especially important in wind farm optimization, where conclusions that appear generic can become misleading if detached from the assumptions of specific model classes or site conditions. For example, the statement that gradient-based methods are advantageous is true only in a qualified sense and depends on the differentiability and smoothness of the problem formulation, as well as the availability of scalable infrastructures and effective treatment of local minima (Baker et al., 2019; Croonenbroeck & Hennecke, 2021; Thomas et al., 2022).

Another methodological principle was attention to non-equivalence among objectives. Some cited studies target pure annual energy production, while others implicitly consider broader efficiency or multidisciplinary design criteria. Some focus on offshore settings, some on general formulations, and others on control-integrated or robust design. Instead of forcing all studies into a single performance metric, the synthesis preserves these distinctions and interprets “optimization success” in context. This matters because an optimizer that excels in unconstrained AEP maximization may not be the best option when yaw control, exclusion zones, or heterogeneous turbines are introduced.

Finally, the methodology recognizes several inherent limitations. Because the article is bounded by the supplied references, it does not claim exhaustive coverage of all work in wind farm optimization. The results are interpretive and conceptual rather than empirical in the statistical sense. No new simulation campaign, field measurement set, or numerical benchmark is conducted here. However, this does not weaken the scholarly value of the study. On the contrary, in a field where knowledge has become segmented across methods and subdisciplines, a deeply comparative synthesis can make a substantial contribution by clarifying why apparently competing methods differ, under what assumptions they converge, and what constitutes a credible research agenda for the next phase of wind farm systems engineering (Herbert-Acero et al., 2014; Meneveau, 2019; Martins & Ning, 2021).

Results

The comparative synthesis of the provided literature reveals that wind farm layout optimization has matured into a field defined by the interplay of four central tensions: fidelity versus tractability, global exploration versus scalable local improvement, nominal performance versus robustness, and geometric idealization versus development realism. These tensions do not represent failures of the field. Rather, they explain why wind farm layout optimization has repeatedly changed methodological direction over time and why no single model or optimizer has permanently dominated the literature. The results presented below organize the major findings across aerodynamic representation, optimization strategy, multimodality treatment, real-world constraint handling, uncertainty, and integrated plant design.

A first major finding is that wake modeling remains the decisive foundation on which all optimization performance ultimately rests. The literature consistently demonstrates that layout optimization cannot be separated from the adequacy of the wake representation used to estimate energy production. This point may sound obvious, but its implications are profound. In optimization studies, much attention is often given to algorithm choice, convergence behavior, or computational speed. Yet the wind plant objective function is itself an aerodynamic abstraction, and if that abstraction misrepresents wake recovery, turbulence growth, yaw effects, or site-specific flow behavior, then even a perfectly converged optimizer can produce a misleading design. Early work by Crespo and Hernández (1996) showed that turbulence characteristics in turbine wakes are essential to understanding downstream flow development. Measurement-based comparisons by Barthelmie et al. (2006) demonstrated both the value and the limitations of wake model simulations when compared against offshore sodar observations. Subsequent offshore studies further quantified wake-induced power losses and underscored the real energy consequences of wake interactions at utility scale (Barthelmie et al., 2009; Barthelmie et al., 2010). Together, these studies establish that wake effects are not a secondary perturbation but one of the primary determinants of farm-level performance.

A related result is that the field has steadily moved from purely empirical or highly simplified wake assumptions toward more flexible and analytically structured representations. Niayifar and Porté-Agel (2016) advanced analytical wind farm modeling in a way that helped bridge physical interpretability and computational usability. Porté-Agel et al. (2020) later synthesized the broader state of wind turbine and wind farm flow research, showing that wake modeling now encompasses a spectrum ranging from engineering approximations to more physically informed models capable of representing deflection, expansion, and farm-scale flow interactions. Bastankhah and Porté-Agel (2016), in particular, broadened the wake optimization design space by clarifying behavior in yawed conditions. This matters because once yaw is allowed as a control variable, the notion of an “optimal layout” becomes conditional on operational strategy. More recent analytical reductions, such as FLOWERS and FLOWERS AEP, show that the latest direction in the field is not simply toward ever more complex wake physics, but toward carefully balanced models that retain enough aerodynamic structure to support optimization while dramatically reducing computational burden (LoCascio et al., 2022; LoCascio et al., 2024). Thus, the result is not that higher fidelity is always better in isolation. Rather, the most useful wake models are those whose fidelity is strategically aligned with the optimization task.

A second major finding is that optimizer performance is highly context dependent, and claims of superiority become unreliable when abstracted away from problem structure. Comparative studies consistently support this conclusion. Wang et al. (2015) showed that wind farm layout optimization methods differ materially in performance and suitability. Baker et al. (2019) went further by emphasizing best practices in selecting wake models and optimization algorithms together, effectively arguing that the optimizer should not be chosen independently of the model and problem setting. Croonenbroeck and Hennecke (2021), through a unified comparison of optimizers, reinforced the notion that there is no universally best method. Instead, performance depends on the number of decision variables, constraint structure, degree of smoothness, computational budget, and susceptibility to local optima. This finding is critical because it challenges a recurring tendency in engineering optimization literature to present new algorithms as generic improvements.

Within this context-dependent landscape, heuristic and metaheuristic methods remain historically important and still relevant in certain classes of problems. Particle swarm optimization and binary particle swarm optimization attracted attention because they can explore nonconvex and discontinuous design spaces without requiring gradient information (Wan et al., 2012; Pookpunt & Ongsakul, 2016). Random search likewise offered a simple but flexible mechanism for traversing complex objective surfaces (Feng & Shen, 2015). These methods proved especially appealing in earlier phases of wind farm optimization research when wake models were often non-differentiable, domain constraints were awkward to encode smoothly, and computational resources limited more elaborate formulations. The results across the literature suggest that heuristic methods remain useful when the design problem contains discrete variables, such as turbine type decisions, binary placement logic, or hard-to-smooth exclusionary rules (Feng & Shen, 2017a; Hou et al., 2016). Their enduring value lies less in raw efficiency and more in their tolerance for rugged design spaces and modeling irregularities.

However, the synthesis shows equally clearly that gradient-based methods have gained substantial and arguably dominant importance for large-scale continuous wind farm optimization. Guirguis et al. (2016) demonstrated the efficiency advantages of using gradient information in layout optimization. Guirguis et al. (2017) further embedded this logic in a multidisciplinary design context, indicating that gradient-based thinking is compatible with broader wind farm engineering objectives rather than layout alone. King et al. (2017) showed the potential of adjoint approaches for wind plant layout optimization, a particularly important contribution because adjoint methods can provide gradients efficiently even when the number of design variables becomes large. These developments align closely with more general principles of engineering design optimization, which emphasize that scalable gradient computation is often the difference between conceptually appealing design formulations and practically solvable ones (Kochenderfer & Wheeler, 2019; Martins & Ning, 2021).

The literature further indicates that the success of gradient-based optimization in wind farm design was not automatic. It required reformulating the problem to reduce pathological multimodality and improve objective smoothness. Thomas and Ning (2018) explicitly addressed the challenge of multi-modality in wind farm layout optimization, showing that the problem landscape can trap local-search methods in poor designs. Stanley and Ning (2019) contributed a simplification framework that made layout optimization more tractable. Thomas et al. (2022) later introduced a wake expansion continuation method, which is especially significant because continuation techniques can start optimization on a smoothed or relaxed version of the problem and progressively recover the original problem complexity. This is a methodological turning point. It suggests that the field's progress has depended not only on better optimizers, but also on engineering the optimization landscape itself so that gradients become more informative and local search becomes less fragile.

A fourth result is that recent work on stochastic gradients and computational acceleration reflects the field's transition from proof-of-concept optimization toward genuinely large wind farm deployment scenarios. Quick et al. (2023) proposed stochastic gradient descent for wind farm optimization, demonstrating that ideas from large-scale machine learning and stochastic optimization can be adapted to wind plant design. Valotta Rodrigues et al. (2024) focused directly on speeding up optimization for large wind farms, underscoring that scalability is now a first-order research goal. This trend is consistent with the integration of large-scale optimization infrastructures such as pyOptSparse, which provides a flexible framework for deploying advanced optimizers efficiently (Wu et al., 2020). These studies collectively suggest that the practical bottleneck in modern wind farm optimization is increasingly not whether an optimum exists conceptually, but whether sufficiently accurate and robust solutions can be obtained within real computational budgets for industrially sized projects.

A fifth major finding concerns the transition from unconstrained geometric optimization toward realistic spatial design. Earlier benchmark studies often used regular domains and minimum-spacing rules as the main constraints. While such formulations were valuable for method development, they do not adequately represent actual development conditions. Hou et al. (2016) showed that offshore wind farm layout optimization in restricted zones changes the problem significantly, making constraint management central to the design. Pillai et al. (2017), in the context of Middelgrunden offshore optimization, reinforced that real project settings impose practical configuration considerations that extend beyond abstract AEP maximization. Serrano González et al. (2018) added another dimension by analyzing neighboring offshore wind farms, showing that spatial optimization cannot always be treated as isolated to a single project boundary. Criado Risco et al. (2024) then incorporated inclusion and exclusion zones directly into gradient-based layout optimization. The result emerging from these studies is that modern optimization frameworks are increasingly judged by their ability to accommodate fragmented and regulation-shaped design spaces, not merely by performance on rectangular test layouts.

The significance of this finding extends beyond engineering elegance. Realistic constraint handling changes the comparative performance of optimizers.

Heuristic methods sometimes retain advantages when constraints are irregular or naturally discrete, but recent gradient-based methods have increasingly closed this gap by embedding constraints into differentiable formulations or using continuation and efficient nonlinear programming frameworks (Gill et al., 2005; Criado Risco et al., 2024). This means the old dichotomy—heuristics for realism, gradients for idealized smooth problems—is becoming less absolute. The field is moving toward hybrid sophistication in which gradient-based methods remain viable even under practical spatial complexity, provided the model architecture is designed carefully.

A sixth result is that robust design and uncertainty treatment remain essential but comparatively underintegrated components of wind farm optimization. Wind farm layouts are typically optimized with respect to annual energy production, yet AEP itself is a model-derived expectation conditioned on wind resource assumptions, wake formulations, and probability distributions over inflow states. Draxl et al. (2015) contributed the WIND Toolkit, which significantly enhanced the quality and accessibility of wind resource characterization. Murcia et al. (2015) examined model evaluation for AEP prediction, demonstrating that different predictive models can yield different energy estimates. Padrón et al. (2019) introduced polynomial chaos for AEP computation, indicating that uncertainty quantification can be incorporated more systematically into the energy evaluation pipeline. Feng and Shen (2017b) brought robustness directly into layout optimization by quantifying performance under changing wind conditions.

The synthesis shows that these contributions collectively undermine the legitimacy of nominal optimization as the sole design objective. A layout that maximizes predicted output for one wind rose or one assumed wake environment may perform poorly if the true wind climate differs or if the model mischaracterizes directional variability. Robustness therefore emerges as more than a supplementary consideration; it is a corrective against overconfident optimization. However, the literature also suggests that robust optimization remains less mature than deterministic AEP maximization. The computational burden of evaluating uncertainty-rich objectives, combined with the already high cost of wake modeling, has slowed full integration. Thus, the field appears to be in transition: aware that uncertainty matters deeply, increasingly equipped with tools to address it, but not yet fully standardized in robust design practice.

A seventh key finding is that wind farm layout optimization is increasingly inseparable from operational control. Fleming et al. (2016) demonstrated that wind plant system engineering can benefit from joint optimization of layout and yaw control. Gebraad et al. (2017) extended this by showing that annual energy production can be improved through simultaneous layout and yaw control optimization. These studies have transformative implications. They show that the layout problem is not static. A turbine arrangement cannot be judged solely by how it performs under greedily optimized individual turbine operation. Wake steering alters the aerodynamic coupling between machines and can make some turbine placements more attractive than they would appear under traditional assumptions. Bastankhah and Porté-Agel's (2016) study of yawed wakes strengthens the physical basis for this design-control integration.

The consequence of this result is conceptual as much as computational. Classical micrositing assumes that placement determines wake exposure and that operation adapts only locally. The newer literature instead suggests that the wind farm should be treated as an actively managed flow-control system. In such a system, layout creates the structural possibilities within which control operates, and control in turn changes the effective value of layout configurations. This co-design perspective is one of the clearest signs that wind farm optimization has become a systems-engineering discipline rather than a narrow placement exercise.

An eighth result concerns turbine heterogeneity. Feng and Shen (2017a) showed that offshore wind farm design can involve multiple turbine types, which broadens the design problem from where to place turbines to what turbines to place. This has important implications for optimization structure because mixed turbine classes introduce discrete or mixed-integer characteristics, alter wake interactions, and change cost-energy tradeoffs. The literature synthesized here suggests that heterogeneity is still less developed than continuous layout optimization, but its appearance is theoretically significant. It signals a move away from a one-technology-fits-all view of wind farm design. Once heterogeneity is admitted, even a purely geometric optimum may cease to be meaningful unless turbine class selection is optimized jointly.

A ninth finding is that terrain complexity and site realism remain major challenges to generalizable optimization. Allen et al. (2020) investigated simulation and layout optimization in complex terrain, demonstrating that terrain-induced flow heterogeneity complicates the extrapolation of methods developed in simpler settings. Much of the literature on wind farm layout optimization has been driven by offshore or simplified onshore contexts where inflow can be represented more cleanly. Yet the result from terrain-aware studies is that optimization quality depends not only on wake interactions but also on how accurately the background flow field is represented. This introduces another layer to the fidelity-versus-tractability tradeoff. If complex terrain requires more expensive flow modeling, then the optimizer and wake model must be chosen in light of a larger simulation stack. Consequently, methods that are efficient in flat or homogeneous conditions may lose their practical advantage when integrated into terrain-aware pipelines.

A tenth finding is that the field increasingly recognizes the importance of standardization, benchmarking, and best-practice guidance. Baker et al. (2019) explicitly addressed wake model and optimization algorithm selection through best-practice recommendations. Croonenbroeck and Hennecke (2021) compared optimizers within a unified standard. Such works indicate that the field is no longer satisfied with isolated case-study claims. Instead, there is a growing expectation that methods be evaluated on comparable grounds, with careful attention to problem formulation and reproducibility. This shift is important because many earlier debates over “best” layout methods were distorted by inconsistent assumptions, distinct test domains, and incomparable objective calculations. Standardization does not eliminate legitimate methodological diversity, but it helps clarify which differences matter and which reflect artifact rather than substance.

Taken together, these results show that wind farm layout optimization is moving toward a mature architecture characterized by differentiable but physically credible wake models, efficient gradient-based large-scale optimization, explicit management of multimodality, realistic geometric constraints, integration of uncertainty, and joint treatment of layout with operational control. Yet the transition is incomplete. The literature still exhibits fragmentation between high-fidelity physics and scalable computation, between deterministic optima and robust design, and between elegant methods and real-world deployment requirements. The most important result of the synthesis is therefore not that one method has won, but that the field now possesses the conceptual ingredients for a next-generation integrated optimization framework.

Discussion

The results of this synthesis suggest that wind farm layout optimization has entered a more mature but also more demanding phase of development. Earlier stages of the field were driven by a relatively straightforward question: how should turbines be placed to minimize wake losses and maximize energy output? That question remains central, but the literature now shows that it can no longer be answered in a narrow way. The optimality of a layout depends on what wake model is used, which wind climate assumptions are adopted, how uncertainty is represented, whether operational wake steering is available, whether multiple turbine classes are allowed, what geometric exclusions apply, and how much computational effort can realistically be expended. In effect, the field has evolved from a single-objective micrositing problem into a layered systems-engineering problem, and this evolution has implications for both scholarship and practice.

One important interpretive lesson from the literature is that methodological progress in wind farm optimization has not been linear. There has been no simple sequence in which cruder methods were replaced by better methods in a universally improving chain. Instead, the field has advanced through cycles of expansion and correction. Simplified wake models enabled early optimization studies but often obscured the physical subtleties revealed by measurement work and turbulence research (Crespo & Hernández, 1996; Barthelmie et al., 2006). Heuristic algorithms made it possible to search rugged landscapes before differentiable infrastructures were mature, but they often suffered from computational inefficiency and uncertain scalability (Wan et al., 2012; Feng & Shen, 2015). Gradient-based methods later gained prominence, but only after the research community developed smoother formulations, continuation strategies, adjoint gradients, and sparse optimization environments capable of handling the structure of large wind farm problems (Guirguis et al., 2016; King et al., 2017; Thomas et al., 2022; Wu et al., 2020). This pattern shows that progress in the field depends less on replacing old ideas wholesale than on re-engineering the assumptions that make particular methods effective.

This insight has direct implications for the frequent debate over fidelity versus tractability. A common but misleading framing is that researchers must choose either physically realistic models that are too expensive to optimize or computationally convenient models that are too simplistic to trust. The literature suggests a more nuanced view. What matters is not maximum fidelity in the abstract, but task-appropriate fidelity. For some early-stage design decisions or large-scale screening analyses, a carefully calibrated analytical or integral wake model may offer the best balance of speed and decision relevance (Niayifar & Porté-Agel, 2016; LoCascio et al., 2022; LoCascio et al., 2024). For terrain-sensitive or control-intensive studies, more detailed flow representation may be necessary, but even then the optimization must be architected so that the computational burden remains manageable (Allen et al., 2020). The real intellectual achievement of recent wind farm optimization research lies precisely in finding models that are not simply more detailed, but more usable without losing interpretive integrity.

The rise of gradient-based methods should be understood through this same lens. It would be an oversimplification to say that gradient methods are inherently superior to heuristics. The literature does not support such a universal statement (Baker et al., 2019; Croonenbroeck & Hennecke, 2021). Instead, gradient methods have become increasingly attractive because the field has created conditions under which their strengths can be realized. Wake models have become smoother or more differentiable. Continuation methods reduce effective multimodality. Sparse, large-scale optimization software makes high-dimensional constrained problems computationally feasible. Adjoint methods reduce gradient cost when design variable counts grow. In other words, gradient dominance is contingent, constructed, and methodological. It is not an intrinsic property of the layout problem itself.

This distinction matters because it prevents a new orthodoxy from replacing an old one. Heuristic methods continue to offer value where discontinuities, mixed-integer features, or irregular constraints dominate the formulation. Multiple turbine types, strongly discrete siting rules, and certain restricted-zone formulations may still benefit from heuristic or hybrid search strategies (Feng & Shen, 2017a; Hou et al., 2016). Moreover, heuristic methods sometimes provide diverse candidate layouts that can be useful for subsequent refinement or decision support even when they do not scale as well as gradient approaches. Thus, the most mature position is not to choose between heuristics and gradients in a doctrinaire way, but to align the algorithm with the structure of the design problem and, where useful, combine broad exploratory search with efficient local improvement.

The literature also points to a deeper conceptual issue: layout optimization is only meaningful relative to an operational doctrine. Classical wind farm design implicitly assumes that each turbine acts to maximize its own power subject to static placement. Yet the emergence of wake steering and yaw control research destabilizes that assumption (Fleming et al., 2016; Gebraad et al., 2017). A layout that is poor under greedy individual operation may become effective when upstream turbines deliberately redirect wakes away from downstream machines. Conversely, a layout designed without considering control may leave value unrealized during plant operation. This implies that the true design variable is not only the location of turbines, but the joint configuration of placement and controllability. Such a shift is highly significant for engineering design philosophy. It reframes the wind farm from a passive arrangement of machines into an actively managed aerodynamic network.

At the same time, control-layout co-design introduces new complexities that the literature has only begun to address fully. Yaw control can increase energy production, but it may also affect structural loads, control effort, maintenance requirements, and the predictability of farm behavior under varying atmospheric conditions. While the provided references focus primarily on AEP improvements and optimization structure, they implicitly raise the broader question of whether energy-only co-design sufficiently captures lifecycle performance (Fleming et al., 2016; Gebraad et al., 2017). This suggests a future research direction in which layout and control are optimized under multidimensional objectives that include fatigue, reliability, and perhaps economic cost structures. The transition to such formulations would further strengthen the integration of wind farm optimization with full plant systems engineering.

Another important discussion point concerns robustness. Nominal optimization remains common because it is computationally simpler and because annual energy production provides a convenient scalar objective. However, the literature makes clear that AEP is itself a compressed representation of uncertainty-rich phenomena. Wind roses are estimates. Atmospheric states vary. Wake models carry structural assumptions. Resource datasets, however valuable, are not exact (Draxl et al., 2015; Murcia et al., 2015). When designers optimize a layout for one assumed distribution of inflow conditions, they may unintentionally optimize against a moving target. Robustness studies such as Feng and Shen (2017b), together with uncertainty-aware AEP computation approaches like Padrón et al. (2019), suggest that the next stage of methodological maturity lies in optimizing not only mean performance but also resilience to uncertainty.

This has important practical implications. In real project development, investors and planners care not merely about maximum predicted output but also about confidence in performance, downside risk, and sensitivity to site uncertainty. A design with marginally lower nominal AEP but substantially greater robustness may be economically preferable. Yet the literature indicates that robust optimization remains less consolidated than deterministic optimization. There is still no broadly standardized framework for incorporating wind climate uncertainty, model-form uncertainty, and operational variability into industrially scalable layout design. The reason is understandable: uncertainty multiplies computational cost, especially when coupled with already expensive wake evaluations. Nonetheless, the conceptual direction is clear. The field cannot remain centered on nominal optima indefinitely without risking a growing mismatch between research outputs and decision-making needs.

Constraint realism raises similar concerns. The movement toward inclusion zones, exclusion zones, neighboring-farm interactions, and restricted offshore areas is one of the most encouraging developments in the literature (Hou et al., 2016; Serrano González et al., 2018; Criado Risco et al., 2024). It indicates a shift from academic abstraction toward project relevance. However, there remains a tension between mathematical elegance and development complexity. Real sites are affected not only by the geometric constraints represented in current optimization studies, but also by seabed conditions, environmental permitting windows, construction logistics, electrical infrastructure placement, and multi-stakeholder planning processes. The provided references do not exhaust these dimensions, but they point strongly in that direction. Therefore, one of the major unresolved challenges for the field is how to preserve the efficiency of advanced optimization methods while embedding progressively richer representations of real-world constraints. This is not merely a software challenge. It is a modeling philosophy challenge about how much complexity to include and at what stage of the design pipeline.

The discussion of terrain complexity underscores this point further. Studies such as Allen et al. (2020) reveal that the generalization of methods from simple to complex flow environments is not guaranteed. In complex terrain, background flow heterogeneity interacts with wakes in ways that can change the relative importance of spacing, alignment, and local placement. This means that some conclusions drawn from offshore or idealized settings may not transfer seamlessly to mountainous or topographically varied sites. The broader implication is that wind farm layout optimization may need a tiered methodological framework rather than a single universal approach. Different site classes may require different balances of flow-model fidelity, optimization structure, and uncertainty treatment. Attempting to force all wind farm design problems into one canonical formulation could therefore be counterproductive.

The literature also invites reflection on the role of standardization. Best-practice and unified-comparison studies are highly valuable because they reduce confusion created by incomparable assumptions and inconsistent benchmarking (Baker et al., 2019; Croonenbroeck & Hennecke, 2021). Yet standardization must be handled carefully. If benchmark problems are too simplified, they may reward methods that perform well in artificial settings but poorly in realistic ones. If benchmark models are too computationally expensive, they may discourage broader comparative work. The ideal standardization strategy is therefore one that includes layered benchmark families: simple problems for diagnostic comparison, intermediate problems for algorithmic realism, and large constraint-rich scenarios for deployability testing. Although the provided references do not formalize such a hierarchy explicitly, the logic of the field strongly supports it.

The computational literature further suggests that scale is becoming the defining practical constraint of the next decade. Large wind farms involve many turbines, many wind directions, many wind speeds, and increasingly many coupled control or uncertainty variables. Without efficient optimization architectures, research risks producing elegant but operationally unusable methods. The importance of pyOptSparse, SNOPT, stochastic gradients, speed-up strategies, and continuation methods should be interpreted in this light (Gill et al., 2005; Wu et al., 2020; Quick et al., 2023; Valotta Rodrigues et al., 2024). Computational scalability is not peripheral; it is what determines whether robust, control-aware, constraint-rich optimization can move from research prototypes into routine design practice.

The present synthesis has several limitations that should be acknowledged openly. First, the analysis is bounded strictly by the supplied references. While these references are rich and collectively representative of many core debates in the field, no claim is made that the article covers every branch of wind farm optimization research. Second, the article is interpretive rather than simulation-based. It does not present new benchmark runs comparing algorithms under a common numerical setup. Third, because the literature itself varies in assumptions, objective functions, and site types, some of the conclusions necessarily remain conditional rather than absolute. These limitations, however, do not weaken the main findings. Instead, they are characteristic of a rigorous synthesis in a multidisciplinary field where the key scholarly contribution lies in organizing knowledge coherently and identifying the most consequential lines of convergence and divergence.

From this synthesis, several future research priorities become evident. One priority is the development of optimization frameworks that integrate differentiable wake models, explicit constraint handling, and uncertainty-aware objectives without prohibitive computational cost. Another is broader integration of layout with operational control, possibly extending beyond yaw to more comprehensive plant-level coordination. A third is the maturation of mixed-design formulations involving heterogeneous turbine classes, lifecycle considerations, and perhaps economic or reliability measures beyond energy alone. A fourth is the creation of more structured benchmark ecosystems that reflect varying levels of site realism. Finally, there is a need for stronger connection between optimization research and development practice, especially in how layouts are evaluated under regulatory, construction, and neighboring-farm realities.

In a broader sense, the literature on wind farm optimization provides an instructive case study in modern engineering design. It shows that optimization is not simply about algorithms, nor simply about physics, nor simply about data. It is about building a coherent relationship among these elements so that a solution is mathematically efficient, physically credible, and practically actionable. Wind farm layout optimization has progressed substantially because researchers have gradually learned to treat it in exactly this way. The most important remaining challenge is to complete that transition and move from fragmented excellence in subproblems to fully integrated wind plant design methodologies.

Conclusion

Wind farm layout optimization has developed from a wake-loss minimization exercise into a sophisticated systems-engineering discipline shaped by aerodynamic fidelity, algorithmic architecture, uncertainty treatment, control integration, and real-world spatial constraints. The literature synthesized in this article demonstrates that wake interactions are foundational to farm performance and that their representation strongly influences the reliability and usefulness of any optimization result (Crespo & Hernández, 1996; Barthelmie et al., 2009; Porté-Agel et al., 2020). It also shows that optimizer choice cannot be separated from wake model structure, problem smoothness, and constraint realism (Baker et al., 2019; Croonenbroeck & Hennecke, 2021).

A central conclusion of this study is that gradient-based methods have become increasingly important because the field has created the modeling and computational conditions necessary for them to succeed. Differentiable wake formulations, adjoint gradients, problem simplification, continuation strategies, and scalable optimization frameworks have together made large-scale constrained optimization far more practical than in earlier phases of research (Guirguis et al., 2016; King et al., 2017; Stanley & Ning, 2019; Thomas et al., 2022; Wu et al., 2020). At the same time, heuristic and metaheuristic methods remain relevant in discrete, rugged, or highly irregular design spaces, especially where mixed turbine types or non-smooth spatial rules dominate the formulation (Wan et al., 2012; Pookpunt & Ongsakul, 2016; Feng & Shen, 2017a).

The synthesis further concludes that the field is moving decisively beyond static deterministic energy maximization. Robustness under changing wind conditions, uncertainty-aware AEP computation, operational wake steering, and inclusion or exclusion zones are no longer peripheral concerns; they are increasingly central to what constitutes a credible layout optimization framework (Feng & Shen, 2017b; Padrón et al., 2019; Gebraad et al., 2017; Criado Risco et al., 2024). This shift is especially important because it narrows the gap between academic optimization studies and the realities of wind plant development.

The most promising future direction is therefore an integrated optimization paradigm that combines physically informed but computationally efficient wake models, scalable gradient-based search, robust treatment of uncertainty, explicit real-world constraints, and joint optimization of layout with operational control. Such a framework would reflect the true character of modern wind farms as actively managed, spatially constrained, uncertainty-exposed energy systems rather than static turbine arrays. The provided literature indicates that the conceptual and methodological foundations for this next stage already exist. The challenge ahead is to unify them into deployable design workflows capable of serving both scientific rigor and industrial practice.

References

1. Allen, J., King, R., & Barter, G. (2020). Wind farm simulation and layout optimization in complex terrain. Journal of Physics: Conference Series, 1452, 012066. https://doi.org/10.1088/1742-6596/1452/1/012066

2. Baker, N. F., Stanley, A. P. J., Thomas, J. J., Ning, A., & Dykes, K. (2019). Best practices for wake model and optimization algorithm selection in wind farm layout optimization. In AIAA Scitech 2019 Forum (pp. 1–18).

3. Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J. G., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E. S., & Chaviaropoulos, P. K. (2009). Modelling and measuring flow and wind turbine wakes in large wind farms offshore. Wind Energy, 12, 431–444. https://doi.org/10.1002/we.348

4. Barthelmie, R. J., Larsen, G. C., Frandsen, S. T., Folkerts, L., Rados, K., Pryor, S. C., Lange, B., & Schepers, G. (2006). Comparison of wake model simulations with offshore wind turbine wake profiles measured by sodar. Journal of Atmospheric and Oceanic Technology, 23, 888–901. https://doi.org/10.1175/JTECH1886.1

5. Barthelmie, R. J., Pryor, S. C., Frandsen, S. T., Hansen, K. S., Schepers, J. G., Rados, K., Schlez, W., Neubert, A., Jensen, L. E., & Neckelmann, S. (2010). Quantifying the impact of wind turbine wakes on power output at offshore wind farms. Journal of Atmospheric and Oceanic Technology, 27, 1302–1317. https://doi.org/10.1175/2010JTECHA1398.1

6. Bastankhah, M., & Porté-Agel, F. (2016). Experimental and theoretical study of wind turbine wakes in yawed conditions. Journal of Fluid Mechanics, 806, 506–541. https://doi.org/10.1017/jfm.2016.595

7. Criado Risco, J., Valotta Rodrigues, R., Friis-Møller, M., Quick, J., Pedersen, M. M., & Réthoré, P.-E. (2024). Gradient-based wind farm layout optimization with inclusion and exclusion zones.

Wind Energy Science, 9, 585–600. https://doi.org/10.5194/wes-9-585-2024

8. Crespo, A., & Hernández, J. (1996). Turbulence characteristics in wind-turbine wakes. Journal of Wind Engineering and Industrial Aerodynamics, 61, 71–85. https://doi.org/10.1016/0167-6105(95)00033-X

9. Croonenbroeck, C., & Hennecke, D. (2021). A comparison of optimizers in a unified standard for wind farm layout optimization. Energy, 216, 119244. https://doi.org/10.1016/j.energy.2020.119244

10. Draxl, C., Clifton, A., Hodge, B.-M., & McCaa, J. (2015). The wind integration national dataset (WIND) toolkit. Applied Energy, 151, 355–366. https://doi.org/10.1016/j.apenergy.2015.03.121

11. Feng, J., & Shen, W. Z. (2015). Solving the wind farm layout optimization problem using random search algorithm. Renewable Energy, 78, 182–192. https://doi.org/10.1016/j.renene.2015.01.005

12. Feng, J., & Shen, W. Z. (2017a). Design optimization of offshore wind farms with multiple types of wind turbines. Applied Energy, 205, 1283–1297. https://doi.org/10.1016/j.apenergy.2017.08.107

13. Feng, J., & Shen, W. Z. (2017b). Wind farm power production in changing wind: Robustness quantification and layout optimization. Energy Conversion and Management, 148, 905–914. https://doi.org/10.1016/j.enconman.2017.06.005

14. Fleming, P. A., Ning, A., Gebraad, P. M. O., & Dykes, K. (2016). Wind plant system engineering through optimization of layout and yaw control. Wind Energy, 19, 329–344. https://doi.org/10.1002/we.1836

15. Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., & Dykes, K. (2017). Maximization of annual energy production through layout and yaw control optimization. Wind Energy, 20, 97–107. https://doi.org/10.1002/we.1993

16. Gill, P. E., Murray, W., & Saunders, M. A. (2005). SNOPT: An SQP algorithm for large-scale constrained optimization. SIAM Review, 47, 99–131. https://doi.org/10.1137/S0036144504446096

17. Guirguis, D., Romero, D. A., & Amon, C. H. (2016). Efficient optimization of wind farm layouts using gradient information. Applied Energy, 179, 110–123. https://doi.org/10.1016/j.apenergy.2016.06.101

18. Guirguis, D., Romero, D. A., & Amon, C. H. (2017). Gradient-based multidisciplinary design of wind farms. Applied Energy, 197, 279–291. https://doi.org/10.1016/j.apenergy.2017.04.030

19. Herbert-Acero, J., Probst, O., Réthoré, P.-E., Larsen, G., & Castillo-Villar, K. (2014). Methodological approaches for wind farm design and optimization. Energies, 7, 6930–7016. https://doi.org/10.3390/en7116930

20. Hou, P., Hu, W., Chen, C., Soltani, M., & Chen, Z. (2016). Optimization of offshore wind farm layout in restricted zones. Energy, 113, 487–496. https://doi.org/10.1016/j.energy.2016.07.062

21. King, R., Dykes, K., Graf, P., & Hamlington, P. E. (2017). Optimization of wind plant layouts using an adjoint approach. Wind Energy Science, 2, 115–131. https://doi.org/10.5194/wes-2-115-2017

22. Kochenderfer, M. J., & Wheeler, T. A. (2019). Algorithms for optimization. MIT Press.

23. LoCascio, M. J., Bay, C. J., Bastankhah, M., Barter, G. E., Fleming, P. A., & Martínez-Tossas, L. A. (2022). FLOWERS: An integral approach to wake modeling. Wind Energy Science, 7, 1137–1151. https://doi.org/10.5194/wes-7-1137-2022

24. LoCascio, M. J., Bay, C. J., Martínez-Tossas, L. A., Bastankhah, M., & Gorlé, C. (2024). FLOWERS AEP: Analytical model for layout optimization. Wind Energy, 27, 1563–1580. https://doi.org/10.1002/we.2954

25. Martins, J. R. R. A., & Ning, A. (2021). Engineering design optimization. Cambridge University Press.

26. Meneveau, C. (2019). Big wind power: Questions for turbulence research. Journal of Turbulence, 20, 2–20. https://doi.org/10.1080/14685248.2019.1584664

27. Murcia, J. P., Réthoré, P. E., Natarajan, A., & Sørensen, J. D. (2015). Model evaluations for predicting AEP. Journal of Physics: Conference Series, 625, 012030. https://doi.org/10.1088/1742-6596/625/1/012030

28. Niayifar, A., & Porté-Agel, F. (2016). Analytical modeling of wind farms. Energies, 9, 741. https://doi.org/10.3390/en9090741

29. Ning, A., & Petch, D. (2016). Integrated design of wind turbines using gradients. Wind Energy, 19, 2137–2152. https://doi.org/10.1002/we.1972

30. Padrón, A. S., Thomas, J., Stanley, A. P. J., Alonso, J. J., & Ning, A. (2019). Polynomial chaos for AEP computation. Wind Energy Science, 4, 211–231. https://doi.org/10.5194/wes-4-211-2019

31. Parada, L., Herrera, C., Flores, P., & Parada, V. (2017). Wind farm layout optimization using Gaussian wake model. Renewable Energy, 107, 531–541. https://doi.org/10.1016/j.renene.2017.02.017

32. Pillai, A. C., Chick, J., Khorasanchi, M., Barbouchi, S., & Johanning, L. (2017). Offshore wind farm optimization at Middelgrunden. Ocean Engineering, 139, 287–297. https://doi.org/10.1016/j.oceaneng.2017.04.049

33. Pookpunt, S., & Ongsakul, W. (2016). Optimal wind farm configuration using binary PSO. Energy Conversion and Management, 108, 160–180. https://doi.org/10.1016/j.enconman.2015.11.002

34. Porté-Agel, F., Bastankhah, M., & Shamsoddin, S. (2020). Wind-turbine and wind-farm flows: A review. Boundary-Layer Meteorology, 174, 1–59. https://doi.org/10.1007/s10546-019-00473-0

35. Quick, J., Réthoré, P.-E., Pedersen, M. M., Rodrigues, R. V., & Friis-Møller, M. (2023). Stochastic gradient descent for wind farm optimization. Wind Energy Science, 8, 1235–1250. https://doi.org/10.5194/wes-8-1235-2023

36. Serrano González, J., Burgos Payán, M., & Riquelme Santos, J. M. (2018). Optimal design of neighbouring offshore wind farms. Applied Energy, 209, 140–152. https://doi.org/10.1016/j.apenergy.2017.10.120

37. Stanley, A. P. J., & Ning, A. (2019). Simplification of wind farm layout optimization. Wind Energy Science, 4, 663–676. https://doi.org/10.5194/wes-4-663-2019

38. Thomas, J. J., & Ning, A. (2018). Reducing multi-modality in wind farm layout optimization. Journal of Physics: Conference Series, 1037, 042012. https://doi.org/10.1088/1742-6596/1037/4/042012

39. Thomas, J. J., McOmber, S., & Ning, A. (2022). Wake expansion continuation method. Wind Energy, 25, 678–699. https://doi.org/10.1002/we.2692

40. Valotta Rodrigues, R., Pedersen, M. M., Schøler, J. P., Quick, J., & Réthoré, P.-E. (2024). Speeding up large wind farm optimization. Wind Energy Science, 9, 321–341. https://doi.org/10.5194/wes-9-321-2024

41. Wan, C., Wang, J., Yang, G., Gu, H., & Zhang, X. (2012). Wind farm micro-siting using particle swarm optimization. Renewable Energy, 48, 276–286. https://doi.org/10.1016/j.renene.2012.04.052

42. Wang, L., Tan, A. C., & Gu, Y. (2015). Comparative study on wind farm layout optimization methods. Journal of Wind Engineering and Industrial Aerodynamics, 146, 1–10. https://doi.org/10.1016/j.jweia.2015.07.009

43. Wu, N., Kenway, G., Mader, C., Jasa, J., & Martins, J. (2020). pyOptSparse: Framework for large-scale optimization. Journal of Open Source Software, 5, 2564. https://doi.org/10.21105/joss.02564