The Joint Optimization of a Wind Farm Layout and Wake Steering

A case study of IJmuiden Ver on the co-design of the wind farm layout and yaw control

Robin de Jong

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

The Joint Optimization of a Wind Farm Layout and Wake Steering

A case study of IJmuiden Ver on the co-design of the wind farm layout and yaw control

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

Robin de Jong

December 28, 2023

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of Technology

The work in this thesis was supported by Eneco.

Delft Center for
Systems and Control

Abstract

Maximizing the extraction of energy from wind farms with ever higher densities is becoming increasingly more important in order to achieve climate targets and simultaneously preserve nature. Improving the yield of a wind farm can be achieved by optimizing the layout, applying control, especially wake steering through yaw control has shown great results, or even combining the optimization of the layout and control into one joint optimization. In this thesis, a case study is performed on the Dutch wind farm 'IJmuiden Ver' to investigate the real-world applicability of joint optimization. The employed method uses the genetic algorithm, capable of handling the discontinuous domain, and an improved version of the geometric yaw relationship, making coupled or nested optimization redundant. In the IJmuiden Ver case, the levelized cost of electricity (LCOE) of a joint optimized layout compared to a sequential optimized layout is around 0.3% better, even remaining around 0.2% to 0.3% better when shrinking the domain to give nature more space. This shows that joint optimization is applicable in practice and has the potential to increase the yield of a wind farm substantially without significantly increasing the computational intensity of the wind farm layout optimization problem (WFLOP).

Acknowledgements

This thesis would not have been possible without the support of my supervisor and colleagues at Eneco, my supervisor and members of the data driven control group of the TU Delft, friends, and family. Special thanks to René Bos, my supervisor from Eneco. He always had time for a meeting, despite his busy schedule, and made me feel positive about the road ahead. With a thesis, things can become blurry, but he helped put everything back into perspective and made sure I kept my eye on the ball. I would also like to emphasize the help of Jan-Willem van Wingerden, my supervisor from TU Delft. He made the collaboration with Eneco possible and guided me during our meetings. He made sure that I paid attention to the aspects that are important for academic research. His vast knowledge on the subject and awareness of current studies resulted in interesting insights. Last but not least I would like to thank Lombard de Leeuw, who supported me in many ways during this thesis.

Table of Contents

Abstract i

Acknowledgements iii

1 Introduction 1
- 1-1 Wake effect . . . 1
- 1-2 Wind farm control methods . . . 2
- 1-3 Wind farm layout . . . 5
- 1-4 Optimization problem . . . 7
- 1-5 Optimization method . . . 7
- 1-6 Wake model . . . 8
- 1-7 Contribution . . . 9

2 Methods 11
- 2-1 Genetic algorithm . . . 11
  - 2-1-1 Selection . . . 12
  - 2-1-2 Crossover . . . 12
  - 2-1-3 Mutation . . . 12
  - 2-1-4 Keep best layouts . . . 12
  - 2-1-5 Implementation . . . 13
- 2-2 Geometric yaw relationship . . . 14
  - 2-2-1 Data generation . . . 14
  - 2-2-2 Data fitting . . . 15
  - 2-2-3 Comparison . . . 16
- 2-3 Objective function . . . 18
  - 2-3-1 Annual energy production (AEP) . . . 18
  - 2-3-2 Levelized cost of electricity (LCOE) . . . 18

2-3-3 Robust objective . . . 19
2-3-4 Ecology . . . 20
2-4 Case design . . . 21
2-4-1 Wind rose . . . 21
2-4-2 Test case . . . 22
2-4-3 IJmuiden Ver case . . . 23

3 Results **25**
3-1 Test case . . . 25
3-2 Case study IJmuiden Ver: AEP and LCOE . . . 27
3-3 Case study IJmuiden Ver: Robustness . . . 30
3-4 Case study IJmuiden Ver: Ecology . . . 32
3-5 Case study IJmuiden Ver: Penalty . . . 34

4 Conclusion **35**

5 Discussion and future work **37**

A Python code **39**

B Density of offshore wind farms **40**

C Wind speeds **41**

D Layouts of recently built wind farms **42**

E Quantitative description of the box plots **43**

Bibliography **47**

Glossary **53**
List of Acronyms . . . 53

Chapter 1

Introduction

Many countries globally, including all member states of the European Union, have made net-zero pledges to combat climate change [1, 2]. This means that greenhouse gas emissions have to be reduced drastically in the coming two to three decades [2, 3]. Since the energy sector is one of the main contributors to greenhouse gas emissions, the energy sector will have to shift from fossil fuels to renewable energy sources [1–3]. In recent years the costs of solar and wind energy have decreased enormously to a level where they have become competitive with fossil fuel costs [3,4]. Consequently, solar and wind energy can potentially make a major contribution to the reduction of greenhouse gas emissions globally [1–3, 5]. When looking at increasing the capacity of offshore wind energy, one of the challenges is getting the most out of suitable areas, i.e. shallow waters with a lot of wind that are not used by other parties [6, 7]. Putting wind turbines on a given wind farm closer together reduces the costs and land use [7–10]. This does however introduce negative influences that individual turbines have on each other's energy production, called the wake effect.

1-1 Wake effect

When a wind turbine extracts energy from the wind, a wake occurs. A wake is the area behind a wind turbine where wind speeds are reduced and turbulence is increased. The wind on the edge of the wake will mix with the free-stream wind. Because of this wake mixing, the wake expands and wind speeds within the wake increase when moving downstream, as shown in Figure 1-1. Eventually, the wind conditions gradually return to that of the free-stream wind [8, 10, 11].

Currently, most wind turbines are built in grid formation [11, 12] and controlled to optimize their individual performance [8,10,13]. In this situation, the wake effect induces a substantial loss of energy production, ranging from 10% up to 30% [5, 8, 9, 13–17]. To reduce this loss, individual turbines can be controlled while taking the wind farm as a whole into account, called wind farm control. Additionally, the layout can be optimized such that the wake effect has less of a negative impact on the energy production [8, 9, 18].

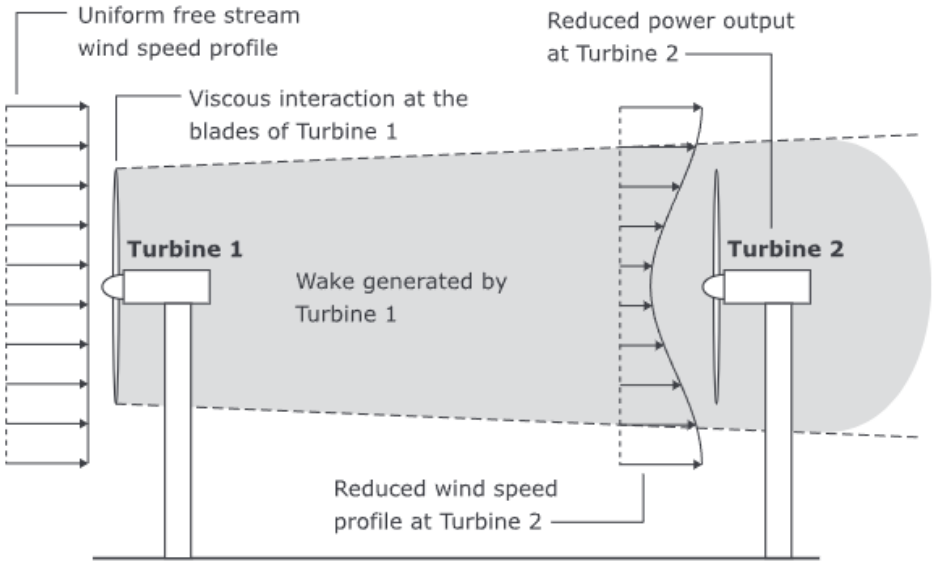

Figure 1-1: The upstream turbine (1) extracts energy from the wind, creating a wake behind it. As a result, the downstream turbine (2) experiences reduced wind speeds and increased turbulence. Because of wake mixing, the wake gradually expands and returns to free-stream conditions. (From [9])

1-2 Wind farm control methods

To minimize the adverse influence of the wake effect, different wind farm control methods have been developed. These control methods make use of adjustments to the generator torque [8] and the pitch, yaw, and tilt angles of the turbine. The amount of mechanical power converted into electricity is controlled by the generator torque. Pitch is the angle of the blades with respect to the incoming wind, yaw is the horizontal angle of the nacelle and tilt is the vertical angle of the nacelle, as can be seen in Figure 1-2. Tilt adjustment is not yet available in practice [8, 9].

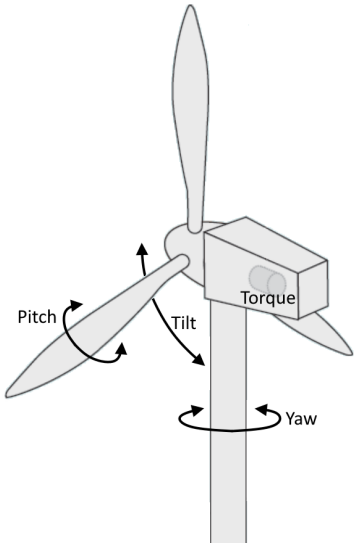

Figure 1-2: Control methods for wind turbines can make use of adjustments to the generator torque and pitch, yaw, and tilt angles. Pitch is the angle of the blades with respect to the incoming wind, yaw is the horizontal angle of the nacelle and tilt is the vertical angle of the nacelle. (From [19])

In wind farm control there are three general control methods: power de-rating control, wake steering, and wake mixing [8–10, 18, 20].

Power de-rating, also known as static induction control, uses the pitch angle or generator torque to reduce a turbine's power capture and therefore its wake, as shown in Figure 1-3. In theory, this can result in such an increase in power capture from the downstream turbine that it exceeds the loss at the upstream turbine, resulting in a net increase in power output [8–10]. Power de-rating control might look promising with low-fidelity models, but the possible power output gains turn out to be a lot smaller with higher-fidelity models, wind tunnel tests or field tests [8–10, 18, 20]. Even negative efficiency rates have been reported [9]. However, power de-rating has proven capable of mitigating the loads that the turbine experiences, extending the lifetime of a turbine [9, 18].

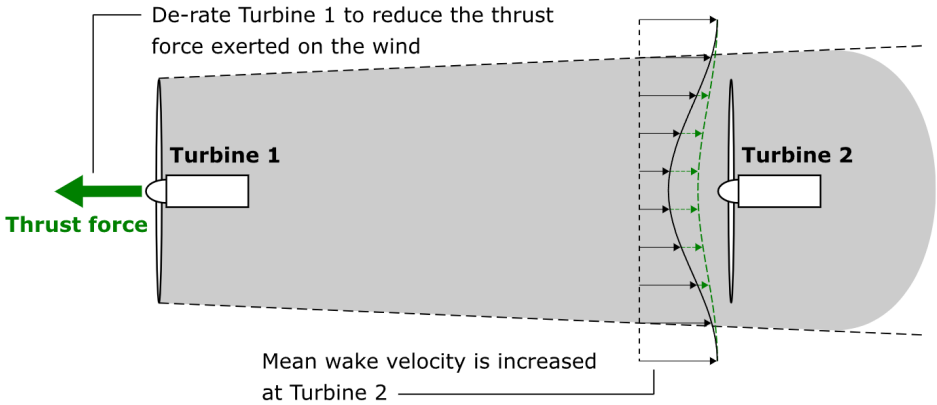

Figure 1-3: Power de-rating is applied to the upstream turbine (1), this reduces the upstream turbine's power capture and therefore its wake. This in turn increases the power capture of the downstream turbine (2). (Adjusted from [9])

Wake steering, also known as wake redirection control, is done by changing the yaw or tilt angle [8]. Tilt-based wake steering is promising but not yet available in practice [8, 9]. Yaw-based wake steering on the other hand gives good results in low- to high-fidelity models as well as in wind tunnel and field tests [9, 18]. Current practice for yaw control is minimizing the yaw angle, thus pointing the wind turbine towards the incoming wind [6, 15]. This strategy, called greedy control, optimizes the individual performance of the turbines without taking its downwind neighbour into account. It has been shown that greedy control of wind turbines leads to suboptimal performance of the wind farm [7, 8, 10, 13]. With a small yaw angle, the wake can be steered away from a downstream turbine, as illustrated in Figure 1-4. Although this does decrease the power generation of the individual turbine, it can increase the power generation of the farm as a whole [10, 15, 21]. This has not only been researched extensively [9], but has also been shown in recent field experiments [15, 22]. Admittedly, the increase in the annual energy production (AEP) is still small, 0.3% in a six-turbine wind farm in Canada, however, yaw control is very promising for specific wind conditions, up to 47% increase in AEP for low wind speeds from certain directions [15].

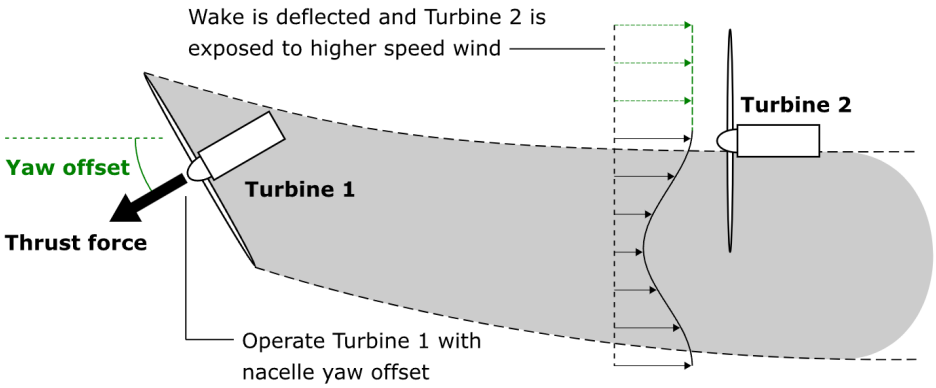

Figure 1-4: The wake of the upstream turbine (1) is steered away from the downstream turbine (2) by using yaw-based wake steering. This in turn increases the power capture of the downstream turbine. (Adjusted from [9])

Wake mixing, also known as dynamic induction control, uses the individual blade pitch angles of the turbine. By continuously changing these pitch angles, a more uneven distribution of reduced wind speeds is generated downstream of the turbine, which results in a faster wake recovery [9, 10, 18]. The pitch angles are currently used in such a way that the power capture is smoothed out and loads on the individual turbine are reduced with minimal power loss [8, 20, 23]. Recently, there has been an increased interest in using the pitch angles to maximize the power output with wake mixing techniques, but not much research has been done for this application [20]. In 2018, a dynamic collective pitch controller was introduced by Munters and Meyers, also known as the Pulse [24]. This control method was tested in a wind tunnel in 2022, showing a 2% to 4% power gain for a three-turbine setup [10]. In 2020 a dynamic individual pitch controller was introduced by Frederik et al., called the Helix approach, which has the potential to increase wake mixing significantly and with this the power production [20]. In Figure 1-5, the Pulse and the Helix methods are both illustrated. The first wind tunnel tests on the Helix approach were conducted in 2023 with positive results; an increase in power output of 15% for a two-turbine setup [18]. While wake mixing is promising, both of these methods were only tested on a small number of turbines for specific wind conditions and have not been tested in the field yet.

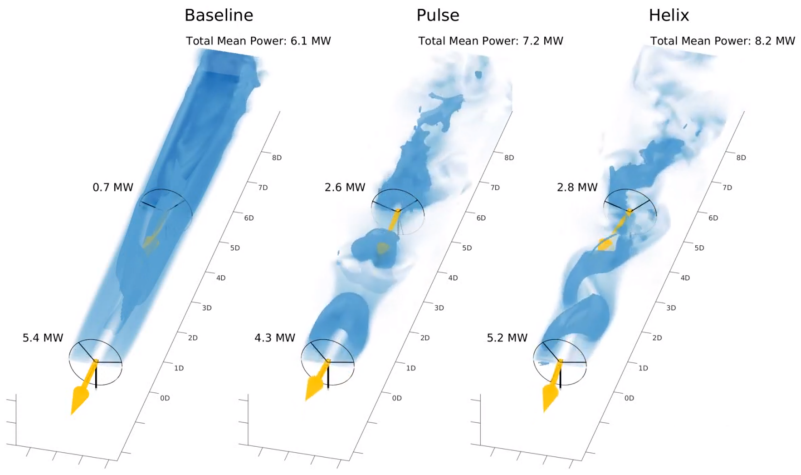

Figure 1-5: Wake mixing using the Pulse and the Helix technique compared to no pitch control. (Adjusted from [25])

1-3 Wind farm layout

The wake loss can also be mitigated by optimizing the wind farm layout. Traditionally most wind farms are built in a grid formation [11, 12]. However in recent wind farms, like Borssele in the Netherlands [26], the grid formation has been abandoned. This is because literature has shown that an optimized layout yields better results than the grid formation [12, 14, 27]. Recently research into the combination of layout and control optimization has gained more interest. Optimizing the wind farm layout and then using yaw control has been shown to give better results than just layout optimization or yaw control [28–31]. An example of possible layout and control improvements to the existing Princess Amalia Wind Park is shown in Figure 1-6.

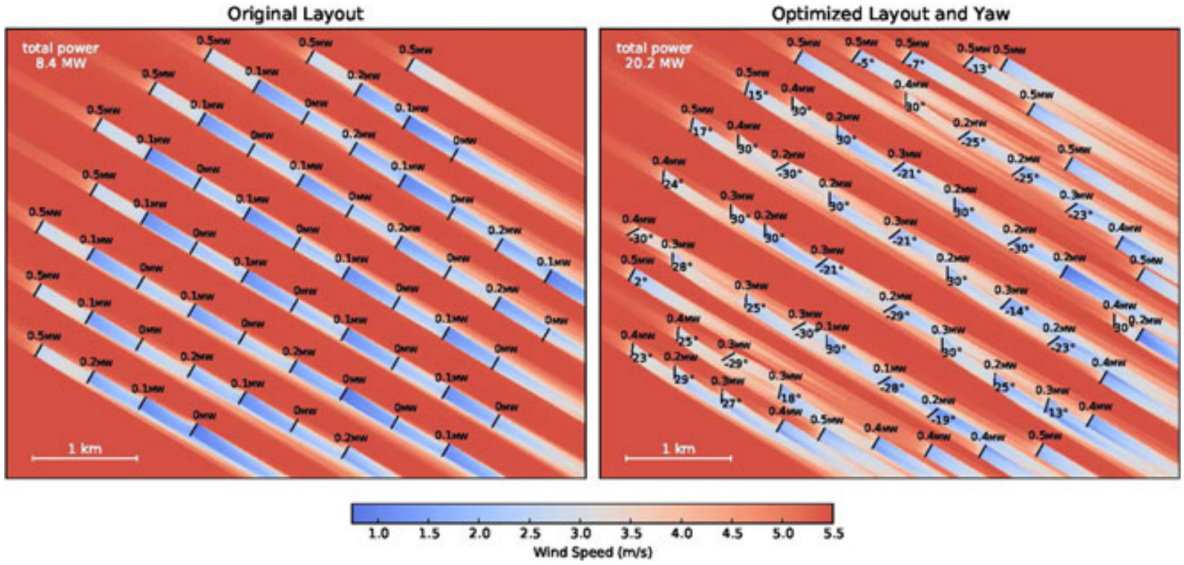

Figure 1-6: On the left the actual layout of the Princess Amalia Wind Park with currently used greedy control. On the right a joint optimized layout with yaw control, increasing the AEP with 5.3%. (From [14])

There has been research that looked into the differences between sequential and joint optimization of layout and control as well [14, 27–32]. With sequential optimization, the layout and control optimizations are treated as two separate problems. The layout and yaw angles are optimized one after the other, making control a retrofit solution. With joint optimization, the layout of the wind farm is optimized taking control into account. Although combining the two optimizations into one can lead to better solutions, it makes the problem more complicated as well. Most studies compared in Table 1-1 find a significant gain when optimizing in a joint way, although they differ considerably [28,29,31,32]. It also seems that when the power density (ρ) becomes higher in a wind farm, there is more to gain from joint optimization [28,31]. This is probably a result of the wake effect having a larger influence with a smaller turbine spacing as the wake effect is more present closer to the turbines [8]. Note that half of the studies used gradient-based methods [14, 29, 30, 32], which are not the most applicable methods in practice as will be discussed in section 1-5. Additionally, two studies adopted de-rating control [27,30]. However, as discussed in section 1-2 de-rating control is not very effective. Moreover, most studies had the AEP as an objective and applied a quite limited wind distribution [14, 28–32].

Table 1-1: The power gain found going from sequential to joint optimization in several studies. Shown are the year of publication, optimization method, control method, wake model, cost function, number of turbines (T), power density (ρ) in terms of MW/km^2, number of wind speeds (v), number of wind directions (α) and gain found.

Reference	Opti.	Control	Cost	T	ρ	v	α	Gain
Fleming et al. (2016) [29]	GB[a]	Yaw	MP[b]	60	21.7	1	72	0.55%
Wang et al. (2016) [27]	GA	De-rat	COE	45	6.8	105	24	-
Gebraad et al. (2017) [14]	GB	Yaw	AEP	60	21.7	20	72	0.091%
Pedersen et al. (2020) [30]	GB	De-rat	AEP	8	row[c]	1	1	0.0045%
Chen et al. (2022) [28]	PSO[d]	Yaw	AEP	16	24.8	1	36	4.3%
Chen et al. (2022) [28]	PSO[d]	Yaw	AEP	16	24.8	1	360	1.4%
Chen et al. (2022) [28]	PSO[d]	Yaw	AEP	80	8.2	1	36	0.95%
Chen et al. (2022) [28]	PSO[d]	Yaw	AEP	80	8.2	1	180	0.38%
Song et al. (2023) [31]	PSO[d]	Yaw	AEP	25	19.5	1	8	0.79%
Song et al. (2023) [31]	PSO[e]	Yaw	AEP	25	19.5	1	8	2.4%
Song et al. (2023) [31]	PSO[e]	Yaw	AEP	25	13.6	1	8	0.45%
Song et al. (2023) [31]	PSO[e]	Yaw	AEP	36	12.5	1	8	0.91%
Stanley et al. (2023) [32]	GB[a]	Yaw	AEP	16	20	1	72	0.8%

[a] Gradient-based
[b] Mean Power (W)
[c] A row of turbines with a spacing of 3.3D is considered
[d] Decomposition-based hybrid method (uses particle swarm optimization)
[e] Adaptive granularity learning distributed particle swarm optimization

The amount of research done on the difference between sequential and joint optimization of layout and control is too small to draw conclusions yet, but it appears that there might be some improvement in energy production possible when applying it to high-density wind farms [28, 29, 31, 32]. Lately putting turbines closer together has become more important[1]. This is because countries want to achieve their renewable energy targets and suitable space for offshore wind farms is limited [6, 33]. On top of that when wind farms are built with a higher density, more space can be left for nature. These ecological issues are playing an increasingly significant role in the construction of wind farms [6, 33, 34].

[1] See Figure B-1

1-4 Optimization problem

The most used objectives for wind farm optimization in literature are the annual energy production (AEP) and levelized cost of electricity (LCOE) [11, 12, 16, 35, 36]. AEP is a quite simple measure, but it is not a good objective from an application point of view [12, 36]. The costs of a wind farm are disregarded completely. A measure that takes this into account but is not too complex is the LCOE [36–38]. This is also the measure used in the wind farm industry [6, 38, 39]. Additionally, the fact that money in the present is worth more than money in the future can be considered in the form of a discount rate [12, 38]. To make the predicted economic gain more realistic, dynamic energy prices can be included. These vary widely and can even become negative when a lot of renewable energy is available at the same time [4, 6, 39]. This makes it hard to accurately predict these energy prices. Recently, the industry has started looking at non-economic factors in optimizing their wind farms as well, for example, taking the impact on the environment into account [4, 6].

A discrepancy between literature and practice is the domain constraints. In practice, the possible building area consists of several separate sub-areas that are not connected to each other as a result of cables, pipelines, maintenance zones, shipping lanes, depth and soil types [6, 12, 16, 39–41]. Due to factors like shipwrecks, plane wrecks, unexploded ordnances and archaeological findings, a lot of small parts are excluded from these building sub-areas [6, 12, 39–42]. What is left are multiple sub-areas with numerous gaps of possible building spots for turbines, making it non-convex [12, 43].

More constraints to take into account are the minimal spacing, the energy connection, nature conservation measures and the turbine types [6, 12, 16, 40, 41]. Wind turbines cannot be built too close together due to turbulence levels that are too high [6, 16]. Additionally, space is often needed in between the turbines for birds or ships to pass through [6, 40]. This is why a minimal turbine spacing is set by the government [40, 41]. The turbine types are constrained by the available technology, possible deals with the manufacturer and government requirements [6, 12, 40, 41]. Nature conservation measurements and the capacity and location of the energy connection are determined by the government too [6, 12, 40, 41].

1-5 Optimization method

The methods used to solve the wind farm layout optimization problem (WFLOP) can be split into two distinct groups: gradient-based methods and metaheuristics methods [11, 12, 14, 36]. While gradient-based optimization methods can give faster and better results [14, 44], this is only with non-realistic conditions. In practice, the gradient-based methods have difficulties with the non-convex domain and discontinuities of most wind farm models [12, 17, 31, 43, 44].

In literature, the most widely used metaheuristic method is the genetic algorithm (GA) followed by partical swarm optimization (PSO) [12, 35–37, 45, 46]. The popularity of the GA can be attributed to its easy implementation and capability of finding the global optimum, although this is not guaranteed [35, 37, 45]. Classic PSO is only suitable for continuous spaces, implementing it for a non-convex domain is not straightforward [47, 48]. One drawback of the GA is that since discretization is needed, there is limited freedom in placing the turbines [36, 49].

Employing a second heuristic to refine the layout has been shown to improve the solution significantly [7, 36, 49]. For instance, the combination of GA with PSO is faster than just PSO and gives a better solution than just GA [49]. In addition to GA and PSO, many other metaheuristics are possible. Multiple studies have shown that simulated annealing (SA) does not perform well for the WFLOP compared to GA and PSO [46–48]. Most other metaheuristics have not yet been thoroughly tested for the WFLOP, such as ant colony optimization, greedy heuristic and random search [12, 35, 37]. Recently machine learning techniques have started to show up as well to improve the computation time [5, 12].

When optimizing the wind farm in a joint way, the yaw angles need to be optimized as well. Traditionally, a nested or coupled optimization of the yaw angles and the wind farm layout is applied, which is very computationally expensive [32]. For example, a gradient-based method like sequential least squares programming (SLSQP) is often used to optimize the yaw angles [50]. In 2022, a much faster optimization method was proposed by Fleming *et al.* [50], called the Serial-Refine method. A year later, in 2023, the geometric yaw relationship was presented by Stanley *et al.* [32]. With this relationship, the optimal yaw angles can be computed directly, making coupled or nested optimization redundant.

1-6 Wake model

Taking yaw control and a realistic domain into account for the WFLOP, a computational-efficient wake model is needed. There are two kinds of wake models. On one side, there are models based on computational fluid dynamics methods, like SOWFA, SP-Wind and UTD-WF [8, 9, 16]. These are accurate but take up a lot of computing power. On the other side, there are analytical models, such as Jensen and Larsen [8, 9, 16]. These are less accurate, but a lot less computationally intensive. Since layout optimization works on large time scales and the wakes need to be computed many times, analytical wake models are a much better fit.

In literature, by far the most used wake model for the WFLOP is Jensen [16, 35, 37], although Gaussian-based wake models have caught some interest as well [45, 51]. These analytical models are assembled in open-source software like the Python packages FLORIS and PyWake [13, 51–53]. A study was done by Yang *et al.* [51] comparing the different wake models in FLORIS for the WFLOP without yaw control. The Curl model was way too computationally intensive for layout optimization. The Gaussian-Curl hybrid model did not perform significantly better than the Gaussian model but did take considerably longer. The multi-zone model was hard to tune, resulting in disappointing layouts. The Jensen model resulted in slightly better layouts compared to the Gaussian model and is around three times as fast. However, the Gaussian model is more accurate, especially in capturing the velocity deficit in the wake [51]. Moreover, the difference in run time will become smaller when accounting for wake deflection [51].

1-7 Contribution

The main contribution of this thesis is threefold. Firstly, a method is created for the joint optimization of a wind farm layout and yaw control that is suitable for real-world application. A non-gradient-based optimization method capable of handling realistic domains, the GA, and a time-efficient way of computing the optimal yaw angles, the geometric yaw relationship, are employed. Secondly, a variation of the geometric yaw relationship, introduced by Stanley *et al.* [32], is designed and its performance is investigated. Thirdly, the effects of joint as opposed to sequential optimization are explored for a test case, comparing the results with literature findings, and explored for a case based on a real wind farm called IJmuiden Ver [34], looking at the objective value as well as the characteristics of the optimized layouts. For the IJmuiden Ver case not only the AEP, but also an objective function that incorporates the building and maintenance costs of a wind farm, the LCOE, is applied. Additionally, the influence of the chosen wind bins and the geometric yaw relationship is examined using a robust objective function. Finally, ecology is taken into account by increasing the distance to the nearby nature reserve, and the penalty for not using yaw control on a joint optimized layout is investigated.

Chapter 2

Methods

In this chapter, the implemented method[1] for the joint optimization of the wind farm layout and yaw control is described. Firstly, the precise workings of the employed optimization method, the genetic algorithm (GA), are discussed in section 2-1. Secondly, a variation of the geometric yaw relationship is introduced in section 2-2. Thirdly, multiple objective functions are analysed in section 2-3. Finally, two cases are discussed in section 2-4 to study the effects of joint as opposed to sequential optimization.

2-1 Genetic algorithm

The classic GA can be described in a few steps [45,47,48], as illustrated in Figure 2-1. In the initialization, a population is generated and the fitness of each individual is calculated. The fitness can be any objective that is set, for example, the annual energy production (AEP) or levelized cost of electricity (LCOE). The first step of producing a new generation of the population is selecting the parents based on their fitness value. They are the individuals that will reproduce. The second step is making crossovers between two parents, combining the characteristics of the parents into a child. This child is added to the next generation if the child is viable, which means that it is a valid layout adhering to the minimum spacing and is not already in the next generation. The final step is some children undergoing mutations, again with viability checks in place. In this context, a mutation implies a turbine is moved to a random empty spot in the layout. When the new generation is complete, the fitness of each layout is computed. Afterwards, a new generation will be created again, as long as the termination condition is not met. In this case, the GA is terminated after a maximum amount of generations has been reached or if there has not been any improvement for a certain amount of generations.

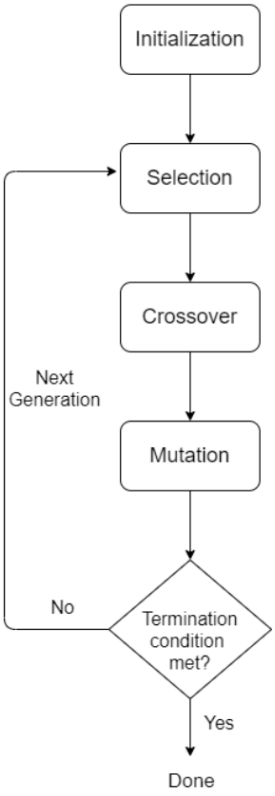

Figure 2-1: The steps of the GA [47].

[1] Code is available; see Appendix A.

2-1-1 Selection

There are many ways to select the parents; three common methods are tested. The first and most straightforward method is called elitist-random. With this method, the $x \cdot n_p$ individuals with the highest fitness values are selected to be parents, where $x \in [0, 1]$ and n_p is defined as the desired number of parents. Then $(1 - x) \cdot n_p$ individuals are randomly selected from the rest of the population. The second method, rank, works with probabilities. All the individuals are ordered from highest to lowest fitness value. Subsequently, a probability is appointed to each individual, where individuals with a higher fitness value are assigned a higher probability. This probability declines linearly to zero and sums up to one. Then, n_p individuals are chosen with these probabilities. The third method is the tournament selection method. All the individuals in the population are randomly divided over n_p tournaments. The winner of each tournament, the individual with the highest fitness value, is added to the parent population.

2-1-2 Crossover

In the crossover step, the characteristics of two parents are combined into a new individual, a child. The two parents are selected at random from the parent population. If both parents have a turbine at a certain position, the child also gets a turbine at that position. The rest of the turbines are chosen at random from the positions where only one of the parents has a turbine. If the child is viable it is added to the next generation.

2-1-3 Mutation

The classic GA can be improved by making informed decisions on the quantity and place mutations are most effective. One way to do this is to change the position of the turbines with the smallest power production. After the parents are selected, the worst turbine, or the set of worst turbines if the number of turbines is big enough, is moved to a random empty spot. In literature, this method is called adaptive GA [54]. Another way of improving the classic GA is making the mutation rate dynamic. This is inspired by the cooling factor used in simulated annealing (SA), which is another metaheuristic method. With a dynamic mutation rate, the GA starts with a high number of mutations, which becomes smaller after a certain amount of generations. When a set minimum number of mutations is reached, the number of worst-turbines-to-move is decreased. Employing the adaptive GA and dynamic mutation rate improved the algorithm significantly.

2-1-4 Keep best layouts

The final explored improvement of the classic GA is duplicating the best layouts of the previous to the next generation without modifications. The best layouts are the layouts with the highest fitness value. A certain amount of layouts, n, is passed on to the next generation without crossovers or mutations. Just before creating a new generation, the worst n layouts of the coming generation are replaced with the n best layouts of the last generation. Doing this ensures that every generation contains a better result than the previous generation. Keeping too many high-performing layouts, however, can result in the algorithm getting stuck in a local optimum due to a lack of diversity.

2-1-5 Implementation

For the population size and initialization of the GA, a trade-off must be made between the required computational power and the performance of the GA. The population size influences the duration and outcome of the GA greatly. A larger population size means more diversity. Therefore, a larger population makes it less likely for the algorithm to get stuck in local optima. On the other hand, a larger population size also means that more computational power is needed. The chosen population size of 100 is a compromise between diversity and the computational power available. For the initialization of the GA, a coarser grid than during the GA is used where each position already adheres to the minimal spacing. This makes generating a random initial generation much faster. However, when the power density in a wind farm is close to its limit, then there is a possibility that not enough unique layouts can be generated in this way.

A tuning process is completed to set the exact parameters in the selection step of the GA. All the parameters of the GA are chosen based on the IJmuiden Ver case described in section 2-4 but with a smaller number of turbines. To make tuning possible, a 16-turbine setup is chosen and the is scaled down accordingly. All three selection methods are tested with varying parent rates, from 0.3 to 0.7, and varying mutation rates, from 0.2 to 0.4. The number of parents is equal to the parent rate multiplied by the population size. The number of mutations is equal to the mutation rate multiplied by the population size. For the elitist-random method, the ratio of high-fitness individuals is also varied, from $\frac{1}{3}$ up to $\frac{4}{5}$. After multiple runs for each variation, the ratio for the elitist-random method that is most effective turned out to be $\frac{1}{2}$. The best-performing selection method throughout the varying rates, of the three methods considered in this thesis, is the tournament method. A parent ratio of 0.6 and mutation rate of 0.3 in combination with the tournament method gives the overall best results.

In the mutation step, several individuals of the population are modified. A set number of individuals are selected at random. Within each selected individual, one turbine is moved to a random empty spot. If the mutated individual is viable, then the individual will be replaced by the mutated version. The static mutation version was compared with the dynamic mutation version. The dynamic-mutation-step parameter indicates after how many generations the mutation rate is lowered with 0.1. When the mutation rate is already 0.1, then the number of worst-turbines-to-move is lowered by 1 with a minimum of 0. An initial mutation rate of 0.4 with an initial worst-turbines-to-move rate of 0.02 and a dynamic mutation step of 50 works best.

The algorithm is terminated when the maximum number of generations, 200, is reached or there has been no improvement for 40 generations. These numbers are chosen based on the results during the tuning process. Higher-density wind farms or wind farms with a large number of turbines need more generations until there is no significant improvement anymore.

2-2 Geometric yaw relationship

In 2023, a new method to optimize the yaw angles of a wind farm was presented by Stanley *et al.* [32], called the geometric yaw relationship. This method builds on the theory that the optimal yaw angle of a turbine is mostly dependent on the relative position of the closest turbine in its wake. This allows a near optimal yaw angle be computed directly, making coupled or nested optimization redundant. The relative position of a turbine to the closest turbine in its wake is expressed in dx and dy, where dx is the distance in rotor diameters in the wind direction and dy is the distance in rotor diameters perpendicular to the wind direction. By processing the data of the optimized yaw angles of over 100,000 turbines from a variety of wind plants with different numbers of turbines, layouts, turbine spacings, and wind speeds, Stanley *et al.* [32] derived the following geometric yaw relationship.

$$f(dx, dy) = \begin{cases} \text{sign}(dy) \cdot 30 \cdot \left(1 - \frac{dx}{25}\right) & \text{if } dx \geq 0 \text{ and } |dy| \leq 1 \\ 0 & \text{otherwise} \end{cases} \tag{2-1}$$

To investigate if the geometric yaw relationship can be further improved, a variation is proposed. First, a new data set is generated. Next, the data is fitted to a different function. Last, a comparison is made between the two versions of the geometric yaw relationship.

2-2-1 Data generation

A new set of optimal yaw angles is computed for approximately 40,000 turbines from a variety of wind plants with different numbers of turbines and turbine spacings. To be more specific, the number of turbines is chosen from $\{5, 10, 20, 30, \ldots, 100\}$ and the turbine spacings from $\{7, 8, 9, 10, 15, 20, 30\}$. With these characteristics, a random layout is generated 10 times for each combination with a minimum spacing of $4D$ instead of $2D$. A 15MW turbine is used instead of a 5MW turbine, the wind speed is taken as 8 m/s, and the Gaussian model is chosen instead of the Cumulative Curl model as a velocity model. The Serial-Refine method in FLORIS is applied to optimize the yaw angles. This data is visualized in Figure 2-2.

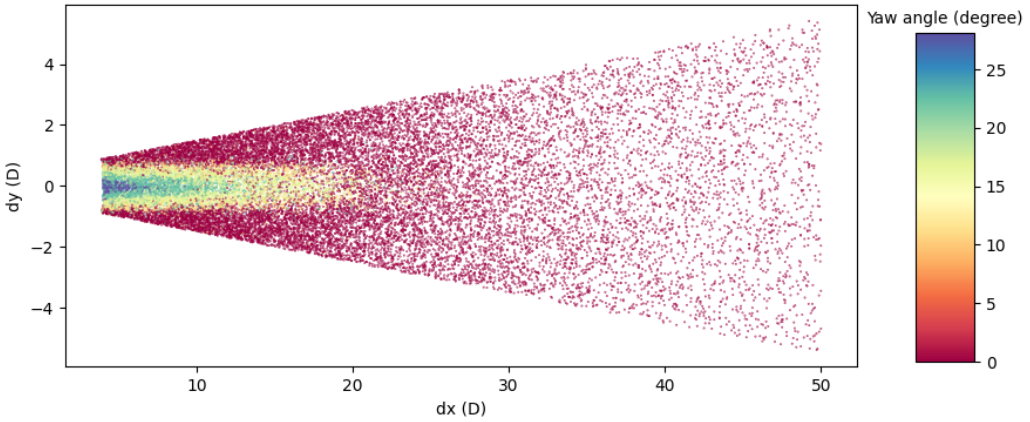

Figure 2-2: The absolute value of a turbine's optimized yaw angle (°) relative to the position of the closest turbine in its wake. This position is expressed in terms of the number of rotor diameters in the wind direction (dx) and the direction perpendicular to the wind direction (dy).

2-2-2 Data fitting

The data is first preprocessed to improve the performance of the new geometric yaw relationship. Only the dx smaller than $28D$ and the dy smaller than $2D$ are taken. This is because, after a dx of $28D$ or a dy of $2D$, no yaw is applied anymore. Furthermore, for dy values close to zero the optimal yaw angle can make a large jump from positive to negative, yet the power production of the turbine is similar. Therefore, the absolute value of the yaw is taken for the data-fitting process. After analysing the data the following function is proposed

$$f(dx, dy) = c_1 \cdot \left(1 - \left[\frac{dx}{c_2}\right]_{\in[0,1]}^{c_3}\right) \cdot \left(1 - \left[\frac{|dy|}{c_4}\right]_{\in[0,1]}^{c_5}\right) \tag{2-2}$$

where c_1, c_2, c_3, c_4 and c_5 are all positive and

$$[x]_{\in[0,1]} = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x > 1 \\ x & \text{otherwise} \end{cases} \tag{2-3}$$

The parameters of the proposed function are fitted to the processed data[2]. c_1, c_2 and c_4 are very close to the values proposed by Stanley *et al.* [32] and therefore taken as such.

$$c_1 = 30 \qquad c_2 = 25 \qquad c_4 = 1 \tag{2-4}$$

After these values are fixed, c_3 and c_5 are fitted again to the data. That gives in the following values.

$$c_3 = 1.336 \qquad c_5 = 1.407 \tag{2-5}$$

When applying the geometric yaw relationship, the yaw is considered positive when dx and dy are positive, negative when dx is positive and dy is negative, and zero when dx is negative. Thus the new geometric yaw relationship becomes

$$\text{yaw} = \begin{cases} f(dx, dy) & \text{if } 0 < dx \text{ and } dy \geq 0 \\ -f(dx, dy) & \text{if } 0 < dx \text{ and } dy < 0 \\ 0 & \text{otherwise} \end{cases} \tag{2-6}$$

where

$$f(dx, dy) = 30 \cdot \left(1 - \left[\frac{1}{25}dx\right]_{\in[0,1]}^{1.336}\right) \cdot \left(1 - [|dy|]_{\in[0,1]}^{1.407}\right) \tag{2-7}$$

[2]Using the `optimize.curve_fit()` function in SciPy

2-2-3 Comparison

As seen in Figure 2-3, there is a clear distinction between the geometric yaw relationship presented by Stanley *et al.* [32], geometric yaw Stanley, and the geometric yaw relationship presented in this thesis, geometric yaw De Jong. Although the (dx, dy) coordinates where a nonzero yaw angle is applied stay identical, the yaw angles themselves differ considerably at these points.

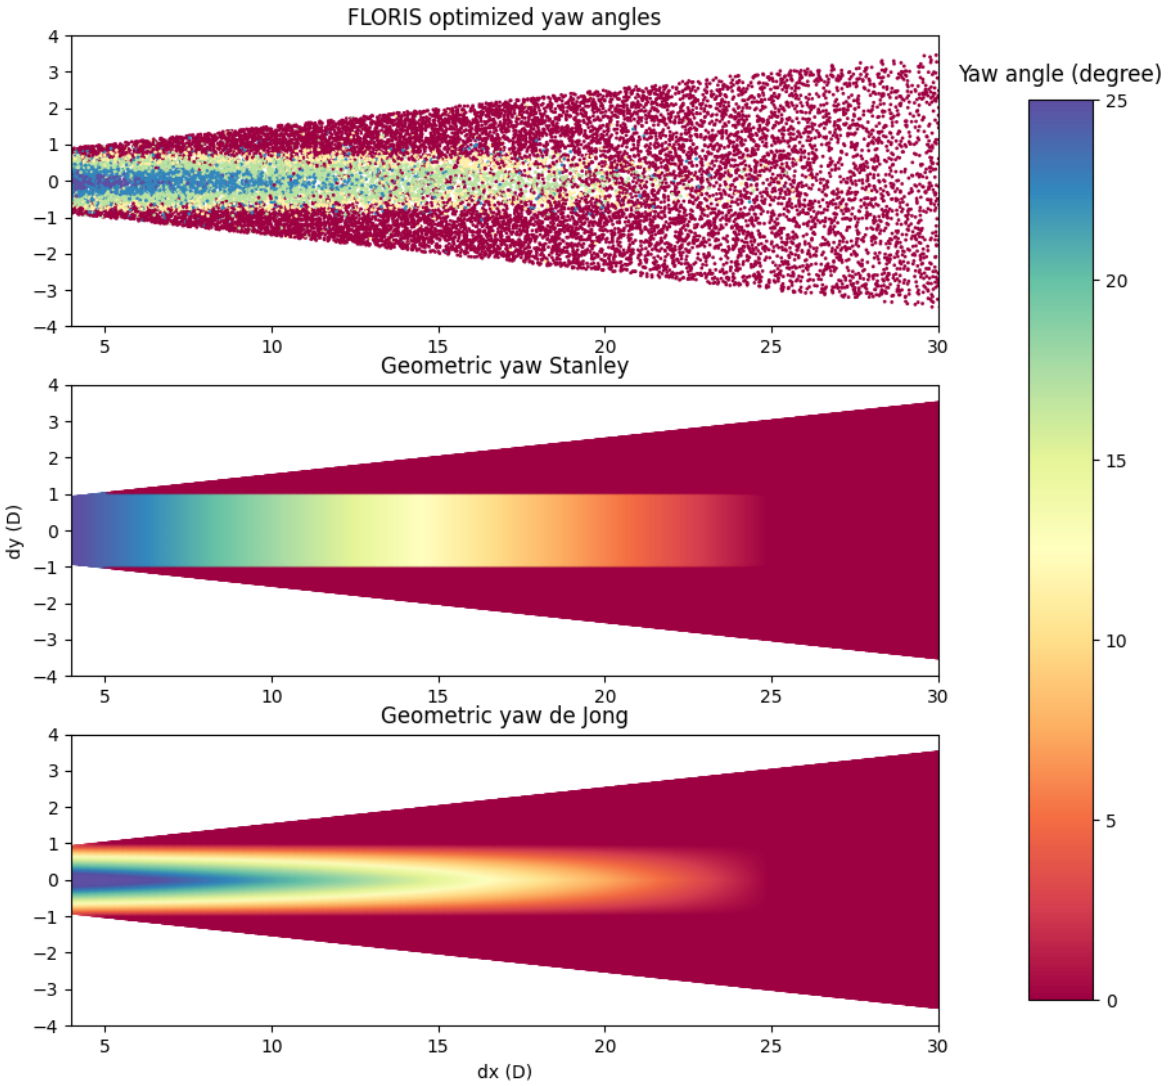

Figure 2-3: A comparison between the data set of FLORIS optimized yaw angles, the geometric yaw relationship presented by Stanley *et al.* [32], geometric yaw Stanley, and the geometric yaw relationship presented in this thesis, geometric yaw De Jong.

To investigate the performance of the geometric yaw relationship presented in this thesis, the different yaw optimizers are applied on 20 sequential optimized layouts of the IJmuiden Ver case presented in section 2-4. The two geometric yaw relationships use similar amounts of computational power. The variation of the geometric yaw relationship proposed in this thesis, GY De Jong, outperforms the method proposed by Stanley *et al.* [32], GY Standley, see Figure 2-4.

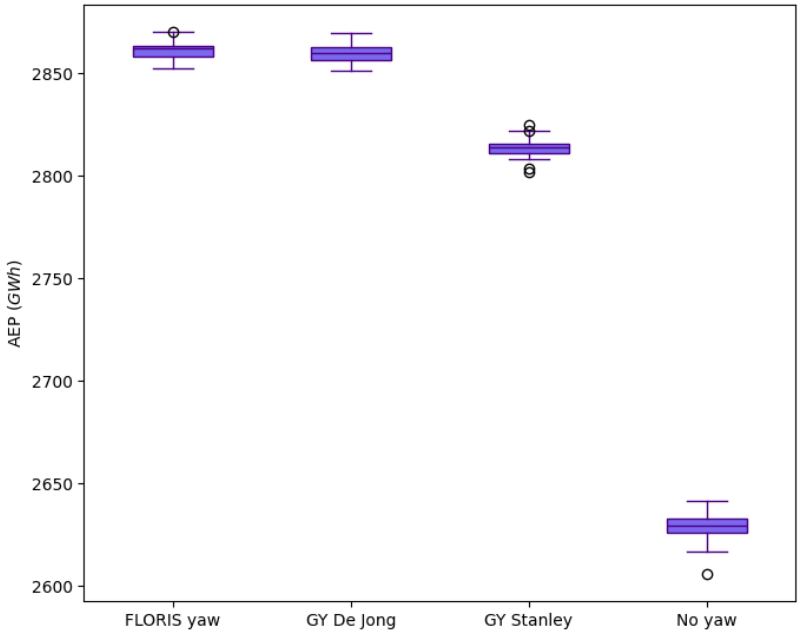

Figure 2-4: A comparison between the AEP computed with the FLORIS optimized yaw angles (FLORIS yaw), the geometric yaw relationship presented in this thesis (GY De Jong), the geometric yaw relationship presented by Stanley *et al.* [32] (GY Stanley), and all yaw angles taken as zero (No yaw). For the precise median and variance values see Table E-1.

2-3 Objective function

The most used objective function for the wind farm layout optimization problem (WFLOP) is the AEP. One can build upon this by taking the costs into account as well, computing the LCOE. The LCOE is also used in practice [6]. To examine the robustness of the found solution, a robust objective is introduced. Ecology is taken into account too, because of its growing importance.

2-3-1 Annual energy production (AEP)

To compute the AEP, the power production of each wind turbine in each wind condition is calculated. A wind condition is a combination of a wind speed (ws) and a wind direction (wd). The power production of each turbine is computed[3] with the IEA 15MW turbine and the Gaussian model as the velocity model. The average power production throughout the year ($\overline{\mathrm{P}}$) is calculated by taking the weighted sum of the turbine powers ($P(T_i)$), where the weights are the corresponding frequencies for each wind condition (f_{ws}^{wd}). Multiplying the average power production with the number of hours in a year (h_{year}) gives the AEP.

$$\begin{aligned} \overline{\mathrm{P}} &= \sum_{ws,wd,i} P(T_i) \cdot f_{ws}^{wd} \\ \mathrm{AEP} &= h_{\mathrm{year}} \cdot \overline{\mathrm{P}} \end{aligned} \tag{2-8}$$

2-3-2 Levelized cost of electricity (LCOE)

With the LCOE, in addition to the energy production, the costs of building, maintaining and decommissioning a wind farm are taken into account. Moreover, the point in time of spending and earning is relevant for the LCOE. Therefore, the costs and energy production are corrected with a discount rate (r) each year. The corrected cost ($C_{\mathrm{discounted}}$) is then divided by the corrected energy production ($E_{\mathrm{discounted}}$) to get the LCOE.

Since the conditions of the wind farm are assumed to be constant, the energy production each year (E_t) stays the same. This is computed in the same way as the AEP, see Equation 2-8. The cost in year t is the sum of the investment cost (I_t) and the maintenance cost (M_t). The maintenance cost per year of the wind farm is constant within its lifetime (t_l). The investment cost consists of the building cost, evenly spread out over the building time (t_b), and the decommission cost, evenly spread out over the decommission time (t_d).

$$C_{\mathrm{discounted}} = \sum_{t=0}^{t_b+t_l+t_d} \frac{I_t + M_t}{(1+r)^t} \tag{2-9}$$

$$E_{\mathrm{discounted}} = \sum_{t=0}^{t_b+t_l+t_d} \frac{E_t}{(1+r)^t} \tag{2-10}$$

$$\mathrm{LCOE} = \frac{C_{\mathrm{discounted}}}{E_{\mathrm{discounted}}} \tag{2-11}$$

[3] Using the `get_turbine_powers()` function in FLORIS

With a set number of turbines, most costs are fixed (C_{fixed}). However, the water depth (z) influences the turbine building cost (C_{turbine}) quite heavily. In addition, the distance to the substation ($d_{\text{substation}}$) influences the cable building cost (C_{cable}).

$$C_{\text{building}} = \sum_{i=0}^{N} C_{\text{fixed}} + C_{\text{turbine}} \cdot z(T_i) + C_{\text{cable}} \cdot d_{\text{substation}}(T_i) \tag{2-12}$$

$$I_t = \begin{cases} \frac{C_{\text{building}}}{t_b} & \text{for } t = 0, 1, \ldots, t_b \\ 0 & \text{for } t = t_b + 1, \ldots, t_b + t_l \\ \frac{C_{\text{decommission}}}{t_d} & \text{for } t = t_b + t_l + 1, \ldots, t_b + t_l + t_d \end{cases} \tag{2-13}$$

The variables that are used for calculating the LCOE based on the 15MW turbine are shown in Table 2-1. Most variables are based on recent UK[4] offshore wind projects [38]. The turbine and cable building costs of the 15MW turbine are selected in consultation with industry experts.

Table 2-1: Variables of the LCOE based on the 15MW turbine

	Unit	Symbol	Value
discount rate		r	0.06
lifetime	year	t_l	27
building time	year	t_b	5
decommission time	year	t_d	1
fixed building cost	€/turbine	C_{fixed}	$37,656,750$
turbine building cost	€/$m_{\text{water depth}}$/turbine	C_{turbine}	$100,000$
cable building cost	€/$m_{\text{distance to substation}}$/turbine	C_{cable}	25
decommission cost	€/turbine	$C_{\text{decommission}}$	$5,692,500$
maintenance cost	€/year/turbine	M_t	$1,293,750$

2-3-3 Robust objective

A robust objective function is introduced to investigate if the algorithm still gives good results when the yaw optimization method or wind conditions are adjusted. In the GA, the geometric yaw relationship is applied on a simplified wind rose where the number of wind bins is considerably reduced. The objective function is required to be time efficient because the objective function is computed around 20,000 times during the GA. The robust objective optimizes the yaw angles with the Serial-Refine method and uses the complete wind rose. Using the FLORIS to optimize the yaw angles takes about 700 times as long as the geometric yaw relationship for a 67-turbine wind farm. On top of that the optimal yaw angles and turbine powers need to be calculated per wind bin. The complete wind rose consists of around $8,000$ wind bins, while the simplified wind rose contains only 72 wind bins. This makes the robust objective function far more computationally intensive. Because of the computational intensity, the robust objective is only applied on the optimized end layout and not during the optimization process.

[4]Pounds are converted to euro by multiplying the amount with 1.15

2-3-4 Ecology

The role of ecology in building a wind farm has increased significantly. This is reflected in the point system of recent wind farm tenders [34, 55]. One approach to build in a nature-conscious way is to increase the distance to neighbouring nature reserves. Since the gain of this approach is not that straightforward, the distance to neighbouring nature reserves is treated as a constraint. To investigate how this constraint influences the solution, the layout is optimized for multiple distances and the resulting LCOE is examined.

2-4 Case design

The developed optimization method is applied on a real wind farm that is going to be built in a few years, called IJmuiden Ver [34]. The 15MW turbine is used, because it is the closest available reference turbine to the turbines that are going to be built [41]. The wind data taken is based on the location of the wind farm. Likewise, the domain that is used is based on IJmuiden Ver. For the test case however, a different domain is adopted to make it more comparable with other studies performed on the joint optimization of the wind farm layout and yaw control [14, 28, 29, 31, 32].

2-4-1 Wind rose

A wind rose is generated based on the data set taken from the WINS50 project [56]. The wind distribution of the most recent data, 2019, in the middle of the IJmuiden Ver lot is taken. The complete wind rose consists of 360 wind direction bins of 1° each and 29 wind speed bins of 1 m/s each, see Figure 2-5a. Note that only the wind speeds between 3 m/s and 25 m/s are actually used in the computations since the 15MW turbine does not produce any energy below a wind speed of 3 m/s, the cut-in speed, or above a wind speed of 25 m/s, the cut-out speed [57]. For computational reasons the wind rose used during the GA is scaled down to 72 wind directions and just one wind speed, as shown in Figure 2-5b. A wind speed of 8 m/s is chosen to optimize the layout because it is the most common speed[5], the mode. Keep in mind that the objective value computed during the GA is worse than if the objective value was calculated with the complete wind rose. This discrepancy arises as a result of the average wind speed being significantly higher, around 10 m/s opposed to 8 m/s. Thus the found AEP and LCOE can be compared for joint and sequential optimization, but are not necessarily realistic values.

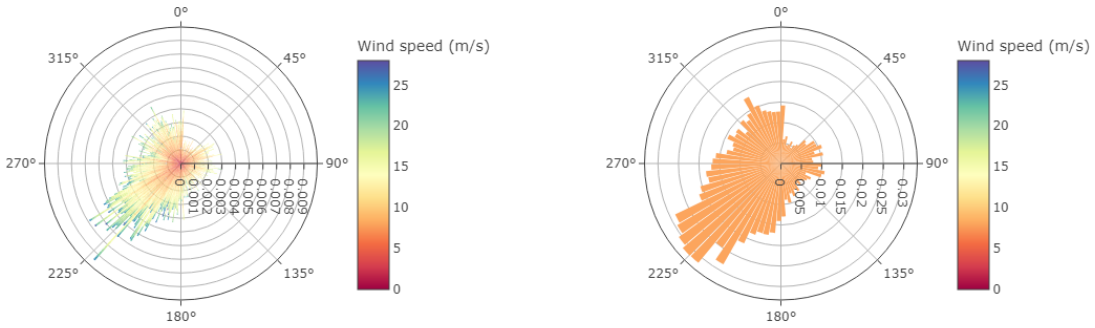

(a) Complete wind rose **(b)** Simplified wind rose

Figure 2-5: The complete wind rose consists of 360 wind direction bins and 29 wind speed bins. For computational reasons the simplified wind rose is used during the GA. This is the complete wind rose scaled down to 72 wind directions and just one wind speed of 8 m/s.

[5] See Figure C-1

2-4-2 Test case

For the test case, the AEP of a 16-turbine wind farm is studied for a variety of power densities in order to make the comparison with literature findings. A 15MW turbine is used and the minimum spacing is set to two rotor diameters[6] ($2D$). The domain is a uniform square, without any discontinuities, where the side lengths are adjusted to achieve the chosen power density. The power densities vary from 8 to 20 W/m^2. A few of the domains that are used are shown with random layouts in Figure 2-6.

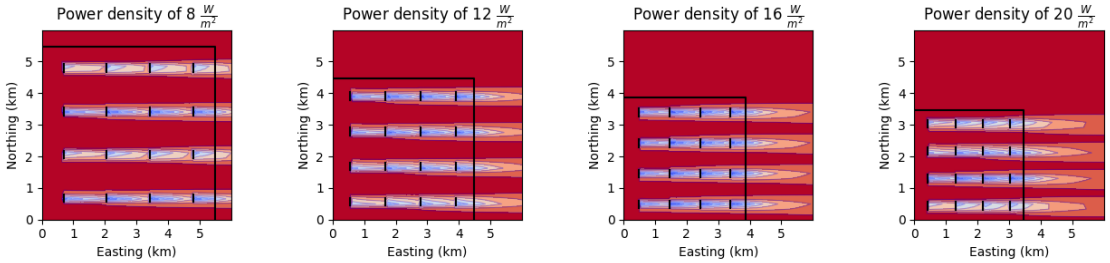

Figure 2-6: For the test case uniform square domains are used on which 16 turbines are placed. The sides of the domain are adjusted to achieve a certain power density. The power densities are chosen $8, 9, \ldots, 20$. Shown above are a few of the domains that are used with a grid formation as a layout. The boundaries of the domain are indicated by the black square.

[6]The 15MW turbine has a rotor diameter of 242.24 meter. ($2D = 484.48$ m)

2-4-3 IJmuiden Ver case

In the IJmuiden Ver case, the chosen parameters and domain are based on the regulation set by the Dutch government [34, 41]. Due to computational reasons, only the Alpha I plot is considered. On this plot, 67 turbines can be placed with a minimum spacing of four rotor diameters[7] ($4D$). The substation is built at 532359.5 meter easting and 5851358.4 meter northing. There is a nature reserve close by, called 'Bruine Bank'. The closest border to the 'Bruine Bank' is the line between $(520457, 5848100)$ and $(533881, 5843854)$. In Figure 2-7, the water depth in meters is shown of the area [58].

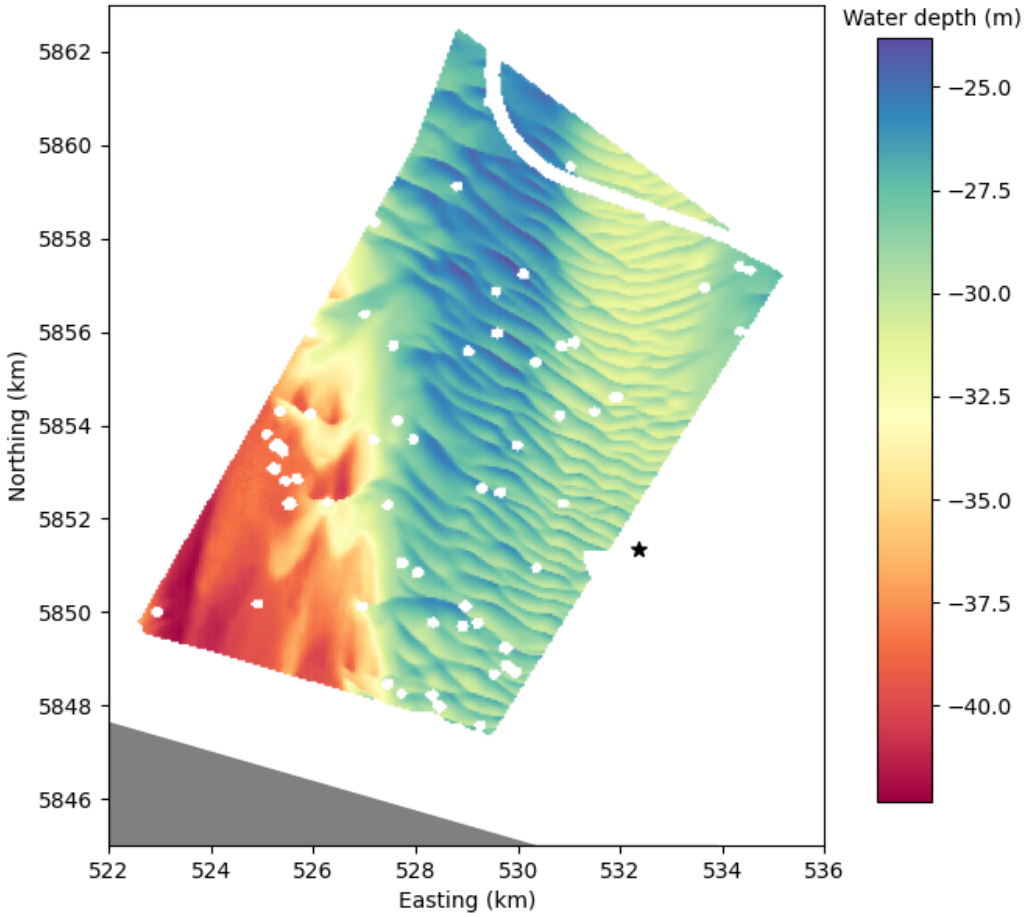

Figure 2-7: For the IJmuiden Ver case, actual real-world data is used. Wind turbines can not be built in the white areas. In the area where wind turbines can be built, the water depth is shown in meters. The substation is shown as a star and the nature reserve 'Bruine Bank' is illustrated in grey.

[7] The 15MW turbine has a rotor diameter of 242.24 meter. ($4D = 968.96$ m)

Chapter 3

Results

This chapter explores the effects of joint as opposed to sequential optimization. Firstly, a test case is discussed in section 3-1, comparing the results with literature findings. Secondly, the layouts of the IJmuiden Ver case are analysed, taking annual energy production (AEP) as well as levelized cost of electricity (LCOE) as objective, in section 3-2. Thirdly, the influence of the chosen wind bins and the geometric yaw relationship is examined in section 3-3. Fourthly, ecology is taken into account by increasing the distance to the nearby nature reserve in section 3-4. Finally, the penalty for not using yaw control on a joint optimized layout is investigated in section 3-5.

3-1 Test case

The goal of the test case is to compare the results with literature findings, see Table 1-1, and illustrate the general workings of the developed genetic algorithm (GA). In the test case, the GA is run 100 times in a joint and 100 times in a sequential way for each power density between 8 and 20. The initial layouts for a given domain are identical for the joint and sequential case. With the optimized layout, the AEP is computed when applying the geometric yaw relationship proposed in section 2-2. The results are shown in Figure 3-1 as a box plot. Each bar of the box plot represents the distribution of the AEP of 100 optimized layouts. Half of the data falls within the limits of the box, with a central line marking the median value. The lines extending from the box indicate the range of the remaining data, except for the outliers shown as dots. The number displayed in each blue bar is the relative improvement of the median AEP of the sequential optimized layouts compared to the joint optimized layouts.

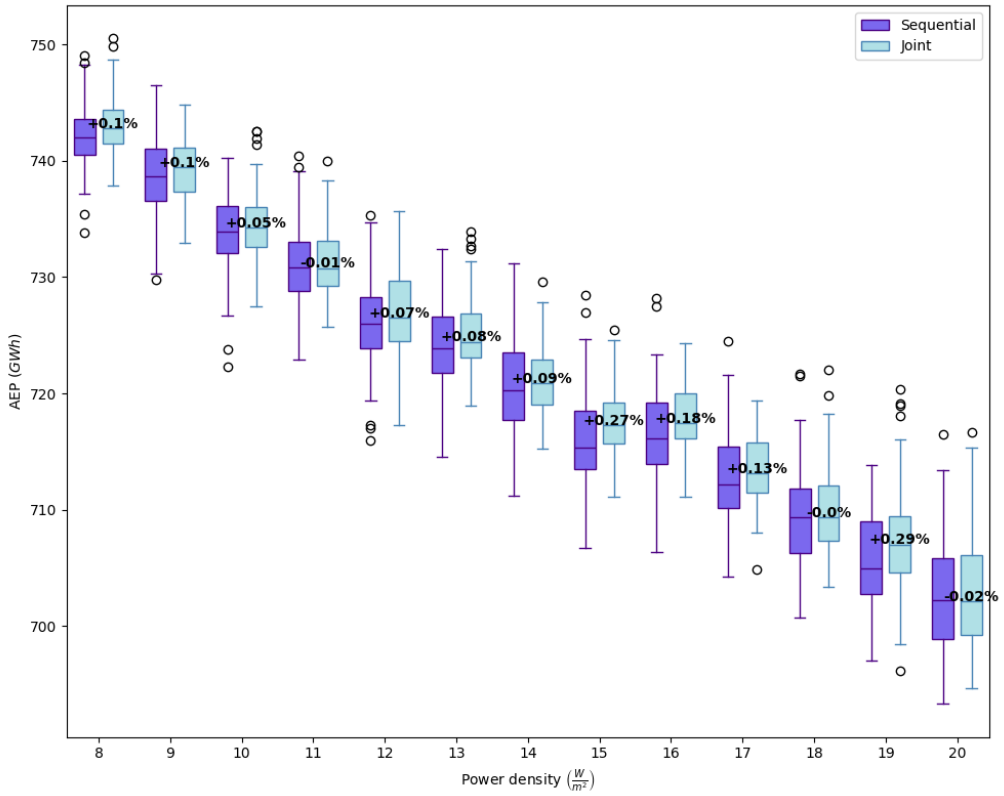

Figure 3-1: The test case is run 100 times in a joint and 100 times in a sequential way for each power density between 8 and 20. These data points are shown in the box plot. The joint optimization method is compared to the sequential optimization method, with the relative improvement of the median AEP annotated in the figure. For the precise median and variance values see Table E-2.

Figure 3-1 indicates that the improvement of the joint optimization compared to the sequential optimization depends substantially on the domain that is adopted when looking at wind farms with a low number of turbines. In some cases, there is an improvement of almost 0.3%, but in other cases, there is no improvement found at all. Furthermore, the variance is quite large, especially for the higher power densities. This could mean that the GA has difficulty finding a solution close to the optimum. In this specific test case, the number of turbines that the algorithm puts on the southern edge and the western edge of the domain is quite dominant in determining the AEP, since the dominant wind direction is the southwest. For the analysed power densities, most of the turbines fit on the southern edge and the western edge. This avoids wake interaction altogether and therefore reduces the need for wake steering.

The improvements of the joint optimization compared to the sequential optimization are generally significantly lower than those in literature findings, see Table 1-1. Most studies find an improvement between 0.55% and 4.3% for high-density wind farms with a power density of 19.5 or more [28, 29, 31, 32], except for a study that reported 0.091% [14]. For lower-density wind farms with a power density of 13.6 or less, the improvements vary from 0.38% up to 0.95% [28, 31]. However, the optimization methods used in literature assume continuous or even convex domains [14, 28, 29, 31, 32]. The method in this thesis is designed for handling realistic domains, consisting of multiple sub-areas with numerous gaps.

3-2 Case study IJmuiden Ver: AEP and LCOE

The IJmuiden Ver case is based on a real wind farm that will be built in a few years [34]. This means the domain is much more complicated than the domain of the test case, see Figure 2-7, and the number of turbines is significantly higher. For this case, not only the AEP is optimized, but the GA is run with the LCOE as objective as well. To compare the joint and sequential optimization of the AEP and the LCOE, the GA is run 50 times for each combination, see Figure 3-2.

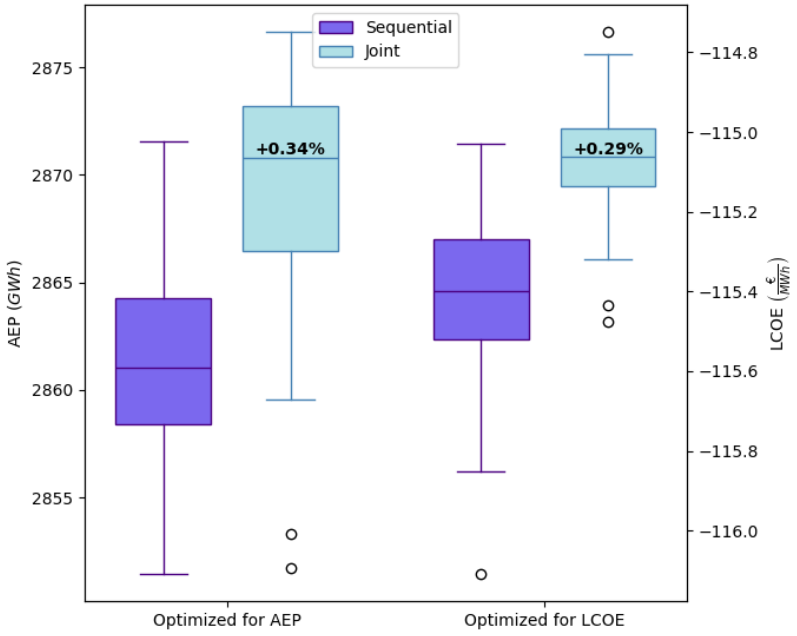

Figure 3-2: The IJmuiden Ver case is run 50 times in a joint and 50 times in a sequential way for both objectives, AEP and LCOE. These data points are shown in the box plot. The joint optimization method is compared to the sequential optimization method, with the relative improvement of the median objective value annotated in the figure. For the precise median and variance values see Table E-3.

The joint optimization method outperforms the sequential optimization method by 0.34%. This is a greater improvement than the improvement with any of the domains in the test case, see Figure 3-1. The influence of optimizing the LCOE as opposed to the AEP on the improvement is small. The results in Figure 3-2 show that the improvement of the LCOE of the joint compared to the sequential optimization, is 0.29%. However, the variance of the joint optimized LCOE is significantly smaller than the sequential optimized LCOE: 0.02 and 0.04 respectively.

To analyse the characteristics of the optimized layouts, a power density metric is applied. This metric visualises where the turbines are placed by plotting the installed power density in the domain. Figure 3-3 shows the average power density of 50 optimized layouts. This is done for the joint optimized AEP, the joint optimized LCOE, the sequential optimized AEP and the sequential optimized LCOE. The power density of 50 initial layouts, which are randomly generated, are shown as well.

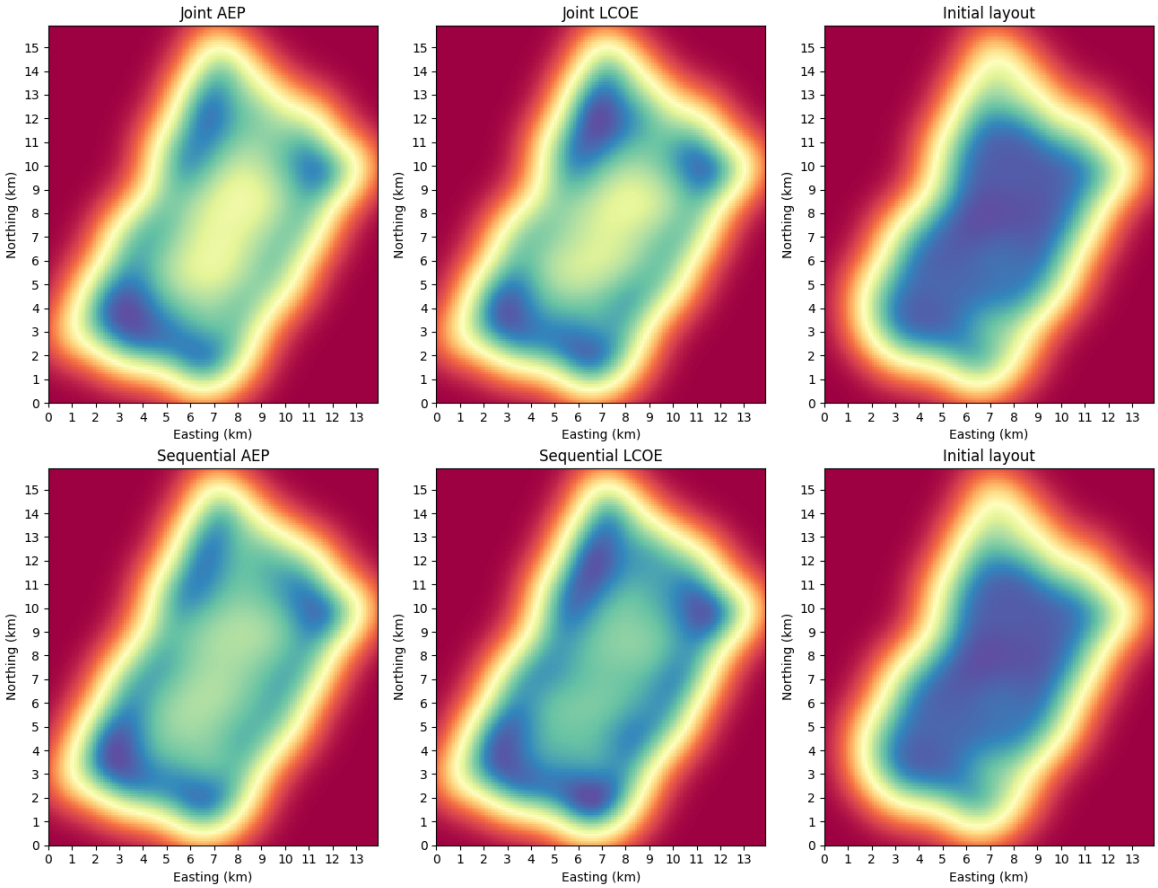

Figure 3-3: A power density metric is applied on 50 optimized layouts, it shows where the turbines are placed. This is done for the joint optimized AEP, the joint optimized LCOE, the sequential optimized AEP, and the sequential optimized LCOE. The power density of the initial population is shown as well. Darker blue means a higher power density. The point $(522, 5847)$ is taken as the origin.

Figure 3-3 illustrates a few differences in the characteristics of the optimized layouts. Namely, from the initial layouts where the turbines are randomly spread out over the domain, the turbines move to the edges with all four of the optimizations. This behaviour is also seen in recently built wind farms with optimized layouts, for example, Hollandse Kust Noord and Dogger Bank[1].

[1] See Appendix D

When looking at the joint optimized layouts as opposed to the sequential optimized layouts, the turbines move even more extremely to the edges. Especially the southwestern edge and the northern and eastern corners. This can be explained by the fact that the dominant wind direction is the southwest, so naturally, the best positions for a turbine are on the southwestern edge. When there is no room on that edge, the northern corner is attractive, since it still receives free stream wind when it comes from the southwest. If there is no room on the southwestern edge as well as the northern corner, the best positions are in the eastern corner. This is because, in the eastern corner, the wakes from the dominant wind direction have the longest recovery time.

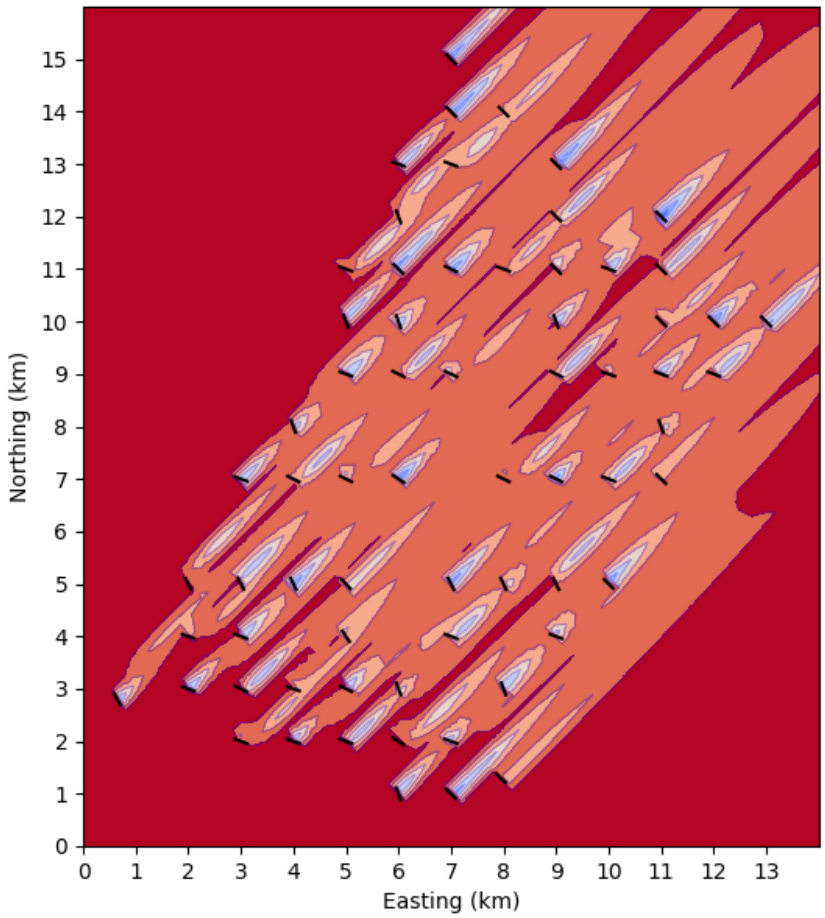

Figure 3-4: The wind speeds throughout the wind farm are shown with the wind coming from the dominant wind direction, the southwest. The shown layout is the best-found layout for the joint optimization of the LCOE. The corresponding LCOE considering all wind directions is 114.75 €/MWh and the AEP is 2878 GWh. The point $(522, 5847)$ is taken as the origin.

The change between the optimized layouts with the AEP and the LCOE as objective, can be explained by the water depth of the domain because the biggest difference between the LCOE and the AEP is the cost of each extra meter of water depth. As shown in Figure 2-7, the west side of the IJmuiden Ver has considerably deeper waters than the rest of the domain. This means the turbines move away from the western corner in the LCOE optimized layout as opposed to the AEP optimized layout, see Figure 3-3.

3-3 Case study IJmuiden Ver: Robustness

To investigate the robustness of the original objective used in the GA, where 72 wind directions, a single wind speed, and the geometric yaw relationship are applied, four robustness checks are considered. Firstly, a FLORIS-yaw check is performed. This means the same wind conditions are taken as the original objective, however, the FLORIS yaw optimizer is employed to get the optimal yaw angles instead of the geometric yaw relationship. Secondly, the objective is calculated using 360 wind directions. The rest is like the original objective: a single wind speed and the geometric yaw relationship are used. Thirdly, 22 wind speeds are applied. Again the other factors are identical to the original objective. Finally, the FLORIS-optimized yaw angles, 360 wind directions and 22 wind speeds are taken for the robust objective. The results of running each robustness check for 20 joint and sequential optimized layouts are shown in Figure 3-5.

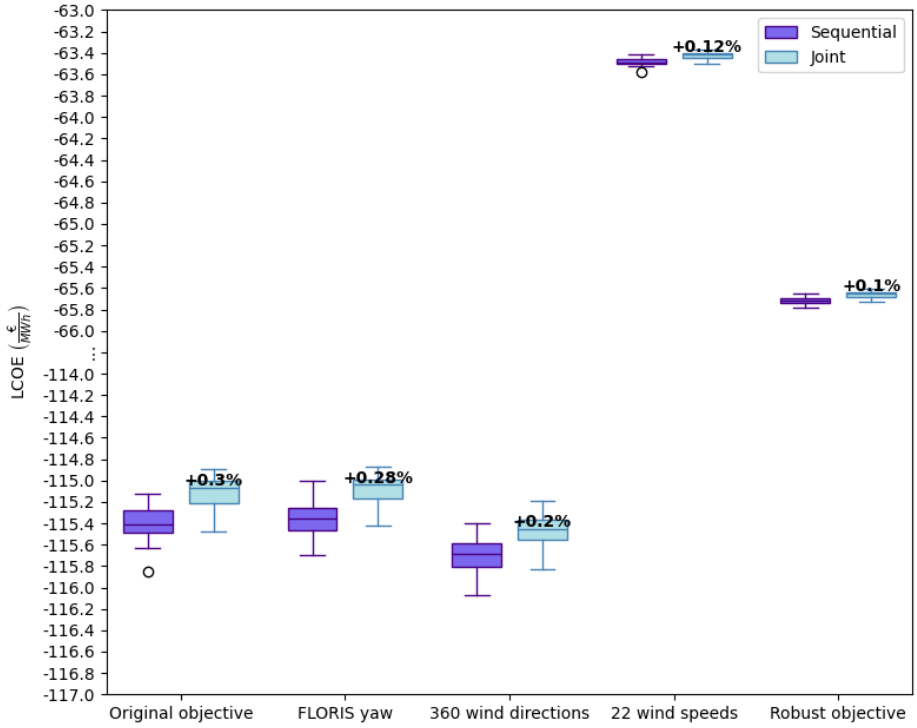

Figure 3-5: Each robustness check is applied on 20 joint and 20 sequential optimized layouts. The original objective is the objective used during the GA. With the FLORIS-yaw check the FLORIS yaw optimizer is employed. The 360-wind-directions check considers all wind directions of the complete wind rose and the 22-wind-speeds check considers all wind speeds. For the robust objective, the FLORIS yaw optimizer and the complete wind rose are applied. The data points are shown in the box plot. The joint optimization method is compared to the sequential optimization method, with the relative improvement of the median LCOE annotated in the figure. For the precise median and variance values see Table E-4.

The FLORIS-yaw check shows a similar improvement when comparing the joint to the sequential optimization as the original objective. The FLORIS-optimized yaw gives a slightly better LCOE. This is expected as the geometric yaw relationship is fitted to the FLORIS-optimized yaw angles. However, the geometric yaw relationship seems to be able to find layouts that perform well with FLORIS-optimized yaw angles too.

When applying four times as many wind directions compared with the original objective, the improvement of the joint compared to the sequential optimization drops. Although there is still a significant improvement of 0.2%, the influence of the wind direction bins is clear. Since the number of wind directions is already taken quite high in the original objective, taking four times as many wind directions might not be desirable for computational reasons. Nonetheless, it is helpful to keep the influence of the number of wind directions in mind.

Looking at the 22-wind-speeds check, there is a considerable difference from the original objective. The original only takes a wind speed of 8 m/s into account. However, the actual average wind speed is around 10 m/s. This makes the energy production of the wind farm significantly higher. Furthermore, the number of wind speeds seems to have a substantial influence on the objective. The improvement when comparing the joint to the sequential optimization is reduced from 0.3% to 0.12% and the variance is around 20 times smaller as well.

For the robust objective, the FLORIS-optimized yaw angles, 360 wind directions and 22 wind speeds are applied. The improvement when comparing the joint to the sequential optimization decreases even further to 0.1%. The variance of the robust objective is small as well, $8.9 \cdot 10^4$ for the joint and $1.34 \cdot 10^3$ for sequential optimization. Even though the objective used in the GA must be simplified for computational reasons, it seems that the wind rose is simplified too much for the GA to give robust results.

3-4 Case study IJmuiden Ver: Ecology

To investigate how ecology goals influence the solution, the LCOE is computed for different ecology constraints. In the IJmuiden Ver case, the minimum distance to the nature reserve 'Bruine Bank' is already set to 2 km [41]. From a distance of 5.5 km and onward , the GA has a hard time finding enough distinctive initial layouts adhering to the minimum spacing of $4D$. For this reason, the distances from 2 up to 5 km with steps of 0.5 km are investigated. These distances correspond to power densities of 10.4 up to 13.5 W/m^2. Note that in the test case, enough distinctive initial layouts could be found for higher energy densities since the minimum spacing was only $2D$. The results of running the GA 50 times in a joint and 50 times in a sequential way for each distance are shown in Figure 3-6.

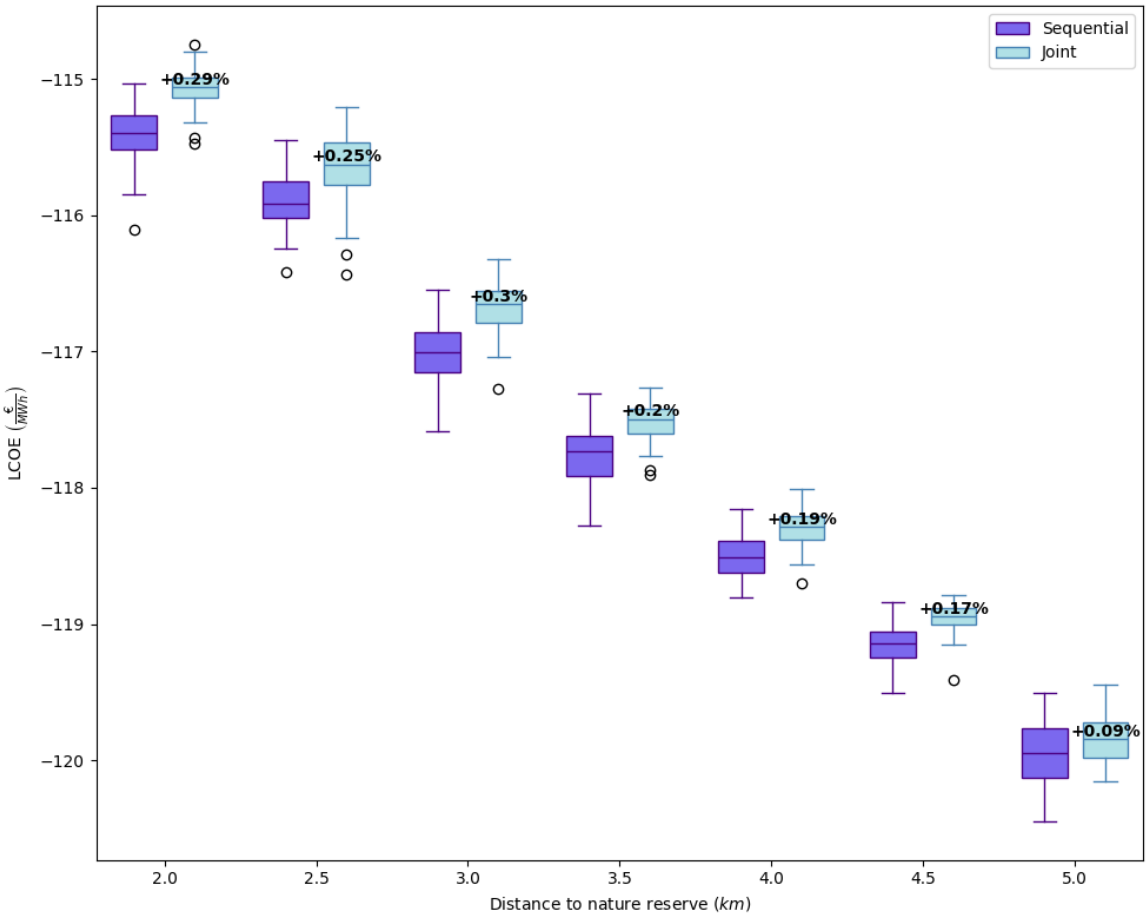

Figure 3-6: The IJmuiden Ver case is run 50 times in a joint and 50 times in a sequential way per distance from 2 up to 5 km to the nature reserve 'Bruine Bank'. These data points are shown in the box plot. The joint optimization method is compared to the sequential optimization method, with the relative improvement of the median LCOE annotated in the figure. For the precise median and variance values see Table E-5.

Figure 3-6 shows that the joint optimization method performs substantially better than the sequential optimization method for different domains. In general, the variance of the joint optimization is smaller as well. However, when the power density of the wind farm gets too high, the difference between the joint and sequential optimization diminishes. This could be due to yaw control being less effective when there is room to steer the wake away from other turbines. Furthermore, when the power density of a wind farm is very high, there is little room to optimize the layout. The underlying characteristics of the joint optimized layout stay consistent for higher density wind farms; the turbines move to the edges, see Figure 3-7.

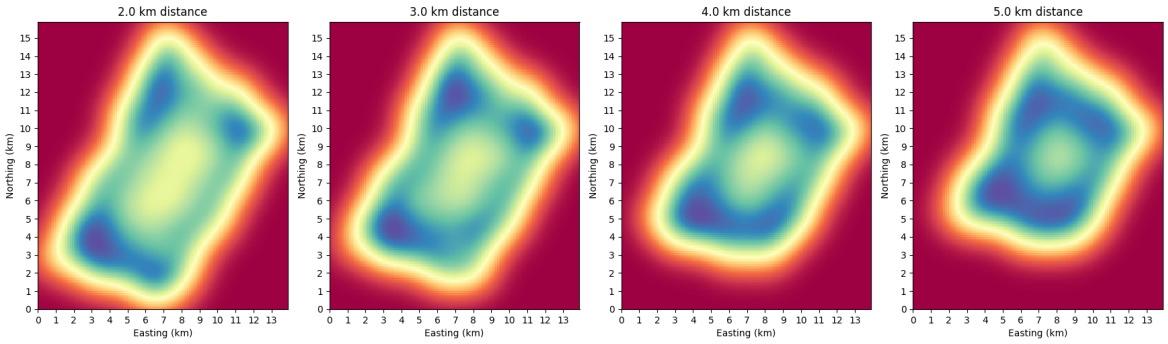

Figure 3-7: The power density metric is applied to multiple ecology cases where the distances to the nature reserve are $2, 3, 4$ and 5 km. Shown is where the turbines of 50 joint optimized layouts for each ecology case are placed. Darker blue means a higher power density. The point $(522, 5847)$ is taken as the origin.

Comparing the results of the IJmuiden Ver case to those of the test case, the improvements when comparing the joint to the sequential optimization in the IJmuiden Ver case are a lot higher than for the corresponding energy densities in the test case. In the IJmuiden Ver case, the improvement is between 0.17% and 0.30% for energy densities up to $13\ W/m^2$, unlike the test case where improvements up to 0.08% were found. In other words, the joint optimization method seems to yield more benefit in wind farms with a large number of turbines, although further research is needed on this topic.

3-5 Case study IJmuiden Ver: Penalty

The gain of optimizing the wind farm layout in a joint way is based on the assumption that yaw control is always available. The penalty for not using yaw control is investigated by computing the LCOE of the sequential and joint optimized layouts with and without yaw control. The 50 joint and 50 sequential optimized layouts for each ecology constraint in the IJmuiden Ver case are used, see section 3-4. The LCOE of these layouts with yaw control is shown in grey, repeating the results of Figure 3-6, and the LCOE without yaw control is shown in colour, see Figure 3-8.

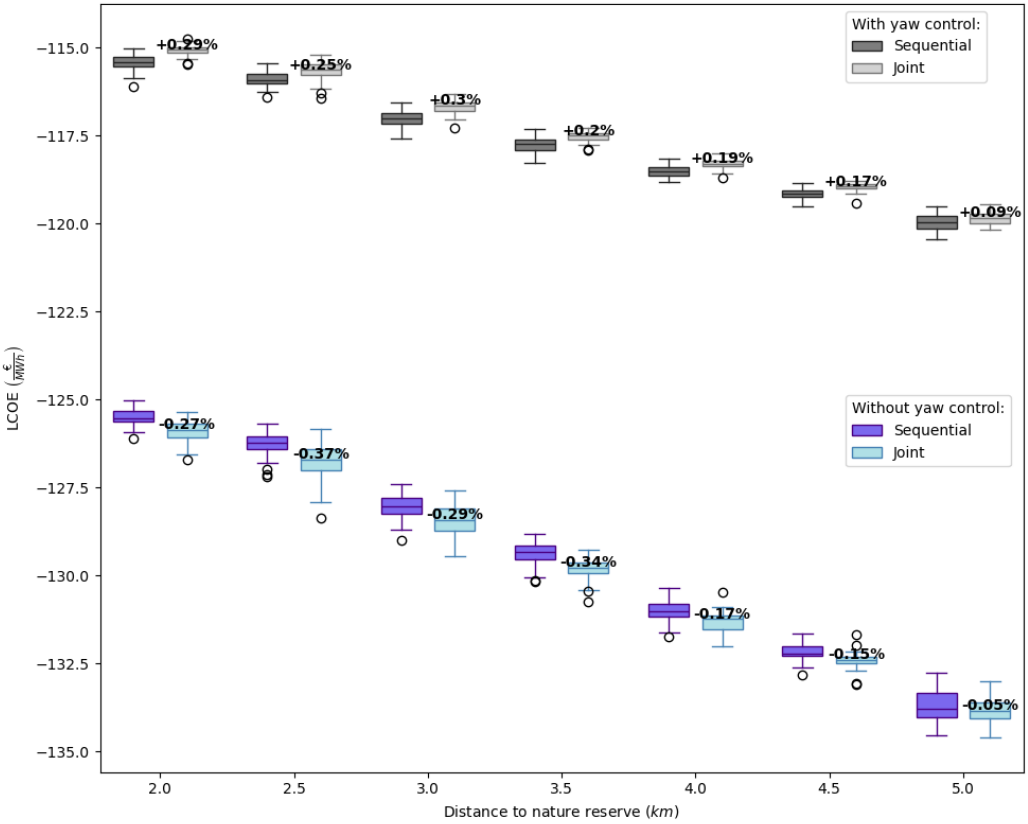

Figure 3-8: The IJmuiden Ver case is run 50 times in a joint and 50 times in a sequential way per distance from 2 up to 5 km to the nature reserve 'Bruine Bank'. The LCOE of the optimized layouts is computed with and without yaw control. These data points are shown in the box plot where the grey bars are with yaw control and the coloured bars are without yaw control. The joint optimization method is compared to the sequential optimization method, with the relative improvement of the median LCOE annotated in the figure. For the precise median and variance values see Table E-6.

Comparing the results with and without yaw control in Figure 3-8 clearly shows the huge benefit of applying yaw control in general. As expected, the sequential optimized layouts perform better without yaw control, than the joint optimized layouts. Without yaw control, the results are almost a mirrored image of the results with yaw control. This means that even though joint optimization gives significantly better results with yaw control, there is a substantial penalty when yaw control is not available.

Chapter 4

Conclusion

For the joint optimization of a wind farm layout and yaw control, the geometric yaw relationship provides a computationally inexpensive approach to take yaw control into account. The optimized layouts that are produced by the genetic algorithm (GA) with the geometric yaw relationship presented in this thesis, perform well with FLORIS-optimized yaw angles too. This makes the step from sequential to joint optimization of the wind farm layout optimization problem (WFLOP) very small since one of the biggest challenges of joint optimization is the computational intensity.

When looking at the effects of joint as opposed to sequential optimization, the improvement of the objective value depends substantially on the number of turbines and on the domain that is adopted. In the test case, where a 16-turbine wind farm is analysed, the improvement differs from almost 0.3% for some power densities to no improvement for other power densities. This is considerably less when compared to literature findings. However, the optimization methods used in literature have the unrealistic assumption that the domain is continuous. In the IJmuiden Ver case, where a 67-turbine wind farm is considered, the improvement when comparing the joint to the sequential optimization is quite significant: 0.34% for the annual energy production (AEP) and 0.29% for the levelized cost of electricity (LCOE). Moreover, the variance of the joint optimized LCOE is significantly smaller than the sequential optimized LCOE in the IJmuiden Ver case, 0.02 and 0.04 respectively. However, if the yaw control is not available in a joint optimized wind farm, the gain turns into loss.

The power density metric shows that the turbines move to the edges with all four of the optimizations: the joint optimized AEP, the joint optimized LCOE, the sequential optimized AEP, and the sequential optimized LCOE. When looking at the joint optimized layouts as opposed to the sequential optimized layouts, the turbines move even more extremely to the edges, especially the edge of the dominant wind direction. The biggest difference between the LCOE and the AEP is the cost of each extra meter of water depth. This means the turbines move away from the deep waters as shown by the power density metric.

The influence of the chosen number of wind bins is significant. When including 360 wind direction bins instead of 72 in LCOE computation of the optimized layout, the improvement when comparing the joint to the sequential optimization drops from 0.3% to 0.2%. Likewise, if 22 wind speed bins are taken as opposed to one, the improvement is reduced from 0.3% to 0.12%. Even though the objective used in the GA must be simplified for computational reasons, it seems that the wind rose is simplified too much for the GA to give robust results.

Joint optimization gives a substantial improvement compared to sequential optimization for the different ecology scenarios with each a distinct distance to the nature reserve 'Bruine Bank', corresponding to power densities of 10.4 up to 13.5 W/m^2. However, when the power density of the wind farm gets too high, the difference between joint and sequential optimization diminishes. This could be due to yaw control being less effective when there is room to steer the wake away from other turbines. Another factor with high-density wind farms is that there is little room to optimize the layout, especially with a minimum spacing of $4D$. In general, the influence of the domain on the improvement is considerably smaller in the IJmuiden Ver case than in the test case. Joint optimization seems to provide more steady improvement in wind farms with a large number of turbines.

Chapter 5

Discussion and future work

Due to a limited amount of time, not everything could be researched in the scope of this thesis. One of the significant choices is the genetic algorithm (GA) as metaheuristic. Because of computational power constraints, the maximum number of iterations was set to 200. Taking a higher number of iterations might improve the results even further. Other metaheuristics, like simulated annealing (SA) or partical swarm optimization (PSO), may give similar or even better results. Moreover, a second heuristic can be applied to refine the solution.

It is possible to improve the levelized cost of electricity (LCOE) as well. Because of limited real-world data, the function is simplified. One of the simplifications is that the annual energy production (AEP) is taken constant throughout the years. Considering future wind farms and climate change makes the results more realistic. Likewise, the cable layout is not taken into account. In this thesis, the distance to the substation is used as an indication of the amount of cable needed. Additionally, dynamic energy prices could be applied. Since the improvement of energy production due to wake steering mostly occurs in low wind conditions, the actual gain may be more significant than expected based on the LCOE.

In this thesis, the geometric yaw relationship is shown to be suited for the joint optimization of a wind farm layout and yaw control. The data, containing the FLORIS optimized yaw angles for varying dx and dy values, produced with the parameters specific to the IJmuiden Ver case is very similar to the data produced by Stanley *et al.* [32], who used a different turbine type, wind rose, velocity model, domain and minimum spacing. This suggests that a general geometric yaw relationship can be designed that is applicable to a large variety of wind farms.

The robust objective showed that if the computational power allows it, taking more wind bins into account is advisable. A more realistic wind rose can be adopted as well by incorporating the effects of neighbouring wind farms and not viewing the wind farm as one point but applying a spatially varying inflow of wind. Furthermore, the robust objective is only the beginning of examining the robustness of the solution. For example, a Monte Carlo simulation of different wind scenarios can be employed. Additionally, uncertainty in the building planning of future wind farms, wake model parameters, turbine maintenance and yaw control availability could be included.

On the ecology front, increasing the distance to the nature reserve 'Bruine Bank' is one of the many possible ecology measures. Other approaches, bird corridors, for instance, are interesting for the wind farm layout optimization problem (WFLOP) as well. This will make the layout optimization even more complex.

When looking at the penalty for not using yaw control on a joint optimized layout, it is clear that this penalty should not be neglected in practice. The penalty is almost as much as the improvement. If it is known that yaw control does not work consistently, then this can be incorporated into the objective.

Furthermore, it would be interesting to see how the joint optimization method performs for different numbers of turbines. It appears the effect of joint optimization is greater in larger wind farms. However, in this thesis, only two cases are compared. Most of the future wind farms are larger than the Alpha I plot of IJmuiden Ver [59]. If the joint optimization turns out to be indeed more effective for larger wind farms, the benefit of using joint optimization will only increase over time.

Although there is still a lot to be improved upon, this thesis showed that joint optimization is applicable in practice. There is a substantial gain when optimizing the layout in a joint way using a realistic domain and LCOE as objective. Naturally, the most energy is produced by a wind farm with a low power density. However, with current and future wind farms being built with higher power densities to preserve nature, the joint optimization of the wind farm layout and yaw control will be important in order to hit climate goals.

Appendix A

Python code

Code available at: `https://github.com/Robin9697/WFLOP_De_Jong`

Appendix B

Density of offshore wind farms

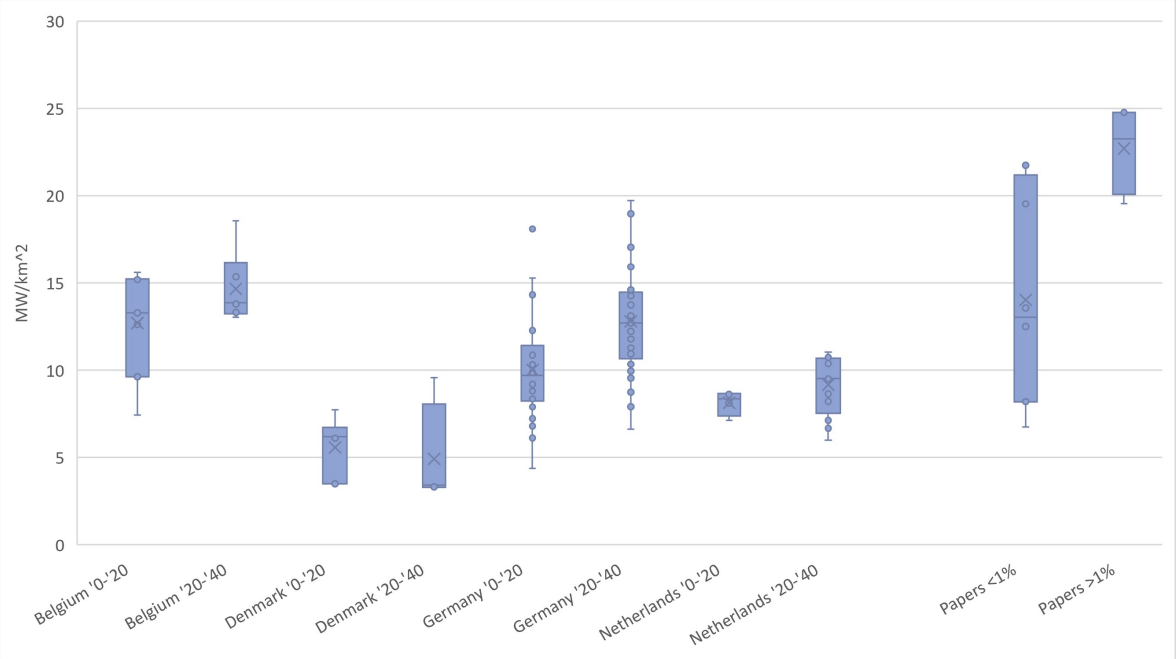

Figure B-1: Power density of offshore wind farms in the North sea in terms of $\frac{W}{m^2}$ per two decades per country, based on data from [59], compared to literature with a gain found below and above 1%, see Table 1-1.

Appendix C

Wind speeds

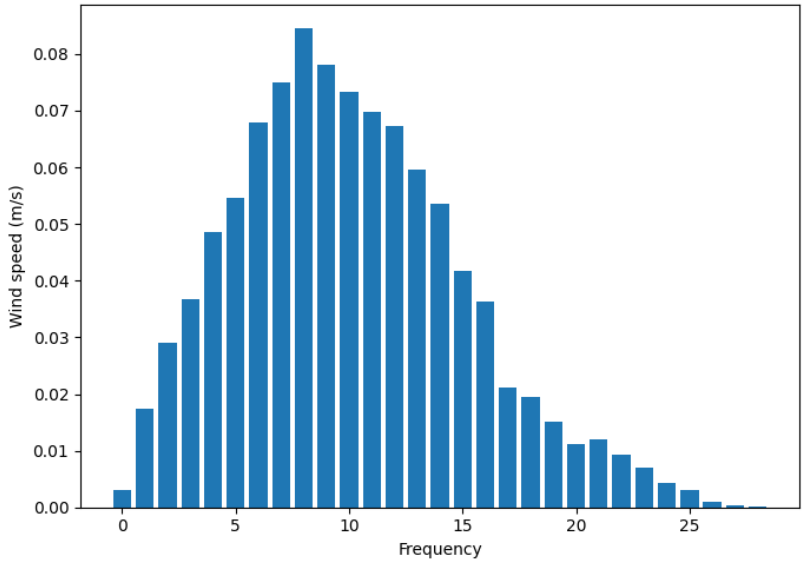

Figure C-1: Shown are the wind speeds of the wind rose and their frequencies. The mode is 8 m/s and the average is around 10 m/s.

Appendix D

Layouts of recently built wind farms

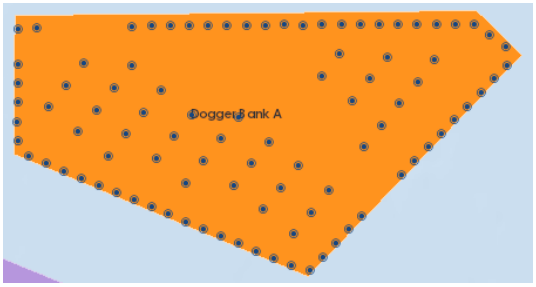

Figure D-1: The layout of Dogger Bank A [60]

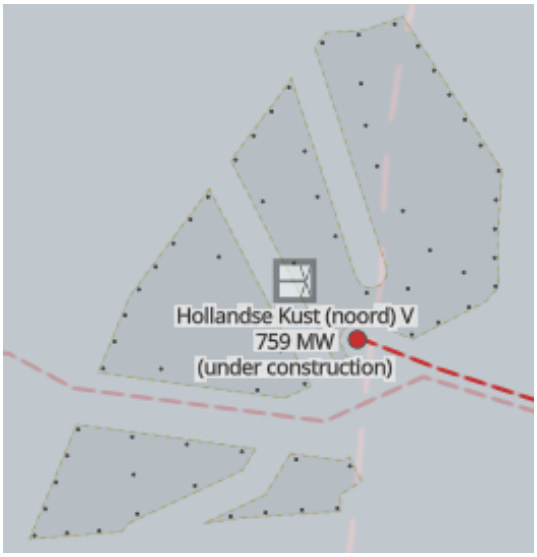

Figure D-2: The layout of Hollandse Kunst Noord [61]

Appendix E

Quantitative description of the box plots

Table E-1: Quantitative description of the box plot in Figure 2-4

Yaw angles	Median	Variance
FLORIS yaw	2861.93414	24.55164
GY De Jong	2860.08455	27.23736
GY Stanley	2813.80870	32.48626
No yaw	2629.33062	65.41379

Table E-2: Quantitative description of the box plot in Figure 3-1

Power density (W/m^2)	Optimization method	Median	Variance
8	Sequential	742.023	7.107
8	Joint	742.776	5.757
9	Sequential	738.652	11.552
9	Joint	739.417	6.864
10	Sequential	733.919	11.356
10	Joint	734.276	8.847
11	Sequential	730.832	10.695
11	Joint	730.776	7.993
12	Sequential	726.006	13.295
12	Joint	726.515	11.694
13	Sequential	723.846	9.740
13	Joint	724.427	9.748
14	Sequential	720.223	15.732
14	Joint	720.851	8.376
15	Sequential	715.299	13.512
15	Joint	717.259	8.750
16	Sequential	716.169	15.985
16	Joint	717.456	8.301
17	Sequential	712.213	14.851
17	Joint	713.139	7.438
18	Sequential	709.371	17.684
18	Joint	709.349	12.553
19	Sequential	704.962	14.426
19	Joint	707.012	18.354
20	Sequential	702.268	23.262
20	Joint	702.123	20.942

Table E-3: Quantitative description of the box plot in Figure 3-2

Objective function	Optimization method	Median	Variance
AEP	Sequential	2861.009	22.434
AEP	Joint	2870.765	28.617
LCOE	Sequential	-115.399	0.042
LCOE	Joint	-115.062	0.023

Table E-4: Quantitative description of the box plot in Figure 3-5

Robustness check	Optimization method	Median	Variance
Original objective	Sequential	-115.40797	0.03496
Original objective	Joint	-115.06637	0.02160
FLORIS yaw	Sequential	-115.35896	0.03040
FLORIS yaw	Joint	-115.04040	0.01971
360 wind directions	Sequential	-115.68020	0.02976
360 wind directions	Joint	-115.45186	0.02322
22 wind speeds	Sequential	-63.48806	0.00151
22 wind speeds	Joint	-63.41196	0.00097
Robust objective	Sequential	-65.71729	0.00134
Robust objective	Joint	-65.65295	0.00089

Table E-5: Quantitative description of the box plot in Figure 3-6

Distance to nature reserve (km)	Optimization method	Median	Variance
2.0	Sequential	-115.399	0.042
2.0	Joint	-115.062	0.023
2.5	Sequential	-115.914	0.042
2.5	Joint	-115.629	0.059
3.0	Sequential	-117.002	0.047
3.0	Joint	-116.650	0.035
3.5	Sequential	-117.728	0.041
3.5	Joint	-117.496	0.020
4.0	Sequential	-118.513	0.027
4.0	Joint	-118.290	0.017
4.5	Sequential	-119.145	0.021
4.5	Joint	-118.946	0.012
5.0	Sequential	-119.946	0.050
5.0	Joint	-119.844	0.031

Table E-6: Quantitative description of the box plot in Figure 3-8

Distance to nature reserve (km)	Optimization method	Yaw control	Median	Variance
2.0	Sequential	Without	-125.51580	0.05464
2.0	Joint	Without	-125.85415	0.08959
2.0	Sequential	With	-115.39943	0.04176
2.0	Joint	With	-115.06167	0.02311
2.5	Sequential	Without	-126.22998	0.11826
2.5	Joint	Without	-126.69116	0.24968
2.5	Sequential	With	-115.91428	0.04205
2.5	Joint	With	-115.62884	0.05852
3.0	Sequential	Without	-128.04100	0.11482
3.0	Joint	Without	-128.41699	0.15844
3.0	Sequential	With	-117.00210	0.04700
3.0	Joint	With	-116.64992	0.03492
3.5	Sequential	Without	-129.32341	0.09691
3.5	Joint	Without	-129.76611	0.08640
3.5	Sequential	With	-117.72813	0.04051
3.5	Joint	With	-117.49627	0.02030
4.0	Sequential	Without	-131.00998	0.09081
4.0	Joint	Without	-131.23526	0.08280
4.0	Sequential	With	-118.51335	0.02671
4.0	Joint	With	-118.28950	0.01720
4.5	Sequential	Without	-132.20046	0.05705
4.5	Joint	Without	-132.39872	0.05198
4.5	Sequential	With	-119.14507	0.02069
4.5	Joint	With	-118.94648	0.01150
5.0	Sequential	Without	-133.76623	0.20236
5.0	Joint	Without	-133.82662	0.15924
5.0	Sequential	With	-119.94582	0.05017
5.0	Joint	With	-119.84373	0.03078

Bibliography

[1] P. Veers, K. Dykes, S. Basu, A. Bianchini, A. Clifton, P. Green, H. Holttinen, L. Kitzing, B. Kosovic, J. Lundquist, J. Meyers, M. O'malley, W. Shaw, and B. Straw, "Grand Challenges: wind energy research needs for a global energy transition," *Wind Energy Science*, vol. 7, no. 6, pp. 2491–2496, 2022.

[2] S. Bouckaert, A. F. Pales, C. McGlade, U. Remme, B. Wanner, L. Varro, D. D'Ambrosio, and T. Spencer, "Net Zero by 2050: A Roadmap for the Global Energy Sector," 2021.

[3] Intergovernmental Panel on Climate Change (IPCC), "Synthesis Report: Climate Change 2023," tech. rep., 2023.

[4] H. Canet, A. Guilloré, and C. L. Bottasso, "The eco-conscious wind turbine: bringing societal value to design," *Wind Energy Science Discussions*, pp. 1–30, June 2022. Publisher: Copernicus GmbH.

[5] S. Anagnostopoulos, J. Bauer, M. C. A. Clare, and M. D. Piggott, "Accelerated wind farm yaw and layout optimisation with multi-fidelity deep transfer learning wake models," Mar. 2023. arXiv:2303.16274 [physics].

[6] R. Bos, "Conversations with the wind turbine and energy yield specialist of the offshore wind department of Eneco," 2023.

[7] X.-Y. Tang, Q. Yang, B. Stoevesandt, and Y. Sun, "Optimization of wind farm layout with optimum coordination of turbine cooperations," *Computers and Industrial Engineering*, vol. 164, 2022.

[8] S. Boersma, B. Doekemeijer, P. Gebraad, P. Fleming, J. Annoni, A. Scholbrock, J. Frederik, and J. van Wingerden, "A tutorial on control-oriented modeling and control of wind farms," in *2017 American Control Conference (ACC)*, pp. 1–18, May 2017. ISSN: 2378-5861.

[9] A. C. Kheirabadi and R. Nagamune, "A quantitative review of wind farm control with the objective of wind farm power maximization," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 192, pp. 45–73, Sept. 2019.

[10] D. van der Hoek, J. Frederik, M. Huang, F. Scarano, C. Simao Ferreira, and J.-W. van Wingerden, "Experimental analysis of the effect of dynamic induction control on a wind turbine wake," *Wind Energy Science*, vol. 7, pp. 1305–1320, June 2022. Publisher: Copernicus GmbH.

[11] B. Wade, R. Pereira, and C. Wade, "Investigation of offshore wind farm layouts regarding wake effects and cable topology," *Journal of Physics: Conference Series*, vol. 1222, p. 012007, May 2019. Publisher: IOP Publishing.

[12] J. F. Herbert-Acero, O. Probst, P.-E. Réthoré, G. C. Larsen, and K. K. Castillo-Villar, "A Review of Methodological Approaches for the Design and Optimization of Wind Farms," *Energies*, vol. 7, pp. 6930–7016, Nov. 2014. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.

[13] P. M. O. Gebraad, F. W. Teeuwisse, J. W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao, "Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study," *Wind Energy*, vol. 19, no. 1, pp. 95–114, 2016. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.1822.

[14] P. Gebraad, J. J. Thomas, A. Ning, P. Fleming, and K. Dykes, "Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control," *Wind Energy*, vol. 20, no. 1, pp. 97–107, 2017. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.1993.

[15] M. F. Howland, S. K. Lele, and J. O. Dabiri, "Wind farm power optimization through wake steering," *Proceedings of the National Academy of Sciences*, vol. 116, pp. 14495–14500, July 2019. Publisher: Proceedings of the National Academy of Sciences.

[16] P. Hou, J. Zhu, K. Ma, G. Yang, W. Hu, and Z. Chen, "A review of offshore wind farm layout optimization and electrical system design methods," *Journal of Modern Power Systems and Clean Energy*, vol. 7, no. 5, pp. 975–986, 2019.

[17] A. Stanley and A. Ning, "Massive simplification of the wind farm layout optimization problem," *Wind Energy Science*, vol. 4, no. 4, pp. 663–676, 2019.

[18] D. van der Hoek, B. van den Abbeele, C. S. Ferreira, and J.-W. van Wingerden, "Maximizing wind farm power output with the helix approach – experimental validation and wake analysis using tomographic PIV," 2023.

[19] W. P. Monthly, "Wind Turbine Control Systems: Exploring the capabilities of the latest systems, and the drivers and challenges for further development.," *Wind Power Monthly*, 2014.

[20] J. A. Frederik, B. M. Doekemeijer, S. P. Mulders, and J. W. van Wingerden, "The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms," *Wind Energy*, vol. 23, no. 8, pp. 1739–1751, 2020. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.2513.

[21] A. Jiménez, A. Crespo, and E. Migoya, "Application of a LES technique to characterize the wake deflection of a wind turbine in yaw," *Wind Energy*, vol. 13, no. 6, pp. 559–572, 2010.

[22] P. Fleming, J. King, K. Dykes, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. van Dam, C. Bay, R. Mudafort, H. Lopez, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake, "Initial results from a field campaign of wake steering applied at a commercial wind farm – Part 1," *Wind Energy Science*, vol. 4, pp. 273–285, May 2019. Publisher: Copernicus GmbH.

[23] M. O. L. Hansen, *Aerodynamics of wind turbines*. London ; Sterling, VA: Earthscan, 2nd ed ed., 2008. OCLC: ocm86172940.

[24] W. Munters and J. Meyers, "Towards practical dynamic induction control of wind farms: Analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines," *Wind Energy Science*, vol. 3, no. 1, pp. 409–425, 2018.

[25] J. W. Van Wingerden, "Wind farm control - wake mixing," Apr. 2021.

[26] R. Z. en Delta, "Windenergiegebied Borssele," 2020. Last Modified: 2022-10-17.

[27] L. Wang, A. Tan, and Y. Gu, "A novel control strategy approach to optimally design a wind farm layout," *Renewable Energy*, vol. 95, pp. 10–21, 2016.

[28] K. Chen, J. Lin, Y. Qiu, F. Liu, and Y. Song, "Joint optimization of wind farm layout considering optimal control," *Renewable Energy*, vol. 182, pp. 787–796, 2022.

[29] P. Fleming, A. Ning, P. Gebraad, and K. Dykes, "Wind plant system engineering through optimization of layout and yaw control," *Wind Energy*, vol. 19, no. 2, pp. 329–344, 2016.

[30] M. Pedersen and G. Larsen, "Integrated wind farm layout and control optimization," *Wind Energy Science*, vol. 5, no. 4, pp. 1551–1566, 2020.

[31] J. Song, T. Kim, and D. You, "Particle swarm optimization of a wind farm layout with active control of turbine yaws," *Renewable Energy*, vol. 206, pp. 738–747, 2023.

[32] A. P. J. Stanley, C. J. Bay, and P. Fleming, "Enabling control co-design of the next generation of wind power plants," *Wind Energy Science*, vol. 8, pp. 1341–1350, Aug. 2023. Publisher: Copernicus GmbH.

[33] M. J. Faijer, M. Sosef, Pondera Consult, L. Perk, and WaterProof, "MER kavel VI Windenergiegebied Hollandse Kust (west)," tech. rep., Ministeries van Economische Zaken en Klimaat, Binnenlandse Zaken en Koninkrijkrelaties Infrastructuur en Waterstaat en Landbouw, Natuur en Voedselkwaliteit, Nov. 2020.

[34] Ministerie van Economische Zaken en Klimaat, "Regeling vergunningverlening kavel Alpha in windenergiegebied IJmuiden Ver," 2023.

[35] S. Tao, Q. Xu, A. Feijóo, G. Zheng, and J. Zhou, "Nonuniform wind farm layout optimization: A state-of-the-art review," *Energy*, vol. 209, 2020.

[36] A. Tesauro, P.-E. Réthoré, and G. Larsen, "State of the art of wind farm optimization," vol. 3, pp. 2020–2030, 2012.

[37] B. Kaya and E. Oğuz, "Investigation of layout optimization for offshore wind farms and a case study for a region in Turkey," *Ocean Engineering*, vol. 266, 2022.

[38] C. O. R. Energy, "Wind farm costs – Guide to an offshore wind farm," 2023.

[39] E. Quaeghebeur, R. Bos, and M. B. Zaaijer, "Wind farm layout optimization using pseudo-gradients," *Wind Energy Science*, vol. 6, pp. 815–839, June 2021. Publisher: Copernicus GmbH.

[40] Ministerie van Economische Zaken en Klimaat, "Kavelbesluit V windenergiegebied Hollandse Kust (noord)," 2019.

[41] Ministerie van Economische Zaken en Klimaat, "Ontwerpkavelbesluit kavel Alpha windenergiegebied IJmuiden Ver," 2023.

[42] S. van den Brenk and R. van Lil, "Archaeological Desk Study: Hollandse Kust (noord) Wind Farm Zone," tech. rep., 2017.

[43] S. Reddy, "An efficient method for modeling terrain and complex terrain boundaries in constrained wind farm layout optimization," *Renewable Energy*, vol. 165, pp. 162–173, 2021.

[44] N. Baker, A. Stanley, J. Thomas, A. Ning, and K. Dykes, "Best practices for wake model and optimization algorithm selection in wind farm layout optimization," AIAA Scitech Forum, 2019.

[45] Z. Liu, S. Fan, Y. Wang, and J. Peng, "Genetic-algorithm-based layout optimization of an offshore wind farm under real seabed terrain encountering an engineering cost model," *Energy Conversion and Management*, vol. 245, 2021.

[46] S. Khan and S. Rehman, "Iterative non-deterministic algorithms in on-shore wind farm design: A brief survey," *Renewable and Sustainable Energy Reviews*, vol. 19, pp. 370–384, 2013.

[47] M. Joshi, M. Gyanchandani, and D. Rajesh Wadhvani, "Analysis Of Genetic Algorithm, Particle Swarm Optimization and Simulated Annealing On Benchmark Functions," in *2021 5th International Conference on Computing Methodologies and Communication (ICCMC)*, pp. 1152–1157, Apr. 2021.

[48] F. Jia and D. Lichti, "A comparison of simulated annealing, genetic algorithm and particle swarm optimization in optimal first-order design of indoor TLS networks," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. IV-2-W4, pp. 75–82, Sept. 2017. Conference Name: ISPRS Geospatial Week 2017 (Volume IV-2/W4) - 18–22 September, Wuhan, China Publisher: Copernicus GmbH.

[49] Q. Yang, H. Li, T. Li, and X. Zhou, "Wind farm layout optimization for levelized cost of energy minimization with combined analytical wake model and hybrid optimization strategy," *Energy Conversion and Management*, vol. 248, 2021.

[50] P. A. Fleming, A. P. J. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doekemeijer, and R. Mudafort, "Serial-Refine Method for Fast Wake-Steering Yaw Optimization," *Journal of Physics: Conference Series*, vol. 2265, p. 032109, May 2022.

[51] P. Yang and H. Najafi, "The Effect of Using Different Wake Models on Wind Farm Layout Optimization: A Comparative Study," *Journal of Energy Resources Technology, Transactions of the ASME*, vol. 144, no. 7, 2022.

[52] J. Meyers, C. Bottasso, K. Dykes, P. Fleming, P. Gebraad, G. Giebel, T. Göçmen, and J. van Wingerden, "Wind farm flow control: Prospects and challenges," *Wind Energy Science*, vol. 7, no. 6, pp. 2271–2306, 2022.

[53] C. M. Nyborg, A. Fischer, P.-E. Réthoré, and J. Feng, "Optimization of wind farm operation with a noise constraint," *Wind Energy Science*, vol. 8, pp. 255–276, Feb. 2023. Publisher: Copernicus GmbH.

[54] X. Ju and F. Liu, "Wind farm layout optimization using self-informed genetic algorithm with information guided exploitation," *Applied Energy*, vol. 248, pp. 429–445, Aug. 2019.

[55] Ministerie van Economische Zaken en Klimaat, "Regeling vergunningverlening windenergiegebied Hollandse Kust (west) kavel VI," 2022.

[56] WINS50, "wins50 dataset," 2023.

[57] NREL, "IEA_15mw_240_rwt turbine model," 2020.

[58] Rijksdienst voor Ondernemend Nederland (RVO), "IJmuiden Ver Morphodynamics and Scour," 2023.

[59] 4C-Offshore, "Data wind farms North sea," 2023.

[60] 4C-offshore, "4C offshore offshore wind map," 2023.

[61] Openinframap, "Open Infrastructure Map," 2023.

Glossary

List of Acronyms

GA	genetic algorithm
AEP	annual energy production
LCOE	levelized cost of electricity
SA	simulated annealing
WFLOP	wind farm layout optimization problem
PSO	partical swarm optimization
SLSQP	sequential least squares programming